

Layout-Guided Novel View Synthesis from a Single Indoor Panorama

Jiale Xu¹ Jia Zheng² Yanyu Xu³ Rui Tang² Shenghua Gao^{1,4*}

¹ShanghaiTech University ²KooLab, Manycore

³Institute of High Performance Computing, A*STAR

⁴Shanghai Engineering Research Center of Intelligent Vision and Imaging

{xujl1, gaoshh}@shanghaitech.edu.cn {jiajia, ati}@qunhemail.com

xu_yanyu@ihpc.a-star.edu.sg

Abstract

Existing view synthesis methods mainly focus on the perspective images and have shown promising results. However, due to the limited field-of-view of the pinhole camera, the performance quickly degrades when large camera movements are adopted. In this paper, we make the first attempt to generate novel views from a single indoor panorama and take the large camera translations into consideration. To tackle this challenging problem, we first use Convolutional Neural Networks (CNNs) to extract the deep features and estimate the depth map from the source-view image. Then, we leverage the room layout prior, a strong structural constraint of the indoor scene, to guide the generation of target views. More concretely, we estimate the room layout in the source view and transform it into the target viewpoint as guidance. Meanwhile, we also constrain the room layout of the generated target-view images to enforce geometric consistency. To validate the effectiveness of our method, we further build a large-scale photo-realistic dataset containing both small and large camera translations. The experimental results on our challenging dataset demonstrate that our method achieves state-of-the-art performance. The project page is at <https://github.com/bluestyle97/PNVS>.

1. Introduction

With the popularity of 360° cameras, panoramas have been widely used in many emerging domains such as Virtual Reality (VR). In a typical VR application, the device displays a 360° virtual scene, which can respond to 6 degree-of-freedom (DoF) head motion and give the user an immersive feeling. However, owing to the tedious image collection process, the panoramas are usually captured at a limited set of locations in practice, which restricts the

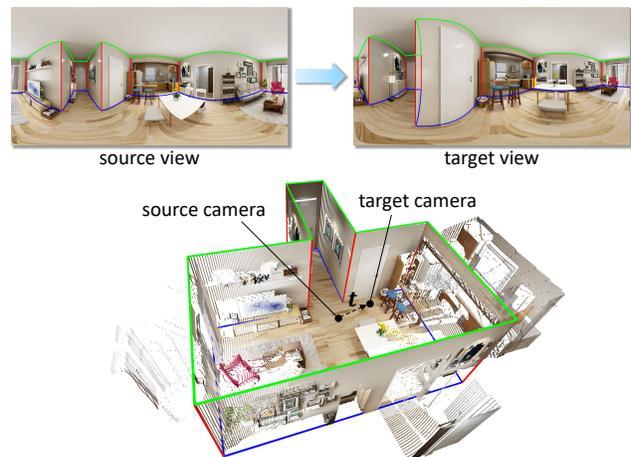


Figure 1. Panoramic novel view synthesis. Our goal is to generate a target-view panorama from the source-view panorama with camera translation t . The green, red, and blue lines represent the ceiling-wall boundaries, wall-wall boundaries, and floor-wall boundaries of the room layout, respectively.

DoF of scene viewing. With the expectation of providing a free-viewpoint scene visualization experience, we make the first attempt to address the problem of panoramic novel view synthesis from a single panorama.

In this paper, we constrain the panoramic view synthesis problem in the indoor scenario on account of its commonness in typical applications. Previous work [30, 36, 39] has shown promising results on novel view synthesis from a single perspective image. However, the performance quickly degrades when larger camera rotations and translations are adopted. Due to the limited field-of-view (FoV) of a pinhole camera, it is arduous to extrapolate the large unseen areas caused by violent camera motion. In contrast, a panorama inherently supports the rotational viewpoint change. Thus, we only need to consider camera translations. Furthermore, 360° FoV provides omnidirectional information, making it possible to consider larger camera translations. By synthe-

*Corresponding author.

sizing panoramic novel views, we can create new 360° contents to achieve 6-DoF scene viewing, which could potentially benefit many applications, such as virtual house tours.

The main challenge of novel view synthesis lies in recovering the missing areas caused by viewpoint change, and the difficulty is amplified when considering large camera translations. Fortunately, a panorama contains more structural information than a perspective image that can be exploited to reduce the difficulty. Previous work on image inpainting [23, 28] has proven the effectiveness of structural information to guide the content generation process. In the indoor scenario, the most common and easy-obtained structural information is the room layout, *i.e.*, the ceiling-wall boundaries, floor-wall boundaries, and wall-wall boundaries. The synthesized images have to keep the room layout reasonable, especially when large camera translations are adopted.

Inspired by the state-of-the-art view synthesis framework [39], we propose a novel method to tackle the panoramic view synthesis problem and exploit the room layout as a prior and geometric constraint. The proposed method is composed of three stages. In the first stage, we use CNNs to extract a dense feature map, a depth map, and room layout from the source-view panorama. In the second stage, we transform the extracted feature map and room layout into the target view with a spherical geometric transformation process and fuse them to synthesize the target panorama. In the final stage, we estimate the room layout of the synthesized panorama and enforce the estimated layout consistent with the transformed target-view layout in the preceding stage.

To validate the effectiveness of our method and facilitate the research on this novel task, we further build a large-scale photo-realistic dataset upon Structured3D dataset [53]. The rendered images are high-fidelity, making the dataset close to realistic application scenarios. Besides the typical settings of previous work, our dataset also considers large camera translations to push the boundaries of the view synthesis task. We split our dataset into an easy set and a hard set according to the camera translation. The easy set contains target panoramas with small camera translations ranging from 0.2 m to 0.3 m, including 13 080 training images and 1791 testing images. The hard set contains target panoramas with large camera translations ranging from 1 m to 2 m, including 17 661 training images and 2279 testing images.

In summary, the main contributions of this paper are as follows: (i) We are the first to tackle the problem of synthesizing panoramic novel views from a single indoor panorama. (ii) We propose a novel layout-guided method to tackle this challenging task, which is able to handle large camera translations. (iii) We build a new high-quality and challenging dataset for this novel task, which contains small and large camera translations. (iv) The experimental results

demonstrate that our method achieves state-of-the-art performance on this novel task and can be generalized to real datasets.

2. Related Work

Novel view synthesis. Previous work on novel view synthesis is based on heterogeneous settings, and we concentrate on learning-based methods here. The most straightforward idea is to perform image generation directly [10, 52]. Instead, some methods [25, 55, 56] estimate the 2D correspondences between the source image and the target image first, *i.e.*, appearance flows, to tackle this problem. More intuitively, many methods adopt the modeling-rendering pattern, which means modeling the scene first and then rendering it to novel views. Following this scheme, a variety of middle representations have been exploited, such as point cloud [24, 39], learned representations [3, 27, 31], layered depth image (LDI) [30, 37], multi-plane images (MPI) [4, 20, 21, 33, 36, 54] and neural radiance fields [22, 49].

Compared with common perspective settings, attempts on view synthesis from panoramas are still very limited so far. Some previous work [8, 29] has tackled the problem of 6-DoF viewing from a pre-captured 360° video to promote VR applications. Huang et al. [8] propose to reconstruct a point cloud from the input 360° video to achieve real-time 6-DoF video playback with a VR device. Serrano et al. [29] present a method for adding parallax and real-time playback of 360° videos, which relies on a layered scene representation. Recently, inspired by the MPI representation, Lin et al. [17] and Attal et al. [2] propose multi-depth panorama (MDP) and multi-sphere image (MSI) representation respectively, to conduct 6-DoF rendering from 360° imagery. However, their settings are quite different from ours. Lin et al. [17] take the images captured by a multi-camera 360° panorama capture rigs as input, while the input of [2] is a 360° stereo video.

Image inpainting. Image inpainting aims to complete the missing region in an image. Traditional patch-based methods [19] and diffusion-based methods [32] are the pioneering work to tackle this problem. In the deep learning era, CNN-based methods [16, 18, 42, 48] and GAN-based methods [14, 44, 45, 46] draw more attention from the research community due to their favorable performance. Several inpainting methods have demonstrated the effectiveness of using structural information. Back in the non-deep-learning era, Sun et al. [35] and Huang et al. [9] have proposed to use line and planar structure to guide the image inpainting process. Various learning-based methods [15, 23, 28, 41] also exploit the structural information. Although the structural information differs in specific forms, *e.g.*, edges, gradients, sketches, or foreground contours, they all act as a global structural prior as well as a geometric constraint and have shown reliable effectiveness.

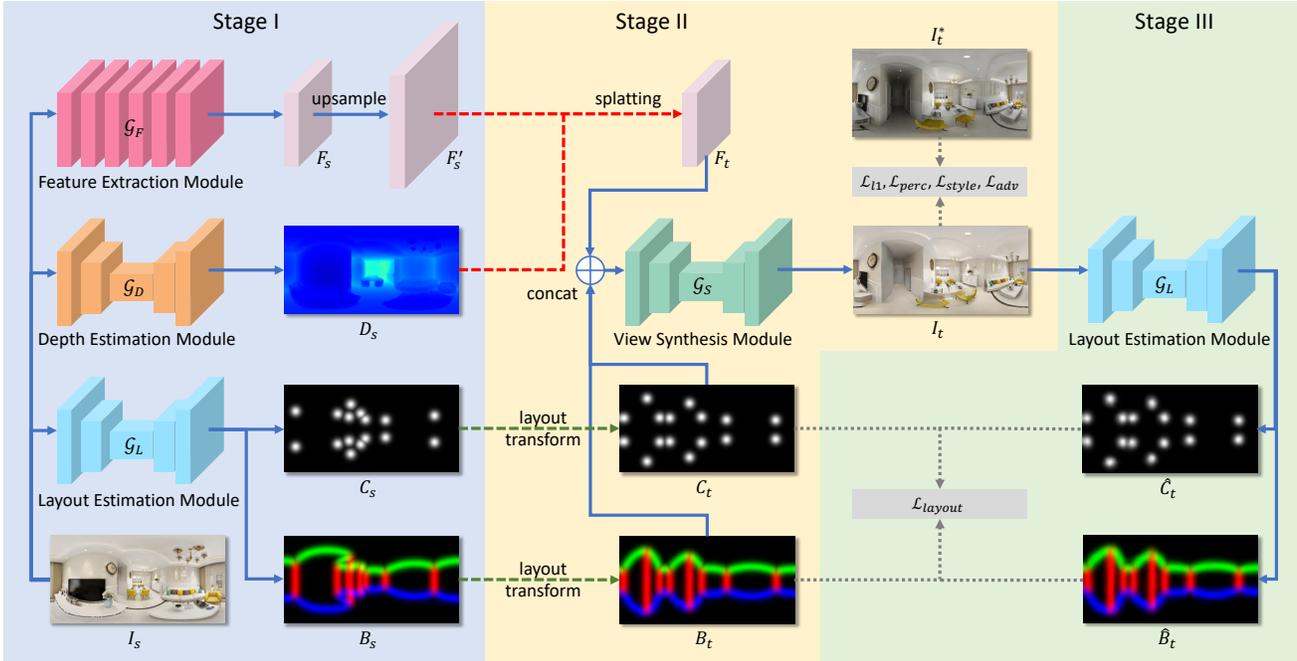


Figure 2. An overview of our pipeline. In the first stage, the network extracts a dense feature map F_s from the source-view panorama I_s as contextual information, and estimates its depth D_s as well as room layout $L_s = \{B_s, C_s\}$ as structural information. In the second stage, F_s and L_s are transformed into the target viewpoint with a forward splatting operation and a layout transformation process to form F_t and $L_t = \{B_t, C_t\}$, respectively. Then, F_t and L_t are fused together to synthesize the target-view panorama I_t . In the final stage, we estimate the room layout of synthesized panorama I_t and enforce it consistent with the transformed layout L_t .

Layout and depth estimation on panoramas. Room layout estimation from a panorama has been sufficiently studied. LayoutNet [59] predicts a boundary probability map and a corner probability map from the input panorama, then estimates the room layout with a Manhattan layout optimizer. HorizonNet [34] further simplifies the layout representation by replacing the 2D probability maps with 1D vectors. DuLa-Net [43] exploits an equirectangular panorama branch and a perspective ceiling-view branch to tackle this problem.

OmniDepth [58] transfers the monocular depth estimation task to panoramas first. Zioulis et al. [57] propose a self-supervised method to estimate panoramic depth, which uses panoramic view synthesis as a proxy task. BiFuse [38] adopts a two-branch architecture to predict panoramic depth. Jin et al. [11] propose to leverage the geometric structure of a scene, *i.e.*, different room layout representations, to conduct this problem. Recently, Zeng et al. [47] propose to jointly learn the panoramic layout and depth since they are tightly intertwined.

3. Method

Given a source-view panorama $I_s \in \mathbb{R}^{H \times W \times 3}$ at the source camera position $p_s \in \mathbb{R}^3$ and a target camera position $p_t \in \mathbb{R}^3$, our goal is to synthesize a target-view

panorama $I_t \in \mathbb{R}^{H \times W \times 3}$. Since the panorama inherently support camera rotations, we can assume that the cameras always face the same direction and only consider the camera translations here.

Our method follows the classical modeling-rendering pattern. We first conduct depth estimation on the source-view image to obtain the 3D scene. Since the estimated 3D scene is inaccurate and noisy, directly rendering new views from it leads to severe shape distortion and pixel misalignment. Inspired by the recent success in room layout estimation [34, 59], we exploit it as a structural prior and geometric constraint to guide the view synthesis process. The three-stage pipeline is shown in Figure 2.

3.1. Feature Extraction and Structure Estimation

In the first stage, we extract contextual and structural information from the source-view image. Concretely, the feature extraction module G_F extracts a dense feature map F_s from I_s as contextual information, and the layout estimation module G_L estimates the room layout L_s from I_s as structural information. To build the scene geometry, the depth estimation module G_D predicts a depth map D_s from I_s .

Previous work [39] has shown that synthesizing novel views from high-level features containing scene semantics instead of simple RGB colors leads to better results. Fol-

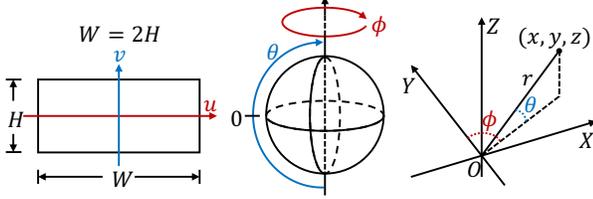


Figure 3. The geometric interpretation of the relationships among \mathcal{P} , \mathcal{S} , and \mathcal{C} . The left and middle pictures explain the relationship between coordinates $(u, v) \in \mathcal{P}$ and $(\phi, \theta) \in \mathcal{S}$, which are borrowed from [40]. The right picture explains the relationship between coordinates $(\phi, \theta, r) \in \mathcal{S}$ and $(x, y, z) \in \mathcal{C}$, O is the camera center.

Following this spirit, our model utilizes a CNN \mathcal{G}_F to extract a dense feature map $\mathbf{F}_s \in \mathbb{R}^{H \times W \times C}$ from the input RGB panorama $\mathbf{I}_s \in \mathbb{R}^{H \times W \times 3}$.

Similar to LayoutNet [59], the layout estimation module \mathcal{G}_L predicts a boundary map $\mathbf{B}_s \in \mathbb{R}^{H \times W \times 3}$ and a corner map $\mathbf{C}_s \in \mathbb{R}^{H \times W}$. With \mathbf{B}_s and \mathbf{C}_s , we follow the standard post-processing procedure of LayoutNet to obtain the 2D positions of room corners $\mathbf{L}_s \in \mathbb{R}^{N \times 2}$.

The feature extraction, depth and layout estimation process can be represented as:

$$\mathbf{F}_s = \mathcal{G}_F(\mathbf{I}_s), \mathbf{D}_s = \mathcal{G}_D(\mathbf{I}_s), \{\mathbf{B}_s, \mathbf{C}_s\} = \mathcal{G}_L(\mathbf{I}_s). \quad (1)$$

\mathcal{G}_F is implemented as a series of ResNet blocks, and C is set to 64. We follow the architectures of Hu et al. [7] and LayoutNet [59] to implement \mathcal{G}_D and \mathcal{G}_L , respectively.

3.2. Viewpoint Transformation and View Synthesis with Layout Prior

In the second stage, we transform the source-view contextual \mathbf{F}_s and structural information \mathbf{L}_s into the target view and synthesize the target panorama \mathbf{I}_t .

The viewpoint transformation is a spherical geometric transformation process. To make it easier to understand, we first clarify several related coordinate systems and show their relationships in Figure 3.

- **Panoramic pixel grid coordinate system \mathcal{P} :** Coordinates $(u, v) \in \mathcal{P}$ represent the pixel at the u -th column and the v -th row pixel on the panoramic image plane, where $u \in [0, W)$ and $v \in [0, H)$.
- **Spherical polar coordinate system \mathcal{S} :** The origin is the camera position. Coordinates $(\phi, \theta, r) \in \mathcal{S}$ represents a point whose longitude is ϕ , latitude is θ , and distance from the origin is r , where $\phi \in [-\pi, \pi]$, $\theta \in [-\pi/2, \pi/2]$, $r > 0$.
- **3D Cartesian camera coordinate system \mathcal{C} :** The origin is the camera position. The X, Y, Z axes points

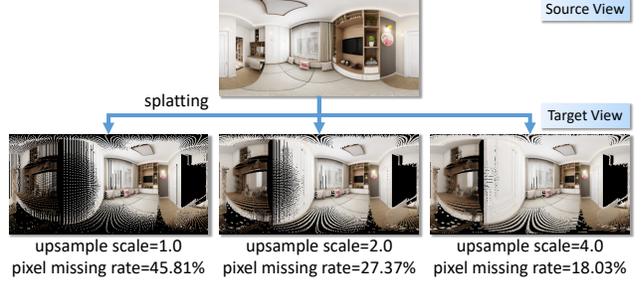


Figure 4. The influence of feature map upsampling. In the splating operation, each source-view pixel contributes to 4 neighboring target pixels. Therefore, using a denser source-view feature map leads to smaller holes on the splatted feature map. Since it is hard to visualize the feature, we show the splating results of RGB values under different upsampling scales instead.

rightward, forward and upward, respectively. Coordinates $(x, y, z) \in \mathcal{C}$ represents the position of a 3D point relative to the origin, where $x, y, z \in \mathbb{R}$.

Feature map view transformation. To transform the source-view feature map \mathbf{F}_s into the target view, we need to map each source-view pixel $(u_s, v_s) \in \mathcal{P}_s$ to a target-view pixel $(u_t, v_t) \in \mathcal{P}_t$, which can be accomplished by a series of coordinate transformations from \mathcal{P}_s to \mathcal{P}_t :

$$f = f_{\mathcal{S}_t \rightarrow \mathcal{P}_t} \circ f_{\mathcal{C}_t \rightarrow \mathcal{S}_t} \circ f_{\mathcal{C}_s \rightarrow \mathcal{C}_t} \circ f_{\mathcal{S}_s \rightarrow \mathcal{C}_s} \circ f_{\mathcal{P}_s \rightarrow \mathcal{S}_s}, \quad (2)$$

where $f_{A \rightarrow B}$ denotes a coordinate transformation from coordinate system A to B , and \circ denotes the composition of transformations. We refer the readers to supplementary material for detailed coordinate transformation equations.

By conducting Eq. (2) on a source-view pixel $(u_s, v_s) \in \mathcal{P}_s$, we can obtain its corresponding target-view pixel position $(u_t, v_t) = f(u_s, v_s) \in \mathcal{P}_t$. With the pixel correspondences, we adopt a differentiable rendering approach [37, 57] to generate target-view feature map $\mathbf{F}_t \in \mathbb{R}^{H \times W \times C}$. Concretely, we splat the feature vector at each pixel of \mathbf{F}_s onto its corresponding pixel position on the target-view panorama plane with the bilinear interpolation. To resolve the conflicts caused by the many-to-one mapping problem, a soft z-buffering is adopted, which can be formulated as:

$$\mathbf{I}_t(u_t, v_t) = \frac{\sum_{(u_s, v_s)} \mathbf{I}_s(u_s, v_s) \exp(-\mathbf{D}_s(u_s, v_s)/d_{\max})}{\sum_{(u_s, v_s)} \exp(-\mathbf{D}_s(u_s, v_s)/d_{\max}) + \epsilon}, \quad (3)$$

where $d_{\max} = 10$ is a pre-defined maximum depth value, ϵ is a small constant for numerical stability.

Feature map upsampling. Since large camera translations are taken into consideration, directly splating leads to large missing areas (*i.e.*, holes) on \mathbf{F}_t , making it difficult to inpaint. Besides, some areas that are supposed to be occluded will be unexpectedly exposed because the areas occluding

them are missing. To tackle this problem, we upsample F_s to F'_s before the forward splatting:

$$F'_s = \text{Conv}(\text{Upsample}(F_s)), \quad (4)$$

where Conv denotes a convolution layer, $\text{Upsample} : \mathbb{R}^{H \times W \times C} \rightarrow \mathbb{R}^{2H \times 2W \times C}$ denotes the nearest upsampling layer. This operation can significantly reduce the missing areas in F_t and make it easier for the network to inpaint. Figure 4 demonstrates such effect.

Layout view transformation. The layout transformation from $L_s \in \mathbb{R}^{N \times 2}$ to $L_t \in \mathbb{R}^{N \times 2}$ is similar to feature map transformation but have some differences. Be aware that we cannot know the depths of layout corners from D_s since they may be occluded by foreground objects. Thus, we estimate the depth of each corner with the camera height h . We provide the details of the layout transformation process in the supplementary material.

To utilize L_t , we draw a boundary map $B_t \in \mathbb{R}^{H \times W \times 3}$ and a corner map $C_t \in \mathbb{R}^{H \times W}$ from L_t with Gaussian blurring. Then, we feed them into the view synthesis module \mathcal{G}_S to serve as structural prior and constrain the synthesis of target-view panorama.

View synthesis with layout prior. With the transformed target-view contextual information F_t and structural information $\{B_t, C_t\}$, the view synthesis module \mathcal{G}_S fuses them all together and synthesizes the target-view panorama I_t :

$$I_t = \mathcal{G}_S(F_t \oplus B_t \oplus C_t), \quad (5)$$

where \oplus denotes the concatenation operation along the channel dimension. We adopt an architecture similar to [23] to implement \mathcal{G}_S .

3.3. Layout Consistency Constraint

In order to maximize the use of room layout guidance, we introduce a layout-consistency loss to force the synthesized panorama I_t to keep the consistency of room layout. Specifically, we feed I_t into the layout estimation module to obtain \hat{B}_t and \hat{C}_t . Then, we compare them with B_t, C_t and calculate the layout consistency loss as:

$$\mathcal{L}_{\text{layout}} = \text{BCE}(\hat{B}_t, B_t) + \text{BCE}(\hat{C}_t, C_t), \quad (6)$$

where BCE represents the binary cross entropy loss.

3.4. Losses

During training, the layout estimation module \mathcal{G}_L and the depth estimation module \mathcal{G}_D are pretrained under the supervision of ground-truth layout and depth, respectively.

Given the synthesized panorama I_t and the ground-truth panorama I_t^* , the rest model is trained with ℓ_1 loss, perceptual loss [12], style loss [5], adversarial loss [6] and layout

consistency loss. Their functions can be formulated as:

$$\mathcal{L}_{\ell_1} = \mathbb{E}[\|I_t - I_t^*\|_1], \quad (7)$$

$$\mathcal{L}_{\text{perc}} = \mathbb{E}\left[\sum_i \|\psi_i(I_t) - \psi_i(I_t^*)\|_1\right], \quad (8)$$

$$\mathcal{L}_{\text{style}} = \mathbb{E}\left[\sum_j \|G_j(I_t) - G_j(I_t^*)\|_1\right], \quad (9)$$

$$\mathcal{L}_{\text{adv}} = \mathbb{E}[\log \mathcal{D}(I_t^*)] + \mathbb{E}[\log(1 - \mathcal{D}(\mathcal{G}(I_s)))]], \quad (10)$$

where ψ_i denotes the activation map of the i -th layer of a pretrained VGG-19, G_j is a $C_j \times C_j$ Gram matrix calculated from ψ_j . \mathcal{G} denotes the generator, i.e., our model, and \mathcal{D} denotes the discriminator, $I_t = \mathcal{G}(I_s)$.

Finally, the total loss is calculated as:

$$\mathcal{L} = \mathcal{L}_{\ell_1} + \mathcal{L}_{\text{perc}} + \lambda \mathcal{L}_{\text{style}} + \mathcal{L}_{\text{adv}} + \mathcal{L}_{\text{layout}}, \quad (11)$$

where λ is set to 100 in our experiments.

4. Experiments

In this section, we conduct experiments to validate the performance of our proposed method. Due to the space limitation, we refer the readers to the supplementary material for extensive qualitative results and failure cases.

4.1. Experimental Setup

Implementation details. Our model is implemented with PyTorch library [26] and trained on two NVIDIA TITAN V GPU devices. We use the Adam [13] optimizer with $\beta_1 = 0.9$ and $\beta_2 = 0.999$. The batch size is set to 4. Specifically, we first train the depth estimation module and the layout estimation module for 30 epochs to make them converge. Then, we freeze them and train the rest model for another 50 epochs. The learning rate for both the generator and the discriminator is set to 1×10^{-4} . After 30 epochs, we reduce the learning rate by 10 times.

Dataset. Our panoramic view synthesis dataset is built upon Structured3D dataset [53]. Each panorama in Structured3D corresponds to a different room. We regard original images as source views and render three target views for each source view. Our dataset is divided into two sets with different target-view camera selection strategies: (i) **an easy set**: the camera translation ranges from 0.2 m to 0.3 m along random directions, which is a typical translational distance in previous view synthesis work. (ii) **a hard set**: the camera translation ranges from 1.0 m to 2.0 m along random directions, which is a very challenging setting and has rarely been considered. To clarify the difficulties of our settings, we visualize the relationship between the pixel missing rate after the splatting operation and the camera translation distance in the supplementary material. The resolution of the

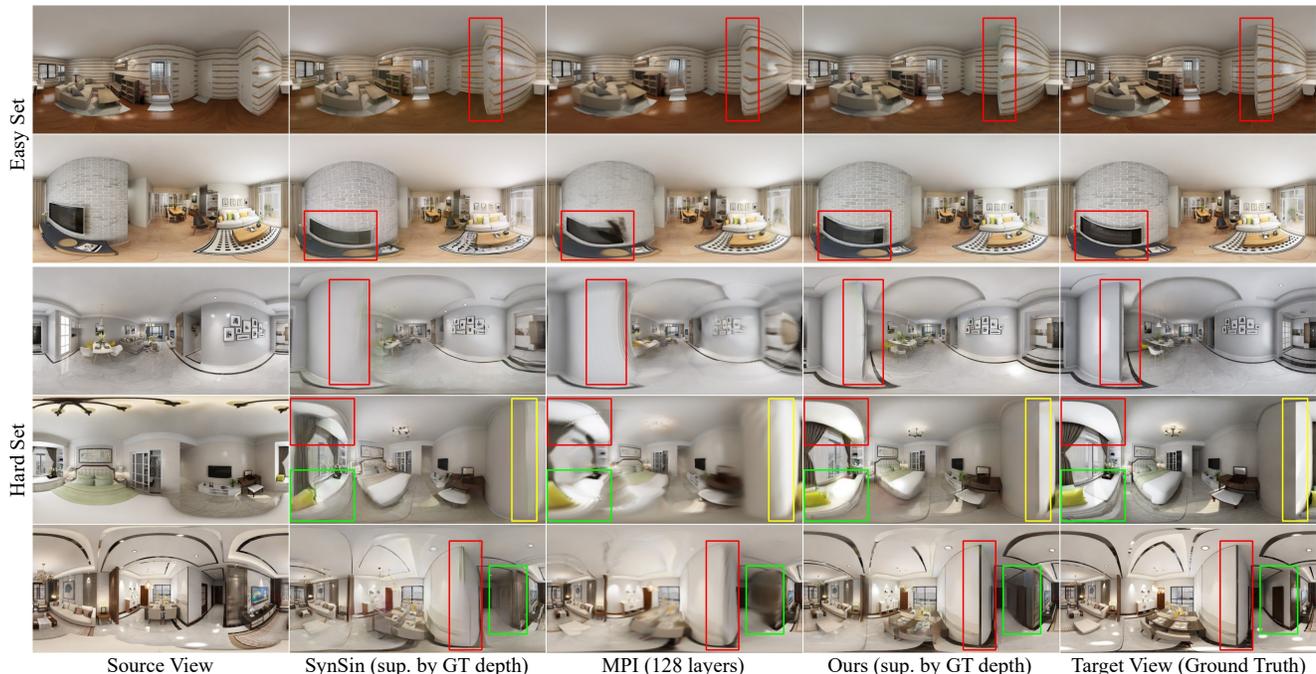


Figure 5. Qualitative view synthesis results on our dataset. The first two rows are from the easy set, while the last three rows are from the hard set. We highlight the major differences using the bounding boxes. More results are shown in the supplementary material.

Methods	Easy Set (0.2 m to 0.3 m)			Hard Set (1.0 m to 2.0 m)		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
SynSin (end-to-end)	16.88	0.7433	0.1946	15.51	0.7298	0.2462
SynSin (supervised by GT depth)	18.04	0.7853	0.1714	17.02	0.7827	0.2119
SynSin (GT depth as input)	18.79	0.8127	0.1559	18.02	0.8181	0.1724
MPI (32 layers)	18.32	0.8044	0.2150	16.53	0.7725	0.3098
MPI (64 layers)	18.08	0.7984	0.2192	16.56	0.7769	0.3051
MPI (128 layers)	18.23	0.8015	0.2170	16.50	0.7776	0.3015
Ours (supervised by GT depth)	19.35	0.8373	0.1351	17.50	0.8148	0.1769
Ours (GT depth as input)	20.52	0.8727	0.1192	18.53	0.8552	0.1544
Ours (GT depth & GT layout as input)	20.83	0.8743	0.1150	18.95	0.8593	0.1454

Table 1. Quantitative results on our dataset.

panorama in our dataset is 512×1024 . In all experiments, we take panoramas of 256×512 as input.

Evaluation metrics. We quantify the performance of our method with three metrics: (i) Peak Signal-to-Noise Ratio (PSNR), (ii) Structural Similarity (SSIM), and (iii) Learned Perceptual Image Patch Similarity (LPIPS) [50].

4.2. Experimental Results

Methods for comparison. We compare our approach with two state-of-the-art single-image view synthesis methods: point-cloud-based method SynSin [39]¹ and MPI-based

method [36]² (in short, MPI). We choose SynSin [39] and MPI [36] as our baselines on account of their good performance as well as the code availability and portability.

We modify SynSin and MPI to make them applicable to panoramas. For SynSin, the perspective projection in the differentiable renderer is replaced with equirectangular projection. Every 3D point is projected to a circular region of the target-view panorama plane with α -compositing. For MPI, we use the same network as [36] to infer a multi-sphere image (MSI) centered at the camera position, which is similar to [2]. Then, we cast the rays from the target view

¹<https://github.com/facebookresearch/synsin>

²https://github.com/google-research/google-research/tree/master/single_view_mpi



Figure 6. The effectiveness of room layout guidance. We highlight the major differences using the bounding boxes.

onto the MSI and use the bilinear interpolation to perform view synthesis.

Quantitative evaluation. SynSin estimates the depth of the source view with an end-to-end training scheme. For a fair comparison, we also train SynSin with ground-truth depth as supervision to meet our setting. Besides, we evaluate SynSin and our model with ground-truth depth as input to investigate their upper-bound performance. For MPI, we set the number of layers as 32, 64, and 128, respectively.

As Table 1 shows, the performance of SynSin and our model increases when using more accurate depth, and the performance of MPI increases when more layers are adopted. Our method outperforms the other two methods in all metrics. When exploiting ground-truth depth as input, our model shows a higher upper bound than SynSin. In addition, we can further boost the performance of our model by adopting the ground-truth layout as input.

Qualitative evaluation. Figure 5 shows the qualitative results of three methods. As one can see, our approach maintains more plausible visual details. Especially the results on the hard set show that our method can maintain the room structure well when large camera translation is adopted. However, the other two approaches have artifacts, such as blurring layout boundaries and distortion.

User study. For more complete qualitative comparison, we further conduct a user study. We first sample 50 images from the easy set and the hard set, respectively. Then, we recruit 50 volunteers and show them the synthesized target views of the three methods in random order, and the ground truth. We ask them to select the closest one to the ground truth and report the percentage of volunteers who prefer a given method. As shown in Table 2, volunteers prefer our method over the other two methods in both sets.

Methods	Easy Set (0.2 m to 0.3 m)	Hard Set (1.0 m to 2.0 m)
SynSin	0.21	0.18
MPI	0.13	0.09
Ours	0.66	0.73

Table 2. User study on our dataset.

4.3. Ablation Studies

We conduct some ablation studies to verify the effectiveness of each component in our proposed method. The results are shown in Table 3 and discussed in detail next.

Feature map upsampling. We try to remove the feature map upsampling from our model. By comparing the quantitative results with the complete model, we can see that the upsampling operation leads to a performance improvement. As shown in Figure 4, the upsampling operation can remarkably reduce the missing pixels after the splatting operation, which abates the contextual information loss and makes the inpainting easier for the view synthesis module.

Room layout. To show the effectiveness of room layout, we remove either the layout prior or layout consistency loss, or both of them. The results show that both the layout prior and layout consistency loss contribute to the performance improvement. When using the layout prior, all metrics increase by a large margin. The layout consistency loss leads to better perceptual quality, which is indicated by the improvement of LPIPS. When using both of them, the performance of the model reaches the peak. Besides, Figure 6 visualizes some target-view results synthesized with or without layout guidance on the hard set. We can see that the model could utilize the structural information provided by

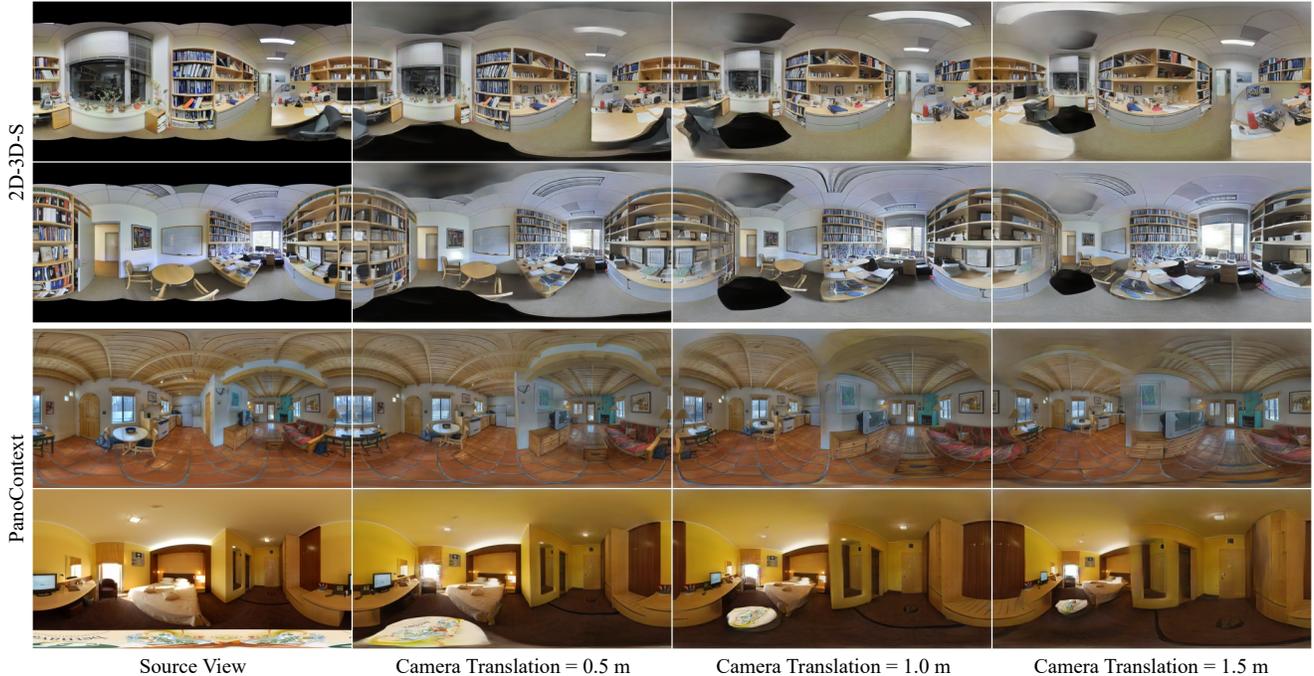


Figure 7. Panoramic view synthesis on 2D-3D-S dataset and PanoContext dataset. More results are shown in the supplementary material.

Components			Easy Set (0.2 m to 0.3 m)			Hard Set (1.0 m to 2.0 m)		
Upsampling	Prior	Consistency	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
			18.44	0.8067	0.1511	16.48	0.7735	0.2007
\checkmark			18.57	0.8110	0.1504	16.70	0.7802	0.1965
\checkmark		\checkmark	18.49	0.8173	0.1503	16.93	0.7981	0.2066
\checkmark	\checkmark		19.40	0.8371	0.1455	17.46	0.8146	0.1917
	\checkmark	\checkmark	19.12	0.8354	0.1390	17.33	0.8090	0.1813
\checkmark	\checkmark	\checkmark	19.35	0.8373	0.1351	17.50	0.8148	0.1769

Table 3. Ablation studies on our dataset.

the room layout to synthesize target-view panoramas with more visual-plausible layout structures.

4.4. Panoramic View Synthesis on Real Datasets

To verify the generalization ability of our method, we also conduct panoramic view synthesis on real datasets. We train the model on our dataset, then directly test it on 2D-3D-S dataset [1] and PanoContext dataset [51]. For each dataset, we set the camera translation distance as 0.5 m, 1.0 m and 1.5 m, along the x -axis or y -axis randomly. Figure 7 shows the qualitative results. The results show that our method generalizes well to the real scenes and has a great potentiality for real-world application. To be noted, the vertical FoV of the panorama in 2D-3D-S dataset does not cover 180° . Thus, there are wavy black regions on the top and the bottom, making the synthesis more challenging and unrealistic to be inpainted completely.

5. Conclusion

In this paper, we explore synthesizing 360° novel views from a single indoor panorama and consider large camera translations. We propose a novel layout-guided method that exploits the room layout as a prior and geometric constraint. We also build a large-scale dataset for this novel task. The experiments show that our method achieves state-of-the-art performance and generalizes well on real datasets. In the future, we plan to exploit more general structures (*e.g.*, planes or wireframes) and extend this idea to outdoor scenes.

Acknowledgements. This work was supported by the National Key R&D Program of China (2018AAA0100704), the National Natural Science Foundation of China (61932020), Science and Technology Commission of Shanghai Municipality (20ZR1436000), and “Shuguang Program” by Shanghai Education Development Foundation and Shanghai Municipal Education Commission.

References

- [1] Iro Armeni, Sasha Sax, Amir R Zamir, and Silvio Savarese. Joint 2d-3d-semantic data for indoor scene understanding. *CoRR*, abs/1702.01105, 2017. [8](#)
- [2] Benjamin Attal, Selena Ling, Aaron Gokaslan, Christian Richardt, and James Tompkin. MatryODShka: Real-time 6DoF video view synthesis using multi-sphere images. In *ECCV*, pages 441–459, 2020. [2](#), [6](#)
- [3] Inchang Choi, Orazio Gallo, Alejandro Troccoli, Min H Kim, and Jan Kautz. Extreme view synthesis. In *ICCV*, pages 7781–7790, 2019. [2](#)
- [4] John Flynn, Michael Broxton, Paul Debevec, Matthew Duvall, Graham Fyffe, Ryan Overbeck, Noah Snavely, and Richard Tucker. Deepview: View synthesis with learned gradient descent. In *CVPR*, pages 2367–2376, 2019. [2](#)
- [5] Leon A. Gatys, Alexander S. Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *CVPR*, pages 2414–2423, 2016. [5](#)
- [6] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NeurIPS*, pages 2672–2680, 2014. [5](#)
- [7] Junjie Hu, Mete Ozay, Yan Zhang, and Takayuki Okatani. Revisiting single image depth estimation: Toward higher resolution maps with accurate object boundaries. In *WACV*, pages 1043–1051, 2019. [4](#)
- [8] Jingwei Huang, Zhili Chen, Duygu Ceylan, and Hailin Jin. 6-dof vr videos with a single 360-camera. In *VR*, pages 37–44, 2017. [2](#)
- [9] Jia-Bin Huang, Sing Bing Kang, Narendra Ahuja, and Johannes Kopf. Image completion using planar structure guidance. *ACM TOG*, 33(4):1–10, 2014. [2](#)
- [10] Rui Huang, Shu Zhang, Tianyu Li, and Ran He. Beyond face rotation: Global and local perception gan for photorealistic and identity preserving frontal view synthesis. In *ICCV*, pages 2439–2448, 2017. [2](#)
- [11] Lei Jin, Yanyu Xu, Jia Zheng, Junfei Zhang, Rui Tang, Shugong Xu, Jingyi Yu, and Shenghua Gao. Geometric structure based and regularized depth estimation from 360° indoor imagery. In *CVPR*, pages 886–895, 2020. [3](#)
- [12] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *ECCV*, pages 694–711, 2016. [5](#)
- [13] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. [5](#)
- [14] Avisek Lahiri, Arnav Kumar Jain, Sanskar Agrawal, Pabitra Mitra, and Prabir Kumar Biswas. Prior guided gan based semantic inpainting. In *CVPR*, pages 13696–13705, 2020. [2](#)
- [15] Jingyuan Li, Fengxiang He, Lefei Zhang, Bo Du, and Dacheng Tao. Progressive reconstruction of visual structure for image inpainting. In *ICCV*, pages 5962–5971, 2019. [2](#)
- [16] Jingyuan Li, Ning Wang, Lefei Zhang, Bo Du, and Dacheng Tao. Recurrent feature reasoning for image inpainting. In *CVPR*, pages 7760–7768, 2020. [2](#)
- [17] Kai-En Lin, Zexiang Xu, Ben Mildenhall, Pratul P. Srinivasan, Yannick Hold-Geoffroy, Stephen DiVerdi, Qi Sun, Kalyan Sunkavalli, and Ravi Ramamoorthi. Deep multi depth panoramas for view synthesis. In *ECCV*, pages 328–344, 2020. [2](#)
- [18] Hongyu Liu, Bin Jiang, Yi Xiao, and Chao Yang. Coherent semantic attention for image inpainting. In *ICCV*, pages 4170–4179, 2019. [2](#)
- [19] Jiaying Liu, Shuai Yang, Yuming Fang, and Zongming Guo. Structure-guided image inpainting using homography transformation. *IEEE TMM*, 20(12):3252–3265, 2018. [2](#)
- [20] Yicun Liu, Jiawei Zhang, Ye Ma, and Jimmy Ren. Self-guided novel view synthesis via elastic displacement network. In *WACV*, pages 164–173, 2020. [2](#)
- [21] Ben Mildenhall, Pratul P Srinivasan, Rodrigo Ortiz-Cayon, Nima Khademi Kalantari, Ravi Ramamoorthi, Ren Ng, and Abhishek Kar. Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *ACM TOG*, 38(4):1–14, 2019. [2](#)
- [22] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, pages 405–421, 2020. [2](#)
- [23] Kamyar Nazeri, Eric Ng, Tony Joseph, Faisal Qureshi, and Mehran Ebrahimi. Edgeconnect: Structure guided image inpainting using edge prediction. In *ICCV Workshop*, 2019. [2](#), [5](#)
- [24] David Novotny, Ben Graham, and Jeremy Reizenstein. Perspectivenet: A scene-consistent image generator for new view synthesis in real indoor environments. In *NeurIPS*, pages 7601–7612, 2019. [2](#)
- [25] Eunbyung Park, Jimei Yang, Ersin Yumer, Duygu Ceylan, and Alexander C Berg. Transformation-grounded image generation network for novel 3d view synthesis. In *CVPR*, pages 3500–3509, 2017. [2](#)
- [26] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. In *NeurIPS Workshop*, 2017. [5](#)
- [27] Eric Penner and Li Zhang. Soft 3d reconstruction for view synthesis. *ACM TOG*, 36(6):1–11, 2017. [2](#)
- [28] Yurui Ren, Xiaoming Yu, Ruonan Zhang, Thomas H. Li, Shan Liu, and Ge Li. Structureflow: Image inpainting via structure-aware appearance flow. In *ICCV*, pages 181–190, 2019. [2](#)
- [29] Ana Serrano, Incheol Kim, Zhili Chen, Stephen DiVerdi, Diego Gutierrez, Aaron Hertzmann, and Belen Masia. Motion parallax for 360 rgbd video. *IEEE TVCG*, 25(5):1817–1827, 2019. [2](#)
- [30] Meng-Li Shih, Shih-Yang Su, Johannes Kopf, and Jia-Bin Huang. 3d photography using context-aware layered depth inpainting. In *CVPR*, pages 8028–8038, 2020. [1](#), [2](#)
- [31] Vincent Sitzmann, Justus Thies, Felix Heide, Matthias Nießner, Gordon Wetzstein, and Michael Zollhofer. Deepvoxels: Learning persistent 3d feature embeddings. In *CVPR*, pages 2437–2446, 2019. [2](#)
- [32] G Sridevi and S Srinivas Kumar. Image inpainting based on fractional-order nonlinear diffusion for image reconstruction. *CSSP*, 38(8):3802–3817, 2019. [2](#)
- [33] Pratul P Srinivasan, Richard Tucker, Jonathan T Barron, Ravi Ramamoorthi, Ren Ng, and Noah Snavely. Pushing the boundaries of view extrapolation with multiplane images. In *CVPR*, pages 175–184, 2019. [2](#)
- [34] Cheng Sun, Chi-Wei Hsiao, Min Sun, and Hwann-Tzong

- Chen. Horizonnet: Learning room layout with 1d representation and pano stretch data augmentation. In *CVPR*, pages 1047–1056, 2019. 3
- [35] Jian Sun, Lu Yuan, Jiaya Jia, and Heung-Yeung Shum. Image completion with structure propagation. *ACM TOG*, pages 861–868, 2005. 2
- [36] Richard Tucker and Noah Snavely. Single-view view synthesis with multiplane images. In *CVPR*, pages 551–560, 2020. 1, 2, 6
- [37] Shubham Tulsiani, Richard Tucker, and Noah Snavely. Layer-structured 3d scene inference via view synthesis. In *ECCV*, pages 302–317, 2018. 2, 4
- [38] 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 *CVPR*, pages 462–471, 2020. 3
- [39] Olivia Wiles, Georgia Gkioxari, Richard Szeliski, and Justin Johnson. Synsin: End-to-end view synthesis from a single image. In *CVPR*, pages 7467–7477, 2020. 1, 2, 3, 6
- [40] Jianxiong Xiao. 3d geometry for panorama, 2012. 4
- [41] Wei Xiong, Jiahui Yu, Zhe Lin, Jimei Yang, Xin Lu, Connelly Barnes, and Jiebo Luo. Foreground-aware image inpainting. In *CVPR*, pages 5840–5848, 2019. 2
- [42] Zhaoyi Yan, Xiaoming Li, Mu Li, Wangmeng Zuo, and Shiguang Shan. Shift-net: Image inpainting via deep feature rearrangement. In *ECCV*, pages 1–17, 2018. 2
- [43] Shang-Ta Yang, Fu-En Wang, Chi-Han Peng, Peter Wonka, Min Sun, and Hung-Kuo Chu. Dula-net: A dual-projection network for estimating room layouts from a single rgb panorama. In *CVPR*, pages 3363–3372, 2019. 3
- [44] Zili Yi, Qiang Tang, Shekoofeh Azizi, Daesik Jang, and Zhan Xu. Contextual residual aggregation for ultra high-resolution image inpainting. In *CVPR*, pages 7508–7517, 2020. 2
- [45] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. Generative image inpainting with contextual attention. In *CVPR*, pages 5505–5514, 2018. 2
- [46] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. Free-form image inpainting with gated convolution. In *ICCV*, pages 4471–4480, 2019. 2
- [47] Wei Zeng, Sezer Karaoglu, and Theo Gevers. Joint 3d layout and depth prediction from a single indoor panorama image. In *ECCV*, pages 666–682, 2020. 3
- [48] Yanhong Zeng, Jianlong Fu, Hongyang Chao, and Baining Guo. Learning pyramid-context encoder network for high-quality image inpainting. In *CVPR*, pages 1486–1494, 2019. 2
- [49] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. Nerf++: Analyzing and improving neural radiance fields. *CoRR*, abs/2010.07492, 2020. 2
- [50] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, pages 586–595, 2018. 6
- [51] Yinda Zhang, Shuran Song, Ping Tan, and Jianxiong Xiao. Panocontext: A whole-room 3d context model for panoramic scene understanding. In *ECCV*, pages 668–686, 2014. 8
- [52] Bo Zhao, Xiao Wu, Zhi-Qi Cheng, Hao Liu, Zequn Jie, and Jiashi Feng. Multi-view image generation from a single-view. In *ACM MM*, pages 383–391, 2018. 2
- [53] Jia Zheng, Junfei Zhang, Jing Li, Rui Tang, Shenghua Gao, and Zihan Zhou. Structured3d: A large photo-realistic dataset for structured 3d modeling. In *ECCV*, pages 519–535, 2020. 2, 5
- [54] Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. Stereo magnification: Learning view synthesis using multiplane images. *ACM TOG*, 37(4), 2018. 2
- [55] Tinghui Zhou, Shubham Tulsiani, Weilun Sun, Jitendra Malik, and Alexei A Efros. View synthesis by appearance flow. In *ECCV*, pages 286–301, 2016. 2
- [56] Hao Zhu, Hao Su, Peng Wang, Xun Cao, and Ruigang Yang. View extrapolation of human body from a single image. In *CVPR*, pages 4450–4459, 2018. 2
- [57] Nikolaos Zioulis, Antonis Karakottas, Dimitrios Zarpalas, Federico Alvarez, and Petros Daras. Spherical view synthesis for self-supervised 360° depth estimation. In *3DV*, pages 690–699, 2019. 3, 4
- [58] Nikolaos Zioulis, Antonis Karakottas, Dimitrios Zarpalas, and Petros Daras. Omnidepth: Dense depth estimation for indoors spherical panoramas. In *ECCV*, pages 448–465, 2018. 3
- [59] Chuhan Zou, Alex Colburn, Qi Shan, and Derek Hoiem. Layoutnet: Reconstructing the 3d room layout from a single RGB image. In *CVPR*, pages 2051–2059, 2018. 3, 4