

Deep Multi Depth Panoramas for View Synthesis

Kai-En Lin¹, Zexiang Xu^{1,3}, Ben Mildenhall², Pratul P. Srinivasan²,
Yannick Hold-Geoffroy³, Stephen DiVerdi³, Qi Sun³, Kalyan Sunkavalli³, and
Ravi Ramamoorthi¹

¹ UC San Diego

² UC Berkeley

³ Adobe Research

Abstract. We propose a learning-based approach for novel view synthesis for multi-camera 360° panorama capture rigs. Previous work constructs RGBD panoramas from such data, allowing for view synthesis with small amounts of translation, but cannot handle the disocclusions and view-dependent effects that are caused by large translations. To address this issue, we present a novel scene representation—Multi Depth Panorama (MDP)—that consists of multiple RGBD α panoramas that represent both scene geometry and appearance. We demonstrate a deep neural network-based method to reconstruct MDPs from multi-camera 360° images. MDPs are more compact than previous 3D scene representations and enable high-quality, efficient new view rendering. We demonstrate this via experiments on both synthetic and real data and comparisons with previous state-of-the-art methods spanning both learning-based approaches and classical RGBD-based methods.

Keywords: 360° panoramas, view synthesis, image-based rendering, virtual reality

1 Introduction

Panoramic images have been widely used to create immersive experiences in virtual environments. Recently, commercial 360° cameras like the Yi Halo and GoPro Odyssey have made panoramic imaging practical. However, the classical panoramic representation only allows for a 3-DoF experience with purely rotational movements; it does not support translational motion, which is necessary for a true 6-DoF immersive experience. While recent work has leveraged panoramic depth to generate 6-DoF motion [23, 27], it is highly challenging for these RGBD-based methods to handle extensive disocclusions and view-dependent effects caused by large movements. Meanwhile, deep learning techniques for view synthesis have demonstrated photo-realistic results [20, 34, 37]; however, these methods are not designed for panoramic inputs and rely on per-viewpoint scene representations that are expensive to store and render.

Our goal is to enable realistic, practical and efficient novel view synthesis that supports translational motion with parallax for complex real scenes. To this end,



Fig. 1. We use 360° image data captured from a Yi-Halo camera (a), which consists of a ring of cameras as shown in (b). We present a learning based approach to reconstruct novel Multi Depth Panoramas (MDP) from these multi-view images, which can synthesize novel view images with both rotational and translational motions. We show panorama results using our MDPs from the center of the device (c) like the standard panorama synthesis, and also from a translated position (d) out of the rig. Note how the camera moves toward the counter. Our method accurately reproduces challenging disocclusion effects as shown in the cropped insets, which are significantly better than previous state-of-the-art methods that are based on RGBD representations [1, 27]. For better details on all the figures, please view the electronic version of this paper.

we propose Multi Depth Panoramas (MDPs)—a novel, compact, geometry-aware panoramic representation inspired by classical layered depth images (LDIs) [28] and learning-based multiple plane images (MPIs) [37]. MDPs consist of a small set of multi-layer $\text{RGBD}\alpha$ (RGB pixel intensity, depth, opacity) panoramas that fully express complex scene geometry and appearance.

We demonstrate the use of MDPs for novel view synthesis from images captured by commercial 360° camera rigs such as the Yi Halo that consist of a sparse array of outward-facing cameras placed in a ring. Previous work proposes limited translational motion [27] using RGBD panoramas reconstructed from this setup [1]. In contrast, we propose an efficient MDP-based rendering scheme that handles challenging occlusions and reflections that cannot be reproduced by state-of-the-art RGBD-based methods (see Fig. 1). Our flexible MDP representation can encode an arbitrary number of layers, degrading gracefully to a regular RGBD panorama when using a single layer. MDPs are also much more compact than previous representations, providing either a similar or significantly better view synthesis quality using fewer layers, as shown in our experiments in Tab. 2. Finally, by encoding the entire 360° panorama in a global representation, MDPs allow for complete panoramic view synthesis, unlike previous methods that focus on synthesizing limited field-of-view images (see Fig. 2).

Our contributions can be summarized as:

- a 360° layered scene representation called Multi Depth Panorama (MDP), offering a more compact and versatile 3D representation than previous work (sec. 3);
- a learning-based method to convert images from common 360° camera setups into our MDP representation (sec. 4.1-4.3);
- an efficient novel view synthesis method based on MDP (sec. 4.4);
- experiments to demonstrate the advantages of the MDP compared to existing representations (sec. 6).

2 Related Work

View synthesis. Novel view synthesis has been extensively studied in computer vision and graphics [7], and has been performed via approaches such as light fields [12, 18] and image-based rendering [5, 10, 29]. Recently, deep learning has been applied in many view synthesis problems; the most successful ones leverage plane sweep volumes to infer depth and achieve realistic geometric-aware view synthesis [11, 16, 34, 37, 9]. We leverage plane sweep volumes to construct per-view MPIs, similar to [37]. By merging multiple MPIs from multiple views in a 360° camera, we construct MDPs for 6-DoF view synthesis.

Layered representation. Layered and volumetric representations have been applied in 3D and view synthesis applications [28, 8, 26, 35, 4, 37]. Unlike a single RGBD image, layered representations make the scene content occluded behind the foreground viewable from side views. In their seminal work, Shade et al. [28] introduce the Layered Depth Images (LDIs) and Sprites with Depth (SwD) representations to render scenes using multiple RGBD images. Our methods shares the RGBD α concept used in SwD, with the distinction that we reason holistically on the 360° scene while SwD decomposes it into multiple independent sprites. LDIs has been extended to a panoramic case [36], where multiple concentric RGBD panoramas are reconstructed. In addition, Zitnick et al. [38] utilizes multiple layers with alpha matting to produce interactive viewpoint videos. In concurrent work, Broxton et al. [3] propose the multi-sphere image (MSI) representation for 6-DoF light field video synthesis, yielding accurate results but involving a memory-heavy process. In contrast with previous representations, we propose a learning-based method to reconstruct multiple memory-efficient RGBD α panoramic images for 6-DoF rendering. Our method uses differentiable rendering to recover smooth object boundaries and specularities that are hard to reproduce by RGBD representations alone.

Zhou et al. [37] leverage a deep network to predict multi plane images (MPIs) [31] for realistic view extrapolation. These MPIs are a dense set of fronto parallel planes with RGBA images at an input view, which enable rendering novel view images locally for the view. Some recent works have extended this local view synthesis technique for large viewing ranges [30, 20]. Mildenhall et al. [20] reconstruct MPIs at multiple views and fuse the 2D renderings from multiple MPIs to enable large viewpoint motion [20]. However, such a 2D image-level fusion tech-




Fig. 2. 3D scene representations: the multiplane image (a) proposes a planar representation of the scene. In contrast, the layered depth image (b) encodes multiple depth pixels along a ray. RGBD panoramas (c) encode the scene distance per-pixel from a reference cylinder. Our multi depth panorama (d) takes inspiration from those previous representations, where depth is encoded on each cylindrical shell.

nique requires expensively storing and rendering multiple view-dependent MPIs. We leverage MPIs for panorama rendering and introduce multi depth panoramas (MDPs) that are view-independent $\text{RGBD}\alpha$ images. We propose to fuse per-view MPIs in a canonical 3D space for MDP reconstruction, which allows for efficient view synthesis from a sparse set of $\text{RGBD}\alpha$ images.

6-DoF rendering. Standard 360° panoramic rendering only allows for 2D experience with 3-DoF rotational motion. Omni-directional stereo [15, 21] has been introduced to provide 3D stereo effects by leveraging panoramic depth [1]. More advanced capture hardware can support 6-DoF rendering of 360° video [23], but there is also research to enable 6-DoF for more accessible hardware. Researchers have designed systems to acquire static scenes with different camera motions over time [14, 19, 2, 13] but they are unable to capture dynamic scenes without artifacts. Other works use ring cameras to capture dynamic scenes and store the scene model in single layer panoramas [32] or as a panorama with a static background [27], but still have significant artifacts around object boundaries, areas of poor depth reconstruction, and for large viewer head motions.

3 Multi Depth Panoramas

Given a 360° camera setup with k cameras, our goal is to infer a layered scene representation that allows for high-quality 6-DoF novel view rendering. To achieve this, we propose the multi depth panorama (MDP) representation.

The MDP representation is inspired by Layered Depth Images (LDI) [28] and Multiplane Images (MPI) [37, 30, 20] (shown in Fig. 2(a) and (b) respectively). LDIs and MPIs have been previously used as *image-based* representations and encode only a limited field-of-view; instead, we are interested in representing the entire 360° scene. Previous work has used RGBD panoramas (Fig. 2(c)) for novel view synthesis for panoramas [27] but this does not encode sufficient information to render the disocclusions that are produced at large translations.

Motivated by these limitations of previous representations, we propose the MDP. An MDP consists of a sparse set of $\text{RGBD}\alpha$ panoramas that divide the 3D space into annulus-like cylindrical regions, as shown in Fig. 2(d). Each layer

l encodes color C_l , depth D_l and pixel transparency α_l of its corresponding region. Consequently, the MDP representation can be denoted by a set of RGBD α layers as $(C_0, D_0, \alpha_0), \dots, (C_L, D_L, \alpha_L)$, where L specifies the maximum number of cylindrical shells. This representation can be thought of as a set of cylindrical shells representing the 3D scene, where in each shell, a pixel is composed of five channels—diffuse RGB color, an alpha channel for occupancy, and a depth channel for finer control over 3D geometry. These components allow for both a more accurate and a more compact representation than previous work. In order to synthesize novel views, we can forward splat each layer onto the target image plane and then perform the standard “over” alpha compositing operation [22].

While similar to an MPI in its use of RGB and α layers, our representation provides the following main benefits over the MPI:

Compactness The main advantage of the proposed MDP is its compactness, yielding an appreciable compression ratio over existing representations. For example, an MPI discretizes the scene into multiple depth planes and requires a large number of such planes (32–128 [37, 20]) for high-quality view synthesis for complex scenes. In contrast, an MDP explicitly stores the depth value allowing it to represent complex scene geometry with a small set of shells; we demonstrate this experimentally in Sec. 6.

Free of depth conflicts Our MDP representation offers a canonical scene representation for merging different views and resolving depth conflicts. MPI planes are created in the viewing frustum of each input camera. In our 360° setup, where individual cameras are placed in an outward-facing ring, there are significant differences between these planes, making it difficult to blend between adjacent per-view MPIs. To address this issue, we adopt a canonical cylindrical coordinate system across all viewpoints, allowing a neural network to automatically resolve depth conflicts and blend viewpoints. This also allows us to construct a single global MDP instead of storing multiple per-view MPIs [20].

Geometry accuracy The depth component grants the MDP the same 3D representational accuracy as a point cloud. This increased accuracy over MPI’s equidistant planes prevents depth quantization artifacts [28], which are typically more noticeable when the scene geometry is observed from grazing angles, and leads to view synthesis results with fewer visual artifacts around geometry edges.

4 Reconstructing and Rendering MDPs

In this section, we describe our method to a) construct MDPs from images captured with a 360° camera setup, and b) render novel viewpoints from this representation. An overview of the full pipeline is shown in Fig. 3.

Given multiple images and their corresponding camera parameters, we first construct per-view MPI volumes (Sec. 4.1). We then project these MPI volumes to a canonical space—a cylindrical coordinate system centered at the center of the 360° camera—and collapse them into a compact per-view MDP (Sec. 4.2).




Fig. 3. Overview of our pipeline. Our network first warps multi-view images to each view to construct per-view PSVs, and leverages a 3D CNN to predict per-view MPIs. These MPIs are projected to a canonical cylindrical coordinate in the world space. We group these per-view MPIs by different radius ranges and collapse them to reconstruct per-view MDPs. The per-view MDPs are finally blended into a single global MDP, which can be used to render new novel view images using forward splatting.

Finally, we blend the different per-view MDPs into a single MDP representation (Sec. 4.3). Given the reconstructed MDP representation, we can render novel views efficiently using forward splatting (Sec. 4.4).

4.1 Reconstructing per-view MPIs from images

In the following, we describe our pipeline for predicting per-view MPI volumes from a 360° camera setup. Our pipeline takes k input views and their corresponding camera parameters as input. For each viewpoint, we create a plane-sweep volume (PSV) by warping images from the four nearest neighboring cameras as shown in Fig. 3. Following previous work [20], we sample the depths linearly in disparity to ensure that the resulting PSV covers accurate object depths.

These PSVs are processed by a 3D CNN that predicts an MPI volume. Please see the supplementary material for a full description of our 3D CNN. Taken as is, these MPI volumes require a large amount of memory to perform novel view synthesis. For example, Mildenhall et al.[20] use 32–128 depth planes for their experiments. We address this by projecting these MPI volumes to canonical cylindrical coordinates and collapsing them into a compact MDP representation.

4.2 Per-view MPIs to Per-view MDPs

The previous step gives us, for every input view v , an L -layer MPI volume. Our goal is to convert these multiple limited field-of-view volumes into a compact and unified representation. We do this by projecting each MPI volume to a canonical




Fig. 4. We show an example of MPI collapsing at a view. In particular, we show three cylindrical shells that divide the space into two bins, and five MPIs in these bins. We mark the MPI points that belong to different bins with different colors.

cylindrical coordinate system with its origin at the center of projection of the entire camera rig. Concretely, for a pixel (x_s, y_s) on layer l of the MPI for view v , its corresponding 3D point in the global world coordinates is given by:

$$\begin{bmatrix} x_w \\ y_w \\ z_w \\ 1 \end{bmatrix} = E_w E_v^{-1} I_v^{-1} \begin{bmatrix} x_s \\ y_s \\ 1/d_l \\ 1 \end{bmatrix}, \quad (1)$$

where d_l represents the depth of layer l , I_v and E_v are the intrinsic and extrinsic matrices of the camera view v , and E_w is the extrinsic matrix for the geometric center of our camera rig. Once projected to the world coordinates, we can calculate their cylindrical coordinates (ρ, ϕ, z) as:

$$\rho_w = \sqrt{x_w^2 + y_w^2}, \quad \phi_w = \arctan(y_w/x_w). \quad (2)$$

Applying the above operation to an MPI volume produces a point cloud of 3D points each with the color, depth, and opacity. Moreover, since each pixel coordinate in the MPI volume has multiple points at different depths, in the cylindrical coordinate system this leads to a set of multiple points along rays originating from the origin. We collapse this large set of points into a more compact set. More specifically, we divide the 3D space into a small number $M \ll L$ of annulus-like 3D cylindrical regions with equidistant radius ranges that cover the entire scene from the nearest to the farthest radius. We bin the 3D MPI points into these M bins based on the radius ρ_w (see Fig. 4).

For each subset of points within the same bin, we use a back-to-front “over” alpha compositing operation (as is used to render novel views with MPIs [37, 20]) to compute a *single* RGBD α value. Because the over operator is associative, we can process each bin individually. Thus, this operation replaces a large set of RGB α values with a single RGBD α value, thereby significantly reducing the data footprint of the representation. In practice, as demonstrated in Tab. 2, we find that even 2–4 layer MDPs produce results better in visual quality than MPIs with 32 layers, leading to a significant compression.

4.3 Per-view MDP Blending

The previous step produces k per-view MDPs with M layers each. Next we blend them all into a single global MDP. Since all the MDPs are in a canonical cylindrical representation, we can blend the individual corresponding layers. If layer m for the view- v MDP is represented as $(C_m^v, D_m^v, \alpha_m^v)$, the blended global MDP is given by:

$$(C_m, D_m, \alpha_m) = \frac{\sum_v w^v \alpha_m^v (C_m^v, D_m^v, \alpha_m^v)}{\sum_i w^v \alpha_m^v}. \quad (3)$$

This represents a weighted average of the color, depth, and opacity of the per-view MDPs where the weights are a product of the opacity and a per-view weight, w^v that we set to the cosine of the angular difference from the optical axis, which gets lower the further a pixel is away from the principal point.

4.4 Differentiable MDP Rendering with Forward Splatting

In this section, we describe our differentiable rendering module. It is achieved by forward splatting and utilizing soft z-buffering to resolve depth conflicts.

One distinction from previous MPI methods is that we cannot do planar homography warping for each layer to synthesize a novel view since each layer now lies on a cylindrical shell with a depth component. In order to render novel viewpoints, we treat the predicted MDP representation as a set of RGB α point clouds and forward-project each point onto the target image plane. Concretely, we can get the world coordinates (x_w, y_w, z_w) by doing the inverse of the cylindrical coordinate transformation (Eqn. 2) and then splat these points onto the target image plane with a bilinear kernel.

Directly projecting the points to their corresponding target location might result in depth conflicts. Similar to rasterization, when a query ray passes through several surfaces along its direction, the resulting pixel might simultaneously have different color information. To ensure only the closest pixel is selected, we use z-buffering to resolve the conflicts. Since z-buffering is non-differentiable, we instead employ soft z-buffering [24, 33] to compute the weighted average pixel value $\bar{C}(x, y)$ of all conflicting pixels. Soft z-buffering can be formulated as:

$$\bar{C}(x, y) = \frac{\sum C(x, y) e^{(d(x, y) - d_{max})\tau}}{\sum e^{(d(x, y) - d_{max})\tau}}, \quad (4)$$

where $C(x, y)$ and $d(x, y)$ denote the pixel value and inverse depth value at (x, y) of a layer, respectively. τ is a scale factor to control the discriminative power of soft z-buffering. The maximum inverse depth across the image d_{max} is subtracted to prevent overflow. As a pixel gets closer, its inverse depth d increases, thus increasing the weight for this pixel. Finally, by resolving the self-occlusion depth conflicts within each layer via soft z-buffering, we can then alpha-composite the projected maps from back to front to produce the final rendering (see Fig. 3).

5 Implementation details

In the following, we describe the datasets used and the training procedure to learn to project images from a 360° camera setup to the MDP representation.

5.1 Dataset

Throughout our experiments, we assume the number of cameras $k = 16$, each with a 100° field of view, similar to commercial cameras such as the Yi Halo and GoPro Odyssey. Each camera viewing direction is 22.5° apart horizontally. This configuration yields a stereo overlap of over 50% between neighboring cameras.

In order to create photorealistic training data, we chose the Unreal Engine as our renderer since it offers complex effects and high quality rendering. We create two datasets of different camera configurations with UnrealCV [25]: the first with a similar sampling scheme as in [20], and the second with a ring-like camera setup similar to the Yi Halo and GoPro Odyssey. Both datasets are generated from the same 21 large-scale scenes, which consist of indoor and outdoor scenes modeled with high resolution textures and complex scene geometry. These datasets offer a large variety of albedo, depth complexity and non-Lambertian specular effects. The first dataset contains thousands of points of interest, from each of which we sample 6 images with a random baseline to ensure various levels of disparity. Image resolution ranges from 320×240 to 640×480 to allow the network to generalize on various resolutions. In the second dataset, the input images follow the predefined 16-camera configuration, while the output images have a random look-at direction and translational movement with a maximum radius of 25cm. Image resolution for the second dataset is adjusted to range from 320×320 to 512×512 in order to better match the field-of-view of 360° cameras.

5.2 Training

To train our method, we adopt a two-step training scheme. During the first step, we train the 3D CNN on the first dataset to output a per-view MPI volume. The first step is performed by selecting 6 neighboring views, using a random set of 5 as inputs and using the last image as a target view for supervision. During the second step, the 3D CNN is fine-tuned end-to-end with our second dataset by synthesizing a target view from its closest 5 cameras. The network is trained using the perceptual loss from [6] and a learning rate of 2×10^{-4} for roughly 60k and 70k iterations for each respective phase.

6 Results

We now evaluate and compare our MDP representation and panorama-based novel view synthesis method.

Evaluation of the numbers of layers. Our approach works for an arbitrary numbers of MDP layers. In this section, we evaluate how the number of layers

impact the representation accuracy using a synthetic test set that consists of two large indoor scenes different from our training set. Several hundreds viewpoints are randomly sampled within each scene. Table 1 shows the quantitative results of our method when varying the number of layers from one to five. Our method consistently improves with an increasing number of layers. This increase typically allows a more accurate representation of scenes with more complex geometry and appearance. Note that an MDP reverts to a standard RGBD panorama when its number of layers is one (see the first row of Tab. 1). In fact, the depth obtained using an MDP with a single layer are significantly better than a state-of-the-art depth estimation method [17] (see Tab. 3). The number of layers is a user-tunable parameter that provides a trade-off between representation accuracy and memory usage. We use five layers for the remaining experiments in this paper. Note that, even with five layers, our MDPs are still compact and more memory-efficient than previous learning-based methods (see Tab. 2).

Layers	PSNR \uparrow	SSIM \uparrow	L1 \downarrow
1	25.75	0.8565	0.0269
2	26.17	0.8628	0.0254
3	26.27	0.8655	0.0251
4	26.35	0.8661	0.0251
5	26.39	0.8664	0.0251

Table 1. Quantitative evaluation of the number of MDP layers. Please see the supplementary material for additional qualitative results.

Comparison with MPIs. Our method effectively converts the costly per-view MPIs to the novel compact view-independent MDPs. MPIs consist of dense planes and only support local view extrapolation, which requires rendering multiple dense sets of planes from multiple views to enable arbitrary rotational motion. In contrast, our MDPs are in a canonical world space, which only requires splatting a single sparse set of depth layers for 360° 6-DoF view synthesis. In Tab. 2, we show the quantitative results of our method with 2 and 5 layer MDPs. We also compare against a naive solution that directly uses the per-view MPIs to do view synthesis with 16-view MPIs and 32 planes per MPI. For this, following the method of Mildenhall et al. [20], we linearly blend the five neighboring per-view MPI renderings with the pixel-wise cosines of the angles between the per-pixel viewing directions and the central direction of each camera. The corresponding memory usage of these MPIs and MDPs are shown in the table. The naive MPI method is not able to effectively blend the multi-view renderings; training it end-to-end to learn the 2D blending process may improve the quality. In contrast, our approach can leverage its priors learned during training to merge the per-view MPIs, which leads to a method that outperforms the MPI approach. Moreover, our method is significantly more memory-efficient, taking an order of magnitude less memory than the MPI method.

Comparisons with other methods on synthetic scenes. We now compare our method with other 360° view synthesis techniques that allow for translational

Methods	PSNR \uparrow	SSIM \uparrow	L1 \downarrow	Dimension	Storage
MPI	24.75	0.8278	0.0306	$16 \times 640 \times 640 \times 32 \times 4$	3.355GB
Ours - 2 Layers	26.17	0.8628	0.0254	$2560 \times 640 \times 2 \times 5$	0.066GB
Ours - 5 Layers	26.39	0.8664	0.0251	$2560 \times 640 \times 5 \times 5$	0.165GB

Table 2. Quantitative results of our method compared to a MPI-based method on synthetic scenes, along with their memory usages. Please see the supplementary material for more qualitative results.

Methods	PSNR \uparrow	SSIM \uparrow	L1 \downarrow
Depth [17] as points	23.06	0.766	0.047
Depth [17] as mesh	23.75	0.780	0.040
Depth [17] + [27]	23.20	0.767	0.043
Our depth + [27]	24.98	0.827	0.031
Our single layer MDP	25.75	0.857	0.027
Our MDPs (5 layers)	26.39	0.866	0.025

Table 3. Comparison on synthetic scenes. We compare with other methods that use a single RGBD panorama for 6DoF rendering. We show results of SSIMs, PSNRs and L1 loss of these methods.

motion and demonstrate quantitative comparisons on our synthetic testing set. The most popular way to do 360° rendering is to reconstruct the depth of a panorama and render the RGBD panorama as a point cloud or mesh, as introduced in [23, 1]. These techniques rely on a pre-computed depth map, and we use a state-of-the-art deep learning based depth estimation method [17] to generate depth for input panoramas. This depth map is used to generate a corresponding RGBD point cloud and mesh that are used in turn to render novel views. We also compare against a state-of-the-art method [27] that is designed to improve the 2D images rendered from the RGBD panorama mesh by resolving the disocclusions around the geometric boundaries.

Disparity (PSNR/SSIM)	32	64	128
Depth [17] + [27]	22.40 / 0.7490	20.20 / 0.6911	18.17 / 0.6322
Our MDPs (5 layers)	26.90 / 0.8839	24.96 / 0.8436	21.94 / 0.7579

Table 4. Quantitative comparison on maximum disparity. Our method consistently outperforms [27] on image quality over all tested disparity levels.

Table 3 shows the quantitative results of these methods. Our MDPs generate significantly better results than the other comparison methods as reflected by the highest PSNRs and SSIMs and the lowest L1 loss. We also show our single layer MDP result in Table 3, which is essentially doing single RGBD panorama-based point cloud rendering. Interestingly, even this performs better than the point cloud rendering using the prior state-of-the-art depth estimation method. This demonstrates that our method can be used as an effective panorama depth estimation technique by estimating a single MDP. We also show that, by giving the depth from our single MDP as input to Serrano et al.’s enhancing tech-

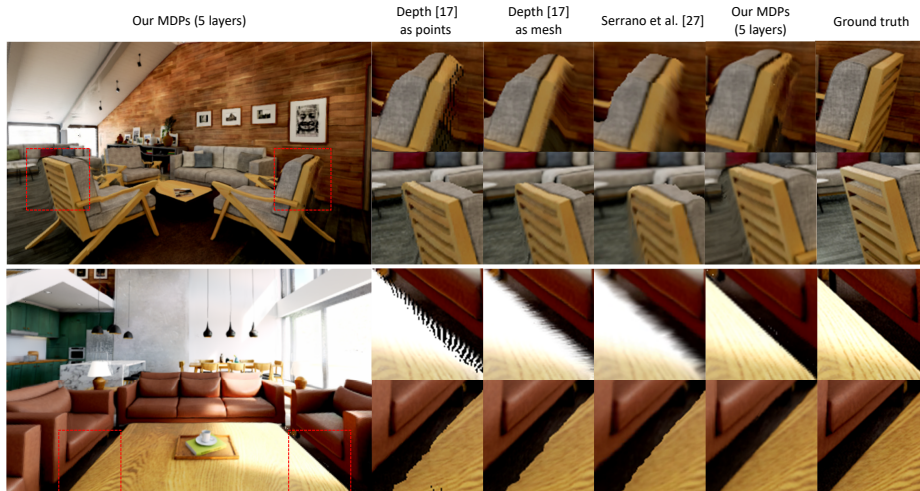


Fig. 5. Visualization of synthetic results. We show two examples of synthetic results to illustrate the visual quality corresponding to the numbers in Tab. 3.

nique [27], it improves their result over using the depth from Lee et al. [17]. Note that, Serrano et al. [27] focus on improving user experience and mitigate unpleasant artifacts, which may decrease the accuracy of direct mesh rendering.

To visualize the comparisons, we show two results on synthetic data in Fig. 5. Our results offers increased visual quality than all other comparison methods, especially around object boundaries. Our results are also the most similar to the ground truth images, which is consistent with the quantitative results in Tab. 3. Note that, a single RGBD representation used in other methods is very limited and cannot well reproduce the challenging disocclusion effects. This leads to holes in the point cloud rendering, and noisy and stretched boundaries in the mesh rendering; Serrano et al. [27] smooth out the boundaries in the mesh rendering, which in turn introduces blurriness and distortion. Our MDPs can describe complex scenes with multiple objects on the line-of-sight and effectively handle challenging boundaries, disocclusions and other appearance effects.

We now analyze the robustness of our method on large disparities. To do so, we translate the target viewpoints in our test set and record the amount of maximum scene disparity with respect to the reference viewpoint. We use a 32-layers MPI to perform the sampling. We compare our method to [27] and report the results in Tab. 4. In this table, we focus on larger disparity levels as these correspond to larger translations. As expected, larger translations yields a more difficult task, which reduces the novel view visual quality. Despite this, our method outperforms [27] across all disparity levels. These observations also match the sampling guidelines in Mildenhall et al. [20].

We also show an example of specular reflection in Fig. 6. Note that the other methods use a single depth, which bakes the reflections in the object color and fails to accurately capture the motion of reflections from different viewpoints.




Fig. 6. Reflection effects. We show one example of specular reflection with zoomed-in insets. Note that, our method recovers the specular reflection motion that matches the ground truth very well, while the reflections in the comparison methods do not move.

In contrast, our novel MDP allows the reflections to be encoded at a different depth, which effectively models the moving reflection.

Comparisons on real scenes. We now evaluate our method on complex real scenes and compare them with the methods using RGBD panoramas. We captured these real scenes with a Yi Halo and used Google Jump Manager [1] to stitch the multi-view images and generate the depth. Figure 7 shows the results of our method with 5-layers MDPs, rendering an RGBD panorama as a point cloud and mesh, and using [27] to improve the mesh rendering. Similar to the synthetic comparisons, our method produces more realistic results than the comparison ones. Note that, the disocclusion effects introduce obvious holes and significant discontinuous noise in the baseline point cloud rendering and mesh rendering techniques. While [27] resolves the noisy boundaries in the mesh rendering for better user experience, the edges are in fact distorted in many regions and still look physically implausible. Our results are significantly more photorealistic. Please refer to supplementary material for more experiments and ablation study on end-to-end training.

Limitations Our proposed MDP representation can be used to represent complex geometry and appearance, given enough layers. While five layers are sufficient for most cases as previously shown, it might not be adapted to more complex and challenging scenes and can exhibit blurriness or distortion. Increasing the number of layers can potentially address this. Besides, our novel view synthesis method is limited to relatively small motions for two reasons. First, large translations will potentially expose large unseen parts of the scene, not encoded in the MDP representation. We hypothesize that future work on generative models might provide a solution to this issue. Second, moving past the first concentric sphere of the MDP would break the alpha ordering.




Fig. 7. Qualitative results on complex real scenes. We show results on novel viewpoints comparing our 5-layers MDP (leftmost) to RGBD based methods.

7 Conclusion

We presented a novel 3D scene representation—Multi Depth Panoramas, or MDP—that represents complex scenes using multiple layers of concentric RGBD_α panoramas. MDPs are more compact than prior scene representations such as MPIs. As such, they can be used to generate high-quality novel view synthesis results with translational motion parallax, using our proposed forward-splatting rendering method. Furthermore, we presented a learning-based method to accurately reconstruct MDPs from commercial 360° camera rigs.

Acknowledgements We would like to thank In-Kyu Park for helpful discussion and comments. This work was supported in part by ONR grants N000141712687, N000141912293, N000142012529, NSF grant 1617234, Adobe, the Ronald L. Graham Chair and the UC San Diego Center for Visual Computing.

References

1. Anderson, R., Gallup, D., Barron, J.T., Kontkanen, J., Snavely, N., Hernández, C., Agarwal, S., Seitz, S.M.: Jump: virtual reality video. *ACM Transactions on Graphics (TOG)* **35**(6), 1–13 (2016)
2. Bertel, T., Campbell, N.D.F., Richardt, C.: Megaparallax: Casual 360° panoramas with motion parallax. *IEEE Transactions on Visualization and Computer Graphics* **25**(5), 1828–1835 (May 2019). <https://doi.org/10.1109/TVCG.2019.2898799>
3. Broxton, M., Flynn, J., Overbeck, R., Erickson, D., Hedman, P., DuVall, M., Dourgarian, J., Busch, J., Whalen, M., Debevec, P.: Immersive light field video with a layered mesh representation **39**(4), 86:1–86:15 (2020)
4. Brunet, T., Merlin, A., Mascaro, B., Zimny, K., Leng, J., Poncelet, O., Aristégui, C., Mondain-Monval, O.: Soft 3D acoustic metamaterial with negative index. *Nature materials* **14**(4), 384–388 (2015)
5. Buehler, C., Bosse, M., McMillan, L., Gortler, S., Cohen, M.: Unstructured lumigraph rendering. In: *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*. pp. 425–432. ACM (2001)
6. Chen, Q., Koltun, V.: Photographic image synthesis with cascaded refinement networks. In: *Proceedings of the IEEE international conference on computer vision*. pp. 1511–1520 (2017)
7. Chen, S.E., Williams, L.: View interpolation for image synthesis. In: *Proceedings of SIGGRAPH*. pp. 279–288 (1993)
8. Cheng, S., Xu, Z., Zhu, S., Li, Z., Li, L.E., Ramamoorthi, R., Su, H.: Deep stereo using adaptive thin volume representation with uncertainty awareness. *arXiv preprint arXiv:1911.12012* (2019)
9. Choi, I., Gallo, O., Troccoli, A., Kim, M.H., Kautz, J.: Extreme view synthesis. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 7781–7790 (2019)
10. Debevec, P.E., Taylor, C.J., Malik, J.: Modeling and rendering architecture from photographs: A hybrid geometry-and image-based approach. In: *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*. pp. 11–20. ACM (1996)
11. Flynn, J., Neulander, I., Philbin, J., Snavely, N.: Deepstereo: Learning to predict new views from the world’s imagery. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 5515–5524 (2016)
12. Gortler, S.J., Grzeszczuk, R., Szeliski, R., Cohen, M.F.: The lumigraph. In: *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*. pp. 43–54. ACM (1996)
13. Hedman, P., Alisan, S., Szeliski, R., Kopf, J.: Casual 3D Photography **36**(6), 234:1–234:15 (2017)
14. Huang, J., Chen, Z., Ceylan, D., Jin, H.: 6-DOF VR videos with a single 360-camera. In: *2017 IEEE Virtual Reality (VR)*. pp. 37–44 (March 2017). <https://doi.org/10.1109/VR.2017.7892229>
15. Ishiguro, H., Yamamoto, M., Tsuji, S.: Omni-directional stereo. *IEEE Transactions on Pattern Analysis & Machine Intelligence* **2**(2), 257–262 (1992)
16. Kalantari, N.K., Wang, T.C., Ramamoorthi, R.: Learning-based view synthesis for light field cameras. *ACM Transactions on Graphics (TOG)* **35**(6), 193 (2016)
17. Lee, J.H., Han, M.K., Ko, D.W., Suh, I.H.: From big to small: Multi-scale local planar guidance for monocular depth estimation. *arXiv preprint arXiv:1907.10326* (2019)

18. Levoy, M., Hanrahan, P.: Light field rendering. In: Proceedings of the 23rd annual conference on Computer graphics and interactive techniques. pp. 31–42. ACM (1996)
19. Luo, B., Xu, F., Richardt, C., Yong, J.H.: Parallax360: Stereoscopic 360° scene representation for head-motion parallax. *IEEE Transactions on Visualization and Computer Graphics* **24**(4), 1545–1553 (Apr 2018). <https://doi.org/10.1109/TVCG.2018.2794071>, <https://doi.org/10.1109/TVCG.2018.2794071>
20. Mildenhall, B., Srinivasan, P.P., Ortiz-Cayon, R., Kalantari, N.K., Ramamoorthi, R., Ng, R., Kar, A.: Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *ACM Transactions on Graphics (TOG)* **38**(4), 1–14 (2019)
21. Peleg, S., Ben-Ezra, M., Pritch, Y.: Omnistere: Panoramic stereo imaging. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **23**(3), 279–290 (2001)
22. Porter, T., Duff, T.: Compositing digital images. In: Proceedings of the 11th annual conference on Computer graphics and interactive techniques. pp. 253–259 (1984)
23. Pozo, A.P., Toksvig, M., Schrager, T.F., Hsu, J., Mathur, U., Sorkine-Hornung, A., Szeliski, R., Cabral, B.: An integrated 6DoF video camera and system design. *ACM Transactions on Graphics (TOG)* **38**(6), 1–16 (2019)
24. Pulli, K., Hoppe, H., Cohen, M., Shapiro, L., Duchamp, T., Stuetzle, W.: View-based rendering: Visualizing real objects from scanned range and color data. In: Dorsey, J., Slusallek, P. (eds.) *Rendering Techniques '97*. pp. 23–34. Springer Vienna, Vienna (1997)
25. Qiu, W., Zhong, F., Zhang, Y., Qiao, S., Xiao, Z., Kim, T.S., Wang, Y., Alan, Y.: UnrealCV: Virtual worlds for computer vision. *ACM Multimedia Open Source Software Competition* (2017)
26. Richter, S.R., Roth, S.: Matryoshka networks: Predicting 3D geometry via nested shape layers. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1936–1944 (2018)
27. Serrano, A., Kim, I., Chen, Z., DiVerdi, S., Gutierrez, D., Hertzmann, A., Masia, B.: Motion parallax for 360 RGBD video. *IEEE transactions on visualization and computer graphics* **25**(5), 1817–1827 (2019)
28. Shade, J., Gortler, S., He, L.w., Szeliski, R.: Layered depth images. In: Proceedings of the 25th annual conference on Computer graphics and interactive techniques. pp. 231–242. Association for Computing Machinery, Inc. (July 1998), <https://www.microsoft.com/en-us/research/publication/layered-depth-images/>
29. Sinha, S., Steedly, D., Szeliski, R.: Piecewise planar stereo for image-based rendering. In: 12th International Conference on Computer Vision (ICCV). pp. 1881–1888. IEEE (2009)
30. Srinivasan, P.P., Tucker, R., Barron, J.T., Ramamoorthi, R., Ng, R., Snavely, N.: Pushing the boundaries of view extrapolation with multiplane images. *Conference on Computer Vision and Pattern Recognition (CVPR)* (2019)
31. Szeliski, R., Golland, P.: Stereo matching with transparency and matting. *International Journal of Computer Vision* **32**(1), 45–61 (May 1999), <https://www.microsoft.com/en-us/research/publication/stereo-matching-with-transparency-and-matting/>
32. Thatte, J., Boin, J.B., Lakshman, H., Girod, B.: Depth augmented stereo panorama for cinematic virtual reality with head-motion parallax. In: 2016 IEEE International Conference on Multimedia and Expo (ICME). pp. 1–6. IEEE (2016)
33. Tulsiani, S., Tucker, R., Snavely, N.: Layer-structured 3D scene inference via view synthesis. In: *ECCV* (2018)

34. Xu, Z., Bi, S., Sunkavalli, K., Hadap, S., Su, H., Ramamoorthi, R.: Deep view synthesis from sparse photometric images. *ACM Transactions on Graphics* **38**(4), 76 (2019)
35. Yao, Y., Luo, Z., Li, S., Fang, T., Quan, L.: Mvsnet: Depth inference for unstructured multi-view stereo. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 767–783 (2018)
36. Zheng, K.C., Kang, S.B., Cohen, M.F., Szeliski, R.: Layered depth panoramas. In: *2007 IEEE Conference on Computer Vision and Pattern Recognition*. pp. 1–8. IEEE (2007)
37. Zhou, T., Tucker, R., Flynn, J., Fyffe, G., Snavely, N.: Stereo magnification: Learning view synthesis using multiplane images. In: *SIGGRAPH* (2018)
38. Zitnick, C.L., Kang, S.B., Uyttendaele, M., Winder, S., Szeliski, R.: High-quality video view interpolation using a layered representation. *ACM transactions on graphics (TOG)* **23**(3), 600–608 (2004)