Statistical Inference in Mapping and Localization for Mobile Robots

Anita Araneda and Alvaro Soto

Pontificia Universidad Catolica de Chile, Santiago 22, Chile
aaraneda@mat.puc.cl, asoto@ing.puc.cl

Abstract. In this paper we tackle the problem of providing a mobile robot with the ability to build a map of its environment using data gathered during navigation. The data correspond to the locations visited by the robot, obtained through a noisy odometer, and the distances to obstacles from each location, obtained from a noisy laser sensor. The map is represented as an occupancy grid. In this paper, we represent the process using a Graphical Representation based on a statistical structure resembling a Hidden Markov model. We determine the probability distributions involved in this Graphical Representation using a Motion Model, a Perception model, and a set of independent Bernoulli random variables associated with the cells in the occupancy grid forming the map. Our formulation of the problem leads naturally to the estimation of the posterior distribution over the space of possible maps given the data. We exploit a particular factorization of this distribution that allows us to implement an Importance Sampling algorithm. We show the results obtained by this algorithm when applied to a data set obtained by a robot navigating inside an office building type of indoor environment.

1 Introduction

Robust navigation in natural environments is an essential capability of truly autonomous mobile robots. Providing robots with this skill, however, has turn to be a difficult problem. This is particularly true, when robots navigate in unknown environments where globally accurate positioning systems, such as GPS, are not available.

In general, robots need a map of their surrounding and the ability to locate themselves within that map, in order to plan their motion and successfully navigate afterwards. This is why robot mapping and localization is now considered a fundamental component of autonomous mobile robots [18] [10].

Some approaches consider the problem of mapping under the assumption that the locations visited by the robot are known [20]. This situation is not realistic, however, if we consider that sensors of location, such as odometers, carry error in the location measurement. On the other hand, some approaches consider the problem of localization, under the assumption that a map of the environment is available [3] [6] [15]. The actual situation in applications, however,

C. Lemaître, C.A. Reyes, and J.A. Gonzalez (Eds.): IBERAMIA 2004, LNAI 3315, pp. 545–554, 2004.

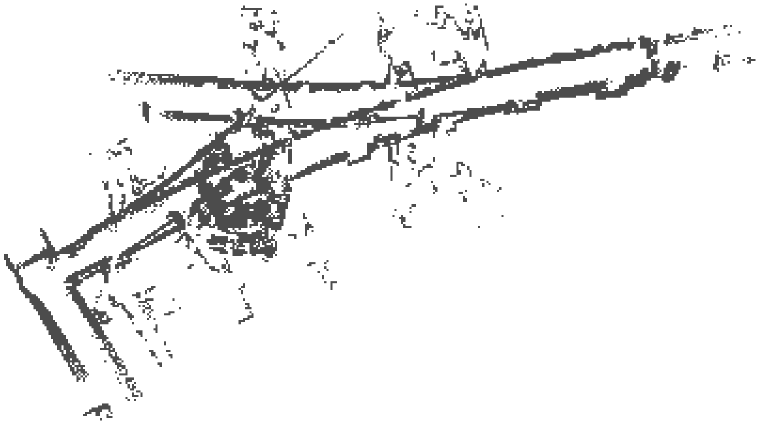

Fig. 1. Map obtained from Raw Data

is that neither the map nor the locations are known. This has led to research on the simultaneous localization and mapping (SLAM) problem using sensor data collected by the robot [13] [21] [19] [10].

As today technology, most of the SLAM approaches consider a robot equipped with an odometer, that collects information about the robot ego-motion, and range sensors, such as sonars or laser range finders, that measure the distances to nearest obstacles. As an example, Figure 1 shows a map drawn from raw odometer and laser readings collected by a mobile robot. The figure shows how odometry error accumulates so that, it seems that the robot has visited two different corridors, instead of just one straight hallway, as it did.

In this paper, we understand SLAM as an estimation problem, where the data correspond to odometer and range sensor readings collected by the robot during its trajectory. The goal is to estimate the posterior distribution of the map and the locations visited by the robot given the data. Odometer readings correspond to rotation and translation measures of the robot movements. Range readings correspond to distances to the closest obstacles, with respect to the robot location. These distances are measured in a set of previously specified directions.

The map is represented by an occupancy grid [5]. In this context, we understand a map of a given environment as a random matrix, each component of which is associated to a spatial location in the environment. The set of associated locations corresponds to a regular grid and each component of the map takes either the value 1 or 0 depending on whether the corresponding location is occupied or not. We express our knowledge of the map through its posterior distribution given the information provided by the robot.

This paper is organized as follows. Section 2 reviews relevant previous work on SLAM. Section 3 discusses the details of our probabilistic approach. Section 4 shows the results of applying our methodology to real data collected by a robot navigating inside an office building type of indoor environment. Finally, Section 5 presents the conclusions of this work.

2 Previous Work

Although there is an extensive research literature touching on mapping or localization for mobile robots[1], the SLAM problem is a relatively newer research area, where most efforts have been made over the last decade. An important family of approaches to SLAM is based on versions of the Kalman filter. The pioneering development in this area is the paper by Smith et al. [16] who basically propose the Hidden Marked Model (HMM) approach widely used today. Smith et al. presents an application of the Kalman filter to the problem of estimating topological maps. They assume a fixed number of landmarks in the environment where these landmarks can be identified by their cartesian coordinates. At a fixed time instant, the set of landmarks coordinates and the location of the robot are assumed to be unobservable or latent variables. As in the Kalman filter, the main assumption is that the posterior distributions of all these variables are Gaussian and that the observations, given the latent variables, can be described as a linear function and a white noise term.

These two assumptions are somewhat restrictive. The Gaussian assumption makes this approach unsuitable for multimodal distributions that arise when the location of the robot is ambiguous. The linearity assumption is not met in general, since the relation between odometry and locations involves trigonometric functions. The Extended Kalman Filter (EKF) [11] partially handles non-linearity using a Taylor approximation.

For the non-Gaussian case, Thrun et al. [22] postulate a general approach that can be used with general distribution functions. Under this approach, however, computing maximum likelihood estimates is computationally too expensive. In [20], Thrun presents an application of the Expectation and Maximization algorithm [4] applied to mapping. The map is treated as the parameter to be estimated while the locations are treated as part of a Hidden Marked Model. Thus, Thrun proposes maximizing the expected log likelihood of the observations and the locations, given the map.

A more recent successful approach to solving the SLAM problem is the Fast-SLAM algorithm [13]. This approach applies to topological maps, and is based on a factorization of the posterior distribution of maps and locations. The models that determine the process are the ones used in [20].

On the other hand, [9] present an approach that is also based on the description in [20], but this one applied to occupancy grids. This approach finds locations iteratively over time. At each point in time, the algorithm estimates the location visited by the robot as the location that maximizes the probability of the current data, given past data and previous location estimates. Next

[1] See [20] for a good overview of the literature in this area.

step finds the map, as the map that maximizes the posterior probability of the estimated locations and the observed data.

Our mapping approach applies to occupancy grid maps of static environments. Our formulation of the problem is based on the approach by Thrun et al. [22]. We build a Graphical Representation of that formulation where the locations are considered unobservable variables determining the observed odometer readings and, together with the map, determining the observed laser readings. The probability model of the whole process is determined by Motion and Perception Models and a prior distribution for the map.

As opposed to previous approaches, our approach propose a fully Bayesian approach where our goal is to estimate the posterior distribution of the map using simulation. Thus, our approach shows a formal description of the entire process and develops a Bayesian solution. The fact that it uses more general Motion Model than the one used by the Kalman Filter approaches, makes it applicable to a wider set of problems. The advantage of this method is that it does not provide a single estimate of the map, as the EM based solution, but it produces multiple maps showing the notion of variability from the expected posterior map. As for localization, we obtain a simulation of the locations visited by the robot from their posterior distribution, as an intermediate step while simulating maps.

3 Our Approach

We describe the SLAM problem in probabilistic terms. We assume that there are true but unobservable locations visited by the robot and that there are true but unobservable distances to obstacles from each of those locations. These locations and distances determine the map of the environment, represented as an occupancy grid. Odometer and laser readings correspond to the observed values of true locations and distances to obstacles, respectively. We assume that the observations are centered at their true counterparts, having random variations around them[2].

Figure 2a) shows a Graphical Representation of the problem where non-observable variables have been circled for clarity. Our Graphical Representation was developed originally in [1] and there is a similar, but less specific Graphical Representation, developed independently in [14]. The Graphical Representation we develop here has been subsequently adopted in [8].

In Figure 2, U_t represents the difference between odometer readings at times $t-1$ and t. Z_t represents the location of the robot at time t. θ_t represents the distances from Z_t to the closest obstacles in front of the robot at time t. It is important to note that $\boldsymbol{\theta}_t$ is fully determined by the location of the robot, Z_t, and the map, $\boldsymbol{M}$. Finally, V_t represents the laser readings at time t.

Writing $\boldsymbol{U}$ and $\boldsymbol{V}$ to denote the matrices of odometer reading differences and laser readings collected from time 1 to T, the SLAM problem can be expressed

[2] From now on we omit the word "true" when referring to true locations and true distances to obstacles.

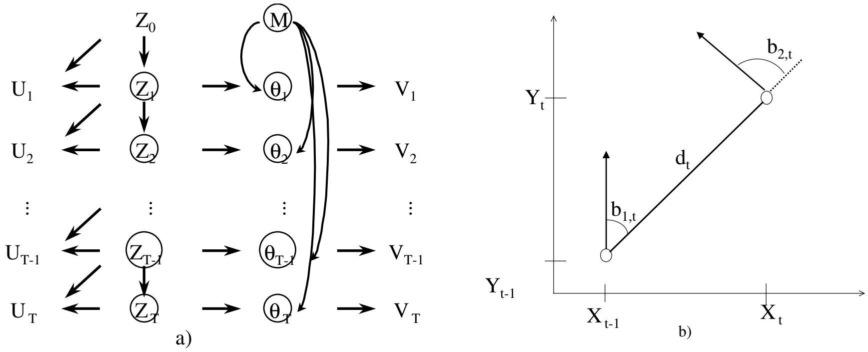

Fig. 2. a) Graphical Representation of the problem. b) Translation and rotation components of movement from Z_{t-1} to Z_t

as the problem of determining the posterior distribution of $\boldsymbol{M}$ given $\boldsymbol{U}$ and $\boldsymbol{V}$, $P(\boldsymbol{M}|\boldsymbol{U},\boldsymbol{V})$.

3.1 Motion and Perception Models

To complete the specification of the process we need to establish the probabilistic models that drive the dependencies, in particular, the so-called Motion and Perception Models. We assume that, to move from Z_{t-1} to Z_t, the robot performs three independent actions: a first rotation , b_{1t}, to face the direction of the translation, a translation , d_t, from Z_{t-1} to Z_t, and a final rotation, b_{2t}, to face the orientation of the robot at time t. These motions are shown in Figure 2b). The elements of U_t correspond to the observed counterparts of $b_{1,t}$, $b_{2,t}$ and d_t, respectively.

Denoting by $\boldsymbol{Z}$ the matrix containing all locations from time 1 to T, and $\boldsymbol{Z}^{t-1}$ the vector of true locations up to time $t-1$, we can write

$$P(\boldsymbol{Z}|\boldsymbol{U}) = \prod_{t=1}^{T} P(Z_t|\boldsymbol{U},\boldsymbol{Z}^{t-1}) = \prod_{t=1}^{T} P(Z_t \mid U_t, Z_{t-1}). \tag{1}$$

The term $P(Z_t \mid U_t, Z_{t-1})$ in the last product in equation (1) is known as the Motion Model. Here we assume a Gaussian Motion Model, which determines that the error of b_{1t}, b_{2t} and d_t with respect to their observed counterparts, correspond to white noise with variances that are proportional to these latter quantities.

Let $\boldsymbol{\theta}$ be the matrix of all distances to obstacles up to time T and $P(\boldsymbol{V} \mid \boldsymbol{\theta})$ be the conditional distribution of laser readings given distances to obstacles. Assuming that laser beams read distances independently from each other and also that these laser readings are independent over time, we write

$$P(\boldsymbol{V}|\boldsymbol{\theta}) = \prod_{t=1}^{T}\prod_{i=1}^{N} P(V_{ti}|\theta_{ti}), \tag{2}$$

where N represents the number of laser beams, and V_{ti} and θ_{ti} represent the laser reading and distance to the closest obstacle, respectively, in the i-th direction at time t. The term $P(V_{ti}|\theta_{ti})$ is known as the Perception Model. We assume that it corresponds to a truncated Gaussian distribution with mean θ_{ti} and with known variance σ^2 determined by the accuracy of the laser sensor. The Gaussian distribution is truncated by 0 at the lower end and the maximum range of the laser at the other.

Finally, we consider that the prior distribution of θ_{ti} corresponds to a geometric distribution, such that $P(\theta_{ti}) = (1-p)^{\theta_{ti}-1}\, p$, where p corresponds to the prior proportion of busy cells in the map. In this work we determine this value empirically by analyzing the data.

3.2 Importance Sampling

In order to find an expression for the posterior distribution of $\boldsymbol{M}$ given $\boldsymbol{U}$ and $\boldsymbol{V}$ we must integrate over pairs $(\boldsymbol{Z},\boldsymbol{\theta})$, as seen in Figure 2a). Alternatively, we note that the map $\boldsymbol{M}$ is fully determined by $\boldsymbol{Z}$ and $\boldsymbol{\theta}$, thus we are interested in the posterior distribution $P(\boldsymbol{Z},\boldsymbol{\theta}|\boldsymbol{U},\boldsymbol{V})$. In the absence of a closed form for this expression, we use Importance Sampling (IS) to sample observations from it.

Importance Sampling [7] [17] is a sampling algorithm used to estimate probability distributions. It has been heavily used in recent years. The main idea in IS is to represent a distribution by a set of observations and a weight associated to each observation in the set. In this way, expected values and other features of the target distribution can be estimated as the weighted average of the observations.

Consider a set of n tuples $(x_j,\omega_j), j = 1,2,\ldots,n$, given by random draws $x_j's$ from a distribution g and corresponding weights $\omega_j's$. Liu and Chen [12] define this set to be *properly weighted* with respect to the distribution π if

$$\lim_{n\to\infty} \frac{\sum_{j=1}^{n} h(x_j)\omega_j}{\sum_{j=1}^{n}\omega_j} = \mathrm{E}_{\pi}(h(X)),$$

for any integrable function h. This means that if $x_1,x_2,\ldots,x_n$ are a sample from g, the set of weights $\omega_j(x_j) = \pi(x_j)/g(x_j)$ properly weights the sample with respect to π.

3.3 Sampling from the Posterior Distribution of Maps

Our IS approach relies on the factorization of $P(\boldsymbol{Z},\boldsymbol{\theta}|\boldsymbol{U},\boldsymbol{V})$ given by $P(\boldsymbol{\theta}|\boldsymbol{U},\boldsymbol{V})\; P(\boldsymbol{Z}|\boldsymbol{\theta},\boldsymbol{U},\boldsymbol{V})$. This decomposition suggests sampling $\boldsymbol{\theta}$ from its posterior distribution, $P(\boldsymbol{\theta}|\boldsymbol{U},\boldsymbol{V})$, first, and sampling $\boldsymbol{Z}$ from $P(\boldsymbol{Z}|\boldsymbol{\theta},\boldsymbol{U},\boldsymbol{V})$ afterwards.

In the case of $P(\boldsymbol{\theta}|\boldsymbol{U},\boldsymbol{V})$, we approximate this distribution by the product of the Perception Model and the geometric prior distribution of $\boldsymbol{\theta}$. Thus, according to the model described above, we sample values of θ_{ti} from a truncated Gaussian distribution centered at the corresponding laser reading V_{ti} and standard deviation σ, and then we associate to each sample a weight given by $\omega(\theta_{ti}) = (1-p)^{\theta_{ti}-1}\, p$. This set of weights properly weight the set of samples, with respect to the posterior distribution of θ.

Importance sampling plays, again, a key role when sampling locations from $P(\boldsymbol{Z} \mid \boldsymbol{\theta}, \boldsymbol{U}, \boldsymbol{V})$. Using the properties of the model described before, we can show that (see [2] for details)

$$P(\boldsymbol{Z}|\boldsymbol{\theta},\boldsymbol{U},\boldsymbol{V}) \,\dot{\propto}\, \prod_{t=1}^{T} P(Z_t|U_t, Z_{t-1})\; P(\theta_t|Z_t, \boldsymbol{Z}^{t-1}, \boldsymbol{\theta}^{t-1}). \tag{3}$$

Equation (3) shows that, at each point in time, we can sample Z_t from the Motion Model and associate a weight

$$\begin{aligned}\omega(Z_t) &= P(\theta_t|Z_t, \boldsymbol{Z}^{t-1}, \boldsymbol{\theta}^{t-1}) \\ &= \prod_{i=1}^{N} P(\theta_{ti}|Z_t, \boldsymbol{Z}^{t-1}, \boldsymbol{\theta}^{t-1}),\end{aligned} \tag{4}$$

to this observation. The term $P(\theta_t|Z_t, \boldsymbol{Z}^{t-1}, \boldsymbol{\theta}^{t-1})$ corresponds to a truncated geometric distribution that represents the degree of agreement between true distances to obstacles at time t, θ_t , and the fact that the robot is at the sampled location, Z_t, within the map built from $\boldsymbol{Z}^{t-1}$ and $\boldsymbol{\theta}^{t-1}$.

4 Results

In this section we show an application of the algorithm described in previous section, to a data set obtained in Wean Hall building at Carnegie Mellon University. The data set was collected by a robot equipped with a laser sensor and an odometer[3].

The robot navigated going back and forth along a hallway. In that journey 3354 measurements were taken, each of them consisting in a pair of odometer reading differences and laser readings. The laser sensor sends beams every degree spanning an angle of 180^o. Thus, there are $N = 180$ distances recorded for each laser reading. A map drawn from raw data is shown in Figure 1. The figure shows how odometry error accumulates so that it seems that the robot has visited two different corridors, instead of just one, as it did. The smoothness of the depicted walls, however, suggests that error in laser sensor readings is small, compared to error in odometer readings.

We sampled distances to obstacles based on the Perception Model, with $\hat{\sigma} = 0.02m$, and sampled locations afterwards, using IS with sample size $n = 100$. Figure 3 shows a path obtained from a set of sampled locations, along with

[3] Data is a courtesy of Nicholas Roy.

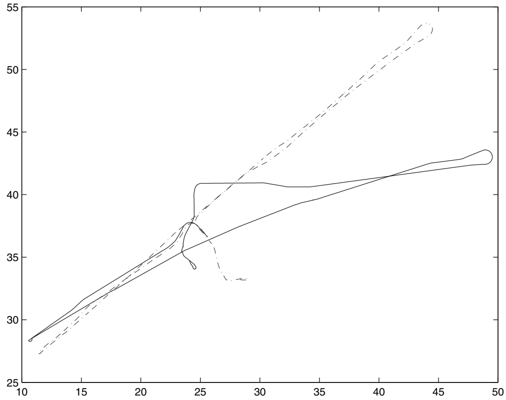

Fig. 3. Path of the robot in Raw data (line) and Sampled Path (dotted)

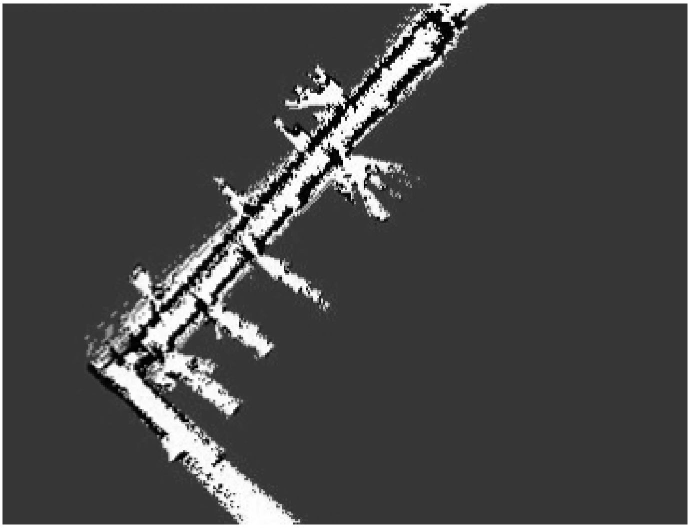

Fig. 4. Average map using the proposed Algorithm

the path obtained from raw odometer readings. The figure shows that sampling using IS allows the robot to recover from odometry errors. Figure 4 shows the average map of the sample.

5 Conclusions

This paper adds to the research in Statistics and Probabilistic Robotics, proposing a complete probabilistic representation of the SLAM problem and obtaining a full Bayesian solution. In particular, it contributes with a new algorithm to sample from the posterior distribution of maps. An intermediate step of the algorithm provides observations from the posterior distribution of locations, that is, we solve the localization problem at the same time.

This paper formalizes the problem of mapping as the problem of learning the posterior distribution of the map given the data. We work on an expression for this distribution and show that there is no closed form for it. Thus, we propose an algorithm based on Importance Sampling for obtaining a sample from the target posterior distribution.

Important Sampling showed to be a computational efficient way to explore the posterior distribution of the map. In addition to this, IS provided with an effective methodology to correct odometry error accumulated over time.

Although we do not have ground truth data to quantify the accuracy of the algorithm, the average of the resulting map samples closely resembles the real map of the environment. In the same way, samples from the robot trajectory resemble the true path followed by the robot.

References

1. A. Araneda. *Statistical Inference in Mapping and Localization for Mobile Robots*. PhD thesis, Thesis Proposal. Dept. of Statistics, Carnegie Mellon University, 2000.
2. A. Araneda. *Statistical Inference in Mapping and Localization for Mobile Robots*. PhD thesis, Department of Statistics, Carnegie Mellon University, 2004.
3. W. Burgard, A.B. Cremers, D. Fox, D. Hähnel, G. Lakemeyer, D. Schulz, W. Steiner, and S. Thrun. Experiences with and interactive museum tour-guide robot. *Artificial Intelligence*, 114:3–55, 1999.
4. A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the EM algorithm (with discussion). *Journal of the Royal Statistical Society, Series B*, 39:1–38, 1977.
5. A. Elfes. *Occupancy Grids: A Probabilistic Framework for Robot Perception and Navigation*. PhD thesis, Department of Electrical and Computer Engineering, Carnegie Mellon University, 1989.
6. D. Fox, W. Burgard, and S. Thrun. Markov localization for mobile robots in dynamic environments. *Journal of Artificial Intelligence Research*, 11, 1999.
7. J. Geweke. Bayesian inference in econometric models using Monte Carlo integration. *Econometrica*, 57:1317–1339, 1989.
8. D. Hähnel, D. Fox, W. Burgard, and S. Thron. An efficient FLOTSAM algorithm for generating maps of large-scale cyclic environments from raw laser range measurements. 2002.
9. D. Hahnel, D. Schulz, and Burgard. W. Map building with mobile robots in populated environments. 2002.
10. P. Jensfelt, H.I. Christensen, and G. Zunino. Integrated systems for Mapping and Localization. IEEE, May 2002.
11. S.J. Julier and J.K. Uhlmann. A new extension of the Kalman Filter to nonlinear systems. 1997.
12. J. Liu and R. Chen. Sequential Monte Carlo methods for dynamic systems. *Journal of the American Statistical Association*, 93, 1998.

13. M. Montemerlo and S. Thrun. Simultaneous localization and mapping with unknown data association using FastSLAM. Submitted for publication, 2002.
14. K. Murphy. *Dynamic Bayesian Networks: Representation, Inference and Learning.* PhD thesis, Computer Science Division. U. C. Berkeley, 2002.
15. R. Simmons and S. Koenig. Probabilistic robot navigation in partially observable environments. pages 1080–1087, 1995.
16. Randall Smith, Matthew Self, and Peter Cheeseman. *Autonomous Robot Vehicles*, chapter Estimating Uncertain Spatial Relationships in Robotics, pages 167–193. Springer-Verlag, 1990.
17. M. Tanner. *Tools for Statistical Inference.* New York: Springer-Verlag, 1996.
18. S. Thrun. Exploration and model building in mobile robot domains. pages 175–180, 1993.
19. S. Thrun. Learning metric-topological maps for indoor mobile robot navigation. *Artificial Intelligence*, 99:21–71, 1998.
20. S. Thrun. Robotic mapping: A survey. Morgan Kaufmann, 2002.
21. S. Thrun, M. Beetz, M. Bennewitz, W. Burgard, A.B. Cremers, F. Dellaert, D. Fox, D. Hähnel, C. Rosenberg, N. Roy, J. Schulte, and D. Schulz. Probabilistic algorithms and the interactive museum tour-guide robot Minerva. *International Journal of Robotics Research*, 19(11):972–999, 2000.
22. Sebastian Thrun, Wolfram Burgard, and Dieter Fox. A probabilistic approach to concurrent mapping and localization for mobile robots. *Machine Learning and Autonomous Robots*, 31/5:1–25, 1998. Joint Issue.