

Non-line-of-sight imaging using phasor-field virtual wave optics

Xiaochun Liu¹, Ibón Guillén², Marco La Manna³, Ji Hyun Nam¹, Syed Azer Reza³, Toan Huu Le¹, Adrian Jarabo², Diego Gutierrez² & Andreas Velten^{1,3*}

Non-line-of-sight imaging allows objects to be observed when partially or fully occluded from direct view, by analysing indirect diffuse reflections off a secondary relay surface. Despite many potential applications^{1–9}, existing methods lack practical usability because of limitations including the assumption of single scattering only, ideal diffuse reflectance and lack of occlusions within the hidden scene. By contrast, line-of-sight imaging systems do not impose any assumptions about the imaged scene, despite relying on the mathematically simple processes of linear diffractive wave propagation. Here we show that the problem of non-line-of-sight imaging can also be formulated as one of diffractive wave propagation, by introducing a virtual wave field that we term the phasor field. Non-line-of-sight scenes can be imaged from raw time-of-flight data by applying the mathematical operators that model wave propagation in a conventional line-of-sight imaging system. Our method yields a new class of imaging algorithms that mimic the capabilities of line-of-sight cameras. To demonstrate our technique, we derive three imaging algorithms, modelled after three different line-of-sight systems. These algorithms rely on solving a wave diffraction integral, namely the Rayleigh–Sommerfeld diffraction integral. Fast solutions to Rayleigh–Sommerfeld diffraction and its approximations are readily available, benefiting our method. We demonstrate non-line-of-sight imaging of complex scenes with strong multiple scattering and ambient light, arbitrary materials, large depth range and occlusions. Our method handles these challenging cases without explicitly inverting a light-transport model. We believe that our approach will help to unlock the potential of non-line-of-sight imaging and promote the development of relevant applications not restricted to laboratory conditions.

We have recently witnessed considerable advances in transient imaging techniques¹⁰ that use streak cameras¹¹, gated sensors⁶, amplitude-modulated continuous waves¹², single-photon avalanche diodes (SPADs)¹³ or interferometry¹⁴. Access to time-resolved image information has led to advances in imaging of objects partially or fully hidden from direct view^{1–3,5–7,15–18}: that is, non-line-of-sight (NLOS) imaging. Other methods are able to use information encoded in the phase of continuous light and do not use the time of flight⁴. In the basic configuration of an NLOS system, light bounces off a relay wall, travels to the hidden scene, then propagates back to the relay wall and finally reaches the sensor.

Recent NLOS reconstruction methods are based on heuristic filtered backprojection^{2,3,6,7,19} or attempt to compute inverse operators of simplified forward light transport models^{5,9,20}. These simplified models do not take into account multiple scattering, surfaces with anisotropic reflectance or, with a few exceptions²⁰, occlusions or clutter in the hidden scene. The depth range that can be recovered is also limited, partially owing to the difference in intensity between first- and higher-order reflections. Existing methods are thus limited to carefully controlled cases, imaging isolated objects of simple geometry with moderate or no occlusion. Whereas the goal of previous works has

been limited to the reconstruction of hidden geometry, we develop a theoretical framework for general NLOS imaging, reconstructing the irradiance at a virtual sensor; this enables applications beyond geometric reconstruction.

Time-of-flight LOS imaging has used a phasor formalism (a phasor, or phase vector, is a complex number representing properties of a light wave) together with Fourier domain ranging¹² to describe the emitted modulated light signal. Kadambi et al.²¹ extended this concept to reconstruct NLOS scenes by using a phasor model along with a non-line-of-sight capture system that uses intensity-modulated light sources and gain-modulated detection. We show that a similar description can be used to model the physics of light transport through the scene. The key insight is that propagation through a scene of intensity-modulated light can be modelled using a Rayleigh–Sommerfeld diffraction (RSD) operator acting on a quantity that we term the phasor field. This allows us to formulate any NLOS imaging problem as a wave imaging problem (Fig. 1) and to transfer well-established insights and techniques from classic optics into the NLOS domain. Given a captured time-resolved dataset of light transport through an NLOS scene, and a choice of a template LOS imaging system, our method provides a recipe that results in an NLOS imaging algorithm mimicking the capabilities of the corresponding LOS system. This template system can be any real or hypothetical wave imaging system that includes a set of light sources and detectors. The resulting algorithms can then be efficiently solved using diffraction integrals such as the RSD, for which various fast exact and approximate solvers exist²². Supplementary Information section A illustrates this.

We start by mathematically defining our phasor field $\mathcal{P}(\mathbf{x}, t)$. Let $\mathcal{E}(\mathbf{x}, t)$ (with units $\sqrt{\text{W m}^{-2}}$) be a quasi-monochromatic scalar field at position $\mathbf{x} \in \mathcal{S}$ and time t , incident on (or reflected from) a Lambertian surface \mathcal{S} , with centre frequency Ω_0 and bandwidth $\Delta\Omega \ll \Omega_0$. We can then define

$$\mathcal{P}(\mathbf{x}, t) \equiv \left\langle \frac{1}{\tau} \int_{t-\tau/2}^{t+\tau/2} |\mathcal{E}(\mathbf{x}, t')|^2 dt' \right\rangle - \left\langle \frac{1}{T} \int_{t-T/2}^{t+T/2} |\mathcal{E}(\mathbf{x}, t')|^2 dt' \right\rangle \quad (1)$$

as the mean subtracted irradiance (in watts per square metre) at point \mathbf{x} and time t . The $\langle \cdot \rangle$ operator denotes spatial speckle averaging (for the reflected case) accounting for laser illumination, and τ represents the averaging of the intensity at a fast detector, with $\tau \ll 1/\Delta\Omega \ll T$. The second integral in equation (1) is a long-term average intensity over an interval $T \gg \tau$ of the signal as seen by a conventional non-transient photodetector. Now, let us define the Fourier component of $\mathcal{P}(\mathbf{x}, t)$ for frequency ω as

$$\mathcal{P}_{0,\omega}(\mathbf{x}) \equiv \int_{-\infty}^{+\infty} \mathcal{P}(\mathbf{x}, t) e^{-i\omega t} dt \quad (2)$$

¹Department of Electrical and Computer Engineering, University of Wisconsin Madison, Madison, WI, USA. ²Graphics and Imaging Lab, Universidad de Zaragoza—I3A, Zaragoza, Spain. ³Department of Biostatistics and Medical Informatics, University of Wisconsin Madison, Madison, WI, USA. *e-mail: velten@wisc.edu

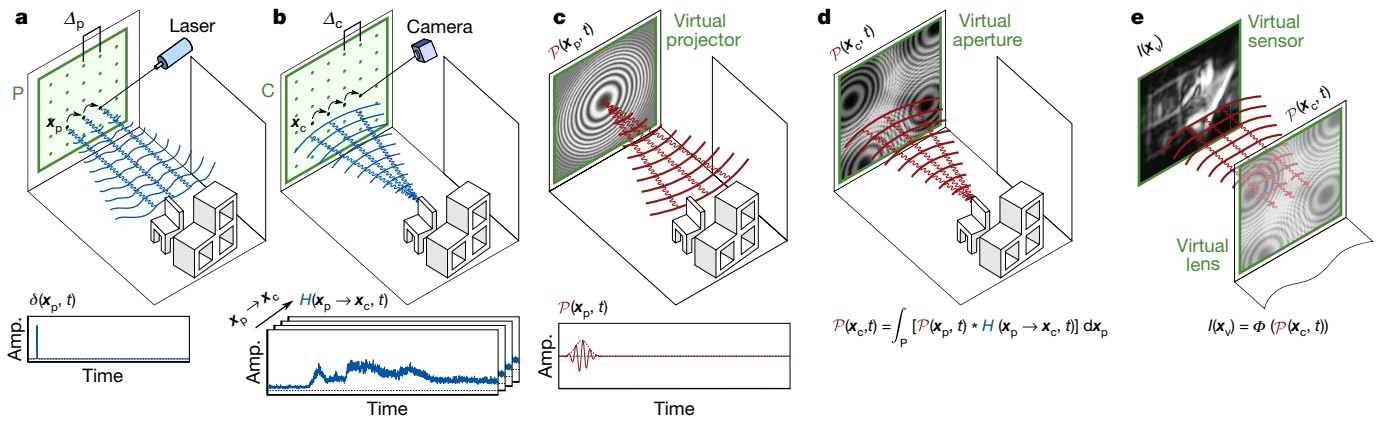


Fig. 1 | NLOS as a virtual LOS imaging system. **a, b,** Capturing scene data. **a,** A pulsed laser sequentially scans a relay wall (green); **b,** the light reflected back from the scene onto the wall is recorded at the sensor, yielding an impulse response H of the scene. **c,** Virtual light source. The phasor-field wave of a virtual light source $\mathcal{P}(x_p, t)$ is modelled after the

wavefront of the light source of the template LOS system. **d,** The scene response to this virtual illumination $\mathcal{P}(x_c, t)$ is computed using H . **e,** The scene is reconstructed from the wavefront $\mathcal{P}(x_c, t)$ using wave diffraction theory. The function $\Phi(\cdot)$ is also taken from the template LOS system. Amp., phasor-field amplitude.

from which we can define a monochromatic component of the phasor field $\mathcal{P}_\omega(\mathbf{x}, t)$ as

$$\mathcal{P}_\omega(\mathbf{x}, t) \equiv \mathcal{P}_{0,\omega}(\mathbf{x}) e^{i\omega t} \quad (3)$$

Using the above, our phasor field $\mathcal{P}(\mathbf{x}, t)$ can be expressed as a superposition of monochromatic plane waves as $\mathcal{P}(\mathbf{x}, t) = \int_{-\infty}^{+\infty} \mathcal{P}_\omega(\mathbf{x}, t) d\omega/2\pi$. Since $\mathcal{P}(\mathbf{x}, t)$ is a real quantity, the Fourier components $\mathcal{P}_{0,\omega}(\mathbf{x})$ are complex and symmetric about $\omega = 0$. Note that, in many places in this Letter, we assign $\mathcal{P}(\mathbf{x}, t)$ an explicitly complex value; in these cases, it is implied that the correct real representation is $\frac{1}{2}(\mathcal{P}(\mathbf{x}, t) + \mathcal{P}^*(\mathbf{x}, t))$. In practice, the complex conjugate can be safely ignored in our calculations. As can be seen in Supplementary Information section B, given an isotropic source plane S and a destination plane D , and assuming that the electric field at S is incoherent, the propagation of its monochromatic component $\mathcal{P}_\omega(\mathbf{x}, t)$ is defined by an RSD-like propagation integral:

$$\mathcal{P}_\omega(\mathbf{x}_d, t) = \gamma \int_S \mathcal{P}_\omega(\mathbf{x}_s, t) \frac{e^{ik|\mathbf{x}_d - \mathbf{x}_s|}}{|\mathbf{x}_d - \mathbf{x}_s|} d\mathbf{x}_s \quad (4)$$

where γ is an attenuation factor, $k = 2\pi/\lambda$ is the wavenumber for wavelength $\lambda = 2\pi/\omega$, $\mathbf{x}_s \in S$ and $\mathbf{x}_d \in D$. Note that, as described in Supplementary Information section B, we approximate γ as a constant over the plane S as $\gamma \approx 1/|\langle S \rangle - \mathbf{x}_d|$; this approximation has a minor effect on the signal amplitude at the sensor but does not change the phase of our phasor field. Although equation (4) is defined for monochromatic signals, it can be used to propagate broadband signals by propagating each monochromatic component independently; this can be efficiently done by time-shifting the phasor field (more details are provided in Supplementary Information section B.1).

The key insight of equation (4) is that, given the assumption that γ is a constant, the propagation of our phasor field is defined by the same RSD operator as any other physical wave. Therefore, to image a scene from a virtual camera with aperture on plane C , we can apply the image formation model of any wave-based LOS imaging system directly over the phasor field $\mathcal{P}(x_c, t)$ at the aperture, with $\mathbf{x}_c \in C$. The challenge is how to compute $\mathcal{P}(x_c, t)$ from an illuminating input phasor field $\mathcal{P}(x_p, t)$, where \mathbf{x}_p is a point in the virtual projector aperture P , given a particular NLOS scene (see Fig. 1).

Because light transport is linear in space and time-invariant^{23,24}, we can characterize light transport through the scene as an impulse

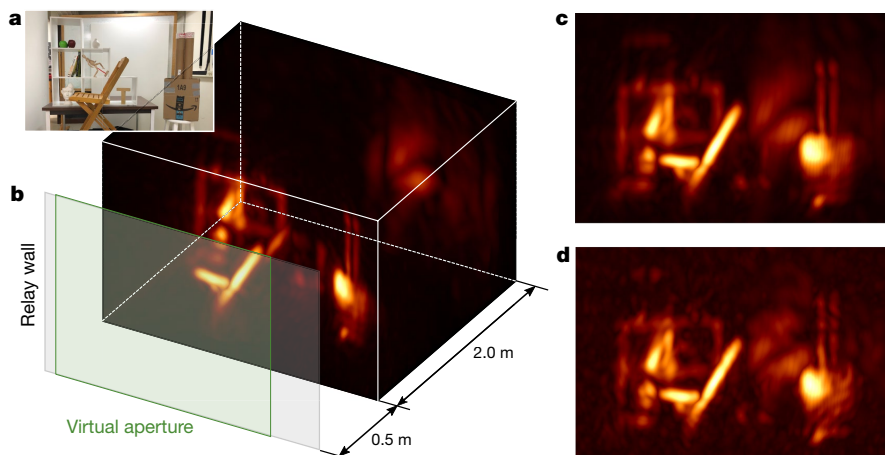


Fig. 2 | Reconstructions of a complex NLOS scene. **a,** Photograph of the scene as seen from the relay wall. The scene contains occluding geometries, with objects towards the front (such as the chair) partially occluding the objects further back; multiple anisotropic surface reflectances; large depth; and strong ambient and multiply scattered light.

b, 3D visualization of the reconstruction with phasor fields ($\lambda = 6$ cm). We include the relay wall location and the coverage of the virtual aperture for illustrative purposes. **c,** Frontal view of the scene, captured with an exposure time of 10 ms per laser position. **d,** Frontal view captured with an exposure time of just 1 ms (24 s for the complete scan).




Fig. 3 | Robustness of our technique. **a**, Reconstruction in the presence of strong ambient illumination (all the lights on during capture). **b**, Hidden scene with a large depth range, leading to very weak signals from objects farther away.

response function $H(\mathbf{x}_p \rightarrow \mathbf{x}_c, t)$, where \mathbf{x}_p and \mathbf{x}_c are the positions of the emitter and detector, respectively. The phasor field at the virtual aperture $\mathcal{P}(\mathbf{x}_c, t)$ can thus be expressed as a function of the input phasor field $\mathcal{P}(\mathbf{x}_p, t)$ and $H(\mathbf{x}_p \rightarrow \mathbf{x}_c, t)$:

$$\mathcal{P}(\mathbf{x}_c, t) = \int_{\mathbf{P}} [\mathcal{P}(\mathbf{x}_p, t) * H(\mathbf{x}_p \rightarrow \mathbf{x}_c, t)] d\mathbf{x}_p \quad (5)$$

where $*$ denotes the convolution operator. Any imaging system can be characterized by its image formation function $\Phi(\cdot)$, which transduces the incoming field into an image

$$I(\mathbf{x}_v) = \Phi(\mathcal{P}(\mathbf{x}_c, t)) \quad (6)$$

where \mathbf{x}_v is the point being imaged (that is, the point at the virtual sensor). This, in turn, can be formulated as an RSD propagator, requiring a diffraction integral to be solved to generate the final image.

In an NLOS scenario, $H(\mathbf{x}_p \rightarrow \mathbf{x}_c, t)$ usually corresponds to five-dimensional transients acquired by an ultrafast sensor focused on \mathbf{x}_c and sequentially illuminating the relay wall with short pulses at different points \mathbf{x}_p (see Fig. 1 and Methods). Points \mathbf{x}_p and \mathbf{x}_c correspond to a virtual LOS imaging system projected on the relay wall. Once $H(\mathbf{x}_p \rightarrow \mathbf{x}_c, t)$ has been captured, both the wavefront $\mathcal{P}(\mathbf{x}_p, t)$ and the imaging operator $\Phi(\cdot)$ can be implemented computationally, so they are not bounded by hardware limitations. We can leverage this to use different $\mathcal{P}(\mathbf{x}_p, t)$ functions from any existing LOS imaging system²⁵ to emulate its characteristics in an NLOS setting.

We illustrate the robustness and versatility of our method by implementing three virtual NLOS imaging systems based on common LOS techniques: a conventional photography camera capable of imaging NLOS scenes without knowledge of the timing or location of the illumination source; a transient photography system capable of capturing transient videos of the hidden scene revealing higher-order interreflections (multiple light bounces between surface elements) beyond third bounce; and a confocal time-gated imaging system robust to interreflections. An in-depth description of these example imaging systems is provided in Supplementary Information section C, including their corresponding $\mathcal{P}(\mathbf{x}_p, t)$ functions and imaging operators, and section D describes some examples of practical integral solvers.

The spatial resolution of our virtual camera is $\Delta_x = 0.61 \lambda L/d$, where d is the virtual aperture diameter and L is the imaging distance. The distance Δ_p between sample points \mathbf{x}_p in \mathbf{P} (see Fig. 1) has to be small enough to sample H at the phasor-field wavelength. We fix $\Delta_p = 1$ cm and, unless stated otherwise, $\lambda = 4$ cm. The minimum sampling rate is $\Delta_p < \lambda/2$; in practice, we found $\Delta_p = \lambda/4$ to provide the best trade-off between reconstruction noise and resolution.

The computational cost of our algorithm is bounded by the RSD solver computing the image formation model $\Phi(\cdot)$. Fast diffraction integral solvers exist²², with complexity $O(N^3 \log N)$. For the particular case of our confocal system, we formulate the algorithm as a backprojection (see Supplementary Information section D.2 for details), and so we are bounded by the computational cost of the backprojection algorithm used.




Fig. 4 | Additional NLOS imaging applications of our method. **a**, NLOS refocusing. The hidden letters (left) are progressively brought in and out of focus as seen from a virtual photography camera at the relay wall, using the exact lens integral (blue border), and the faster Fresnel approximation (red border). **b**, NLOS transient video. Example frames of light travelling

through a hidden office scene when illuminated by a pulsed laser. Timestamps indicate the propagation time from the relay wall. Frames with a green border show third-bounce objects, frames with an orange border show fourth- and fifth-bounce effects.

One common application of NLOS imaging is the reconstruction of hidden geometry. Figure 2 shows the result for a complex scene imaged with our virtual confocal camera. This challenging scene contains multiple objects with occlusions distributed over a large depth, a wide range of surface reflectances and albedos, and strong interreflections. Our method is able to image many details of the scene, at the correct depths, even with an ultra-short (1 ms) exposure. More analysis on the robustness of our method to capture noise can be found in the Methods. For simpler scenes (no occlusions, limited depth, controlled reflectance and no interreflections), our method yields results on par with current techniques, which already approach theoretical limits for reconstruction quality (see Methods).

In Fig. 3, we demonstrate the robustness of our method when dealing with other challenging scenarios, including strong multiple scattering and ambient illumination (Fig. 3a), or a high dynamic range from objects spanning a large range of depths (Fig. 3b). Finally, our method allows new NLOS imaging systems and applications to be implemented, making use of the wealth of tools and processing methods available in LOS imaging. Figure 4a demonstrates NLOS refocusing with our virtual photography camera, computed using both the exact RSD operator and a faster Fresnel approximation, while Fig. 4b shows frames of NLOS femto-photography reconstructed using our virtual transient photography system, revealing fourth- and fifth-bounce components in the scene. The first, second and fourth frames, in green, show how light first illuminates the chair, then propagates to the shelf and finally hits the back wall 3 m away. The frames in orange show higher-order bounces. The third frame shows that the chair is illuminated again by light bouncing back from the relay wall, and the last two frames show how the pulse of light travels from the wall back to the scene (see Supplementary Video 1). A description of the Fresnel approximation to the RSD operator, as well as the LOS projector-camera functions used in these examples, appear in Supplementary Information sections D.1 and C.2.

In the Methods, we include comparisons against ground truth for two synthetic scenes, inside a corridor of $2\text{ m} \times 2\text{ m} \times 3\text{ m}$ to create interreflections, simulated using an open-source transient renderer²⁶; these scenes are included in a publicly available database²⁷. We analyse the robustness of our method with and without such interreflections; the reconstruction mean square error (MSE) does not increase, remaining below 5 mm. Finally, we progressively vary the specularity of the hidden geometry, from purely Lambertian to highly specular; again, the quality of the reconstructions does not vary significantly (MSE of about 2 mm).

The examples shown highlight the primary benefit of our approach. By turning NLOS into a virtual LOS system, the intrinsic limitations of previous approaches no longer apply, enabling a class of NLOS imaging methods that take advantage of existing wave-based imaging methods. Formulating NLOS light propagation as a wave does not impose limitations on the types of problems that can be addressed, nor the datasets that can be used. Any signal can be represented as a superposition of phasor-field waves; our formulation can thus be viewed as a choice of basis to represent any kind of NLOS data. Expressing the NLOS problem this way allows a direct analogy to LOS imaging, which can be exploited to derive suitable imaging algorithms and to implement them efficiently.

We have shown three imaging algorithms derived from our method. Our results include more complex scenes than in NLOS reconstructions shown so far in the literature, as well as new applications. In addition, our approach is flexible, fast, memory-efficient and lacks computational complexity since it does not require inverting a light transport model. We anticipate that it can be applied to other LOS imaging systems, for instance to separate light transport into direct and global components, or to use the phase of \mathcal{P}_ω for enhanced depth resolution. Our virtual imaging system could also be used to create a virtual imaging system to see around two corners, assuming the presence of a secondary relay Lambertian surface in the hidden scene, or to select and manipulate individual light paths to isolate specific aspects of the light transport in

different NLOS scenes. In that context, combining our theory with light transport inversions, via, for example, an iterative approach, could potentially lead to better results and is an interesting avenue for future work.

Online content

Any methods, additional references, Nature Research reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at <https://doi.org/10.1038/s41586-019-1461-3>.

Received: 18 October 2018; Accepted: 21 May 2019;

Published online 5 August 2019.

- Kirmani, A., Hutchison, T., Davis, J. & Raskar, R. Looking around the corner using ultrafast transient imaging. *Int. J. Comput. Vis.* **95**, 13–28 (2011).
- Gupta, O., Willwacher, T., Velten, A., Veeraraghavan, A. & Raskar, R. Reconstruction of hidden 3D shapes using diffuse reflections. *Opt. Express* **20**, 19096–19108 (2012).
- Velten, A. et al. Recovering three-dimensional shape around a corner using ultrafast time-of-flight imaging. *Nat. Commun.* **3**, 745 (2012).
- Katz, O., Small, E. & Silberberg, Y. Looking around corners and through thin turbid layers in real time with scattered incoherent light. *Nat. Photon.* **6**, 549–553 (2012).
- Heide, F., Xiao, L., Heidrich, W. & Hullin, M. B. Diffuse mirrors: 3D reconstruction from diffuse indirect illumination using inexpensive time-of-flight sensors. In *IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 3222–3229 (IEEE, 2014).
- Laurenzis, M. & Velten, A. Nonline-of-sight laser gated viewing of scattered photons. *Opt. Eng.* **53**, 023102 (2014).
- Buttafava, M., Zeman, J., Tosi, A., Eliceiri, K. & Velten, A. Non-line-of-sight imaging using a time-gated single photon avalanche diode. *Opt. Express* **23**, 20997–21011 (2015).
- Arellano, V., Gutierrez, D. & Jarabo, A. Fast back-projection for non-line of sight reconstruction. *Opt. Express* **25**, 11574–11583 (2017).
- O’Toole, M., Lindell, D. B. & Wetzstein, G. Confocal non-line-of-sight imaging based on the light-cone transform. *Nature* **555**, 338–341 (2018).
- Jarabo, A., Masia, B., Marco, J. & Gutierrez, D. Recent advances in transient imaging: a computer graphics and vision perspective. *Visual Informatics* **1**, 65–79 (2017).
- Velten, A. et al. Femto-photography: capturing and visualizing the propagation of light. *ACM Trans. Graph.* **32**, 44 (2013).
- Gupta, M., Nayar, S. K., Hullin, M. B. & Martin, J. Phasor imaging: a generalization of correlation-based time-of-flight imaging. *ACM Trans. Graph.* **34**, 156 (2015).
- O’Toole, M. et al. Reconstructing transient images from single-photon sensors. In *2017 IEEE Int. Conf. Computational Photography (CVPR)*, 1539–1547 (IEEE, 2017).
- Gkioulekas, I., Levin, A., Durand, F. & Zickler, T. Micron-scale light transport decomposition using interferometry. *ACM Trans. Graph.* **34**, 37 (2015).
- Xin, S. et al. A theory of Fermat paths for non-line-of-sight shape reconstruction. In *IEEE Int. Conf. Computer Vision and Pattern Recognition (CVPR)*, 6800–6809 (IEEE, 2019).
- Tsai, C., Sankaranarayanan, A. & Gkioulekas, I. Beyond volumetric albedo a surface optimization framework for non-line-of-sight imaging. In *IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 1545–1555 (IEEE, 2019).
- Liu, X., Bauer, S. & Velten, A. Analysis of feature visibility in non-line-of-sight measurements. In *IEEE Intl Conf. Computer Vision and Pattern Recognition (CVPR)* 10140–10148 (IEEE, 2019).
- Wu, R. et al. Adaptive polarization-difference transient imaging for depth estimation in scattering media. *Opt. Lett.* **43**, 1299–1302 (2018).
- Laurenzis, M. & Velten, A. Feature selection and back-projection algorithms for nonline-of-sight laser-gated viewing. *J. Electron. Imaging* **23**, 063003 (2014).
- Heide, F. et al. Non-line-of-sight imaging with partial occluders and surface normals. *ACM Trans. Graph.* **38**, 22 (2019).
- Kadambi, A., Zhao, H., Shi, B. & Raskar, R. Occluded imaging with time-of-flight sensors. *ACM Trans. Graph.* **35**, 15 (2016).
- Shen, F. & Wang, A. Fast-Fourier-transform based numerical integration method for the Rayleigh–Sommerfeld diffraction formula. *Appl. Opt.* **45**, 1102–1110 (2006).
- Sen, P. et al. Dual photography. *ACM Trans. Graph.* **24**, 745–755 (2005).
- O’Toole, M. et al. Temporal frequency probing for 5D transient analysis of global light transport. *ACM Trans. Graph.* **33**, 87 (2014).
- Goodman, J. *Introduction to Fourier Optics* 3rd edn (Roberts, 2005).
- Jarabo, A. et al. A framework for transient rendering. *ACM Trans. Graph.* **33**, 177 (2014).
- Galindo, M. et al. A dataset for benchmarking time-resolved non-line-of-sight imaging. In *IEEE Intl Conf Computational Photography* <https://graphics.unizar.es/nlos> (IEEE, 2019).

Publisher’s note: Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

© The Author(s), under exclusive licence to Springer Nature Limited 2019

METHODS

Details on data acquisition. Hardware configuration. Our capture system, shown in Extended Data Fig. 1, consists of a Onefive Katana HP amplified diode laser (1 W at 532 nm, and a pulse width of about 35 ps used at a repetition rate of 10 MHz) and a gated SPAD detector processed by a time-correlated single-photon counter (PicoQuant HydraHarp), with a time resolution of about 30 ps and a dead time of 100 ns. Two additional charge-coupled device (CCD) cameras are used to calibrate the laser's position. The measured time resolution of our system is approximately 65 ps, a combination of the pulse width of the laser and the time jitter of the system.

NLOS measurement geometry. We obtain an impulse response function $H(\mathbf{x}_p \rightarrow \mathbf{x}_c, t)$ of the scene by sequentially illuminating points \mathbf{x}_p on the relay wall with a short pulse and detecting the signal returning at points \mathbf{x}_c .

Our hardware device is located 2.5 m from the relay wall, with the NLOS scenes hidden from direct view. The field of view is 25°. The walls are made of standard white styrofoam. The scanning area in the relay wall (virtual camera) is 1.8 m \times 1.3 m, with laser points \mathbf{x}_p spaced by $\Delta_p = 1$ cm in each direction. The SPAD is focused at a position near the centre of the grid. We avoid scanning a small square region around the SPAD focused position (the confocal position) because the signal becomes very noisy at this location. Figures 2, 3 provide additional details for the specific scenes shown.

Exposure time. Our capture set-up includes CCD cameras (Extended Data Fig. 1) to confirm the 3D position of every laser during the measurement; these are a limiting factor in the speed of our experiments. Because the capture process runs in parallel, we use a very long (1 s) exposure time per laser position for some datasets. They are used for all results unless otherwise specified. In addition, we capture scenes without the additional CCD photographs that can be collected much faster and with much shorter exposure times. In Fig. 2, we show datasets of an office scene captured with exposure times of 1 ms to 10 ms per laser position, which results in a total capture time as low as 24 s. Further reconstructions of a shelf dataset are shown later as additional results, showing that we can reduce exposure times at least down to 50 ms per data point without a significant loss in quality, even with ambient light. This results in less than 20 min of total capture time. In our current prototype, we capture data sequentially with a single SPAD. Prototype SPAD arrays are currently under development, and it seems likely that a 16 \times 16 array will be available by the end of the year. We thus expect to be able to capture 256 data points in less than 0.1 s in the near future.

Collected data. In total (counting captures with different lighting and exposure times as different sets), we use 12 experimental and two simulated datasets. All experimental datasets use a single SPAD location and 180 by 130 laser positions. The datasets and exposure times are:

- An office scene collected with 1 s exposure per laser position. This dataset is used to create the video shown in Supplementary Video 1, frames of which are shown in Fig. 4b. A photograph and reconstruction of this scene is also shown in the Supplementary Video. The data are analysed in Extended Data Fig. 3 and Extended Data Table 1.
- An office scene collected with exposures of 10 ms, 5 ms and 1 ms, used in Fig. 2, Extended Data Figs. 6–8 and Extended Data Table 1.
- A scene of a bookshelf used in Fig. 3a and in Extended Data Table 1.
- A scene of a bookshelf captured with various exposure times and ambient light conditions, shown in Extended Data Fig. 2 and Extended Data Fig. 5.
- A scene with letters distributed over a large depth, used in Fig. 3b and Extended Data Table 1.
- A scene of the letters NLOS in a plane, used in Fig. 4a and Extended Data Table 1.

To provide further insight into the noise and artefacts present in our data, we go through an analysis of the raw data from our 1-s-exposure office scene. We compare the maximum and average number of photons per second and laser position \mathbf{x}_p for our captured scenes in Extended Data Table 1. The dark count rate of our detector is 10 photons per second. We do not explicitly subtract dark counts nor ambient light or backgrounds. The high total photon numbers in the transient responses (Extended Data Table 1) are due to the long responses associated with the large depth and volume of the scenes, and not due to a particularly bright signal. Example data for a scene of a shelf are shown in Extended Data Fig. 2 (whose reconstruction can be found on Extended Data Fig. 5). In this scene, our longest (1 s) exposure time peaks at about 150 photons per second (such peaks are probably due to the presence of specular surfaces), and the captured signal is extremely noisy. In comparison, the recent method by O'Toole et al.⁹ acquires a brighter, cleaner signal in 0.1 s, peaking at about 600 photons per second, owing to the use of retroreflective paint applied on the hidden objects (data from their data_resolution_chart_40cm dataset).

Let us further analyse the captured data. In Extended Data Fig. 3a, we show a visualization of our data matrix for the 1-s-exposure office scene using the Matlab function *imagesc*, in which each row is the data collected for a different location of the laser illumination spot, and each column contains a different time

bin. The first time bin corresponds to the time when the illumination laser pulse leaves the relay wall. In the images, we do not show time bins 10,001 to 15,000 as they are mostly empty, owing to the closing of the gate. As can be seen, there are some sparse, very large peaks in the dataset that saturate the counting registers of our time-correlated single-photon counter ($2^{16} - 1$ counts). As we will see, these artefacts in the data are likely to be due to imperfections in the gating or optical set-up.

Let us focus on the first instants of the captured data shown in Extended Data Fig. 3a, which reveal features that look like straight diagonal lines in the first few time bins. The fact that there are straight lines in this plot indicates that they are likely related to a first-bounce signal, rather than the scene response. NLOS signals should show up as hyperbolas or sections of hyperbolas in this type of visualization, and the curvature of the hyperbolas should be highest at the earliest time bins. The image contains many more features that look like straight lines that do not appear to have the correct hyperbolic curvature to be NLOS signals. Many of them also appear identically in the other datasets, which is another hint that they are probably not real NLOS data but artefacts related to the measurement system. Our algorithm is completely agnostic to the presence of these artefacts. The brightest peaks also appear too early in the data to be associated with a NLOS object. To see this, consider that the closest object in any of our scenes is the chair in the office scene, and it is more than 50 cm away from the wall. Consequently, the first time response from an actual object cannot arrive at the SPAD earlier than 3.3 ns after the laser illuminates the relay wall. Time bins are 4 ps wide. Any data before time bin 833 therefore can only be an artefact. We will speculate more about the origin of these artefacts later.

If we ignore those first 833 time bins that contain no useful data, we obtain a dataset that can yield some meaningful statistics about the data. In this dataset, the largest photon count in all our over 200 million time bins is smaller than 1,400 photons. As we show below, this 1,400 maximum is probably still due to a gate artefact that happened to occur slightly later than 3 ns into the dataset. Statistics for all datasets are shown in Extended Data Table 1.

Maximum photon counts usually come from the objects in the scene closest to the wall. Considering the large depth and specularities of our scenes, most of the reconstructed scene volume is using signals much weaker than the maximum signal, as voxels are further from the wall. Signals from a given surface are expected to drop in magnitude with distance L as $1/L^4$. An object generating 100 photon peaks at 50 cm distance in the front of our scene would therefore only create 100/8 photons if placed at 1 m and $100/625 = 0.2$ photons at 2.5 m towards the end of our office scene. This ability to handle scenes with large dynamic range in the data is another advantage of our algorithm.

In Extended Data Fig. 3b, we show a plot of the photon counts over time bins for the laser position that received the most total photons. We again see the extreme peak of $2^{16} - 1$ counts in the beginning of the dataset. Again, this peak cannot be a real third-bounce signal as it would require the pulse to travel between the laser position and SPAD position much faster than the speed of light. The actual NLOS data start around time bin 1,000 and peak at just above 50 photons.

Finally, we show a plot of the laser position that received the total photon count closest to the median of all laser positions (Extended Data Fig. 3c). We can see that the count generally stays below 150 photons, with what are probably specular peaks reaching 200 photons and a large (450 photon) peak at the beginning of the dataset that is either a specular peak or another gate artefact. Note that as we illuminate only a grid of points at the wall, we do not capture all the specular peaks in our data. To see a specular reflection peak from a scene surface, we have to be lucky enough to illuminate the exact spot on the wall that results in the specular reflection that overlaps with the SPAD position (see Supplementary Fig. 2 for an illustration). Therefore, specular peaks in our measurements can vary greatly, depending on how close to the peak the laser sampled the wall. Again, we point out that this type of uncontrolled artefact does not affect our algorithm.

As we stated above, the time bin with the highest photon count when ignoring obvious early artefacts contains about 1,400 photons. Next we plot the laser position that contains this time bin (Extended Data Fig. 3d). Note that zero on the x axis here corresponds to time bin 834. As we see, the 1,400 photon peak appears very close to the beginning of the transient and may be a gating artefact that occurs in the data just after the opening of the gate. This type of data distortion is described further below. If not a gating artefact, the peak is probably a specular reflection, as it is very narrow and could only be caused by a small isolated diffuse patch or a specular surface in the scene. Peaks from extended diffuse surfaces are necessarily longer in duration.

We conclude that although our data contain artefacts, the photon counts useful for reconstructions are no higher or cleaner than in previous methods. Note that the removal of early artefacts is only done here to generate Extended Data Fig. 3b–d, to allow visualization. All reconstructions shown in the manuscript contain the full recorded data without the removal of any potential artefacts or time bins.

Even though an understanding of the origin of the artefacts is not needed for our method, we can speculate on the sources of some of them.

(1) Many of the early peaks in our data are likely to be related to imperfections in our gating method. When the SPAD gate opens just after the laser pulse has passed, photoelectrons in the SPAD may cause a detection event that is not due to a photon but to electrons excited by the first-bounce light and trapped in long-lived states in the SPAD. Even though these electrons are not amplified, they need to be transported off the SPAD junction or they can cause counts as soon as the gate opens.

(2) The gate may not block the pulse for some laser positions. The gate has to be positioned such that it blocks the laser in all laser positions while not blocking any signal. This is not always possible, and we do not re-adjust the gate for each position while scanning.

(3) Effects inside the imaging system can keep light trapped long enough to cause a peak at the time when the NLOS data arrive. This can be due to multiple reflections between lenses, multiphoton fluorescence in the glass or coating of the lenses, or stray light reflecting off a random surface at the right distance. We have confirmed some of these effects but suspect there are many more.

(4) In particular, we can see light that travels from the laser spot to the SPAD, reflects off the surface of the SPAD pixel, is imaged back to the relay wall and comes back to the SPAD. In confocal or near-confocal configurations, this can create a peak that is many times brighter than the data.

Retroreflective targets can be used to reduce many of these artefacts, most of which are created either by the laser or a first-bounce reflection of the laser. If the hidden target is retroreflecting, the ratios between the brightness of the laser and its first bounce and the brightness of the third-bounce NLOS data are reduced by multiple orders of magnitude.

Helmholtz reciprocity. Ideally, we would capture $H(\mathbf{x}_p \rightarrow \mathbf{x}_c, t)$ sampling points on both the projector aperture $\mathbf{x}_p \in P$ and the camera aperture $\mathbf{x}_c \in C$. In our current set-up with a single SPAD, we only sample a single point for \mathbf{x}_c . From Helmholtz reciprocity, we can interpret these datasets as having a single \mathbf{x}_p and array of \mathbf{x}_c . The choice of capture arrangement is made for convenience, as it is easier to calibrate the position of the laser spot on the wall. Improved results are anticipated once array sensors become available (currently under development).

Additional validation and discussion. *Resolution limits.* The resolution limit for NLOS imaging systems with an aperture diameter d at imaging distance L is closely related to the Rayleigh diffraction limit⁷: $\Delta_x = 1.22c\sigma L/d$, with c the speed of light in vacuum, for a pulse of full width at half maximum σ . O'Toole et al.⁹ derive a criterion for a resolvable object based on the separability of the signal in the raw data, not in the reconstruction, resulting in a similar formula, $\Delta_x = 0.5c\sigma L/d \approx 0.5\lambda L/d$.

In our virtual LOS imaging system, we can formulate a resolution limit that ensures a minimum contrast in the reconstruction, based on the well-known resolution limits of wave-based imaging systems. The resolution limit therefore depends on the particular choice of virtual imaging system. For an imaging system that uses focusing only on the detection or illumination side, this limit is approximated by the Rayleigh criterion. For an imaging system that provides focusing on both the light source and the detector side, the resolution doubles (as it does, for example, in a confocal or structured illumination microscope) and the resolution limit becomes $\Delta_x = 0.61\lambda L/d$.

Effect of strong interreflections. To confirm the presence and effect of strong interreflections in our captured data, we compare the data qualitatively with primary data from a synthetic bookshelf scene, with and without interreflections. The bookshelf is placed in a corridor of $2\text{ m} \times 2\text{ m} \times 3\text{ m}$, with only a single lateral aperture of $1\text{ m} \times 2\text{ m}$ to allow the hidden scene to be imaged. The shelf has a size of $1.4\text{ m} \times 0.5\text{ m}$, placed at 1.7 m from the relay wall and 0.3 m from the lateral walls. The virtual aperture has a size of $1.792\text{ m} \times 1.792\text{ m}$ and a granularity of 256×256 laser points; we use $\lambda = 4\Delta_p$ and $\Delta_p = 0.7\text{ cm}$.

As can be seen in Extended Data Fig. 4, the synthetic data clearly show how the presence of interreflections adds, as expected, low-frequency information resembling echoes of light. The same behaviour can be seen in the real captured data, revealing the presence of strong interreflections.

Additionally, we evaluate the robustness of our method in the presence of such interreflections. Similar to recent work⁹, we compare between a voxelization of the ground-truth geometry and a reconstructed voxel-grid obtained from our irradiance reconstructions, with and without including interreflections; the resulting MSE is as follows: without interreflections (Extended Data Fig. 4a), MSE 4.93 mm; with interreflections (Extended Data Fig. 4b), MSE = 4.66 mm.

Effect of exposure time. *Ambient light.* To analyse how well our technique works in ambient light and with much shorter exposure times, we perform several additional measurements using progressively shorter exposure times, showing that we can reduce exposure times at least down to 50 ms per data point without a significant loss in quality (see Extended Data Fig. 5). Extended Data Fig. 2 shows raw data for one of the laser positions. In particular, it shows the number of photons per second accumulated in each time bin (that is, the collected histogram divided

by the integration time in seconds). As expected, all three curves appear to follow the same mean but have a larger variance for lower exposure times. The raw data thus become noisier as exposure time decreases. The effects on our reconstruction, however, are minor, as Extended Data Fig. 5 shows.

Short-exposure captured data. Extended Data Fig. 6 shows the reconstruction of the office scene (Fig. 2) for short exposure times of 10 ms, 5 ms and 1 ms for each of the roughly 24,000 laser positions. This leads to total capture times of about 4 min, 2 min and 24 s respectively. Plots showing raw data from those datasets are given in Extended Data Fig. 7.

We compare the results of our reconstructions on the 1 ms data against filtered backprojection with a Laplacian filter³, as well as the Laplacian-of-Gaussian (LOG)-filtered backprojection¹⁹, which generally achieves better results. We are not aware of any reconstruction method that consistently outperforms a LOG-filtered backprojection. Extended Data Fig. 8 shows the result of this comparison. **Non-Lambertian surfaces.** To validate the robustness of our method in the presence of non-Lambertian materials in the hidden scene, we have created a synthetic scene made up of two letters, R and D, one partially occluding the other, placed in a corridor of $2\text{ m} \times 2\text{ m} \times 3\text{ m}$, with only a single lateral aperture of $1\text{ m} \times 2\text{ m}$ to allow imaging the hidden scene. The letters have a size of $0.75\text{ m} \times 0.8\text{ m}$, placed at 1.25 m and 1.7 m from the relay wall, respectively, and 0.5 m from the lateral walls (see Extended Data Fig. 9a). The virtual aperture has a size of $1.792\text{ m} \times 1.792\text{ m}$ and a granularity of 128×128 laser points; we use $\lambda = 4\Delta_p$ with $\Delta_p = 1.4\text{ cm}$. We start with purely Lambertian targets and progressively increase their specularly. We use the Ward BRDF model²⁸, decreasing the surface roughness, using available transient rendering software²⁶. The simulation includes up to the fifth indirect bounce.

Extended Data Fig. 9b shows the resulting irradiance reconstructions. Because our method does not make any assumption about the surface properties of the hidden scene, the changes in material appearance do not significantly affect our irradiance reconstructions. Similar to recent work⁹, we compare a voxelization of the ground-truth geometry and the reconstructed voxel-grid; the resulting MSE for each of the different reflectances is as follows: for a surface roughness of 1 (perfect Lambertian), MSE = 2.1 mm; for a surface roughness of 0.4, 2.2 mm; for a surface roughness of 0.2, MSE = 2.2 mm.

Reconstruction comparison with other methods. Our imaging system allows hidden geometry to be reconstructed. For this application, we show a comparison using the publicly available confocal dataset⁹. This set can be reconstructed using different NLOS methods; we show results for confocal NLOS deconvolution⁹, filtered backprojection⁷ and our proposed method. For these confocal measurements, backprojection can be expressed as a convolution with a pre-calculated kernel, and thus all three methods are using the same backprojection operator. Neither our method nor filtered backprojection is limited to confocal data, and both can be acquired by making use of simpler devices and capture configurations. They can thus be applied to a broader set of configurations and considerably more complex scenes. For the confocal NLOS deconvolution method⁹, we leave the optimal parameters unchanged. For our proposed virtual wave method, we use the aperture size and its spatial sampling grid (see Supplementary Information) to calculate the optimal phasor-field wavelength. For the filtered backprojection, it is important to choose a good discrete approximation of the Laplacian operator in the presence of noise. Previous works implicitly do the denoising step by adjusting the reconstruction grid size to approximately match the expected reconstruction quality^{2,3,7}, or by downsampling across the measurements⁹. If used correctly, all of these methods result in a high-quality reconstruction from a Laplacian filter. To provide a fair comparison without changing the reconstruction grid size, we convolve a Gaussian denoising kernel with the Laplacian kernel, resulting in a LOG filter, which we apply over the backprojected volume.

Note that a large improvement in reconstruction quality for the simple scenes included in the dataset (isolated objects with no interreflections) is not to be expected, since existing methods already deliver reconstructions approaching their resolution limits. We nevertheless achieve improved contrast and cleaner contours in our wave camera method, due to our better handling of multiply scattered light, which pollutes the reconstructions in the other methods (see Extended Data Fig. 10).

In the noisy datasets (Extended Data Fig. 11), filtered backprojection fails. confocal NLOS includes a Wiener filter that performs well at removing uniform background noise, although a noise level must be explicitly estimated. Our phasor-field virtual wave method, on the other hand, performs well automatically, without the need to explicitly estimate a noise level. This is important in complex scenes with interreflections, where the background is not uniform across the scene, and the noise level cannot be reliably estimated.

Nevertheless, our main contribution is not that of improving the reconstruction for simple, third-bounce scenes. Instead, our method allows a new class of NLOS algorithms to be derived, which can successfully handle scenes of much greater complexity.

Data availability

The measured data and the phasor-field NLOS code supporting the findings of this study are available in the figshare repository <https://doi.org/10.6084/m9.figshare.8084987>. Additional data and code are available from the corresponding authors upon request.

Code availability

Our data and reconstruction code can be found in the figshare repository <https://doi.org/10.6084/m9.figshare.8084987>.

28. Ward, G. J. Measuring and modeling anisotropic reflection. *Comput. Graph.* **26**, 265–272 (1992).

Acknowledgements This work was funded by DARPA through the DARPA REVEAL project (HR0011-16-C-0025), the NASA Innovative Advanced Concepts (NIAC) Program (NNX15AQ29G), the Air Force Office of Scientific Research (AFOSR) Young Investigator Program (FA9550-15-1-0208), the Office of Naval Research (ONR, N00014-15-1-2652), the European Research Council (ERC) under the EU's Horizon 2020 research and innovation programme (project CHAMELEON, grant no. 682080), the Spanish Ministerio de Economía y Competitividad (project TIN2016-78753-P) and the BBVA Foundation

(Leonardo Grant for Researchers and Cultural Creators). We thank J. Teichman for insights and discussions in developing the phasor-field model. We also acknowledge M. Buttafava, A. Tosi and A. Ingle for help with the gated SPAD detector, and B. Masia, S. Malpica and M. Galindo for careful reading of the manuscript.

Author contributions X.L., S.A.R., M.L.M. and A.V. conceived the method. X.L., I.G., M.L.M. and J.H.N. implemented the reconstruction. M.L.M., X.L., J.H.N. and T.H.L. built and calibrated the system. I.G., D.G. and A.J. developed the simulation system. A.J., D.G. and A.V. coordinated the project. All authors contributed to writing the paper.

Competing interests The authors declare no competing interests.


Additional information

Supplementary information is available for this paper at <https://doi.org/10.1038/s41586-019-1461-3>.


Correspondence and requests for materials should be addressed to A.V.

Peer review information *Nature* thanks Jeffrey H. Shapiro, Ashok Veeraraghavan and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

Reprints and permissions information is available at <http://www.nature.com/reprints>.

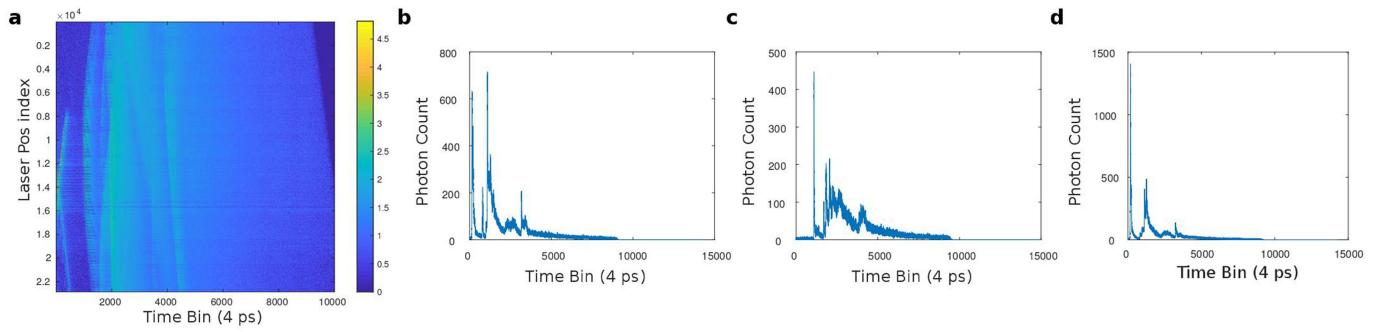


Extended Data Fig. 1 | Capture hardware used for the results shown in this Letter.



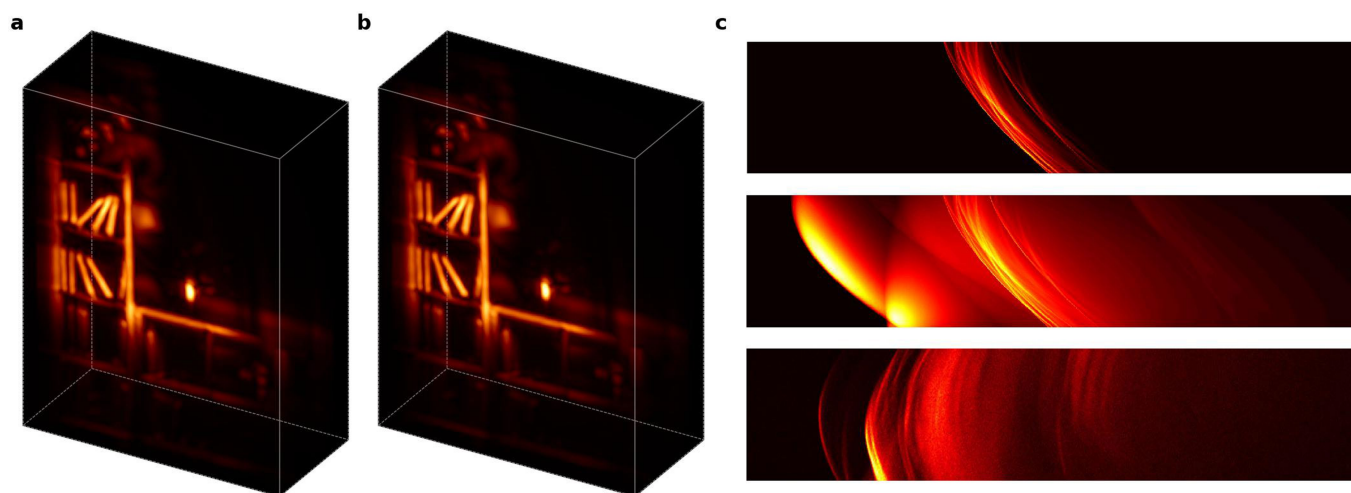
Extended Data Fig. 2 | Data comparison. **a**, Raw data for one of the laser positions x_p . Shown is the number of photons per second accumulated in each time bin (that is, the collected histogram divided by the integration time in seconds). Time bins are 4 ps wide. As expected, all three curves appear to follow the same mean, but there is a larger variance for lower

exposure times. The raw data thus become noisier as exposure time decreases. The effects on the reconstruction are minor, as Extended Data Fig. 4 shows. T_{acq} , acquisition time. **b**, Example dataset from ref. ⁹ for comparison.




Extended Data Fig. 3 | Visualization of the raw data for our long-exposure office scene. **a**, Base-10 logarithm of the photon counts in all time bins. Pos index, laser position index; the 24,000 laser positions on the wall are labelled with these consecutive numbers. **b–d**, After removal of the first 833 time bins in each dataset, the plots show: the photon counts

for the laser position that received the largest total number of photons in the dataset (**b**); the counts for the laser position that received the median number of photon counts (**c**); and the counts for the laser position that contains the time bin with the global maximum count in the entire set (**d**).




Extended Data Fig. 4 | Robustness to multiple reflections. Result for the synthetic bookshelf scene. **a**, Without interreflections. **b**, Including high-order interreflections. The quality of the results is very similar. **c**, Primary data (streak images) from the same scene without (top), and with interreflections (middle). The synthetic data clearly show how the presence of interreflections adds, as expected, low-frequency information

resembling echoes of light. The bottom image shows primary data captured from the real office scene in Fig. 2. It follows the same behaviour as the middle image, revealing the presence of strong interreflections. Colours refer to numerical values from Matlab's 'fire colormap', in arbitrary units.



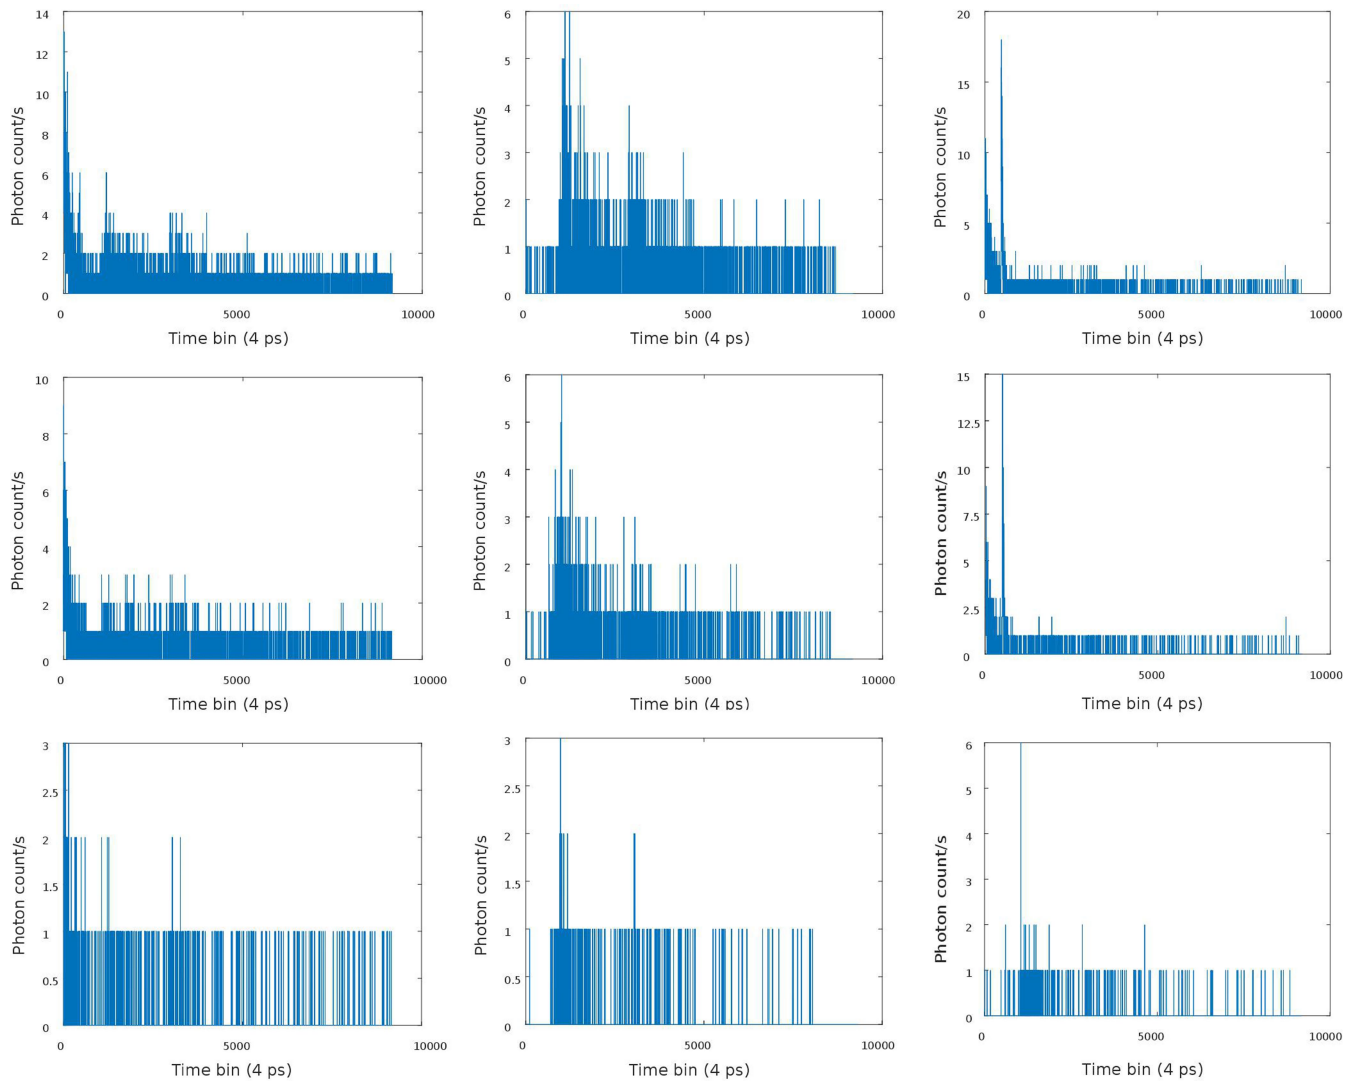
Extended Data Fig. 5 | Robustness to ambient light and noise. **a**, Hidden bookshelf. **b**, Imaging results with increasingly higher exposure times; even at 50 ms, there is no significant loss in quality. Top row, image using only the pulsed laser as illumination source. Bottom row, on adding a large

amount of ambient light (same conditions as the photograph in **a**), the quality remains constant. **c**, Difference between the 50-ms- and 1,000-ms-exposure captures for the lights-off case.




Extended Data Fig. 6 | Short-exposure reconstructions. Reconstruction of the office scene using very short capture times. **a**, Photograph of the captured scene. **b**, From left to right, reconstructions for data captured

with 10 ms, 5 ms and 1 ms exposure time per laser. The total capture time was about 4 min, 2 min and 24 s, respectively.




Extended Data Fig. 7 | Short-exposure data. Photon counts in the raw data for our office scene for 10 ms (top row), 5 ms (centre row) and 1 ms (bottom row) exposure times per laser position. After removing the first 833 time bins in each dataset, the columns show: the photon counts for

the laser position that received the largest total number of photons in the dataset (left); the counts for the laser position that received the median number of photon counts (centre); and the laser position that contains the time bin with the global maximum count in the entire set (right).




Extended Data Fig. 8 | Comparison to prior methods. Reconstruction of the office scene using very short capture times of 1 ms per laser (24 s in total). **a**, Filtered backprojection using the Laplacian filter. **b**, LOG-filtered backprojection. **c**, Our method.




Extended Data Fig. 9 | Robustness to scene reflectance. **a**, Geometry of our experimental set-up. **b**, From left to right, imaging results for the Lambertian targets (roughness 1) and increasingly specular surfaces

(roughness 0.4 and roughness 0.2). The reconstructed irradiance is essentially the same for all cases.



Extended Data Fig. 10 | Reconstruction comparison on a public dataset. From left to right: confocal NLOS deconvolution, filtered (LOG) backprojection (FBP) and our proposed method. A large improvement in reconstruction quality for the simple scenes included in the dataset

(isolated objects with no interreflections) is not to be expected, as existing methods already deliver reconstructions approaching their resolution limits. Nevertheless, our method achieves improved contrast and cleaner contours, owing to better handling of multiply scattered light.



Extended Data Fig. 11 | Reconstruction comparison (noisy data).

From left to right: confocal NLOS deconvolution, FBP and our proposed method. Top row represents a non-retroreflective object; bottom row represents a retroreflective object captured in sunlight. In the presence of noisy data, FBP fails. Confocal NLOS includes a Wiener filter that needs to

be explicitly estimated. Our phasor-field virtual wave method yields better results automatically. This is particularly important in complex scenes with interreflections, where the background is not uniform across the scene, and the noise level cannot be reliably estimated.

Extended Data Table 1 | Photon statistics for captured data used in the paper and in Supplementary Information

	Total Photons	Photons/bin	Max. bin	Avg./ laser
Large depth scene	3215722952	9.7	552	13742
NLOS letters	6502986696	19.6	2889	27791
Shelf	6158590767	18.6	2074	26319
Office Scene	6201680972	18.7	1406	26503
Office Scene 10 ms	48017499	0.14	18	2716
Office Scene 5 ms	24012257	0.072	15	1026
Office Scene 1 ms	4801568	0.014	6	205

The first four scenes were captured with 1 s exposure time. The first column shows the total photons counted, the second shows the average photon count per time bin, the third is the maximum count over all time bins, and the last contains the average number of photons collected in each laser position in the dataset.