AUTOMATIC REGISTRATION OF TERRESTRIAL LASER SCANNER POINT CLOUDS USING NATURAL PLANAR SURFACES

P. W. Theiler*, K. Schindler

Institute of Geodesy and Photogrammetry, ETH Zurich, CH-8093, Zurich, Switzerland
pascal.theiler@geod.baug.ethz.ch, konrad.schindler@geod.baug.ethz.ch

Commission III, WG III/2

KEY WORDS: Laser Scanning, TLS, Point Cloud, Registration, Automation, Planar Surfaces

ABSTRACT:

Terrestrial laser scanners have become a standard piece of surveying equipment, used in diverse fields like geomatics, manufacturing and medicine. However, the processing of today's large point clouds is time-consuming, cumbersome and not automated enough. A basic step of post-processing is the registration of scans from different viewpoints. At present this is still done using artificial targets or tie points, mostly by manual clicking. The aim of this registration step is a coarse alignment, which can then be improved with the existing algorithm for fine registration. The focus of this paper is to provide such a coarse registration in a fully automatic fashion, and without placing any target objects in the scene. The basic idea is to use virtual tie points generated by intersecting planar surfaces in the scene. Such planes are detected in the data with RANSAC and optimally fitted using least squares estimation. Due to the huge amount of recorded points, planes can be determined very accurately, resulting in well-defined tie points. Given two sets of potential tie points recovered in two different scans, registration is performed by searching for the assignment which preserves the geometric configuration of the largest possible subset of all tie points. Since exhaustive search over all possible assignments is intractable even for moderate numbers of points, the search is guided by matching individual pairs of tie points with the help of a novel descriptor based on the properties of a point's parent planes. Experiments show that the proposed method is able to successfully coarse register TLS point clouds without the need for artificial targets.

1. INTRODUCTION

Terrestrial Laser Scanners (TLS) are increasingly used the last 10 years, in a continuously growing number of applications ranging from cultural heritage documentation, surveying, industry and manufacturing to medicine. Many applications have the goal to generate high-quality geometric models. The processing of the resulting large datasets with billions of points is time-consuming, cumbersome and not automated enough.

Due to the fact that TLS is a line-of-sight instrument with limited range of coverage, most applications require a series of scans to obtain a complete model. Hence, the scans need to be registered in a common coordinate system.

2. RELATED WORK

To speed up the registration process, numerous approaches have been published over the last couple of years, which aim at the automatic registration of two or more point clouds. In this section we review the state-of-the-art.

The registration amounts to finding the relative orientation of two or more point clouds. If the scan orientations are known from external sensors (e.g. from GNSS and INS), the problem does not arise. Such direct georeferencing is often applied in mobile laser scanner applications (e.g. Asai et al., 2005). Due to the uncertainty of the additional sensors, the resulting registration is often rather coarse. Furthermore, the costs for the sensor system increase.

The standard technique to achieve a coarse registration without additional sensors is to place artificial targets in the scene and identify them in different scans. The targets are typically objects, which are invariant to the scanner viewpoint (e.g. planes, spheres, cones). Often, they are retro-reflective to simplify their automatic recognition (e.g. Bornaz et al., 2002). Many commercial software packages support automatic registration with reflective targets. A related approach finds non-reflective spheres (Franaszek et al., 2009).

To avoid the need to place targets in the scene, a logical next step is to base the registration on natural geometric elements in the scene such as points, lines or surfaces (Goshtasby, 2005). Thus, the main challenge is to reliably extract and match corresponding features in different scans. Roth (1999) describes a method based on point features, which are extracted from the laser intensity image with an interest operator and transferred to 3D using the range information. The matching is accomplished by an exhaustive search for congruent tie point triangles. Seo et al. (2005) investigate the possibility to apply standard interest point operators to register laser point clouds. Similar approaches are presented in Bendels et al. (2004) and Böhm and Becker (2007). They apply the SIFT operator (Lowe, 2003) on the intensity image to find adequate tie points. Likewise, Moldovan et al. (2009) use intensity information and the SIFT operator to register multiple scans. To further improve the registration with the SIFT operator Mateo and Binefa (2009) first extract dominant planes in the scan and apply the interest point operator on them. Plane extraction is done by calculating the local normal for each point with singular value decomposition. Robust matching of the extracted point features is often achieved with RANSAC (Fischler and Bolles, 1981). Kang (2009) present an approach based on pixel-to-pixel correspondence in the intensity image, followed by outlier detection and computation of the transformation parameters in the 3D space.

Instead of exploiting the intensity data, Basdogan and Oztireli (2008) apply a geometric descriptor, which is based on the distance from each point to the centre of mass of its neighbours, at every single 3D point. Thus, the method is only suitable for small point clouds. Barnea and Fillin (2007) present a tie point extractor by applying a 3D corner detector to the range image. Johnson and Herbert (1999) use spin images as descriptor to find correspondences. This descriptor is generated by computing a local basis at an oriented point (3D point with local surface normal), and storing the resulting two-parameter description of nearby points to describe the local geometry. Shan et al. (2004) improves that approach with geometric constraints between simple configurations of multiple tie points.

He et al. (2005) propose an approach to register range images using Complete Plane Patches. The main idea is to extract planar surfaces and pick out only those, which are complete (not occluded). Correspondence is established using an interpretation tree and several constrains. Dold and Brenner (2006) have developed a similar approach based on planar patches found by region growing. Patches are matched using their area, boundary length, bounding box and mean intensity value. Additionally, an on-board image sensor is used to improve the registration. Instead of constraining the matching based on the plane properties, Dold and Brenner (2007) use the intersection angle to prune the number of possible plane triples to calculate the coarse transformation.

The approaches described so far are primarily designed for coarse registration. The results are then polished with a fine-registration algorithm. The most established method is the Iterative Closest Point (ICP) algorithm (Besl and McKay 1992, Chen and Medioni 1992). The algorithm iteratively minimizes the distances of all points in one scan to the nearest point or plane in the other. There are many variants and extensions of the initial algorithm (e.g. Masuda and Yokoya, 1995, Bergevin et al., 1996, Bae and Lichti, 2004), aimed to increase computational efficiency, robustness, convergence etc. ICP is a local optimization scheme and requires a good initial approximation of the transformation. An alternative is Least Squares 3D Surface Matching (Gruen and Akca 2005). Again, the method finds a local optimum and needs good initial values to converge to the correct solution.

To summarize, the crucial step is the coarse registration, whereas the refinement of an approximate registration can be considered solved.

3. PROPOSED METHOD

The goal of our method is a coarse registration of two terrestrial laser point clouds without the need for artificial objects. The principle workflow is shown in Figure 1. The focus of the paper is on tie point extraction and matching. To complete the registration process, a fine registration with a standard algorithm (e.g. ICP, LS3D) is necessary. As the goal is coarse registration, the paper focusses on pairwise registration (i.e. two scans), which is sufficient to generate approximate transformations for further processing.

Our method is based on virtual tie points which are generated by intersecting triples of scene planes. The planes are detected with RANSAC, embedded in a multi-scale pyramid, which on the one hand increases the chance to find the dominant planes and on the other hand reduces the computation time.
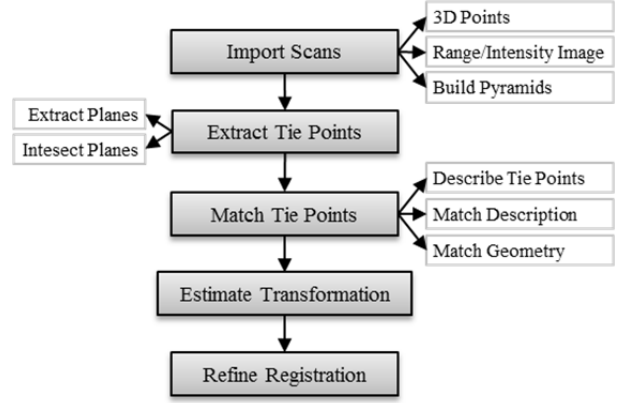

Figure 1: Workflow of proposed method for autonomous coarse registration of terrestrial laser scans

The reasoning behind the apparent detour of generating virtual tie points – rather than directly matching the planes - is the following: plane matching so far proved to be rather unreliable, because the geometric properties of planar segments tend to vary a lot across viewpoints. Points, corresponding to triples of intersecting planes, have additional geometric invariants, which allow for more powerful local descriptors.

Matching based on geometric constraints between tie points is very reliable, but of combinatorial complexity and thus intractable for useful numbers of tie points. On the contrary, matching points individually using local descriptors is efficient, but error-prone and often ambiguous. We thus opt to combine the two steps in a two-stage procedure to get the best of both worlds. First the combinatorial set of putative correspondences is filtered with a novel geometric descriptor, by discarding the large majority of correspondences whose descriptors are very different. Then, rather than taking final decisions based on descriptors, matching within that reduced set of candidate matches is accomplished by searching the largest subset, for which all point-to-point distances are consistent between the two scans.

4. TIE POINT EXTRACTION

The tie points are virtual 3D points, which are generated by intersection of three non-parallel scene planes. These planes are extracted from the scan data by means of RANSAC on a multi-scale pyramid.

Individual laser scans are represented as range images, which due to the polar measurement principle of the scanner does not entail any loss of information. For each scan pyramids are created by repeated down sampling of the range image by a factor of 0.5. For our purposes it is more important to avoid smoothing over range discontinuities, whereas geometric aliasing is not a concern, thus we use nearest-neighbour resampling. The required number of pyramid levels depends on the point density and on the geometry of the scene – denser scans and larger number of dominant planes require more levels.

Plane extraction starts at the highest pyramid level (i.e. the one with the smallest range image, respectively the lowest 3D point density). Planes are found by iterative RANSAC. To increase the chance of finding correct planes we constrain the random sampling to points within a maximal radius. As usual, the

parameters of detected planes are re-estimated with all inlier points. Laser scans have an absolute scale, such that thresholds can be specified in metric world units.

With a set of points (3 or more) plane fitting is accomplished in a total least squares sense by estimating the normal vector with singular value decomposition (SVD). Given the normal vector N the orthogonal distance to the origin D is found by projecting the 3D points' centre of mass $\bar{X}$ onto the normal vector, $D = \bar{X}^T \cdot N$. We point out that in range scans the point density in world coordinates decreases with the distance from the scanner. The threshold t for the minimum support of a plane is thus adapted to the range. Given the total number of points in the scan S_0, the current pyramid level l and a user-specified proportion p, as well as the mean range of the scan R_0, the threshold for a plane with the mean range R_i is computed as

$$S_i > t = \frac{p \cdot S_0}{l} \cdot \frac{R_0}{R_i} \qquad (1)$$

Planes which fulfil the condition are accepted. When no further planes can be detected, the extraction continues on the next-lower pyramid level. The scalar p specifies what fraction of the total scan a plane shall minimally cover. The appropriate value depends on the size of the planes in the recorded scene. The linear scaling with the pyramid level adapts the threshold to the point density of a given level while still biasing the extraction process towards dominant planes.
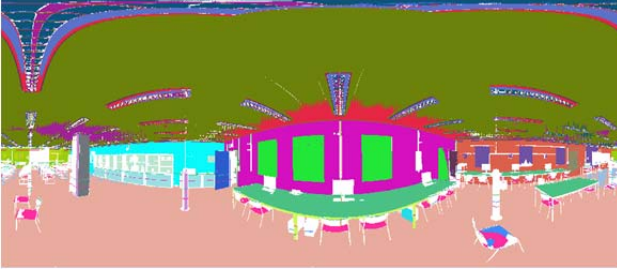

Figure 2: Panorama image with the extracted planes in an indoor environment
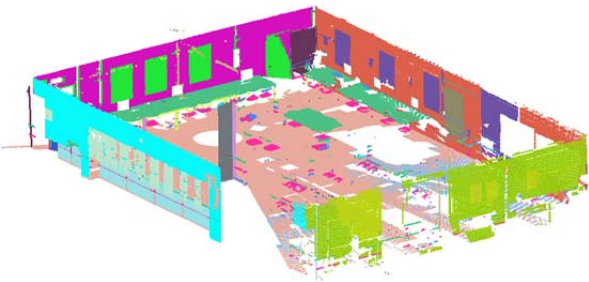

Figure 3: Extracted planes in the 3D space in an indoor environment (without the ceiling)

Figure 2 shows an example scan with the detected planes coded by different point colours. The same result in a 3D view is presented in Figure 3, with the ceiling plane removed for better visibility. The figure demonstrates the completeness and correctness of the plane extraction – almost all planes of reasonable point count have been found, with only few spurious planes.

For all triples of detected planes, the intersection point x is computed analytically to yield a virtual tie point.

$$x = \begin{pmatrix} N_1 \\ N_2 \\ N_3 \end{pmatrix}^{-1} \cdot \begin{pmatrix} D_1 \\ D_2 \\ D_3 \end{pmatrix} = A^{-1} \cdot b \qquad (2)$$

The quality of each tie point is calculated using the reciprocal condition number of the matrix A. A well-conditioned matrix indicates that the planes are not near-parallel and all intersection angles are large enough. Consequently the tie point is well defined. Very badly conditioned tie points are discarded.

The order of the parent planes is based on the z-value of the plane normal. Ambiguous order was treated by taking all solutions as possible tie points.

5. TIE POINT MATCHING

Laser scans have an absolute Euclidean scale and thus distances between tie points are directly comparable across scans. For any two pairs of corresponding tie points the point-to-point distances in both scans must therefore be the same. This geometric constraint is exploited in our scheme. Scan matching is formulated as finding the largest possible set of correspondences, for which all pairwise distances are the same (up to noise). This is an instance of the maximum-clique problem, which is NP-hard. Even though there are faster approximations, the problem is computationally too expensive for large point sets. To make it tractable, we therefore conservatively prune the exhaustive set of possible matches with a novel descriptor before geometric matching, which greatly reduces the number of putative correspondences and hence the search space for the maximum clique.

5.1 Construction of the Description Vector

For each tie point we construct a descriptor vector, which encodes local properties of the point, respectively the planes from which it has been constructed. The vector has the following entries:

- Reciprocal condition number (1 value)
- Intersection angles between planes (3 values)
- Extent of planar segments (6 values)
- Smoothness of planes (3 values)

We go onto explain the individual entries of the 13-dimensional descriptor in more detail.

Reciprocal Condition Number

This scalar is already calculated when intersecting the three planes. It encodes the geometric quality of the intersection.

Intersection angles

The three parent planes of a tie point give rise to three pairwise intersection angles, calculated simply as the scalar product of the two unit normal vectors. By convention we use the smaller of the two intersection angles and divide angles by pi/2, such that the values are between 0 and 1.

Extent of Segments

The tie point's parent planes are constructed by fitting a set of 3D scan points found with RANSAC. For each plane we calculate the bounding rectangle of those points. The width and height of the rectangle are found by principal component analysis in the 2D coordinate system of the plane. Outliers

outside the three-sigma interval are ignored. To obtain values between 0 and 1, the width and height values are divided by the maximum possible range, which equates to twice the maximum measurement range of the scanner.

Plane Smoothness

Although all planes are extracted with the same absolute threshold on residuals (in our experiments 1 cm), their average residuals are quite different due to varying material properties of the planar surfaces. To preserve that information, the mean residuals of the three plane fits are added to the descriptor. The values are scaled by the inlier threshold to map them to the range [0...1].

5.2 Descriptor-based Pruning

So far the relative scaling of the descriptor entries is arbitrary. To obtain meaningful similarities between descriptors, the right relative scaling between the descriptor dimensions is needed. We have experimentally determined suitable weights. The experiments showed that the most reliable values are the intersection angles and the reciprocal condition number. This can be explained by the fact that the angles and thus the condition number depend only on the quality of the normal vectors, which are very accurate due to the large number of points used to calculate them. On the other hand, the extent of the planar segments depends on the viewpoint and the associated occlusions. The smoothness is also less reliable, because it is sensitive to the range – more distant planes with the same surface properties are noisier. A possible solution would be to normalize the smoothness value with the mean range value of a plane. However this has not yet been tested, and is left for future work.

In Table 1 the empirical weights are shown which we have used in our experiments.

Property	Weight
Reciprocal Condition Value	10
Intersection Angle	100
Plane Extent	1
Plane Smoothness	5

Table 1: Weighting of the description vector

After applying the weights, Euclidean distances are used for descriptor matching, and all matches below a threshold are declared putative correspondences. Note that, contrary to most descriptor-based schemes, we only prune very dissimilar matches. We do not attempt to disambiguate matches below the threshold based on their descriptors, since in our experience the descriptor distances are too unreliable to do so and vary across different scanning environment. A variable threshold is used for accepting candidate matches, and adjusted such that the number of putative matches is not too high for the subsequent geometric verification.

5.3 Geometric constraint matching

The Euclidean distance between tie points is invariant across scans and can thus be used to reject false correspondences. Finding the largest possible subset of the reduced pool of matching candidates is a maximum clique problem, which has exponential complexity. Intuitively speaking, it can only be

solved to global optimality by (nearly) exhaustive search. At present we solve the problem with a greedy multi-start heuristic. Although it works well, we are planning to switch to a more sophisticated approximate optimization scheme in the future. The current scheme works as follows. For all pairs of candidate matches the point-to-point distances are computed in both scans and compared to a threshold (in the experiments 10 cm). Pairs which do not exceed the threshold are flagged as compatible, resulting in a symmetric binary matrix of pairwise compatibilities. Each candidate match in turn (i.e. each line in the matrix) is chosen as seed point, and all candidates are removed for which the point-to-point distances to the seed are incompatible. In the remaining candidate set the one with the smallest number of compatible matches (i.e., the smallest number of '1' entries) is removed, and that step is iterated, until no more incompatibilities remain.

The cliques of compatible matches for every seed are sorted by decreasing size (number of correspondences), and starting from the biggest one, the rigid transformation is estimated. If the residuals of the transformation are above a threshold, the clique is discarded and the next smaller one is tested, until a valid transformation has been found. The last step is required since testing pairwise distances does not account for long-range error accumulation. The output of the scheme are two sets of tie points in the two scans, which have the same (high) cardinality and the same geometric configuration, i.e. they can be matched without geometric contradictions.
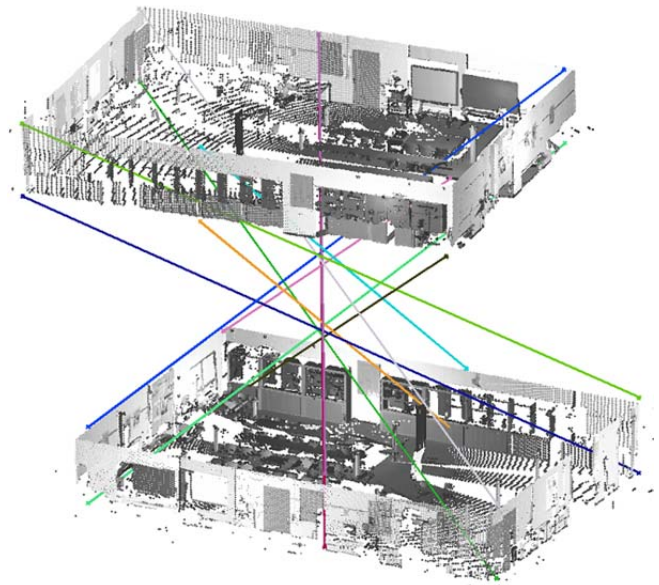

Figure 4: Set of matched tie points in two scans

Figure 4 shows a subset of matched tie points. While some points are purely virtual, others correspond to existing objects (e.g. room corners).

6. EXPERIMENTAL EVALUATION

6.1 Test Set Up

The proposed method has been tested with an indoor data set consisting of 4 scans. The rectangular room (~15 x 10 metres) has desks and tables, chairs, round pillars and some whiteboards on the walls. The scans have been acquired with the Zoller + Fröhlich TLS Imager 5006i and have a field of view of 360°

horizontal and 150° vertical. Each scan consists of 2.7 million points (2502 x 1076).

The following parameters were used: 5 pyramid levels for plane fitting; outlier threshold 1 cm; minimum plane cardinality 0.1% of all scan points; minimal reciprocal condition number to accept a tie point 0.1; maximum number of candidate matches after descriptor matching 5000 (corresponding to descriptor distances between 0.01 and 0.001); threshold for compatible geometric distances 10 cm; maximum allowable mean residual of rigid transformation 10 cm. Except when testing the noise sensitivity, these parameters were left unchanged.

6.2 Success rate

The success rate has been computed for all pairs of scans in the test data set, see Table 2. All experiments were repeated 50 times with the same parameters, because of the randomness induced by RANSAC plane fitting. The estimated transformation parameters were compared to manually registered ground truth.

Master \ Slave	1	2	3	4	Total
1		86%	80%	88%	87%
2	84%		100%	88%	81%
3	82%	100%		100%	94%
4	86%	88%	98%		91%
Total	84%	91%	93%	92%	**90%**

Table 2: Success rate of proposed method using test sets

As shown in Table 2, 90% of the scan pairs were correctly matched. Most of the mistakes are due to the high degree of symmetry of the scene: regarding only the main walls, several registrations are equally correct. The remaining dominant planes (e.g. tables) mostly reduce the ambiguity to two cases, which differ by a rotation of 180° around the vertical axis.

Further tests were performed to assess, how the success rate depends on the number of extracted planes. In Figure 5, the minimum, maximum and mean number of extracted planes in two scans is visualized. As expected, with more planes, the number of successful registrations increases, at the price of increased computational cost. Note the timing is based on an unoptimized Matlab implementation on a current 8-core machine including the whole working process (from import until the resulting transformation).
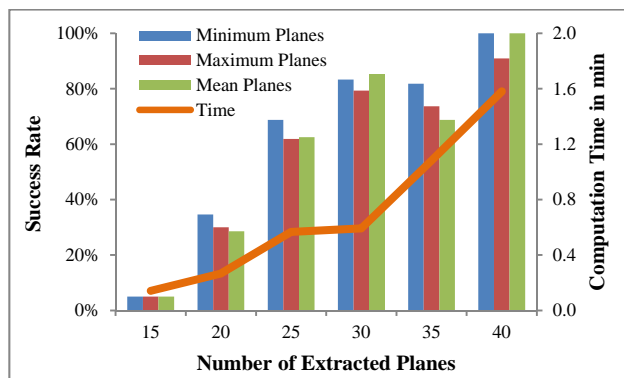

Figure 5: Correctness of registration and computation time for varying numbers of extracted planes

As an extreme case to test the limits of the method, the registration of an additional scan was attempted, taken from the ground, and thus leading to massive occlusions by the furniture. The success rate decreases to 39%, and the most dissimilar scan pairs could not be matched at all.

6.3 Sensitivity to Noise

Noise affects the proposed method only during plane extraction, since all subsequent steps are based on the planes and not on the 3D points themselves. To test the sensitivity to noise in the point cloud, we have repeated the test with different levels of synthetically added Gaussian i.i.d. noise (see Figure 6). In these experiments, the outlier threshold of RANSAC was set to 3 sigma. For completeness we also show the results with constants thresholds of 1 cm (low) and 10 cm (high).
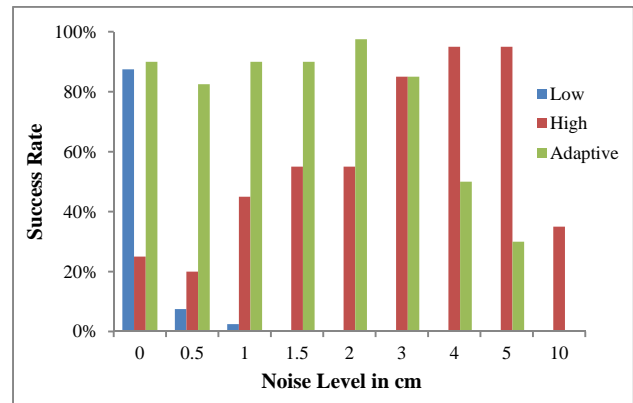

Figure 6: Successful registrations with different noise level regarding the outlier threshold

As expected, a low threshold quickly causes the method to break down. A higher threshold is less efficient in case of lower noise, but surpasses the lower and even the adaptive threshold with increasing noise (> 3 cm), since it has a higher chance of finding very noisy planes, even though the inlier/outlier separation might not be accurate. Overall, when the correct threshold (respectively the measurement uncertainty) is known, the results start to deteriorate for sigma bigger than 3 cm. Given that laser scanners typically have measurement uncertainties below 1 cm, the robustness of our algorithm should be sufficient for practical applications.

6.4 Contribution of Descriptors and Geometric Constraints

As already mentioned in Chapter 5.2 the number of candidate matches was limited (in our tests to 5000). The mean number of candidates during a correct registration process was 4000, of which in average 45 were accepted. The geometric matching thus can successfully discover correct cliques which cover only 1% of the candidates. Reducing the maximum candidate number results in lower success rates, by wrongly pruning correct correspondences. This confirms the claim that descriptor matching alone is not yet reliable enough.

7. CONCLUSIONS

Our proposed algorithm is able to perform coarse registration of two laser scans in a fully automated fashion, such that registration can be completed with a suitable algorithm such as ICP.

Our preliminary tests have shown the potential of the proposed method, although further improvements are necessary, especially regarding the efficiency of the descriptor. So far, only geometrical information is used to describe a tie point. We plan to also use intensity measurements of the scanner itself, and possibly also colour information of a camera with known orientation, to construct stronger descriptors. Also, splitting the matching into two sequential steps might allow one to better exploit weaker features like the plane extent, which at present have almost no influence due to their excessively low weight.

Furthermore we will also test the method in outdoor environments. We expect that extensions to other geometric elements may become necessary, since fewer planes are available outdoors.

REFERENCES

Asai, T., Kanbara, M. and Yokoya, N., 2005. 3D modeling of outdoor environments by integrating omnidirectional range and color images. In: Proceedings of the 5[th] International Conference on 3-D Digital Imaging and Modeling, Ottawa, Ontario, Canada, pp. 447–454.

Bae, K., Lichti, D., 2004. Automated registration of unorganised point clouds from terrestrial laser scanners. International Archives of Photogrammetry and Remote Sensing 35 (Part B5), pp. 222–227.

Barnea, S., Filin, S., 2008. Keypoint based autonomous registration of terrestrial laser point-clouds. ISPRS Journal of Photogrammetry and Remote Sensing, Theme Issue: Terrestrial Laser Scanning 63, pp. 19–35.

Basdogan, C., Oztireli, A.C., 2008. A new feature-based method for robust and efficient rigid-body registration of overlapping point clouds. The Visual Computer, vol. 24, 2008, pp. 7–9.

Bendels, G. H., Degener, P., Wahl, R., Körtgen, M., Klein, R., 2004. Image-based registration of 3D-range data using feature surface elements. In: The 5[th] International Symposium on Virtual Reality, Archaeology and Cultural Heritage, Eurographics, pp. 115–124.

Bergevin, R., Soucy, M., Gagnon, H., Laurendeau, D., 1996. Towards a general multi-view registration technique. IEEE Transactions on Pattern Analysis and Machine Intelligence, 18 (5), pp. 540–547.

Besl, P. J., McKay, N. D., 1992. A method for registration of 3-d shapes. IEEE Transactions on Pattern Analysis and Machine Intelligence 14(2), pp. 239–256.

Böhm, J., Becker, S., 2007. Automatic marker-free registration of terrestrial laser scans using reflectance features. In: 8[th] Conference on Optical 3D Measurement Techniques, Zürich, pp. 338–344.

Bornaz, L., Lingua, A., Rinaudo, F., 2002. A new software for the automatic registration of 3D digital models acquired using laser scanner devices. CIPA WG6 International Workshop for Cultural Heritage Recording, Corfu, September 1–2, pp. 52–57.

Brenner, C., Dold, C., 2007. Automatic relative orientation of terrestrial laser scans using planar structures and angle constraints. Proc. Laser Scanning 2007 and SilviLaser 2007, IAPRS Vol. 36 Part 3/W52, pp. 84–89.

Chen, Y., Medioni, G., 1991. Object modelling by registration of multiple range images. In: International Conference on Robotics and Automation, Vol. 3, pp. 2724–2729.

Dold, C., Brenner, C., 2006. Registration of terrestrial laser scanning data using planar patches and image data. International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences 36 (Part 5), pp. 25–27.

Fischler, M. A., Bolles, R. C., 1981. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. Communications of the ACM 24(6), pp. 381–393.

Franaszek, M., Cheok, G., Witzgall, C., 2009. Fast automatic registration of range images from 3D imaging systems using sphere targets, Automation in Construction 18 (2009), pp. 265–274.

Goshtasby, A. A., 2005. 2-D And 3-D Image Registration: for Medical, Remote Sensing, and Industrial Applications. Wiley–Interscience.

Gruen, A., Akca, D., 2005. Least squares 3D surface and curve matching. ISPRS Journal of Photogrammetry and Remote Sensing 59 (3), pp. 151–174.

He, W., Ma, W. and Zha, H., 2005. Automatic registration of range images based on correspondence of complete plane patches. In: Proceedings of the 5[th] International Conference on 3-D Digital Imaging and Modeling, Ottawa, Ontario, Canada, pp. 470–475.

Johnson, A. E., Hebert, M., 1999. Using spin images for fficient object recognition in cluttered 3D scenes. IEEE Transactions on Pattern Analysis and Machine Intelligence, 21 (5), pp. 433–449.

Kang, Z., Li, J., Zhang, L., Zhao, Q., Zlatanova, S., 2009. Automatic registration of terrestrial laser scanning point clouds using Panoramic Reflectance Images. Sensors 2009, *9*, pp. 2621–2646.

Lowe, D., 2003. Distinctive image features from scale-invariant keypoints. In: International Journal of Computer Vision, Vol. 20, pp. 91–110.

Masuda, T., Yokoya, N., 1995. A robust method for registration and segmentation of multiple range images. Computer Vision and Image Understanding, 61 (3), pp. 295–307.

Moldovan, D., Yano, S., Inoue, N., 2009. Automatic registration of multiple range views based on feature matching and feature-depth correspondences. IEEE 5[th] International Conference on Intelligent Computer Communication and Processing, Cluj-Napoca, August 27–29, pp 169–172.

Roth, G., 1999. Registration two overlapping range images. IEEE International Conference on 3D Imaging and Modeling, Ottawa, October 4–8, pp. 191–200.

Seo J. K., Sharp G. C., Lee S. W., 2005. Range Data Registration Using Photometric Features. IEEE Computer Society Conference on Computer Vision and Pattern Recognition, June 20-25, Vol. 2, pp. 1140–1145.

Shan, Y., Matei, B., Sawhney, H.S., Kumar, R., Huber, D., Hebert, M., 2004. Linear model hashing and batch RANSAC for rapid and accurate object recognition. IEEE Computer Society Conference on Computer Vision and Pattern Recognition, Vol. 2, pp. 121–128.