

Highly Accurate Maximum Likelihood Laser Mapping by Jointly Optimizing Laser Points and Robot Poses

Michael Ruhnke

Rainer Kümmerle

Giorgio Grisetti

Wolfram Burgard

Abstract—In this paper we describe an algorithm for learning highly accurate laser-based maps that treats the overall mapping problem as a joint optimization problem over robot poses and laser points. We assume that a laser range finder senses points sampled from a regular surface and we utilize an improved likelihood function that accounts for two phenomena affecting the laser measurements that are often neglected: the conic shape of the laser beam and the incidence angle. To solve the entire problem we apply an optimization procedure that jointly adjusts the position of all the robot poses and all points in the scans. As a result, we obtain highly accurate maps. We evaluated our approach using simulated and real-world data and we show that utilizing the estimated maps greatly improves the localization accuracy of robots. The results furthermore suggest that the accuracy of the resulting map can exceed the resolution of the laser sensors used.

I. INTRODUCTION

The problem of learning environment models is well known in the domains of robotics and of computer vision. In robotics, the problem is known as simultaneous localization and mapping (SLAM), and it has been deeply investigated in the last decades. The idea of SLAM is to simultaneously estimate the robot position and a map of the environment from the measurements acquired by a robot. A wide variety of approaches has been proposed to solve SLAM and they differ by the underlying estimation technique as well as by the representation of the map [25], [15], [11], [10], [20], [18]. Many state-of-the-art SLAM algorithms seek to find the minimum of a graph-based error function [14], [13], [8], [21], [12]. Every node in the graph represents a robot position, while an edge represents the relative transformation between two robot poses. This relative transformation is usually modeled as a Gaussian distribution and it is determined by matching observations acquired at nearby poses. The task of a SLAM algorithm consists of (simultaneously) constructing the graph from the raw sensor measurements and to find the poses of the nodes that best explain the spatial constraints encoded by the edges. This is usually done by utilizing nonlinear minimization techniques that search in the space of the node configurations.

Traditional laser-based SLAM algorithms typically treat the laser scans as rigid bodies. They calculate the constraints by estimating the relative transformation between adjacent scans once and keeping this fixed during the optimization. The overall map is then obtained by integrating these scans into an occupancy grid map, given the obtained maximum

All authors are with the Autonomous Systems Lab, Department of Computer Science, University of Freiburg. G. Grisetti is also with Sapienza, University of Rome. This work has partly been supported by the EC under contract numbers FP7-231888-EUROPA and FP7-248873-RADHAR.

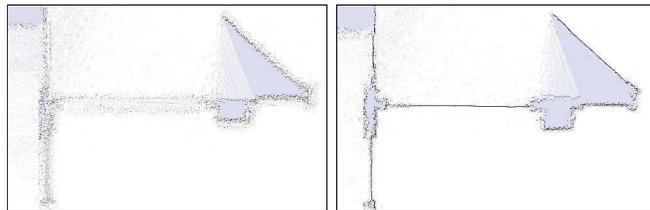


Fig. 1. This figure illustrates the accuracy of the maps achievable by our approach. The left image shows a magnified view of a map obtained with a state-of-the-art traditional on-line SLAM algorithm and the right image shows a magnified view of a map rendered from the result of the joint optimization of the robot poses and the individual scan points.

likelihood robot positions. Whereas this has been proven to yield sufficiently accurate maps, but it does not improve the data associations given the results of the optimization procedure and furthermore does not optimize the scan-points themselves, which may be affected by the sensor noise.

The goal of the approach presented in this paper is jointly to determine the optimal position of the points in the scans and of the robot poses. It relies on the assumption that the laser scanner samples points from a surface that is locally smooth in most cases. It utilizes a dedicated observation model that treats the laser beams as cones and takes into account the incidence angle of a beam to the surface to compute the uncertainty of a measure. Our algorithm performs a full optimization of all the points in the scans and all robot poses based on this refined observation model. Despite the high number of variables involved, the problem can be efficiently solved by using a highly efficient optimization technique for sparse systems.

Figure 1 shows an example illustrating the advantages of our approach compared to previous techniques. We show a particular view of a map estimated with a standard approach and with the method proposed in this paper, which appears to be substantially more accurate. We evaluated the accuracy of the maps generated by our approach both in simulated and real world datasets. Furthermore, we carried on localization experiments using these maps to evaluate the benefit of the increased accuracy in the maps for standard navigation tasks. In experiments we furthermore found out that the resulting maps have a higher resolution than the range scanner used.

II. RELATED WORK

Several approaches have been proposed in the past to characterize the error affecting the result of scan-matching algorithms and thus to accurately estimate the constraints of the optimization problem. Bengtsson *et al.* [1], proposed to analyze the residual of the re-projection error around the

minimum by either sampling the error function or by approximating it by a quadratic form. Subsequently, Censi [5] presented a closed form minimization algorithm to determine this covariance. All these approaches rely on point-to-point correspondences. Olson [23] proposed a hierarchical correlative algorithm derived from [16] that can compute the histogram of the possible robot positions around the minimum. This method does not rely on specific point-to-point correspondences. Instead it uses a histogram approximation of the likelihood function on a grid. Segal *et al.* [24] presented a variant of 3D ICP scan matching that minimizes the matching error between corresponding planar patches extracted from the input scans. Biber *et al.* [3] proposed to solve the scan-matching problem by approximating the reference scan by a set of Gaussians. In this way, the alignment can be computed in closed form given the data association. Compared to the point-to-point criterion minimized by traditional IPC scan matchers [17], [2] using the Gaussians allows to weight the error along different directions based on the shape of the matched surfaces. Subsequently, Magnusson *et al.* [19] extended this approach to 3D.

All these methods focus on accurately registering two scans, but they treat them as rigid bodies. They can be used to construct the pairwise constraints of a graph-based problem. However, they do not attempt to refine the points of the individual scans. In this paper we present an approach to determine the optimal position of both the points in the scans and the robot poses. Our method is similar in spirit to traditional bundle adjustment problems in computer vision [26]. A bundle adjustment algorithm seeks to find the configuration of a set of world points and camera poses that minimizes the re-projection error over the measured sequence of images. The main difference between a bundle-adjustment algorithm and our method is that we do not explicitly rely on point-to-point correspondences as it is done in traditional BA or ICP because a scan point is not a distinguishable physical point as usually in BA but just a sample of the surface. We globally minimize a local plane-to-plane error metric similar to Segal *et al.* [24]. Furthermore, we refine the data associations after every optimization run. This yields maps whose resolution goes beyond the resolution of the sensor.

III. ESTIMATING IMPROVED ENVIRONMENT MODELS BY USING SURFACE-BASED SENSOR MODELS

Our approach aims at obtaining models of the environment that are more accurate than models obtained by considering the plain data provided by the sensor. To this end, we exploit the fact that a range sensor measures the distance to a surface in the environment. The model that we use for our sensor utilizes the fact that in man-made environments the surface is usually smooth and therefore can be approximated in many regions by the corresponding tangent line segment. In the remainder of this paper we will refer to these tangent segments as “surface patches” or simply “patches”.

The overall goal of our approach is to determine the robot positions and the global positions of the surface patches that are maximally consistent. To this end, we construct an

optimization problem that seeks to minimize the distance between corresponding surface patches. This leads to a sparse optimization problem that can be efficiently solved as described in Section III-C. To determine the correspondences in the optimization problem, we account for the local characteristics of the surfaces, like the normal and the tangent direction, thus reducing the number of outliers. This procedure is described in Section III-D.

After every optimization run, we re-estimate the properties of the surfaces based on the new configuration of the points. For efficiency reasons we skip the sensor model update for changed surface point configurations since the changes in direction of the well defined normals are typically small. Subsequently, we construct a new optimization problem by re-computing the correspondences based on the new configuration. We repeat this process until convergence or a maximum number of iterations is reached.

A. Surface Model

We represent a range scan as a set of surface tangents: one for each point in the scan. The local characteristics of the tangents around the sampled points are captured by Gaussian distributions. These Gaussians represent the orientation of the normal and how well the local surface can be represented by a segment.

In our approach, we assume the sensed surface to be piecewise smooth, and that we are able to determine the normals in small regions of this surface. To this end, we determine an approximation of the characteristics of the surface by considering the neighboring surface patches. More precisely, we describe the local characteristics of every point in a scan by a Gaussian. The Gaussian estimates the location μ_{nk} of a surface patch that is measured by the k^{th} noisy range reading r_k that originates from the n^{th} robot pose \mathbf{x}_n . Its covariance is given by considering the scan points patches within a given region around μ_{nk} .

Once the Gaussian is computed, the estimated normal of the surface $\hat{\mathbf{n}}_k$ is the eigenvector of the smallest eigenvalue of the covariance matrix in the direction of the laser scanner. The ratio between the two eigenvalues tells us if the normal is well defined. In case the ratio between the smaller and the bigger eigenvalue is below a given threshold, the points are distributed along a “flat” area, thus the normal is well defined. If this is not the case, we do not compute any local characteristic for that point.

The left picture of Figure 2 illustrates our surface model. To compute it, we start from a laser scan (see left image in Figure 2) and then extract the local characteristics of the surface by grouping neighboring points. This leads to a set of Gaussian distributions that locally approximate the surface.

B. Sensor Model

A laser scanner measures the length of a set of light beams emitted by a laser source. The length of a beam is obtained by measuring the time elapsed between the emission of a laser pulse and the return of its reflection. Due to the optical characteristics of the laser, a beam has a conic shape. For this

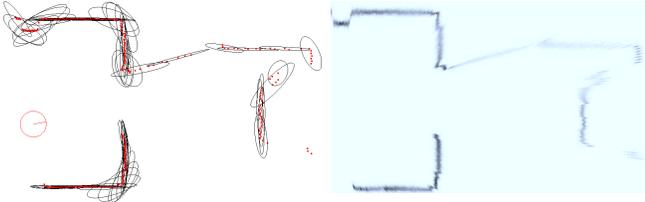


Fig. 2. Evaluation of the likelihood of the points. The left image shows the confidence ellipses representing local planar segments extracted from a scan by applying our surface model. The right image depicts the likelihood of the surface points, given the scan.

reason, a single laser beam does not measure the surface at a specific point, but rather returns some aggregated measure of the distances of the surface within the spot of the beam. This effect is usually negligible when one is interested in low-resolution maps or the robot operates in narrow environments only, but it becomes evident in all other situations.

Furthermore, the incidence angle on the surface usually plays a major role in the error affecting a laser beam. This is another consequence of the conic shape of the beam: the area of the spot on the surface increases with the angle between the normal of the surface and the center of the beam. Since the distance measure is obtained by averaging over a larger region, it tends to be less accurate.

The uncertainty affecting a single laser beam hitting a surface typically depends on the quantization error of the device and the diameter of the spot at a given range. The quantization error η^{quant} is usually uniformly distributed over a small range (e.g., ± 1 cm for the SICK LMS). The diameter d_k of the beam's spot on the surface increases with the length of the measured beam and it depends on the incidence angle α_k to the surface. The spot diameter is proportional to the beam's aperture k_a and to the norm $\|\mathbf{r}_k\|$ of the beam's vector \mathbf{r} . The diameter will then be $k_a\|\mathbf{r}_k\|$, and its projection onto the surface will be

$$d_k \simeq k_a\|\mathbf{r}_k\| \cdot \tan(|\alpha_k|). \quad (1)$$

If we have a rough estimate of the normal $\hat{\mathbf{n}}_k$ from our surface model and we obtain a laser measurement \mathbf{r}_k , the potential measurements will be distributed approximately as a Gaussian oriented along the beam's direction. The mean will lie at the center of the surface patch and the uncertainty is represented by the covariance Σ^{meas} . The standard deviation σ_{11} along the beam's direction is proportional to the projection of the spot d_k on the surface along the beam's direction plus a quantity proportional to the quantization error

$$\sigma_{11} = k_{11}d_k \sin |\alpha_k| + \eta^{\text{quant}}. \quad (2)$$

The standard deviation σ_{22} along the direction orthogonal to \mathbf{r}_k will depend on the diameter of the spot:

$$\sigma_{22} = k_{22}k_a\|\mathbf{r}_k\|. \quad (3)$$

In the above equations the constant factors k_{11} and k_{22} are laser dependent parameters and σ_{11} and σ_{22} are the respective entries in the covariance matrix Σ^{meas} . Figure 3 illustrates how we compute the distribution of the point on a surface that generated a range measurement.

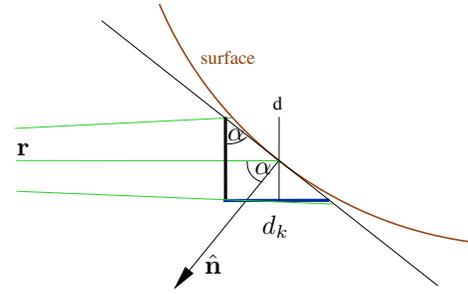


Fig. 3. The sensor model of smooth surfaces. For convenience we dropped the indices. \mathbf{r} : the beam of the laser vector. $\hat{\mathbf{n}}$ the normal of the surface. α : incidence angle. k_a : the aperture of the beam's cone. d : diameter of the beam's spot on the surface \mathbf{r} . σ_{11} : standard deviation along the beam's direction (not shown) and σ_{22} : standard deviation along the beam's tangent direction (not shown).

C. Objective Function and Optimization

The goal of our approach is to determine the map and the configuration of robot poses that is maximally consistent with the measurement. The map \mathcal{M} is represented as a set of tangent segments, or points where the tangent is not defined. Each tangent is described by a Gaussian $\langle \mu_{nk}, \Sigma_{nk} \rangle$, where μ_{nk} denotes the k^{th} beam of the n^{th} robot pose \mathbf{x}_n , as described in Section III-A. The robot positions $\mathbf{x}_{1:N}$ are constrained by the odometry measurements $\mathbf{u}_{1:N}$.

If we know the pairwise correspondences between two surface patches extracted from different scans, we can define the error vector for this correspondence as the difference between the centroids of the ellipsoids representing the tangents:

$$\mathbf{e}_{ij}^{\text{surf}}(\mu_{n_i k_i}, \mu_{n_j k_j}) = \mu_{n_i k_i} - \mu_{n_j k_j} = \Delta \mu_{ij}. \quad (4)$$

When constructing a quadratic optimization problem, we need to weigh the errors between the surface patches according to the orientation of the surfaces. More specifically we want to allow the surfaces to "slide" along the tangential direction, but we want them to be more rigid along the direction of their normal. This can be effectively accounted by the sum of the inverses of the covariance matrices of the ellipsoids.

$$\Omega_{ij} = \Sigma_{n_i k_i}^{-1} + \Sigma_{n_j k_j}^{-1}. \quad (5)$$

The quadratic error introduced by this correspondence will then be:

$$\mathbf{e}_{ij}^{\text{surf}} = \Delta \mu_{ij}^\top \Omega_{ij} \Delta \mu_{ij}. \quad (6)$$

Figure 4 illustrates the calculation of this component of our error function.

Each surface patch $\langle \mu_{nk}, \Sigma_{nk} \rangle$ is connected to the laser pose by a measured laser beam \mathbf{r}_{nk} . This error is distributed according to the covariance Σ^{meas} computed in Section III-B:

$$\mathbf{e}_{nk}^{\text{meas}} = ((\mu_{nk} \ominus \mathbf{x}_n) - \mathbf{r}_{nk})^\top (\Sigma_{nk}^{\text{meas}})^{-1} ((\mu_{nk} \ominus \mathbf{x}_n) - \mathbf{r}_{nk}). \quad (7)$$

An odometry measurement \mathbf{u}_n between two consecutive robot poses \mathbf{x}_n and \mathbf{x}_{n+1} contributes to the error function by the following term:

$$\mathbf{e}_n^{\text{od}} = (\mathbf{u}_n \ominus (\mathbf{x}_{n+1} \ominus \mathbf{x}_n))^\top \Sigma_n^{-1} (\mathbf{u}_n \ominus (\mathbf{x}_{n+1} \ominus \mathbf{x}_n)). \quad (8)$$

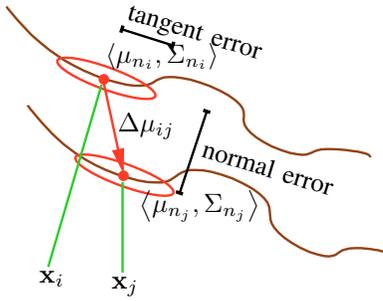


Fig. 4. This figure illustrates the computation of the error function between two corresponding surface patches extracted from two different scans. \mathbf{x}_{n_i} and \mathbf{x}_{n_j} are the two positions of the laser scans. $\mu_{n_i k_i}, \mu_{n_j k_j}$ are the centroids of the surface patches. $\Sigma_{n_i k_i}^{-1}$ and $\Sigma_{n_j k_j}^{-1}$ are the covariance matrices of the ellipses and $\Delta \mu_{ij}$ is the error vector.

Here, \ominus is the usual inverse motion composition operator as described in [25] and Σ_n^{-1} is the covariance matrix that captures the uncertainty of the odometry.

Assuming a known set of correspondences, we can set up a least squares minimization problem that seeks to find the configuration of robot poses $\mathbf{x}_{1:n}^*$ and surface patches \mathcal{M}^* as:

$$\langle \mathbf{x}_{1:n}^*, \mathcal{M}^* \rangle = \underset{\mathbf{x}_{1:n}, \mathcal{M}}{\operatorname{argmin}} \sum_{n=1}^N \mathbf{e}_n^{\text{od}} + \sum_{\langle i,j \rangle} \mathbf{e}_{ij}^{\text{surf}} + \sum_{\langle n,k \rangle} \mathbf{e}_{nk}^{\text{meas}}. \quad (9)$$

To find this minimum we utilize the Gauss-Newton algorithm. Whereas the number of variables to optimize is typically large (in the order of one million of elements), the resulting linear system is typically sparse. The objective function is the sum of factors involving only pairs of state variables. Thus the approximated Hessian contains a number of non-zero entries that is proportional to the number of constraints. Since the range of the sensor is limited, this results in a sparse approximated Hessian. We can achieve the desired performance by solving the linear system by sparse Cholesky decomposition using the CHOLMOD algorithm [7]. Given this algorithm, our current system can perform one iteration of non-linear optimization of a system consisting of 172,522 surface patches acquired from 616 robot positions, 996,451 surface constraints, and 671,550 constraints between surface patches in less than 5 seconds using one core of a Core Quad running at 2.6 Ghz. Figure 5 shows the typical non-zero pattern of the sparse Hessian.

In the Gauss-Newton algorithm we do not directly optimize the covariances of the patches $\Sigma_{n_i k_i}$, since we internally store them relative to the robot position from where they have been acquired. In this way, when a robot position is updated, we implicitly rotate the covariances of the patches that have been seen from that position. However, the error functions for the landmarks (Eq 6) require these covariances to be expressed in the global reference frame to compute the information matrices Ω_{ij} . We carry on this operation in the linearization step of the Gauss-Newton algorithm.

D. Data Association

In the previous section we assumed the correspondences between the surface patches to be known. This assumption is

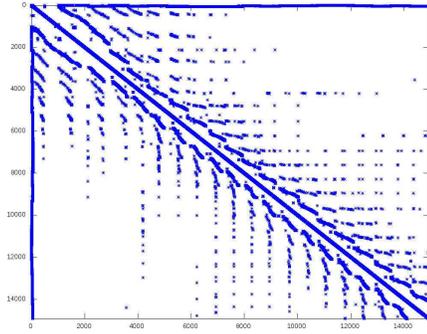


Fig. 5. Non-zero entries of the approximated Hessian of your optimization procedure. The “dense” bands on the top and on the right correspond to constraints between the robot poses and the patches.

obviously not true in reality. In this section we will explain how we determine the potential correspondence between two patches in different scans.

Given an initial configuration of two patches $\langle \mu_i, \Sigma_i \rangle$ and $\langle \mu_j, \Sigma_j \rangle$, where the normals are well defined, we utilize the “normal-shooting” heuristic proposed in Chen *et al.* [6]. The idea is to consider every surface patch where the normal is well defined and to search along its normal direction for the closest patch of another scan whose ellipsoid has a similar shape as well as a similar orientation of the normal. If such a patch is found, we add a constraint between them. Whenever two robot poses are connected by an odometry constraint, we apply this heuristic to introduce a constraint between the surface patches of the scans. Furthermore, we consider the surface patches that are closer than 0.2 m to each other. If a surface patch has more than one neighbor, we only add a constraint to the one having the smallest index. By doing so, we enforce a high degree of sparsity of the Hessian without a substantial decrease in the final map accuracy. The data association is updated after every optimization run. We consider a system converged if changes in the χ^2 error are below a given threshold for at least 5 optimization runs.

IV. EXPERIMENTS

In this section, we present experiments carried out to evaluate the performances of our approach. Throughout our experiments we used the Freiburg indoor building 079, the Intel Research Lab, the MIT CSAIL Building, and the ACES Building data sets. We choose those datasets because they are publicly available and well known in the SLAM community. Another important advantage of this dataset collection is the fact that they were acquired with different laser sensors and give an intuition of the generality of our sensor model. For further quantitative evaluations we furthermore used a simulated map resembling the Freiburg indoor building 079 map.

A. Entropy on Real World Data

In the first experiment we evaluated the impact of the combined pose and observation optimization on a set of real world datasets. The main purpose of our method is to produce accurate maps, which typically results in a low entropy of the corresponding occupancy grid map, which

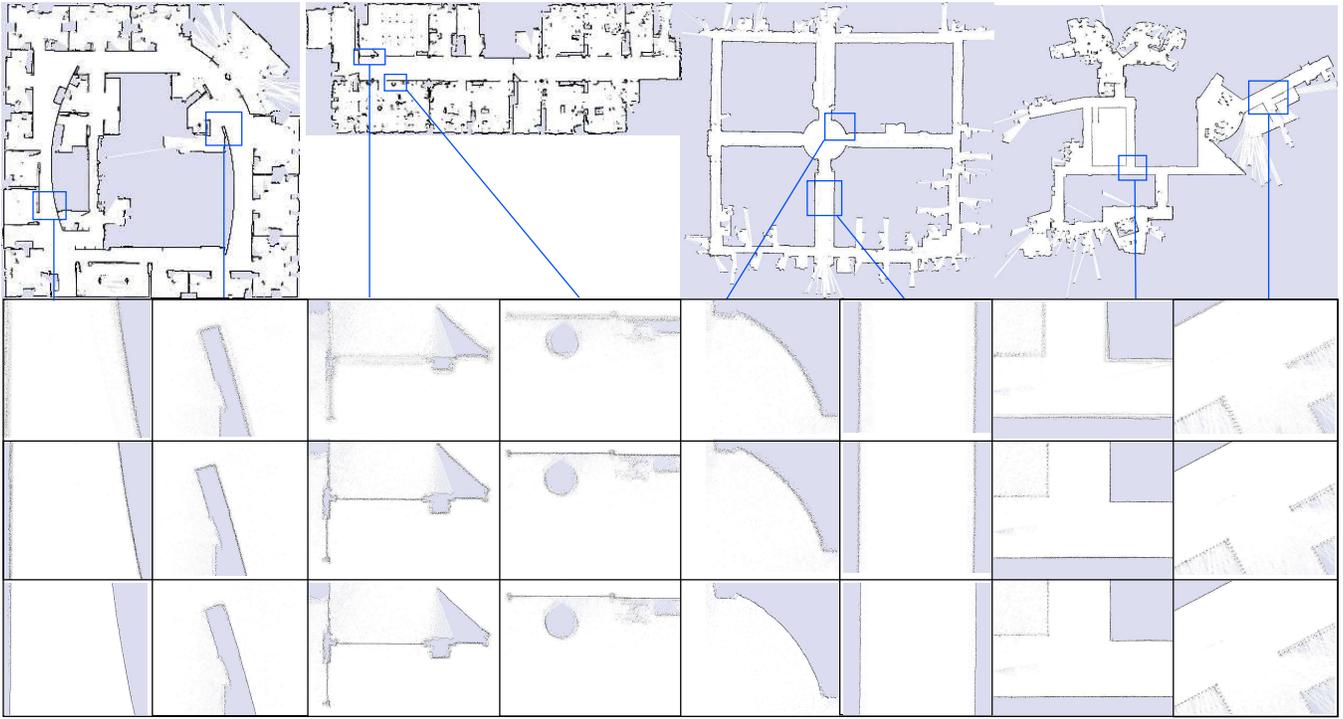


Fig. 6. The top row of this figure shows the resulting maps of our method. We enlarged some regions to visualize the differences in the SLAM results used as input (second row), the pose-only optimized result (third row), and the output of our method (bottom row). The bottom row shows only few blur and sharp borders between free and occupied cells.

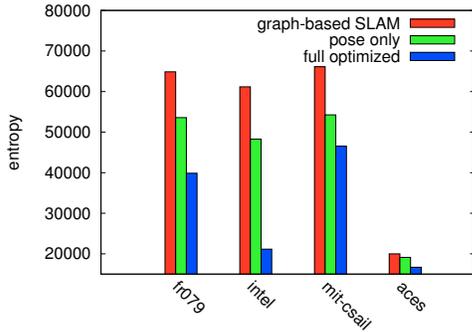


Fig. 7. This figure shows the computed entropy for the first experiment. The entropy of the optimized result is lower than the entropy of the input data on all maps.

essentially measures the uncertainty about the state of the environment. Our approach is not intended as a solution for the entire SLAM problem but rather serves as an post-optimization step that maximizes the accuracy of the map. In this experiment, we therefore computed a standard graph-based SLAM solution for each dataset and applied our method to the resulting estimates. The occupancy maps generated based on the estimates of the standard SLAM algorithm serves as a baseline for the quality of the maps obtained with our approach. To analyze the impact of the laser point optimization we additionally generated maps using only the robot poses estimated by our approach but generating the maps from the original scans. We generated grid maps of 5 mm resolution for both methods as well as our approach and analyzed the resulting maps visually and numerically by measuring the entropy. The first row of

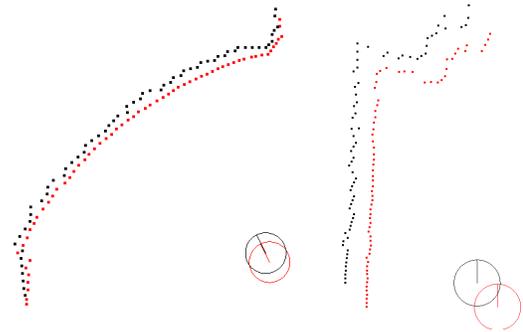


Fig. 8. Two partial scans of the ACES Building dataset before optimization (black) and after optimization (red/gray). In both cases the black points show typical sensor errors. In contrast the red/gray points give a more likely explanation of the surface structure.

Figure 6 shows the resulting maps. The second row shows the SLAM results used as input, the third row shows the same part of the map for the pose-only optimization case and the bottom row shows the same part of the map generated by our method. As can be seen from the figure, our approach yields maps showing the highest accuracy compared to both other approaches. The entropy of the resulting is displayed in Figure 7. As can be seen, our approach yields a substantially lower entropy, which is consistent with the visual inspection the maps.

The lower entropy comes from the fact that the optimization of laser points reduces the impact of sensor noise. Figure 8 illustrates the differences before and after optimization on two partial laser readings taken from the ACES Building dataset. The optimized points (red) give a visually better

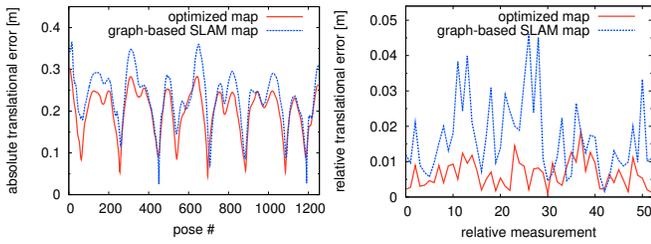


Fig. 9. Localization error on a graph-based SLAM (blue-dotted) and an optimized map (red) in the simulator experiment. The left plot shows the absolute errors for all poses computed during the evaluation of the RMS. It gives an upper bound for the global error of the underlying map. The right plot shows the relative errors between pose pairs that were selected sequentially on the trajectory with a distance of at least 1 m.

description of the observed structure than the original sensor reading (black).

B. Localization and Map Accuracy

Since we lack precise ground truth for the real world datasets we evaluated our approach in a simulation scenario. To evaluate this, we measured the root of mean squared error (RMS) as described in the work of Olson *et al.* [22]. Additionally, we computed the relative error between localization pose pairs as additional measure. We refer to this measure as Relative Pose Error (RPE) in the following [4].

As first step of the experiment we manually constructed a map with a cell size of 5 mm and recorded two datasets on this map using a simulator. From the first dataset we constructed a map with graph-based SLAM and a post-optimized map with our method. For both maps we ran a standard Monte-Carlo-Localization (MCL) [9] given the second dataset and evaluated the corresponding localization accuracy. The left plot in Figure 9 shows the absolute error for all poses calculated during the computation of the RMS. This error plot represents the accumulated map and localization error and gives an upper bound to the global error in the map. The RMS of the SLAM result is 0.0636 m and the RMS of the optimized map is 0.0469 m. The right plot in Figure 9 shows the RPE of localization pose pairs that were selected sequentially on the robot trajectory and have a distance of at least 1 m to each other. Both plots demonstrate that the localization gets substantially better on a map optimized with our approach than on a standard SLAM map.

V. CONCLUSIONS

In this paper we described an approach that improves the results of traditional SLAM approaches by jointly optimizing robot poses and laser data. Our approach is based on the assumption that a laser range finder senses points from a regular surface and utilizes two phenomena of laser measurements: the conic shape of the beam and the influence of the incidence angle. Experimental results demonstrate that our approach yields substantial improvements for maps obtained with a standard SLAM technique. The software and all data sets are available at <http://www.openslam.org>.

REFERENCES

- [1] O. Bengtsson and A.-J. Baereldt, "Robot localization based on scan-matching – estimating the covariance of for the icp algorithm," *Journal of Robotics & Autonomous Systems*, vol. 44, pp. 29–40, 2003.
- [2] P. J. Besl and N. D. McKay, "A method for registration of 3-d shapes," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 14, no. 2, pp. 239–256, 1992.
- [3] P. Biber and W. Strasser, "The normal distributions transform: A new approach to laser scan-matching," in *Proc. of the Int. Conf. on Intelligent Robots and Systems (IROS)*, 2003.
- [4] W. Burgard, C. Stachniss, G. Grisetti, B. Steder, R. Kümmerle, C. Dornhege, M. Ruhnke, A. Kleiner, and J. D. Tardós, "A comparison of SLAM algorithms based on a graph of relations," in *Proc. of the Int. Conf. on Intelligent Robots and Systems (IROS)*, 2009.
- [5] A. Censi, "The achievable accuracy for range finder localization," *IEEE Trans. on Robotics*, 2009.
- [6] Y. Chen and G. Medioni, "Object modeling by registration of multiple range images," in *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 1991.
- [7] Y. Chen, T. A. Davis, W. W. Hager, and S. Rajamanickam, "Algorithm 887: Cholmod, supernodal sparse cholesky factorization and update/downdate," *ACM Trans. Math. Softw.*, vol. 35, no. 3, pp. 1–14, 2008.
- [8] F. Dellaert, "Square Root SAM," in *Proc. of Robotics: Science and Systems (RSS)*, Cambridge, MA, USA, 2005, pp. 177–184.
- [9] 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)*, 1999.
- [10] G. Dissanayake, S. Williams, H. Durrant-Whyte, and T. Bailey, "A computationally efficient solution to the simultaneous localisation and map building (SLAM) problem," *Journal of Autonomous Robots*, 2000.
- [11] A. Doucet, J. de Freitas, K. Murphy, and S. Russell, "Rao-Blackwellized particle filtering for dynamic bayesian networks," in *Proc. of the Conf. on Uncertainty in Artificial Intelligence (UAI)*, Stanford, CA, USA, 2000, pp. 176–183.
- [12] U. Frese and T. Duckett, "A multigrid approach for accelerating relaxation-based SLAM," in *Proc. of the IJCAI Workshop Reasoning with Uncertainty in Robotics*, Acapulco, Mexico, 2003.
- [13] G. Grisetti, R. Kümmerle, C. Stachniss, U. Frese, and C. Hertzberg, "Hierarchical optimization on manifolds for online 2d and 3d mapping," in *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2010.
- [14] G. Grisetti, C. Stachniss, and W. Burgard, "Non-linear constraint network optimization for efficient map learning," *IEEE Transactions on Intelligent Transportation Systems*, 2009.
- [15] K. Konolige, "A gradient method for realtime robot control," in *Proc. of the Int. Conf. on Intelligent Robots and Systems (IROS)*, 2000.
- [16] K. Konolige and K. Chou, "Markov localization using correlation," in *Proc. of the Int. Conf. on Artificial Intelligence (IJCAI)*, 1999.
- [17] F. Lu and E. Milios, "Robot pose estimation in unknown environments by matching 2d range scans," in *IEEE Computer Vision and Pattern Recognition Conference (CVPR)*, 1994, pp. 935–938.
- [18] —, "Globally consistent range scan alignment for environment mapping," *Autonomous Robots*, vol. 4, no. 4, pp. 333–349, 1997.
- [19] M. Magnusson, T. Duckett, and A. J. Lilienthal, "3d scan registration for autonomous mining vehicles," *Journal on Field Robotics*, vol. 24, no. 10, pp. 803–827, 2007.
- [20] A. Nüchter, K. Lingemann, J. Hertzberg, and H. Surmann, "6d SLAM with approximate data association," in *Proc. of the 12th Int. Conference on Advanced Robotics (ICAR)*, 2005, pp. 242–249.
- [21] E. Olson, J. Leonard, and S. Teller, "Fast iterative optimization of pose graphs with poor initial estimates," in *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2006, pp. 2262–2269.
- [22] E. Olson and M. Kaess, "Evaluating the performance of map optimization algorithms," in *RSS Workshop on Good Experimental Methodology in Robotics*, 2009.
- [23] E. B. Olson, "Real-time correlative scan matching," in *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2009.
- [24] A. V. Segal, D. Haehnel, and S. Thrun, "Generalized-ICP," in *Proc. of Robotics: Science and Systems (RSS)*, 2009.
- [25] R. Smith, M. Self, and P. Cheeseman, "Estimating uncertain spatial relationships in robotics." in *Autonomous Robot Vehicles*, I. Cox and G. Wilfong, Eds. Springer Verlag, 1990, pp. 167–193.
- [26] B. Triggs, P. Mclauchlan, H. Richard, and A. Fitzgibbon, "Bundle adjustment – a modern synthesis," 2000.