

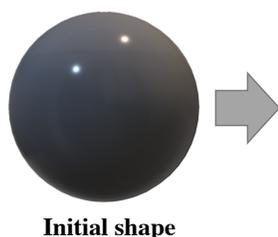
Physics-Based Inverse Rendering using Combined Implicit and Explicit Geometries

G. Cai^{1,2}, K. Yan^{1,2}, Z. Dong², I. Gkioulekas³, S. Zhao¹

¹University of California, Irvine

²Meta Reality Labs Research

³Carnegie Mellon University



Optimized shapes

Figure 1: We introduce a new technique that combines implicit and explicit geometric representations to solve physics-based inverse-rendering problems. When optimizing object geometry, our technique is capable of deforming simple spheres into detailed shapes with complex topologies. In this figure, we show re-renderings of 3D models (i.e., objects on the table and the chair behind them) reconstructed individually using our technique.

Abstract

Mathematically representing the shape of an object is a key ingredient for solving inverse rendering problems. Explicit representations like meshes are efficient to render in a differentiable fashion but have difficulties handling topology changes. Implicit representations like signed-distance functions, on the other hand, offer better support of topology changes but are much more difficult to use for physics-based differentiable rendering. We introduce a new physics-based inverse rendering pipeline that uses both implicit and explicit representations. Our technique enjoys the benefit of both representations by supporting both topology changes and differentiable rendering of complex effects such as environmental illumination, soft shadows, and interreflection. We demonstrate the effectiveness of our technique using several synthetic and real examples.

1. Introduction

Inverse rendering—the inference of object shape and appearance from 2D images—has been a long-standing problem in computer vision and graphics.

A key ingredient to solving inverse-rendering problems is the mathematical description of object surfaces. Previously, most vision and graphics applications utilize *explicit* representations like polygonal meshes that enjoys high flexibility and easy editability. Additionally, state-of-the-art physics-based differentiable rendering techniques—which requires Monte Carlo sampling over

not only object surfaces, but also discontinuity (e.g., visibility) boundaries—have largely focused on mesh-based representations. However, when solving inverse-rendering problems using gradient-based methods, explicit representations suffer from one major limitation: the difficulty in handling topology changes. This causes most existing mesh-based inverse rendering technique to require good initial geometry and/or frequent remeshing during optimization.

On the contrary, *implicit* geometric representations define a surface as the zero level set of some function (such as signed-distance

functions). Compared to their explicit counterparts, these representations enjoy a major advantage of allowing easy topology changes and, therefore, have been gaining popularity for solving inverse-rendering problems. However, implicit representations suffer from their own limitations. First, computing the intersection between a ray and an implicit surface (using, for example, sphere tracing [Har96]) can be much more expensive than computing intersections with polygons (like triangles). More importantly, since implicitly expressed surfaces are known to be difficult to parameterize, Monte Carlo sampling over these surfaces (or curves embedded in them) is also challenging. This makes it difficult to adopt state-of-the-art physics-based differentiable rendering methods [LADL18,ZWZ*19,ZMY*20] to handle implicit surfaces.

In this paper, we introduce a new physics-based inverse rendering technique that uses both implicit and explicit representations and, thus, enjoys the benefits of both worlds. Taking as input a collection of (geometrically calibrated) input images, our technique is capable of generating high-quality 3D reconstructions of object shape and appearance without the need of good initial geometries. The output of our technique is in “standard” format (i.e., a mesh with spatially varying BRDF maps) and enjoys easy editability and wide applicability.

Concretely, our contributions include:

- A physics-based differentiable rendering method for implicit surfaces (§4.1). Leveraging differentiable iso-surface extraction [RLR*20,GRL*21], we first convert an implicit surface to the corresponding explicit form (i.e., a mesh), and then apply physics-based differentiable rendering to the mesh. We demonstrate that, compared to existing differentiable-sphere-tracing based methods (e.g., [NMOG20,YKM*20]), our pipeline is not only more general—by supporting more complex light-transport effects such as environmental illumination, soft shadow, and interreflection—but also computationally more efficient.
- A two-stage inverse-rendering pipeline (§4.1 and §4.2) that enables high-fidelity 3D reconstruction of objects with complex topologies.

We demonstrate the effectiveness of our technique via several synthetic and real examples.

2. Related Works

Shape and reflectance reconstruction. Reconstructing the shape and/or reflectance of an object has been a central problem in computer vision. Many techniques like multi-view stereo (MVS) [SD99,VTC05,SZPF16], shape from shading (SfS) [Hor70,IH81,QMC*17], and photometric stereo (PS) [Woo80,HLHZ08,ZT10,QMD16] have been developed to recover object geometry and reflectance. Unfortunately, these techniques typically only recover diffuse albedo for reflectance. Further, they rely on assumptions about object appearance—such as diffuse-dominated reflectance or being sufficiently textured—and can produce low-quality reconstructions if these assumptions are violated.

Recently, Luan et al. [LZBD21] have demonstrated that, using physics-based differentiable rendering [LADL18,ZWZ*19,ZMY*20], the reconstruction of object shape and reflectance can

be formulated as an *inverse rendering* (aka. *analysis by synthesis*) problem and solved using gradient-based methods. This technique offers high generality and robustness as it relies on few assumptions about object shape or appearance. On the other hand, since their method is fully mesh-based, it has difficulties refining object topology and, therefore, requires nontrivial initialization of object geometry (by, for example, using COLMAP [SZPF16]).

In this paper, we address this problem by using both implicit and explicit representations for object geometry.

Differentiable rendering of meshes. Specialized differentiable renderers have long existed in computer graphics and vision [GZB*13, GLZ16, ALKN19, TSG19, CLZ*20]. Recently, several general-purpose ones like `renderer` [LADL18], `Mitsuba 2` [NDVZJ19], and `psdr-cuda` [ZMY*20] have been developed. A key technical challenge in differentiable rendering is to estimate gradients with respect to object geometry (e.g., positions of mesh vertices). To this end, several approximated methods [LLCL19, LHJ19, RRN*20] have been proposed. Unfortunately, inaccuracies introduced by these techniques can lead to degraded result quality, as demonstrated by Luan et al. [LZBD21]. On the contrary, recent techniques specifically handling visibility boundaries [LADL18,ZWZ*19,ZMY*20] provides unbiased gradient estimates capable of producing higher-quality reconstructions.

Differentiable rendering of implicit surfaces. Forward rendering of implicit surfaces has been studied long ago, and techniques like sphere tracing [Har96] have allowed efficient rendering of implicit geometries. Differentiable rendering implicit surfaces, on the other hand, used to be largely under-explored. Several recent works (e.g. `DVR` [NMOG20] and `IDR` [YKM*20]) have derived differentiable sphere tracing—which differentiates the ray-implicit-surface intersection computations—and introduced inverse-rendering solutions based on this technique. Unfortunately, physics-based rendering is beyond ray-surface intersection: To support effects like full-pixel anti-aliasing, area and environmental illumination, soft shadows and interreflection, contributions of rays need to be integrated (over a pixel or the hemisphere). Differentiating these integrals for implicit surfaces is challenging. In a concurrent work, Vicini et al. [VSJ22] have introduced a new technique for differentiable rendering SDFs. We consider comparisons between this approach and our technique—and potential combinations of the two—interesting future topics.

We overcome this challenge by using meshes as an intermediate representation, which can then be rendered using state-of-the-art mesh-based methods [LADL18,ZWZ*19,ZMY*20]. To our knowledge, although the high-level idea of using explicit and implicit geometries has been experimented before in different contexts (e.g., generative 3D modeling [PFAK20]), our technique is one of the first (among concurrent works like [MHS*22]) to realize this idea for physics-based inverse rendering.

Differentiable mesh generation. Recently, several techniques have been developed to generate meshes in a differentiable fashion. Remelli et al. [RLR*20], for instance, have introduced

MeshSDF—a differentiable marching cube algorithm that converts signed-distance functions (SDFs) into meshes. This technique was later extended to handle general implicit representations [GRL*21]. Additionally, Peng et al. [PJL*21] have presented a technique that creates meshes from point clouds. We utilize MeshSDF in our differentiable rendering pipeline.

3. Preliminaries on Differentiable Rendering

In what follows, we briefly revisit some key aspects in physics-based differentiable rendering. For more comprehensive introductions, please refer to recent surveys and course materials (e.g., [KBM*20, ZJL20]).

Physics-based differentiable rendering. At the core of physics-based (forward) rendering is the *rendering equation* [Kaj86]. This integral equation states that, on a non-emissive surface \mathcal{M} , the radiance L leaving some point $x \in \mathcal{M}$ with direction $\omega_o \in \mathbb{S}^2$ is given by:

$$L(x, \omega_o) = \int_{\mathbb{S}^2} L_i(x, \omega_i) f_r(x, \omega_i, \omega_o) d\sigma(\omega_i), \quad (1)$$

where L_i denotes incident radiance; f_r is the cosine-weighted BRDF (given by the original BRDF multiplied with the dot product of the incident direction ω_i and the surface normal at x); and σ is the solid-angle measure.

Recent works in physics-based differentiable rendering [LADL18, ZWZ*19, ZMY*20] have shown that, in general, differentiating the rendering equation (1) with respect to some scene parameter θ yields the *differential rendering equation*:

$$\begin{aligned} \frac{d}{d\theta} L(x, \omega_o) &= \int_{\mathbb{S}^2} \frac{d}{d\theta} [L_i(x, \omega_i) f_r(x, \omega_i, \omega_o)] d\sigma(\omega_i) \\ &+ \int_{\partial\mathbb{S}^2} v_{\perp}(\omega_i) \Delta L_i(x, \omega_i) f_r(x, \omega_i, \omega_o) d\ell(\omega_i), \end{aligned} \quad (2)$$

where the right-hand side consists of two terms. The first *interior* integral is simply provided by differentiating the integrand of the rendering equation (1). The second *boundary* integral, on the other hand, is unique to differentiable rendering: The domain $\partial\mathbb{S}^2$ is comprised of all jump discontinuity points of L_i with respect to ω_i (assuming the BRDF f_r to be continuous); $v_{\perp}(\omega_i)$ captures the change rate (with respect to θ) of the discontinuity boundary at ω_i ; and ΔL_i denotes the difference in L_i across the boundary. We note that the discontinuity of the incident radiance L_i can arise from visibility boundaries—which always exist regardless of how smooth the surface geometries are.

In practice, when the parameter θ controls the scene geometry (such as the pose of an object or the position of a mesh vertex), the boundary integral in Eq. (2) should not be neglected as it can be the major contributor to the resulting gradient $dL/d\theta$.

Geometric representations. Physics-based forward rendering methods have largely been agnostic to geometric representations: Most algorithms (such as path tracing) only requires ray-surface intersection to be efficiently computed—which can be achieved using acceleration data structures (e.g., Kd-trees) for meshes and sphere tracing [Har96] for implicit surfaces.

Physics-based differentiable rendering techniques, on the contrary, are more dependent on geometric representations due to their need for estimating the boundary integral in Eq. (2). When the scene geometries are expressed using meshes, the discontinuities curves $\partial\mathbb{S}^2$ are simply (projected) face edges. When the scene geometries are implicit, on the other hand, so are the discontinuities curves. Although there are prior works capable of generating (approximately) uniform point samples on implicit curves/surfaces (e.g., [WH94]), they are generally unsuitable for estimating integrals since the probability distributions remain unknown. To overcome this obstacle, we introduce a new pipeline in §4.1 that combines implicit and explicit representations.

4. Our Method

We now present our technique that combines implicit and explicit geometries for physics-based inverse rendering.

Our optimization pipeline involves the following two main stages. The first *implicit* stage utilizes implicit representations like signed-distance functions (SDFs) for the object geometry. The main objective of this stage is to obtain a coarse reconstruction of the object's 3D shape with the right *topology*. Since implicit geometries are used in this stage, the optimization enjoys the flexibility of allowing topology changes.

Our second *explicit* stage, on the contrary, is purely mesh-based. Specifically, this stage starts with taking the implicit geometry from the previous stage and applying (non-differentiable) iso-surface extraction to obtain the corresponding explicit expression. Then, we jointly optimize the object shape (by changing mesh vertex positions) and appearance (by changing pixel values in SVBRDF maps).

In what follows, we detail each of the two stages.

4.1. Implicit Stage

Our implicit stage computes coarse reconstructions of an object's shape and appearance. In this stage, the object surface $\mathcal{M} \subset \mathbb{R}^3$ is represented *implicitly* as the zero level set of some function ϕ that is, in turn, controlled by some abstract set of *geometric parameters* θ_g . Namely,

$$\mathcal{M}(\theta_g) := \{x \in \mathbb{R}^3 : \phi(x; \theta_g) = 0\}. \quad (3)$$

Additionally, to describe the appearance of the object, we utilize the microfacet BRDF model. For any surface point $x \in \mathcal{M}(\theta_g)$, we specify the diffuse albedo k_d , specular albedo k_s , and surface roughness α at this point using a *reflectance field* \mathcal{B} controlled by an abstract set of *optical parameters* θ_a :

$$\mathcal{B}(x; \theta_a) = (k_d \quad k_s \quad \alpha). \quad (4)$$

Differentiable rendering. To render the surface represented with \mathcal{M} and \mathcal{B} , as opposed to prior works that directly compute differentiable ray intersection against the SDF ϕ , we instead (i) apply differentiable iso-surface extraction (e.g., MeshSDF [RLR*20]) to obtain a triangle mesh $\tilde{\mathcal{M}} \approx \mathcal{M}$, and (ii) use physics-based differentiable rendering to render this mesh. As illustrated in Figure 2, this

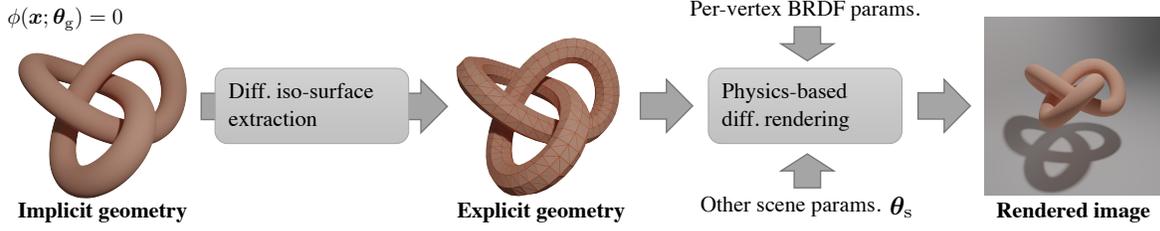


Figure 2: Our pipeline for *physics-based differentiable rendering of implicit surfaces*: Instead of directly rendering implicit geometry using techniques like differentiable sphere tracing, we leverage differentiable iso-surface extraction to covert implicit geometries into explicit ones (i.e., meshes) and render the latter using state-of-the-art differentiable rendering techniques capable of handling complex light-transport effects like environmental/area lighting, soft shadows, and interreflection.

process effectively implements a *render function* $\mathcal{R}(\theta_g, \theta_a, \theta_s)$ that takes as input the geometric parameters θ_g and optical ones θ_a —among other scene parameters θ_s such as illumination and viewing conditions—and returns an image.

Note that, since the conversion from the implicit representation ϕ to the mesh $\tilde{\mathcal{M}}$ is differentiable, the gradients of mesh vertex positions with respect to the geometric parameters θ_g are automatically available during differentiable rendering. This makes the entire rendering function \mathcal{R} end-to-end differentiable (with respect to θ_g).

Inverse-rendering optimization. Similar to prior works, we take as input n images \mathcal{I}_j of an object with known illumination and viewing conditions θ_s^j , for $j = 1, 2, \dots, n$. Then, we reconstruct the shape and appearance of the object by solving the following inverse-rendering optimization:

$$(\theta_g^*, \theta_a^*) = \arg \min_{\theta_g, \theta_a} \mathcal{L}(\theta_g, \theta_a), \quad (5)$$

where the loss \mathcal{L} is further given by

$$\mathcal{L}(\theta_g, \theta_a) = \mathcal{L}_{\text{img}}(\theta_g, \theta_a) + \lambda_{\text{reg}} \mathcal{L}_{\text{reg}}(\theta_g, \theta_a). \quad (6)$$

In Eq. (6), \mathcal{L}_{img} denotes the image loss given by

$$\mathcal{L}_{\text{img}}(\theta_g, \theta_a) = \sum_{j=1}^n \left\| \mathcal{R}(\theta_g, \theta_a, \theta_s^{(j)}) - \mathcal{I}_j \right\|_1, \quad (7)$$

and \mathcal{L}_{reg} is a regularization term with a hyper-parameter $\lambda_{\text{reg}} \in \mathbb{R}_{\geq 0}$ controlling its weight. Optionally, when each input image \mathcal{I}_j is supplemented with an anti-aliased mask image \mathcal{S}_j , our optimization can get mask supervision by including an extra loss term:

$$\lambda_{\text{mask}} \mathcal{L}_{\text{mask}}(\theta_g) = \sum_{j=1}^n \left\| \mathcal{R}_{\text{mask}}(\theta_g, \theta_s^{(j)}) - \mathcal{S}_j \right\|_1, \quad (8)$$

where $\mathcal{R}_{\text{mask}}$ is similar to the rendering function \mathcal{R} but returns only an anti-aliased mask. In practice, $\mathcal{R}(\theta_g, \theta_a, \theta_s^{(j)})$ and $\mathcal{R}_{\text{mask}}(\theta_g, \theta_s^{(j)})$ can be computed simultaneously from one rendering pass.

Discussion. By using both implicit and explicit geometries, our approach enjoys the benefit of both representations: It not only allows topological changes during the optimization, but also is capable of

capturing light transport effects—such as environmental illumination and soft shadows—that cannot be easily handled by soft rasterizers or differentiable sphere tracers. We will demonstrate the practical advantage of our technique in §5.

4.2. Explicit Stage

Based on the coarse reconstruction provided by our implicit stage (§4.1), our explicit stage further refines the result using a purely mesh-based approach. To start, we convert the implicit representation returned by the previous stage into its explicit counterpart. Specifically, we apply iso-surface extraction (e.g., marching cube) to the optimized function $\phi(\cdot; \theta_g^*)$ returned by the previous stage and obtain a mesh $\tilde{\mathcal{M}}_0$. We note that this extraction step does not need to be differentiable since it is essentially a preprocessing step for the explicit stage.

Provided the initial mesh $\tilde{\mathcal{M}}_0$, we then use the boundary first flattening (BFF) method [SC17] to UV-parameterize the mesh and, in turn, generate the spatially varying BRDF (SVBRDF) maps based on the reflectance field $\mathcal{B}(\cdot; \theta_a^*)$ given by the implicit stage.

Using this mesh and the SVBRDF maps for the initial geometry and appearance, respectively, we then apply a second inverse-rendering optimization. Our pipeline is based on the mesh-based framework introduced by Luan et al. [LZBD21] and jointly optimizes the position of each mesh vertex and the value of each pixel in the SVBRDF maps. Additionally, we utilize the recent work by Nicolet et al. [NJJ21] to update mesh vertex positions—as opposed to using the `e1Topo` geometric processing library [BB09] like Luan et al. did. This has led to significantly better performance in practice, since it does not require expensive (continuous) collision detection and handling.

4.3. Implementation Details

For our implicit stage (§4.1), we adopt the neural networks from the `idr` work by Yariv et al. [YKM*20]. Specifically, we use their occupancy network for our function ϕ and adopt their rendering network for our reflectance field \mathcal{B} . In this case, the geometric parameters θ_g and optical ones θ_a become weights of edges in the two neural networks, respectively.

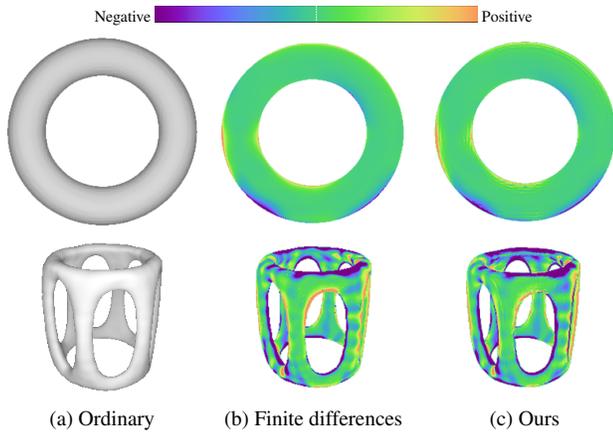


Figure 3: We validate the correctness of gradients computed by our method (§4.1) by comparing to finite-difference references. Each derivative image is computed with respect to the weight of one edge in the occupancy network (which specifies the implicit surface).

Further, we set the regularization loss as

$$\mathcal{L}_{\text{reg}}(\theta_g, \theta_a) = \sum_i (\|\nabla_x \phi(x_i; \theta_g)\| - 1)^2, \quad (9)$$

that encourages ϕ to be a signed-distance function (i.e., with unit-length gradients).

We utilize the MeshSDF technique [RLR*20] for differentiable iso-surface extraction and cache the BRDF parameter values (i.e., diffuse albedo k_d , specular albedo k_s , and surface roughness α) at each mesh vertex (as opposed to evaluate the neural reflectance field \mathcal{B} per pixel on the fly) for better rendering performance. Additionally, during optimization, we reuse one mesh produced by MeshSDF to render multiple camera views (from the same minibatch), further reducing the amortized computational overhead.

Lastly, we build our Monte Carlo differentiable renderer—which is used in both stages—based on Zhang et al.’s path-space technique [ZMY*20] and use the Adam method [KB14] implemented in pyTorch for our inverse-rendering optimizations in both stages—except for updating mesh vertex positions in the explicit stage where we use Nicolet et al.’s approach [NJJ21].

5. Results

In what follows, we validate in §5.1 our differentiable rendering of implicit surfaces. Then, we demonstrate the effectiveness of our inverse-rendering technique in §5.2. Please refer to the supplement for more results.

5.1. Validation

We validate our differentiable rendering of implicit surfaces (described in §4.1) by comparing gradient images (with respect to the weight of one edge in the occupancy network) generated with our method and finite differences (FD), as shown in Figure 3. In this figure, each derivative image is computed with respect to the weight

Table 1: Performance statistics for our inverse-rendering results in Figures 4, 5, and 6. The “MC” column shows the grid resolution used for differentiable iso-surface extraction. The “Time” column contains per-iteration optimization time (in seconds) measured on a workstation with an NVIDIA Titan RTX graphics card.

Example	# Input images	Implicit stage			Explicit stage	
		MC	# Iter.	Time	# Iter.	Time
Bunny temple	96	96 ³	1000	1.19	1000	0.73
Chair	196	128 ³	300	9.24	500	2.71
Frame	96	128 ³	3500	11.24	700	1.35
Pegasus	396	96 ³	2000	5.69	200	3.14
Chess	74	128 ³	400	2.49	200	3.34
Teapot	135	128 ³	1000	4.66	200	3.51
Leopard	132	96 ³	1500	6.32	500	3.47
Head	126	96 ³	750	5.97	350	3.38

of one edge in the occupancy neural network representing ϕ . When generating the ordinary and finite-difference (FD) results, we use sphere tracing to compute ray-surface intersection directly (without converting the implicit representations to meshes). The results demonstrate that gradients estimated by our method closely match the FD references. The stair-case-like artifacts in our results are caused by the use of marching cube with relatively low grid resolutions. Fortunately, as we will demonstrate in the rest of this section, these approximations do not affect the performance of our inverse-rendering solution.

5.2. Inverse-Rendering Results

We now demonstrate the effectiveness of our technique using a few inverse-rendering results. Please refer to Table 1 for performance numbers and the supplement for additional error visualizations and recovered SVBRDF maps.

Synthetic results. We compare our inverse-rendering pipeline with a few baselines using several synthetic examples in Figure 4. The first baseline—which we denote as IDR*—is a modified version of the IDR technique [YKM*20] where the neural rendering network is replaced with a neural reflectance field (NeRF) that outputs per-pixel BRDF parameters. Using these parameters with ray-surface intersection results provided by differentiable sphere tracing, we implemented a (differentiable) physics-based shading step that produces the final rendered image. Our second baseline is a purely mesh-based method that directly applies our explicit stage (§4.2). We set the initial geometry for IDR* and our implicit stage using network weights from IDR that lead to sphere-like shapes (see Figure 4-b). For the mesh-based baseline, we obtain the initial mesh by applying (non-differentiable) iso-surface extraction to the initial implicit geometry. As the IDR* baseline supports only one-bounce light transport, we configure all methods to render direct illumination only.

The **bunny temple** example in Figure 4 has a bunny sitting within a temple-like structure, leading to a fairly complex overall topology. In this example, we use 96 target images under collocated configurations (that is, with a point light source collocated with the camera) for all methods and show renderings from one novel

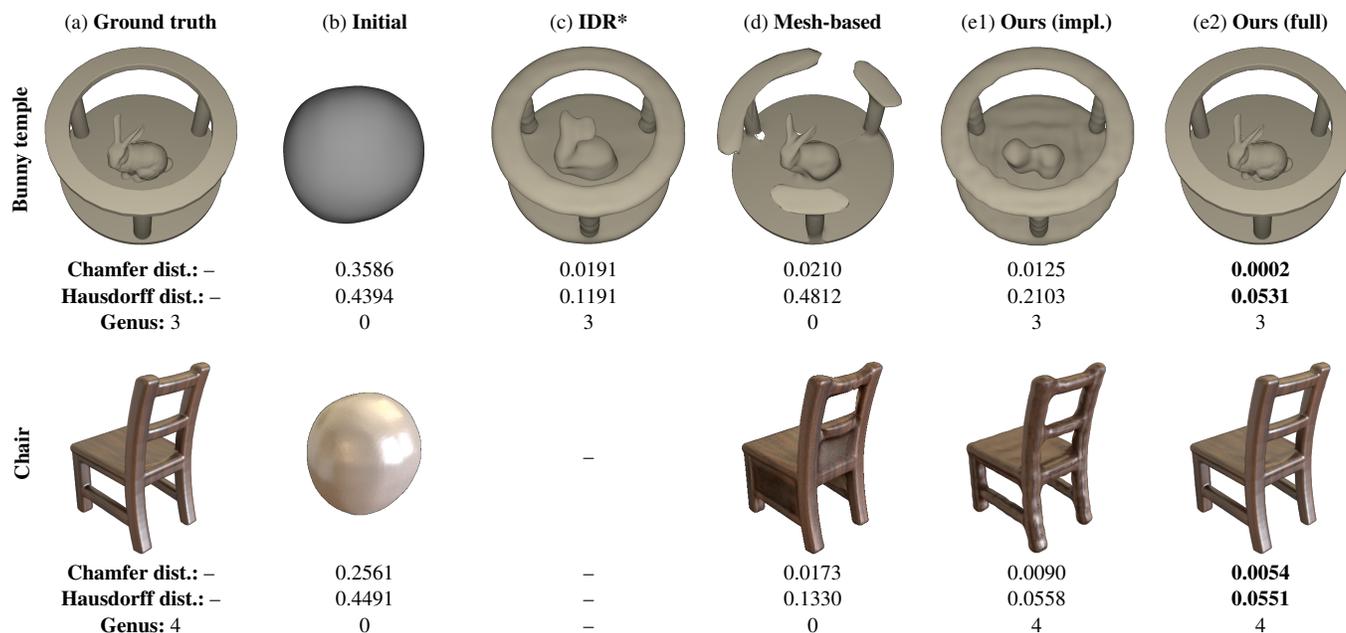


Figure 4: Inverse-rendering comparisons: We show reconstruction results generated using IDR*—a modified version of IDR that uses physics-based shading—in (c), Luan et al.’s mesh-based method [LZBD21] in (d), our implicit stage in (e1), and our full pipeline in (e2). All methods shared identical initializations shown in (b). The numbers below each reconstruction result indicate the Chamfer [BTBW77] and Hausdorff [ASCE02] distances between the reconstructed and groundtruth geometries (normalized so that the GT has a unit bounding box). Additionally, we show genus numbers of all results to validate their topologies.

viewing direction in the figure. Initialized with a sphere-like geometry (Figure 4-b), IDR* successfully recovers the overall shape of the object but fails to recover some geometric details of the bunny (Figure 4-c). The mesh-based method, on the contrary, successfully recovers some of the details but fails to obtain correct overall topology (Figure 4-d).

Our implicit stage behave similarly to IDR*: It recovers slightly fewer geometric details (Figure 4-e1) but runs approximately $3.2\times$ faster: under identical configurations, our method takes 1.19 seconds per iteration while IDR* takes 3.90 seconds. This is because (i) we use relatively coarse (i.e., 96^3) grids for iso-surface extractions, introducing low computational overhead; and (ii) our underlying differentiable renderer is much faster than IDR’s sphere-tracing-based implementation. Initialized with results from our implicit stage, our explicit stage successfully recovers the geometric details and produces a reconstruction with the best quality (Figure 4-e2).

The **chair** example in Figure 4 involves a chair with spatially varying reflectance under environmental lighting. We use 196 target images and show renderings under one novel view (top) and one novel illumination (bottom). The IDR* baseline does not apply here due to the need to estimate the boundary integrals in Eq. (2)—which is challenging since the domain of integration ∂S^2 is implicit. The mesh-based method is still applicable. But similar to the previous example, it has difficulties recovering the overall topology and introduces self intersections (Figure 4-d). Our implicit stage—thanks to its use of meshes as an intermediate representation—is

capable of perform differentiable rendering under environmental illumination and successfully returns a coarse reconstruction of the chair (Figure 4-e1). This result is further refined by our explicit stage, leading to a high-fidelity reconstruction (Figure 4-e2).

We show two additional inverse-rendering results in Figure 5. For both examples, our implicit stage successfully produces coarse reconstructions with correct overall topology. These reconstructions are then refined by our explicit stage to recover the geometric and reflectance details.

Real results. We further apply our technique to reconstruct the shape and reflectance of real objects, as shown in Figure 6. We take photographs of these objects under indoor illumination with two bright area lights to better show specular highlights. Our technique accurately recovers the fine textures for the **chess**, the detailed surface geometries for the **teapot**, the glossiness for the **chess** and the **leopard**, and the large hole for the **head** model—all based on sphere-like initial shapes.

Proof-of-concept example. To further motivate the advantage of our physics-based differentiable rendering of implicit surface, we show three proof-of-concept synthetic examples with prominent global-illumination effects in Figure 7. The top and middle examples show Cornell boxes containing two glossy kitty models and two rough glass spheres, respectively. The bottom example involves the same rough-glass object lit by a small area light. For all examples, we take as input multiple images (100 for the top, and 26

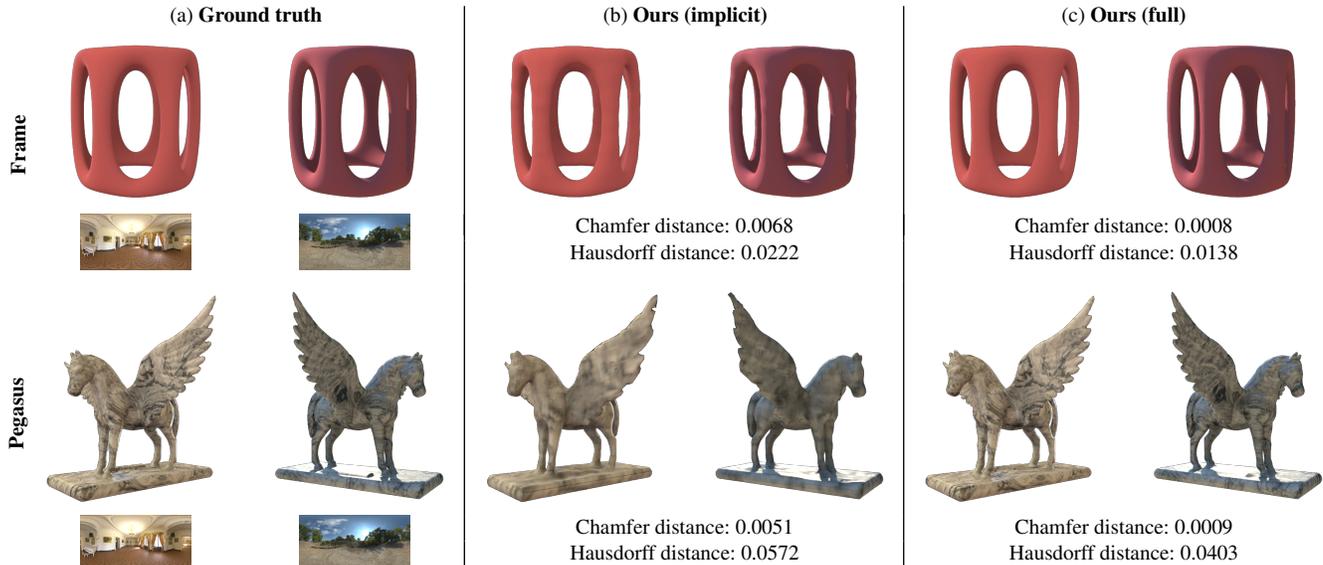


Figure 5: Inverse-rendering results (synthetic): Renderings of groundtruth (a) and reconstructed models generated by our implicit stage (b) and full pipeline (c). The left image in each column is rendered using the training illumination and a novel viewing angle; The right image is rendered using novel illumination and viewing conditions. All optimizations use the same initializations for the occupancy network (that yield sphere-like shapes similar to Figure 4-b) and neural reflectance fields (that produce near-constant BRDF parameters).

for the middle and bottom examples) of the object (with varying known orientations) and solve for the shape of the object. We note that differentiating these renderings cannot be handled by existing differentiable renderers: Soft rasterizers and simple ray tracers have difficulties handling light-transport effects like soft shadows and interreflection; existing physics-based differentiable renderers, on the other hand, cannot be easily adopted to render implicit geometries.

Thanks to our new differentiable rendering pipeline described in §4.1, for all examples, our inverse-rendering technique manages to recover the target shapes using initializations with varying topologies. We believe that the new level of generality enabled by our technique will enable interesting future applications in computer graphics, vision, and computational imaging.

6. Discussion and Conclusion

Limitations and future work. When jointly optimizing shape and spatially varying reflectance, the inverse-rendering problem can be under-constrained, causing illumination or geometric details to be “baked” into textures. Our technique, likely many (if not most) inverse-rendering techniques, can suffer from this problem—specifically when provided relatively few input images. Overcoming this challenge requires devising discriminative losses and effective regularizations, which we believe is an important topic for future research.

Additionally, the MeshSDF library [RLR*20] relies on marching cube with regular grids. Improving this step to more efficiently capture geometric details will allow our implicit stage to produce finer reconstructions, reducing the amount of refinement needed for

the explicit stage. Also, extending our pipeline to use other forms of differentiable mesh generation (e.g., [PJM*21]) is worth exploring.

Conclusion. We introduced a new physics-based inverse rendering technique that uses both implicit and explicit representations of object geometry. A key component of our technique is a new physics-based differentiable renderer for implicit surfaces: Instead of using differentiable sphere tracing, our method leverages differentiable iso-surface extraction to produce an intermediate mesh, and render this mesh using physics-based differentiable rendering. Our technique enjoys the benefits from both implicit and explicit representations by allowing easy topology changes and supporting complex light-transport effects like environmental illumination, soft shadows, and interreflection.

Acknowledgments. We thank the anonymous reviewers for their constructive feedback. This work was supported in part by NSF grants 1900783, 1900849, and 1900927. Ioannis Gkioulekas was supported by a Sloan Research Fellowship.

References

- [ALKN19] AZINOVIC D., LI T.-M., KAPLAYAN A., NIESSNER M.: Inverse path tracing for joint material and lighting estimation. In *Proc. IEEE/CVF CVPR* (2019), pp. 2447–2456. 2
- [ASCE02] ASPERT N., SANTA-CRUZ D., EBRAHIMI T.: Mesh: Measuring errors between surfaces using the hausdorff distance. In *Proceedings. IEEE international conference on multimedia and expo* (2002), vol. 1, IEEE, pp. 705–708. 6
- [BB09] BROCHU T., BRIDSON R.: Robust topological operations for dynamic explicit surfaces. *SIAM Journal on Scientific Computing* 31, 4 (2009), 2472–2493. 4



Figure 6: Inverse-rendering results (real): Photographs of real-world objects (a) and renderings of reconstructed models generated by our implicit stage (b) and full pipeline (c) under two novel views. Similar to the synthetic results, all optimizations use the same sphere-like initial geometry and near-constant BRDF parameter values.

[BTBW77] BARROW H. G., TENENBAUM J. M., BOLLES R. C., WOLF H. C.: *Parametric correspondence and chamfer matching: Two new techniques for image matching*. Tech. rep., SRI INTERNATIONAL MENLO PARK CA ARTIFICIAL INTELLIGENCE CENTER, 1977. 6

[CLZ*20] CHE C., LUAN F., ZHAO S., BALA K., GKIOULEKAS I.: Towards learning-based inverse subsurface scattering. *ICCP* (2020), 1–12. 2

[GLZ16] GKIOULEKAS I., LEVIN A., ZICKLER T.: An evaluation of computational imaging techniques for heterogeneous inverse scattering. In *ECCV* (2016), Springer, pp. 685–701. 2

[GRL*21] GUILLARD B., REMELLI E., LUKOIANOV A., RICHTER S., BAGAUDINOV T., BAQUE P., FUA P.: Deepmesh: Differentiable iso-surface extraction. *arXiv preprint arXiv:2106.11795* (2021). 2

[GZB*13] GKIOULEKAS I., ZHAO S., BALA K., ZICKLER T., LEVIN A.: Inverse volume rendering with material dictionaries. *ACM Trans. Graph.* 32, 6 (2013), 1–13. 2

[Har96] HART J. C.: Sphere tracing: A geometric method for the antialiased ray tracing of implicit surfaces. *The Visual Computer* 12, 10 (1996), 527–545. 2, 3

[HLHZ08] HOLROYD M., LAWRENCE J., HUMPHREYS G., ZICKLER T.: A photometric approach for estimating normals and tangents. *ACM Trans. Graph.* 27, 5 (2008), 133:1–133:9. 2

[Hor70] HORN B. K.: Shape from shading: A method for obtaining the shape of a smooth opaque object from one view. 2

[IH81] IKEUCHI K., HORN B. K.: Numerical shape from shading and occluding boundaries. *Artificial intelligence* 17, 1-3 (1981), 141–184. 2

[Kaj86] KAJIYA J. T.: The rendering equation. In *Proceedings of the 13th annual conference on Computer graphics and interactive techniques* (1986), pp. 143–150. 3

[KB14] KINGMA D. P., BA J.: Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014). 5

[KBM*20] KATO H., BEKER D., MORARIU M., ANDO T., MATSUOKA T., KEHL W., GAIDON A.: Differentiable rendering: A survey. *arXiv preprint arXiv:2006.12057* (2020). 3

[LADL18] LI T.-M., AITTALA M., DURAND F., LEHTINEN J.: Differentiable Monte Carlo ray tracing through edge sampling. *ACM Trans. Graph.* 37, 6 (2018), 1–11. 2, 3

[LHJ19] LOUBET G., HOLZSCHUCH N., JAKOB W.: Reparameterizing discontinuous integrands for differentiable rendering. *ACM Trans. Graph.* 38, 6 (2019), 1–14. 2

[LLCL19] LIU S., LI T., CHEN W., LI H.: Soft rasterizer: A differentiable renderer for image-based 3D reasoning. In *ICCV* (2019), pp. 7708–7717. 2

[LZBD21] LUAN F., ZHAO S., BALA K., DONG Z.: Unified shape and

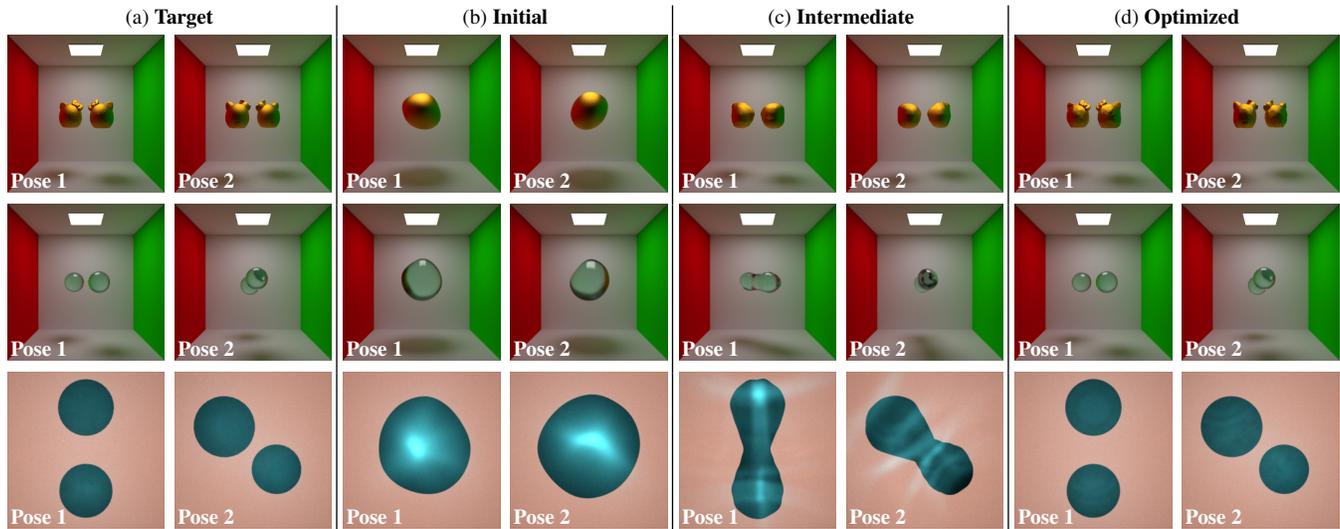


Figure 7: Proof-of-concept examples: The top and middle examples consist of Cornell boxes containing two glossy kitty models and two rough glass spheres, respectively. In the bottom example, we use cast shadow (and caustics) images of two rough glass spheres lit by a small area light. For all examples, we infer the shape of the object using multiple input images (with two shown) and corresponding object orientations. Using one sphere as the initial shape, our inverse-rendering pipeline successfully recovers the target shapes with different topologies as the initial.

- svbrdf recovery using differentiable monte carlo rendering. *Computer Graphics Forum* 40, 4 (2021), 101–113. 2, 4, 6
- [MHS*22] MUNKBERG J., HASSELGREN J., SHEN T., GAO J., CHEN W., EVANS A., MÜLLER T., FIDLER S.: Extracting triangular 3d models, materials, and lighting from images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022), pp. 8280–8290. 2
- [NDVZ19] NIMIER-DAVID M., VICINI D., ZELTNER T., JAKOB W.: Mitsuba 2: A retargetable forward and inverse renderer. *ACM Trans. Graph.* 38, 6 (2019), 1–17. 2
- [NJJ21] NICOLET B., JACOBSON A., JAKOB W.: Large steps in inverse rendering of geometry. *ACM Transactions on Graphics* 40, 6 (2021), 248:1–248:13. 4, 5
- [NMOG20] NIEMEYER M., MESCHEDER L., OECHSLE M., GEIGER A.: Differentiable volumetric rendering: Learning implicit 3D representations without 3D supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2020), pp. 3504–3515. 2
- [PFAK20] POURSAEED O., FISHER M., AIGERMAN N., KIM V. G.: Coupling explicit and implicit surface representations for generative 3D modeling. In *European Conference on Computer Vision* (2020), Springer, pp. 667–683. 2
- [PJJ*21] PENG S., JIANG C. M., LIAO Y., NIEMEYER M., POLLEFEYS M., GEIGER A.: Shape as points: A differentiable poisson solver. In *Advances in Neural Information Processing Systems (NeurIPS)* (2021). 2, 7
- [QMC*17] QUÉAU Y., MÉLOU J., CASTAN F., CREMERS D., DUROU J.-D.: A variational approach to shape-from-shading under natural illumination. In *International Workshop on Energy Minimization Methods in Computer Vision and Pattern Recognition* (2017), pp. 342–357. 2
- [QMD16] QUÉAU Y., MECCA R., DUROU J.-D.: Unbiased photometric stereo for colored surfaces: A variational approach. In *Proc. IEEE CVPR* (2016), pp. 4359–4368. 2
- [RLR*20] REMELLI E., LUKOIANOV A., RICHTER S., GUILLARD B., BAGAUTDINOV T., BAQUE P., FUA P.: MeshSDF: Differentiable isosurface extraction. In *Advances in Neural Information Processing Systems* (2020), Larochelle H., Ranzato M., Hadsell R., Balcan M. F., Lin H., (Eds.), vol. 33, pp. 22468–22478. 2, 3, 5, 7
- [RRN*20] RAVI N., REIZENSTEIN J., NOVOTNY D., GORDON T., LO W.-Y., JOHNSON J., GKIOXARI G.: Accelerating 3D deep learning with PyTorch3D. *arXiv preprint arXiv:2007.08501* (2020). 2
- [SC17] SAWHNEY R., CRANE K.: Boundary first flattening. *ACM Trans. Graph.* 37, 1 (2017), 5:1–5:14. 4
- [SD99] SEITZ S. M., DYER C. R.: Photorealistic scene reconstruction by voxel coloring. *International Journal of Computer Vision* 35, 2 (1999), 151–173. 2
- [SZPF16] SCHÖNBERGER J. L., ZHENG E., POLLEFEYS M., FRAHM J.-M.: Pixelwise view selection for unstructured multi-view stereo. In *ECCV* (2016). 2
- [TSG19] TSAI C.-Y., SANKARANARAYANAN A. C., GKIOULEKAS I.: Beyond volumetric albedo—a surface optimization framework for non-line-of-sight imaging. In *Proc. IEEE/CVF CVPR* (2019), pp. 1545–1555. 2
- [VSJ22] VICINI D., SPEIERER S., JAKOB W.: Differentiable signed distance function rendering. *ACM Trans. Graph.* 41, 4 (2022), 125:1–125:18. 2
- [VTC05] VOGIATZIS G., TORR P. H., CIPOLLA R.: Multi-view stereo via volumetric graph-cuts. In *Proc. IEEE CVPR* (2005), vol. 2, pp. 391–398. 2
- [WH94] WITKIN A. P., HECKBERT P. S.: Using particles to sample and control implicit surfaces. In *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques* (1994), SIGGRAPH '94, p. 269–277. 3
- [Woo80] WOODHAM R. J.: Photometric method for determining surface orientation from multiple images. *Optical engineering* 19, 1 (1980), 191139. 2
- [YKM*20] YARIV L., KASTEN Y., MORAN D., GALUN M., ATZMON M., RONEN B., LIPMAN Y.: Multiview neural surface reconstruction

- by disentangling geometry and appearance. *Advances in Neural Information Processing Systems* 33 (2020). [2](#), [4](#), [5](#)
- [ZJL20] ZHAO S., JAKOB W., LI T.-M.: Physics-based differentiable rendering: from theory to implementation. In *ACM SIGGRAPH 2020 Courses*. 2020, pp. 1–30. [3](#)
- [ZMY*20] ZHANG C., MILLER B., YAN K., GKIOULEKAS I., ZHAO S.: Path-space differentiable rendering. *ACM Trans. Graph.* 39, 4 (2020), 143:1–143:19. [2](#), [3](#), [5](#)
- [ZT10] ZHOU Z., TAN P.: Ring-light photometric stereo. In *European Conference on Computer Vision* (2010), pp. 265–279. [2](#)
- [ZWZ*19] ZHANG C., WU L., ZHENG C., GKIOULEKAS I., RAMAMOORTHY R., ZHAO S.: A differential theory of radiative transfer. *ACM Trans. Graph.* 38, 6 (2019), 1–16. [2](#), [3](#)