

Virtual Mirrors: Non-Line-of-Sight Imaging Beyond the Third Bounce

DIEGO ROYO, Universidad de Zaragoza – I3A, Spain

TALHA SULTAN, University of Wisconsin–Madison, USA

ADOLFO MUÑOZ, Universidad de Zaragoza – I3A, Spain

KHADIJEH MASUMNIA-BISHEH, University of Wisconsin–Madison, USA

ERIC BRANDT, University of Wisconsin–Madison, USA

DIEGO GUTIERREZ, Universidad de Zaragoza – I3A, Spain

ANDREAS VELTEN, University of Wisconsin–Madison, USA

JULIO MARCO, Universidad de Zaragoza – I3A, Spain

Non-line-of-sight (NLOS) imaging methods are capable of reconstructing complex scenes that are not visible to an observer using indirect illumination. However, they assume only third-bounce illumination, so they are currently limited to single-corner configurations, and present limited visibility when imaging surfaces at certain orientations. To reason about and tackle these limitations, we make the key observation that planar diffuse surfaces behave specularly at wavelengths used in the computational wave-based NLOS imaging domain. We call such surfaces *virtual mirrors*. We leverage this observation to expand the capabilities of NLOS imaging using illumination beyond the third bounce, addressing two problems: imaging single-corner objects at limited visibility angles, and imaging objects hidden behind two corners. To image objects at limited visibility angles, we first analyze the reflections of the known illuminated point on surfaces of the scene as an estimator of the position and orientation of objects with limited visibility. We then image those limited visibility objects by computationally building secondary apertures at other surfaces that observe the target object from a direct visibility perspective. Beyond single-corner NLOS imaging, we exploit the specular behavior of virtual mirrors to image objects hidden behind a second corner by imaging the space behind such virtual mirrors, where the mirror image of objects hidden around two corners is formed. No specular surfaces were involved in the making of this paper.

CCS Concepts: • **Computing methodologies** → **3D imaging**; **Computational photography**.

Additional Key Words and Phrases: non-line-of-sight imaging, time-of-flight imaging, wave-based imaging, computational photography

ACM Reference Format:

Diego Royo, Talha Sultan, Adolfo Muñoz, Khadijeh Masumnia-Bisheh, Eric Brandt, Diego Gutierrez, Andreas Velten, and Julio Marco. 2023. Virtual Mirrors: Non-Line-of-Sight Imaging Beyond the Third Bounce. *ACM Trans. Graph.* 42, 4, Article 140 (August 2023), 15 pages. <https://doi.org/10.1145/3592429>

Authors' addresses: Diego Royo, droyo@unizar.es, Universidad de Zaragoza – I3A, Zaragoza, Spain; Talha Sultan, tsultan@wisc.edu, University of Wisconsin–Madison, Madison, WI, USA; Adolfo Muñoz, adolfo@unizar.es, Universidad de Zaragoza – I3A, Zaragoza, Spain; Khadijeh Masumnia-Bisheh, khadijeh.masumnia-bisheh@wisc.edu, University of Wisconsin–Madison, Madison, WI, USA; Eric Brandt, elbrandt@wisc.edu, University of Wisconsin–Madison, Madison, WI, USA; Diego Gutierrez, diego@unizar.es, Universidad de Zaragoza – I3A, Zaragoza, Spain; Andreas Velten, velten@wisc.edu, University of Wisconsin–Madison, Madison, WI, USA; Julio Marco, juliom@unizar.es, Universidad de Zaragoza – I3A, Zaragoza, Spain.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2023 Copyright held by the owner/author(s).

0730-0301/2023/8-ART140

<https://doi.org/10.1145/3592429>




Fig. 1. We image geometry around two corners. The displayed room only has diffuse surfaces. A single laser source illuminates the relay surface, then the resulting indirect illumination is captured and used as input to the problem. The main observation in our work shows that, by only applying computations to the captured data, surfaces that are diffuse in the real world can exhibit specular properties, based on well-known wave interference principles. We leverage this observation in the room example shown, using one of the diffuse surfaces as a *virtual mirror* to image a T-shaped geometry, that would otherwise not be visible by only looking around one corner.

1 INTRODUCTION

Non-line-of-sight (NLOS) imaging methods retrieve information of scenes that are hidden from the observer, including geometric reconstructions [Velten et al. 2012a; Wu et al. 2021], position detection [Bouman et al. 2017; Yi et al. 2021], or motion tracking [Garipey et al. 2016], with many applications in fields such as remote sensing, autonomous driving or biological imaging. Under this regime, methods that image scenes hidden around a corner have shown promising results thanks to ultra-fast imaging devices (e.g., [Buttafava et al. 2015; Shin et al. 2016; Velten et al. 2013]). Such time-gated NLOS imaging methods provide detailed reconstructions of hidden scenes by triangulating geometric positions using the time of flight of round-trip third-bounce illumination paths between a visible relay surface and the hidden scene [Lindell et al. 2019b; O’Toole et al. 2018; Velten et al. 2012a; Xin et al. 2019].

These methods operate under the assumption of third-bounce-only illumination, with higher-order illumination usually degrading the reconstructions due to ambiguities in the time of flight of light. Recent wave-based NLOS imaging methods have shown how such higher-order bounces can be isolated from third-bounce illumination and visualized, based on an imaging paradigm that interprets the time-resolved illumination captured at a relay surface as light arriving at a virtual aperture [Liu et al. 2020, 2019b; Marco et al.

2021; Nam et al. 2021], effectively transforming the relay surface into a virtual line-of-sight (LOS) imaging system.

In our work, we demonstrate how such higher-order bounces can be used to expand the capabilities of existing NLOS imaging systems, and overcome some of its current limitations. In particular, we draw a parallelism between Huygens' principle and the recent wave-based phasor-field NLOS imaging formulation [Liu et al. 2019b]. We intuitively show how, due to well-known wave interference principles, surfaces that are diffuse under visible light can behave like mirrors during the computational NLOS wave-based imaging process; we call such surfaces *virtual mirrors*. It is thus important to understand that virtual mirrors only show specular behavior in the computational domain, since real NLOS capture systems still receive diffuse illumination.

From this observation, we show how (in the computational domain) specular reflections in the form of *fourth-* and *fifth-bounce* illumination actually encode useful information about the scene, then leverage these higher-order bounces to address two longstanding problems in NLOS imaging: visibility issues due to the missing-cone problem, and looking around *two* corners.

The missing-cone problem is inherent to all third-bounce NLOS imaging methods, and refers to the fact that certain scene configurations and features cannot be accessed by NLOS measurements depending on their position and orientation with respect to the relay surface [Liu et al. 2019a]. The corresponding surfaces are said to be inside the null-reconstruction space of third-bounce methods.

To image surfaces inside such null-reconstruction space, we first image the mirror reflections of a known illuminated point on the relay surface, produced by all surfaces of the scene, including those surfaces inside the null-reconstruction space. By analyzing these reflections, we infer the position and orientation of the hidden surfaces inside the null-reconstruction space that produced such mirror reflections. To avoid ambiguities introduced by inference, we introduce a novel procedure to directly image the hidden surfaces by creating a *second* virtual aperture at other scene surfaces.

In addition, we propose a second novel procedure to image objects hidden behind *two* corners (Figure 1). For this, we show how to use fifth-bounce illumination to image the space behind a diffuse hidden surface, which effectively acts as a mirror as seen from the relay surface, allowing us to observe a mirror image of the object hidden behind two corners. Our key insight for this second procedure is selecting the location of the volume being imaged, which is in principle orthogonal to the particular imaging algorithm used. Given our virtual mirror surfaces, we target the reflected space *behind* such surfaces, where mirror images are formed (just like with real mirrors). We demonstrate our procedure works even if the intermediate surface itself falls inside the null-reconstruction space.

In summary, we demonstrate how to extend the capabilities of existing NLOS imaging methods, by i) imaging surfaces inside the null-reconstruction space by leveraging fourth-bounce illumination, and ii) imaging surfaces hidden around two corners using fifth-bounce illumination. We validate our findings both in simulation and using real captured data. Last, all of our experiment data, simulation and imaging software are publicly available¹.

¹https://graphics.unizar.es/projects/VirtualMirrors_2023

2 RELATED WORK

NLOS imaging methods analyze indirect illumination from paths that scatter one or multiple times in the target hidden scene to obtain information about it. They can be divided into active and passive methods. Active methods use controlled light sources to illuminate the hidden scene [Cao et al. 2022; Katz et al. 2014; Luesia et al. 2022; Velten et al. 2013], while passive methods rely on ambient illumination [Bouman et al. 2017; Krška et al. 2022] or light emitted by hidden objects themselves [Saunders et al. 2019]. In our work, an active laser source emits light pulses to illuminate the hidden scene.

Time-gated NLOS imaging. The ability to capture time-gated illumination using time-of-flight detectors at picosecond resolution has nurtured a wide range of NLOS imaging methods [Faccio et al. 2020; Jarabo et al. 2017; Maeda et al. 2019a; Pediredla et al. 2019]. The first methods to demonstrate high-quality 3D reconstructions employed ellipsoidal backprojection [Buttafava et al. 2015; Velten et al. 2012b] acquired for non-confocal optical paths. These methods operate on the time domain of the captured light transport and are computationally expensive. By restricting data acquisition to confocal optical light paths, NLOS reconstructions can be formulated as a closed-form, deconvolution-based linear inverse problem which can be solved efficiently in the frequency domain [Lindell et al. 2019b; O'Toole et al. 2018]. Since they are based on frequency-space deconvolutions, they require the use of a regular sampling grid on a planar relay surface. Furthermore, in contrast to our work, both ellipsoidal-based and deconvolution-based techniques are unable to account for light bounces beyond the third.

Phasor-field formulation. The phasor-field formulation [Liu et al. 2019b] provides wave-based models to propagate virtual waves into the hidden scene, which allows to turn a visible relay surface into a virtual LOS camera. This allows to build fundamental wave-optics parallelisms between forward operators in LOS imaging and the backprojection operators that drive time-gated NLOS imaging approaches [Dove and Shapiro 2020a,b,c; Guillén et al. 2020; Laurenzis and Christnacher 2022; Reza et al. 2019a,b; Teichman 2019]. The phasor-field formalism allows to image hidden scenes using either confocal or non-confocal setups [Liu et al. 2019b]. Subsequent implementations have gained efficiency by working in the frequency domain [Liu et al. 2020], which has led to interactive and real-time reconstructions of dynamic hidden scenes [Liao et al. 2021; Nam et al. 2020] at the cost of using a regular sampling grid. Moreover, it allows to use both planar and non-planar relay surfaces [La Manna et al. 2020], and to leverage known occlusions in the reconstructions [Dove and Shapiro 2019]. For memory-constrained applications, different propagation operators can be implemented, such as zone plates [Luesia-Lahoz et al. 2023]. Last, Marco et al. [2021] leveraged the phasor-field formulation to separate direct and indirect illumination of hidden scenes by combining exhaustive scans of both laser and sensor positions. In this work, we build on top of the virtual-wave LOS parallelisms and reason about the virtual reflectivity of hidden surfaces. We then leverage higher-order illumination to image geometry hidden around two corners, as well as to directly estimate hidden objects that have limited visibility to classic third-bounce methods.

NLOS with specular reflections. Prior works utilized actual specular reflections by using centimeter-scale acoustic waves for NLOS reconstructions [Lindell et al. 2019a], or millimeter-scale radio waves for tracking hidden objects [Scheiner et al. 2020]. Using specular reflections requires directly sampling the path from the hidden scene to the relay surface, needing strong assumptions about surface albedo and orientation [Lindell et al. 2019a], or placing the scanning system far from the relay surface [Scheiner et al. 2020]. Specular reflections produced by infrared wavelengths have also been utilized for passive NLOS [Kaga et al. 2019; Maeda et al. 2019b]. These systems are also restricted to reconstructions of planar scenes or to object tracking inside the hidden scene. In our work we observe that, when computationally transforming the temporal profile of diffuse surfaces into the frequency domain (following the phasor-field model [Liu et al. 2019b]), they exhibit mirror-like reflectance properties. We leverage higher-order illumination bounces produced by these virtual mirrors to overcome classical NLOS visibility challenges and to look around an additional second corner.

3 BACKGROUND AND INSIGHTS

In the following, we cover the background on wave-based NLOS imaging that forms the basis of our work, as well as related insights to understand our contributions. In Section 4 we describe the two key existing challenges that we address in this paper.

3.1 Wave-based NLOS image formation

In a classic NLOS setup, the capture device illuminates and measures indirect light at a diffuse relay surface, lacking direct line of sight to the target scenes being imaged. The phasor-field formulation [Liu et al. 2019b] brings the NLOS problem into a *virtual* LOS domain, creating computational imaging devices at the relay surface which can directly illuminate and capture the hidden scene. To understand our work, it is key to distinguish between the real and computational imaging domains in the phasor-field formulation. First, in the real domain, the acquisition process involves a laser device that emits illumination pulses towards locations \mathbf{x}_l on the relay surface (Figure 2a). The resulting indirect illumination produced by the hidden scene is then captured by an ultra-fast sensing device at points $\mathbf{x}_s \in \mathcal{S}$ on the relay surface (Figure 2b), yielding a time-resolved impulse response function $H(\mathbf{x}_l, \mathbf{x}_s, t)$, where t represents time. In the second stage, the phasor-field framework operates on H to *computationally* illuminate the hidden scene (Figure 2c) and then compute images of the hidden scene under such illumination (Figure 2d). Note that, while the first stage illuminates and senses the scene in the real domain, the second stage (which we describe in the rest of this section) is entirely computational.

The time of flight between the laser device and \mathbf{x}_l (Figure 2a, red) and between $\mathbf{x}_s \in \mathcal{S}$ and the sensing device (Figure 2b, blue) introduces temporal delays on the illumination captured in $H(\mathbf{x}_l, \mathbf{x}_s, t)$. In practice, we shift the temporal dimension of H so that t represents the time of flight of light paths that start at \mathbf{x}_l , scatter in the hidden scene, and end at \mathbf{x}_s . Placing the origin of the time reference system at the relay surface instead of at the laser and sensing devices is common practice in NLOS imaging and does not affect the algorithms [Liu et al. 2019b; Marco et al. 2021].

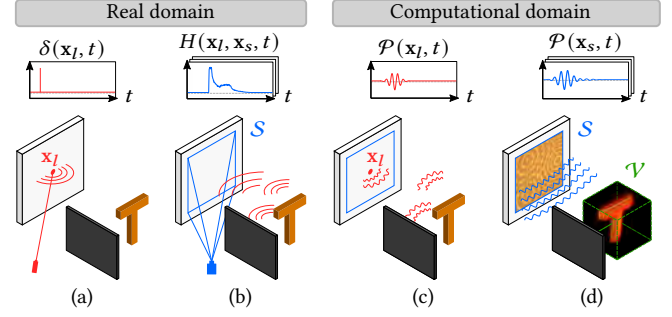


Fig. 2. (a) A laser device emits a delta light pulse $\delta(\mathbf{x}_l, t)$ that illuminates a point \mathbf{x}_l on the relay surface. (b) An ultra-fast sensing device captures time-resolved illumination $H(\mathbf{x}_l, \mathbf{x}_s, t)$ on multiple visible points \mathbf{x}_s of the imaging aperture \mathcal{S} . (c) By convolving the impulse response $H(\mathbf{x}_l, \mathbf{x}_s, t)$ with an illumination function $\mathcal{P}(\mathbf{x}_l, t)$, the phasor-field framework obtains the response of the hidden scene $\mathcal{P}(\mathbf{x}_s, t)$ to any arbitrary illumination function using only computation. (d) Lastly, the relay surface acts as a computational lens that focuses (propagates and adds) the response $\mathcal{P}(\mathbf{x}_s, t)$ from all points \mathbf{x}_s at each point \mathbf{x}_v in the bounding volume \mathcal{V} of the hidden scene, effectively transforming the NLOS problem into a virtual LOS problem. For this example, we compute the time-resolved image $f_{cc}(\mathbf{x}_v, t)$ using the confocal camera model which, when evaluated at $t = 0$, shows the T-shaped object.

Computationally illuminating the hidden scene. Points \mathbf{x}_l in the relay surface reflect delta illumination pulses emitted by the laser. In the computational domain, these points \mathbf{x}_l are now the *emitters*. Given the impulse response function $H(\mathbf{x}_l, \mathbf{x}_s, t)$, captured from a delta illumination pulse $\delta(\mathbf{x}_l, t)$, we can compute the response of the scene to any other arbitrary time-resolved illumination function $\mathcal{P}(\mathbf{x}_l, t)$ emitted from \mathbf{x}_l (Figure 2c). The resulting time-resolved response $\mathcal{P}(\mathbf{x}_s, t)$ at points \mathbf{x}_s on the relay surface is computed as

$$\mathcal{P}(\mathbf{x}_s, t) = \int_{\mathcal{L}} \mathcal{P}(\mathbf{x}_l, t) * H(\mathbf{x}_l, \mathbf{x}_s, t) d\mathbf{x}_l, \quad (1)$$

where $*$ represents a convolution in time. Throughout our paper, we only use a single point \mathbf{x}_l located at the center of the relay surface. Following the work by Liu et al. [2019b], for the illumination function $\mathcal{P}(\mathbf{x}_l, t)$ we use a pulse wave with a Gaussian envelope (Figure 2c) with wavelength λ_c and standard deviation σ :

$$\mathcal{P}(\mathbf{x}_l, t) = e^{i2\pi \frac{t}{\lambda_c} - \frac{1}{2} \left(\frac{t}{\sigma} \right)^2}. \quad (2)$$

We discuss the choice of λ_c and σ and their implications in the imaging process in Section 5.

The virtual camera analogy. The phasor-field framework treats the relay surface as the computational lens in a virtual camera that directly observes the hidden scene, with an aperture \mathcal{S} defined by points $\mathbf{x}_s \in \mathcal{S}$. The lens operators used in the phasor-field framework are well-known wave-based lens imaging operators [Liu et al. 2020, 2019b] defined as a function of each frequency component Ω of the signal. Because of this, we transform $\mathcal{P}(\mathbf{x}_s, t)$ to the frequency domain by applying a Fourier transform over the time domain, obtaining the complex-valued field $\hat{\mathcal{P}}(\mathbf{x}_s, \Omega) = \mathcal{F}\{\mathcal{P}(\mathbf{x}_s, t)\}$. Each complex value $\hat{\mathcal{P}}(\mathbf{x}_s, \Omega)$ represents a wave (i.e., the values form a

phasor field) resulting from illuminating the scene with the phasor $\widehat{\mathcal{P}}(\mathbf{x}_I, \Omega) = \mathcal{F}\{\mathcal{P}(\mathbf{x}_I, t)\}$, with illumination frequency Ω . In the camera analogy, $\widehat{\mathcal{P}}(\mathbf{x}_s, \Omega)$ can be understood as out-of-focus illumination from the hidden scene, and the goal of the computational lens is to focus the phasors $\widehat{\mathcal{P}}(\mathbf{x}_s, \Omega)$ to form an in-focus image of the hidden scene (Figure 2d). This focusing operation is specific to the chosen imaging operator Φ , which defines the properties of the virtual camera:

$$\widehat{f}(\mathbf{x}_v, \Omega) = \Phi(\mathbf{x}_v, \widehat{\mathcal{P}}(\mathbf{x}_s, \Omega)), \quad (3)$$

where $\widehat{f}(\mathbf{x}_v, \Omega)$ is the resulting image of the hidden scene computed for an imaging frequency Ω at points \mathbf{x}_v on the focal plane of the virtual camera. Applying the focusing operation from the phasor field $\widehat{\mathcal{P}}(\mathbf{x}_s, \Omega)$ at a point on a plane behind the relay surface is equivalent to focusing at symmetric locations in front of the relay surface. For clarity in the explanations, we use the latter approach in the rest of the paper, where \mathbf{x}_v denotes a point in the bounding volume \mathcal{V} of the hidden scene (Figure 2d). Due to the large size of the aperture \mathcal{S} at the relay surface, imaging the hidden scene with a single focal plane (as in conventional photography) results in a very shallow depth of field, which yields out-of-focus illumination from objects outside the focal plane [Marco et al. 2021]. To mitigate this problem, we use the computational lens to sweep the focal plane across the hidden scene, capturing a sequence of planar images taken at several focal distances (i.e., creating a focal stack). We arrange these planar images to form a volumetric image of the scene contained in \mathcal{V} .

Lens imaging operators. The analogy of a computational lens focusing at any point \mathbf{x}_v in \mathcal{V} is defined, in practice, by the propagation of the phasors $\widehat{\mathcal{P}}(\mathbf{x}_s, \Omega)$ from all points \mathbf{x}_s in the aperture \mathcal{S} to \mathbf{x}_v . In general, when light travels from any point \mathbf{x}_a to another point \mathbf{x}_b , the phasor $\widehat{\mathcal{P}}(\mathbf{x}_a, \Omega)$ at \mathbf{x}_a undergoes a phase shift and attenuation modeled by the Rayleigh-Sommerfeld Diffraction (RSD) operator. The phasor at \mathbf{x}_b , $\widehat{\mathcal{P}}(\mathbf{x}_b, \Omega)$, is then computed as

$$\widehat{\mathcal{P}}(\mathbf{x}_b, \Omega) = \widehat{\mathcal{P}}(\mathbf{x}_a, \Omega) \frac{e^{ik|\mathbf{x}_b - \mathbf{x}_a|}}{|\mathbf{x}_b - \mathbf{x}_a|}, \quad (4)$$

where $|\mathbf{x}_b - \mathbf{x}_a|$ is the optical distance between \mathbf{x}_a and \mathbf{x}_b . The wavenumber $k = 2\pi\Omega/c$ (where c is the speed of light) is the conversion factor from optical distance to phase. The phase shift (numerator) and attenuation (denominator) form the RSD operator. The phasor-field formulation uses RSD operators to define its imaging operators Φ , and capture images $\widehat{f}(\mathbf{x}_v, \Omega)$ of the hidden scene with different characteristics. These images can be defined under the following general expression:

$$\widehat{f}(\mathbf{x}_v, \Omega) = \int_{\mathcal{S}} \frac{e^{ikt_s c}}{t_s c} \int_{\mathcal{L}} \frac{e^{ikt_l c}}{t_l c} \widehat{\mathcal{P}}(\mathbf{x}_I, \Omega) \widehat{H}(\mathbf{x}_I, \mathbf{x}_s, \Omega) d\mathbf{x}_I d\mathbf{x}_s, \quad (5)$$

where t_s and t_l are parameters of the RSD operators, representing time of flight, and are used to implement the two different camera models in our work, as explained in Section 3.2. Note that the imaging frequency Ω is included in the wavenumber $k = 2\pi\Omega/c$

of these operators. $\widehat{\mathcal{P}}(\mathbf{x}_I, \Omega)$ and $\widehat{H}(\mathbf{x}_I, \mathbf{x}_s, \Omega)$ are the frequency-domain counterparts of $\mathcal{P}(\mathbf{x}_I, t)$ and $H(\mathbf{x}_I, \mathbf{x}_s, t)$, respectively, obtained via Fourier transform. The innermost integral over \mathcal{L} is the frequency-domain counterpart of the illumination function (Equation 1), but including an RSD operator. This allows for more general imaging operators by interpreting \mathcal{L} as an illumination aperture where another lens on the relay surface can focus the emitted illumination at specific locations \mathbf{x}_v in the scene. In our case, with only a single point \mathbf{x}_I in \mathcal{L} , this focus operation simply applies a phase shift and attenuation to the illumination phasor $\widehat{\mathcal{P}}(\mathbf{x}_I, \Omega)$; note that we can apply a different focus operation at each location \mathbf{x}_v in the scene. In the time domain, this phase shift effectively shifts the time instant at which each point \mathbf{x}_v in the scene is computationally illuminated using $\mathcal{P}(\mathbf{x}_I, t)$. In the following, we summarize the two camera models used in our work, which result from choosing specific values for the parameters t_s and t_l in Equation 5.

3.2 Time-resolved camera models

Throughout our paper we implement two camera models introduced by previous work [Liu et al. 2019b]: the *transient* camera model and the *confocal* camera model. These two models allow us to address different challenges in NLOS imaging in this work, such as the missing-cone problem and imaging objects hidden around two corners, by analyzing captured images of the scene at different locations and time instants. These models compute time-resolved images of the hidden scene via inverse Fourier transform over the frequency domain. We use $f_{tc}(\mathbf{x}_v, t) = \mathcal{F}^{-1}\{\widehat{f}_{tc}(\mathbf{x}_v, \Omega)\}$ and $f_{cc}(\mathbf{x}_v, t) = \mathcal{F}^{-1}\{\widehat{f}_{cc}(\mathbf{x}_v, \Omega)\}$ to refer to time-resolved images computed using the transient and confocal camera models $\widehat{f}_{tc}(\mathbf{x}_v, \Omega)$ and $\widehat{f}_{cc}(\mathbf{x}_v, \Omega)$, respectively. These models result from choosing specific values of t_s and t_l in Equation 5.

As previously discussed, we mitigate depth of field issues coming from the large camera aperture by computing multiple planar images at different focal distances that cover the volume \mathcal{V} , which form a focal stack; we arrange these planar images to form a volumetric image $\widehat{f}(\mathbf{x}_v, \Omega)$ for all points \mathbf{x}_v of the hidden scene. Any hidden scene element located at \mathbf{x}_v will therefore be in focus in one planar image that forms $\widehat{f}(\mathbf{x}_v, \Omega)$. Due to the time-frequency correspondence, the resulting time-resolved volumetric image $f(\mathbf{x}_v, t) = \mathcal{F}^{-1}\{\widehat{f}(\mathbf{x}_v, \Omega)\}$ is composed of a set of time-resolved planar images in the focal stack. In conventional photography, a focal stack contains multiple planar images of the same scene captured at different focal distances. The case of $f(\mathbf{x}_v, t)$ is analogous, but also including the temporal dimension: each frame at t of $f(\mathbf{x}_v, t)$ combines multiple planar images that capture the hidden scene at a different focal distances and at different time instants. Due to this effect, light transport events that occurred at the same time in the scene are captured in a different order on each planar image that forms the time-resolved volumetric image $f(\mathbf{x}_v, t)$. Consequently, each frame of $f(\mathbf{x}_v, t)$ may simultaneously show multiple light transport events of the hidden scene, despite these occurred at different time instants. In our work we compute time-resolved volumetric images $f_{tc}(\mathbf{x}_v, t)$ and $f_{cc}(\mathbf{x}_v, t)$ to address different challenges of NLOS imaging by analyzing light transport events in the hidden scene. These events have occurred at different instants, but each frame t captures them simultaneously.

Transient camera. The transient camera model implements a computational lens located at \mathcal{S} , focused at hidden scene points \mathbf{x}_v with $t_s = |\mathbf{x}_s - \mathbf{x}_v|/c$, and it does not implement any lens for the illumination aperture \mathcal{L} . For this, the RSD operator in \mathcal{L} of Equation 5 is ignored, resulting in

$$\widehat{f}_{tc}(\mathbf{x}_v, \Omega) = \int_{\mathcal{S}} \frac{e^{ik|\mathbf{x}_s - \mathbf{x}_v|}}{|\mathbf{x}_s - \mathbf{x}_v|} \int_{\mathcal{L}} \widehat{\mathcal{P}}(\mathbf{x}_l, \Omega) \widehat{H}(\mathbf{x}_l, \mathbf{x}_s, \Omega) d\mathbf{x}_l d\mathbf{x}_s. \quad (6)$$

The computed images $f_{tc}(\mathbf{x}_v, t) = \mathcal{F}^{-1} \{ \widehat{f}_{tc}(\mathbf{x}_v, \Omega) \}$ of this camera model resemble the captures of existing time-resolved cameras [Heide et al. 2013; Velten et al. 2013] if the hidden scene were illuminated using $\mathcal{P}(\mathbf{x}_l, t)$. However, as previously discussed, the computed images $f_{tc}(\mathbf{x}_v, t)$ capture the hidden scene at a different time for every point \mathbf{x}_v . As a result from the focusing operation in Equation 6, the hidden scene elements that are in focus at \mathbf{x}_v will be captured in $f_{tc}(\mathbf{x}_v, t)$ in the frame with time t corresponding to the time of flight between \mathbf{x}_l and the location of the *actual* scene element captured at \mathbf{x}_v . Importantly, as we deal with mirror reflections in this work, note that the location of the actual scene element may not correspond with \mathbf{x}_v if $f_{tc}(\mathbf{x}_v, t)$ captures a mirror reflection of such element at \mathbf{x}_v . For actual elements of the hidden scene located at \mathbf{x}_v , the computed images are a good approximation of time-resolved light transport at such elements under computational illumination. In our work, we use this camera model to analyze mirror reflections of known elements (e.g., the known illuminated point \mathbf{x}_l) in the hidden scene, and show how to leverage them to address the missing-cone problem.

Confocal camera. We also implement a confocal camera model, which creates two computational lenses located at \mathcal{L} and \mathcal{S} , focused at the same hidden scene point \mathbf{x}_v , respectively using $t_l = |\mathbf{x}_v - \mathbf{x}_l|/c$ and $t_s = |\mathbf{x}_s - \mathbf{x}_v|/c$, resulting in

$$\widehat{f}_{cc}(\mathbf{x}_v, \Omega) = \int_{\mathcal{S}} \frac{e^{ik|\mathbf{x}_s - \mathbf{x}_v|}}{|\mathbf{x}_s - \mathbf{x}_v|} \int_{\mathcal{L}} \frac{e^{ik|\mathbf{x}_v - \mathbf{x}_l|}}{|\mathbf{x}_v - \mathbf{x}_l|} \widehat{\mathcal{P}}(\mathbf{x}_l, \Omega) \widehat{H}(\mathbf{x}_l, \mathbf{x}_s, \Omega) d\mathbf{x}_l d\mathbf{x}_s. \quad (7)$$

Throughout our work, we use a single illuminated point \mathbf{x}_l , which yields a special case of $f_{cc}(\mathbf{x}_v, t) = \mathcal{F}^{-1} \{ \widehat{f}_{cc}(\mathbf{x}_v, \Omega) \}$ for the frame at $t = 0$. The phase shift defined by $t_l = |\mathbf{x}_v - \mathbf{x}_l|/c$ in the illumination aperture \mathcal{L} is equivalent to a temporal shift of the illumination function $\mathcal{P}(\mathbf{x}_l, t)$ corresponding to the time of flight from \mathbf{x}_l to each location \mathbf{x}_v . Similar to the transient camera model, the computed images $f_{cc}(\mathbf{x}_v, t)$ capture the hidden scene at a different time for every imaged point \mathbf{x}_v . For this camera model, the frame at $t = 0$ of $f_{cc}(\mathbf{x}_v, t)$ represents direct illumination from the emitter at scene elements that are in focus at \mathbf{x}_v . This is equivalent to the imaging models used by the vast majority of existing time-gated NLOS imaging methods [Ahn et al. 2019; Lindell et al. 2019b; Liu et al. 2019b; O'Toole et al. 2018; Velten et al. 2012a; Xin et al. 2019] to reconstruct single-corner hidden scenes, which consider light with the shortest path $\mathbf{x}_l \rightarrow \mathbf{x}_v \rightarrow \mathbf{x}_s$ a good estimator of the hidden geometry at \mathbf{x}_v . In our work, we use this camera model as an intermediate step to address the missing-cone problem, and we show how to leverage it to image objects hidden behind *two* corners.

4 CURRENT NLOS IMAGING LIMITATIONS

In this paper we address two of the fundamental limitations of NLOS imaging: the missing-cone problem and single-corner imaging, which we summarize in the following.

4.1 The missing-cone problem

The missing cone is a fundamental problem of NLOS imaging where, for a relay surface of a given size, certain hidden surfaces cannot be accessed by NLOS measurements (i.e., the impulse response H), and thus cannot be reconstructed regardless of the imaging method employed. This problem is inherent to other well-established imaging methods as well e.g., computed tomography [Benning et al. 2015; Delaney and Bresler 1998], which assume three-bounce transport. We use the term *null-reconstruction space* to denote the set of such surfaces that cannot be accessed by NLOS measurements. Note that surfaces inside the null-reconstruction space may still reflect some light towards the relay surface. However, their response H only changes its unmodulated components (frequency $\Omega = 0$), lacking the required modulated changes (components with frequency $\Omega > 0$) to be reconstructed using third-bounce time-of-flight information. The in-depth analysis by Liu et al. [2019a] shows that the missing-cone problem is universal across any NLOS measurement, and therefore affects both the transient camera f_{tc} and the confocal camera f_{cc} models. In Section 6, we intuitively address the missing-cone problem from our virtual mirror analogy, and use fourth-bounce illumination to image objects that are inside the null-reconstruction space of third-bounce methods.

4.2 Single-corner imaging

The existing time-resolved NLOS imaging methods are limited to reconstructing objects hidden behind a single corner (Figure 2), based on third-bounce illumination assumptions. However, general scenes typically contain objects hidden behind several occluders, creating higher-order illumination at the relay surface that is equally captured on the impulse response function H by time-resolved sensors, degrading the imaging results of third-bounce methods. In our work, we show how to leverage such higher-order illumination with applications such as imaging objects hidden behind two corners in Section 7.

5 DIFFUSE SURFACES AS VIRTUAL MIRRORS

A key observation in our work is that surfaces that are diffuse under visible light may still exhibit specular properties in the computational NLOS wave imaging domain. As we will show, leveraging this observation allows us to extend the current range of NLOS imaging capabilities, including imaging scenes that are hidden behind *two* corners.

In wave-based methods, time-resolved transport $\mathcal{P}(\mathbf{x}, t)$ at any point \mathbf{x} in the hidden scene becomes a phasor $\widehat{\mathcal{P}}(\mathbf{x}, \Omega)$ in the frequency domain. According to Huygens' principle, when reaching a surface M , this spherical wavefront will in turn generate multiple secondary spherical wavefronts. As an example, consider a point light at \mathbf{x}_l whose emission is defined by a phasor $\widehat{\mathcal{P}}(\mathbf{x}_l, \Omega)$ that illuminates points \mathbf{x}_m on a planar surface M , resulting in phasors $\widehat{\mathcal{P}}(\mathbf{x}_m, \Omega)$. We can then compute the resulting phasor at any point \mathbf{x}_v




Fig. 3. Specular behavior of a diffuse surface in the computational wave-based domain. (a) During the capture process, each point \mathbf{x}_m on a diffuse surface M reflects incoming light (red) isotropically (reflected rays in grey, reflection shape in cyan). On the right, the light from \mathbf{x}_l reflected on all points \mathbf{x}_m in the surface M is simulated for each point \mathbf{x}_v in \mathcal{V} using ray optics and Monte Carlo and the steady-state time average is shown. In the Fourier domain this is the component with frequency $\Omega = 0$. (b) When considering only modulated light (i.e., choosing a Fourier component with frequency $\Omega > 0$), as predicted by Huygens' principle, the resulting wavefront (cyan) from the reflected waves at each point \mathbf{x}_m (grey) follows the specular direction. The simulation on the right takes into account wave propagation and interference, by solving Equation 8 for all points \mathbf{x}_v in \mathcal{V} using Monte Carlo integration.

in a volume \mathcal{V} as a superposition of phasors from \mathbf{x}_m by extending Equation 4 as

$$\begin{aligned} \widehat{\mathcal{P}}(\mathbf{x}_v, \Omega) &= \int_M \widehat{\mathcal{P}}(\mathbf{x}_m, \Omega) \frac{e^{ik|\mathbf{x}_v - \mathbf{x}_m|}}{|\mathbf{x}_v - \mathbf{x}_m|} d\mathbf{x}_m \\ &= \int_M \widehat{\mathcal{P}}(\mathbf{x}_l, \Omega) \frac{e^{ik(|\mathbf{x}_v - \mathbf{x}_m| + |\mathbf{x}_m - \mathbf{x}_l|)}}{|\mathbf{x}_v - \mathbf{x}_m| |\mathbf{x}_m - \mathbf{x}_l|} d\mathbf{x}_m. \end{aligned} \quad (8)$$

For diffuse surfaces which are planar with respect to the illumination wavelength $\lambda = \Omega^{-1}$, the newly generated phasors $\widehat{\mathcal{P}}(\mathbf{x}_v, \Omega)$ result in a specular reflection of the incoming wavefront from $\widehat{\mathcal{P}}(\mathbf{x}_l, \Omega)$. In practice, this means that while the surface M may reflect visible light in all directions, the transient modulations (components with frequency $\Omega > 0$) that we need to image the scene propagate in the specular direction of the reflected computational wave.

This is shown in Figure 3a; we illustrate the capture process (real domain) of incoming light from \mathbf{x}_l that reaches points \mathbf{x}_m on a diffuse surface M which, as the simulation on the right shows for points \mathbf{x}_v in the volume \mathcal{V} , reflects light isotropically in all directions. In Figure 3b, we illustrate the wave-based computational light transport of the same scene, ignoring unmodulated light (frequency $\Omega = 0$) and considering only the part of the signal containing transient-modulated intensity (i.e., choosing a Fourier component with frequency $\Omega > 0$). Also in Figure 3b, the first schematic shows a planar light wavefront (red) at the time that it reaches the surface. The second schematic shows the reflected wavefront (cyan) at a later time instant, resulting from the superposition of spherical wavefronts (grey) at points \mathbf{x}_m in M . This produces a specular reflection, as predicted by Huygens' principle, shown in the simulation on the right. Both simulations on the right of Figure 3 have been generated using Monte Carlo integration, using standard ray optics (Figure 3a) and wave optics as described by Equation 8 (Figure 3b).

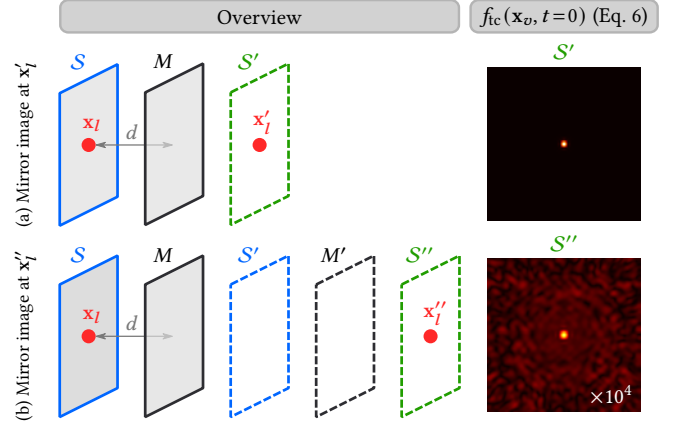


Fig. 4. (a) A planar surface M coplanar to the aperture S at a distance d creates a mirror image at \mathbf{x}'_l of the illuminated point \mathbf{x}_l that we image from a third-bounce reconstruction method at a distance $2d$; we show $f_{ic}(\mathbf{x}_v, t=0)$ for points \mathbf{x}_v in the plane S' , the mirror image of S . (b) Since the planar relay surface with aperture S also behaves as a mirror in the wave domain, another reflection of \mathbf{x}_l appears at a distance $4d$, only changing the imaged volume to points \mathbf{x}_v in the plane S'' (enhanced by a factor of 10^4 for visualization purposes). Both images use $\lambda_c = \sigma = 3$ cm.

Infinity mirror experiment. To further illustrate how these specular reflections take place in a NLOS scenario, we set up a simple simulated scene made up of a diffuse hidden surface M in front of the relay surface with aperture S , at a distance d (see Figure 4). We illuminate a point \mathbf{x}_l on the relay surface using a laser device, and obtain the impulse response $H(\mathbf{x}_l, \mathbf{x}_s, t)$ at points \mathbf{x}_s in S on the relay surface using transient rendering simulations [Jarabo et al. 2014; Royo et al. 2022]. The relay surface (with aperture S) and M are planar diffuse surfaces, thus behave like virtual mirrors in the computational wave domain. Looking at M , light emitted from \mathbf{x}_l is reflected by M , forming a mirror image of \mathbf{x}_l behind M like any conventional mirror. To capture in-focus images of the mirror reflection of \mathbf{x}_l , we place the focal plane of the virtual camera behind M . In particular, we use the impulse response $H(\mathbf{x}_l, \mathbf{x}_s, t)$ and implement a transient camera model (Equation 6) to obtain $f_{ic}(\mathbf{x}_v, t)$. We compute the time-resolved image $f_{ic}(\mathbf{x}_v, t)$ on the specific locations \mathbf{x}_v where the mirror images are formed. In Figure 4a, we place the focal plane of the camera at a distance $2d$, on a plane S' (green) which denotes the mirror image of S behind M . The plane S' also contains the mirror image at \mathbf{x}'_l of \mathbf{x}_l produced by M , captured in our result on the right as a bright spot. As stated in Section 3.2, each frame at t of $f_{ic}(\mathbf{x}_v, t)$ combines points \mathbf{x}_v at different times in the hidden scene: the transient camera model captures events in the frame at t equal to the time of flight from \mathbf{x}_l and the actual scene element corresponding to the mirror image at \mathbf{x}'_l , which is also \mathbf{x}_l in this case. Consequently, the mirror image at \mathbf{x}'_l is captured in the frame at $t = 0$.

We can image a higher-order mirror reflection pushing this effect even further, as shown in Figure 4b. Since the diffuse relay surface that contains the aperture S also behaves like a mirror in the computational wave domain, the specular interactions between M




Fig. 5. The visibility of surfaces in NLOS setups can be explained on the basis of their specular behavior in the wave domain. Our example follows the original missing cone explanation and uses the image (d) from work by Liu et al. [2019a], obtained using the confocal camera model f_{cc} . From an illuminated point x_l in the relay surface, the visibility of the three surfaces M_{1-3} depends on their position and orientation with respect to such relay surface (which also acts as the aperture S of a virtual camera). As (a-c) show, only the reflected wavefronts from M_1 and M_2 reach the aperture S . As a result (d), the imaging system cannot see surface M_3 .

and S create additional mirror reflections of the illuminated point x_l at locations further behind M . This effect is analogous to the real situation where we observe multiple reflections of an object placed between two confronted, real mirrors. We showcase this effect in Figure 4b, using the same impulse response $H(x_l, x_s, t)$ and implementing the same transient camera model (Equation 6), but adjusting the focal plane at a distance $4d$ to match points x_v in the plane S'' . Looking at the computed image on the right, this yields a clear but dimmer spot that corresponds to the second mirror image at x_l'' of the illuminated point x_l .

Computational wavelength. The computational specular behavior of a real diffuse surface is explained by wave optics (illustrated in Figure 3), and depends on the computational wavelengths λ used through the imaging process (λ is the inverse of the imaging frequencies Ω). In practice, the computational wavelengths used through the imaging process depend on the frequency spectrum of the illumination function $\hat{P}(x_l, \Omega)$. In our work, we determine the spectrum of frequencies Ω of the imaging process through the central wavelength λ_c and standard deviation σ in Equation 1. The choice of these frequencies introduces a trade-off: lower values of λ_c and σ can properly image geometric features with more detail, but they may also introduce unwanted high-frequency noise. Following previous works, we use values for λ_c from 3 cm to 14 cm, and values for σ proportional to λ_c . We specify the particular values used in each experiment. For reference, all experiments share the same aperture size of 2×2 m. Within this range of values, planar surfaces behave specularly during the NLOS imaging process, which we leverage to address different challenges of NLOS imaging methods. We do not use wavelengths larger than this range, as these may degrade the specular behavior of surfaces due to the ratio between the surface size and the wavelength.

Missing-cone problem through virtual mirrors. The frequency-space specular behavior of diffuse surfaces allows us to intuitively explain the missing-cone problem from our virtual mirrors perspective: for a point x_v in the hidden scene, if light from x_l does not

reach any point x_s in the aperture S after a specular reflection in x_v , then the point x_v is inside the null-reconstruction space of third-bounce imaging methods and cannot be reconstructed. Figure 5 illustrates this for third-bounce methods using a scene with three diffuse surfaces M_{1-3} . However, only two surfaces are visible on the image shown in Figure 5d, computed using existing third-bounce NLOS imaging methods. Wave propagation in the computational NLOS imaging domain is illustrated in Figure 5a to Figure 5c for M_1 to M_3 , respectively. The reflected wavefronts from surfaces M_1 and M_2 reach the sensor S . However, given its particular position and orientation, this is not the case for M_3 and thus cannot be imaged. The surface M_3 is said to be inside the null-reconstruction space.

Even if a surface is inside the null-reconstruction space of third-bounce imaging methods (such as M_3 in Figure 5), we observe that the combination of several surfaces may produce higher-order illumination bounces that actually do reach S . In Section 6, we show how to leverage such higher-order illumination to infer and to directly image objects that are inside the null-reconstruction space of third-bounce methods. We achieve this by analyzing mirror images of other objects produced by the surface inside the null-reconstruction space to infer its position and orientation. Beyond inference, we provide a procedure to translate our imaging system to a secondary surface that directly observes the target object.

Moreover, in Section 7 we show another way to leverage virtual mirrors in NLOS imaging, and image objects hidden behind *two* corners using existing imaging models. For this we rely on fifth-bounce specular reflections in the computational domain to image the space behind virtual mirror surfaces, where the mirror image of such objects would appear.

6 ADDRESSING THE MISSING CONE

In the following, we propose a procedure to address the missing-cone problem. Our key insight is that we can obtain information of surfaces inside the null-reconstruction space by analyzing the fourth-bounce illumination that reaches the relay surface through multiple interreflections with other surfaces in the hidden scene. First, we show how to analyze elements captured in time-resolved images of the hidden scene to infer the position and orientation of surfaces inside the null-reconstruction space of third-bounce methods. As this inference is limited by ambiguities, we propose a procedure to directly image such surfaces without requiring inference. All experiments in this section rely on simulated data using transient rendering [Jarabo et al. 2014; Royo et al. 2022].

Problem statement. Figure 6a shows an example scene, consisting of two surfaces M and G , the NLOS imaging aperture S , and the illuminated point x_l . In the real domain, some light paths will bounce on the illuminated point x_l , then on G and finally on a point in the aperture S , for a total of three bounces. However, in the NLOS computational domain, G reflects third-bounce illumination specularly, away from S (pink, note that not all three bounces are shown). Therefore, imaging G is limited if using third-bounce methods. Figure 7a illustrates the procedure of third-bounce methods, showing that G is inside the null-reconstruction space.

Our key observation is that four-bounce paths (purple, note that not all four bounces are shown) that reach S through a specular-like




Fig. 6. Illustration of our methodology that targets surfaces inside the null-reconstruction space of third-bounce NLOS imaging methods by using higher-order illumination. (a) Example NLOS scene with an imaging aperture S , an illuminated point x_I , and surfaces G and M . Surface G is inside the null-reconstruction space of classic NLOS imaging methods due to lost third-bounce paths (pink). We use higher-order bounces (purple), which reach S bouncing on both M and G . Note that both pink and purple paths do not show all three and four bounces, respectively. (b) Surfaces M and G produce mirror reflections of x_I and M after one or two specular reflections, denoted by their superscripts. We infer the position and orientation of G as the plane between a point and its mirror image observed by the camera. For example, points x_I and x_I^G , or points x_I^M and x_I^{MG} are the reflection of each other produced by the surface G , denoted by the dotted lines in between. (c) Our direct imaging procedure: To avoid inference ambiguities, we instead propose a procedure to directly image x_I^G and G by translating the aperture from S to M . For this, we obtain the transient response at points $x_v \in \mathcal{V}$, $f_{tc}(x_v, t)$, in a region \mathcal{V} (green) containing M , and evaluate $f_{tc}(x_v, t)$ at points $x_v \in M$. (d) We then implement imaging models at a secondary aperture M (blue) that observes elements in the volume \mathcal{W} (green) that were inside the null-reconstruction space from S . Using M as aperture, x_I^G is captured by the transient camera, and G is captured by the confocal camera (paths $x_I \rightarrow G \rightarrow M$, purple).

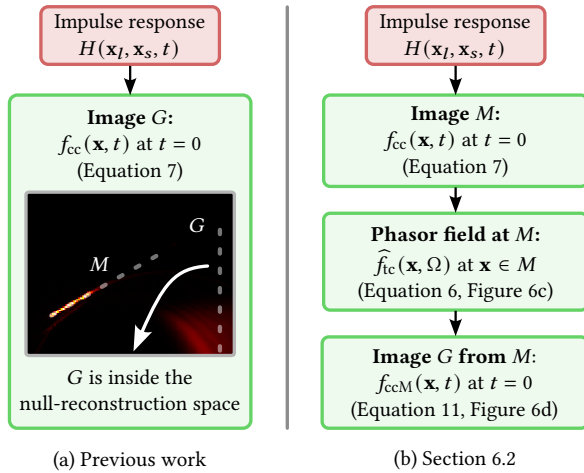


Fig. 7. Flowchart of the computational procedures applied to the impulse response $H(x_I, x_S, t)$ captured for the scene in Figure 6a. (a) Following existing procedures, trying to image G using the confocal camera model $f_{cc}(x, t)$, the frame at $t = 0$ shows that G is inside the null-reconstruction space. (b) In Section 6.2 we propose a procedure to address this problem. We use the confocal camera model with its aperture at M instead of S , denoted as f_{ccM} , which allows to image G directly. For this second aperture at M , we compute the phasors $\hat{\mathcal{P}}(x, \Omega) \equiv \hat{f}_{tc}(x, \Omega)$ for selected points $x \in M$ using the transient camera model.

reflection on M ($x_I \rightarrow G \rightarrow M \rightarrow S$) provide valuable information about G . In particular, we show how to analyze illumination in time-resolved images of the scene to infer the position and orientation of G through mirror images of scene elements produced by G (Section 6.1), and then propose a procedure to directly image G by creating secondary apertures at surfaces that are visible by third-bounce methods such as M (Section 6.2).

6.1 Inferring surfaces from mirror images

In this section, we show how to infer the position and orientation of G from the impulse response function $H(x_I, x_S, t)$. The key insight is that if we can observe mirror images of scene elements created by G , it is because G lies on the plane between such scene element and its mirror image, just like what happens with a real mirror. Figure 6b points the location of mirror images of the illuminated point x_I and M , produced by M and G in our example scene. Superscripts denote the surface (or surfaces, in order) that produces the mirror images of each scene element: x_I is mirrored by M and G at x_I^M and x_I^G , respectively; M is mirrored by G at M^G ; x_I^G is mirrored by M at x_I^{GM} ; and x_I^M is mirrored by G at x_I^{MG} .

Observing mirror images from S . Previously, we described the location of mirror images of the illuminated point x_I produced by surfaces in the hidden scene, as shown in Figure 6b. To show which of these mirror images are visible from the aperture S and use them to infer the location and orientation of G , we implement a transient camera model f_{tc} (Equation 6), and evaluate it at different time instants. In the resulting time-resolved image $f_{tc}(x_v, t)$, light emitted from the illuminated point x_I is captured in the frame at $t = 0$. We show the frame $f_{tc}(x_v, t = 0)$ in Figure 8a, which captures several bright spots at the locations of the mirror images of x_I that are visible from aperture S over a volume that covers the whole scene. In particular, the aperture S can observe the mirror images at x_I^M (produced by three-bounce paths $x_I \rightarrow M \rightarrow S$), at x_I^{MG} and at x_I^{GM} (both produced by four-bounce paths $x_I \rightarrow M \rightarrow G \rightarrow S$ and $x_I \rightarrow G \rightarrow M \rightarrow S$). We cannot, however, observe the mirror image at x_I^G , since three-bounce paths from x_I to G do not reach S in the computational domain (Figure 6a, pink). Evaluating $f_{tc}(x_v, t)$ at frames with $t > 0$ may show other scene elements and their mirror images, similarly to x_I , based on the time of flight from x_I to each scene element. For example, light from the illuminated point x_I




Fig. 8. (a) Mirror images at \mathbf{x}_l^M , \mathbf{x}_l^{GM} and \mathbf{x}_l^{MG} of \mathbf{x}_l produced by surfaces M and G , computed with the transient camera $f_{ic}(\mathbf{x}, t)$ (Equation 6) for points \mathbf{x} that cover the region depicted in Figure 6b. Although these reflections happen at different times, our computed time-resolved image composes them in the same frame at $t = 0$. (b) When light from \mathbf{x}_l reaches the center point $\bar{\mathbf{x}}_m$ of M at $\bar{t}_m = |\bar{\mathbf{x}}_m - \mathbf{x}_l|/c$, we can see both M and its mirror image M^G produced by G . Both images are computed using $\lambda_c = \sigma = 3$ cm.

will reach the central point of M , denoted as $\bar{\mathbf{x}}_m$ (Figure 8a), at the time of flight $\bar{t}_m = |\bar{\mathbf{x}}_m - \mathbf{x}_l|/c$. We can therefore identify reflections at points near the center of the plane M by looking at the frame at $t = \bar{t}_m$ of $f_{ic}(\mathbf{x}_v, t)$ (Figure 8b), which captures not only M and M^G , but many other mirror images produced by M and G .

Both images in Figure 8 show bright areas outside the points or surfaces mentioned before. This is mainly because the impulse response function $H(\mathbf{x}_l, \mathbf{x}_s, t)$ combines coupled information from paths of different bounces and optical lengths, which introduces out-of-focus, low-frequency artifacts at each imaged location \mathbf{x}_v . Also, note that during the NLOS imaging process we only know the location of the illuminated point at \mathbf{x}_l , and we do not have any prior knowledge of other hidden scene elements. Consequently, there exist ambiguities when identifying the captured bright spots as mirror images of \mathbf{x}_l , so we cannot guarantee that e.g., \mathbf{x}_l^{MG} is a mirror image of \mathbf{x}_l^M ; instead e.g., there may be two physical surfaces located at M^G and M , respectively.

Inferring the position and orientation of G . By assuming that \mathbf{x}_l^{MG} is a mirror image of \mathbf{x}_l^M —i.e., \mathbf{x}_l has undergone two reflections produced by M and then G —the surface G that produced \mathbf{x}_l^{MG} should lie in the perpendicular plane between \mathbf{x}_l^M and \mathbf{x}_l^{MG} . We obtain a point \mathbf{c}_G on such plane and its normal vector \mathbf{n}_G as

$$\mathbf{c}_G = \frac{\mathbf{x}_l^M + \mathbf{x}_l^{MG}}{2}, \quad \mathbf{n}_G = \frac{\mathbf{x}_l^M - \mathbf{x}_l^{MG}}{|\mathbf{x}_l^M - \mathbf{x}_l^{MG}|}, \quad (9)$$

which define the position and orientation of G , respectively. We could also infer the position and orientation of G from M and its mirror reflection M^G . Such inferences require assumptions on the number of reflections undergone by the observed patterns (Figure 8). For this particular inference, \mathbf{x}_l^M can be identified as the reflection of \mathbf{x}_l from the visible orientation of M , and \mathbf{x}_l^{MG} cannot be produced by M , so we assume it is a second mirror reflection by another hidden surface.

In general, there are ambiguities on recognizing the source of every identified reflection, which may introduce errors when inferring hidden surfaces. To avoid errors due to such ambiguities, in the following we propose a procedure to directly image G without requiring inference.

6.2 Imaging surfaces from secondary apertures

Here we show a procedure to directly image G using fourth-bounce illumination. The key idea is that, while the surface G is inside the null-reconstruction space of imaging systems created at S (Figure 6a, pink), G is not inside the null-reconstruction space of imaging systems created at M , since there exist three-bounce paths $\mathbf{x}_l \rightarrow G \rightarrow M$ that reach M (Figure 6d, purple). Based on this observation, we show how to computationally translate our imaging system from S to M (Figure 6c) to directly image G using M as a secondary aperture (Figure 6d).

Our procedure is illustrated in Figure 7b: first, how to obtain a phasor field at M , and then how to use this response to generate computational cameras with the aperture located at M . The rest of this section describes this procedure in detail along with its results.

Phasor field at M . The phasor-field formulation uses phasors $\hat{\mathcal{P}}(\mathbf{x}_s, \Omega)$ at points \mathbf{x}_s , which encode the response of the scene to the illumination function, to implement imaging models with a camera aperture S . To translate the camera aperture from points \mathbf{x}_s to points $\mathbf{x}_m \in M$, we need to compute their corresponding phasors $\hat{\mathcal{P}}(\mathbf{x}_m, \Omega)$. The transient camera model f_{ic} achieves this goal, since $\hat{\mathcal{P}}(\mathbf{x}_m, \Omega) \equiv \hat{f}_{ic}(\mathbf{x}_m, \Omega)$ propagates phasors from all points \mathbf{x}_s to each point \mathbf{x}_m . To determine where M is, we first implement a confocal camera model to capture $f_{cc}(\mathbf{x}_v, t)$ from S —equivalent to existing NLOS imaging models—to image all points \mathbf{x}_v in the region \mathcal{V} (Figure 6c). The computed frame at $t = 0$ can be seen in Figure 7a. We can estimate all points \mathbf{x}_m in M by thresholding the image $f_{cc}(\mathbf{x}_v, t = 0)$. We then implement a transient camera model to obtain $\hat{f}_{ic}(\mathbf{x}_m, \Omega)$, yielding a phasor field at M that can be used to implement new imaging systems at M .

M as a secondary aperture. Based on RSD propagation principles, we implement a lens at M which focuses the phasors $\hat{f}_{ic}(\mathbf{x}_m, \Omega)$ at points \mathbf{x}_w in a volume \mathcal{W} that contains G (Figure 6d), equivalent to what the transient camera model does. In practice, we include an RSD propagator from points \mathbf{x}_m to points \mathbf{x}_w , yielding the transient camera model with aperture M (denoted as f_{icM}) as

$$\hat{f}_{icM}(\mathbf{x}_w, \Omega) = \int_M \frac{e^{ik|\mathbf{x}_w - \mathbf{x}_m|}}{|\mathbf{x}_w - \mathbf{x}_m|} \hat{f}_{ic}(\mathbf{x}_m, \Omega) d\mathbf{x}_m. \quad (10)$$

The time-domain version $f_{icM}(\mathbf{x}_w, t) = \mathcal{F}^{-1} \left\{ \hat{f}_{icM}(\mathbf{x}_w, \Omega) \right\}$ represents a time-resolved image of the scene as captured from M . The frame at $t = 0$ of $f_{icM}(\mathbf{x}_w, t)$ captures the initial light up of the illuminated point \mathbf{x}_l and its mirror image at \mathbf{x}_l^G . Similarly to the mirror images captured from S (Figure 8), given \mathbf{x}_l and its mirror image \mathbf{x}_l^G (captured by the aperture at M) we can infer the position and orientation of the surface G that produced \mathbf{x}_l^G (Equation 9).

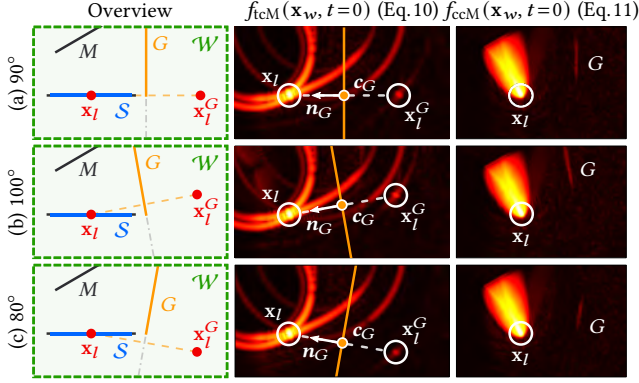


Fig. 9. **Left column:** We rotate G in Figure 6a to show that it acts as a virtual mirror surface. In (a-c), G is rotated 90° , 100° and 80° with respect to the relay surface. Rotating G changes the position of the mirror image at \mathbf{x}_l^G of the illuminated point \mathbf{x}_l . We image all points \mathbf{x}_w in the volume \mathcal{W} . **Middle column:** From the known position of \mathbf{x}_l and the computed position of \mathbf{x}_l^G , we can use the transient camera model with aperture M (f_{icM}) and infer the position \mathbf{c}_G and orientation \mathbf{n}_G of surface G . **Right column:** We directly image G using the confocal camera model with aperture M (f_{ccM}). All images use $\lambda_c = \sigma = 3$ cm.

Imaging G from M . Instead of approximating G through geometric inference, we *directly* image G by extending Equation 10 to implement a confocal camera model at M . Under a single illuminated point \mathbf{x}_l , this is equivalent to incorporating an RSD operator to Equation 10 that propagates the phasor $\widehat{\mathcal{P}}(\mathbf{x}_m, \Omega) \equiv \widehat{f}_{ic}(\mathbf{x}_m, \Omega)$ accounting for the distance between \mathbf{x}_l and \mathbf{x}_w , yielding a confocal camera model with aperture M (denoted as f_{ccM}) as

$$\begin{aligned} \widehat{f}_{ccM}(\mathbf{x}_w, \Omega) &= \widehat{f}_{icM}(\mathbf{x}_w, \Omega) \frac{e^{ik|\mathbf{x}_w - \mathbf{x}_l|}}{|\mathbf{x}_w - \mathbf{x}_l|} \\ &= \int_M \frac{e^{ik(|\mathbf{x}_w - \mathbf{x}_m| + |\mathbf{x}_w - \mathbf{x}_l|)}}{|\mathbf{x}_w - \mathbf{x}_m||\mathbf{x}_w - \mathbf{x}_l|} \widehat{f}_{ic}(\mathbf{x}_m, \Omega) d\mathbf{x}_m. \end{aligned} \quad (11)$$

The time-resolved image $f_{ccM}(\mathbf{x}_w, t) = \mathcal{F}^{-1} \{ \widehat{f}_{ccM}(\mathbf{x}_w, \Omega) \}$ of this confocal camera model captures G in the frame at $t = 0$. This follows the basis of classic NLOS reconstruction methods, which directly image G from the time of flight of three-bounce paths $\mathbf{x}_l \rightarrow G \rightarrow M$, but instead using four-bounce paths $\mathbf{x}_l \rightarrow G \rightarrow M \rightarrow S$ in the captured impulse response $H(\mathbf{x}_l, \mathbf{x}_s, t)$.

6.3 Results

We illustrate results of our procedure using the scene in Figure 6a for both inference based on \mathbf{x}_l and its mirror image \mathbf{x}_l^G , and direct imaging of G in Figure 9. We use different orientations of the plane G to show how it affects the resulting mirror image \mathbf{x}_l^G and the direct image of G . In the left column we show an overview of the scenes. Note that while inference (Section 6.1) and direct imaging (Section 6.2) are two separate procedures, here we illustrate inference using the mirror image at \mathbf{x}_l^G captured through the secondary aperture at M of the direct imaging procedure (Section 6.2).

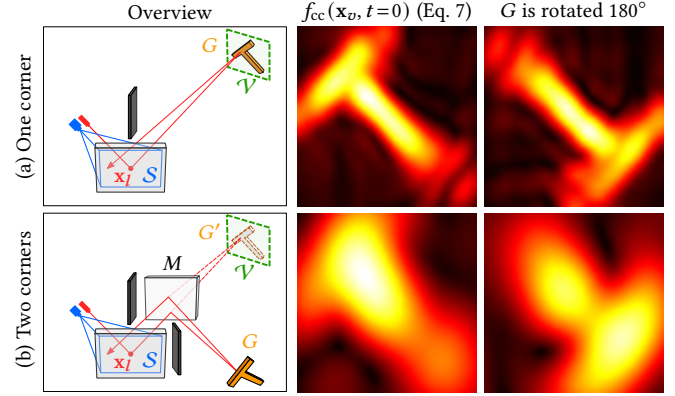


Fig. 10. (a) A T-shaped object is hidden behind one corner. Focusing the confocal camera at points \mathbf{x}_v in \mathcal{V} (using $\lambda_c = 5$ cm, $\sigma = 6$ cm) uses three-bounce illumination produced by G , and the T-shaped object is visible in the image when $t = 0$. (b) Our two-corner imaging setup. None of the points \mathbf{x}_s in \mathcal{S} , or \mathbf{x}_l , have a direct line of sight towards G , that is, three-bounce illumination cannot contain information about G . Focusing the confocal camera at points \mathbf{x}_v in \mathcal{V} (using $\lambda_c = 12$ cm, $\sigma = 14$ cm) captures the mirror image at G' of the geometry G produced by the diffuse surface M . We also rotate the T-shaped object by 180° , to show that it also affects the resulting image.

We show results for inference, identifying both \mathbf{x}_l and \mathbf{x}_l^G in the computed images $f_{icM}(\mathbf{x}_w, t)$ (Equation 10) at $\mathbf{x}_w \in \mathcal{W}$ and $t = 0$ (Figure 9, middle column). From \mathbf{x}_l and \mathbf{x}_l^G we infer a point \mathbf{c}_G and the normal \mathbf{n}_G of a plane corresponding to the surface G .

To illustrate our direct imaging procedure, we compute Equation 11 to directly image G , evaluating $f_{ccM}(\mathbf{x}_w, t)$ at $\mathbf{x}_w \in \mathcal{W}$ and $t = 0$ (Figure 9, right column). This yields a clear image of the surface G which is entirely on the null-reconstruction space of third-bounce methods (Figure 7a). Note only part of G is visible, the rest is inside the null-reconstruction space of the imaging system at M as some fourth-bounce illumination paths do not reach S through specular bounces on M in the computational domain. The results also show a bright region near \mathbf{x}_l . In this case, we use our fourth-bounce imaging method, but the impulse response function $H(\mathbf{x}_l, \mathbf{x}_s, t)$ combines coupled information from paths of different bounces and optical lengths, which translates into out-of-focus, low-frequency artifacts.

7 LOOKING AROUND TWO CORNERS

Here we leverage our virtual mirror reflections in the wave domain to image objects around *two* corners. The fundamental idea of our approach is exploiting our observation that diffuse planar surfaces behave like virtual mirrors (Section 5); we image the region where their mirror image is formed using only the confocal camera model (Equation 7). Note this imaging model is similar to those used by third-bounce NLOS imaging methods, analogous to the procedure detailed in Figure 7a. However, unlike previous work, we exploit fifth-bounce illumination in the impulse response function $H(\mathbf{x}_l, \mathbf{x}_s, t)$ by shifting the imaged region based on our observed behavior of virtual mirrors. In the following we detail this procedure using simulated data.

We illustrate this with the simulated scene depicted in Figure 10b, composed by the relay surface with an illuminated point \mathbf{x}_l and the aperture \mathcal{S} , a diffuse T-shaped object hidden behind two corners denoted as G (which we aim to image), and a diffuse surface M hidden behind a single corner, which behaves as a virtual mirror during the computational imaging process. We also place two black occluders to ensure that M is not directly visible to the NLOS imaging device, and G is not directly visible neither to the NLOS device nor to the imaging aperture \mathcal{S} or illuminated point \mathbf{x}_l .

Our goal is to image the object G . Due to the location and orientation of M in our scene, the aperture \mathcal{S} and the object G are at each other's specular direction with respect to M , so that M forms a mirror image of G in the space behind M that is captured from \mathcal{S} . Specifically, we place the diffuse surface M coplanar to the aperture \mathcal{S} with a lateral shift so it reflects light specularly towards G and back to \mathcal{S} through another bounce in M . This creates fifth-bounce paths with the form $\mathbf{x}_l \rightarrow M \rightarrow G \rightarrow M \rightarrow \mathcal{S}$ (marked in red in the schematic of Figure 10b), which we leverage to image G . Also, note that in our setup M is in the null-reconstruction space of third-bounce methods, as third-bounce specular paths do not reach \mathcal{S} in the computational domain. However, fifth-bounce specular paths actually reach \mathcal{S} , and therefore we can image G even when M is in the null-reconstruction space.

In our experiment setup, we first obtain a simulated impulse response function $H(\mathbf{x}_l, \mathbf{x}_s, t)$ of the scene using transient rendering [Jarabo et al. 2014; Royo et al. 2022] to mimic a real acquisition process of the relay surface. We then implement a confocal camera model (Equation 7) and compute the frame at $t = 0$ of $f_{cc}(\mathbf{x}_v, t)$ at points \mathbf{x}_v in the imaged plane on \mathcal{V} (Figure 10b, green), where the mirror image of G produced by M would be formed (G' in the schematic). The computed images on Figure 10b show the result of this imaging process for two orientations of the T-shaped geometry, showing that the shape's structure is preserved on both, even after this second corner. For reference, we configure a single-corner scene (Figure 10a) by removing the surface M and placing the object G at the position marked by G' , where the mirror image should be formed for the two-corner case. The images of the T-shaped object appear blurrier when imaged around two corners. Following our observations in Section 5, this effect is mainly caused by two factors. First, even if M is perfectly diffuse, the mirror-like behavior of M in the computational domain is not perfectly specular, and the resolution of the mirror images that M produces is limited by diffraction. Second, we use larger wavelengths on the two-corner case (i.e., values for λ_c and σ are higher in Figure 10b than Figure 10a). We discuss this mirror behavior and its effects further on Section 9.

8 RESULTS IN REAL SCENES

In the following we illustrate and validate our methods in real scenarios, imaging diffuse planar surfaces inside the null-reconstruction space of third-bounce methods in single-corner configurations, and then imaging scenes hidden behind two corners.

Hardware details. Our NLOS imaging system consists of a SPAD array sensor, a laser emitter and a two-mirror galvanometer (Figure 11). The galvanometer guides the laser towards multiple points on the relay surface, while the detector is aimed at a fixed position




Fig. 11. Our hardware setup used to validate our procedures in real scenarios. (a) 16x16 Single Photon Avalanche Diode (SPAD) array focused at a point on the relay surface. (b) Laser source which can emit 35 picosecond pulses. (c) Two-mirror galvanometer that guides the laser to scan the relay surface.

on the relay surface. A PM-1.03-25TM laser from Polar Laser Laboratories is used as an illumination source. The laser is combined with a frequency doubler to emit 515 nm pulses with a maximum pulse width of 35 ps, an average power of 375 mW, and at an average repetition rate of 5 MHz. A two-mirror Thorlabs galvanometer (Thorlabs GVS012) is used to scan a relay surface at 1 cm spacing, with a total scan area around 1.9×1.9 m. Our detector is a 16x16 Single Photon Avalanche Diode (SPAD) array [Riccardo et al. 2022] focused at a 7.1 cm by 4.7 cm area on the relay surface using a Canon EF 85 mm f/1.8 USM Lens. The temporal resolution of the array has a Full-Width at Half Maximum (FWHM) of around 60 ps and a deadtime of less than 100 ns. All scene surfaces are diffuse expanded polystyrene foam and unfinished drywall, with no retroreflective properties.

Addressing the missing cone. We design a scene similar to Figure 6a to test our procedure to address the missing-cone problem with our hardware setup. A photograph of the scene is displayed in Figure 12a, with two hidden surfaces M and G . Surface G is in the null-reconstruction space of third-bounce methods since third-bounce illumination from G falls outside our imaging aperture \mathcal{S} in the computational domain. We aim to first infer the position and orientation of G , then to directly image G . We capture the impulse response function $H(\mathbf{x}_l, \mathbf{x}_s, t)$ for an illuminated point \mathbf{x}_l and points \mathbf{x}_s of the aperture \mathcal{S} at the relay surface. The surface M produces fourth-bounce illumination at \mathbf{x}_s from our target diffuse surface G . We experiment with three different orientations of G (90° , 100° and 80°) with respect to the relay surface, which is in all cases located at 50 cm from the illuminated point \mathbf{x}_l (in all three cases, G cannot be imaged using existing NLOS algorithms). The surface M is tilted at 30° and separated 1.5 m from \mathbf{x}_l at its center point. Figure 12b and c show photographs and top-view schematics of the different orientations, respectively. Figure 12d shows how, for all three orientations, the position \mathbf{c}_G and orientation \mathbf{n}_G of G are accurately inferred from the illuminated point \mathbf{x}_l and its reflection \mathbf{x}_l^G captured with the transient camera from the aperture M (Equation 10), in $f_{icM}(\mathbf{x}_w, t = 0)$. For the 90° , 100° and 80° cases, our inferred \mathbf{c}_G is 6 cm, 4 cm and 5 cm away from G , and \mathbf{n}_G has an orientation error of 0.3° , 2.8° and 0.1° , respectively. Additionally, to directly image plane G we turn M into a secondary aperture where

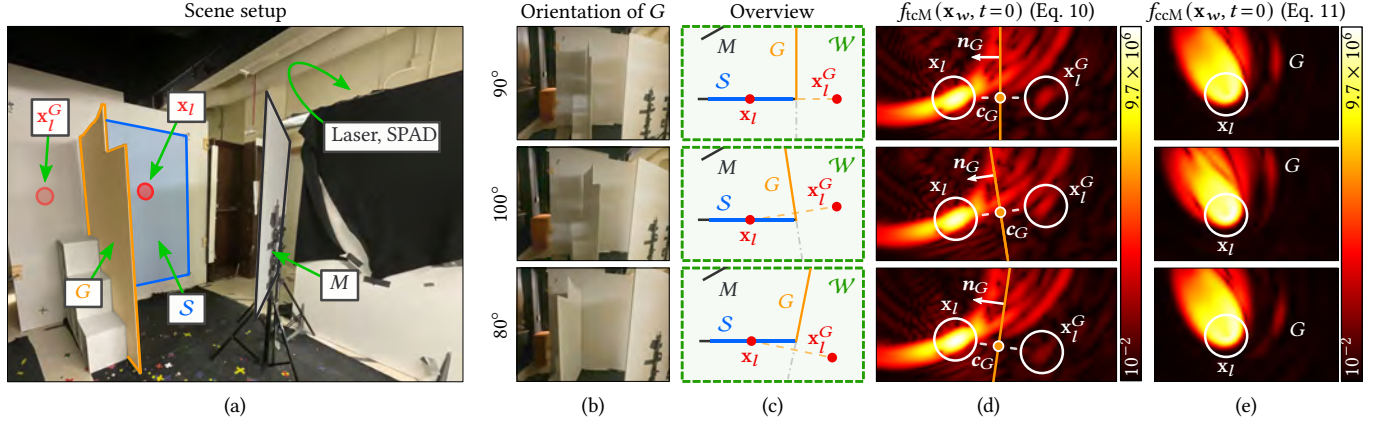


Fig. 12. Results for the inference of the position and orientation of surfaces inside the null-reconstruction space, and direct imaging of such surfaces (Section 6, same setup) using real captures. (a) Photograph of the general scene setup. Laser and SPAD array are on the right. (b) Photographs of different orientations of the surface G inside the null-reconstruction space, to be estimated for three independent experiments, one per row. (c) Overview of the scene setups (changing the orientation of G) and the imaged volume \mathcal{W} with points $\mathbf{x}_w \in \mathcal{W}$. (d) Inference of the position \mathbf{c}_G and orientation \mathbf{n}_G of surface G , from the illuminated point \mathbf{x}_l and its mirror image at \mathbf{x}_l^G produced by the virtual mirror surface G . Computed using $\lambda_c = \sigma = 7$ cm. (e) Imaging of the surface G from the secondary aperture M , using $\lambda_c = 10$ cm and $\sigma = 7$ cm. Both the inferred position and orientation of G and the direct image of G taken from M are a close match with respect to the capture setup in all three cases, even though it is not visible from the relay surface for classic NLOS imaging methods. The colorbar is displayed in a logarithmic scale.

we implement a confocal camera (Equation 11) to obtain the image $f_{ccM}(\mathbf{x}_w, t = 0)$. The results are shown in Figure 12e. Similar to Figure 9, the f_{ccM} model produces a bright region near \mathbf{x}_l due to coupled illumination in the impulse response $H(\mathbf{x}_l, \mathbf{x}_s, t)$.

Looking around two corners. In this experiment we use the scene shown in Figure 13a, where we image objects hidden behind two corners. The scene is made up of a diffuse surface M at a 45° angle with respect to the relay surface, and several hidden geometries G (Figure 13b, top) oriented at 90° with respect to the relay surface. Two occluders ensure that the geometry G hidden around two corners is not directly visible from points $\mathbf{x}_s \in \mathcal{S}$, \mathbf{x}_l , or the capture hardware itself. For each geometry, we image points \mathbf{x}_v in \mathcal{V} where the mirror images are produced by plane M , using the confocal camera model (Equation 7) with the impulse response $H(\mathbf{x}_l, \mathbf{x}_s, t)$ to obtain $f_{cc}(\mathbf{x}_v, t)$. The frame at $t = 0$ of $f_{cc}(\mathbf{x}_v, t)$ captures the objects hidden around two corners (Figure 13b, bottom) by leveraging five-bounce specular paths $\mathbf{x}_l \rightarrow M \rightarrow G \rightarrow M \rightarrow \mathcal{S}$ in the computational domain. The resulting images appear blurrier than single-corner reconstructions since the mirror behavior at NLOS imaging frequencies is not perfectly specular (Section 9).

Capture noise. Previous NLOS imaging methods that relied on single-pixel SPAD sensors suffered from low signal-to-noise ratio, requiring long capture times. The implementation of gated SPAD array sensors [Riccardo et al. 2022], which we use in our work, significantly mitigates this issue, and can enable imaging speeds of up to five frames per second [Nam et al. 2021]. While the signal degrades with the number of bounces, in our experiments we observed only minor changes (noise) in our computed images over multiple measurements of the same scene, for the computational wavelengths previously specified. This suggests that our imaging procedures are mainly affected by other limiting factors (e.g., surface

size and reflectance) than by capture noise; in fact, in our work we had to lower the power of the laser to prevent overexposing our SPAD array sensor. Additionally, we compared the photon count of our two-corner experiments (Figure 13) based on the fifth bounce, and their third-bounce counterparts with the target object G placed at the mirror location G' in a single-corner configuration. Under the same exposure time, the total photons captured in our fifth-bounce setups is an order of magnitude higher than their third-bounce counterparts (around 10^9 and 10^8 photons, respectively), which suggests third- and fifth-bounce setups are similar in terms of capture noise when imaging similar regions of the hidden scene.

9 DISCUSSION AND FUTURE WORK

We have established a connection between the surface reflectance defined by well-known wave propagation principles and a wave-based NLOS imaging formulation, showing how diffuse planar surfaces become virtual mirrors at NLOS imaging wavelengths. We have then introduced a procedure to address the missing-cone problem: by analyzing mirror images produced by other virtual mirrors, and then showing how to directly image such surfaces using secondary apertures. Moreover, our insights have allowed us to image objects hidden behind two corners by imaging the space behind virtual mirrors, where we have observed mirror images of objects hidden around two corners.

Mirror reflections under existing imaging methods. In our work, we have showed how to image mirror reflections of different scene elements to address current limitations of NLOS imaging methods. To image objects around two corners, our key idea is to reason about the location of the imaged volume based on our analysis of mirror-like behavior of planar surfaces in the computational domain. The evaluation of the confocal camera model at $t = 0$ (which we use to obtain the image of the object hidden behind two corners) is

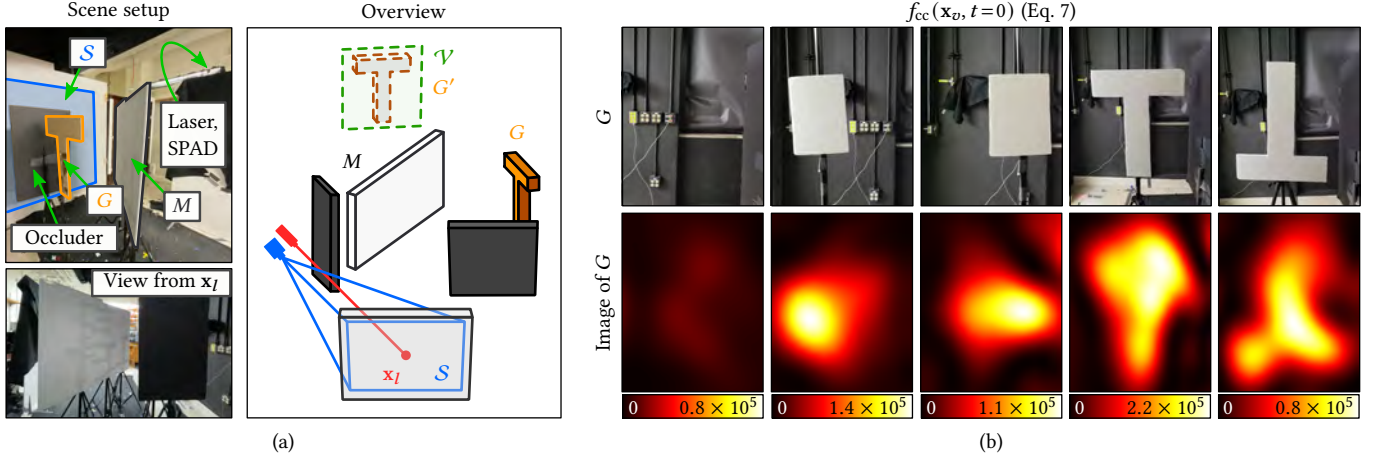


Fig. 13. Results for two-corner imaging (Section 7, similar setup) from real captured data. (a) Photographs and overview of the setup. The geometry G is hidden behind two corners: the relay surface with aperture S , and a diffuse surface M that is a virtual mirror in the computational domain. This diffuse surface M is oriented so that the specular reflection from x_I reaches G , and the specular reflection from G reaches S . The geometry G is not directly visible from x_I or any point in the aperture S as it is covered by an occluder. (b) **Top**: Photographs of the objects hidden around two corners. **Bottom**: Imaging results for different geometries G , one per column, placing the focal plane of the virtual camera at \mathcal{V} . We can clearly identify shape, locations and orientations despite the geometries being hidden after a second corner. Computed using $\lambda_c = 14$ cm and $\sigma = 7.5$ cm in our illumination function.

equivalent to third-bounce imaging used by existing single-corner NLOS imaging methods. We have described our procedure using the wave-based phasor-field formulation of such imaging model. Nevertheless, our methodology and observations could, in principle, generalize to existing single-corner methods that use similar models in order to extend them to two-corner scenes, providing an interesting path for future research.

Fourth-bounce assumptions and higher-order bounces. To address the missing-cone problem, we have used fourth-bounce illumination paths. Looking at Figure 6a, the scene has to meet two conditions to be able to compute an image of G , which is inside the null-reconstruction space of third-bounce methods. First, the hidden scene must contain another surface that is not in the null-reconstruction space for third-bounce methods. In our scene, the surface M has this purpose, which then can be used as a secondary aperture. Second, there must exist a four-bounce path that reaches both the surface M where the secondary aperture is located, and the target surface G . This fourth bounce must be able to reach S when following the specular bounce direction in the computational NLOS imaging domain, else both surfaces would be in the null-reconstruction space of fourth-bounce methods too. Note that imaging surfaces with third-bounce illumination already requires similar assumptions. This could in principle generalize to fifth- or even higher-order bounces, allowing to create additional higher-order apertures to observe further into hidden scenes. A more thorough exploration of the potential of these higher-order apertures is thus an interesting avenue of future work.

Fifth-bounce assumptions and multiple virtual mirrors. We have demonstrated how to image a surface G hidden around two corners from fifth-bounce illumination, using a diffuse surface M as a virtual mirror. For this to work, light from the illuminated point x_I has to follow specular paths in the computational domain that

must reach G after one bounce on M , and must reflect back to the aperture S after another bounce on M , yielding a five-bounce path $x_I \rightarrow M \rightarrow G \rightarrow M \rightarrow S$. Thus, imaging G depends on the location and orientation of M with respect to G and S . Note that this is no different from classic third-bounce NLOS setups, where objects must be located and oriented in regions outside of the null-reconstruction space to be imaged. Our method could in principle generalize to more cluttered scenarios, where specular reflections between different planar surfaces would increase the coverage of NLOS imaging. To explore this, an exhaustive analysis of the connection between imaging wavelength, surface size and features, and their reflectance properties at different imaging frequencies would be necessary.

Coverage of the missing-cone problem using higher-order bounces. In our work we have demonstrated how to image surfaces inside the null-reconstruction space of third-bounce methods. However, as can be seen in Figure 9 (simulated) and Figure 12 (captured), only a part of the surface G reflects fourth-bounce illumination towards S in the computational domain. Thus, some parts of G remain inside the null-reconstruction space of our fourth-bounce imaging procedure. The third-bounce analysis of the missing-cone problem by Liu et al. [2019a] concludes that the visibility of a point in the scene only depends on the position of points on the visible relay surface. This is not the case for the null-reconstruction space of higher-order imaging methods, where the visibility of a point in the scene also depends on other hidden scene elements. Thus, a thorough analysis of the coverage of the missing-cone is an open challenging contribution in NLOS imaging regarded as future work.

Mirror behavior. While our experiments showed that diffuse planar surfaces produce specular reflections at NLOS imaging frequencies, these reflections do not follow exactly a delta function due to diffraction effects. This happens likely because the surfaces we use are not much larger than the wavelength of the computational wave.

Image quality therefore depends significantly on the size and position of the mirror surface. In our experimental results, we observed that reflections through such diffuse surfaces appear blurry, mainly due to diffraction artifacts. Finally, some higher-order paths may have the same time of flight as third-bounce paths, and thus are coupled in the impulse response $H(\mathbf{x}_l, \mathbf{x}_s, t)$ introducing undesired artifacts in the imaging process, which is the case already for all existing NLOS imaging methods. To what extent diffraction, imperfect mirror behavior, noise, and coupling between bounces enter this problem is an interesting topic for future research.

In conclusion, our virtual mirrors framework addresses two of the most limiting problems of current NLOS imaging algorithms, leveraging fourth- and fifth-bounce illumination to compute images of surfaces inside the null-reconstruction space of existing methods, and even hidden behind two corners. We hope that our work spurs further research in this direction to explore the full potential of the field.

ACKNOWLEDGMENTS

We want to thank the anonymous reviewers for their time and insightful comments, and the members of the Graphics and Imaging Lab for their help with the manuscript. Our work was funded by the European Union's European Defense Fund Program through the ENLIGHTEN project under grant agreement No. 101103242, by the Gobierno de Aragón (Departamento de Ciencia, Universidad y Sociedad del Conocimiento) through project BLINDSIGHT (ref. LMP30_21), by MCIN/AEI/10.13039/501100011033 through Project PID2019-105004GB-I00, by the Air Force Office for Scientific Research (FA9550-21-1-0341), and by the National Science Foundation (1846884). Additionally, Diego Royo was supported by a Gobierno de Aragón predoctoral grant.

REFERENCES

- Byeongjoo Ahn, Akshat Dave, Ashok Veeraraghavan, Ioannis Gkioulekas, and Aswin C Sankaranarayanan. 2019. Convolutional approximations to the general non-line-of-sight imaging operator. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 7889–7899.
- Martin Benning, Christoph Brune, Marinus Jan Lagerwerf, and Carola-Bibiane Schönlieb. 2015. TGV sinogram inpainting for limited angle tomography. *Proceedings of the Royal Society A* (2015).
- Katherine L Bouman, Vickie Ye, Adam B Yedidia, Frédo Durand, Gregory W Wornell, Antonio Torralba, and William T Freeman. 2017. Turning corners into cameras: Principles and methods. In *Proceedings of the IEEE International Conference on Computer Vision*. 2270–2278.
- Mauro Buttafava, Jessica Zeman, Alberto Tosi, Kevin Eliceiri, and Andreas Velten. 2015. Non-line-of-sight imaging using a time-gated single photon avalanche diode. *Optics Express* 23, 16 (2015), 20997–21011.
- Ruizhi Cao, Frederic de Goumoens, Baptiste Blochet, Jian Xu, and Changhui Yang. 2022. High-resolution non-line-of-sight imaging employing active focusing. *Nature Photonics* 16, 6 (2022), 462–468.
- Alexander H Delaney and Yoram Bresler. 1998. Globally convergent edge-preserving regularized reconstruction: an application to limited-angle tomography. *IEEE Transactions on Image Processing* 7, 2 (1998), 204–221.
- Justin Dove and Jeffrey H Shapiro. 2019. Paraxial theory of phasor-field imaging. *Optics Express* 27, 13 (2019), 18016–18037.
- Justin Dove and Jeffrey H Shapiro. 2020a. Nonparaxial phasor-field propagation. *Optics Express* 28, 20 (2020), 29212–29229.
- Justin Dove and Jeffrey H Shapiro. 2020b. Paraxial phasor-field physical optics. *Optics Express* 28, 14 (2020), 21095–21109.
- Justin Dove and Jeffrey H Shapiro. 2020c. Speckled speckled speckle. *Optics Express* 28, 15 (2020), 22105–22120.
- Daniele Faccio, Andreas Velten, and Gordon Wetzstein. 2020. Non-line-of-sight imaging. *Nature Reviews Physics* 2, 6 (2020), 318–327.
- Genevieve Garipey, Francesco Tonolini, Robert Henderson, Jonathan Leach, and Daniele Faccio. 2016. Detection and tracking of moving objects hidden from view. *Nature Photonics* 10, 1 (2016), 23–26.
- Ibón Guillén, Xiaochun Liu, Andreas Velten, Diego Gutierrez, and Adrian Jarabo. 2020. On the Effect of Reflectance on Phasor Field Non-Line-of-Sight Imaging. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 9269–9273.
- Felix Heide, Matthias B Hullin, James Gregson, and Wolfgang Heidrich. 2013. Low-budget transient imaging using photonic mixer devices. *ACM Transactions on Graphics* 32, 4 (2013), 1–10.
- Adrian Jarabo, Julio Marco, Adolfo Muñoz, Raul Buisan, Wojciech Jarosz, and Diego Gutierrez. 2014. A framework for transient rendering. *ACM Transactions on Graphics* 33, 6 (2014), 1–10.
- Adrian Jarabo, Belen Masia, Julio Marco, and Diego Gutierrez. 2017. Recent advances in transient imaging: A computer graphics and vision perspective. *Visual Informatics* 1, 1 (2017), 65–79.
- Masaki Kaga, Takahiro Kushida, Tsuyoshi Takatani, Kenichiro Tanaka, Takuya Funatomi, and Yasuhiro Mukaigawa. 2019. Thermal non-line-of-sight imaging from specular and diffuse reflections. *IPJ Transactions on Computer Vision and Applications* 11, 1 (2019), 1–6.
- Ori Katz, Pierre Heidmann, Mathias Fink, and Sylvain Gigan. 2014. Non-invasive single-shot imaging through scattering layers and around corners via speckle correlations. *Nature Photonics* 8, 10 (2014), 784–790.
- William Krska, Sheila W Seidel, Charles Saunders, Robinson Czajkowski, Christopher Yu, John Murray-Bruce, and Vivek Goyal. 2022. Double your corners, double your fun: the doorway camera. In *Proceedings of the IEEE International Conference on Computational Photography*. IEEE, 1–12.
- Marco La Manna, Ji-Hyun Nam, Syed Azer Reza, and Andreas Velten. 2020. Non-line-of-sight-imaging using dynamic relay surfaces. *Optics Express* 28, 4 (2020), 5331–5339.
- Martin Laurenzis and Frank Christnacher. 2022. Time domain analysis of photon scattering and Huygens-Fresnel back projection. *Optics Express* 30, 17 (2022), 30441–30454.
- Zhengpeng Liao, Deyang Jiang, Xiaochun Liu, Andreas Velten, Yajun Ha, and Xin Lou. 2021. FPGA Accelerator for Real-Time Non-Line-of-Sight Imaging. *IEEE Transactions on Circuits and Systems I: Regular Papers* 69, 2 (2021), 721–734.
- David B Lindell, Gordon Wetzstein, and Vladlen Koltun. 2019a. Acoustic non-line-of-sight imaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 6780–6789.
- David B Lindell, Gordon Wetzstein, and Matthew O'Toole. 2019b. Wave-based non-line-of-sight imaging using fast f - k migration. *ACM Transactions on Graphics* 38, 4 (2019), 1–13.
- Xiaochun Liu, Sebastian Bauer, and Andreas Velten. 2019a. Analysis of feature visibility in non-line-of-sight measurements. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 10140–10148.
- Xiaochun Liu, Sebastian Bauer, and Andreas Velten. 2020. Phasor field diffraction based reconstruction for fast non-line-of-sight imaging systems. *Nature Communications* 11, 1 (2020), 1645.
- Xiaochun Liu, Ibón Guillén, Marco La Manna, Ji Hyun Nam, Syed Azer Reza, Toan Huu Le, Adrian Jarabo, Diego Gutierrez, and Andreas Velten. 2019b. Non-line-of-sight imaging using phasor-field virtual wave optics. *Nature* 572, 7771 (2019), 620–623.
- Pablo Luesia, Miguel Crespo, Adrian Jarabo, and Albert Redo-Sanchez. 2022. Non-line-of-sight imaging in the presence of scattering media using phasor fields. *Optics Letters* 47, 15 (2022), 3796–3799.
- Pablo Luesia-Lahoz, Diego Gutierrez, and Adolfo Muñoz. 2023. Zone Plate Virtual Lenses for Memory-Constrained NLOS Imaging. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE.
- Tomohiro Maeda, Guy Satat, Tristan Swedish, Lagnojita Sinha, and Ramesh Raskar. 2019a. Recent Advances in Imaging Around Corners. *arXiv preprint arXiv:1910.05613* (2019).
- Tomohiro Maeda, Yiqin Wang, Ramesh Raskar, and Achuta Kadambi. 2019b. Thermal non-line-of-sight imaging. In *Proceedings of the IEEE International Conference on Computational Photography*. IEEE, 1–11.
- Julio Marco, Adrian Jarabo, Ji Hyun Nam, Xiaochun Liu, Miguel Ángel Cosculluela, Andreas Velten, and Diego Gutierrez. 2021. Virtual light transport matrices for non-line-of-sight imaging. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2440–2449.
- Ji Hyun Nam, Eric Brandt, Sebastian Bauer, Xiaochun Liu, Marco Renner, Alberto Tosi, Eftychios Sifakis, and Andreas Velten. 2021. Low-latency time-of-flight non-line-of-sight imaging at 5 frames per second. *Nature Communications* 12, 1 (2021), 1–10.
- Ji Hyun Nam, Eric Brandt, Sebastian Bauer, Xiaochun Liu, Eftychios Sifakis, and Andreas Velten. 2020. Real-time Non-line-of-Sight imaging of dynamic scenes. *arXiv preprint arXiv:2010.12737* (2020).

- Matthew O'Toole, David B Lindell, and Gordon Wetzstein. 2018. Confocal non-line-of-sight imaging based on the light-cone transform. *Nature* 555, 7696 (2018), 338.
- Adithya Pediredla, Akshat Dave, and Ashok Veeraraghavan. 2019. Snlos: Non-line-of-sight scanning through temporal focusing. In *Proceedings of the IEEE International Conference on Computational Photography*. IEEE, 1–13.
- Syed Azer Reza, Marco La Manna, Sebastian Bauer, and Andreas Velten. 2019a. Phasor field waves: A Huygens-like light transport model for non-line-of-sight imaging applications. *Optics Express* 27, 20 (2019), 29380–29400.
- Syed Azer Reza, Marco La Manna, Sebastian Bauer, and Andreas Velten. 2019b. Phasor field waves: experimental demonstrations of wave-like properties. *Optics Express* 27, 22 (2019), 32587–32608.
- Simone Riccardo, Enrico Conca, Vincenzo Sesta, Andreas Velten, and Alberto Tosi. 2022. Fast-Gated 16×16 SPAD Array With 16 on-Chip 6 ps Time-to-Digital Converters for Non-Line-of-Sight Imaging. *IEEE Sensors Journal* 22, 17 (2022), 16874–16885.
- Diego Royo, Jorge García, Adolfo Muñoz, and Adrian Jarabo. 2022. Non-line-of-sight transient rendering. *Computers & Graphics* 107 (2022), 84–92.
- Charles Saunders, John Murray-Bruce, and Vivek K Goyal. 2019. Computational periscopy with an ordinary digital camera. *Nature* 565, 7740 (2019), 472–475.
- Nicolas Scheiner, Florian Kraus, Fangyin Wei, Buu Phan, Fahim Mannan, Nils Appenrodt, Werner Ritter, Jurgen Dickmann, Klaus Dietmayer, Bernhard Sick, et al. 2020. Seeing around street corners: Non-line-of-sight detection and tracking in-the-wild using doppler radar. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2068–2077.
- Donggeek Shin, Feihu Xu, Dheera Venkatraman, Rudi Lussana, Federica Villa, Franco Zappa, Vivek K Goyal, Franco NC Wong, and Jeffrey H Shapiro. 2016. Photon-efficient imaging with a single-photon camera. *Nature Communications* 7, 1 (2016), 12046.
- Jeremy A Teichman. 2019. Phasor field waves: a mathematical treatment. *Optics Express* 27, 20 (2019), 27500–27506.
- Andreas Velten, Thomas Willwacher, Otakrist Gupta, Ashok Veeraraghavan, Mouni G Bawendi, and Ramesh Raskar. 2012a. Recovering three-dimensional shape around a corner using ultrafast time-of-flight imaging. *Nature Communications* 3, 1 (2012), 745.
- Andreas Velten, Di Wu, Adrian Jarabo, Belen Masia, Christopher Barsi, Chinmaya Joshi, Everett Lawson, Mouni G Bawendi, Diego Gutierrez, and Ramesh Raskar. 2013. Femto-Photography: Capturing and Visualizing the Propagation of Light. *ACM Transactions on Graphics* 32, 4 (2013), 1–8.
- Andreas Velten, Di Wu, Adrian Jarabo, Belen Masia, Christopher Barsi, Everett Lawson, Chinmaya Joshi, Diego Gutierrez, Mouni G. Bawendi, and Ramesh Raskar. 2012b. Relativistic ultrafast rendering using time-of-flight imaging. In *ACM SIGGRAPH 2012 Talks*.
- Cheng Wu, Jianjiang Liu, Xin Huang, Zheng-Ping Li, Chao Yu, Jun-Tian Ye, Jun Zhang, Qiang Zhang, Xiankang Dou, Vivek K Goyal, et al. 2021. Non-line-of-sight imaging over 1.43 km. *Proceedings of the National Academy of Sciences* 118, 10 (2021).
- Shumian Xin, Sotiris Nouisias, Kiriakos N Kutulakos, Aswin C Sankaranarayanan, Srinivasa G Narasimhan, and Ioannis Gkioulekas. 2019. A theory of Fermat paths for non-line-of-sight shape reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 6800–6809.
- Shinyoung Yi, Donggun Kim, Kiseok Choi, Adrian Jarabo, Diego Gutierrez, and Min H Kim. 2021. Differentiable Transient Rendering. *ACM Transactions on Graphics* 40, 6 (2021), 1–11.