PLANAR FEATURE EXTRACTION IN TERRESTRIAL LASER SCANS USING GRADIENT BASED RANGE IMAGE SEGMENTATION

Ben Gorte

Department of Earth Observation and Space Systems (DEOS)
Delft University of Technology, Kluyverweg 1, 2629 HS Delft, The Netherlands
b.g.h.gorte@tudelft.nl

Commission WG V/3

KEY WORDS: Terrestrial Laser Scanning, Range Images, Segmentation.

ABSTRACT:
The paper presents a new segmentation algorithm, to be applied to terrestrial lasers scans of urban environments. The algorithm works directly in a range image. This is the fundamental laser scan data structure as a laser recording can be regarded as a 2-dimensional grid of range measurements. The horizontal and vertical axes of the grid denote the horizontal and vertical angles at which the scanner emits the laser beam, receives the reflections, and measures the distance (the range) between the instrument and the reflecting surface at those angles.

The presented algorithm estimates for each measurement (pixel) in the range image three parameters of the 3D plane that contains the pixel: two angles (horizontal and vertical) and the distance between the plane and the origin. The estimates are based on the scan angles (the horizontal and vertical angles at which the laser beam was emitted from the scanner) and the image gradients, i.e. the rate of change in the distance that is observed between adjacent measurements. Since the three estimated parameters characterize a plane in 3D space, regions of adjacent pixels with similar parameter values are likely to be part of the same plane. Such pixels are grouped into segments by a region-growing image segmentation step, which takes the three parameters into account simultaneously.

The overall algorithm uses two complementary strategies to deal with the measurement noise affecting the gradients, during the gradient calculation and the region growing steps respectively.

1. INTRODUCTION

Range image segmentation has a long tradition in the computer vision research community. For example, Hoover *et al.* (1996) already set up a framework for experimental comparison of range image segmentation algorithm results, flowed up by Xiang *et al* (2000). At the time, range images were produced by photogrammetric interpretation of stereo imagery, by structured light techniques, or by early laser equipment, such as the Perceptron laser ranger. Commonly, those data were recorded in well-controlled conditions. Nevertheless, they tended to be noisy and have low point densities. They described relatively simple close-range scenes, to be used in industrial applications and in robot vision experiments.

Recent advancements in terrestrial laser scanning cause a renewed interest in segmentation of range data. Terrestrial laser scanners are being used in a multitude of applications, for example in 3D model reconstruction of complex outdoor scenes in urban environments. Nowadays scanners are able to record datasets with millions of points, and with recording speeds of several hundreds of thousands of points per second. The distance range has been increased to approx. 100 m for phase scanners, and to several kilometres for time-of-flight scanners, with spatial accuracies in the cm-range (Staiger, 2007). These changes pose new requirements to segmentation algorithms.

The common perception of terrestrial laser scanning is that it results in a 3-dimensional point cloud, *i.e.* a collection of (x,y,z) coordinates, corresponding to locations in the scene were the laser beam was reflected by a surface. Additionally, most laser scanners record the intensity of the reflected beam as it is recorded by the instrument. Some types of laser equipment record a colour image of the scene, more or less at the same time and from approximately the same position as the laser scan. This image can be used to "colour" the point cloud, *i.e.* to assign (R,G,B) values to the (x,y,z) points of the laser scan. Many authors have reported on point cloud segmentation algorithms, amongst whom Rabbani (2006), who starts off with a Hough transform using (θ,ϕ,ρ) parameterization (see section 2). He also gives an overview of previous methods, stating that these either resample the data in a 2D or 3D grid, or build a topology on the point cloud using triangulation.

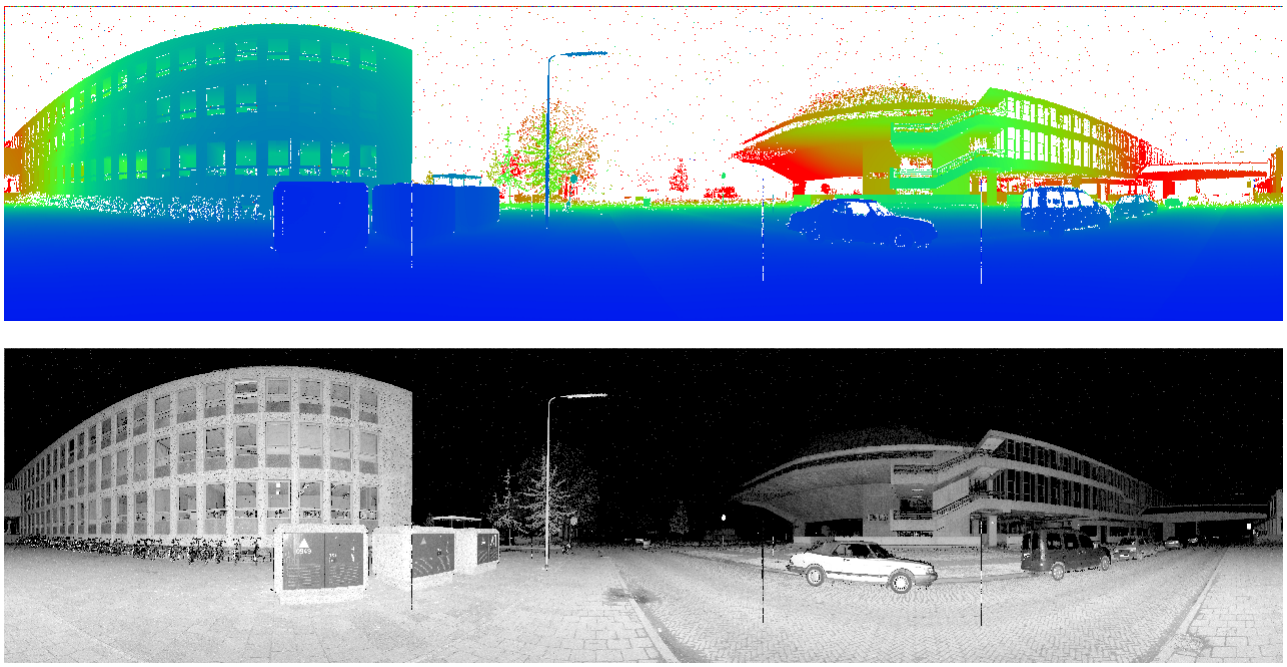

Figure 1. Upper: range image, lower: intensity image. The data were recorded by a FARO 880 phase scanner. Image size is approx. 8100 x 2200 pixels, covering 360° horizontally and 100° vertically. This corresponds to 1/5th of the scanner's maximum resolution.

As an alternative to the 3D point-cloud notion, and better in accordance with scanner operation, a laser recording can be regarded as a 2-dimensional grid of range measurements. The horizontal and vertical axes of the grid denote the horizontal and vertical angles (let these be called α and β respectively) at which the scanner emits the laser beam, receives the reflections, and measures the distance (the range) R between the instrument and the reflecting surface at those angles. The angles α and β are sampled at regular intervals $\Delta\alpha$ and $\Delta\beta$ (the angular resolutions). Therefore, laser scanning results in a 2D (often termed "2.5D") range image R $[i_\alpha, i_\beta]$, with $\alpha = i_\alpha \Delta\alpha$ and $\beta = i_\beta \Delta\beta$. (Figure 1).

Figure 1 also shows the intensity image, which is recorded simultaneously with the range image. (The remainder of this paper only concerns range measurements and does not consider intensity or (R,G,B)-information.)

Pulli (1993) already described range image segmentation with some similarity to the method presented here, working with normal vectors and 3-feature image segmentation.

2. PLANES IN RANGE IMAGES

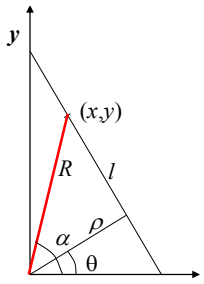

Figure 2. A line in 2D containing point (x,y), and its normal vertor

We will first establish a relation between the equation of a plane in 3D Cartesian coordinates, and the representation of that plane in a range image with 2d spherical image coordinates. The situation resembles the equation of a line l in 2D passing through a point (x, y), as it is often used in Hough transforms:

$$\rho = x \cos\theta + y \sin\theta, \qquad (1)$$

where θ is the direction of the normal vector of the line and ρ is the distance between the line and the origin. With varying θ (and therefore ρ) this yields all lines passing though a given point (x, y) (Figure 2). The point (x,y) may have been measured by a laser scanner at scanning angle α and range (distance from the scanner) R.

In analogy, Figure 3 shows that a parametric form of a plane in 3D, containing a point (x,y,z) is given by

$$\rho = x\cos\theta\cos\varphi + y\sin\theta\cos\varphi + z\sin\varphi\,. \qquad (2)$$

We will use this equation to establish the distance ρ between a plane and the origin of the Cartesian coordiate syste, after having estimated the two angles θ and ϕ, denoting the orientation of the plane from the range image gradients. This will be done for every pixel in the range image. The resulting three "image features" will then be submitted to a three-band image segmentation algorithm.

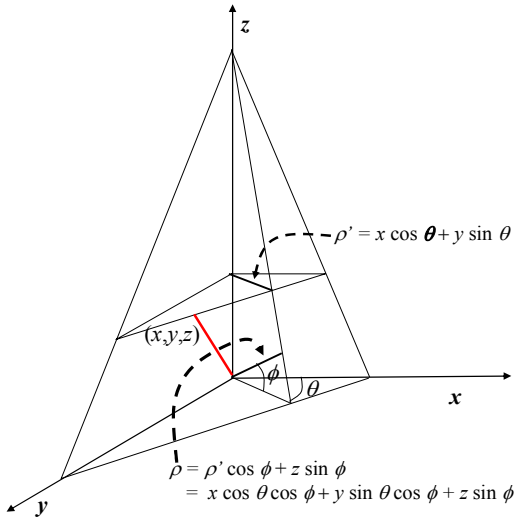

Figure 3. A plane in 3D containing point (x,y,z) and the normal vector to the plane.

Assuming that a certain pixel belongs to a plane P we want to compute the orientation angles θ and ϕ of the normal vector of P, θ being the angle between the x-axis and the projection of the normal vector in the xy-plane, and ϕ being the angle between the normal vector and the xy-plane. We will do this by computing the difference angles $\Delta\theta$ and $\Delta\phi$ between the normal vector and the scan angle, which is given by the coordinates α and β of the range image. This is illustrated in Figure 6, which shows the following relationship between scan angle and normal vector orientation:

$$\theta = \alpha - \Delta\theta$$
$$\phi = \beta - \Delta\phi\,. \qquad (3)$$

3. RANGE IMAGE GRADIENTS

The computation of the difference angles $\Delta\theta$ and $\Delta\phi$ between the normal vector of a plane and the scan angle is based on the gradients of the range image. The horizontal gradient is the change in the image (range) value that is observed when going one pixel to the right; the vertical gradient is the observed change when going one pixel up. Since the image coordinates are related to the scan angles α and β, the gradients can be considered estimates of the partial differences of the range with respect to the horizontal and vertical components of the scan angle, $\Delta R/\Delta\alpha$ and $\Delta R/\Delta\beta$, respectively, with $\Delta\alpha$ and $\Delta\beta$ denoting the angular resolution of the scanner in the`horizontal and vertical directions.

To obtain the corresponding spatial resolution (the perpendicular distance between to neighbouring points), the angular resolution has to be multiplied by the range R $[i_\alpha, i_\beta]$ itself. Finally, the arctangents of $\Delta R/(R\times\Delta\alpha)$ and $\Delta R/(R\times\Delta\beta)$ yield the required difference angles. The procedure is illustrated for the vertical angle ϕ in Figure 4.

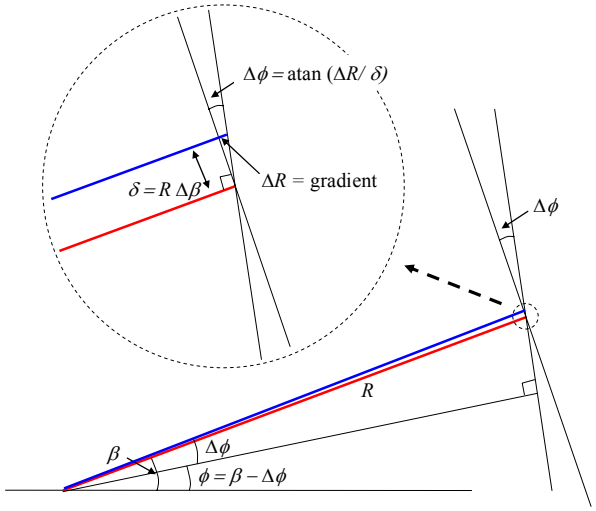

Figure 4. Computing difference angles from range image gradients

Unfortunately, laser scanner measurements are not entirely accurate. In a range image $R[\alpha,\beta]$, both the range measurement R , as well as the scan angles α and β, contain noise, which may severely affect the range image gradients and propagate into the derived estimates of difference angles and plane orientations.

Gradients are computed using convolution filtering with so-called gradient kernels. Examples are Sobel filters (Mather 1999). The noise problem is addressed firstly by using larger filtering kernels, having an smoothing effect, for the gradient computations. In the current example (see Figure 1) we used 5 x 5 kernels,

-1	-2	0	2	1
-2	-3	0	3	2
-3	-4	0	4	3
-2	-2	0	3	2
-1	-1	0	2	1

and

1	2	3	2	1
2	3	4	3	2
0	0	0	0	0
-2	-3	-4	-3	-2
-1	-2	-3	-2	-1

for horizontal and vertical gradients respectively.

A second countermeasure against laser measurement noise is applied during the segmentation step; see Section 4.

The last step in the feature extraction phase is the determination of the third parameter ρ, the distance between the plane containing the pixel at image coordinate $[i_\alpha, i_\beta]$ and the origin, using the plane equation

$$\rho = x\cos\theta\cos\varphi + y\sin\theta\cos\varphi + z\sin\varphi \qquad (4)$$

This equation contains for each range image pixel the 3D coordinate (x,y,z) of the point where the laser beam was reflected. It can be computed from the range image as:

$$x = R[i_\alpha, i_\beta] \cos(i_\alpha \Delta\alpha) \cos(i_\beta \Delta\beta)$$
$$y = R[i_\alpha, i_\beta] \sin(i_\alpha \Delta\alpha) \cos(i_\beta \Delta\beta)$$
$$z = \sin(i_\beta \Delta\beta) . \quad (5)$$

It is very important to note that a large plane, such as the façade of the building at the left in Figure 1, contains a large variety in range measurements, as the different colors indicate. However, after the above-described transformations, even such large plane should become rather homogeneous in the plane parameters θ, ϕ and ρ. This can be seen in Figure 6, which displays colour coded images of these three parameters.

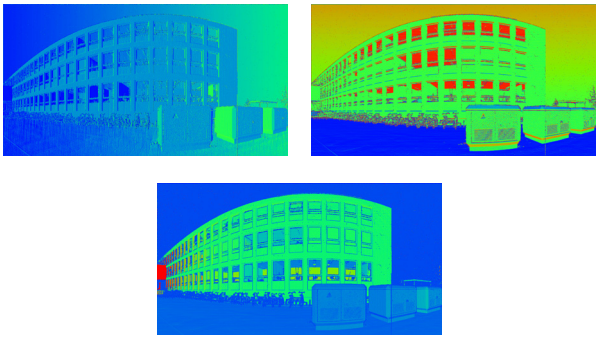

Figure 5. Images of plane parameters θ, ϕ and ρ.

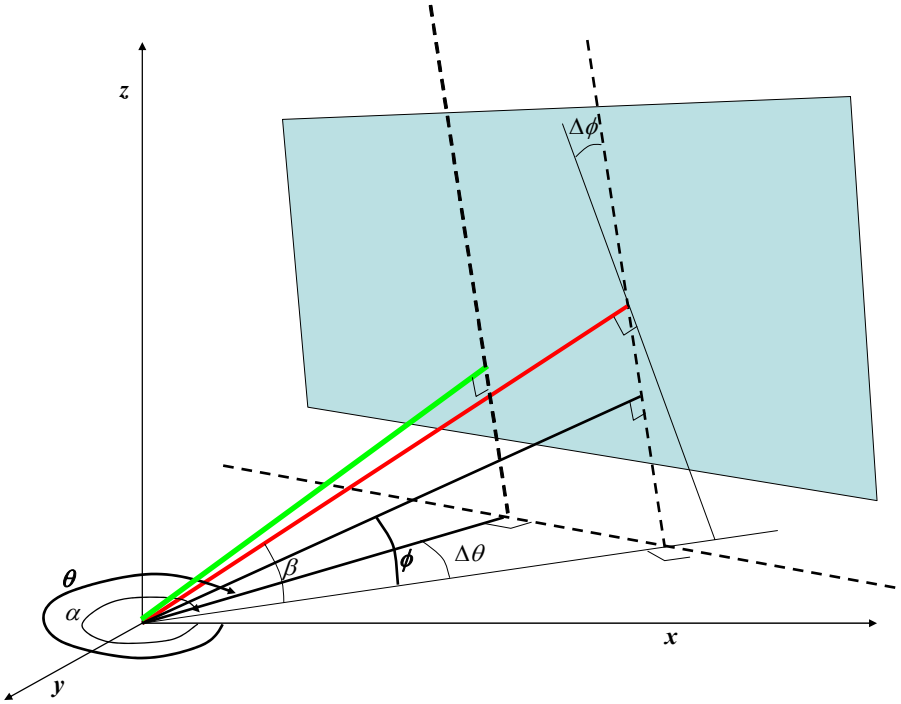

Figure 6. Relationship between scan angle and normal vector orientation

4. IMAGE SEGMENTATION

The purpose of image segmentation is to subdivide an image into adjacent groups of pixels, called segments, which (hopefully) coincide with meaningful objects in the scene. Image segmentation algorithms can be roughly subdivided into region based methods, where pixels within each segment obey some homogeneity criterion, and edge based methods, looking explicitly for boundaries between segments. Within the region based methods popular approaches are region growing (starting from seed pixels, pixels are added to regions as long as homogeneity is sufficiently maintained) and region merging (of adjacent regions that are similar enough).

The image segmentation algorithm used in this study is a region merging method (Gorte, 1999). It was designed for multi-spectral image segmentation, taking three image bands into account simultaneously. It is a quadtree-based method that works bottom-up: merging pixels (quadtree leaves) into segments, and adjacent segments into larger segments while maintaining for each segment a mean feature vector, and as long as two criteria are satisfied:

a) the Euclidian distance between feature vectors of adjacent segments should not exceed a threshold
b) the elements of the variance-covariance matrix within a segment *after* merging should not exceed a threshold.

It has been previously shown that as the second threshold value the squared of the first one can be used, so only one value has to be specified. However, to prevent 'order dependency', the algorithm performs best when applied iteratively with a number of threshold values in a steadily increasing sequence (each iteration being recursive as previously stated). Therefore, the algorithm needs, in addition to the final threshold value, the number of iterations to be performed (usually 3 or 4).

Instead of using three bands of a multi-spectral image, we will now submit the three plane parameters images of Figure 5 to the segmentation algorithm. It will create segments of adjacent pixels that have similar values for θ, ϕ and ρ, and therefore are likely to belong to the same plane in the scene – the word 'similar' indicating that the values may still be contaminated by noise. The amount of noise expected within a single plane determines the threshold value.

Figure 7 shows the result of segmentation. Segments are displayed in arbitrary colours, just to distinguish them from each other. It should be noted, however, that each segment has its "mean feature vector", containing average values for the plane parameters θ, ϕ and ρ. Also the size of each segments (the number of pixels) is known, and the pixels within very small segments can be considered not part of any plane, and can easily be removed from the result in a post-processing step, or assigned to a larger neighbour if required (Figure 8).

5. CONCLUSION AND OUTLOOK

The paper presents a new segmentation algorithm that subdivides a range image created by terrestrial laser scanning into segments that correspond to planar surfaces. Because it works in the 2.5D image domain, rather than in the 3D point cloud domain, the algorithm is quite straightforward and can be implement very efficiently in a suitable image processing environment: it only requires standard processing steps: convolution, image calculation and multi-spectral segmentation.

Working in the original range image data, as delivered by the scanner, the algorithm cannot be applied to point clouds that are created by co-registering multiple scans. A future research topic may be, however, the usefulness of the segmentation method for extracting segments to be used in a feature (or object) based registration process.

A major obstacle in any segmentation effort is posed by measurement noise. It is clear that the effect of noise can be much better quantified, for example by regarding the specifications of the manufacturer (distinguishing between noise in angle and range measurements) and by taking the dependence of noise on the range into account.

REFERENCES

A. Hoover et al., An experimental comparison of range image segmentation algorithms, *IEEE Transaction on Pattern Analysis and Machine Intelligence*, 18(7), 673-689, 1996

X. Jiang, K. Bowyer, Y.Morioka, S.Hiura, K.Sato, S.Inokuchi, M.Bock, C.Guerra, R.E.Loke, J.M.H.duBuf. Some Further Results of Experimental Comparison of Range Image Segmentation Algorithms, *15th Int. Conference on Pattern Recognition*, Spain, Sept. 2000.

Kari Pulli and Matti Pietikainen, Range Image Segmentation Based on Decomposition of Surface Normals, *Scandinavian Conference on Image Analysis (SCIA)*, Norway, 1993.

R. Staiger, Selection criteria for terrestrial laser scanners, *FIG working week*, Hong Kong, May 2007.

Paul M. Mather, *Computer processing of remotely sensed images: an introduction,* second edition, Chichester, NY, Wiley, 1999.

Tahir Rabbani Shah, *Automatic reconstruction of industrial installations using point clouds and images,* PhD thesis, Delft, 2006.

Ben Gorte, *Probabilistic segmentation of remotely sensed images,* PhD Thesis, ITC publication number 63, Enschede, 1999.

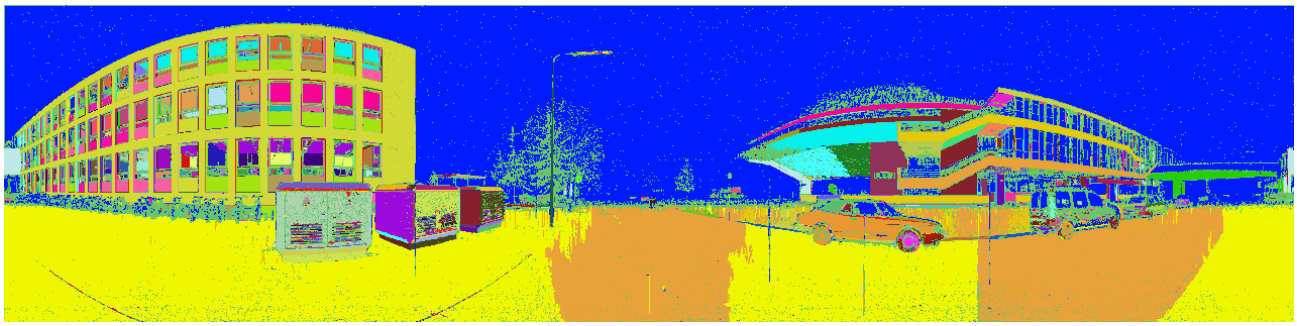

Figure 7. Segmentation result. Oversegmentation in the horizontal plane occurs because the angle θ is not well-defined

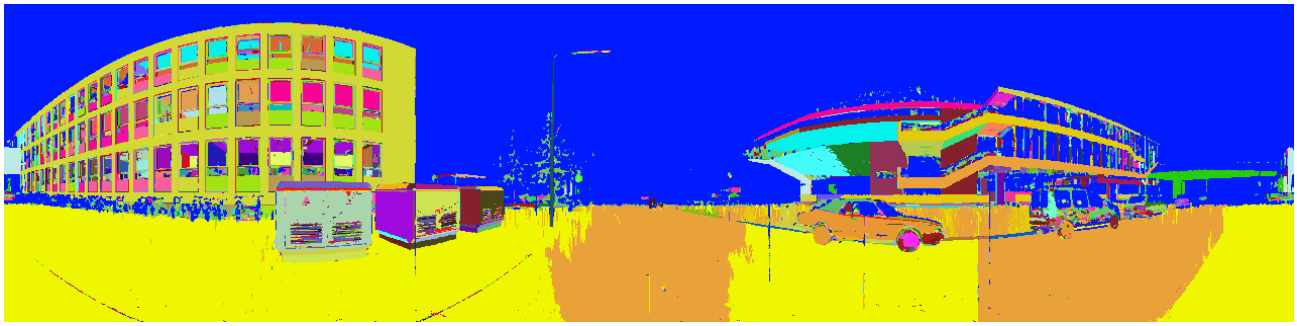

Figure 8. Segmentation result after post-processing to remove small segments