Chapter 6
Semantic Information in Exploration and Localization*

6.1 Introduction

The work presented in the previous chapters showed how to augment the representation of indoor environments using semantic information about places. In this chapter we describe how robots can use the intrinsic information of human-made environments to improve their actions. In particular, we apply the semantic labeling of places to two robotic tasks: multi-robot exploration, and localization. In both cases the performance of the robot increases when it takes into account the classification of its location.

Exploration and localization belong to the fundamental problems in mobile robotics [22]. In the exploration task a mobile robot is controlled in a way that maximizes the information about its environment. A typical goal in exploration consists of creating a map of a previously unseen environment. Moreover, the use of multiple robots is often suggested to have advantages over a single robot during exploration, since cooperating robots have the potential to accomplish a task faster than a single robot [7]. In the localization task, a mobile robot has to determine its pose relative to a map of a given environment.

In this chapter, we first present an approach to include semantic information about places to better distribute the robots in human-made environments during an exploration task. As we have seen in previous chapters, indoor environments constructed by humans contain structures like corridors, rooms or offices. Moreover, corridors are connected to several rooms and provide more branchings to new unexplored areas. The key idea is then to assign higher rewards to robots that first explore corridors. As a result, the overall completion time of an exploration task can be significantly reduced.

In a second approach, we use the semantic labeling to localize a robot in an indoor environment using the Monte Carlo localization approach [3]. The main idea here is to take as observation model the semantic classification of the current pose of the mobile robot.

* The work presented in this chapter originated from a collaboration with Cyrill Stachniss.

Ó.M. Mozos: Semantic Labeling of Places with Mobile Robots, STAR 61, pp. 71–81.
springerlink.com

The rest of the chapter is organized as follows. The next section presents the approach for multi-robot exploration using semantic information. In Sect. 6.3, we introduce the Monte Carlo approach for localization using semantic labels. Section 6.4 presents experimental results. We discuss related work in Sect. 6.5. Finally, conclusions are presented in Sect. 6.6.

6.2 Multi-robot Exploration Using Semantic Information

In multi-robot exploration, a team of robots is distributed in a new environment with the objective of accumulating information to create a map. In this task, we are interested in finding good assignments of goal positions for the robots in the team. In our case, we assign target locations with the aim of minimizing the time needed to complete the exploration.

6.2.1 Classifying Target Locations

We assume that the knowledge about the environment is represented by an occupancy grid map. In this representation, target locations are found at the frontier between known and unknown areas [24]. As an example, Fig. 6.1(a) shows a map together with the frontiers detected there (dashed lines). For each of the frontiers, a target location is generated.

The goal now is to classify each potential target location into a semantic class. One possible solution to classify a target location is to simulate an observation at its position, and then classify this observation using the approach presented in Chap. 3. However, the target position is located at a frontier, which means that part of the neighboring areas are not known. Therefore, the laser observations simulated at frontier cells contain a significant number of maximum-range readings, which can lead to high missclassification rates.

To increase the classification rate in these cases, a short virtual trajectory to the desired goal location is generated. We then simulate laser range observations at different poses along this virtual trajectory using the partially known map. These poses

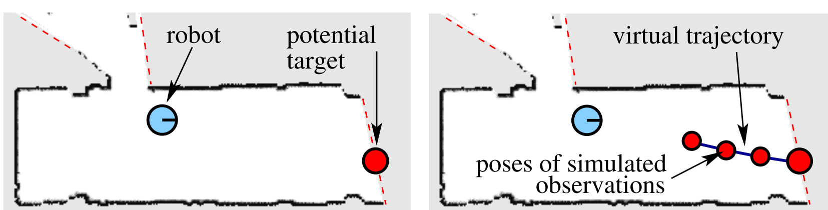

(a) Frontier extraction (b) Virtual trajectory

Fig. 6.1 (**a**) The image shows a situation in which a robot has extracted the frontiers of the occupancy grid map (*dashed lines*). Additionally, a target location is shown for one of the frontiers. (**b**) A virtual trajectory to the target is generated by the robot.

are generated selecting cells in the occupancy grid which are as far away as possible from the unknown locations in the current map using the euclidean distance transformation [14]. The reason for this selection is that cells having more information about its surroundings will have a lower error in their semantic classification, since their simulated range scans will contain fewer maximum-readings. Then an A* planner is used to generate the virtual trajectory to the target location. An example illustrating this process is shown in Fig. 6.1(b).

Once the virtual trajectory is created, we follow the approach presented in Chap. 5 to classify trajectories. This method applies a hidden Markov model (HMM) to maintain a posterior $bel(t_t)$ about the type y_t of the place the virtual sensor is located at as

$$bel(y_t) = \eta \, Pr(z_t \mid y_t) \sum_{y_{t-1}} Pr(y_t \mid y_{t-1}, u_{t-1}) bel(y_{t-1}). \tag{6.1}$$

The different components of this model are calculated in the same way as explained in Chap. 5. Using (6.1), we classify the target location $bel(y_{\text{target}})$ using the classification of the positions leading to it.

6.2.2 Target Assignment Using Semantic Place Labeling

As indicated above, the main idea is to give priority to target locations that are located in corridors, because they lead to a higher number of unknown areas. This is achieved by using the algorithm used for target assignment shown in Fig. 6.2.

- Determine the set of frontier cells.
- Compute for each robot i the cost V_t^i for reaching each frontier cell t.
- Assign to each frontier cell t a semantic labeling L_t.
- Set the utility U_t of all frontier cells t to $U_{init}(L_t, n)$ according to their semantic labeling L_t and the size n of the team.
- While there is one robot left without a target point

  1. Determine a robot i and a frontier cell t which satisfy

     $$(i,t) = \underset{(i',t')}{\text{argmax}} \left(U_{t'} - V_{t'}^{i'} \right).$$

  2. Reduce the utility of each target point t' in the visibility area according to

     $$U_{t'} \leftarrow U_{t'} - Pr_{vis}(t,t').$$

Fig. 6.2 The algorithm for the assignation of target locations to the different robots in a team.

The algorithm first calculates the set of frontier cells in the current submap. For each robot i in the team, the algorithm then calculates the cost V_t^i of reaching each frontier cell t. This cost is based on the distance the robot has to travel to reach the cell. Additionally, the algorithm estimates the semantic label y_t of the target location t using the HMM-based approach presented in Sect. 6.2.1.

Using the label y_t together with the number n of robots in a team, an initial utility function $U_{init}(y_t, n)$ is assigned to each target location t. In this function, the target locations classified as corridors get an initial utility U_{init} which is β times higher than other target locations. In our implementation, we selected a value of $\beta = 5$, since after several experiments we found that this value led to the best results in different runs of the algorithm.

The algorithm continues with an iterative process which combines two steps. First, the best combination of robot i and target t is selected. This selection is done maximizing the utility function U_t. A frontier cell is discounted as soon as it is selected. In this way, we avoid the situation in which several robots received the same frontier cell. We additionally discount target locations which can potentially be observed by other robots that already have a target assigned. This is done by using the utility function

$$U(t_n \mid t_1, \ldots, t_{n-1}) = U_{t_n} - \sum_{i=1}^{n-1} Pr_{vis}(t_n, t_i) \,, \tag{6.2}$$

with $Pr_{vis}(t_n, t_i)$ being the probability that the frontier t_n can be observed by a robot moving to frontier cell t_i. We approximate this probability density using a linear function.

The algorithm of Fig. 6.2 reduces the interference of robots during the exploration taking into account the visibility constraints, which are included in the utility function. Moreover, the inclusion of semantic information about the target locations improves the distribution of robots, giving preference to corridor places when selecting goal position for exploring unknown areas. As a result, the robots are better distributed and the time needed to explore the environment is significantly reduced.

The reduction of time during the exploration is not significant when the number of members in the team is small. This fact can be explained by considering the single-robot exploration scenario. To create a map of the environment, a single robot has to explore the whole environment, and it makes no sense to focus first on a particular place like a corridor. In our experiments, the exploration time does not decrease if the team has less than five robots.

6.3 Localization Using Place Recognition

A second problem which benefits from the semantic classification of places is the localization of a robot in an indoor environment. In this task the robot is given a representation of an indoor environment in the form of an occupancy grid map [4, 15]. Each cell in this grid additionally contains a semantic label about its place.

In our case the set of possible places to be recognized is composed of a corridor, doorways, offices, a kitchen, a seminar room, and a laboratory. This set of places is similar to the one used in Chap. 5.

The method applied to localize the robot in the environment is the popular Monte Carlo Localization (MCL) approach [3]. This localization method applies a recursive Bayesian scheme to maintain a posterior about the location of the robot x_t given the map m of the environment, the odometry information $u_{0:t-1}$, and the observations $z_{1:t}$ as

$$bel(x_t \mid m, z_{1:t}, u_{0:t-1}) = \eta \cdot Pr(z_t \mid m, x_t) \cdot Pr(x_t \mid m) \cdot \cdot \int_{x'} Pr(x_t \mid x', u_{t-1}) \cdot bel(x' \mid m, z_{1:t-1}, u_{0:t-2})\, dx'. \tag{6.3}$$

In MCL, the posterior about the robot positions is estimated using particle filters. The belief $bel(x_t)$ is represented by a set of random samples or particles in a non parametric form. This representation can approximate a broad set of distributions [22].

Our implementation of MCL is characterized by the observation model which uses place semantic information. As observations $z_{1:t}$, we use the output of the classifier the robot uses for place labeling. This classifier is the same as the one introduced in Sect. 5.2, and applies a probabilistic decision list in which each element contains a binary classifier. The quantity $Pr(z_t \mid m, x_t)$ is then determined as $Pr(z_t \mid y_t)$, where y_t is the class assigned to x_t in m. To estimate $Pr(z_t \mid y_t)$, we generated statistics about the output of the sequential multi-class classifier given the robot was at a place corresponding to y_t. Additionally, we weight the particles inversely proportional to the occupancy probability at x_t in m.

6.4 Experimental Results

To show the improvements obtained in the exploration and localization tasks when using semantic information about places we carried out several experiments using simulation and real robots. The approach for multi-robot exploration was implemented using teams with different number of robots. Due to the high numbers of robots used, we evaluated our technique only in simulation experiments. For the robot localization we used an ActivMedia Pioneer II robot with two lasers.

6.4.1 Multi-robot Exploration

The following experiments were designed to show the improvements of the multi-robot exploration technique making use of semantic place information. We evaluate the approach on simulation due to the large number of robots that form the teams. Further evaluations can be found in [21].

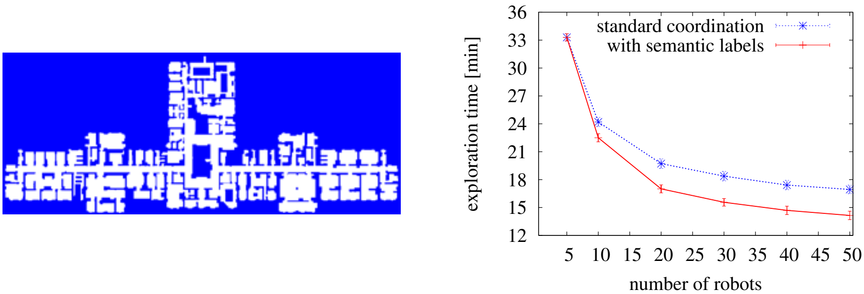

(a) Fort Sam Huston hospital

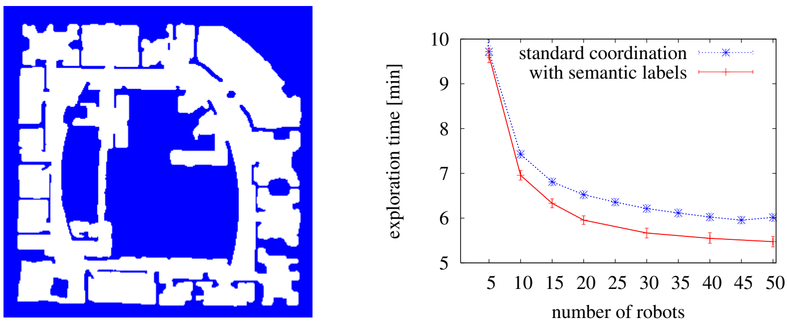

(b) Intel Research Lab

Fig. 6.3 (**a**) Map of the Fort Sam Huston hospital, and performance when the semantic information is taken into account in comparison to the case where no label information is used. (**b**) Map of the Intel research Lab with its corresponding performance plot.

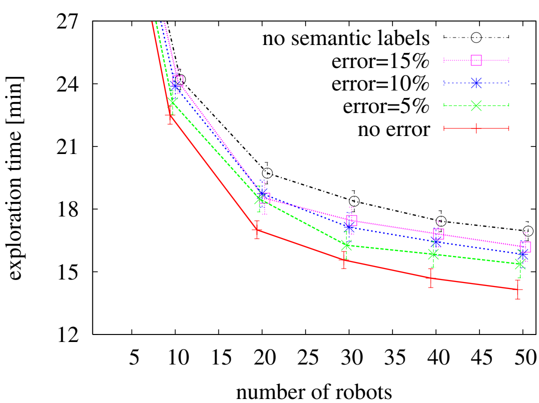

Fig. 6.4 Exploration performance with different classification errors.

In a first experiment we used the map of the Fort Sam Huston hospital (cf. Fig. 6.3(a)), which contains several corridors together with rooms adjacent to them. In the experiment we apply our method for coordinating several robots using

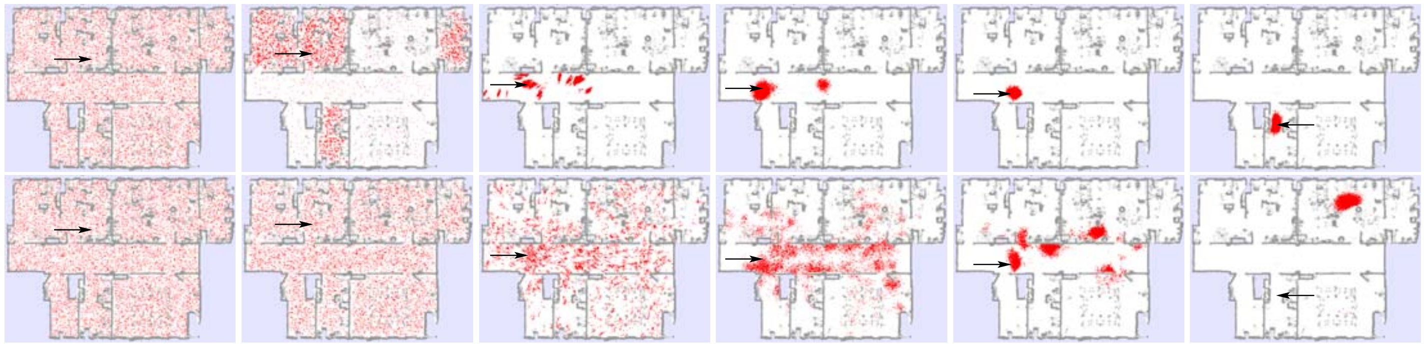

Fig. 6.5 Global localization using semantic information and odometry (first row) compared to an approach using only the odometry information (second row). The images in the same column represent the corresponding filter at the same time. The arrow indicates the ground truth position. As the results indicate, semantic information can be used to speed up global localization.

semantic information about places, and compared it to the case in which no place information is used. In addition, Fig. 6.3(a) shows the performance when the number of robots in the team varies from 5 to 50 members. For each team size, we repeated the experiments 50 times. In all the experiments the robots started from the same initial position. As the plot shows, the time needed to explore the complete environment is significantly reduced at the confidence level of 0.05 when using semantic place information. A similar experiment was carried out using the map of the Intel Research Lab. Again we observed a significant reduction in the exploration time as shown in Fig. 6.3(b).

In the previous experiments we assumed that the semantic classification of the target locations had no errors. In real situations, however, errors usually appear during the labeling process (see experimental results in Chaps. 3 and 5). To study the performance of the method according to the classification errors, we carried out an experiment in which we randomly misclassified different percentages of target locations and measure the exploration time according to them. Figure 6.4 shows the resulting performance using different team sizes. When the error in the classification exceeds 15% the improvement using semantic information is not significant anymore.

6.4.2 Localization

The last experiment is designed to illustrate how semantic information about places can be used to improve the localization of a mobile robot in its environment. In this experiment, we used an ActivMedia Pioneer II robot equipped with two laser range finders covering 360° around the robot (cf. Fig 3.8(a)). Note that the laser data is only fed into the semantic classifier and not used for metric localization.

We apply a Monte Carlo localization approach following the model in (6.3). Figure 6.5 illustrates the evolution of two particle sets over time. In the first row, the semantic information was available whereas in the second row only the odometry information was used. Both filters were initialized with a uniform distribution with $10,000$ particles. The robot initially was located in the second left office, north of the corridor. Therefore, particles located in offices received higher importance weights compared to the other samples. Whereas the approach utilizing semantic information converges quickly to the correct solution, the particle filter that relies only on the odometry information $Pr(x_t \mid m)$ finally diverges.

6.5 Related Work

Different aspects of multi-robot exploration have been studied in the past. For example, Yamauchi [23] presented a technique to learn maps with a team of mobile robots. In this approach, the robots exchange information about the map that is continuously updated whenever new sensor input arrives. Koenig et al. [12] analyze different terrain coverage methods for simple robots with limited sensing and computational capabilities.

In the work by Zlot et al. [25], a technique is presented in which the exploration is guided by a market economy. It considers sequences of potential target locations for each robot and tasks are traded between the robots using single-item first-price sealed-bid auctions. Singh and Fujimura [20] present a method for heterogeneous robot teams. In this approach, if a robot is too big to pass through a narrow passage, it informs other robots to do this task. Howard et al. [9] introduce an incremental deployment approach that explicitly deals with situations in which the path of one robot is blocked by another. Also Matarić and Sukhatme [13] introduce different strategies for allocating tasks in robot teams, and analyze their performance in different experiments. Finally, the Hungarian method to compute the assignments of frontier cells to robots is introduced by Ko et al. [11].

The coordination technique presented is this chapter is an extension of the work by Burgard et al. [2]. We also discount the utility of target locations if they are visible from a goal location already assigned to a robot. In contrast to [2], our approach estimates and incorporates background knowledge about environmental structure into the goal point assignment procedure.

The semantic labels used to improve multi-robot coordination can be seen as background knowledge about spacial structures. Fox et al. [6] presented a technique which learns background knowledge in typical indoor environments and later on use that knowledge for map building. They apply their approach to decide whether the robot is seeing a previously built portion of a map, or is exploring new terrain.

Localization is a typical problem in mobile robotics, and different approaches have been applied to solve it. The grid-based Monte Carlo localization was introduced by Simmons and Koenig [18]. This approach approximates the posterior of the robot pose using a histogram over the possible discrete poses. Several authors have successfully applied grid-based Monte Carlo localization in their work, as for example Burgard et al. [1], Hertzberg and Kirchner [8], and Simmons et al. [19]. Multi-hypothesis extended Kalman filters is another approach for localization used by different authors, as for instance Jensfelt and Kristensen [10], Roumeliotis and Bekey [17], and Reuter [16]. Finally, particle filter approaches were introduced by Dellaert et al. [3] and Fox et al. [5]. In this chapter we use this approach but additionally include the semantic classification of places.

For related work about the semantic labeling of places we refer the reader to Sects. 3.7 and 5.6.

6.6 Conclusion

In this chapter, we have shown how the semantic information helps to improve several robotic tasks. In particular, we proposed a technique that takes into account semantic information about places in the context of coordinated multi-robot exploration. The main idea is that mobile robots can use the intrinsic information of human-made environments to improve their actions. This improvement is obtained by selecting the best target locations according to their semantic classification. The semantic labeling of the target locations is done using an AdaBoost-based classifier.

Additionally, a hidden Markov model is applied to improve the classification in a virtual trajectory to the target position.

Alternatively we have seen how the semantic information about places can be used to localize the robot in an indoor environment using the Monte Carlo localization approach. In this case, the observation model of the robot corresponds to the semantic classification of its position. Experimental results indicate that this approach can be used to speed up global localization.

Both methods demonstrated that the semantic information can improve the performance in different tasks using autonomous mobile robots.

References

1. Burgard, W., Cremers, A.B., Fox, D., Hähnel, D., Lakemeyer, G., Schulz, D., Steiner, W., Thrun, S.: Experiences with an interactive museum tour-guide robot. Artificial Intelligence 114(1-2) (2000)
2. Burgard, W., Moors, M., Stacniss, C., Schneider, F.: Coordinated multi-robot exploration. IEEE Transactions on Robotics 21(3), 376–378 (2005)
3. Dellaert, F., Fox, D., Burgard, W., Thrun, S.: Monte carlo localization for mobile robots. In: Proceedings of the IEEE International Conference on Robotics and Automation (1999)
4. Elfes, A.: Using occupancy grids for mobile robot perception and navigation. Computer 22(6), 46–57 (1989)
5. Fox, D., Burgard, W., Dellaert, F., Thrun, S.: Monte carlo localization: Efficient position estimation for mobile robots. In: Proc. of the National Conference on Artificial Intelligence (1999)
6. Fox, D., Ko, J., Konolige, K., Stewart, B.: A hierarchical bayesian approach to the revisiting problem in mobile robot map building. In: Proceedings of the International Symposium of Robotics Research, Siena, Italy (2003)
7. Guzzoni, D., Cheyer, A., Julia, L., Konolige, K.: Many robots make short work. AI Magazine 18(1), 55–64 (1997)
8. Hertzberg, J., Kirchner, F.: Landmark-based autonomous navigation in sewerage pipes. In: Proceedings of the First Euromicro Workshop on Advanced Mobile Robots (1996)
9. Howard, A., Mataric, M.J., Sukhatme, G.S.: An incremental deployment algorithm for mobile robot teams. In: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, Lausanne, Switzerland, pp. 2849–2854 (2002)
10. Jensfelt, P., Kristensen, S.: Active global localisation for a mobile robot using multiple hypothesis tracking. IEEE Transactions on Robotics and Automation 17(5), 748–760 (2001)
11. Ko, J., Stewart, B., Fox, D., Konolige, K., Limketkai, B.: A practical, decision-theoretic approach to multi-robot mapping and exploration. In: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, Las Vegas, NV, USA, pp. 3232–3238 (2003)
12. Koenig, S., Szymanski, B., Liu, Y.: Efficient and inefficient ant coverage methods. Annals of Mathematics and Artificial Intelligence 31, 41–76 (2001)
13. Matarić, M.J., Sukhatme, G.S.: Task-allocation and coordination of multiple robots for planetary exploration. In: Proceedings of the International Conference on Advanced Robotics, Budapest, Hungary, pp. 61–70 (2001)

14. Meijster, A., Roerdink, J.B.T.M., Hesselink, W.H.: A general algorithm for computing distance transforms in linear time. In: Mathematical Morphology and its Applications to Image and Signal Processing, pp. 331–340. Kluwer Academic Publishers, Dordrecht (2000)
15. Moravec, H.P.: Sensor fusion in certainty grids for mobile robots. AI Magazine 9, 61–74 (1988)
16. Reuter, J.: Mobile robot self-localization using PDAB. In: Proceedings of the IEEE International Conference on Robotics and Automation, San Francisco, CA, USA, pp. 3512–3518 (2000)
17. Roumeliotis, S.I., Bekey, G.A.: Bayesian estimation and kalman filtering: a unified framework for mobile robot localization. In: Proceedings of the IEEE International Conference on Robotics and Automation, San Francisco, CA, USA, pp. 2985–2992 (2000)
18. Simmons, R., Koenig, S.: Probabilistic robot navigation in partially observable environments. In: Proceedings of the International Joint Conference on Artificial Intelligence, pp. 1080–1087 (1995)
19. Simmons, R., Lopez Fernandez, J., Goodwin, R., Koenig, S., O'Sullivan, J.: Lessons learned from Xavier. Robotics and Automation Magazine, 733–739 (2000)
20. Singh, K., Fujimura, K.: Map making by cooperating mobile robots. In: Proceedings of the IEEE International Conference on Robotics and Automation, Atlanta, GA, USA, pp. 254–259 (1993)
21. Stachniss, C., Mozos, O.M., Burgard, W.: Speeding-up multi-robot exploration by considering semantic place information. In: Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), Orlando, FL, USA, pp. 1692–1697 (2006)
22. Thrun, S., Burgard, W., Fox, D.: Probabilistic Robotics. MIT Press, Cambridge (2005)
23. Yamauchi, B.: Frontier-based exploration using multiple robots. In: Proceedings of the International Conference on Autonomous Agents, Minneapolis, MN, USA, pp. 47–53 (1998)
24. Yamauchi, B., Schultz, A., Adams, W.: Integrating exploration and localization for mobile robots. Adaptive Behavior 7(2), 217–229 (1999)
25. Zlot, R., Stentz, A., Dias, M.B., Thayer, S.: Multi-robot exploration controlled by a market economy. In: Proceedings of the IEEE International Conference on Robotics and Automation, Washington, DC, USA (2002)