Time-of-Flight Imaging

lidar, single-photon imaging, non-line-of-sight imaging

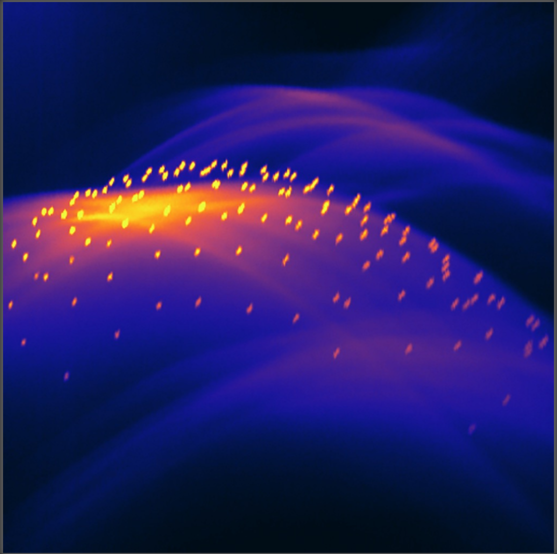

CSC2529

David Lindell

University of Toronto

cs.toronto.edu/~lindell/teaching/2529

*slides adapted from Matt O'Toole, Gordon Wetzstein, Yannis Gkioulekas

transient imaging (a.k.a. femtophotography)

Raskar et al. [2011]

overview

- Time-resolved imaging
- Single-photon avalanche diodes (SPADs)
- Single-photon lidar
- Non-line-of-sight imaging
- Imaging through scattering media

echolocation

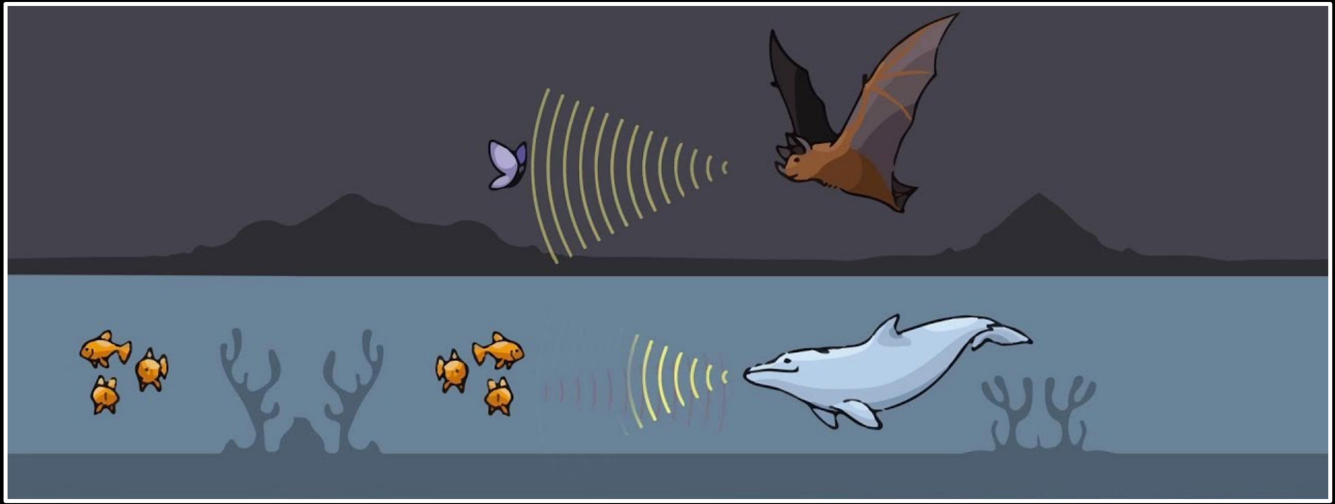

speed of sound in air: 343 meters / sec

in water: 1480 meters / sec


Light takes 1.255 seconds to travel from the earth to the moon
speed of light in a vacuum: **299,792,458** meters / sec
(Light travels approximation **1 MILLION** times faster than sound!)

transient imaging

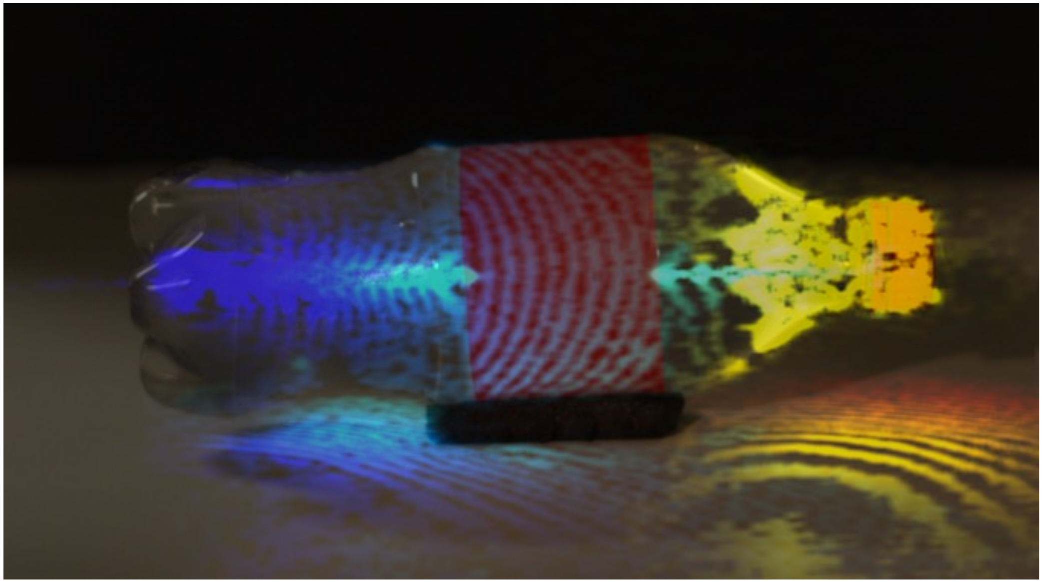

speed of light in a vacuum: **299,792,458** meters / sec

(Light travels approximation **1 MILLION** times faster than sound!)

direct and **indirect** time-of-flight sensors for transient imaging

Velodyne VLS-128

Direct time-of-flight sensor


Microsoft Kinect v2

Indirect time-of-flight sensor

direct and indirect time-of-flight sensing

Velodyne VLS-128

Microsoft Kinect v2

Direct time-of-flight sensor

Indirect time-of-flight sensor

direct and indirect time-of-flight sensing

Velodyne VLS-128

Direct time-of-flight sensor

direct time-of-flight principle

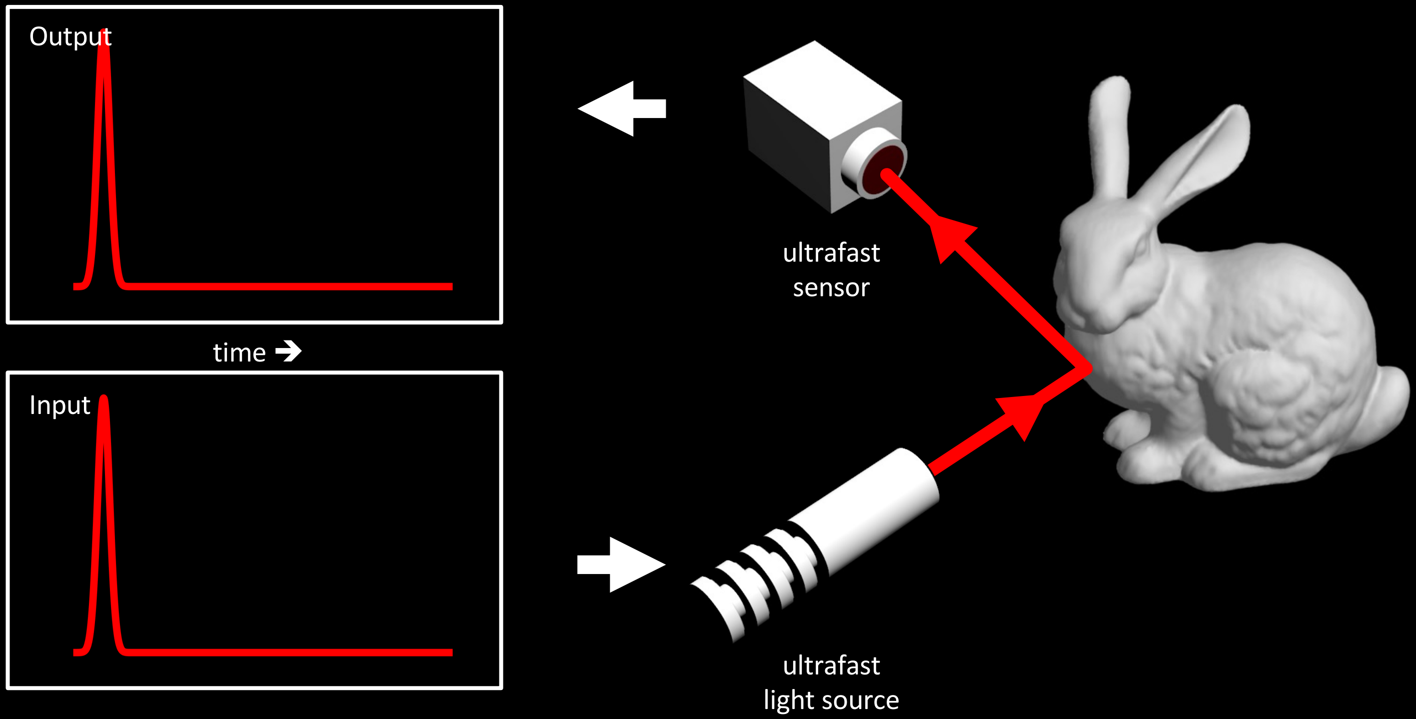

direct time-of-flight principle

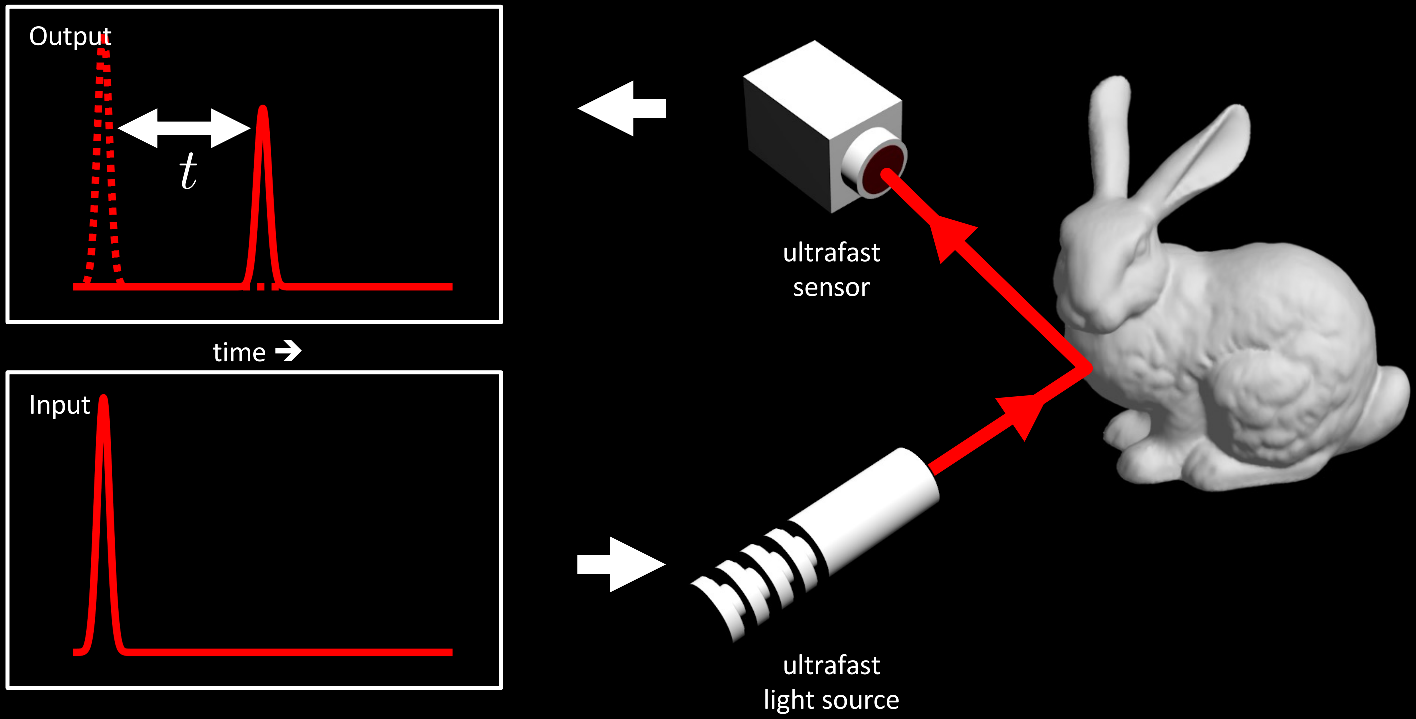

direct time-of-flight principle

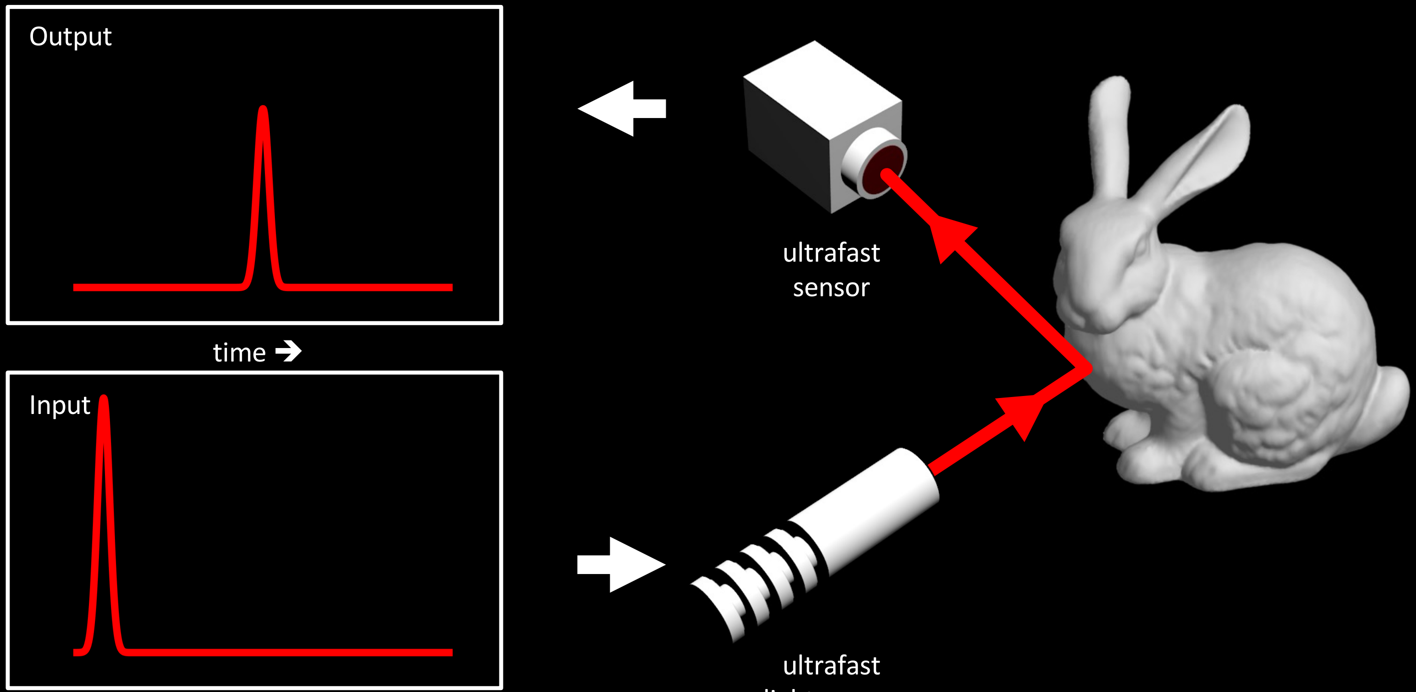

direct time-of-flight principle

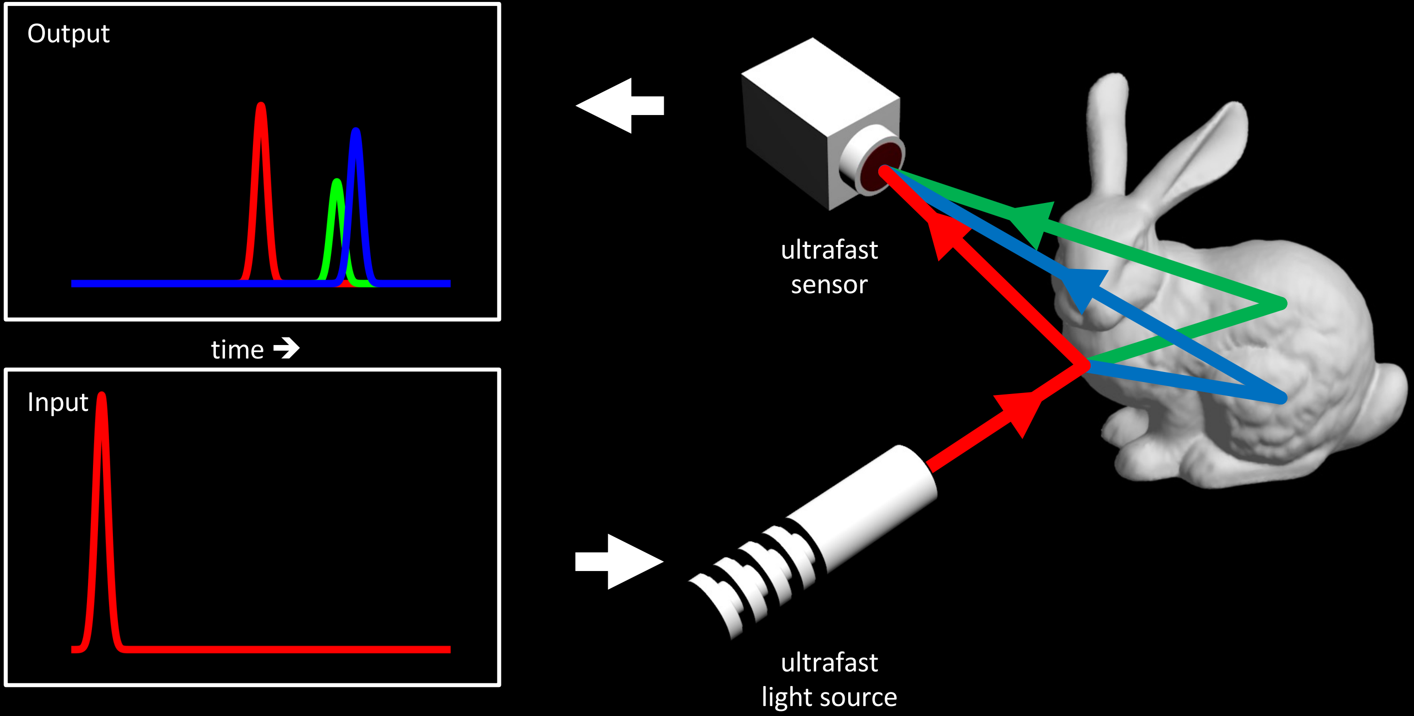

direct time-of-flight principle

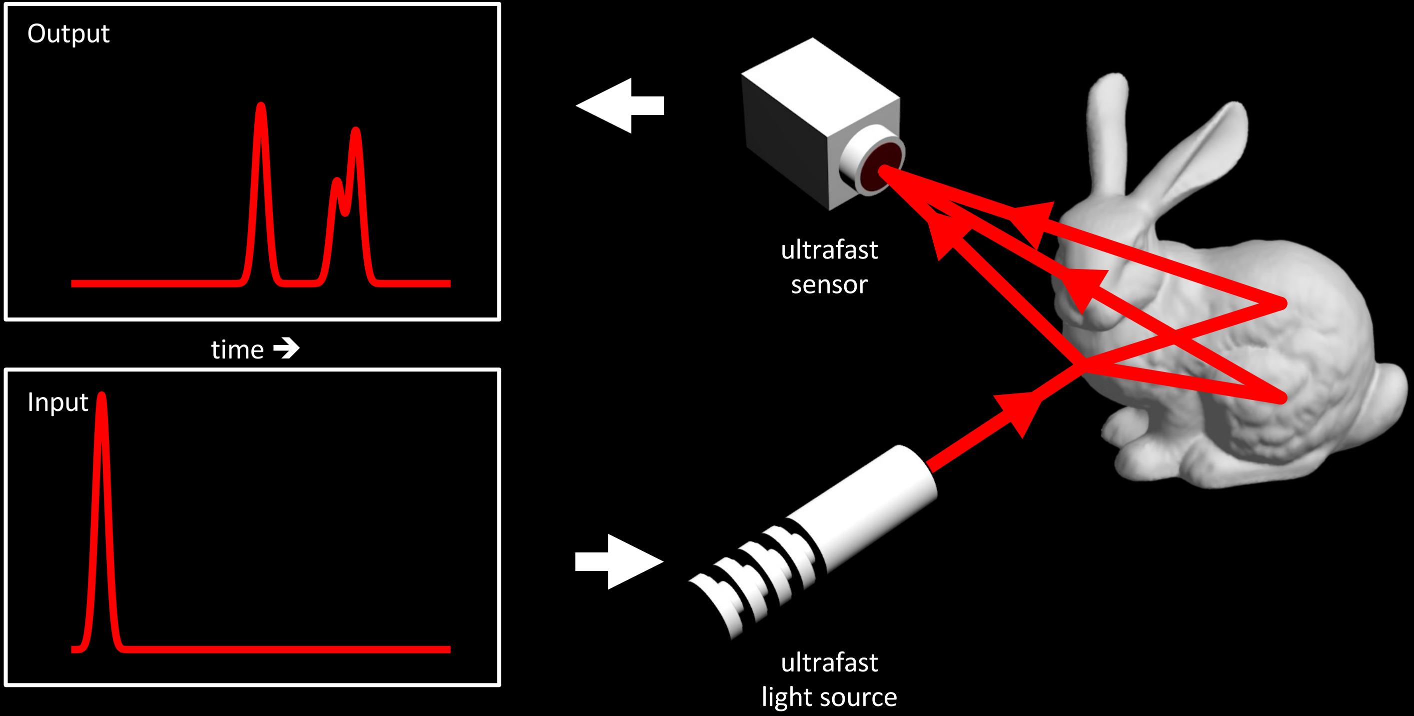

direct time-of-flight principle

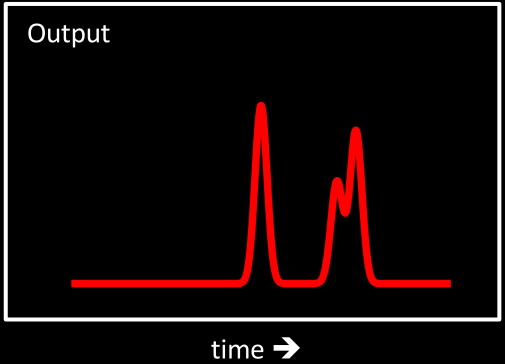

transient measurement

direct and indirect time-of-flight sensing

Velodyne VLS-128

Microsoft Kinect v2

Direct time-of-flight sensor

Indirect time-of-flight sensor

direct and indirect time-of-flight sensing

Microsoft Kinect v2

Indirect time-of-flight sensor

indirect time-of-flight principle

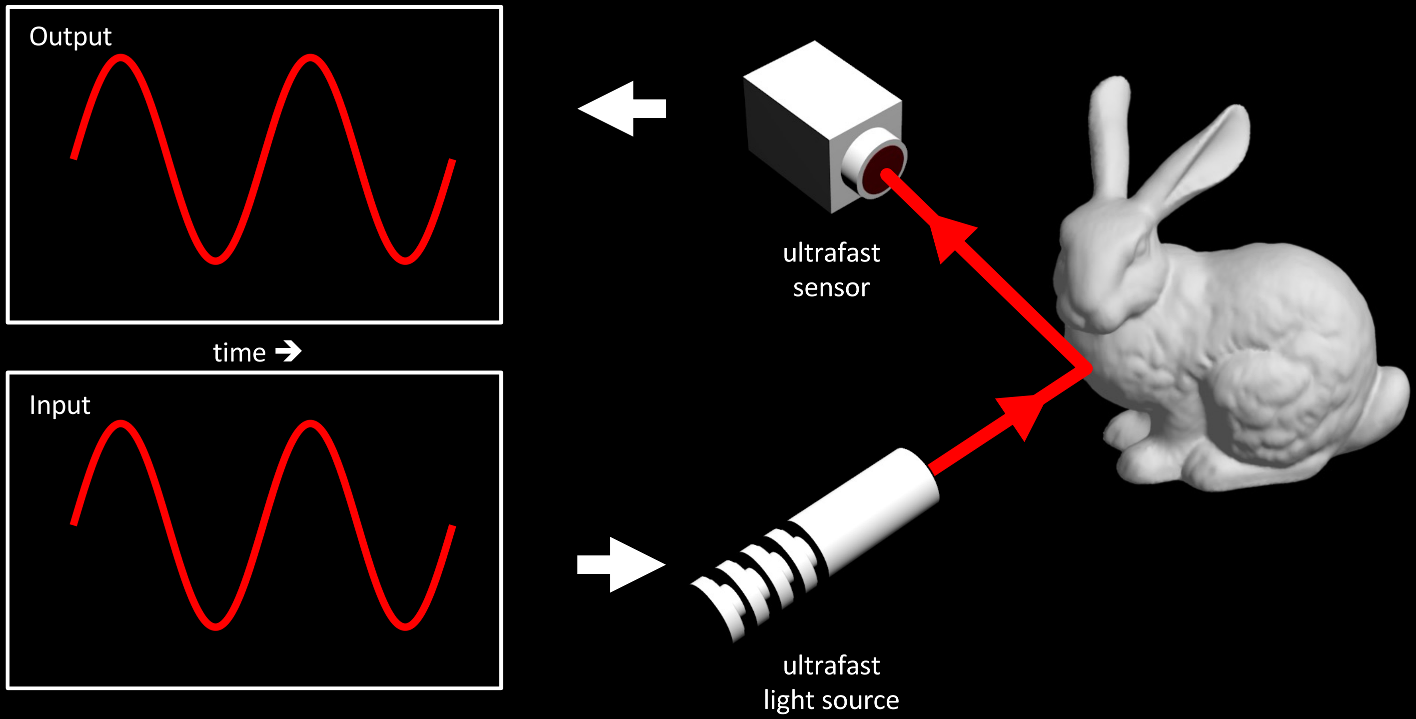

indirect time-of-flight principle

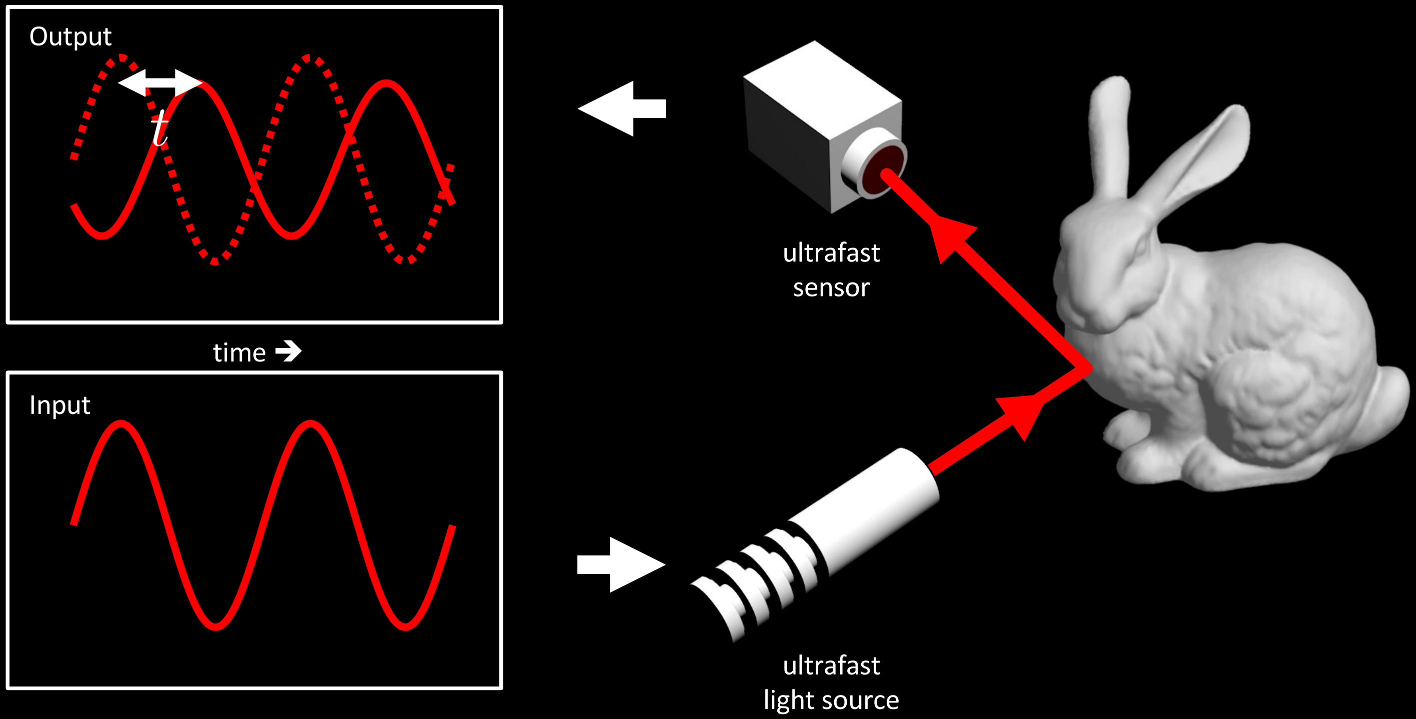

indirect time-of-flight principle

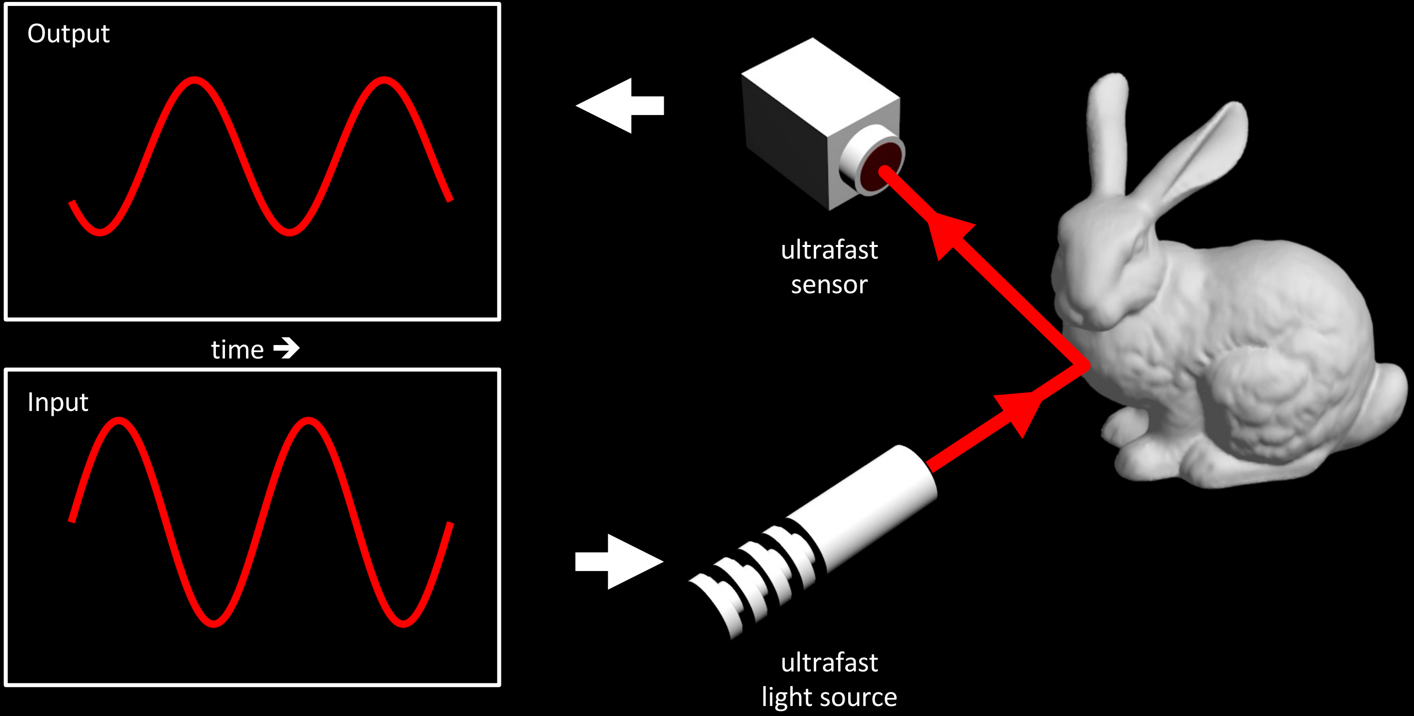

indirect time-of-flight principle

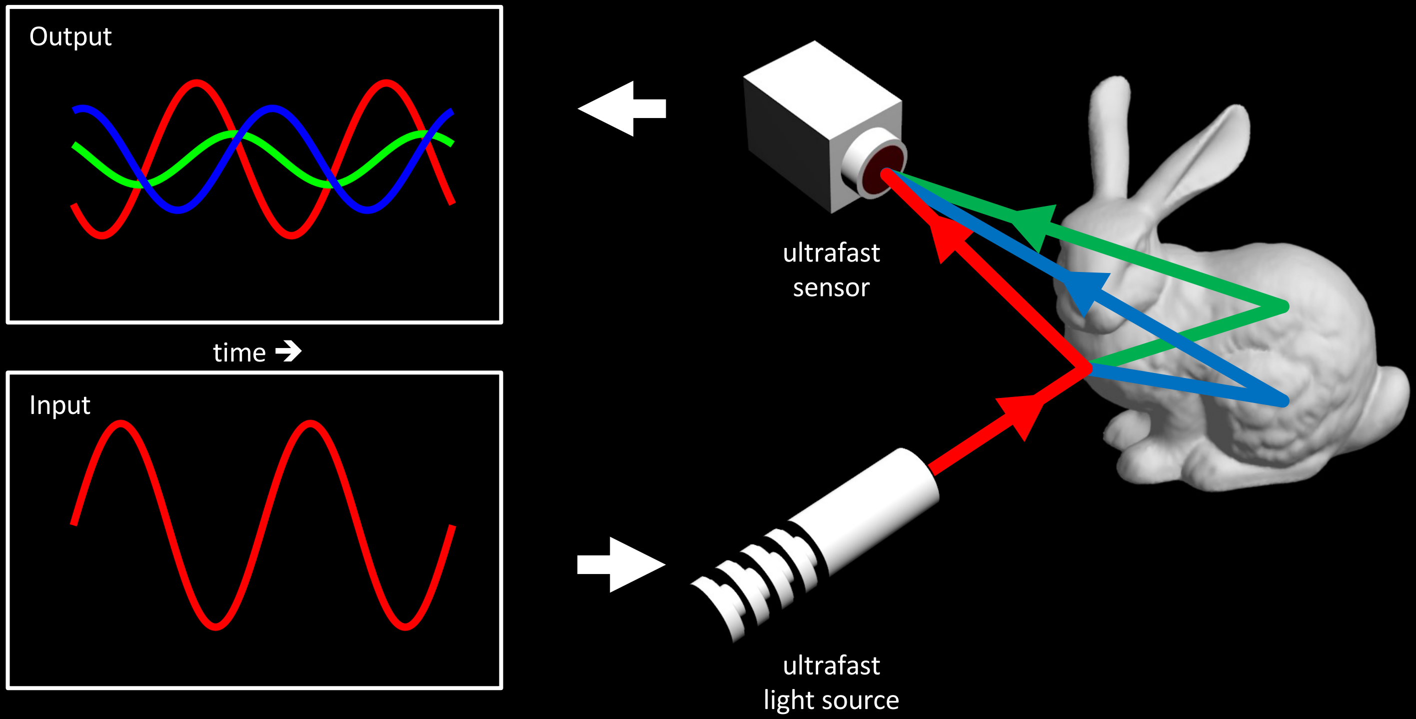

indirect time-of-flight principle

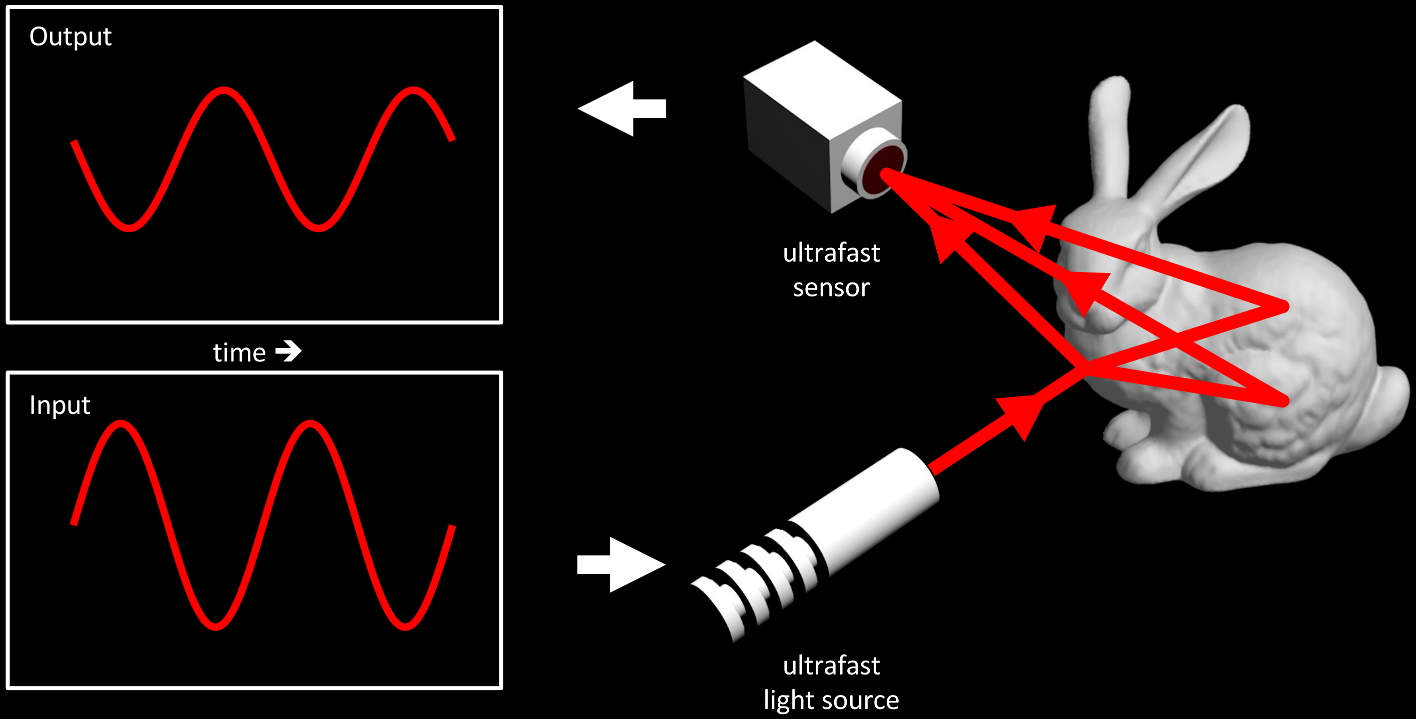

indirect time-of-flight principle

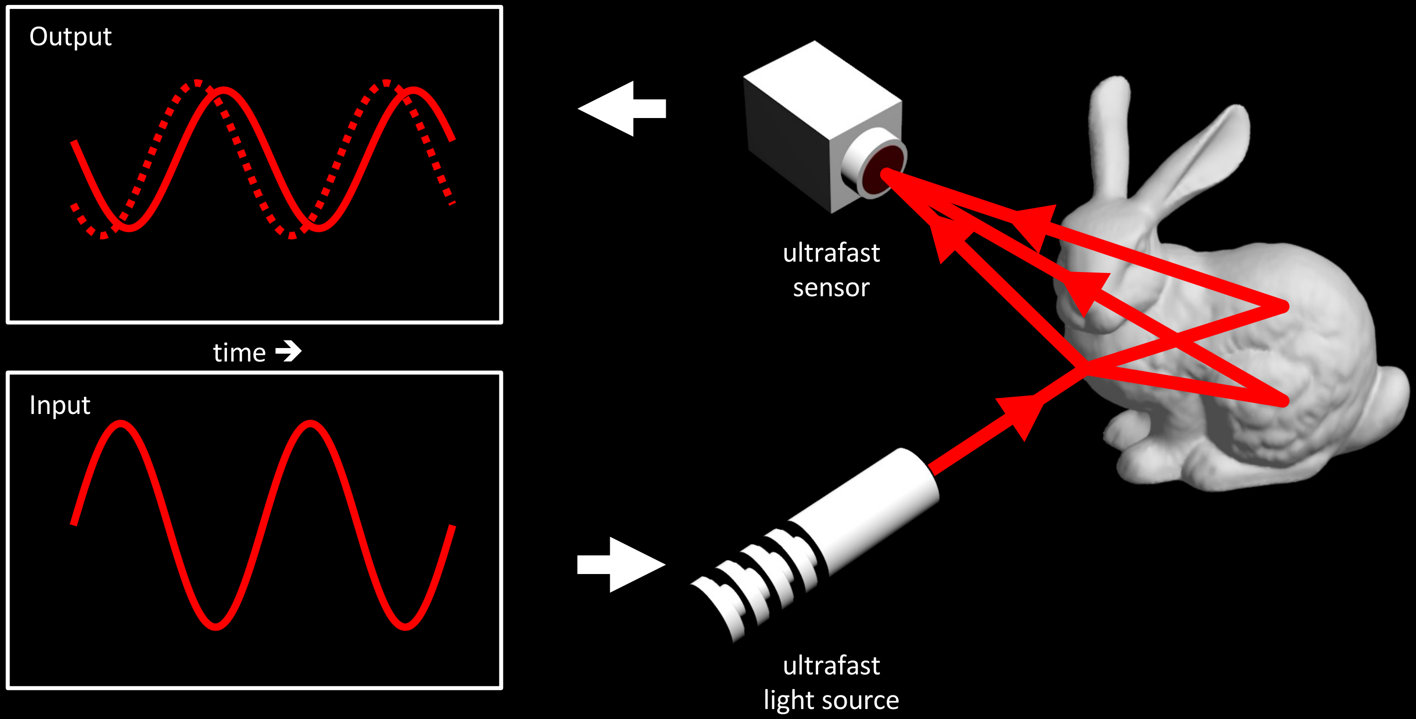

indirect time-of-flight principle

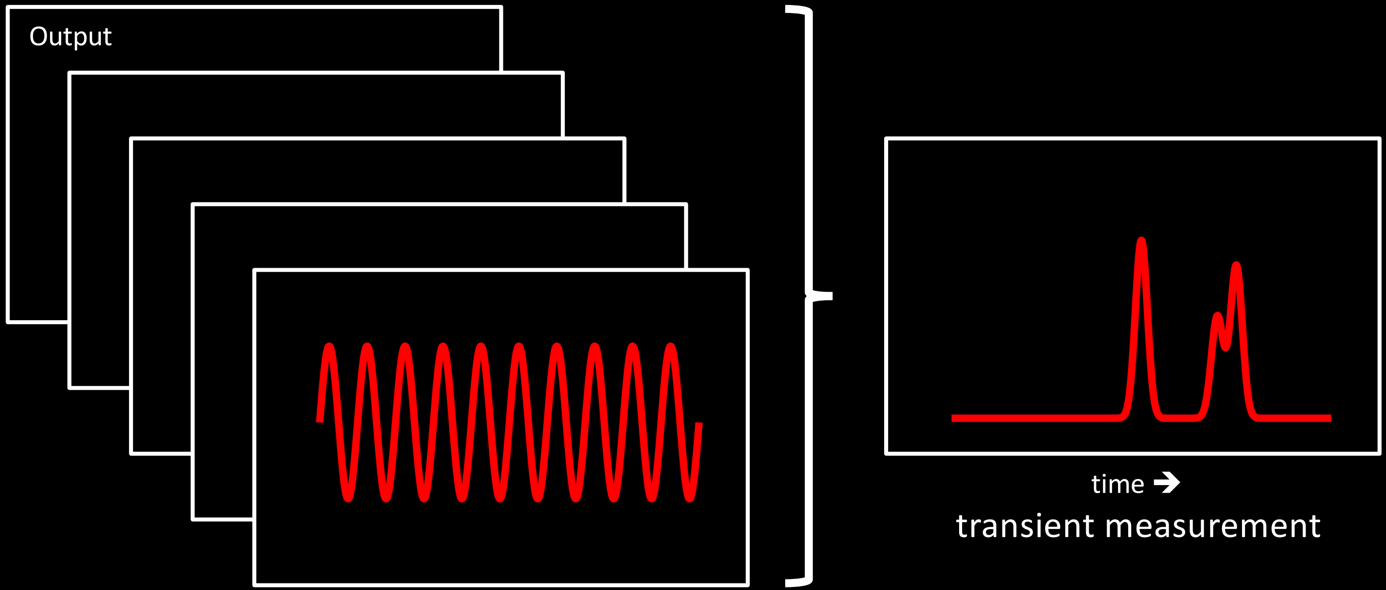

transient sensing technologies

	optical coherence tomography	streak camera	single-photon avalanche diodes	time-of-flight cameras	avalanche photodiode
temporal resolution	1 femtosecond (10^{-15} secs)	1 picosecond (10^{-12} secs)	100 picosecond (10^{-10} secs)	1 nanosecond (10^{-9} secs)	10 nanoseconds (10^{-8} secs)
frame rate	quadrillion fps	trillion fps	10 billion fps	billion fps	100 million fps
distance travelled	1 micron (10^{-6} meters)	1 millimeter (10^{-3} meters)	10 centimeters (10^{-1} meters)	1 meter (10^{0} meters)	10 meters (10^{1} meters)

transient sensing technologies

	optical coherence tomography	streak camera	single-photon avalanche diodes	time-of-flight cameras	avalanche photodiode
temporal resolution	1 femtosecond (10^{-15} secs)	1 picosecond (10^{-12} secs)	100 picosecond (10^{-10} secs)	1 nanosecond (10^{-9} secs)	10 nanoseconds (10^{-8} secs)
frame rate	quadrillion fps	trillion fps	10 billion fps	billion fps	100 million fps
distance travelled	1 micron (10^{-6} meters)	1 millimeter (10^{-3} meters)	10 centimeters (10^{-1} meters)	1 meter (10^{0} meters)	10 meters (10^{1} meters)

transient sensing technologies

	optical coherence tomography	streak camera	single-photon	time-of-flight	avalanche
temporal resolution	1 femtosecond (10^{-15} secs)	1 picosecond (10^{-12} secs)			
frame rate	quadrillion fps	trillion fps			
distance travelled	1 micron (10^{-6} meters)	1 millimeter (10^{-3} meters)	10 centimeters (10^{-1} meters)	1 meter (10^{0} meters)	10 meters (10^{1} meters)

Velten et al. [2012]

transient sensing technologies

Micro Photon Devices

	optical coherence tomography	streak camera	single-photon avalanche diodes	time-of-flight cameras	avalanche photodiode
temporal resolution			100 picosecond (10^{-10} secs)	1 nanosecond (10^{-9} secs)	10 nanoseconds (10^{-8} secs)
frame rate			10 billion fps	billion fps	100 million fps
distance travelled	1 micron (10^{-6} meters)	1 millimeter (10^{-3} meters)	10 centimeters (10^{-1} meters)	1 meter (10^{0} meters)	10 meters (10^{1} meters)

transient sensing technologies

	optical coherence tomography	streak camera	single-photon	time-of-flight	avalanche photodiode
temporal resolution	1 femtosecond (10^{-15} secs)	1 picosecond (10^{-12} secs)			10 nanoseconds (10^{-8} secs)
frame rate	quadrillion fps	trillion fps			100 million fps
distance travelled	1 micron (10^{-6} meters)	1 millimeter (10^{-3} meters)	10 centimeters (10^{-1} meters)	1 meter (10^{0} meters)	10 meters (10^{1} meters)

Velodyne VLS-128

transient sensing technologies

	optical coherence tomography	streak camera	single-photon avalanche diodes	time-of-flight cameras	avalanche photodiode
temporal resolution	1 femtosecond (10^{-15} secs)	1 picosecond (10^{-12} secs)	100 picosecond (10^{-10} secs)	1 nanosecond (10^{-9} secs)	10 nanoseconds (10^{-8} secs)
frame rate	quadrillion fps	trillion fps	10 billion fps	billion fps	100 million fps
distance travelled	1 micron (10^{-6} meters)	1 millimeter (10^{-3} meters)	10 centimeters (10^{-1} meters)	1 meter (10^{0} meters)	10 meters (10^{1} meters)

transient sensing technologies

	optical coherence tomography	streak camera	single-photon	time-of-flight	avalanche photodiode
temporal resolution	1 femtosecond (10^{-15} secs)				0 nanoseconds (10^{-8} secs)
frame rate	quadrillion fps				100 million fps
distance travelled	1 micron (10^{-6} meters)	1 millimeter (10^{-3} meters)	10 centimeters (10^{-1} meters)	1 meter (10^{0} meters)	10 meters (10^{1} meters)

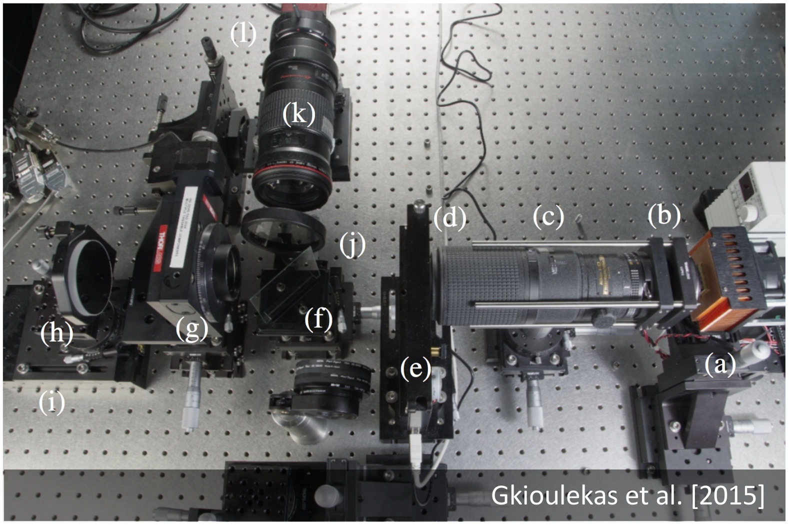

Gkioulekas et al. [2015]

transient sensing technologies

	optical coherence tomography	streak camera	single-photon avalanche diodes	time-of-flight cameras	avalanche photodiode
temporal resolution				1 nanosecond (10^{-9} secs)	10 nanoseconds (10^{-8} secs)
frame rate				billion fps	100 million fps
distance travelled	1 micron (10^{-6} meters)	1 millimeter (10^{-3} meters)	10 centimeters (10^{-1} meters)	1 meter (10^{0} meters)	10 meters (10^{1} meters)

Heide et al. [2013]

spectrum of transient sensing technologies

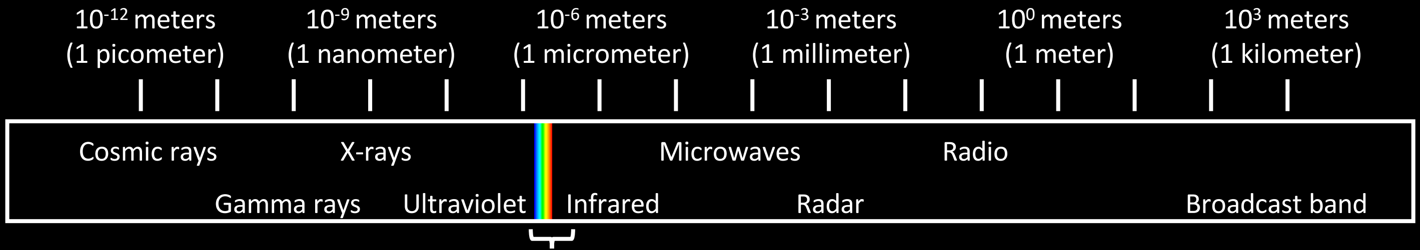

- **optical coherence tomography**
- **streak camera**
- **single-photon avalanche diodes**
- **time-of-flight cameras**
- **avalanche photodiode**

spectrum of transient sensing technologies

10^{-12} meters (1 picometer)	10^{-9} meters (1 nanometer)	10^{-6} meters (1 micrometer)	10^{-3} meters (1 millimeter)	10^{0} meters (1 meter)	10^{3} meters (1 kilometer)

Cosmic rays | X-rays | Microwaves | Radio

Gamma rays | Ultraviolet | Infrared | Radar | Broadcast band

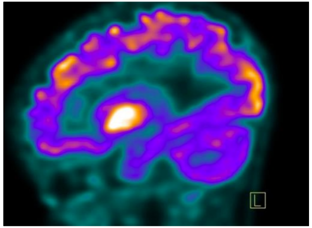

Positron Emission Tomography (PET) Scan

- optical coherence tomography
- streak camera
- single-photon avalanche diodes
- time-of-flight cameras
- avalanche photodiode

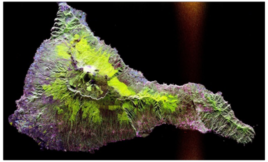

Synthetic Aperture Radar (SAR)

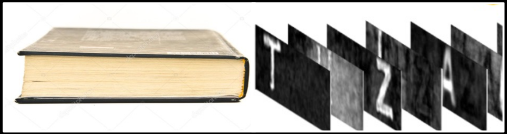

Terahertz Time-gated Imaging [Redo-Sanchez et al., 2015]

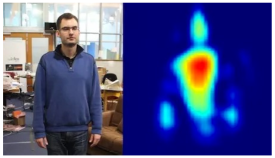

Seeing through Walls [Adib et al., 2015]

spectrum of transient sensing technologies

10^{-12} meters (1 picometer)	10^{-9} meters (1 nanometer)	10^{-6} meters (1 micrometer)	10^{-3} meters (1 millimeter)	10^{0} meters (1 meter)	10^{3} meters (1 kilometer)

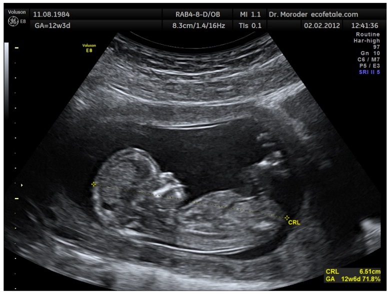

Ultrasound Imaging

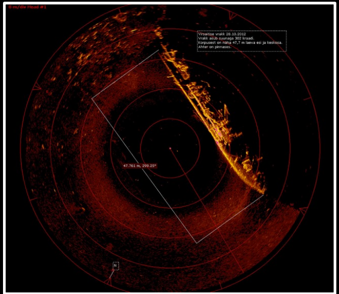

SONAR
(Sound Navigation and Ranging)

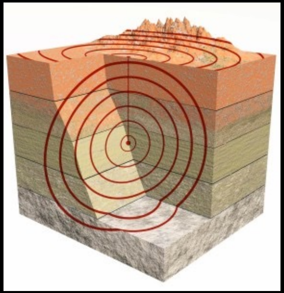

Seismic Imaging

overview

- Time-resolved imaging
- Single-photon avalanche diodes (SPADs)
- Single-photon lidar
- Non-line-of-sight imaging
- Imaging through scattering media

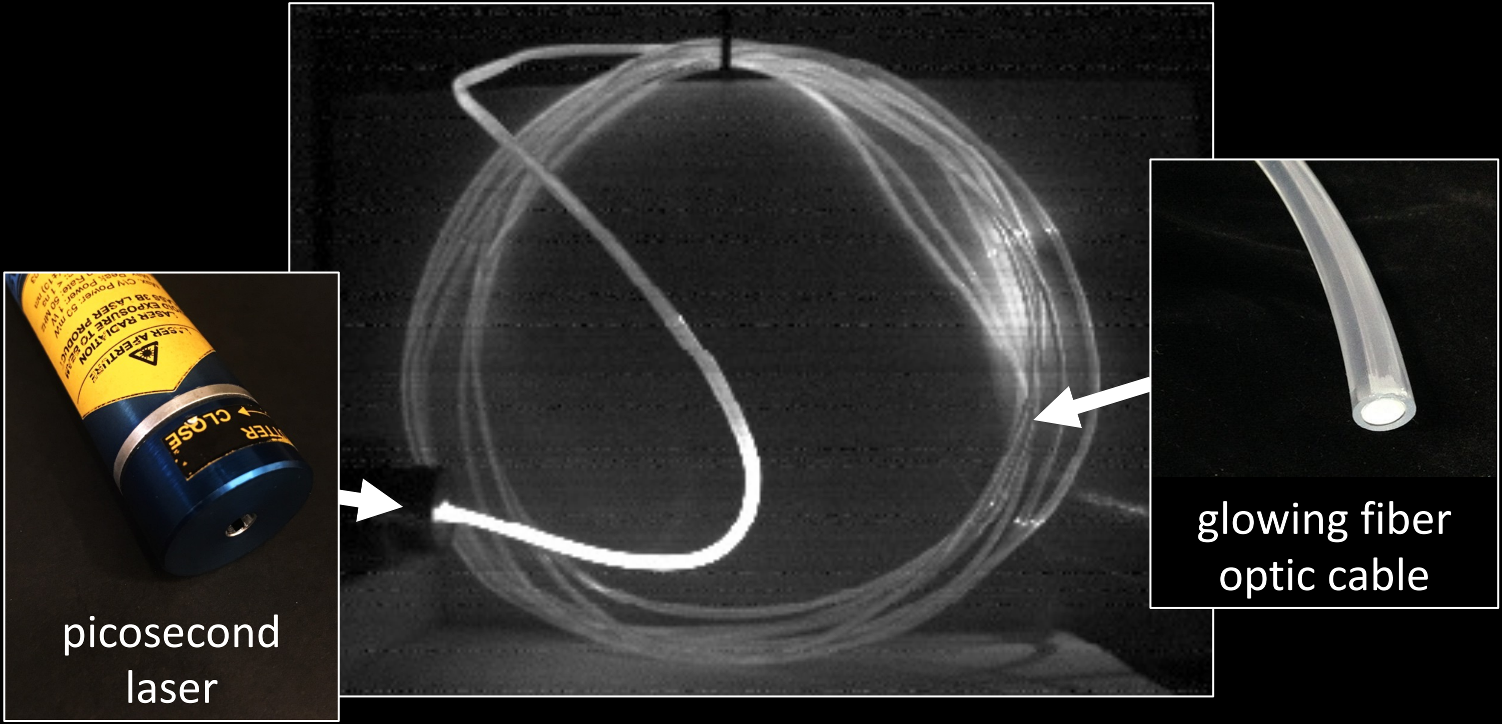

regular image

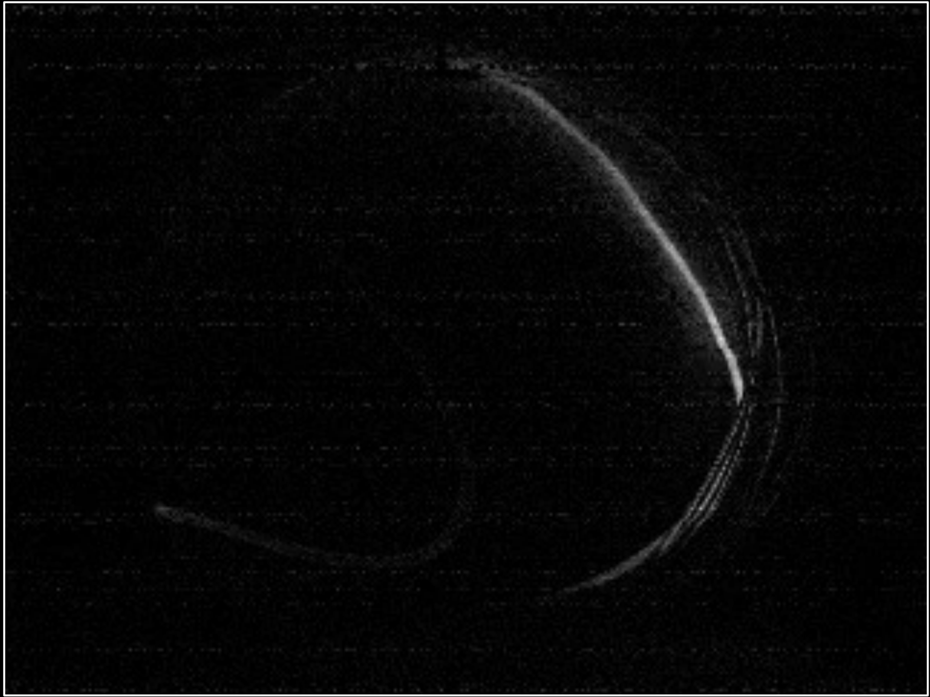

transient image

Applications

Autonomous Navigation
Image by Wikimedia Commons

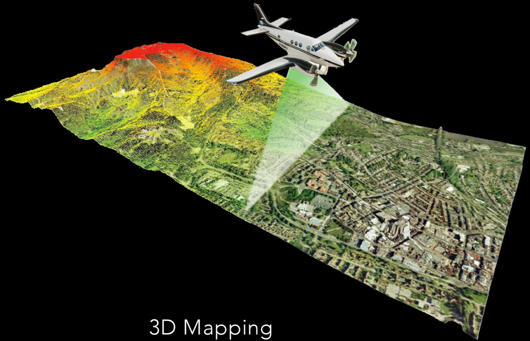

3D Mapping
Image by LIDAR-America

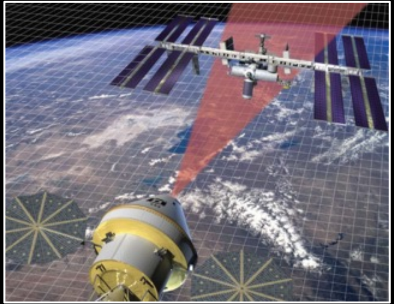

Space Station Docking
Image by NASA

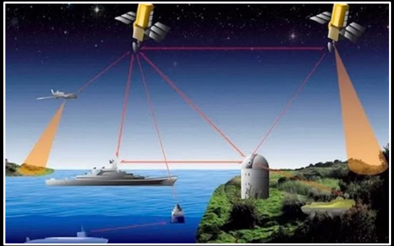

Optical Communications
Image by Siasat Daily

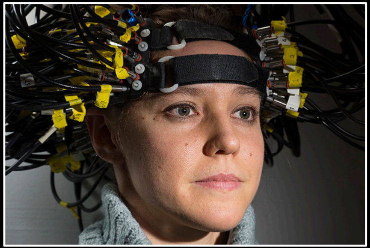

Biomedical Imaging
Image by Washington University

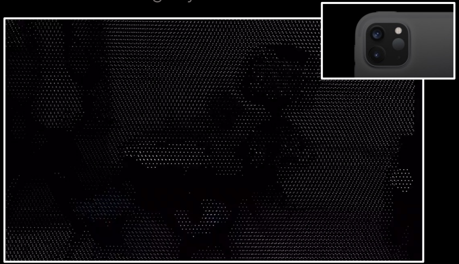

Consumer Electronics (2020 iPad Pro)
Video by Tim Fields

Appl

Au
Ima

king

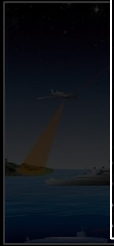

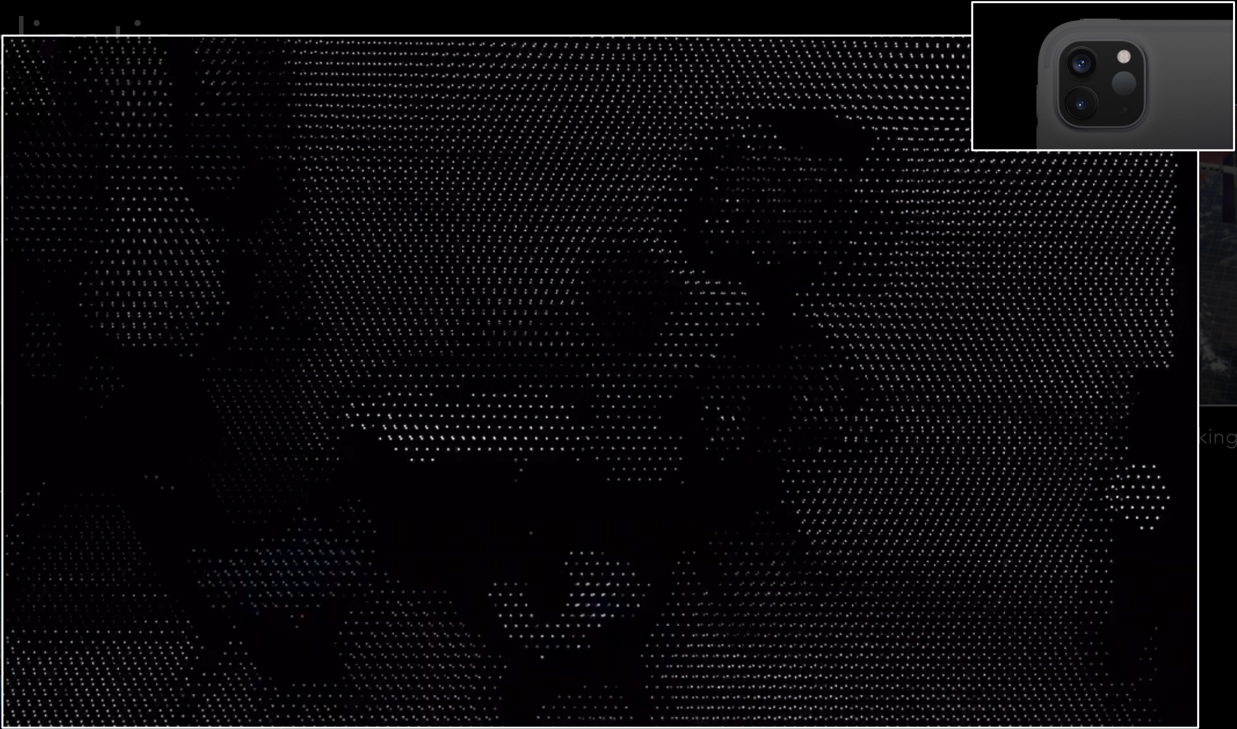

Optical Communications
Image by Siasat Daily

Biomedical Imaging
Image by Washington University

Consumer Electronics (2020 iPad Pro)
Video by Tim Fields

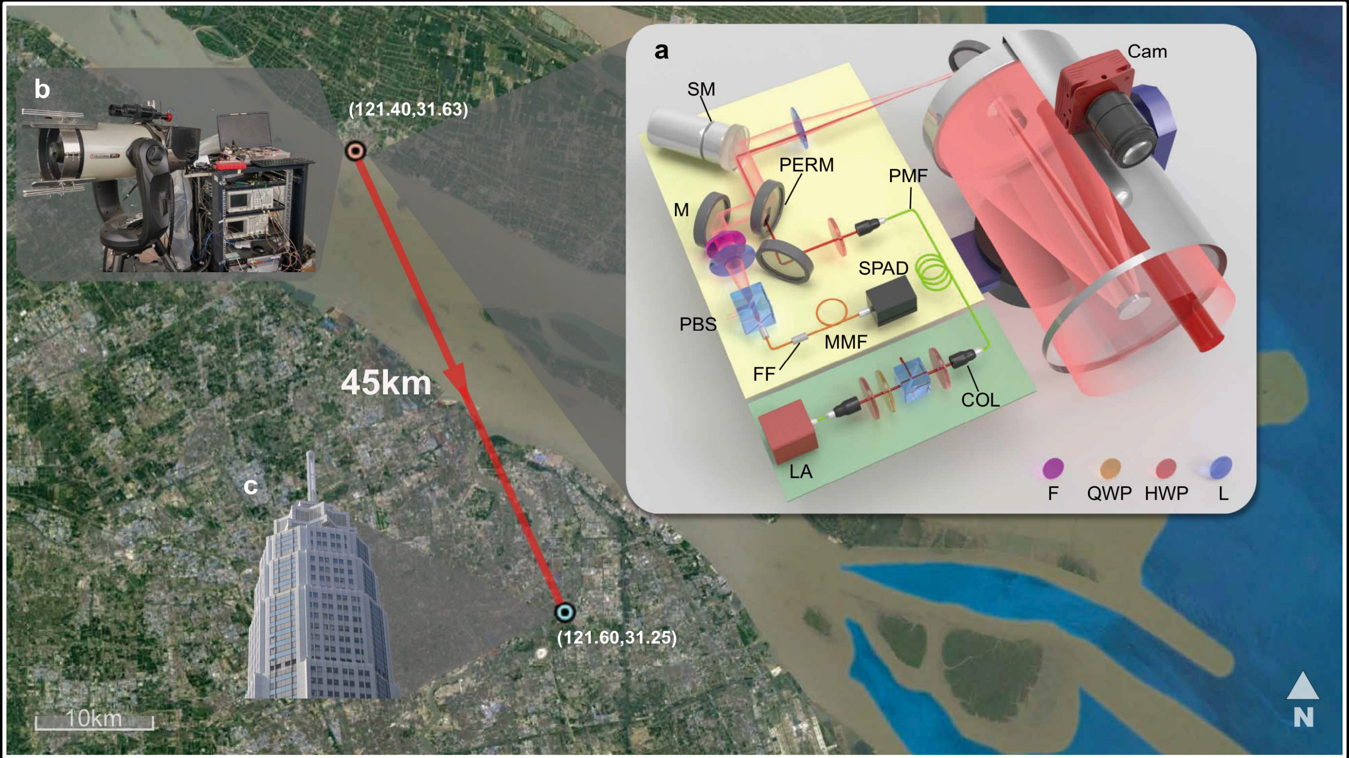

a
Cam
SM
PERM
PMF
M
SPAD
PBS
MMF
FF
COL
LA
F
QWP
HWP
L
b
(121.40,31.63)
45km
c
(121.60,31.25)
10km
N

International Laser Ranging Service (ILRS)

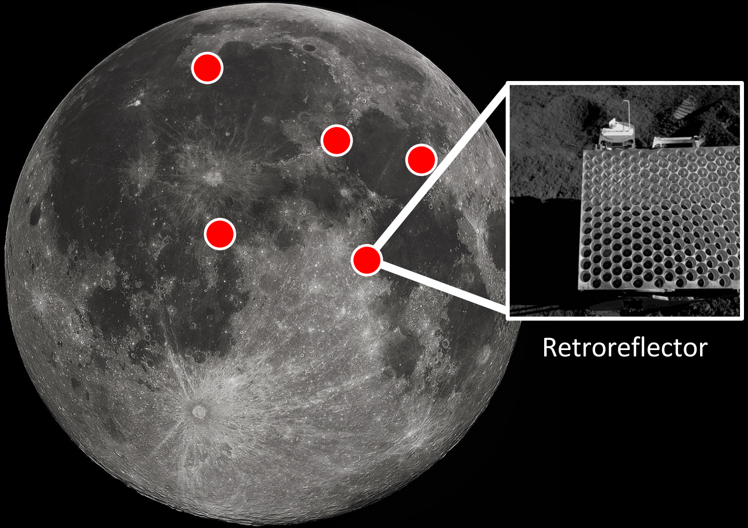

Lunar Laser Ranging (LLR)

- Location of Lunar Retroreflector

single-photon avalanche diode (SPAD)

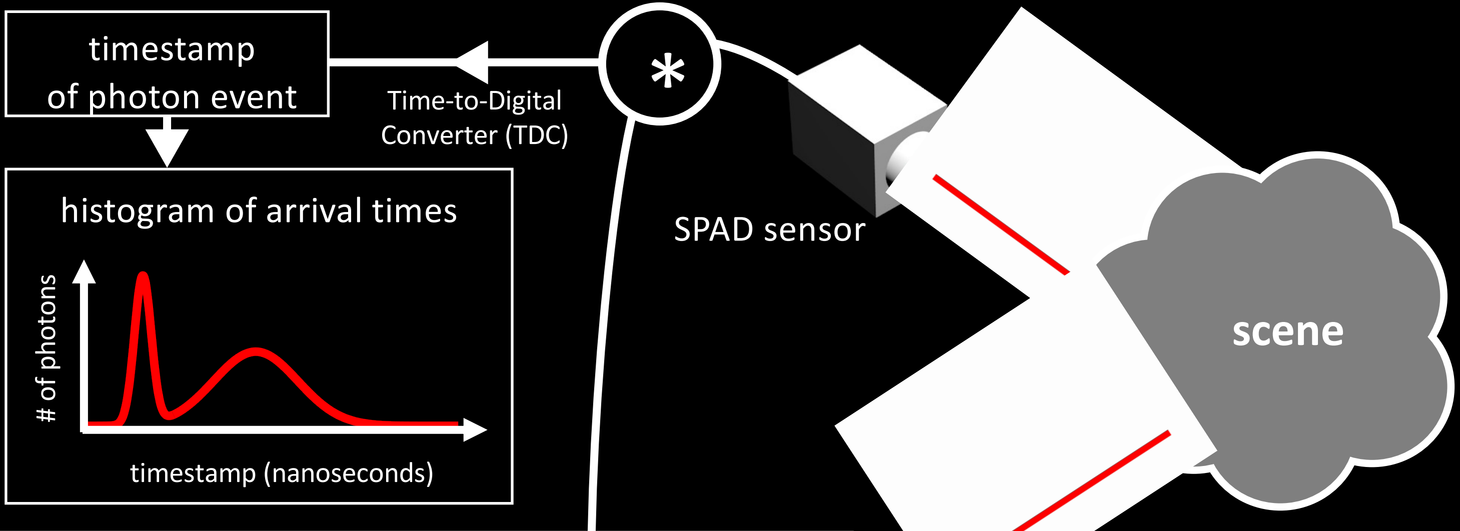

SPAD properties:

- Each photon timestamped with 60 ps precision
- Measure up to 10 million photons a second

linear array of SPADs

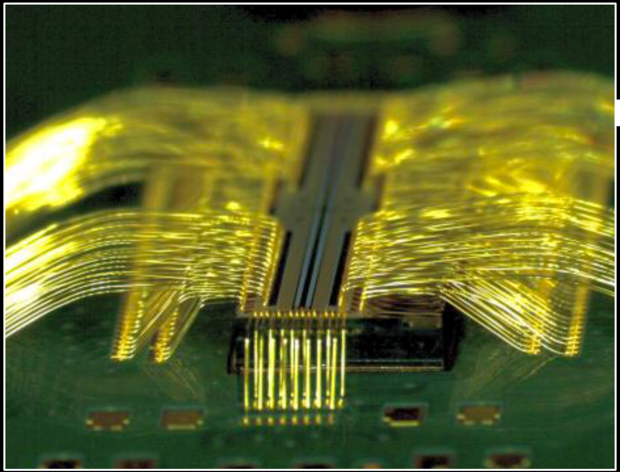

LinoSPAD from FastTree 3D

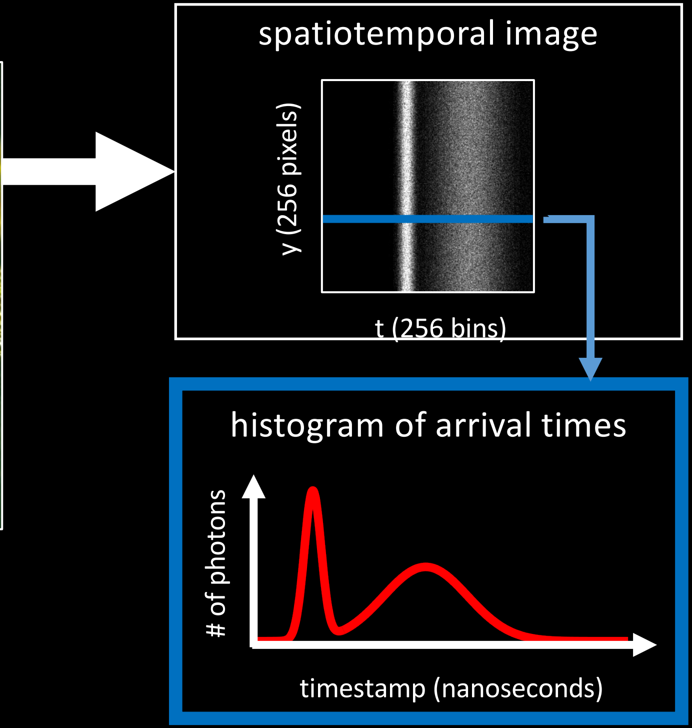

scanning procedure

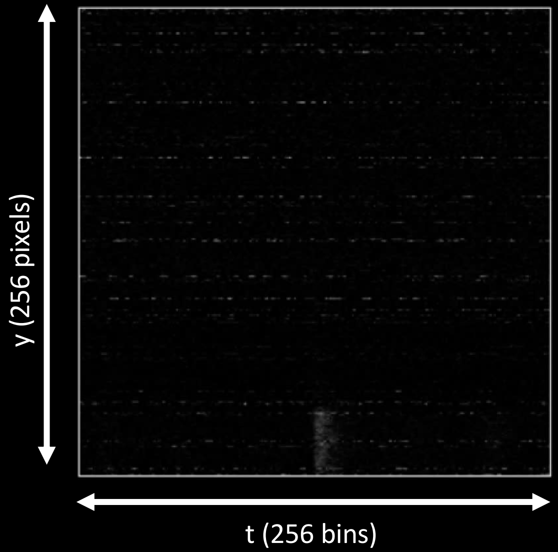

SPAD output

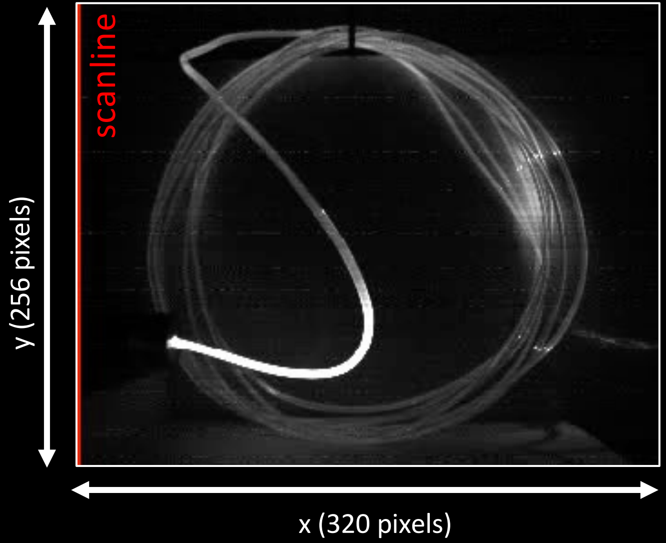

regular image

scanning procedure

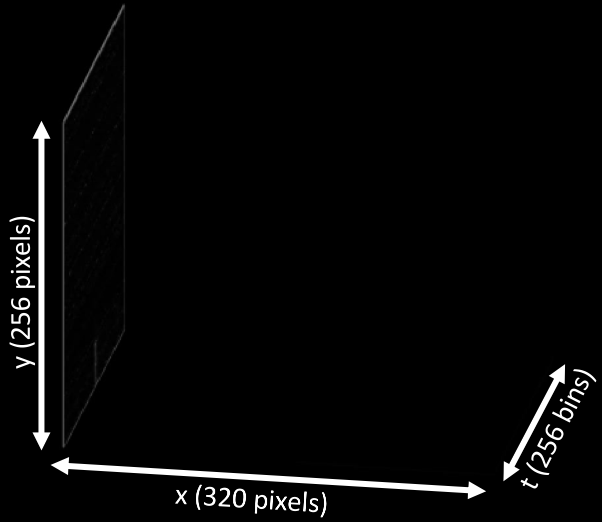

transient image

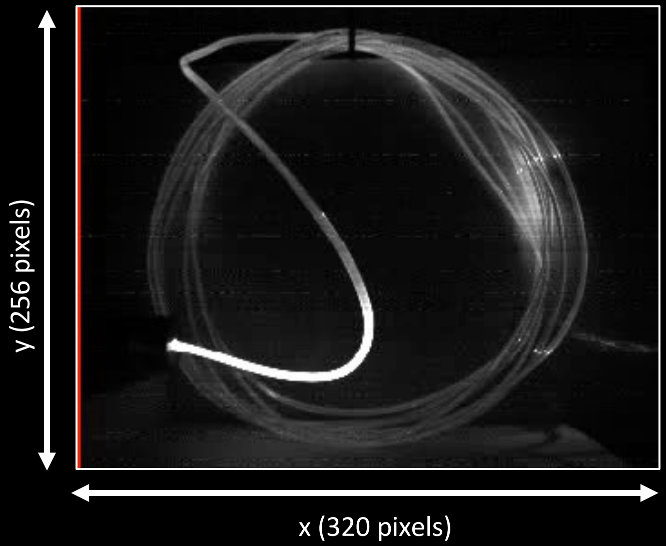

regular image

scanning procedure

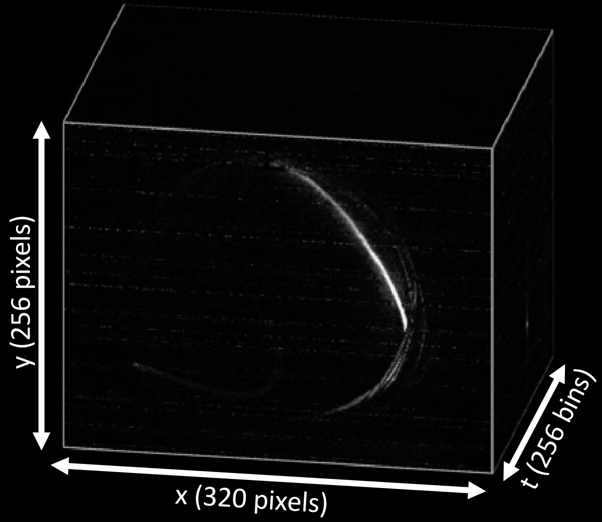

transient image

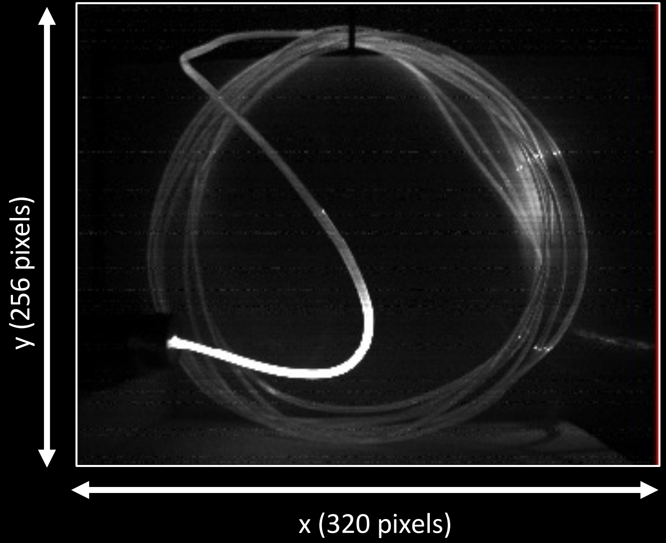

regular image

overview

- Time-resolved imaging
- Single-photon avalanche diodes (SPADs)
- Single-photon lidar
- Non-line-of-sight imaging
- Imaging through scattering media

LiDAR

LiDAR

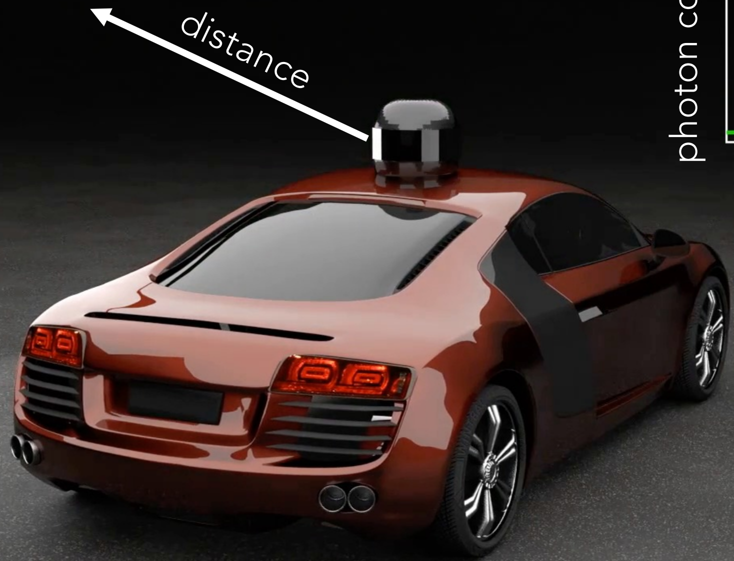

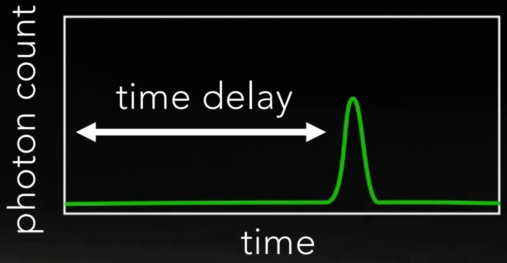

LiDAR

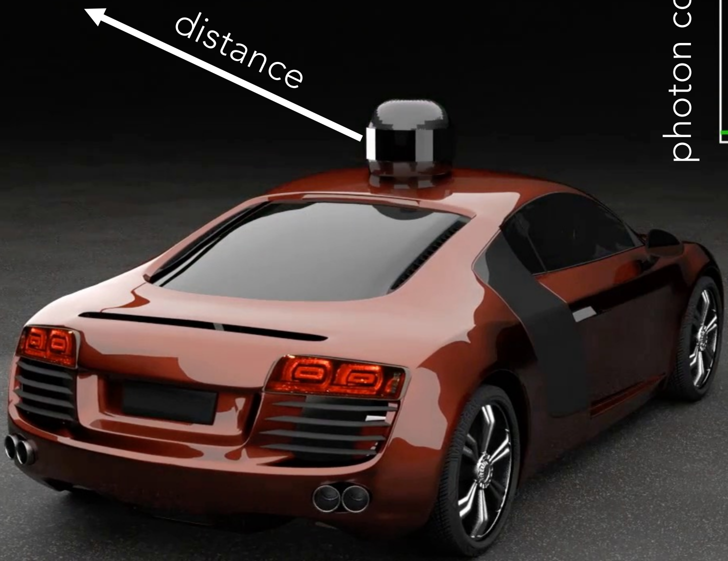

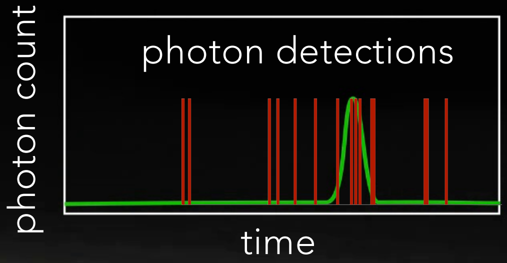

LiDAR

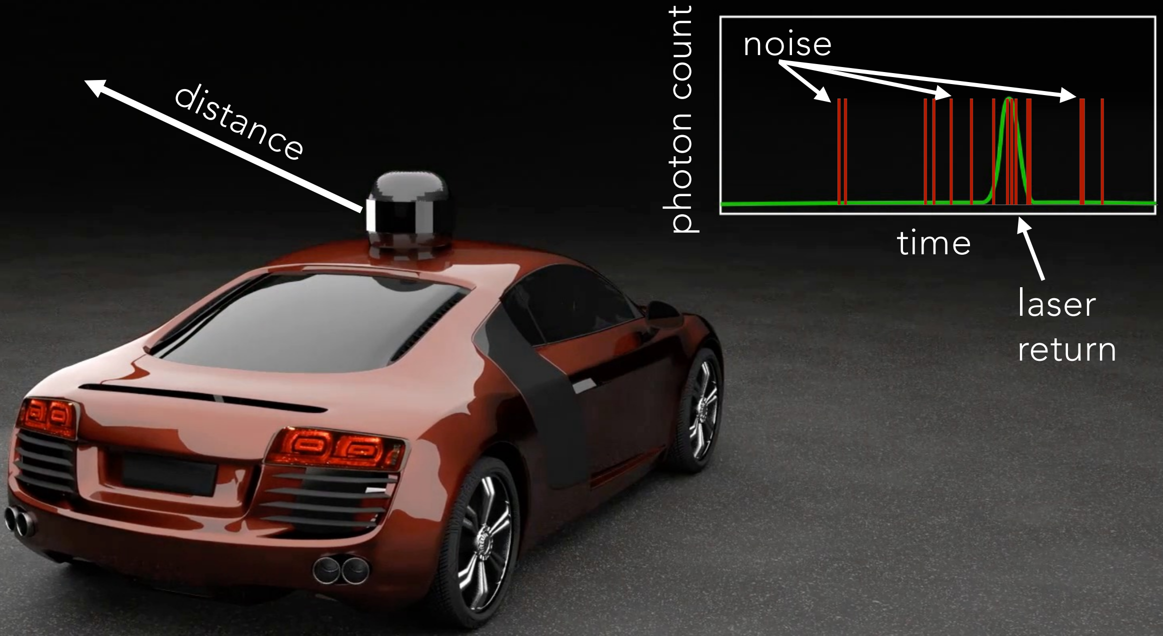

LiDAR

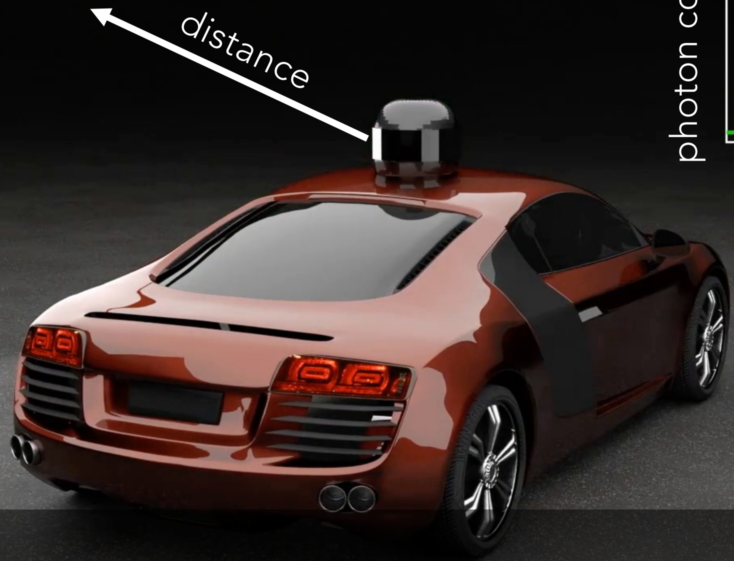

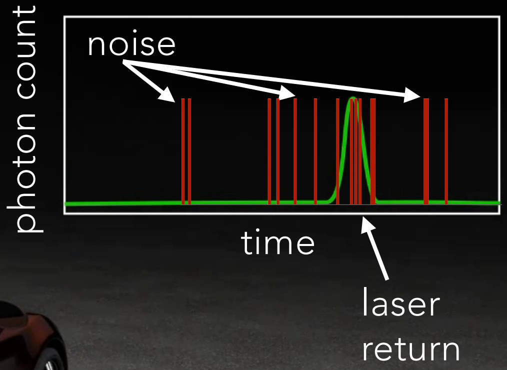

Challenges

1. Light efficiency / photon sensitivity (determines range)
2. High-speed time stamping (determines accuracy)
3. Computational algorithms (determines range and accuracy)

Challenges

1. Light efficiency / photon sensitivity (determines range)

- intensity of returned light falls off with $1/d^2$, i.e. very quickly!
- emit as much light as possible - ***fundamentally limited by eye safety*** (in most applications)
- detect as much light as possible, ideally individual photons

Challenges

2. High-speed time stamping (determines accuracy)

- speed of light is ~300,000,000 m/s
- 1 m = 3.3 ns; 1 cm = 33 ps; 1 mm = 3.3 ps
- need picosecond-accurate time-stamping → usually high-end electronics, but also done with ASICs, FPGAs

Challenges

3. Computational algorithms (determines range and accuracy)

   • robust depth estimation from single photon per pixel!

Kirmani et al. "First-photon Imaging", Science 2014

conventional method first-photon imaging

(Single-photon) Avalanche Photodiodes

<u>Linear mode</u> (i.e., avalanche photodiode or *APD*):

acts like a conventional photodiode with extremely high gain or amplification

time resolution >300 ps – 10 ns

<u>Geiger mode</u> (i.e., single-photon avalanche photodiode *SPAD*):

500x more sensitive, i.e. single-photon sensitive

time resolution ~50 ps

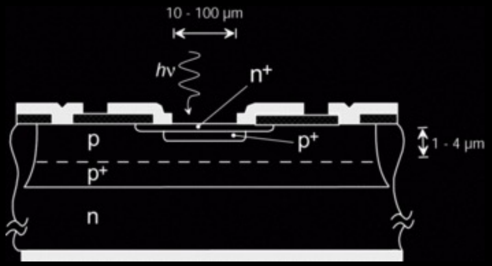

Semiconductor devices

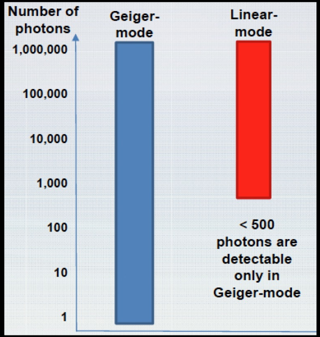

image by Princeton Lightwave

Single-Photon Avalanche Diodes

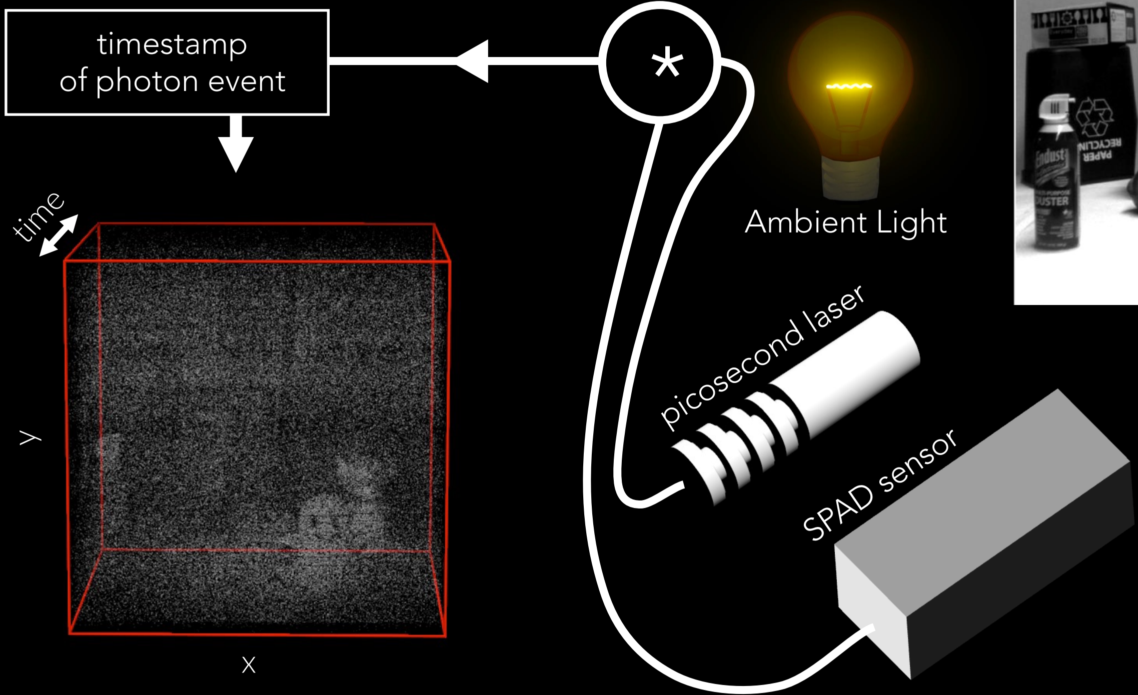

Measurements

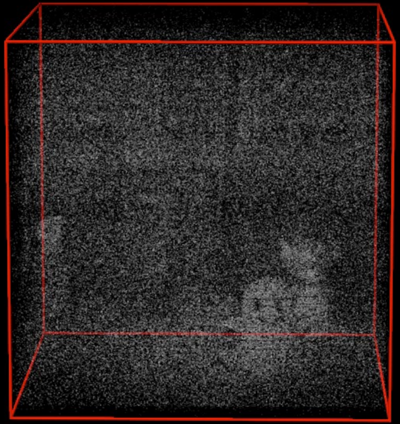

SPAD measurements
(256 x 256 x 1536)

Intensity image
(1024 x 1024)

Noisy <1 photon per pix., low spatial resolution

clean, high spatial resolution

How to fuse information from both?

Lindell et al., SIGGRAPH 2018

CNN Architecture for Depth Estimation (1 of 3)

SPAD measurements only

Input

SPAD measurements

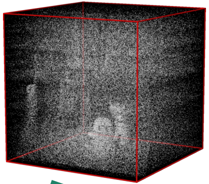

Intensity image

Output (1 of 3)

Geometry

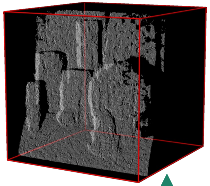

CNN Architecture (1 of 3)

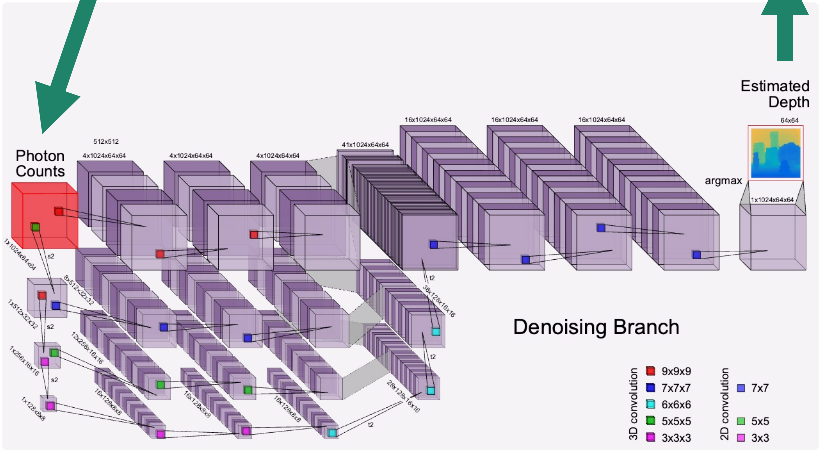

CNN Architecture for Depth Estimation (2 of 3)

Improved denoising by sensor fusion

Input

SPAD measurements

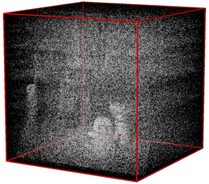

Intensity image

Output (1 of 3)

Geometry

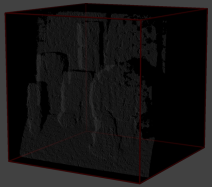

Output (2 of 3)

Geometry

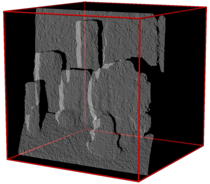

CNN Architecture (2 of 3)

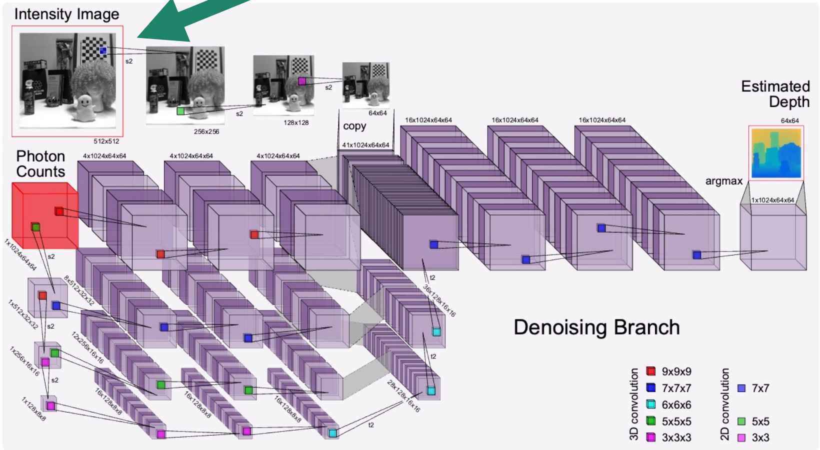

CNN Architecture for Depth Estimation (3 of 3)

Guided upsampling by sensor fusion

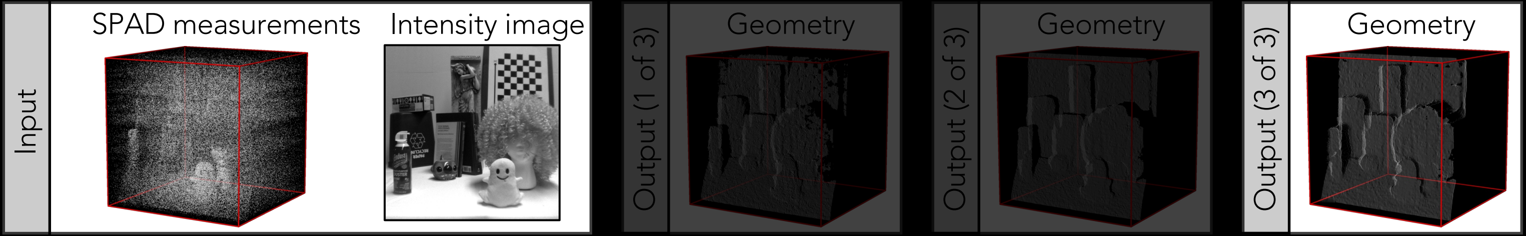

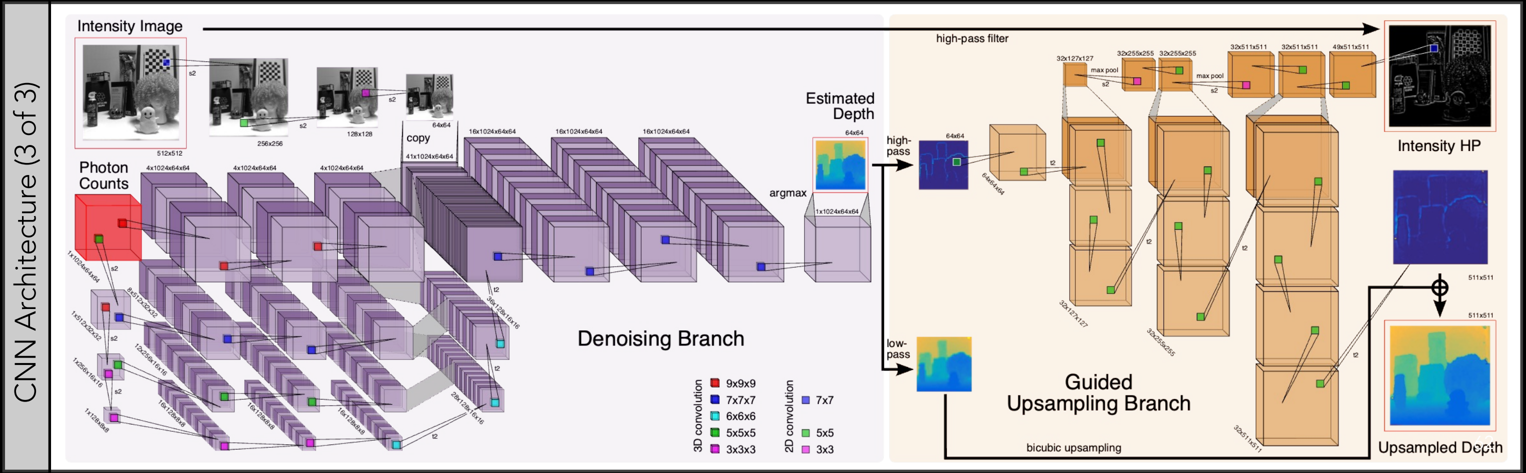

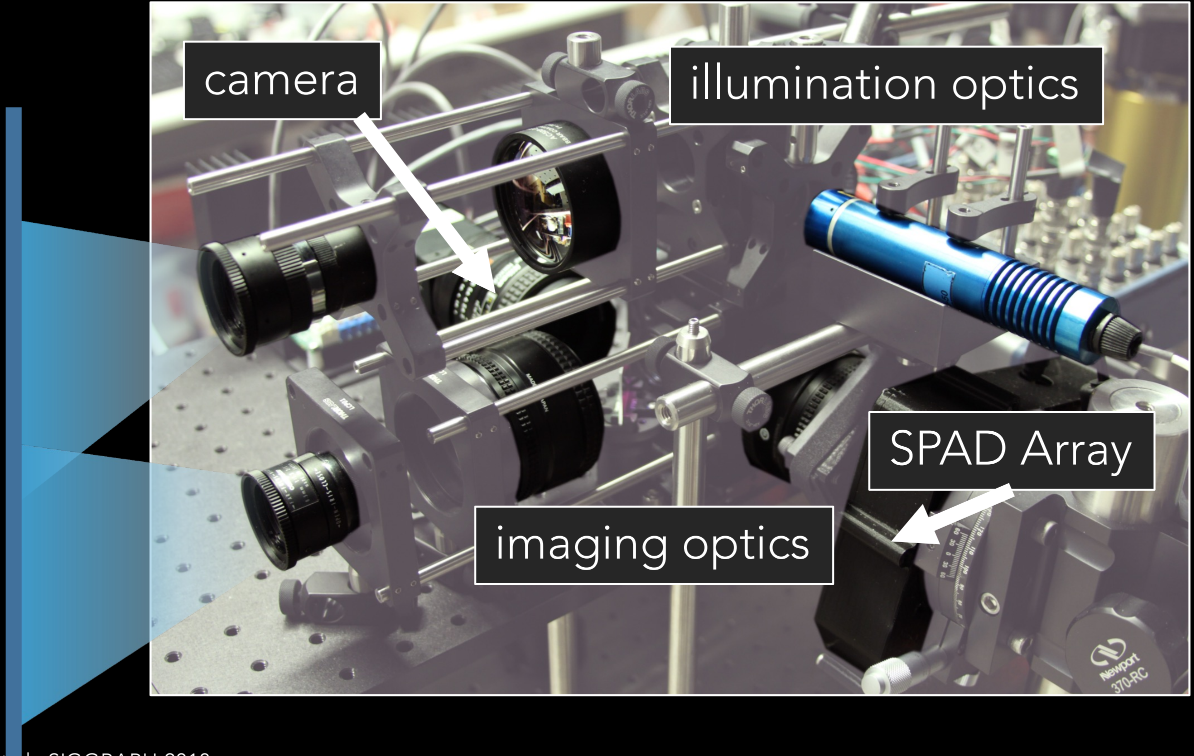

Lindell et al., SIGGRAPH 2018

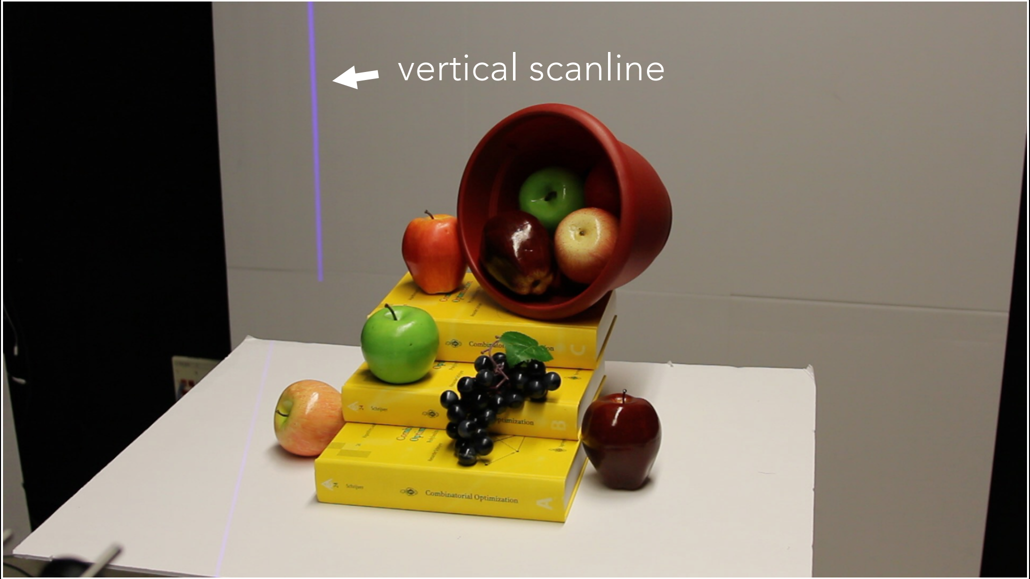

scan rate: 20 Hz

lights on

(note: laser illumination is too weak to observe visually while scanning under ambient light)

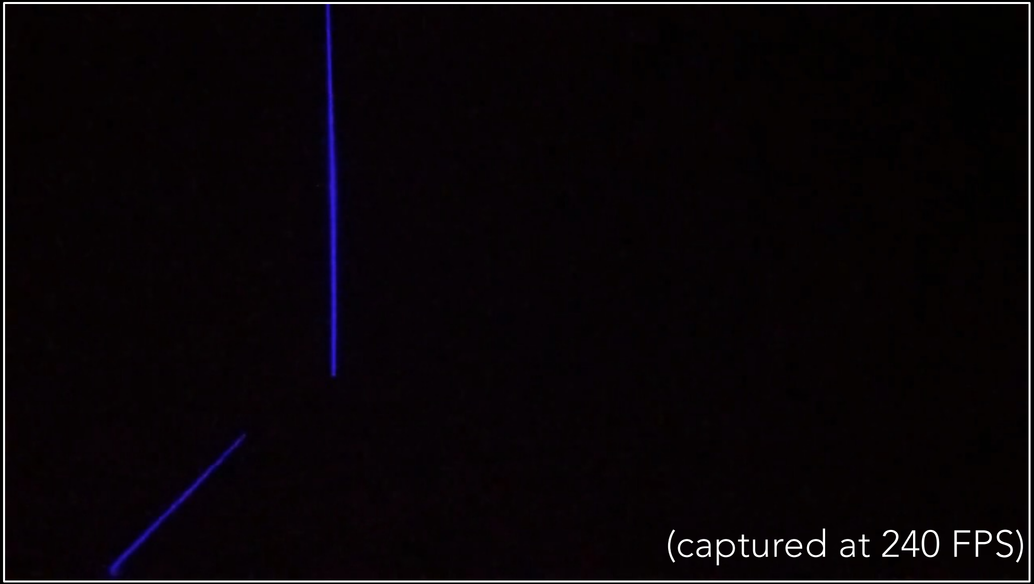

scan rate: 20 Hz

lights off

Lindell et al., SIGGRAPH 2018

Intensity image

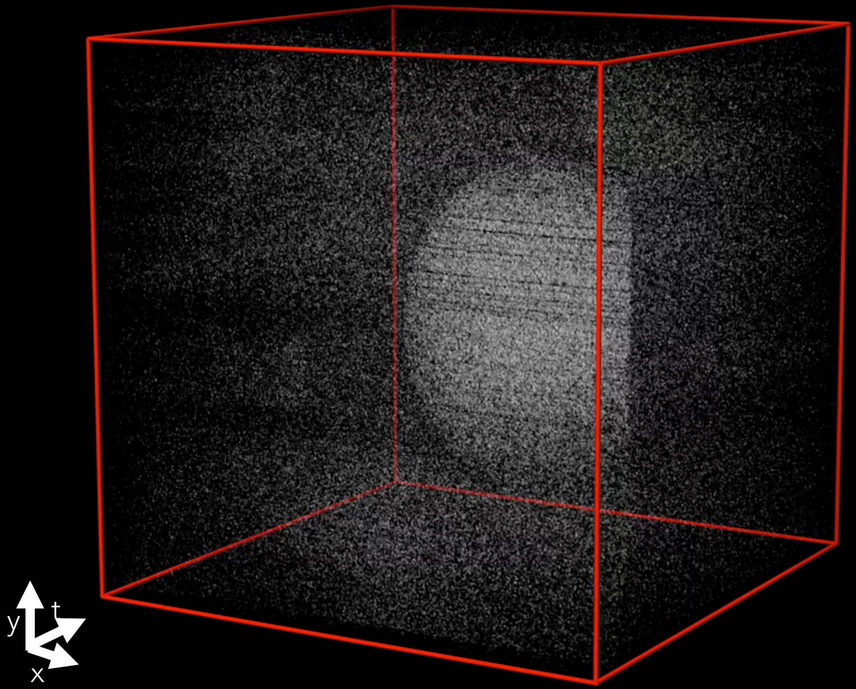

SPAD measurements (20 Hz)

Average per spatial position
0.64 Signal Detections
0.87 Background Detections

Lindell et al., SIGGRAPH 2018

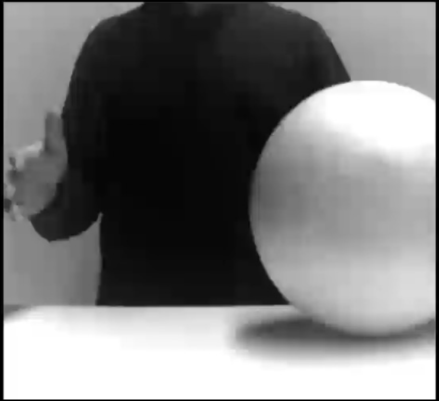

Intensity image

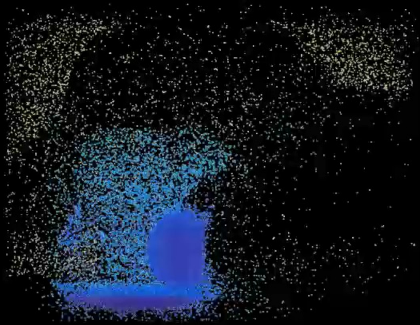

Log-matched filter

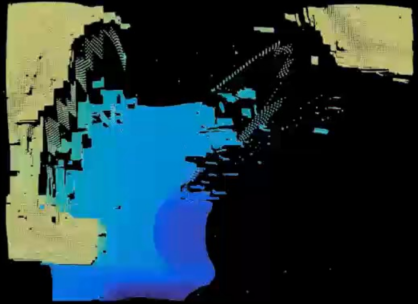

[Rapp and Goyal 2017]

Why is Deep Learning Useful Here?

- fusing complementary information from different sensors is not straightforward, but we can learn the mapping

- idea extends to other sensors: radar, thermal, …

- inverse method is fully differentiable → can attach higher-level tasks, such as classification (car, pedestrian, biker, …), and train end-to-end from photon counts to class label or control

overview

- Time-resolved imaging
- Single-photon avalanche diodes (SPADs)
- Single-photon lidar
- Non-line-of-sight imaging (part 1)
- Imaging through scattering media

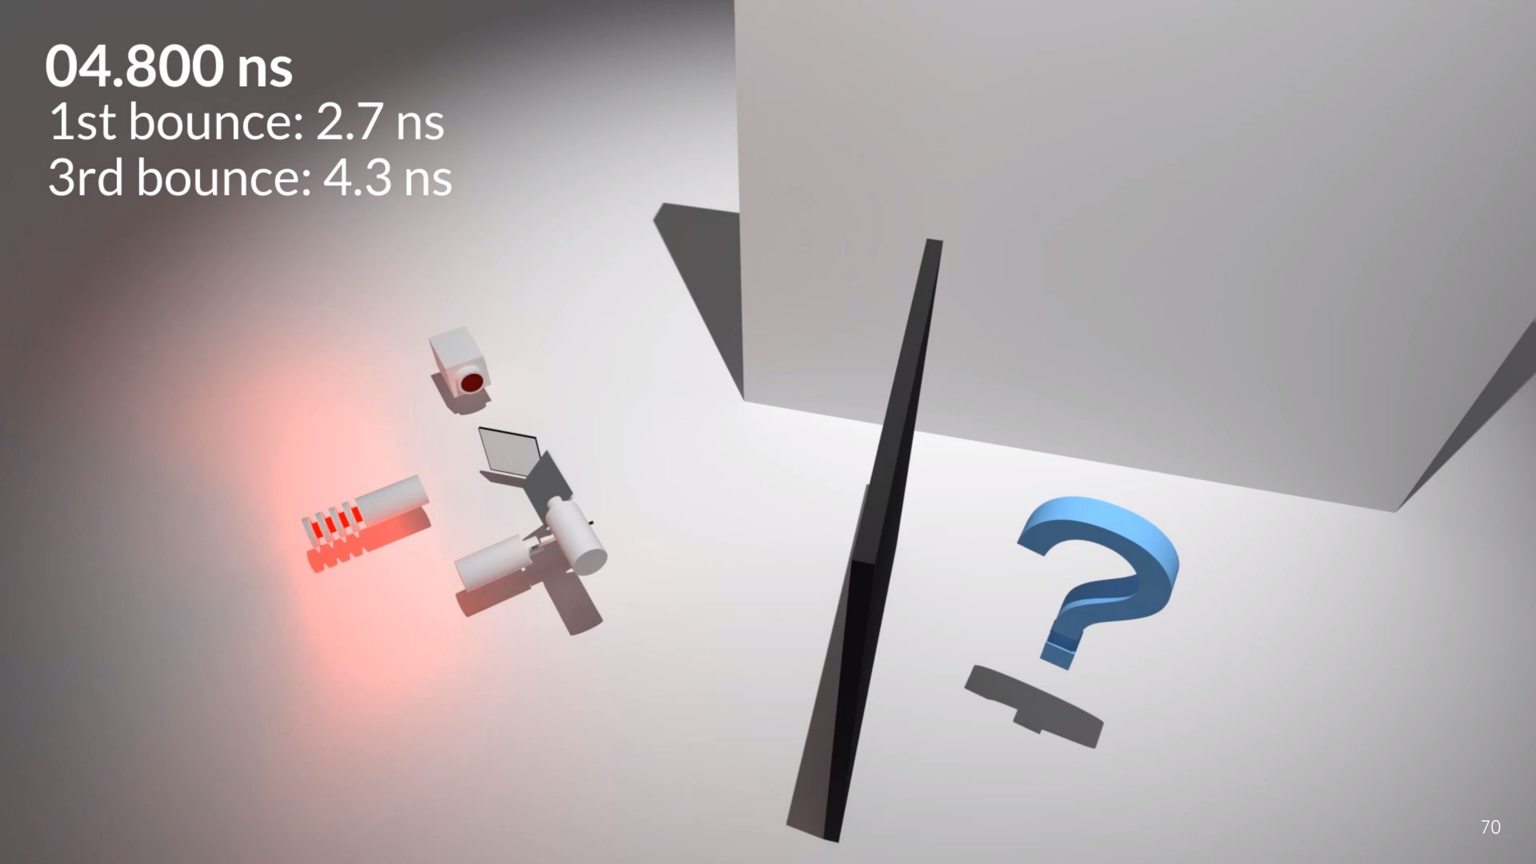
04.800 ns
1st bounce: 2.7 ns
3rd bounce: 4.3 ns

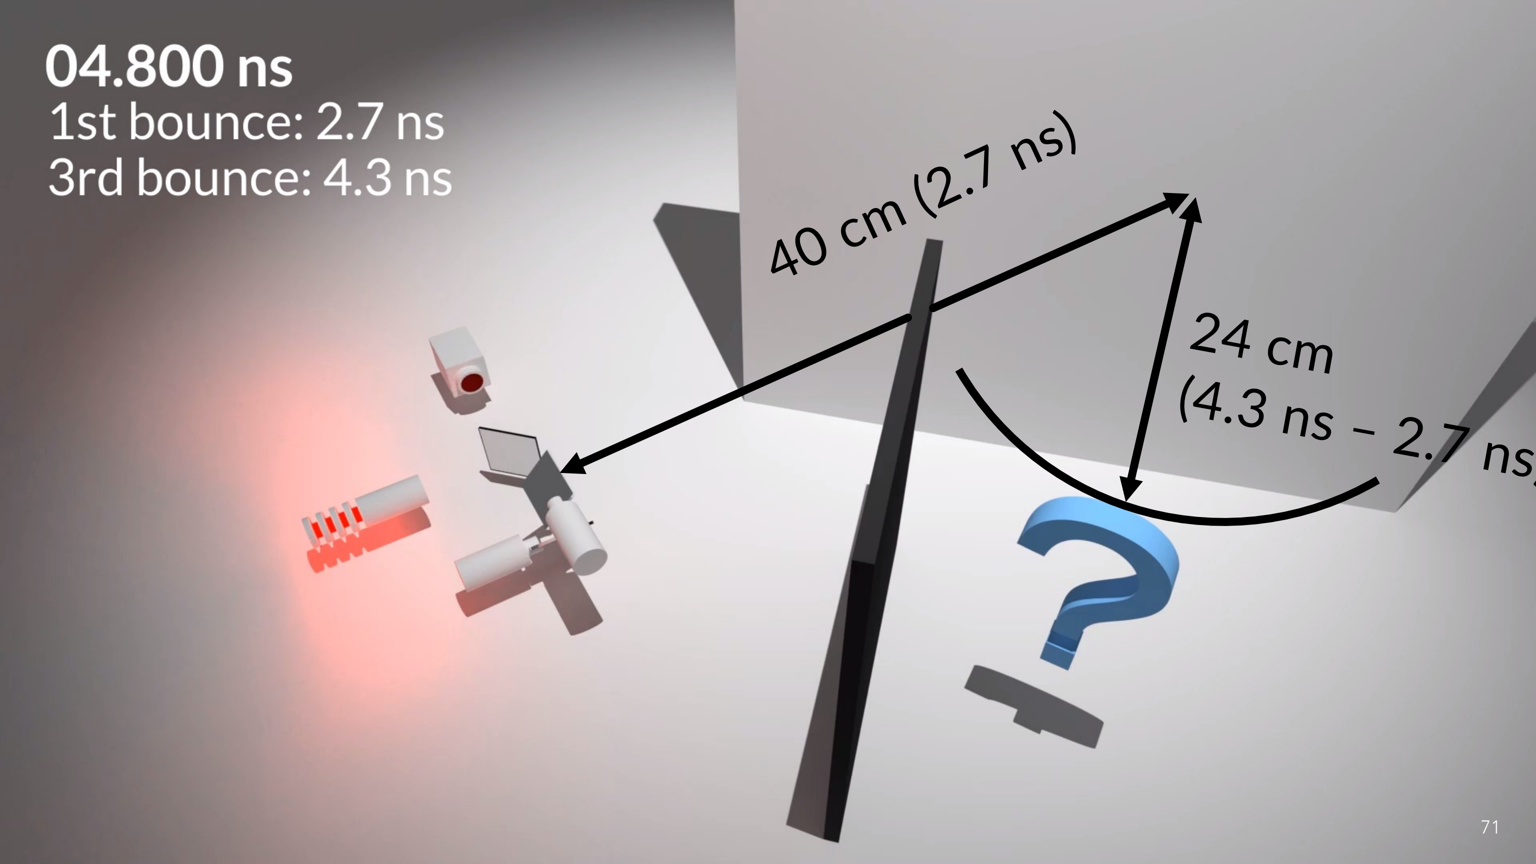
04.800 ns
1st bounce: 2.7 ns
3rd bounce: 4.3 ns
40 cm (2.7 ns)
24 cm
(4.3 ns − 2.7 ns

RAW histogram (10 FPS)

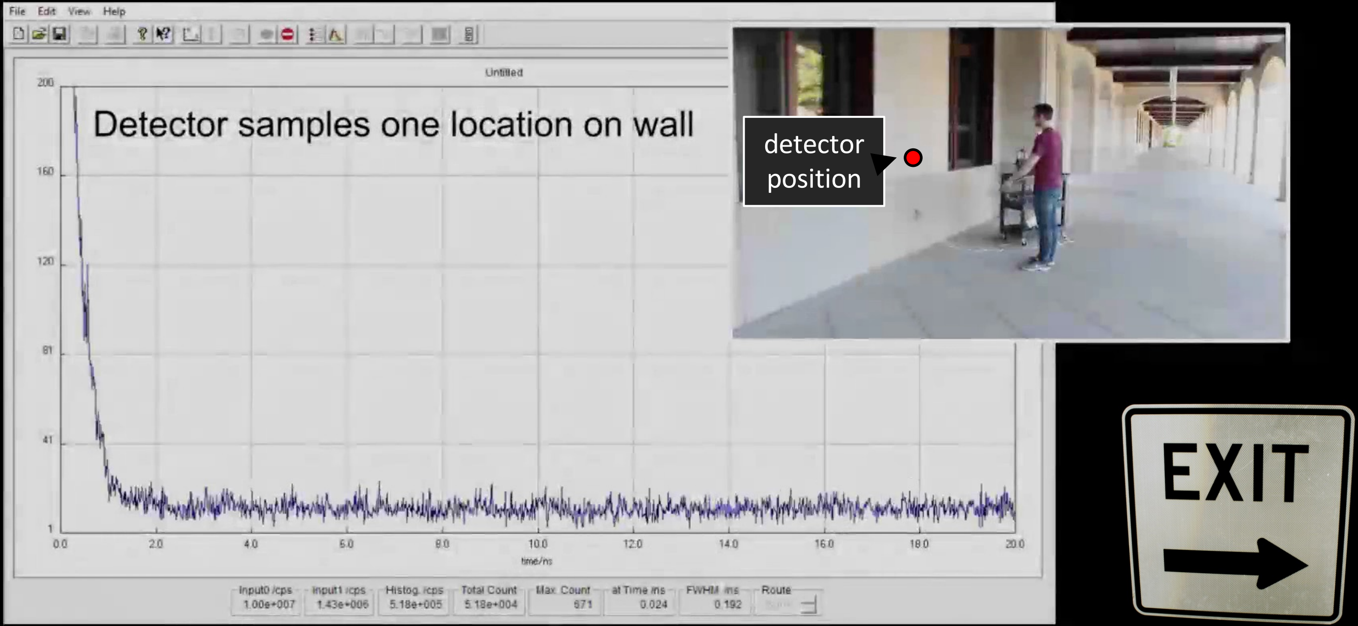

object

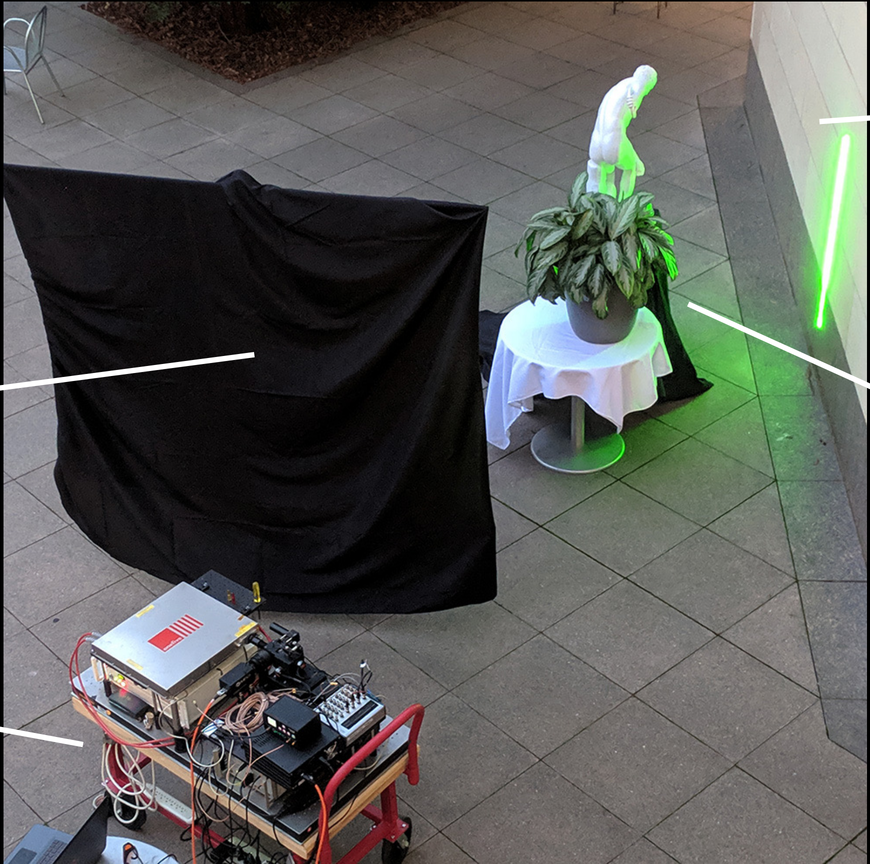

occluder

NLOS
imaging
system

wall

hidden
scene


scanning mirrors
laser
detector
PICOSECOND DELAYER
INPUT
NIM OUT
TTL OUT


resolution: 128 x 128
area: 2 m × 2 m

scene photo

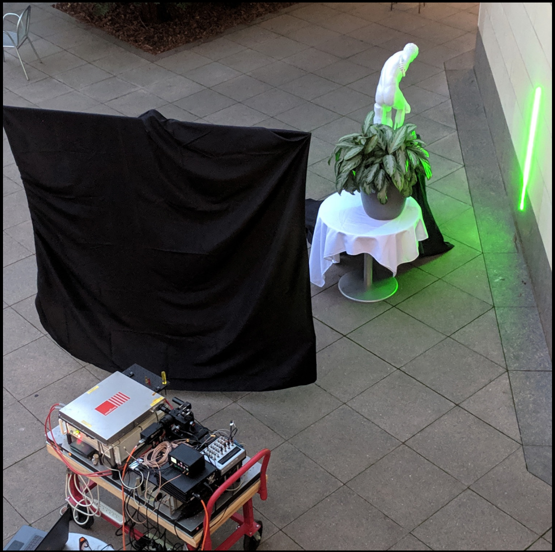

measurements

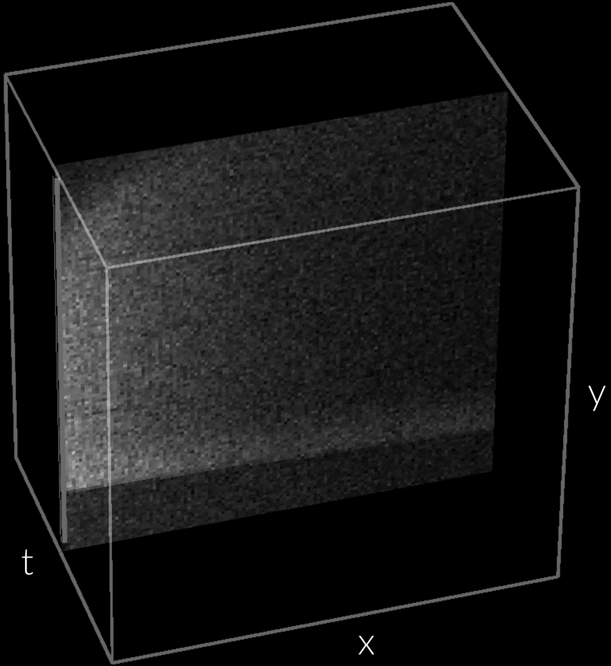

Lindell et al., SIGGRAPH 2019

scene photo

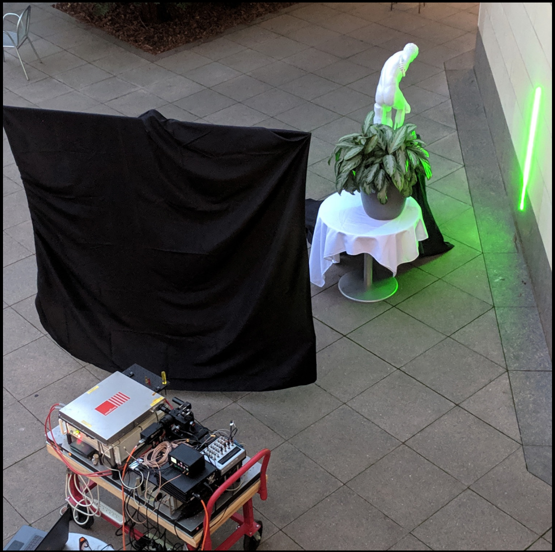

reconstruction

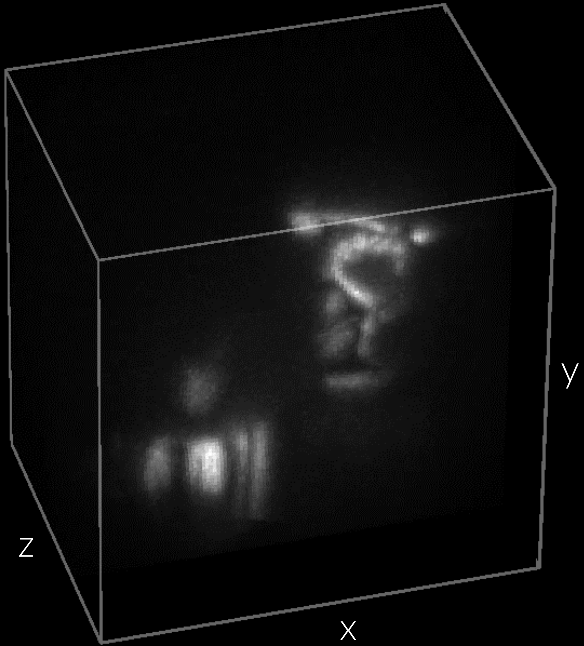

Lindell et al., SIGGRAPH 2019

Dimensions: 2 m x 2 m x 1.5 m

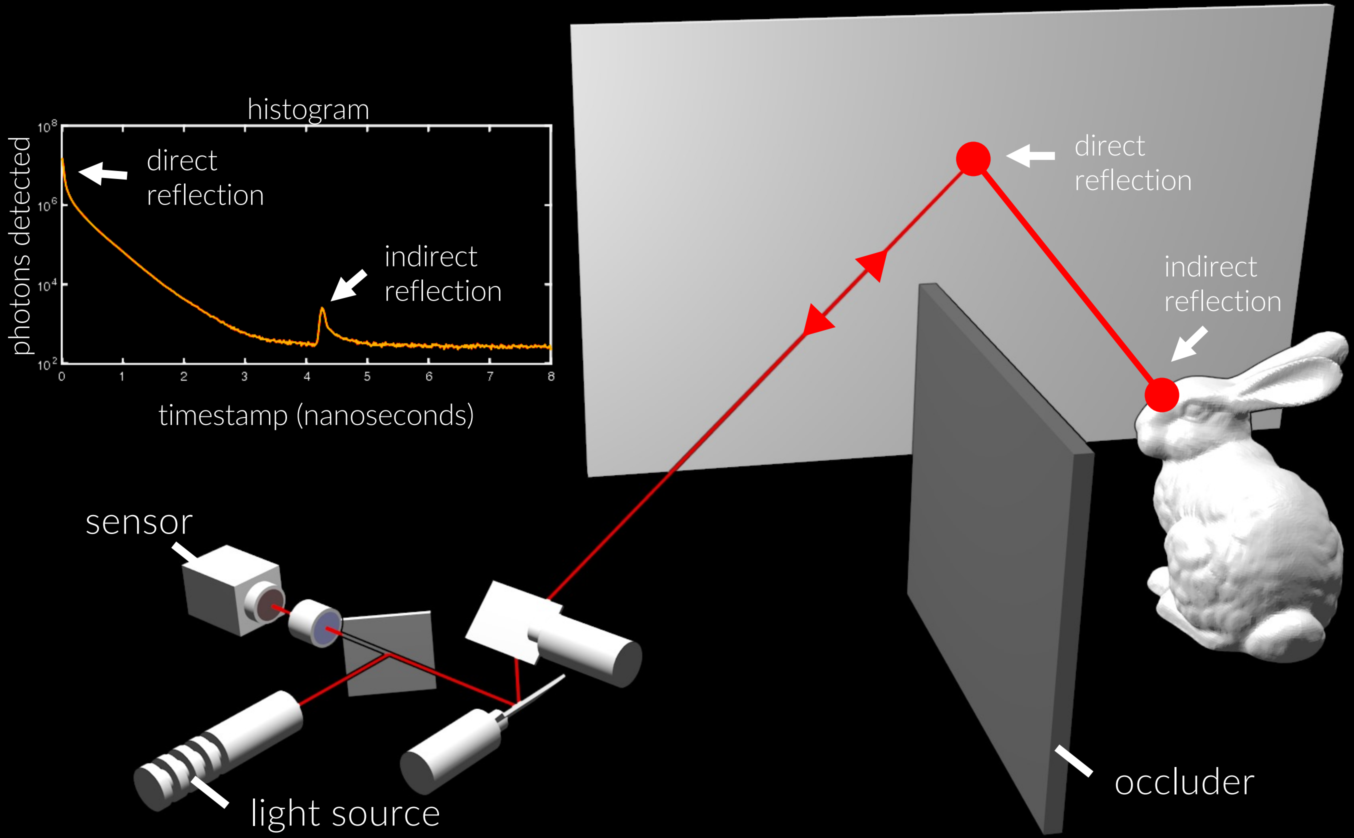
histogram
photons detected
timestamp (nanoseconds)
direct
reflection
indirect
reflection
10^8
10^6
10^4
10^2
0
1
2
3
4
5
6
7
8
direct
reflection
indirect
reflection
sensor
light source
occluder

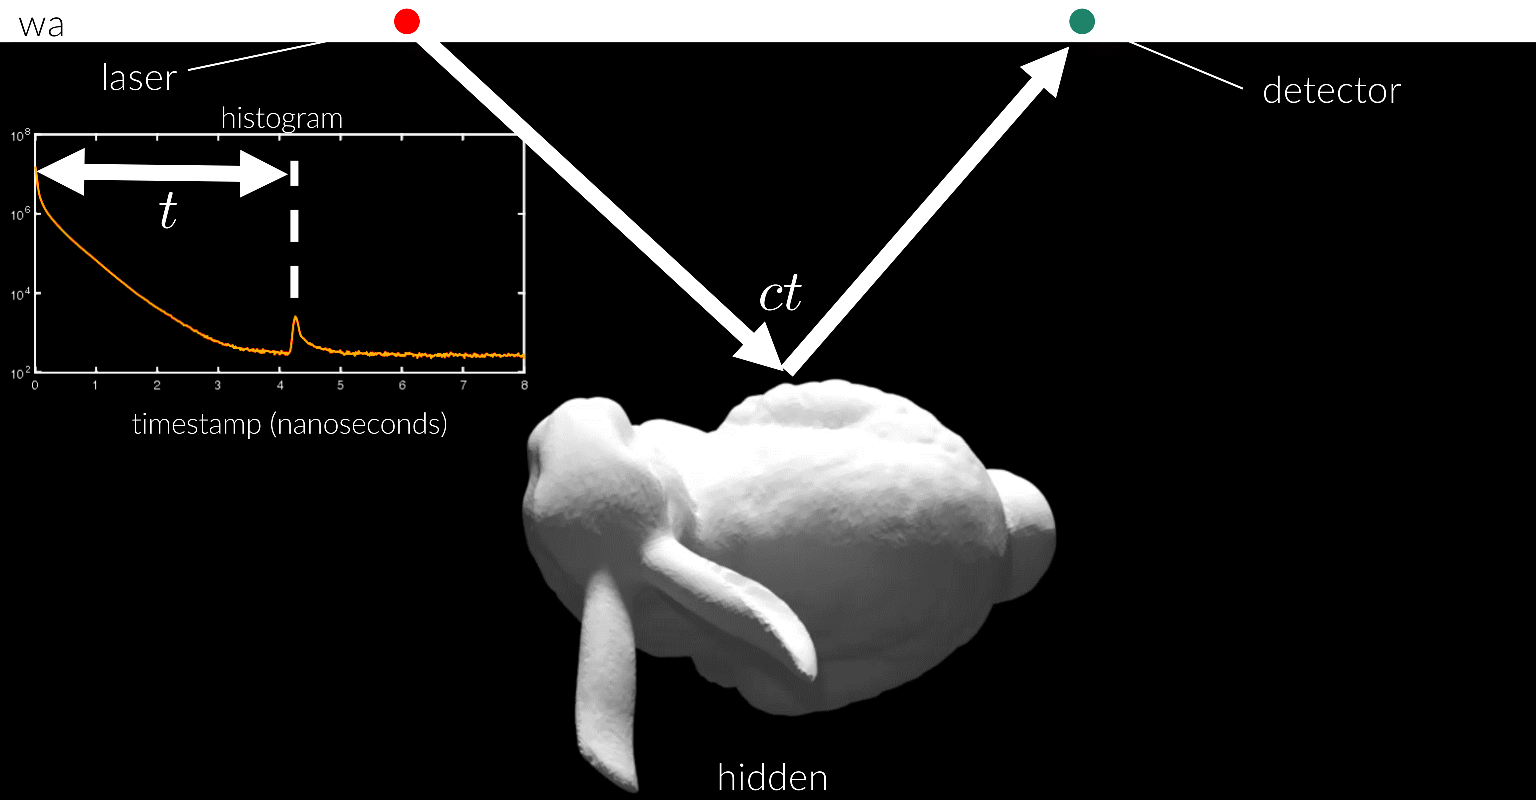

non-confocal sampling

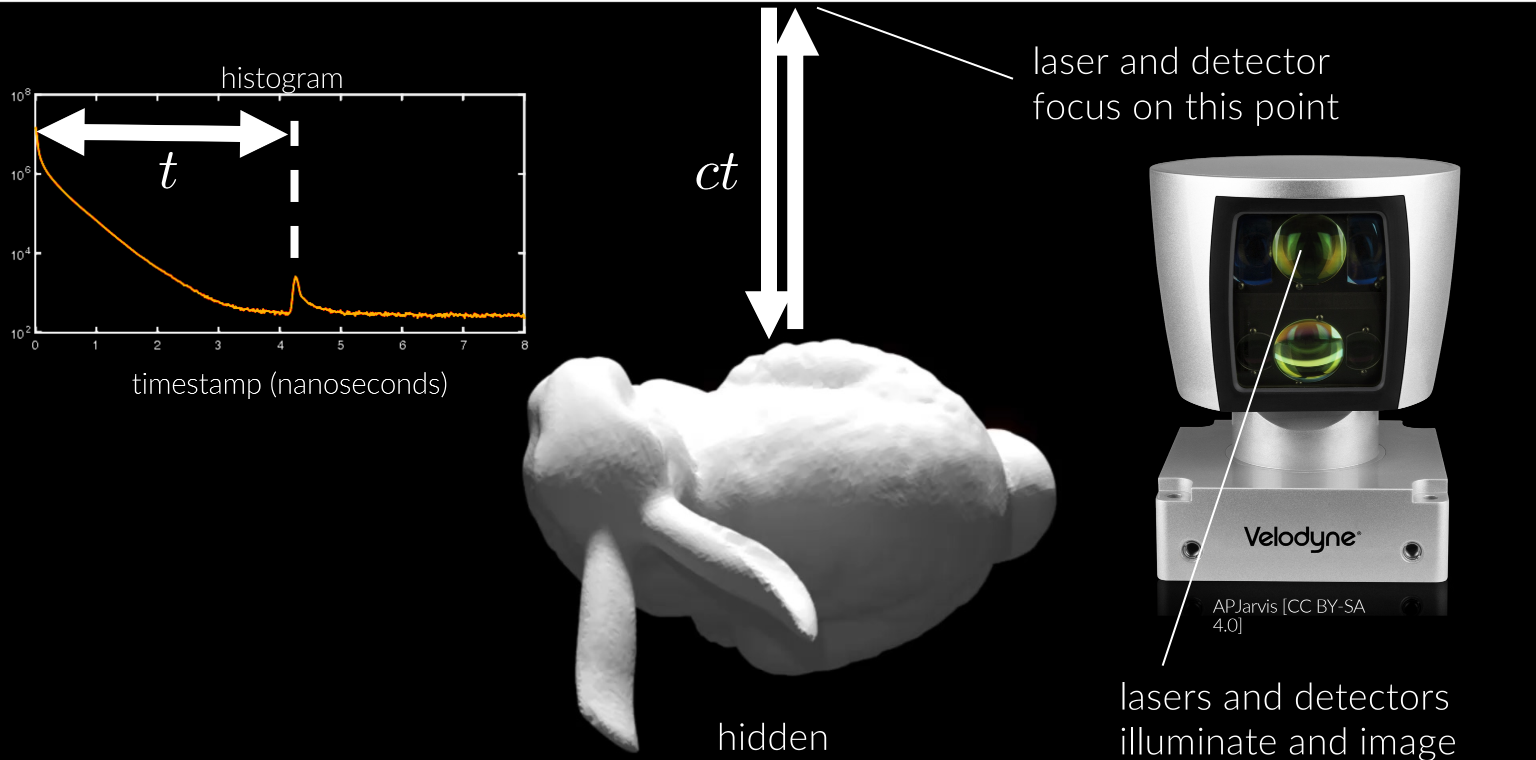

confocal sampling

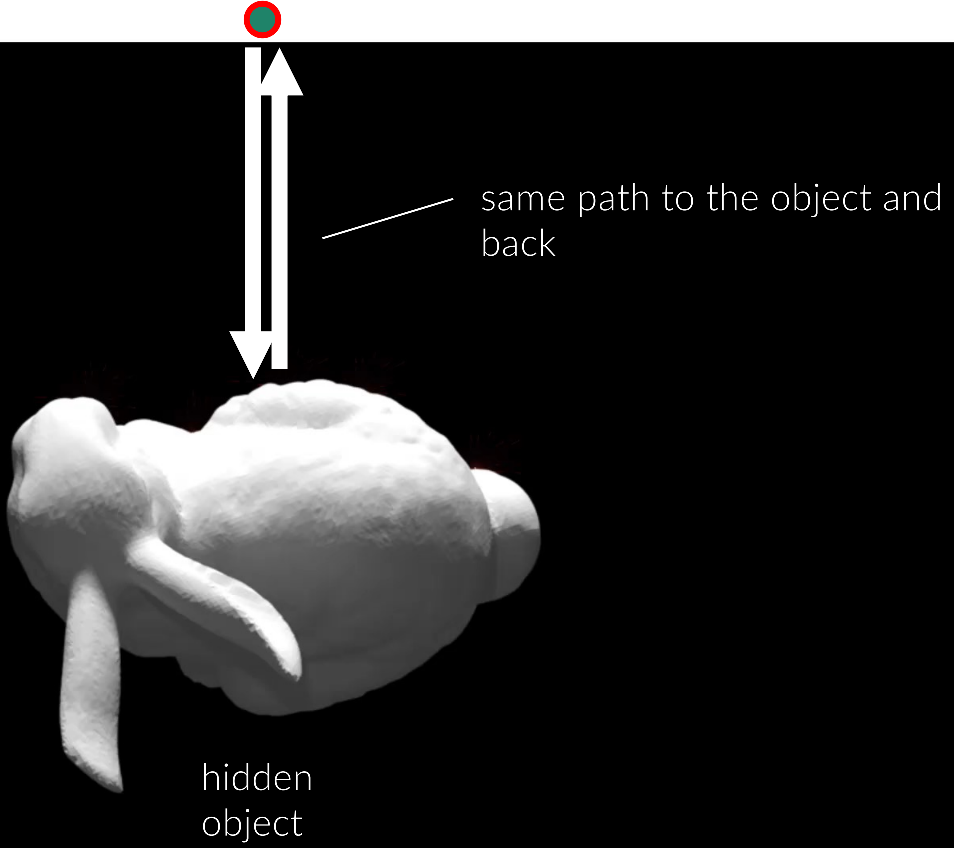

confocal sampling

wa

- simplified NLOS mathematical model
- enables efficient NLOS reconstruction

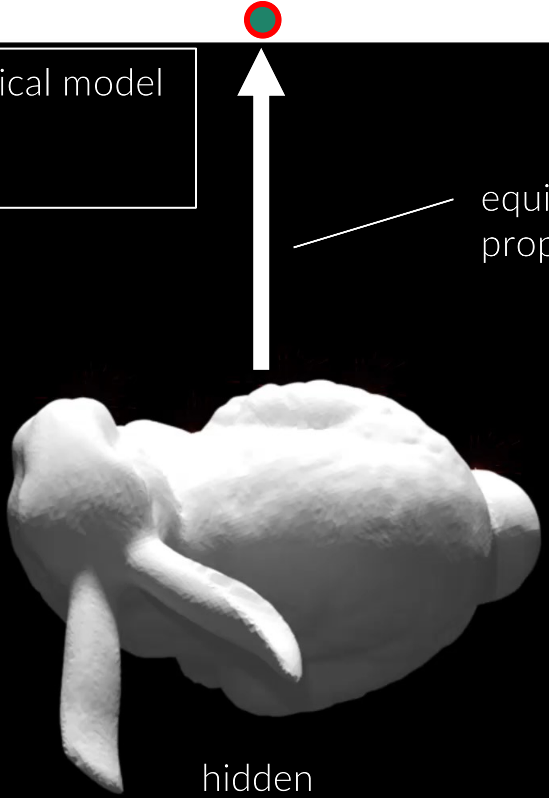

equivalent to one-way propagation at half-speed

hidden object

confocal sampling

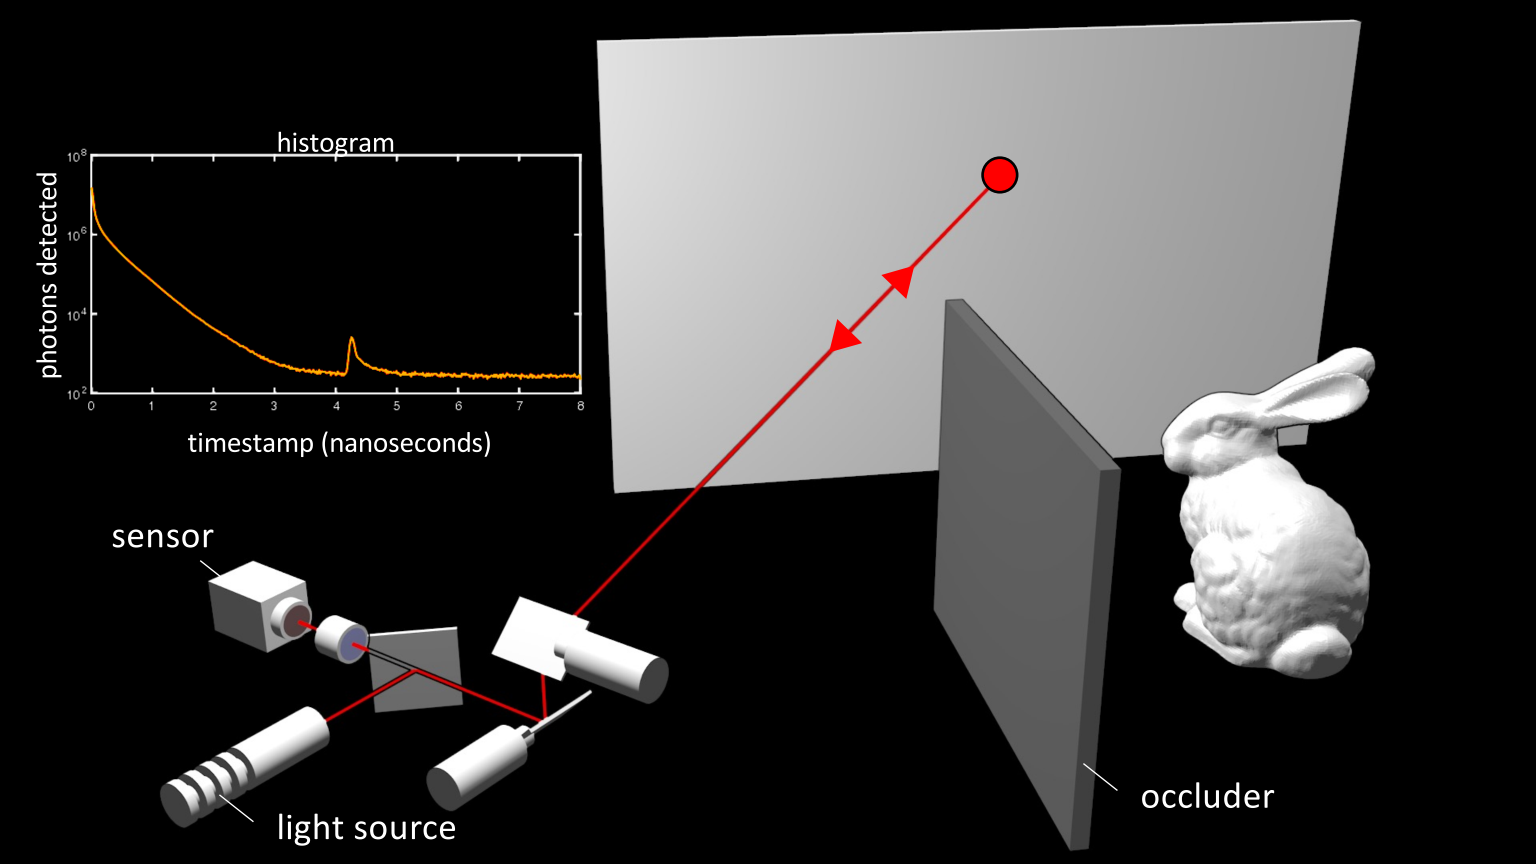
histogram
photons detected
timestamp (nanoseconds)
sensor
light source
occluder

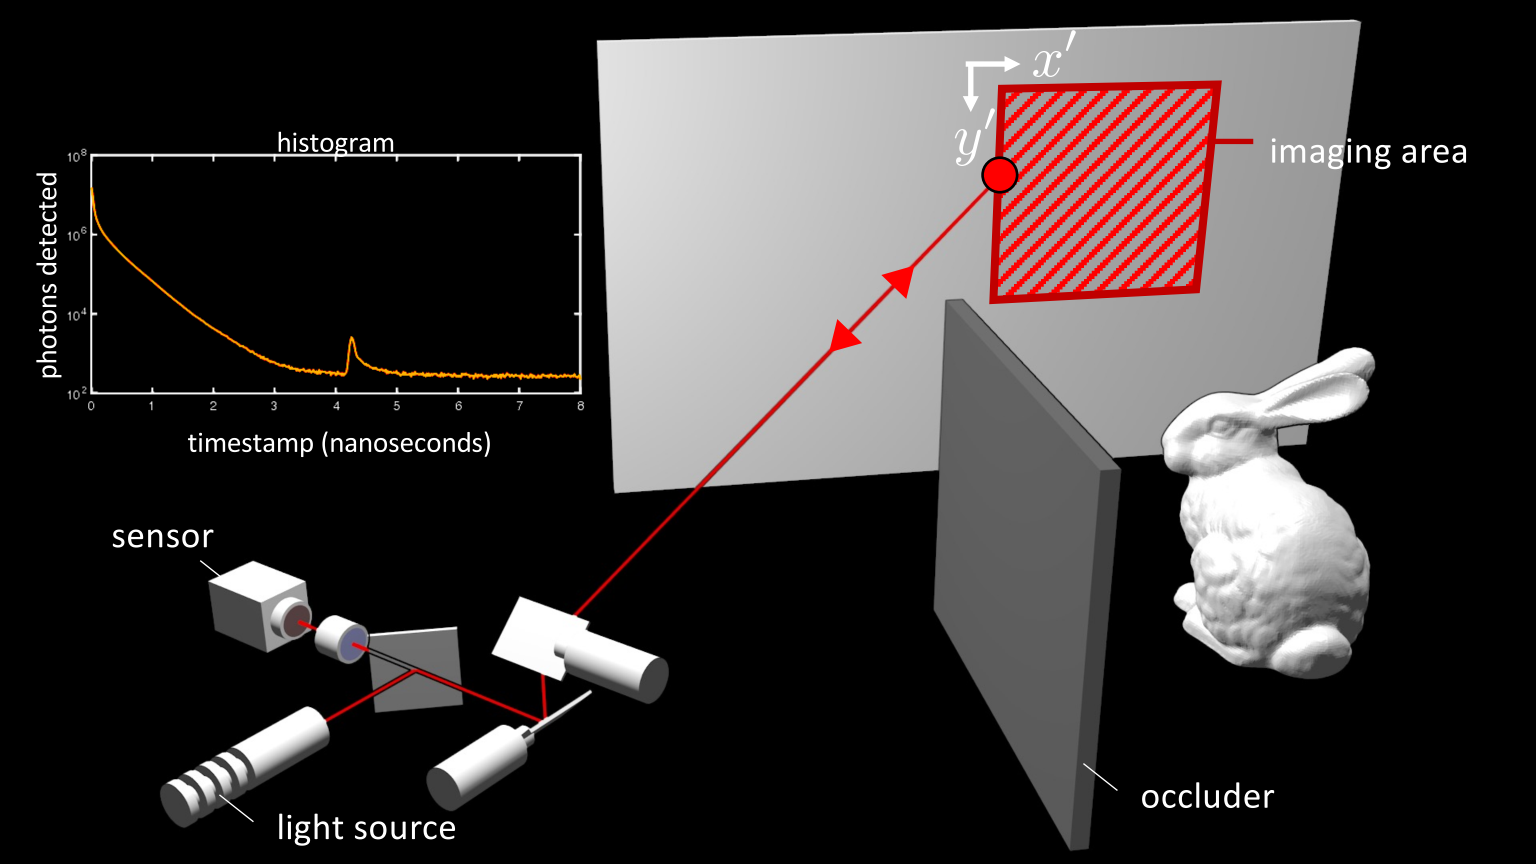

histogram
photons detected
timestamp (nanoseconds)
imaging area
sensor
light source
occluder

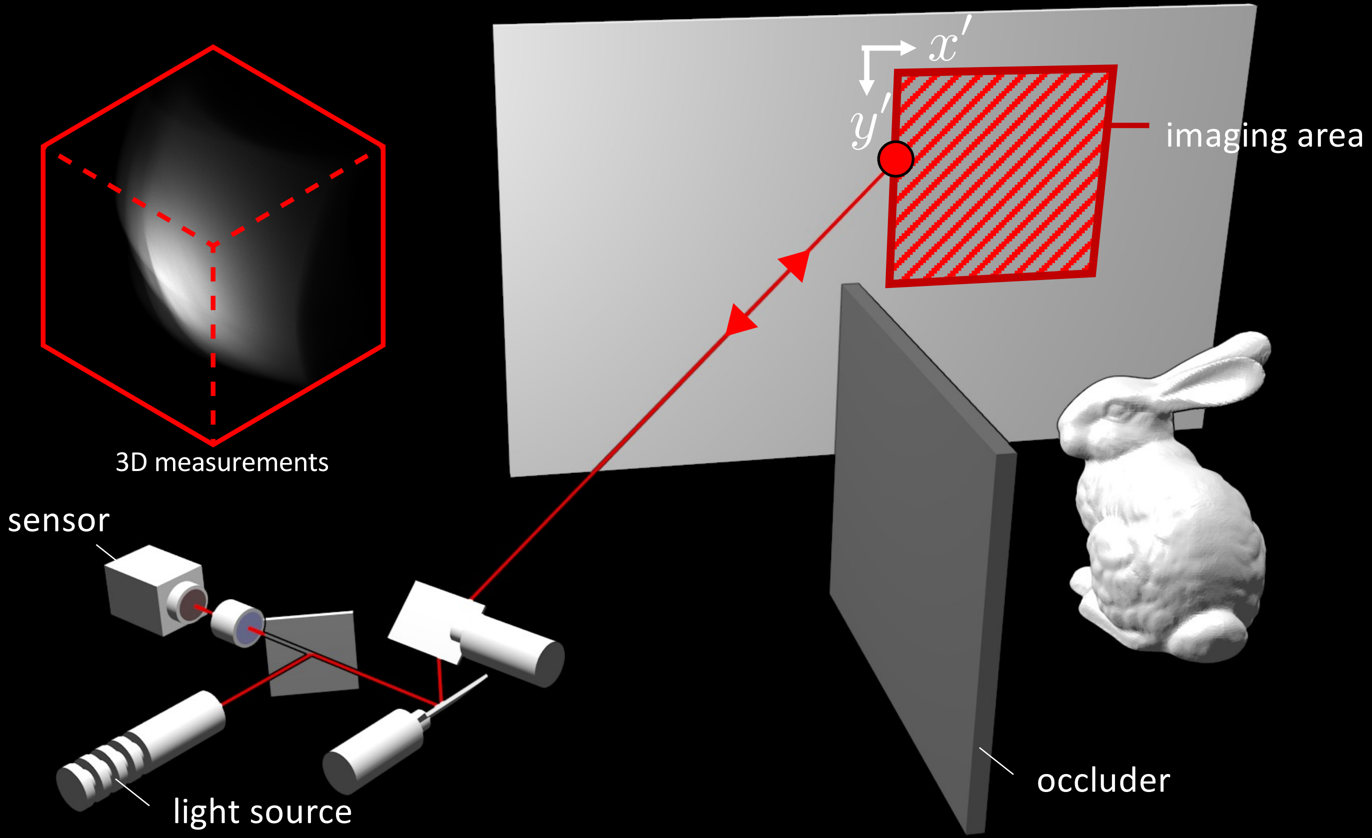

3D measurements
x'
y'
imaging area
sensor
occluder
light source

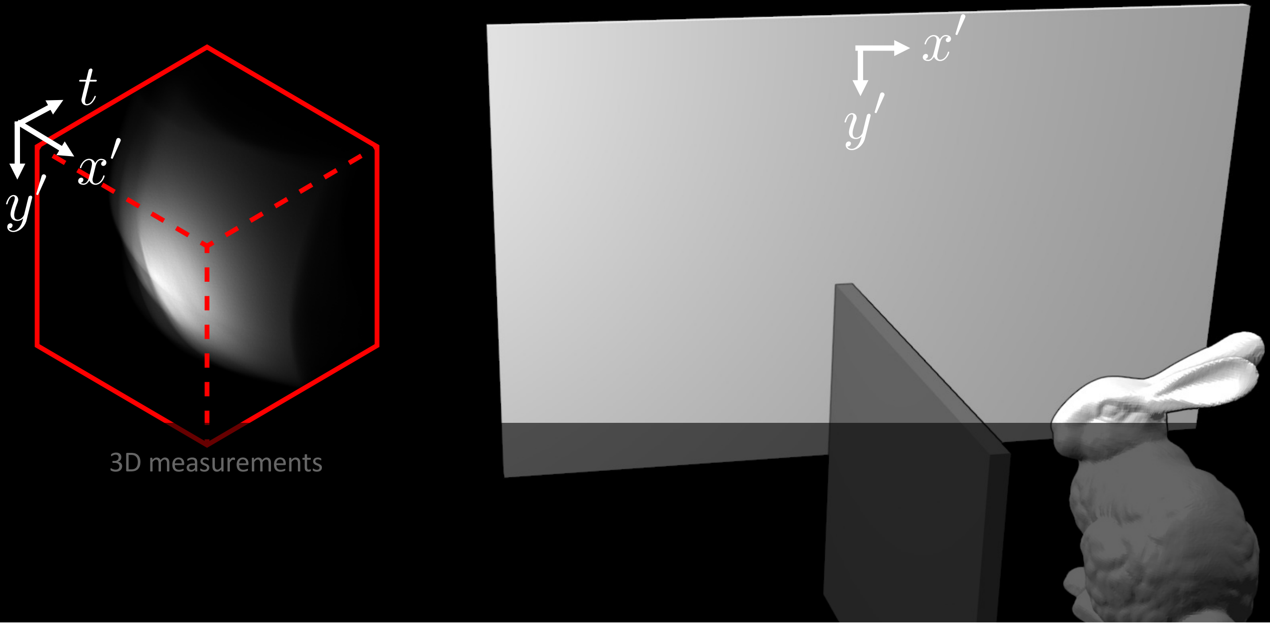

$$\underbrace{\tau(x', y', t)}_{\text{3D measurements}} = \iiint_{\Omega} \underbrace{\frac{1}{r_l^2 r^2}}_{\text{radiometric term}} \underbrace{\delta\left(r_l + r - tc\right)}_{\text{geometric term}} \cdot \underbrace{\rho(x, y, z)}_{\text{hidden 3D volume}} \; dx \; dy \; dz$$

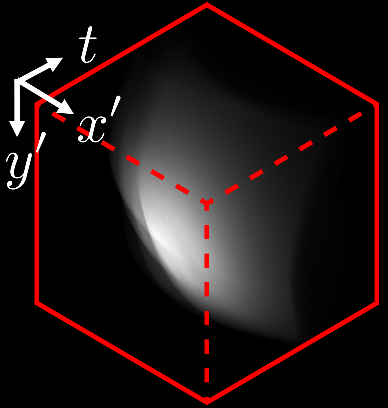

3D measurements

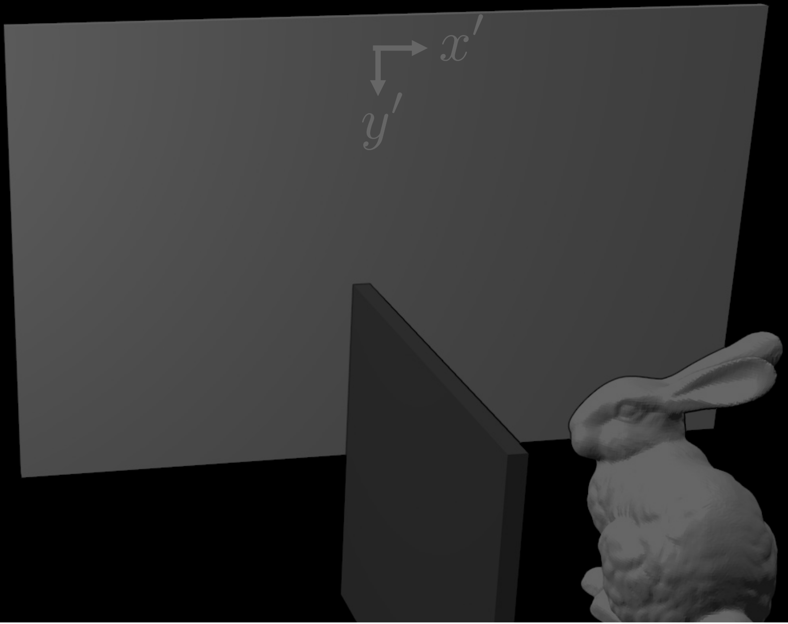

$$\underbrace{\tau(x',y',t)}_{\text{3D measurements}} = \iiint_\Omega \underbrace{\frac{1}{r_l^2 r^2}}_{\text{radiometric term}} \underbrace{\delta\left(r_l + r - tc\right)}_{\text{geometric term}} \cdot \underbrace{\rho(x,y,z)}_{\text{hidden 3D volume}} \; dx \; dy \; dz$$

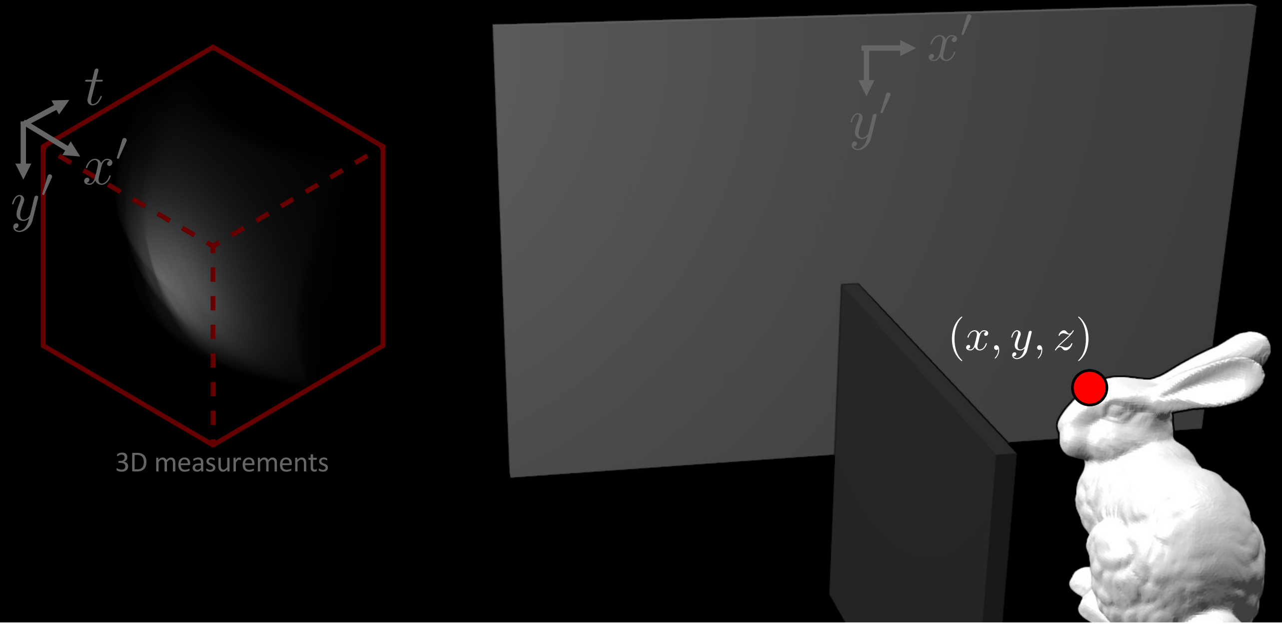

$$\underbrace{\tau(x',y',t)}_{\text{3D measurements}} = \iiint_\Omega \underbrace{\frac{1}{r_l^2 r^2}}_{\text{radiometric term}} \underbrace{\delta\left(r_l + r - tc\right)}_{\text{geometric term}} \cdot \underbrace{\rho(x,y,z)}_{\text{hidden 3D volume}}\; dx\; dy\; dz$$

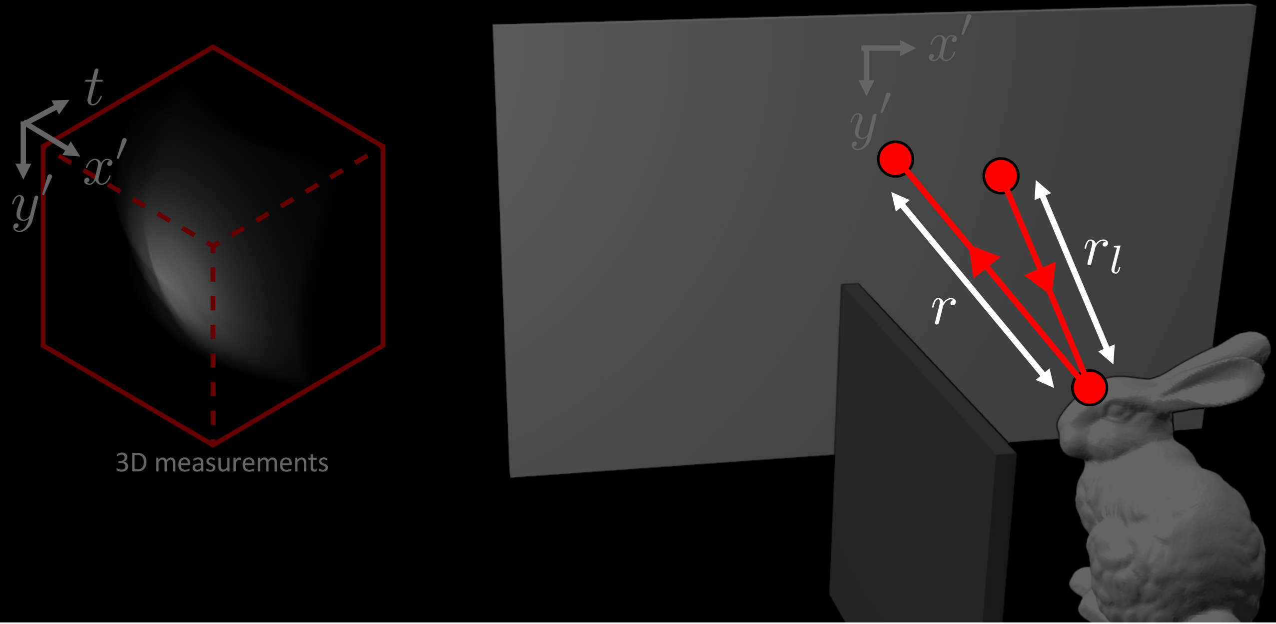

$$\underbrace{\tau(x', y', t)}_{\text{3D measurements}} = \iiint_{\Omega} \underbrace{\frac{1}{r_l^2 r^2}}_{\text{radiometric term}} \underbrace{\delta\left(r_l + r - tc\right)}_{\text{geometric term}} \cdot \underbrace{\rho(x, y, z)}_{\text{hidden 3D volume}} \, dx \; dy \; dz$$

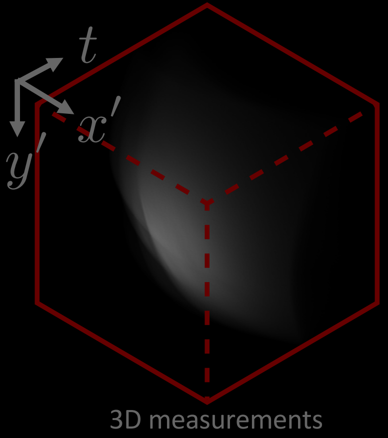

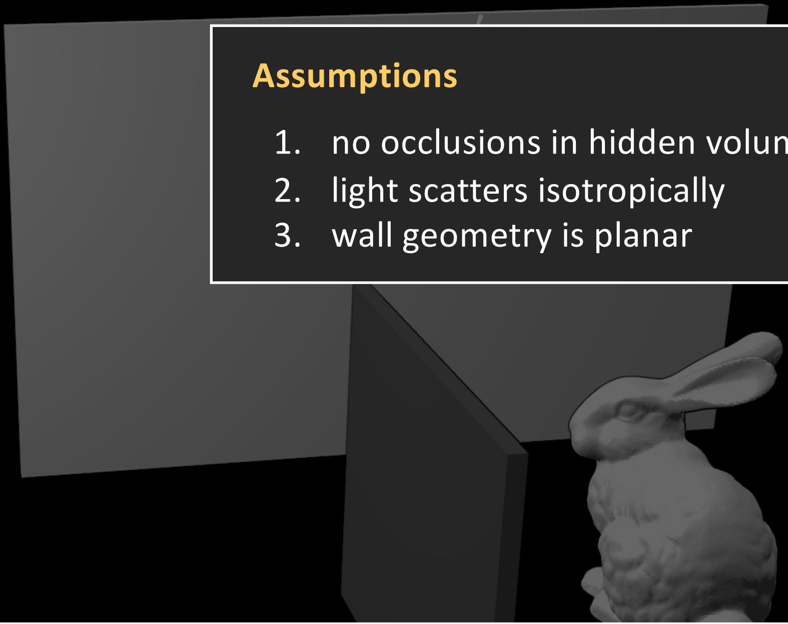

Assumptions

1. no occlusions in hidden volume
2. light scatters isotropically
3. wall geometry is planar

$$\underbrace{\tau(x', y', t)}_{\text{3D measurements}} = \iiint_\Omega \underbrace{\frac{1}{r_l^2 r^2}}_{\text{radiometric term}} \underbrace{\delta\left(r_l + r - tc\right)}_{\text{geometric term}} \cdot \underbrace{\rho(x, y, z)}_{\text{hidden 3D volume}} \, dx \; dy \; dz$$

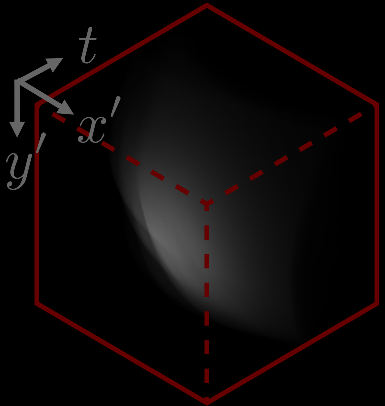

3D measurements

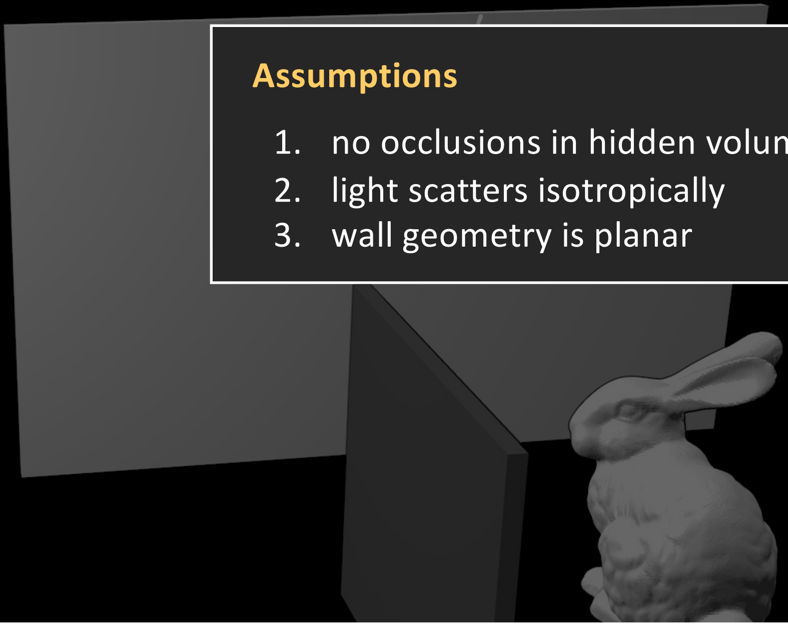

Assumptions

1. no occlusions in hidden volume
2. light scatters isotropically
3. wall geometry is planar

$$\underbrace{\tau(x',y',t)}_{\tau} \underset{=}{=} \iiint_\Omega \underbrace{\frac{1}{r_l^2 r^2}\delta\left(r_l + r - tc\right)}_{\mathbf{A}} \underset{\times}{\cdot} \underbrace{\rho(x,y,z)}_{\rho}\ dx\ dy\ dz$$

NLOS image formation mode:

$$\tau = \mathbf{A}\rho$$

measurements $n^3 \times 1$

transport matrix $n^3 \times n^3$

unknown volume $n^3 \times 1$

PROBLEM: $\mathbf{A}$ extremely large in practice (e.g., for n = 100, $\mathbf{A}$ has 1 trillion elements)

Backpropagation [Velten 12, Buttafava 15]

Flops: $O(n^5)$

Memory: $O(n^3)$

Runtime: Approx. 10 min.

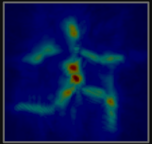

Iterative Inversion [Gupta 12, Heide 13]

Flops: $O(n^5)$ per iter.

Memory: $O(n^5)$

Runtime: > 1 hour

Our approach

express image formation model as a 3D convolution, by:

1. confocalizing measurements
2. performing a change of variables (set $z = \sqrt{u}$, $t = 2\sqrt{v}/c$)

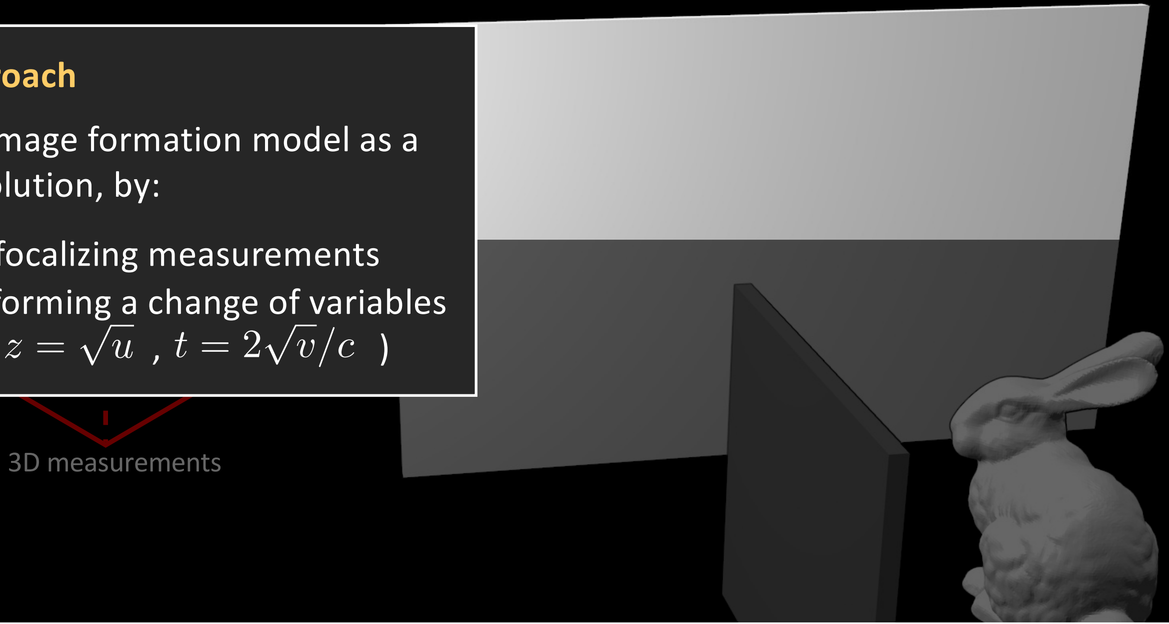

$$\tau(x', y', t) = \iiint_{\Omega} \frac{1}{r_l^2 r^2} \delta\left(r_l + r - tc\right) \cdot \rho(x, y, z)\ dx\ dy\ dz$$

Our approach

express image formation model as a 3D convolution, by:

1. confocalizing measurements
2. performing a change of variables (set $z = \sqrt{u}$, $t = 2\sqrt{v}/c$)

3D measurements

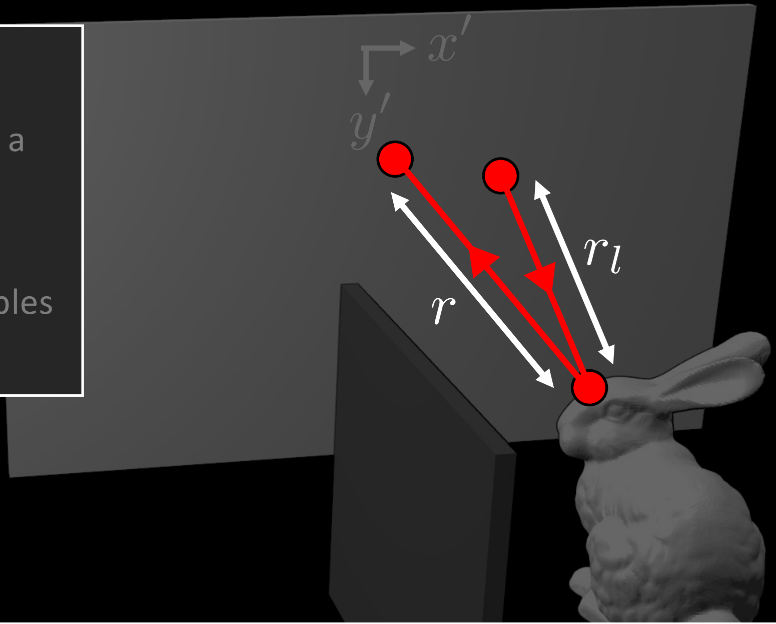

$$\tau(x', y', t) = \iiint_\Omega \frac{1}{r_l^2 r^2} \delta\left(r_l + r - tc\right) \cdot \rho(x, y, z)\ dx\ dy\ dz$$

Our approach

express image formation model as a 3D convolution, by:

1. confocalizing measurements
2. performing a change of variables (set $z = \sqrt{u}$, $t = 2\sqrt{v}/c$)

3D measurements

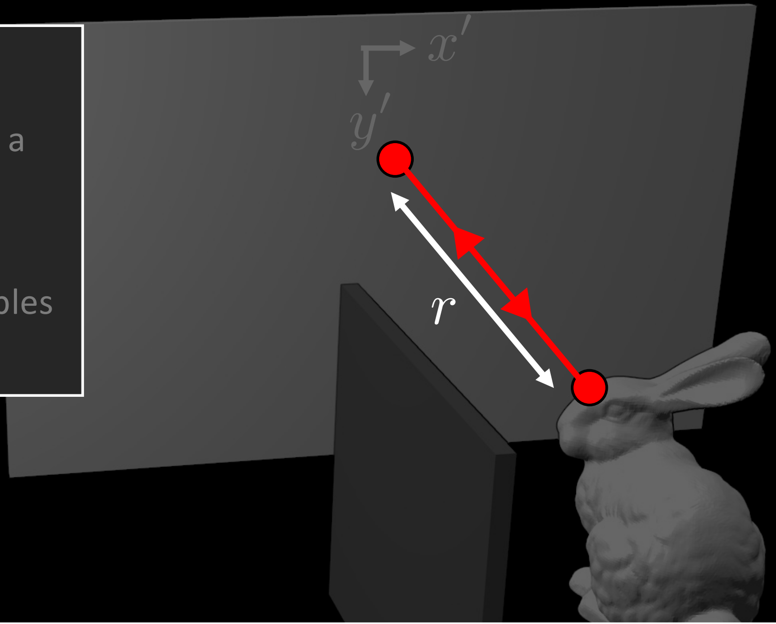

$$\tau(x', y', t) = \iiint_\Omega \frac{1}{r^4} \delta\left(2r - tc\right) \cdot \rho(x, y, z)\ dx\ dy\ dz$$

Our approach

express image formation model as a 3D convolution, by:

1. confocalizing measurements
2. performing a change of variables (set $z = \sqrt{u}$, $t = 2\sqrt{v}/c$)

3D measurements

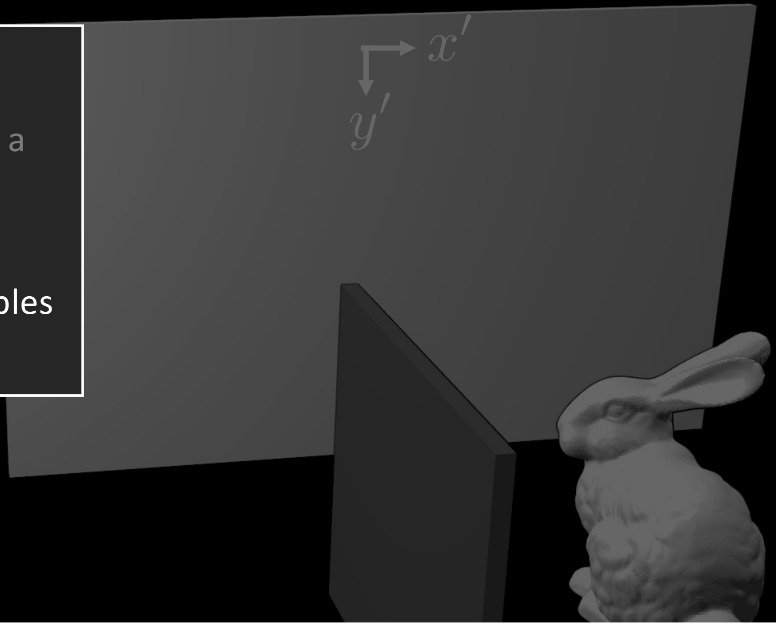

$$\tau(x', y', t) = \iiint_{\Omega} \frac{1}{r^4} \delta\left(2r - tc\right) \cdot \rho(x, y, z)\; dx\; dy\; dz$$

Our approach

express image formation model as a 3D convolution, by:

1. confocalizing measurements
2. performing a change of variables (set $z = \sqrt{u}$, $t = 2\sqrt{v}/c$)

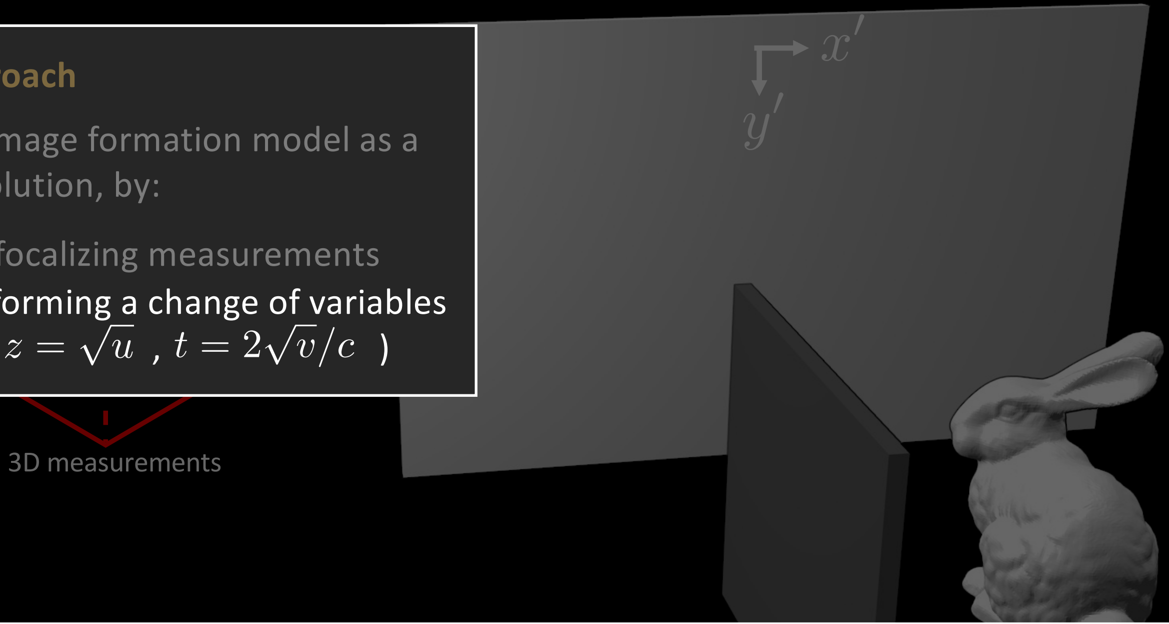

$$v^{3/2}\tau(x', y', \frac{2}{c}\sqrt{v}) = \iiint_\Omega \frac{1}{2\sqrt{u}} \delta\left((x'-x)^2 + (y'-y)^2 + u - v\right) \cdot \rho(x, y, \sqrt{u})\,dx\,dy\,du$$

Our approach

express image formation model as a 3D convolution, by:

1. confocalizing measurements
2. performing a change of variables (set $z = \sqrt{u}$, $t = 2\sqrt{v}/c$)

3D measurements

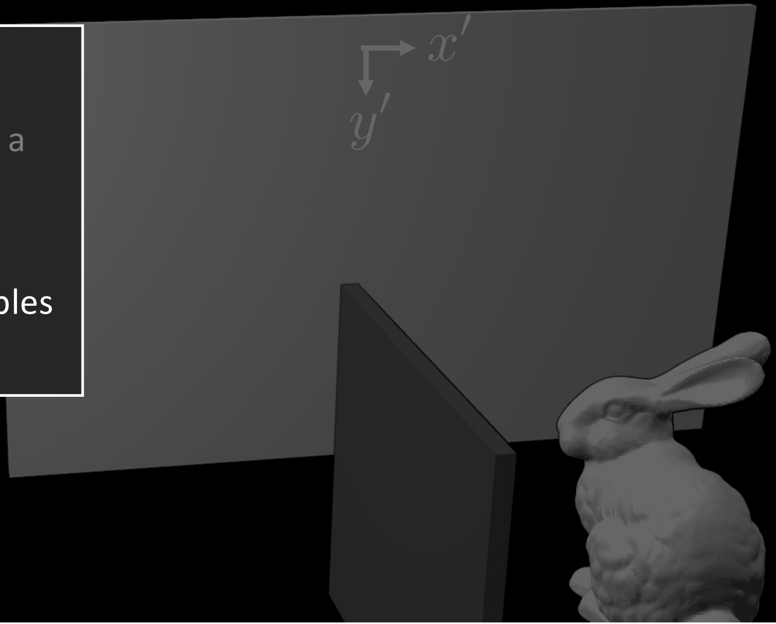

$$\underbrace{v^{3/2}\tau(x', y', \tfrac{2}{c}\sqrt{v})}_{\tau} = \underbrace{\iiint_{\Omega} \frac{1}{2\sqrt{u}} \delta\left((x'-x)^2 + (y'-y)^2 + u - v\right)}_{\mathbf{a}} \cdot \underbrace{\rho(x, y, \sqrt{u})}_{\rho} dx\,dy\,du$$

$$\tau = \mathbf{a} * \rho$$

NLOS image formation mode:

$$\tau = \mathbf{A}\rho$$

measurements $n^3 \times 1$

transport matrix $n^3 \times n^3$

unknown volume $n^3 \times 1$

Backpropagation [Velten 12, Buttafava 15]

Flops: $O(n^5)$

Memory: $O(n^3)$

Runtime: Approx. 10 min.

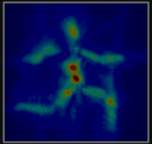

Iterative Inversion [Gupta 12, Heide 13]

Flops: $O(n^5)$ per iter.

Memory: $O(n^5)$

Runtime: > 1 hour

Confocal NLOS image formation mode:

$$\tau = \mathbf{a} * \rho$$

measurements $n \times n \times n$

blur kernel $n \times n \times n$

unknown volume $n \times n \times n$

NLOS image formation mode:

$$\tau = \mathbf{A}\rho$$

measurements $n^3 \times 1$

transport matrix $n^3 \times n^3$

unknown volume $n^3 \times 1$

Confocal NLOS image formation mode:

$$\tau = \mathbf{a} * \rho$$

measurements $n \times n \times n$

blur kernel $n \times n \times n$

unknown volume $n \times n \times n$

Backpropagation [Velten 12, Buttafava 15]

Flops: $O(n^5)$

Memory: $O(n^3)$

Runtime: Approx. 10 min.

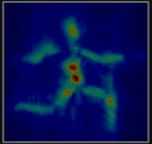

Iterative Inversion [Gupta 12, Heide 13]

Flops: $O(n^5)$ per iter.

Memory: $O(n^5)$

Runtime: > 1 hour

3D Deconvolution (with Light Cone Transform)

[O'Toole et al. 2018]

Flops: $O(n^3 \log(n))$

Memory: $O(n^3)$

Runtime: < 1 second

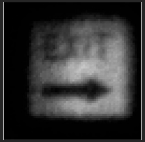

measurements

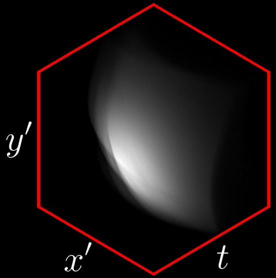

O'Toole et al., Nature 2018

measurements

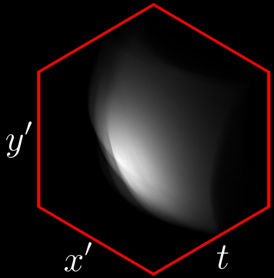

Step 1: resample and attenuate along t-axis

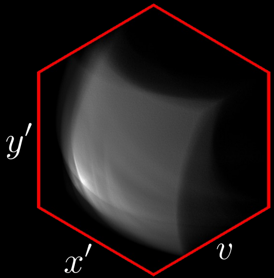

O'Toole et al., Nature 2018

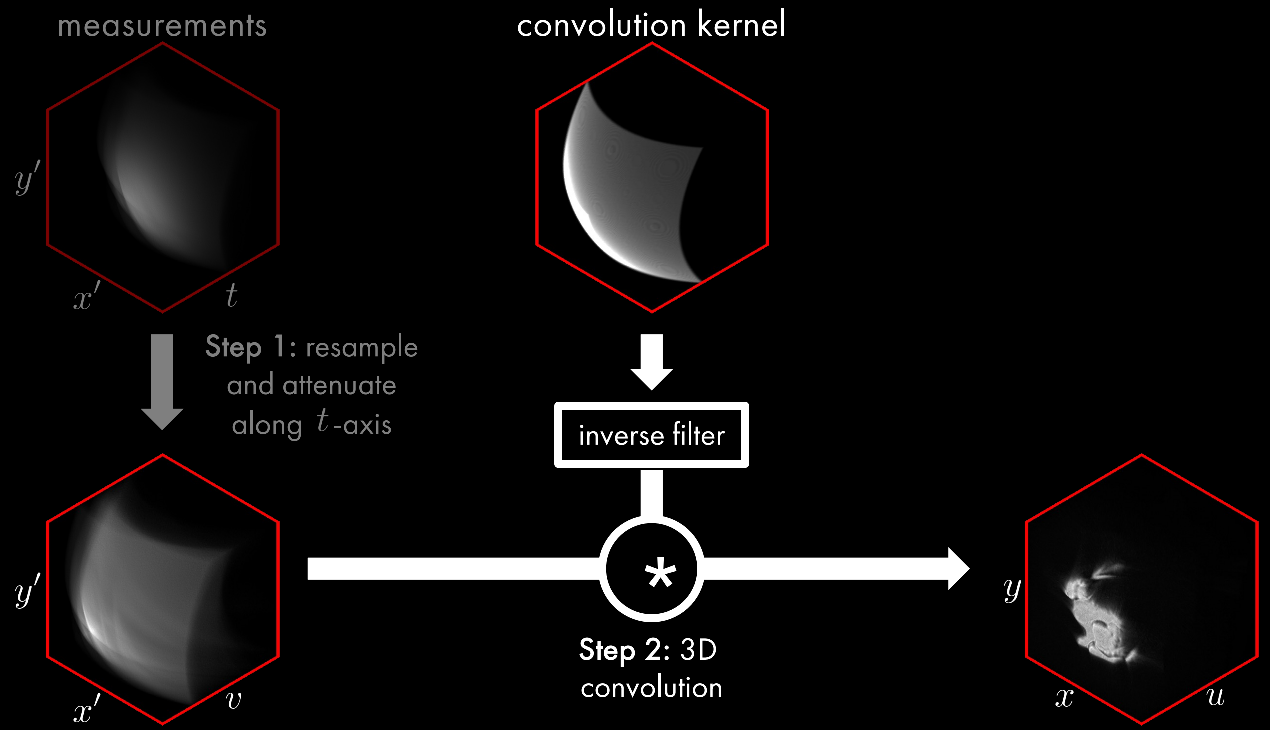

O'Toole et al., Nature 2018

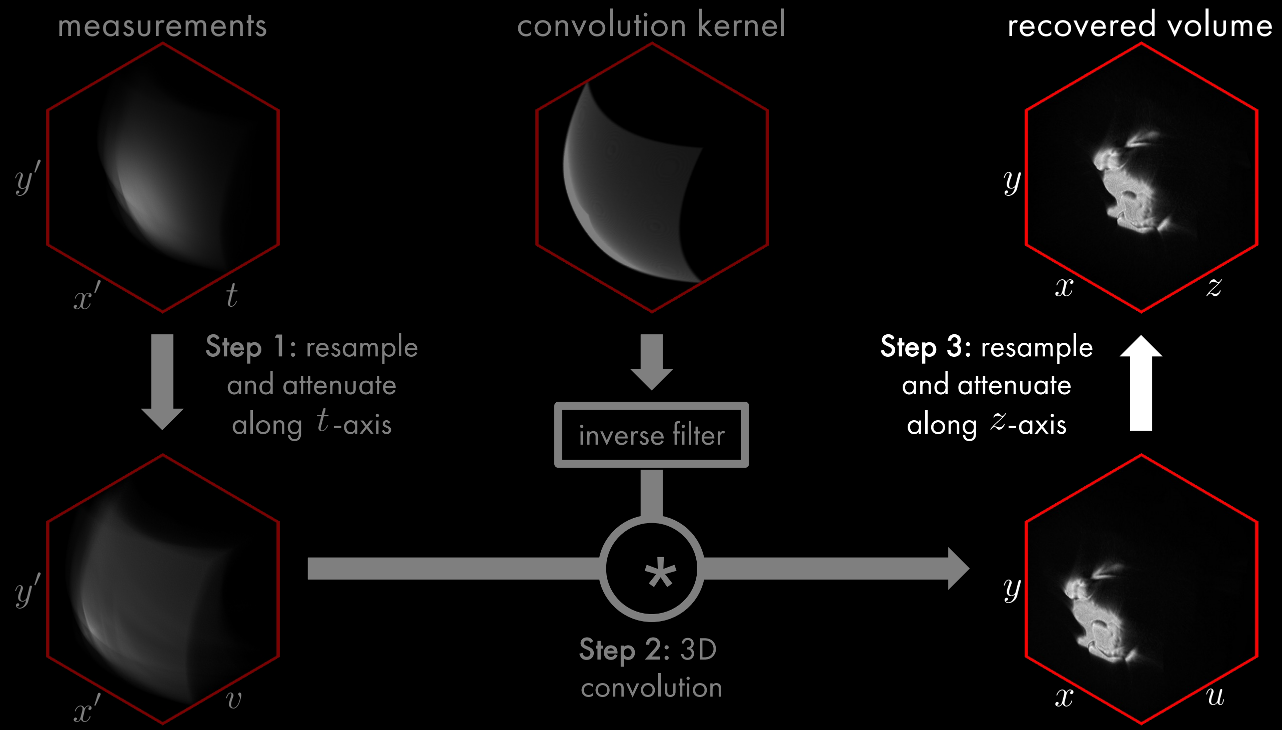

O'Toole et al., Nature 2018

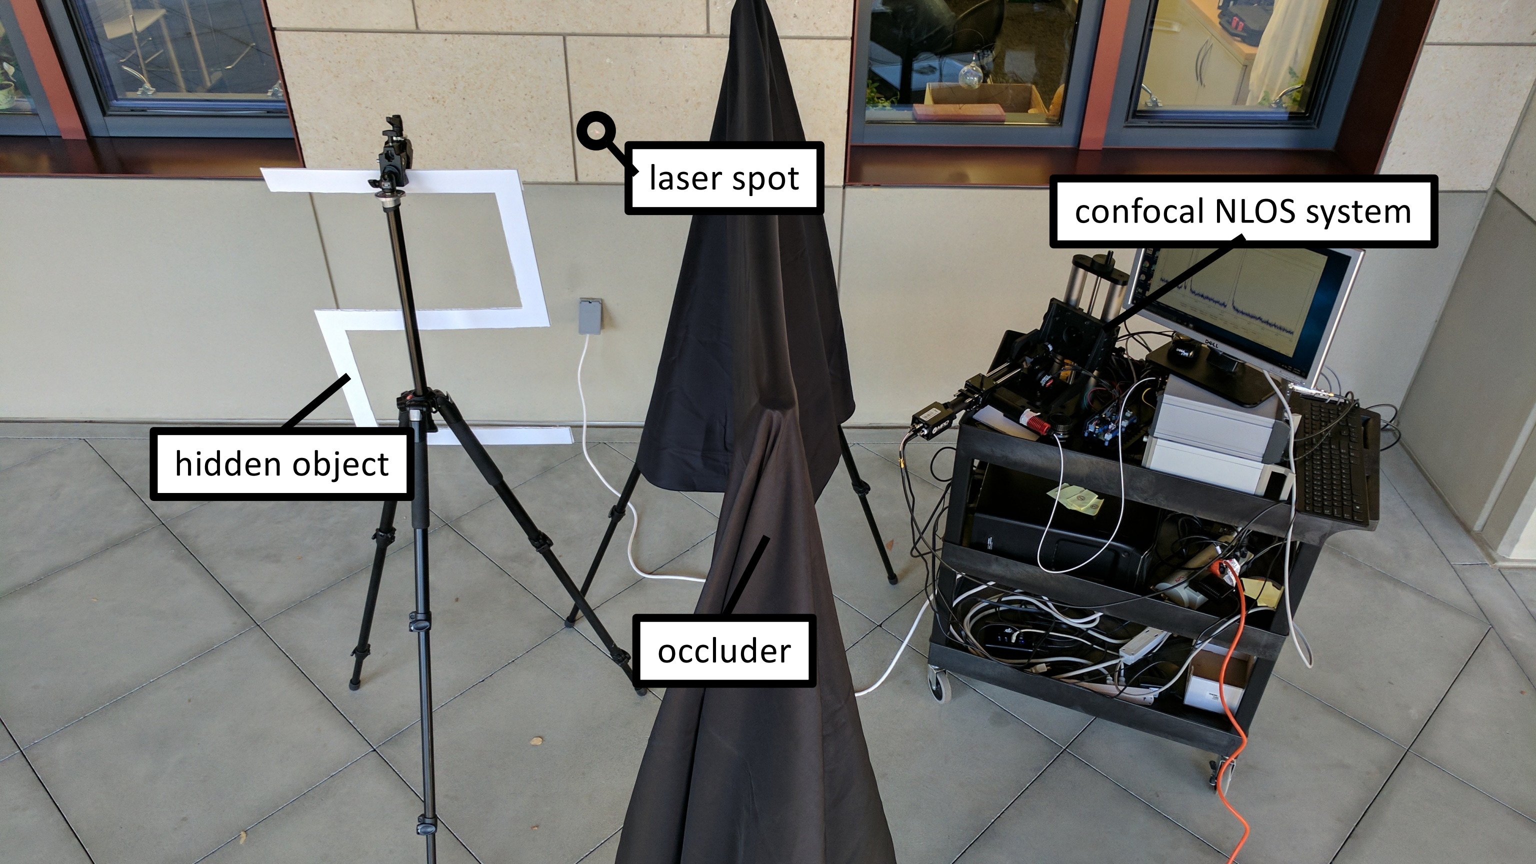
laser spot
confocal NLOS system
hidden object
occluder


1.25 meters

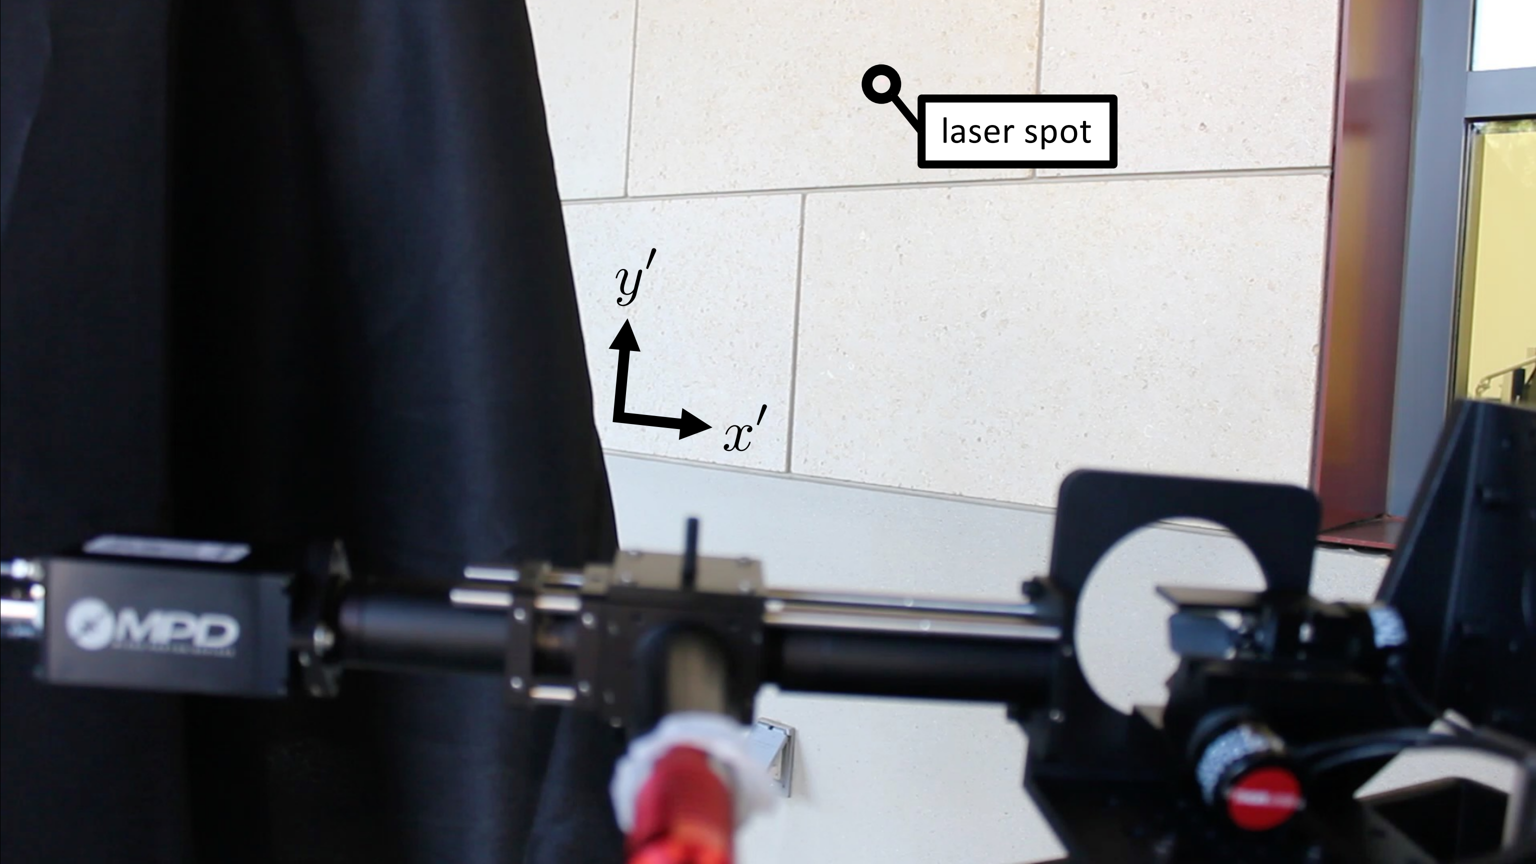

laser spot
y'
x'
MPD

measurements

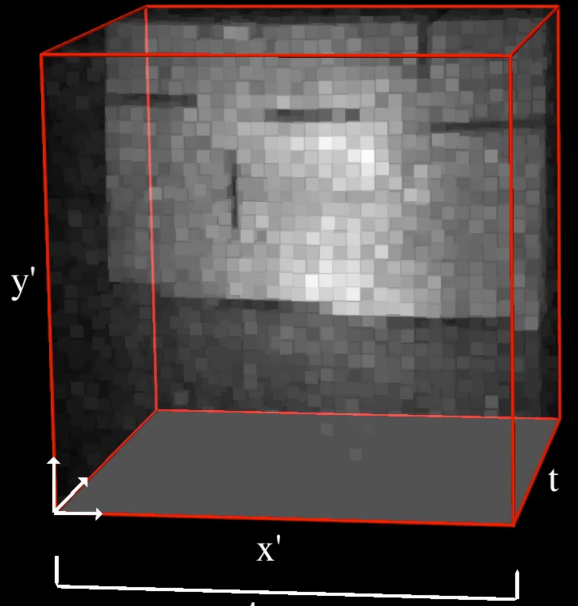

Maximum Intensity Projection

NLOS image formation model:

$$\tau = \mathbf{A}\rho$$

measurements $n^3 \times 1$ | transport matrix $n^3 \times n^3$ | unknown volume $n^3 \times 1$

Confocal scanning and Light-Cone Transform

$$\tilde{\tau} = a * \tilde{\rho}$$

measurements $n \times n \times n$ | blur kernel $n \times n \times n$ | unknown volume $n \times n \times n$

Backprojection [Velten 12, Buttafava 15]

Flops: $O(n^5)$

Memory: $O(n^3)$

Runtime: Approx. 10 min.

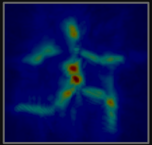

Iterative Inversion [Gupta 12, Wu 12, Heide 13]

Flops: $O(n^5)$ per iter.

Memory: $O(n^5)$

Runtime: > 1 hour

3D Deconvolution (with Light-Cone Transform) [O'Toole et al. 2018]

Flops: $O(n^3 \log(n))$

Memory: $O(n^3)$

Runtime: < 1 second

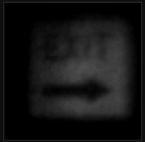

Assumption:

- Isotropic scattering (only diffuse or retroreflective objects)

overview

- Time-resolved imaging
- Single-photon avalanche diodes (SPADs)
- Single-photon lidar
- Non-line-of-sight imaging (part 2)
- Imaging through scattering media

$$\nabla^2\Psi - \frac{1}{v^2}\frac{\partial^2\Psi}{\partial t^2} = 0$$

wa

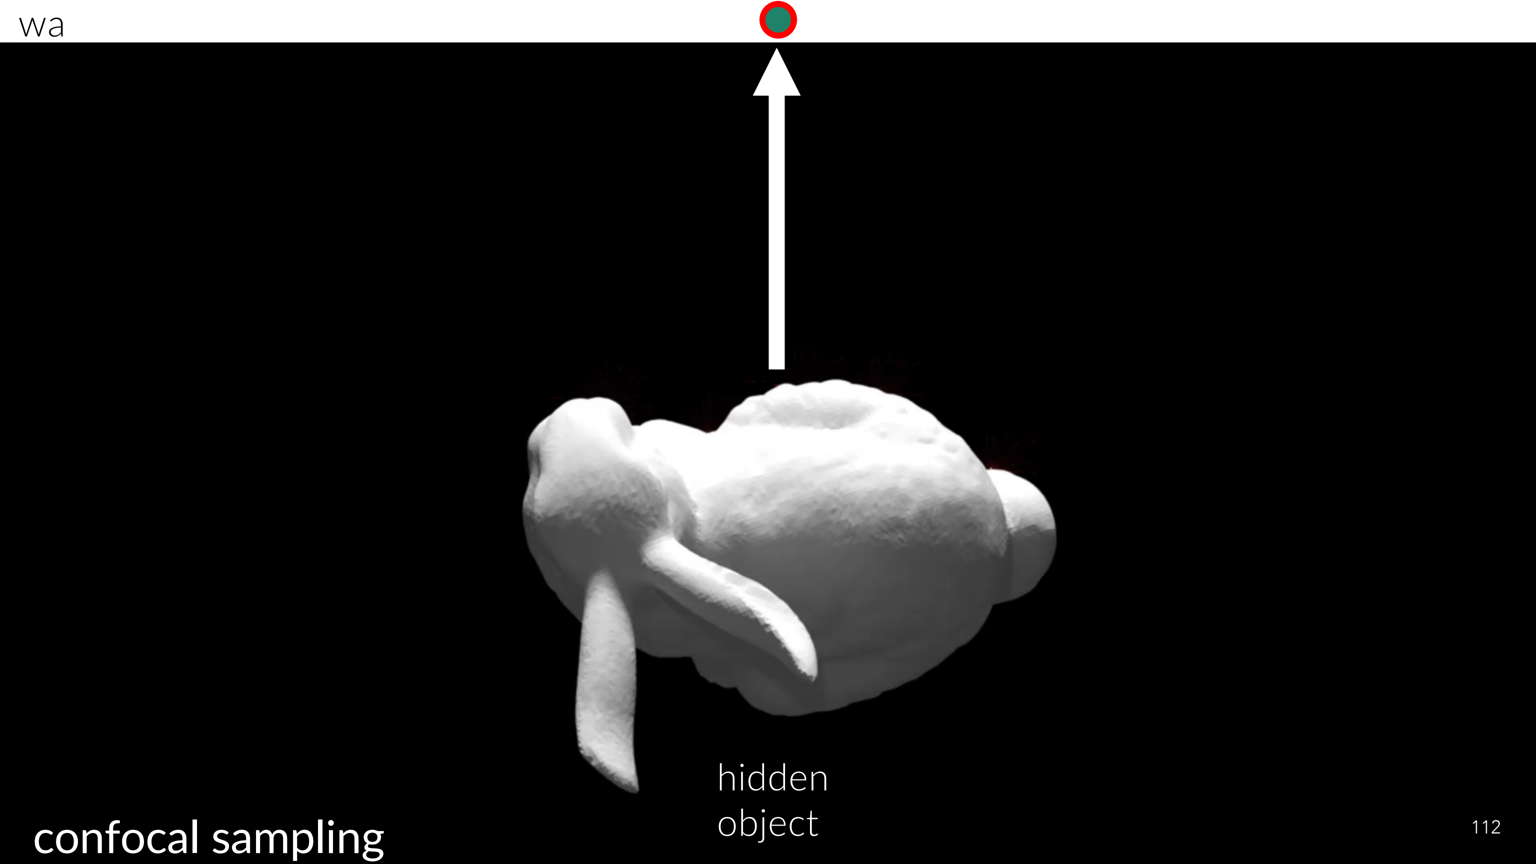

confocal sampling

image formation model

$$\Psi(x, z, t)$$

wavefield

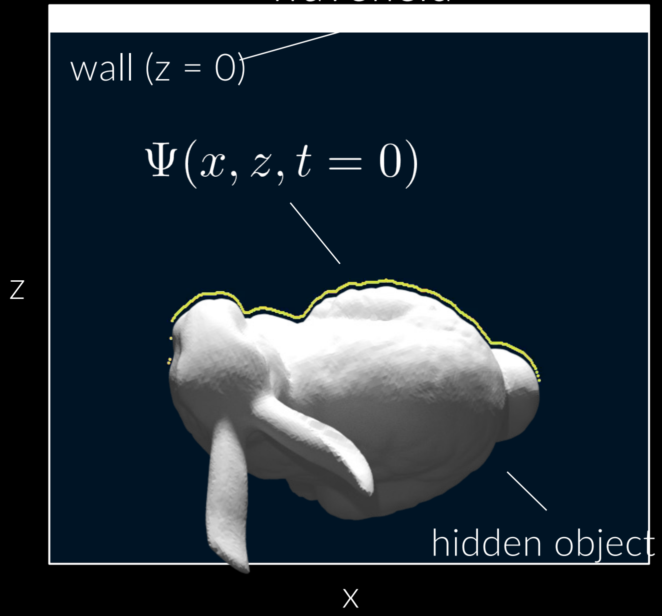

$$\Psi(x, z = 0, t)$$

confocal measurements

general solution (time reversal)

$$\Psi(x, z, t = \textbf{03.000 ns})$$

wavefield

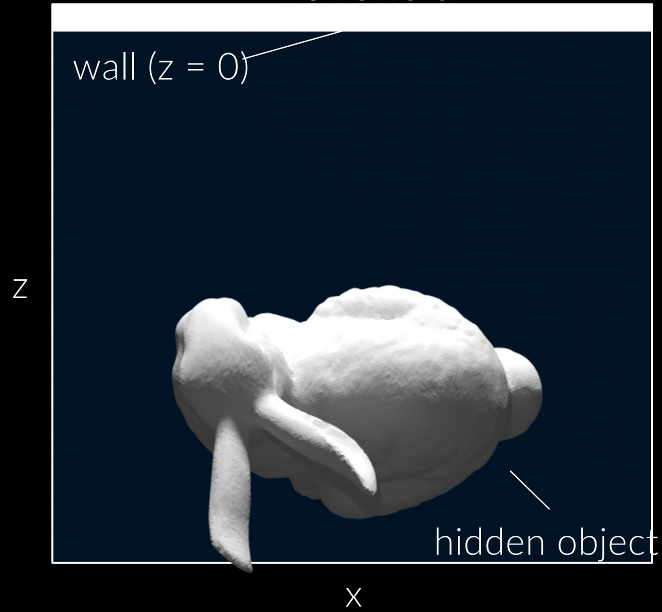

$$\Psi(x, z = 0, t)$$

confocal measurements

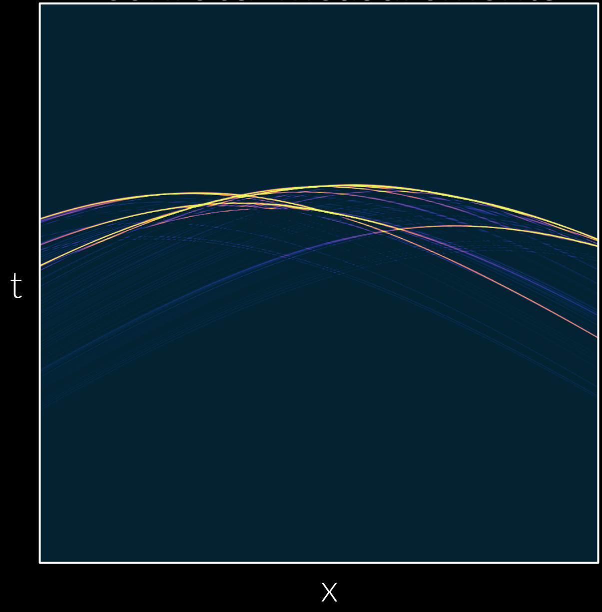

general solution (time reversal)

finite-difference time-domain method

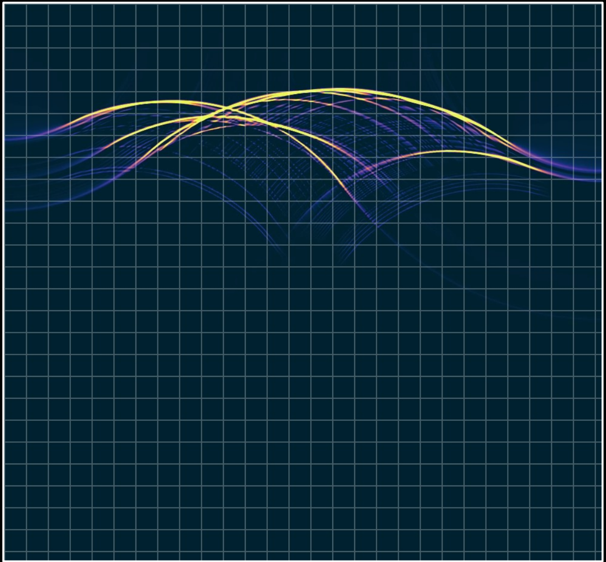

1. approximate wave equation with finite differences

$$\frac{\partial^2 \Psi}{\partial t^2} \approx \frac{\Psi_i^{n+1} - 2\Psi_i^n + \Psi_i^{n-1}}{(\Delta t)^2}$$

2. solve for previous timestep

$$\Psi_i^{n-1} = f\left(\Psi^n, \Psi^{n+1}\right)$$

3. repeatedly update Ψ at all grid cells

general solution (time reversal)

finite-difference time-domain method

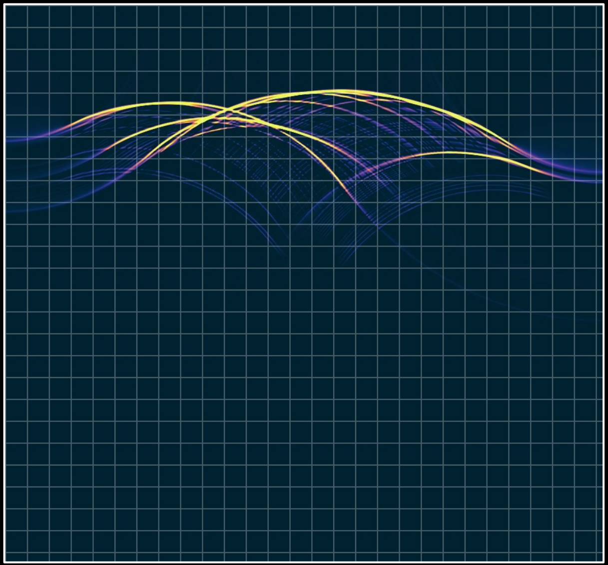

1. approximate wave equation with finite differences

$$\frac{\partial^2 \Psi}{\partial t^2} \approx \frac{\Psi_i^{n+1} - \cdots{}^{n-1}}{\cdots}$$

2. solve … us timestep

$$\Psi_i^{n-1} = f\left(\Psi^n, \Psi^{n+1}\right)$$

3. repeatedly update Ψ at all grid cells

Slow to get t=0 at high-resolution!

frequency–wavenumber (*f*–*k*) Migration

$$\Psi(x, z, t = 0)$$

wavefield

$$\Psi(x, z = 0, t)$$

confocal measurements

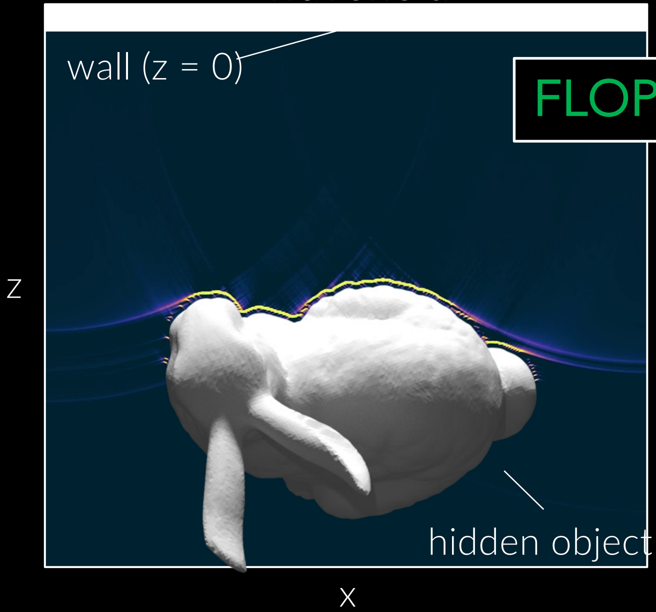

FLOPS: $O(n^3 \log n)$

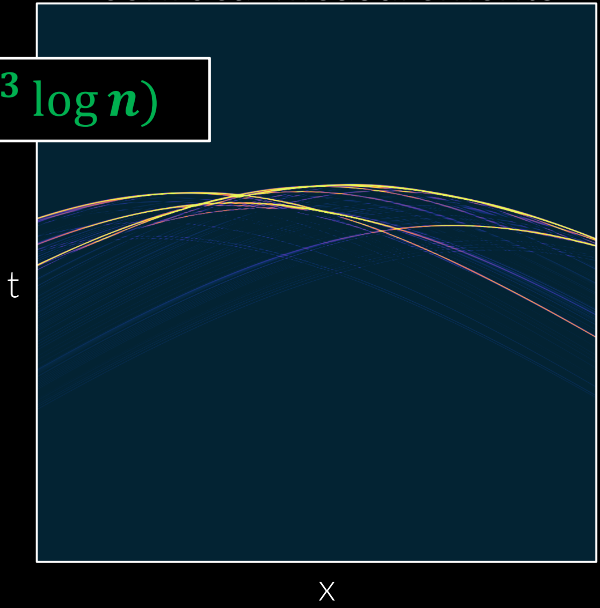

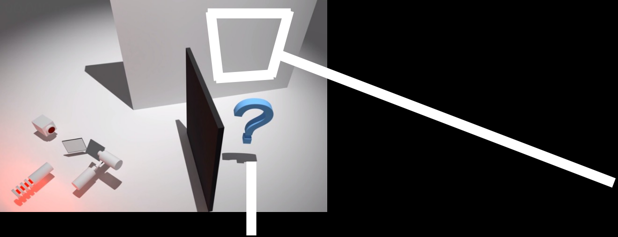

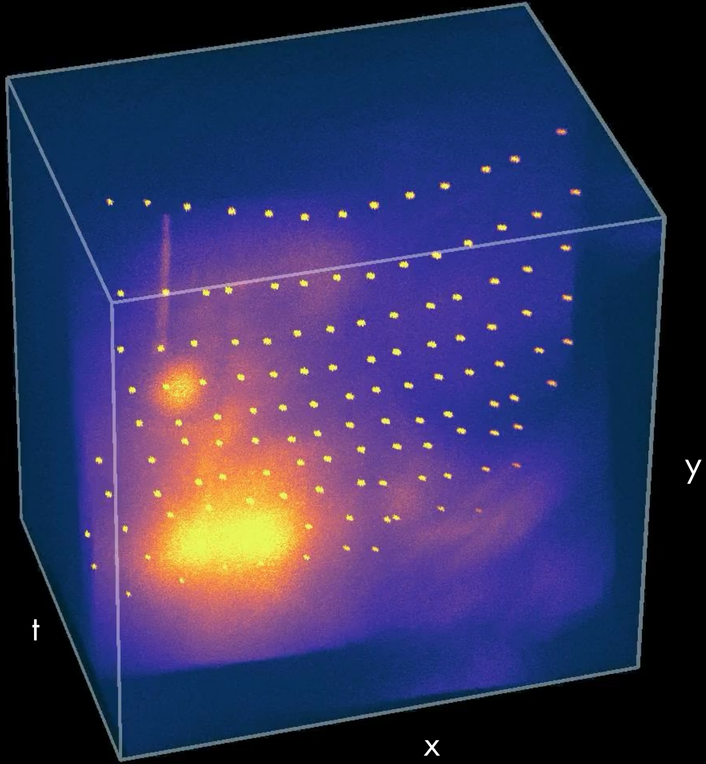

Captured Measurements

Lindell et al., SIGGRAPH 2019

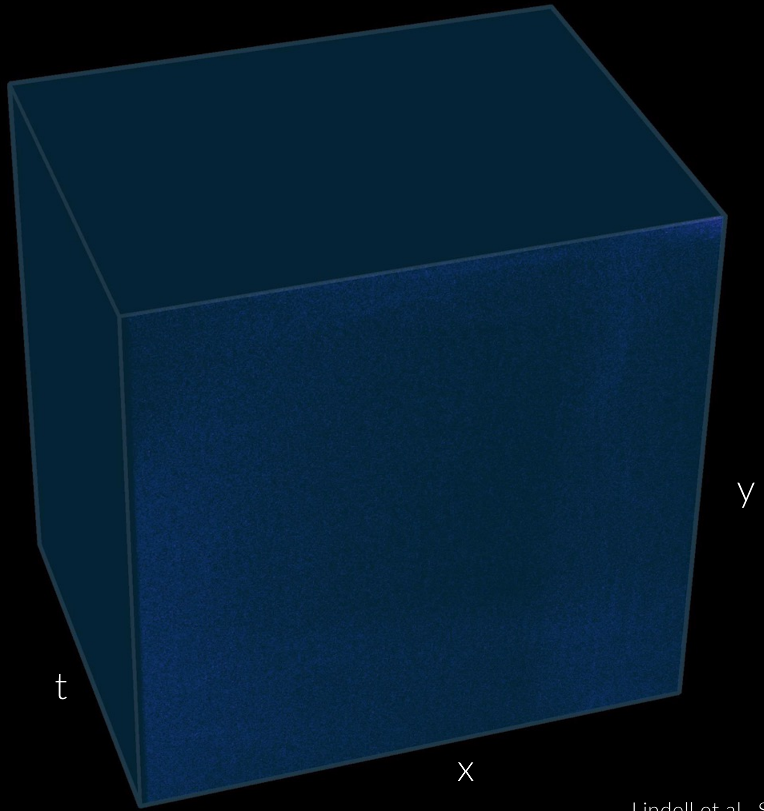

Lindell et al., SIGGRAPH 2019

f-k Migration

$$\Psi(x, y, t)$$

Measurements (z=0)

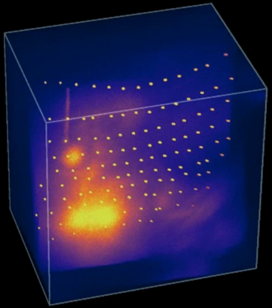

$$\bar{\Phi}(k_x, k_y, f)$$

Spectrum

$$\Phi(k_x, k_y, k_z)$$

Interpolated Spectrum

$$\Psi(x, y, z)$$

Hidden Volume (t=0)

$\rightarrow$ $\mathcal{F}$

$\rightarrow$ Resamp

$\rightarrow$ $\mathcal{F}^{-1}$

f-*k* Migration

Express wavefield as function of measurement spectrum (plane wave decomposition)

$$\Psi(x, y, z, t) = \iiint \bar{\Phi}(k_x, k_y, f) e^{2\pi i (k_x x + k_y y + k_z z - f t)} dk_x dk_y df$$

wavefield

Fourier transform of measurements

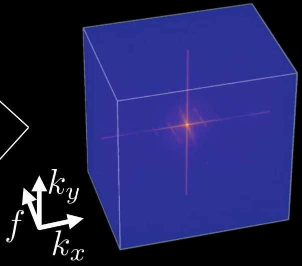

Set t=0 to get migrated solution

$$\Psi(x, y, z, t = 0) = \iiint \bar{\Phi}(k_x, k_y, f) e^{2\pi i (k_x x + k_y y + \underbrace{k_z z})} dk_x dk_y \underbrace{df}$$

Almost an inverse Fourier Transform!

f-*k* Migration

Set t=0 to get migrated solution

$$\Psi(x,y,z,t=0) = \iiint \bar{\Phi}(k_x,k_y,f)e^{2\pi i(k_x x+k_y y+\underbrace{k_z z}_{})}dk_x dk_y \underbrace{df}_{}$$

Almost an inverse Fourier Transform!

Use dispersion relation[1] to perform substitution of variables

$$f = v\sqrt{k_x^2+k_y^2+k_z^2}$$

$$\boxed{f \Rightarrow k_z}$$

[1]Georgi, Howard. *The physics of waves*. Englewood Cliffs, NJ: Prentice Hall, 1993.

Use dispersion relation[1] to perform substitution of variables

$$f = v\sqrt{k_x^2 + k_y^2 + k_z^2}$$

$$\boxed{f \Rightarrow k_z}$$

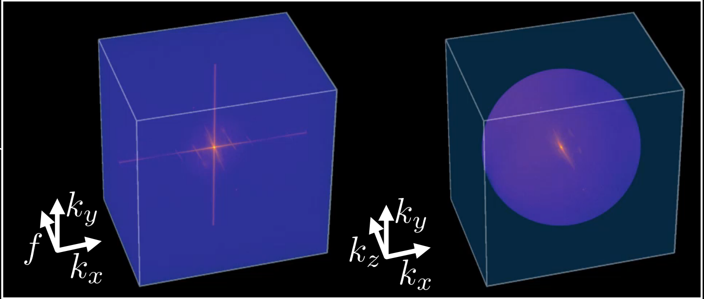

Use dispersion relation[1] to perform substitution of variables

$$f = v\sqrt{k_x^2 + k_y^2 + k_z^2}$$

$$\boxed{f \Rightarrow k_z}$$

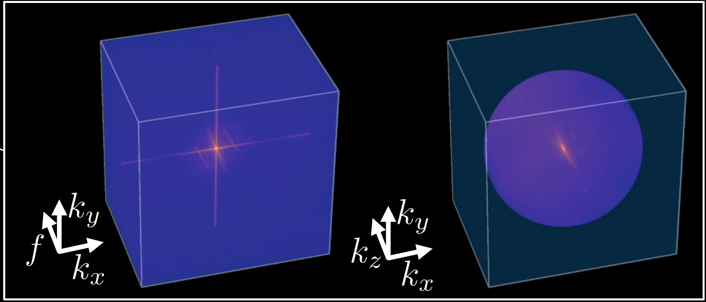

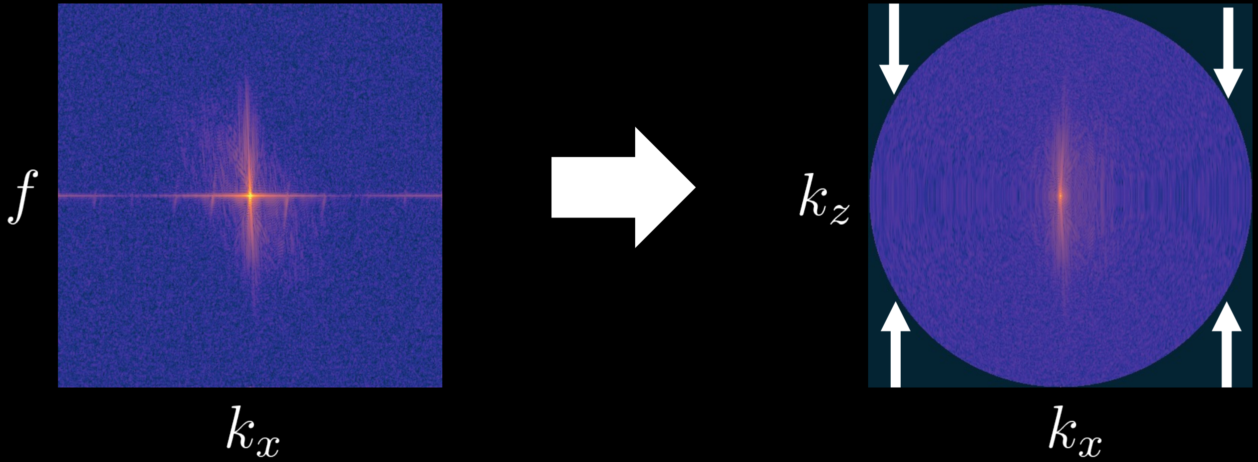

Use dispersion relation[1] to perform substitution of variables

$$f = v\sqrt{k_x^2 + k_y^2 + k_z^2}$$

$$\boxed{f \Rightarrow k_z}$$

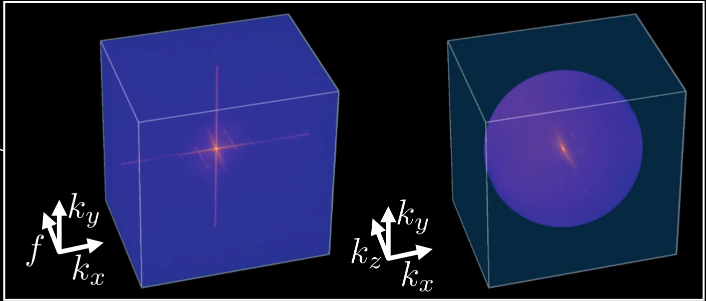

The migrated solution is an inverse Fourier Transform!

$$\boxed{\Psi(x, y, z, t = 0) = \iiint \Phi(k_x, k_y, k_z) e^{2\pi i(k_x x + k_y y + k_z z)} dk_x dk_y dk_z}$$

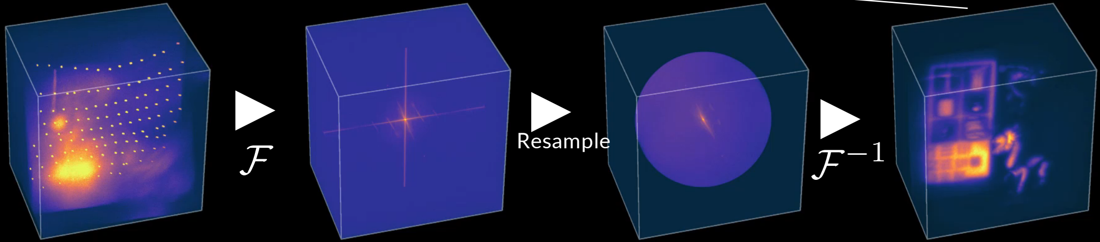

f–k Migration

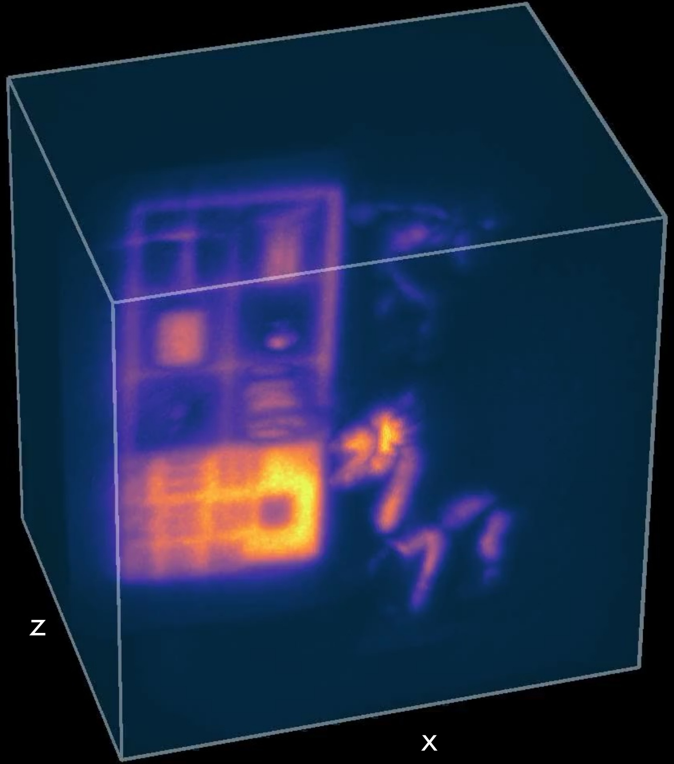

Dimensions: 2 x 2 m

Exposure: 180 min

Reconstruction time: ~90 sec (CPU)

Lindell et al., SIGGRAPH 2019

Reconstruction Comparison

dimensions: 2 m x 2 m x 1.5 m

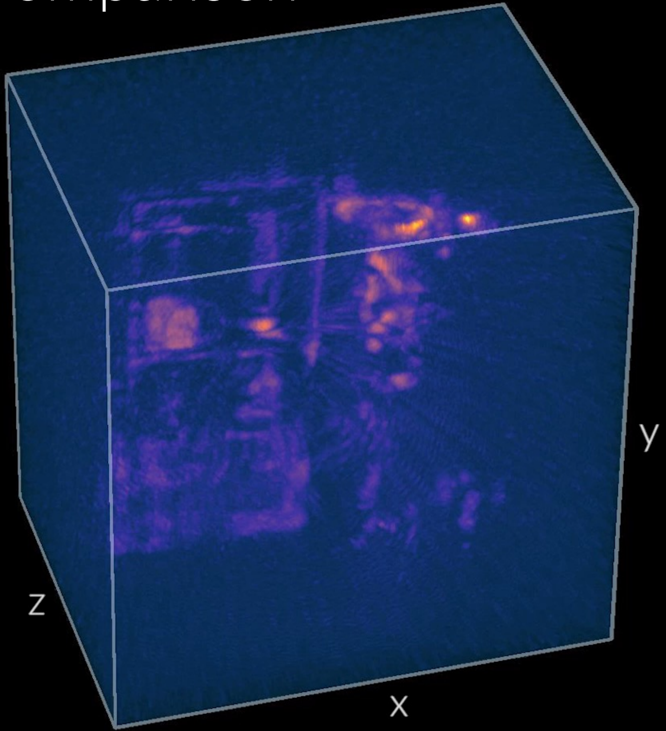

Filtered Backprojection

hardware prototype

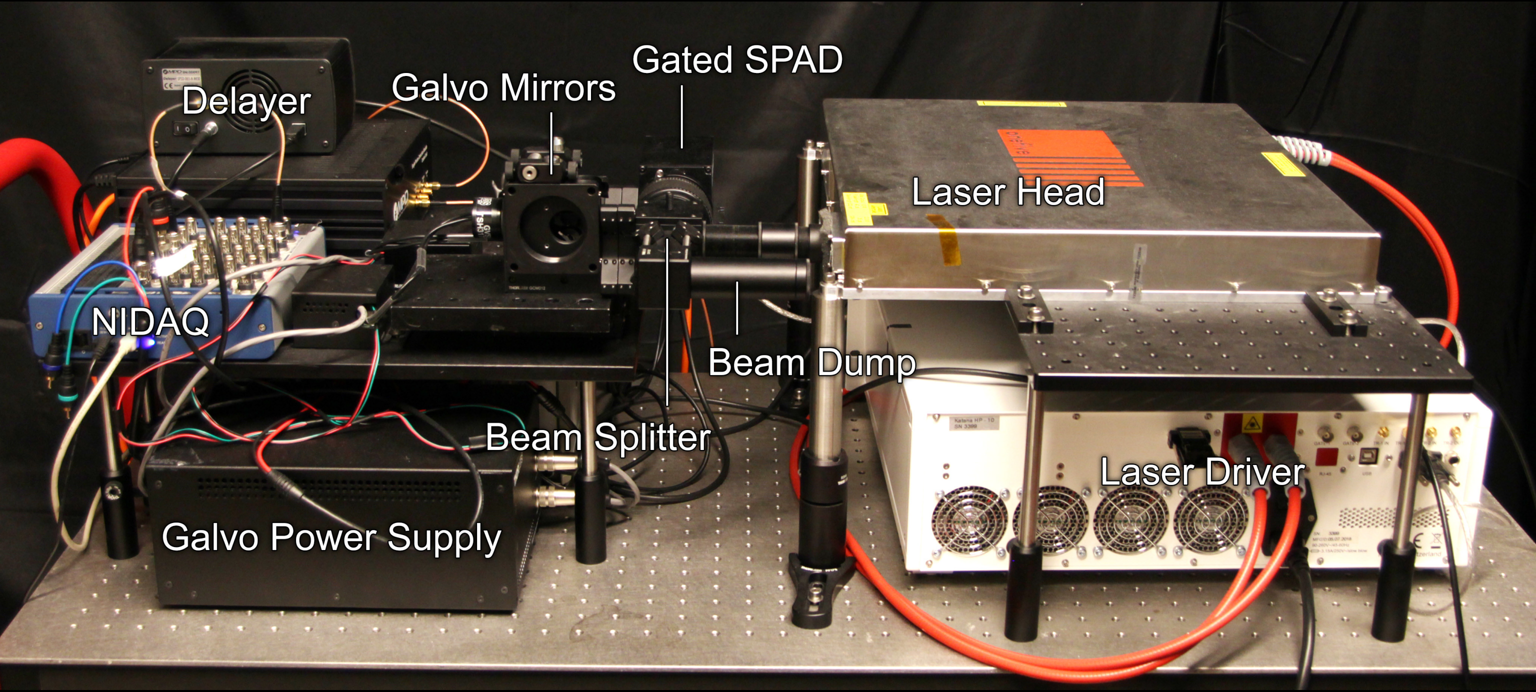

hardware prototype

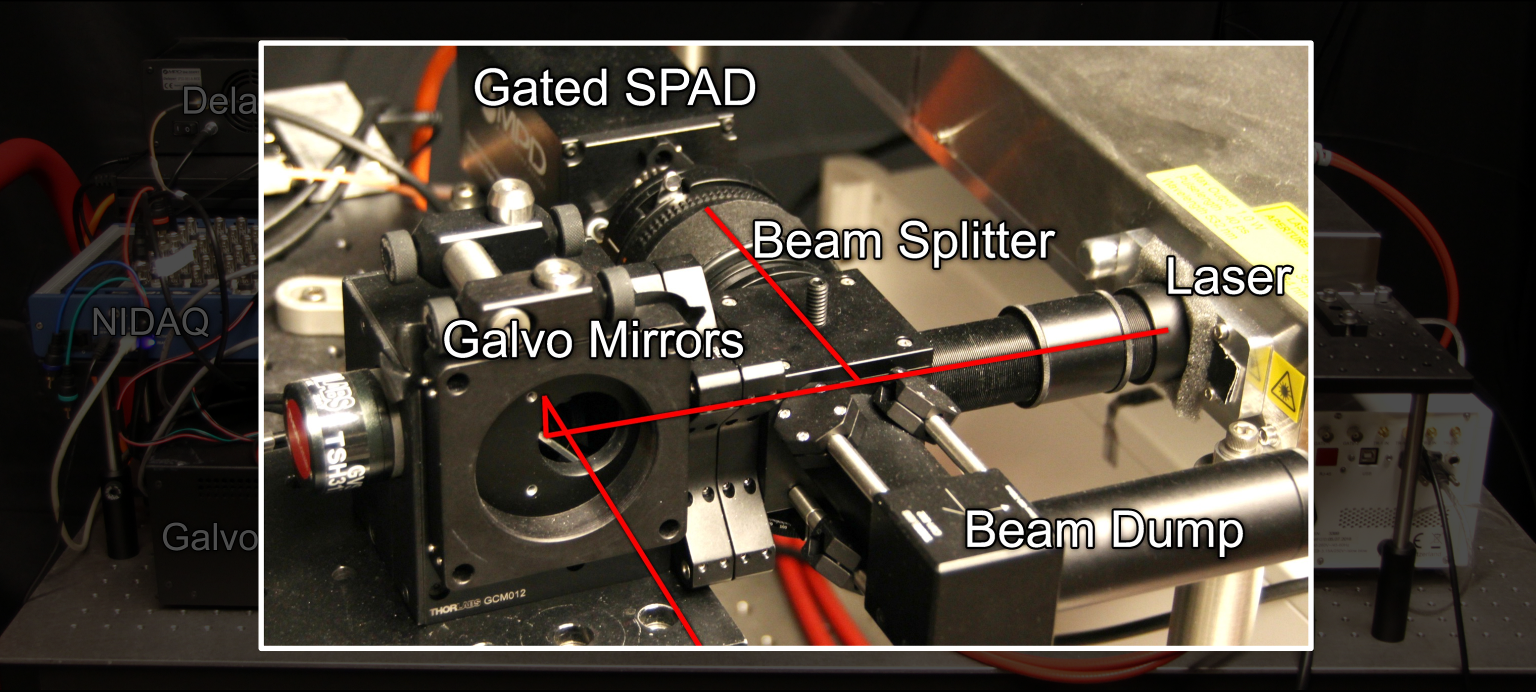

real-time scanning

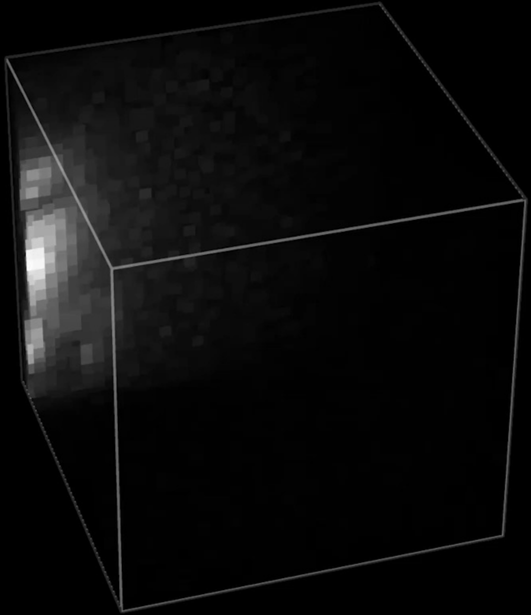

Framerate: 4 Hz
Resolution: 32 x 32
Dimensions: 2 m x 2 m x 2 m
Reconstruction time: ~1 s per frame

Lindell et al., SIGGRAPH 2019

Outlook

Directional Light-Cone Transform

hidden scene

Recovered surface

[Young et al., CVPR 2020]

Outlook

Keyhole NLOS Imaging

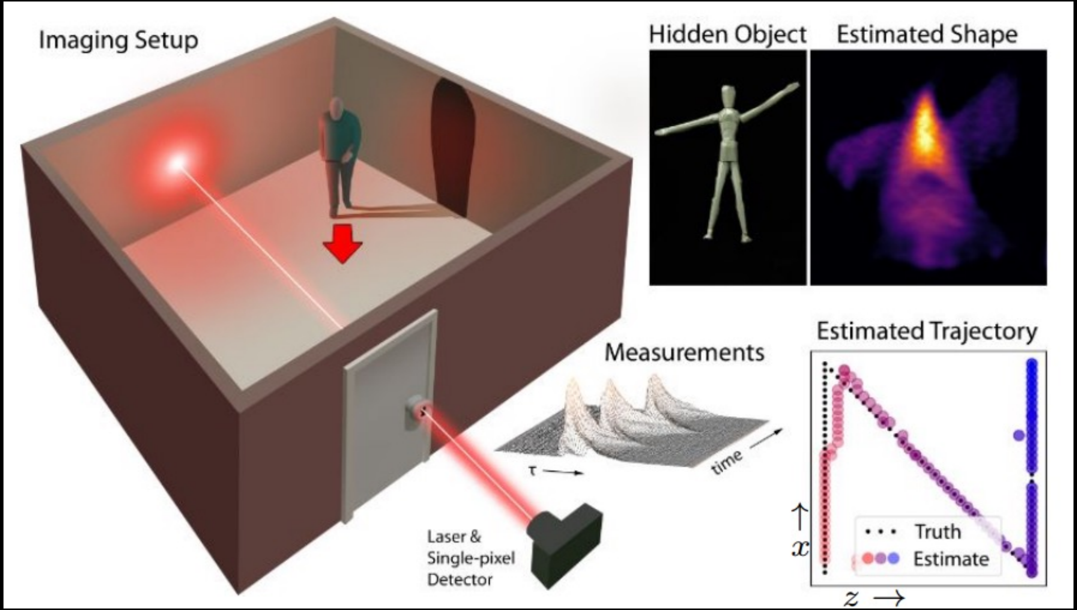

[Metzler et al., IEEE TCI 2021]

overview

- Time-resolved imaging
- Single-photon avalanche diodes (SPADs)
- Single-photon lidar
- Non-line-of-sight imaging (part 2)
- Imaging through scattering media

Time-resolved active imaging

Challenges

- very few returning photons
- information is 'scrambled' by scattering

Imaging through scattering media

light enters here and begins scattering

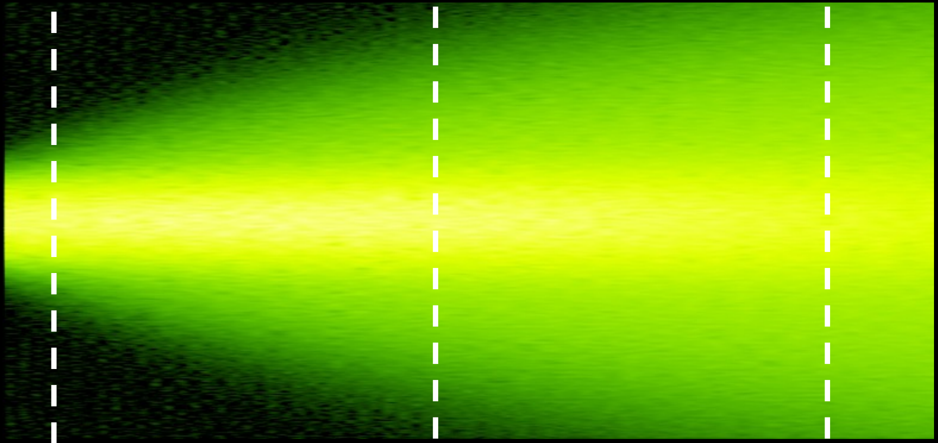

ballistic regime

increased scattering

random walk

Imaging through scattering media

light enters here and begins scattering

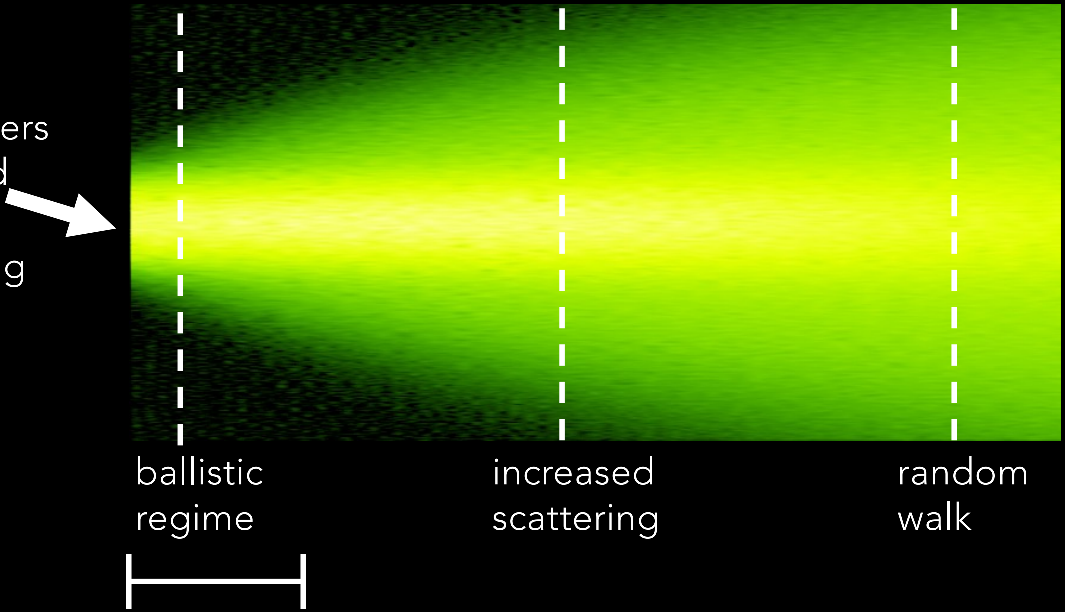

Imaging through scattering media

light enters here and begins scattering

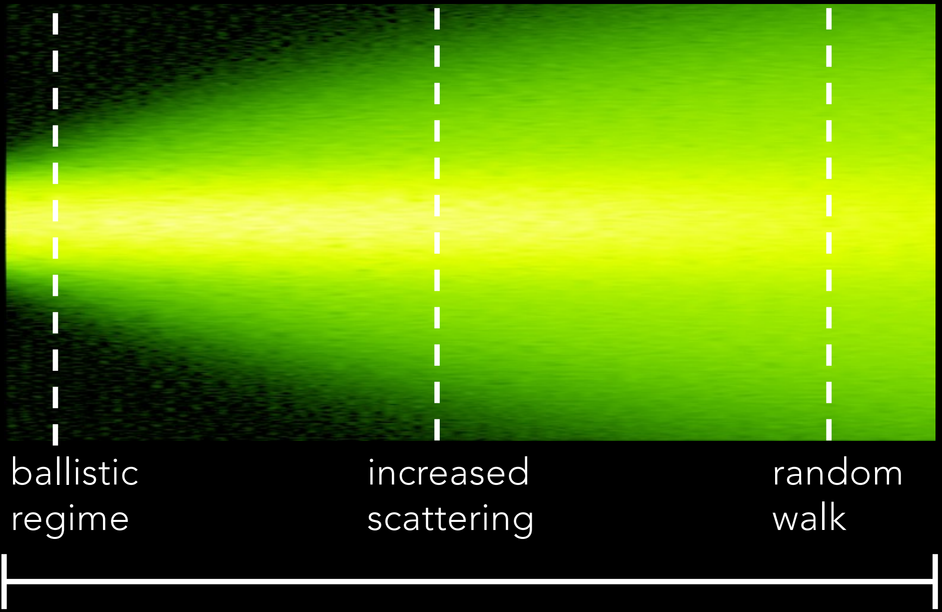

Imaging through scattering media

this work

> 6 transport mean free paths (TMFP)

Related Work

Ballistic imaging

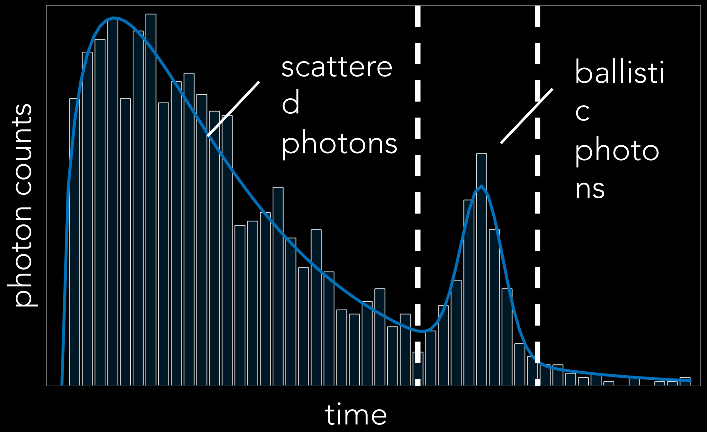

[Wang '91], [Redo-Sanchez '16], [Satat '18], …

Related Work

Ballistic imaging

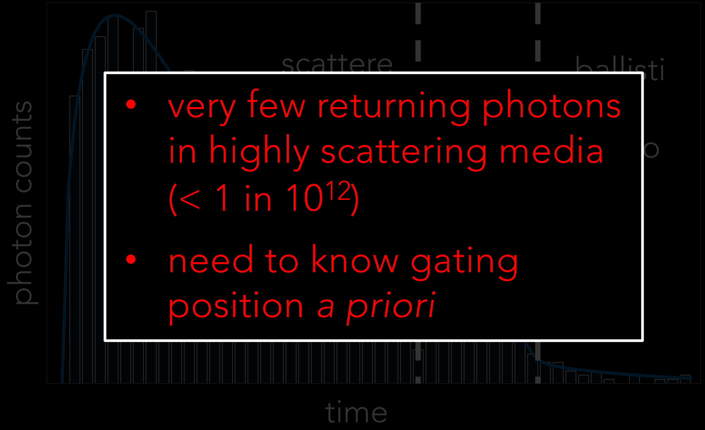

- very few returning photons in highly scattering media (< 1 in 10^{12})
- need to know gating position *a priori*

[Wang '91], [Redo-Sanchez '16], [Satat '18], …

Related Work

Ballistic imaging

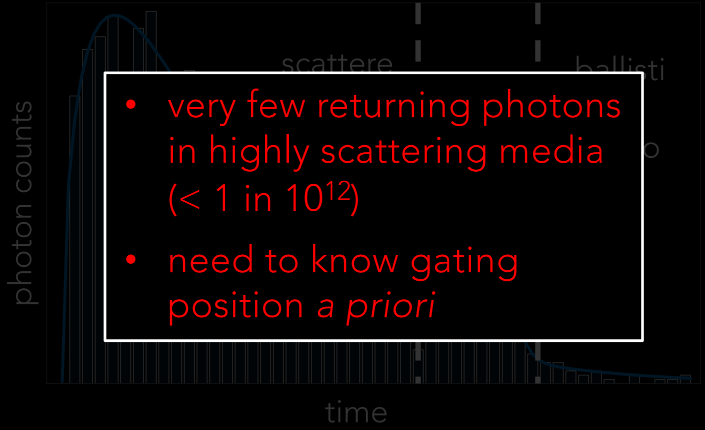

- very few returning photons in highly scattering media (< 1 in 10^{12})
- need to know gating position *a priori*

[Wang '91], [Redo-Sanchez '16], [Satat '18], …

Diffuse Optical Tomography

2D

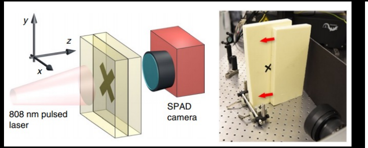

[Lyons '19]

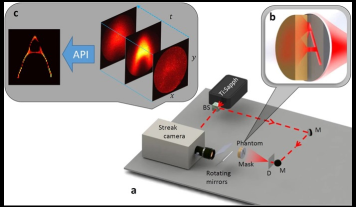

[Satat '16]

3D

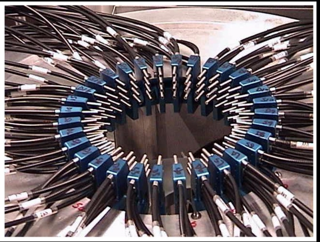

[Hajihashemi '12]

Related Work

Ballistic imaging

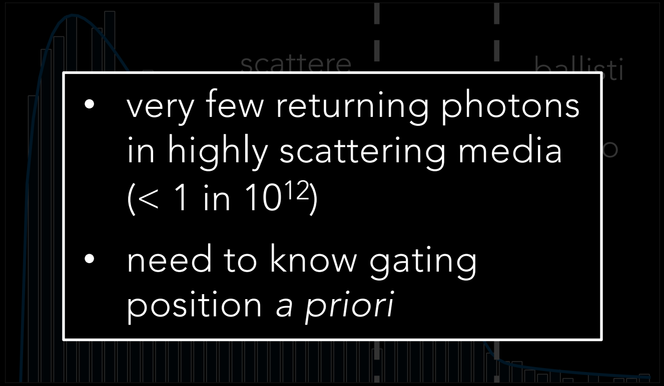

- very few returning photons in highly scattering media (< 1 in 10^{12})
- need to know gating position *a priori*

[Wang '91], [Redo-Sanchez '16], [Satat '18], …

Diffuse Optical Tomography

2D

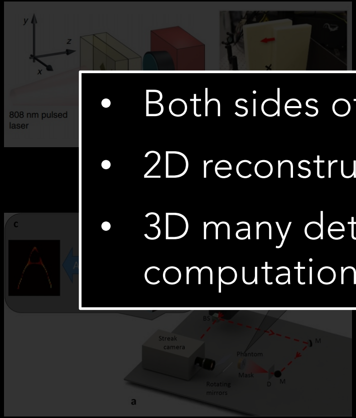

3D

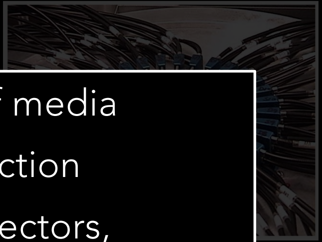

- Both sides of media
- 2D reconstruction
- 3D many detectors, computationally expensive

[Satat '16]

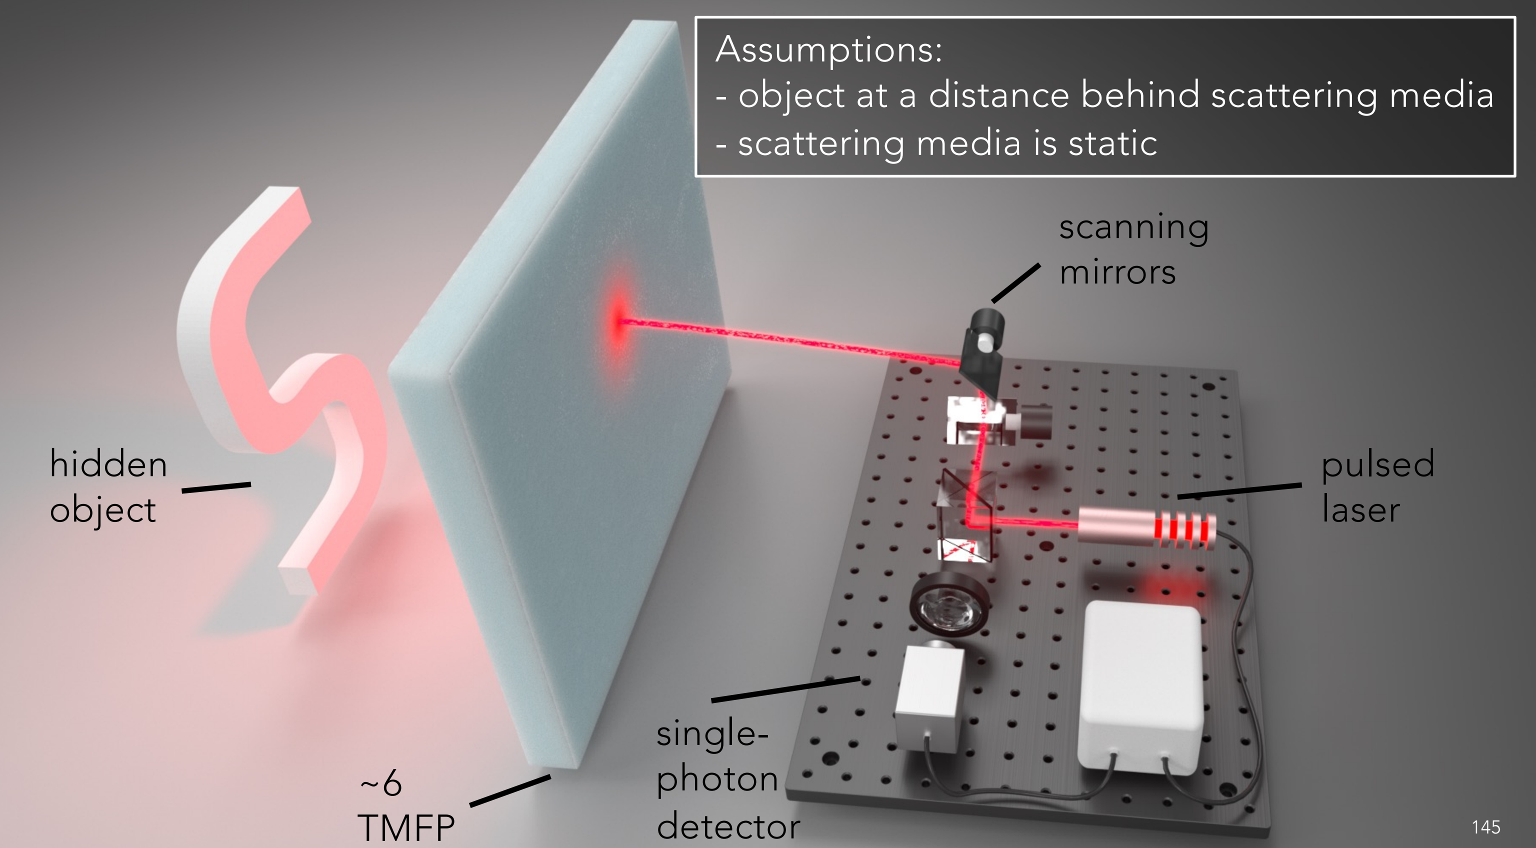
Assumptions:
- object at a distance behind scattering media
- scattering media is static
scanning
mirrors
hidden
object
pulsed
laser
single-
photon
detector
~6
TMFP

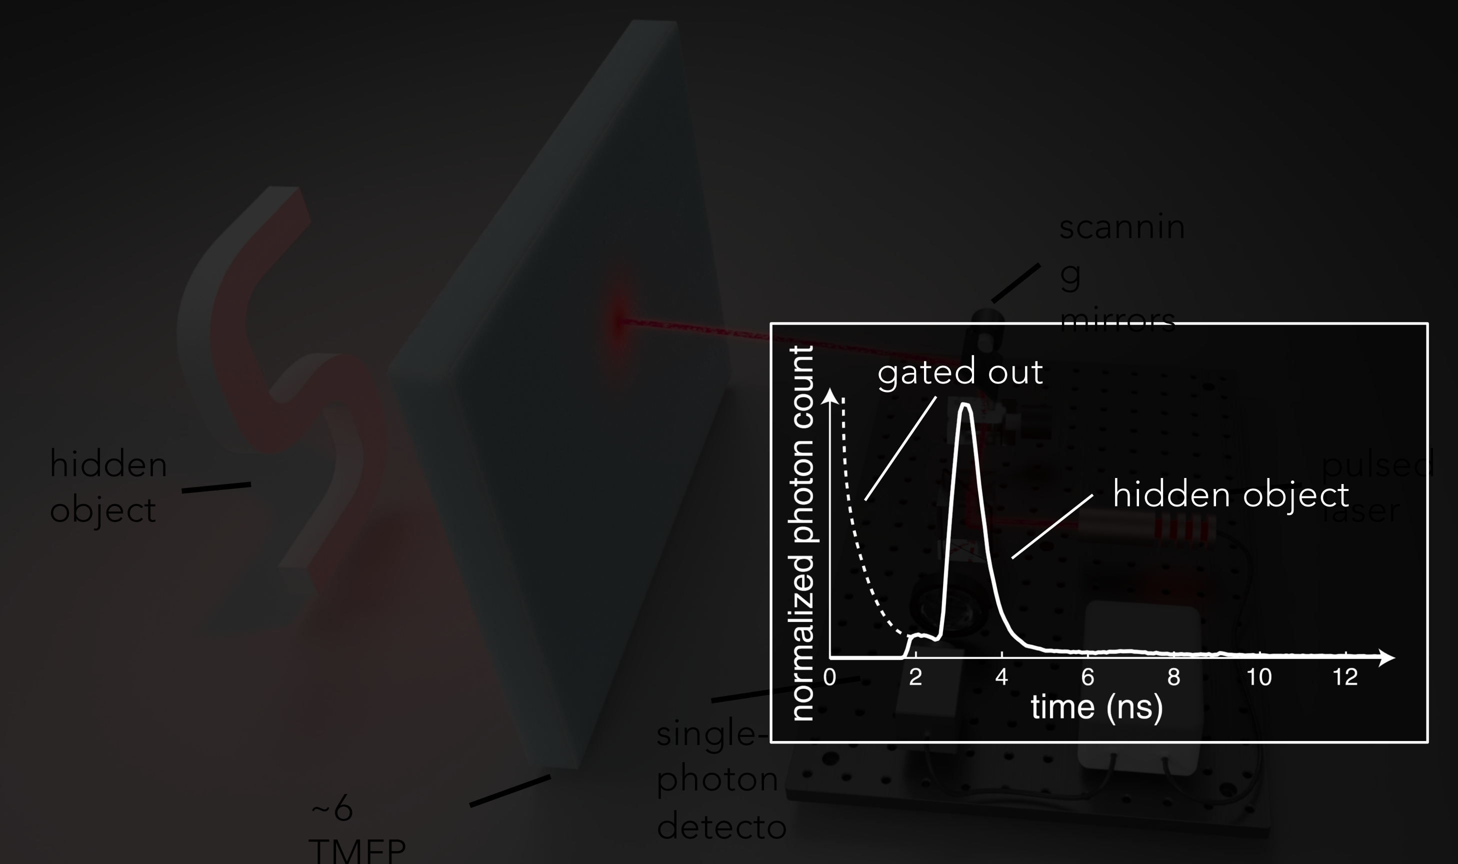
scannin
g
mirrors
hidden
object
pulsed
laser
~6
TMFP
single-
photon
detecto
gated out
hidden object
normalized photon count
0
2
4
6
8
10
12
time (ns)

Results

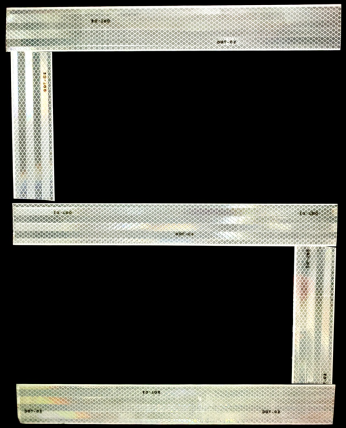

hidden object

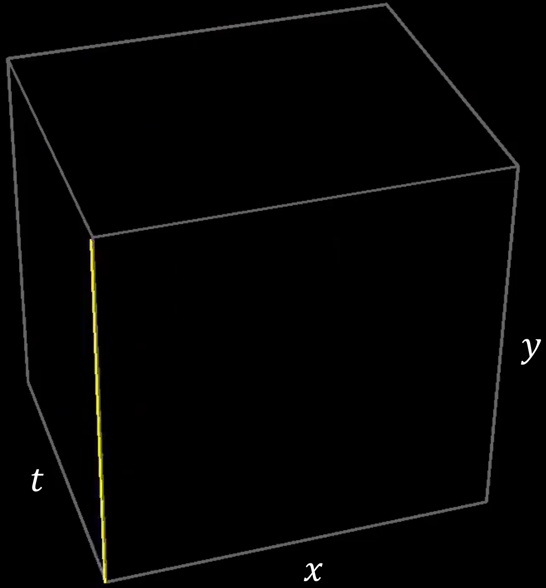

captured measurements
(0.6 m × 0.6 m × 3.3 ns)

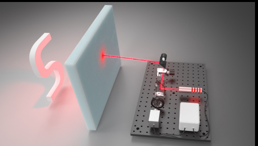

total acquisition time: 1 min. (60 ms/sample)

Results

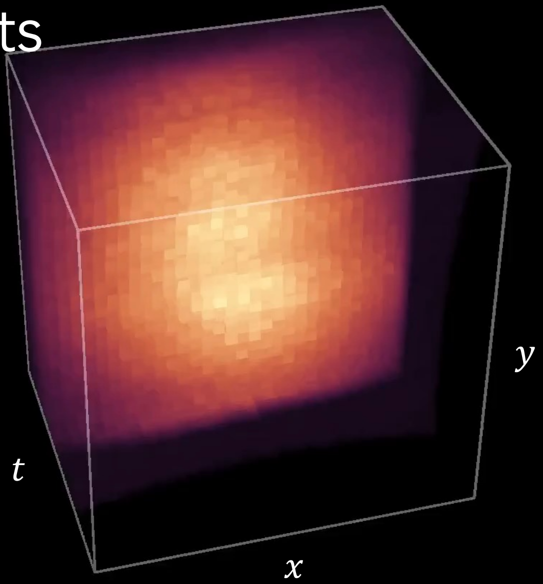

captured measurements

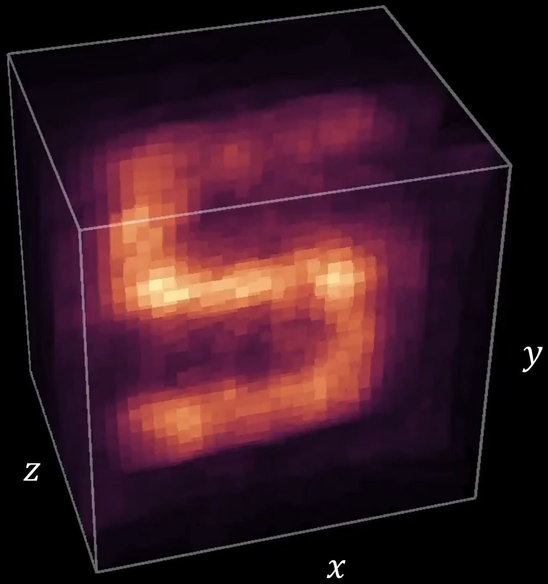

reconstruction (50 ms)
(0.6 m × 0.6 m × 0.5 m)

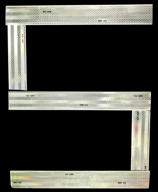

total acquisition time: 1 min. (60 ms/sample)

Method

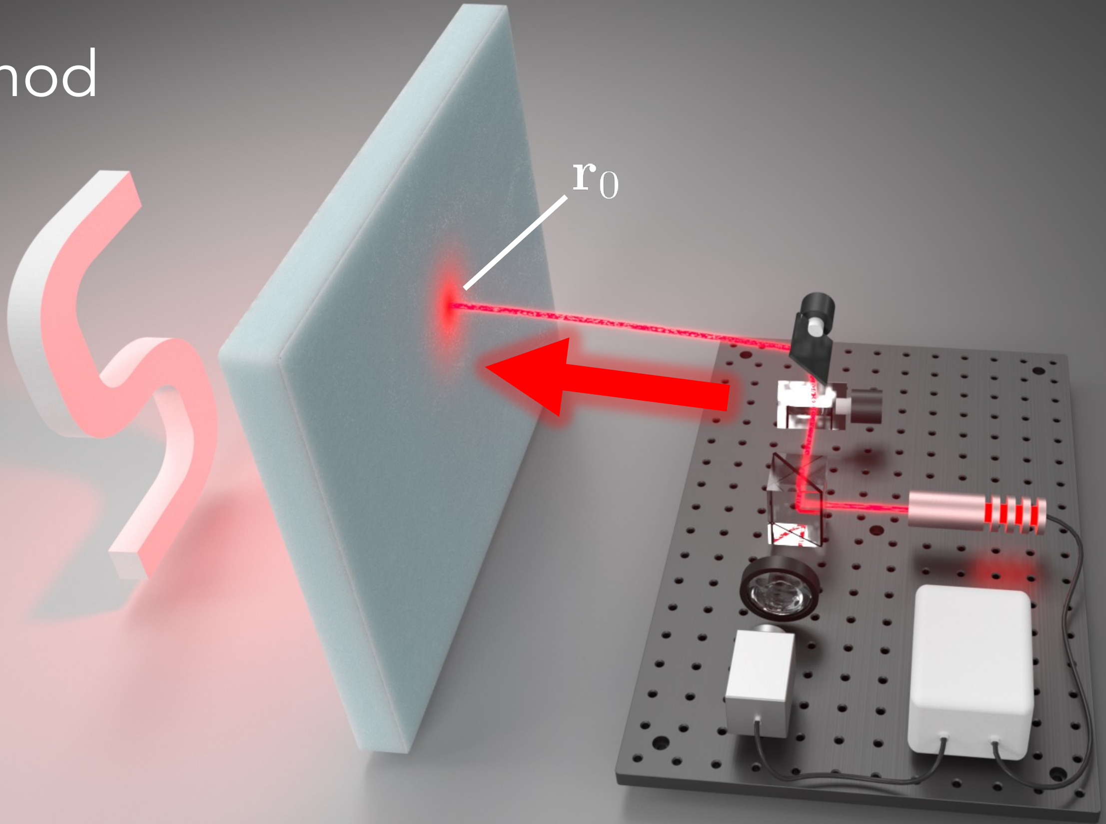

Method

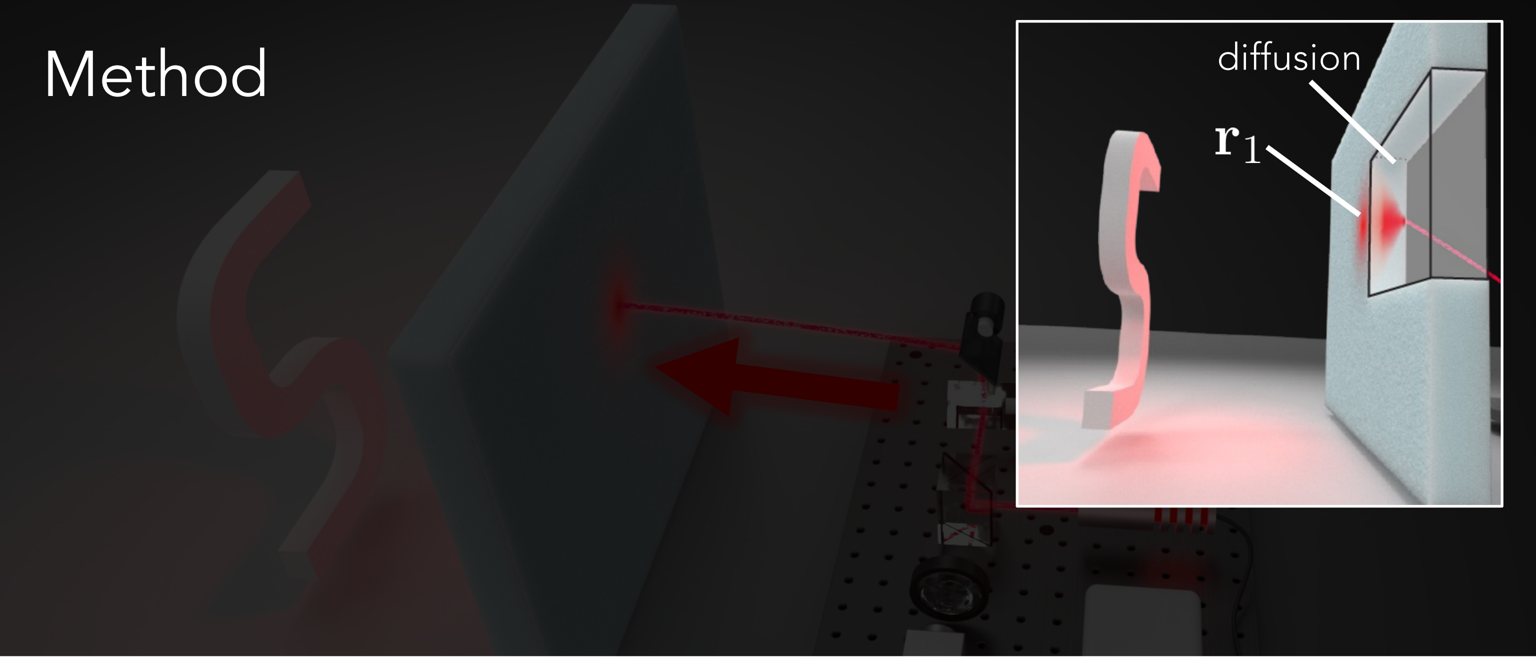

$$\phi(t, \mathbf{r_0}, \mathbf{r_1}) = \frac{c}{(4\pi Dct)^{3/2}} \exp\left(-\frac{\|\mathbf{r_1} - \mathbf{r_0}\|_2^2}{4Dct} - \mu_a ct\right)$$

Method

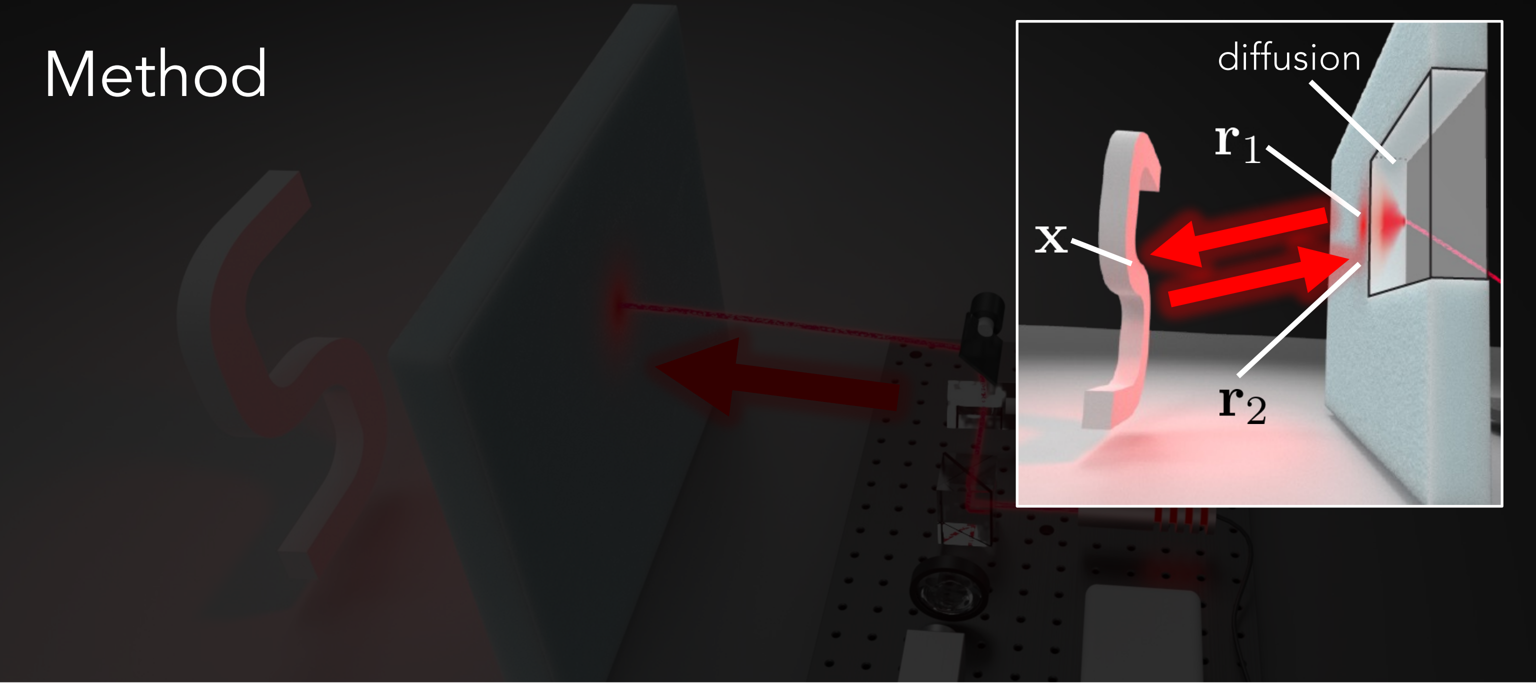

$$I(t', \mathbf{r_1}, \mathbf{r_2}) = \int_\Psi \boxed{f(\mathbf{x}, \mathbf{r_1}) f(\mathbf{x}, \mathbf{r_2})}\boxed{\delta(ct' - \|\mathbf{x} - \mathbf{r_1}\| - \|\mathbf{x} - \mathbf{r_2}\|)} d\mathbf{x}$$

Method

Need evaluate all paths $\mathbf{r_1} \rightarrow \mathbf{x} \rightarrow \mathbf{r_2}$

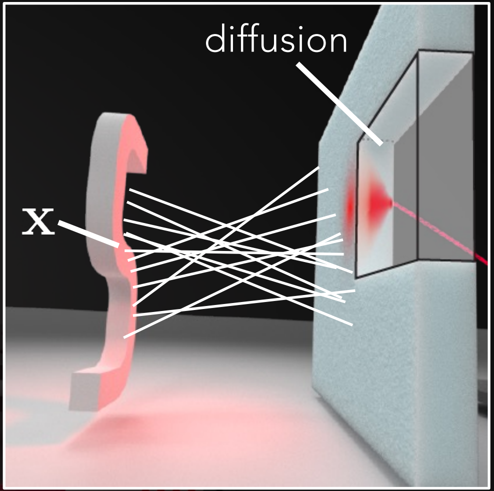

$$I(t', \mathbf{r_1}, \mathbf{r_2}) = \int_{\Psi} f(\mathbf{x}, \mathbf{r_1}) f(\mathbf{x}, \mathbf{r_2}) \delta(ct' - \|\mathbf{x} - \mathbf{r_1}\| - \|\mathbf{x} - \mathbf{r_2}\|) d\mathbf{x}$$

Method

non-line-of-sight imaging

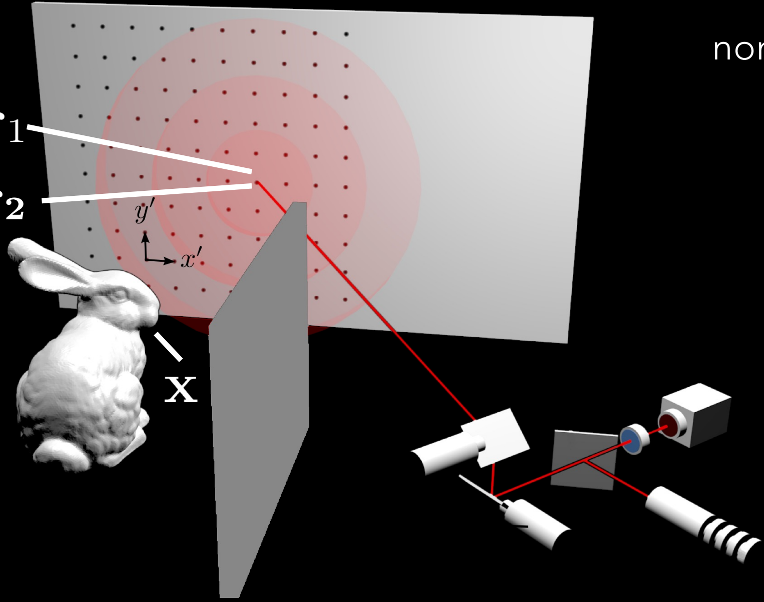

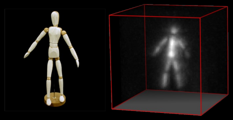

[O'Toole '18]

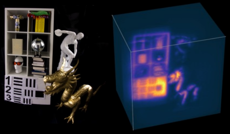

[Lindell '19]

$$I(t', \mathbf{r_1}, \mathbf{r_2}) = \int_{\Psi} f(\mathbf{x}, \mathbf{r_1}) f(\mathbf{x}, \mathbf{r_2}) \delta(ct' - \|\mathbf{x} - \mathbf{r_1}\| - \|\mathbf{x} - \mathbf{r_2}\|) d\mathbf{x}$$

Method

non-line-of-sight imaging

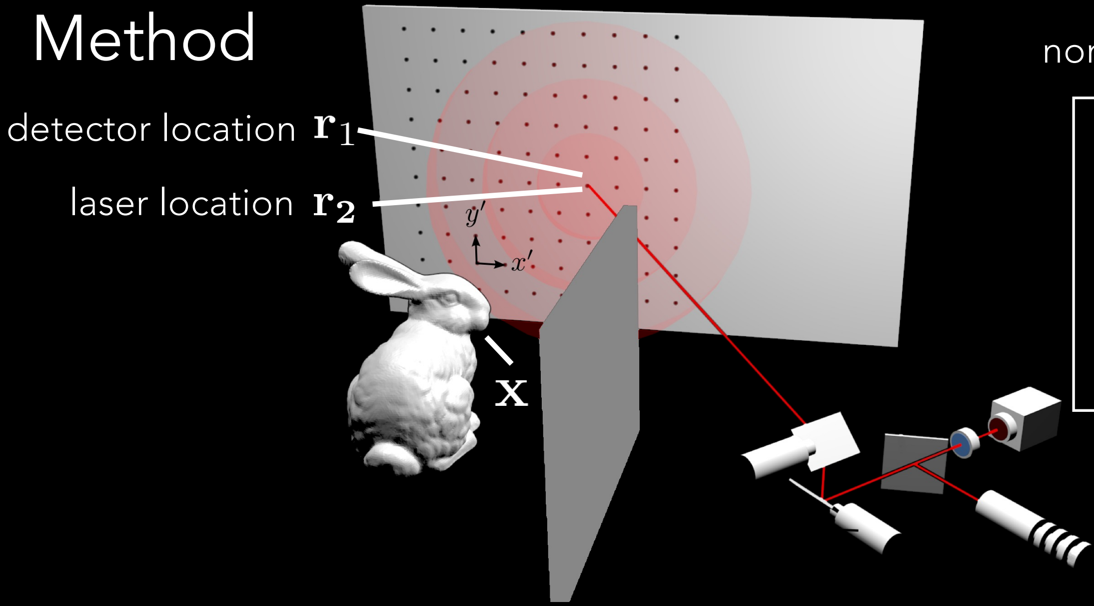

Efficient inversion if:

$$\mathbf{r_1} = \mathbf{r_2}$$

$$I(t, \mathbf{r_1}, \mathbf{r_1})$$

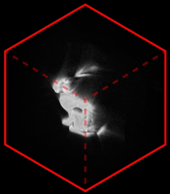

$$I(t', \mathbf{r_1}, \mathbf{r_2}) = \int_{\Psi} f(\mathbf{x}, \mathbf{r_1}) f(\mathbf{x}, \mathbf{r_2}) \delta(ct' - \|\mathbf{x} - \mathbf{r_1}\| - \|\mathbf{x} - \mathbf{r_2}\|) d\mathbf{x}$$

Method

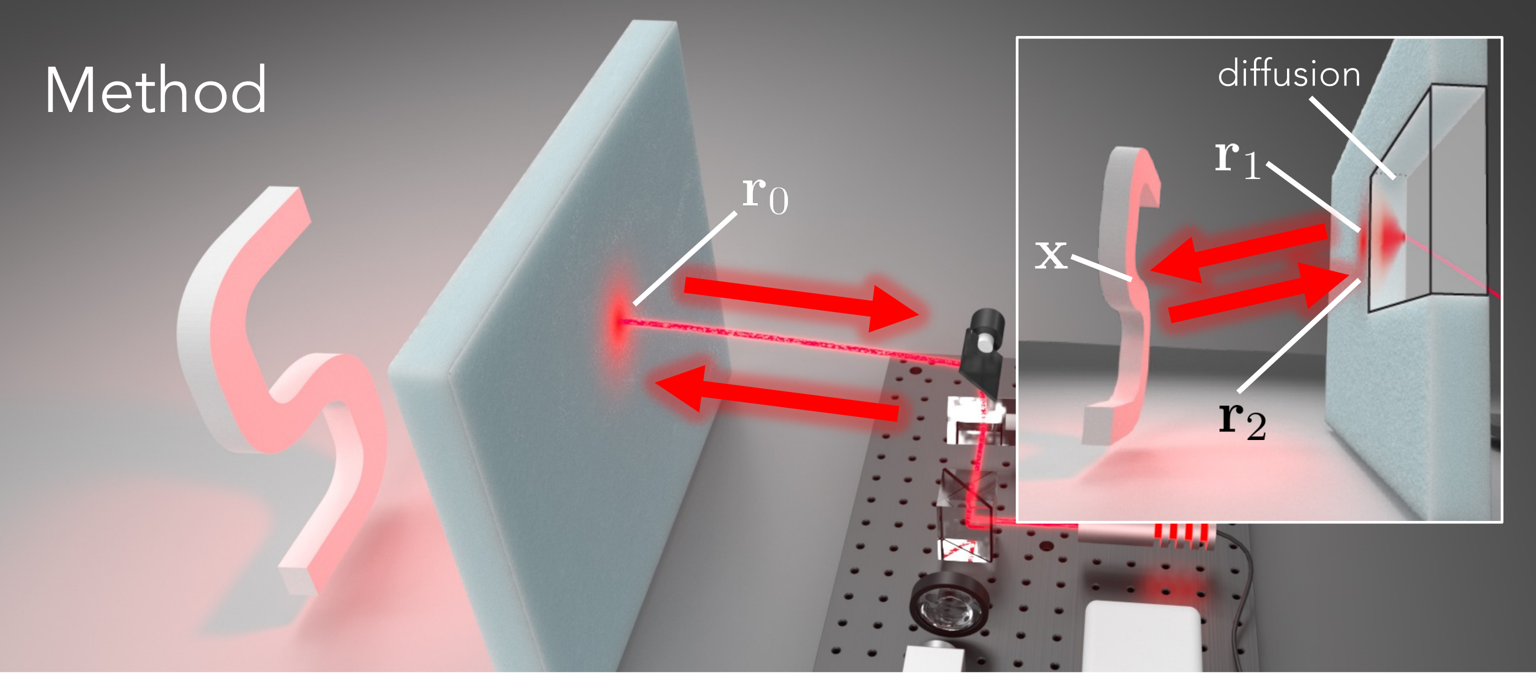

$$\phi(t, \mathbf{r_0}, \mathbf{r_1}) = \frac{c}{(4\pi Dct)^{3/2}} \exp\left(-\frac{\|\mathbf{r_1} - \mathbf{r_0}\|_2^2}{4Dct} - \mu_a ct \right)$$

Method

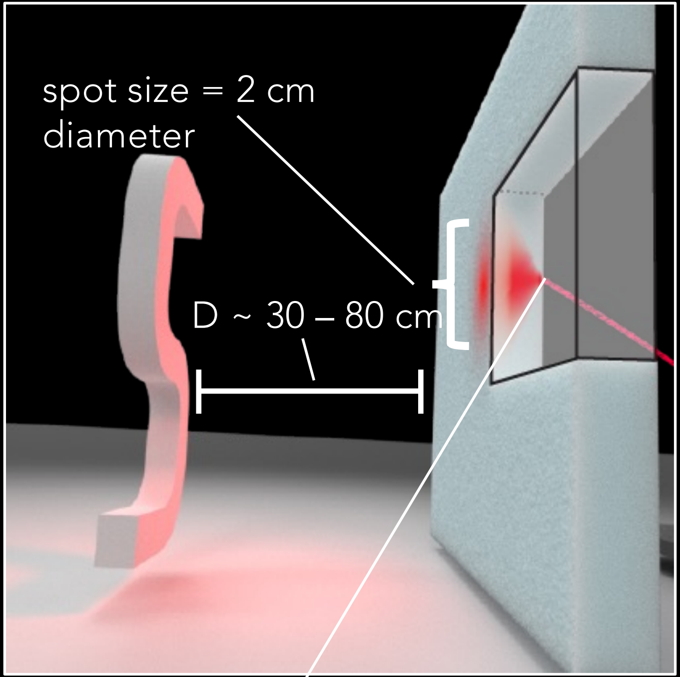

confocal: illuminate and image here

Method

Approximation:
Approximate measured light as scattering back to the same spot. $\mathbf{r_1} = \mathbf{r_2}$

Error ~ (spot size)^2 / (2 * distance) << 1 cm

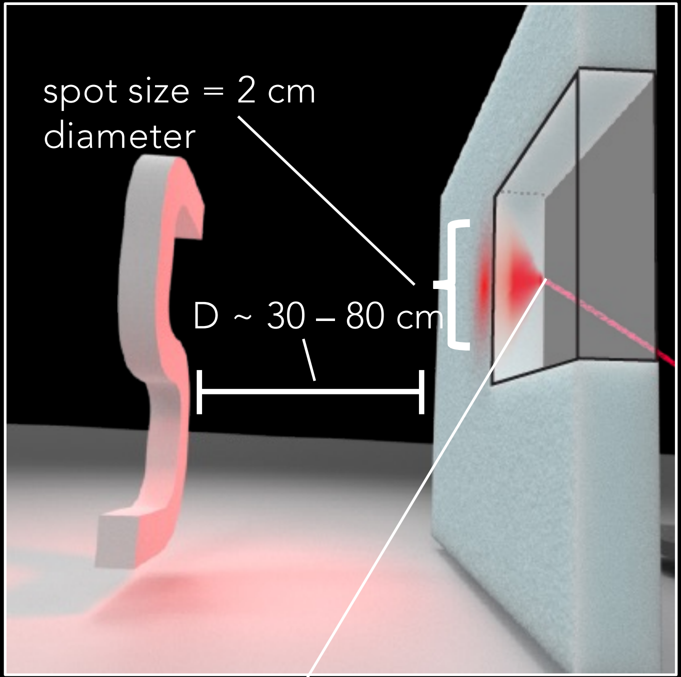

Method

Approximation:
Approximate measured light as scattering back to the same spot. $\mathbf{r_1} = \mathbf{r_2}$

Error ~ (spot size)^2 / (2 * distance) << 1 cm

measurements

$$\hat{\tau}(t, \mathbf{r_0}) =$$

$$\phi(t, \mathbf{r_0}, \mathbf{r_1}) * \phi(t, \mathbf{r_0}, \mathbf{r_1}) * I(t, \mathbf{r_1}, \mathbf{r_1})$$

diffusion kernels NLOS model

Can use efficient NLOS inversion!

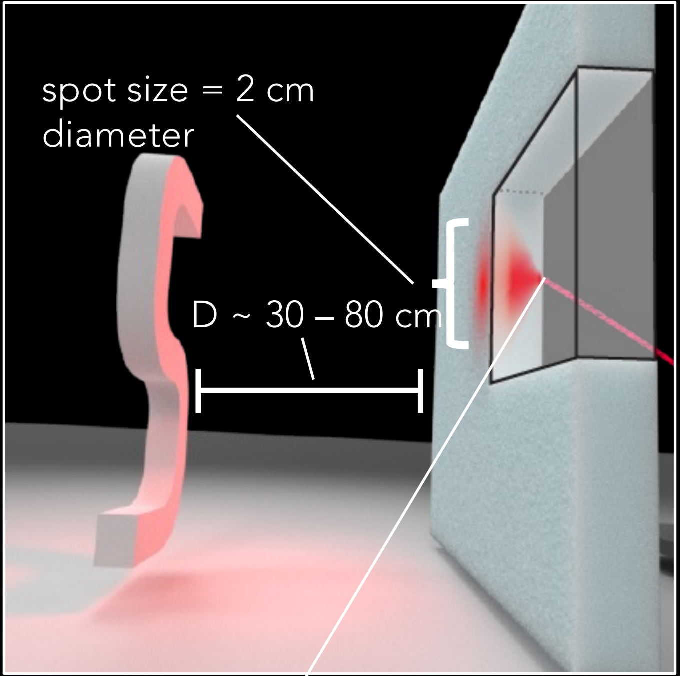

confocal: illuminate and image here

Reconstruction

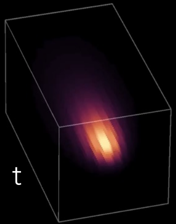

diffusion filter

$$\left[\phi(t, \mathbf{r_0}, \mathbf{r_1}) * \phi(t, \mathbf{r_0}, \mathbf{r_1})\right]^{-1}$$

inverse filter

$*$

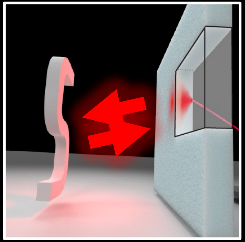

$$\left[I(t, \mathbf{r_1}, \mathbf{r_1})\right]^{-1}$$

inverse light transport

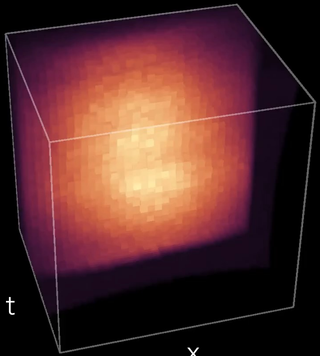

measurements

Hardware

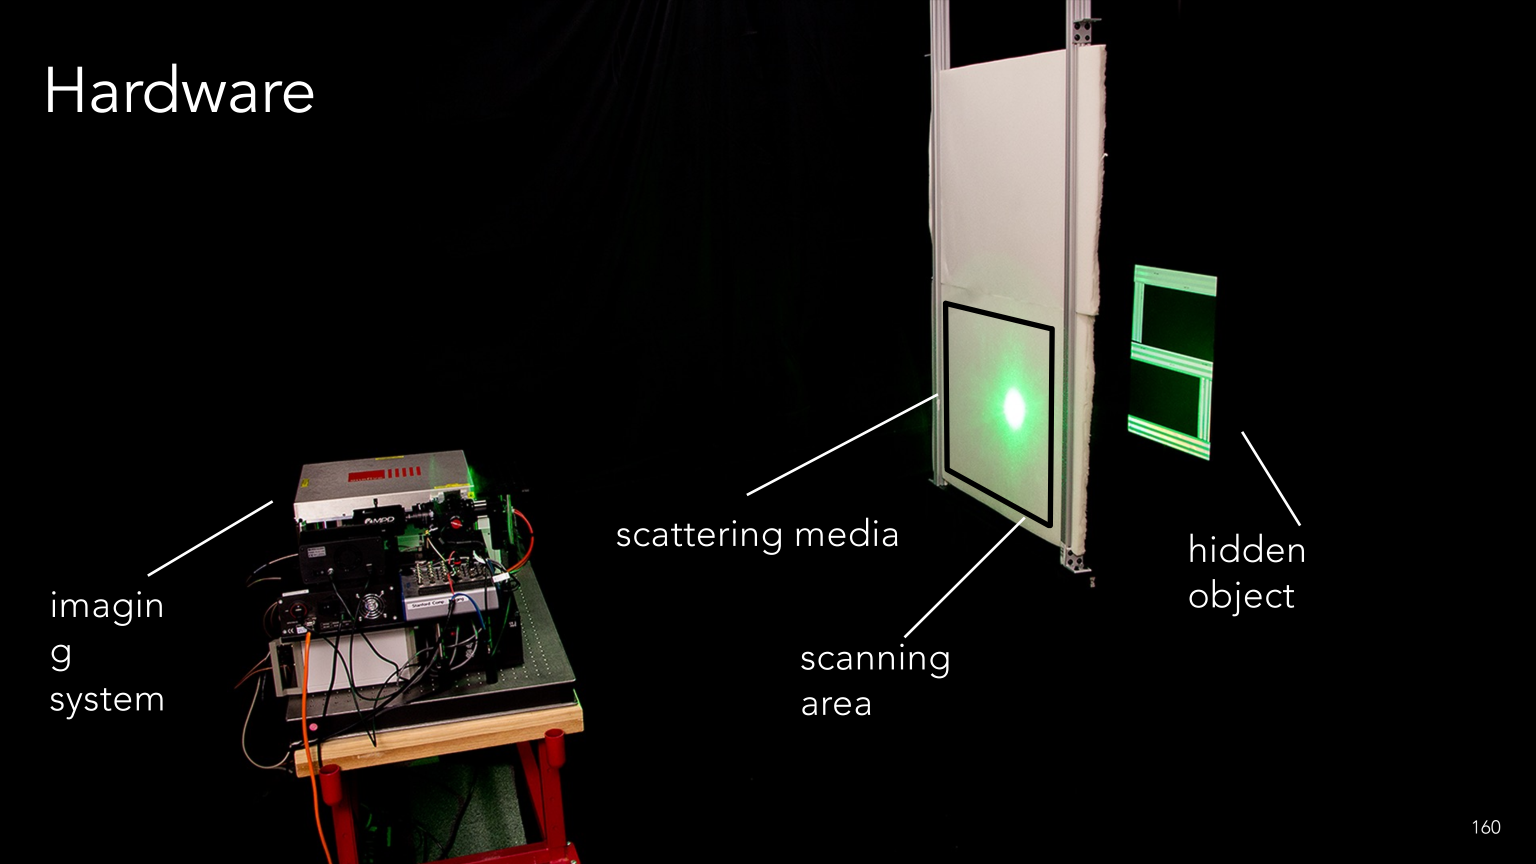
imagin
g
system
scattering media
scanning
area
hidden
object

Hardware

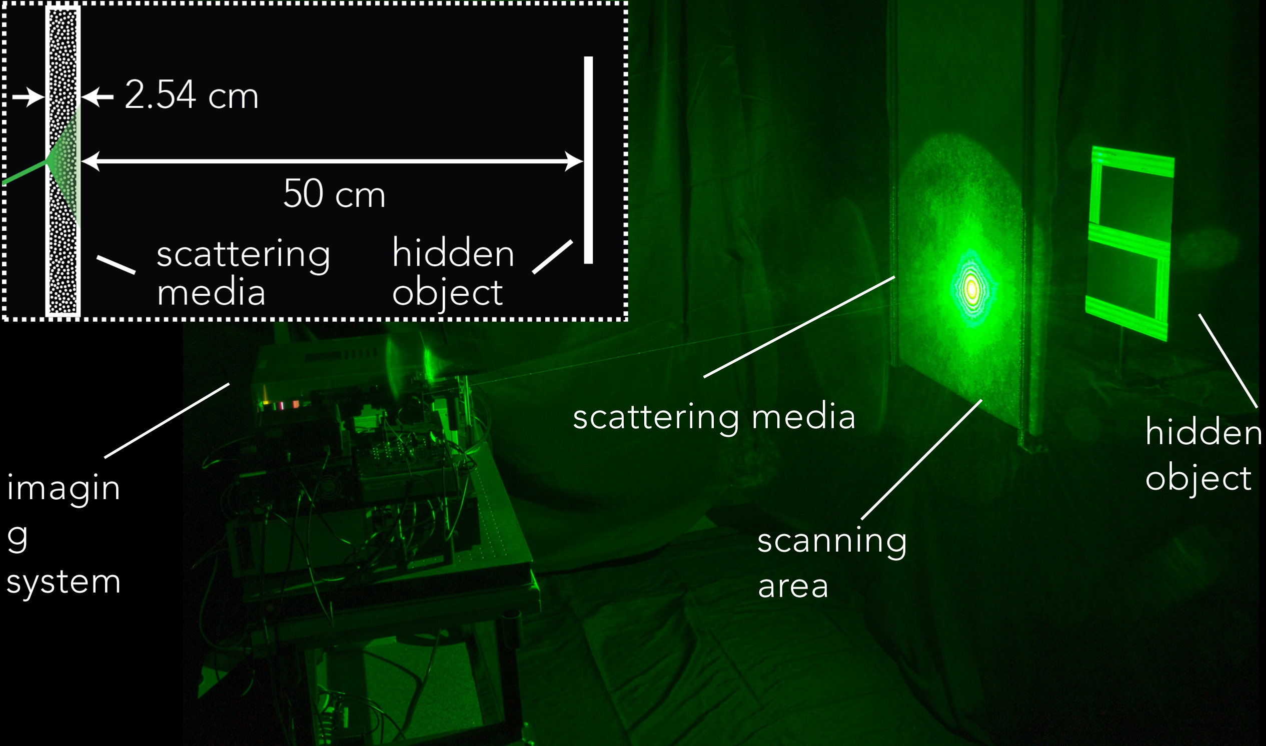
2.54 cm
50 cm
scattering media
hidden object
imagin
g
system
scattering media
scanning
area
hidden
object

Hardware

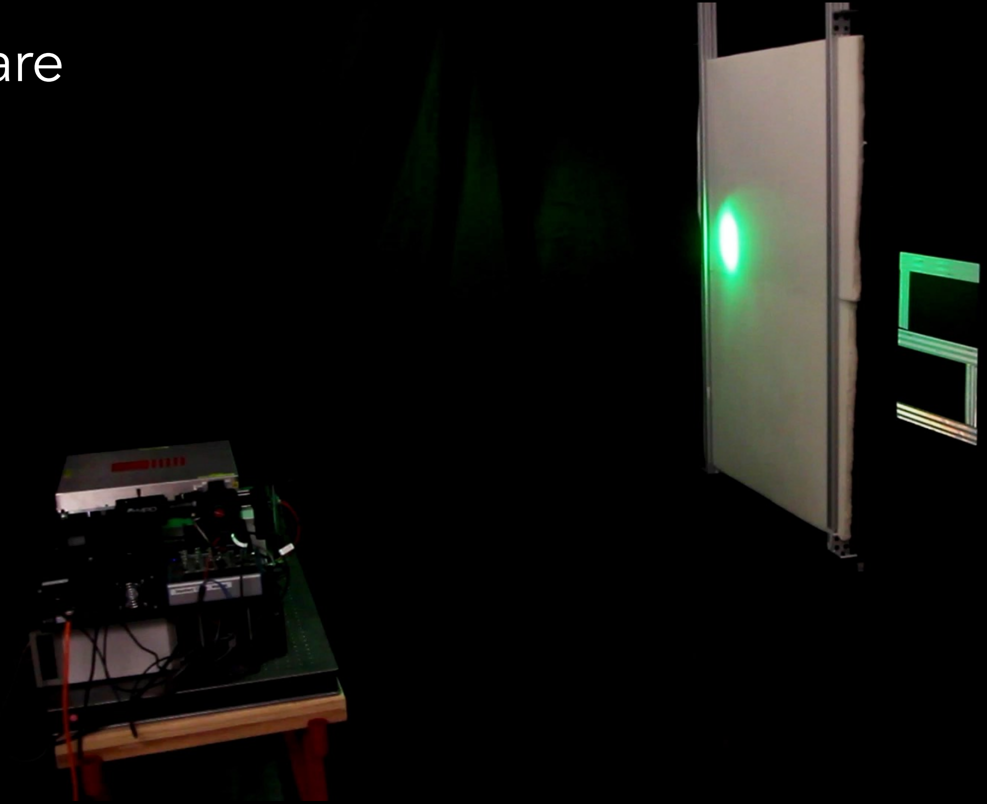

Results

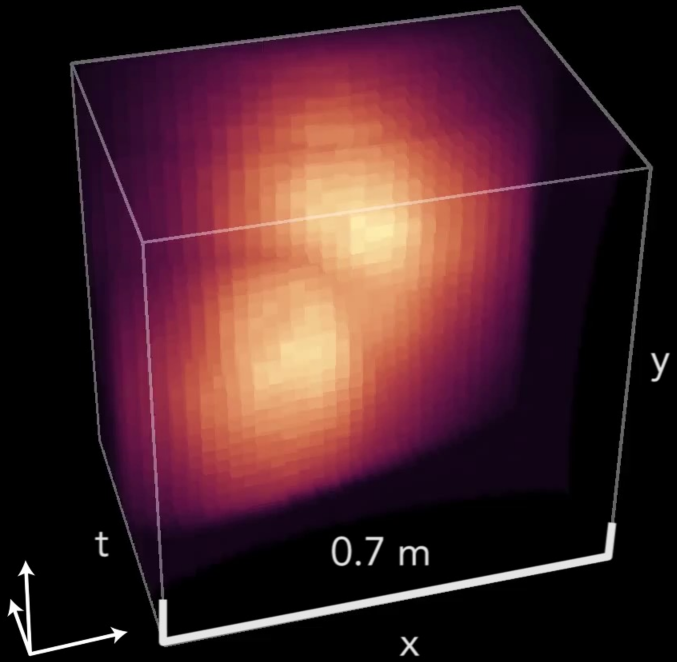

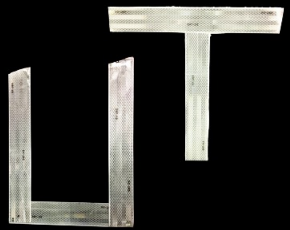

measurements

Lindell et al., Nat. Commun. 2020

Results

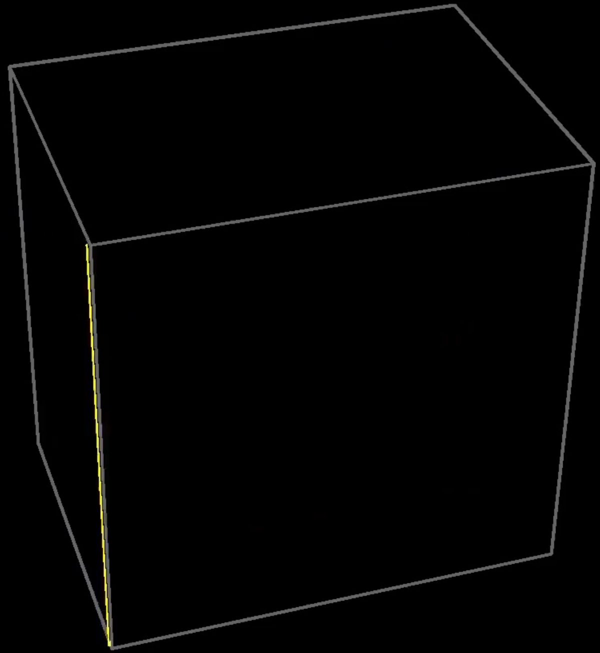

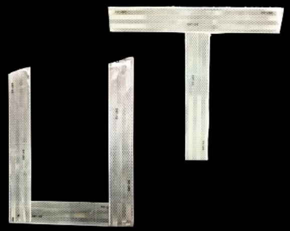

measurements

Lindell et al., Nat. Commun. 2020

Results

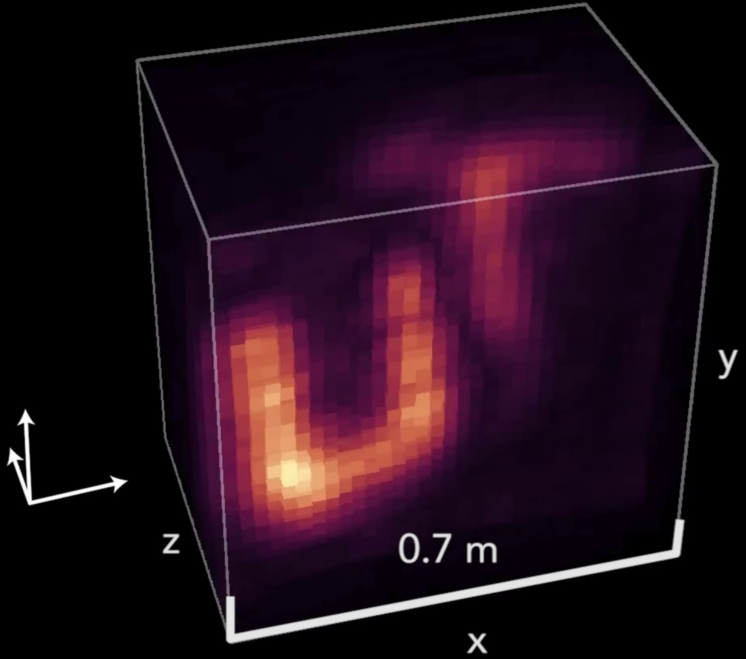

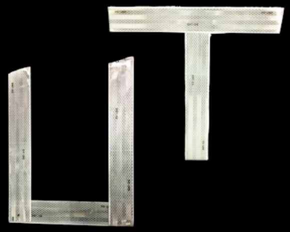

reconstruction

Lindell et al., Nat. Commun. 2020

Results

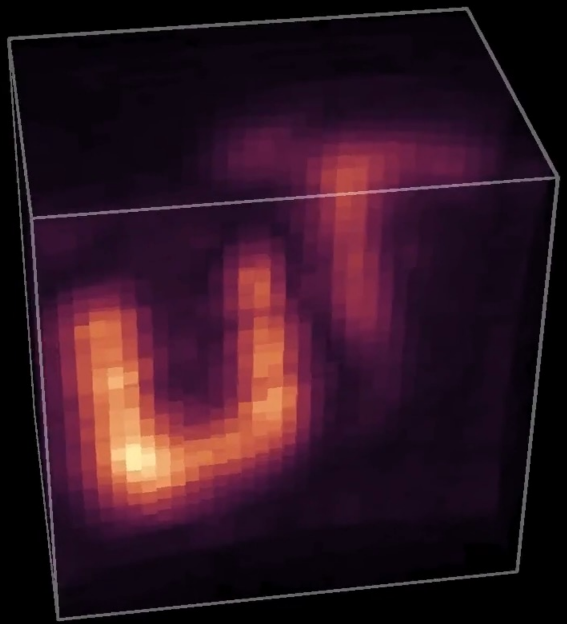

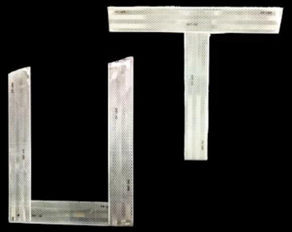

reconstruction

Lindell et al., Nat. Commun. 2020

Results

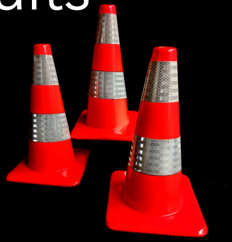

traffic
cones

reflective mannequin

diffuse letter

Lindell et al., Nat. Commun. 2020

Results

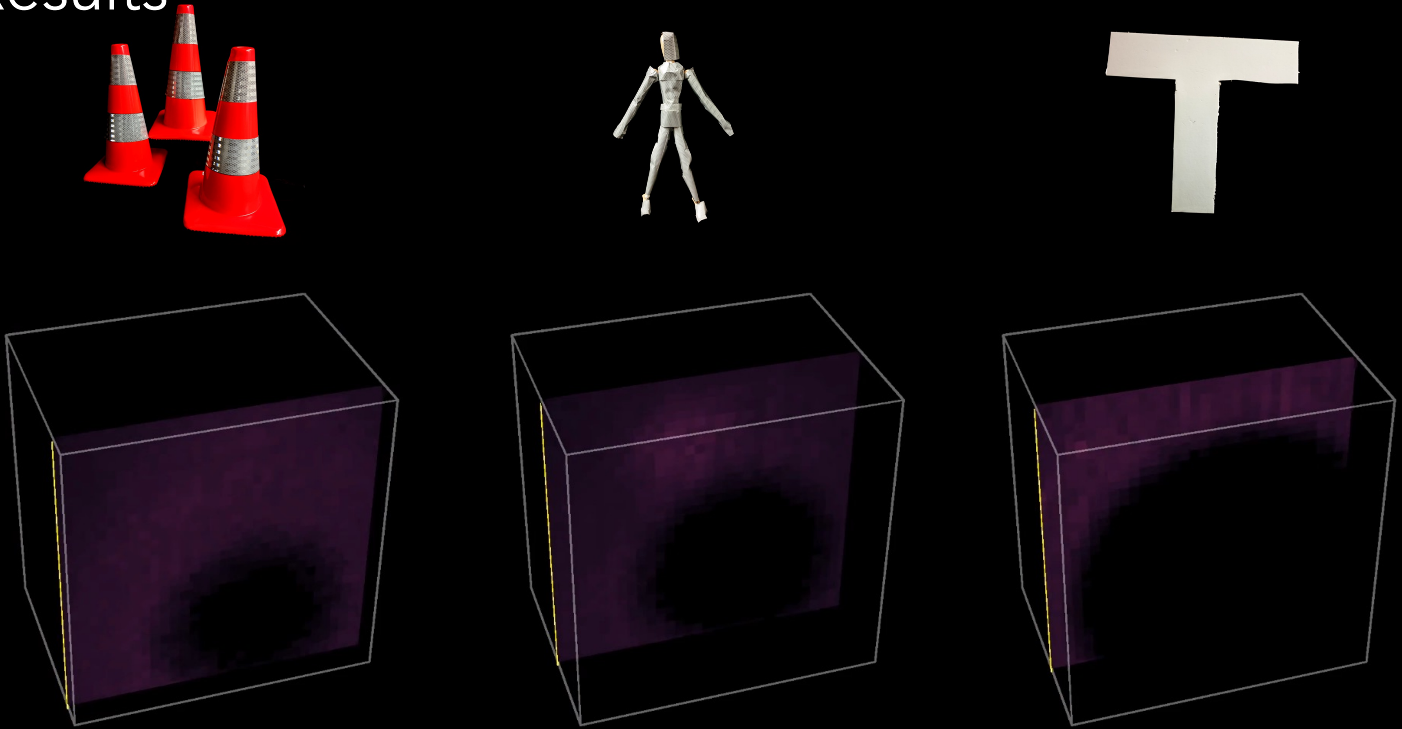

Lindell et al., Nat. Commun. 2020

Results

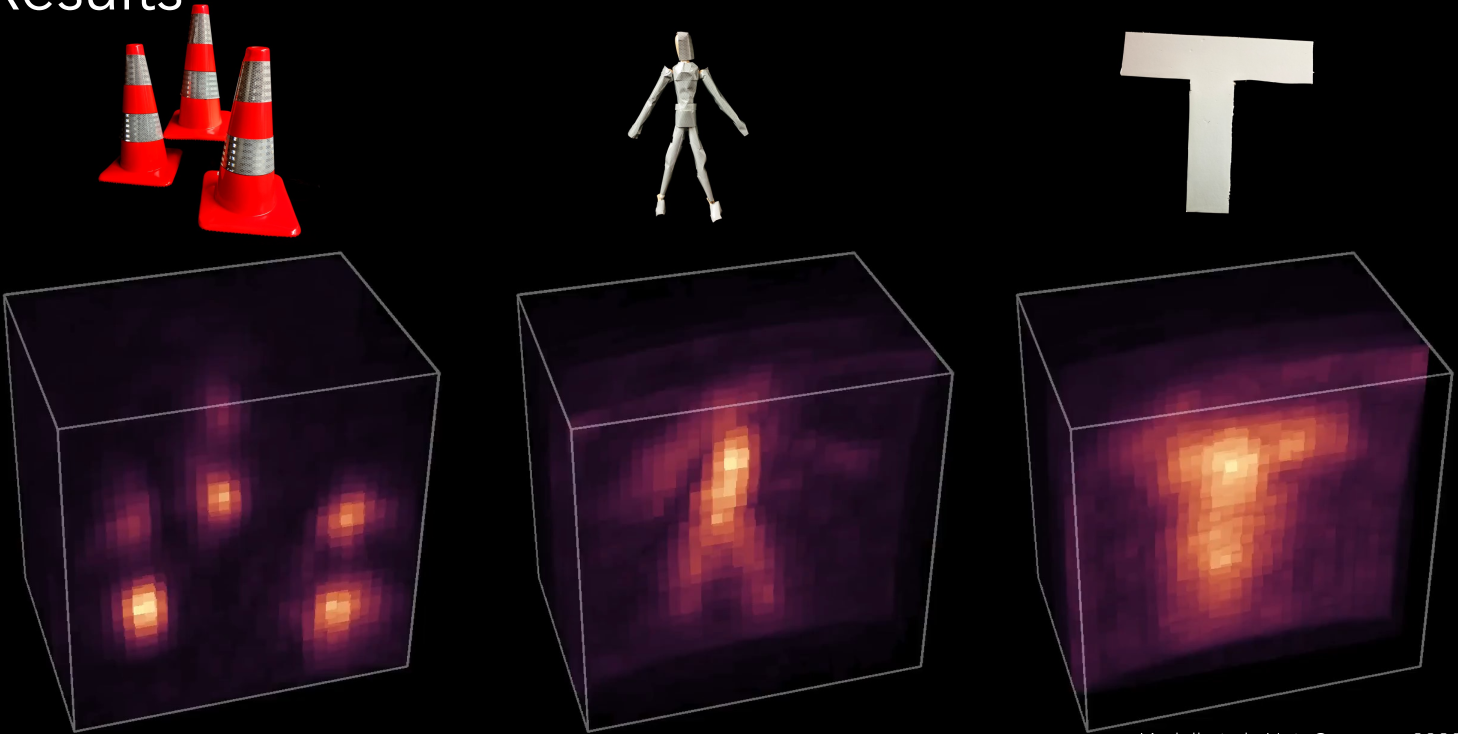

Lindell et al., Nat. Commun. 2020

Conclusion

- Time-resolved imaging + closed-form solution for imaging through scattering media
- future work
  - embedded, anisotropic media
  - priors, machine learning

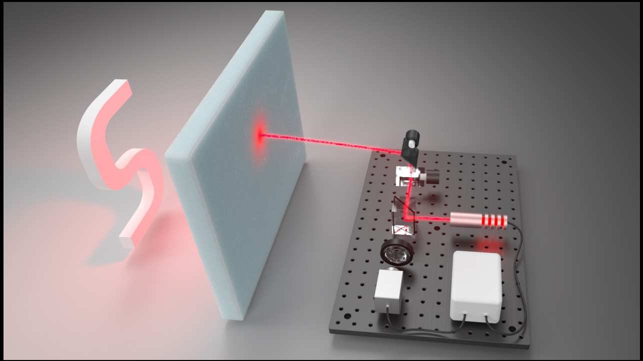

concluding remarks

- Many applications for time-of-flight imaging
  - Lidar
  - Non-line-of-sight
  - Transient imaging
  - Imaging through scatter

- New capabilities through combining emerging sensors with computation!

Next time…

Representing & processing signals with neural networks

$\mathbf{x}$

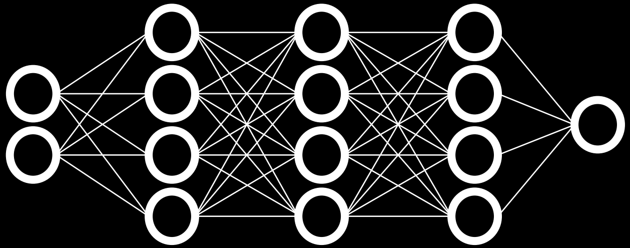

$$\Phi(\mathbf{x}) : \mathbb{R}^2 \mapsto \mathbb{R}$$

References

- D. B. Lindell and G. Wetzstein, “Three-dimensional imaging through scattering media based on confocal diffuse tomography,” Nature Communications, vol. 11, no. 4517, 2020. [J10] C. A. Metzler, D. B.
- C. A. Metzler, D. B. Lindell, G. Wetzstein, “Keyhole imaging: Non-line-of-sight imaging and tracking of moving objects along a single optical path at long standoff distances,” IEEE Trans. Comput. Imag., 2020, (Accepted).
- F. Heide, M. O’Toole, K. Zang, D. B. Lindell, S. Diamond, G. Wetzstein, “Non-line-of-sight imaging with partial occluders and surface normals,” ACM Trans. Graph., 2019
- D. B. Lindell, G. Wetzstein, M. O’Toole, “Wave-based non-line-of-sight imaging using fast f-k migration,” ACM Trans. Graph. (SIGGRAPH), vol. 38, no. 4, 2019.
- F. Heide, S. Diamond, D. B. Lindell, G. Wetzstein, “Sub-picosecond photonefficient 3D imaging using single-photon sensors,” Scientific Reports, vol. 8, no. 17726, 2018.
- D. B. Lindell, M. O’Toole, G. Wetzstein, “Single-photon 3D imaging with deep sensor fusion,” ACM Trans. Graph. (SIGGRAPH), vol. 37, no. 4, 2018.
- M. O’Toole, D. B. Lindell, G. Wetzstein, “Confocal non-line-of-sight imaging based on the light cone transform,” Nature, vol. 555, no. 7696, p. 338, 2018.
- S. I. Young, D. B. Lindell, B. Girod, D. Taubman, G. Wetzstein, “Non-line-ofsight surface reconstruction using the directional light-cone transform,” in IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2020, (Oral).
- D. B. Lindell, M. O’Toole, G. Wetzstein, “Towards transient imaging at interactive rates with single-photon detectors,” in IEEE International Conference on Computational Photography (ICCP), 2018.
- M. O’Toole, F. Heide, D. B. Lindell, K. Zang, S. Diamond, G. Wetzstein, “Reconstructing transient images from single-photon sensors,” in IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017, (Spotlight).