

Depth Map Decomposition for Monocular Depth Estimation

Jinyoung Jun¹[0000-0002-8256-6846], Jae-Han Lee²[0000-0002-3674-4023],
Chul Lee³[0000-0001-9329-7365], and Chang-Su Kim¹[0000-0002-4276-1831]

¹ School of Electrical Engineering, Korea University, Seoul, Korea

² Gauss Labs Inc.

³ Department of Multimedia Engineering, Dongguk University, Seoul, Korea
jyjun@mcl.korea.ac.kr, jaehan.lee@gausslabs.ai,
chullee@dongguk.edu, changsukim@korea.ac.kr

Abstract. We propose a novel algorithm for monocular depth estimation that decomposes a metric depth map into a normalized depth map and scale features. The proposed network is composed of a shared encoder and three decoders, called G-Net, N-Net, and M-Net, which estimate gradient maps, a normalized depth map, and a metric depth map, respectively. M-Net learns to estimate metric depths more accurately using relative depth features extracted by G-Net and N-Net. The proposed algorithm has the advantage that it can use datasets without metric depth labels to improve the performance of metric depth estimation. Experimental results on various datasets demonstrate that the proposed algorithm not only provides competitive performance to state-of-the-art algorithms but also yields acceptable results even when only a small amount of metric depth data is available for its training.

Keywords: Monocular depth estimation, relative depth estimation, depth map decomposition

1 Introduction

Monocular depth estimation is a task to predict a pixel-wise depth map from a single image to understand the 3D geometry of a scene. The distance from a scene point to the camera provides essential information in various applications, including 2D-to-3D image/video conversion [52], augmented reality [35], autonomous driving [8], surveillance [22], and 3D CAD model generation [20]. Since only a single camera is available in many applications, monocular depth estimation, which infers the 3D information of a scene without additional equipment, has become an important research topic.

Recently, learning-based monocular depth estimators using convolutional neural networks (CNNs) have shown significant performance improvements, overcoming the intrinsic ill-posedness of monocular depth estimation by exploiting a huge amount of training data [1, 3, 6, 7, 14, 25, 28, 54, 56]. Existing learning-based monocular depth estimators can be classified into two categories according to



Fig. 1: Metric depth maps and their normalized depth maps.

the properties of estimated depth maps: relative depth estimation and metric depth estimation. Relative depth estimation predicts the relative depth order among pixels [2, 34, 51, 58]. Metric depth estimation, on the other hand, predicts the absolute distance of each scene point from the camera [3, 6, 14, 25, 54], which is a pixel-wise regression problem.

To estimate a metric depth map, a network should learn both the 3D geometry of the scene and the camera parameters. This implies that a metric depth estimator should be trained with a dataset obtained by a specific camera. In contrast, a relative depth estimator can be trained with heterogeneous datasets, *e.g.*, disparity maps from stereo image pairs or even manually labeled pixel pairs. Thus, relative depth estimation is an easier task than metric depth estimation is. Moreover, note that the geometry of a scene can be easily estimated when extra cues are available. For example, depth completion [38, 39, 55], which recovers a dense depth map from sparse depth measurements, can be performed more accurately and more reliably than monocular depth estimation is. Based on these observations, metric depth estimation algorithms using relative depths as extra cues have been developed via fitting [34, 43] or fine-tuning [42].

In this paper, we propose a monocular metric depth estimator that decomposes a metric depth map into a normalized depth map and scale features. As illustrated in Fig. 1, a normalized depth map contains relative depth information, and it is less sensitive to scale variations or camera parameters than a metric depth map is. The proposed algorithm consists of a single shared encoder and three decoders, G-Net, N-Net, and M-Net, which estimate gradient maps, a normalized depth map, and a metric depth map, respectively. M-Net learns to estimate metric depth maps using relative depth features extracted by G-Net and N-Net. To this end, we progressively transfer features from G-Net to N-Net and then from N-Net to M-Net. In addition, we develop the mean depth residual (MDR) block for M-Net to utilize N-Net features more effectively. Because the proposed algorithm learns to estimate metric depths by exploiting gradient maps and relative depths, additional datasets containing only relative depths can be used to improve the metric depth estimation performance further. Experimental results show that the proposed algorithm is competitive with state-of-the-art metric depth estimators, even when it is trained with a smaller metric depth dataset.

This paper has the following contributions:

- We propose a novel monocular depth estimator, which decomposes a metric depth map into a normalized depth map and relative depth features and then exploits those relative features to improve the metric depth estimation performance.
- The proposed algorithm can be adapted to a new camera efficiently since it can be trained with a small metric depth dataset together with camera-independent relative depth datasets.
- The proposed algorithm provides competitive performance to conventional state-of-the-art metric depth estimators and can improve the performance further through joint training using multiple datasets.

2 Related Work

2.1 Monocular Metric Depth Estimation

The objective of monocular metric depth estimation is to predict pixel-wise absolute distances of a scene from a camera using a single image. Since different 3D scenes can be projected onto the same 2D image, monocular depth estimation is ill-posed. Nevertheless, active research has been conducted due to its practical importance. To infer depths, early approaches made prior assumptions on scenes, *e.g.* box blocks [9], planar regions [44], or particular layout of objects [10]. However, they may provide implausible results, especially in regions with ambiguous colors or small objects.

With recent advances in deep learning, CNN techniques for monocular depth estimation have been developed, yielding excellent performance. Many attempts have been made to find better network architecture [3, 6, 14, 25, 54] or to design more effective loss functions [2, 5, 16, 25]. It has been also demonstrated that the depth estimation performance can be improved by predicting quantized depths through ordinal regression [7], by employing Fourier domain analysis [26], by enforcing geometric constraints of virtual normals [56], or by reweighting multiple loss functions [28]. Recently, the vision transformer [4] was employed for monocular depth estimation [1], improving the performance significantly.

2.2 Relative Depth Estimation

The objective of relative depth estimation is to learn the pairwise depth order [58] or the rank of pixels [2, 51] in an image. Recently, listwise ranking, instead of pairwise ranking, was considered for relative depth estimation [34]. Also, scale-invariant loss [6] and its variants [32, 33, 43, 48] have been used to alleviate the scale ambiguity of depths, thereby improving the performance of relative depth estimation.

Unlike metric depths, relative depth information — or depth order information — is invariant to camera parameters. Therefore, even though a training set

is composed of images captured by different cameras, it does not affect the performance of relative depth estimation adversely. Therefore, heterogeneous training data, such as disparity maps from stereo image pairs [48, 50, 51] or videos [43], structure-from-motion reconstruction [32, 33], and ordinal labels [2], have been used to train relative depth estimators.

2.3 Relative vs. Metric Depths

A metric depth map contains relative depth information, whereas relative depth information is not sufficient for reconstructing a metric depth map. However, relative-to-metric depth conversion has been attempted by fitting relative depths to metric depths [34, 43] or by fine-tuning a relative depth estimator for metric depth estimation [42].

On the other hand, relative and metric depths can be jointly estimated to exploit their correlation and to eventually improve the performance of metric depth estimation. To this end, ordinal labels are used with a ranking loss in [2]. Also, in [27], relative and metric depth maps at various scales are first estimated and then optimally combined to yield a final metric depth map.

The proposed algorithm also estimates relative depth information, in addition to metric depths, to improve the performance of metric depth estimation. However, differently from [2, 27], the proposed algorithm better exploits the correlation between relative and metric depths by decomposing a metric depth map. Furthermore, the proposed algorithm can provide promising results even with a small metric depth dataset by exploiting a relative depth dataset additionally.

3 Proposed Algorithm

Fig. 2 is an overview of the proposed algorithm, which consists of a shared encoder and three decoders — G-Net, N-Net, and M-Net. The shared encoder extracts common features that are fed into the three decoders. Then, G-Net predicts horizontal and vertical gradients of depths, while N-Net and M-Net estimate a normalized depth map and a metric depth map, respectively. Note that features extracted by G-Net are fed into N-Net to convey edge information, and those by N-Net are, in turn, fed into M-Net to provide relative depth features. Finally, via the MDR block, M-Net exploits the relative depth features to estimate a metric depth map more accurately.

3.1 Metric Depth Decomposition

Given an RGB image $I \in \mathbb{R}^{h \times w \times 3}$, the objective is to estimate a metric depth map $M \in \mathbb{R}^{h \times w}$. However, this is ill-posed because different scenes with different metric depths can be projected onto the same image. Moreover, scale features of depths are hard to estimate from the color information only since they also depend on the camera parameters. To address this issue, we decompose a metric

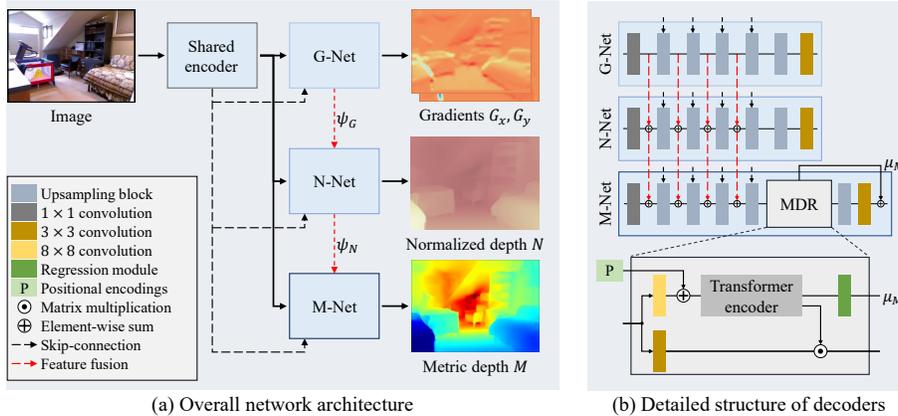


Fig. 2: (a) Overall network architecture of the proposed algorithm and (b) detailed structure of decoders. The proposed algorithm consists of a shared encoder and three decoders: G-Net, N-Net, and M-Net. G-Net predicts horizontal and vertical gradient maps, while N-Net and M-Net estimate normalized and metric depth maps, respectively. Note that G-Net features are fed into N-Net, and N-Net features are fed into M-Net.

depth map M into a normalized depth map N and scale parameters. The normalized depth map N contains relative depth information, so it is less sensitive to scale variations or camera parameters than the metric depth map M is.

There are several design choices for normalizing a metric depth map, including min-max normalization or ranking-based normalization [11]. However, the min-max normalization is sensitive to outliers, and the ranking-based normalization is unreliable in areas with homogeneous depths, such as walls and floors. Instead, we normalize a metric depth map using the z -score normalization. Given a metric depth map M , we obtain the normalized depth map N by

$$N = \frac{M - \mu_M U}{\sigma_M} \quad (1)$$

where μ_M and σ_M , respectively, denote the mean and standard deviation of metric depths in M . Also, U is the unit matrix whose all elements are 1.

N-Net, denoted by f_N , estimates this normalized depth map, and its estimate is denoted by \hat{N} . When the scale parameters μ_M and σ_M are known, the metric depth map M can be reconstructed by

$$\hat{M}_{\text{direct}} = \sigma_M \hat{N} + \mu_M U. \quad (2)$$

In practice, μ_M and σ_M are unknown. Conventional methods in [34, 42, 43] obtain fixed σ_M and μ_M for all images based on the least-squares criterion. In such a case, the accuracy of \hat{M}_{direct} in (2) greatly depends on the accuracy of σ_M and μ_M . In this work, instead of the direct conversion in (2), we estimate the

metric depth map by employing the features $\psi_N(I)$, which are extracted by the proposed M-Net, f_N , during the estimation of \hat{N} . In other words, the proposed M-Net, f_M , estimates the metric depth map by

$$\hat{M} = f_M(I, \psi_N(I)). \quad (3)$$

For metric depth estimation, structural data (*e.g.* surface normals or segmentation maps) have been adopted as additional cues [5, 29, 37, 40], or relative depth features have been used indirectly via loss functions (*e.g.* pairwise ranking loss [2] or scale-invariant loss [1, 6, 30]). In contrast, we convert a metric depth map to a normalized depth map. Then, the proposed N-Net estimates the normalized depth map to extract the features ψ_N , containing relative depth information. Then, the proposed M-Net uses ψ_N for effective metric depth estimation.

Similarly, we further decompose the normalized depth map N into more elementary data: horizontal and vertical gradients. The horizontal gradient map G_x is given by

$$G_x = \nabla_x N \quad (4)$$

where ∇_x is the partial derivative operator computing the differences between horizontally adjacent pixels. The vertical gradient map G_y is obtained similarly. The proposed G-Net is trained to estimate these gradient maps G_x and G_y . Hence, G-Net learns edge information in a scene, and its features ψ_G are useful for inferring the normalized depth map. Therefore, similar to (3), N-Net estimates the normalized depth map via

$$\hat{N} = f_N(I, \psi_G(I)) \quad (5)$$

using the gradient features $\psi_G(I)$.

3.2 Network Architecture

For the shared encoder in Fig. 2, we adopt EfficientNet-B5 [47] as the backbone network. G-Net and N-Net have an identical structure, consisting of five upsampling blocks. However, G-Net outputs two channels for two gradient maps G_x and G_y , while N-Net yields a single channel for a normalized depth map N . M-Net also has a similar structure, except for the MDR block, which will be detailed in Section 3.3. MDR predicts the mean μ_M of M separately, which is added back at the end of M-Net.

The encoder features are fed into the three decoders via skip-connections [13], as shown in Fig. 2(a). To differentiate the encoder features for the different decoders, we apply 1×1 convolution to the encoder features before feeding them to each decoder. Also, we apply the channel attention [15] before each skip-connection to each decoder.

We transfer features unidirectionally from G-Net to N-Net and also from N-Net to M-Net to exploit low-level features for the estimation of high-level data. To this end, we fuse features through element-wise addition before each of the first four upsampling blocks in N-Net and M-Net, as shown in Fig. 2(b).

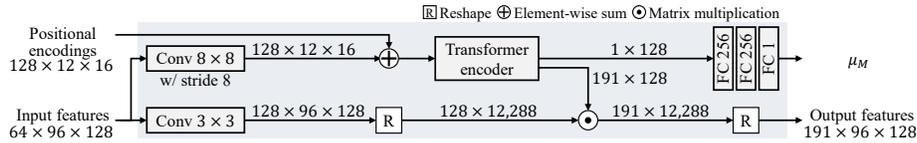


Fig. 3: Detailed structure of the MDR block.

Specifically, let ψ_G^{out} and ψ_N^{out} denote the output features of G-Net and N-Net, respectively. Then, the input feature ψ_N^{in} to the next layer of N-Net is given by

$$\psi_N^{\text{in}} = \omega_G \psi_G^{\text{out}} + \omega_N \psi_N^{\text{out}} \quad (6)$$

where ω_G and ω_N are pre-defined weights for ψ_G^{out} and ψ_N^{out} to control the relative contributions of the two features. For M-Net, the features from N-Net are added similarly. In order to fuse features, we use addition, instead of multiplication or concatenation, for computational efficiency.

3.3 MDR Block

We design the MDR block to utilize the features ψ_N of N-Net more effectively for the metric depth estimation in M-Net. Fig. 3 shows the structure of the MDR block, which applies patchwise attention to an input feature map and estimates the mean μ_M of M separately using the transformer encoder [4]. Note that the transformer architecture enables us to use one of the patchwise-attended feature vectors to regress μ_M .

Specifically, MDR first processes input features using an 8×8 convolution layer with a stride of 8 and a 3×3 convolution layer, respectively. The patchwise output of the 8×8 convolution is added to the positional encodings and then input to the transformer encoder [4]. The positional encodings are learnable parameters, randomly initialized at training. Then, the transformer encoder generates 192 patchwise-attended feature vectors of 128 dimensions. We adopt the mini-ViT architecture [1] for the transformer encoder. The first vector is fed to the regression module, composed of three fully-connected layers, to yield μ_M . The rest 191 vectors form a matrix, which is multiplied with the output of the 3×3 convolution layer to generate $191 \times 96 \times 128$ output features through reshaping. Finally, those output features are fed to the next upsampling block of M-Net. Also, the estimated μ_M is added back at the end of M-Net in Fig. 2(b), which makes the remaining parts of M-Net focus on the estimation of the mean-removed depth map $M - \mu_M U$ by exploiting the N-Net features ψ_N .

3.4 Loss Functions

Let us describe the loss functions for training the three decoders. For G-Net, we use the ℓ_1 loss

$$\mathcal{L}_G = \frac{1}{T} (\|\hat{G}_x - G_x\|_1 + \|\hat{G}_y - G_y\|_1) \quad (7)$$

where \hat{G}_x and \hat{G}_y are predictions of the ground-truth gradient maps G_x and G_y , respectively. Also, T denotes the number of valid pixels in the ground-truth.

For N-Net, we use two loss terms: the ℓ_1 loss and the gradient loss. The ℓ_1 loss is defined as

$$\mathcal{L}_N = \frac{1}{T} \|\hat{N} - N\|_1 \quad (8)$$

where \hat{N} and N are predicted and ground-truth normalized depth maps. Note that scale-invariant terms are often adopted to train monocular depth estimators [1, 30, 42]. However, we do not use such scale-invariant losses since normalized depth maps are already scale-invariant. Next, the gradient losses [16, 33, 51] for N in the horizontal direction are defined as

$$\mathcal{L}_{Nx} = \frac{1}{T \cdot s^2} \|\nabla_x \hat{N}_s - \nabla_x N_s\|_1 \quad (9)$$

where \hat{N}_s and N_s are the bilinearly scaled \hat{N} and N with a scaling factor s . We compute the gradient losses at various scales, as in [33, 51], by setting s to 0.5, 0.25, and 0.125. The losses L_{Ny} in the vertical direction are also used.

Similarly, for M-Net, we use the loss terms \mathcal{L}_M , \mathcal{L}_{Mx} , and \mathcal{L}_{My} . In addition, we use two more loss terms. First, \mathcal{L}_{μ_M} is defined to train the MDR block, which is given by

$$\mathcal{L}_{\mu_M} = \|\mu(\hat{M}) - \mu_M\|_1 \quad (10)$$

where $\mu(\hat{M})$ denotes the mean of depth values in \hat{M} . Second, we define the logarithmic ℓ_1 loss,

$$\mathcal{L}_{\log M} = \frac{1}{T} \|\log \hat{M} - \log M\|_1. \quad (11)$$

In this work, we adopt inverse depth representation of metric depths to match the depth order with a relative depth dataset [51]. In this case, theoretically, a metric depth can have a value in the range $[0, \infty)$. Thus, when a metric depth is near zero, its inverted value becomes too large, which interferes with training. We overcome this problem through a simple modification. Given an original metric depth m_o , its inverted metric depth m is defined as

$$m = 1/(m_o + 1). \quad (12)$$

In this way, inverted metric depth values are within the range of $(0, 1]$ and also are more evenly distributed.

However, using the ℓ_1 loss \mathcal{L}_M on inverse depths has a disadvantage in learning distant depths. Suppose that $\hat{\chi}$ and χ are predicted and ground-truth metric depth values for a pixel, respectively. Then, the ℓ_1 error E is given by

$$E = \left| \frac{1}{\hat{\chi}} - \frac{1}{\chi} \right|. \quad (13)$$

As χ gets larger, E becomes smaller for the same $|\hat{\chi} - \chi|$. This means that the network is trained less effectively for distant regions. This problem is alleviated by employing $\mathcal{L}_{\log M}$ in (11).

4 Experimental Results

4.1 Datasets

We use four depth datasets: one for relative depths [51] and three for metric depths [23, 45, 46]. When relative depth data are used in training, losses are generated from the loss terms for N-Net and G-Net only because the loss terms for M-Net cannot be computed.

HR-WSI [51]: It consists of 20,378 training and 400 test images. The ground-truth disparity maps are generated by FlowNet 2.0 [19]. We use only the training data of HR-WSI. We normalize the disparity maps by (1) and regard them as normalized depth maps.

NYUv2 [45]: It contains 120K video frames for training and 654 frames for test, together with the depth maps captured by a Kinect v1 camera. We use the NYUv2 dataset for both training and evaluation. We construct three training datasets of 51K, 17K, and 795 sizes. Specifically, we extract the 51K and 17K images by sampling video frames uniformly. For the 795 images, we use the official training split of NYUv2. We fill in missing depths using the colorization scheme [31], as in [45].

DIML-Indoor [23]: It consists of 1,609 training images and 503 test images, captured by a Kinect v2 camera.

SUN RGB-D [46]: It consists 5,285 training images and 5,050 test images, obtained by four different cameras: Kinect v1, Kinect v2, RealSense, and Xtion.

4.2 Evaluation Metrics

Metric depths: We adopt the four evaluation metrics in [6], listed below. Here, M_i and \hat{M}_i denote the ground-truth and predicted depths of pixel i , respectively. $|\cdot|$ denotes the number of valid pixels in a depth map. For the NYUv2 dataset, we adopt the center crop protocol [6].

$$\begin{aligned} \text{RMSE} &: \frac{1}{|M|} \left(\sum_i (\hat{M}_i - M_i)^2 \right)^{0.5} \\ \text{REL} &: \frac{1}{|M|} \sum_i |\hat{M}_i - M_i| / M_i \\ \log 10 &: \frac{1}{|M|} \sum_i |\log_{10}(\hat{M}_i) - \log_{10}(M_i)| \\ \delta_k &: \% \text{ of } M_i \text{ that satisfies } \max \left(\frac{\hat{M}_i}{M_i}, \frac{M_i}{\hat{M}_i} \right) < 1.25^k, \quad k \in \{1, 2, 3\} \end{aligned}$$

Relative depths: we use two metrics for relative depths. First, WHDR (weighted human disagreement rate) [2, 43, 51] measures the ordinal consistency between point pairs. We follow the evaluation protocol of [51] to randomly sample 50,000 pairs in each depth map. However, WHDR is an unstable protocol, under which the performance fluctuates with each measurement. We hence use Kendall’s

Table 1: Comparison of depth estimation results on the NYUv2 dataset. ‘#’ is the number of training images, and † means that additional data is used for training. The best results are **boldfaced**. Lower RMSE, REL, and log 10 indicate better results, while higher δ_k values are better ones.

	#	Encoder backbone	RMSE	REL	log 10	δ_1	δ_2	δ_3
Eigen <i>et al.</i> [6]	120K	-	0.641	0.158	-	0.769	0.950	0.988
Laina <i>et al.</i> [25]	12K	ResNet-50 [13]	0.573	0.127	0.055	0.811	0.953	0.988
Hao <i>et al.</i> [12]	13K	ResNet-101 [13]	0.555	0.127	0.053	0.841	0.966	0.991
Fu <i>et al.</i> [7]	120K	ResNet-101 [13]	0.509	0.115	0.051	0.828	0.965	0.992
Hu <i>et al.</i> [16]	50K	SENet-154 [15]	0.530	0.115	0.050	0.866	0.975	0.993
Chen <i>et al.</i> [3]	50K	SENet-154 [15]	0.514	0.111	0.048	0.878	0.977	0.994
Yin <i>et al.</i> [56]	29K	ResNeXt-101 [53]	0.416	0.108	0.048	0.875	0.976	0.994
Lee <i>et al.</i> [30]	24K	DenseNet-161 [17]	0.392	0.110	0.047	0.885	0.978	0.994
Hyunh <i>et al.</i> [18]	50K	DRN-D-22 [57]	0.412	0.108	-	0.882	0.980	0.996
Lee and Kim [28]	58K	PNASNet-5 [36]	0.430	0.119	0.050	0.870	0.974	0.993
Bhat <i>et al.</i> [1]	50K	EfficientNet-B5 [47]	0.364	0.103	0.044	0.903	0.984	0.997
Proposed	51K	EfficientNet-B5 [47]	0.362	0.100	0.043	0.907	0.986	0.997
Wang <i>et al.</i> [49] [†]	200K	-	0.745	0.220	0.094	0.605	0.890	0.970
Ramam. and Lepetit [41] [†]	400K	ResNet-50 [13]	0.502	0.139	0.047	0.836	0.966	0.993
Ranftl <i>et al.</i> [42] [†]	1.4M	ViT-Hybrid [4]	0.357	0.110	0.045	0.904	0.988	0.998
Proposed [†]	71K	EfficientNet-B5 [47]	0.355	0.098	0.042	0.913	0.987	0.998

τ [21] additionally, which considers the ordering relations of all pixel pairs. Given a ground-truth normalized depth map D and its prediction \hat{D} , Kendall’s τ is defined as

$$\tau(\hat{D}, D) = \frac{\alpha(\hat{D}, D) - \beta(\hat{D}, D)}{\binom{|D|}{2}} \quad (14)$$

where $\alpha(\hat{D}, D)$ and $\beta(\hat{D}, D)$ are the numbers of concordant pairs and discordant pairs between D and \hat{D} , respectively. Note that Kendall’s τ can measure the quality of a metric depth map, as well as that of a relative one.

4.3 Implementation Details

Network architecture: We employ EfficientNet-B5 [47] as the encoder backbone. The encoder takes an 512×384 RGB image and generates a 16×12 feature with 2,048 channels. The output feature is used as the input to the three decoders. G-Net and N-Net consist of 5 upsampling blocks, each of which is composed of a bilinear interpolation layer and two 3×3 convolution layers with the ReLU activation. Also, in addition to the 5 upsampling blocks, M-Net includes the MDR block, located between the fourth and fifth upsampling blocks. For feature fusion in (6), $\omega_G = \omega_N = 1$.

Training: We train the proposed algorithm in two phases. First, we train the network, after removing M-Net, for 20 epochs with an initial learning rate of 10^{-4} . The learning rate is decreased by a factor of 0.1 at every fifth epoch. Second, we train the entire network, including all three decoders, jointly for 15 epochs with an initial learning rate of 10^{-4} , which is decreased by a factor of 0.1

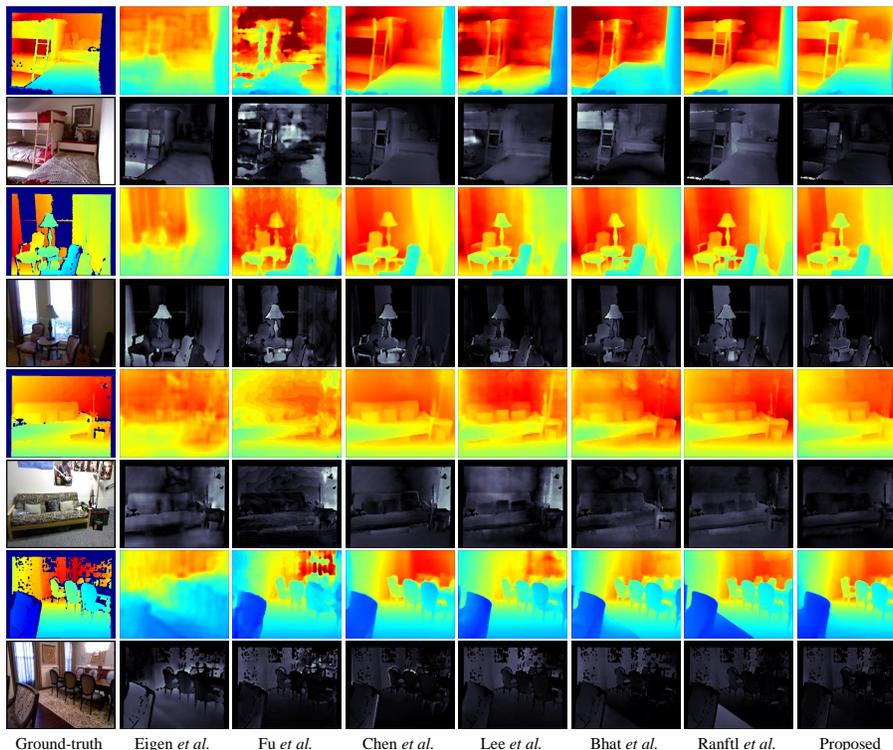


Fig. 4: Qualitative comparison of the proposed algorithm with the conventional algorithms. For each depth map, the error map is also provided, in which brighter pixels correspond to larger errors.

at every third epoch. We use the Adam optimizer [24] with a weight decay of 10^{-4} . If a relative depth is used in the second phase, losses are calculated from the loss terms for N-Net and G-Net only.

4.4 Performance Comparison

Table 1 compares the proposed algorithm with conventional ones on NYUv2 dataset. Some of the conventional algorithms use only NYUv2 training data [1, 3, 6, 7, 12, 15, 18, 25, 28, 30, 56], while the others use extra data [41, 42, 49]. For fair comparisons, we train the proposed algorithm in both ways: ‘Proposed’ uses NYUv2 only, while ‘Proposed[†]’ uses both HR-WSI and NYUv2. The following observations can be made from Table 1.

- ‘Proposed’ outperforms all conventional algorithms in all metrics with no exception. For example, ‘Proposed’ provides a REL score of 0.100, which is 0.003 better than that of the second-best algorithm, Bhat *et al.* [1]. Note that both algorithms use the same encoder backbone of EfficientNet-B5 [47].

Table 2: Comparison of depth estimation results on various datasets. ‘#’ is the size of metric depth dataset, and † means that 20K HR-WSI data are additionally used for training. A lower Kendall’s τ indicates a better result.

	#	Setting	RMSE	REL	log 10	δ_1	δ_2	δ_3	Kendall’s τ
NYUv2	795	Baseline	0.487	0.147	0.061	0.809	0.963	0.991	0.738
		Proposed	0.468	0.142	0.059	0.824	0.969	0.992	0.762
		Proposed†	0.417	0.122	0.052	0.868	0.977	0.995	0.800
	17K	Baseline	0.400	0.113	0.048	0.880	0.981	0.996	0.803
		Proposed	0.370	0.103	0.045	0.903	0.986	0.997	0.829
		Proposed†	0.362	0.100	0.043	0.909	0.987	0.997	0.835
	51K	Baseline	0.386	0.109	0.047	0.888	0.980	0.995	0.813
		Proposed	0.362	0.100	0.043	0.907	0.986	0.997	0.837
		Proposed†	0.355	0.098	0.042	0.913	0.987	0.998	0.840
DIML-Indoor	1.6K	Baseline	0.589	0.247	0.099	0.701	0.879	0.968	0.492
		Proposed	0.537	0.180	0.075	0.719	0.943	0.986	0.696
		Proposed†	0.517	0.171	0.072	0.742	0.949	0.989	0.742
SUN RGB-D	5.3K	Baseline	0.306	0.132	0.055	0.847	0.971	0.992	0.761
		Proposed	0.303	0.129	0.055	0.850	0.973	0.993	0.776
		Proposed†	0.301	0.127	0.054	0.853	0.973	0.992	0.784

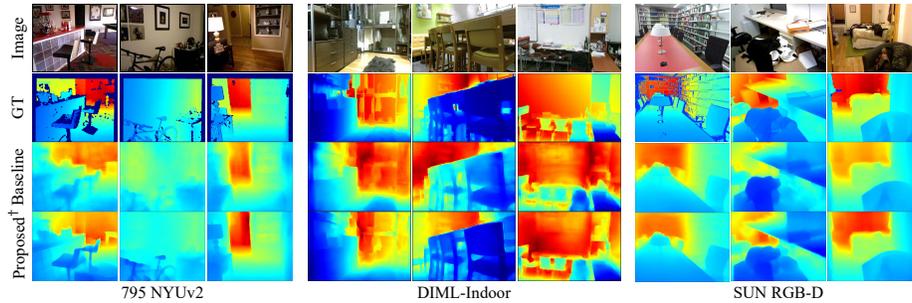


Fig. 5: Qualitative comparison of the proposed algorithm with the baseline using the NYUv2 (795), DIML-Indoor, and SUN RGB-D datasets.

- ‘Proposed†’ provides the best results in five out of six metrics. For δ_2 , the proposed algorithm yields the second-best score after Ranftl *et al.* [42]. It is worth pointing out that Ranftl *et al.* uses about 20 times more training data than the proposed algorithm does.

Fig. 4 compares the proposed algorithm with the conventional algorithms [1, 3, 6, 7, 30, 42] qualitatively. We see that the proposed algorithm estimates the depth maps more faithfully with smaller errors.

4.5 Various Datasets

Table 2 verifies the effectiveness of the proposed algorithm on various datasets. The first two columns in Table 2 indicate the metric depth dataset and its size.

Table 3: Ablation studies of the proposed algorithm using the NYUv2 (17K) dataset.

M	N	G	MDR*	MDR	†	RMSE	REL	log 10	δ_1	δ_2	δ_3	Kendall's τ	WHDR(%)
✓	-	-	-	-	-	0.400	0.113	0.048	0.880	0.981	0.996	0.803	14.95
✓	✓	-	-	-	-	0.389	0.111	0.047	0.888	0.982	0.996	0.814	14.19
✓	✓	✓	-	-	-	0.387	0.109	0.047	0.888	0.982	0.997	0.817	14.01
✓	✓	✓	✓	-	-	0.381	0.108	0.046	0.894	0.984	0.997	0.824	13.54
✓	✓	✓	-	✓	-	0.370	0.103	0.045	0.903	0.986	0.997	0.829	13.18
✓	✓	✓	-	✓	✓	0.362	0.100	0.043	0.909	0.987	0.997	0.835	12.72

‘Baseline’ is a bare encoder-decoder for monocular depth estimation. Specifically, we remove G-Net and N-Net, as well as the MDR block in M-Net, from the proposed algorithm in Fig. 2 to construct ‘Baseline.’ For its training, only three loss terms \mathcal{L}_M , \mathcal{L}_{Mx} , and \mathcal{L}_{My} are used. ‘Proposed’ means the proposed algorithm without employing the 20K HR-WSI training data, while ‘Proposed[†]’ means using the HR-WSI data additionally. The following observations can be made from Table 2.

- By comparing ‘Proposed’ with ‘Baseline,’ we see that G-Net and N-Net help M-Net improve the performance of metric depth estimation by transferring edge information and relative depth information. Also, ‘Proposed[†]’ meaningfully outperforms ‘Proposed’ by leveraging relative depth training data in HR-WSI, which contain no metric depth labels.
- Even when only the 795 NYUv2 images are used, the proposed algorithm provides acceptable results. For example, the RMSE score of 0.417 is similar to that of the Hyunh *et al.*’s estimator [18] in Table 1, which uses 50K metric depth map data. In contrast, the proposed algorithm uses the 795 metric depth maps only.
- The proposed algorithm also exhibits similar trends in the DIML-Indoor and SUN RGB-D datasets, which are collected using different cameras: the proposed algorithm can be trained effectively even with a small number of metric depth images. This is advantageous in practical applications in which an algorithm should be adapted for various cameras.

Fig. 5 compares ‘Baseline’ and ‘Proposed[†]’ qualitatively using the 795 NYUv2, DIML-Indoor, and SUN RGB-D datasets. For all datasets, ‘Proposed[†]’ provides more accurate and more detailed depth maps, especially around chairs, tables, and desks, than ‘Baseline’ does.

4.6 Analysis

Ablation studies: We conduct ablation studies that add the proposed components one by one in Table 3. Here, the 17K images from NYUv2 are used for training. M, N, and G denote the three decoders. MDR* is the MDR block with μ_M deactivated. † indicates the use of relative depth data. We see that all

Table 4: Effectiveness of the two-phase training scheme.

Setting	RMSE	REL	log 10	δ_1	δ_2	δ_3	Kendall’s τ	WHDR(%)
Single-phase	0.386	0.107	0.046	0.892	0.984	0.997	0.820	13.74
Proposed	0.362	0.100	0.043	0.909	0.987	0.997	0.835	12.72

Table 5: Complexity comparison.

	Ranftl <i>et al.</i> [42]	Proposed
# Params	130M	102M
Speed (fps)	13.4	34.7

Table 6: Complexity of each component.

	Encoder	G	N	M	MDR
# Params	55M	15M	15M	15M	1.7M
Speed (fps)	50.9	475	534	474	447

components lead to performance improvements, especially in terms of the two relative depth metrics Kendall’s τ and WHDR.

Table 4 shows the effectiveness of the two-phase training scheme of the proposed algorithm. The proposed algorithm, which trains G-Net and N-Net first, shows better results than the single-phase scheme, which trains the entire network at once.

Complexities and inference speeds: Table 5 compares the complexities of the proposed algorithm and the Ranftl *et al.*’s algorithm [42]. The proposed algorithm performs faster with a smaller number of parameters than the Ranftl *et al.*’s algorithm [42] does. This indicates that the performance gain of the proposed algorithm is not from the increase in complexity but from the effective use of relative depth features. Table 6 lists the complexity of each component of the proposed algorithm. The encoder spends most of the inference time, while the three decoders are relatively fast.

5 Conclusions

We proposed a monocular depth estimator that decomposes a metric depth map into a normalized depth map and scale features. The proposed algorithm is composed of a shared encoder with three decoders, G-Net, N-Net, and M-Net, which estimate gradient maps, a normalized depth map, and a metric depth map, respectively. G-Net features are used in N-Net, and N-Net features are used in M-Net. Moreover, we developed the MDR block for M-Net to utilize N-Net features and improve the metric depth estimation performance. Extensive experiments demonstrated that the proposed algorithm provides competitive performance and yields acceptable results even with a small metric depth dataset.

Acknowledgements

This work was supported by the National Research Foundation of Korea (NRF) grants funded by the Korea government (MSIT) (No. NRF-2021R1A4A1031864 and No. NRF-2022R1A2B5B03002310).

References

1. Bhat, S.F., Alhashim, I., Wonka, P.: AdaBins: Depth estimation using adaptive bins. In: CVPR. pp. 4009–4018 (2021) [1](#), [3](#), [6](#), [7](#), [8](#), [10](#), [11](#), [12](#)
2. Chen, W., Fu, Z., Yang, D., Deng, J.: Single-image depth perception in the wild. In: NIPS. pp. 730–738 (Dec 2016) [2](#), [3](#), [4](#), [6](#), [9](#)
3. Chen, X., Chen, X., Zha, Z.J.: Structure-aware residual pyramid network for monocular depth estimation. In: IJCAI. pp. 694–700 (Aug 2019) [1](#), [2](#), [3](#), [10](#), [11](#), [12](#)
4. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houshy, N.: An image is worth 16x16 words: Transformers for image recognition at scale. In: ICLR (2021) [3](#), [7](#), [10](#)
5. Eigen, D., Fergus, R.: Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In: ICCV. pp. 2650–2658 (Dec 2015) [3](#), [6](#)
6. Eigen, D., Puhrsch, C., Fergus, R.: Depth map prediction from a single image using a multi-scale deep network. In: NIPS. pp. 2366–2374 (2014) [1](#), [2](#), [3](#), [6](#), [9](#), [10](#), [11](#), [12](#)
7. Fu, H., Gong, M., Wang, C., Batmanghelich, K., Tao, D.: Deep ordinal regression network for monocular depth estimation. In: CVPR. pp. 2002–2011 (2018) [1](#), [3](#), [10](#), [11](#), [12](#)
8. Godard, C., Aodha, O.M., Brostow, G.J.: Unsupervised monocular depth estimation with left-right consistency. In: CVPR. pp. 270–279 (2017) [1](#)
9. Gupta, A., Efros, A.A., Hebert, M.: Blocks world revisited: Image understanding using qualitative geometry and mechanics. In: ECCV. pp. 482–496 (2010) [3](#)
10. Gupta, A., Hebert, M., Kanade, T., Blei, D.: Estimating spatial layout of rooms using volumetric reasoning about objects and surfaces. In: NIPS (2010) [3](#)
11. Han, J., Pei, J., Kamber, M.: Data mining: concepts and techniques. Elsevier (2011) [5](#)
12. Hao, Z., Li, Y., You, S., Lu, F.: Detail preserving depth estimation from a single image using attention guided networks. In: 3DV. pp. 304–313 (2018) [10](#), [11](#)
13. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR. pp. 770–778 (2016) [6](#), [10](#)
14. Heo, M., Lee, J., Kim, K.R., Kim, H.U., Kim, C.S.: Monocular depth estimation using whole strip masking and reliability-based refinement. In: ECCV. pp. 36–51 (Sep 2018) [1](#), [2](#), [3](#)
15. Hu, J., Shen, L., Sun, G.: Squeeze-and-excitation networks. In: CVPR. pp. 7132–7141 (2018) [6](#), [10](#), [11](#)
16. Hu, J., Ozay, M., Zhang, Y., Okatani, T.: Revisiting single image depth estimation: Toward higher resolution maps with accurate object boundaries. In: WACV. pp. 1043–1051 (Jan 2019) [3](#), [8](#), [10](#)
17. Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: CVPR. pp. 4700–4708 (2017) [10](#)
18. Huynh, L., Nguyen-Ha, P., Matas, J., Rahtu, E., Heikkilä, J.: Guiding monocular depth estimation using depth-attention volume. In: ECCV. pp. 581–597 (2020) [10](#), [11](#), [13](#)
19. Ilg, E., Mayer, N., Saikia, T., Keuper, M., Dosovitskiy, A., Brox, T.: FlowNet 2.0: Evolution of optical flow estimation with deep networks. In: CVPR. pp. 2462–2470 (2017) [9](#)
20. Izadinia, H., Shan, Q., Seitz, S.M.: IM2CAD. In: CVPR. pp. 5134–5143 (2017) [1](#)

21. Kendall, M.G.: A new measure of rank correlation. *Biometrika* **30**(1/2), 81–93 (1938) [9](#)
22. Kim, H., Lee, S., Kim, Y., Lee, S., Lee, D., Ju, J., Myung, H.: Weighted joint-based human behavior recognition algorithm using only depth information for low-cost intelligent video-surveillance system. *Expert Syst. Appl.* **45**, 131–141 (2016) [1](#)
23. Kim, Y., Jung, H., Min, D., Sohn, K.: Deep monocular depth estimation via integration of global and local predictions. *IEEE Trans. Image Process.* **27**(8), 4131–4144 (2018) [9](#)
24. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: *ICLR* (2015) [11](#)
25. Laina, I., Rupprecht, C., Belagiannis, V., Tombari, F., Navab, N.: Deeper depth prediction with fully convolutional residual networks. In: *3DV*. pp. 239–248 (Oct 2016) [1](#), [2](#), [3](#), [10](#), [11](#)
26. Lee, J.H., Heo, M., Kim, K.R., Kim, C.S.: Single-image depth estimation based on Fourier domain analysis. In: *CVPR*. pp. 330–339 (Jun 2018) [3](#)
27. Lee, J.H., Kim, C.S.: Monocular depth estimation using relative depth maps. In: *CVPR*. pp. 9729–9738 (Jun 2019) [4](#)
28. Lee, J.H., Kim, C.S.: Multi-loss rebalancing algorithm for monocular depth estimation. In: *ECCV*. pp. 785–801 (Aug 2020) [1](#), [3](#), [10](#), [11](#)
29. Lee, J.H., Lee, C., Kim, C.S.: Learning multiple pixelwise tasks based on loss scale balancing. In: *ICCV*. pp. 5107–5116 (2021) [6](#)
30. 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) [6](#), [8](#), [10](#), [11](#), [12](#)
31. Levin, A., Lischinski, D., Weiss, Y.: Colorization using optimization. *ACM Trans. Graph.* **23**(3), 689–694 (Aug 2004) [9](#)
32. Li, Z., Dekel, T., Cole, F., Tucker, R., Snavely, N., Liu, C., Freeman, W.T.: Learning the depths of moving people by watching frozen people. In: *CVPR*. pp. 4521–4530 (2019) [3](#), [4](#)
33. Li, Z., Snavely, N.: MegaDepth: Learning single-view depth prediction from internet photos. In: *CVPR*. pp. 2041–2050 (2018) [3](#), [4](#), [8](#)
34. Lienen, J., Hullermeier, E., Ewerth, R., Nommensen, N.: Monocular depth estimation via listwise ranking using the Plackett-Luce model. In: *CVPR*. pp. 14595–14604 (2021) [2](#), [3](#), [4](#), [5](#)
35. Liu, C., Yang, J., Ceylan, D., Yumer, E., Furukawa, Y.: PlaneNet: Piece-wise planar reconstruction from a single RGB image. In: *CVPR*. pp. 2579–2588 (2018) [1](#)
36. Liu, C., Zoph, B., Neumann, M., Shlens, J., Hua, W., Li, L.J., Fei-Fei, L., Yuille, A., Huang, J., Murphy, K.: Progressive neural architecture search. In: *ECCV*. pp. 19–34 (2018) [10](#)
37. Liu, S., Johns, E., Davison, A.J.: End-to-end multi-task learning with attention. In: *CVPR*. pp. 1871–1880 (2019) [6](#)
38. Ma, F., Karaman, S.: Sparse-to-dense: Depth prediction from sparse depth samples and a single image. In: *ICRA*. pp. 4796–4803 (2018) [2](#)
39. Park, J., Joo, K., Hu, Z., Liu, C.K., So Kweon, I.: Non-local spatial propagation network for depth completion. In: *ECCV*. pp. 120–136 (2020) [2](#)
40. Qi, X., Liao, R., Liu, Z., Urtasun, R., Jia, J.: GeoNet: Geometric neural network for joint depth and surface normal estimation. In: *CVPR*. pp. 283–291 (2018) [6](#)
41. Ramamonjisoa, M., Lepetit, V.: SharpNet: Fast and accurate recovery of occluding contours in monocular depth estimation. In: *ICCVW* (2019) [10](#), [11](#)
42. Ranftl, R., Bochkovskiy, A., Koltun, V.: Vision transformers for dense prediction. In: *ICCV*. pp. 12179–12188 (2021) [2](#), [4](#), [5](#), [8](#), [10](#), [11](#), [12](#), [14](#)

43. Ranftl, R., Lasinger, K., Hafner, D., Schindler, K., Koltun, V.: Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE Trans. Pattern Anal. Mach. Intell.* (2020) [2](#), [3](#), [4](#), [5](#), [9](#)
44. Saxena, A., Sun, M., Ng, A.Y.: Make3D: Learning 3D scene structure from a single still image. *IEEE Trans. Pattern Anal. Mach. Intell.* **31**(5), 824–840 (May 2008) [3](#)
45. Silberman, N., Hoiem, D., Kohli, P., Fergus, R.: Indoor segmentation and support inference from RGBD images. In: *ECCV*. pp. 746–760 (2012) [9](#)
46. Song, S., Lichtenberg, S.P., Xiao, J.: SUN RGB-D: A RGB-D scene understanding benchmark suite. In: *CVPR*. pp. 567–576 (2015) [9](#)
47. Tan, M., Le, Q.: EfficientNet: Rethinking model scaling for convolutional neural networks. In: *ICML*. pp. 6105–6114 (2019) [6](#), [10](#), [11](#)
48. Wang, C., Lucey, S., Perazzi, F., Wang, O.: Web stereo video supervision for depth prediction from dynamic scenes. In: *3DV*. pp. 348–357. *IEEE* (2019) [3](#), [4](#)
49. Wang, P., Shen, X., Lin, Z., Cohen, S., Price, B., Yuille, A.L.: Towards unified depth and semantic prediction from a single image. In: *CVPR*. pp. 2800–2809 (2015) [10](#), [11](#)
50. Xian, K., Shen, C., Cao, Z., Lu, H., Xiao, Y., Li, R., Luo, Z.: Monocular relative depth perception with web stereo data supervision. In: *CVPR*. pp. 311–320 (2018) [4](#)
51. Xian, K., Zhang, J., Wang, O., Mai, L., Lin, Z., Cao, Z.: Structure-guided ranking loss for single image depth prediction. In: *CVPR*. pp. 611–620 (2020) [2](#), [3](#), [4](#), [8](#), [9](#)
52. Xie, J., Girshick, R., Farhadi, A.: Deep3D: Fully automatic 2D-to-3D video conversion with deep convolutional neural networks. In: *ECCV*. pp. 842–857 (2016) [1](#)
53. Xie, S., Girshick, R., Dollár, P., Tu, Z., He, K.: Aggregated residual transformations for deep neural networks. In: *CVPR*. pp. 1492–1500 (2017) [10](#)
54. Xu, D., Ricci, E., Ouyang, W., Wang, X., Sebe, N.: Multi-scale continuous CRFs as sequential deep networks for monocular depth estimation. In: *CVPR*. pp. 5354–5362 (2017) [1](#), [2](#), [3](#)
55. Xu, Y., Zhu, X., Shi, J., Zhang, G., Bao, H., Li, H.: Depth completion from sparse LiDAR data with depth-normal constraints. In: *ICCV*. pp. 2811–2820 (2019) [2](#)
56. Yin, W., Liu, Y., Shen, C., Yan, Y.: Enforcing geometric constraints of virtual normal for depth prediction. In: *ICCV*. pp. 5684–5693 (Oct 2019) [1](#), [3](#), [10](#), [11](#)
57. Yu, F., Koltun, V., Funkhouser, T.: Dilated residual networks. In: *CVPR*. pp. 472–480 (2017) [10](#)
58. Zoran, D., Isola, P., Krishnan, D., Freeman, W.T.: Learning ordinal relationships for mid-level vision. In: *ICCV*. pp. 388–396 (2015) [2](#), [3](#)