# LiDAR Observations by Motion Compensation and Scan Accumulation

Antoine Lima, Anthony Welte, Philippe Bonnifait, Philippe Xu

**HAL Id: hal-02946634**

**https://hal.archives-ouvertes.fr/hal-02946634**

Submitted on 27 Aug 2021

# LiDAR Observations by Motion Compensation and Scan Accumulation

Antoine Lima*, Anthony Welte*, Philippe Bonnifait* and Philippe Xu*

*Abstract*— In dynamic localization problems, the observations used from exteroceptive sensors are usually obtained from a single measurement. However, there are cases where the current measurement is not sufficient to detect the referenced landmark or to get a sufficient level of accuracy. In this study, a point cloud accumulation strategy is used to improve the resolution of a LiDAR sensor along its sparse axis. In particular, we are interested in the detection of markings transverse to the road axis in order to improve the accuracy of localization of an autonomous vehicle when approaching intersections or roundabouts. We present a method that allows the construction of an accurate observation with an associated observation model based on a High-Definition (HD) map through an accumulation of scans as the vehicle moves, by compensating the vehicle motion. The parameters of the accumulator are studied in terms of detection and accuracy. The quality of the observations and their impact on the localization quality are analyzed using real experiments carried out with an experimental vehicle equipped with a low-cost GNSS receiver, dead-reckoning sensors and a ground truth system.

## I. INTRODUCTION

Accurate localization is a prerequisite in many navigation tasks of autonomous vehicles ranging from vehicle control to path following. A Global Navigation Satellite System (GNSS) receiver coupled with Dead-Reckoning (DR) is often used as a basis for the localization scheme. Unfortunately, GNSS signals can be degraded by non line-of-sight and multi-path in particular in urban environments, and a DR pose estimate drifts over time. A solution is to embed perception sensors such as cameras, RADARs or LiDARs to detect visual features in the surrounding environment and to use them to localize the vehicle. These features are usually either mapped in a previously acquired map or localized and mapped simultaneously, *i.e.*, SLAM-based approaches.

Our approach falls in the former category, in particular, we make use of a High-Definition (HD) map containing accurate geo-referenced locations of road markings. Because the reflectivity of road markings is higher than the rest of the road, as illustrated in Figure 1b, a LiDAR sensor can extract such features using the intensity of reflected laser beams. Simple shapes such as lane markings are usually detected by single point cloud methods and used as observations to improve the lateral localization but, more complex markings, in particular in the longitudinal direction, accumulation of scans is required since a single scan does not contain enough information. Our proposal aims at improving the longitudinal resolution of LiDAR data by scan accumulation while compensating the motion of the vehicle during the

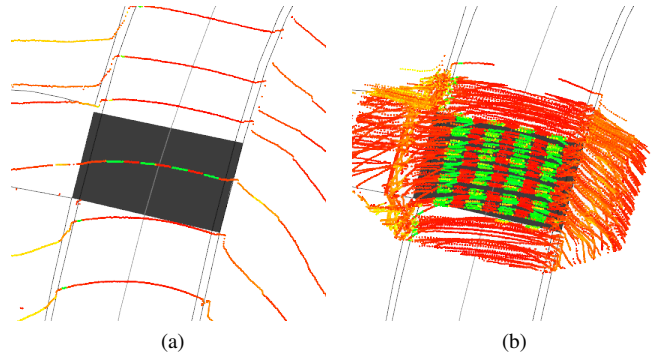Université de Technologie de Compiègne, CNRS UMR 7253 Heudiasyc, France.

Fig. 1: Comparison between a single LiDAR scan and an accumulation around crosswalks. Colors from red to green represent low to high reflectivity. In the background are the road borders and crosswalk areas of an HD map.

scanning thanks to its odometry. To study the accuracy of this approach, we focus on pedestrian crossings, as their shape does not allow for single frame detection while multi-frame provides longitudinal information, as illustrated in Figure 1a. Note that throughout this paper, only zebra crossings, which are designed with several strips parallel to the lane, are considered, even though detection of other types of marking such as stop or yield lines, or even road arrows could also be helped by this approach. The detector applies ground and markings filtering and incorporates a coarse pre-filtering around crosswalk areas that makes use of the HD map. Strips are then detected and used to compute the crosswalk pose. Detections are integrated in a localization scheme by an observation model in which measurements constrain the position along the crosswalks axis, *i.e.*, in the longitudinal direction of the road.

The main contributions of this paper are:

- An accumulation method to get an accurate observation in the sparse direction of the LiDAR sensor;
- A detector and a model associated to this observation;
- A discussion on strategies to integrate such temporally correlated observations in a localization scheme.

The paper is organized as follows. In Section II, a state of the art of localization using visual cues from LiDARs is exposed. Then in Section III, the crosswalk detection process is explained. The localization scheme and crosswalk observation model are presented in Section IV and then experimental results and parameter tuning are given in Section V.

## II. State of the Art

The field of localization using visual features is widely studied but remains an open issue. Two sensors are usually used as sources of information for visual cues: cameras and LiDARs. Localization using LiDARs can be separated into two families: dense and sparse. In dense methods, there is no feature detection and measurement are registered against a dense representation. Levinson and Thrun [10] used a grid to store intensities as Gaussian distributions and formulated the registration problem as finding a transformation between the measured point cloud and the grid in the estimated body frame. An extension of this idea was introduced in [19] by storing both intensity and height as Gaussian mixtures then using a branch-and-bound registration algorithm. More generally, dense methods store a 3D map previously built by a LiDAR and formulate the localization problem as a registration one in that map using algorithms such as Iterative Closest Point (ICP) [14], Normal Distributions Transform (NDT) [12] or derivatives.

Sparse approaches, on the contrary, extract features from measurements and match them to a prior map to refine a global pose given by a GNSS receiver, for example. Feature detection on LiDAR can once again be grouped into two categories: 2D grid-based or 3D. In grid-based approaches, the point cloud is projected onto a 2D grid placed on the ground. These grids are assimilated to bird-eye-view images and interpolation such as Inverse Distance Weighting is used to fill in gaps between impacts. Castorena and Agarwal [2] used the intensity gradient to extract markings and register them against a prior map. In [20], markings are extracted by computing an integral image ([1]) and clustered. They are then detected either by testing their size for linear markings, or by matching patterns for more complex markings. Yu *et al.* [21] extracted markings using an extension of the Otsu dynamic thresholding [13] incorporating the distance to the sensor and the point density. Markings are then clustered and refined before being classified using a combination of deep learning for small markings and Principal Component Analysis (PCA) for rectangular markings. Finally, in [6], Hough transforms are applied to the grid to detect along-track and cross-track features that are then tracked using a Kalman filter. The use of grids was introduced to reduce the number of points and simplifying the problem from a 3D to 2D one, but this has the disadvantage of losing accuracy depending on the desired cell size.

Finally, in 3D-based detection approaches, the scanning pattern can be used to detect changes such as curbs. In [8], Hata and Wolf detected curbs using ring compression robust to obstacles and extracted markings using a modified robust version of Otsu. An occupancy grid is then used to integrate these detections to a Monte Carlo localization. In [16], a collision-based curb detection is introduced and detections are accumulated using an Inertial Measurement Unit (IMU) before being refined by a two-pass beam model. Non-SLAM point cloud accumulation methods have been explored in studies such as [15], where the combination of successive frames as well as the change between them has been used
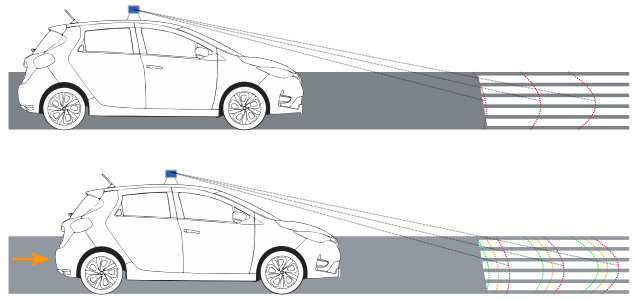


Fig. 2: Sweeping of the crosswalk as the vehicle moves.

to detect pedestrians.

Most of these methods use the intensity measure, but its sensitivity is a common LiDAR issues. Out of the factory, LiDARs are generally not sufficiently well calibrated and intensities depend on the beam which did the measurement. The intensity can also vary with distance, as the incident angle diminishes. Those issues can be tackled by offline calibrations such as the ones proposed by Levinson and Thrun in [10], [11]. However, the major issue appears in bad weather. While rain and snow can produce light noise in the geometrical measurements, it can greatly degrade the intensity readings, in particular on pavement, as studied in [9].

## III. Crosswalk detector

### A. Problem Statement

Detection of complex markings on the ground requires good along-road measurement accuracy and resolution in order to accurately measure where the markings begins and stops. This can be an issue with low density LiDARs or certain sensor placements. For example, despite their good measurement accuracy and along-ring resolution, LiDARs, when placed on the roof parallel to the ground, can have a ring resolution up to several meters when hitting the ground because of their construction using stacks of lasers, as can be observed in Figure 1a. In this section, a crosswalk detector capable of computing observables on crosswalks from LiDAR scans thanks to DR information and buffers is detailed. We chose not to use grids nor machine learning in order to build an accurate observation model for the localization filter.

### B. Pre-Processing

Typical scans being composed of tens of thousands of points, a buffer can have to hold millions of points, making it hard to guarantee real-time processing with any subsequent algorithm. We apply ground filtering, based on the ground plane fitting algorithm from [22], and coarse road filtering that uses a prior HD map and a localization estimate to remove points outside the road closed polygon. For this, an area-based test is realized for each point, where sums of triangle areas are compared to that of the polygon. This significantly reduces the number of points without too much information loss, contrary to ray-casting approaches, polygon enlargement [5] or others detailed in [7]. As our approach

uses ground points in the close vicinity, we only focus on lower rings. We have noticed that the factory intensity calibration was sufficient for this use case. Therefore, we have not used a calibration method such as described in [11].

### C. Buffering

The classical solution to accumulate or "buffer" multiple point clouds is to transform the point clouds in a static working frame and taking their union. The issue with this method is that bad transformations (i.e., imprecise localization) lead to incoherent point cloud. To dampen this issue, grids have been developed [10], [8] and work by projecting points in a 2D grid placed on the ground.

Our proposal is to remain in the vehicle frame by storing untransformed point clouds alongside their relative motion provided by the vehicle DR sensors, wheel encoders and yaw rate gyro. Point clouds are accumulated by bringing past scans to the present using the stored motion and merging them together, resulting in a denser point cloud referenced in the current sensor frame. The advantage of this method is that the vehicle does not have to know its global pose and can rely exclusively on its on-board DR sensors. In the absence of accurate DR sensors, other approaches could use ICP or derivatives to register these scans prior to ground and road filtering. However, as our DR model, described in [17], is fairly stable over short periods of time (e.g: two seconds), we had no need for such registrations.

This has the effect of progressively scanning the surface of the road as illustrated in Figure 2, using movement to produce an instantaneous measurement. The buffer duration, or number of accumulated point clouds, can be changed to match the resolution and placement of the LiDAR, the accuracy of the DR model or the application of its output, depending on whether short bursts of accurate points are preferable to long term but less accurate points. This is detailed in Section V-B but for our use, an accumulation of 1 to 3 seconds is sufficient to drastically increase the longitudinal resolution using a LiDAR scanning at 10 Hz. The vehicle speed is also an important factor, as it directly impacts the resulting longitudinal resolution. Too slow and the rolling buffer will have lost the start by the time the end is reached and too fast and the distance between consecutive point clouds will be greater than the marking. This limits possible speeds between 10 and 50 km/h with our sensor placement.

### D. Strips Detection

In our approach, crosswalk strips are first detected from the accumulated point cloud filtered around crosswalks before being fused in a single crosswalk detection. This allows for a more robust detection as strips constitute redundant information and can be easily rejected based on their shape.

The LiDAR points hitting the strips are extracted using the traditional Otsu dynamic thresholding [13]. This method works by maximizing a weighted between-class variance inside a bimodal histogram. Assuming measured intensities are separated into a "pavement" and "marking" distributions, which is reasonable given the previous filtering steps, markings can be extracted by retaining points whose intensity is above the computed threshold. If some strips are weathered, they will thus be classified as pavement and not be used for crosswalk estimation, which is not an issue as long as there are other good strip detections.

These markings are then clustered together in order to process each strip independently. The chosen algorithm for this is called DBSCAN [4], and is similar to classic Euclidean distance-based clustering, to the difference of a density constraint. With DBSCAN, a class is propagated only when there is enough neighbors, otherwise, a point is considered noise. This choice is driven by the presence of between-strips outliers caused by rectification errors. The propagation of one strip to its neighbors is thus prevented by the low point density between strips.

Oriented bounding boxes are then fitted on each cluster using PCA. It returns the cluster center $(x_s, y_s)$ and principal orientation $\theta_s$ that we will assume to be that of the strip, as strips tend to be longer than they are wide. However, in order to reject obvious outliers, the length $L_S$ and width $W_S$ of the bounding box are compared to their theoretical shape $L$, $W$ and $\theta_R$ by using confidence intervals. We consider that, after the successive outlier rejection mechanisms, the strip poses are used to compute the pose of the crosswalk as the average of its constituent strips, $C = \begin{bmatrix} x_C, y_C, \theta_C \end{bmatrix}^T$ in the body frame of the car. Depending on sensor placement, markings can be seen from tens of meters ahead to tens of meters behind. Thanks to the accumulation, observation can be made as long as there are enough points, up to several seconds even with temporary occlusion such as when the vehicle passes over the crosswalk.

## IV. LOCALIZATION FILTER

In order to evaluate the quality of crosswalk detection and to test the observation model, a state filtering scheme has been set up. The state and evolution model are taken from [17]. An Extended Kalman Filter (EKF) has been implemented with measurements from a low-cost Ublox GNSS receiver that provides a global but inaccurate pose, the four wheel-encoders and the yaw rate gyro of the vehicle.

### A. Marking Representation

In our HD map, illustrated in the background of Figure 1a, crosswalks are represented by a single closed polygon that covers all the crosswalk strips (and not each strips independently), as visible in gray on Figure 1a. It is attached to a series of nodes describing the center of the lane from which the orientation of the crosswalks can be extracted. In the following section, the polygon center $[x_O, y_O]$ alongside the road orientation $\theta_O$, expressed in the working frame, are used. Their accuracy is studied in [18] and is in the order of centimeters. It is worth noting that this information is also accessible in OpenStreetMap (OSM), where crosswalks are simply points referenced by lanes.

## B. Crosswalk Observation Model

In order to incorporate the crosswalk detection into a localization filter, we developed an observation model based on the line-to-line constraint proposed in [3]. In this paper, representing lines as distances and orientations came naturally as a consequence of using Hough transform. However, in our case, we chose this representation because a crosswalk (as well as most other markings) can only give accurate information in one of its directions, here along the strips and this representation thus serves as a 1D constraint over this direction.

In this observation model, illustrated in Figure 3, we assume the LiDAR to be rigidly attached to the vehicle body frame and properly calibrated so that the output of the detector is expressed in the body frame $M$. $X$ is the state to estimate that defines the frame $M$ and is composed of its position $x, y$ and its heading $\theta$ defined in the world frame as well. Let

$$z = {}^{M}\begin{bmatrix} \rho_C \\ \theta_C \end{bmatrix} = {}^{M}\begin{bmatrix} x_C \cos\theta_C + y_C \sin\theta_C \\ \theta_C \end{bmatrix} \quad (1)$$

be the measurement, as described in previous sections and

$$h(x) = {}^{M}\begin{bmatrix} \rho_O \\ \theta_O - \theta \end{bmatrix}$$
$$= {}^{M}\begin{bmatrix} (x_O - x)\cos\theta_O + (y_O - y)\sin\theta_O \\ \theta_O - \theta \end{bmatrix} \quad (2)$$

be the estimated measurement projected in the body frame.

Then, $H$, the Jacobian of $h(X)$ can be calculated as

$$H = \begin{bmatrix} -\cos\theta_O & -\sin\theta_O & 0 \\ 0 & 0 & -1 \end{bmatrix} \quad (3)$$

and letting $\sigma_\rho$ and $\sigma_\theta$ respectively be the longitudinal and orientation standard deviations (determined from the detector uncertainty analysis of Section V-B), the observation covariance matrix $R$ of the filter can be defined as:

$$R = \begin{bmatrix} \sigma_\rho^2 & 0 \\ 0 & \sigma_\theta^2 \end{bmatrix} \quad (4)$$

Because observations are produced from overlapping buffers, observation errors for the same crosswalk are correlated. This has been mitigated by enlarging the evolution model noise.

## V. EXPERIMENTAL RESULTS

### A. Experimental Setup

The method has been evaluated on data recorded with a Renault ZOE experimental vehicle from the Heudiasyc laboratory. The vehicle was equipped with a Velodyne VLP-32C LiDAR sensor. The vehicle internal sensors (gyro, wheel encoders) were accessed through its CAN bus. Ground truth positioning has been obtained by a Novatel SPAN-CPT system providing localization by fusing high accuracy IMU and Real Time Kinematics (RTK) corrections. The sensors were interfaced and recorded using the Robot Operating System (ROS) framework.
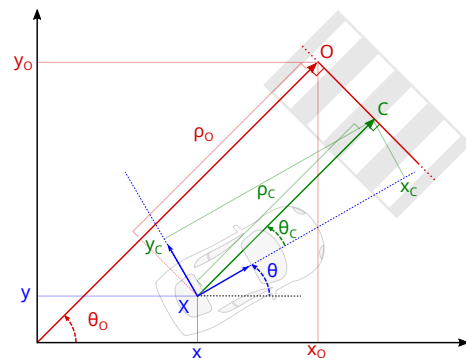


Fig. 3: Illustration of the observation model. In blue are vehicle-related frames and coordinates, while green is for measurement and red for map-relative information.



Fig. 4: Satellite view and trajectory of the experiment. Crosswalks are highlighted in red.

The trajectory used in these experiments is located in Rambouillet, France and is illustrated in Figure 4. It consists of a stretch of 400 m of mostly straight roads where 5 crosswalks can be observed. This section is followed by 450 m of trajectory on a succession of roundabouts where over five crosswalks can be observed. Speeds during the recording were about 40 km/h on straight roads and 25 km/h in the vicinity of roundabouts.

### B. Accuracy analysis of the crosswalk observations

The accuracy and number of observations of the detector depend on several factors, such as the buffer length or the odometry accuracy. We studied these factors through crosswalk detection using ground truth as position so that the observation error can be assumed to be entirely caused by the detector. We evaluate using the observation residual instead of using the classical precision and recall metrics by lack of observation ground truth.

*1) Influence of buffer length:* The buffer length is also a critical parameter of the detector. Intuitively, one may expect that longer buffer would lead to better accuracy and more observations. The experimental results in Figure 5 shows that on the contrary, a slight bias appears above 5 s. This behavior can be explained by several reasons. Longer buffers means that more DR measurements are accumulated, which in turns also accumulates more of their respective errors, increasing the overall one. Moreover, a long buffer results in observations being produced even though the crosswalk has
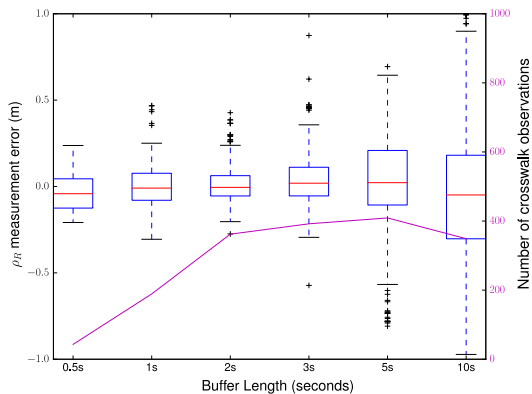
Fig. 5: Error distribution of the crosswalk detector and number of detections for different buffer lengths.
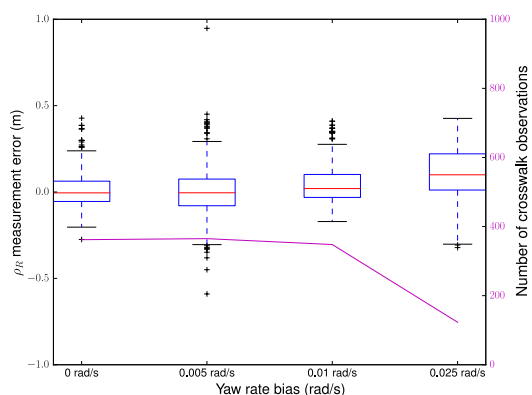


Fig. 6: Error distribution of the crosswalk detector and number of detections for yaw rate biases.

not been observed for some time. Therefore, it is preferred to have a shorter buffer but sufficiently large to be able to detect crosswalks and provide an acceptable number of observations to the localization process. The small number of crosswalks detected at length $0.5$ s confirms the benefits of accumulating point clouds. For the rest of the experiments, the buffer size has been chosen at 2 seconds as it is well balanced for crosswalks of 2 or 3 meters at nominal speeds.

*2) Influence of DR measurements accuracy:* The vehicle motion between two scans used for accumulating is sensitive to the accuracy of the DR system. This is all the more relevant since automobile grade sensors can be biased, especially gyros. We evaluate the effect of bad inertial measurements on the detector by artificially adding a yaw rate bias to the measurements. Figure 6 presents the loss of accuracy and availability that has been observed for different bias values (for a 2 s buffer).

For this size of buffer, the detector experiences little effect of a gyro bias up until a bias of $0.025$ rad/s. This remains acceptable since typical observed biases are around $0.004$ rad/s [17]. A 2 second buffer can therefore be assumed short enough not to be significantly affected by inertial errors. It is important to note that the longer the buffer, the more sensitive it is to this effect.

## C. Localization accuracy and consistency analysis

The contribution of marking information to global localization is evaluated here. In the following results, depicted in Figure 7, the ground truth positioning is only used to evaluate the localization accuracy of the filter. Here, the point clouds are accumulated using speeds and yaw rates computed using the vehicle on-board sensors. Since the filter here uses inaccurate GNSS positioning, the accuracy gained from crosswalks is easy to observe. In general, a notable improvement in localization error can be observed when crosswalks are detected, such as at times $t = 20$, $42$ and $105$ s, where the 2D error drops from $0.6$ m to $0.1$ m on average. At times $t = 46$, $117$ and $130$ s, crosswalks are observed while turning, as can be seen from the significant angle change in the bottom figure of Figure 7. This results in the marking being seen from different angles which does not improve localization, as the DR model is more sensitive to rotation error than translation error. When navigating in roundabouts (starting at $t = 95$ s), observing a single crosswalk does not have a significant impact on the localization error, but when two are seen at the same time, the state is constrained in two different directions, resulting in significant error reduction. On the whole trajectory, the average 2D error is reduced from $0.68$ m to $0.49$ m and the longitudinal error from $0.30$ m to $0.19$ m when in the presence of crosswalks. Finally, the filter and thus the accumulation and detection approach are consistent, as the East and North errors are always inside their $3\sigma$ confidence interval.

## VI. CONCLUSION AND PERSPECTIVES

The local accumulation of LiDAR scans along a sparse axis of the sensor has shown to be efficient for accurately using geo-referenced markings for localization. We have proposed a method that uses movement to produce a punctual observation that can be used to localize a crosswalk in real-time. Our accumulation takes place in the body frame and produces point clouds that can be used to detect sparse features without loss of accuracy. It has been experimentally shown to improve the localization error in the vicinity of crosswalks, correcting the localization DR drift and reducing uncertainty while maintaining consistency up to $3\sigma$. This indicates that LiDARs can provide accurate along-track measurements of road surface markings.

We have observed that incorporating an observation obtained through buffering of LiDAR scans must be done carefully. Observations produced in such a way can have correlated errors. The degree of correlation is difficult to assess as new scans are not guaranteed to change the measurement, e.g. if the new points are inside the already buffered strips. There are several ways to deal with this problem. The solution we have studied is to enlarge the evolution model noise to diminish the correlation impact. Another solution would be to use a filter capable of handling unknown correlation. The first time a crosswalk is detected, a Kalman update would be done. For all subsequent measurements on this crosswalk, a Covariance Intersection update would address the problem.
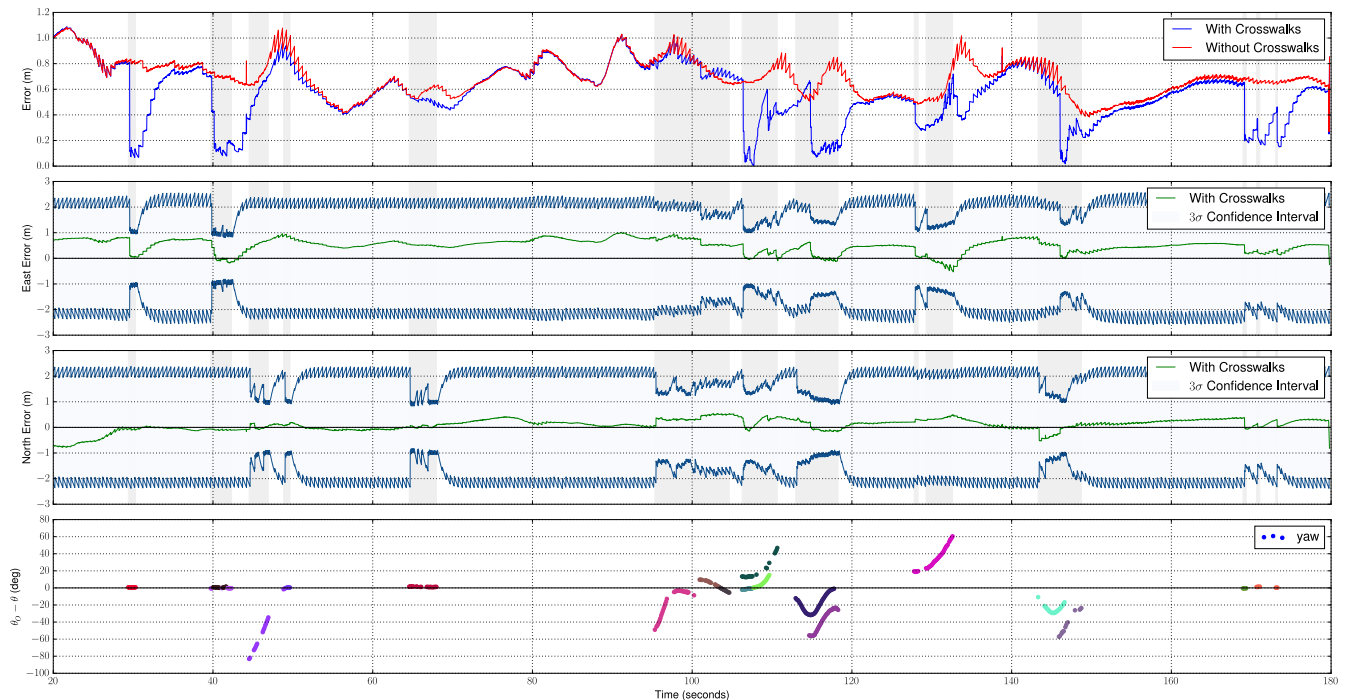
Fig. 7: At the top, absolute 2D localization error without crosswalk observations (red), and using crosswalk observations (blue). Below are the 2D error decomposition in East and North components, along with their associated $3\sigma$ confidence interval. Gray areas indicate when a crosswalk is observed. At the bottom are the angle at which crosswalks are observed relative to the vehicle, where colors represent a given crosswalk.

## REFERENCES

[1] D. Bradley and G. Roth. Adaptive Thresholding using the Integral Image. *Journal of Graphics Tools*, pages 13–21, 2007.

[2] J. Castorena and S. Agarwal. Ground-Edge-Based LIDAR Localization Without a Reflectivity Calibration for Autonomous Driving. *Robotics and Automation Letters*, pages 344–351, 2018.

[3] Y-H. Choi, T-K. Lee, and S-Y. Oh. A line feature based SLAM with low grade range sensors using geometric constraints and active exploration for mobile robot. *Autonomous Robots*, pages 13–27, 2008.

[4] M. Ester, H-P. Kriegel, J. Sander, and X. Xu. A density-based algorithm for discovering clusters in large spatial databases with noise. In *International Conference on Knowledge Discovery and Data Mining*, pages 226–231, 1996.

[5] G. Frisch, P. Xu, and E. Stawiarski. High integrity lane level localization using multiple lane markings detection and horizontal protection levels. In *International Conference on Control, Automation, Robotics and Vision*, pages 1496–1501, 2018.

[6] F. Ghallabi, F. Nashashibi, G. El-Haj-Shhade, and M-A. Mittet. LIDAR-Based Lane Marking Detection For Vehicle Positioning in an HD Map. In *International Conference on Intelligent Transportation Systems*, pages 2209–2214, 2018.

[7] E. Haines. Point in Polygon Strategies. In *Graphics Gems*, pages 24–46. 1994.

[8] A-Y. Hata and D-F. Wolf. Feature Detection for Vehicle Localization in Urban Environments Using a Multilayer LIDAR. *Transactions on Intelligent Transportation Systems*, pages 420–429, 2016.

[9] M. Kutila, P. Pyykönen, H. Holzhüter, M. Colomb, and P. Duthon. Automotive LiDAR performance verification in fog and rain. In *International Conference on Intelligent Transportation Systems*, pages 1695–1701, 2018.

[10] J. Levinson and S. Thrun. Robust vehicle localization in urban environments using probabilistic maps. In *International Conference on Robotics and Automation*, pages 4372–4378, 2010.

[11] J. Levinson and S. Thrun. Unsupervised Calibration for Multi-beam Lasers. In *Experimental Robotics*, pages 179–193. 2014.

[12] M. Magnusson. The Three-Dimensional Normal-Distributions Transform — an Efficient Representation for Registration, Surface Analysis, and Loop Detection. *PhD thesis, Örebro University*, 2009.

[13] N. Otsu. A Threshold Selection Method from Gray-Level Histograms. *Transactions on Systems, Man, and Cybernetics*, pages 62–66, 1979.

[14] A. Segal, D. Hähnel, and S. Thrun. Generalized-ICP. In *Robotics: Science and Systems*, 2009.

[15] Y. Tatebe, D. Deguch., Y. Kawanishi, I. Ide, H. Murase, and U. Sakai. Can we detect pedestrians using low-resolution lidar? integration of multi-frame point-clouds. In *International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications*, pages 157–164, 2017.

[16] Z. Wang, J. Fang, X. Dai, H. Zhang, and L. Vlacic. Intelligent vehicle self-localization based on double-layer features and multilayer LIDAR. *Transactions on Intelligent Vehicles*.

[17] A. Welte, P. Xu, and P. Bonnifait. Four-Wheeled Dead-Reckoning Model Calibration using RTS Smoothing. In *International Conference on Robotics and Automation*, pages 312–318, 2019.

[18] A. Welte, P. Xu, P. Bonnifait, and C. Zinoune. Estimating the reliability of georeferenced lane markings for map-aided localization. In *Intelligent Vehicles Symposium*, pages 1097–1103, 2019.

[19] R. Wolcott and R. Eustice. Robust LIDAR localization using multiresolution Gaussian mixture maps for autonomous driving. *The International Journal of Robotics Research*, pages 292–319, 2017.

[20] L. Yao, Q. Chen, C. Qin, H. Wu, and S. Zhang. Automatic extraction of road markings from mobile laser-point cloud using intensity data. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, pages 2113–2119, 2018.

[21] Y. Yu, J. Li, H. Guan, F. Jia, and C. Wang. Learning Hierarchical Features for Automated Extraction of Road Markings From 3-D Mobile LiDAR Point Clouds. *Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, pages 709–726, 2015.

[22] D. Zermas, I. Izzat, and N. Papanikolopoulos. Fast segmentation of 3D point clouds: A paradigm on LiDAR data for autonomous vehicle applications. In *International Conference on Robotics and Automation*, pages 5067–5073, 2017.