

Towards Reflectometry from Interreflections

Kfir Shem-Tov^{*1} Sai Praveen Bangaru^{*2} Anat Levin¹ Ioannis Gkioulekas²

Abstract—Reflectometry is the task for acquiring the bidirectional reflectance distribution function (BRDFs) of real-world materials. The typical reflectometry pipeline in computer vision, computer graphics, and computational imaging involves capturing images of a convex shape under multiple illumination and imaging conditions; due to the convexity of the shape, which implies that all paths from the light source to the camera perform a single reflection, the intensities in these images can subsequently be analytically mapped to BRDF values. We deviate from this pipeline by investigating the utility of higher-order light transport effects, such as the interreflections arising when illuminating and imaging a concave object, for reflectometry. We show that interreflections provide a rich set of constraints on the unknown BRDF, significantly exceeding those available in equivalent measurements of convex shapes. We develop a differentiable rendering pipeline to solve an inverse rendering problem that uses these constraints to produce high-fidelity BRDF estimates from even a single input image. Finally, we take first steps towards designing new concave shapes that maximize the amount of information about the unknown BRDF available in image measurements. We perform extensive simulations to validate the utility of this reflectometry from interreflections approach.

Index Terms—bidirectional reflectance distribution function, reflectometry, interreflections, differentiable rendering

1 INTRODUCTION

Bidirectional reflectance distribution functions (BRDF) [1] have served as one of the main ways for describing the reflectance properties of (opaque) materials in computer vision, computer graphics, and computational imaging. They have found use in applications ranging from physically accurate rendering, to simultaneous shape and material acquisitions, and from virtual and augmented reality, to semantic scene parsing [2], [3], [4], [5], [6], [7].

Given the importance of BRDFs, there has been extensive research on techniques that can measure the BRDF of real materials [2], [8], [9], [10], [11], [12], [13], [14]. The most common approach for this acquisition task is to fabricate a convex object (e.g., a plane or sphere) of known shape whose surface is characterized by the unknown BRDF, and then use a goniometer to capture images of that object from multiple (and potentially multiplexed) illumination and viewing directions. The intensities in these images can then directly be mapped to BRDF values.

A key challenge with this approach is the sheer number of images that must be acquired to obtain high-fidelity estimates of the BRDF: Each intensity measurement provides only a single sample of a, in full-generality, four-dimensional function, requiring several dozens of images to produce a dense sampling of this function. This has motivated significant research towards reducing the required number of measurements. This has led to the discovery of inherent symmetries and general-purpose parametric forms that can be used to model real-world BRDFs [15], [16], [17], [18], and to the design of reduced sampling schemes that take advantage of this prior knowledge [11], [19], [20], [21].

In this paper, we approach the problem of reducing

the acquisition effort of reflectometry from an orthogonal direction: Our main insight is that the intensity of light paths corresponding to interreflections (e.g., those produced when illuminating and imaging a shape with concavities) is a function of more than one samples of the BRDF—one sample for each reflection event that takes place along the path. Therefore, by capturing images of concave objects, we effectively *multiplex* multiple measurements of the BRDF into a smaller set of images than what would be required if we were using a convex object. This multiplexing does not come for free: Given the infinitely many possible interreflection paths contributing to each image pixel, and the fact that the intensity of each path is a *non-linear* value of the BRDF, extracting BRDF values from images containing interreflections is a challenging *inverse rendering* problem. We show that we can use recently-developed *differentiable rendering* technologies [22], [23], [24], [25] to efficiently solve this problem, and obtain high-fidelity BRDF estimates from a much smaller number of images than what would be required using interreflection-free convex shapes. To provide initial evidence for the feasibility and utility of this *reflectometry from interreflections* approach, we perform extensive simulations using images synthesized with physically-accurate rendering, for various convex and concave shapes. Last but not least, we use the insights drawn from these simulations to take first steps towards designing new concave shapes that produce an optimal set of measurements for BRDF estimation.

2 RELATED WORK

Reflectance acquisition. Measuring BRDFs has a long history in computer vision and graphics [7]. Laboratory measurements often involve using a goniometer to image and illuminate a known, simple shape from multiple directions using a goniometer or a light stage [2], [8], [9], [10], or

• ¹ Department of Electrical Engineering, Technion.

• ² Robotics Institute, Carnegie Mellon University.

• kfirshemtov@gmail.com, sbangaru@mit.edu, anat.levin@ee.technion.ac.il, igkioule@cs.cmu.edu

• ^{*} Equal contribution authors

different forms of coded illumination [12], [13], [14]. Typically, these techniques require a large number of measurements, which has motivated significant research towards sufficient [26], [27], optimal [19], [20], [21], perceptually-uniform [28], [29], or adaptive [11] sampling schemes. We take a different approach towards reducing the required number of image measurements, by showing that images of concave objects can provide sufficient samples for detailed BRDF recovery thanks to interreflections.

BRDF estimation from a single or few images has also been explored in contexts where laboratory measurements are not practical, for example under passive, potentially unknown, illumination [18], [27], [30], [31], [32]. More recently, data-driven approaches have become popular as a means for recovering low-dimensional parametric forms of even spatially-varying BRDFs from such limited measurements [33], [34], [35]. We take inspiration from these works to show that high-fidelity reflectance acquisition can also be achieved from a single or few images in the active illumination setting.

Interreflections as a source of information. Interreflections have previously been shown to be a rich source of information for shape estimation [36], [37], [38]. Most of these works assume that the shape has a Lambertian BRDF, which simplifies the mathematical expressions for the underlying light transport quantities.

In the context of reflectance acquisition, interreflections have been used in two different ways. Naik et al. [39], [40] used time-of-flight measurements of interreflections between a sample of unknown reflectance and Lambertian surfaces to perform single-view BRDF recovery. Tsai et al. [41] showed that time-of-flight measurements of two-bounce interreflections between surfaces of unknown shape and reflectance provide dense samples of the unknown BRDF; they additionally developed an algorithm for simultaneous shape and BRDF recovery. We draw inspiration from these works to develop an algorithm for reflectometry from interreflections that: (i) works using measurements from a steady-state intensity camera, without the need for time-of-flight information; (ii) requires very few image measurements, without the need to acquire high-dimensional light transport matrices; and (iii) uses information from light transport of arbitrary number of bounces, without the need to isolate only two-bounce transport.

Differentiable rendering. The ability to estimate derivatives of radiometric quantities (e.g., the images a camera would capture) with respect to arbitrary scene parameters has recently emerged as a key computational tool for solving inverse problems in computer vision and graphics. Differentiable renderers based on mathematically-tractable direct shading models [42], [43] have become common in deep learning pipelines that attempt to infer physical scene parameters, including reflectance [44], [45]. More general differentiable renderers that can handle higher-order light transport effects were first introduced in the context of inverse scattering [46], [47], [48], [49], [50], where the goal was to recover material scattering parameters rather than BRDFs. More recently, these have been extended to support differentiation with respect to arbitrary scene parameters [22], [23], [24], [25] and have been used to infer low-parametric reflectance model under known [51] or unknown

geometry [52], [53]. Here we take inspiration from this work, to show the utility of differentiable rendering as a computational tool for high-quality reflectometry.

3 PROBLEM SETTING

We denote a BRDF as a 4D function $f_r(\omega_i, \omega_o)$, where ω_i, ω_o are unit vectors describing the incoming and outgoing directions (each having two degrees of freedom). For any two such vectors, $f_r(\omega_i, \omega_o)$ denotes the ratio of radiance reflected along ω_o , to incident irradiance from direction ω_i . The directions ω_i, ω_o are defined with respect to the local normal \mathbf{n} of the imaged surface point. We assume that we have available a set of M input images I_1, \dots, I_M of the same object, captured with a fixed orthographic camera, under different *directional* illuminations from directions $\mathbf{l}_1, \dots, \mathbf{l}_M$. We assume that the geometry, including the 3D shape of the object, light directions, and camera pose are known, and that the target object is characterized by a spatially-uniform BRDF. The selection of the shape of the object will be a key theme of our paper as discussed in later sections.

To recover the unknown BRDF f_r , we use an *analysis-by-synthesis* (also known as *inverse rendering* [54], [55]) approach: We search for the BRDF that, when used to render synthetic images, most closely matches the input images. Concretely, we can express this as the following minimization problem:

$$\min_{f_r} \sum_{m=1}^M \|I_m - \mathcal{R}(\mathbf{l}_m, f_r)\|^2, \quad (1)$$

where $\mathcal{R}(\cdot)$ is the rendering operator, implicitly including the scene geometry.

3.1 BRDF parameterization

Before showing how to efficiently solve the optimization problem of Eq. (1), we first discuss how we parameterize the space of possible BRDFs f_r .

Isotropic BRDFs. We begin by adopting a commonly-used change of variables to express the BRDF as a function of not the incoming and outgoing directions ω_i, ω_o , but the half-vector \mathbf{h} and difference vector \mathbf{d} [15]. The half vector is the unit vector half way between the incoming and outgoing directions, and the difference vector is simply the incident direction in a frame of reference in which the half vector is at the north pole. Concretely:

$$\mathbf{h} = \frac{\omega_i + \omega_o}{\|\omega_i + \omega_o\|}, \quad \mathbf{d} = \text{Rot}_y(-\theta_h) \text{Rot}_z(-\phi_h) \omega_i, \quad (2)$$

where $\theta_h, \phi_h, \theta_d, \phi_d$ are the spherical coordinates of the vectors \mathbf{h}, \mathbf{d} .

This so-called half-vector parameterization has been shown to facilitate access to several properties exhibited by real-world BRDFs. In particular, it has been observed that many real world BRDFs are *isotropic* [15], meaning that they are invariant to rotations of the half vector around the normal. In the half-vector parametrization, this implies that the BRDF is invariant to angle ϕ_h , and therefore becomes a 3D function. We adopt this widely-employed assumption

as well, and restrict our attention to the space of isotropic BRDFs.

Dictionary representation. There is a large literature of models for developing parametric forms for isotropic BRDFs [11], [17], [21], [56], [57], [58], [59], [60], [61], [62], [63], [64], [65]. Here, we follow Hui et al. [32] and adopt a *dictionary* representation, for reasons of both parsimony (many different BRDFs can be accurately expressed using this model) and mathematical tractability (such a representation helps make the optimization problem of Eq. (1) easier to solve). In particular, we assume that the unknown BRDF can be expressed as a convex combination of a dictionary $\{f_r^n, n = 1, \dots, N\}$ of BRDFs,

$$f_r(\alpha) = \sum_{n=1}^N \alpha_n f_r^n, \quad s.t. \alpha_n \geq 0, \|\alpha\|_1 = 1, \quad (3)$$

and hence Eq. (1) reduces to finding the mixture weights

$$\min_{\alpha: \alpha_n \geq 0, \|\alpha\|_1 = 1} \sum_{m=1}^M \left\| I_m - \mathcal{R} \left(\mathbf{l}_m, \sum_{n=1}^N \alpha_n f_r^n \right) \right\|^2, \quad (4)$$

For our dictionary, we combine different microfacet BRDFs using the isotropic GGX parametric model [59]. This model uses two parameters: a roughness parameter α , where α approaching 0 leads to very specular BRDFs and $\alpha \approx 0.5$ has a very wide specular lobe; and the index of refraction η controlling the Fresnel reflection coefficient. We selected a dictionary of 40 GGX functions, for different values of α and η , shown in Fig. 1, that best explains the BRDFs in the MERL BRDF database [10]. We additionally augment this dictionary with a unit-albedo Lambertian BRDF, as well as a BRDF of zero albedo (completely-absorbing surface), for a total of $N = 42$ dictionary elements. Given that all atoms in our dictionary are physically-accurate BRDFs, requiring the weights in α to be non-negative and sum to one ensures that all BRDFs represented by our dictionary model will satisfy the positivity, energy-conservation, and Helmholtz reciprocity properties required of physical BRDFs. The use of the fully-absorptive dictionary element is needed in order to allow our BRDFs to have non-unit albedo. Finally, as all elements of the dictionary are isotropic BRDFs, the BRDFs from our dictionary model will also be isotropic.

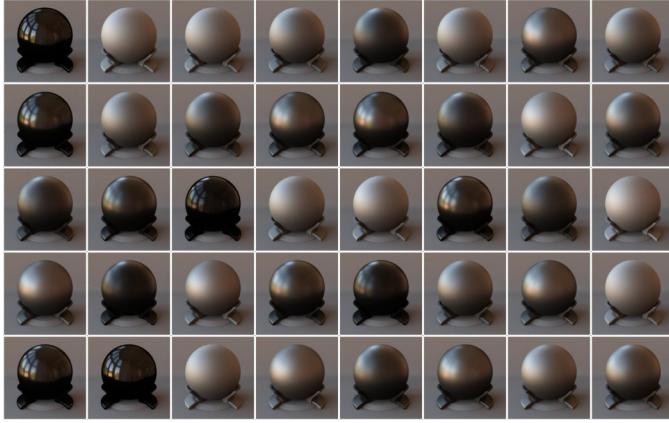


Fig. 1. The 40 microfacet BRDFs in our dictionary.

Bivariate BRDFs. Romeiro and Zickler [16] noted that many real-world BRDFs are approximately independent of ϕ_d ,

and can therefore be reduced to 2D functions of only θ_h, θ_d . We do not assume this so-called *bivariate representation* for our recovery experiments, but we will use it in later sections to visualize BRDFs and show sampling patterns.

4 OPTIMIZATION WITH DIFFERENTIABLE RENDERING

Given the large number of unknowns, efficiently minimizing the analysis-by-synthesis loss of Eq. (4) requires using a gradient-based optimization algorithm. Computing gradients requires differentiating the rendering operator $\mathcal{R}(\cdot)$ with respect to BRDF parameters. For convex shapes, this operator describes light paths involving a single reflection event; in such cases the rendering operator \mathcal{R} is a simple linear function of the unknown BRDF dictionary parameters, making differentiation trivial. However, for concave shapes, the rendering operator includes contributions from paths with multiple reflection events, due to interreflections. This makes computing both the rendering operator itself and its derivatives considerably more complicated, requiring techniques such as physically-accurate Monte Carlo forward and differentiable rendering, respectively. We provide a brief overview of these techniques, deferring to [22], [23], [24], [25], [66] for details.

Differentiable rendering. Our starting point is to use the so-called *path integral formulation* of light transport [67], in order to express the rendering operator as the sum of throughput contributions from all valid light paths in the scene. Concretely:

$$\mathcal{R}(\mathbf{l}_m, f_r) = \int_{\vec{\mathbf{x}} \in \mathbb{P}} f_p(\vec{\mathbf{x}}; f_r) d\vec{\mathbf{x}}, \quad (5)$$

where $\vec{\mathbf{x}} = \mathbf{x}_0 \rightarrow \mathbf{x}_1 \rightarrow \dots \mathbf{x}_K \rightarrow \mathbf{x}_{K+1}$ is a piecewise-linear path starting at the illumination $\mathbf{x}_0 = \mathbf{l}_m$, ending on the sensor $\mathbf{x}_{K+1} = \mathbf{i}^1$, and passing through K intermediate points on the unknown object, $K \geq 0$, and \mathbb{P} is the space of all such paths. The path throughput f_p describes how radiance is reflected at the intermediate points, and can be written in product form as

$$f_p(\vec{\mathbf{x}}; f_r) = G(\mathbf{x}_0, \mathbf{x}_1) \prod_{k=1}^{K-1} G(\mathbf{x}_k, \mathbf{x}_{k+1}) f_r(\widehat{\mathbf{x}_k \mathbf{x}_{k+1}}), \quad (6)$$

where $\widehat{\mathbf{x}_k \mathbf{x}_{k+1}}$ denotes the unit vector in the direction $\mathbf{x}_{k+1} - \mathbf{x}_k$. The geometric term $G(\mathbf{x}_k, \mathbf{x}_{k+1})$ accounts for visibility, foreshortening, and light fall-off effects,

$$\frac{\langle \mathbf{n}_k \mathbf{x}_k, \widehat{\mathbf{x}_{k+1}} \rangle \langle \mathbf{x}_k \widehat{\mathbf{x}_{k+1}}, \mathbf{n}_k \rangle}{\|\mathbf{x}_{k+1} - \mathbf{x}_k\|^2} V(\mathbf{x}_k, \mathbf{x}_{k+1}), \quad (7)$$

where \mathbf{n}_k is the normal of the surface at surface point \mathbf{x}_k , and $V(\mathbf{x}_k, \mathbf{x}_{k+1})$ denotes binary visibility between points \mathbf{x}_k and \mathbf{x}_{k+1} .

By differentiating (5) with respect to a BRDF dictionary weight α_n , and after a short calculation, we can write:

$$\frac{\partial \mathcal{R}(\mathbf{l}_m, f_r)}{\partial \alpha_n} = \int_{\vec{\mathbf{x}} \in \mathbb{P}} f_p(\vec{\mathbf{x}}; f_r) S_p(\vec{\mathbf{x}}; f_r, \alpha_n) d\vec{\mathbf{x}}, \quad (8)$$

where S_p is the path score function defined as,

$$S_p(\vec{x}; f_r, \alpha_n) = \sum_{k=1}^{K-1} \frac{\partial f_r(\widehat{\mathbf{x}_k \mathbf{x}_{k-1}}, \widehat{\mathbf{x}_k \mathbf{x}_{k+1}}) / \partial \alpha_n}{f_r(\widehat{\mathbf{x}_k \mathbf{x}_{k-1}}, \widehat{\mathbf{x}_k \mathbf{x}_{k+1}})} \quad (9)$$

$$= \sum_{k=1}^{K-1} \frac{f_r^n(\widehat{\mathbf{x}_k \mathbf{x}_{k-1}}, \widehat{\mathbf{x}_k \mathbf{x}_{k+1}})}{f_r(\widehat{\mathbf{x}_k \mathbf{x}_{k-1}}, \widehat{\mathbf{x}_k \mathbf{x}_{k+1}})}, \quad (10)$$

and in Eq. (10) we used the dictionary parameterization of Eq. (3).

Given the expressions of Eq. (5) and Eq. (8), Monte Carlo forward and differentiable rendering approximates the rendering operator and its derivative, respectively, by first sampling P paths from a distribution $q(\vec{x})$, and then forming the estimators:

$$\mathcal{R}(\mathbf{l}_m, f_r) \approx \frac{1}{P} \sum_{p=1}^P \frac{f_p(\vec{x}_p)}{q(\vec{x}_p)}, \quad (11)$$

$$\frac{\partial \mathcal{R}(\mathbf{l}_m, f_r)}{\partial \alpha_n} \approx \frac{1}{P} \sum_{p=1}^P \frac{f_p(\vec{x}_p) S_p(\vec{x}_p; f_r, \alpha_n)}{q(\vec{x}_p)}. \quad (12)$$

To reduce the variance of these estimators, $q(\vec{x})$ should be a good approximation of the integrands. In our implementation, we use bidirectional path tracing [68] to sample paths, that we use to form both estimates. We use Mitsuba [69] for rendering, which we have modified to implement the above differentiable rendering procedure. We will make our implementation publicly available upon acceptance of the paper.

Optimization procedure. Given access to *stochastic* estimates of the derivative of the rendering operator, and therefore the loss function of Eq. (4), we can optimize for the unknown mixing weights α using gradient-based optimization algorithms. We use the *exponentiated gradient* algorithm [70], [71], modified to use Adam-style updates for scheduling the learning rate and incorporating momentum [72]. The use of the exponentiated gradient descent algorithm guarantees that the resulting mixing weights will satisfy the constraints of the optimization problem of Eq. (4) (non-negativity and a sum equal to one). To initialize the gradient descent procedure, we use an initial estimate of the mixing weights α , produced by assuming that the input images only contain single-bounce paths. Under this assumption, the optimization problem of Eq. (3) reduces to a linear least-squares problem that can be solved analytically (see Sec. 5 and Eq. (13)).

5 THE BENEFIT OF INTERREFLECTIONS

We now have available to us all the tools we need to demonstrate the benefits of using interreflections for BRDF estimation. To this end, we will be comparing the results of the optimization problem of Eq. (3), for measurements captured from convex and concave versions of the same shape. In the convex case, where all light paths have just one reflection event, this problem can be solved analytically—one can first render images of the object under each of the BRDFs in the dictionary, then recover the mixing weights α by solving a linear least-squares problem:

$$\min_{\alpha: \alpha_n \geq 0, \|\alpha\|=1} \sum_{m=1}^M \left\| I_m - \sum_{n=1}^N \alpha_n \mathcal{R}(\mathbf{l}_m, f_r^n) \right\|^2, \quad (13)$$

By contrast, in the concave case, because of the presence of multi-bounce paths, solving this optimization requires performing the computationally expensive gradient descent procedure we discussed in the previous section.

The computational tractability of the convex case comes at the cost of reduced information about the unknown BRDF: For fixed viewing and illumination directions, the number of unique constraints on the BRDF provided by one image is equal to number of unique visible normals on the shape for a general BRDF, and a subset of those for an isotropic BRDF. By contrast, in the concave case, each reflection event in each of the multi-bounce paths in integral Eq. (5) provides information about a different point sample of the underlying BRDF, resulting in a significantly larger number of constraints that can be used for inference.

Simple demonstration. To demonstrate the above described trade-off between computational complexity and amount of information, we first show experiments on a simple tetrahedron shape, comprising three planar facets (see Fig. 2). Under a single directional illumination and orthographic camera, the convex shape provides us with only three constraints on the underlying BRDF, one for each facet. If the number of unknown parameters we need to determine in the optimization problem Eq. (3) is greater than three (i.e., the dictionary has $N > 3$ elements), then the optimization problem is ill-posed. This is shown in Fig. 3, where we see that a single measurement of the convex shape can be used to successfully recover a BRDF described by a three-element dictionary, but does not provide sufficient information to recover higher-dimensional BRDFs. By contrast, in the concave case, the multi-bounce paths contributing to the single image measurement provide a much larger number of BRDF constraints. As shown in the figure, these are sufficient to recover even a 40-dimensional BRDF, with estimation error remaining essentially constant for varying dictionary sizes.

Quantitative evaluation. To more systematically evaluate the utility of interreflections for BRDF recovery, we consider convex and concave versions of a sphere. We select this shape because its convex version is commonly used for reflectometry, and already provides a much richer set of constraints on the unknown BRDF. Throughout the rest of this section, we also show results for an additional shape of our design. We defer discussion of this shape to Sec. 6.

We perform reconstruction experiments using BRDFs from the MERL database [10], assuming only *one* input measurement image. In Figs. 6 and 7, we compare renderings produced for a novel shape and environment map using the BRDFs recovered from our optimization procedure. In all cases, the BRDF recovered from the concave shape matches the appearance of the groundtruth more closely than the one recovered from the convex shape. We note that, even under perfect measurements, the fit will be imperfect, because most BRDFs in the MERL dataset are outside the span of the dictionary we are using. To demonstrate this, we include a *reference* image rendered with the BRDF produced by directly projecting the target BRDF on the dictionary space. To visualize the groundtruth and reconstructed BRDFs, in Fig. 8 we compare their projections to the 2D bivariate space discussed in Sec. 3.1.

We also consider how different shapes perform as the number of available measurements increases. In Fig. 4,

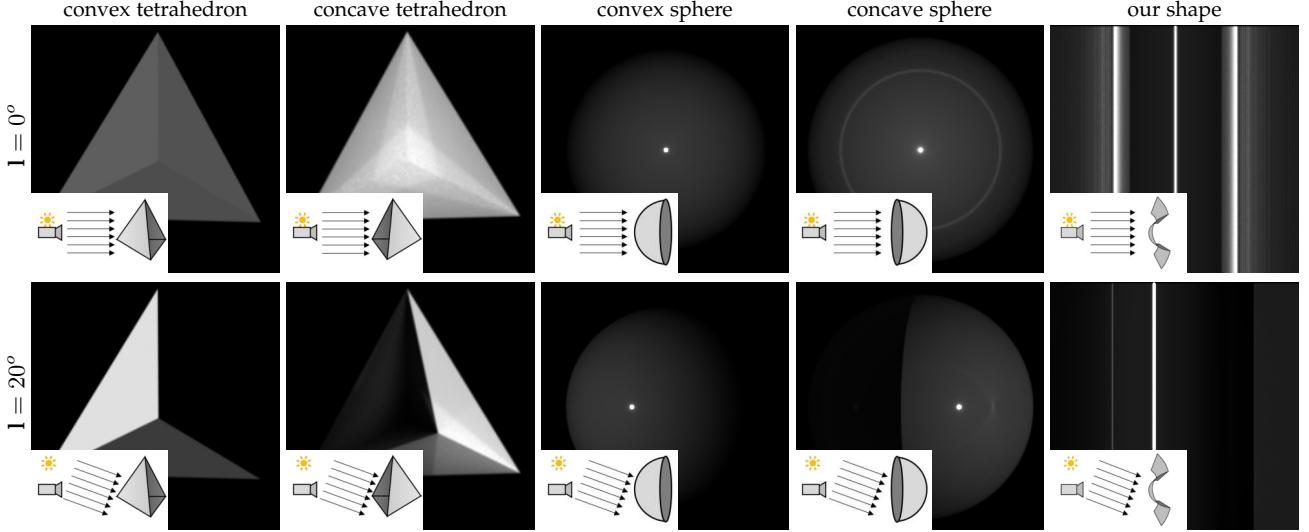


Fig. 2. Input images produced from the five shapes we evaluate, rendered under two different directional illuminations.

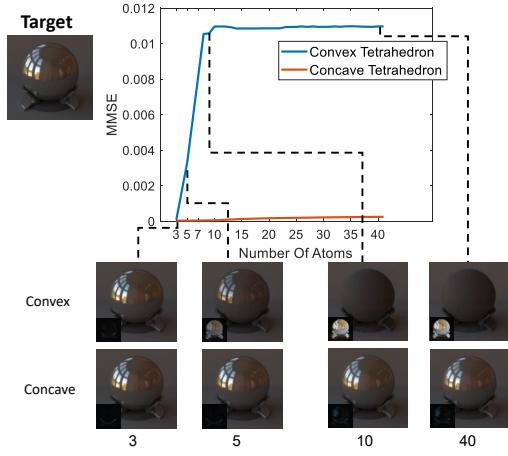


Fig. 3. Simple demonstration of the information available about the BRDF when imaging convex and concave versions of a simple tetrahedron shape. We capture one image of each shape, and use it to solve the optimization problem of Eq. (3), for BRDFs represented by dictionaries of different sizes.

we plot the numerical mean squared BRDF reconstruction error averaged across four different environment maps (all different from the input illumination conditions) and across multiple BRDFs. In Fig. 9, we show corresponding visual comparisons. In both figures, we observe that the concave shape can provide reasonable reconstruction from even just one input image, whereas the convex shape only reaches comparable performance at higher numbers of images. We note that we did not process concave shapes with more than 4 input images, as we noticed the result does not improve much further, and because of the increased computational complexity of the concave cases.

Visualizing the BRDF constraints. To better understand the BRDF constraints provided by the different shapes, we display in Fig. 10 a histogram of the number of times each entry of the 2D bivariate BRDF projection is visited when rendering images for each of the target shapes. In the case of the convex tetrahedron, only three entries of the BRDF

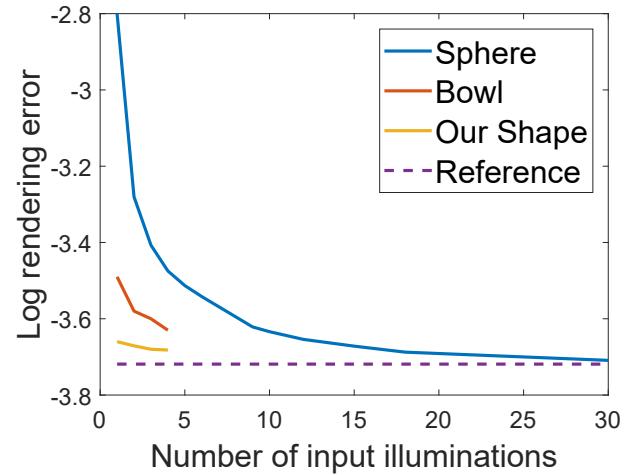


Fig. 4. Reconstruction errors evaluated by rendering multiple BRDFs from the MERL database under multiple environment maps. As a reference point stating the inherent limits of the dictionary representation, we plot the error of the best projection of the MERL BRDFs on the dictionary. Note that when many input illuminations are provided the convex shapes achieve optimal reconstruction. However, our designed concave shape achieves almost the same result with one input image.

table are non-zero as expected, explaining the results of Fig. 3. Even with a convex sphere, only one column of the table is non-zero. This is because the definition of the difference vector \mathbf{d} , and thus θ_d , are both invariant to surface normal and only depend on the input and output directions ω_i, ω_o . As we use orthographic cameras and directional light sources, only one \mathbf{d} value is explaining the full image. By contrast, when rendering the concave shapes we trace a much larger number of paths and hence visit a larger number of entries from the BRDF table. To better visualize this, we show in Fig. 10 separate histograms for single-bounce paths, as well as both single-bounce and two-bounce paths. Even if we ignore higher-order paths, considering the two-bounce paths already provides a much richer set of constraints, as had previously been reported [41].

6 DESIGNING AN INFORMATIVE SHAPE FOR BRDF RECOVERY

The results of the previous section suggest that shapes that include concavities can provide significantly more information about the object's BRDF. This motivates the question: what are optimal shapes for BRDF recovery? Our differentiable rendering framework allows considering more general shapes than previous works that considered this question, which were restricted to convex shapes [73]. In this section, we take some first steps towards devising shapes that facilitate BRDF recovery, focusing on the specific setting of recovery from a single input image.

For this, we start from the histograms in Fig. 10, where we can see that, even with the concave shapes, there are large parts of the bivariate space that are not sampled. We can therefore attempt to design a shape that will maximize this coverage of the bivariate space. Additionally, as paths of low number of bounces carry more energy and are easier to invert, ideally we want to cover most of the space with paths of up to 2 bounces. As a third desideratum, we can emphasize paths where both bounces correspond to specular reflections, as those typically have strong contributions and are very informative about the specular lobe of the BRDF. In specular reflections, ω_i and ω_o are symmetric around the normal, meaning that $\mathbf{h} = \mathbf{n}$ and $\theta_h = 0$. Therefore, these paths correspond to the top row of the BRDF table. Fig. 10 shows that even the concave bowl does not sample all angles of this row. To simplify the analysis, we assume that we are

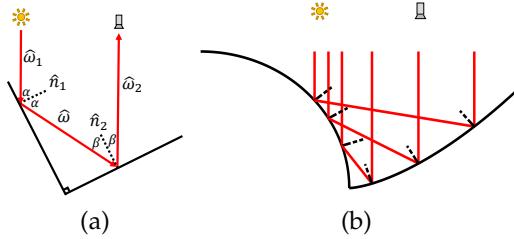


Fig. 5. Designing an informative shape for BRDF recovery. (a) Assuming co-axial illumination and sensing $\omega_{in} = \omega_{out} = [0, 0, 1]$, a two-bounce path reflects at the mirror direction only if the normals at the two points it reflects are orthogonal to each other. (b) Using the observation of (a) we design a new shape that produces mirror reflection paths for all angles. We start from a convex sphere on the left. For each point on this sphere, we trace the mirror reflection path, then set the normal n_2 at the location of the second bounce on the right part of the shape to be orthogonal to the normal n_1 at the location of the first bounce. Finally, we form the full shape by numerical integration of the normals.

illuminating and imaging with a co-axial configuration, with $\omega_i = \omega_o = [0, 0, 1]$. Consider a two-bounce path reflecting at two planar facets with normals $\mathbf{n}_1, \mathbf{n}_2$, as illustrated in Fig. 5(a). Under our illumination and imaging conditions, in order for both reflections to be specular, it is easy to prove that the two facets must be orthogonal to each other, that is, $\langle \mathbf{n}_1, \mathbf{n}_2 \rangle = 0$.

We can use this intuition to construct a curved surface that satisfies this property everywhere. The construction is shown in Fig. 5(b): The surface is composed of a convex part on the left, which we select to be a sphere, and a concave part on the right whose normals are selected to match the above orthogonality constraint. By integrating the normals, we can numerically compute the shape of the concave part,

as shown in Fig. 5(b). The full 3D shape can be produced by extruding or revolving the derived composite 2D profile.

In the last column of Fig. 10, we see that using this new shape indeed allows us to sample all θ_d values of the top row. In the reconstruction results of Fig. 4 and Figs. 6-9, we also observe that this shape provides more accurate estimations of the target BRDF, even from only one input image. Visually, the main improvements come in the form of better reproduction of the BRDF's diffuse color, and better matching of the highlights at the rim of the sphere, which are due to grazing angle reflections.

7 DISCUSSION

We have shown how differentiable rendering can be used for reflectometry from concave shapes producing multiple interreflections. Such shapes have previously been avoided for reflectometry, because of the difficulty in inverting the interreflections to infer BRDF values. We have demonstrated that, at the cost of increased computation, the use of concave shapes can help produce high-fidelity BRDF estimates from much fewer image measurements than what is required when using convex shapes, due to the rich set of constraints provided on the BRDF when considering the higher-order interreflection paths. We have also taken first steps towards designing shapes that maximize the number of constraints on the BRDF available in the captured images.

We expect that our differentiable rendering framework will motivate follow-up research on reflectometry from interreflections. One promising direction for exploration is the design of shapes that, by combining convex and concave parts, facilitate reflectometry. The results of Sec. 6 take a first step in this direction, by highlighting one useful property that a shape well-suited for reflectometry should have: it should provide coverage of all angles in the bivariate BRDF representation. Designing shapes for reflectometry will require devising optimality criteria that take into account not just this property, but also fabrication and application-specific constraints. Finally, we hope that our findings will encourage the exploration and development of diverse real-world reflectometry systems that utilize interreflections for the efficient and accurate collection of reflectance data.

ACKNOWLEDGMENTS

This work was supported by ERC 635537, ISF 1046-14, Oldendorff Minerva Center of the Technion, NSF Expeditions award 1730147, NSF award IIS-1900849, and a gift from the AWS Cloud Credits for Research program.

REFERENCES

- [1] F. E. Nicodemus, "Reflectance nomenclature and directional reflectance and emissivity," *Applied Optics*, 1970.
- [2] W. Matusik, H. Pfister, M. Brand, and L. McMillan, "Efficient isotropic brdf measurement," in *EGRW*, 2003.
- [3] W. Matusik, "A data-driven reflectance model," Ph.D. dissertation, Massachusetts Institute of Technology, 2003.
- [4] R. Vidaurre, D. Casas, E. Garces, and J. Lopez-Moreno, "Brdf estimation of complex materials with nested learning," in *(WACV)*, 2019.
- [5] D. Antensteiner and S. Stolc, "Full brdf reconstruction using cnns from partial photometric stereo-light field data," in *(CVPR)*, 2017.

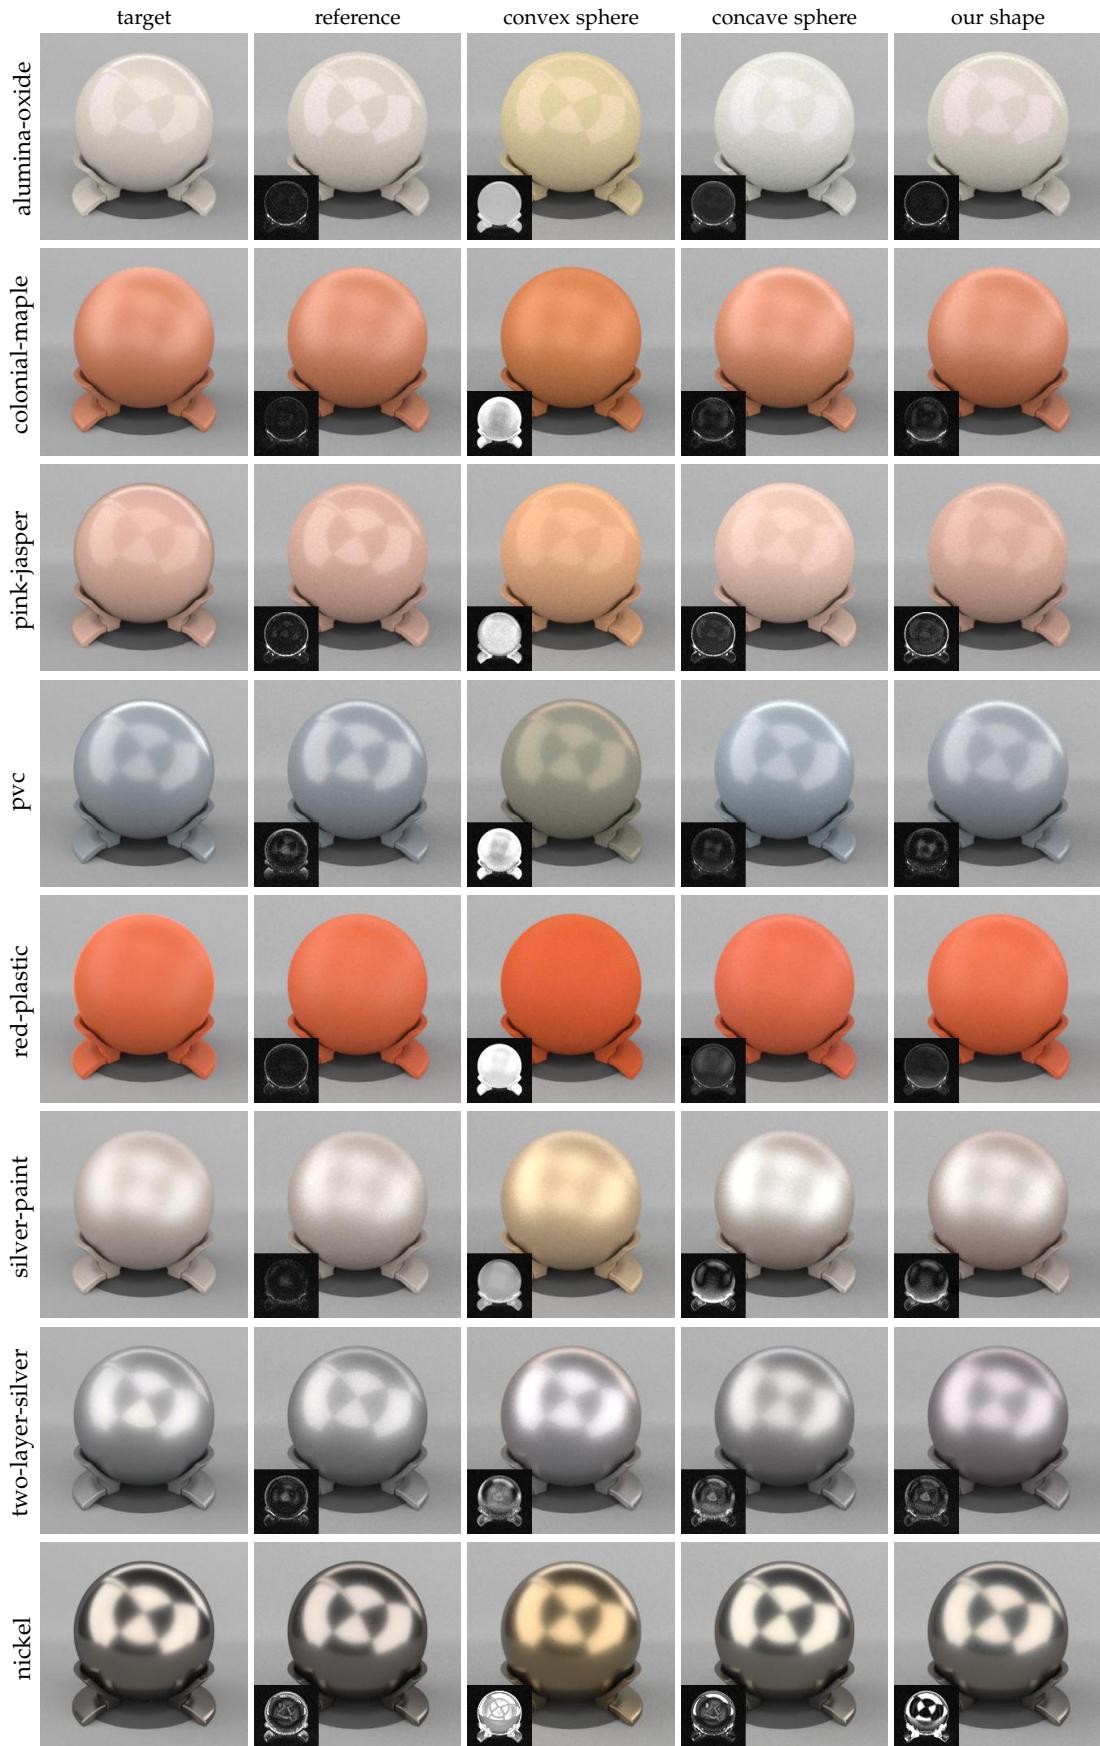


Fig. 6. Reconstructed BRDF part 1, rendered under the checkerboard environment map. Insets show error maps with respect to the target.



Fig. 7. Reconstructed BRDF part 2, rendered under the museum environment map. Insets show error maps with respect to the target.

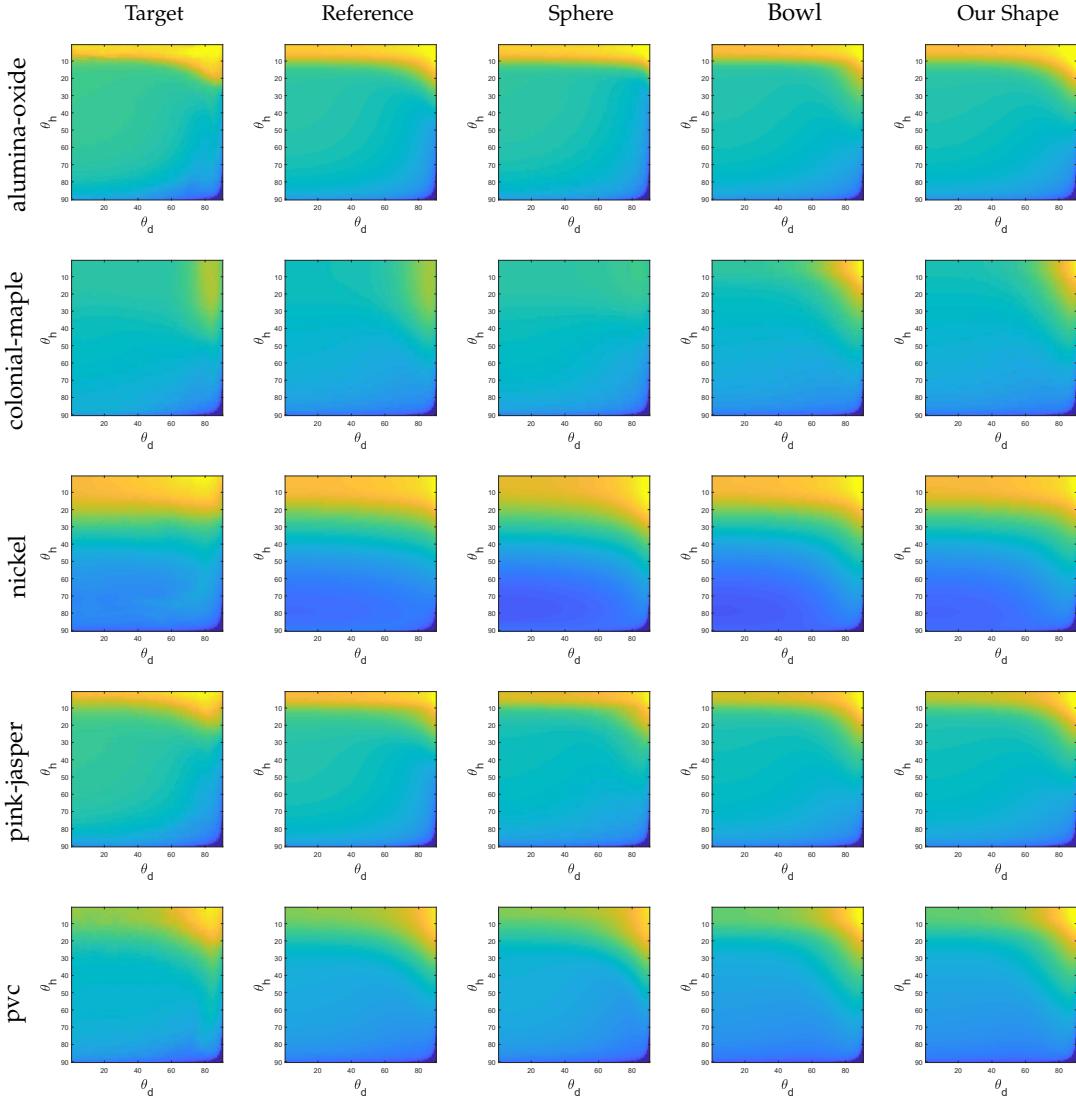


Fig. 8. Reconstructed bivariate tables of the BRDFs.

- [6] A. e. a. Meka, "Deep reflectance fields: High-quality facial reflectance field inference from color gradient illumination," *ACM Trans. Graph.*, 2019.
- [7] T. Weyrich, J. Lawrence, H. P. Lensch, S. Rusinkiewicz, T. Zickler *et al.*, "Principles of appearance acquisition and representation," *Foundations and Trends® in Computer Graphics and Vision*, 2009.
- [8] W. Erb, "Computer-controlled gonioreflectometer for the measurement of spectral reflection characteristics," *Applied optics*, 1980.
- [9] S. R. Marschner, S. H. Westin, E. P. F. Lafontaine, K. E. Torrance, and D. P. Greenberg, "Image-based BRDF measurement including human skin," in *EGWR*, 1999.
- [10] W. Matusik, H. Pfister, M. Brand, and L. McMillan, "A data-driven reflectance model," *ACM Transactions on Graphics*, vol. 22, no. 3, pp. 759–769, Jul. 2003.
- [11] J. Dupuy and W. Jakob, "An adaptive parameterization for efficient material acquisition and rendering," *ACM Transactions on graphics (TOG)*, 2019.
- [12] A. Ghosh, S. Achutha, W. Heidrich, and M. O'Toole, "Brdf acquisition with basis illumination," *ICCV*, 2007.
- [13] B. Tunwattanapong, G. Fyffe, P. Graham, J. Busch, X. Yu, A. Ghosh, and P. Debevec, "Acquiring reflectance and shape from continuous spherical harmonic illumination," *ACM Transactions on graphics (TOG)*, 2013.
- [14] M. Aittala, T. Weyrich, and J. Lehtinen, "Practical svbrdf capture in the frequency domain," *ACM Trans. Graph.*, 2013.
- [15] S. Rusinkiewicz, "A new change of variables for efficient BRDF representation," in *Rendering Techniques*, 1998.
- [16] F. Romeiro and T. Zickler, "Inferring reflectance under real-world illumination," *Technical Report, Harvard University*, 2010.
- [17] K. Nishino, "Directional statistics brdf model," in *ICCV*, 2009.
- [18] M. Chandraker and R. Ramamoorthi, "What an image reveals about material reflectance," *ICCV*, 2011.
- [19] Z. Xu, J. B. Nielsen, J. Yu, H. W. Jensen, and R. Ramamoorthi, "Minimal brdf sampling for two-shot near-field reflectance acquisition," *ACM Transactions on Graphics (TOG)*, 2016.
- [20] J. B. Nielsen, H. W. Jensen, and R. Ramamoorthi, "On optimal, minimal brdf sampling for reflectance acquisition," *ACM Transactions on Graphics (TOG)*, 2015.
- [21] J. Lawrence, S. Rusinkiewicz, and R. Ramamoorthi, "Efficient brdf importance sampling using a factored representation," 2004.
- [22] T.-M. Li, M. Aittala, F. Durand, and J. Lehtinen, "Differentiable monte carlo ray tracing through edge sampling," in *SIGGRAPH Asia*, 2018.
- [23] M. Nimier-David, D. Vicini, T. Zeltner, and W. Jakob, "Mitsuba 2: A retargetable forward and inverse renderer," *ACM Transactions on Graphics (TOG)*, 2019.
- [24] C. Zhang, L. Wu, C. Zheng, I. Gkioulekas, R. Ramamoorthi, and S. Zhao, "A differential theory of radiative transfer," *ACM Transactions on Graphics (TOG)*, 2019.
- [25] C. Che, F. Luan, S. Zhao, K. Bala, and I. Gkioulekas, "Inverse transport networks," *arXiv*, 2018.
- [26] Z. Hui, K. Sunkavalli, J.-Y. Lee, S. Hadap, J. Wang, and A. C. Sankaranarayanan, "Reflectance capture using univariate sampling of BRDFs," *ICCV*, 2017.

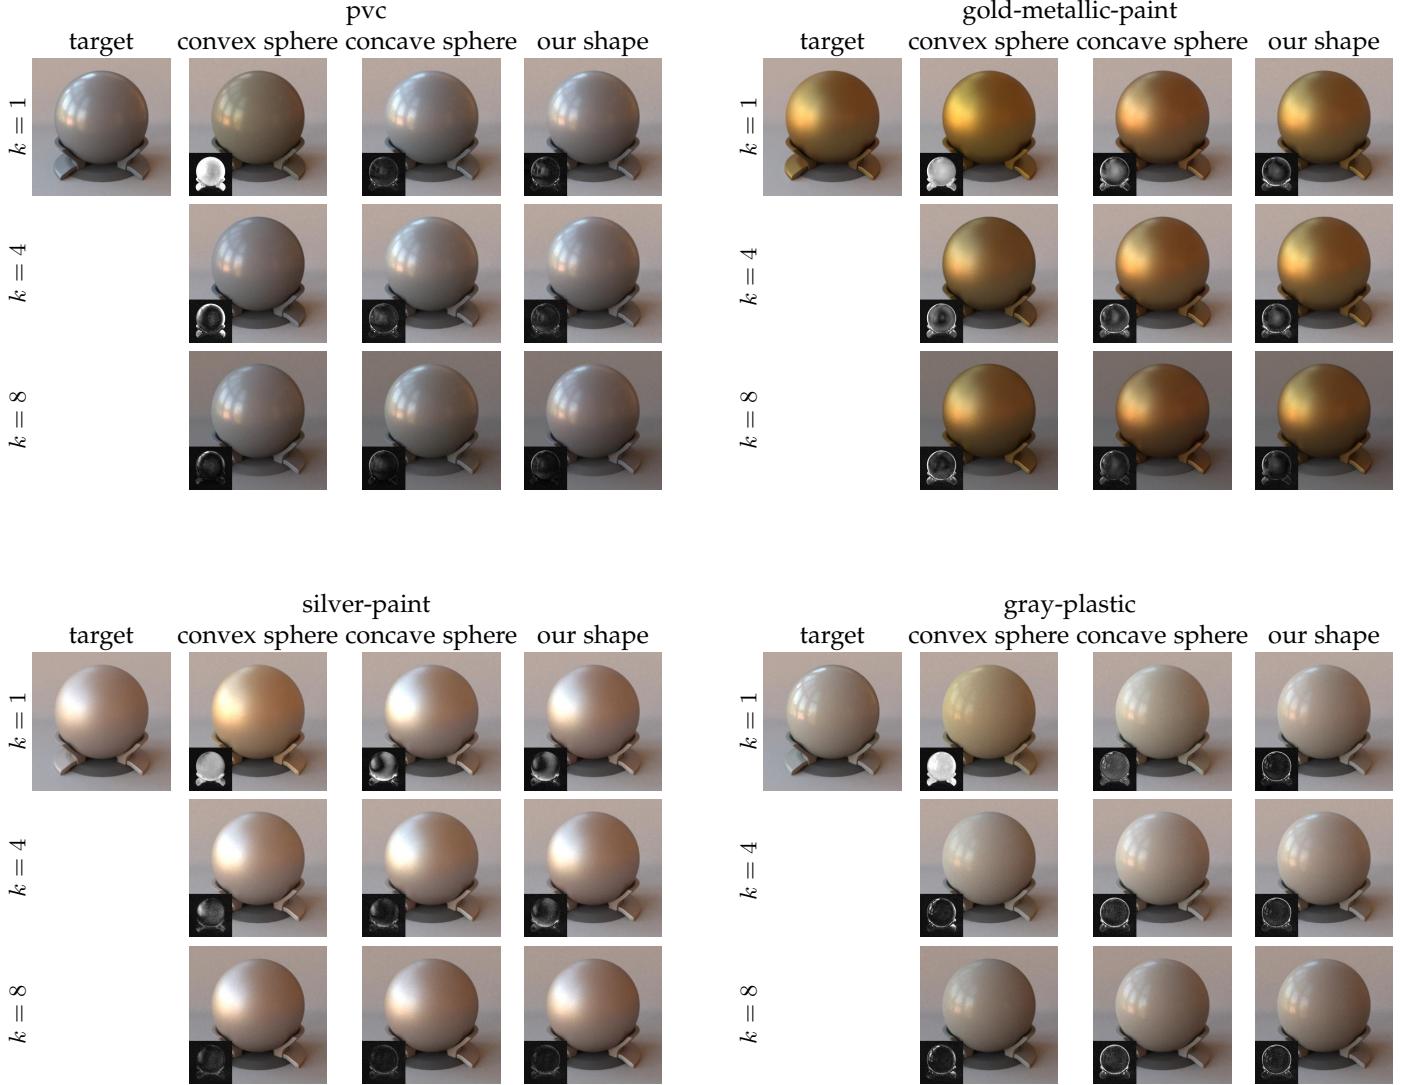


Fig. 9. Reconstruction with multiple input images. The results using the convex sphere improve significantly when more input images are provided, whereas for the concave shapes the improvement is minor.

- [27] F. Romeiro, Y. Vasilyev, and T. Zickler, "Passive reflectometry," *ECCV*, 2008.
- [28] V. Havran, J. Filip, and K. Myszkowski, "Perceptually motivated brdf comparison using single image," in *Computer graphics forum*, vol. 35, no. 4. Wiley Online Library, 2016, pp. 1-12.
- [29] J. Wills, S. Agarwal, D. Kriegman, and S. Belongie, "Toward a perceptual space for gloss," *ACM Transactions on graphics (TOG)*, vol. 28, no. 4, pp. 1-15, 2009.
- [30] F. Romeiro and T. Zickler, "Blind reflectometry," *ECCV*, 2010.
- [31] S. Lombardi and K. Nishino, "Reflectance and illumination recovery in the wild," *IEEE transactions on pattern analysis and machine intelligence*, 2015.
- [32] Z. Hui and A. C. Sankaranarayanan, "Shape and spatially-varying reflectance estimation from virtual exemplars," *IEEE Trans. Pattern Analysis and Machine Intelligence (PAMI)*, 2017.
- [33] Z. Li, Z. Xu, R. Ramamoorthi, K. Sunkavalli, and M. Chandraker, "Learning to reconstruct shape and spatially-varying reflectance from a single image," *ACM Transactions on Graphics (TOG)*, 2019.
- [34] K. Rematas, T. Ritschel, M. Fritz, E. Gavves, and T. Tuytelaars, "Deep reflectance maps," *CVPR*, 2016.
- [35] X. Li, Y. Dong, P. Peers, and X. Tong, "Modeling surface appearance from a single photograph using self-augmented convolutional neural networks," *ACM Transactions on Graphics (TOG)*, 2017.
- [36] M. K. Chandraker, F. Kahl, and D. J. Kriegman, "Reflections on the generalized bas-relief ambiguity," *CVPR*, 2005.
- [37] S. Liu, T.-T. Ng, and Y. Matsushita, "Shape from second-bounce of light transport," *ECCV*, 2010.
- [38] S. K. Nayar, K. Ikeuchi, and T. Kanade, "Shape from interreflections," *IJCV*, 1991.
- [39] N. Naik, S. Zhao, A. Velten, R. Raskar, and K. Bala, "Single view reflectance capture using multiplexed scattering and time-of-flight imaging," in *Proceedings of the 2011 SIGGRAPH Asia Conference*, 2011, pp. 1-10.
- [40] N. Naik, C. Barsi, A. Velten, and R. Raskar, "Estimating wide-angle, spatially varying reflectance using time-resolved inversion of backscattered light," *JOSA A*, vol. 31, no. 5, pp. 957-963, 2014.
- [41] C.-Y. Tsai, A. Veeraraghavan, and A. C. Sankaranarayanan, "Shape and reflectance from two-bounce light transients," *ICCP*, 2016.
- [42] M. M. Loper and M. J. Black, "Opendr: An approximate differentiable renderer," in *European Conference on Computer Vision*. Springer, 2014.
- [43] H. Kato, Y. Ushiku, and T. Harada, "Neural 3d mesh renderer," 2017.
- [44] V. Deschaintre, M. Aittala, F. Durand, G. Drettakis, and

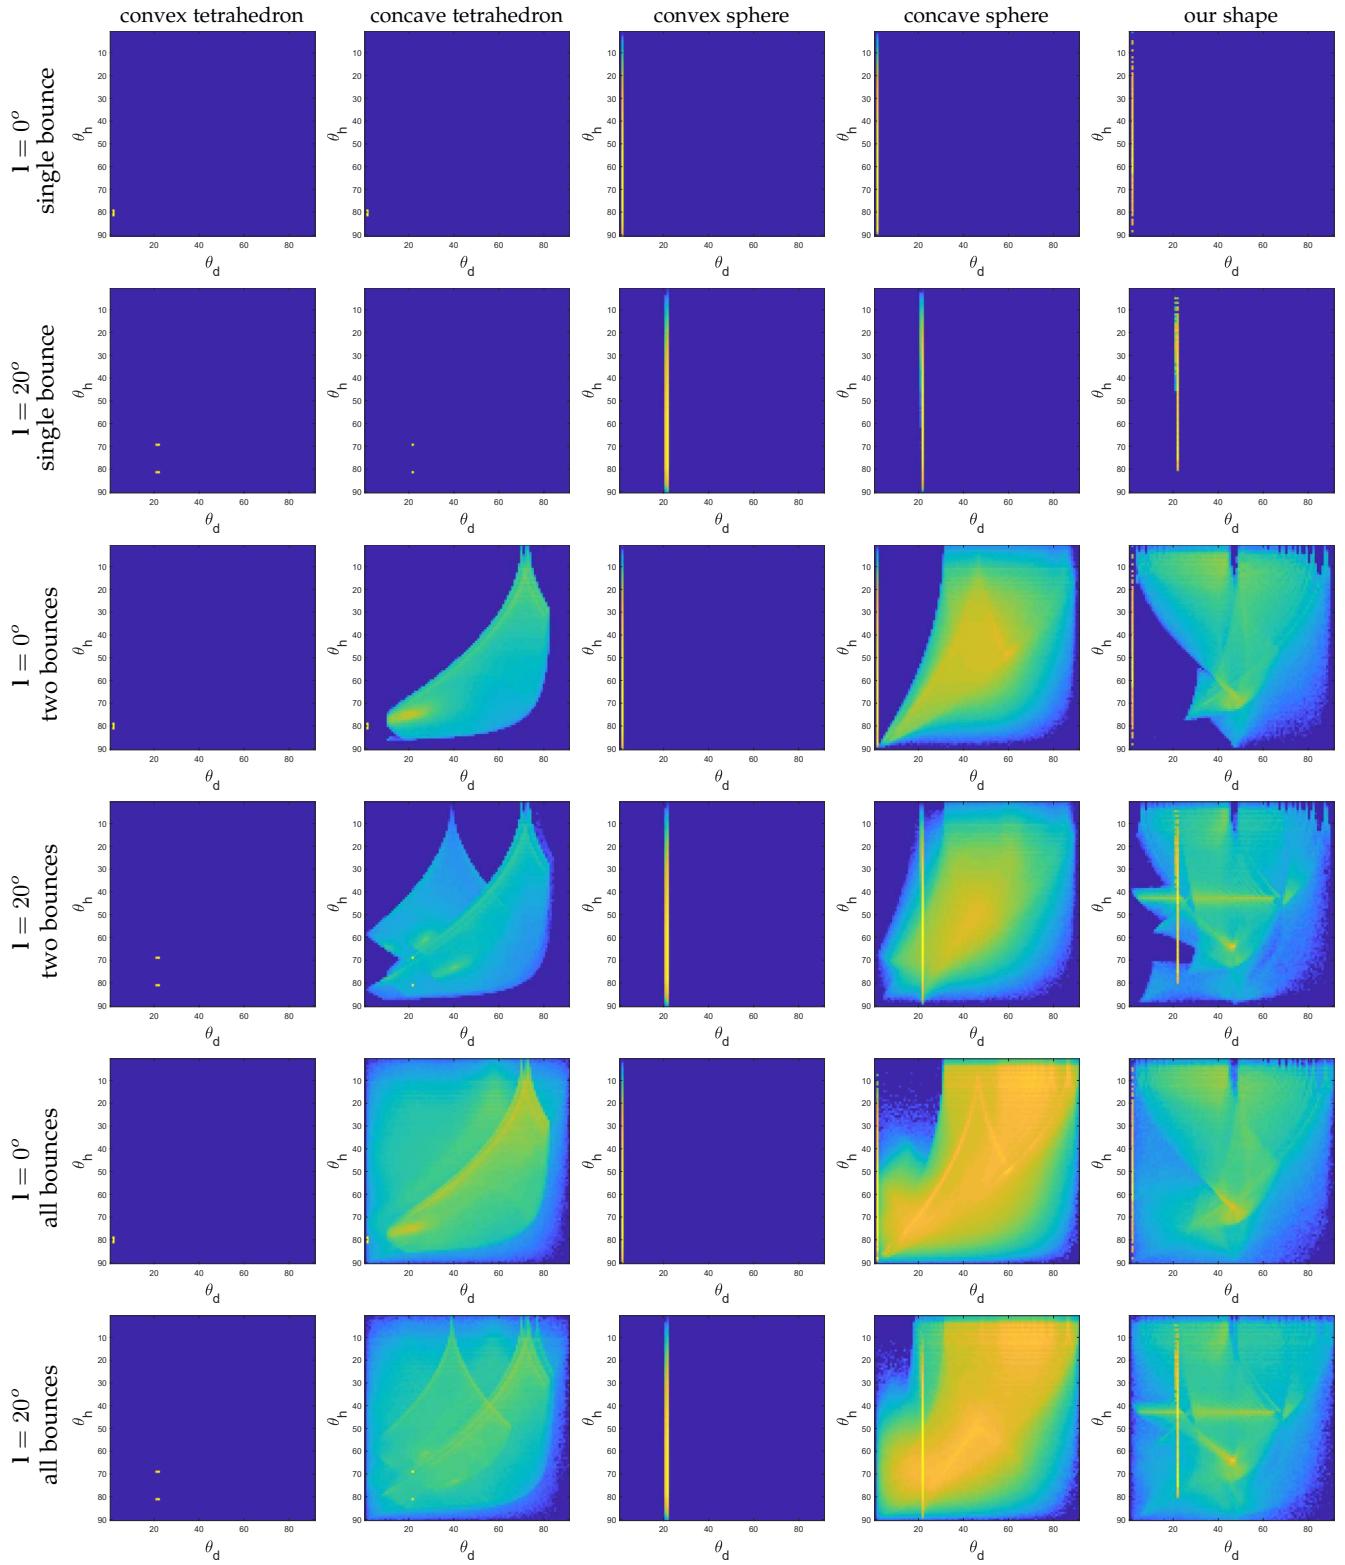


Fig. 10. Histogram of the number of times each BRDF entry is used when rendering input images. The figures visualize the log of the number of visits.

A. Bousseau, "Single-image svbrdf capture with a rendering-aware deep network," *ACM TOG*, 2018.

[45] G. Liu, D. Ceylan, E. Yumer, J. Yang, and J.-M. Lien, "Material editing using a physically based rendering network," in *ICCV*, 2017.

[46] I. Gkioulekas, S. Zhao, K. Bala, T. Zickler, and A. Levin, "Inverse volume rendering with material dictionaries," *ACM SIGGRAPH Asia*, 2013.

[47] I. Gkioulekas, A. Levin, and T. Zickler, "An evaluation of computational imaging techniques for heterogeneous inverse scattering," *ECCV* 2016.

[48] S. Zhao, L. Wu, F. Durand, and R. Ramamoorthi, "Downsampling scattering parameters for rendering anisotropic media," *ACM Transactions on Graphics (TOG)*, 2016.

- [49] P. Khungurn, D. Schroeder, S. Zhao, K. Bala, and S. Marschner, "Matching real fabrics with micro-appearance models." *ACM Trans. Graph.*, 2015.
- [50] A. Levis, Y. Y. Schechner, A. Aides, and A. B. Davis, "Airborne three-dimensional cloud tomography," *CVPR*, 2015.
- [51] D. Azinović, T.-M. Li, A. Kaplanyan, and M. Nießner, "Inverse path tracing for joint material and lighting estimation," *CVPR*, 2019.
- [52] S. Lombardi and K. Nishino, "Radiometric scene decomposition: Scene reflectance, illumination, and geometry from rgbd images," *3DV*, 2016.
- [53] C.-Y. Tsai, A. C. Sankaranarayanan, and I. Gkioulekas, "Beyond volumetric albedo—a surface optimization framework for non-line-of-sight imaging," *CVPR*, 2019.
- [54] S. R. Marschner, "Inverse rendering in computer graphics," *PhD thesis*, 1998.
- [55] G. Patow and X. Pueyo, "A survey of inverse rendering problems," *Computer graphics forum*, 2003.
- [56] K. Nishino and S. Lombardi, "Directional statistics-based reflectance model for isotropic bidirectional reflectance distribution functions," *Optical Society of America*, 2011.
- [57] J. Dupuy, E. Heitz, J.-C. Iehl, P. Poulin, and V. Ostromoukhov, "Extracting microfacet-based brdf parameters from arbitrary materials with power iterations," *Comput. Graph. Forum*, 2015.
- [58] E. Heitz, "Understanding the masking-shadowing function in microfacet-based brdfs," *(JCGT)*, 2014.
- [59] B. Walter, S. R. Marschner, H. Li, and K. E. Torrance, "Microfacet models for refraction through rough surfaces," *Proceedings of the 18th Eurographics conference on Rendering Techniques*, 2007.
- [60] A. Ngan, F. Durand, and W. Matusik, "Experimental validation of analytical brdf models," 2004.
- [61] K. E. Torrance and E. M. Sparrow, "Theory for off-specular reflection from roughened surfaces," *J. Opt. Soc. Am.*, 1967.
- [62] M. Ashikmin, S. Premož, and P. Shirley, "A microfacet-based brdf generator," in *SIGGRAPH*, 2000.
- [63] J. F. Blinn, "Models of light reflection for computer synthesized pictures," *SIGGRAPH Comput. Graph.*, 1977.
- [64] G. J. Ward, "Measuring and modeling anisotropic reflection," *SIGGRAPH Comput. Graph.*, 1992.
- [65] C. Schlick, "An inexpensive brdf model for physically-based rendering," *Computer Graphics Forum*, 1994.
- [66] P. Dutre, P. Bekaert, and K. Bala, *Advanced global illumination*. AK Peters/CRC Press, 2006.
- [67] E. Veach, *Robust monte carlo methods for light transport simulation*. Stanford University PhD thesis, 1997, no. 1610.
- [68] E. Veach and L. Guibas, "Bidirectional estimators for light transport," in *Photorealistic Rendering Techniques*. Springer, 1995, pp. 145–167.
- [69] W. Jakob, "Mitsuba renderer," 2010, <http://www.mitsuba-renderer.org>.
- [70] J. Kivinen and M. K. Warmuth, "Exponentiated gradient versus gradient descent for linear predictors," *information and computation*, 1997.
- [71] A. Beck and M. Teboulle, "Mirror descent and nonlinear projected subgradient methods for convex optimization," *Operations Research Letters*, vol. 31, pp. 167–175, 05 2003.
- [72] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [73] J. Filip and F. Maile, "In the search of an ideal measurement geometry for effect coatings," *Pigment and Colour Science Forum*, 2018.



Sai Praveen Bangaru is a PhD student at MIT CSAIL since 2019, advised by Prof. Fredo Durand. He received his Masters degree from the Computer Science Department at Carnegie Mellon University, mentored by Prof. Ioannis Gkioulekas, where he based his thesis on differentiable rendering. Previously, he completed his undergraduate degree at the Indian Institute of Technology, Madras.



Anat Levin is an Associate Prof. at the department of Electrical Engineering, Technion, Israel, doing research in the field of computational imaging. She received her Ph.D. from the Hebrew University at 2006. During the years 2007-2009 she was a postdoc at MIT CSAIL, and during 2009-2016 she was an Assistant and Associate Prof. at the department of Computer Science and Applied Math, the Weizmann Inst. of Science.



Ioannis Gkioulekas is an Assistant Professor at the Robotics Institute, Carnegie Mellon University. He received M.Sc. and Ph.D. degrees from Harvard University, and a Diploma of Engineering from the National Technical University of Athens. His research interests are on computational imaging, physics-based rendering, and differentiable rendering. He has received the Best Paper Award at CVPR 2019, and a Sloan Research Fellowship.



Kfir Shem Tov is a MSc graduate from the department of Electrical Engineering, Technion, Israel, supervised by Prof Anat Levin. He received his B.Sc in Electrical Engineering from Bar Ilan University, Israel. Currently an algorithm engineer at MultiVu - Advanced Photography Technologies.