

BUILDING ROOF RECONSTRUCTION FROM LIDAR DATA AND AERIAL IMAGES THROUGH PLANE EXTRACTION AND COLOUR EDGE DETECTION

Angelina Novacheva^{a,*}

^aInstitute of Photogrammetry and Remote Sensing, Technical University of Dresden, 01062 Dresden, Germany
angelina. - novacheva@mailbox.tu-dresden.de, <http://www.tu-dresden.de/ipf/photo>

KEY WORDS: Laser Sanning (LiDAR), Building Rconstruction, Aerial Photogrammetry, Data Integration, Urban planning, Image Pocessing

ABSTRACT:

In this paper a strategy for 3D reconstruction of building roofs from airborne laser scanning and aerial images is discussed. In order to keep it as general as possible, no predefined primitives or ground plans are required. The processing is done directly on the raw LiDAR point cloud, so as to avoid any loss of information due to interpolation. Computations involving local surface normals, which are usually rather noisy in dense datasets, are avoided. Only roofs composed of planar patches are considered. The guiding principle is to select thresholds that can be derived from the data itself and to make the algorithms largely independent of their exact values. The main purpose of image integration is the refinement of the building outline. In that the importance of utilising the available chromatic information is shown.

1. INTRODUCTION

3D city modelling has recently been a lively research area within the photogrammetric community. Buildings, as the most prominent features of the urban landscape, receive special attention. The new developments in sensor technology allow for increased automation in their reconstruction. The improved accuracy and density of airborne laser scanning (LiDAR), as well as the availability of simultaneously recorded height data and colour frame imagery have directed the attention of many researchers towards the extensive application of LiDAR and the integration of aerial images.

Most authors, e.g. [Rottensteiner & Briese, 2003] concentrate on the usage of raster DSMs, obtained through the interpolation of the laser scanning data to a regular grid. That allows for the application of available image processing software and fast segmentation methods, but has the disadvantage of decreasing the information content.

Another work, focused on building reconstruction through the combination of image and height data [Haala, 1996], uses a raster DSM, obtained through matching from aerial images along with the image data itself. 3D intensity or DSM edges are compared to a building model. The data integration is mainly limited to the detection of regions, corresponding to buildings in the height data.

In the following, LiDAR data with density of about 5.3 points/ m^2 is considered along with colour aerial images of 10 cm ground resolution.

2. ROOF RECONSTRUCTION FROM LIDAR

2.1 Prerequisites

There are two important assumptions related to the proposed algorithm. First, a rough segmentation of the point cloud should be available. Second, only roofs consisting of planar faces can

be reliably reconstructed. An overview of the current processing pipeline is given in Figure1.

The segmentation could be done by pre-processing the data as described in [P. Axelsson, 1999], which does not require additional information like multiple return or intensity values. In the following processing steps it is expected to be neither error free nor complete. However, in order to accurately extract the separate buildings, at least three neighbouring vertices from each roof plane should be present in a connected component of segments labelled as belonging to class “buildings”.

Further each single building with its immediate surroundings is handled separately. The standard deviation (RMS) of plane fit is also determined a priori and considered uniform within the data set.

2.2 Roof Segment Identification

At first the 2.5D Delaunay triangulation of a building region is computed, which becomes the basis for the definition of the neighbourhood relations within the point cloud. In that, as well as to support further development works, the data structures and the functionality of the Computational Geometry Algorithms Library [CGAL, 2006] are employed.

A procedure based on the Hough transform is responsible for the successive detection, verification and refinement of planar segments. First a modified version of the Hough transform is performed, generating a 3D parameter space for plane detection, based on the perpendicular distance to the origin and the polar coordinates of the plane's normal vector [A. Novacheva, 2007]. In that, special care is taken to provide uniform sampling of the Gaussian sphere. As the parameters acquired in such a way are only approximate, they are refined in the next step.

Outliers are removed given a predefined threshold for the orthogonal distances of the measurements to the plane based on the standard deviation. Empirically the value of 1.2 x RMS was found appropriate.

Given the Delaunay structure, a connected component analysis with single linkage is performed and only the largest cluster is further regarded as a correct detection, with the plane parameters being adjusted to it. That is necessary, not only because of the preceding outlier removal, but also because the Hough transform itself is prone to finding spurious features.

Due to the discrete nature of the Hough space there is no guarantee that all the measurements belonging to a given plane will be fully contained within a single accumulator cell. Thus, singly connected adjacent points to those belonging to the outline of the largest component are also examined against the orthogonal distance criteria applied above, and possibly associated with the plane. In that way the correct spatial extent of the detected segment is obtained.

Once a roof segment has been extracted, its corresponding measurements are excluded from processing and the algorithm iterates till no face of class "building" is left.

After all segments have been identified, the assignment of vertices along their common borders is reconsidered in order to eliminate the effects of the sequence in which the processing took place.

The final result is a set of roof planes with their optimal form and orientation parameters - i.e. the number of segments has been determined and each vertex is assigned to the one, to which it is most likely to belong.




Figure1. Procedure for roof reconstruction from LiDAR data.

2.3 Modelling of the Roof Structure

The correct interpretation of the roof structure is the most challenging task. First the common border of a segment with a given neighbouring segment is detected. Usually that is a single straight line, but if not it has to be subdivided in such. That is achieved through Ramer's algorithm for polygon simplification [Ramer, 1972] in the following way: a straight line between the first and last vertex (or the two at a largest distance from each other in case of a closed polygon) of a border segment is intersected with the 2.5D triangulation of the point cloud; the largest of the orthogonal 2D distances between the line and the vertices incident to the triangulation edges intersected by it, is selected as the split threshold (Figure2). That suggests a relatively uniform point distribution. The threshold, selected in that way is rather stringent and usually the result is over-segmentation. That however, is a desired effect, due to the sparse nature of the LiDAR data, which hinders the selection of the optimal split vertex at every step. The following aggregation is done in two steps. First sequences of linear segments, with less than k vertices, where k is some threshold are merged. Then single short lines are merged with the preceding or succeeding, depending on the orientation difference.




Figure 2. Ramer's algorithm – border vertices of a segment with another (black), initial approximation (red line), vertices used to derive the split threshold (red) and final result (black lines).

The roof can contain two types of discontinuities - break lines and step edges. Those could be distinguished based on the height differences between the border points, belonging to the corresponding straight line segments of two neighbouring planes. Such distinction is important as each type of feature is modelled in a different way. The position of a break line could accurately be determined as the intersection of a pair of planes. Its length is obtained through the orthogonal projection of the end border points between two segments. Step edges are handled similarly to the gutter.

2.4 Determination of the Building Outlines

Through the outlier removal and the segment growing, during the roof segment identification, misdetections in the classification have been compensated for. That means that at his point an improved estimation for the building gutter is available.

It is extracted as the outline of the combination of all roof structures. For that a gift wrapping like algorithm has been developed, which uses the neighbourhood information of the triangulation. Because of the discreteness of the LiDAR measurements, the building footprint will contain many short, insignificant edges and needs to be further generalized. To tackle that problem Ramer's algorithm has also been used here (Figure 3). It could be seen that the outline is successfully smoothed, but the edges still deviate from their true position. Because of the higher spatial resolution of the image data it can be used in improving the accuracy of the gutter, e.g. through the extraction of edges.



Figure 3. Building outline (green) and its generalization through Ramer's algorithm (blue) projected onto the image.

3. COLOUR EDGE DETECTION

Here the applicability of different edge detection approaches to colour aerial images is considered. Usually, when processing digital colour images the first thing to do is to disregard the colour information by converting the data to intensity. That makes the application of conventional edge detectors straightforward. However, if neighbouring regions differ only in their chromatic characteristics, they could not be separated. Notice e.g. that the edge between the building roof and the pavement in Figure 4 is almost completely lost.




Figure 4. Canny intensity edges.

Often colour edge detection is performed individually on the different components of a colour image, on the analogy of intensity edge detection. Though such methods are rather popular, it is not clear how the results should be integrated. If the vector sum of the gradients is used, the result is exactly the same as the one obtained by calculating the gradient of the intensity, i.e. the problem of isoluminance due to opposing derivative vectors is not solved in that way. A better approach is the calculation of gradient magnitudes on an individual channel basis. That, according to [Gonzalez & Woods, 2008], yields acceptable results in most practical applications, but an extreme case of misperformance occurs when gradients pointing in perpendicular directions in the separate components are treated as if they were collinear.

Clearly a different definition of the gradient, suitable for vector-valued functions is needed if the information content is to be explored to its full scale. One possibility is to view the image as a two dimensional vector field and assign to each pixel a vector from a given colour space. A simple procedure would be to take a 3x3 mask and compare the distances from the centre to the neighbouring pixels – e.g. based on RGB coordinates with an angular Euclidian metric. The largest distance then corresponds to the gradient magnitude and the gradient direction is one of the eight possible within the mask [Plataniotis & Venetsanopoulos].

A more elaborate approach was proposed by [Di Zenzo, 1986], who derives the gradient direction and magnitude based on the colour structure tensor. The required partial derivatives in that case can be obtained e.g. through the Sobel operator, which includes Gaussian smoothing and is thus relatively robust to noise. The result of the Canny operator with gradient magnitudes obtained in that way is shown in figure. 5.



Figure 5. Canny colour edges.

In [Weijer et al., 2006], additional photometric information is used to reason about the physical causes for image discontinuities. The authors proposed the so called photometric quasiinvariants, which like the full invariants (e.g. normalized RGB) provide robustness against shadows and specular reflections, but do not possess the same instabilities. Figure 6 shows the result of Canny edge detection in the RGB colour space with shadow-shading and Figure 7 – with shadow-shading-specular quasiinvariants. It was noticed, that in both cases less luminance edges are found, but the accuracy of the location and orientation of the edge segments is inferior compared to the one with the colour gradient. Those displacements are probably due to the general attenuation of the signal associated with invariants.

Due to the unfavourable behaviour of quasiinvariants it is not advisable to integrate them in the edge detection for this particular type of application. In addition the vector gradient approaches to colour edge detection already possess a certain degree of invariance.



Figure 6. Canny colour edges with shadow-shading



Figure 7. Canny colour edges with shadow-shading-specular quasiinvariants.

4. USING THE LIDAR DATA IN GUIDING THE EDGE DETECTION

The approximate building outlines obtained from the LiDAR data can be used in the determination of search regions for the corresponding image edges in the following way: a 2D line is fitted to each segment, obtained through Ramer's algorithm; its parameters are used to define a region in the image where edge detection is carried out; only edge elements, whose gradient direction is consistent with the edge approximation from the LiDAR data are considered; neighbouring segments can then be examined for collinearity and merged if necessary;

5. CONCLUSIONS AND FUTURE WORK

The possibility of using LiDAR data in guiding the edge detection process was discussed. It was shown, based on a practical example, that colour information should not be disregarded in that. Furthermore it was noted that quasiinvariants are not suitable for this type of applications.

In future it would be interesting to evaluate the applicability of level set methods to the extraction of the gutter from the images. In that the contours could directly be initialized with the outline of the triangulated roof, thus skipping the generalization by Ramer's algorithm. Further it should be possible to use not only discontinuities, but also to integrate regional properties in the functional optimization.

REFERENCES

Axelsson, P. (1999). Processing of Laser Scanner Data – Algorithms and Applications. *ISPRS Journal of Photogrammetry & Remote Sensing*, vol. 54.

CGAL Editorial Board, 2006. *CGAL-3.2 User and Reference Manual*.

- Di Zenzo, S., 1986. A Note on the Gradient of a Multi-Image. *Computer Vision, Graphics and Image Processing*, vol. 33.
- Gonzalez, R. & Woods, R., 2008. *Digital Image Processing*. Prentice Hall.
- Haala, N., 1996. *Gebäuderekonstruktion durch Kombination von Bild- und Höhendaten*, PhD Dissertation, Stuttgart.
- Novacheva, A. 2007. Towards Automatic 3D Roof Reconstruction from Laser Scanner Point Clouds. *4-th Conference "Recent Problems in Geodesy and Related Fields"*, Sofia, Bulgaria.
- Plataniotis, K. N. & Venetsanopoulos, A.N., 2000. *Color Image Processing and Applications*. Springer-Verlag, Heidelberg.
- Ramer, U., 1972. An Iterative Procedure for the Polygonal Approximation of Plane Curves, *Computer Graphics and Image Processing*, vol. 1, part 3, pp. 244 - 256.
- Rottensteiner, F. & Briesse, C., 2003. Automatic Generation of Building Models from LIDAR Data and the Integration of Aerial Images, *ISPRS workshop "3-D reconstruction from airborne laserscanner and InSAR data"*, Dresden, Germany.
- Weijer, J. van den, Gevers, T., & Smeulders, A.W.M, 2006. Robust Photometric Invariant Features From the Color Tensor. *IEEE Transactions on Image Processing*, vol. 15, No. 1

