

Gaussian Mixture Models for Probabilistic Localization

Patrick Pfaff Christian Plagemann Wolfram Burgard

Abstract—One of the key tasks during the realization of probabilistic approaches to localization is the design of a proper sensor model, that calculates the likelihood of a measurement given the current pose of the vehicle and the map of the environment. In the past, range sensors have become popular for mobile robot localization since they directly measure distance. However, in situations in which the robot operates close to edges of obstacles or in highly cluttered environments, small changes in the pose of the robot can lead to large variations in the acquired range scans. If the sensor model used does not appropriately characterize the resulting fluctuations, the performance of probabilistic approaches may substantially degrade. A common solution is to artificially smooth the likelihood function or to only integrate a small fraction of the measurements. In this paper we present a more fundamental and robust approach which uses mixtures of Gaussians to model the likelihood function for single range measurements. In practical experiments we compare our approach to previous methods and demonstrate that it yields a substantially increase in robustness.

I. INTRODUCTION

The ability to robustly localize a mobile robot given a map of its environment belongs to the fundamental problems in mobile robotics. One of the key challenges in this context is to design a likelihood function or observation model $p(z | x, m)$ which specifies how to compute the likelihood of an observation or measurement z given the robot is at pose x in a given map m . For probabilistic approaches like Monte Carlo localization (MCL) the proper design of the likelihood function is essential. For example, too optimistically specified sensor models might make the vehicle overly confident in its position, might cause a depletion of the correct particles, and finally might lead to a divergence of the filter. On the other hand, too conservative models might lead to a high uncertainty or even prevent the robot from localizing itself as the sensor information cannot compensate for the uncertainty introduced by the motion of the vehicle.

In the past, sophisticated sensor models have been developed for probabilistic approaches to robot localization. Some of them take into account various aspects such as objects not contained in the map or cross-talk. Whereas such approaches have been proven to be extremely robust in real-world applications, they do not appropriately take into account potential effects not stemming from the measurement process itself but caused by the fact that the map is only an approximation of the real world or that not all objects can be appropriately modeled, for example, due to discretization errors. Additionally, such sensor models

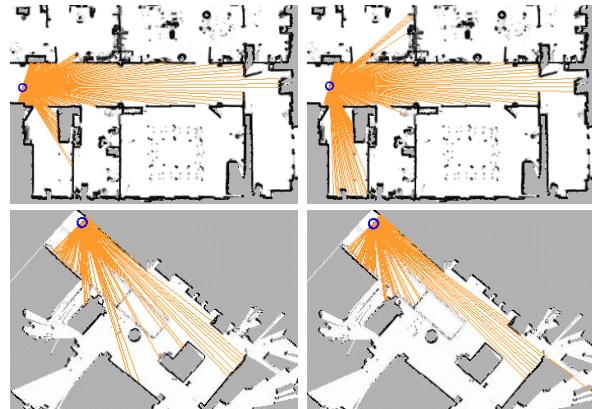


Fig. 1. In mobile robot localization, small variations in the robot pose can cause large changes of the range measurements. This leads to multi-modal distributions of beam-lengths in the local neighborhood of a pose hypothesis.

are sensitive to discontinuities in the map. For example, when the environment is cluttered, slight changes in the map might lead to huge differences in the length of the expected measurement at that particular location. This will lead to extremely small likelihoods of the measurement and then in turn might result in a divergence of the filter. To cope with such situations, likelihood fields have been proposed [17]. These models provide multi-modal likelihood functions to better deal with clutter in the environment but ignore the information along the individual beam of a range measurement. Therefore, likelihood fields are less effective in situations in which the robot has to perform a global localization.

In this paper, we present a novel sensor model that applies mixtures of Gaussians to better represent the likelihood function at each individual place. The key idea of our approach is to simulate a series of range measurements in the local neighborhood of each particular location and to approximate the likelihood function by a mixture of Gaussians using the Expectation Maximization (EM) algorithm. The advantage of this approach is that the resulting likelihood function is location-dependent and correctly captures the effects of clutter and discontinuities on the measurements. As a result, the localization process becomes more robust. In practical experiments carried out with data obtained with a real robot, we demonstrate that our new model substantially outperforms existing sensor models.

This paper is organized as follows. After discussing related work in the next section, we briefly describe Monte Carlo localization in Section III and the principle of likelihood models in Section IV. In Section V, we introduce our novel likelihood model based on mixtures of Gaussians and finally,

in Section VI, we present experimental results illustrating that our sensor model outperforms other popular likelihood models.

II. RELATED WORK

In the literature, various techniques for computing the likelihood function for probabilistic localization methods with proximity sensors have been introduced [2], [8], [17], [18]. These approaches either directly approximate the physical characteristics of the sensor or try to provide smooth likelihood models to increase the robustness of the localization process. In the past, it has been observed that the likelihood function can have a major influence on the performance of Monte Carlo Localization. Pitt and Shepard [14], for example, as well as Thrun *et al.* [19] observed that more particles are required if the likelihood function is peaked. In the limit, i.e., for a perfect sensor, the number of required particles becomes infinite. To deal with this problem, Lenser and Veloso [10] and Thrun *et al.* [19] introduced techniques to directly sample from the observation model and in this way ensure that there is a critical mass of samples at the important parts of the state space. Unfortunately, this approach depends on the ability to sample from observations, which can often only be done in an approximate, inaccurate way. Another way of dealing with the limitations of the sample-based representation is to dynamically adapt the number of particles, as done in KLD sampling [7]. However, for highly accurate sensors, even such an adaptive technique might require a huge number of samples in order to achieve a sufficiently high particle density during global localization. Alternatively, one can artificially inflate the measurement uncertainty to achieve a regularization of the likelihood function, e.g., see the *Scaling Series* approach presented by Petrovskaya *et al.* [11]. Also Kalman filters have limitations in highly non-linear systems and in the case of multi-modal likelihood functions. To overcome this problem several researchers used Gaussian mixture models. Duckett and Nehmzow [6], for example, introduced a multi-modal generalization of the Kalman filter based on mixtures of Gaussians. Recently, Upcroft *et al.* [20] introduced a fast re-parameterization of Gaussian mixture models to represent the probability distribution. Takamasa *et al.* [9] use Gaussian mixture models to fuse odometry and sonar and to reduce the localization error in the case of noisy sensors.

The focus of this paper is to model possible multimodalities in likelihood functions for laser range measurements using Gaussian mixture models. Our approach improves the robustness of probabilistic localization approaches like MCL especially in situations in which small changes of the robot's pose can have drastic effects on the range measurements.

III. MONTE CARLO LOCALIZATION USING RANGE SENSORS

Throughout this paper, we consider the problem of estimating the pose $\mathbf{x} = (x, y, \theta)$ of a robot relative to a given map \mathbf{m} using a particle filter. The key idea of this approach

is to maintain a probability density $p(\mathbf{x}_t | \mathbf{z}_{1:t}, \mathbf{u}_{0:t-1})$ of the location \mathbf{x}_t of the robot at time t given all observations $\mathbf{z}_{1:t}$ up to time t and all control inputs $\mathbf{u}_{0:t-1}$ up to time $t - 1$. This probability is calculated recursively as

$$p(\mathbf{x}_t | \mathbf{z}_{1:t}, \mathbf{u}_{0:t-1}) = \alpha \cdot p(\mathbf{z}_t | \mathbf{x}_t) \int p(\mathbf{x}_t | \mathbf{u}_{t-1}, \mathbf{x}_{t-1}) \cdot p(\mathbf{x}_{t-1}) d\mathbf{x}_{t-1} \quad (1)$$

Here, α is a normalizing constant ensuring that $p(\mathbf{x}_t | \mathbf{z}_{1:t}, \mathbf{u}_{0:t-1})$ sums up to one over all \mathbf{x}_t . The terms to be described in Eqn. (1) are the *prediction model* $p(\mathbf{x}_t | \mathbf{u}_{t-1}, \mathbf{x}_{t-1})$ and the *sensor model* $p(\mathbf{z}_t | \mathbf{x}_t)$ respectively.

For the implementation of the described filtering scheme, we use a sample-based approach which is commonly known as *Monte Carlo localization (MCL)* [3]. Monte Carlo localization is a variant of particle filtering [5] where each particle corresponds to a possible robot pose and has an assigned weight w_i . The *belief update* from Eqn. (1) is performed by the following two alternating steps:

- 1) In the **prediction step**, we draw for each particle with weight w_i a new particle according to w_i and to the prediction model $p(\mathbf{x}_t | \mathbf{u}_{t-1}, \mathbf{x}_{t-1})$.
- 2) In the **correction step**, a new observation \mathbf{z}_t is integrated. This is done by assigning a new weight w_i to each particle according to the sensor model $p(\mathbf{z}_t | \mathbf{x}_t)$.

The likelihood model $p(\mathbf{z} | \mathbf{x})$ plays a crucial role in the correction step of the particle filter and its proper design is essential for the robustness of the filter. In the following section we will describe typical likelihood models for range sensors. Afterwards, we will present our model that uses mixtures of Gaussians to represent multi-modalities in the likelihood function.

IV. LIKELIHOOD MODELS

A laser scan \mathbf{z}_t is a vector of beams $\mathbf{z}_t = (z_t^1, \dots, z_t^N)^T$, which have fixed orientations relative to the sensor. *Beam-based* sensor models (see Fox *et al.* [8] for a typical example) consider each value z_t^i of the measurement vector \mathbf{z} as a separate range measurement and represent its one-dimensional distribution by a parametric function depending on the expected distance in the respective beam direction. Such models are closely linked to the geometry and the physics involved in the measurement process. They are sometimes also called *ray cast* models because they rely on ray casting operations within the map of the environment, e.g., an occupancy grid map, to calculate the expected beam lengths. Another popular measurement model for range finder sensors are the so-called *likelihood fields* (AKA *end point model*) [17]. This correlation-based method measures the correlation between the measurement and the map. Here, the likelihood of a range measurement is a function of the distance of the respective end point of the beam to the closest obstacle in the environment. This model lacks physical explanation as it can basically “see through walls”, but in the case of position tracking it is efficient and works well in practice. The reader may notice that likelihood fields are

less effective in situations in which the robot has to perform global localization. In all above mentioned approaches, the individual beams are treated independently. Due to this, the likelihood $p(z_t | x_t, m)$ of the scan z_t given the position x_t and the map m can be calculated by

$$p(z_t | x_t, m) = \prod_{i=1}^N p(z_t^i | x_t, m). \quad (2)$$

Given that laser range finders typically provide between 181 and 540 measurements with a resolution from 0.25 to 1.0 degrees, the independence assumption leads to highly peaked likelihoods. In practice, this problem is dealt with by sub-sampling of measurements [18], by introducing minimal likelihoods for beams, or by other means of regularization of the resulting likelihoods, see e.g. Arulampalam *et al.* [1]. In our previous work [12], we addressed this problem by adapting the “peakedness” of the beam model using learned heuristics. In other previous work [13], [15], we introduced *scan-based* likelihood models. These models allow to directly calculate the likelihood of entire scans rather than individual beams only. The common idea of these two approaches is that they are location-dependent and explicitly take the approximation error coming from the sample-based representation into account. More precisely, they estimate $p(z | x)$ based on the local environment $\mathcal{U}(x)$ of a pose x ,

$$p(z | x) = \int_{\tilde{x} \in \mathcal{U}(x)} p(z | \tilde{x}) p(\tilde{x}) d\tilde{x}. \quad (3)$$

This is based on the observation that laser range finders are extremely accurate sensors with a low level of measurement noise. Thus, if one learns $p(z | x)$ directly for exact sensor poses x , e.g., with a mobile robot that is not moved during training, one gets an extremely peaked model with $p(z | x + \delta) \ll p(z | x)$ already for small pose perturbations δ . This peakedness, in turn, leads to problems when only a finite number of pose hypotheses can be evaluated, as it is the case, for example, with particle filters where the number of particles is limited. The model described in Eqn. 3, however, is able to explicitly consider the sampling density by adjusting the spatial extent of the local neighborhoods $\mathcal{U}(x)$ accordingly. The hard task is indeed to estimate and represent the distributions of range scans $p(z | x)$ from a given number of training scans from $\mathcal{U}(x)$, which are typically simulated using the map of the environment m . Our previous approaches [13], [15] deal with the high dimensionality of this distribution by modeling it as an N -dimensional Gaussian. While this leads to an increased performance compared to approaches that only consider the current pose of the particle, it does not capture the multi-modality of the likelihood function. In the following section, we describe how to learn a Gaussian mixture model for the distribution obtained by Eqn. 3 to improve the robustness of the localization process.

V. PLACE-DEPENDENT GAUSSIAN MIXTURE MODELS

Figure 1 illustrates the drastic effects that small changes of the robot’s pose can have on the measured range scans.

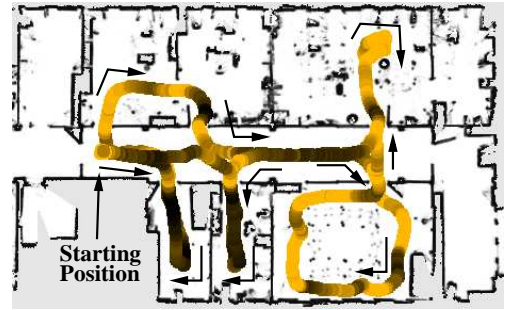


Fig. 2. Traditional sensor models which ignore the multi-modality in the distribution of beam lengths become inaccurate in the proximity of doors, corners, and clutter. In this diagram, dark parts of the trajectory mark locations, where the true pose receives a high likelihood relative to neighboring ones, while orange/brightly colored parts highlight the less accurately modeled locations.

The distribution of measured distances that arises when the robot pose is varied locally as described in the previous section is only unimodal in a perfectly convex world. In general, however, there can be large jumps in perceived range measurements when the sensor pose is changed only slightly. Typically, such multi-modalities arise in the proximity of doorways, corners, and cluttered areas of the environment.

In contrast to former approaches [8], [12], [13], [15] which modeled the likelihood functions as unimodal distributions for single beams or entire scans we now consider to model each beam independently as a mixture of K Gaussian distributions [16]. In such a mixture model, the likelihood of the i -th beam of z_t becomes

$$p(z_t^i | x_t, m) = \sum_{j=1}^K p(z_t^i | j) P(j), \quad (4)$$

where

$$p(z_t^i | j) = \frac{1}{\sqrt{2\pi}\sigma_j^i} \cdot \exp\left(-\frac{(z_t^i - \mu_j^i)^2}{2\sigma_j^i{}^2}\right). \quad (5)$$

Here, the individual mixture components are indexed by j and their relative mixing weights are denoted – with a slight abuse of notation for better readability – as $P(j) =: \alpha_j^i$. To determine these weights α_j^i with $\sum_{j=1}^K \alpha_j^i = 1$, $0 \leq \alpha_j^i \leq 1$, as well as the parameters μ_j^i and σ_j^i of the individual Gaussians, we cluster the simulated ranges \mathcal{D}^i using the expectation-maximization (EM) algorithm [4]. Concretely, for each pose hypothesis x_t , we simulate L complete range scans $\mathcal{D} = \{d_1, \dots, d_L\}$ at locations drawn uniformly from $\mathcal{U}(x_t)$ using the given map m of the environment. The simulation of the laser range scans $\mathcal{D} = \{d_1, \dots, d_L\}$ takes into account the geometry and the physics involved in the measurement process. It relies on ray casting operations within an occupancy grid map to calculate the expected beam lengths. The set of ranges simulated in direction of the i -th laser beam is denoted as $\mathcal{D}^i = \{d_1^i, \dots, d_L^i\}$. The EM algorithm iteratively assigns these distances to the mixture components and optimizes their parameters in the following manner. Consider that θ' denotes the current estimate of parameters μ_j^i , σ_j^i , and α_j^i . In the E-Step, we calculate the

expected value of the complete log-likelihood

$$Q(\theta, \theta') = E[\log\{p(\mathcal{D}^i, Y^i | \theta)\} | \mathcal{D}^i, \theta'] \quad (6)$$

$$= \int_{y^i} \log\{p(\mathcal{D}^i, y^i | \theta)\} p(y^i | \mathcal{D}^i, \theta') dy^i, \quad (7)$$

where Y^i denotes data associations of the simulated data points \mathcal{D}^i to one of the Gaussian mixture components. Visually speaking, we estimate the assignment likelihoods of the individual samples to the clusters while keeping the other model parameters fixed. Then, in the M-Step, we fix the data associations and optimize the expected value of the complete log-likelihood

$$\theta'' = \operatorname{argmax}_{\theta} Q(\theta, \theta') \quad (8)$$

by updating the cluster parameters according to

$$\alpha_j^i = \frac{1}{L} \sum_{l=1}^L P(j | d_l^i, \theta'), \quad (9)$$

$$\mu_j^i = \frac{\sum_{l=1}^L P(j | d_l^i, \theta') d_l^i}{\sum_{l=1}^L P(j | d_l^i, \theta')}, \quad (10)$$

$$(\sigma_j^i)^2 = \frac{\sum_{l=1}^L P(j | d_l^i, \theta') (d_l^i - \mu_j^i)^2}{\sum_{l=1}^L P(j | d_l^i, \theta')}. \quad (11)$$

We now set $\theta' \leftarrow \theta''$ and iterate this procedure until the amount of improvement per iteration falls below a specified threshold.

VI. EXPERIMENTS

The approach described above has been implemented and tested on data obtained with a mobile robot and by simulation. To evaluate our approach we performed several experiments. We first show that the pose uncertainty of the robot can result in serious problems during a localization process, especially when the multi-modality of the beams is not considered. Then we analyze our Gaussian mixture model in a global localization task in which in multi-modal situations frequently occur and compare it to alternative models that do not take into account the multi-modality. In particular, we compared the performance of the following sensor models:

- GM*: Our place-dependent beam-based Gaussian mixture sensor model as detailed in Section V.
- IB*: The standard beam-based sensor model that assumes independent beams with an additive white noise component.
- EP*: The end-point sensor model [17] that calculates the likelihood of a range measurement as a function of the distance of the end point of the respective beam to the closest obstacle in the environment.
- EC*: The scan-based place-dependent model with learned covariance matrix as detailed in our previous work [13].
- DC*: The same model as *EC* with cross-correlation components ignored, which means that only the diagonal entries of the covariance matrix are learned.

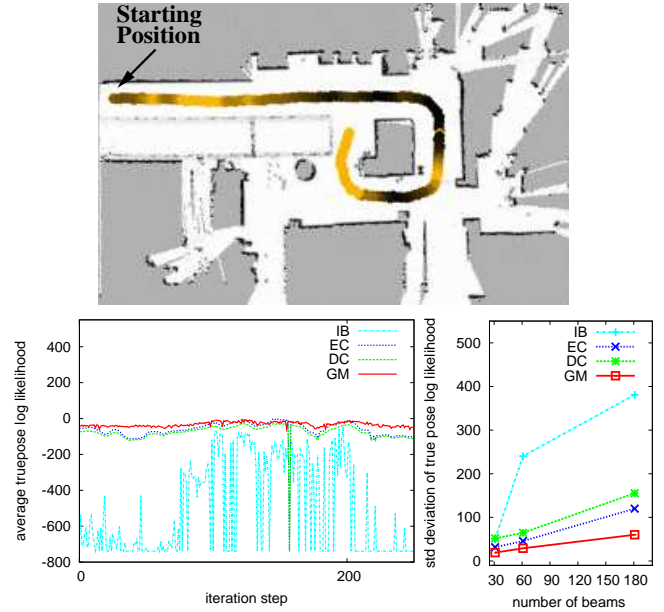


Fig. 3. Experiment using 61 beams of laser range data collected in the Heinz-Nixdorf-Forum in Paderborn. The experiment shows that the standard ray cast model (*IB*) performs sub-optimally in regions close to corners and that it produces highly fluctuating likelihood estimates. Our Gaussian mixture model (*GM*) outperforms the other sensor models and produces much less variance in the evaluated likelihoods.

A. Likelihood Evaluation

In the first set of experiments we evaluated the likelihood of the true position of the robot in different data sets. We therefore compared our Gaussian mixture model (*GM*) to other likelihood models which are also based on ray casting operations (*IB*, *EC*, and *DC*). This set of experiments is designed to investigate the case that the robot is not able to localize itself at different locations with the same robustness. Figures 2 and 3 show two experiments using maps built from a real data. Figure 5 shows an artificial data set with strong discontinuities in the right part of the map. During these experiments, we simulated laser range scans with an opening angle of 180° on a simulated robot trajectory. Then we calculated for different sensor models (*GM*, *IB*, *EC*, and *DC*) the likelihood of the simulated range scan given the true position of the robot. Figure 2, 3, and 5 depict the trajectories of the robot. The likelihoods of the scans at the true robot poses are represented by the different colors (grey scale values) between orange (light grey) and black. While orange (light grey) marks regions where the scans at the true robot location are assigned low likelihoods, high likelihood areas are printed in black. The figures show that whenever the robot traverses regions close to obstacles, doorways, or clutter the likelihood of the true position decreases. In the case of global localization using a particle filter this leads to serious problems because the particles at these positions have a high risk of being depleted. Figure 4 shows the mean likelihoods for 31, 61, and 181 laser beams and different sensor models. We evaluated the likelihood at 847 robot poses in our office environment depicted in Figure 2 and averaged over 50 runs. As can be seen from the figures, our Gaussian mixture model (*GM*) yields much less variance in

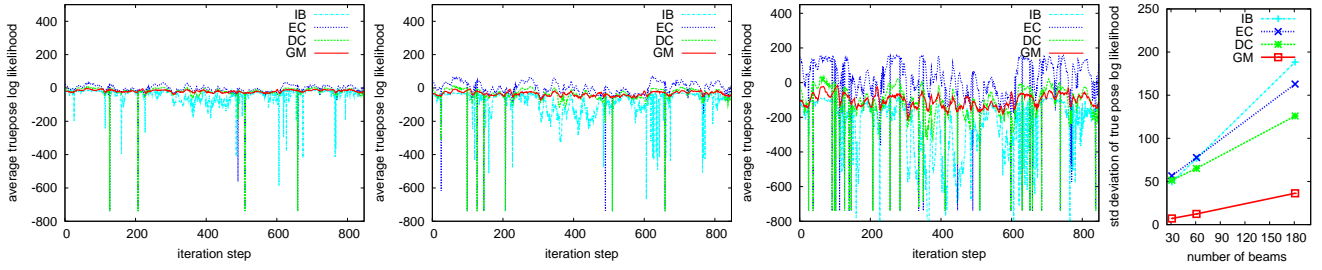


Fig. 4. Evaluated likelihood for 31, 61, and 181 laser beams (from left to right) and different sensor models at 847 robot poses in our office environment depicted in Figure 2. The figures show that the likelihood of our Gaussian mixture model (*GM*) yields much less variance in the estimated likelihood of the true pose of the robot than the other sensor models. The right most diagram depicts the standard deviations of the different sensor models for 31, 61, and 181 laser beams.

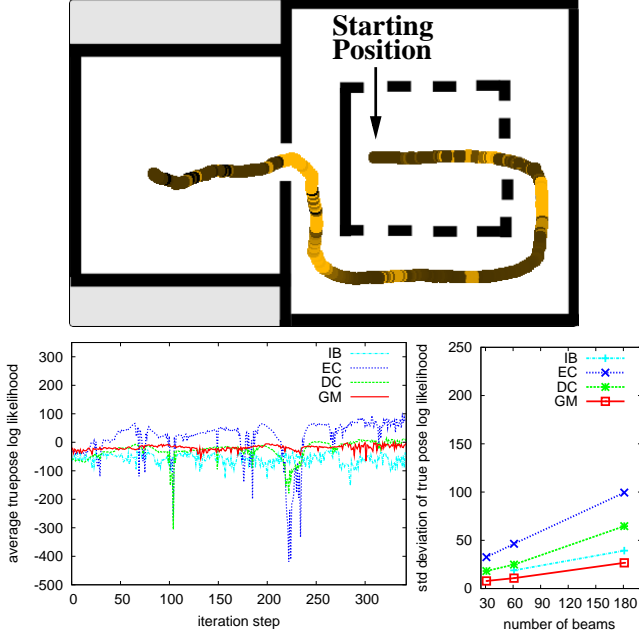


Fig. 5. Experiment carried out to evaluate the likelihood of 348 simulated robot poses using an artificial data set including strong discontinuities like corners and doorways. The left diagram depicts the average likelihoods for 61 laser beams and different sensor models. The figures show that the likelihood of our Gaussian mixture model (*GM*) yields much less variance in the likelihood of the true pose of the robot than the other sensor models. The right diagram shows the standard deviations of the different sensor models for 31, 61, and 181 laser beams.

the estimated likelihood of the true pose compared to the other sensor models. The rightmost diagram shows the standard deviation of the different sensor models for 31, 61, and 181 laser beams. Figure 3 shows the same experiment using data collected in the Heinz-Nixdorf-Forum in Paderborn for 61 laser beams. The experiment shows that the standard ray cast model (*IB*) performs sub-optimally in regions close to corners and that it produces highly fluctuating likelihood estimates. Our Gaussian mixture model (*GM*) outperforms the other sensor models and has much lower variance in the estimated likelihoods. In a final experiment documented in Figure 5 we observed a similar behavior of the different sensor models in an artificial data set which produces strong discontinuities because of corners and doorways.

B. Localization

The second set of experiments is designed to illustrate that our new sensor model (*GM*) which takes the multi-modality

of measurements into account achieves more robust and accurate localization than the other sensor models. The left part of Figure 6 shows the six positions in a real environment where we obtained the highest probability that the global localization fails. These probabilities have been determined by random restarts of the localization procedure during 50 complete runs on the data set. The marked positions directly correspond to orange (light grey) marked regions in Figure 2 where the likelihoods of the true poses are extremely low due to the multi-modality of the measurements. To evaluate the properties of the different sensor models we performed 20 global localization runs at each position and compared the average success rates. In these experiments, we assumed that the localization was achieved when the mean of the particles differed by at most 50 cm from the true location of the robot. The central diagram of Figure 6 shows the number of successful localizations after ten integrations of 61 measurements of each scan for different models. The experiments show that our Gaussian mixture model (*GM*) allows us to more robustly localize the robot in situations in which the other models frequently fail. It also illustrates that the endpoint model (*EP*) which shows good performance for position tracking in cluttered environments is not able to solve the global localization task in the marked regions of our environment. Additionally, we analyzed the robustness of the different sensor models with respect to the computation time (see the right diagram of Figure 6). The x-axis of this diagram represents the product of the number of particles used and the average computation time per particle for the different sensor models. As can be seen from this diagram, our Gaussian mixture model (*GM*) outperforms the other models also with respect to the computational complexity. Whereas the time per iteration is higher compared to the other approaches, it requires considerably less particles for a successful localization run and thus achieves a higher robustness relative to the required computational resources. We also carried out experiments, in which we analyzed the accuracy of our model (*GM*) when the system is tracking the pose of the vehicle. We compared our sensor model to various other models and evaluated the average localization error of the individual particles. The right diagram of Figure 6 depicts the average localization error for a position tracking experiment with 61 laser beams. It can be seen that the two beam-based ray cast sensor models (*IB*) and (*DC*) diverge while our beam

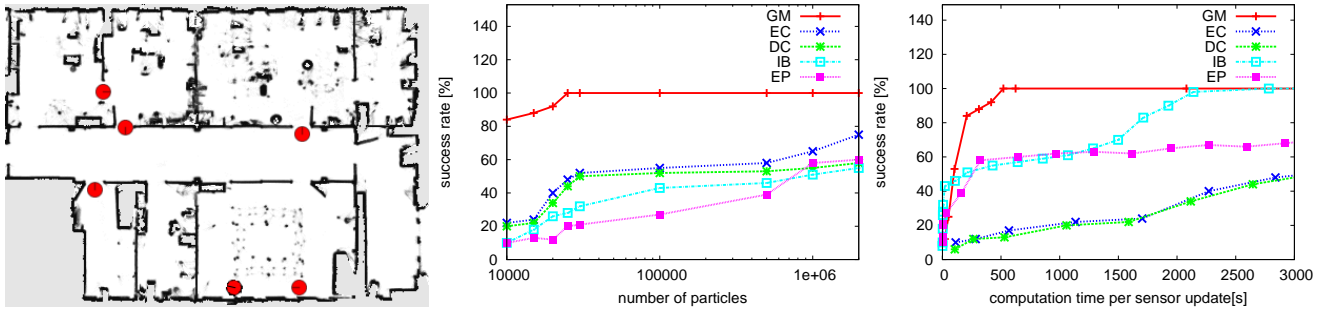


Fig. 6. The six positions with the highest probability that the global localization in the office environment fails (left). The positions marked in this image directly correspond to the orange (light grey) positions in Figure 2. At these positions the likelihoods of the true poses are extremely low due to the multi-modality of the measurements. The central diagram shows the number of successful localizations after ten integrations of 61 measurements at these locations. The right diagram also depicts the success rate but now plotted over the product of the corresponding numbers of particles and average computation time per particle for the different sensor models in these situations.

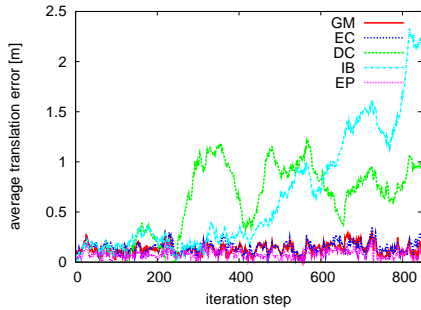


Fig. 7. Average localization error in meters for a position tracking experiment with 61 laser beams.

based Gaussian mixture model (GM) performs as well as the endpoint model (EP) and the scan-based place-dependent model (EC).

VII. CONCLUSIONS

In this paper, we presented a novel beam-based sensor model for probabilistic localization techniques that explicitly takes multi-modalities in the distribution of beam lengths into account. In contrast to other location-independent models our approach adapts the likelihood evaluation according to the local environment of each evaluated pose hypothesis to achieve a natural and accurate form of regularization. By learning a Gaussian mixture model for the resulting distribution of possible range measurements using the EM algorithm, our approach is able to outperform the state-of-the-art approaches in terms of localization accuracy and robustness also relative to the required computational resources.

ACKNOWLEDGMENT

This work has partly been supported by the DFG within the Research Training Group 1103 and under contract number SFB/TR-8, by the EU under FP6-004250-CoSy, and by the German Ministry for Education and Research (BMBF) through the DESIRE project.

REFERENCES

[1] S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp. A tutorial on particle filters for on-line non-linear/non-gaussian bayesian tracking. In *IEEE Transactions on Signal Processing*, volume 50, pages 174–188, February 2002.

[2] H. Choset, K.M. Lynch, S. Hutchinson, G. Kantor, W. Burgard, Kavraki L.E., and S. Thrun. *Principles of Robot Motion Planning*. MIT-Press, 2005.

[3] Frank Dellaert, Dieter Fox, Wolfram Burgard, and Sebastian Thrun. Monte carlo localization for mobile robots. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, pages 99–141, 1998.

[4] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society*, 1977.

[5] A. Doucet, N. de Freitas, and N. Gordon, editors. *Sequential Monte-Carlo Methods in Practice*. Springer Verlag, 2001.

[6] T. Duckett and U. Nehmzow. Mobile robot self-localization using occupancy histograms and a mixture of Gaussians location hypotheses. *Robotics and Autonomous Systems*, 34(2-3):119–130, 2001.

[7] D. Fox. Adapting the sample size in particle filters through KLD-sampling. *Int. Journal of Robotics Research*, 22(12):985 – 1003, 2003.

[8] Dieter Fox, Wolfram Burgard, and Sebastian Thrun. Markov localization for mobile robots in dynamic environments. *Journal of Artificial Intelligence Research*, 11:391–427, 1999.

[9] T. Koshizen, P. Bartlett, and A. Zelinsky. Sensor fusion of odometry and sonar sensors by the Gaussian mixture Bayes’ technique in mobile robot position estimation. In *IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, 1999.

[10] S. Lenser and M. Veloso. Sensor resetting localization for poorly modelled mobile robots. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2000.

[11] A. Petrovskaya, O. Khatib, S. Thrun, and A.Y. Ng. Bayesian estimation for autonomous object manipulation based on tactile sensors. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2006.

[12] P. Pfaff, W. Burgard, and D. Fox. Robust monte-carlo localization using adaptive likelihood models. In H.I. Christensen, editor, *European Robotics Symposium 2006*, volume 22, pages 181–194. Springer-Verlag Berlin Heidelberg, Germany, 2006.

[13] P. Pfaff, C. Plagemann, and W. Burgard. Improved likelihood models for probabilistic localization based on range scans. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2007.

[14] M. K. Pitt and N. Shephard. Filtering via simulation: auxiliary particle filters. *Journal of the American Statistical Association*, 94(446), 1999.

[15] C. Plagemann, K. Kersting, P. Pfaff, and W. Burgard. Gaussian beam processes: A nonparametric bayesian measurement model for range finders. In *Robotics: Science and Systems (RSS)*, June 2007.

[16] R. A. Redner and H. F. Walker. Mixture densities, maximum likelihood, and the em algorithm. *SIAM Review*, 26(2):195–239, 1984.

[17] S. Thrun. An online mapping algorithm for teams of mobile robots. *Int. Journal of Robotics Research*, 20(5):335–363, 2001.

[18] S. Thrun, W. Burgard, and D. Fox. *Probabilistic Robotics*. MIT-Press, 2005.

[19] S. Thrun, D. Fox, W. Burgard, and F. Dellaert. Robust Monte Carlo localization for mobile robots. *Artificial Intelligence*, 128(1-2), 2001.

[20] B. Uprocroft, S. Kumar, M.F. Ridley, L. Ong, and H.F. Durrant-Whyte. Fast re-parameterisation of Gaussian mixture models for robotics applications. In *Australasian Conference on Robotics and Automation*, 2004.