



# Detection and mapping of specular surfaces using multibounce LiDAR returns

CONNOR HENLEY,<sup>1,2,\*</sup> SIDDHARTH SOMASUNDARAM,<sup>1</sup>  
JOSEPH HOLLMANN,<sup>2</sup> AND RAMESH RASKAR<sup>1</sup>

<sup>1</sup>MIT Media Lab, Massachusetts Institute of Technology, 75 Amherst St., Cambridge, MA 02139, USA

<sup>2</sup>The Charles Stark Draper Laboratory, 555 Technology Sq., Cambridge, MA 02139, USA

\*co24401@mit.edu

**Abstract:** We propose methods that use specular, multibounce LiDAR returns to detect and map specular surfaces that might be invisible to conventional LiDAR systems that rely on direct, single-scatter returns. We derive expressions that relate the time- and angle-of-arrival of these multibounce returns to scattering points on the specular surface, and then use these expressions to formulate techniques for retrieving specular surface geometry when the scene is scanned by a single beam or illuminated with a multi-beam flash. We also consider the special case of transparent specular surfaces, for which surface reflections can be mixed together with light that scatters off of objects lying behind the surface.

© 2023 Optica Publishing Group under the terms of the [Optica Open Access Publishing Agreement](#)

## 1. Introduction

Although LiDAR is widely used for mapping the 3D geometry of surfaces, the technology has historically been challenged by *specular*, or mirror-like, surfaces that typically scatter very little light directly back to the receiver. This inability to detect and localize specular surfaces can result in the failure to detect navigational obstacles like mirrors and windows, or hazards such as wet or icy patches on the ground. It can also result in incomplete scans of cityscapes or man-made interior environments in which glass and metal surfaces are relatively common, and in the complete inability to digitize artifacts that are made of glass or that present a polished metal or chrome finish.

The range to a surface is typically computed using the LiDAR range equation, which is only valid for single-scatter time-of-flight measurements. Thus, when the direct, single-scatter return from a specular surface is too faint to detect, it becomes impossible to compute the range to that surface via conventional means. Nevertheless, the presence of specular surfaces is often revealed by intense *multibounce* returns. For instance, when one directly illuminates a *diffusely* reflecting surface, nearby specular surfaces may produce mirror images of the true laser spot (also referred to as *highlights*) that appear just as bright as the original. Alternatively, when a specular surface is illuminated directly, the beam is often deflected such that it lands on a diffusely reflecting surface nearby—thus producing a clearly visible spot. These multibounce signals are easy to observe when one waves a laser pointer around any room containing mirrors and windows, and yet they have been largely overlooked by the LiDAR sensing community as useful sources of information.

In this work we demonstrate that multibounce returns are both an important cue that reveals the presence of specular surfaces, as well as an information source that can be used to estimate a specular reflector's shape. We review the geometry of multibounce returns that are encountered in scenes that contain both diffuse and specular reflectors. Using our knowledge of this geometry, we propose criteria for unambiguously detecting the presence of specular multibounce signals, as well as a set of equations that relates the time- and angle-of-arrival of multibounce impulses to the position and orientation of points on the specular surface.

We apply our findings in a series of experiments in which we use multibounce LiDAR measurements to acquire the 3D shape of various specular objects. These objects include planar reflectors such as mirrors and windows, as well as polished metal and thin glass objects with freeform shapes. For the special case of transparent surfaces, we propose criteria that can be used to distinguish between multibounce reflections off of the surface, and single-scatter returns from objects that lie behind the surface. We assess the accuracy of our specular surface retrieval method when it is applied to a large mirror surface, and also highlight some of the most important conditions that cause our method to fail.

Our methods are implemented most efficiently (in terms of acquisition time) with a wide field of view (FOV) receiver such as a single-photon avalanche diode (SPAD) array, however they can be implemented using any LiDAR system for which the transmit and receive axes can be steered independently. For most of our experiments we use a pencil-beam illumination source that must be scanned to acquire a point cloud of the full scene. However we also propose an algorithm that can be used to map the surface of planar reflectors when the scene is illuminated by many beams simultaneously. Multi-beam illumination enables faster scene acquisition, and is employed by several commercial LiDAR scanners [1].

## 2. Related work

### 2.1. Detecting mirrors and windows using LiDAR

Specular surfaces typically reflect most light away from the LiDAR receiver, resulting in very faint direct returns. When specular reflections *are* directed towards the receiver, the signal can be so intense that it saturates one's detectors (this event is sometimes called "glare"). The presence of mirrors in a scene can also produce false detections of the *mirror images* of real objects, which typically appear behind the mirror.

To overcome these challenges, Diosi and Kleeman [2] use an ultrasonic scanner to detect specular surfaces that LiDAR scanners cannot see. Yang et al. [3] infer that framed gaps in detected surfaces likely contain mirrors or windows. After classifying mirrors in this way, they identify mirror image detections and reflect them across the mirror plane to the position of the true object. Foster et al. use sparse glare events to detect specular surfaces, but accumulate glare information over time in an occupancy grid as their laser scanner moves through a space [4]. Like Yang et al., Tibebu et al. [5] search for frames that are indicative of windows. However, to avoid detecting windowless holes, they use measurable pulse broadening caused by transmission through glass as a second heuristic to make their window detector more robust.

Unlike this previous work, our method does not rely on contextual cues like frame or mirror image detection that might be produced by non-specular scenes, we do not rely on rare glare events, and we do not require additional sensing modalities like ultrasound. The work that comes closest to ours was a speculative report by Raskar and Davis [6], who suggested that the time-of-flight of two-bounce returns could be used to localize specular surfaces, but did not demonstrate their method, and did not (as we do) consider the scenario where the specular surface is illuminated directly. Henley et al. [7] and Tsai et al. [8] estimate scene shape from two-bounce time-of-flight measurements, but neither consider the special conditions imposed by specular reflectors.

### 2.2. Specular surface estimation using cameras

There is an extensive body of literature that investigates methods for estimating the geometry of specular surfaces using conventional cameras that cannot capture time-of-flight information. A relatively recent review of this research was provided by Ihrke et al. [9]. A camera that observes a specular surface will see a distorted (if the mirror is curved) reflection of the scene that surrounds the surface. In most camera-based methods for specular geometry estimation, features in the

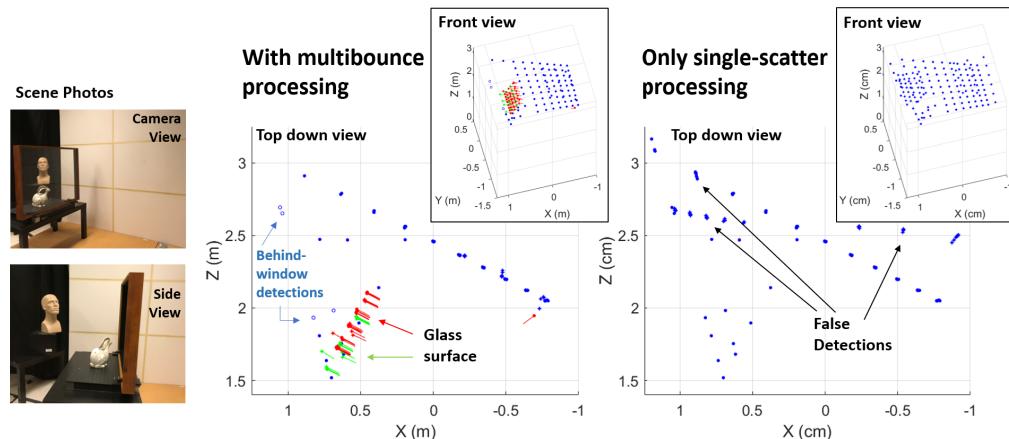
distorted, reflected image are matched to features in the true scene. If the positions of the camera and true feature are known, then the *depth* and *surface normal* of the surface point that reflects the distorted feature can be determined up to a one-dimensional ambiguity. This ambiguity can be resolved in a variety of ways. In *shape from specularity* methods, the reflected scene features are point sources that produce specular highlights in the camera image [10,11]. In *shape from distortion* methods, a camera observes how specular reflection distorts the reflected image of a reference pattern [12–14]. In *specular flow* methods a camera observes how the distorted reflection of an *uncalibrated* scene appears to move across the surface of a specular object as the camera, object, or scene are moved [15].

Because we observe the specular reflection of laser spots, our method can be interpreted as a *shape from specularity* technique that uses time-of-flight constraints to resolve the depth-normal ambiguity. In future work, it would be interesting to explore the combination of time-of-flight constraints with other imaging strategies such as *shape from distortion* or *specular flow*.

### 3. Methods

#### 3.1. Geometry of specular multibounce signals

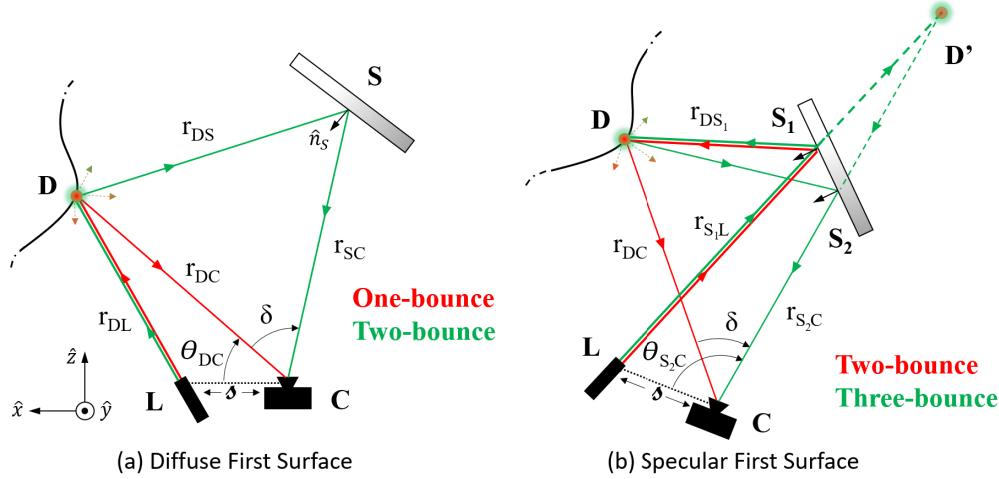
Consider the scenario illustrated in Fig. 2, where a mirror has been placed in an environment that is otherwise composed of diffusely reflecting surfaces. The mirror is a perfect specular reflector, which means that it will reflect all light from an incident beam in a single direction that is determined by the light's direction of incidence and the mirror's surface normal orientation. Our LiDAR system consists of a transmitter at position  $L$  that emits a focused, pulsed beam in a single direction. A receiver array is placed at position  $C$ , which may be displaced from  $L$  by a baseline distance  $s$ .



**Fig. 1.** We acquire the shape of a glass window (green and red points) and the scene surrounding it (blue points), which includes objects behind the window (white points, photographs on left). When single-scatter paths are assumed (right), the window surface cannot be retrieved and numerous false points are generated.

In this scenario the transmitter can illuminate either a diffuse reflector or the mirror directly. These two cases introduce distinct multibounce geometries that must be treated separately. In either case, our objective is to compute the position of all directly and indirectly illuminated points in the scene using the time-of-flight and angle-of-arrival of all observable direct and multibounce returns.

**First surface is a diffuse reflector.** In the first case, shown in Fig. 2(a), the transmitter directly illuminates a point  $D$  on the diffusely reflecting surface. The receiver subsequently observes two



**Fig. 2.** (a) A diffusely reflecting surface is illuminated directly at  $D$ . The receiver observes one-bounce returns incident from  $D$  and two-bounce returns incident from point  $S$  on the mirror surface. (b) The specular surface is illuminated directly (at  $S_1$ ) and the single-bounce return is too faint to observe. The receiver observes two-bounce returns incident from  $D$  and three-bounce returns incident from  $S_2$ .

laser spots: the *true* spot, which (correctly) appears to be located at  $D$ , and a *mirror image* of that spot. Light from the true spot has propagated along the single-bounce path  $LDC$ , whereas light that appears to originate from the spot's mirror image has propagated along the two-bounce path  $LDSC$ . Here  $S$  is the point at which light reflects off of the mirror. We use a method that was originally introduced in [7] to estimate the positions of  $D$  and  $S$  from the times- and angles-of-arrival of the one- and two-bounce returns. Angular positions are described by a clockwise rotation by  $\theta$  about  $\hat{y}$  followed by a clockwise rotation by  $\phi$  about  $\hat{x}$ .

Light incident from the direction of the true spot  $D$  will arrive at the angle  $(\theta_{DC}, \phi_{DC})$  at time  $t_1 = \frac{1}{c}(r_{DL} + r_{DC})$ , where  $r_{DL}$  and  $r_{DC}$  are the distances to  $D$  from  $L$  and  $C$ , respectively. The distance to  $r_{DC}$  can be computed using the following bi-static range equation:

$$r_{DC} = \frac{1}{2} \frac{c^2 t_1^2 - s^2}{ct_1 - scos(\theta_{DC})}. \quad (1)$$

Light incident from the direction of  $S$  will arrive at the angle  $(\theta_{SC}, \phi_{SC})$  at time  $t_2 = \frac{1}{c}(r_{DL} + r_{DS} + r_{SC})$ . Here  $r_{DS}$  is the distance from  $S$  to  $D$  and  $r_{SC}$  is the distance from  $C$  to  $S$ . From simple arithmetic we see that  $r_{DS} = c(t_2 - t_1) + r_{DC} - r_{SC}$ . We substitute this expression into the law of cosines for the triangle  $DCS$  to obtain

$$r_{SC} = \frac{c}{2} \cdot \frac{\Delta t_{12} [\Delta t_{12} + \frac{2r_{DC}}{c}]}{\Delta t_{12} + (1 - \cos \delta) \frac{r_{DC}}{c}}. \quad (2)$$

Here  $\Delta t_{12} = t_2 - t_1$  and  $\delta$  is the apparent angular separation of the true spot and its mirror image, with  $\cos \delta = \cos \theta_{DC} \cos \theta_{SC} + \sin \theta_{DC} \sin \theta_{SC} \cos(\phi_{DC} - \phi_{SC})$ .

We have now completely determined the positions of  $S$  and  $D$ . We can additionally compute the surface normal at  $S$  which, by the law of reflection, must be the normalized bisector of the angle formed by line segments  $DS$  and  $CS$ .

**First surface is a specular reflector.** In the second case, shown in Fig. 2(b), the mirror is illuminated directly at  $S_1$ . Because the mirror is a perfect specular reflector, no light that scatters

at  $S_1$  will travel directly back to the receiver, and so we do not observe a one-bounce return. Instead, all light in the beam is deflected such that it illuminates a spot  $D$  on a diffusely reflecting surface. This laser spot is visible to the receiver. Once again, however, the receiver also sees a mirror image of the true laser spot. The mirror image of  $D$  appears to lie behind (assuming that the mirror is not strongly concave) the mirror at point  $D'$ . This time, light that arrives from the true spot has traveled along the two-bounce path  $LS_1DC$ , and light that arrives from the spot's mirror image follows the *three*-bounce path  $LS_1DS_2C$ .

Here we show that it is possible to compute the positions of  $D$ ,  $S_1$  and  $S_2$  if  $S_1$  and  $S_2$  lie on the same plane (or, more precisely, if there is a single plane that is tangent to the mirror surface at both  $S_1$  and  $S_2$ ). This condition is automatically satisfied if the specular surface is itself a plane, or if the baseline  $s = 0$  (in which case  $S_1 = S_2$ ).

Our derivation hinges upon the observation that the light that has actually followed the three-bounce trajectory  $LS_1DS_2C$  will *appear*, from the perspective of the receiver, to have followed the one-bounce trajectory  $LD'C$ , which includes a single scattering event at  $D'$ . We can thus compute the apparent range to  $D'$  from the three-bounce travel time  $t_3 = \frac{1}{c}(r_{S_1L} + r_{DS_1} + r_{DS_2} + r_{S_2C}) = \frac{1}{c}(r_{D'L} + r_{D'C})$ . Using the range equation from Eq. (1), we obtain

$$r_{D'C} = r_{DS_2} + r_{S_2C} = \frac{1}{2} \frac{c^2 t_3^2 - s^2}{ct_3 - s \cos(\theta_{S_2C})}. \quad (3)$$

The range to the illusory point  $D'$  is useful because it can be directly related to the range of the true point  $D$  using the following expression:

$$r_{DC} = r_{D'C} - c(t_3 - t_2), \quad (4)$$

where  $t_2 = \frac{1}{c}(r_{S_1L} + r_{DS_1} + r_{DC})$  is the two-bounce travel time. Once we've obtained  $r_{DC}$  we compute the range to  $S_2$  by substituting  $r_{DS_2} = c(t_3 - t_2) + r_{DC} - r_{S_2C}$  into the law of cosines for the triangle  $DCS_2$  to obtain

$$r_{S_2C} = \frac{c}{2} \cdot \frac{\Delta t_{23} [\Delta t_{23} + \frac{2r_{DC}}{c}]}{\Delta t_{23} + (1 - \cos \delta) \frac{r_{DC}}{c}}, \quad (5)$$

where  $\Delta t_{23} = t_3 - t_2$  and  $\delta$  once again refers to the apparent angular separation of the true spot and its mirror image. We note that if there are multiple specular surfaces in the scene then there may be multiple mirror image spots that are visible to the receiver. Only light from one of those mirror images can be used to compute  $r_{DC}$  via Eq. (4). This will be the mirror image that appears to lie along the transmitted beam (in the direction of  $S_1L$ ). However, the range to the reflection points on other specular surfaces can be computed once  $r_{DC}$  is known. This is accomplished by plugging the angle-of-arrival and three-bounce time-of-flight associated with these other mirror image spots into Eq. (5).

If the direction of the transmitted beam is known then we can also compute the position of the directly illuminated point  $S_1$ . We begin by computing the apparent distance  $r_{D'L}$  from the laser to  $D'$ :

$$r_{D'L} = ct_3 - r_{D'C} = ct_2 - r_{DC}. \quad (6)$$

The distance from  $L$  to  $S_1$  can then be computed using the identity  $r_{DS_1} = r_{D'L} - r_{S_1L}$  and the law of cosines from the triangle  $S_1LD$ . This distance is

$$r_{S_1L} = \frac{1}{2} \frac{r_{D'L}^2 - r_{DL}^2}{r_{D'L} - r_{DL} \cos \delta_L}, \quad (7)$$

where  $r_{DL} = \|D - L\|_2$  and  $\cos \delta_L = \frac{1}{r_{DL}} (D - L) \cdot \hat{d}_{S_1L}$ . Here  $\hat{d}_{S_1L}$  is the direction of the transmitted beam. Finally, from the law of reflection, we can compute the surface normal vectors at  $S_1$  and  $S_2$  once  $D$ ,  $S_1$  and  $S_2$  are known.

**Identifying the true spot and the first surface.** Although we have derived expressions that could in theory be used to compute the positions of all diffuse and specular scattering points within the scene, there are two ambiguities that need to be resolved before these expressions can be applied. First, we must determine whether the directly illuminated surface is a diffuse or a specular reflector. Doing so from raw measurements is not as straightforward as it might seem because in either case our receiver will see at least two spots, and one of these spots will appear to lie along the transmitter's boresight. Second, regardless of which surface was illuminated first, we must also determine which observed spot is the true laser spot and which are mirror images.

Resolving the second ambiguity is always straightforward if we have time-of-flight information. *Light from the true spot always arrives before the light from the mirror images.* This can be confirmed by inspecting Fig. 2. In this figure all propagation paths travel through  $D$ . However, light from the true spot travels directly from  $D$  to  $C$ , whereas light from the mirror images must follow a longer, indirect path that includes an additional reflection. The first-surface ambiguity can be resolved once the true spot has been identified. *If the true spot lies along the transmitter boresight, then a diffuse reflector was illuminated first. If it doesn't, then the specular surface was illuminated first.* Once both ambiguities have been resolved, the appropriate range equations can be applied to determine the scattering points.

It is worth noting that we are only able to disambiguate these two cases because we have time-of-flight information. If we had instead been viewing the scene with a regular camera that only measured angles and intensities, then there would be no direct way to determine which spot was the true spot and which surface was illuminated first. Instead, we would have to rely on contextual cues to determine, for instance, which pixels seemed to have mirrors in them.

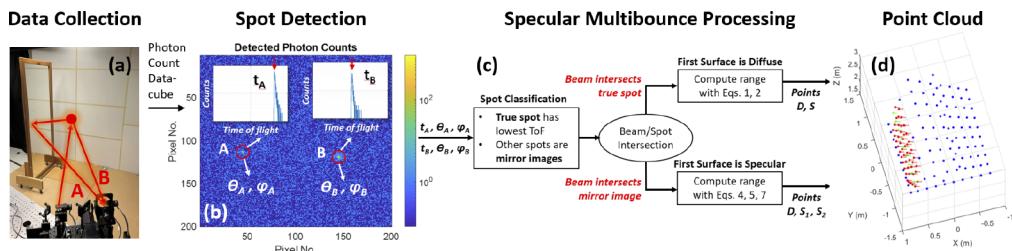
### 3.2. Single-beam specular surface scanning method

The principles described in Section 3.1 can be used to define a simple method for scanning specular surfaces. The main steps of this method are summarized in Fig. 3. First, a single, pulsed, pencil-beam source illuminates the scene of interest, and a wide FOV receiver measures backscattered light. The time- and angle-of-arrival of all visible laser spots are extracted from these measurements. From these extracted metrics, the true laser spot  $D$  is identified and other spots are classified as mirror images. If the true spot appears to lie along the transmitted beam vector, Eqs. (1) and (2) are used to compute the range to all spots. Otherwise, Eqs. (4), (5), and (7) are used. Finally, the computed ranges are combined with angle-of-arrival information (or the transmitted beam direction, for  $S_1$  points) to determine the 3D position of all diffuse and specular scattering points. The beam can be scanned to build up a point cloud representation of diffuse and specular surfaces in the scene.

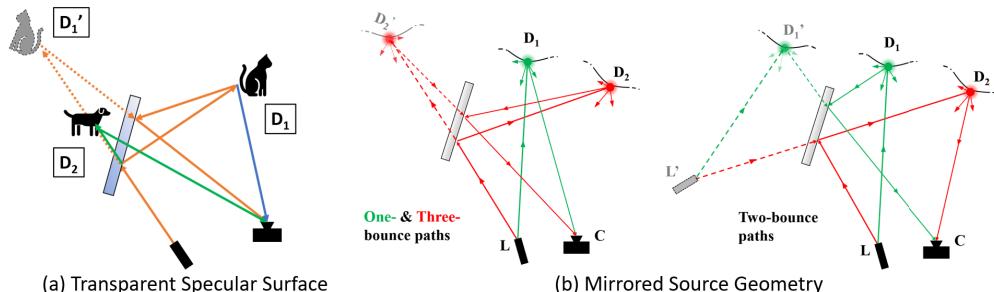
### 3.3. Transparent specular surfaces

Many of the most common specular surfaces—glass windows, for example—are *partially transparent*. Transparency introduces an additional source of ambiguity into our measurements, which is highlighted in Fig. 4(a). Here we see that, as was true for the mirror surface, three-bounce returns will typically resemble spots that lie behind the specular surface. When that surface is also transparent, however, these three-bounce returns may be mixed together with one-bounce returns that have scattered diffusely off of something that lies behind (or on) the specular surface. In this section we propose several criteria that can be used to detect and disambiguate these two kinds of signals. We note that these criteria need only be applied when the specular surface is illuminated directly. When the first surface is a diffuse reflector, the ambiguities described here will not occur.

**Multiple detections along beam.** Our disambiguation logic is triggered when we detect multiple spots lying along the transmitted beam vector, and one spot that does not lie along the beam vector. When this occurs, it suggests that we may have detected the mirror image of a true



**Fig. 3.** A summary of the proposed single-beam specular surface scanning method. (a) A pencil-beam source illuminates the scene of interest, and a wide FOV receiver observes all backscattered light. (b) Time- and angle-of-arrival of all visible laser spots are extracted. (c) The true spot and first-illuminated surface are determined, and this information is used to apply the appropriate range equations, and thus compute the position of all diffuse and specular scattering points. (d) The beam is scanned to build up a point cloud representation of diffuse and specular surfaces in the scene.



**Fig. 4.** (a) When a transparent specular surface is illuminated directly, multiple laser spots may be detected along the beam vector—a true spot  $D_2$  that lies behind the surface, and a mirror image ( $D'_1$ ) of a spot  $D_1$  that lies in front of the surface. (b) Two-bounce paths (right) appear to originate at  $L'$ , a mirror image of the true source. One- and three-bounce paths (left) appear to originate at the true source  $L$ .

spot that lies in the scene in front of the window, in addition to a one-bounce return from a true spot that lies behind the window or on its surface.

**Comparison to two-bounce travel time.** When the window is illuminated directly, light that arrives from the true spot in front of the window follows a *two-bounce* path, and the spot will not lie on the beam vector. As was explained in Sec. 3.1, the mirror image in this scenario must correspond to a three-bounce return, and the associated three-bounce time-of-flight must be *greater* than the two-bounce time-of-flight from the true spot. Thus, if the time-of-flight associated with any of the spots that lie along the beam vector is *less* than the two-bounce time-of-flight, we deduce that they cannot be mirror image returns and, thus, must be single-scatter returns from on or behind the window surface.

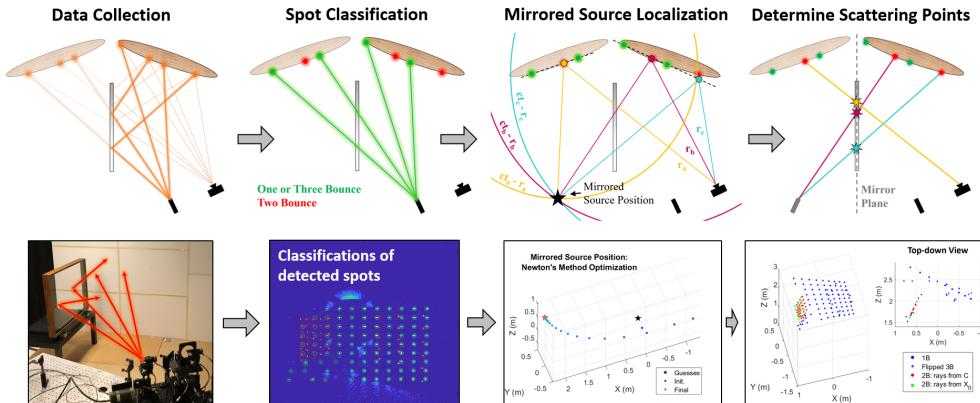
**Multiple reflections diminishes intensity.** If the time-of-flight associated with more than one in-beam spot is greater than the two-bounce time-of-flight, we infer that the *dimmest* spot corresponds to the mirror image. The reasoning behind this is that three-bounce signals are reflected twice by the window, whereas one-bounce signals are transmitted twice. Common transparent, specular materials (such as glass) typically transmit more light than they reflect, although this assumption can break down at glancing incidence angles. To choose the dimmest spot, we rank in-beam detections by their range-adjusted intensity ( $\sim r^2 I$ ). For this computation, we use the *apparent* range of each spot, computed using Eq. (1).

**Summary of transparent surface processing** Our transparent surface processing logic can be summarized as follows:

1. Transparent surface processing is triggered when at least two spots are detected that lie along the beam vector, and one spot is detected that does not lie along the beam vector.
2. The spot that does not lie along the beam vector is assumed to be  $D_i$ —the true spot that lies inside the window, as shown in Fig. 4(a). The time-of-flight associated with this spot is labeled  $t_{D_i}$ .
3. Spots along the beam vector with time-of-flight  $t \leq t_{D_i}$  are classified as one-bounce, diffuse returns
4. If there are spots along beam vector with time-of-flight  $t > t_{D_i}$ , then, within this set of spots, the spot with the dimmest range-adjusted intensity is determined to be  $D'_i$ , the mirror image of  $D_i$ . The rest are classified as diffuse, one-bounce returns.
5. Compute range to all diffuse, one-bounce returns using Eq. (1).
6. If  $D'_i$  is detected, compute range to  $D_i$  using Eq. (4), and to  $D'_i$  using Eq. (5). Otherwise discard detection associated with the true spot  $D_i$ .

### 3.4. Illumination with multiple beams

It is straightforward to employ the concepts introduced in Secs. 3.1 and 3.3 to acquire specular-surface geometry if the scene can be scanned by transmitting a single beam at a time. Unfortunately, the time required to do so may be unacceptably long for some applications. In this section we propose an algorithm that permits the mapping of specular surface geometry without any knowledge of spot-to-beam associations, and can thus be implemented when the scene is illuminated by a multi-beam flash. The main steps of this algorithm are visualized in Fig. 5.



**Fig. 5.** A summary of the proposed multi-beam specular surface scanning method. First, data is collected using a wide-FOV receiver and a multi-beam flash transmitter. Second, one- and three-bounce spots are identified by intersecting apparent spot positions with transmitted beam vectors. Third, the positions of two-bounce scattering points are approximated by interpolating from the nearest one- or three-bounce spots. The position of the mirrored source is then computed by solving a multilateration problem, with two-bounce travel times to compute ranges to the mirrored source. Finally, the mirrored source position defines the plane of the mirror, and scattering points on the mirror are determined by intersecting the mirror plane with rays drawn to two-bounce spots from the receiver, and from the mirrored source.

**Mirrored source geometry** If a transmitter emits a pulse at time  $t$  in a scene that contains a flat mirror, then a receiver that is pointed at the mirror will observe a mirror image of the transmitter that also appears to emit a pulse at  $t$ . The direction that the mirrored pulse propagates will be flipped across the mirror plane. This geometry is visualized in Fig. 4(b), where the true source is at  $L$ , and its mirror image is at  $L'$ . The position of  $L'$  is significant because the plane of the mirror perfectly bisects the line segment  $LL'$ , and the plane's surface normal is parallel to  $LL'$ . This means that estimating the position of  $L'$  (relative to  $L$ ) is equivalent to estimating the mirror plane. This is the key principle that underlies our multi-beam surface estimation algorithm.

The one-, two-, and three-bounce signals described in previous sections can also be interpreted under a mirrored space model. From the perspective of the receiver, one- and three-bounce returns appear to originate from the true source  $L$ . Two-bounce returns appear to originate from the *mirrored* source, at  $L'$ . From Fig. 4(b), we see that the order of the two bounces determines where the scattering event will appear to occur. If light bounces off of the mirror first, then the scattering event appears to occur in the true space at  $D$ . If it bounces off a diffuse reflector first, then light appears to scatter in the mirrored space at  $D'$ .

Because spots produced by one- or three-bounce light paths appear to originate from the true source, they must also lie along the axis of one of the true, transmitted beams. Spots that result from two-bounce light paths, on the other hand, will not in general lie along a transmitted beam vector. Because of this it is straightforward to determine which detected spots correspond to two-bounce light paths. Images of spot detections classified in two multi-beam data collections are shown in Fig. 10.

**Source localization using multilateration** If we could compute the apparent positions of at least three two-bounce spots, then we could compute the ranges  $r_{DL'} = ct_2 - r_{DC}$  or  $r_{D'L'} = ct_2 - r_{D'C}$ , and determine  $L'$  by solving a multilateration problem. Although we cannot compute these positions directly without knowledge of spot-to-beam associations, we can approximate them. From one-bounce returns, we compute the positions of several points in the true space using Eq. (1). We also compute the positions of a number of points in the mirrored space from three-bounce returns, using Eq. (3). We then approximate the positions of the two-bounce spots by interpolating from the positions of the nearest (in apparent angle) one- and three-bounce spots.

Once we've computed these approximate positions we can estimate the mirrored source position  $L' = [x_0, y_0, z_0]^T$  by solving the following optimization problem

$$\{x_0, y_0, z_0\} = \arg \min_{x_0, y_0, z_0} \frac{1}{2} \sum_i \left[ (x_i - x_0)^2 + (y_i - y_0)^2 + (z_i - z_0)^2 - (ct_i - r_i)^2 \right]^2. \quad (8)$$

Here  $(x_i, y_i, z_i)$  is the approximate position (in  $C$ -centered coordinates) of the  $i^{th}$  detected two-bounce spot,  $r_i = \sqrt{x_i^2 + y_i^2 + z_i^2}$  and  $t_i$  is the two-bounce travel-time associated with the  $i^{th}$  spot. The objective function is twice-differentiable, so we can solve the problem using Newton's method. In practice, misclassifications of two-bounce spots as one- or three-bounce spots, or vice-versa, produce large errors in the source localization result. To make our algorithm more robust to misclassification errors, we solve Eq. (8) using a RANSAC [16] approach that is robust to outliers. A more detailed description of this approach is provided in [Supplement 1](#).

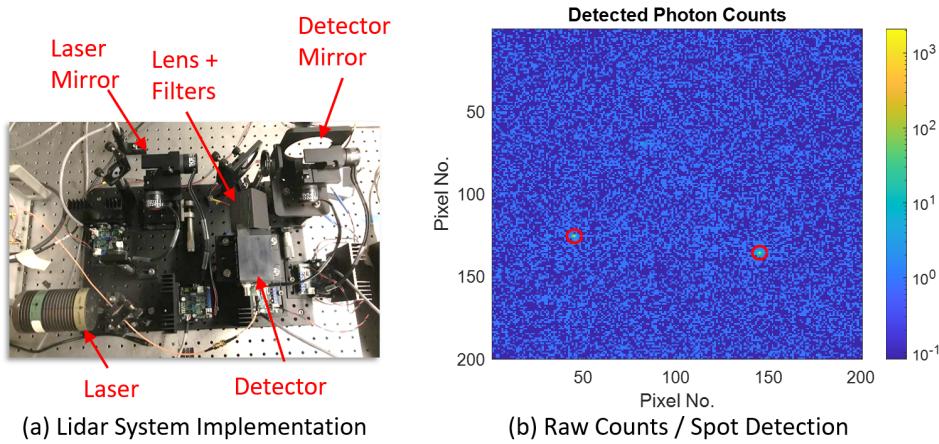
**Determining scattering points** As can be seen in Fig. 4(b), rays drawn from the receiver to two-bounce points lying *behind* the mirror plane will intersect the mirror plane at reflection points, as will rays drawn from the mirrored source to two-bounce spots *in front* of the mirror plane. Thus, once the mirror plane has been computed from our estimate of  $L'$ , we can find all points of specular reflection. From these reflection points we can also estimate the mirror's boundary. This allows us to determine which spots correspond to three-bounce paths. The apparent positions of three-bounce spots can then be reflected across the mirror plane to retrieve the true diffuse scattering positions. This can be achieved even when the true scattering points lie outside of the receiver's field-of-view (e.g. if they are hidden around a corner).

It is important to note that Eq. (8) only applies to planar specular surfaces, and so this algorithm will not produce accurate reconstructions of non-planar specular surfaces.

## 4. Results

### 4.1. Implementation

**LiDAR System.** A photograph of our LiDAR system is provided in Fig. 6(a). The transmitter consists of a focused, pulsed laser source (640nm wavelength) that is scanned using a two-axis mirror galvanometer. The receiver is a single-pixel SPAD detector with a focused FOV that can be scanned independently from the laser using a second set of galvo mirrors. The overall instrument response function (IRF) of our system was measured to be 128 ps (full width at half maximum). Details concerning the specific equipment used can be found in [Supplement 1](#).



**Fig. 6. Implementation.** (a) Our LiDAR system consists of a focused, pulsed galvo-scanned laser and a single-pixel SPAD that can be scanned independently of the laser. In our experiments the SPAD FOV is scanned across a dense, uniform grid to mimic the angular sampling pattern of a SPAD array camera. (b) Per-pixel map of detected photon counts. Detected spots are circled in red.

For each beam pointing direction, we reproduce the angular sampling pattern of a SPAD array camera by scanning the FOV of the detector across a dense, uniform grid. For experiments reported in this paper, the per-pixel dwell time was either 5 or 10 ms. The laser was operated at a pulse repetition frequency of 20 MHz and an average transmitted power of 5  $\mu$ W.

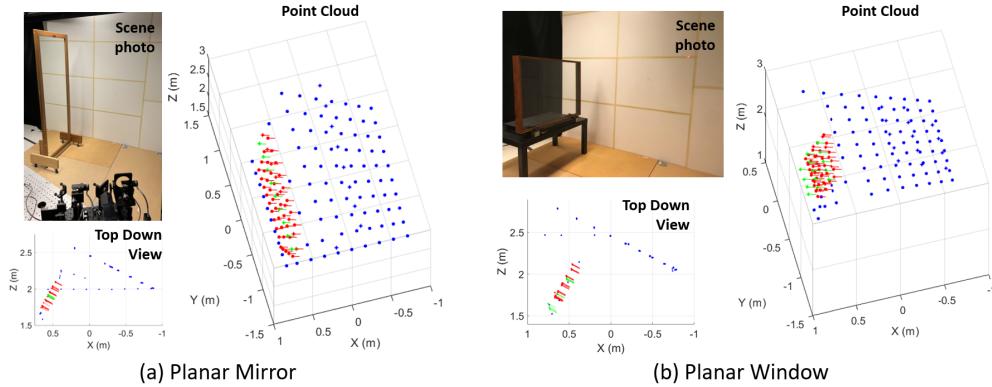
**Spot extraction.** Raw photon count measurements are sorted into a data cube that is binned by photon time-of-arrival and detector scan angle. From these raw measurements we detect spot-like returns and then extract the time-of-flight, angle-of-arrival, and returned energy of each spot. An image of per-pixel detected photon counts collected from a single detector scan is shown in Fig. 6(b). Detected spots are circled in red. A more detailed description of our spot detection and low-level signal processing pipelines is provided in [Supplement 1](#).

### 4.2. Planar surface scans

#### 4.2.1. Planar mirror

We scanned a scene (shown in Fig. 7(a)) that contained a tall, flat mirror, a wooden floor, and a matte white wall. The mirror was placed approximately 2m away from the LiDAR scanner. We scanned the laser beam along a  $10 \times 10$  grid of pointing directions. 14 beams illuminated the mirror directly. The remainder illuminated the floor, wall, or the mirror's wooden frame. For

each beam direction the detector FOV was scanned across a  $200 \times 200$  grid that subtended a  $\pm 30^\circ$  angle of view in both vertical and horizontal directions. The per-pixel dwell time was 5 ms. Total acquisition time was 5 hours and 33 minutes, although we note that an equivalent data collection made using a  $200 \times 200$  pixel SPAD array would have been captured in 0.5 seconds.



**Fig. 7.** Scans of (a) a large flat mirror and (b) a single-pane glass window. Red and green points indicate points on specular surfaces, blue points indicate points on diffusely reflecting surfaces. Circles indicate diffuse-first illumination geometry and asterisks indicate specular-first geometry. Surface normals are also plotted at specular points.

Our results are shown in Fig. 7(a). Here, blue points represent points on diffuse surfaces, red points are points on the mirror as seen from the receiver ( $S$  or  $S_2$  in Fig. 2), and green points are points on the mirror that were illuminated directly ( $S_1$  in Fig. 2). Points computed using diffuse-first equations are marked by circles, whereas points computed using specular-first equations are marked with asterisks. Surface normals are also plotted for specular surface points.

It is evident from Fig. 7(a) that the point cloud accurately captures the dimensions of the scene. From a comparison to a ground truth scan collected on the mirror's frame, we determined that the points on the mirror surface had a root-mean-square (RMS) displacement of 9.5mm with respect to the ground truth plane, and the surface normals had an RMS tilt of  $0.65^\circ$  with respect to the true surface normal. A complete evaluation of the scan's accuracy is provided in Section 5.1.

#### 4.2.2. Planar glass window

We also scanned the shape of a single-pane glass window, which is shown in Fig. 7(b). The window was placed in the same position and orientation as the previously scanned mirror. For this collection we did *not* place any objects behind the window, and so the primary challenge was that the lower reflectance of the glass resulted in fainter multibounce returns, particularly for three-bounce light. We found that each reflection off the window reduced the range-adjusted intensity of a spot by approximately a factor of 10.

We compensated for the lower reflectance by doubling the per-pixel dwell time, to 10 ms. The laser scan pattern was changed to an  $11 \times 9$  grid that more densely sampled the portion of the scene that contained the window. The detector scan pattern was identical to that used in the mirror experiment. The total acquisition time was 11 hours, although we note that an equivalent data collection could have been captured by a SPAD array in just under one second.

Our results are shown in Fig. 7. The window pane can be seen clearly, and appears to have been mapped accurately. Despite the lower reflectance of the glass, we are able to detect all two-bounce and three-bounce returns with no misses or false alarms. However, when the point cloud is viewed from the top-down it is clear that the points computed from three-bounce returns (marked

by red and green asterisks) are noisier than two-bounce points. This was likely a consequence of the lower relative intensity of three-bounce returns.

#### 4.3. Detecting objects behind a window

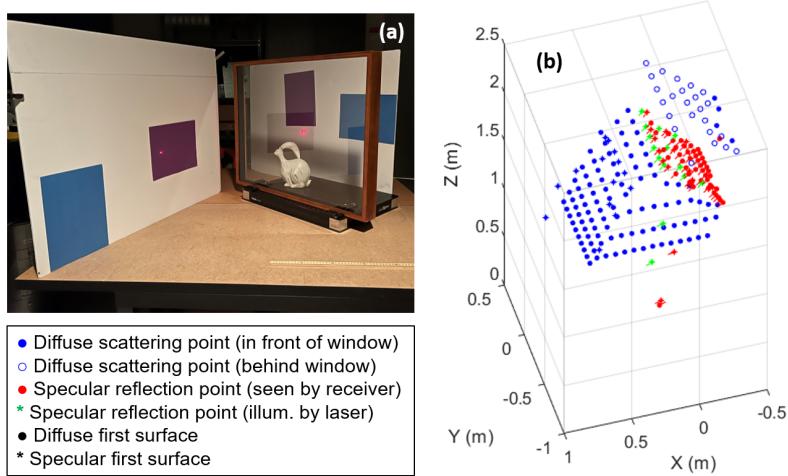
##### 4.3.1. Multiple objects

We placed two objects behind the window and repeated the window scan described previously. Our results are shown in Fig. 1. Behind-window detections are marked by white dots circled in blue. From these results we see that our disambiguation logic accurately classified four spots that scattered off of the objects behind the window. We note that all other points that appear to lie behind the window in the top-down view of the point cloud in fact lie above or below the window aperture.

**Comparison to naïve single-bounce processing.** On the right side of Fig. 1 we show the point cloud that would have been generated if the positions of all spots had been naïvely computed using the conventional one-bounce range equation. Notably, this naïve processing detects *no* points on the window's surface. Detections that might have been mapped to the window surface are instead mapped to erroneous points. Only true one-bounce returns are mapped correctly.

##### 4.3.2. Test of transparent surface processing

We performed a second experiment to rigorously test the disambiguation criteria described in Section 3.3. In this experiment the scene (shown in Fig. 8(a)) contained two large, diffusely reflecting walls that were respectively placed in front of and behind a single-pane glass window. Dark pieces of construction paper were pasted onto each wall to produce surfaces that varied between high and low reflectance. In some look directions, the mirror image of the front wall appeared to be closer to the detector than the true plane of the back wall, whereas in other look directions it appeared to be more distant. This assured that both the time-of-flight and reflectance criteria proposed in Sec. 3.2 were put to the test. For added scene complexity, we also placed a small rabbit figurine immediately behind the window.



**Fig. 8.** (a) Photograph of scanned scene which included a glass window as well as surfaces in front of and behind the window that had varying reflectance. (b) Point cloud acquired for the scene. Marker key is shown on the bottom left.

We scanned the scene using a pattern of  $14 \times 10$  beams scanned in a uniformly spaced rectangular grid. The detector scan was  $190 \times 140$  pixels that spanned  $50.3^\circ \times 37.3^\circ$  with a

per-pixel dwell time of 10ms. Total scan time was 11 hours, 7 minutes, and 20 seconds, although a SPAD array could have made an equivalent measurement in 1.4 seconds.

In contrast to previous experiments, here there were many occasions where the three-bounce reflection off of the window surface was too faint to detect. This frequently occurred when the laser beam was deflected onto a dark surface on the front wall. Additional processing logic was used to reduce false mapping of detections in this scenario. This logic is described in detail in [Supplement 1](#).

Our results are shown in Fig. 8(b). We are able to retrieve the surface of the window as well as the surface of the two large walls and the floor. A single detection on the rabbit figurine was also registered. However, there are a few artifacts in our point cloud that are worth mentioning. Three (of 140) beams produced points in false positions. These points can be seen floating in the space between the LiDAR system and the window, and floating behind the front wall. In all three cases, the true three-bounce returns could not be detected. When both the three-bounce and the behind-window returns were detected, our discrimination was always correct. Many points on the back wall are also missing. Most of these missing points correspond to portions of the wall that lay in the shadow of the rabbit figurine. The rest correspond to the falsely mapped points described previously. Finally, four spots lying behind the window were classified as lying in front of the window. This occurred during exposures for which the spot behind the window was the only spot detected, and was thus treated as a regular spot on a diffusely reflecting surface.

#### 4.4. Non-planar object scans

##### 4.4.1. Polished metallic object

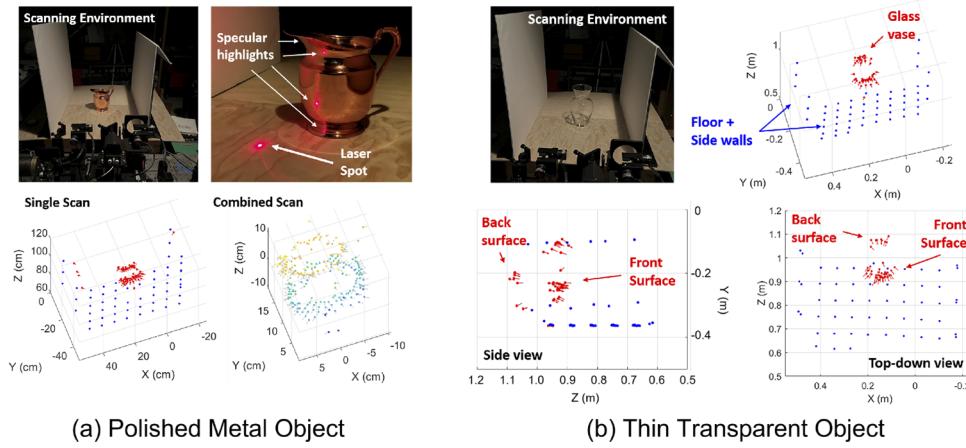
Our method can also be used to acquire the shape of non-planar specular surfaces. We demonstrated this by scanning the shape of a copper pitcher. A photograph of this pitcher and the results of our scan are shown in Fig. 9(a). The pitcher was placed on a wooden floor and in between two white side walls. We illuminated a sequence of 60 laser spots on these surfaces, and observed two-bounce returns that reflected off of the pitcher.

Laser spots and two-bounce highlights were acquired using a  $200 \times 100$  pixel scan with a 5 ms per-pixel dwell time. We repeated the scan four times, rotating the pitcher  $90^\circ$  between each scan so that we could acquire its front, back, and side-facing surfaces. The four point clouds were aligned to a pitcher-centered coordinate system and then combined. From the combined point cloud in Fig. 9(a), we see that we are able to recover the general shape of the pitcher, although there are several gaps that correspond to regions that did not produce any detectable highlights due to their surface orientation.

We point out two distinctive features of the non-planar surface scanning process. First, because the pitcher's surface alternated between convex, concave, and hyperbolic curvature, one laser spot would produce multiple highlights, each of which could be used to locate a point on the pitcher's surface. This can be seen in the photograph on the top-right of Fig. 9(a). Second, Eqs. (4) and (5) are only correct when the tangent planes of surface points  $S_1$  and  $S_2$  are coincident. This is automatically satisfied for monostatic ( $s = 0$ ) LiDAR systems, for which  $S_1 = S_2$ . However, because our scanner was bi-static, these equations could not be applied for non-planar surface mapping. Consequently, all points in Fig. 9(a) had to be acquired by directly illuminating spots on diffuse reflectors, and observing two-bounce returns reflected by the pitcher. Specular-first returns could not be used. However, using the criteria outlined in Sec. 3.1, specular-first returns could be detected automatically and discarded.

##### 4.4.2. Thin transparent object

In addition to scanning a metal pitcher, we also scanned a glass vase. We used the same detector and laser scan patterns that were used to acquire the shape of the copper pitcher in Sec. 4.3 of the



**Fig. 9.** We scan the shape of a copper pitcher (a) and a glass vase (b) by directly illuminating a sequence of 60 laser spots on nearby, diffusely reflecting surfaces and measuring the time-of-flight of two-bounce returns that reflect off of the specular object. For single scans, blue dots represent points on the floor and side-walls, red dots are points on the vase, and red arrows are the surface normal vectors associated with those points. For the copper pitcher (a), we repeated the scan four times, rotating the pitcher 90° between scans, to acquire the full shape (a, bottom right). Notably, because the vase (b) was transparent, we were able to map points on the front and back surfaces of the vase without moving our LiDAR, or rotating the vase.

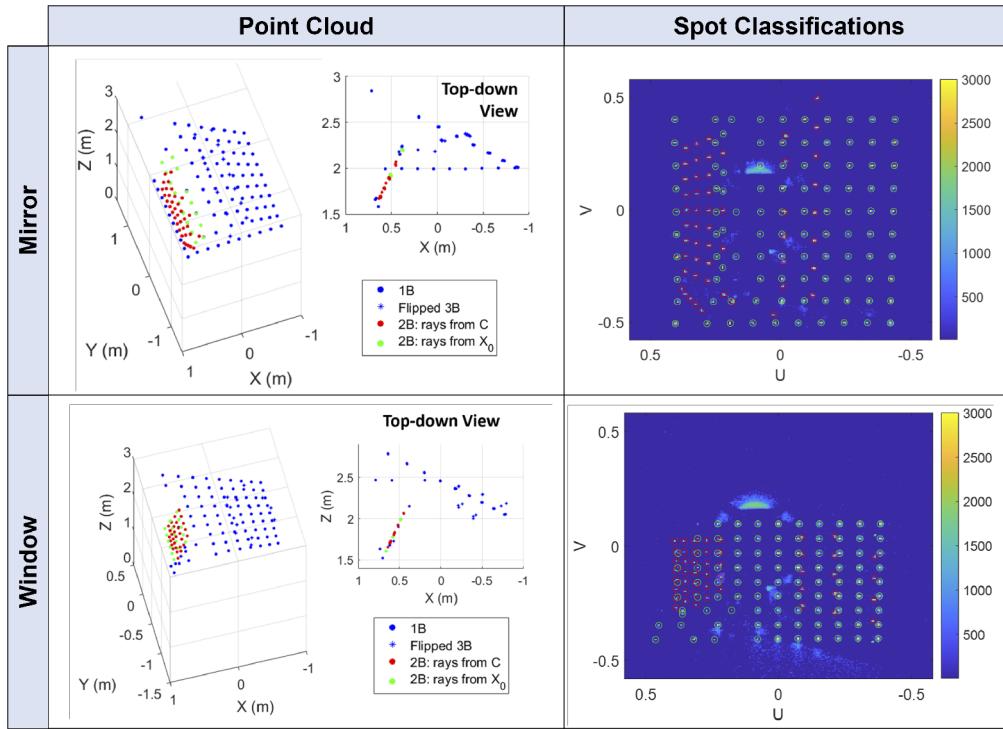
main text, but increased the per-pixel dwell time of our detector to 10ms to account for the lower reflectivity of the glass.

Our result is shown in Fig. 9. Due to the rotational symmetry of the vase, we did not rotate it to measure the shape of all sides. Interestingly, after scanning the vase from just a single vantage point, we noticed that we were able to measure points on both the front *and* the back surface of the vase, due to the fact that the vase was transparent. Furthermore, because the vase is a thin glass object, we expect that the effect of distortion of the returned signals due to refraction should be small. Thus, the position of these back surface points should be relatively accurate. Surface normals for back-surface points are identified as inward-facing, whereas surface normals on the front surface are outward facing. This is technically correct, and in each case corresponds to the true direction of the surface normal vector at the point of specular reflection.

#### 4.5. Multiple-beam illumination

To test our multi-beam surface mapping algorithm, we summed the single-beam photon count histograms acquired during the planar mirror and window scans described in Sec. 4.2. This produced the equivalent of a single multi-beam flash exposure for each scene. Detected per-pixel energy values from this multi-beam data are shown on the right-hand side of Fig. 10. Detected laser spots that were classified as one- or three-bounce spots are circled in green, and those classified as two-bounce spots are circled in red.

Our results are shown on the left-hand side of Fig. 10. By visual inspection we see that the computed points match the single-beam results shown in Fig. 7 very closely. Points in the multi-beam results in fact appear less noisy, although this is because they are intersections with a single, fitted plane. Many reflected three-bounce points appear in erroneous positions. Most of these points correspond to two-bounce spots that were incorrectly classified as three-bounce spots during the beam intersection step of our algorithm.



**Fig. 10.** (Left) Point clouds generated using multi-beam surface mapping algorithm on mirror (top) and glass window (bottom) dataset (see Fig. 7). (Right) Detected per-pixel energy values of multi-beam data collection, overlaid with spot positions. Red circles mark two-bounce spots, and green circles mark one- or three-bounce spots.

Interestingly, because our multi-beam algorithm did not require spot-to-beam associations, it was able to reclaim certain points that were not computed during the original single-beam scan of the mirror. These additional points corresponded to points directly illuminated near the edge of the mirror's surface ( $S_1$  in Fig. 2(b)). In the original scan, no highlight ( $S_2$ ) was observed for these beams because the point of reflection would have been beyond the mirror's edge, and thus  $S_1$  was not computed. However, in the multi-beam algorithm, the intersections of all beam vectors with the fitted mirror plane are computed, including the points that had no associated highlight measurement.

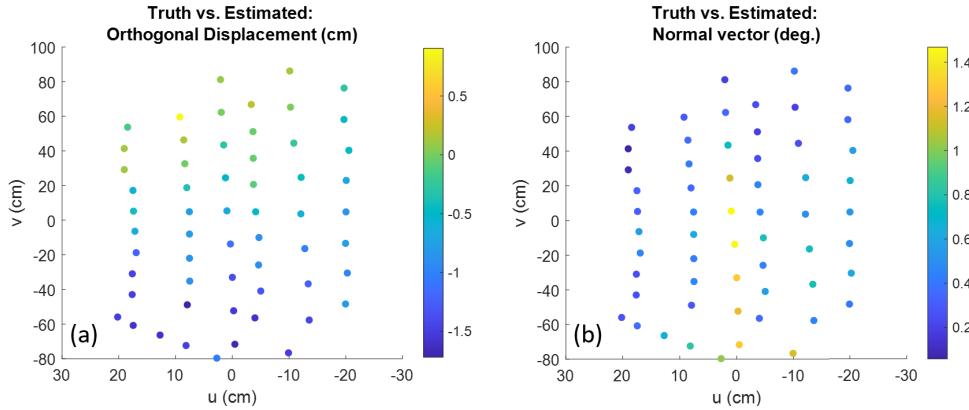
## 5. Analysis

### 5.1. Analysis of errors in mirror scan experiment

To validate our method, we investigated the accuracy of the large mirror scan that was presented in Sec. 4.2.1. We fit a plane to a set of 22 ground truth points on the mirror's frame. These ground truth points were acquired during a separate data collection that used a dwell time of 25ms per pixel. This longer dwell time allowed us to collect more photons per point and thus make more precise position estimates. All points on the mirror frame reflected light diffusely, so their positions could be computed from the conventional single-scatter range equation. The ground truth plane fit was then adjusted to account for the fact that the mirror surface was recessed 1.6cm behind the plane of the frame.

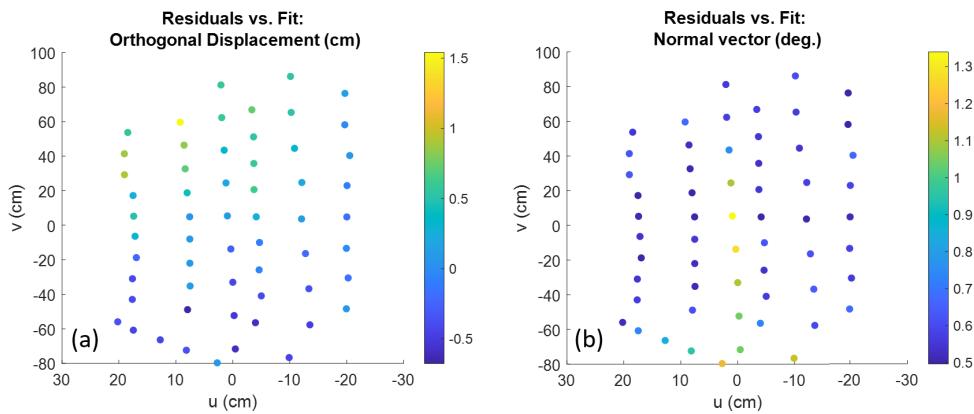
We compared the position and surface normal of mirror surface points computed from the original scan data to this ground truth plane. In Fig. 11(a) we plot the orthogonal displacement

of all estimated mirror plane points as a function of their projected position on the ground truth plane. The RMS displacement of these points was  $9.4\text{mm}$  and the mean displacement was  $-7.1\text{mm}$ . In Fig. 11(b) we show the angular displacement of the estimated normal vectors from the ground truth normal vector. The RMS angular displacement was  $0.63^\circ$  and the mean displacement was  $0.54^\circ$ . The presence of small but significant mean errors suggest a systematic bias in our results. This bias could have resulted from many factors. First, the angular calibration of our system was not perfect, and errors in angle-of-arrival estimates of diffuse surface returns may have been coupled into specular surface positions. Second, the ground truth scan was taken 24 hrs after the initial scan, and it is possible that the mirror moved or tilted slightly in between scans. Finally, the presence of a glass first-surface may have also contributed to the bias. The mirror's protective glass was  $6.35\text{mm}$  thick, and we did not account for the reduced speed of light within the glass while estimating the positions of points on the mirror surface. To do so would have been non-trivial because the increase in travel time would have been dependent on angles of incidence and exitance with respect to the glass surface. Another prominent artifact is the larger normal vector errors associated with points computed from specular first-surface measurements as opposed to diffuse first-surface measurements. These points form a vertical line down the center of the plane which can be plainly seen in Fig. 11(b).



**Fig. 11. Mirror scan error analysis.** Plots of (a) distance of points from ground truth mirror plane and (b) angular displacement of surface normal vectors from the ground truth plane. Errors are plotted as a function of the projected position of each point onto the ground truth plane.

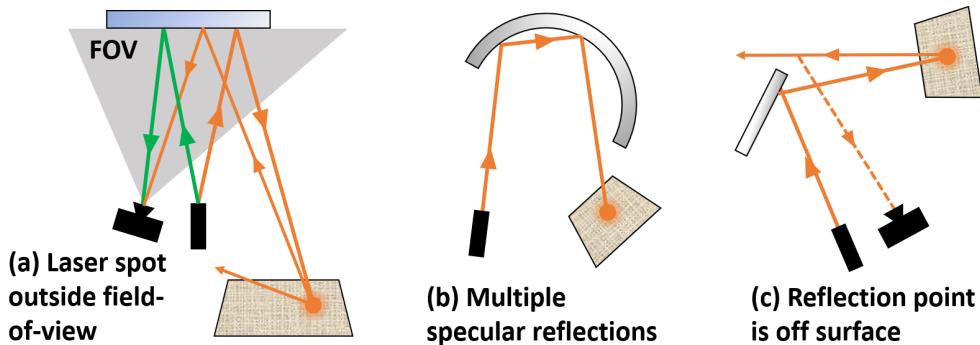
To assess the precision of our estimates, we fit a plane to the collection of mirror surface points and then compared individual points to the plane fit. Because each oriented point includes a position and a surface normal vector, each oriented point defines a plane. Thus, to "fit" a plane to these oriented points, we simply need to compute the plane parameters for each individual point and then take the average of these parameters. We define a plane by the identity  $\hat{n}\mathbf{x} = d$ . The fit plane had the parameters  $\hat{n}_{fit} = [-.8797 - .0048 - .4754]^T$  and  $d_{fit} = -1.406\text{m}$ . For comparison, the ground truth plane had parameters  $\hat{n}_{truth} = [-.8825 - .0010 - .4704]^T$  and  $d_{truth} = -1.389\text{m}$ . We plot the orthogonal displacement of all estimated mirror plane points with respect to the fitted plane in Fig. 12(a). We show the angular displacement of the estimated normal vectors from the mean normal vector in Fig. 12(b). The RMS orthogonal distance was  $4.7\text{mm}$  and the RMS angular displacement was  $0.70^\circ$ .



**Fig. 12. Mirror scan residuals analysis.** Plots of (a) distance of points from the fitted mirror plane and (b) angular displacement of surface normal vectors from the fitted plane. Residuals are plotted as a function of the projected position of each point onto the fitted plane.

### 5.2. Common failure cases

Our methods can detect and localize specular surfaces that might be invisible to a conventional LiDAR system that relies on measurements of direct, single-scatter returns. Nevertheless, there are specific circumstances that cause our method to fail. We depict some of these cases in Fig. 13.



**Fig. 13.** Illustrations of three common failure cases of the proposed specular surface mapping method.

#### 5.2.1. Only one spot is detected

When the scene is illuminated with a single beam, the position of the specular surface can only be determined if *both* the true laser spot and its mirror image are detected. There are many scenarios in which this condition is not met. The most obvious is when the laser beam reflects off of a specular surface but never lands on a diffusely reflecting surface. This scenario is more common in wide-open outdoor environments in which the beam might get reflected towards the sky. A scenario that is more common in enclosed environments is when the beam is deflected onto a surface that lies outside the receiver's FOV. This scenario is illustrated in Fig. 13(a). This scenario often occurs when a specular surface is illuminated at near-normal incidence, such that the beam lands on a diffuse reflector that is behind the LiDAR. One remedy in this context is to use a receiver with a very wide FOV. Although such receivers may not be typical, many

LiDAR systems used by autonomous vehicles have a 360° horizontal FOV [17]. Additionally, while near-normal incidence angles challenge methods that rely on specular multibounce returns, they are also more conducive to observations of the direct, single-scatter returns that can be interpreted using conventional processing. Lastly, it is also possible that an occluding surface will block the receiver's line-of-sight to the true laser spot.

The consequences of detecting only one of two spots depend upon which spot is detected and which surface was illuminated first. If a diffuse reflector is illuminated first and only one-bounce returns are visible, then the diffuse reflection point can be determined accurately using Eq. (1) but we will learn nothing about the specular reflector. If the specular reflector is illuminated first and only the three-bounce return is observed, then the three-bounce return may be interpreted as a one-bounce return, causing a false point to be mapped behind the specular surface. Regardless of which surface was illuminated first, if only the two-bounce return is detected then it can be identified and discarded without producing false points. This is because the direction that two-bounce light returns from will not, in general, match the direction of the transmitted beam.

### 5.2.2. Multiple specular reflections

Our method does not consider propagation paths that include two or more specular reflections in a row. We've depicted one of these propagation paths in Fig. 13(b). This scenario can occur when scanning a free-form specular object with concavities, or when a scene has multiple planar reflectors that face each other. Although we do not treat such cases in this work, it is likely that the additional specular points can be mapped in these circumstances if the additional highlights that they produce are observed. Alternatively, certain cues such as chirality flips of an asymmetric beam pattern might be used to identify and discard multi-specular returns.

### 5.2.3. Reflection point is off surface

If there is a non-zero baseline separation between the transmitter and the receiver (that is, if the LiDAR system is bi-static), then there will be a disparity between the point that the transmitter directly illuminates on a specular surface ( $S_1$ ) and the second point of reflection ( $S_2$ ), which corresponds to the specular highlight that is visible to the receiver. This sometimes leads to the scenario depicted in Fig. 13(c) when a specular surface is illuminated near its edge. Here, the point that *would* have reflected light back towards the receiver is just empty space—it's beyond the edge of the specular surface. In this case no three-bounce return is observed, but a two-bounce return is still produced. In such a scenario,  $S_1$  can be recovered by finding the point at which the beam vector intersects a plane (or other shape) fit to *other* specular surface points that were computed using measurements corresponding to different beam vectors. Once  $S_1$  is retrieved, the position of the diffuse (two-bounce) spot can also be determined using angle and time-of-flight constraints.

## 6. Conclusion

We have demonstrated methods that use multibounce specular LiDAR returns to detect and map specular surfaces that might otherwise be invisible to LiDAR systems that rely on single-scatter measurements. We considered the cases of single-beam and multiple-beam illumination, and demonstrated our methods by scanning planar, non-planar, and transparent surfaces.

Because specular surfaces are relatively common, our work could be used in most domains that LiDAR scanning is applied to, including autonomous navigation, mapping of indoor or outdoor spaces, and object scanning. Future work might address failure cases, such as propagation paths with consecutive specular reflections, or geometries for which the true or mirror image spot lies outside the receiver FOV. We would also be interested in extending our multi-beam algorithm to the mapping of curved surfaces, and to allow dense flash illumination.

**Funding.** Office of Naval Research (N00014-21-C-1040).

**Acknowledgments.** Connor Henley was supported by a Draper Scholarship. This material is based upon work supported by the Office of Naval Research under Contract No. N00014-21-C-1040. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Office of Naval Research.

**Disclosures.** The authors declare no conflicts of interest.

**Data availability.** Data underlying the results presented in this paper are available at [18].

**Supplemental document.** See [Supplement 1](#) for supporting content.

## References

1. Apple, Inc., “Apple unveils new ipad pro with breakthrough lidar scanner and brings trackpad support to ipados,” <https://www.apple.com/newsroom/2020/03/apple-unveils-new-ipad-pro-with-lidar-scanner-and-trackpad-support-in-ipados/> (2020). Accessed: 2021-10-28.
2. A. Diosi and L. Kleeman, “Advanced sonar and laser range finder fusion for simultaneous localization and mapping,” in *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (IEEE Cat. No.04CH37566)*, vol. 2 (2004), pp. 1854–1859.
3. S.-W. Yang and C.-C. Wang, “On solving mirror reflection in lidar sensing,” *IEEE/ASME Trans. Mechatron.* **16**(2), 255–265 (2011).
4. P. Foster, Z. Sun, J. J. Park, and B. Kuipers, “Visagge: Visible angle grid for glass environments,” in *2013 IEEE International Conference on Robotics and Automation*, (2013), pp. 2213–2220.
5. H. Tibebu, J. Roche, V. De Silva, and A. Kondoz, “Lidar-based glass detection for improved occupancy grid mapping,” *Sensors* **21**(7), 2263 (2021).
6. R. Raskar and J. Davis, “5 d time-light transport matrix : What can we reason about scene properties?” Tech. rep., MIT, Cambridge, MA (2008).
7. C. Henley, J. Hollmann, and R. Raskar, “Bounce-flash lidar,” *IEEE Trans. Comput. Imaging* **8**, 411–424 (2022).
8. C. Tsai, A. Veeraraghavan, and A. C. Sankaranarayanan, “Shape and reflectance from two-bounce light transients,” in *2016 IEEE International Conference on Computational Photography (ICCP)*, (2016), pp. 1–10.
9. I. Ihrke, K. N. Kutulakos, H. P. A. Lensch, M. Magnor, and W. Heidrich, “Transparent and specular object reconstruction,” *Comput. Graph. Forum* **29**(8), 2400–2426 (2010).
10. A. Zisserman, P. Giblin, and A. Blake, “The information available to a moving observer from specularities,” *Image Vision Comput.* **7**(1), 38–42 (1989).
11. A. Blake and G. Brelstaff, “Geometry from specularities,” in *Proceedings of the 2nd International Conference on Computer Vision*, (IEEE, 1988), pp. 394–403.
12. S. Savarese, M. Chen, and P. Perona, “Local shape from mirror reflections,” *Int. J. Comput. Vis.* **64**(1), 31–67 (2005).
13. T. Whelan, M. Goesele, S. Lovegrove, J. Straub, S. Green, R. Szeliski, S. Butterfield, S. Verma, and R. Newcombe, “Reconstructing scenes with mirror and glass surfaces,” *ACM Trans. Graph.* **37**(4), 1–11 (2018).
14. T. Bonfort and P. Sturm, “Voxel carving for specular surfaces,” in *Proceedings of the Ninth International Conference on Computer Vision*, (IEEE, 2003).
15. S. Roth and M. Black, “Specular flow and the recovery of surface structure,” in *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’06)*, vol. 2 (2006), pp. 1869–1876.
16. M. A. Fischler and R. C. Bolles, “Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography,” *Commun. ACM* **24**(6), 381–395 (1981).
17. S. Royo and M. Ballesta-Garcia, “An overview of lidar imaging systems for autonomous vehicles,” *Appl. Sci.* **9**(19), 4093 (2019).
18. C. Henley, “Imaging glass and mirrors,” <https://github.com/co24401/Imaging-Glass-and-Mirrors> (2021).