TOWARDS ROAD MODELLING FROM TERRESTRIAL LASER POINTS

Eloïse Denis, Renaud Burck and Caroline Baillard

Siradel
3, allée Adolphe Bobierre
CS 14343, 35043 Rennes, France
edenis@siradel.com, renaud.burck@ensg.eu, cbaillard@siradel.com
http://www.siradel.com

Commission III, WG III/2 and WG III/4

KEY WORDS: Urban, Terrestrial, Laser scanning, Point Cloud, Modelling, City, Feature, Measurement

ABSTRACT:

This paper presents a complete pipeline for extracting and modeling urban roads as 3D surfaces, using laser points acquired from a terrestrial mobile mapping system (MMS) and existing 3D road axes derived from aerial imagery. First ground points are extracted from the laser cloud with an adapted surface growing method. Then the pavement edges are detected by analysing the elevation gradient, then ordered and connected. After registration to the laser data, the existing road axes are associated to the detected pavement edge components. Finally, the road surface and the road width are estimated for each road segment. Our method is illustrated on test data sets describing narrow historical streets. It shows good results, despite the significant slope of the street and the numerous parked cars occluding the curbstones. The road surface is delineated all along the street, and the estimated width is very close to ground truth.

1 INTRODUCTION

City modelling is an expanding subject of study, stimulated by the advent of virtual worlds such as the ones provided by Google Earth or Visual Earth. Modelling an urban scene consists in describing not only buildings (location, textures, roofs, façade details, etc.) and vegetation, but also ground objects such as roads, parking lots, pavement, etc. Thus road modelling has become increasingly complex and generates more and more interest within the scientific community.

This paper presents a new approach for modelling urban roads as 3D surfaces, using vehicle-based laser points and existing road axes. More precisely, the aim of this study is to locate and connect pavement edges, to delineate the road surface and to compute the road width. This information will be used as an input to a more complex road modelling system.

The laser point clouds used for this work have been acquired with a mobile mapping vehicle equipped with two lateral laser scanners. Each one is mounted on a side of the vehicle and produces vertical scan lines, describing any object located between the ground near the vehicle and the top of the lateral buildings. For each data set, right and left laser point clouds are provided. The mobile mapping vehicle is equipped with a very precise georeferencing system composed of a Global Positionning System (GPS), an Inertial Measurement Unit (IMU) and an odometer, that provides precisely geo-referenced laser point clouds. Moreover, as a data provider, Siradel has a huge amount of 3D geographical data bases at an accuracy better than 1 meter, derived from aerial imagery. These data include a geometric description of the road axes, which will be used as input data of our road modelling process.

Section 2 presents some previous work and section 3 describes the principles of our method. Our algorithm was tested on experimental data sets and the results are presented in section 4. Section 5 finally concludes this paper and presents some future evolutions.

2 RELATED WORK

2.1 Ground point detection from terrestrial laser points

Rao et al. have presented two methods to classify ground and non-ground points in a vehicle-based laser point cloud (Rao et al., 2006). They are both based on point projection onto a horizontal plane. The first method consists in generating an accumulation image of the point cloud. Each cell stores the data related to the projected points. The cells containing few points, with small height deviation, are considered as ground cells, and the corresponding points are classified as ground points. This method can unfortunately also detect points belonging to roof slopes or trees. However, as these points float in mid-air, they can easily be detected and removed. The second method takes advantage of the local spatial characteristics of the point cloud. Within the projection plane, points are grouped into large cells and the algorithm iteratively finds regions with low height variations that are spatially correlated. With a modified minimum filter, laser points are classified into ground and non-ground points. This method accurately extracts points belonging to the ground but also points located at the foot of objects (bottom of tires, trees or walls).

Goulette et al. have developed a real-time algorithm to segment vehicle-based laser points into four classes: ground, façades, trees and other (Goulette et al., 2007, Abuhadrous, 2005). Their method is based on prior knowledge on the different classes: the ground is a horizontal plane with a high density of points, the façades are vertical planes, and the trees are free shapes whose projections onto the ground have a specific width. The histograms of each scan line are studied to get a first segmentation of the points. Then fuzzy logic helps to refine the segmentation.

2.2 Pavement edge detection and vectorization

Most work about pavement detection is either based on edge detection in an image of point accumulation or based on histogram analysis.

Shi et al. have worked on automatic extraction of road boundaries and road marks from terrestrial images and laser range data (Shi et al., 2008). They detect pavement edges by analysing individual scan lines, assuming that the road surface is flat and lower than the pavement surface. Aufrère et al. have detected curbstones with the help of a dedicated vehicle (Aufrere et al., 2003). A laser scanner and a camera are located on the front bumper of the vehicle and oriented towards the pavement. For each scan line, the range histogram is thresholded to find the main class. This class represents the curbstone points. This laser-based detection is refined by a camera-based detection. The curbstone positions detected in the laser point cloud are projected on the camera image. A region of interest (ROI) is defined. Within the ROI, a Canny edge detection is performed. The extracted new edge allows to define a new ROI to repeat this operation and extend the detected pavement edge as much as possible. An occlusion detection module, based on laser points, increases robustness in case of occlusions.

Jaakkola et al. have worked on road surface modelling from terrestrial laser point clouds (Jaakkola et al., 2008). The curbstones are detected studying gradients in a height image. Hernández and Marcotegui have also worked with a terrestrial mobile mapping vehicle and have detected curbstones using laser point clouds (Hernández and Marcotegui, 2009). The clouds are first subdivided into building blocks, and artifacts (cars, pedestrians, etc.) are removed. Then the ground points are detected using a cloud segmentation. As some ground points are removed by the artifacts suppression module, a recovery step is performed. All the points belonging to the estimated ground plane are recovered. An elevation image is finally derived from the new cloud of ground points, then segmented into planar regions with a watershed algorithm. The boundaries between planar regions finally give the pavement edges.

Vosselman and Zou have worked on curbstone detection from aerial lidar (Vosselman and Zhou, 2009). Points located on or near the road are first selected, then a threshold is applied on their elevation to classify them as "high points" or "low points". Each high point is associated to the nearest low point and each low point is associated to the nearest high point. Pairs are selected only if they belong to both computed sets of pairs. For each selected pair, the median point is computed and stored as a curbstone point. The curbstone points are linked into a single component when they are near enough. Each component is finally vectorized using a line detection technique based on RANSAC.

3 METHOD DESCRIPTION

We propose a method to delineate the road surface associated to a given road axis. It relies on a combination of existing methods that were adapted to be robust to the road slope and to occlusions caused by parked cars. It also takes advantage of the existing road axes. The whole road surface is delineated and the road width is computed. The approach is composed of four stages, presented in this section:

1. ground point detection;
2. curbstone detection;
3. pavement edge ordering and connection;
4. road surface modelling and width computation.

These stages are independently applied on the right and the left laser point clouds, using the same road axes.

3.1 Ground point detection

The curbstone detection method starts with the extraction of ground points from the laser point cloud. The right and the left clouds are processed independently. The elevation histogram of each vertical scan line is computed. Most ground points belong to the lowest significant peak of the elevation histogram. These points are therefore recorded as seed points then used within an iterative clustering algorithm applied to the whole cloud. Clustering techniques are commonly used for ground detection from airborne lidar data Digital Elevation Models (Vosselman et al., 2004, Tovari and Pfeifer, 2008, Badea and Jacobsen, 2008, Chehata et al., 2008). Our clustering algorithm works as follows. The ground seeds detected at the previous step are iteratively and chronologically stored into a small fixed size queue (typically 500 points), the acquisition order corresponding to the progression along the street.

At each iteration, a mean square plane is computed over the stored points. All the points of the cloud belonging to this plane are marked as ground points. The last acquired ground points are used for updating the queue. Thus, the seed location moves along the street at each iteration and follows the ground curvature. A small queue size will be enable high ground curvature, at the risk of including low objects.

This process is first performed chronologically (seeds are stored following the point acquisition times) then in the inverse direction.

A drawback of this method lies in the fact that it keeps points belonging to the bottom of objects lying on the floor like car tires, pedestrian feet or façade bottoms.

3.2 Curbstone detection

The pavement edge detection is based on the analysis of the elevation gradient computed over the ground point cloud. Thus, it is guaranteed to detect only differences in elevation occurring on the ground level. An image of accumulation is first computed, whose pixel value is the minimal elevation of the corresponding ground points. In order to remove noise, the elevation image is smoothed with a Gaussian filter. Then gradients are computed with a Sobel filter and edges are determined by a hysteresis filter. Thresholds are chosen assuming that pavement height is about 10 cm.

To remove false detections due to object feet as tires, pedestrians or façades, the gradient is also computed over the elevation image of the whole point cloud. The pixels selected at the previous step are studied in the new gradient image. As illustrated in Figure 1, if the difference between the new gradient and the previous one is larger than a predetermined threshold, then the pixel is rejected, because the height variation is too important to correspond to a pavement edge. The threshold value is chosen equal to 20 cm.

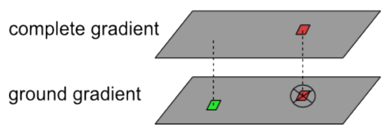

Figure 1: Illustration of the method for removing false detections.

The resulting pixels give a coarse position of the curbstones. They are labelled as connected components representing sections of pavement edges. Each component is then processed individually. The accurate curbstone positions are computed with a method adapted from (Vosselman and Zhou, 2009), described at the end of section 2.2. The ground points are classified as high or low, and pairs of corresponding high and low points are selected. Median points are stored as curbstone point. Figure 2 illustrates its principle. This method was designed for aerial lidar data but works well on terrestrial data. Furthermore, the terrestrial data give direct information about the altitude of the road associated to the pavement, thanks to the vehicle altitude. By subtracting the road altitude to the laser point altitude, it is possible to compute the relative pavement altitude. Thus, the determination of the threshold delimiting high and low points is not influenced by the road slope. At the end, the curbstone points are stored as a set of connected points.

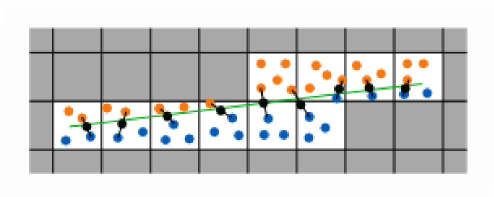

Figure 2: Principle of the curbstone determination method. The real pavement edge is in green (straight line), enlighten pixels of the connected component are in white, high points are in light orange and low points are in dark blue. Selected pairs and their median points (detected curbstone points) are in black.

3.3 Pavement edge ordering and connection

As the curbstone points can be irregularly distributed within a connected component, they need to be ordered and connected within each component. Each component is processed independently. The distances between all curbstone points within a connected component are computed and the two farthest points are stored as pavement edge extremities. A path between these extremities is then found by iteratively selecting the nearest point that has not already been stored as a path point.

As laser points are acquired with a mobile vehicle, the relative acquisition times of the points give the ordered pavement edge sequence along the road. Then different pavement edges are connected if they are near enough. The distance is chosen to connect pavement edges separated by the shadow of a car. The direction of each component is given by the chronological order.

3.4 Road surface delineation

The aim of the road surface modelling stage is first to associate each pavement edge to the available road axes. The road axes are derived from aerial imagery, simultaneously to 3D building models. The road axes are registered to the laser point cloud via the registration of the 3D building models, using the method described in (Denis and Baillard, 2010).

The existing road model is composed of a series of 3D polylines describing road axes. Each polyline is made of "road segments" connecting two successive points. Each segment describes a road portion. A road intersection always implies a node in the model, and consequently a segment end point. However, a road portion without intersection can be described by several segments. The 3D road model brings more information than the vehicle track recorded during data acquisition. The street intersection positions are not provided by the vehicle track unless all the streets are covered.

Each detected pavement edge is associated to the nearest road segment, if the following criteria are respected:

- the segment is close to the pavement edge (typically less than 10 meters);
- the angle between the road segment and the main direction of the pavement edge is small (typically less than 40°).

In the road model, one road can be represented by several segments, even if no crossroad occurs. Thus, a single pavement edge can correspond to several segments. At the previous step, an edge is associated to the nearest segment. It is now necessary to split pavement edges according to the road segments, and to assign the curbstone points to the right segments. Figure 3 illustrates the result expected at this step.

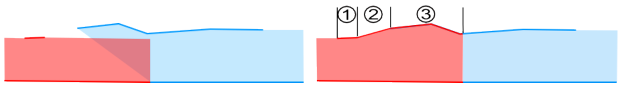

(a) (b)

Figure 3: Scheme illustrating the road segment / pavement edge association. (a) before edge split and merge, (b) after edge split and merge. 1. initial segment edge, 2. extended edge, 3. edge from neighbouring segment.

For this purpose, the points of each pavement edge are projected onto the corresponding road segment axis. Only points projecting on the road segment are kept. Points projecting out of the segment are reassigned to another road segment, according to the following rules:

- the segment is connected to the current one;
- the angle between the road segment and the main direction of the pavement edge is small;
- if several segments respect the two previous conditions, the nearest is selected.

If not any road segment can be chosen, the edge points are lost. If several pavement edges are associated to the same road segment, they are extended and merged. As no road intersection can occur within a road segment, the pavement edges are assumed to be continuous along a road segment.

Finally, when a road segment is associated to an incomplete pavement edge, this latter is extended from its last point assuming a constant road width.

At the end of the process, the pavement edges derived from the right and the left laser point clouds are associated to the corresponding road segments of the existing model. Then, the road surface associated to each road segment can easily be delineated.

3.5 Average road width

The final purpose of this work is to provide a complex road model with geometrical information. In particular, the average width is an information that must be computed and recorded for all the road segments.

The road surface delineation allows to compute the average width of each road segment. The right and the left sides of the segments are processed independently.

Let (x_i, y_i) be the local coordinates of the curbstone point number i, where the system origin is the road segment extremity and the X direction is the road segment direction. The half-width is defined as follows:

$$halfWidth = \sum_{i=1}^{n-1} \frac{y_i - y_{i-1}}{2} \cdot W_i, \quad (1)$$

where:

- n is the number of curbstone points for the current side, ordered and numbered from 0 to $n-1$;
- W_i is a weight given by:

$$W_i = \frac{|x_i - x_{i-1}|}{L}, \quad (2)$$

where L is the road segment length.

The average road segment width is the sum of the left half-width and the right half-width.

4 EXPERIMENTATION AND RESULTS

The algorithm was tested on a data set acquired in Cavell street, a narrow street in the historical town centre of Rennes, in France. This street is characterized by a significant slope and the presence of several parking lots and parked cars (see Figure 4). Both the left and the right laser point clouds were available on this street, but not on perpendicular streets. Figure 5 shows both laser point clouds used for this study. There are many cars occluding the curbstones on the right side. The results at each step of the method, as well as the final result, are presented in this section.

Figure 4: Real view of Cavell street.

4.1 Ground points extraction

First, the ground laser points is extracted from both the right and the left point clouds, using the method presented in section 3.1. Figure 6 shows the resulting ground points. Despite the street slope, the ground points are correctly extracted all along the street.

4.2 Curbstone point detection

Curbstones are then independently extracted from the ground point clouds located on the left and on the right sides of the vehicle, following the method proposed in section 3.2. Figure 7 shows the gradient images associated to the two point clouds. The pixel size is about 13.3 centimetres. These images are binarized in order to provide pavement pixels and corresponding laser points.

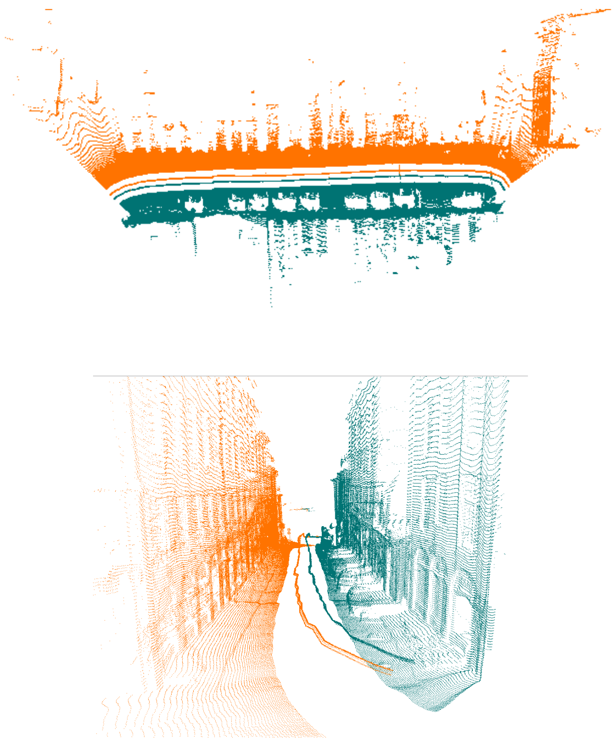

Figure 5: Top and side views of the left (light orange) and the right (dark blue) original point clouds from Cavell street.

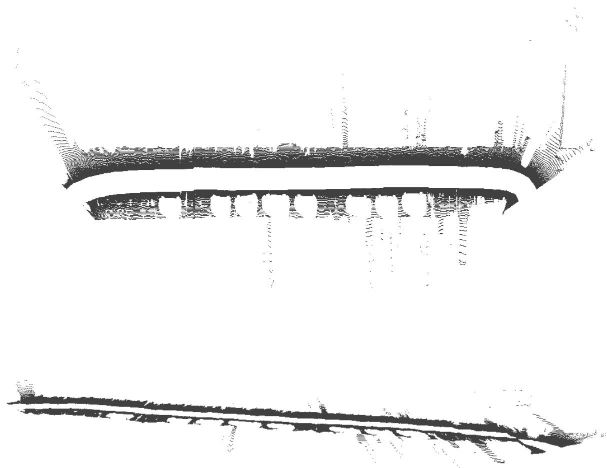

Figure 6: Top and side views of the ground points extracted on Cavell street. The second one shows the street slope.

The curbstone points are then derived from these images, as it is shown in Figure 8. The right laser point cloud contains many parked cars occluding the pavement edges. Some curbstone points are detected between cars, that is sufficient to recover pavement edges despite occlusions.

4.3 Pavement edge ordering and connection

After curbstone points are detected, they are ordered and connected into pavement edges (cf. section 3.3). The result of this step on our test data is presented in Figure 9. The connected components have been linked when they were near enough but the lack of detected curbstone points on the right cloud has produced a break in the right pavement edge.

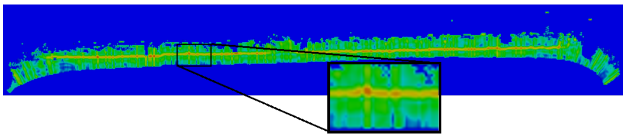

(a) Gradient image derived from the left ground point cloud.

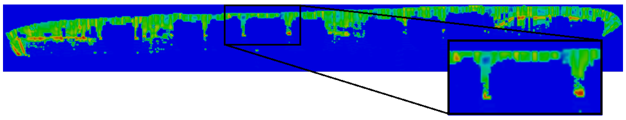

(b) Gradient image derived from the right ground point cloud.

Figure 7: Gradient images derived from the two point clouds. High gradients are yellow and red.

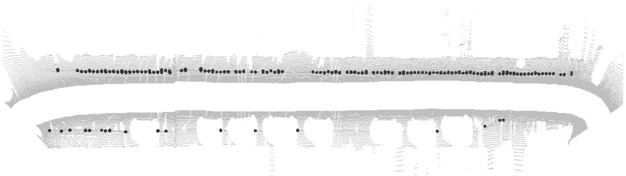

Figure 8: Detected curbstone points (black balls) superimposed with ground points.

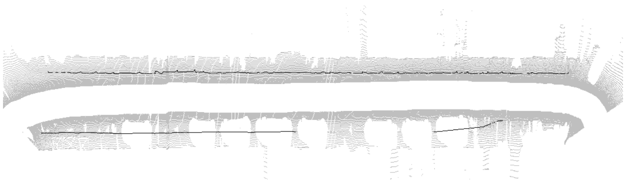

Figure 9: Connected pavement edges.

Note that the computed pavement edge can be locally irregular. This is due to the width of the line derived from the gradient image. For example, it is about two pixels wide for the gradient image of Figure 7(a).

4.4 Road surface modelling

The laser point clouds obtained with the mobile mapping vehicle are compared to the road and building data from the 3D database (see Figure 10(a)). The building façades are first registered to the laser point clouds and the registration parameters are applied to the road segments (see Figure 10(b)).

The pavement edges from the left and the right clouds are then assigned to the nearest road segments to delineate the road surface. The final result of this experimentation is presented in Figure 11, that shows that the street is correctly delineated.

4.5 Average road width

The average width is computed for every processed road segment. The green segment of Figure 11 (the northest one) has an average width of 7.68 meters whereas the red segment (the southest one) has an average width of 7.01 meters. The green segment is wider because it contains more parking lots than the red one, as it can be clearly seen in Figure 11(a). Indeed, it is the roadway width that is estimated rather than the lane width. These results are consistent with manual image-based measurements. Manual measurements show an average value of 7.6 meters for the green segment and 7.0 meters for the red one, which shows an accuracy of a few centimetres.

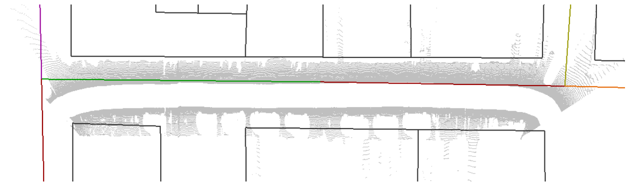

(a) before registration.

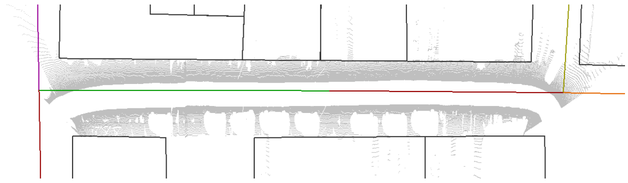

(b) after registration.

Figure 10: 2D view of road segments, façades and ground points.

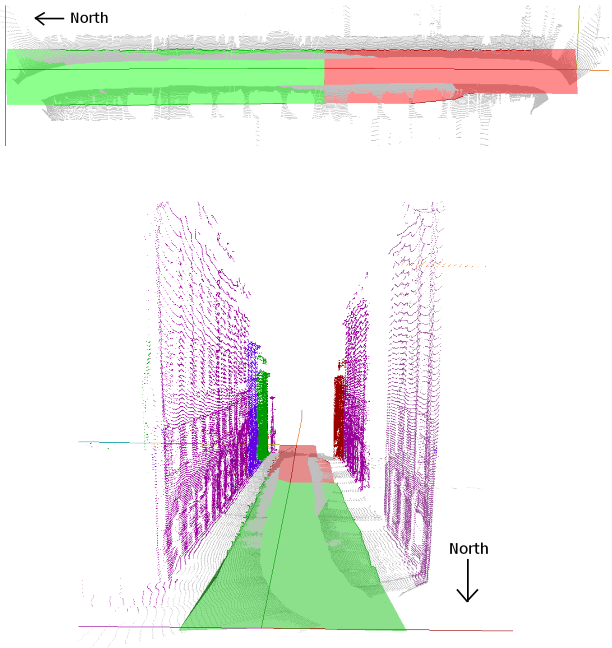

Figure 11: Top and side views of final road surface delineation.

4.6 Other results

The results of the pavement edge splitting step are illustrated on another data set of the same type (cf. section 3.4). Two pavement edge components are detected on the right cloud of the street whereas it is composed of four road segments in the model, as it can be seen in Figure 12(a). Figure 12(b) shows the result of the splitting step. The blue edge has been divided into four parts and each part has been associated to the right road segment. The part associated to the red segment has been rightly merged to the neighbouring pavement edge.

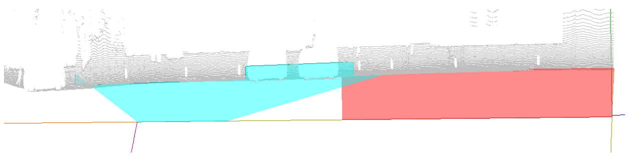

(a) before pavement edge segmentation.

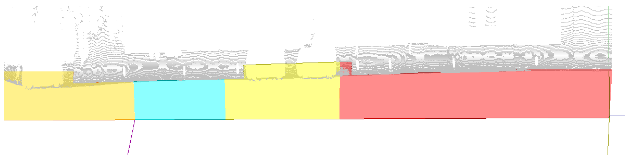

(b) after pavement edge segmentation.

Figure 12: Pavement edges from the right cloud and their associated road segments.

5 CONCLUSIONS AND FUTURE WORK

In this paper, we have proposed a complete approach to delineate the road surface and to estimate the road width from mobile mapping vehicle laser point clouds. The ground detection module successfully extracts ground points, despite significant ground curvature or slope. Curbstone points are detected with a method initially designed for aerial lidar data (Vosselman and Zhou, 2009). We adapted it to terrestrial laser point clouds and proved that it can be used on such points. In particular, this kind of data gives access to the altitude of the laser points relatively to the road altitude, which provides good results even in case of a strong slope.

We managed to obtain connected pavement edges despite parked cars occluding the curbstones. The road surface correctly follows the curbstones visible in the laser point clouds. The validation of this work is mainly graphical and subject to human interpretation. Quality measurement methods should be defined and implemented to improve this work.

The method presented in this paper finally gives several results that can be used in different ways. The ground point cloud obtained at the end of the first stage can be used for façade refinement based on laser point cloud analysis, to find the bottom boundary of the façades (Denis and Baillard, 2010). The road delineation as well as the road width can be used for integration into a complex road model for driving simulators (Platsim project website, 2007).

The algorithm described in the paper is only a first draft. It must be improved and tested on various road configurations to ensure its robustness. Future work will also focus on the detection of changes in the road width and on the determination of the parking lots location.

REFERENCES

Abuhadrous, I., 2005. Système embarqué temps réel de localisation et de modélisation 3D par fusion multi-capteur. PhD thesis, Mines ParisTech (ENSMP).

Aufrere, R., Mertz, C. and Thorpe, C., 2003. Multiple sensor fusion for detecting location of curbs, walls, and barriers. In: Proceedings of the IEEE Intelligent Vehicles Symposium (IV2003), Colombus, Ohio, USA.

Badea, D. and Jacobsen, K., 2008. Filtering process of lidar data. In: Proceedings of the XXI ISPRS Congress, International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences, Vol. XXXVII (Part B3), Beijing, China.

Chehata, N., David, N. and Bretar, F., 2008. Lidar data classification using hierarchical k-means clustering. In: Proceedings of the XXI ISPRS Congress, International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences, Vol. XXXVII (Part B3), Beijing, China.

Denis, E. and Baillard, C., 2010. Refining existing 3D building models with terrestrial laser points acquired from a mobile mapping vehicle. In: Proceedings of the ISPRS CRIMT Conference, International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences, Vol. XXXIX (Part 5), Newcastle upon Tyne, England.

Goulette, F., Nashashibi, F., Abuhadrous, I., Ammoun, S. and Laurgeau, C., 2007. An integrated on-board laser range sensing system for on-the-way city and road modelling. Revue française de photogrammétrie et de télédétection 185, pp. 78–83.

Hernández, J. and Marcotegui, B., 2009. Filtering of Artifacts and Pavement Segmentation from Mobile LiDAR Data. In: Proceedings of the ISPRS Laser scanning Conference, International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences, Vol. XXXVIII (Part 3/W8), Paris, France, pp. 329–333.

Jaakkola, A., Hyyppä, J., Hyyppä, H. and Kukko, A., 2008. Retrieval algorithms for road surface modelling using laser-based mobile mapping. Sensors 8(9), pp. 5238–5249.

Platsim project website, 2007. `http://www.images-et-reseaux.com/en/les-projets/fiche-projets-finances.php?id=114`.

Rao, R., Konda, A., Opitz, D. and Blundell, S., 2006. Ground surface extraction from side-scan (vehicular) lidar. In: Proceedings of the MAPPS/ASPRS Fall Conference, San Antonio, Texas, USA.

Shi, Y., Shibasaki, R. and Shi, Z., 2008. An efficient method for extracting road lane mark by fusing vehicle-based stereo image and laser range data. In: Proceeding of IEEE International Workshop on Earth Observation and Remote Sensing Applications (EORSA), Beijing, China, pp. 1–5.

Tovari, D. and Pfeifer, N., 2008. Segmentation based robust interpolation - a new approach to laser data filtering. In: Proceedings of the ISPRS Laser Scanning, International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences, Vol. XXXVI (Part 3/W19), Enschede, the Netherlands, pp. 79–84.

Vosselman, G. and Zhou, L., 2009. Detection of curbstones in airborne laser scanning data. In: Proceedings of the ISPRS Laser scanning Conference, International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences, Vol. XXXVIII (Part 3/W8), Paris, France, pp. 111–116.

Vosselman, G., Gorte, B. G. H., Sithole, G. and Rabbani, T., 2004. Recognising structure in laser scanner point clouds. In: Proceedings of the ISPRS NATSCAN Conference, International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences, Vol. XLVI (Part 8/W2), Freiburg, Germany, pp. 33–38.