End-to-end 6DoF Pose Estimation from Monocular RGB Images

Wenbin Zou, Di Wu, Shishun Tian, Canqun Xiang, Xia Li and Lu Zhang

Abstract—We present a conceptually simple framework for 6DoF object pose estimation, especially for autonomous driving scenarios. Our approach can efficiently detect the traffic participants from a monocular RGB image while simultaneously regressing their 3D translation and rotation vectors. The proposed method 6D-VNet, extends the Mask R-CNN by adding customised heads for predicting vehicle’s finer class, rotation and translation. It is trained end-to-end compared to previous methods. Furthermore, we show that the inclusion of translational regression in the joint losses is crucial for the 6DoF pose estimation task, where object translation distance along longitudinal axis varies significantly, *e.g.* in autonomous driving scenarios. Additionally, we incorporate the mutual information between traffic participants via a modified non-local block to capture the spatial dependencies among the detected objects. As opposed to the original non-local block implementation, the proposed weighting modification takes the spatial neighbouring information into consideration whilst counteracting the effect of extreme gradient values. We evaluate our method on the challenging real-world Pascal3D+ dataset and our 6D-VNet reaches the 1st place in ApolloScape challenge 3D Car Instance task [1], [2].

I. INTRODUCTION

The increase in consumer demand for safer vehicles greatly promote the development of the advanced driver-assistance systems (ADASs) and autonomous driving. Recently, vision-based systems [3] have been drawn extensive attention in autonomous driving and ADASs due to their great potential in roadway-environment understanding [4], such as traffic light detection [5], road semantic segmentation [6] *etc.*, of which one crucial component is to detect, estimate and reconstruct the 3D shape of vehicles directly from the captured RGB video *cf.* Fig. 1.

The 2D object detection has gained significant improvement in recent years thanks to the development of deep learning, while the estimation of object shape and pose in 3D remains a challenging problem. The current state-of-the-art RGB-based 6DoF pose estimation methods [7], [8], [9] are two-staged: the first stage is to detect the object with 3D rotation via a trained network, the second stage is to estimate the full 3D translation

Fig. 1: 6D-VNet is trained end-to-end to estimate vehicles’ six degree of freedom poses from a single monocular image. The output will precisely estimates the vehicle’s voxels in 3D occupancy.

via projective distance estimation. The aforementioned two-staged systems primarily focus on the industry-relevant bin-picking tasks. Typically, a robot needs to grasp a single arbitrary instance of the required object, *e.g.* a component such as a bolt or nut, and operate with it. In such scenario, the surface alignment in the Z dimension, *i.e.* the optical axis of the camera, is less important than the alignment in the X and Y dimensions. Such industrial setting requires accurate estimation of rotation, whereas the translation tolerance can be relaxed. However, in autonomous driving, translation distance of traffic participants along longitudinal axis varies significantly. The translation estimation is thus more challenging and the estimation of vehicle’s translation is more critical than that of orientation.

Traditional methods usually leave the translational estimation as a separate procedure after the object class prediction and rotation estimation by using a geometric projection method. However, the geometric projection method assumes that: (i) the object centre in 3D will be projected to the object bounding box centre in the 2D image; (ii) the predicted object class and rotation vector is correctly estimated. Therefore, by using geometric projection as a post-processing step, the error from object class estimation and rotation regression will be aggregated in the following projective distance estimation.

Wenbin Zou, Shishun Tian, Canqun Xiang and Xia Li are with Shenzhen Key Laboratory of Advanced Machine Learning and Applications, Guangdong Key Laboratory of Intelligent Information Processing, Institute of Artificial Intelligence and Advanced Communication, College of Electronics and Information Engineering, Shenzhen University, Shenzhen, 518060, China. (wzou@szu.edu.cn; stian@szu.edu.cn; xiangcanqun2018@email.szu.edu.cn; lixia@szu.edu.cn)

Di Wu is with Visual Computing Group, PingAn Insurance. (steven-wudi@gmail.com)

Lu Zhang is with National Institute of Applied Sciences of Rennes (INSA de Rennes) and Institut d’Électronique et des Technologies du numÉrique (IETR), Rennes, France. (lu.ge@insa-rennes.fr)

Fig. 2: The core component of the proposed network: following the detection head in 2D space, the 6D-VNet regresses the 6DoF rotation and translation vector simultaneously in 3D space.

To accommodate the requirement for accurate translation estimation in autonomous driving, we propose a framework, termed 6D-VNet, aiming at regressing the vehicle's rotation and translation simultaneously *cf.* Fig. 2. 6D-VNet streamlines the vehicle's 6DoF via the intermediate outputs from the Region Proposal Network (RPN) [10]. The detection part of the network is the canonical 2D object detection network (Mask R-CNN). The 6DoF estimation part of the network takes the intermediate output from the detection head. The challenging aspect of learning 6DoF vehicle pose is to design a loss function which is able to learn both rotation and translation. The model learns a complementary representation when supervised by both translation and orientation signals. Moreover, traffic participants exert mutual influence among their neighbours. Therefore, we introduce a weighted non-local block, which is a modified version of [11], to capture the collective information between traffic participants with interpretable self-attention map.

Specifically, the network is trained end-to-end by joint losses designed with solid geometric ground. Experimental results show that the proposed method outperforms the state-of-the-art two-staged systems. We list our contributions as follows:

- To our best knowledge, this is the first work which successfully regresses the rotation and translation simultaneously for deep learning-based object 6DoF pose estimation. And we show the effectiveness of the translation head inclusion into end-to-end training scheme (Sec. III-A).
- With a grounding in geometry, we investigate several joint losses that function synergistically (Sec. III-B).
- We capture the densely spatial dependencies by introducing a weighted non-local operation with interpretable self-attention map (Sec. III-C).

A Preliminary version of this manuscript was published previously [12]. Since then, we have conducted experiments on the Pascal3D+ dataset and extensive ablation studies.

II. RELATED WORK

Monocular-based 3D object detection were helped by the early work on face detection to popularise bounding box object detection. Later, PASCAL VOC [13], MS-COCO [14], [15] datasets pushed the detection towards a more diverse, challenging task. A tracking and estimation integrated model is proposed in [16] to determine spatial configuration of body parts in each frame. 3D convolutional neural network is presented in [17] for real-time 3D hand pose estimation, and more relevant works are given in [18] and [19], [20]. 3D head pose estimation [21], [22] is also a challenge. 3D head pose and facial actions in monocular video sequences that can be provided by low quality cameras are initialized and tracked by a 3-D pose estimator and 2-D face detector [23]. Recognizing the visual focus of attention of meeting participants based on their head pose is addressed in [24] using a Gaussian mixture model. Studies of human pose can be referred from [25]. KITTI dataset [26], [27] propelled the research for traffic participants under autonomous driving scenario. However, the 3D object detection task in the KITTI dataset primarily focuses on using point clouds data from Velodyne laser scanner, which is an expensive apparatus. In addition, the KITTI 3D object detection task only has half degree of freedom for rotation overlooking vehicle's heading direction.

Camera pose estimation is the problem of determining the position and orientation of a calibrated camera, which is to infer where you are and is key to the applications of mobile robotics, navigation and augmented reality, where localization is crucial for performance. [28] propose a robust and real-time monocular 6DoF relocalization framework, called PoseNet. It trains a convolutional neural network to regress the 6DoF camera pose from a single RGB image in an end-to-end manner with no requirement of additional engineering or graph optimisation. It is robust to difficult lighting, motion blur and unknown camera intrinsics, where point based SIFT registration fails. However, it was trained using a naive loss function, with hyper-parameters which require expensive tuning. To address this issue, a more fundamental theoretical treatment is given in [29] by exploring a number of loss functions based on geometry and scene re-projection error. [30] proposes a unified framework to tackle self-localisation and camera pose estimation simultaneously. Instead of using a single RGB image, it integrates the signals from multiple sensors to achieve high efficiency and robustness. The problem of camera pose estimation is egocentric, in other words, a single vector of 6 dimensions will suffice to relocalise the camera pose.

6DoF object detection [7] is also essential for robotic manipulation [31] and augmented reality applications [32]. The BOP benchmark [33] consists of eight datasets in a unified format that cover different practical scenarios and shows that the methods based on point-pair features currently outperform the methods based on template matching, learning-based and 3D local features. Encouraging results have been shown in recent researches [7], [8], [9] which either RGB or RGB-D images are used to detect 3D model instances and estimate their 6DoF poses. Especially, [7] proposes a

two-staged 6DoF object detection pipeline: firstly, a Single Shot Multibox Detector (SSD) [34] is applied to provide object bounding boxes and identifiers. Then, an Augmented Autoencoder (AAE) is applied to estimate the object rotation using a Domain Randomisation [35] Strategy. However, the aforementioned methods focuses on the industry-relevant objects with neither significant texture nor discriminative colour or reflectance properties. Moreover, objects of interests are lying on a uniform ground plane (*e.g.*, in T-LESS dataset [36], the range of object distance is from 650mm to 940mm). Hence, the tolerance of rotation needs to stay low, whereas translation can be relaxed.

III. MODEL

6D-VNet is conceptually intuitive and structurally hereditary: Faster R-CNN has two outputs for each candidate object, a class label and a bounding-box offset, Mask R-CNN adds a third branch that outputs the object mask. Likewise, 6D-VNet streamlines the objects 6DoF prediction via the intermediate outputs from the Region Proposal Network (RPN). However, in order not to break already learnt functionality of the pre-trained network, careful design choices must be customised for the end-to-end training. Next, we present the overall architecture in Sec. III-A. Particularly, we introduce the end-to-end training paradigm with translation estimation integration which greatly outperforms other two staged frameworks in terms of translation estimation accuracy. The design choices for the joint losses are presented in Sec. III-B. Lastly we show the spatial relationship between traffic participants can be incorporated via a modified weighted non-local block in Sec. III-C.

A. Network Architecture

6D-VNet is built upon the canonical object detection network as shown in Fig. 3. The system is a two-staged network which is trained end-to-end to estimate the 6DoF pose information for object of interest. The first stage of the network is a typical 2D object detection network (Mask R-CNN). The second stage of the network is the customised heads to estimate the object 6DoF pose information.

The 6DoF pose estimation branch is the main novelty of the model and is split into two parts: the first part only takes RoIAlign [37] from each candidate box if the candidate is of the vehicle class and performs sub-class categorisation and rotation estimation. Since in-plane rotation is unique for a given vehicle class, all vehicles share similar rotational features for the same yaw, pitch, and roll angles. Therefore, the fixed-size visual cue from RoIAlign layer is sufficient for estimating the candidate sub-category and rotation.

The second part takes both RoIAlign feature and bounding box information (in world unit as described in Sec. III-B) into consideration via a concatenation operation to estimate the 3 dimensional translational vector. To our knowledge, this novel formulation is the first of its kind to regress the translational vector directly. The joint feature combination scheme implicitly encodes the object class and rotation information via the concatenation operation ($\oplus$ in Fig. 3). The

translation regression head functions in synergy when it is combined with the joint loss from sub-category classification and quaternion regression. We show in the experiment that our novel formulation for translation regression produces much more accurate position estimation comparing to the methods that treat the translation estimation as a post-processing step. This accurate estimation of translational vector is particularly crucial for the applications where the distance of the objects are of primary importance (*e.g.*, in the autonomous driving scenario).

B. Joint Losses

We minimise the following loss $\mathcal{L}$ to train our network in an end-to-end fashion: $\mathcal{L} = \mathcal{L}_{det} + \mathcal{L}_{inst}$, where $\mathcal{L}_{det}$ denotes the multi-task loss as in a canonical detection network: $\mathcal{L}_{det} = \mathcal{L}_{cls} + \mathcal{L}_{box} + \mathcal{L}_{mask}$. The classification loss $\mathcal{L}_{cls}$, 2D bounding box loss $\mathcal{L}_{box}$ and 2D mask loss $\mathcal{L}_{mask}$ are identical as those defined in [37]. In order to accelerate the network training and keep the functionality of the pre-trained multi-task module (*e.g.*, mask head for instance segmentation), we can freeze these heads and their corresponding child nodes, *i.e.*, the convolutional backbone and set the $\mathcal{L}_{det}$ to zero during back propagation phase. $\mathcal{L}_{inst}$ denotes the individual instance loss for 6DoF estimation with sub-class categorisation. Specifically, it is defined as a triple-loss: $\mathcal{L}_{inst} = \lambda_{sub_cls} \mathcal{L}_{sub_cls} + \lambda_{rot} \mathcal{L}_{rot} + \lambda_{trans} \mathcal{L}_{trans}$, where λ_{sub_cls} , λ_{rot} , λ_{trans} are hyper-parameters used to balance their corresponding loss. Next we explain the design choices for the above triple losses.

Sub-category classification loss $\mathcal{L}_{sub_cls}$. Sub-category denotes the finer class of the vehicle corpus: *e.g.*, *Audi-A6*, *BMW-530*, *Benze-ML500*, *etc.* In order to balance the rare cases for infrequently appearing cars in the training images, weighted cross entropy is used for sub-category classification loss.

Rotation loss $\mathcal{L}_{rot}$. There are generally three representations for providing orientation information: Euler angles, $SO(3)$ rotation matrices and Quaternions. Euler angles are easily understandable and interpretable parametrisation of 3D rotation. However, there are two issues when directly regressing the Euler angles: (1) non-injectivity: the same angle could be represented by multiple values due to the wrapping around 2π radians, which make the regression a non uni-modal task; (2) Gimbal lock: possible loss of one degree of freedom does not make Euler angles invalid but makes them unsuited for practical applications. Given 3D models of the objects, one way to work around the problem is to rotate each view at fixed intervals to cover the whole $SO(3)$ and then find the nearest neighbour [7] or closest viewpoint [9], which treat the rotation estimation problem as a classification problem. But this requires a complete CAD model of the object and a discretisation step of orientation angles. To estimate rotation matrix directly, [8] propose a LieNet to regress a Lie algebra based on rotation representation. However, a 3×3 orthogonal matrix is over-parameterised and enforcing the orthogonality is non-trivial.

Quaternions are favourable due to the universality mapping from 4 dimensional values to legitimate rotations. This is

Fig. 3: System pipeline: 6D-VNet takes a monocular image as input and performs the vehicles' 6DoF estimation. The grey box represents a canonical instance segmentation network and the dark blue branch is for estimating object 6DoF pose and its sub-category.

a simpler process than the orthonormalisation of rotation matrices. Quaternions are continuous and smooth, lying in a unit manifold, which can be easily enforced through back-propagation. Therefore, the rotation head in our network focuses on the regression of quaternion representation. However, the main problem with quaternions is that they are not injective: the quaternion $\mathbf{q}$ and $-\mathbf{q}$ represent the same rotation because two unique values (from each hemisphere) map to a single rotation. To address the issue, we constrain all quaternions to one hemisphere such that there is a unique value for each rotation as follows:

Enforcement of Quaternions to one hemisphere: Let $\mathbb{Q}$ represents the set of quaternions, in which each quaternion, $\mathbf{q} \in \mathbb{Q}$, is represented as $\mathbf{q} = a + bi + cj + dk$, and $a, b, c, d \in \mathbb{R}$. A quaternion can be considered as a four-dimensional vector. The symbols i , j , and k are used to denote three “imaginary” components of the quaternion. The following relationships are defined: $i^2 = j^2 = k^2 = ijk = -1$, from which it follows that $ij = k$, $jk = i$, and $ki = j$.

The quaternion $\mathbf{q}$ and $-\mathbf{q}$ represent the same rotation because a rotation of θ in the direction v is equivalent to a rotation of $2\pi - \theta$ in the direction $-v$. One way to force uniqueness of rotations is to require staying in the “upper half” of $\mathbb{S}^3$. For example, require that $a \geq 0$, as long as the boundary case of $a = 0$ is handled properly because of antipodal points at the equator of $\mathbb{S}^3$. If $a = 0$, then require that $b \geq 0$. However, if $a = b = 0$, then require that $c \geq 0$ because points such as $(0, 0, -1, 0)$ and $(0, 0, 1, 0)$ are the same rotation. Finally, if $a = b = c = 0$, then only $d = 1$ is allowed.

Hence, for the rotation head, given the ground truth unique quaternion $\mathbf{q}$ and the predicted $\hat{\mathbf{q}}$, the rotation loss is defined

as:

$$\mathcal{L}_{rot}(\mathbf{q}, \hat{\mathbf{q}}) = \|\mathbf{q} - \frac{\hat{\mathbf{q}}}{\|\hat{\mathbf{q}}\|}\|_{\gamma} \quad (1)$$

An important choice for regressing in Euclidean space is the regression norm $\|\cdot\|_{\gamma}$. Typically, deep learning models use $L_1 = \|\cdot\|_1$ or $L_2 = \|\cdot\|_2$. With the datasets used in this paper, we found the L_1 norm performs better: the error does not increase quadratically with magnitude nor over-attenuate large residuals.

Translation loss $\mathcal{L}_{trans}$. Regressing translation vector in world unit instead of pixel unit stabilises the loss. The transformation of the detected object takes 2D bounding box centre, height and width u_p, v_p, h_p, w_p in pixel space and then outputs their corresponding u_w, v_w, h_w, w_w in world unit as:

$$u_w = \frac{(u_p - c_x)z_s}{f_x}, v_w = \frac{(v_p - c_x)z_s}{f_y}, h_w = \frac{h_p}{f_x}, w_w = \frac{w_p}{f_y}$$

where the matrix $[f_x, 0, c_x; 0, f_y, c_y; 0, 0, 1]$ is the camera intrinsic calibration matrix.

Huber loss is adopted to describe the penalty in translation estimation: give ground truth 3 dimensional translation vector $\mathbf{t}$ and the prediction $\hat{\mathbf{t}}$, the translation loss is:

$$\mathcal{L}_{trans}(\mathbf{t}, \hat{\mathbf{t}}) = \begin{cases} \frac{1}{2}(\mathbf{t} - \hat{\mathbf{t}})^2 / \delta & \text{if } |\mathbf{t} - \hat{\mathbf{t}}| < \delta, \\ |\mathbf{t} - \hat{\mathbf{t}}| - \frac{1}{2}\delta & \text{otherwise.} \end{cases} \quad (2)$$

where the hyperparameter δ controls the boundary of outliers. If δ is set to 1, then it becomes the smooth-L1 loss used in [10]. In this paper, δ is set as 2.8 which is the cut off threshold for translational evaluation as described in Sec. IV-B.

C. Weighted Non-local neighbour embedding

In order to capture spatial dependencies among detected objects of interest, we introduce a non-local block with a

weighted operation. We reason that the dependencies among neighbouring objects will assist the network to regularise the 6DoF pose estimation collectively better than treating them individually. For example, neighbouring cars on the same lane will follow almost the same orientation and maintain certain distance. There are several advantages of using a weighted non-local operations comparing with other social embedding schemes (i) non-local operations capture long-range dependencies directly by computing interactions between any two positions, regardless of their positional distance; (ii) non-local operations maintain the variable input sizes and can be easily combined with other operations; (iii) our proposed weighted operation renders it possible to associate the output maps with self-attention mechanisms for better interpretability.

The non-local means (NL-means) is first introduced in [38], based on a non-local averaging of all pixels in the image. Later [11] introduced the *non-local* operations as an efficient and generic component for capturing long-range dependencies with deep neural networks. The non-local operations maintain the variable input sizes. Intuitively, a non-local operation computed the response at a position as a weighted sum of the features at all positions in the input feature maps. The generic non-local operation in deep neural network is defined as:

$$y_i = \frac{1}{\mathcal{C}(x)} \sum_{\forall j} f(x_i, x_j) g(x_j) \quad (3)$$

where i is the index of an output position (here in space) whose response is to be computed and j is the index that enumerates all possible positions. x is the input signal and y is the output signal of the same size as x . A pairwise function f computes a scalar (representing relationship such as affinity) between i and j . The unary function g computes a representation of the input signal at the position j . The response is normalised by a factor $\mathcal{C}(x)$. The non-local models are not sensitive to the design choices of f and g . For simplicity and fast computation, we consider g in the form of a linear embedding: $g(x_j) = W_g x_j$, where W_g is a weight matrix to be learnt. The pairwise function f is implemented in the form of *embedding Gaussian* as: $f(x_i, x_j) = e^{\theta(x_i)^T \phi(x_j)}$, where $\theta(x_i) = W_\theta x_i$ and $\phi(x_j) = W_\phi x_j$ are two embeddings. And $\mathcal{C}(x) = \sum_{\forall j} f(x_i, x_j)$. The recently proposed self-attention module [39] is a special case of non-local operations in the embedding Gaussian version: when, for a given i , $\frac{1}{\mathcal{C}(x)} \sum_{\forall j} f(x_i, x_j)$ becomes *softmax* computation along the dimension j . So we have $y = \text{softmax}(x^T W_\theta^T W_\phi x) g(x)$.

However, we found out that when the input dimension d (for a feature map of $H \times W \times C$ where C is the channel number, $d = H \times W$) gets large, the dot products grow large in magnitude, pushing the softmax function into regions where it has extreme gradients. As a result, y will have extreme value as well.¹ To counteract the effect, we propose to use a *weighted non-local operation* for calculating the self-attention map $\mathcal{A}$

as:

$$\mathcal{A} = \text{softmax}\left(\frac{x^T W_\theta^T W_\phi x}{\sqrt{d}}\right) \quad (4)$$

So that $y = \mathcal{A} \cdot g(x)$. The weighted non-local operation scales the dot-product attention variance to unit 1, which consequently does not push the softmax operation to extreme, saturated values. Intuitively, the weighted operation has the same form of expression. However, the suitable temperature in softmax is finicky to determine. Alternatively, we scale the dot-product input variance to unit 1. Consequently, the output map after softmax operation will provide a justifiable interpretation in the form of a self-attention formulation.

IV. EXPERIMENTS

We benchmark our method on the challenging PAS-CAL3D+ [40] dataset and perform comprehensive studies on the challenging Apolloscape dataset [2]. For both Pascal3D+ and Apolloscape dataset, we focus our experiments on the “Car” category in urban scene.

Implementation Details: ResNet-101 is adopted as the convolutional body with Feature Pyramid Networks (FPN) as the detection backbone. The instance segmentation head is pre-trained using the Apolloscape scene dataset² with “car, motorcycle, bicycle, pedestrian, truck, bus, and tricycle” as 8 instance-level annotations.

The hyperparameters λ_{sub_cls} , λ_{rot} , λ_{trans} in Eqn. III-B are set to 1.0, 1.0, 0.1 to scale the loss accordingly. To decrease the translational outlier penalty and stabilise the network training, the hyperparameter δ in Eqn. III-B is set to 2.8 metres as the loose end of the translational metric in Sec. IV-B. The base learning rate starts from 0.01 with warm start-up scheme and the models are trained up to 5×10^4 iterations with learning rate divided by 10 at 1.5×10^4 th and 3×10^4 th iterations. We use a weight decay of 0.0001 and a momentum of 0.9. The RoIAlign takes a feature map of 7×7 from each RoI. The weighted non-local block is plugged into the last layer (5th) of the convolutional body with a receptive field of 32×32 . During training, the images are resized with the largest side randomly in [2000, 2300]. Due to memory limitation, batch size is set as one image per single GPU. The incorporation of non-local block will increase the memory requirement, hence the training images are resized with the largest side randomly in [1500, 2000] when non-local block is plugged in. Top 1000 regions are chosen as per FPN level with 100 batch size per image. During testing phase, 0.1 is chosen as the detection threshold in the Faster R-CNN head with multi-scale augmentation.

A. Analysis on Pascal3D+ dataset

We first evaluate our method on the task of joint detection and viewpoint estimation. The object orientation is represented in terms of viewpoint: azimuth, elevation and tilt angles. The results are reported using Average Viewpoint Precision (AVP) under different quantisation of angles, as proposed by [40]. The results are shown in Tab. I. Our proposed framework improves almost all previous methods apart from the recent

¹To illustrate why the dot products get large, assume that the components of $\mathbf{x}$ are random variables with mean 0 and variance 1. Then its self dot product, $\mathbf{x}^T \cdot \mathbf{x} = \sum_i^d x_i^2$, has mean 0 and variance d (d is the input dimension).

²<http://apolloscape.auto/scene.html>

Method	Car			
	AVP_4	AVP_8	AVP_{16}	AVP_{24}
Pepik et al. [41]	36.9	36.6	29.6	24.6
Viewpoints & KeyPoints [42]	55.2	51.5	42.8	40.0
RenderForCNN [43]	41.8	36.6	29.7	25.5
Poirson et al. [44]	51.4	45.2	35.4	35.7
Massa et al. [45]	58.3	55.7	46.3	44.2
Xiang et al. [46]	48.7	37.2	31.4	24.6
3D-RCNN [47]	71.8	65.5	55.6	52.1
6D-VNet	72.2	63.0	54.8	48.7

TABLE I: Joint detection and viewpoint evaluation on Pascal3D+ dataset for “Car” category.

Method	$Acc_{\frac{\pi}{6}} \uparrow$	$MedErr \downarrow$
Viewpoints & KeyPoints(TNET) [42]	0.89	8.8°
Viewpoints & KeyPoints [42]	0.90	9.1°
RenderForCNN [43]	0.88	6.0°
Deep3DBox [48]	0.90	5.8°
3D-RCNN(VGG16) [47]	0.94	3.4°
3D-RCNN(ResNet50) [47]	0.96	3.0°
6D-VNet	0.95	3.3°

TABLE II: Evaluation of viewpoint estimation with ground-truth detections on Pascal3D+ for “Car” category.

3D-RCNN [47] where a differentiable Render-and-Compare loss is introduced. This loss is complementary to our framework and we believe the inclusion of the loss that allows 3D shape and pose to be learned with 2D supervision can further improve the accuracy of the pose estimation.

We follow [47] [46] to evaluate viewpoint on ground-truth boxes in Tab. II which provides an upper-bound of viewpoint accuracy independent of the object detector adopted. The viewpoint estimation error is measured as geodesic distance over the rotation group $SO(3)$. The $Acc_{\pi/6}$ that estimates accuracy at $\frac{\pi}{6}$ and the median angular error $MedErr$ which was used in [46]. Our method performs on par with the state-of-the-art on both the $\frac{\pi}{6}$ and median angular error metrics.

B. Analysis on Apolloscape 3D Car Instance dataset

The Apolloscape 3D Car Instance challenge contains a diverse set of stereo video sequences recorded in the street scenes from different cities. There are 3941/208/1041 high quality annotated images in the training/validation/test set.³ The monocular RGB images are of pixel size 2710×3384 . It is worth noticing the high resolution of the images: the total number of pixels of a single image is 100 times than those of other canonical image datasets (e.g., MS-COCO, Mapillary Vistas, ImageNet). The camera intrinsic parameters are provided in the form of camera focal lengths (f_x, f_y) and optical centres expressed in pixels coordinates (c_x, c_y). Car models are provided in the form of triangle meshes. The mesh models have around 4000 vertices and 5000 triangle faces. One example mesh model is shown as in Fig. 2. There are total 79 car models (Fig. 4) in three categories (sedan1, sedan2, SUV) with only 34 car models appearing in the training set. In addition, ignored marks are provided as unlabelled regions and

³After manual examination, we have deleted visually distinguishable wrongly labelled images, leaving us with 3888/206 images for training/validation.

we only use the ignored masks to filter out detected regions during test.

Fig. 4: 79 car model meshes for the Apolloscape 3D Car Instance dataset.

	AP	AP_{50}	AP_{75}	AP_S	AP_M	AP_L	AP_{XL}
Single Scale	0.57	0.87	0.62	0.34	0.50	0.65	0.73
Multiple Scale	0.59	0.89	0.64	0.34	0.51	0.68	0.84

TABLE III: 2D bounding box detection mAP from the Faster R-CNN head. It is the upperbound for 3D detection result. Subscripts of AP represent squared object one side size in pixels as: S:28-56, M:56-112, L: 112-512, XL:512+.

Evaluation Metrics: The evaluation metrics follow similar instance mean AP as the MS-COCO [14]. However, due to 3D nature, the 3D car instance evaluation has its own idiosyncrasies: instead of using 2D mask IoU to judge a true positive, the 3D metric used in this dataset contains the perspective of shape (s), 3D translation (t) and 3D rotation (r). The shape similarity score is provided by an $N_{car} * N_{car}$ matrix where N_{car} denotes the number of car models. For 3D translation and 3D rotation, the Euclidean distance and arccos distance are used for measuring the position and orientation difference respectively.

Specifically, given an estimated 3D car model in an image $C_i = \{s_i, t_i, r_i\}$ and ground truth model $C_i^* = \{s_i^*, t_i^*, r_i^*\}$, the evaluation for these three estimates are as follows: for 3D shape, reprojection similarity is considered by putting the model at a fix location and rendering 10 views (v) by rotating the object. Mean IoU is computed between the two poses (P) rendered from each view. Formally, the metric is defined as: $c_{shape} = \frac{1}{|V|} \sum_{v \in V} IoU(P(s_i), P(s_i^*))_v$, where V is a set

of camera views. For 3D translation and rotation, the evaluation metric follows that of the canonical self-localisation: $c_{trans} = \|t_i - t_i^*\|^2$ and $c_{rot} = \arccos(|q(r_i) \cdot q(r_i^*)|)$. Then, a set of 10 thresholds from loose criterion to strict criterion ($c_0, c_1, \dots, c_9$) is defined as:

$$\begin{aligned} shapeThrs &= [.5 : .05 : .95] \\ rotThrs &= [50 : 5 : 5] \\ transThrs &= [2.8 : .3 : 0.1] \end{aligned}$$

where the most loose metric c_0 : 0.5, 50, 2.8 means shape similarity > 0.5 , rotation distance $< 50^\circ$ and translation distance < 2.8 metres, and stricter metrics can be interpreted correspondingly: all three criterion must be satisfied simultaneously so as to be counted as a true positive.

It is worth noting the strict translation distance threshold of 2.8 metres: it requires that the detected vehicle's distance from the camera centre needs to be correctly estimated within a 2.8 metres threshold even if the vehicle is hundreds metres away from the camera, otherwise the detection will be counted as a false positive. The precise translational estimation requirement is the major factor for the network to produce incorrect false positive, which is a challenging task from a human perspective as well.

1) *Main Results & Ablation Studies*: We present the results of the proposed 6D-VNet which reaches the 1st place in Apolloscape challenge 3D Car Instance task. Furthermore, we perform comprehensive evaluations to analyse the “bells and whistles” in 6D-VNet, which further improve state-of-the-art *cf.* Tab. VII. Our ablation studies gradually incorporated all components and are detailed as follows.

Effect of End-to-End Training. We first provide the 2D bounding box mAP from Faster R-CNN head to serve as the upper bound by 2D object detection as in Tab. III. We can find that small objects are more challenging to detect. Small object also indicates the object longitudinal axis distance is typically far away from the camera. The accurate estimation of large translational distance value is thus more challenging.

In Tab. IV we show that the translational head is crucial for improving the mAP in our end-to-end training scheme. The projective distance estimation was the defacto approach in previous state-of-the-art methods [7], [9], [8] as a second stage to measure the translational distance as described as follows:

Fig. 5: Projective Distance Estimation.

Projective Distance Estimation: Fig. 5 illustrates the projective distance estimation via geometric method which were adopted in previous state-of-the-art methods [7], [9], [8]. For each object we precomputed the 2D bounding box and centroid.

To this end, the object is rendered at a canonical centroid distance z_r (z_r should be set larger than the object length in longitudinal axis so that the entire object can be projected onto the image plane). Subsequently, the object distance z_s can be inferred from the projective ratio according to $z_s = \frac{l_r}{l_s} z_r$, where l_r denotes diagonal length of the precomputed bounding box and l_s denotes the diagonal length of the predicted bounding box on the image plane. Given its depth component z_s , the complete translational vector can be recovered geometrically as:

$$x_s = \frac{(u - c_x)z_s}{f_x}, y_s = \frac{(v - c_y)z_s}{f_y}$$

where $[u, v]$ is the bounding box centre, and the matrix $[f_x, 0, c_x; 0, f_y, c_y; 0, 0, 1]$ is the camera intrinsic calibration matrix. The formulation assumes that: (i) the object centre in 3D will be projected to the object bounding box in the 2D image; (ii) the predicted object class and rotation vector is correctly estimated.

The depth component estimation is treated as an entirely independent process given the rotation pose estimation. Hence, the geometric post-processing method achieves around only 3.8% mAP due to the crude translation estimation. We show that using the same bounding box information (in world unit) to train a translation regression head improves the mAP to 8.8%.

Fig. 6: Losses w/o & w/ weighted non-local block (NL).

Effect of Joint Losses. We then investigate the synergistic training of using visual information to regress translational vector as the concatenation operation $\oplus$ in Fig. 3. The w/o $\oplus$ row in Tab. IV represents the translation head using only normalised world unit bounding box information, the $\oplus$ row represents the translation head that combines the intermediate visual information from RoIAlign with the bounding box information. The overall mAP is further improved by 2% using the RoIAlign intermediate feature. It is worth noting the synergy of the joint losses reflecting through the *intermediate training value* columns in Tab. IV: the shape similarity, rotation and translational score improve by 0.01, 1.4° and 2.3 metres respectively. By connecting the translation head with intermediate RoIAlign branch, the translational losses are jointly back propagated with the losses from the sub-categorisation and rotation. The improvement of translational estimation synergistically boost the accuracy for shape and rotation estimation. It shows that the network is able to learn the implicit information that is shared amongst the object class, rotation and translation.

Effect of Weighted Non-local Block. The weighed non-local block is inserted into the last layer of the ResNet to

Method	Intermediate training value			val mAP	test mAP
	shape sim	rot dist	trans dist		
projective distance	0.84	13.5°	-	0.038	0.0371
w/o $\oplus$	0.86	11.9°	4.6	0.088	0.0882
$\oplus$	0.87	10.5°	2.3	0.108	0.1049
$\oplus$ (fine-tune)	0.90	8.8°	1.3	0.128	0.1223

TABLE IV: Effect of end-to-end training and joint losses. The projective distance is the most commonly adopted approach in previous state-of-the-art methods [7], [9], [8] as a second stage to measure the translational distance. w/o $\oplus$ denotes the network without the concatenation operation in Fig. 3. $\oplus$ represents the branch trained with joint losses. $\oplus$ (fine-tune) denotes, when training the network, the learnt parameters in convolutional body and Faster R-CNN head are unfreezed.

Method	Inference Time (second)				val mAP	test mAP
	Det Head	Triple Head	Misc	Total		
w/o Weighted Non-local (Single Scale)	0.966	1.322	0.002	2.40	0.128	0.1223
w/o Weighted Non-local (Multi-Scale)	5.535	1.381	0.004	6.86	0.143	0.1412
Weighted Non-local (Multi-Scale)	5.810	1.473	0.003	7.29	0.146	0.1435

TABLE V: Effect of Weighted Non-local Block and runtime analysis.

Fig. 7: Examples of the behaviours of a weighted non-local block. These two examples are from held-out validation images. On the left images, the starting point (pink square) represents one x_i . Since we insert the weighted non-local block in res5 layer, one pink square denotes a receptive field of pixel size 32×32 . We position the starting points on one of the vehicles. The end points of arrows represent x_j . The 10 highest weighted arrows for each x_i are visualised. The arrows clearly indicate that the weighted non-local block attends to the neighbouring traffic participants. Note that in the second illustration, end points are able to locate neighbouring vehicles' wheels and rear-view mirrors which are critical clues to position and orient the vehicle. The right images are the visualisations of self-attention map $\mathcal{A}$ of x_i from Eqn. 4. These visualisations show how the weighted non-local block finds interpretable, relevant indications of neighbouring vehicles to adjust the pose estimation.

Fig. 8: 3D renderings of predicted vehicles with camera coordinate axis. The bottom renderings framed by silver plate represent the results from the model with weighted non-local block. The vehicles in colour silver, green, red represent ground truth, true positive and false positive respectively. Due to the strictly fixed 2.8-metre translation criterion, vehicles that are farther way from the cameras are more difficult to measure. The pink arrows highlight the predictions that have been adjusted according to their neighbouring vehicles when weighted non-local block is plugged in.

encode dense spatial dependencies. Fig. 6 shows that by plugging in the weighted non-local block, both the rotational and translational training losses further decrease comparing with the previously converged model, which greatly boost up

Module	number of parameters
mask_rcnn_r101	242MB
triplet head	54MB
non_local block	518kb

TABLE VI: Number of parameters of each component of the proposed method.

the training procedure. Tab. V shows that the incorporation of weighted non-local block consistently improves mAP on both validation and test set with marginally increased inference runtime. According to Tab. VI, the backbone Mask RCNN and the triplet heads contain most parameters (242Mb+54Mb), the additional parameters are caused by the proposed weighted non-local block only achieve 518kb which is extremely low compared to the backbone. Thus, similar conclusion can be drawn that the proposed method can improve the performance with marginally increased number of parameters.

We visualize the self-attention map $\mathcal{A}$ of Eqn. 4 in the weighted non-local block in Fig. 7. The heat maps on the right exhibit meaningful attention positions (without the weighted operation, the heat map will exhibit extreme high temperature with hard to interpret positions on the attention map), which demonstrates that the weighted non-local block can learn to find meaningful relational clues regardless of the distance in space. In Fig. 8, we visualise the predicted vehicles in 3D space. When incorporated with weighted non-local block, the model is able to capture the spatial dependencies so that it adjusts the predictions according to the distance and orientation of neighbouring vehicles.

V. CONCLUSIONS

We proposed an end-to-end network 6D-VNet for 6DoF pose estimation of vehicles from monocular RGB images. It can not only detect the traffic participants, but also generate their translations and rotations. To the best of our knowledge, the incorporation of translation regression into the network is the first of its kind and greatly improve the pose estimation accuracy compared to those methods which treat the translation estimation as a post-processing step. In addition, we design the joint losses according to solid grounding in geometry, which are crucial to achieve accurate pose estimation. The experiments show that the proposed method reaches the first place in Apolloscape challenge 3D Car Instance task and achieves the performance of current state-of-the-art framework on the PASCAL3D+ dataset. Particularly, a large improvement for position estimation is observed when the translation head is trained from both visual clues and bounding box information. Furthermore, we demonstrate that the spatial dependencies among neighbouring vehicles can be incorporated via a weighted non-local block and an interpretable self-attention map. It can help regularise the 6DoF object pose estimation collectively which is better than treating them individually. In this paper, the 6DoF pose is directly generated from the network, post-refinement is not considered. In the future work we will try to further improve the 6DoF pose estimation by using post-processing techniques such as iterative closet point based algorithms or iterative refinement network.

REFERENCES

- [1] "Apolloscape website," http://apolloscape.auto/leader_board.html, 2018, [Online; accessed 1-February-2021].
- [2] X. Huang, X. Cheng, Q. Geng, B. Cao, D. Zhou, P. Wang, Y. Lin, and R. Yang, "The apolloscape dataset for autonomous driving," in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2018, pp. 1067–10676, doi:10.1109/CVPRW.2018.00141.
- [3] D. Bećirbašić, M. Molnar, . Lukač, and D. Samardžija, "Video-processing platform for semi-autonomous driving over 5g networks," in *2017 IEEE 7th International Conference on Consumer Electronics - Berlin (ICCE-Berlin)*, 2017, pp. 42–46, doi: 10.1109/ICCE-Berlin.2017.8210584.
- [4] Shubham, M. Reza, S. Choudhury, J. K. Dash, and D. S. Roy, "An ai-based real-time roadway-environment perception for autonomous driving," in *2020 IEEE International Conference on Consumer Electronics - Taiwan (ICCE-Taiwan)*, 2020, pp. 1–2, doi: 10.1109/ICCE-Taiwan49838.2020.9258145.
- [5] D. Vitas, M. Tomic, and M. Burul, "Traffic light detection in autonomous driving systems," *IEEE Consumer Electronics Magazine*, vol. 9, no. 4, pp. 90–96, 2020, doi: 10.1109/MCE.2020.2969156.
- [6] T. Kim and H. Shim, "Road semantic segmentation oriented dataset for autonomous driving," in *2020 IEEE International Conference on Consumer Electronics - Asia (ICCE-Asia)*, 2020, pp. 1–3, doi: 10.1109/ICCE-Asia49877.2020.9277270.
- [7] M. Sundermeyer, Z. Marton, M. Durner, and R. Triebel, "Implicit 3d orientation learning for 6d object detection from rgb images," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, doi:10.1007/978-3-030-01231-1_43.
- [8] M.-T. Do, T. Pham, M. Cai, and I. Reid, "Real-time monocular object instance 6d pose estimation," in *British Machine Vision Conference (BMVC)*, 2018.
- [9] W. Kehl, F. Manhardt, F. Tombari, S. Ilic, and N. Navab, "Ssd-6d: Making rgb-based 3d detection and 6d pose estimation great again," in *Proceedings of the International Conference on Computer Vision (ICCV)*, 2017, doi: 10.1109/ICCV.2017.169.
- [10] R. Girshick, "Fast r-cnn," in *2015 IEEE International Conference on Computer Vision (ICCV)*, 2015, pp. 1440–1448, doi: 10.1109/ICCV.2015.169.
- [11] X. Wang, R. Girshick, A. Gupta, and K. He, "Non-local neural networks," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, doi: 10.1109/CVPR.2018.00813.
- [12] D. Wu, Z. Zhuang, C. Xiang, W. Zou, and X. Li, "6d-vnet: End-to-end 6-dof vehicle pose estimation from monocular rgb images," *IEEE conference on computer vision and pattern recognition, workshop of autonomous driving (WAD)*, 2019, doi: 10.1109/CVPRW.2019.00163.
- [13] M. Everingham, S. A. Eslami, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman, "The pascal visual object classes challenge: A retrospective," *International journal of computer vision*, 2015, doi: 10.1007/s11263-014-0733-5.
- [14] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft coco: Common objects in context," in *European conference on computer vision (ECCV)*, 2014, doi: 10.1007/978-3-319-10602-1_48.
- [15] S. Aslan, G. Ciocca, and R. Schettini, "Semantic food segmentation for automatic dietary monitoring," in *2018 IEEE 8th International Conference on Consumer Electronics - Berlin (ICCE-Berlin)*, 2018, pp. 1–6, doi: 10.1109/ICCE-Berlin.2018.8576231.
- [16] L. Zhao, X. Gao, D. Tao, and X. Li, "Learning a tracking and estimation integrated graphical model for human pose tracking," *IEEE Transactions on Neural Networks and Learning Systems*, 2015, doi: 10.1109/TNNLS.2015.2411287.
- [17] L. Ge, H. Liang, J. Yuan, and D. Thalmann, "Real-time 3d hand pose estimation with 3d convolutional neural networks," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019, doi: 10.1109/TPAMI.2018.2827052.
- [18] M. Patacchiola and A. Cangelosi, "Head pose estimation in the wild using convolutional neural networks and adaptive gradient methods," *Pattern Recognition*, vol. 71, pp. 132–143, 2017, doi: 10.1016/j.patcog.2017.06.009.
- [19] P. Bao, A. I. Maqueda, C. R. del-Blanco, and N. García, "Tiny hand gesture recognition without localization via a deep convolutional network," *IEEE Transactions on Consumer Electronics*, vol. 63, no. 3, pp. 251–257, 2017, doi: 10.1109/TCE.2017.014971.

PD	TH	TV	FT	MS	NL	IM	QS	mAP	c_0	c_5	AP^S	AP^M	AP^L	AR^1	AR^{10}	AR^{100}	AR^S	AR^M	AR^L
✓								0.037	0.125	0.042	0.072	0.042	0.076	0.016	0.091	0.128	0.072	0.133	0.245
	✓							0.067	0.190	0.089	0.069	0.057	0.194	0.026	0.116	0.139	0.069	0.138	0.305
	✓	✓						0.088	0.237	0.122	0.100	0.077	0.246	0.030	0.141	0.174	0.100	0.169	0.360
	✓	✓	✓					0.121	0.331	0.164	0.126	0.115	0.297	0.035	0.162	0.231	0.126	0.246	0.424
	✓	✓	✓	✓				0.141	0.351	0.198	0.128	0.131	0.357	0.040	0.179	0.243	0.128	0.253	0.477
	✓	✓	✓	✓	✓			0.144	0.352	0.199	0.131	0.133	0.360	0.040	0.187	0.249	0.131	0.261	0.481
	✓	✓	✓	✓	✓	✓		0.147	0.357	0.206	0.117	0.137	0.366	0.041	0.190	0.246	0.117	0.262	0.489
	✓	✓	✓	✓	✓	✓	✓	0.148	0.353	0.209	0.115	0.138	0.371	0.042	0.191	0.244	0.115	0.259	0.490

TABLE VII: Performance on test set in terms of mAP. c_0 is the most loose criterion for evaluating AP and c_5 is in the middle of criterion. Superscript S, M, L of average precision (APs) and average recall (ARs) represent the object sizes. Superscript number 1, 10, 100 represent the total number of detections for calculating recalls. Projective distance (PD) [8] estimation is adopted in the state-of-the-art methods [8], [9], [7]. Triple head (TH) is the baseline 6D-VNet. Translation head is then concatenated with visual branch (TV), represented by $\oplus$ operation in Fig. 3. Fine-tuning (FT) the convolutional body and detection head gives the task specific network a 3% boost in mAP. Multi-scale testing (MS) further increases the accuracy. The incorporation of weighted non-local block (NL) improves both precision and recall. Using ignore mask (IM) to filter 2D detection bounding box with 0.5 IoU as threshold improves the precision, however, slightly degrades the recall. Finally, enforcing the quaternions to one hemisphere (QS) achieves the current state-of-the-art.

- [20] M. Al-Hammadi, G. Muhammad, W. Abdul, M. Alsulaiman, and M. S. Hossain, "Hand gesture recognition using 3d-cnn model," *IEEE Consumer Electronics Magazine*, vol. 9, no. 1, pp. 95–101, 2020, doi: 10.1109/MCE.2019.2941464.
- [21] K. Chen, K. Jia, H. Huttunen, J. Matas, and J.-K. Kämäräinen, "Cumulative attribute space regression for head pose estimation and color constancy," *Pattern Recognition*, vol. 87, pp. 29–37, 2019, doi: 10.1016/j.patcog.2018.10.015.
- [22] J. T. Rui Li, Zhenyu Liu, "A survey on 3d hand pose estimation: Cameras, methods, and datasets," *Pattern Recognition*, 2019, doi: 10.1016/j.patcog.2019.04.026.
- [23] F. Dornaika and B. Raducanu, "Three-dimensional face pose detection and tracking using monocular videos: Tool and application," *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, vol. 39, no. 4, pp. 935–944, 2009, doi: 10.1109/TSMCB.2008.2009566.
- [24] S. O. Ba and J. Odobez, "Recognizing visual focus of attention from head pose in natural meetings," *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 2009, doi: 10.1109/TSMCB.2008.927274.
- [25] P. O. O. Duc Thanh Nguyen, Wanqing Li, "Human detection from images and videos: A survey," *Pattern Recognition*, 2019, doi: 10.1016/j.patcog.2015.08.027.
- [26] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? the kitti vision benchmark suite," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012, doi: 10.1109/CVPR.2012.6248074.
- [27] Y. Fan, B. Wu, C. Huang, and Y. Bai, "Environment detection of 3d lidar by using neural networks," in *2019 IEEE International Conference on Consumer Electronics (ICCE)*, 2019, pp. 1–2, doi: 10.1109/ICCE.2019.8662037.
- [28] A. Kendall, M. Grimes, and R. Cipolla, "Posenet: A convolutional network for real-time 6-dof camera relocalization," in *Proceedings of the IEEE international conference on computer vision (ICCV)*, 2015, doi: 10.1109/ICCV.2015.336.
- [29] A. Kendall, R. Cipolla *et al.*, "Geometric loss functions for camera pose regression with deep learning," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, doi: 10.1109/CVPR.2017.694.
- [30] P. Wang, R. Yang, B. Cao, W. Xu, and Y. Lin, "Dels-3d: Deep localization and segmentation with a 3d semantic map," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, doi: 10.1109/CVPR.2018.00614.
- [31] S. Lee, S. Jang, J. Kim, and B. Choi, "A hardware architecture of face detection for human-robot interaction and its implementation," in *2016 IEEE International Conference on Consumer Electronics-Asia (ICCE-Asia)*, 2016, pp. 1–2, doi: 10.1109/ICCE-Asia.2016.7804752.
- [32] C. Cuevas, D. Berjon, F. Moran, and N. Garcia, "Moving object detection for real-time augmented reality applications in a gpgpu," *IEEE Transactions on Consumer Electronics*, vol. 58, no. 1, pp. 117–125, 2012, doi: 10.1109/TCE.2012.6170063.
- [33] T. Hodan, F. Michel, E. Brachmann, W. Kehl, A. GlentBuch, D. Kraft, B. Drost, J. Vidal, S. Ihrke, X. Zabulis, C. Sahin, F. Manhardt, F. Tombari, T.-K. Kim, J. Matas, and C. Rother, "Bop: Benchmark for 6d object pose estimation," in *The European Conference on Computer Vision (ECCV)*, September 2018, doi: 10.1007/978-3-030-01249-6_2.
- [34] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg, "Ssd: Single shot multibox detector," in *European conference on computer vision (ECCV)*, 2016, doi: 10.1007/978-3-319-46448-0_2.
- [35] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel, "Domain randomization for transferring deep neural networks from simulation to the real world," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2017, pp. 23–30, doi: 10.1109/IROS.2017.8202133.
- [36] T. Hodan, P. Haluza, Š. Obdržálek, J. Matas, M. Lourakis, and X. Zabulis, "T-less: An rgb-d dataset for 6d pose estimation of textureless objects," in *Applications of Computer Vision (WACV)*, 2017, doi: 10.1109/WACV.2017.103.
- [37] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask r-cnn," in *Proceedings of the International Conference on Computer Vision (ICCV)*, 2017, doi: 10.1109/ICCV.2017.322.
- [38] A. Buades, B. Coll, and J.-M. Morel, "A non-local algorithm for image denoising," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2005, doi: 10.1109/CVPR.2005.38.
- [39] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems (NIPS)*, 2017.
- [40] Y. Xiang, R. Mottaghi, and S. Savarese, "Beyond pascal: A benchmark for 3d object detection in the wild," in *IEEE Winter Conference on Applications of Computer Vision*, 2014, doi: 10.1109/WACV.2014.6836101.
- [41] B. Pepik, M. Stark, P. Gehler, and B. Schiele, "Teaching 3d geometry to deformable part models," in *IEEE conference on computer vision and pattern recognition*, 2012, doi: 10.1109/CVPR.2012.6248075.
- [42] S. Tulsiani and J. Malik, "Viewpoints and keypoints," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015, doi: 10.1109/CVPR.2015.7298758.
- [43] H. Su, C. R. Qi, Y. Li, and L. J. Guibas, "Render for cnn: Viewpoint estimation in images using cnns trained with rendered 3d model views," in *Proceedings of the IEEE International Conference on Computer Vision*, 2015, doi: 10.1109/ICCV.2015.308.
- [44] P. Poirson, P. Ammirato, C.-Y. Fu, W. Liu, J. Kosecka, and A. C. Berg, "Fast single shot detection and pose estimation," in *2016 Fourth International Conference on 3D Vision (3DV)*. IEEE, doi: 10.1109/3DV.2016.78.
- [45] F. Massa, R. Marlet, and M. Aubry, "Crafting a multi-task cnn for viewpoint estimation," *British Machine Vision Conference*, 2016, doi: 10.5244/C.30.91.
- [46] Y. Xiang, W. Choi, Y. Lin, and S. Savarese, "Subcategory-aware convolutional neural networks for object proposals and detection," in *2017 IEEE winter conference on applications of computer vision (WACV)*, doi: 10.1109/WACV.2017.108.
- [47] A. Kundu, Y. Li, and J. M. Rehg, "3d-rnn: Instance-level 3d object reconstruction via render-and-compare," in *Proceedings of the IEEE*

- Conference on Computer Vision and Pattern Recognition*, 2018, doi: 10.1109/CVPR.2018.00375.
- [48] A. Mousavian, D. Anguelov, J. Flynn, and J. Kosecka, “3d bounding box estimation using deep learning and geometry,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2017, doi: 10.1109/CVPR.2017.597.