

# MonoRUN: Monocular 3D Object Detection by Reconstruction and Uncertainty Propagation

Hansheng Chen, Yuyao Huang, Wei Tian,\* Zhong Gao, Lu Xiong

Institute of Intelligent Vehicles, School of Automotive Studies, Tongji University

hanshengchen97@gmail.com {huangyuyao, tian-wei, 1931604, xiong\_lu}@tongji.edu.cn

## Abstract

Object localization in 3D space is a challenging aspect in monocular 3D object detection. Recent advances in 6DoF pose estimation have shown that predicting dense 2D-3D correspondence maps between image and object 3D model and then estimating object pose via Perspective-n-Point (PnP) algorithm can achieve remarkable localization accuracy. Yet these methods rely on training with ground truth of object geometry, which is difficult to acquire in real outdoor scenes. To address this issue, we propose MonoRUN, a novel detection framework that learns dense correspondences and geometry in a self-supervised manner, with simple 3D bounding box annotations. To regress the pixel-related 3D object coordinates, we employ a regional reconstruction network with uncertainty awareness. For self-supervised training, the predicted 3D coordinates are projected back to the image plane. A Robust KL loss is proposed to minimize the uncertainty-weighted reprojected error. During testing phase, we exploit the network uncertainty by propagating it through all downstream modules. More specifically, the uncertainty-driven PnP algorithm is leveraged to estimate object pose and its covariance. Extensive experiments demonstrate that our proposed approach outperforms current state-of-the-art methods on KITTI benchmark.<sup>1</sup>

## 1. Introduction

Monocular 3D object detection is an active research area in computer vision. Although deep-learning-based 2D object detection has achieved remarkable progress [3, 30], the 3D counterpart still poses a much greater challenge on accurate object localization, since a single image cannot provide explicit depth information. To address this issue, a large number of works leverage geometrical priors and solve the object pose (position and orientation in camera frame) via

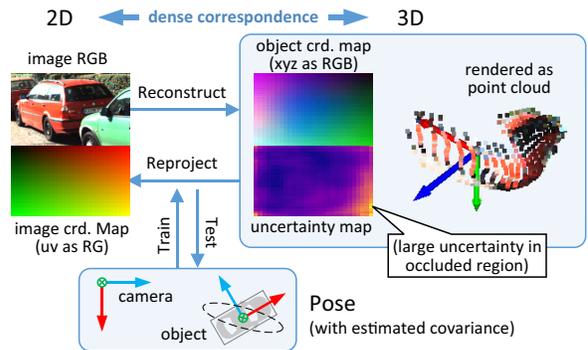


Figure 1. 3D reconstruction is conducted by regressing the pixel-related object coordinate map, which can be visualized as local point cloud in object space. For self-supervision, the object coordinates are reprojected to recover the image coordinate map. To focus on foreground pixels, we estimate the aleatoric uncertainty of network prediction. The coordinate uncertainty can be further propagated to estimate the pose covariance.

2D-3D constraints. These constraints either require extra keypoint annotations [4, 13], or exploit centers, corners and edges of ground truth bounding boxes [21, 27]. Yet the accuracy largely depends on the number and quality of available constraints, and the performance degrades in occlusion and truncation cases with fewer visible keypoints. A more robust approach is using dense 2D-3D correspondences, in which every single foreground pixel is mapped to a 3D point in the object space. This has proven successful in monocular 6DoF pose estimation tasks [15].

Current state-of-the-art dense correspondence methods [22, 28, 40] rely on both ground truth pose and 3D object model, so that target 3D coordinate map and object mask can be rendered for training supervision. This requirement restricts the training data to synthetic or simple laboratory scenes, where exact or pre-reconstructed 3D models are readily available. However, 3D object detection in real scenes (*e.g.*, driving scenes) mostly deals with category-level objects, where acquiring accurate 3D models for all instances is impractical. An intuitive idea could be

\*Corresponding author: Wei Tian.

<sup>1</sup>Code: <https://github.com/tjiiiv-cprg/MonoRUN>.

using LiDAR points to generate sparse coordinate maps for supervision. However, the persisting challenge is the deficiency of LiDAR points on specific objects or parts, *e.g.*, on distant objects or reflective materials.

A workaround to the lack of ground truth is leveraging self-supervision. Typically, Wang *et al.* [35] adopted a self-supervised network to directly learn object pose with the given ground truth 3D geometry. Our work, on the contrary, adopts the opposite idea: learning the 3D geometry from ground truth pose in a self-supervised manner during training, and then solving the object pose via 2D-3D correspondences during testing.

In this paper, we propose the *MonoRUn*, a novel **monocular** 3D object detection method using self-supervised reconstruction with **uncertainty** propagation. MonoRUn can extend off-the-shelf 2D detectors by appending a 3D branch to the region of interest (RoI) within each predicted 2D box. The 3D branch regresses dense 3D object coordinates in the RoI, which effectively builds up the geometry and 2D-3D correspondences.

To overcome the need for supervised foreground segmentation, we estimate the uncertainty of the predicted coordinates and adopt an uncertainty-driven PnP algorithm, which focuses on the low-uncertainty foreground. Furthermore, by forward propagating the uncertainty through PnP module, we can estimate the pose covariance matrix, which is used for scoring the detection confidence.

Self-supervision is conducted by projecting the predicted 3D coordinates back to the image via ground truth object pose and camera intrinsic parameters. To minimize the reprojection error with uncertainty awareness, we propose the *Robust KL loss* that minimizes the KL divergence between the predicted Gaussian distribution and the ground truth Dirac distribution. This novel loss function is the key to the state-of-the-art performance for the MonoRUn network.

To summarize, our main contributions are as follows:

- We propose a novel monocular 3D object detection network with uncertainty awareness, which can be trained without extra annotations (*e.g.*, keypoints, 3D models, masks). To the best of our knowledge, this is the first dense correspondence method employed for 3D detection in real driving scenes.
- We propose the Robust KL loss for general deep regression with uncertainty awareness, and demonstrate its superiority over the plain KL loss in previous work.
- Extensive evaluation on KITTI [10] benchmark shows significant improvement of our approach compared to current state-of-the-art methods.

## 2. Related Work

**Monocular 3D Object Detection** The majority of previous methods can be roughly divided into the following two

categories, based on how depth information is derived.

**1) With off-the-shelf monocular depth estimators.** Representatively, Pseudo-LiDAR [39] converts the estimated depth map to 3D point cloud and takes advantage of existing LiDAR-based 3D object detection pipeline. D<sup>4</sup>LCN [8] uses depth map as guidance to generate dynamic depth-wise convolutional filters, which can extract 3D information from RGB image more effectively. These methods largely benefit from the pre-trained depth estimator, *e.g.*, DORN [9], which may have generalization issues.

**2) With 2D-3D geometrical constraints.** DeepMANTA [4] annotates the training data with 36-keypoint 3D vehicle templates. A network is trained to find the best-matched template and regress 2D keypoint coordinates, and vehicle pose is computed via EPnP [20] algorithm. RTM3D [21] detects virtual keypoints (corners and center of the 3D bounding box) using a CenterNet-like [41] network. Another commonly used constraint is 2D box and 3D box consistency, first used by Mousavian *et al.* [27]. The above methods suffer from constraint deficiency by truncation or occlusion.

**Dense Correspondence and 3D Reconstruction** Most existing work uses ground truth geometry to train deep correspondence mapping networks. Nevertheless, some have explored end-to-end training via differentiable PnP algorithm without ground truth geometry.

**1) With geometry supervision.** Pix2Pose [28] directly regresses normalized object coordinates (NOC) of each object pixel. DPOD [40] predicts two-channel UV correspondences that map the object surface to 3D coordinates. Regarding category-level objects, Wang *et al.* [36] demonstrated that the scale-invariant NOC can handle unseen instances in a given category. These methods are only tested on synthetic or simple indoor data.

**2) Without geometry supervision.** Brachmann and Rother [1] proposed an approximate PnP back-propagation method to train an end-to-end network for Structure from Motion (SfM) problem. A further development is the BPnP [5], which is an exact PnP back-propagation approach. However, both methods train the networks in conjunction with reprojection loss as regularization term, which is essentially self-supervision. The potential of self-supervision alone is not investigated in these researches.

**Uncertainty Estimation** The uncertainty in deep learning consists of *aleatoric* and *epistemic* uncertainty [18]. The former captures noise inherent in observations, while the latter represents the uncertainty of model parameters. Kendall and Gal [18] introduced heteroscedastic regression for deep networks to directly output data-dependent aleatoric uncertainty, which can be learned with a *KL loss* [14, 18] function. Yet the plain KL loss is sensitive to outliers and is not well-balanced against other loss functions, leaving room for improvement.

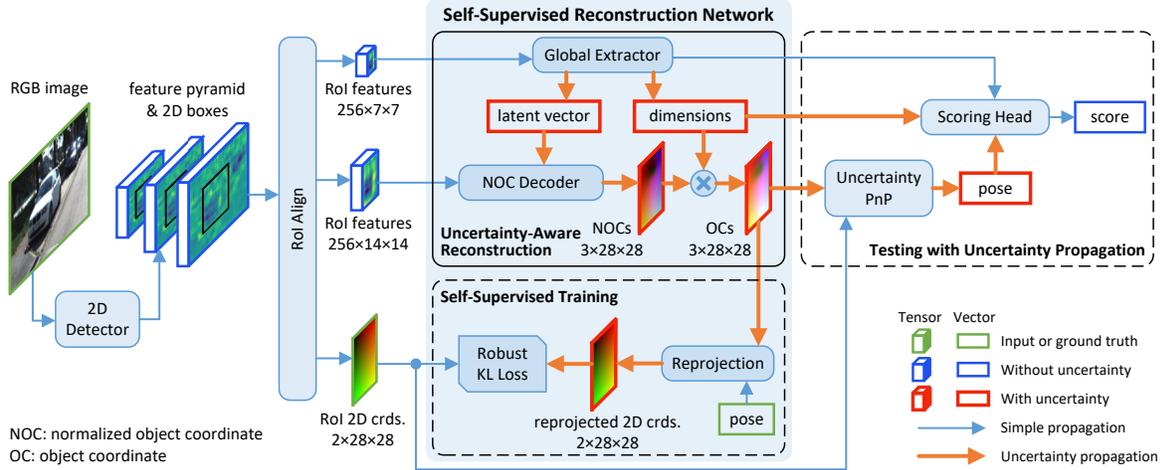


Figure 2. **Training and testing pipeline of MonoRUN.** The uncertainty-aware variables in red are modeled by Monte Carlo approach or by probabilistic models (e.g., Gaussian models).

### 3. Proposed Approach

#### 3.1. Problem Formulation and Approach Overview

Given an RGB image, the aim of 3D object detection is to localize and classify all objects of interest, yielding a 3D bounding box with class label for each object. The 3D box can be parameterized with dimensions  $\mathbf{d} = [l \ h \ w]^T$  and pose  $\mathbf{p} = [\beta \ t_x \ t_y \ t_z]^T$ , where  $\beta$  is the object yaw angle and  $t_x, t_y, t_z$  is the bottom center of the box in camera coordinate system.

Based on off-the-shelf 2D object detectors, we aim to predict a 3D object coordinate map using the RoI features within each 2D box. For self-supervision, given the ground truth pose  $\mathbf{p}_{gt}$  and camera model, we can project the 3D coordinates back to the image, obtaining the reprojected 2D coordinates  $(u_{rp}, v_{rp})$ . The objective is to recover the original 2D coordinates  $(u, v)$ . However, if we simply minimize the reprojection error without foreground segmentation, the network will be disrupted by large errors (hence large uncertainty) on irrelevant background points. Therefore, we design an *uncertainty-aware reconstruction* module that estimates the aleatoric uncertainty of  $(u_{rp}, v_{rp})$ , and the network is optimized by minimizing the uncertainty-weighted reprojection error using the proposed Robust KL loss.

During inference, we adopt the *uncertainty-driven PnP* module, through which the network uncertainty is propagated to object pose, represented with a multivariate normal distribution. This distribution is used by the scoring head to compute the detection confidence.

#### 3.2. Self-Supervised Reconstruction Network

To deal with category-level objects of various sizes, we employ two network branches to predict the 3D dimensions and the dimension-invariant normalized object coordinates

(NOC) [36] respectively. Then, the object coordinate vector  $\mathbf{x}^{OC}$  is the element-wise product (denoted by  $\circ$ ) of NOC vector  $\mathbf{x}^{NOC}$  and dimension vector  $\mathbf{d}$ :

$$\mathbf{x}^{OC} = \mathbf{x}^{NOC} \circ \mathbf{d}. \quad (1)$$

The first branch is called the *global extractor*, which extracts the global understanding of an object and predicts the 3D dimensions. The second branch is called the *NOC decoder*, which predicts the dense NOC map using convolutional layers. Since convolutional layers have limited capability of understanding the global context, we let the global extractor predict an additional global latent vector to enhance the NOC decoder. This latent vector potentially encodes the occlusion, truncation and shape of the object, and is found to be beneficial for aleatoric uncertainty estimation (as shown in Section 4.4). Details about the network are presented as follows.

**Global Extractor** As shown in Fig. 2, a  $7 \times 7$  RoI feature map is extracted from a higher level of the feature pyramid. The features are then flattened and fed into the global extraction head, which extracts a 16-channel global latent vector and predicts the 3D dimensions. The dimensions can be directly supervised by the annotated 3D bounding box sizes. For this network branch, we adopt two 1024-channel fully connected layers, as shown in Fig. 3.

**NOC Decoder** The NOC decoder network is designed to aggregate the global latent vector and local convolutional features for NOC prediction. This is realized by incorporating the *Excitation* operation from the Squeeze-Excitation Network [16]. As shown in Fig. 4, the channel size of the latent vector is first expanded to 256. Then, a channel-wise addition is conducted before the upsampling layer. Apart from predicting the three-channel NOC map, the NOC de-

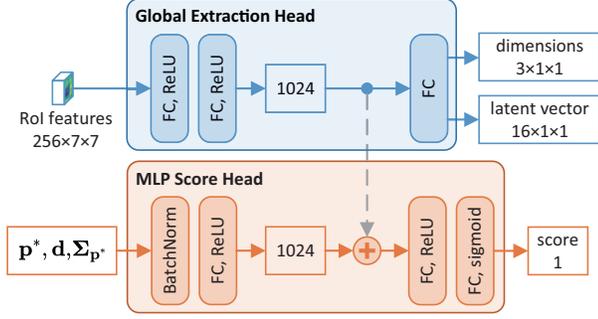


Figure 3. The **global extraction head** predicts the dimensions and latent vector, while the **MLP score head** exploits both the detection results and the global feature to predict the detection score.

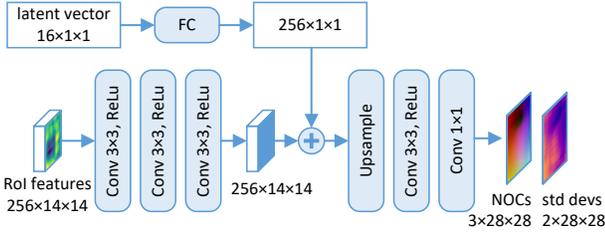


Figure 4. The **NOC decoder network** aggregates the local features and global embeddings by channel-wise addition, and outputs the dense NOCs as well as standard deviations of reprojected coordinates. For upsampling, we adopt the CARAFE layer [37], which is more efficient than the deconvolutional layer often used in instance segmentation.

coder network is also responsible for estimating the two-channel aleatoric uncertainty map, as explained in the next paragraph.

**Self-Supervision with Aleatoric Uncertainty** Given the ground truth of object pose, the predicted object coordinates can be projected back to the image plane. The reprojection error of the pixel  $(u, v)$  is formulated as:

$$\mathbf{r}_{(u,v)} = K(\mathbf{R}\mathbf{x}_{(u,v)}^{\text{OC}} + \mathbf{t}) - \begin{bmatrix} u \\ v \end{bmatrix} \quad (2)$$

with the camera projection function  $K(\cdot)$ , rotation matrix  $\mathbf{R}$ , and translation vector  $\mathbf{t} = [t_x \ t_y \ t_z]^T$ . To focus on minimizing foreground error without instance segmentation, we introduce the aleatoric uncertainty in our work. More specifically, we use univariate Gaussian distribution to model the reprojected 2D coordinates, and let the network predict the means and standard deviations, which are optimized using the Robust KL loss. Formally we could follow the uncertainty propagation path in Fig. 2 and predict the uncertainty of intermediate variables at first. Practically, leveraging the flexibility of deep networks, we take a shortcut by letting the NOC decoder directly regress the standard deviations of reprojected 2D coordinates, as shown in Fig. 4.

**Additional Epistemic Uncertainty** Estimating the epistemic uncertainty is essential for safety-critical applications such as autonomous driving. Following [18], we compute the mean and variance of  $\mathbf{x}^{\text{OC}}$  using Monte Carlo dropout during inference. We insert a channel dropout layer [34] after RoI Align [11] and a 1D dropout layer after each FC layer. Since PnP algorithms handle 2D projection variances more efficiently, we first transform the 3D variances of object coordinates into 2D variances of reprojected coordinates, which is then combined with the aleatoric uncertainty. Details are shown in the supplementary materials.

### 3.3. Robust KL Loss

By definition, KL loss is derived from the KL divergence between the predicted distribution and target distribution. Assuming Gaussian priors, the KL divergence is as follows:

$$D_{\text{KL}}(\mathcal{N}_{\text{tgt}} \parallel \mathcal{N}_{\text{pred}}) = \frac{1}{2} \left( \frac{\sigma_{\text{tgt}}^2}{\sigma_{\text{pred}}^2} + \frac{(\mu_{\text{pred}} - \mu_{\text{tgt}})^2}{\sigma_{\text{pred}}^2} - 1 + \log \frac{\sigma_{\text{pred}}^2}{\sigma_{\text{tgt}}^2} \right). \quad (3)$$

For fixed target distribution,  $\log \sigma_{\text{tgt}}^2$  is constant and can be omitted in the minimization. Assuming narrow target (Dirac-like),  $(\mu_{\text{pred}} - \mu_{\text{tgt}})^2$  dominates much more than  $\sigma_{\text{tgt}}^2$ . Let  $y = \mu_{\text{tgt}}$ , the minimization objective is simplified as:

$$L_{\text{KL}} = \frac{1}{2\sigma_{\text{pred}}^2} (\mu_{\text{pred}} - y)^2 + \frac{1}{2} \log \sigma_{\text{pred}}^2. \quad (4)$$

We call Eq. (4) the Gaussian KL loss. Hereafter we omit the subscript  $(\cdot)_{\text{pred}}$  for brevity. To capture heteroscedastic aleatoric uncertainty in regression, Kendall and Gal [18] proposed to directly predict the data-dependent mean  $\mu$  and log variance  $\log \sigma^2$  using deep networks, which is optimized by Eq. (4). Apparently, the first term in Eq. (4) is a weighted L2 loss, where errors with higher uncertainty are less punished.

Despite its probabilistic origin, the Gaussian KL loss has two defects when applied to a deep regression model:

- As a generalization form of L2 loss, the Gaussian KL loss is not robust to outliers;
- The gradient w.r.t.  $\mu$  tends to increase as the denominator  $2\sigma^2$  decays during training, whereas regular L2 or L1 losses have decreasing or constant gradient, leading to loss imbalance in multi-task learning.

Regarding the first problem, an alternative KL loss derived from Laplacian distribution is used in work [7, 18]:

$$L_{\text{LapKL}} = \frac{\sqrt{2}}{\sigma} |\mu - y| + \log \sigma. \quad (5)$$

Same as the L1 loss, this function is not differentiable at  $\mu = y$ . To overcome this issue, we design a mixed KL loss,

written as a function of weighted error  $e = (\mu - y)/\sigma$  and standard deviation  $\sigma$ :

$$L_{\text{mKL}} = \begin{cases} \frac{1}{2}e^2 + \log \sigma, & |e| \leq \sqrt{2}, \\ \sqrt{2}|e| - 1 + \log \sigma, & |e| > \sqrt{2}. \end{cases} \quad (6)$$

It can be verified that this function is differentiable w.r.t. both  $\mu$  and  $\sigma$  (on condition that  $\sigma > 0$ ). The mixed KL loss can be regarded as an extension of Huber loss (smooth L1), which is robust to outliers and easier to optimize.

The second problem is caused by the increasing weight  $1/\sigma$  as the denominator  $\sigma$  decays during training. This can be mitigated by normalizing the weight such that its batch-wise mean equals unity. Inspired by Batch Normalization [17], we perform online estimate of the mean weight  $\hat{w} = \mathbb{E}[1/\sigma]$  using exponential moving average:

$$\hat{w} \leftarrow \alpha \hat{w} + (1 - \alpha) \frac{1}{N} \sum_{i=1}^N \frac{1}{\sigma_i}, \quad (7)$$

where  $\alpha$  is the momentum,  $N$  is the number of samples in a batch. The finalized Robust KL loss is simply the weight-normalized mixed KL loss:

$$L_{\text{RKL}} = \frac{1}{\hat{w}} L_{\text{mKL}}. \quad (8)$$

In practice, directly optimizing  $\sigma$  can lead to gradient explosion. Therefore, we let the network predict the logarithmic standard deviation  $\log \sigma$  as an alternative.

### 3.4. Uncertainty-Driven PnP

**Maximum Likelihood Estimation** To solve PnP with uncertainty is to perform the maximum likelihood estimation (MLE) of pose  $\mathbf{p}$ , in which the negative log likelihood (NLL) function is the sum of squared reprojection error  $\mathbf{r}_{(u,v)}$  measured in Mahalanobis distance:

$$\mathbf{p}^* = \arg \min_{\mathbf{p}} \frac{1}{2} \sum_{(u,v) \in \text{RoI}} \mathbf{r}_{(u,v)}^T \Sigma_{(u,v)}^{-1} \mathbf{r}_{(u,v)} \quad (9)$$

with  $\Sigma_{(u,v)} = \text{diag}[\sigma_{u_p}^2, \sigma_{v_p}^2] |_{(u,v)}$ , where  $\sigma_{u_p}, \sigma_{v_p}$  are the predicted standard deviations of the reprojected coordinates. This minimization can be solved efficiently using the Levenberg-Marquardt algorithm.

**Covariance Estimation** For probabilistic object localization, we also need to estimate the covariance of  $\mathbf{p}^*$ , which can be approximated by the inverse Hessian of the NLL at  $\mathbf{p}^*$  [29]:

$$\text{Cov}[\mathbf{p}^*] \approx \mathbf{H}(\mathbf{p}^*)^{-1}. \quad (10)$$

To avoid computing the second derivative during inference, we approximate the exact Hessian using the Gauss-Newton matrix  $\mathbf{J}(\mathbf{p}^*)^T \mathbf{J}(\mathbf{p}^*)$ , with  $\mathbf{J}(\mathbf{p}^*) = \partial \mathbf{r}_{\text{all}} / \partial \mathbf{p}^T |_{\mathbf{p}^*}$ , where  $\mathbf{r}_{\text{all}} = [r_1/\sigma_1 \quad r_2/\sigma_2 \quad \dots \quad r_{2n}/\sigma_{2n}]^T$  (flattened vector of all weighted reprojection errors).

**Online Covariance Calibration** In practice, using Eq. (10) can result in a much lower covariance than the actual one. This is mainly because Eq. (9) assumes that the reprojection error of each point is independent, whereas the network outputs are apparently correlated. Therefore, we further conduct online covariance calibration using a 4x1 learnable calibration vector  $\mathbf{k}$ :

$$\Sigma_{\mathbf{p}^*} = \exp(\text{diag } \mathbf{k}) (\mathbf{J}(\mathbf{p}^*)^T \mathbf{J}(\mathbf{p}^*))^{-1} \exp(\text{diag } \mathbf{k}). \quad (11)$$

The calibration vector can be learned by applying the multivariate Gaussian KL loss:

$$L_{\text{calib}} = \frac{1}{2} (\mathbf{p}^* - \mathbf{p}_{\text{gt}})^T \Sigma_{\mathbf{p}^*}^{-1} (\mathbf{p}^* - \mathbf{p}_{\text{gt}}) + \frac{1}{2} \log \det \Sigma_{\mathbf{p}^*}, \quad (12)$$

where  $\mathbf{p}^*$  is detached and only the calibration vector is optimized. Despite the defects of Gaussian KL loss stated in Section 3.3, it is sufficient for the simple calibration task.

### 3.5. Scoring Head

The confidence score of a detected object, denoted by  $c$ , can be decomposed into localization score  $\Pr[Fg]$  (the probability of detecting a foreground object) and classification score  $\Pr[Cls]$  (the probability of predicting the correct class label). For 3D object detection, the score can be expressed as the product of 3D localization conditional probability and 2D score:

$$c_{3D} = \underbrace{\Pr[Fg_{3D}|Fg_{2D}]}_{c_{3DLoc}} \underbrace{\Pr[Fg_{2D}]}_{c_{2D}} \Pr[Cls]. \quad (13)$$

The 2D score  $c_{2D}$  is given by the 2D detection module. Thus, we only need to predict the 3D localization score  $c_{3DLoc}$ . Since the 3D branch is trained only with positive 2D samples, the predicted  $c_{3DLoc}$  is naturally conditional.

By sampling object poses from the estimated distribution  $\mathcal{N}(\mathbf{p}^*, \text{Cov}[\mathbf{p}^*])$ , the 3D localization score can be computed via Monte Carlo integration w.r.t. 3D IoU, as we show in the supplementary material. Due to slow 3D IoU calculation, Monte Carlo scoring has adverse effects on inference time. Therefore, we practically adopt the multi-layer perceptron (MLP) approach (Fig. 3), which is faster, end-to-end trainable, and capable of fusing both pose uncertainty and network features to predict a more reliable score. To train the MLP scoring branch, we use the same binary cross-entropy loss as in [31, 32]:

$$L_{\text{score}} = -c_{\text{tgt}} \log c_{3DLoc} - (1 - c_{\text{tgt}}) \log(1 - c_{3DLoc}), \quad (14)$$

where  $c_{3DLoc}$  is the output of MLP, which is bounded to 0~1 by logistic activation function,  $c_{\text{tgt}}$  is a clamped linear function w.r.t. the 3D IoU between prediction and ground truth:

$$c_{\text{tgt}} = \max(0, \min(1, 2IoU_{3D} - 0.5)). \quad (15)$$

The performance comparison between Monte Carlo and MLP is presented in the supplementary materials.

### 3.6. Network Training

The proposed MonoRUn network can be trained with three different setups, which are compared in the experiments.

**Fully Self-Supervised Reconstruction (Without Extra Supervision)** In this mode, neither the LiDAR point supervision nor the end-to-end PnP is used. The 3D reconstruction process is trained in a fully self-supervised manner, except that 3D dimensions are directly supervised. The overall loss function is formulated as:

$$L_{\text{self}} = L_{2D} + L_{\text{proj}} + L_{\text{dim}} + L_{\text{score}} + \lambda L_{\text{calib}}, \quad (16)$$

where  $L_{2D}$  is the 2D detection loss,  $L_{\text{proj}}$  is a Robust KL loss on self-supervised reprojection error,  $L_{\text{dim}}$  is a smooth L1 loss on dimension error, and  $\lambda$  is a hyperparameter for calibration loss, which is set to 0.01.

**LiDAR Supervision** Direct NOC loss can be imposed by converting foreground LiDAR points into sparse ground truth of NOC map. In this case, the aleatoric uncertainty is unnecessary. Thus, we adopt the weighted smooth L1 loss:

$$L_{\text{NOC}} = \frac{1}{\sum_i w_i} \sum_i w_i L_{\text{SmoothL1}}(x_{i,\text{pred}}^{\text{NOC}} - x_{i,\text{LiDAR}}^{\text{NOC}}), \quad (17)$$

where  $x_i^{\text{NOC}}$  denotes the  $i$ -th element of the NOC tensor,  $w_i$  equals 1 where  $x_{i,\text{LiDAR}}^{\text{NOC}}$  is available and 0 elsewhere. The overall loss becomes:

$$L_{\text{LiDAR}} = L_{\text{self}} + L_{\text{NOC}}. \quad (18)$$

Without specific statement, we train all the models using this setup.

**End-to-End Training** Incorporating the PnP back-propagation approach in [5], we apply smooth L1 loss on the Euclidean errors of estimated translation vector and yaw angle. Details of the PnP derivatives and loss functions are presented in the supplementary materials. Since end-to-end training is unstable at the beginning, we use a similar training protocol to [1], *i.e.*, applying end-to-end training as refinement after self-supervised training. **This setup is only investigated in ablation studies for pure comparison.**

## 4. Experiments

### 4.1. Dataset

We evaluate the proposed model on the KITTI-Object benchmark [10]. It consists of 7481 training images and 7518 test images as well as the corresponding point clouds, comprising a total of 80256 labeled objects in eight classes. Each object is assigned to one of three difficulty levels according to truncation, occlusion and 2D box height. The training images are further split into 3712 training and 3769

validation images [6]. The official benchmark evaluates detection performance on three classes: *Car*, *Pedestrian* and *Cyclist*. Evaluation metrics are based on precision-recall curves with 3D IoU threshold of 0.7 or 0.5. We adopt the official metric that computes the 40-point interpolated average precision (AP) [33].

### 4.2. Implementation Details

**2D Detector** We use pre-trained Faster R-CNN [30] with ResNet-101 [12] as backbone. We adopt a six-level feature pyramid network [23], in which an additional upsampled level is added.

**Reconstruction Module** We set the dropout rate to 0.5 for 1D dropout layers and 0.2 for channel dropout layers. Network outputs (dimensions, NOCs) are normalized w.r.t. the mean and standard deviation calculated from training data. When training with multiple classes, we predict a set of class-specific latent vectors, dimensions and NOCs.

**Data Augmentation** During training, we apply random flip and photometric distortion augmentation. We set two NOC decoder branches at the last  $1 \times 1$  convolutional layer for original and mirrored objects respectively.

**Training Schedule** The network is trained by the AdamW [24] optimizer with a weight decay of 0.01. We take a batch size of 6 on two Nvidia RTX 2080 Ti GPUs, and train the network using cosine learning rate decay with a base learning rate of 0.0002. Total epochs are 32 for the full training set and 50 for the split training set. For end-to-end training, we append a second cycle of 15 epochs with a reduced base learning rate of 0.00003.

**Testing Configuration** For epistemic uncertainty, we use 50 Monte Carlo dropout samples as in [18]. By default, we only sample the global extractor to estimate the dimension uncertainty (see discussions in Section 4.4). During post-processing, we use 3D non-maximum suppression (NMS) with an IoU threshold of 0.01.

### 4.3. Comparison to the State of the Art

We evaluate two variations of our model (fully self-supervised and plus LiDAR supervision) on the official test set and the validation split. Table 1 lists the top methods from the official leaderboard. We observe that: (1) When trained **with LiDAR supervision**, our method outperforms the state of the art by a wide margin. Note that the top three competitors also use extra supervision and even extra data (by using depth estimators pre-trained on the much larger KITTI-Depth dataset). (2) When trained **without extra supervision**, our method still outperforms the non-depth methods (RTM3D [21], MonoPair [7]) on the official test set. (3) Our method achieves state-of-the-art accuracy within a reasonable runtime (0.070 s, including Monte Carlo and PnP), whereas the top three competitors spend

Method	Test AP <sub>IoU≥0.7</sub>			Val AP <sub>IoU≥0.5</sub>			Val AP <sub>IoU≥0.7</sub>			Time (sec)
	Easy	Mod.	Hard	Easy	Mod.	Hard	Easy	Mod.	Hard	
M3D-RPN [2]	14.76	9.71	7.42	48.53	35.94	28.59	14.53	11.07	8.65	0.16
MonoPair [7]	13.04	9.99	8.65	55.38	42.39	37.99	16.28	12.30	10.42	0.057
RTM3D [21]	14.41	10.34	8.77	-	-	-	-	-	-	0.055
AM3D* [26]	16.50	10.74	9.52	-	-	-	28.31	15.76	12.24	0.4
PatchNet* [25]	15.68	11.12	10.17	-	-	-	31.6	16.8	13.8	0.4
D <sup>4</sup> LCN* [8]	16.65	11.72	9.51	-	-	-	22.32	16.20	12.30	0.2
Ours (w/o extra supv.)	16.04	10.53	9.11	55.88	40.03	35.59	17.26	12.27	10.41	0.070
Ours (+ LiDAR supv.)	19.65	12.30	10.58	59.71	43.39	38.44	20.02	14.65	12.61	0.070

Table 1. **3D detection performance of Car category** on KITTI official test set and validation set. The 1<sup>st</sup>, 2<sup>nd</sup> and 3<sup>rd</sup> place are color coded. \* indicates using the pre-trained depth estimator DORN [9]. Wang *et al.* [38] pointed out that the training data of DORN overlaps with KITTI-Object validation data, causing the 3D detectors to overfit. We use gray to indicate the values affected by overfitting.

Method	Ped. AP <sub>IoU≥0.5</sub>			Cycl. AP <sub>IoU≥0.5</sub>		
	Easy	Mod.	Hard	Easy	Mod.	Hard
M3D-RPN [2]	4.92	3.48	2.94	0.94	0.65	0.47
MonoPSR [19]	6.12	4.00	3.30	8.37	4.74	3.68
MonoPair [7]	10.02	6.68	5.53	3.79	2.12	1.83
Ours	10.88	6.78	5.83	1.01	0.61	0.48

Table 2. **3D detection performance of Pedestrian and Cyclist categories** on KITTI official test set. Red indicates the best.

more than 0.2 s (not counting the 0.5 s depth estimation time of DORN [9]).

For Pedestrian and Cyclist, we present the detection performance in Table 2. Our method (with LiDAR supervision) achieves the best performance in Pedestrian detection, yet underperforms in Cyclist detection, presumably due to inadequate training samples.

#### 4.4. Ablation Studies

In this section, all models are trained and evaluated with the train/val split. We show the mean value of all six AP metrics (mAP) for 3D Car detection. All results are presented in Table 3.

**Self-Supervision versus LiDAR Supervision** While self-supervision alone can achieve the state-of-the-art performance (28.57), using only LiDAR supervision leads to very poor performance (18.84), which demonstrates the importance of self-supervision. Nonetheless, the overall best result is reached with both supervisions together (31.21). As revealed in Fig. 5b, the self-supervised geometry does not necessarily match the true surface and is therefore prone to overfitting, which can be alleviated by the shape-regularizing effect from LiDAR supervision.

**Robust KL Loss** By comparing the performance of smooth L1 loss (26.35), Laplacian KL loss (29.47), mixed KL loss (30.05) and Robust KL loss (31.21), we observe incremental improvements. The largest performance gap is between  $L_{SmoothL1}$  and  $L_{LapKL}$ , which reveals the contribution of aleatoric uncertainty. Moreover, both Gaussian

	$L_{proj}$ (Self)	$L_{NOC}$ (LiDAR)	E2E	Epistemic	Latent vector	mAP
$L_{RKL}$	✓				✓	31.21
$L_{RKL}$					✓	28.57
		✓			✓	18.84
$L_{SmoothL1}$	✓				✓	26.35
$L_{LapKL}$	✓				✓	29.47
$L_{mKL}$	✓				✓	30.05
$L_{LapKL}$	✓		✓		✓	29.73
$L_{RKL}$	✓		✓		✓	31.09
$L_{RKL}$	✓			dim	✓	31.47
$L_{RKL}$	✓			full	✓	31.16
$L_{RKL}$	✓					29.78

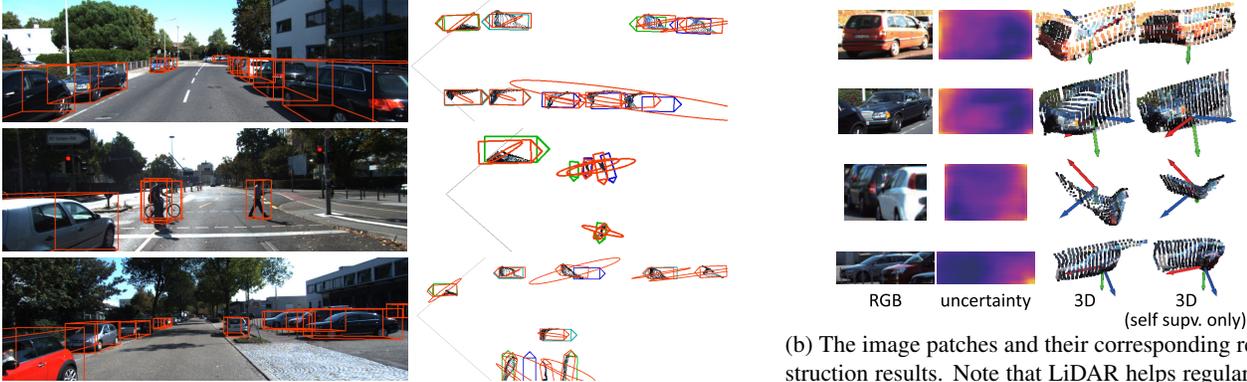
Table 3. Results of **ablation studies** on reprojection loss function, LiDAR supervision, end-to-end training, epistemic uncertainty and latent vector.

Laplacian mixture and weight normalization make important contributions to the Robust KL loss, totaling an mAP increase of 1.74 compared to the Laplacian KL loss.

**End-to-End Refinement** We observe that the performance of end-to-end refinement is strongly related to the baseline performance. For the baseline trained with Laplacian KL loss (29.47), end-to-end refinement boosts the mAP by 0.26 (29.73). For the baseline trained with Robust KL loss (31.21), however, end-to-end refinement slightly worsens the performance (31.09). This validates that self-supervised training with Robust KL loss can better optimize the network than end-to-end training via differentiable-PnP.

**Epistemic Uncertainty** Estimating the epistemic uncertainty of box dimensions alone shows improvement to the baseline (31.47 vs 31.21). When sampling the full reconstruction network, we observe adverse effect on the detection performance (31.16).

**Latent vector** When disabling the latent vector, the performance drops significantly (29.78 vs 31.21). To find out how the latent vector impacts the network performance, we evaluate the network sensitivity to the latent vector by reset-



(a) 3D bounding boxes in images and bird’s-eye views (BEV). Red indicates detected boxes (along with 95% confidence ellipse), ground truth boxes are color coded by their occlusion levels: **fully visible**, **partly occluded**, **difficult to see**.

Figure 5. **Visualization of detection and reconstruction results** on the KITTI validation set.

$L_{proj}$	mAP (original)	mAP (zero latent)	Diff
$L_{SmoothL1}$ (w/o aleatoric)	26.35	26.33	-0.02
$L_{RKL}$ (w/ aleatoric)	31.21	29.23	-1.98

Table 4. **Performance sensitivity to the latent vector.**

ting it to zero during inferring. As shown in Table 4, only the network with aleatoric uncertainty is sensitive to the latent vector. This implies that the latent vector encodes very important information for the estimation of aleatoric uncertainty.

#### 4.5. Reliability of the Localization Uncertainty

As the online covariance calibration is conducted on the training data, the model tends to be overconfident on the testing data. To assess the reliability of the localization uncertainty (w.r.t. the translation vector  $\mathbf{t}$ ), we discretize the ground truth object distance  $t_z$  into a number of bins. For the samples in each bin, we calculate the mean covariance matrix of prediction ( $\bar{\Sigma}_{\mathbf{t}^*}$ ), as well as the covariance matrix of the actual localization error ( $\text{Cov}[\mathbf{t}^* - \mathbf{t}_{gt}]$ ). Ideally, these two covariance matrices should be equal. For comparison, we calculate the Gaussian entropy of the two covariance matrices respectively:

$$H_{\text{pred}} = \frac{1}{2} \log \det(2\pi e \bar{\Sigma}_{\mathbf{t}^*}), \quad (19)$$

$$H_{\text{actual}} = \frac{1}{2} \log \det(2\pi e \text{Cov}[\mathbf{t}^* - \mathbf{t}_{gt}]). \quad (20)$$

As shown in Fig. 6, on the train split, the predicted uncertainty is very close to the actual error. This demonstrates the effectiveness of covariance calibration. For the unseen validation data, since epistemic sampling does not cover the full network, the model generally predicts overconfident results,

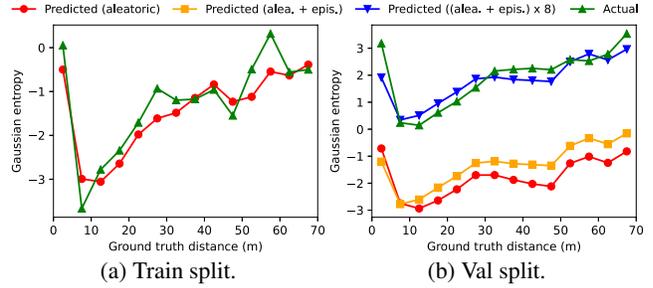


Figure 6. **Gaussian entropy of predicted and actual covariance matrices** vs object distance, tested on the Car category of train and val splits.

which can be roughly corrected by applying an empirical covariance scaling factor.

## 5. Conclusion

We presented the MonoRUN framework, a novel monocular 3D object detector with state-of-the-art performance and high practicality. To employ dense correspondence method for 3D detection in real driving scenes, we overcame the deficiency of geometry supervision by self-supervised reconstruction with uncertainty awareness. Meanwhile, we made uncertainty-aware deep regression networks easier to optimize by proposing the Robust KL loss. Finally, we are among the first to explore probabilistic 3D object localization by uncertainty propagation through PnP, which may open up new possibilities for downstream tasks such as robust tracking and motion prediction.

**Acknowledgements** This research was supported by the National Natural Science Foundation of China (Grant No. 52002285), the Shanghai Pujiang Program (Grant No. 2020PJD075) and the Shenzhen Future Intelligent Network Transportation System Industry Innovation Center (Grant No. 17092530321).

## References

- [1] Eric Brachmann and Carsten Rother. Learning less is more - 6d camera localization via 3d surface regression. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [2] Garrick Brazil and Xiaoming Liu. M3d-rpn: Monocular 3d region proposal network for object detection. In *IEEE International Conference on Computer Vision (ICCV)*, 2019.
- [3] Zhaowei Cai and Nuno Vasconcelos. Cascade r-cnn: Delving into high quality object detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [4] Florian Chabot, Mohamed Chaouch, Jaonary Rabarisoa, Céline Teulière, and Thierry Chateau. Deep manta: A coarse-to-fine many-task network for joint 2d and 3d vehicle analysis from monocular image. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1827–1836, 2017.
- [5] Bo Chen, Alvaro Parra, Jiewei Cao, Nan Li, and Tat-Jun Chin. End-to-end learnable geometric vision by backpropagating pnp optimization. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [6] Xiaozhi Chen, Kaustav Kundu, Yukun Zhu, Andrew G Berneshawi, Huimin Ma, Sanja Fidler, and Raquel Urtasun. 3d object proposals for accurate object class detection. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems (NIPS)*, pages 424–432. Curran Associates, Inc., 2015.
- [7] Yongjian Chen, Lei Tai, Kai Sun, and Mingyang Li. Monopair: Monocular 3d object detection using pairwise spatial relationships. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [8] Mingyu Ding, Yuqi Huo, Hongwei Yi, Zhe Wang, Jianping Shi, Zhiwu Lu, and Ping Luo. Learning depth-guided convolutions for monocular 3d object detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [9] Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, and Dacheng Tao. Deep Ordinal Regression Network for Monocular Depth Estimation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [10] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.
- [11] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *IEEE International Conference on Computer Vision (ICCV)*, pages 2980–2988, 2017.
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016.
- [13] Tong He and Stefano Soatto. Mono3d++: Monocular 3d vehicle detection with two-scale 3d hypotheses and task priors. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2019.
- [14] Yihui He, Chenchen Zhu, Jianren Wang, Marios Savvides, and Xiangyu Zhang. Bounding box regression with uncertainty for accurate object detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [15] Stefan Hinterstoisser, Vincent Lepetit, Slobodan Ilic, Stefan Holzer, Gary Bradski, Kurt Konolige, and Nassir Navab. Model based training, detection and pose estimation of texture-less 3d objects in heavily cluttered scenes. In *Asian Conference on Computer Vision (ACCV)*, 2012.
- [16] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [17] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 448–456, 2015.
- [18] Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? In *Conference on Neural Information Processing Systems (NIPS)*, 2017.
- [19] Jason Ku, Alex D. Pon, and Steven L. Waslander. Monocular 3d object detection leveraging accurate proposals and shape reconstruction. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11859–11868, 2019.
- [20] Vincent Lepetit, Francesc Moreno-Noguer, and Pascal Fua. Epnnp: An accurate o(n) solution to the pnp problem. *International Journal Of Computer Vision*, 81:155–166, 2009.
- [21] Peixuan Li, Huaici Zhao, Pengfei Liu, and Feidao Cao. Rtm3d: Real-time monocular 3d detection from object keypoints for autonomous driving. In *European Conference on Computer Vision (ECCV)*, 2020.
- [22] Zhigang Li, Gu Wang, and Xiangyang Ji. Cdpn: Coordinates-based disentangled pose network for real-time rgb-based 6-dof object pose estimation. In *IEEE International Conference on Computer Vision (ICCV)*, pages 7677–7686, 2019.
- [23] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 936–944, 2017.
- [24] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations (ICLR)*, 2019.
- [25] Xinzhu Ma, Shinan Liu, Zhiyi Xia, Hongwen Zhang, Xingyu Zeng, and Wanli Ouyang. Rethinking pseudo-lidar representation. In *European Conference on Computer Vision (ECCV)*, 2020.
- [26] 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 *IEEE International Conference on Computer Vision (ICCV)*, 2019.
- [27] Arsalan Mousavian, Dragomir Anguelov, John Flynn, and Jana Košecká. 3d bounding box estimation using deep learning and geometry. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.

- [28] Kiru Park, Timothy Patten, and Markus Vincze. Pix2pose: Pixel-wise coordinate regression of objects for 6d pose estimation. In *IEEE International Conference on Computer Vision (ICCV)*, 2019.
- [29] Yudi Pawitan. *In All Likelihood: Statistical Modelling and Inference Using Likelihood*. Oxford University Press, 2001.
- [30] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(6):1137–1149, 2017.
- [31] Shaoshuai Shi, Chaoxu Guo, Li Jiang, Zhe Wang, Jianping Shi, Xiaogang Wang, and Hongsheng Li. Pv-rcnn: Point-voxel feature set abstraction for 3d object detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10526–10535, 2020.
- [32] Shaoshuai Shi, Zhe Wang, Jianping Shi, Xiaogang Wang, and Hongsheng Li. From points to parts: 3d object detection from point cloud with part-aware and part-aggregation network. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- [33] Andrea Simonelli, Samuel Rota Bulò, Lorenzo Porzi, Manuel López-Antequera, and Peter Kotschieder. Disentangling monocular 3d object detection. In *IEEE International Conference on Computer Vision (ICCV)*, pages 1991–1999, 2019.
- [34] Jonathan Tompson, Ross Goroshin, Arjun Jain, Yann LeCun, and Christoph Bregler. Efficient object localization using convolutional networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 648–656, 2015.
- [35] Gu Wang, Fabian Manhardt, Jianzhun Shao, Xiangyang Ji, Nassir Navab, and Federico Tombari. Self6d: Self-supervised monocular 6d object pose estimation. In *European Conference on Computer Vision (ECCV)*, 2020.
- [36] He Wang, Srinath Sridhar, Jingwei Huang, Julien Valentin, Shuran Song, and Leonidas J. Guibas. Normalized object coordinate space for category-level 6d object pose and size estimation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [37] Jiaqi Wang, Kai Chen, Rui Xu, Ziwei Liu, Chen Change Loy, and Dahua Lin. Carafe: Content-aware reassembly of features. In *IEEE International Conference on Computer Vision (ICCV)*, pages 3007–3016, 2019.
- [38] Xinlong Wang, Wei Yin, Tao Kong, Yuning Jiang, Lei Li, and Chunhua Shen. Task-aware monocular depth estimation for 3d object detection. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2020.
- [39] Yan Wang, Wei-Lun Chao, Divyansh Garg, Bharath Hariharan, Mark Campbell, and Kilian Weinberger. Pseudo-lidar from visual depth estimation: Bridging the gap in 3d object detection for autonomous driving. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [40] Sergey Zakharov, Ivan Shugurov, and Slobodan Ilic. Dpod: 6d pose object detector and refiner. In *IEEE International Conference on Computer Vision (ICCV)*, 2019.
- [41] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. *CoRR*, abs/1904.07850, 2019.