

RA-Depth: Resolution Adaptive Self-Supervised Monocular Depth Estimation

Mu He, Le Hui, Yikai Bian, Jian Ren, Jin Xie*, and Jian Yang*

Key Lab of Intelligent Perception and Systems for High-Dimensional Information of
Ministry of Education

Jiangsu Key Lab of Image and Video Understanding for Social Security
PCA Lab, School of Computer Science and Engineering
Nanjing University of Science and Technology, China
{muhe, le.hui, yikai.bian, renjian, csjxie, csjyang}@njjust.edu.cn

Abstract. Existing self-supervised monocular depth estimation methods can get rid of expensive annotations and achieve promising results. However, these methods suffer from severe performance degradation when directly adopting a model trained on a fixed resolution to evaluate at other different resolutions. In this paper, we propose a resolution adaptive self-supervised monocular depth estimation method (RA-Depth) by learning the scale invariance of the scene depth. Specifically, we propose a simple yet efficient data augmentation method to generate images with arbitrary scales for the same scene. Then, we develop a dual high-resolution network that uses the multi-path encoder and decoder with dense interactions to aggregate multi-scale features for accurate depth inference. Finally, to explicitly learn the scale invariance of the scene depth, we formulate a cross-scale depth consistency loss on depth predictions with different scales. Extensive experiments on the KITTI, Make3D and NYU-V2 datasets demonstrate that RA-Depth not only achieves state-of-the-art performance, but also exhibits a good ability of resolution adaptation. Source code is available at <https://github.com/hmhemu/RA-Depth>.

Keywords: Self-Supervised Learning, Monocular Depth Estimation, Resolution Adaptation

1 Introduction

Monocular depth estimation [38,5,21,25,20,6] recovers a pixel-wise depth map from a single image, which is a challenging but essential task in computer vision. It has promoted the development of various applications such as autonomous driving, robot navigation, 3D scene reconstruction, and augmented reality. Due to the difference of acquisition devices in different application scenarios, the resolutions of the acquired images are different. Thus, to reduce resource consumption, how to make a single trained model of monocular depth estimation

* Corresponding authors.

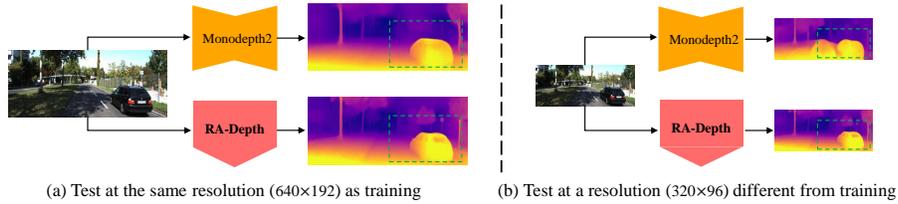


Fig. 1. Depth predictions for different resolutions. We compare the depth predictions of Monodepth2 [10] and our RA-Depth at different resolutions. Both Monodepth2 and RA-Depth are trained at the fixed resolution of 640×192 , and then tested at 640×192 (a) and 320×96 (b). RA-Depth exhibits the good resolution adaptation ability when tested at a resolution different from training.

adaptive to the changes of image resolution in a self-supervised manner is a valuable and open problem.

Existing self-supervised monocular depth estimation methods [9,28,49,10,27] usually use geometrical constraints on stereo pairs or monocular sequences as the sole source of supervision and have made great progress. However, these supervision methods cannot adapt to the changes of image resolution at inference, whose performance drops severely when the test resolution is inconsistent with the training resolution. As shown in Fig. 1, we compare the depth prediction results of existing advanced work Monodepth2 [10] and our RA-Depth. Both Monodepth2 and RA-Depth are trained at the fixed resolution of 640×192 and then tested at two resolutions 640×192 and 320×96 . It can be seen that Monodepth2 predicts inconsistent and poor depth maps when tested at different resolutions. To test at different resolutions, existing methods train an individual model for each resolution. As a result, this not only requires unavoidable training overhead and high storage cost, but also limits the flexibility and practicality of monocular depth estimation models.

In this paper, we propose a resolution adaptive self-supervised monocular depth estimation method named as RA-Depth. First, we propose an arbitrary-scale data augmentation method to generate images with arbitrary scales for the same scene. Specifically, we change the scale of the input image while maintaining the image resolution/size (*i.e.*, width \times height) through randomly resizing, cropping, and stitching (details in Section 3.2). These generated images with arbitrary scales can mimic the scale variations at different image resolutions, thus prompting the model to implicitly learn the scale invariance of scene depth. Then, we propose an efficient monocular depth estimation framework with multi-scale feature fusion. In our framework, we use the superior high-resolution representation network HRNet [40] as encoder, and design a decoder with efficient multi-scale feature fusion. As a result, the encoder and decoder form a dual HRNet architecture, which can fully exploit and aggregate multi-scale features to infer the accurate depth. Thanks to this efficient design, we can achieve better depth estimation performance with the lower network overhead. Finally, we propose a novel cross-scale depth consistency loss to explicitly learn

the scale invariance of scene depth. Thus, the monocular depth estimation model can predict the consistent depth maps for the same scene, even if the scales of input images are different. Extensive experiments on the KITTI [8], Make3D [38], and NYU-V2 [39] datasets demonstrate that RA-Depth achieves state-of-the-art performance on monocular depth estimation and shows a good ability of resolution adaptation.

To summarize, the main contributions of our work are as follows:

- To the best of our knowledge, our RA-Depth is the first work to tackle the image resolution adaptation for self-supervised monocular depth estimation.
- We propose an arbitrary-scale data augmentation method to promote the model to learn depths from images with different scales.
- We develop an efficient dual high-resolution network with multi-scale feature fusion for monocular depth estimation, which is trained with a novel cross-scale depth consistency loss.
- Extensive experiments demonstrate that our RA-Depth achieves state-of-the-art performance on monocular depth estimation, and has a good generalization ability to different resolutions.

2 Related Work

2.1 Supervised Monocular Depth Estimation

Monocular depth estimation aims to predict the scene depth from single color image. Early works [5,24,21,43,16,22,44] estimated continuous depth maps pixel by pixel using convolutional neural networks by minimizing the error between the prediction and ground-truth depths. Furthermore, the geometry [44,31] or multi-task constraints [35,47] are imposed on the depth estimation network to predict more accurate depth maps. However, these continuous regression methods usually suffer from slow convergence and unsatisfactory local solutions. Deep ordinal regression network [6] thus introduced a spacing-increasing discretization strategy to discretize depth, where the depth estimation problem is converted into the ordinal depth regression problem. AdaBins [1] proposed a Transformer based depth estimation scheme by dividing the depth range into different bins. Recently, some works [29,23] employed the multimodel learning for depth estimation and further improved its performance.

2.2 Self-Supervised Monocular Depth Estimation

Self-supervised methods convert the monocular depth estimation task into an image reconstruction problem, and they use the photometric loss on stereo pairs or monocular sequences as the supervisory signal to train networks.

One type of self-supervised monocular depth estimation methods employ the geometric constraints on stereo pairs to learn depth. Garg *et al.* [7] first propose a self-supervised monocular depth estimation framework, where depth is estimated from synchronized stereo image pairs. Following this framework,

Monodepth [9] designs a left-right disparity consistency loss to produce more accurate prediction results, while 3Net [34] introduces a trinocular stereo assumption to solve the occlusion problem between stereo pairs. Based on these methods, Depth Hints [42] uses the off-the-shelf Semi-Global Matching (SGM) algorithm [14,15] to obtain complementary depth hints, which can enhance existing photometric loss and guide the network to learn more accurate weights. Recently, EPCDepth [32] adopted a data grafting technique to make the model focus on non-vertical image positions for accurate depth inference.

Another type of depth estimation methods with self-supervision come from unlabelled monocular videos. SfM-Learner [49] is a pioneering work that jointly learns monocular depth and ego-motion from monocular videos in a self-supervised way. Based on SfM-Learner, SC-SfMLearner [2] proposes a geometry consistency loss to ensure the scale-consistency of depth. Monodepth2 [10] introduces a minimum reprojection loss and an auto-masking loss to solve the problem of occlusions and moving objects. Poggi *et al.* [33] explore the uncertainty of self-supervised monocular depth estimation and propose a novel peculiar technique specifically designed for self-supervised depth estimation. Besides, many recent works [3,19,4,50,12] utilize extra semantic supervision to aid self-supervised monocular depth estimation. More recently, some works [27,48] have developed high-performance monocular depth estimation models through efficient network architecture designs.

3 Method

3.1 Problem Formulation

Self-supervised methods exploit geometric constraints on stereo pairs or monocular videos as the supervision signal for training. In this paper, our proposed framework uses unlabelled monocular videos for training, each training instance contains a target frame \mathbf{I}_t and two source frames \mathbf{I}_s ($s \in \{t-1, t+1\}$) adjacent to \mathbf{I}_t temporally. With the depth \mathbf{D}_t of the target view and the camera pose $\mathbf{T}_{t \rightarrow s}$ between the target view and source view, the image synthesis process from the source view to the target view can be expressed as $\mathbf{I}_{s \rightarrow t} = \mathbf{I}_s \langle \text{proj}(\mathbf{D}_t, \mathbf{T}_{t \rightarrow s}, \mathbf{K}) \rangle$, where \mathbf{K} represents the known camera intrinsics, $\langle \rangle$ is the sampling operator [17] and $\text{proj}()$ is the coordinate projection operation [49]. The photometric error is composed of L_1 and SSIM [41]:

$$pe(\mathbf{I}_a, \mathbf{I}_b) = \frac{\alpha}{2}(1 - SSIM(\mathbf{I}_a, \mathbf{I}_b)) + (1 - \alpha) \|\mathbf{I}_a - \mathbf{I}_b\|_1. \quad (1)$$

Following [10], in order to handle the issue of occlusion, we adopt the per-pixel minimum reprojection loss as our photometric loss:

$$\mathcal{L}_{ph}(\mathbf{I}_t, \mathbf{I}_{s \rightarrow t}) = \min_s pe(\mathbf{I}_t, \mathbf{I}_{s \rightarrow t}). \quad (2)$$

The edge-aware smoothness loss is used to deal with disparity discontinuities:

$$\mathcal{L}_{sm}(\mathbf{D}_t, \mathbf{I}_t) = |\partial_x \mathbf{D}_t^*| e^{-|\partial_x \mathbf{I}_t|} + |\partial_y \mathbf{D}_t^*| e^{-|\partial_y \mathbf{I}_t|} \quad (3)$$

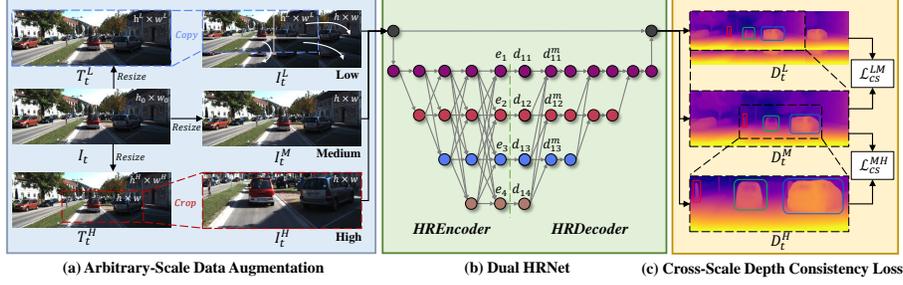


Fig. 2. Overview of RA-Depth. (a) First, given an original instance $\{\mathbf{I}_t, \mathbf{I}_s\}$ ($s \in \{t-1, t+1\}$, t for target image and s for source image), we use our arbitrary-scale data augmentation to generate three training instances with different scales: $\{\mathbf{I}_t^L, \mathbf{I}_s^L\}$, $\{\mathbf{I}_t^M, \mathbf{I}_s^M\}$, and $\{\mathbf{I}_t^H, \mathbf{I}_s^H\}$. For simplicity, only target images \mathbf{I}_t^L , \mathbf{I}_t^M , and \mathbf{I}_t^H are shown in this figure. (b) Then, we use our proposed efficient Dual HRNet to separately predict depth maps (\mathbf{D}_t^L , \mathbf{D}_t^M , \mathbf{D}_t^H) for these target images. Note that Dual HRNet only inputs a single target image and then outputs the corresponding depth map at a time, and all target images share the same Dual HRNet. (c) Finally, we compute the proposed cross-scale depth consistency losses \mathcal{L}_{cs}^{LM} (between \mathbf{D}_t^L and \mathbf{D}_t^M) and \mathcal{L}_{cs}^{MH} (between \mathbf{D}_t^M and \mathbf{D}_t^H) to constrain the network to learn the scale invariance of the scene depth.

where ∂_x and ∂_y are the gradients in the horizontal and vertical direction respectively. Besides, $\mathbf{D}_t^* = \mathbf{D}_t / \overline{\mathbf{D}_t}$ is the mean-normalized inverse depth from [10] to discourage shrinking of the estimated depth.

3.2 Resolution Adaptive Self-Supervised Framework

Overview. An overview of our proposed method can be seen in Fig. 2. Given an original instance $\{\mathbf{I}_t, \mathbf{I}_s\}$ ($s \in \{t-1, t+1\}$, t for target image and s for source image), we employ our proposed arbitrary-scale data augmentation to expand three training instances: $\{\mathbf{I}_t^L, \mathbf{I}_s^L\}$, $\{\mathbf{I}_t^M, \mathbf{I}_s^M\}$, and $\{\mathbf{I}_t^H, \mathbf{I}_s^H\}$, where L , M and H denote the low, middle and high scales, respectively. For simplicity, we only show \mathbf{I}_t^L , \mathbf{I}_t^M , and \mathbf{I}_t^H in Fig. 2(a). Then, we use our proposed efficient Dual HRNet to separately predict depth maps (\mathbf{D}_t^L , \mathbf{D}_t^M , \mathbf{D}_t^H) of these target views, and use the pose network to predict the relative 6D camera pose $\mathbf{T}_{t \rightarrow s}$ from the target view to the source view. With \mathbf{D}_t^L and $\mathbf{T}_{t \rightarrow s}$, we can synthesize image $\tilde{\mathbf{I}}_{s \rightarrow t}^L$ by interpolating the source image \mathbf{I}_s^L . In the same way, we can synthesize images $\tilde{\mathbf{I}}_{s \rightarrow t}^M$ and $\tilde{\mathbf{I}}_{s \rightarrow t}^H$. Then we compute the photometric losses between real images ($\mathbf{I}_t^L, \mathbf{I}_t^M, \mathbf{I}_t^H$) and synthesized images ($\tilde{\mathbf{I}}_{s \rightarrow t}^L, \tilde{\mathbf{I}}_{s \rightarrow t}^M, \tilde{\mathbf{I}}_{s \rightarrow t}^H$). Note that all image synthesis processes use the same camera pose $\mathbf{T}_{t \rightarrow s}$ predicted from \mathbf{I}_t^M to \mathbf{I}_s^M . Finally, the proposed cross-scale depth consistency losses \mathcal{L}_{cs}^{LM} (between \mathbf{D}_t^L and \mathbf{D}_t^M) and \mathcal{L}_{cs}^{MH} (between \mathbf{D}_t^M and \mathbf{D}_t^H) are computed. For clarity, the pose network, photometric loss, and smoothness loss are not shown in Fig. 2.

Algorithm 1 Arbitrary-Scale Data Augmentation.

-
- Input:** An original image \mathbf{I} of shape (c, h_0, w_0) .
Output: Three training images $\mathbf{I}^L, \mathbf{I}^M, \mathbf{I}^H$ of shape (c, h, w) .
- 1: Random initializing the scale factors s^L and s^H with respective ranges of $[0.7, 0.9]$ and $[1.1, 2.0]$; Zero initializing $\mathbf{I}^L, \mathbf{I}^M, \mathbf{I}^H$ of shape (c, h, w) ;
 - 2: $(h^L, w^L), (h^M, w^M), (h^H, w^H) = \text{int}((h, w) \times s^L), (h, w), \text{int}((h, w) \times s^H)$;
 - 3: Generating the low-resolution image (\mathbf{I}^L);
 - (a) Resizing the original image \mathbf{I}
 $\mathbf{T}^L \leftarrow \text{Resize}(\mathbf{I})$ to (c, h^L, w^L)
 - (b) Slicing and copying \mathbf{T}^L
 $\mathbf{I}^L[:, : h^L, : w^L] \leftarrow \mathbf{T}^L[:, :, :]$
 $\mathbf{I}^L[:, h^L : h, : w^L] \leftarrow \mathbf{T}^L[:, (2 \times h^L - h) : h^L, : w^L]$
 $\mathbf{I}^L[:, : h^L, w^L : w] \leftarrow \mathbf{T}^L[:, : h^L, (2 \times w^L - w) : w^L]$
 $\mathbf{I}^L[:, h^L : h, w^L : w] \leftarrow \mathbf{T}^L[:, (2 \times h^L - h) : h^L, (2 \times w^L - w) : w^L]$
 - 4: Generating the middle-resolution image (\mathbf{I}^M);
 - (a) Resizing the original image \mathbf{I}
 $\mathbf{T}^M \leftarrow \text{Resize}(\mathbf{I})$ to (c, h^M, w^M)
 - (b) Copying \mathbf{T}^M
 $\mathbf{I}^M[:, :, :] \leftarrow \mathbf{T}^M[:, :, :]$
 - 5: Generating the high-resolution image (\mathbf{I}^H);
 - (a) Resizing the original image \mathbf{I}
 $\mathbf{T}^H \leftarrow \text{Resize}(\mathbf{I})$ to (c, h^H, w^H)
 - (b) Randomly cropping \mathbf{T}^H
 $\mathbf{I}^H[:, :, :] \leftarrow \text{RandomCrop}(\mathbf{T}^H)$ to (c, h, w)
 - 6: Return $\mathbf{I}^L, \mathbf{I}^M, \mathbf{I}^H$;
-

Arbitrary-Scale Data Augmentation. The details of arbitrary-scale data augmentation for a single input are shown in Algorithm 1. For a single original image \mathbf{I} of shape (c, h_0, w_0) , we first obtain three images ($\mathbf{T}^L, \mathbf{T}^M$, and \mathbf{T}^H) with different resolutions by the *Resize* operation. The resolution of \mathbf{T}^M is fixed to (h, w) , which denotes the middle scale. The resolutions of \mathbf{T}^L and \mathbf{T}^H vary randomly over a continuous range, controlled by the scale factors s^L and s^H , respectively. The ranges of s^L and s^H are $[0.7, 0.9]$ and $[1.1, 2.0]$, respectively, making the resolution of \mathbf{T}^L lower than (h, w) and the resolution of \mathbf{T}^H higher than (h, w) . Then we can generate a low scale image \mathbf{I}^L from \mathbf{T}^L by means of image stitching (represented by white arrows in Fig. 2(a)), and directly treat \mathbf{T}^M as the middle scale image \mathbf{I}^M , and generate the high scale image \mathbf{I}^H from \mathbf{T}^H by means of *RandomCrop* (represented by red dashed box in Fig. 2(a)). Finally, we obtain three images $\mathbf{I}^L, \mathbf{I}^M$, and \mathbf{I}^H with the same shape of (c, h, w) at different scales, which served as the input of the network. Note, in this paper, (c, h_0, w_0) is $(3, 375, 1242)$, and (c, h, w) is $(3, 192, 640)$.

Dual HRNet for Monocular Depth Estimation. As shown in Fig. 2(b), we develop an efficient dual high-resolution network (Dual HRNet) with multi-scale feature fusion for monocular depth estimation. We use the superior high-resolution representation network HRNet18 [40] as the encoder, named

HREncoder. The core idea of HRNet [40] is to perform efficient multi-scale feature fusion while maintaining high-resolution feature representations, so as to obtain semantically richer and spatially more precise multi-scale features. Inspired by HRNet [40], we design a decoder with efficient multi-scale feature fusion named HRDecoder. Specifically, HRDecoder first inherits the multi-scale features of the encoder, and then gradually fuses low-scale features while maintaining the high-resolution feature representations. Let e_i denote the feature of the encoder at the i -th level, d_{ji} denote the feature of the decoder at the i -th level and the j -th stage, and d_{ji}^m denote the feature after multi-scale feature fusion at the i -th level and the j -th stage. Let us take the calculation of d_{1i} and d_{1i}^m (refer to Fig. 2(b)) at the first stage as an example:

$$\begin{cases} d_{1i} = \text{CONV}_{3 \times 3}(e_i), & i = 1, 2, 3, 4 \\ d_{1i}^m = d_{1i} + [\text{CONV}_{1 \times 1}(\mu(d_{1k}))] & i = 1, 2, 3, k = i + 1, \dots, 4 \end{cases} \quad (4)$$

where $\text{CONV}_{3 \times 3}$ represents a 3×3 convolution layer, and $\text{CONV}_{1 \times 1}$ represents an 1×1 convolution layer. In addition, $\mu(\cdot)$ is an upsampling operator, and $[\cdot]$ is a summation operator. By exploiting the HRDecoder, the multi-scale features can be fully fused from the low-resolution feature map to the high-resolution feature map. It is desired that the Dual HRNet can effectively learn depth information of different scales for accurate depth estimation at different test resolutions.

Cross-Scale Depth Consistency Loss. As shown in Fig. 2(c), we compute the cross-scale depth consistency losses \mathcal{L}_{cs}^{ML} (between D_t^M and D_t^L) and \mathcal{L}_{cs}^{MH} (between D_t^M and D_t^H) to constrain the model to explicitly learn the scale invariance of scene depth. For example, the same pillar or car among D_t^L , D_t^M and D_t^H should have the same depth as shown in Fig. 2(c). Specifically, according to the known scale factors s^L and s^H as well as the cropped position of *RandomCrop* in Algorithm 1, we can obtain the pixel correspondence among D_t^L , D_t^M and D_t^H . For D_t^M and D_t^H , we first find the corresponding position of D_t^H in D_t^M , which is represented by the black dashed box (named as \tilde{D}_t^M) in D_t^M . Then we resize D_t^H to keep the same size as \tilde{D}_t^M to obtain \tilde{D}_t^H . Finally, we calculate the scene depth consistency loss between \tilde{D}_t^M and \tilde{D}_t^H :

$$\mathcal{L}_{cs}^{MH} = \frac{\alpha}{2}(1 - SSIM(\tilde{D}_t^M, \tilde{D}_t^H)) + (1 - \alpha) \left\| \tilde{D}_t^M - \tilde{D}_t^H \right\|_1. \quad (5)$$

Similarly, for D_t^L and D_t^M , we first find the corresponding position of D_t^M in D_t^L , which is represented by the black dashed box (named as \hat{D}_t^L) in D_t^L . Then we resize D_t^M to keep the same size as \hat{D}_t^L to obtain \hat{D}_t^M . Finally, we calculate the cross-scale depth consistency loss between \hat{D}_t^L and \hat{D}_t^M :

$$\mathcal{L}_{cs}^{LM} = \frac{\alpha}{2}(1 - SSIM(\hat{D}_t^L, \hat{D}_t^M)) + (1 - \alpha) \left\| \hat{D}_t^L - \hat{D}_t^M \right\|_1. \quad (6)$$

Final Training Loss. Our overall loss function can be formulated as follows:

$$\begin{aligned} \mathcal{L}_f = & \gamma(\mathcal{L}_{ph}(\mathbf{I}_t^L, \tilde{\mathbf{I}}_{s \rightarrow t}^L) + \mathcal{L}_{ph}(\mathbf{I}_t^M, \tilde{\mathbf{I}}_{s \rightarrow t}^M) + \mathcal{L}_{ph}(\mathbf{I}_t^H, \tilde{\mathbf{I}}_{s \rightarrow t}^H)) \\ & + \lambda(\mathcal{L}_{sm}(\tilde{\mathbf{D}}_t^L, \mathbf{I}_t^L) + \mathcal{L}_{sm}(\tilde{\mathbf{D}}_t^M, \mathbf{I}_t^M) + \mathcal{L}_{sm}(\tilde{\mathbf{D}}_t^H, \mathbf{I}_t^H)) \\ & + \beta(\mathcal{L}_{cs}^{LM} + \mathcal{L}_{cs}^{MH}) \end{aligned} \quad (7)$$

Here, γ , λ , and β are the hyper-parameters of the photometric loss, smoothness loss, and our proposed scene depth consistency loss.

4 Experiments

4.1 Implementation Details

Our proposed method is implemented in PyTorch [30] and trained on a single Titan RTX GPU. Our models are trained for 20 epochs using Adam optimizer [18], with a batch size of 12. The learning rate for the first 15 epochs is 10^{-4} and then dropped to 10^{-5} for the last 5 epochs. The hyper-parameters γ , λ , and β of the final training loss in Eq. (7) are set to 1.0, 0.001, and 1.0, respectively. We set $\alpha = 0.85$ in Eq. (1), Eq. (5), and Eq. (6). For the monocular depth estimation network, We implement our proposed Dual HRNet described in Section 3.2, which uses HRNet18 [40] as the encoder and performs efficient multi-scale feature fusion on the decoder. For the pose estimation network, we use the same architecture as Monodepth2 [10], which contains ResNet18 [13] followed by several convolutional layers. Similar to [10,48], we use the weights pretrained on ImageNet [37] to initialize HRNet18 and ResNet18. During training, the input resolution of networks is 640×192 . To improve the training speed, we only output a single-scale depth for the depth estimation network and compute the loss on the single-scale depth, instead of computing losses on multi-scale outputs of the depth estimation network (4 scales in [49,10] and 6 scales in [36]).

4.2 Results

Basic Results. We evaluate the monocular depth estimation performance of our single model on the KITTI dataset [8] using the data split of Eigen *et al.* [5]. Following Zhou *et al.* [49], we remove static frames, using 39810, 4424, and 697 images for training, validation, and test, respectively. During the evaluation, we cap depth to 80m and adopt the same conventional metrics as [5]. Since we use monocular videos for training, we perform median scaling introduced by [49] on our predicted depths when evaluating. Table 1 shows the basic quantitative results and all models are tested at the same resolution as training. Our proposed method outperforms all existing self-supervised monocular depth estimation approaches significantly. In particular, RA-Depth achieves the SqRel error of 0.632, the error reduction rate of 17% over the best-published result 0.764.

Table 1. Quantitative monocular depth estimation results on the KITTI dataset [8] using the Eigen split [5]. We divide compared methods into three categories based on training data. In each category, the best results are in **bold** and the second are underlined. The test resolution is the same as the training resolution for all these methods. S: trained on stereo pairs; M: trained on monocular videos; Se: trained with semantic labels.

Method	Train	Resolution	Error Metric ↓				Accuracy Metric ↑		
			AbsRel	SqRel	RMSE	RMSElog	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$
Monodepth [9]	S	512×256	0.148	1.344	5.927	0.247	0.803	0.922	0.964
3Net [34]	S	512×256	0.142	1.207	5.702	0.240	0.809	0.928	0.967
Chen [3]	S+Se	512×256	0.118	0.905	5.096	0.211	0.839	0.945	0.977
Monodepth2 [10]	S	640×192	0.109	0.873	4.960	0.209	0.864	0.948	0.975
Depth Hints [42]	S	640×192	0.106	0.780	4.695	0.193	0.875	0.958	0.980
EPCDepth [32]	S	640×192	0.099	0.754	4.490	0.183	0.888	0.963	0.982
Zhan [46]	M+S	608×160	0.135	1.132	5.585	0.229	0.820	0.933	0.971
EPC++ [26]	M+S	832×256	0.128	0.935	5.011	0.209	0.831	0.945	0.979
Monodepth2 [10]	M+S	640×192	0.106	0.818	4.750	0.196	0.874	0.957	0.979
HR-Depth [27]	M+S	640×192	0.107	0.785	4.612	0.185	0.887	0.962	0.982
Depth Hints [42]	M+S	640×192	0.105	0.769	4.627	0.189	0.875	0.959	0.982
DIFFNet [48]	M+S	640×192	0.101	0.749	4.445	0.179	0.898	0.965	0.983
SfMLearner [49]	M	640×192	0.183	1.595	6.709	0.270	0.734	0.902	0.959
EPC++ [26]	M	832×256	0.141	1.029	5.350	0.216	0.816	0.941	0.976
SC-SfMLearner [2]	M	832×256	0.114	0.813	4.706	0.191	0.873	0.960	0.982
Monodepth2 [10]	M	640×192	0.115	0.903	4.863	0.193	0.877	0.959	0.981
SGDDepth [19]	M+Se	640×192	0.113	0.835	4.693	0.191	0.879	0.961	0.981
SAFENet [4]	M+Se	640×192	0.112	0.788	4.582	0.187	0.878	0.963	0.983
PackNet-SfM [11]	M	640×192	0.111	0.785	4.601	0.189	0.878	0.960	0.982
Mono-Uncertainty [33]	M	640×192	0.111	0.863	4.756	0.188	0.881	0.961	0.982
HR-Depth [27]	M	640×192	0.109	0.792	4.632	0.185	0.884	0.962	0.983
DIFFNet [48]	M	640×192	0.102	0.764	4.483	0.180	0.896	0.965	0.983
RA-Depth	M	640×192	0.096	0.632	4.216	0.171	0.903	0.968	0.985

Resolution Adaptation Results. As shown in Table 2, we compare our single model RA-Depth with existing three advanced approaches including Monodepth2 [10], HR-Depth [27], and DIFFNet [48] in terms of resolution adaptation on the KITTI dataset [8] using the Eigen split [5]. All these models are trained at the resolution of 640×192 and then tested at four different resolutions including 416×128, 512×160, 832×256, and 1024×320. Note, we directly resize the original image to generate four resolution images without using the scale augmentation strategy in the training process. The results show that existing methods perform poorly in terms of resolution adaptation. That is to say, existing methods cannot adapt to the changes of image resolution during testing. In contrast, our single model RA-Depth maintains high performance across different test resolutions, significantly outperforming existing self-supervised methods. Fig. 3 also shows the resolution adaptation results of different models on two error metrics SqRel and RMSE and one accuracy metric $\delta < 1.25$. It can be seen that our single model RA-Depth achieves the lowest error on two error metrics at all resolutions, even lower than the errors of these individual models trained separately at each test resolution. Moreover, our single model achieves the highest accuracy at all resolutions compared with other single models and is comparable to these individual models. Fig. 4 illustrates the visualization results

Table 2. Resolution adaptive comparison results on the KITTI dataset [8] using the Eigen split [5]. All these models are trained at the resolution of 640×192 and then tested at four different resolutions including 416×128 , 512×160 , 832×256 , and 1024×320 .

Method	Test Resolution	Error Metric ↓				Accuracy metric ↑		
		AbsRel	SqRel	RMSE	RMSElog	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$
Monodepth2 [10]	416×128	0.184	1.365	6.146	0.268	0.719	0.911	0.966
HR-Depth [27]	416×128	0.157	1.120	5.669	0.231	0.787	0.938	0.976
DIFFNet [48]	416×128	0.142	1.068	5.552	0.218	0.826	0.948	0.979
RA-Depth	416×128	0.111	0.723	4.768	0.187	0.874	0.961	0.984
Monodepth2 [10]	512×160	0.132	1.062	5.318	0.209	0.846	0.951	0.978
HR-Depth [27]	512×160	0.128	0.978	4.999	0.199	0.855	0.957	0.981
DIFFNet [48]	512×160	0.121	1.007	5.107	0.196	0.870	0.958	0.981
RA-Depth	512×160	0.101	0.658	4.373	0.175	0.895	0.967	0.985
Monodepth2 [10]	832×256	0.131	0.992	5.451	0.222	0.815	0.947	0.978
HR-Depth [27]	832×256	0.136	0.920	5.321	0.218	0.817	0.949	0.980
DIFFNet [48]	832×256	0.128	0.896	5.328	0.214	0.834	0.950	0.980
RA-Depth	832×256	0.095	0.613	4.106	0.170	0.906	0.969	0.985
Monodepth2 [10]	1024×320	0.193	1.335	6.058	0.271	0.673	0.921	0.972
HR-Depth [27]	1024×320	0.183	1.241	6.054	0.267	0.687	0.920	0.972
DIFFNet [48]	1024×320	0.163	1.153	6.011	0.251	0.743	0.928	0.974
RA-Depth	1024×320	0.097	0.608	4.131	0.174	0.901	0.968	0.985

of our RA-Depth and Monodepth2 [10] for resolution adaptation in the depth estimation. It can be observed that the depth maps predicted by the single model Monodepth2 are inconsistent at different resolutions (see green rectangles). In contrast, our single model RA-Depth shows credible and consistent depth predictions at different resolutions, which demonstrates that RA-Depth can learn the scale-invariance of scene depth.

Generalization Results Across Datasets. In Table 3, we report the generalization performance on the Make3D [38] and NYU-V2 [39] datasets using models trained at the resolution of 640×192 on the KITTI dataset [8]. For a fair comparison, we adopt the same evaluation criteria, cropping approach [10] and data preprocessing strategy [45] for all models. When the test resolution is the same as the training resolution (640×192), we achieve state-of-the-art results whether generalizing to the outdoor dataset Make3D or the indoor dataset NYU-V2. Moreover, we also test our model at four different resolutions (416×128 , 512×160 , 832×256 , 1024×320) and take the average results as the resolution adaptation results across datasets. These results demonstrate that our model still has the good resolution adaptation ability when tested across datasets.

4.3 Ablation Study

To better comprehend how the components of our method contribute to the overall performance as well as the resolution adaptation ability, we perform sufficient ablation experiments in this section. As shown in Table 4, we perform

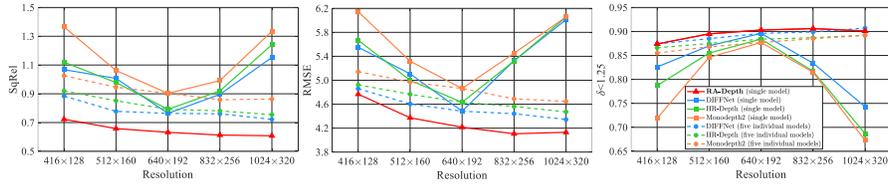


Fig. 3. Resolution adaptation results on the error metrics (SqRel, RMSE) and the accuracy metric ($\delta < 1.25$). Solid lines denote the resolution adaptation results tested at five different resolutions using a single model trained on 640×192 . Dotted lines denote results tested at five different resolutions using five individual models trained on each test resolution. Our single model RA-Depth outperforms other single models significantly on all metrics and is comparable to these individual models trained separately at each test resolution. This figure is best viewed in a zoomed-in document.

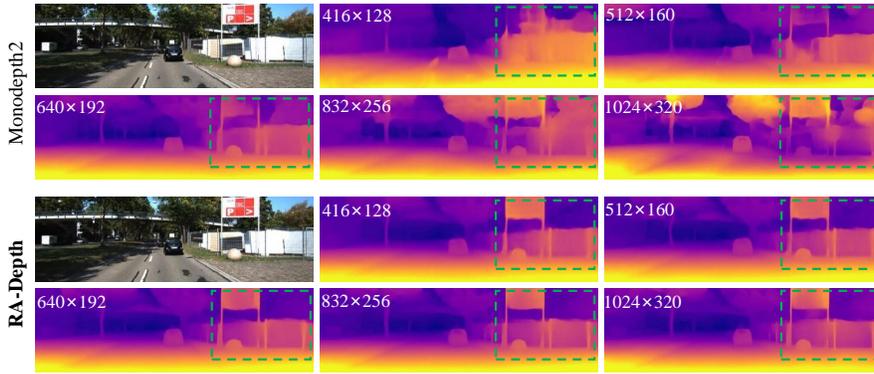


Fig. 4. Visualization results of resolution adaptation. We use models trained at the resolution of 640×192 to predict depths at various test resolutions.

ablation studies by changing various components of our method, where all models are trained at the resolution of 640×192 on the KITTI dataset [8].

We conduct ablation experiments to verify the effectiveness of each component on overall performance at the same resolution as training (640×192). Moreover, we also verify the effectiveness of each component on the resolution adaptation ability at four different resolutions (416×128 , 512×160 , 832×256 , 1024×320), where the average results are reported at four resolutions. For the baseline model, we use HRNet18 [40] as the encoder and use the same decoder as Monodepth2 [10]. The ablation results show that our proposed arbitrary-scale data augmentation (AS-Aug), Dual HRNet (D-HRNet), and cross-scale depth consistency loss (CS-Loss) can bring obvious improvements individually. Since the CS-Loss is imposed on the AS-Aug, experiments involving the CS-Loss must include the AS-Aug operation. All combined components can yield the best performance (RA-Depth). Besides, even without using ImageNet pre-trained weights, our model can also obtain the large gain over the baseline model.

Table 3. Generalization results on the Make3D [38] and NYU-V2 [39] datasets. All models are trained at the resolution of 640×192 on the KITTI dataset [8]. The upper part of this table reports the generalization results at the same resolution as training. The bottom half of this table reports the average results tested at four different resolutions (416×128 , 512×160 , 832×256 , 1024×320), which represent the results of resolution adaptation across datasets.

Dataset	Method	Error Metric ↓				Accuracy Metric ↑		
		AbsRel	SqRel	RMSE	RMSElog	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$
(a) Results At The Same Resolution As Training								
Make3D	Monodepth2 [10]	0.322	3.589	7.418	0.163	0.559	0.799	0.909
	HR-Depth [27]	0.315	3.208	7.024	0.159	0.562	0.798	0.913
	DIFFNet [48]	0.309	3.313	7.008	0.155	0.575	0.815	0.919
	RA-Depth	0.277	2.703	6.548	0.143	0.612	0.846	0.933
NYU-V2	Monodepth2 [10]	0.377	0.778	1.388	0.414	0.442	0.728	0.879
	HR-Depth [27]	0.321	0.521	1.150	0.367	0.490	0.774	0.913
	DIFFNet [48]	0.307	0.481	1.095	0.352	0.518	0.796	0.921
	RA-Depth	0.250	0.286	0.843	0.293	0.605	0.861	0.955
(b) Average Results Of Resolution Adaptation At Four Different Resolutions								
Make3D	Monodepth2 [10]	0.342	3.862	7.921	0.172	0.509	0.776	0.901
	HR-Depth [27]	0.340	3.527	7.577	0.171	0.502	0.773	0.903
	DIFFNet [48]	0.333	3.753	7.700	0.167	0.528	0.790	0.909
	RA-Depth	0.278	2.722	6.655	0.146	0.601	0.844	0.932
NYU-V2	Monodepth2 [10]	0.380	0.778	1.384	0.413	0.445	0.730	0.882
	HR-Depth [27]	0.325	0.520	1.147	0.377	0.491	0.778	0.915
	DIFFNet [48]	0.322	0.508	1.167	0.364	0.500	0.781	0.913
	RA-Depth	0.265	0.326	0.890	0.304	0.586	0.848	0.948

Effects of Arbitrary-Scale Data Augmentation. As shown in Table 4(a), our proposed arbitrary-scale data augmentation (AS-Aug) significantly improves performance when tested at the same resolution as training. Moreover, results in Table 4(b) demonstrate that AS-Aug brings large gains over the baseline in terms of resolution adaptation performance. In particular, when combining AS-Aug to the baseline model, we achieve the error reduction rates of 29.6%, 41.4%, and 22.2% in terms of AbsRel, SqRel, and RMSE, respectively, and the accuracy gain of 11.4% in terms of $\delta < 1.25$. As described in Algorithm 1, given an original image, AS-Aug generates three training images with three different scales to let the model implicitly learn the scale invariance of the scene depth. When generating only one training image from an original image, we compare our AS-Aug component with the data augmentation (dubbed Bian-Aug) proposed by Bian *et al.* [2] as shown in Table 5(a). Although Bian-Aug also uses randomly resizing and cropping to change the scale of images, the differences between us are two-fold. On the one hand, Bian-Aug first downsizes the original image and then upsamples the downsized image to high resolution for cropping. However, our AS-Aug directly crops on the original image so that the detailed information of the images can be remained. On the other hand, [2] only inputs one-scale image to the network, while our method inputs three-scale images to the network so that three kinds of camera intrinsics can be simulated. By imposing the cross-scale depth

Table 4. Ablation results for each component of our method. All models are trained at the resolution of 640×192 on the KITTI dataset [8] using the Eigen split [5]. AS-Aug: arbitrary-scale data augmentation. D-HRNet: Dual HRNet. CS-Loss: cross-scale depth consistency loss. pt: pretrained weights on the ImageNet dataset. (a) We test the effects of each component on overall performance at the same resolution as training. (b) We report the average results tested at four different resolutions (416×128 , 512×160 , 832×256 , 1024×320) to verify the effectiveness of each component on the resolution adaptation ability.

	Pretrained	AS-Aug	D-HRNet	CS-Loss	Error Metric ↓				Accuracy Metric ↑		
					AbsRel	SqRel	RMSE	RMSElog	$\delta < 1.25$	$\delta < 1.25^*$	$\delta < 1.25^*$
(a) Results At The Same Resolution As Training											
Baseline	✓				0.107	0.866	4.680	0.183	0.890	0.963	0.982
Baseline + AS-Aug	✓	✓			0.102	0.716	4.395	0.178	0.899	0.965	0.984
Baseline + D-HRNet	✓		✓		0.102	0.746	4.466	0.178	0.897	0.965	0.983
RA-Depth w/o D-HRNet	✓	✓		✓	0.097	0.651	4.246	0.172	0.902	0.967	0.984
RA-Depth w/o CS-Loss	✓	✓	✓		0.101	0.732	4.382	0.176	0.900	0.967	0.984
RA-Depth	✓	✓	✓	✓	0.096	0.632	4.216	0.171	0.903	0.968	0.985
Baseline w/o pt					0.122	0.957	4.955	0.199	0.860	0.954	0.979
RA-Depth w/o pt		✓	✓	✓	0.115	0.806	4.633	0.189	0.873	0.959	0.982
(b) Average Results Of Resolution Adaptation At Four Different Resolutions											
Baseline	✓				0.152	1.239	5.745	0.225	0.799	0.945	0.978
Baseline + AS-Aug	✓	✓			0.107	0.730	4.470	0.182	0.890	0.964	0.983
Baseline + D-HRNet	✓		✓		0.145	0.994	5.332	0.220	0.801	0.950	0.980
RA-Depth w/o D-HRNet	✓	✓		✓	0.102	0.664	4.375	0.177	0.893	0.966	0.985
RA-Depth w/o CS-Loss	✓	✓	✓		0.107	0.724	4.467	0.182	0.891	0.965	0.984
RA-Depth	✓	✓	✓	✓	0.101	0.651	4.345	0.177	0.894	0.966	0.985
Baseline w/o pt					0.209	2.179	7.285	0.290	0.678	0.890	0.955
RA-Depth w/o pt		✓	✓	✓	0.124	0.856	4.844	0.198	0.856	0.955	0.981

consistency loss among them, we can learn the scale invariance of different image resolutions. Therefore, the proposed AS-Aug significantly outperforms the Bian-Aug in terms of resolution adaptation of depth estimation. Table 5(b) reports how the number of images generated by AS-Aug affects the resolution adaptation performance. Since the performance is better when generating multiple images and the proposed cross-scale depth consistency loss needs to be implemented based on multiple images, we generate three training images for each input.

Effects of Cross-Scale Depth Consistency Loss. As shown in Table 4, whether testing at the same resolution as training or four different resolutions, the addition of cross-scale depth consistency loss (CS-Loss) brings significant improvements. Moreover, as shown in Table 5(c), for different settings of AS-Aug, combining CS-Loss can obviously improve the resolution adaptation performance of the model. These experiments demonstrate that the explicit constraint of CS-Loss can further help the model learn the scale invariance of scene depth, so as to enhance the resolution adaptation ability of the model. This verifies that our depth estimation model can explicitly learn the scale invariance of scene depth with our proposed cross-scale depth consistency loss.

Effects of Dual HRNet. Table 6 reports the parameter complexity (#Params) and computation complexity (GFLOPs) of our model RA-Depth. Compared with existing state-of-the-art models as well as the baseline model, our model is the most efficient in terms of parameter complexity and computation complexity.

Table 5. Additional ablation studies for AS-Aug and CS-Loss. Bian-Aug represents the data augmentation method of Bian *et al.* [2]. All models are trained at the resolution of 640×192 on the KITTI dataset [8]. We report the average results of resolution adaptation tested at four different resolutions (416×128 , 512×160 , 832×256 , 1024×320).

		Error Metric ↓				Accuracy Metric ↑		
		AbsRel	SqRel	RMSE	RMSElog	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$
(a)	Baseline	0.152	1.239	5.745	0.225	0.799	0.945	0.978
	Baseline + Bian-Aug (1 image)	0.128	0.847	4.770	0.200	0.850	0.960	0.981
	Baseline + AS-Aug (1 image)	0.108	0.742	4.491	0.184	0.885	0.961	0.982
(b)	Baseline + AS-Aug (1 image)	0.108	0.742	4.491	0.184	0.885	0.961	0.982
	Baseline + AS-Aug (2 images)	0.107	0.734	4.478	0.182	0.888	0.963	0.983
	Baseline + AS-Aug (3 images)	0.107	0.730	4.470	0.182	0.890	0.964	0.983
(c)	Baseline + AS-Aug (2 images) + CS-Loss	0.103	0.670	4.383	0.178	0.891	0.965	0.984
	Baseline + AS-Aug (3 images) + CS-Loss	0.102	0.664	4.375	0.177	0.893	0.966	0.985

Table 6. Parameters and GFLOPs of the depth estimation network.

Method	Encoder	#Params	GFLOPs	Error Metric ↓				Accuracy Metric ↑		
				AbsRel	SqRel	RMSE	RMSElog	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$
Monodepth2 [10]	ResNet50	34.57M	16.63	0.110	0.831	4.642	0.187	0.883	0.962	0.982
HR-Depth [27]	ResNet18	14.61M	11.27	0.109	0.792	4.632	0.185	0.884	0.962	0.983
DIFF-Net [48]	HRNet18	10.87M	15.77	0.102	0.764	4.483	0.180	0.896	0.965	0.983
Baseline	HRNet18	11.31M	10.99	0.107	0.866	4.680	0.183	0.890	0.963	0.982
RA-Depth	HRNet18	9.98M	10.78	0.096	0.632	4.216	0.171	0.903	0.968	0.985

That is, thanks to the efficient design of depth network Dual HRNet with multi-scale feature fusion, our model can achieve the best performance with the lowest network overhead.

5 Conclusions

We proposed RA-Depth, a self-supervised monocular depth estimation method to efficiently predict depth maps of images when the resolutions of images are inconsistent during training and testing. Based on our proposed arbitrary-scale data augmentation method, we developed a dual high-resolution network to efficiently fuse multi-scale features for the accurate depth prediction. RA-Depth achieves state-of-the-art results on various public datasets, and exhibits the good ability of resolution adaptation. Moreover, benefiting from the constructed dual HRNet, our RA-Depth model is efficient in terms of the parameter and computation complexity.

Acknowledgments

The authors would like to thank reviewers for their detailed comments and instructive suggestions. Thank Beibei Zhou and Kun Wang for their valuable suggestions and discussions. This work was supported by the National Science Fund of China (Grant Nos. U1713208, 61876084).

References

1. Bhat, S.F., Alhashim, I., Wonka, P.: AdaBins: Depth estimation using adaptive bins. In: CVPR (2021)
2. 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. In: NeurIPS (2019)
3. Chen, P.Y., Liu, A.H., Liu, Y.C., Wang, Y.C.F.: Towards scene understanding: Unsupervised monocular depth estimation with semantic-aware representation. In: CVPR (2019)
4. Choi, J., Jung, D., Lee, D., Kim, C.: Safenet: Self-supervised monocular depth estimation with semantic-aware feature extraction. arXiv preprint arXiv:2010.02893 (2020)
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)
6. Fu, H., Gong, M., Wang, C., Batmanghelich, K., Tao, D.: Deep ordinal regression network for monocular depth estimation. In: CVPR (2018)
7. Garg, R., Bg, V.K., Carneiro, G., Reid, I.: Unsupervised cnn for single view depth estimation: Geometry to the rescue. In: ECCV (2016)
8. Geiger, A., Lenz, P., Stiller, C., Urtasun, R.: Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research* (2013)
9. Godard, C., Mac Aodha, O., Brostow, G.J.: Unsupervised monocular depth estimation with left-right consistency. In: CVPR (2017)
10. Godard, C., Mac Aodha, O., Firman, M., Brostow, G.J.: Digging into self-supervised monocular depth estimation. In: ICCV (2019)
11. Guizilini, V., Ambrus, R., Pillai, S., Raventos, A., Gaidon, A.: 3D packing for self-supervised monocular depth estimation. In: CVPR (2020)
12. Guizilini, V., Hou, R., Li, J., Ambrus, R., Gaidon, A.: Semantically-guided representation learning for self-supervised monocular depth. arXiv preprint arXiv:2002.12319 (2020)
13. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016)
14. Hirschmuller, H.: Accurate and efficient stereo processing by semi-global matching and mutual information. In: CVPR (2005)
15. Hirschmuller, H.: Stereo processing by semiglobal matching and mutual information. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2007)
16. Hu, J., Ozay, M., Zhang, Y., Okatani, T.: Revisiting single image depth estimation: Toward higher resolution maps with accurate object boundaries. In: WACV (2019)
17. Jaderberg, M., Simonyan, K., Zisserman, A., Kavukcuoglu, K.: Spatial transformer networks. arXiv preprint arXiv:1506.02025 (2015)
18. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
19. Klingner, M., Termöhlen, J.A., Mikolajczyk, J., Fingscheidt, T.: Self-supervised monocular depth estimation: Solving the dynamic object problem by semantic guidance. In: ECCV (2020)
20. Ladicky, L., Shi, J., Pollefeys, M.: Pulling things out of perspective. In: CVPR (2014)
21. Laina, I., Ruppel, C., Belagiannis, V., Tombari, F., Navab, N.: Deeper depth prediction with fully convolutional residual networks. In: 3DV (2016)

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)
23. Leistner, T., Mackowiak, R., Ardizzone, L., Köthe, U., Rother, C.: Towards multimodal depth estimation from light fields. In: CVPR (2022)
24. Liu, F., Shen, C., Lin, G.: Deep convolutional neural fields for depth estimation from a single image. In: CVPR (2015)
25. Liu, F., Shen, C., Lin, G., Reid, I.: Learning depth from single monocular images using deep convolutional neural fields. IEEE Transactions on Pattern Analysis and Machine Intelligence (2015)
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. IEEE Transactions on Pattern Analysis and Machine Intelligence (2019)
27. Lyu, X., Liu, L., Wang, M., Kong, X., Liu, L., Liu, Y., Chen, X., Yuan, Y.: Hr-depth: High resolution self-supervised monocular depth estimation. AAAI (2020)
28. Mahjourian, R., Wicke, M., Angelova, A.: Unsupervised learning of depth and ego-motion from monocular video using 3D geometric constraints. In: CVPR (2018)
29. Parida, K.K., Srivastava, S., Sharma, G.: Beyond image to depth: Improving depth prediction using echoes. In: CVPR (2021)
30. 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)
31. Patil, V., Sakaridis, C., Liniger, A., Van Gool, L.: P3Depth: Monocular depth estimation with a piecewise planarity prior. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1610–1621 (2022)
32. Peng, R., Wang, R., Lai, Y., Tang, L., Cai, Y.: Excavating the potential capacity of self-supervised monocular depth estimation. In: ICCV (2021)
33. Poggi, M., Aleotti, F., Tosi, F., Mattocchia, S.: On the uncertainty of self-supervised monocular depth estimation. In: CVPR (2020)
34. Poggi, M., Tosi, F., Mattocchia, S.: Learning monocular depth estimation with unsupervised trinocular assumptions. In: 3DV (2018)
35. Qi, X., Liao, R., Liu, Z., Urtasun, R., Jia, J.: GeoNet: Geometric neural network for joint depth and surface normal estimation. In: CVPR (2018)
36. 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: CVPR (2019)
37. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. International journal of computer vision (2015)
38. Saxena, A., Sun, M., Ng, A.Y.: Make3D: Learning 3D scene structure from a single still image. IEEE transactions on pattern analysis and machine intelligence (2008)
39. Silberman, N., Hoiem, D., Kohli, P., Fergus, R.: Indoor segmentation and support inference from rgb-d images. In: ECCV (2012)
40. Wang, J., Sun, K., Cheng, T.: Deep high-resolution representation learning for visual recognition. IEEE Transactions on Pattern Analysis and Machine Intelligence (2019)
41. 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 (2004)
42. Watson, J., Firman, M., Brostow, G.J., Turmukhambetov, D.: Self-supervised monocular depth hints. In: ICCV (2019)

43. Xu, D., Wang, W., Tang, H., Liu, H., Sebe, N., Ricci, E.: Structured attention guided convolutional neural fields for monocular depth estimation. In: CVPR (2018)
44. Yin, W., Liu, Y., Shen, C., Yan, Y.: Enforcing geometric constraints of virtual normal for depth prediction. In: ICCV (2019)
45. Yu, Z., Jin, L., Gao, S.: P²net: Patch-match and plane-regularization for unsupervised indoor depth estimation. In: ECCV (2020)
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)
47. Zhang, Z., Cui, Z., Xu, C., Yan, Y., Sebe, N., Yang, J.: Pattern-affinitive propagation across depth, surface normal and semantic segmentation. In: CVPR (2019)
48. Zhou, H., Greenwood, D., Taylor, S.: Self-supervised monocular depth estimation with internal feature fusion. In: BMVC (2021)
49. Zhou, T., Brown, M., Snavely, N., Lowe, D.G.: Unsupervised learning of depth and ego-motion from video. In: CVPR (2017)
50. Zhu, S., Brazil, G., Liu, X.: The edge of depth: Explicit constraints between segmentation and depth. In: CVPR (2020)