

6DOF Semi-Rigid SLAM for Mobile Scanning

Jan Elseberg, Dorit Borrmann, and Andreas Nüchter

Abstract—The terrestrial acquisition of 3D point clouds by laser range finders has recently moved to mobile platforms. Measuring the environment while simultaneously moving the vehicle demands a high level of accuracy from positioning systems such as the IMU, GPS and odometry. We present a novel semi-rigid SLAM algorithm that corrects the global position of the vehicle at every point in time, while simultaneously improving the quality and accuracy of the entire acquired map. Using the algorithm the temporary failure of positioning systems or the lack thereof can be compensated for. We demonstrate the capabilities of our approach on a wide variety of systems and data sets.

I. INTRODUCTIONS

Mobile laser scanning provides an efficient way to actively acquire accurate and dense 3D point clouds of object surfaces or environments and are currently used for modeling in architecture as well as urban and regional planning. Modern systems like the Riegl VMX-450 and the Lynx Mobile Mapper by Optech combine a high precision GPS, a highly accurate Inertial Measurement Unit (IMU) and the odometry of the vehicle to compute the fully timestamped trajectory. Using motion compensation this trajectory is then used to “unwind” the laser range measurements that were acquired by the 2D laser scanner mounted on the vehicle. The precision of the resulting point cloud depends on several factors:

- the precision of the calibration of the system;
- the accuracy of the positioning sensors, i.e., the GPS, IMU and odometry;

J. Elseberg, D. Borrmann, and A. Nüchter are with the Automation and Machine Vision Group at the Jacobs University Bremen gGmbH, 28759 Bremen, Germany {j.elseberg|d.borrmann|a.nuechter}@jacobs-university.de. The authors would like to thank Vaibhav Bajpai, Vitali Bashko, Mohammad Faisal, Markus Harz, Bojan Koccev, and Vladislav Perelman for their contribution for setting up the car, and TopScan GmbH for providing the Dortmund data set. This work was partially supported by the AUTOSCAN project (BMW KF2470003DF0).

- the availability of the GPS, as it may suffer temporary blackouts in tunnels, between high-rises etc.;
- the precision of the laser scanner itself.

The deficiency of the GPS in constricted spaces is an exceptionally limiting factor for mobile scanning. Without the dependance on a GPS device, applications such as construction of tunnels and mines and facility management would be open to the use of mobile scanning systems.

In this paper we propose an algorithm to overcome the problem of erroneous positioning of the vehicle in a novel fashion. We present an algorithm for computing a corrected trajectory that produces point clouds that are a more accurate representation of the measured environment. To this purpose we constructed mobile platforms that use a modified approach to mobile laser scanning. Instead of the common 2D laser scanner in a mobile scanning system, we employ a 3D laser scanner as depicted in Fig. 2. This design retains the high degree of automation and speed of common mobile laser scanning systems. Furthermore, only a single laser range sensor is required to produce models with minimal data shadowing. The additional rotation with respect to the platform during the scanning process reveals the errors caused by incorrect positioning to a higher degree than the common design and also allows us to correct them.

We evaluate our algorithms on data sets acquired by our own mobile platform Irma3D as well as on data acquired on a car with only minimal means of pose estimation.

II. RELATED WORK

The area of non-rigid registration is largely unexplored but in the medical imaging community where it is widespread due to the need to align multi-modal data [6], [10].

Williams et al. [14] describe an extension of a rigid registration algorithm that includes point estimation to

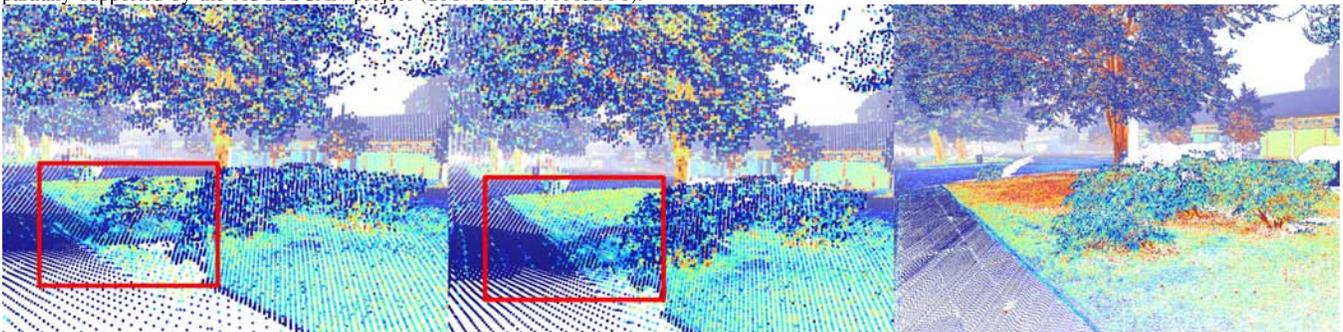


Fig. 1: A part of the campus scanned by a mobile laser scanner before and after our novel optimization. Note the erroneously duplicated vegetation in the initial point cloud (left) which is corrected in the optimized point cloud (middle), so that it now appears much closer to the ground truth acquired by terrestrial laserscanning (right). Points are colored by reflectance information.



Fig. 2: Two mobile laser scanning systems used in our experiments equipped with a Riegl VZ-400 3D laser scanner and an Xsens MTi IMU. The Riegl completes a rotation around the vertical axis every 6 seconds and acquires 750k points per rotation with a vertical opening angle of 100 degrees. Top: The robot Irma3D receives odometry estimates from the volksbot motor encoders. Bottom: No odometry for the car is available. An On Board Diagnosis device (OBDII) acquires speed estimates with a precision of at most $1 \frac{km}{h}$.

compensate for noisy sensor data. This technically constitutes a non-rigid registration algorithm designed for low scale high frequency deformations. Similarly, Pitzer et al. [12] provide a probabilistic mapping algorithm for 2D range scans, where point measurements are also estimated.

Chui and Rangarajan [5] proposed a point matching algorithm that is capable of aligning point clouds with each other. These approaches are usually time expensive due to the enlarged state space.

Brown and Rusinkiewicz developed a global non-rigid registration procedure [4]. They introduced a novel Iterative Closest Point (ICP) variant to find correspondences. Though the registration requires extreme subsampling the deformation is successfully generalized onto the entire scan. Unfortunately, this technique is not fully applicable to laser scans acquired by mobile robots [8].

Stoyanov et al. [13] presented a non rigid optimization for a mobile laser scanning system. They optimize point cloud quality by matching the beginning and the end of a single scanner rotation using ICP. The estimate of the 3D pose difference is then used to optimize the robot trajectory in between. Similarly, Bosse and Zlot [3] use a modified ICP with a custom correspondence search to optimize the pose of 6 discrete points in time of the trajectory of a robot during a single scan rotation. The trajectory in between is modified by distributing the errors with a cubic spline.

The approach presented in this paper optimizes the point cloud using full 6D poses and is not limited to a single scanner rotation. We improve scan quality globally in all 6 degrees of freedom for the entire trajectory.

III. SEMI-RIGID OPTIMIZATION FOR SLAM

We developed an algorithm that improves the entire trajectory of the vehicle simultaneously, and is unlike previous

algorithms [3], [13] not restricted to purely local improvements. We make no rigidity assumptions, except for the computation of the point correspondences, require no explicit vehicle motion model, although such information may be incorporated at no additional cost and the algorithm requires no high-level feature computation, i.e., we require only the points themselves.

In addition to one or multiple 2D or 3D laser scanners mobile laser scanning systems consist of a mobile base equipped with sensors for trajectory estimation such as IMU and GPS in addition to the odometry of the vehicle.

The movement of the vehicle between time t_0 and t_n creates a trajectory $S = \{V_0, \dots, V_n\}$, where $V_i = (T_x, T_y, T_z, \theta_x, \theta_y, \theta_z)$ is the 6DOF pose of the vehicle at time t_i with $t_0 \leq t_i \leq t_n$. Using the trajectory of the vehicle a 3D representation of the environment can be obtained with the laser measurements M to create the final map P . However, sensor errors in odometry, IMU and GPS as well as systematic calibration errors and the accumulation of pose errors during temporary GPS outages degrade the accuracy of the trajectory and therefore the point cloud quality.

The current state of the art for improving overall map quality of mobile mappers in the robotics community is to coarsely discretize the time. This results in a partition of the trajectory into subscans that are treated rigidly. Then rigid registration algorithms like the ICP and other solutions to the SLAM problem are employed. Obviously, trajectory errors within a subscan cannot be improved in this fashion. Applying rigid pose estimation to this non-rigid problem is also erroneous because it only approximates the solution. Consequently, overall map quality suffers as a result.

We employ a much finer discretization of the time, at the level of individual 2D scan slices or individual points. This results in the set of measurements $M = \{m_0, \dots, m_n\}$ where $m_i = (m_{x,i}, m_{y,i}, m_{z,i})$ is a point acquired at time t_i in the local coordinate system of V_i . Applying the pose transformation \oplus we derive the point $p_i = V_i \oplus m_i$ in the global coordinate frame and thereby also the map $P = \{p_0, \dots, p_n\}$. Given M and S we find a new trajectory $S' = \{V'_1, \dots, V'_n\}$ with modified poses so that P generated via S' more closely resembles the real environment.

A. Pose Estimation

Our algorithm incorporates pose estimations from many sources, such as odometry, IMU and GPS. For the results in this paper we use sequential pose estimates $V_{i,i+1}$ that were combined from all those sources:

$$\bar{V}_{i,i+1} = V_i \ominus V_j. \quad (1)$$

using a constant covariance $C_{i,i+1}$. In addition to these default pose estimates, that may also be enhanced by separating all pose sensors into their own estimates as well as the proper covariances, we estimate differences between poses via the point cloud P .

We employ nearest neighbor search sped up by a fast and compact representation of the entire point cloud into a single octree. This allows us to compute the nearest neighbor for a

point at comparable speed to state of the art k -d tree search libraries. The octree also compresses the point cloud so it can be easily stored and processed. For each measurement \mathbf{p}_i , we find a closest measurement \mathbf{p}_j with $|i - j| > \delta$, where δ is the minimal amount of time that must have elapsed for the laser scanner to have measured the same point on the surface again. Temporally close measurements are usually spatially close as well, so they must be excluded since they cannot represent the same physical surface. We employ the nearest neighbor search algorithm as described in [9] that has been modified to enforce the time constraint.

Points are stored in the global coordinate frame as defined by the estimated trajectory S . Closest points are accepted if $\|\mathbf{p}_i - \mathbf{p}_j\| \leq d_{\max}$. The positional error of two poses \mathbf{V}_i and \mathbf{V}_j is then given by

$$E_{i,j}(\mathbf{V}) = \sum_{k=i-m}^{i+m} \|\mathbf{V}_i \oplus \mathbf{m}_k - \mathbf{V}_j \oplus \mathbf{m}'_k\|^2 \quad (2)$$

Here, \mathbf{V} is the concatenation of all poses, $\mathbf{m}_k, \mathbf{m}'_k$ is the pair of closest points written in their respective local coordinate system, and m defines a small neighborhood of points taken in the order of hundreds of milliseconds that is assumed to have negligible pose error. After a Taylor expansion of $E_{i,j}$ with respect to \mathbf{V} , we obtain its minimum and the corresponding covariance by

$$\bar{\mathbf{V}}_{i,j} = (\mathbf{M}^T \mathbf{M})^{-1} \mathbf{M}^T \mathbf{Z} \quad (3)$$

$$\mathbf{C}_{j,k} = s^2 (\mathbf{M}^T \mathbf{M}). \quad (4)$$

Here \mathbf{Z} is the concatenation of all $\bar{\mathbf{V}}_i \oplus \mathbf{m}_k - \bar{\mathbf{V}}_j \oplus \mathbf{m}_k$ into a vector and \mathbf{M} is a concatenation of:

$$\mathbf{M}_i = \begin{pmatrix} 1 & 0 & 0 & 0 & -m_{y,i} & -m_{z,i} \\ 0 & 1 & 0 & m_{z,i} & m_{x,i} & 0 \\ 0 & 0 & 1 & -m_{y,i} & 0 & m_{x,i} \end{pmatrix}. \quad (5)$$

The unbiased estimate of the covariance of the independent and identically distributed errors of $\mathbf{E}_{i,j}$ is computed as

$$s^2 = (\mathbf{Z} - \mathbf{M} \bar{\mathbf{V}}_{i,j})^T (\mathbf{Z} - \mathbf{M} \bar{\mathbf{V}}_{i,j}) / (2m - 3). \quad (6)$$

B. Pose Optimization

We maximize the likelihood of all pose estimates and their respective covariances via the Mahalanobis distance

$$W = \sum_i \sum_j (\bar{\mathbf{V}}_{i,j} - (\mathbf{V}'_i - \mathbf{V}'_j)) \mathbf{C}_{i,j}^{-1} (\bar{\mathbf{V}}_{i,j} - (\mathbf{V}'_i - \mathbf{V}'_j)),$$

or, with the incidence matrix \mathbf{H} in matrix notation:

$$W(\mathbf{V}) = (\bar{\mathbf{V}} - \mathbf{H}\mathbf{V})^T \mathbf{C}^{-1} (\bar{\mathbf{V}} - \mathbf{H}\mathbf{V}). \quad (7)$$

The minimization of W is accomplished via solving the following linear equation system:

$$(\mathbf{H}^T \mathbf{C}^{-1} \mathbf{H}) \mathbf{V} = \mathbf{H}^T \mathbf{C}^{-1} \bar{\mathbf{V}}. \quad (8)$$

Computing the optimized trajectory is then reduced to inverting a positive definite matrix [11] that is sparse due to a large number of empty correspondences. Thus, we make use of the sparse Cholesky decompositions by Davis [7].

The complete semi-rigid registration algorithm proceeds as follows: Given a trajectory estimate, we compute the

point cloud P in the global coordinate system and create our octree for fast nearest neighbor search. Then, after computing the estimates $\bar{\mathbf{V}}_{i,j}$ of pose differences and their respective covariances $\mathbf{C}_{i,j}$ we optimize the trajectory S . The process is iterated until convergence, i.e., until the change in the trajectory falls below a threshold. Establishing point correspondences is the most time consuming step in the process with $O(n \log n)$, where n is the number of points. To deal with massive amount of data in a reasonable time frame, we employ 2 strategies. First, we uniformly and randomly reduce the point cloud by using only a constant number of points per volume, typically to 1 point per 3 cm^3 . The octree is ideally suited for this type of subsampling. Second, in initial stages of the algorithm estimates $\bar{\mathbf{V}}_{i,j}$ are only computed for a subset of poses $\mathbf{V}_0, \mathbf{V}_m, \mathbf{V}_{2m}, \dots$, with m in the order of hundreds of milliseconds. In every iteration m is decreased so that the trajectory can be optimized on a finer scale.

IV. EXPERIMENTS AND RESULTS

We demonstrate the efficacy of the presented algorithms on two data sets acquired by two different robotic systems. In addition we also evaluate our algorithms in an indoor data set where ground truth data is available.

A. Indoor data set

For a direct evaluation of the algorithm with comparison to ground truth data we acquired a data set using the robot Irma3D in an empty basement room (see Fig. 2). Ground truth data of the room is available in the form of a geodetically measured model acquired by a Riegl VZ-400 using terrestrial laser scanning. The accuracy of the scanner and therefore the model is 5 mm.

This data set with about 4 million points was acquired in continuous mode, i.e., the laser scanner rotates around its vertical axis while the robot moves simultaneously. The robot moved in a ‘‘serpentine’’ trajectory several meters in length, i.e., taking left and right turns and segments where the heading remains unchanged. To evaluate the quality of the resulting point cloud we compare it to the high precision ground truth model of the room. The point cloud is matched to the model using ICP [2] from the 3D Toolkit (3DTK [1]). Then we compute point to plane distances on the ceiling, floor and each of the 4 walls.

A qualitative comparison of the results of our algorithm is presented in Fig. 3. The point clouds obtained with Irma3D in the enclosed room are shown in Fig. 3. The figure presents the initial point cloud, the rigid registration obtained with ICP, the semi-rigid registration computed by the novel algorithm and the ground truth data. The results of the direct comparison between the the point clouds before and after automatic semi-rigid registration and the model of the room are shown in Fig. 4. The deviations between model and point cloud are plotted in color coded images, i.e., green for absolute errors less than 1 cm, yellow to red for large positive errors and cyan to blue for large negative

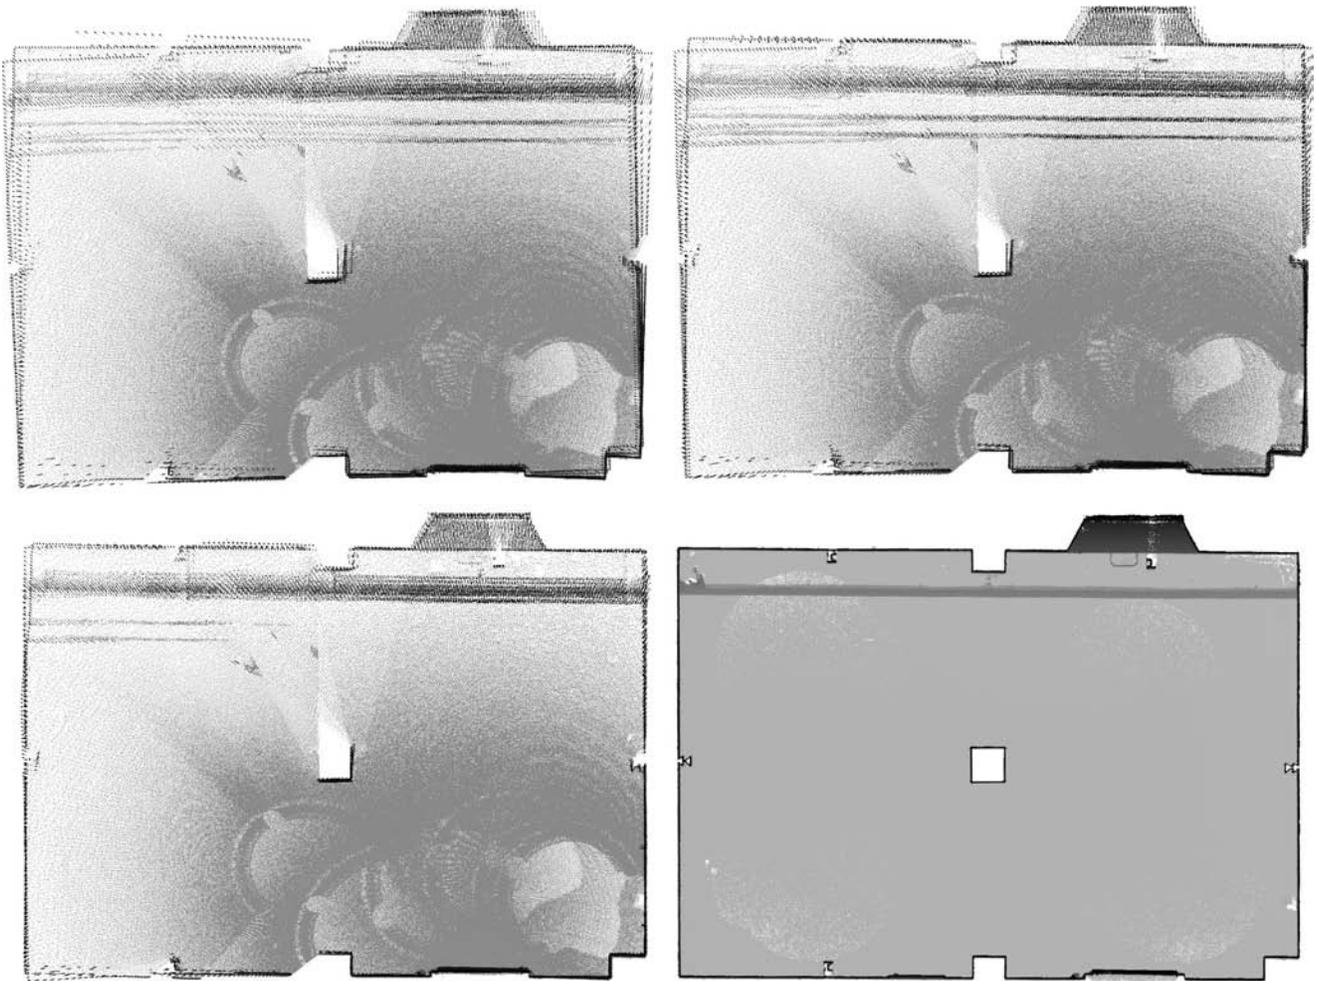


Fig. 3: An overview on the room data set used for quantitative evaluation of the semi-rigid registration. Top: The initial data set with no registration (left) and with rigid registration via ICP&SLAM (right). Bottom left: The result of the novel semi rigid registration procedure. Bottom right: The model of the room acquired with an absolute precision of 5mm.

errors. White areas indicate that no point was measured at the corresponding location.

The proposed semi-rigid registration produces point clouds that more closely resemble the ground truth model. Although rigid registration procedures can improve map quality, it is clear that the error within one scanner rotation is too large to allow for a good match. The reduction of the inner error is even more obvious in the quantitative evaluation as seen in Fig. 4. The left side shows the point-to-model error of each of the subscans before semi-rigid registration. The right column displays the error after the data set has been corrected. Apart from some small regions in the top scan, the errors are significantly reduced by the semi-rigid registration.

B. Outdoor data set

To further prove the viability of the proposed algorithm in other environments we acquired a second data set using the car in Fig. 2. The vehicle traveled alongside the campus buildings, took a U-turn, traveled straight back on the other side of the road and took another U-turn to return to the beginning. Due to the large inaccuracies in the initial data set, a ground truth comparison similar to the one before is not applicable. Nonetheless we acquired a ground truth data

set of the campus in the same fashion as before to provide a point of reference.

The second data set is a larger one with 12 million points. It is also far more erroneous since no odometry of wheel encoder measurements was available. Indeed, the error is so large that any attempts to improve upon the map quality by rigid registration methods failed. The initial and the semi-rigidly corrected point cloud as well as the ground truth model of the campus is shown in Fig. 5. The semi-rigid registration shows a remarkable improvement on the input data. The resulting point cloud exhibits no sign of error accumulation and reflects the ground truth quite well.

C. Further Experiments

We also tested the algorithm on the Dortmund data set acquired by TopScan GmbH via the Lynx Mobile Mapper by Optech. The Lynx mobile mapper employs twin 2D laser scanners and pose estimation by the Applanix POS L with integrated odometry, IMU and a GPS unit. The video supplement to this paper can be seen in high quality under <http://youtu.be/L28C2YmUPWA>.

Processing time for all data sets presented in this paper is in the order of several minutes on a consumer laptop.

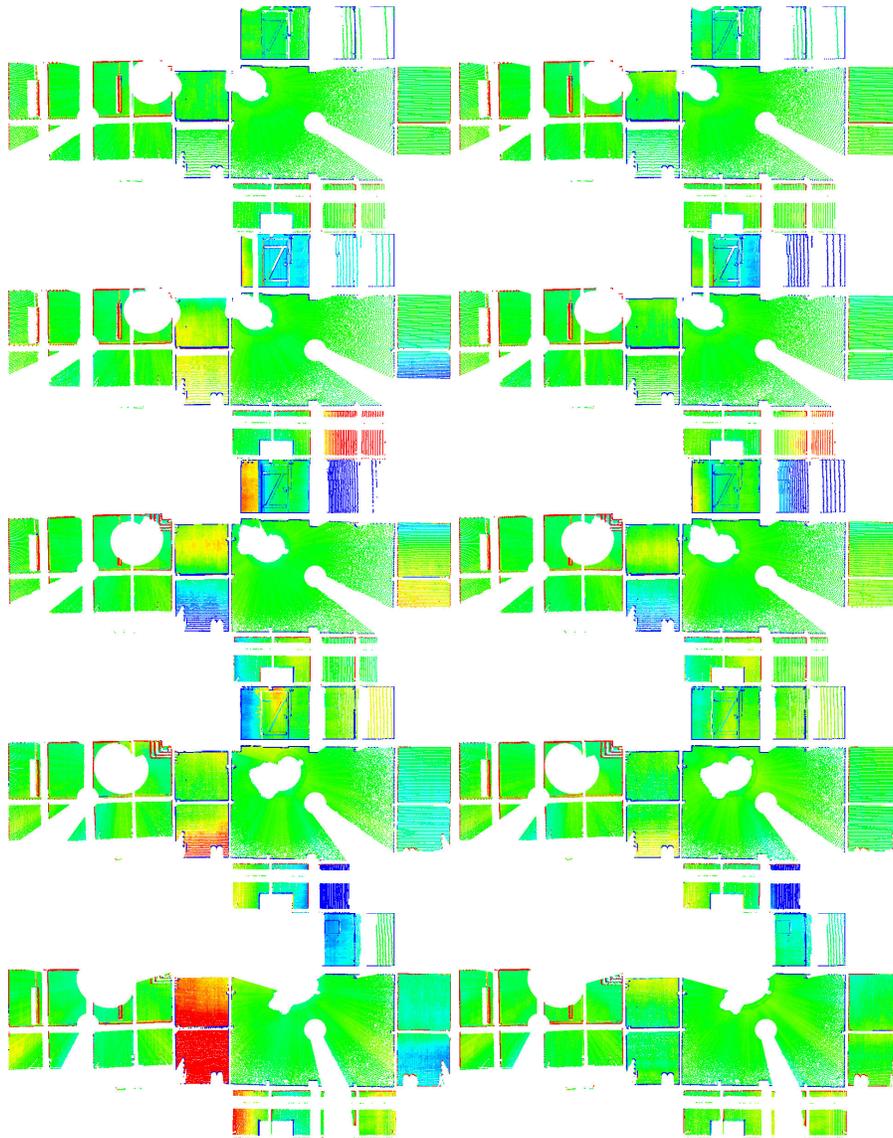


Fig. 4: Comparison of the acquired laser scans with the model using the initial (left) and the automatic semi rigid registration (right). Deviations in cm are color coded as indicated on the right. Best viewed in color.

V. CONCLUSIONS

The proposed semi rigid registration algorithms has shown that it is capable of processing and significantly improving upon a variety of data sets. It exceeds current state of the art rigid registration techniques not only when comparing the final maps produced, but also when comparing the inner deformation of subsections of the point cloud with the ground truth.

REFERENCES

- [1] Andreas Nüchter et al. 3DTK – The 3D Toolkit. <http://slam6d.sourceforge.net/>, February 2011.
- [2] P. Besl and N. McKay. A Method for Registration of 3-D Shapes. *IEEE Trans. on PAMI*, 14(2):239 – 256, February 1992.
- [3] Michael Bosse and Robert Zlot. Continuous 3d scan-matching with a spinning 2d laser. In *IEEE ICRA '09*, pages 4312 –4319, may 2009.
- [4] B. J. Brown and S. Rusinkiewicz. Global Non-Rigid Alignment of 3-D Scans. *ACM Transactions on Graphics*, 26(3):21, 2007.
- [5] H. Chui and A. Rangarajan. A New Point Matching Algorithm for Non-Rigid Registration. *CVIU*, 89(2-3):114–141, 2003.
- [6] L. D. Cohen and I. Cohen. Deformable Models for 3D Medical Images using Finite Elements and Balloons. In *IEEE CVPR*, 1992.
- [7] T. A. Davis. Algorithm 849: A concise sparse Cholesky factorization package. *ACM Trans. Math. Softw.*, 31(4), 2005.
- [8] J. Elseberg, D. Borrmann, K. Lingemann, and A. Nüchter. Non-rigid registration and rectification of 3d laser scans. In *IEEE/RSJ Int. Conf. IROS*, pages 1546 –1552, oct. 2010.
- [9] J. Elseberg, D. Borrmann, and A. Nüchter. Efficient processing of large 3d point clouds. In *ICAT 2011 XXIII International Symposium on*, pages 1 –7, oct. 2011.
- [10] F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens. Multimodality Image Registration by Maximization of Mutual Information. *IEEE Transactions on Medical Imaging*, 1997.
- [11] A. Nüchter, J. Elseberg, P. Schneider, and D. Paulus. Study of Parameterizations for the Rigid Body Transformations of The Scan Registration Problem. *CVIU*, 2010.
- [12] B. Pitzer and C. Stiller. Probabilistic mapping for mobile robots using spatial correlation models. In *IEEE ICRA*, pages 5402–5409, May 2010.
- [13] Todor Stoyanov and Achim J. Lilienthal. Maximum likelihood point cloud acquisition from a mobile platform. In *Proc. of the IEEE ICAR*, June 22–26 2009.
- [14] J. A. Williams, M. Bennamoun, and S. Latham. Multiple View 3D Registration: A Review and a New Technique. In *Proc. of the Int. Conf. on Systems, Man, and Cybernetics*, 1999.

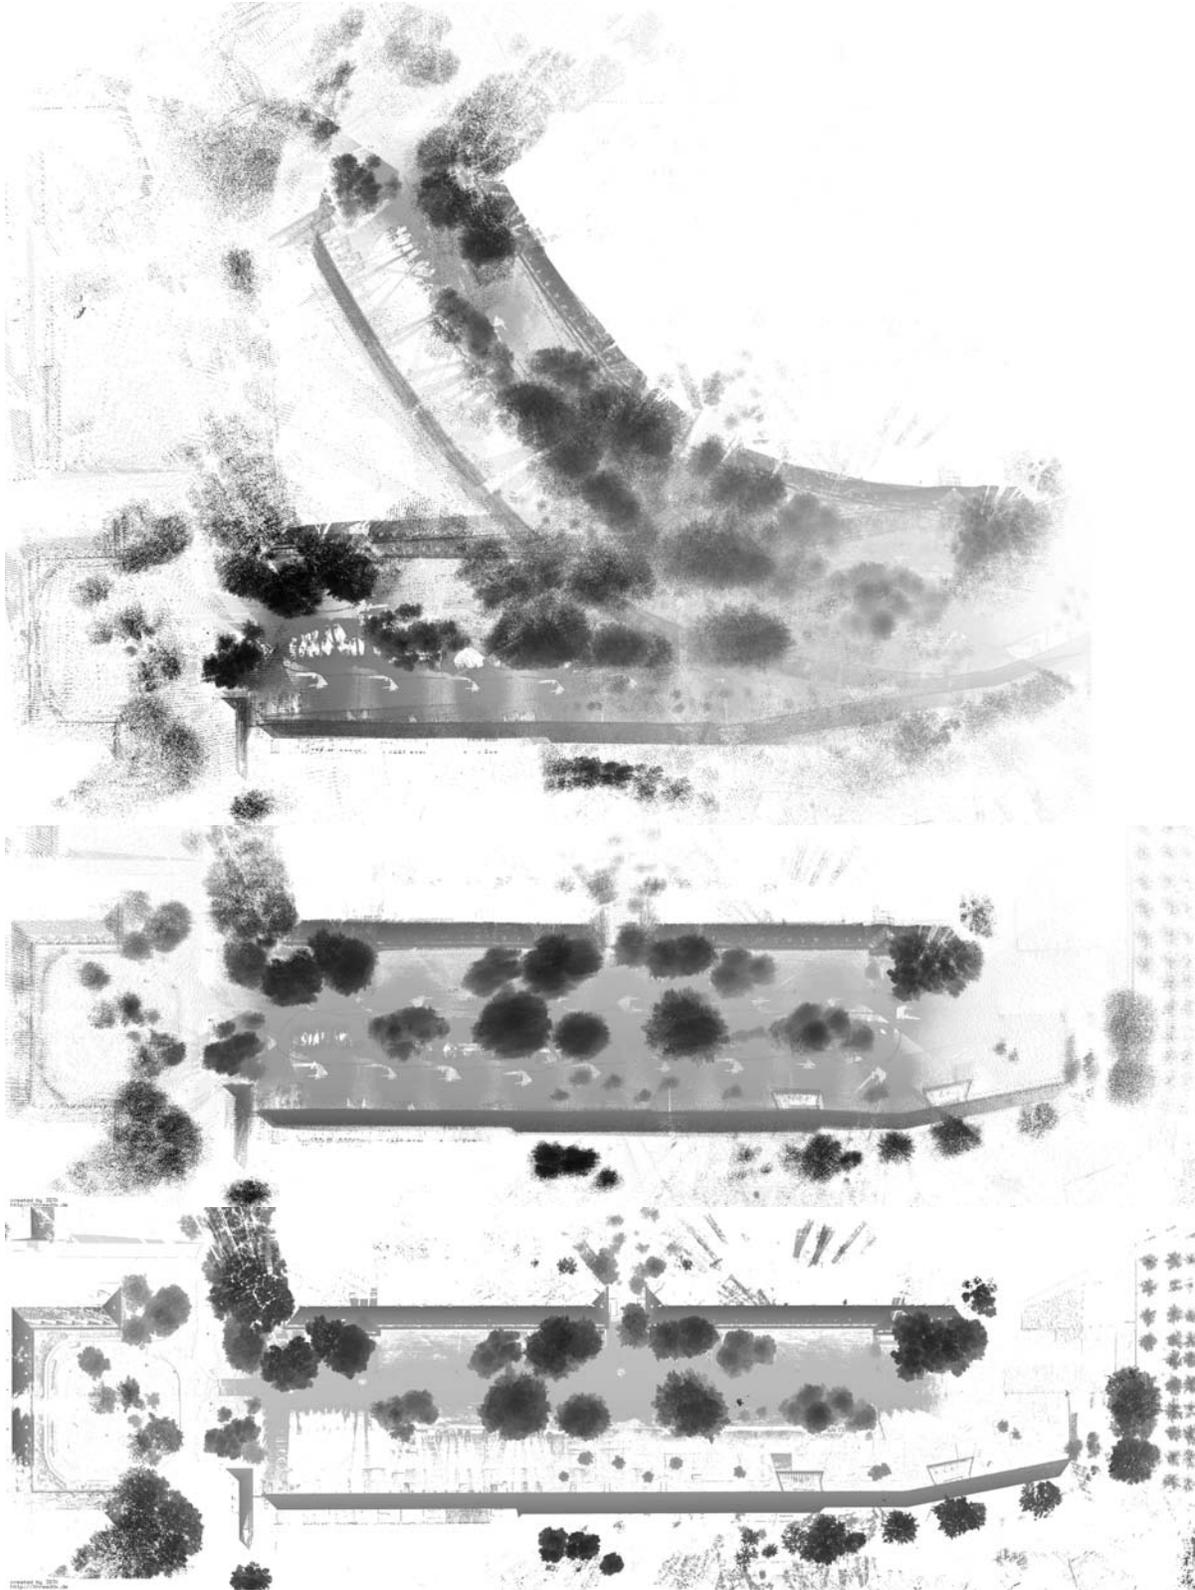


Fig. 5: A comparison of the initial and optimized point cloud acquired by our mobile laser scanner with the ground truth data acquired by a terrestrial laser scanner. Top: The initial point cloud of the car data set with about 12 million points. Large scale errors in the pose estimation prevent rigid scan matching algorithms from being successful. Middle: The result of our novel semi rigid registration algorithm. Bottom: The ground truth data set acquired with the Riegl VZ-400 mounted on a tripod at 12 positions with about 250 million points.