

# DevNet: Self-supervised Monocular Depth Learning via Density Volume Construction

Kaichen Zhou<sup>1</sup>, Lanqing Hong<sup>2†</sup>, Changhao Chen<sup>3</sup>, Hang Xu<sup>2</sup>, Chaoqiang Ye<sup>2</sup>,  
Qingyong Hu<sup>1</sup>, and Zhenguo Li<sup>2</sup>

<sup>1</sup> University of Oxford

<sup>2</sup> Huawei Noah's Ark Lab

<sup>3</sup> National University of Defense Technology

kc.zhou2020@hotmail.com,

{honglanqing, yechaoqiang, xu.hang, li.zhenguo}@huawei.com,

{changhao.chen66, huqingyong15}@outlook.com

**Abstract.** Self-supervised depth learning from monocular images normally relies on the 2D pixel-wise photometric relation between temporally adjacent image frames. However, they neither fully exploit the 3D point-wise geometric correspondences, nor effectively tackle the ambiguities in the photometric warping caused by occlusions or illumination inconsistency. To address these problems, this work proposes Density Volume Construction Network (DevNet), a novel self-supervised monocular depth learning framework, that can consider 3D spatial information, and exploit stronger geometric constraints among adjacent camera frustums. Instead of directly regressing the pixel value from a single image, our DevNet divides the camera frustum into multiple parallel planes and predicts the pointwise occlusion probability density on each plane. The final depth map is generated by integrating the density along corresponding rays. During the training process, novel regularization strategies and loss functions are introduced to mitigate photometric ambiguities and overfitting. Without obviously enlarging model parameters size or running time, DevNet outperforms several representative baselines on both the KITTI-2015 outdoor dataset and NYU-V2 indoor dataset. In particular, the root-mean-square-deviation is reduced by around 4% with DevNet on both KITTI-2015 and NYU-V2 in the task of depth estimation. Code is available at <https://github.com/gitkaichenzhou/DevNet>.

**Keywords:** Depth Estimation, Monocular Camera, Self-supervised Learning, Occlusion Probability Density, Volume Rendering

## 1 Introduction

Vision-based depth estimation (VDE) attracts attentions due to its significance in understanding the geometry of a 3D scene. It is the basis of higher-level 3D

---

<sup>†</sup> Corresponding author.

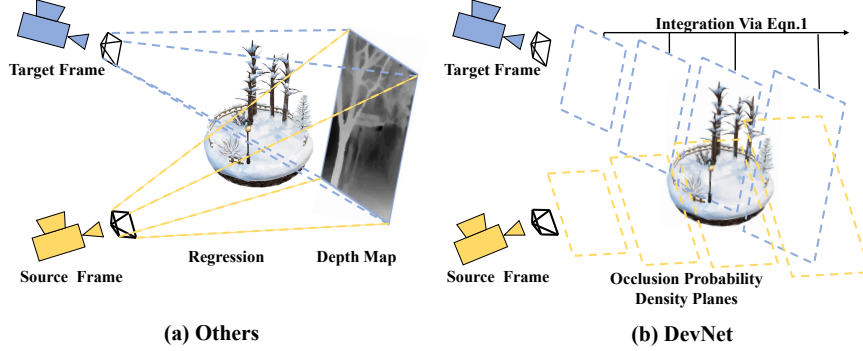


Fig. 1: **A comparison between our proposed DevNet and other depth learning approaches [13, 14, 32, 36, 51].** (a) Most existing depth learning methods predict depth maps directly from input images [10, 11], and the correspondences among adjacent frames are built only with 2D pixel-wise depth maps. (b) Our DevNet renders depth maps by predicting the density information of parallel planes, and the correspondences among adjacent frames are built with 3D frustums. The pyramid denotes the frustum.

tasks, e.g., scene reconstruction [25] and object detection [38], and supports a number of cutting-edge applications, from autonomous driving [16, 47] to augmented reality [21].

Recently learning-based VDE becomes a focus, that can be generally divided into two categories – supervised [1, 4] and self-supervised approaches [46, 50, 52]. Supervised VDE usually requires high-precision ground-truth depth as training labels, whose process would be costly and time-consuming. Self-supervised methods learn depth maps from monocular images [50] or stereo image pairs [52]. Self-supervised monocular VDE [11, 36, 51], exploits the photometric loss based on the warping relation between temporally adjacent image frames [50]. Along with per-pixel depth predictions, it also produces ego-motion estimation. Though self-supervised monocular VDE has seen great progresses, there still exists a large performance gap between self-supervised and supervised methods [31]. This gap comes from the fact that it is unreliable to optimize the framework only based on the photometric loss without understanding the geometry of the whole scene. The assumption that lower photometric loss denotes higher depth accuracy does not hold in many cases, e.g., the occlusions, illumination inconsistency, non-Lambertian surfaces, and textureless regions.

To tackle these problems, we propose DevNet, a novel density volume construction based monocular depth learning framework. *By exploiting volume rendering [20], 3D spatial information along the ray can be integrated into pixel-wise depth estimation, which contributes to building stronger correspondences among adjacent camera frustums.* Instead of directly producing a depth map from a single RGB image, DevNet discretizes the scene into multiple parallel planes, as shown in Fig. 1(b), and predicts the occlusion probability density information of each point on planes. Then, the pixelwise depth value is computed by integrat-

ing the probability density values of intersections between the ray and planes. During the training phase, our DevNet mitigates the limitations of photometric loss by adopting an occlusion regularization and a brightness regularization. Finally, a depth consistency loss based on occupancy probability density volume is introduced to maintain the geometric consistency between temporally adjacent frames. Extensive experiments conducted on both the KITTI-2015 [9] and NYU-V2 [37] datasets demonstrate the effectiveness of our proposed DevNet.

Our main contributions are summarized as below:

- We introduce DevNet, a novel self-supervised depth learning framework that renders depth maps via predicting the occlusion probability density of sampled points in the scene.
- We present a novel depth consistency loss using the predicted density volume, to ensure the geometric consistency between the depth maps of adjacent frames, which can further contribute to performance improvement.
- A occlusion regularization is introduced to detect and reduce the photometric ambiguity caused by camera motion and dynamic objects.

## 2 Related Works

### 2.1 Supervised Monocular Depth Learning

In monocular depth learning, supervised approaches [4, 22, 31, 49] usually take a single image as input and use depth data measured with range sensors (e.g. RGB-D camera or LIDAR) as ground truth labels during training. [35] first proposed a learning-based method to learn depth from single monocular images in an end-to-end manner and the final algorithm could recover even accurate prediction for unstructured scenes. With the development of deep learning, [4] constructs a monocular depth estimator based on the convolutional architecture to infer the corresponding depth with a global estimation layer, followed by a local refinement layer. [34] further improved neural network-based depth estimation by using a predicted categorical depth distribution for each pixel. Although the depth estimation network trained with ground-truth labels can achieve high accuracy in this task, obtaining high-resolution ground-truth labels from different domain/scenes is costly, time-consuming and not always applicable, which limits the potentials of these supervised methods.

### 2.2 Self-supervised Depth Learning

Self-supervised methods realize depth learning based on a series of unlabeled images. [8] firstly made use of a pair of images to estimate corresponding by minimizing a photometric loss between the left input image and the warped right image as the target of optimization. [50] stated that only based on image reconstruction loss would lead to a poor result. They further improved the performance of depth estimation by proposing a new network architecture and a novel left-right consistency loss. However, due to the convenience of self-supervised

monocular depth learning, it attracts more attention. [50] designed an end-to-end self-supervised monocular depth estimation network for depth estimation and motion estimation of monocular unlabeled video. By using the estimated pose, the loss is constructed based on warping nearby views to the target view. [44] further improved the accuracy of the self-supervised depth estimation task by jointly optimizing depth estimation, optic flow, and camera ego-motion. Predictions from three modules are used to construct the target view and form the loss function for both static background and dynamic components. [11] proposes lots of useful strategies such as a robust reprojection loss, multi-scale sampling, and an automasking loss to further improve the accuracy of monocular depth estimation. To improve the performance at non-contiguous regions and motion boundaries. [2] found that image reconstruction loss ignores the influence of dynamic components and proposes to exploit the geometric consistency among different frames for realizing scale-consistency prediction results. [19] further introduces the self-attention mechanism and takes the discrete disparity prediction into the framework for monocular depth estimation, which could learn more general contextual information and generate a more robust prediction.

### 2.3 Neural Rendering

Rendering represents the process of generating 2D images from a 2D/3D model. Classic rendering contains two common approaches: Rasterization and Raytracing. Recent neural rendering generally denotes the combination of the recent proposed generative model with traditional computer graphics rendering techniques, which generates images of 3D scenarios from different camera views based on learned implicit representation. Based on this conception, NeRF [30] is proposed to encode both the appearance and the geometry information of a 3D scenario with a Multilayer Perceptron (MLP), which realizes the view synthesis of real-life scenes. After that, numerous studies further adapt NERF to various scenarios, including images in wild [27], dynamic scenes with monocular or stereo videos [3, 24, 33], post estimation [43], and depth estimation [41]. Despite diverse applications of neural rendering, we are the first to integrate this technique in the monocular depth estimation process. Considering the characteristics of this task and neural rendering, we further propose a novel occlusion regularization and depth consistency loss function to improve its performance in depth learning.

## 3 Density Volume Rendering Based Depth Learning

Most previous self-supervised monocular VDEs [13, 14, 32, 36, 51] adopt the encoder-decoder structure to predict depth map directly from a single RGB image. However, it is not always reliable to optimize these models only by minimizing the photometric loss constructed by warping adjacent frames as in Fig. 1(a) without considering the 3D spatial constraints. DevNet aims to build stronger geometric constraints among adjacent frames by learning the density volume of corresponding camera frustums. The framework of DevNet consists of



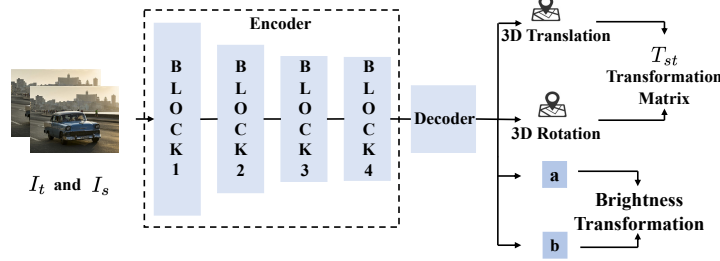


Fig. 3: **The architecture of the Pose Module.** This module would take concatenated two frames as input and output the pose transformation and the brightness transformation parameters.

### 3.2 Framework

DevNet framework contains a depth module for the occlusion probability density volume construction in Fig. 2, and a pose module predicting the motion transformation between the target frame and source frame in Fig. 3. Specifically, DevNet discretizes the scene into multiple parallel planes as in Fig. 1(b) and predicts the occlusion probability density of each point on planes. The pixelwise depth value is calculated through rendering over the ray passing through both the camera origin and the pixel location.

**Density Volume Prediction:** Given a RGB image  $I_i \in \mathbb{R}^{C \times H \times W}$ , instead of directly producing the depth  $Z_i$  as in the previous research (Fig. 1(a)), DevNet divides the camera frustum into  $K$  multiple parallel planes  $\{\Pi_i^k\}_{k=1}^K$  located at the depth  $\{d_i^k\}_{k=1}^K$  (Fig. 1(b)). Our depth module produces the occlusion probability density of planes  $V(\sigma_i) = \{\sigma_i^k\}_{k=1}^K, \sigma_i^k \in \mathbb{R}^{1 \times H \times W}$  based on  $I_i$ . Finally, the  $Z_i$  map is predicted through volume rendering Eqn. (2).

**Rendering-Based Depth Estimation Module:** The framework of depth module is shown in Fig. 2, including an image encoder and a depth decoder. The image encoder is based on the ResNet structure [15] which takes  $I_i$  as input and generates hierarchical feature maps  $\{f_i^j\}_{j=1}^N, f_i^j \in \mathbb{R}^{C^j \times \frac{H}{2^j} \times \frac{W}{2^j}}$ , where we adopt  $N = 4$  in our experiment. The the embedding strategy [30] is introduced to process depth information of each plane, written as

$$e(d_i^k) = (\sin(2^0 \pi d_i^k), \cos(2^0 \pi d_i^k), \dots, \sin(2^{(\frac{E}{2}-1)} \pi d_i^k), \cos(2^{(\frac{E}{2}-1)} \pi d_i^k)), \quad (3)$$

where  $E$  is dimension of embedded depth. After that, feature maps are concatenated with the corresponding embedded depth to obtain multi-scale feature maps  $\{\text{cat}(e(d_i)^j, f_i^j)\}_{j=1}^N$ . These feature maps are skip-connected with the decoder to generate the probability density volume  $V(\sigma_i)$ . During the training phase, DevNet outputs multi-scale density volumes  $V(\sigma_i) = \{V^j(\sigma_i)\}_{j=1}^N$  all of which are used in constructing the training loss. During the inference stage, DevNet only uses the biggest layer density volume.

**Pose Estimation Module:** This module also adopts encoder-decoder structure as in Fig. 3. The pose encoder adopts the ResNet structure and takes concatenated target frame  $I_t$  and source frame  $I_s$  as input. The pose decoder takes the

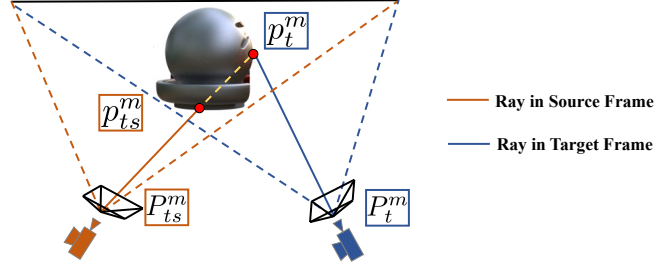


Fig. 4: **The illustration of the occlusion regularization.** This figure shows the intuition behind proposed occlusion regularization where camera is moving from left to right.

feature map from the last layer of the encoder as input. It outputs 3-dimensional translation and 3-dimensional rotation vectors, which are used to construct the transformation matrix  $T_{ts}$ , and the 2-dimensional brightness transformation parameters presented in Sec. 3.3.

### 3.3 Regularizations

The self-supervised depth learning algorithms [10, 11, 36, 42] generally rely on the warping relation between the pixel  $P_s^m(x_s^m, y_s^m, 1)$  in the source frame and its counterpart  $P_{st}^m(x_t^m, y_t^m, 1)$  in the target frame, which can be written as

$$P_{st}^m = Z_t(P_{st}^m)^{-1} K_t T_{st} Z_s(P_s^m) K_s^{-1} P_s^m, \quad (4)$$

where  $K_t, K_s$  are the intrinsic matrices for the target frame and the source frame;  $T_{st}$  is the transformation matrix from the source frame to the target frame;  $Z_s(P_s^m)$  is the depth of pixel  $P_s^m$  in source camera frustum and  $Z_t(P_{st}^m)$  is the depth of pixel  $P_{st}^m$  with the respect to the target camera frustum. However, this relation is destroyed in the following cases, such as moving objects, the camera motion, the illumination inconsistency, non-Lambertian surfaces, etc. To mitigate these influences of aforementioned occasions, we propose the following regularization strategies.

**Regularization for Occlusion:** Due to dynamic objects and camera motion, occlusions are inevitable in application scenarios. These occlusions would cause ambiguities in the photometric warping between source and target frames. To minimize its influence, we design a novel occlusion regularization strategy to mitigate this problem as illustrated in Fig. 4. During the training phase, after predicting the depth map  $Z_t$  of the target frame, the intersection  $p_t^m(x, y, z)$  between the rays of pixel  $P_t^m(X, Y, 1)$  and the surfaces in the target camera frustum can be formulated as

$$p_t^m = Z_t(P_t^m) K_t^{-1} P_t^m, \quad (5)$$

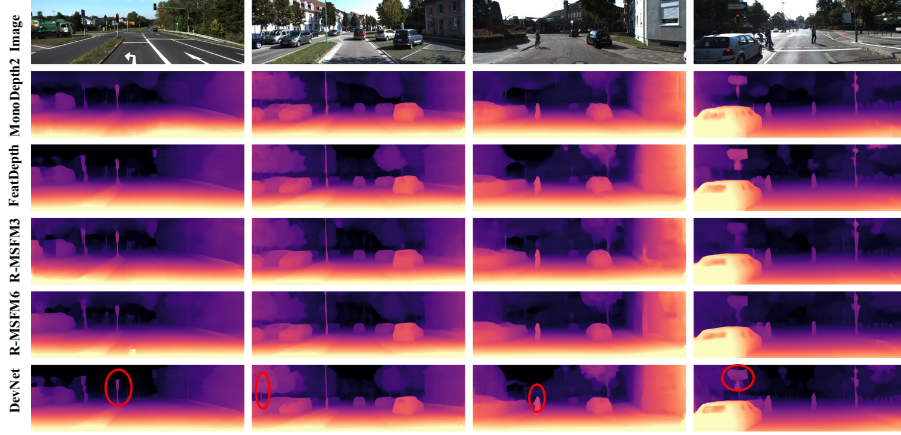


Fig. 5: **The qualitative results of depth estimation on the KITTI-2015.** The results generated by DevNet are shown in the last row, which are sharper than results estimated by listed methods. This comparison is consistent with quantitative results in Tab. 1.

whose depth with respect to the source frame is denoted as  $Z_s(p_t^m)$ . Then the corresponding pixel  $P_{ts}^m$  in the source pixel frame is:

$$P_{ts}^m = \frac{1}{Z_s(p_t^m)} K_t T_{ts} p_t^m, \quad (6)$$

whose depth is denoted as  $Z_s(P_{ts}^m)$  and corresponding points  $p_{ts}^m$  in the source frame could also be found. Due to the movement of object or the motion of camera, point  $p_{ts}^m$  would be different from point  $p_t^m$ . This difference could be used to compute the occlusion mask, which could be written as:

$$\mathcal{M}_o = (|Z_s(p_t^m) - Z_s(P_{ts}^m)| < \gamma), \quad (7)$$

where  $\gamma$  is the threshold in terms of the depth difference. Details could be found in our provided code.

**Regularization for Brightness:** With the help of the Eqn. (4), we project the source frame  $I_s$  to the target frame and get the reconstructed target image  $I_{st}$ , which are used in photometric loss function construction introduced in Sec. 3.4. However, despite the proposed occlusion regularization, the adjacent frames of outdoor scenarios often suffer from illumination inconsistency. Directly forcing reconstructed pixel values to equal to original pixel values will inevitably deteriorate the prediction performance. Hence, we apply the following regularization to transform the reconstructed image  $I_{st}$ :

$$I_{st}^{ab} = a_{st} I_{st} + b_{st}, \quad (8)$$

where  $a$  and  $b$  are the brightness transformation parameters as in [6, 7, 18, 42].



### 3.4 Training Loss

Our training loss function consists of four parts, including a photometric loss  $\mathcal{L}_p$ , a smoothness loss  $\mathcal{L}_s$ , a depth synthesis loss  $\mathcal{L}_d$  and a regularization loss  $\mathcal{L}_r$ :

$$\mathcal{L} = \mathcal{L}_p + \alpha\mathcal{L}_s + \beta\mathcal{L}_d + \eta\mathcal{L}_r, \quad (9)$$

where  $\alpha$ ,  $\beta$  and  $\eta$  are corresponding weights to balance these losses.

**The photometric loss function** consists of L1 and SSIM [10, 11]:

$$\begin{aligned} \mathcal{L}_p = & 0.15 * \sum_p \mathcal{M}_o \odot \mathcal{M}_i \odot |I_t - I_{st}^{ab}| \\ & + 0.85 * \frac{1 - \text{SSIM}(\mathcal{M}_o \odot \mathcal{M}_i \odot I_t, \mathcal{M}_o \odot \mathcal{M}_i \odot I_{st}^{ab})}{2}, \end{aligned} \quad (10)$$

where SSIM is used to quantify the quality degradation in the image processing [40] and  $\mathcal{M}_i$  is the identity mask provided in MonoDepth2 [11].

**The smoothness loss function** is applied to the target depth map  $Z_t$  as:

$$\mathcal{L}_s = |\partial_x z_t^*| e^{-|\partial_x I_t|} + |\partial_y z_t^*| e^{-|\partial_y I_t|}, \quad (11)$$

where  $z_t^* = z_t / \bar{z}_t$  and  $\bar{z}_t$  is the average inverse depth of target frame [39]. This loss function is used to perverse the edge in the depth reconstruction.

**The depth consistency loss.** Based on this Eqn. (4), our DevNet reconstructs the depth map of target frame  $Z_{st}$  based on the density volume of source frame  $V^j(\sigma_s)$ , as the density of each point  $\sigma_t(p_t)$  in target frame can be found by querying the density information of its counterpart in source frame  $\sigma_s(T_{st}p_t)$ . The depth consistency loss is formulated as:

$$\mathcal{L}_d = \frac{1}{M'} \sum_m \left( \frac{|\mathcal{M}_o \odot Z_{st}^m - \mathcal{M}_o \odot Z_t^m|}{\mathcal{M}_o \odot Z_{st}^m + \mathcal{M}_o \odot Z_t^m} \right), \quad (12)$$

where  $M'$  is the number of estimated valid depth values.

**The regularization loss  $\mathcal{L}_r$**  regularizes brightness transformation parameters:

$$\mathcal{L}_r = (a - 1)^2 + b^2, \quad (13)$$

which is proposed under the assumption that there is no significant illumination inconsistency between adjacent frames.

### 3.5 Discussions

**The Relation with NeRF [30] Similarity:** Both of us reconstruct the color image or the depth map by discretizing the scenario into thousands of points. The color or depth value of each pixel is calculated by integrating along with its rays. **Difference:** (1) Instead of running the neural network *thousands of times* to obtain point-wise information for the whole scene, DevNet only passes through the network *once* to acquire the density information of the whole scene.

Methods	Type	Lower is Better				Larger is Better		
		Abs Rel	Sq Rel	RMSE	RMSE log	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$
SIGNet [29]	M	0.133	0.905	5.181	0.208	0.825	0.947	0.981
LearnK [13]	M	0.128	0.959	5.230	0.212	0.845	0.947	0.976
DualNet [49]	M	0.121	0.837	4.945	0.197	0.853	0.955	0.982
MonoDepth2(R50) [11]	M	0.110	0.831	4.642	0.187	0.883	0.962	0.982
Johnston [19]	M	0.106	0.861	4.699	0.185	0.889	0.962	0.982
FeatDepth(R50) [36]	M	0.104	0.729	4.481	0.179	<b>0.893</b>	0.965	0.984
PackNet-SfM [14]	M	0.107	0.802	4.538	0.186	0.889	0.962	0.981
R-MSFM6 [51]	M	0.108	0.748	4.470	0.185	0.889	0.963	0.982
w/o Pretrain(R50)(P24)	M	0.114	0.836	4.888	0.191	0.877	0.958	0.981
DevNet(R50)(P16)	M	0.103	0.713	4.459	0.177	0.890	0.965	0.982
DevNet(R50)(P24)	M	<b>0.100</b>	<b>0.699</b>	<b>4.412</b>	<b>0.174</b>	<b>0.893</b>	<b>0.966</b>	<b>0.985</b>
DFR [46]	MS	0.135	1.132	5.585	0.229	0.820	0.933	0.971
EPC++ [26]	MS	0.128	0.935	5.011	0.209	0.831	0.945	0.979
MonoDepth2(R50) [11]	MS	0.106	0.818	4.75	0.196	0.874	0.957	0.979
FeatDepth(R50) [36]	MS	0.099	0.697	4.427	0.184	0.889	0.963	0.982
R-MSFM6 [51]	MS	0.108	0.753	4.469	0.185	0.888	0.963	0.982
w/o Pretrain(R50)(P24)	MS	0.107	0.751	4.461	0.189	0.883	0.963	0.981
DevNet(R50)(P24)	MS	<b>0.095</b>	<b>0.671</b>	<b>4.365</b>	<b>0.174</b>	<b>0.895</b>	<b>0.970</b>	<b>0.988</b>
NeWCRFs [45]	MSup	0.052	0.155	2.129	0.079	0.974	0.997	0.999

Table 1: **The quantitative results of depth estimation on the KITTI-2015.** Comparison between DevNet and other depth estimation methods on KITTI-2015 with the Eigen split. Best results are in bold. M: trained with monocular sequence. MS: trained with both monocular sequence and stereo images. MSup: supervised trained with monocular sequence.

(2) DevNet concentrates on self-supervised depth estimation and this network can be used in multiple scenarios, while NeRF focuses on learning an implicit representation in a certain scenario.

**The Relation with FLA [12] Similarity:** Both of us adapt the auto-encoder structure and output multi-channel geometry information used for depth map construction. **Difference:** (1) FLA outputs multi-channel disparity logit volume which constructs depth through *SoftMax* calculation. DevNet outputs multi-channel density volume which generates depth through *integration*. (2) FLA pre-defines the position of each channel and does not consider the position information. At each time, DevNet randomly takes samples from a uniform distribution, and the position is used to construct embedded depth volume  $V_i^j$ . (3) Unlike FLA, DevNet does not need the two-stage training strategy and realizes end-to-end training. (4) DevNet focuses on self-supervised *monocular* depth learning, while FLA focuses on self-supervised *stereo* depth learning.

## 4 Experiments

Our experiments are mainly conducted on the KITTI-2015 [9] and NYU-V2 [37]. For KITTI-2015, the Eigen split [4] is used to train and test our model. For a fair comparison, following the same data selection strategy in previous research [50, 51], we remove all static images from the training set, after which 39,810 training images and 4,424 test images are generated. We adapt the camera matrix setting in [11], which uses the same intrinsic matrix for all images. The

length focal is set as the average value. For NYU-V2, we adapt the official train and test split used in [48], where the train set contains 20000 images and the test set consists of 654 labeled images. In this section, the performance of DevNet is firstly tested in terms of depth prediction. Depth prediction results are compared with following widely used evaluation metrics [5, 11]:

$$\begin{aligned}
- \text{Abs Rel} &= \frac{1}{|M|} \sum_{m=1}^M \frac{|\hat{z}^m - z^m|}{\hat{z}^m}; \\
- \text{Sq Rel} &= \frac{1}{|M|} \sum_{m=1}^M \frac{\|\hat{z}^m - z^m\|_2}{\hat{z}^m}; \\
- \text{RMSE} &= \sqrt{\frac{1}{|M|} \sum_{m=1}^M \|\hat{z}^m - z^m\|_2}; \\
- \text{RMSE log} &= \sqrt{\frac{1}{|M|} \sum_{m=1}^M \|\log(\hat{z}^m) - \log(z^m)\|_2}; \\
- \frac{1}{M} |\{\delta = \max(\frac{\hat{z}^m}{z^m}, \frac{z^m}{\hat{z}^m}) < 1.25^t, m = 1, \dots, M\}|;
\end{aligned}$$

where  $\hat{z}^m$  is the ground truth;  $z^m$  is predicted depth value and  $t = 1, 2, 3$ . Then, to evaluate each component in DevNet, ablation studies are also conducted by varying the sample number, removing mask and loss functions, and changing the backbone. Besides the depth estimation, the DevNet is capable of achieving visual odometry estimation, which is compared to baselines. Finally, the generalization of DevNet trained on KITTI-2015 is directly tested on the NYU-V2 [37].

Model	Type	Abs Rel	Sq Rel	RMSE	RMSE log	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$
MonoDepth2(R18) [11]	M	0.160	0.152	0.601	0.186	0.767	0.949	0.988
Monoindoor(R18) [17]	M	0.134	-	0.526	-	0.823	0.958	0.989
<b>DevNet(R18)(P24)</b>	M	<b>0.129</b>	<b>0.098</b>	<b>0.492</b>	<b>0.153</b>	<b>0.843</b>	<b>0.965</b>	<b>0.992</b>

Table 2: **The quantitative results of depth estimation on the NYU-V2.** Comparison between DevNet and other depth estimation methods on NYU-V2 with the official split. Best results are in bold.

#### 4.1 Depth Estimation

**KITTI-2015:** In this experiment, the depth encoder adapts ResNet50(R50), while the pose encoder adapts ResNet18(R18) as base model. During the training phase, our DevNet takes the image with the resolution of  $320 \times 1024$  as input. When testing our DevNet, we set  $80m$  as the max depth value and  $0.1m$  as the minimum depth value [11, 18]. Same as in [50], the median value of the ground-truth image is used to calculate the scale factor. The comparison between the performance of DevNet and that of state-of-art algorithms is demonstrated in Tab. 1. It is clear to see that DevNet generally performs better no matter using monocular images, or a combination of monocular and stereo images. Compared to [19] that also learns a disparity volume used to generate depth map through SoftMax, DevNet outperforms them largely by 5.7% in terms of Abs Rel. Feat-Depth [36] proposes to pre-train a feature encoder for low-texture region learning. DevNet outperforms them by 3.8% in terms of Abs Rel without pretraining the feature auto-encoder. A qualitative comparison of depth predictions is presented in Fig. 5, showing that DevNet detects the details of moving objects

and thin objects, highlighted by red circles. One thing worth mentioning, even though DevNet could achieve good performance with unsupervised setting, supervised algorithms, e.g., NeWCRRFs [45], still largely outperform unsupervised algorithms. *Size&Speed*: One thing worth mentioning is that DevNet slightly increases the parameter size and the running time only, as DevNet only modifies skip-connects and the last layer channel number of MonoDepth2 (R50) backbone. While DevNet significantly improves the performance.

**NYU-V2:** For this indoor dataset, during the training phase, DevNet is fed with a resolution of  $288 \times 384$ . We set  $10.0m$  as the max depth value and  $0.1m$  as the minimum depth value. Considering the lower resolution of NYU-V2, we adapt the ResNet18 backbone for both depth and pose modules. The result in Tab. 2 shows that DevNet also outperforms its counterparts in the indoor dataset. Specifically, DevNet outperforms Monoindoor by around 4% in terms of Abs Rel.

Algorithm	Seq9		Seq10	
	Translation	Rotation	Translation	Rotation
DFR [46]	11.93	3.91	12.45	3.46
MonoDepth2 [11]	10.85	2.86	11.60	5.72
NeuralBundler [23]	8.10	2.81	12.90	<b>3.71</b>
SC-SfMlearner [50]	8.24	2.19	10.70	4.58
FeatDepth [36]	8.75	2.11	10.67	4.91
<b>DevNet(R50)(P24)</b>	<b>8.09</b>	<b>2.05</b>	<b>10.05</b>	4.32

Table 3: **The quantitative results of odometry estimation on the KITTI odometry dataset.** Comparison between DevNet and others in terms of Average translational root mean square error drift and average rotational root mean square error drift.

## 4.2 Odometry Estimation

To test the performance of odometry estimation, we adapt the split strategy used in [46, 50]. This strategy uses the Sequences 00-08 of the KITTI odometry dataset as the training set and the Sequences 09-10 as the test set, as shown in Tab. 3. The averaged translational and rotational root mean square error (RMSE) are reported. Our DevNet outperforms other representative learning-based methods, e.g. MonoDepth2 and FeatDepth. This is because our DevNet focuses on learning the geometric correspondences between adjacent frames and its accurate depth estimation benefits odometry estimation when jointly optimizing both the odometry module and depth estimation module.

## 4.3 Ablation Study

In this section, we analyze the performance of DevNet by studying the influence of different components and training strategies on the depth prediction result.

**Training Strategies:** To evaluate the contribution of each training strategy towards our DevNet framework, we conducted ablation studies under several settings, including the DevNet without embedding strategy applied on position

Regularization	Type	Abs Rel	Sq Rel	RMSE	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$
Baseline(R50)(P24)	M	0.118	0.881	4.897	0.872	0.950	0.971
w/o Embedding(R50)(P24)	M	0.102	0.711	4.439	0.889	0.963	0.983
w/o Color Reg(R50)(P24)	M	0.103	0.719	4.462	0.889	0.962	0.979
w/o Occlusion Mask(R50)(P24)	M	0.106	0.791	4.547	0.882	0.959	0.978
w/o Depth Loss(R50)(P24)	M	0.107	0.784	4.487	0.880	0.959	0.967
DevNet(R18)(P24)	M	0.108	0.801	4.712	0.874	0.951	0.973
DevNet(R50)(P08)	M	0.108	0.746	4.471	0.884	0.949	0.969
DevNet(R50)(P16)	M	0.103	0.713	4.459	0.890	0.965	0.982
DevNet(R50)(P24)	M	0.100	0.699	4.412	<b>0.893</b>	0.966	0.985
DevNet(R50)(P32)	M	<b>0.100</b>	<b>0.697</b>	<b>4.408</b>	0.891	<b>0.968</b>	<b>0.988</b>

Table 4: **The ablation study on depth estimation above the KITTI-2015.** The ablation studies concern the influences of different components, training strategies, and different sample number of planes.

information, without color regularization, without occlusion mask, and without depth loss. The baseline algorithm denotes the DevNet with only photometric loss, smooth loss, and general min-reprojection mask. Based on the analysis in Sec. 3, adding the proposed components in the training process is useful to improving estimation accuracy, which is also verified in Tab. 4. Hence, the main concern here is which component contributes largest towards the model. By comparing the results of DevNet without certain components, it indicates that the position embedding strategy and the color regularization are less influential compared with the occlusion mask and depth loss. The performances of DevNet w/o occlusion mask and DevNet w/o depth mask are similar, as both of two terms focus on the spatial correspondence. However, the highest performance achieved by DevNet shows that the occlusion mask and the depth loss are beneficial towards each other. **Small Backbone:** Resnet50 is used as our backbone. To study the influence of backbone with different size, we also test the performance DevNet with Resnet18 in Tab. 4. We could notice that replacing the backbone R50 with the R18 could decline DevNet’s performance, while DevNet still outperforms the presented methods on KITTI-2015. **Sampled Planes Number:** DevNet splits the camera frustum into multiple parallel planes, therefore we are also interested in the influence of the number of plane samples towards the performance of DevNet. As demonstrated in Tab. 4, increasing the number of samples improves the estimation accuracy, which is noticed by comparing the performance of DevNet(P08), DevNet(P16), DevNet(P24) and DevNet(P32). However, it can be observed that when the number of samples is already large enough, increasing the number of samples will not result in a high margin in terms of prediction accuracy.

#### 4.4 Generalization

To test the generalization ability of our DevNet, we adapt the DevNet model trained on the KITTI-2015, to directly evaluate it above the NYU-V2 [37] without fine-tuning shown in Tab. 5 and Fig. 6. Despite huge differences, indoor and outdoor scenes still share some features, e.g., color inconsistency probably indicates depth difference, which makes this generalization possible. The network is

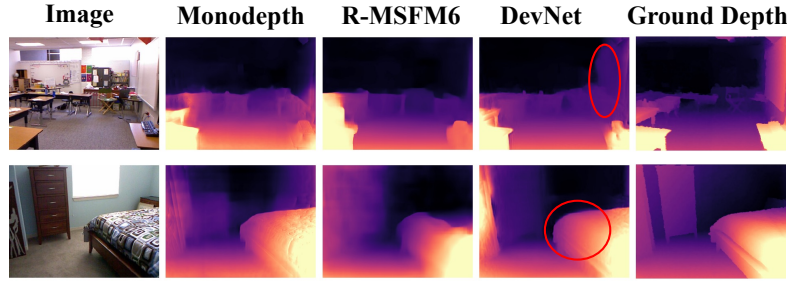


Fig. 6: **The qualitative results of depth estimation on the NYU-V2.** All methods are trained on KITTI-2015 and tested on NYU-V2.

fed with images with a resolution of  $288 \times 384$ . The comparison among FeatDepth, R-MSFM3, and R-MSFM6 indicates that better performance on the KITTI-2015 doesn't necessarily mean better performance on the NYU-V2. It also could be noticed that our model outperforms other monocular self-supervised methods in all evaluation metrics, which is a shred of strong evidence to support that our model has better generalization capability and could effectively learn the geometric information based on appearance information.

Algorithm	Type	Abs Rel	Sq Rel	RMSE	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$
MonoDepth2 [11]	M	0.391	0.898	1.458	0.415	0.724	0.874
FeatDepth [36]	M	0.353	0.701	1.231	0.516	0.776	0.904
R-MSFM3 [51]	M	0.355	0.672	1.276	0.476	0.755	0.898
R-MSFM6 [51]	M	0.372	0.803	1.344	0.494	0.746	0.884
DevNet(R50)(P24)	M	<b>0.333</b>	<b>0.605</b>	<b>1.142</b>	<b>0.541</b>	<b>0.790</b>	<b>0.911</b>

Table 5: **The quantitative results of depth estimation on NYU-V2.** Algorithms are trained on monocular video of KITTI-2015 and tested on NYU-V2.

## 5 Conclusion

In this work, we propose DevNet, an effective self-supervised monocular depth learning framework, that can accurately predict depth maps by constructing the density volume of the scene. To mitigate the influences of occlusions and illumination inconsistency, a novel occlusion mask and a brightness regularization are designed to benefit the training process of this framework. Moreover, the depth consistency loss function is designed to maintain the consistency of density volume for successive frames in the temporal domain. We show that using the training strategy consisting of them provides a simple and efficient model to produce monocular depth and odometry estimation accurately for both indoor and outdoor scenarios, which also possesses high generalization ability.

**Acknowledgements.** This work was partially supported by the Huawei UK AI Fellowship.

## References

1. Bhat, S.F., Alhashim, I., Wonka, P.: Adabins: Depth estimation using adaptive bins. In: CVPR (2021) [2](#)
2. Bian, J., Li, Z., Wang, N., Zhan, H., Shen, C., Cheng, M.M., Reid, I.: Unsupervised scale-consistent depth and ego-motion learning from monocular video. In: NeurIPS (2019) [4](#)
3. Du, Y., Zhang, Y., Yu, H.X., Tenenbaum, J.B., Wu, J.: Neural radiance flow for 4d view synthesis and video processing. In: ICCV (2021) [4](#)
4. Eigen, D., Fergus, R.: Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In: ICCV (2015) [2](#), [3](#), [10](#)
5. Eigen, D., Puhrsch, C., Fergus, R.: Depth map prediction from a single image using a multi-scale deep network. arXiv preprint arXiv:1406.2283 (2014) [11](#)
6. Engel, J., Koltun, V., Cremers, D.: Direct sparse odometry. TPAMI **40**(3), 611–625 (2017) [8](#)
7. Engel, J., Stückler, J., Cremers, D.: Large-scale direct slam with stereo cameras. In: IROS (2015) [8](#)
8. Garg, R., Bg, V.K., Carneiro, G., Reid, I.: Unsupervised cnn for single view depth estimation: Geometry to the rescue. In: ECCV (2016) [3](#)
9. Geiger, A., Lenz, P., Urtasun, R.: Are we ready for autonomous driving? the kitti vision benchmark suite. In: CVPR (2012) [3](#), [10](#)
10. Godard, C., Mac Aodha, O., Brostow, G.J.: Unsupervised monocular depth estimation with left-right consistency. In: CVPR (2017) [2](#), [7](#), [9](#)
11. Godard, C., Mac Aodha, O., Firman, M., Brostow, G.J.: Digging into self-supervised monocular depth estimation. In: ICCV (2019) [2](#), [4](#), [7](#), [9](#), [10](#), [11](#), [12](#), [14](#)
12. GonzalezBello, J.L., Kim, M.: Forget about the lidar: Self-supervised depth estimators with med probability volumes. In: NeurIPS (2020) [10](#)
13. Gordon, A., Li, H., Jonschkowski, R., Angelova, A.: Depth from videos in the wild: Unsupervised monocular depth learning from unknown cameras. In: ICCV (2019) [2](#), [4](#), [10](#)
14. Guizilini, V., Ambrus, R., Pillai, S., Raventos, A., Gaidon, A.: 3d packing for self-supervised monocular depth estimation. In: CVPR (2020) [2](#), [4](#), [10](#)
15. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016) [6](#)
16. Huo, Y., Yi, H., Wang, Z., Shi, J., Lu, Z., Luo, P., et al.: Learning depth-guided convolutions for monocular 3d object detection. In: CVPR Workshop (2020) [2](#)
17. Ji, P., Li, R., Bhanu, B., Xu, Y.: Monoindoor: Towards good practice of self-supervised monocular depth estimation for indoor environments. In: ICCV (2021) [11](#)
18. Jin, H., Favaro, P., Soatto, S.: Real-time feature tracking and outlier rejection with changes in illumination. In: ICCV (2001) [8](#), [11](#)
19. Johnston, A., Carneiro, G.: Self-supervised monocular trained depth estimation using self-attention and discrete disparity volume. In: CVPR (2020) [4](#), [10](#), [11](#)
20. Kajiya, J.T.: The rendering equation. In: Annual Conference on Computer Graphics and Interactive Techniques (1986) [2](#)
21. Kalia, M., Navab, N., Salcudean, T.: A real-time interactive augmented reality depth estimation technique for surgical robotics. In: ICRA (2019) [2](#)
22. Lee, J.H., Han, M.K., Ko, D.W., Suh, I.H.: From big to small: Multi-scale local planar guidance for monocular depth estimation. arXiv preprint arXiv:1907.10326 (2019) [3](#)



23. Li, Y., Ushiku, Y., Harada, T.: Pose graph optimization for unsupervised monocular visual odometry. In: ICRA (2019) 12
24. Li, Z., Niklaus, S., Snavely, N., Wang, O.: Neural scene flow fields for space-time view synthesis of dynamic scenes. In: CVPR (2021) 4
25. Liu, H., Tang, X., Shen, S.: Depth-map completion for large indoor scene reconstruction. *Pattern Recognition* **99**, 107112 (2020) 2
26. Luo, C., Yang, Z., Wang, P., Wang, Y., Xu, W., Nevatia, R., Yuille, A.: Every pixel counts++: Joint learning of geometry and motion with 3d holistic understanding. *TPAMI* **42**(10), 2624–2641 (2019) 10
27. Martin-Brualla, R., Radwan, N., Sajjadi, M., Barron, J.T., Dosovitskiy, A., Duckworth, D.: NeRF in the wild: Neural radiance fields for unconstrained photo collections. *arXiv preprint arXiv:2008.02268* (2020) 4
28. Max, N.: Optical models for direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics* **1**(2), 99–108 (1995) 5
29. Meng, Y., Lu, Y., Raj, A., Sunarjo, S., Guo, R., Javidi, T., Bansal, G., Bhargava, D.: Signet: Semantic instance aided unsupervised 3d geometry perception. In: CVPR (2019) 10
30. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: ECCV (2020) 4, 5, 6, 9
31. Peng, R., Wang, R., Lai, Y., Tang, L., Cai, Y.: Excavating the potential capacity of self-supervised monocular depth estimation. In: ICCV (2021) 2, 3
32. Pillai, S., Ambruş, R., Gaidon, A.: Superdepth: Self-supervised, super-resolved monocular depth estimation. In: ICRA (2019) 2, 4
33. Pumarola, A., Corona, E., Pons-Moll, G., Moreno-Noguer, F.: D-nerf: Neural radiance fields for dynamic scenes. In: CVPR (2021) 4
34. Reading, C., Harakeh, A., Chae, J., Waslander, S.L.: Categorical depth distribution network for monocular 3d object detection. In: CVPR (2021) 3
35. Saxena, A., Chung, S.H., Ng, A.Y., et al.: Learning depth from single monocular images. In: *NeurIPS* (2005) 3
36. Shu, C., Yu, K., Duan, Z., Yang, K.: Feature-metric loss for self-supervised learning of depth and egomotion. In: ECCV (2020) 2, 4, 7, 10, 11, 12, 14
37. Silberman, N., Hoiem, D., Kohli, P., Fergus, R.: Indoor segmentation and support inference from rgb-d images. In: ECCV (2012) 3, 10, 11, 13
38. Sun, J., Chen, L., Xie, Y., Zhang, S., Jiang, Q., Zhou, X., Bao, H.: Disp R-CNN: Stereo 3d object detection via shape prior guided instance disparity estimation. In: CVPR (2020) 2
39. Wang, C., Buenaposada, J.M., Zhu, R., Lucey, S.: Learning depth from monocular videos using direct methods. In: CVPR (2018) 9
40. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing* **13**(4), 600–612 (2004) 9
41. Wei, Y., Liu, S., Rao, Y., Zhao, W., Lu, J., Zhou, J.: Nerfingmvs: Guided optimization of neural radiance fields for indoor multi-view stereo. In: ICCV (2021) 4, 5
42. Yang, N., Stumberg, L.v., Wang, R., Cremers, D.: D3vo: Deep depth, deep pose and deep uncertainty for monocular visual odometry. In: CVPR (2020) 7, 8
43. Yen-Chen, L., Florence, P., Barron, J.T., Rodriguez, A., Isola, P., Lin, T.Y.: iNERF: Inverting neural radiance fields for pose estimation. *arXiv preprint arXiv:2012.05877* (2020) 4



44. Yin, Z., Shi, J.: Geonet: Unsupervised learning of dense depth, optical flow and camera pose. In: CVPR. pp. 1983–1992 (2018) [4](#)
45. Yuan, W., Gu, X., Dai, Z., Zhu, S., Tan, P.: New crfs: Neural window fully-connected crfs for monocular depth estimation. arXiv preprint arXiv:2203.01502 (2022) [10](#), [12](#)
46. Zhan, H., Garg, R., Weerasekera, C.S., Li, K., Agarwal, H., Reid, I.: Unsupervised learning of monocular depth estimation and visual odometry with deep feature reconstruction. In: CVPR (2018) [2](#), [10](#), [12](#)
47. Zhang, D., Lo, F.P.W., Zheng, J.Q., Bai, W., Yang, G.Z., Lo, B.: Data-driven microscopic pose and depth estimation for optical microrobot manipulation. *Acs Photonics* **7**(11), 3003–3014 (2020) [2](#)
48. Zhao, W., Liu, S., Shu, Y., Liu, Y.J.: Towards better generalization: Joint depth-pose learning without posenet. In: CVPR (2020) [11](#)
49. Zhou, J., Wang, Y., Qin, K., Zeng, W.: Unsupervised high-resolution depth learning from videos with dual networks. In: ICCV (2019) [3](#), [10](#)
50. Zhou, T., Brown, M., Snavely, N., Lowe, D.G.: Unsupervised learning of depth and ego-motion from video. In: CVPR (2017) [2](#), [3](#), [4](#), [10](#), [11](#), [12](#)
51. Zhou, Z., Fan, X., Shi, P., Xin, Y.: R-MSFM: Recurrent multi-scale feature modulation for monocular depth estimating. In: ICCV (2021) [2](#), [4](#), [10](#), [14](#)
52. Zhu, S., Brazil, G., Liu, X.: The edge of depth: Explicit constraints between segmentation and depth. In: CVPR (2020) [2](#)