

Robust Geometry-Preserving Depth Estimation Using Differentiable Rendering

Chi Zhang^{1*}, Wei Yin^{2*}, Gang Yu^{1†}, Zhibin Wang¹, Tao Chen³,
 Bin Fu¹, Joey Tianyi Zhou⁵, Chunhua Shen⁴

¹ Tencent ² DJI Technology ³ Fudan University ⁴ Zhejiang University

⁵ Centre for Frontier AI Research, A*STAR ⁵ Institute of High Performance Computing, A*STAR

¹{johnczhang, skicyyu, billzbwang, brianfu}@tencent.com

²yvanwy@outlook.com ³eetchen@fudan.edu.cn ⁴chunhua@me.com ⁵joey.tianyi.zhou@gmail.com

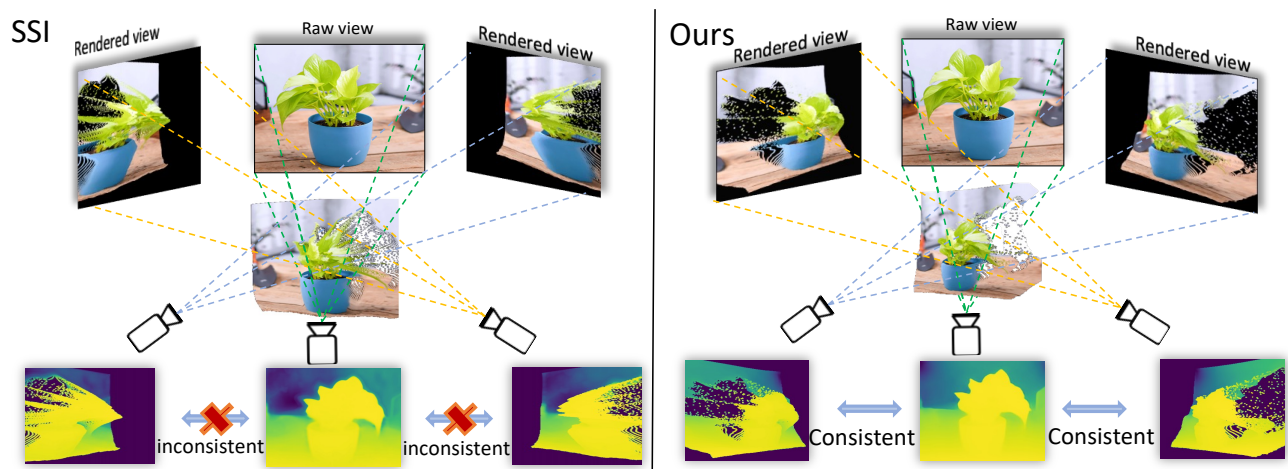


Figure 1: An illustration of our motivation. Zero-shot depth estimation models trained with scale-and-shift invariant (SSI) loss [25] produce geometry-incomplete depth prediction, which results in distorted 3D structures (left). A good 3D model should look realistic across different views, such that the depth estimation results of them are consistent. Based on this intuition, we render novel views of the generated 3D structure and design loss functions to promote the consistency of depth estimation across different views, which produces realistic 3D structures (right).

Abstract

In this study, we address the challenge of 3D scene structure recovery from monocular depth estimation. While traditional depth estimation methods leverage labeled datasets to directly predict absolute depth, recent advancements advocate for mix-dataset training, enhancing generalization across diverse scenes. However, such mixed dataset training yields depth predictions only up to an unknown scale and shift, hindering accurate 3D reconstructions. Existing solutions necessitate extra 3D datasets or geometry-complete depth annotations, constraints that limit their versatility. In this paper, we propose a learning framework that

trains models to predict geometry-preserving depth without requiring extra data or annotations. To produce realistic 3D structures, we render novel views of the reconstructed scenes and design loss functions to promote depth estimation consistency across different views. Comprehensive experiments underscore our framework’s superior generalization capabilities, surpassing existing state-of-the-art methods on several benchmark datasets without leveraging extra training information. Moreover, our innovative loss functions empower the model to autonomously recover domain-specific scale-and-shift coefficients using solely unlabeled images.

*Equal contributions.

†Corresponding author.

1. Introduction

Recovering 3D geometries of scenes from monocular images has become an area of significant interest, driven by recent advances in monocular depth estimation, with wide-ranging applications such as 3D photography [31]. 3D scene recovery [45, 44] of in-the-wild monocular images relies on a powerful depth estimator [24, 48, 45, 25, 49] that can accurately predict the geometry of diverse scenes. State-of-the-art depth estimation models [25, 45, 24, 49] now advocate mix-dataset training [43, 25], which can generate robust depth predictions across diverse scenes. This opens up the possibility of large-scale pre-training and deploying only a single model in various application scenarios.

To enable mix-dataset training, scale-and-shift invariant (SSI) losses [25, 45] are designed to normalize depth representations explicitly, thereby removing scale-and-shift changes between different data sources. As a result, datasets with various depth representations, such as metric depth, uncalibrated disparity maps, and relative depth up to scale (UTS), can be jointly utilized for training. However, despite the strong generalization capabilities across scenes, mix-dataset training comes with its own set of drawbacks. Unlike previous depth estimation models that produce absolute depth or relative depth up to scale, which can be directly unprojected to 3D structures given the intrinsic camera parameters, models trained with SSI losses predict depth up to unknown scale and shift factors (UTSS), which is geometrically incomplete [22] for reconstructing 3D models. Although scaling depth maps typically adheres to the original 3D scene recovery’s geometric integrity, the unknown shifts may introduce structural distortions. Depth estimation models that are optimized to produce absolute depth or relative depth up to scale do not suffer from this problem, but they require geometry-complete depth annotations, such as metric depth or relative depth from multi-view stereo, for learning. Leres [45] offers a potential solution by rectifying the distorted point cloud via a separately optimized post-processing module, which, however, necessitates additional 3D datasets. Unfortunately, the extra 3D data or geometry-complete depth annotations are significantly less diverse than the geometrically incomplete data used in the original mix-data training, and as a result, their generalization ability on in-the-wild images is limited.

In this research, our primary objective is to develop depth estimation models that can predict geometry-preserving depth up to a scale for 3D scene recovery without requiring extra data or annotations through mix-dataset training. To achieve this goal, we propose a novel framework based on differentiable rendering. Specifically, we reconstruct 3D point clouds based on the predicted depth and use a differentiable renderer to generate novel views of the 3D model. We then predict the depth of the synthesized views

with the same model and employ loss functions to ensure that the depth predictions of the rendered views are consistent. Fig. 1 illustrates our motivation. In this process, the network is optimized to produce undistorted 3D structures from depth using the informative gradients from the differentiable renderer. This ensures that the rendered images from different views look realistic and their depth estimations are consistent. Compared with previous works, our method can produce geometry-preserving predictions without relying on extra annotations or 3D datasets, enabling us to make full use of mixed datasets collected from various resources to improve generalization. Our loss functions can also recover domain-specific scale and shift coefficients of a trained UTSS model, such as Midas [25] and HDN [49], in a self-supervised manner using unlabelled images from the same domain. Moreover, we demonstrate that our proposed self-supervised loss can be used to predict intrinsic camera parameters, such as focal length, by selecting the parameter from a few options that minimize the proposed consistency losses. Our extensive experiments on multiple benchmark datasets validate the effectiveness of our design. Our main contributions are summarized as follows:

- We propose a novel depth estimation learning framework that can produce geometry-preserving depth without relying on extra datasets or annotations.
- Our proposed consistency loss can recover domain-specific affine coefficients of a trained model in a self-supervised manner using unlabelled images from the same domain.
- We demonstrate that the self-supervised loss can also be used to roughly estimate camera intrinsic parameters.
- Experiments on multiple benchmark datasets show that our method can better recover scene structures of diverse images both quantitatively and qualitatively.

2. Related Work

In this section, we review related works on monocular depth estimation based on deep neural networks and literature on 3D scene reconstruction from single-view images.

Monocular depth estimation. Deep neural networks have made significant progress in many computer vision tasks, including monocular depth estimation (MDE) [46, 42, 20, 3, 4, 47, 8, 9], which is the focus of this paper. Most deep-learning-based MDE models aim to learn a pixel-wise depth predictor by fitting a labeled training set. However, recent works [25, 45, 8, 5, 36] have shown that such models can be domain-sensitive and have poor generalization capability across datasets due to dataset bias. Moreover, since collecting diverse labeled depth data at scale is difficult, the training datasets [5, 47, 43, 37] are often small,

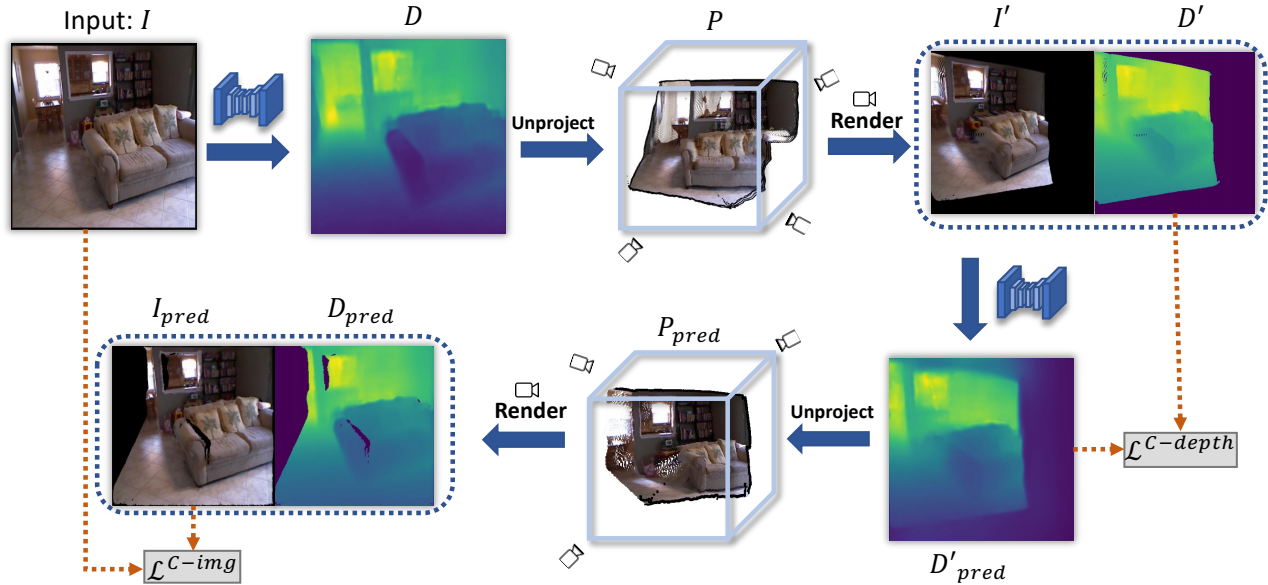


Figure 2: Overview of our framework for geometry-preserving depth estimation. Given an input image, we reconstruct the point cloud based on the depth estimation of the image. Then, a new view of the 3D structure is rendered and the depth is estimated again using the same model. We then reconstruct the point cloud based on the depth estimation of the new view and render it back to the original view. Finally, loss functions are employed to promote consistency between the outputs from different views.

in sharp contrast to many other vision tasks, such as image segmentation [17, 18, 41, 16] and object detection [16, 30], which consistently benefit from the increasing number of diverse training data. Recent studies [40, 25, 5, 24, 45, 49] advocate mix-dataset training, which allows datasets from various sources to be jointly utilized for training. In particular, it is cheap to collect diverse stereo images at scale from the Internet [36] or 3D movies [25]. However, as the stereo cameras are not calibrated, the disparity map generated by optical flow algorithms can only recover inverse depths up to unknown scales and shifts. Therefore, there exists a scale-and-shift gap between data annotations from various sources.

Self-supervised learning has been widely investigated in the area of depth estimation and 3D reconstruction [50, 11, 21]. Compared with previous works, our focus is on generating robust 3D point clouds from monocular depth estimation through mix-data training, where we do not use additional information, such as video sequences or stereo images for self-supervision, and only a single monocular image is provided.

To enable mix-data training with various forms of annotations, such as metric depth [32, 10, 7] obtained from RGB-D cameras, disparity map from stereo matching [2, 6], and relative depth from multi-view stereo [19], state-of-the-art works [24, 25, 45, 49] normalize the depth annotations explicitly, such that the scale-and-shift changes between depth annotations can be removed. In particular, HDN [49] proposes to normalize depth annotations hierarchically to

improve scale-and-shift invariant losses. A significant advantage of mix-dataset training is that the learned model can be directly evaluated on various benchmarks without using their individual training sets, enabling zero-shot transfer [25]. However, the predicted depth is up to an unknown scale and shift, which needs to be aligned with the ground truth using least squares to find the optimal scale and shift, such that the standard evaluation metrics can be used for comparison. Alternatively, the zero-shot model can be fine-tuned with the training set of each benchmark for the evaluation of metric depth [25].

Recover 3D structure based on depth. Depth estimation is a crucial step in recovering the structure of a scene from monocular images. Our goal is to leverage the impressive generalization ability of zero-shot depth estimation models for 3D scene reconstruction in real-world settings. While mix-dataset training has yielded good results in depth estimation, the learned model cannot be used directly for 3D scene reconstruction. This is because an unknown shift in depth can distort the structures. LeRes [44] addresses this problem by training a separate rectification module to correct the point cloud generated from raw depth maps. This module learns to shift the raw point cloud and predict the focal length given an initial value. However, this module requires an additional 3D point cloud dataset for training, which is hard to obtain for outdoor scenes and thus limits generalization across scenes. GP2 [22] demonstrates that using an extra scale-invariant loss on a portion of the dataset can help the model output geometry-preserving depth pre-

dictions that are automatically shifted. Adding a scale-invariant loss is the most straightforward and effective strategy to train a scale-invariant depth predictor. However, this loss requires extra geometry-complete depth annotations for supervision, which largely limits the model’s generalization ability in diverse scenes. Consequently, the learned model is still domain-specific, which is contrary to the goal of learning a generic zero-shot model across scenes. Additionally, GP2 [22] cannot estimate the focal length, which is essential for 3D reconstruction.

In contrast, our framework presented in this paper is complementary to previous mix-dataset training pipelines and can directly train a model that makes geometry-preserving depth predictions up to a scale without seeking extra datasets or annotations, which shows a strong generalization capability.

Other 3D reconstruction in the literature. While our work focuses on 3D shape recovery from single images through depth estimation, there are various other methods that have been proposed to address this problem. Early works in this field relied mostly on hand-crafted features, such as local segments, shadings, and edges, as priors [13, 23, 28]. More recent approaches, such as those proposed in [38, 26, 39, 27], are data-driven and use an end-to-end learning approach. However, these methods are typically limited to specific types of scenes or objects, such as human bodies [27], faces, cars, *etc.*, and require different forms of supervision, such as ground truth 3D structures represented by a mesh [35]. *One key advantage of our framework is that it is scene-agnostic, making it applicable to any real-world images.*

3. Preliminary

Before introducing our framework, we first present some preliminary concepts in monocular depth estimation and 3D scene reconstruction from depth.

Scale-and-shift invariant loss. To train a depth estimator that can be applied to various scenes, we utilize the scale-and-shift invariant loss [25] that eliminates the scale and shift variances between different data representations. Specifically, we first estimate the shift μ_D and scale σ_D for a given depth map D and a binary mask M that indicates valid pixels, as follows:

$$\mu_D = \text{median}(D), \quad \sigma_D = \frac{1}{|M|} \sum_{i=1}^{|M|} |D(i) - \mu_D|, \quad (1)$$

where $D(i)$ is the depth value of a pixel location i , and $|M|$ is the number of valid pixels. We then remove the scale and shift for both the predicted depth map and the ground truth annotations, as shown below:

$$\hat{D} = \frac{D - \mu_D}{\sigma_D}, \quad \hat{D}^* = \frac{D^* - \mu_{D^*}}{\sigma_{D^*}}. \quad (2)$$

Next, we apply a standard ℓ_1 loss between the normalized depth representations to supervise the training, as follows:

$$\mathcal{L}^{\text{SSI}} = |\hat{D} - \hat{D}^*|. \quad (3)$$

This ensures that any affine changes to the predicted depth maps or the annotations will not affect the losses. Thus, uncalibrated disparity maps can be used for training.

3D reconstruction from depth maps. Given the predicted depth D , we can reconstruct the point cloud P from the image coordinate based on the pinhole camera model, as shown below:

$$\begin{cases} x = \frac{u - u_0}{f} d \\ y = \frac{v - v_0}{f} d \\ z = d \end{cases}, \quad (4)$$

where (x, y, z) and (u, v) are the 3D and 2D coordinates, (u_0, v_0) is the optical center, and f is the focal length. As we can see, scaling the depth leads to uniform changes to the scene, while shifting causes non-uniform changes, which leads to distortion. Since the predicted depth maps are up to unknown scale and shift coefficients, the reconstructed 3D structure from them is likely to be distorted from inappropriate affine changes.

4. Method

This section presents a detailed description of our framework. Our framework comprises three crucial steps: depth estimation, point cloud reconstruction, and differentiable rendering of new views. Once we obtain a novel view of the structure, we repeat the same operations on the generated image and render it back to the original view. We then employ loss functions to promote consistency between the multi-view predictions. An overview of our framework is depicted in Fig. 2.

4.1. Pipeline

Render a novel view of a point cloud. To obtain a novel view of a point cloud, we begin by reconstructing the point cloud P using the predicted depth map D of the input image I , as specified in Eq. (4). Next, we render a depth map D' and an RGB image I' of the point cloud from a different perspective. Specifically, we rotate the camera horizontally by an angle θ and shift the camera along the z axis by a distance t . To ensure that the scene remains in the rendered image, we rotate the camera around $T_{\text{center}} = [0, 0, \min_z(P)]^T$ instead of rotating around the origin of coordinates. Here, $\min_z(P)$ returns the smallest depth value along the z axis. In this way, the structure always appears in the middle of the rendered image during rotation. More specifically, the rotation matrix R and transition matrix T are given below:

$$R = \begin{bmatrix} \cos(\theta) & 0 & \sin(\theta) \\ 0 & 1 & 0 \\ -\sin(\theta) & 0 & \cos(\theta) \end{bmatrix}, \quad T = \begin{bmatrix} 0 \\ 0 \\ t \end{bmatrix}. \quad (5)$$

Here, the shift t and rotation angle θ are randomly sampled from $[-\min_z(P), 2 \cdot \min_z(P)]$ and $[-30^\circ, 30^\circ]$, respectively. Finally, we use the differentiable renderer [34], which is a function of the point cloud data, the rotation matrix, and the transition matrix, to render the new image I' and the depth map D' as follows:

$$I', D' = \text{Render}(P - T_{\text{center}}, R, T + T_{\text{center}}), \quad (6)$$

To implement the change of the rotation center conveniently, we shift the point cloud data using $P - T_{\text{center}}$ and compensate for the shift by adding it to the shift matrix after rotation: $T + T_{\text{center}}$.

Render the raw view. Given the rendered image I' and the depth map D' of new views, we can recover the same structure P and accurately render back the raw view, except for occluded regions. Based on this observation, we estimate the depth of the rendered image I' with the same model to obtain D'_{pred} , and use it together with I' to render back the raw view. Our intuition is that we can achieve accurate recovery of the structure and rendering of the raw view only if the predicted depth D'_{pred} of the rendered image matches the rendered depth map D' . To achieve this goal, both the reconstructed structure P and the rendered image I' must be geometrically accurate and visually realistic. Therefore, we render back the raw view based on I' and D'_{pred} . However, due to the scale-and-shift changes of the depth predictions output by the model trained with SSI loss, we need to align D'_{pred} with D' using the least squares method to ensure consistency:

$$(a, b) = \arg \min_{a, b} \sum_{i=1}^{|M'|} (aD'_{\text{pred}}(i) + b - D'(i)), \quad (7)$$

$$\hat{D}'_{\text{pred}} = aD'_{\text{pred}} + b. \quad (8)$$

Specifically, we solve for a and b in the equation above, where $|M'|$ is the number of valid pixel locations. We can then reconstruct the structure again based on I' and \hat{D}'_{pred} to obtain P_{pred} . Finally, we render the depth map D_{pred} and the image I_{pred} of the raw view based on the point cloud P_{pred} :

$$I_{\text{pred}}, D_{\text{pred}} = \text{Render}(P_{\text{pred}}, R_{\text{inv}}, T_{\text{inv}}), \quad (9)$$

$$R_{\text{inv}} = R^T, \quad (10)$$

$$T_{\text{inv}} = -R^T(T + T_{\text{center}}) + T_{\text{center}}. \quad (11)$$

4.2. Loss Functions

Consistency Loss. To ensure consistency during the rendering process, we employ two loss functions: the depth consistency loss $\mathcal{L}^{\text{C-depth}}$ and the image consistency loss $\mathcal{L}^{\text{C-img}}$. The overall consistency loss is defined as follows:

$$\mathcal{L}^{\text{C}} = \mathcal{L}^{\text{C-depth}} + \alpha \mathcal{L}^{\text{C-img}}, \quad (12)$$

where α is a weight term. For the image consistency loss, we compute the mean pixel-wise ℓ_1 loss (L1) between the rendered image I_{pred} and the original input image I over the valid region M_{img} : $\mathcal{L}^{\text{C-img}} = \text{L1}_{M_{\text{img}}}(I_{\text{pred}}, I)$. For the depth consistency loss, we apply a scale-and-shift invariant loss (SSI) in Eq. 3 between the rendered depth map D' of the novel view and the new prediction D'_{pred} over the valid regions M_{depth} : $\mathcal{L}^{\text{C-depth}} = \text{SSI}_{M_{\text{depth}}}(D', D'_{\text{pred}})$. Here, M_{img} and M_{depth} include the pixels that are not occluded during rendering.

Multi-Focal-Length (MFL) losses. Our pipeline and loss computation rely on knowing the focal length f of the input image to unproject the depth map. Although the intrinsic camera parameters are provided in some datasets, they are often unknown in many cases, such as web images. Availability of focal length can definitely make our design more robust but it meanwhile limits the training in the mix-dataset scenario. Therefore, we assume that the focal length of training images is unknown during training, which allows us to train our algorithm with more diverse data. To handle the unknown and varied focal length of training images, we further develop a multi-focal-length loss for our algorithm. A focal length f can be roughly estimated from the horizontal field of view (FOV), such as 60° , by

$$f = \frac{W}{2 \cdot \tan(\text{FOV}/2)}, \quad (13)$$

where W is the width of the image. Our proposed consistency loss based on the focal length f is denoted as \mathcal{L}_f^{C} , and we can compute multiple consistency losses by selecting different focal length values from a set \mathcal{F} that corresponds to different FOVs. At training time, we only keep the minimum one as the final loss of each input image $\mathcal{L}_{f^*}^{\text{C}}$ for optimization, where

$$f^* = \arg \min_{f \in \mathcal{F}} \mathcal{L}_f^{\text{C}}. \quad (14)$$

Intuitively, the appropriate focal length can more accurately reconstruct the structure and thus leads to small rendering losses. Finally, we add the raw scale-and-shift invariant loss between the original predicted depth map D and its ground truth D^* , and the final loss is given by:

$$\mathcal{L} = \mathcal{L}^{\text{SSI}} + \beta \mathcal{L}_{f^*}^{\text{C}}, \quad (15)$$

where β is a hyper-parameter to balance two terms.

5. Experiments

We conduct extensive experiments on multiple benchmarks to validate the effectiveness of our algorithms, including evaluations of the geometry-preserving depth estimation and the accuracy of the recovered point clouds. Additionally, we conduct several ablation studies to analyze each component of our design.

Method	Extra Info.	NYU		ScanNet		ETH3D		KITTI	
		AbsRel↓	δ_1 ↑	AbsRel↓	δ_1 ↑	AbsRel↓	δ_1 ↑	AbsRel↓	δ_1 ↑
SSI	-	27.9	52.5	28.5	51.8	26.3	62.5	21.0	63.9
SSI + PCM [45]	3D dataset	13.5	83.4	12.4	85.2	20.6	71.2	33.7	39.2
SSI + Ours	-	10.5	89.2	12.2	85.4	11.3	87.1	12.9	81.4
GP2 [22]	Metric Depth	9.8	90.3	11.4	87.2	14.0	85.4	13.4	81.2
GP2 + Ours	Metric Depth	9.2	91.2	10.9	88.0	10.9	88.4	11.6	84.9

Table 1: Comparison of different methods for geometry-preserving depth estimation with scale alignment only. Our method outperforms previous methods without using any additional information.

Method	KITTI	2D3D	NYU
		RMSE↓	
PCM [45]	7.05	0.67	0.38
RenderFOV w/ $\mathcal{L}^{C\text{-img}}$	5.95	0.72	0.43
RenderFOV w/ $\mathcal{L}^{C\text{-depth}}$	3.04	0.48	0.36

Table 2: Comparison of different methods for point cloud reconstruction. Our method with $\mathcal{L}^{C\text{-depth}}$ for selecting the FOV has the optimal results on indoor and outdoor benchmarks.

$\mathcal{L}^{C\text{-img}}$	$\mathcal{L}^{C\text{-depth}}$	MFL	AbsRel↓
✓	✓	✓	11.7
✓	✓		13.3
✓		✓	13.4
	✓	✓	12.2
			26.0

Table 3: Ablation study on loss functions. We report the mean AbsRel over four benchmarks for evaluation. Our designs can provide remarkable performance improvement.

Method	Eval	NYU	KITTI
		AbsRel↓	
SSI	SSA	9.3	14.0
SSI	SA	27.9	21.0
SSL w/ $\mathcal{L}^{C\text{-img}}$	SA	10.8	13.3
SSL w/ $\mathcal{L}^{C\text{-depth}}$	SA	10.8	13.1

Table 4: Self-supervised learning (SSL) experiments. We only learn the affine coefficients to scale and shift the prediction of a model trained by the SSI loss with the unlabelled training images in individual benchmarks. **SSA** and **SA** denote scale-and-shift alignment and scale alignment, respectively, for evaluation.

Implementation details. We follow Leres [45] and GP2 [22] to construct datasets for mix-dataset training, which contain 121K images from DIML [14] dataset, 114K images from taskonomy [47] dataset, 48K images from Holopix50K [12], and 20K images from HRWSI [37] dataset. We withhold 200 images from each dataset for validation. For ablation study and analysis, we sample a subset of 16,000 images evenly from different datasets to train the models. All baselines in our experiments and our model em-

ploy the DPT [24] depth estimation network, which is currently state-of-the-art, with an input image size of 384×384 pixels. During training, we apply random cropping and horizontal flipping for data augmentation. Our training process consists of two stages: first, we fully optimize our model for 20 epochs using the SSI loss, as in previous mix-dataset training methods. Second, we finetune the model by adding our proposed losses for 2 additional epochs. Following DPT [24], we use the Adam optimizer [15], with a learning rate of 10^{-5} and the batch size of 16. We construct mini-batches by sampling equal numbers of data from each dataset. The weight terms α and β are set to 1.0 and 0.1, respectively. We select three different fields of view (FOV) from $\{50^\circ, 60^\circ, 70^\circ\}$ during training and compute their corresponding focal length values based on Eq. 13 to construct the focal length set \mathcal{F} .

Evaluation Metrics. We evaluate our model on five benchmark datasets, including NYU V2 [32], ScanNet [7], KITTI [10], ETH3D [29], and 2D3D [1]. For evaluation of geometry-preserving depth estimation, we follow Leres [45] to first align the scale of the predicted depth map D and the ground truth D^* by multiplying the prediction with a factor s , which is computed by:

$$s = \text{median}(D^*/D).$$

We use two common metrics to evaluate the accuracy: the absolute relative error (AbsRel):

$$(1/|M|) \sum_{i=1}^{|M|} |D(i) - D(i)^*|/D(i)^*.$$

and the percentage of pixels with

$$\delta_1 = \max(D(i)/D^*(i), D^*(i)/D(i)) < 1.25.$$

For evaluation of the reconstructed point clouds, we compute the Root Mean Square Error (RMSE) between the point clouds unprojected from the aligned depth prediction and the ground-truth following [22].

5.1. Results

Evaluation of geometry-preserving depth estimation.

To demonstrate the advantages of our design, we compare our model with several existing methods, including:

Method	NYU		KITTI		ScanNet		ETH3D	
	AbsRel↓	δ_1 ↑	AbsRel↓	δ_1 ↑	AbsRel↓	δ_1 ↑	AbsRel↓	δ_1 ↑
MegaDepth [19]	19.4	71.4	20.1	66.3	26.0	64.3	39.8	52.7
WSVD [33]	22.6	65.0	24.4	60.2	18.9	71.4	26.1	61.9
DiverseDepth [40]	11.7	87.5	19.0	70.4	10.8	88.2	22.8	69.4
MiDaS [25]	11.1	88.5	23.6	63.0	11.1	88.6	18.4	75.2
Leres [45]	9.0	91.6	14.9	78.4	9.5	91.2	17.1	77.7
DPT [†] [24]	8.8	92.7	12.7	84.9	9.6	91.3	16.1	77.6
Ours (SA)	7.7	93.6	11.1	86.2	9.0	91.8	9.1	91.4

Table 5: Comparison with the state-of-the-art depth estimation models on zero-shot benchmarks. Our geometry-preserving model with only scale alignment (SA) significantly outperforms previous works with scale-and-shift alignment for evaluation. [†] denotes our re-implementation with our data.

1) **SSI**. Our first baseline is the model trained with the standard SSI loss. As the output from it is up to unknown scale and shift coefficients, we can observe the influence when the shift is omitted.

2) **SSI + PCM**. We incorporate the point cloud module (PCM) from Leres [45] into our SSI baseline to predict the shifts required to rectify the distorted point clouds generated by depth estimation models. In this analysis, we used their released PCM model to shift the depth maps generated by our SSI baseline.

3) **GP2**. Although our learning approach focuses on developing geometry-preserving depth estimators without relying on geometry-complete depth annotations, we also compared our method to GP2, which directly utilizes scale-invariant (SI) losses, requiring geometry-complete depth annotations. Since the metric depth annotations were only available in the Taskonomy dataset, we applied the SI loss to the data from this dataset in each mini-batch during training. The comparison of the results is presented in Table 1.

The results indicate that directly omitting the shift term of the SSI baseline resulted in poor performance across all benchmarks, implying that the model trained with SSI alone cannot be used to reconstruct 3D structures directly. Our proposed design outperforms PCM [45] across all benchmarks, especially in those with outdoor images such as KITTI and ETH3D. This is because PCM is trained with 3D indoor datasets and is only effective in indoor datasets, such as NYU and ScanNet. Utilizing scale-invariant losses with metric depth labels can effectively improve the performance on indoor benchmarks, but is less effective than our model on outdoor benchmarks. In particular, our result in terms of AbsRel outperforms GP2 by 19.2% on the ETH3D dataset. Our design can further improve the performance when they are combined, which generates optimal results across all benchmarks.

Evaluation of reconstructed point clouds. We next evaluate the accuracy of the reconstructed point clouds. Given the predicted depth map, the focal length of the image needs to be estimated for point cloud reconstruction. Here we use the same depth estimation model learned by

our losses and compare the following strategies to generate the focal length values.

1) **PCM** from Leres [45]. In Leres [45], they reconstruct the point clouds using an initial focal length value, which corresponds to a FOV of 60°. Then, the PCM takes the point cloud as input and predicts shifts of the points, as well as focal length values. These values are used to reconstruct the point cloud again for evaluation.

2) **RenderFOV**. We use the proposed consistency losses as an indicator to choose the focal length values from a few selections during testing. Here we provide multiple FOV values from {40°, 50°, 60°, 70°, 80°}, and compute their corresponding focal length values. We render multiple images to compute the mean consistency losses, by selecting the camera shifts t from $\{-0.5 \cdot \min_z(P), 0.5 \cdot \min_z(P)\}$ and rotation angles θ from $\{-20^\circ, 20^\circ\}$. We choose the focal length value that produces minimal consistency loss. We report results using the image consistency loss and depth consistency losses, respectively.

Table 2 shows the comparison of the methods. We find that the depth consistency loss is more effective than the image consistency loss for selecting the appropriate focal length values. Therefore, using depth consistency loss becomes our default choice for point cloud reconstruction. Our proposed method outperforms the previous method PCM [45], particularly in the outdoor dataset KITTI.

Self-supervised learning for scale-and-shift recovery. We explore self-supervised learning with our framework using unlabelled images. Here we fix the parameters of a fully trained model with SSI loss, and learn the affine parameters to transform the raw outputs up to scale and shift coefficients to geometry-preserving outputs up to a scale. We observe the optimal scales and shifts of different benchmarks are different, hence we learn a group of affine parameters for each benchmark based on their individual training sets. Specifically, we conduct experiments on two datasets: NYU, an indoor dataset containing 795 training images, and KITTI, an outdoor dataset containing 21,591 training images. We also report the results of the pre-trained model with scale-and-shift alignment for evaluation, which can be




Figure 3: Qualitative comparison of point cloud reconstruction. We render a new view of the reconstructed point clouds. Adding our loss in training can effectively improve the baseline.

considered as the upper bound for the results of our models with only scale alignment. The results are presented in Table 4.

Our self-supervised learning with both losses yields good results on both benchmarks, indicating the effectiveness of our proposed losses and suggesting that predictions of images from the same domain have similar affine coefficients. Notably, our self-supervised learning models on the KITTI dataset even outperform the SSI pre-trained model with scale-and-shift alignment. This result suggests that using least squares to find the affine coefficients is less effective

than our approach in this case. One possible explanation is that least-squares regression requires accurate predictions for each image to obtain the affine parameters for alignment. When the raw prediction is poor, the errors may be amplified by alignment, resulting in sub-optimal results. In contrast, optimizing the affine parameters with our losses based on the images from the same domain yields better results.

Qualitative results. We next compare the 3D structures reconstructed by the model trained with SSI loss and our losses through visualization in Fig. 3. Specifically, we show

the rendered images of novel views. Our proposed losses can effectively eliminate distortions in the 3D models, as shown in the visualization.

Comparison with the state-of-the-art methods. We compare our best model with state-of-the-art methods on four benchmark datasets: NYU, KITTI, ScanNet, and ETH3D. We use all the training data in our mixed datasets and add the scale-invariant loss [22] during training. As the training data of DPT [24] are not publicly available, we re-implement DPT with our datasets and training hyperparameters. Previous methods align the scale and the shift for evaluation, while our methods only require alignment of the scale. As shown in Table 5, our method outperforms the state-of-the-art with significant performance advantages. Notably, our performance on the challenging ETH3D dataset surpasses the second-best result by 47% (AbsRel). Our model demonstrates strong generalization ability in various benchmarks while also providing geometry-preserving predictions, which show practical values.

6. Conclusion

In this paper, we propose a new geometry-preserving depth estimation framework that supports large-scale mixed-dataset training without requiring extra information. The proposed multi-view consistency loss can recover the domain-specific affine parameters of a trained model in a self-supervised manner. It can also roughly estimate the focal length by selecting appropriate values from a few selections. Our experiments on multiple benchmarks validate the effectiveness of our design.

Acknowledgements

Part of this work was supported by the National Key R&D Program of China (No. 2022ZD0118700).

References

- [1] Iro Armeni, Sasha Sax, Amir R Zamir, and Silvio Savarese. Joint 2d-3d-semantic data for indoor scene understanding. *arXiv: Comp. Res. Repository*, page 1702.01105, 2017. 6
- [2] Zuria Bauer, Francisco Gomez-Donoso, Edmanuel Cruz, Sergio Orts-Escolano, and Miguel Cazorla. Uasol, a large-scale high-resolution outdoor stereo dataset. *Scientific data*, 6(1):1–14, 2019. 3
- [3] Shariq Farooq Bhat, Ibraheem Alhashim, and Peter Wonka. Adabins: Depth estimation using adaptive bins. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, pages 4009–4018, 2021. 2
- [4] Jiawang Bian, Zhichao Li, Naiyan Wang, Huangying Zhan, Chunhua Shen, Ming-Ming Cheng, and Ian Reid. Unsupervised scale-consistent depth and ego-motion learning from monocular video. *Advances in neural information processing systems*, 32, 2019. 2
- [5] Weifeng Chen, Shengyi Qian, David Fan, Noriyuki Kojima, Max Hamilton, and Jia Deng. Oasis: A large-scale dataset for single image 3d in the wild. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, pages 679–688, 2020. 2, 3
- [6] Jaehoon Cho, Dongbo Min, Youngjung Kim, and Kwanghoon Sohn. DIML/CVL RGB-D dataset: 2m RGB-D images of natural indoor and outdoor scenes. *arXiv: Comp. Res. Repository*, 2021. 3
- [7] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, pages 5828–5839, 2017. 3, 6
- [8] Ainaz Eftekhari, Alexander Sax, Jitendra Malik, and Amir Zamir. Omnidata: A scalable pipeline for making multi-task mid-level vision datasets from 3d scans. In *Proc. IEEE Int. Conf. Comp. Vis.*, pages 10786–10796, 2021. 2
- [9] Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, and Dacheng Tao. Deep Ordinal Regression Network for Monocular Depth Estimation. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2018. 2
- [10] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, pages 3354–3361. IEEE, 2012. 3, 6
- [11] Clément Godard, Oisín Mac Aodha, and Gabriel J Brostow. Unsupervised monocular depth estimation with left-right consistency. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 270–279, 2017. 3
- [12] Yiwen Hua, Puneet Kohli, Pritish Uplavikar, Anand Ravi, Saravana Gunaseelan, Jason Orozco, and Edward Li. Holopix50k: A large-scale in-the-wild stereo image dataset. In *IEEE Conf. Comput. Vis. Pattern Recog. Worksh.*, June 2020. 6
- [13] Olga Karpenko and John Hughes. Smoothsketch: 3d free-form shapes from complex sketches. In *ACM Trans. Graphics*, pages 589–598. 2006. 4
- [14] Youngjung Kim, Hyunjoo Jung, Dongbo Min, and Kwanghoon Sohn. Deep monocular depth estimation via integration of global and local predictions. *IEEE Trans. Image Process.*, 27(8):4131–4144, 2018. 6
- [15] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [16] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Alexander Kolesnikov, Tom Duerig, and Vittorio Ferrari. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *Int. J. Comput. Vision*, 2020. 3
- [17] John Lambert, Zhuang Liu, Ozan Sener, James Hays, and Vladlen Koltun. MSeg: A composite dataset for multi-domain semantic segmentation. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2020. 3
- [18] Boyi Li, Kilian Q Weinberger, Serge Belongie, Vladlen Koltun, and Rene Ranftl. Language-driven semantic seg-

- mentation. In *Proc. Int. Conf. Learn. Representations*, 2022. 3
- [19] Zhengqi Li and Noah Snavely. Megadepth: Learning single-view depth prediction from internet photos. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, pages 2041–2050, 2018. 3, 7
- [20] Zhenyu Li, Xuyang Wang, Xianming Liu, and Junjun Jiang. Binsformer: Revisiting adaptive bins for monocular depth estimation. *arXiv preprint arXiv:2204.00987*, 2022. 2
- [21] Reza Mahjourian, Martin Wicke, and Anelia Angelova. Unsupervised learning of depth and ego-motion from monocular video using 3d geometric constraints. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5667–5675, 2018. 3
- [22] Nikolay Patakin, Anna Vorontsova, Mikhail Artemyev, and Anton Konushin. Single-stage 3d geometry-preserving depth estimation model training on dataset mixtures with uncalibrated stereo data. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, pages 1705–1714, 2022. 2, 3, 4, 6, 9
- [23] Emmanuel Prados and Olivier Faugeras. Shape from shading: a well-posed problem? In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, volume 2, pages 870–877. IEEE, 2005. 4
- [24] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. In *Proc. IEEE Int. Conf. Comp. Vis.*, pages 12179–12188, 2021. 2, 3, 6, 7, 9
- [25] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2020. 1, 2, 3, 4, 7
- [26] Jeremy Reizenstein, Roman Shapovalov, Philipp Henzler, Luca Sbordone, Patrick Labatut, and David Novotny. Common objects in 3d: Large-scale learning and evaluation of real-life 3d category reconstruction. In *Proc. IEEE Int. Conf. Comp. Vis.*, pages 10901–10911, 2021. 4
- [27] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, pages 2304–2314, 2019. 4
- [28] Ashutosh Saxena, Min Sun, and Andrew Y Ng. Make3d: Learning 3d scene structure from a single still image. *IEEE Trans. Pattern Anal. Mach. Intell.*, 31(5):824–840, 2008. 4
- [29] Thomas Schops, Johannes L Schonberger, Silvano Galliani, Torsten Sattler, Konrad Schindler, Marc Pollefeys, and Andreas Geiger. A multi-view stereo benchmark with high-resolution images and multi-camera videos. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, pages 3260–3269, 2017. 6
- [30] Shuai Shao, Zeming Li, Tianyuan Zhang, Chao Peng, Gang Yu, Xiangyu Zhang, Jing Li, and Jian Sun. Objects365: A large-scale, high-quality dataset for object detection. In *Proc. IEEE Int. Conf. Comp. Vis.*, pages 8430–8439, 2019. 3
- [31] Meng-Li Shih, Shih-Yang Su, Johannes Kopf, and Jia-Bin Huang. 3d photography using context-aware layered depth inpainting. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, pages 8028–8038, 2020. 2
- [32] Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor segmentation and support inference from rgbd images. In *Proc. Eur. Conf. Comp. Vis.*, pages 746–760. Springer, 2012. 3, 6
- [33] Chaoyang Wang, Simon Lucey, Federico Perazzi, and Oliver Wang. Web stereo video supervision for depth prediction from dynamic scenes. In *Int. Conf. 3D. Vis.*, pages 348–357. IEEE, 2019. 7
- [34] Olivia Wiles, Georgia Gkioxari, Richard Szeliski, and Justin Johnson. Synsin: End-to-end view synthesis from a single image. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, pages 7467–7477, 2020. 5
- [35] Shangzhe Wu, Christian Rupprecht, and Andrea Vedaldi. Unsupervised learning of probably symmetric deformable 3d objects from images in the wild. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, pages 1–10, 2020. 4
- [36] Ke Xian, Chunhua Shen, Zhiguo Cao, Hao Lu, Yang Xiao, Ruibo Li, and Zhenbo Luo. Monocular relative depth perception with web stereo data supervision. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, pages 311–320, 2018. 2, 3
- [37] Ke Xian, Jianming Zhang, Oliver Wang, Long Mai, Zhe Lin, and Zhiguo Cao. Structure-guided ranking loss for single image depth prediction. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, pages 611–620, 2020. 2, 6
- [38] Yuliang Xiu, Jinlong Yang, Dimitrios Tzionas, and Michael J Black. Icon: Implicit clothed humans obtained from normals. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, pages 13286–13296. IEEE, 2022. 4
- [39] Gengshan Yang, Minh Vo, Natalia Neverova, Deva Ramanan, Andrea Vedaldi, and Hanbyul Joo. Banmo: Building animatable 3d neural models from many casual videos. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2022. 4
- [40] Wei Yin, Yifan Liu, and Chunhua Shen. Virtual normal: Enforcing geometric constraints for accurate and robust depth prediction. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2021. 3, 7
- [41] Wei Yin, Yifan Liu, Chunhua Shen, Anton van den Hengel, and Baichuan Sun. The devil is in the labels: Semantic segmentation from sentences. *arXiv: Comp. Res. Repository*, page 2202.02002, 2022. 3
- [42] Wei Yin, Yifan Liu, Chunhua Shen, and Youliang Yan. Enforcing geometric constraints of virtual normal for depth prediction. In *Proc. IEEE Int. Conf. Comp. Vis.*, 2019. 2
- [43] Wei Yin, Xinlong Wang, Chunhua Shen, Yifan Liu, Zhi Tian, Songcen Xu, Changming Sun, and Dou Renyin. Diversedepth: Affine-invariant depth prediction using diverse data. *arXiv: Comp. Res. Repository*, page 2002.00569, 2020. 2
- [44] Wei Yin, Jianming Zhang, Oliver Wang, Simon Niklaus, Simon Chen, Yifan Liu, and Chunhua Shen. Towards accurate reconstruction of 3d scene shape from a single monocular image. *IEEE Trans. Pattern Anal. Mach. Intell.*, pages 1–21, 2022. 2, 3
- [45] Wei Yin, Jianming Zhang, Oliver Wang, Simon Niklaus, Long Mai, Simon Chen, and Chunhua Shen. Learning to recover 3d scene shape from a single image. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2021. 2, 3, 6, 7
- [46] Weihao Yuan, Xiaodong Gu, Zuozhuo Dai, Siyu Zhu, and Ping Tan. New crfs: Neural window fully-connected

- crfs for monocular depth estimation. *arXiv preprint arXiv:2203.01502*, 2022. [2](#)
- [47] Amir Zamir, Alexander Sax, , William Shen, Leonidas Guibas, Jitendra Malik, and Silvio Savarese. Taskonomy: Disentangling task transfer learning. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.* IEEE, 2018. [2](#), [6](#)
- [48] Amir R Zamir, Alexander Sax, Nikhil Cheerla, Rohan Suri, Zhangjie Cao, Jitendra Malik, and Leonidas J Guibas. Robust learning through cross-task consistency. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, pages 11197–11206, 2020. [2](#)
- [49] Chi Zhang, Wei Yin, Zhibin Wang, Gang Yu, Bin Fu, and Chunhua Shen. Hierarchical normalization for robust monocular depth estimation. *Proc. Advances in Neural Inf. Process. Syst.*, 2022. [2](#), [3](#)
- [50] Tinghui Zhou, Matthew Brown, Noah Snavely, and David G Lowe. Unsupervised learning of depth and ego-motion from video. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, pages 1851–1858, 2017. [3](#)