

# Seeing the World in a Bag of Chips


Jeong Joon Park

Aleksander Holynski

Steven M. Seitz

University of Washington, Seattle

{jjpark7,holynski,seitz}@cs.washington.edu



(a) Input images

(b) Estimated environment (top), ground truth (bottom)

(c) Zoom-in

Figure 1: From a hand-held RGBD sequence of an object (a), we reconstruct an image of the surrounding environment (b, top) that closely resembles the real environment (b, bottom), entirely from the specular reflections. Note the reconstruction of fine details (c) such as a human figure and trees with fall colors through the window. We use the recovered environment for novel view rendering.<sup>†</sup>

## Abstract

We address the dual problems of novel view synthesis and environment reconstruction from hand-held RGBD sensors. Our contributions include 1) modeling highly specular objects, 2) modeling inter-reflections and Fresnel effects, and 3) enabling surface light field reconstruction with the same input needed to reconstruct shape alone. In cases where scene surface has a strong mirror-like material component, we generate highly detailed environment images, revealing room composition, objects, people, buildings, and trees visible through windows. Our approach yields state of the art view synthesis techniques, operates on low dynamic range imagery, and is robust to geometric and calibration errors.

## 1. Introduction

The glint of light off an object reveals much about its shape and composition – whether it’s wet or dry, rough or polished, round or flat. Yet, hidden in the pattern of highlights is also an image of the environment, often so distorted that we don’t even realize it’s there. Remarkably, images of the shiny bag of chips (Fig. 1) contain sufficient clues to be able to reconstruct a detailed image of the room, including the layout of lights, windows, and even objects outside that are visible through windows.

In their *visual microphone* work, Davis *et al.* [13] showed how sound and even conversations can be reconstructed from the minute vibrations visible in a bag of chips. Inspired by their work, we show that the same bag of chips can be used to reconstruct the environment. Instead of high speed video, however, we operate on RGBD video, as obtained with commodity depth sensors.

Visualizing the environment is closely connected to the problem of modeling the scene that reflects that environment. We solve both problems; beyond visualizing the room, we seek to predict how the objects and scene appear from any new viewpoint — i.e., to virtually explore the scene as if you were there. This view synthesis problem has a large literature in computer vision and graphics, but several open problems remain. Chief among them are 1) specular surfaces, 2) inter-reflections, and 3) simple capture. In this paper we address all three of these problems, based on the framework of *surface light fields* [68].

Our environment reconstructions, which we call *specular reflectance maps (SRMs)*, represent the distant environment map convolved with the object’s specular BRDF. In cases where the object has strong mirror-like reflections, this SRM provides sharp, detailed features like the one seen in Fig. 1. As most scenes are composed of a mixture of materials, each scene has multiple basis SRMs. We therefore reconstruct a global set of SRMs, together with a weighted

<sup>†</sup>Video URL: [https://youtu.be/9t\\_Rx6n1HGA](https://youtu.be/9t_Rx6n1HGA)

material segmentation of scene surfaces. Based on the recovered SRMs, together with additional physically motivated components, we build a neural rendering network capable of faithfully approximating the true surface light field.

A major contribution of our approach is the capability of reconstructing a surface light field with the same input needed to compute shape alone [52] using an RGBD camera. Additional contributions of our approach include the ability to operate on regular (low-dynamic range) imagery, and applicability to general, non-convex, textured scenes containing multiple objects and both diffuse and specular materials. Lastly, we release RGBD dataset capturing reflective objects to facilitate research on lighting estimation and image-based rendering.

We point out that the ability to reconstruct the reflected scene from images of an object opens up real and valid concerns about privacy. While our method requires a depth sensor, future research may lead to methods that operate on regular photos. In addition to educating people on what’s possible, our work could facilitate research on privacy-preserving cameras and security techniques that actively identify and scramble reflections.

## 2. Related Work

We review related work in environment lighting estimation and novel-view synthesis approaches for modeling specular surfaces.

### 2.1. Environment Estimation

**Single-View Estimation** The most straightforward way to capture an environment map (image) is via light probes (e.g., a mirrored ball [15]) or taking photos with a 360° camera [56]. Human eye balls [54] can even serve as light probes when they are present. For many applications, however, light probes are not available and we must rely on existing cues in the scene itself.

Other methods instead study recovering lighting from a photo of a general scene. Because this problem is severely under-constrained, these methods often rely on human inputs [34, 77] or manually designed “intrinsic image” priors on illumination, material, and surface properties [35, 6, 5, 7, 44].

Recent developments in deep learning techniques facilitate data-driven approaches for single view estimation. [19, 18, 63, 40] learn a mapping from a perspective image to a wider-angle panoramic image. Other methods train models specifically tailored for outdoor scenes [29, 28]. Because the single-view problem is severely ill-posed, most results are plausible but often non-veridical. Closely related to our work, Georgoulis *et al.* [20] reconstruct higher quality environment images, but under very limiting assumptions; textureless painted surfaces and manual specification of materials and segmentation.

**Multi-View Estimation** For the special case of planar reflectors, layer separation techniques [65, 62, 73, 25, 24, 31, 76] enable high quality reconstructions of reflected environments, e.g., from video of a glass picture frame. Inferring reflections for general, curved surfaces is dramatically harder, even for humans, as the reflected content depends strongly and nonlinearly on surface shape and spatially-varying material properties,

A number of researchers have sought to recover *low-frequency* lighting from multiple images of curved objects. [81, 55, 46] infer spherical harmonics lighting (following [59]) to refine the surface geometry using principles of shape-from-shading. [60] jointly optimizes low frequency lighting and BRDFs of a reconstructed scene. While suitable for approximating light source directions, these models don’t capture detailed images of the environment.

Wu *et al.* [69], like us, use a hand-held RGBD sensor to recover lighting and reflectance properties. But the method can only reconstruct a single, floating, convex object, and requires a black background. Dong *et al.* [16] produces high quality environment images from a video of a single rotating object. This method assumes a laboratory setup with a mechanical rotator, and manual registration of an accurate geometry to their video. Similarly, Xia *et al.* [70] use a robotic arm with calibration patterns to rotate an object. The authors note highly specular surfaces cause trouble, thus limiting their real object samples to mostly rough, glossy materials. In contrast, our method operates with a hand-held camera for a wide-range of multi-object scenes, and is designed to support specularly.

### 2.2. Novel View Synthesis

Here we focus on methods capable of modeling *specular reflections* from new viewpoints.

**Image-based Rendering** Light field methods [23, 42, 10, 68, 12] enable highly realistic views of specular surfaces at the expense of laborious scene capture from densely sampled viewpoints. Chen *et al.* [8] regresses surface light fields with neural networks to reduce the number of required views, but requires samples across a full hemisphere captured with a mechanical system. Park *et al.* [56] avoid dense hemispherical view sampling by applying a parametric BRDF model, but assume known lighting.

Recent work applies convolutional neural networks (CNN) to image-based rendering [17, 49]. Hedman *et al.* [27] replaced the traditional view blending heuristics of IBR systems with a CNN-learned blending weights. Still, novel views are composed of existing, captured pixels, so unobserved specular highlights cannot be synthesized. More recently, [2, 66] enhance the traditional rendering pipeline by attaching learned features to 2D texture maps [66] or 3D point clouds [2] and achieve high quality view synthesis results. The features are nonetheless specifically optimized

to fit the input views and do not extrapolate well to novel views. Recent learning-based methods achieve impressive local (versus hemispherical) light field reconstruction from a small set of images [50, 64, 11, 33, 78].

**BRDF Estimation Methods** Another way to synthesize novel views is to recover intrinsic surface reflection functions, known as BRDFs [53]. In general, recovering the surface BRDFs is a difficult task, as it involves inverting the complex light transport process. Consequently, existing reflectance capture methods place limits on operating range: e.g., an isolated single object [69, 16], known or controlled lighting [56, 14, 41, 79, 72], single view surface (versus a full 3D mesh) [21, 43], flash photography [1, 39, 51], or spatially constant material [48, 37].

**Interreflections** Very few view synthesis techniques support interreflections. Modeling general multi-object scene requires solving for global illumination (e.g. shadows or interreflections), which is difficult and sensitive to imperfections of real-world inputs [4]. Similarly, Lombardi *et al.* [45] model multi-bounce lighting but with noticeable artifacts and limit their results to mostly uniformly textured objects. Zhang *et al.* [74] require manual annotations of light types and locations.

### 3. Technical Approach

Our system takes a video and 3D mesh of a static scene (obtained via Newcombe *et al.* [52]) as input and automatically reconstructs an image of the environment along with a scene appearance model that enables novel view synthesis. Our approach excels at specular scenes, and accounts for both specular interreflection and Fresnel effects. A key advantage of our approach is the use of easy, casual data capture from a hand-held camera; we reconstruct the environment map and a surface light field with the same input needed to reconstruct the geometry alone, e.g., using [52].

Section 3.1 formulates surface light fields [68] and define the specular reflectance map (SRM). Section 3.2 shows how, given geometry and diffuse texture as input, we can jointly recover SRMs and material segmentation through an end-to-end optimization approach. Lastly, Section 3.3, describes a *scene-specific* neural rendering network that combines recovered SRMs and other rendering components to synthesize realistic novel-view images, with interreflections and Fresnel effects.

#### 3.1. Surface Light Field Formulation

We model scene appearance using the concept of a surface light field [68], which defines the color radiance of a surface point in every view direction, given approximate geometry, denoted  $\mathcal{G}$  [52].

Formally, the surface light field, denoted  $SL$ , assigns an RGB radiance value to a ray coming from surface point  $\mathbf{x}$  with outgoing direction  $\boldsymbol{\omega}$ :  $SL(\mathbf{x}, \boldsymbol{\omega}) \in \text{RGB}$ . As is

common [57, 67], we decompose  $SL$  into diffuse (view-independent) and specular (view-dependent) components:

$$SL(\mathbf{x}, \boldsymbol{\omega}) \approx D(\mathbf{x}) + S(\mathbf{x}, \boldsymbol{\omega}). \quad (1)$$

We compute the diffuse texture  $D$  for each surface point as the minimum intensity of across different input views following [65, 56]. Because the diffuse component is view-independent, we can then render it from arbitrary viewpoints using the estimated geometry. However, textured 3D reconstructions typically contain errors (e.g., silhouettes are enlarged, as in Fig. 2), so we refine the rendered texture image using a neural network (Sec. 3.2).

For the specular component, we define the specular reflectance map (SRM) (also known as *lumisphere* [68]) and denoted  $SR$ , as a function that maps a reflection ray direction  $\boldsymbol{\omega}_r$ , defined as the vector reflection of  $\boldsymbol{\omega}$  about surface normal  $\mathbf{n}_x$  [68] to specular reflectance (i.e., radiance):  $SR(\boldsymbol{\omega}_r) : \Omega \mapsto \text{RGB}$ , where  $\Omega$  is a unit hemisphere around the scene center. This model assumes distant environment illumination, although we add support for specular inter-reflection later in Sec. 3.3. Note that this model is closely related to prefiltered environment maps [36], used for real-time rendering of specular highlights.

Given a specular reflectance map  $SR$ , we can render the specular image  $S$  from a virtual camera as follows:

$$S(\mathbf{x}, \boldsymbol{\omega}) = V(\mathbf{x}, \boldsymbol{\omega}_r; \mathcal{G}) \cdot SR(\boldsymbol{\omega}_r), \quad (2)$$

where  $V(\mathbf{x}, \boldsymbol{\omega}_r; \mathcal{G})$  is a shadow (visibility) term that is 0 when the reflected ray  $\boldsymbol{\omega}_r := \boldsymbol{\omega} - 2(\boldsymbol{\omega} \cdot \mathbf{n}_x)\mathbf{n}_x$  from  $\mathbf{x}$  intersects with known geometry  $\mathcal{G}$ , and 1 otherwise.

An SRM contains distant environment lighting convolved with a particular specular BRDF. As a result, a single SRM can only accurately describe one surface material. In order to generalize to multiple (and spatially varying) materials, we modify Eq. (2) by assuming the material at point  $\mathbf{x}$  is a linear combination of  $M$  basis materials [21, 3, 80]:

$$S(\mathbf{x}, \boldsymbol{\omega}) = V(\mathbf{x}, \boldsymbol{\omega}_r; \mathcal{G}) \cdot \sum_{i=1}^M W_i(\mathbf{x}) \cdot SR_i(\boldsymbol{\omega}_r), \quad (3)$$

where  $W_i(\mathbf{x}) \geq 0$ ,  $\sum_{i=1}^M W_i(\mathbf{x}) = 1$  and  $M$  is user-specified. For each surface point  $\mathbf{x}$ ,  $W_i(\mathbf{x})$  defines the weight of material basis  $i$ . We use a neural network to approximate these weights in image-space, as described next.

#### 3.2. Estimating SRMs and Material Segmentation

Given scene shape  $\mathcal{G}$  and photos from known viewpoints as input, we now describe how to recover an optimal set of SRMs and material weights.

Suppose we want to predict a view of the scene from camera  $P$  at a pixel  $\mathbf{u}$  that sees surface point  $\mathbf{x}_u$ , given known SRMs and material weights. We render the diffuse




Figure 2: The role of diffuse network  $u_\phi$  to correct geometry and texture errors of RGBD reconstruction. The bottle geometry in image (a) is estimated larger than it actually is, and the background textures exhibit ghosting artifacts (faces). The use of the refinement network corrects these issues (b). Best viewed digitally.

component  $D_P(\mathbf{u})$  from the known diffuse texture  $D(\mathbf{x}_\mathbf{u})$ , and similarly the blending weight map  $W_{P,i}$  from  $W_i$  for each SRM using standard rasterization. A reflection direction image  $R_P(\mathbf{u})$  is obtained by computing per-pixel  $\omega_r$  values. We then compute the specular component image  $S_P$  by looking up the reflected ray directions  $R_P$  in each SRM, and then combining the radiance values using  $W_{P,i}$ :

$$S_P(\mathbf{u}) = V(\mathbf{u}) \cdot \sum_{i=1}^M W_{P,i}(\mathbf{u}) \cdot SR_i(R_P(\mathbf{u})), \quad (4)$$

where  $V(\mathbf{u})$  is the visibility term of pixel  $\mathbf{u}$  as used in Eq. (3). Each  $SR_i$  is stored as a 2D panorama image of resolution 500 x 250 in spherical coordinates.

Now, suppose that SRMs and material weights are unknown; the optimal SRMs and combination weights minimize the energy  $\mathcal{E}$  defined as the sum of differences between the real photos  $G$  and the rendered composites of diffuse and specular images  $D_P, S_P$  over all input frames  $\mathcal{F}$ :

$$\mathcal{E} = \sum_{P \in \mathcal{F}} \mathcal{L}_1(G_P, D_P + S_P), \quad (5)$$

where  $\mathcal{L}_1$  is pixel-wise  $L1$  loss.

While Eq. (5) could be minimized directly to obtain  $W_{P,i}$  and  $SR_i$ , two factors introduce practical difficulties. First, specular highlights tend to be sparse and cover a small percentage of specular scene surfaces. Points on specular surfaces that don't see a highlight are difficult to differentiate from diffuse surface points, thus making the problem of assigning material weights to surface points severely under-constrained. Second, captured geometry is seldom perfect, and misalignments in reconstructed diffuse texture can result in incorrect SRMs. In the remainder of this section, we describe our approach to overcome these limiting factors.

**Material weight network.** To address the problem of material ambiguity, we pose the material assignment problem as a statistical pattern recognition task. We compute the 2D weight maps  $W_{P,i}(\mathbf{u})$  with a convolutional neural network  $w_\theta$  that learns to map a diffuse texture image patch to the blending weight of  $i$ th material:  $W_{P,i} = w_\theta(D_P)_i$ .

This network learns correlations between diffuse texture and material properties (i.e., shininess), and is trained on each scene by jointly optimizing the network weights and SRMs to reproduce the input images.

Since  $w_\theta$  predicts material weights in image-space, and therefore per view, we introduce a view-consistency regularization function  $\mathcal{V}(W_{P_1}, W_{P_2})$  penalizing the pixel-wise  $L1$  difference in the predicted materials between a pair of views when cross-projected to each other (i.e., one image is warped to the other using the known geometry and pose).

**Diffuse refinement network.** Small errors in geometry and calibration, as are typical in scanned models, cause misalignment and ghosting artifacts in the texture reconstruction  $D_P$ . Therefore, we introduce a refinement network  $u_\phi$  to correct these errors (Fig. 2). We replace  $D_P$  with the refined texture image:  $D'_P = u_\phi(D_P)$ . Similar to the material weights, we penalize the inconsistency of the refined diffuse images across viewpoints using  $\mathcal{V}(D'_{P_1}, D'_{P_2})$ . Both networks  $w_\theta$  and  $u_\phi$  follow the encoder-decoder architecture with residual connections [32, 26], while  $w_\theta$  has lower number of parameters. We refer readers to supplementary for more details.

**Robust Loss.** Because a pixel-wise loss alone is not robust to misalignments, we define the image distance metric  $\mathcal{L}$  as a combination of pixel-wise  $L1$  loss, perceptual loss  $\mathcal{L}_p$  computed from feature activations of a pretrained network [9], and adversarial loss [22, 30]. Our total loss, for a pair of images  $I_1, I_2$ , is:

$$\mathcal{L}(I_1, I_2; d) = \lambda_1 \mathcal{L}_1(I_1, I_2) + \lambda_p \mathcal{L}_p(I_1, I_2) + \lambda_G \mathcal{L}_G(I_1, I_2; d), \quad (6)$$

where  $d$  is the discriminator, and  $\lambda_1 = 0.01$ ,  $\lambda_p = 1.0$ , and  $\lambda_G = 0.05$  are balancing coefficients. The neural network-based perceptual and adversarial loss are effective because they are robust to image-space misalignments caused by errors in the estimated geometry and poses.

Finally, we add a sparsity term on the specular image  $\|S_P\|_1$  to regularize the specular component from containing colors from the diffuse texture.

Combining all elements, we get the final loss function:

$$SR^*, \theta^*, \phi^* = \arg \min_{SR, \theta, \phi} \max_d \sum_{P \in \mathcal{F}} \mathcal{L}(G_P, D'_P + S_P; d) + \lambda_S \|S_P\|_1 + \lambda_V \mathcal{V}(W_P, W_{P_r}) + \lambda_T \mathcal{V}(D'_P, D'_{P_r}), \quad (7)$$

where  $P_r$  is a randomly chosen frame in the same batch with  $P$  during each stochastic gradient descent step.  $\lambda_S$ ,  $\lambda_T$  and  $\lambda_V$  are set to 1e-4. An overview diagram is shown in Fig. 3. Fig. 5 shows that the optimization discovers coherent material regions and a detailed environment image.

### 3.3. Novel-View Neural Rendering

With reconstructed SRMs and material weights, we can synthesize specular appearance from any desired viewpoint

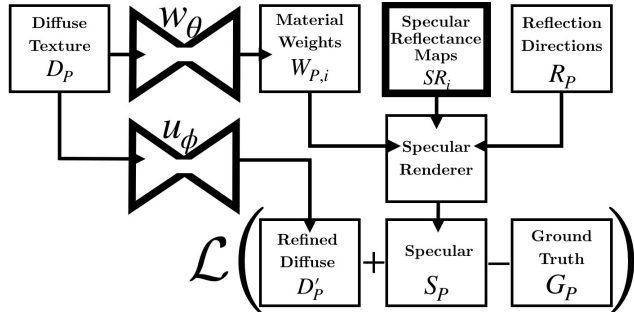


Figure 3: The components of our SRM estimation pipeline (optimized parameters shown in bold). We predict a view by adding refined diffuse texture  $D'_P$  (Fig. 2) and the specular image  $S_P$ .  $S_P$  is computed, for each pixel, by looking up the basis SRMs ( $SR_i$ 's) with surface reflection direction  $R_P$  and blending them with weights  $W_{P,i}$  obtained via network  $w_\theta$ . The loss between the predicted view and ground truth  $G_P$  is backpropagated to jointly optimize the SRM pixels and network weights.




Figure 4: Modeling interreflections. First row shows images of an unseen viewpoint rendered by a network trained with direct (a) and with interreflection + Fresnel models (b), compared to ground truth (c). Note accurate interreflections on the bottom of the green bottle (b). (d), (e), and (f) show first-bounce image (FBI), reflection direction image ( $R_P$ ), and Fresnel coefficient image (FCI), respectively. Best viewed digitally.

via Eq. (2). However, while the approach detailed in Sec. 3.2 reconstructs high quality SRMs, the renderings often lack realism (shown in supplementary), due to two factors. First, errors in geometry and camera pose can sometimes lead to weaker reconstructed highlights. Second, the SRMs do not model more complex light transport effects such as interreflections or Fresnel reflection. This section describes how we train a network to address these two limitations, yielding more realistic results.

Simulations only go so far, and computer renderings will never be perfect. In principle, you could train a CNN to ren-

der images as a function of viewpoint directly, training on actual photos. Indeed, several recent neural rendering methods adapt image translation [30] to learn mappings from projected point clouds [49, 58, 2] or a UV map image [66] to a photo. However, these methods struggle to extrapolate far away from the input views because their networks don't have built-in physical models of specular light transport.

Rather than treat the rendering problem as a black box, we arm the neural renderer with knowledge of physics – in particular, diffuse, specular, interreflection, and Fresnel reflection, to use in learning how to render images. Formally, we introduce an adversarial neural network-based generator  $g$  and discriminator  $d$  to render realistic photos.  $g$  takes as input our best prediction of diffuse  $D_P$  and specular  $S_P$  components for the current view (obtained from Eq. (7)), along with interreflection and Fresnel terms  $FBI$ ,  $R_P$ , and  $FCI$  that will be defined later in this section.

Consequently, the generator  $g$  receives  $C_P = (D_P, S_P, FBI, R_P, FCI)$  as input and outputs a prediction of the view, while the discriminator  $d$  scores its realism. We use the combination of pixelwise  $L_1$ , perceptual loss  $L_p$  [9], and the adversarial loss [30] as described in Sec. 3.2:

$$g^* = \arg \min_g \max_d \lambda_G \bar{\mathcal{L}}_G(g, d) + \lambda_p \bar{\mathcal{L}}_p(g) + \lambda_1 \bar{\mathcal{L}}_1(g), \quad (8)$$

where  $\bar{\mathcal{L}}_p(g) = \frac{1}{|\mathcal{F}|} \sum_{P \in \mathcal{F}} \mathcal{L}_p(g(C_P), G_P)$  is the mean of perceptual loss across all input images, and  $\bar{\mathcal{L}}_G(g, d)$  and  $\bar{\mathcal{L}}_1(g)$  are similarly defined as an average loss across frames. Note that this renderer  $g$  is *scene specific*, trained only on images of a particular scene, as commonly done in the neural rendering community [49, 66, 2].

**Modeling Interreflections and Fresnel Effects** Eq. (2) models only the direct illumination of each surface point by the environment, neglecting interreflections. While modeling full, global, diffuse + specular light transport is intractable, we can approximate first order interreflections by ray-tracing a first-bounce image (FBI) as follows. For each pixel  $\mathbf{u}$  in the virtual viewpoint to be rendered, cast a ray from the camera center through  $\mathbf{u}$ . If we pretend for now that every scene surface is a perfect mirror, that ray will bounce potentially multiple times and intersect multiple surfaces. Let  $\mathbf{x}_2$  be the second point of intersection of that ray with the scene. Render the pixel at  $\mathbf{u}$  in FBI with the diffuse color of  $\mathbf{x}_2$ , or with black if there is no second intersection (Fig. 4(d)).

Glossy (imperfect mirror) interreflections can be modeled by convolving the FBI with the BRDF. Strictly speaking, however, the interreflected image should be filtered in the *angular domain* [59], rather than image space, i.e., convolution of incoming light following the specular lobe whose center is the reflection ray direction  $\omega_r$ . Given  $\omega_r$ , angular domain convolution can be approximated in image

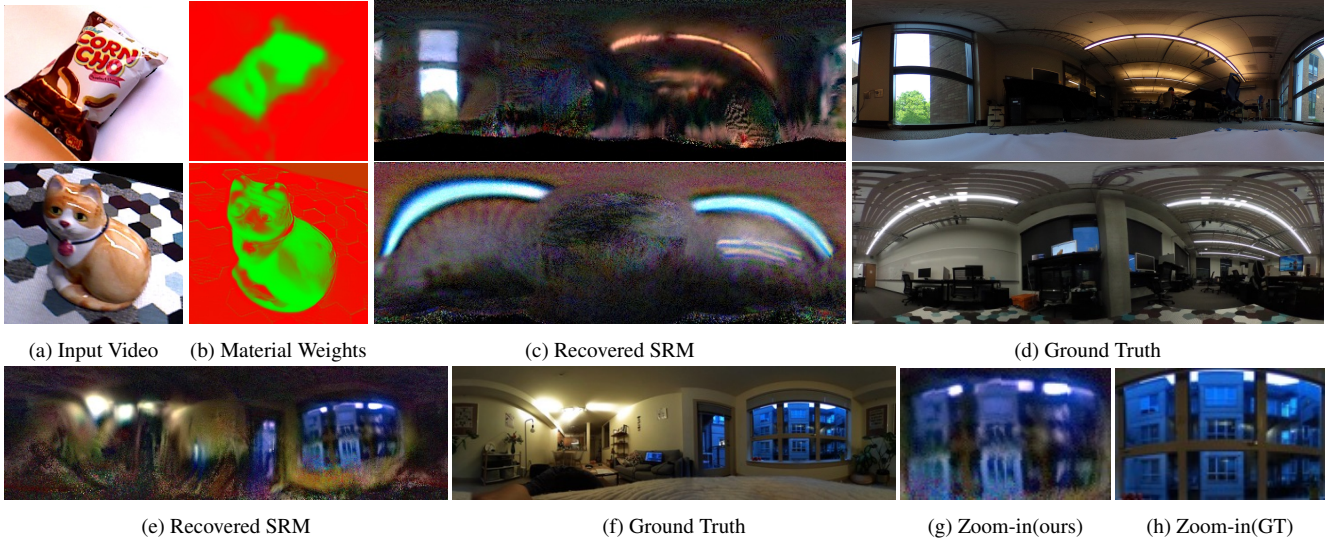


Figure 5: Sample results of recovered SRMs and material weights. Given input video frames (a), we recover global SRMs (c) and their linear combination weights (b) from the optimization of Eq. (7). The scenes presented here have two material bases, visualized with red and green channels. Estimated SRMs (c) corresponding to the shiny object surface (green channel) correctly capture the light sources of the scenes, shown in the reference panorama images (d). For both scenes the SRMs corresponding to the red channel is mostly black, thus not shown, as the surface is mostly diffuse. The recovered SRM of (c) overemphasizes blue channel due to oversaturation in input images. Third row shows estimation result from a video of the same bag of chips (first row) under different lighting. Close inspection of the recovered environment (g) reveals many scene details, including floors in a nearby building visible through the window.

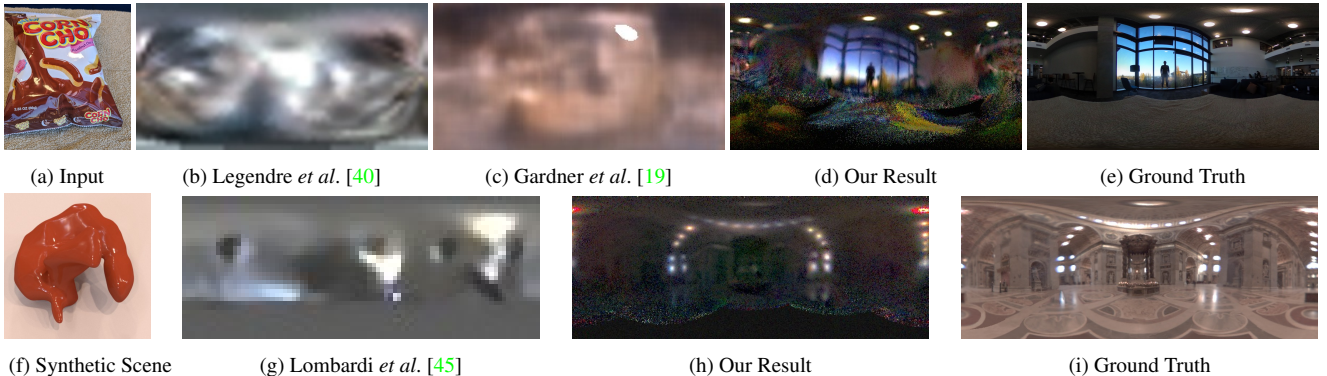


Figure 6: Comparisons with existing single-view and multi-view based environment estimation methods. Given a single image (a), DeepLight [40] (b), and Gardner *et al.* [18] (c), do not produce accurate environment reconstructions, relative to what we obtain from an RGBD video (d) which better matches ground truth (e). Additionally, from a video sequence and noisy geometry of a synthetic scene (f), our method (h) more accurately recovers the surrounding environment (i) compared to Lombardi *et al.* (g).

space by convolving the FBI image weighted by  $\omega_r$ . However, because we do not know the specular kernel, we let the network infer the weights using  $\omega_r$  as a guide. We encode the  $\omega_r$  for each pixel as a three-channel image  $R_P$  (Fig. 4(e)).

Fresnel effects make highlights stronger at near-glancing view angles and are important for realistic rendering. Fresnel coefficients are approximated following [61]:  $R(\alpha) = R_0 + (1 - R_0)(1 - \cos\alpha)^5$ , where  $\alpha$  is the angle between the surface normal and the camera ray, and  $R_0$  is a material-specific constant. We compute a Fresnel coefficient image

(FCI), where each pixel contains  $(1 - \cos\alpha)^5$ , and provide it to the network as an additional input, shown in Fig. 4(f).

In total, the rendering components  $C_P$  are now composed of five images: diffuse and specular images, FBI image,  $R_P$ , and FCI.  $C_P$  is then given as input to the neural network, and our network weights are optimized as in Eq. (8). Fig. 4 shows the effectiveness of the additional three rendering components for modeling interreflections.

### 3.4. Implementation Details

We follow [32] for the generator network architecture, use the PatchGAN discriminator [30], and employ the loss of LSGAN [47]. We use ADAM [38] with learning rate  $2e-4$  to optimize the objectives. Data augmentation was essential for viewpoint generalization, by applying random rotation, translation, flipping, and scaling to each input and output pair. More details can be found in supplementary.

### 3.5. Dataset

We captured ten sequences of RGBD video with a hand-held Primesense depth camera, featuring a wide range of materials, lighting, objects, environments, and camera paths. The length of each sequence ranges from 1500 to 3000 frames, which are split into train and test frames. Some of the sequences were captured such that the test views are very far from the training views, making them ideal for benchmarking the extrapolation abilities of novel-view synthesis methods. Moreover, many of the sequences come with ground truth HDR environment maps to facilitate future research on environment estimation. Further capture and data-processing details are in supplementary.

## 4. Experiments

We describe experiments to test our system’s ability to estimate images of the environment and synthesize novel viewpoints, and ablation studies to characterize the factors that most contribute to system performance.

We compare our approach to several state-of-the-art methods: recent single view lighting estimation methods (DeepLight [40], Gardner *et al.* [19]), an RGBD video-based lighting and material reconstruction method [45], an IR-based BRDF estimation method [56] (shown in supplementary), and two leading view synthesis methods capable of handling specular highlights – DeepBlending [27] and Deferred Neural Rendering (DNS) [66].

### 4.1. Environment Estimation

Our computed SRMs demonstrate our system’s ability to infer detailed images of the environment from the pattern and motion of specular highlights on an object. For example from 5(b), we can see the general layout of the living room, and even count the number of floors in buildings visible through the window. Note that the person capturing the video does not appear in the environment map because he is constantly moving. The shadow of the moving person, however, causes artifacts, e.g. the fluorescent lighting in the first row of Fig. 5 is not fully reconstructed.

Compared to state-of-the-art single view estimation methods [40, 19], our method produces a more accurate image of the environment, as shown in Fig. 6. Note our reconstruction shows a person standing near the window and

autumn colors in a tree visible through the window.

We compare with a multi-view RGBD based method [45] on a synthetic scene containing a red object, obtained from the authors. As in [45], we estimate lighting from the known geometry with added noise and a video of the scene rendering, but produce more accurate results (Fig. 6).

### 4.2. Novel View Synthesis

We recover specular reflectance maps and train a generative network for each video sequence. The trained model is then used to generate novel views from held-out views.

In the supplementary, we show novel view generation results for different scenes, along with the intermediate rendering components and ground truth images. As view synthesis results are better shown in video form, we strongly encourage readers to watch the *supplementary video*.

**Novel View Extrapolation** Extrapolating novel views far from the input range is particularly challenging for scenes with reflections. To test the operating range of our and other recent view synthesis results, we study how the quality of view prediction degrades as a function of the distance to the nearest input images (in difference of viewing angles) (Fig. 8). We measure prediction quality with perceptual loss [75], which is known to be more robust to shifts or misalignments, against the ground truth test image taken from same pose. We use two video sequences both containing highly reflective surfaces and with large differences in train and test viewpoints. We focus our attention on parts of the scene which exhibit significant view-dependent effects. That is, we mask out the diffuse backgrounds and measure the loss on only central objects of the scene. We compare our method with DeepBlending [27] and Thies *et al.* [66]. The quantitative (Fig. 8) and qualitative (Fig. 7) results show that our method is able to produce more accurate images of the scene from extrapolated viewpoints.

### 4.3. Robustness

Our method is robust to various scene configurations, such as scenes containing multiple objects (Fig. 7), spatially varying materials (Fig. 9), and concave surfaces (Fig. 10). In the supplementary, we study how the loss functions and surface roughness affect our results.

## 5. Limitations and Future work

Our approach relies on the reconstructed mesh obtained from fusing depth images of consumer-level depth cameras and thus fails for surfaces out of the operating range of these cameras, e.g., thin, transparent, or mirror surfaces. Our recovered environment images are filtered by the surface BRDF; separating these two factors is an interesting topic of future work, perhaps via data-driven deconvolution (e.g. [71]). Last, reconstructing a room-scale photorealistic appearance model remains a major open challenge.

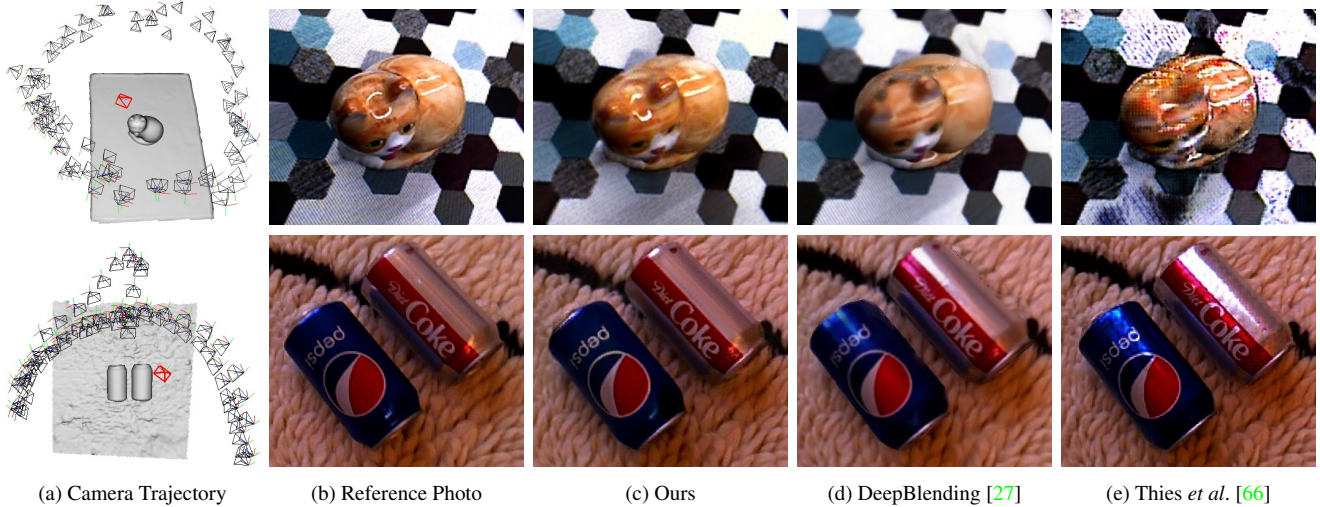


Figure 7: View extrapolation to extreme viewpoints. We evaluate novel view synthesis on test views (red frusta) that are furthest from the input views (black frusta) (a). The view predictions of DeepBlending [27] and Thies *et al.* [66] (d,e) are notably different from the reference photographs (b), e.g., missing highlights on the back of the cat, and incorrect highlights at the bottom of the cans. Thies *et al.* [66] shows severe artifacts, likely because their learned UV texture features overfits to the input views, and thus cannot generalize to very different viewpoints. Our method (c) produces images with highlights appearing at correct locations.




Figure 8: Quantitative comparisons for novel view synthesis. We plot the perceptual loss [75] between a novel view rendering and the ground truth test image as a function of its distance to the nearest training view (measured in angle between the view vectors). We compare our method with two leading NVS methods [27, 66] on two scenes. On average, our results have lowest error.

## Acknowledgement

This work was supported by funding from Facebook, Google, Futurewei, and the UW Reality Lab.

## References

- [1] Miika Aittala, Tim Weyrich, Jaakko Lehtinen, et al. Two-shot svbrdf capture for stationary materials. 2015. 3
- [2] Kara-Ali Aliev, Dmitry Ulyanov, and Victor Lempitsky. Neural point-based graphics. *arXiv preprint arXiv:1906.08240*, 2019. 2, 5
- [3] Neil Alldrin, Todd Zickler, and David Kriegman. Photometric stereo with non-parametric and spatially-varying re-




Figure 9: Image (a) shows a synthesized novel view using neural rendering (Sec. 3.3) of a scene with multiple glossy materials. The spatially varying materials (SRM blending weights) of the wooden tabletop and the laptop are accurately estimated by our algorithm (Sec. 3.2), as visualized in (b).




Figure 10: Concave surface reconstruction. The appearance of highly concave bowls is realistically reconstructed by our system. The rendered result (b) captures both occlusions and highlights of the ground truth (a).

flectance. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2008. 3

- [4] Dejan Azinović, Tzu-Mao Li, Anton Kaplanyan, and Matthias Nießner. Inverse path tracing for joint material and



- lighting estimation. *arXiv preprint arXiv:1903.07145*, 2019. 3
- [5] Jonathan T Barron and Jitendra Malik. Shape, albedo, and illumination from a single image of an unknown object. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 334–341. IEEE, 2012. 2
- [6] 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. 2
- [7] Sai Bi, Xiaoguang Han, and Yizhou Yu. An l1 image transform for edge-preserving smoothing and scene-level intrinsic decomposition. *ACM Transactions on Graphics (TOG)*, 34(4):78, 2015. 2
- [8] Anpei Chen, Minye Wu, Yingliang Zhang, Nianyi Li, Jie Lu, Shenghua Gao, and Jingyi Yu. Deep surface light fields. *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, 1(1):14, 2018. 2
- [9] Qifeng Chen and Vladlen Koltun. Photographic image synthesis with cascaded refinement networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1511–1520, 2017. 4, 5
- [10] Wei-Chao Chen, Jean-Yves Bouguet, Michael H Chu, and Radek Grzeszczuk. Light field mapping: Efficient representation and hardware rendering of surface light fields. In *ACM Transactions on Graphics (TOG)*, volume 21, pages 447–456. ACM, 2002. 2
- [11] Inchang Choi, Orazio Gallo, Alejandro Troccoli, Min H Kim, and Jan Kautz. Extreme view synthesis. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7781–7790, 2019. 3
- [12] Abe Davis, Marc Levoy, and Fredo Durand. Unstructured light fields. In *Computer Graphics Forum*, volume 31, pages 305–314. Wiley Online Library, 2012. 2
- [13] Abe Davis, Michael Rubinstein, Neal Wadhwa, Gautham Mysore, Fredo Durand, and William T. Freeman. The visual microphone: Passive recovery of sound from video. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 33(4):79:1–79:10, 2014. 1
- [14] Paul Ernest Debevec. *Modeling and rendering architecture from photographs*. University of California, Berkeley, 1996. 3
- [15] Paul E. Debevec and Jitendra Malik. Recovering high dynamic range radiance maps from photographs. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '97*, pages 369–378, New York, NY, USA, 1997. ACM Press/Addison-Wesley Publishing Co. 2
- [16] Yue Dong, Guojun Chen, Pieter Peers, Jiawan Zhang, and Xin Tong. Appearance-from-motion: Recovering spatially varying surface reflectance under unknown lighting. *ACM Transactions on Graphics (TOG)*, 33(6):193, 2014. 2, 3
- [17] John Flynn, Ivan Neulander, James Philbin, and Noah Snavely. Deepstereo: Learning to predict new views from the world’s imagery. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5515–5524, 2016. 2
- [18] Marc-André Gardner, Yannick Hold-Geoffroy, Kalyan Sunkavalli, Christian Gagné, and Jean-François Lalonde. Deep parametric indoor lighting estimation. *arXiv preprint arXiv:1910.08812*, 2019. 2, 6
- [19] 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. *arXiv preprint arXiv:1704.00090*, 2017. 2, 6, 7
- [20] Stamatis 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. 2
- [21] Dan B Goldman, Brian Curless, Aaron Hertzmann, and Steven M Seitz. Shape and spatially-varying brdfs from photometric stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(6):1060–1071, 2010. 3
- [22] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014. 4
- [23] Steven J Gortler, Radek Grzeszczuk, Richard Szeliski, and Michael F Cohen. The lumigraph. 1996. 2
- [24] Xiaojie Guo, Xiaochun Cao, and Yi Ma. Robust separation of reflection from multiple images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2187–2194, 2014. 2
- [25] Byeong-Ju Han and Jae-Young Sim. Reflection removal using low-rank matrix completion. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5438–5446, 2017. 2
- [26] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 4
- [27] Peter Hedman, Julien Philip, True Price, Jan-Michael Frahm, George Drettakis, and Gabriel Brostow. Deep blending for free-viewpoint image-based rendering. In *SIGGRAPH Asia 2018 Technical Papers*, page 257. ACM, 2018. 2, 7, 8
- [28] Yannick Hold-Geoffroy, Akshaya Athawale, and Jean-François Lalonde. Deep sky modeling for single image outdoor lighting estimation. *arXiv preprint arXiv:1905.03897*, 2019. 2
- [29] Yannick Hold-Geoffroy, Kalyan Sunkavalli, Sunil Hadap, Emiliano Gambaretto, and Jean-François Lalonde. Deep outdoor illumination estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7312–7321, 2017. 2
- [30] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017. 4, 5, 7
- [31] Jan Jachnik, Richard A Newcombe, and Andrew J Davison. Real-time surface light-field capture for augmentation of pla-

- nar specular surfaces. In *2012 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pages 91–97. IEEE, 2012. 2
- [32] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016. 4, 7
- [33] Nima Khademi Kalantari, Ting-Chun Wang, and Ravi Ramamoorthi. Learning-based view synthesis for light field cameras. *ACM Transactions on Graphics (TOG)*, 35(6):193, 2016. 3
- [34] Kevin Karsch, Varsha Hedau, David Forsyth, and Derek Hoiem. Rendering synthetic objects into legacy photographs. In *ACM Transactions on Graphics (TOG)*, volume 30, page 157. ACM, 2011. 2
- [35] Kevin Karsch, Kalyan Sunkavalli, Sunil Hadap, Nathan Carr, Hailin Jin, Rafael Fonte, Michael Sittig, and David Forsyth. Automatic scene inference for 3d object compositing. *ACM Transactions on Graphics (TOG)*, 33(3):32, 2014. 2
- [36] Jan Kautz, Pere-Pau Vázquez, Wolfgang Heidrich, and Hans-Peter Seidel. A unified approach to prefiltered environment maps. In *Rendering Techniques 2000*, pages 185–196. Springer, 2000. 3
- [37] Kihwan Kim, Jinwei Gu, Stephen Tyree, Pavlo Molchanov, Matthias Nießner, and Jan Kautz. A lightweight approach for on-the-fly reflectance estimation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 20–28, 2017. 3
- [38] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 7
- [39] Joo Ho Lee, Adrian Jarabo, Daniel S Jeon, Diego Gutierrez, and Min H Kim. Practical multiple scattering for rough surfaces. In *SIGGRAPH Asia 2018 Technical Papers*, page 275. ACM, 2018. 3
- [40] Chloe LeGendre, Wan-Chun Ma, Graham Fyffe, John Flynn, Laurent Charbonnel, Jay Busch, and Paul Debevec. Deep-light: Learning illumination for unconstrained mobile mixed reality. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5918–5928, 2019. 2, 6, 7
- [41] Hendrik Lensch, Jan Kautz, Michael Goesele, Wolfgang Heidrich, and Hans-Peter Seidel. Image-based reconstruction of spatial appearance and geometric detail. *ACM Transactions on Graphics (TOG)*, 22(2):234–257, 2003. 3
- [42] Marc Levoy and Pat Hanrahan. Light field rendering. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 31–42. ACM, 1996. 2
- [43] Zhengqin Li, Zexiang Xu, Ravi Ramamoorthi, Kalyan Sunkavalli, and Manmohan Chandraker. Learning to reconstruct shape and spatially-varying reflectance from a single image. In *SIGGRAPH Asia 2018 Technical Papers*, page 269. ACM, 2018. 3
- [44] Stephen Lombardi and Ko Nishino. Reflectance and natural illumination from a single image. In *European Conference on Computer Vision*, pages 582–595. Springer, 2012. 2
- [45] Stephen Lombardi and Ko Nishino. Radiometric scene decomposition: Scene reflectance, illumination, and geometry from rgb-d images. In *2016 Fourth International Conference on 3D Vision (3DV)*, pages 305–313. IEEE, 2016. 3, 6, 7
- [46] Robert Maier, Kihwan Kim, Daniel Cremers, Jan Kautz, and Matthias Nießner. Intrinsic3d: High-quality 3d reconstruction by joint appearance and geometry optimization with spatially-varying lighting. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3114–3122, 2017. 2
- [47] Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2794–2802, 2017. 7
- [48] Abhimitra Meka, Maxim Maximov, Michael Zollhoefer, Avishek Chatterjee, Hans-Peter Seidel, Christian Richardt, and Christian Theobalt. Lime: Live intrinsic material estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6315–6324, 2018. 3
- [49] Moustafa Meshry, Dan B. Goldman, Sameh Khamis, Hugues Hoppe, Rohit Pandey, Noah Snively, and Ricardo Martin-Brualla. Neural rerendering in the wild. *CoRR*, abs/1904.04290, 2019. 2, 5
- [50] Ben Mildenhall, Pratul P Srinivasan, Rodrigo Ortiz-Cayon, Nima Khademi Kalantari, Ravi Ramamoorthi, Ren Ng, and Abhishek Kar. Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *arXiv preprint arXiv:1905.00889*, 2019. 3
- [51] Giljoo Nam, Joo Ho Lee, Diego Gutierrez, and Min H Kim. Practical svbrdf acquisition of 3d objects with unstructured flash photography. In *SIGGRAPH Asia 2018 Technical Papers*, page 267. ACM, 2018. 3
- [52] Richard A Newcombe, Shahram Izadi, Otmar Hilliges, David Molyneaux, David Kim, Andrew J Davison, Pushmeet Kohli, Jamie Shotton, Steve Hodges, and Andrew W Fitzgibbon. Kinectfusion: Real-time dense surface mapping and tracking. In *ISMAR*, volume 11, pages 127–136, 2011. 2, 3
- [53] Fred E Nicodemus. Directional reflectance and emissivity of an opaque surface. *Applied optics*, 4(7):767–775, 1965. 3
- [54] Ko Nishino and Shree Nayar. Eyes for relighting. *ACM Trans. Graph.*, 23:704–711, 08 2004. 2
- [55] Roy Or-El, Guy Rosman, Aaron Wetzler, Ron Kimmel, and Alfred M Bruckstein. Rgb-d-fusion: Real-time high precision depth recovery. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5407–5416, 2015. 2
- [56] Jeong Joon Park, Richard Newcombe, and Steve Seitz. Surface light field fusion. In *2018 International Conference on 3D Vision (3DV)*, pages 12–21. IEEE, 2018. 2, 3, 7
- [57] Bui Tuong Phong. Illumination for computer generated pictures. *Communications of the ACM*, 18(6):311–317, 1975. 3
- [58] Francesco Pittaluga, Sanjeev J Koppal, Sing Bing Kang, and Sudipta N Sinha. Revealing scenes by inverting structure from motion reconstructions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 145–154, 2019. 5

- [59] 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. ACM, 2001. 2, 5
- [60] Thomas Richter-Trummer, Denis Kalkofen, Jinwoo Park, and Dieter Schmalstieg. Instant mixed reality lighting from casual scanning. In *2016 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pages 27–36. IEEE, 2016. 2
- [61] Christophe Schlick. An inexpensive brdf model for physically-based rendering. In *Computer graphics forum*, volume 13, pages 233–246. Wiley Online Library, 1994. 6
- [62] Sudipta N Sinha, Johannes Kopf, Michael Goesele, Daniel Scharstein, and Richard Szeliski. Image-based rendering for scenes with reflections. *ACM Trans. Graph.*, 31(4):100–1, 2012. 2
- [63] Shuran Song and Thomas Funkhouser. Neural illumination: Lighting prediction for indoor environments. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6918–6926, 2019. 2
- [64] Pratul P Srinivasan, Tongzhou Wang, Ashwin Sreelal, Ravi Ramamoorthi, and Ren Ng. Learning to synthesize a 4d rgbd light field from a single image. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2243–2251, 2017. 3
- [65] Richard Szeliski, Shai Avidan, and P Anandan. Layer extraction from multiple images containing reflections and transparency. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition. CVPR 2000 (Cat. No. PR00662)*, volume 1, pages 246–253. IEEE, 2000. 2, 3
- [66] Justus Thies, Michael Zollhöfer, and Matthias Nießner. Deferred neural rendering: Image synthesis using neural textures. *arXiv preprint arXiv:1904.12356*, 2019. 2, 5, 7, 8
- [67] Gregory J Ward et al. Measuring and modeling anisotropic reflection. *Computer Graphics*, 26(2):265–272, 1992. 3
- [68] Daniel N Wood, Daniel I Azuma, Ken Aldinger, Brian Curless, Tom Duchamp, David H Salesin, and Werner Stuetzle. Surface light fields for 3d photography. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 287–296. ACM Press/Addison-Wesley Publishing Co., 2000. 1, 2, 3
- [69] Hongzhi Wu, Zhaotian Wang, and Kun Zhou. Simultaneous localization and appearance estimation with a consumer rgb-d camera. *IEEE transactions on visualization and computer graphics*, 22(8):2012–2023, 2015. 2, 3
- [70] Rui Xia, Yue Dong, Pieter Peers, and Xin Tong. Recovering shape and spatially-varying surface reflectance under unknown illumination. *ACM Transactions on Graphics (TOG)*, 35(6):187, 2016. 2
- [71] Li Xu, Jimmy SJ Ren, Ce Liu, and Jiaya Jia. Deep convolutional neural network for image deconvolution. In *Advances in neural information processing systems*, pages 1790–1798, 2014. 7
- [72] Zexiang Xu, Sai Bi, Kalyan Sunkavalli, Sunil Hadap, Hao Su, and Ravi Ramamoorthi. Deep view synthesis from sparse photometric images. *ACM Transactions on Graphics (TOG)*, 38(4):76, 2019. 3
- [73] Tianfan Xue, Michael Rubinstein, Ce Liu, and William T Freeman. A computational approach for obstruction-free photography. *ACM Transactions on Graphics (TOG)*, 34(4):79, 2015. 2
- [74] Edward Zhang, Michael F Cohen, and Brian Curless. Emptying, refurbishing, and relighting indoor spaces. *ACM Transactions on Graphics (TOG)*, 35(6):174, 2016. 3
- [75] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 586–595, 2018. 7, 8
- [76] Xuaner Zhang, Ren Ng, and Qifeng Chen. Single image reflection separation with perceptual losses. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4786–4794, 2018. 2
- [77] Youyi Zheng, Xiang Chen, Ming-Ming Cheng, Kun Zhou, Shi-Min Hu, and Niloy J Mitra. Interactive images: cuboid proxies for smart image manipulation. *ACM Trans. Graph.*, 31(4):99–1, 2012. 2
- [78] Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. Stereo magnification: Learning view synthesis using multiplane images. *arXiv preprint arXiv:1805.09817*, 2018. 3
- [79] Zhiming Zhou, Guojun Chen, Yue Dong, David Wipf, Yong Yu, John Snyder, and Xin Tong. Sparse-as-possible svbrdf acquisition. *ACM Transactions on Graphics (TOG)*, 35(6):189, 2016. 3
- [80] Todd Zickler, Sebastian Enrique, Ravi Ramamoorthi, and Peter Belhumeur. Reflectance Sharing: Image-based Rendering from a Sparse Set of Images. In Kavita Bala and Philip Dutre, editors, *Eurographics Symposium on Rendering (2005)*. The Eurographics Association, 2005. 3
- [81] Michael Zollhöfer, Angela Dai, Matthias Innmann, Chenglei Wu, Marc Stamminger, Christian Theobalt, and Matthias Nießner. Shading-based refinement on volumetric signed distance functions. *ACM Transactions on Graphics (TOG)*, 34(4):96, 2015. 2