

Speeding-Up Multi-Robot Exploration by Considering Semantic Place Information

Cyrill Stachniss Óscar Martínez Mozos Wolfram Burgard

University of Freiburg, Department of Computer Science, D-79110 Freiburg, Germany

Abstract—In this paper, we consider the problem of exploring an unknown environment with a team of mobile robots. One of the key issues in multi-robot exploration is how to assign target locations to the individual robots. To better distribute the robots over the environment and to avoid redundant work, we take into account the type of place a potential target is located in (e.g., a corridor or a room). To determine the type of a place, we apply a classifier learned with AdaBoost which additionally considers spatial dependencies between nearby locations. Our approach to incorporate the type of places in the coordination of the robots has been implemented and tested in different environments. The experiments demonstrate that our system effectively distributes the robots over the environment and allows them to accomplish their mission faster compared to approaches that ignore the semantic place labels.

I. I

The problem of exploring an environment belongs to the fundamental problems in mobile robotics. There are several applications like planetary exploration [2], reconnaissance [9], or cleaning [11] in which the complete coverage of a terrain belongs to the integral parts of a robotic mission. The use of multiple robots is often suggested to have several advantages over single robot systems [4, 8], since cooperating robots have the potential to accomplish a task faster than a single robot. However, if robots operate in teams there is the risk of interferences between them. Several exploration techniques dealing with the problem of appropriate collaboration between robots were presented in the past [3, 12, 20, 24, 26].

A popular approach to exploration can be summarized in the following way. First, a set of potential target locations or target areas is determined. Second, the robots choose target locations and then include their observations obtained along the paths to the targets into a map. This process is repeated, until the environment has been fully covered with the sensors of the robots. In the context of multi-robot exploration, it is important to achieve a collaboration behavior so that the robots avoid traveling unnecessary long distances, avoid doing redundant work, and avoid interference with other team-mates.

Indoor environments constructed by humans often contain certain structures like corridors with adjacent rooms or offices. However, it is mainly unexplored how robots can utilize such background information to more efficiently solve the exploration task. One of our observations is that the more unexplored target locations are known when assigning targets to robots, the faster the team can explore the environment. This is due to the fact that the robots can be better distributed over the environment. In this way, the amount of redundant

work is reduced and interferences occur less likely. It therefore makes sense to focus on areas first which are likely to provide a large number of new target locations in order to obtain a better assignment of targets to robots.

The contribution of this paper is a technique to estimate and utilize semantic information during collaborative multi-robot exploration. In our approach, the robots get a higher reward for exploring corridors since they typically provide more branchings to unexplored areas like adjacent rooms compared to rooms itself. This is especially useful in case of large robot teams, because if more target locations are available the robots can be better distributed over the environment. As a result, the overall completion time of an exploration task can be significantly reduced.

II. R W

The various aspects of the problem of exploring unknown environments with mobile robots have been studied intensively in the past. For example, Yamauchi [24] 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. To acquire knowledge about the environment, all robots move to the closest frontier cell. Koenig *et al.* [14] analyze different terrain coverage methods for ants which are simple robots with limited sensing and computational capabilities. Furthermore, there has been research on how to deal with limited communication in the context of multi-robot exploration [3, 20].

One approach towards cooperation between heterogeneous robot systems has been presented by Singh and Fujimura [23]. If a robot is too big to pass through a narrow passage, it informs other robots about this task. Howard *et al.* [10] presented an incremental deployment approach that explicitly deals with obstructions, i.e., situations in which the path of one robot is blocked by another. Zlot and colleagues [26] proposed an architecture for mobile robot teams in which the exploration is guided by a market economy. They consider sequences of potential target locations for each robot and trade tasks between the robots using single-item first-price sealed-bid auctions. Such auction-based techniques have also been applied by Gerkey and Mataric [7] to efficiently solve the task allocation problem with a group of robots.

Mataric and Sukhatme [16] consider different strategies for task allocation in robot teams and analyze the performance of the team in extensive experiments. Parker [19] described a project to perform reconnaissance and surveillance task. It

investigates how to jointly accomplish a task with heterogeneous robots that cannot be solved by a robot individually. Ko *et al.* [12] present an approach that uses the Hungarian method to compute the assignments of frontier cells to robots. In contrast to our work, Ko *et al.* mainly focuses on finding a common frame of reference in case the start locations of the robots are not known.

The coordination technique presented in this paper is an extension of our previous work [3]. We also discount the utility of target locations if they are visible from a goal location already assigned to a robot. In contrast to [3], the approach presented in this paper estimates and incorporates background knowledge about environmental structure into the goal point assignment procedure.

In order to improve the navigation, we use semantic place labels learned from sensor data. A series of authors addressed the problem of learning environmental structures with mobile robots. For example, Koenig and Simmons [13] use a pre-programmed routine to detect doorways from range data. Althaus and Christensen [1] use line features to detect corridors and doorways. Some authors also apply learning techniques to localize the robot or to identify distinctive states in the environment. For example, Oore *et al.* [18] train a neural network to estimate the location of a mobile robot in its environment using the odometry information and ultrasound data. In our work, we apply a technique originally proposed by Martínez Mozos *et al.* [15]. This technique uses simple features extracted from laser range scans to train a set of classifiers and in this way are able to label a place given a single 2d laser range observation. Furthermore, our smoothing technique bears resemblance with our previous approach [21], in which a hidden Markov model is applied to improve the classification result. In contrast to the work presented here, we combined in [21] laser data and visual information to obtain more features and in this way are able to distinguish between more classes.

The semantic labels used to improve multi-robot coordination can be seen as background knowledge about spacial structures. Fox *et al.* [5] presented a technique which aims to learn 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.

Due to the best of our knowledge, there is no work that investigates how semantic information about places in the environment can be used to optimize the collaboration behavior of a team of robots. The contribution of this paper is an approach that estimates and explicitly uses semantic information in order to more efficiently spread the robots over the environment. This results in an more balanced target location assignment with less interferences between robots. As a result, the overall time needed to cover the whole environment can be significantly reduced.

III. S P L

This section explains, how semantic place labels can be obtained with mobile robots based on laser range observations.

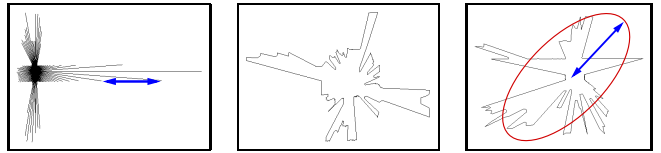


Fig. 1. Examples for features generated from three different laser scans, namely the average distance between two consecutive beams, the perimeter of the area covered by a scan, and the mayor axis of the ellipse that approximates the polygon described by the scan.

The goal is to learn a classifier, that is able to distinguish corridors from other kinds of indoor structure. To obtain such a classifier, we apply the AdaBoost algorithm introduced by Freund and Schapire [6].

The key idea of AdaBoost is to form a strong binary classifier given a set of weak classifiers. Thereby, the weak classifiers h_j only need to be better than random guessing and are constructed using simple, single-value features $f_j \in \mathbb{R}$

$$h_j(x) = \begin{cases} 1 & \text{if } p_j \cdot f_j(x) < p_j \cdot \theta_j \\ 0 & \text{otherwise,} \end{cases} \quad (1)$$

where x is an example, θ_j is a threshold value, and p_j is either -1 or $+1$ and thus represents the direction of the inequality. The AdaBoost algorithm determines for each weak classifier h_j the optimal parameter tuple (θ_j, p_j) , such that the number of misclassified training examples is minimized.

In our approach, the features f_j are directly extracted from observations. Examples for features extracted from laser range data are depicted in Figure 1. Such features are the average distance between consecutive beams, the area covered by a range scan, or the perimeter of that area. Most of the features are standard geometrical features used in shape analysis [22]. The full list of features is provided in [15].

The input to the AdaBoost algorithm is a set of labeled, positive and negative training examples. In our case, this is a set of laser-range observations recorded in a corridor and a second set taken outside corridors. In a series of T rounds, the algorithm repeatedly selects a weak classifier h_j based on a distribution D over the training examples. This distribution specifies an importance weight for each example in the current round. The selected weak classifier is expected to have a small classification error on the training data. The idea of the algorithm is to modify the distribution D by increasing the weights of the most difficult training examples in each round. The final strong classifier H is a weighted majority vote of the best T weak classifiers

$$H(x) = \sum_{t=1}^T w_t \cdot h_t(x). \quad (2)$$

In our system, the resulting strong classifier takes as input a single 360 degree laser range scan recorded by a robot and is able to determine whether or not the position from which the scan was taken belongs to the class *corridor*.

IV. E L G L

The idea described in the previous section is well-suited to determine the type of place the robot is currently in. In the context of exploration, however, we are interested in

classifying potential targets of the robot. Typically, target locations are located at the frontier between known and unknown areas. According to our grid-based representation, such frontier cells [25] can be easily extracted. In our approach, we generate a potential target location for each group of frontier cells lying on the same frontier. This procedure is repeated for each frontier. As an example, the left image of Figure 2 depicts a potential target location. It was generated from the frontier in the corridor (the targets for the other two frontiers are not shown in that image).

One solution to classify a place which is not the current pose of the robot is to simulate a laser range scan at this place given the (partial) map constructed so far. However, since large neighboring areas of frontier cells have not been observed so far, classifying such a frontier cells with the approach presented in the previous section leads to a high misclassification rate. In the following, we therefore introduce a HMM-based technique that takes into account spacial dependencies between nearby locations in order to obtain a lower error rate for places like frontier cells.

Due to the structure of man-made environments, the semantic class does not change randomly between nearby poses. Therefore, it makes sense to consider a smoothing between places located close together [21]. To do so, we generate a short virtual trajectory to the desired goal location. We simulate laser range observations within the partially know map along the virtual trajectory. Whenever the ray-casting operation which is used to simulate a beam reaches an unknown cell in the grid map, the virtual sensor reports a maximum-range reading. We then apply a hidden Markov model (HMM) and maintain a posterior $Bel(L_x)$ about the type L_x of the place x the virtual sensor is currently at

$$Bel(L_x) = \alpha \cdot P(c_x | L_x) \cdot \sum_{L_{x'}} P(L_x | L_{x'}) \cdot Bel(L_{x'}). \quad (3)$$

In this equation, c_x is the result of the classifier learned with AdaBoost for the place x and α is a normalizing constant ensuring that the left-hand side sums up to one over all semantic labels.

To implement this HMM, three components need to be known. First, we need to specify the observation model $P(c_x | L_x)$ which is the likelihood that the classification output is c_x given the actual class is L_x . The observation model is learned based on 5.000 observations, recorded in different environments combined with the corresponding manually created ground truth labeling.

Second, we need to specify the transition model $P(L_x | L_{x'})$ which defines the probability that the virtual sensor moves from class $L_{x'}$ to class L_x . To determine the motion model, we evaluated typical trajectories obtained during exploration. We can directly compute $P(L_x | L_{x'})$ by counting the transitions between places, which have been manually labeled.

Furthermore, we need to specify how the belief $Bel(L_{start})$ is initialized. In our current system, we choose a uniform distribution, which means that all classes (here corridor and non-corridor) have the same likelihood.

Finally, we have to describe how the virtual trajectory is generated. The endpoint of the trajectory is the frontier cell to

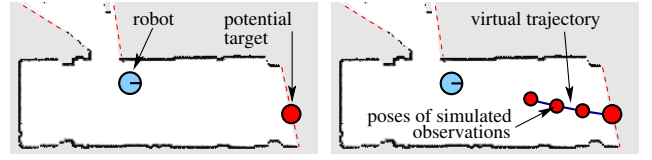


Fig. 2. This figure illustrates the generation of the virtual trajectory used for the HMM smoothing. The left image depicts the current location of the robot, the frontiers (dashed lines), and a potential target location to be evaluated. To do so, the robot generates a virtual trajectory as shown in the right image and simulates observations at several positions located on the trajectory. These sequence of observations is used as the input of the HMM in order to obtain a more robust classification result.

be classified. Since locations which have less unknown grid cells in their surroundings can typically be classified with a higher success rate, the other positions on that trajectory should be as far away from the unknown locations as possible. Therefore, we apply the euclidian distance transformation [17] with respect to unknown and occupied cells in the local area of the frontier. We select the pose in the free space within that local area with the highest distance to unknown areas. Then an A* planner is used to generate the virtual trajectory to the target location. An illustrating example is depicted in Figure 2.

V. U	S	P	I	E
	M	-R	E	

The goal of collaborative multi-robot tasks is to share the load between the members of a team in order to accomplish the task faster. As discussed in the related work section, different approaches exist that assign target locations to robots using job-shop-scheduling techniques, bidding algorithms, or decision theoretic approaches. In the approach described here, we discount frontiers based on visibility constraints as in [3]. The approach works in a centralized fashion but can also deal with limited communication. Typically, one robot calculates the assignments. In case the whole team splits up into several teams due to the restricted communication range, one member of each sub-team becomes a leader and executes the target assignment procedure.

For each robot i in a team, the algorithm computes the cost V_i^t to each target location t based on the distance to be traveled in order to reach that location. To avoid that several robots focus on the same frontier, each target location is discounted after being assigned to one robot. In this way, the robots get distributed over the environment and do not focus on the same local area. Additionally, target locations which can potentially be observed by other robots already assigned are discounted. This is done by introducing a utility function $U(t)$ given by

$$U(t_n | t_1, \dots, t_{n-1}) = U_{t_n} - \sum_{i=1}^{n-1} P_{vis}(t_n, t_i), \quad (4)$$

where $P_{vis}(t_n, t_i)$ describes the probability that the frontier t_n can be observed by a robot moving to t_i . In our approach, this probability density is approximated by a linear function.

To determine appropriate target points for all robots, we apply an iterative approach. In each round, the tuple (i, t) , where i is a robot and t a frontier cell, with the best overall

evaluation $U_t - V_t^i$ is chosen. One then recomputes the utilities of all frontier cells according to Eq. (4) given the new and all previous assignments. Then the process is repeated for the remaining robots.

The knowledge about the semantic labels is integrated into the utility function. All places which are supposed to provide several branchings to adjacent places are initialized with a high utility. In our current implementation, all corridor locations get a γ times higher initial utility (U_{init}) compared to all other potential target locations. In this way, the robots prefer targets in corridors and eventually make slight detours in order to explore them first. To determine the actual value of γ , we performed exploration runs in different environments with varying γ . We figured out that we obtained the best results using a γ -value of around 5. Algorithm 1 depicts the resulting coordination technique used in our current system.

Algorithm 1 Target Assignment Algorithm Using Semantic Place Labels.

- 1: Determine the set of frontier cells.
 - 2: Compute for each robot i the cost V_t^i for reaching each frontier cell t .
 - 3: Estimate for each frontier cell t the semantic labeling L_t (according to Section IV).
 - 4: 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 (see text below).
 - 5: **while** there is one robot left without a target point **do**
 - 6: Determine a robot i and a frontier cell t which satisfy:
 $(i, t) = \operatorname{argmax}_{(i, t)} (U_t - V_t^i)$.
 - 7: Reduce the utility of each target point t' in the visibility area according to $U_{t'} \leftarrow U_{t'} - P_{vis}(t, t')$.
 - 8: **end while**
-

Our approach distributes the robots in a highly efficient manner over the environment and reduces the amount of redundant work by taking into account visibility constraints between targets and their semantic labels. The labels are used to focus the exploration on unexplored corridors, because they typically provide more branchings to adjacent rooms than other places. The high number of branchings results in a higher number of potential target locations that are available in the assignment process. This typically leads to a more balanced distribution of robots over the environment. As we will demonstrate in the experiments, the integration of such semantic labels helps to reduce the overall exploration time of multi-robot exploration approaches for large robot teams.

Please note that for very small teams of robots we do not achieve a reduction of the exploration time using our technique. This fact can be explained by considering the single-robot exploration scenario. In this case, it makes no sense to focus on exploring the corridors first, since the robot has to cover the overall environment with its sensor. Moving through the corridors first will in general lead to an increased trajectory length and in this way will increase the overall exploration time. We observed this effect for robot teams smaller than five robots.

To prevent a loss of performance compared to approaches



Fig. 3. Maps of the Fort Sam Huston hospital and the Intel Research Lab.

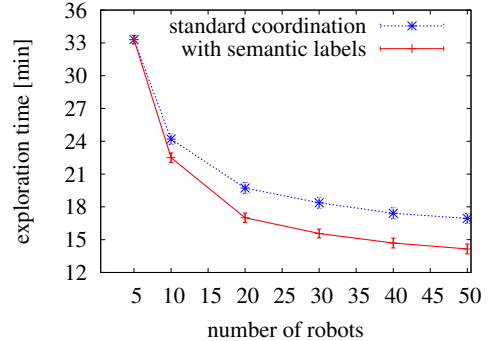


Fig. 4. Coordination results obtained in the Fort Sam Huston hospital map employing the coordination strategy with and without the use of semantic place labels.

which do not consider semantic place information for small robot teams, we trigger the influence of the semantic place information depending on the size of the team. We linearly decrease the influence γ for teams smaller than 10 robots. The linear interpolation of the influence of the semantic labels is encoded in the utility function $U_{init}(L_t, n)$, where n denotes the number of robots, in Algorithm 1.

VI. E

This section is designed to evaluate the improvements of our multi-robot coordination technique which makes use of semantic place information. Due to this big numbers of robots, we evaluated our collaboration technique only in simulation experiments.

A. Performance Improvement using Semantic Place Information

The first experiment has been carried out in the map of the Fort Sam Huston hospital, which is depicted in the left image of Figure 3. This environment contains a long horizontal corridor, vertical corridors, and several rooms adjacent to the corridors.

We varied the size of the robot team from 5 to 50 robots and applied the coordination technique with and without taking into account semantic information about places. Like in all our experiments, the group of robots started from the same initial position which was chosen randomly for the individual runs. For each setting, we carried out 50 runs. Figure 4 depicts the result of the exploration experiment by plotting the exploration time versus the number of robots. The error bars in that plot indicate the 0.05 confidence level. As can be seen, our technique significantly outperforms the collaboration scheme that does not consider the place information. This significant reduction of exploration time is due to the fact that the robots

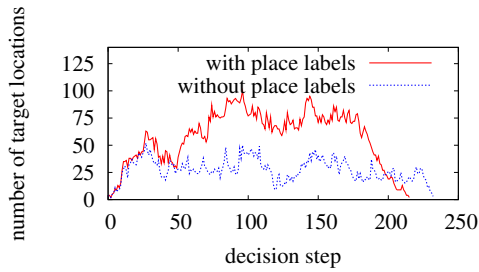


Fig. 5. The number of potential target locations at the different decision steps during exploration.

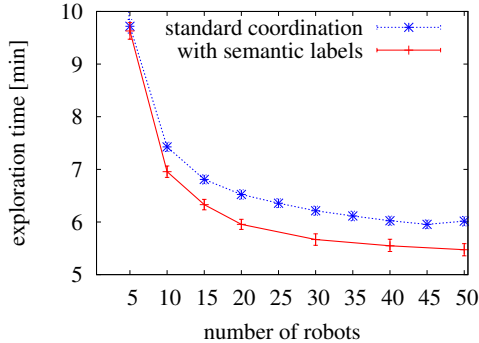


Fig. 6. Results obtained in the Intel Research Lab.

focus on exploring the corridors first. As a result, a big number of frontiers typically emerges due to numerous adjacent rooms. Especially in the context of large teams, this results in a better distribution of robots over the environment and thus speeds up the exploration process. This effect can be observed in Figure 5. The graphs plot the number of potential target locations over time during an exploration task carried out using the Fort Sam Houston map. Using our approach, more target locations are available in the decision process most of the time. This leads to a better assignment of target locations to robots and as a result the amount of redundant work is reduced.

Furthermore, we observed a reduction of interferences between robots when they plan their paths through the environment. In our simulator, interferences result in a reduced travel speed, since the robots often block paths of other robots. Therefore, reducing the number of interferences allows the robots to accomplish their task faster. In our experiments, we observed a reduction of robot-robot interferences of up to 20%.

We performed similar experiments in different environments, like for example in the Intel Research Lab depicted in the right image of Figure 3. The result is comparable to the previous experiment and again the knowledge about the semantic categories of places allows the robots to complete the exploration task more efficiently. The actual evolution of the exploration time in this experiment is depicted in Figure 6.

B. Influence of Noise in the Semantic Place Information

In the experiments presented above, we assumed that the robots are able to correctly classify the different target locations into the semantic categories. This assumption, however, is typically not justified. In this experiment, we evaluate the performance of our approach for different classification error rates. We evaluated the exploration time for a classifier

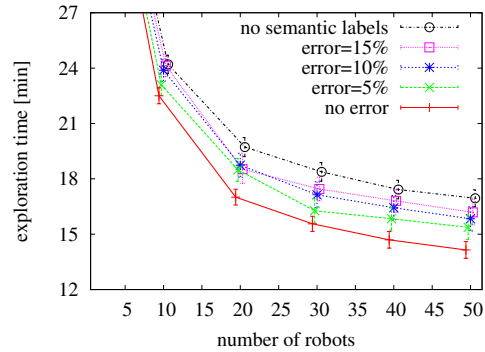


Fig. 7. Exploration results with wrongly labeled places.

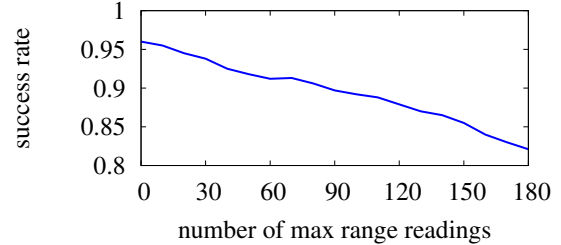


Fig. 8. This plot illustrates the classification performance of the standard classifier depending on how many consecutive beams of a 360 degree observation (1 degree angular resolution) are maximum range readings.

which randomly misclassified 5%, 10%, and 15% of the places. Figure 7 depicts a plot comparing the different error rates. As can be seen, even at a high error of 10%, our approach significantly outperforms the coordination technique that ignores the semantic information. When the error of the classification exceeds 15%, the exploration time is still reduced, although this result is not significant anymore.

C. Improvements of the HMM Smoothing and Error Analysis of the Classifier

In this section, we want to analyze the actual error of our place classification system and illustrate the improvements of the HMM smoothing. To do so, we labeled an environment, trained a corridor classifier using AdaBoost, and used a test set to evaluate the success rate. Whenever a single full 360 degree laser range scan was available, we obtained accurate classification results in different office environments. In this case, the error-rate was typically between 2% and 4%.

Figure 8 depicts the result of our classifier depending on the number of invalid readings caused by unknown grid cells close to frontiers. The x-axis shows the size of a continuous block of maximum range measurements (with an angular resolution of the laser of 1 degree). As can be seen, if only half of the observations are available, the classification error rate is between 18% and 19%.

In the final experiment, we determined the success rate of our HMM based smoothing method to determine the semantic labels. First, we determined the success rate without the HMM smoothing. In this case, the average classification rate was 81.2%. By considering the exploration speed-up depending on the classification rate depicted in Figure 7, such a high error rate is not sufficient to obtain a significant improvement.

Second, we applied our HMM-based smoothing approach that generates virtual trajectories towards the frontier and in this way incorporates the spatial dependencies between nearby locations. As a result, we obtained an average success rate of 92.8%. This is a good result considering that we obtained an average success rate of 96.2% (see Figure 8) if all observations are perfectly known. This fact illustrates that the HMM is a useful tool to improve the place labeling especially if not the full 360 degree range scan is available. It allows us to estimate the semantic labels with a comparably low error rate.

In sum, our experiments demonstrate that semantic place information can significantly reduce the exploration time even under larger classification errors.

VII. C

In this paper, we proposed a novel technique that takes into account semantic information about places in the context of coordinated multi-robot exploration. Since indoor environments are made by humans, they typically consist of structures like corridors and rooms. The knowledge about the type of place of potential target locations allows us to better distribute teams of robots over the environment and to reduce redundant work as well as the risk of interference between the robots. The semantic labels are determined by learning a classifier using AdaBoost in combination with an HMM to consider spatial dependencies.

Our approach has been implemented and tested in extensive simulation runs with up to 50 robots. Experiments presented in this paper illustrate that a team of robots using our approach can complete their exploration mission in a significantly shorter period of time. Furthermore, we believe that our technique for utilizing semantic information during exploration is not restricted to our exploration method and that it can be readily integrated into other, state-of-the-art coordination approaches.

In future work, we plan to learn the place labels in an unsupervised fashion. In this way, the system might be able to determine on its own what kind of spacial structures are useful for coordinated exploration and does not rely on manually defined labels.

A

This work has partly been supported by the DFG under contract number SFB/TR-8 (A3) and by the EC under contract number FP6-004250-CoSy. The Fort Sam Houston map was obtained from Radish, thanks go to Richard Vaughan for providing this data. Furthermore, the authors would like to thank Dirk Haehnel for providing the Intel Research Lab dataset.

R

- [1] P. Althaus and H.I. Christensen. Behaviour coordination in structured environments. *Advanced Robotics*, 17(7):657–674, 2003.
- [2] D. Apostolopoulos, L. Pedersen, B. Shamah, K. Shillcutt, M.D. Wagner, and W.R.L. Whittaker. Robotic antarctic meteorite search: Outcomes. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, pages 4174–4179, Seoul, Korea, 2001.

- [3] W. Burgard, M. Moors, C. Stachniss, and F. Schneider. Coordinated multi-robot exploration. *IEEE Transactions on Robotics*, 21(3):376–378, 2005.
- [4] G. Dudek, M. Jenkin, E. Milius, and D. Wilkes. A taxonomy for multi-agent robotics. *Journal of Autonomous Robots*, 3(4):375–397, 1996.
- [5] D. Fox, J. Ko, K. Konolige, and B. Stewart. A hierarchical bayesian approach to the revisiting problem in mobile robot map building. In *Proc. of the Int. Symposium of Robotics Research (ISRR)*, Siena, Italy, 2003.
- [6] Y. Freund and R.E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 55(1):119–139, 1997.
- [7] B.P. Gerkey and M.J. Mataric. Sold!: Auction methods for multirobot coordination. *IEEE Transactions on Robotics and Automation*, 18(5):758–768, 2002.
- [8] D. Guzzoni, A. Cheyer, L. Julia, and K. Konolige. Many robots make short work. *AI Magazine*, 18(1):55–64, 1997.
- [9] D.F. Houghton, S. Benjaafar, J.C. Bonney, J.R. Budenske, M. Dvorak, M. Gini, H. French, D.G. Krantz, P.Y. Li, F. Malver, B. Nelson, N. Papanikolopoulos, P.E. Rybski, S.A. Stoeter, R. Voyles, and K.B. Yesin. A miniature robotic system for reconnaissance and surveillance. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, pages 501–507, San Francisco, CA, USA, 2000.
- [10] A. Howard, M.J. Mataric, and S. Sukhatme. An incremental deployment algorithm for mobile robot teams. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, pages 2849–2854, Lausanne, Switzerland, 2002.
- [11] M. Jäger and B. Nebel. Dynamic decentralized area partitioning for cooperating cleaning robots. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, pages 3577–3582, Washington, DC, USA, 2002.
- [12] J. Ko, B. Stewart, D. Fox, K. Konolige, and B. Limketkai. A practical, decision-theoretic approach to multi-robot mapping and exploration. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, pages 3232–3238, Las Vegas, NV, USA, 2003.
- [13] S. Koenig and R. Simmons. Xavier: A robot navigation architecture based on partially observable markov decision process models. In D. Kortenkamp, R. Bonasso, and R. Murphy, editors, *Artificial Intelligence Based Mobile Robotics: Case Studies of Successful Robot Systems*, pages 91–122. MIT Press, 1998.
- [14] S. Koenig, B. Szymanski, and Y. Liu. Efficient and inefficient ant coverage methods. *Annals of Mathematics and Artificial Intelligence*, 31:41–76, 2001.
- [15] O. Martínez Mozos, C. Stachniss, and W. Burgard. Supervised learning of places from range data using adaboost. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2005.
- [16] M.J. Mataric and G. Sukhatme. Task-allocation and coordination of multiple robots for planetary exploration. In *Proc. of the Int. Conf. on Advanced Robotics (ICAR)*, pages 61–70, Budapest, Hungary, 2001.
- [17] A. Meijster, J.B.T.M. Roerdink, and W.H. Hesselink. *Mathematical Morphology and its Applications to Image and Signal Processing*, chapter A General Algorithm for Computing Distance Transforms in Linear Time, pages 331–340. Kluwer Academic Publishers, 2000.
- [18] S. Oore, G.E. Hinton, and G. Dudek. A mobile robot that learns its place. *Neural Computation*, 9(3):683–699, 1997.
- [19] L.E. Parker. The effect of heterogeneity in teams of 100+ mobile robots. In *Multi-Robot Systems Volume II: From Swarms to Intelligent Automata*. Kluwer Academic Publishers, Boston, 2003.
- [20] I. Rekleitis, V. Lee-Shue, A. Peng New, and H. Choset. Limited communication, multi-robot team based coverage. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, pages 3462–3468, New Orleans, LA, USA, 2004.
- [21] A. Rottmann, O. Martínez Mozos, C. Stachniss, and W. Burgard. Place classification of indoor environments with mobile robots using boosting. In *Proc. of the National Conference on Artificial Intelligence (AAAI)*, pages 1306–1311, Pittsburgh, PA, USA, 2005.
- [22] J.C. Russ. *The Image Processing Handbook*. CRC Press, 1992.
- [23] K. Singh and K. Fujimura. Map making by cooperating mobile robots. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, pages 254–259, Atlanta, GA, USA, 1993.
- [24] B. Yamauchi. Frontier-based exploration using multiple robots. In *Proc. of the Second International Conference on Autonomous Agents*, pages 47–53, Minneapolis, MN, USA, 1998.
- [25] B. Yamauchi, A. Schultz, and W. Adams. Integrating exploration and localization for mobile robots. *Adaptive Behavior*, 7(2):217–229, 1999.
- [26] R. Zlot, A.T. Stenz, M.B. Dias, and S. Thayer. Multi-robot exploration controlled by a market economy. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, Washington, DC, USA, 2002.