

AUTOMATIC GENERATION OF BUILDING MODELS FROM LIDAR DATA AND THE INTEGRATION OF AERIAL IMAGES

F. Rottensteiner¹ and Ch. Briese²

¹School of Surveying and Spatial Information Systems, The University of New South Wales, Sydney, Australia
f.rottensteiner@unsw.edu.au

²Institute of Photogrammetry and Remote Sensing, Vienna University of Technology, Austria
cb@ipf.tuwien.ac.at

KEY WORDS: 3D building models, Building extraction, Data acquisition, LIDAR processing, Multi-sensor processing.

ABSTRACT

In this paper, we present a method for the automated generation of 3D building models from directly observed point clouds generated by LIDAR. First, building regions are detected automatically. After that, roof planes are detected by applying a curvature-based segmentation technique. These roof planes are grouped in order to create polyhedral building models. In this grouping process, the shapes of the roof plane boundaries are determined, and the whole model is improved by an overall adjustment using all available sensor information and, if possible, geometric constraints. We describe the existing modules and those still in the implementation phase, and we discuss the issue of integrating aerial images into the reconstruction process to increase the geometric quality of the reconstructed models. Intermediate results are presented for a test site located in the City of Vienna, Austria.

1 INTRODUCTION

1.1 Motivation and Goals

The progress in sensor technology for airborne laser scanning has rendered possible the acquisition of very dense point clouds, so that the generation of 3D building models from LIDAR data has become feasible. It is possible not only to detect buildings and their approximate outlines, but also to extract planar roof faces and, thus, to create models that correctly resemble the roof structure. In many approaches for building extraction from LIDAR data, ground plans of the buildings are required to give the precise location of the building outlines (Haala et al., 1998, Brenner, 2000, Vosselman and Dijkman, 2001). If no such data are used, the algorithms rely on regular structures such as orthogonal walls, thus restricting their applicability (Vosselman, 1999).

It is the goal of this paper to present a method for the automatic creation of polyhedral building models in densely built-up areas from high-resolution LIDAR data without using ground plans or geometric restrictions with respect to the mutual relations of neighbouring planes. It is an improved version of the method described in (Rottensteiner and Briese, 2002). Roof planes are extracted by a hierarchic segmentation technique based on the statistic behaviour of the normal vectors of a digital surface model (DSM) computed from the LIDAR data. The neighbourhood relations of these roof planes are detected, and the roof polygons are computed by analysing the mutual configurations of pairs of roof planes. We present a method for the consistent estimation of all model parameters by hybrid adjustment, optionally introducing geometric constraints for a regularisation of the resulting model. We also plan to integrate digital aerial images into the process to improve the geometric quality of the building models. The examples presented here were computed using the LIDAR data captured by TopoSys in the course of a pilot project for the Municipality of Vienna (Figure 1a). Though our method is still work in progress, our preliminary results, obtained in a densely built-up area characterised by irregular building shapes with non-rectangular walls and roofs that are not symmetric, show its potential.

1.2 Related Work

It is the first step in building extraction to detect building candidate regions in the sensor data. This is essentially a classification task, commonly solved by applying morphologic filters for computing a digital terrain model (DTM) and applying a threshold to the height differences between the DSM and the DTM (Weidner, 1997). The initial classification has to be improved in order to remove vegetation areas. In (Rottensteiner and Briese, 2002), we have presented an algorithm for building detection from LIDAR points based on DTM generation by hierarchical robust linear prediction (Briese et al., 2002) and the evaluation of a texture roughness criterion for the elimination of vegetation (Figure 1).

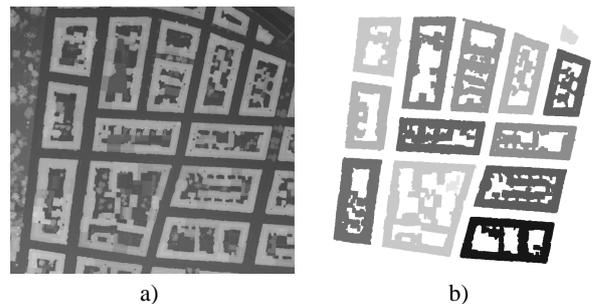


Figure 1: Building detection in a test site in Vienna. Original resolution: 0.1 m (in-flight) by 0.9 m (cross-flight). a) DSM ($0.5 \times 0.5 \text{ m}^2$; extent: $410 \times 435 \text{ m}^2$). b) Building regions.

The geometrical reconstruction of the buildings in previously detected regions of interest has been tackled in two ways. First, parametric primitives can be instantiated and fit to the data if sufficient evidence is found, e.g. in rectangular regions derived by splitting a 2D ground plan (Brenner, 2000, Vosselman and Dijkman, 2001). Second, planar segments can be detected by a segmentation of a DSM created from the LIDAR points, e.g. by an analysis of surface curvature (Brenner, 2000), and polyhedral building models can be derived by grouping these planar segments. As the building outlines are difficult to be located precisely, ground plans are often used, too (Haala et al., 1998, Brenner, 2000, Vosselman and Dijkman, 2001). (Weidner, 1997)

has tackled the problem of precise determination of the building outlines from LIDAR data by applying the minimum description length principle for deciding on regularisations.

As soon as planar segments have been detected, neighbouring segments are grouped, and the polygons delineating the borders of these segments have to be found (Baillard et al., 1999). These border polygons have to be combined in order to obtain consistent building models, and the models are completed by vertical walls and by a floor. A coarse-to-fine strategy can be applied by first searching for the most relevant structures in the data and using refined methods for modelling the buildings in regions not being “explained” sufficiently by the initial models (Vosselman and Dijkman, 2001). (Vosselman, 1999) describes a method for the geometric reconstruction of buildings from LIDAR data alone. He relies on assumptions on the orthogonality of walls to determine the building outlines.

It is another possibility to improve the results of building extraction to combine LIDAR data with aerial images. (Schenk and Csatho, 2002) give a method for fusing LIDAR data and aerial images to obtain a more complete surface description without explicitly reconstructing buildings. (Ameri, 2000) propagates a segmentation in image space to object space using LIDAR data, and he matches the resulting object edges with image edges to improve their geometric quality. A similar technique for wire frame fitting has been used for measurement of parametric primitives in semi-automatic building extraction (Rottensteiner, 2001). It is general enough to be applicable to any polyhedral model.

1.3 Overview of the Proposed Method

In this paper, we only deal with the geometric reconstruction of buildings. As far as the LIDAR data are concerned, except for parameter estimation, we use a regular DSM grid derived from the original LIDAR points by linear prediction using a straight line as the covariance function, thus with a very low degree of smoothing (Rottensteiner and Briese, 2002). Sampling the DSM in a regular grid renders possible the application of fast segmentation algorithms. The grid width of the DSM should correspond to the resolution of the LIDAR data. In our example, the original LIDAR points were characterised by point distances of 0.1 m in-flight and 0.9 m cross-flight. The grid width of 0.5 m chosen for the DSM was a compromise, resulting in an interpolation in cross-flight direction and in a higher degree of smoothing in flight direction.

The work flow for building reconstruction consists of four steps:

1. Detection of roof planes: In the building regions, a segmentation based on an analysis of the variations of the DSM normal vectors is applied to find planar segments which are expanded by region growing algorithms. This method is described in section 2.

2. Grouping of roof planes and model generation: Neighbouring roof segments being co-planar have to be merged, and a further analysis of neighbourhood relations gives hypotheses for lines of intersection and/or step edges. The resulting polyhedral model consisting of a conglomerate of (mostly: roof) planes is completed by walls and by a floor. Grouping and model generation are described in section 3.

3. Consistent estimation of the model parameters: Having generated initial building models, their geometric parameters have to be improved by a consistent estimation procedure taking into account all the available (sensor) information.

4. Model regularisation: Finally, the models are improved by introducing hypotheses about geometric constraints between planes, and parameter estimation is repeated. Parameter estimation and model regularisation are described in section 4.

The level of detail of the building models thus derived and the accuracy of the positions of step edges are limited by the resolution of the LIDAR data, which is still below the resolution of aerial images. In section 5 we will discuss strategies for integrating digital aerial images into the work flow of building extraction in order to overcome these limitations.

2 DETECTION OF ROOF PLANES

(Fuchs, 1998) describes a method for polymorphic feature extraction aiming at a classification of texture as being “homogeneous”, “linear”, or “point-like”, by an analysis of the first derivatives of a digital image. That classification is based on a significance test of the squared gradient of the digital images, the most important parameter being the significance level α of that test, from

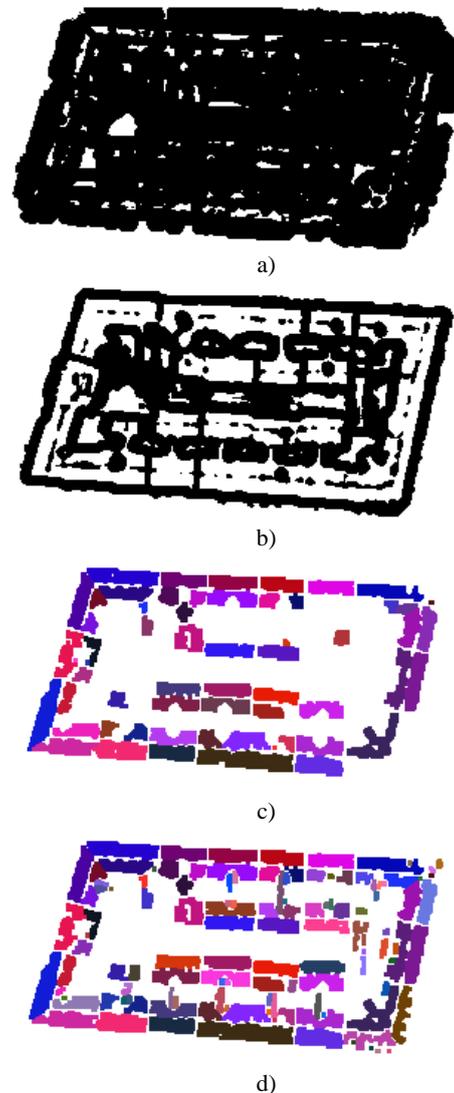


Figure 2: Detection of planar segments in one of the building regions from Figure 1. a) and b) Classification results using different significance levels. White: “homogeneous” pixels. c) Planar regions obtained by applying a connected component algorithm and region growing to a). d) Final segment label image.

which the threshold for classification is derived. We apply this method to the first derivatives of the DSM. “Homogeneous” pixels are surrounded by pixels having similar components of the surface normal vectors, thus, by co-planar points. Figures 2a and 2b show the homogeneous pixels in one of the building regions from Figure 1, derived by using two different significance levels. In (Rottensteiner and Briese, 2002), we applied a connected component analysis to the binary image of homogeneous pixels to detect seed regions for region growing. In order to compensate for classification errors, we proposed an iterative strategy for roof plane detection, splitting the connected components of homogeneous pixels into smaller parts by morphologic filtering and only allowing well-fitting planes to grow. Although the results thus achieved were already quite satisfactory, the procedure failed if the initial segments were shaped in a way that parts belonging to different planes in object space could not be separated.

We have improved the method to put it on a better statistical basis. Again, roof plane detection is carried out in an iterative way, but this time we iterate over different significance levels in texture classification. We start by texture classification using a very tight threshold, thus a very small significance level α . Statistically this means that we accept a large percentage of homogeneous pixels erroneously classified as non-homogeneous, and a small percentage of non-homogeneous pixels erroneously classified as homogeneous. Only the pixels in the “most significantly planar” regions are actually classified as “homogeneous” (Figure 2a). We use connected regions of homogeneous pixels as seed regions for region growing, allowing only regions achieving a good planar fit (indicated by a low r.m.s. error) to grow. New pixels adjacent to a region are added if their distances from the adjusting plane is below a certain threshold. This step is repeated until no additional pixel can be added to any segment. Figure 2c shows the segment label image derived from growing the seed regions from Figure 2a. About 50% of the pixels of the building region have been assigned to one of these planes, each plane having an r.m.s. error of planar fit better than ± 15 cm. After that, the classification is repeated, using another threshold (thus, accepting a larger percentage of non-homogeneous pixels erroneously classified as homogeneous) and applying our procedure for seed region selection to connected components of homogeneous pixels not yet assigned to a planar segment. In this way, a certain number of iterations is carried out, in each iteration taking into account less and less significantly homogeneous areas. Figure 2b shows the classification results of the last iteration, and Figure 2d shows the resulting segment label image. About 64% of the pixels in the building region are assigned to one of these segments. All segments have an r.m.s. error of planar fit better than ± 15 cm.

Figure 3a shows the pixels of the building region in Figure 2d not yet assigned to a planar segment. Most of these “unclassified” pixels are pixels at the borders of the planar regions, but there are also larger patches corresponding either to small roof structures not yet detected, to roof parts not being planar (e.g. roof gardens), or to regions not being parts of the building at all (trees). We try to find additional planes in the unclassified regions and to eliminate objects not belonging to the building. First, we search for connected components of unclassified pixels having a small height gradient strength. As the kernel size for computing the height gradient is much smaller than the size required for texture classification, it is possible to find smaller seed regions. We apply region growing to these regions, allowing larger residuals, so that the additional planar segments have a worse planar fit than the original ones. Thus, we might get segments generalising the actual shape of a building part, a limitation caused by the resolution of the LIDAR data. After that, we eliminate connected components of pixels not belonging to any planar segment and

having a great percentage of pixels classified as “point-like” in polymorphic feature extraction, an indicator for vegetation areas (Rottensteiner and Briese, 2002). Figure 3b shows the final segment label image, and Figure 3c shows the pixels not assigned to any planar segment after eliminating the vegetation areas.

The results of roof plane detection for all the building regions detected in Figure 1 are summarised in Table 1. 69.9% of all “building pixels” are assigned to a planar segment. 92.1% of these pixels (or 64.4% of all building pixels) are in a planar segment having an r.m.s. error of planar fit better than ± 15 cm. The remaining 9.9% of the pixels assigned to a planar segment mostly correspond to planes added to the model in the second segmentation phase, generalising more detailed building shapes.

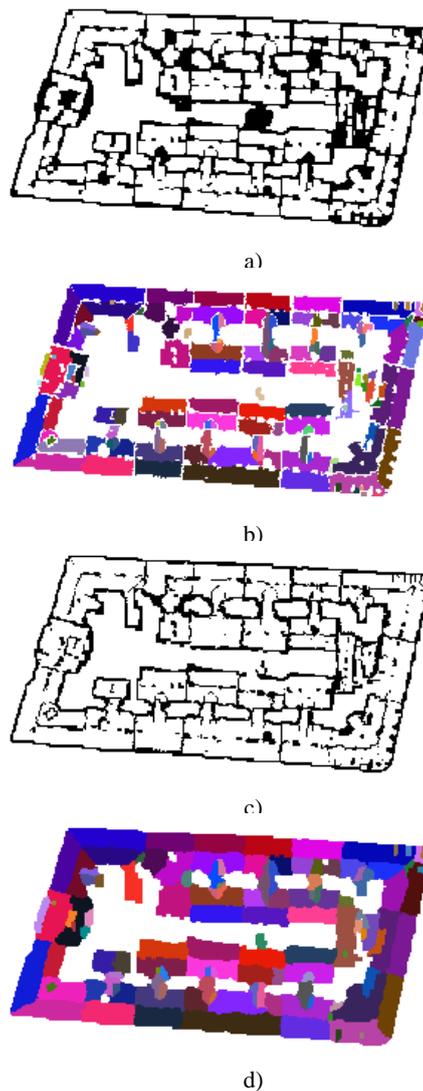


Figure 3: Post-processing of the segmentation results. a) Pixels not yet assigned to a planar segment in the building region in Figure 2. b) Segment label image with additional planar segments after post-processing. c) Pixels not assigned to a planar segment after post-processing. d) Voronoi diagram of b) after merging co-planar segments.

3 GROUPING OF ROOF PLANES AND GENERATION OF ROOF POLYGONS

The neighbourhood relations of the planar segments are derived from a Voronoi diagram of the segment label image (Ameri, 2000):

r.m.s. [m]	0.00-0.05	0.05-0.10	0.10-0.15	0.15-0.20	0.20-0.25	0.25-0.30	0.30-0.35	0.35-0.40	0.40-0.45
Planes [%]	19.5	57.3	15.3	6.5	0.8	0.3	0.2	0.0	0.1
All [%]	13.6	40.1	10.7	4.5	0.7	0.2	0.1	0.0	0.0

Table 1: Distribution of the r.m.s. errors of the planar fit for all building regions in Figure 1. Planes: percentage of pixels in planar segments having the respective r.m.s. error compared to the number of all pixels in all planar regions. All: percentage of pixels in planar segments having the respective r.m.s. error compared to the number of all pixels inside the building regions. 69.9% of the pixels in the building candidate regions are classified as belonging to a planar segment.

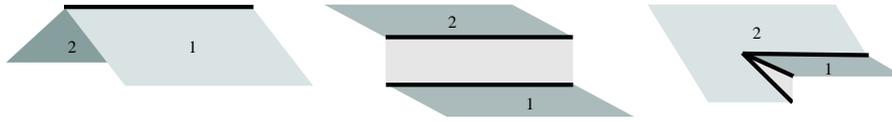


Figure 4: Mutual relations between two planes ε_1 and ε_2 . Left: intersection. Centre: step edge. Right: intersection and step edge.

each pixel inside the region of interest not yet assigned to a planar segment is assigned to the nearest segment. After having found the neighbourhood relations, co-planar segments are merged. The co-planarity test is based on a Fisher test comparing the r.m.s. errors of planar fit obtained by applying two different mathematical models (two separate planes vs. one compound plane). Figure 3d shows the Voronoi diagrams after merging co-planar segments.

The neighbourhood relations of the remaining segments are analysed. The boundary polygons of the Voronoi regions give a first estimate for the planimetric positions of the segment outlines (Figure 5a). Each portion of the boundary separating two planar segments has to be classified according to the geometric configuration of these segments (Figure 4): There might be an intersection, a step edge, or both an intersection and a step edge.

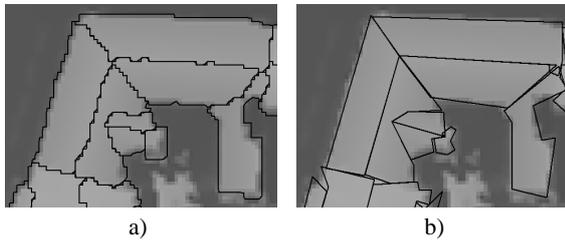


Figure 5: A part of the roof polygons of the building in Figure 2 super-imposed to the DSM. a) The original pixel chains from Figure 3d. b) After finding intersections and step edges.

In order to perform this classification, the intersection lines of all pairs of neighbouring planar segments are computed. For each pair we look for parts of the original edge pixel chain separating the two segments that fit to the intersection line (i.e., where the r.m.s. error of the original edge pixels with respect to the intersection line is below a certain threshold). If the whole edge pixel chain fulfills that criterion, the intersection line is the actual border between these planar segments, and the edge pixel chain is replaced by the intersection line. If the criterion is fulfilled for parts of the original edge pixel chain, the edge pixel chain is split into several parts, each of them being handled individually. For the parts that fit to the intersection line, the edge pixel chain is replaced by the intersection line, too.

For all the other parts and for edge pixel chains not having a part that fits to the intersection line, we assume a step edge. To locate the step edge, we approximate the original edge pixel chain by a polygon and then search for the positions of maximum height change in the direction perpendicular to the polygon. We evaluate the norm of the inner product of the DSM gradient and the normal vector of the polygon (i.e. the gradient component perpendicular

to the proposed step edge) for that purpose. We obtain an improved edge pixel chain. If this edge pixel chain can be approximated by a straight line segment, it is replaced by that segment. Otherwise, it has to be approximated by a polygon, the polygon giving the final position of the step edge. Currently, no further efforts are performed to improve the shapes of these polygons that might be rather irregular, which is partly caused by outliers (large gradients actually corresponding to another step edge).

Having replaced the original edge pixel chains separating two planar segments either by intersection lines or by polygons corresponding to step edges, the polygon delineating a planar segment has to be put together from these individual parts. We compute the intersection points of neighbouring straight line or polygon segments. Where the intersection point is consistent with the model (in the sense that it is not too far away from the original vertices), it is accepted to be a vertex of the boundary polygon. Where it is inconsistent, a straight line segment is used as an edge connecting the two original vertices. Thus, for each planar segment, we get one or more closed border polygon(s).

After having generated the roof polygons, the polyhedral models have to be created. The planar segments become the roof faces of the model, the vertices of the boundary polygons become the building vertices, and the edges of these polygons become the edges of the polyhedral model. A vertical face corresponding to a wall is inserted for each polygon edge classified as a step edge, especially at the building outlines. Finally, a floor is added to complete the polyhedral model.

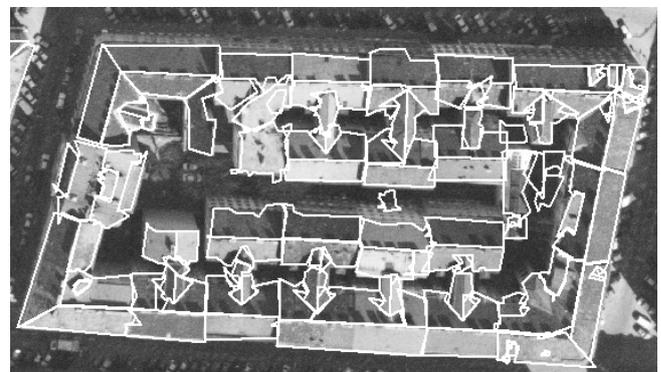


Figure 6: The roof polygons of the building in Figure 2 back-projected to an aerial image. Ground pixel size: 15 cm.

The modules for grouping and model generation are still work in progress. Figure 5b shows preliminary results of computing intersections and step edges. The polygon parts being the results of

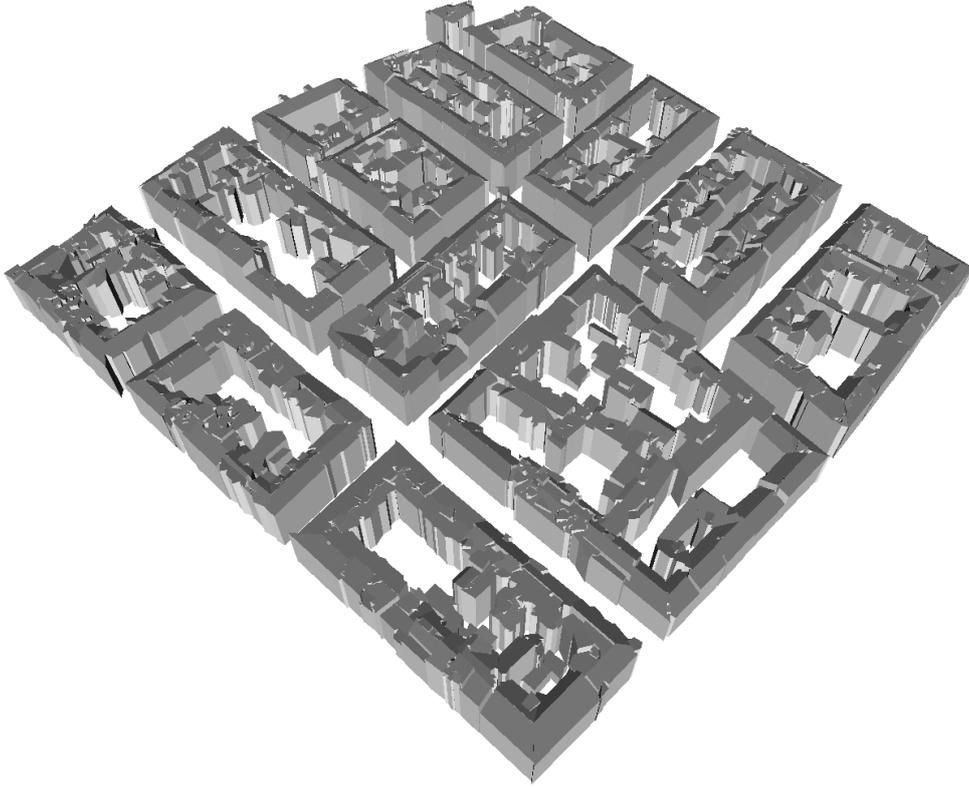


Figure 7: VRML visualisation of a model created from vertical prisms bordered by the boundary polygons and the floors.

step edge extraction still have to be improved by eliminating outliers and by more sophisticated methods for thinning, e.g. similar to the one proposed by (Weidner, 1997) for simplifying the shapes of the polygons describing the building outlines.

Figure 6 shows the reconstructed roof polygons super-imposed to an aerial image of scale 1:7000. The roof structure is resembled quite well by the model. Smaller roof details such as the chimneys could not be reconstructed given the resolution of the LIDAR data, and there are still planimetric offsets in the images that correspond to about 5 image pixels, thus 1-1.5 times the grid width of the DSM. The somewhat ragged appearance of some polygon segments is caused by outliers in step edge detection.

As the computation of intersection lines and step edges moves the positions of the borders between neighbouring faces, the neighbourhood relations might change. This means that it would be desirable to repeat the analysis of neighbourhood relations after generating the polygons.

Figure 7 shows a visualisation of a 3D model created from vertical prisms bounded by the roof polygons and the floor. The structure of the roofs is correctly modelled. These preliminary results are very encouraging, and they show the high potential of high-resolution LIDAR data for building extraction.

4 ESTIMATION OF THE BUILDING PARAMETERS

In the previous phases of reconstruction, the parameters of the planar segments were determined individually. The building vertices were determined from the intersections of adjacent planes or by step edge extraction. An overall adjustment including all available sensor information is important to get geometrically consistent building models and to consider geometric constraints.

We use our hybrid adjustment program ORIENT (Kager, 1989) for the estimation of the model parameters. It is the idea of our method to find a “mapping” between the boundary representation (B-rep) of a building and a system of “GESTALT” observations representing that B-rep in adjustment. GESTALT observations are observations of a point \mathbf{P} being situated on a polynomial surface. The polynomial is parameterised in an observation co-ordinate system (u, v, w) related to the object co-ordinate system by a shift \mathbf{P}_0 and three rotations $\theta = (\omega, \varphi, \kappa)^T$. The actual observation is a fictitious observation of the largest component of the distance vector to be 0. Using $(u_R, v_R, w_R)^T = \mathbf{R}^T(\theta) \cdot (\mathbf{P} - \mathbf{P}_0)$, with $\mathbf{R}^T(\theta)$ being a transposed rotational matrix parameterised by θ , and restricting ourselves to vertical planes for walls and tilted planes for roofs, there are three possible formulations of GESTALT observation equations:

$$\begin{aligned} r_u &= m_u \cdot u_R + a_{00} + a_{01} \cdot (m_v \cdot v_R) \\ r_v &= m_v \cdot v_R + b_{00} + b_{10} \cdot (m_u \cdot u_R) \\ r_w &= m_w \cdot w_R + c_{00} + c_{10} \cdot (m_u \cdot u_R) + c_{01} \cdot (m_v \cdot v_R) \end{aligned} \quad (1)$$

In equation 1, r_i are the corrections of the fictitious observations of co-ordinate i and $m_i \in \{-1, 1\}$ are mirror coefficients. An application is free to decide which of the parameters (\mathbf{P} , \mathbf{P}_0 , θ , surface parameters a_{jk}, b_{ik}, c_{ij}) are to be determined in adjustment and how to parameterise a surface. In addition, different GESTALTs can refer to identical transformation or surface parameters, which can be used to enforce geometric constraints. For building extraction, we declare the rotations to be 0 and constant. \mathbf{P}_0 is situated inside the building and constant. For each face of the B-rep of the building model, we define a set of GESTALT observations, taking one of the first two equations 1 for walls and the third one for roofs. The unknowns to be determined are the object co-ordinates of each point \mathbf{P} and the plane parameters (a_{jk}, b_{ik}, c_{ij}). As each building vertex is neighbored by at least three faces, the object coordinates of the vertices can be deter-

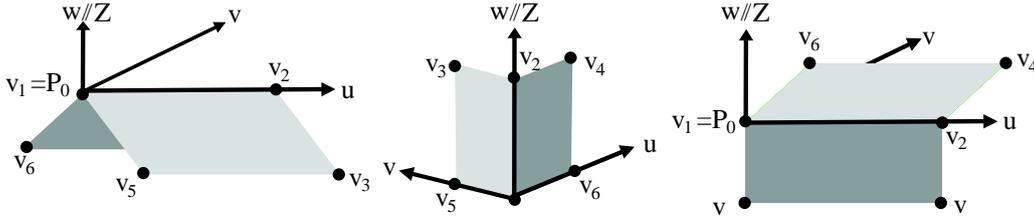


Figure 8: Three possible geometrical constraints between two planes ε_1 and ε_2 . Left: a horizontal ridge. Centre: two orthogonal walls. Right: a horizontal eave.

mined from GESTALT observations for these vertices. We use the sensor data to determine the unknown plane parameters for all the planes of the building:

- *LIDAR points* are assigned to roof planes. For each LIDAR point, one GESTALT observation will be introduced, as well as three direct observations for the point's LIDAR co-ordinates to avoid singularities with respect to the object co-ordinates P .
- *Direct observations* for the planimetric co-ordinates of building vertices at step edges (the results of step edge detection) are introduced in order to make the parameters of the wall faces determinable.

Using this mathematic model, we get a highly redundant adjustment system. It is possible to apply robust estimation (Kager, 1989) to eliminate false or contradicting observations. The weights of the observations depend on their a priori standard deviations. The most important parameter of the stochastic model is the a priori standard deviation of the GESTALT observations, because in robust estimation, it decides about the "rigidity" of the building model. Typically, we select it in the range of a few centimetres.

4.1 Regularisation of Polyhedral Models

Up to now, we have not yet introduced geometric constraints. In the future, we will develop a module that has to analyse the polyhedral model and introduce geometric constraints where sufficient evidence for their occurrence is found. If two planes ε_1 and ε_2 are found to fulfil a geometric condition, additional GESTALT observations will be added to adjustment, taking advantage of specific definitions of the observation co-ordinate system and specific parameterisations of the planes. Figure 8 shows three such constraints. In all cases, one of the axes of the observation co-ordinate system is defined to be the intersection of ε_1 and ε_2 , and one of the vertices of the intersection line is the reference point P_0 of both planes. The rotations ω and φ are 0 and constant, but there is one additional unknown rotational angle κ . For each vertex v_i of the planes, GESTALT observations are added for ε_1 and/or ε_2 . These planes are parameterised in a specific way:

- The intersection of two planes is a horizontal ridge:
 $\varepsilon_1 : r_w = w_R + c_{01}^1 \cdot v_R$ and $\varepsilon_2 : r_w = w_R - c_{01}^1 \cdot v_R$.
 There is only one tilt parameter c_{01}^1 . Symmetry is enforced by selecting $m_v = -1$ for ε_2 .
- Two walls are perpendicular: $\varepsilon_1 : r_u = u_R, \varepsilon_2 : r_v = v_R$.
 There is no additional surface parameter to be determined.
- A wall and a roof plane intersect in a horizontal eave:
 $\varepsilon_1 : r_w = w_R + c_{01}^1 \cdot v_R$ and $\varepsilon_2 : r_v = v_R$. There is an additional unknown roof tilt c_{01}^1 .

The stochastic model of these constraints is described by the a priori standard deviations of the GESTALT observations. After creating these additional observations, the estimation of the model parameters has to be repeated.

5 INTEGRATION OF LIDAR AND AERIAL IMAGES

The quality of the building models derived from LIDAR data is restricted by the ground resolution of the LIDAR sensor. As the resolution of aerial images is still much better than that, it would be desirable to integrate aerial images into the process of building extraction. As we have seen in the previous sections, the most relevant roof structures can be extracted well from LIDAR data. There are two stages where the aerial images can help to improve the quality of the resulting models: They can help to detect additional planar segments, and they can be used to improve the geometric quality of the model edges by matching model and image edges.

5.1 Detection of Planar Segments

In section 2, we have described that we accept segments having r.m.s. errors planar fit that indicate that small roof structures have been generalised, and there are still some unclassified pixels. In both cases, it is possible to improve the segmentation results using the information provided by the aerial images.

We use polymorphic feature extraction for finding segments of homogeneous grey levels in each digital image. Figure 9 shows one of the areas containing pixels in the LIDAR DSM not yet assigned to a planar segment (from Figure 3a), the corresponding area in an aerial image, and the results of a grey level segmentation in that aerial image. In the LIDAR DSM, that area only has a width of about 5 pixels. It is impossible to separate the two roof planes obviously connecting two larger building parts on the upper and the lower margins of the figure, because the sizes of the filters involved in seed region selection limit the size of detectable

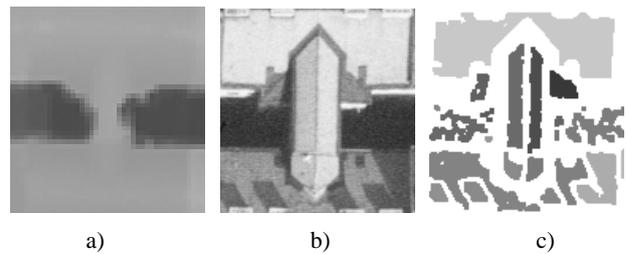


Figure 9: a) The DSM in one of the areas not yet assigned to a planar segment in Figure 3a. b) The corresponding area in an aerial image. c) Result of a grey level segmentation in that aerial image.

planes. However, in the aerial image, the grey level segmentation gives two segments corresponding to two roof planes. Thus, Figure 9 shows that it makes sense to use the image segmentation results in the areas where no planes can be detected in the LIDAR data. It is possible to project the image segments to the DSM to get a first estimate for the plane parameters. In addition, homologous segments from different images can be matched, so that the segmentation results from different images can be combined. In the future, the most significant segments, (those obtaining support in multiple images), will be selected for further processing. They can be added directly to the list of planar segments, or they can be used as new seed regions for region growing in areas where no such regions could be derived from the LIDAR data alone due to the restrictions imposed by the sizes of the involved filters.

5.2 Wire frame fitting

Wire frame fitting is a good approach for improving the geometric quality of the polyhedral models created from LIDAR data: Unlike with LIDAR data, the grey level edges corresponding to the building outlines can be extracted very precisely in the images. It is the idea of wire frame fitting to project the polyhedral model derived in the way described in sections 2 and 3 back to the images (Figure 6), where its edges can be matched with image edges. In each image, the projected position of an object edge gives us a region of interest, and we search for image edges inside that region of interest that are almost parallel to the projected object edge. Image edges fulfilling this condition are supposed to be matching candidates, and they are considered in the estimation of the model parameters by an expansion of the adjustment model described in section 4: In addition to the observations described there, the image co-ordinates of the end points of the image edge segments are assigned to a roof edge. For each end point, we get two image co-ordinates and two GESTALT observations (one for each face being neighbour of the object edge), and three additional unknown object co-ordinates P .

Whereas we search for matching candidates for all roof edges independently in all images, parameter estimation is performed in an overall robust estimation process in a hybrid adjustment of the GESTALT observations representing the object model, the LIDAR points, the positions of the step edges, and the image co-ordinates of image features. Robust estimation is applied in order to determine false matches between image and object edges. We have first used this model fitting algorithm for automatic fine measurement of parametric building models in a system for semi-automatic building extraction from aerial images (Rottensteiner, 2001). In a test project in the context of semi-automatic building extraction (image scale 1:4500, six-fold overlap) we have shown that results with an accuracy in the range of a few centimetres can be achieved using that method.

6 CONCLUSION AND FUTURE WORK

We have presented a method for building extraction from high-resolution LIDAR data in densely built-up areas. Preliminary results from a test site in the City of Vienna have shown the high potential of the method, and we have also discussed the issue of integrating aerial images into the work flow to obtain better results by what could be called "multi-sensor-grammetry" rather than photogrammetry. Future work will not only include the implementation of the modules still missing, but also the assessment of quality parameters for the results. With respect to the building outlines, this can be accomplished by a comparison to existing GIS data of our test area.

ACKNOWLEDGEMENTS

This research has been supported by the Austrian Science Foundation (FWF) under project no. P15789 and by the Australian Research Council (ARC) under Discovery Project no. RM01105.

REFERENCES

- Ameri, B., 2000. Automatic Recognition and 3D Reconstruction of Buildings from Digital Imagery. PhD thesis, University of Stuttgart. DGK Volume 526.
- Baillard, C., Schmid, C., Zisserman, A. and Fitzgibbon, A., 1999. Automatic Line Matching and 3D Reconstruction of Buildings from Multiple Views. IAPRS, Vol. XXXII/3-2W5, Munich, Germany, pp. 69–80.
- Brenner, C., 2000. Dreidimensionale Gebäuderekonstruktion aus digitalen Oberflächenmodellen und Grundrissen. PhD thesis, University of Stuttgart. DGK Volume 530.
- Briese, C., Pfeifer, N. and Dorninger, P., 2002. Applications of the Robust Interpolation for DTM determination. IAPRSIS, Vol. XXXIV/3A, Graz, Austria, pp. 55 – 61.
- Fuchs, C., 1998. Extraktion polymorpher Bildstrukturen und ihre topologische und geometrische Gruppierung. PhD thesis, University of Bonn. DGK Volume 502.
- Haala, N., Brenner, C. and Anders, K.-H., 1998. Urban GIS from Laser Altimeter and 2D Map Data. IAPRS, Vol. XXXII-3/1, Columbus, OH, pp. 339–346.
- Kager, H., 1989. ORIENT: A Universal Photogrammetric Adjustment System. In: A. Grün and H. Kahmen (eds), Optical 3-D Measurement, Herbert Wichmann Verlag, Karlsruhe, Germany, pp. 447–455.
- Rottensteiner, F., 2001. Semi-automatic extraction of buildings based on hybrid adjustment using 3D surface models and management of building data in a TIS. PhD thesis, Vienna University of Technology. Vol. 56 of Geowissenschaftliche Mitteilungen der TU Wien.
- Rottensteiner, F. and Briese, C., 2002. A New Method for Building Extraction in Urban Areas from High-Resolution LIDAR Data. IAPRSIS, Vol. XXXIV/3A, Graz, Austria, pp. 295 – 301.
- Schenk, T. and Csatho, B., 2002. Fusion of LIDAR data and aerial imagery for a more complete surface description. IAPRSIS, Vol. XXXIV/3A, Graz, Austria, pp. 310 – 317.
- Vosselman, G., 1999. Building Reconstruction Using Planar Faces in Very High Density Height Data. IAPRS, Vol. XXXII/3-2W5, Munich, Germany, pp. 87–92.
- Vosselman, G. and Dijkman, S., 2001. 3D building model reconstruction from point clouds and ground plans. IAPRSIS, Vol. XXXII-3/W4, Annapolis, Maryland, pp. 37–43.
- Weidner, U., 1997. Gebäudeerfassung aus digitalen Oberflächenmodellen. PhD thesis, University of Bonn. DGK Volume 474.