

A Neural Rendering Framework for Free-Viewpoint Relighting

Zhang Chen^{1,2,3}Anpei Chen¹Guli Zhang¹Chengyuan Wang⁴Yu Ji⁵Kiriakos N. Kutulakos⁶Jingyi Yu¹¹ ShanghaiTech University² Shanghai Institute of Microsystem and Information Technology³ University of Chinese Academy of Sciences⁴ Shanghai University⁵ DGene, Inc.⁶ University of Toronto

{chenzhang, chenap, zhanggl, yujingyi}@shanghaitech.edu.cn

ericw385@i.shu.edu.cn

yu.ji@dgene.com

kyros@cs.toronto.edu

github.com/LansburyCH/relightable-nr

Abstract

We present a novel *Relightable Neural Renderer (RNR)* for simultaneous view synthesis and relighting using multi-view image inputs. Existing neural rendering (NR) does not explicitly model the physical rendering process and hence has limited capabilities on relighting. RNR instead models image formation in terms of environment lighting, object intrinsic attributes, and light transport function (LTF), each corresponding to a learnable component. In particular, the incorporation of a physically based rendering process not only enables relighting but also improves the quality of view synthesis. Comprehensive experiments on synthetic and real data show that RNR provides a practical and effective solution for conducting free-viewpoint relighting.

1. Introduction

Neural rendering (NR) has shown great success in the past few years on producing photorealistic images under complex geometry, surface reflectance, and environment lighting. Unlike traditional modeling and rendering techniques that rely on elaborate setups to capture detailed object geometry and accurate surface reflectance properties, often also with excessive artistic manipulations, NR can produce compelling results by using only images captured under uncontrolled illumination. By far, most existing NR methods have focused on either free-viewpoint rendering under fixed illumination or image-based relighting under fixed viewpoint. In this paper, we explore the problem of simultaneous novel view synthesis and relighting using NR.

State-of-the-art deep view synthesis techniques follow the pipeline that first extracts deep features from input images and 3D models, then projects the features to the image space via traditional camera projection, and finally ap-




Figure 1. Results from our Relightable Neural Renderer (RNR). Top row shows the relighting results for a synthetic sphere composed of complex materials. Bottom row shows free-viewpoint relighting results for real captured data.

plies a rendering network to render the projected features to a RGB image. Such approaches exploit learnable components to both encode 3D representations and model the rendering process. Approaches such as neural point cloud [2], neural volume [63] and neural texture [69] utilize deep representations for 3D content. With rich training data, these methods can tolerate inaccuracies in geometry and maintain reasonable rendering quality. For example, DeepVoxels [63] uses a learnable volume as an alternative to standard 3D representation while combining physically based forward/backward projection operators for view synthesis.

Using NR to produce visually plausible free-viewpoint relighting is more difficult compared with changing viewpoints under fixed illumination. This is because under fixed illumination, existing NRs manage to model 2D/3D geometry as learnable components to directly encode appearance of different views. Relighting, in contrast, requires further separating appearance into object intrinsic attributes and illumination. From a NR perspective, the final rendering step

in existing approaches cannot yet achieve such separation.

In this paper, we present a novel Relightable Neural Renderer (RNR) for view synthesis and relighting from multi-view inputs. A unique step in our approach is that we model image formation in terms of environment lighting, object intrinsic attributes, and light transport function (LTF). RNR sets out to conduct regression on these three individual components rather than directly translating deep features to appearance as in existing NR. In addition, the use of LTF instead of a parametric BRDF model extends the capability of modeling global illumination. While enabling relighting, RNR can also produce view synthesis using the same network architecture. Comprehensive experiments on synthetic and real data show that RNR provides a practical and effective solution for conducting free-viewpoint relighting.

2. Related Work

Image-based Rendering (IBR). Traditional IBR methods [17, 37, 24, 5, 7, 86, 23, 57] synthesize novel views by blending pixels from input images. Compared with physically based rendering, which requires high-resolution geometry and accurate surface reflectance, they can use lower quality geometry as proxies to produce relatively high quality rendering. The ultimate rendering quality, however, is a trade-off between the density of sampled images and geometry: low quality geometry requires dense sampling to reduce artifacts; otherwise the rendering exhibits various artifacts including ghosting, aliasing, misalignment and appearance jumps. The same trade-off applies to image-based relighting, although for low frequency lighting, sparse sampling may suffice to produce realistic appearance. Hand-crafted blending schemes [9, 35, 8, 23, 57] have been developed for specific rendering tasks but they generally require extensive parameter tuning.

Deep View Synthesis. Recently, there has been a large corpus of works on learning-based novel view synthesis. [68, 13] learn an implicit 3D representation by training on synthetic datasets. Warping-based methods [88, 55, 67, 90, 28, 11] synthesize novel views by predicting the optical flow field. Flow estimation can also be enhanced with geometry priors [87, 45]. Kalantari *et al.* [30] separate the synthesis process into disparity and color estimations for light field data. Srinivasan *et al.* [66] further extend to RGB-D view synthesis on small baseline light fields.

Eslami *et al.* [14] propose Generative Query Network to embed appearances of different views in latent space. Disentangled understanding of scenes can also be conducted through interpretable transformations [82, 36, 77], Lie groups-based latent variables [15] or attention modules [6]. Instead of 2D latent features, [72, 52, 20] utilize volumetric representations as a stronger multi-view constraint

whereas Sitzmann *et al.* [64] represent a scene as a continuous mapping from 3D geometry to deep features.

To create more photo-realistic rendering for a wide viewing range, [22, 70, 10, 63, 47, 69, 2, 61, 49, 79] require many more images as input. Hedman *et al.* [22] learn the blending scheme in IBR. Thies *et al.* [70] model the view-dependent component with self-supervised learning and then combine it with the diffuse component. Chen *et al.* [10] apply fully connected networks to model the surface light field by exploiting appearance redundancies. Volume-based methods [63, 47] utilize learnable 3D volume to represent scene and combine with projection or ray marching to enforce geometric constraint. Thies *et al.* [69] present a novel learnable neural texture to model rendering as image translation. They use coarse geometry for texture projection and offer flexible content editing. Aliev *et al.* [2] directly use neural point cloud to avoid surface meshing. Auxiliary information such as poses can be used to synthesize more complex objects such as human bodies [61].

To accommodate relighting, Meshry *et al.* [49] learn an embedding for appearance style whereas Xu *et al.* [79] use deep image-based relighting [81] on multi-view multi-light photometric images captured using specialized gantry. Geometry-differentiable neural rendering [58, 46, 44, 40, 32, 48, 84, 43, 27, 71, 51] can potentially handle relighting but our technique focuses on view synthesis and relighting without modifying 3D geometry.

Free-Viewpoint Relighting. Earlier free-viewpoint relighting of real world objects requires delicate acquisitions of reflectance [18, 75, 76] while more recent low-cost approaches still require controlled active illumination or known illumination/geometry [50, 25, 89, 78, 16, 83, 42, 31, 12, 41, 80]. Our work aims to use multi-view images captured under single unknown natural illumination. Previous approaches solve this ill-posed problem via spherical harmonics (SH) [85] or wavelets [19] or both [39] to represent illumination and a parametric BRDF model to represent reflectance. Imber *et al.* [26] extract pixel-resolution intrinsic textures. Despite these advances, accurate geometry remains as a key component for reliable relighting whereas our RNR aims to simultaneously compensate for geometric inaccuracy and disentangle intrinsic properties from lighting. Tailored illumination models can support outdoor relighting [59, 21, 60] or indoor inverse rendering [3] whereas our RNR uses a more generic lighting model for learning the light transport process. Specifically, our work uses a set of multi-view images of an object under fixed yet unknown natural illumination as input. To carry out view projection and texture mapping, we assume known camera parameters of the input views and known coarse 3D geometry of the object, where standard structure-from-motion and multi-view stereo reconstruction can provide reliable estimations.




Figure 2. The neural rendering pipeline of RNR.

3. Image Formation Model

Under the rendering equation [29], the radiance \mathbf{I} emitting from point \mathbf{x} at viewing direction ω_o is computed as:

$$\mathbf{I}(\mathbf{x}, \omega_o) = \int_{S^2} f_r(\mathbf{x}, \omega_i, \omega_o) v(\mathbf{x}, \omega_i) L(\mathbf{x}, \omega_i) \mathbf{n} \cdot \omega_i d\omega_i, \quad (1)$$

where $L(\mathbf{x}, \omega_i)$ is the radiance that arrives at point \mathbf{x} from direction ω_i . $v(\mathbf{x}, \omega_i)$ denotes the visibility of \mathbf{x} from direction ω_i and $f_r(\mathbf{x}, \omega_i, \omega_o)$ is the bidirectional reflectance distribution function (BRDF) that describes the ratio of outgoing radiance over the incident irradiance. S^2 is the upper hemisphere surrounding the surface point. For distant illumination, $L(\mathbf{x}, \omega_i)$ can be replaced with $L(\omega_i)$.

Instead of separately conducting regression to recover each individual term in Eq. 1, we learn light transport function (LTF) $\mathbf{T}(\mathbf{x}, \omega_i, \omega_o) = f_r(\mathbf{x}, \omega_i, \omega_o) v(\mathbf{x}, \omega_i) \mathbf{n} \cdot \omega_i$. By further separating view-independent albedo $\rho(\mathbf{x})$ from LTF (we still refer to this counterpart with albedo factored out as LTF in this paper for brevity), we have

$$\mathbf{I}(\mathbf{x}, \omega_o) = \int_{S^2} \rho(\mathbf{x}) \mathbf{T}(\mathbf{x}, \omega_i, \omega_o) L(\omega_i) d\omega_i. \quad (2)$$

The key observation here is that, for static objects, the LTF can be decoupled from illumination. This allows us to decompose photometric attributes into albedo, light transport and illumination for conducting relighting. Specifically, our RNR uses a network to represent the LTF $\mathbf{T}(\cdot)$. Learning the LTF instead of the BRDF has several advantages. First, it can compensate for outlier effects such as incorrect visibility caused by inaccurate 3D proxy common in IBR. Second, under distant illumination, since LTF can be regarded as the total contribution (with all light paths taken into account) from incoming radiance along a direction to the outgoing radiance, it can potentially encode non-local effects such as inter-reflection. Finally, it reduces the computation when evaluating the radiance of a pixel. It is worth

noting that inferring the LTF can be viewed as the inverse problem of precomputed radiance transfer (PRT) [65] which is widely used in physically based rendering.

Same with previous relighting techniques, we assume illumination can be modelled using Spherical Harmonics (SH) up to order 10. The implicit assumption here is that the object cannot be too specular or mirror like. Following common practices, we further decompose into diffuse and specular components, which gives:

$$\mathbf{I}(\mathbf{x}, \omega_o) = \int_{S^2} \rho_d(\mathbf{x}) \mathbf{T}_d(\mathbf{x}, \omega_i, \omega_o) \sum_k c_k Y_k(\omega_i) d\omega_i + \int_{S^2} \rho_s(\mathbf{x}) \mathbf{T}_s(\mathbf{x}, \omega_i, \omega_o) \sum_k c_k Y_k(\omega_i) d\omega_i, \quad (3)$$

where ρ_d and \mathbf{T}_d are the albedo and LTF of diffuse component, ρ_s and \mathbf{T}_s are the albedo and LTF of specular component, Y_k is the k th SH basis and c_k its coefficient.

Illumination Initialization. Our SH representation contains 121 coefficients for each color channel. We first exploit the background regions of multi-view images to initialize illumination. We assume that background pixels lie faraway, so we establish the image-to-panorama mapping and fill in the environment map with image pixels. We take the median of the image pixels that map to the same position in environment map to reduce ghosting artifacts. We then project the environment map onto SH basis to obtain the initial value of SH coefficients.

Neural Texture. Neural texture [69] provides an efficient encoding of latent properties of 3D scenes. It can be seen as an extension of traditional texture-space data such as color texture, normal map, displacement map, etc. While these data record certain hand-crafted properties of 3D content,

neural texture is learnable and can be trained to encode the critical information for a given task (e.g., novel view synthesis). We use the first 3 channels of neural texture as diffuse albedo and second 3 channels as specular albedo. For the rest of the channels, we leave them unconstrained so as to encode latent properties. To project neural texture to image space, we first rasterize 3D proxy using camera parameters to obtain uv map (texel-to-pixel mapping) and use bilinear interpolation to sample features from neural texture. Following [69], we use a 4-level mipmap Laplacian pyramid for neural texture and set the resolution of the top level as 512×512 . We also evaluate the first 9 SH coefficients at per-pixel view direction and multiply with channel 7-15 of projected neural texture (neural image).

4. Relightable Neural Renderer (RNR)

Next, we set out to simultaneously estimate the albedos $\rho_d(\cdot)$, $\rho_s(\cdot)$, the LTFs $\mathbf{T}_d(\cdot)$, $\mathbf{T}_s(\cdot)$ and the SH coefficients c_k . We use the neural texture [69] to encode the albedo and additional latent properties of the object. We then propose sampling schemes for the light directions used in evaluating Eq. 3. Next, we propose a Light-Transport-Net (LTN) to predict light transport at the sampled light directions for each pixel. Note that the entire process is differentiable and only requires 2D supervision from input multi-view images. Fig. 2 shows our pipeline.

4.1. Light Direction Sampling

Instead of densely sampling light directions for each vertex (high angular resolution but low spatial resolution), we resort to sparsely sampling light directions for each pixel (low angular resolution but high spatial resolution). In this case, high rendering quality can be achieved even with coarse 3D proxy. We argue that under SH lighting, using sparse light direction sampling only leads to minor inaccuracy on the radiance evaluated in Eq. 3, which can be effectively compensated by LTN.

Since diffuse and specular light transport behaves differently based on light direction and view direction, we utilize different sampling schemes, as shown in Fig. 3. For the diffuse component, we first construct k_d cones centered around the surface normal, with half angles of $\{\theta_1^d, \theta_2^d, \dots, \theta_{k_d}^d\}$. Then we uniformly sample directions on each cone. This is motivated by the fact that diffuse light transport (ignoring visibility and other effects) follows a cosine attenuation based on the angle between light direction and surface normal. Therefore, light directions nearer to the surface normal are more likely to contribute more to the radiance at the surface point. For the specular component, we similarly construct k_s cones around the surface normal, and uniformly sample on these cones to obtain halfway directions. Then we reflect view direction around these halfway directions to obtain sampled light directions. This is motivated by the




Figure 3. Light direction sampling schemes for diffuse and specular components.

microfacets theory which models surfaces as collections of perfect mirror microfacets. The normals of these microfacets follow a normal distribution function, which we assume to cluster around the macro surface normal.

We carry out the above light direction sampling in tangent space and then transform to world space by

$$\omega_i(\mathbf{x}) = \mathbf{R}_{TBN}(\mathbf{x}) \cdot \omega'_i(\mathbf{x}), \quad (4)$$

where $\omega'_i(\mathbf{x})$ is the sampled directions in tangent space and $\mathbf{R}_{TBN}(\mathbf{x})$ is the rotation matrix from tangent space to world space. By stacking the sampled light directions $\{\omega_i\}_d$, $\{\omega_i\}_s$ of the two components along the channel dimension, we form a light direction map, which is then input to LTN.

4.2. Light Transport Estimation

Our LTN consists of a graph convolutional network (GCN) to extract global geometric features and a modified U-Net to predict per-pixel light transports at the sampled light directions for diffuse and specular components.

We first concatenate neural image with view direction map, normal map and light direction map as input to the U-Net. As 2D convolutional network does not fully exploit information in a non-Euclidean structure data, we further augment U-Net with a GCN [34, 4] to extract global features of the 3D geometry. Inspired by [62, 74], we use dynamic edge connections during the training of GCN to learn a better graph representations. But different from [74], which changes the edge connection by finding the nearest neighbour of each vertex, we follow the step of [73, 38] and apply a dilated K-NN method on feature space to search the neighborhood of each vertex. Moreover, rather than using naive GCN, we utilize ResGCN - a much deeper GCN with residual blocks [38] to gather higher-level features of each vertex. At the end of the ResGCN, a fully connected layer is applied to fuse all the features into global geometric features. We repeat and concatenate this feature vector with the U-Net feature map after the first downsample layer. This allows light transport estimation to incorporate global geometric information rather than being limited to features within a single view.

The output of the U-Net is a light transport map, which contains per-pixel light transport at each sampled light di-

rection. To render an image, we retrieve the illumination radiance on each sampled light direction, and then integrate with albedo and light transport following Eq. 3. Notice that we carry out the integration separately for the diffuse and specular components and then sum these two components to obtain the final image.

4.3. Loss Functions

We use ℓ_1 loss for the difference between rendered images and ground-truth images:

$$\mathcal{L}_{im} = \frac{1}{n} \sum_{\mathbf{x}} \|\mathbf{I}(\mathbf{x}) - \mathbf{I}_{render}(\mathbf{x})\|_1, \quad (5)$$

where n is the number of image pixels. However, with ℓ_1 loss alone, we cannot guarantee correct relighting. This is due to the ambiguity between albedo, light transport and illumination in Eq. 3: the network can overfit training images with an incorrect combination of the three components. Therefore, we need to apply additional losses to ensure the network learns a physically plausible interpretation.

Chromaticity of Light Transport To constrain the learned LTFs, we propose a novel loss on the chromaticity of light transports. For a pixel, while its light transports at different light directions differ in intensity, they usually share similar chromaticity. An exception is that for pixels with low intensities, their light transports may contain a visibility of 0 and hence do not have valid chromaticities. Therefore, we formulate a weighted chromaticity loss on light transport as:

$$\mathcal{L}_{chr} = \frac{1}{nm} \sum_{\mathbf{x}} \sum_{\omega_i} w(\mathbf{x}) (1 - \mathbf{T}'(\mathbf{x}, \omega_i, \omega_o) \cdot \mathbf{T}'_{mean}(\mathbf{x}, \omega_o)), \quad (6)$$

where m is the number of sampled light directions, $w(\mathbf{x}) = \min(20 \cdot \|\mathbf{I}(\mathbf{x})\|_2, 1)$ is a weight depending on image intensity. $\mathbf{T}'(\mathbf{x}, \omega_i, \omega_o) = \frac{\mathbf{T}(\mathbf{x}, \omega_i, \omega_o)}{\|\mathbf{T}(\mathbf{x}, \omega_i, \omega_o)\|_2}$ is the chromaticity of a light transport and $\mathbf{T}'_{mean}(\mathbf{x}, \omega_o)$ is the mean chromaticity for the light transports at pixel \mathbf{x} . We compute the loss separately for the diffuse and specular components and then sum together.

Illumination Although the initial environment map stitched from input multi-view images contains artifacts such as ghosting, the corresponding SH-based environment map is smooth and relatively accurate. Therefore, we would like to constrain our final estimated illumination to be close to the initial one within the regions that are initially covered. We first uniformly sample 4096 directions in the unit sphere and then compute loss based on the SH radiance on these directions:

$$\mathcal{L}_{illum} = \frac{1}{p} \sum_{\mathbf{p}} \sum_k \|c_k Y_k(\mathbf{p}) - c'_k Y_k(\mathbf{p})\|_1, \quad (7)$$

where p is the number of directions within initial covered regions, c_k is estimated SH coefficients and c'_k is initial SH coefficients.

Albedo From Eq. 3, we can see that there is a scale ambiguity between albedo and light transport. Hence, we include a regularization on albedo so that its mean is close to 0.5:

$$\mathcal{L}_{alb} = \frac{1}{q} \sum_{\mathbf{x}} \|\rho(\mathbf{x}) - 0.5\|_1, \quad (8)$$

where q is the number of texels. This loss is applied to both diffuse and specular albedo.

Our total loss is a weighted composition of the above losses:

$$\mathcal{L} = \mathcal{L}_{im} + \lambda_{chr} \mathcal{L}_{chr} + \lambda_{illum} \mathcal{L}_{illum} + \lambda_{alb} \mathcal{L}_{alb}. \quad (9)$$

5. Experimental Results

We implement our method in PyTorch [56]. Before training, we precompute uv map along with view direction map, normal map and per-pixel tangent space transformation matrix for each training view. We remove the parts in the initial 3D proxy that correspond to background and use a downsampled mesh with 7,500 vertices per model as input to ResGCN. For neural texture, we use 24 channels. For light direction sampling, we set the half angles of cones to $\{20^\circ, 40^\circ\}$ for the diffuse component and $\{5^\circ, 10^\circ\}$ for the specular component. We train our end-to-end network using Adam [33] as optimizer, with a learning rate of 0.001, $\beta_1 = 0.9, \beta_2 = 0.999$. We set $\lambda_{chr} = \lambda_{illum} = \lambda_{alb} = 1$ and train our models for 20k to 50k iterations based on object complexity.

5.1. Evaluations on Synthetic Data

We first evaluate RNR on synthetic data for both novel view synthesis and relighting. We choose 4 objects with different geometry complexity, for each we render 200 randomly sampled views under 2 different illuminations. We purposely set the illuminations to have different brightness and color tones. We use a physically based renderer - Tungsten [1] and render at a resolution of 1024×1024 . We further use different material configurations for each object, ranging from nearly diffuse to moderately specular, and from single material to multiple materials. Example images are shown in the first 4 rows of Fig. 4. As aforementioned, our technique cannot handle highly specular objects.

Novel View Synthesis. We compare our approach with two state-of-the-art view synthesis methods: DeepVoxels [63] and Deferred Neural Rendering (DeferredNR) [69]. The two methods and ours all require per-scene training. For each object, we randomly select 180 images under the first illumination as training views and use the rest 20 for testing. We downsample the images to 512×512 before




Figure 4. Comparisons on view synthesis. The top 4 rows are rendered synthetic data and the bottom 3 are captured real data. For each set of close-up views, the first row shows zoomed-in rgb patches while the second row shows corresponding error maps.

feeding into the three methods and set an equal batch size of 1. At each iteration, DeepVoxels takes one source view and two additional target views as input whereas DeferredNR and ours only require one view as input. For DeepVoxels, we use their implementation and default hyperparameters. For DeferredNR, we implement our own version since it is not yet open source. We increase the number of neural texture channels as well as the feature channels in the rendering network to match the number of parameters with ours. We notice slight improvements with this modification. Since our goal is to synthesize views of the object instead of the entire scene, we only compute the loss for the pix-

els on the object for all three methods. For each object, we train our network for an equal or smaller number of iterations than the other two methods. The left 4 columns in Table 1 compare the PSNR and SSIM on the test views. Our proposed method outperforms the two state-of-the-art by a noticeable margin in all cases. Qualitative comparisons, close-up views, and error maps are shown in the first 4 rows of Fig. 4. Compared to ours, DeepVoxels produces over-smoothed results whereas DeferredNR introduces higher errors near specular highlights, as shown in the 1st and 3rd rows. This illustrates the benefits of encoding the image formation model in rendering process.

Table 1. Quantitative comparisons (PSNR/SSIM) of our RNR vs. DeepVoxels [63] and DeferredNR [69] on view synthesis.

Method	Bunny	Horse	Material	Earth	Beauty	Apple	Dyck
DeepVoxels	26.67/0.86	27.98/0.89	28.92/0.93	21.00/0.75	22.05/0.81	19.39/0.75	29.75/0.94
DeferredNR	31.53/0.93	36.44/0.97	30.93/0.93	30.13/0.96	28.12/0.87	26.05/0.89	36.36/0.98
RNR (Ours)	39.08/0.98	38.48/0.98	36.18/0.98	31.39/0.97	32.82/0.97	28.29/0.93	37.62/0.99



Figure 5. Relighting results of RNR on synthetic data. The top row shows ground truth, the second row relighting results, and the bottom row error maps.

Table 2. Quantitative evaluation (PSNR/SSIM) of RNR (w/ GCN) and RNR (no GCN) on relighting synthetic scenes.

Method	Earth	Bunny	Material	Horse
w/ GCN	26.29/0.94	25.13/0.92	28.04/0.89	29.56/0.94
no GCN	25.87/0.93	24.96/0.91	27.67/0.81	28.76/0.93

Free-Viewpoint Relighting. For each object, we use the model trained under the first illumination to carry out free-viewpoint relighting, verified by using the second illumination rendered at a novel viewpoint. We compare the synthesized results with the rendered ground truth in Fig. 5. We also conduct an ablation study on the effectiveness of using GCN to augment U-Net. Table 2 shows the evaluation metrics for w/ and w/o GCN, from which we can see that using GCN leads to moderate performance improvement. We further analyze the importance of each loss in Fig. 6. Without light transport chromaticity loss or illumination loss, we observe the learnable components will overfit the training data and lead to incorrect relighting results. This illustrates the importance of our regularization terms.

Number of Training Views. To illustrate the effectiveness of RNR in encoding geometric and photometric representations, we further carry out an experiment using sparse input views (20 views in our case). Table 3 shows the PSNR and SSIM measure for view synthesis and relighting. We observe that both DeepVoxels and DeferredNR degrades drastically with sparse training views. In contrast, RNR is less affected in both tasks. This reveals the effective-



Figure 6. Ablation study on losses. On top left shows PSNR and SSIM of each case.

Table 3. Quantitative comparisons (PSNR/SSIM) of DeepVoxels [63], DeferredNR [69] and our RNR when using sparse inputs.

Method	Bunny	Earth
DeepVoxels	17.97/0.76	16.44/0.54
DeferredNR	24.38/0.81	22.10/0.86
RNR	30.87/0.94	25.32/0.91
RNR (Relight)	22.47/0.85	25.41/0.90

ness of encoding the image formation model. Specifically, compared with a black box solution, RNR can interpret object appearance following the actual physically based rendering model, thus boosting its generalization to unseen views and lighting.

5.2. Evaluations on Real Data

We have also compared RNR with DeepVoxels and DeferredNR on 3 real scenes: Beauty, Apple, Dyck. We captured the first two scenes using a handheld DSLR, with the objects positioned on a tripod. Dyck is directly adopted from DeepVoxels, captured as a video sequence. Beauty and Apple contain 151 and 144 views. For Dyck, we first remove images that contain excessive motion blurs and use the remaining 224 views. We use structure-from-motion software Agisoft Photoscan to estimate camera parameters as well as 3D proxy. Similar to synthetic data, we use 90% of the images for training and 10% for testing. The right 3 columns in Table 1 show the performance of each method, where our method performs the best for all cases. The last 3 rows of Fig. 4 show visual comparisons. DeferredNR produces similar results as RNR although RNR manages to better preserve sharp details in Beauty. Fig. 7 shows the view extrapolation results of DeferredNR and RNR. We observe that DeferredNR exhibits visual artifacts such as incorrect highlights and color blocks whereas RNR produces more coherent estimations. Please refer to the supplementary video and material for additional results.

We further apply relighting to Beauty, Apple and Dyck in Fig. 8. It is worth noting that Dyck only contains views from




Figure 7. Comparisons of our RNR vs. DeferredNR [69] on view extrapolation. On the left are two example training views.




Figure 8. Relighting results of RNR on real data.

the front, i.e., the initial illumination stitched by our method only covers a small portion of the entire environment map. Yet RNR manages to produce reasonable relighting results.

To evaluate relighting accuracy, we use an additional Pig scene from Multi-view Objects Under the Natural Illumination Database [53]. The data contains HDR images captured under 3 different illuminations, each with about 16 calibrated views. We use the images captured in “outdoor” illumination for training. Since the source images are tightly cropped at the object, we are not able to stitch the initial illumination. Hence we use the ground truth illumination in this experiment. The reconstructed geometry in [53] is not publicly available, so we use the laser-scanned mesh followed by smoothing and simplification as our 3D proxy. For testing, we synthesize with the camera parameters and illumination corresponding to a novel view under “indoor”




Figure 9. Comparisons on view synthesis and relighting using data from [53].

illumination. The rightmost column of Fig. 9 shows our synthesized results vs. [54] and the ground truth. We observe the results of RNR appear more realistic than [54], although RNR incurs inaccuracy in highlight and color. This is partially attributed to the low number of training views as well as inaccurate camera parameters provided by the dataset. The left 3 columns in Fig. 9 show that our view synthesis is also more reliable than DeferredNR.

6. Conclusions and Future Work

We have presented a new neural rendering scheme called Relightable Neural Renderer (RNR) for simultaneous view synthesis and relighting. RNR has exploited the physically based rendering process and separates appearance into environment lighting, object intrinsic attributes, and light transport function (LTF). All three components are learnable through deep networks. In particular, we have shown that by incorporating rendering constraints, our method not only enables relighting but also produces better generalization for novel view synthesis.

Our current approach cannot yet refine geometry or adaptively sample light directions. When 3D proxy contains severe artifacts, they also negatively impact rendering quality. We refer readers to supplementary material for failure cases. We also do not explicitly handle the lack of dynamic range during data capture, which may influence relighting quality. A possible way is to learn the conversion from LDR inputs to the HDR ones. In addition, RNR cannot handle highly specular objects. In the future, all-frequency lighting representations can be used in conjunction with LTF for free-viewpoint relighting of highly specular objects.

Acknowledgments

This work is supported by the National Key Research and Development Program (2018YFB2100500), the programs of NSFC (61976138 and 61977047), STCSM (2015F0203-000-06), and SHMEC (2019-01-07-00-01-E00003).

References

- [1] <https://github.com/tunabrain/tungsten>.
- [2] Kara-Ali Aliev, Dmitry Ulyanov, and Victor Lempitsky. Neural point-based graphics. *arXiv preprint arXiv:1906.08240*, 2019.
- [3] Dejan Azinovic, Tzu-Mao Li, Anton Kaplanyan, and Matthias Nießner. Inverse path tracing for joint material and lighting estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2447–2456, 2019.
- [4] Michael M Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, and Pierre Vandergheynst. Geometric deep learning: going beyond euclidean data. *IEEE Signal Processing Magazine*, 34(4):18–42, 2017.
- [5] Chris Buehler, Michael Bosse, Leonard McMillan, Steven Gortler, and Michael Cohen. Unstructured lumigraph rendering. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 425–432. ACM, 2001.
- [6] Christopher P Burgess, Loic Matthey, Nicholas Watters, Rishabh Kabra, Irina Higgins, Matt Botvinick, and Alexander Lerchner. Monet: Unsupervised scene decomposition and representation. *arXiv preprint arXiv:1901.11390*, 2019.
- [7] Joel Carranza, Christian Theobalt, Marcus A Magnor, and Hans-Peter Seidel. Free-viewpoint video of human actors. *ACM transactions on graphics (TOG)*, 22(3):569–577, 2003.
- [8] Rodrigo Ortiz Cayon, Abdelaziz Djelouah, and George Drettakis. A bayesian approach for selective image-based rendering using superpixels. In *2015 International Conference on 3D Vision*, pages 469–477. IEEE, 2015.
- [9] Gaurav Chaurasia, Sylvain Duchene, Olga Sorkine-Hornung, and George Drettakis. Depth synthesis and local warps for plausible image-based navigation. *ACM Transactions on Graphics (TOG)*, 32(3):30, 2013.
- [10] 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.
- [11] Xu Chen, Jie Song, and Otmar Hilliges. Nvs machines: Learning novel view synthesis with fine-grained view control. *arXiv preprint arXiv:1901.01880*, 2019.
- [12] Valentin Deschaintre, Miika Aittala, Fredo Durand, George Drettakis, and Adrien Bousseau. Single-image svbrdf capture with a rendering-aware deep network. *ACM Transactions on Graphics (TOG)*, 37(4):128, 2018.
- [13] Alexey Dosovitskiy, Jost Tobias Springenberg, Maxim Tatarchenko, and Thomas Brox. Learning to generate chairs, tables and cars with convolutional networks. *IEEE transactions on pattern analysis and machine intelligence*, 39(4):692–705, 2016.
- [14] SM Ali Eslami, Danilo Jimenez Rezende, Frederic Besse, Fabio Viola, Ari S Morcos, Marta Garnelo, Avraham Ruder, Andrei A Rusu, Ivo Danihelka, Karol Gregor, et al. Neural scene representation and rendering. *Science*, 360(6394):1204–1210, 2018.
- [15] Luca Falorsi, Pim de Haan, Tim R Davidson, Nicola De Cao, Maurice Weiler, Patrick Forré, and Taco S Cohen. Explorations in homeomorphic variational auto-encoding. *arXiv preprint arXiv:1807.04689*, 2018.
- [16] Duan Gao, Xiao Li, Yue Dong, Pieter Peers, Kun Xu, and Xin Tong. Deep inverse rendering for high-resolution svbrdf estimation from an arbitrary number of images. *ACM Transactions on Graphics (TOG)*, 38(4):134, 2019.
- [17] Steven J Gortler, Radek Grzeszczuk, Richard Szeliski, and Michael F Cohen. The lumigraph. In *Siggraph*, volume 96, pages 43–54, 1996.
- [18] Darya Guarnera, Giuseppe Claudio Guarnera, Abhijeet Ghosh, Cornelia Denk, and Mashhuda Glencross. Brdf representation and acquisition. In *Computer Graphics Forum*, volume 35, pages 625–650. Wiley Online Library, 2016.
- [19] Tom Haber, Christian Fuchs, Philippe Bekaer, Hans-Peter Seidel, Michael Goesele, and Hendrik PA Lensch. Relighting objects from image collections. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 627–634. IEEE, 2009.
- [20] Adam W Harley, Fangyu Li, Shrinidhi K Lakshmikanth, Xian Zhou, Hsiao-Yu Fish Tung, and Katerina Fragkiadaki. Embodied view-contrastive 3d feature learning. *arXiv preprint arXiv:1906.03764*, 2019.
- [21] Daniel Cabrini Hauagge, Scott Wehrwein, Paul Upchurch, Kavita Bala, and Noah Snavely. Reasoning about photo collections using models of outdoor illumination. In *BMVC*, 2014.
- [22] 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.
- [23] Peter Hedman, Tobias Ritschel, George Drettakis, and Gabriel Brostow. Scalable inside-out image-based rendering. *ACM Transactions on Graphics (TOG)*, 35(6):231, 2016.
- [24] Benno Heigl, Reinhard Koch, Marc Pollefeys, Joachim Denzler, and Luc Van Gool. Plenoptic modeling and rendering from image sequences taken by a hand-held camera. In *Mustererkennung 1999*, pages 94–101. Springer, 1999.
- [25] Zhuo Hui, Kalyan Sunkavalli, Joon-Young Lee, Sunil Hadap, Jian Wang, and Aswin C Sankaranarayanan. Reflectance capture using univariate sampling of brdfs. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5362–5370, 2017.
- [26] James Imber, Jean-Yves Guillemaut, and Adrian Hilton. Intrinsic textures for relightable free-viewpoint video. In *European Conference on Computer Vision*, pages 392–407. Springer, 2014.
- [27] Eldar Insafutdinov and Alexey Dosovitskiy. Unsupervised learning of shape and pose with differentiable point clouds. In *Advances in Neural Information Processing Systems*, pages 2802–2812, 2018.
- [28] Shi Jin, Ruiyong Liu, Yu Ji, Jinwei Ye, and Jingyi Yu. Learning to dodge a bullet: Concyclic view morphing via deep learning. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 218–233, 2018.
- [29] James T Kajiya. The rendering equation. In *ACM Siggraph Computer Graphics*, volume 20, pages 143–150. ACM, 1986.

- [30] 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.
- [31] Kaizhang Kang, Zimin Chen, Jiaping Wang, Kun Zhou, and Hongzhi Wu. Efficient reflectance capture using an autoencoder. *ACM Trans. Graph.*, 37(4):127–1, 2018.
- [32] Hiroharu Kato, Yoshitaka Ushiku, and Tatsuya Harada. Neural 3d mesh renderer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3907–3916, 2018.
- [33] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014.
- [34] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.
- [35] Johannes Kopf, Michael F Cohen, and Richard Szeliski. First-person hyper-lapse videos. *ACM Transactions on Graphics (TOG)*, 33(4):78, 2014.
- [36] Tejas D Kulkarni, William F Whitney, Pushmeet Kohli, and Josh Tenenbaum. Deep convolutional inverse graphics network. In *Advances in neural information processing systems*, pages 2539–2547, 2015.
- [37] 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.
- [38] Guohao Li, Matthias Muller, Ali Thabet, and Bernard Ghanem. Deepgcns: Can gcns go as deep as cnns? In *Proceedings of the IEEE International Conference on Computer Vision*, pages 9267–9276, 2019.
- [39] Guannan Li, Chenglei Wu, Carsten Stoll, Yebin Liu, Kiran Varanasi, Qionghai Dai, and Christian Theobalt. Capturing relightable human performances under general uncontrolled illumination. In *Computer Graphics Forum*, volume 32, pages 275–284. Wiley Online Library, 2013.
- [40] Tzu-Mao Li, Miika Aittala, Frédo Durand, and Jaakko Lehtinen. Differentiable monte carlo ray tracing through edge sampling. In *SIGGRAPH Asia 2018 Technical Papers*, page 222. ACM, 2018.
- [41] Xiao Li, Yue Dong, Pieter Peers, and Xin Tong. Modeling surface appearance from a single photograph using self-augmented convolutional neural networks. *ACM Transactions on Graphics (TOG)*, 36(4):45, 2017.
- [42] Zhengqin Li, Kalyan Sunkavalli, and Manmohan Chandraker. Materials for masses: Svbrdf acquisition with a single mobile phone image. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 72–87, 2018.
- [43] Chen-Hsuan Lin, Chen Kong, and Simon Lucey. Learning efficient point cloud generation for dense 3d object reconstruction. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [44] Hsueh-Ti Derek Liu, Michael Tao, and Alec Jacobson. Pappazzi: surface editing by way of multi-view image processing. *ACM Trans. Graph.*, 37(6):221–1, 2018.
- [45] Miaomiao Liu, Xuming He, and Mathieu Salzmann. Geometry-aware deep network for single-image novel view synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4616–4624, 2018.
- [46] Shichen Liu, Tianye Li, Weikai Chen, and Hao Li. Soft rasterizer: A differentiable renderer for image-based 3d reasoning. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7708–7717, 2019.
- [47] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: Learning dynamic renderable volumes from images. *ACM Transactions on Graphics (TOG)*, 38(4):65, 2019.
- [48] Matthew M Loper and Michael J Black. Opendr: An approximate differentiable renderer. In *European Conference on Computer Vision*, pages 154–169. Springer, 2014.
- [49] Moustafa Meshry, Dan B Goldman, Sameh Khamis, Hugues Hoppe, Rohit Pandey, Noah Snavely, and Ricardo Martin-Brualla. Neural rerendering in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6878–6887, 2019.
- [50] 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.
- [51] Thu H Nguyen-Phuoc, Chuan Li, Stephen Balaban, and Yongliang Yang. Rendernet: A deep convolutional network for differentiable rendering from 3d shapes. In *Advances in Neural Information Processing Systems*, pages 7891–7901, 2018.
- [52] Kyle Olszewski, Sergey Tulyakov, Oliver Woodford, Hao Li, and Linjie Luo. Transformable bottleneck networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7648–7657, 2019.
- [53] Geoffrey Oxholm and Ko Nishino. Multiview shape and reflectance from natural illumination. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2155–2162, 2014.
- [54] Geoffrey Oxholm and Ko Nishino. Shape and reflectance estimation in the wild. *IEEE transactions on pattern analysis and machine intelligence*, 38(2):376–389, 2015.
- [55] Eunbyung Park, Jimei Yang, Ersin Yumer, Duygu Ceylan, and Alexander C Berg. Transformation-grounded image generation network for novel 3d view synthesis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3500–3509, 2017.
- [56] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017.
- [57] Eric Penner and Li Zhang. Soft 3d reconstruction for view synthesis. *ACM Transactions on Graphics (TOG)*, 36(6):235, 2017.
- [58] Felix Petersen, Amit H Bermano, Oliver Deussen, and Daniel Cohen-Or. Pix2vex: Image-to-geometry reconstruction using a smooth differentiable renderer. *arXiv preprint arXiv:1903.11149*, 2019.
- [59] Julien Philip, Michaël Gharbi, Tinghui Zhou, Alexei A Efros, and George Drettakis. Multi-view relighting using a

- geometry-aware network. *ACM Transactions on Graphics (TOG)*, 38(4):78, 2019.
- [60] Qi Shan, Riley Adams, Brian Curless, Yasutaka Furukawa, and Steven M Seitz. The visual turing test for scene reconstruction. In *2013 International Conference on 3D Vision-3DV 2013*, pages 25–32. IEEE, 2013.
- [61] Aliaksandra Shysheya, Egor Zakharov, Kara-Ali Aliev, Renat Bashirov, Egor Burkov, Karim Isakov, Aleksei Ivakhnenko, Yury Malkov, Igor Pasechnik, Dmitry Ulyanov, et al. Textured neural avatars. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2387–2397, 2019.
- [62] Martin Simonovsky and Nikos Komodakis. Dynamic edge-conditioned filters in convolutional neural networks on graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3693–3702, 2017.
- [63] Vincent Sitzmann, Justus Thies, Felix Heide, Matthias Nießner, Gordon Wetzstein, and Michael Zollhofer. Deepvoxels: Learning persistent 3d feature embeddings. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2437–2446, 2019.
- [64] Vincent Sitzmann, Michael Zollhöfer, and Gordon Wetzstein. Scene representation networks: Continuous 3d-structure-aware neural scene representations. In *Advances in Neural Information Processing Systems*, pages 1119–1130, 2019.
- [65] Peter-Pike Sloan, Jan Kautz, and John Snyder. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. In *ACM Transactions on Graphics (TOG)*, volume 21, pages 527–536. ACM, 2002.
- [66] Pratul P Srinivasan, Tongzhou Wang, Ashwin Sreelal, Ravi Ramamoorthi, and Ren Ng. Learning to synthesize a 4d rgb light field from a single image. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2243–2251, 2017.
- [67] Shao-Hua Sun, Minyoung Huh, Yuan-Hong Liao, Ning Zhang, and Joseph J Lim. Multi-view to novel view: Synthesizing novel views with self-learned confidence. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 155–171, 2018.
- [68] Maxim Tatarchenko, Alexey Dosovitskiy, and Thomas Brox. Single-view to multi-view: Reconstructing unseen views with a convolutional network. *arXiv preprint arXiv:1511.06702*, 6, 2015.
- [69] Justus Thies, Michael Zollhöfer, and Matthias Nießner. Deferred neural rendering: Image synthesis using neural textures. *ACM Transactions on Graphics (TOG)*, 38(4):1–12, 2019.
- [70] Justus Thies, Michael Zollhöfer, Christian Theobalt, Marc Stamminger, and Matthias Nießner. Ignor: Image-guided neural object rendering. *arXiv preprint arXiv:1811.10720*, 2018.
- [71] Shubham Tulsiani, Tinghui Zhou, Alexei A Efros, and Jitendra Malik. Multi-view supervision for single-view reconstruction via differentiable ray consistency. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2626–2634, 2017.
- [72] Hsiao-Yu Fish Tung, Ricson Cheng, and Katerina Fragkiadaki. Learning spatial common sense with geometry-aware recurrent networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2595–2603, 2019.
- [73] Diego Valsesia, Giulia Fracastoro, and Enrico Magli. Learning localized generative models for 3d point clouds via graph convolution. In *ICLR*, 2019.
- [74] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E Sarma, Michael M Bronstein, and Justin M Solomon. Dynamic graph cnn for learning on point clouds. *ACM Transactions on Graphics (TOG)*, 38(5):146, 2019.
- [75] Michael Weinmann and Reinhard Klein. Advances in geometry and reflectance acquisition (course notes). In *SIGGRAPH Asia 2015 Courses*, page 1. ACM, 2015.
- [76] Tim Weyrich, Jason Lawrence, Hendrik PA Lensch, Szymon Rusinkiewicz, Todd Zickler, et al. Principles of appearance acquisition and representation. *Foundations and Trends® in Computer Graphics and Vision*, 4(2):75–191, 2009.
- [77] Daniel E Worrall, Stephan J Garbin, Daniyar Turmukhambetov, and Gabriel J Brostow. Interpretable transformations with encoder-decoder networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5726–5735, 2017.
- [78] 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.
- [79] 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.
- [80] Zexiang Xu, Jannik Boll Nielsen, Jiyang Yu, Henrik Wann Jensen, and Ravi Ramamoorthi. Minimal brdf sampling for two-shot near-field reflectance acquisition. *ACM Transactions on Graphics (TOG)*, 35(6):188, 2016.
- [81] Zexiang Xu, Kalyan Sunkavalli, Sunil Hadap, and Ravi Ramamoorthi. Deep image-based relighting from optimal sparse samples. *ACM Transactions on Graphics (TOG)*, 37(4):126, 2018.
- [82] Jimei Yang, Scott E Reed, Ming-Hsuan Yang, and Honglak Lee. Weakly-supervised disentangling with recurrent transformations for 3d view synthesis. In *Advances in Neural Information Processing Systems*, pages 1099–1107, 2015.
- [83] Wenjie Ye, Xiao Li, Yue Dong, Pieter Peers, and Xin Tong. Single image surface appearance modeling with self-augmented cnns and inexact supervision. In *Computer Graphics Forum*, volume 37, pages 201–211. Wiley Online Library, 2018.
- [84] Wang Yifan, Felice Serena, Shihao Wu, Cengiz Öztireli, and Olga Sorkine-Hornung. Differentiable surface splatting for point-based geometry processing. *ACM Transactions on Graphics (TOG)*, 38(6):1–14, 2019.
- [85] Tianli Yu, Hongcheng Wang, Narendra Ahuja, and Wei-Chao Chen. Sparse lumigraph relighting by illumination and reflectance estimation from multi-view images. In *ACM SIGGRAPH 2006 Sketches*, page 175. ACM, 2006.

- [86] Ke Colin Zheng, Alex Colburn, Aseem Agarwala, Maneesh Agrawala, David Salesin, Brian Curless, and Michael F Cohen. Parallax photography: creating 3d cinematic effects from stills. In *Proceedings of Graphics Interface 2009*, pages 111–118. Canadian Information Processing Society, 2009.
- [87] 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.
- [88] Tinghui Zhou, Shubham Tulsiani, Weilun Sun, Jitendra Malik, and Alexei A Efros. View synthesis by appearance flow. In *European conference on computer vision*, pages 286–301. Springer, 2016.
- [89] 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.
- [90] Hao Zhu, Hao Su, Peng Wang, Xun Cao, and Ruigang Yang. View extrapolation of human body from a single image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4450–4459, 2018.