

Rope3D: The Roadside Perception Dataset for Autonomous Driving and Monocular 3D Object Detection Task

Xiaoqing Ye^{1*} Mao Shu^{1*} Hanyu Li¹ Yifeng Shi¹
Yingying Li¹ Guangjie Wang² Xiao Tan^{1†} Errui Ding¹

¹Baidu Inc. ²China University of Mining and Technology

Abstract

Concurrent perception datasets for autonomous driving are mainly limited to frontal view with sensors mounted on the vehicle. None of them is designed for the overlooked roadside perception tasks. On the other hand, the data captured from roadside cameras have strengths over frontal-view data, which is believed to facilitate a safer and more intelligent autonomous driving system. To accelerate the progress of roadside perception, we present the first high-diversity challenging **Roadside Perception 3D dataset- Rope3D** from a novel view. The dataset consists of 50k images and over 1.5M 3D objects in various scenes, which are captured under different settings including various cameras with ambiguous mounting positions, camera specifications, viewpoints, and different environmental conditions. We conduct strict 2D-3D joint annotation and comprehensive data analysis, as well as set up a new 3D roadside perception benchmark with metrics and evaluation devkit. Furthermore, we tailor the existing frontal-view monocular 3D object detection approaches and propose to leverage the geometry constraint to solve the inherent ambiguities caused by various sensors, viewpoints. Our dataset is available on <https://thudair.baai.ac.cn/rope>.

1. Introduction

Autonomous driving plays a crucial role in helping reduce traffic accidents and improve transportation efficiency. Current perceptual systems mainly equip the moving vehicle with LiDAR or camera sensors. Owing to the movement, the vehicle perceptual system can not observe surroundings for a long period. In addition, since the mounted sensor is relatively low (usually on the top of a vehicle), the perceptual range is comparatively limited and is vulnerable to occlusion. On the contrary, the data captured from road-




Figure 1. The comparison of (a) frontal view and (b) roadside camera view with a pitch angle. The car view focuses more on the frontal area whereas the roadside camera observes the scene in a long-term and large-range manner. Vehicles can be easily occluded by closer objects in frontal view but the roadside view alleviates the risk. For example, for car-view (a), the white van is occluded by the black jeep whereas in roadside view (b) they are both visible, corresponding to the white and pink 3D boxes in (c). The triangle mark denotes the same LiDAR-mounted vehicle.

side cameras has its inherent strengths in terms of robustness to occlusion and long-time event prediction, since they are collected from cameras mounted on poles a few meters above the ground. The comparisons between two different views of data are depicted in Fig. 1.

The importance of roadside perception is listed as follows: (1) Cooperative to Autonomous driving (AD). AD still faces safety challenges and uncontrolled threats due to blind spots. Instead, the roadside view can cover the blind spots for two extra advantages over car views: a long-range global perspective to extend vehicles' perception field spatially and temporally and global trajectory prediction for

*indicates equal contribution

†Xiao Tan (tanxchong@gmail.com) is the corresponding author.

safety. For example, a pedestrian walking behind a parked vehicle might suddenly crash into a moving vehicle since vehicle sensors fail to detect abrupt changes in the environment owing to the limited perceptual range or heavy occlusion. On the contrary, the roadside view is capable of behavior prediction timely. (2) Global perception. Further objects are occluded (even with 360° sensors) by closer objects in existing car-view datasets, causing blind spots. Thanks to roadside cameras mounted overhead, the invisible region is now visible. Besides, Autonomous vehicles (AV) can be informed to choose a faster lane when having a dead car in the queue since the roadside view perceives globally. (3) Cost-efficient. In terms of cost, it is worthy for ensuring safety by cooperative perception and cost-efficient since information from roadside cameras can broadcast to all surrounding AVs. (4) Intelligent traffic control. The roadside perception also facilitates smart traffic control and flow management. The critical contribution of roadside perceptual systems in facilitating a safer and more intelligent autonomous driving system has been acknowledged in many works [8, 34, 41].

However, existing researches on the roadside perceptual ability focus only on 2D tasks such as 2D detection and tracking, the ability of 3D localization is still under-explored [29, 30, 39]. In this work, we focus on monocular 3D detection that localizes objects in 3D space from a single image. Although abundant perception datasets have been published to fuel the development in autonomous driving from vehicle view, such as KITTI [12], nuScenes [5], A*3D [32] and Waymo [38], none of them is designed particularly for the overlooked roadside 3D perception task. We hence release the first large-scale high-diversity **Roadside Perception Dataset (Rope3D)**, with the hope of bridging this gap. Compared with the existing vehicle view datasets, the roadside perceptual data can be different in three ways. First, the ambiguity lies everywhere due to various cameras specifications such as distinct pitch angles of the viewpoint, mounting heights as well as various roadside environments, which increases the difficulty of monocular 3D detection tasks to a great extent. Second, since the roadside cameras are mounted on the poles instead of on top of the vehicle in frontal view, thus the assumption of the camera’s optical axis being parallel to the ground is no more valid, leading to the incompatibility of directly applying the existing monocular 3D detection approaches using this prior. Third, due to a much larger sensible range of the roadside perceptual system, a larger number of objects are expected to observe in roadside view, increasing the density and difficulty of a perceptual system. All these differences prevent directly applying most existing 3D detection methods. We hence tailor existing monocular 3D object detection methods to the roadside application.

To summarize, our contributions are as follows:

- We present the first challenging high-diversity road-

side dataset termed “Rope3D”, consisting of 50k images and over 1.5M 3D objects collected across a variety of lighting conditions (daytime / night / dusk), different weather conditions (rainy / sunny / cloudy), and distinct road scenes with different camera specifications like focal length and viewpoints.

- We specially tailor current frontal-view monocular 3D detection methods to deal with the roadside view data and conduct a comprehensive study with the new 3D detection metrics particularly designed for roadside 3D detection tasks, hoping to facilitate the development of monocular 3D perception tasks in roadside scenarios.

2. Related work

2.1. Frontal-view Autonomous Driving Datasets

Recent years have witnessed great progress in autonomous driving, thanks to a great number of large-scale traffic scene datasets. [2, 10, 11, 17, 26, 32, 44] As a pioneer work, KITTI [12] provides multimodal data and opens a leader board for a variety of tasks. Although the raw point clouds from LiDAR is 360° around the collecting vehicle, the annotation is only conducted within the overlapping frontal view of the camera. To tackle the limitation, the Honda Research Institute 3D Dataset (H3D) [31] supplies a total of 1.1M 3D boxes in full 360 view, which accelerates full-surround multi-object detection and tracking. Another 360° view multimodal 3D detection dataset is introduced in nuScenes [5], providing over 1.4M annotated 3D boxes in 1000 scenes, including nighttime and rainy conditions. The ApolloScape and ApolloCar3D [16, 37] focus on the pixel-wise semantic segmentation task, including scene parsing, 3D car instance, lane segmentation tasks. The Argoverse dataset [7] designs for vehicle perception tasks such as 3D tracking and motion forecasting. The Waymo Open Dataset [38] consists of over 1000 scenes and 12M 3D boxes in urban and suburban scenarios, under various weather and lighting conditions. Although the above-mentioned datasets have fueled the innovation on autonomous driving, they are all designed for vehicle view perception. However, there is a lack of a dataset helping us to effectively carry out research on 3D localization under roadside surveillance cameras. A related work is from BoxCars [36], which performs fine-grained vehicle recognition by estimating the projected vertices of 3D bounding box on the image rather than real-world location, size, and orientation. Another contemporary work [49] focuses on the LiDAR-based 3D detection task. The comparisons between our roadside 3D dataset and the previous AD datasets are listed in Table 1.

2.2. Monocular 3D Object Detection

Though challenging, monocular-based 3D detection is still an attractive solution especially in autonomous driving



Figure 2. The examples of collected samples under different weather and conditions. From left to right, each column corresponds to clear/sunny, rainy, night and dawn/dusk. More samples can be found in supplementary material.

| View | Dataset | RGB frames | Scenes | LiDAR channel | 3D Boxes | 2D Boxes | RGB resolution | Cls | Year | Diversity | | | Range (m) |
|----------|--------------------|------------|--------|-----------------|-------------------|----------|----------------|------|------|-----------|-------|------|-----------|
| | | | | | | | | | | Rain | Night | Dawn | |
| Frontal | KITTI [12] | 15k | 22 | 64 | 80k | 80k | 1392×512 | 8 | 2013 | | | | 70 |
| | Apollo Scape [16] | 144k | / | 64 | 70k | 0 | 3384×2710 | 8-35 | 2019 | | ✓ | | 420 |
| | AS lidar [26] | 90k | / | 64 | 475k | 0 | 1920×1080 | 8 | 2019 | | | | 70 |
| | Lyft Level 5 [17] | 46k | 366 | 40 | 1.3M | 0 | 1920×1080 | 9 | 2019 | | | | / |
| | A2D2 [13] | 12k | / | 16 | 9k | 0 | 1928×1208 | 38 | 2019 | | | | 100 |
| | Argoverse [7] | 22k | 113 | 32 | 993k | 0 | 1920×1200 | 15 | 2019 | ✓ | ✓ | | 200 |
| | H3D [31] | 27.7k | 160 | 64 | 1M | 0 | 1920×1200 | 8 | 2019 | | | | 100 |
| | A*3D [32] | 39k | / | 64 | 230k | 0 | 2048×1536 | 7 | 2020 | ✓ | ✓ | | 100 |
| | CityScapes 3D [11] | 5k | 1150 | no [†] | 27k | 0 | 2048×1024 | 8 | 2020 | | | | 150 |
| | nuScenes [5] | 1.4M | 1000 | 32 | 1.4M | 0 | 1600×900 | 23 | 2020 | ✓ | ✓ | | 75 |
| Roadside | Waymo Open [38] | 230k | 1150 | 64 | 12M | 9.9M | 1920×1080 | 4 | 2020 | ✓ | ✓ | ✓ | 75 |
| | ONCE [27] | 7M | 1M | 40 | 417k | 0 | 1920×1020 | 5 | 2021 | ✓ | ✓ | | 200 |
| | BoxCars116k [36] | 116k | 137 | no | 116k [‡] | 0 | ~128×128 | 6 | 2018 | | | | / |
| | Ours | 50k | 26 | 40/300 | 1.5M | 670k | 1920×1080 | 12 | / | ✓ | ✓ | ✓ | 200 |

Table 1. The comparison of 3D AD datasets. The top and bottom parts indicate the front view and roadside view datasets, respectively. LiDAR channel means the beam number of the LiDAR laser. The 2D boxes number denotes those who only have 2D box annotations. ‘/’ denotes unknown information. (†): No lidar sensors to obtain ground truth, use stereo instead. (‡): For BoxCars116k dataset, only the projected eight corner points instead of the 3D bounding boxes are provided. In other words, the location, dimension, and orientation of 3D bounding boxes are unknown. Besides, only cropped images around 128×128 rather than the full images are provided.

systems, where the method predicts the 3D bounding boxes from a single image [1, 9, 18–20, 22, 24, 28, 45]. Monocular 3D detection methods can be divided into three categories. **(1) Anchor-based.** Methods in this category exploit a series of predefined 3D bounding box with a location called “anchor” and estimate the offset w.r.t the anchor. M3D-RPN [3] leverages a 3D region proposal network and the geometric constrains of 2D and 3D perspectives to directly regress the 3D location and size. Kinematic3D [4] further extends M3D-RPN by leveraging 3D kinematics from monocular videos to improve the overall localization. **(2) Keypoint-based.** Many attempts [23, 25, 33, 40, 50] have been made to directly regress the keypoints, and then estimate 3D bounding box size and location from the image position of keypoints by optimization e.g., RTM-3D [21] and MonoGR-Net [33]. **(3) Pseudo-Lidar / depth based.** Extra depth estimation modules and/or point cloud guidance are employed to alleviate the lack of accurate depth information. The pioneering work pseudo-LiDAR [42, 43, 47] imitates the

LiDAR-based methods by utilizing off-the-shelf depth estimators to convert image pixels into pseudo-LiDAR point clouds, and employs LiDAR-based approaches for further detection. DA-3Ddet [46] adapts the features from unsound image-based pseudo-LiDAR domain to reliable LiDAR domain for guidance to boost the monocular performance. UrbanNet [6] utilizes the urban 3D map, including driving lanes, elevation, and slope as prior for assisting the 3D detection task. The existing monocular 3D detection methods are mainly designed for processing vehicle view data. Due to the domain gap and the distribution shift, a question raise naturally is whether these methods still applicable for roadside applications and if not how can we tailor these methods to the new scenario.

3. The Roadside Perception Dataset

3.1. Specifications

Sensors Setup. The roadside data collection is conducted by two sets of sensors, one is the roadside cameras mounted

on the pole or beside the traffic light; the other is the LiDAR equipped on a parked/driving vehicle to obtain the 3D point clouds of the same scene. For sensor synchronization, we adopt the nearest time matching strategy to find the pairs of image and point clouds within 5 milliseconds error.

- Roadside cameras: RGB with 1920×1080 resolution, 30-60Hz capture frequency and 1/1.8" CMOS sensor.
- LiDAR: (1) HESAI Pandar 40P, 40 laser beams, 10/20Hz capture frequency, $\pm 2cm$ accuracy, 360° horizontal FOV, $-25^\circ \sim +15^\circ$ vertical FOV, $\leq 200m$ range. (2) Jaguar Prime from Innovusion: 300 beams, 6-20 FPS with $\pm 3cm$ accuracy, 100° horizontal FOV, 40° vertical FOV, $\leq 280m$ range.

Coordinate Systems and Calibration. There are three coordinate systems used in the dataset: the World Coordinate (*i.e.*, the Universal Transverse Mercator coordinate system (UTM Coord.)), the Camera Coordinate, as well as the LiDAR Coordinate.

To obtain the reliable ground truth 2D-3D joint annotation, the calibration between different sensors is required. First, the camera is calibrated to obtain the intrinsics by detecting the chessboard patterns. Then the Lidar-to-World calibration is conducted by the vehicle localization module to obtain the high definition (HD) map in UTM Coord. For World-to-Camera calibration, we first project the HD map which contains lane and crosswalk endpoints onto the 2D image to obtain the raw transformation. A bundle adjustment refinement is followed to derive the final transformation. Then Lidar-to-Camera transformation can be obtained by simply multiplying Lidar-to-World and World-to-Camera transformations.

After obtaining the transformation between the three coordinate systems, we can easily compute the ground equation $G(\alpha, \beta, \gamma, d)$ by fitting the ground points $[x, y, z]$ to the ground plane in the camera coordinate, with $\alpha x + \beta y + \gamma z + d = 0$.

3.2. Data Collection and Annotation

After obtaining the intrinsics as well as the LiDAR-to-Camera transformation, we can collect the 2D-3D data. We choose various roadside cameras and let a LiDAR-equipped vehicle park or drive around. To keep the high diversity and complexity of the real environment, we collect more than 50k image frames at different times (daytime, night, dawn/dusk), different weather conditions (sunny, cloudy, rainy), different densities (crowded, normal, less traffic), different distributions of traffic elements and so on. There are totally 13 object classes with their corresponding category, 2D properties (occlusion, truncation) and the 7-DOF 3D bounding box: Location (x, y, z), Size (width- W , length- L , height- H), Orientation (the heading angle- θ). The full pipeline is illustrated in Fig. 3. (1) First after obtaining the 3D point clouds and the 2D image (they are within

the same space but differ in the viewpoint), we first annotate the 3D bounding boxes directly on the 3D point clouds. (2) Simultaneously, the annotated 3D bounding boxes will be projected on the 2D image plane, see the top part of Fig. 3(c). We adjust the 3D parameters so that the projected points align with the 2D instance and mainly cover it. (3) For 2D box annotations, if the objects are scanned by the laser, their 2D box labels in the image are the minimum bounding box of the amodal projections of the eight 3D corners. For objects that are heavily occluded or too far to be detected by laser, 2D complementary labeling is conducted to label 2D bounding boxes directly in the image and leave its 3D annotations empty, see the bottom part of Fig. 3 (c).

3.3. Statistics and Analysis

The collected images have high diversity and inherent ambiguity due to different settings of pitch angles, height, and camera types as well as the manifold scenes. Two levels of categories are adopted in the dataset. The coarse-grained level mainly focuses on the most common traffic elements: Car, Big Vehicle, Pedestrian, and Cyclist. To be more fine-grained, Car includes car and van, Big Vehicle can be further divided into truck and bus, and meanwhile, Cyclist can be subdivided into cyclist, motorcyclist, barrow, and tricyclist, since they are driving non-motor vehicles. We have annotated 13 classes, in addition to the above-mentioned categories, there are four extra classes: ‘traffic cones’, ‘triangle_plate’, ‘unknown-unmovable’, ‘unknown-movable’. The following statistics are mainly on the coarse and fine-grained classes.

Quantity distribution. We first give an overview analysis of the dataset on the number of 2D and 3D objects in Fig. 4. As is stated, 2D objects are more than 3D objects since some objects are not scanned by the LiDAR laser, so they only have 2D annotations. We give the detailed number of coarse level and fine-grained levels of the categories, corresponding to (a) and (b) in Fig. 4.

Depth distribution. Besides, we analyze the depth distribution of coarse-categories in Fig. 5. The depth of the captured 3D objects can range from within 10m to over 140m. Most objects lie between 60 and 80 meters.

Density. The density can be a key factor that affects the capability of 3D perception. Thus, we analyze the density of the dataset from two aspects in Fig. 6. From the global level, we show the 2D and 3D annotated number of each image in (a). The samples can be up to more than one hundred. Compared to other datasets whose densities are KITTI [12]: 5.3, nuScenes [5]: 9.7 and A*3D [32]: 5.9, our dataset has a much higher density (34 and 24 for 2D/3D per image). From the view of coarse categories, the numbers of 3D samples per frame are shown in Fig. 6 (b). The ‘Car’ category is relatively evenly distributed in densities whereas there are less than 10 big vehicles in each image.

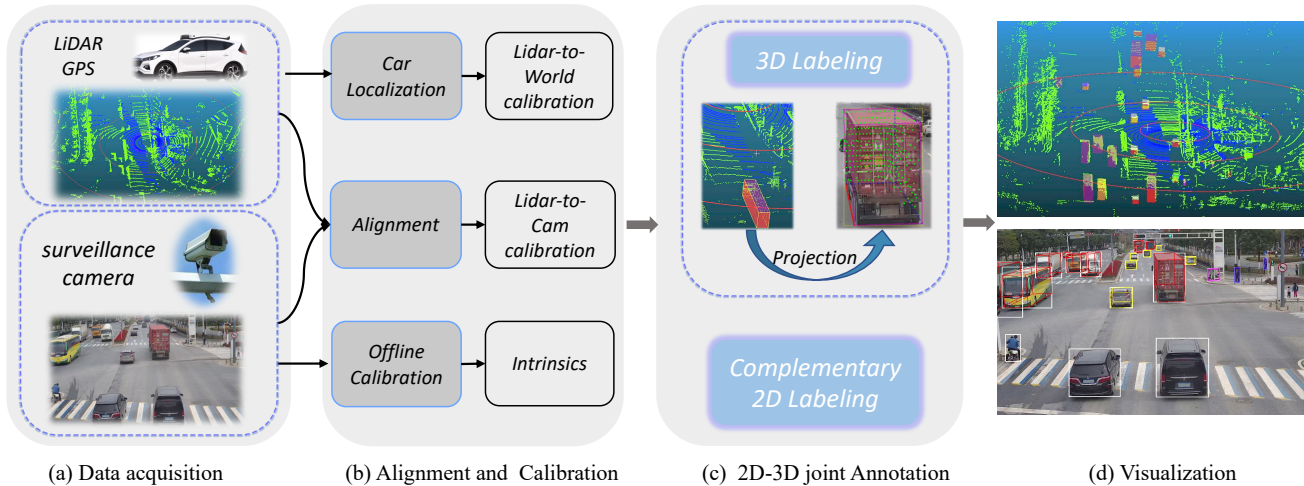


Figure 3. The data acquisition and labeling pipeline. Our platform takes the captured roadside images and the point clouds scanned by the LiDAR mounted on a parked/driving vehicle as input. After calibration and alignment between various sensors, the transformations between LiDAR, the world, and the camera are obtained, as well as the ground plane equation and intrinsics. The 2D-3D joint annotation is carried out by projecting the point clouds onto the images and adjusting the 3D bounding boxes manually to fit the 2D instance. For objects that are not scanned by the laser, the 2D complementary labeling is performed on the images only. For example in (d), some objects only have white 2D bounding boxes and no 3D colored annotations due to a lack of 3D points.

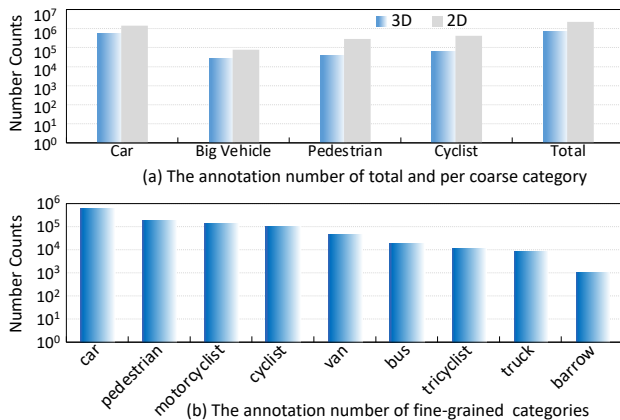



Figure 4. The quantities of: (a) coarse level, (b) fine-grained level.



Occlusion and truncation Analysis. Next, we annotate three levels for occlusion and truncation attributes. For occlusion, Level 0 denotes no occlusion, 1 and 2 means less / more than 50% occlusion. For truncation attribute, Level 0 means no truncation, 1 and 2 denote the horizontal and vertical truncation in the image border. The statistics are shown in Fig. 7. More than half of the objects are partially or heavily occluded while the occlusion percentage of KITTI is be-

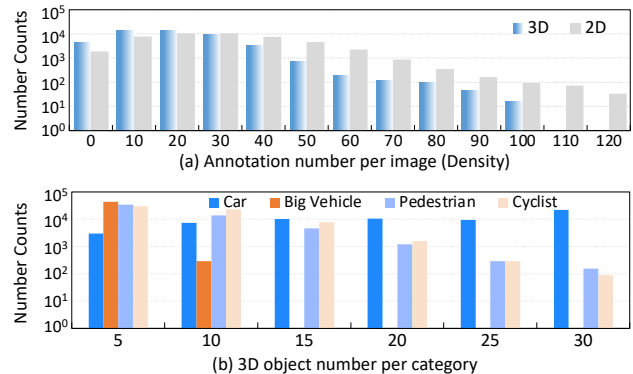


Figure 6. Top: The number of total annotated objects per image. Bottom: The number of coarse-category 3D objects per image.

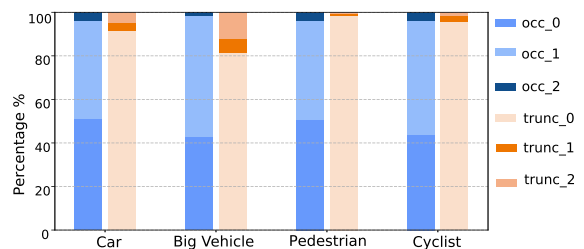


Figure 7. The occlusion and truncation distribution of coarse categories. Over half of the objects are partially or heavily occluded.

tween 5% to 30%, which reflects the difficulty of our 3D perception dataset and task.

The ambiguity analysis. The roadside dataset has inherent ambiguity due to the adopted various cameras with different settings of camera specifications, mounting heights, the

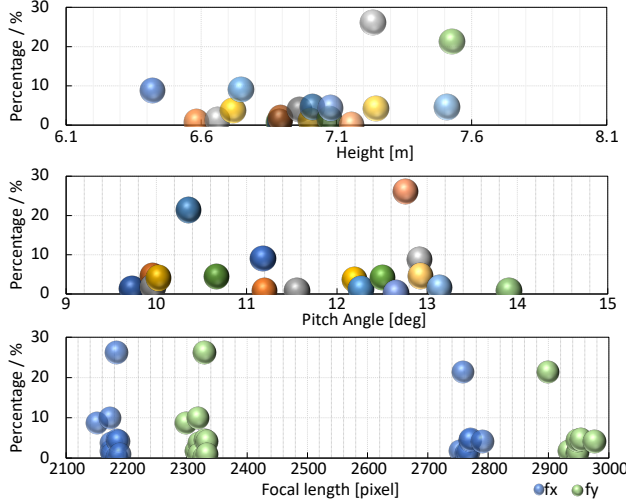


Figure 8. The diversity of roadside cameras. From top to bottom are the distribution of mounting heights, pitch angles of the cameras and the focal lengths over the dataset, respectively.

pitch angles of viewpoint, so on. Thus we analyze the diversity distribution of settings in Fig. 8.

4. Task and Metrics

4.1. The task overview and metrics.

Task Overview. The monocular 3D perception task is to localize the objects given ambiguous images captured under various settings and scenes, including different camera specifications, viewpoints, and mounting positions.

Evaluation Metrics. There are multiple evaluation metrics for AD datasets. In KITTI [12], the 11-point Interpolated Average Precision metric ($AP|_{R_{11}}$) and the average orientation similarity (AOS) are proposed to assess the localization and orientation performance. [35] suggests 40 recall positions instead of 11 positions for a more fair comparison. The nuScenes [5] consolidates the mean AP and five True Positive (TP) error types into the nuScenes detection score (NDS), including translation, scale, orientation, velocity and attribute error types. Inspired by the AP metric, we adopt $AP|_{R_{40}}$ [35], *i.e.*,

$$AP|_R = \frac{1}{|R|} \sum_{r \in R} \max_{\tilde{r} \geq r} \rho(\tilde{r}) \quad (1)$$

where $\rho(\tilde{r})$ is the precision at a certain recall threshold $r \in \{1/40, 2/40, \dots, 1\}$. To facilitate a comprehensive evaluation *w.r.t.* a certain factor such as orientation, we decouple the consolidated metric into several sub-metrics.

Average Ground Center Similarity. The distance between objects on the ground plane indicates the risk of collision, we hence compute the ground Euclidean distance by pro-

jecting the object center points onto the ground plane:

$$ACS = \frac{1}{|D|} \sum_{s \in D} \left(1 - \min \left(1, \frac{\Delta_s^c}{C_s} \right) \right) \quad (2)$$

where D is set of true positive samples, C_s is the norm of the GT ground center, Δ_s^c is the Euclidean distance between predicted ground center and GT ground center of sample s . $|D|$ is the total number of true positive objects.

Average Orientation Similarity. The Average Orientation Similarity (AOS) is introduced to measure how well the orientation is estimated, which is defined similarly as [38] by,

$$AOS = \frac{1}{|D|} \sum_{s \in D} \frac{1 + \cos(2 * \Delta_s^\theta)}{2} \quad (3)$$

where Δ_s^θ is the angle difference of sample s , and $\cos(2 * \Delta_s^\theta)$ means that during evaluation, we don't distinguish whether the head or tail of the object is facing the camera.

Average Area Similarity. We measure the ground occupancy of the prediction *w.r.t.* the GT in term of area, where Δ_s^A is the absolute area difference and A_s is the ground truth area.

$$AAS = \frac{1}{|D|} \sum_{s \in D} \left(1 - \min \left(1, \frac{\Delta_s^A}{A_s} \right) \right) \quad (4)$$

Average Four Ground Points Distance and Similarity. We also compute the average distance of four ground vertices of the 3D bounding box (AGD), since it consolidates the location, orientation, and width/length together.

$$AGD = \frac{1}{|D|} \sum_{s \in D} \left(\frac{1}{K} \sum_{g=0}^{K-1} |s_g - \hat{s}_g| \right) \quad (5)$$

where \hat{s}_g and s_g are the g th predicted and GT ground points of sample s , respectively. $K = 4$ is the total number of ground points. To be consistent with other similarity metrics, we define the AGS (ground points similarity) as:

$$AGS = \frac{1}{|D|} \sum_{s \in D} \left(1 - \min \left(1, \frac{1}{K} \sum_{g=0}^{K-1} \frac{|s_g - \hat{s}_g|}{|\hat{c}|} \right) \right) \quad (6)$$

Assume $S = (ACS + AOS + AAS + AGS)/4$, we consolidate into $Rope_{score}$ by reweighting the 3D AP and the proposed similarities metrics with $\omega_1 = 8$ and $\omega_2 = 2$.

$$Rope_{score} = (\omega_1 * AP + \omega_2 * S) / (\omega_1 + \omega_2) \quad (7)$$

4.2. 3D Roadside Perception Task

As is illustrated in Sec. 1, due to the inherent ambiguity of roadside data caused by diverse camera specifications (various intrinsics and mounting positions, et.al.), the existing frontal-view monocular 3D object detection approaches can not be directly applied to the Rope3D Dataset. Hence

we make simple-and-effective attempts to alleviate the ambiguity problem by utilizing camera specifications and encoding the ground knowledge. Two modifications (early-fusion and deep-fusion) to incorporate the depth map of the ground plane with RGB image and two kinds (integrate and multi-gridded) of ground planes are made to alleviate the multi-focal ambiguity.

Adaptations by leveraging ground planes. We adopt the ground plane equation $G(\alpha, \beta, \gamma, d)$ and camera intrinsic $K^{3 \times 3}$ to generate the depth map D_G of the ground plane with the same size as the image.

$$\begin{cases} Z [x, y, 1]^T = K^{3 \times 3} [X, Y, Z]^T \\ G^{1 \times 4} [X, Y, Z, 1]^T = 0 \end{cases} \quad (8)$$

where $[x, y]$ is the pixel in the image coordinates, $[X, Y, Z]$ is the corresponding 3d point in camera coordinate that lies on the ground plane. Thus the depth Z can be derived with the known 2d image points and the ground plane equation G . We incorporate the ground depth map with the RGB appearance feature by early fusion and deep fusion. The first one is directly concatenating the depth map with the original RGB channels as input, and the second is adopting another siamese network for depth feature extraction and further weighted fusion of the two depth predictions. The performances of these two methods are similar and we hence only report the results by concatenation on the anchor-based M3D-RPN and the keypoint-based MonoDLE and MonoFlex approaches. We believe more sophisticated approaches might further improve the performance, which is out of the scope of this paper. In addition, two different formats of ground planes are attempted. One is to fit the entire ground within the visual field to a single plane, which is represented by the ground plane equation. Another is to divide the entire ground into multiple small grids, and each grid is represented by a ground equation.

5. Experiments

5.1. Experimental Setup

Our roadside 3D perception dataset contains 50k images, with the training and validation ratio set to 8:2. We offer two kinds of splitting the training and validation set, \mathbb{I} : Homologous, for each scene we select 70% images and combine them for training, and leave all the rest images for validation. \mathbb{III} : Heterologous, we select 80% of the cameras with the collected images for training and leave the remaining unseen 20% (different camera specifications) for validation, which can be used for validating the generalization ability of the monocular 3D object detection approaches.

Implementation Detail. (1) For M3D-RPN [3], we experiment on vanilla and improved approaches with ResNet34 [14] backbone, (2) Kinematic3D [4] is a monocular video-based 3D object detector with DenseNet121 [15] backbone,

we only implement the first phase without video knowledge. (3) MonoDLE [25] is based on the anchor-free one stage detector CenterNet [51] with backbone DLA34 [48]. (4) MonoFlex [50] is a keypoint based method with modified DLA34 [48] backbone. The training image resolutions are adjusted to fit our dataset.

For objects that contain only 2D annotations without 3D annotations, we compute only losses on 2D attributes. For those objects having 3D labeling, the training loss weights of the 2D and 3D are both set to 1.

5.2. Main Results and Analysis

Performance of vanilla and improved approaches. The performances of monocular 3D detection approaches on the Rope3D Dataset are depicted in Table 2. Approaches with suffix (G) denote that we customize the corresponding approach with ground plane function to reconnect the 3D locations and 2D projected points even when the optical axis of cameras is not parallel to the ground plane because of the pitch angle. The improved approach is noted with suffix (D). We adopt the 3D $AP|_{R40}$ as well as the proposed metrics for evaluation. We find that most approaches have an obvious performance decline from homologous to the heterologous validation set. However, the performance drop is insignificant when applying a 3D detection model trained on the vehicle-view ONCE dataset to nuScenes dataset [27]. This phenomenon indicates the domain gap caused by various camera specifications and setting positions can not be ignored, which might be a distinguishable difference between vehicle-view and roadside view applications and should be taken carefully. By leveraging the depth map of ground planes, we observe an obvious improvement in most methods, even on the heterologous set where training and validating sets have different camera specifications.

Performance of different ground plane formats. We further analyze the performance by adopting two different formats of ground planes, *i.e.*, by fitting the entire ground within the visual field to a single plane or by dividing the entire ground into multiple $5m \times 5m$ grids piecwisely. We carry out the experiment on KM3D [20], a method that predicts 2D keypoints and solves 3D position by minimizing the re-projection error. In other words, it relies on the differentiable geometric constraint to recover 3D location rather than direct prediction, which heavily depends on the accuracy of ground plane. As is shown in Table 3, KM3D- (GG) , which takes advantage of gridded ground planes, shows better performance on most fine-grained categories. The main reason might be that the piecwisely-gridded planes better fit the actual the ground plane.

Performance of different ranges. We further analyze the performance of the models within different ranges from 0 to 120m. As is shown in Table 4, with the depth range increases, most of the performances decrease, especially for

| Setting | Method | Backbone | Branch | IoU = 0.5 | | | | IoU = 0.7 | | | |
|---------|------------------------------|-------------|--------|----------------------|-----------------------|----------------------|-----------------------|----------------------|-----------------------|----------------------|-----------------------|
| | | | | Car | | Big Vehicle | | Car | | Big Vehicle | |
| | | | | AP _{3D R40} | Rope _{score} | AP _{3D R40} | Rope _{score} | AP _{3D R40} | Rope _{score} | AP _{3D R40} | Rope _{score} |
| II | M3D-RPN-(<i>G</i>) [3] | ResNet34 | A | 54.19 | 62.65 | 33.05 | 44.94 | 16.75 | 32.90 | 6.86 | 24.19 |
| | M3D-RPN-(<i>D</i>) [3] | ResNet34 | A | 67.17 | 73.14 | 39.06 | 49.95 | 33.94 | 46.45 | 11.28 | 28.12 |
| | Kinematic3D-(<i>G</i>) [4] | DenseNet121 | A | 50.57 | 58.86 | 37.60 | 48.08 | 17.74 | 32.99 | 6.10 | 22.88 |
| | MonoDLE-(<i>G</i>) [25] | DLA-34 | K | 51.70 | 60.36 | 40.34 | 50.07 | 13.58 | 29.46 | 9.63 | 25.80 |
| | MonoDLE-(<i>D</i>) [25] | DLA-34 | K | 77.50 | 80.84 | 49.07 | 57.22 | 54.53 | 62.48 | 17.25 | 32.00 |
| | MonoFlex-(<i>G</i>) [50] | DLA-34 | K | 60.33 | 66.86 | 37.33 | 47.96 | 33.78 | 46.12 | 10.08 | 26.16 |
| | MonoFlex-(<i>D</i>) [50] | DLA-34 | K | 59.78 | 66.66 | 59.81 | 66.07 | 35.64 | 47.43 | 24.61 | 38.01 |
| III | M3D-RPN-(<i>G</i>) [3] | ResNet34 | A | 21.75 | 36.40 | 21.49 | 35.49 | 6.05 | 23.84 | 2.78 | 20.82 |
| | M3D-RPN-(<i>D</i>) [3] | ResNet34 | A | 36.33 | 48.16 | 24.39 | 37.81 | 11.09 | 28.17 | 3.39 | 21.01 |
| | Kinematic3D-(<i>G</i>) [4] | DenseNet121 | A | 23.56 | 37.05 | 13.85 | 28.58 | 5.82 | 23.06 | 1.27 | 18.92 |
| | MonoDLE-(<i>G</i>) [25] | DLA-34 | K | 19.08 | 33.72 | 19.76 | 33.07 | 3.77 | 21.42 | 2.31 | 19.55 |
| | MonoDLE-(<i>D</i>) [25] | DLA-34 | K | 31.33 | 43.68 | 23.81 | 36.21 | 12.16 | 28.39 | 3.02 | 19.96 |
| | MonoFlex-(<i>G</i>) [50] | DLA-34 | K | 32.01 | 44.37 | 13.86 | 28.47 | 10.86 | 27.39 | 0.97 | 18.18 |
| | MonoFlex-(<i>D</i>) [50] | DLA-34 | K | 37.27 | 48.58 | 47.52 | 55.86 | 11.24 | 27.79 | 13.10 | 28.22 |

Table 2. Overall performance of the monocular 3D object detection approaches on the Rope3D Dataset with IoU = 0.5 and 0.7 under two train-val splitting settings: the homologous (II) and the heterologous (III). -(*G*) denotes adapting the ground plane, -(*D*) means using the depth map of ground. The abbr. in the branch column denotes: A: anchor-based, K: keypoint-based.

| Setting | Method | Backbone | AP _{3D R40} [Mod] / Rope _{score} | | | | | | | | | |
|---------|-------------------------|----------|--|---------------|---------------|---------------|---------------|---------------|---------------|---------------|--|--|
| | | | car | van | bus | truck | cyclist | motorcyclist | tricyclist | pedestrian | | |
| II | KM3D [20] | ResNet34 | 8.97 / 25.09 | 7.77 / 23.79 | 8.07 / 23.89 | 4.94 / 20.59 | 1.81 / 17.34 | 3.61 / 19.34 | 14.39 / 27.85 | 0.37 / 17.93 | | |
| | KM3D-(<i>G</i>) [20] | ResNet34 | 9.83 / 26.60 | 13.16 / 29.26 | 4.19 / 22.05 | 18.42 / 32.40 | 11.35 / 27.24 | 11.45 / 27.24 | 19.13 / 33.50 | 9.90 / 26.28 | | |
| | KM3D-(<i>GG</i>) [20] | ResNet34 | 9.86 / 26.64 | 15.71 / 31.30 | 7.66 / 24.48 | 12.67 / 27.71 | 13.23 / 28.94 | 15.08 / 30.14 | 19.97 / 34.13 | 11.92 / 27.90 | | |
| III | KM3D [20] | ResNet34 | 5.89 / 22.91 | 2.91 / 20.26 | 21.20 / 34.30 | 25.86 / 37.46 | 1.36 / 17.14 | 4.67 / 20.03 | 2.40 / 19.48 | 0.29 / 17.93 | | |
| | KM3D-(<i>G</i>) [20] | ResNet34 | 17.39 / 32.71 | 30.48 / 43.22 | 21.25 / 35.20 | 34.93 / 45.84 | 24.98 / 38.30 | 14.49 / 29.47 | 47.47 / 56.79 | 12.61 / 28.44 | | |
| | KM3D-(<i>GG</i>) [20] | ResNet34 | 23.70 / 37.90 | 31.37 / 44.04 | 19.99 / 34.61 | 37.65 / 48.34 | 26.38 / 39.58 | 16.58 / 30.82 | 54.03 / 62.46 | 12.81 / 28.46 | | |

Table 3. Performance of KM3D with different ground plane formats on the Rope3D Dataset under two train-val splitting settings: the homologous (II) and the heterologous (III). -(*G*) denotes adapting the ground plane equation, -(*GG*) means using the gridded ground planes. IoU = 0.25 for non-motor vehicles and pedestrian, IoU = 0.5 for motor vehicles.

| Method | Range (m) | AP _{3D R40} / Rope _{score} | | | |
|-----------------------|-----------|--|---------------|---------------|---------------|
| | | Car | Big Vehicle | Cyclist | Pedestrian |
| MonoDLE-(<i>G</i>) | all | 19.08 / 33.72 | 19.76 / 33.07 | 10.93 / 26.44 | 3.72 / 21.42 |
| | 0-30 | 31.43 / 43.12 | 37.36 / 46.69 | 19.83 / 33.28 | 10.26 / 26.61 |
| | 30-60 | 10.42 / 26.68 | 8.68 / 24.02 | 8.11 / 24.23 | 3.90 / 21.66 |
| | 60-90 | 18.60 / 33.42 | 32.46 / 44.15 | 9.29 / 25.23 | 2.37 / 20.28 |
| | 90-120 | 11.84 / 28.05 | 10.29 / 26.09 | 9.84 / 25.11 | 2.88 / 20.58 |
| MonoFlex-(<i>G</i>) | all | 32.01 / 44.37 | 13.86 / 28.47 | 44.27 / 53.58 | 25.48 / 39.04 |
| | 0-30 | 15.49 / 30.49 | 27.68 / 39.16 | 61.94 / 67.37 | 37.25 / 48.50 |
| | 30-60 | 45.69 / 55.33 | 12.18 / 27.48 | 50.70 / 58.78 | 35.74 / 47.37 |
| | 60-90 | 46.72 / 56.41 | 19.34 / 33.45 | 30.65 / 42.82 | 8.94 / 25.65 |
| | 90-120 | 14.19 / 30.15 | 1.30 / 18.82 | 9.43 / 25.20 | 4.79 / 22.35 |
| M3D-RPN-(<i>D</i>) | all | 36.33 / 48.16 | 24.39 / 37.81 | 11.22 / 27.54 | 3.93 / 21.54 |
| | 0-30 | 52.07 / 60.60 | 24.07 / 37.18 | 16.77 / 31.80 | 5.19 / 22.53 |
| | 30-60 | 33.57 / 46.30 | 25.22 / 38.84 | 14.38 / 30.36 | 8.09 / 24.99 |
| | 60-90 | 24.07 / 38.60 | 39.17 / 50.42 | 5.80 / 23.02 | 1.10 / 19.16 |
| | 90-120 | 11.19 / 28.23 | 6.55 / 24.06 | 4.23 / 21.48 | 0.13 / 18.02 |

Table 4. The performance within different ranges on the heterologous (III) set. We set IoU = 0.5 for Car / Big Vehicle and IoU = 0.25 for Cyclist and Pedestrian, following KITTI's hard-modality.

90-120 meters. The reason is owing to two aspects: too small area in the image to extract strong features for learning and much less 3D annotations in far-away regions due to occlusion. MonoFlex-(*G*) shows better performance on cyclists and pedestrians whereas inferior in motor vehicles.

6. Conclusion

We propose the first high-diversity challenging roadside monocular 3D perception dataset - Rope3D. Rope3D is collected from the roadside view with joint 2D-3D annotations, making it unique from any previously released datasets and is particularly designed for the roadside 3D perception. Furthermore, we specially tailor the existing monocular 3D object detection approaches to the novel dataset, due to its unique viewpoint and inherent ambiguity lying in the various camera specifications and diverse road scenes. We hope to raise the attention to the special view - roadside view, so as to facilitate a safer and more intelligent autonomous driving system.

Ethical concerns and Limitation. To prevent from being utilized for illegal surveillance, all the images in the dataset are time-discrete and not allowed for tracking tasks. Note that all the sensitive information including license plates, human faces, names of bus stops, roads, and buildings are totally masked. The attempts for adapting the concurrent vehicle-view 3D detection approaches need further study.

References

- [1] Yousef Atoum, Joseph Roth, Michael Bliss, Wende Zhang, and Xiaoming Liu. Monocular video-based trailer coupler detection using multiplexer convolutional neural network. In *ICCV*, 2017. 3
- [2] Apratim Bhattacharyya, Daniel Olmeda Reino, Mario Fritz, and Bernt Schiele. Euro-pvi: Pedestrian vehicle interactions in dense urban centers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6408–6417, 2021. 2
- [3] Garrick Brazil and Xiaoming Liu. M3d-rpn: Monocular 3d region proposal network for object detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 9287–9296, 2019. 3, 7, 8
- [4] Garrick Brazil, Gerard Pons-Moll, Xiaoming Liu, and Bernt Schiele. Kinematic 3d object detection in monocular video. In *European Conference on Computer Vision*, pages 135–152. Springer, 2020. 3, 7, 8
- [5] 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 multi-modal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020. 2, 3, 4, 6
- [6] Juan Carrillo and Steven Waslander. Urbannet: Leveraging urban maps for long range 3d object detection. In *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*, pages 3799–3806. IEEE, 2021. 3
- [7] Ming-Fang Chang, John W Lambert, Patsorn Sangkloy, Jagjeet Singh, Slawomir Bak, Andrew Hartnett, De Wang, Peter Carr, Simon Lucey, Deva Ramanan, and James Hays. Argoverse: 3d tracking and forecasting with rich maps. In *CVPR*, 2019. 2, 3
- [8] Qi Chen, Sihai Tang, Qing Yang, and Song Fu. Cooper: Cooperative perception for connected autonomous vehicles based on 3d point clouds. In *2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS)*, pages 514–524. IEEE, 2019. 2
- [9] Xiaozhi Chen, Kaustav Kundu, Ziyu Zhang, Huimin Ma, Sanja Fidler, and Raquel Urtasun. Monocular 3d object detection for autonomous driving. In *CVPR*, 2016. 3
- [10] Yukyung Choi, Namil Kim, Soonmin Hwang, Kibaek Park, Jae Shin Yoon, Kyoungwan An, and In So Kweon. Kaist multi-spectral day/night data set for autonomous and assisted driving. *IEEE Transactions on Intelligent Transportation Systems*, 19(3):934–948, 2018. 2
- [11] Nils Gähler, Nicolas Jourdan, Marius Cordts, Uwe Franke, and Joachim Denzler. Cityscapes 3d: Dataset and benchmark for 9 dof vehicle detection. *arXiv preprint arXiv:2006.07864*, 2020. 2, 3
- [12] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE conference on computer vision and pattern recognition*, pages 3354–3361. IEEE, 2012. 2, 3, 4, 6
- [13] Jakob Geyer, Yohannes Kassahun, Mentar Mahmudi, Xavier Ricou, Rupesh Durgesh, Andrew S Chung, Lorenz Hauswald, Viet Hoang Pham, Maximilian Mühlegg, Sebastian Dorn, et al. A2d2: Audi autonomous driving dataset. *arXiv preprint arXiv:2004.06320*, 2020. 3
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 7
- [15] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017. 7
- [16] Xinyu Huang, Peng Wang, Xinjing Cheng, Dingfu Zhou, Qichuan Geng, and Ruigang Yang. The apolloscape open dataset for autonomous driving and its application. *IEEE transactions on pattern analysis and machine intelligence*, 42(10):2702–2719, 2019. 2, 3
- [17] R. Kesten, M. Usman, J. Houston, T. Pandya, K. Nadhamuni, A. Ferreira, M. Yuan, B. Low, A. Jain, P. Ondruska, S. Omari, S. Shah, A. Kulkarni, A. Kazakova, C. Tao, L. Platinisky, W. Jiang, and V. Shet. Lyft level 5 av dataset 2019. 2019. 2, 3
- [18] Jason Ku, Alex D Pon, and Steven L Waslander. Monocular 3d object detection leveraging accurate proposals and shape reconstruction. In *CVPR*, 2019. 3
- [19] Buyu Li, Wanli Ouyang, Lu Sheng, Xingyu Zeng, and Xiaogang Wang. Gs3d: An efficient 3d object detection framework for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1019–1028, 2019. 3
- [20] Peixuan Li and Huaici Zhao. Monocular 3d detection with geometric constraint embedding and semi-supervised training. *IEEE Robotics and Automation Letters*, 6(3):5565–5572, 2021. 3, 7, 8
- [21] Peixuan Li, Huaici Zhao, Pengfei Liu, and Feidao Cao. Rtm3d: Real-time monocular 3d detection from object keypoints for autonomous driving. *arXiv preprint arXiv:2001.03343*, 2020. 3
- [22] Lijie Liu, Jiwen Lu, Chunjing Xu, Qi Tian, and Jie Zhou. Deep fitting degree scoring network for monocular 3d object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1057–1066, 2019. 3
- [23] Zechen Liu, Zizhang Wu, and Roland Tóth. Smoke: Single-stage monocular 3d object detection via keypoint estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 996–997, 2020. 3
- [24] Xinzhu Ma, Zhihui Wang, Haojie Li, Pengbo Zhang, Wanli Ouyang, and Xin Fan. Accurate monocular 3d object detection via color-embedded 3d reconstruction for autonomous driving. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 6851–6860, 2019. 3
- [25] Xinzhu Ma, Yinmin Zhang, Dan Xu, Dongzhan Zhou, Shuai Yi, Haojie Li, and Wanli Ouyang. Delving into localization errors for monocular 3d object detection. In *CVPR*, pages 4721–4730, 2021. 3, 7, 8
- [26] Yuexin Ma, Xinge Zhu, Sibozhang, Ruigang Yang, Wenping Wang, and Dinesh Manocha. Trafficpredict: Trajectory

- prediction for heterogeneous traffic-agents. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6120–6127, 2019. 2, 3
- [27] Jiageng Mao, Minzhe Niu, Chenhan Jiang, Xiaodan Liang, Yamin Li, Chaoqiang Ye, Wei Zhang, Zhenguo Li, Jie Yu, Chunjing Xu, et al. One million scenes for autonomous driving: Once dataset. 2021. 3, 7
- [28] Arsalan Mousavian, Dragomir Anguelov, John Flynn, and Jana Kosecka. 3d bounding box estimation using deep learning and geometry. In *CVPR*, pages 7074–7082, 2017. 3
- [29] Milind Naphade, Zheng Tang, Ming-Ching Chang, David C Anastasiu, Anuj Sharma, Rama Chellappa, Shuo Wang, Pranamesh Chakraborty, Tingting Huang, Jenq-Neng Hwang, et al. The 2019 ai city challenge. In *CVPR Workshops*, volume 8, 2019. 2
- [30] Milind Naphade, Shuo Wang, David C Anastasiu, Zheng Tang, Ming-Ching Chang, Xiaodong Yang, Yue Yao, Liang Zheng, Pranamesh Chakraborty, Christian E Lopez, et al. The 5th ai city challenge. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4263–4273, 2021. 2
- [31] Abhishek Patil, Srikanth Malla, Haiming Gang, and Yi-Ting Chen. The h3d dataset for full-surround 3d multi-object detection and tracking in crowded urban scenes. In *ICRA*, pages 9552–9557. IEEE, 2019. 2, 3
- [32] Quang-Hieu Pham, Pierre Sevestre, Ramanpreet Singh Pahwa, Huijing Zhan, Chun Ho Pang, Yuda Chen, Armin Mustafa, Vijay Chandrasekhar, and Jie Lin. A* 3d dataset: Towards autonomous driving in challenging environments. In *ICRA*, pages 2267–2273. IEEE, 2020. 2, 3, 4
- [33] Zengyi Qin, Jinglu Wang, and Yan Lu. Monogmet: A geometric reasoning network for monocular 3d object localization. In *AAAI*, 2019. 3
- [34] Andreas Rauch, Felix Klanner, Ralph Rasshofer, and Klaus Dietmayer. Car2x-based perception in a high-level fusion architecture for cooperative perception systems. In *2012 IEEE Intelligent Vehicles Symposium*, pages 270–275. IEEE, 2012. 2
- [35] Andrea Simonelli, Samuel Rota Buló, Lorenzo Porzi, Manuel López-Antequera, and Peter Kotschieder. Disentangling monocular 3d object detection. In *ICCV*, pages 1991–1999, 2019. 6
- [36] Jakub Sochor, Jakub Špaňhel, and Adam Herout. Boxcars: Improving fine-grained recognition of vehicles using 3-d bounding boxes in traffic surveillance. *IEEE transactions on intelligent transportation systems*, 20(1):97–108, 2018. 2, 3
- [37] Xibin Song, Peng Wang, Dingfu Zhou, Rui Zhu, Chenye Guan, Yuchao Dai, Hao Su, Hongdong Li, and Ruiqiang Yang. ApolloCar3d: A large 3d car instance understanding benchmark for autonomous driving. In *CVPR*, pages 5452–5462, 2019. 2
- [38] Pei Sun, Henrik Kretschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for autonomous driving: Waymo open dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2446–2454, 2020. 2, 3, 6
- [39] Xiao Tan, Zhigang Wang, Minyue Jiang, Xipeng Yang, Jian Wang, Yuan Gao, Xiangbo Su, Xiaoqing Ye, Yuchen Yuan, Dongliang He, et al. Multi-camera vehicle tracking and re-identification based on visual and spatial-temporal features. In *CVPR Workshops*, pages 275–284, 2019. 2
- [40] Tai Wang, Xinge Zhu, Jiangmiao Pang, and Dahua Lin. Fcos3d: Fully convolutional one-stage monocular 3d object detection. *arXiv preprint arXiv:2104.10956*, 2021. 3
- [41] Tsun-Hsuan Wang, Sivabalan Manivasagam, Ming Liang, Bin Yang, Wenyuan Zeng, and Raquel Urtasun. V2vnet: Vehicle-to-vehicle communication for joint perception and prediction. In *European Conference on Computer Vision*, pages 605–621. Springer, 2020. 2
- [42] Yan Wang, Wei-Lun Chao, Divyansh Garg, Bharath Hariharan, Mark Campbell, and Kilian Q Weinberger. Pseudo-lidar from visual depth estimation: Bridging the gap in 3d object detection for autonomous driving. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8445–8453, 2019. 3
- [43] Xinshuo Weng and Kris Kitani. Monocular 3d object detection with pseudo-lidar point cloud. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 0–0, 2019. 3
- [44] Yu Xiang, Roozbeh Mottaghi, and Silvio Savarese. Beyond pascal: A benchmark for 3d object detection in the wild. In *IEEE winter conference on applications of computer vision*, pages 75–82. IEEE, 2014. 2
- [45] Bin Xu and Zhenzhong Chen. Multi-level fusion based 3d object detection from monocular images. In *CVPR*, 2018. 3
- [46] Xiaoqing Ye, Liang Du, Yifeng Shi, Yingying Li, Xiao Tan, Jianfeng Feng, Errui Ding, and Shilei Wen. Monocular 3d object detection via feature domain adaptation. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IX 16*, pages 17–34. Springer, 2020. 3
- [47] Yurong You, Yan Wang, Wei-Lun Chao, Divyansh Garg, Geoff Pleiss, Bharath Hariharan, Mark Campbell, and Kilian Q Weinberger. Pseudo-lidar++: Accurate depth for 3d object detection in autonomous driving. *arXiv preprint arXiv:1906.06310*, 2019. 3
- [48] Fisher Yu, Dequan Wang, Evan Shelhamer, and Trevor Darrell. Deep layer aggregation. In *CVPR*, pages 2403–2412, 2018. 7
- [49] Haibao Yu, Yizhen Luo, Mao Shu, Yiyi Huo, Zebang Yang, Yifeng Shi, Zhenglong Guo, Hanyu Li, Xing Hu, Jirui Yuan, and Zaiqing Nie. Dair-v2x: A large-scale dataset for vehicle-infrastructure cooperative 3d object detection. In *IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, June 2022. 2
- [50] Yunpeng Zhang, Jiwen Lu, and Jie Zhou. Objects are different: Flexible monocular 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3289–3298, 2021. 3, 7, 8
- [51] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. *arXiv preprint arXiv:1904.07850*, 2019. 7