2023 IEEE International Conference on Robotics and Automation (ICRA 2023)
May 29 - June 2, 2023. London, UK

LiDAR-based Indoor Localization with Optimal Particle Filters using Surface Normal Constraints

Heruka Andradi[1], Sebastian Blumenthal[2], Erwin Prassler[1], Paul G. Plöger[1]

***Abstract*— Accurate and robust localization systems are often highly desired in autonomous mobile robots. Existing LiDAR-based localization systems generally use standard particle filters which suffer from the well-known particle degeneracy problem. Furthermore, standard particle filters are ill-suited for handling discrepancies between maps and the actual operating environments. In this work, we present an effective LiDAR-based indoor localization system which addresses these two issues. The particle degeneracy problem is tackled with an efficient implementation of an optimal particle filter. Map discrepancies are then handled with the use of a high-fidelity observation model for accurate particle propagation and a separate low-fidelity observation model for robust weight update. Evaluations were carried out against a standard particle filter baseline on both real-world and simulated data from challenging indoor environments. The proposed system was found to show significantly better performance in-terms of accuracy, robustness to ambiguity, and robustness to map discrepancies. These performance gains were observed even with more than ten times smaller particle set sizes than in the baseline, while the increase in the computation time per particle was only around 20%.**

I. INTRODUCTION

Self-localization is an essential component of almost every autonomous mobile robot. For indoor environments, LiDAR based localization systems are often used as they do not rely on any external infrastructure or require modification of the operating environment. Furthermore, they benefit from the high depth accuracy, long range, and invariance to lighting conditions of typical LiDAR sensors. Due to ambiguities in the environment, its partial observability, and measurement noise, it is generally not possible to infer the robot's pose based on a single LiDAR scan. Particle filters are therefore widely used for localization, with wheel odometry as an additional source of input. In particle filters, the belief of the robot's pose is represented by a set of weighted samples which are updated recursively through Bayesian filtering. This allows for the representation of arbitrary multimodal uncertainty distributions that may arise when operating in challenging environments, unlike with Kalman filters which are restricted to unimodal Gaussians [1].

The most commonly used particle filtering algorithm for robot localization is the Sampling Importance Resampling (SIR) filter, which consists of three steps for each time step t. First, for each particle i with pose x_{t-1}^i from the previous time step, a new pose is drawn from some proposal distribution $q(x_t \mid x_{t-1}^i, z_t, u_t)$, where u_t and z_t are the control input and observation at time step t, respectively. Next, the particle weights are updated as

$$w_t^i = \frac{p(z_t \mid x_t^i) p(x_t^i \mid x_{t-1}^i, u_t)}{q(x_t^i \mid x_{t-1}^i, z_t, u_t)}. \tag{1}$$

Finally, the particles are resampled according to their weights to better represent the robot's pose. In the context of robot localization, $p(x_t \mid x_{t-1}^i, u_t)$ corresponds to the robot's motion model and $p(z_t \mid x_t)$ corresponds to the observation model. In LiDAR-based localization systems, the robot's motion model is a common choice for the proposal distribution as it simplifies (1) to $w_t^i = p(z_t \mid x_t^i)$. Similar to [2], we refer to particle filters with this choice of proposal distribution as standard particle filters.

There are two main problems with standard particle filters which adversely affect localization performance. The first is the well-known *particle degeneracy* problem. As illustrated in Fig. 1, robot motion models often tend to be much noisier in comparison to observation models from precise sensors. Standard particle filters are highly inefficient in such cases as most of the samples drawn from the motion model will be assigned negligible weights based on the observation model, and thereby get discarded during resampling. This results in reduced accuracies and poor robustness to ambiguity.

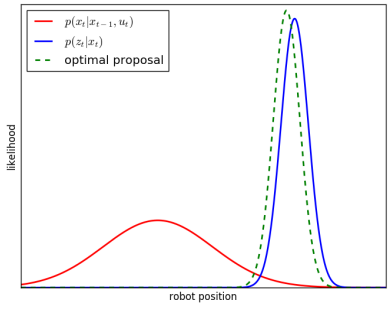

Fig. 1. Illustration of an optimal proposal distribution (green) for a typical noisy robot motion model (red) and a more precise observation model (blue)

The second problem with standard particle filters is the trade-off between accuracy and robustness that arises when attempting to handle map discrepancies. Discrepancies may be present between maps and real-world environments due to mapping errors, incompleteness, or changes in the operating environment. However due to the high computational complexity of additionally estimating the actual state of the environment, only the robot pose is usually included in the estimated state variable and the available map is assumed to be exact [1]. Therefore, map discrepancies can lead to

[1]Bonn-Rhein-Sieg University of Applied Sciences, Germany heruka.andradi@smail.inf.h-brs.de erwin.prassler@h-brs.de paul.ploeger@h-brs.de
[2]KELO Robotics GmbH, Germany blumenthal@kelo-robotics.com

localization failures due to the resulting incorrect observation likelihood estimates. While robustness to map discrepancies can be improved by modifying the observation model to alleviate their impact, this leads to a less informative observation model which thereby degrades accuracy [3], [4].

The main contribution of this work is an accurate and robust indoor localization system for wheeled mobile robots equipped with 3D LiDARs. As the localization system was primarily developed for ground-based logistics robots, we focus on 3 Degree of Freedom (3DOF) localization in dynamic and ambiguous environments. The proposed system employs an efficient implementation of an optimal particle filter to address the two aforementioned problems with standard particle filters. Extensive evaluation results are provided to verify that significant performance gains can be achieved by effectively tackling the two issues.

II. Related Work

It has been shown in [5] that the optimal proposal distribution which minimizes the variance of particle weights (and thereby particle degeneracy) is given by

$$q(x_t \mid x_{t-1}^i, z_t, u_t) = \frac{p(z_t \mid x_t)p(x_t \mid x_{t-1}^i, u_t)}{p(z_t \mid x_{t-1}^i, u_t)}. \quad (2)$$

An illustration of the optimal proposal can be seen in Fig. 1. There are two major challenges in using the optimal proposal distribution. It requires the ability to draw samples from the product distribution $p(z_t \mid x_t)p(x_t \mid x_{t-1}^i, u_t)$ and to evaluate its integral over the state space. As outlined in [6], there are two cases where the use of the optimal proposal distribution is tractable: if the state-space of the robot is discrete and finite, or if $p(x_t \mid x_{t-1}^i, z_t, u_t)$ is Gaussian.

While SIR particle filtering with the optimal proposal distribution has been effectively applied to landmark-based localization in [7], it is non-trivial for LiDAR -based localization where observation likelihoods can often only be evaluated point-wise in a non-parametric form [8]. Therefore in [9], Grisseti et al. presents a sampling-based approach of obtaining a Gaussian approximation of the optimal proposal distribution with scan matching. An alternative approach which does not rely on scan matching or Gaussian approximation has been presented by Blanco et al. in [8]. However, both approaches have high computational costs as they contain additional sampling steps which involve multiple observation likelihood evaluations per particle.

Similar to this work, a particle filter based approach with a closed-form Gaussian approximation of the optimal proposal is presented in [10]. A 2D probabilistic occupancy grid based map representation is used, and the optimal proposal is obtained using an Extended Kalman Filter (EKF). However, details on the computation of the observation Jacobians have not been provided. An alternative standard particle filter based approach is presented in [11], which utilizes a 2D vector map and a differentiable observation model based on point-to-line correspondences. A refinement step is introduced after new particles poses are drawn from the motion model, where the sampled poses are refined through several iterations of gradient descent using the observation model. Although this alleviates particle degeneracy, the gradient descent step adds significant computational burden.

To account for map discrepancies, Jianchao et al. proposes the use of a motion detection module in [12] to discard scan points from dynamic obstacles. However, this method only handles unmapped objects that are in motion during localization. Therefore in [13], temporary maps of the environment are constructed and used during localization such that observations of semi-static objects can be better utilized. A similar approach is provided in [14], where the temporary maps can additionally distinguish between semi-static and dynamic regions. However, both approaches rely on accurate estimates from a particle filter to generate the temporary maps, and therefore cannot be used when there is uncertainty in the robot's pose. In [15], Tipaldi et al. presents a multi-session Simultaneous Localization And Mapping (SLAM) approach where in addition to the robot pose, the environment is also included in the estimated state variable. However, this additional map generation leads to significantly higher computational requirements.

Unlike in existing work, the proposed approach employs a closed-form Gaussian approximation of the optimal proposal. This is obtained by utilizing surface normal information in LiDAR scans and maps to derive an efficient Gaussian approximation of the observation model. Map discrepancies are then effectively handled with the use of two separate observation models in the particle propagation and weight update steps of the filter. This allows for significant improvements in both accuracy and robustness with only a slight increase in computational costs.

III. Our Approach

The proposed localization system relies on a 3D point cloud map with reasonably accurate ($<0.1rad$) surface normal estimates. During mapping, surface normals of scan points are estimated through Principal Component Analysis (PCA) of neighbouring points. In simulated environments the maps are constructed with robot poses from the simulator while in real-world environments the output poses from the LeGO-LOAM mapping system are used as reference. [16].

A. Motion Model

In this work, we employ a simple Gaussian wheel odometry motion model similar to the omni-directional motion model in the amcl ROS package [17]. Based on odometry input, the relative motion of the robot $\mathbf{u_t} = [\Delta x_t \ \Delta y_t \ \Delta\theta_t]^T$ since the previous time step is first computed. Given the pose $\mathbf{x_{t-1}^i} = [x_{t-1} \ y_{t-1} \ \theta_{t-1}]^T$ of particle i in the previous time step, the prior distribution of its current pose is modeled as

$$\mathbf{x_t^i} \sim p(\mathbf{x_t}|\mathbf{x_{t-1}^i}, \mathbf{u_t}) = \underbrace{\mathbf{x_{t-1}^i} + \mathbf{R_z}(\theta_{t-1}^i)\mathbf{u_t}}_{\boldsymbol{\mu}_{mm,t}^i} + \mathbf{w_t^i}, \quad (3)$$

where $\mathbf{R_z}(\theta_{t-1}^i)$ is a 3D rotational matrix representing a θ_{t-1}^i radian rotation about the z-axis, and $\mathbf{w_t^i} \sim \mathcal{N}(0, \boldsymbol{\Sigma}_{mm,t}^i)$

is zero mean Gaussian process noise with the covariance $\boldsymbol{\Sigma}^{\boldsymbol{i}}_{\boldsymbol{mm,t}}$ computed as follows:

$$\begin{bmatrix}\sigma_x^2\\ \sigma_y^2\\ \sigma_\theta^2\end{bmatrix} = \begin{bmatrix}\alpha_1 & \alpha_2 & \alpha_3\\ \alpha_2 & \alpha_1 & \alpha_3\\ \alpha_4 & \alpha_4 & \alpha_5\end{bmatrix}\begin{bmatrix}|\Delta x_t|\\ |\Delta y_t|\\ |\Delta\theta_t|\end{bmatrix}, \tag{4}$$

$$\boldsymbol{\Sigma}^{\boldsymbol{i}}_{\boldsymbol{mm,t}} = \mathbf{R}(\theta^{\mathbf{i}}_{\mathbf{t-1}})\begin{bmatrix}\sigma_x^2 & 0 & 0\\ 0 & \sigma_y^2 & 0\\ 0 & 0 & \sigma_\theta^2\end{bmatrix}\mathbf{R}(\theta^{\mathbf{i}}_{\mathbf{t-1}})^T. \tag{5}$$

The α parameter values above are assigned based on the uncertainty in the wheel odometry estimates of the robot.

B. Observation Model

The observation model $p(\mathcal{Z}_t|\mathbf{x_t})$ of a particle i is generally a multimodal distribution over the entire state space. However since the motion model $p(\mathbf{x_t}|\mathbf{x^i_{t-1}}, \mathbf{u_t})$ is a unimodal Gaussian, only the mode closest to the motion model mean $\boldsymbol{\mu}^{\boldsymbol{i}}_{\boldsymbol{mm,t}}$ has a significant impact on the optimal proposal distribution (2) in general. The following steps describe the procedure for computing a closed-form Gaussian approximation of this mode by utilizing surface normal information.

1) Observation: For each point of a 3D LiDAR scan, its surface normal is first estimated as the cross product of the vector between adjacent points of the same laser beam and the vector between the points in the adjacent laser beams. Since the focus is on 3DOF localization, scan points with surface normals which are not closely parallel to the ground plane are discarded. Next, the resulting point cloud is downsampled to N (=300 in this work) points to limit the computational load. The downsampled point cloud $\mathcal{Z}_t = \{\langle \mathbf{z^1}, \hat{\mathbf{n}}^{\mathbf{1}}_{\mathbf{z}}\rangle, \langle \mathbf{z^2}, \hat{\mathbf{n}}^{\mathbf{2}}_{\mathbf{z}}\rangle, \cdots, \langle \mathbf{z^N}, \hat{\mathbf{n}}^{\mathbf{N}}_{\mathbf{z}}\rangle\}$ is then taken as the observation at time t, where $\mathbf{z^j}$ is the 3D position of the scan point j relative to the robot frame and $\hat{\mathbf{n}}^{\mathbf{j}}_{\mathbf{z}}$ is its surface normal.

2) Point-to-Plane Correspondences: Next, the points are transformed to the map frame with rotation $\mathbf{R}$ and translation $\mathbf{t}$ based on the mean pose $\boldsymbol{\mu}^{\boldsymbol{i}}_{\boldsymbol{mm,t}}$ of the particle's motion model distribution. For each transformed scan point, the corresponding nearest neighbor map point $\langle \mathbf{m^j}, \hat{\mathbf{n}}^{\mathbf{j}}\rangle$ is then found, where $\mathbf{m^j}$ is the 3D position of the point and $\mathbf{n^j}$ is it's surface normal. Correspondences with high angular deviation in surface normals (i.e. $\cos^{-1}(\mathbf{R}\hat{\mathbf{n}}^{\mathbf{j}}_{\mathbf{z}} \cdot \hat{\mathbf{n}}^{\mathbf{j}}) > \theta_{max}$) or large point-to-plane distances (i.e. $(\mathbf{R}\mathbf{z^j} + \mathbf{t} - \mathbf{m^j}) \cdot \hat{\mathbf{n}}^{\mathbf{j}} > D_{max}$) are considered invalid. In our implementation, we use values of $0.2rad$ and $0.3m$ for θ_{max} and D_{max} respectively.

3) Closed-Form Gaussian Approximation: Given the correspondences $\mathcal{M} = \{\langle \mathbf{m^1}, \hat{\mathbf{n}}^{\mathbf{1}}\rangle, \langle \mathbf{m^2}, \hat{\mathbf{n}}^{\mathbf{2}}\rangle, \cdots, \langle \mathbf{m^j}, \hat{\mathbf{n}}^{\mathbf{j}}\rangle\}$ of a particle i, the observation likelihood is computed as

$$p(\mathcal{Z}_t|\mathbf{x_t}) = \prod_{j=1}^{N} p(\mathbf{z^j}|\mathbf{x_t}) = \prod_{j=1}^{N} \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(d^j)^2}{2\sigma^2}}, \tag{6}$$

where σ is the sensor noise and d^j is the point-to-plane correspondence distance of the scan point j. Unlike with point-to-point correspondences, point-to-plane correspondences tend to remain unchanged for small changes in the robot's pose. Therefore for a pose $\mathbf{x_t}$ with a small relative motion $\Delta\mathbf{x_t} = [\Delta x\ \Delta y\ \Delta\theta]^T$ from $\boldsymbol{\mu}^{\boldsymbol{i}}_{\boldsymbol{mm,t}}$, the point-to-plane correspondence distance of scan point j can be computed as

$$d^j = \left(\mathbf{m^j} - \mathbf{R}\mathbf{R_z}(\boldsymbol{\Delta}\theta)\mathbf{z^j} - \begin{bmatrix}\Delta x\\ \Delta y\\ 0\end{bmatrix} - t\right) \cdot \hat{\mathbf{n}}^{\mathbf{j}}. \tag{7}$$

Using small angle assumptions this can then be simplified to

$$\begin{aligned} d^j = &\underbrace{-\begin{bmatrix}\hat{n}^j_x & \hat{n}^j_y & (\hat{\mathbf{n}}^{\mathbf{j}})^T\mathbf{R}\begin{bmatrix}-z^j_y\\ z^j_x\\ 0\end{bmatrix}\end{bmatrix}}_{\mathbf{a_j}}\Delta\mathbf{x_t}\\ &+\underbrace{\left(\mathbf{m^j} - R\mathbf{z^j} - t\right)\cdot\hat{\mathbf{n}}^{\mathbf{j}}}_{b_j}. \end{aligned} \tag{8}$$

It should be noted here that invalid correspondences from the previous step are assigned a constant distance value with $\mathbf{a_j} = \mathbf{0}$ and $b_j = D_{max}$. Substituting (8) to (6) then gives

$$p(\mathcal{Z}_t|\Delta\mathbf{x_t}) = \prod_{j=1}^{N} \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(\mathbf{a_j}\Delta\mathbf{x_t} + b_j)^2}{2\sigma^2}}. \tag{9}$$

Since this is a product of normal distributions, $p(\mathcal{Z}_t|\Delta\mathbf{x_t})$ will in turn follow a Gaussian distribution of the form $K^i\mathcal{N}(\bar{\boldsymbol{\mu}}^{\boldsymbol{i}}_{\boldsymbol{om,t}}, \boldsymbol{\Sigma}^{\boldsymbol{i}}_{\boldsymbol{om,t}})$. The mean $\bar{\boldsymbol{\mu}}^{\boldsymbol{i}}_{\boldsymbol{om,t}}$ can be obtained by minimizing the negative log-likelihood objective function $f(\Delta\mathbf{x_t}) = -\ln p(\mathcal{Z}_t|\Delta\mathbf{x_t})$, and thereby computed efficiently by solving the resulting linear least squares formulation

$$\bar{\boldsymbol{\mu}}^{\boldsymbol{i}}_{\boldsymbol{om,t}} = \underset{\Delta\mathbf{x_t}}{\arg\min} \sum_{j=1}^{N} (\mathbf{a_j}\Delta\mathbf{x_t} + b_j)^2 \tag{10}$$

through Singular Value Decomposition (SVD). The covariance $\boldsymbol{\Sigma}^{\boldsymbol{i}}_{\boldsymbol{om,t}}$ can be computed as the inverse Hessian $\mathbf{H}$ of the objective function [18], [19],

$$\boldsymbol{\Sigma}^{\boldsymbol{i}}_{\boldsymbol{om,t}} = \mathbf{H^{-1}} = \sigma^2\left(\sum_{j=1}^{N} \mathbf{a_j}^T\mathbf{a_j}\right)^{-1}. \tag{11}$$

Substituting $\boldsymbol{\Delta\mathbf{x_t}} = \bar{\boldsymbol{\mu}}^{\mathbf{i}}_{\mathbf{om,t}}$ to (9), the scaling factor K^i can then be evaluated as

$$K^i = \sqrt{(2\pi)^3\left|\boldsymbol{\Sigma}^{\boldsymbol{i}}_{\boldsymbol{om,t}}\right|}\ \prod_{j=1}^{N} \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(b_j - \mathbf{a_j}\bar{\boldsymbol{\mu}}^{\mathbf{i}}_{\mathbf{om,t}})^2}{2\sigma^2}}, \tag{12}$$

resulting in the Gaussian observation model $p(\mathcal{Z}_t|\mathbf{x_t}) = K^i\mathcal{N}(\boldsymbol{\mu}^{\boldsymbol{i}}_{\boldsymbol{om,t}}, \boldsymbol{\Sigma}^{\boldsymbol{i}}_{\boldsymbol{om,t}})$ with mean $\boldsymbol{\mu}^{\boldsymbol{i}}_{\boldsymbol{om,t}} = \boldsymbol{\mu}^{\boldsymbol{i}}_{\boldsymbol{mm,t}} + \bar{\boldsymbol{\mu}}^{\boldsymbol{i}}_{\boldsymbol{om,t}}$.

4) Accounting for Noisy Surface Normals: Noisy surface normals in maps can lead to over-confident covariance estimates at ambiguous locations. For example, even along a completely ambiguous long corridor, this could lead to a variance of around $\frac{100\ \sigma^2}{N}$ for an observation of N scan points with surface normal deviations of 0.1 radians. To tackle this problem, we perform an Eigen decomposition of the covariance matrix $\boldsymbol{\Sigma}^{\boldsymbol{i}}_{om,t}$ and set all eigenvalues greater than $\frac{100\ \sigma^2}{N}$ to a very large value (e.g. 10^6). The covariance matrix $\boldsymbol{\Sigma}^{\boldsymbol{i}}_{om,t}$ is then reconstructed with the new

eigenvalues such that the uncertainty of the robot's position is better represented. Similarly to better represent rotational ambiguity, we nullify the effect of rotation on the point-to-plane distance of observation points with near parallel surface normals vectors (i.e $\cos^{-1}(\mathbf{z^j} \cdot \hat{\mathbf{n}}^{\mathbf{j}}_{\mathbf{z}}) < 0.1rad$). This is done by setting the third component of $\mathbf{a_j}$ to zero for such points when computing the observation covariance in (11).

C. Optimal Particle Filter

To minimize the effects of particle degeneracy, an optimal particle filter based localization system is employed. The following three steps are performed at each time step of the filter. At the end of each time step, the robot's pose is estimated as the average of the particle poses.

1) Predict: For each particle i, new poses $\mathbf{x^i_t}$ are first sampled from the optimal proposal distribution,

$$q(\mathbf{x_t}|\mathbf{x^i_{t-1}}, \mathcal{Z}_t, \mathbf{u_t}) = \frac{p(\mathcal{Z}_t|\mathbf{x_t})p(\mathbf{x_t}|\mathbf{x^i_{t-1}}, \mathbf{u_t})}{\int p(\mathcal{Z}_t|\mathbf{x'_t})p(\mathbf{x'_t}|\mathbf{x^i_{t-1}}, \mathbf{u_t})\mathbf{dx'_t}}. \quad (13)$$

With the Gaussian motion and observation models computed above, the mean vector $\boldsymbol{\mu^i_t}$ and covariance $\boldsymbol{\Sigma^i_{om,t}}$ of the resulting Gaussian proposal distribution can be computed as follows:

$$\boldsymbol{\Sigma^i_t} = \left[\left(\boldsymbol{\Sigma^i_{mm,t}}\right)^{-1} + \left(\boldsymbol{\Sigma^i_{om,t}}\right)^{-1}\right]^{-1}, \quad (14)$$

$$\boldsymbol{\mu^i_t} = \boldsymbol{\Sigma^i_t}\left[\left(\boldsymbol{\Sigma^i_{mm,t}}\right)^{-1}\ \boldsymbol{\mu^i_{mm,t}} + \left(\boldsymbol{\Sigma^i_{om,t}}\right)^{-1}\ \boldsymbol{\mu^i_{om,t}}\right]. \quad (15)$$

Thereby for each particle i, its new pose $\mathbf{x^i_t}$ is sampled from the normal distribution $\mathcal{N}(\boldsymbol{\mu^i_t}, \boldsymbol{\Sigma^i_t})$.

2) Update: The weight of each particle i is then updated based on the integral of the product of the motion and observation distributions over the entire state space as

$$\begin{aligned} w^i_t &= \int p(\mathcal{Z}_t|\mathbf{x'_t})p(\mathbf{x'_t}|\mathbf{x^i_{t-1}}, \mathbf{u_t})\mathbf{dx'_t} \\ &= p(\mathcal{Z}_t|\boldsymbol{\mu^i_t})\ p(\boldsymbol{\mu^i_t}|\mathbf{x^i_{t-1}}, \mathbf{u_t})\ \sqrt{(2\pi)^3|\boldsymbol{\Sigma^i_t}|}. \end{aligned} \quad (16)$$

3) Resample: Finally, the particles are resampled using the low variance sampling algorithm presented in [20].

D. Handling of Map Discrepancies

As described previously, attempting to account for map discrepancies through the observation model degrades accuracy, as it lowers the confidence in the observations and affects the effective compensation of odometry noise. Unlike with standard particle filters however, the observation model appears both in the `prediction` and `update` steps of optimal particle filters. Therefore, we leverage on this trait to improve robustness to map discrepancies without sacrificing on accuracy. A high-fidelity observation model assuming no map discrepancies is used in the `prediction` step which enables each particle to accurately track its pose at unambiguous locations by sampling from an observation dominant proposal distribution. A lower-fidelity observation model is then used in the `update` step, such that any uncertainty in the robot's pose can be better represented by the particles while being less susceptible to map discrepancies.

For the high-fidelity observation model, we use the observation model described in section III-B, with a sensor noise value of $\sigma = 0.02m$ according to our LiDAR specifications. For the choice of the low-fidelity observation model, we present two options that have been implemented and evaluated as separate localization systems:

1) $mcl\_optimal$: A simple and computationally efficient choice is to use the previously described Gaussian observation model with inflated sensor noise. In our implementation, a noise value of $\sigma_{lf} = 2m$ is used.
2) $mcl\_optimal\_raycast$: To better cope with incomplete maps and dynamic objects in the environment, a ray casting based observation model can be used. Here, only particles with scan points that appear to pass through static surfaces of the map are penalized in the weight update step. At each time step t, a set of downsampled scan points $\mathcal{Y} = \{\mathbf{y^1_t}, \mathbf{y^2_t}, \cdots, \mathbf{y^L_t}\}$ of size L are first obtained. For each particle i, ray casting is then carried out with these points from the particle's pose $\boldsymbol{x^i_t}$. The particle weights are then updated as follows:

$$p(\mathbf{y^k_t}|\boldsymbol{x^i_t}) = \begin{cases} 1 & |\mathbf{y^k_t}| < d^k + D_{thresh} \\ P_{miss} & \text{otherwise} \end{cases}, \quad (17)$$

$$w^i_t \propto p(\mathcal{Z}_t|\boldsymbol{x_t}) = \prod_{k=1}^{L} p(\mathbf{y^k_t}|\boldsymbol{x^i_t}), \quad (18)$$

where d^k is the ray casted distance, D_{thresh} is an allowance threshold for mapping errors, and P_{miss} is a penalty value between zero and one. In our implementation, we use $P_{miss} = 0.9$, $D_{thresh} = 0.1m$ and $L = 10$ as defaults.

IV. Evaluation

In this section we provide evaluation results for accuracy, robustness, and runtime on both real-world and simulated challenging indoor environments. The localization tests were carried out using an omni-directional mobile robot with wheel odometry and a 16 channel 3D LiDAR (Robosense RS-LiDAR-16 for real-world tests). To assess the effectiveness of tackling particle degeneracy, a 10% systematic error was introduced to odometry estimates during all experiments. For fair comparison, the proposed system was evaluated against a standard particle filter based baseline implementation (`mcl_standard`) which uses the same motion and observation models given in (3) and (6).

A. Accuracy

Accuracy evaluations were carried out at a roughly 7m long square room with planar walls. Due to the unavailability of ground-truth, localization errors were estimated against reference poses obtained by performing scan matching on the full LiDAR scan point clouds. The translational accuracy results can be seen in Fig. 2. The results of $mcl\_optimal\_raycast$ has been excluded since it has

the same proposal distribution (and thereby accuracy) as $mcl\_optimal$. It can be seen that even with a single particle, $mcl\_optimal$ achieves significantly higher accuracy than $mcl\_standard$ with 500 particles. Furthermore, it should also be noted here that the accuracy of $mcl\_standard$ degrades as the σ (sensor noise) value used in the observation model is inflated. Together with the following results on robustness, this highlights the trade-off between accuracy and robustness in standard particle filters. Similar accuracy results were also observed for the rotational estimates in the real-world environment, as well as on a simulated Gazebo environment with ground truth.

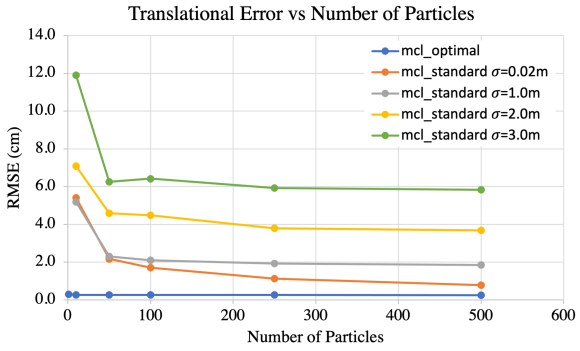

Fig. 2. Translational accuracy results in the square room environment

B. Robustness to Ambiguity

Robustness to ambiguity was evaluated with simulated data from three challenging environments: a 100m long featureless corridor, a set of three identical square rooms of 10m length, and a large circular room of 20m diameter. The path taken by the robot in the circular room environment is shown in Fig. 3. A large unmapped partition occludes the exit way during localization, such that there is roto-translational ambiguity in the generated LiDAR scans. Ideally, the localization system must capture this uncertainty through appropriate dispersion of particles such that the robot is able to recover its exact pose as it reaches the exit way.

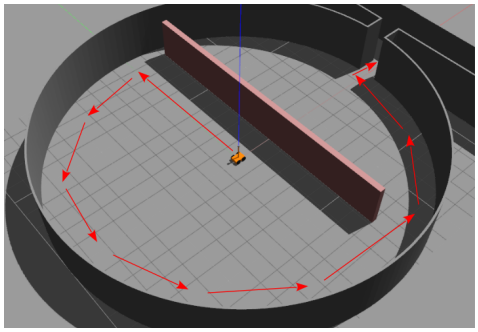

Fig. 3. Circular room environment with large unmapped partition (in red)

Localization results over 10 trials can be seen in Fig. 4. A trial is considered "fault-free" if the localization system maintains at least a single particle within a meter to the robot's true pose throughout the run, and recovers the exact pose at the end. It can be seen from the results that the robustness of $mcl\_standard$ improves to a certain extent when the sensor noise σ is inflated. However, as observed previously, this is at the cost of lower accuracy. Despite inflating the sensor noise in $mcl\_standard$, it is still found to be significantly outperformed by $mcl\_optimal$. Even with just 50 particles, $mcl\_optimal$ is seen to be more robust than $mcl\_standard$ run with up to 500 particles. Similar results were obtained for the experiments in the identical rooms and corridor environments. The particle distribution during a localization run is illustrated in Fig. 5. It can be seen how $mcl\_optimal$ captures the ambiguity more effectively, allowing for greater robustness with fewer particles.

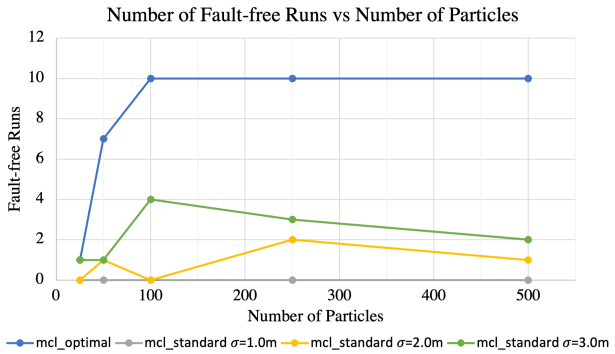

Fig. 4. Robustness results in the circular room environment over 10 trials

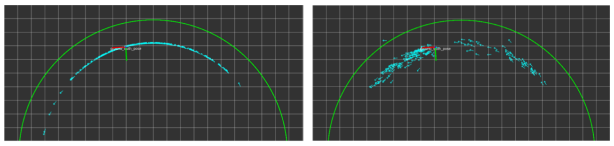

(a) mcl optimal (b) mcl standard $\sigma = 3m$

Fig. 5. Particle distribution (blue arrows) during localization in circular room environment

C. Robustness to Map Discrepancies

A simulation test was carried out in a simple environment consisting of a planar corridor and a door, as shown in Fig. 6. The door is left open during mapping, and closed during localization. Localization results over 10 trials can be seen in Fig. 7. A trial is considered successful if pose estimate at the end of the run is within 1m of the true pose. As expected, inflating the σ (sensor noise) values in $mcl\_standard$ is found to improve robustness as it dampens the impact of scan points from the unmapped door. $mcl\_optimal$ is found to perform worse than $mcl\_standard$ with $\sigma = 3m$, as the low fidelity observation model of $mcl\_optimal$ uses a σ value of 2m. Counter-intuitively, the performance of $mcl\_optimal$ is seen to drop with increasing particles. This is due to the higher chance of particles being drawn at wrong locations where the scans are more consists with the false map. $mcl\_optimal\_raycast$ on the other hand achieves a 100% success rate with just 25 particles, as it can identify the door as a dynamic object and account for it accordingly.

Thereby it outperforms the baseline systems even with 20 times fewer particles.

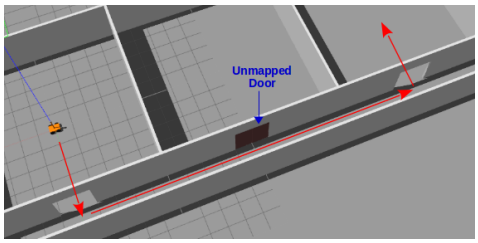

Fig. 6. Closed door environment

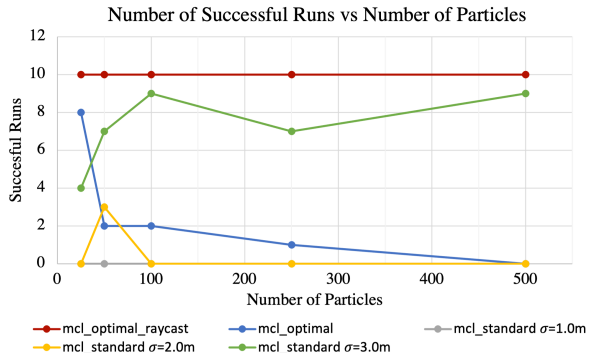

Fig. 7. Robustness results in the closed door environment over 10 trials

Robustness to ambiguity and map discrepancies was also evaluated on real-world data collected for the ROPOD project at a challenging hospital environment. The environment consists of 60m long featureless corridor with both dynamic and semi-static objects as seen in Fig. 8. The localization results are shown in Fig. 9. Again, $mcl\_optimal\_raycast$ is seen to have a 100% success rate with just 25 particles, which outperform the baseline systems having even 20 times as many particles. Unlike with the closed door tests, $mcl\_optimal$ is also found to outperform $mcl\_standard$ (with $\sigma = 3m$) due to the less challenging impact of map discrepancies and greater ambiguity in the environment.

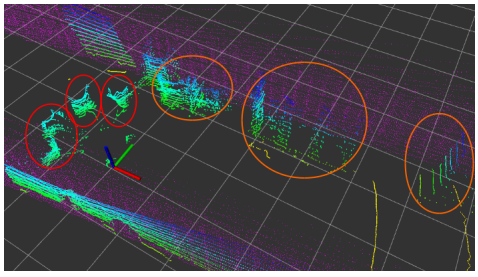

Fig. 8. Map points (purple) and scan at hospital corridor with occlusion from carts (orange circles) and people (red circles).

D. Runtime

The computation time per particle to perform a single `prediction` and `update` step is shown in Fig. 10. It can be seen that the computational costs of the proposed systems are only 10-20% higher than the baseline. This is due to the

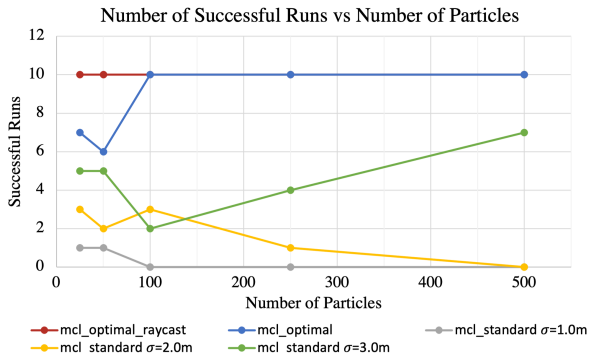

Fig. 9. Robustness results in the hospital environment over 10 trials

additional computation time of the optimal proposal being relatively small in comparison to the $N = 300$ costly nearest neighbor search operations in the observation model. Furthermore, since ray casting is performed on a heavily downsampled set ($L = 10$) of points, $mcl\_optimal\_raycast$ is found to be only slightly more expensive than $mcl\_optimal$.

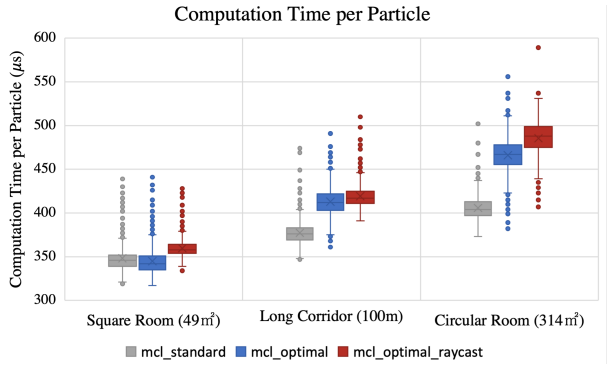

Fig. 10. Computation time per particle

V. Conclusion

In this paper, we have presented an effective indoor localization system for 3DOF wheeled mobile robots equipped with 3D LiDAR sensors. The proposed system addresses two of the major issues with the typically used standard particle filter based localization systems. The particle degeneracy problem was tackled with an efficient implementation of an optimal particle filter by utilizing surface normal constraints. The accuracy-robustness trade-off when dealing with map discrepancies was handled with the use of separate observation models in the prediction and update steps of the filter. Evaluation results on both simulated and real-world data confirm that the proposed solutions enable significantly higher accuracy, robustness to ambiguity, and robustness to map discrepancies with only slightly higher per particle computational costs. In future work, the proposed system needs to be extended and evaluated with 6DOF localization.

Acknowledgment

The authors acknowledge the support of the Bonn-Rhein-Sieg University of Applied Sciences and the Bonn-Aachen International Center for Information Technology.

REFERENCES

[1] F. Dellaert, D. Fox, W. Burgard, and S. Thrun, "Monte carlo localization for mobile robots," in *Proceedings 1999 IEEE International Conference on Robotics and Automation (Cat. No. 99CH36288C)*, vol. 2. IEEE, 1999, pp. 1322–1328.

[2] J. L. Blanco-Claraco, F. Mañas-Alvarez, J. L. Torres-Moreno, F. Rodriguez, and A. Gimenez-Fernandez, "Benchmarking particle filter algorithms for efficient velodyne-based vehicle localization," *Sensors*, vol. 19, no. 14, p. 3155, 2019.

[3] D. Fox, W. Burgard, and S. Thrun, "Markov localization for mobile robots in dynamic environments," *Journal of artificial intelligence research*, vol. 11, pp. 391–427, 1999.

[4] A. Q. Li, M. Xanthidis, J. M. O'Kane, and I. Rekleitis, "Active localization with dynamic obstacles," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2016, pp. 1902–1909.

[5] A. Doucet, "On sequential simulation-based methods for bayesian filtering," 1998.

[6] M. S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, "A tutorial on particle filters for online nonlinear/non-gaussian bayesian tracking," *IEEE Transactions on signal processing*, vol. 50, no. 2, pp. 174–188, 2002.

[7] M. Montemerlo, S. Thrun, D. Koller, B. Wegbreit, *et al.*, "Fastslam 2.0: An improved particle filtering algorithm for simultaneous localization and mapping that provably converges," in *IJCAI*, vol. 3, 2003, pp. 1151–1156.

[8] J.-L. Blanco, J. González, and J.-A. Fernández-Madrigal, "Optimal filtering for non-parametric observation models: applications to localization and slam," *The International Journal of Robotics Research*, vol. 29, no. 14, pp. 1726–1742, 2010.

[9] G. Grisetti, C. Stachniss, and W. Burgard, "Improved techniques for grid mapping with rao-blackwellized particle filters," *IEEE transactions on Robotics*, vol. 23, no. 1, pp. 34–46, 2007.

[10] Y. Xiao, Y. Ou, and W. Feng, "Localization of indoor robot based on particle filter with ekf proposal distribution," in *2017 IEEE international conference on cybernetics and intelligent systems (CIS) and IEEE conference on robotics, automation and mechatronics (RAM)*. IEEE, 2017, pp. 568–571.

[11] J. Biswas, B. Coltin, and M. Veloso, "Corrective gradient refinement for mobile robot localization," in *2011 IEEE/RSJ international conference on Intelligent Robots and Systems*. IEEE, 2011, pp. 73–78.

[12] S. Jianchao, Z. Xuebo, S. Lei, and L. Jingtai, "Map-based robust localization for indoor mobile robots," in *2017 36th Chinese Control Conference (CCC)*. IEEE, 2017, pp. 6945–6950.

[13] D. Meyer-Delius, J. Hess, G. Grisetti, and W. Burgard, "Temporary maps for robust localization in semi-static environments," in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2010, pp. 5750–5755.

[14] R. Valencia, J. Saarinen, H. Andreasson, J. Vallvé, J. Andrade-Cetto, and A. J. Lilienthal, "Localization in highly dynamic environments using dual-timescale ndt-mcl," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2014, pp. 3956–3962.

[15] G. D. Tipaldi, D. Meyer-Delius, and W. Burgard, "Lifelong localization in changing environments," *The International Journal of Robotics Research*, vol. 32, no. 14, pp. 1662–1678, 2013.

[16] T. Shan and B. Englot, "Lego-loam: Lightweight and ground-optimized lidar odometry and mapping on variable terrain," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 4758–4765.

[17] G. Brian, "amcl," https://github.com/ros-planning/navigation/commits/noetic-devel/amcl, 2009.

[18] S. Bonnabel, M. Barczyk, and F. Goulette, "On the covariance of icp-based scan-matching techniques," in *2016 American Control Conference (ACC)*. IEEE, 2016, pp. 5498–5503.

[19] A. Censi, "On achievable accuracy for range-finder localization," in *Proceedings 2007 IEEE International Conference on Robotics and Automation*. IEEE, 2007, pp. 4170–4175.

[20] S. Thrun, "Probabilistic robotics," *Communications of the ACM*, vol. 45, no. 3, pp. 52–57, 2002.