Advanced Point Cloud Processing

GEORGE VOSSELMAN, Enschede, the Netherlands

ABSTRACT

The high pulse frequencies of today's airborne, mobile and terrestrial laser scanners enable the acquisition of point clouds with densities from some 20-50 points/m^2 for airborne scanners to several thousands points/m^2 for mobile and terrestrial scanners. For the (semi-)automated extraction of geo-information from point clouds these high point densities are very beneficial. The large number of points on the surfaces of objects to be extracted describe the surface geometry with a high redundancy. This allows the reliable detection of such surfaces in a point cloud. In this paper various examples are presented on how point cloud segmentations can be used to automatically extract geo-information. The paper focusses on the extraction of man-made objects in the urban environment. The examples include the processing of point clouds acquired by airborne, mobile as well as terrestrial laser scanners. The usage of generic knowledge on the objects to be mapped is shown to play a key role in the automation of the point cloud interpretation.

1. INTRODUCTION

The past two decades have shown a very rapid development and acceptance of various laser scanning technologies for a wide variety of purposes. After a start with airborne laser profilers combined with GPS and inertial navigation systems in 1988 (Lindenberger, 1989), the first scanning airborne laser rangers were introduced in the early 1990's. Since then airborne laser scanning technology advanced from 2 to 250 kHz pulse frequencies, from single echo recordings to multiple echo and full waveform recordings, and from a few decimetre to a few centimetre accuracy (Mallet and Bretar, 2009). Terrestrial laser scanning entered the market in the late 1990's. Terrestrial laser scanners based on the principles of time-of-flight measurement of pulses, phase measurement with continuous waves, or optical triangulation offer solutions for point cloud acquisition at various ranges and with different accuracies (Fröhlich and Mettenleiter, 2004). A few years later mobile laser scanners were added to the spectrum to enable corridor mapping from terrestrial platforms like cars, trains or vessels (Barber et al., 2008).

All three types of laser scanning systems, airborne, mobile and terrestrial, are used to acquire point clouds. The same application may sometimes be served by different sensor platforms. For example, surveying of road and rail road environments is done both by airborne and mobile laser scanning. Likewise, both mobile and terrestrial laser scanners are used in projects to reconstruct building façade models.

Although the point densities and accuracies vary with the type of scanner and distance to the scanned surface, the processing of point clouds acquired by airborne, mobile and terrestrial scanners shows many similarities. In most cases the point clouds will be used to extract geo-referenced information. Common steps in the information extraction procedures are the detection of planar or smooth surfaces and the classification of points or point clusters based on local point patterns, echo intensity and echo count information (Vosselman et al., 2004; Darmawati, 2008).

Point cloud segmentation and classification will briefly be discussed in the next section. The paper then continues with an overview on various point cloud processing projects performed at ITC, Enschede. They show how the results of segmentations, classifications and other point cloud processing can be used for the extraction of geo-information on man-made objects like buildings and roads. The processing of point clouds for the purpose of accurate geo-referencing, DTM production or forestry and engineering applications will be left outside the scope of this paper (Pfeifer and Briese, 2007; Vosselman and Maas, 2009). The applications presented in this paper are grouped according to the platform used for data acquisition. They demonstrate the high quality of point clouds that can nowadays be acquired with laser scanning and show the large potential for the automated and semi-automated extraction of geo-information from this data.

2. POINT CLOUD SEGMENTATION AND CLASSIFICATION

The geometry of man-made objects can often be described by a set of planar surfaces. To a large extent the terrain can be described by smooth surfaces. For extracting the geometry of man-made objects and the terrain algorithms are required that recognise planar and smooth surfaces in a point cloud (Vosselman et al., 2004). The mostly used algorithms are those that first try to find a small set of nearby points with a good fit to a plane. This set of points then constitutes the seed segment for a surface growing procedure in which adjacent non-classified points are added to the segment if their distance to the plane or locally defined smooth surface is below some threshold. Once there are no points left that satisfy this condition, further seed segments are selected and expanded until all points have been assigned to a segment. This procedure is very similar to the well-known region growing algorithm used for image segmentation (Ballard and Brown, 1982). The analysis whether or not a local set of points contains a large percentage of co-planar points that can be used as seed segment is performed with a 3D Hough transform or RANSAC plane detection. These plane estimation methods are very robust and quick if only a small set of points needs to be analysed (Vosselman and Maas, 2009). As the point densities of today's laser scanning surveys typically result into many points on the object surfaces, the detection of these surfaces in the point cloud is usually very reliable.

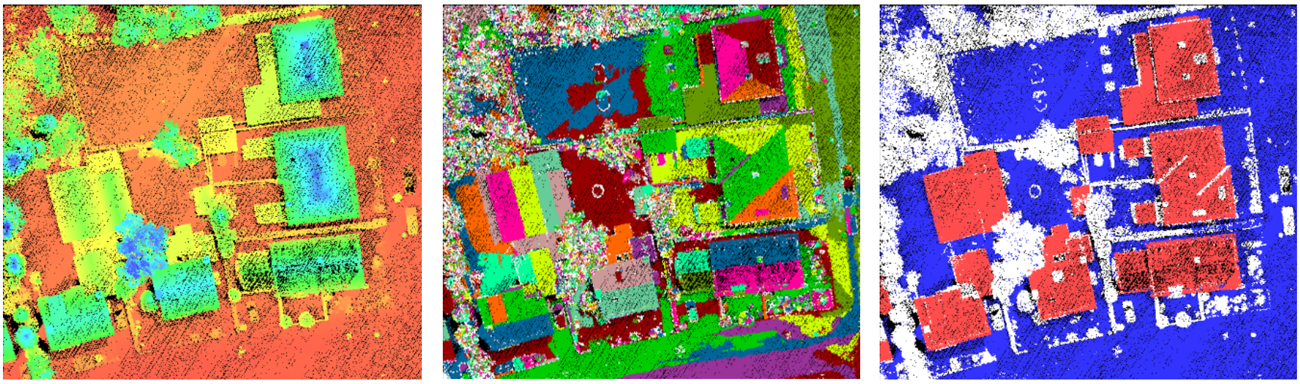

(a) (b) (c)

Fig. 1: (a) Point cloud with colour coded heights, (b) Point cloud segmented into planar surfaces, (c) Segments classified based on various segment properties (red (light grey): building, blue (dark grey): terrain: white: other classes).

Once a point cloud has been segmented, segment attributes can be collected to classify the segments. Like in image processing, a segment-wise classification is more reliable than a point-wise (or pixel-wise) classification. An example of a segment-wise classification is shown in Fig. 1. The point cloud of Fig. 1(a) was segmented into planar surfaces using a 3D Hough transform and a surface growing algorithm. The segmentation result (Fig. 1(b)) shows that most roof planes have a one-to-one relationship with a segment. In the areas with vegetation, many small segments are generated that consist of points in the tree canopy that are approximately co-planar. The segmentation algorithm used a minimum segment size of ten points. Many points in the vegetation could not be grouped to segments of this size and were left without a segment number. These points are shown in white in Fig. 1(b). The terrain surface is represented by various large segments, because it was not planar enough to be described by a single plane. Three segment attributes were

used to obtain the classification of Fig. 1(c): the number of points in a segment, the percentage of last echo points in a segment, and the average height of a segment above a local minimum height. Next to the segment size, the percentage of last echo points is very useful to discriminate between roof segments and vegetation segments (Darmawati, 2008). The points on a roof plane usually are the last echo of a laser pulse (except for a few points on the roof edge). In contrast, the points on a vegetation segment are usually not corresponding to the last echo. Note that the classification is done on segments in a 3D point cloud. This implies that multiple segments may be present on top of each other such that roof segments and terrain segments below vegetation can also be detected.
Other attributes than the ones used above example could be added to further improve the segmentation and classification accuracy. Such attributes include the reflectance strength of the laser pulse, the width of a pulse as extracted from a recorded full waveform (Rutzinger et al, 2008), and multispectral information obtained from a simultaneous recording with an optical camera (Rottensteiner et al., 2005).

3. PROCESSING AIRBORNE LASER SCANNING POINT CLOUDS

Airborne laser scanning has its main application in the production of digital terrain models. With the high point densities of current airborne laser scanners, applications that require higher planimetric accuracies and higher detail become feasible. In this section two projects are described that process point clouds with 20 points/m^2 to extract 3D roof landscapes and road sides.

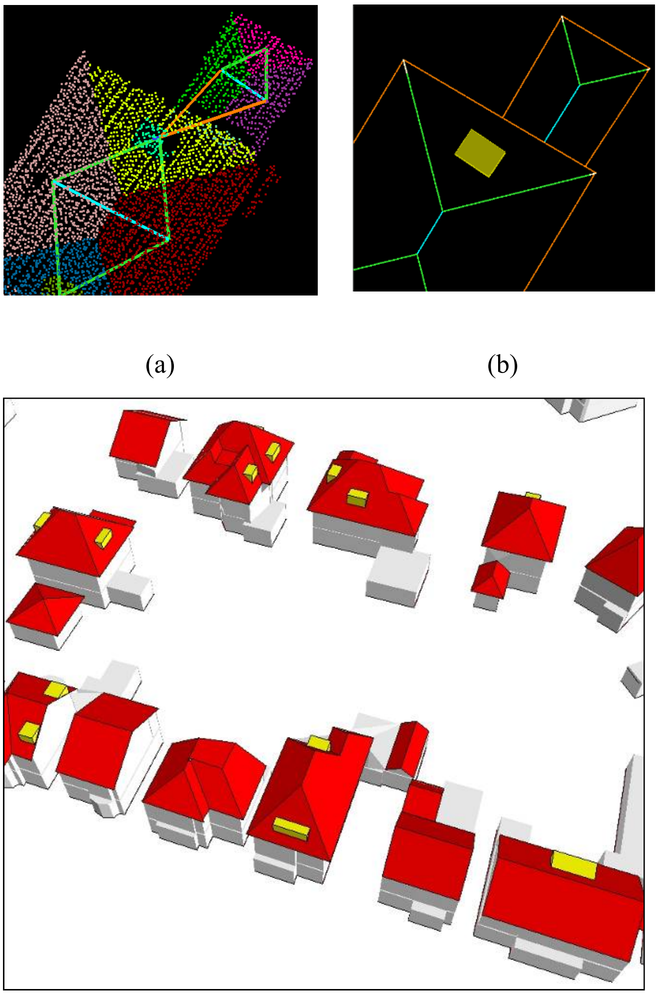

(a) (b)

(c)

Fig. 2: (a) Roof segments with classified relationships. (b) Reconstructed outlines. (c) Reconstructed building shapes.

3.1. Modelling of roof landscapes

The extraction of 3D building models from laser scanning data has been the focus of a large number of studies in the past years (Brenner, 1999, 2009; Rottensteiner, 2003; Vosselman and Dijkman, 2001). Most approaches either are strictly model driven or data driven. Oude Elberink (2009) tries to combine these approaches by utilising graph matching to recognise the topology of common roof shapes. After the segmentation into planar faces, the topology of the roof segments is described by a graph (Fig 2(a)). The detected segments are the nodes of the graph and the edges correspond to pairs of adjacent segments. These edges are labelled by the type of edge (e.g. horizontal intersection line, height jump edge, sloped convex intersection line). Subgraph isomorphisms are sought between the graph of the building segments and a library of roof shapes. The graph matching may result into incomplete matches. In this case, hypotheses are generated for segments that may not have been found in the laser scanning data, but are required to make a topologically correct description. After determining the best match on the topology, the geometry of the roof segments is reconstructed (Fig. 2(b)). The graphs in the roof shape library only describe simple shapes. For more complex roof shapes the graph matching will result in multiple subgraph isomorphisms. The topologies of these matching subgraphs are then combined to one graph and allow the reconstruction of the complex roof shapes (Fig. 2(c)). For suburban areas of a complexity as shown in Fig. 2(c), the topology of roof parts larger than 1.5 m^2 is reconstructed correctly in 75% of the buildings. The largest problems are caused by missing laser data on flat roof parts that were covered by water. As water absorbs infrared light, laser pulses are not or hardly reflected on these roof parts.

The building models of Fig. 2(c) combine the roof outlines as extracted from the point cloud with building outlines obtained from a large scale map. Because the map lines represent the location of the walls, roof extensions are introduced in the models when the point cloud segments extend beyond the walls. The map lines are also used to outline flat roof parts with water, as the lack of points on those roof parts does not allow a data driven reconstruction of the roof geometry.

3.2. Outlining of road sides

In urban areas curbstones are often used to separate the street from the pavement. The height difference between the street and the raised pavement can very well be seen in point clouds acquired by airborne laser scanning. Fig 3(a) shows a point cloud with a colour cycle length of 0.5 m in height. Height differences of 5 cm therefore already appear in different colours. This makes the small height differences at road sides and traffic islands visible. As the noise in the distance measurements by the laser ranger is in the order of 1-2 cm (σ), the signal-to-noise ratio at curbstones enables an automated extraction of curbstone locations. Because all points within a radius of a few meters are acquired within a fraction of a second, the locations of these points all depend on the interpolation between the same two GPS observations. Noise in these observations (σ of 2-3 cm) does therefore not have an effect on the signal-to-noise ratio of local height jumps.

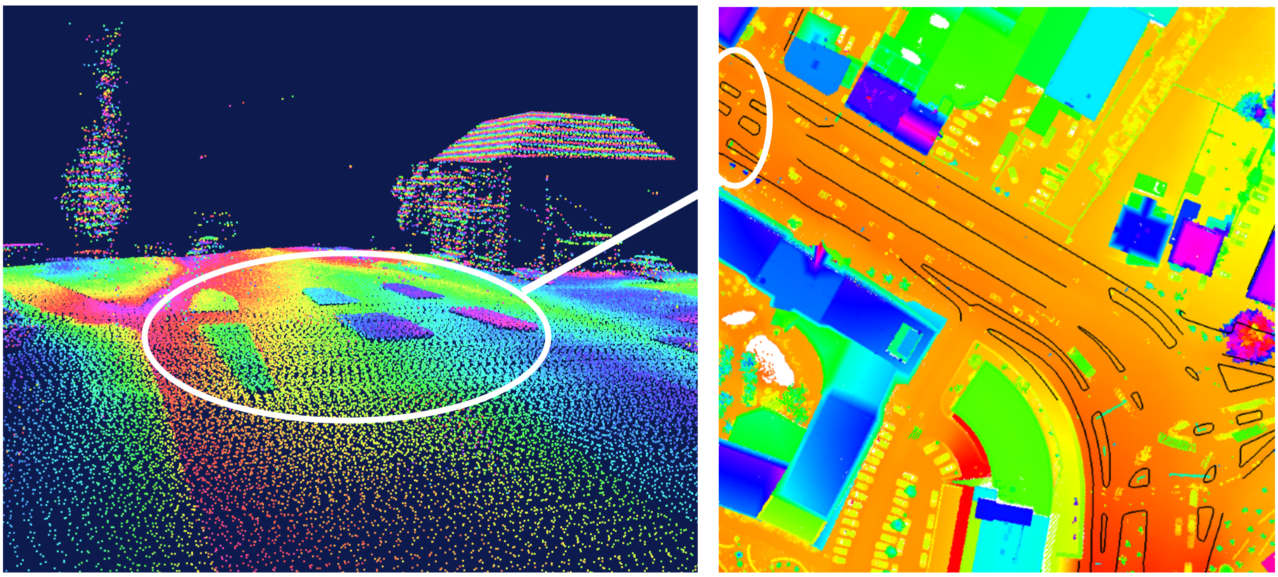

(a) (b)

Fig. 3: (a) Perspective view of colour coded points showing small height differences of traffic islands.
(b) Automatically extracted curbstone locations.

The road sides shown in Fig. 3(b) were extracted with the following procedure (Vosselman and Zhou, 2009). First, a coarse DTM is produced by selecting all low segments of the segmented point cloud. Second, all pairs of nearby points that show a small height difference and are close to the DTM are selected. Third, the mid points of the edges between the nearest selected point pairs are taken as locations of the curbstone. These points are put in a sequence in order to define a polygon. Short sequences can safely be ignored as road sides are expected to be relatively long. Finally, smooth curves are fitted to the point sequences and smaller gaps between collinear curves are closed. These gaps were usually caused by either locally lowered pavement (wheel chair crossings) or curbstones that were occluded by parked cars.

Both the completeness and correctness of the extracted road sides are between 85 and 90% for scenes like the one in Fig. 3. In case many cars occlude larger parts of a road side, the completeness may drop to 50%. This does not affect the correctness of the extracted road sides. The missing road sides in Fig. 3(b) were caused by some larger road side parts without curbstones that could not be bridged automatically. The accuracy of the extracted road sides was estimated to be 0.18 m in a comparison with GPS reference measurements. With some modifications of the smoothing step, accuracies around 0.10 cm seem feasible. This shows the potential of using height data for the extraction of this often required type of topographic information.

4. PROCESSING MOBILE LASER SCANNING POINT CLOUDS

Mobile laser scanners can acquire data at a much higher point density then airborne laser scanners. Point densities may be in the order of 100-1000 points/m^2, depending on the distance to the reflecting object. The absolute accuracy is better than 5 cm standard deviation and comparable to airborne laser scanning at low altitudes. However, this accuracy can only be obtained if a sufficient number of GPS satellites can be tracked. In urban areas with high buildings this may be a problem.

Curbstones as extracted above from airborne laser scanning data are just one of the features that can be acquired as well by surveying the street environment with a mobile laser scanner. Other objects that are acquired for road inventory surveys include traffic lights, traffic signs, road markings, street lights, and buildings and vegetation near the street.

Two examples of processing mobile laser scanning point clouds are presented in this section. The first one is extracting road markings from the reflection strength of the laser scanner. The second example shows the potential of extracting building walls. Brenner (2009) describes a further feature extraction process to automatically locate poles and use those for the relative positioning of cars to their environment.

4.1. Extraction of road markings

Laser scanners usually also record the strength of the reflected laser pulse. This property can be used to distinguish road markings from the road surface. Figs. 4(a) and (b) show a mobile laser scanning data set to which two different thresholds on the reflectance strength were applied. With the higher threshold (Fig. 4(a)) only points on road markings near the path of the vehicle are selected. The lower threshold (Fig. 4(b)) leads to the selection of points on more distant road markings, but also to the selection of many points on the road surface close to the vehicle. The reflectance strength, of course, depends on the distance of the reflecting surface to the laser scanner. To properly select the points based on reflectance strength, the threshold should be specified as a function of the distance, or the reflectance should be normalised prior to the thresholding. With large variations in distances this normalisation is much more important than for airborne laser scanning (Höfle and Pfeifer, 2007).

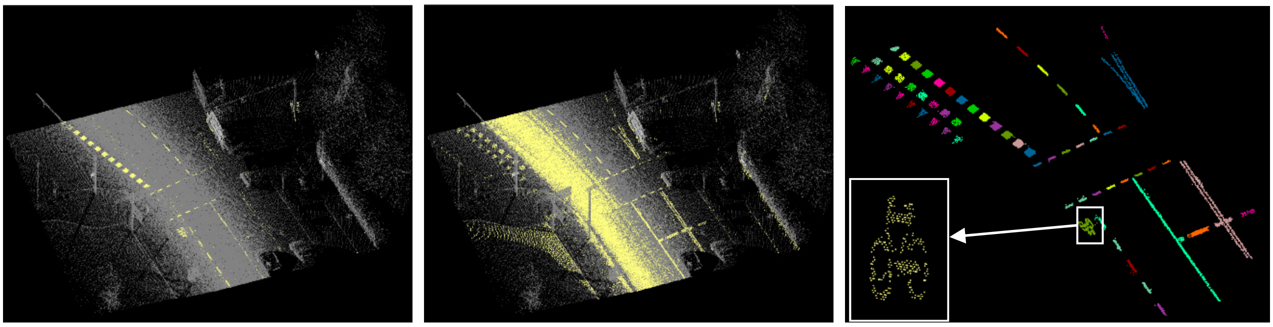

(a) (b) (c)

Fig. 4: (a) Point selected with high threshold on reflectance. (b) Points selected with low threshold on reflectance. (c) Detected road markings after distance dependent thresholding and connected component analysis.

With a distance dependent threshold all points on road marks can be selected successfully. Points with high reflectance values on cars and other objects are easily removed by including the distance to the ground level as a selection criterion. With a connected component analysis of the remaining points, points can be grouped to mark segments (Fig. 4 (c)). Full outlining of the markings can then be obtained by fitting predefined shapes to those segments. Alternatively, one could also use the location of the detected marks as an approximate value for a more accurate outlining in photographs. The detection of markings is, however, easier done in a point cloud. Segmentation of a photograph may also be used to detect bright spots on a dark background, but the point cloud processing enables an easier determination of the distance of such a bright spot to the ground and will therefore have a lower false alarm rate.

4.2. Extraction of walls

Beside the usage in typical corridor mapping applications, mobile laser scanning can also be used to acquire the building façade geometry to generate realistic building models for visualisations at street level. Rutzinger et al. (2009) presented a first analysis on the automated extraction of building walls from mobile laser scanning data. Building walls were extracted from the point cloud by a segmentation of the point cloud into planar faces, followed by a selection of segments that could be walls. Segments were classified as wall segment when the inclination was less than 3°, the segment dimensions were larger than 2 m in height and 0.5 m in width and the segment contained more than 1000 points.

The results of this study are presented in Fig. 5. The left figure shows the building outlines of the large scale map (yellow, grey) and the black lines are the walls that are theoretically visible from the perspective of the survey path (shown by the dots). The right figure shows the extracted vertical point cloud segments that were classified as wall, again overlaid on the building outlines of the large scale map. It shows that the walls facing the street could be identified well. Side walls, however, proved to be difficult to extract. In total the extracted walls are only 56% of those that are theoretically visible. Further analysis of this result has to be performed, but it is assumed that the major reason for the low detection rate of the side walls is the occlusion by fences and vegetation.

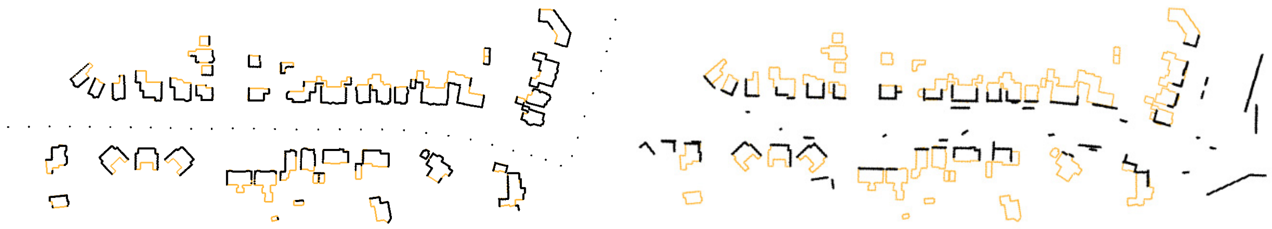

(a) (b)

Fig. 5: (a) Theoretically visible walls and (b) automatically detected walls of a mobile laser scanning survey. The survey path is shown by dots. The yellow (light grey) lines are taken from the building layer of a large scale map.

5. PROCESSING TERRESTRIAL LASER SCANNING POINT CLOUDS

Building façades have also been modelled from data acquired by terrestrial laser scanners mounted on a tripod. The point densities obtained with these scanners may be even higher than those of mobile laser scanners. Within the data recorded at a single scan position, object dimensions can be obtained more accurately than with mobile laser scanning as these measurements are not influenced by platform positioning errors. Yet, for recording different sides of an object multiple scan positions are required and the point clouds of the scans have to be registered with respect to each other. The first example in this section shows the processing of a terrestrial laser scanning point cloud for the detailed reconstruction of a building façade. The second example shows how a building model can be aligned to optical imagery in order to improve the photorealistic rendering.

5.1. Extraction of façade models

Like for the detection of walls in mobile laser scanning data, the first step in the processing is the segmentation into planar pieces as most building parts can be described by planar surfaces. The segmentation result in Fig. 6(a) shows that even though doors, window frames or curtains may only be a few centimetres behind the plane of the wall, they are detected as different segments (Pu and Vosselman, 2009a). By formulating rules on the possible sizes, relative position, and orientation of the building components, segments can be classified into the categories of wall, roof faces, doors, protrusions, and ground surface (Fig 6(b)). As window frames are often partially occluded, their detection is the point cloud is not reliable. However, windows can be easily found by outlining the gaps in the wall segment that are not explained by doors and protrusions. Texturing the model with registered imagery and artificial textures for less visible parts is used to obtain the visualisation of Fig. 6(c).

(a) (b) (c)

Fig. 6: (a) Segmented point cloud of a building façade. (b) Segments classified as wall (blue, dark grey), door (orange, grey), roof (yellow, light grey), protrusion (turquoise, light grey) or ground (green, grey). (c) Textured building model.

5.2. Matching models with imagery for accurate texturing

Photo textures are often applied to obtain a realistic visualisation. Unfortunately, the human eye is very sensitive to discrepancies between the geometric model and the texture. Such misfits may be caused by incorrect orientation parameters of the photograph or model, image distortions or errors in the model. Fig. 7(a) shows a part of a panoramic photograph that was used for texturing a building model reconstructed from a terrestrial laser scanning point cloud. The texture applied in Fig. 7(b) includes a small part of the sky. Producers of realistic visualisations often spend considerable time to repair these kinds of errors. To automatically obtain a better alignment edges can be extracted from the image and matched against the edges of the reconstructed (3D) model (Fig. 7(c), Pu and Vosselman, 2009b). For the generation of Fig. 7(d) it was assumed that the misfit was caused by errors in the model. The outlines of the front façade were therefore shifted in the plane of the façade in order to obtain good correspondence with the projected image edges.

(a) (b) (c) (d)

Fig. 7: (a) Image used for texturing, (b) Initial texture projection, (c) Extracted edges from model (blue) and image (red, pink), (d) Texture projection after adapting the model to the image edges.

6. CONCLUSIONS

The high pulse frequencies of current laser scanners allows to generate point clouds in which object surfaces are captured by hundreds to thousands of points. As described in the above examples, these high point densities enable a reliable segmentation of point clouds into planar and smooth surfaces. Such segmentations form the basis for semi-automated or automated geometric modelling of man-made objects.

The interpretation of segments requires the modelling of knowledge on the objects that are to be reconstructed. Progress in knowledge modelling will be the key to further automation in mapping from point clouds. In this sense, point cloud understanding encounters the same problem as image understanding, even though the 3D features (segments) of point clouds may be richer than the points and edges extracted from imagery.

Photogrammetric workstations are equipped with software tools that have been optimised over several decades. Only in recent years commercial software packages have become available for information extraction from point clouds. It is likely that there is still much room for improvement and that the next years will show more advanced tools for semi-automated mapping in point clouds. As imagery is often acquired together with laser scanning surveys, it would be desirable that tools would be developed that enable the operator to simultaneously work with both data sources, or at least to easily select the source that is most suitable for mapping the object at hand.

7. ACKNOWLEDGEMENTS

The mobile laser scanning data was kindly provided by TopScan GbmH. The terrestrial laser scanning data was kindly provided by Oranjewoud B.V. Panoramic imagery was kindly provided by Cyclomedia B.V.

8. REFERENCES

Ballard, D. H., Brown, C. M. (1982): Computer Vision, Prentice Hall, Englewood Cliffs, USA.

Barber, D., Mills, J., Smith-Voysey, S. (2008): Geometric validation of a ground-based mobile laser scanning system. ISPRS Journal of Photogrammetry and Remote Sensing 63 (1), 128-141.

Brenner, C. (1999): Interactive Modeling tools for 3D Building Reconstruction. In: Photogrammetric Week '99, Fritsch, D., Spiller, R., (Eds.), Herbert Wichmann Verlag, Heidelberg, pp. 23-34.

Brenner, C. (2009): Building extraction. In: Airborne and Terrestrial Laser Scanning, Vosselman, G., Maas, H.-G. (Eds.), Whittles Publishing. To appear in autumn 2009.

Brenner, C. (2009): Extraction of Features from Mobile Laser Scanning Data for Future Driver Assistance Systems. In: Sester, M., Bernard, L., Paelke, V. (Eds.), Advances in GIScience, Lecture Notes in Geoinformation and Cartography, Springer, pp. 25-42.

Darmawati, A. T. (2008): Utilization of multiple echo information for classification of airborne laser scanning data. Master's Thesis, International Institute for Geo-Information Science and Earth Observation (ITC), Enschede, the Netherlands.

Fröhlich, C., Mettenleiter, M. (2004): Terrestrial laser scanning – new perspectives in 3D surveying. The International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences, vol. 36, part 8/W2, Freiburg, Germany, pp. 7-13.

Höfle, B., Pfeifer, N. (2007): Correction of laser scanning intensity data: Data and model-driven approaches. ISPRS Journal of Photogrammetry and Remote Sensing 62 (6), 415-433.

Lindenberger, J. (1989): Test results of laser profiling for topographic terrain survey. Proceedings 42nd Photogrammetric Week, Stuttgart, pp. 25-39.

Mallet, C., Bretar, F. (2009): Full-waveform topographic lidar: State-of-the-art. ISPRS Journal of Photogrammetry and Remote Sensing 64 (1), 1-16.

Oude Elberink, S. (2009): Target graph matching for building reconstruction. The International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences, vol. 38, part 3/W8, Paris, 1-2 September, accepted.

Pfeifer, N., Briese, C. (2007): Geometrical aspects of airborne laser scanning and terrestrial laser scanning. The International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences, vol. 36, part 3/W52, pp. 311-319.

Pu, S., Vosselman, G. (2009a): Knowledge based reconstruction of building models from terrestrial laser scanning data. ISPRS Journal of Photogrammetry and Remote Sensing, Article in Press, doi:10.1016/j.isprsjprs.2009.04.001.

Pu, S., Vosselman, G. (2009b): Building façade reconstruction by fusing terrestrial laser points and images. Sensors 9 (6), 4525-4542.

Rottensteiner, F. (2003): Automatic generation of high-quality building models from LIDAR Data. IEEE Computer Graphics and Applications 23 (6), 42-50.

Rottensteiner, F., Trinder, J., Clode, S., Kubik, K. (2005): Using the Dempster-Shafer method for the fusion of LIDAR data and multi-spectral images for building detection. Information fusion 6 (4), 283-300.

Rutzinger, M., Höfle, B., Hollaus, M., Pfeifer, N. (2008): Object-Based Point Cloud Analysis of Full-Waveform Airborne Laser Scanning Data for Urban Vegetation Classification. Sensors 8 (8), 4505-4528.

Rutzinger, M., Oude Elberink, S., Pu, S., Vosselman, G. (2009): Automatic extraction of vertical walls from mobile and airborne laser scanning data. The International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences, vol. 38, part 3/W8, Paris, 1-2 September, accepted.

Vosselman, G., Dijkman, S. (2001): 3D Building Model Reconstruction from Point Clouds and Ground Plans. The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, vol. 34, part 3/W4, Annapolis, MA, USA, October 22-24, pp. 37- 44.

Vosselman, G., Gorte, B.G.H., Sithole, G., Rabbani, T. (2004): Recognising structure in laser scanner point clouds. The International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences, vol. 46, part 8/W2, Freiburg, Germany, 4-6 October, pp. 33-38.

Vosselman, G., Maas, H.-G., Eds. (2009): Airborne and Terrestrial Laser Scanning. Whittles Publishing. To appear in autumn 2009.

Vosselman, G., Zhou, L. (2009): Detection of curbstones in airborne laser scanning data. The International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences, vol. 38, part 3/W8, Paris, 1-2 September, accepted.