LUND UNIVERSITY

Extending the Occupancy Grid Concept for Low-Cost Sensor Based SLAM

Nordh, Jerker; Berntorp, Karl

Published in:
10th IFAC Symposium on Robot Control

DOI:
[10.3182/20120905-3-HR-2030.00079](https://doi.org/10.3182/20120905-3-HR-2030.00079)

2012

[Link to publication](#)

Citation for published version (APA):
Nordh, J., & Berntorp, K. (2012). Extending the Occupancy Grid Concept for Low-Cost Sensor Based SLAM. In *10th IFAC Symposium on Robot Control* (pp. 151-156). IFAC. <https://doi.org/10.3182/20120905-3-HR-2030.00079>

Total number of authors:
2

General rights

Unless other specific re-use rights are stated the following general rights apply:
Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Read more about Creative commons licenses: <https://creativecommons.org/licenses/>

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

LUND UNIVERSITY

PO Box 117
221 00 Lund
+46 46-222 00 00

Extending the Occupancy Grid Concept for Low-Cost Sensor-Based SLAM

Jerker Nordh* Karl Berntorp*

* *Lund University, Lund, Sweden (e-mail:
Firstname.Lastname@control.lth.se).*

Abstract: The simultaneous localization and mapping problem is approached by using an ultrasound sensor and wheel encoders. To account for the low precision inherent in ultrasound sensors, the occupancy grid notion is extended. The extension takes into consideration with which angle the sensor is pointing, to compensate for the issue that an object is not necessarily detectable from all positions due to deficiencies in how ultrasonic range sensors work. A mixed linear/nonlinear model is derived for future use in Rao-Blackwellized particle smoothing.

Keywords: SLAM, Particle filtering, Kalman filtering, Occupancy grids, Mobile robots, Rao-Blackwellized

1. INTRODUCTION

The problem of having a robot simultaneously localize itself and learn its map is commonly referred to as simultaneous localization and mapping (SLAM), and is still considered a challenging problem. The problem is often solved using odometry readings in combination with vision or range sensors. In mobile robotics it has been studied extensively over the last three decades. For surveys and tutorials of the SLAM problem and its different solutions up to recently, see for example (Thrun, 2002) or (Durrant-Whyte and Bailey, 2006).

At least since the early 1990s the approach to SLAM has been probabilistic, and one of the earliest works on this was presented in (Smith et al., 1990), where extended Kalman filtering (EKF) was used for state estimation. One of the problems with using Kalman filtering is that the nonlinearities that typically are present tend to lead to divergence of the state estimates. For example, the kinematics of a planar robot is typically nonlinear in the heading angle, and the consequent linearizations that the EKF uses for estimating the odometry may lead to instability. To remedy this, particle filtering was introduced as a means to solve the SLAM problem; the reader is referred to (Grisetti et al., 2005) and (Grisetti et al., 2007) for state of the art algorithms.

Several approaches exist of how to represent the map, where two possible approaches are metric and topological. In the metric approaches, which this paper will focus on, the maps capture the geometric properties of the environment, while the topological maps try to describe the connectivity of different places using graphs, see (Thrun, 1998). Perhaps the most popular representative of the metric approaches is known as occupancy grid mapping. In this representation the space is described by evenly spaced

grids, see (Siciliano and Khatib, 2008) for an introduction. The grids are considered to be either occupied or free, with some probability distribution associated with the grid. A possible usage of occupancy grid maps is when utilizing range sensors, such as laser sensors or sonar sensors. Both types of sensors have noise and may occasionally give severe measurement errors. Since laser sensors have very high spatial resolution, thus giving a sharp probability distribution, they appear to be the most common solution, see (Hhnel et al., 2003), (Eliazar and Parr, 2003), and (Grisetti et al., 2007) for some examples. In contrast, sonar sensors have the problem of covering a cone in space, which typically makes it impossible to determine from a single measurement whether a certain cell is occupied or not because of the low spatial resolution in the tangential direction. Also, ultrasound sensors are very sensitive to the angle of an objects surface relative to the sensor. This leads to the problem that measuring the same surface from slightly different angles may render different results. Obviously, this could potentially lead to estimation errors. See Fig. 1 and Section 2 for a more detailed description of the problems encountered with ultrasound based range sensors.

In this paper the SLAM problem is approached using only wheel encoder readings and one ultrasonic sensor. To get rid of, or at least attenuate, the problems inherent in ultrasonic sensors described earlier, a new approach to grid mapping is developed. This should be seen as an extension to the notion of occupancy grids described in (Siciliano and Khatib, 2008), in the sense that the angle with which the sensor is facing the cell is now taken into account. Particle filtering is used for position estimation, and each particle represents a possible robot position and a map. It is known that using particle filters for SLAM tends to destroy the map over time caused by sample depletion, see (Kwak et al., 2007) for an investigation. However, an idea is that particle smoothing could be a way to get rid of this problem. To prepare for exploiting the ideas of Rao-Blackwellized particle smoothing established in (Lindsten

* This work was sponsored by the ELLIIT excellence center and by the Swedish Foundation for Strategic Research through the project ENGROSS. We thank professor Bo Bernhardsson for the discussions and feedback he has provided.

and Schön, 2011), a mixed linear/nonlinear state-space model is developed. Compared to the regular occupancy grid this formulation also represents the variance of each probability estimate.

Using sonar sensors for SLAM has been studied before; an example is (Burgard et al., 1999), in which an offline expectation maximization algorithm was used for occupancy grid mapping using 24 Polaroid sensors with 15 degrees opening angle. An early work is (Rencken, 1993), where the SLAM problem was solved in simulation using 24 ultrasonic sensors by estimating the errors introduced in the localization and mapping parts, respectively, and correcting for them using a modified Kalman filter approach. A third example is (Leonard et al., 1992) in which the SLAM problem was solved using a feature based approach with the aid of servo-mounted ultrasonic sensors. A more recent work is (Ahn et al., 2008), where ultrasonic sensors and a stereo camera is used in an EKF-SLAM setting.

The paper is organized as follows: The SLAM problem is introduced in Section 2, where the difficulties with ultrasonic sensors is explained in more detail. Section 3 details the scope of the work presented in this article. In Section 4 the kinematics, sensor, map modeling, as well as the SLAM algorithm are introduced. Implementational aspects are discussed in Section 5. The validation results are shown in Section 6. Finally, the paper is concluded in Section 7.

2. PROBLEM FORMULATION

As previously mentioned there are two major systematic issues with ultrasonic sensor that have been observed:

- (1) Large field of view; making it uncertain from which point a measurement originated. This leads to a fundamental limitation on the resolution of the sensor, which varies with the distance to the object. Because of the typically large opening angle of ultrasonic sensors the resolution is of the same order as the distance to the detected object. That is, when detecting an object at a range of 1 meter, the spatial resolution of the sensor is roughly 1 meter in the tangential direction.
- (2) Angle of incidence; for angles above a certain threshold the object becomes more or less invisible for the sensor. Thus an object can only be detected from certain directions. This problem becomes apparent when the sensor is close to a wall and measuring along it. Typically the wall will be inside the field of view of the sensor, but the sensor will not detect it because of a too narrow angle of incidence. If the same wall is then measured from another position within the room where the sensor faces it perpendicularly, it will be detected. Note that the angle is dependent on the material of the observed object.

See Fig. 1 for an illustration.

The basic premise for the SLAM problem is that it is possible to reobserve parts of the environment, thus relating the current position to those before. Therefore if it is not possible to observe objects that have been detected previously, it will lead to inconsistencies in the map as well

Fig. 1. Properties of ultrasonic sensors. The sensor has an opening angle, β . The sound only reflect back to the sensor if the angle of incidence is less than α . The figure illustrates a straight wall. Sections marked 3 are outside the field of view, the section marked 2 is inside the field of view but the angle of incidence is so large that it is not visible to the sensor. The section marked 1 is inside the field of view. Also, the angle of incidence is steep enough for it to be detected by the sensor. The range returned by the sensor will be the distance to the closest point within section 1, marked with an X .

as inaccuracy in the position estimate. This article aims to provide a method for dealing with these systematic errors.

3. ARTICLE SCOPE

The method presented in this paper extends the concept of occupancy grids to take the angle of incidence into consideration and partitioning each cell of the map into several parts, where each part is visible only from certain regions within the environment. This reduces the problem of conflicting measurements. The drawback is that it also reduces the correlation between measurements, and thus the underlying SLAM algorithm will require more data to converge.

3.1 Computational Complexity

Particle methods are sensitive to the number of states in the model as the number of particles needed to represent the probability density function explodes with the number of states. Therefore a conditionally linear model is very beneficial for reducing the computational burden. The method presented in this paper is conditionally linear given the position and orientation of the sensor, thus for the planar case only 3 nonlinear states are needed. The number of linear states depends on the size and resolution of the map.

3.2 Evaluation

To evaluate the model, simulated data corresponding to different sensors characteristics were generated. The data sets were used with different levels of subdivision of the grid cells, showing the methods strengths and weaknesses for a selection of sensor characteristics. The focus is on investigating how the angle-of-incidence limitations on the sensor affect the position estimate of the SLAM algorithm.

The goal is position estimation, the map is merely a tool. Therefore the map estimates are not presented.

The results presented are generated using more pessimistic noise values than what could be expected for even very inexpensive sensor, typically available from hobby electronics suppliers for a few tens of dollars. The amount of noise introduced by the wheel encoders is also believed to be exaggerated. The interesting quantity to study is the relative performance of using the simulated ultrasonic sensor for SLAM, with angle-of-incidence dependence, compared with relying only on dead reckoning. The absolute positioning error is therefore not as relevant and all the results presented in the article are normalized against the dead-reckoning scheme.

The model presented contains a large number of parameters, but the work in this paper does not address how to optimally choose them. Rather, the same set of parameters for all parametrizations of the model are used to show the relative merits of the method. It is surely possible to improve the results by better parameter choices and more sophisticated particle methods, but that is outside the scope of this paper. More effective methods for storing the map could be implemented. This is however not done as this is currently only a proof of concept.

The concept of modeling the angle of incidence could also be applied to methods for the SLAM problem which are not particle based.

4. MODELING

4.1 State-Space Model

The robot used is a differential-driven mobile robot, with the sensor placed at the front end. Since the robot moves in a plane, only three states are needed to describe the motion. Using the position variables p_x and p_y , as well as the heading θ as state variables, the robot's kinematics is

$$\dot{x}_r = \begin{pmatrix} \dot{p}_x \\ \dot{p}_y \\ \dot{\theta} \end{pmatrix} = \frac{R}{2} \begin{pmatrix} \cos \theta & \cos \theta \\ \sin \theta & \sin \theta \\ -2B^{-1} & 2B^{-1} \end{pmatrix} \left[\begin{pmatrix} \omega_l \\ \omega_r \end{pmatrix} + \begin{pmatrix} v_l \\ v_r \end{pmatrix} \right] \quad (1)$$

which in short is written as $\dot{x}_r = f(x_r, \omega, v)$. Here, $\omega_{l,r}$ is the left and right wheel angular velocity which have noise $v_{l,r}$, R is the wheel radius, and B is the distance between them. To be implementationally useful, this model is discretized using a second order approximation.

By assuming that the map is slowly time varying, the map can be modeled as a constant position model. Every cell in the occupancy grid representation is a state. Since each state represents a probability it should only take values between zero and one, representing the probability that a cell is occupied. The discretized occupancy grid model is on the form

$$x_{t+1}^m = x_t^m + v_m, \quad (2)$$

where each x^m is a cell in the map and v_m is the noise on the states. This representation allows that the map varies over time meaning that old measurements should not be given as much weight as more recent measurements. Taking the angle of incidence into account gives rise to additional states. As an example, assume that a map with two cells is used, and assume further that the angle

dependence has a resolution of 45 degrees. This means that the cell can be viewed from 8 different angles, which gives that there will be 16 states in total for representing this simple map.

4.2 Measurement Model

Assume that an ultrasound measurement returns a distance, and that the opening angle of the sensor is β degrees. The field of view is then a closed cone with aperture 2β . The cells that are inside the field of view can now be calculated, given that a position estimate exists. Assuming that a method exists for directly measuring the states of each cell inside the field of view the measurement equation would be linear, and the C -matrix in the equation $y = Cx$ would be sparse with a single 1 per row, and the same number of rows as the number of cells inside the field of view.

The question of how to generate the measurements, y , is not trivial. A common way to generate the measurements is to create a probability function with a peak in the center of the cone, depending on the distance from the robot that the sensor returned. The probability function then decays with the angle from the center of the cone and the distance from the most probable cell, reaching the nominal value of occupancy at the end of the cone. This is called evidence mapping. Here, the focus is instead on exploiting variance. Each measurement is given a variance, dependent on which distance is returned from the sensor. If the distance is short the variance should be low, since fewer cells are visible. Furthermore, the cells' angles with respect to the center of the cone is also influencing the variance. For cells at a given distance the probability is the same. This is of course an approximation, as a single distance measurement does not contain information about the individual cells. However, it provides the conditionally linear formulation that is desired.

A typical C matrix for the example in Section 4.1 could be

$$y_t = \begin{pmatrix} 1/2 \\ 1/2 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 1 & 0 & \cdots & 0 \end{pmatrix} x_t^m + w_t = C_t x_t^m + w_t. \quad (3)$$

In this particular example both cells were equally probable to be occupied since they were at the same distance from the sensor. The cells not inside the field of view at time t does not generate any measurements at all for time t . The measurement noise at time t , w_t , is a tuning parameter, parametrized by the angle and distance as previously discussed.

4.3 Map Estimation

Since the approximated measurement model is linear, a Kalman filter approach will yield the optimal result. The Kalman filter equations are written out next with the notation from (2) and (3). The reader is referred to (Anderson and Moore, 1979) for a thorough investigation on linear filtering.

$$\hat{x}_{t|t-1}^m = \hat{x}_{t-1|t-1}^m \quad (4)$$

$$P_{t|t-1} = P_{t-1|t-1} + Q_n \quad (5)$$

$$S_t = C_t P_{t|t-1} C_t^T + R_n \quad (6)$$

$$K_t = P_{t|t-1} C_t^T S_t^{-1} \quad (7)$$

$$e_t = y_t - C_t \hat{x}_{t|t-1}^m \quad (8)$$

$$\hat{x}_{t|t}^m = \hat{x}_{t|t-1}^m + K_t e_t \quad (9)$$

$$P_{t|t} = (I - K_t C_t) P_{t|t-1} \quad (10)$$

Here, $\hat{x}^m$ are the linear states and P is the estimated covariance matrix. The matrices Q_n and R_n are the variances of the process noise and measurement noise, respectively.

This is a convenient result. First of all, it means that the solution is analytic. Secondly, the computation of the distribution function, characterized by the mean $\hat{x}_{t|t}^m$ and the variance $P_{t|t}$, is relatively cheap compared to if a particle filter would be used for map estimation.

4.4 Pose Estimation

As already mentioned, particle filtering is used to estimate the robot's pose. A good tutorial on particle filters is (Arulampalam et al., 2002); only the parts relevant for this paper will be explained next: The key idea with the particle filter is to approximate the density function p as a weighted set of samples

$$p_N(x_t|y_t) = \sum_{i=1}^N w_t^i \delta(x_t - x_t^i), \sum_{i=1}^N w_t^i = 1, w_t^i \geq 0, \forall i, \quad (11)$$

where δ is the Dirac function used to represent the empirical probability density function with support in N discrete samples. The weights w^i reflects the importance of each particle. These values are updated when new measurements arrive.

Each of the N particles represent a filtered trajectory estimate. To proceed in the algorithm the particles are simulated forwards one step using the input signals, yielding

$$\hat{x}_t^i = f(x_{t-1}^i, \omega_{t-1}, v_{t-1}), \quad i = 1, \dots, N. \quad (12)$$

The weights are normally calculated according to the measurement density error function $p(y_t|x_t)$. However, due to the parametrization used for the measurements, the weights are now instead updated as

$$w_t^i = (e_t^i)^T S_t^{-1} e_t^i, \quad i = 1, \dots, N, \quad (13)$$

$$w_t^i = w_t^i / \sum_{i=1}^N w_t^i. \quad (14)$$

The weighting in (13) stemming from (4)-(10) reflects how likely the particle was, given the ultrasound measurement. The state estimate may now be formed as

$$\hat{x}_t = \sum_{i=1}^N w_t^i \hat{x}_t^i.$$

Since all particles but a few will have negligible weights after a while, new particles are drawn with replacement according to the sampling importance resampling principle

$$P(x_t^i = \hat{x}_t^j) = w_t^j, \quad j = 1, \dots, N. \quad (15)$$

Note that this step is only performed given that the effective sampling number $1/(\sum_{i=1}^N (w_t^i)^2) < N_{\text{eff}}$, where

N_{eff} is a number between zero and N . To summarize: The SLAM algorithm is run by first calculating (3) as explained in Section 4.2, and then running (4)-(15).

5. IMPLEMENTATION DETAILS

There are a number of parameters to be chosen when implementing the model described in Section 4. In addition to this a number of minor tweaks to help with the implementation have been made.

Each particle contains a Kalman filter that estimates the probability of occupancy for each cell in the grid. Instead of storing this value directly it is first converted to log-odds format and thus spans the entire range of real numbers instead of only the interval $(0, 1)$. All measurements are also converted to this format before being passed to the Kalman filter. The covariance matrix grows quadratically with the map size, therefore it is critical to exploit the sparsity that occurs. Under the assumption that all cells are independent only the diagonal elements will be nonzero.

Because of the large field of view of the sensor it rarely provides accurate information on the position of objects. However, it provides information on which cells are likely to be empty. To exploit this fact a threshold value has been added, and whenever the measured range exceeds this value only the cells that are detected as empty are included in the measurement update. For closer distances also the cells that are on the edge of the cone and thus likely to be occupied are included in the measurement update.

Since the model not only contains an estimate of the probability of occupancy, but also an estimate of how certain that estimate is, it is possible to have a very high prior probability of occupancy since the map will quickly converge to the new estimate once a measurement has been made. This allows the algorithm to successively clear areas of the map, relying more on the information on which areas are empty rather than those that are occupied. This is a better match for the characteristics of an ultrasound sensor.

6. RESULTS

The focus of this section is to show the relative merits of the method for different sensor characteristics, and to give an indication of what performance different number of angular subdivisions of each cell provide. Therefore the figures are normalized such that the absolute position error of the dead reckoning is 1 at time 1.

The data used for the simulations were generated by driving a robot in a simulated environment consisting of a number of walls with different orientations and lengths. The trajectory was a circle repeated roughly twenty times, with the sensor sweeping back and forth through 180° in front of the robot.

The data presented in this section are the absolute errors of the position estimates. As a reference comparison for the SLAM algorithm the same number of particles (200) was simulated using the same noise model, but only using the information from the wheel encoders. The position of the dead-reckoning estimate was then taken as the mean of these particles. For the SLAM estimate the position was

Fig. 2. Comparison of the absolute position error for 3 different sensors with different characteristics and number of angular subdivisions for each grid cell. All the simulated sensors have an opening angle of 60° . The maximum angle of incidence is denoted α , the number of subdivisions is denoted d . Note that $d = 1$ is similar to a normal occupancy grid approach. All figures are normalized such that the total accumulated error of the dead reckoning is 1 at time 1. The red line is the accumulated error for the dead-reckoning. For each simulation the dead reckoning is evaluated as the mean of 200 particles simulated using the same noise model as that provided to the SLAM algorithm. The blue line corresponds to the same quantities when using the SLAM algorithm. The position of the SLAM estimate is taken as the weighted average of 200 simulated particles. The dot-dashed line is the 1 standard-deviation interval estimated by repeatedly running the Monte-Carlo simulation but with different noise realizations for each iteration. From the results shown it can be seen that for all the sensor characteristics it is beneficial to increase the number of subdivisions, d , for each cell. Counter-intuitively this also seems to hold for the 'perfect' sensor where the angle of incidence, α , does not affect the measurement.

taken as a weighted mean with the weight at each instant being proportional to how good the particular particle is deemed to be, with the particles being resampled whenever the weights are considered to be too unevenly distributed.

The results presented in Fig. 2 are the mean of the estimates and standard deviations calculated by repeated Monte-Carlo simulations using different noise perturbations of the correct input. As can be seen in the figure the SLAM estimate is clearly improved by subdividing each cell taking the angle of incidence into consideration. Most notable is the significant decrease in the variance of the estimate, the mean error is roughly the same for all the different subdivisions.

Unexpectedly this method also seems to improve the estimates when using a simulated sensor with no dependence on the angle of incidence, which can be seen in the right-most column in Fig. 2. It is believed that this could be explained from the fact that the wide field of view of the sensor is still a limiting factor of the performance, and that using the angular subdivision scheme increases the resolution by separating the measurements from different positions. This effect, however, might not always be desirable since it also decreases the correlation between different robot poses, making it harder for the particle filter to converge.

7. CONCLUSIONS AND FUTURE WORK

The results presented in this paper show that by extending the concept of the regular occupancy grid to model the fact that objects are not necessarily detected from all positions within a room, the performance for a particle filtering based SLAM algorithm can be significantly improved. The method presented here can thus be seen as a way to trade sensor performance against computational resources. In a world with computing power becoming cheaper day by day, this tradeoff will become more relevant as time progresses, perhaps allowing the use of SLAM algorithms in inexpensive consumer products.

The model presented in this paper is clearly suited for Rao-Blackwellized particle methods due to its conditional linearity, and in the future the implementation will most likely be extended with smoothing methods, such as those presented by (Lindsten and Schön, 2011). It is believed that smoothing will decrease the position error since a single measurement contains fairly little information of the robot pose. Also, the problem with particle depletion should be possible to solve with particle smoothing. Therefore a method which takes the trajectory into consideration is likely to perform significantly better, but with larger computational burden

More extensive simulations and individual tweaking of the parameters for the different number of subdivision of the cells would be beneficial for clearly establishing the methods merits, but our initial results seem promising.

We will gather actual measurement data under conditions where we have access to a reliable ground truth. In this way it is possible to provide proper estimates on the accuracy of the method under real-world conditions.

REFERENCES

- Ahn, S., Choi, J., Doh, N.L., and Chung, W.K. (2008). A practical approach for EKF-SLAM in an indoor environment: fusing ultrasonic sensors and stereo camera. *Auton. Robots*, 24, 315–335.
- Anderson, B.D.O. and Moore, J.B. (1979). *Optimal Filtering*. Prentice Hall.
- Arulampalam, M., Maskell, S., Gordon, N., and Clapp, T. (2002). A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking. *Signal Processing, IEEE Trans. on*, 50(2), 174–188.
- Burgard, W., Fox, D., Jans, H., Matenar, C., and Thrun, S. (1999). Sonar-based mapping with mobile robots using EM.
- Durrant-Whyte, H. and Bailey, T. (2006). Simultaneous localisation and mapping (SLAM): Part 2. *IEEE ROBOTICS AND AUTOMATION MAGAZINE*, 2, 2006.
- Eliazar, A. and Parr, R. (2003). DP-SLAM: Fast, robust simultaneous localization and mapping without predetermined landmarks. In *in Proc. 18th Int. Joint Conf. on Artificial Intelligence (IJCAI-03)*, 1135–1142. Morgan Kaufmann.
- Grisetti, G., Stachniss, C., and Burgard, W. (2005). Improving grid-based SLAM with Rao-Blackwellized particle filters by adaptive proposals and selective resampling. In *In Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2443–2448.
- Grisetti, G., Stachniss, C., and Burgard, W. (2007). Improved techniques for grid mapping with Rao-Blackwellized particle filters. *IEEE Trans. on Robotics*, 23, 2007.
- Hhnel, D., Burgard, W., Fox, D., and Thrun, S. (2003). A highly efficient FastSLAM algorithm for generating cyclic maps of large-scale environments from raw laser range measurements, 206–211.
- Kwak, N., Kim, I.K., Lee, H.C., and Lee, B.H. (2007). Analysis of resampling process for the particle depletion problem in FastSLAM. *Robot and Human interactive Communication, 2007. RO-MAN 2007. The 16th IEEE Int. Symposium on*, 200–205.
- Leonard, J.J., Durrant-Whyte, H.F., and Cox, I.J. (1992). Dynamic map building for an autonomous mobile robot. *Int. J. Rob. Res.*, 11, 286–298.
- Lindsten, F. and Schön, T. (2011). Rao-Blackwellised particle smoothers for mixed linear/nonlinear state-space models. Technical Report.
- Rencken, W. (1993). Concurrent localisation and map building for mobile robots using ultrasonic sensors. In *Proc. of the 1993 IEEE/RSJ Int. Conf. on IROS '93*, volume 3, 2192–2197 vol.3.
- Siciliano, B. and Khatib, O. (eds.) (2008). *Springer Handbook of Robotics*. Springer, Berlin, Heidelberg.
- Smith, R., Self, M., and Cheeseman, P. (1990). *Estimating uncertain spatial relationships in robotics*, 167–193. Springer-Verlag New York, Inc., New York, NY, USA.
- Thrun, S. (1998). Learning metric-topological maps for indoor mobile robot navigation. *Artif. Intell.*, 99(1), 21–71.
- Thrun, S. (2002). Robotic mapping: A survey. In *Exploring Artificial Intelligence in the New Millenium*. Morgan Kaufmann.