

AUTOMATIC DETECTION OF ZENITH DIRECTION IN 3D POINT CLOUDS OF BUILT-UP AREAS

Wolfgang von Hansen

FGAN-FOM, Gutleuthausstr. 1, 76275 Ettlingen, Germany
wvhansen@fom.fgan.de

KEY WORDS: Automation, Orientation, Algorithms, Urban, LIDAR, Point Cloud

ABSTRACT:

3D modeling of built-up areas has become an important issue during the past years. One technique for model creation is LIDAR. Terrestrial laser systems have become popular, but due to limited range and occlusions, many scanner positions are needed for data acquisition, requiring co-registration and fusion of the resulting models.

Two datasets can be co-registered by translation and rotation which are independent of each other. In this paper, we focus on the rotation and propose a method for the automatic detection of the zenith direction and thereby providing the ground plane. This yields an implicit surface for each dataset so that only one additional surface correspondence is needed to solve for the unknown rotation. The method exploits the geometric characteristics specific to built-up areas: Many vertical walls exist as well as a roughly horizontal ground surface. Results are presented for tests on a total of 26 datasets.

1 INTRODUCTION

1.1 Overview

3D modeling of built-up areas has become an important issue during the past years. One technique for model creation is laser scanning (LIDAR), which can be subdivided into airborne and terrestrial platforms. While airborne laser scanners cover large areas, their geometric resolution is low and only roof surfaces are well captured.

Terrestrial laser systems overcome these disadvantages, but due to limited range and occlusions, many scanner positions are needed for data acquisition, requiring co-registration and fusion of the resulting models. Many approaches exist for an automatic coarse registration of the scan data. (Akca, 2003, Bae, 2006, Dold, 2005, Dold and Brenner, 2006, von Hansen, 2006, He et al., 2005, Liu and Hirzinger, 2005, Rabbani and van den Heuvel, 2005, Ripperda and Brenner, 2005, Wendt, 2004) The fine registration is usually solved through the ICP (*iterative closest point*) algorithm (Besl and McKay, 1992, Rusinkiewicz and Levoy, 2001).

Segmentation techniques include clustering based on local surface normal analysis (Bretar and Roux, 2005, Liu and Hirzinger, 2005), region growing using scan geometry and point neighborhoods (Dold and Brenner, 2004), and a split-and-merge scheme applying an octree structure (Wang and Tseng, 2004). The objects are often represented by planar elements recovered through RANSAC schemes (Bretar and Roux, 2005) or least squares adjustment (Wang and Tseng, 2004).

Two datasets can be co-registered by one translation and one rotation which both are independent of each other. Three corresponding surfaces are sufficient for translation and two for rotation. However, as there is only a small overlap between two models and many surfaces are only partially seen due to occlusions, it is not easy to determine correct correspondences.

In this paper, we focus on the rotation and propose a method for the automatic detection of the zenith direction and thereby the ground plane. As this yields an implicit surface for each dataset, only one pair of corresponding surfaces is needed to solve for the unknown rotation between two datasets. As this can be

done independently for each dataset, this is an important aid for the coarse registration process. Our method exploits geometric characteristics specific to built-up areas: Many vertical walls exist as well as a large and roughly horizontal ground surface.

1.2 Methodology

According to (Dold and Brenner, 2004), both components of registration – rotation and translation – can be carried out independently. When planar surfaces have been extracted from the point clouds, two corresponding planes define the rotation about all three axes and three corresponding planes define the translation in space. In fact, three correspondences would solve the complete problem, but if there is only a little overlap between two datasets, many of the possible combinations are false.

Therefore, a bottom up approach seems more feasible. If a part of the problem can be solved more easily, it leads to additional information that can be used in following steps. This paper proposes an automatic orientation of a single point cloud from outdoor scenes in built-up areas, so that the zenith direction is aligned with the z -axis, i. e. is pointing upwards. This is done separately for each single point cloud – therefore no correspondence search is necessary at this stage. After orientation, only a single plane correspondence is needed to solve for rotation.

In a preprocessing step, the raw point cloud is converted into a set of plane elements. Then, a RANSAC based scheme is applied to their normal vectors, partitioning them into ground, walls and other elements. This allows a robust computation of the zenith direction. As by-product, the segmentation into these three classes is returned.

We have not found any reference to such single dataset orientations during our literature research.

1.3 Geometrical axioms

Our approach is based on some simple axioms on planar structures that are valid for built-up areas. These axioms give constraints that allow to segment planes into ground, building walls and other objects:

1. The ground surface is large compared to other surfaces.

Ambiguity interval	53.5 m
Resolution range	16 Bit 1 mm/lb
Range noise 10 m	1.3–3.0 mm rms
Range noise 25 m	3.0–9.0 mm rms
Laser output power	23 mW (red)
Beam divergence	0.22 mrad
Field of view vertical	310°
Field of view horizontal	360°
Resolution vertical	0.018°
Resolution horizontal	0.01°
Accuracy	0.02° rms
Number of pixels in dataset	200 million

Table 1: Specifications of the Z+F Imager 5003 laser scanner.

2. The normal vector of the ground surface is approximately – but not exactly – orthogonal to the horizontal plane and pointing roughly to the zenith direction.
3. Building walls have a normal vector that lies exactly in the horizontal plane.
4. Roof surfaces are inclined and point upwards.
5. Other surfaces, like e. g. trees or small structures, have normal vectors pointing to random directions.

The distinction between roofs and other surfaces is not a simple decision, as it requires statistical knowledge about the surroundings of a position. Here, we will not discriminate between these two types and regard roofs as other surfaces.

The objective to find the zenith direction is replaced by the detection of the horizontal plane. From the axioms we see two complementary hints at the correct horizontal plane: First, there are many normal vectors lying in this plane. Their orientation/direction is not important. Every building wall – which we assume vertical – will contribute to this criterion. Second, the normal vector of the ground already is an estimate for the zenith direction. But we have to acknowledge, that these estimates may contain systematic errors because the terrain may be sloped even in built-up areas. Even the steepest roads typically have an inclination of less than 20% which is roughly 11°. This can be used as cut-off for ground surfaces. A terrestrial system is usually placed below the roof, so that flat roofs normally will not appear in the datasets and therefore can not be confused with the ground.

1.4 Laser scanner and test data

The laser scanner used for data acquisition is a Zoller+Fröhlich Imager 5003 – some of its technical data is given in Tab. 1. It has an omnidirectional field of view, so that there is no tilt mechanism necessary to point the laser scanner to the object of interest. Instead, the scanner is leveled prior to use by means of a circular level. This way, the true zenith direction is contained in the data and can be utilized for the evaluation of our results in the sense that gross errors can be detected. It should be emphasized, that this information is *not* exploited by the algorithm. The datasets that we dispose of are 26 overlapping outdoor scans from a village scene.

2 METHODOLOGY

2.1 Generation of planar surface elements

As a preprocessing step we estimate locally delimited plane elements from the point cloud as proposed by (von Hansen, 2006).

Pos	planes	Pos	planes
1	6841	14	3510
2	5163	15	4955
3	3741	16	4072
4	5065	17	4982
5	3685	18	4360
6	2616	19	4995
7	2955	20	5503
8	4629	21	7078
9	5498	22	5184
10	4726	23	3290
11	7160	24	3567
12	4769	25	2859
13	5117	26	9971

Table 2: Number of planes extracted for each position.

The measured point cloud is split into a regular 3D raster and the plane that is supported by most of the points in each raster cell is computed in a robust way. Large object surfaces are thereby split into several coplanar elements. An alternate approach would be to try to reconstruct the complete surfaces, but the advantage of many small planes is the introduction of an implicit weighting factor: As large object planes will be divided into many plane elements, they will contribute a lot to the result.

We have implemented the generation of surface elements via the well known RANSAC strategy (Fischler and Bolles, 1981): From three randomly chosen, non collinear points the uniquely defined plane is computed. Then for all points, their distance to the plane is computed, counting those with a distance below a certain threshold as inliers. This procedure is repeated, finally returning the plane parameters for the plane with the largest inlier count. A plane can be represented by the Hesse normal form

$$\mathbf{n}^T \mathbf{x} + d = 0 \quad (1)$$

where \mathbf{n} is the plane’s normal vector, d its distance to the origin and \mathbf{x} the set of all points on the plane. Since we are only interested in the rotation, it is sufficient to keep just the normal vectors. The set of all normal vectors \mathbf{n}_i of a dataset will be denoted as \mathcal{N} .

2.2 Zenith direction detection

We will present a two step approach to the determination of the zenith direction. The first is to determine an approximate zenith direction from axioms 1 and 2, that is to find a maximum number of parallel normal vectors. The second step is the determination of the exact zenith direction from axioms 2 and 3: Using the approximate zenith direction, the normal vectors of the walls are identified and from these the final solution is estimated.

2.2.1 Robust detection of approximate ground plane From axiom 1 we assert that the majority of the data is representing the ground surface and from axiom 2 we see that its normal vector already is a first estimate for the solution. The task therefore is to find a maximum number of parallel normal vectors. For fast computation, a RANSAC scheme is used.

A single normal vector \mathbf{n}_i is randomly drawn from \mathcal{N} and tested with all other vectors $\mathbf{n}_j \in \mathcal{N}$ for parallelity to generate the i -th inlier set

$$\mathbf{N}_i := \{\mathbf{n}_j \in \mathcal{N} : \mathbf{n}_i \parallel \mathbf{n}_j, j = 1 \dots |\mathcal{N}|\} \quad (2)$$

where \parallel denotes parallelity of two vectors which is implemented as

$$\mathbf{a} \parallel \mathbf{b} :\Leftrightarrow \mathbf{a}^T \mathbf{b} = \cos \alpha \stackrel{!}{\geq} \theta_{\parallel} \quad (3)$$




Figure 1: 3D model of position 22, showing the segmentation into different surface types.

where \mathbf{a} and \mathbf{b} are normalized vectors and θ_{\parallel} is a predefined threshold parameter for parallelity given by a maximum angle α between \mathbf{a} and \mathbf{b} . The vector $\hat{\mathbf{n}}$ yielding the largest inlier set

$$\hat{\mathbf{N}} := \arg \max(|\mathbf{N}_i|) \quad (4)$$

is kept as the approximate zenith direction \mathbf{z}_0 .

2.2.2 Estimation of exact zenith direction An approximate zenith direction \mathbf{z}_0 has already been determined as shown in the previous section. The ground surface usually is not horizontal but sloped, so that \mathbf{z}_0 is not guaranteed to be correct. A precise solution can be retrieved by taking into account that most walls are built vertically and therefore their normal vector can be utilized to compute a result (axiom 3).

The set of normal vectors of the walls \mathcal{W} can be determined through axiom 1 as the set of all vectors roughly orthogonal to \mathbf{z}_0

$$\mathcal{W} := \{\mathbf{n}_i \in \mathcal{N} : \mathbf{n}_i \perp \mathbf{z}_0, i = 1 \dots |\mathcal{N}|\} \quad (5)$$

where \perp denotes orthogonality of two vectors which is implemented as

$$\mathbf{a} \perp \mathbf{b} :\Leftrightarrow |\mathbf{a}^T \mathbf{b}| = |\cos \alpha| = |\sin \beta| \leq \theta_{\perp} \quad (6)$$

where \mathbf{a} and \mathbf{b} are normalized vectors and θ_{\perp} is a predefined threshold parameter for orthogonality given by a small angle β for the maximum difference to a right angle. θ_{\perp} should not be chosen too small because we have to account for sloped terrain.

All vectors $\mathbf{w}_i \in \mathcal{W}$ lie in the horizontal plane, typically pointing to different directions. The final zenith direction $\hat{\mathbf{z}}$ is orthogonal to all \mathbf{w}_i and can be computed as the null space of the matrix \mathbf{W} composed from all vectors of \mathcal{W} via singular value decomposition.

There is a sign ambiguity in the returned vector as both zenith and nadir direction are valid solutions. This can easily be corrected by changing the sign if necessary to ensure that $\hat{\mathbf{z}}$ points into the same direction as \mathbf{z}_0 .

2.2.3 Segmentation With the known zenith direction $\hat{\mathbf{z}}$, an additional result can be obtained without much further computation. After rotation of the dataset such that $\hat{\mathbf{z}}$ points upwards, all normal vectors are put into some classes based on the axioms and their inclination with respect to the horizontal plane. This can help in the construction of a 3D model from the data.

3 EXPERIMENTS AND RESULTS

The laser scanner is already aligned to the horizontal plane prior to operation so that the “true” zenith direction is already known, however this knowledge is *not* exploited by the algorithm. The advantage here is, that the correctness of a result can immediately be checked. As a by-product, it is possible to get an idea on the quality of the instrument’s leveling which had been carried out only roughly using the circular level of the laser scanner.

All datasets have been converted to plane elements using a distance threshold of 3 cm. Tab. 2 shows the number of planes generated for each position. The different scene complexities are mainly due to the presence of natural objects as many planes are generated to describe e. g. the volume of a tree. Especially position 26 is located at the end of the village where several trees were in the range of the scanner.

An example for a resulting dataset is presented in Fig. 1, showing a view across a larger town square. The buildings at the far end are located on top of a small hill. The different classes have been color coded in the images. Green represents the ground, cyan marks building walls and the roofs are colored in red. Other surfaces are shown in light gray. Some of the wall elements are shown in gray which is due to a tight threshold and some surface roughness. The texture as acquired by the laser scanner is mapped onto the surfaces. The square pattern that can be seen on all object surfaces is a result of boundary effects from the spatial data partitioning into cubes.

The algorithm for the estimation of the zenith direction had been applied successfully to all datasets. This experimentally proves that the assumptions defined through the axioms were valid:

- The ground surface is represented by the largest group of parallel normal vectors. Other object planes like building facades may have a combined area that is larger, but their normals are pointing to different directions.
- In real data, there is enough variation in the wall's normal vectors to define the zenith direction. In a perfectly U-shaped street canyon, the wall's vectors would be collinear with an ambiguous solution \mathbf{z}_0 for the zenith direction. In this rare case, the approximate solution from the first step would still be available. In addition, it could also be used to check for such situations.

Some numerical results are given in Tab. 3. The column "Inliers" refers to the number of normal vectors of the ground plane with respect to the total number of normal vectors. This directly gives the percentage of the ground surface area in the scene. On the average, about 20% of structures in the acquired scenes are ground surface. In a certain way, this contradicts axiom 1, because on the average, only one fifth, and at maximum only one third, of the normal vectors are from the ground. But obviously, these small but systematic areas are sufficient to outweigh wall or roof areas. The worst case is position 11, where only one ninth of the scene is detected as ground.

The next column lists the number of iterations required for the RANSAC scheme. These numbers are computed from the inlier rate such that there is a 99.9% probability that one correct sample had been drawn. This is the reason for the strong correlation between inlier rate and number of iterations. The computation itself is sufficiently fast – less than 0.5 s on an ordinary computer – as the test for one iteration step amounts to n scalar products to check for parallelity where n is the number of planes as shown in Tab. 2.

The last column shows the angular deviation given in mrad between the scanner's zenith direction from the leveling and the estimated zenith direction. The values are distributed uniformly leading to the interpretation that both the leveling and the algorithm delivered good results. The maximum deviation is 41.2 mrad which amounts to 2.36° . Results of the segmentation are shown as different colors in Fig. 1 and Fig. 2. The cyan color denotes those wall elements that are used for the computation of the zenith direction. Other surfaces that have not been put into a specific class are colored in light gray and could be considered as *rejects*. The majority of the walls has been segmented correctly, while there are still some surface elements that have been rejected. This is acceptable as the threshold had been chosen rather tight here in order to include only good wall elements.

Red is assigned to roof surfaces, but there are also other surfaces assigned to this class. The main reason is that the allowed inclination for the roof surface is a large interval. As the absolute height above the ground or neighborhoods have not been considered, other upwards pointing surfaces have been marked as roofs as well. However, all real roof elements have been assigned to the correct class.

The ground is colored in green and has been segmented very well.

4 SUMMARY AND CONCLUSIONS

We have shown a new method that determines the zenith direction in a 3D point cloud of a built-up area. In a two step procedure, the raw point cloud is first transformed into a set of locally delimited object plane elements. An initial zenith direction is then recovered from the set of plane normal vectors through a RANSAC

Pos	Inlier/%	Iterations	Angle/mrad
1	28.8	22	2.2
2	27.5	23	9.0
3	23.1	28	10.0
4	25.8	24	3.9
5	23.7	27	5.5
6	14.5	45	1.5
7	15.5	42	6.8
8	12.0	50	23.9
9	22.9	28	17.0
10	13.7	48	39.6
11	10.7	61	34.3
12	20.1	32	24.7
13	21.7	29	14.2
14	17.5	37	13.6
15	22.9	28	17.4
16	22.4	28	7.7
17	22.6	27	2.7
18	18.9	34	13.9
19	21.1	30	5.9
20	22.5	28	12.4
21	19.4	33	41.2
22	29.1	21	5.7
23	27.4	23	12.4
24	32.6	19	13.6
25	26.9	23	18.7
26	21.2	30	2.6

Table 3: Results zenith directions.

scheme that clusters parallel vectors in order to detect the ground plane. From this, the walls can be identified and used to estimate the final zenith direction as the vector orthogonal to their normal vectors.

Results have been shown for a total of 26 datasets. Being successful for all positions shows that the automatically recovered ground plane is correct. The small angular deviations are a combined quality measure for both algorithm and the scanner's circular level. No gross errors exist so that the automatically found walls can be deemed suitable for zenith direction estimation.

The main objective was to detect a global plane normal from a single 3D point cloud as one corresponding plane for the relative rotation of two datasets. For all of the available datasets, the proposed method was successful, making it useful for pre-rotation of point clouds in order to help coarse registration of multiple positions. Results also show that the leveling of the laser scanner had been carried out very precisely, so that the upwards direction of the dataset could already be regarded as zenith direction.

REFERENCES

- Akca, D., 2003. Full automatic registration of laser scanner point clouds. In: *Optical 3-D Measurements VI*, pp. 330–337.
- Bae, K.-H., 2006. Automated Registration of Unorganised Point Clouds from Terrestrial Laser Scanners. PhD thesis, Curtin University of Technology. URL: adt.curtin.edu.au/theses/available/adt-WCU20060921.094236/.
- Besl, P. J. and McKay, N., 1992. A Method for Registration of 3-D Shapes. *PAMI* 14(2), pp. 239–256.
- Bretar, F. and Roux, M., 2005. Hybrid Image Segmentation Using LIDAR 3D Planar Primitives. In: G. Vosselman and C. Brenner (eds), *Laser scanning 2005*, IAPRS, Vol. XXXVI-3/W19. URL: www.commission3.isprs.org/laserscanning2005/papers/072.pdf.




Figure 2: 3D model of position 25.

Dold, C., 2005. Extended Gaussian Images for the Registration of Terrestrial Data. In: G. Vosselman and C. Brenner (eds), Laser scanning 2005, IAPRS, Vol. XXXVI-3/W19. URL: www.commission3.isprs.org/laserscanning2005/papers/180.pdf.

Dold, C. and Brenner, C., 2004. Automatic Matching of Terrestrial Scan Data as a Basis for the Generation of Detailed 3D City Models. In: O. Altan (ed.), Proc. of the XXth ISPRS Congress, IAPRS, Vol. XXXV-B3. URL: www.isprs.org/istanbul2004/comm3/papers/429.pdf.

Dold, C. and Brenner, C., 2006. Registration of terrestrial laser scanning data using planar patches and image data. In: H.-G. Maas and D. Schneider (eds), Image Engineering and Vision Metrology, IAPRS, Vol. XXXVI Part 5. URL: www.isprs.org/commission5/proceedings06/paper/DOLD_637.pdf.

Fischler, M. A. and Bolles, R. C., 1981. Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography. *Comm. of the ACM* 24(6), pp. 381–395.

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

Liu, R. and Hirzinger, G., 2005. Marker-free Automatic Matching of Range Data. In: R. Reulke and U. Knauer (eds), Panoramic Photogrammetry Workshop, IAPRS, Vol. XXXVI-5/W8. URL: www.informatik.hu-berlin.de/sv/pr/PanoramicPhotogrammetryWorkshop2005/Paper/PanoWS_Berlin2005_Rui.pdf.

Rabbani, T. and van den Heuvel, F., 2005. Automatic point cloud registration using constrained search for corresponding objects. In: 7th Conference on Optical 3-D Measurements.

Ripperda, N. and Brenner, C., 2005. Marker-free Registration of Terrestrial Laser Scans Using the Normal Distribution Transform. In: S. El-Hakim, F. Remondino and

L. Gonzo (eds), 3D-ARCH 2005, IAPRS, Vol. XXXVI-5/W17. URL: www.commission5.isprs.org/3darch05/pdf/33.pdf.

Rusinkiewicz, S. and Levoy, M., 2001. Efficient variants of the icp algorithm. In: Proceedings of the Third Intl. Conf. on 3D Digital Imaging and Modeling, pp. 142–152.

von Hansen, W., 2006. Robust automatic marker-free registration of terrestrial scan data. In: W. Förstner and R. Steffen (eds), Photogrammetric Computer Vision, IAPRS, Vol. XXXVI Part 3. URL: www.isprs.org/commission3/proceedings06/singlepapers/O_08.pdf.

Wang, M. and Tseng, Y.-H., 2004. LIDAR Data Segmentation and Classification Based on Octree Structure. In: O. Altan (ed.), Proc. of the XXth ISPRS Congress, IAPRS, Vol. XXXV-B3. URL: www.isprs.org/istanbul2004/comm3/papers/286.pdf.

Wendt, A., 2004. On the automation of the registration of point clouds using the metropolis algorithm. In: O. Altan (ed.), Proc. of the XXth ISPRS Congress, IAPRS, Vol. XXXV-B3. URL: www.isprs.org/istanbul2004/comm3/papers/250.pdf.

