Chapter 1
Introduction

Robots need to understand their environment in order to be able to perform different tasks within it. A robot's interface with the external world is usually composed of several sensors that gather data. How to understand and interpret this sensor information is one of the fundamental problems in mobile robots.

Building maps of the environment is a typical example of how a robot uses sensor information to interpret the world. Some of the most usual maps in mobile robotics represent the parts in the environment which are occupied by objects such as occupancy grid maps [3, 8]. The occupancy information enables the robots to navigate without collisions, or to localize itself inside the environment. To build an occupancy map, a mobile robot usually moves along a trajectory while gathering information with some proximity sensor able to detect obstacles. Creating maps of the environment is extensively studied in the research area of *Simultaneous Localization And Mapping* (SLAM) — or *Concurrent Mapping and Localization* (CML)— [2, 6, 9, 11, 13, 14].

Occupancy maps are useful models for localization and navigation, however, other robotic tasks require a higher level of knowledge about the environment. Spatial concepts like the distinction between different locations, connectivity between them, or global topologies can be acquired by a robot from its experience inside the environment as shown by Kuipers and colleagues in the *Spatial Semantic Hierarchy* and its extensions [1, 5, 10].

In the particular case of indoor environments, we can find typical divisions of space such as corridors, rooms, or doorways. These divisions can be grouped into categories represented by a *semantic term*. A semantic term relates a category to some functionality, objective situation, or possible affordance in the place it represents. For example, the term *corridor* refers to the places which include doorways leading to other rooms. Furthermore, the term *doorway* indicates the places representing transitions between two different rooms, or between a room and a corridor. The process of applying a semantic term to some division of the environment is also known as *semantic place labeling*.

For a lot of applications, robots can improve their service and human-robot communication if they are able to obtain a semantic categorization of the environment in

Ó.M. Mozos: Semantic Labeling of Places with Mobile Robots, STAR 61, pp. 1–5.
springerlink.com

which they work. As an example, a robot that possesses information about the type of places can be instructed to "open the door to the corridor." Moreover, semantic terms like *corridor* or *room* give the human user a more intuitive idea of the position of the robot in comparison to the raw 2D coordinates inside a map. In addition, the semantic place information about places can improve the performance of the robot in other tasks such as localization or exploration.

In this book, we consider the problem of semantically categorizing the different locations of indoor environments using a mobile robot. An example is given in Fig. 1.1. The image in Fig. 1.1(a) shows the occupancy grid map corresponding to the ground floor of building 52 at the University of Freiburg. In this map, only information about occupied and free space is given. However, some natural divisions can be extracted from this environment, as for example rooms, doorways and a corridor. The original occupancy map is then extended with information about these place categories as shown in Fig. 1.1(b).

As explained before, a robot gathers information about the environment through its sensors. Thus, the main question is how a mobile robot can recognize the different places of an indoor environment using sensor data. In this book we solve this problem using different classification techniques. In our specific case, a classifier is a function that maps sensor information into place categories. A mobile robot uses this function to determine the place in which it is located based on its current sensor observation. Before performing this mapping, the data coming from sensors is transformed into a feature vector which is fed into the classifier. This general approach comprises the basics of the different methods that will be presented along the book.

It is important to note that we will present methods that allow a robot to recognize places by categories. The difference between categories and instances is important. In our case, an instance represents a concrete place in the environment such as "room 1015 in building 79 at the University of Freiburg". A category, however, represents a set of instances. For example the category *room* would contain all indoor places that are likely to be a room. In general, categorization is considered a harder problem than instantiation because the classifier needs to be general enough to include all instances in one category, while at the same time needs to be specific enough to distinguish instances between different categories. In addition, when creating a classifier, the robot needs to find common features for all instances belonging to the same category.

All classifiers presented in this book are learned using supervision. Supervised learning is an approach which needs a set of already labeled examples to construct the mapping function. These examples are usually provided to the robot by an external agent or tutor. An example of a tutor can be a person showing the robot the different places in a house.

Opposite to supervised learning approaches are unsupervised techniques. A robot learns in a unsupervised fashion when there is no tutor indicating the different categories of places that can be found in the environment. In this case it is the robot itself who learns the different places according to some distinction measurement extracted from its sensor data [1, 4, 5, 16]. Typical unsupervised approaches use

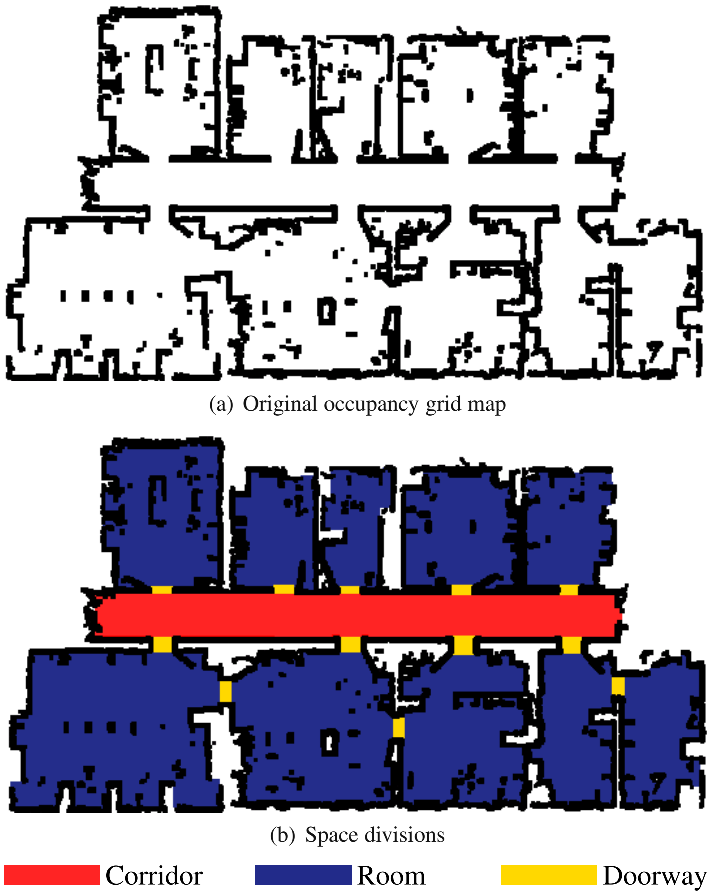

(a) Original occupancy grid map

(b) Space divisions

Fig. 1.1 Natural divisions of typical indoor environments. (**a**) shows the original occupancy grid map corresponding to the ground floor of building 52 at the University of Freiburg. (**b**) depicts some natural divisions extracted from the environment: rooms, doorways, and a corridor.

clustering [12] to separate the instances corresponding to each category. Unsupervised learning methods are outside the scope of this book.

The rest of the book is organized as follows. Chapter 2 gives a short introduction to supervised learning and presents the AdaBoost algorithm. For a more extensive introduction to supervised learning we refer the reader to [7, 12, 15].

In Chap. 3 we introduce the application of AdaBoost to assign semantic labels to different places in indoor environments using a mobile robot. The main idea is to classify each pose of a mobile robot into one of the semantic categories according to the laser range observation the robot gathered at that position. The classification is carried out using geometrical features extracted from the laser beams. Additionally, the boosting approach allows us to determine which are the most informative features used to recognize each place.

The method presented in Chap. 3 is extended in Chap. 4 to extract topological maps from indoor environments. The key idea is to apply the semantic classification to all possible poses of the robot in a map. In this way we obtain a complete categorization of the free space. Neighboring poses with similar classification are then grouped into regions which form the different nodes in the final topological map. Previous to the grouping, a smoothing method is applied that takes into account spatial dependencies between different labels.

The previous two methods cover the semantic classification of the different poses of a mobile robot, but they do not take into account the movement of the robot along a trajectory. When operating in indoor environments, the robot usually have a moderate velocity and a relatively continuous movement. That means, that observations obtained by a mobile robot at nearby poses are typically very similar. Based on this assumptions, Chap. 5 describes a method that takes into account previous classifications when classifying a new pose of a mobile robot along a trajectory. These spatial dependencies are modeled using a hidden Markov model. The robot used for the experiments in this chapter is equipped with a camera in addition to the laser sensor. The visual information is composed of objects extracted from the images. Using both sensors, the robot is able to distinguish a wider set of categories including kitchens, laboratories, offices, and seminar rooms.

The semantic labeling can be applied not only to improve the human-robot communication, but also to better carry out some other specific tasks for autonomous mobile robots. Chapter 6 presents the exploration of environments with a team of robots using place information. We will show how the semantic labeling of places can improve the distribution of the robots during the exploration. The main idea here is that corridors are better exploration targets as they lead to other rooms. In a second application, we will see how to accelerate the localization process of a single robot using the semantic classification of the different rooms.

Chapter 7 presents the semantic labeling of places as part of a high level conceptual representation of indoor environments called *multi-layer conceptual map*. This representation extends the semantic classification of places adding upper layers which include more complex conceptual terms, such as living rooms. The terms not only represent places but also objects such as TV sets or couches, and are used to create a human-friendly dialogue while interacting with people.

The above mentioned techniques are used to augment the information about environments with semantic terms. However, the AdaBoost-based classifiers can also be used to include semantic information in sensor data. Chapter 8 will show an approach to semantically label the beams of a laser scan. The main idea is to assign each beam the class of the object it hits. In this chapter, we restrict the classification to the labels *person* and *non-person*, although the method can be easily extended to use additional ones.

References

1. Beeson, P., Jong, N.K., Kuipers, B.: Towards autonomous topological place detection using the extended voronoi graph. In: Proceedings of the IEEE International Conference on Robotics and Automation (2005)
2. Dissanayake, G., Durrant-Whyte, H.F., Bailey, T.: A computationally efficient solution to the simultaneous localisation and map building (SLAM) problem. In: Proceedings of the IEEE International Conference on Robotics and Automation, San Francisco, CA, USA, April 2000, pp. 1009–1014 (2000)
3. Elfes, A.: Using occupancy grids for mobile robot perception and navigation. Computer 22(6), 46–57 (1989)
4. Fabrizi, E., Saffiotti, A.: Extracting topology-based maps from gridmaps. In: Proceedings of the IEEE International Conference on Robotics and Automation, San Francisco, CA, pp. 2972–2978 (2000)
5. Kuipers, B.: The spatial semantic hierarchy. Artificial Intelligence 119(1-2), 191–233 (2000)
6. Leonard, J.J., Durrant-Whyte, H.F.: Mobile robot localization by tracking geometric beacons. IEEE Transactions on Robotics and Automation 7(3), 376–382 (1991)
7. Mitchell, T.M.: Machine Learning. McGraw-Hill, New York (1997)
8. Moravec, H.P.: Sensor fusion in certainty grids for mobile robots. AI Magazine 9, 61–74 (1988)
9. Moutarlier, P., Chatila, R.: Stochastic multisensory data fusion for mobile robot location and environmental modelling. In: Fifth International Symposium of Robotics Research, Tokyo, Japan, pp. 85–94 (1989)
10. Remolina, E., Kuipers, B.: Towards a general theory of topological maps. Artificial Intelligence 152(1), 47–104 (2004)
11. Smith, R.C., Cheeseman, P.: On the representation and estimation of spatial uncertainty. The International Journal of Robotics Research 5(4), 56–68 (1986)
12. Theodoridis, S., Koutroumbas, K.: Pattern Recognition, 3rd edn. Academic Press, London (2006)
13. Thrun, S.: Robotic mapping: A survey. In: Lakemeyer, G., Nebel, B. (eds.) Exploring Artificial Intelligence in the New Millenium, pp. 1–34. Morgan Kaufmann, San Francisco (2002)
14. Thrun, S., Burgard, W., Fox, D.: Probabilistic Robotics. MIT Press, Cambridge (2005)
15. Witten, I.H., Frank, E.: Data Mining: Practical Machine Learning Tools and Techniques, 2nd edn. Morgan Kaufmann, San Francisco (2005)
16. Zivkovic, Z., Bakker, B., Kröse, B.: Hierarchical map building using visual landmarks and geometric constraints. In: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (2005)