Guiding Monocular Depth Estimation Using Depth-Attention Volume

 $\begin{array}{l} {\rm Lam \; Huynh^{1[0000-0002-8311-1288]}, \; Phong \; Nguyen-Ha^{1[0000-0002-9678-0886]}, \; Jiri \\ {\rm Matas^{2[0000-0003-0863-4844]}, \; Esa \; Rahtu^{3[0000-0001-8767-0864]}, \; and \; Janne \\ {\rm Heikkil\ddot{a}^{1[0000-0003-0073-0866]}} \end{array}$

¹ Center for Machine Vision and Signal Analysis, University of Oulu, Finland {lam.huynh,phong.nguyen,janne.heikkila}@oulu.fi

 ² Center for Machine Perception, Czech Technical University, Czech Republic matas@fel.cvut.cz
 ³ Computer Vision Group, Tampere University, Finland

esa.rahtu@tuni.fi

Abstract. Recovering the scene depth from a single image is an ill-posed problem that requires additional priors, often referred to as monocular depth cues, to disambiguate different 3D interpretations. In recent works, those priors have been learned in an end-to-end manner from large datasets by using deep neural networks. In this paper, we propose guiding depth estimation to favor planar structures that are ubiquitous especially in indoor environments. This is achieved by incorporating a non-local coplanarity constraint to the network with a novel attention mechanism called depth-attention volume (DAV). Experiments on two popular indoor datasets, namely NYU-Depth-v2 and ScanNet, show that our method achieves state-of-the-art depth estimation results while using only a fraction of the number of parameters needed by the competing methods. Code is available at: https://github.com/HuynLam/DAV

Keywords: Monocular depth \cdot Attention mechanism \cdot Depth estimation.

1 Introduction

Depth estimation is a fundamental problem in computer vision due to its wide variety of applications including 3D modeling, augmented reality and autonomous vehicles. Conventionally it has been tackled by using stereo and structure from motion techniques based on multiple view geometry [11,32]. In recent years, the advances in deep learning have made monocular depth estimation a compelling alternative [2,5,8,10,13,19,20,24,26,27,28,40,44].

In learning-based monocular depth estimation, the basic idea is simply to train a model to predict a depth map for a given input image, and to hope that the model can learn those monocular cues that enable inferring the depth directly from the pixel values. This kind of a brute-force approach requires a huge amount of training data and leads to large network architectures. It has been a common practice to use a deep encoder such as VGG-16 [5], ResNet-50 [19,26,27], ResNet-101 [8], ResNext-101 [40], SeNet-154 [2,13] followed by some upsampling and fusion strategy including the up-projection module [19], multi-scale feature fusion [13] or adaptive dense feature fusion [2] that all result in bulky networks with a large number of parameters. Because high computational complexity and memory requirements limit the use of these networks in practical applications, also fast monocular depth estimation models such as FastDepth [36] have been proposed, but their speed increase comes with the price of reduced accuracy. Moreover, despite of good results achieved with standard benchmark datasets such as NYU-Depth-v2, it still remains questionable if these networks are able to generalize well to unseen scenes and poses that are not present in the training data.

Instead of trying to learn all the monocular cues blindly from the data, in this paper, we investigate an approach where the learning is guided by exploiting a simple coplanarity constraint for scene points that are located on the same planar surfaces. Coplanarity is an important constraint especially in indoor environments that are composed of several non-parallel planar surfaces such as walls, floor, ceiling, tables, etc. We introduce a concept of depth-attention volume (DAV) to aggregate spatial information non-locally from those coplanar structures. We use both fronto-parallel and non-fronto-parallel constraints to learn the DAV in an end-to-end manner.

It should be noticed that plane approximations have already been used previously in monocular depth estimation, for example, in PlaneNet [24], where 3D planes were explicitly segmented and estimated from the images, but in contrast to these works, we embed the coplanarity constraint *implicitly* to the model by using the DAV, which is a building block inspired by the non-local neural networks [35]. Unlike the convolutional operation, it operates non-locally and produces a weighted average of the features across the whole image paying attention on planar structures, and favoring depth values that are originating from those planes. By using the DAV we not only incorporate an efficient and important geometric constraint to the model, but also enable shrinking the size of the network considerably without sacrificing the accuracy. To summarize, our key contributions include:

- A novel attention mechanism called depth-attention volume that captures non-local depth dependencies between coplanar points.
- An end-to-end neural network architecture that implicitly learns to recognize planar structures from the scene and use them as priors in monocular depth estimation.
- State-of-the-art depth estimation results on NYU-Depth-v2 and ScanNet datasets with a model that uses considerably less parameters than previous methods achieving similar performance.



Fig. 1. Visualization of depth-attention maps. The input image with four query points is shown on the left. The corresponding ground-truth and predicted depth maps are in the middle. Because of the coplanarity prior the depth of the textureless white wall can be accurately recovered. The ground-truth and predicted depth-attention maps for the query points are on the right. Warm colour indicates strong depth prediction ability for the query point.

2 Related work

Learning-based monocular depth estimation: Saxena et al. [29] is one of the first studies using Markov Random Field (MRF) to predict depth from a single image. Later on Eigen et al. proposed method to estimate depth using multi-scale deep network [6] and a multi-task learning model [5]. Since then, various studies using deep neural networks (DNNs) have been introduced. Laina et al. [19] employed a fully convolutional residual network (FCRN) as the encoder and four up-projection modules as the decoder to up-sample the depth map resolution. Fu et al. [8] successfully formulated monocular depth estimation as an ordinal regression problem. Qi et al. [26] proposed a network called GeoNet that investigate the duality between depth map and surface normal. The DNNs from Ren et al. [28] classified input images as indoor or outdoor before estimating the depth values. Lee et al. [20] suggested an idea of using a DNNs to estimate the relative depth between pairs of pixels. The proposed method from Jiao et al. [15] incorporated object segmentation into the training to increase depth estimation accuracy. Hu et al. [13] introduced an architecture that included an encoder, a decoder, a multi-scale feature fusion (MFF), and a new loss term for preserving edge structures. Inspired by [13], Chen et al. [2] used adaptive dense feature fusion (ADFF), and residual pyramid decoder in their network. The study by Facil et al. [7] proposed a DNNs that aims to learn calibration-aware patterns to improve the generalization capabilities of monocular depth prediction. Recently, Ramamonjisoa et al. [27] presented SharpNet that exploits occluding contours as an additional driving factor to optimize the depth estimation model besides the depth and the surface normal.

Plane-based approaches: Liu et al. [24] was the first study to consider using the planar constraint to predict depth maps from single images. Later the same authors published an incremental study to refine the quality of plane segmentation [23]. Yin et al. [40] formed a geometric constraint called virtual normal to predict the depth map as well as a point cloud and surface normals. Note that methods by Liu et al. focused explicitly on estimating a set of plane parameters



Fig. 2. Depth-attention volume (DAV) is a collection of depth-attention maps (Eq. 3, Figure 1) obtained using each image location as a query point at a time. Therefore, the DAV for an image of size $8H \times 8W$ is a 4D tensor of size $H \times W \times H \times W$.

and planar segmentation masks, while Yin et al. calculated a large virtual plane to train a DNNs that is robust to noise in the ground truth depth.

Attention mechanism: Attention was initially used in machine translation and it was brought to computer vision by Xu et al. [39]. Since then, attention mechanism has evolved and branched into channel-wise attention [12,33], spatial-wise attention [1,35] and mix attention [34] in order to tackle object detection and image classification problems. Some recent monocular depth estimation studies also followed this line of work. Xu et al. [38] proposed multi-scale spatial-wise attention to guide a Conditional Random Fields (CRFs) model. Li et al. [22] proposed a discriminative depth estimation model using channel-wise attention. Kong et al. [18] embedded a discrete binary mask, namely the pixel-wise attentional gating unit, into a residual block to modulate learned features.

In this paper, we propose using depth-attention volume (DAV) to encode non-local geometric dependencies. It can be seen as an attention mechanism that guides depth estimation to favor depth values originating from planar surfaces that are ubiquitous in man-made scenes. In contrast to previous plane-based approaches, we do not train the network to segment the planes explicitly, but instead, we let the network to learn the coplanarity constraint implicitly.

3 Proposed Method

This section describes the proposed depth estimation method. The first subsection defines the depth-attention volume and the following two subsections outline the network architecture and the loss functions. Further details are provided in the supplementary material.

3.1 Depth-attention volume

Given two image points $P_0 = (x_0, y_0)$ and $P_1 = (x_1, y_1)$ with corresponding depth values d_0 and d_1 , we define that the depth-attention $A(P_0, P_1)$ is the

ability of P_1 to predict the depth of P_0 . This ability is quantified as a confidence in the range [0, 1] so that 0 means no ability and 1 represents maximum certainty of being a good *predictor*.

To estimate A we make the assumption that the scene contains multiple nonparallel planes, which is common particularly in indoor environments. The depth values of all points belonging to the same plane are linearly dependent. Hence, they are good depth predictors of each other. To exploit this property, we detect N prominent planes from the training images and parameterize each plane with $S = (n_x, n_y, n_d, c)$, where (n_x, n_y, n_d) is the plane normal and c is the orthogonal distance from the origin. We construct the first-order depth-attention volumes for all N planes:

$$A_i(P_0, P_1) = 1 - \sigma(|S_i \cdot X_0| + |S_i \cdot X_1|), \quad i = 1, \dots, N$$
(1)

where σ is the sigmoid function, $X_0 = (x_0, y_0, d_0, 1)$ and $X_1 = (x_1, y_1, d_1, 1)$. These volumes are represented as 4-D tensors of size $H \times W \times H \times W$, where H and W are the vertical and horizontal sizes, respectively. In practice, one needs to subsample the volumes to keep the memory requirements reasonable. In all our experiments, we used a subsampling factor of 8.

In addition, we assume that all points located on the same fronto-parallel plane are good depth predictors of each other, because they share the same depth value. We use the ground-truth depths, and create a zero-order depthattention volume (DAV) for every training image

$$A_0(P_0, P_1) = 1 - \sigma(|d_0 - d_1|).$$
⁽²⁾

Finally, we combine these volumes by taking the maximum attention value of all volumes:

$$A_D(P_0, P_1) = \max(A_i(P_0, P_1)), \quad i = 0, \dots, N$$
(3)

It is easy to observe that DAV is a symmetric function, i.e. $A_D(P_0, P_1) = A_D(P_1, P_0)$.

If we consider P_0 to be a query point in the image as illustrated in Figure 1 (left), we can visualize the DAV as a two-dimensional *attention map* shown in Figure 1 (right). Figure 2 provides an example of a depth-attention volume generated from the ground truth depth map.

3.2 Network Architecture

Figure 3 gives an overview of our model that includes three main modules: an encoder, a non-local depth-attention module, and a decoder.

We opt to use a simplified dilated residual networks (DRN) with 22 layers (DRN-D-22) [41,42] as our encoder, which extracts high-resolution features and downsamples the input image only 8 times. The DRN-D-22 is a variation of DRN that completely removes max-pooling layers as well as smoothly distributes the dilation to minimize the gridding artifacts. This is crucial to our network,

because to make training feasible, the non-local depth-attention module needs to operate on a sub-sampled feature space. However, to capture meaningful spatial relationships this feature space also needs to be large enough.

The decoder part of our network contains a straightforward up-scaling scheme that increases the spatial dimension from 29×38 to 57×76 and then to 114×152 . Upsampling consists of two bilinear interpolation layers followed by convolutional layers with a kernel size of 3×3 . Two convolutional layers with a kernel size of 5×5 are then used to estimate the final depth map.

The non-local depth-attention module is located between the encoder and the decoder. It maps the input features \mathbf{X} to the output features \mathbf{Y} of the same size. The primary purpose of the module is to add the non-local information embedded in the depth-attention volume (DAV) to \mathbf{Y} , but it is also used to predict and learn the DAV based on the ground-truth data. The structure of the module is presented in Figure 4.

We implement the DAV-predictor by first transforming **X** into green and blue embeddings using 1×1 convolution. We exploit the symmetry of DAV, and maximize the correlation between these two spaces by applying cross-denormalization on both green and blue embeddings. Cross-denormalization is a conditional normalization technique [4] that is used to learn an affine transformation from the data. Specifically, the green embedding is first normalized to zero mean and unit standard deviation using batch-normalization (BN). Then, the blue embedding is convolved to create two tensors that are multiplied and added the normalized features from the green branch, and vise versa. The denormalized representations are then activated with ReLUs and transformed by another 1×1 convolution before multiplying with each others. Finally, the DAV is activated using the sigmoid function to ensure that the output values are in range [0, 1]. We empirically verified that applying cross-modulation in two embedding spaces is superior than using a single embedding with double the number of features.

Furthermore, \mathbf{X} is fed into the orange branch and multiplied with the estimated DAV to amplify the effect of the input features. Finally, we add a residual connection (red) to prevent the vanishing gradient problem when training our network.



Fig. 3. The pipeline of our proposed network. An image is passed through the encoder, then the non-local depth-attention module, and finally the decoder to produce the estimated depth map. The model is trained using $\mathcal{L}_{attention}$ and \mathcal{L}_{depth} losses, which are described in Subsection 3.3.



Fig. 4. Structure of the non-local depth-attention module. " \bigcirc " presents element-wise multiplication, " \bigoplus " presents element-wise sum, and " \bigotimes " is the outer product.

3.3 Loss Function

As illustrated in Figure 3 our loss function consists of two main components: attention loss and depth loss.

Attention loss: The primary goal of this term is to minimize the error between the estimated (output of the DAV-predictor in Figure 4) and the ground-truth DAV. The \mathcal{L}_{mae} is defined as the mean absolute error between the predicted and the ground truth depth-attention values:

$$\mathcal{L}_{mae} = \frac{1}{(HW)^2} \sum_{i} \sum_{j} |\hat{A}_{i,j} - A_{i,j}|$$

$$\tag{4}$$

where $\hat{A}_{i,j} \equiv \hat{A}_D(P_i, P_j)$ and $A_{i,j} \equiv A_D(P_i, P_j)$ are the predicted and ground truth depth-attention volumes.

In addition, we minimize the angle between the predicted and the ground truth depth-attention maps for all query positions i and j:

$$\mathcal{L}_{ang} = \frac{1}{HW} \left(\sum_{i} \left| 1 - \frac{\sum_{j} \hat{A}_{i,j} A_{i,j}}{\sqrt{\sum_{j} \hat{A}_{i,j}^{2} \sum_{j} A_{i,j}^{2}}} \right| + \sum_{j} \left| 1 - \frac{\sum_{i} \hat{A}_{i,j} A_{i,j}}{\sqrt{\sum_{i} \hat{A}_{i,j}^{2} \sum_{i} A_{i,j}^{2}}} \right| \right)$$
(5)

The full attention loss is defined by

$$\mathcal{L}_{attention} = \mathcal{L}_{mae} + \lambda \mathcal{L}_{ang} \tag{6}$$

where $\lambda \in \mathbb{R}^+$ is the weight loss coefficient.

Depth loss: Moreover, we define depth loss as a combination of three terms \mathcal{L}_{log} , \mathcal{L}_{grad} and \mathcal{L}_{norm} that were originally introduced in [13]. The \mathcal{L}_{log} loss is a variation of the L_1 norm that is calculated in the logarithm space and defined as

$$\mathcal{L}_{log} = \frac{1}{M} \sum_{i=1}^{M} F(|\hat{d}_i - d_i|)$$
(7)

where M is the number of valid depth values, d_i is the ground truth depth, \hat{d}_i is the predicted depth, and $F(x) = \log(x + \alpha)$ with α set to 0.5 in our experiments.

Another loss term is \mathcal{L}_{grad} , which is used to penalize sudden changes of edge structures in both x and y directions. It is defined by

$$\mathcal{L}_{grad} = \frac{1}{M} \sum_{i=1}^{M} F(\Delta_x(|\hat{d}_i - d_i|)) + F(\Delta_y(|\hat{d}_i - d_i|))$$
(8)

where Δ_x and Δ_y is the gradient of the error with respect to x and y. Finally, we use \mathcal{L}_{norm} to emphasize small details by minimizing the angle between the ground truth (n_i) and the predicted (\hat{n}_i) surface normals:

$$\mathcal{L}_{norm} = \frac{1}{M} \sum_{i=1}^{M} |1 - \hat{n}_i \cdot n_i|.$$
(9)

where surface normals are estimated as $n \equiv (-\nabla_x(d), -\nabla_y(d), 1)$ using Sobel filter, like [13]. The depth loss is then defined by

$$\mathcal{L}_{depth} = \mathcal{L}_{log} + \mu \mathcal{L}_{grad} + \theta \mathcal{L}_{norm} \tag{10}$$

where $\mu, \theta \in \mathbb{R}^+$ are weight loss coefficients. Our full loss is

$$\mathcal{L} = \mathcal{L}_{attention} + \gamma \mathcal{L}_{depth} \tag{11}$$

where $\gamma \in \mathbb{R}^+$ is a weight loss coefficient. Subsection 4.2 describes in detail how the network is trained using these loss functions.

4 Experiments

In this section, we evaluate the performance of the proposed method by comparing it against several baselines. We start by introducing datasets, evaluation metrics, and implementation details. The last three subsections contain the comparison to the state-of-the-art, ablation studies, and a cross-dataset evaluation. Further results are available in the supplementary material.

4.1 Datasets and evaluation metrics

Datasets: We assess the proposed method using NYU-Depth-v2 [30] and Scan-Net [3] datasets. NYU-Depth-v2 contains ~ 120K RGB-D images obtained from 464 indoor scenes. From the entire dataset, we use 50K images for training and the official test set of 654 images for evaluation. ScanNet dataset comprises of 2.5 million RGB-D images acquired from 1517 scenes. For this dataset, we use the training subset of ~ 20K images provided by the Robust Vision Challenge 2018 [9] (ROB). Unfortunately, the ROB test set is not available, so we report the results on the Scannet official test set of 5310 images instead. SUN-RGBD is yet another indoor dataset consisting of ~ 10K images collected with four different sensors. We do not use it for training, but only for cross-evaluating the pre-trained models on the test set of 5050 images. **Evaluation metrics:** The performance is assessed using the standard metrics provided for each dataset. That is, for NYU-Depth-v2 [30] we calculate the mean absolute relative error (REL), root mean square error (RMS), and thresholded accuracy (δ_i). For the ScanNet and SUN-RGBD dataset, we provide the mean absolute relative error (REL), mean square relative error (sqREL), scale-invariant mean squared error (SI), mean absolute error (iMAE), and root mean square error (iRMSE) of the inverse depth values. For iBims-1 benchmark [17], we compute 5 similar metrics as for NYU-Depth-v2 plus the root mean square error in log-space (log10), planarity errors ($\epsilon^{plan}, \epsilon^{orie}$), depth boundary errors ($\epsilon^{acc}, \epsilon^{comp}$), and directed depth error ($\epsilon^0, \epsilon^-, \epsilon^+$). Detailed definitions of the metrics are provided in the supplementary material.

4.2 Implementation Details

The proposed model is implemented with the Pytorch [25] framework, and trained using a single Tesla-V100, batch size of 32 images, and Adam optimizer [16] with $(\beta_1, \beta_2, \epsilon) = (0.9, 0.999, 10^{-8})$. The training process is split into three parts. During the first phase, we replace the DAV-predictor (Figure 4) with the DAVs computed from the ground truth depth maps. We train the model for 200 epochs using only the depth loss (Eq. 10) and the learning rate of 10^{-4} . In the second phase, we add the DAV-predictor to the model, freeze the weights of other parts of the model, and train for 200 epochs with the learning rate of 7.0×10^{-5} . In the last phase, we train the entire model for 300 epochs using the learning rate of 7.0×10^{-5} for the first 100 epochs and then reduce it at the rate of 5% per 25 epochs. The last two stages employ the full loss function defined in Equation (11). We set all the weight loss coefficients λ, μ, θ , and γ as 1.

We augment the training data using random scaling ([0.875, 1.25]), rotation ([-5.0, +5.0] degrees), horizontal flip, rectangular window droppings, and colorization. Planes, required for training, are obtained by fitting a parametric model to the back-projected 3D point cloud using RANSAC with the inlier threshold of 1 cm. We select at most the best N-planes in terms of the inlier count with a maximum of 100 iterations. Furthermore, we keep only planes that cover more than 7% of the image area.

4.3 Comparison with the state-of-the-art

In this section, we compare the proposed approach with the current state-of-theart monocular depth estimation methods.

NYU-Depth-v2: Table 1 contains the performance metrics on the official NYU-Depth-v2 test set for our method and for [2,5,8,10,13,19,20,24,26,27,28,40,44] In addition, the table shows the number of model parameters for each method. The performance figures for the baselines are obtained using the pre-trained models provided by the authors [2,8,13,24,27,40] or from the original papers if the model was not available [5,10,19,20,26,28,44]. Methods indicated with ** and [‡] are trained using the entire training set of 120K images or with external data,

Table 1. Evaluation results on the NYU-Depth-v2 dataset. Metrics with \downarrow mean lower is better and \uparrow mean higher is better. Timing is the average over 1000 images using a NVIDIA GTX-1080 GPU, in frames-per-second (FPS).

Methods	#params	Memory	FPS	$ \mathbf{REL}\downarrow$	$ \mathbf{RMS}\downarrow$	$\delta_1\uparrow$	δ_2^{\uparrow}	δ_3^{\uparrow}
Eigen'15 [5]**	141.1M	-	-	0.215	0.907	0.611	0.887	0.971
Laina'16 [19]**	63.4M	-	-	0.127	0.573	0.811	0.953	0.988
Liu'18 [24] [‡]	47.5M	124.6MB	93	0.142	0.514	0.812	0.957	0.989
Fu'18 [8] **	110.0M	489.1MB	42	0.115	0.509	0.828	0.965	0.992
Qi'18 [26]	67.2M	-	-	0.128	0.569	0.834	0.960	0.990
Hao'18 [10]	60.0M	-	-	0.127	0.555	0.841	0.966	0.991
Lee'19 [20]	118.6M	-	-	0.131	0.538	0.837	0.971	0.994
Ren'19 [28] **	49.8M	-	-	0.113	0.501	0.833	0.968	0.993
Zhang'19 [44]	95.4M	-	-	0.121	0.497	0.846	0.968	0.994
Ramam.'19 [27] [‡]	80.4M	336.6MB	47	0.139	0.502	0.836	0.966	0.993
Hu'19 [13]	157.0M	679.7MB	15	0.115	0.530	0.866	0.975	0.993
Chen'19 [2]	210.3M	$1250.9 \mathrm{MB}$	12	0.111	0.514	0.878	0.977	0.994
Yin'19 [40]	114.2M	437.6MB	37	0.108	0.416	0.875	0.976	0.994
Ours	25.1M	96.1MB	218	0.108	0.412	0.882	0.980	0.996

respectively. For instance, Ramamonjisoa et al. [27] trained the method using synthetic dataset PBRS [43] before fine-tuning on NYU-Depth-v2. The best performance is achieved by the proposed model that also contains the least amount of parameters. The best performing baselines, Yin et al. [40], Hu et al. [13], and Chen et al. [2], have 4.5, 6.2, and 8.3 times more parameters compared to ours, respectively. Figure 5 provides an additional illustration of the model parameters with respect to the performance.

Figure 6 shows qualitative examples of the obtained depth maps. In this case, the maps for the baseline methods are produced using the pre-trained models provided by the authors. The method by Eigen and Fergus [5] performs well on uniform regions, but has difficulties in detailed structures. Laina et al. [19] produces overly smoothed depth maps losing many small details. In contrast, Fu et



Fig. 5. Analyzing the accuracy $\delta_1(\%)$ and mean absolute relative error (%) with respect to the number of parameters (*millions*) for recent monocular depth estimation methods on NYU-Depth-v2. The left picture presents the thresholded accuracy where higher values are better, while the right picture shows the absolute relative error where lower values are better.



11

Fig. 6. Qualitative results on the official NYU-Depth-v2 [30] test set from different methods. The color indicates the distance where red is far and blue is close. Our estimated depth maps are closer to the ground truth depth when comparing with state-of-art methods.

al. [8] returns many details, but with the expense of discontinuities inside objects or smooth areas. The depth images by Ramamonjisoa et al. [27] contain noise and are prone to miss fine details. Yin et al. [40], Hu et al. [13], and Chen et al. [2] provide the best results among the baselines. However, they have difficulties e.g. on the third (near the desk and table) and the fourth examples from the left (wall area). We provide further qualitative examples in the supplementary material.

ScanNet: Table 2 contains the performance figures on the official ScanNet test set for our method, Ren et al. [28] (taken from the original paper), Hu et al. [13] and Chen et al. [2]. We use the public code from [2,13] to train their models. Unfortunately, the other baselines do not provide the results for ScanNet official

Architecture	#params	REL	sqREL	SI	iMAE	iRMSE	Test set
CSWS_E_ROB [21]	65.8M	0.150	0.060	0.020	0.100	0.130	
DORN_ROB [8]	110.0M	0.140	0.060	0.020	0.100	0.130	ROB
DABC_ROB [22]	56.6M	0.140	0.060	0.020	0.100	0.130	
Hu'19 [13]	157.0M	0.139	0.081	0.016	0.100	0.105	
Chen'19 [2]	210.3M	0.134	0.077	0.015	0.093	0.100	Official
Ren'19 [28]	49.8M	0.138	0.057	-	-	-	Oniciai
Ours	$25.1\mathrm{M}$	0.118	0.057	0.015	0.089	0.097	

Table 2. Evaluation results on ScanNet [3].

test set. Moreover, the test set used in the Robust Vision Challenge (ROB) is not available at the moment and we are unable to report our performance on that. Nevertheless, we have included the best methods from the ROB challenge in Table 2 to provide indicative comparison. Note that all methods are trained with the same ROB training split. The proposed model outperforms [28] with a clear margin in terms of REL. The results are also substantially better compared to ROB challenge methods, although the comparison is not strictly fair due to different test splits. Figure 7 provides qualitative comparison between our method and [2,13,22], using the sample images provided in [22]. The geometric structures and details are clearly better extracted by our method.



Fig. 7. Predicted depth maps from our model with baselines on the official ScanNet [3] test set.

Table 3. The iBims-1 benchmark

Method	$\mathbf{REL}\downarrow$	$\log 10 \downarrow$	$\mathbf{RMS} \downarrow$	$\delta_1\uparrow$	$\delta_2\uparrow$	$\delta_3\uparrow$	$\epsilon^{plan} \downarrow$	$\epsilon^{orie} \downarrow$	$\epsilon^{acc}\downarrow$	$\epsilon^{comp}\downarrow$	$\epsilon^{0}\uparrow$	$\epsilon^{-\downarrow}$	$\epsilon^+\downarrow$
Liu'18 [24]	0.29	0.17	1.45	0.41	0.70	0.86	7.26	17.24	4.84	8.86	71.24	28.36	0.40
Ramam.'19 [27]	0.26	0.11	1.07	0.59	0.84	0.94	9.95	25.67	3.52	7.61	84.03	9.48	6.49
Ours	0.24	0.10	1.06	0.59	0.84	0.94	7.21	18.45	3.46	7.43	84.36	6.84	6.27

Planarity error analysis: We also evaluated our method on the iBims-1 benchmark [17] and compared it with two recent works [24,27]. The results, shown in Table 3, indicate that we outperform the baselines in most of the metrics, including plane related ones. Extensive planarity analysis is provided in the supplementary material.

4.4 Ablation studies

Firstly, we assess how the number of prominent planes, used to estimate the ground truth DAVs in the training phase, affects the performance (see Sec. 3.1). To this end, we train our model using the fronto-parallel planes (see Eq. 2) plus three, five, and seven non-fronto-parallel planes (N in Eq. 1). The corresponding results for the NYU-Depth-v2 test set are provided in Table 4. One can observe that the results improve by increasing the number of planes up to five and decrease after that. Possible explanation for this could be that the images used in the experiments do not typically contain more than five significant planes that can predict the depth values reliably. We also re-trained our model without the non-local depth attention (DAV) module (and any planes) and the performance degraded substantially as shown in Table 4.

Secondly, we study the impact of the attention loss term (Eq. 6). For this purpose, we first train our model with and without the attention loss, and then continue training by dropping the attention loss after convergence. We report the results in Table 5. The model without the attention loss has clearly inferior performance indicating the importance of this loss term. Furthermore, continuing training by dropping the attention loss also degrades the performance.

DAV-types	$\mathbf{REL}\!\!\downarrow$	$\mathbf{RMS}{\downarrow}$	$\delta_1\uparrow$	$\delta_2\uparrow$	$\delta_3\uparrow$
w/o DAV-module	0.140	0.577	0.827	0.960	0.989
-Plane-DAV	0.116	0.442	0.867	0.976	0.995
3-Plane-DAV	0.110	0.421	0.879	0.978	0.995
5-Plane-DAV	0.108	0.412	0.882	0.980	0.996
7-Plane-DAV	0.111	0.447	0.851	0.970	0.993

Table 4. Performance of our model using different types of depth-attention volume.

Table 5. Ablation studies of models without and with the attention loss on the NYU-Depth-v2. This shows the importance of the DAV in guiding the monocular depth model.

Training	$\mathbf{REL}\downarrow$	$\mathbf{RMS}{\downarrow}$	$\delta_1\uparrow$	$\delta_2\uparrow$	$\delta_3\uparrow$
w/o $\mathcal{L}_{attention}$	0.126	0.540	0.841	0.967	0.992
w/ full loss	0.108	0.412	0.882	0.980	0.996
continue w/o $\mathcal{L}_{attention}$	0.109	0.415	0.882	0.979	0.995

Table 6. Cross-dataset evaluation with training on NYU-Depth-v2 and testing onSUN-RGBD.

Models	#params	REL	sqREL	SI	iMAE	iRMSE
w/o DAV-module	17.5M	0.254	0.416	0.035	0.111	0.091
Hu'19 [13]	157.0M	0.245	0.389	0.031	0.108	0.087
Chen'19 [2]	210.3M	0.243	0.393	0.031	0.102	0.069
Ours	25.1M	0.238	0.387	0.030	0.104	0.075



Fig. 8. Direct results on SUN RGB-D dataset [31] without fine-tuning. Some regions in the white boxes show missing or incorrect depth values from the ground truth data.

4.5 Cross-dataset evaluation

To assess the generalisation properties of the model, we perform a cross-dataset evaluation, where we train the network using NYU-Depth-v2 and test with SUN-RGBD [14,31,37] without any fine-tuning. We also evaluate the baseline methods from [2,13] and report the results in Table 6. As can be seen our model performs favourably compared to the other methods. Figure 8 contains a few examples of the results with the SUN-RGBD dataset. One can observe that our model is able to well estimate the geometric structures and details of the scene despite the differences in data distributions between the training and testing sets. Moreover, we evaluated our model without the DAV-module in the same cross-dataset setup. The results, shown in Table 6, clearly demonstrates that the DAV-module improves the generalization.

5 Conclusions

This paper proposed a novel monocular depth estimation method that incorporates a non-local coplanarity constraint with a novel attention mechanism called depth-attention volume (DAV). The proposed attention mechanism encourages depth estimation to favor planar structures, which are common especially in indoor environments. The DAV enables more efficient learning of the necessary priors, which results in considerable reduction in the number of model parameters. The performance of the proposed solution is state-of-the-art on two popular benchmark datasets while using 2-8 times less parameters than competing methods. Finally, the generalisation ability of the method was further demonstrated in cross dataset experiments.

References

- Bello, I., Zoph, B., Vaswani, A., Shlens, J., Le, Q.V.: Attention augmented convolutional networks. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 3286–3295 (2019)
- Chen, X., Chen, X., Zha, Z.J.: Structure-aware residual pyramid network for monocular depth estimation. In: Proceedings of the 28th International Joint Conference on Artificial Intelligence. pp. 694–700. AAAI Press (2019)
- Dai, A., Chang, A.X., Savva, M., Halber, M., Funkhouser, T., Nießner, M.: Scannet: Richly-annotated 3d reconstructions of indoor scenes. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5828–5839 (2017)
- Dumoulin, V., Shlens, J., Kudlur, M.: A learned representation for artistic style. In: International Conference on Learning Representations ICLR (2017)
- Eigen, D., Fergus, R.: Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In: Proceedings of the IEEE international conference on computer vision. pp. 2650–2658 (2015)
- Eigen, D., Puhrsch, C., Fergus, R.: Depth map prediction from a single image using a multi-scale deep network. In: Advances in neural information processing systems. pp. 2366–2374 (2014)
- Facil, J.M., Ummenhofer, B., Zhou, H., Montesano, L., Brox, T., Civera, J.: Camconvs: camera-aware multi-scale convolutions for single-view depth. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 11826– 11835 (2019)
- Fu, H., Gong, M., Wang, C., Batmanghelich, K., Tao, D.: Deep ordinal regression network for monocular depth estimation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2002–2011 (2018)
- Geiger, A. and Nie
 ßner, M. and Dai, A.: Robust Vision Challenge CVPR Workshop (2018)
- Hao, Z., Li, Y., You, S., Lu, F.: Detail preserving depth estimation from a single image using attention guided networks. In: 2018 International Conference on 3D Vision (3DV). pp. 304–313. IEEE (2018)
- 11. Hartley, R., Zisserman, A.: Multiple view geometry in computer vision. Cambridge university press (2003)
- Hu, J., Shen, L., Sun, G.: Squeeze-and-excitation networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7132–7141 (2018)
- Hu, J., Ozay, M., Zhang, Y., Okatani, T.: Revisiting single image depth estimation: Toward higher resolution maps with accurate object boundaries. In: IEEE Winter Conf. on Applications of Computer Vision (WACV) (2019)
- Janoch, A., Karayev, S., Jia, Y., Barron, J.T., Fritz, M., Saenko, K., Darrell, T.: A category-level 3d object dataset: Putting the kinect to work. In: Consumer depth cameras for computer vision, pp. 141–165. Springer (2013)
- Jiao, J., Cao, Y., Song, Y., Lau, R.: Look deeper into depth: Monocular depth estimation with semantic booster and attention-driven loss. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 53–69 (2018)
- Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
- Koch, T., Liebel, L., Fraundorfer, F., Körner, M.: Evaluation of cnn-based singleimage depth estimation methods. In: Leal-Taix, L., Roth, S. (eds.) European Conference on Computer Vision Workshop (ECCV-WS). pp. 331–348. Springer International Publishing (2018)

- 16 L. Huynh et al.
- Kong, S., Fowlkes, C.: Pixel-wise attentional gating for scene parsing. In: 2019 IEEE Winter Conference on Applications of Computer Vision (WACV). pp. 1024– 1033. IEEE (2019)
- Laina, I., Rupprecht, C., Belagiannis, V., Tombari, F., Navab, N.: Deeper depth prediction with fully convolutional residual networks. In: 2016 Fourth international conference on 3D vision (3DV). pp. 239–248. IEEE (2016)
- Lee, J.H., Kim, C.S.: Monocular depth estimation using relative depth maps. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 9729–9738 (2019)
- Li, B., Dai, Y., He, M.: Monocular depth estimation with hierarchical fusion of dilated CNNs and soft-weighted-sum inference. Pattern Recognition 83, 328–339 (2018)
- Li, R., Xian, K., Shen, C., Cao, Z., Lu, H., Hang, L.: Deep attention-based classification network for robust depth prediction. In: Asian Conference on Computer Vision. pp. 663–678. Springer (2018)
- Liu, C., Kim, K., Gu, J., Furukawa, Y., Kautz, J.: Planercnn: 3d plane detection and reconstruction from a single image. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 4450–4459 (2019)
- Liu, C., Yang, J., Ceylan, D., Yumer, E., Furukawa, Y.: Planenet: Piece-wise planar reconstruction from a single rgb image. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2579–2588 (2018)
- 25. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S.: Pytorch: An imperative style, high-performance deep learning library. In: Advances in Neural Information Processing Systems 32, pp. 8024-8035. Curran Associates, Inc. (2019), http://papers.neurips.cc/paper/ 9015-pytorch-an-imperative-style-high-performance-deep-learning-library pdf
- Qi, X., Liao, R., Liu, Z., Urtasun, R., Jia, J.: Geonet: Geometric neural network for joint depth and surface normal estimation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 283–291 (2018)
- 27. Ramamonjisoa, M., Lepetit, V.: Sharpnet: Fast and accurate recovery of occluding contours in monocular depth estimation. The IEEE International Conference on Computer Vision (ICCV) Workshops (2019)
- Ren, H., El-khamy, M., Lee, J.: Deep robust single image depth estimation neural network using scene understanding. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops. pp. 37–45 (2019)
- Saxena, A., Chung, S.H., Ng, A.Y.: Learning depth from single monocular images. In: Advances in neural information processing systems. pp. 1161–1168 (2006)
- Silberman, N., Hoiem, D., Kohli, P., Fergus, R.: Indoor segmentation and support inference from rgbd images. In: European Conference on Computer Vision. pp. 746–760. Springer (2012)
- Song, S., Lichtenberg, S.P., Xiao, J.: Sun rgb-d: A rgb-d scene understanding benchmark suite. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 567–576 (2015)
- Szeliski, R.: Structure from motion. In: Computer Vision, pp. 303–334. Springer (2011)
- 33. Tan, M., Chen, B., Pang, R., Vasudevan, V., Sandler, M., Howard, A., Le, Q.V.: Mnasnet: Platform-aware neural architecture search for mobile. In: Proceedings of

the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2820–2828 (2019)

- 34. Wang, F., Jiang, M., Qian, C., Yang, S., Li, C., Zhang, H., Wang, X., Tang, X.: Residual attention network for image classification. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3156–3164 (2017)
- Wang, X., Girshick, R., Gupta, A., He, K.: Non-local neural networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 7794–7803 (2018)
- Wofk, D., Ma, F., Yang, T.J., Karaman, S., Sze, V.: Fastdepth: Fast monocular depth estimation on embedded systems. In: 2019 International Conference on Robotics and Automation (ICRA). pp. 6101–6108. IEEE (2019)
- Xiao, J., Owens, A., Torralba, A.: Sun3d: A database of big spaces reconstructed using sfm and object labels. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1625–1632 (2013)
- Xu, D., Wang, W., Tang, H., Liu, H., Sebe, N., Ricci, E.: Structured attention guided convolutional neural fields for monocular depth estimation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3917– 3925 (2018)
- 39. Xu, K., Ba, J., Kiros, R., Cho, K., Courville, A., Salakhudinov, R., Zemel, R., Bengio, Y.: Show, attend and tell: Neural image caption generation with visual attention. In: International conference on machine learning. pp. 2048–2057 (2015)
- 40. Yin, W., Liu, Y., Shen, C., Yan, Y.: Enforcing geometric constraints of virtual normal for depth prediction. In: The IEEE International Conference on Computer Vision (ICCV) (2019)
- 41. Yu, F., Koltun, V.: Multi-scale context aggregation by dilated convolutions. In: International Conference on Learning Representations (ICLR) (2016)
- Yu, F., Koltun, V., Funkhouser, T.: Dilated residual networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 472–480 (2017)
- Zhang, Y., Song, S., Yumer, E., Savva, M., Lee, J.Y., Jin, H., Funkhouser, T.: Physically-based rendering for indoor scene understanding using convolutional neural networks. The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017)
- Zhang, Z., Cui, Z., Xu, C., Yan, Y., Sebe, N., Yang, J.: Pattern-affinitive propagation across depth, surface normal and semantic segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 4106–4115 (2019)