REGISTERING POINTCLOUDS OF POLYHEDRAL BUILDINGS TO 2D MAPS

K. Khoshelham, B.G.H. Gorte

Optical and Laser Remote Sensing, Delft University of Technology, 2629 HS Delft, The Netherlands.
{k.khoshelham, b.g.h.gorte}@tudelft.nl

Commission V, WG V/4

KEY WORDS: Terrestrial laser scanning, Automation, Building, Map, Referencing, PointCloud, Plane matching

ABSTRACT:

This paper presents a method for automated pointcloud-to-map registration using a plane matching technique. The registration is based on estimating a transformation between a set of planes inferred from the map and their corresponding planes extracted from the pointcloud. A plane matching algorithm is developed to search for corresponding planes in the pointcloud and map coordinate systems, and estimate the transformation between the corresponding planes. The search for correspondences takes advantage of an initiate-and-extend strategy to avoid high computational cost of an extensive search. The search strategy is further strengthened by using a linear model for the estimation of the transformation. The plane matching algorithm is shown to perform robustly in the presence of outlier and missing planes, and achieve accuracies in the order of centimetres as the mean distance between the transformed pointcloud planes and the map planes.

1. INTRODUCTION

Virtual reality models of urban environments are attractive means for municipalities and government bodies involved in town planning. A recent trend in the generation of virtual city models is the integration of panoramic images, terrestrial laser scanner data, and existing 2d maps. The integration of these data sources essentially involves a map-referencing step, where texture-mapped pointclouds of buildings are referenced to a large-scale map containing building boundaries. Automated methods for the registration of pointclouds to maps are of great interest as they can speed up the generation of virtual city models and reduce the costs.

The common procedure for transforming a pointcloud to a reference coordinate system is georeferencing by using GPS or other position and attitude measurement instruments. There are two main approaches to the georeferencing of the terrestrial laser scanner data: the indirect approach and the direct approach. In the indirect georeferencing approach, a minimum of three control points within the scanned area are signalized and measured in a reference coordinate system using a GPS receiver and/or a total station, and the position and attitude of the scanner with respect to the reference coordinate system is determined using a 3D resection method (Gordon and Lichti, 2004). This approach requires additional work during the scanning process for signalizing and measuring the targets, and further recognition of the targets in the data, which is difficult to automate. In the direct georeferencing approach, the position of the scanner is directly measured by a GPS receiver, and the attitude of the scanner with respect to a reference coordinate system is determined using a backsight target (Lichti and Gordon, 2004) or supplementary sensors such as a digital compass and a tilt sensor (Böhm et al., 2005; Schuhmacher and Bohm, 2005). The direct georeferencing approach also involves difficulties with the placement of the GPS receiver and the supplementary sensors.

In this paper, we propose a method for automated pointcloud-to-map registration using a plane matching technique. The advantage of using a map for georeferencing the pointclouds is that no additional workload is imposed during the scanning process. In addition, georeferencing methods can result in misalignments between the pointcloud and the map, which can be minimized when the pointcloud is directly registered to the map. Using planar surfaces as the corresponding entities in the registration leads to a linear least-squares model for the estimation of the transformation parameters. The benefit of a linear estimation model is that no initial approximations are required, and a fine registration can be performed in a single pass. Moreover, the linear model results in a more efficient search for correspondences.

The paper proceeds with an overview of the proposed method for pointcloud-to-map registration. The core of the method is the plane matching technique, which is described in Section 3. In Section 4, results of the experimental evaluation of the method are presented. Conclusions appear in Section 5.

2. POINTCLOUD-TO-MAP REGISTRATION: AN OVERVIEW OF THE PLANE MATCHING METHOD

The proposed method for pointcloud-to-map registration is based on the assumption that buildings are polyhedral objects, and thus, planar segments can be extracted from the pointcloud of a building. The registration is then formulated as a plane matching problem, where a transformation between a set of planes inferred from the map and their corresponding planes extracted from the pointcloud is estimated. The plane matching algorithm, therefore, performs two main tasks. The first task is to search for the largest set of corresponding planes in two sets of planes respectively in the scanner and the map coordinate system. The second task is to estimate a transformation between the corresponding planes, and to transform the pointcloud to the map coordinate system using the estimated transformation.

The search for corresponding planes is a crucial process and can be very expensive as the search space can grow very large when the number of planes increases. To allow an efficient search, we adopt an initiate-and-extend search strategy similar to that of SCERPO vision system (Lowe, 1987). In this strategy, initial correspondences are hypothesized and are then extended with new planes if an extension criterion is satisfied. The plane matching algorithm combines the search for correspondences and the transformation estimation in the following steps:

- Set an initial number of corresponding planes k;
- Create all combinations of k planes in the two sets that satisfy the constraints;
- For each set of k corresponding planes,
  - Insert the planes in an initial match set;
  - Estimate the transformation H;
  - Extend the match set with new planes that satisfy the extension criterion;
- The match set that has the maximum number of matching planes and minimum residuals is the best match.

In the following section a more detailed description of the processes within the plane matching algorithm is presented.

3. INFERRING A TRANSFORMATION FROM TWO SETS OF PLANES

Most of the existing registration methods are based on point-to-point correspondences (Salvi et al., 2007), although surfaces have also been used for establishing correspondence (Tarel et al., 1998). In this section, we show the advantage of plane-to-plane correspondences for developing a linear model to estimate the transformation.

3.1 Plane extraction

Planar surfaces can be extracted from laser scanner data, using range image segmentation methods (Hoover et al., 1996), or directly from the pointcloud, using clustering (Vosselman, 1999) or Hough transform (Rabbani Shah, 2006; Vosselman et al., 2004). In this work, we assume that a global registration of multiple scans of a building into a single pointcloud has been performed beforehand, and the registration parameters are available. To obtain the planar faces of the pointcloud, we use a gradient-based range image segmentation method (Gorte, 2007) to extract the planes from each scan. The extracted planes are then transformed to the pointcloud coordinate system by applying the registration parameters.

Using the building polygon in the map, planes are reconstructed as walls upon the polygon edges and a floor that contains the polygon. Assume that a polygon edge is given as a 2D line segment with endpoints $p_1 = (x_1, y_1)$ and $p_2 = (x_2, y_2)$. The direction vector of this line segment is defined as: $\boldsymbol{l} = (x_2 - x_1, y_2 - y_1, 0)^{\mathrm{T}}$. Let $\boldsymbol{\pi}$ be a plane defined by the direction of its normal vector $\mathbf{n} = (n_1, n_2, n_3)^T$, and the perpendicular distance d from the origin of the coordinate system. The vertical plane constructed on the line segment has a normal vector perpendicular to the direction vector of the line segment:

$$\mathbf{n} = \left(\frac{(y_2 - y_1)}{\|l\|}, \quad \frac{-(x_2 - x_1)}{\|l\|}, \quad 0 \right)^T \qquad (1)$$

where $\|l\|$ is the length of the line segment. The distance d from the origin to the plane is equivalent to the perpendicular distance from the origin to the line segment (Khoshelham and Li, 2004):

$$d = \frac{|x_1 y_2 - x_2 y_1|}{\|l\|} \qquad (2)$$

A horizontal floor plane is independent of the building polygon, and is defined as: $\mathbf{n} = (0, 0, 1)^{\boldsymbol{T}}$ and $d = 0$. Fig. 1 illustrates the construction of a vertical plane on a line segment.

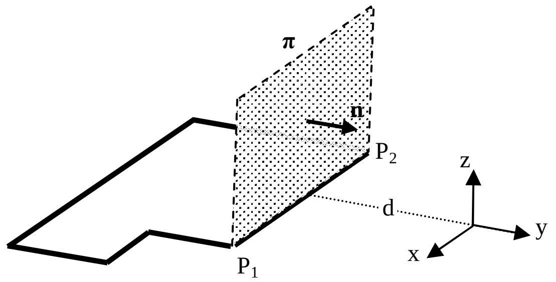

Fig. 1. Parameters of a vertical plane constructed upon a line segment of a building polygon.

3.2 Estimation of the transformation

As mentioned, a plane is identified with the direction of its normal vector and its distance to the origin. The condition that a point $\mathbf{x}$ lies on a plane $\boldsymbol{\pi}$ is expressed in homogenous representation as (Hartley and Zisserman, 2003):

$$\pi^T \mathbf{x} = 0 \qquad (3)$$

where $\boldsymbol{\pi}$ is $(n_1, n_2, n_3, -d)^{\mathrm{T}}$ and $\mathbf{x} = (x_1, x_2, x_3, 1)^{\mathrm{T}}$. The transformation of a set of n points from the scanner coordinate system s to the map coordinate system m is expressed as:

$$\mathbf{x}^m_{(4x\mathrm{n})} = \mathbf{H}_{(4x4)} \mathbf{x}^s_{(4x\mathrm{n})} \qquad (4)$$

where H is normally a similarity transformation when pointcloud-to-map registration is concerned:

$$\mathbf{H} = \begin{bmatrix} s\mathbf{R}_{3x3} & \mathbf{t}_{3x1} \\ 0^T_{1x3} & 1 \end{bmatrix} \qquad (5)$$

consisting of a scale s (normally s=1), a 3D rotation $\mathbf{R}$ and a translation vector $\mathbf{t}$. The corresponding transformation for transforming the plane containing the points from the scanner to the map coordinate system is:

$$\pi^m = \mathbf{H}^{-T} \pi^s \qquad (6)$$

because the points-on-plane condition is invariant to transformation, i.e. :

$$\pi^{m^T} \mathbf{x}^m = \left(\mathbf{H}^{-T} \pi^s \right)^T \mathbf{H}\mathbf{x}^s = \pi^{s^T} \mathbf{H}^{-1} \mathbf{H} \mathbf{x}^s = \pi^{s^T} \mathbf{x}^s = 0 \qquad (7)$$

If a plane i is available in the two coordinate systems, the plane matching equation can be derived from Eq. 6 as:

$$\pi_i^{s^T} = \pi_i^{m^T}\mathbf{H} \tag{8}$$

H can be estimated from a set of equations of the form given in Eq. (8) when a minimum of 3 corresponding planes in the scanner and the map coordinate system are available. To obtain an estimation of **H** that minimizes the norm $\left\| \pi_i^{s^T} - \pi_i^{m^T}\mathbf{H} \right\|$ we rearrange Eq. (8) as:

$$\pi_i^{m^T}\begin{bmatrix} \mathbf{h}_1^T \\ \mathbf{h}_2^T \\ \mathbf{h}_3^T \\ 0 \quad 0 \quad 0 \quad 1 \end{bmatrix} = \pi_i^{s^T} \tag{9}$$

where h_1^T, h_2^T, and h_3^T are the first three rows of the transformation matrix H. Stacking these in a vector yields:

$$\begin{bmatrix} n_{1i}^m\mathbf{I}_4 & n_{2i}^m\mathbf{I}_4 & n_{3i}^m\mathbf{I}_4 \end{bmatrix}_{4x12} \begin{bmatrix} \mathbf{h}_1 \\ \mathbf{h}_2 \\ \mathbf{h}_3 \end{bmatrix}_{12x1} = \begin{bmatrix} n_{1i}^s \\ n_{2i}^s \\ n_{3i}^s \\ -d_i^s + d_i^m \end{bmatrix}_{4x1} \tag{10}$$

Having n >= 3 corresponding planes, a set of linear equations is obtained, which can be expressed in matrix form as:

$$\mathbf{A}_{4\mathrm{n}x12} \cdot \mathbf{X}_{12x1} = \mathbf{B}_{4\mathrm{n}x1} \tag{11}$$

where X is the vector of unknown parameters of the similarity transformation H, A is the coefficient matrix formed by plane normals in the map coordinate system, and B is the vector containing plane parameters in both the scanner and the map coordinate system. The solution for **X** that minimizes $\left\| \mathbf{AX} - \mathbf{B} \right\|$ is given as:

$$\mathbf{X} = (\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T\mathbf{B} \tag{12}$$

Since the equations are linear no initial values for the unknown parameters are needed, and a solution is obtained in a single iteration.

3.3 Search for correspondences

The plane matching algorithm begins with creating all combinations of *k* planes in two sets of *m* and *n* planes respectively in the scanner and the map coordinate system. Starting with all combinations as initial match sets guarantees a robust performance of the algorithm as it ensures that always the best set of corresponding planes is found. However, depending on the values for *m*,*n*, and *k* the number of initial combinations can be very large. In principle, the number of possible ways to choose *k* corresponding planes from two sets of *m* and *n* planes is expressed as (Gorte et al., 2008):

$$n_c = C_k^m \cdot P_k^n = \frac{m!}{k!(m-k)!} \times \frac{n!}{(n-k)!} \tag{13}$$

where *C* and *P* denote combination (unordered selection) and permutation (ordered selection) functions respectively. Assume that the matching is initiated with combinations of $k = 4$ corresponding planes in $m = n = 10$ planes in the scanner and the map respectively. The total number of combinations will be in the order of 10^6, which is very large and results in a very expensive matching as each initial match set is extended within the matching process.

To make plane matching a practical method for pointcloud-to-map registration at a reasonable cost, reducing the number of initial combinations is inevitable. To this end, we impose a number of constraints to limit the search space. The first constraint is the inclusion of only near-vertical and near-horizontal planes in the initial combinations. This constraint is based on the assumption that the walls and the floor in the laser scanner data have relatively small deviations from vertical and horizontal planes respectively. This is a realistic assumption since in the terrestrial laser scanning the scanner is most often mounted on a tripod and is levelled using a bubble level.

The second constraint is the so-called relative orientation constraint (Gorte et al., 2008), which states that the relative orientation of plane normals in a set is invariant to rotation, and therefore, two sets of planes in the scanner and the map coordinate system are corresponding only if the relative orientation of the normals within the first set is the same as that of the second set. The concept of the relative orientation constraint is illustrated in Fig. 2.

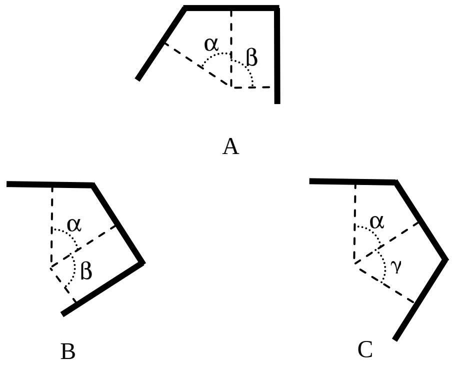

Fig. 2. Relative orientation constraint. Correspondence can be established between planes in set A and set B, but not between planes in set A and set C.

A third constraint can be applied when the number of corresponding plane pairs *k* in the initial combinations is larger than 3. In this case, a vector of residuals can be computed for each combination by estimating the transformation between the planes. The norm of this residual vector is an indication of how well the planes fit into the transformation model. A combination with non-corresponding plane pairs results in a large norm residual, and can be eliminated from the search.

Imposing the constraints described above greatly reduces the number of initial combinations (often by a factor of 10^3), which in turn leads to an efficient extension of the search. Once the initial combinations (match sets) are created, they are extended with new planes if an extension criterion is satisfied. Each match set is extended by taking a new plane from the pointcloud (that is not already in the initial match set) and transforming it to the map coordinate system (using the transformation estimated by the initial match set). Then in the map coordinate system the nearest plane (having the most similar parameters) to the transformed plane is found. This plane and the plane taken from the pointcloud are added to the match set and a new transformation is estimated. If the residuals obtained by this new transformation are within an acceptable small range, the match set is extended with the two new planes; otherwise, the planes are removed from the match set. The extension of each match set is continued until no new planes satisfy the extension criterion.

4. EXPERIMENTS

The plane matching method was tested with a simulated dataset as well as a real set of pointcloud and map data. The simulated pointcloud was generated using an airborne laser dataset of a building roof as a prototype. The ground plan of the building was manually drawn along the boundaries of the roof. Points on the walls and the floor were generated with a noise level of 20cm and a density of about 0.5m (equal to the density of the points on the roof). The simulated pointcloud was then transformed with an arbitrary set of rotation and translation parameters. A total of 21 planes were extracted from the transformed pointcloud, while from the map also 21 planes were derived. The planes were input to the plane matching algorithm, and the computed transformation was applied to the pointcloud. The result is shown in Fig. 3. As can be seen, the transformation parameters obtained by the plane matching method result in a registration of the pointcloud to the ground plan, where no visually noticeable misalignments can be observed.

As a measure of the accuracy of the registration, the mean and the maximum distance between the transformed pointcloud planes and the map planes were computed. Table 1 summarizes these accuracy measures for the simulated dataset. Since the pointcloud was simulated in this test, no superfluous or missing planes were present, and a relatively large proportion of the planes take part in the final match set. The mean distance of 27cm between the map planes and the transformed pointcloud planes is consistent with the level of noise which was added to the simulated points.

For the second test an H-shaped building within the campus of TU Delft was chosen (known as Logistic building), and a set of 7 scans of the building were obtained using a FARO LS880 laser scanner. A manual registration of the scans was carried out to create a pointcloud of the entire building. To obtain planes from this pointcloud, each individual scan was exported into a range image and a segmentation of each range image into planar segments was performed. These planes were then transformed to the coordinate system of the pointcloud using the registration parameters. A total of 32 planes from the pointcloud were taken for plane matching. Fig. 4 depicts the reflectance image and the segmentation of one of the scans.

The ground plane of the Logistic building was taken from the large-scale basis map of the Netherlands (GBKN). An interface was developed to allow the user to select the polygon of interest from the map (using a simple point-in-polygon operation). Fig. 5 (A) demonstrates the selection of the ground plan from the map. A total of 13 planes (1 floor and 12 walls) were derived from the ground plan of the Logistic building. The plane matching method was applied to these planes and those extracted from the pointcloud. The registration results are shown in Fig. 5 (B).

The set of planes extracted from the range images of the Logistic building contained a large number of superfluous planes, which had no correspondence in the ground plan. These are evident in Fig. 5 (B) as the walls of the smaller buildings (which were not of interest) around the main building. Also, what can be seen at the centre of the figure is an interior wall, which was recorded by the laser beams through the windows. In addition to these superfluous planes, a number of walls were completely covered by vegetation while a few others appeared to have very low point density due to their large distance from the scanner and the large incidence angle of the laser beams. This resulted in a number of missing planes in the pointcloud, for which corresponding planes were already derived from the ground plan. Despite the presence of the superfluous and missing planes, the results show a robust performance of the plane matching method.

Table1. Registration accuracy of the simulated pointcloud

	Nr of planes in: pointcloud	map	Final match set	Mean distance (m)	Max distance (m)
Simulated dataset	21	21	18	0.27	1.03

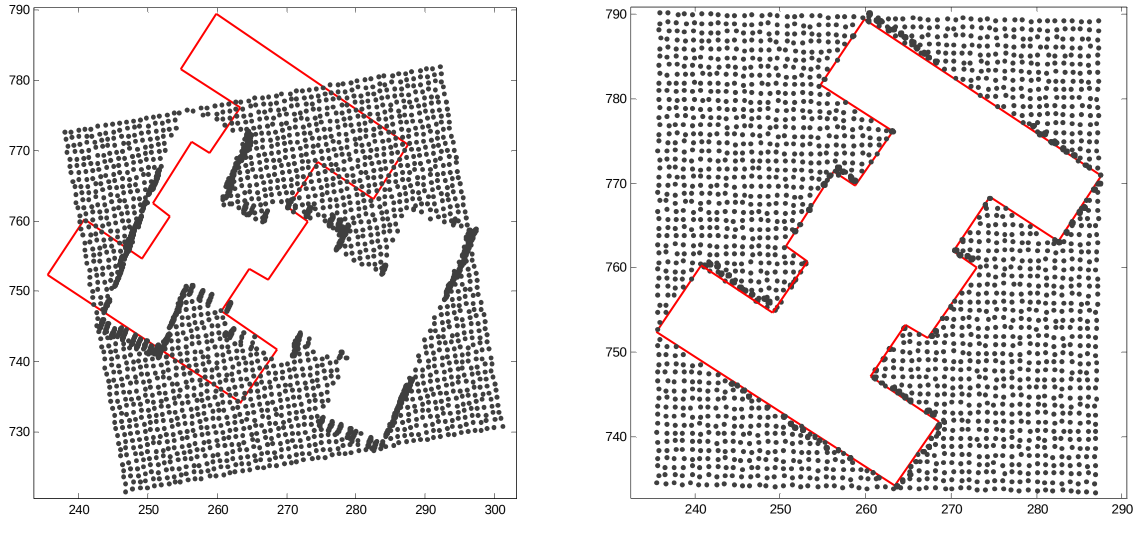

A B

Fig. 3. Automated registration of the simulated building pointcloud to its 2D ground plan. A. Top view of the pointcloud transformed with arbitrary parameters superimposed on the ground plan; B. The pointcloud is registered to the ground plan using transformation parameters obtained by plane matching.

Table 2 summarizes the accuracy measures obtained for the map-registration of the Logistic pointcloud. As can be seen, the mean distance and the max distance between the map planes and the transformed pointcloud planes are very close to those of the test with simulated pointcloud. Also, the number of planes in the final match set indicates that a relatively large proportion of the planes have contributed to the estimation of the registration parameters.

The efficiency of the method was evaluated by timing the process on a Pentium D CPU with 3.2 GHz speed and 2.00 GB memory. Table 3 shows the number of initial combinations and processing times for different initial settings. The processing times are directly related to the number of initial combinations, which is a function of the number of plane pairs in each initial combination (k) and the number of planes used to create the initial combinations (m,n). Obviously, by initiating the matching with a small subset of the planes one can reduce the number of initial combinations. However, the number of initial planes should be sufficiently large to ensure that at least one initial combination contains a set of correct correspondences. Experiments with various numbers of initial planes showed that a correct match will always be found if the matching is initiated with at least 7 initial planes. In practice, however, the number of initial planes should be increased with the complexity of the ground plan.

It is also worth noting that the efficiency of the method relies to a great extent to the application of the constraints. Fig. 6 illustrates the number of initial combinations with and without constraints. It is evident that plane matching without the constraints can be a very expensive process for complex buildings, and only after the application of the constraints the computational cost of the method becomes affordable for practical applications.

Table2. Registration accuracy of the real pointcloud

	Nr of planes in			Mean distance (m)	Max distance (m)
	pointcloud	Map	Final match set		
Logistic building dataset	32	13	10	0.34	1.03

A

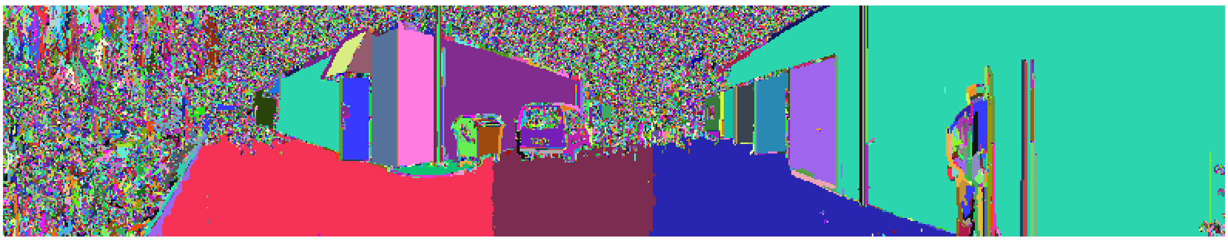

B

Fig. 4. One of the seven scans of the Logistic building. A. Reflectance image; B. Segmented range image.

Table 3: Number of initial combinations and processing times for different initial settings in the test with simulated data.

		Nr of planes to initiate matching	m=7 n=7	m=8 n=8	m=9 n=9	m=10 n=10	m=11 n=11	m=12 n=12
Number of initial combinations	k = 3	Without constraints	7350	18816	42336	86400	163350	290400
		Plane orientation constraint	582	1039	1859	3035	4838	7206
	k = 4	Plane orientation constraint + Transformation residual constraint	49	109	276	617	1661	3982
Processing time for k = 4 (seconds)			2.8	6.6	18.0	46.6	121.0	302.5

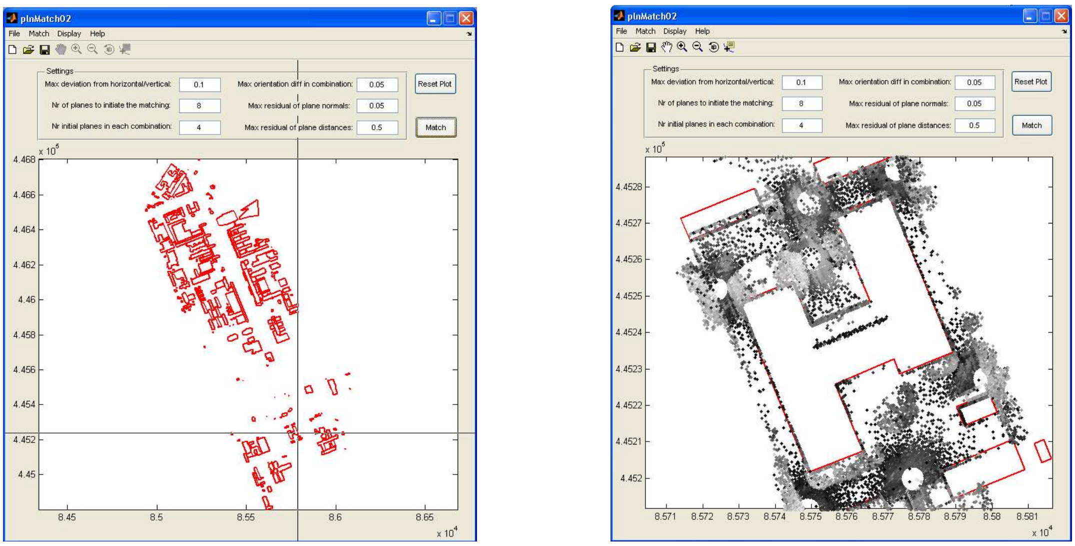

A **B**

Fig. 5. A. Automated registration of the pointcloud of the Logistic building to its 2D ground plan. A. Interface for selecting the ground plan from the map of the area; B. The pointcloud is registered to the ground plan using transformation parameters obtained by plane matching.

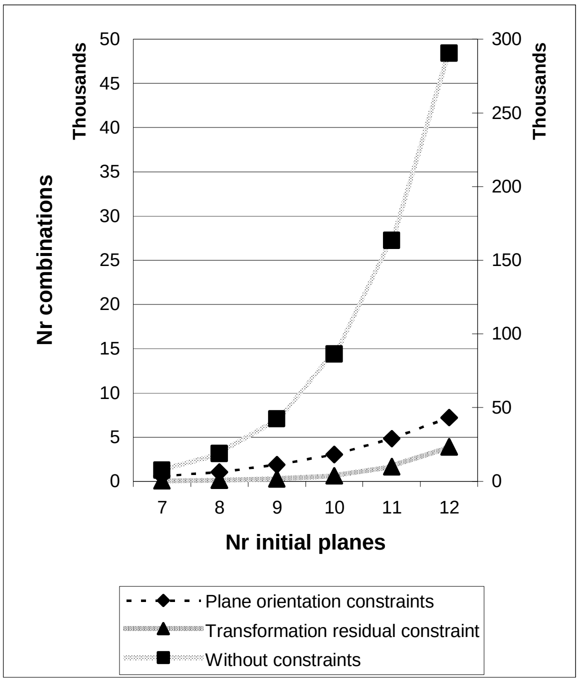

Fig. 6. A. Role of constraints in reducing the number of initial combinations.

5. CONCLUSIONS

In this paper, a plane matching method was presented for automated map-referencing of terrestrial pointclouds. The method was shown to perform robustly in the presence of outlier and missing planes, and achieve accuracies in the order of centimetres as the mean distance between the pointcloud and map planes. The computational expense of the automated process was also shown to be affordable for practical applications.

The advantage of using a map for the georeferencing of laser scanner data is that no additional workload is imposed during the scanning process. Using planes for registration leads to a linear least-squares model for the estimation of the transformation parameters, which is beneficial as it requires no initial approximations, and results in a more efficient search for correspondences.

The application of the plane matching method can be extended to pair-wise registration of multiple scans, provided that polyhedral objects are present in the overlapping area between the scans. It is also possible to combine the map-referencing process with global registration of multiple scans of a polyhedral building. The plane matching paradigm also provides a method for simultaneous scanning and navigation in indoor environments. Future research can be directed towards exploring the potential of the plane matching method in different applications.

ACKNOWLEDGEMENTS

The authors would like to thank Giorgi Mikadze for his assistance during the scanning and registration of the data. This research was performed as a part of the project **RGI-013** ‘Virtual reality for urban planning and security’ funded by the BSIK programme of the Dutch Government.

REFERENCES

Böhm, J., Haala, N. and Alshawabkeh, Y., 2005. Automation in laser scanning for cultural heritage applications, International Workshop on Recording, Modeling and Visualization of Cultural Heritage, Ascona, Switzerland, pp. 8.

Gordon, S.J. and Lichti, D.D., 2004. Terrestrial laser scanners with a narrow field of view: The effect on 3D resection solutions. Survey Review, 37(292): 448-468.

Gorte, B., 2007. Planar feature extraction in terrestrial laser scans using gradient based range image segmentation, ISPRS Workshop on Laser Scanning 2007 and SilviLaser 2007, Espoo, Finland, pp. 173-177.

Gorte, B., Khoshelham, K. and Verbree, E., 2008. Indoor navigation by using segmentation of range images obtained by laser scanners, International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences. Vol. XXXVII. Part B1., Beijing, China, pp. 971-976.

Hartley, R. and Zisserman, A., 2003. Multiple view geometry in computer vision, 2nd edition. Cambridge University Press, Cambridge, UK, 655 pp.

Hoover, A. et al., 1996. An experimental comparison of range image segmentation algorithms. IEEE transactions on pattern analysis and machine intelligence, 18(7): 673-689.

Khoshelham, K. and Li, Z.L., 2004. A model-based approach to semi-automated reconstruction of buildings from aerial images. Photogrammetric Record, 19(108): 342-359.

Lichti, D.D. and Gordon, S.J., 2004. Error Propagation in Directly Georeferenced Terrestrial Laser Scanner Point Clouds for Cultural Heritage Recording, FIG Working Week 2004, Athens, Greece, pp. 16.

Lowe, D.G., 1987. Three-dimensional object recognition from single two-dimensional images. Artificial intelligence, 31(3): 355-395.

Rabbani Shah, T., 2006. Automatic reconstruction of industrial installations using point clouds and images. PhD Thesis, Delft University of Technology, Delft, 154 pp.

Salvi, J., Matabosch, C., Fofi, D. and Forest, J., 2007. A review of recent range image registration methods with accuracy evaluation. Image and Vision Computing, 25(2007): 578-596.

Schuhmacher, S. and Bohm, J., 2005. Georeferencing of terrestrial laser scanner data for applications in architectural modeling. In: S. El-Hakim, F. Remondino and L. Gonzo (Editors), 3D-ARCH 2005: Virtual Reconstruction and Visualization of Complex Architectures, Mestre-Venice, Italy.

Tarel, J.-P., Civi, H. and Cooper, D.B., 1998. Pose Estimation of Free-Form 3D Objects without Point Matching using algebraic surface Models, Proceedings of IEEE Worshop Model Based 3D Image Analysis, Mumbai, India, pp. 13--21.

Vosselman, G., 1999. Building reconstruction using planar faces in very high density height data, ISPRS Conference on Automatic Extraction of GIS Objects from Digital Imagery, Munich, pp. 87-92.

Vosselman, G., Gorte, B.G., Sithole, G. and Rabbani Shah, T., 2004. Recognising structure in laser scanner point cloud, International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences, Vol. 46, Part 8/W2, Freiburg, Germany, pp. 33-38.