

Mapping and Exploration with Mobile Robots using Coverage Maps

Cyrill Stachniss Wolfram Burgard

University of Freiburg
Department of Computer Science
D-79110 Freiburg, Germany
{stachnis|burgard}@informatik.uni-freiburg.de

Abstract— Exploration and mapping belongs to the fundamental tasks of mobile robots. In the past, many approaches have used occupancy grid maps to represent the environment during the map building process. Occupancy grids, however, are based on the assumption that each cell is either occupied or free. In this paper we introduce coverage maps as an alternative way of representing the environment of a robot. Coverage maps store for each cell of a given grid a posterior about the amount the corresponding cell is covered by an obstacle. We also present a model that allows us to update coverage maps upon input obtained from proximity sensors. We furthermore describe how to use coverage maps for a decision theoretic approach to exploration. Finally we present experimental results illustrating that coverage maps can be used to efficiently learn highly accurate models even if noisy sensors such as ultrasounds are used.

I. INTRODUCTION

Generating maps is one of the fundamental tasks of mobile robots [5, 8, 9, 16]. Many successful robotic systems use maps of the environment to perform their tasks. The questions of how to represent environments and how to acquire models using this representation therefore is an active research area. Throughout this paper we focus on the problem of how to represent the environment of a mobile robot, of how to update the map whenever new sensory input arrives, and of how to guide the vehicle to efficiently build accurate maps. In particular 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 [13], in which each cell is considered as either occupied or free, coverage 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 the left images of 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 top left picture of this figure shows the resulting occupancy probabilities (black represents high likelihood that the cell is occupied). Since the object does only cover 20% of this cell, a coverage value of .2 (as shown in the lower left image of Figure 1) would be a better approximation of the true situation. In addition to




Fig. 1. Typical occupancy map obtained in situations in which cells are only partly occupied (top left) and a coverage map containing the corresponding coverage values (lower left). The coverage posterior for the cell containing the obstacle depicted in the right image.

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.

Exploration is the task of guiding a vehicle during mapping such that it covers the environment with its sensors. In addition to the mapping task, efficient exploration strategies are also relevant for surface inspection, mine sweeping, or surveillance [3, 11]. In the past, several strategies for exploration have been developed. A popular technique for exploration is to extract frontiers between known and unknown areas [2, 6, 18] and to visit the nearest unexplored place. Recently Koenig et al. [10] have shown that such a strategy, which guides the vehicle to the closest unexplored point, keeps the traveled distance reasonably small. Most approaches applying such a technique solely 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. [7] 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. [12] as well as Bourgault et al. [1] use occupancy grids [13] and compute the entropy of each cell in the grid to determine the utility of scanning from a certain location. Whaite and Ferrie [17] 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. Additionally,

several researchers focus on the problem of simultaneous localization and mapping during exploration [1, 4, 5], an aspect that we do not address in this paper.

This paper is organized as follows. In the next section we introduce coverage maps. In Section III we present a sensor model that allows us to update a given coverage map upon sensory input. In Section IV we describe a decision-theoretic approach to exploration based on coverage maps. Finally, we present experiments illustrating the various properties of our approach. We present accurate maps learned by a real robot and discuss the advantages of our technique over existing approaches.

II. 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 for each cell a posterior about its coverage. 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, where each bin contains the probability that the corresponding grid cell has the particular coverage.

In the beginning of the exploration task all cells are typically initialized using a equal distribution in order to represent the maximum uncertainty about the actual state of the cell. The right image of Figure 1 shows a typical coverage posterior we frequently obtain for partly covered cells. Such a scenario is depicted in the lower left image of this figure. In contrast to occupancy grids (top left image) the resulting map represents the situation more precisely. In the next Section we describe how we can update coverage maps based on sensory input.

III. UPDATING COVERAGE MAPS UPON SENSORY INPUT

To update a coverage map based whenever sensor data arrives, we apply a Bayesian update scheme similar to that of occupancy grids. Throughout this paper we assume that our sensor provides distance information. Thus, we need a

formalism 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:

$$\begin{aligned} p(c | d_0, \dots, d_T) &\stackrel{\text{Bayes}}{=} \frac{p(d_0, \dots, d_T | c) \cdot p(c)}{p(d_0, \dots, d_T)} \quad (1) \\ &= \beta \cdot p(c) \cdot p(d_0, \dots, d_T | c) \quad (2) \\ &= \beta \cdot p(c) \cdot \prod_{t=0}^T p(d_t | c). \quad (3) \end{aligned}$$

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:

$$\begin{aligned} p(c | d_0, \dots, d_T) &\stackrel{\text{Bayes}}{=} \beta \cdot p(c) \cdot \prod_{t=0}^T \frac{p(c | d_t) \cdot p(d_t)}{p(c)} \quad (4) \\ &= \beta \cdot \frac{\xi^{T+1}}{p(c)^T} \cdot \prod_{t=0}^T p(c | d_t) \quad (5) \\ &= \beta' \cdot \prod_{t=0}^T p(c | d_t). \quad (6) \end{aligned}$$

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 . We assume that the individual cells of a coverage map are independent. This assumption is frequently used in the context of occupancy maps. We would like to refer to a recent work by Thrun [15] on how to better deal with the dependency between cells. We finally obtain:

$$\begin{aligned} p(c | d_0, \dots, d_T) &= \beta' \cdot \prod_{t=0}^T \prod_{l=0}^L p(c_l | d_t) \quad (7) \\ &= \beta' \cdot \prod_{l=0}^L \prod_{t=0}^T p(c_l | d_t). \quad (8) \end{aligned}$$

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 :

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




Fig. 2. This picture shows our sensor model $P(c_l = x | d)$ for Albert's ultrasound sensors (here for a measured distance $d = 100\text{cm}$).

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:

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

with a grid resolution of r . 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 is also 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 various functions in our sensor model (see Equation (9)) we apply the maximum likelihood principle. We used data sets recorded with our B21r robot Albert, depicted in Figure 4, in our department building. We then compared the resulting maps with a ground-truth map obtained by applying a highly accurate scan-alignment procedure [9] on the laser range information and by manually extracting geometric objects from this data. Given these geometric primitives we can easily compute the exact coverage of each cell of a given discretization by straightforward geometric operations. We evaluate a particular set of parameters




Fig. 3. Coverage map learned from ultrasound data (left image) and ground truth map (right image).

by computing the likelihood of the ground truth map given the corresponding coverage map and apply local search techniques to determine a parameter setting that maximizes the likelihood of the ground truth map.

The image of Figure 2 depicts a fraction of the resulting sensor model $P(c_l = x | d)$ for a distance of $d = 100\text{cm}$ for the ultrasound sensors of our robot. As the plot illustrates, for a measured distance of 1m , cells close to the robot are with high likelihood unoccupied. However, cells close the measured distance are covered with a high likelihood. The maximum likelihood coverage map obtained with this model is shown in the left image of Figure 3. The size of the environment is $17 \times 2.6\text{m}^2$ and the resolution of the here presented map is 5cm . The image also shows the trajectory of the robot during the exploration task. The right image of this figure shows the ground truth map. As can be seen from the figure, the similarity between the learned map and the ground truth is quite high.

IV. EXPLORATION WITH COVERAGE MAPS

One of the key problems during exploration is to choose appropriate vantage points. In general there are two different aspects that are relevant. 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.

Coverage maps are well-suited to support a decision-theoretic approach to exploration. 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



Fig. 4. The right image depicts the B21r robot Albert used to carry out the experiments. Albert is equipped with a SICK PLS laser range sensor and 24 ultrasound sensors. Both images in the middle show photographs taken within the corridor of our office environment. The right image depicts a coverage map learned by Albert using its ultrasound sensors in this environment.

histograms in the coverage map. Furthermore, we can specify, when the exploration task has been completed. Suppose the environment is of limited size. 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. 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.

To take into account the accuracy of the information provided by the sensor we compute the expected information gain which is the expected change of entropy given that the robot obtains a measurement at a certain location in the map. For a given cell c_l and measurement d , the information gain is defined as:

$$I(h(c_l) | d) = H(h(c_l)) - H(h'_d(c_l)). \quad (12)$$

Here $h'_d(c_l)$ 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 scans the environment at a certain position l , we have to integrate over all possible measurements to compute the expected information gain for that view-point:

$$E[I(l)] = \sum_d p(d | c) \cdot \sum_{c_i \in C(l,d)} I(h(c_i) | d). \quad (13)$$

Here $C(l,d)$ is the set of cells covered by a measured distance of d from location l . To efficiently compute the likelihood of an observation $p(d | c)$ we apply a ray-tracing technique similar to Moravec and Elfes [13] using the current maximum likelihood coverage map.

In extensive experiments we figured out that an approach that purely relies on the information gained at particular vantage points usually minimizes the number of measurements needed to learn a map, it has the major disadvantage that it does not take into account the overall

path length of the resulting trajectory. On the other hand, the popular strategy, which guides the robot to the closest unexplored point, minimizes the length of the trajectory traveled by the robot but increases the number of necessary measurements. The strategy used by our robot therefore computes a tradeoff between the utility of vantage points L_λ and the costs $d_c(l, x)$ of reaching them:

$$l_{next} = \operatorname{argmax}_{l \in L} \left[\alpha \cdot \frac{E[I(l)]}{\max_{l' \in L} E[I(l')]} - (1 - \alpha) \cdot \frac{d_c(l, x)}{\max_{l' \in L} d_c(l', x)} \right]. \quad (14)$$

This way it combines the advantages of both. It reduces the distance to be traveled by the robot and the number of measurements necessary to achieve the desired level of certainty. Please note that by adapting the weight α the user can easily influence the behavior of a robot and optimize the robot's performance for a special task. A value close to zero results in a greedy behavior. A value close to 1, in contrast leads to a strategy that only considers the information gain. A more detailed discussion about different the exploration strategies including experimental comparisons can be found in [14].

V. EXPERIMENTAL RESULTS

Our techniques described above have been implemented and evaluated using data gathered with a real robot and in simulation runs. In our experiments the use of coverage maps has shown an advantage over standard occupancy grids for the decision about which locations need further consideration. When the robot has to actively control its motions in order to acquire all relevant information necessary to generate an accurate map, the uncertainty representation is of utmost importance. The experiments described in this section are designed to illustrate that coverage maps in combination with our sensor model can be used to learn high-quality maps. They furthermore illustrate that they facilitate a decision-theoretic control of the robot during exploration for generating highly accurate maps with noisy sensors.




Fig. 5. The top image depicts a coverage map build from recorded sonar data at the University of Washington. The lower image shows 2% of all sonar scans used to build the map above and illustrates the high noise in the measurement data.

A. Mapping with Noisy Sensors

The first experiment is designed to illustrate that we obtain highly accurate coverage maps using our sensor model. In this real world experiment we used sonar data gathered with our mobile robot Albert in our office environment. Albert traveled along the corridor and entered three rooms of our lab. The two middle images of Figure 4 show pictures of this environment. As can be seen, there are lots of glass panes which are hard to map with ultrasounds because of their smooth surface. The resulting coverage map is shown in the right image of Figure 4. We would like to emphasize that even smaller details such as the narrow pillars at the walls are visible in the resulting map.

Another example for a coverage map build from real sonar data is depicted in the top image of Figure 5. The sonar data (see lower image of the same figure) has been recorded while the robot was controlled manually using a joystick. Since the robot was not performing an exploration task it did not enter any of the doorways in this environment.

B. Advantage over Scan Counting

The next 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 counting techniques store for each cell the number of times it has been intercepted by a measurement. Several exploration techniques [2, 6, 18] assume that a place is explored if it has been scanned once. This is problematic especially when the underlying sensors are noisy. Figure 6 shows a typical occupancy grid map of our laboratory environment obtained from real sonar data and using this approach. Although this map reveals the structure of the environment it lacks several details that are contained in the corresponding coverage map (see Figure 3). Since the exploration process stops as soon as all reachable locations were intercepted by a




Fig. 6. This image depicts an occupancy grid map obtained with scan counting ($n = 1$).

technique	path length	$ \{c_i \mid H(h(c_i)) > \epsilon\} $
coverage maps	89.1m	0%
counting ($n = 1$)	26.6m	21%
counting ($n = 50$)	90.6m	1.5%

TABLE I

THIS TABLE SHOWS THE PATH LENGTH AND NUMBER OF CELLS WITH HIGH ENTROPY FOR DIFFERENT EXPLORATION STRATEGIES.

measurement, many cells of the resulting occupancy map have a high uncertainty. Especially if noisy sensors are used the robot has to scan cells multiple times. This leads to an extension of scan counting in which one assumes that each cell has to be covered n times and not only once. A candidate value for n could be the maximum number of measurements necessary for obtaining a coverage map that fulfills the entropy threshold criterion.

To analyze the quality of occupancy maps obtained for different values of n we performed several experiments. In these experiments we additionally discounted longer beams in order to account for the fact that range sensors provide fewer information for distant places. The results for $n = 1$ and $n = 50$ (in practical experiments we found that $n = 50$ yields coverage maps that typically fulfill the entropy criterion for ultrasound sensors) are summarized in Table I. The right column of this table contains the percentage of cells in c for which the entropy exceeds the given threshold. As can be seen from the figure, more than 20% of the cells do not fulfill the entropy criterion if $n = 1$. In the case of $n = 50$ still 1.5% of the cells are above this threshold. In contrast to this, our approach considers the uncertainty in the coverage of the individual cells so that the resulting maps are more accurate. As this experiment demonstrates, 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.

To analyze the relationship between the overall distance traveled and the percentage of sufficiently explored cells, we performed 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. We counted the number of cells m that were 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 7. The cross on the right side indicates the path length obtained when using our exploration strategy for coverage maps with a low




Fig. 7. This image shows the resulting path length for scan counting obtained using a simulator. The cross shows the average path length when using coverage maps.

value of α (see Equation (14)). If one requires that 85% or more of the cells c_i should satisfy $H(h(c_i)) < \epsilon = 0.65$, a decision-theoretic exploration strategy yields shorter trajectories than extended scan counting.

VI. CONCLUSIONS

In this paper we introduced coverage maps as a new representation scheme for grid-based maps built with mobile robots from sensor data. Coverage maps store in each cell a posterior about the coverage of that cell. This way they offer the opportunity to reason about the uncertainty of the robot about each particular point in the environment. We also developed a sensor model designed to update coverage maps upon sensory input. Finally, we presented a decision-theoretic approach to guide a vehicle during exploration.

The technique has been implemented and evaluated in extensive simulation runs and in real world experiments. The experiments illustrate that by using coverage maps it is possible to build accurate maps. Additionally they demonstrate that coverage maps can be used to control a robot in order to obtain maps not exceeding a given level of uncertainty which is auxiliary especially if the robot possesses noisy sensors such as ultrasounds.

VII. ACKNOWLEDGMENTS

This work has partly been supported by the German Science Foundation (DFG) under contract number SFB/TR8-03 (project A3) and by the EC under contract number IST-2000-29456.

VIII. 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 *Proc. of the Dagstuhl Seminar on Plan-based Control of Robotic Agents*. 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. In *ICRA'2000 Workshop on Mobile Robot Navigation and Mapping*, 2000.
- [6] 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.
- [7] H.H. González-Baños and J.C. Latombe. Navigation strategies for exploring indoor environments. *International Journal of Robotics Research*, 2001.
- [8] 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.
- [9] 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.
- [10] S. Koenig, C. Tovey, and W. Halliburton. Greedy mapping of terrain. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2001.
- [11] N. Massios, L. Dorst, and F. Voorbraak. A strategy for robot surveillancs using the geometrical structure of the environment,. In *IJCAI Workshop "Reasoning with Uncertainty in Robotics"*, 2001.
- [12] 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.
- [13] H.P. Moravec and A.E. Elfes. High resolution maps from wide angle sonar. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 1985.
- [14] C. Stachniss and W. Burgard. Exploring unknown environments with mobile robots using coverage maps. In *Proc. of the Int. Conf. on Artificial Intelligence (IJCAI)*, 2003.
- [15] S. Thrun. Learning occupancy grids with forward sensor models. *Autonomous Robots*, 2002. To appear.
- [16] S. Thrun. An online mapping algorithm for teams of mobile robots. *International Journal of Robotics Research*, 2001.
- [17] P. Whaite and F. P. Ferrie. Autonomous exploration: Driven by uncertainty. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(3):193–205, 1997.
- [18] B. Yamauchi, A. Schultz, and W. Adams. Integrating exploration and localization for mobile robots. *Adaptive Systems*, 7(2), 1999.