Proc. of the Fourteenth National Conference on Artificial Intelligence (AAAI-96)

Estimating the Absolute Position of a Mobile Robot Using Position Probability Grids

Wolfram Burgard and **Dieter Fox** and **Daniel Hennig** and **Timo Schmidt**

Institut für Informatik III, Universtiät Bonn
Römerstr. 164
D-53117 Bonn, Germany
{wolfram,fox,hennig,schmidt1}@uran.cs.uni-bonn.de

Abstract

In order to re-use existing models of the environment mobile robots must be able to estimate their position and orientation in such models. Most of the existing methods for position estimation are based on special purpose sensors or aim at tracking the robot's position relative to the known starting point. This paper describes the position probability grid approach to estimating the robot's absolute position and orientation in a metric model of the environment. Our method is designed to work with standard sensors and is independent of any knowledge about the starting point. It is a Bayesian approach based on certainty grids. In each cell of such a grid we store the probability that this cell refers to the current position of the robot. These probabilities are obtained by integrating the likelihoods of sensor readings over time. Results described in this paper show that our technique is able to reliably estimate the position of a robot in complex environments. Our approach has proven to be robust with respect to inaccurate environmental models, noisy sensors, and ambiguous situations.

Introduction

In order to make use of environmental models mobile robots always must know their current position and orientation[1] in their environment. Therefore, the ability of estimating their position is one of the basic preconditions for the autonomy of mobile robots. The methods for position estimation can be roughly divided into two classes: *relative* and *absolute* position estimation techniques (Feng, Borenstein, & Everett 1994). Members of the first class track the robot's relative position according to a known starting point. The problem solved by these methods is the correction of accumulated dead reckoning errors coming from the inherent inaccuracy of the wheel encoders and other factors such as slipping. Absolute position estimation techniques attempt to determine the robot's position without a priori information about the starting position. These approaches of the second class can be used to initialize the tracking techniques belonging to the first class.

[1] In the remainder of this paper we use the notion "position" to refer to "position and orientation" if not stated otherwise.

This paper addresses the problem of estimating the absolute position of a mobile robot operating in a known environment. There are two reasons why we consider this problem as relevant:

1. Whenever the robot is switched on, it should be able to re-use its model of the environment. For this purpose, it first has to localize itself in this model.
2. If the position tracking has failed, i.e. the robot has lost its position in the environment, it should be able to perform a repositioning.

To avoid modifications of the environment and expensive special purpose sensors we are interested in map-matching techniques, which match measurements of standard sensors against the given model of the environment. We have the following requirements to such a method:

1. **The method should be able to deal with uncertain information.** This is important because sensors are generally imperfect. This concerns wheel encoders as well as proximity sensors such as ultrasonic sensors or laser range-finders. Moreover, models of the environment are generally inaccurate. Possible reasons for deviations of the map from the real world come from imperfect sensors, measuring errors, simplifications, open or closed doors, or even moving objects such as humans or other mobile robots.
2. **The method should be able to deal with ambiguities.** Typical office environments contain several places which cannot be distinguished with a single measurement. As example consider a long corridor, where changes of the position due to the limited range of the sensors do not necessarily result in changes of the measured values. Thus, the set of possible positions of the robot is a region in that corridor.
3. **The method should allow the integration of sensor readings from different types of sensors over time.** Sensor fusion improves reliability while the integration

over time compensates noise and is necessary to resolve ambiguities.

Position probability grids simultaneously address all these desiderata. They allow a mobile robot to determine its position in typical office environments within a short time. Moreover, our method is able to deal with uncertain sensor information and ambiguous situations.

The approach described in this paper is based on the construction of certainty grid maps described in (Moravec & Elfes 1985). Certainty grid maps have been proven to be a powerful means for the solution of different problems. Originally, they were designed to provide a probabilistic model of the robot's environment. In the past such occupancy probability maps or variants of them have been successfully used for collision avoidance (Borenstein & Koren 1990; 1991) and path planning (Buhmann *et al.* 1995; Moravec 1988). This paper issues a further application area of this technique, namely the estimation of the absolute position of a robot. The principle of our approach is to accumulate in each cell of the position probability grid the posterior probability of this cell referring to the current position of the robot. Because we have to consider a discrete set of possible orientations in addition to the discretization of the two-dimensional environment, position estimation is a three-dimensional problem. This extension, however, does not result in any principle problems, because the certainty grid concept can easily be extended to problems with higher dimensionality (Moravec & Martin 1994).

Related work

Various techniques for the estimation of the position of mobile vehicles by matching sensor readings against a given model of the environment have been developed in the past (Cox & Wilfong 1990; Feng, Borenstein, & Everett 1994). Most of them address the problem of tracking the current position and orientation of the robot given its initial configuration. Recently, more and more probabilistic techniques are applied to position estimation problems. These approaches can be distinguished by the type of maps they rely on.

Techniques based on metric or grid-based representations of the environment generally generate unimodal or Gaussian distributions representing the estimation of the robot's position. (Weiß, Wetzler, & von Puttkamer 1994) store angle histograms constructed out of range-finder scans taken at different locations of the environment. The position and orientation of the robot is calculated by maximizing the correlation between histograms of new measurements and the stored histograms. (Schiele & Crowley 1994) compare different strategies to track the robot's position based on occupancy grid maps. They use two different maps: a local grid computed using the most recent sensor readings, and a global map built during a previous exploration of the environment or by an appropriate CAD-tool. The local map is matched against the global map to produce a position and orientation estimate. This estimate is combined with the previous estimate using a Kalman filter (Maybeck 1990), where the uncertainty is represented by the width of the Gaussian distribution. Compared to the approach of Weiß et al., this technique allows an integration of different measurements over time rather than taking the optimum match of the most recent sensing as a guess for the current position.

Other researchers developed positioning techniques based on topological maps. (Nourbakhsh, Powers, & Birchfield 1995) apply Markov Models to determine the node of the topological map which contains the current position of the robot. Different nodes of the topological map are distinguished by walls, doors or hallway openings. Such items are detected using ultrasonic sensors, and the position of the robot is determined by a "state-set progression technique", where each state represents a node in the topological map. This technique is augmented by certainty factors which are computed out of the likelihoods that the items mentioned above will be detected by the ultrasonic sensors. (Simmons & Koenig 1995) describe a similar approach to position estimation. They additionally utilize metric information coming from the wheel encoders to compute state transition probabilities. This metric information puts additional constraints on the robot's location and results in more reliable position estimates. (Kortenkamp & Weymouth 1994) combine information obtained from sonar sensors and cameras using a Bayesian network to detect gateways between nodes of the topological map. The integration of sonar and vision information results in a much better place recognition which reduces the number of necessary robot movements respectively transitions between different nodes of the topological map.

Due to the separation of the environment into different nodes the methods based on topological maps, in contrast to the methods based on metric maps described above, allow to deal with ambiguous situations. Such ambiguities are represented by different nodes having high position probabilities. However, the techniques based on topological maps provide a limited accuracy because of the low granularity of the discretization. This restricted precision is disadvantageous if the robot has to navigate fast through its environment or even grasp for objects.

The position probability grid method described here allows to estimate the robot's position up to a few centimeters. This is achieved by approximating a position probability function over a discrete metric space defining possible positions in the environment. It therefore can be used to provide an initial estimate for the tracking techniques. But even the methods based on topological maps could be augmented by our approach. If the nodes of the topological map additionally contain metric information, our approach

could be used to position the robot within a node.

Building position probability grids

The certainty grid approach was originally designed by Elfes and Moravec as a probabilistic grid model for the representation of obstacles. The basic idea is to accumulate in each cell of a rectangular grid field the probability that this cell is occupied by an obstacle. Whereas Moravec and Elfes construct a model of the environment given the position of the robot and sensor readings, we go the opposite direction estimating the position given the environmental model and the sensor readings. For this purpose, we construct a *position probability grid* P containing in each field the posterior probability that this field includes the current position of the robot. For a grid field x this certainty value is obtained by repeatedly firing the robot's sensors and accumulating in x the likelihoods of the sensed values supposed the center of x currently is the position of the robot in the environment model m. Each time the robot's sensors are fired, the following two steps are carried out:

1. Update P according to the movement of the robot since the last update. This includes a processing of P to deal with possible dead-reckoning errors.

2. For each grid field x of P and each reading s, compute the likelihood of s supposed x is the current position of the robot in m, and combine this likelihood with the probability stored in x to obtain a new probability for x.

The basic assumptions for our approach are:

- *The robot must have a model m of the world the sensor readings can be matched against.* Such models can either come from CAD-drawings of the environment or can themselves be grid representations of occupancy probabilities.

- *The robot does not leave the environmental model.* This assumption allows us to use the same size for the position probability grid P as for the environmental model m, and to set the probability for positions outside of P to 0.

In the remainder of this section we describe how to integrate different sensor readings into position probabilities. Furthermore we show how to keep track of the robot's movements with explicit consideration of possible dead-reckoning errors.

Integrating multiple sensor readings

In order to give reliable position estimates we have to integrate the information of consecutive sensor readings. Suppose m is the model of the environment, and $p(x \mid s_1 \wedge \ldots \wedge s_{n-1} \wedge m)$ is the (posterior) probability that x refers to the current position of the robot, given m and the sensor readings $s_1, \ldots, s_{n-1}$. Then, to update the probability for x given new sensory input s_n we use the following update formula (Pearl 1988):

$$p(x \mid s_1 \wedge \ldots \wedge s_{n-1} \wedge s_n \wedge m) = \alpha \cdot p(x \mid s_1 \wedge \ldots \wedge s_{n-1} \wedge m) \cdot p(s_n \mid x \wedge m) \quad (1)$$

The term $p(s_n \mid x \wedge m)$ is the likelihood of measuring the sensory input s_n given the world model m and assuming that x refers to the current position of the robot. The constant α simply normalizes the sum of the position probabilities over all x up to 1.

To initialize the grid we use the a priori probability $p(x \mid m)$ of x referring to the actual position of the robot given m. The estimation of $p(x \mid m)$ and $p(s_n \mid x \wedge m)$ depends on the given world model and the type of the sensors used for position estimation. Below we demonstrate how to use occupancy probability maps for position estimation and how sensor readings of ultrasonic sensors are matched against such maps.

Integrating the movements of the robot

In order to update the grid according to the robot's movements and to deal with possible dead reckoning errors we use a general formula coming from the domain of Markov chains. We regard each cell in P as one possible state of the robot, and determine a state transition probability $p(x \mid \tilde{x} \wedge \tau \wedge t)$ for each pair x, $\tilde{x}$ of cells in P, which depends on the trajectory τ taken by the robot and the time t elapsed since the previous update. This transition probability should also model how the trajectory τ fits into the environment. For example, a trajectory leading the robot through free space has a higher probability than a trajectory blocked by an obstacle. Thus, the new probability of a grid field after a movement of the robot is:

$$P[x] \quad := \quad \alpha \cdot \sum_{\tilde{x} \in P} P[\tilde{x}] \cdot p(x \mid \tilde{x} \wedge \tau \wedge t) \quad (2)$$

where α is a normalizing constant. At any time the field of P with the highest probability represents the best estimation for the current position of the robot. The confidence in this estimation depends on the absolute value of the probability and on the difference to the probabilities in the remaining fields of the grid. Thus, ambiguities are represented by different fields having a similar high probability.

Position estimation with occupancy probability maps as world model

In this section we describe the application of this approach by matching ultrasonic sensors against occupancy grid maps.

Matching sonar sensor readings against occupancy grids

To compute the likelihood $p(s \mid x \wedge m)$ that a sensor reading s is received given the position x and an occupancy grid

map m we use a similar approach as described in (Moravec 1988). We consider a discretization $R_1, \ldots, R_k$ of possible distances measured by the sensor. Consequently, $p(R_i \mid x \wedge m)$ is the likelihood that the sonar beam is reflected in R_i.

Suppose $p(r(\tilde{x}) \mid x \wedge m)$ is the likelihood that the cell $\tilde{x}$ reflects a sonar beam, given the position x of the robot and the map m. Furthermore suppose that $\tilde{x}$ belongs to R_i. Assuming the reflection of the sonar beam by $\tilde{x}$ being conditionally independent of the reflection of the other cells in R_i, the likelihood that R_i reflects a sonar beam is

$$p(R_i \mid x \wedge m) = 1 - \prod_{\tilde{x} \in R_i} (1 - p(r(\tilde{x}) \mid x \wedge m)) \qquad (3)$$

Before the beam reaches R_i, it traverses $R_1, \ldots, R_{i-1}$. Supposed that the sonar reading s is included by range R_i, the likelihood $p(s \mid x \wedge m)$ equals the likelihood that R_i reflects the sonar beam given that none of the ranges $R_{<i}$ reflects it. Thus, we have

$$p(s \mid x \wedge m) = p(R_i \mid x \wedge m) \cdot \prod_{j=1}^{i-1} (1 - p(R_j \mid x \wedge m)) \qquad (4)$$

Computing position estimates using occupancy grids

It remains to estimate the initial probability $p(x \mid m)$ that the field x of m contains the current position of the robot. We assume that this probability directly depends on the occupancy probability $m(x)$ of the field x in m: the higher the occupancy probability, the lower is the position probability and vice versa. Therefore, the value $p(x \mid m)$ is computed as follows:

$$p(x \mid m) \quad = \quad \frac{1 - m(x)}{\sum_{\tilde{x} \in m} (1 - m(\tilde{x}))} \qquad (5)$$

Experiments

In this section we show the results from experiments carried out with our robot *RHINO* in real world environments such as typical offices and the AAAI '94 mobile robot competition arena. For the position estimation we match sensor readings coming from ultrasonic sensors against occupancy grid maps.

Implementation aspects

For the sake of efficiency we implemented a simplified model of sonar sensors to compute the likelihood of a reading: instead of considering all cells of the grid covered by the sonar wedge as done in (Moravec 1988) we only consider the cells on the acoustic axis of the sensor. This rough simplification has already been applied successfully in (Borenstein & Koren 1991) to realize a fast collision avoidance technique for mobile robots.

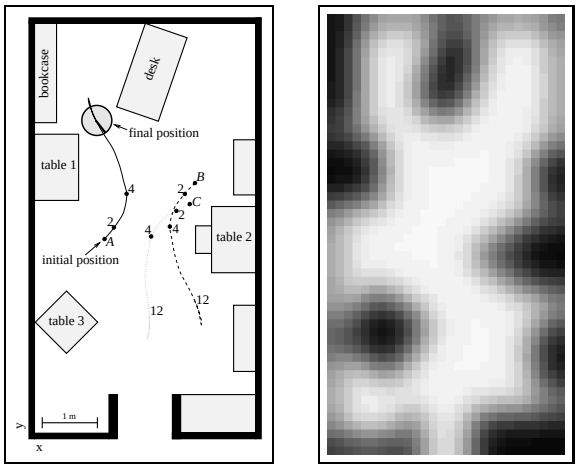

Figure 1: Outline and occupancy grid map of the office

To compute the transition probability $p(x \mid \tilde{x} \wedge \tau \wedge t)$ we assume the dead reckoning errors to be normally distributed. Supposing frequent updates of the position information we simply approximate the probability of the robot's trajectory τ by $p(y|m)$ where y is the position corresponding to x. Thus we have

$$p(x \mid \tilde{x} \wedge \tau \wedge t) \quad = \quad \omega_{x,\tau,t}(\tilde{x}) \cdot p(y \mid m) \qquad (6)$$

where $\omega_{x,\tau,t}$ is a Gaussian distribution.

Position estimation in a typical office

To evaluate the capabilities of our approach we used the task of estimating the position in a typical office of our department. Figure 1 shows an outline of this office, which has a size of $4 \times 7m^2$ and the occupancy grid map used to compute the likelihoods of the sensor readings. For the position estimation we used only 8 of the 24 ultrasonic sensors our robot is equipped with. The size of one grid field is 15×15 cm^2, while we consider 180 possible orientations. For this grid and 8 sensor readings per step, the update of the grid takes about 6 seconds on a Pentium 90 computer.

Figure 1 also shows the initial and final position of the path taken by the robot. At the beginning the robot turned to the left and moved between the bookcase and the desk. At the end of the trajectory the robot turned and started to leave the corner. On this trajectory, which is illustrated by the solid line, 12 sweeps of sonar readings were taken for the position estimation. In addition to the real trajectory *A* two alternative paths *B* and *C* are shown. Figure 2 shows plots of the maximum probabilities for the first, second, fourth, and twelfth reading sets for each position of the map. For the sake of simplicity only the maximal probability over all orientations at each position is shown. Note that the z-axes of the four figures have different scales. The probabilities of the corresponding points belonging to the trajectories *A*, *B*, and *C* are highlighted by vertical lines.

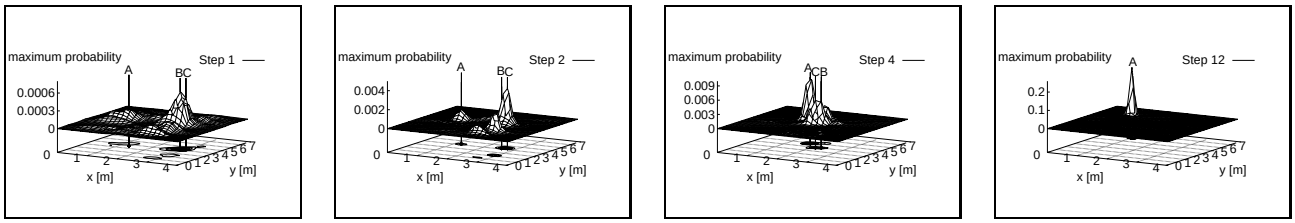

Figure 2: Position probability distribution after 1, 2, 4, and 12 steps

After the first reading we obtain a multi-modal distribution with several small local maxima. At the correct position we observe only a small peak, which is dominated by the starting position of trajectory *B*. After interpreting the second set of readings the probabilities become more concentrated. We observe four small peaks which now have their maximum in position 2 of trajectory *C*. The third and fourth reading sets support the initial position so that the position on trajectory *A* gets the maximum probability. There are two peaks where the smaller one is a super-imposition of two different peaks for the trajectories *B* and *C*. After evaluating 12 sonar sweeps all ambiguities are resolved, and the result is a significant and unique peak with probability 0.26 for the final point of trajectory *A*. This position in fact refers to the real position of the robot.

Dealing with large environments

In the previous example ambiguities appeared as several peaks in the position probability distribution. In large environments we have to expect that due to the limited range of the proximity sensors ambiguities spread out over complete regions. In order to demonstrate the capability of our approach to deal with such complex situations we applied it to the arena of the AAAI '94 mobile robot competition (Simmons 1995). The size of this arena amounts $20 \times 30m^2$. Figure 3 shows the occupancy grid map of this arena constructed with the map-building tool described in (Thrun 1993). The sonar sensor measurements were recorded during an exploration run in this arena. The trajectory of the robot and the 12 positions at which the sensors were fired are included in Figure 3. Again we only used 8 of the 24 sonar sensors and the same resolution for the position probability grid as in the previous example.

Figures 4 and 5 show logarithmic density plots of the maximum position probabilities for all directions after interpreting 6 and 12 sets of sensor readings. Although the information obtained after the first 6 sensor readings does not suffice to definitely determine the current position of the robot, it is obvious that the robot must be in a long corridor. After 12 steps the position of the robot is uniquely determined. The corresponding grid cell has a probability of 0.96 while the small peak at the bottom of Figure 5 has a maximum of 8e-6.

Conclusions

We presented the position probability grid approach as a robust Bayesian technique to estimate the position of a mobile robot. Our method allows the integration of sensor readings from different types of sensors over time. We showed that this method is able to find the position of a robot even if noisy sensors such as ultrasonic sensors and approximative environmental models like occupancy grid maps are used. Our approach has any-time characteristic, because it is able to give an estimation for the current position of the robot already after interpreting the first sensor reading. By incorporating new input this estimation is continuously improved. Position probability grids allow to represent and to deal with ambiguous situations. These ambiguities are resolved if sufficient sensory information is provided. Our technique has been implemented and tested in several complex real-world experiments.

The only precondition for the applicability of the position probability grid approach is an environmental model which allows to determine the likelihood of a sensor reading at a certain position in the environment. In our implementation we used occupancy probability grids as world model in combination with ultrasonic sensors. Alternatively one could use a CAD-model of the environment and cameras for edge detection or integrate simple features like the color of the floor.

Using the currently implemented system our robot needs about one minute to determine its position in a typical office. Although the computation time depends linearly on the grid size, very large environments such as the $600m^2$ wide AAAI '94 robot competition arena do not impose any principle limitations on the algorithm. We are convinced that different optimizations will make our approach applicable online even in such large environments. The most important source for speed-up lies in the pre-analysis of the environmental model. This includes computing and storing the likelihoods of all possible sensor readings for all positions. Additionally, in orthogonal environments the reduction of possible orientations to the alignment of the walls drastically reduces the complexity of the overall problem. Furthermore, the application of a map resolution hierarchy as proposed in (Moravec 1988) can be used to produce rough position estimates which are refined subsequently.

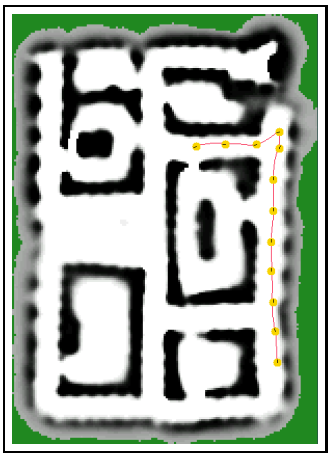

Figure 3: Occupancy grid map of the AAAI '94 mobile robot competition

Figure 4: Density plot after 6 steps

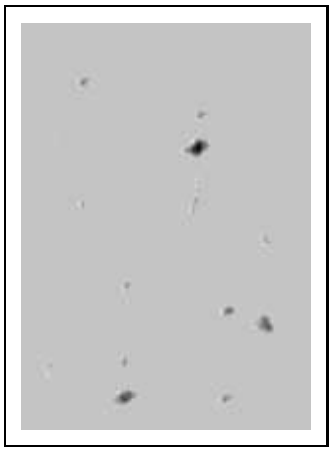

Figure 5: Density plot after 12 steps

Despite these encouraging results there are several warrants for future research. This concerns optimizations as described above as well as active exploration strategies. Such strategies will guide the robot to points in the environment, which provide the maximum information gain with respect to the current knowledge.

References

Borenstein, J., and Koren, Y. 1990. Real-time obstacle avoidance for fast mobile robots in cluttered environments. In *Proc. of the IEEE International Conference on Robotics and Automation*, 572–577.

Borenstein, J., and Koren, Y. 1991. The vector field histogram - fast obstacle avoidance for mobile robots. *IEEE Transactions on Robotics and Automation* 7(3):278–288.

Buhmann, J.; Burgard, W.; Cremers, A. B.; Fox, D.; Hofmann, T.; Schneider, F.; Strikos, J.; and Thrun, S. 1995. The mobile robot Rhino. *AI Magazine* 16(2):31–38.

Cox, I., and Wilfong, G., eds. 1990. *Autonomous Robot Vehicles*. Springer Verlag.

Feng, L.; Borenstein, J.; and Everett, H. 1994. "Where am I?" Sensors and Methods for Autonomous Mobile Robot Positioning. Technical Report UM-MEAM-94-21, University of Michigan.

Kortenkamp, D., and Weymouth, T. 1994. Topological mapping for mobile robots using a combination of sonar and vision sensing. In *Proc. of the Twelfth National Conference on Artificial Intelligence*, 979–984.

Maybeck, P. S. 1990. The Kalman filter: An introduction to concepts. In Cox and Wilfong (1990).

Moravec, H. P., and Elfes, A. 1985. High resolution maps from wide angle sonar. In *Proc. IEEE Int. Conf. Robotics and Automation*, 116–121.

Moravec, H. P., and Martin, M. C. 1994. Robot navigation by 3D spatial evidence grids. Mobile Robot Laboratory, Robotics Institute, Carnegie Mellon University.

Moravec, H. P. 1988. Sensor fusion in certainty grids for mobile robots. *AI Magazine* 61–74.

Nourbakhsh, I.; Powers, R.; and Birchfield, S. 1995. DERVISH an office-navigating robot. *AI Magazine* 16(2):53–60.

Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann Publishers, Inc.

Schiele, B., and Crowley, J. L. 1994. A comparison of position estimation techniques using occupancy grids. In *Proc. of the IEEE International Conference on Robotics and Automation*, 1628–1634.

Simmons, R., and Koenig, S. 1995. Probabilistic robot navigation in partially observable environments. In *Proc. International Joint Conference on Artificial Intelligence*.

Simmons, R. 1995. The 1994 AAAI robot competition and exhibition. *AI Magazine* 16(2):19–29.

Thrun, S. 1993. Exploration and model building in mobile robot domains. In *Proceedings of the ICNN-93*, 175–180. San Francisco, CA: IEEE Neural Network Council.

Weiß, G.; Wetzler, C.; and von Puttkamer, E. 1994. Keeping track of position and orientation of moving indoor systems by correlation of range-finder scans. In *Proceedings of the International Conference on Intelligent Robots and Systems*, 595–601.