

Adversarial View-Consistent Learning for Monocular Depth Estimation

Yixuan Liu*¹
liuyixua17@mails.tsinghua.edu.cn

Yuwang Wang²
yuwwan@microsoft.com

Shengjin Wang¹
wsgsj@tsinghua.edu.cn

¹ Tsinghua University
Beijing, China

² Microsoft Research Asia
Beijing, China

Abstract

This paper addresses the problem of Monocular Depth Estimation (MDE). Existing approaches on MDE usually model it as a pixel-level regression problem, ignoring the underlying geometry property. We empirically find this may result in sub-optimal solution: while the predicted depth map presents small loss value in one specific view, it may exhibit large loss if viewed in different directions. In this paper, inspired by multi-view stereo (MVS), we propose an Adversarial View-Consistent Learning (AVCL) framework to force the estimated depth map to be all reasonable viewed from multiple views. To this end, we first design a differentiable depth map warping operation, which is end-to-end trainable, and then propose a pose generator to generate novel views for a given image in an adversarial manner. Collaborating with the differentiable depth map warping operation, the pose generator encourages the depth estimation network to learn from hard views, hence produce view-consistent depth maps. We evaluate our method on NYU Depth V2 dataset and the experimental results show promising performance gain upon state-of-the-art MDE approaches.

1 Introduction

Learning to predict depth from 2D images has been an actively studied problem over recent years, with widespread applications in 3D vision, such as scene modeling, 3D object detection and segmentation. Conventional approaches usually require multiple views of the observation or limited environment, so that the application is restricted. Therefore, learning to predict depth from a single image, *i.e.*, *Monocular Depth Estimation (MDE)*, is now attracting more research attention.

MDE is a typical ill-posed problem, because one 2D images may correspond to many reasonable 3D scenes. However, structural prior knowledge of natural images (texture / object shapes / object continuity) can provide rich information to overcome this inherent ambiguity, and is leveraged in early works on MDE [10, 17, 31, 32].

Recently, along with the rise of deep learning, MDE performance is significantly facilitated by Deep CNN-based models [8, 9, 18, 26, 30, 37], in which the structural prior

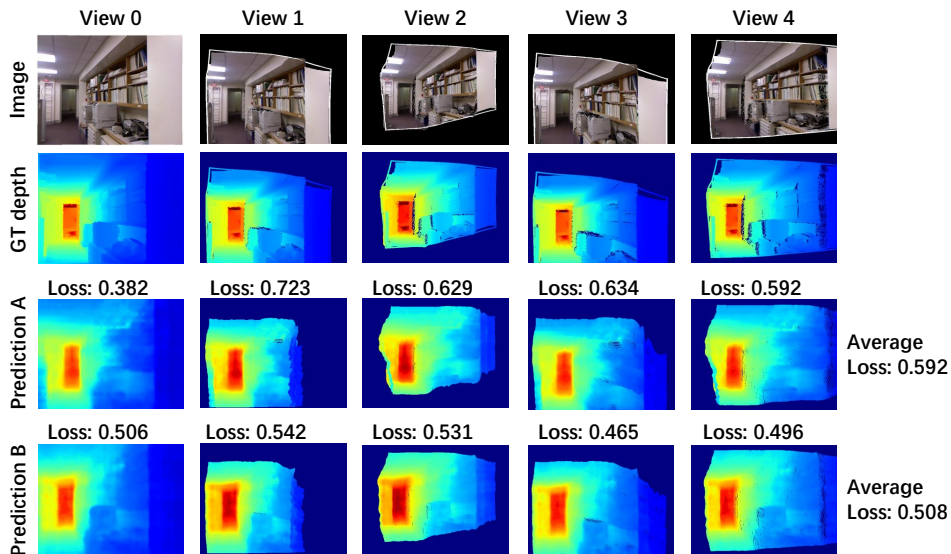


Figure 1: An illustration on view-consistent depth map (prediction B) and view-inconsistent depth map (prediction A). Viewed in multiple direction, prediction B presents small loss value in average. In contrast, prediction A may present small loss view in some specific view (view 0), but the average loss across multiple views is large.

knowledge is learned from training data. Most of the existing CNN-based methods focus on designing more suitable network structures for learning depth [4, 5, 39], or proposing new learning objectives [6, 13] to replace the vanilla loss functions (e.g. L_1 or L_2). However, no matter with which form of network structure or loss function, depth prediction in every single pixel has equal contribution to the loss value. As a consequence, two depth map prediction with the same loss value, may present quite different results (Fig. 1).

To address this problem, we empirically find that applying multiple supervisions on multiple generated views of the same scene helps the model to learn a more reasonable depth map. Fig. 1 shows an illustration: A good depth prediction, no matter viewed in which direction, always has small loss values with regard to the ground truth, which means the prediction is *view-consistent*. In contrast, if a bad depth prediction happens to have small loss in one specific view, it is unlikely the loss values in all views are small as well. In other words, the prediction is not view-consistent.

Moreover, the transformation pose of the view point also matters during training. With limited training time, it is impractical to enumerate all possible transformation poses. Also, not all views are valuable for training the depth model: if a generated view still has small loss, it means the cross-view regularization does not work much in this view. Therefore, it is of great importance to efficiently mine *hard* views, which benefit the training more.

In this work, we propose an Adversarial View-Consistent Learning (AVCL) framework for monocular depth estimation (Fig. 2). In AVCL, the predicted depth map for an input image is transformed to another view by a differentiable depth map warping operation, and an additional supervision is applied upon the warped depth map to regularize the predicted depth map to be view-consistent. Furthermore, in order to efficiently mine hard views for

the additional supervision, we design a warping pose generation module to automatically generate the warping pose, which is trained adversarially against the objective of the warped branch. Experiments on NYU Depth v2 dataset show both the view-consistent (multiple supervisions) learning and also the adversarial pose generation module bring significant performance gain on baselines in terms of 7 commonly used metrics. Moreover, as AVCL is complementary to existing methods, the performance gain is addable.

The contribution of this work can be summarized as follows:

- (1) We show that applying multiple supervisions on top of multiple views of the depth map leads the model to a better optimum, where the predicted depth maps are view-consistent.
- (2) We propose a differentiable depth map warping operation to generate multiple views of the predicted depth map, so that additional supervision can be simply applied on the warped branch, and the entire framework is end-to-end trainable.
- (3) We propose an adversarial warping pose generation module to automatically mine the hard views in view-consistent learning, which further benefits the learning of depth.

2 Related Work

Monocular Depth Estimation. Early works [22, 31] in MDE usually base on hand-crafted features [7, 23, 32] and probabilistic graph models [27, 31, 44], learn from structural prior knowledge of natural images such as texture and object shapes, while suffer in unconstrained scenes due to excessively strong assumptions on scene geometry. Recent advances in MDE [5, 14, 15, 18, 20, 24, 30] mostly employ deep CNN to learn depth from large amount of RGB-D data, mitigating previous strong geometry assumptions, and dramatically increase performance. Most of the existing deep MDE approaches focus on either designing adaptive network architectures [9, 39], or re-formulate the learning objective by introducing new loss functions [5, 18]. Although these methods present considerable accuracy, to our knowledge, there are few existing efforts on learning view-consistent depth maps.

Multi-view stereo (MVS) methods reconstruct 3D model of the scene from multi-view images [6]. Recent methods [9, 35] focus on depth map fusion and achieve competitive results in practice *w.r.t.* early volumetric-based methods [16, 33]. Deep learning methods [13, 40] shows great success on this task with the large datasets are available. Even though single depth estimation task does not belong to MVS, but MVS methods take depth map as an important intermediate to fusion the 3D representation of the target scene. Our method is inspired by MVS since multiviews act as an effective constraints on the geometry of scenes. The error of depth map estimated from one view can be magnified when observed from other viewpoints.

Adversarial learning is first proposed for generative models [8] and has been extensively studied and leveraged in applications like image synthesis [11, 29] and manipulation [11, 37, 43]. Despite its success on image generation, the idea of adversarial learning is also applied in discriminative models for representation learning, where the mainly usage it to mine or to generate hard training samples. For instance, Wang *et al.* [38] employ adversarial learning to generate occlusion and spatial transformation on images for object detection; BIER [28] adopt adversarial loss to ensure the diversity of embeddings so that the ensemble yields better accuracy. In this work, we follow these ideas and apply adversarial learning to generate hard views for training.

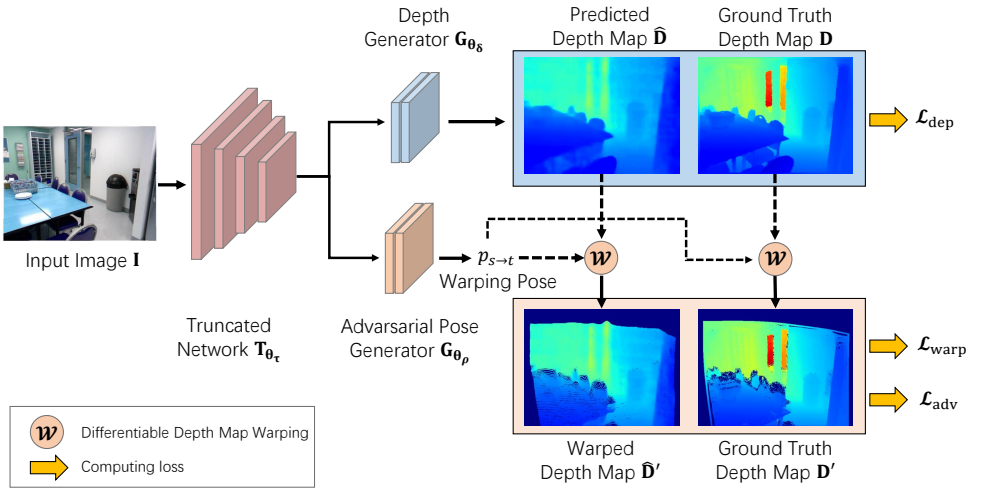


Figure 2: An overview of the proposed Adversarial View-Consistent Learning framework.

3 Approach

In this section we introduce the framework of the proposed adversarial view-consistent learning (AVCL). As shown in Fig. 2, AVCL involves two key components, namely *differentiable depth map warping* (Sec. 3.1) and *adversarial pose generation* (Sec. 3.2). In differentiable depth map warping, given an initial depth map and a specific warping pose, a novel view of the depth map is generated, and all the operations are differentiable so that back-propagation can be applied. Resort to the differentiable warping operation, we are able to impose multiple supervision signals on the generated novel views of the depth map. With supervisions on multiple views, the model is expected to output more view-consistent depth map. However, how to choose the warping pose here remains a problem. To this end, we design an adversarial pose generation module, which is also differentiable, and output an adversarial warping pose. Warped with the adversarial pose, the novel view of depth map is expected to have large loss compared to the warped ground truth depth map. In other words, the module adversarially mines a *hard* view to train the depth generator.

The combination of a depth generator, the differentiable depth map warping operation and an adversarial pose generation module forms the AVCL framework and will now be described in more detail in the following sections.

3.1 Differentiable Depth Map Warping

The differentiable depth map warping operation $\mathcal{W}(\mathbf{D}_s, p_{s \rightarrow t})$ takes a source depth map $\mathbf{D}_s \in \mathbb{R}^{H \times W}$ and a 6-DoF pose vector $p_{s \rightarrow t}$ as input and outputs a target depth map $\mathbf{D}_t \in \mathbb{R}^{H \times W}$. The target depth map is a novel view with the camera moves with pose $p_{s \rightarrow t}$.

Let $p_s = (x_s, y_s, 1)^T$ denote the homogeneous coordinates of a pixel in the source depth map \mathbf{D}_s . For every p_s , we calculate the corresponding target coordinate $p_t = (x_t, y_t, 1)^T$ and depth z_t by,

$$\begin{pmatrix} x_t z_t \\ y_t z_t \\ z_t \end{pmatrix} = K (R_{s \rightarrow t} \mathbf{D}_s(p_s) K^{-1} p_s + T_{s \rightarrow t}) \quad (1)$$




Figure 3: Gradient flow during training in AVCL.

where $\mathbf{D}_s(p_s)$ is the depth value of the pixel, and K denotes the camera intrinsics matrix. $R_{s \rightarrow t} \in \mathbb{R}^{3 \times 3}$ and $T_{s \rightarrow t} \in \mathbb{R}^3$ are the rotation matrix and translation vector from the source view to the target view, respectively, calculated from $p_{s \rightarrow t}$. The whole process can be broken down into three procedures. First, the source pixel coordinates are mapped into world coordinates, *i.e.*, $p_w = \mathbf{D}_s(p_s)K^{-1}p_s$. Second, the camera moves with pose $p_{s \rightarrow t}$, so that the world coordinates are transformed to $p'_w = R_{s \rightarrow t}p_w + T_{s \rightarrow t}$. Finally, the world coordinates are being transformed into a different camera coordinate frame.

Note that the coordinates $p_t = (x_t, y_t, 1)^T$ here are continuous coordinates, in order to obtain the pixel-wise depth map \mathbf{D}_t , we need to fill the depth values into corresponding pixels. Different from the cases in STN [42] and SfMLearner [42] where bilinear sampling is employed to deal with the one-to-many mapping, in our case this is a many-to-one mapping since different coordinates may fall into the same pixel. To ensure the differentiability, we use the following approximation,

$$\mathbf{D}_t(i, j) = \min_{x, y} z(x, y), \quad \forall (x, y) \in \{(x_t, y_t) \mid \lfloor x_t \rfloor = i, \lfloor y_t \rfloor = j, (x_t, y_t) \in \mathcal{P}_t\} \quad (2)$$

where i and j are the discrete pixel coordinates, \mathcal{P}_t is the collection of all the projected target coordinates, and $z(x, y)$ represents the depth value of the point (x, y) . We first map the continuous coordinates into discrete ones with the floor operation. Accordingly, if multiple points fall into the same pixel, we compare their depth value, and select the minimum value among them as the depth of this pixel. This is reasonable because if we see two objects coincide, the front object always occludes the back one. Moreover, if a pixel received no depth value, this pixel is marked as ignored and do not participate in computing the loss.

As a summary, the warping operation $\mathcal{W}(\mathbf{D}_s, p_{s \rightarrow t})$ generates a transformed view of the depth map. The computation therein is either matrix manipulation (the projection procedure) or simple value assignment (the coordinate discretization procedure), so that the entire operation is differentiable and able to be embedded in the network for end-to-end training.

3.2 Adversarial Pose Generation

With the help of the differentiable warping operation, we are able to generate multiple views of the predicted / ground-truth depth map. Accordingly, multiple supervisions can be applied to restrict the depth generator to output more view-consistent depth maps. Nevertheless, a question arises here: How should we set the warping pose, so that the model benefits the most from learning such warped depth map? The simplest solution is to randomly sample

the warping pose. As will be illustrated in Sec. 4.2, random warping pose does help, but the performance gain is limited. To address this issue, we propose an adversarial pose generation module to automatically generate the warping pose. The pose generation module is trained adversarially against the depth generator, hence is expected to generate a pose that results in the *hardest* view of the predict depth map. In other words, viewed in the generated pose, the predicted depth map is expected to have the largest error *w.r.t.* the warped ground truth.

Formally, let us denote the truncated network as $\mathbf{T}_{\theta_\tau}(\mathbf{I})$, where \mathbf{I} is the input image, and the output $\mathbf{z} = \mathbf{T}_{\theta_\tau}(\mathbf{I})$ is the intermedia feature map. The depth generator $\mathbf{G}_{\theta_\delta}(\mathbf{z})$, consisting of a 3×3 , an 1×1 convolution layer and a fc layer, takes \mathbf{z} as input and outputs the predicted depth $\hat{\mathbf{D}}$. The adversarial pose generation module $\mathbf{G}_{\theta_p}(\mathbf{z})$ is another branch comprised of a stack of several convolution layers and a fully-connected layer, then followed by a pre-defined scaled and shifted sigmoid function to restrict the range of the output pose vector $p = \mathbf{G}_{\theta_p}(\mathbf{z}) \in \mathbb{R}^6$. The pose vector p is then used to warp the predicted depth map $\hat{\mathbf{D}}$ and the ground truth depth map \mathbf{D} , yielding $\hat{\mathbf{D}}'$ and \mathbf{D}' respectively. Accordingly, we compute L_1 loss for both $\hat{\mathbf{D}}$ and $\hat{\mathbf{D}}'$,

$$\mathcal{L}_{\text{dep}} = \|\hat{\mathbf{D}} - \mathbf{D}\|_1, \quad \mathcal{L}_{\text{warp}} = \|\hat{\mathbf{D}}' - \mathbf{D}'\|_1 \quad (3)$$

During back-propagation, we update the parameters in the shared network \mathbf{T}_{θ_τ} with gradient from both the two branches, $\frac{\partial \mathcal{L}_{\text{dep}} + \mathcal{L}_{\text{warp}}}{\partial \theta_\tau}$ and update the parameters of the depth generator $\mathbf{G}_{\theta_\delta}$ with only $\frac{\partial \mathcal{L}_{\text{dep}}}{\partial \theta_\delta}$. In contrast, we expect the pose generator to output hard poses, so we optimize the parameters of \mathbf{G}_{θ_p} to the opposite direction by inverting the gradient, *i.e.*, using the negative gradient $-\frac{\partial \mathcal{L}_{\text{warp}}}{\partial \theta_p}$. The adversarial loss is only applied to the pose generator and can be written as,

$$\mathcal{L}_{\text{adv}} = -\mathcal{L}_{\text{warp}} + \lambda \cdot (p \odot p) \quad (4)$$

where λ is a vector of hyperparameters that regularize the value of the warping pose not to be too large. For an intuitive understanding, we also visualize the gradient flows during back-propagation, and the visualization is shown in Fig 3.

Note that the adversarial pose generation module is only used in the training phase. During inference, the pose generation module and warping procedure are discarded and only a single-pass depth generator is used.

4 Experiments

4.1 Experiment Setup


Dataset. We adopt NYU Depth v2 dataset [K14] for both training and evaluation. NYU v2 dataset is a large-scale indoor scene understanding dataset that contains about 120K images and corresponding depth maps captured by RGB-D cameras. The testing data consists of 694 images from 215 different scenes. For training, we explore two settings, namely `all`, which means all the training data (120K) are used, and `down10`, which means we down-sample the training set by 10 times so that the total number of training data is 1.2K.

Evaluation Metrics. For quantitative evaluation we employ several commonly used metrics: the accuracy under threshold ($\delta < 1.25^i, i = 1, 2, 3$), mean absolute relative error (rel), mean log 10 error (log 10), root mean squared error (rms) and rms(log).

Method	AVCL	Data	Higher is better			Lower is better			
			$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$	rel	\log_{10}	rms	rms ₁₀
L1 Loss		down10	0.800	0.951	0.987	0.147	0.061	0.512	0.185
berHu Loss [18]		down10	0.811	0.955	0.988	0.143	0.060	0.504	0.180
DORN Loss [9]		down10	0.817	0.959	0.989	0.138	0.059	0.523	0.179
L1 Loss	✓	down10	0.810	0.959	0.989	0.141	0.060	0.525	0.182
berHu Loss [18]	✓	down10	0.822	0.960	0.989	0.121	0.059	0.515	0.177
DORN Loss [9]	✓	down10	0.827	0.963	0.990	0.120	0.053	0.505	0.176
DORN Loss [9]		all	0.828	0.965	0.992	0.115	0.051	0.509	0.176
DORN Loss [9]	✓	all	0.836	0.967	0.992	0.110	0.051	0.498	0.174

Table 1: Comparison on NYU Depth v2 between with and without AVCL under settings of different loss functions and training data.

Implementation Details. To fairly demonstrate the generality of the proposed AVCL, we follow the network architectures and loss functions in previous works when compared with them. For instance, when compared with DORN [9] we use ResNet-101 [14] and DORN loss, while compared with Deeper Depth [18] we use ResNet-50 and berHu loss. All training images are down-sampled to half size with invalid border cropped, and the predicted depth map are up-sampled to original size at the end, following [9, 21, 18]. The hyperparameter λ is set to (1, 1, 1, 1, 1) in all experiments.



δ_1	V1	V2	V3	V4
w/ AVCL	0.825	0.836	0.816	0.826
wo AVCL	0.818	0.827	0.808	0.820

Table 2: Quantitative comparison between with and without AVCL in different views.

Figure 4: Qualitative comparison between with and without AVCL in different views.

4.2 Experimental Results

In this section, we will show three significant merits of the proposed AVCL: First, with AVCL the depth generator outputs more view-consistent depth maps. Second, the demand of training data amount in AVCL is largely decreased when the same performance is reached. Last, the proposed AVCL consistently brings performance gain when combined with state-of-the-art MDE approaches [9, 18], which means AVCL is complementary to these methods. Finally, we demonstrate by an ablation study that the performance gain comes from both warping-based multi-supervision (View-Consistent Learning), and the adversarial pose generation module.

View-Consistent. We perform qualitative and quantitative analysis to validate that the proposed AVCL generates more view-consistent depth map. Firstly, we compare several views of the predicted depth map with and without AVCL in Fig. 4. As it can be observed, when the network is trained with AVCL, the output depth maps have sharper object edges, and consistently presents smaller error *w.r.t.* the ground truth across different views. With fixed

Method	Higher is better			Lower is better			
	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$	rel	\log_{10}	rms	rms _{log}
Make3D[30]	0.447	0.745	0.897	0.349	-	1.214	-
DepthTransfer[14]	-	-	-	0.35	0.131	0.12	-
Liu <i>et al.</i> [27]	-	-	-	0.335	0.127	1.06	-
Zhuo <i>et al.</i> [14]	0.525	0.838	0.962	0.305	0.122	1.04	-
Ladicky <i>et al.</i> [17]	0.542	0.829	0.941	-	-	-	-
Wang <i>et al.</i> [37]	0.605	0.890	0.970	0.220	-	0.824	-
Eigen <i>et al.</i> [9]	0.611	0.887	0.971	0.215	-	0.907	0.285
Li <i>et al.</i> [19]	0.621	0.886	0.968	0.232	0.094	0.821	-
Liu <i>et al.</i> [24]	0.650	0.906	0.976	0.213	0.087	0.759	-
Eigen <i>et al.</i> [8]	0.769	0.950	0.988	0.158	-	0.641	0.214
Li <i>et al.</i> [20]	0.788	0.958	0.991	0.143	0.063	0.635	-
Chakrabarti [2]	0.806	0.958	0.987	0.149	-	0.620	-
Laina <i>et al.</i> [18]	0.811	0.953	0.988	0.127	0.055	0.573	0.195
MS-CRF [20]	0.811	0.954	0.987	0.121	0.052	0.586	-
DORN [5]	0.828	0.965	0.992	0.115	0.051	0.509	-
DORN + AVCL	0.836	0.967	0.992	0.110	0.050	0.498	0.174

Table 3: Comparisons with existing monocular depth estimation methods on NYU Depth v2 dataset.

warping poses, we also quantitatively compute the accuracy on the testing data, and compare in Tabel. 2 the δ_1 accuracy of these views with and without AVCL. Experimental results show that, with AVCL, the accuracy of the predicted depth map in almost all the views significantly increases, which means the predicted depth maps are more consistent across different views.

Experiments with different losses. We test with several widely-used loss functions, namely L_1 loss, berHu loss [18] and DORN loss [5], to demonstrate the generality of the proposed AVCL. To use berHu loss in AVCL, it is straightforward that one only needs to replace L_1 loss in \mathcal{L}_{dep} , $\mathcal{L}_{\text{warp}}$ and \mathcal{L}_{adv} with berHu loss. As for DORN loss, the case is different. DORN transforms the ordinal regression into multiple binary classification, and takes the summation of all the predicted ones as the output depth. Therefore, there is no such thing as a continuous depth map that can be warped and applied supervision on. To address this problem, we adopt a soft approximation of the thresholding-and-summation operation by a summation of shifted sigmoid functions $\sum_i^D \text{sigmoid}(100(x - 0.5))$, where D is the number of binary classifiers. In this way, the warped depth map is continuous and differentiable, and we apply L_1 loss on the warped branch.

Table. 1 shows comparisons between different loss functions with and without AVCL. As expected, DORN loss presents better performance in terms of all the metrics than berHu loss, and the L_1 loss is the worst. Trained with AVCL, performances with all loss functions consistently increase, even the accuracy of the DORN baseline is already quite high. These results suggest that our method is complementary to these existing works on monocular depth estimation, which is reasonable because none of them considers the view-consistency. **Experiments on small training set.** In one training pass, AVCL applies two branches of supervisions to train the depth generator, which can be understood as a kind of guided data augmentation. As a consequence, the demand of training data amount is dynamically decreased in AVCL. As shown in Fig. 4, when trained with the $\frac{1}{10}$ size training set down10,

Method	Higher is better			Lower is better			
	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$	rel	\log_{10}	rms	rms_{\log}
Baseline	0.817	0.951	0.987	0.119	0.062	0.523	0.185
Fixed pose A	0.822	0.960	0.989	0.117	0.059	0.506	0.177
Fixed pose B	0.809	0.957	0.984	0.121	0.061	0.510	0.186
Fixed pose C	0.802	0.956	0.983	0.124	0.062	0.530	0.187
Random pose	0.825	0.961	0.990	0.115	0.052	0.503	0.176
Adversarial pose	0.836	0.962	0.990	0.110	0.050	0.498	0.174

Table 4: Comparisons with existing monocular depth estimation methods on NYU Depth v2 dataset.

the performance gain of AVCL is larger than trained with the entire training set. Furthermore, trained with `down10` and AVCL, the performance is almost the same level as trained with `all` and without AVCL, or even slightly better. These results show that, with the same amount of training data, the proposed AVCL framework utilizes data more efficiently, so that the model is not prone to be overfitting.

Comparison with existing methods. Tabel. 3 showcases several state-of-the art monocular depth estimation methods on NYU Depth v2 dataset. In fact, AVCL is complementary to existing methods so that the performance gain from AVCL and existing methods is addable. For instance, DORN+AVCL outperforms DORN by 0.8% δ_1 accuracy trained with all training data. This results may seem incremental, but one should note that when trained with $\frac{1}{10}$ data the gain is more remarkable, *i.e.*, from 0.810 to 0.827 with +1.7% δ_1 accuracy.

Ablation study. In order to analyze the details of the proposed AVCL, we conduct several ablation studies, and results are shown in Tabel 4. In all experiments we employ L_1 loss and use the `down10` training set. We start with a simple baseline without any bells and whistles (Baseline). Accordingly, we design two warping strategy to test the effectiveness of multi-view supervision: the first is to adopt a fixed warping pose (Fixed pose A-C), and the second is to randomly generate the warping pose (Random pose). Results show both fixed warping pose and random generated pose outperforms the baseline. However, with different fixed pose, the performance gain is sensitive to the pose, ranging from -0.015 to $+0.005$. Randomly generated poses are superior to fixed poses, since the view point is more diverse, thus more information/supervision is introduced. Finally, we add the adversarial pose generation module, so that the warping pose is expected to be the hardest under the recent state. Results show that the adversarial poses outperforms randomly generated poses, which proves the effectiveness and necessity of the adversarial pose generation module.

5 Conclusion

In this paper, we empirically find existing MDE approaches often result in sub-optimal solution that the output depth maps present inconsistent losses viewed from different directions. We therefore propose an Adversarial View-Consistent Learning framework, force the estimated depth map to be all reasonable viewed from multiple views. To this end, we design a differentiable depth map warping operation to generate multiple views of the depth map given a transformation pose, and meanwhile propose an adversarial pose generator to generate the hard poses. Experiments show learning from multiple views leads to more view-consistent predictions, and the adversarial generated poses further increases the accuracy.

References

- [1] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018.
- [2] Ayan Chakrabarti, Jingyu Shao, and Greg Shakhnarovich. Depth from a single image by harmonizing overcomplete local network predictions. In *NIPS*. 2016.
- [3] David Eigen and Rob Fergus. Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In *CVPR*, 2015.
- [4] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *NIPS*. 2014.
- [5] Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, and Dacheng Tao. Deep ordinal regression network for monocular depth estimation. In *CVPR*, 2018.
- [6] Yasutaka Furukawa, Carlos Hernández, et al. Multi-view stereo: A tutorial. *Foundations and Trends® in Computer Graphics and Vision*, 9(1-2):1–148, 2015.
- [7] Silvano Galliani, Katrin Lasinger, and Konrad Schindler. Massively parallel multi-view stereopsis by surface normal diffusion. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 873–881, 2015.
- [8] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NIPS*, 2014.
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [10] Derek Hoiem, Alexei A Efros, and Martial Hebert. Recovering surface layout from an image. *International Journal of Computer Vision*, 75(1):151–172, 2007.
- [11] Rui Huang, Shu Zhang, Tianyu Li, and Ran He. Beyond face rotation: Global and local perception gan for photorealistic and identity preserving frontal view synthesis. In *ICCV*, 2017.
- [12] Max Jaderberg, Karen Simonyan, Andrew Zisserman, et al. Spatial transformer networks. In *NIPS*, 2015.
- [13] Mengqi Ji, Juergen Gall, Haitian Zheng, Yebin Liu, and Lu Fang. SurfacerNet: An end-to-end 3d neural network for multiview stereopsis. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2307–2315, 2017.
- [14] K. Karsch, C. Liu, and S. B. Kang. Depth transfer: Depth extraction from video using non-parametric sampling. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(11):2144–2158, 2014.
- [15] Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? In *nips*. 2017.

- [16] Kiriakos N Kutulakos and Steven M Seitz. A theory of shape by space carving. *International journal of computer vision*, 38(3):199–218, 2000.
- [17] Lubor Ladicky, Jianbo Shi, and Marc Pollefeys. Pulling things out of perspective. In *CVPR*, 2014.
- [18] Iro Laina, Christian Rupprecht, Vasileios Belagiannis, Federico Tombari, and Nassir Navab. Deeper depth prediction with fully convolutional residual networks. In *2016 Fourth international conference on 3D vision (3DV)*, 2016.
- [19] Bo Li, Chunhua Shen, Yuchao Dai, Anton van den Hengel, and Mingyi He. Depth and surface normal estimation from monocular images using regression on deep features and hierarchical crfs. In *CVPR*, 2015.
- [20] Jun Li, Reinhard Klein, and Angela Yao. A two-streamed network for estimating fine-scaled depth maps from single rgb images. In *ICCV*, 2017.
- [21] Jun Li, Reinhard Klein, and Angela Yao. A two-streamed network for estimating fine-scaled depth maps from single rgb images. In *ICCV*, 2017.
- [22] B. Liu, S. Gould, and D. Koller. Single image depth estimation from predicted semantic labels. In *cvpr*, 2010.
- [23] Ce Liu, Jenny Yuen, Antonio Torralba, Josef Sivic, and William T. Freeman. Sift flow: Dense correspondence across different scenes. In David Forsyth, Philip Torr, and Andrew Zisserman, editors, *ECCV*, 2008.
- [24] F. Liu, C. Shen, G. Lin, and I. Reid. Learning depth from single monocular images using deep convolutional neural fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(10):2024–2039, 2016.
- [25] Fayao Liu, Chunhua Shen, Guosheng Lin, and Ian Reid. Learning depth from single monocular images using deep convolutional neural fields. *IEEE transactions on pattern analysis and machine intelligence*, 38(10):2024–2039, 2016.
- [26] Fayao Liu, Chunhua Shen, Guosheng Lin, and Ian Reid. Learning depth from single monocular images using deep convolutional neural fields. *IEEE transactions on pattern analysis and machine intelligence*, 38(10):2024–2039, 2016.
- [27] Miaomiao Liu, Mathieu Salzmann, and Xuming He. Discrete-continuous depth estimation from a single image. In *CVPR*, 2014.
- [28] Michael Opitz, Georg Waltner, Horst Possegger, and Horst Bischof. Deep metric learning with bier: Boosting independent embeddings robustly. *IEEE transactions on pattern analysis and machine intelligence*, 2018.
- [29] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.
- [30] Anirban Roy and Sinisa Todorovic. Monocular depth estimation using neural regression forest. In *CVPR*, 2016.

- [31] A. Saxena, M. Sun, and A. Y. Ng. Make3d: Learning 3d scene structure from a single still image. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(5): 824–840, 2009.
- [32] Ashutosh Saxena, Sung H Chung, and Andrew Y Ng. Learning depth from single monocular images. In *NIPS*, pages 1161–1168, 2006.
- [33] Steven M Seitz and Charles R Dyer. Photorealistic scene reconstruction by voxel coloring. *International Journal of Computer Vision*, 35(2):151–173, 1999.
- [34] Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor segmentation and support inference from rgb-d images. In *ECCV*, 2012.
- [35] Engin Tola, Christoph Strecha, and Pascal Fua. Efficient large-scale multi-view stereo for ultra high-resolution image sets. *Machine Vision and Applications*, 23(5):903–920, 2012.
- [36] Luan Tran, Xi Yin, and Xiaoming Liu. Disentangled representation learning gan for pose-invariant face recognition. In *CVPR*, 2017.
- [37] Peng Wang, Xiaohui Shen, Zhe Lin, Scott Cohen, Brian Price, and Alan L. Yuille. Towards unified depth and semantic prediction from a single image. In *CVPR*, 2015.
- [38] Xiaolong Wang, Abhinav Shrivastava, and Abhinav Gupta. A-fast-rcnn: Hard positive generation via adversary for object detection. In *CVPR*, 2017.
- [39] Junyuan Xie, Ross Girshick, and Ali Farhadi. Deep3d: Fully automatic 2d-to-3d video conversion with deep convolutional neural networks. In *ECCV*, 2016.
- [40] Dan Xu, Elisa Ricci, Wanli Ouyang, Xiaogang Wang, and Nicu Sebe. Multi-scale continuous crfs as sequential deep networks for monocular depth estimation. In *CVPR*, 2017.
- [41] Yao Yao, Zixin Luo, Shiwei Li, Tian Fang, and Long Quan. Mvsnet: Depth inference for unstructured multi-view stereo. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 767–783, 2018.
- [42] Tinghui Zhou, Matthew Brown, Noah Snavely, and David G Lowe. Unsupervised learning of depth and ego-motion from video. In *CVPR*, 2017.
- [43] Jun-Yan Zhu, Philipp Krähenbühl, Eli Shechtman, and Alexei A Efros. Generative visual manipulation on the natural image manifold. In *ECCV*, 2016.
- [44] Wei Zhuo, Mathieu Salzmann, Xuming He, and Miaomiao Liu. Indoor scene structure analysis for single image depth estimation. In *CVPR*, 2015.