# JPerceiver: Joint Perception Network for Depth, Pose and Layout Estimation in Driving Scenes

Haimei Zhao<sup>1</sup>, Jing Zhang<sup>1</sup>, Sen Zhang<sup>1</sup>, and Dacheng Tao<sup>2,1</sup>

<sup>1</sup> The University of Sydney, 6 Cleveland St, Darlington, NSW 2008, Australia <sup>2</sup> JD Explore Academy, Beijing, China {hzha7798,szha2609}@uni.sydney.edu.au jing.zhang1@sydney.edu.au dacheng.tao@gmail.com

Abstract. Depth estimation, visual odometry (VO), and bird's-eyeview (BEV) scene layout estimation present three critical tasks for driving scene perception, which is fundamental for motion planning and navigation in autonomous driving. Though they are complementary to each other, prior works usually focus on each individual task and rarely deal with all three tasks together. A naive way is to accomplish them independently in a sequential or parallel manner, but there are three drawbacks, i.e., 1) the depth and VO results suffer from the inherent scale ambiguity issue; 2) the BEV layout is usually estimated separately for roads and vehicles, while the explicit overlay-underlay relations between them are ignored; and 3) the BEV layout is directly predicted from the front-view image without using any depth-related information, although the depth map contains useful geometry clues for inferring scene layouts. In this paper, we address these issues by proposing a novel joint perception framework named JPerceiver, which can simultaneously estimate scale-aware depth and VO as well as BEV layout from a monocular video sequence. It exploits the cross-view geometric transformation (CGT) to propagate the absolute scale from the road layout to depth and VO based on a carefully-designed scale loss. Meanwhile, a crossview and cross-modal transfer (CCT) module is devised to leverage the depth clues for reasoning road and vehicle layout through an attention mechanism. JPerceiver can be trained in an end-to-end multi-task learning way, where the CGT scale loss and CCT module promote inter-task knowledge transfer to benefit feature learning of each task. Experiments on Argoverse, Nuscenes and KITTI show the superiority of JPerceiver over existing methods on all the above three tasks in terms of accuracy, model size, and inference speed. The code and models are available at https://github.com/sunnyHelen/JPerceiver.

Keywords: Depth Estimation, Visual Odometry, Layout Estimation

# 1 Introduction

Autonomous driving has witnessed great progress in recent years, where deep learning is playing an increasing role in perception [37,27,61], planning [6,48],



Fig. 1: The proposed JPerceiver can predict BEV semantic layout (middle top, white for roads, blue for cars and cyan for the ego car), scale-aware depth map (middle bottom) and VO result (bottom right) simultaneously from a monocular video sequence (top left). The drivable area is visualized in green (bottom left) by projecting the BEV road layout onto the image plane. The ground truth BEV layout is shown in top right.

navigation [2,44], and decision making [18,21]. Among them, scene perception is the basis for other subsequent procedures in autonomous driving [56], which includes various sub-tasks for different perception purposes, e.g., depth estimation for 3D measurement [14,17], ego motion estimation for localization and visual odometry (VO) [49,25,62,58], as well as bird's-eye-view (BEV) or front-view (FV) layout estimation for detecting obstacles and drivable areas [31,36,54,10]. Although these tasks have underlying relations to each other intuitively, they are usually tackled separately in prior works [17,54]. The joint estimation for all these tasks has not drawn enough attention so far, and the benefits and challenges of doing so remain unclear, which is the focus of this paper.

Depth and VO estimation are two closely related computer vision tasks that have been studied for decades [45.33]. Recent self-supervised learning methods use the photometric consistency between consecutive frames to achieve the simultaneous estimation of scene depth and VO from monocular video sequences. where no ground truth depth labels are required [60,17,42]. On the other hand, BEV scene layout estimation refers to the task of estimating the semantic occupancy of roads and vehicles in the metric-scale BEV plane directly from FV images [31,36,54,20]. Though significant progress has been made in each individual task, they still suffer from some inherent problems: i.e., (1) the scale ambiguity in monocular depth and VO estimation since the photometric error between corresponding pixels is equivalent up to an arbitrary scaling factor w.r.t. depth and translation, and (2) the lack of geometry priors for predicting complex BEV layout. Consequently, monocular depth and VO predictions need to be rescaled with a scaling ratio derived from ground truth [17,4], which is not appealing in real-world applications. And previous BEV methods [31,36,54] usually predict the BEV layout of different semantic categories separately and ignore potentially useful geometry clues such as the depth order between cars.

In this paper, we propose to handle these three tasks simultaneously and provide complementary information for each other to address the aforementioned issues. We are inspired by the following two key observations. First, we note that the BEV road layout can provide absolute scale under the weak assumption that the road is flat, which allows us to exploit the cross-view geometric transformation and obtain a depth map with an absolute scale corresponding to the distance field that existed in the layout. As a result, the absolute scale can be introduced into our depth and VO predictions, resolving the scale ambiguity problem. Second, the learned depth predictions can provide useful priors about scene geometry (e.g., the relationship between near and far as well as overlay and underlay between objects and roads in the scene) to help solve the challenges (e.g., occlusions) in BEV layout estimation.

To this end, we propose a novel joint perception network named JPerceiver that can estimate scale-aware depth and VO as well as BEV layout of roads and vehicles simultaneously, as shown in Fig. 1. JPerceiver follows the multi-task learning framework, consisting of three networks for depth, pose and layout, respectively, which can be efficiently trained in an end-to-end manner. Specifically, we design a cross-view geometric transformation-based (CGT) scale loss to propagate the absolute scale from the road layout to depth and VO. Meanwhile, a cross-view and cross-modal transfer (CCT) module is devised to leverage the depth clues for inferring the road and vehicle layouts through an attention mechanism. Our proposed scale loss and CCT module not only promote inter-task knowledge transfer but also benefit the feature learning of each task via network forward computation and gradient back-propagation.

The contributions of this paper are summarized as follows: 1) we propose the first joint perception framework JPerceiver for depth, VO and BEV layout estimation simultaneously; 2) we design a CGT scale loss to leverage the absolute scale information from the BEV layout to achieve scare-aware depth and VO; 3) we devise a CCT module that leverages the depth clues to help reason the spatial relationships between roads and vehicles implicitly, and facilitates the feature learning for BEV layout estimation; and 4) we conduct extensive experiments on public benchmarks and show that JPerceiver outperforms the state-of-the-art methods on the above three tasks by a large margin.

# 2 Related Work

Self-supervised depth estimation and VO. SfMLearner [60] is one of the first works that propose to optimize depth and pose jointly in a self-supervised manner, utilizing the photometric consistency among continuous frames. Though this self-supervised learning scheme has drawn great attention from researchers and achieved promising results [47,29,17,42,17], current monocular unsupervised methods still suffer from the scale ambiguity problem. McCraith et at. [32] fit sample points to get the road plane estimation in the 3D world during test to obtain scale hint, but the hard formulation limits its general applicability. DNet [52] proposed to recover the scale by calculating the ratio of the estimated camera height and a given one, which requires a visible ground plane to be detected during inference. Wagstaff and Kelly [46] also use the camera height as the scale hint by training a plane segmentation network, which is the most similar work to ours. However, they use a three-stage training strategy to train networks separately which is much more complex than our end-to-end method.

BEV-based environment perception. Due to the limited field of view (FOV) of FV cameras, BEV representation is commonly used in environment perception and motion planning for autonomous driving [50,39]. Traditional methods usually predict depth and segmentation from front images, and then warp them to BEV through inverse perspective mapping (IPM) [30,43], which loses a large amount of information and cause distortions due to potential occlusions. Recently deep learning-based methods have been developed to estimate the road and vehicle layout in the orthographic BEV plane, taking the advantage of the strong hallucination ability of CNN [26,31,54,36,20,1]. The newly released self-driving datasets like Argoverse [7] and Nuscenes [5] provide a large number of BEV maps that contain annotations of drivable areas, which makes it possible to train models for BEV perception using real-world data. Compared with prior methods, we explore the incorporation of self-supervised depth learning explicitly, which provides an important perception output with useful geometric clues for BEV layout estimation. Dwivedi et al. [13] also conduct explicit depth estimation but just take it as an intermediate process to model 3D geometry rather than a joint learning perception task. Besides, prior works usually predict different semantic categories separately, while JPerceiver exploits the synergy of different semantics and predicts the layouts of all categories simultaneously.

**MTL-based environment perception.** Recently, some multi-task learning works [55,63,38,59,24,11,41,12] propose to combine related perception tasks with depth estimation and VO to exploit complementary information such as segmentation [38,24,41] and optical flow [55,63,12], which effectively boost the network performance. However none of them tackles the scale ambiguity problem of monocular depth and VO via multi-task learning, which is one of our key purposes in this paper.

# 3 Method

### 3.1 Overview of JPerciver

As shown in Fig. 2, JPerceiver consists of three networks for depth, pose and layout, respectively, which are all based on the encoder-decoder architecture. The depth network aims to predict the depth map  $D_t$  of the current frame  $I_t$ , where each depth value indicates the distance between a 3D point and the camera. And the goal of the pose network is to predict the pose transformations  $T_{t\to t+m}$  between the current frame and its adjacent frames  $I_{t+m}$ . The layout network targets to estimate the BEV layout  $L_t$  of the current frame, i.e. semantic occupancy of roads and vehicles in the top-view Cartesian plane. The three networks are jointly optimized during training. The overall objective function consists of the loss items of all the three tasks and can be formulated as:

$$\ell_{total} = \ell_{dp} + \ell_{layout},\tag{1}$$

where  $\ell_{dp}$  is the loss of depth and VO estimation in the self-supervised learning scheme, and  $\ell_{layout}$  is the loss of the layout estimation task. We explain the details of  $\ell_{dp}$  and  $\ell_{layout}$  in Sec. 3.2 and Sec. 3.3, respectively.

5



Fig. 2: JPerceiver consists of three networks for depth, pose and layout estimation, and is trained in the end-to-end manner.  $F_d$ ,  $F_f$ ,  $F_p$  represent the feature learned for three tasks and CCT denotes the cross-view and cross-modal transfer module.

### 3.2 Self-supervised Depth Estimation and VO

We adopt two networks to predict depth and pose, respectively, which are jointly optimized using the photometric loss and the smoothness loss in a self-supervised manner, following the baseline method [17]. We additionally devise a CGT scale loss to address the scale ambiguity problem of monocular depth and VO estimation. We describe the loss items of depth and pose networks in this section.

Self-supervised monocular depth and pose estimation is achieved by leveraging the geometry consistency among continuous frames. During training, the depth  $D_t$  of current frame  $I_t$  and the poses  $\{T_{t\to t-1}, T_{t\to t+1}\}$  between  $I_t$  and its adjacent frames  $\{I_{t-1}, I_{t+1}\}$  are used to obtain the reconstructed current frames  $\{\hat{I}_{t-1\to t}, \hat{I}_{t+1\to t}\}$  via the differentiable warping function  $\omega$  from  $\{I_{t-1}, I_{t+1}\}$ :

$$\hat{I}_{t+m\to t} = \omega(KT_{t\to t+m}D_tK^{-1}I_{t+m}), \ m \in \{-1,1\}.$$
(2)

Then, the photometric differences between  $I_t$  and its reconstructed counterparts  $\hat{I}_{t-1\to t}, \hat{I}_{t+1\to t}$  are minimized to train the depth and pose networks. We quantify the photometric differences using the SSIM and L1 losses:

$$\ell_{ph} = \min_{m \in \{-1,1\}} \frac{\alpha (1 - SSIM(I_t, I_{t+m \to t}))}{2} + (1 - \alpha) |I_t - \hat{I}_{t+m \to t}|, \qquad (3)$$

where  $\alpha$  is set to 0.85. Following our baseline method [17], we also take the perpixel minimum of the photometric loss and adopt the auto-masking strategy.

To overcome the discontinuity of the predicted depth map, a smoothness loss [16,17] is adopted based on the gradient of  $I_t$ :

$$\ell_{sm} = |\partial_x \mu_{D_t}| e^{-|\partial_x I_t|} + |\partial_y \mu_{D_t}| e^{-|\partial_y I_t|}, \tag{4}$$

where  $\mu_{D_t}$  denotes the normalized inverse depth. By minimizing the above losses, the depth network and the pose network can be optimized simultaneously.



Fig. 3: The demonstration of CGT scale loss. Fig. 4: The structure of cross-view and cross-modal transfer (CCT) module.

**CGT Scale Loss.** To accomplish scale-aware environment perception, we propose the cross-view geometric transformation-based (CGT) scale loss for depth estimation and VO by utilizing the scale information in the BEV layout. Since BEV layouts demonstrate the semantic occupancy in the BEV Cartesian plane, covering the range of Z meters in front of the ego vehicle and horizon-tally covering  $\frac{Z}{2}$  meters to the left and right, respectively. It provides a natural distance field z with each pixel having a metric distance value  $z_{ij}$  with respect to the ego vehicle, as shown in Fig. 3. By assuming that the BEV plane is a flat plane on the ground with its origin just below the origin of the ego vehicle coordinate system, the BEV plane can then be projected to FV using the camera extrinsic parameters via a homography transformation:

$$H_{bev}^{cam} = K T_{eqo}^{cam} T_{bev}^{ego}, \quad T_{eqo}^{cam} \text{ and } T_{bev}^{ego} \in \mathcal{SE}(3), \tag{5}$$

where cam, ego, bev represent the camera coordinate system, the ego vehicle coordinate system and the BEV ground system, respectively.  $T_{bev}^{ego}$  and  $T_{ego}^{cam}$  are the  $\mathcal{SE}(3)$  transformations that transform the BEV plane coordinate system to the ego-vehicle coordinate system and then transform to the camera coordinate system, respectively. Therefore, the BEV distance field z can be projected into FV as  $z^{fv}$  as shown in Fig. 3, which is then utilized to regulate the predicted depth d, leading to our proposed CGT scale loss:

$$\ell_{CGT} = \frac{1}{h_d w_d} \sum_{j=1}^{h_d} \sum_{i=1}^{w_d} \frac{|z_{ij}^{fv} - d_{ij}|}{z_{ij}^{fv}}, \quad if \ z_{ij}^{fv} \neq 0.$$
(6)

To learn the scale-aware depth and pose in a self-supervised manner, we take the weighted sum of  $\ell_{ph}$ ,  $\ell_{sm}$ , and  $\ell_{CGT}$  as the final depth-and-pose objective:

$$\ell_{dp} = \ell_{ph} + \ell_{sm} + \beta \cdot \ell_{CGT},\tag{7}$$

where  $\beta$  is a hyper-parameter and set to 0.1 empirically.

#### 3.3 BEV Layout Estimation

For layout estimation, an encoder-decoder network structure is adopted, following the prior work [54]. It is noteworthy that we use one shared encoder as the feature extractor and different decoders to learn BEV layouts of different semantic categories simultaneously instead of training networks for each category individually as in prior works [31,54,36]. In addition, a CCT module is designed to strengthen feature interaction and knowledge transfer between tasks, and impose 3D geometry information for the spatial reasoning in BEV. To regularize the layout network, we combine various loss items to form a hybrid loss and achieve a balanced optimization for different categories.

CCT Module. To enhance feature interaction and impose 3D geometry information for BEV perception, we devise CCT to investigate the correlation between the FV feature  $F_f$ , the BEV layout feature  $F_b$ , the retransformed front feature  $F'_{f}$  and the FV depth feature  $F_{d}$ , and refine the layout feature accordingly, as shown in Fig. 4. We describe CCT in two parts, i.e. CCT-CV and CCT-CM for the cross-view module and the cross-modal module, respectively. CCT is inspired by prior work [54], but uses different structures and modal information. In CCT,  $F_f$  and  $F_d$  are extracted by the encoders of the corresponding perception branches, while  $F_b$  is obtained by transforming  $F_f$  to BEV with a view projection MLP, which is then re-transformed to  $F'_{f}$  using the same MLP constrained by a cycle loss, following prior work [54]. All the features are set to the same size, i.e.,  $F_f, F_d, F_b, F'_f \in \mathbb{R}^{H \times W \times D}$ . In CCT-CV, a cross-attention mechanism is used to discover the geometry correspondence between FV and BEV features, which is then utilized to guide the FV information refinement and prepared for BEV reasoning. To fully exploit the FV image features,  $F_b$  and  $F_f$  are projected to patches  $Q_{bi} \in Q_b(i \in [1, ..., HW])$  and  $K_{bi} \in K_b(i \in [1, ..., HW])$ , acting as the Query and Key respectively. Then each location in the FV will retrieve the correlation from every location in BEV to form a correlation matrix:

$$C_b = \frac{Q_b K_b^T}{\sqrt{D}} \in \mathbb{R}^{H \times W \times H \times W},\tag{8}$$

where the normalization factor  $\frac{1}{\sqrt{D}}$  is used to restrict the value range. The crossview correspondence can be identified by finding the location with the largest correlation value in  $C_b$ , which is differentiable by using the softmax operation:

$$M_b = softmax(C_b) \in \mathbb{R}^{H \times W \times H \times W}.$$
(9)

Since  $F'_f$  is obtained by first transforming the FV features to BEV and then back to FV, it contains both FV and BEV information. We thus project  $F'_f$  to patches  $V_{bi} \in V_b(i \in [1, ..., HW])$  and concatenate it with  $F_f$  to provide the Value for the cross-view correlation after a convolution layer:

$$F_{cv} = Conv(Concat(F'_f, V_b)) \odot M_b.$$
<sup>(10)</sup>

Except for utilizing the FV features, we also deploy CCT-CM to impose 3D geometry information from  $F_d$ . Since  $F_d$  is extracted from FV images, it is reasonable to use  $F_f$  as the bridge to reduce the cross-modal gap and learn the correspondence between  $F_d$  and  $F_b$ . Thus, similar to CCT-CV,  $F_b$  and  $F_f$  are

regarded as the Query  $Q_d$  and the Key  $K_d$  to calculate the correlation matrix:

$$C_d = \frac{Q_d K_d^T}{\sqrt{D}} \in \mathbb{R}^{H \times W \times H \times W}.$$
(11)

 $F_d$  plays the role of Value so that we can acquire the valuable 3D geometry information correlated to the BEV information and further improve the accuracy of layout estimation. The final CCT-CM feature  $F_{cm}$  is then derived as:

$$M_d = softmax(C_d) \in \mathbb{R}^{H \times W \times H \times W},$$
  

$$F_{cm} = F_d \odot M_d.$$
(12)

In the end, the original FV feature  $F_f$ , the cross-view correlated feature  $F_{cv}$  and the cross-modal correlated feature  $F_{cm}$  are summed up as the input of the layout decoder branch to conduct subsequent learning:  $F_{out} = F_f + F_{cv} + F_{cm}$ .

**Hybrid Loss.** BEV layout estimation is a binary classification problem for each semantic category in the BEV grid to determine whether an area belongs to roads, cars, or backgrounds. Thus, this task is usually regularized by minimizing the difference  $Diff(\cdot)$  between the predictions  $L_{pred}$  and the ground truth  $L_{gt}$ using Cross-Entropy (CE) or L2 loss in prior works [54,20]:

$$\ell_{layout}^c = Diff(L_{pred} - L_{gt}), \ c \in \{c_{road}, c_{vehicle}\}.$$
(13)

In the process of exploring our joint learning framework to predict different layouts simultaneously, we observe that a great difference exists in the characteristics and distributions of different semantic categories. For characteristics, the road layout in driving scenes usually needs to be connected, while different vehicles instead must be separated. And for distributions, more scenes with straight roads are observed than scenes with turns, which is reasonable in realworld datasets. Such difference and imbalance increase the difficulty of BEV layout learning, especially for predicting different categories jointly, due to the failure of the simple CE or L1 losses in such circumstances. Thus, we incorporate several kinds of segmentation losses including the distribution-based CE loss, the region-based IoU loss, and the boundary loss into a hybrid loss to predict the layout for each category. First, the Weighted Binary Cross-Entropy Loss is adopted, which is most commonly used in semantic segmentation tasks:

$$\ell_{WBCE} = \frac{1}{h_l w_l} \sum_{n=1}^{h_l w_l} \sum_{m=1}^{M} -w_m [y_n^m \cdot \log x_n^m + (1 - y_n^m) \cdot \log(1 - x_n^m)], \quad (14)$$

where  $x_n$  and  $y_n$  denote the *n*-th predicted category value and the counterpart ground truth in a layout of size  $h_l w_l$ , respectively. M = 2 means whether a pixel belongs to foreground roads or cars, while  $w_m$  is the hyperparameter for tackling the issue of sampling imbalance between different labels. We set  $w_m$  to 5 and 15 for roads and vehicles respectively following [54].

Since the CE loss treats each pixel as an independent sample and thus neglects the interactions between nearby pixels, we then adopt the Soft IoU Loss  $\ell_{IoU}$  to ensure the connectivity within the region:

$$\ell_{IoU} = -\frac{1}{M} \sum_{m=1}^{M} \frac{\sum_{n=1}^{h_l w_l} x_n^m y_n^m}{\sum_{n=1}^{h_l w_l} (x_n^m + y_n^m + x_n^m y_n^m)}, \ M = 2.$$
(15)

For the integrity of the region edges, we further use the Boundary Loss  $\ell_{Bound}$  to constrain the learning of the boundary predictions, which has proven effective for mitigating issues related to regional losses in highly unbalanced segmentation problems such as medical image segmentation [22,28,53]. It is calculated as the Hadamard product of the signed distance (SDF) map of the ground truth layout:

$$\ell_{Bound} = M_{SDF}(L_{gt}) \odot L_{pred},$$

$$M_{SDF}(L_{gt}) = \begin{cases} -\inf_{\substack{y \in L_{gt}^{b} \\ y \in L_{gt}^{b}}} ||x - y||_{2}, x \in L_{gt}^{in}, \\ +\inf_{\substack{y \in L_{gt}^{b} \\ y \in L_{gt}^{b}}} ||x - y||_{2}, x \in L_{gt}^{out}, \\ 0, x \in L_{gt}^{b}, \end{cases}$$
(16)

where  $L_{gt}^{in}, L_{gt}^{out}$ , and  $L_{gt}^{b}$  represent regions inside, outside and at the foreground object boundaries in the ground truth, respectively.  $||x - y||_2$  is the Euclidian distance between x and y.

The final loss of the layout estimation for each category then reads:

$$\ell_{layout}^{c_i} = \ell_{WBCE}^{c_i} + \lambda \cdot \ell_{IoU}^{c_i} + \lambda \cdot \ell_{Bound}^{c_i}, \tag{17}$$

where  $\lambda = 20$  and  $c_i \in \{c_{road}, c_{vehicle}\}$ .

Different from prior works [54,36], our joint learning framework predict all semantic categories simultaneously instead of training a network separately for each category. The final optimization loss for our layout network reads:

$$\ell_{layout} = \sum_{c_i} \ell_{layout}^{c_i}, \ c_i \in \{c_{road}, c_{vehicle}\}.$$
(18)

# 4 Experiments

Since there is no previous work that accomplishes depth estimation, visual odometry and BEV layout estimation simultaneously, we evaluate the three tasks on their corresponding benchmarks and compared our method with the SOTA methods of each task. In addition, extensive ablation studies are performed to verify the effectiveness of our joint learning network architecture and loss items.

#### 4.1 Datasets.

We evaluate our JPerceiver on three driving scene datasets, i.e., Argoverse [8], Nuscenes [5] and KITTI [15]. Argoverse and Nuscenes are relatively newly published autonomous driving datasets that provide high-resolution BEV semantic

Mathada	Argover	se Road	Argoverse Vehicle		
Methods	mIoU(%)	mAP(%)	mIoU(%)	mAP(%)	
VED [26]	72.84	78.11	24.16	36.83	
VPN [34]	71.07	86.83	16.58	39.73	
Monolay [31]	73.25	84.56	32.58	51.06	
PYVA [54]	76.51	87.21	48.48	64.04	
JPeceiver("1-1")	77.86	90.59	49.94	65.44	
JPerceiver("1-2")	77.50	90.21	49.45	65.84	
Baseline("1-1")	76.66	87.17	46.97	63.36	
+CCT-CV	77.76	88.42	49.33	64.05	
+CCT-CV+CCT-CM	77.80	89.00	49.39	64.86	
Ours("1-1"+CCT+HLoss)	77.86	90.59	49.94	65.44	
Baseline("1-2")	76.52	86.54	42.77	59.39	
+CCT-CV	76.91	87.19	46.46	61.02	
+CCT-CV+CCT-CM	77.38	88.40	47.19	61.43	
$\overline{Ours("1-2"+CCT(S)+HLoss)}$	76.81	89.39	48.06	63.61	
Ours("1-2"+CCT+HLoss)	77.50	90.21	49.45	65.84	

Table 1: Quantitative comparisons (top part) and ablation study (bottom part) on Argoverse [8]. "CCT-CV" and "CCT-CM" denote the cross-view and cross-modal part in CCT.

occupancy labels for roads and vehicles. We evaluate the performance of BEV layout estimation on Argoverse with 6,723 training images and 2,418 validation images within the range of  $40m \times 40m$ , and on Nuscenes with 28,130 training samples and 6,019 validation samples under two settings (in Supplementary). The ablation study for layout estimation is performed on Argoverse. Two semantic categories are included in the evaluated BEV layouts, i.e. roads and vehicles. We adopt the mean of Intersection-over-Union (mIoU) and Average Precision (mAP) as the evaluation metrics, following prior works [31,54]. Due to the insufficient annotations in KITTI, we follow prior works to use three splits for the tasks, i.e., the KITTI Odometry split (15,806 and 6,636 items for training and validation) for depth, VO, as well as road layout estimation, the KITTI Raw split (10,156 and 5,074 items for training and validation) for road layout estimation, and the KITTI 3D Object split (3,712 and 3,769 items for training and validation) for vehicle layout estimation, all within the range of  $40m \times 40m$ .

**Implementation Details.** We adopt the encoder-decoder structure for the three networks, all using pre-trained ResNet18 [19] as the encoder backbone except that we modify the pose network to take two-frame pairs as input. Following the prior method [54], the input sizes of the depth and layout networks are both set to  $1024 \times 1024$ , while a smaller input size  $192 \times 640$  is used for the pose network to save computation resources since its outputs are not pixel-wise. Our model is implemented in PyTorch [35] and trained for 80 epochs using Adam [23], with a learning rate of  $10^{-4}$  for the first 50 epochs and  $10^{-5}$  for the remaining epochs. The details of the network structure are presented in Supplementary. And there is a potential promising improvement using more powerful network backbones (e.g. transformer [51,57]) in our method.



Fig. 5: The qualitative results of depth and layout on the three Fig. 6: Visualization of features datasets. White and blue regions indicate road and vehicle and attention maps aligned with layouts. The corresponding views.

#### 4.2 Layout Estimation

**Argoverse.** We first quantitatively compare JPerceiver with SOTA methods on Argoverse [8], including VPN [34], Monolay [31] and PYVA [54]. We compare two variants of our method, i.e., (1) JPerceiver("1-1") that trains two layout estimation networks separately for each category, which is also the common practice in the compared methods, and (2) JPerceiver("1-2") that predicts the two kinds of layouts jointly with one shared encoder. As shown in the top part of Table 1, the two variants both show superiority over other approaches in the road and vehicle layout estimation. Moreover, JPerceiver("1-2") achieves comparable performance with JPerceiver("1-1"), even if JPerceiver("1-2") only uses a shared encoder and thus is more efficient in terms of memory and computation.

**Ablation.** To investigate the effect of the component and the hybrid loss in our model, we conduct ablation studies on Argoverse and report the results in the bottom part of Table 1. We take a basic "1-1" structure as our baseline, i.e., an encoder-decoder structure trained for each semantic category with a CE loss. We then use CCT and jointly train the depth network, pose network, as well as layout network together. Specifically, we ablate the cross-view module CCT-CV and the cross-modal module CCT-CM, respectively. As can be seen, both the cross-view and the cross-modal modules improve the performance of the baseline, and the complete CCT brings a gain of 1.14% mIoU and 1.83% mAP for road layout estimation and 2.42% mIoU and 1.5% mAP for vehicle layout estimation, respectively. After using the hybrid loss (HLoss), the performance reaches the best, i.e., 77.86% mIoU and 90.59% mAP for roads as well as 49.94% mIoU and 65.44% mAP for vehicles, respectively. We then conduct the same ablation study on the "1-2" structures. It is observed that using one encoder to learn representations for both two semantic layouts significantly decreases the performance of baseline models, especially for the vehicle layout, i.e., from 46.97% mIoU and 63.36% mAP to 42.77% mIoU and 59.39% mAP. After using the proposed CCT module and hybrid loss, the performance drop can be recovered, where the final model achieves comparable results as the "1-1" structure. We further investigate the performance of using a shared CCT module for the two semantic layouts, denoted as "CCT(S)" in Table 1. As can be seen, its results are much worse

than using separate CCT for each category. It is probably due to the distinct geometry characteristics of road and vehicle, where different regions should be paid attention to as illustrated in Fig. 6.

**KITTI.** We train our layout network for roads on the KITTI Odometry and the KITTI Raw splits, and the layout for vehicles on the KITTI Object split. As shown in Table 2 (left), our performance on KITTI Odometry is superior to other methods w.r.t. both mIoU and mAP. The evaluation results on KITTI Raw are listed in Table 2 (middle). We report our reproduced result of PYVA [54] within parentheses because their reported results are trained with processed ground truth. Even so, our method still outperforms their reported results with a gain of 5.39% in mAP. We observe a performance degradation from KITTI Odometry to KITTI Raw, potentially because the ground truth of the latter comes from registered semantic segmentation of Lidar scans while the former obtains the ground truth from the more accurate Semantic KITTI dataset [3], both collected by [31]. For vehicle layout estimation, we show the quantitative results in Table 2 (right). Our method exceeds other works by a large margin, i.e., 2.06% and 6.97% w.r.t. mIoU and mAP.

#### 4.3 Depth Estimation and Visual Odometry

KITTI presents the most commonly used dataset for depth estimation and VO. And we report our scale-aware depth and VO results on KITTI Odometry.

**Depth Estimation.** We compare with several self-supervised depth estimation methods on the KITTI Odometry test set, shown in Table 3. The scaling factor is calculated as the average of all depth map scale ratios, which is the ratio of the median of depth values and the median of ground truth values. We use Monodepth2 [17] as the baseline and DNet [52] as a representative competitor of the scale-aware methods. Though Monodepth2 [17] achieves good up-to-scale accuracy, however, without the scaling factor, its performance significantly degrades. DNet [52] predicts a camera height during inference and calculates the ratio between the ground truth camera height and the predicted one to get the scaling factor. However, its output depths still need to be scaled. Differently, thanks to the CGT scale loss, our depth prediction naturally contains the absolute metric scale and does not require any scaling operation. As shown in the top part of Table 3, our scale factor computed during inference is 1.065 with a variance of 0.071, while a comparable precision is also achieved.

Ablation. We further conduct the ablation study for depth and report the results in the bottom part of Table 3 with the input resolution of  $512 \times 512$ . Introducing CCT on the baseline structure boosts the depth estimation results no matter with or without scaling, and the scaling factor is similar to the baseline counterpart. We then add the CGT scale loss to the baseline to validate its feasibility, resulting in a nearly perfect scaling ratio. However, since our CGT scale loss only takes regional pixels in ground areas into account instead of using all pixels and is based on an assumption that the ground plane is flat, a less accurate result is observed for the overall prediction. Our full model achieves comparable depth performance with the baseline but up to a metric scale. Of note

Table 2: Quantitative comparisons of BEV layout estimation results on KITTI Odometry, KITTI Raw and KITTI 3D Object.

Methoda	KITTI Odometry Road		KITTI R	aw Road	KITTI 3D Object		
Methous	mIoU(%) $mAP(%)$		mIoU(%)	mAP(%)	mIoU(%)	mAP(%)	
VED [26]	65.74	67.84	58.41	66.01	20.45	22.5	
Mono3D [9]	-	-	-	-	17.11	26.62	
OFT [40]	-	-	-	-	25.24	34.69	
VPN [34]	66.81	81.79	59.58	79.07	16.80	35.54	
Monolay [31]	76.81	85.25	66.02	75.73	30.18	45.91	
PYVA [54]	77.49	86.69	68.34(65.52)	80.78 (79.52)	38.79	50.26	
JPeceiver	78.13	89.57	66.39	86.17	40.85	57.23	

Table 3: Quantitative comparisons and ablation study for depth estimation. "w" and "w/o" denote evaluation results with or without rescaling by the scale factor, which is calculated during inference.

Methods	Resolution	Scaling	Abs Rel $(\downarrow)$	Sq $\operatorname{Rel}(\downarrow)$	$RMSE(\downarrow)$	RMSE $\log(\downarrow)$	Scale factor
Monodepth2 [17]	$1024 \times 1024$	w	0.113	0.526	3.656	0.181	$42.044 \pm 0.076$
		w/o	0.976	13.687	17.128	3.754	_
DNet [52]	$1024 \times 1024$	w	0.121	0.582	3.762	0.192	$34.393 \pm 0.077$
Divet [02]	1024 × 1024	w/o	0.970	13.528	17.028	3.545	-
IPorcoivor	$1024 \times 1024$	W	0.116	0.517	3.573	0.180	$1.065\pm0.071$
JFerceiver		w/o	0.112	0.559	3.817	0.196	_
Baseline	$512 \times 512$	w	0.120	0.550	3.670	0.184	$39.452 \pm 0.077$
		w/o	0.974	13.616	17.073	0.179	_
+CCT	$512 \times 512$	w	0.108	0.505	3.574	0.179	$37.711 \pm 0.067$
		w/o	0.973	13.616	17.083	3.645	-
+scale loss	$512 \times 512$	w	0.135	0.633	3.860	0.194	$1.088 \pm 0.093$
		w/o	0.125	0.643	4.092	0.211	_
Ours	$512 \times 512$	w	0.128	0.574	3.739	0.189	$1.099 \pm 0.085$
		w/o	0.122	0.628	3.952	0.205	–

is that all variants with our scale loss obtain lower Abs Rel error without scaling while worse results in other metrics compared with the results with scaling. It may be because the calculation of our Scale loss is the same as the Abs Rel metric. Different ways of calculating and utilizing the scale loss on other baselines might be explored to further improve the estimation accuracy in future work.

**Visual Odometry.** We train Perceiver on the KITTI Odometry sequences 01-06 and 08-09, and use the sequences 07 and 10 as our test set for evaluating our model for VO. We compare with several self-supervised visual odometry methods in Table 4, including SfMLearner [60], GeoNet [55], Monodepth2 [17], SC-Sfmlearner [4], which are trained on sequence 00-08. "Scaling" means the scaling method is used during inference. "GT" means the scaling factor for correcting the predictions comes from the ground truth. Dnet [52], LSR [46] and Ours all borrow information from the road plane to recover the scale but in different ways. Of note is that Dnet [52] needs invisible ground plane to predict camera height during inference for scale recovery, which is not required by our method. While LSR [46] incorporates a front-view ground plane estimation task, their networks are trained in a serial way, i.e. unscaled depth and VO network, plane segmentation network, and then scaled depth and VO network. In comparison, our method can produce scaled VO results during inference without any hint by using our CGT scale loss. In addition, our method is superior to other competitors w.r.t. the average translational and rotational RMSE drift metrics. The

Table 5: The Analysis of model complexity with input resolution  $512 \times 512$  using one single GPU.

Methods	Task	Params(M)	Flops(G)	FPS(BS=1)	FPS(BS=6)
Monodepth2 [17]	depth & pose	39.33	26.93	30.3	36.9
PYVA [54]	layout	29.73	20.42	65.2	72
	depth & pose & layout	69.06	47.35	15.6	18.6
JPerceiver	depth & pose & layout	57.15	37.69	19.8	26.9

comparison of VO trajectories with other methods without rescaling is shown in Fig. 7, which further proves the effectiveness of our scale loss.



Fig. 7: The comparison of VO trajectories on sequence 07.

Table 4: The comparison of Visual Odometry. $t_{err}$ is the average
translational RMSE drift (%) on length from 100, 200 to 800 m,
and $r_{err}$ is average rotational RMSE drift ( $^{\circ}/100m$ ) on length
from 100, 200 to 800 m.

Methode	Scoling	Seque	nce 07	Sequence 10		
wiethous	Scanng	$t_{err}$	$r_{err}$	$t_{err}$	$r_{err}$	
SfMLearner [60]	GT	12.61	6.31	15.25	4.06	
GeoNet [55]	GT	8.27	5.93	20.73	9.04	
Monodepth2 [17]	GT	8.85	5.32	11.60	5.72	
SC-Sfmlearner [4]	GT	8.29	4.53	10.74	4.58	
Dnet [52]	Camera height	-	-	13.98	4.07	
LSR [46]	None	-	-	10.54	4.03	
JPerveiver	None	4.57	2.94	7.52	3.83	

## 4.4 Model Complexity

We compare the complexity of our JPerciever with the single-task competitors and the simply combined model in Table 5. Since our method can predict the depth, VO, and BEV layouts of the two semantic categories simultaneously, it is less complex in terms of parameters (M), computations (FLOPs), and inference speed (FPS) compared with the combined model while producing better prediction in all the three tasks. Besides, our JPerciever benefits more from parallel acceleration as shown in the last column, where a batch size of 6 is used.

## 5 Conclusion

In this paper, we propose a joint perception framework named JPerceiver for the autonomous driving scenarios, which accomplishes scale-aware depth estimation, visual odometry, and also BEV layout estimation of multiple semantic categories simultaneously. To realize the joint learning of multiple tasks, we introduce a cross-view and cross-modal transfer module and fully make use of the metric scale from the BEV layout to devise a cross-view geometry transformation-based scale loss to obtain scale-aware predictions. Our method achieves better performance towards all the above three perception tasks using less computation resource and training time. We hope our work can provide valuable insight to the future study of designing more effective joint environment perception model.

Acknowledgement Ms. Haimei Zhao, Dr. Jing Zhang, and Mr. Sen Zhang are supported by ARC FL-170100117, DP-180103424, IC-190100031, and LE-200100049.

# References

- 1. Translating images into maps. arXiv preprint arXiv:2110.00966 (2021)
- Badki, A., Gallo, O., Kautz, J., Sen, P.: Binary ttc: A temporal geofence for autonomous navigation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 12946–12955 (June 2021)
- Behley, J., Garbade, M., Milioto, A., Quenzel, J., Behnke, S., Stachniss, C., Gall, J.: Semantickitti: A dataset for semantic scene understanding of lidar sequences. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9297–9307 (2019)
- Bian, J.W., Li, Z., Wang, N., Zhan, H., Shen, C., Cheng, M.M., Reid, I.: Unsupervised scale-consistent depth and ego-motion learning from monocular video. Advances in Neural Information Processing Systems (2019)
- Caesar, H., Bankiti, V., Lang, A.H., Vora, S., Liong, V.E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., Beijbom, O.: nuscenes: A multimodal dataset for autonomous driving. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 11621–11631 (2020)
- Casas, S., Sadat, A., Urtasun, R.: Mp3: A unified model to map, perceive, predict and plan. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 14403–14412 (June 2021)
- Chang, M.F., Lambert, J., Sangkloy, P., Singh, J., Bak, S., Hartnett, A., Wang, D., Carr, P., Lucey, S., Ramanan, D., Hays, J.: Argoverse: 3d tracking and forecasting with rich maps. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019)
- Chang, M.F., Lambert, J., Sangkloy, P., Singh, J., Bak, S., Hartnett, A., Wang, D., Carr, P., Lucey, S., Ramanan, D., et al.: Argoverse: 3d tracking and forecasting with rich maps. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8748–8757 (2019)
- Chen, X., Kundu, K., Zhang, Z., Ma, H., Fidler, S., Urtasun, R.: Monocular 3d object detection for autonomous driving. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2147–2156 (2016)
- Chen, Z., Zhang, J., Tao, D.: Progressive lidar adaptation for road detection. IEEE/CAA Journal of Automatica Sinica 6(3), 693–702 (2019)
- Chen, Z., Wang, C., Yuan, B., Tao, D.: Puppeteergan: Arbitrary portrait animation with semantic-aware appearance transformation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13518–13527 (2020)
- Chi, C., Wang, Q., Hao, T., Guo, P., Yang, X.: Feature-level collaboration: Joint unsupervised learning of optical flow, stereo depth and camera motion. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2463–2473 (June 2021)
- Dwivedi, I., Malla, S., Chen, Y.T., Dariush, B.: Bird's eye view segmentation using lifted 2d semantic features. In: British Machine Vision Conference (BMVC). pp. 6985–6994 (2021)
- Fu, H., Gong, M., Wang, C., Batmanghelich, K., Tao, D.: Deep ordinal regression network for monocular depth estimation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2002–2011 (2018)
- Geiger, A., Lenz, P., Urtasun, R.: Are we ready for autonomous driving? the kitti vision benchmark suite. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3354–3361. IEEE (2012)

- 16 Haimei Zhao et al.
- Godard, C., Mac Aodha, O., Brostow, G.J.: Unsupervised monocular depth estimation with left-right consistency. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 270–279 (2017)
- 17. Godard, C., Mac Aodha, O., Firman, M., Brostow, G.J.: Digging into selfsupervised monocular depth prediction (October 2019)
- Hang, P., Lv, C., Xing, Y., Huang, C., Hu, Z.: Human-like decision making for autonomous driving: A noncooperative game theoretic approach. IEEE Transactions on Intelligent Transportation Systems 22(4), 2076–2087 (2020)
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 770–778 (2016)
- Hu, A., Murez, Z., Mohan, N., Dudas, S., Hawke, J., Badrinarayanan, V., Cipolla, R., Kendall, A.: FIERY: Future instance segmentation in bird's-eye view from surround monocular cameras. In: Proceedings of the International Conference on Computer Vision (ICCV) (2021)
- Huang, C., Lv, C., Hang, P., Xing, Y.: Toward safe and personalized autonomous driving: Decision-making and motion control with dpf and cdt techniques. IEEE/ASME Transactions on Mechatronics 26(2), 611–620 (2021)
- Kervadec, H., Bouchtiba, J., Desrosiers, C., Granger, E., Dolz, J., Ayed, I.B.: Boundary loss for highly unbalanced segmentation. In: International conference on medical imaging with deep learning. pp. 285–296. PMLR (2019)
- Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
- Klingner, M., Termöhlen, J.A., Mikolajczyk, J., Fingscheidt, T.: Self-supervised monocular depth estimation: Solving the dynamic object problem by semantic guidance. In: European Conference on Computer Vision. pp. 582–600. Springer (2020)
- Li, R., Wang, S., Long, Z., Gu, D.: Undeepvo: Monocular visual odometry through unsupervised deep learning. In: IEEE International Conference on Robotics and Automation. pp. 7286–7291. IEEE (2018)
- 26. Lu, C., van de Molengraft, M.J.G., Dubbelman, G.: Monocular semantic occupancy grid mapping with convolutional variational encoder-decoder networks. IEEE Robotics and Automation Letters 4(2), 445–452 (2019)
- Luo, C., Yang, X., Yuille, A.: Self-supervised pillar motion learning for autonomous driving. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3183–3192 (2021)
- Ma, J., Wei, Z., Zhang, Y., Wang, Y., Lv, R., Zhu, C., Gaoxiang, C., Liu, J., Peng, C., Wang, L., et al.: How distance transform maps boost segmentation cnns: an empirical study. In: Medical Imaging with Deep Learning. pp. 479–492. PMLR (2020)
- Mahjourian, R., Wicke, M., Angelova, A.: Unsupervised learning of depth and egomotion from monocular video using 3d geometric constraints. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5667–5675 (2018)
- Mallot, H.A., Bülthoff, H.H., Little, J., Bohrer, S.: Inverse perspective mapping simplifies optical flow computation and obstacle detection. Biological cybernetics 64(3), 177–185 (1991)
- Mani, K., Daga, S., Garg, S., Narasimhan, S.S., Krishna, M., Jatavallabhula, K.M.: Monolayout: Amodal scene layout from a single image. In: The IEEE Winter Conference on Applications of Computer Vision. pp. 1689–1697 (2020)

17

- McCraith, R., Neumann, L., Vedaldi, A.: Calibrating self-supervised monocular depth estimation. arXiv preprint arXiv:2009.07714 (2020)
- Nistér, D., Naroditsky, O., Bergen, J.: Visual odometry. In: Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004. vol. 1, pp. I–I. Ieee (2004)
- Pan, B., Sun, J., Leung, H.Y.T., Andonian, A., Zhou, B.: Cross-view semantic segmentation for sensing surroundings. IEEE Robotics and Automation Letters 5(3), 4867–4873 (2020)
- Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., Lerer, A.: Automatic differentiation in pytorch (2017)
- Philion, J., Fidler, S.: Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. In: European Conference on Computer Vision. pp. 194–210. Springer (2020)
- Phillips, J., Martinez, J., Barsan, I.A., Casas, S., Sadat, A., Urtasun, R.: Deep multi-task learning for joint localization, perception, and prediction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4679–4689 (June 2021)
- Ranjan, A., Jampani, V., Balles, L., Kim, K., Sun, D., Wulff, J., Black, M.J.: Competitive collaboration: Joint unsupervised learning of depth, camera motion, optical flow and motion segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 12240–12249 (2019)
- Reading, C., Harakeh, A., Chae, J., Waslander, S.L.: Categorical depth distribution network for monocular 3d object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8555–8564 (2021)
- 40. Roddick, T., Kendall, A., Cipolla, R.: Orthographic feature transform for monocular 3d object detection. arXiv preprint arXiv:1811.08188 (2018)
- Schön, M., Buchholz, M., Dietmayer, K.: Mgnet: Monocular geometric scene understanding for autonomous driving. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 15804–15815 (2021)
- 42. Shu, C., Yu, K., Duan, Z., Yang, K.: Feature-metric loss for self-supervised learning of depth and egomotion. In: ECCV (2020)
- Simond, N., Parent, M.: Obstacle detection from ipm and super-homography. In: 2007 IEEE/RSJ International Conference on Intelligent Robots and Systems. pp. 4283–4288. IEEE (2007)
- 44. Thavamani, C., Li, M., Cebron, N., Ramanan, D.: Fovea: Foveated image magnification for autonomous navigation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 15539–15548 (October 2021)
- 45. Torralba, A., Oliva, A.: Depth estimation from image structure. IEEE Transactions on pattern analysis and machine intelligence **24**(9), 1226–1238 (2002)
- 46. Wagstaff, B., Kelly, J.: Self-supervised scale recovery for monocular depth and egomotion estimation. arXiv preprint arXiv:2009.03787 (2020)
- 47. Wang, C., Miguel Buenaposada, J., Zhu, R., Lucey, S.: Learning depth from monocular videos using direct methods. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2022–2030 (2018)
- Wang, H., Cai, P., Fan, R., Sun, Y., Liu, M.: End-to-end interactive prediction and planning with optical flow distillation for autonomous driving. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops. pp. 2229–2238 (June 2021)
- 49. Wang, S., Clark, R., Wen, H., Trigoni, N.: Deepvo: Towards end-to-end visual odometry with deep recurrent convolutional neural networks. In: 2017 IEEE in-

ternational conference on robotics and automation (ICRA). pp. 2043–2050. IEEE (2017)

- Wang, Y., Mao, Q., Zhu, H., Zhang, Y., Ji, J., Zhang, Y.: Multi-modal 3d object detection in autonomous driving: a survey. arXiv preprint arXiv:2106.12735 (2021)
- Xu, Y., Zhang, Q., Zhang, J., Tao, D.: Vitae: Vision transformer advanced by exploring intrinsic inductive bias. Advances in Neural Information Processing Systems 34, 28522–28535 (2021)
- Xue, F., Zhuo, G., Huang, Z., Fu, W., Wu, Z., Ang, M.H.: Toward hierarchical self-supervised monocular absolute depth estimation for autonomous driving applications. In: 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 2330–2337. IEEE (2020)
- 53. Xue, Y., Tang, H., Qiao, Z., Gong, G., Yin, Y., Qian, Z., Huang, C., Fan, W., Huang, X.: Shape-aware organ segmentation by predicting signed distance maps. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 34, pp. 12565–12572 (2020)
- 54. Yang, W., Li, Q., Liu, W., Yu, Y., Ma, Y., He, S., Pan, J.: Projecting your view attentively: Monocular road scene layout estimation via cross-view transformation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15536–15545 (2021)
- 55. Yin, Z., Shi, J.: Geonet: Unsupervised learning of dense depth, optical flow and camera pose. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1983–1992 (2018)
- 56. Zhang, J., Tao, D.: Empowering things with intelligence: a survey of the progress, challenges, and opportunities in artificial intelligence of things. IEEE Internet of Things Journal 8(10), 7789–7817 (2020)
- 57. Zhang, Q., Xu, Y., Zhang, J., Tao, D.: Vitaev2: Vision transformer advanced by exploring inductive bias for image recognition and beyond. arXiv preprint arXiv:2202.10108 (2022)
- Zhang, S., Zhang, J., Tao, D.: Towards scale consistent monocular visual odometry by learning from the virtual world. In: 2022 IEEE International Conference on Robotics and Automation (ICRA) (2022)
- Zhao, H., Bian, W., Yuan, B., Tao, D.: Collaborative learning of depth estimation, visual odometry and camera relocalization from monocular videos. In: IJCAI. pp. 488–494 (2020)
- Zhou, T., Brown, M., Snavely, N., Lowe, D.G.: Unsupervised learning of depth and ego-motion from video. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1851–1858 (2017)
- Zhuang, Z., Li, R., Jia, K., Wang, Q., Li, Y., Tan, M.: Perception-aware multisensor fusion for 3d lidar semantic segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 16280–16290 (October 2021)
- Zou, Y., Ji, P., Tran, Q.H., Huang, J.B., Chandraker, M.: Learning monocular visual odometry via self-supervised long-term modeling. In: European Conference on Computer Vision. pp. 710–727. Springer (2020)
- Zou, Y., Luo, Z., Huang, J.B.: Df-net: Unsupervised joint learning of depth and flow using cross-task consistency. In: The European Conference on Computer Vision. pp. 36–53 (2018)