An Experimental Comparison of Localisation Methods, the MHL Sessions

Steen Kristensen
DaimlerChrysler Research and Technology
Cognition and Robotics
Alt-Moabit 96A, D-10559 Berlin
Germany
steen.kristensen@daimlerchrysler.com

Patric Jensfelt
Centre for Autonomous Systems
Royal Institute of Technology
SE-100 44 Stockholm
Sweden
patric@nada.kth.se

***Abstract*— In this paper we compare multi hypothesis localisation (MHL)—which is a mobile robot localisation method based on multi hypothesis tracking—with six other methods reported in the literature. The comparison is performed using a standard set of test data and corresponding evaluation tools, thus facilitating a direct comparison of the obtained results. The experiments show that MHL compares favourably to all other methods in terms of recovering when the robot has been kidnapped. When using a validation gate for filtering out noisy measurements, MHL and the standard extended Kalman filter both perform as well as all other reported methods in terms of accuracy while being faster to compute.**

I. Introduction

Localisation is one of the most important problems in mobile robotics. Thus a large number of localisation methods have been proposed over the last years. However, as is often the case in robotics, scientific comparisons and thus progress has been hampered by researches using different sensors and locomotion systems and by having tested their algorithms in different environments.

A noteworthy exception was therefore when Gutmann *et al.* [9], [10] decided to test a number of localisation algorithms on the same problem. They did that by recording a one hour run of a robot in a RoboCup-like environment with known landmarks and a system for recording ground truth. Using this recording, six different localisation methods were tested as reported in [10]. The tests were conducted to determine accuracy, robustness to noisy and sparse data and the ability to handle the *kidnapped robot problem*.

Since it is unrealistic to expect one pair of researchers to test all available localisation methods, and since it is probably scientifically sounder to let people test their "own" methods, Gutmann and Fox explicitly left out testing multi hypothesis localisation (MHL) methods which use a set of Gaussians (a Gaussian mixture) to represent the pose[1] of the robot. Instead they encouraged other researches to perform this test. This paper presents the results of exactly that endeavour.

[1]position and orientation

Gutmann and Fox kindly provided us with the data files they used for testing their algorithms and with the corresponding evaluation tools. We therefore believe that the results presented in this paper can directly be compared with the results presented in [10].

II. Related Work

As mentioned in the Introduction, a large body of research towards solving the mobile robot localisation problem exists. The two most commonly used groups of probabilistic localization methods are the Kalman filter based methods and the Monte Carlo localisation methods. Some examples of Kalman filer based localisation are found in e.g., [5], [11], [3], [16]. For examples of the latter category we refer to [17], [7]. For a general overview of the most common methods for mobile robot localisation see for example [2], [12].

Since it is outside the scope of this paper to review and evaluate all of this, we will concentrate on the methods presented in [10] which are the ones we compare us with and which furthermore represent well state of the art in probabilistic robot localisation.

Gutmann and Fox present and test the following methods:

Extended Kalman filter (EKF) The EKF is the "classical" probabilistic method for integrating uncertain information for the purpose of pose tracking [1], [16]. The main drawback of the Kalman filter method is that its probability density function is unimodal (a Gaussian) and thus it is not well suited for representing situations where the robot pose is ambiguous. Also, since it is a model-based approach, it requires that the process can be reasonably well modelled. Advantages of the Kalman filter is that it is a precise, simple and fast method which is rather well understood.

Markov localisation combined with EKF (ML-EKF) This method combines grid based Markov Localisation (ML) with an extended Kalman filter [8]. The idea is to use the coarse but robust grid-based method to keep the overview over the situation and to "supervise" a Kalman filter which is thereby made less sensitive to model errors and the problems with its unimodal probability distribution. The ML part of this approach only maintains a 2D probability grid for the positions.

Monte Carlo Localisation (MCL) As an alternative to the fixed discretisation offered by a grid as in ML—with its associated high computational cost—sample based techniques have emerged [6]. The probability distribution is here represented by a set of weighted samples. The weight of the samples, which tell how well it matches sensor data, determine how likely they are to be re-sampled into the next iteration. The stronger ones will survive and the resources are shifted to where they are needed. Without modifications the MCL algorithm suffers from the problem that unless there already are samples in an area, this area cannot attract new samples. This leads to a situation where many samples are needed to do global localisation and where the kidnapped robot situation typically fails.

Three methods with different flavours of MCL were tested in [10] as described below:

Sensor Resetting Localisation (SRL) In SRL [15] the idea is to draw a fraction of the samples in the re-sampling not from the previous set, but instead directly based on where the measurements indicate that there should be samples. Two different parameter settings are evaluated in the experiment in [10], and are referred there to as SRL1 and SRL2. The difference in setting is the amount of samples added based on observations.

Adaptive MCL (A-MCL) Adaptive MCL [4] extends SRL with a schema for adaptively determining how many samples should be added.

Mixture MCL (Mix-MCL) Mixture MCL [18] also draws samples from the observations but the samples are properly weighted, with the probability assigned to the position where the sample is placed. This probability is typically estimated based on a grid approximation.

In previous work [13] we have shown that MHT based localisation (MHL) is efficient not only for pose tracking but also for solving the global localisation problem, i.e. to determine the robot pose without any prior information given about this pose. The method was tested in two different real world environments on two different robots. It was never experimentally verified that MHL could handle kidnapping, but this will be tested in this paper. In the work presented in this paper we made slight adaptations to the method presented in [13], mainly to account for the somewhat different environment. The resulting method is presented in detail in the following section.

III. Multi Hypothesis Localisation

The basic idea behind MHL is to alleviate the inherent problems of using a single Gaussian by using a number of these to represent the robot pose, thus effectively creating a mixture of Gaussians enabling the representation of any given probability distribution of the robot pose. Normally, each EKF tracks one hypothesis about the robot. By adding and removing such hypotheses it can furthermore be achieved that the robot pose is robustly tracked although the behaviour of the robot and the environment violates the model assumptions of the single EKFs.

Therefore, with MHL it is in principle possible to maintain the advantages of EKFs while not having to suffer from the drawbacks. The price for this is that we need a sensible method for adding and deleting hypotheses and for estimating the probability of each hypothesis being the correct one, i.e. its weight in the Gaussian mixture.

A. Data Association

One of the prerequisites for using the EKF is the that the measurement noise is zero-mean. The data association problem makes this a true challenge in localisation. Unless it is possible to correctly establish the correspondence between a measurement and a feature in the map the estimation process is at risk of breaking down.

A common way to tackle the data association problem is to use a *validation gate*. The gate defines a region within which the measurement must fall to be associated with a certain map feature. Let $\nu_{k,i}$ be the innovation and $S_{k,i}$ the measurement covariance given when matching a measurement with the ith map feature at time k.

Using the Mahalanobis distance, $\rho_{k,i}$, between the measurement and the ith map feature, the gate can be defined as

$$\rho_{k,i} = \nu_{k,i} S_{k,i}^{-1} \nu_{k,i}^T < \gamma, \quad (1)$$

where γ gives the size of the gate.

B. Hypothesis Generation

Since all hypotheses in the mixture have to be updated in each move-observation cycle it is preferable to have as few Gaussians as possible. Therefore new hypotheses should not be introduced randomly but only where there is reason to believe that the robot is, i.e. where the corresponding weight of the Gaussian would be larger that some ε, where $\varepsilon > 0$. This is only possible if the observations are used to guide the insertion of new hypotheses.

To make the description more concrete we consider the case where the landmarks in the map are of point type. Each observation, o, is a triple, (L, r, b), consisting of the landmark id, L, the measured range to the landmark, r, and the bearing, b. The position, $\mathbf{x}_L$, of each landmark is known from the map $\mathscr{M}$. If the robot from one position observes two different landmarks, A and B, it is possible to estimate the pose of the robot by calculating the intersection of the two circles with centres $\mathbf{x}_A$ and $\mathbf{x}_B$ and radii r_A and r_B, respectively. When such an intersection falls inside the area where the robot is known to move, we thus add a new hypothesis at this point. Given the position of the new hypothesis and the observation bearings, b_A and b_B, it is also possible to estimate the hypothesis angle.

Since the robot however only rarely observes two landmarks at the time, we relaxed this constraint and also generate new hypotheses when two different landmarks are seen in two consecutive observations. Here we simply ignore the movement between the two observations assuming that the resulting error on the intersections is negligible compared to the error due to the range uncertainties.

To simplify matters, each new hypothesis is initiated with a standard covariance although it is in principle possible to estimate this from the landmark positions and the observation uncertainties.

In the experiments with MHL described in Section IV all hypotheses were introduced using this mechanism, i.e. initially no pose hypotheses exist.

C. Hypothesis Probability Estimation

In this section we describe how we estimate the weight of the Gaussians, i.e. the probability that a given hypothesis is correct.

If $P(H_i)$ describes the probability of the i'th hypothesis being correct, the new probability after the receipt of an observation, o_L, indicating that landmark L has been detected, can—using Bayes's inversion formula—be written as:

$$P(H_i|o_L) = \frac{P(o_L|H_i)P(H_i)}{P(o_L)} \tag{2}$$

where $P(o_L)$ is effectively a scale factor which ensures that the $P(H_i|o_L)$'s sum to 1. This implicitly assumes that *one* of the hypotheses is correct which, however, cannot always be guaranteed. We therefore introduced a pseudo hypothesis, H_0, which is a hypothesis accounting for the probability that all the real hypotheses are incorrect. For more details on H_0, please see [13].

In equation 2, the term $P(o_L|H_i)$ expresses the probability of making the observation, o_L, given that the robot is at the position described by H_i. This is normally estimated using the map, $\mathscr{M}$, and can be re-written as:

$$\begin{aligned} P(o_L|H_i) \quad = \quad & P(o_L|f_L)P(f_L|H_i,\mathscr{M}) + \\ & P(o_L|\neg f_L)P(\neg f_L|H_i,\mathscr{M}) \end{aligned} \tag{3}$$

where $P(o_L|f_L)$ is the probability that an observation of L is generated given that landmark L is seen by a sensor and $P(f_L|H_i,\mathscr{M})$ is the probability that landmark L can be seen given that the robot is at the position described by H_i. In other words, $P(o_L|f_L)$ is a model of the reliability of the recogniser extracting landmarks from sensor data.

We dynamically estimate $P(o_L|f_L)$ as the ratio between the number of successfully and the totally tried landmark matches of the currently best hypothesis. The logic behind this is that if all observations are correct, they will all match the best (hopefully correct) hypothesis, and vice versa.

If a match between a hypothesis, H_i, and a landmark, L, has been established, i.e. if the observation passed the validation gate, we estimate $P(f_L|H_i,\mathscr{M})$ as follows: A temporary hypothesis, C, is generated with centre, $\mathbf{z}$:

$$\mathbf{z} = \mathbf{x}_L + r\vec{d}_{\mathscr{L}} \tag{4}$$

where $\vec{d}_{\mathscr{L}}$ is a unit direction vector for the line $\mathscr{L}$ connecting L and $\mathbf{x}_i$. The angle of C relative to H_i is the same as the angle between $\mathscr{L}$ and the observation bearing, b. This means that if the landmark is exactly where it would be expected according to H_i, C and H_i will coincide. The larger the deviation the further C and H_i will be apart. The covariance of C is set equal to the observation covariance with the modification that the angle uncertainty is doubled to account for the fact that forcing C onto $\mathscr{L}$ may increase the angle deviation.

Given H_i and C we can estimate of $P(f_L|H_i,\mathscr{M})$ as [14], [13]:

$$P(f_L|H_i,\mathscr{M}) = \exp^{-\frac{1}{2}(\hat{\mathbf{x}}_i-\mathbf{z})\Sigma_i^{-1}(\hat{\mathbf{x}}_i-\mathbf{z})^T} \tag{5}$$

When a H_i has not been matched to any landmark we assume that we have a random observation which has a very small probability and use $P(f_L|H_i,\mathscr{M}) = \delta$, where δ was fixed at 0.01 in our experiments.

D. Track Splitting

A standard component of MHT is track splitting. The idea is to get robust against spurious measurements by splitting hypotheses prior to updating them (whereby only the original is updated), effectively making a "backup". If the copy and the original are still close together after the update, they are merged again. If the observation was noisy, this will normally not be the case, and due to the backup, track is not lost.

E. Hypothesis Pruning

We use two standard mechanisms to prune the number of hypotheses. We delete hypotheses for which $P(H_i) < \eta$ and return its probability mass $P(H_i)$ to $P(H_0)$. We also check if two hypotheses are virtually representing the same pose. This happens for example basically every time when the track is split and the data association was correct, i.e. there is no split. To see if two hypotheses represent the same pose, we check if the Mahalanobis distance between the two hypotheses is below a certain value. If this is the case, we remove the weaker of two hypotheses, i.e. the one with the lowest $P(H_i)$. If everything is well-behaved, this may eventually lead to the situation where only one hypothesis exists.

IV. Experiments and Results

To complement the comparison made in [10] in a fair way with the results using the MHL technique, the same set of data is used. In short, the experiments are designed to test accuracy, robustness to noisy and sparse data as well as the ability to deal with the kidnapped robot problem. The latter is particulary hard as the robot falsely believes it knows where it is and thus first has to realise that it is lost.

To further ease the comparison, results will be presented for the standard EKF as well, just like in [10]. We will present results of both EKF and MHL with and without using a validation gate. From now on, let EKF and MHL denote the case when we run the methods without a validation gate and correspondingly EKF-V and MHL-V when running with it.

All experiments were conducted with a single program with one set of movement and sensor models, i.e. no parameters were changed between tests. Two switches were used: one to turn on and off the validation gate and one to turn off the hypothesis generation and track splitting, thus turning MHL into EKF. Since, in the EKF experiments, no hypotheses were automatically generated, one initial hypothesis was introduced at the known start pose of the robot.

In the experiments we used the same sensor error model as Gutmann and Fox which was 0.15 m per meter range and a fixed angle of 10 degrees.

A. The Experimental Data

The experimental data were collected in a custom made 3x2 m environment with 6 landmarks with different colour codes (see [10] for more details). In the experiments a Sony AIBO ERS 2100 equipped with a CMOS camera for detecting and identifying the landmarks was used. The robot was joysticked around in the environment for almost an hour collecting odometry and processed sensor data in the form of landmark id, range and bearing in a log file. All tests were carried out using files with data derived from this original log file. To evaluate the accuracy of the methods, a tag was added in the logfile when the robot crossed special markers on the field. In total, 48636 landmarks, 8333 motion steps and 143 ground truth positions were recorded.

B. Results

1) Sparse Data: To test performance with sparse data, data files where created with only fractions of the observations left, ranging from all observations down to 1/256th of the observations.

In Figure 1 the results from using the sparse data sets are shown. The distance errors are shown with estimated 95% error bounds. The performance degrades gracefully for all methods just like in [10]. The difference in performance between the four methods is negligible.

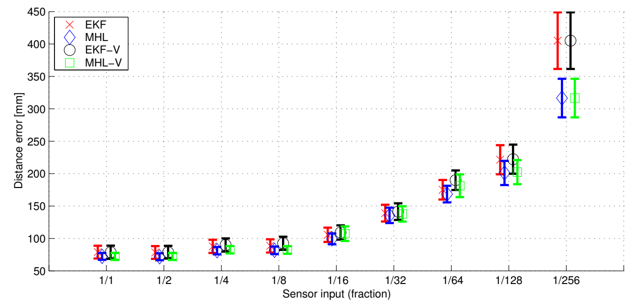

Fig. 1. The accuracy of the different methods when only a fraction of the data is used.

2) Noisy Data: Since the log files are gathered with real sensors, there is already a certain amount of naturally occuring noise in the data. However, data sets were also created where a procentage (from 10% to 80%) of the observations were replaced by random landmark observations. Note that the observations are replaced by complete outliers and not just made more uncertain by adding for example Gaussian noise to them. In reality, complete outliers might come from clutter that are recognized as a landmark, for example when there are many robots on the field.

Figure 2 shows the results for the four different methods when applied to the noisy data sets. Notice how well the EKF-V and MHL-V handle the noisy data. The importance of the validation gates can here clearly be seen when comparing with the results in [10]. It is also clear from the figure that MHL, even when used without a validation gate, improves robustness to noise. This is due the the splitting of hypotheses which reduces the risk of loosing track. The difference in performance for the EKF compared to [10] might be explained by the use of different motion models.

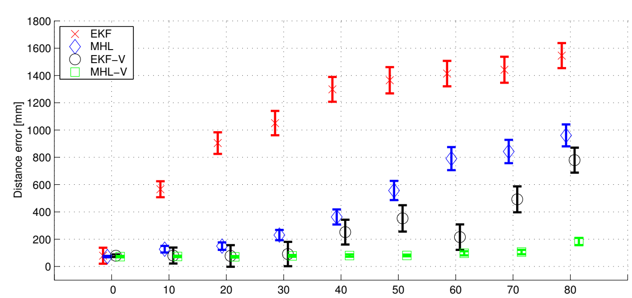

Fig. 2. The localisation accuracy for different levels of noise.

The accuracy of EKF-V and MHL-V are comparable to the results obtained by the best method tested by Gutmann and Fox (A-MCL).

Method	mean time [s]	95% interval [s]
EKF	1.85	0.39
MHL	0.74	0.19
EKF-V	12.2	3.9
MHL-V	0.67	0.13

TABLE I

THE TIME NEEDED TO RECOVER FROM KIDNAPPING.

3) Recovering from Kidnapping: To test the ability to recover from kidnapping one data file was created where, at 22 places, data from several seconds was removed (after which the odometry was re-calculated to cover this up). Looking at the odometry in the resulting log file it is not possible to detect the kidnappings.

Table I shows the results. We can see that the MHL methods recover quite quickly. This is probably due to the observation-driven addition of hypotheses and to the probabilistic formulation of hypothesis weight, which enables the method to quickly "switch" to a new hypothesis. What is interesting to see is that EKF recovers much faster than EKF-V which is due to the fact that the validation gate will filter out all measurements until the Kalman filter has accumulated a significant pose uncertainty from the robot movements. So where the validation gate improves robustness to noise it hampers recovery. The difference in recover time for the EKF compared to the one reported in [10] could be a result of using different motion models.

The times for MHL to recover are about 2.5 times faster than the best times reported by Gutmann and Fox (ML-EKF and A-MCL).

4) Computational resources: One of the major advantages of the EKF is that it is computationally efficient. The claim we made in the beginning of the paper was that MHL would inherit this property. To test this thesis we have measured the computation time needed by the prediction step, the update step and the total computation time for processing the entire data set just like in [10]. The timing test was carried out on a Pentium-III 1GHz laptop computer. The computation time for the EKF was 3.8μs and 8.5μs for the time per iteration to do prediction and update respectively. The total computation time for the whole log file was 0.7 s. The times are similar to the values reported in [10] which is an indication that the rest of the values can be compared with those results as well.

As most of the time is spent in the update step in all algorithms we take a closer look at the update time as a function of the amount of noise in the data in Figure 3. We can see that the MHL method required a bit more time than the EKF which comes from having to update more than one hypothesis. Since there is always only one hypothesis when running EKF, whether or not a gate is used, there is no difference between those two. There is however a significant difference between MHL and MHL-V. When not using the validation gate the track splitting produces more hypotheses that are not pruned away by merging, which increases the computational cost.

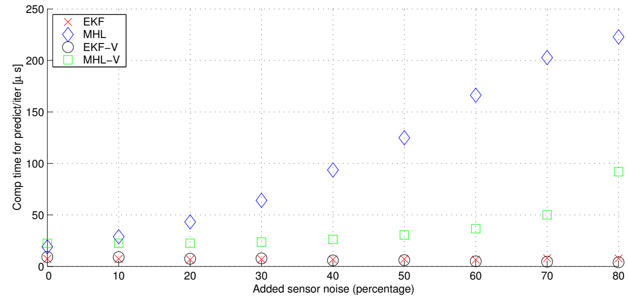

Fig. 3. The computation time for the update step for the different methods as a function of the noise level.

The plot in Figure 3 nicely illustrates the dynamic way MHL works: When everything is well-behaved, only very few hypotheses exist (in the no-noise experiment an average of 1.3) making the method computationally about as efficient as the EKF. When encountering noisy observations or some irregularities as e.g. kidnapping, more hypotheses will (at least for a while) exist in parallel demanding more resources but ensuring good accuracy and fast recovery.

The computation times for EKF, EKF-V and MHL-V (at low noise levels) are almost an order of magnitude lower than those reported in [10] for the ML-EKF and MCL methods. Depending on the implementation (e.g. fixed sample set size or not) the computation time for the methods reported in [10] will also depend on the noise level.

V. DISCUSSION AND CONCLUSION

Localisation is one of the most important problems in mobile robot navigation and benchmark tests are important for scientific and practical progress. In this paper we have taken up the ball from Gutmann *et al.* and tested the EKF and MHL with and without a validation gate on a standard set of data.

The test data were designed to reveal the methods' ability to cope with noise, sparse data and kidnapping. With respect to sparse data, all methods tested here and in [10] exhibit good, effectively identical performance.

On noisy data, the methods without a validation gate as expected perform rather poorly, MHL however being somewhat more robust than the EKF due to the track splitting technique. With the validation gate, both methods perform quite well, MHL being just as good as the best of the methods reported in [10] but requiring less computation time. Drawback of the validation gate is that it also efficiently prevents the EKF to recover once track is lost, e.g. due to kidnapping. This, however, is not a problem with the MHL method that will quickly generate new hypotheses to recover track.

In the kidnapping experiment, MHL and MHL-V recovered about 2.5 times faster than the best of the methods in [10]. We think this is mainly due to the observation-driven addition of hypotheses and to the probabilistic formulation of hypothesis weights.

The timing analysis proved that MHL is potentially as computationally efficient as the EKF, however dynamically demanding more resources when needed due to e.g. noisy observations.

On the basis of these results we think it is safe to conclude that MHL is a viable, computationally efficient alternative to the grid- and sample-based relocalisation methods. The cost for this is that MHL relies more on models of the process than the other methods which may tend to make the algorithms somewhat more elaborate.

Although we have in [13] shown that MHL performs well in normal indoor scenes, the less model-based approaches may have some advantages in environments which are harder to model and where ambiguous landmarks etc. are common. To examine this, however, is the subject of future work.

Acknowledgements

We would like to thank Steffen Gutmann and Dieter Fox for providing the data, the evaluation tools and for constructive comments. Thanks also to Henrik Christensen for valuable comments.

Patric's work is sponsored by the Swedish Foundation for Strategic Research through the Centre for Autonomous Systems, support is also given from the Royal Swedish Academy of Sciences (KVA). Both sources of funding is gratefully acknowledged.

VI. References

[1] Y. Bar-Shalom and T. E. Fortmann. *Tracking and Data Association*. Academic Press, 1988.

[2] J. Borenstein, H.R. Everett, and Liqiang Feng. *Navigating Mobile Robots: System and Techniques*. A K Peters, Ltd., 1996.

[3] Ingemar J. Cox. Blanche: Position estimation for an autonomous robot vehicle. In *Proc. of the International Workshop on Intelligent Robots and Systems*, pages 432–439, Tsukuba, Japan, September 4-6, 1989. IEEE/RSJ.

[4] Z. Crisman, E. Curre, C. Kwok, N. Ratliff, L. Tsybert, and D. Fox. Team description: Uw huskies-02. In G.Kaminka, P. Lima, and R. Rojas, editors, *RoboCup-2002: Robot Soccer World Cup VI*. Springer Verlag, 2003.

[5] J.L. Crowley. Navigation for an intelligent mobile robot. *IEEE Journal of Robotics and Automation*, 1(1):31–41, March 1985.

[6] F. Dellaert, D. Fox, W. Burgard, and S. Thrun. Monte carlo localization for mobile robots. In *Proc. of the IEEE International Conference on Robotics and Automation (ICRA'99)*, volume 2, pages 1322 – 1328, Detroit, Michigan, USA, May 10-15, 1999.

[7] Dieter Fox, Sebastian Thrun, Wolfram Burgard, and Frank Dellaert. *Sequential Monte Carlo Methods in Practice*, chapter Particle Filters for Mobile Robot Localization. Springer Verlag, New York, 2000.

[8] J.-S. Gutmann. Markov-kalman localization for mobile robots. In *Proc. of the International Conference on Pattern Recognition (ICPR'02)*, Quebec, Canada, August 2002.

[9] J.-S. Gutmann, W. Burgard, D. Fox, and K. Konolige. An experimental comparison of localization methods. In *Proc. of the 1998 IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS'98)*, 1998.

[10] Jens-Steffen Gutmann and Dieter Fox. An experimental comparison of localization methods continued. In *Proc. of the 2002 IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS'02)*, pages 454–459, Lausanne, Switzerland, October 2002.

[11] R. Hinkel and T. Knieriemen. Environment perception with a laser radar in a fast moving robot. In *Symposium on Robot Control*, pages 68.1–68.7, October 1988.

[12] P. Jensfelt. *Approaches to Mobile Robot Localization in Indoor Environments*. PhD thesis, Signal, Sensors and Systems (S3), Royal Institute of Technology, SE-100 44 Stockholm, Sweden, 2001. http://www.cas.kth.se/~ patric/publications/phd.html.

[13] P. Jensfelt and S. Kristensen. Active global localisation for a mobile robot using multiple hypothesis tracking. *IEEE Trans. on Robotics and Automation*, 17(5):748–760, October 2001.

[14] U. Larsson, J. Forsberg, and Å. Wernersson. On robot navigation using identical landmarks: Integrating measurements from a time–of–flight laser. In *Proceedings of the 1994 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems*, pages 17–26, Las Vegas, Nevada, October 1994. IEEE.

[15] Scott Lenser and Manuela Veloso. Sensor resetting localization for poorly modelled mobile robots. In *Proc. of the IEEE International Conference on Robotics and Automation (ICRA'00)*, volume 2, pages 1225–1232, April 2000.

[16] J.J. Leonard and H.F. Durrant-Whyte. Mobile robot localization by tracking geometric beacons. *IEEE Transactions on Robotics and Automation*, 7(3):376–382, 1991.

[17] Reid Simmons and Sven Koenig. Probabilistic navigation in partially observable environments. In *Proc. of the International Joint Conference on Artificial Intelligence (IJCAI'95)*, pages 1080–1087, 1995.

[18] Sebastian Thrun, Dieter Fox, and Wolfram Burgard. Monte carlo localization with mixture proposal distribution. In *Proc. of the National Conference on Artificial Intelligence (AAAI)*, Austin, Texas, USA, 2000.