

High-level camera-LiDAR fusion for 3D object detection with machine learning

Gustavo A. Salazar-Gomez *
Department of Automatics and Electronics
Universidad Autónoma de Occidente
gustavo.salazar@uao.edu.co

Miguel A. Saavedra-Ruiz *
Department of Automatics and Electronics
Universidad Autónoma de Occidente
miguel.saavedra@uao.edu.co

Victor A. Romero-Cano
Department of Automatics and Electronics
Universidad Autónoma de Occidente
varomero@uao.edu.co

Abstract

This paper tackles the 3D object detection problem, which is of vital importance for applications such as autonomous driving. Our framework uses a Machine Learning (ML) pipeline on a combination of monocular camera and LiDAR data to detect vehicles in the surrounding 3D space of a moving platform. It uses frustum region proposals generated by State-Of-The-Art (SOTA) 2D object detectors to segment LiDAR point clouds into point clusters which represent potentially individual objects. We evaluate the performance of classical ML algorithms as part of an holistic pipeline for estimating the parameters of 3D bounding boxes which surround the vehicles around the moving platform. Our results demonstrate an efficient and accurate inference on a validation set, achieving an overall accuracy of 87.1%.

1. Introduction

Over the preceding years self-driving vehicles have received attention among the research community as a result of their potential of improving mobility, safety and reliability of transportation systems [1]. However, one of the core capabilities needed to unveil the complete potential of self-driving vehicles is the ability to perceive the objects surrounding it in the 3D space [2].

3D object detection allows autonomous agents to estimate the relative pose of multiple objects neighbouring an ego-vehicle. Modern Deep Learning (DL) methods have been extensively wielded to address this issue. Some of the most common methods work directly over point clouds with

convoluted deep neural network architectures [3, 4] or by creating a frustum region proposal, traditionally employing a RGB camera and a depth sensor [1, 5, 6, 7]. Notwithstanding the astonishing results presented by these models, their implementation is usually occluded by the vast need of computational resources required to deploy them [8]. Furthermore, substantial amounts of labeled datasets such as nuScenes [9] are needed to obtain acceptable accuracy levels.

2. Research problem and motivation

In this paper we present a framework to address the mentioned issues by combining SOTA deep learning algorithms for 2D detection with low-complexity, classical ML algorithms. Particularly, we show how these techniques can leverage the use of camera and LiDAR information to create a frustum region proposal [5] and deliver 3D object detections with few data samples in real-time. Classic ML algorithms have been exploited to resolve unsupervised learning problems like clustering sparse point clouds [10, 11] or supervised learning ones as pose parameter regression [12]. Nevertheless, there has not been significant research efforts towards the employment of these techniques into an high-level camera-LiDAR fusion for 3D object detection system.

The aforementioned considerations thoroughly motivate the conceptualization of this work. This paper is driven by the following research question: How to develop a ML pipeline to detect vehicles in the 3D space leveraging the utilization of mature 2D object detectors, using camera-LiDAR data, to estimate 3D bounding boxes in real-time for self-driving applications?

*Equal contribution

Set	x	y	z	ψ	w	l	h	Avg.	Avg. 3D	Avg. BEV
Training	98.8	98.2	99.9	78.0	96.8	94.2	99.8	95.1	62.0	68.4
Test	96.6	97.8	95.4	55.7	80.7	88.4	95.0	87.1	42.7	47.8

Table 1: Evaluation metrics of the proposed method.

3. Technical contribution

In order to address this problem, we wield a set of different classic ML algorithms to estimate the 3D bounding box parameters of a given vehicle. Initially, we adopt a similar approach as the one proposed in [5], where a frustum region proposal is assembled, taking advantage of SOTA 2D object detectors [13] reason by which we will only focus on the subsequent steps. Subsequently, the point cloud instance inside the frustum proposal is segmented using the DBSCAN [14] algorithm. Finally, a global feature representation encoding the relevant information of the given segmented instance is used as input for a Support Vector Regressor (SVR) [15] to estimate the 3D bounding box parameters. The final goal is to estimate the x, y, z centroid coordinates, the box dimensions w, l, h and its heading ψ .

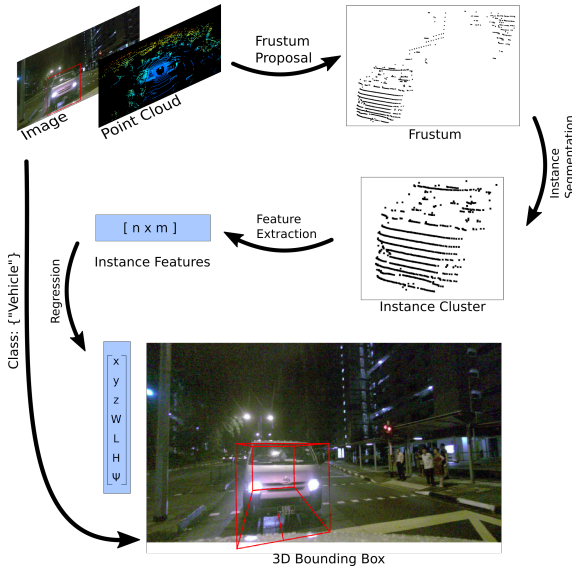


Figure 1: Proposed framework for 3D object detection.

Other algorithms were tested like Random Forest [16], XGBoost [17] for regression and KMeans [18] for clustering. Over different experiments, the SVR algorithm exhibited the best results in terms of accuracy at a similar computational cost and less parameterization as the other regression algorithms. Likewise, KMeans constrained the clustering stage of the process as a specific number of cluster had to be set regardless that a scene may have different objects

within. For the sake of space, only the assessment of the pipeline with SVR and DBSCAN is presented in this work.

Our proposed framework is presented in Fig 1 and it follows the steps previously mentioned. The nuScenes dataset is composed of 1000 scenes of 20 seconds each. In this work we employed a small version of this dataset called “Mini” which contains a total of 10 scenes [9]. Our model was trained and tested using a total of 1420 image samples with the original image size provided in the dataset of 1600×900 pixels. All the images used have at least one vehicle within and the dataset was split in 80% for training and 20% for testing. To assess performance, the 3D Intersection Over Union (IoU) was used to measure the percentage of volume intersected between the predicted bounding box and the ground truth.

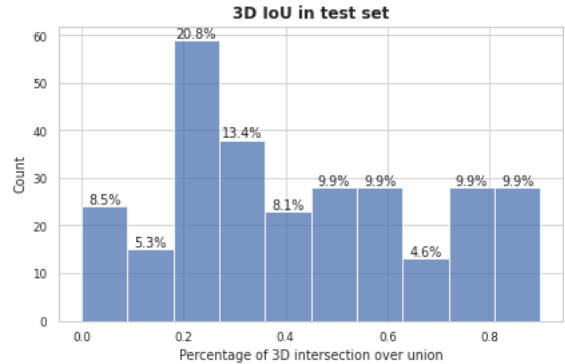


Figure 2: 3D intersection over union results of ground truth versus predictions for the test set.

As shown in Fig 2 approximately the 44% of predictions have an IoU above the 50%. The thorough assessment of the 3D IoU and Birds-Eye-View (BEV) metrics can be seen in Table 1.

To obtain a thoroughly evaluation and establish which parameters are affecting the 3D IoU score, the accuracy of each one is presented in Table 1. From there it is possible to see that the proposed framework is capable of accurately estimating the centroid coordinates x, y, z and the bounding box dimensions w, l, h except for the width which presents an accuracy of approximately 80% in the test set. The low performance achieved in ψ is due to the difficulty estimating orientations with supervised learning techniques.

Module	Instance Segmentation	Feature Extraction	Regression	Total
Training	$11.1s \pm 2.3s$	$19.1s \pm 0.5s$	$2.6s \pm 8.3ms$	$32.8s \pm 2.35s$
Inference	$4.7ms \pm 4.8ms$	$13ms \pm 1.4ms$	$0.7ms \pm 1.2ms$	$18.4ms \pm 5.1ms$

Table 2: Processing times through training and inference stages per module.

Additionally, in Table 1 is shown how with few data samples our proposed framework¹ is capable to achieve an overall accuracy of 87.1% for the validation set with an average inference time of $18.4ms$ per image and point cloud pair, using a 3.20 GHz CPU for training and inference. The aforementioned results validates how using a small number of samples from the original dataset, classical ML algorithms are capable to produce promising results with limited data and computational results. In fact, deep learning based SOTA algorithms are generally trained with huge datasets as the nuScenes which is composed of approximately 1.4 million of images and 390 thousand LiDAR sweeps in order to produce accurate results [9].

To assess the processing time of our high-level architecture, Table 2 presents the training and inference time within each stage of the process. There it is possible to notice how our method is capable to train the whole system with 1136 images and LiDAR sweeps in roughly $32.8s$ and process a new data sample in approximately $18.4ms$ or 55FPS in a CPU-only setup. Compared with SOTA methods such like [19] which has 11FPS using a Titan RTX GPU, our method excels in processing time using limited amounts of data and computational resources.

Method	Cam	Rad	LiDAR	mASE	mAOE
CenterFusion	✓	✓	-	0.142	0.085
Ours	✓	-	✓	0.573	0.840

Table 3: Performance comparison with Baseline for 3D object detection on nuScenes Dataset for the Car category.

We use the CenterFusion [20] model as baseline, in order to compare our results. This was submitted for nuScenes detection challenge in CVPR 20, and represents the SOTA of models that rely in frustum region to perform 3D object detection, what makes it the most suitable method to compare with. The authors stated that the model was trained using nuScenes dataset where two Nvidia P5000 GPUs were employed and the images size were reduced to 800×450 pixels in order to increase computational speed. The evaluation metrics used in the comparison are; Average Scale Error (ASE) that is calculated as $1 - \text{IoU}$ after aligning cen-

¹https://github.com/MikeS96/3d_obj_detection

ters and orientation, and Average Orientation Error (AOE) that is the smallest yaw angle difference between prediction and ground-truth in radians, where values close to zero are better. The results used for this assessment were the ones provided by the authors in the original paper [20].

Table 3 presents the results for the Car category. It shows whether the methods use radar or LiDAR for measuring depth. CenterFusion obtained 0.142 on ASE metric and compared with our proposed method which obtained an ASE of 0.573, we found that this difference may part from the robustness of the first one and enhancements such as Data Augmentation. Furthermore, our model is based on classical ML algorithms which makes the models less complex and faster to compute, necessary assets if working on scarce computational resources.

By comparing the AOE metrics we found a significant contrast with the baseline that shows 0.085 and ours that obtained 0.840, due to our proposed method does not use techniques such as MultiBin [21] architecture for orientation estimation, commonly used to solve this task.

Although the provided results do not surpass the current SOTA methods in 3D object detection in terms of accuracy, the given performance can be considered positive in scenarios where there are few amounts of data and the computational resources are limited. Additionally, comparing the results of our framework with the baseline, it is possible to notice that our models is capable to train and perform inference using only a CPU whereas the baseline relies on the use of expensive GPUs, reason by which our high-level framework could be deployed within low-cost settings. With improvements such as the implementation of MultiBin for heading estimation, the evaluation metrics could be considerable boosted toward better results.

4. Conclusion

In this work, we propose a framework that is capable of predicting 3D bounding boxes for vehicles and shows promising results estimating its parameters using classic ML techniques. Comparing our method with CenterFusion, it is possible to notice that even if the accuracy does not surpass the baseline's, the performance in training and inference stages can be considered positive, and allows our framework to be deployed within low-cost settings and its use in real-time applications.

References

- [1] M. Liang, B. Yang, Y. Chen, R. Hu, and R. Urtasun. Multi-task multi-sensor fusion for 3d object detection. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7337–7345, 2019.
- [2] Victor Romero-Cano, Nicolas Vignard, Christian Laugier. XDvision: Dense outdoor perception for autonomous vehicles. In *IEEE Intelligent Vehicles Symposium (IV)*, pages 752–757, 2017.
- [3] Charles Ruizhongtai Qi, Hao Su, Kaichun Mo, and Leonidas J. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. *CoRR*, abs/1612.00593, 2016.
- [4] Bin Yang, Wenjie Luo, and Raquel Urtasun. PIXOR: real-time 3d object detection from point clouds. *CoRR*, abs/1902.06326, 2019.
- [5] Charles Ruizhongtai Qi, Wei Liu, Chenxia Wu, Hao Su, and Leonidas J. Guibas. Frustum pointnets for 3d object detection from RGB-D data. *CoRR*, abs/1711.08488, 2017.
- [6] Xiaoke Shen and Ioannis Stamos. Frustum voxnet for 3d object detection from rgb-d or depth images, 2020.
- [7] P. Cao, H. Chen, Y. Zhang, and G. Wang. Multi-view frustum pointnet for object detection in autonomous driving. In *2019 IEEE International Conference on Image Processing (ICIP)*, pages 3896–3899, 2019.
- [8] Yin Zhou and Oncel Tuzel. Voxelnet: End-to-end learning for point cloud based 3d object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [9] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. *CoRR*, abs/1903.11027, 2019.
- [10] Dominik Kellner, Jens Klappstein, and Klaus Dietmayer. Grid-based dbscan for clustering extended objects in radar data. pages 365–370, 06 2012.
- [11] Chunxiao Wang, Min Ji, Jian Wang, Wei Wen, Ting Li, and Yong Sun. An improved dbscan method for lidar data segmentation with automatic eps estimation. *Sensors*, 19:172, 01 2019.
- [12] S. Ando, Y. Kusachi, A. Suzuki, and K. Arakawa. Appearance based pose estimation of 3d object using support vector regression. In *IEEE International Conference on Image Processing 2005*, volume 1, pages I–341, 2005.
- [13] Z. Zhao, P. Zheng, S. Xu, and X. Wu. Object detection with deep learning: A review. *IEEE Transactions on Neural Networks and Learning Systems*, 30(11):3212–3232, 2019.
- [14] Anant Ram, Jalal Sunita, Anand Jalal, and Kumar Manoj. A density based algorithm for discovering density varied clusters in large spatial databases. *International Journal of Computer Applications*, 3, 06 2010.
- [15] Debasish Basak, Srimanta Pal, and Dipak Patranabis. Support vector regression. *Neural Information Processing – Letters and Reviews*, 11, 11 2007.
- [16] Breiman, L. Random Forests. *Machine Learning* 45, 5–32 (2001). <https://doi.org/10.1023/A:1010933404324>
- [17] Chen, T. and Carlos Guestrin. “Chen, Tianqi, and Carlos Guestrin. “Xgboost: A scalable tree boosting system.” *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, 2016.
- [18] T. Kanungo, D. M. Mount, N. S. Netanyahu, C. D. Piatko, R. Silverman and A. Y. Wu, “An efficient k-means clustering algorithm: analysis and implementation,” in *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 24, no. 7, pp. 881–892, July 2002, doi: 10.1109/TPAMI.2002.1017616.
- [19] Tianwei Yin, Xingyi Zhou, and Philipp Krahenbuhl. Center-based 3d object detection and tracking. *arXiv:2006.11275*, 2020.
- [20] Ramin Nabati and Hairong Qi. Centerfusion: Center-based radar and camera fusion for 3d object detection. *arXiv preprint arXiv:2011.04841*, 2020.
- [21] Arsalan Mousavian, Dragomir Anguelov, John Flynn, and Jana Kosecka. 3d bounding box estimation using deep learning and geometry. *CoRR*, abs/1612.00496, 2016.