

Efficiently Learning High-dimensional Observation Models for Monte-Carlo Localization using Gaussian Mixtures

Patrick Pfaff Cyrill Stachniss Christian Plagemann Wolfram Burgard

Abstract—Whereas probabilistic approaches are a powerful tool for mobile robot localization, they heavily rely on the proper definition of the so-called observation model which defines the likelihood of an observation given the position and orientation of the robot and the map of the environment. Most of the sensor models for range sensors proposed in the past either consider the individual beam measurements independently or apply uni-modal models to represent the likelihood function. In this paper, we present an approach that learns place-dependent sensor models for entire range scans using Gaussian mixture models. To deal with the high dimensionality of the measurement space, we utilize principle component analysis for dimensionality reduction. In practical experiments carried out with data obtained from a real robot, we demonstrate that our model substantially outperforms existing and popular sensor models.

I. INTRODUCTION

In the past, probabilistic localization techniques have been demonstrated to be a robust approach to mobile robot localization as they allow the vehicles to globally localize themselves, to efficiently keep track of their position, and to even recover from localization failures. One of the key challenges in context of probabilistic localization, however, lies in the design of the so-called observation model $p(z | x, m)$ which is a likelihood function that specifies how to compute the likelihood of an observation z given the robot is at pose x in a given map m . For probabilistic approaches the proper design of the likelihood function is essential. Too optimistically specified sensor models might make the vehicle overly confident in its position. In the context of Monte-Carlo-Localization [4], this can lead to a depletion of particles and finally might cause a divergence of the filter. Too conservative models, in contrast, might result in a high pose 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 sensor cross-talk. Whereas such approaches have been proven to be robust even in real-world situations, they do not appropriately take into account potential effects not stemming from the measurement process itself. This, for example, regards the fact that the map typically is only an approximation of the real world. Additionally, such sensor models are sensitive to discontinuities in the map. For




Fig. 1. In mobile robot localization, small variations in the robot pose can cause large variations of the range measurements. This leads to multi-modal distributions of beam-lengths even in small areas around a potential pose.

example, when the environment is cluttered, slight changes in the pose of the robot in the map might lead to huge differences in the length of the expected measurement at that particular location. This fact is illustrated in Figure 1. The figure shows two different scans obtained with a laser range scanner with a robot that passes a doorway. Whereas the two positions, at which the scans were taken, are close to each other, the obtained scans differ largely. Accordingly, even small errors in the pose estimate can lead to an extremely small likelihood of the measurement. This in turn can lead to a divergence of the filter.

One sensor model that has been especially designed to cope with such situations are the so-called likelihood fields [17]. This “beam-end-point-model” provides smooth and multi-modal likelihood functions to better deal with clutter in the environment but ignores 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 global localization. Most observation models furthermore assume the independency of the individual beams of a laser range scan. However, the more beams a scan has, the more this assumption is violated, which then might lead to overly confident estimates. Recently, some techniques have been developed that explicitly consider the dependencies of individual beams [14], [16]. However, these techniques have the drawback that they assume an uni-modal distribution.

In this paper, we propose a novel probabilistic observation model for proximity sensors such as laser range finders. Our model has two advantages over most previous approaches. First, it explicitly considers the dependencies between the individual beams of a range scan, and second, it accounts for

The authors are with the University of Freiburg, Department of Computer Science, Georges-Koehler-Allee 79, 79110 Freiburg, Germany. {pfaff, stachnis, plagem, burgard} @ informatik.uni-freiburg.de

the multi-modal nature of the observation function. It does so while still considering that the observation is obtained from a time-of-flight proximity sensor (such as laser range finders or sonars). This is achieved by considering place-dependent measurement models and utilizing a Gaussian mixture model together with a dimensionality reduction technique. 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 in Section III Monte Carlo localization and the principle of likelihood models. In Section IV, we introduce our novel likelihood model based on high-dimensional mixtures of Gaussians and finally, in Section V, 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 [3], [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. Thrun *et al.* [19], for example, 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, Lense 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. An alternative strategy to deal with this problem is to 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.* [12].

The classical Kalman filter has limitations since it requires Gaussian observation likelihoods and thus cannot handle multi-modalities or ambiguous situations. To overcome this problem, several researchers used Gaussian mixture models. Duckett and Nehmzow [7], 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.

Recently, Limketkai *et al.* [11] proposed an approach for learning the motion and sensor model for MCL using conditional random fields. This allows for considering the dependencies between the individual beams of a laser range

scan. The approach is furthermore reported to provide better results than generatively learned models but it requires ground truth location information of a robot to carry out discriminative learning.

The focus of this paper is to model possible multi-modalities 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

A. Pose Estimation using a Particle Filter

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)* [4]. Monte Carlo localization is a variant of particle filtering [6] 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 $p(\mathbf{z}_t | \mathbf{x}_t)$.

B. Likelihood Models for Range Sensors

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. Due to that, in this paragraph we will describe typical likelihood models for range sensors and we shortly will introduce the likelihood models of our previous work [14], [15]. Afterwards, we will present our new high dimensional Gaussian mixture model that is able to represent multi-modalities in the likelihood function as well as dependencies between the individual laser beams.

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], which are a correlation-based method that 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 endpoint 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 and the likelihood $p(\mathbf{z}_t | \mathbf{x}_t, m)$ of the entire scan \mathbf{z}_t is calculated as the product of the individual beam likelihoods $p(z_t^i | \mathbf{x}_t, m)$.

Given the high resolution of typical laser range finders (.25 to 1 degrees), the independence assumption leads to highly peaked likelihoods. In practice, this problem is dealt with by sub-sampling the measurements [18], by introducing minimal likelihoods for beams, or by other means of regularization of the resulting likelihoods (see Arulampalam *et al.* [1]). One way to overcome the peakedness is to consider that the likelihood models are location-dependent as well as that the location of the robot is modeled by set finite set of pose hypotheses (particles). Therefore, we estimate $p(\mathbf{z} | \mathbf{x})$ based on the local environment $\mathcal{U}(\mathbf{x})$ around a pose hypothesis \mathbf{x} by

$$p(\mathbf{z} | \mathbf{x}) = \int_{\tilde{\mathbf{x}} \in \mathcal{U}(\mathbf{x})} p(\mathbf{z} | \tilde{\mathbf{x}}) p(\tilde{\mathbf{x}}) d\tilde{\mathbf{x}}. \quad (2)$$

Here, $\mathcal{U}(\mathbf{x})$ is a particle-dependent, circular area. The diameter of that area is given by the distance to the closest neighboring particle which can be efficiently determined using a *kd*-tree. This model is able to represent the fact that pose hypotheses given by the samples are typically less accurate than the measurements obtained by a (SICK) laser range finder.

Thus, if one learns $p(\mathbf{z} | \mathbf{x})$ directly for exact sensor poses \mathbf{x} , e.g., with a mobile robot that is not moved during training, one gets an extremely peaked model with $p(\mathbf{z} | \mathbf{x} + \delta) \ll p(\mathbf{z} | \mathbf{x})$ 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




Fig. 2. Obtained dimensionality reduction over a full robot trajectory in a real world experiment using Principal component analysis (PCA).

of particles is limited. The model described in Eqn. 2, however, is able to explicitly consider the sampling density by adjusting the spatial extent of the local neighborhoods $\mathcal{U}(\mathbf{x})$ accordingly. The hard task is indeed to estimate and represent the distributions of range scans $p(\mathbf{z} | \mathbf{x})$ from a given number of training scans from $\mathcal{U}(\mathbf{x})$, which are typically simulated using the map of the environment m .

In our previous approaches, we modeled the observation likelihood as either unimodal distributions for single beams [8], [13] or for entire scans [14], [16]. Alternatively, our recently proposed method [15] is able to consider the multi-modality of the observation model but is unable to represent the dependency between the individual scans. In contrast to this, the novel method presented in this paper combines the advantages of these previous models: It considers the dependency between the beams of a range scan and it correctly handles the multi-modal nature of the likelihood function and at the same time can be computed efficiently. As we will demonstrate in the experiments, this more sophisticated model significantly improves the ability of a mobile robot to localize itself.

In the following section, we describe how to efficiently learn a high-dimensional Gaussian mixture model for the distribution obtained by Eqn. 2 to improve the robustness of the localization process.

IV. LEARNING HIGH DIMENSIONAL GAUSSIAN MIXTURE OBSERVATION MODELS

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

One way to model the different modes in the distribution of the expected range observation for each laser beam is to explicitly consider the multi-model nature [15]. Whereas this technique yields appropriate, multi-model distributions for individual beams, it is unable to model the dependency between in individual beams.

The straightforward extension that considers also the dependency between the individual beams is to learn a Gaussian mixture model (GMM) based on full laser scans and not individual beams. Most clustering techniques based on the Gaussian mixture model, however, show a suboptimal performance if the size of the training dataset is too small

compared to the number its dimensionality (the parameters to estimate). Typically, this leads to serious over-fitting. It is therefore necessary to find a good balance between the number of parameters to estimate and the generality of the model.

One way to overcome this problem is to apply k-means clustering since it does not estimate the covariance matrix and thus less parameters need to be estimated. However, the dependencies between beams are then neglected when estimating the clusters. Alternatively, the high dimensional data clustering (HDDC) approach recently proposed by Bouveyron *et al.* [2] can be applied. This technique combines dimensionality reduction with the expectation-maximization (EM) algorithm [5] to learn a Gaussian mixture model. By assuming certain dependencies in the covariance matrix, the learned clusters can be easily re-projected onto the original space yielding robust clusters with a significantly reduced risk of over-fitting.

Our approach can be seen as a reduced version of the method of Bouveyron *et al.* [2]. We perform an EM-based Gaussian mixture clustering in a reduced space and then use the obtained class coefficients to compute the mixture model in the high dimensional space.

A. Dimensionality Reduction

In dimensionality reduction techniques, one is interested in finding a mapping from the original, n -dimensional inputs space to a new space with $k < n$ -dimensions with a minimal loss of information. Principal component analysis (PCA) is an un-supervised technique that maximizes the variance in the data in the new space.

Let Σ be the covariance matrix of the input data D . PCA computes the eigenvalues λ_i and eigenvectors of Σ . Let Q be the matrix of the eigenvectors sorted according to the eigenvalues. We can then compute a matrix $\Delta = Q^T \Sigma Q$ so that Δ is a diagonal matrix with the eigenvalues on the diagonal in descending order. Let $\lambda_i \geq \lambda_j$ for all $i < j$. We consider only the first k dimensions for clustering that cover 95% of the variance which is given by $\sum_{i=1}^k \lambda_i [\sum_{i=1}^n \lambda_i]^{-1} \geq 0.95$.

By considering only the first k dimensions, we obtain an approximative but compact representation for laser range scans. Figure 2 depicts the obtained dimensionality reduction in real world settings.

Concretely, for each pose hypothesis \mathbf{x}_t , we simulate L complete range scans $\mathcal{D} = \{\mathbf{d}_1, \dots, \mathbf{d}_L\}$ at locations drawn uniformly from $\mathcal{U}(\mathbf{x}_t)$ using the given map \mathbf{m} of the environment. The simulation of the laser range scans \mathcal{D} 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 elements of the set \mathcal{D} of laser range scans are used to compute the PCA and thus lead to a projection into a reduced, k -dimensional space.

B. Clustering in the Reduced Space

Let $\tilde{\cdot}$ refer to quantities computed in the reduced, k -dimensional space. Thus, $\tilde{\mathcal{D}} = \{\tilde{\mathbf{d}}_1, \dots, \tilde{\mathbf{d}}_L\}$ are the elements of the set \mathcal{D} projected to the low-dimensional space

given the transformation matrices described in the previous section. In the reduced space, we are now able to efficiently cluster the range scans while reducing the risk of over-fitting (compare [2]).

To estimate the clusters in the low-dimensional space, we apply the EM algorithm to efficiently learn the mixture distribution. The EM algorithm iteratively assigns the reduced data scans in $\tilde{\mathcal{D}}$ to the mixture components and optimizes their parameters in the following manner. Consider that θ' denotes the current estimate of parameters $\tilde{\mu}_j$, $\tilde{\Sigma}_j$, and α_j . In the E-Step, we calculate the expected value of the complete log-likelihood

$$Q(\theta, \theta') = E \left[\log\{p(\tilde{\mathcal{D}}, Y | \theta)\} | \tilde{\mathcal{D}}, \theta' \right] \quad (3)$$

$$= \int_y \log\{p(\tilde{\mathcal{D}}, y | \theta)\} p(y | \tilde{\mathcal{D}}, \theta') dy, \quad (4)$$

where Y denotes data associations of the projected simulated data points $\tilde{\mathcal{D}}$ 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

$$\theta'' = \operatorname{argmax}_{\theta} Q(\theta, \theta') \quad (5)$$

by updating the cluster parameters according to

$$\alpha_j = \frac{1}{L} \sum_{l=1}^L P(j | \tilde{\mathbf{d}}_l, \theta'), \quad (6)$$

$$\tilde{\mu}_j = \frac{\sum_{l=1}^L P(j | \tilde{\mathbf{d}}_l, \theta') \tilde{\mathbf{d}}_l}{\sum_{l=1}^L P(j | \tilde{\mathbf{d}}_l, \theta')}, \quad (7)$$

$$\tilde{\Sigma}_j = \frac{\sum_{l=1}^L P(j | \tilde{\mathbf{d}}_l, \theta') (\tilde{\mathbf{d}}_l - \tilde{\mu}_j)(\tilde{\mathbf{d}}_l - \tilde{\mu}_j)^T}{\sum_{l=1}^L P(j | \tilde{\mathbf{d}}_l, \theta')}. \quad (8)$$

We now set $\theta' \leftarrow \theta''$ and iterate this procedure until the amount of improvement per iteration falls below a specified threshold. To determine the actual number of clusters in the resulting model, we apply the Bayesian information criterion and choose the model with the best score.

C. Transferring the Mixture Components to the Measurement Space

After identifying the individual clusters and the corresponding probabilities $P(j | \tilde{\mathbf{d}}_l, \theta')$, we can compute our mixture model in the high dimensional space. This can be easily achieved if we assume that the corresponding probabilities are identical in the reduced space as well as in the measurements space. Thus, the mixture in the high-dimensional space is given by

$$p(z_t | \mathbf{x}_t, \mathbf{m}) = \sum_{j=1}^J \alpha_j \mathcal{N}(\mathbf{x}_t, \mu_j, \Sigma_j), \quad (9)$$




Fig. 3. Evaluated likelihood for 61, and 181 laser beams (from left to right) for different sensor models (upper diagrams) and the two sensor models (*HDGM*, *GM*) which take the multi-modalities in the laser measurements into account (lower diagrams) at 847 robot poses in our office environment depicted in Figure 5.




Fig. 4. Standard deviations of the different sensor models for 31, 61, and 181 laser beams (left). Comparison of the standard deviation of the two sensor models (*HDGM*, *GM*) which take the multi-modalities in the expected laser measurements into account.

where J is the number of clusters and $\mathcal{N}(\mathbf{x}, \boldsymbol{\mu}, \boldsymbol{\Sigma})$ refers to the n -dimensional Gaussian evaluated at \mathbf{x} with mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$ according to

$$\boldsymbol{\mu}_j = \frac{\sum_{l=1}^L P(j | \tilde{\mathbf{d}}_l, \boldsymbol{\theta}') \mathbf{d}_l}{\sum_{l=1}^L P(j | \tilde{\mathbf{d}}_l, \boldsymbol{\theta}')} \quad (10)$$

$$\boldsymbol{\Sigma}_j = \frac{\sum_{l=1}^L P(j | \tilde{\mathbf{d}}_l, \boldsymbol{\theta}') (\mathbf{d}_l - \boldsymbol{\mu}_j) (\mathbf{d}_l - \boldsymbol{\mu}_j)^T}{\sum_{l=1}^L P(j | \tilde{\mathbf{d}}_l, \boldsymbol{\theta}')} \quad (11)$$

In contrast to former approaches which modeled the likelihood functions as unimodal distributions for single beams [8], [13] or entire scans [14], [16], or as a multi-modal distributions for single beams [15], we now consider high-dimensional, multi-modal mixture models. This allows us to take the dependency between the individual beams as well as the multi-modal nature of the distribution into account. As we will demonstrate in the experiments, this more sophisticated model significantly improves the ability of a mobile robot to localize itself.

V. EXPERIMENTS

The approach described above has been implemented and tested on data obtained with a mobile robot. 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.

Additionally, we show the improvements achieved by also considering the dependencies between the individual laser beams. Then in the second set of experiments, we analyze our high-dimensional Gaussian mixture model in a global localization task in which multi-modal situations frequently occur. We therefore compare it to alternative models, which do not simultaneously take into account the multi-modality and the dependencies between the individual laser beams. In particular, we compared the performance of the following sensor models:

HDGM: Our high-dimensional Gaussian mixture model as detailed in Section IV.

GM: The place-dependent beam-based Gaussian mixture sensor model as detailed in our previous work [15].

IB: The standard beam-based sensor model that assumes independent beams with an additive white noise component.

EC: The scan-based, place-dependent model with learned covariance matrix as detailed in our earlier previous work [14].

A. Likelihood Evaluation

In the first set of experiments, we evaluated the likelihood of the true position of the robot in a data set acquired using a real robot. We therefore compared our high-dimensional Gaussian mixture model (*HDGM*) to other likelihood models which are also based on ray casting operations (*GM*, *IB*, and *EC*). 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. In our previous work [15], we demonstrated 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. 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. The two upper diagrams of Figure 3 show the evaluated likelihood for 61, and 181 laser beams (from left to right) for different sensor models at 847 robot poses in our office environment depicted in Figure 5. The lower diagrams show the same for the two sensor models (*HDGM*, *GM*) which take the multi-modalities in the expected laser measurements into account. As can be seen from Figure 4, our high-dimensional Gaussian mixture model (*HDGM*) yields a smaller variance in the estimated likelihood of the true pose compared to the other sensor models. Additionally, the right diagram of this Figure illustrates that our novel high-dimensional model (*HDGM*) yields even a smaller variance in the estimated likelihood compared to the beam-based Gaussian mixture model (*GM*) especially when the number of integrated laser beams is increased. This higher variance in the estimated likelihoods, which is caused by the independence assumption of the beam-based sensor model might lead to a divergence of the probabilistic localization even in the case of position tracking.

B. Localization

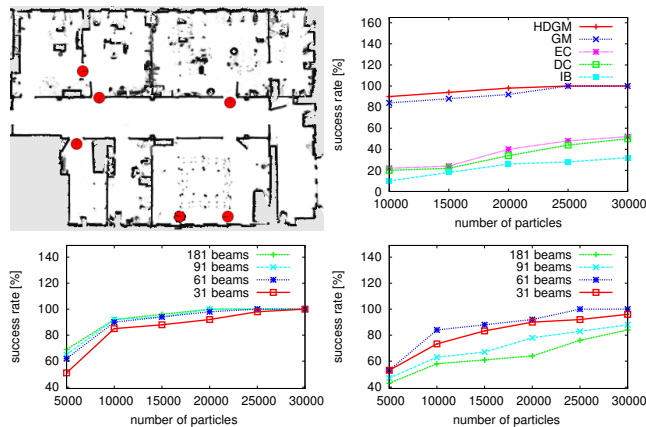


Fig. 5. The six positions with the highest probability that the global localization in the office environment fails (upper left). The upper right diagram shows the number of successful localizations after ten integrations of 61 measurements at these locations. The lower diagrams show the same experiment for the two multi-modal sensor models (left: *HDGM*, right: *GM*) for different beam numbers.

The second set of experiments is designed to illustrate that our new high-dimensional sensor model (*HDGM*), which takes the multi-modality as well as the dependencies of measurements into account, achieves a more robust and accurate localization than the other sensor models. The upper left image in Figure 5 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. At these positions typically 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 upper diagram of Figure 5 shows the number of successful localizations after ten integrations of 61 measurements at these locations. The lower diagrams show the same experiment for different beam numbers for the two sensor models (*HDGM*, *GM*) which take the multi-modalities in the expected laser measurements into account. The experiments show that our high-dimensional Gaussian mixture model (*GM*) allows us to more robustly localize the robot in situations in which the other models frequently fail. Additionally, it demonstrates that we are able to integrate more measurements to achieve a higher accuracy of the filtering process without losing robustness.

VI. CONCLUSIONS

In this paper, we presented a novel place-dependent sensor model for range scans that considers entire scans instead of individual beams and in this way overcomes the independence assumption underlying popular alternative models. At

the same time, it utilizes Gaussian mixture models to represent potential multi-modalities of the likelihood function. To reduce the dimensionality of the measurement space it applies the principle component analysis.

Our approach has been implemented and extensively tested on data obtained with mobile robots equipped with laser range finders. In our experiments, our new model showed superior performance over other popular models proposed in the past.

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, 2002.
- [2] C. Bouveyron, S. Girard, and C. Schmid. High-dimensional data clustering. *Computational Statistics and Data Analysis*, 52(1):502–519, 2007.
- [3] 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.
- [4] F. Dellaert, D. Fox, W. Burgard, and S. Thrun. Monte carlo localization for mobile robots. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, pages 99–141, 1998.
- [5] 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.
- [6] A. Doucet, N. de Freitas, and N. Gordon, editors. *Sequential Monte-Carlo Methods in Practice*. Springer Verlag, 2001.
- [7] 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.
- [8] D. Fox, W. Burgard, and S. 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] B. Limketkai, D. Fox, and L. Liao. Crf-filters: Discriminative particle filters for sequential state estimation. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2007.
- [12] 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.
- [13] 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 of *STAR Springer tracts in advanced robotics*, pages 181–194. Springer Verlag, 2006.
- [14] 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.
- [15] P. Pfaff, C. Plagemann, and W. Burgard. Gaussian mixture models for probabilistic localization. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2008.
- [16] 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.
- [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. Upercroft, S. Kumar, M.F. Ridley, L. Ong, and H.F. Durrant-Whyte. Fast re-parameterisation of Gaussian mixture models for robotics applications. In *Australian Conference on Robotics and Automation*, 2004.