Casual Indoor HDR Radiance Capture from Omnidirectional Images

Pulkit Gera¹ pulkit.gera@research.iiit.ac.in Mohammad Reza Karimi Dastjerdi² mohammad.karimi-dastjerdi.1@ulaval.ca Charles Renaud²

charles.renaud.1@ulaval.ca

P. J. Narayanan¹ pjn@iiit.ac.in Jean-François Lalonde² jflalonde@gel.ulaval.ca

- ¹ Centre for Visual Information Technology & Kohli Centre on Intelligent Systems (KCIS) International Institute of Information Technology (IIIT), Hyderabad, India
- ² Université Laval Quebec City, Canada

Abstract

We present PanoHDR-NeRF, a neural representation of the full HDR radiance field of an indoor scene, and a pipeline to capture it casually, without elaborate setups or complex capture protocols. First, a user captures a low dynamic range (LDR) omnidirectional video of the scene by freely waving an off-the-shelf camera around the scene. Then, an LDR2HDR network uplifts the captured LDR frames to HDR, which are used to train a tailored NeRF++ model. The resulting PanoHDR-NeRF can render full HDR images from any location of the scene. Through experiments on a novel test dataset of real scenes with the ground truth HDR radiance captured at locations not seen during training, we show that PanoHDR-NeRF predicts plausible HDR radiance from any scene point. We also show that the predicted radiance can synthesize correct lighting effects, enabling the augmentation of indoor scenes with synthetic objects that are lit correctly. Datasets and code are available at https://lvsn.github.io/PanoHDR-NeRF/.

1 Introduction

Capturing the incoming radiance of a given scene is an important step for many augmented reality applications, since it allows immersive exploration and realistic scene augmentation. To accurately measure the radiance, Debevec [\Box] pioneered *image-based lighting* which involves photographing a chrome sphere—called a *light probe*—at different exposures and merging them into a single, high dynamic range (HDR) image [\Box]. HDR light probing was later extended to use wide angle lenses or 360° cameras and is at the heart of the lighting capture necessary to achieve special effects in movies today¹. Because of the close proximity between light sources and objects in the scene, the HDR radiance field of a typical indoor scene varies rapidly: lighting near a window is vastly different from the center of the room. Accurately

^{© 2022.} The copyright of this document resides with its authors.

It may be distributed unchanged freely in print or electronic forms.

¹https://www.fxguide.com/fxfeatured/the-definitive-weta-digital-guide-to-ibl/



a) Casual capture

b) Novel view synthesis

c) HDR radiance estimation

Figure 1: We capture the continuous HDR radiance of an indoor scene. Our PanoHDR-NeRF approach takes a) casually captured LDR images from an off-the-shelf camera (shown in inset) as input, and performs b) novel view synthesis of the indoor scene. c) As opposed to existing techniques such as NeRF++ [\square] (left), PanoHDR-NeRF (right) properly estimates the HDR radiance of the scene, visualized by relighting virtual test objects.

capturing indoor radiance involves moving the apparatus and repeating the operation many times, limiting scalability. Approximations such as reprojecting the measured radiance onto a proxy 3D model [1] must be used. To simplify the capture process, one could use inverse tonemapping techniques [11], [19], [11] which estimate the true dynamic range information from low dynamic range (LDR) inputs. This yields radiance estimates only where images are taken. Can we capture and learn a representation of the scene from which HDR radiance can be synthesized at any point?

Novel view synthesis of an object or scene from multiple views has garnered attention recently, especially since the emergence of neural radiance fields (NeRF) [1]. In this approach, a set of images are captured, co-registered using structure-from-motion (SfM), and used to train a deep network that learns to predict the color and opacity along any 3D ray using volumetric rendering. Recent flavors of NeRF model large scenes [3, 53], handle omnidirectional input images $[\square]$, and decompose the scene into intrinsic components $[\square, \square]$. Most NeRF-based methods accept LDR images as input and do not accurately model the true HDR radiance of indoor environments. There are recent attempts to address this limitation [22],], which require multiple LDR images at different exposures to model the full dynamic range of indoor lighting. Capturing 11 exposures at a point of the scene, which is necessary to reconstruct a reasonable HDR panorama spanning over 22 f-stops, takes approximately two minutes with a conventional 360° camera (Ricoh Theta Z1). Thus to capture HDR lighting for a large indoor scene, we need over 200 images bringing the capture time to over 3.5 hours. In addition, an elaborate setup would be required to avoid ghosting artifacts. These methods to capture the full dynamic range of lighting for large indoor scenes require specialized setup and are tedious and cumbersome.

In this paper, we present PanoHDR-NeRF, a neural representation of the plausible full HDR radiance field of an indoor scene. We also present a method to generate it from casually captured images of the scene. HDR radiance from any novel viewpoint in the scene can subsequently be estimated (fig. 1) from PanoHDR-NeRF. Our method does not require any special equipment or complicated capture protocols. It accepts as input a video sequence captured by freely moving a commercial 360° camera around the scene. As output, it can produce the HDR radiance at any given location in the scene. To do so, we leverage two deep neural networks: 1) an LDR2HDR model that predicts the HDR radiance from a single LDR panorama captured by the camera and 2) a modified NeRF++ model trained on the predicted

HDR outputs of the first network. To evaluate our proposed method, we capture a set of six different indoor scenes, which we augment with a set of ground truth HDR light probes at each scene. Our experiments demonstrate that, despite the simplicity of the capture procedure, PanoHDR-NeRF can accurately predict HDR radiance in a variety of challenging conditions. Our approach can render 360° spatially varying HDR light probes, which can be used to provide correct lighting effects when the scene is augmented with virtual objects. Compared to previous methods, our approach reduces the complexity and time required to capture the full HDR radiance field of the scene.

2 Related work

Inverse tonemapping Inverse tonemapping aims to recover missing information in the over- and under-saturated areas of an LDR image. While earlier methods [53] relied on heuristics, several deep learning architectures have been proposed recently [11, 52]. These include a 2D encoder and 3D decoder with skip connections to generate bracketed LDR image stacks (over exposed and under exposed) [11], an encoder-decoder to reconstruct the HDR image directly [23], and a multiscale autoencoder architecture to learn multilevel features from LDR image which are merged to reconstruct HDR images [51]. Lee *et al.* [24] generate LDR stacks with a two-branch network. DeepHDR [12] mask out the saturated areas to reduce sub-optimal features from well-exposed and saturated pixels. HallucinationNet [22] individually models the main components of the imaging pipeline: dynamic range clipping, camera response function, and quantization. LANet [52] introduces a multi-task network with a luminance attention and HDR reconstruction streams. We build on the LANet architecture and augment it with an additional rendering loss. This combination outperforms alternatives at predicting the high dynamic range radiance in indoor environments.

Classical methods reconstruct an explicit 3D model of the scene [9], Novel view synthesis □, □]. Recent works utilize off-the-shelf SfM techniques to generate a coarse geometry and use neural networks to render photorealistic novel views [19, 52, 53]. Breaking the scene into Multi-Plane Image (MPI) representation to render novel views by blending has been tried [1, 3, 5]. DeepVoxels [1] learn a voxel-based volumetric representation of the scene using Gated Recurrent Units (GRUs) [6]. Free View Synthesis and Stable View Synthesis [III, III] map the encoded features from the source images into the target view and blend them via a neural network directly or in a geometric space. Novel view synthesis from omnidirectional images are fairly limited so far. Huang et al. [2] reconstruct a point cloud from 360° videos to achieve real time video playback on a VR device. Serrano et al. [13] design a layered scene representation that facilitates parallax and real time playback of 360° video. [III, ZII] present Multi-Depth Panorama (MDP) and multi-sphere images to create 6-DoF renderings. This requires an elaborate setup and does not accommodate free viewpoint synthesis. Bertel et al. [1] propose a fast, casual and robust capture of immersive real-world VR experience. However it takes a lot of memory and poor proxy geometry causes warping artifacts. Xu *et al.* [**D**] estimate the entire indoor scene from a single image using a CNN but need room layout priors and depth that are challenging to obtain for real scenes.

Neural radiance fields Neural Radiance Fields (NeRF) [52] learn an implicit volumetric scene representation with a Multi-Layer Perceptron (MLP) that receives viewing positions and directions as input and predicts the RGB colour and opacity as output. The original formulation was later extended in several ways relevant to our work. NeRF++ [52], which we leverage, models an unbounded scene by splitting it into foreground/background, learning



Figure 2: Overview of our pipeline. At training time (left), the captured panoramas are linearized (calibrated using a color checker, not shown) and processed by a pre-trained LDR2HDR network to obtain HDR estimates. The HDR panoramas, along with the camera poses obtained with OpenSfM [23], are used to train the PanoHDR-NeRF network, which learns to synthesize HDR scene radiance at any point in the scene. At inference time (right), we simply provide the novel camera pose and obtain the corresponding novel HDR panorama.

each separately. Mega-NeRF [[1]] divides the scene into smaller sections. Block-NeRF [[1]] splits city-scale scenes into blocks, trains a NeRF for each block, and combines the results. Mip-NeRF [[1], [1]] replaces rays with anti-aliased conical frustums for speed and accuracy. NeRF with HDR images has been explored in two works. NeRF in the dark [[1]] train directly on linear RAW images with higher effective dynamic range (14-16 bits compared to more typical 8-bit cameras). Huang *et al.* [[2]] train from a set of LDR images captured at alternating exposures. Both methods require multiple LDR images at different exposures, which is time-consuming to capture and requires specialized setup. In contrast, our method plausibly predicts the full dynamic range from a single casually captured video. Finally, panoramic images have been explored in [[1], [2]], who present a method to synthesize panoramas from a single omnidirectional input. They require depth as input and novel views can only be rendered on a straight line path. CylindricalNeRF [[2]] proposes cylindrical sampling for unbounded scenes captured in circular trajectory. OmniNeRF [[2]] synthesizes novel fish-eye projection images, using spherical sampling to improve the quality of results as we do.

HDR scene capture Zhang *et al.* [5] leverage an RGBD scan of an indoor scene and estimate scene parameters including HDR radiance over scene geometry. Walton *et al.* [5] combine a depth sensor with a fisheye camera in a SLAM-based approach to recover geometry and lighting. Tarko *et al.* [5] takes an omnidirectional video as input and uses inverse tone mapping [5] to convert it to HDR. Yang *et al.* [5] learn background and objects as separate NeRF models and combine them with a captured panoramic image into a single scene. Our method synthesizes panoramas from novel viewpoints with full dynamic range.

3 Method

Given a set of LDR panoramas $\{I_k\}_{k=1}^N$ captured freely using an off-the-shelf 360° camera and associated camera poses $\{\mathbf{P}_k\}_{k=1}^N$ obtained with SfM, our objective is to predict the HDR radiance of the scene at any novel viewpoint. We achieve this by recovering HDR values from LDR frames using learning-based inverse tone mapping and then using them as the supervision for novel view synthesis. An overview of our method is given in fig. 2.



Figure 3: Representative images from each test scene used in the experiments.

3.1 High dynamic range with LDR2HDR network

Extrapolating HDR from LDR inputs is typically framed as recovering values in the over- and under-exposed regions. Our work focuses on recovering the over-exposed regions exclusively, with the goal of predicting accurate HDR radiance values (specially the intensities of light sources) for realistic virtual object insertion.

Network architecture We borrow the network architecture proposed in Luminance Attentive Networks (LANet) [53], which is designed as a multi-task network with two *luminance attention* and *HDR reconstruction* streams. The former attempts to create a spatially-weighted attention map of the over-exposed regions in the input image, while the latter uses the attention map to estimate the HDR images. Note that we do not use their proposed adaptation for panoramas since it did not improve the performance in our case. So, we train the model on equirectangular images directly.

Loss functions For the LDR2HDR module, we use the same loss function ℓ_{lanet} as [53], combining scale invariant and luminance segmentation losses. The HDR panoramas are radiance values used to light the scene. To match rendering quality, we use a rendering loss

$$\ell_{\text{rend}} = ||\mathbf{T} \mathbf{y}_{\text{HDR}} - \mathbf{T} \mathbf{t}_{\text{HDR}}||_2, \qquad (1)$$

where **T** is a pre-computed transport [13] matrix for a Lambertian scene (single bounce), \mathbf{y}_{HDR} is the predicted HDR panorama, and \mathbf{t}_{HDR} is the ground truth. A top-down view of a "spiky sphere" on a plane is used for **T**. The rendered images guide the network about the color and direction of HDR lighting in the scene. The final loss for training the LDR2HDR network is an equally weighted combination $\ell_{hdr} = \ell_{lanet} + \ell_{rend}$.

Datasets We pretrain the LDR2HDR network on the Laval Indoor HDR Dataset [\square], which consists of 2,400 HDR panoramas captured in a variety of indoor settings, with a train, validation and test split as 80:10:10. Since the sensor used to capture it (Canon 5D Mark iii camera) and the sensor we use to capture indoor scenes (Ricoh Theta Z1) are different, a domain gap was observed. To alleviate this, we finetune the LDR2HDR network on a small dataset of 78 HDR panoramas captured at different locations, different from the test scenes, using the target sensor (Theta Z1).

3.2 Continuous HDR radiance with PanoHDR-NeRF

Network architecture We take inspiration from and combine several recent works on NeRF. First, we employ the NeRF++ architecture as the base. Second, similar to [**1**], we incorporate the anti-aliased conical frustums from Mip-NeRF [**1**]. We train PanoHDR-NeRF by sampling rays in spherical coordinates instead of pixel coordinates (more details in supplementary). **Loss functions** We train PanoHDR-NeRF using supervision from LDR2HDR network. Traditional volume rendering methods are used to predict the radiance and densities of points

	Input LDR			LDR2HDR pre-trained			LDR2HDR finetuned		
Dataset	PU-PSNR↑	$\text{RMSE}{\downarrow}$	HDR-VDP3↑	PU-PSNR↑	$\text{RMSE}{\downarrow}$	HDR-VDP3↑	PU-PSNR↑	$\text{RMSE}{\downarrow}$	HDR-VDP3↑
Chess room	31.659	0.051	8.067	33.994	0.048	8.234	36.995	0.005	8.492
Stairway	31.964	0.224	7.881	33.297	0.213	8.016	33.685	0.019	8.489
Cafeteria	25.299	5.378	6.098	26.664	5.268	6.418	28.499	4.061	7.164
Spotlights	23.939	3.489	6.001	25.118	3.438	6.097	28.966	0.877	7.667
Dark class	30.657	0.364	7.429	32.125	0.340	7.592	32.594	0.262	8.135
Small class	32.353	2.162	8.018	34.095	1.889	8.267	35.465	0.221	8.678
Overall	28.399	1.703	7.243	30.913	1.634	7.406	33.651	0.696	8.209

Table 1: Quantitative comparison of different strategies for recovering radiance across captured scenes. "Input LDR" are on the images captured by the camera, "LDR2HDR pre-trained" is on our network pre-trained on the Laval Indoor Dataset, and "LDR2HDR finetuned" is after the network finetuned to test camera. As expected, finetuning helps bridge the domain gap and significantly improves the results. Results are averaged over different viewpoints.



Figure 4: Qualitative comparison of different strategies for recovering radiance across captured scenes. For each example, the figure shows virtual test objects relit to demonstrate the dynamic range. Note that despite some color imbalance (e.g. "Spotlights"), fine-tuning helps bridge the domain gap between the training data and the captured images. Images tonemapped for display with $\gamma = 2.2$.

sampled on the ray as in [\square , \square]. We define NeRF loss ℓ_{nerf} between predicted radiance \hat{e} and ground truth HDR e as

$$\ell_{\text{nerf}} = \sum_{r \in \mathcal{R}(\mathbf{P})} \|\hat{E}(\mathbf{r}) - E(\mathbf{r})\|^2,$$
(2)

where E = log(1 + e) and $\mathcal{R}(\mathbf{P})$ is the set of camera rays at pose **P**. The photographer, who inevitably is in the images, is segmented using an off-the-shelf segmentation algorithm [II], and the corresponding pixels are ignored in the loss.

Datasets To obtain training data for a given indoor scene, we casually capture a set of LDR panoramas using a commercial 360° camera (Ricoh Theta Z1). We attach the camera to a portable tripod and capture a 360° video while waving the camera around the scene to cover the entire volume as much as possible for a few minutes (typically 3–5 minutes for the scenes used in the experiments). Approximately 200 frames are then extracted from the video at even intervals. The camera parameters of the input LDR panoramas are recovered using OpenSFM[2] and given as input to PanoHDR-NeRF.

4 Evaluation

In this section, we evaluate our approach against a set of challenging real scenes, where the ground truth HDR radiance is also captured at several locations. We further establish the sensitivity of design choices and compare them to closely related techniques.

4.1 Radiance capture evaluation dataset

For evaluation, we capture different real indoor scenes from a variety of different environments (fig. 3). For each scene, we first capture a 360° video using the Ricoh Theta Z1 panoramic camera as described in sec. 3.2. We also capture a set of HDR panoramas for evaluation to use as ground truth. To obtain these, we set the camera on a tripod at certain locations in the scene (between 3–10 locations per scene) and program it using exposure bracketing to capture 11 exposures spanning over 22 f-stops that are subsequently merged to HDR using the PTGui Pro commercial software. We also capture a short video with the same camera parameters as the captured video at the same location as the HDR. We then extract a single frame from that video, allowing us to have an LDR image at the exact same location as its HDR counterpart. The resulting LDR image is linearized using a pre-calibrated camera response function obtained with a Macbeth color checker. In total, we capture six different scenes containing a total number of 10 LDR/HDR ground truth panorama pairs.

4.2 LDR2HDR evaluation

The LDR2HDR network presented in sec. 3.1 is evaluated on the test set described in sec. 4.1. Tab. 1 compares the performance obtained by: using the LDR images directly, the LDR2HDR network pre-trained on the Laval Indoor HDR Dataset [13], and after fine-tuning on the 78 HDR dataset captured with the same test camera. For evaluation, we use the "PU-PSNR" [23], which is a perceptually-uniform PSNR adjusted for HDR images. In addition, the "RMSE" corresponds to the rendering loss in eq. (1). Finally, we also report the HDR-VDP3 [52], where a value of 10 indicates a perfect match. Here, color encoding is set as "rgb-bt.709" for HDR evaluation, assuming a 24-inch display, 1920×1080 resolution, and a viewing distance of 1 meter.

As shown in tab. 1 and illustrated qualitatively in fig. 4, there exists a significant domain gap between the training dataset and the test camera: simply pre-training on [13] works marginally better than the input LDR image itself, but finetuning results in a significant performance gain on all metrics. Visually, finetuning produces renderings that look very similar to the ground truth (fig. 4).

4.3 PanoHDR-NeRF evaluation

To evaluate how well PanoHDR-NeRF works in terms of capturing the high dynamic range radiance field, we use the same set of ground truth HDR images as described in sec. 4.1. We infer environment maps at the locations of HDR panoramas and use them to render a synthetic scene. We modified NeRF++ to work with equirectangular image representation and trained directly on the LDR frames of the input video. The rendering results of NeRF++ appear dark compared to ground truth, and the lighting is not realistic (fig. 5). In addition, the generated shadows are soft and faded. In contrast, PanoHDR-NeRF produces well-lit results, with sharp shadows that are similar to ground truth.

	w/o re	nder loss	render loss		
	PU-PSNR↑	HDR-VDP3↑	PU-PSNR↑ HDR-VDP3↑		
Hall. Net [🗖]	30.06	7.54	31.51	8.12	
LANet [🔂]	32.43	8.26	38.20	9.67	

Table 2: Quantitative comparison of two single image HDR estimation architectures on the Laval Indoor HDR test set, with and withough the rendering loss ℓ_{render} while training the network. Render Loss with LANet significantly improves the results.

Dataset	Linear PU-PSNR↑	loss RMSE	PanoHDR PU-PSNR↑	-NeRF RMSE	NeRF-LDR2HDR PU-PSNR↑ RMSE		
Dunoer	1010101	THINDLY	1010101	Π	10 I blint	iuno±γ	
Chess room	35.152	0.011	36.941	0.012	35.991	0.006	
Stairway	31.810	0.055	33.169	0.056	32.707	0.055	
Cafeteria	24.139	4.376	28.029	4.179	26.537	5.298	
Spotlights	26.719	0.909	28.657	0.431	27.324	1.619	
Dark class	28.431	1.367	30.687	1.509	29.819	0.621	
Small class	36.829	0.043	37.687	0.054	38.529	0.006	
Overall	29.725	1.071	32.528	1.038	31.650	1.301	

Table 3: Quantitative comparison between linear (left) and log (middle) losses used for training. Comparison of PanoHDR-NeRF with NeRF then LDR2HDR is at right

4.4 Ablation study

LDR2HDR network architecture We compare between two recent SOTA single image HDR estimation architectures, namely HallucinationNet [22] and LANet [53], with and without our rendering loss. For this, we evaluate on the Laval Indoor HDR [12] test set and report results in sec. 4.4. With its saturation-driven attention, the LANet architecture outperforms HallucinationNet. In addition, the use of the rendering loss ℓ_{rend} (eq. (1)) helps the network focus on the bright light sources, which is crucial for accurate radiance reproduction.

Planar vs spherical sampling We compare the impact of planar vs spherical sampling (check supplementary) for training our network on LDR images. In both cases, we follow the hierarchical sampling strategy of NeRF [1] and train the network with 64 coarse samples and 128 fine samples. We observe that sphere sampling outperforms planar sampling for omnidirectional images as it does not oversample points at the poles but does so uniformly on the sphere. (more details in supplementary)

Loss in log space We evaluate the importance of computing the loss in log space (c.f. sec. 3.2) in sec. 4.4, which suggests that the network is able to estimate the high dynamic range equally well with or without the log-space loss. However, fig. 6 shows that training in linear space results in more floating artifacts and blurrier images than those obtained by the log loss. This is because HDR intensities vary significantly and log space helps us restrict the range which in turn makes it easier for NeRF-MLP to predict the radiance values.

Order of operations PanoHDR-NeRF uses a NeRF network trained on HDR images. It is also possible to reverse the order by training the NeRF on LDR images and pass its output through the LDR2HDR network (dubbed "NeRF-LDR2HDR"). We compare these options in sec. 4.4 and fig. 5. Though the metrics given in the table don't differ much, the supplementary video shows PanoHDR-NeRF produces temporally more stable results. This could be due to the averaging that naturally happens within the NeRF network.



Figure 5: Comparing input LDR, NeRF++, NeRF-LDR2HDR, PanoHDR-NeRF (ours), and GT panoramas. For each example, the figure shows a virtual test object relit to show the dynamic range. While NeRF++ is able to model the scene correctly, it is unable to capture the radiance of the scene. PanoHDR-NeRF is able to faithfully capture the radiance of the scene. We compare it with NeRF-LDR2HDR which estimates HDR from NeRF++ outputs. Although it is able to closely estimate the radiance, it leads to flickering between consecutive frames. Images tonemapped for display with $\gamma = 2.2$.



Figure 6: Comparing panoramas generated by PanoHDR-NeRF after loss in linear space and log space. For each example, the figure shows (top) the panorama with (bottom) virtual test objects relit to show the dynamic range. While the network learns the high dynamic range in both cases, we observe that taking loss in linear space leads to poor visual quality and floating artifacts in the output panoramas. PanoHDR-NeRF produces better results when trained in log space, consistent with [23, 43].

5 Discussion and conclusion

The main contribution of this work is a novel pipeline to predict the full HDR radiance of an indoor scene without using special hardware, careful scanning of the scene, or intricately calibrated camera configurations. Our pipeline can work with a single, off-the-shelf 360° camera that is moved around the scene. Although [22, 53] have demonstrated recovery of HDR intensities of forward facing scenes, they require multiple sets of LDR images at various exposures. Recovering HDR intensities for large indoor scenes using these methods requires an elaborate setup which is tedious and cumbersome. PanoHDR-NeRF can render novel 360° views from any point within an unbounded indoor scene in high dynamic range from a casually captured video. We show their use for the realistic relighting of virtual objects in real scenes, hopefully getting one step closer to democratizing augmented reality.

Limitations and future work Blurriness of the NeRF results is a big limitation of this work, despite using cone-casting from Mip-NeRF. We believe further improvements such as [3] can help in reconstructing sharper estimates. Another limitation is that the photographer capturing the scene ends up modifying the light field ever so slightly by casting shadows, creating reflections off of shiny surfaces, etc. Unfortunately, the intensity changes this creates are too soft for existing shadow detectors [1]. Methods modeling transient changes [1] could potentially be of help. The finetuning for each camera introduces additional effort though only once per camera. Future research can improve generalization across cameras, perhaps using multiple cameras for training, or through other data augmentation techniques. Our approach learns radiance and view synthesis in two independent steps by specialized networks. Exploring how both can be done simultaneously, potentially in conjunction with geometry and material estimation $[\mathbf{\overline{b}}, \mathbf{\overline{by}}]$, is an exciting direction for future work. PanoHDR-NeRF assumes that the scene is static, so the lighting should remain unchanged for the inference time. In addition, PanoHDR-NeRF is an offline method and therefore the novelty of the proposed method is limited to only capturing process and not rendering. Recent efforts have demonstrated how to significantly shorten training $[\underline{\mathbf{M}}, \underline{\mathbf{M}}]$ and inference $[\underline{\mathbf{M}}, \underline{\mathbf{M}}]$ times of NeRF-based approaches, these could straightforwardly be incorporated into our framework. By reducing the time between capture and visualization, PanoHDR-NeRF can be used for AR/VR applications such as virtual tours and VFX generation.

6 Acknowledgments

This research was supported by NSERC grant RGPIN-2020-04799, Compute Canada, and a MITACS Globalink internship to Pulkit Gera. The authors thank Bowei Chen for his early work on the project and David Ibarzabal for his help with data capture. We also thank Yohan Poirier-Ginter and Jinsong Zhang for their help.

References

- Benjamin Attal, Selena Ling, Aaron Gokaslan, Christian Richardt, and James Tompkin. Matryodshka: Real-time 6dof video view synthesis using multi-sphere images. In *Eur. Conf. Comput. Vis.*, 2020.
- [2] Jonathan T. Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-

Brualla, and Pratul P. Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *IEEE/CVF Int. Conf. Comput. Vis.*, 2021.

- [3] Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, 2022.
- [4] Tobias Bertel, Mingze Yuan, Reuben Lindroos, and Christian Richardt. Omniphotos: Casual 360° vr photography. ACM Trans. Graph., 2020.
- [5] Mark Boss, Raphael Braun, V. Jampani, Jonathan T. Barron, Ce Liu, and Hendrik P. A. Lensch. Nerd: Neural reflectance decomposition from image collections. *IEEE/CVF Int. Conf. Comput. Vis.*, 2021.
- [6] Kyunghyun Cho, Bart van Merrienboer, Çaglar Gülçehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using RNN encoder-decoder for statistical machine translation. In *Conf. Emp. Meth. Nat. Lang. Proc.*
- [7] Paul Debevec. Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *Conf. Comp. Graph. Int. Tech.*, SIGGRAPH, pages 189–198, 1998.
- [8] Paul Debevec and Jitendra Malik. Recovering high dynamic range radiance maps from photographs. In *ACM SIGGRAPH*, pages 369–378, 1997.
- [9] Paul Debevec, Yizhou Yu, and George Boshokov. Efficient view-dependent ibr with projective texture-mapping. In *EG Rend. Works.*, 1998.
- [10] Yuki Endo, Yoshihiro Kanamori, and Jun Mitani. Deep reverse tone mapping. *ACM Trans. Graph.*, 2017.
- [11] John Flynn, Michael Broxton, Paul E. Debevec, Matthew DuVall, Graham Fyffe, Ryan S. Overbeck, Noah Snavely, and Richard Tucker. Deepview: View synthesis with learned gradient descent. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, 2019.
- [12] Stephan J. Garbin, Marek Kowalski, Matthew Johnson, Jamie Shotton, and Julien Valentin. Fastnerf: High-fidelity neural rendering at 200fps. In *IEEE/CVF Int. Conf. Comput. Vis.*, 2021.
- [13] Marc-André Gardner, Kalyan Sunkavalli, Ersin Yumer, Xiaohui Shen, Emiliano Gambaretto, Christian Gagné, and Jean-François Lalonde. Learning to predict indoor illumination from a single image. ACM Trans. Graph., 2017.
- [14] Steven J. Gortler, Radek Grzeszczuk, Richard Szeliski, and Michael F. Cohen. The lumigraph. ACM, 1996.
- [15] Kai Gu, Thomas Maugey, Sebastian Knorr, and Christine Guillemot. Omni-nerf: neural radiance field from 360° image captures. In *Int. Conf. Multi. Expo*, 2022.
- [16] Takayuki Hara and Tatsuya Harada. Enhancement of novel view synthesis using omnidirectional image completion. ArXiv, abs/2203.09957, 2022.

- [17] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask R-CNN. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, 2017.
- [18] Peter Hedman, Tobias Ritschel, George Drettakis, and Gabriel J. Brostow. Scalable inside-out image-based rendering. *ACM Trans. Graph.*, 2016.
- [19] Peter Hedman, Julien Philip, True Price, Jan-Michael Frahm, George Drettakis, and Gabriel J. Brostow. Deep blending for free-viewpoint image-based rendering. *ACM Trans. Graph.*, 2018.
- [20] Ching-Yu Hsu, Cheng Sun, and Hwann-Tzong Chen. Moving in a 360 world: Synthesizing panoramic parallaxes from a single panorama. *ArXiv*, 2021.
- [21] Jingwei Huang, Zhili Chen, Duygu Ceylan, and Hailin Jin. 6-DOF VR videos with a single 360-camera. In *IEEE Virt. Real.*, 2017.
- [22] Xin Huang, Qi Zhang, Feng Ying, Hongdong Li, Xuan Wang, and Qing Wang. Hdr-nerf: High dynamic range neural radiance fields. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, 2022.
- [23] Wesley Khademi and Jonathan Ventura. View synthesis in casually captured scenes using a cylindrical neural radiance field with exposure compensation. SIGGRAPH '21. Association for Computing Machinery, 2021.
- [24] Siyeong Lee, Gwon Hwan An, and Suk-Ju Kang. Deep chain hdri: Reconstructing a high dynamic range image from a single low dynamic range image. *IEEE Access*, 2018.
- [25] Jinghui Li and Peiyu Fang. Hdrnet: Single-image-based hdr reconstruction using channel attention cnn. In *Int. Conf. Mult. Sign. Proc.*, 2019.
- [26] 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 *Eur. Conf. Comput. Vis.*, 2020.
- [27] Yu-Lun Liu, Wei-Sheng Lai, Yu-Sheng Chen, Yi-Lung Kao, Ming-Hsuan Yang, Yung-Yu Chuang, and Jia-Bin Huang. Single-image hdr reconstruction by learning to reverse the camera pipeline. *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, 2020.
- [28] Rafał K. Mantiuk and Maryam Azimi. Pu21: A novel perceptually uniform encoding for adapting existing quality metrics for hdr. *Pict. Coding Symp.*, 2021.
- [29] Mapillary. Opensfm. https://github.com/mapillary/OpenSfM, 2022.
- [30] Demetris Marnerides, Thomas Bashford-Rogers, Jonathan Hatchett, and Kurt Debattista. Expandnet: A deep convolutional neural network for high dynamic range expansion from low dynamic range content. *Comput. Graph. Forum (CGF)*, 2018.
- [31] Ricardo Martin-Brualla, Noha Radwan, Mehdi SM Sajjadi, Jonathan T Barron, Alexey Dosovitskiy, and Daniel Duckworth. Nerf in the wild: Neural radiance fields for unconstrained photo collections. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, 2021.
- [32] Moustafa Meshry, Dan B. Goldman, S. Khamis, Hugues Hoppe, Rohit Pandey, Noah Snavely, and Ricardo Martin-Brualla. Neural rerendering in the wild. *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, 2019.

- [33] Ben Mildenhall, Pratul P. Srinivasan, Rodrigo Ortiz Cayon, Nima Khademi Kalantari, Ravi Ramamoorthi, Ren Ng, and Abhishek Kar. Local light field fusion. ACM Trans. Graph., 2019.
- [34] 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 *Eur. Conf. Comput. Vis.*, 2020.
- [35] Ben Mildenhall, Peter Hedman, Ricardo Martin-Brualla, Pratul P. Srinivasan, and Jonathan T. Barron. Nerf in the dark: High dynamic range view synthesis from noisy raw images. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, 2022.
- [36] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.*, 2022.
- [37] Manish Narwaria, Rafał K. Mantiuk, Matthieu Perreira Da Silva, and Patrick Le Callet. Hdr-vdp-2.2: a calibrated method for objective quality prediction of high-dynamic range and standard images. *Journ. Elect. Imag.*, 2015.
- [38] Christian Reiser, Songyou Peng, Yiyi Liao, and Andreas Geiger. Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps. In *IEEE/CVF Int. Conf. Comput. Vis.*, 2021.
- [39] Allan G. Rempel, Matthew Trentacoste, Helge Seetzen, H. David Young, Wolfgang Heidrich, Lorne Whitehead, and Greg Ward. LDR2HDR: On-the-fly reverse tone mapping of legacy video and photographs. ACM Trans. Graph., 26(3), 2007.
- [40] Gernot Riegler and Vladlen Koltun. Free view synthesis. In *Eur. Conf. Comput. Vis.*, 2020.
- [41] Gernot Riegler and Vladlen Koltun. Stable view synthesis. *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, 2021.
- [42] Marcel Santana Santos, Ing Ren Tsang, and Nima Khademi Kalantari. Single image hdr reconstruction using a cnn with masked features and perceptual loss. *ACM Trans. Graph.*, 2020.
- [43] Ana Serrano, Incheol Kim, Zhili Chen, Stephen DiVerdi, Diego Gutierrez, Aaron Hertzmann, and Belen Masia. Motion parallax for 360° rgbd video. *IEEE Trans. Vis. Comput. Graph.*, 2019.
- [44] Vincent Sitzmann, Justus Thies, Felix Heide, Matthias Nießner, Gordon Wetzstein, and Michael Zollhöfer. Deepvoxels: Learning persistent 3d feature embeddings. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, 2019.
- [45] Peter-Pike J. Sloan, Jan Kautz, and John M. Snyder. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. *Conf. Comp. Graph. Int. Tech.*, 2002.
- [46] Matthew Tancik, Vincent Casser, Xinchen Yan, Sabeek Pradhan, Ben Mildenhall, Pratul P. Srinivasan, Jonathan T. Barron, and Henrik Kretzschmar. Block-nerf: Scalable large scene neural view synthesis. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, 2022.

- [47] Joanna Tarko, James Tompkin, and Christian Richardt. Real-time virtual object insertion for moving 360° videos. In *Int. Conf. Virt. Real. Cont. App. Indus.*, 2019.
- [48] Justus Thies, Michael Zollhöfer, and Matthias Nießner. Deferred neural rendering. *ACM Trans. Graph.*, 2019.
- [49] Haithem Turki, Deva Ramanan, and Mahadev Satyanarayanan. Mega-nerf: Scalable construction of large-scale nerfs for virtual fly-throughs. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, 2022.
- [50] David R. Walton and Anthony Steed. Dynamic hdr environment capture for mixed reality. *ACM Symp. Virt. Real. Soft. Tech.*, 2018.
- [51] Tianyu Wang, Xiaowei Hu, Qiong Wang, Pheng-Ann Heng, and Chi-Wing Fu. Instance shadow detection. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, 2020.
- [52] Jiale Xu, Jia Zheng, Yanyu Xu, Rui Tang, and Shenghua Gao. Layout-guided novel view synthesis from a single indoor panorama. *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, 2021.
- [53] Bangbang Yang, Yinda Zhang, Yijin Li, Zhaopeng Cui, S. Fanello, Hujun Bao, and Guofeng Zhang. Neural rendering in a room: Amodal 3d understanding and freeviewpoint rendering for the closed scene composed of pre-captured objects. ACM Trans. Graph., 41(4), 7 2022.
- [54] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. PlenOctrees for real-time rendering of neural radiance fields. In *IEEE/CVF Int. Conf. Comput. Vis.*, 2021.
- [55] Hanning Yu, Wentao Liu, Chengjiang Long, Bo Dong, Qin Zou, and Chunxia Xiao. Luminance attentive networks for hdr image and panorama reconstruction. *Comput. Graph. Forum (CGF)*, 2021.
- [56] Edward Zhang, Michael F. Cohen, and Brian Curless. Emptying, refurnishing, and relighting indoor spaces. *ACM Trans. Graph.*, 2016.
- [57] Jinsong Zhang and Jean-François Lalonde. Learning high dynamic range from outdoor panoramas. In *IEEE/CVF Int. Conf. Comput. Vis.*, 2017.
- [58] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. Nerf++: Analyzing and improving neural radiance fields. *CoRR*, abs/2010.07492, 2020.
- [59] Xiuming Zhang, Pratul P. Srinivasan, Boyang Deng, Paul Debevec, William T. Freeman, and Jonathan T. Barron. Nerfactor: Neural factorization of shape and reflectance under an unknown illumination. *ACM Trans. Graph.*, 2021.
- [60] Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. Stereo magnification: learning view synthesis using multiplane images. ACM Trans. Graph., 2018.