The 2010 IEEE/RSJ International Conference on
Intelligent Robots and Systems
October 18-22, 2010, Taipei, Taiwan

Non-Rigid Registration and Rectification of 3D Laser Scans

Jan Elseberg, Dorit Borrmann, Kai Lingemann, and Andreas Nüchter

Abstract— **Three dimensional point clouds acquired by range scanners often do not represent the environment precisely due to noise and errors in the acquisition process. These latter systematical errors manifest as deformations of different kinds in the 3D range image. This paper presents a novel approach to correct deformations by an analysis of the structures present in the environment and correcting them by non-rigid transformations. The resulting algorithms are used for creating high-accuracy 3D indoor maps.**

I. Introductions

Many robotic tasks require the creation of a model of the robot's surrounding. Then the robot uses its sensors to acquire spatial information. Any sensor and any sensing process is erroneous. This cannot be avoided. Thus, the engineering tasks are to reduce sensor errors as good as technically possible, and to regard noisy and faulty sensor readings in the robot programming.

Recently, 3D laser scanners have become popular in mobile robotics for making 3D maps of the environment [3], [11], [14], [19], [22]. Typical 3D scanners used in robotics consist of a high-precision laser measurement unit either using a pulsed laser measuring the time-of-flight or using a continuous laser wave measuring the phase shift. 2D scanners are built using a rotating mirror. This measurement system is actuated by a servo to yield 3D scans i.e. a 3D point cloud is assembled from several 2D laser scans at a single robot position and different servo positions. As result, ranges can be measured with high accuracy, but systematic errors are the result of imprecise servo movements.

J. Elseberg, D. Borrmann, and A. Nüchter are with School of Engineering and Science of the Jacobs University Bremen gGmbH, 28759 Bremen, Germany {j.elseberg|d.borrmann|a.nuechter}@jacobs-university.de

K. Lingemann is with the Department of Computer Science, University of Osnabrück, 49069 Osnabrück, Germany lingemann@informatik.uni-osnabrueck.de

Different 3D point cloud improvement algorithms are introduced and evaluated in this paper. They establish their non-rigid deformation by computing *correct* points in a global frame of reference. These correct points are either computed directly, or with help of plane information obtained by the 3D Hough transform [2]. Therefore, presented algorithms differ in the way how the desired point positions are computed.

The registration of laser scans is usually done with the goal of finding a set of rigid transformations that brings the scans into a single coordinate system. The implicit assumption is that the range images are a precise representation of the world. However, at some point rigid transformations are no longer sufficient to consistently register the scans so that non-rigid methods need to be adopted.

Preceding all non-rigid improvement of the point clouds, we therefore employ the well-known Iterative Closest Point (ICP) algorithm to align the scans sequentially [1], [19]. After pairwise matchings have been computed the global rigid alignment according to Lu and Milios' relaxation algorithm [3] is used to create a globally consistent map. As a first step of our non-rigid registration and rectification of 3D laser scans a Hough transform is executed to find a set of planes describing the scan. These planes are then optimized so that they comply with a general geometrical model and fit the registered map as good as possible. This is done by using the scheme laid out in section III. From the so produced structures we extract global points, whose positions correspond to the true underlying geometry. A summary of the overall algorithm is given in Algorithm 1.

The last two step of the algorithm are variable and allow two configurations of the non-rigid deformation and two global point extraction methods. The most obvious choice for the computation of global point coordinates is to exploit the point-to-plane correspondences, for example by choosing the

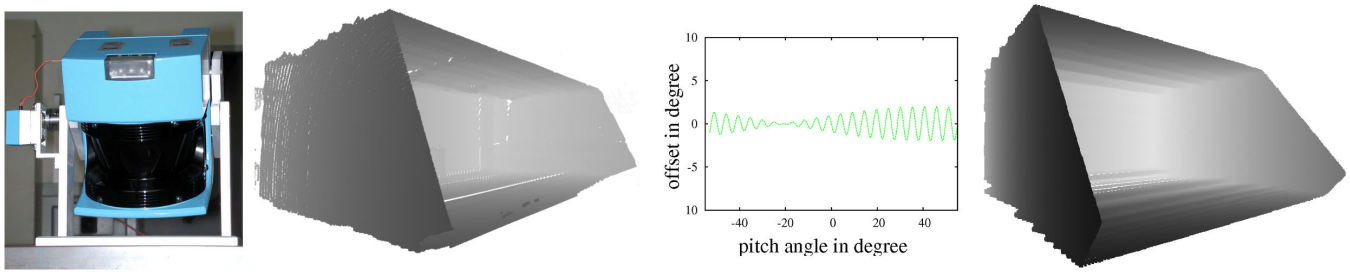

Fig. 1: From left to right: (1) 3D laser scanner. Its technical basis is a SICK 2D laser range finder (LMS-200). (2) A scan from an office environment with deformations as they occur due to the pitching laser scanner. (3) The offset that was generated to simulate an oscillation of the rotating scanner. (4) A simulated scan with randomly generated servo errors in a similar perspective to 2.

Algorithm 1 Map Improvement using Plane Extraction

1: Rigidly register all scans using the globally consistent alignment framework as given in [3]
2: Use Hough Transform to extract planes in every scan [2].
3: Improve the Plane Model (Section III).
4: Extract global point information.
5: Deform each scan using either the TPS (Section VI) or the minimization in Section V.

projection of a point to its plane as its desired global point. Since architectural shapes of environments follow standard conventions arising from tradition or utility [10] we can exploit knowledge for improving 3D scans made in indoor environments, e.g., plane walls. The information about global points is fed into the Thin Plate Spline (TPS) or the slice based deformation as given in Section V.

II. Related Work

The area of non-rigid registration is largely unexplored in anything but the medical imaging community where it is widespread [7], [15].

Williams et al. [21] describe an extension of a rigid registration algorithm that includes point estimation to compensate for noisy sensor data. This technically constitutes a non-rigid registration algorithm designed for low scale high frequency deformations. Similarly, Haehnel et al. [12] modified the ICP algorithm to estimate the position of each measured point. Chui and Rangarajan [6] proposed a point matching algorithm that is capable of aligning point clouds with each other. All these approaches have in common that they only allow for a small amount of points. Subsampling the point clouds leads to less accurate correspondences, thus reducing registration quality. Aiming to reduce noise Unnikrishnan [20] locally fit high-order polynomials to 3D data. Large scale deformations will not be corrected by such a local approach.

Mitra et al. [16] rely on a high number of scans which contain a slowly deforming object. Instead of computing point correspondences they estimate the deformation by the local space-time neighborhood. Consequently, higher point density is not a difficulty but actually a benefit for the algorithm. On the other hand the algorithm is inapplicable to problems with large deformations and movements between individual scans.

The most relevant work has been done by the Stanford Graphics Laboratory. Ikemoto et al. [13] dice laser scans into a set of overlapping regions, which are rigidly transformed. This is iterated until convergence. Extreme care has to be taken when and how to dice each scan so as not to destroy the internal consistency of a scan. Brown and Rusinkiewicz developed a global non-rigid registration process [4]. In [5] they introduced a novel ICP variant to find very accurate correspondences. Even though the registration requires extreme subsampling the deformation is generalized onto the entire scan. Although the computation time of this algorithm is reported to be several days for some data sets, we believe this to be the most promising non-rigid alignment process. This paper gives a closer examination of this algorithm and propose an improvement to greatly speed up the the computations involved.

III. Desired Points by Plane Extraction

Planes are extracted from the point clouds using the randomized Hough Transform as described in [2] or similar plane detection techniques. Each plane i is represented using the Hesse normal form, i.e. by a unit normal $\boldsymbol{n}_i$ and a distance from the origin ρ_i. Systematic errors are still present in the estimated parameters. This has several reasons:

- The plane conforms to the data itself, so that deformations in the scans carry over to the plane.
- The same surface can be detected several times due to registration errors.

For example, in a deformed scan floors and walls may be slightly non perpendicular to each other.

The first step for improving the position and orientation of a given set of planes is to establish geometrical relations between planes as they appear in an indoor environment. Then all plane parameters are optimized and planes are joined if necessary.

The plane classification problem can be stated as follows: Given a set of planes we want to label each as either *floor*, *ceiling*, *wall* or as a *non-feature*. Small planar objects within the scans are regarded as irrelevant to the global plane model and should be classified as non-feature.

Nüchter et al. [18] employed a semantic net as simple Prolog program that encodes semantic relations between objects to solve this problem. With a larger number of planes this becomes intractable and therefore, we propose a different strategy.

It is obvious that the planes we want to classify usually aggregate around a few points in the spherical coordinate system, due to regularities in the architecture. The normals represented in spherical coordinates are inserted into an octree. The spatial subdivision of the octree can now be used to quickly cluster "close" planes together. The distance between two planes i, j is defined as $\measuredangle(\boldsymbol{n}_i, \boldsymbol{n}_j)$. The clustering is controlled by a parameter θ, where planes with $\measuredangle(\boldsymbol{n}_i, \boldsymbol{n}_j) \leq \theta$ will belong to the same cluster. After clustering, each cluster is represented by a normal $\boldsymbol{n}'_i$:

$$\boldsymbol{n}'_i = \frac{\sum_{j=1}^{m} m_j \boldsymbol{n}_j}{\sum_{j=1}^{m} m_j},$$

where m_j is the number of points that lie on plane $\boldsymbol{p}_j$. We assume that the feature planes that are to be classified are joined in their respective clusters. It follows that we need to find 3 clusters that are orthogonal to each other. It is a trivial task to find the biggest three $\boldsymbol{n}'_i$,$\boldsymbol{n}'_j$,$\boldsymbol{n}'_k$ for which:

$$\boldsymbol{n}'_i \cdot \boldsymbol{n}'_j = \boldsymbol{n}'_j \cdot \boldsymbol{n}'_k = \boldsymbol{n}'_k \cdot \boldsymbol{n}'_i = 0$$

holds. *Floor* and *ceiling* planes are now easily identified by their relation to the 3D scanner. The other two clusters

represent the walls of the environment. All remaining planes are classified as *non-feature*.

With the planes correctly classified, their poses are now optimized by minimizing the positional error

$$E = \sum_{\boldsymbol{p}_i} \sum_{j=1}^{m_i} \left| \boldsymbol{n}_i \boldsymbol{x}_j - \rho_i \right|^2, \tag{1}$$

under the side conditions:

$$\begin{aligned} \boldsymbol{n}_i \cdot \boldsymbol{n}_j &= 0 \qquad \text{if } i \text{ and } j \text{ ought to be orthogonal} \\ \boldsymbol{n}_i &= \boldsymbol{n}_j \qquad \text{if } i \text{ and } j \text{ ought to be parallel} \\ |\boldsymbol{n}_i| &= 1. \end{aligned}$$

For each i the m_i points $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_{m_i}$ lie on plane p_i.

There is no closed-form solution to the above equation system. Nüchter [18] introduced an additional error term that replaces the side conditions. This new error metric is then minimized by non-linear optimization algorithms, such as the Downhill Simplex Method or Powell's Method. This approach is likely to lead to a remaining residual error in the planes.

To improve these results the error metric is reformulated.

$$E(\boldsymbol{R}) = \sum_{i} \sum_{j=1}^{m_i} \left| (\boldsymbol{R} \cdot \boldsymbol{n}_i)\, \boldsymbol{x}_j - \rho_i \right|^2. \tag{2}$$

where $\boldsymbol{R}$ is a 3×3 rotation matrix and the normals n_i is the corresponding base vector for the ith plane. The number of parameters is only $3 + n'$, where n' is the total number of classified planes, instead of $3 \cdot n'$. A 3-dimensional rotation matrix is non-linear, so that the metric needs to be linearized.

Assuming that only small rotations are necessary to adjust the planes, the linearized rotation matrix is given by:

$$\boldsymbol{R} \approx \boldsymbol{I}_3 + \begin{pmatrix} 0 & -\theta_z & \theta_y \\ \theta_z & 0 & -\theta_x \\ -\theta_y & \theta_x & 0 \end{pmatrix}.$$

This can now be minimized by solving a linear equation system, for details see [17]. This will not yield the exact solution that minimizes the original metric. However, the optimization is iterated until it converges to a stable set of normals $\boldsymbol{n}'_i$.

After the pose of each plane has been corrected, two planes are joined if their classification is equal and the spatial distance between them is below a threshold. The ρ' of the joined planes is computed as a weighted average by:

$$\rho' = \frac{\sum_{\boldsymbol{p}_i \in E} \sum_{j=1}^{m_i} (\boldsymbol{n}_i \boldsymbol{x}_j)}{\sum_{\boldsymbol{p}_i \in E} m_i},$$

where E is a set containing the equal planes.

IV. Desired Points for Unstructured Environments

The algorithm presented in this section was first described by Brown and Rusinkiewicz [4], [5]. We propose an improvement that is later evaluated on simulated as well as on real world data.

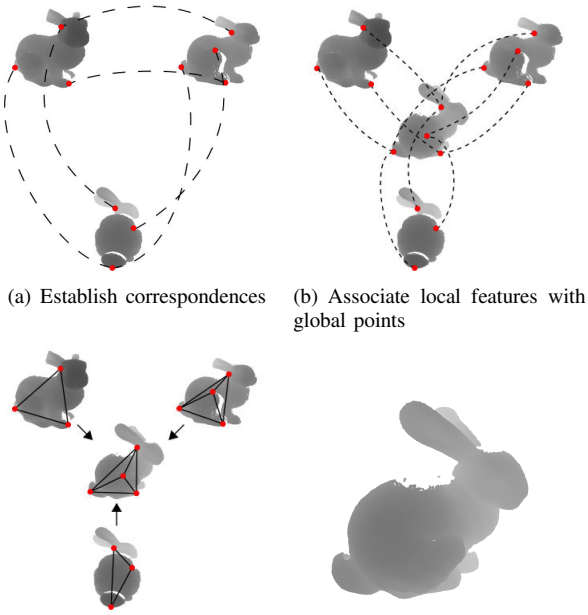

(a) Establish correspondences (b) Associate local features with global points

(c) Optimize global positions (d) Warp scans to global points

Fig. 2: The individual steps of the global non-rigid alignment algorithm. First, the pairwise correspondences between features selected in the laser scans are computed. Then correspondences of local to global features are computed and the global feature positions are optimized. After a consistent global point set is computed, the scans are warped.

The algorithm is an extended scan matching procedure where non-rigid deformations are allowed. Non-rigid registration is a particularly hard problem due to the summation of errors. Small errors in the matching process accumulate and result in large scale incorrect warps that are applied to scans later in the sequence.

The algorithm relies on features that were randomly selected from a scan. First, the correspondences between features in pairs of scans are established by a variant of the ICP algorithm. Second, for each of the features a desired point, i.e., a global position for this feature is computed from the correspondences by minimizing an error metric. Now that each scan has a set of features with associated desired points a non-rigid transformation is calculated to deform the scans. Fig. 2 gives an overview over the complete process.

A. Finding Features and their Correspondences

Feature points should ideally cover the entire model and should be chosen so as to yield stable correspondences between scans. It is intuitive to pick features at significant places in the environment. This would however require some knowledge about the geometry of the entire environment and does not guarantee a sufficient coverage.

A fixed percentage of points (approx. $1\,\%$) are selected using uniform random sampling to ensure coverage everywhere in the environment.

For each feature $\boldsymbol{m}_i^k$ in scan k and each overlapping scan l we search for a corresponding point $\boldsymbol{m}_i^l$. This is achieved by locally weighted ICP, which gives higher importance

to points surrounding the feature $\boldsymbol{m}_i^k$. Instead of roughly estimating correspondences for all points we wish to compute the most precise estimate for each feature. This is equivalent to increasing the accuracy of the alignment in areas close to $\boldsymbol{m}_i^k$. After the algorithm converges, the closest point to $\boldsymbol{m}_i^k$ on scan l is selected as its corresponding point.

The locally weighted ICP is based on the standard ICP baseline algorithm with a point-to-point error metric and closest-point computations using k-D trees for speed up. The only change to the standard algorithm is in the point-selection stage. The set of corresponding points that are considered in the minimization step is sampled using a probability distribution that makes it more likely that points near $\boldsymbol{m}_i$ are used in the estimation process:

$$p(\boldsymbol{x}) = \frac{1}{|\boldsymbol{x} - \boldsymbol{m}_i|^2}.$$

The probability distribution is normalized and numerically inverted to transform the random variable into samples to be selected. This is called the inversion method for sub-sampling [8].

Because an initial registration of the laser scans has already been performed, only few iterations of the locally weighted ICP are required for convergence. In addition, the computationally most expensive part of the ICP, i.e., building the k-D tree is done only once per scan.

B. Computing the Global Point Positions

In order to calculate consistent deformations for all scans, a global position $\boldsymbol{d}_i$ for each feature i is computed. These positions are computed in such a way as to minimize unnecessary deformations.

The arrangement of the global point positions should be as similar to the arrangement of the corresponding local features as possible. Under the assumption that the scans contain no deformation or errors of any kind and the correspondences have been computed correctly, the arrangements of both points are identical. If the scans are slightly deformed this correlation will only hold approximately.

Guided by the intuition that global point positions should be affected by their geometric configuration relative to neighboring features, Brown and Rusinkiewicz optimize the positions $\boldsymbol{d}_i$ by attempting to preserve their relative distances as indicated by the computed correspondences within the scans. Consequently the points are no longer represented as simple positions in a global coordinate system but relative to each other. A "spring" with non-zero rest length is placed between all pairs of features i and j. Assuming feature i was originally selected on scan k and feature j on l, the rest length is set to

$$\frac{1}{2}\left(\left|\boldsymbol{m}_i^k - \boldsymbol{m}_j^k\right| + \left|\boldsymbol{m}_i^l - \boldsymbol{m}_j^l\right|\right).$$

Then an error metric is formulated:

$$E = \sum_k \sum_i \sum_{\substack{j \\ j>i}} w_{i,j} \left(\left|\boldsymbol{d}_i - \boldsymbol{d}_j\right| - \left|\boldsymbol{m}_i^k - \boldsymbol{m}_j^k\right|\right)^2. \quad (3)$$

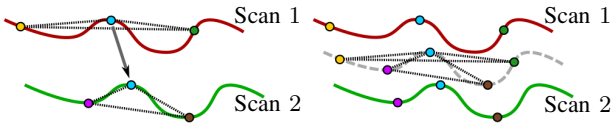

Fig. 3: Global point positions as computed by the error metrics. Left: Local feature points are represented relative to each other. Each of these "springs" is one summand in the error metric. Right: Global point positions are then adjusted to minimize the total spring energy.

When computing this it is necessary to consider that the inner sum over i and j is dependent on k. This means the inner sum is not over all i, j but over all i, j such that

- both features i and j were originally selected on k, or
- feature i was selected on k and feature j, selected in some other scan, has a valid correspondence m_j^k on k.

See Fig. 3 for a closer look on how the sum is computed.

The error metric can be conceptualized as a measure of the discrepancy between the distances of the global points and the distances of the corresponding local features. Minimizing this discrepancy clearly leads to a geometric configuration of the global points that closely resembles the features $\boldsymbol{m}_i^k$. Brown and Rusinkiewicz apply gradient descent to minimize the error.

We propose a novel alternative to the optimization of the global points with the aim of considerably improving the performance of the algorithm. Instead of considering the distance between feature points we will also consider the relative orientation of each feature point to each other. In addition, the problem is modeled in probabilistic terms allowing complex "weights". This approach enables us to compute a globally optimal solution with respect to the error metric much more efficiently.

We view each scan k as a set of measurements of the features $\boldsymbol{m}_i^k$ within that scan. To be more precise, for any feature pair $\boldsymbol{m}_i^k$,$\boldsymbol{m}_j^k$ we assume scan k measured the relative displacement of $\boldsymbol{d}_j$ from $\boldsymbol{d}_i$, i.e., $\boldsymbol{D}_{i,j} = \boldsymbol{d}_i - \boldsymbol{d}_j$. The measurement $\bar{\boldsymbol{D}}_{i,j}^k = \boldsymbol{m}_i^k - \boldsymbol{m}_j^k$ is associated with an uncertainty $\boldsymbol{C}_{i,j}^k$ that replaces the weights $w_{i,j}$ in (3).

Given all the measurements, we wish to maximize the likelihood of the global point positions by minimizing the following Mahalanobis distance.

$$\begin{aligned} E &= \sum_k \sum_{i,j} \left(\boldsymbol{D}_{i,j} - \bar{\boldsymbol{D}}_{i,j}^k\right)^T \boldsymbol{C}_{\boldsymbol{i,j}}^{\boldsymbol{k}\ -1} \left(\boldsymbol{D}_{i,j} - \bar{\boldsymbol{D}}_{i,j}^k\right) \quad (4) \\ &= \sum_k \sum_{i,j} \left(\boldsymbol{d}_i - \boldsymbol{d}_j - \bar{\boldsymbol{D}}_{i,j}^k\right)^T \boldsymbol{C}_{\boldsymbol{i,j}}^{\boldsymbol{k}\ -1} \left(\boldsymbol{d}_i - \boldsymbol{d}_j - \bar{\boldsymbol{D}}_{i,j}^k\right). \end{aligned}$$

Under the assumption that $\boldsymbol{C}_{\boldsymbol{i,j}}^{\boldsymbol{k}\ -1} = w_{i,j}\boldsymbol{I}_3$ minimizing (4) is equivalent to minimizing:

$$E = \sum_k \sum_i \sum_{\substack{j \\ j>i}} w_{i,j} \left(\left(\boldsymbol{d}_i - \boldsymbol{d}_j\right) - \left(\boldsymbol{m}_i^k - \boldsymbol{m}_j^k\right)\right)^2,$$

which is a restated (3). The minimum to (4) can be easily computed in a single step by concatenating the global point

positions into the $(n-1) \times 3$-dimensional matrix $\boldsymbol{D} = (d_2, \ldots, d_n)^T$. The first feature is not included as it defines the global reference system. (4) is now represented in matrix form:

$$E = \left(\bar{\boldsymbol{D}} - \boldsymbol{H}\boldsymbol{D}\right)^T \boldsymbol{C}^{-1} \left(\bar{\boldsymbol{D}} - \boldsymbol{H}\boldsymbol{D}\right),$$

where $\boldsymbol{C}$ is a diagonal matrix containing all $w_{i,j}$ and $\bar{\boldsymbol{D}}$ is the concatenation of all $\bar{\boldsymbol{D}}_{i,j}$. $\boldsymbol{H}$ is an incidence matrix representing the graph created by the edges i, j. The minimum $\boldsymbol{D}$ is given by:

$$\begin{aligned}\boldsymbol{D} &= \left(\boldsymbol{H}^T \boldsymbol{C}^{-1} \boldsymbol{H}\right)^{-1} \boldsymbol{H}^T \boldsymbol{C}^{-1} \bar{\boldsymbol{D}} \\ &= \boldsymbol{G}^{-1} \boldsymbol{B},\end{aligned}$$

where $\boldsymbol{G}$ and $\boldsymbol{B}$ are given by:

$$\begin{aligned}G_{i,i} &= \sum_k \sum_j w_{i,j}^{-1} \\ G_{i,j} &= \sum_k w_{i,j}^{-1} \qquad \text{for} \quad (i \neq j) \\ B_i &= \sum_k \sum_{i,j} w_{i,j}^{-1} \left(\boldsymbol{m}_i^k - \boldsymbol{m}_j^k\right).\end{aligned}$$

Estimating the optimal positions of the global points is therefore reduced to computing the inverse of an $(n-1) \times (n-1)$ matrix. As this is in the order of $O(n^3)$ this algorithm is much more efficient than the gradient descent method.

V. Deformation on Basis of an Error Model

We improve the scan quality for scanners with systematical errors after a set of desired points has been computed. The rotation axis of the scanner is assumed to be the x-axis. For each slice s, we aim to compute the correct orientation α_s the scanner had by minimizing the summed point-to-plane distances.

$$E = \sum_s \sum_{(\boldsymbol{d}_j, \boldsymbol{m}_j)} \left|\boldsymbol{R}_{\alpha_s} \boldsymbol{m}_j - \boldsymbol{d}_j\right|^2 + \lambda \sum_s \left|\alpha_s\right|^2,$$

If the desired position of the measurements is given by a set of planes this is restated as:

$$E = \frac{1}{n} \sum_{p_i} \sum_{\boldsymbol{x}_j \in p_i} \left|\left(\boldsymbol{R}_{\alpha_s} \boldsymbol{x}_j\right) \boldsymbol{n}_i - \rho_i\right|^2,$$

where n is the total number of points associated to planes. The rotation matrix $\boldsymbol{R}_{\alpha_s}$ describes a rotation around the x-axis by α_s. The error metric is easily minimized in an iterative fashion by using the small angle approximations $\sin(\alpha) \approx \alpha$ and $\cos(\alpha) \approx 0$.

VI. Deformation on Basis of TPS

The Thin Plate Spline (TPS) was first formulated by J. Duchon in 1976 [9] as a way to model the deformation of a thin metal sheet. The TPS is able to represent any perceivable deformation of a scan, while guaranteeing the maximum amount of smoothness. It is robust to measurement errors as well as errors in the correspondences The TPS function

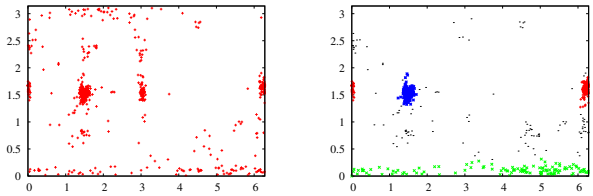

(a) Extracted normals (b) Classification by clustering

Fig. 4: (a) Plotted in spherical coordinates, the normals are concentrated around a few points. (b) The classification of the planes as computed by the clustering approach. All normals have been flipped to the side of the unity sphere where the corresponding cluster has first been identified. The algorithm classified a total amount of 568 planes. 122 planes remain unclassified. These are mostly normals that lie between the 3 main clusters and are of dubious relevance to the optimization.

T is separable into an affine part $\boldsymbol{A} \in \mathbb{R}^{e \times d}$ and a non-affine part, the warping parameters $\boldsymbol{W} = (\boldsymbol{w}_1, \ldots, \boldsymbol{w}_n)^T$ with $\boldsymbol{W} \in \mathbb{R}^{e \times n}$. $\boldsymbol{T}$ is given by:

$$\boldsymbol{T}(\boldsymbol{x}) = \boldsymbol{A}\boldsymbol{x} + \sum_{i=1}^{n} \boldsymbol{w}_i \left|\boldsymbol{d}_i - \boldsymbol{x}\right|,$$

where $\boldsymbol{x}$ is a point in homogenous coordinates. Concatenating the desired points into $\boldsymbol{D}$ and their corresponding feature points into the $\boldsymbol{M}$, the TPS can be computed by solving the linear equation system:

$$\begin{aligned}\boldsymbol{A}\boldsymbol{D} + \boldsymbol{W}(\boldsymbol{K} + n\lambda\boldsymbol{I}) &= \boldsymbol{M} \\ \boldsymbol{W}\boldsymbol{D}^T &= 0.\end{aligned}$$

$\boldsymbol{D}$ is a $4 \times n$ matrix, since the desired points are represented in homogenous coordinates and $\boldsymbol{M}$ is a $3 \times n$ matrix. $\boldsymbol{K}$ is given by $K_{i,j} = |\boldsymbol{d}_i - \boldsymbol{d}_j|$.

VII. Experimental Evaluation

A. Plane Optimization

First we evaluate the quality of the classification algorithm using a showcase data set with 63 laser scans from the fourth floor of the university building in Osnabrück. After extracting almost 700 planes, the two classification algorithms were started. The results are depicted in Fig. 4(b). Many of the detected planes do not belong to either floors, ceilings or walls. The clustering algorithm is preferable because it produces a more realistic labeling.

The Prolog program does not perform well when the amount of planes increases to a more realistic number ($>$ 40). As Prolog exhausts every possible assignment when looking for an exact solution its computation time can reach extreme heights. The clustering algorithm clearly outperforms by several magnitudes (ms instead of min). Even when using several hundred planes as input the clustering is computed in less than a second.

The required computation time for the improvement step after classification never exceeded 30 ms. The usual number of iterations required for convergence is 5.

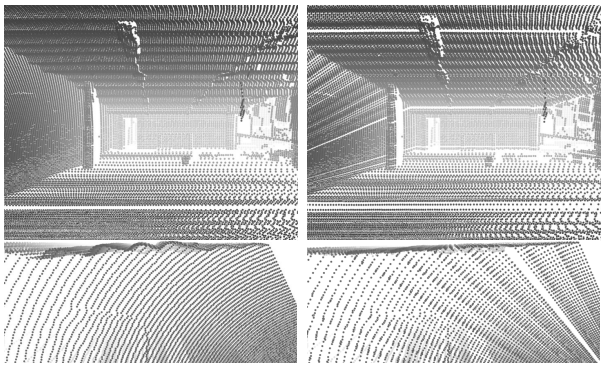

Fig. 5: A scan of the soccer field on the fourth floor in the AVZ building in Osnabrück before and after improvement. Top: The scan before(left) and after(right) correction. Bottom : Close-up view of the ceiling of the initial scan and the corrected scan.

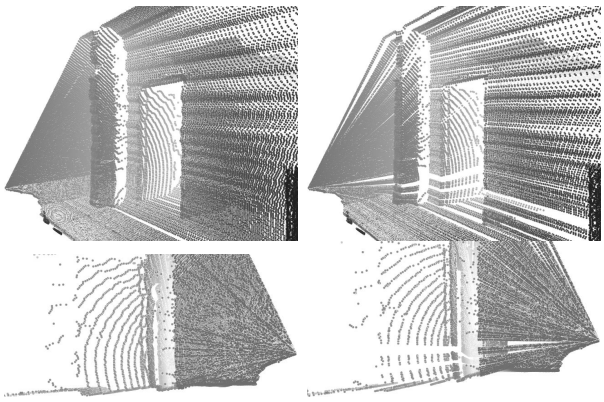

Fig. 6: A second scan of the soccer field before and after improvement. In this scan the wall in front of the scanner is more prominent while other planes are much less apparent. The algorithm introduced some new error into the scan. Top: The scan before(left) and after(right) correction. Bottom: Close-up view of the introduced error in the floor. The floor in general is improved. However, slices immediately behind the door are incorrectly positioned.

B. Deformation using planes

The improvement algorithm from section V is evaluated on real as well as on simulated data. The computation time of the improvement step (minus plane extraction) in all examples was less than $1\,\mathrm{ms}$ per scan on modern hardware.

We applied the algorithm to every one of the 63 scans from the fourth floor of the AVZ university building in Osnabrück. The errors that produce wave-like effects on floor and ceiling were greatly reduced and the scans were visibly improved. Fig. 5 shows a representative scan.

In some cases however (in 8 of 63 scans) the algorithm encountered difficult geometry as in Fig. 6 and parts of the scans were misaligned. This often occurred on slices that were acquired shortly before the scanner reached a parallel position to the floor. Points from these slices tend to be mistaken for points stemming from the floor. The following improvement steps then find an α that moves these points closer to the floor so that the slice is misaligned. However, even in these worst cases, most slices of the scan were brought into a visually more appealing position.

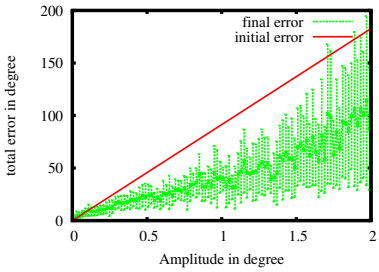

Fig. 7: Results of the improvement using simulated data. The error bars indicate the variance for each trial.

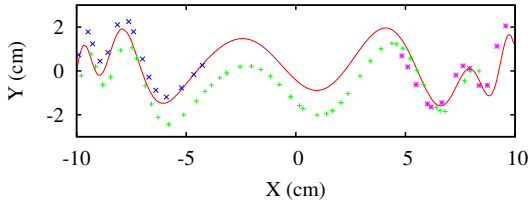

Fig. 8: The simulated environment with 3 scans. The solid curve represents the environment that is to be estimated. The points constitute the scans. Rigid transformation errors simulating a misalignment in the registration process have been applied.

To objectively demonstrate the effectiveness of the improvement, a laser scanner with a variable amount of oscillation error was simulated in a simple environment. An example is given in Fig. I. The function generating the error is given by $E(\varphi) = A\sin(p_1\varphi)\sin(p_2\varphi + P)$, where the phase offset P was randomly generated from a uniform distribution $P \in \left[-\frac{\pi}{2}, \frac{\pi}{2}\right]$. In each trial the amplitude A was varied and 20 scans were generated. The mean error of the optimized scans, i.e., the summed absolute difference between the estimated and the correct α is plotted against the amplitude A in Fig. 7. Clearly, the effectiveness of the improvement suffers from higher initial errors.

C. Deformations with TPS

In order to objectively assess the quality of the different error metrics an artificial 2-dimensional test scenario is used. The complete non-rigid alignment is also qualitatively evaluated on a real data set.

A number of global points that lie on a curve as seen in Fig. 8 are generated. Then a set of 2-dimensional range scans that sampled the curve are simulated. To each of the local scan points a positional error as well as a randomly generated rigid transformation is applied. The effect of the noise and rigid errors are depicted in Fig. 9. Finally, for each scan a non-rigid deformation is simulated by transforming a point $\boldsymbol{p}$ as follows:

$$p'_y = p_y + A\sin L.$$

The effect of the amplitude A on the registration accuracy is depicted in Fig. 9. The registrations accuracy is given by sum of distances:

$$E = \sum_i \left|\boldsymbol{d}_i - \bar{\boldsymbol{d}}_i\right|,$$

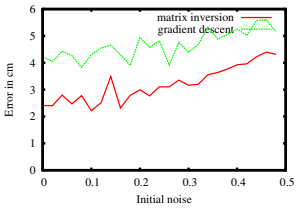

(a) Noise

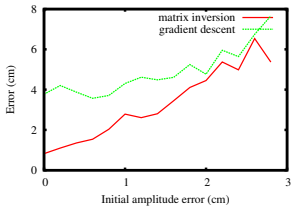

(b) Non-Rigid errors

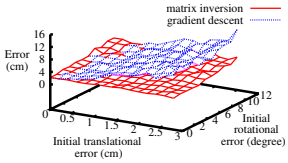

(c) Rigid errors

Fig. 9: Evaluation of the global estimation algorithms with different amount and types of error. All experiments were carried out with 100 global points. Left: The error of both algorithm with varying amount of noise. Right: In this test the parameter pertaining to the non-rigid deformation was varied. Bottom: The accuracy of both approaches with different rigid errors.

where $\bar{\boldsymbol{d}}_i$ is the estimation of the ith global point $\boldsymbol{d}_i$.

Each of the trials were repeated 20 times with randomly generated initial errors and scans. Both algorithms received the same input. The results are shown in Fig. 9. For small and proportional errors the matrix inversion usually fares better. The accuracy of the novel estimation algorithm decreases when more rotational error is introduced in the scans. This is due to the new metric encoding the relative position of the points. Rotating a scan leads to an increase of the error, while this is not the case for Brown's metric.

The second point of interest is the efficiency of both algorithms. As seen in Fig. 10 the newly proposed estimation algorithm is much faster than the gradient descent. Also shown is the error of the estimated global points in relation to their total number. As expected the error increases with more points.

VIII. Conclusions and Future Work

This paper presented a novel map improvement approach for structured environments that shows promising results. Possible application are low-cost laser scanners as employed on mobile robots and kinematic laser scanning. We also presented a significant improvement in the efficiency of the non-rigid alignment procedure by Brown and Rusinkiewicz. Needless to say, much work remains to be done. We plan on strongly generalizing the improvement algorithm to cope with arbitrary structures and variable laser scanners. In the process robustness against incorrectly classified structures will be enhanced. For the non-rigid registration in unstructured environments we plan to incorporate further parameters for the rotation of each scan. This should sacrifices a little speed while gaining more flexibility and accuracy.

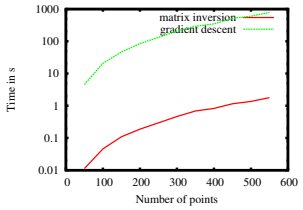

(a)

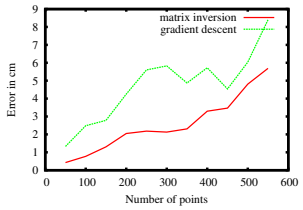

(b)

Fig. 10: Evaluation of the estimation algorithms on the synthetic data set with a variable number of points. Left: The computation time required for both algorithms in a plot of logarithmic scale. Right: Comparison of the accuracy of both algorithms.

References

[1] P. Besl and N. McKay. A Method for Registration of 3–D Shapes. *IEEE Trans. PAMI*, 14(2):239 – 256, Feb. 1992.

[2] D. Borrmann, J. Elseberg, K. Lingemann, and A. Nüchter. A Data Structure for the 3d Hough Transform for Plane Detection. In *Proc. of IAV*, 2010 (accepted).

[3] D. Borrmann, J. Elseberg, K. Lingemann, A. Nüchter, and J. Hertzberg. Globally Consistent 3D Mapping with Scan Matching. *Journal of Robotics and Autonomous Systems*, 56(2):130–142, 2008.

[4] B. J. Brown and S. Rusinkiewicz. Non-Rigid Global Alignment using Thin-Plate Splines. *ACM Trans. on Graphics*, 26(31):16, 2005.

[5] B. J. Brown and S. Rusinkiewicz. Global Non-Rigid Alignment of 3-D Scans. *ACM Trans. on Graphics*, 26(3):21, 2007.

[6] H. Chui and A. Rangarajan. A New Point Matching Algorithm for Non-Rigid Registration. *CVIU*, 89(2-3):114–141, 2003.

[7] L. D. Cohen and I. Cohen. Deformable Models for 3D Medical Images using Finite Elements and Balloons. In *CVPR*, '92.

[8] L. Devroye. *Non-Uniform Random Variate Generation*. Springer-Verlag, 1986.

[9] J. Duchon. Interpolation des Fonctions de deux Variables suivant le Principe de la Flexion des Plaques Minces. *RAIRO Analyse Numerique*, 10:5 – 12, 1976.

[10] R. B. Fisher. Applying Knowledge to Reverse Engeneering Problems. In *Proc. of GMP '02*, pages 149 – 155, Riken, Japan, July 2002.

[11] A. Georgiev and P. K. Allen. Localization Methods for a Mobile Robot in Urban Environments. *Trans. on RA*, 20(5):851 – 864, October 2004.

[12] D. Hähnel, S. Thrun, and W. Burgard. An Extension of the ICP Algorithm for Modeling Nonrigid Objects with Mobile Robots. In *Proc. of IJCAI*, Acapulco, Mexico, 2003. IJCAI.

[13] L. Ikemoto, N. Gelfand, and M. Levoy. A Hierarchical Method for Aligning Warped Meshes. In *Proceedings of the Fourth International Conference on 3-D Digital Imaging and Modeling (3DIM)*, 2003.

[14] P. Kohlhepp, M. Walther, and P. Steinhaus. Schritthaltende 3D-Kartierung und Lokalisierung für mobile Inspektionsroboter. In *Proc. of AMS 2003, 18. Fachgespraeche*, December 2003.

[15] F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens. Multimodality Image Registration by Maximization of Mutual Information. *T-MI*, 1997.

[16] N. J. Mitra, S. Flöry, M. Ovsjanikov, N. Gelfand, L. Guibas, and H. Pottmann. Dynamic Geometry Registration. *Proc. SGP* 2007.

[17] A. Nüchter, J. Elseberg, P. Schneider, and D. Paulus. Linearization of Rotations for Globally Consistent n-Scan Matching. In *Proc. of ICRA 10*, Anchorage, AK, USA, 2010.

[18] A. Nüchter, H. Surmann, K. Lingemann, and J. Hertzberg. Semantic Scene Analysis of Scanned 3D Indoor Environments. In *Proc. VMV '03*, pages 215 – 222, Munich, Germany, November 2003.

[19] H. Surmann, K. Lingemann, A. Nüchter, and J. Hertzberg. 6D SLAM - Preliminary Report on Closing the Loop in Six Dimensions. In *Proc. IAV '04*, Lissabon, Portugal, June 2004.

[20] R. Unnikrishnan and M. Hebert. Denoising Manifold and Non-Manifold Point Clouds. In *BMVC*, 2007.

[21] J. A. Williams, M. Bennamoun, and S. Latham. Multiple View 3D Registration: A Review and a New Technique. In *Proc. of SMC*, 1999.

[22] O. Wulf, K. O. Arras, H. I. Christensen, and B. A. Wagner. 2D Mapping of Cluttered Indoor Environments by Means of 3D Perception. In *Proc. ICRA '04*, pages 4204 – 4209, New Orleans, USA, April 2004.