

A Theory of Fermat Paths for Non-Line-of-Sight Shape Reconstruction

Shumian Xin¹, Sotiris Nousias^{2,3}, Kiriakos N. Kutulakos², Aswin C. Sankaranarayanan¹,
Srinivasa G. Narasimhan¹, and Ioannis Gkioulekas¹

¹Carnegie Mellon University ²University of Toronto ³University College London

Abstract

We present a novel theory of Fermat paths of light between a known visible scene and an unknown object not in the line of sight of a transient camera. These light paths either obey specular reflection or are reflected by the object's boundary, and hence encode the shape of the hidden object. We prove that Fermat paths correspond to discontinuities in the transient measurements. We then derive a novel constraint that relates the spatial derivatives of the path lengths at these discontinuities to the surface normal. Based on this theory, we present an algorithm, called Fermat Flow, to estimate the shape of the non-line-of-sight object. Our method allows, for the first time, accurate shape recovery of complex objects, ranging from diffuse to specular, that are hidden around the corner as well as hidden behind a diffuser. Finally, our approach is agnostic to the particular technology used for transient imaging. As such, we demonstrate mm-scale shape recovery from pico-second scale transients using a SPAD and ultrafast laser, as well as micron-scale reconstruction from femto-second scale transients using interferometry. We believe our work is a significant advance over the state-of-the-art in non-line-of-sight imaging.

1. Introduction

Most computer vision research assumes that the scene of interest is directly visible to the camera. In other words, the photons from a source that reach a camera are assumed to have interacted with only the visible scene. However, some of the source photons are reflected by the visible scene toward parts—say, the back of an object facing a camera, an object around a corner, or an object viewed through a diffuser—that are hidden from the direct line of sight of the camera. In turn, the hidden scene scatters the photons back toward the visible scene, which then redirects photons toward the camera. Imaging and understanding the scene hidden from the camera's view is of significant importance to many security and safety applications.

Capturing *non-line-of-sight* (NLOS) photons is chal-

lenging as they are vastly outnumbered by line-of-sight (LOS) photons. Passive approaches analyze the subtle umbra and penumbra of the shadow cast by the hidden scene to estimate rough motion and structure [6, 2, 38], or use coherence properties of light to localize hidden objects [5, 3]. These approaches do not have sufficient information to compute precise 3D shape of an unknown arbitrary hidden scene. Extracting additional information about the hidden scene is possible by using active illumination, including coherent lighting [42, 21, 4, 22] and steady-state intensity sources [25, 44, 51, 45]. The majority of approaches for reconstructing hidden *shape* information employ fast modulated light sources together with time-resolved sensors (e.g., continuous-wave ToF [15, 20], ultrafast photodiodes [24], streak cameras [49, 48, 13], and single-photon avalanche photodetectors (SPADs) [11, 32]). These sensors record not only the number of incident photons (intensity) but also their arrival times, at a range of temporal resolutions (milli- to femto-seconds) [49, 12, 31, 32, 11, 18]. Such measurements are called *transients* and the approach is called *transient NLOS imaging*.

By measuring transients at various locations of a known visible scene, most active techniques perform a volumetric 3D reconstruction by attempting to invert the time-resolved radiometric image formation process. Examples include elliptic backprojection [48, 7, 1, 26, 36], regularized linear system approaches [13, 15, 14, 20], the light-cone transform [33], and analysis-by-synthesis using rendering [35, 47]. These methods have two fundamental disadvantages: (1) they rely on radiometric information and existing SPADs produce poor intensity estimates due to effects such as pile-up and after-pulsing [16], as well as due to extreme sensitivity to photon noise and ambient lighting; and (2) to simplify the inverse problem, all existing reconstruction techniques rely on an assumption of Lambertian reflectance for the NLOS object.

In this paper, we overcome the above limitations by developing techniques that use only *geometric*, rather than intensity, constraints derived from transient measurements of an NLOS scene. For this, we present a new theory of NLOS photons that follow specific geometric paths, called *Fermat*




Figure 1: **Non-line-of-sight imaging.** We consider the problem of reconstructing surfaces that are: (a) outside the field of view of sensor, or (b) occluded from it by a diffuser. We develop an algorithm that can use transient imaging measurements to accurately reconstruct the shape of the non-line-of-sight surface. The figure shows example reconstructions of a US quarter from measurements captured by a femtosecond-scale transient imaging system. In (c), we compare our reconstructions against groundtruth, obtained using a direct depth scan of the object with the same transient imaging system.

paths between the LOS and NLOS scene. Based on Fermat’s principle [43], we observe that these paths follow either the law of specular reflection or reflect at specific points at the object’s boundary. We then prove that Fermat paths correspond to *discontinuities* in the transient measurements. The temporal locations of the discontinuities are a function of only the shape of the NLOS object and not its reflectance (BRDF). We additionally show that the shape of the transient around the discontinuity is related to the curvature of the hidden surface. This theory generalizes previous work on the paths of first-returning photons [46], which are a special case of Fermat paths.

We use the above theory to derive an algorithm, called *Fermat flow*, for accurate NLOS shape reconstruction. We show that the spatial derivative of the Fermat pathlength provides a simple constraint that uniquely determines both the depth and normal of a hidden scene point. This derivative is estimated numerically by fitting a smooth pathlength function to a sparser set of measurements. We then apply a final refinement step that computes a smooth mesh by combining both the depth and normal information [23, 9]. While most previous approaches reconstruct an albedo volume of the NLOS object, our approach is one of the few that reconstruct its surface. Compared to alternative surface reconstruction algorithms based on analysis-by-synthesis from intensity measurements [47], our approach uses only geometric constraints, which makes it BRDF-invariant and robust to imperfections in intensity measurements.

Our theory is agnostic to the specific transient imaging technology used. We validate our theory and demonstrate results at both pico-second and femto-second temporal scales, using a pulsed laser and SPAD for the former and interferometry for the latter. Hence, for the first time, we are able to compute millimeter-scale and micrometer-scale NLOS shapes of curved objects with BRDFs ranging from purely diffuse to purely specular. In addition, our theory ap-

plies to both reflective NLOS (looking around the corner) and transmissive NLOS (seeing through a diffuser) scenarios. Figure 1 shows the estimated micrometer-scale relief of a coin seen around the corner as well as through thick paper (diffuser). The obtained height profiles compare well with the reconstruction of the coin when imaged in the *line of sight*. This result demonstrates the significant theoretical and practical contribution of this work to active NLOS imaging, pushing the boundary of what is possible.

2. Fermat Paths in NLOS Transients

Problem setup. We consider a *transient imaging* system [19], comprising a light source and detector, located at points $s, d \in \mathbb{R}^3$, respectively. Our theory is agnostic to the specific transient imaging technology used, and in Section 4 we describe implementations, one based on a pulsed laser and a picosecond detector, and another based on interferometry. The *visible* scene $\mathcal{V} \subset \mathbb{R}^3$ is the union of surfaces contained within the common line of sight of the source and detector. In addition to \mathcal{V} , we assume that there exist surfaces outside their line of sight; this could be because either these surfaces are outside the field of view, or they are inside it but occluded by another surface. We are only interested in such surfaces that can *indirectly* receive light from the light source by means of a single reflection or transmission through the visible scene, and can *indirectly* send light to the detector in a likewise manner. We call the union of such surfaces the *non-line-of-sight* (NLOS) scene \mathcal{X} . Some situations where these conditions apply, and which will be relevant to our experiments, are shown in Figure 1.

We assume that the light source and detector are illuminating and imaging the same *visible point* $v \in \mathcal{V}$, which can be any point in the visible scene. This corresponds to the *confocal* scanning scheme, proposed by O’Toole et al. [33]. We emphasize that this assumption is only to simplify exposition: All of our theory generalizes to the non-confocal

case, as we briefly discuss in various places throughout the paper, and detail in the supplement. In particular, in Section 4, we show results from non-confocal experiments.

The detector records a *transient* $I(t; \mathbf{v})$, which equals the irradiance from photons with time of flight t . We assume that all recorded photons follow paths of the form $\mathbf{s} \rightarrow \mathbf{v} \rightarrow \mathbf{x} \rightarrow \mathbf{v} \rightarrow \mathbf{d}$, where $\mathbf{x} \in \mathcal{X}$. This *three-bounce* assumption is commonplace in NLOS imaging applications, for two reasons: First, NLOS transient imaging systems typically have time-gating mechanisms that can be used to remove *direct* photons that only interact with the visible scene. Second, photons with more than one interactions with the NLOS scene \mathcal{X} have greatly reduced signal-to-noise ratio, and in practice are difficult to detect [7, 35].

Finally, we assume that we have calibrated the distance $\tau_{\mathcal{V}}(\mathbf{v}) \triangleq \|\mathbf{s} - \mathbf{v}\| + \|\mathbf{d} - \mathbf{v}\|$ from the source to the visible point, and from there to the detector. Then, we can use the pathlength traveled in \mathcal{X} , $\tau \triangleq ct - \tau_{\mathcal{V}}(\mathbf{v})$ where c is the speed of light, to uniquely reparameterize transients as $I(\tau; \mathbf{v})$. Under these assumptions, we can write [10, 37]:

$$I(\tau; \mathbf{v}) = \int_{\mathcal{X}} f(\mathbf{x}; \mathbf{v}) \delta(\tau - \tau(\mathbf{x}; \mathbf{v})) dA(p, q) \quad (1)$$

where $\tau(\mathbf{x}; \mathbf{v}) \triangleq 2 \cdot \|\mathbf{x} - \mathbf{v}\|$, $(p, q) \in [0, 1]^2$ is a parameterization of the NLOS surface \mathcal{X} , $A(p, q)$ is the corresponding area measure, and the *throughput* f absorbs inverse-square fall-off, shading, reflectance, and visibility.

2.1. Fermat paths

We assume that the NLOS scene \mathcal{X} is formed as the union of *smooth* surfaces, and $\partial\mathcal{X} \subset \mathcal{X}$ is the set of points on the NLOS surface where a surface normal is not defined. We will be referring to $\partial\mathcal{X}$ as the *boundary* of \mathcal{X} for simplicity, but note that, in addition to boundary points, it includes points at discontinuous intersections of the smooth surfaces that make up \mathcal{X} . Then, we will be focusing on specific distinguished points $\mathbf{x} \in \mathcal{X}$ as follows.

Definition 1. *For any visible point \mathbf{v} :*

- The specular set $\mathcal{S}(\mathbf{v}) \subset \mathcal{X}$ consists of all points $\mathbf{x} \in \mathcal{X} \setminus \partial\mathcal{X}$ such that the vector $\mathbf{v} - \mathbf{x}$ is orthogonal to the tangent plane $T_{\mathbf{x}}\mathcal{X}$ of \mathcal{X} at \mathbf{x} .
- The boundary set $\mathcal{B}(\mathbf{v}) \subset \partial\mathcal{X}$ consists of all points $\mathbf{x} \in \partial\mathcal{X}$ such that the vector $\mathbf{v} - \mathbf{x}$ is orthogonal to the tangent vector $\hat{\mathbf{t}}(\mathbf{x})$ of $\partial\mathcal{X}$ at \mathbf{x} .
- The Fermat set $\mathcal{F}(\mathbf{v}) \subset \mathcal{X}$ is the union of these two sets, $\mathcal{F}(\mathbf{v}) \triangleq \mathcal{S}(\mathbf{v}) \cup \mathcal{B}(\mathbf{v})$.

Definition 1 implies that, at points $\mathbf{x} \in \mathcal{S}(\mathbf{v})$, the vector $\mathbf{v} - \mathbf{x}$ is also parallel to the surface normal $\hat{\mathbf{n}}(\mathbf{x})$. Equivalently, the path $\mathbf{p}(\mathbf{x}; \mathbf{v}) \triangleq \mathbf{v} \rightarrow \mathbf{x} \rightarrow \mathbf{v}$ corresponds to a

specular reflection at \mathbf{x} , explaining the name *specular set*. The name *Fermat set* is due to the following classical proposition of geometric optics [17, 8, 29, 43].

Proposition 2. *Let $(p, q) \in [0, 1]^2$ be a parameterization of the NLOS surface \mathcal{X} . Then, for any visible point \mathbf{v} ,*

$$\mathcal{S}(\mathbf{v}) = \{\mathbf{x} \in \mathcal{X} : \nabla_{(p,q)}\tau(\mathbf{x}(p, q); \mathbf{v}) = \mathbf{0}\}. \quad (2)$$

Let $r \in [0, 1]$ be a parameterization of the NLOS surface boundary $\partial\mathcal{X}$. Then, for any visible point \mathbf{v} ,

$$\mathcal{B}(\mathbf{v}) = \{\mathbf{x} \in \partial\mathcal{X} : \partial\tau(\mathbf{x}(r); \mathbf{v})/\partial r = 0\}. \quad (3)$$

For completeness, we provide a proof in the supplement. Proposition 2 is known as *Fermat's principle*, and characterizes paths of stationary length with respect to their local variations. We note that, even though Fermat's principle is often described as the "shortest path principle", it allows for longest or saddle-point light paths, as we demonstrate later in this section. Depending on \mathcal{X} , $\mathcal{F}(\mathbf{v})$ will have at least one, and potentially multiple, points, as shown in Figure 2.

We will associate each point $\mathbf{x} \in \mathcal{F}(\mathbf{v})$ with the sphere $\text{Sph}(\tau(\mathbf{x}; \mathbf{v})/2; \mathbf{v})$ of center \mathbf{v} and radius $\tau(\mathbf{x}; \mathbf{v})/2$. We call this the *tangent sphere*, because Proposition 2 implies that, for $\mathbf{x} \in \mathcal{S}(\mathbf{v})$ or $\mathbf{x} \in \mathcal{B}(\mathbf{v})$, the $\text{Sph}(\tau(\mathbf{x}; \mathbf{v})/2; \mathbf{v})$ is tangent to \mathcal{X} or $\partial\mathcal{X}$, respectively, at \mathbf{x} [29, 43].

Relationship to first-returning photons. Fermat paths are a superset of the paths of the first-returning photons described by Tsai et al. [46]. In particular, the pathlength of the first-returning photon is the *global* minimum of $\tau(\mathbf{x}; \mathbf{v})$. Then, Proposition 2 implies that $\mathbf{x} \in \mathcal{F}(\mathbf{v})$. Observations 2 and 3 of Tsai et al. [46], which make an assumption of local smoothness, correspond to the case where additionally $\mathbf{x} \in \mathcal{S}(\mathbf{v})$: Observation 3 describes the specular path $\mathbf{p}(\mathbf{x}; \mathbf{v})$, and Observation 2 describes the tangent sphere.

2.2. Fermat pathlengths as transient discontinuities

Except when the BRDF of the \mathcal{X} surface is perfectly specular, transients $I(\tau; \mathbf{v})$ will include contributions from photons that follow both Fermat and non-Fermat paths $\mathbf{p}(\mathbf{x}; \mathbf{v})$. Without prior knowledge of the scene, it would seem impossible to identify parts of the transient due to Fermat paths. However, we make the following observation.

Proposition 3. *Assume that the BRDF of the \mathcal{X} surface is non-zero in the specular direction. Then, for all $\mathbf{x} \in \mathcal{F}(\mathbf{v})$, the transient $I(\tau; \mathbf{v})$ will have a discontinuity at pathlength $\tau(\mathbf{x}; \mathbf{v})$. If $\mathbf{x} \in \mathcal{S}(\mathbf{v})$, then $I(\tau; \mathbf{v})$ will additionally have a vertical asymptote at $\tau(\mathbf{x}; \mathbf{v})$.*

Proof sketch. We sketch a proof for the specular case, and provide the full proof in the supplement. Let $\text{Sph}(\rho; \mathbf{v})$ be the sphere of center \mathbf{v} and radius ρ . Let the curve $\mathcal{C}(\rho; \mathbf{v})$ be the intersection of $\text{Sph}(\rho; \mathbf{v})$ with \mathcal{X} , parameterized by




Figure 2: **Theory of Fermat paths.** We illuminate and image an NLOS surface \mathcal{X} from a point v on a visible surface \mathcal{V} . (We show the camera and light sources in Figure 1.) Among all points $x \in \mathcal{X}$, some points on the surface ($x_{F,2}, x_{F,3}$) and boundary ($x_{F,1}$) will create paths that satisfy Fermat’s principle, corresponding to local minima ($x_{F,1}, x_{F,2}$) or maxima ($x_{F,3}$) of the pathlength function $\tau(x; v)$ (bottom right). The paths for the non-boundary points ($x_{F,2}, x_{F,3}$) will additionally be specular. We can identify the lengths of these *Fermat paths* from the fact that the transient $I(\tau; v)$ (bottom left) will be discontinuous at the corresponding pathlengths ($\tau_{F,1}, \tau_{F,2}, \tau_{F,3}$).

$r \in [0, 1]$. Then, we can use $(r, \rho) \in [0, 1] \times [0, \infty)$ to reparameterize \mathcal{X} , and rewrite the integral of Equation (1):

$$I(\tau; v) = \int_{\mathcal{X}} f(x; v) \delta(\tau - \tau(x; v)) \left| \mathcal{J}_{(p,q)}^{(r,\rho)}(x) \right|^{-1} dA(r, \rho), \quad (4)$$

where $\mathcal{J}_{(p,q)}^{(r,\rho)}(x)$ is the Jacobian of the transformation $(p, q) \mapsto (r, \rho)$. We now consider a point $x_S \in \mathcal{S}(v)$. Recognizing that $\rho(x_S) = \tau(x_S; v)/2$, we have from Equation (2) that $\nabla_{(p,q)} \rho(x_S) = \mathbf{0}$. Consequently,

$$\left| \mathcal{J}_{(p,q)}^{(r,\rho)}(x_S) \right| = \frac{\partial \rho(x_S)}{\partial p} \frac{\partial r(x_S)}{\partial q} - \frac{\partial \rho(x_S)}{\partial q} \frac{\partial r(x_S)}{\partial p} = 0. \quad (5)$$

Then, from Equation (4), at $\tau = \tau(x_S; v)$, the transient converges to infinity, resulting in a discontinuity. \square

Figure 2 visualizes this proposition for a *two-dimensional* Lambertian scene \mathcal{X} , and a visible point v such that $\mathcal{S}(v) = \{x_{F,2}, x_{F,3}\}$, $\mathcal{B}(v) = \{x_{F,1}\}$. We note that, in two dimensions, the boundary $\partial\mathcal{X}$ is not a curve but just

isolated points, and therefore the tangency property of Definition 1 and the tangent sphere are not meaningful.

BRDF invariance. Proposition 3 implies that the pathlengths where the transient $I(\tau; v)$ is discontinuous are determined completely by the function $\tau(x; v)$. In turn, $\tau(x; v)$ depends only on the *geometry* of v and \mathcal{X} . Therefore, the discontinuity pathlengths are independent of the BRDF of the NLOS surface \mathcal{X} . The BRDF is included in the throughput term f in Equation (4), and thus only affects the intensity of the transient at the discontinuity pathlength. Figure 3 demonstrates this reflectance invariance property.

Identifying type of stationarity. Proposition 3 allows us to identify the lengths of all Fermat paths that contribute to a transient $I(\tau; v)$, as the pathlengths where $I(\tau; v)$ is discontinuous. From Proposition 2, each of these pathlengths is a stationary point of the function $\tau(x; v)$. When the BRDF of \mathcal{X} is not perfectly specular, we can additionally identify the type of stationarity from the shape of the transient at a neighborhood of the discontinuity. We use Figure 2 for intuition, and refer to the supplement for details.

Specifically, let $\tau_{\mathcal{F}}$ be a pathlength where the transient is discontinuous. If $\tau_{\mathcal{F}}$ is a local maximum, the discontinuity in the transient $I(\tau; v)$ occurs at the limit from the left, $\tau \rightarrow \tau_{\mathcal{F}}^-$, and the transient decreases to the right of $\tau_{\mathcal{F}}$ (Figure 2, $\tau_{F,3}$). Conversely, when $\tau_{\mathcal{F}}$ is a local minimum, the discontinuity occurs at the limit from the right, $\tau \rightarrow \tau_{\mathcal{F}}^+$ (Figure 2, $\tau_{F,1}, \tau_{F,2}$). Finally, when $\tau_{\mathcal{F}}$ is a saddle point, the discontinuity and intensity rise are two-sided. An example of this is shown in Figure 3 (paraboloid case).

Identifying the stationarity type of *specular* discontinuities provides us with curvature information about \mathcal{X} .

Proposition 4. *Let a transient $I(\tau; v)$ have a specular discontinuity at τ_S , corresponding to a point $x_S \in \mathcal{S}(v)$. If $\kappa_{min}, \kappa_{max}$ are the principal curvatures of \mathcal{X} at x_S , then:*

- If τ_S is a local minimum of $\tau(x; v)$, $2/\tau_S < \kappa_{min}$.
- If τ_S is a local maximum of $\tau(x; v)$, $\kappa_{max} < 2/\tau_S$.
- If τ_S is a saddle point of $\tau(x; v)$, $\kappa_{min} \leq 2/\tau_S \leq \kappa_{max}$.

We provide the proof in the supplement, but we can use Figure 2 to provide intuition: The pathlength $\tau_{F,2}$ of point $x_{F,2}$ is a local minimum. All \mathcal{X} points in the neighborhood of $x_{F,2}$ are at a distance from v greater than $\tau_{F,2}$, and therefore outside the tangent sphere $\text{Sph}(\tau_{F,2}/2; v)$. This implies that the (minimal, in 3D) principal radius of curvature is greater than $\tau_{F,2}/2$. And conversely for the pathlength $\tau_{F,3}$ of point $x_{F,3}$, which is a local maximum. We note that a *sufficient* condition for \mathcal{X} to produce only locally-minimum specular pathlengths is that \mathcal{X} is convex. However, this is not a *necessary* condition: As explained in Proposition 4, it is possible for \mathcal{X} to contain concavities and still produce only locally-minimum specular pathlengths.




Figure 5: **Reconstruction pipeline.** (a) We first collect transient measurements $I(\tau, \mathbf{v})$ at multiple points \mathbf{v} on the visible surface \mathcal{V} . (b) For each measured transient, we detect pathlengths where the transient is discontinuous. These correspond to samples of the multi-valued Fermat pathlength function $\tau_{\mathcal{F}}(\mathbf{v})$. In the example shown, $\tau_{\mathcal{F}}(\mathbf{v})$ has two branches $\tau_{\mathcal{F},1}(\mathbf{v})$ and $\tau_{\mathcal{F},2}(\mathbf{v})$, shown in blue and red respectively. (c) Within each branch of $\tau_{\mathcal{F}}(\mathbf{v})$, we interpolate to compute the gradient $\nabla_{\mathbf{v}}\tau_{\mathcal{F}}(\mathbf{v})$. (d) Finally, by applying the Fermat flow equation (7), we reconstruct from each branch a set of points, either on the boundary (branch $\tau_{\mathcal{F},1}(\mathbf{v})$, blue) or at the interior (branch $\tau_{\mathcal{F},2}(\mathbf{v})$, red) of the NLOS shape.

We first define the *Fermat pathlength function* $\tau_{\mathcal{F}}(\mathbf{v})$:

$$\tau_{\mathcal{F}}(\mathbf{v}) = \{\tau : I(\tau; \mathbf{v}) \text{ is discontinuous}\}. \quad (6)$$

As each transient $I(\tau; \mathbf{v})$ can be discontinuous at more than one pathlengths, $\tau_{\mathcal{F}}(\mathbf{v})$ is a *multi-valued* function. We now prove the following property for this function.

Proposition 5. *Consider a branch of the Fermat pathlength function $\tau_{\mathcal{F}}(\mathbf{v})$ evaluated at $\mathbf{v} \in \mathcal{V}$. Assume that there is a unique point $\mathbf{x}_{\mathcal{F}} \in \mathcal{F}(\mathbf{v})$ with $\tau(\mathbf{x}_{\mathcal{F}}; \mathbf{v}) = \tau_{\mathcal{F}}(\mathbf{v})$. Then,*

$$\nabla_{\mathbf{v}}\tau_{\mathcal{F}}(\mathbf{v}) = -2 \frac{\mathbf{x}_{\mathcal{F}} - \mathbf{v}}{\|\mathbf{x}_{\mathcal{F}} - \mathbf{v}\|}. \quad (7)$$

We provide the proof in the supplement. It is instructive to consider the case of the branch of $\tau_{\mathcal{F}}(\mathbf{v})$ corresponding to the *global minimum* of $\tau(\mathbf{x}; \mathbf{v})$. This branch is the pathlength of the first-returning photons of Tsai et al. [46], and is also equal to twice the *distance function* $\mathcal{D}(\mathbf{v}, \mathcal{X})$, where

$$\mathcal{D}(\mathbf{v}, \mathcal{X}) \triangleq \min_{\mathbf{x} \in \mathcal{X}} \|\mathbf{x} - \mathbf{v}\|. \quad (8)$$

Then, Equation (7) is equivalent to the well-known *eikonal equation* for distance functions [30, 40, 39]. In this context, the uniqueness requirement of Proposition 5 is equivalent to a requirement that the minimizer of Equation (8) be unique, which is also a requirement for the eikonal equation. Proposition 5 generalizes the eikonal equation to apply to all branches of $\tau_{\mathcal{F}}(\mathbf{v})$, corresponding to all stationary points of $\tau(\mathbf{x}; \mathbf{v})$ and not only those for first-returning photons.

Using $\tau_{\mathcal{F}}(\mathbf{v}) = 2\|\mathbf{x}_{\mathcal{F}} - \mathbf{v}\|$, we rewrite Equation (7) as

$$\mathbf{x}_{\mathcal{F}} = \mathbf{v} - (\tau_{\mathcal{F}}(\mathbf{v})/4)\nabla_{\mathbf{v}}\tau_{\mathcal{F}}(\mathbf{v}). \quad (9)$$

In this form, the Fermat flow equation states that an NLOS point $\mathbf{x}_{\mathcal{F}} \in \mathcal{F}(\mathbf{v})$ generating a Fermat path can be uniquely reconstructed from the corresponding visible point \mathbf{v} , the length $\tau_{\mathcal{F}}(\mathbf{v})$, and the gradient $\nabla_{\mathbf{v}}\tau_{\mathcal{F}}(\mathbf{v})$. This reconstruction can be done with a simple geometric operation,

by intersecting the sphere $\text{Sph}(\tau_{\mathcal{F}}(\mathbf{v})/2; \mathbf{v})$ with the line $\mathbf{v} - \lambda\nabla_{\mathbf{v}}\tau_{\mathcal{F}}(\mathbf{v})/2$. If the Fermat path is also specular, $\mathbf{x}_{\mathcal{F}} \in \mathcal{S}(\mathbf{v})$, then we can also reconstruct the normal at $\mathbf{x}_{\mathcal{F}}$ as $\hat{\mathbf{n}}(\mathbf{x}_{\mathcal{F}}) = -\nabla_{\mathbf{v}}\tau_{\mathcal{F}}(\mathbf{v})/2$. This is shown in Figure 4.

Non-confocal case. We describe here a similar proposition for the non-confocal case, which we prove in the supplement. In this case, the Fermat pathlength function has two arguments, $\tau_{\mathcal{F}}(\mathbf{v}_s, \mathbf{v}_d)$. Then, as shown in Figure 4,

$$\nabla_{\mathbf{v}_i}\tau_{\mathcal{F}}(\mathbf{v}_s, \mathbf{v}_d) = -\frac{\mathbf{x}_{\mathcal{F}} - \mathbf{v}_i}{\|\mathbf{x}_{\mathcal{F}} - \mathbf{v}_i\|}, \quad (10)$$

where \mathbf{v}_i can be either of the two visible points \mathbf{v}_s and \mathbf{v}_d .

Gradient estimation. Using Equation (9) requires knowing the gradient $\nabla_{\mathbf{v}}\tau_{\mathcal{F}}(\mathbf{v})$, computed in the local coordinate frame of the visible surface \mathcal{V} at \mathbf{v} . We cannot measure this gradient directly, but we can estimate it through interpolation. For simplicity, we discuss here only the case when the visible surface \mathcal{V} is planar, deferring the general case for the supplement. When the visible surface \mathcal{V} is planar, the local coordinate system at \mathbf{v} can be assumed to be the $x-y$ plane of the global coordinate system. Given a Fermat pathlength $\tau_{\mathcal{F}}(\mathbf{v})$ at a point \mathbf{v} , we can estimate its partial derivatives $\partial\tau_{\mathcal{F}}/\partial x$ and $\partial\tau_{\mathcal{F}}/\partial y$ by locally interpolating Fermat pathlengths from transients measured at nearby points on the plane. We can then infer the derivative with respect to z by noting that Equation (7) implies that $\|\nabla_{\mathbf{v}}\tau_{\mathcal{F}}(\mathbf{v})\| = 2$,

$$\nabla_{\mathbf{v}}\tau_{\mathcal{F}}(\mathbf{v}) = \left(\frac{\partial\tau_{\mathcal{F}}}{\partial x}, \frac{\partial\tau_{\mathcal{F}}}{\partial y}, \sqrt{4 - \left(\frac{\partial\tau_{\mathcal{F}}}{\partial x}\right)^2 - \left(\frac{\partial\tau_{\mathcal{F}}}{\partial y}\right)^2} \right) \Bigg|_{\mathbf{v}}. \quad (11)$$

We note that Equation (11) fails when the uniqueness condition of Proposition 5 is not satisfied. In practice, this can only happen at isolated points \mathbf{v} on the visible surface, corresponding to non-generic symmetries of the NLOS surface \mathcal{X} . Additionally, we note that the above interpolation procedure needs to be performed separately for each branch of the Fermat pathlength function $\tau_{\mathcal{F}}(\mathbf{v})$.




Figure 6: **Comparison with groundtruth.** We perform one-dimensional scans of 3D-printed objects, in a looking-around-the-corner configuration. For each object, we show a photograph under ambient light (left), and reconstruction results (red points) superimposed against the groundtruth mesh used to fabricate the object (middle and right).

Surface fitting. The above procedure produces an *oriented* point cloud, of density comparable to the density of measurements on \mathcal{V} . We can then use algorithms that take advantage of normal information to fit a surface representation (e.g., triangular mesh) to the point cloud with increased accuracy [23]. Given such an initial surface reconstruction, in the supplement we describe an optimization procedure, based on the theory of specular path perturbations [9, 17], that refines the fitted surface to account for possible errors due to inaccurate estimation of the gradients $\nabla_{\mathbf{v}}\tau_{\mathcal{F}}(\mathbf{v})$.

4. Experiments

We discuss results from NLOS reconstruction experiments we have performed to validate and evaluate the Fermat flow algorithm. All experiments are based on measurements captured with two transient imaging setups, one operating at picosecond and the other at femtosecond temporal scales. We show additional experiments in the supplement.

4.1. Picosecond-scale experiments

Imaging system. We use a SPAD-based transient imaging system [11, 33, 31], consisting of a picosecond laser (NKT SuperK EXW-12), a SPAD detector (MPD module), and a time-correlated single photon counter (TCSPC, PicoQuant PicoHarp). The temporal binning resolution of the TCSPC unit is 4 ps, for an absolute upper bound in depth resolution of 1.2 mm. In practice, the resolution is lower, because of laser and TCSPC jitter. We use galvo mirrors to independently control viewpoint and illumination direction, and perform both confocal and non-confocal scanning in the looking-around-the-corner setting of Figure 1(a).




Figure 7: **Comparison of point cloud and surface reconstructions.** We scan a rough specular kettle, shown in the left under ambient light. We reconstruct an oriented point cloud, shown in the middle from two views, where points are colored according to their normal. Finally, we fit a surface to the point cloud, shown to the right under two views.

Comparison with groundtruth. We fabricated small objects from CAD meshes, providing us with ground-truth shape for comparison. The objects were painted with matte white paint to create Lambertian reflectance. The objects are approximately 15 cm in each dimension, and are placed at a distance 25 cm from a planar visible surface. All objects are ruled surfaces, to allow reconstruction of their profile from only one-dimensional scans. We capture measurements under a non-confocal setting, by fixing the point imaged by the SPAD and scanning the point illuminated by the source along multiple horizontal lines on the visible wall. Along each line, we scan 200 points, at a distance of approximately 1 mm from each other. Figure 6 shows point clouds reconstructed from these measurements using the Fermat flow procedure, superimposed against the meshes used for fabrication. The reconstructions closely reproduce the shape of the objects, including their concave and convex surfaces, and match the groundtruth within 2 mm.

Table-top objects. We scanned a variety of every day objects (Figures 7-8), with convex and concave geometry of different BRDFs, including translucent (plastic jug), glossy (bowl, vase), rough specular (kettle) and smooth specular (sphere). Most of the objects have a major dimension of approximately 20 – 30 cm, and are placed at a distance of 80 cm from the visible wall. We use confocal scanning with a grid of 64×64 points distributed in an area of $80 \text{ cm} \times 80 \text{ cm}$ on the visible wall.

Figure 7 visualizes point cloud, normal, and final surface reconstruction for one of the objects, an electrical kettle with rough-specular reflectance. We observe that our reconstruction procedure produces a point cloud that closely matches the shape of the object, including accurate normals on its front surface. We note that we do not reconstruct normals at the handle of the object: This is expected, because these parts of the object produce Fermat paths of boundary, rather than specular, type, and such paths do not provide normal information. The final fitted surface further improves the reconstruction quality. Figure 8 shows recon-




Figure 8: **Table-top objects.** We scan objects that span a variety of shapes (convex, concave) and reflectances (translucent, glossy, specular). For each object, we show a photograph under ambient light, and two views of its surface reconstruction.

structions for the remaining table-top objects. In all cases, the reconstruction closely matches the object shape, demonstrating the ability of our algorithm to handle a variety of complex geometry and reflectance combinations.

4.2. Femtosecond-scale experiments

Imaging system. We use a time-domain, full-frame optical coherent tomography system [12]. We use this system to perform confocal scans under both the looking-around-the-corner and looking-through-diffuser settings (Figure 1). We use spatially and temporally incoherent LED illumination, which allows us to combine transient imaging with diagonal probing [34]. In the context of confocal scanning, this means that we can simultaneously collect transients $I(\tau, \mathbf{v})$ at all points on the visible surface without scanning, as transient measurements taken at one point will not be contaminated with light emanating from a different point. Our implementation has depth resolution of $10 \mu\text{m}$.

Coin reconstructions. We perform experiments in both the looking-around-the-corner and looking-through-diffuser settings (Figure 1), where for the diffuser we use a thin sheet of paper. In both cases, the NLOS object is a US quarter, with the obverse side facing the visible surface. We place the coin at a distance of 10 mm from the visible surface, and collect transient measurements on an area of about $40 \text{ mm} \times 40 \text{ mm}$, at an 1 MPixel grid of points. For validation, we additionally use the same setup to directly scan the coin without occlusion. Figure 1 shows our results. In both cases, we can reconstruct fine detail on the coin, sufficient to infer its denomination. The reconstructed detail is also in close agreement with the groundtruth shape measured with the coin directly in the line-of-sight.

5. Discussion

We discuss some limitations of our approach. Our reconstruction procedure does not require radiometric calibration, as it does not use intensity information, instead relying on estimation of the pathlengths where measured transients are discontinuous. Consequently, our reconstructions can be sensitive to inaccurate discontinuity detection.

Reconstruction quality can additionally suffer if we do not have sufficiently dense measurements for estimating Fermat pathlength gradients through interpolation. Finally, using only pathlength information provides BRDF invariance, but it also means that we do not take advantage of information available in measured intensities about the NLOS scene.

The theory we developed offers new insights into the NLOS imaging problem, linking it to classical areas such as specular and differential geometry, and providing ample opportunity for transfer of ideas from these areas to the NLOS imaging setting. By allowing us to treat NLOS reconstruction from a purely geometric perspective, our theory introduces a new methodology for tackling this problem, distinct from but complementary to approaches such as (elliptic) backprojection [48, 33] and analysis-by-synthesis [47], which focus on the radiometric aspects of the problem.

Interestingly, concurrent work [28] has uncovered an intriguing link between our and backprojection approaches, by showing that the latter cannot reconstruct NLOS points not on Fermat paths, even if those points otherwise contribute to measured transients. Therefore, both approaches reproduce the same part of the NLOS scene. Further exploration of connections between the geometric and backprojection approaches can help shed light into their fundamental limits and strengths, potentially by allowing us to derive results applicable to both classes of approaches using whichever mathematical framework (geometric, radiometric) is more convenient for analysis.

More broadly, an exciting future direction of research is combining the two classes of approaches, not only for NLOS imaging, but also for other related applications, including acoustic and ultrasound imaging [27], lensless imaging [50], and seismic imaging [41].

Acknowledgments. We thank Chia-Yin Tsai for valuable discussions. This work was supported by the DARPA REVEAL program under contract HR0011-16-C-0025. SX, ACS, SGN, and IG were additionally supported by NSF Expeditions award CCF-1730147. KNK was supported by the NSERC RGPIN and RTI programs.

References

- [1] Victor Arellano, Diego Gutierrez, and Adrian Jarabo. Fast back-projection for non-line of sight reconstruction. *Optics Express*, 25(10):11574–11583, 2017. [1](#)
- [2] Manel Baradad, Vickie Ye, Adam B Yedidia, Frédo Durand, William T Freeman, Gregory W Wornell, and Antonio Torralba. Inferring light fields from shadows. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6267–6275, 2018. [1](#)
- [3] Mufeed Batarseh, Sergey V Sukhov, Zhiqin Shen, Heath Gemar, Reza Rezvani, and Aristide Dogariu. Passive sensing around the corner using spatial coherence. *Nature communications*, 9(1):3629, 2018. [1](#)
- [4] Jacopo Bertolotti, Elbert G van Putten, Christian Blum, Ad Lagendijk, Willem L Vos, and Allard P Mosk. Non-invasive imaging through opaque scattering layers. *Nature*, 491(7423):232, 2012. [1](#)
- [5] Jeremy Boger-Lombard and Ori Katz. Non line-of-sight localization by passive optical time-of-flight. *ArXiv e-prints*, Aug. 2018. [1](#)
- [6] Katherine L. Bouman, Vickie Ye, Adam B. Yedidia, Frédo Durand, Gregory W Wornell, Antonio Torralba, and William T Freeman. Turning corners into cameras: Principles and methods. In *ICCV*, 2017. [1](#)
- [7] Mauro Buttafava, Jessica Zeman, Alberto Tosi, Kevin Eliceiri, and Andreas Velten. Non-line-of-sight imaging using a time-gated single photon avalanche diode. *Optics express*, 23(16):20997–21011, 2015. [1](#), [3](#)
- [8] Min Chen and James Arvo. Perturbation methods for interactive specular reflections. *IEEE Transactions on Visualization and Computer Graphics*, 6(3):253–264, 2000. [3](#)
- [9] Min Chen and James Arvo. Theory and application of specular path perturbation. *ACM Transactions on Graphics (TOG)*, 19(4):246–278, 2000. [2](#), [7](#)
- [10] Philip Dutré, Kavita Bala, and Philippe Bekaert. *Advanced global illumination*. AK Peters, Ltd., 2006. [3](#)
- [11] Genevieve Garipey, Nikola Krstajić, Robert Henderson, Chunyong Li, Robert R Thomson, Gerald S Buller, Barmak Heshmat, Ramesh Raskar, Jonathan Leach, and Daniele Faccio. Single-photon sensitive light-in-flight imaging. *Nature communications*, 6:6021, 2015. [1](#), [7](#)
- [12] Ioannis Gkioulekas, Anat Levin, Frédo Durand, and Todd Zickler. Micron-scale light transport decomposition using interferometry. *ACM Transactions on Graphics*, 2015. [1](#), [8](#)
- [13] Otkrist Gupta, Thomas Willwacher, Andreas Velten, Ashok Veeraraghavan, and Ramesh Raskar. Reconstruction of hidden 3d shapes using diffuse reflections. *Optics express*, 20(17):19096–19108, 2012. [1](#)
- [14] Felix Heide, Matthew O’Toole, Kai Zhang, David Linnell, Steven Diamond, and Gordon Wetzstein. Robust non-line-of-sight imaging with single photon detectors. *arXiv preprint arXiv:1711.07134*, 2017. [1](#)
- [15] Felix Heide, Lei Xiao, Wolfgang Heidrich, and Matthias B Hullin. Diffuse mirrors: 3d reconstruction from diffuse indirect illumination using inexpensive time-of-flight sensors. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3222–3229, 2014. [1](#)
- [16] Quercus Hernandez, Diego Gutierrez, and Adrian Jarabo. A computational model of a single-photon avalanche diode sensor for transient imaging. *arXiv preprint arXiv:1703.02635*, 2017. [1](#), [5](#)
- [17] Wenzel Jakob and Steve Marschner. Manifold exploration: A markov chain monte carlo technique for rendering scenes with difficult specular transport. *ACM Transactions on Graphics*, 2012. [3](#), [7](#)
- [18] Adrian Jarabo, Julio Marco, Adolfo Muñoz, Raul Buisan, Wojciech Jarosz, and Diego Gutierrez. A Framework for Transient Rendering. *ACM Trans. Graph.*, 33(6):177:1–177:10, Nov. 2014. [1](#)
- [19] Adrian Jarabo, Belen Masia, Julio Marco, and Diego Gutierrez. Recent advances in transient imaging: A computer graphics and vision perspective. *Visual Informatics*, 1(1):65–79, 2017. [2](#)
- [20] Achuta Kadambi, Hang Zhao, Boxin Shi, and Ramesh Raskar. Occluded imaging with time-of-flight sensors. *ACM Transactions on Graphics (ToG)*, 35(2):15, 2016. [1](#)
- [21] Ori Katz, Pierre Heidmann, Mathias Fink, and Sylvain Gigan. Non-invasive single-shot imaging through scattering layers and around corners via speckle correlations. *Nature photonics*, 8(10):784, 2014. [1](#)
- [22] Ori Katz, Eran Small, and Yaron Silberberg. Looking around corners and through thin turbid layers in real time with scattered incoherent light. *Nature photonics*, 6(8):549–553, 2012. [1](#)
- [23] Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. Poisson surface reconstruction. In *Proc. Eurographics Symposium on Geometry Processing*, pages 61–70, 2006. [2](#), [7](#)
- [24] Ahmed Kirmani, Tyler Hutchison, James Davis, and Ramesh Raskar. Looking around the corner using ultrafast transient imaging. *International journal of computer vision*, 95(1):13–28, 2011. [1](#)
- [25] Jonathan Klein, Christoph Peters, Jaime Martín, Martin Laurenzis, and Matthias B Hullin. Tracking objects

- outside the line of sight using 2d intensity images. *Scientific reports*, 6, 2016. 1
- [26] Marco La Manna, Fiona Kine, Eric Breitbach, Jonathan Jackson, and Andreas Velten. Error back-projection algorithms for non-line-of-sight imaging. 1
- [27] David Lindell, Gordon Wetzstein, and Vladlen Koltun. Acoustic non-line-of-sight imaging. In *CVPR*, 2019. 8
- [28] Xiaochun Liu, Sebastian Bauer, and Andreas Velten. Analysis of feature visibility in non-line-of-sight measurements. In *CVPR*, 2019. 8
- [29] Don Mitchell and Pat Hanrahan. Illumination from curved reflectors. In *ACM SIGGRAPH Computer Graphics*, volume 26, pages 283–291. ACM, 1992. 3
- [30] Stanley Osher and Ronald Fedkiw. *Level set methods and dynamic implicit surfaces*, volume 153. Springer Science & Business Media, 2006. 6
- [31] Matthew O’Toole, Felix Heide, David B Lindell, Kai Zang, Steven Diamond, and Gordon Wetzstein. Reconstructing transient images from single-photon sensors. *CVPR*, 2017. 1, 7
- [32] Matthew O’Toole, Felix Heide, Kai Lindell, David B. and Zang, Steven Diamond, and Gordon Wetzstein. Reconstructing transient images from single-photon sensors. In *CVPR*, 2017. 1
- [33] Matthew O’Toole, David B Lindell, and Gordon Wetzstein. Confocal non-line-of-sight imaging based on the light-cone transform. *Nature*, 555(7696):338, 2018. 1, 2, 7, 8
- [34] M. O’Toole, R. Raskar, and K.N. Kutulakos. Primal-dual coding to probe light transport. *ACM Transactions on Graphics*, 2012. 8
- [35] Adithya Pediredla, Mauro Buttava, Alberto Tosi, Oliver Cossairt, and Ashok Veeraraghavan. Reconstructing rooms using photon echoes: A plane based model and reconstruction algorithm for looking around the corner. *ICCP*, 2017. 1, 3
- [36] Adithya Pediredla, Akshat Dave, and Ashok Veeraraghavan. Snlos: Non-line-of-sight scanning through temporal focusing. *ICCP*, 2019. 1
- [37] Adithya Pediredla, Ashok Veeraraghavan, and Ioannis Gkioulekas. Elliptic path sampling for time-gated rendering. *ACM Transactions on Graphics (TOG)*, 2019. 3
- [38] Charles Saunders, John Murray-Bruce, and Vivek K Goyal. Computational periscopy with an ordinary digital camera. *Nature*, 565(7740):472, 2019. 1
- [39] James A Sethian. Fast marching methods. *SIAM review*, 41(2):199–235, 1999. 6
- [40] James A Sethian. *Level set methods and fast marching methods: evolving interfaces in computational geometry, fluid mechanics, computer vision, and materials science*, volume 3. Cambridge university press, 1999. 6
- [41] James A Sethian and A Mihai Popovici. 3-d travelttime computation using the fast marching method. *Geophysics*, 64(2):516–523, 1999. 8
- [42] Brandon M. Smith, Matthew O’Toole, and Mohit Gupta. Tracking multiple objects outside the line of sight using speckle imaging. In *CVPR*, 2018. 1
- [43] Orestes Stavroudis. *The optics of rays, wavefronts, and caustics*. Elsevier, 1972. 2, 3
- [44] Matthew Tancik, Guy Satat, and Ramesh Raskar. Flash photography for data-driven hidden scene recovery. *arXiv preprint arXiv:1810.11710*, 2018. 1
- [45] Christos Thrampoulidis, Gal Shulkind, Feihu Xu, William T Freeman, Jeffrey Shapiro, Antonio Torralba, Franco Wong, and Gregory Wornell. Exploiting occlusion in non-line-of-sight active imaging. *IEEE Transactions on Computational Imaging*, 2018. 1
- [46] Chia-Yin Tsai, Kiriakos N Kutulakos, Srinivasa G Narasimhan, and Aswin C Sankaranarayanan. The geometry of first-returning photons for non-line-of-sight imaging. In *CVPR*, 2017. 2, 3, 6
- [47] Chia-Yin Tsai, Aswin C. Sankaranarayanan, and Ioannis Gkioulekas. Beyond volumetric albedo—a surface optimization framework for non-line-of-sight imaging. In *CVPR*, 2019. 1, 2, 8
- [48] Andreas Velten, Thomas Willwacher, Otkrist Gupta, Ashok Veeraraghavan, Mounqi G Bawendi, and Ramesh Raskar. Recovering three-dimensional shape around a corner using ultrafast time-of-flight imaging. 2012. 1, 8
- [49] Andreas Velten, Di Wu, Adrian Jarabo, Belen Masia, Christopher Barsi, Chinmaya Joshi, Everett Lawson, Mounqi Bawendi, Diego Gutierrez, and Ramesh Raskar. Femto-photography: Capturing and Visualizing the Propagation of Light. *ACM Trans. Graph.*, 32(4):44:1–44:8, July 2013. 1
- [50] Di Wu, Gordon Wetzstein, Christopher Barsi, Thomas Willwacher, Qionghai Dai, and Ramesh Raskar. Ultrafast lensless computational imaging through 5d frequency analysis of time-resolved light transport. *International journal of computer vision*, 110(2):128–140, 2014. 8
- [51] Feihu Xu, Gal Shulkind, Christos Thrampoulidis, Jeffrey H Shapiro, Antonio Torralba, Franco NC Wong, and Gregory W Wornell. Revealing hidden scenes by photon-efficient occlusion-based opportunistic active imaging. *Optics express*, 26(8):9945–9962, 2018. 1