

LIDAR Based Vehicle Classification

Benjamin Coifman, PhD
Associate Professor
The Ohio State University
Joint appointment with the Department of Civil and Environmental Engineering and Geodetic Science,
and the Department of Electrical and Computer Engineering
Hitchcock Hall 470
2070 Neil Ave, Columbus, OH 43210
Phone: (614) 292-4282
E-mail: Coifman.1@OSU.edu

Ho Lee, PhD Candidate
Graduate Research Associate
Department of Civil and Environmental Engineering and Geodetic Science
The Ohio State University
Columbus, OH 43210
E-mail: lee.2406@osu.edu



DISCLAIMER

The contents of this report reflect the views of the authors, who are responsible for the facts and the accuracy of the information presented herein. This document is disseminated under the sponsorship of the Department of Transportation University Transportation Centers Program, in the interest of information exchange. The U.S. Government assumes no liability for the contents or use thereof.

ACKNOWLEDGEMENTS

This material is based upon work supported in part by NEXTRANS the USDOT Region V Regional University Transportation Center, The Ohio Transportation Consortium University Transportation Center, and the Ohio Department of Transportation. We are particularly grateful for the assistance and input from Dave Gardner and Lindsey Pflum at the Ohio Department of Transportation.

ABSTRACT

Vehicle classification data are used for numerous transportation applications. Most of the classification data come from permanent in-pavement sensors or temporary sensors mounted on the pavement. Moving out of the right-of-way, this study develops a LIDAR (light detection and ranging) based classification system with the sensors mounted in a side-fire configuration next to the road. The first step is to distinguish between vehicle returns and non-vehicle returns, and then cluster the vehicle returns into individual vehicles. The algorithm examines each vehicle cluster to check if there is any evidence of partial occlusion from another vehicle. Several measurements are taken from each non-occluded cluster to classify the vehicle into one of six classes: motorcycle, passenger vehicle, passenger vehicle pulling a trailer, single-unit truck, single-unit truck pulling a trailer, and multi-unit truck. The algorithm was evaluated at six different locations under various traffic conditions. Compared to concurrent video ground truth data for over 27,000 vehicles on a per-vehicle basis, 11% of the vehicles are suspected of being partially occluded. The algorithm correctly classified over 99.5% of the remaining, non-occluded vehicles. This research also uncovered emerging challenges that likely apply to most classification systems: differentiating commuter cars from motorcycles.

Occlusions are inevitable in this proof of concept study since the LIDAR sensors were mounted roughly 6 ft above the road, well below the tops of many vehicles. Ultimately we envision using a combination of a higher vantage point (in future work), and shape information (begun herein) to greatly reduce the impacts of occlusions.

INTRODUCTION

Vehicle classification data are used in many transportation applications, including: pavement design, environmental impact studies, traffic control, and traffic safety [1]. There are several classification methods, including: axle based (e.g., pneumatic tube and piezoelectric detectors), vehicle length based (e.g., dual loop and some wayside microwave detectors), as well as emerging machine vision based detection. As noted by the Traffic Monitoring Guide [1], each sensor technology has its own strengths and weaknesses regarding costs, accuracy, performance, and ease of use.

In the present study we add another technology to the mix and develop a vehicle classification algorithm for LIDAR (Light detection and ranging) sensors mounted in a side-fire configuration. Our prototype system consists of two LIDAR sensors mounted on the driver's side of a probe vehicle parked alongside the roadway. Each LIDAR scans a vertical plane across the roadway, providing a rich view of the passing vehicles. In practice, the LIDAR sensors could be mounted on a temporary deployment platform like this system, or permanently mounted on a pole adjacent to the roadway.

To classify vehicles, first we segment them from the background, next we look for possible occlusions using algorithms developed herein, and then we measure several features of size and shape for each vehicle. These features are subsequently used for classification into six categories. The classification algorithm is evaluated by comparing the individual vehicle results against concurrent video. Occlusions are inevitable in this proof of concept study since the LIDAR sensors were mounted roughly 6 ft above the road, well below the tops of many vehicles. The present work focuses primarily on the non-occluded vehicles. Ultimately we envision using a combination of a higher vantage point in future work (similar to wayside microwave detectors), and shape information (begun herein) to greatly reduce the impacts of occlusions.

LIDAR technology has been applied in various transportation applications, such as highway safety [2,3] and highway design [4,5]. There have been a few demonstrations of LIDAR or related optical range finding technologies to monitor traffic and sometimes classify the vehicles. The most notable example being the Schwartz Autosense [6], which consisted of a sensor mounted over the lane of travel; though this basic approach pre-dates the Autosense system [7]. While the overhead view eliminates occlusions, the need to mount the sensor over the roadway makes deployment more difficult. Others have contemplated using airborne LIDAR platforms for traffic monitoring [8,9]. For example, [9] collected LIDAR imagery data over transportation corridors, segmented individual vehicle from the road surface, and then extracted six parameters of vehicle shape and size for each vehicle. They classified vehicles in three categories (passenger vehicles, multi-purpose vehicles, and trucks) using principle component analysis. Finally, our group has also contemplated the use of LIDAR to classify vehicles from a moving platform [10,11].

The remainder of this paper is organized as follows. First the process of collecting the LIDAR data and the procedure of segmenting the vehicles from the background are presented. Next, the LIDAR based vehicle classification algorithm is developed. Third, the algorithm is evaluated on a per-vehicle-basis against concurrent video ground truth from field data at six directional locations, exhibiting various traffic conditions, distance between LIDAR and target vehicles, and road type (freeway and arterial road). The evaluation dataset includes over 25,000 vehicles (23,000 non-occluded). Then, the paper closes with conclusions.

LIDAR MEASUREMENTS AND VEHICLE DETECTION

Figure 1(b) shows an overhead schematic of the prototype deployment. The two LIDAR sensors are each mounted at a height of about 6.7 ft above ground and they are 4.6 ft apart from one another. Each LIDAR sensor scans a vertical plane across the roadway at roughly 37 Hz. Each scan sweeps 180 degrees, returning the distance to the nearest object (if any) at 0.5 degree increments with a ranging resolution of

0.1 inch and a maximum range of 262 ft. So each scan returns 361 samples in polar coordinates (range and angle) relative to the LIDAR sensor and these data are transformed into a Cartesian coordinate system (lateral distance and relative height) for analysis.

Using these LIDAR data, vehicle segmentation is split into two steps. First we distinguish between vehicle returns and non-vehicle returns (e.g., pavement, foliage, barriers, etc.). Then we cluster the vehicle returns into discrete vehicles.

To segment the vehicle and non-vehicle returns we adapt background subtraction techniques from conventional image processing. The LIDAR are fixed, so when no vehicles are present they will return nearly identical scans of the *background*. Thus, the background's range at a given angle is (roughly) constant and over time the background returns are the dominant range reported at each angle. Formalizing this concept to extract the background from the LIDAR data, we set the background equal to the median range at each angle as observed over an extended time period with largely free flowing traffic. Whenever a vehicle is present, the vehicle's returns can only be at a range that is closer than the background range for the given angle. So in the absence of free flowing traffic, one could instead take the distribution of observed ranges at a given angle and set the background equal to the furthest mode of the distribution.

Figure 2(a) shows the data collection on I-270 southbound, on the west side of Columbus, Ohio. The probe vehicle was parked just off of the right hand shoulder to collect LIDAR data. Figure 2(b) shows the corresponding background that was extracted from the LIDAR data. Because the van would occasionally roll a small amount about its central axis as personnel entered or exited the vehicle during the data collection, returns from the background did not always fall on the measured background curve. All returns falling beyond the background curve as well as any returns that were within an inch above the background curve were considered non-vehicle returns and excluded from further analysis. However, if the low-lying returns prove critical to a subsequent application, one could estimate the LIDAR's instantaneous angle relative to the shoulders (0 - 11.8 ft, and 47.2 - 65.6 ft in Figure 2(b)) and normalize this angle across scans.

Only vehicle returns should remain after removing the background, however, these returns still need to be clustered into individual vehicles and we take the following steps to do so. First we establish the lane boundaries by looking at the distribution of the lateral distance across the vehicle returns. We expect to see one distinct mode per travel lane, corresponding to the near side of the vehicles when traveling in the given lane since the vertical edges on the vehicles will generally yield many returns at the same lateral distance; though, there will be other returns in the distribution from horizontal vehicle surfaces, vehicles changing lanes, and so forth. Provided the LIDAR sensors are not moved, this step only needs to be done once, using a few minutes of data.

Second, in each scan we segment the LIDAR returns by lane using the lane boundaries from the previous step. As long as a vehicle travels within a lane, all of the returns from that vehicle will fall between the respective lane boundaries in the given scan. In most cases even a single return in the lane will be taken as that lane being occupied in that scan. However, in the relatively rare cases when a vehicle changes lanes as it passes the LIDAR, that vehicle's returns may fall into two adjacent lanes (we saw this event occur 253 times out of 27,450 vehicles). To find the cases when a single vehicle is seen in adjacent lanes, we explicitly look for concurrent returns in neighboring lanes. When this occurs, we take the mode of lateral distance in the near lane and the far lane, respectively. Again, the nearside of a vehicle is characterized by a large number of returns at a given lateral distance, i.e., the mode lateral distance within the lane. If in the given scan the difference between the modes in successive lanes is less than the maximum feasible vehicle width (set to 8.5 ft, the maximum width of commercial motor vehicles [12]), the vehicle returns in the adjacent lanes are assumed to come from a single vehicle and are grouped together in the lane corresponding to the median lateral distance among the set of returns in question. Otherwise, the two modes are too far apart to come from a single vehicle and the groups are kept separate. Obviously this approach assumes that at most one vehicle can occupy a lane in a given scan; although we know that it is not always the case, e.g., when two motorcycles pass side by side within a lane, we have yet to observe any such exceptions in the LIDAR data so addressing these exceptions is left to future research.

Third, taking the temporal sequence by lane, the returns are clustered into vehicles. After each scan is processed, whenever a given lane is occupied, if there is not already an open vehicle cluster in that lane then a new vehicle cluster is begun with the corresponding returns; otherwise, the corresponding returns are added to the open vehicle cluster in that lane. On the other hand, if there is an open vehicle cluster and the lane has not been occupied for at least 1/4 sec (roughly 9 scans) then the open vehicle cluster is closed. To be retained, a closed vehicle cluster must span at least two scans and at least two of the scans must have different heights, otherwise, the vehicle cluster is discarded. Because the returns in a scan are grouped by lane before the clustering step and we make the above correction for vehicles changing lanes, it is theoretically possible for two neighboring vehicles to be erroneously clustered together. Though we have not seen this problem occur, to safeguard against it, if the net width of a closed cluster is greater than the maximum feasible vehicle width then the cluster is split in two, by lane. On the other hand, it is possible for a vehicle changing lanes to be assigned to different lanes at different time steps, resulting in separate clusters in each lane. To catch these breakups, when a cluster ends in one lane, we check the next scan to see if a new cluster begins in an adjacent lane a small distance away, i.e., if the difference between the mode lateral distance is less than 3.5 ft, the two clusters are merged together and assigned to the lane with the larger cluster. The segmentation and clustering steps are repeated for each lane across each successive LIDAR scan until all of the vehicle returns have been clustered into discrete vehicles.

Occlusion Reasoning

A key step in classifying a given vehicle is determining whether the entire vehicle was seen or if there was evidence of a partial occlusion. Table 1 shows that the latter case arose for about 12% of the vehicles observed on the multilane facilities. The frequency is so small because the spacing between vehicles is typically much larger than one might think, e.g., according to the HCM [13], LOS F on a freeway begins at 46 passenger cars per mile per lane or 117 ft per passenger car and passenger cars are generally on the order of 10-20 ft long. In any event, partially occluded vehicles are likely to be misclassified in our algorithm if the occlusion is not identified and handled separately from the non-occluded vehicles. Of course from the LIDAR data stream we cannot detect completely occluded vehicles, though we found these errors occurred between 3-6% in the three multilane datasets that had an independent detector to monitor occluded lanes (I-71 and I-270 in Table 1), and as one might expect, most of these occluded vehicles were PV. A higher vantage point or monitoring from the median of the roadway should reduce the frequency of completely occluded vehicles.

For any given vehicle cluster we suspect a partial occlusion occurred unless we see at least one non-vehicle return on all sides of the cluster (both temporally and spatially). To automatically detect partially occluded vehicles, first we check the vehicles seen in each scan of the LIDAR. If we cannot see the background curve between a given pair of vehicles the further vehicle is suspected of being partially occluded by the closer vehicle. Second, we check successive scans, if one vehicle is seen at a given angle in scan i , and a different vehicle is seen at the same angle in scan $i+1$, whichever vehicle cluster is further away is considered to be occluded.

LIDAR BASED VEHICLE CLASSIFICATION ALGORITHM

In this section we develop an algorithm to classify the vehicle clusters extracted from the LIDAR data in the previous section. The core algorithm focuses on the non-occluded vehicles and sorts them into six vehicle classes: motorcycle (MC), passenger vehicle (PV), PV pulling a trailer (PVPT), single-unit truck/bus (SUT), SUT pulling a trailer (SUTPT), and multi-unit truck (MUT). These classes are a refinement of commonly used length based classes (as noted in [1], a user might not need the full 13 axle based classes and three or four simple categories may suffice). After classifying the non-occluded vehicles we separately handle the partially occluded vehicles, taking care to address the uncertainty about what went unobserved.

We derived the vehicle classification algorithm using a ground truth development dataset that consists of 24 min of free flow data collected across four lanes on I-71 southbound in Columbus, Ohio, between 11th Ave and 17th Ave on July 9, 2009. There were 1,502 non-occluded vehicles in this dataset and all of the vehicle classifications were manually verified from the video ground truth data. The two primary vehicle features used by the classification algorithm are length and height measured from the individual vehicle clusters, as shown in Figure 3(a). Compared to using length alone, as would be done from loop detectors (see, e.g., [14]), vehicle height helps separate different vehicle classes (e.g., SUT and PVPT). However, the boundaries of various classes still overlap in the length-height plane, so we calculate six additional measurements of the vehicle's shape (for a total of eight shape measurements), as enumerated below and explained in the following subsections.

- Vehicle length (VL)
- Vehicle height (VH)
- Detection of middle drop (DMD)
- Vehicle height at middle drop (VHMD)
- Front vehicle height (FVH)
- Front vehicle length (FVL)
- Rear vehicle height (RVH)
- Rear vehicle length (RVL)

Vehicle Length (VL) and Vehicle Height (VH)

The two side LIDAR sensors are mounted in a “speed-trap” configuration with 4.6 ft spacing. Any moving target will appear at different times in the two views, thereby allowing for speed measurement. Figure 1(a) shows a hypothetical example of the time-space diagram as a vehicle passes by the two LIDAR sensors and Figure 1(b) shows the corresponding schematic on the same distance scale. In this study a vehicle passes the rear LIDAR sensor first and then the front LIDAR sensor. Figure 1(c) and (d) show the vehicle returns from each of the two LIDAR sensors as the vehicle passes, where FT and LT respectively denote the first and last time samples in which the vehicle was scanned by the given LIDAR (subscript “r” for rear and “f” for front). OnT_r and OnT_f indicate the duration of time that a vehicle is scanned by the given LIDAR sensor, i.e., the on-time, where $OnT_x = LT_x - FT_x$, and x is either “r” or “f”. Meanwhile, the traversal time is defined as the difference between the first scan time at the two sensors, i.e., $TT_{FT} = FT_f - FT_r$, or the last scan time, i.e., $TT_{LT} = LT_f - LT_r$. Speed is calculated via Equation (1) from the LIDAR spacing, D, and the traversal time. Vehicle length (VL) is calculated from the mean of V_{FT} and V_{LT} , multiplied by OnT_r (we arbitrarily select the rear LIDAR on-time in this study), yielding Equation (2). Finally, vehicle height (VH) is directly measured from the difference of the highest relative height and the lowest relative height across all of the returns in the given vehicle cluster from the rear LIDAR, yielding Equation (3). By using the difference in cluster heights, this step accounts for the fact that the road cross-section is not flat, each lane may be at a different height relative to the LIDAR sensor.

$$V_{FT} = \frac{D}{TT_{FT}}, V_{LT} = \frac{D}{TT_{LT}} \quad (1)$$

$$VL = \text{mean}(V_{FT}, V_{LT}) \times OnT_r \quad (2)$$

$$VH = \max(h(t)) - \min(h(t)), \forall t \in [FT_r, LT_r], \forall h(t) \in \text{cluster} \quad (3)$$

where $h(t)$ is height of a LIDAR return relative to a height of LIDAR sensor at time t .

Figure 3(a) shows a scatter plot of vehicle height versus vehicle length for the 1,502 non-occluded vehicles from the development dataset sorted by the six vehicle classes. The VH for almost all of the MC, PV and PVPT are below 8 ft, while VH for almost all of the SUT, SUTPT, and MUT are above 8ft. As will be discussed shortly, the height of the trailer (or its load) is sometimes the tallest point on a PVPT or SUTPT and thus is reflected in VH for that vehicle. The observed VL are distributed between 5 ft and 89 ft, with a clear but overlapping progression from MC to PV to PVPT, and similarly from SUT to SUTPT to MUT. Based on this plot, we select $VL = 7.5$ ft as the dividing line between MC and PV. To segregate the remaining classes, we look for a characteristic "gap" before the start of a trailer (PVPT, SUTPT, and MUT) as follows.

Detection of a Middle Drop in a Vehicle (DMD)

The vertically scanning LIDAR captures the profile shape of the passing vehicles. This profile is useful to distinguish between vehicle classes with overlapping VL and VH ranges, e.g., SUT and MUT. For vehicles in these ranges, we look for the presence of a gap that is indicative of the start of a trailer, as manifest as one or more scans with a "drop" in the number of returns somewhere in the middle of the vehicle cluster. To determine whether a vehicle has such a *middle drop*, we first tally the number of LIDAR returns, nLR , as a function of each scan (i.e., time) that the vehicle cluster was seen, yielding $nLR(t)$. For example, Figure 4(a) shows the image of a pickup truck pulling a trailer (an example of PVPT) as it passes by the LIDAR sensors while Figure 4(b) shows the corresponding LIDAR returns from the vehicle cluster. Figure 4(c) shows the $nLR(t)$ curve for the vehicle cluster. The curve does a good job highlighting the point where the trailer is connected to the pickup truck via the low $nLR(t)$. Note that we deliberately use $nLR(t)$ rather than the height of the vehicle because there are some trailers that have a return near the top of the gap even though most of the gap is open (e.g., tree trimming trucks).

Formalizing the process, once the $nLR(t)$ curve is obtained, the set of local minimum points on the curve are considered as potential locations of a middle drop in the vehicle's shape, where $nLR(t_i^*)$ denotes the i -th minima. Since the middle drop should correspond to relatively few LIDAR returns in the given scan (but not necessarily zero due to the connecting link, e.g., the hitch in Figure 4(a)), we assume that nLR at a middle drop must be less than the average of $nLR(t)$ for the cluster across all times, \overline{nLR} . So, we ignore i -th local minimum if it is greater than \overline{nLR} . Formalizing this process, a given scan is considered a possible middle drop if it satisfies all of the conditions in Equation (4).

$$\begin{aligned} nLR(t_i^*) &< nLR(t_i^* - 1) \\ nLR(t_i^*) &< nLR(t_i^* + 1) \\ nLR(t_i^*) &< \overline{nLR} \end{aligned} \quad (4)$$

For each minima at t_i^* , we take the difference of $nLR(t)$ and $nLR(t_i^*)$ over all times, $t \in (FT_r, LT_r)$, denoted $\Delta n(t, t_i^*)$. We find $\max(\Delta n(t, t_i^*))$ over the α ft ahead of the scan at t_i^* ($\alpha = 4$ feet in this study), add it to $\max(\Delta n(t, t_i^*))$ for α ft behind the scan and divide the sum by $nLR(t_i^*)$, yielding $SRD(t_i^*)$ via Equation (5) at each t_i^* . The use of distance rather than time is to make the algorithm robust to slow moving vehicles. Next we select the $\max SRD(t_i^*)$ and call this value the *Detection of Middle Drop* (DMD) indicator, as expressed via Equation (6), and set t^* equal to the corresponding t_i^* . Figure 3(b) shows the cumulative distribution function of DMD for the 1,502 non-occluded vehicles by vehicle class in the development dataset. As expected PVPT, SUTPT and MUT have a wider range of DMD than MC, PV, and SUT. The latter three classes usually present zero DMD, indicative of a vehicle without a middle drop.

$$\begin{aligned} \text{SRD}(t_i^*) &= \frac{\Delta n(t_a, t_i^*) + \Delta n(t_b, t_i^*)}{n\text{LR}(t_i^*)} \\ &= \frac{n\text{LR}(t_a) + n\text{LR}(t_b) - 2 \times n\text{LR}(t_i^*)}{n\text{LR}(t_i^*)} \end{aligned} \quad (5)$$

where,

$$\begin{aligned} t_a &= \max\left(t_i^* - \frac{\alpha}{\text{mean}(V_{\text{FT}}, V_{\text{LT}})}, \text{FT}_r\right) \\ t_b &= \min\left(t_i^* + \frac{\alpha}{\text{mean}(V_{\text{FT}}, V_{\text{LT}})}, \text{LT}_r\right) \end{aligned}$$

$$\text{DMD} = \max(\text{SRD}(t_i^*)) \quad (6)$$

Based on the distributions in Figure 3(b), if $\text{DMD} < 1$, the vehicle is presumed to be a single unit vehicle that is not pulling a trailer. Only 1 vehicle out of 16 vehicles pulling a trailer had $\text{DMD} < 1$ (a PVPT with zero DMD), or 6%. In addition, 40 out of 48 MUT (83%) had $\text{DMD} > 1$. Figure 3(b) also shows that 6% of PV and 15% of SUT had $\text{DMD} > 1$. From the development dataset, most of the vehicles with $\text{DMD} < 1$ can be correctly classified based on VL and VH. Correctly classifying the vehicles with $\text{DMD} > 1$ is the topic of the next section.

Additional Measurements of a Vehicle with Middle Drop

To correctly classify the vehicle clusters where $\text{DMD} > 1$, we segment a vehicle with middle drop into the front part of the vehicle (from the front bumper to the middle drop) and rear part of the vehicle (from the middle drop to the rear bumper). We then calculate the length of the front (FVL), height of the front (FVH), length of the rear (RVL), height of the rear (RVH), and the height of the vehicle at the middle drop (VHMD), as illustrated in Figure 4(b). Note that VH of a vehicle with middle drop corresponds to the maximum of FVH and RVH.

Vehicle Height at Middle Drop (VHMD)

The VHMD due to the hitch in PVPT or SUTPT should usually be lower than the VHMD due to the rear portion of a semi-trailer tractor in a MUT. We set a threshold height of the connection to be 2 ft. If VHMD is lower than this threshold the vehicle will be classified as either PVPT or SUTPT (depending on the FVL, discussed below). Otherwise, we need to check the other measurements to classify the vehicle. The VHMD is calculated via Equation (7) applied to the returns in the vehicle cluster.

$$\text{VHMD} = \max(h(t^*)) - \min(h(t)), \forall t \in [\text{FT}_r, \text{LT}_r] \quad (7)$$

Front Vehicle Height (FVH) and Front Vehicle Length (FVL)

The front part of a vehicle cluster with a true middle drop is either a PV, SUT, or the tractor of a MUT. As was shown in Figure 3(a), VH of SUT and MUT is usually higher than 8 ft, while VH of PV is usually lower than 8 ft. So we use FVH calculated via Equation (8) to capture height of the front portion of the cluster and if the height is below 8 ft, the vehicle is classified as PVPT. Otherwise, we need to check the other measurements to classify the vehicle.

$$\text{FVH} = \max(h(t)) - \min(h(t)), \forall t \in [\text{FT}_r, t^*] \quad (8)$$

In the case of PVPT or SUTPT the FVL calculated via Equation (9) is the VL of the PV or SUT portion of the cluster. From the development dataset we found the minimum length of the PV portion of the PVPT is above 15ft. If the FVL is below 15ft, we conclude that the middle drop is not due to a trailer and the vehicle is a single unit, PV or SUT.

$$FVL = V \times (t^* - FT_r) \quad (9)$$

Rear Vehicle Height (RVH) and Rear Vehicle Length (RVL)

The rear part of a vehicle with a true middle drop is trailer in a PVPT, SUTPT or MUT. If the RVH calculated via Equation (10) is sufficiently low (below 2.4 ft in the algorithm based on the development dataset), it is considered to be an empty flatbed trailer behind a PV or SUT and the complete cluster will be classified as either PVPT or SUTPT depending on the other measurements. If the RVH is sufficiently high (above 12ft in the algorithm based on the development dataset), it is considered to be a semi-trailer and the complete cluster will be classified as a MUT. Otherwise, we need to check the other measurements to classify the vehicle. The trailer length is captured by RVL, Equation (11). If RVL is below 28 ft, we assume this trailer cannot come from a semi-trailer truck.

$$RVH = \max(h(t)) - \min(h(t)), \forall t \in [t^*, LT_r] \quad (10)$$

$$RVL = V \times (LT_r - t^*) \quad (11)$$

The LIDAR Based Vehicle Classification Algorithm

The eight shape measurements and various tests described above are combined into the classification tree shown in Figure 5. This figure shows our classification algorithm for non-occluded vehicles that we produced, based on the development dataset. As noted above, before applying this algorithm we automatically differentiate between non-occluded and partially occluded vehicles. For the latter group we cannot be as precise as Figure 5 for our classification, as follows.

Classifying Partially Occluded Vehicles

While some information is missing about the partially occluded vehicles, the intersection between the occluded and the occluder dimensions bound the size of the occluded vehicle, i.e., the size of the occluded part of the vehicle is no larger than the size of the occluder vehicle. Being careful not to double count scans where both the occluder and occluded are seen "overlapping", the length of the non-overlapping portion of the occluder vehicle is measured and the length of the occluded vehicle is bounded by Equation (12). Overlapping is not an issue for height, and the height of the occluded vehicle is bounded by Equation (13).

$$VL_o \leq VL_o^{est} \leq VL_o + VL_{c-NOL} \quad (12)$$

Where,

VL_o^{est} = estimation of unknown actual vehicle length of the occluded vehicle,

VL_o = vehicle length seen from the occluded vehicle,

VL_{c-NOL} = vehicle length of non-overlapping portion of the occluder vehicle.

$$VH_o \leq VH_o^{est} \leq \text{Max}(VH_o, VH_c) \quad (13)$$

Where,

VH_o^{est} = estimation of unknown actual vehicle height of the occluded vehicle,

VH_o = vehicle height seen from the occluded vehicle,

VH_c = vehicle height from the occluder vehicle.

In this proof of concept study we only attempt to classify occlusions that involve two vehicles, though the principles could easily be extended to more complicated multi-vehicle occlusions. When classifying a partially occluded vehicle, the six classes are defined by static boundaries in the vehicle length and vehicle height plane, as shown in Figure 3(c). Because VL_o^{est} and VH_o^{est} each span a range, it is possible for a partially occluded vehicle to be associated with more than one class.

EVALUATION OF THE LIDAR BASED VEHICLE CLASSIFICATION ALGORITHM

Thus far this research has used a single development dataset collected on July 9, 2009 to derive the classification algorithm. In addition to the development dataset, we collected three additional freeway datasets and three arterial datasets for evaluation. We used a total of just over 15 hrs of data: 24 min for development and the rest for evaluation. All of the datasets were collected in the Columbus metropolitan area. All locations were visited a single time in this study except for I-71, which we visited twice. The facility, number of lanes, date, time period, duration, and distance between the LIDAR sensors and travel lanes are shown in the first few columns of Table 1. The next four columns show the average speed over all vehicles seen in the data collection period, the number of vehicles seen, the number of vehicles that our algorithm labeled as partially occluded, and the number of totally occluded vehicles as counted by the detectors. Among the freeway datasets two come from free flow (5.4 hrs) and two from mild congestion (3.5 hrs). All of the data sets come from clear weather conditions.

Overall the algorithm suspected 2,938 out of 27,450 vehicles (11%) are partially occluded and these vehicles are excluded from the classification algorithm performance evaluation in Tables 1 and 2. Instead, we separately evaluate the classification performance on partially occluded vehicles at the end of this section. The highest rate of partially occluded vehicles occurred at the I-71 site on Nov 19, 2009 under mildly congested conditions (22.6%), while the lowest rate of partially occluded vehicles on the freeway segments occurred on SR 315 (9.6%). Not surprisingly, across the four freeway datasets the percentage of partially occluded vehicles increased as the number of lanes increased and at the I-71 location, as congestion increased (17.2% in free flow and 22.6% in mild congestion).

The vehicle class was manually reduced from the video ground truth data for all 27,450 vehicles in these datasets and the partial occlusions were verified at that time (see [15] for an example of the data reduction tool). We also ran the classification algorithm from Figure 5 on the datasets and the last three columns of Table 1 evaluate the performance of the algorithm against the ground truth data. The errors are tallied on a per-vehicle basis, and thus, are not allowed to cancel one another across vehicles. Collectively, the algorithm correctly classifies 24,390 out of 24,512 non-occluded vehicles (99.5%) and misclassifies 122 vehicles (0.5%). The error rate was low across all seven datasets taken separately, the largest error rate was only 0.7%. The distance between the LIDAR and the roadway does not appear to have a large effect even though the further away a target vehicle is the smaller portion of the LIDAR field of view it occupies (and thus, the fewer angles in a LIDAR scan that provide vehicle returns). Among the freeway datasets the performance appears to degrade slightly as the average speed increases due to the 37 Hz sampling rate, but with only four datasets, the number is not large enough to draw any firm conclusions.

Table 2 shows the classification results by class against the ground truth data for all six evaluation datasets combined. The cells on the diagonal tally the number of vehicles where the LIDAR classification is the same as the ground truth classification, while the off-diagonal cells tally the incorrect vehicle classifications. The *false positive* column is the row sum, counting the number of times that the algorithm

erroneously classified vehicles as the given class, while the *false negative* row is the column sum, counting the number of times that the algorithm assigned a vehicle from the given class to the wrong class. The last row tallies the number of partially occluded vehicles-by-class that are excluded from the non-occluded LIDAR based vehicle classification. Often an operating agency will group PVPT with PV and SUTPT with MUT, for reference, these supersets are shown in the table, denoted PV* and MUT*, respectively. If using the two supersets, 14% of the errors (16 vehicles) in Table 2 and 16% of the errors (20 vehicles) in Table 1 would be eliminated. Overall, the number of false positives must be equal to the number of false negatives, and the algorithm incorrectly classified a total of 114 out of 23,010 vehicles (0.5%) in the evaluation datasets.

The most common errors are between PV and SUT because the length and height ranges of these vehicles overlap (30 SUT misclassified as PV and 15 PV misclassified as SUT), accounting for 39% of all errors. Also of note, we see 10 PV misclassified as MC. All of these PV were confirmed to have exceptionally short length, e.g., a 7.5 ft long commuter car (Smart Car). As with PV/SUT the MC/PV problem arises because the length and height ranges overlap between the two classes (3 MC were also misclassified as PV). This problem is not unique to LIDAR, the relatively new commuter cars will likely degrade the performance of most classification technologies when segregating MC. However, with the higher vantage point envisioned in our future research, the LIDAR should also be able to measure vehicle width, which should distinguish MC from commuter cars.

Finally, the algorithm for classifying partially occluded vehicles was applied to 1.5 hrs of the I-270 dataset. There were 465 partially occluded vehicles detected and of these, 219 are placed into a single feasible class (47% of partially occluded vehicles) and only six of these (3%) are incorrectly classified. The remaining 246 partially occluded vehicles are assigned two more feasible vehicle classes. Within this set, 34 (14%) were assigned all six classes. Out of the remaining 212 vehicles, 96% had the correct class among the two or more classes assigned to the given vehicle.

CONCLUSIONS

This paper developed and tested a side-fire LIDAR based vehicle classification algorithm. The algorithm includes up to eight different measurements of vehicle shape to sort vehicles into six different classes. The algorithm was tested over seven datasets (including one development dataset) collected at various locations. The results were compared against the concurrent video-recorded ground truth data on a per-vehicle basis. Overall, 2,938 out of 27,450 vehicles (11%) are suspected of being partially occluded and these vehicles are classified separately. Occlusions are inevitable given the low vantage point of the sensors in this proof of concept study. In future research we will investigate higher views (comparable to typical microwave radar detector deployments) to mitigate the impact of occlusions. These higher views should also provide additional features, e.g., vehicle width. Unlike video, a vehicle's width and height are easily separable in the LIDAR ranging data. The algorithm correctly classifies 24,390 of the 24,512 non-occluded vehicles (99.5%). While most side-fire detectors have challenges with occluded vehicles, the algorithms developed by this project are able to work around those problems. When a vehicle was partially occluded, we calculate the range of feasible length and height. These ranges are then used to assign one or more feasible vehicle classes to the given vehicle. Among these partially occluded vehicles, 47% were assigned a single class and 97% of these were correct.

Finally, this work also uncovered an emerging challenge facing most vehicle classification technologies: separating commuter cars from motorcycles. The two groups have similar lengths, axle spacing and height, though they differ in width and likely in weight.

REFERENCES

- [1] FHWA, *Traffic Monitoring Guide*. Federal Highway Administration, 2001
- [2] Khattak, A. J., S. Hallmark, and R. Souleyrette. Application of Light Detection and Ranging Technology to Highway Safety. *Transportation Research Record: Journal of the Transportation Research Board*, No. 1836, Transportation Research Board of the National Academies, Washington, D.C., 2003, pp. 7-15
- [3] Tsai, Y., Q. Yang, and Y. Wu. Identifying and Quantifying Intersection Obstruction and Its Severity Using LiDAR Technology and GIS Spatial Analysis. *90th Annual Meeting of the Transportation Research Board*, Washington, D.C., 2011
- [4] Veneziano, D., R. Souleyrette, and S. Hallmark. Integration of Light Detection and Ranging Technology with Photogrammetry in Highway Location and Design. *Transportation Research Record: Journal of the Transportation Research Board*, No. 1836, Transportation Research Board of the National Academies, Washington, D.C., 2003, pp. 1-6
- [5] Souleyrette, R., S. Hallmark, S. Pattnaik, M. O'Brien, and D. Veneziano, *Grade and Cross Slope Estimation from LIDAR-Based Surface Models*. MTC Project 2001-02, FHWA, U.S. Department of Transportation, 2003
- [6] Chenoweth, A. J., R. E. McConnell, and R. L. Gustavson. Overhead Optical Sensor for Vehicle Classification and Axle Counting. *Mobile robots XIV: proceedings of SPIE, vol. 3838*, Boston, MA, 1999
- [7] Cunagin, W. D. and D. J. Vitello Jr. Development of an Overhead Vehicle Sensor System, *Transportation Research Record: Journal of the Transportation Research Board*, No. 1200, Transportation Research Board of the National Academies, Washington, D.C., 1988, pp. 15-23
- [8] Yao, W., S. Hinz, and U. Stilla. Traffic Monitoring from Airborne Full-Waveform LIDAR – Feasibility, Simulation and Analysis. *International Archives of the Photogrammetry, Remote Sensing and Spatial Geoinformation Sciences*, Vol. 37(B3B), pp. 593-598, Beijing 2008
- [9] Grejner-Brzezinska, D. A, C. Toth, and M. McCord. *Airborne LiDAR: A New Source of Traffic Flow Data*. FHWA/OH-2005/14, FHWA, U.S. Department of Transportation, 2005
- [10] Yang, R., *Vehicle Detection and Classification from a LIDAR Equipped Probe Vehicle*, Masters Thesis, The Ohio State University, 2009
- [11] Redmill, K., Coifman, B., McCord, M., Mishalani, R., Using Transit or Municipal Vehicles as Moving Observer Platforms for Large Scale Collection of Traffic and Transportation System Information, *Proc. of the 14th International IEEE Conference on Intelligent Transportation Systems*, Oct 5-7, 2011, Washington, DC.
- [12] FHWA, *Federal Size Regulations for Commercial Motor Vehicles*, Federal Highway Administration, 2000
- [13] TRB, *Highway Capacity Manual*. Transportation Research Board of the National Academies, Washington, DC, 2000

- [14] Coifman, B. and S. Kim. Speed Estimation and Length Based Vehicle Classification from Freeway Single Loop Detectors. *Transportation Research Part-C*, Vol 17, No 4, 2009, pp 349-364
- [15] Coifman, B. and H. Lee. *Diagnosing Chronic Errors in Freeway Loop Detectors from Existing Field Hardware*. TS-514 Final Report, Caltrans, report number CA10-0978, 2011, pp. 139

TABLE 1 Summary of LIDAR data collected to evaluate the algorithm and the performance of the algorithm by each dataset

Data type	Road type	Location (direction)	Number of lanes	Date	Time period (Start time ~ End time)	Duration (hr: min)	Distance between LIDAR sensor and the nearest travelled lane (ft)	Average of the LIDAR speeds over the duration (mph)	Number of vehicles seen by LIDAR	Number of partially occluded vehicles	Number of totally occluded vehicles	Performance of the algorithm		% errors
												Success	Errors	
Development	Free-way	I-71 (SB)	4	July 9, 2009	18:09 ~ 18:33	0:24	58	63	1,813	311	65	1,494	8	0.5%
Evaluation	Free-way	I-71 (SB)	4	Nov 19, 2009	07:41 ~ 08:09	0:28	58	47	2,619	591	145	2,021	7	0.3%
		I-270 (SB)	3	Nov 2, 2010	09:29 ~ 14:29	5:00	15	65	13,397	1,376	422	11,934	87	0.7%
		SR-315 (NB)	2	Aug 12, 2010	14:57 ~ 17:57	3:00	2	41	6,900	660	n/a	6,230	10	0.2%
		Subtotal of Evaluation Freeway				-	8:28	-	-	22,916	2,627	567	20,185	104
	Arterial Rd.	Dublin Rd (SB)	1	Oct 28, 2010	07:32 ~ 08:57 14:30 ~ 15:55	2:50	2	36	1,344	-	-	1,337	7	0.5%
		Wilson Rd (NB)	1	Oct 28, 2010	09:08 ~ 09:56 16:02 ~ 16:54	1:40	2	36	666	-	-	664	2	0.3%
		Wilson Rd (SB)	1	Oct 28, 2010	10:18 ~ 10:58 17:00 ~ 18:00	1:40	2	38	711	-	-	710	1	0.1%
		Subtotal of Arterial Rd.				-	6:10	-	-	2,721	-	-	2,711	10
Evaluation data total					-	14:38	-	-	25,637	2,627	567	22,896	114	0.5%
Overall total					-	15:02	-	-	27,450	2,938	632	24,390	122	0.5%

TABLE 2 Comparison of LIDAR based vehicle classification and actual vehicle class from the six evaluation ground truth datasets

From six evaluation datasets		Ground truth data						Number of vehicles from LIDAR vehicle classification	False positive	
		MC	PV*		SUT	MUT*				
			PV	PVPT		SUTPT	MUT			
LIDAR vehicle classification	MC	31	10	0	0	0	0	41	10	
	PV*	PV	3	20,762	2	30	0	0	20,797	35
		PVPT	0	3	192	4	6	3	208	16
	SUT	0	15	6	688	2	9	720	32	
	MUT*	SUTPT	0	0	3	4	31	5	43	12
		MUT	0	0	1	2	6	1,192	1,201	9
Number of vehicles from ground truth data		34	20,790	204	728	45	1,209	23,010	114	
False negative		3	28	12	40	14	17	114		
Number of partially occluded vehicles that are excluded from LIDAR based vehicle classification		4	2,366	25	61	2	169	2,627		

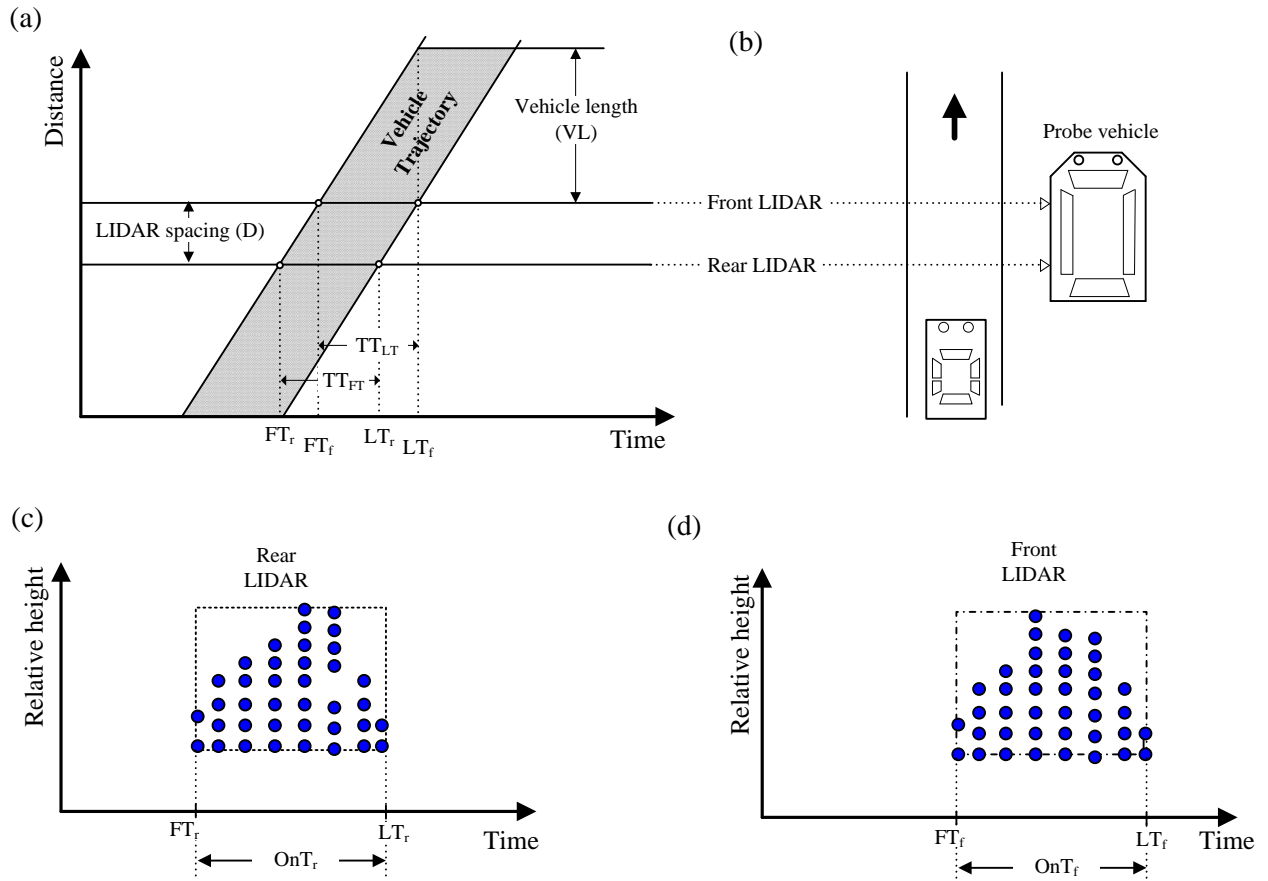


FIGURE 1 A hypothetical example of a vehicle passing by the two side-fire LIDAR sensors: (a) in time-space plane; (b) a top-down schematic of the scene; and the corresponding returns from the vehicle from (c) the rear LIDAR sensor and (d) the front LIDAR sensor.

(a)



(b)

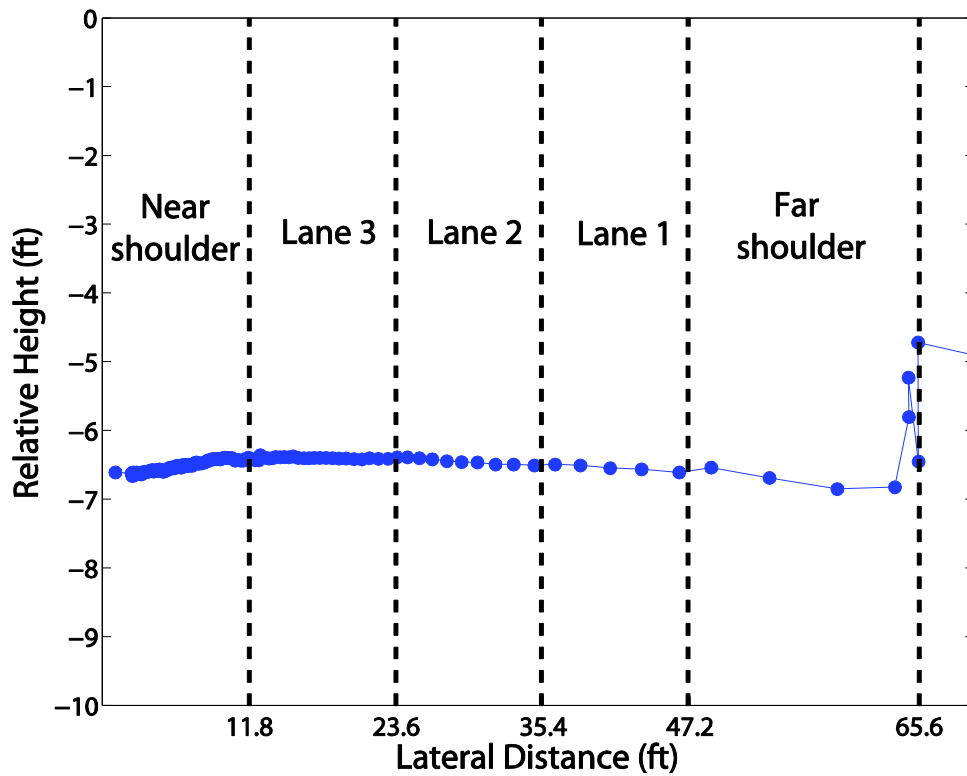


FIGURE 2 (a) the LIDAR data collection on I-270 southbound, on the west side of Columbus, Ohio; and (b) the corresponding background curve extracted from the data.

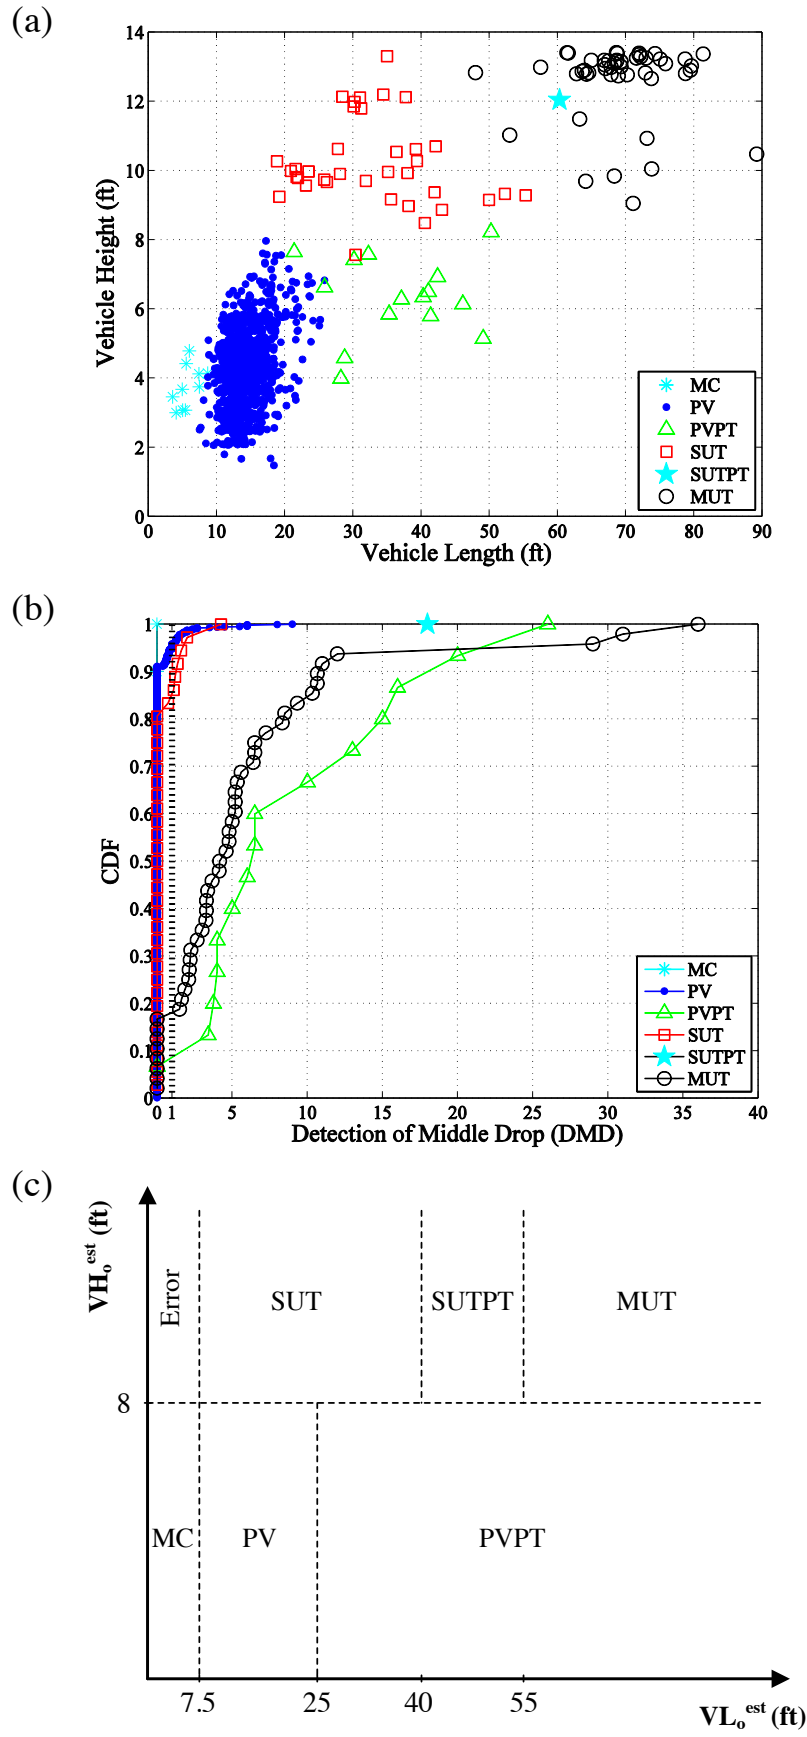


FIGURE 3 (a) A scatter plot of vehicle height and vehicle length of 1,502 non-occluded vehicles from the development dataset; (b) the cumulative distribution of DMD for these vehicles; and (c) the classification space for partially occluded vehicles.

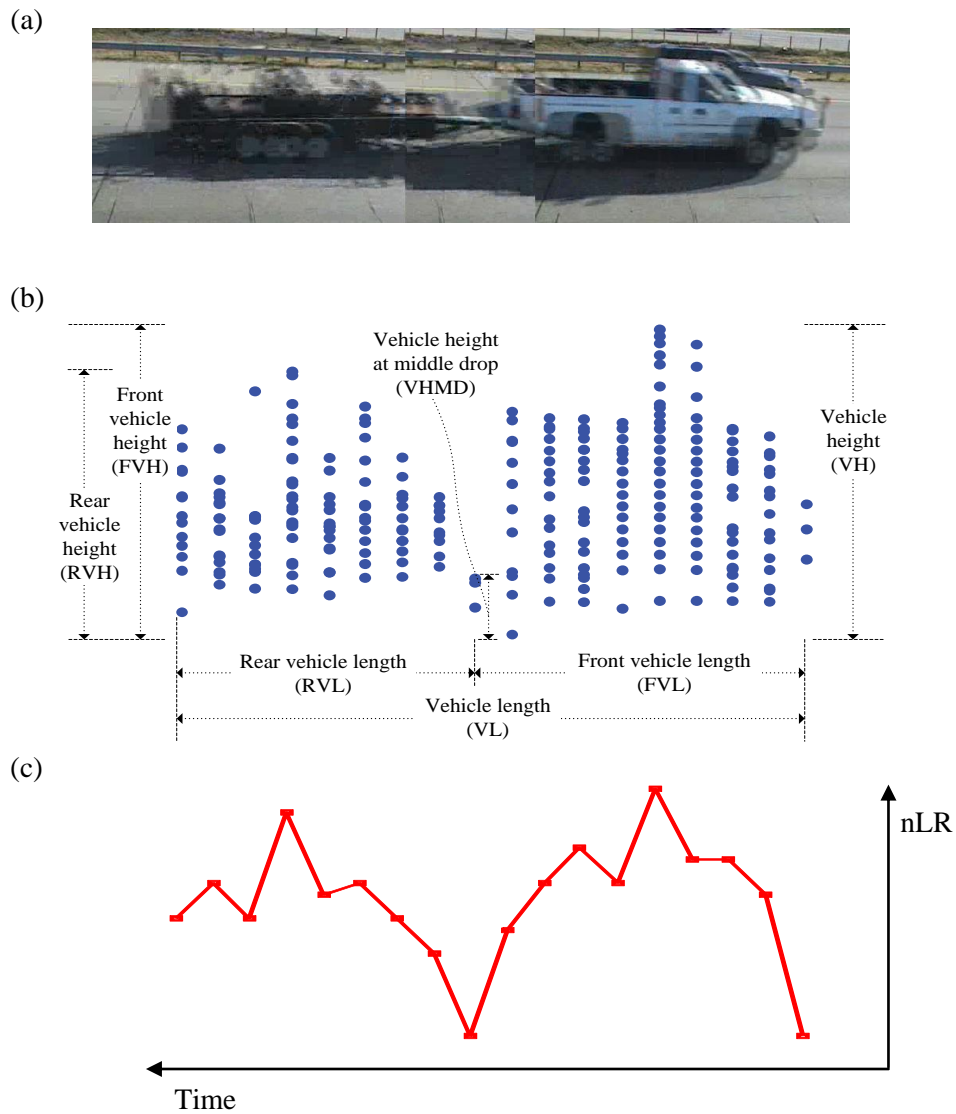


FIGURE 4 (a) A pickup truck pulling a trailer; (b) the corresponding vehicle cluster of returns and the various measurements used for vehicle classification; and (c) the number of returns by scan, capturing the vehicle shape. Note that time in Figure 4(c) is increasing to the left in this plot because the front of the vehicle is seen first and the vehicle orientation is presented consistent with the rest of the figure.

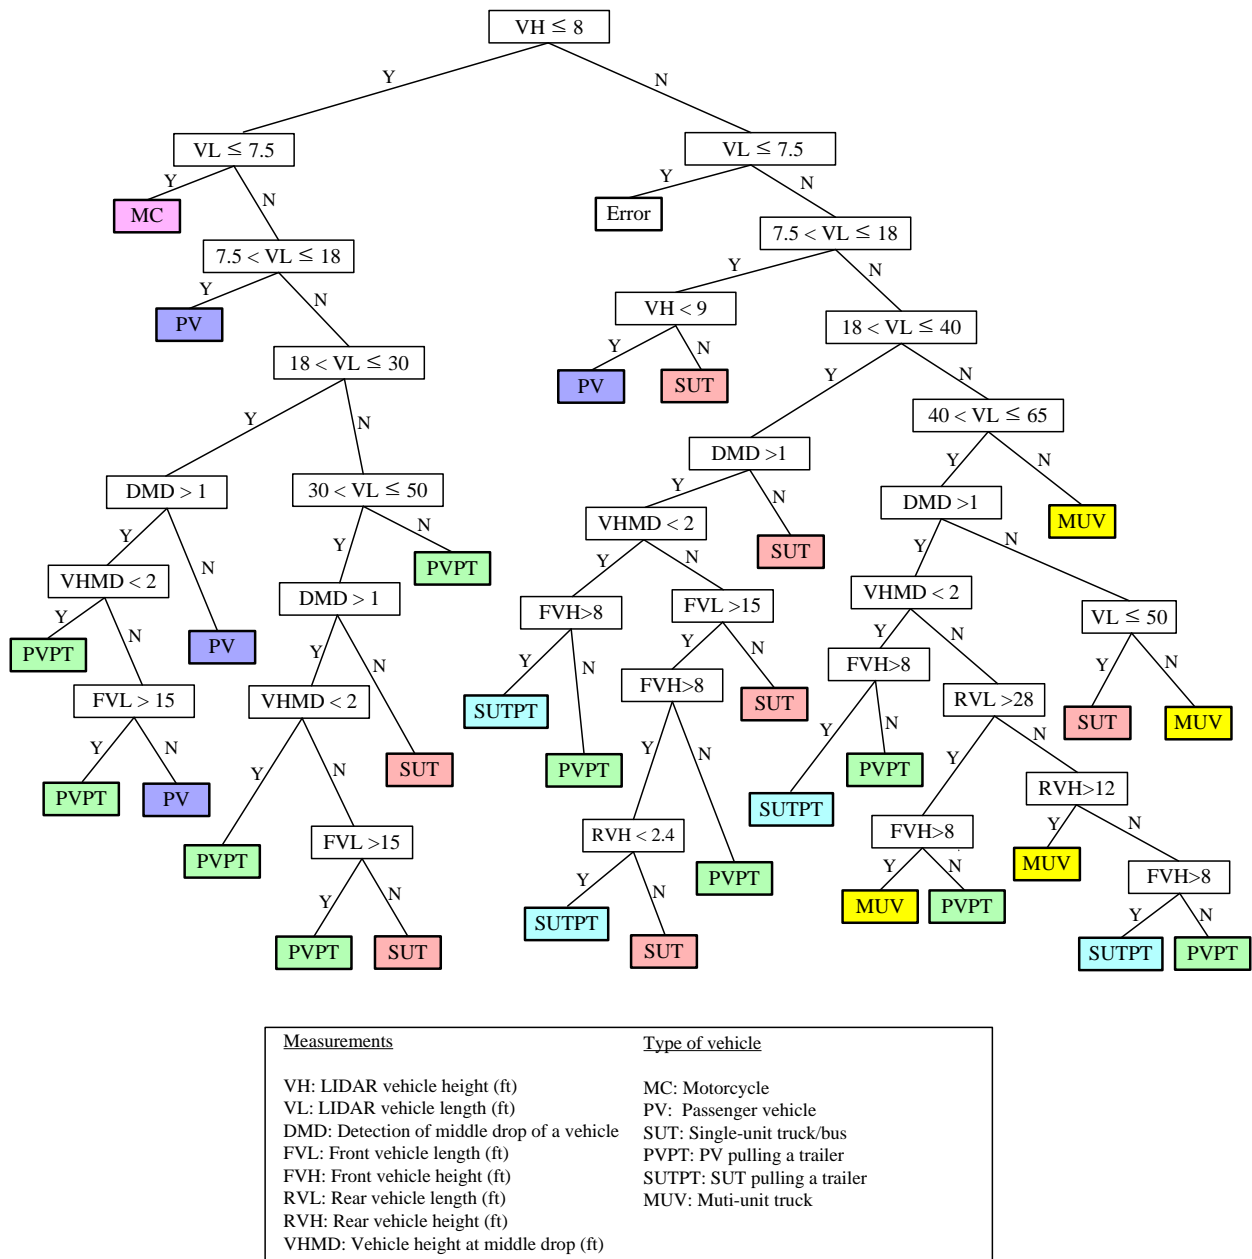


FIGURE 5 The decision tree underlying the non-occluded LIDAR based vehicle classification algorithm.