Visual Attention-based Self-supervised Absolute Depth Estimation using Geometric Priors in Autonomous Driving

Jie Xiang¹, Yun Wang², Lifeng An², Haiyang Liu², Zijun Wang² and Jian Liu²

Abstract—Although existing monocular depth estimation methods have made great progress, predicting an accurate absolute depth map from a single image is still challenging due to the limited modeling capacity of networks and the scale ambiguity issue. In this paper, we introduce a fully Visual Attentionbased Depth (VADepth) network, where spatial attention and channel attention are applied to all stages. By continuously extracting the dependencies of features along the spatial and channel dimensions over a long distance, VADepth network can effectively preserve important details and suppress interfering features to better perceive the scene structure for more accurate depth estimates. In addition, we utilize geometric priors to form scale constraints for scale-aware model training. Specifically, we construct a novel scale-aware loss using the distance between the camera and a plane fitted by the ground points corresponding to the pixels of the rectangular area in the bottom middle of the image. Experimental results on the KITTI dataset show that this architecture achieves the state-of-the-art performance and our method can directly output absolute depth without post-processing. Moreover, our experiments on the SeasonDepth dataset also demonstrate the robustness of our model to multiple unseen environments.

Index Terms—Deep learning for visual perception, computer vision for transportation, range sensing.

I. INTRODUCTION

ONOCULAR Depth Estimation (MDE) is a fundamen-tal problem in the field of refers to predicting the corresponding depth map from a single image. Depth Estimation has broad application prospects in autonomous driving, such as 3D object detection [1], scene understanding [2], and obstacle avoidance [3]. Compared to other ways of measuring depth [4], [5], monocular depth estimation has unique advantages in obtaining dense depth maps at low cost. Therefore, MDE has aroused the interest of many researchers.

Manuscript received: May 1, 2022; Revised August 2, 2022; Accepted September 17, 2022.

This paper was recommended for publication by Editor Cesar Cadena Lerma upon evaluation of the Associate Editor and Reviewers' comments. This work was supported by the National Key Research and Development Program of China (No. 2021YFB2501403). (Corresponding author: Jie Xi-

ang.) ¹Jie Xiang is with the Institute of Microelectronics, Chinese Academy of Sciences, Beijing 100029, China, and with the School of Electronic, Electrical and Communication Engineering, University of Chinese Academy of Sciences, Beijing 100049, China. (Email: xiangjie@ime.ac.cn).

²Yun Wang, Lifeng An, Haiyang Liu, Zijun Wang and Jian Liu are with the Institute of Microelectronics, Chinese Academy of Sciences, Beijing 100029, China. (Email: {wangyun, anlifeng, liuhaiyang, wangzijun, liujian}@ime.ac.cn).

Digital Object Identifier (DOI): 10.1109/LRA.2022.3210298.

In the past few years, many deep learning-based MDE methods [6]-[23] have emerged and continual improvements [7]–[23] have been made. These methods can be divided into supervised methods [6]-[9] and self-supervised methods [10]-[23]. Unlike supervised methods, self-supervised methods do not need to rely on ground truth depth for training. Furthermore, existing self-supervised MDE methods can be divided into two categories: self-supervised stereo training methods [10]–[12] and monocular training methods [13]– [23]. The former use stereo image pairs while the latter use monocular videos as training data. In contrast to stereo training, monocular training is a more general form of selfsupervised methods and easier to obtain training data.

However, existing self-supervised monocular training methods still have difficulties in predicting accurate metric depth. This unsatisfied performance may come from two reasons: the limited modeling power of the model and the limited supervision of loss functions. Many existing self-supervised depth estimation networks [16]–[18], [23], [24] adopt residual networks [25] based on an encoder-decoder architecture [26]. Such a depth network shows a powerful capacity for capturing local information and representing hierarchical abstract features, but is insufficient to distinguish important details from interfering information. Due to these characteristics of the model, it is difficult for these methods to predict accurate depth estimation results on patterned surfaces or tiny structures. Thus, we introduce a fully Visual Attention-based Depth (VADepth) network. Specifically, we build a decoder based on the designed visual attention blocks and use a visual attention network followed by a dual attention module as the encoder. Benefiting from the ability to dynamically handle long-range dependencies brought by combining channel attention and spatial attention at all stages of the encoder-decoder architecture, our model not only continuously preserves important details but also effectively filters out noise.

On the other hand, most depth networks [13]-[22] trained on monocular videos only output relative depth due to lacking scale constraints. The commonly used scale recovery method called median scaling [13] requires ground truth depth data that is unavailable at test time in many practical applications. Training a model that directly produces absolute depth in metric units can omit the scale recovery operation like median scaling. An intuitive idea to this end is to use object size [27] but it suffers from seeking ubiquitous fixed-size objects and detecting the appearance size of detected objects. A simple but effective approach is to use camera height with respect to ground surface because the ground truth camera height in autonomous driving scenarios is usually fixed and known in advance. Furthermore, we also observe that the middle bottom regions of almost all images captured in these scenarios belong to the ground surface. Considering that not all points on the ground surface are coplanar, we propose to detect all coplanar points on the ground surface from the estimated depth, based on the assumption that a small rectangular region in the middle bottom of any image is part of the ground surface. Using the height errors calculated based on detection results, we construct an absolute scale loss. Combining the scale loss with photometric reconstruction loss [24], the visual attention-based depth network is able to output accurate absolute depth map without post-processing.

In summary, we make the contributions as follows:

- A novel depth network architecture that fully utilizes the long-range perception properties of visual attention to better perceive the scene structure and predict more accurate depth maps.
- An absolute scale loss that leverages the geometric priors in autonomous driving scenarios to train a scale-aware depth estimation model.
- We conduct extensive experiments to verify the effectiveness of our network structure and loss function. Experimental results [28] show that our method outperforms previous state-of-the-art methods on the KITTI dataset and our model generalizes well on the unseen Season-Depth dataset [29] that contains multiple challenging environments without fine-tuning.

II. RELATED WORK

A. Monocular Depth Estimation

Early deep learning-based works [6], [7] on monocular depth estimation mainly focused on supervised methods and dramatically improved the state-of-the-art performance. However, ground truth depth maps for supervised methods are hard to obtain. To avoid the heavy work of collecting ground truth depth, Garg et al. [10] proposed a self-supervised method inspired by view synthesis task [30], which leverages stereo image pairs to construct the reconstruction loss. Compared to stereo image pairs, monocular videos are much easier to obtain. Therefore, SfMLearner [13] leveraged monocular videos to train MDE models. Unlike stereo training approaches [10]–[12], monocular training methods [13], [16] need to simultaneously estimate ego-motion and depth at training time.

Recently, monocular training methods have been extensively studied. Some works [14], [15], [17], [20], [24], [31] explored to design more effective loss functions based on the photometric reconstruction loss. For example, edge-aware depthnormal consistency loss was used in [14], [15]. Godard et al. [24] proposed a per-pixel minimum reprojection loss to solve the occlusion problem, a multi-scale loss for alleviating the local gradient problem, and the auto-masking method to filter out stationary pixels. Based on Mondepth2 [24], the use of feature-metric loss [31], temporal geometric consistency [17] and semantic-depth consistency [20] further improved monocular depth estimation accuracy. Our method also follows the general pipeline of Monodepth2 [24].

Besides, some works [22], [32] proposed more powerful convolutional networks to enhance the accuracy of depth estimation. Beyond ResNet-based U-Net, Packnet-sfm [22] leveraged 3D convolutions to improve representational capacity of the network. HR-Depth [32] introduced more dense skip-connections to better fuse multi-scale feature maps. Because of the improvement of loss functions and network architectures, the performance of self-supervised monocular depth estimation has been improved a lot. However, these methods still face challenges in predicting metrically accurate depth due to the scale ambiguity issue and lacking the ability to distinguish important features from noise.

B. Scale Ambiguity

Because of lacking scale constraints, monocular vision suffers from the scale ambiguity issue. In order to produce absolute depth at test time, some MDE works [13]–[22] leveraged post-processing for scale recovery. The typical post-processing technique of median scaling proposed in [13] scales the relative depth map by the ratio of the median of the ground truth depth values to the median of the estimated depth values. Median scaling has been adopted by many subsequent works [17]–[22], [24], [31]. Nevertheless, ground truth depth maps required by median scaling are not always available at test time in many cases. As the use of camera height in classical monocular visual odometry [33], DNet [16] leveraged a dense geometrical constraints module to determine the scale factor by the estimated camera height from every ground point.

Compared to post-processing techniques, training a scaleaware model is more attractive due to its simplicity of testing. To learn metrically accurate depth estimation, [34] pretrained depth network on stereo data of one dataset and then finetuned the model on monocular videos of another dataset to preserve the absolute scale, and PackNet-sfm [22] leveraged the camera's velocity and [35] utilized 4-beam Lidar data to construct scale-aware loss. All of these works require additional sensor data for training. In order to avoid using other sensor data, [23] proposed a camera height-based loss function to enforce metrically-scaled depth during training. Similar to [23], we also leverage camera height to construct scale constraints. However, our method directly determines the pixels corresponding to the ground plane based on predicted depth instead of relying on an extra pretrained ground plane segmentation model as in [23]. Thus, we simplify the training procedure and improve depth estimation accuracy by mitigating the effects of the limited accuracy of the ground plane segmentation model and the uneven ground surface.

C. Attention Mechanism

Attention mechanism is a dynamic process that diverts attention to important features [36]. The pioneering work of visual attention is the spatial attention network called RAM [37]. After that, more attention mechanisms [38]–[41] were proposed, such as self-attention [42] and channel attention [38]. Nowadays, attention mechanisms have been applied in

many visual tasks [37]-[41], [43], [44]. As for monocular depth estimation, the attention-based network has been applied not only to supervised methods [48], [49] but also to selfsupervised methods [19]-[21]. In [19], an attention module was used to better perceive contextual information. FSRE-Depth [20] utilized an attention module to fuse depth features and semantic features. CADepth-Net [21] applied channel attention for structure perception as well as skip-connection. However, these networks [48], [49], [19]-[21] just added several attention modules to the original network, and the body of the depth network is still the usual localized convolutional network. Therefore, we introduce the visual attention mechanism throughout the processing of the network, so that our model can continuously and adaptively extract important features and remove noise interference at all stages for more accurate depth predictions.

III. METHOD

In this section, we first introduce the basic paradigm for self-supervised monocular depth estimation. Then, we describe the details of our VADepth network. Finally, we introduce a simple method for detecting coplanar points and construct a scale loss based on the planar geometric prior.

A. Self-Supervised Monocular Depth Estimation

During the self-supervised MDE model training, the model is optimized by minimizing the difference between the target image and the image synthesized by the source image and the predicted target depth map. In the training process using monocular videos, the source frame is chosen from the same monocular image sequence as the target frame. Following [13], the source frame I_s is an adjacent frame to the target frame I_t , i.e. $s = t \pm 1$. Since the relative pose between two adjacent frames is unknown, we need to estimate not only the depth but also the ego-motion during training. The depth network takes an RGB image I_t as input and outputs a depth map D_t . Simultaneously, the pose network predicts the relative pose $T_{t\to s}$ between I_s and I_t . Let K denote the known camera intrinsic matrix. Then we can project the target pixel coordinates (u, v) to the source view and obtain the 2D coordinates of the projected depths $D_t(u, v)$ in I_s as follows:

$$(u_s, v_s) = KT_{t \to s} P_t(u, v), \tag{1}$$

and
$$P_t(u, v) = D_t(u, v)K^{-1}(u, v, 1)^T$$
. (2)

Here, (u, v) and (u_s, v_s) are the projected coordinates in I_t and I_s of the same 3D points $P_t(u, v)$, respectively. For simplify of notation, the conversions between inhomogeneous and homogeneous coordinates are omitted in (1). Because the calculated coordinates (u_s, v_s) are not always integer values, we use bilinear interpolation to compute pixel values $I_s(u_s, v_s)$ as in [45]. Thus, we can synthesize the target image:



Fig. 1. Network structure of VADepth. The encoder extracts feature maps X_i^e at different resolutions. The Dual Attention Module (DAM) takes the last level feature map X_4^e as input and emphasizes the important features to produce the initial input of the decoder. The decoder has five successive Visual Attention Blocks (VABs) and outputs multi-scale depth maps for the top four VABs.

where $\langle \cdot \rangle$ denotes bilinear sampling operation. Following [11], we combine the L1 loss and structural similarity index measure (SSIM) [46] to formulate the photometric error:

$$PE(I_a, I_b) = \alpha ||I_a - I_b||_1 + (1 - \alpha) \frac{1 - SSIM(I_a, I_b)}{2}.$$
 (4)

To avoid the false photometric error of occluded pixels in the source frame, we apply per-pixel minimum photometric loss [24], i.e.

$$L_{ph} = \min_{s \in \{t-1,t+1\}} PE(I_t, I_{s \to t}).$$
 (5)

Besides, we adopt the auto-masking method [24] to filter out pixels that do not change the appearance between I_s and I_t . The binary mask μ is computed as

$$\mu = [\min_{s} PE(I_t, I_{s \to t}) < \min_{s} PE(I_t, I_s)], \qquad (6)$$

where $[\cdot]$ is the Inverse bracket. Like [11], we also use the edge-aware smoothness loss:

$$L_{sm} = |\partial_u d_t^*| e^{-|\partial_u I_t|} + |\partial_v d_t^*| e^{-|\partial_v I_t|}, \tag{7}$$

where $d_t^* = d_t/\overline{d_t}$ is the mean-normalized inverse depth. Considering the gradient locality of the bilinear sampler, we predict multi-scale depth maps in the decoder and compute the individual losses at each scale following [24]. Thus, the final baseline loss with a hyperparameter λ_{sm} is defined as

$$L_{baseline} = \frac{1}{S} \sum_{i=0}^{S-1} \mu L_{ph} + \lambda_{sm} L_{sm}, \qquad (8)$$

 $I_{s \to t}(u, v) = I_s \langle (u_s, v_s) \rangle, \tag{3}$

where S refers to the number of multi-scale depth maps.



Fig. 2. Network details of Visual Attention Block (VAB). A Spatial Attention Module (SAM) and a Channel Attention Module (CAM) are used in the VAB.

B. VADepth Network Architecture

As shown in Fig. 1, VADepth adopts the improved U-shape encoder-decoder architecture, which takes a single RGB image as input and outputs multi-scale depth maps. Unlike U-net architecture [26], the encoder and decoder in VADepth are not completely symmetrical, and an additional Dual Attention Module (DAM) head (see details in [40]) is used to connect the encoder and the decoder.

For the encoder, we employ a visual attention network (VAN) [44] to extract multi-scale feature maps $X_i^e, i = 1, 2, 3, 4$. A VAN has four stages, where spatial adaptability and channel adaptability are efficiently implemented by the large kernel attention.

Then the internal features of the encoder are fed into the decoder with skip connections and the lowest resolution features are taken as the input of the DAM. The DAM adaptively selects the discriminative features as the initial input of the decoder by learning spatial and channel interdependencies of features in parallel.

As for the decoder, we design a novel Visual Attention Block (VAB). There are five successive VABs in the decoder. As illustrated in Fig. 2, the *i*-th VAB has two inputs, the internal features X_{i-1}^e from the encoder and the features X_i^d from the previous VAB. X_{i-a}^e is fed into a Spatial Attention Module (SAM), which emphasizes the important regions and suppresses the unimportant or noisy regions as in [39]. In the SAM, the initial spatial attention maps are formed by max pooling and average pooling along the channel axis. Then the concatenated initial attention maps are passed through a convolutional layer with a sigmoid function to obtain the final spatial attention map. The output of the SAM is the element-wise multiplication of the broadcasted final attention map and X_{i-1}^e . At the same time, X_i^d is accordingly passed through a convolutional layer and a nearest upsampling layer with a factor of 2 to recover the resolution. Then a Channel Attention Module (CAM) is used to concatenate and fuse the low-level feature from the SAM and the high-level feature from the upsampling operation. For the CAM, as in [38], the concatenated features are processed through the global average pooling layer, fully-connected (FC) layer with ReLU, and FC layer with sigmoid function to generate the channel attention map. Similar to the SAM, the output of the CAM is obtained by element-wise multiplication of the attention map and the original features. After the process of CAM, a convolutional layer is used to output X_{i-1}^d . Finally, a convolutional layer with sigmoid output σ can be converted to the depth map D_i with $1/(a\sigma + b)$, which is omitted in Fig. 2 for simplicity. Note that the SAM is removed from VAB when the internal feature from the encoder are unavailable.

C. Co-planar Points Detection on Ground Surface

The depth estimation model obtained by minimizing the baseline loss defined in (8) can only directly output the relative depth. In order to train a depth estimation model that can output absolute depth maps, we need to utilize the geometric information that contains absolute scale. In most self-driving scenarios, the ubiquitous camera height can be considered as a known constant value. Therefore, the prior knowledge about camera height is suitable for constructing scale constraints by computing the camera height errors derived from the estimated depth.

To estimate the camera height from the predicted depth map, we need to detect the co-planar points on the ground surface. We observe that the middle-lower region of the image is part of the ground surface in almost all self-driving cases. Based on such geometric prior knowledge, we suppose that a small stationary rectangular region in image I_t is the ground plane. We use two hyperparameters $\alpha_u = 0.075$ and $\alpha_v = 0.875$ to predefine an $H \times W$ binary mask M_{rect} (as shown in Fig. 5) to distinguish whether a pixel is in the rectangular region.

$$M_{rect}(u,v) = \begin{cases} 1 & \text{if } |0.5 - \frac{u}{W}| < \alpha_u \text{ and } \frac{v}{H} > \alpha_v, \\ 0 & \text{others.} \end{cases}$$
(9)

Let $P_{rect} \in R^{N_{rect} \times 3}$ be a matrix for a set of 3D points $\{P_t(u,v)|M_p(u,v)=1\}$ in the predefined rectangular region, where $N_{rect} = 2\alpha_u W(1-\alpha_v)H$ and $P_t(u,v)$ is calculated from the estimated depth according to (2). Based on these 3D points, we can fit a plane Pl_{rect} that does not pass through the origin by solving the equation: $P_{rect}n = 1$. We take the Moore-Penrose inverse of P_{rect} to obtain the least-squares solution of the contradictory equations:

$$\boldsymbol{n} = P_{rect}^+ \boldsymbol{1}. \tag{10}$$

After plane fitting, we need to detect all 3D points co-planar with the plane Pl_{rect} in the target view using a binary mask:

$$M_p(u,v) = \begin{cases} 1 & \text{if } |P_t(u,v)\boldsymbol{n} - 1| < \delta, \\ 0 & \text{others.} \end{cases}$$
(11)

In M_p , an element whose value is equal to one indicates that the corresponding point is on the plane Pl_{rect} . The threshold

 TABLE I

 QUANTITATIVE RESULTS ON THE EIGEN SPLIT OF KITTI DATASET

Mathad	Train	Scale Factor	The lower the better				The higher the better			#D
Method			Abs Rel	Sq Rel	RMSE	RMSE log	δ_1	δ_2	δ_3	#Param.
SfMlearner [13]†	М	GT	0.183	1.595	6.709	0.270	0.734	0.902	0.959	31.6M
Monodepth2 [24]	М	GT	0.115	0.903	4.863	0.193	0.877	0.959	0.981	14.3M
DNet [16]	М	GT	0.113	0.864	4.812	0.191	0.877	0.960	0.981	14.3M
SGDepth [18]	M+Sem	GT	0.113	0.835	4.693	0.191	0.879	0.961	0.981	14.3M
Packnet-sfm [22]	М	GT	0.111	0.785	4.601	0.189	0.878	0.960	0.982	128M
HR-Depth [32]	М	GT	0.109	0.792	4.632	0.185	0.884	0.962	0.983	14.6M
Johnston et al. [19]	М	GT	0.106	0.861	4.699	0.185	0.889	0.962	0.982	-
Wang et al. [17]	М	GT	0.109	0.779	4.641	0.186	0.883	0.962	0.982	-
FSRE-Depth [20]	M+Sem	GT	0.105	0.722	4.547	0.182	0.886	0.964	0.984	25.2M
CADepth [21]	М	GT	0.105	0.769	4.535	0.181	0.892	0.964	0.983	58.3M
VADepth (Ours) w/o Las	М	GT	0.104	0.774	4.552	0.181	0.892	0.965	0.983	18.8M
DNet [16]	М	Cam. Height	0.118	0.925	4.918	0.199	0.862	0.953	0.979	14.3M
Packnet-sfm [22]	M+v	None	0.111	0.829	4.788	0.199	0.864	0.954	0.980	128M
Wagstaff et al. [23]	М	None	0.123	0.996	5.253	0.213	0.840	0.947	0.978	-
Monodepth2 [24] with L_{as}	М	None	0.112	0.875	4.905	0.199	0.863	0.955	0.980	14.3M
VADepth (Ours) with L_{as}	М	None	0.109	0.785	4.624	0.190	0.875	0.960	0.982	18.8M

All models are tested with the resolution of 192×640 unless otherwise specified, for the maximum depth of 80m. † means the newer results from github with the resolution of 128×416 . The best scores for each category are in **bold**. In the "Train" column, we list the supervision for each method with M — Self-supervised monocular supervision, Sem — Semantic supervision, v — velocity supervision. In the "Scale Factor" column, "GT" refers to determining a per-image scale factor using the ground truth depth for median scaling based scale recovery, "Cam. Height" indicates leveraging the known camera height w.r.t the ground surface to recover scaled depth, while "None" means test without any post-processing.

 δ is set to 0.01, which is used to judge whether the points are coplanar or not.

D. Absolute Scale Loss

After co-planar points detection on the ground surface, we construct an absolute scale loss term using the camera height. To form a scale constraint, we need to compute the camera height from the estimated depth results:

$$H_{cam}(u,v) = P_t(u,v)\boldsymbol{n}_e,\tag{12}$$

where $n_e = \frac{n}{||n||_1}$ is the normalized vector of n. Using the estimated heights H_{cam} and the ground truth height h_{gt} , we define the absolute scale loss as follows:

$$L_{as} = \frac{1}{||M_p||_1} ||M_p \circ (H_{cam} - h_{gt})||_1,$$
(13)

where \circ is Hadamard product operator. Combining L_{as} with $L_{baseline}$, we have the final objective function:

$$L = \frac{1}{S} \sum_{i=0}^{S-1} \mu L_{ph} + \lambda_{sm} L_{sm} + \lambda_{as} L_{as}.$$
 (14)

In our setting, λ_{sm} , λ_{as} and S are set to 0.001, 0.01 and 4, respectively.

IV. EXPERIMENTAL RESULTS

A. Implementation details

We use a single Nvidia GeForce RTX 2080 Ti to implement our models in Pytorch. In our setting, we use the VADepth network as our depth network with an input/output resolution of 192×640 . For pose estimation, we use a modified ResNet18 [25] to predict a single 6-DoF ego-motion of two adjacent frames following [24]. The pose network adopts the same input resolution as the depth network. Both the encoder of the depth network and the encoder of the pose network are initialized with weights pretrained on ImageNet [47]. The depth network and pose network are jointly trained with Adam optimizer for 20 epochs. The learning rate is set to 5.0×10^{-5} for the first 15 epochs and then drops by half for the remaining epochs. The source code and models are available at https://github.com/xjixzz/vadepth-net.

B. Evaluation on KITTI

KITTI dataset [28] is the commonly used dataset for autonomous driving. We use the Eigen split [6] of KITTI Stereo dataset to evaluate our method with the metrics proposed in [6]. Following [13], we also remove the static frames before training, which results in 39810 monocular triplets for training, 4424 images for validation and 697 images for test. The resolution of these images is approximately 375×1242 . All images are required to resize to 192×640 for training and evaluation. At test time, the depth network only outputs the depth map with the resolution of 192×640 , which is then resized to the full resolution of the original RGB image for evaluation.

To evaluate the performance of our depth network, we first train VADepth network only with the baseline loss function. We report the results of applying median scaling at the top of Table I, where we can see that our model achieves the best scores in the metrics of Abs Rel, RMSE log, δ_1 , and δ_2 , and ranks second or third in the remaining metrics. Among these methods, CADepth [21] achieves close performance to VADepth but requires much more parameters, due to the underutilization of visual attention, especially ignoring the channel and spatial interdependencies of low-level features.

Besides, we also combine our absolute scale loss to train our model so that the depth network can directly output metrically accurate depth without post-processing that rely on the ground



Fig. 3. Quantitative results on the Eigen split of KITTI dataset. Compared to other methods [20], [21], [24], our model not only outputs smoother depth in patterned planar regions, but also preserves better details for small objects. White boxes highlight the difference. Best viewed in color and zoom in.

					SeasonDepth test						
Method	Train	Resolution	KITTI Eigen test		Average		Variance (10^{-2})		Relative Range		
			$AbsRel\downarrow$	$a_1 \uparrow$	$AbsRel\downarrow$	$a_1 \uparrow$	$AbsRel\downarrow$	$a_1\downarrow$	$AbsRel\downarrow$	$1 - a_1 \downarrow$	
Eigen et al. [6]‡	D	-	0.203	0.702	1.093	0.340	0.346	0.0170	0.206	0.0746	
VNL(ResNext101) [8]‡	D	-	0.072	0.938	0.306	0.527	0.126	0.166	0.400	0.290	
BTS(DenseNet161) [9]‡	D	-	0.060	0.955	0.676	0.209	0.545	0.0650	0.405	0.129	
SfMLearner [13]‡	M	128×416	0.181	0.733	0.360	0.495	0.0801	0.0628	0.269	0.182	
Monodepth2 [24]	M	192×640	0.115	0.877	0.266	0.611	0.0410	0.0457	0.266	0.202	
FSRE-Depth [20]	M	192×640	0.105	0.886	0.256	0.624	0.0288	0.0283	0.227	0.158	
CADepth-Net [21]	M	192×640	0.105	0.892	0.257	0.625	0.0447	0.0725	0.265	0.278	
VADepth (Ours)	M	192×640	0.104	0.892	0.230	0.667	0.0158	0.0215	0.205	0.179	

 TABLE II

 GENERALIZATION PERFORMANCE ON SEASONDEPTH

In the Train Column, D refers to ground truth depth supervision and M refers to self-supervised monocular supervision. \downarrow means the lower the better while \uparrow means the higher the better. Methods marked with \ddagger indicate that the corresponding test results are taken from [29].

truth depth maps. As listed at the bottom of Table I, we compare our methods with existing self-supervised absolute depth estimation methods that do not rely on the ground truth depth. Among these methods, our method achieves the best results on all metrics without requiring extra pretrained road segmentation network [23] or velocity supervision [22] at training time, or post-processing [16] at test time.

The qualitative results are presented in Fig. 3. Compared with [20], [21], [24], our model predicts smoother and more accurate depth for patterned surfaces (e.g. the traffic sign in row 1, the banner in row 3, and the bus in row 5), which validates the effectiveness of our model to suppress unnecessary information. Meanwhile, our model also works well for thin things, such as the thin poles in row 2 and the traffic sign in row 4, which reflects the ability to highlight the important details. Taken together, these results illustrate that incorporating the attention mechanism throughout the whole depth network can help to improve the ability of the model to automatically distinguish important contextual information.

C. Generalization performance

SeasonDepth dataset [29] is an MDE dataset that contains the same multi-traverse routes under 12 different environmental conditions. The test set of SeasonDepth contains 17225 RGB images and corresponding ground truth depth maps with the resolution of 768×1024 . We use the model trained on KITTI to evaluate the generalization ability of our method on the test set of SeasonDepth without fine-tuning. Results in Table II show that our model generalizes better than existing supervised methods and self-supervised monocular training methods, which proves that our model is more capable of learning transferable feature representation.

We also provide the quantitative comparison in Fig. 4. The inferior results of CADepth [21] and FSRE [20] suggest that applying channel attention in the decoder is not enough to obtain high-level features that are robust to lighting and weather changes. Despite extreme light conditions or challenging weather conditions, the output of our model still correctly reflects the scene structure. This further verifies that robust low-level feature representations extracted from our encoder and SAM are also important for forming robust high-level features.

D. Ablation Study

To verify the effectiveness of each component in our method, we perform an ablation study on the KITTI dataset.

Ablations of the VADepth network are reported in Table III, where all models are trained with the loss in (14). We use



Fig. 4. Quantitative results on the test set of SeasonDepth dataset. Our method generates depth maps that are more consistent with the scene structure in challenging environments.

 TABLE III

 Ablation study of our network structure

VAN	DAM	VAB		The lower the better					
		SAM	CAM	AbsRel	SqRel	RMSE	RMSE log		
				0.112	0.875	4.905	0.199		
\checkmark				0.111	0.876	4.762	0.195		
\checkmark	\checkmark			0.108	0.815	4.686	0.191		
\checkmark		\checkmark	\checkmark	0.110	0.821	4.681	0.193		
\checkmark	\checkmark	\checkmark		0.106	0.789	4.672	0.190		
\checkmark	\checkmark		\checkmark	0.109	0.798	4.631	0.192		
\checkmark	\checkmark	\checkmark	\checkmark	0.109	0.785	4.624	0.190		

Monodepth2 [24] as the baseline, which adopts a ResNet18based encoder-decoder without any attention mechanisms. The model of row 2 is composed of a VAN [44] based encoder and a convolutional decoder similar to [24], which performs better than the baseline. Adding either DAM or VAB leads to improvements on all metrics. Combining DAM with SAM or CAM can further improve the performance. Incorporating all modules together results in the best overall performance.

We also ablate the absolute scale loss L_{as} using different network architectures. It can be seen from Table IV that adding L_{as} with proper α_u and α_v leads to the similar performance or even better performance compared to the baseline loss when ground truth depth maps are used for median scaling at test time. However, when testing without median scaling, the model trained without absolute scale loss performs poorly. When trained with absolute scale loss, both network architectures achieve good performance without any post-processing. These results indicate that the scale-aware depth network can be obtained with the camera height-based scale constraints. Besides, we also list the results of selecting different values for α_u and α_v , which determine the size of the rectangular area defined in (9), in the lower part of Table IV. Selecting a too large or too small rectangle degrades performance. Thus, we empirically select α_u and α_v as 0.075 and 0.875, respectively.

TABLE IV Ablation study on absolute scale loss

Scale	$\begin{array}{c c} L_{as} \\ \hline \alpha_u & \alpha_v \end{array}$			S. 1							
Factor			AbsRel	SqRel	RMSE	RMSElog	01				
Monodepth2											
GT	w/o Las		0.115	0.903	4.863	0.193	0.877				
01	0.075	0.875	0.112	0.864	4.804	0.190	0.878				
None	w/o Las		0.968	15.094	19.176	3.459	0.000				
None	0.075 0.875		0.112	0.875	4.905	0.199	0.863				
VADepth											
GТ	w/o Las		0.104	0.774	4.552	0.181	0.892				
01	0.075 0.875		0.105	0.757	4.501	0.180	0.891				
	w/o	L_{as}	0.966	15.045	19.147	3.398	0.000				
	0.025	0.950	0.127	0.932	4.920	0.207	0.850				
	0.050	0.900	0.110	0.870	4.719	0.191	0.876				
	0.050	0.875	0.110	0.792	4.654	0.192	0.872				
None	0.075	0.900	0.109	0.877	4.702	0.190	0.879				
None	0.075	0.875	0.109	0.785	4.624	0.190	0.875				
	0.075	0.850	0.109	0.827	4.709	0.192	0.875				
	0.100 0.875		0.109	0.825	4.680	0.190	0.876				
	0.100 0.800		0.109	0.818	4.674	0.192	0.875				
	0.200	0.600	0.126	0.863	4.858	0.209	0.850				
RGB VADepth w/o L _{as} VADepth with L _{as}											
Citra							-				
Contrast.	ie d			0	-						
/		No. of Concession, Name		and the second second							
CADepth-Net			Mon	ith Las							
		2			C. State						
	1-2					pro tri to					
-		No.	-	and the second se			and the second s				

Fig. 5. Quantitative results for ablation study of absolute scale loss on the KITTI dataset. The green box in the RGB image is the rectangular region defined by (9). Using the absolute scale loss, the predicted results at the white solid line are smoother and more accurate.

Moreover, we find that absolute scale loss helps alleviate the impact of lane lines on depth estimation. As shown in Fig. 5, although these networks can not predict smooth depth at the solid white line unless absolute scale loss is used, both VADepth and Monodepth2 [24] trained with the absolute scale loss output more accurate depths at the lane line, which suggests that adopting supervision of camera height can lead to smooth depth estimates of the ground surface, in addition to absolute scaled depth. This further demonstrates that our network can effectively learn the geometric knowledge of autonomous driving scenarios with the proposed loss.

V. CONCLUSION

In this work, we presented a novel deep learning architecture for self-supervised monocular depth estimation and proposed an absolute scale loss to supervise the scale of estimated depth maps. Experimental results show that our method achieves state-of-the-art performance with or without median scaling on the KITTI dataset and generalizes well on the unseen changing environments. In the future, we plan to further improve the generalization capacity and interpretability of the depth estimation network by visualizing the attention map and using the results of semantic segmentation to guide the network to strengthen the responses of different channels to specific semantic information or improve the distinction between important features and noises.

REFERENCES

- D. Park, R. Ambrus, V. Guizilini, J. Li, and A. Gaidon, "Is pseudolidar needed for monocular 3d object detection?" in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, October 2021, pp. 3142–3152.
- [2] M. Schön, M. Buchholz, and K. Dietmayer, "Mgnet: Monocular geometric scene understanding for autonomous driving," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2021, pp. 15 804–15 815.
- [3] M. Mancini, G. Costante, P. Valigi, and T. A. Ciarfuglia, "J-mod 2: Joint monocular obstacle detection and depth estimation," *IEEE Robot. Autom. Lett.*, vol. 3, no. 3, pp. 1490–1497, 2018.
- [4] J. Uhrig, N. Schneider, L. Schneider, U. Franke, T. Brox, and A. Geiger, "Sparsity invariant cnns," in *Proc. Int. Conf. 3D Vis.*, 2017, pp. 11–20.
- [5] R. Hartley and A. Zisserman, *Multiple view geometry in computer vision*. Cambridge Univ. Press, 2003.
- [6] D. Eigen, C. Puhrsch, and R. Fergus, "Depth map prediction from a single image using a multi-scale deep network," *Proc. Adv. Neural Inf. Process. Syst.*, vol. 27, pp. 2366–2374, 2014.
- [7] D. Eigen and R. Fergus, "Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture," in *Proc. IEEE Int. Conf. Comput. Vis.*, December 2015, pp. 2650–2658.
- [8] W. Yin, Y. Liu, C. Shen, and Y. Yan, "Enforcing geometric constraints of virtual normal for depth prediction," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2019, pp. 5684–5693.
- [9] J. H. Lee, M.-K. Han, D. W. Ko, and I. H. Suh, "From big to small: Multi-scale local planar guidance for monocular depth estimation," *arXiv* preprint arXiv:1907.10326, 2019.
- [10] R. Garg, V. K. Bg, G. Carneiro, and I. Reid, "Unsupervised cnn for single view depth estimation: Geometry to the rescue," in *Proc. Eur. Conf. Comput. Vis.*, 2016, pp. 740–756.
- [11] C. Godard, O. Mac Aodha, and G. J. Brostow, "Unsupervised monocular depth estimation with left-right consistency," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 270–279.
- [12] R. Peng, R. Wang, Y. Lai, L. Tang, and Y. Cai, "Excavating the potential capacity of self-supervised monocular depth estimation," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2021, pp. 15560–15569.
- [13] T. Zhou, M. Brown, N. Snavely, and D. G. Lowe, "Unsupervised learning of depth and ego-motion from video," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 1851–1858.
- [14] Z. Yang, P. Wang, W. Xu, L. Zhao, and R. Nevatia, "Unsupervised learning of geometry from videos with edge-aware depth-normal consistency," in *Proc. AAAI Conf. Artif. Intell.*, 2018, pp. 7493–7500.
- [15] Z. Yang, P. Wang, Y. Wang, W. Xu, and R. Nevatia, "Lego: Learning edge with geometry all at once by watching videos," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 225–234.
- [16] F. Xue, G. Zhuo, Z. Huang, W. Fu, Z. Wu, and M. H. Ang, "Toward hierarchical self-supervised monocular absolute depth estimation for autonomous driving applications," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2020, pp. 2330–2337.
- [17] L. Wang, Y. Wang, L. Wang, Y. Zhan, Y. Wang, and H. Lu, "Can scale-consistent monocular depth be learned in a self-supervised scaleinvariant manner?" in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2021, pp. 12727–12736.
- [18] M. Klingner, J.-A. Termöhlen, J. Mikolajczyk, and T. Fingscheidt, "Selfsupervised monocular depth estimation: Solving the dynamic object problem by semantic guidance," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 582–600.
- [19] A. Johnston and G. Carneiro, "Self-supervised monocular trained depth estimation using self-attention and discrete disparity volume," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 4756–4765.
- [20] H. Jung, E. Park, and S. Yoo, "Fine-grained semantics-aware representation enhancement for self-supervised monocular depth estimation," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2021, pp. 12642–12652.
- [21] J. Yan, H. Zhao, P. Bu, and Y. Jin, "Channel-wise attention-based network for self-supervised monocular depth estimation," in *Proc. Int. Conf. 3D Vis.*, 2021, pp. 464–473.
- [22] V. Guizilini, R. Ambrus, S. Pillai, A. Raventos, and A. Gaidon, "3d packing for self-supervised monocular depth estimation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 2485–2494.
- [23] B. Wagstaff and J. Kelly, "Self-supervised scale recovery for monocular depth and egomotion estimation," in *IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2020, pp. 2620–2627.
- [24] C. Godard, O. Mac Aodha, M. Firman, and G. J. Brostow, "Digging into self-supervised monocular depth estimation," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2019, pp. 3828–3838.

- [25] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 770–778.
- [26] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assisted Intervention*. Cham: Springer International Publishing, 2015, pp. 234–241.
- [27] E. Sucar and J.-B. Hayet, "Probabilistic global scale estimation for monoslam based on generic object detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshops*, 2017, pp. 48–56.
- [28] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? the kitti vision benchmark suite," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2012, pp. 3354–3361.
- [29] H. Hu, B. Yang, Z. Qiao, D. Zhao, and H. Wang, "Seasondepth: Cross-season monocular depth prediction dataset and benchmark under multiple environments," arXiv preprint arXiv:2011.04408, 2020.
- [30] J. Xie, R. Girshick, and A. Farhadi, "Deep3d: Fully automatic 2d-to-3d video conversion with deep convolutional neural networks," in *Proc. Eur. Conf. Comput. Vis.*, 2016, pp. 842–857.
- [31] C. Shu, K. Yu, Z. Duan, and K. Yang, "Feature-metric loss for selfsupervised learning of depth and egomotion," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 572–588.
- [32] X. Lyu, L. Liu, M. Wang, X. Kong, L. Liu, Y. Liu, X. Chen, and Y. Yuan, "Hr-depth: High resolution self-supervised monocular depth estimation," in *Proc. AAAI Conf. Artif. Intell.*, vol. 35, no. 3, 2021, pp. 2294–2301.
- [33] X. Wang, H. Zhang, X. Yin, M. Du, and Q. Chen, "Monocular visual odometry scale recovery using geometrical constraint," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2018, pp. 988–995.
- [34] T. Roussel, L. Van Eycken, and T. Tuytelaars, "Monocular depth estimation in new environments with absolute scale," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2019, pp. 1735–1741.
- [35] F. Bartoccioni, É. Zablocki, P. Pérez, M. Cord, and K. Alahari, "Lidartouch: Monocular metric depth estimation with a few-beam lidar," *arXiv* preprint arXiv:2109.03569, 2021.
- [36] M.-H. Guo, T.-X. Xu, J.-J. Liu, Z.-N. Liu, P.-T. Jiang, T.-J. Mu, S.-H. Zhang, R. R. Martin, M.-M. Cheng, and S.-M. Hu, "Attention mechanisms in computer vision: A survey," *Comput. Vis. Media*, Mar 2022.
- [37] V. Mnih, N. Heess, A. Graves et al., "Recurrent models of visual attention," Proc. Adv. Neural Inf. Process. Syst., vol. 27, pp. 1–9, 2014.
- [38] J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, June 2018, pp. 7132–7141.
- [39] S. Woo, J. Park, J.-Y. Lee, and I. S. Kweon, "Cbam: Convolutional block attention module," in *Proc. Eur. Conf. Comput. Vis.*, September 2018, pp. 3–19.
- [40] J. Fu, J. Liu, H. Tian, Y. Li, Y. Bao, Z. Fang, and H. Lu, "Dual attention network for scene segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, June 2019, pp. 3146–3154.
- [41] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, "An image is worth 16x16 words: Transformers for image recognition at scale," in *Proc. Int. Conf. Learn. Represent.*, 2021.
- [42] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," *Proc. Adv. Neural Inf. Process. Syst.*, vol. 30, pp. 5998–6008, 2017.
- [43] R. Ranftl, A. Bochkovskiy, and V. Koltun, "Vision transformers for dense prediction," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2021, pp. 12179–12188.
- [44] M.-H. Guo, C.-Z. Lu, Z.-N. Liu, M.-M. Cheng, and S.-M. Hu, "Visual attention network," arXiv preprint arXiv:2202.09741, 2022.
- [45] M. Jaderberg, K. Simonyan, A. Zisserman et al., "Spatial transformer networks," Proc. Adv. Neural Inf. Process. Syst., vol. 28, 2015.
- [46] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, 2004.
- [47] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2009, pp. 248–255.
- [48] Y. Chen, H. Zhao, Z. Hu, and J. Peng, "Attention-based context aggregation network for monocular depth estimation," *Int. J. Mach. Learn. Cybern.*, vol. 12, no. 6, pp. 1583–1596, 2021.
- [49] M. Lee, S. Hwang, C. Park, and S. Lee, "Edgeconv with attention module for monocular depth estimation," in *Proc. IEEE/CVF Winter Conf. Appl. of Comput. Vis. (WACV)*, January 2022, pp. 2858–2867.