

Closed-Form Full Map Posteriors for Robot Localization with Lidar Sensors

Lukas Luft, Alexander Schaefer, Tobias Schubert, Wolfram Burgard

Abstract—A popular class of lidar-based grid mapping algorithms computes for each map cell the probability that it reflects an incident laser beam. These algorithms typically determine the map as the set of reflection probabilities that maximizes the likelihood of the underlying laser data and do not compute the full posterior distribution over all possible maps. Thereby, they discard crucial information about the confidence of the estimate. The approach presented in this paper preserves this information by determining the full map posterior. In general, this problem is hard because distributions over real-valued quantities can possess infinitely many dimensions. However, for two state-of-the-art beam-based lidar models, our approach yields closed-form map posteriors that possess only two parameters per cell. Even better, these posteriors come for free, in the sense that they use the same parameters as the traditional approaches, without the need for additional computations. An important use case for grid maps is robot localization, which we formulate as Bayesian filtering based on the closed-form map posterior rather than based on a single map. The resulting measurement likelihoods can also be expressed in closed form. In simulations and extensive real-world experiments, we show that leveraging the full map posterior improves the localization accuracy compared to approaches that use the most likely map.

I. INTRODUCTION

Robot mapping and localization are probabilistic processes. Therefore, it is desirable to determine the posterior probability distribution over all possible maps given all observations rather than to determine a particular map. For some types of grid maps, it is well-known how to compute this distribution.

Grid maps are a popular representation of the environment of a robot. In their basic formulation, each voxel of the map holds a binary value which expresses whether the voxel is occupied or not. For these so-called occupancy grids, the posterior distribution over each occupancy state is characterized by one real-valued parameter. Moravec [8] and Elfes [3], [4] show how to compute this parameter.

In real-world scenarios, however, map voxels are not always completely free or completely occupied. They often contain structures smaller than the grid resolution. As occupancy grids are not capable of representing these structures, it makes sense to use real-valued maps.

Lukas Luft and Alexander Schaefer contributed equally to this work.

This work has been partially supported by the European Commission in the Horizon 2020 framework program under grant agreement 644227-Flourish; by the Graduate School of Robotics in Freiburg; and by the State Graduate Funding Program of Baden-Württemberg.

All authors are with the Department of Computer Science, University of Freiburg, Germany.

{luft, aschaefer, tobschub, burgard}
@cs.uni-freiburg.de

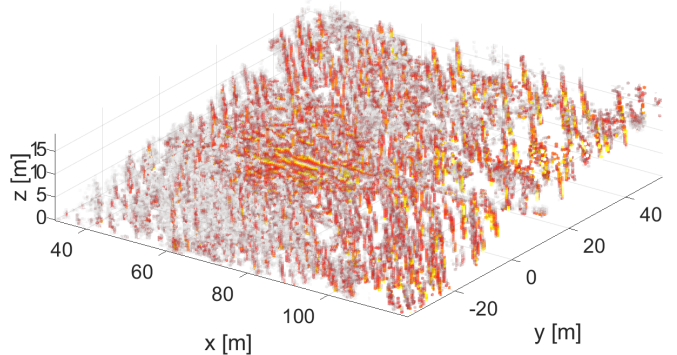


Fig. 1: Perspective view of a section of the reflection map built from the forest dataset. The map encodes the reflection probability of each voxel by the voxel color: Bright yellow corresponds to low reflection probability, dark red corresponds to high reflection probability. Although the map is highly cluttered, one can clearly recognize a large number of tree trunks.

A popular example of real-valued maps are so-called reflection maps [5]. Another method which also characterizes cells by real values builds so-called decay-rate maps [10]. In contrast to posteriors over discrete map values, posteriors over real-valued maps, like reflection maps and decay-rate maps, can contain infinitely many parameters. Therefore, one typically only computes the mode of the map posterior and uses it as map – with few exceptions [7].

In this paper, we present a method to derive the full posterior over real-valued grid maps based on data provided by lidar sensors. Our approach, which is applicable to a broad class of forward sensor models, relies on a rigorous Bayesian formulation of the mapping process. For the reflection model and the decay-rate model in particular, the proposed approach leads to closed-form map posteriors. These posteriors come for free: The most likely map already contains the required parameters.

In addition, we leverage the full map posterior for robot localization, the process of estimating the belief over the robot pose. Just like approaches that use a single given map, we can express the recursive Bayesian update in closed-form. Although our approach possesses the same computational complexity as the former ones, we demonstrate that it yields higher accuracies in extensive localization experiments.

II. RELATED WORK

The proposed approach computes the posterior over real-valued grid maps. Therefore, we structure our overview

of the related work as follows: We start with grid-based mapping approaches that compute the full posterior, move on to approaches which compute the most likely real-valued grid map, and close with posteriors over feature-based maps.

In robotics, occupancy grid maps are widely used. To derive the posterior over their binary values, most approaches assume that the individual voxels are independent. Then, the binary Bayes filter allows to recursively update the map posterior based upon the inverse sensor model, as shown by Moravec [8] and Elfes [3], [4]. To obtain a full posterior over a discrete map in the context of SLAM, Doucet et al. [2] and Tipaldi et al. [14] employ a Rao-Blackwellized particle filter [9]. Each of the particles represents not only a pose hypothesis, but also holds a distribution over a discrete map. Thrun [12] uses a forward sensor model to compute posteriors over grid maps. He drops the assumption of voxel independence by accounting for measurement noise. Marks et al. [7] present an approach to compute posterior distributions over real-valued grid maps; in their case, each map voxel represents the height variance of the surface.

Other than Marks et al. [7], most approaches for real-valued grid maps compute the most likely map only: Hähnel et al. [5], for example, extend the reflection model by introducing a binary variable that expresses whether a reflection is caused by a dynamic object or a static object. The recently introduced decay-rate model [10] also produces real-valued grid maps. These maps represent decay rates of the laser ray instead of reflection probabilities. All approaches in this paragraph have in common that they consistently leverage the forward sensor model for both mapping and localization.

Instead of using voxels, maps can also be represented by a finite set of landmarks. Extended Kalman Filtering techniques assume the robot pose and the positions of these landmarks to be normally distributed and compute the full posterior over the robot pose and the map in closed form. For one example among a wide range of publications in this context, see [1]. Extended Kalman Filtering is also a popular choice for collaborative localization, where robots use their teammates as moving landmarks. The belief of the joint pose state can then be interpreted as posterior over a dynamic landmark map, see for example [6].

III. APPROACH

This section describes how to calculate the full posterior over a real-valued grid map and how to use it for localization. Our approach is applicable to a broad class of beam-based sensor models, which we define in III-A. We call them factorizing models. In this section, we also recall the formulas for two examples of this class: the reflection model and the decay-rate model. In III-B, we derive a recursive update equation to compute the posteriors over grid maps based on factorizing models. We leverage this equation in III-C to derive closed-form posteriors over reflection maps and decay-rate maps. Once the posteriors are established, we move on to perform robot localization: Section III-D establishes a general recursive Bayesian update equation for

localization based on map posteriors rather than based on a single given map. For the reflection model and the decay-rate model, this update equation possesses closed form, as presented in III-E. Table I provides an overview over the notation used throughout the paper.

i	voxel index
v_i	i^{th} voxel
H_i	total number of hits in v_i during mapping
M_i	total number of misses in v_i during mapping
λ_i	decay rate in v_i
μ_i	reflection probability in v_i
x	sensor pose with respect to map frame
m	map
m_i	map value in voxel v_i
r	radius of a laser ray
r_i	distance that a ray travels inside v_i
R_i	total distance that all rays travel in v_i during mapping
\mathcal{I}	set of all voxels
$N = \mathcal{I} $	number of all voxels
$\mathcal{I}(r, x)$	set of voxels entered by a beam with r and x
X_m	sensor poses during mapping
Z	sensor measurements recorded during localization
Z_m	sensor measurements recorded during mapping
z	most recent measurement

TABLE I: Notation

A. Factorizing Forward Sensor Models

The present paper deals with mapping and localizing based on lidar data. The formalism presented later on is valid for a broad class of sensor models which we call factorizing models. They are characterized by the following property:

$$p(r \mid x, m) = \prod_{i \in \mathcal{I}(r, x)} f(r_i, m_i, \delta_i), \quad (1)$$

where r is the length of the measured laser ray, x denotes the sensor pose, and m is a fixed map, $\mathcal{I}(r, x)$ is the set of indices of all voxels which the beam enters, r_i is the radius that the beam travels within voxel v_i , m_i denotes the map value of voxel i , and $\delta_i(r, x)$ tells whether or not voxel i reflects the ray:

$$\delta_i = \delta_i(r, x) = \begin{cases} 1 & \text{if } v_i \text{ reflects the ray} \\ 0 & \text{else} \end{cases} \quad (2)$$

The reflection model [5] defines $f(r_i, m_i, \delta_i)$ as a binomial distribution over the event δ_i :

$$f(r_i, \mu_i, \delta_i) = \mu_i^{\delta_i} (1 - \mu_i)^{1 - \delta_i} = f(\mu_i, \delta_i). \quad (3)$$

It disregards the information about how far the ray travels inside each voxel: (3) does not depend on r_i .

In contrast, the decay-rate model [10] incorporates this information, as it computes f as follows:

$$f(r_i, \lambda_i, \delta_i) = \lambda_i^{\delta_i} e^{-\lambda_i r_i}. \quad (4)$$

Here, λ_i is the decay rate within voxel i . If we fix the value of r_i , Equation (4) yields a binary Poisson distribution over δ_i .

Conversely, for $\delta_i = 1$, it yields an exponential distribution over r_i .

B. Recursive Map Update

This section addresses the core of our approach. We show how to calculate a posterior distribution over all maps. This distribution is called the belief $bel(m) := p(m | Z_m, X_m)$, where Z_m denotes the set of all measurements recorded during the mapping process, and where X_m denotes the set of corresponding sensor poses. In order to calculate the full map posterior, we first need to introduce the following definitions.

$$bel(m) = p(m | Z_m, X_m) \quad (5)$$

$$\overline{bel}(m) = p(m | \overline{Z}_m, X_m) \quad (6)$$

$$bel(m_i) = p(m_i | Z_m, X_m) \quad (7)$$

$$\overline{bel}(m_i) = p(m_i | \overline{Z}_m, X_m) \quad (8)$$

Here, $\overline{Z}_m = Z_m \setminus \{z\}$ represents the set of mapping measurements without the most recent measurement z . Note that when referring to general sensor models, we call the sensor output z ; only in the context of factorizing models, we write r for radius. Proposition 1 now shows how to recursively compute the full map posterior from the latest belief and the latest measurement.

Proposition 1: Assuming a factorizing sensor model

$$p(r | x, m) = \prod_{i \in \mathcal{I}(r, x)} f(r_i, m_i, \delta_i), \quad (9)$$

and mutual independence of the individual voxels

$$bel(m) = \prod_{i=1}^N bel(m_i), \quad (10)$$

the belief over each map value is recursively updated according to

$$bel(m_i) = \eta_i \overline{bel}(m_i) f(r_i, m_i, \delta_i), \quad (11)$$

where η_i is a normalizing constant independent of m_i .

Proof: To prove the above proposition, we define the following notation:

$$m_{\mathcal{I} \setminus i} := m \setminus \{m_i\} \quad (12)$$

and

$$\int_{m_1} \dots \int_{m_{i-1}} \int_{m_{i+1}} \dots \int_{m_N} (\cdot) dm_N \dots dm_{i+1} dm_{i-1} \dots dm_1.$$

Using this notation, we derive Proposition 1 as follows:

$$\begin{aligned} bel(m_i) &= p(m_i | Z_m, X_m) \\ &\stackrel{(8)}{=} \eta p(z | m_i, \overline{Z}_m, X_m) \overline{bel}(m_i) \\ &= \eta \overline{bel}(m_i) \int_{m_{\mathcal{I} \setminus i}} p(z, m_{\mathcal{I} \setminus i} | m_i, \overline{Z}_m, X_m) dm_{\mathcal{I} \setminus i} \\ &\stackrel{(12)}{=} \eta \overline{bel}(m_i) \int_{m_{\mathcal{I} \setminus i}} p(z | m, \overline{Z}_m, X_m) \\ &\quad p(m_{\mathcal{I} \setminus i} | m_i, \overline{Z}_m, X_m) dm_{\mathcal{I} \setminus i} \\ &= \eta \overline{bel}(m_i) \int \prod_{j \in \mathcal{I}(r, x)} f(r_j, m_j, \delta_j) \\ &\quad \prod_{k \in \mathcal{I} \setminus \{i\}} \overline{bel}(m_k) dm_{\mathcal{I} \setminus i} \\ &= \eta_i \overline{bel}(m_i) f(r_i, m_i, \delta_i). \end{aligned} \quad (13)$$

η and η_i are normalizing constants independent of m_i . To obtain (13), we make use of both (9) and (10). To transition from (13) to the last line, we pull $f(r_i, m_i, \delta_i)$ out of the integral and merge the remaining integral, which is independent of m_i , with the normalizer. ■

The update equation (11) in Proposition 1 might look familiar: It is a generalization of the well-known map update $bel(m) = \eta \overline{bel}(m) p(z | x, m)$, which can be derived from Bayes rule in a straight-forward manner, see equation (7) in [11]. Another update equation which is related to (11) is the voxel-wise update for binary occupancy maps, see (18) in [11]. In contrast to the proposed update equation, the latter employs the inverse sensor model.

C. Closed-Form Map Posteriors

In this section, we leverage Proposition 1 to derive the closed-form map posteriors for the reflection model and for the decay-rate model. For the reflection model, update equation (11) yields the following posterior over μ_i :

$$\begin{aligned} bel(\mu_i) &\propto \prod_{Z_m} f(\mu_i, \delta_i) p(\mu_i) \\ &\stackrel{(3)}{\propto} \mu_i^{H_i} (1 - \mu_i)^{M_i} p(\mu_i) \\ &\propto \text{Beta}(H_i + 1, M_i + 1) p(\mu_i). \end{aligned} \quad (14)$$

Here, $p(\mu_i)$ denotes the prior distribution over the map, H_i tells how many rays are reflected in v_i , and M_i is the number of rays that penetrate v_i without reflection. $\text{Beta}(\cdot)$ denotes a beta distribution.

If the prior is a beta distribution $p(\mu_i) = \text{Beta}(\alpha, \beta)$, which is the conjugate prior for the binomial distribution $f(\mu_i, \delta_i)$, Equation (14) yields

$$bel(\mu_i) = \text{Beta}(H_i + \alpha, M_i + \beta). \quad (15)$$

The most likely reflection map can easily be derived from (15): Assuming a uniform prior distribution $p(\mu_i) = \text{Beta}(1, 1) = 1$ and computing the mode of the resulting beta posterior distribution yields the same result

as formulated by Hähnel et al. [5] for maximum likelihood reflection maps:

$$\mu_i^* = \frac{H_i}{H_i + M_i}. \quad (16)$$

For the decay-rate model, the update equation (11) becomes

$$\begin{aligned} \text{bel}(\lambda_i) &\propto \prod_{Z_m} f(r_i, \lambda_i, \delta_i) p(\lambda_i) \\ &\stackrel{(4)}{\propto} \lambda_i^{H_i} e^{-\lambda_i R_i} p(\lambda_i) \\ &\propto \text{Gamma}(H_i + 1, R_i) p(\lambda_i). \end{aligned}$$

$\text{Gamma}(\cdot)$ denotes a gamma distribution, $p(\lambda_i)$ is the prior map distribution, and R_i is the sum of the distances all rays travel within v_i .

If $p(\lambda_i)$ is a gamma distribution $\text{Gamma}(\alpha, \beta)$, which is the conjugate prior for the Poisson distribution and for the exponential distribution, we obtain the gamma-distributed belief¹

$$\text{bel}(\lambda_i) = \text{Gamma}(H_i + \alpha, R_i + \beta). \quad (17)$$

Setting $\alpha = 1$ and $\beta = 0$ leads to the so-called uninformative prior. Plugging this prior into (17) and computing the mode of the resulting posterior leads to the decay rates of the maximum likelihood approach as given in [10]:

$$\lambda_i^* = \frac{H_i}{R_i}. \quad (18)$$

For both the reflection model and the decay-rate model, the parameters α and β of the prior distribution need to be estimated during mapping. In Section IV, we explain how we obtain the prior parameters used throughout the experiments.

D. Localization with Map Posteriors

In this section, we formulate robot localization on the basis of the full posterior over the map rather than on the basis of a fixed map. In contrast to the well-known approach [13] which computes the belief over the robot pose on the basis of the given map m as

$$\text{bel}_m(x) = \eta \overline{\text{bel}}(x) p(z | x, m), \quad (19)$$

we leverage the full posterior $\text{bel}(m) = p(m | X_m, Z_m)$ instead of m . Consequently, we derive the pose belief as follows:

$$\begin{aligned} \text{bel}(x) &= p(x | Z, Z_m, X_m) \\ &= \eta p(z | x, \overline{Z}, Z_m, X_m) \underbrace{p(x | \overline{Z}, Z_m, X_m)}_{=: \text{bel}(x)} \\ &= \eta \overline{\text{bel}}(x) \int p(z | x, m) p(m | \overline{Z}, Z_m, X_m) dm \\ &= \eta \overline{\text{bel}}(x) \underbrace{\int p(z | x, m) \text{bel}(m) dm}_{=: L(z, x)}. \end{aligned} \quad (20)$$

¹In the context of height maps, Marks et al. [7] also obtain gamma-shaped posteriors over grid values.

Here, X_m and Z_m are the poses and measurements recorded during the mapping process, respectively, Z denotes the measurements recorded during localization, and $\overline{Z} = Z \setminus \{z\}$ is the set of all measurements but the most recent one.

Equations (19) and (20) differ only in the last term, which, in (19), is called the measurement likelihood. Analogously, we call $L(z, x)$ in (20) the measurement likelihood based on the map posterior. The next section presents closed-form solutions of $L(z, x)$ for the reflection model and the decay-rate model.

E. Closed-Form Measurement Likelihoods

For two particular factorizing sensor models, the reflection model and the decay-rate model, the solution of the integral contained in the measurement likelihood $L(z, x)$ leads to closed-form expressions. To compute those, we need to factorize $L(z, x)$ first:

$$\begin{aligned} L(z, x) &= \int_m p(z | x, m) \text{bel}(m) dm \\ &= \int_m \prod_{i \in \mathcal{I}(r, x)} f(r_i, m_i, \delta_i) \prod_{j=1}^N \text{bel}(m_j) dm \\ &= \int_{m_{\mathcal{I}(r, x)}} \prod_{i \in \mathcal{I}(r, x)} f(r_i, m_i, \delta_i) \prod_{j \in \mathcal{I}(r, x)} \text{bel}(m_j) dm_{\mathcal{I}(r, x)} \\ &= \prod_{i \in \mathcal{I}(r, x)} l(r_i, \delta_i) \end{aligned}$$

with

$$l(r_i, \delta_i) := \int f(r_i, m_i, \delta_i) \text{bel}(m_i) dm_i. \quad (21)$$

Now, we evaluate this integral for both sensor models. For the reflection model with beta-shaped prior $p(\mu_i) = \text{Beta}(\alpha, \beta)$, we take the distribution $f(\mu_i, \delta_i)$ from (3) and the posterior from (15), plug them into (21), and solve the integral:

$$l_{\text{ref}}(r_i, \delta_i) = l_{\text{ref}}(\delta_i) = \frac{(H_i + \alpha)^{\delta_i} (M_i + \beta)^{1-\delta_i}}{H_i + \alpha + M_i + \beta}. \quad (22)$$

If the posterior $\text{bel}(\mu_i)$ is set to a Dirac delta distribution $\text{Delta}(\mu_i - \mu_i^*)$, Equation (21) reproduces the measurement likelihood based on the most likely map, as derived in [5]:

$$l_{\text{ref}}(\delta_i) = f(\delta_i) = \frac{H_i^{\delta_i} M_i^{1-\delta_i}}{H_i + M_i}. \quad (23)$$

Equation (23) is only valid for $H_i + M_i > 0$. If no ray has visited the voxel during mapping, the maximum likelihood approach has to assign some initial value. In contrast, (22) is also valid for voxels that have not been visited by any ray during mapping.

Note that in the special case of $\alpha = \beta$, Equation (23) can be transformed into (22) using Laplace smoothing – see for example Equation (13) in [15].

In order to obtain closed-form solutions for the decay-rate model, we assume a gamma-shaped prior

$p(\lambda_i) = \text{Gamma}(\alpha, \beta)$ and plug (4) and (17) into (21). This leads to

$$l_{\text{dec}}(r_i, \delta_i) = \left(\frac{R_i + \beta}{R_i + \beta + r_i} \right)^{H_i + \alpha} \left(\frac{H_i + \alpha}{R_i + \beta + r_i} \right)^{\delta_i}$$

Analogously to the reflection model, the measurement likelihood based on the most likely map as derived in [10] can be reproduced from (21) by setting $\text{bel}(\lambda_i) = \text{Delta}(\lambda_i - \lambda_i^*)$:

$$l_{\text{dec}}(r_i, \delta_i) = f(r_i, \delta_i) = e^{-\frac{H_i}{R_i} r_i} \left(\frac{H_i}{R_i} \right)^{\delta_i}.$$

Until now, we have assumed that the lidar sensor reports a reflection for each emitted beam. In practice, however, the sensor range is always limited by a lower bound r_{\min} and an upper bound r_{\max} . The following equations show for both measurement models how to attribute probabilities to these out-of-range measurements:

$$P(r < r_{\min} \mid x, Z_m, X_m) = 1 - \prod_{i \in \mathcal{I}(r_{\min}, x)} l(r_i, \delta_i = 0),$$

$$P(r > r_{\max} \mid x, Z_m, X_m) = \prod_{i \in \mathcal{I}(r_{\max}, x)} l(r_i, \delta_i = 0).$$

IV. EXPERIMENTS

In the previous section, we have derived all the necessary equations to compute the posterior over a grid map and to leverage this posterior for localization. Now, to demonstrate that localization benefits from employing the full map posterior instead of the most likely map, we perform a simulation and extensive real-world experiments.

As shown in Section III-C, the computation of the map posterior for the reflection model and for the decay-rate model requires the specification of the two parameters α and β of the prior distribution. In our experiments, we determine them as follows. First, we compute the empirical mean and variance of the map value over all voxels. Then, we choose α and β such that the mean and variance of the parameterized distributions equal these empirical values. To that end, we solve the known equations for the mean and the variance of the beta and gamma distribution for α and β . For the reflection model, we obtain

$$\alpha = -\frac{\text{E}[\mu] \left(\text{E}[\mu]^2 - \text{E}[\mu] + \text{var}[\mu] \right)}{\text{var}[\mu]},$$

$$\beta = \frac{\text{E}[\mu] - \text{var}[\mu] + \text{E}[\mu] \text{var}[\mu] - 2 \text{E}[\mu]^2 + \text{E}[\mu]^3}{\text{var}[\mu]}.$$

For the decay-rate model, the parameters are

$$\alpha = \frac{\text{E}[\lambda]^2}{\text{var}[\lambda]},$$

$$\beta = \frac{\text{E}[\lambda]}{\text{var}[\lambda]}.$$

To strictly avoid the repeated use of the same information, one has to estimate these parameters for each voxel separately, computing the empirical mean and variance over all

voxels but the considered one. However, the sheer number of observations in our datasets render this effect negligible. Hence, it is sufficient to compute α and β only once.

Throughout the experiments section, MLM denotes the maximum likelihood approach, while FMP refers to the proposed approach, which leverages the full map posterior. REF and DEC denote the reflection model and the decay-rate model, respectively.

A. Localization in Simulation

We simulate the following scenario to evaluate the localization performance with both maximum likelihood maps and full posteriors: A mobile robot is located in a corridor that is modeled by a row of $N = 100$ consecutive voxels. Each experiment run consists of two phases: During the mapping phase, the robot visits every voxel n times and for each voxel collects the measurements H_i , M_i , and R_i , which it uses to build both a reflection map and a decay-rate map. In the localization phase, which consists of 100 iterations, the robot traverses the corridor from start to end knowing its motor commands. In each iteration, the robot moves to the next voxel, fires its sensor once, and updates the belief over its pose. Therefore, with the maximum likelihood approach it uses (19), with the proposed approach it uses (20). After the robot has arrived at the last voxel of the corridor, we evaluate the pose belief at the true pose averaged over each iteration:

$$\rho := \mathbf{E}(\text{bel}(x_{\text{true}})).$$

We perform 10,000 such localization runs for each $n \in \{1, 2, 3, 4, 5, 10, 20, 50, 100, 200\}$.

In each run, we synthesize a new map m by drawing samples from a uniform distribution for the reflection model or from a gamma distribution $\text{Gamma}(1, 1)$ for the decay-rate model. This map is hidden from the localization algorithms. We use m to simulate the measurements the robot records during mapping and localization by sampling from the distribution described in Equation (3) for the reflection model or from (4) for the decay-rate model. The belief over the robot pose is initialized with a uniform distribution over x .

For both measurement models, we compare three algorithms: MLM, FMP with uniform prior, and FMP with conjugate prior. The results in Figure 2 indicate that the proposed method – FMP with conjugate prior – yields higher localization accuracy than the compared methods for both measurement models. Moreover, we observe that the benefit of the proposed method is greater for the reflection model than for the decay-rate model. We perform a one-tailed, paired-sample t-test, which validates these observations: For the reflection model, the proposed method outperforms the other two approaches for $n \leq 100$ with a probability greater than 0.9999. For the decay-rate model, we obtain the same significance level for $n \leq 4$.

B. Real-World Localization

In order to validate the findings of the simulation in the real world, we test our approach on three datasets in

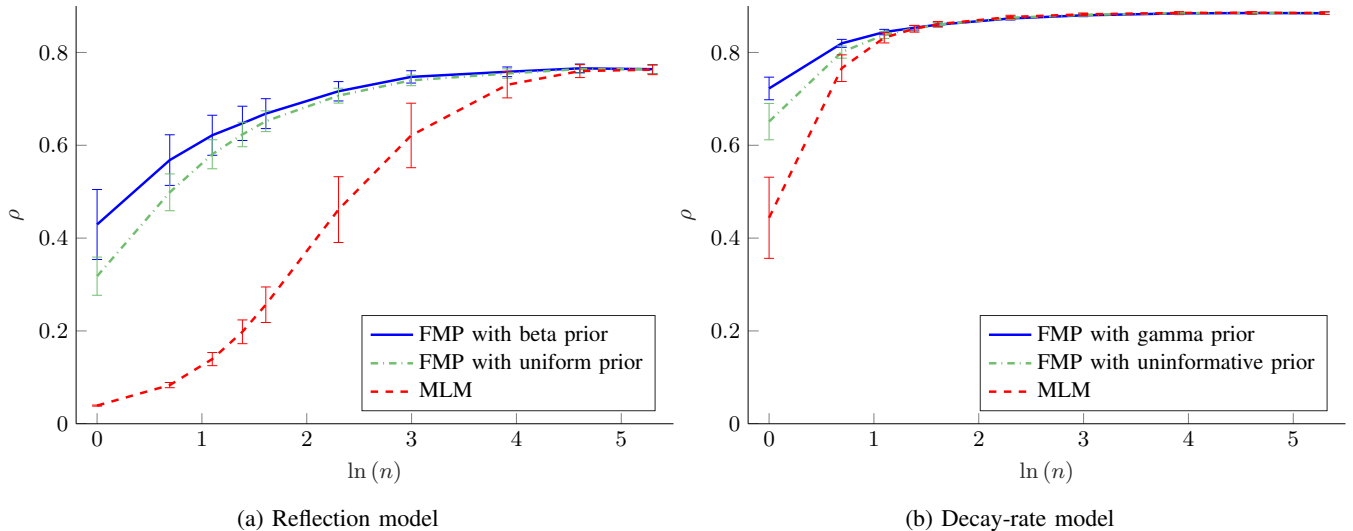


Fig. 2: Localization accuracy in a simulated environment. n is the number of observations per voxel during the mapping process. The data points correspond to $n \in \{1, 2, 3, 4, 5, 10, 20, 50, 100, 200\}$. The accuracy measure ρ shows the average probability which the algorithm assigns to the ground truth position. The error bars represent the variances over 10,000 runs.

which our mobile off-road robot navigates different kinds of environments: In the campus dataset, the robot drives around on the campus of the University of Freiburg. The forest dataset, a section of which is shown in Figure 1, was recorded on a small trail in the middle of a forest, while the park dataset contains a long trajectory along a broad road through an open forest. We tessellate the environment into cubic axis-aligned voxels with edge length 0.5 m. Consequently, the campus maps contain $444 \times 406 \times 43$ voxels, the park maps contain $515 \times 561 \times 41$ voxels, and the forest maps contain $393 \times 403 \times 86$ voxels.

We use the following hardware: The off-road robot VIONA by Robot Makers carries a Velodyne HDL-64E lidar sensor and an Applanix POS LV localization system, which provides a centimeter-accurate estimate of the robot pose by fusing the data from multiple GPS sensors, an IMU, and odometry. Due to its accuracy, we use the Applanix data as pose ground truth during mapping and in the evaluation of all localization estimates.

To assess the localization performance using full posteriors versus maximum-likelihood maps, we employ different measures:

1) *Measurement Likelihood*: We employ the logarithm of the measurement likelihood at the true pose as a measure of how well the approaches predict the measurements given the true pose. For MLM, the measurement likelihood is defined as $p(z | x, m)$. For FMP, it is defined by $L(z, x)$ in (20). We sum up the likelihood values of all measurements in a dataset and divide the result for MLM by the result for FMP. The ratios are presented in Table II. For both measurement models and all three datasets, these values are greater than one, which means that FMP achieves better prediction of the measurements than MLM.

	REF	DEC
campus	1.21	1.16
forest	1.22	1.31
park	1.25	1.27

TABLE II: Log-likelihood ratios of the measurements given the ground-truth pose. For each dataset, we show the ratio of the cumulated measurement log-likelihoods for MLM to the ones for FMP. Values greater than one mean that FMP attributes higher likelihoods to measurements given the true robot pose.

2) *Kullback-Leibler Divergence*: As a measure of dissimilarity, we employ the Kullback-Leibler (KL) divergence from the estimated position distribution p to the ground-truth pose distribution p_{gt} :

$$D_{KL}(p_{gt}(x) \parallel p(x)) = \int p_{gt}(x) \log \left(\frac{p_{gt}(x)}{p(x)} \right) dx. \quad (24)$$

We assume that p_{gt} is a bivariate normal distribution with standard deviation $\sigma_x = \sigma_y = 0.05 \text{ m}$, chosen according to the specifications of the Applanix localization system. The formula for the pose likelihood p depends on the approach:

$$p = \begin{cases} p(x | z, m) \propto p(z | x, m) & \text{for MLM} \\ p(x | z, X_m, Z_m) \propto L(z, x) & \text{for FMP} \end{cases} \quad (25)$$

To evaluate Equation (24), we perform a Monte-Carlo integration over x :

$$D_{KL}(p_{gt} \parallel p) \approx \sum \log \left(\frac{p_{gt}}{p} \right),$$

where p_{gt} and p are normalized over the samples. We compute the results shown in Table III by summing up

	REF	DEC
campus	1.24	1.25
forest	1.27	1.82
park	1.37	1.59

TABLE III: Ratios of the Kullback-Leibler divergence from the position distribution estimated by MLM and FMP to the ground truth distribution. For each dataset, we show the ratio of the cumulated KL divergences for the MLM approach to the corresponding value for the FMP approach. Ratios greater than one indicate that the distribution estimated by FMP is closer to the ground-truth than the one estimated by MLM.

the approximated D_{KL} values for all measurements in one dataset and dividing the result for MLM by the result of FMP. All ratios are greater than one, which implies that the distribution estimated by FMP is closer to the ground truth than the distribution estimated by MLM. Moreover, all ratios are higher than the corresponding ratios of the measurement likelihoods in Table II.

3) *Monte-Carlo Localization*: The results of the two aforementioned measures are entirely reproducible, in the sense that they are independent of localization algorithm design. In the corresponding experiments, the proposed approach always yields better results than the maximum likely approach. To demonstrate the consequences of this behavior in a real application, we evaluate the localization accuracy of two particle filters. Each filter employs 3000 particles to localize the off-road robot in six dimensions in the park scenario with the decay-rate model. The initial variance is 0.1 m in the translational dimensions and 0.1 rad in the rotational dimensions. The filters only differ in the method used to weight the particles: One uses the likelihood derived from the most likely map, the other uses the likelihood derived from the full map posterior. Figure 3 shows the corresponding localization errors averaged over ten runs.

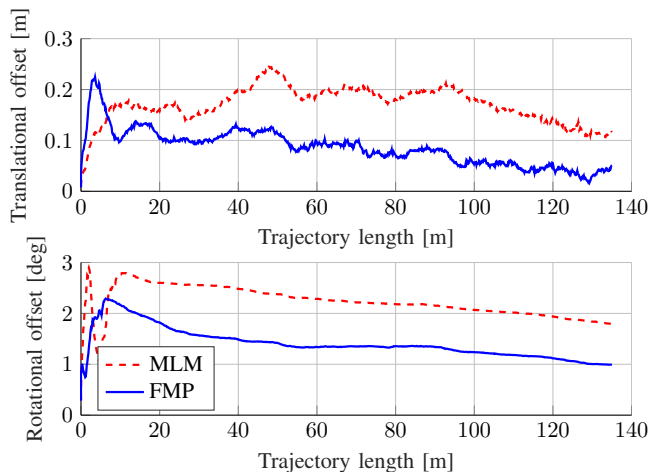


Fig. 3: Localization accuracy with the decay-rate model on the park dataset averaged over ten runs.

V. CONCLUSION AND FUTURE WORK

In this paper, we present an approach to compute posterior distributions over real-valued grid maps from lidar observations. We demonstrate that for the well-established reflection sensor model and the recently introduced decay-rate sensor model, the posterior distributions can be represented in closed form. Our approach requires the same measurement information and has the same computational demands as approaches which only determine the mode of the distribution. Simulations and extensive real-world experiments show that taking into account the full map posterior improves the accuracy of robot localization.

In the future, we plan to relax the assumption that all map cells are independent by accounting for measurement noise. Moreover, we will embed the presented approach into a SLAM framework, which we will then use to localize our off-road robot during operation.

REFERENCES

- [1] J. A. Castellanos, J. M. M. Montiel, J. Neira, and J. D. Tardos. The spmap: a probabilistic framework for simultaneous localization and map building. *IEEE Transactions on Robotics and Automation*, 15(5):948–952, Oct 1999.
- [2] A. Doucet, N. de Freitas, K. P. Murphy, and S. J. Russell. Rao-blackwellised particle filtering for dynamic bayesian networks. In *Proceedings of the 16th Conference on Uncertainty in Artificial Intelligence*, UAI '00, pages 176–183, San Francisco, CA, USA, 2000. Morgan Kaufmann Publishers Inc.
- [3] A. Elfes. *Occupancy grids: a probabilistic framework for robot perception and navigation*. PhD thesis, Carnegie Mellon University, 1989.
- [4] A. Elfes. Using occupancy grids for mobile robot perception and navigation. *Computer*, 22(6):46–57, June 1989.
- [5] D. Hähnel, R. Triebel, W. Burgard, and S. Thrun. Map building with mobile robots in dynamic environments. In *2003 IEEE International Conference on Robotics and Automation. Proceedings*, volume 2, pages 1557–1563, Sept. 2003.
- [6] L. Luft, T. Schubert, S. I. Roumeliotis, and W. Burgard. Recursive decentralized collaborative localization for sparsely communicating robots. In *Proceedings of Robotics: Science and Systems*, Ann Arbor, Michigan, June 2016.
- [7] T. K. Marks, A. Howard, M. Bajracharya, G. W. Cottrell, and L. Matthies. Gamma-slam: Using stereo vision and variance grid maps for slam in unstructured environments. In *2008 IEEE International Conference on Robotics and Automation*, pages 3717–3724, May 2008.
- [8] H. Moravec. Sensor fusion in certainty grids for mobile robots. *AI Mag.*, 9(2):61–74, July 1988.
- [9] K. Murphy. Bayesian map learning in dynamic environments. In *In Neural Info. Proc. Systems (NIPS)*, pages 1015–1021. MIT Press, 2000.
- [10] A. Schaefer, L. Luft, and W. Burgard. An analytical lidar sensor model based on ray path information. *IEEE Robotics and Automation Letters*, PP(99):1–1, 2017.
- [11] S. Thrun. Robotic mapping: A survey. In *Exploring Artificial Intelligence in the New Millenium*. Morgan Kaufmann, 2002.
- [12] S. Thrun. Learning occupancy grid maps with forward sensor models. *Auton. Robots*, 15(2):111–127, September 2003.
- [13] S. Thrun, W. Burgard, and D. Fox. *Probabilistic Robotics (Intelligent Robotics and Autonomous Agents)*. The MIT Press, 2005.
- [14] G. D. Tipaldi, D. Meyer-Delius, and W. Burgard. Lifelong localization in changing environments. 32(14), December 2013.
- [15] D. Valcarce, J. Parapar, and Á. Barreiro. Additive smoothing for relevance-based language modelling of recommender systems. In *Proceedings of the 4th Spanish Conference on Information Retrieval*, CERI '16, pages 9:1–9:8, New York, NY, USA, 2016. ACM.