

ORCa: Glossy Objects as Radiance-Field Cameras

Kushagra Tiwary^{1*}, Akshat Dave^{2*}, Nikhil Behari¹, Tzofi Klinghoffer¹,
 Ashok Veeraraghavan², Ramesh Raskar¹
¹Massachusetts Institute of Technology, ²Rice University
 {ktiwary, behari, tzofi, raskar}@mit.edu, {akshat, vashok}@rice.edu

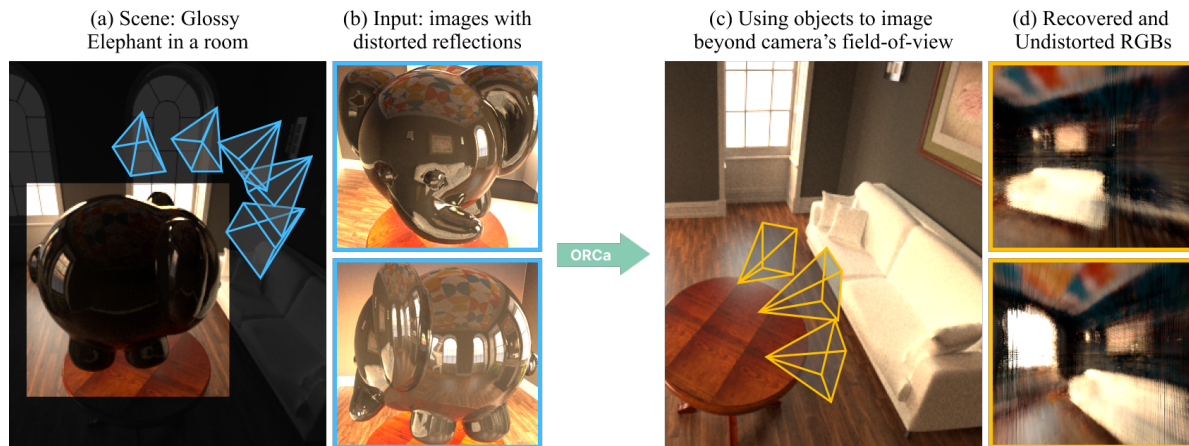


Figure 1. **Objects as radiance-field cameras.** We convert everyday objects with unknown geometry (a) into radiance-field cameras by modeling multi-view reflections (b) as projections of the 5D radiance field of the environment. We convert the object surface into a virtual sensor to capture this radiance field (c), which enables depth and radiance estimation of the surrounding environment. We can then query this radiance field to perform beyond field-of-view novel view synthesis of the environment (d).

Abstract

Reflections on glossy objects contain valuable and hidden information about the surrounding environment. By converting these objects into cameras, we can unlock exciting applications, including imaging beyond the camera’s field-of-view and from seemingly impossible vantage points, e.g. from reflections on the human eye. However, this task is challenging because reflections depend jointly on object geometry, material properties, the 3D environment, and the observer’s viewing direction. Our approach converts glossy objects with unknown geometry into radiance-field cameras to image the world from the object’s perspective. Our key insight is to convert the object surface into a virtual sensor that captures cast reflections as a 2D projection of the 5D environment radiance field visible to and surrounding the object. We show that recovering the environment radiance fields enables depth and radiance estimation from the object to its surroundings in addition to beyond field-of-view

novel-view synthesis, i.e. rendering of novel views that are only directly visible to the glossy object present in the scene, but not the observer. Moreover, using the radiance field we can image around occluders caused by close-by objects in the scene. Our method is trained end-to-end on multi-view images of the object and jointly estimates object geometry, diffuse radiance, and the 5D environment radiance field. For more information, visit our [website](#).

1. Introduction

Imagine that you’re driving down a city street that is packed with lines of parked cars on both sides. Inspection of the cars’ glass windshields, glossy paint, and plastic reveal sharp, but faint and distorted views of the surroundings that might be otherwise hidden from you. Humans can infer depth and semantic cues about the occluded areas in the environment by processing reflections visible on reflective objects, internally decomposing the object geometry and radiance from the specular radiance being reflected onto it. Our aim is to decompose the object from its reflections to “see” the world from the object’s perspective, effectively turning

*Equal contribution

(a) Learning Neural Implicit Surfaces [Sec. 3.2]

(b) Objects Surface as Virtual Sensor [Sec. 3.3]

(c) Learning Environment Radiance Fields [Sec. 3.4]

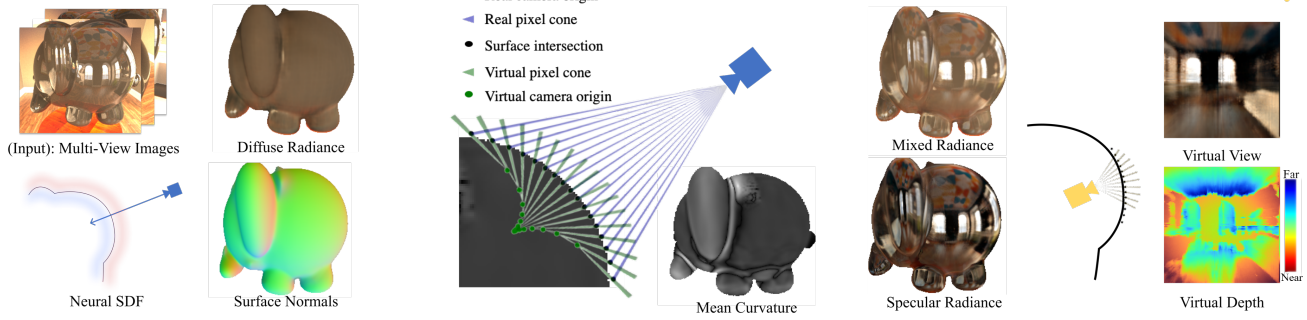


Figure 2. **ORCa Overview.** We jointly estimate the object’s geometry and diffuse along with the environment radiance field estimation through a three-step approach. First, we model the object as a neural implicit surface (a). We model the reflections on virtual viewpoints (b) estimated analytically from surface properties. We model the environment as a radiance field queried on these viewpoints (c). Both neural implicit surface and environment radiance field are trained jointly on multi-view images of the object using a photometric loss.

the object into a camera that images its environment. However, reflections pose a long-standing challenge in computer vision as the reflections are a 2D projection of an unknown 3D environment that is distorted based on the shape of the reflector.

To capture the 3D world from the object’s perspective, we model the object’s surface as a virtual sensor that captures the 2D projection of a 5D environment radiance field surrounding the object. This environment radiance field consists largely of areas only visible to the observer through the object’s reflections. Our use of environment radiance fields not only enables depth and radiance estimation from the object to its surroundings but also enables beyond *field-of-view* novel-view synthesis, i.e. rendering of novel views that are only directly visible to the glossy object present in the scene but not the observer. Unlike conventional approaches that model the environment as a 2D map, our approach models it as a 5D field without assuming the scene is infinitely far away. Moreover, by sampling the 5D radiance field, instead of a 2D map, we can capture depth and images around occluders, such as close-by objects in the scene, as shown in Fig. 3. These applications cannot be done from a 2D environment map.

We aim to decompose reflections on the object’s surface, from its surface and exploit those reflections to construct a radiance field surrounding the object, therefore capturing the 3D world in the process. This is a challenging task because the reflections are extremely sensitive to local object geometry, viewing direction and inter-reflections due to the object’s surface. To capture this radiance field, we convert glossy objects with unknown geometry and texture into radiance-field cameras. Specifically, we exploit neural rendering to estimate the local surface of the object viewed

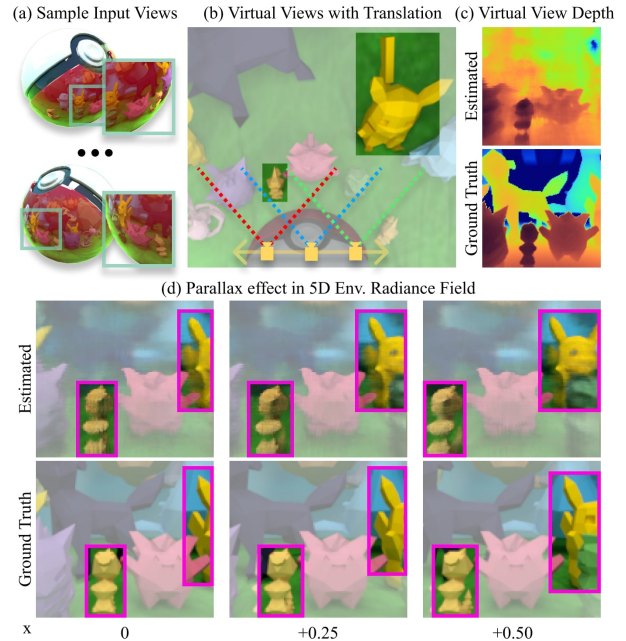


Figure 3. **Advantages of 5D environment radiance field.** Modeling reflections on object surfaces (a) as a 5D env. radiance field enables beyond *field-of-view* novel-view synthesis, including rendering of the environment from translated virtual camera views (b). Depth (c) and environment radiance of translated and parallax views can further enable imaging behind occluders, for example revealing the tails behind the primary Pokemon occluders (d).

from each pixel of the real camera. We then convert this local surface into a virtual pixel that captures radiance from the environment. This virtual pixel captures the environment radiance as shown in Fig 5. We estimate the outgo-

ing frustum from the virtual pixel as a cone that samples the scene. By sampling the scene from many virtual pixels on the object surface, we construct an environment radiance field that can be queried independently of the object surface, enabling beyond *field-of-view* novel-view synthesis from previously unsampled viewpoints.

Our approach jointly estimates object geometry, diffuse radiance, and the environment radiance field from multi-view images of glossy objects with unknown geometry and diffuse texture in three steps. First, we use neural signed distance functions (SDF) and an MLP to model the glossy object’s geometry as a neural implicit surface and diffuse radiance, respectively, similar to PANDORA [10]. Then, for every pixel on the observer’s camera, we estimate the virtual pixels on the object’s surface based on the estimated local geometry from the neural SDF. We analytically compute the parameters of the virtual cone through the virtual pixel. Lastly, we use the cone formulation in MipNeRF [5] to cast virtual cones from the virtual camera to recover the environment radiance .

To summarize, we make the following contributions:

- We present a method to convert implicit surfaces into virtual sensors that can image their surroundings using virtual cones. (Sec. 3.3)
- We jointly estimate object geometry, diffuse radiance, and estimate the 5D environment radiance field surrounding the object. (Fig. 7 & 9)
- We show that the environment radiance field can be queried to perform *beyond-field-of-view* novel viewpoint synthesis, i.e render views only visible to the object in the scene (Section 3.4)

Scope. We only model glossy objects with low roughness as such specular reflections tend to have a high signal-to-noise ratio, therefore, are a sharper estimate of the environment radiance field. However, we note that the virtual cone computation can be extended to model the cone radius as a function of surface roughness. Deblurring approaches can further improve the resolution of estimated environment. In addition, we approximate the local curvature using mean curvature, which fails for objects with varying radius of curvature along the tangent space. We explain how our virtual cone curvature estimation can be extended to handle general shape operators in the supplementary material. Lastly, similar to other multi-view approaches, our approach relies on a sufficient virtual baseline between virtual viewpoints to recover the environment radiance field.

2. Related Work

2.1. Modeling reflections

Catadioptric imaging systems incorporate reflections on curved reflective mirrors to expand the field of view of con-

ventional cameras [2, 26], to increase the baseline of light field cameras [11, 40] and to perform novel view synthesis from a single capture [44]. These works assume the geometry of the reflecting surface is known or calibrated while we create a catadioptric imaging system from everyday glossy objects of unknown geometry. Grossberg et al. [15] propose a generalized model for light transport through imaging systems, including catadioptric systems. Our work focuses on light transport reflecting off a general object in the scene.

Light field imaging is shown to be effective for reflection removal [21, 27], reconstructing specular surfaces [18], intrinsic decomposition [1], and neural rendering [35]. These works typically consider planar reflections or require training on synthetic datasets, while our approach models reflections on complex geometry and is unsupervised. Prior works have also utilized additional light properties such as polarization [9, 10, 20] and time-of-flight [17, 33] for the separation of the reflected component and specular surface reconstruction. While the input of our approach is RGB images, there is scope for improving reconstruction quality by supplementing the algorithm with these additional cues.

Recent progress in neural radiance fields has enabled impressive novel view rendering and geometry reconstruction from multi-view images [24]. MipNeRF [5] demonstrates better novel view synthesis by modeling outgoing rays as cones to enable anti-aliasing. RefNeRF [41] proposes Integrated Directional Embeddings for improved novel view synthesis of reflections. NeRFReN [16] separates diffuse and specular radiance by using separate neural networks. Neural Catacaustics [19] propose a neural warping method to model reflections by learning the caustics of the surface. While these works focus on novel-view synthesis of the scene from the primary camera, we perform view synthesis that is beyond the line-of-sight of the primary camera, i.e., rendering views only visible to the objects present in the scene, while jointly estimating object geometry and separating diffuse and specular radiance.

2.2. Environment Estimation

Recovering underlying scene properties from multiple images is inherently ill-posed [32], but can be regularized using the natural statistics of scene properties as a prior [4, 34]. Recent works exploit this prior through deep neural networks and demonstrate inverse rendering of indoor scenes from a single image [12, 22, 43, 48]. However, these techniques typically recover only coarse representations of lighting and cannot reconstruct fine details of the environment. Lombardi *et al.* [23] recover environment and reflectance, assuming the scene is composed of known geometry and uniform material. Georgolis *et al.* [13] recover the environment map behind the camera from a single image of a glossy object, assuming the object is composed of textureless materials and using ground truth seg-

Approach	Input	Scene Geometry	Algorithm Type	Environment Dimension
Neural Illum	Single RGB	Normals	Supervised	2D
Lighthouse	Stereo RGB pair	Multi-plane Image	Supervised	5D
NeRFFactor	Multi-view RGB	Volumetric Albedo	Semi-supervised	2D
RefNeRF	Multi-view RGB	Volumetric Albedo	Self-supervised	2D
PANDORA	Multi-view RGB+Pol.	Neural SDF	Self-supervised	2D
ORCa (Ours)	Multi-view RGB	Neural SDF	Self-supervised	5D

Figure 4. **Comparison of environment estimation approaches.** Approaches such Neural Illum [36], Lighthouse [37], NeRFFactor [47] train on datasets of natural illumination maps to regularize ill-posedness of environment estimation. PANDORA [10] and RefNeRF [41] exploit multi-view reflections on objects but assume the surrounding environment is infinitely far away and is modeled with a flat 2D map. From multi-view reflections, ORCa converts the object surface into a virtual sensor and extracts the 5D radiance field of the environment.

mentation masks. Song *et al.* [36] estimate plausible environment maps by mapping reflections in the image and inpainting unmapped regions. Srinivasan *et al.* [37] capture stereo image pairs and estimate plausible spatially-coherent environment maps. NeRD [7], NeRFFactor [47] and NeuralPIL [8] employ data-driven priors for lighting and BRDF in a NeRF-based approach for radiance decomposition. Park *et al.* [31] use RGB-D videos to estimate environment map. Swedish *et al.* [39] recover high-frequency illumination map from the shadows of an object with known geometry. PhysSG [46] and Munkberg *et al.* [25] perform inverse rendering from multi-view images by modeling the surface as signed distance functions. PANDORA [10] performs radiance decomposition from polarized RGB images.

3. Learning environment radiance fields from multi-view reflections

3.1. Overview

Reflections on glossy objects offer a glimpse into the surrounding environment beyond the camera’s field of view. From multi-view images of a glossy object with unknown geometry and albedo, we aim to recover the 5D radiance field of the surrounding environment. The mapping from images captured by the observer to the surrounding environment depends on the glossy object’s surface properties, in particular, the surface normals and curvature. We first cast a cone from the observer camera’s center-of-projection through each pixel viewing the scene. When the cone intersects the object’s surface, it reflects, causing the cone to be transformed (Fig. 5). The transformed cone, referred to as a virtual cone, samples the environment and is primar-




Figure 5. **Virtual Sensor.** We image the world through the object by modeling each pixel’s specular radiance as a projection of the 5D radiance field of the environment onto the object’s surface. We capture the radiance field by treating the surface area on the object that the pixel views, dS_t , as a single-pixel virtual camera with its center-of-projection at v_o . We cast virtual cones through the virtual sensor to capture the 5D radiance field of the environment.

ily responsible for the specular radiance observed on the glossy object. Our key insight is that the reflections captured by the observer’s camera can be modeled as a projection of the environment radiance field onto the object’s surface. By modeling the reflected rays as a cone and computing the parameters of the cone, we can more accurately estimate the projected environment radiance field onto the object surface, as shown in Fig. 10.

ORCa is composed of three steps: modeling the object’s geometry as a neural implicit surface (Sec. 3.2), converting the object’s surface into a virtual sensor (Sec. 3.3), and modeling the environment radiance field as a projection along these virtual cones (Sec. 3.4). The learned environment radiance field can then be queried on novel viewpoints to show occluded areas in the scene. Fig. 2 depicts our output for each component on a scene rendered with a complex glossy object and 3D environment. Fig. 6 shows our system architecture. Next, we describe each step in detail.

3.2. Learning Neural implicit Surfaces

Neural Signed Distance Function We model the object geometry as a neural signed distance function (SDF). $f : \mathbb{R}^3 \rightarrow \mathbb{R}$. SDFs provide a helpful inductive bias for learning smooth surface geometry [30, 42, 45] that assists downstream tasks in our pipeline. Moreover, the surface properties crucial for our framework, surface normals and curvature, can be conveniently computed from SDFs in a differentiable manner. Consider the 3D spatial coordinates, \mathbf{x} , in the scene. The glossy object surface, \mathcal{S} is then represented by the zero-level set of the SDF

$$\mathcal{S} = \{f_{\mathcal{S}}(\mathbf{x}) = 0 | \mathbf{x} \in \mathbb{R}^3\} \quad (1)$$




Figure 6. Overview of our proposed architecture.

Similar to Yariv *et al.* [45], we model the SDF f_S as a coordinate-based MLP.

Surface Normals Gradients of the SDF at the zero level set point \mathcal{S} towards the surface normals \mathcal{S} ,

$$\mathbf{n}(\mathbf{x}) = \frac{\nabla_{\mathbf{x}} f_S(\mathbf{x})}{\|\nabla_{\mathbf{x}} f_S(\mathbf{x})\|} \quad \mathbf{x} \in \mathcal{S} \quad (2)$$

Surface Curvature We employ differential geometry techniques developed by Novello *et al.* [29] to estimate curvature for neural implicit surfaces. In particular, we estimate the mean curvature $K(\mathbf{x})$ for the implicit surface from the divergence, ∇ of the surface normals

$$K(\mathbf{x}) = \frac{\nabla \cdot \mathbf{n}(\mathbf{x})}{2} \quad (3)$$

Mean curvature approximates the surface with an osculating sphere. Our approach also works for more generalized notions of curvature through the shape operator, at the cost of higher computational complexity. We refer our readers to the supplement for the general case.

Diffuse Radiance We separate the captured radiance at the observer camera with diffuse radiance, which depends on the glossy object’s albedo, and specular radiance which depends on the environment radiance. The diffuse radiance does not have any view dependence and only depends on surface point \mathbf{x} . We denote the diffuse radiance as f_d and model it using a coordinate-based MLP (Fig. 6).

Volume Rendering As proposed in [45], we perform volumetric rendering on the SDF. We define the volume density $\sigma(\mathbf{x})$ as the cumulative distribution function (CDF), denoted as $\Psi(s)$, applied to f_S :

$$\sigma(\mathbf{x}) = \alpha \Psi_{\beta}(f_S) \quad (4)$$

In contrast to [45], however, we only aim to recover the diffuse radiance of the object along a particular ray. We define a function f_d that estimates the diffuse radiance at each point, \mathbf{x} , along the ray. To get the final diffuse radiance along a given primary ray, $\mathbf{r}_p(t)$, we perform volumetric rendering:

$$\hat{c}_d(\mathbf{r}) = \int_0^{\infty} f_d(\mathbf{r}(t), f_S^k(\mathbf{r}(t))) \tau(t) dt \quad (5)$$

Note that there is no view dependence in Eq. 5 and intermediate features, f_S^k , are used as input. $\tau(t)$ is the accumulated transmittance along the ray.




Figure 7. Qualitative comparisons of diffuse-specular separation and geometry estimation on rendered dataset. The environment contains nearby objects with complex occlusions when seen through reflections on the glossy object. RefNeRF fails to perform accurate diffuse-specular separation and PANDORA blurs the nearby objects in the specular map. ORCA can model the complex specular reflections through environment radiance field.

3.3. Objects Surface as Virtual Sensor

Each pixel, \mathbf{p} , with a finite surface area, $d\mathbf{A}_p$, on the real-camera sensor views the surface of the object through a frustum originating at that pixel. The object then samples the environment radiance field through this finite surface converting the finite surface into a virtual pixel with surface area, $d\mathbf{S}$. Through this model, we can interpret the object surface as a virtual sensor consisting of many virtual pixels that sample radiance from the environment field based on the geometry of the object and observer viewing direction. We now formulate a virtual pixel based on real camera pose and implicit surface geometry. Please refer to Fig. 5 for a visualization of the virtual sensor.

Consider a real camera origin as \mathbf{o} and a pixel on the real sensor $\mathbf{p}_{i,j}$ that corresponds to ray direction \mathbf{d} . The primary ray for pixel $\mathbf{p}_{i,j}$ is parameterized with ray length t as $\mathbf{r}_p(t) = \mathbf{o} + t\mathbf{d}$

Casting Real Cones We can approximate the outgoing conical frustum from pixel $\mathbf{p}_{i,j}$ as a cone originating at \mathbf{o} with axis-of-direction \mathbf{d} and radius \dot{r} , equivalent to half the distance of the pixel in the x and y directions. We represent the real cone as a parametric volume,

$$\mathbf{r}_{\text{cone}}(\dot{r}, s, \theta) = \dot{r}s \cos(\theta) \hat{\mathbf{e}}_u + \dot{r}s \sin(\theta) \hat{\mathbf{e}}_v + \dot{r}s \mathbf{d}, \quad (6)$$

where $\hat{\mathbf{e}}_u$ and $\hat{\mathbf{e}}_v$ are basis vectors in the plane perpendicular

to \mathbf{d} , $\theta \in [0, \pi]$ and $s \in [0, t_{max}]$

Virtual Pixel. Virtual pixels are characterized by the intersection of the real cone with the object surface. In Sec. 3.2, we model local surface properties using mean curvature which enables efficient analytical computations for the virtual pixel parameters even though our approach works with general shape operators. For a sampled point t_i along the ray, we have the surface normals $\mathbf{n}(t_i)$ from Eq. 2 and estimated mean curvature $K(t_i)$ from Eq. 3. The local object surface at t_i , can be approximated with an osculating sphere, $\mathcal{O}(t_i)$, centered at $\hat{\mathbf{o}}_{\mathbf{S}}(t_i)$ with radius, $R(t_i)$ as

$$R(t_i) = \frac{2}{K_{t_i}} \hat{\mathbf{o}}_{\mathbf{S}}(t_i) = \mathbf{r}_p(t_i) + R(t_i) \cdot \hat{\mathbf{n}}(t_i)$$

For concave surfaces $K_{t_i} < 0$. So, $\hat{\mathbf{o}}_{\mathbf{S}}$ will lie outside the object. For $K_{t_i} > 0$, $\hat{\mathbf{o}}_{\mathbf{S}}$ will lie inside the object.

The edges of the virtual pixel for $\mathbf{r}_p(t_i)$ would lie at the intersection of the osculating sphere $\mathcal{O}(t_i)$ and the primary cone given by \mathbf{r}_{cone} . Computing exact cone-sphere intersections are computationally expensive so we approximate the cone-sphere intersection using rays bound cone-sphere intersectional surface \mathbf{dS} . We consider four rays that bound the cone and sample them at $\theta_j \in \{0, \pi/2, \pi, 3\pi/2\}$ with Eq. 6. We perform intersections of the corresponding bounding rays with the osculating sphere $\mathcal{O}(t_i)$ to get corners of the virtual pixel \mathbf{ds}_j . These ray sphere intersections can be computed analytically in an efficient manner.

Virtual Cone Origin. From the virtual pixel surface area, we can now compute the virtual cone that samples the environment. We first compute normal vectors at virtual pixel corners \mathbf{ds}_j from the center of osculating sphere $\hat{\mathbf{o}}_{\mathbf{S}}$

$$\hat{\mathbf{n}}_j = \frac{\mathbf{ds}_j - \hat{\mathbf{o}}_{\mathbf{S}}}{\|\mathbf{ds}_j - \hat{\mathbf{o}}_{\mathbf{S}}\|} \quad (7)$$

At each virtual pixel corner, we compute the reflected ray directions, ω_j^r , by computing the dot product between the incoming ray directions, ω_k^i , and the normals, $\hat{\mathbf{n}}_k$, where ω_0^r is the primary ray's reflected vector.

$$\omega_0^r = \mathbf{d} - (\mathbf{d} \cdot \hat{\mathbf{n}}(t_i)) \hat{\mathbf{n}}(t_i) \quad (8)$$

$$\omega_j^r = \mathbf{d}_j - (\mathbf{d}_j \cdot \hat{\mathbf{n}}_j(k)) \hat{\mathbf{n}}_j(k) \quad (9)$$

\mathbf{d}_j are the incident directions to the virtual pixel corners \mathbf{ds}_j . The virtual cone origin is the intersection of these reflected rays at the pixel corners and pixel center. However, these rays might not intersect at a single point so we approximate a virtual origin to be the point that minimizes the sum of distances to the reflected rays ω_j .

$$\mathbf{v}_o = \operatorname{argmin}_{\mathbf{v}} \sum_j |(\mathbf{v} - \mathbf{ds}_j) \times \omega_j^r| \quad (10)$$

We pose this as a linear least squares problem and estimate the virtual cone origin efficiently through pseudo-inverse.

Virtual Cones Direction. The reflected ray at the center of the virtual pixel reflects the object surface along the direction ω_0^r from Eq. 8. We consider this as the direction-of-axis of the virtual cone.

$$\hat{\mathbf{v}}_d = \omega_0^r \quad (11)$$

Virtual Cone Radius. We compute the radius of the cone by treating the reflection vectors of the bounding rays as the neighboring ‘‘pixel’’ directions. Similar to [5], we can compute the distance between $\{\omega_{k\theta}^r\}_{\theta=0}^{2\pi}$ and the primary reflected ray ω_0^r in the (x, y) components (omitted below for clarity).

$$\hat{\mathbf{v}}_r = \|\{\omega_{k\theta}^r\}_{\theta=0}^{2\pi} - \omega_0^r\| \quad (12)$$

Finally, for each sampled point t_i , we can characterize our single-pixel virtual sensor located at the object surface \mathbf{dS} as a virtual cone with $\hat{\mathbf{v}}_o$ as its apex, $\hat{\mathbf{v}}_d$ as axis-direction, $\hat{\mathbf{v}}_r$ as the radius.

Connections to caustics. Our work takes inspiration from catadioptric imaging systems. To convert objects into cameras, we compute the surface and find a corresponding center-of-projection for this surface-as-sensor. However, unlike conventional perspective cameras, objects don't have a fixed center-of-projection, other than in a few special configurations [3], but a locus of viewpoints that vary with object geometry and viewing direction. These viewpoints lie on the ‘‘caustic surface’’ of the object. While typical works in catadioptric imaging analytically compute the caustic surface by assuming known geometry [14, 38], or making assumptions about placement of the observer [40], our formulation approximates the caustic surface of unknown geometry through the intersection of reflected rays on virtual pixels. We empirically show in the supplement that as the surface area of the virtual pixel goes to 0, $d\mathbf{S} \rightarrow 0$, our method estimates the true caustic of the object without assuming geometry. Our method also has applications in estimating the caustic surface of the unknown geometry.

3.4. Environment Radiance Fields

Our goal is to capture a 5D environment radiance field of the scene by imaging the world through these single-pixel virtual sensors located at the object's surface. We use our formulation of virtual cones to recover 5D environment radiance fields. We define an environment radiance field as $f_{\mathcal{E}} : (\hat{\mathbf{v}}_o, \hat{\mathbf{v}}_d) \rightarrow (\sigma^{Env}, c_s)$,

where $f_{\mathcal{E}}$ outputs opacity and radiance along sampled virtual cones. We note that this view-dependent radiance is equivalent to the specular radiance at point t_i sampled




Figure 8. **Evaluation of environment depth from reflections.** For the scene in Fig. 3, ORCa can faithfully estimate the depth of the environment, especially for nearby surroundings.

Approach	Diffuse Radiance		Specular Radiance		Mixed Radiance		Normals MAE ↓ (°)
	PSNR ↑ (dB)	SSIM ↑	PSNR ↑ (dB)	SSIM ↑	PSNR ↑ (dB)	SSIM ↑	
Ref-NeRF	18.80	0.7304	16.99	0.6633	20.19	0.7890	43.690
PANDORA	18.25	0.7260	16.25	0.6483	18.90	0.7284	7.606
ORCA	19.84	0.7893	20.74	0.7535	22.00	0.7947	2.339

Table 1. **Average evaluation metrics on rendered scenes.** We compare ORCa to other neural rendering techniques that model reflections, including Ref-NeRF and PANDORA, on six simulated scenes. ORCa provides consistent improvements in geometry estimation, diffuse-specular separation and novel view synthesis. Please refer to the supplement for additional metrics.

along the primary-camera ray $\mathbf{r}_p(t)$. We can render the final specular radiance at pixel $\mathbf{p}_{i,j}$ as follows:

$$\hat{\mathbf{c}}_s(\mathbf{r}) = \int_0^\infty f_{\mathcal{E}}(\hat{\mathbf{v}}_o, \hat{\mathbf{v}}_d) \tau(t) dt$$

$$\hat{\mathbf{c}} = \hat{\mathbf{c}}_d + \hat{\mathbf{c}}_s$$

Intuitively, $f_{\mathcal{E}}$ learns the 5D radiance field by sampling single-pixel virtual sensors from the object surface and must learn geometry and environment radiance that is consistent with the multi-view reflections. Moreover, we can query $f_{\mathcal{E}}$ to render novel viewpoints and associated depths that are beyond the field-of-view of the real camera. We volume render each virtual cone by dividing them into conical frustums using Integrated-Positional Encoding as proposed in MipNeRF [5]. Our formulation of virtual cones works well with Mip-Nerf’s rays-as-cones method.

4. Experiments

4.1. Implementation Details

As in PANDORA, we parameterize f_S with an 8-layer MLP to estimate the surface, and, as in MipNeRF, f_d with 4-layer MLP with input geometric features of size 512 from f_S . We follow the SDF-to-opacity conversion and the iterative sampling of the ray proposed in [45]. To aid the network to learn geometry quickly, we also train f_S with a mask-net as proposed in [10]. We use five losses in our architecture: photometric loss, mask loss [10], normal loss [41], eikonal loss [45], and distortion loss [6]. Additional training details are discussed in the supplement.

Scene: Ball-cup in hallway			Scene: Owl in hallway		
Approach	Mixed Radiance		Approach	Mixed Radiance	
	PSNR ↑ (dB)	SSIM ↑		PSNR ↑ (dB)	SSIM ↑
Ref-NeRF	32.75	0.9617	Ref-NeRF	26.65	0.8890
PANDORA	28.83	0.9758	PANDORA	27.24	0.9343
ORCA	30.86	0.9799	ORCA	26.84	0.9299

Table 2. **Quantitative metrics for real-world, captured scenes.** ORCa demonstrates comparable performance in novel view synthesis on scenes from the PANDORA real dataset.

4.2. Datasets

We conduct experiments on both simulated and real-world datasets. Simulated datasets are rendered in Mitsuba2 [28]. Simulated datasets contain a range of increasingly complex object geometries (elephant, Pokeball, and orca) and scenes (living room and Pokemon). We train with 200 views for simulated datasets. We also show results for a real-world dataset [10] capturing a glossy cup with a black vase sitting atop it using 35 views. All datasets will be publicly released upon publication.

4.3. Advantages of Environment Radiance Fields

Other neural rendering techniques that handle reflections, Ref-NeRF and PANDORA, estimate the environment as a 2D map, while ORCa recovers a 5D environment radiance field. Fig. 3 shows the advantages of recovering a 5D environment radiance field. Close-by surrounding objects often cause occlusions that cannot be modeled by 2D environment maps. By estimating the radiance field, we can image behind occluders through sampling novel viewpoints such as the translated viewpoints shown in Fig. 3. Moreover, we can also show the depth to the surroundings from these virtual viewpoints. In Fig. 8, we quantitatively evaluate the estimated depth of the surroundings for the synthesized virtual views in Fig. 3 and show that ORCa provides reliable depth estimates especially for nearby objects. We provide additional examples of depth estimation and beyond *field-of-view* novel-view synthesis in the supplement.

4.4. Comparisons with Baselines

We compare ORCa with Ref-NeRF and PANDORA on synthetic and real datasets. As shown in Fig. 7 and 9, ORCa estimates more accurate surface normals than other methods. While the total radiance predicted by Ref-NeRF and PANDORA are visually similar, the surface normals are less smooth than ORCa. We also observe that ORCa is able to achieve better diffuse and specular radiance separation than PANDORA, which is evident in the Pokeball surface normals Fig. 7. In these examples, PANDORA recovers blurry specular radiance. We see that ORCa’s predicted depth is highly interpretable and matches the underlying geometry of the environment, as shown in Fig. 3. Even on cylindri-

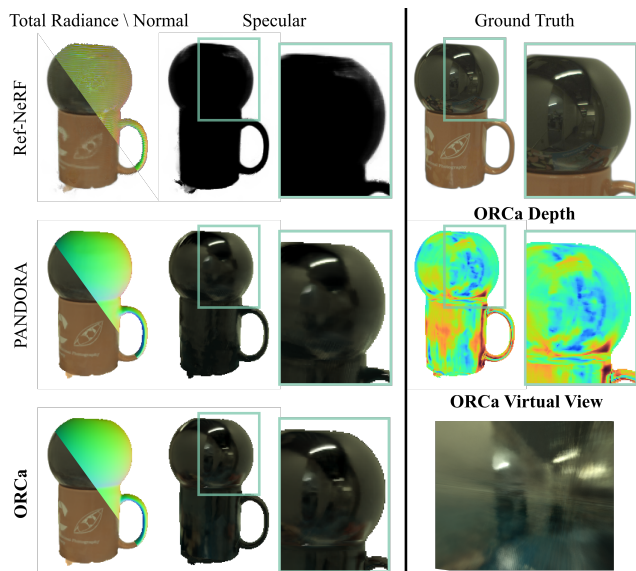


Figure 9. **Comparisons on real dataset.** Using a real-world dataset with only 35 views, ORCa can model the sharp specularities on the ball arriving from regions of the nearby scene, such as the table, and far away scene regions, such as the hallway, to learn an environment radiance field. We query the radiance field for the depth of the far hallway (blue) and the nearby objects, such as the table (red). We also render novel viewpoints that are beyond the field of view of the observer camera and show that ORCa interpolates well between those views.

cal real-world datasets, such as the black vase in Fig. 9, the nearby hallway is visible in both the virtual view and depth, despite never being in the field of view of the primary camera. Unlike Ref-NeRF, our primary objective is not to perform novel-view synthesis, but instead to capture the environment radiance field from the object surface. As shown in Table 1, ORCa improves over Ref-NeRF and PANDORA in estimating diffuse radiance, specular radiance, mixed radiance, and normals. Here we report the image quality metrics, PSNR and SSIM, and surface normal metric, MAE, averaged over six synthetic scenes with varying environments, camera view distribution, and object complexity. Please refer to the supplement for individual metrics on each scene. ORCa also shows competitive performance for novel-view synthesis metrics on real-world scenes from PANDORA dataset, as shown in Table 2.

4.5. Impact of Virtual Cone Computation

We base our method on a physically accurate formulation by modeling ray-cone intersections and using the surface as a virtual sensor, as described in Sec. 3.3. We demonstrate the importance of this step by setting up two ablation experiments. In the first experiment, which we term Naive Virtual Cone, we place the origin of the cone at the object surface instead of computing the virtual cone origin based

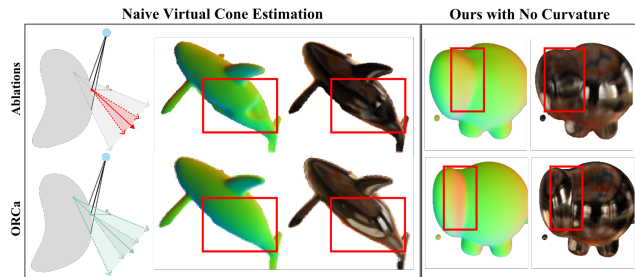


Figure 10. **Ablations on virtual cone computation** Accurate estimation of the virtual cones (Sec. 3.3) is crucial for ORCa. If the virtual cone origin is assumed to be at the object surface (left column) or if the surface is assumed to locally have no curvature (right column), the surface normal and specular radiance outputs suffer from artifacts (red boxes).

on the curvature. In the second experiment, which we term No Curvature, we assume that locally the surface is like a flat mirror and has no curvature. In Fig. 10, we show that for both of these ablation experiments, we see worse performance as demonstrated by the artifacts in estimated surface normals and specular components (shown as red boxes).

5. Conclusion

We present a method to convert glossy objects with unknown geometry and texture into radiance-field cameras that capture the environment’s radiance field around them. Our method recovers object geometry and diffuse radiance, in addition to capturing the depth and radiance of the object’s surroundings from its perspective. Our modeling of the environment as a radiance field is effective in recovering close-by objects (Fig. 7) and is occlusion aware (Fig. 3). From the recovered environment radiance field, we can perform beyond *field-of-view* novel-view synthesis. Our work can unleash applications in virtual object insertion and 3D perception, e.g. inferring information beyond the line-of-sight of the camera using predicted virtual views and depth. Our formulation goes beyond the conventional direct-line-of-sight radiance fields and can enable further areas of research to extract more information from multi-view images directly from the environment and the objects present in it.

Acknowledgments

This research was supported by the SMART Contract IARPA Grant 2021-20111000004. We would also like to thank Systems Technology Research (STR). AD and AV are supported by NSF Expeditions-Award IIS-1730574 and NSF CAREER-Award IIS-1652633.

References

- [1] Anna Alperovich, Ole Johannsen, and Bastian Goldluecke. Intrinsic light field decomposition and disparity estimation with deep encoder-decoder network. In *2018 26th European Signal Processing Conference (EUSIPCO)*, pages 2165–2169. IEEE, 2018. 3
- [2] S. Baker and S.K. Nayar. A theory of catadioptric image formation. In *Sixth International Conference on Computer Vision (IEEE Cat. No.98CH36271)*, pages 35–42, 1998. 3
- [3] Simon Baker and Shree K. Nayar. A theory of single-viewpoint catadioptric image formation. *International Journal of Computer Vision*, 35:175–196, 2004. 6
- [4] Jonathan T Barron and Jitendra Malik. Shape, illumination, and reflectance from shading. *IEEE transactions on pattern analysis and machine intelligence*, 37(8):1670–1687, 2014. 3
- [5] 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. *ICCV*, 2021. 3, 6, 7
- [6] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5470–5479, 2022. 7
- [7] Mark Boss, Raphael Braun, Varun Jampani, Jonathan T Barron, Ce Liu, and Hendrik Lensch. Nerf: Neural reflectance decomposition from image collections. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12684–12694, 2021. 4
- [8] Mark Boss, Varun Jampani, Raphael Braun, Ce Liu, Jonathan Barron, and Hendrik Lensch. Neural-pil: Neural pre-integrated lighting for reflectance decomposition. *Advances in Neural Information Processing Systems*, 34:10691–10704, 2021. 4
- [9] Akshat Dave, Yannick Hold-Geoffroy, Miloš Hašan, Kalyan Sunkavalli, and Ashok Veeraraghavan. Snapshot polarimetric diffuse-specular separation. *Optics Express*, 30(19):34239–34255, 2022. 3
- [10] Akshat Dave, Yongyi Zhao, and Ashok Veeraraghavan. Pandora: Polarization-aided neural decomposition of radiance. *arXiv preprint arXiv:2203.13458*, 2022. 3, 4, 7
- [11] Michael De Zeeuw and Aswin C Sankaranarayanan. Wide-baseline light fields using ellipsoidal mirrors. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. 3
- [12] Mathieu Garon, Kalyan Sunkavalli, Sunil Hadap, Nathan Carr, and Jean-François Lalonde. Fast spatially-varying indoor lighting estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6908–6917, 2019. 3
- [13] Stamatios Georgoulis, Konstantinos Rematas, Tobias Ritschel, Mario Fritz, Tinne Tuytelaars, and Luc Van Gool. What is around the camera? In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5170–5178, 2017. 3
- [14] J. Gluckman and S.K. Nayar. Planar catadioptric stereo: geometry and calibration. In *Proceedings. 1999 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (Cat. No PR00149)*, volume 1, pages 22–28 Vol. 1, 1999. 6
- [15] Michael D Grossberg and Shree K Nayar. The raxel imaging model and ray-based calibration. *International Journal of Computer Vision*, 61(2):119, 2005. 3
- [16] Yuan-Chen Guo, Di Kang, Linchao Bao, Yu He, and Song-Hai Zhang. Nerfren: Neural radiance fields with reflections. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18409–18418, June 2022. 3
- [17] Connor Henley, Siddharth Somasundaram, Joseph Hollmann, and Ramesh Raskar. Detection and mapping of specular surfaces using multibounce lidar returns. *Optics Express*, 31(4):6370–6388, 2023. 3
- [18] Jan Jachnik, Richard A Newcombe, and Andrew J Davison. Real-time surface light-field capture for augmentation of planar specular surfaces. In *2012 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pages 91–97. IEEE, 2012. 3
- [19] Georgios Kopanas, Thomas Leimkühler, Gilles Rainer, Clément Jambon, and George Drettakis. Neural point caustics for novel-view synthesis of reflections. *ACM Transactions on Graphics*, 41(6):Article–201, 2022. 3
- [20] Rui Li, Simeng Qiu, Guangming Zang, and Wolfgang Heidrich. Reflection separation via multi-bounce polarization state tracing. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIII 16*, pages 781–796. Springer, 2020. 3
- [21] Tingtian Li, Daniel PK Lun, Yuk-Hee Chan, et al. Robust reflection removal based on light field imaging. *IEEE Transactions on Image Processing*, 28(4):1798–1812, 2018. 3
- [22] Zhengqin Li, Mohammad Shafiei, Ravi Ramamoorthi, Kalyan Sunkavalli, and Manmohan Chandraker. Inverse rendering for complex indoor scenes: Shape, spatially-varying lighting and svbrdf from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2475–2484, 2020. 3
- [23] Stephen Lombardi and Ko Nishino. Reflectance and natural illumination from a single image. In *European Conference on Computer Vision*, pages 582–595. Springer, 2012. 3
- [24] 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. *Communications of the ACM*, 65(1):99–106, 2021. 3
- [25] Jacob Munkberg, Jon Hasselgren, Tianchang Shen, Jun Gao, Wenzheng Chen, Alex Evans, Thomas Müller, and Sanja Fidler. Extracting triangular 3d models, materials, and lighting from images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8280–8290, 2022. 4
- [26] Shree K Nayar and Simon Baker. Catadioptric image formation. In *Proceedings of the 1997 DARPA Image Understanding Workshop*, pages 1431–1437, 1997. 3

- [27] Yun Ni, Jie Chen, and Lap-Pui Chau. Reflection removal on single light field capture using focus manipulation. *IEEE Transactions on Computational Imaging*, 4(4):562–572, 2018. [3](#)
- [28] Merlin Nimier-David, Delio Vicini, Tizian Zeltner, and Wenzel Jakob. Mitsuba 2: A retargetable forward and inverse renderer. *ACM Transactions on Graphics (TOG)*, 38(6):1–17, 2019. [7](#)
- [29] Tiago Novello, Guilherme Schardong, Luiz Schirmer, Vinicius da Silva, Helio Lopes, and Luiz Velho. Exploring differential geometry in neural implicit, 2022. [5](#)
- [30] Michael Oechsle, Songyou Peng, and Andreas Geiger. Unisurf: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5589–5599, 2021. [4](#)
- [31] Jeong Joon Park, Aleksander Holynski, and Steven M Seitz. Seeing the world in a bag of chips. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1417–1427, 2020. [4](#)
- [32] Ravi Ramamoorthi and Pat Hanrahan. A signal-processing framework for inverse rendering. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 117–128, 2001. [3](#)
- [33] Raskar Ramesh and James Davis. 5d time-light transport matrix: What can we reason about scene properties? Technical report, 2008. [3](#)
- [34] Fabiano Romeiro and Todd Zickler. Blind reflectometry. In *European conference on computer vision*, pages 45–58. Springer, 2010. [3](#)
- [35] Vincent Sitzmann, Semon Rezhikov, Bill Freeman, Josh Tenenbaum, and Fredo Durand. Light field networks: Neural scene representations with single-evaluation rendering. *Advances in Neural Information Processing Systems*, 34:19313–19325, 2021. [3](#)
- [36] Shuran Song and Thomas Funkhouser. Neural illumination: Lighting prediction for indoor environments. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6918–6926, 2019. [4](#)
- [37] Pratul P Srinivasan, Ben Mildenhall, Matthew Tancik, Jonathan T Barron, Richard Tucker, and Noah Snavely. Lighthouse: Predicting lighting volumes for spatially-coherent illumination. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8080–8089, 2020. [4](#)
- [38] R. Swaminathan, M.D. Grossberg, and S.K. Nayar. Caustics of catadioptric cameras. In *Proceedings Eighth IEEE International Conference on Computer Vision. ICCV 2001*, volume 2, pages 2–9 vol.2, 2001. [6](#)
- [39] Tristan Swedish, Connor Henley, and Ramesh Raskar. Objects as cameras: Estimating high-frequency illumination from shadows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2593–2602, 2021. [4](#)
- [40] Yuichi Taguchi, Amit Agrawal, Ashok Veeraraghavan, Srikumar Ramalingam, and Ramesh Raskar. Axial-cones: Modeling spherical catadioptric cameras for wide-angle light field rendering. *ACM Trans. Graph.*, 29(6), dec 2010. [3](#), [6](#)
- [41] Dor Verbin, Peter Hedman, Ben Mildenhall, Todd Zickler, Jonathan T. Barron, and Pratul P. Srinivasan. Ref-NeRF: Structured view-dependent appearance for neural radiance fields. *CVPR*, 2022. [3](#), [4](#), [7](#)
- [42] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *arXiv preprint arXiv:2106.10689*, 2021. [4](#)
- [43] Zian Wang, Jonah Philion, Sanja Fidler, and Jan Kautz. Learning indoor inverse rendering with 3d spatially-varying lighting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12538–12547, 2021. [3](#)
- [44] Ziyu Wang, Liao Wang, Fuqiang Zhao, Minye Wu, Lan Xu, and Jingyi Yu. Mirrornerf: One-shot neural portrait radiance field from multi-mirror catadioptric imaging. In *2021 IEEE International Conference on Computational Photography (ICCP)*, pages 1–12. IEEE, 2021. [3](#)
- [45] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. In *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021. [4](#), [5](#), [7](#)
- [46] Kai Zhang, Fujun Luan, Qianqian Wang, Kavita Bala, and Noah Snavely. Physg: Inverse rendering with spherical gaussians for physics-based material editing and relighting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5453–5462, 2021. [4](#)
- [47] 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 Transactions on Graphics (TOG)*, 40(6):1–18, 2021. [4](#)
- [48] Rui Zhu, Zhengqin Li, Janarbek Matai, Fatih Porikli, and Manmohan Chandraker. Irisformer: Dense vision transformers for single-image inverse rendering in indoor scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2822–2831, 2022. [3](#)