

VEHICLE RECOGNITION FROM LIDAR DATA

C. K. Toth^a, A. Barsi^b, T. Lovas^b

^aOSU, Center for Mapping, 1216 Kinnear Road, Columbus, OH 43212-1154, USA – toth@cfm.ohio-state.edu

^bBUTE, Department of Photogrammetry and Geoinformatics, Budapest, Hungary – ([barasi.arpad](mailto:barasi.arpad@fmt.bme.hu), [lovas.tamas](mailto:lovas.tamas@fmt.bme.hu))@fmt.bme.hu

Commission III

KEY WORDS: LiDAR, vehicle extraction, classification

ABSTRACT:

This paper focuses on the potential of using airborne laser scanning technology for transportation applications, especially for identifying moving objects on roads. An adaptive thresholding algorithm is used to segment the LiDAR point cloud, which is followed by a selection process to extract the vehicles. The LiDAR data are capable of measuring the vertical profile of a vehicle, and hence provide a base for distinguishing major vehicle types. Various techniques, such as the use of statistical, neural, and rule-based classifiers, were used to recognize the vehicle classes. The classification is based on features derived from a principal component transformation. Thereafter the extracted vehicles were classified into main categories, such as passenger cars, multi-purpose vehicles, and trucks. The feasibility of the developed method to effectively extract vehicles from LiDAR data has been demonstrated on several datasets. The proposed technique makes LiDAR suitable for new transportation applications, such as collecting data for traffic flow monitoring and management, including data on the vehicle count, traffic density, and velocity.

1. INTRODUCTION

Airborne laser mapping is an emerging technology in the field of remote sensing that is capable of rapidly generating high-density, geo-referenced digital elevation data with an accuracy equivalent to traditional land surveys but significantly faster than traditional airborne surveys (Flood, 1999). Despite the initial high price, these systems have made remarkable market penetration, and recent technical and methodological advancements have further improved the capabilities of this remote sensing technology (Wehr and Lohr, 1999). In addition to the conventional DSM/DEM products, the latest high-performance LiDAR systems can deliver very dense and accurate point clouds and thus provide data for more sophisticated applications. At the APSRS 2003 Convention, Optech introduced the ALTM 30/70, a 70 kHz system and soon after that, LHS announced the 58 kHz version of its system, which provides excellent support for corridor mapping. These developments make LiDAR technology capable of acquiring transportation application-specific information beyond conventional mapping and thus supporting tasks such as extracting moving objects. In this paper we investigate the potential of using airborne laser scanning technology for traffic monitoring and other transportation applications.

Road transportation systems have undergone considerable increases in complexity and at the same time traffic congestion has continued to increase. In particular, surface vehicle ownership and the use of vehicles are growing at rates much higher than the rate at which roads and other infrastructures are being expanded. Transportation authorities are increasingly turning to existing and new technologies to acquire timely spatial information of traffic flow to preserve mobility, improve road safety, and minimize congestion, pollution, and environmental impact (Zhao 1997). Besides the widely used conventional traffic data collection techniques, such as detection loops, roadside beacons, travel probes and driver input, the state-of-the-art remote sensing technologies, such as LiDAR and high-resolution digital cameras can provide traffic flow data over large areas without

ground-based sensors. It is expected that the use of modern airborne sensors supported by state-of-the-art georeferencing and image processing technologies will enable fast, reliable, and accurate data capture for traffic flow information retrieval with high spatial and temporal resolution. In particular, the following data will be supported: vehicle count/type, vehicle velocity and travel time estimation, origin-destination flows, highway densities (passenger car per unit distance per lane) and exit flow monitoring, intersection turning volumes, detection of congested/incident areas in real-time to support traffic redirection decision-making, platoon dispersion/condensation monitoring (which can be effectively accomplished only by remote sensing methods), and incident detection and response (Toth et.al., 2003).

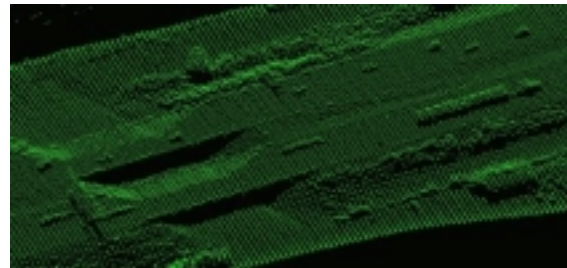


Figure 1. The LiDAR dataset captured over a freeway.

This paper discusses the use of LiDAR data for extracting moving vehicles over the transportation corridors and grouping them into broad classes. The method includes a filtering process of identifying vehicles, the selection of a parameterization to describe the LiDAR point cloud of vehicles, the optimization of the parameter representation, and the classification process. Using three datasets obtained from typical LiDAR surveys, three classification techniques have been tested to assess the performance of the vehicle grouping. Figure 1 shows a typical road segment with various vehicles clearly identifiable from the LiDAR point cloud.

2. VEHICLE EXTRACTION, MODELING AND REPRESENTATION

To support the vehicle recognition process, vehicles must be extracted from the LiDAR point cloud and their geometry should be adequately modeled to provide a good parameter space for the classifier. In order to extract the vehicles from LiDAR data a simple thresholding method can be applied. Since normally the road surface is flat, one threshold value may easily separate the LiDAR points reflected from a vehicle from the points reflected from the road surface. However, for longer vehicles, such as an 18-wheeler, assumption of the road flatness and horizontality may not hold, especially for higher-grade levels (steep roads). Therefore, an adaptive thresholding technique should be used, which adjusts the threshold level as the surface around the vehicle is changing. The applied method, similar to the techniques used in image segmentation is based on the dataset histograms, here created by using the height values instead of the intensities (Pitas, 2000). The problem of the steep slopes can be bypassed using windows moving over the dataset, hence taking limited amount of data into consideration at any thresholding step. Figure 2 shows the processing steps of vehicle extraction from the LiDAR point cloud.

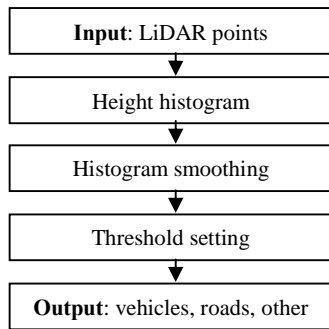


Figure 2. Design architecture and data processing flow.

In order to distinguish major vehicle types, characteristic parameters have to be chosen; here we used a six-parameter representation that includes the size of the vehicle footprint and then four vertical parameters (average height values computed over the four equally sized regions) as shown in Figure 3.

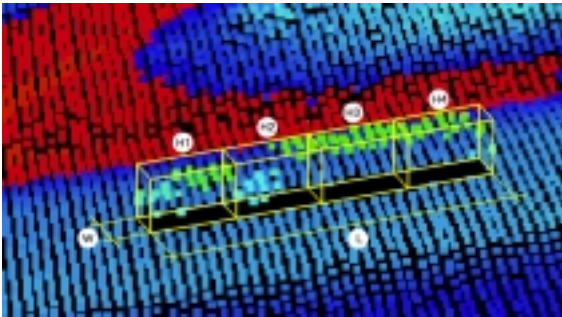


Figure 3. Parameterization of LiDAR points representing a vehicle.

To support the vehicle classification study in using LIDAR data for traffic flow extraction, Woolpert LLP from Dayton,

OH provided a dataset, obtained from flights done for regular mapping purposes. The point density was 1.5 point/m², which was certainly adequate for topographic mapping and could be considered at best minimal for vehicle identification. The LiDAR data covered a freeway section of State Route 35 (East of Dayton), packed with vehicles, and was used as a training dataset for developing the classifiers. 72 vehicles were chosen and processed in an interactive way, the regions containing vehicles were selected by an operator and the vehicles were automatically extracted by the thresholding method presented earlier. All the vehicles were parameterized and then categorized into three main groups: passenger cars, MPVs (multi-purpose vehicles such as SUVs, minivans, light trucks), and trucks/eighteen-wheelers.

An important aspect of the input data selected for testing is the relative velocity between the airborne data acquisition platform and the vehicles to be observed. The aircraft speed for the Dayton survey was known from the GPS/INS navigation solution and the average speed of the LiDAR sensor was about 200 km/h during the survey. This roughly translates into the relative velocity range of 100-300 km/h between the data acquisition platform and the moving objects observed. Figure 1 clearly shows the impact of the relative speed as the vehicles traveling at faster relative speed (opposite direction) have smaller footprints while the smaller relative velocity (airplane and vehicles are moving in the same direction) results in elongated vehicle footprints. The extreme of zero relative velocity, such as the vehicle moving with same ground speed as the aircraft, the LiDAR-sensed vehicle size would be infinite; the vehicle would become practically not detectable. In this paper, the estimation of the vehicle velocity is not considered, some aspects of this process are discussed in (Toth et.al., 2003).

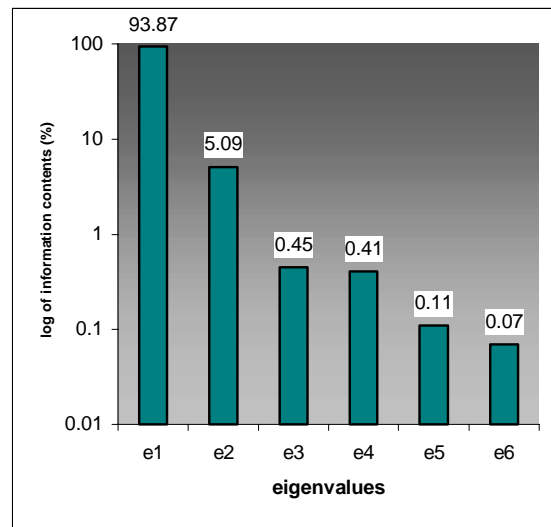


Figure 4: The eigenvalues and information contents of the training data set, which consisted of 72 vehicles.

To reduce the dimensionality of the parameter space, Principal Component Analysis (PCA) was then performed. PCA is an effective tool for handling data representation/classification problems where there is a significant correlation among the parameters describing the object patterns. By training the datasets, the correlation can

be determined and a reduced parameter set can be defined that can both represent the information in a more compact way and can support an efficient classification in the reduced feature space. The clear advantage of the method is that it does not require any physical modeling of the data; of course, the selection of the input parameters has some importance. Provided that a rich set of input parameters is defined, however, the method will effectively identify the redundancy and thus usually results in a quite reduced parameter representation. In our investigations the 72 vehicles provided a statistically meaningful dataset for the PCA process. The eigenvalues computed from the covariance matrix and ordered monotonically are shown in Figure 4.

In analyzing the results, it is quite striking to see that more than 98% percent of the original information content is preserved if only the two largest eigenvalue components are used for data representation. To assess the classification performance, for which high information contents do not necessarily give guarantees, the 72 vehicles converted into the two-dimensional feature space as plotted in Figure 5. Cars are marked with \circ , MPVs with $+$, and trucks with $*$, respectively; vehicle direction with respect to sensor motion is coded in red and blue. Figure 6 shows the results if only the height parameters (4) were used as input in the PCA.

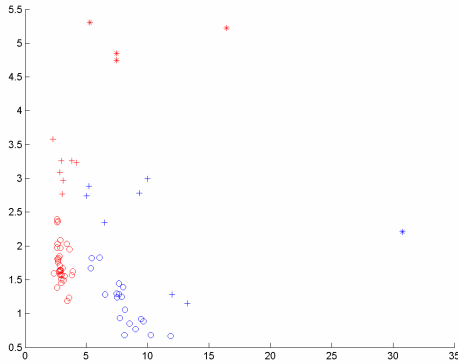


Figure 5. Vehicle distribution in the two-dimensional feature space (6-parameter input data-based PCA).

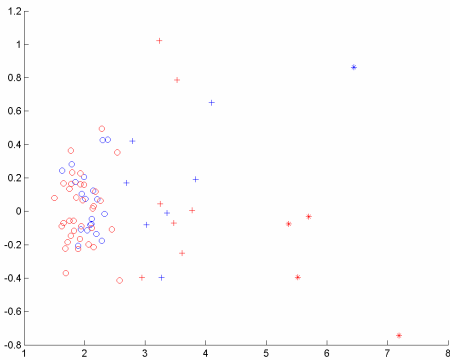


Figure 6. Vehicle distribution in the two-dimensional feature space if only height parameters (4) were used in the PCA process.

Comparing Figures 5 and 6, it is apparent that vehicle categories can be effectively separated using solely vehicle height parameters. In fact, the vertical profile of the vehicles itself seems to be sufficient for the vehicle group classification. Obviously, not using the length information means that the vehicle travel directions become indistinguishable. Why the width has no significant impact is probably explained by two facts. First the variations between the three vehicle groups are rather small – the difference between the mean vehicle widths is about 0.5 m. Second, the footprint of the LiDAR, the area that one pulse will reach is about 25 cm (diameter of the ellipse). Given the spacing between the LiDAR pulses, which is at least 0.5 m, it is apparent that the measuring accuracy of the vehicle width is rather poor and consequently the information content of this parameter is rather insignificant. The vehicle travel direction, however, can be recovered from the 6-parameter model.

3. VEHICLE RECOGNITION

For vehicle classification, three methods have been considered. The main goal here is to classify the vehicles into the given categories: passenger cars (P), multi-purpose vehicles (MPV) and trucks (T). Each category has two subclasses (along and against) considering the traffic direction relative to the flight direction. Therefore, the recognition process is expected to separate the vehicles into six groups, identified as follows:

<i>ID</i>	<i>Category</i>
1	P along
2	P against
3	MPV along
4	MPV against
5	T along
6	T against

The recognition process, including the derivation of the classifier's parameters, was performed by using the Ohio data set (72 vehicles) for all three different classification methods.

Rule-based classifier

The first method was a rule-based classifier, which contains decision rules derived from the PCA transformed features. As depicted in Figure 7, a clear separation, in other words, clustering of samples with identical labels can be easily made between the groups by using straight lines. These lines are of course specified by two variables, which are determined by simple calculations.

For example, Category 1 (passenger cars traveling along the flight direction) is bounded by *Line A* and *C*, furthermore by the coordinate axis *x*. *Line A* can be defined by (1):

$$y = a_A x + b_A = \frac{0.5 - 3}{15} x + 3 \quad (1)$$

where *x* and *y* are the two first principal components.

Similarly *Line C* is defined by (2):

$$x = 4.5 \quad (2)$$

The rule for the category is thereafter:

$$(y < -\frac{2.5}{15}x + 3) \text{ AND } (x > 4.5) \text{ AND } (y > 0) \quad (1)$$

Category 3 (MPV traveling along the flight direction) represents a more complex cluster boundary, which can be described as:

$$(y > a_A x + b_A) \text{ AND } (y < a_B x + b_B) \text{ AND} \quad (3)$$

$$(x > x_C) \text{ AND } (y > c_D) \text{ AND } (y > 0)$$

where the indices show which parameters correspond to which lines.

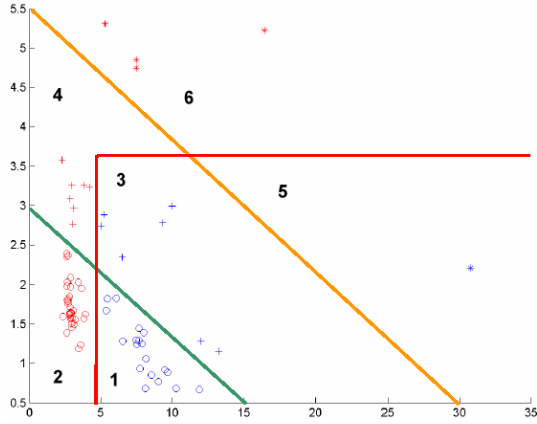


Figure 7. Segmentation of the two-dimensional feature space of the training vehicles

The determination of all parameters and subsequent creating of all the rules is a rather straightforward task. However, the introduction of new observations (new features) usually requires the refinement of the rules. Applying the rules to an unknown feature vector is obviously simple and fast.

Minimum-distance method

The second investigated classifier was a fundamental statistical technique: the minimum distance method. This classifier is based on a class description involving the class centers, which are calculated by averaging feature components of each class. An unknown pattern is classified by computing the distances between the pattern and all class centers and the smallest distance determines in which class the pattern will be classified. The distance calculation based on the Euclidean measure in our two-dimensional case is (Duda, 2001):

$$D_j = \sqrt{(x - \bar{x}_j)^2 + (y - \bar{y}_j)^2} \quad (5)$$

where the class center of class j is given by \bar{x}_j and \bar{y}_j . The classification is based on the evaluation of (6):

$$C = \arg \min_j (D_j) \quad j = 1, 2, \dots, 6 \quad (6)$$

This method is simple and the algorithm executes fast. As new vehicles are added to the training set, the class centers have to be recalculated but the decision formula remains unchanged. Class centers and boundaries, which form a Voronoi tessellation, are shown in Fig. 8.

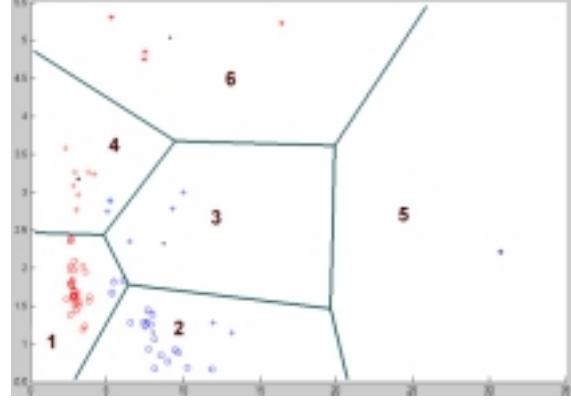


Figure 8. Segmentation of the two-dimensional feature space of training vehicles by the minimum-distance method (Voronoi tessellation)

Artificial neural network classifier

The third method in the vehicle recognition investigation was based on an artificial neural network classifier. The feed-forward (back-propagation) neural networks have to be trained by the features. As it is commonly agreed (Brause, 1995; Rojas, 1993), most practical works require 3-layer networks; hence such a structure was implemented in our tests. In order to get the simplest network, the following strategy was applied: a network with a small number of neurons was created and then trained. If the network's recognition accuracy reached the required value, the design phase was stopped, otherwise a neuron was added. The additional neurons were given firstly to the second layer, then to the first one. This successive method ensures the minimal balanced structure of an acceptable network. All the designed networks had a logistic sigmoid transfer function on the first and second layers, and a linear transfer function on the third (output) layer. This structure is capable of directly producing the required class identifiers. The training method was the Levenberg-Marquard algorithm (Demuth, 1998), the maximal number of training steps (epochs) was 70, and the required error goal value was 0.1. The network error was calculated by the mean square error (MSE) method. At the end, the output of the neural network was rounded to the nearest integer.

In our experiments, the above strategy was applied with an initial network structure of 2-2-1. The network addition was stopped at 6-6-1, at which point eight networks were found which had smaller recognition error than 10 (13.8 %). From that network set the simplest was chosen, which had a structure of 3-4-1. The introduction of additional vehicles in neural networks means the repetition of the entire training procedure. The use of the trained network (network simulation), however, is relatively fast.

4. RESULTS

The three developed vehicle recognition techniques were tested on the training data set of Ohio (1), on the data set containing vehicles from Ohio and Michigan, (2) and on combined dataset, including the Ontario data (3). The first test (in-sample test) was only an internal check of the algorithms. Table 1 shows a performance comparison of the three techniques.

<i>Data set (total number of vehicles)</i>	<i>Rule-based</i>	<i>Minimum distance</i>	<i>Neural network</i>
Ohio (72 vehicles)	0 (0%)	8 (11.1%)	2 (2.8%)
Ohio + Michigan (87)	2 (2.3%)	12 (13.8%)	8 (9.2%)
Ohio + Michigan + Ontario (102)	2 (2%)	17 (16.7%)	16 (15.7%)

Table 1. The comparison of the three recognition techniques; number of misclassification errors (percentage)

<i>Data set (total number of vehicles)</i>	<i>Rule-based</i>	<i>Minimum distance</i>	<i>Neural network</i>
Ohio (72 vehicles)	0 (0%)	4 (5.6%)	2 (2.8%)
Ohio + Michigan (87)	2 (2.3%)	8 (9.2%)	8 (9.2%)
Ohio + Michigan + Ontario (102)	2 (2.3%)	10 (9.8%)	14 (13.7%)

Table 2. The misclassification errors of the methods without considering the vehicle travel directions

The rule-based method has perfectly identified the features, while the other two methods have small recognition errors. In all methods, the most frequent misclassification error type was the mismatch of the Ps and the MPVs in the along direction, since passenger cars can have shape and length very similar to MPVs. Ignoring the relative traveling direction, in other words classifying into three classes instead of six, the results are somewhat different as shown in Table 2.

The tests with the combined Ohio, Michigan and Ontario data show strong out-of-sample performance, which is a good indication of the applicability of the proposed vehicle recognition method. Obviously, more tests with a variety of data are needed to confirm the ultimate potential of using LiDAR data as a source for traffic flow estimates.

5. CONCLUSIONS

Considering the fact that all the three classification methods used have produced rather good results, it is fair to say that LiDAR data can be used to support traffic flow applications. All three methods were able to recognize the vehicle categories with accuracy better than 80 %. This high recognition rate proves that a classifier designed and parameterized by an adequate training dataset can be successfully applied on other, unknown data sets. Furthermore, the results are even more encouraging if the relatively modest LiDAR point density is factored in (1.5 point/m²). The state-of-the-art LiDAR systems can easily provide a 3-5 times denser point cloud and consequently better classification performance can be expected.

The developed method has demonstrated that LiDAR data contain valuable information to support vehicle extraction, including vehicle grouping and localizations. The classification performance showed strong evidence that the major vehicle categories can be efficiently separated. With the anticipated improvements in LiDAR technology, such as denser point cloud and smaller pulse footprint, the classification efficiency is expected to grow further. The price of LiDAR, however, is prohibitive at this point to support real-life applications. Nevertheless, collecting data over transportation corridors during regular surveys already offers a no-cost opportunity to obtain important traffic data. In addition, the advantage of the moving platform is that it can be freely deployed more or less any time and anywhere.

6. ACKNOWLEDGEMENTS

This research was partially supported by the NCRST-F program. The authors would like to thank Woolpert LLC and Optech International for providing the LiDAR datasets.

7. REFERENCES

- Baltsavias, E. P. – Gruen, A. – L. V. Gool (Eds.) (2001): Automatic Extraction of Man-Made Objects from Aerial and Space Images (III), Alkema Publishers, Lisse
- Bässmann, H. – Kreys, J. (1998): Bildverarbeitung AdOculos, Springer Verlag, Berlin
- Brause, R. (1995): Neuronale Netze, B. G. Teubner, Stuttgart
- Demuth, H. – Beale, M. (1998): Neural Network Toolbox, Matlab User's Guide, The MathWorks, Natick
- Duda, R. O. – Hart, P. E. – Stork, D. G. (2001): Pattern Classification, Wiley, New York
- Flood, M. (1999): www.airbornelasermapping.com
- Jähne, B. – Haußecker, H. – Geißler, P. (Eds.) (1999): Handbook of Computer Vision and Applications I-II-III. CD-ROM Set, Academic Press
- Lillesand, T. M. – Kiefer, R. W. (1994): Remote Sensing and Image Interpretation, Wiley, New York
- Pitas, I. (2000): Digital Image Algorithms and Applications, John Wiley & Sons, Inc.
- Rojas, R. (1993): Theorie der neuronalen Netze – Eine systematische Einführung, Springer Verlag, Berlin
- Rumelhart, D. E. – McClelland, J. L. (Eds.) (1988): Parallel Distributed Processing, Explorations in the Microstructure of Cognition, Vol. 1: Foundations, MIT Press, Cambridge
- Russ, J. C. (1995): The Image Processing Handbook, CRC Press, Boca Raton
- Toth C. – Grejner-Brzezinska D. and Lovas T.: Traffic Flow Estimates from LiDAR Data, Proc. ASPRS Annual Conference, May 5-9, 2003, pp. 203-212, CD ROM.
- Wehr A. – Lohr U. (1999): Airborne Laser Scanning – and Introduction and Overview, *ISPRS Journal of Photogrammetry and Remote Sensing*, 54, pp.68-82.
- Zhao, Y. (1997): Vehicle Location and Navigation Systems, Artech House, Inc., Boston