

# NeARportation: A Remote Real-time Neural Rendering Framework

Yuichi Hiroi  
Sony CSL Kyoto  
The University of Tokyo  
Kyoto / Tokyo, Japan  
yhiroii@g.ecc.u-tokyo.ac.jp

Yuta Itoh  
The University of Tokyo  
RIKEN AIP  
Tokyo, Japan  
yuta.itoh@iii.u-tokyo.ac.jp

Jun Rekimoto  
Sony CSL Kyoto  
The University of Tokyo  
Kyoto / Tokyo, Japan  
rekimoto@acm.org

## ABSTRACT

While presenting a photorealistic appearance plays a major role in immersion in Augmented Reality environment, displaying that of real objects remains a challenge. Recent developments in photogrammetry have facilitated the incorporation of real objects into virtual space. However, reproducing complex appearances, such as subsurface scattering and transparency, still requires a dedicated environment for measurement and possesses a trade-off between rendering quality and frame rate.

Our NeARportation framework combines server–client bidirectional communication and neural rendering to resolve these trade-offs. Neural rendering on the server receives the client’s head posture and generates a novel-view image with realistic appearance reproduction that is streamed onto the client’s display. By applying our framework to a stereoscopic display, we confirm that it can display a high-fidelity appearance on full-HD stereo videos at 35-40 frames per second (fps) according to the user’s head motion.

## CCS CONCEPTS

• **Human-centered computing** → **Mixed / augmented reality**.

## KEYWORDS

augmented virtuality, appearance reproduction, real-time rendering, neural rendering, remote rendering

### ACM Reference Format:

Yuichi Hiroi, Yuta Itoh, and Jun Rekimoto. 2022. NeARportation: A Remote Real-time Neural Rendering Framework. In *28th ACM Symposium on Virtual Reality Software and Technology (VRST '22)*, November 29-December 1, 2022, Tsukuba, Japan. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3562939.3565616>

## 1 INTRODUCTION

Photorealistic reproduction of the appearance and shape of real-world objects in virtual form is important in enhancing immersion and value judgments in a wide range of VR and AR applications.

---

This is a preprint to be presented at  
VRST '22, November 29-December 1, 2022, Tsukuba, Japan

However, it remains difficult to reproduce the photorealistic appearance aspects of real objects, such as gloss, transparency, and subsurface scattering, on displays.

Recent mobile photogrammetry technology allows nonexperts to measure the shape and color of real objects and easily import them into virtual space. However, these photogrammetry technologies assume that the surface is constant in radiance from any view angle. Thus, the surfaces of the reproduced subjects often possess a plaster-like texture. Researchers often utilize special hardware such as dome-shaped setups with structural light and multiple cameras [6, 16] to measure the complex, viewpoint-dependent light scattering on its surface, which is both costly and inaccessible. Moreover, even if complex light-scattering models can be measured, high-quality rendering models that can represent light scattering, such as ray tracing, create a trade-off between image quality and frame rate.

To resolve these trade-offs in measurement and rendering, we propose NeARportation, a framework that combines server-client bidirectional communication and neural rendering. NeARportation interactively displays the photorealistic appearances of real objects measured from readily available color images. In this framework, an AR/VR device communicates with a remote GPU server running neural radiance fields (NeRF) [17] to generate and transmit images according to the user’s viewpoint. We apply our framework to a naked-eye stereoscopic display and show that our framework can render photorealistic full-HD images at around 35 fps from stereoscopic viewpoints, tracking the user’s viewpoint position.

To the best of our knowledge, this is the first demonstration that neural rendering has the potential to interactively benefit AR/VR applications at reasonable frame rates and latencies on high-end display devices. Our major contributions include the following:

- Proposing NeARportation, a remote, real-time neural rendering framework for displaying photorealistic appearance in AR/VR devices.
- Demonstrating a proof-of-concept system of the NeARportation framework on a naked-eye stereoscopic display.
- Evaluating the reproducibility of appearance, achievable frame rate, and motion-to-photon latency in stereoscopic rendering.
- Providing discussions and future research directions for the current NeARportation framework.

## 2 RELATED WORK

Reproducing the photorealistic appearance of real objects has long been a research challenge. Conventional rendering approaches explicitly estimate a reflective model of an object’s surface in reproducing a photorealistic appearance. Hardware approaches include

domed structural lighting and cameras [3, 6, 16] or light field cameras [18, 21, 32], yet these dedicated camera systems maintain extensive measurement and reproduction costs. While deep neural network-based methods can estimate reflection models from color cameras [1, 14, 30], it remains difficult to generalize explicit reflection estimations to various object appearances.

Real-time rendering is a continuous research theme, mainly in game graphics [11]. When rendering a realistic appearance, ray casting in the rendering process can dramatically improve the expressive power of the image. However, state-of-the-art full ray-casting real-time rendering requires over 11 GB of GPU memory [13], making it unsuitable for AR/VR devices with relatively low computing power and less energy for extended use.

NeRF [17] employs volume rendering to compose new view images and can handle light attenuation and transparent objects. NeRF has developed remarkably within the past few years [36], and new models have further extended appearance representation [27, 33, 35] and realized real-time inference [4, 7, 12, 19, 31]. Our framework aims to achieve multiview, realistic appearance, and high-frame rate rendering in AR/VR by leveraging both the real-time performance and appearance representation of NeRF.

While recent NeRF architectures can generate various appearances in real time for AR/VR applications, the trade-offs of resolution, number of viewpoints, and computing resources remain unresolved. Our framework introduces a remote rendering system [25] that facilitates the process of rendering 3D graphics when the computing power of the local device is insufficient. In remote rendering, point cloud streaming [23, 29] and video streaming rendered by ray tracing [10, 24] exist as 3D content delivery. However, these works do not fully consider the application of neural rendering in remote rendering, except for a concurrent work [15] that uses remote rendering using NeRF with game engines. While the focus of this concurrent work is NeRF interaction in VR, we focus on stereoscopic and high-resolution presentation in the present study.

### 3 SYSTEM OVERVIEW

Figure 1 shows a system overview diagram of NeARportation. Our framework utilizes WebRTC-based remote rendering between a client with low computational resources (Fig 1, top) and a GPU server (Fig 1, bottom).

The client consists of a sensor (e.g., a camera) to acquire the user’s head pose, displays for rendered images streamed from the server, RTC peers, and an application such as a game engine. The server consists of the main process that mediates data transmission and reception, NeRF-based renderers, and RTC peers. The main process on the server includes a single receiver and senders equal to the number of peers. These peers and NeRF renderers are launched with the same number of displays on the client. Each RTC peer connects to internal processes with interprocess communication (IPC) to mediate between the processes and the P2P connection.

Before remote rendering, the client uploads a video of the object to be reconstructed in advance to train the NeRF model, which corresponds to the design of a virtual 3D scene in conventional AR/VR applications. Using the trained NeRF model, the server and client peers establish P2P connections. After connecting, the server sends the client the intrinsic parameters of the camera used to capture the video. The client sets the virtual cameras using the

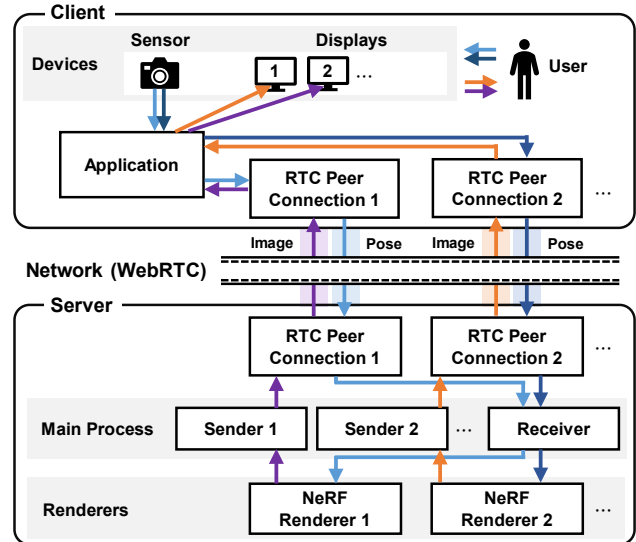


Figure 1: An overview of the NeARportation framework.

received parameters to reproduce the actual size of the real objects. If there are multiple viewpoints on the client (e.g., stereopsis), the virtual cameras are set for all viewpoints.

When the client completes the configuration of the virtual cameras at every frame, the application sends the pose matrix of each viewpoint camera to the server. Each packet received also contains labels for each viewpoint (e.g., for the left and right eye in the case of a stereoscopic display). The receiver on the server reads this label and distributes the pose matrix to the corresponding renderer process.


Each renderer on the server includes a buffer to store the pose matrix and a NeRF process. The pose matrix sent from the receiver to the renderer is stored in the pose matrix buffer. The matrix on the pose buffer is always overwritten with the latest one, and then the NeRF acquires the latest pose matrix and starts a new rendering immediately after the rendering is finished. After the NeRF completes the image rendering, it sends the image to the sender corresponding to each virtual camera label. The sender then delivers the rendered images to the client through MediaStream on WebRTC. Finally, the client draws the received images on the texture of a planar object that aligns exactly with each virtual camera’s field of view to present them on the displays.

## 4 IMPLEMENTATION

We implement a proof-of-concept system for displaying photorealistic appearance using NeARportation by applying our framework to a stereoscopic display.

### 4.1 Server Architecture

We use a desktop machine with Ubuntu 20.04 OS, 2x Intel (R) Core (TM) i9-10920X CPU @ 3.50 GHz, and 2x NVIDIA GeForce RTX 3090 GPUs as a server-side machine. This machine is connected to a 1-GbE wired network.



**Figure 2: (Left)** The naked-eye stereoscopic display used in our prototype (SONY ELF-SR1). The camera above the display detects the user’s head pose. **(Right)** A user viewing an image on a stereoscopic display. The head posture detected by the camera is synchronized with the virtual camera to show the same perspective image on display.

We adopt instant neural graphics primitive (instant-ngp) [19] for the NeRF renderer. We launch instant-ngp instances on two different GPUs. Each NeRF renderer process accesses instant-ngp through Python binding.

We use Momo and Ayame Labo (Shiguredo Corp.) as WebRTC clients and the signaling server, respectively. Momo clients encode and stream the videos from the virtual camera devices activated on the server. We choose VP 8 as an encoding format. As an IPC pipeline, we use *nanomsg*<sup>1</sup> as a messaging library, which establishes an IPC pipeline via port-to-port TCP connections.

## 4.2 Client Architecture

We use a desktop machine with Windows 10, an AMD Ryzen 5950X CPU @ 3.40GHz, and an NVIDIA GeForce RTX 3090 GPU as a client-side machine. This machine is connected to a 2.4-GHz Wi-Fi network in a 1-GbE LAN environment. The server and client are connected to different network segments.

As a naked-eye stereoscopic display, we use a Sony ELF-SR1 (Fig. 2). The display comprises a  $3840 \times 2160$ -pixel LCD panel angled at  $45^\circ$  and a high-speed camera sensor for tracking the user’s head pose. Micro-optical lenses are installed on the front of the panel to guide two images to the left and right eyes to achieve stereoscopic viewing. Corresponding to the training images, we use the half-resolution of this panel (full HD).

We implement the virtual application in Unity 2020.3.36f and Sony Spatial Display SDK for Unity. The head posture detected by the sensor is immediately reflected in the anchors of the virtual cameras in the left and right eyes. The images rendered from the virtual cameras are fused according to the microlens geometry. We use Microsoft MixedReality-WebRTC<sup>2</sup> as a WebRTC client. The client applies the video received from the server to a planar object.

## 4.3 Training Image Acquisition

For stereopsis, a stereo image pair must use the correct disparity according to the user’s interpupillary distance. Although NeRF generally uses COLMAP [22] to roughly determine initial camera poses, COLMAP does not match the actual size of objects or their vertical orientation with reality. Instead, we place the ChArUcO [8] marker under the object and capture training images around the object.

<sup>1</sup><https://nanomsg.org/>

<sup>2</sup><https://github.com/microsoft/MixedReality-WebRTC>



**Figure 3: Qualitative performance overview of the NeARportation framework. (a)** Captured images of real objects. **(b)** Real capture images of the reproduced appearances displayed on ELF-SR1, taken from two different viewpoints. Note that parallax artifacts appear in the images capturing the stereoscopic display because each image is taken directly with a monocular camera. This framework is interactive and is best viewed on video; thus, we urge readers to view our supplemental video.

We train instant-ngp in  $3 \times 10^4$  steps with a batch size of  $2^{18}$ , which takes approximately 4 minutes. When rendering, we crop only the area above the ChArUcO to improve the frame rate.


## 5 EXPERIMENTS

We evaluated the reproducibility of complex appearances when displaying the captured objects and the performance of our prototype.

### 5.1 Appearance Reproduction

We captured real objects with complex appearances and displayed the reproduced stereo images with our NeARportation prototype. To obtain training images, we first captured the object with the iPhone SE 2 for 3-5 minutes and cropped the video as images at 2 fps. These images were trained in instant-ngp and displayed on the application using our NeARportation framework. Figure 3 (a) shows examples of our training images. The objects used in the experiment and the number of training images are as follows: a kiwi cut in half (267 images), agar jelly with fruits (288 images), a tea whisk (274 images), a metallic cup (317 images), and a CD (440 images).

Figure 3 (b) shows the reproduced appearances. From the 1st and 2nd rows of Fig. 3, we confirmed that NeARportation could



**Figure 4: Transitions of (top) frame rate and (bottom) RTL and rendering latency when we keep NeARportation running for 10 seconds. The gray area represents when the application checks the communication between the client and the server at regular intervals.**

interactively display the sparkling at the cross-section of a kiwi or the complex refraction in agar from novel viewpoints. Additionally, we confirmed that our framework has sufficient reproducibility for subtle, gaze-dependent appearances following the user’s head motion, such as very thin objects (3rd row), mirror reflections (4th row), and structural colors (5th row).

## 5.2 Performance Evaluation

Our NeARportation framework is intended to perform neural rendering interactively. To evaluate the performance toward this goal, we measure the current prototype’s frame rate and round-trip latency (RTL) for 10 seconds. To measure these metrics, we added millisecond timestamps to packets sent by the client. The server then returned the timestamp received from the client intact; at the same time, it returned the corresponding image. The client measured the RTL by the difference between the current time and the received timestamp and the frame rate by the timestamp interval. We also added the rendering time of instant-ngp for each image to the returned packet to determine the percentage of NeRF rendering in the RTL. During the measurement, the users move their heads to see the effect on rendering time.

Figure 4 shows the results of the performance measurement. In the current implementation, the application regularly checks whether P2P communication between the client and server is sustained. Therefore, video streaming is momentarily interrupted during this period (approximately 200 ms). Both frame rate and latency are restored to their original values after the application confirms communication with the server.

From Figure 4 (a), the frame rate of our prototype is around 35 ~ 40 fps, which confirms that our prototype can provide sufficient frame rates for interactive applications such as games and videos. In contrast, we need further improvements to achieve the frame rates required for modern VR headsets (>90 fps) [5]. We will discuss the ways to extend our framework to higher frame rates in terms of cloud computing and improved neural network models.

From Figure 4 (b), the RTL of this prototype oscillates at about 100 ~ 400 ms. Additionally, we see that the rendering time of instant-ngp is almost constant regardless of the head position (30 ~ 40 ms). From the results, we confirmed that most of the RTL on this system depends on round-trip latency on the network. This RTL could be improved with future network bandwidth improvements. Later, we discuss possible research directions to compensate for network latency by predicting head motion.

## 6 DISCUSSION AND FUTURE WORK

We investigate the remaining issues of the current NeARportation framework and discuss the prospects of future research directions.

*Head Motion Prediction.* An RTL exceeding 20 ms is perceived to have detrimental effects [20]. However, due to the nature of our framework, which uses remote rendering, it is difficult to keep latency below 20 ms. Compensating for this latency is currently out of scope, as this paper focuses on proving the concept for our framework. One future research direction involves compensating for this RTL using motion prediction, such as the Kalman filter-based approaches [9, 28] or a time-sequential neural network model [34].

*Improving Frame Rate.* Reducing rendering time is essential for improving the frame rate. A simple approach to reduce rendering time is distributed rendering with multiple GPUs, in which poses stored in a buffer are split and rendered on multiple GPUs. Another future research direction involves a novel NeRF network, which can render images from multiple poses simultaneously on a single GPU. In this direction, an efficient method of caching scenes between viewpoints [12] can be applied.

*Perceptive Alignment with Real Environment.* The appearance reproduced by our framework will be improved by blending the image with the environment in which the display is placed, such as color and dynamic range correction, depth-of-field reproduction, and relighting. Although various works regarding relighting in NeRF exist [2, 26, 37], to the best of our knowledge, no model has been proposed to perform it in real time.

## 7 CONCLUSION

We proposed NeARportation, a framework for displaying photo-realistic appearance with remote rendering and a neural renderer. The proof-of-concept system on a face-tracking stereoscopic display showed that our framework can interactively present objects’ complex appearances (e.g., specular reflection, structural color, and subsurface scattering) from the easily available color images. We also described the prospects for improving such remote rendering for photorealistic appearance. Our work aims to inspire the community along this framework to further improve the system, such as higher frame rate and more accurate head motion tracking. We believe this work opens the door to AR/VR experiences that are *NeAR* reality.

## ACKNOWLEDGMENTS

This project was partially supported by JST Moonshot R&D Grant Number JPMJMS2012, the commissioned research by National Institute of Information and Communications Technology (NICT) Japan, JST PRESTO Grant Number JPMJPR17J2, JSPS KAKENHI Grant Number JP20H05958, JP20H04222, and JP22J01340, Japan.

## REFERENCES

- [1] Louis-Philippe Asselin, Denis Laurendeau, and Jean-François Lalonde. 2020. Deep SVBRDF Estimation on Real Materials. In *2020 International Conference on 3D Vision (3DV)*. IEEE Computer Society, Los Alamitos, CA, USA, 1157–1166. <https://doi.org/10.1109/3DV50981.2020.00126>
- [2] Mark Boss, Raphael Braun, Varun Jampani, Jonathan T. Barron, Ce Liu, and Hendrik P. A. Lensch. 2021. NeRD: Neural Reflectance Decomposition from Image Collections. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*. IEEE Computer Society, Los Alamitos, CA, USA, 12664–12674. <https://doi.org/10.1109/ICCV48922.2021.01245>
- [3] Michael Broxton, John Flynn, Ryan Overbeck, Daniel Erickson, Peter Hedman, Matthew DuVall, Jason Dourgian, Jay Busch, Matt Whalen, and Paul Debevec. 2020. Immersive Light Field Video with a Layered Mesh Representation. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 39, 4 (2020), 86:1–86:15.
- [4] Hwann-Tzong Chen Cheng Sun, Min Sun. 2022. Direct Voxel Grid Optimization: Super-fast Convergence for Radiance Fields Reconstruction. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, Los Alamitos, CA, USA, 5449–5459. <https://doi.org/10.1109/CVPR52688.2022.00538>
- [5] Eduardo Cuervo, Krishna Chintalapudi, and Manikanta Kotaru. 2018. Creating the Perfect Illusion : What will it take to Create Life-Like Virtual Reality Headsets?. In *HotMobile'18 (hotmobile'18 ed.)*. Association for Computing Machinery. <https://doi.org/10.1145/3177102.3177115>
- [6] Paul Debevec, Tim Hawkins, Chris Tchou, Haarm-Pieter Duiker, Westley Sarokin, and Mark Sagar. 2000. Acquiring the Reflectance Field of a Human Face. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '00)*. ACM Press/Addison-Wesley Publishing Co., USA, 145–156. <https://doi.org/10.1145/344779.344855>
- [7] Stephan J Garbin, Marek Kowalski, Matthew Johnson, Jamie Shotton, and Julien Valentin. 2021. FastNeRF: High-Fidelity Neural Rendering at 200FPS. *arXiv preprint arXiv:2103.10380* (2021).
- [8] Sergio Garrido-Jurado, Rafael Muñoz-Salinas, Francisco J. Madrid-Cuevas, and Manuel J. Marín-Jiménez. 2014. Automatic Generation and Detection of Highly Reliable Fiducial Markers Under Occlusion. *Pattern Recognition* 47, 6 (2014), 2280–2292. <https://doi.org/10.1016/j.patcog.2014.01.005>
- [9] Serhan Gül, Sebastian Bosse, Dimitri Podborski, Thomas Schierl, and Cornelius Hellge. 2020. Kalman Filter-Based Head Motion Prediction for Cloud-Based Mixed Reality. In *Proceedings of the 28th ACM International Conference on Multimedia*. Association for Computing Machinery, New York, NY, USA, 3632–3641. <https://doi.org/10.1145/3394171.3413699>
- [10] Serhan Gül, Dimitri Podborski, Thomas Buchholz, Thomas Schierl, and Cornelius Hellge. 2020. Low-Latency Cloud-Based Volumetric Video Streaming Using Head Motion Prediction. In *Proceedings of the 30th ACM Workshop on Network and Operating Systems Support for Digital Audio and Video (Istanbul, Turkey) (NOSSDAV '20)*. Association for Computing Machinery, New York, NY, USA, 27–33. <https://doi.org/10.1145/3386290.3396933>
- [11] Eric Haines, Naty Hoffman, et al. 2018. *Real-time rendering*. Crc Press.
- [12] T. Hu, S. Liu, Y. Chen, T. Shen, and J. Jia. 2022. EfficientNeRF - Efficient Neural Radiance Fields. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, Los Alamitos, CA, USA, 12892–12901. <https://doi.org/10.1109/CVPR52688.2022.01256>
- [13] Huang Jensen. 2020. NVIDIA Omniverse Now in Open Beta. (Oct. 2020). NVIDIA GPU Technology Conference (GTC) 2020 Keynote.
- [14] Alen Joy and Charalambos Poollis. 2022. Multi-view Gradient Consistency for SVBRDF Estimation of Complex Scenes under Natural Illumination. <https://doi.org/10.48550/ARXIV.2202.13017>
- [15] Naruya Kondo, So Kuroki, Ryosuke Hyakuta, Yutaka Matsuo, Shixiang Shane Gu, and Yoichi Ochiai. 2022. Immersive Real World through Deep Billboards. In *ACM SIGGRAPH 2022 Immersive Pavilion (Vancouver, BC, Canada) (SIGGRAPH '22)*. Association for Computing Machinery, New York, NY, USA, Article 8, 2 pages. <https://doi.org/10.1145/3532834.3536210>
- [16] Wan-Chun Ma, Tim Hawkins, Pieter Peers, Charles-Felix Chabert, Malte Weiss, and Paul Debevec. 2007. Rapid Acquisition of Specular and Diffuse Normal Maps from Polarized Spherical Gradient Illumination. In *Proceedings of the 18th Eurographics Conference on Rendering Techniques (Grenoble, France) (EGSR'07)*. Eurographics Association, Goslar, DEU, 183–194.
- [17] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. 2020. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In *ECCV*.
- [18] Tim Milliron, Chrissy Szczupak, and Orin Green. 2017. Hallelujah: The World's First Lytro VR Experience. In *ACM SIGGRAPH 2017 VR Village (Los Angeles, California) (SIGGRAPH '17)*. Association for Computing Machinery, New York, NY, USA, Article 7, 2 pages. <https://doi.org/10.1145/3089269.3089283>
- [19] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. 2022. Instant Neural Graphics Primitives with a Multiresolution Hash Encoding. *ACM Trans. Graph.* 41, 4, Article 102 (July 2022), 15 pages. <https://doi.org/10.1145/3528223.3530127>
- [20] Juraj Obert, J. M. P. van Waveren, and Graham Sellers. 2012. Virtual Texturing in Software and Hardware. In *ACM SIGGRAPH 2012 Courses (Los Angeles, California) (SIGGRAPH '12)*. Association for Computing Machinery, New York, NY, USA, Article 5, 29 pages. <https://doi.org/10.1145/2343483.2343488>
- [21] Ryan S. Overbeck, Daniel Erickson, Daniel Evangelakos, Matt Pharr, and Paul Debevec. 2018. A System for Acquiring, Processing, and Rendering Panoramic Light Field Stills for Virtual Reality. *ACM Trans. Graph.* 37, 6, Article 197 (dec 2018), 15 pages. <https://doi.org/10.1145/3272127.3275031>
- [22] J. L. Schonberger and J. Frahm. 2016. Structure-from-Motion Revisited. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, Los Alamitos, CA, USA, 4104–4113. <https://doi.org/10.1109/CVPR.2016.445>
- [23] Sebastian Schwarz, Marius Preda, Vittorio Baroncini, Madhukar Budagavi, Pablo Cesar, Philip A. Chou, Robert A. Cohen, Maja Krivokuća, Sébastien Lasserre, Zhu Li, Joan Llach, Khaled Mammou, Rufael Mekuria, Ohji Nakagami, Ernestasia Siahaan, Ali Tabatabai, Alexis M. Tourapis, and Vladyslav Zakharchenko. 2019. Emerging MPEG Standards for Point Cloud Compression. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems* 9, 1 (2019), 133–148. <https://doi.org/10.1109/JETCAS.2018.2885981>
- [24] Woong Seo, Sanghun Park, and Insung Ihm. 2022. Efficient Ray Tracing of Large 3D Scenes for Mobile Distributed Computing Environments. *Sensors* 22, 2 (2022). <https://doi.org/10.3390/s22020491>
- [25] Shu Shi and Cheng-Hsin Hsu. 2015. A Survey of Interactive Remote Rendering Systems. *ACM Comput. Surv.* 47, 4, Article 57 (may 2015), 29 pages. <https://doi.org/10.1145/2719921>
- [26] Pratul P. Srinivasan, Boyang Deng, Xiuming Zhang, Matthew Tancik, Ben Mildenhall, and Jonathan T. Barron. 2021. NeRV: Neural Reflectance and Visibility Fields for Relighting and View Synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 7495–7504.
- [27] Mohammed Suhail, Carlos Esteves, Leonid Sigal, and Ameesh Makadia. 2022. Light Field Neural Rendering. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, Los Alamitos, CA, USA, 8259–8269. <https://doi.org/10.1109/CVPR52688.2022.00809>
- [28] Nicholas T. Swafford, Bastiaan J. Boom, Kartic Subr, David Sinclair, Darren Cosker, and Kenny Mitchell. 2014. Dual Sensor Filtering for Robust Tracking of Head-Mounted Displays. In *Proceedings of the 20th ACM Symposium on Virtual Reality Software and Technology (Edinburgh, Scotland) (VRST '14)*. Association for Computing Machinery, New York, NY, USA, 221–222. <https://doi.org/10.1145/2671015.2675694>
- [29] Jeroen van der Hooft, Tim Wauters, Filip De Turck, Christian Timmerer, and Hermann Hellwagner. 2019. Towards 6DoF HTTP Adaptive Streaming Through Point Cloud Compression. In *Proceedings of the 27th ACM International Conference on Multimedia (Nice, France) (MM '19)*. Association for Computing Machinery, New York, NY, USA, 2405–2413. <https://doi.org/10.1145/3343031.3350917>
- [30] Giuseppe Vecchio, Simone Palazzo, and Concetto Spampinato. 2021. SurfaceNet: Adversarial SVBRDF Estimation from a Single Image. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*. 12820–12828. <https://doi.org/10.1109/ICCV48922.2021.01260>
- [31] Liao Wang, Jiakai Zhang, Xinhang Liu, Fuqiang Zhao, Yanshun Zhang, Yingliang Zhang, Minye Wu, Jingyi Yu, and Lan Xu. 2022. Fourier PlenOctrees for Dynamic Radiance Field Rendering in Real-Time. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 13524–13534.
- [32] Ting-Chun Wang, Jun-Yan Zhu, Nima Khademi Kalantari, Alexei A. Efros, and Ravi Ramamoorthi. 2017. Light Field Video Capture Using a Learning-Based Hybrid Imaging System. *ACM Trans. Graph.* 36, 4, Article 133 (jul 2017), 13 pages. <https://doi.org/10.1145/3072959.3073614>
- [33] Suttisak Wizadwongsa, Pakkapon Phongthawee, Jiraphon Yenphraphai, and Supasorn Suwajanakorn. 2021. NeX: Real-time View Synthesis with Neural Basis Expansion. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, Los Alamitos, CA, USA, 8530–8539. <https://doi.org/10.1109/CVPR46437.2021.00843>
- [34] Erwin Wu and Hideki Koike. 2019. FuturePose - Mixed Reality Martial Arts Training Using Real-Time 3D Human Pose Forecasting With a RGB Camera. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*. 1384–1392. <https://doi.org/10.1109/WACV.2019.00152>
- [35] Liwen Wu, Jae Yong Lee, Anand Bhattad, Yuxiong Wang, and David Forsyth. 2021. DiVer: Real-time and Accurate Neural Radiance Fields with Deterministic Integration for Volume Rendering. [arXiv:2111.10427](https://arxiv.org/abs/2111.10427) [cs.CV]
- [36] Yiheng Xie, Towaki Takikawa, Shunsuke Saito, Or Litany, Shiqin Yan, Nupur Khan, Federico Tombari, James Tompkin, Vincent Sitzmann, and Srinath Sridhar. 2022. Neural Fields in Visual Computing and Beyond. *Computer Graphics Forum* 41, 2 (2022), 641–676. <https://doi.org/10.1111/cgf.14505> <https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.14505>
- [37] Yuanqing Zhang, Jiaming Sun, Xingyi He, Huan Fu, Rongfei Jia, and Xiaowei Zhou. 2022. Modeling Indirect Illumination for Inverse Rendering. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, Los Alamitos, CA, USA, 18622–18631. <https://doi.org/10.1109/CVPR52688.2022.01809>