

Exploring Unknown Environments with Mobile Robots using Coverage Maps

Cyrril Stachniss Wolfram Burgard

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

Abstract

In this paper we introduce coverage maps as a new way of representing the environment of a mobile robot. Coverage maps store for each cell of a given grid a posterior about the amount the corresponding cell is covered by an obstacle. Using this representation a mobile robot can more accurately reason about its uncertainty in the map of the environment than with standard occupancy grids. We present a model for proximity sensors designed to update coverage maps upon sensory input. We also describe how coverage maps can be used to formulate a decision-theoretic approach for mobile robot exploration. We present experiments carried out with real robots in which accurate maps are build from noisy ultrasound data. Finally, we present a comparison of different view-point selection strategies for mobile robot exploration.

1 Introduction

Generating maps is one of the fundamental tasks of mobile robots and many researchers have focused on the problem of how to represent the environment as well as how to acquire models using this representation [5, 9, 10, 14, 17]. The mapping problem itself has several aspects that have been studied intensively in the past. Some of the most important aspects are the localization of the vehicle during mapping, appropriate models of the environment and the sensors, as well as strategies for guiding the vehicle. In literature, the localization problem plays an important role, since the quality of the resulting map depends strongly on the accuracy of the pose estimates during the mapping process. However, the accuracy of the map also depends on the choice of view-points during exploration. Especially if noisy sensors are used, the map will be quite inaccurate in areas which have been sensed a few times only or maybe even from disadvantageous view-points.

Exploration is the task of guiding a vehicle in such a way that it covers the environment with its sensors. Efficient exploration strategies are also relevant for surface inspection,

mine sweeping, or surveillance [3, 12]. In the past, several strategies for exploration have been developed. One group of approaches deals with the problem of simultaneous localization and mapping [1, 4], an aspect that we do not address in this paper. A common technique for exploration strategies is to extract frontiers between known and unknown areas [2, 7, 20] and to visit the nearest unexplored place. These approaches only distinguish between scanned and un-scanned areas and do not take into account the actual information gathered at each view-point. To overcome this limitation, Gonzales et al. [8] determine the amount of unseen area that might be visible to the robot from possible view-points. To incorporate the uncertainty of the robot about the state of the environment Moorehead et al. [13] as well as Bourgault et al. [1] use occupancy grids [14] and compute the entropy of each cell in the grid to determine the utility of scanning from a certain location. Whaite and Ferrie [19] present an approach that also uses the entropy to measure the uncertainty in the geometric structure of objects that are scanned with a laser range sensor. In contrast to the work described here they use a parametric representation of the objects to be scanned. Edlinger and Puttkamer [7] developed a hierarchical exploration strategy for office environments. Their approach first explores rooms and then traverses through doorways to explore other parts of the environment. Taylor and Kriegman [16] describe a system for visiting all landmarks in the environment of the robot. Their robot maintains a list of unvisited landmarks that are approached and mapped by the robot. Dudek et al. [6] propose a strategy for exploring an unknown graph-like environment. Their algorithm does not consider distance metrics and is designed for robots with very limited perceptual capabilities. Recently Koenig has shown, that a strategy, which guides the vehicle to the closest point that has not been covered yet, keeps the traveled distance reasonably small [11]. However, as experiments reported in this paper illustrate, such techniques can lead to a serious increase of measurements necessary to build an accurate map if the robot is not able to incorporate measurements on-the-fly while it is moving. This might be the case, for example, for robots extracting distance information from camera images.

In this paper we introduce coverage maps as a new probabilistic way to represent the belief of the robot about the state of the environment. In contrast to occupancy grids [14], in which each cell is considered as either occupied or free, cov-

*This work has partly been supported by the German Science Foundation (DFG) under contract number SFB/TR8-03.




Figure 1: Typical occupancy map obtained in situations in which cells are only partly occupied (left image) and a coverage map containing the corresponding coverage values (right image).

erage maps represent in each cell of a given discretization a posterior about the amount this cell is covered by an object. As an example consider the situation depicted in Figure 1 in which a cell is partly covered by an obstacle. With the standard occupancy algorithm the probability that this cell is occupied will converge to 1 if the sensors of the robot repeatedly detect the obstacle. The left picture of this figure shows the resulting occupancy probabilities (black represents high likelihood). Since the object does only cover 20% of this cell, a coverage value of .2 (as shown in the right image of Figure 1) would be a better approximation of the true situation. In addition to the representation aspect, we also present a sensor model that allows the robot to appropriately update a coverage map upon sensory input and describe how coverage maps can be used to realize a decision-theoretic approach to exploration of unknown environments.

This paper is organized as follows. In the next section we introduce coverage maps. In Section 3 we present a probabilistic technique to update a given coverage map upon sensory input. In Section 4 we describe a decision-theoretic approach to exploration based on coverage maps. In Section 5 we present experiments illustrating that our approach allows a mobile robot can learn accurate maps from noisy range sensors. Additionally, we present experiments comparing different view-point selection strategies for exploration.

2 Coverage Maps

As already mentioned above, occupancy grids rest on the assumption that the environment has binary structure, i.e. that each grid cell is either occupied or free. This assumption, however, is not always justified. For example, if the environment contains a wall that is not parallel to the x- or y-axis of the grid there must be grid cells which are only partly covered. In occupancy grids the probability that such cells are occupied will inevitably converge to one (see Figure 1). Coverage maps overcome this limitation by storing a posterior about its coverage for each cell. Coverage values range from 0 to 1. Whereas a coverage of 1 means that the cell is fully occupied, an empty cell has a coverage of 0. Since the robot usually does not know the true coverage of a grid cell, it maintains a probabilistic belief $p(c_l)$ about the coverage of the cell c_l . In principle, there are different ways of representing $p(c_l)$. They range from parametric distributions such as (mixtures of) Gaussians or non-parametric variants such as histograms. Throughout this paper we assume that each $p(c_l)$ is given by a histogram over possible coverage values. More precisely, we store a histogram for each grid cell, in which each bin contains the probability that the corresponding grid cell has the particular coverage. A typical example is depicted in Figure 2. It shows the posterior for the cell containing the obstacle in the situation illustrated in Figure 1 obtained after 30 measurements.




Figure 2: Coverage posterior for the cell containing the obstacle in the situation depicted in Figure 1.

3 Updating Coverage Maps

To update a coverage map based on sensory input, we apply a Bayesian update scheme similar to that of occupancy grids. Throughout this paper we assume that our sensor provides distance information. Accordingly, we have to convert the distance information to coverage values. What we need to know is the coverage map c that has the highest likelihood under all distance measurements d_0, \dots, d_T . If we use Bayes rule and assume that consecutive measurements are independent given that we know the map c , we obtain:

$$p(c | d_0, \dots, d_T) \stackrel{\text{Bayes}}{=} \frac{p(d_0, \dots, d_T | c) \cdot p(c)}{p(d_0, \dots, d_T)} \quad (1)$$

$$= \beta \cdot p(c) \cdot p(d_0, \dots, d_T | c) \quad (2)$$

$$= \beta \cdot p(c) \cdot \prod_{t=0}^T p(d_t | c) \quad (3)$$

Next we need to know how to determine the likelihood $p(d_t | c)$ of measuring d_t given the map c . Again we apply Bayes rule and obtain:

$$p(c | d_0, \dots, d_T) \stackrel{\text{Bayes}}{=} \beta \cdot p(c) \cdot \prod_{t=0}^T \frac{p(c | d_t) \cdot p(d_t)}{p(c)} \quad (4)$$

$$= \beta \cdot \frac{\xi^{T+1}}{p(c)^T} \cdot \prod_{t=0}^T p(c | d_t) \quad (5)$$

$$= \beta' \cdot \prod_{t=0}^T p(c | d_t) \quad (6)$$

Equation (5) is obtained from Equation (4) by assuming that $p(c)$ is constant and that $\xi = p(d_t)$ is constant for every t . The variables β and β' represent normalization constants ensuring that the left-hand side sums up to one over all c . Under the strong assumption that the individual cells of a coverage map are independent¹ we finally obtain:

$$p(c | d_0, \dots, d_T) = \beta' \cdot \prod_{t=0}^T \prod_{l=0}^L p(c_l | d_t) \quad (7)$$

$$= \beta' \cdot \prod_{l=0}^L \prod_{t=0}^T p(c_l | d_t) \quad (8)$$

¹This independence is frequently assumed in the context of occupancy maps. We would like to refer to a recent work by Thrun [18] on how to better deal with the dependency between cells.

Thus, to update a map given a measurement d_t we simply have to multiply the current belief about the coverage of each cell c_l by the belief about the coverage of this cell resulting from d_t . Additionally the maximum likelihood coverage map is obtained by choosing the mode of the coverage histogram for each cell c_l .

It remains to describe how we actually compute $p(c_l | d_t)$, i.e. how we determine the distribution about the potential coverage values of a cell c_l with distance d^l to the sensor given a measurement d_t . In our current system, we use a mixture of a Gaussian $\mathcal{N}(\mu, \sigma)$ and a uniform distribution γ to compute the probability $p(c_l = x | d_t)$ that the coverage of c_l is x :

$$p(c_l = x | d_t) = \gamma(d^l, d_t) + \mathcal{N}(\mu(d^l - d_t), \sigma(d^l, d_t), x). \quad (9)$$

The value of the uniform distribution is computed using the function $\gamma(d^l, d_t)$ which increases monotonously in d^l and d . It reflects a typical behavior of proximity sensors like sonars, because the accuracy of a measurement decreases with the distance to the obstacle.

The mean $\mu(d^l - d)$ of the Gaussian is computed in the following way:

$$\mu(d^l - d) = \begin{cases} 0, & (d^l - d) < -\frac{r}{2} \\ \frac{1}{2} + \frac{d^l - d}{r}, & |d^l - d| < \frac{r}{2} \\ 1, & (d^l - d) > \frac{r}{2} \end{cases} \quad (10)$$

where r is the grid resolution of the map. Note that we distinguish three situations, depending on whether the measurement ends in c_l or not. Suppose that the measurement does not end in c_l and the distance d^l is shorter than d . In this case we have $d^l - d < -\frac{r}{2}$. In such a situation, the mean of the Gaussian is zero, since it is more likely that a cell covered by a range measurement that does not end in it is completely empty. The second line of Equation (10) represents the situation in which d ends within c_l . In this case the mean is inverse proportional to the amount the cell is covered by d . Finally, cells lying up to $20cm$ behind a cell, in which the measurement ends, are most likely completely occupied so that the mean is 1.

The value of the standard deviation $\sigma(d^l, d)$ of the Gaussian also is a function that is monotonously increasing in d^l and d except when $|d^l - d| < \frac{r}{2}$. In this range $\sigma(d^l, d)$ has a constant value that exceeds all values outside of this interval.

To obtain the optimal parameters for the functions in our sensor model we apply the maximum likelihood principle. We first apply a highly accurate scan-alignment procedure [10] on laser range information. Next we manually extract geometric objects from the corresponding range data. Given these geometric primitives we use straightforward geometric projections to compute the ground truth information, i.e., the exact coverage of each cell of a given discretization. We evaluate a particular setting for the parameters of our model by determining the likelihood of the ground truth given a coverage map obtained using this setting. To maximize the likelihood we apply local search techniques in the parameter space. Figure 3 plots $p(c_l = x | d)$ for a measured distance of $d = 1m$ obtained for ultrasound data recorded with our B21r robot Albert, which is depicted in picture (a) of Figure 4. As the plot illustrates, for a measured distance of $1m$, cells close




Figure 3: Sensor model $p(c_l = x | d)$ for a measured distance $d = 100cm$.

to the robot are unoccupied with high likelihood. However, cells close to the measured distance are covered with high probability. Figure 4 shows two coverage maps build from real sonar data obtained by a real robot. Image (b) depicts a map of the Sieg Hall at the University of Washington and (c) a map of the office environment of our laboratory.

4 Strategies for Choosing the Next View-Point

One of the key problems during exploration is to choose appropriate vantage points. At the selected location the robot will perform the next measurement to retrieve new information about its environment. In this section we will present four methods to choose an appropriate position. In general there are two different aspects that are relevant for the view-point selection. On the one hand, the uncertainty of the robot in the map should be as small as possible, and on the other hand, the number of measurements to be incorporated as well as the distance traveled should be minimized.

To determine the uncertainty in the state of a particular cell we consider the entropy of the posterior for that cell. Entropy is a general measure for the uncertainty of a belief. The entropy H of a histogram h consisting of n bins h_i ($i = 1, \dots, n$) is defined as:

$$H(h) = - \sum_{i=1}^n p(h_i) \cdot \log p(h_i). \quad (11)$$

H is maximal in case of a uniform distribution. The minimal value zero is obtained if the system is absolutely certain about the state of the corresponding cell. Thus, if we want to minimize the uncertainty in the current map, all we need to do is to reduce the entropy of the histograms in the coverage map. Furthermore, we can specify, at which moment the exploration task has been completed. Suppose the size of the environment is limited. Then the goal of the exploration process for a coverage map c has been achieved if $H(h(c_l)) < \epsilon$ for all cells $c_l \in c$ that can be reached by the robot. Here $h(c_l)$ corresponds to the histogram representing the coverage of c_l . Additionally, the system has to detect a situation in which the robot is unable to reduce the entropy of a cell below ϵ to ensure the termination of the exploration task. In our system this is achieved by monitoring the change of entropy. If this change is below .001 for five consecutive measurements, the cell is regarded as explored enough.




Figure 4: Picture (a) shows our B21r robot Albert equipped with a ring of 24 sonar sensors. Image (b) depicts a coverage map generated for data collected in the Sieg Hall building of the University of Washington. Picture (c) displays the coverage map of our laboratory environment.

4.1 Closest Location (CL)

A simple and frequently used strategy is to drive to the closest location at which the robot can gather information about a cell that has not been explored well enough. This way the question whether a cell has been explored well enough depends on the actual measure used to quantify the uncertainty about individual grid cells. In our approach this measure is the entropy. A cell is regarded as been sensed accurately enough if the entropy of the coverage belief does not exceed ϵ or if it does not change any longer. This strategy CL does not take into account how much information will be obtained at a particular view-point. Rather it seeks to minimize the distance to that location:

$$l_{next} = \operatorname{argmin}_{l \in L(c)} d_c(l, x), \quad (12)$$

where $L(c)$ is the set of cells which have a grid cell with high entropy in its visible range and $d_c(l, x)$ is the distance between the locations l and x given the current map c of the environment.

4.2 Maximum Information Gain (IG) and (IG_WIN)

The second strategy is solely governed by the information gain that can be obtained about the environment at a specific view-point. The information gain is defined as the change of entropy introduced by incorporating the measurement obtained at that location into the map. If we integrate a single measurement d into a cell c_l , the information gain is defined as:

$$I(h(c_l) | d) = H(h(c_l)) - H(h'_d(c_l)) \quad (13)$$

where $h'_d(c_l)$ is the histogram of cell c_l after integrating measurement d according to our sensor model. The information gain of a measurement is then computed as the sum of the information gains for all cells covered by that measurement. Since we do not know which measurement we will receive if the robot measures at a certain position l , we have to integrate over all possible measurements to compute the expected information gain for that view-point:

$$E[I(l)] = \sum_d p(d | c) \cdot \sum_{c_i \in C(l, d)} I(h(c_i) | d) \quad (14)$$

Here $C(l, d)$ is the set of cells covered by measurement d . To efficiently compute the likelihood of an observation $p(d | c)$ we apply a ray-tracing technique similar to Moravec and Elfes [14] using the current maximum likelihood coverage map.

Since the complexity of Equation (14) depends exponentially on the number of dimensions of the measurement, we consider all measurements independently. For example, for our robot equipped with 24 ultrasound sensors we compute the average information gain over all 24 sensors.

The next view-point is then defined as:

$$l_{next} = \operatorname{argmax}_{l \in L(c)} E[I(l)] \quad (15)$$

One of the disadvantages of this strategy is that it does not take into account the distance to be traveled by the robot. To deal with this problem we also consider the strategy IG_WIN which restricts the search for potential vantage points to a local window until this has been explored. Once this has been done, there is no need for the robot to return to this area again.

4.3 Combination of IG and CL (IG-CL)

The final strategy discussed in this paper tries to combine the properties of the strategies CL and IG. The goal is to find an optimal tradeoff between the evaluation functions (12) and (15):

$$l_{next} = \operatorname{argmax}_{l \in L(c)} \left[\alpha \cdot \frac{E[I(l)]}{\max_{l' \in L(c)} E[I(l')]} - \frac{d_c(l, x)}{\max_{l' \in L(c)} d_c(l', x)} \right] \quad (16)$$

By adapting the weight α the user can easily influence the behavior of a robot and optimize its performance for a special task. A value close to zero results in a behavior similar to the strategy CL. For high values of α the strategy converges to the strategy IG. Please note that functions of this type have successfully been applied in the past for coordinating multiple robots during exploration [2, 15].

5 Experiments

We implemented the models and exploration strategies described above and performed a series of exploration runs in different environments. The goal of the experiments presented in this section is to illustrate that a robot can build accurate maps using our models. Additionally we describe experiments in which we analyze the properties of the four view-point selection techniques described above.

5.1 Advantage over Scan Counting

The first experiment is designed to illustrate that an approach which considers the uncertainty in the belief about the coverage of a cell to select view-points yields more accurate maps than techniques relying on scan counting approaches. Scan




Figure 5: These images depict maps and trajectories of the robot obtained in the corridor of our office environment. The left map (a) is a typical occupancy grid map obtained when every cell needs to be covered just once. Image (b) shows a coverage map obtained with the strategy CL. The rightmost image (c) shows an occupancy map obtained using scan counting with a threshold of $n = 50$.

technique	path length	$\{c_i \mid H(h(c_i)) > 0.65\}$
coverage maps	$89.1m$	0%
counting	$26.6m$	21%
ext. counting	$90.6m$	1.5%

Table 1: Data from the experiments shown in Figure 5. To compute the entropy for the scan counting results we generated coverage maps from the recorded sonar data.

counting techniques count the number of times a cell is intercepted by a measurement. Several exploration techniques [2, 7, 20] assume that a place is explored if it has been scanned once. This is problematic especially when the underlying sensors are noisy. Figure 5 (a) shows a typical occupancy grid map of our laboratory environment obtained from real sonar data when using the scan counting technique. Since the exploration process is stopped as soon as all reachable locations were covered by a measurement, many cells of the resulting occupancy map have a high uncertainty (see Table 1). This is due to the high amount of sensor noise, which sonars often produce. On the other hand, if we use our approach and consider the uncertainty in the coverage of individual cells to select view-points, the resulting maps are more accurate (see Figure 5(b)).

Obviously, a straightforward extension of the scan counting would be to assume that each cell has to be covered n times and not only once. An occupancy grid obtained by this extension can be seen in Figure 5 (c). In this experiment we additionally discounted longer beams in order to account for the fact that range sensors provide fewer information for distant places. Whereas this map looks similar as the corresponding coverage map, the robot’s uncertainty about the state of the environment is higher. This is illustrated by the values given in Table 1. The right column of this table contains the percentage of cells in c for which the entropy exceeds the given threshold. Thus, even extended scan counting does not guarantee that in the end every cell is explored well enough. Typically, some cells will be measured too often, others not often enough.

Figure 6 shows the result of a series of 50 simulation experiments. In these experiments we forced the robot to reach a scan count of n where n varied between 1 and 130.² We counted the number of cells m that would be explored well enough given the entropy criterion for coverage maps and plotted the length of the overall path against m . The resulting graph is shown in Figure 6. The cross on the right side indi-

²In practice, a good candidate value of n would be the maximum number of measurements necessary to obtain a coverage map that fulfills the entropy threshold criterion




Figure 6: Average path length of the robot depending on the number of sufficiently explored cells (left image) and the simulation environment used during this experiment (right image).

cates the path length obtained when using the strategy CL for coverage maps. As can be seen, if more than 85% of the map should be explored well enough, a decision-theoretic exploration strategy yields shorter trajectories than extended scan counting.

5.2 A Comparison of the View-Point Selection Strategies

Robots performing 2-d exploration tasks with sonars oder laser scanners normally integrate every sensor measurement because the amount of data is reasonably small and easy to integrate. In this section we consider the situation that analyzing a measurement produces high costs. This might be the case if e.g. the distance information needs to be extracted from stereo images. In such a situation the number of measurements needed for the exploration task is a value of interest. As mentioned above, one of the major advantages of our coverage maps is that they allow the integration of the uncertainty into the selection process of the next view-point. The experiments in this section are designed to compare the performance of the different strategies. To carry out the experiments, we varied the size of the local window when using IG_WIN and the weight α in the evaluation function of IG_CL (see Equation (16)). In Figure 7 (a) and (b) the numbers behind IG_CL show the value of the weight α and the numbers behind IG_WIN indicate the radius of a circle which defines the local window. The results have been obtained using 20 runs per strategy in the environment shown in Figure 7 (d). Please note that further experiments carried out in alternative environments showed similar results and are omitted for the sake of brevity. The maximum allowed entropy during all experiments described in this section was set to .6.

Figure 7 (a) shows the average number of measurements necessary to complete the exploration task for each strategy. As can be seen from the figure, the strategy IG needs the minimum amount of measurements. The strategy IG_CL with $\alpha = 1.0$ needs approximately the same number of measure-




Figure 7: Figure (a) shows the average number of measurements needed by different strategies, whereas in (b) depicts the average path length of the exploration for each strategy. The value behind IG_WIN shows the size of the local window and behind IG_CL the value of the parameter α . The error-bars show the .05 significance interval. The images (c) and (d) show the paths driven during the exploration experiment. Figure (c) shows a sample path of method IG and the (d) a path generated by CL.

ments as IG. The strategy CL requires the maximum number of measurements compared to all other strategies considered here. The reason is that it only seeks to minimize the path length without considering the information gained at particular locations. The error-bars correspond to the 5% confidence interval.

In our experiments we found that a nearest neighbor view-point selection strategy like CL outperforms an approach considering information gain if the robot is allowed to integrate measurements while it is moving (assuming that the acquisition and integration of measurements can be done fast). This can be seen in Figure 7 (b), which plots the average path length driven by the robot during the exploration task for all different strategies. With respect to the path length the strategy CL shows the best behavior as the resulting trajectories are shorter than those of all other techniques. Thus, the CL strategy efficiently covers the terrain. In contrast to that, the IG strategy ignores the distance to be driven and therefore produces an extremely long path which results in the worst behavior of all strategies (see also Figure 7 (c) & (d)). The IG_CL strategy with $\alpha = .3$ appears to yield a good trade-off between number of measurements and overall path length. According to the experiments it slightly outperforms the IG_WIN strategy.

6 Conclusions

In this paper we have introduced a new representation scheme for maps build with mobile robots from sensor data. In contrast to standard occupancy maps our coverage maps store a posterior for each cell about its coverage. This offers the opportunity to more accurately compute the uncertainty of the robot about the corresponding area in the environment. Additionally, we have presented a sensor model designed to update these maps upon sensory input. Finally, we have compared decision-theoretic approaches to guide a vehicle during exploration.

The technique has been implemented and evaluated in extensive simulation runs and real world applications. The experiments illustrate that by using coverage maps it is possible to build accurate maps even if noisy sensors are used. Experiments analyzing different exploration strategies indicate that a technique combining the maximum uncertainty reduction and the distance to be traveled yields the best trade-off between the number of necessary measurements and the length of the resulting paths.

References

- [1] F. Bourgout, A. A. Makarenko, S. B. Williams, B. Grocholsky, and F. Durrant-Whyte. Information based adaptive robotic exploration. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2002.
- [2] W. Burgard, M. Moors, and F. Schneider. Collaborative exploration of unknown environments with teams of mobile robots. In *dagstuhl*. Springer Verlag, 2002.
- [3] H. Choset. Coverage for robotics - a survey of recent results. *Annals of Mathematics and Artificial Intelligence*, 31, 2001.
- [4] H. Choset. Topological simultaneous localization and mapping (slam): Toward exact localization without explicit localization. *IEEE Transactions on Robotics and Automation*, 17(2), April 2001.
- [5] G. Dissanayake, H. Durrant-Whyte, and T. Bailey. A computationally efficient solution to the simultaneous localisation and map building (SLAM) problem. Working notes of ICRA'2000 Workshop W4: Mobile Robot Navigation and Mapping, April 2000.
- [6] G. Dudek, M. Jenkin, E. Miliotis, and D. Wilkes. Robotic exploration as graph construction. *IEEE Transactions on Robotics and Automation*, 7(6):859–865, 1991.
- [7] T. Edlinger and E. von Puttkamer. Exploration of an indoor-environment by an autonomous mobile robot. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 1994.
- [8] H.H. González-Baños and J.C. Latombe. Navigation strategies for exploring indoor environments. *International Journal of Robotics Research*, 2001.
- [9] J.-S. Gutmann and K. Konolige. Incremental mapping of large cyclic environments. In *Proc. of the International Symposium on Computational Intelligence in Robotics and Automation (CIRA)*, 2000.
- [10] D. Hähnel, D. Schulz, and W. Burgard. Map building with mobile robots in populated environments. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2002.
- [11] S. Koenig, C. Tovey, and W. Halliburton. Greedy mapping of terrain. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2001.
- [12] N. Massios, L. Dorst, and F. Voorbraak. A strategy for robot surveillances using the geometrical structure of the environment,. In *IJCAI Workshop "Reasoning with Uncertainty in Robotics"*, 2001.
- [13] S.J. Moorehead, R. Simmons, and W.L. Whittaker. Autonomous exploration using multiple sources of information. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2001.
- [14] Hans P. Moravec and A.E. Elfes. High resolution maps from wide angle sonar. In *Proc. IEEE Int. Conf. Robotics and Automation*, pages 116–121, 1985.
- [15] R. Simmons, D. Apfelbaum, W. Burgard, D. Fox, M. Moors, S. Thrun, and H. Younes. Coordination for multi-robot exploration and mapping. In *Proc. of the National Conference on Artificial Intelligence (AAAI)*, 2000.
- [16] C.J. Taylor and D.J. Kriegman. Exploration strategies for mobile robots. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, pages 248–253, 1993.
- [17] S. Thrun. An online mapping algorithm for teams of mobile robots. *International Journal of Robotics Research*, 2001.
- [18] S. Thrun. Learning occupancy grids with forward sensor models. *Autonomous Robots*, 2002. To appear.
- [19] P. Whaite and F. P. Ferrie. Autonomous exploration: Driven by uncertainty. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(3):193–205, 1997.
- [20] B. Yamauchi, A. Schultz, and W. Adams. Integrating exploration and localization for mobile robots. *Adaptive Systems*, 7(2), 1999.