

Differentiable Point-based Inverse Rendering

Hoon-Gyu Chung Seokjun Choi Seung-Hwan Baek
 POSTECH

Abstract

We present differentiable point-based inverse rendering, DPIR, an analysis-by-synthesis method that processes images captured under diverse illuminations to estimate shape and spatially-varying BRDF. To this end, we adopt point-based rendering, eliminating the need for multiple samplings per ray, typical of volumetric rendering, thus significantly enhancing the speed of inverse rendering. To realize this idea, we devise a hybrid point-volumetric representation for geometry and a regularized basis-BRDF representation for reflectance. The hybrid geometric representation enables fast rendering through point-based splatting while retaining the geometric details and stability inherent to SDF-based representations. The regularized basis-BRDF mitigates the ill-posedness of inverse rendering stemming from limited light-view angular samples. We also propose an efficient shadow detection method using point-based shadow map rendering. Our extensive evaluations demonstrate that DPIR outperforms prior works in terms of reconstruction accuracy, computational efficiency, and memory footprint. Furthermore, our explicit point-based representation and rendering enables intuitive geometry and reflectance editing.

1. Introduction

Inverse rendering aims to estimate geometry and reflectance of real-world objects from a set of images, with applications including relighting, augmented and virtual reality, and object digitization. While it has been a longstanding challenge in computer vision and graphics, inverse rendering is still notoriously difficult. A common approach to solve inverse rendering is to take the analysis-by-synthesis principle. That is, synthesized images using forward rendering are compared to input captured images, and through the iterative minimization of the difference between these images, geometry and reflectance are updated. Hence, the forward rendering plays a critical role in the analysis-by-synthesis inverse rendering framework.

Mesh-based rendering methods combined with path tracing are prevalent for photo-realistic forward rendering. However, they often struggle in inverse rendering due to

discontinuity issues [24]. Volumetric rendering has recently been recognized for view synthesis, particularly in variants of NeRFs [26, 39, 40]. Inspired by such success, many state-of-the-art inverse rendering methods have adopted volumetric rendering as a forward-rendering engine, demonstrating compelling reconstruction of geometry and reflectance [11, 37, 46, 50, 53]. However, challenges remain, such as the excessive computational costs due to multiple sampling for each ray, and ambiguities in BRDF and normal reconstructions given limited light-view angular samples.

Point-based rendering has been extensively studied for forward rendering, where points serve as compact scene representations [12, 14, 33, 42]. Point-based forward rendering involves splatting these points onto a viewpoint using circular disks, ellipsoids, or surfels [30, 59–61]. Point-based rendering, especially in the analysis-by-synthesis framework, have recently gained renewed interest [14, 15, 33, 42, 52, 54]: compared to volumetric rendering, point-based rendering is more efficient since it avoids the need for multiple samplings per ray. Recent methods proposed by Kopanas et al.[13, 14], Point-NeRF[45], Zhang et al.[52], and Kerbl et al.[12] have shown the effectiveness of point-based rendering in modeling radiance fields for novel-view synthesis.

Inspired by this success, we present differentiable point-based inverse rendering, DPIR, that exploits point-based forward rendering for inverse rendering. DPIR processes either multi-view multi-light images or photometric images captured by multi-view photometric setups [19] and flash photography, respectively. Using point-based forward rendering confronts several challenges critical for inverse rendering: (1) The discrete point representation hinders the reconstruction of both smooth and detailed surface. (2) The inherent difficulty in jointly reconstructing geometry and spatially-varying BRDF from limited light-view samples remains. (3) There is a need for efficient shadow consideration to ensure precise inverse rendering.

To address these challenges, we develop a hybrid point-volumetric geometry representation and a regularized basis-BRDF representation. The hybrid geometric representation enjoys the benefits of both point-based and volumet-

ric geometry representations, ensuring efficient rendering through point-based splatting, while benefiting from the intricate and stable modeling capability inherent to signed-distance functions (SDFs). Note that points are our only rendering primitives. The regularized basis-BRDF representation consists of a per-position diffuse albedo and a specular reflectance component. The specular component is modeled with weighted specular-basis BRDFs by exploiting the spatial coherency of specular reflectance. This allows for overcoming the ill-posedness of inverse rendering under limited light-view angular samples. Also, we tackle the shadow detection problem through an efficient point-based shadow detection method, bypassing the volumetric integration often used in learning-based inverse rendering methods [37, 46, 53, 55]. DPIR then jointly optimizes point locations, point radii, diffuse albedo, specular basis BRDFs, specular coefficients, and SDF, by leveraging point-based splatting as a forward renderer in the analysis-by-synthesis framework. Our extensive evaluations show that DPIR outperforms previous state-of-the-art inverse rendering methods [11, 46, 50, 51], in accuracy, training speed, and memory footprint. Furthermore, the explicit point representation and rendering of DPIR enable convenient scene editing.

In summary, our contributions are:

- DPIR, an inverse rendering method that exploits differentiable point-based rendering as a forward renderer, outperforming prior methods in reconstruction accuracy, training speed, and memory footprint.
- Hybrid point-volumetric geometry representation and a regularized basis-BRDF representation that enable efficient and high-quality optimization of geometry and spatially-varying BRDFs through point-based rendering.
- Point-based shadow detection method merged into our inverse rendering framework, enabling efficient and accurate shadow detection for inverse rendering.

2. Related Work

Inverse Rendering Inverse rendering is a long-standing problem in computer vision and graphics [2, 5, 6, 9, 17, 28, 35, 44, 57]. Learning-based single-image inverse rendering enables accurate reconstruction for planar samples [4, 20–22, 34, 36, 41, 49]. For non-planar objects, Li et al.[18] proposed an implicit neural representation that estimates geometry and reflectance from single-view multi-light images, albeit suffering from depth ambiguity. For multi-view inputs captured under a constant lighting, learning-based inverse rendering methods [37, 53] have recently shown promising results by exploiting volumetric rendering as in NeRF [26]. Incorporating a signed distance function (SDF) as a geometric representation into volumetric rendering further improves normal-reconstruction quality [23, 50, 51, 55]. However, these methods struggle with low reconstruction accuracy of spatially-varying reflectance and excessive training

time due to the multiple sampling per each ray in the volumetric rendering process. Monte Carlo differentiable rendering can be also used for inverse rendering, however it suffers from complex scene geometry [3, 8, 25] and limited light-view angular samples [8, 27, 38], due to Monte Carlo noise and mesh representations.

Using multi-view and multi-light inputs improves reconstruction quality as demonstrated in PS-NeRF [46]. However, volumetric rendering used in PS-NeRF still results in excessive training time, which is further complicated by its multi-stage pipeline that requires a pretrained photometric stereo network. MVPSNet [56] demonstrated faster training time by extracting lighting features for multi-view photometric stereo. However, only geometry is estimated that cannot be used for inverse rendering applications where reflectance is necessary.

Our DPIR provides highest reconstruction quality, fastest training time, and lowest memory footprint. To this end, we use points as rendering primitives, resulting in fast training with splatting-based rendering. To handle the limited light-view angular samples, we gather multiple samples from points with similar specular appearance by using regularized basis BRDFs. Also, we handle shadow using efficient point-based shadow rendering.

Point-based Rendering Point-based rendering uses points as a compact scene representation for rendering [7]. Splatting with circular disks, ellipsoids, or surfels [30, 59–61] enables high-quality and efficient point-based rendering. To model geometry and outgoing radiance distribution of a scene, various point-based methods have been proposed. Combining point-based rendering with image-refinement neural networks enables high-quality novel view synthesis [1, 13–15, 29, 31, 33, 42, 58]. Recently, Kopanas et al.[13] achieve high-quality novel-view synthesis of highly specular surfaces using point-based rendering and explicitly modeling reflection. Zhang et al.[52] employ efficient point-based rendering and spherical-harmonics point radiance for efficient novel view synthesis, without using refinement neural networks. Kerbl et al.[12] use anisotropic Gaussians and tile-based optimization method for efficient and accurate novel-view synthesis of in-the-wild scenes. Yifan et al.[48] employ geometric regularizers that lead to high-quality geometry reconstruction with point-based rendering.

In summary, point-based rendering has proven itself as a promising technique that has fast rendering speed, low memory footprint, and high rendering quality. In this paper, we combine the point-based rendering with the analysis-by-synthesis inverse rendering framework. By resolving associated challenges, DPIR demonstrates high-quality and efficient inverse rendering.

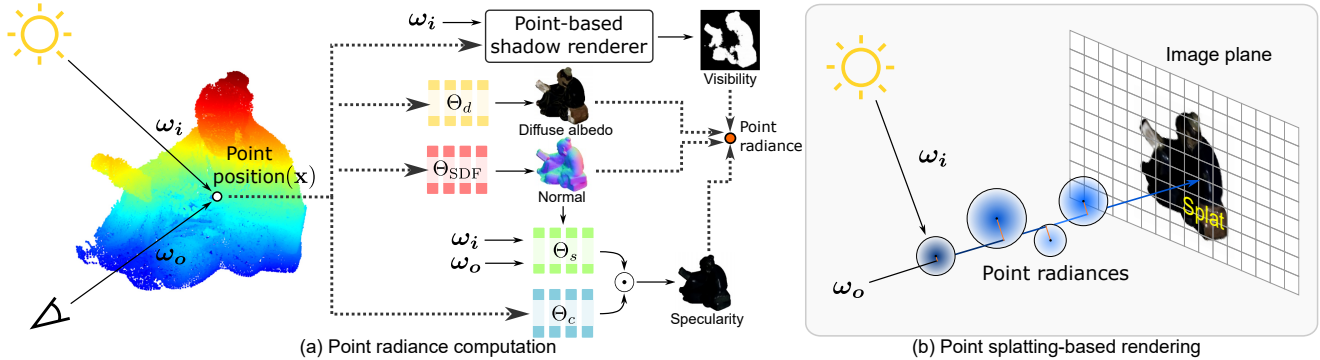


Figure 1. **Overview of differentiable forward rendering.** (a) For each 3D point, its position is used as a query for the diffuse-albedo MLP Θ_d , SDF MLP Θ_{SDF} , and specular-basis coefficient MLP Θ_c . The specular-basis BRDF MLP Θ_s models specular-basis reflectance, given the incident and outgoing directions ω_i and ω_o . The point-based shadow renderer estimates the point visibility from a light source per each image. By using the diffuse albedo, normals, specular reflectance, and visibility, we compute the radiance for each point. (b) The radiance is then projected onto a camera plane to render the pixel color through splatting-based differentiable forward rendering.

3. Method

The proposed DPIR estimates points, surface normals, spatially-varying BRDF, and visibility from the images of a static object captured either by flash photography [50] or multi-view multi-light imaging [46]. We use object masks as additional inputs similar to NeuS [40] and Zhang et al. [52]. DPIR compares rendered images with the input images to update our scene representation. Figure 1 shows a detailed overview of our differentiable forward rendering process, which is pivotal for achieving efficient and high-quality inverse rendering.

3.1. Scene Representation

Hybrid Point-volumetric Geometric Representation

We use a set of 3D points, where each point possesses two parameters: position $\mathbf{x} \in \mathbb{R}^{3 \times 1}$ and radius $r \in \mathbb{R}^{1 \times 1}$. Using the points as geometric primitives allows for fast splatting-based rendering by bypassing the per-ray integration used in volumetric rendering. However, using the points only often makes its surface normals noisy as the discrete points are non-uniformly distributed through space. To represent both detailed and smooth geometry with accurate surface normals \mathbf{n} for each point at position \mathbf{x} , we use SDF, represented as a coordinate-based MLP, Θ_{SDF} , as follows:

$$\mathbf{n} = \nabla_{\mathbf{x}} \Theta_{\text{SDF}}(\mathbf{x}), \quad (1)$$

where $\nabla_{\mathbf{x}}$ is the differentiation operator [40, 47].

Our hybrid point-volumetric representation allows *discrete points* to move in space while the surface normals of the points can be sampled from the *continuous SDF*.

Regularized Basis BRDF Representation Reconstructing per-point BRDF from limited light-view angular samples is an ill-posed problem. Thus, we propose to exploit




Figure 2. **Estimated spatially-varying BRDFs.** DPIR accurately reconstruct the BRDFs of the specular gold appearance and the red diffuse appearance. We visualize the BRDFs on unit spheres illuminated by a point light source.

spatial coherency of specular reflectance by using the basis BRDF representation [16, 18, 28]. Specifically, we use three MLPs for diffuse albedo, specular coefficients, and specular basis BRDFs, describing the total reflectance f_r as

$$f_r(\omega_o, \omega_i, \mathbf{x}, \mathbf{n}) = \Theta_d(\mathbf{x}) + \Theta_s(\mathbf{h}, \mathbf{n}) \Theta_c(\mathbf{x}), \quad (2)$$

where Θ_d , Θ_s , and Θ_c are the MLPs for the diffuse albedo, specular basis, and regularized specular-basis coefficients, respectively. The outputs of the MLPs are in the following dimensions: $\Theta_d(\mathbf{x}) \in \mathbb{R}^{3 \times 1}$, $\Theta_s(\mathbf{h}, \mathbf{n}) \in \mathbb{R}^{3 \times K}$, $\Theta_c(\mathbf{x}) \in \mathbb{R}^{K \times 1}$. K is the number of basis BRDFs. ω_i and ω_o are the incident and outgoing light directions. \mathbf{h} is the half-way vector: $\mathbf{h} = (\omega_i + \omega_o) / \|\omega_i + \omega_o\|$.

Unlike previous inverse rendering methods utilizing basis BRDFs [18, 28, 46], we found that enforcing a positive constraint in the *specular basis coefficients* and optimizing under an l_1 -norm with a lower bound of ϵ significantly enhances the accuracy of reflectance estimation. We set $\epsilon = 1$ for highly glossy objects, $\epsilon = 0.5$ otherwise. Figure 2 shows the learned BRDFs consisting of per-point diffuse albedo and regularized specular basis BRDFs.

3.2. Point-based Visibility Test

Identifying shadow on which illumination does not reach to a point is critical for robust inverse rendering [11, 37, 46, 53,




Figure 3. **Point-based visibility test.** To determine the visibility of each point, we compute the depth using the z-buffer from a virtual camera positioned at the light source, and compare the depth with the distance from each point to the virtual camera.

55]. Departing from computationally extensive volumetric-rendering approaches for visibility test in recent learning-based methods [18, 46, 53], we use a simple method compatible with our point representation: a shadow map technique used in rasterization-based graphics [10, 32, 43]. Figure 3 shows the schematics of our visibility test. For each input image and a known light source, we place a virtual orthographic camera at the location of the directional light source. We then perform splatting-based rendering of the points to the virtual camera. The resulting z-buffer stores the depth z of each point with respect to the virtual camera. The visibility function f_v then can be estimated from the z-buffer as

$$f_v(\omega_i, \mathbf{x}) = \sigma(\tau + z_0(\mathbf{x}) - z(\mathbf{x})), \quad (3)$$

where $\sigma(\cdot)$ is the step function that returns 1, meaning visible from the light source if the input is positive, and 0, which means invisible, otherwise. $z_0(\mathbf{x})$ denotes the depth value of the first intersection point and τ is the threshold, set to be 0.1 in our experiments. For photometric dataset, visibility of every point is set to 1 as the light source and the camera are co-located.

Our shadow handling method efficiently and accurately compute visibility of each point with point-based rasterization, which brings benefits over existing learning-based shadow estimation that requires computationally expensive volumetric ray marching [18].

3.3. Optimization via Point Splatting

Point-based Rendering For rendering each image, we use the point location \mathbf{x} , normals \mathbf{n} sampled from the SDF, reflectance from the learned BRDFs $f_r(\cdot)$, and visibility from the point-based test $f_v(\cdot)$. We compute the radiance R for each point \mathbf{x} along the direction ω_o towards a camera

pixel as follows:

$$R(\omega_o, \mathbf{x}) = L(\omega_i, \mathbf{x}) f_v(\omega_i, \mathbf{x}) f_r(\omega_o, \omega_i, \mathbf{x}, \mathbf{n})(\omega_i \cdot \mathbf{n}), \quad (4)$$

where $L(\omega_i, \mathbf{x})$ is the incident radiance from the direction ω_i to the point \mathbf{x} . Equation (4) is a simplified rendering equation by exploiting that a scene point in each input image is contributed from a single incident light ray, which is valid for both multi-view multi-light images [19] and photometric images. Note that extending the rendering equation to handle multiple incident light rays is feasible by simply integrating the computed radiance, which we demonstrate in the environment-map rendering results.

Once the radiance of each point is computed, we project the points to the camera viewpoints from which images are captured, using point splatting. The rendered pixel intensity of a camera pixel u amounts to the results of α -blending of the projected point radiance R :

$$I(u) = \sum_{i=1}^N R_i \alpha_i(u) \prod_{j=1}^{i-1} (1 - \alpha_j(u)), \quad (5)$$

$$\alpha_i(u) = 1 - \frac{(p_i - u)^2}{r_i^2}, \quad (6)$$

where N is the number of points, R_i is the radiance of i -th point computed from Equation (4), p_i is the projected pixel location, and r_i is the radius of the point. α_i is the corresponding weight. Points are sorted according to their distances to the camera. This point-based rendering equation, also depicted in Figure 1, can be efficiently evaluated, making DPIP outperform previous volumetric inverse rendering methods in efficiency.

Optimization We propose to optimize point positions \mathbf{x} , point radii r , and MLPs for SDF Θ_{SDF} , diffuse albedo Θ_d , specular coefficients Θ_c , and specular-basis BRDFs Θ_s by minimizing the following loss:

$$\mathcal{L}_2 + \lambda_{\text{ssim}} \mathcal{L}_{\text{ssim}} + \lambda_{\text{SDF}} \mathcal{L}_{\text{SDF}} + \lambda_c \mathcal{L}_c + \lambda_m \mathcal{L}_m, \quad (7)$$

where \mathcal{L}_2 and $\mathcal{L}_{\text{ssim}}$ are the l_2 loss and the SSIM loss for the rendered image I and the captured image I' .

\mathcal{L}_{SDF} promotes the zero-level set of the SDF lies near the point positions: $\mathcal{L}_{\text{SDF}} = \|\Theta_{\text{SDF}}(\mathbf{x})\|^2$. \mathcal{L}_c regularizes l_1 -norm of per point specular coefficient to be ϵ : $\mathcal{L}_c = \|\Theta_c(\mathbf{x})\|_1 - \epsilon$. \mathcal{L}_m is defined as the l_2 loss between the estimated mask and the ground-truth mask, aligning the rendered mask from the points to the input masks. For stable optimization, we employ the mask-based initialization of the point locations, pruning, and upsampling schemes following Zhang et al. [52]. We refer to the Supplemental Document for the details of the loss functions and optimization techniques.




Figure 4. **Reconstruction results.** DPIR enables accurate and efficient reconstruction of geometry and reflectance for various objects.

4. Results

DPIR enables accurate and efficient reconstruction of geometry and reflectance as demonstrated in Figure 4. In the following, we evaluate DPIR compared to previous methods and provide detailed ablation results.

4.1. Comparison

Multi-view Multi-light Images We evaluate DPIR on DiLiGenT-MV [19], a multi-view multi-light dataset, compared to state-of-the-art neural inverse rendering methods: PhysSG [51], TensoIR [11] and PS-NeRF [46]. Table 1 and Table 2 show quantitative evaluations of novel-view relight-

Config.	Method	Bear			Buddha			Cow			Pot2			Reading		
		PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
Single light	PhySG	24.52	0.9590	0.041	20.92	0.9229	0.0624	23.64	0.9604	0.0342	25.21	0.9609	0.0241	19.05	0.9056	0.0817
	TensoIR	24.81	0.9597	0.0510	25.40	0.9521	0.0374	27.09	0.9766	0.0260	26.89	0.9741	0.0309	27.02	0.9607	0.0277
	PS-NeRF	35.19	0.9925	0.0061	32.85	0.9783	0.0074	36.57	0.9942	0.0026	38.40	0.9943	0.0019	33.73	0.9792	0.0077
	Ours	43.21	0.9944	0.0037	37.62	0.9858	0.0048	38.23	0.9946	0.0022	39.32	0.9951	0.0022	35.75	0.9843	0.0063
Multi light	PS-NeRF	34.27	0.9802	0.0127	31.58	0.9637	0.0114	36.03	0.9871	0.0066	37.76	0.9851	0.0041	31.16	0.9736	0.0202
	Ours	39.78	0.9821	0.0083	34.88	0.9726	0.0090	37.64	0.9890	0.0041	38.86	0.9885	0.0034	32.50	0.9788	0.0204

Table 1. Quantitative comparison of novel-view rendering and relighting on DiLiGenT-MV dataset.

Method	Bear MAE \downarrow	Buddha MAE \downarrow	Cow MAE \downarrow	Pot2 MAE \downarrow	Reading MAE \downarrow	Train time \downarrow	Mem. \downarrow
PhySG	11.35	27.20	16.10	11.98	26.67	20h	13MB
TensoIR	35.93	37.54	30.10	25.68	33.67	6h	68MB
PS-NeRF	4.68	11.92	5.89	7.55	10.64	22h	40MB
Ours	4.35	11.10	4.61	6.71	9.03	2h	5MB

Table 2. Quantitative comparison of normal accuracy, average training time, and memory footprint on DiLiGenT-MV dataset.




Figure 5. Comparison of novel view rendering and estimated normal on DiLiGenT-MV dataset. Our DPIR recovers detailed surface normals and reproduces accurate appearance.

ing, normal accuracy, training time, and memory footprint. DPIR not only outperforms the compared methods in rendering and normal accuracy, but also offers $10\times$ faster training and $8\times$ lower memory footprints than PS-NeRF, which is the only competitive method in rendering accuracy. Figure 5 shows novel-view rendering images and estimated normals. Note that DPIR does not require using any pre-trained network contrary to PS-NeRF that uses a pre-trained photometric-stereo network.




Figure 6. Comparison of novel-view flashlight relighting. DPIR successfully reproduces the ground-truth appearance, outperforming IRON.

	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	MAE \downarrow	Train \downarrow	Mem. \downarrow
IRON	31.91	0.9557	0.0446	9.38	12h	28MB
Ours	35.56	0.9734	0.0285	8.74	2h	5MB

Table 3. Quantitative comparison of novel-view relighting, normal accuracy, training time, and memory footprint on the photometric dataset.

For the above experiments, we have two comparison configurations. First, since PhySG and TensoIR assume constant environment illumination, we use the multi-light averaged image for each view to simulate a virtual environment map [46]. PhySG and TensoIR are trained on 15 views lighting-averaged images and tested on 5 novel views under the same environment map. As PS-NeRF and our DPIR take varying illumination images, we use 15 views and 16 lightings for training and use 5 views and 96 lightings for testing. For the comparison with PhySG and TensoIR, we compute the average image across 96 lightings for PS-NeRF and DPIR on each testing view. This comparison protocol is adopted from the PS-NeRF [46]. Second, we make a comparison only between PS-NeRF and DPIR to evaluate true novel-view relighting. That is, we directly compare the rendered images under a novel view and lighting without any averaging. Our DPIR achieves the best performance over the previous methods in both configurations.

Photometric Images We then evaluate DPIR on a photometric dataset, which we render using four objects at 300




Figure 7. **Importance of point-based visibility test.** Our efficient point-based visibility test enables accurate normal reconstruction for occluded regions.




Figure 8. **Impact of hybrid shape representation.** Using points only as shape representation often results in inaccurate reconstruction. Our hybrid shape representation based on points and a SDF allows for recovering detailed appearance. We recommend viewing this figure on a monitor.

views with co-located point lights, following the configuration of mobile flash photography [28]. 200/100 views are used for training/testing, respectively. Table 3 and Figure 6 show that DPIR outperforms IRON [50], the-state-of-the-art inverse rendering method for photometric images, with higher novel-view relighting quality, $6\times$ faster training, and $5\times$ lower memory footprint. IRON is a two-stage method that uses volumetric and surface renderings for each stage, resulting in not only lengthy training time, but also smoothed geometry and saturated reflectance. DPIR is based on one-stage point-based analysis-by-synthesis framework that directly optimizes geometry and reflectance, enabling fast training and accurate reconstruction.

4.2. Ablation Study

Point-based Shadow Detection The point scene representation and rendering of DPIR allows for efficient and accurate shadow estimation, which is crucial for reconstructing objects with complex geometry. Our visibility test only requires $0.1\times$ additional training time. Figure 7 shows that our point-based shadow detection leads to notable improve-




Figure 9. **Impact of basis-BRDF regularization.** Without using the regularization, diffuse appearance is baked into the specular image. Regularized specular coefficients enable accurate appearance decomposition.




Figure 10. **Impact of point radius optimization.** Learning point radius allows for reconstructing geometry with both low- and high-frequency geometric features.

ments of surface normal accuracy for the occluded region.

Hybrid Shape Representation DPIR utilizes the hybrid point-volumetric shape representation. Figure 8 shows that only using point representation leads to unstable reconstruction. Our hybrid point-volumetric representation improves reconstruction quality by learning delicate geometry and appearance of both smooth and detailed surfaces.

Regularization for Specular Basis Coefficients DPIR constrains specular coefficients with the regularization loss. Figure 9 shows that using the regularizer leads to accurate appearance decomposition between diffuse and specular components. Without using the regularizer, diffuse appearance is baked into the estimated specular image.

Dynamic Point Radius DPIR learns not only the location of points but also their radii, used for α -blending in Equation (5). Figure 10 show that learning the point radii enables accurate geometry reconstruction by adjusting the point size to the spatial frequencies of a target scene.




Figure 11. **Ablation on mask inputs.** Using mask inputs improves reconstruction quality while DPIR can obtain plausible results even without mask inputs. Without using mask, optimization results near surface suffers from reconstruction artifacts.

Dependency on Masks DPIR uses mask inputs to initialize the point locations and compute the mask loss. Mask-based point initialization provides coarse point geometry as a starting step and the mask loss provides further shape guidance for low-light intensity region. In Figure 11, we demonstrate that using mask inputs improves reconstruction quality. However, DPIR often achieves plausible reconstruction even without using the masks that show a potential application of DPIR for larger-scale scenes.

4.3. Applications

The explicit point representation and rendering of DPIR enable convenient scene editing by simply editing points and its attributes. Figure 12 shows four applications enabled by DPIR. We swap the diffuse and specular basis BRDF of an object with that of another object we scanned. Also, simply editing the points from the reconstruction allows for intuitive geometric editing as we demonstrated in the geometry removal and the object merging. Lastly, DPIR supports environment-map rendering by accumulating reflected radiance for each light source in the environment map. Environment-map rendering time is linearly proportional to the number of illumination sources.

5. Conclusion

In this paper, we introduced DPIR, a point-based inverse rendering method that integrates efficient differentiable point-splatting forward rendering into the analysis-by-synthesis inverse rendering framework. To achieve this,




Figure 12. **Applications using point-based scene representation rendering.** DPIR enables convenient scene editing. We demonstrate reflectance change, geometry removal, object merging, and environment relighting.

we have developed a hybrid point-volumetric geometry representation, introduced regularized basis BRDFs, and used a point-based visibility detection method. DPIR jointly optimizes the point locations, radii, surface normals, and reflectance in a single stage without using any pre-trained network. Through evaluations, we demonstrate DPIR outperforms state-of-the-art inverse rendering methods in accuracy, training speed, and memory footprint.

Limitations Many interesting open questions remain. First, DPIR only models direct reflection without considering global-illumination effects such as inter-reflections. Modeling the global-illumination effects with an efficient global rendering method would be an interesting future work. Second, DPIR does not explicitly model transmission. Generalizing inverse rendering with bidirectional scattering distribution functions may open up the applicability of DPIR for more diverse materials.

Acknowledgements This work was supported by Korea NRF (RS-2023-00211658, 2022R1A6A1A03052954), Samsung Advanced Institute of Technology, Samsung Research Funding & Incubation Center for Future Technology grant (SRFC-IT1801-52), Samsung Electronics, and Korea IITP MSIT (RS-2023-00215700).

References

- [1] Kara-Ali Aliev, Artem Sevastopolsky, Maria Kolos, Dmitry Ulyanov, and Victor Lempitsky. Neural point-based graphics. In *Eur. Conf. Comput. Vis.*, pages 696–712. Springer, 2020. [2](#)
- [2] Jonathan T Barron and Jitendra Malik. Shape, illumination, and reflectance from shading. *IEEE Trans. Pattern Anal. Mach. Intell.*, 37(8):1670–1687, 2014. [2](#)
- [3] Guangyan Cai, Kai Yan, Zhao Dong, Ioannis Gkioulekas, and Shuang Zhao. Physics-based inverse rendering using combined implicit and explicit geometries. In *Comput. Graph. Forum*, volume 41, pages 129–138. Wiley Online Library, 2022. [2](#)
- [4] Wenzheng Chen, Joey Litalien, Jun Gao, Zian Wang, Clement Fuji Tsang, Sameh Khamis, Or Litany, and Sanja Fidler. Dib-r++: learning to predict lighting and material with a hybrid differentiable renderer. *Adv. Neural Inform. Process. Syst.*, 34:22834–22848, 2021. [2](#)
- [5] Yue Dong, Guojun Chen, Pieter Peers, Jiawan Zhang, and Xin Tong. Appearance-from-motion: Recovering spatially varying surface reflectance under unknown lighting. *ACM Trans. Graph.*, 33(6):1–12, 2014. [2](#)
- [6] Dan B Goldman, Brian Curless, Aaron Hertzmann, and Steven M Seitz. Shape and spatially-varying brdfs from photometric stereo. *IEEE Trans. Pattern Anal. Mach. Intell.*, 32(6):1060–1071, 2009. [2](#)
- [7] Jeffrey P Grossman and William J Dally. Point sample rendering. In *Rendering Techniques' 98: Proceedings of the Eurographics Workshop in Vienna, Austria, June 29–July 1, 1998* 9, pages 181–192. Springer, 1998. [2](#)
- [8] Jon Hasselgren, Nikolai Hofmann, and Jacob Munkberg. Shape, light, and material decomposition from images using monte carlo rendering and denoising. *Adv. Neural Inform. Process. Syst.*, 35:22856–22869, 2022. [2](#)
- [9] Carlos Hernandez, George Vogiatzis, and Roberto Cipolla. Multiview photometric stereo. *IEEE Trans. Pattern Anal. Mach. Intell.*, 30(3):548–554, 2008. [2](#)
- [10] J Charles Hourcade and Alain Nicolas. Algorithms for antialiased cast shadows. *Computers & Graphics*, 9(3):259–265, 1985. [4](#)
- [11] Haian Jin, Isabella Liu, Peijia Xu, Xiaoshuai Zhang, Songfang Han, Sai Bi, Xiaowei Zhou, Zexiang Xu, and Hao Su. Tensor: Tensorial inverse rendering. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 165–174, 2023. [1](#), [2](#), [3](#), [5](#)
- [12] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):1–14, 2023. [1](#), [2](#)
- [13] Georgios Kopanas, Thomas Leimkühler, Gilles Rainer, Clément Jambon, and George Drettakis. Neural point catacaustics for novel-view synthesis of reflections. *ACM Trans. Graph.*, 41(6):1–15, 2022. [1](#), [2](#)
- [14] Georgios Kopanas, Julien Philip, Thomas Leimkühler, and George Drettakis. Point-based neural rendering with per-view optimization. In *Comput. Graph. Forum*, volume 40, pages 29–43. Wiley Online Library, 2021. [1](#)
- [15] Christoph Lassner and Michael Zollhofer. Pulsar: Efficient sphere-based neural rendering. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1440–1449, 2021. [1](#), [2](#)
- [16] Jason Lawrence, Aner Ben-Artzi, Christopher DeCoro, Wojciech Matusik, Hanspeter Pfister, Ravi Ramamoorthi, and Szymon Rusinkiewicz. Inverse shade trees for non-parametric material representation and editing. *ACM Trans. Graph.*, 25(3):735–745, 2006. [3](#)
- [17] Jason Lawrence, Szymon Rusinkiewicz, and Ravi Ramamoorthi. Efficient brdf importance sampling using a factored representation. *ACM Trans. Graph.*, 23(3):496–505, 2004. [2](#)
- [18] Junxuan Li and Hongdong Li. Neural reflectance for shape recovery with shadow handling. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 16221–16230, 2022. [2](#), [3](#), [4](#)
- [19] Min Li, Zhenglong Zhou, Zhe Wu, Boxin Shi, Changyu Diao, and Ping Tan. Multi-view photometric stereo: A robust solution and benchmark dataset for spatially varying isotropic materials. *IEEE Trans. Image Process.*, 29:4159–4173, 2020. [1](#), [4](#), [5](#)
- [20] Zhengqin Li, Mohammad Shafiei, Ravi Ramamoorthi, Kalyan Sunkavalli, and Manmohan Chandraker. Inverse rendering for complex indoor scenes: Shape, spatially-varying lighting and svbrdf from a single image. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2475–2484, 2020. [2](#)
- [21] Zhengqin Li, Zexiang Xu, Ravi Ramamoorthi, Kalyan Sunkavalli, and Manmohan Chandraker. Learning to reconstruct shape and spatially-varying reflectance from a single image. *ACM Trans. Graph.*, 37(6):1–11, 2018.
- [22] Daniel Lichy, Jiaye Wu, Soumyadip Sengupta, and David W Jacobs. Shape and material capture at home. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 6123–6133, 2021. [2](#)
- [23] Yuan Liu, Peng Wang, Cheng Lin, Xiaoxiao Long, Jiepeng Wang, Lingjie Liu, Taku Komura, and Wenping Wang. Nero: Neural geometry and brdf reconstruction of reflective objects from multiview images. *arXiv preprint arXiv:2305.17398*, 2023. [2](#)
- [24] Guillaume Loubet, Nicolas Holzschuch, and Wenzel Jakob. Reparameterizing discontinuous integrands for differentiable rendering. *ACM Trans. Graph.*, 38(6):1–14, 2019. [1](#)
- [25] Fujun Luan, Shuang Zhao, Kavita Bala, and Zhao Dong. Unified shape and svbrdf recovery using differentiable monte carlo rendering. In *Comput. Graph. Forum*, volume 40, pages 101–113. Wiley Online Library, 2021. [2](#)
- [26] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. [1](#), [2](#)
- [27] Jacob Munkberg, Jon Hasselgren, Tianchang Shen, Jun Gao, Wenzheng Chen, Alex Evans, Thomas Müller, and Sanja Fidler. Extracting triangular 3d models, materials, and lighting from images. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 8280–8290, 2022. [2](#)
- [28] Giljoo Nam, Joo Ho Lee, Diego Gutierrez, and Min H Kim. Practical svbrdf acquisition of 3d objects with unstructured flash photography. *ACM Trans. Graph.*, 37(6):1–12, 2018. [2](#), [3](#), [7](#)
- [29] Julian Ost, Issam Laradji, Alejandro Newell, Yuval Bahat, and Felix Heide. Neural point light fields. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 18419–18429, 2022. [2](#)
- [30] Hanspeter Pfister, Matthias Zwicker, Jeroen Van Baar, and

- Markus Gross. Surfels: Surface elements as rendering primitives. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 335–342, 2000. 1, 2
- [31] Ruslan Rakhimov, Andrei-Timotei Ardelean, Victor Lempitsky, and Evgeny Burnaev. Npbg++: Accelerating neural point-based graphics. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 15969–15979, 2022. 2
- [32] William T Reeves, David H Salesin, and Robert L Cook. Rendering antialiased shadows with depth maps. In *Proceedings of the 14th annual conference on Computer graphics and interactive techniques*, pages 283–291, 1987. 4
- [33] Darius Rückert, Linus Franke, and Marc Stamminger. Adop: Approximate differentiable one-pixel point rendering. *ACM Trans. Graph.*, 41(4):1–14, 2022. 1, 2
- [34] Shen Sang and Manmohan Chandraker. Single-shot neural relighting and svbrdf estimation. In *Eur. Conf. Comput. Vis.*, pages 85–101. Springer, 2020. 2
- [35] Carolin Schmitt, Simon Donne, Gernot Riegler, Vladlen Koltun, and Andreas Geiger. On joint estimation of pose, geometry and svbrdf from a handheld scanner. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3493–3503, 2020. 2
- [36] Soumyadip Sengupta, Jinwei Gu, Kihwan Kim, Guilin Liu, David W Jacobs, and Jan Kautz. Neural inverse rendering of an indoor scene from a single image. In *Int. Conf. Comput. Vis.*, pages 8598–8607, 2019. 2
- [37] Pratul P Srinivasan, Boyang Deng, Xiuming Zhang, Matthew Tancik, Ben Mildenhall, and Jonathan T Barron. Nerv: Neural reflectance and visibility fields for relighting and view synthesis. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 7495–7504, 2021. 1, 2, 3
- [38] Cheng Sun, Guangyan Cai, Zhengqin Li, Kai Yan, Cheng Zhang, Carl Marshall, Jia-Bin Huang, Shuang Zhao, and Zhao Dong. Neural-pbir reconstruction of shape, material, and illumination. *Int. Conf. Comput. Vis.*, 2023. 2
- [39] Dor Verbin, Peter Hedman, Ben Mildenhall, Todd Zickler, Jonathan T Barron, and Pratul P Srinivasan. Ref-nerf: Structured view-dependent appearance for neural radiance fields. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5481–5490. IEEE, 2022. 1
- [40] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *arXiv preprint arXiv:2106.10689*, 2021. 1, 3
- [41] Xin Wei, Guojun Chen, Yue Dong, Stephen Lin, and Xin Tong. Object-based illumination estimation with rendering-aware neural networks. In *Eur. Conf. Comput. Vis.*, pages 380–396. Springer, 2020. 2
- [42] Olivia Wiles, Georgia Gkioxari, Richard Szeliski, and Justin Johnson. Synsin: End-to-end view synthesis from a single image. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 7467–7477, 2020. 1, 2
- [43] Lance Williams. Casting curved shadows on curved surfaces. In *Proceedings of the 5th annual conference on Computer graphics and interactive techniques*, pages 270–274, 1978. 4
- [44] Rui Xia, Yue Dong, Pieter Peers, and Xin Tong. Recovering shape and spatially-varying surface reflectance under unknown illumination. *ACM Trans. Graph.*, 35(6):1–12, 2016. 2
- [45] Qiangeng Xu, Zexiang Xu, Julien Philip, Sai Bi, Zhixin Shu, Kalyan Sunkavalli, and Ulrich Neumann. Point-nerf: Point-based neural radiance fields. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5438–5448, 2022. 1
- [46] Wenqi Yang, Guanying Chen, Chaofeng Chen, Zhenfang Chen, and Kwan-Yee K Wong. Ps-nerf: Neural inverse rendering for multi-view photometric stereo. In *Eur. Conf. Comput. Vis.*, pages 266–284. Springer, 2022. 1, 2, 3, 4, 5, 6
- [47] Lior Yariv, Yoni Kasten, Dror Moran, Meirav Galun, Matan Atzmon, Basri Ronen, and Yaron Lipman. Multiview neural surface reconstruction by disentangling geometry and appearance. *Adv. Neural Inform. Process. Syst.*, 33:2492–2502, 2020. 3
- [48] Wang Yifan, Felice Serena, Shihao Wu, Cengiz Öztireli, and Olga Sorkine-Hornung. Differentiable surface splatting for point-based geometry processing. *ACM Trans. Graph.*, 38(6):1–14, 2019. 2
- [49] Ye Yu and William AP Smith. Inverserendernet: Learning single image inverse rendering. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3155–3164, 2019. 2
- [50] Kai Zhang, Fajun Luan, Zhengqi Li, and Noah Snavely. Iron: Inverse rendering by optimizing neural sdf and materials from photometric images. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5565–5574, 2022. 1, 2, 3, 7
- [51] Kai Zhang, Fajun Luan, Qianqian Wang, Kavita Bala, and Noah Snavely. Physg: Inverse rendering with spherical gaussians for physics-based material editing and relighting. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5453–5462, 2021. 2, 5
- [52] Qiang Zhang, Seung-Hwan Baek, Szymon Rusinkiewicz, and Felix Heide. Differentiable point-based radiance fields for efficient view synthesis. In *SIGGRAPH Asia 2022 Conference Papers*, pages 1–12, 2022. 1, 2, 3, 4
- [53] Xiuming Zhang, Pratul P Srinivasan, Boyang Deng, Paul Debevec, William T Freeman, and Jonathan T Barron. Nerfactor: Neural factorization of shape and reflectance under an unknown illumination. *ACM Trans. Graph.*, 40(6):1–18, 2021. 1, 2, 3, 4
- [54] Yanshu Zhang, Shichong Peng, Alireza Moazeni, and Ke Li. Papr: Proximity attention point rendering. *arXiv preprint arXiv:2307.11086*, 2023. 1
- [55] Yuanqing Zhang, Jiaming Sun, Xingyi He, Huan Fu, Rongfei Jia, and Xiaowei Zhou. Modeling indirect illumination for inverse rendering. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 18643–18652, 2022. 2, 4
- [56] Dongxu Zhao, Daniel Lichy, Pierre-Nicolas Perrin, Jan-Michael Frahm, and Soumyadip Sengupta. Mvpsnet: Fast generalizable multi-view photometric stereo. In *Int. Conf. Comput. Vis.*, pages 12525–12536, 2023. 2
- [57] Zhenglong Zhou, Zhe Wu, and Ping Tan. Multi-view photometric stereo with spatially varying isotropic materials. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1482–1489, 2013. 2
- [58] Yiming Zuo and Jia Deng. View synthesis with sculpted neural points. *arXiv preprint arXiv:2205.05869*, 2022. 2
- [59] Matthias Zwicker, Mark Pauly, Oliver Knoll, and Markus Gross. Pointshop 3d: An interactive system for point-based surface editing. *ACM Trans. Graph.*, 21(3):322–329, 2002.

1, 2

- [60] Matthias Zwicker, Hanspeter Pfister, Jeroen Van Baar, and Markus Gross. Surface splatting. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 371–378, 2001.
- [61] Matthias Zwicker, Jussi Rasanen, Mario Botsch, Carsten Dachsbacher, and Mark Pauly. Perspective accurate splatting. In *Proceedings-Graphics Interface*, number CONF, pages 247–254, 2004. 1, 2