Shape and Reflectance from Two-Bounce Light Transients

Chia-Yin Tsai[†], Ashok Veeraraghavan[‡] and Aswin C. Sankaranarayanan[†]
[†] Dept. of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA
[‡] Dept. of Electrical and Computer Engineering, Rice University, Houston, TX

Abstract

Computer vision and image-based inference have predominantly focused on extracting scene information by assuming that the camera measures direct light transport (i.e., single-bounce light paths). As a consequence, strong multi-bounce effects are treated typically as sources of noise and, in many scenarios, the presence of such effects can result in gross errors in the estimates of shape and reflectance. This paper provides the theoretical and algorithmic foundations for shape and reflectance estimation from two-bounce light transients, i.e., scenarios where photons from a light source interact with the scene exactly twice before reaching the sensor. We derive sufficient conditions for exact recovery of shape and reflectance given lengths and intensities associated with two-bounce light paths. We also develop algorithms for recovery of shape and reflectance, and validate these on a range of simulated scenes.

1. Introduction

The vast majority of shape estimation techniques rely on information derived from direct or single-bounce light paths, where photons from a source are observed after a single reflection/scattering event. This includes geometric triangulation [6], as used in stereo and structured light [19], time of flight (ToF), as used in LIDAR, photonic mixer devices, radar, and RF-based depth acquisition, as well as photometric methods, such as photometric stereo and polarization-based 3D. All of these methods, while extremely popular, are limited to settings where the multi-bounce effects are minimal and can effectively be ignored.

Real-world scenes often exhibit spatial geometries and reflectances that lead to strong multi-bounce interactions. Even seemingly simple shapes like concave bowls, v-grooves, and corners, as well as reflectances associated with common materials like wood, metal, and ceramic, result in multi-bounce interactions that can sometimes overwhelm single-bounce photons. This corrupts the information encoded in single-bounce light paths, resulting in gross errors in shape estimation [3, 4].

Traditional methods for shape estimation treat multi-path effects as a source of noise that must be tolerated or minimized. In contrast, we argue that multi-path effects provide a rich encoding of the 3D properties of the scene and that this information can be exploited. Further, a framework of shape and reflection estimation from multi-bounce light paths can provide capabilities that are unprecedented and not achievable using just single-bounce light paths. The focus of this paper is the estimation of shape and reflectance information from the intensity and the time of flight of each observed second-bounce light path.

Motivation. There are several important benefits to be obtained by going beyond single-bounce light paths.

(i) Number of light-paths. There are significantly larger number of light paths with two bounces than with a single bounce. For a scene with N surface elements, there are at most N single-bounce light paths while there can potentially be as many as $\binom{N}{2}$ individual second-bounce light paths. Incorporating second-bounce information can provide additional constraints on the depth recovery problem which often leads to increased robustness in shape estimates.

(ii) Ability to handle arbitrary reflectance functions. Many specular and shiny materials concentrate incident light along certain directions. This often results in the absence of single-bounce light paths between the source and the sensor (commonly observed as holes in 3D scans). Here, the use of second and higher-order bounces can potentially alleviate the lack of information in the single-bounce light paths.

(iii) Angular sampling for reflectometry. Single-bounce based reflectometry often requires light stages that completely encompass the target in order to obtain a diverse angular sampling of the scene's bidirectional reflectance distribution function (BRDF) [18]. In contrast, reflectometry using two-bounce light paths can exploit the large angle subtended by scene points with one another; hence, we can potentially obtain reflectance estimates from two-bounce light paths even with a collocated source and sensor.

Contributions. *This paper provides the theoretical and algorithmic foundations for shape and reflectance estimation from two-bounce light paths, i.e., light paths where photons from a light source interact with the scene exactly twice before reaching the sensor.* While there is a rich body of work associated with all aspects of shape and reflectance estimation from single-bounce light paths, little is known in terms of the capabilities and limitations of two-bounce light paths. This paper is among the very first to address this problem. We assume that for each observed second-bounce light path, the intensity and the time of travel along that path is available. Our specific contributions are as follows.

- *Two-bounce shape estimation.* We propose a formulation for the systematic study of shape estimation from two-bounce light paths and develop a graph-based framework to represent the information in such light paths.
- *Uniqueness of shape estimation.* We establish sufficient conditions for uniqueness of shape estimation given only second-bounce light paths. Our sufficient conditions are intricately tied to the topology of the graph that characterize the available second-bounce light paths.
- *Algorithms for shape estimation.* We provide novel algorithms that estimate shape given lengths of second-bounce light paths.
- *Algorithms for reflectometry.* Given the radiance associated with second-bounce light paths, we propose novel algorithms for BRDF estimation.

While the focus of this paper is to exclusively characterize the shape and reflectance information embedded in the two-bounce light paths, we believe that in the long term, holistic methods that simultaneously exploit the single- as well as multi-bounce ideas will result in significant performance improvements.

2. Prior work

There are three areas of research that are intimately related to our results: ToF imaging, shape estimation from ToF, and reflectometry.

Time-of-flight imaging. ToF sensors measure the time that a pulse of light takes to travel from a source to the scene and back. Since the velocity of light is large (3×10^8m/s), the associated time scales of travel times are invariable small, necessitating either ultra-fast streak cameras that provide direct measurement of travel time [22] or photonic mixer devices (PMDs) that measure phase shift of the modulated light source signal with the received signal [7]. ToF using PMDs is popular because of the longer exposure time and low cost of the sensor; however, the maximum temporal resolution achieved by these devices is currently limited to about 100 picoseconds. More recently, Gkioulekas et al. [2] use the interferometric techniques along with ToF measurements to obtain depth measurements with resolution in tens of microns.

Transient light transport. The ideas introduced in this paper are intimately related to the transient light transport operator. The concept of 5D transient light transport was first introduced in [17]. For an arbitrary incident direction (two dimensions), and an arbitrary exitant direction (two dimensions), the transient light transport is the one dimensional impulse response function of the scene. Thus, the transient light transport matrix is five-dimensional and contains information about all the multiple bounces of light between the illumination source and the camera sensor. The availability of this 5D light transport matrix is the building block of our proposed shape and reflectance estimation. While O'Toole et al. [17] use the direct component in the 5D matrix to estimate the shape of the scene, we explore utilizing the second-bounce information from the 5D matrix to estimate both shape and reflectance.

Mitigating multi-bounce effects. Most traditional computer vision techniques assume that the light transport is dominated by direct or single-bounce light paths. Scenes and materials that contribute to multiple bounces (sub-surface scattering, specular reflections, V-grooves, corners, etc.) usually result in gross errors in traditional shape estimation techniques. An early approach to mitigating multi-bounce effects was to estimate the shape of a concave object iteratively while progressively accounting for inter-reflections [14]. Nayar et al. [15] propose an approach to separate the direct components of light from global (multi-bounce) components and show successful photometric stereo even under inter-reflections. Gupta et al. [4] show that with an appropriate choice of spatial codes in structured light, the effects of both sub-surface scattering and inter-reflections can be minimized resulting in robust shape recovery for material exhibiting moderate global light transport. In a similar approach, Gupta et al. [5] show that with the appropriate choice of temporal modulation frequencies in a ToF sensor, the effects of global light transport effects can be significantly suppressed. While these methods focus on reducing the impact of multi-bounce light transport on the estimated shape, our approach is fundamentally different: *we show that second-bounce light transport actually contains sufficient information to reliably estimate shape from exclusively such light paths.*

Shape from multi-bounce light paths. There has been very little prior work in understanding the relationship between scene shape and the associated multi-bounce light paths. Liu et al. [11] recover the shape of a Lambertian scene from the form factor of two-bounce light paths. By using local planarity and common elements in different two-bounce light paths, the scene geometry can be recovered

without scaling ambiguity. For a transparent object, light changes direction twice, Kutulakos and Steger [10] show that the shape of the transparent object can be recovered. In [9,21], three-bounce light paths, with first and third bounces being visible, can be used to reconstruct a non-line-of-sight object. From ToF imaging, the path length from the first bounce to the third bounce is known. Therefore, by using multiple three-bounce light paths, the shape of the hidden surface can be revealed using elliptic tomography.

Reflectance from multi-bounce light paths. Naik et al. [13] use three-bounce light paths to probe the BRDF of the material. The availability of time-of-flight data limits the location of participating patches, thus makes the BRDF estimation tractable. Also, the use of three-bounce light path enlarges the coverage of the parameters of bi-variate model of BRDF. However, in their setup, the first and third bounces are assumed to be Lambertian, whereas in our setup, no restriction on the material is needed.

3. Overview

In this section, we provide a brief overview of the problem setup and formulate the "light-path graph", a construct critical to deriving the main results of this paper.

Problem setup. Given a scene, we first discretize the scene into N points with 3D locations denoted as $\{\mathbf{v}_1, \ldots, \mathbf{v}_N\}$, in a right-handed camera-centric coordinate system. The principal point of the camera is at the origin and its optical axis is aligned to the positive z-axis. The 3D location of each scene point can, hence, be expressed as $\mathbf{v}_k = d_k \mathbf{i}_k$, where $d_k \geq 0$ is the depth of the scene point and $\mathbf{i}_k$ is the unit-norm vector that provides the orientation of the scene point from the viewpoint of the camera.

Two-bounce light-path measurements. We assume that we can measure the following two signals:

- *Path lengths.* We can measure the time taken for light along single and two-bounce paths. Specifically, $\tau_{pk}, p, k \in \{1, \ldots, N\}$ measures the time taken to traverse the two-bounce light path from the origin to scene point p to scene point k and back to the origin (see Figure 1). Note that, by Helmholtz's reciprocity, $\tau_{pk} = \tau_{kp}$. Further, the length associated with these paths is simply given as $c\tau_{pk}$, where c is the speed of light.
- *Light transport.* We can measure light transport along two-bounce paths; we can measure L_{pk}, the fraction of light that is observed along $\mathbf{i}_k$ via the two-bounce light path defined by the scene points $\mathbf{v}_p$ and $\mathbf{v}_k$, when a unit-lux of light is shined along $\mathbf{i}_p$.

Obtaining light-path measurements. Both measurements, path lengths and light transport associated with two-bounce paths, can be acquired using an imaging system that is capable of measuring 5D transient light transport [17]. In particular, depending upon the desired spatial and temporal resolutions, we can either use a femto-second laser and a streak camera (as in [21]) or use a laser diode and a photonic mixer device (as in [8]). Consider a co-located projector-camera system with the transient response observed at a pixel p when a different pixel k is illuminated (with a temporal dirac), the first peak of the received time profile often corresponds to a two-bounce path. This is because the two-bounce path is the shortest path linking the patch illuminated by pixel p, and the patch observed by pixel k. The same approach works for non-co-located systems as well. However, the first peak could be from either a single or two-bounce. Single-bounce paths can be identified using: (i) projector pixel p and camera pixel k have to lie on epipolar lines, (ii) the distance measured by triangulation should be consistent with the time-of-flight for the first peak. If these two are simultaneously satisfied, then that corresponds to a single-bounce path between p and k; otherwise, it is most likely to be a two-bounce light path.

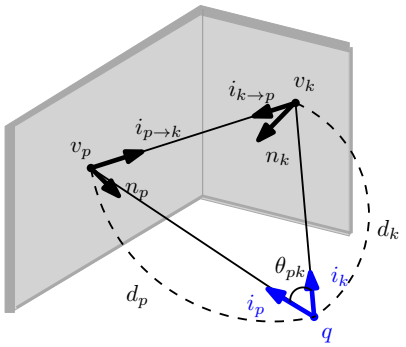

Figure 1. **Shape estimation using two-bounce time of flight information.** A light path originates from a point $\mathbf{q}$ first bounces off a scene point $\mathbf{v}_p$ then traverses to another scene point $\mathbf{v}_k$ and enters the sensor at $\mathbf{q}$. The variables labeled in blue are known and form the inputs to the shape estimation problem.

Problem statement. Given the path lengths and light transport associated with a collection of two-bounce light paths, our goal is to estimate the scene depths $\{d_1, \ldots, d_N\}$ as well as the reflectance function of the scene, under the assumption that all scene points have the same BRDF. To better describe the problem, we characterize the inputs using a graph.

Light-path graph. We construct a graph $G = (V, E)$, where the vertex set $V = \{1, \ldots, N\}$ corresponds to the scene points. An edge (p, k) suggests that the path-length and the light transport associated with the two-bounce path involving the p and k can be observed. The edge $(p, k) \in E$ is also endowed with two observations: the path-length $c\tau_{pk}$ and the light transport L_{pk}, both corresponding to the two-

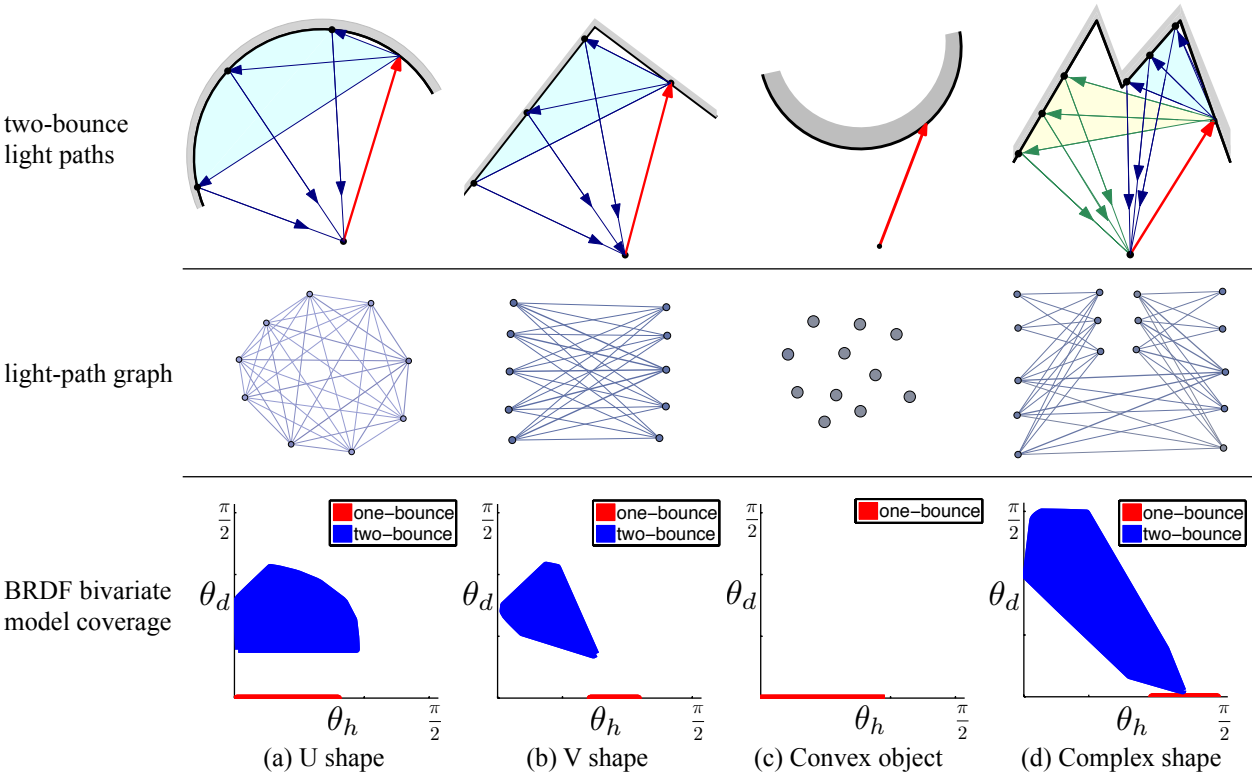

Figure 2. Sample light-path graph and BRDF half/difference angle coverage for objects with different shapes. Top row: observable two-bounce light paths of different objects. Middle row: observable light paths can be represented by a light-path graph. Bottom row: the BRDF coverage of one-bounce and two-bounce light paths.

bounce light path. Note that, by Helmholtz reciprocity, the graph is non-directional, i.e., the observations associated with the edge (a, b) are identical to that of (b, a). We refer to the graph G as the light-path graph.

It is illustrative to consider the topology of light-path graphs associated with a few classes of scenes (see Fig. 2).

- *Convex objects.* Convex objects have no two-bounce light paths. Hence, the graph G is simply a collection of vertices with no edges (see Fig. 2 (c)).

- *Concave objects.* For a generic concave shape — for example a cup — a two-bounce light path can exist between any two scene points. Hence, in the general setting, the graph is fully-connected or a clique (see Fig. 2 (a)). However, in practice, in addition to the shape, the reflectance of the scene also plays an important role in determining the topology. For example, for a concave shape with mirror reflectance, the topology of G reduces to isolated edges, i.e., if an edge (p, k) exists, then there are no other edges involving the vertices p and k.

- *V-grooves.* For a generic V-groove, two-bounce light paths can only exist between scene points that belong to opposite sides of the groove. Hence, we can divide the vertex set V into two non-intersecting sub-sets V_1 and V_2 such that edges can only occur between a vertex in V_1 and a vertex in V_2. Thus V-grooves result in bipartite light-path graphs (see Fig. 2 (b)).

The light-path graph provides a succinct characterization of all available inputs and its topology is central to the results derived in this paper.

4. Uniqueness of shape recovery

In this section, we provide guarantees for uniqueness of shape recovery given lengths of two-bounce light paths. Specifically, given a light-path graph (or equivalently, a collection of two-bounce light paths and their lengths), can there be multiple shapes that satisfy the path length constraints? The following theorem provides a sufficient condition for uniqueness of shape given a light-path graph.

Theorem 1. *If each connected component*[1] *of the light-path graph G contains either a cycle with odd number of vertices or two cycles with even number of vertices, then the depth of all vertices can be determined uniquely from the available two-bounce path lengths.*

The rest of this section is devoted to the proof of Theorem 1. Note that this result is for the noiseless case where the two-bounce light path lengths are known exactly.

[1] A connected component of a graph is defined as a maximally connected sub-graph.

Encoding of depths in two-bounce path length. We now derive the relationships between the depths of scene points and the length of a two-bounce light path.

Consider the two-bounce light path associated with two scene points $\mathbf{v}_p$ and $\mathbf{v}_k$ with unknown depths d_p and d_k, and known directions $\mathbf{i}_p$ and $\mathbf{i}_k$, respectively (see Fig. 1). From law of cosines, the total light-path length is given as

$$c\tau_{pk} = d_p + d_k + \sqrt{d_p^2 + d_k^2 - 2d_p d_k \cos\theta_{pk}}. \quad (1)$$

Squaring and rearranging (1), we can find the following constraint on d_p and d_k:

$$2d_p d_k(-\cos\theta_{pk} - 1) + 2c\tau_{pk}(d_p + d_k) = c^2\tau_{pk}^2. \quad (2)$$

Therefore, d_k can be expressed as a function of d_p as

$$d_k = \frac{c^2\tau_{pk}^2 - 2c\tau_{pk}d_p}{2d_p(-\cos\theta_{pk} - 1) + 2c\tau_{pk}} = f_{pk}(d_p), \quad (3)$$

$$d_k = \frac{c\tau_{pk}(c\tau_{pk} - 2d_p)}{2(c\tau_{pk} - d_p(1 + \cos\theta_{pk}))}.$$

Since the light path of interest undergoes two bounces, the total path length $c\tau_{pk}$ will be larger than $2d_p$ or $2d_k$. That is, we can bound d_p using

$$d_p < \frac{c\min_{k\neq p}(\tau_{pk})}{2}. \quad (4)$$

From (4), both the numerator and the denominator of the right-hand side of the expression in (3) is positive, so d_k is guaranteed to be positive. Also, the first derivative of $f_{pk}(d_p)$ defined in (3) is:

$$f'_{pk}(d_p) = \frac{2c^2\tau_{pk}^2(\cos\theta_{pk} - 1)}{[2\tau_{pk} - 2d_p(1 + \cos\theta_{pk})]^2}.$$

Since we only consider two-bounce light path, where $p \neq k$, and $\cos\theta_{pk} \neq 1$, we can conclude that

$$f'_{pk}(d_p) < 0. \quad (5)$$

Thus $f_{pk}(\cdot)$ is a strictly decreasing function.

There are three important consequences to the derivation above. First, given d_p, the value of d_k is uniquely known. Second, f_{pk} is a fractional linear transformation (FLT). Note that the composition of multiple FLTs remain fractional linear. Third, there are infinite pairs of (d_k, d_p) that satisfy the constraint in (3) that relates the two-bounce path length of two individual points. Hence, we cannot solve the depths of two scene points by considering a single two-bounce light path connecting them. Further, it can be shown that the ambiguity inherent in (3) cannot be resolved if we consider a light-path graph whose topology is a tree since the relationship between depths is an FLT. This motivates us to explicitly consider light-path graph topologies that include cycles, with the hope that cycles in the graph will help resolve these shape ambiguities. The following Lemma provides a concrete result along this direction.

Lemma 2. *There are at most two solutions for a connected light-path graph whose topology is a single cycle.*

Proof. Since the light-path graph is connected and has a single cycle, all vertices must be part of the cycle. We assume that the cycle in consideration is given by the following edges: $(q_1, q_2), (q_2, q_3), \ldots, (q_{N-1}, q_N)$, and (q_N, q_1). By applying (3) to the edges in succession, we can obtain the following relationship:

$$d_{q_1} = f_{q_N q_1}(\cdots f_{q_2,q_3}(f_{q_1 q_2}(d_{q_1}))) = T(d_{q_1}). \quad (6)$$

Given that each $f_{q_i q_j}$ is an FLT, the RHS of (6) is also an FLT of the form

$$d_{q_1} = \frac{C_1 + C_2 d_{q_1}}{C_3 + C_4 d_{q_1}}, \quad (7)$$

where C_1, C_2, C_3, and C_4 are dependent on the individual two-bounce path lengths. We can rearrange (7) to obtain a second-order polynomial equation in d_{q_1}, which has two roots and hence, two potential solutions for d_{q_1}. For each solution, we can estimate the depths of all other vertices uniquely via (3). □

The implications of Lemma 2 are promising since it restricts the solution space associated with a cycle. Further, while there are two potential solutions, it is entirely possible that one of them is infeasible or both solutions are exactly the same. We show that this is indeed the case for cycles with odd number of vertices.

Proposition 3. *There is exactly one solution for a light-path graph whose topology is a single cycle with odd number of vertices.*

Proof. We prove the uniqueness of the solution by contradiction. Suppose there are two feasible solutions for d_{q_1} from (7), $d^{(1)}$ and $d^{(2)}$ with $d^{(1)} < d^{(2)}$. That is,

$$T\left(d^{(1)}\right) = d^{(1)}, \quad T\left(d^{(2)}\right) = d^{(2)}. \quad (8)$$

By chain rule, the first derivative of $T(\cdot)$ is

$$T'(d) = f'_{q_N q_1}(\cdots f_{q_2 q_3}(f_{q_1 q_2}(d)))\cdots f'_{q_1 q_2}(d).$$

Since the cycle contains an odd number of edges and $f'(d) < 0$ from (5), $T'(d) < 0$. Hence, $T(d_p)$ is a strictly decreasing function. Recall that we assumed $d^{(1)} < d^{(2)}$, thus,

$$T(d^{(2)}) < T(d^{(1)}) = d_p^{(1)} < d_p^{(2)}. \quad (9)$$

We see a contradiction in (8) and (9). Therefore, $d^{(1)} = d^{(2)}$ and hence, there is only one feasible solution. □

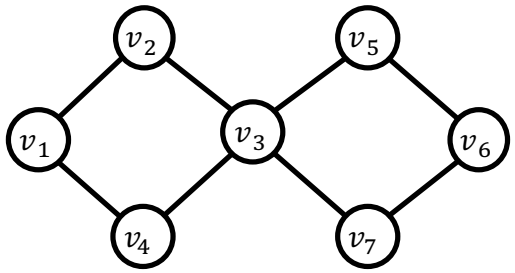

Figure 3. Sample graph to explain Proposition 4. We form one set of estimation from cycle $\{v_1, v_2, v_3, v_4\}$ and the other set from cycle $\{v_3, v_5, v_6, v_7\}$. There will only be one common estimate for d_3, which is the true depth of vertex v_3.

Proposition 4. *For a connected light-path graph consisting of two intersecting cycles, each with even number of vertices, then depth of all vertices can be uniquely recovered.*

Proof. From Lemma 2, we can find the two potential depths for each of the vertices on the first cycle. Similarly, we can obtain two potential depths for each of the vertices for the second cycle. Since the cycles are intersecting, we can intersect the solutions at a vertex that belongs to both cycles and resolve the ambiguity. □

As an example, consider the graph in Figure 3, the depth of all scene patches can be uniquely recovered. We can first find the two solutions of v_3, $\{a_1^3, b_1^3\}$ by considering cycle $\{v_1, v_2, v_3, v_4\}$. Similarly, from cycle $\{v_3, v_5, v_6, v_7\}$, we find $\{a_2^3, b_2^3\}$. One element from each set will be consistent, which is the depth estimation for v_3. After finding d_3, we can propagate the depth estimation to all vertices in the graph.

Proof of Theorem 1. We now have all the components for the proof of Theorem 1. From Lemma 2, the presence of cycle in the light-path graph limits the solution space to two. We further prove in Proposition 3 that for a cycle with odd number of vertices, only one of the solution is feasible. In Proposition 4, when considering multiple cycles, even though each cycle gives two solutions, we can find the unique solution by intersecting the solution across cycles. Therefore, Theorem 1 provides sufficient conditions for shape recovery from light-path graph topology.

Implications of Theorem 1. The sufficient conditions in Theorem 1 only depends on the the topology of the light-path graph but not the actual light-path lengths. The topology of the light-path graph depends only on the existence of two-bounce light paths between surface patches, which is completely determined by the geometry of the scene. Hence, Theorem 1 infers uniqueness of shape estimates for a broad class of scenes.

5. Two-bounce shape estimation

In this section, we introduce two algorithms to recover shape from two-bounce light paths. Our goal is in leveraging information over the entire light-path graph to robustly estimate the depth of the scene in the presence of noise corrupting the measured path lengths.

Algorithm 1 — Multi-cycle clustering algorithm. One of the main results from Section 4 is that there are at most two depth estimates associated with a cycle in the light-path graph. With presence of noise, the depth estimation from different sub-graphs will be different. Therefore, we proposed a clustering method to find a solution that is most consistent across different cycles.

We randomly select Q cycles from the light-path graph. By Lemma 2, each cycle provides us with two candidate depths for each of its vertices. We denote the candidate depths for vertex k from cycle q as a_q^k and b_q^k. Repeating this process for each of the selected cycles, we have $2Q$ estimates for each vertex. In the absence of noise, Q out of these $2Q$ solutions (one out of 2 solutions from each cycle containing the vertex) would be identical and that would indicate the correct solution for the depth of that vertex. In the presence of noise, the $2Q$ solutions are all potentially different and hence, we need an alternative technique to identify the correct solution.

Let $\mathbf{x} = [x_1, \ldots, x_Q]$ be a Q-dimensional indicator vector, where $x_q = 1$ implies that, in cycle q, the candidates in $\{a_q^k\}$ are selected as the estimate and $x_q = -1$ implies that the candidates $\{b_q^k\}$ are selected. Therefore, the depth estimate for vertex k provided by cycle q is

$$d_q^k = \frac{1 + x_q}{2} a_q^k + \frac{1 - x_q}{2} b_q^k.$$

We now minimize the variance of the depth estimates:

$$\min_{\mathbf{x} \in \{-1,1\}^Q} \sum_{k=1}^{N} \left[\frac{1}{Q} \sum_{q=1}^{Q} (d_q^k)^2 - \left(\frac{1}{Q} \sum_{q=1}^{Q} d_q^k \right)^2 \right]$$

This is a combinatorial problem due to the feasible set being $\mathbf{x} \in \{-1, 1\}^Q$. If we relax this constraint to let $\mathbf{x}$ take real values, and observing that the objective is quadratic in $\mathbf{x}$, we can obtain a closed form solution for $\mathbf{x}$. By thresholding this solution at zero, we get our estimator of the indicator vector $\widehat{\mathbf{x}}$. We use the mean of candidate depths at each vertex selected by $\widehat{\mathbf{x}}$ as our estimate.

The clustering of estimates over a select few cycles is sub-optimal since it does not simultaneously leverage information in all the two-bounce light paths. This is addressed by the second algorithm which exploits all the observed two-bounce path lengths in a holistic framework to estimate the scene depth.

Algorithm 2 — Rank-constrained optimization. Recall that the relationship between length of a two-bounce light path to the depth of scene points is bilinear (see (2)). A common approach to handle bilinear systems is to over-parameterize the unknowns using a rank-1 matrix $\mathbf{D}$, as described below. Specifically, defining $\mathbf{d} = [1, d_1, \cdots, d_N]$, we formulate the depth estimation problem in terms of the rank-1 matrix $\mathbf{D} = \mathbf{d}^\mathbf{T}\mathbf{d}$. Denoting $D_{i,j}$ as the (i,j)-th entry of $\mathbf{D}$, the path-length constraint associated with scene points p and k in (2) can be expressed as

$$c\tau_{pk}(D_{1,p+1} + D_{1,k+1}) - D_{p+1,k+1}(1 + \cos\theta_{pk}) = \frac{c^2\tau_{pk}^2}{2}.$$

The linearity of the relationship allows us to formulate the depth estimation problem as:

$$\min_{\mathbf{D}\succeq\mathbf{0}} \|\mathcal{A}(\mathbf{D}) - \mathbf{b}\|^2 \quad \text{s.t.} \quad \text{rank}(\mathbf{D}) = 1, D_{1,1} = 1, \quad (10)$$

where $\mathcal{A}$ and $\mathbf{b}$ can be set from (2). We can also incorporate single-bounce depth estimates into the optimization framework above by associating them with the diagonal elements of $\mathbf{D}$. We use a projected gradient descent method to optimize (10):

$$\mathbf{D}^{k+1} = P_\Omega(\mathbf{D}^k - \gamma\mathcal{A}^T(\mathcal{A}(\mathbf{D}^k) - \mathbf{b})),$$

where γ is the step size and $P_\Omega(.)$ is a projection operator to the space of rank-1 matrices. Since (10) is non-convex, the solution is greatly affected by the initialization. We initialize $\mathbf{D}$ by the results from shape from the multi-cycle clustering method (Algorithm 1).

Simulations. In Figure 4, we show the recovery results on a 1D slice of an object for varying amount of additive Gaussian noise on the measured path lengths. We compare single-bounce estimates, the multi-cycle clustering and the rank-constrained method initialized with the multi-cycle clustering estimates. We use $Q = 100$ cycles for the multi-cycle method. Both two-bounce methods outperform single-bounce techniques by a large margin, in part, due to the availability of a larger number of light paths.

In Figure 5, we show the reconstruction signal-to-noise ratio (SNR) when the observed total path-length is perturbed by different levels of noise for four different shapes. For the U-shape object, the depth range is $49\sim55$cm. The ToF for two-bounce paths is around 3ns. For SNRs of 40, 30, 20 dB, the standard deviations of the noise added to the path length are 42, 133, 419 ps, respectively. Note that the temporal resolution of streak cameras are around $1\sim2$ ps while those of PMD sensors are around 100 ps, making these practical noise levels. From our numerical experiments, we observed that light-path graph topology does affect the quality of the reconstruction under noise. For a fully connected light-path graph (U-shape objects), the results are significantly better. However, for a bipartite graph (V-shape object) which only has even cycles, the reconstruction quality is degraded when using multi-cycle clustering; surprisingly, this happens only at higher measurement SNRs. This indicates that the shape estimation algorithms purely from even cycles can be unstable; however, an analysis of stability under noise is beyond the scope of this paper.

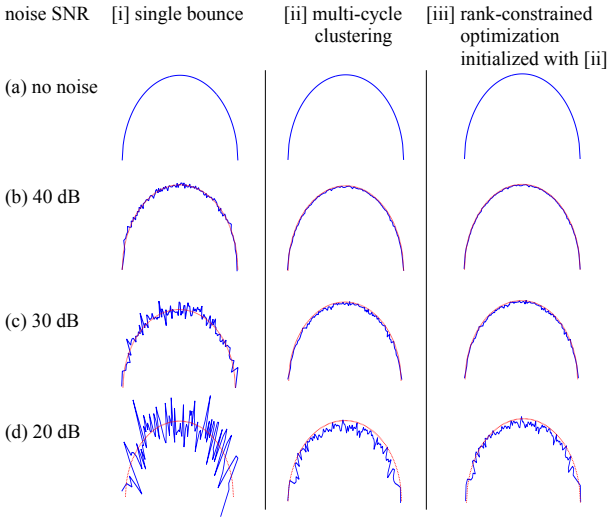

Figure 4. **2D shape recovery results.** We perturbed the input using different level of noise. The blue line shows our reconstructed surface and the red dotted line is the ground truth.

6. Reflectance from two-bounce light paths

In this section, we describe how information from light transport associated with the two-bounce light paths can be used to recover the BRDF of the surface material. We assume knowledge of the shape of the scene, in terms of its depth and surface normals; this could be obtained by using the techniques outlined in Section 5.

The main advantage of using two-bounce light paths to recover reflectance is the coverage of different incident and outgoing ray directions. This is because, the two-bounce light paths increase the diversity of angular coverage both for incident and exiting light paths. Note that this increase in angular diversity depends upon the scene geometry. In Figure 2, we show the coverage of bi-variate BRDF representation [18] for two-bounce light paths. In contrast to the single-bounce light paths that only include $\theta_d = 0$ instances, two-bounce light paths includes greater coverage of the half/diff angle. Also, angle that is close to mirror reflection direction $\theta_h = 0$ is also covered.

Relationship between two-bounce light transport and BRDF. Estimating the BRDF of the scene requires the relative light transport strength for each of the two-bounce light paths, i.e., given a two-bounce light path passing through vertices p and k, we need to know the fraction of light directed at the vertex p that eventually returns to the

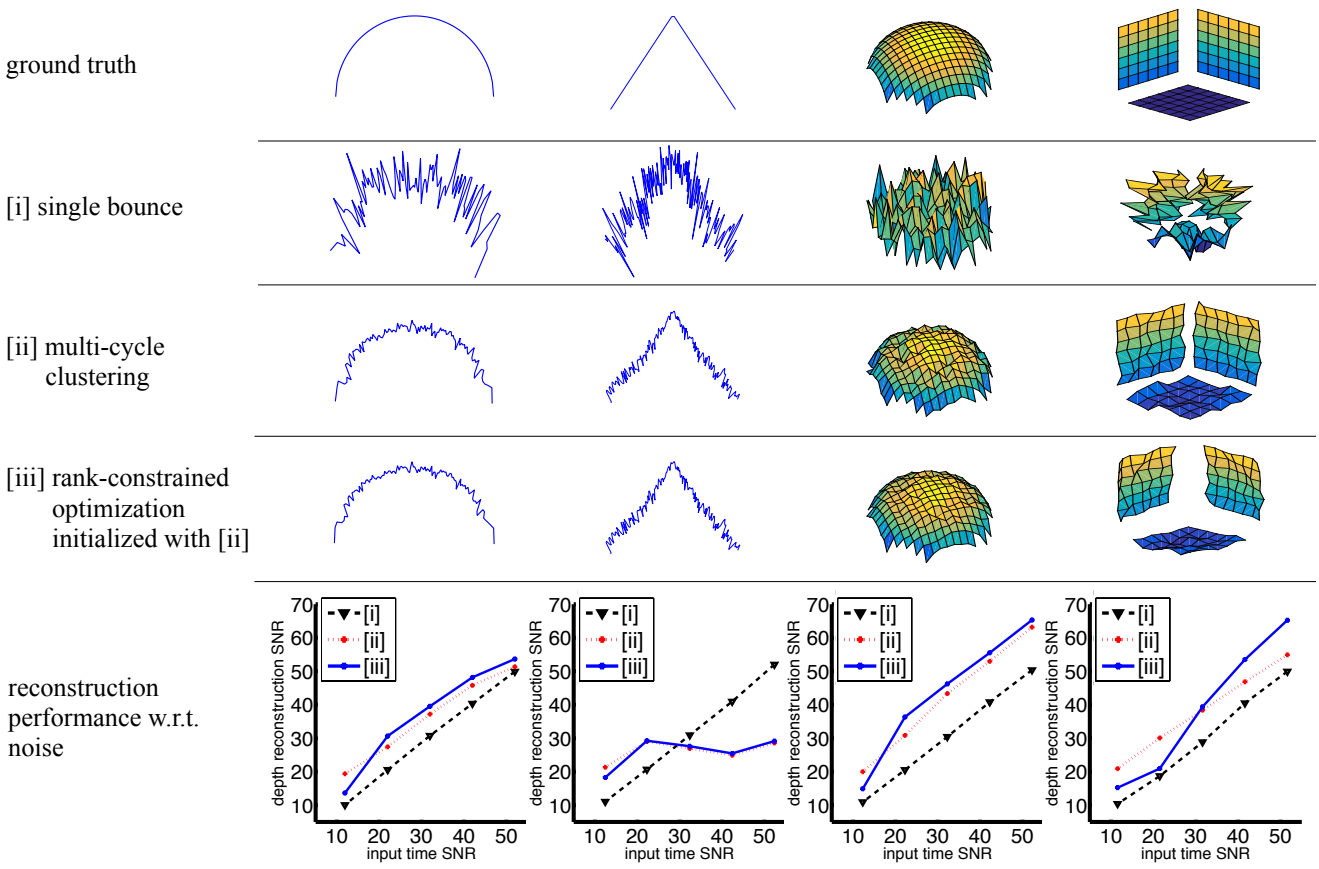

(a) U-shape in 2D (b) V-shape in 2D (c) U-shape in 3D (d) Corner in 3D

Figure 5. **Shape recovery results.** We compare the following three methods: [i] Shape from single-bounce. [ii] Multi-cycle clustering. [iii] Rank-constrained optimization initialized with results from [ii]. Notice that [ii] and [iii] are using only two-bounce light paths information. In second to fourth rows, we show the shape estimation results of 20dB added noise.

camera through vertex k. The two-bounce light transport between two scene points can be given as product of the two BRDF entries and a form factor term that is set by the geometry of the scene [11]. The radiance of the two-bounce light path involving two scene patches can be written as:

$$L_{pk} = L_p B(p,k) V(p,k) G(p,k),$$

where L_p is the intensity of incident light at point p, $B(p,k)$ represents the multiplicative effect of the two BRDF terms, $V(p,k)$ represents a binary visibility term, and $G(p,k)$ represents the geometric scaling term and these can be expanded out as,

$$\begin{aligned} B(p,k) &= \boldsymbol{\rho}_p(-\mathbf{i}_p, \mathbf{n}_p, \mathbf{i}_{p\to k}) \max(0, -\mathbf{i}_p^T \mathbf{n}_p) \\ &\quad \boldsymbol{\rho}_k(\mathbf{i}_{k\to p}, \mathbf{n}_k, -\mathbf{i}_k) \max(0, -\mathbf{i}_{k\to p}^T \mathbf{n}_k) \\ G(p,k) &= \frac{(\mathbf{n}_p^T \mathbf{i}_{p\to k})(\mathbf{n}_k^T \mathbf{i}_{k\to p})}{\|\mathbf{v}_p - \mathbf{v}_k\|^2}, \end{aligned}$$

where $\boldsymbol{\rho}_p$ and $\boldsymbol{\rho}_k$ are the BRDFs at $\mathbf{v}_p$ and $\mathbf{v}_k$, respectively. Further, if the two-bounce path is visible, then

$$\begin{aligned} I_{pk} &= \frac{L_{pk}}{L_p(-\mathbf{i}_p^T \mathbf{n}_p)(-\mathbf{i}_{k\to p}^T \mathbf{n}_k)} \frac{1}{G(p,k)} \\ &= \boldsymbol{\rho}_p(-\mathbf{i}_p, \mathbf{n}_p, \mathbf{i}_{p\to k}) \boldsymbol{\rho}_k(\mathbf{i}_{k\to p}, \mathbf{n}_k, -\mathbf{i}_k) \\ &= \mathcal{M}_{-\mathbf{i}_p, \mathbf{n}_p, \mathbf{i}_{p\to k}} \boldsymbol{\rho}_p \mathcal{M}_{-\mathbf{i}_k, \mathbf{n}_k, \mathbf{i}_{k\to p}} \boldsymbol{\rho}_k, \end{aligned}$$

where $\mathcal{M}_{\mathbf{v},\mathbf{n},\mathbf{l}}$ is a linear operator that probes BRDF entry corresponding to the incident direction, $\mathbf{l}$, normal, $\mathbf{n}$, and viewing direction, $\mathbf{v}$.

Parametric reflectance model. Similar to [13], we use the Ashikhmin-Shirley model [1], which is composed of Lambertian diffuse term and specular lobe term, to parameterize the scene BRDF.

$$\rho(\mathbf{v}, \mathbf{n}, \mathbf{l}) = \rho_d(\mathbf{n}, \mathbf{l}) + \rho_s(\mathbf{v}, \mathbf{n}, \mathbf{l})$$

$$\rho_d(\mathbf{n}, \mathbf{l}) = \frac{k_d}{\pi}$$

$$\rho_s(\mathbf{v},\mathbf{n},\mathbf{l}) = k_s \frac{k_n+1}{8\pi} \frac{\langle \mathbf{n},\mathbf{h}\rangle^{k_n}}{\langle \mathbf{v},\mathbf{h}\rangle \max(\langle \mathbf{n},\mathbf{l}\rangle, \langle \mathbf{n},\mathbf{v}\rangle)} F(F_0, \langle \mathbf{v},\mathbf{h}\rangle)$$

Here, $\mathbf{h}$ is the half vector. The Fresnel effect term [20] is used to improve the accuracy of the model.

$$F(F_0, \langle \mathbf{v},\mathbf{h}\rangle) = F_0 + (1-F_0)(1-\langle \mathbf{v},\mathbf{h}\rangle)^5,$$

Thus, the parameters of this model are k_d, k_s, k_n and F_0. Estimating the material BRDF now amounts to estimating the 4 parameters of the Ashikhmin-Shirley model for each color channel of the scene patch.

The image formation for a two-bounce light path can be rewritten as:

$$I_{pk} = [\rho_d(\mathbf{n}_p, -\mathbf{i}_p) + \rho_s(\mathbf{i}_{p\to k}, \mathbf{n}_p, -\mathbf{i}_p)] \; [\rho_d(\mathbf{n}_k, \mathbf{i}_{k\to p}) + \rho_s(-\mathbf{i}_k, \mathbf{n}_k, \mathbf{i}_{k\to p})], \quad (11)$$

where $\mathbf{h}_{pk}$ and $\mathbf{h}_{kp}$ are the half vector. The unknowns for each surface patch are k_d, k_s, k_n and F_0. We make two more assumptions that further simplify the model. First, similar to [16], we assume the three color channels of the material share the same specular lobe and Fresnel effect coefficient. This reduces the number of parameters to 8 per scene patch. In addition, if the object of interest only contains one material, we only need to recover 8 parameters to estimate the reflectance for that material.

Simulations. Collecting all two-bounce light path intensity, we use the lsqcurvefit function from Matlab to find the 8 BRDF parameters of the non-linear least squares problem defined in (11). In Fig. 6(b-e), we show the reconstructed BRDF rendered as a sphere and illuminated using the Grace Cathedral environment map.[2] We model two sources of noise, read noise and photon noise, thus the added noise to each observed intensity is $\nu = \nu_{\text{read}} + \nu_{\text{photon}}$. The read noise is $\nu_{\text{read}} = 0.001 \max(I)\xi$ and the photon noise is $\nu_{\text{photon}} = \eta I_{pk}\xi$, where ξ is a random number from the normal distribution and η is a noise parameter. In the paper, we set $\eta = 0.1$. In Fig. 6(f), we show the estimation error over all the materials in the MERL BRDF database [12]. Given the ground truth BRDF ρ and the estimate $\widehat{\rho}$, the BRDF error is computed as

$$E = \sqrt{\frac{\sum_i w[(\widehat{\rho}_i - \rho_i)\cos\theta_i]^2}{\sum_i w}},$$

where θ_i is the elevation angle of the incident ray. In our experiment, we set $w = 1$.

From Fig. 6(b-e), we observe that when using two-bounce light paths, the mirror direction is covered, thus the parameter controlling the specular lobe will be updated. While when only single-bounce light paths are used, the specular lobe shape is not updated at all, only the diffuse term will be updated to fit the observed intensity. This is especially clear when viewing the polar graph of BRDF. From Fig. 6(f), we observe that the performance of our proposed work is roughly the same as the parameter fitting using the whole BRDF [16].

[2] http://www.pauldebevec.com/Probes/

7. Conclusions

In this paper, we have developed a framework for the study of shape and reflectance information encoded in the two-bounce light paths. While the focus of our results have been reliant exclusively on the two-bounce light paths, we hope that hybrid methods for 3D shape and reflectance estimation that use the results in this paper in conjunction with other techniques that rely on single-bounce light paths will emerge and be able to overcome the limitations of traditional techniques to complex reflectance and geometries.

Acknowledgments

We thank Zhuo Hui for discussions on the BRDF estimation. CYT was supported, in part, by the Bertucci Graduate Fellowship. CYT and ACS gratefully acknowledge support by a gift from Adobe Research. ACS and AV were partially supported by the DARPA REVEAL program.

References

[1] M. Ashikhmin and P. Shirley. An anisotropic phong brdf model. *J. Graphics Tools*, 5(2):25–32, 2000.

[2] I. Gkioulekas, A. Levin, F. Durand, and T. Zickler. Micron-scale light transport decomposition using interferometry. *ACM Trans. Graphics*, 34(4):37, 2015.

[3] J. P. Godbaz, A. A. Dorrington, and M. J. Cree. Understanding and ameliorating mixed pixels and multipath interference in AMCW lidar. In *TOF Range-Imaging Cameras*, pages 91–116. Springer, 2013.

[4] M. Gupta, A. Agrawal, A. Veeraraghavan, and S. G. Narasimhan. A practical approach to 3d scanning in the presence of interreflections, subsurface scattering and defocus. *Intl. J. Computer Vision*, 102(1-3):33–55, 2013.

[5] M. Gupta, S. K. Nayar, M. B. Hullin, and J. Martin. Phasor imaging: A generalization of correlation-based time-of-flight imaging. *ACM Trans. Graphics*, 34(5):156, 2015.

[6] R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, ISBN: 0521540518, second edition, 2004.

[7] F. Heide, M. B. Hullin, J. Gregson, and W. Heidrich. Low-budget transient imaging using photonic mixer devices. *ACM Trans. Graphics*, 32(4):45, 2013.

[8] F. Heide, L. Xiao, W. Heidrich, and M. Hullin. Diffuse mirrors: 3d reconstruction from diffuse indirect illumination using inexpensive time-of-flight sensors. In *CVPR*, 2014.

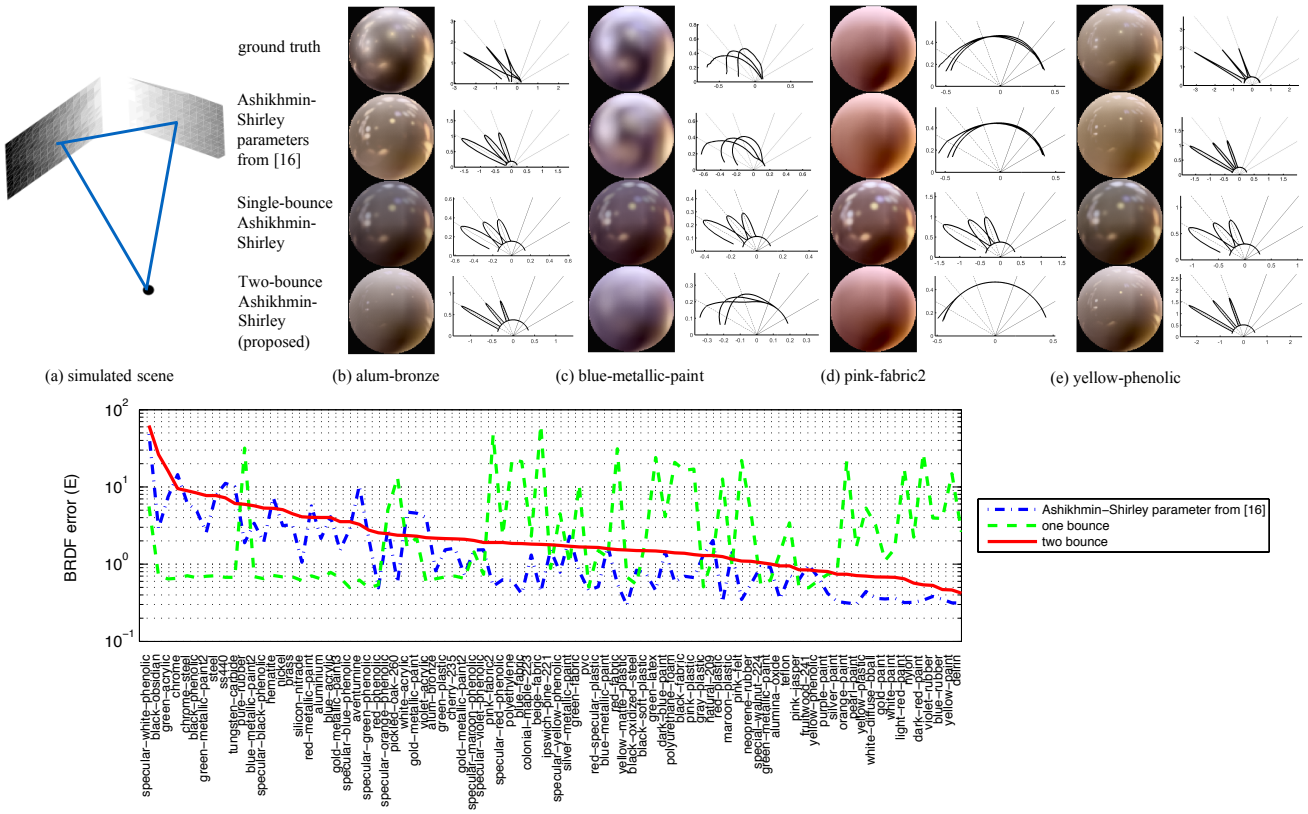

(f) BRDF estimation as a function of material BRDF

Figure 6. **BRDF estimation results.** (a) We present reflectance estimation results for a 3D V-groove. (b-e) We show the BRDF estimation when setting $\eta = 0.1$ by rendering the material as a sphere and illuminate it with the Grace Cathedral environment map. Also, we show the polar graph of incident directions at $20°, 40°, 60°$ with cubic root applied. (b-e) First row: ground truth rendering using the full 4374000-dimensional BRDF measurement. Second row: Ashikhmin-Shirley parameters from [16]. Third row: Ashikhmin-Shirley parameters estimated using single-bounce light paths. Last row: our proposed parametric reflectance estimation using two-bounce light paths. (f) BRDF estimation error over materials in the MERL BRDF database.

[9] A. Kirmani, T. Hutchison, J. Davis, and R. Raskar. Looking around the corner using transient imaging. In *ICCV*, 2009.

[10] K. N. Kutulakos and E. Steger. A theory of refractive and specular 3d shape by light-path triangulation. *Intl. J. Computer Vision*, 76(1):13–29, 2008.

[11] S. Liu, T.-T. Ng, and Y. Matsushita. Shape from second-bounce of light transport. In *ECCV*, 2010.

[12] W. Matusik, H. Pfister, M. Brand, and L. McMillan. A data-driven reflectance model. *ACM Trans. Graphics*, 22(3):759–769, 2003.

[13] N. Naik, S. Zhao, A. Velten, R. Raskar, and K. Bala. Single view reflectance capture using multiplexed scattering and time-of-flight imaging. *ACM Trans. Graphics*, 30(6):171, 2011.

[14] S. K. Nayar, K. Ikeuchi, and T. Kanade. Shape from interreflections. *Intl. J. Computer Vision*, 6(3):173–195, 1991.

[15] S. K. Nayar, G. Krishnan, M. D. Grossberg, and R. Raskar. Fast separation of direct and global components of a scene using high frequency illumination. *ACM Trans. Graphics*, 25(3):935–944, 2006.

[16] A. Ngan, F. Durand, and W. Matusik. Experimental analysis of brdf models. *Eurographics Symp. Rendering*, 2005.

[17] M. O'Toole, F. Heide, L. Xiao, M. B. Hullin, W. Heidrich, and K. N. Kutulakos. Temporal frequency probing for 5d transient analysis of global light transport. *ACM Trans. Graphics*, 33(4):87, 2014.

[18] F. Romeiro, Y. Vasilyev, and T. Zickler. Passive reflectometry. In *ECCV*, 2008.

[19] D. Scharstein and R. Szeliski. High-accuracy stereo depth maps using structured light. In *CVPR*, 2003.

[20] C. Schlick. An inexpensive brdf model for physically-based rendering. *Computer graphics forum*, 13(3):233–246, 1994.

[21] A. Velten, T. Willwacher, O. Gupta, A. Veeraraghavan, M. G. Bawendi, and R. Raskar. Recovering three-dimensional shape around a corner using ultrafast time-of-flight imaging. *Nature Comm.*, 3:745, 2012.

[22] A. Velten, D. Wu, A. Jarabo, B. Masia, C. Barsi, C. Joshi, E. Lawson, M. Bawendi, D. Gutierrez, and R. Raskar. Femto-photography: Capturing and visualizing the propagation of light. *ACM Trans. on Graphics*, 32(4):44, 2013.