Sensory Augmentation for Increased Awareness of Driving Environment

Pranay Agrawal
John M. Dolan

Dec. 12, 2014

Technologies for Safe and Efficient Transportation (T-SET) UTC
The Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213

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 U.S. Department of Transportation's University Transportation Centers Program, in the interest of information exchange. The U.S. Government assumes no liability for the contents or use thereof.

Table of Contents

Problem 4
Approach/Methodology 4
Findings 9
Conclusions/Recommendations 10
Works Cited 11

Problem

The goal of this project was to detect road boundaries and stationary obstacles on the road using low-cost automotive-grade LiDAR scanners for the purpose of lateral positioning of the vehicle along the road. For autonomous driving, this information is especially useful on roads without proper lane markings or regions with poor GPS reception, such as tunnels, dense urban environment or under a heavy tree canopy. For manual driving, it can similarly be used as input to lane-centering assist systems.

Approach/Methodology

For the purpose of road boundary detection, we use the 6 IBEO LiDAR scanners installed in the GM-CMU Autonomous Driving Lab's Cadillac SRX vehicle, which give 360-degree coverage around the car [1]. Figure 1 shows the location of the LiDAR scanners and radar in the car, and Figure 2 shows the shows the range and field of view of these scanners along with their blind spots.

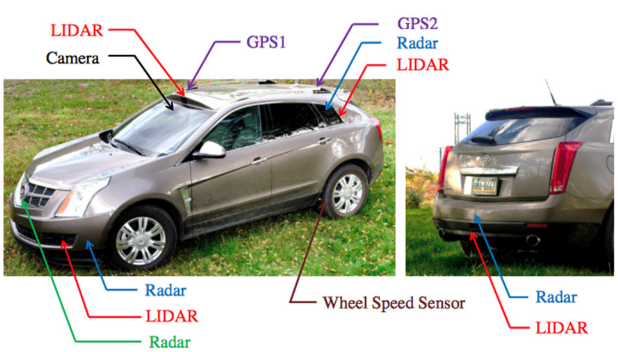

Figure 1: Sensor installation

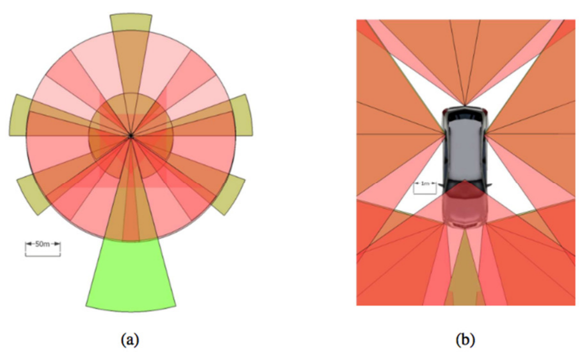

Figure 2: Perception system coverage with all sensors integrated. a) Maximum ranges b) Sensor Field-of-View (FoV) and coverage close to the car

Using these data, we try to model the road boundaries in the form of a clothoid [2], which defines a function of road curvature given road (longitudinal) length as $\kappa(l) = c_0 + c_1 l = \int_0^l \kappa(\varsigma) d\varsigma$, where c_0 is the initial curvature and c_1 is the curvature derivative. They are numerically represented as their 3rd-order Taylor series expansion given by

$$x = \frac{1}{6}c_1 y^3 + \frac{1}{2}c_0 y^2 + \beta y + x_{offset}$$

where x_{offset} is the lateral offset between the boundary and center of the vehicle, β is the heading angle with respect to the vehicle's driving direction, c_0 is the curvature, and c_1 is the curvature rate. Here, x is in the vehicle's lateral direction and y is the vehicle's longitudinal direction. The road is assumed to be curved in x as a cubic function of y.

Prior work on road edge detection using LiDAR, and LiDAR in combination with vision, used downward-looking scanners with scan planes making a large angle with the road [3, 4]. These scans are then integrated over time to create a 3-D rolling window to derive the road boundary. In our system, the scanning planes are almost parallel to the ground and each LiDAR scanner has 4 scanning planes per sensor with a total of 3.2-degrees vertical field-of-view. Because of this, we are able to scan a long section of the road-boundary with a single LiDAR scan, as shown in Figure 3, and obtain its shape. These road boundaries can then be tracked over time to account for inaccuracies in the LiDAR sensing and car's odometer, and obtain a robust road boundary even if we are not able to obtain the correct boundary in every frame.

A limitation of our sensing capability is that we are unable to detect curbs that are lower than ~8 inches in height, as we do not get the adequate number of points required for a good detection. Most of the high curbs, guardrails, Jersey barriers and tunnel walls are detected well.

Figure 4 shows the steps taken to obtain the road boundaries from each frame and track them over time. To obtain the road boundaries from each LiDAR frame, we first preprocess the IBEO data to obtain a database of all points which could potentially be associated with the road-boundary. This can be done based on its position relative to the car, the sensor from which the reading is coming, and the IBEO's internal filtering.

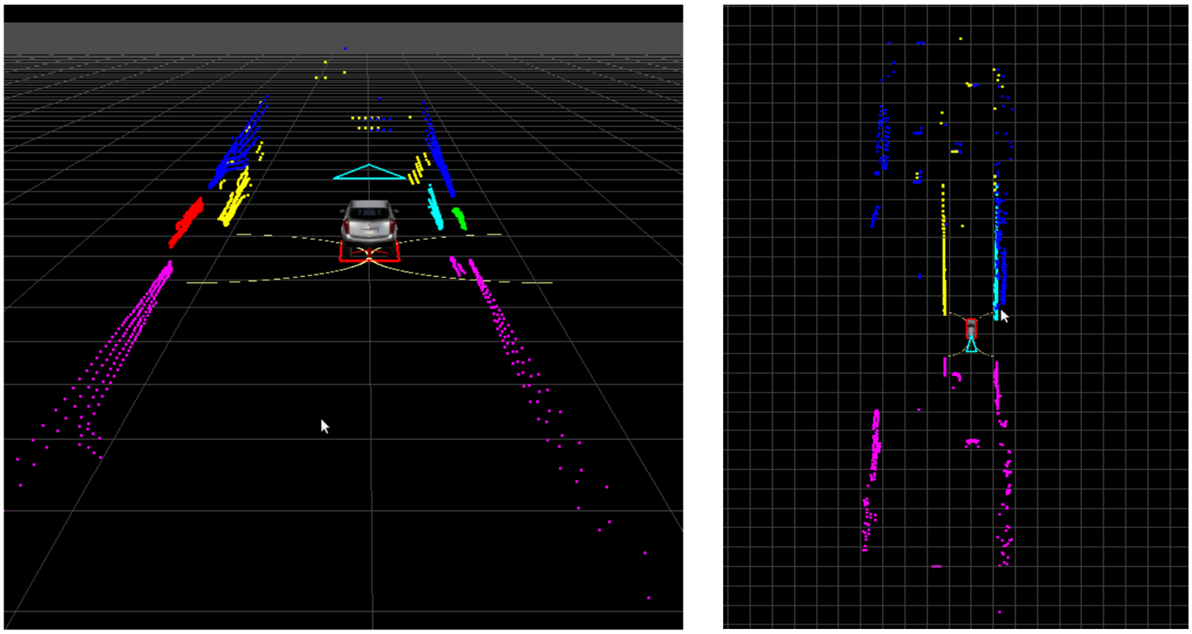

Figure 3. Raw IBEO scan at a single time with visible road boundary (a) inside a tunnel (b) on a highway

Figure 4 shows the steps taken to obtain the road boundaries from each frame and track them over time. To obtain the road boundaries from each LiDAR frame, we first preprocess the IBEO data to obtain a database of all points which could potentially be associated with the road-boundary. This can be done based on its position relative to the car, the sensor from which the reading is coming, and the IBEO's internal filtering.

With the set of preprocessed points, we want to group points into different clusters such that all the points representing a particular boundary are grouped into a single cluster. To do so, we find the direction of maximum variance for the data set using Principal Component Analysis (PCA). This direction of maximum variance is assumed to correspond to the direction of the road. The points are then compressed along the PCA direction and a k-means clustering is applied to it where the number of clusters k = $\alpha\sqrt{n}$, n being the total number of points and α being a suitable constant. Figure 4(a) shows a set of clustered preprocessed points. Following that, we find a cubic polynomial for each cluster using a RANSAC search. Using RANSAC over a curve-fitting technique ensures that the curve fits a section of the points very precisely while ignoring the outliers. Data points from all the clusters that fail to give a good cubic fit and the outliers from other clusters are removed from the dataset, giving us a filtered set of points that may correspond to a road boundary.

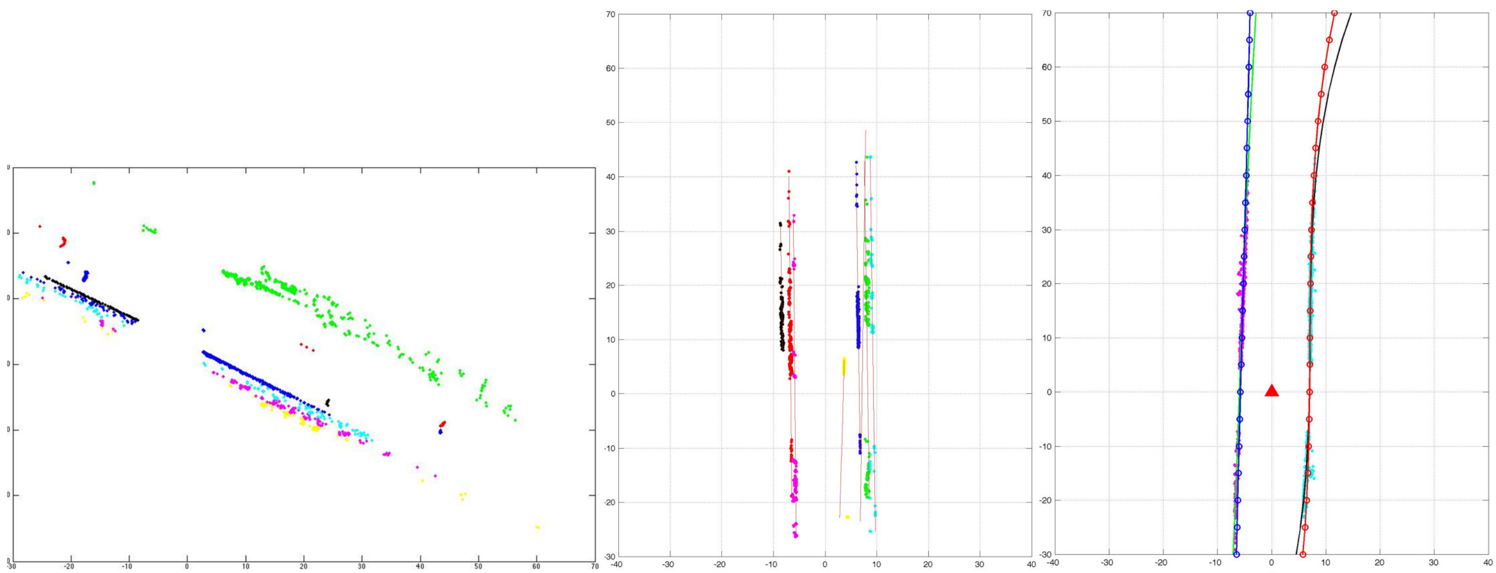

Figure 4. (a) Raw IBEO scan grouped into different clusters, shown by different colors. (b) Line fits generated using RANSAC for each cluster shown in red. Out of them, 3 lie on left of the road, 4 on the right of the road and 1 inside the road. (c) Final cubic road boundaries after UKF shown in blue and red which take into account the boundaries over previous time intervals and are more robust to sudden changes. The green and black lines are the cubic current-time-instant cubic fits.

On this set of filtered points, we again apply the PCA compression and k-means clustering, and we rotate the points with the road's direction as *y* and fit lines to each new cluster, as shown in Figure 4(b). This is done in order to identify the clusters that give us the most probable road boundary, as in many cases we get several good curve fits on both sides of the road. Out of this set of new lines, we rate the probability of their corresponding to the actual road boundary based on their distance to the car's center, their slopes, the range and density of points in the cluster, the number of points representing the curve, etc. and we find the lines on both sides of the car with the highest probabilities. The clusters to which those fits belong are kept while the others are removed from the dataset of LiDAR points. We then transform these points to an ego-car perspective and find the best cubic fits for both sides of the roads using RANSAC and determine their probability of representing the actual road boundary. If high, we designate those curves as the road boundaries for that time.

In order to track the left and right boundaries of the road over time, we use an unscented Kalman Filter (UKF). The UKF uses an unscented transformation to model nonlinear dynamics [5], which often works better than an Extended Kalman filter, which uses a linearization technique to approximate non-linear dynamics using the Jacobian. Here, the state is the road boundary expressed as x = $[x_{offset}, \beta, c_0, c_1]$ and the current road boundary, expressed as a set of points $\{x_i^L, y_i^L, x_i^R, y_i^R\}_{i=1:n}$ is taken as the observation. Figure 4(c) shows the road boundary after implementing the UKF tracking.

For the purpose of detecting static obstacles on the road, such as parked cars, trailers, etc., we pick out all the points that lie inside the current road boundaries in the original dataset of LiDAR scans. Out of these points, we find clusters of points with a high density and find line-fits perpendicular to the road boundaries. These clusters and lines are then tracked over time to determine their velocities relative to the car. The obstacles that appear to be stationary are then included in the static road map.

Findings

In this work, our goal was to reliably detect road boundaries on highways and tunnels using relatively inexpensive automotive-grade LiDAR sensors compared to high-cost, high-capability sensors such as the $70,000 Velodyne spinning laser sensor used in order to get a high-density 3D range point cloud. Such a sensor is not practical for production automobiles from a cost or appearance/integration standpoint.

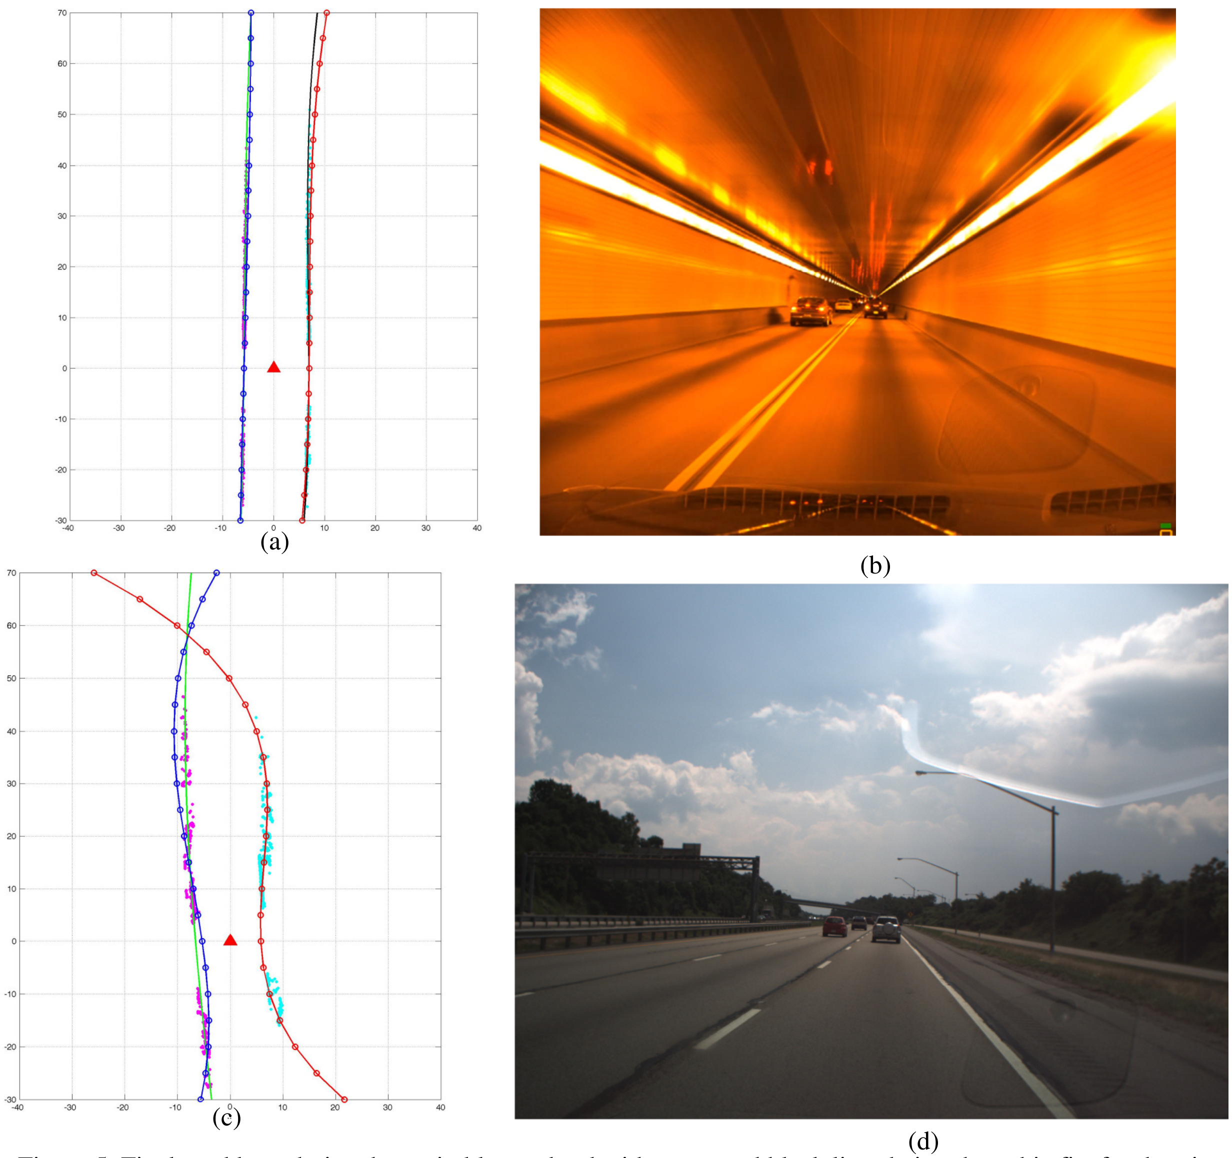

Figure 5. Final road boundaries shown in blue and red with green and black lines being the cubic fits for that time instant for (a) tunnel and (c) highway with guardrails, and the corresponding visual data shown in (b) and (d).

Figure 5 (a) shows an example of a road boundary detected inside a tunnel (b). Due to the lack of noise from inside the tunnel and good data points from the high curb, we are able to get accurate road boundaries. Figure 5 (c) shows an example of a road boundary detected on a highway with guardrails (d). Due to the lack of a good boundary on the right, we are not able to identify it accurately.

More generally, for regions where the road boundary is high and not very far apart from the car (<10 meters on either side) such as inside a tunnel or on a highway with high Jersey barriers, the road boundary is accurately detected at every time step and smoothly tracked over time. In regions with no boundary clearly observable by a LiDAR, such as a short guardrail or a curb, the boundaries are only detected approximately 65% of the time and even with the UKF, they cannot be tracked well at most times.

Conclusions/Recommendations

This project lays the groundwork for the detection of road boundaries along highways and tunnels using relatively inexpensive automotive-grade LiDAR. This is very useful for the purpose of lateral positioning of the vehicle along the road, especially on roads without proper lane markings or regions with poor GPS reception. We are expanding this into detecting and mapping stationary obstacle on the road including parked cars and build a static map of the road.

Future work should include:

- Refining the road-boundary detection and using adaptive algorithms that deal with varying scenarios
- Understanding the various failure scenarios and modifying our approach to handle them better

- Use radar along with LiDAR to track vehicles along the road and differentiate parked cars from stationary cars
- Build a static road map by combining the road-boundary and stationary obstacle inside the road.

Works Cited

[1] Wei, Junqing, et al. "Towards a viable autonomous driving research platform." *Intelligent Vehicles Symposium (IV), 2013 IEEE*. IEEE, 2013.

[2] E. D. Dickmanns and B. D. Mysliwetz, “Recursive 3-d road and relative ego-state recognition,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 14, no. 2, pp. 199–213, 1992.

[3] Zhang, Wende. "Lidar-based road and road-edge detection." *Intelligent Vehicles Symposium (IV), 2010 IEEE*. IEEE, 2010.

[4] Qin, Baoxing, et al. "Road detection and mapping using 3D rolling window." *Intelligent Vehicles Symposium (IV), 2013 IEEE*. IEEE, 2013.

[5] S.J.JulierandJ.K.Uhlmann,“An extension of the kalman filter to nonlinear systems,” in *Proceedings of SPIE Signal Processing, Sensor Fusion, and Target Recognition*, 1997.