SEGMENTATION OF LASER ALTIMETER DATA FOR BUILDING RECONSTRUCTION: DIFFERENT PROCEDURES AND COMPARISON

Roland GEIBEL, Uwe STILLA
FGAN-FOM Research Institute for Optronics and Pattern Recognition
D 76275 Ettlingen, Germany
{gei,usti}@fom.fgan.de

Working Group III/3

KEY WORDS: Segmentation, Laser, Building, Evaluation, Surface Reconstruction.

ABSTRACT

One important step in the reconstruction of buildings from laser altimeter data is the segmentation step. In this contribution four procedures known from literature are investigated and compared. One procedure, which was made for detecting straight lines, which produces a segmentation as an intermediate result, was used. Three other procedures were installed, which are based on region growing. For each of the procedures experiments with an extensive variation of the parameters were carried out. For the particular task of roof reconstruction and the speciality of height data an own procedure was developed. In this procedure the distance measure is not computed from derived properties but directly from the sets of points. Moreover it also uses the knowledge about the separation of background and foreground and it is regulated by only one parameter. In order to make the segmentation results comparable known comparison measures for single segments (the relative correspondence related to the model and the result) were taken. For the estimation of the whole segmentation several aspects of the evaluation of single segments are discussed and a complete evaluation function is developed. For a quantitative measurement some buildings were chosen as examples.

1 MOTIVATION

1.1 City models

Airborne height data recorded with laser scanners are gaining increasing importance in generating three dimensional city models. There is a particular demand for such models in visualisation, mission planning, disaster management and as a base for simulations, e.g. the area of environment protection and telecommunication. In industrialised countries in recent years many maps were recorded digitally and they are more and more available as vector maps. However large scale topographic maps and cadastral maps show only the footprints of buildings without information about the height or the roof shape. So far information about the height of buildings was derived by manual measurement or from stereo image pairs.

At the present time height data of airborne laser scanner systems are commercially available (e.g. [Lohr, 1998] [Huising & Gomes Pereira, 1998]). The sensed surface points, scattered over a stripe of 250-500m width, allow the production of a geocoded 2D-raster with height data for each cell (height image). Single flight strips are combined into a consistent digital height model of the area investigated.

A main component of city models is the vector description of the buildings with their roofs. For analysing the roof shapes the height data belonging to each building are masked using a digital map. For the reconstruction of the roof shape the segmentation of the height data as the first processing step plays an important role. For the segmentation of range data, several procedures have already been developed. In a former paper procedures for laser scanner images and structured light images of some polyhedral bodies were investigated in a laboratory-like environment. It still is an unsolved problem, which of these procedures is most appropriate for the task presented here and the specific data of natural scenes with artefacts.

1.2 Related Work

Using height images produced by laser scanner data, a big part of the roof shapes usually found in urban areas can be reconstructed. Thus [Stilla et al., 1998] dealt with the reconstruction of different roof shapes (sloped and flat, with and without superstructure). Therein a building with sloped and flat roof parts could be reconstructed using a segmentation of the height image and of the gradient image. Thereon it was used that roofs of rectangular buildings usually have

planar parts which are oriented in few directions. This assumption allowed to use orientation histograms for segmentation.

If also objects shall be recognised which do not fulfil this assumption, then more general segmentations must be considered. Since in general an essential step in reconstructing objects from height data consists in a good segmentation of the height image into geometrically well matching surface parts, other procedures for the segmentation of range images known from literature were investigated. Yet it shows that the question of what is a good segmentation and what is a good segmentation procedure is too general in this wording. In the restricted area of the segmentation of laser altimeter images, we present a further development of the evaluation measure defined in [Hoover et al., 1996].

2 SEGMENTATION PROCEDURES

In the contribution presented here, four procedures for the segmentation of range images, known from literature are investigated and compared. Three region growing based procedures which were available via the WWW were installed. Two procedures are based on an iterative region growing from seed regions. One procedure uses a clustering approach. An other procedure which serves for extracting straight lines yields as an intermediate result also a region segmentation. Furthermore an own simple region growing procedure was developed. It is regulated by only one parameter and in addition produces the neighbourhood graph of the segments. In each of the procedures the result of the segmentation is a label image and for each segment the description of the plane by a plane equation.

2.1 Burns

In the procedure of Burns et al [1986] straight lines are extracted from an intensity image by investigating the gradient. Straight edges have in their neighbourhood areas with similar gradient orientation, which is perpendicular to the edge. Thus first connected regions of equal gradient orientation are searched. In height images planar areas show gradients of equal size and direction. Applying the procedure of Burns to height data, the partitioning of the inspected image into connected areas of equal orientation can also be considered as a segmentation into regions.

2.2 The UB-procedure (University of Bern)

In the forefront of the procedure by Jiang and Bunke [1994], [Hoover et al., 1996] there is the line-like grouping of points in range images recorded in rows. For this purpose first each row of the image is recursively split into most possible straight line segments, then a region growing is performed on these line segments instead of the single pixels. As seed regions for the region growing process those triples of neighbouring line segments are used for which (i) all three line segments have at least a requested minimal length, (ii) the intersection of two adjacent line segments has a certain percentage of the total length of each of the three line segments and (iii) each pair of adjacent points on two line segments does not exceed a certain distance threshold.

Among all of these possible seed regions the one with the highest total length is chosen as an optimal seed region. Then other line segments are added to this region if the distance of both of the endpoints of the line segment from the plane of the growing region remains inside a threshold. This process of adding a further line segment is repeated until no further line segments can be added. Then the next best seed region is used to start the same process. This is repeated until all seed regions are processed. In a post-processing step small regions are eliminated.

2.3 The WSU-procedure (Washington State University)

The procedure of [Flynn and Jain, 1991], [Hoffman and Jain, 1987] is not restricted to objects with planar surfaces, but can also describe elliptic surfaces (spheres, cylinders, cones). It proceeds in the following 9 steps: (i) First jump edge pixels are identified by z-coordinate differences in the 8-neighbourhood. (ii) For the pixels distant enough from these jump edge pixels the surface normal is estimated. (iii) For the set of pixels with surface normal a clustering is performed (hoffcluster). This clustering produces connected sets of similar orientation. (iv) The pixels belonging to the same cluster are labelled identically. (v) Investigating the labelled sets with respect to their local neighbourhood segments are produced. (vi) In order to counteract over-segmentation an edge-based method for joining regions is applied. Adjacent regions are joined together, if the difference in angle of all normals along the common border line is small enough. (vii) Each segment is checked for planarity by a regression procedure. This is the first time that plane equations are calculated for the segments. Non planar segments are taken out of further consideration. (viii) Neighbouring planar segments are merged if their parameters "angle of surface normal" and "distance term" of the plane equation are similar. (ix) Unlabeled pixels at the border of each segment are added to the segment if their distance to the plane lies inside a threshold. The steps (vii) to (ix) are repeated until the result is stable. In a following post-processing step small regions are deleted.

2.4 The USF-procedure (University of South Florida)

The procedure of Hoover et al. [1996] proceeds in two steps: First for each pixel the normal vector is estimated, then the region growing itself starts. For the investigation a window of size at least 5x5 is laid around each pixel. For the estimation of the normal for each pixel a sub-window is made centred around the pixel and eight other sub-windows in horizontally, vertically and diagonally shifted directions. For all of these sub-windows the normal vector and a residual is computed. The normal vector with the smallest residual is taken as the normal vector of the centre pixel. For pixels in the interior of a segment this residual will be small. Noise or a near by edge produce a big residual. Thus a small residual can be a hint for the interiorness inside a planar segment.

The pixel with the smallest residual is taken as seed point for the region growing process. Criteria for the addition of a further pixel to a growing region are (i) the four-neighbourhood, (ii) the angle between the pixels normal and the region grown so far, (iii) the distance between the new pixel to be added and the plane equation of the region grown so far and (iv) the point-to-point distance between the pixel and its neighbours in the scene. This way a region grows pixel-wise at its border. If no more pixels can be added to a region the next best pixel available (based on its residual) is taken as the seed pixel for the next region. Too small regions are discarded and their pixels play no further role.

2.5 Qualitative valuation of the procedures by an example

For a first qualitative valuation the procedures were installed and tested with the data in hand. For that purpose it was necessary to modify source code and the data to make the data types, the data formats and the different image formats compatible. Several interfaces for camera models present in the original source code could not be used. The results of the segmentation by the different procedures are displayed in Fig. 1.


Fig. 1: Results of the segmentation procedures: a) Burns, b) UB, c) WSU, d) USF

The result of the procedure Burns (Fig. 1a) shows a heavy sectioning of the areas, caused by the sensitivity of the gradient to noise by the small 2x2 window. Since the procedure Burns makes the segmentation only based on gradient orientation adjacent areas of same orientation but different slope can not be discerned. Since gradients are only considered if they have a certain absolute size horizontal areas can not be detected either. Since the aim is to find edges the decision about intersecting concurrent segments is determined by their length and not the number of pixels which would be desirable. Remarkable in Fig. 1b) are the horizontally directed segments, resulting from a supposed line structure of the scanning system which however was not present. The result in Fig. 1d) shows extraordinarily many small segments in the border area and a very big part at the border of the segmented object and between the segments considered as unsure. This is due to the operating mask being small with respect to the resolution of the data. Since the results produced with the data present seem not satisfactory a further segmentation procedure was developed and implemented.

2.6 The FOM procedure

This procedure is based on the stepwise growing of similar areas from single pixels up to big regions. Beginning with a masking made in advance the knowledge about the separation into background and foreground is used. The procedure is marked out by a special distance measure and a continuing recalculation of the plane equations. Thereon the distance measure is not determined from derived features of the segments but directly from their pixel sets. For the calculation of the distance measure between two adjacent segments P, Q the Euclidean distance d_E is determined for all pixels p of segment P to the plane of the other segment Q. The distance d_1 of one segment P to Q is calculated as

$$d_1(P,Q) := \max_{p \in P} \{d_E(p, \text{Plane}(Q))\} \qquad (1)$$

whereas the symmetric distance measure between P and Q is taken to be the minimum of both values:

$$d (P,Q) := \min (d_1 (P,Q), d_1 (Q,P)) \tag{2}$$

Then that pair of segments Po, Qo is fused which has the minimal distance:

$$d(Po,Qo) = \min \{d(P,Q) \} . \tag{3}$$

Pairs of adjacent segments with minimal distance are fused until this minimal distance d becomes bigger than a threshold d_{th}. This threshold d_{th} is the only parameter to control the procedure. During the initialisation phase the plane equations are determined either by all pixels of the 8-neighbourhood or from the best fit with 3 adjacent pixels which form a square with the centre pixel. In a post-processing phase very small segments like for example single border pixels are added to an adjacent segment. The result of the procedure is the partitioning of the image into disjoint segments and the neighbourhood graph of the segments.


a   b   c   d

Fig. 2: Stepwise fusing of segments after a) 2690, b) 2860, c) 2940, and d) 3050 steps

The four images in Fig. 2 show intermediate results after 2690, 2860 and 2940 iterations and the final result after 3050 iterations. Since each iteration fuses at least one pair of segments there can be no more iterations than the number of segments in the initial phase, i.e.: the number of pixels of the object. The number of comparisons in each iteration is bounded by the number of pixels of the object too. The total effort of the procedure is of quadratic order with respect to the number of pixels.

3 TEST ENVIRONMENT

The valuation of one or more algorithms requires to process relevant data and to evaluate the results. In order to process and valuate a big number of data sets it is suggested to automate the components of the test environment as far as possible. To gain a qualitative statement about a procedure task oriented measures must be defined [Hoover et al., 1996].

3.1 Data

3.1.1 Input data

In general we discriminate test data by their origin into (i) data of synthetic objects, (ii) data of model objects and (iii) data of real objects. In this contribution height data of real objects are used which were recorded by a laser scanner. Nowadays airborne laser scanner systems yield sensing ratios of one point per $6m^2$ up to 9 points per 1 m^2 [Huising & Gomes Pereira, 1998]. In this investigation we used data which were recorded with a ratio of 4 points per 1 m^2. They were used to derive a geocoded equidistant grid which can be depicted as an image.

3.1.2 Ground truth data

Ground truth data describe an object in a scene or in image data. If scene descriptions are used as ground truth data, then a comparison will also contain the special properties of the sensor. Thereon it may be, that objects of the scene are not visible in the image because of some properties of the sensor. If the task is to investigate only the procedure alone, then the excerption of ground truth data from the image data is suggested. This case is called sensed truth [Klausmann et al., 1999]. While for input data of synthetic objects and model objects ground truth data can often be obtained automatically input data of real objects usually require to derive the ground truth data manually. For valuating the segmentation of height data plane areas were extracted manually following the visible shape, with areas down to 4 square meters being recorded. Fig. 3 shows an example for a manual ground truth segmentation.


Fig. 3: Ground truth

3.2 Performance measurement

3.2.1 Measures for comparing ground truth and machine segments

In the following segments of the sensed truth data are denoted as T-segments and segments which are produced by a segmentation procedure, are denoted as machine segments (M-segments). In the ideal case the segmentation procedure produces one M-segment for each T-segment and both segments are exactly equal. Since the borderlines can not be found exactly because of distortions in the real case the T- and M-segments will only intersect partly even if the relation is one-by-one. To meter the correspondence the intersection area of T-segment and M-segment is taken as measure. In order to denote the intersection ratio independently of the size of the intersection area it is normed by the size of the segment. This produces the measures

$$S_T := A_S/A_T \quad \text{and} \quad S_M := A_S/A_M \qquad (4)$$

which are related to the area of the T-segment A_T and of the M-segment A_M.

Correct classification. If the intersection ratio of M- and T-segment is very low, then it should not be assumed that they are corresponding. As a criterion for a correct classification a minimal relative intersection

$$S_T > S_{th} \quad \text{and} \quad S_M > S_{th} \qquad (5)$$

is required. To ensure that a T-segment can correspond to only on M-segment as correct, $S_{th} > 0.5$ is required. The example in Fig. 4a shows with $S_T=15/24=0.625$ and $S_M=15/20=0.75$ and thus a correct classification, while in Fig. 4b $S_T=8/24=0.333$ and $S_M=8/20=0.4$ and thus the segments are not classified as correct.


a

b

Fig. 4: One M-Segment.
a) correct, b) not correct

Over-segmentation. It is possible that the expected T-segment is represented by several parts of M-segments. The intersection area A_S then results from the total area of the M-segments ($S = T \cap (M1 \cup M2 ...$). The relative intersection is defined analogously to equation (4) as

$$S_{TO} := A_S/A_T \quad \text{and} \quad S_{MO} := A_S / A_{M1 \cup M2} ... \qquad (6)$$

If more than one M-segment intersects with the T-segment, then only one M-segment can be marked as classified correctly. If no M-segment was assigned as corresponding correctly, then a T-segment is classified as over-segmented, if

$$S_{TO} > S_{th} \quad \text{and} \quad S_{MO} > S_{th} \qquad (7)$$

hold. If an M-Segment was assigned a correct classification, then the T-segment is classified as over-segmented, if in addition the average of the measures of the new correspondence is bigger than that of the measures of the one M-segment classified as correct.

$$\frac{S_{TO} + S_{MO}}{2} > \frac{S_T + S_M}{2} \qquad (8)$$

If the condition in equation (8) is not satisfied, then the classification of the one correct segment will be kept and the other M-segments are not associated to the T-segment.


a

b

c

d

Fig. 5: Two M-segments. a) not correct, not over-segmented b) not correct, over-segmented c) correct, over-segmented, d) correct not over-segmented

Fig. 5 shows cases, in which the T-segment is intersected by two M-segments. In Fig. 5a-b the M-segments are so small, that in both cases the M-segments could not be classified as correct. In Fig. 5a the T-segment is neither over-segmented, since both of the M-segments together do not fulfil the condition of equation (7), while in Fig. 5b it holds. In Fig. 5c-d M2 was assigned to the T-segment as a correct classification. Since by the position of M1 in Fig .5c

equation (7) and (8) are fulfilled, the T-segment is denoted as over-segmented. In Fig. 5d M1 is intersected by T only to a little part, such that equation (8) does not hold and only M2 is assigned to T as correct. T is not over-segmented.

Under-segmentation. If one M-segment intersects with more than one T-segment, we can expand our definition of the measures S_T and S_M to:

$$S_{TU} := A_S / A_{T1 \cup T2} \quad \text{and} \quad S_{MU} := A_S / A_M \quad (9)$$

Like in the case of over-segmentation a segment could already be classified as correct, now one of the T-segments can already be classified as correct or as over-segmented. If no classification is assigned and if

$$S_{TU} > S_{th} \text{ and } S_{MU} > S_{th}, \quad (10)$$

holds, then the M-segment is classified as under-segmented. If however one of the T-segments was already classified as correct or over-segmented, then in addition to equation (10) the average of the measures for the under-segmentation must be bigger than the average of the measures for the former classification (cp. equation (8)). In Fig. 6 M is first assigned to T1 as a correct classification, then to T1$\cup$T2 as under-segmentation.

Missed and Noise. T-Segments, which then are neither classified as 'correct', nor as 'over-segmented' nor as 'under-segmented' are classified as missed (Fig. 7a). M-segments, which then are neither classified as 'correct', nor as 'over-segmented' nor as 'under-segmented' are classified as noise (Fig. 7b).

3.3 Segmentation of multiple areas

The above classification is carried out for all T-segments and all M-segments of an object (e.g.: a building). Thereon one T-segment can be related to many M-segments and one M-segment can be related to many T-segments. For instance the classification of the segmentation result shown in Fig. 2d (43 segments) based on the ground truth with 27 segments (Fig. 3) results in 10 correct T- and M-segments, 5 over-segmented T-segments, 4 under-segmented M-segments, 4 missed T-segments and 8 noise M-segments. In investigating multiple objects of a scene for a valuation of different procedures, such a classification is obtained for each object.

3.4 Valuation of the segmentation

In order to receive a quality measure for single objects (e.g.: buildings) or a complete scene (e.g.: city model), the measurements and classifications of single object parts (roof parts) must be evaluated. A simple measure of quality can be derived from the number of correctly found segments $|T_K|$ or from its ratio to all T-segments $|T_A|$.

$$q = |T_K| / |T_A| \quad (11)$$

However since the intersection with the T-segments can differ even if they have been classified as found correctly the measure S_T should be taken into account as a weight factor for each of the T-segments.

$$q = \sum_{T \in T_K} S_T / |T_A| \quad (12)$$

Since the segments which were classified as over-segmented or even those which were classified as under-segmented can be used for further processing and make a contribution to the segmentation quality as compared to missed or noise segments they should be recognised in calculating the quality measure.


Fig. 6: Two T-segments and one M-segment (correct, under-segmented)

Fig. 7 a) T-segment (missed), b) M-segment (noise)

$$q = (\sum_{T \in T_K} S_T + \sum_{T \in T_O} k_O \cdot S_T + \sum_{M \in M_U} k_U \cdot S_{T_l \cup \ldots \cup T_M}) / |T_A| \qquad (13)$$

For the over-segmented and for the under-segmented segments the ratio of intersection S_{TO} and S_{TU} are taken into account as weight factors too. However their contribution is reduced by the additional factors $k_O < 1$ and $k_U < 1$. The influence of over-segmented and under-segmented segments should be different. For an over-segmentation it is possible that in one of the following processing steps M-segments are combined if they belong together. For an under-segmentation a subsequent splitting based on the features of the fused segment may be difficult. For visualisation tasks also the effect of an over-segmentation is not as bad as that of an under-segmentation if the normal vectors of the segmented parts were determined correctly. As weight factors for k_O and k_U we propose:

$$k_O = (2n - 1)/n^2 \qquad \text{with n = number of M-segments associated with the T-segment} \qquad (14)$$

and

$$k_U = 1/m^2 \qquad \text{with m = number of T-segments associated with the M-segment.} \qquad (15)$$

The segments classified as missed worsen the quality q because they make no contribution. Segments classified as noise worsen the result and therefore should be considered in reducing the quality. For example the sum in the numerator of equation (12) can be decreased for the number $|M_N|$ of missed segments. However since the size of the segments was not yet recognised a few small noise segments would have much more effect than one big segment. In general also the size of segments should be recognised in their importance for the total area. As a measure of quality we propose:

$$q_{L0} = (\sum_{T \in T_K} S_T \cdot A_T + \sum_{T \in T_O} k_O \cdot S_T \cdot A_T + \sum_{M \in M_U} k_U \cdot S_{T_l \cup \ldots \cup T_M} \cdot A_{T_l \cup \ldots \cup T_M} - \sum_{M \in M_N} A_M) \bigg/ \sum_{T \in T_A} A_T \qquad (16)$$

$$q_L = \max (0, q_{L0}) \qquad (17)$$

For the evaluation of a complete scene the quality number q_G has to be derived from the results of all of the segments of the scene or from the results q_L of the objects. For the evaluation of a scene we use

$$q_G = \sum_{\text{local objects}} q_L \cdot A_L \bigg/ \sum_{\text{local objects}} A_L \quad \text{with} \quad A_L = \text{area of the local object} \qquad (18)$$

4 EXPERIMENTS

4.1 Quantitative evaluation of the procedures by an example

In order to test and evaluate the segmentation procedures laser altimeter data of real objects in the area of the town of Karlsruhe were used. The resolution of the data is 1 m^2 on ground and the resolution in altitude is 0.06 m. The scenes were recorded in "first-pulse-mode". For adjusting the segmentation procedures many experiments were made for each of the procedures. The parameter adjustment was done for one selected object of the test area. The results of the investigated procedures with the parameters valued as best are depicted in Fig. 1 and 2d. The evaluation of the results is shown in table 1. The quality measure q_L was determined using the weight factors k_U and k_O proposed in equation (14) and (15).

Number of segments	Burns	USF	UB	WSU	FOM
correct	8	7	1	5	10
over	1	3	10	3	5
under	5	0	3	3	4
missed	2	17	9	2	4
noise	7	9	4	5	8
q	0.33	0.54	0.38	0.23	0.93

Table 1: Number of classified segments by different procedures

4.2 Irritability by rotation

It would be supposed that a segmentation procedure produces the same output if the same input data is presented in different order. Practice shows however that rotation of the data leads to more or less different outputs in many of the considered procedures. To test this dependency the images were turned, segmented and turned back. Fig. 8b-d shows the results with the input image turned for 90°, 180° and 270°. The images show small variations of the contours of small segments and the appearing or vanishing of small segments. Bigger segments stay stable.

FOM rotation	0	90	180	270
correct	10	11	13	14
over	5	4	5	6
under	4	5	4	2
missed	4	2	1	3
noise	8	10	5	4
q	0.93	0.93	0.94	0.93

Table 2: Number of classified segments with rotation of the image

Fig. 8: Rotation of the height image
a) 0°, b) 90°, c) 180°, d) 270°

By a rotation of the same image data for 90, 180 or 270 degrees each of the investigated procedures showed more ore less changes in the segmentation. For the procedure Burns the rotation had the least effect.

5 DISCUSSION

During the experiments there arose difficulties in the determination of optimal parameters. In the procedures USF and WSU this was caused by the big number of parameters and their interdependencies. For this reason it could not be ensured that the used parameter set was optimal. Some of the investigated procedures assume different sensor specific recording geometries which however were not present in the data used. In the procedure UB the assumed line-like sensing order lead to a stripe-like over-segmentation. The own procedure produced the best results. However it also

Fig. 9: Segmentation results of the FOM procedure with variation of the parameter d_{th}
a) section of aerial image, b) height image, c) d_{th}=0.45 m, d) d_{th}=1.0 m, e) d_{th}=1.5 m, f) d_{th}=2.5 m

needed most computation time. The evaluation function used was weighted subjectively. Thereon an over-segmentation was tolerated more than an under-segmentation. Depending on the task given and the further processing expected other orders are possible.

In general because of the low horizontal resolution of 1m on ground not all of the details of the buildings (e.g.: dormers, small building parts) could be segmented. Just buildings which have many such structures on their roof (e.g. dormers), which disturb the course of a plane, cause a specific problem. Fig. 9a shows such a building in an aerial image. The segmentation result of the FOM procedure with the parameter adjustment as used in Fig. 2d is depicted in Fig. 9c. The segments keep so small that no connected area can be recognised. Varying the parameter d_{th}, as shown in Fig. 9d-f , bigger deviations are tolerated in the plane approximation and bigger segments are produced. If however the same parameter adjustment as used in Fig. 9f were used for all objects, then small details of buildings get lost for other objects (see Fig. 2d). In principle also the quality of the data (noise, filtering by pre-processing) could be taken into account during the parameter adjustment.

For the man made objects considered here it was assumed that the roof areas are plane. Curved areas which appear on other man made objects (e.g. industrial complexes, fuel depots) were not respected. Since the investigation was done with the data of only one flight, for a validation of the results data sets of other flights at different seasons and with different sensors are to be examined. In order to judge the problems in the reconstruction of whole buildings caused by faulty segmentations it is attempted to also evaluate the resulting vector descriptions comparing to a ground truth in a test bed. Further work will also deal with the segmentation of unmasked height data.

REFERENCES

Amerinex Applied Imaging, Inc. (1996) The KBVision system task reference manual, Release 3.2,

Burns JB, Hanson AR und Riseman EM (1986) Extracting straight lines. IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI), 8(4):425-455.

Flynn PJ, Jain AK (1988) Surface classification: Hypothesis testing and parameter estimation. Proc. IEEE CS Conf. Computer Vision and pattern recognition (CVPR '88), 261-267,

Flynn PJ, Jain AK (1991) BONSAI: 3D object recognition using constrained search. IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI), 13(10):1066-1075.

Goldgof DB, Huang DW, Lee H (1989) A curvature-based approach to terrain recognition. IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI), 11(10):1213-1217.

Hoffman RL, Jain AK (1987) Segmentation and classification of range images. IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI), 9(5):608-620.

Hoover A, Jean-Baptiste G, Jiang X, Flynn PJ, Bunke H, Goldgof DB, Bowyer K, Eggert DW, Fitzgibbon A, Fisher RB (1996) An experimental comparison of range image segmentation algorithms. IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI), 18(7):673-689.

Huising EJ, Gomes Pereira LM (1998) Errors and accuracy estimates of laser data acquired by various laser scanning systems for topographic applications. ISPRS Journal of photogrammetry and remote sensing, 53:245-261

Jiang XY, Bunke H (1994), Fast segmentation of range images into planar regions by scan line grouping. Machine vision and applications, 7(2):115-122

Klausmann P, Fries S, Willersinn D, Stilla U, Thönnessen U (1999) Application-oriented assessment of computer vision algorithms. In: Jaehne et al. (eds) Handbook of computer vision and applications, 3:133-152

Lohr U (1998) Laserscan DEM for various applications. In: Fritsch D, Englich M, Sester M (eds) GIS - Between vision and applications. 353-356

Stilla U, Geibel R, Repasi E (1998) Generation and update of 3D-city models from maps and elevation data. In: Ikeuchi K (ed): Proceedings of MVA'98, 42-45.

Stilla U, Soergel U, Jaeger K (2000) Generation of 3D-city models and their utilisation in image. IAPRS Vol XXXIII Amsterdam