Quality Assessment of Automatically Generated Feature Maps for Future Driver Assistance Systems

Sabine Hofmann
Institute of Cartography and Geoinformatics
Leibniz Universität Hannover
Appelstr. 9a, 30167 Hannover, Germany
sabine.hofmann@ikg.uni-hannover.de

Claus Brenner
Institute of Cartography and Geoinformatics
Leibniz Universität Hannover
Appelstr. 9a, 30167 Hannover, Germany
claus.brenner@ikg.uni-hannover.de

ABSTRACT

Future driver assistance systems will require highly accurate positioning. One way to achieve this is by using on-board sensors to measure the relative location of landmarks for which the absolute coordinates are known.

This paper investigates the use of mobile laser scanning for the fully automatic generation of such landmark maps. Starting from a 21.7 km scanned trajectory with a total of 70.7 million scanned points, we extract pole-like structures, such as signposts, traffic and street lights, and tree trunks. The location of all those structures then forms our landmark map to be used later by on-board systems for positioning. The focus of this paper is on the extraction and quality assessment of the features, including a description of different types of error sources and approaches to reduce false positives among the extracted poles.

Categories and Subject Descriptors

H.2.8 [**Database Management**]: Database applications – *Spatial databases and GIS;* I.2.9 [**Artificial Intelligence**]: Robotics – *Autonomous vehicles;*. I.4.8 [**Image Processing and Computer Vision**]: Scene Analysis – *Object recognition, range data;* I.5.4 [**Pattern Recognition**]: Applications – *Computer Vision.*

General Terms

Algorithms, Measurement, Performance, Experimentation.

Keywords

Mobile laser scanning, feature extraction, localization, autonomous vehicles, driver assistance systems, landmark based maps.

1. INTRODUCTION

In recent years, research on safety systems has revealed that a major improvement beyond the current safety level requires a transition from passive (such as anti-locking brakes) to active systems, which prevent dangerous situations altogether, instead of just mitigating the consequences of accidents. The so-called advanced driver assistance systems (ADAS) operate at a very detailed scale, requiring large scale maps. In the 'Enhanced Digital Mapping' project, it was concluded that contemporary mapping techniques would be too expensive to provide appropriate maps [4]. The key observation is that these findings assume a 'traditional' map production, i.e. the map is a highly abstract representation in terms of a vector description of the geometry, with added attributes. However, depending on the application, this is not always necessary. For example, the accurate positioning of vehicles using relative measurements to existing map features does not necessarily require a vector map in the usual sense. Rather, any map representation is acceptable as long as it serves the purpose of accurate positioning.

Terrestrial laser scanning has made huge progress during the last years, with systems now being able to measure at accuracies of a few millimeters (using a stationary scanner) and at a rate of more than 100.000 points per second. Being combined with GPS and inertial sensors, mobile laser scanning (MLS) systems available today are well suited for the production of large scale maps, since they capture the environment in great detail and reach a relative accuracy of down to a few centimeters [7].

The problem of building a representation of the environment has also been investigated in robotics. Iconic representations, such as occupancy grids, have been used as well as symbolic representations consisting of line maps or landmark based maps [3]. One of the major problems to solve is the simultaneous localization and mapping (SLAM), which incrementally builds a map while a robot drives (and measures) in unknown terrain [9]. It is probable that it will still take some time until full 3D scanners are available in vehicles. What can be expected in the near future, though, are scanners which scan in several planes such as the close-to-production IBEO Lux four-plane scanner [5].

Apparently, it would be unreasonable to provide dense georeferenced point clouds for the entire road network, as this would imply huge amounts of data to be stored and transmitted. On the other hand, obtaining a highly abstracted representation usually requires manual editing, which would be too expensive. Following the work in robotics, a landmark based map is a suitable approach. Using landmarks such as poles of traffic signs and traffic lights for positioning has been investigated earlier, e.g. by Weiss et al. [10]. They combine GPS, odometry, and a (four-plane) laser scanner to estimate the position of a vehicle, given a previously defined map of poles. Their results show that, given such a map, a much higher positioning accuracy can be obtained

ACM GIS '09 , November 4-6, 2009. Seattle, WA, USA (c) 2009 ACM ISBN 978-1-60558-649-6/09/11...$10.00

than using the standard GPS/ odometry solution alone, with reported accuracies down to a few centimeters. It is the purpose of this paper to describe a fully automatic pole extraction based on mobile laser scanning data, and to assess its results in terms of the error rate. Investigations regarding the achievable positioning accuracy and the global uniqueness of local pole patterns were presented in [1] and [2], respectively.

2. DATA ACQUISITION AND POLE EXTRACTION

For our experiments, we obtained a dense laser scan of a set of roads in Hannover, Germany, acquired using the Steetmapper mobile mapping system [6]. The scan was acquired with a configuration of four scanners, each with a maximum range of 150 m and a ranging accuracy of 25 mm. All scanners were operated simultaneously at the maximum scanning angle of 80° and scanning rate of about 10.000 points/s. Positioning was accomplished using IGI's TERRAcontrol GNSS/IMU system which consists of a GNSS receiver, a fiber optic gyro IMU operating at 256 Hz, an odometer, and a control computer. Using the calibrated relative orientations of the scanners and the time synchronization, a georeferenced point cloud is obtained. Single point absolute accuracies of 10 cm and better can be expected under good conditions, however, in cities, with prolonged occlusions from trees and buildings, it can be 1 m and worse. It is important to note that even in this case, the *relative* accuracy (which is the crucial parameter for relative positioning) of the points will still be around a few centimeters. Figure 1 shows an overview of the captured scene. Note that the scanned area contains streets in densely built-up regions as well as highway like roads. The total length of the scanned trajectory is 21.7 kilometers, captured in 48 minutes, which is an average speed of 27 km/h. During that time, 70.7 million points were captured, corresponding to an effective measurement rate of 24,500 points per second. On average, each road meter is covered by more than 3,200 points.

The automatic extraction of simple shapes (like cylinders) from laser scanning data has already been investigated in other contexts, e.g. the alignment of multiple terrestrial scans in industrial environments [8]. However, single poles are not hit by very many scan rays. Thus, methods which rely on the extraction of the surface, of surface normal vectors, or even of curvature, are not applicable. Therefore, we use a simple geometric model for pole extraction, namely that the basic characteristic of a pole is that it is upright, there is a kernel region of radius r_1 where laser scan points are required to be present (the pole), and a hollow cylinder, between r_1 and r_2 ($r_1 < r_2$) where no points are allowed to be present (the area around the pole, see Figure 2, left). The structure is analyzed in stacks of hollow cylinders. A pole is confirmed when a certain minimum number of stacked cylinders is found (Figure 2, right). The method also extracts some tree trunks of diameter smaller than r_1, which we do not attempt to discard since they are useful for positioning purposes as well.

For the entire 22 km scene, a total of 2,658 poles was found fully automatically, which is one pole every 8 meters on average. In terms of data reduction, this is one extracted pole per 27,000 original scan points. Although the current implementation is not optimized, processing time is uncritical and yields several poles per second on a standard PC. Of course, poles are not distributed uniformly, and from Figure 3 it can be seen that there are too few along the highway in the lower left corner.

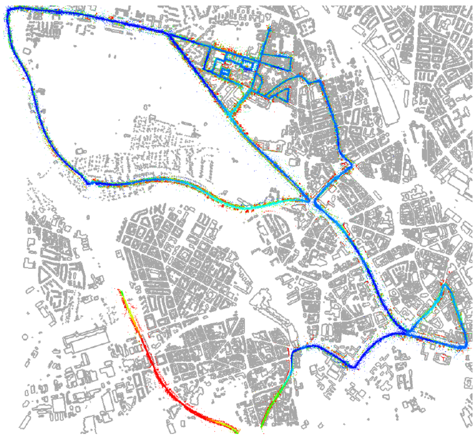

Figure 1. All scanned points along the trajectory (color encodes absolute height using a temperature scale) with ground plans from the cadastral map overlaid.

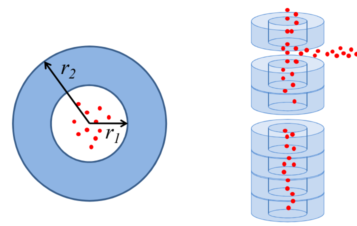

Figure 2. Left: Geometric pole detection using two radii. Right: Analysis in terms of cylindrical stacks.

Figure 3. Extracted poles (blue dots) in a larger area with buildings from the cadastral map (green). The thin dotted line is the trajectory of the Streetmapper van.

3. REDUCTION OF FALSE POSITIVES

To create a robust and reliable pole searching algorithm it is not only important to know how many of the extracted poles were found correctly but also what causes failings. Therefore, we classified the extracted poles manually (46% trees, 21% street lights, 5% traffic lights, 5% tram poles, 8% sign posts and 9% others). The overall error rate was about 6%. In densely built-up areas the error rate is higher than in non-urban regions. The number of errors only includes false positives, as we did not attempt to map all potential objects manually.

3.1 Types of Errors

The false positives can be divided into three groups. The main group (about 60% of the false positives) results from a variable point density within the dataset. The point density is dependent on the driving speed, the distance between the object and trajectory of the scanning vehicle and the angle between the laser beam and the object normal. The distance between the trajectory and the scanned object especially plays a role when the heading of the vehicle changes (Figure 4, left). Furthermore, it is important to take into account the angle between the laser beam and the object normal. The highest point density is achieved when scanning parallel to the object normal (Figure 4, right). Thinking of the pole extraction algorithm, a low point density causes several different effects on the results. If points are scanned in vertical, pole-like structures may also be detected within vertical planes, e.g. facades (Figure 8).

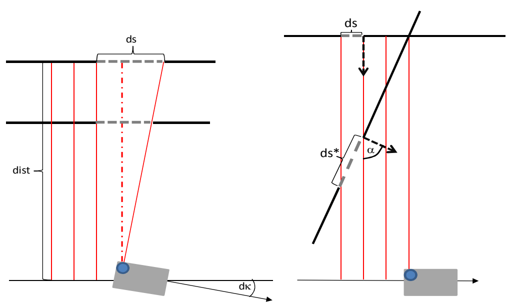

Figure 4. Left: The distance *ds* of neighboring points depends on the distance between trajectory and measured object and the heading difference. Right: As the angle between the object normal (dashed black arrow) and the laser beam increases, the point density decreases by a factor of cos(α).

The second group of false positives contains objects which appear to lie inside of buildings, so that they should actually not be visible to the laser scanner. However, as shown in Figure 5 there are different reasons why it is possible to get a signal from behind walls.

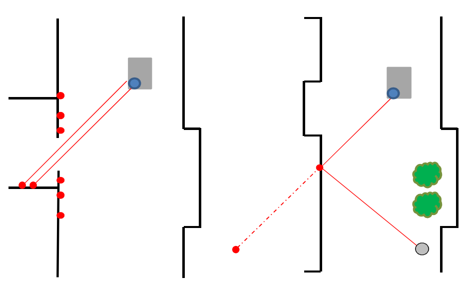

Figure 5. Left: Measurements through a window generate pole-like objects behind facades, when the scanning direction is vertical. Right: Measurements on a reflective surface. The scanned point is assumed to be located at the end of the dashed red line, although it belongs to an object on the opposite side of the road.

The third group of errors includes hovering poles, i.e. poles which are detected at a certain height but do not extend down to the ground. Some of these may not be actual errors, e.g. pole-like objects on rooftops. So even though these objects may be poles according to the given definition, they would probably not be visible for an in-vehicle scanner system with only one horizontal scan level. In the following, two different approaches are shown to reduce false positives using the point density and one approach regarding hovering poles.

The first approach is to compute the *beam density* for any pole. In this case, the beam density is defined as the number of laser beams intersecting a given area surrounding any extracted pole. With regard to the basic searching algorithm for poles, the surrounding area is defined as the area between the inner (r_1) and the outer (r_2) searching radii. The main idea is to ascertain that there is a reasonable amount of rays reflected from behind a detected pole. False positives e.g. within or behind facades normally have only a small beam density (Figure 6). However, the approach fails for objects which stand in front of an open space. This problem is actually due to the fact that in the current setup, we are only able to reconstruct rays which belong to a measured point. We are unable to count rays which pass the surrounding area but do not hit any object at all.

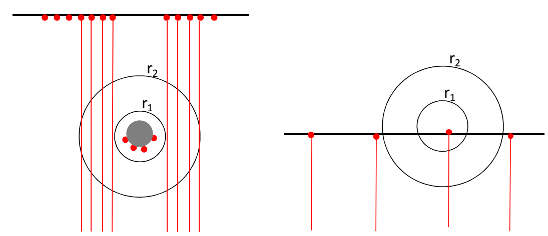

Figure 6. Left: Usual configuration: a pole (solid grey circle) lying in front of e.g. a facade. Right: In case of a low point density, the algorithm may erroneously detect poles within facades (black line). Since there is no object behind the pole, there is no laser beam intersecting the area around the pole.

The second approach is to compute a *theoretic point density* around every extracted pole and exclude all poles located in an area of marginal density. The theoretic point density is computed by using the position, velocity and orientation of the mobile mapping vehicle relative to the scanned object (as illustrated in Figure 4). With a fixed threshold, areas with a low point density are completely excluded.

Finding *hovering poles* in the dataset is done by analyzing the distance between the lowest level of any extracted pole and the ground height at this position. Subsequently, the region below the extracted pole is analyzed. In case there are no points inside the inner searching radius the pole is marked as hovering and removed from the dataset. Assuming that a pole-like object is on top of another object, one may find points below the extracted pole. To make sure the extracted object has to be excluded this pole is again analyzed using the algorithm described earlier, however, a modified search radius is employed. While the kernel radius, r_1, remains constant, the radius of the outer cylinder, r_2, is decreased. This method helps to prevent the exclusion of valid poles in cases where objects are in the vicinity, e.g. pedestrians.

3.2 RESULTS

The elimination of false positives was done separately using the different approaches described in the preceding section. There were two areas chosen from the complete dataset, to analyze the reduction of false positives. The first area is located along a 3.5 km long trajectory, which contains a street in a built-up area as well as a highway like road (Area 1). The second region is a densely built-up area along a trajectory with a length of 1.3 km (Area 2, Figure 7).

Figure 7. Classified poles in a densely built-up area. False positives are shown as black dots. Trees are marked in green, street lights in blue, sign posts in light beige, unidentified objects in red.

For both areas, the true number of correctly extracted poles and the number of false positives were obtained by manual inspection. The results for the two test areas are summarized in Table 1.

Table 1. Results of the elimination of false positives for the two test areas (hovering: hovering poles excluded, point (beam) density: areas with low point (beam) density excluded). The error rate gives the number of errors among the extracted poles before excluding false positives.

		Area 1		Area 2	
No. of extracted poles	all poles	299	100%	131	100%
	hovering	200	67%	57	44%
	point density	81	27%	55	42%
	beam density	230	77%	86	66%
Error rate	false positives	22	7%	44	34%

Apparently, the different algorithms mark too many objects as being false positives. When using the algorithm searching for hovering poles, also temporarily occluded poles may be detected as errors. Taking into account the theoretic point density, e.g. poles on the outside of the driven curve radius lie in areas with a low point density. As the point density decreases with increasing driving speed, most poles along the highway are excluded. The best results were obtained by computing the beam density for each extracted pole.

4. DISCUSSION AND OUTLOOK

In this paper, we propose an asymmetric scheme for accurate vehicle positioning. A feature map – in our case, a map of pole-like objects – is extracted from a dense 3D point cloud of the road environment, obtained using an MLS system. The quality of the feature map is then analyzed. The main problem the algorithm has to deal with is the low point density in some regions of the point cloud. Other major problems arise from occlusions, which lead to erroneously detected poles behind facades and hovering poles.

There are a number of extensions left for future work. First, we notice that although the scans are quite dense, single objects are often hit by a low number of points, thus calling for an even denser point cloud acquisition. Second, our analysis could benefit from using the raw data directly, i.e. the original laser strips, instead of the (preprocessed) point cloud. Third, we intend to replace our simple feature extraction method by a more elaborated classification method and to extend our set of features to include, for example, planar areas (facades) and guard rails.

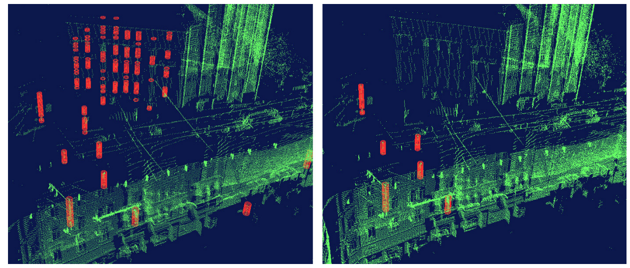

Figure 8. 3D view of a road scenery with extracted poles (red cylinders). Left: Original set of extracted poles with several false positives in the top middle area. Right: Same scene with poles excluded which do not meet the beam density criterion.

5. REFERENCES

[1] Brenner, C. (2009) Extraction of Features from Mobile Laser Scanning Data for Future Driver Assistance Systems, M. Sester et al. (eds.), Advances in GIScience, Lecture Notes in Geoinformation and Cartography, Springer, 25-42.

[2] Brenner, C. (2009), Global Localization of Vehicles using Local Pole Patterns, Proc. DAGM 2009, 31st Annual Symposium of the German Association for Pattern Recognition, Springer LNCS 5748, 61-70.

[3] Burgard, W., Hebert, M. (2008) World Modeling, in: Springer Handbook of Robotics, Springer, 853-869.

[4] EDMap (2004), Enhanced Digital Mapping Project Final Report, Technical report, United States Department of Transportation, Federal Highway Administration and National Highway Traffic and Safety Administration, http://ntl.bts.gov/ lib/ jpodocs/ repts_te/ 14161.htm, 189p. Last accessed June 9, 2009.

[5] IBEO (2009) Ibeo Lux laser scanner, www.ibeo-as.de. Last accessed June 9, 2009.

[6] Kremer, J., Hunter, G. (2007) Performance of the Streetmapper Mobile LIDAR Mapping System in 'Real World' Projects. In: Photogrammetric Week 2007, Wichmann, 215-225.

[7] Kukko, A., Andrei, C.-O., Salminen, V.-M., Kaartinen, H., Chen, Y., Rönnholm, P. Hyyppä, H., Hyyppä, J., Chen, R. Haggrén, H., Kosonen, I., Čapek, K. (2007) Road environment mapping system of the finnish geodetic institute – FGI Roamer, Proc. Laser Scanning 2007 and SilviLaser 2007, IAPRS Vol. 36 Part 3/W52, 241-247.

[8] Rabbani, T., Dijkman, S., van den Heuvel, F., Vosselman, G. (2007) An integrated approach for modeling and global registration of point clouds. ISPRS Journal of Photogrammetry and Remote Sensing 61 (6), 355–370.

[9] Thrun, S., Burgard, W., Fox, D. (2005) Probabilistic Robotics, The MIT Press, Cambridge, Mass.

[10] Weiss, T., Kaempchen, N., Dietmayer, K. (2005) Precise ego localization in urban areas using lasedrscanner and high accuracy feature maps. Proc. 2005 IEEE Intelligent Vehicles Symposium, Las Vegas, USA, 284-289.