

Real-time Dynamic Object Detection for Autonomous Driving using Prior 3D-Maps

B Ravi Kiran¹, Luis Roldão^{1,2}, Beñat Irastorza¹, Renzo Verastegui¹, Sebastian Süss³, Senthil Yogamani⁴, Victor Talpaert¹, Alexandre Lepoutre¹, and Guillaume Trehard¹

¹ R&D Department AKKA Technologies, 78280 Guyancourt, France
 {kiran.BANGALORE-RAVI, luis.ROLDAO, benat.UGALDE, renzo.VERASTEGUI,
 victor.TALPAERT, alexandre.LEPOUTRE, guillaume.TREHARD}@akka.eu

² Robotics and Intelligent Transportation Systems (RITS) Team, INRIA Paris

³ Spleenlab, Germany

⁴ Valeo Vision Systems, Ireland

Abstract. Lidar has become an essential sensor for autonomous driving as it provides reliable depth estimation. Lidar is also the primary sensor used in building 3D maps which can be used even in the case of low-cost systems which do not use Lidar. Computation on Lidar point clouds is intensive as it requires processing of millions of points per second. Additionally there are many subsequent tasks such as clustering, detection, tracking and classification which makes real-time execution challenging. In this paper, we discuss real-time dynamic object detection algorithms which leverages previously mapped Lidar point clouds to reduce processing. The prior 3D maps provide a static background model and we formulate dynamic object detection as a background subtraction problem. Computation and modeling challenges in the mapping and online execution pipeline are described. We propose a rejection cascade architecture to subtract road regions and other 3D regions separately. We implemented an initial version of our proposed algorithm and evaluated the accuracy on CARLA simulator.

Keywords: Prior Maps, 3D obstacles, Inlier rejection

1 Introduction

Autonomous driving systems are seldom complete nowadays without a controller, motion planning and perception stack. Pre-curated maps of their environments can be further used to improve the robustness and completeness. A modern mapping system, often referred as HD Vector Maps or Prior Maps, is typically comprised of the following stages:

- **Static Map extraction:** This is an offline step called mapping where a geometric 3D model is built for the static environment.
- **Localization:** The vehicle localizes itself independently with a GPS and measures it's movement with an IMU. Alternatively the vehicle can localize itself by map aligning within the known 3D map.

- **Dynamic events representation:** Besides providing a map of the environment, modern mapping systems also provide contextual information such as: speed limits, drivable directions and space, lane markings and distance to intersections. Dynamic critical events like accidents, roadworks and lane closures are additionally provided. Some systems are able to analyze curated driver behavior profiles to predict its dynamic behavior.

The current day mapping technologies are advancing rapidly, however a formal description of the pipeline/sequence of algorithmic operations is still lacking. Companies such as Civil Maps, HERE, DeepMap, TomTom and Mobileye have various implementations of 3D maps of the environment augmented with various temporal and spatial meta-data. Moreover, crowd-sourcing to learn the environment [10] has become an important trend in modeling and updating/maintaining these maps.

We briefly describe the outline of the paper. Section 1 describes why prior 3D maps are required and their essential characteristics. Section 2 Reviews the current state of the art of Static 3D Maps and HD Maps, and the basic point cloud representations. Section 3 describes the 3D mapping pipeline while reviewing the literature regarding the various constituent steps. Section 4 describes our experiments with the CARLA simulator : creation of a evaluation framework for dynamic object extraction, creating a semantic point cloud, and finally a visual demonstration of a bounding box based point cloud inlier rejection. We conclude the paper with future work and other operations essential to the completeness of 3D Maps.

1.1 Motivation

Vehicle perception and interpretation of an unknown environment is a difficult and computationally demanding task, where any prior information that contributes to augment the information retrieved by the vehicle would aid to improve the performance. As automated driving vehicles might be commissioned to drive over closed loop routes and repeat pre-recorded paths, having a consistent parametric representation of the static environment is very useful. This is particularly the case with automated driver-less shuttles or taxi services. Current experimental driver-less prototypes commonly use globally registered 3D prior-maps with centimetric-precision. They are created using Lidar sensors and GPS+INS systems [21].

Point clouds retrieved by Lidar provide an accurate geometrical representation of the environment. However, they do not explicitly provide information about unknown areas or free space. They are also memory-intensive and lack an inherent mechanism to adapt to changes in the environment [42].

For these reasons, a pre-recorded 3D representation of the static surroundings of the environment is a useful prior that can be used for localization, obstacle detection and tracking tasks. Such representation would ideally be characterized by the following features:

Accurate 3D Model representation of the static background environment, including driving road surfaces, buildings planes/facades, lamps, roundabouts, traffic signs among others. Here, obstacles can be classified into two categories based on their stationarity across the training and test set :

1. **Stationary Objects (SO)** : Completely stationary scene components such as the road surface and buildings.
2. **Non-Stationary Static Objects (NSSO)** : Static objects that appeared or disappeared between mapping and re-localization steps.

An online background estimation and removal would help to reject a considerable portion of static points in the cloud to localize dynamic objects in the scene.

Computation Speedup : Lidar point clouds (64 Layers) usually contain around 100K points per frame, with approximately 10 frames/sec. Computation can be reduced by performing clustering of point clouds into simpler higher level models like planes. For example, plane models can be fit during mapping stage for road surfaces, building facades, etc and bounding boxes can be used for volumetric clusters estimated. This way, obstacle detection and tracking steps in the perception pipeline can be performed faster. The computational speedup is also dependent on the data structures adapted and used to store the data, we refer the readers to work done in [44].

Updatability. Apart from the cost of 3D Lidar sensors, a significant cost of a successful automated driving system goes into building and maintaining an updated 3D prior map. Surveying environments on a large scale requires keeping account of changes in the road surface, buildings, parked vehicles, local vegetation and seasonal variation, which is prohibitively expensive for a map provider. Any economic benefits of selling automated driving systems may be offset by the constant need to re-survey the environment where they operate in. Changing environments are handled in [9,23], where instead of building a single static global map, each vehicle independently creates and maintains a set of partial updates. The confidence of the parametric representation increases with the number of observations and updates done by each vehicle.

2 Prior Maps and Point Cloud Representation

2.1 Types of Maps

Mapping is one of the key pillars of automated driving. The first reliable demonstrations of automated driving by Google were primarily reliant on localization to pre-mapped areas. Because of the scale of the problem, traditional mapping techniques are augmented by semantic object detection for reliable disambiguation. In addition, localized high definition maps (HD maps) can be used as a prior for object detection. More details of mapping and usage of modern deep learning methods are discussed in [25].

Private Small Scale Maps: There are three primary reasons for the use of customized small scale maps. The first reason is privacy where it is not legally



Fig. 1. Example of High Definition (HD) map from TomTom RoadDNA (Reproduced with permission of the copyright owner)

allowed to map the area, for example, private residential area. The second reason is that HD maps still do not cover most of the areas. The third reason is the detection of dynamic structures, that may differ from global measurements. This is typically obtained by classical semi-dense point cloud maps or landmark based maps.

Large Scale HD Maps: There are two types of HD maps namely Dense Semantic Point Cloud Maps and Semantic Landmark based Maps. Semantic Landmarked based maps are an intermediate solution to dense semantic point cloud and likely to become redundant.

1. Landmark based Maps are based on semantic objects instead of generic 3D point clouds. Thus it works primarily for camera data. Companies like Mobileye and HERE follow this approach. In this method, object detection is leveraged to provide an HD map and the accuracy is improved by aggregating over several observations from different cars.
2. Dense Semantic Point Cloud Maps: This version is the best representation where all the semantics and dense point cloud are available at high accuracy. Google and TomTom follow this approach. Figure 1 demonstrates dense semantic 3D point cloud from TomTom and alignment to an image. However this approach is computationally expensive and needs large memory requirements. In this case, mapping is treated as a stronger cue than perception. If there is good alignment with the map, all the static objects (road, lanes, curb, traffic signs) are obtained from the map already and dynamic objects are obtained via background subtraction. In this work, we propose an efficient solution to obtain dynamic objects.

2.2 Review On Grid-based Representations

Grid-based representations split the space into equally sized cells in order to represent the state of specific portions of the environment. Moravec et al. [26], create a 2D grid based occupancy map with a sonar. Herbert et al. [19] proposed to extend this representation by assigning an additional variable to store at each cell the height of objects above the ground level, this approach is usually referred

as an elevation map. Moreover, a multiple surface representation was presented by [35], where several height values can be stored within the cells.

The main advantage of grid-based representations is that free, occupied and unknown space can be represented from the measurements obtained by a range sensor, usually a Lidar. These methods are commonly known as occupancy-grids, where a ray casting operation [6] is performed along each measurement in order to decrease the occupancy probability of traversed cells and increase it for the the impacted ones [34]. This is however a costly operation. Each cell is recursively updated by applying a static state binary Bayes filter as introduced in [15], where the occupancy probability of a cell v being occupied $p(v|\mathbf{z}_{t,1:n})$ is defined as:

$$p(v|\mathbf{z}_{t,1:n}) = \left[1 + \frac{1 - p(v|\mathbf{z}_{t,n})}{p(v|\mathbf{z}_{t,n})} \frac{1 - p(v|\mathbf{z}_{t,1:n-1})}{p(v|\mathbf{z}_{t,1:n-1})} \beta \right]^{-1}, \quad \beta = \frac{p(v)}{1 - p(v)} \quad (1)$$

where $\mathbf{z}_{t,1:n}$ is the complete set of sensor measurements $\{\mathbf{z}_{t,1}, \dots, \mathbf{z}_{t,n}\}$ obtained from the sensor returns at time t . The term β depends on the prior knowledge about the state of the cell v , as shown in Eq. (1). If this initial state is unknown, then $p(v) = 0.5$ and $\beta = 1$. The updates are then performed by following:

$$p(v|\mathbf{z}_{t,i}) = \begin{cases} p_{free}, & \text{the ray } \mathbf{z}_{t,i} \text{ traverses through } v \\ p_{occupied}, & \text{the ray } \mathbf{z}_{t,i} \text{ impacts within } v \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

where $p(v|\mathbf{z}_{t,i})$ corresponds to the probability update given by observation $\mathbf{z}_{t,i}$ and p_{free} and $p_{occupied}$ are the assigned probabilities when a cell is completely traversed or impacted respectively. A wide range of approaches [36,11,43] chose to perform this update using a log-odds ratio formulation L , where instead of multiplying terms from prior, past and current measurements as shown in Eq. (1), logarithmic properties allow to perform this update with simple additions:

$$L(v|\mathbf{z}_{t,1:n}) = \left(\sum_{i=1}^n L(v|\mathbf{z}_{t,i}) \right) - L(v), \quad \text{where } L(a) = \log \left(\frac{p(a)}{p(\neg a)} \right) \quad (3)$$

Similarly as in Eq. (2), log-odds probability values l_{free} and $l_{occupied}$ are defined for the update of $L(v|\mathbf{z}_{t,i})$. The advantage of this approach is that truncation issues associated with probabilities close to 0 or 1 can be avoided.

3 3D-Prior Maps Pipeline for Online Rejection

3D-Prior Maps built in various pipelines essentially contain two steps: Initially, an offline **Mapping** stage is performed, during which the computationally heavy parameter estimation is carried out. An online **Driving** stage constitutes the prediction of planar and volumetric inliers w.r.t the mapping phase estimated

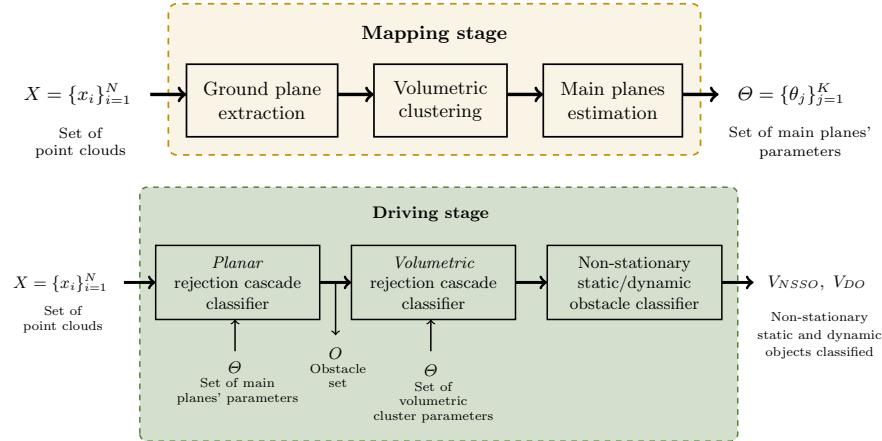


Fig. 2. A Standard clustering, Ground Plane extraction, obstacle detection and tracking pipeline in an autonomous driving pipeline.

parameters. Figure 2 demonstrates a concrete example of parameter estimation during the offline mapping stage and online model inlier rejection during the online driving stage. Similar pipelines are developed in [3] and [2]. In [33] a voxelization step is included to reduce the time complexity for the clustering/detection in the down-streams stages. Further on, we refer to the parameter estimation stage as training, and the inlier rejection/model fitting stage as testing. The point cloud dataset acquired for training is assumed to have no dynamic obstacles, while the test dataset shall contain point clouds acquired at new vehicle poses (thus incurring errors in GPS localization and vehicle orientation) with a set of dynamic objects observable in each frame.

The obstacle detection pipeline can be then considered as a binary classification problem, with the background being constituted by the ground and building facades planes as well as volumetric static obstacles (vegetation, parked cars among others), while the foreground is formed by the dynamic obstacles not belonging to the background. Since the parametric estimation for geometric and volumetric obstacles is performed on a training set containing no moving obstacles, while prediction is performed on unseen test set, one could detect non-stationary static obstacles (NSSO) in the test set which have appeared/disappeared in between the mapping and driving steps.

We categorize the 3D object mapping problem in two categories : First is to consider the detection of 3D obstacles on prior maps as a stationary background estimation problem, which is well known in the background subtraction community. The second is to separate the background objects in geometrical and volumetric structures such as planes, road surfaces and dense clusters. We shall explore the second approach in our paper. The key steps performed in this approach are:

1. **Global Coordinate Point Cloud Registration:** Each point cloud frame produced by the vehicles across different poses is registered into a common global origin, yielding a single point cloud representation of the entire environment. This is usually achieved by applying rigid transformations to each point cloud frame with the aid of GPS/INS and Inertial Measurement Systems (IMU). Registration algorithms are also applied for this purpose, where Iterative Closest Point (ICP) algorithm is the most popular approach [4].
2. **Ground Surface Approximation:** Surface reconstruction of roads and normal estimation to obtain their local orientation. Planar ground assumptions are usually considered to estimate a planar surface parameters. Non planar surfaces can be approximated with several planes as in [3].
3. **Planar surface extraction:** Extract other dominant planes present in the geometric representation, commonly linked to building facades.
4. **Clustering and Classification:** Clustering of the remaining volumetric features (vegetation, parked vehicles, lamp posts, etc.) in the global map coordinate frame. The classification step is usually performed over the clusters to obtain a class for each cluster : Vehicles, Buildings, Vegetation and others. This provides semantic information regarding the type of obstacle.
5. **Long-term occupancy probability estimation:** In this step the environment is classified in the mentioned classes: Stationary Objects (SO), Non-Stationary Static Objects (NSSO) and dynamic obstacles. The precision of the classification will depend on the amount of gathered data. NSSO clusters are removed or updated based on user intervention. For example, a construction site might have appeared during the mapping stage, while disappeared during the online driving stage. In such a case, the pipeline classifies the empty construction site as a dynamic obstacle.

3.1 Ground surface approximation

Road surface extraction allows the vehicle to detect the drivable space while removing a considerable proportion of points for the subsequent obstacle detection and classification steps. Random sampling methods (RANSAC), normals estimation or Hough based methods are among the most popular approaches for this task.

Recently, Chen et al. [8] build a depth image by using the Lidar spherical coordinates, where each row (i.e., each fixed azimuth value ϕ) corresponds to a given laser layer of the sensor. The columns of the image correspond to the horizontal polar angles (θ) and the image intensities are represented by the radial distances (d). An example of the depth image can be seen at Fig. 6 top left. Following the assumption that for a given sensor layer (a given ϕ) all the points in the road will be at a same distance from the sensor along the x axis, the v -disparity of a given pixel (p, q) is calculated as:

$$\Delta_{p,q} = \frac{1}{x(p,q)} = \frac{1}{d_{p,q} \cos(\phi_{p,q}) \cos(\theta_{p,q})} \quad (4)$$

Since all road points in a given row q will share a similar value x , they will fall into a same box in the histogram and, assuming the road is flat, the high-intensity values in the histogram will draw a line (since the value of $\Delta_{p,q}$ will change linearly as we traverse the values of φ). Aside from road plane parameter estimation, this method can also be used for online obstacle detection, due to the lightness of its computations. All objects placed on the road will be closer than the road to the sensor, and will therefore have a higher disparity value. Thus, they will be located above the line that corresponds to the road in the histogram (see fig. 6).

3.2 Static Background Estimation in 3D-Environments

Plane extraction : RANSAC based methods A standard plane extraction algorithm is the Random sample consensus (RANSAC) method. This algorithm aims to find a plane in the scene, where the resulting plane should "contain" the most possible points with respect to a threshold distance from plane to point. In the self-driving context, the road is therefore assimilated to a plane and we use a large threshold to accommodate for any potential curvature. Fast Ground Segmentation extraction methods on Lidar Data with the use of Squeeze-net Architectures are capable of real-time prediction performances [37].

Geometric approximation: Plane extraction methods can be categorized into Hough-based, region growing, or RANSAC-based approaches. Authors [17] describe a novel Hough-based voting scheme for finding planes in 3D point clouds. They assume an important geometrical property of the Lidar, i.e. planes in the 3D intersect with Lidar rays at a fixed azimuth angle ϕ in conic sections of varying curvature dependent upon the inclination angle of the Lidar and the relative orientation of the plane. A fast approximation of regions where for a 1D scan $d_\phi(\theta)$ smoothly varying segments are extracted, followed by accumulator framework similar to the one in randomized Hough transform is utilized to vote the planes that were extracted. Background Subtraction based approaches have been studied in accessibility analysis using Evidence Grids [1].

3.3 Review on point cloud representations for Mapping

Here we shortly review the different classes of methods to evaluate a static 3D Map. Readers are directed towards a recent comprehensive review on road object extraction using laser point cloud data by authors [22]. **Multi-resolution and multi-scale methods:** In [40] authors model each cell of an occupancy grid with a multi-resolution Gaussian Mixture models (GMMs) for obstacle mapping as well as localization in prior maps. In [13], a dynamic method to adequate the resolution of the grid mapping-cell sizes by using an octree structure is presented. This could be applied on the prior maps to help compress the amount of data and to improve the obstacles detection performance.

Compressing Information : Authors in [27] provide a principled study on how to compress map information for autonomous exploration robots. To reduce the

computational cost of exploration, they develop an information theoretic strategy for simplifying a robots representation of the environment, in turn allowing information-based reward to be evaluated more efficiently. To remain effective for exploration, their strategy must adapt the environment model in a way that sacrifices a minimal amount of information about expected future sensor measurements.

Dynamic Obstacle Classification methods: We shortly describe methods that do no rely on estimating an accurate 3D obstacle prior map, but instead extract features that classify the scene into foreground and background. Authors in [36] create a logistic classifier is built on a binary feature map created from the voxelized point cloud grid. The training of the binary classifier is performed on the KITTI dataset with positive class for dynamic obstacles given by the bounding box tracklets, while the background is any window not containing a dynamic obstacle. [32] use the region proposal network from YOLO to perform bounding box location and orientation prediction. MODnet [31] learns to directly extract moving objects using motion and appearance features and could be used to augment existing dynamic object extraction methods.

Multi-model approaches: In [28], multi-modal background subtraction is performed by classifying image data as background/foreground. This is done by employing a per-pixel Gaussian mixture model estimated on a background without dynamic obstacles. The lidar point clouds are reprojected into the camera's field of view and consequently masked as foreground or background by applying the pre-existing correspondence between pixels and lidar-points. This methodology provides a direct bridge to a family of background subtraction methods [5] that could now be applied on the image data to perform binary classification of lidar points.

Authors in [41] utilize a lidar based mapping stage that not only use the Lidar point clouds coordinates x, y, z , but also the reflectance α . They localize a monocular camera during the driving stage within a 3D prior map by evaluating the normalized mutual information between the illuminance with surface reflectivities obtained during the mapping stage, and localize themselves given an initial pose belief.

Deep Learning based methods Authors in [39] propose a deep learning based system for fusing multiple sensors, i.e. RGB images, customer-grade GPS/IMU, and 3D semantic maps, which improves the robustness and accuracy for camera localization and scene parsing. Authors also propose the Zpark dataset along with a study that includes dense 3D semantically labeled point clouds, ground truth camera poses and pixel-level semantic labels of video camera images, which will be released in order to benefit related researches.

Terrestrial point clouds provide a larger field of view as well as point cloud density. 3D point cloud semantic segmentation produced with terrestrial scanners have been evaluated by authors in [18]. The 3D prior maps mapping architecture could also be augmented by aligning ground scanned point clouds with Airborne point clouds. Authors in [14] register aerially scanned point cloud and local ground scanned point clouds within Airborne point clouds using Deep

stacked Autoencoders. Authors in [30] propose a full convolutional volumetric autoencoder that learns volumetric representation from noisy data by estimating the voxel occupancy grids.

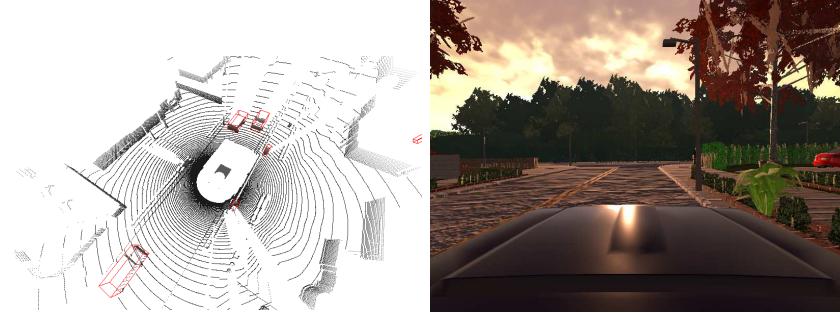


Fig. 3. Figure (left) shows tracklets extracted from CARLA simulator. Figure (right) shows the corresponding view of camera.

4 Experiments

We use the CARLA simulated cities with dynamic objects as our test framework. During the **Mapping stage**, point clouds at each vehicle pose (GPS coordinate with vehicle orientations yaw, pitch and roll) are obtained from the CARLA simulator. In the driving stage, we perform three distinct steps : Firstly, we extract a parametric representation of the road, namely the RANSAC and the Lidar Histogram as described before. Secondly, we cluster the current frame with the hierarchical DBSCAN clustering algorithm [24,7]. The third step involves the estimation of the final rejection cascade model using robust bounding boxes for planar clusters, while minimum volume bounding boxes for volumetric clusters.

4.1 CARLA Setup

The CARLA simulator [12] has recently created a lidar point cloud generator that enables users to parametrise the simulated sensor with different parameters such as its angular resolution, number of layers among others. Virtual KITTI based point clouds are now made available by authors in [16]. On the other hand, [45] now provide simulated Lidar point clouds and object labels, aiming to augment the datasets to be used for supervised learning using deep convolutional networks.

In our study, we use the CARLA simulated lidar point cloud stream, with the localization and IMU information generated for the simulated vehicles to map the environment. CARLA provides a reliable environment where the performance of the dynamic objects detection can be tested, since the user has access to the

location, speed and bounding boxes of all agents in the scene. Figure 3 shows an example of the object tracklets with the associated camera image. The semantic segmentation of the scene is also provided.

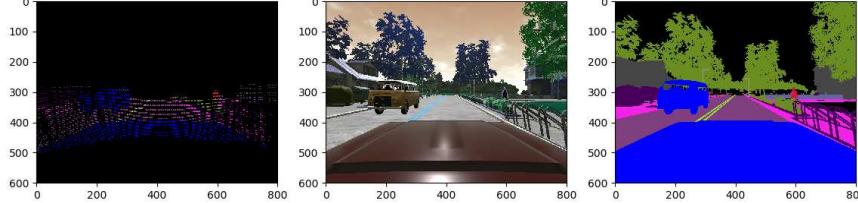


Fig. 4. Left to right : Pointcloud with labels in camera view, image from camera, semantic segmentation of the scene. Mapping the semantic segmentation labels to 3D point cloud from the Lidar. These initial results point towards the possibility of obtaining a completely labeled semantic point cloud for downstream classification tasks and parameter cross-validation purposes.

4.2 Clustering and Road Extraction

The road extraction in CARLA is performed using a frame based RANSAC plane estimation. The Lidar Histogram has been observed to be useful in cases where the road surface is not a plane and contains positive and negative slopes.

The clustering process involves two steps : An initial frame-based clustering is achieved using the HDBSCAN algorithm [7], as shown in figure 5, and a subsequent clustering of the centroids of each frame's cluster, to obtain a globally consistent cluster label. This second clustering step, termed as super-clustering is performed again using the HDBSCAN. Our primary motivation in using HDBSCAN is to handle the variable density of point-clouds due to the geometrical projection produced by the Lidar sensor.

After each step of the frame based clustering a the labelled point cloud is saved back for each frame. Once all frames have been labelled, a global association step is performed in the final step when the centroids across all frames are clustered to obtain a consistently labelled point cloud. In the future, a hierarchical data structure could alleviate this two-level clustering task.

4.3 Alternative road extraction with Lidar-histogram

In this section we show the result of road segmentation on readings from the KITTI dataset, results are shown in figures 6. CARLA simulated roads in the default environments are flat, while KITTI provides roads with changes in gradient.

Since the Velodyne points in this dataset are given in Cartesian coordinates, they were transformed to spherical ones first. We used a naive approach to

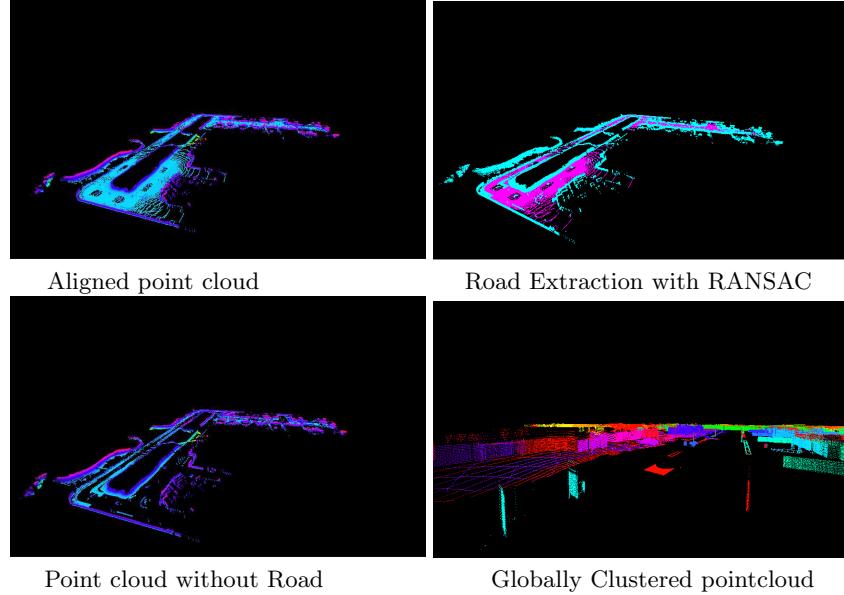


Fig. 5. Joined cloud in global coordinates, Road extraction using RANSAC, Road removed point cloud, Clustered point cloud using the hierarchical DBSCAN (HDBSCAN) on each frame followed super-clustering step on centroids from the previous step. [24].

approximate the line representing the road, selecting a desired number of the most intense values and fitting a line with the least squared error. Although this proved to be precise enough, the line could also be fitted using a RANSAC-based method or even Hough transforms. In order to get the road points from the line, both the disparity and x values were computed for each row and all values within a threshold were accepted. The pixels corresponding to the road have correctly been identified, as can be seen in both figures, without overriding the obstacles present in the lane. If desired, the classified point clouds in Cartesian coordinates could then be retrieved by putting the points in spherical coordinates back to Cartesian ones.

4.4 Planar and Volumetric cluster classification

In our third step after road extraction using RANSAC and frame based clustering using HDBSCAN, we describe briefly how we obtain an inlier rejection bounding box for such clusters. The planar and volumetric clusters are decided by a threshold on average planarity feature. The planarity is obtained from eigenvalues of the 3D structure tensor given by $P_\lambda = (\lambda_2 - \lambda_3)/\lambda_1 \in [0, 1]$, for each point in the point cloud. We show two types of clusters with their respective bounding boxes in figure 7.

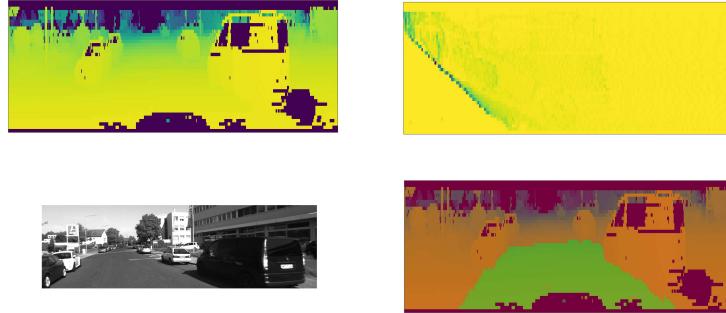


Fig. 6. Example of the LIDAR-histogram method at work, applied to a frame from the KITTI dataset (drive 9, frame number 228). The pictures above show the depth image (left) and histogram (right). The images below are the original picture and the result of the the classification of the pixels (the ones corresponding to the road in green).

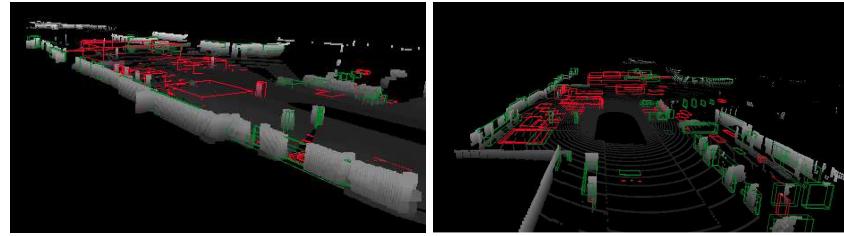


Fig. 7. A prior 3D map constituting of two classes of bounding boxes : Robust planar bounding boxes for planes (green), and volumetric minimum volume bounding box (red). In future studies we plan to use these bounding boxes to achieve a rejection cascade.

5 Conclusions

3D-Prior maps for obstacle detection has become a key engineering problem in today's autonomous driving system, moving a large part of the detection problem into the construction of precise 3D probabilistic descriptions of the environment. In this study, we reviewed the basic steps involved in the construction of prior maps : Road extraction, clustering and subsequent plane extraction, super-clustering for globally consistent clusters.

We demonstrated the results of applying Lidar Histogram to various tracks in the KITTI dataset. We showed it to be an efficient parametric representation of the road surface, while modeling positive and negative (holes) obstacles. It also provides the decision making system with the estimation of free space around dynamic obstacles. Decomposing the road segmentation algorithm into off-line

mapping and online driving stages is key to obtain real-time performance while not conceding accuracy.

Estimating the frequency of occurrence of different geometrical and volumetric features enables us to envisage an efficient implementation following a rejection cascade employed first by [38]. Our implementation follows background subtraction based cascade as developed in [20].

Furthermore maintaining the temporal relevance of obstacles in 3D-Maps is a costly and an important question. Long-term mapping solutions that assume a semi-static environment are now in development [29]. Real-time implementations should include a third phase aside Mapping and Driving, that is an efficient change-point detection (in the point clouds associated with non-stationary static obstacles) and parameter update mechanism.

References

1. Anderson-Sprecher, P., Simmons, R., Huber, D.: Background subtraction and accessibility analysis in evidence grids. In: Robotics and Automation (ICRA), 2011 IEEE International Conference on. pp. 3104–3110. IEEE (2011)
2. Asvadi, A., Peixoto, P., Nunes, U.: Two-stage static/dynamic environment modeling using voxel representation. In: Robot 2015: Second Iberian Robotics Conference. pp. 465–476. Springer (2016)
3. Asvadi, A., Premebida, C., Peixoto, P., Nunes, U.: 3d lidar-based static and moving obstacle detection in driving environments: An approach based on voxels and multi-region ground planes. *Robotics and Autonomous Systems* **83**, 299–311 (2016)
4. Bellekens, B., Spruyt, V., Berkvens, R., Weyn, M.: A survey of rigid 3d pointcloud registration algorithms. In: AMBIENT 2014: the Fourth International Conference on Ambient Computing, Applications, Services and Technologies, August 24-28, 2014, Rome, Italy. pp. 8–13 (2014)
5. Bouwmans, T.: Traditional and recent approaches in background modeling for foreground detection: An overview. *Computer Science Review* **11**, 31–66 (2014)
6. Bresenham, J.E.: Algorithm for computer control of a digital plotter. *IBM Systems Journal* **4**(1), 25–30 (1965)
7. Campello, R.J.G.B., Moulavi, D., Zimek, A., Sander, J.: Hierarchical density estimates for data clustering, visualization, and outlier detection. *ACM Trans. Knowl. Discov. Data* **10**(1) (July 2015)
8. Chen, L., Yang, J., Kong, H.: Lidar-histogram for fast road and obstacle detection. In: Robotics and Automation (ICRA), 2017 IEEE International Conference on. pp. 1343–1348. IEEE (2017)
9. Churchill, W., Newman, P.: Practice makes perfect? managing and leveraging visual experiences for lifelong navigation. In: Robotics and Automation (ICRA), 2012 IEEE International Conference on. pp. 4525–4532. IEEE (2012)
10. Dabeer, O., Ding, W., Gowaiker, R., Grzechnik, S.K., Lakshman, M.J., Lee, S., Reitmayr, G., Sharma, A., Somasundaram, K., Sukhavasi, R.T., et al.: An end-to-end system for crowdsourced 3d maps for autonomous vehicles: The mapping component. In: Intelligent Robots and Systems (IROS), 2017 IEEE/RSJ International Conference on. pp. 634–641. IEEE (2017)
11. Doherty, K., Wang, J., Englot, B.: Bayesian generalized kernel inference for occupancy map prediction. 2017 IEEE International Conference on Robotics and Automation (ICRA) pp. 3118–3124 (2017)

12. Dosovitskiy, A., Ros, G., Codevilla, F., Lopez, A., Koltun, V.: CARLA: An open urban driving simulator. In: Proceedings of the 1st Annual Conference on Robot Learning. pp. 1–16 (2017)
13. Einhorn, E., Schröter, C., Gross, H.M.: Finding the adequate resolution for grid mapping-cell sizes locally adapting on-the-fly. In: Robotics and Automation (ICRA), 2011 IEEE International Conference on. pp. 1843–1848. IEEE (2011)
14. Elbaz, G., Avraham, T., Fischer, A.: 3d point cloud registration for localization using a deep neural network auto-encoder. In: Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference on. pp. 2472–2481. IEEE (2017)
15. Elfes, A.: Using occupancy grids for mobile robot perception and navigation. Computer **22**(6), 46–57 (June 1989)
16. Francis, E., Theodora, K., Alexander, H., Bastian, L.: Exploring spatial context for 3d semantic segmentation of point clouds. In: IEEE International Conference on Computer Vision, 3DRMS Workshop, ICCV (2017)
17. Grant, W.S., Voorhies, R.C., Itti, L.: Finding planes in lidar point clouds for real-time registration. In: Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on. pp. 4347–4354. IEEE (2013)
18. Hackel, T., Savinov, N., Ladicky, L., Wegner, J.D., Schindler, K., Pollefeys, M.: Semantic3d.net: A new large-scale point cloud classification benchmark. In: ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences. vol. IV-1-W1, pp. 91–98 (2017)
19. Herbert, M., Caillas, C., Krotkov, E., Kweon, I.S., Kanade, T.: Terrain mapping for a roving planetary explorer. In: Proceedings, 1989 International Conference on Robotics and Automation. pp. 997–1002 vol.2 (May 1989)
20. Kiran, B.R., Yogamani, S.: Real-time background subtraction using adaptive sampling and cascade of gaussians. arXiv preprint arXiv:1705.09339 (2017)
21. Levinson, J., Askeland, J., Becker, J., Dolson, J., Held, D., Kammel, S., Kolter, J.Z., Langer, D., Pink, O., Pratt, V., et al.: Towards fully autonomous driving: Systems and algorithms. In: Intelligent Vehicles Symposium (IV), 2011 IEEE. pp. 163–168. IEEE (2011)
22. Ma, L., Li, Y., Li, J., Wang, C., Wang, R., Chapman, M.A.: Mobile laser scanned point-clouds for road object detection and extraction: A review. Remote Sensing **10**(10) (2018)
23. Maddern, W., Pascoe, G., Newman, P.: Leveraging experience for large-scale lidar localisation in changing cities. In: Robotics and Automation (ICRA), 2015 IEEE International Conference on. pp. 1684–1691. IEEE (2015)
24. McInnes, L., Healy, J., Astels, S.: hdbscan: Hierarchical density based clustering. The Journal of Open Source Software **2**(11) (mar 2017). <https://doi.org/10.21105/joss.00205>, <https://doi.org/10.21105%2Fjoss.00205>
25. Milz, S., Arbeiter, G., Witt, C., Abdallah, B., Yogamani, S.: Visual slam for automated driving: Exploring the applications of deep learning. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops. pp. 247–257 (2018)
26. Moravec, H., Elfes, A.E.: High resolution maps from wide angle sonar. In: Proceedings of the 1985 IEEE International Conference on Robotics and Automation. pp. 116 – 121 (March 1985)
27. Nelson, E.A.: Environment model adaptation for autonomous exploration. Robotics Institute, Pittsburgh, PA, Tech. Rep. CMU-RI-TR-15-12 (2015)
28. Ortega Jiménez, A.A., Andrade-Cetto, J.: Segmentation of dynamic objects from laser data. In: ECMR 5th European Conference on Mobile Robots. pp. 115–121 (2011)

29. Rosen, D.M., Mason, J., Leonard, J.J.: Towards lifelong feature-based mapping in semi-static environments. In: *Robotics and Automation (ICRA), 2016 IEEE International Conference on*. pp. 1063–1070. IEEE (2016)
30. Sharma, A., Grau, O., Fritz, M.: Vconv-dae: Deep volumetric shape learning without object labels. In: *European Conference on Computer Vision*. pp. 236–250. Springer (2016)
31. Siam, M., Mahgoub, H., Zahran, M., Yogamani, S., Jagersand, M., El-Sallab, A.: Modnet: Moving object detection network with motion and appearance for autonomous driving. *arXiv preprint arXiv:1709.04821* (2017)
32. Simon, M., Milz, S., Amende, K., Gross, H.M.: Complex-yolo: Real-time 3d object detection on point clouds. *arXiv preprint arXiv:1803.06199* (2018)
33. Sixta, T.: LIDAR Based Sequential Registration and Mapping for Autonomous Vehicles. Master’s thesis, Czech technical university in Prague (2017)
34. Thrun, S., Burgard, W., Fox, D.: *Probabilistic Robotics (Intelligent Robotics and Autonomous Agents)*. The MIT Press (2005)
35. Triebel, R., Pfaff, P., Burgard, W.: Multi-level surface maps for outdoor terrain mapping and loop closing. In: *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*. pp. 2276–2282 (Oct 2006)
36. Ushani, A.K., Wolcott, R.W., Walls, J.M., Eustice, R.M.: A learning approach for real-time temporal scene flow estimation from lidar data. In: *Robotics and Automation (ICRA), 2017 IEEE International Conference on*. pp. 5666–5673. IEEE (2017)
37. Velas, M., Spanel, M., Hradis, M., Herout, A.: Cnn for very fast ground segmentation in velodyne lidar data. In: *Autonomous Robot Systems and Competitions (ICARSC), 2018 IEEE International Conference on*. pp. 97–103. IEEE (2018)
38. Viola, P., Jones, M.: Rapid object detection using a boosted cascade of simple features. In: *Computer Vision and Pattern Recognition, 2001. CVPR 2001. Proceedings of the 2001 IEEE Computer Society Conference on*. vol. 1, pp. I–I. IEEE (2001)
39. Wang, P., Yang, R., Cao, B., Xu, W., Lin, Y.: Dels-3d: Deep localization and segmentation with a 3d semantic map. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 5860–5869 (2018)
40. Wolcott, R.W.: Robust localization in 3d prior maps for autonomous driving. (2016)
41. Wolcott, R.W., Eustice, R.M.: Visual localization within lidar maps for automated urban driving. In: *Intelligent Robots and Systems (IROS 2014), 2014 IEEE/RSJ International Conference on*. pp. 176–183. IEEE (2014)
42. Wurm, K.M., Hennes, D., Holz, D., Rusu, R.B., Stachniss, C., Konolige, K., Burgard, W.: Hierarchies of octrees for efficient 3d mapping. In: *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on*. pp. 4249–4255. IEEE (2011)
43. Wurm, K.M., Hornung, A., Bennewitz, M., Stachniss, C., Burgard, W.: Octomap: A probabilistic, flexible, and compact 3d map representation for robotic systems. In: *In Proc. of the ICRA 2010 workshop* (2010)
44. Yin, H., Berger, C.: Mastering data complexity for autonomous driving with adaptive point clouds for urban environments. In: *Intelligent Vehicles Symposium (IV), 2017 IEEE*. pp. 1364–1371. IEEE (2017)
45. Yue, X., Wu, B., Seshia, S.A., Keutzer, K., Sangiovanni-Vincentelli, A.L.: A lidar point cloud generator: from a virtual world to autonomous driving. In: *Proceedings of the 2018 ACM on International Conference on Multimedia Retrieval*. pp. 458–464. ACM (2018)