

# OmniFusion: 360 Monocular Depth Estimation via Geometry-Aware Fusion

Yuyan Li<sup>1\*</sup> Yuliang Guo<sup>2\*</sup> Zhixin Yan<sup>2</sup> Xinyu Huang<sup>2</sup> Ye Duan<sup>1</sup> Liu Ren<sup>2</sup>

<sup>1</sup>University of Missouri <sup>2</sup>Bosch Research North America

{y1235, duanye}@umsystem.edu {yuliang.guo2, zhixin.yan2, xinyu.huang, liu.ren}@us.bosch.com

## Abstract

A well-known challenge in applying deep-learning methods to omnidirectional images is spherical distortion. In dense regression tasks such as depth estimation, where structural details are required, using a vanilla CNN layer on the distorted 360 image results in undesired information loss. In this paper, we propose a 360 monocular depth estimation pipeline, *OmniFusion*, to tackle the spherical distortion issue. Our pipeline transforms a 360 image into less-distorted perspective patches (i.e. tangent images) to obtain patch-wise predictions via CNN, and then merge the patch-wise results for final output. To handle the discrepancy between patch-wise predictions which is a major issue affecting the merging quality, we propose a new framework with the following key components. First, we propose a geometry-aware feature fusion mechanism that combines 3D geometric features with 2D image features to compensate for the patch-wise discrepancy. Second, we employ the self-attention-based transformer architecture to conduct a global aggregation of patch-wise information, which further improves the consistency. Last, we introduce an iterative depth refinement mechanism, to further refine the estimated depth based on the more accurate geometric features. Experiments show that our method greatly mitigates the distortion issue, and achieves state-of-the-art performances on several 360 monocular depth estimation benchmark datasets. Our code is available at <https://github.com/yuyanli0831/OmniFusion>.

## 1. Introduction

A 360 image provides a comprehensive view of the scene with its wide field of view (FoV), which is beneficial in understanding the scene holistically. However, commonly used 360 image representation format such as the equirectangular projection (ERP) image can introduce geometric distortions. The distortion factor varies in the vertical direction and may degrade the performance of regular convolutional layers designed for non-distorted perspective images.

\*Equal contribution




Figure 1. Our method, *OmniFusion*, produces high-quality dense depth (shown as the image on the right) from a monocular ERP input (shown as an image wrapped on a unit sphere on the left). Our method uses a set of  $N$  perspective patches (i.e. tangent images) to represent the ERP image (top branch), and fuse the image features with 3D geometric features (bottom branch) to improve the estimation of the merged depth map. The corresponding camera poses of the tangent images are shown in the middle row.

Many studies have been proposed to address the distortion issue. [4, 29, 39] proposed distortion-aware convolutions or spherical customized kernels. However, it remains unclear how effective such spherical convolutions are, especially in deeper layers [29, 35]. Some spherical CNNs [8, 37] defined convolution in the spectral domain, with potentially heavier computation overhead. Attempts have also been made to tackle the ERP distortion via other less-distorted formats. BiFuse [35] and UniFuse [20] took complementary properties from ERP and cubemap. Several works [7, 30] applied regular CNNs repeatedly to multiple perspective projections of the 360 image. Recently, Eder et al. [13] proposed to use a set of subdivided icosahedron tangent images, and demonstrated that using tangent image representation can facilitate the network transfer between perspective and 360 images.

It is advantageous to use tangent images [13] as it has less distortion, and can make good use of the large pool of pre-trained CNNs developed for perspective imaging. Additionally, the tangent image representation inherits a superior scalability to handle high resolution inputs compared to those holistically method. However, the vanilla pipeline [13] has some limitations. First, severe discrepancies occur between perspective views since the same object may appear differently from multiple views (an example is shown

in Figure 3). This issue is especially problematic for the depth regression task, since the inconsistent depth scale estimated from individual tangent images creates undesired artifacts during merging. Second, the advantage of estimating depth from holistic 360 image is unfortunately lost, because of the decomposition of the global scene into local tangent images. The predictions from the tangent images are independent of each other and there is no information exchange between tangent images.

In this paper, we present *OmniFusion*, a 360 monocular depth estimation framework with geometry-aware fusion (see Figure 1). We proposed the following three key components to solve the aforementioned discrepancy issue and merge the depth results of tangent images seamlessly. First, we use a geometric embedding module to provide additional features to compensate for the discrepancy between 2D features from patch to patch. For each patch, we calculate the 3D points located on the spherical surface that correspond to the patch pixels, encode them and the patch center coordinate through shared Multi-layer Perceptron (MLP), and add the geometric features to the corresponding 2D features. Second, to regain the holistic power in understanding the entire scene, we incorporate a self-attention-based transformer in our pipeline. With the transformer, patch-wise information is globally aggregated to enhance the estimation of the global scale of depth, and to improve the consistency between patch-wise results. Third, we introduce an iterative refining mechanism, where more accurate 3D information from the predicted depth maps is fed back to the geometric embedding module to further improve the depth quality in an iterative manner.

We test *OmniFusion* on three benchmark datasets: Stanford2D3D [1], Matterport3D [3], and 360D [43]. Experimental results show that our method outperforms state-of-the-art methods by a significant margin on all of these datasets.

Our contributions can be summarized as follows:

- We present a 360 monocular depth prediction pipeline that addresses the distortion issue via geometry-aware fusion and achieves the state-of-the-art performance.
- We introduce a geometric embedding network to provide 3D geometric features to mitigate the discrepancy in patch-wise image features.
- We incorporate a self-attention-based transformer to globally aggregate patch-wise information which enhances the estimation of the physical scale of depth.
- We propose an iterative mechanism to further improve the depth estimation with structural details.

## 2. Related Work

### 2.1. Monocular depth estimation

Monocular depth estimation, which takes a single RGB image as input to predict pixel-wise depth value, has been extensively investigated due to its broad applications. Early works mainly focused on network architecture and supervision [14, 19, 22]. Recently, researchers have been investigating the use of unsupervised learning on stereo pairs [16, 17, 36] or monocular video streams [18, 41] to expand training data to unlabelled image sequences for broader applications. However, such approaches are still sensitive to many factors (e.g. camera intrinsic changes), and very challenging to be generalizable to new scenes. To improve the robustness and scalability, some methods utilize additional sensor input such as LiDAR and RGBD camera [6, 24]. However, the extra computation or power consumption are not welcomed in many practical scenarios.

### 2.2. 360 depth estimation

Monocular depth estimation from 360 images has been investigated from a variety of perspectives. Zioulis et al. [42] explored the spherical stereo geometry and estimated depth from monocular ERP input via view synthesis. PanDepth [23] leveraged 360 stereo constraints to improve monocular depth performance. Eder et al. [11] and Zeng et al. [38] explored joint learning from different modalities (e.g. layout, normal, semantics, etc.). HoHoNet [31] proposed to utilize latent horizontal feature representation to encode ERP image features. To handle the irregular distortion of ERP images, several distortion-aware convolutions [4, 15, 28, 29, 39] have been proposed. For example, Fernandez et al. [15] introduced EquiConv which applied deformable convolution to accommodate spherical geometry. Tateno et al. [32] proposed to apply regular CNN to perspective images during training, and distortion-aware convolution during testing. Instead of directly tackling the distortion of ERP, several approaches proposed to use other representations with less distortion, such as cubemap [5, 34], fusion between ERP and cubemap [20, 35], and multiple perspective projections of 360 images [7, 30]. A recent work by Eder et al. [13] proposed to use tangent images, a set of oriented, low-distortion images rendered tangent to faces of the icosahedron, to represent a 360 image. It is advantageous to use tangent images since it has less distortion and can effectively leverage pre-trained CNN models developed for perspective imaging. However, discrepancies between tangent images are not addressed in [13], which leads to a downgrade of the final merged result. In this work, we follow the paradigm proposed in [13] of using tangent images, but simplified and adapted it for depth estimation. In addition, we successfully address the discrepancy issue by incorporating geometry-aware fusion and the transformer.

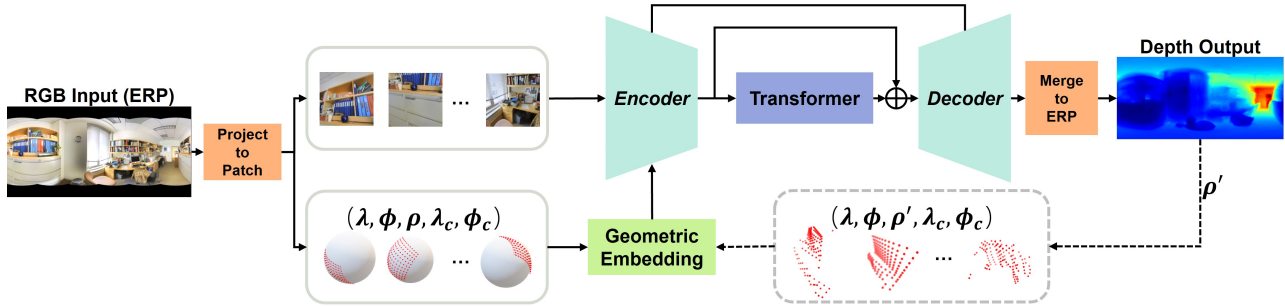


Figure 2. Overview of our proposed *OmniFusion*. Our method takes a monocular ERP RGB as input, projects it onto multiple patches at multiple viewpoints, and processes each distortion-free patch with an encoder-decoder network to produce patch-wise depth maps (top-stream). The patch-wise outputs are merged into a final ERP depth map. Meanwhile, the corresponding points located on the spherical surface are sampled and passed through a geometric embedding network to produce geometric features (bottom-stream). The geometric features are fused into the image encoder to compensate for the patch-wise discrepancy and to improve the quality of the merged result. For each sampled point, we use its spherical coordinates  $(\lambda, \phi, \rho)$ , together with the tangent plane center coordinates  $(\lambda_c, \phi_c)$  as input attributes to the geometric embedding network which provides the necessary information to align 2D features. A transformer architecture is integrated to conduct global aggregation of the deep patch-wise feature which further improves the consistency of patch-wise outputs. Moreover, we incorporate an iterative refining mechanism (visualized in dashes), to further improve the depth recovery. In particular,  $\rho$  value is updated according to the depth estimated from the previous iteration.

### 2.3. Transformer

Originally proposed in natural language processing [33], the transformer architecture has since been widely used in computer vision tasks such as image classification [10], depth estimation [26], object detection [2], and semantic segmentation [27, 40]. The visual transformer has a natural fit with monocular depth estimation as long-range context can be explicitly exploited by the self-attention module. When applying transformer to 360 images, the distortion however, can decrease the power of the transformer in exploiting the pairwise correlation between patches. In this work, we feed the transformer with distortion-free, and geometry-aware input, so that the transformer can focus on the global aggregation of patch-wise information.




Figure 3. (a) An example of tangent image projection. Two tangent images are projected from two different viewpoints. The corresponding areas are highlighted with the same color in both ERP and tangent patches. As illustrated, there usually exist overlapping areas between two neighboring patches, and the same object may appear differently in different patches. (b) The illustration of the gnomonic projection. A point  $P_s(\lambda, \phi)$  located on the spherical sphere is projected onto a point  $P_t(x_t, y_t)$  on the flat plane which is tangent to a point  $P_c(\lambda_c, \phi_c)$ .

## 3. Method

Figure 2 shows an overview of the full pipeline of the proposed *OmniFusion* framework. First, an ERP input image is transformed into a set of tangent images via gnomonic projection (Figure 3). The projected distortion-free tangent images are then passed through an encoder-decoder network to produce patch-wise depth estimations, which are later fused into an ERP depth output. To ease the patch-wise discrepancy, we introduce a novel geometric embedding module that encodes the spherical coordinate associated with each tangent image pixel, providing additional geometric features to facilitate the integration of patch image features. To further improve the consistency between patch-wise predictions and to better estimate the global depth scale, the features from the deepest level of the encoder are globally aggregated through a self-attention-

based transformer. Finally, an iterative refining mechanism is adopted to further improve the depth quality. We update the spherical coordinates iteratively based on the more accurate estimation obtained from the previous iteration. We train our network in an end-to-end fashion, with the only supervision being the final merged depth compared to the ground truth.

### 3.1. Depth estimation from tangent images

Our method relies on the less-distorted tangent images to address the irregular distortion in 360 images. A tangent image is generated via *gnomonic projection* of a sphere surface onto a flat, rectangular plane surface. The gnomonic projection [9] (Figure 3) is a mapping obtained by projecting points  $P_s(\lambda, \phi)$  on the surface of sphere from a sphere’s




Figure 4. First row: (a) An example of an ERP RGB input image, (b) the final merged predicted ERP depth map, (c) the ground truth ERP depth. Second row: (d) RGB tangent image patches generated from (a), (e) the patch-wise estimated depth maps, (f) the patch-wise estimated confidence maps that are used as weights and facilitates the ERP depth merging.

center  $O$  to point  $P_t(x_t, y_t)$  in a plane that is tangent to a point  $P_c(\lambda_c, \phi_c)$ . We use  $(\lambda, \phi)$  to indicate the longitude and latitude, respectively, and  $(x_t, y_t)$  to indicate a 2D point position on the tangent image. The detailed formulas are included in the Appendix.

In our experiments, we use a set of  $N = 18$  tangent images for a balance of speed and performance (A related ablation study can be found in Section 4.4). Tangent images are sampled at four different latitudes:  $-67.5^\circ$ ,  $-22.5^\circ$ ,  $22.5^\circ$ ,  $67.5^\circ$ , and we sample 3, 6, 6, 3 patches on each of these latitudes, respectively (Figure 4). All tangent images share the same resolution and FoV. We chose this non-uniform sampling based on the fact that tangent images of the same resolution can cover different ranges of longitude when the centered at different latitudes. To ensure the sampled patches near the poles do not overlap to an extreme extent, we take fewer samples to cover the near-pole area in the ERP space. Since the generated tangent images are distortion-free, we can easily apply regular encoder-decoder CNN architectures to predict a depth map from each tangent image. For better convergence and accuracy, we leverage high-performance pre-trained networks (e.g., ResNet [19]) when initializing our encoder. We pass all  $N$  tangent images simultaneously through the encoder, and obtain  $N$  feature maps that will be used as tokens later in the transformer. For the decoder, we use a stack of up-sampling layers followed by  $3 \times 3$  convolutions, with skip-connections from the encoder.

The baseline presented so far can be considered as a customized version of [13]. We adopt different rendering of tangent images and the network architecture from [13] to make the baseline method more effective and efficient. Note that for our baseline, no transformer, geometric fusion, or confidence map is used, the output depth is the average of

all patches.

### 3.2. Geometry-aware feature fusion

The simplicity of predicting depth maps from tangent images nonetheless comes with a cost. As the depth estimation is now conducted independently, a globally consistent depth scale is no longer guaranteed. Furthermore, as shown in Figure 3 (a) and Figure 4 (d), an object (e.g. the painting on the wall in Figure 3 (a)) will be projected onto multiple tangent images from various angles and therefore will be encoded differently in different tangent images. Discrepancies between patch depth estimations, especially in overlapping areas, can result in significant artifacts in the final merged ERP depth map (Figure 5 (e)).

To compensate for the differences between patch-wise image features, we introduce a *geometric embedding* network (see Figure 2) to provide additional geometric information. For a pixel  $P_t(x_t, y_t)$  located on a tangent image, we use its corresponding spherical coordinates located on the unit sphere,  $P_s(\lambda, \phi, \rho)$ , together with the center of the tangent image  $P_c(\lambda_c, \phi_c)$ , as the input attributes of the geometric embedding network.  $P_s$  makes the embedding aware of the global position, e.g., to tell whether two image pixels from two patches relate to the same spherical coordinates. However, geometric features out of  $P_s$  alone can not align different 2D features. To this end,  $P_c$  is taken as additional attributes to make the embedding able to differentiate from patch to patch, such that the learned geometric features can make the patch features tend to be consistent. Through the combination of the tangent image features and the geometric features as well as an end-to-end learned network, the adjusted features lead to a much cleaner merged depth. As observed in Figure 5 (d), the extracted image features with geometric embedding show much better consistency in the






Figure 5. Illustration of the effectiveness of geometry-aware feature fusion. An ERP RGB image is shown in (a), the ground truth depth is shown in (b). Visualizations of the feature map and the final depth map from the baseline are shown in (c) and (e) respectively. For comparison, (d) and (f) show the feature map and the final depth map out of the proposed *OmniFusion*, where the image features are fused with geometric features. Observe that our method yields a more self-consistent feature map and a more structural depth map compared to the baseline, especially in regions highlighted in rectangles.

feature map merged in the ERP space, compared to features without geometric embedding as shown in Figure 5 (c). Consequently, the final depth map out of *OmniFusion* shown in Figure 5 (f) appears to be much cleaner compared to the baseline depth map shown in Figure 5 (e).

The geometric embedding network consists of two layers of MLPs, and encodes the 5-channel spherical attributes into 64-channel feature maps. We fuse this geometric embedding with image features at the same pixel location in the encoder via element-wise summation. In order to maintain more structural details, early fusion is adopted. Geometric features are added to the *layer1* of the ResNet encoder where we experimentally achieved the best performance. It is worth noting that the additional computational cost associated with the geometric embedding module is minimal compared to the original encoder-decoder (Table 2). The geometric features for the first iteration are fixed once learned, since they are independent from image inputs. Only the second iteration requires to re-compute the geometric features.

### 3.3. Global aggregation with transformer

When decomposing the ERP into a sequence of tangent images, we no longer have the holistic view of the 3D environment. To make up for this loss, we leverage the transformer architecture to aggregate information from the patches in a global fashion. The global aggregation is ex-

pected to improve the consistency of depth estimations from patches, and to better regress the global scale of depth out of a larger FoV.

Using the feature maps extracted from the encoder, we first apply a  $1 \times 1$  convolution layer to reduce channel dimensions for better efficiency. Then we flatten the feature maps into  $N$  1-D feature vectors  $X_0 = [x^1, x^2, \dots, x^N] \in R^{N \times d}$  which will be used as tokens in the transformer. The learnable positional embedding  $E_{pos} \in R^{N \times d}$  are added to the feature tokens to retain positional information in a similar way as proposed in [10]. Through the self-attention architecture, the transformer learns to globally aggregate the information from all the patches to adjust the features from each patch, where the aggregation weights account for the pairwise correlation both from the visual features and the positional features. The architecture of the multi-head attention transformer follows [33].

### 3.4. Depth merging with learnable confidence map

The aforementioned geometric embedding and transformer modules significantly reduce discrepancies among different patch-wise depth estimations. Yet, the depth merging does not achieve a pixel-level seamless fusion. To further improve the merging (Figure 4 (b)), we ask the network to simultaneously predict a confidence map for each patch besides depth regression. The merged depth is then computed as a weighted average of all patch depth predictions with confidence scores used as weights. In detail, two separate regression layers are appended to the decoder, one for depth regression, the other for confidence score regression. Both the depth maps (Figure 4 (e)) and confidence maps (Figure 4 (f)) are mapped to ERP domain following the inverse gnomonic transformation before merging. (More details are included in the Appendix.)

### 3.5. Iterative depth refinement

The geometric embedding utilizes the spherical coordinates  $(\lambda, \phi, \rho)$  corresponding to tangent image pixels for geometry-aware fusion.  $\rho$  is initially fixed as no depth information is available. The depth information will be available after one iteration, which can be used to update  $\rho$  and provide more accurate geometry information for the geometric embedding module. Based on this observation, we propose an iterative depth refinement scheme (see Figure 2).

In the first iteration (Section 3.2), the spherical coordinates  $(\lambda, \phi, \rho)$  of points located on the unit sphere are used for geometric embedding. For the subsequent iterations, we update  $\rho \rightarrow \rho'$ , using the new depth value estimated from the previous iteration (the depth of ERP image is defined as the distance from the real-world point to the camera center). The updated attributes with more accurate geometry will be passed into the geometric embedding network in the next iteration. An ablation study is presented in Section 4

to demonstrate the effectiveness of more accurate geometric embedding.

## 4. Experiments

### 4.1. Datasets

*OmniFusion* is tested on three well-known benchmark datasets: Stanford2D3D [1], Matterport3D [3], 360D [43]. **Stanford2D3D** [1] dataset consists of 1,413 real world panorama images from six large-scale indoor areas. We follow the official train-test split which uses the fifth area for testing, and others for training. We use resolution  $512 \times 1024$ .

**Matterport3D** [3] contains a total of 10,800 indoor panorama RGBD images. We follow the official split which takes 61 rooms for training and the rest for testing. We use resolution  $512 \times 1024$  in our experiments.

**360D** [43] is a RGBD panorama benchmark provided by Zioulis et al. [43]. It is composed of two other synthetic datasets (SunCG and SceneNet), and two real world datasets (Stanford2D3D and Matterport3D). There are 35,977 photo-realistic panorama RGBD images in the 360D that are rendered from the aforementioned four datasets. We follow the default train-test splits and use resolution  $256 \times 512$ .

### 4.2. Implementation details

We adopt the same quantitative evaluation metrics as used in [22, 43], including Absolute Relative Error (Abs Rel), Root Mean Squared Error (RMSE), Root Mean Squared Error in logarithmic space (RMSE(log)) and accuracy with a threshold  $\delta_t$ , where  $t \in 1.25, 1.25^2, 1.25^3$ . Arrows next to the metric indicate the direction of better performance in all tables. We implement our network using PyTorch and train it on two Nvidia RTX GPUs. We use the default setting of Adam optimizer [21] and a initial learning rate of 0.0001 with cosine annealing [25] learning rate policy. We train Stanford2D3D [1] for 80 epochs, and 60 epochs for Matterport3D [3] and 360D [43]. The default number of patches we use is 18. The default patch size we use for Stanford2D3D [1] and matterport [3] is  $256 \times 256$ , the patch FoV is  $80^\circ$ . For 360D [43], we use  $128 \times 128$  as patch size. We leverage pre-trained ResNet [19] as image encoder in these experiments. The network is trained end-to-end, the same model is used for all iterations. For the loss function, following [32, 35], we adopt BerHu loss [22] for depth supervision. The final loss is the summation of depth losses from all iterations.

### 4.3. Overall performance

We present our model performances and compare it to the existing methods in Table 1. We omit the methods that use supervision signals other than depth [11, 38] and the

self-supervised approaches [42] for fair comparison. For all datasets, we show our results with 1-iteration (1-iter) and 2-iterations (2-iter). We demonstrate in Table 1 that even with 1-iter setting, our method is able to outperform all the existing methods on Matterport3D [3], and achieve on par performance with current state-of-the-arts on 360D. With 2-iter setting, our method outperforms BiFuse [35] by 21.4% (Abs Rel) on Stanford2D3D, 56.1% (Abs Rel) on Matterport3D, 30% (Abs Rel) on 360D. Comparing to UniFuse [20], our method improves by 6.3% (Abs Rel) on Stanford2D3D, 15.3% (Abs Rel) on Matterport3D, 7.7% (Abs Rel) on 360D. Note that compared to ODE-CNN [6] which used additional sensor input, our method reduces Abs Rel by 7.9%. Qualitative results of our method can be visualized in Figure 6. As observed, our method (1-iter and 2-iter) improves the baseline, a direct customization from [13], significantly in producing less erroneous depth maps with sharper boundaries and smoother surfaces recovered.

### 4.4. Ablation studies

**Individual component study.** We investigate the effectiveness of our method by adding one key component at a time (Table 2 and Figure 7). We form our baseline experiment with ResNet34 as encoder without the transformer or the geometric fusion. We experiment on Stanford2D3D, using the configuration of 18 patches,  $256 \times 256$  patch size,  $80^\circ$  FoV. As observed from Table 2, the geometry-aware fusion, which only adds less than 2K parameters, is able to improve Abs Rel significantly by 9.7%. While being extremely light-weighted, the geometric fusion part proves to be quite beneficial. The incorporation of the transformer, which increases around 19M parameters, leads to another boost of performance by 5.7% (Abs Rel). Together with transformer and geometric fusion, the performance is significantly improved by 15.4% (Abs Rel) with 1-iter setting, and 16.4% (Abs Rel) with 2-iter setting. Qualitative results are shown in Figure 7. As observed, as we add more modules into our pipeline, the output depth map appears to show fewer artifacts and more structural details. In the meantime, the visualized error maps clearly show the decreasing trend of estimation errors.

**Patch size and number of patches.** Patch size and the number of patches affect both the accuracy and the efficiency of the method. In this study, we aim to find an optimal balance between efficiency and performance. Theoretically, neither a large patch size nor a large number of patches is desired since they both lead to higher computational complexity. However, table 3 also indicates the patch size can not be too small, since the monocular depth estimation requires large-enough FoV to hypothesis the depth scale. We also observe that keep increasing the number of patches (e.g.,  $\geq 26$ ) can degrade the performance, since a larger number of patches also increases the overlapping

Datasets	Methods	Abs Rel↓	Sq Rel ↓	RMSE↓	RMSE(log)↓	$\delta_1 \uparrow$	$\delta_2 \uparrow$	$\delta_3 \uparrow$
Stanford2D3D [1]	FCRN [22]	0.1837	-	0.5774	-	0.7230	0.9207	0.9731
	RectNet [43]	0.1996	-	0.6152	-	0.6877	0.8891	0.9578
	BiFuse with fusion [35]	0.1209	-	0.4142	-	0.8660	0.9580	0.9860
	UniFuse with fusion [20]	0.1114	-	0.3691	-	0.8711	0.9664	0.9882
	HoHoNet [31]	0.1014	-	0.3834	-	<b>0.9054</b>	0.9693	0.9886
	<b>OmniFusion, Ours (1-iter)</b>	0.0961	0.0543	0.3715	0.1699	0.8940	0.9714	0.9900
	<b>OmniFusion, Ours (2-iter)</b>	<b>0.0950</b>	<b>0.0491</b>	<b>0.3474</b>	<b>0.1599</b>	0.8988	<b>0.9769</b>	<b>0.9924</b>
Matterport3D [3]	FCRN [22]	0.2409	-	0.6704	-	0.7703	0.9714	0.9617
	RectNet [43]	0.2901	-	0.7643	-	0.6830	0.8794	0.9429
	BiFuse with fusion [35]	0.2048	-	0.6259	-	0.8452	0.9319	0.9632
	UniFuse with fusion [35]	0.1063	-	0.4941	-	0.8897	0.9623	0.9831
	HoHoNet [31]	0.1488	-	0.5138	-	0.8786	0.9519	0.9771
	<b>OmniFusion, Ours (1-iter)</b>	0.0980	0.0611	0.4536	0.1587	0.9040	0.9757	0.9919
	<b>OmniFusion, Ours (2-iter)</b>	<b>0.0900</b>	<b>0.0552</b>	<b>0.4261</b>	<b>0.1483</b>	<b>0.9189</b>	<b>0.9797</b>	<b>0.9931</b>
360D [43]	FCRN [22]	0.0699	0.2833	-	-	0.9532	0.9905	0.9966
	RectNet [43]	0.0702	0.0297	0.2911	0.1017	0.9574	0.9933	0.9979
	Mapped Convolution [12]	0.0965	0.0371	0.2966	0.1413	0.9068	0.9854	0.9967
	BiFuse with fusion [35]	0.0615	-	0.2440	-	0.9699	0.9927	0.9969
	UniFuse with fusion [35]	0.0466	-	0.1968	-	0.9835	0.9965	0.9987
	ODE-CNN [6]	0.0467	0.0124	<b>0.1728</b>	0.0793	0.9814	0.9967	0.9989
	<b>OmniFusion, Ours (1-iter)</b>	0.0469	0.0127	0.1880	0.0792	0.9827	0.9963	0.9988
	<b>OmniFusion, Ours (2-iter)</b>	<b>0.0430</b>	<b>0.0114</b>	0.1808	<b>0.0735</b>	<b>0.9859</b>	<b>0.9969</b>	<b>0.9989</b>

Table 1. Quantitative Results for depth estimation on Stanford2D3d [1], Matterport3D [3], 360D [43] datasets. Notably, our method *OmniFusion* achieves state-of-the-art performances in all datasets, outperforming the existing works by a significant margin.

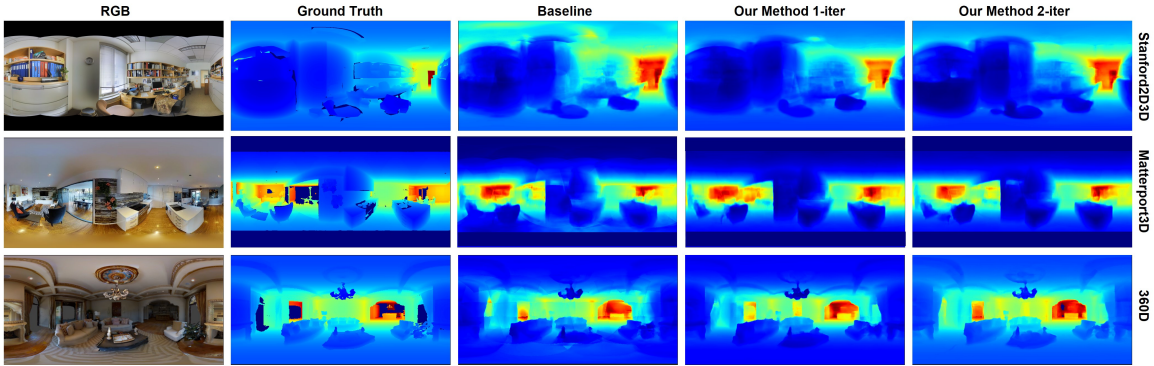


Figure 6. Qualitative results on Stanford2D3D [1], Matterport3D [3] and 360D [43]. From left to right: ERP RGB input, ground truth depth, depth output from the baseline, depth output from our method 1-iter and 2-iter. In comparison to the baseline (described in Section 3.1, our method (1-iter, 2-iter) leads to more structural depth, which appear sharp along those object boundaries and smooth within surfaces.

Methods	#Params	FPS↑	Abs Rel↓	Sq Rel↓	RMSE↓
Baseline	23.5M	9.4	0.1136	0.0638	0.3894
Baseline + geometric fusion (1-iter)	23.5M (+1.3K)	9.3	0.1026	0.0588	0.3812
Baseline + geometric fusion + transformer (1-iter)	42.3M (+18.8M)	9.2	0.0961	0.0543	0.3715
<b>Baseline + geometric fusion + transformer (2-iter)</b>	42.3M (+18.8M)	4.6	<b>0.0950</b>	<b>0.0491</b>	<b>0.3474</b>

Table 2. The ablation study for individual components. Starting from a baseline method with no geometric fusion or transformer, we add each component one at a time. We use ResNet34 for all the experiments.

area, which in turn may intensify the discrepancy problem. As a result, we choose to use a relatively small number of patches  $N = 18$  with a relatively large resolution  $256 \times 256$  to balance between efficiency and performance.

**Image encoder and number of iterations.** We compare the performance of leveraging different image encoders. As listed in Table 4, ResNet34 [19] outperforms ResNet18 with more complexity. This indicates the potential of our



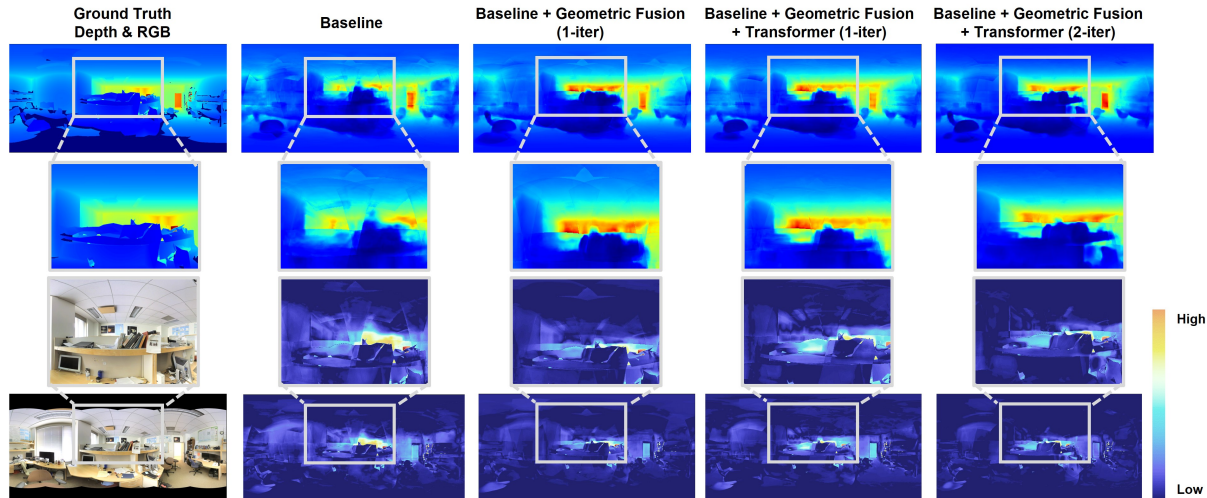


Figure 7. Qualitative comparisons regarding individual components. The top row shows the visual comparisons in depth maps, and the bottom row shows the visual comparisons of the corresponding error maps between the predicted depth maps. The middle two rows show the close-up views of the highlighted areas in the top and bottom rows, respectively. As we add more modules into the pipeline (Figure 2), the depth estimation becomes more accurate with lower errors, sharper object boundaries and smoother surfaces. The trend of the change in errors can be directly observed from the error maps.

#patch	Patch size	Patch FoV	Abs Rel↓	Sq Rel↓	RMSE↓
10	256x256	120	0.1067	0.0571	0.3788
18	128x128	80	0.1178	0.0666	0.4018
<b>18</b>	<b>256x256</b>	<b>80</b>	<b>0.1037</b>	<b>0.0589</b>	<b>0.3686</b>
26	256x256	60	0.1104	0.0679	0.3955
46	128x128	50	0.1181	0.0680	0.4101

Table 3. The ablation study for patch size and number of patches.

Encoder	#iters	FPS↑	Abs Rel↓	Sq Rel↓	RMSE↓
ResNet18	1	<b>9.8</b>	0.1037	0.0589	0.3686
ResNet18	2	4.6	0.0979	0.0539	0.3702
ResNet18	3	3.1	<b>0.0981</b>	0.0521	<b>0.3699</b>
ResNet18	4	1.5	0.0983	<b>0.0519</b>	0.3700
ResNet34	1	<b>9.2</b>	0.0961	0.0543	0.3715
ResNet34	2	4.6	0.0950	0.0491	<b>0.3474</b>
ResNet34	3	2.9	<b>0.0894</b>	<b>0.0482</b>	0.3498
ResNet34	4	1.4	0.0899	0.0485	0.3491

Table 4. The ablation study for different encoder models and different number of iterations.

method, as one can incorporate a more sophisticated encoder network. We also study the influence of iterations. We use the 2-iteration framework for the training since we expect the trained network to handle different types of 3D coordinates. While for testing, we compare 1-4 iterations respectively on the two backbones. As seen from Table 4, there is an evident improvement from 1-iter to 2-iter, a slighter improvement from 2-iter to 3-iter, and no gain from 3-iter to 4-iter. Considering the trade-off in performance and the speed, we opt to choose 1-iter or 2-iter settings.

## 5. Conclusion

In this paper, we propose a novel pipeline, *OmniFusion*, for 360 monocular depth estimation. To address the spherical distortion presented in 360 images, as well as to improve the scalability to high-resolution inputs, we use gnomonic projection-based tangent image presentation. To alleviate the discrepancy between patches, we introduce a geometry-aware fusion mechanism which fuse 3D geometric features with the image features. A self-attention transformer is integrated into our pipeline to globally aggregate information from patches, which leads to more consistent patch-wise predictions. We further extend the geometry-aware fusion with an iterative refining scheme which further improves the depth estimation with more structural details. We show that *OmniFusion* effectively mitigates distortion, and significantly improves the depth estimation performance. Our experiments show that our method achieves state-of-the-art performances on several datasets.

## Acknowledgments

The research of Yuyan Li and Ye Duan were partially supported by the National Science Foundation under award CNS-2018850, National Institute of Health under awards NIBIB-R03-EB028427 and NIBIB-R01-EB02943, and U.S. Army Research Laboratory W911NF2120275. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the U. S. Government or agency thereof.



## References

- [1] Iro Armeni, Sasha Sax, Amir R Zamir, and Silvio Savarese. Joint 2d-3d-semantic data for indoor scene understanding. *arXiv preprint arXiv:1702.01105*, 2017. 2, 6, 7
- [2] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European Conference on Computer Vision*, pages 213–229. Springer, 2020. 3
- [3] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niessner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3d: Learning from rgb-d data in indoor environments. *International Conference on 3D Vision (3DV)*, 2017. 2, 6, 7
- [4] Hong-Xiang Chen, Kunhong Li, Zhiheng Fu, Mengyi Liu, Zonghao Chen, and Yulan Guo. Distortion-aware monocular depth estimation for omnidirectional images. *IEEE Signal Processing Letters*, 28:334–338, 2021. 1, 2
- [5] Hsien-Tzu Cheng, Chun-Hung Chao, Jin-Dong Dong, Hao-Kai Wen, Tyng-Luh Liu, and Min Sun. Cube padding for weakly-supervised saliency prediction in 360 videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1420–1429, 2018. 2
- [6] Xinjing Cheng, Peng Wang, Yanqi Zhou, Chenye Guan, and Ruigang Yang. Omnidirectional depth extension networks. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 589–595. IEEE, 2020. 2, 6, 7
- [7] Shih-Han Chou, Yi-Chun Chen, Kuo-Hao Zeng, Hou-Ning Hu, Jianlong Fu, and Min Sun. Self-view grounding given a narrated 360 video. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018. 1, 2
- [8] Taco S Cohen, Mario Geiger, Jonas Köhler, and Max Welling. Spherical cnns. *arXiv preprint arXiv:1801.10130*, 2018. 1
- [9] Harold Scott Macdonald Coxeter. Introduction to geometry. 1961. 3
- [10] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 3, 5
- [11] Marc Eder, Pierre Moulon, and Li Guan. Pano popups: Indoor 3d reconstruction with a plane-aware network. In *2019 International Conference on 3D Vision (3DV)*, pages 76–84. IEEE, 2019. 2, 6
- [12] Marc Eder, True Price, Thanh Vu, Akash Bapat, and Jan-Michael Frahm. Mapped convolutions. *arXiv preprint arXiv:1906.11096*, 2019. 7
- [13] Marc Eder, Mykhailo Shvets, John Lim, and Jan-Michael Frahm. Tangent images for mitigating spherical distortion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12426–12434, 2020. 1, 2, 4, 6
- [14] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. In *Advances in neural information processing systems*, pages 2366–2374, 2014. 2
- [15] Clara Fernandez-Labrador, José M Fácil, Alejandro Perez-Yus, Cédric Demonceaux, Javier Civera, and José J Guerrero. Corners for layout: End-to-end layout recovery from 360 images. *arXiv:1903.08094*, 2019. 2
- [16] Ravi Garg, B. G. Vijay Kumar, Gustavo Carneiro, and Ian D. Reid. Unsupervised CNN for single view depth estimation: Geometry to the rescue. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part VIII*, volume 9912, pages 740–756, 2016. 2
- [17] 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. 2
- [18] Clément Godard, Oisín Mac Aodha, Michael Firman, and Gabriel J Brostow. Digging into self-supervised monocular depth estimation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3828–3838, 2019. 2
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 2, 4, 6, 7
- [20] Hualie Jiang, Zhe Sheng, Siyu Zhu, Zilong Dong, and Rui Huang. Unifuse: Unidirectional fusion for 360° panorama depth estimation. *IEEE Robotics and Automation Letters*, 2021. 1, 2, 6, 7
- [21] Diederick P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015. 6
- [22] 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)*, pages 239–248. IEEE, 2016. 2, 6, 7
- [23] Yuyan Li, Zhixin Yan, Ye Duan, and Liu Ren. Panodepth: A two-stage approach for monocular omnidirectional depth estimation. In *2021 International Conference on 3D Vision (3DV)*, pages 648–658. IEEE, 2021. 2
- [24] Juan-Ting Lin, Dengxin Dai, and Luc Van Gool. Depth estimation from monocular images and sparse radar data. In *IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2020, Las Vegas, NV, USA, October 24, 2020 - January 24, 2021*, pages 10233–10240. IEEE, 2020. 2
- [25] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016. 6
- [26] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 12179–12188, October 2021. 3
- [27] Robin Strudel, Ricardo Garcia, Ivan Laptev, and Cordelia Schmid. Segmenter: Transformer for semantic segmentation. *arXiv preprint arXiv:2105.05633*, 2021. 3

- [28] Yu-Chuan Su and Kristen Grauman. Learning spherical convolution for fast features from 360 imagery. In *Advances in Neural Information Processing Systems*, pages 529–539, 2017. [2](#)
- [29] Yu-Chuan Su and Kristen Grauman. Kernel transformer networks for compact spherical convolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9442–9451, 2019. [1](#), [2](#)
- [30] Yu-Chuan Su, Dinesh Jayaraman, and Kristen Grauman. Pano2vid: Automatic cinematography for watching 360 videos. In *Asian Conference on Computer Vision*, pages 154–171. Springer, 2016. [1](#), [2](#)
- [31] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Hohonet: 360 indoor holistic understanding with latent horizontal features. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2573–2582, 2021. [2](#), [7](#)
- [32] Keisuke Tateno, Nassir Navab, and Federico Tombari. Distortion-aware convolutional filters for dense prediction in panoramic images. In *The European Conference on Computer Vision (ECCV)*, September 2018. [2](#), [6](#)
- [33] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017. [3](#), [5](#)
- [34] Fu-En Wang, Hou-Ning Hu, Hsien-Tzu Cheng, Juan-Ting Lin, Shang-Ta Yang, Meng-Li Shih, Hung-Kuo Chu, and Min Sun. Self-supervised learning of depth and camera motion from 360 videos. In *Asian Conference on Computer Vision*, pages 53–68. Springer, 2018. [2](#)
- [35] Fu-En Wang, Yu-Hsuan Yeh, Min Sun, Wei-Chen Chiu, and Yi-Hsuan Tsai. Bifuse: Monocular 360 depth estimation via bi-projection fusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 462–471, 2020. [1](#), [2](#), [6](#), [7](#)
- [36] Junyuan Xie, Ross B. Girshick, and Ali Farhadi. Deep3d: Fully automatic 2d-to-3d video conversion with deep convolutional neural networks. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part IV*, volume 9908, pages 842–857, 2016. [2](#)
- [37] Jiachen Yang, Tianlin Liu, Bin Jiang, Wen Lu, and Qinggang Meng. Panoramic video quality assessment based on non-local spherical cnn. *IEEE Transactions on Multimedia*, 23:797–809, 2020. [1](#)
- [38] Wei Zeng, Sezer Karaoglu, and Theo Gevers. Joint 3d layout and depth prediction from a single indoor panorama image. In *European Conference on Computer Vision*, pages 666–682. Springer, 2020. [2](#), [6](#)
- [39] Qiang Zhao, Chen Zhu, Feng Dai, Yike Ma, Guoqing Jin, and Yongdong Zhang. Distortion-aware cnns for spherical images. In *IJCAI*, pages 1198–1204, 2018. [1](#), [2](#)
- [40] Sixiao Zheng, Jiachen Lu, Hengshuang Zhao, Xiatian Zhu, Zekun Luo, Yabiao Wang, Yanwei Fu, Jianfeng Feng, Tao Xiang, Philip HS Torr, et al. Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6881–6890, 2021. [3](#)
- [41] Tinghui Zhou, Matthew Brown, Noah Snavely, and David G. Lowe. Unsupervised learning of depth and ego-motion from video. In *CVPR*, 2017. [2](#)
- [42] Nikolaos Zioulis, Antonis Karakottas, Dimitrios Zarpalas, Federico Alvarez, and Petros Daras. Spherical view synthesis for self-supervised 360 depth estimation. In *2019 International Conference on 3D Vision (3DV)*, pages 690–699. IEEE, 2019. [2](#), [6](#)
- [43] Nikolaos Zioulis, Antonis Karakottas, Dimitrios Zarpalas, and Petros Daras. Omnidepth: Dense depth estimation for indoors spherical panoramas. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 448–465, 2018. [2](#), [6](#), [7](#)