Time-of-Flight Imaging & Single-Photon Imaging

lidar, non-line-of-sight imaging, ultrafast imaging

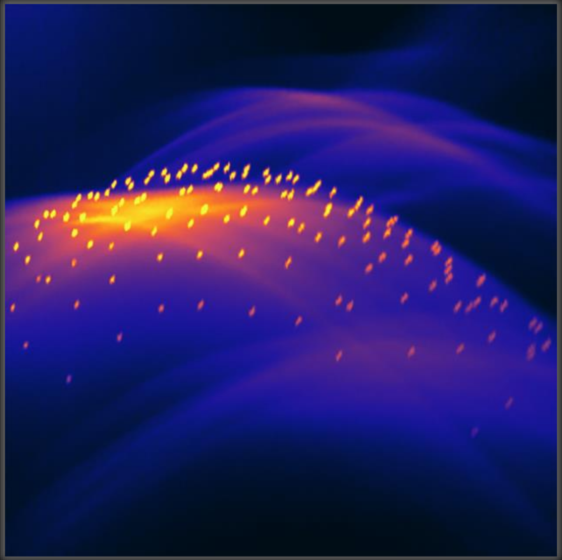

CSC2529

David Lindell

University of Toronto

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

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

Poster Session

- **Tuesday Nov 26 2-4 pm**
  - Bahen Atrium
  - you MUST attend in person or make alternative arrangements with me (only in extenuating circumstances)
- **You are responsible for making sure that your poster is printed on time!**
  - use the “same-day” or ”next-day” ordering option from https://utposter.com/
  - use the offer code on Quercus (“pages”) so you don’t have to pay
  - Pick up poster from 339 Bloor St W
  - bring it to the poster session and put it up before poster session begin

- You have from now until Nov 25 (the day before the poster session) to print your poster using this service.

Product Code: Next day service

Price: $3.99

Available Options

* **Shorter Side (24-60 inch):**

24

* **Longer Side (36-200 inch):**

36

* **Materials:**

High Quality Glossy Photo Paper

Tube (Mandatory for Paper Poster Shipping):

☐ Yes (+$3.99)

How Many Letter Size Handout Copies (Optional):

--- Please Select ---

* **Pick Up /Shipping Options:**

Pick Up @ store Mon-Sat 9am-6pm. Exclude Long Weekend Sat & Mon

Use the offer code on Quercus

Poster Session

- All proposals should have received feedback!
    - messaged you on Piazza

- Check Quercus->Assignments->Final Project for rubric for the poster session, code, report

Toronto Vision Seminar

UNIVERSITY OF TORONTO
FACULTY OF ARTS & SCIENCE

FIELDS

Mohammad Norouzi

Ideogram

"Diffusion Models and The Future of Creative Expression"

Wednesday, Nov 6th

2:00pm–3:00pm

BA5187

Reception to follow

Scan for Zoom link

Holographic displays

Manu Gopakumar
Ph.D. Candidate, Stanford

Tuesday Nov 12, 2-4 PM (Zoom)
(Lecture 11)

Single photon imaging

Sotiris Nousias
Postdoc, University of Toronto
Incoming Assistant Professor,
Purdue University

Tuesday Nov 19, 2-4 PM, here
(Lecture 11)

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
- Neural rendering for propagating light

echolocation

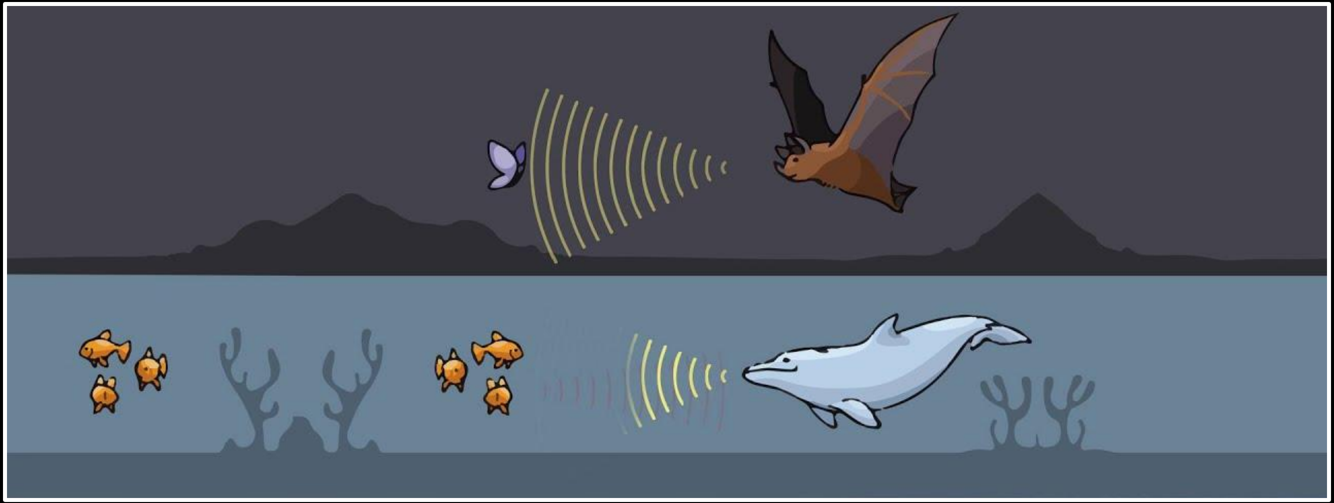

speed of sound in air: 343 meters / sec
in water: 1480 meters / sec


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

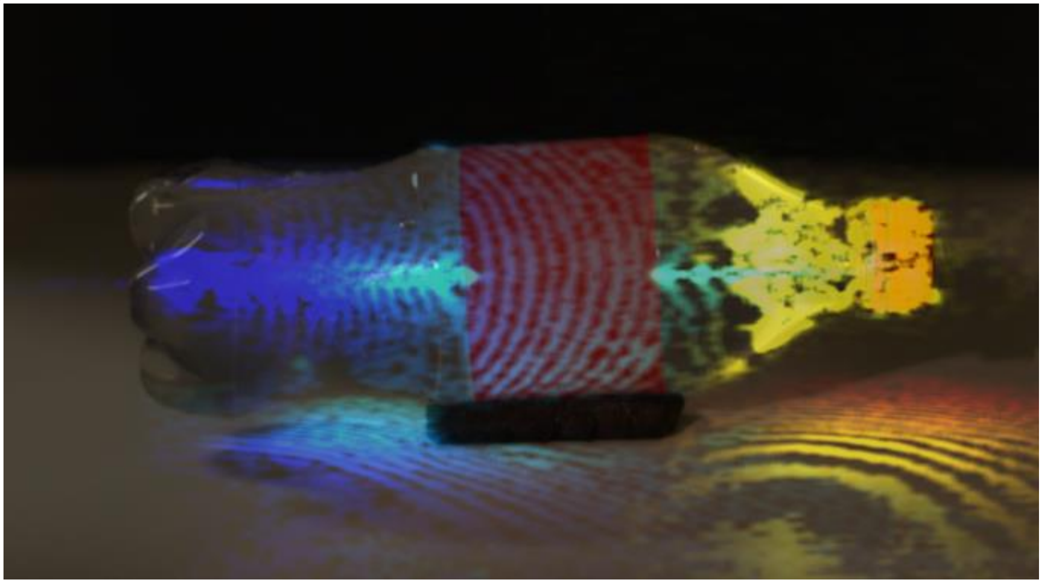

speed of light in a vacuum: **299,792,458** meters / sec

(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

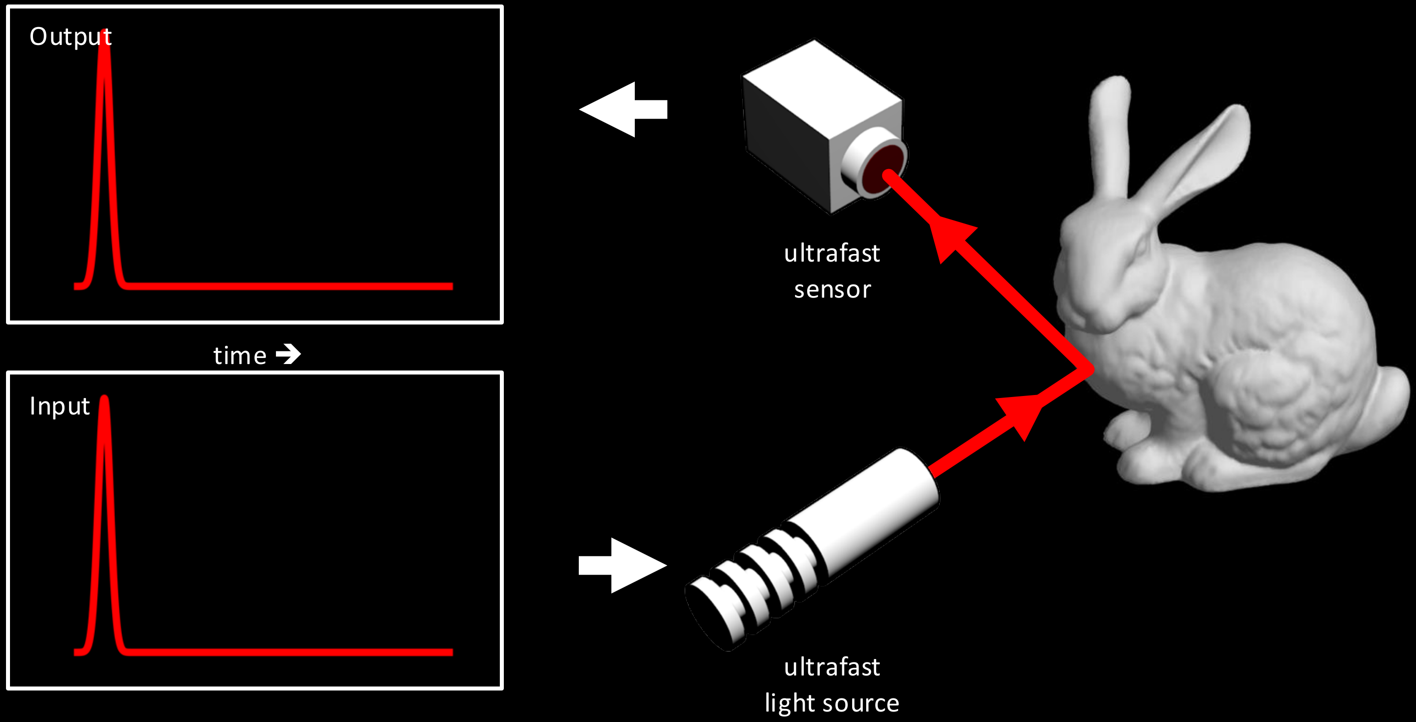

direct time-of-flight principle

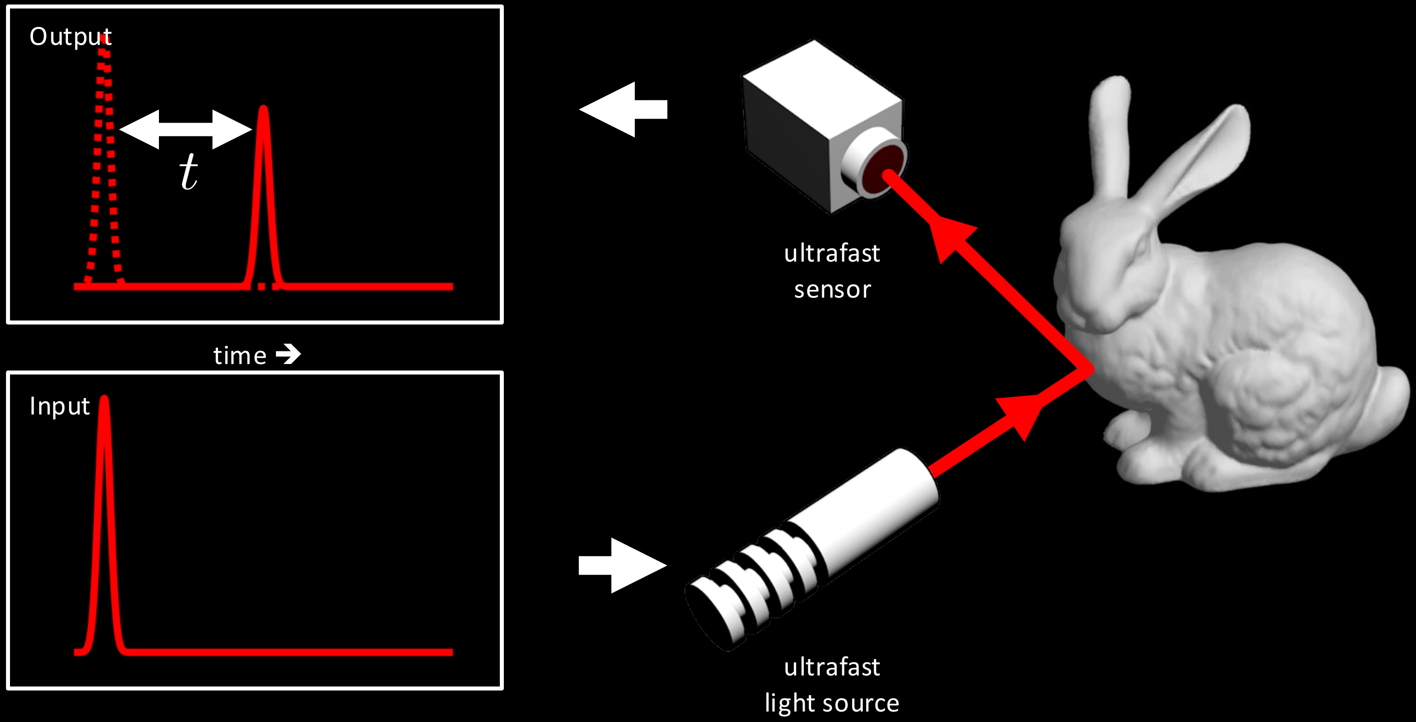

direct time-of-flight principle

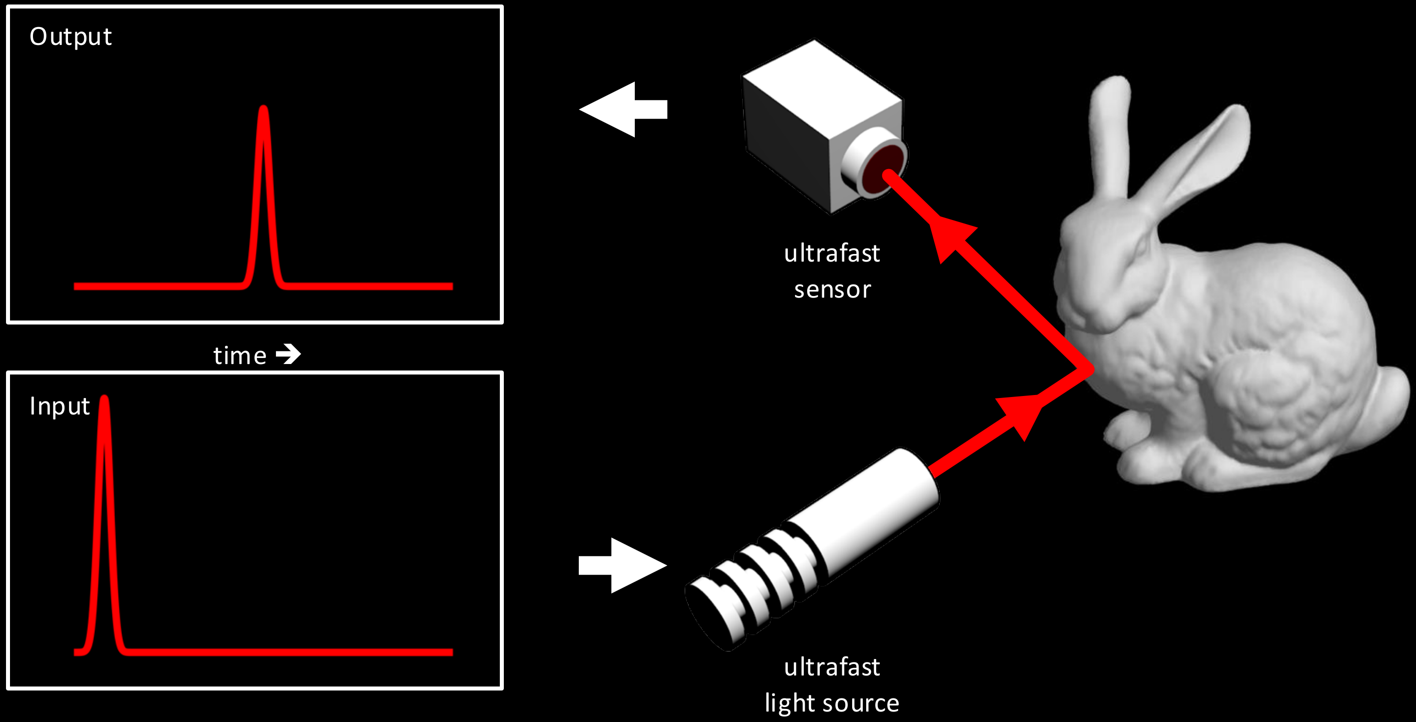

direct time-of-flight principle

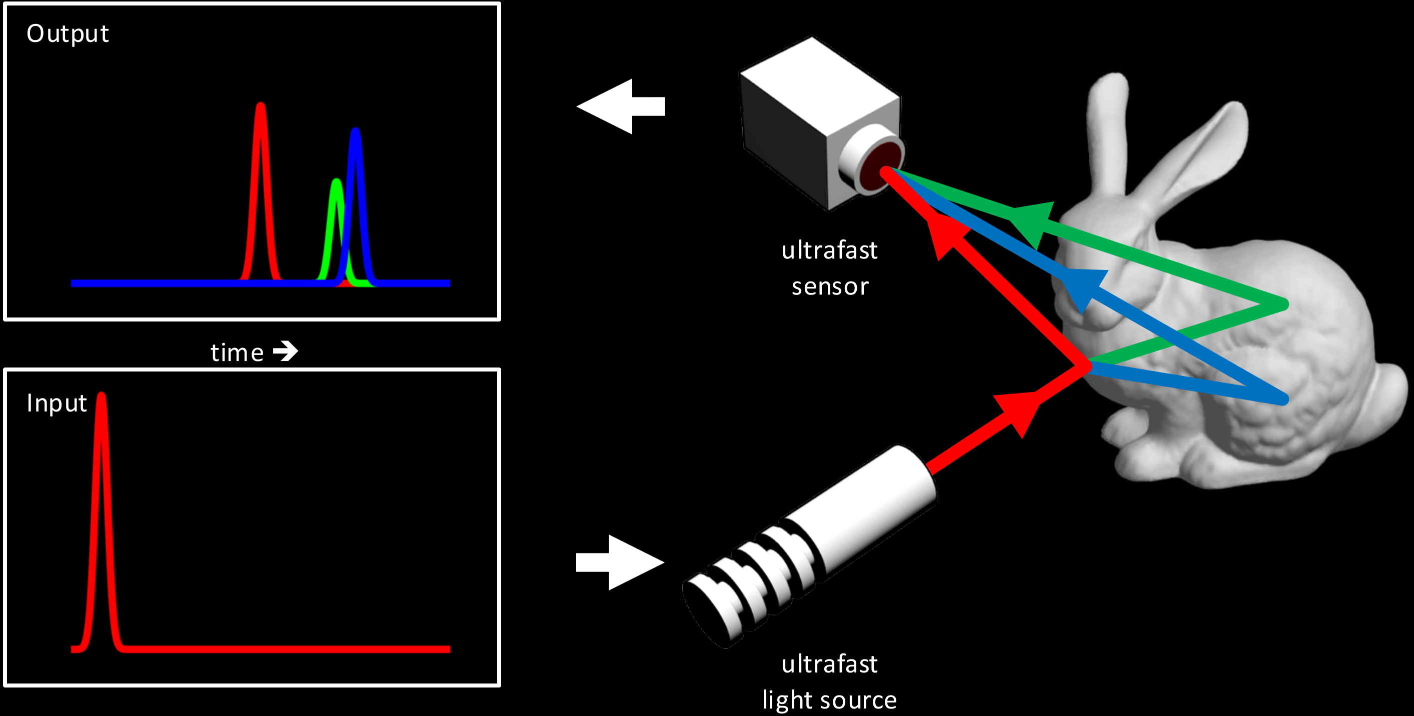

direct time-of-flight principle

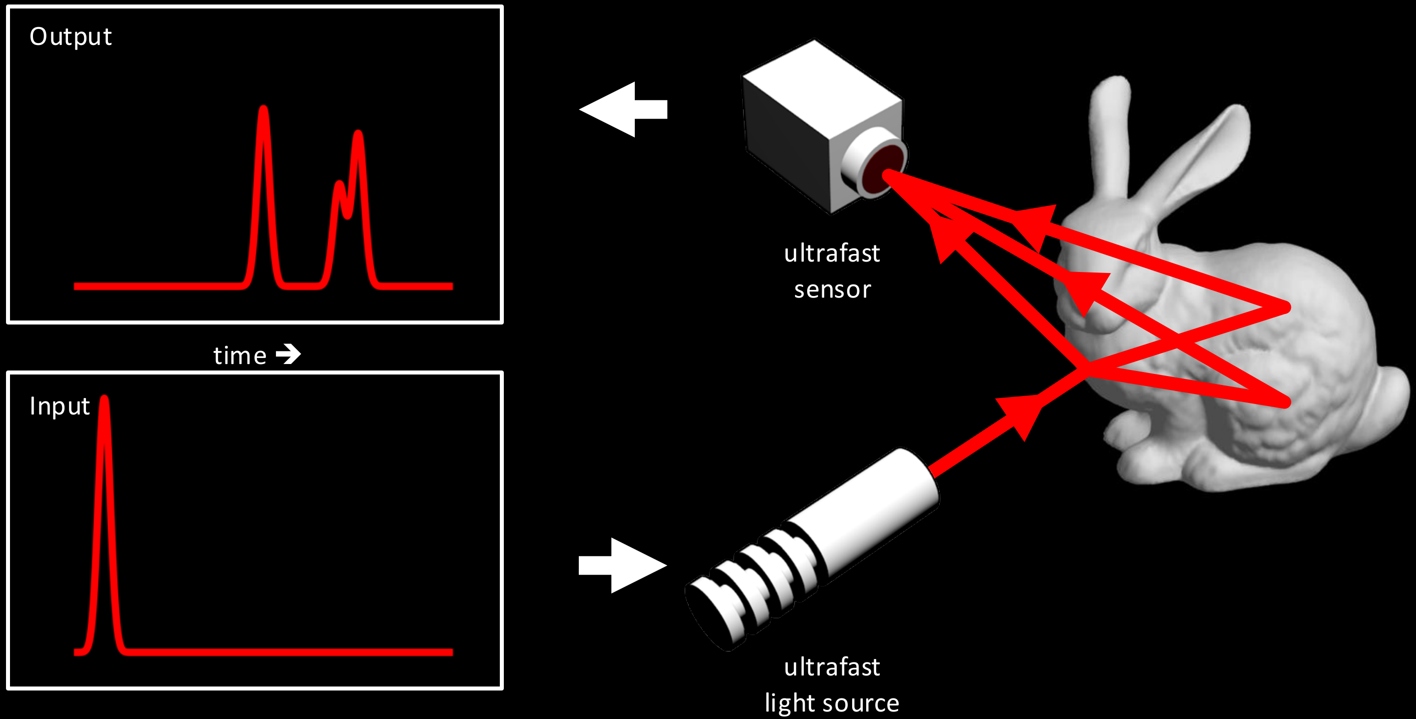

direct time-of-flight principle

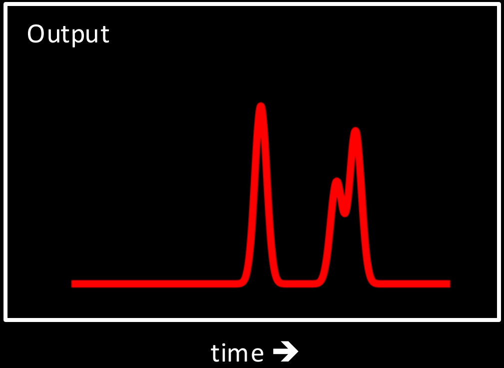

transient measurement

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

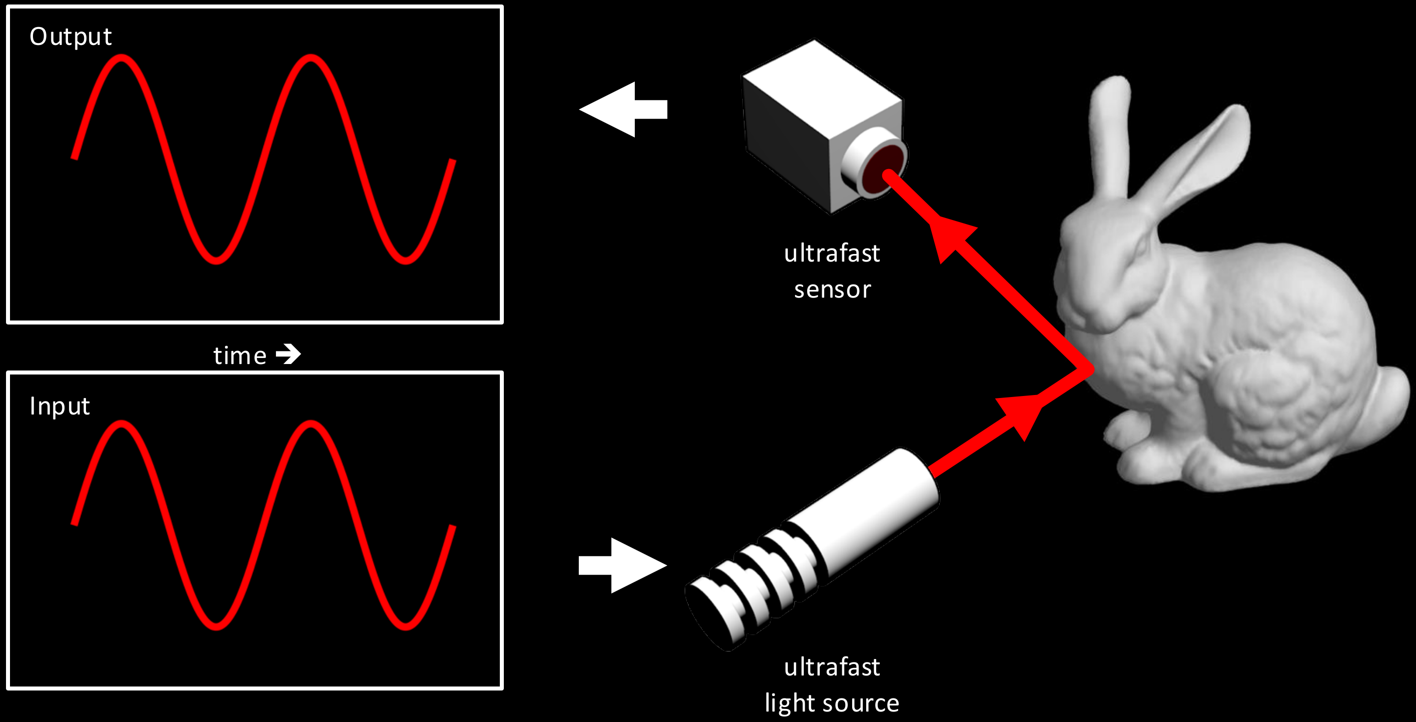

indirect time-of-flight principle

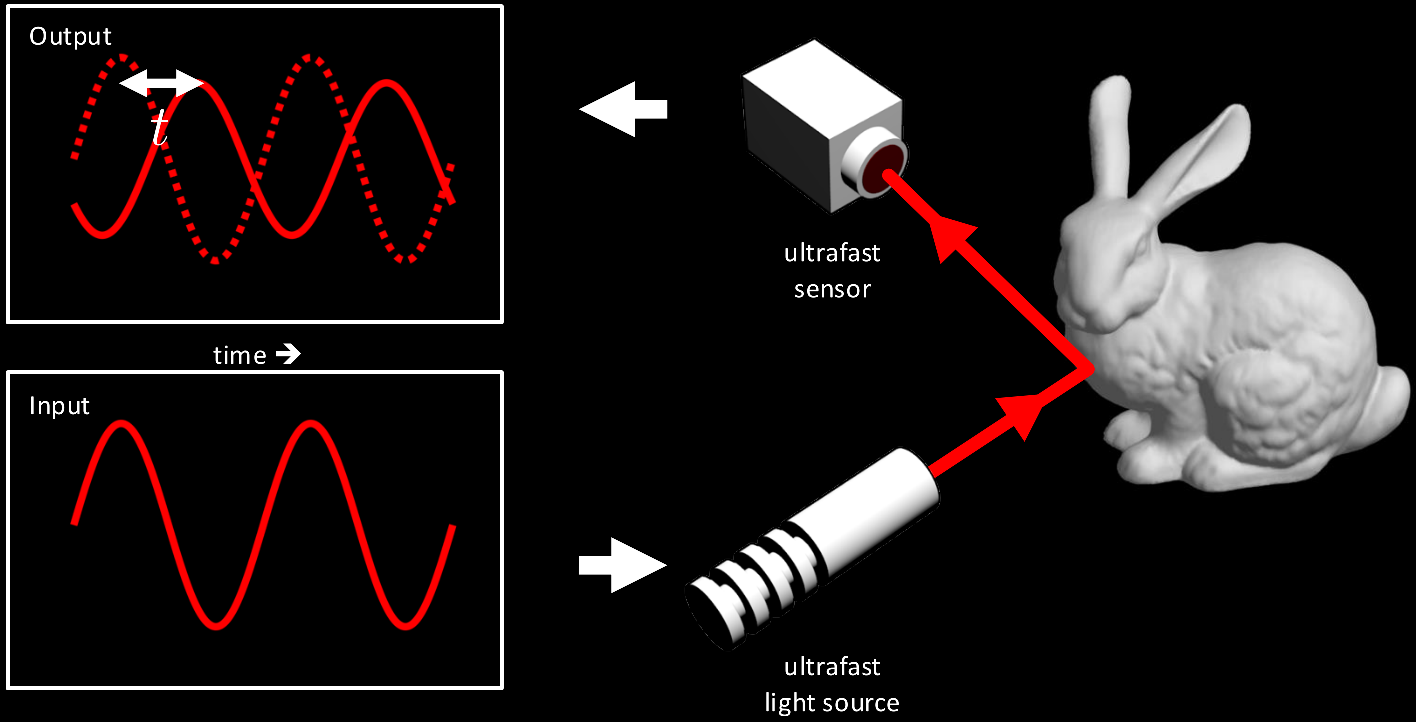

indirect time-of-flight principle

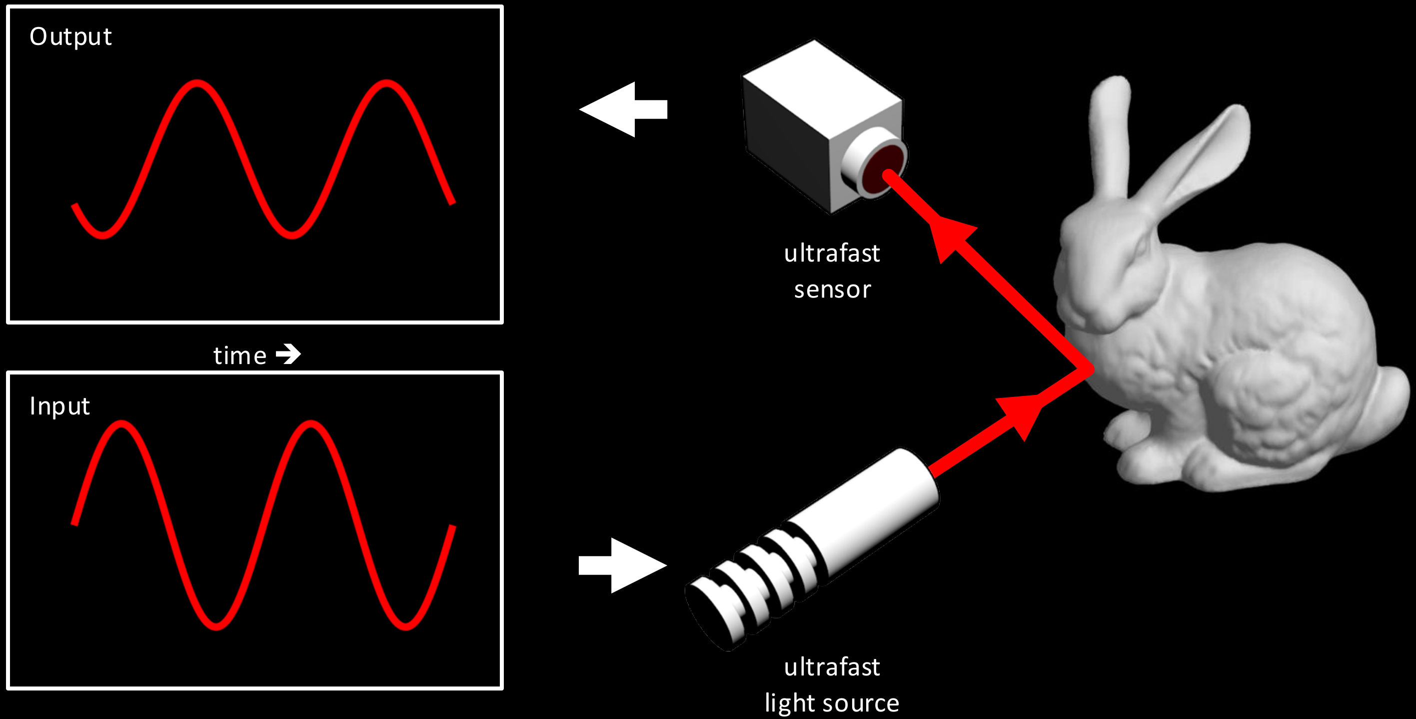

indirect time-of-flight principle

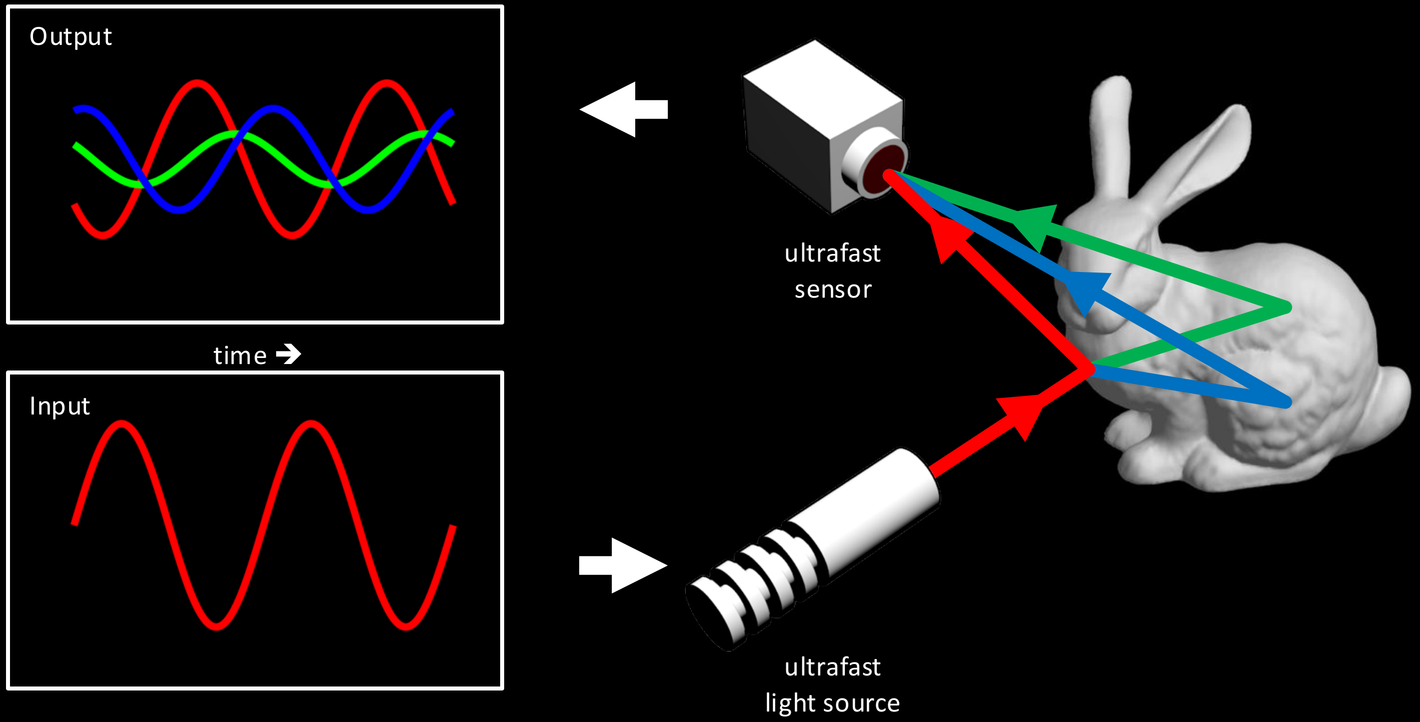

indirect time-of-flight principle

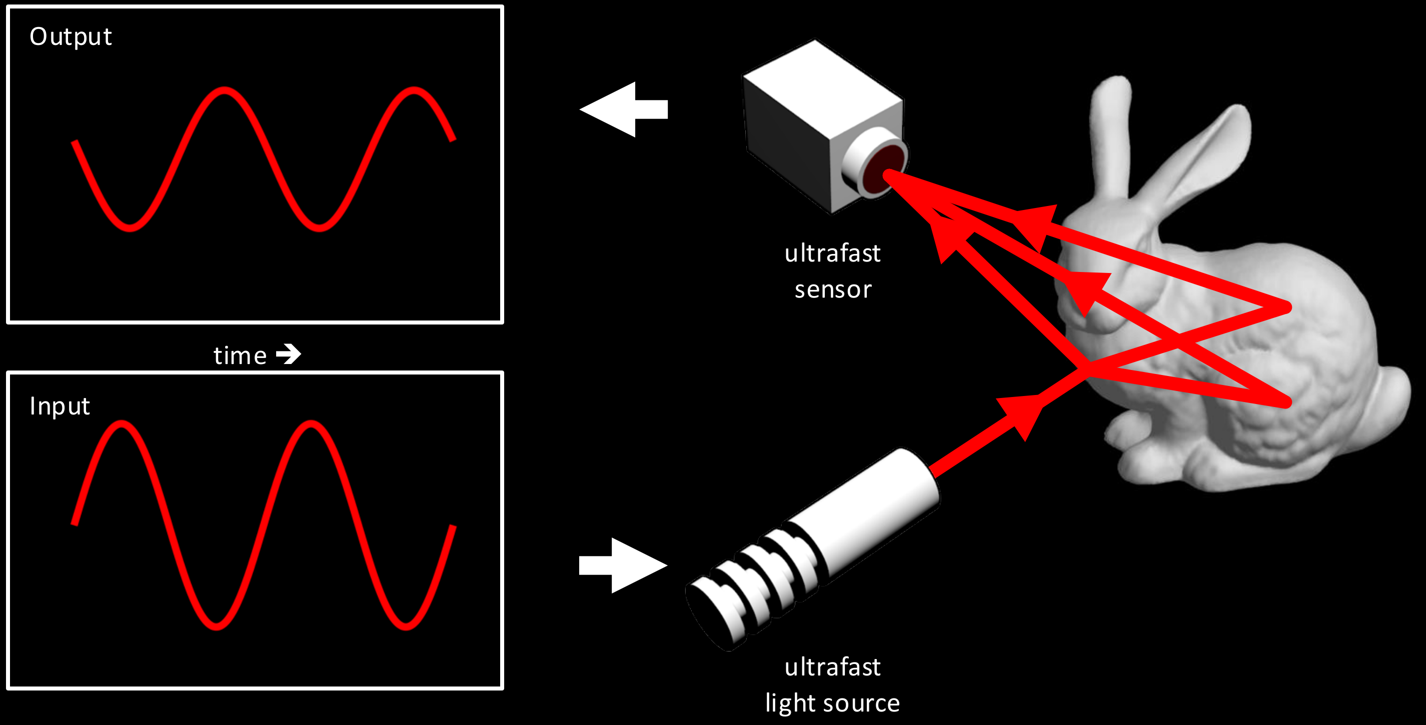

indirect time-of-flight principle

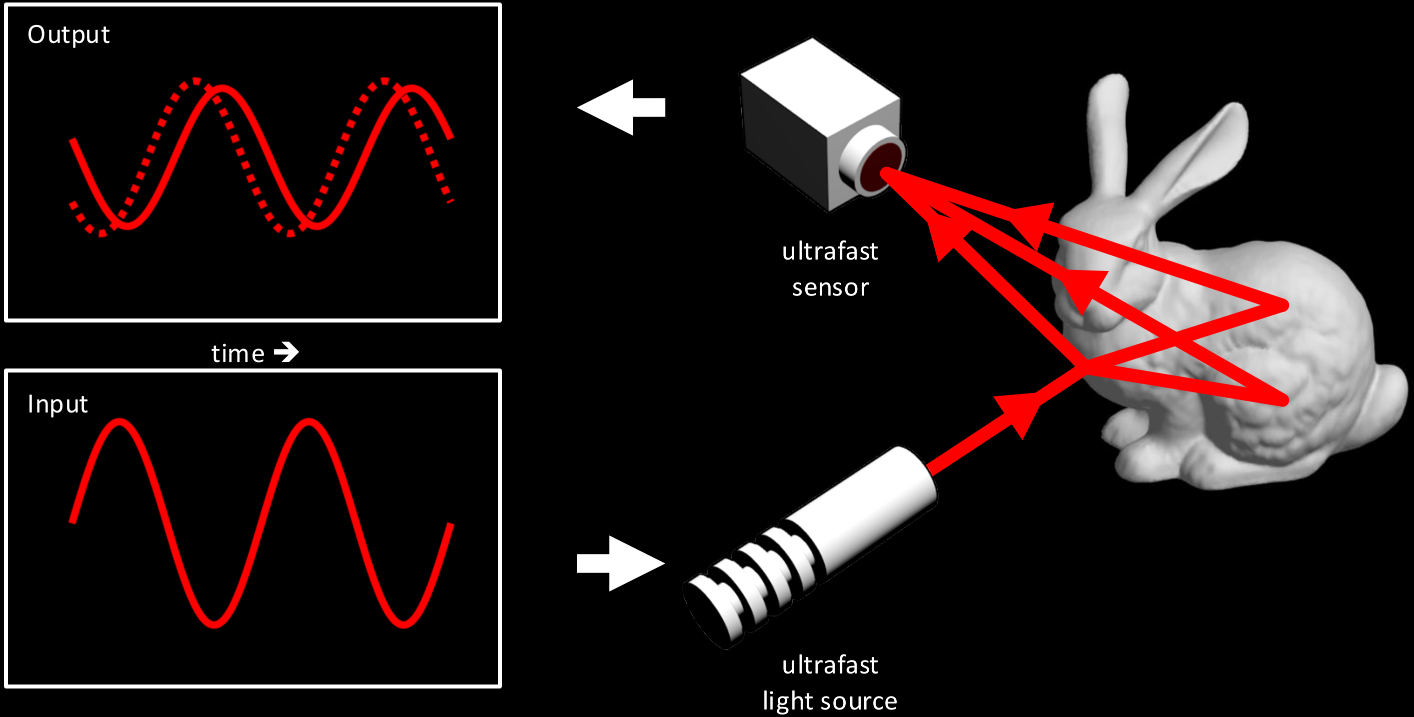

indirect time-of-flight principle

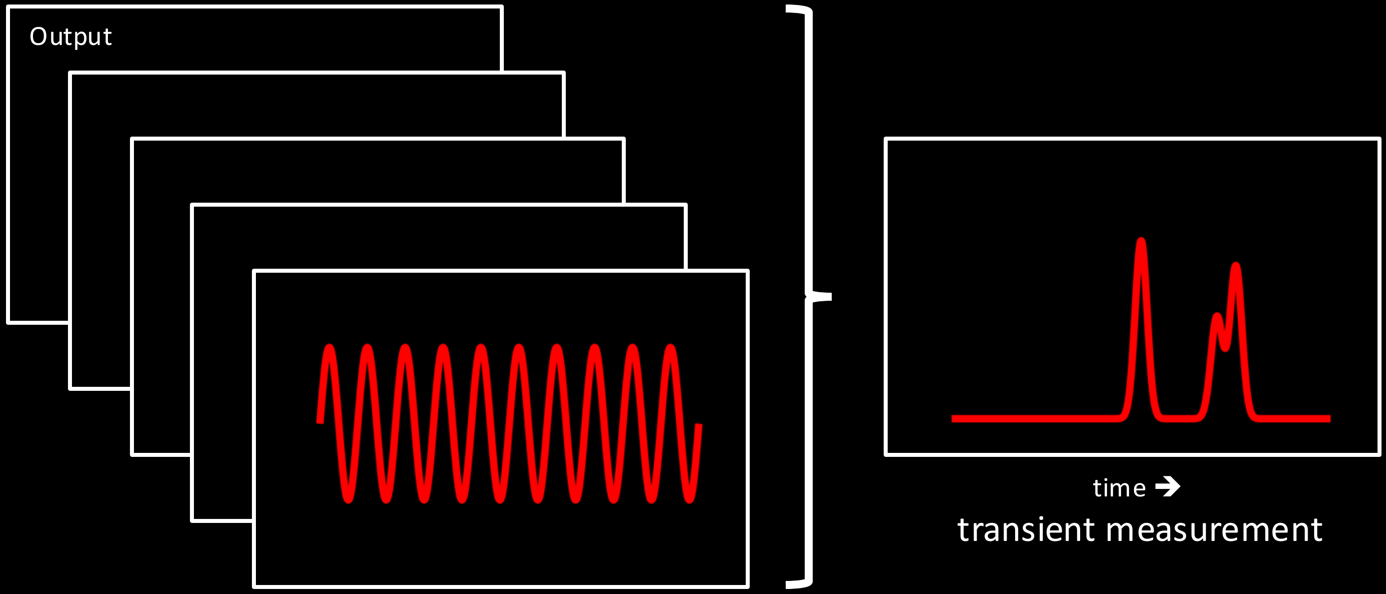

transient sensing technologies

	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)

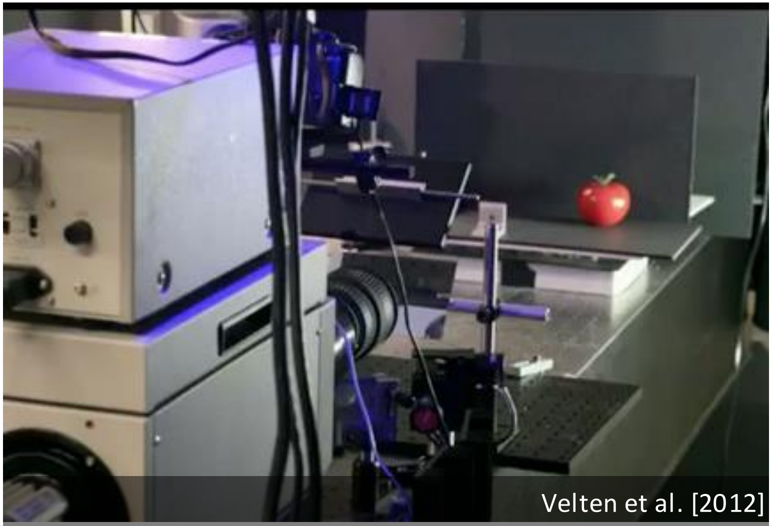

Velten et al. [2012]

transient sensing technologies

	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)

Micro Photon Devices

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)

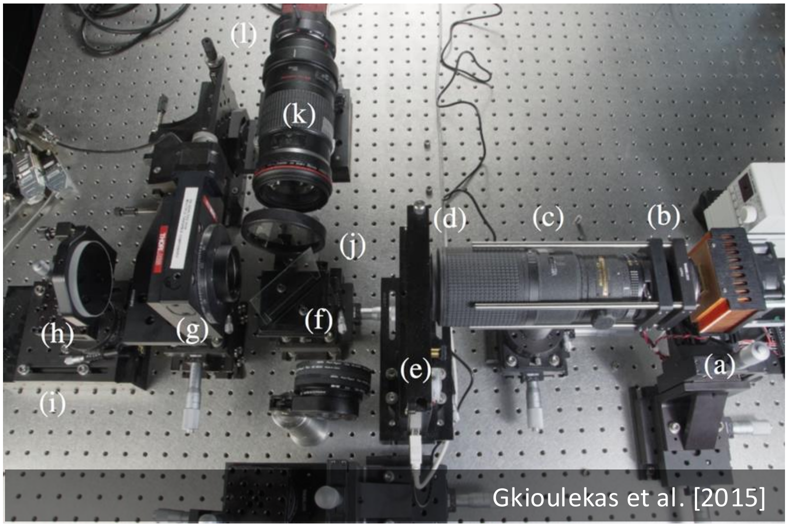

Gkioulekas et al. [2015]

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

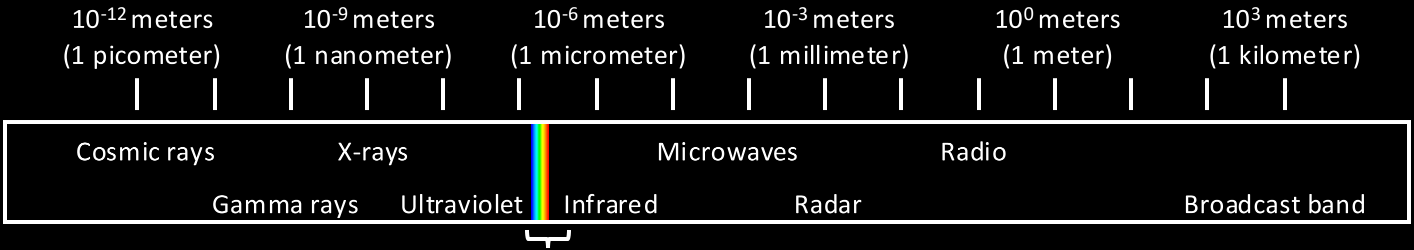

- **optical coherence tomography**
- **streak camera**
- **single-photon avalanche diodes**
- **time-of-flight cameras**
- **avalanche photodiode**

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 | Gamma rays | X-rays | Ultraviolet | Infrared | Microwaves | Radar | Radio | Broadcast band

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

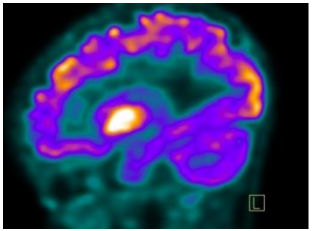

Positron Emission Tomography (PET) Scan

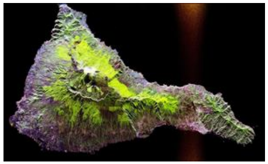

Synthetic Aperture Radar (SAR)

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

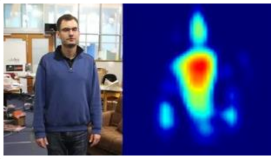

Seeing through Walls [Adib et al., 2015]

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

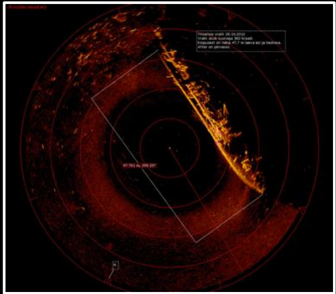

SONAR
(Sound Navigation and Ranging)

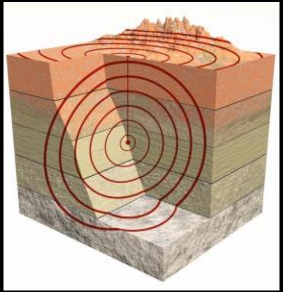

Seismic Imaging

overview

- Time-resolved imaging
- Single-photon avalanche diodes (SPADs)
- Single-photon lidar
- Non-line-of-sight imaging
- neural rendering for propagating light

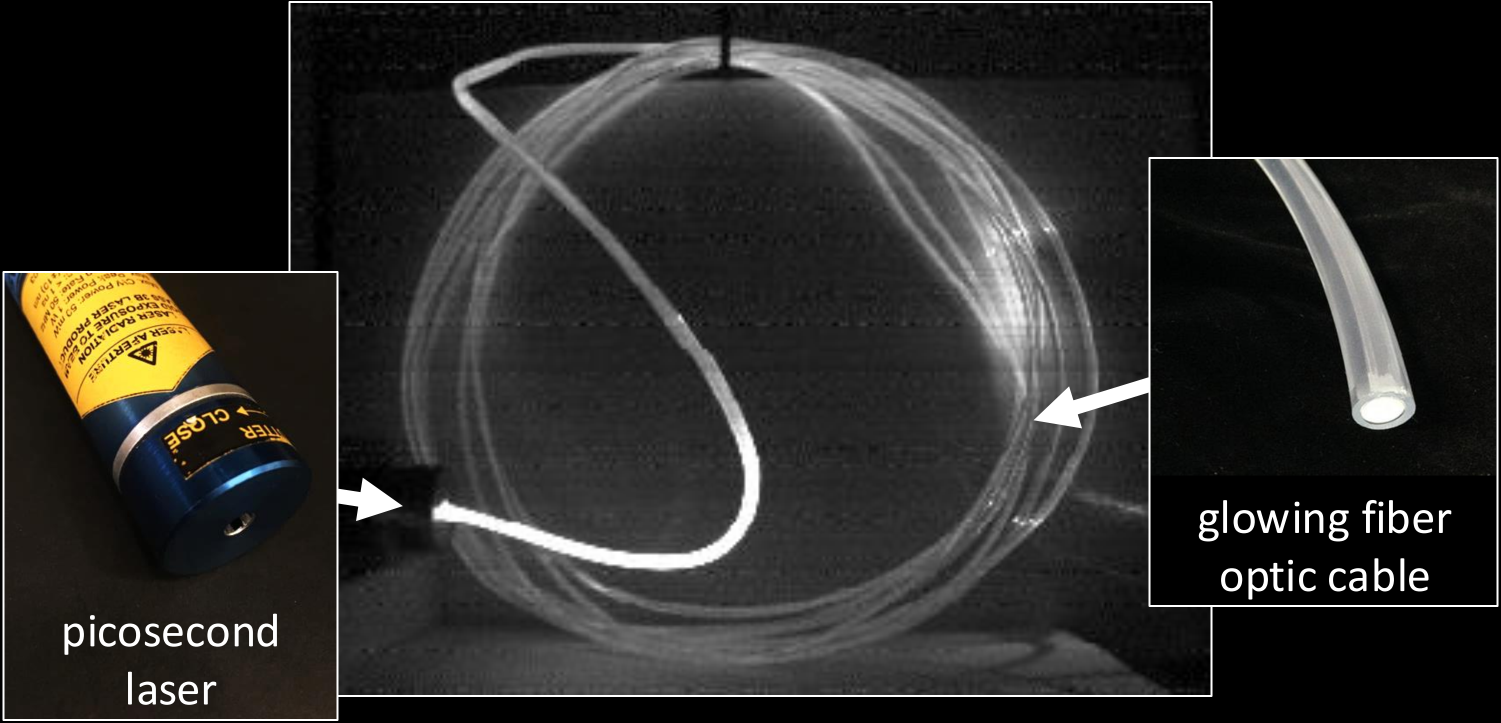

regular image

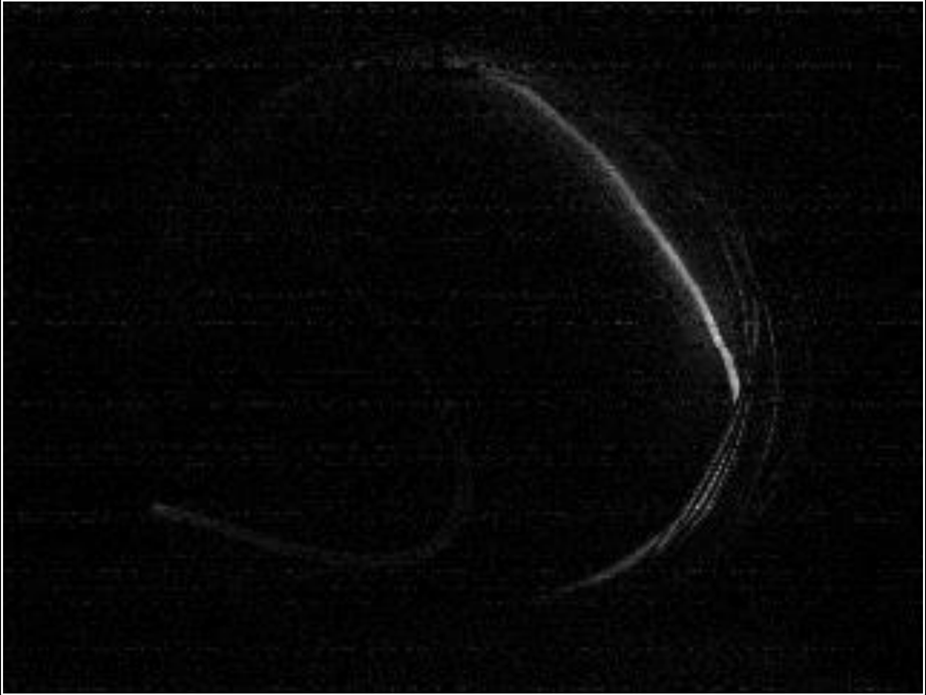

transient image

Applications

Autonomous Navigation
Image by Wikimedia Commons

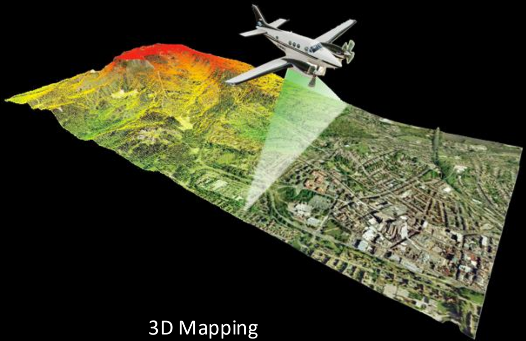

3D Mapping
Image by LIDAR-America

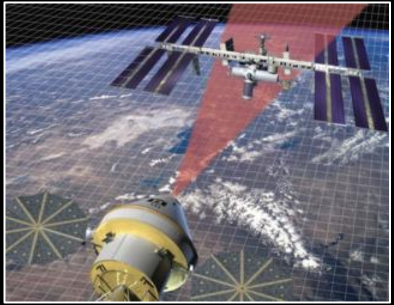

Space Station Docking
Image by NASA

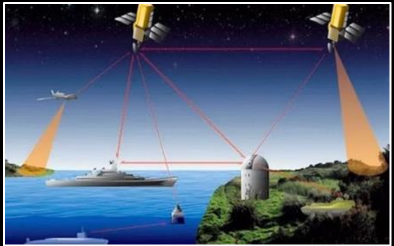

Optical Communications
Image by Siasat Daily

Biomedical Imaging
Image by Washington University

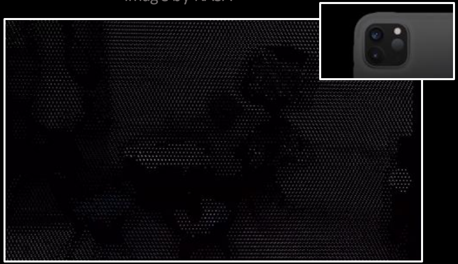

Consumer Electronics (2020 iPad Pro)
Video by Tim Fields

Appl

Au
Ima

ng

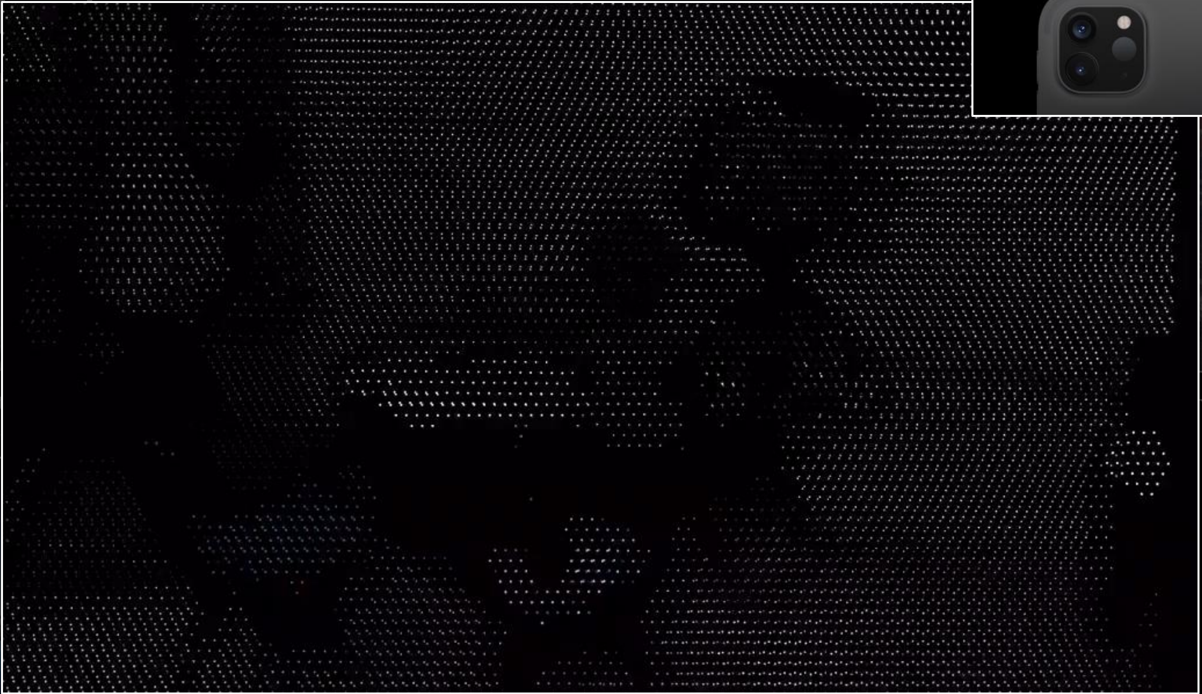

Optical Communications
Image by Siasat Daily

Biomedical Imaging
Image by Washington University

Consumer Electronics (2020 iPad Pro)
Video by Tim Fields

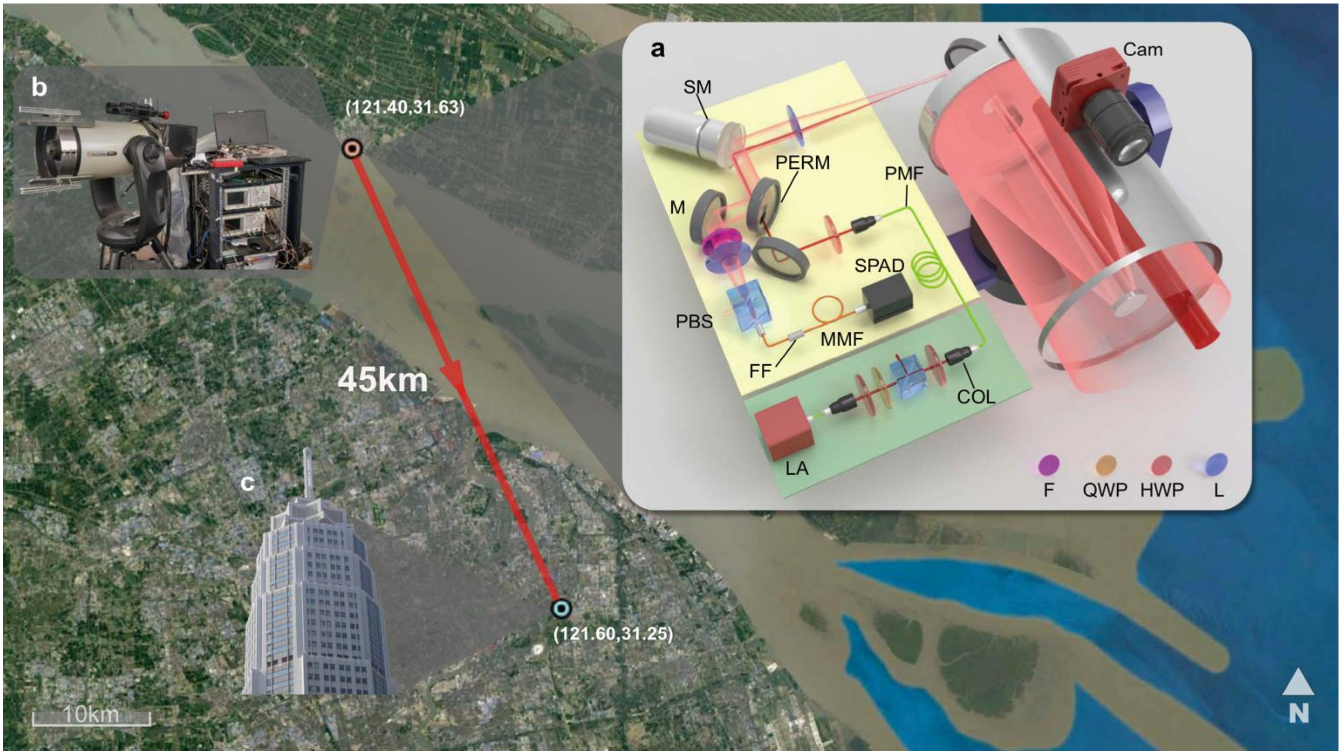
a
SM
Cam
PERM
PMF
M
SPAD
PBS
MMF
FF
COL
LA
F
QWP
HWP
L
b
(121.40,31.63)
45km
c
(121.60,31.25)
10km
N

International Laser Ranging Service (ILRS)

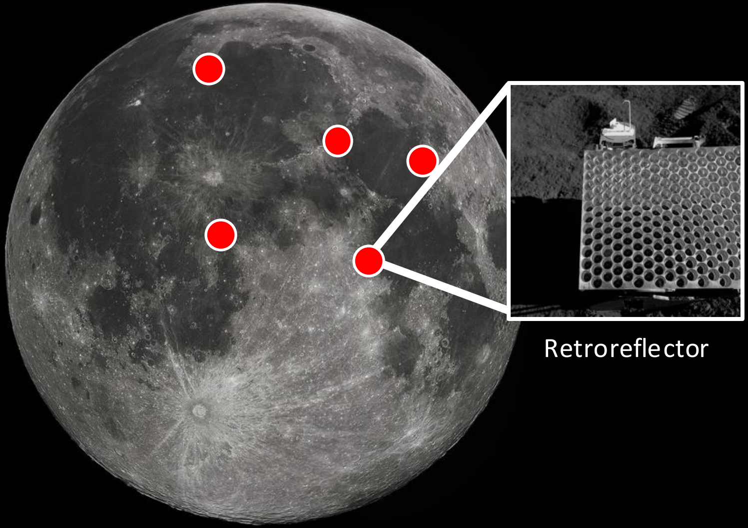

Lunar Laser Ranging (LLR)

● - Location of Lunar Retroreflector

single-photon avalanche diode (SPAD)

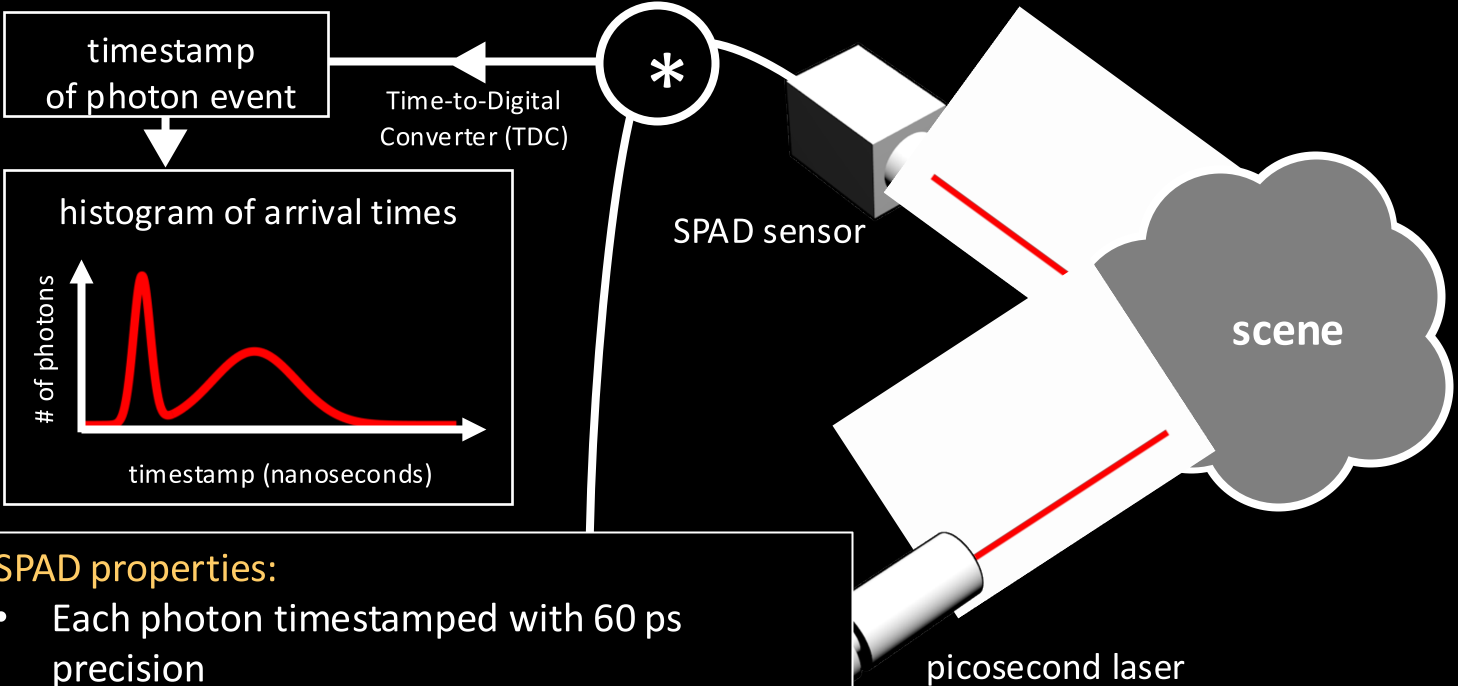

SPAD properties:

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

linear array of SPADs

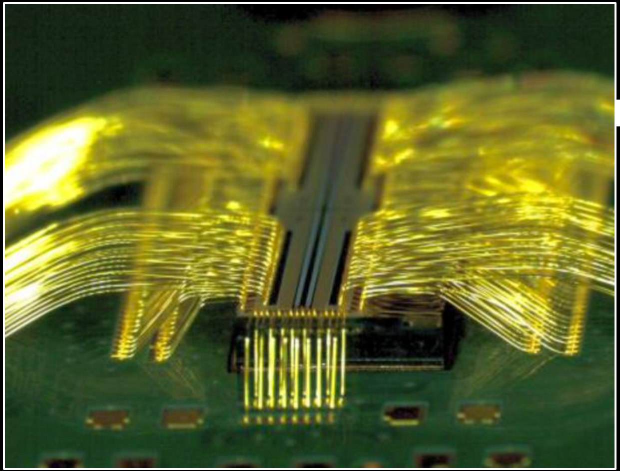

LinoSPAD from FastTree 3D

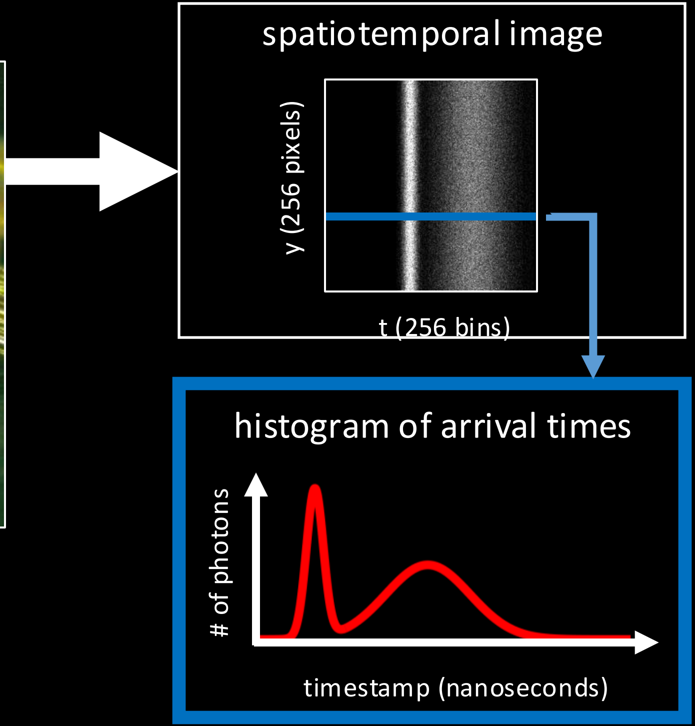

scanning procedure

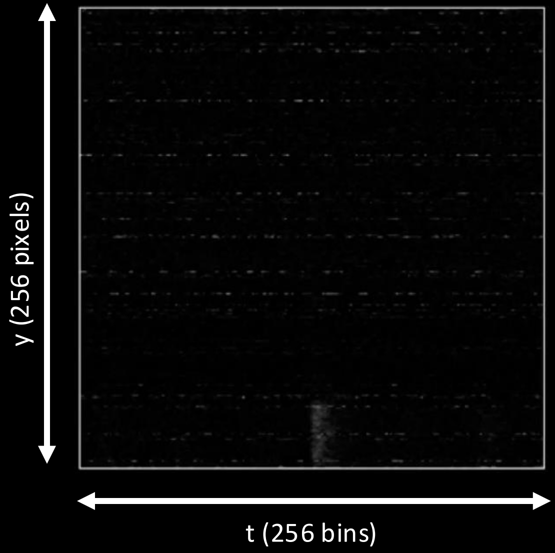

SPAD output

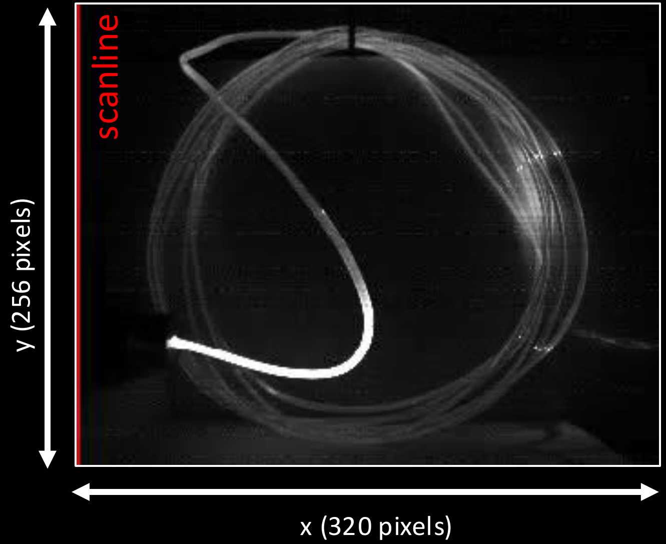

regular image

scanning procedure

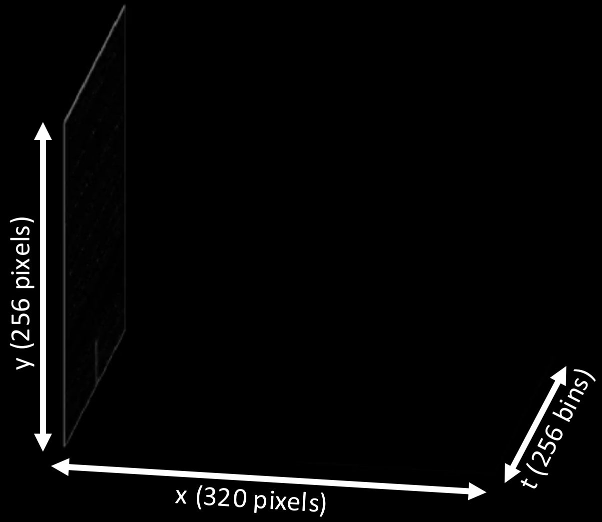

transient image

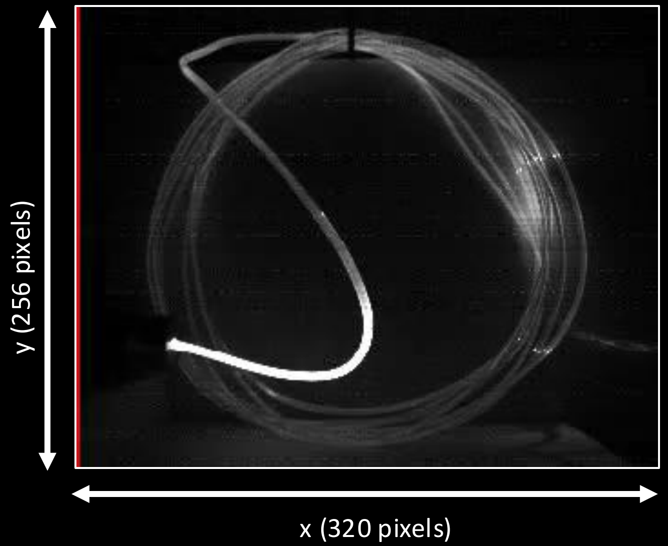

regular image

scanning procedure

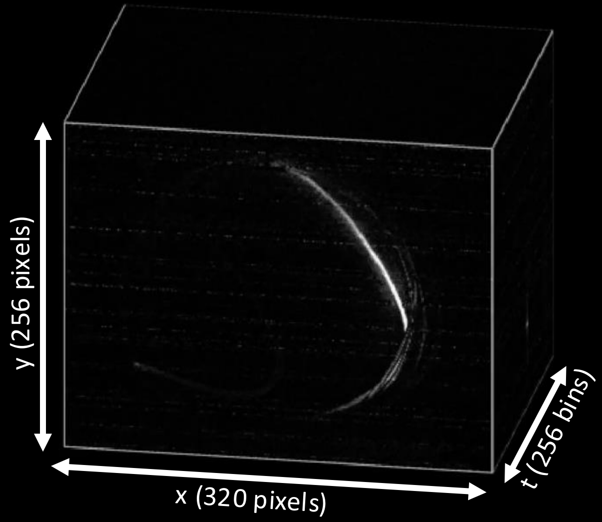

transient image

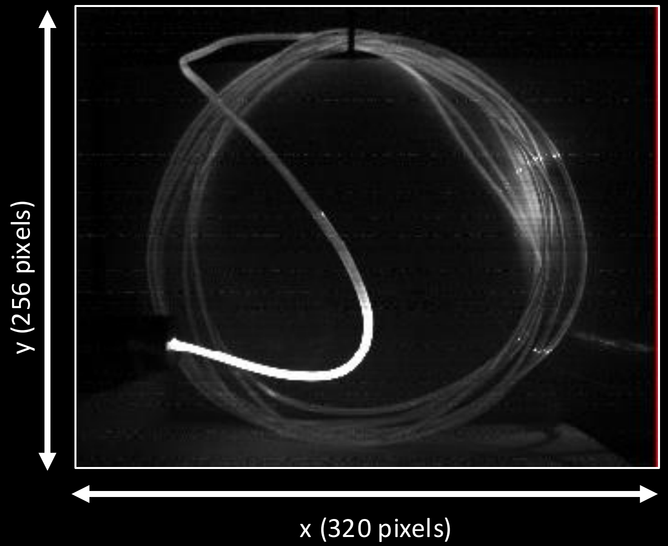

regular image

overview

- Time-resolved imaging
- Single-photon avalanche diodes (SPADs)
- Single-photon lidar
- Non-line-of-sight imaging
- passive ultra-wideband sensing

LiDAR

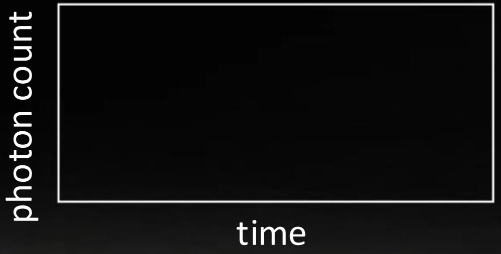

LiDAR

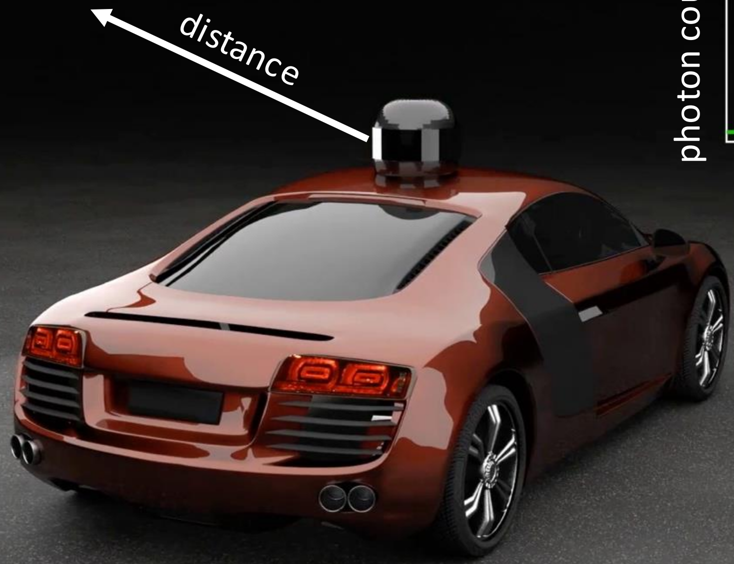

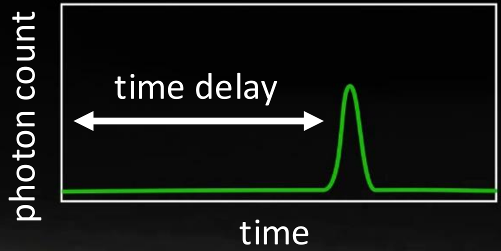

LiDAR

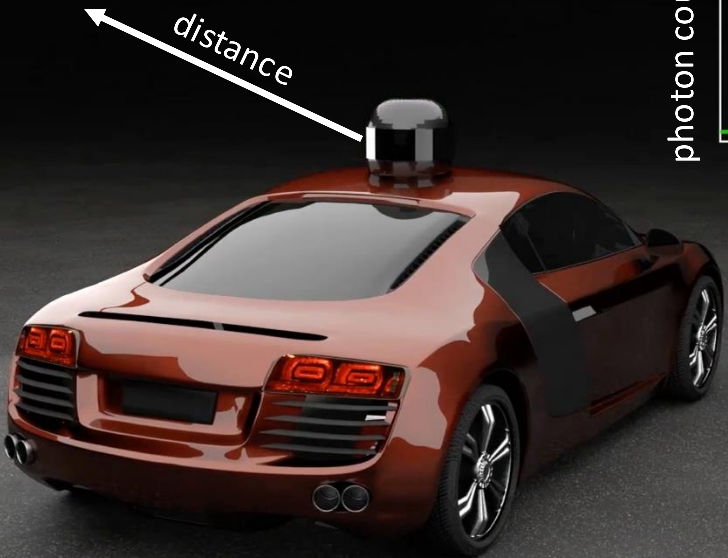

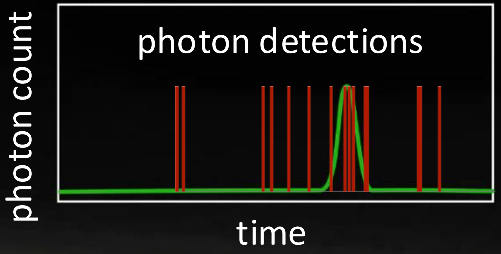

LiDAR

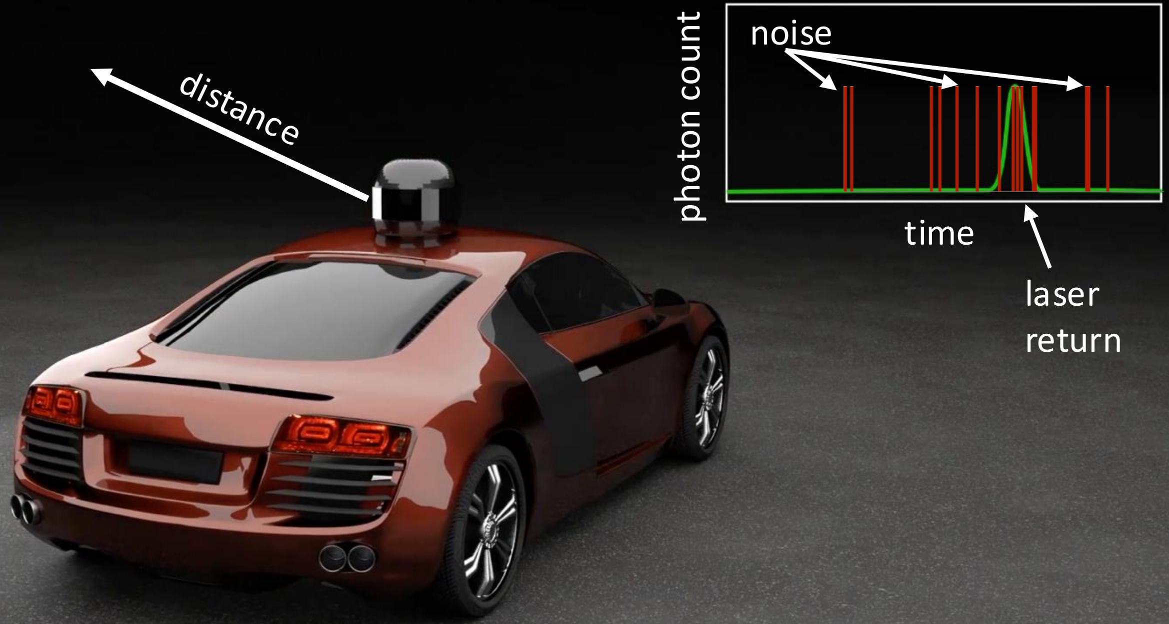

LiDAR

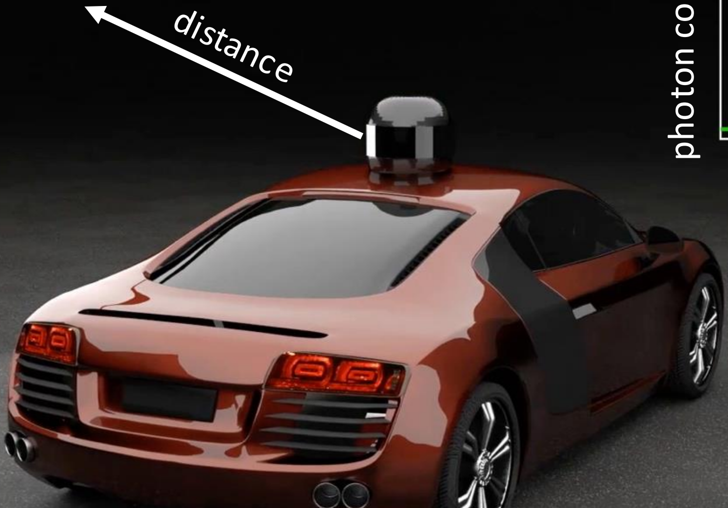

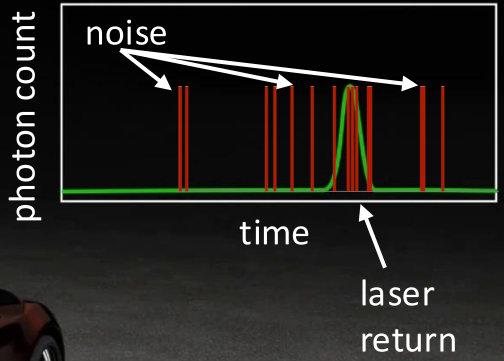

Challenges

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

Linear mode (i.e., avalanche photodiode or ***APD***):

acts like a conventional photodiode with extremely high gain or amplification

time resolution >300 ps – 10 ns

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

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

time resolution ~50 ps

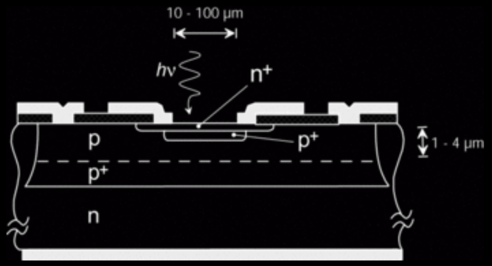

Semiconductor devices

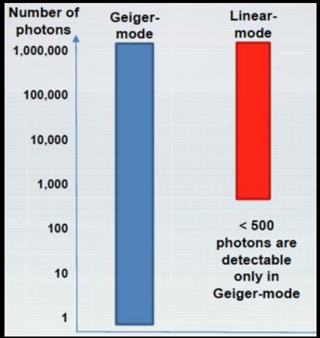

image by Princeton Lightwave

Single-Photon Avalanche Diodes

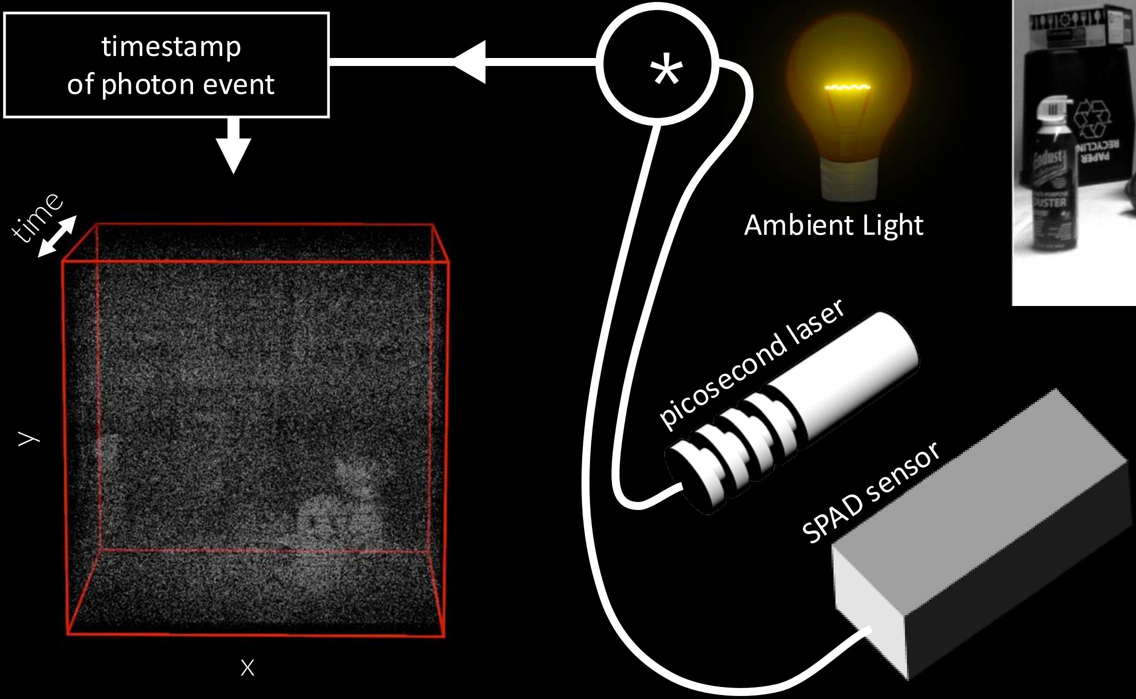

CNN Architecture for Depth Estimation

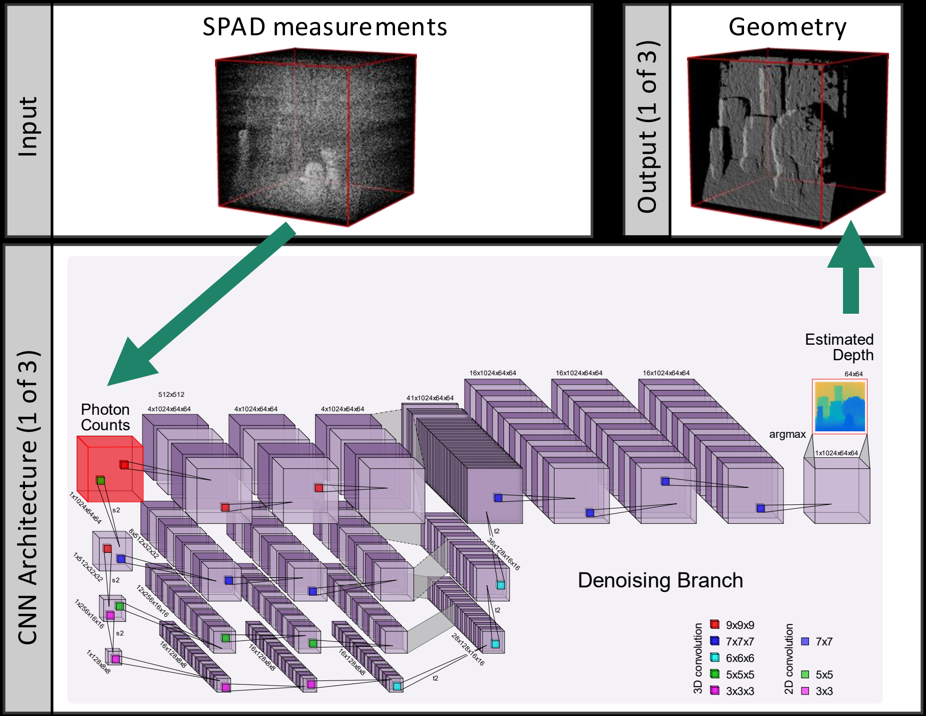

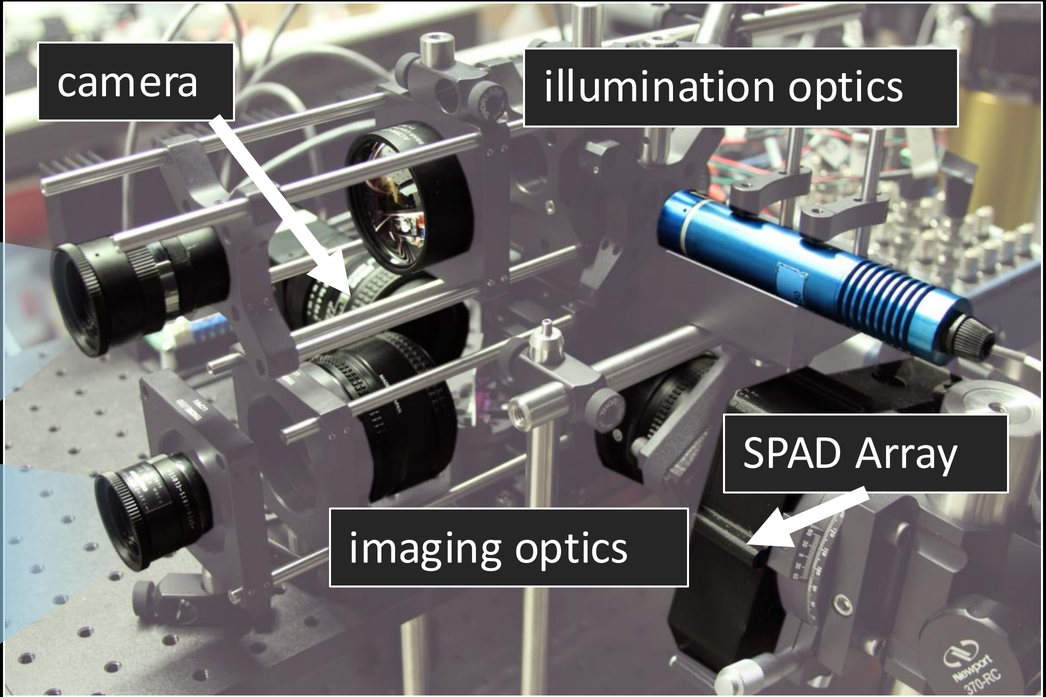

vertical scanline

Lindell et al., SIGGRAPH 2018

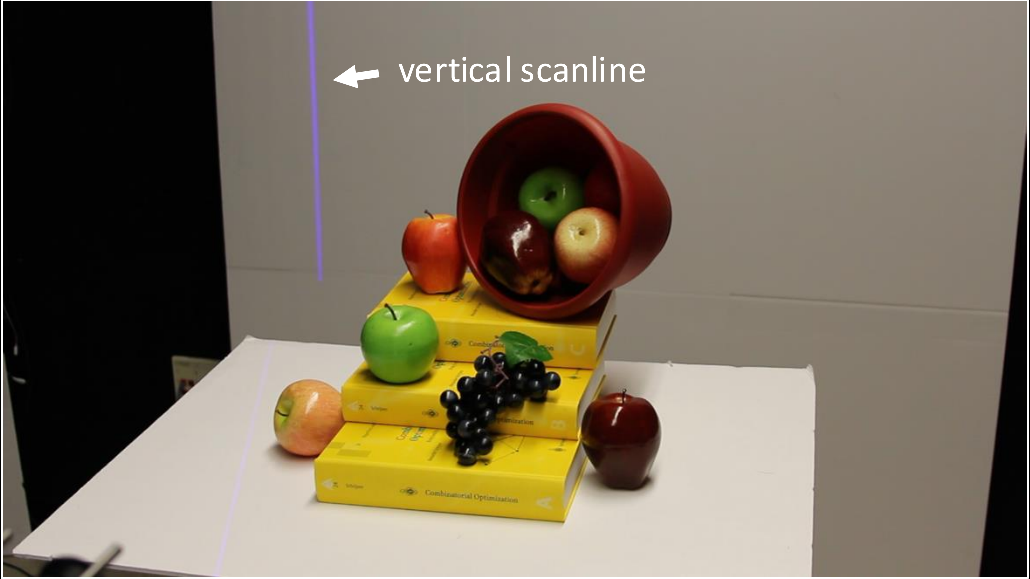

scan rate: 20 Hz

lights on

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

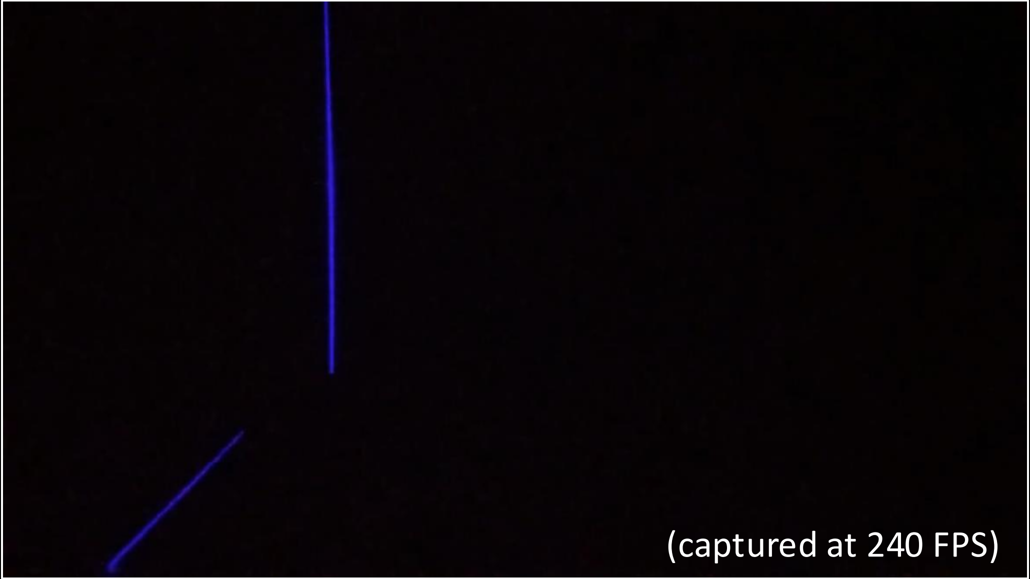

scan rate: 20 Hz

lights off

Lindell et al., SIGGRAPH 2018

Intensity image

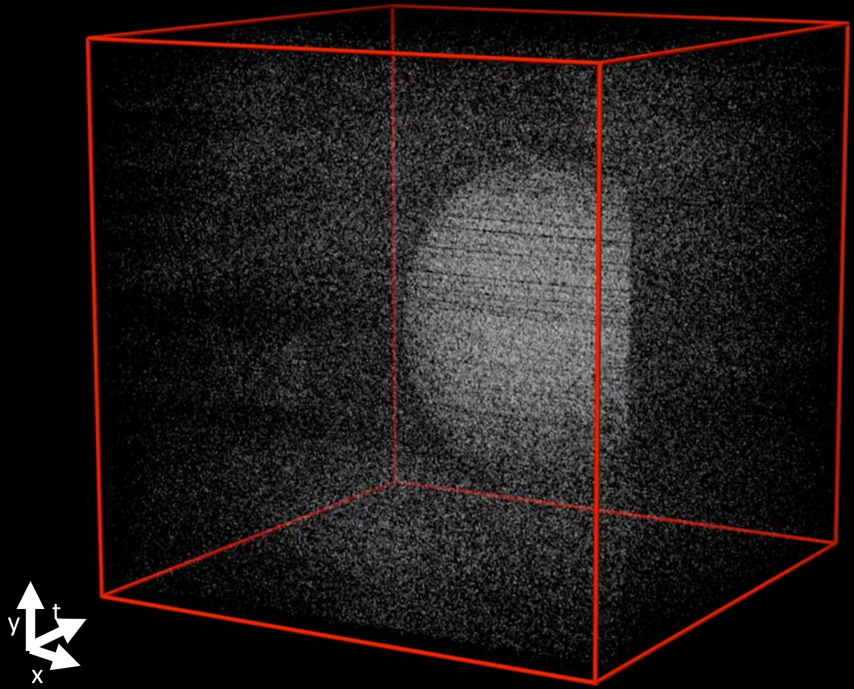

SPAD measurements (20 Hz)

Average per spatial position
0.64 Signal Detections
0.87 Background Detections

Lindell et al., SIGGRAPH 2018

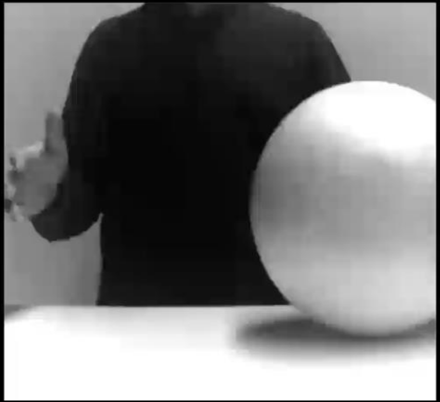

Intensity image

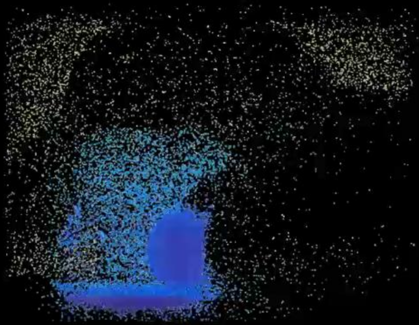

Log-matched filter

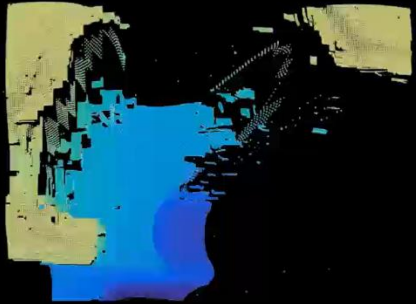

[Rapp and Goyal 2017]

CASPI: collaborative photon processing for active single-photon imaging

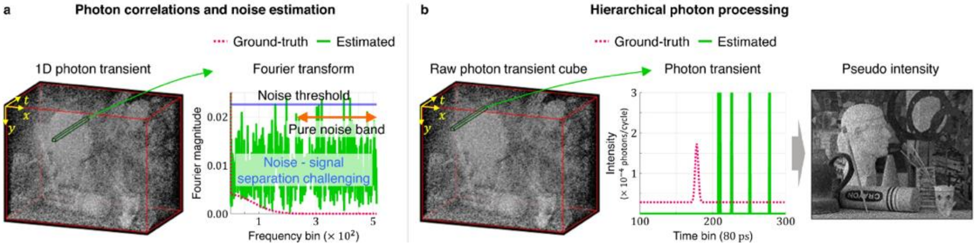

CASPI: collaborative photon processing for active single-photon imaging

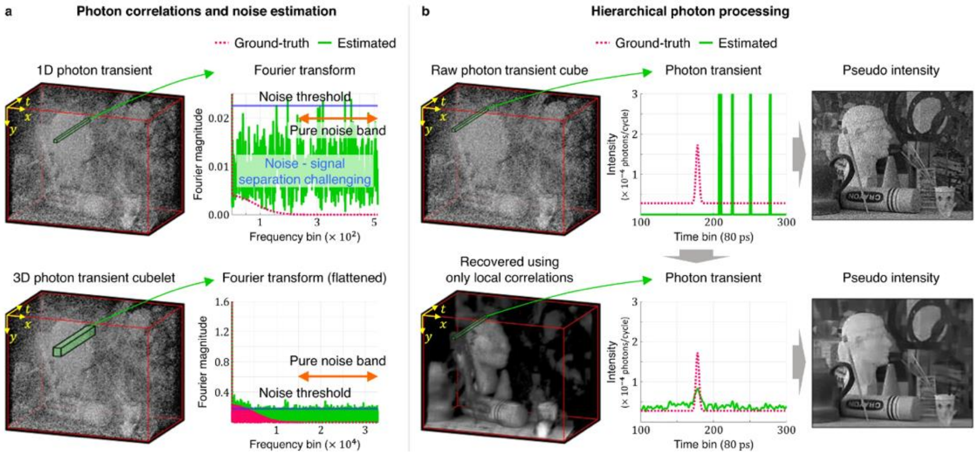

CASPI: collaborative photon processing for active single-photon imaging

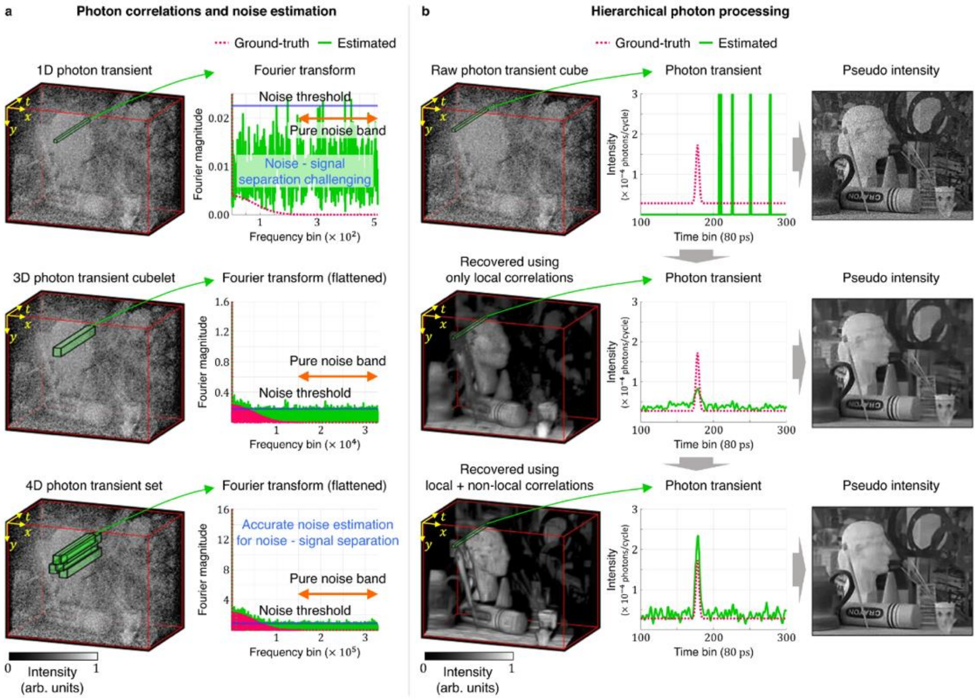

CASPI: collaborative photon processing for active single-photon imaging

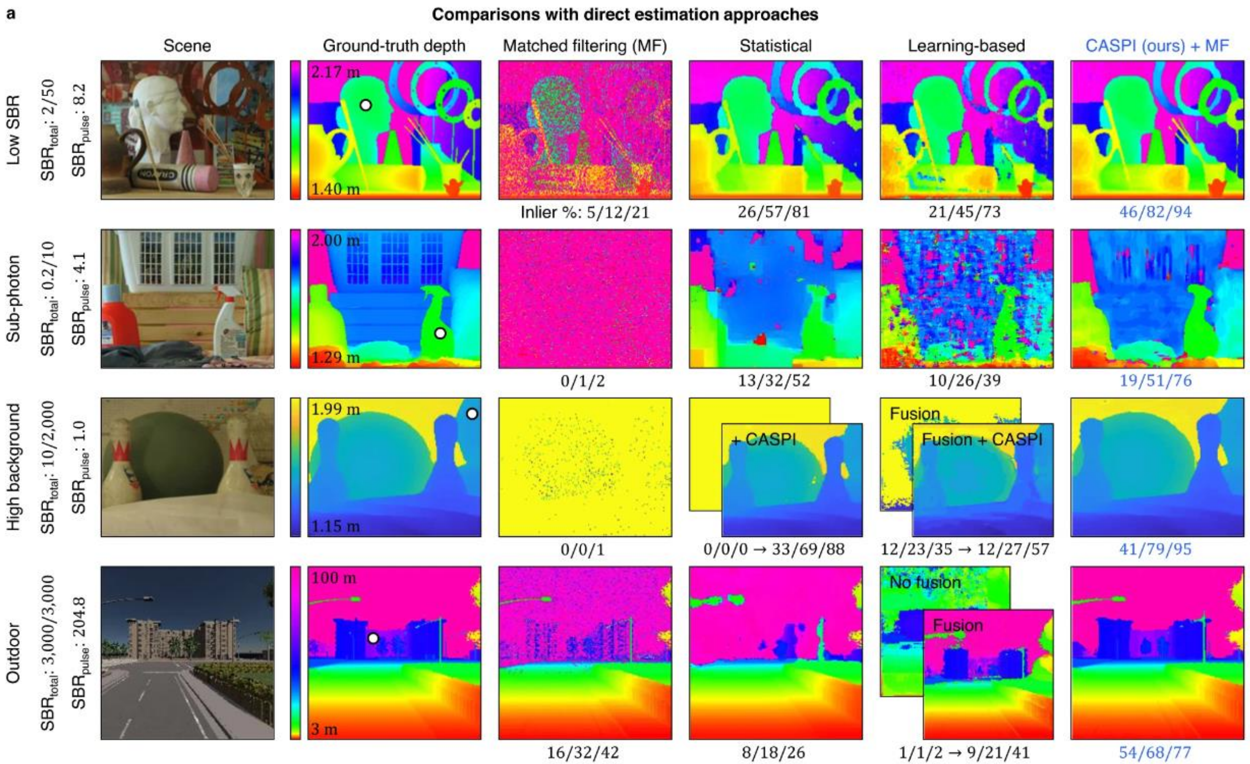

CASPI: collaborative photon processing for active single-photon imaging

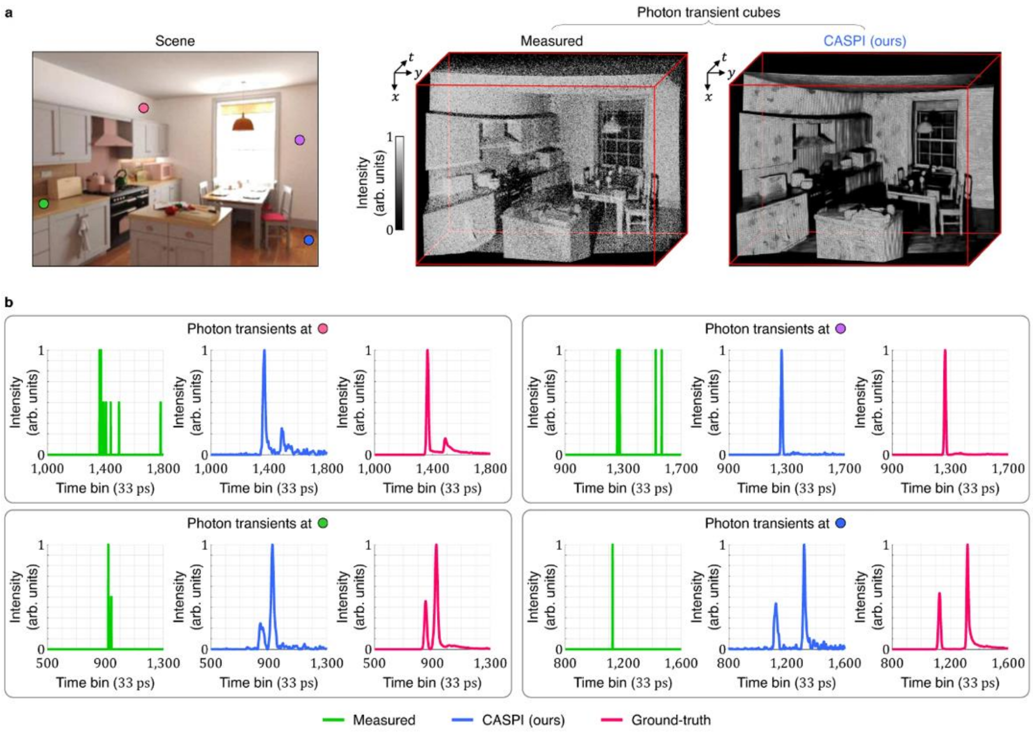

overview

- Time-resolved imaging
- Single-photon avalanche diodes (SPADs)
- Single-photon lidar
- Non-line-of-sight imaging (part 1)
- neural rendering for propagating light

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

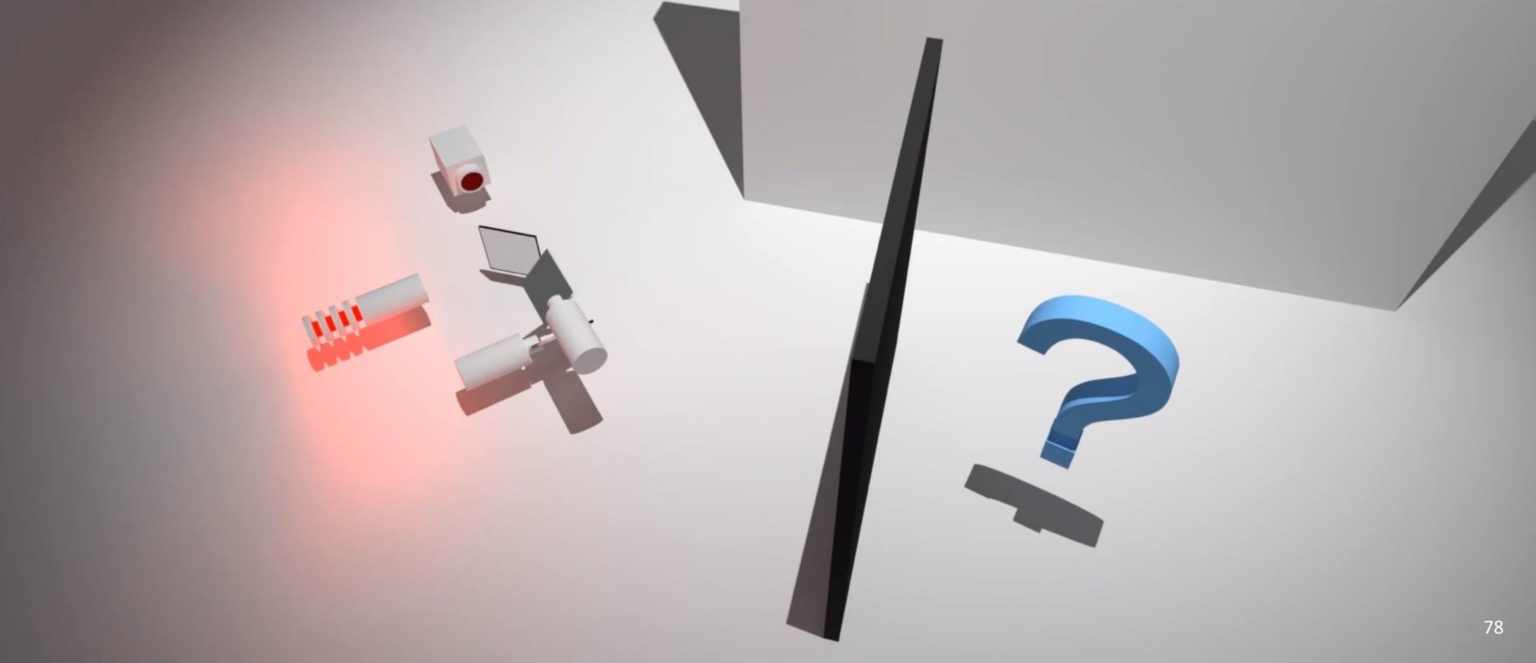

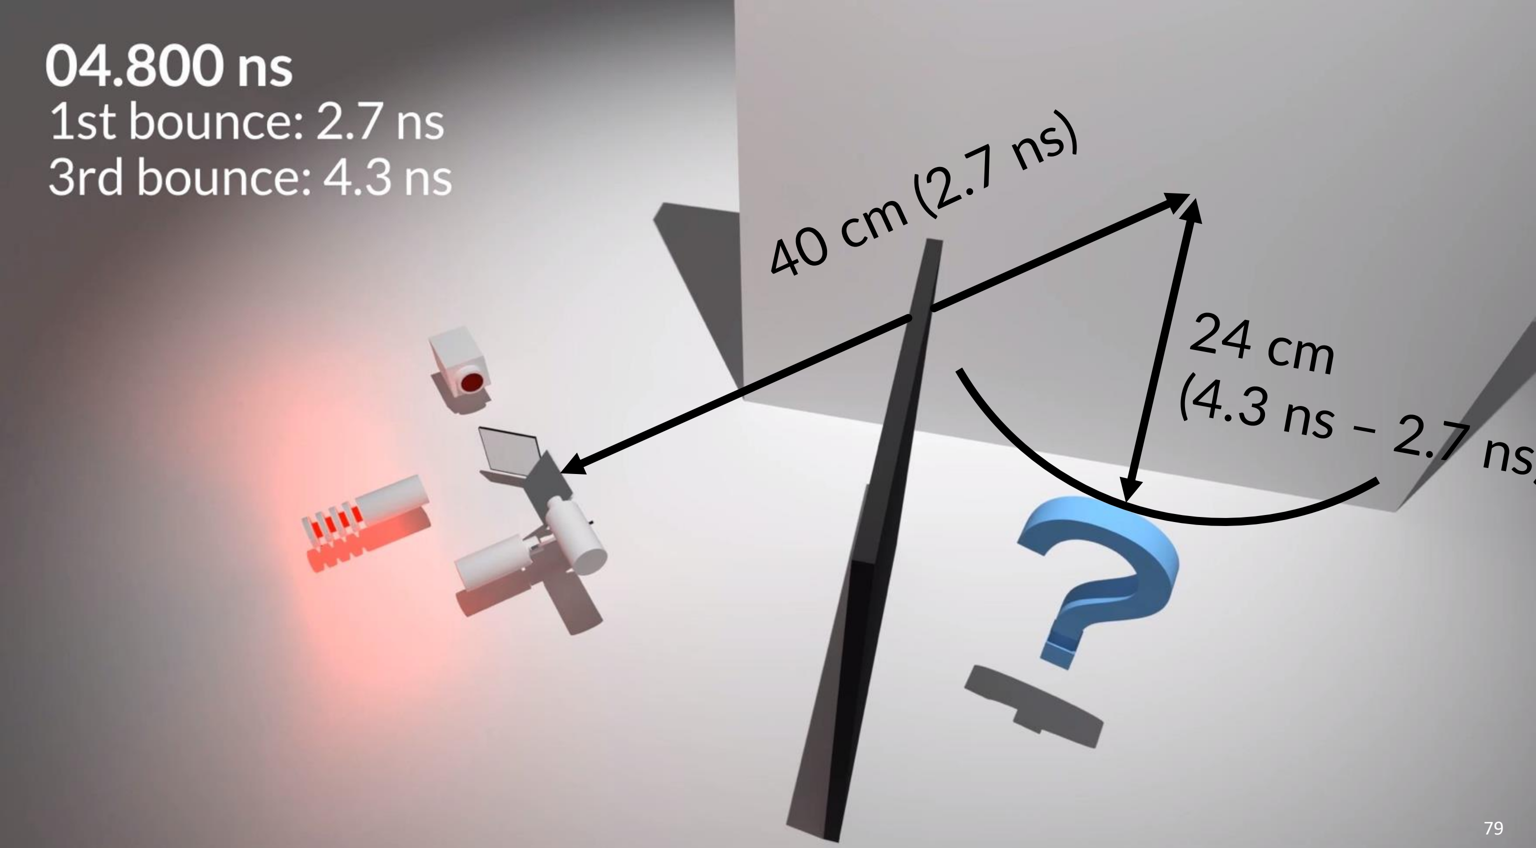
04.800 ns
1st bounce: 2.7 ns
3rd bounce: 4.3 ns
40 cm (2.7 ns)
24 cm
(4.3 ns – 2.7 ns)

RAW histogram (10 FPS)

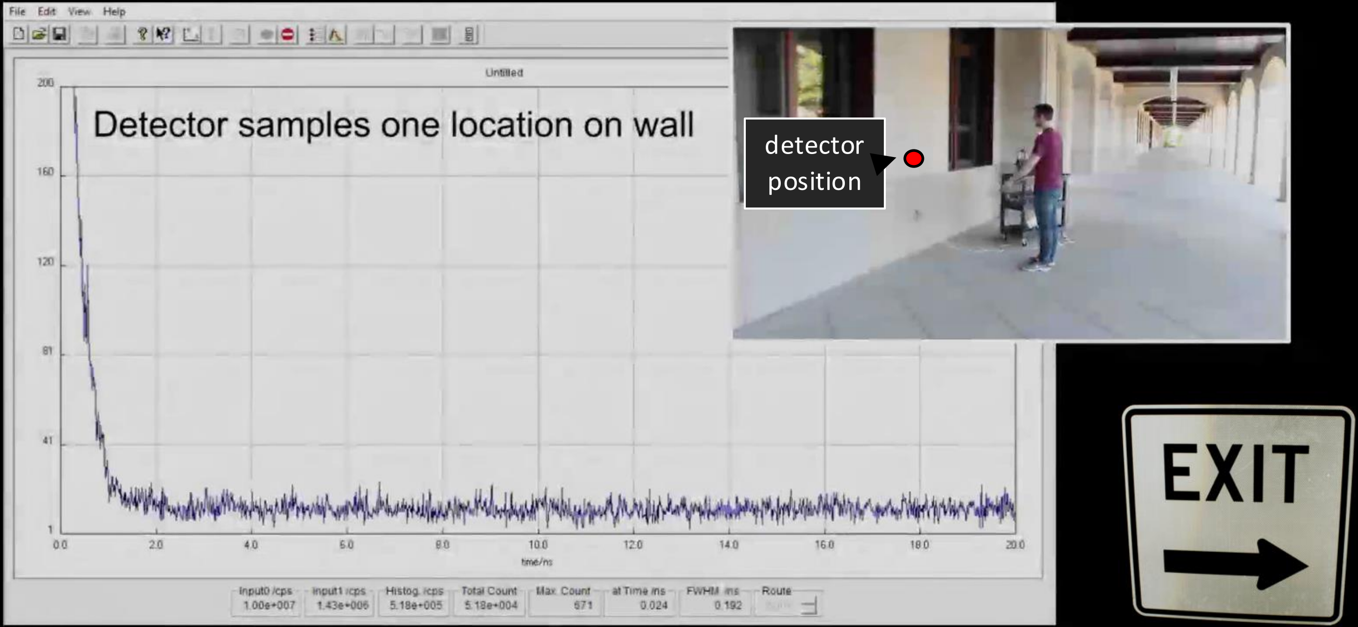

object

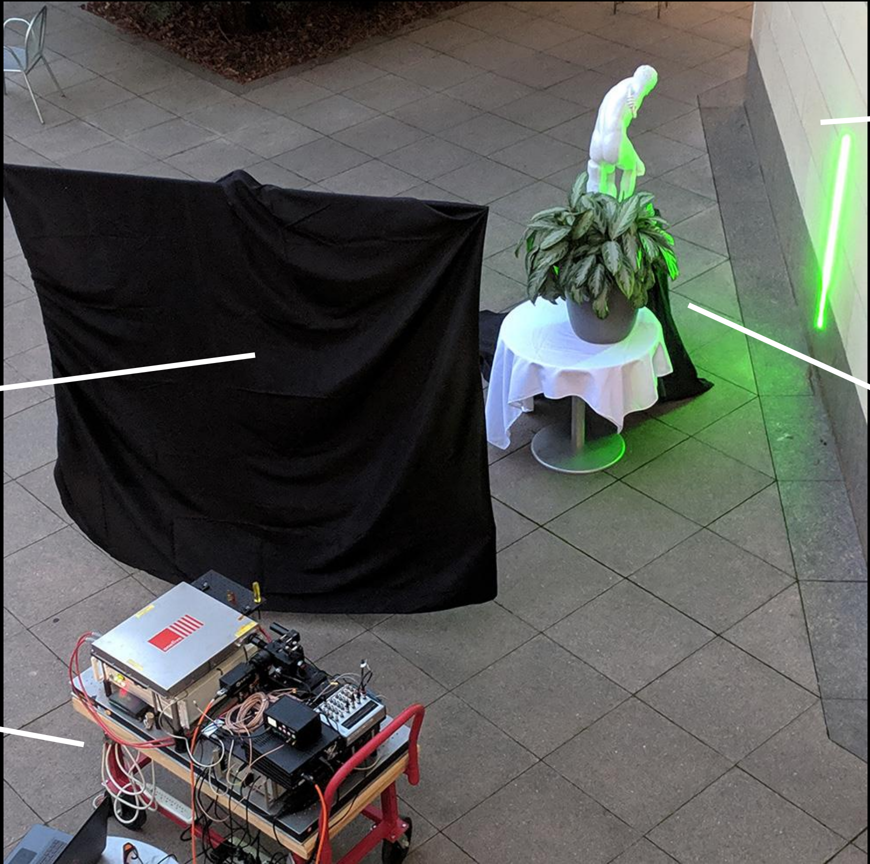
wall
occluder
hidden scene
NLOS imaging system

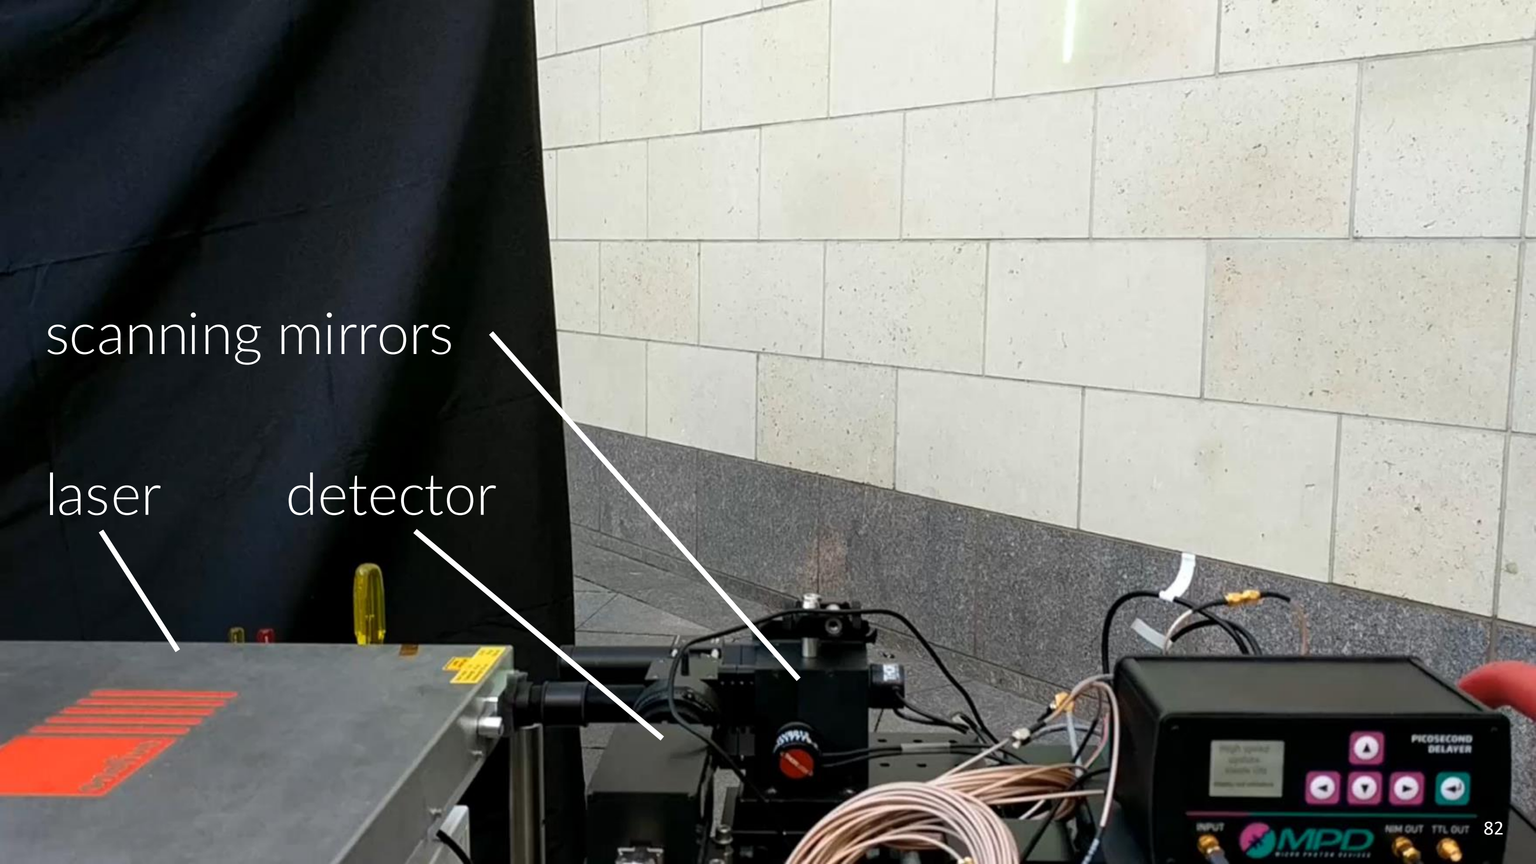
scanning mirrors
laser
detector
PICOSECOND DELAYER
MPD
INPUT
NIM OUT
TTL OUT


resolution: 128 x 128
area: 2 m × 2 m

scene photo

measurements

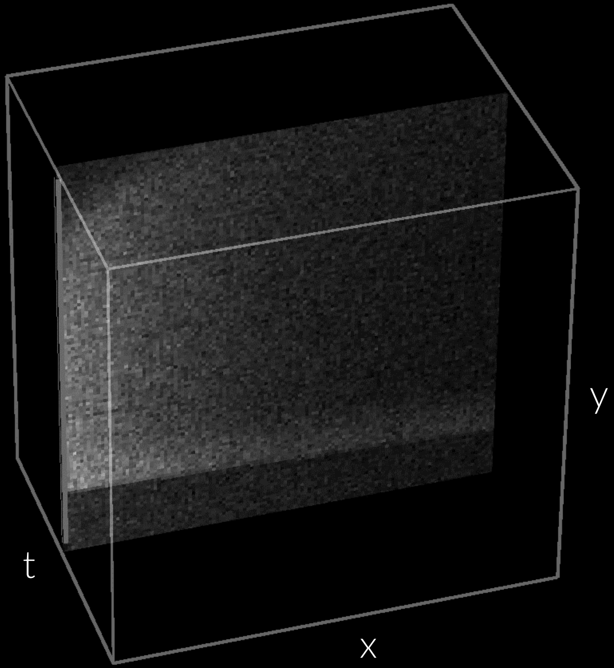

Lindell et al., SIGGRAPH 2019

scene photo

reconstruction

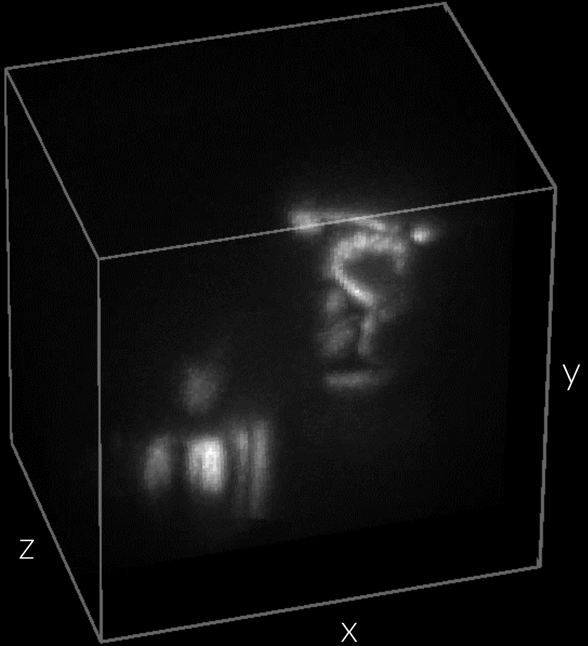

Lindell et al., SIGGRAPH 2019

Dimensions: 2 m x 2 m x 1.5 m

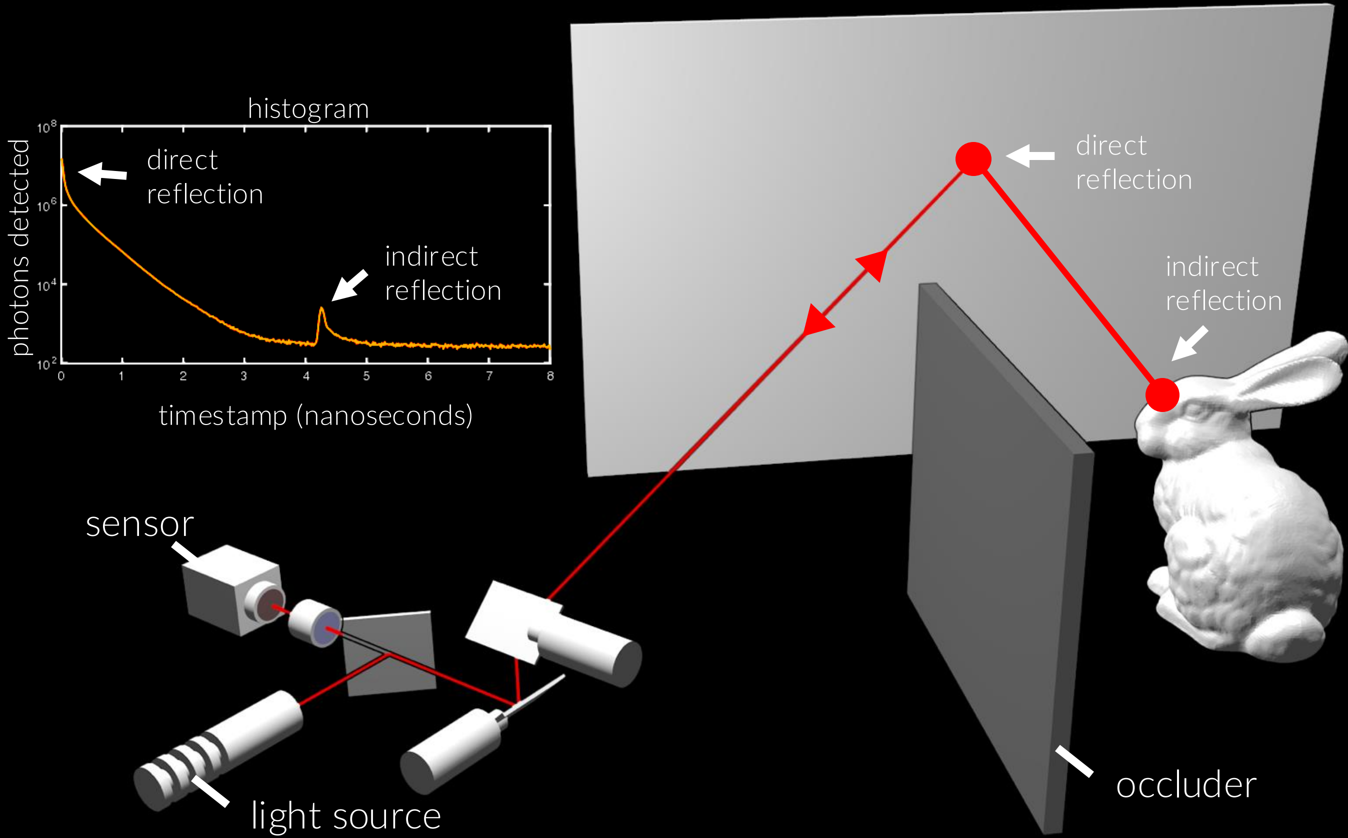
histogram
photons detected
10^8
10^6
10^4
10^2
direct reflection
indirect reflection
0 1 2 3 4 5 6 7 8
timestamp (nanoseconds)
direct reflection
indirect reflection
sensor
light source
occluder

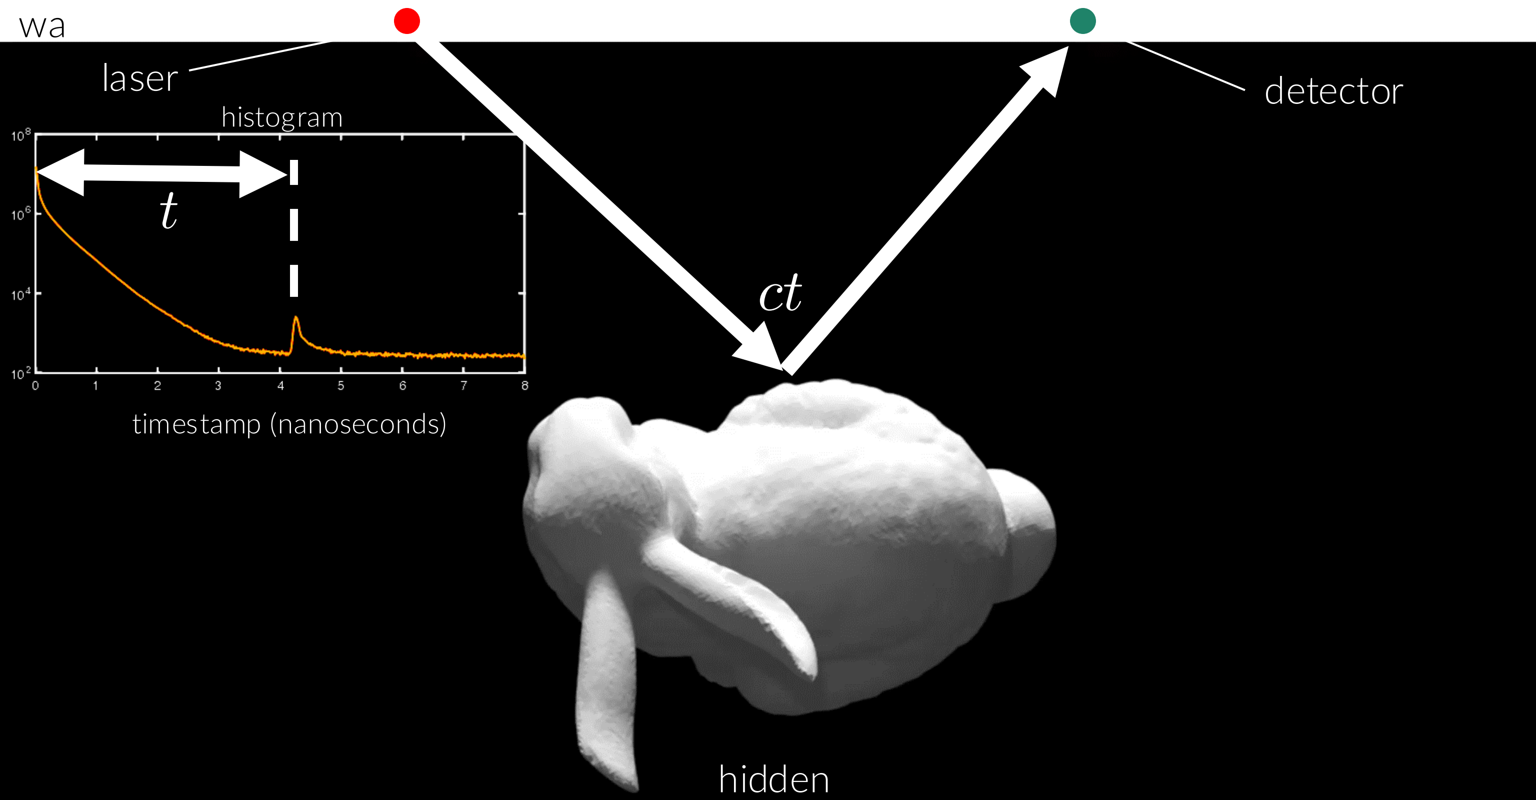

non-confocal sampling

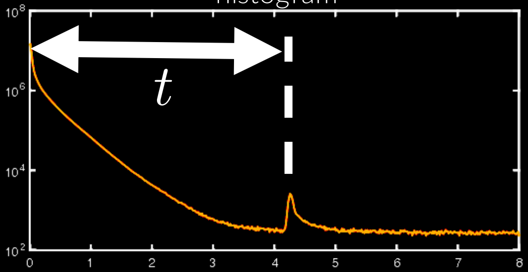

ct

laser and detector focus on this point

hidden object

APJarvis [CC BY-SA 4.0]

lasers and detectors illuminate and image same points

confocal sampling

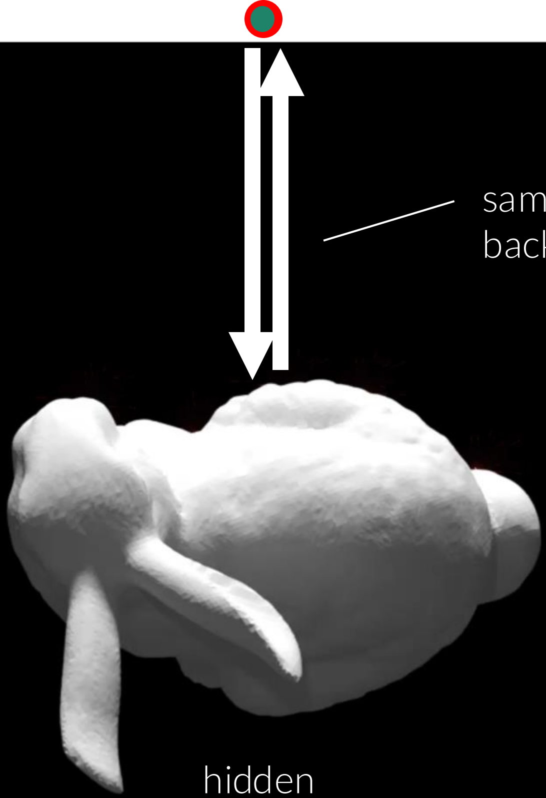

confocal sampling

wa

- simplified NLOS mathematical model
- enables efficient NLOS reconstruction

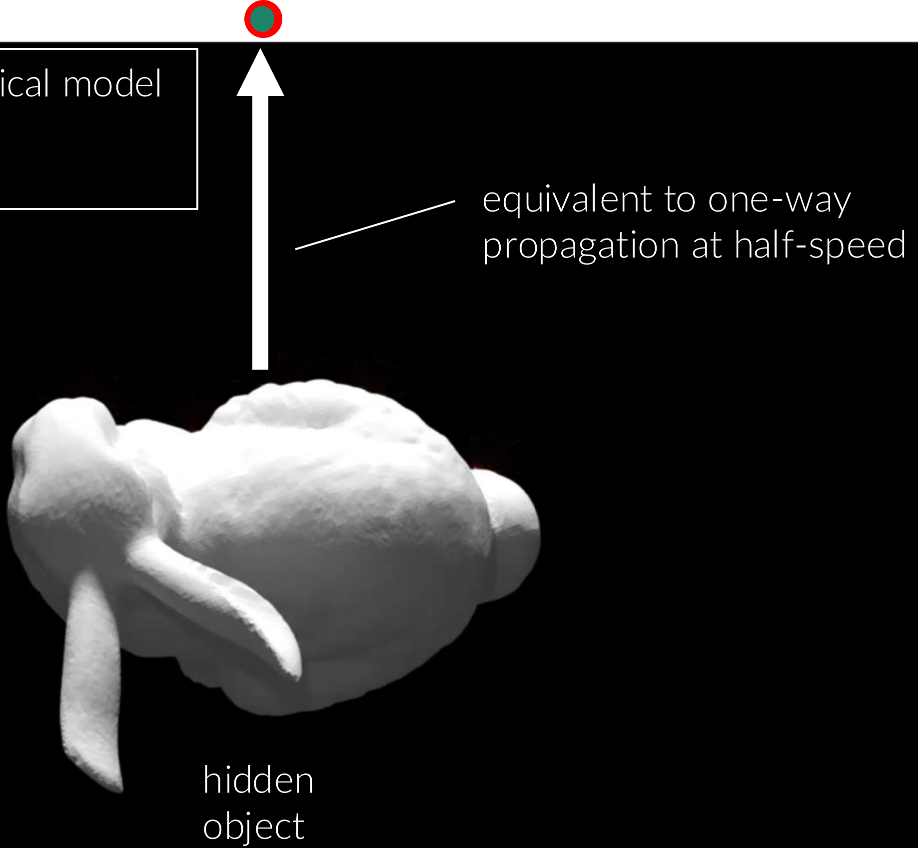

confocal sampling

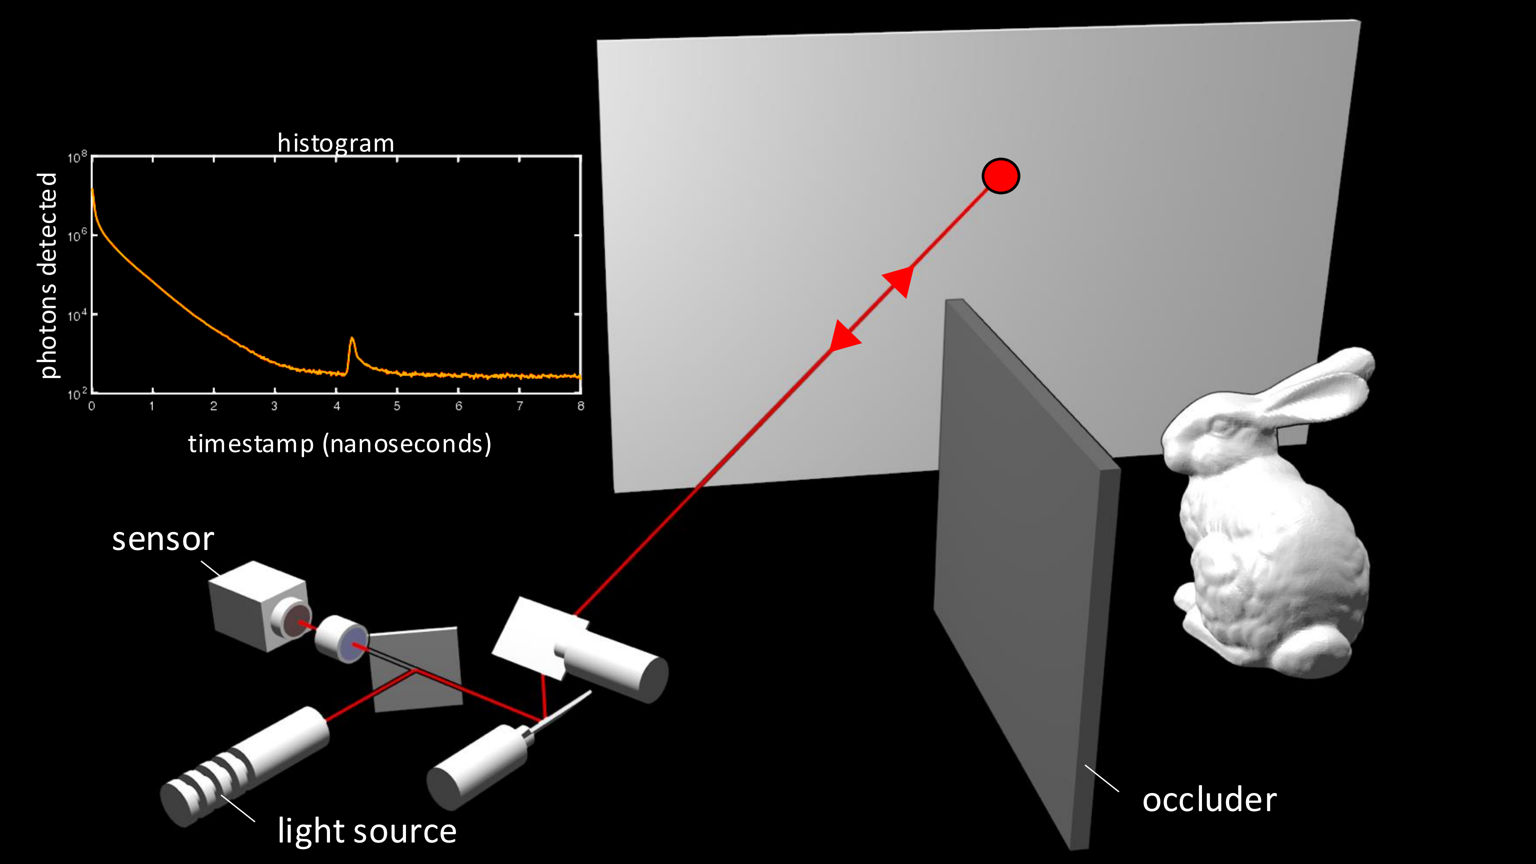
histogram
photons detected
timestamp (nanoseconds)
sensor
light source
occluder

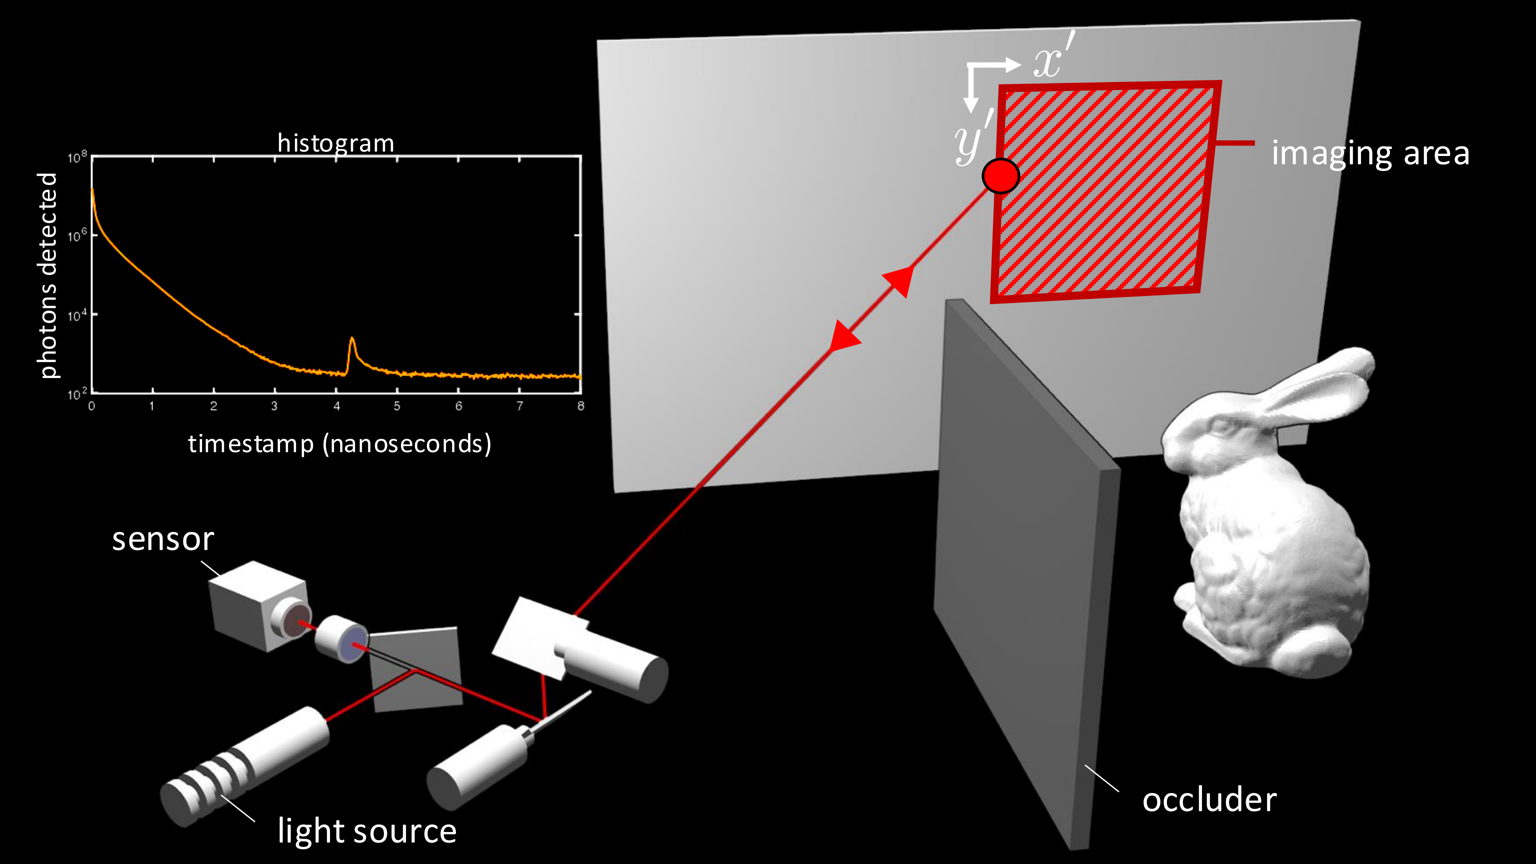

histogram
photons detected
timestamp (nanoseconds)
x'
y'
imaging area
sensor
light source
occluder

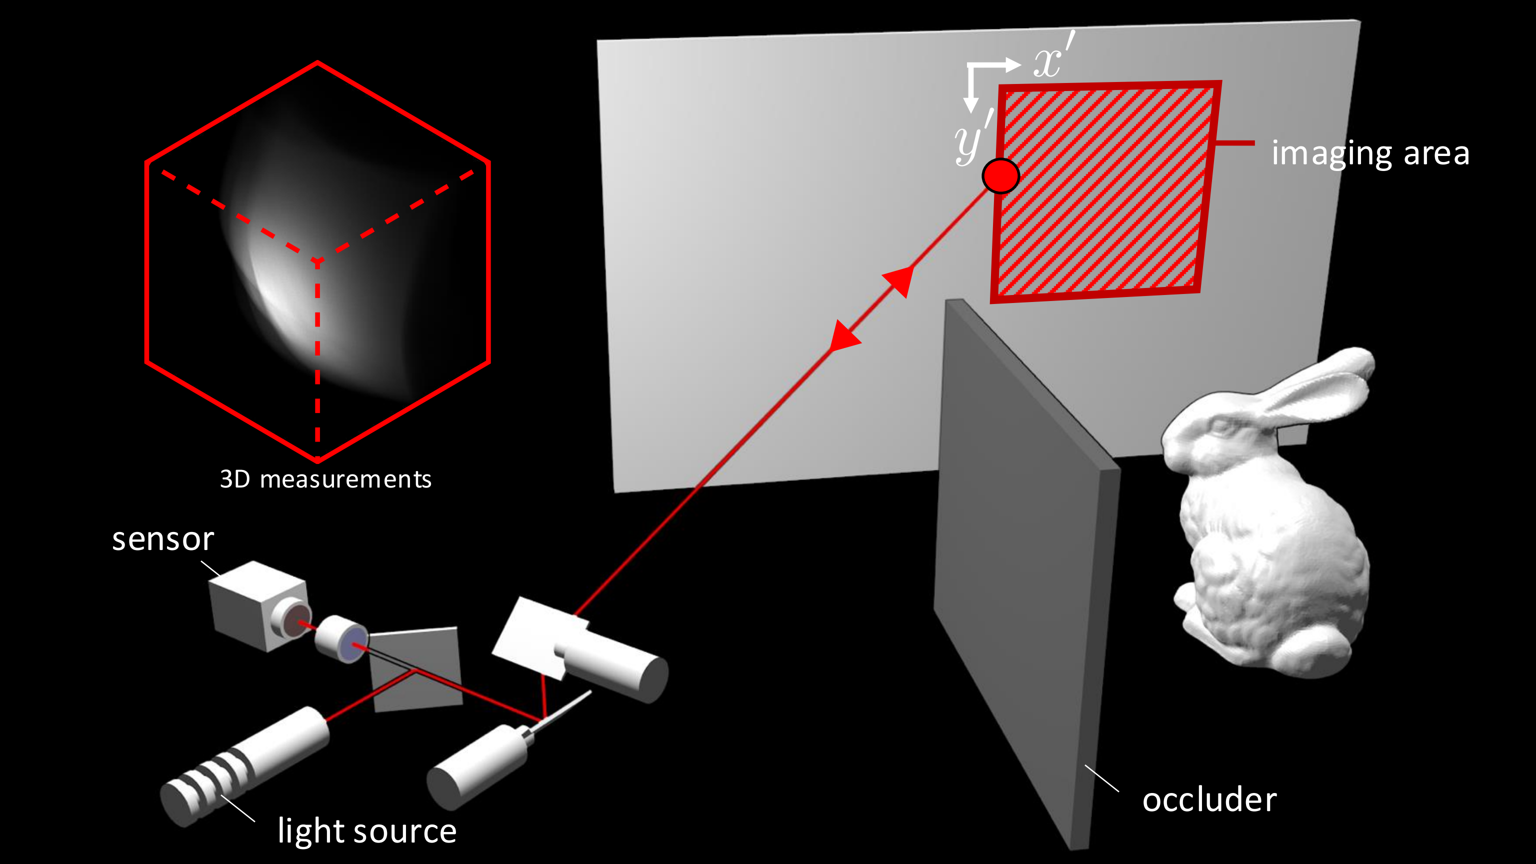

x'
y'
imaging area
3D measurements
sensor
occluder
light source

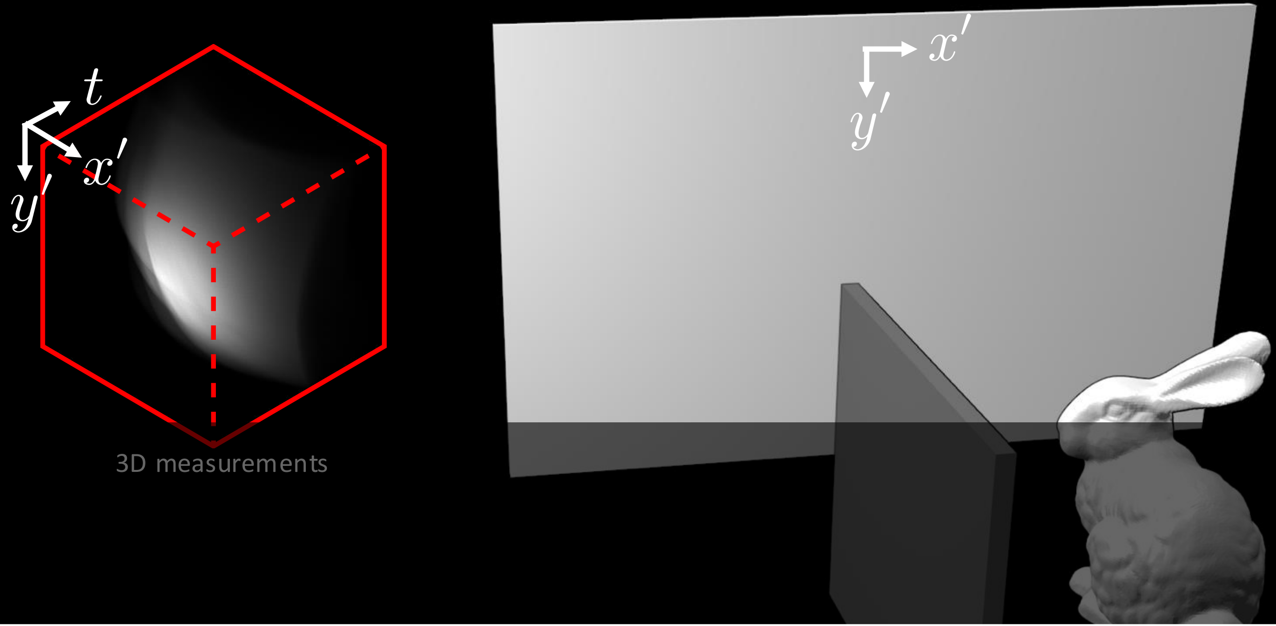

$$\underbrace{\tau(x', y', t)}_{\text{3D measurements}} = \iiint_{\Omega} \underbrace{\frac{1}{r_l^2 r^2}}_{\text{radiometric term}} \underbrace{\delta\left(r_l + r - tc\right)}_{\text{geometric term}} \cdot \underbrace{\rho(x, y, z)}_{\text{hidden 3D volume}} \; dx \; dy \; dz$$

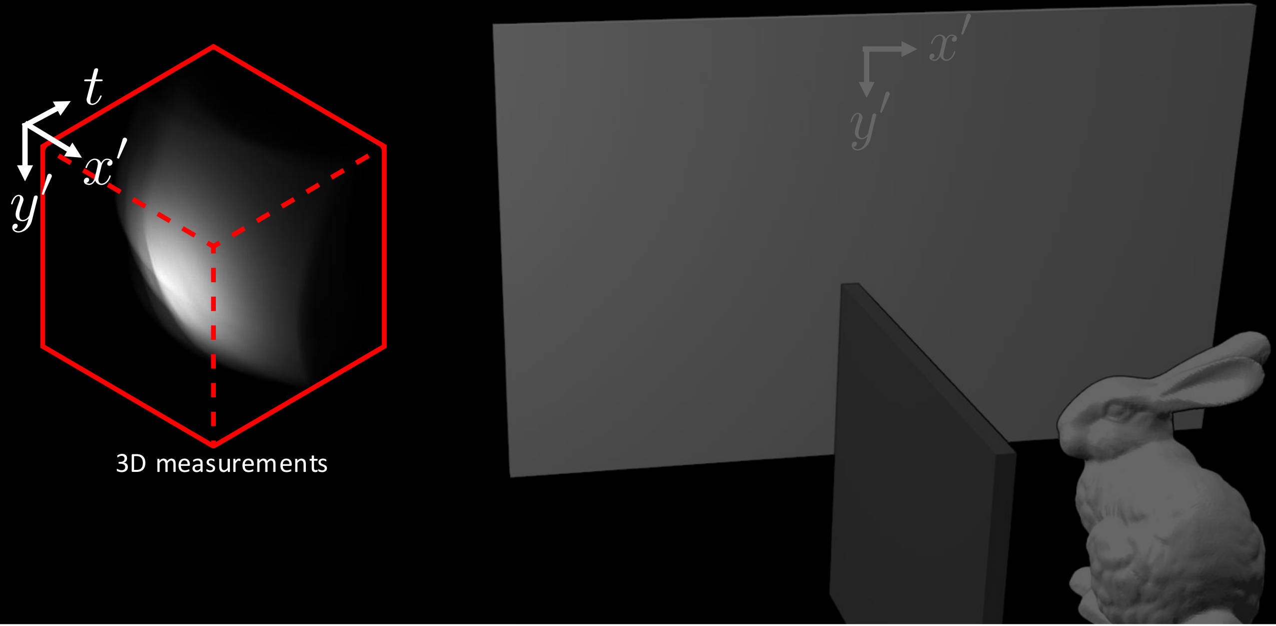

$$\underbrace{\tau(x', y', t)}_{\text{3D measurements}} = \iiint_{\Omega} \underbrace{\frac{1}{r_l^2 r^2}}_{\text{radiometric term}} \underbrace{\delta\left(r_l + r - tc\right)}_{\text{geometric term}} \cdot \underbrace{\rho(x, y, z)}_{\text{hidden 3D volume}} \, dx \, dy \, dz$$

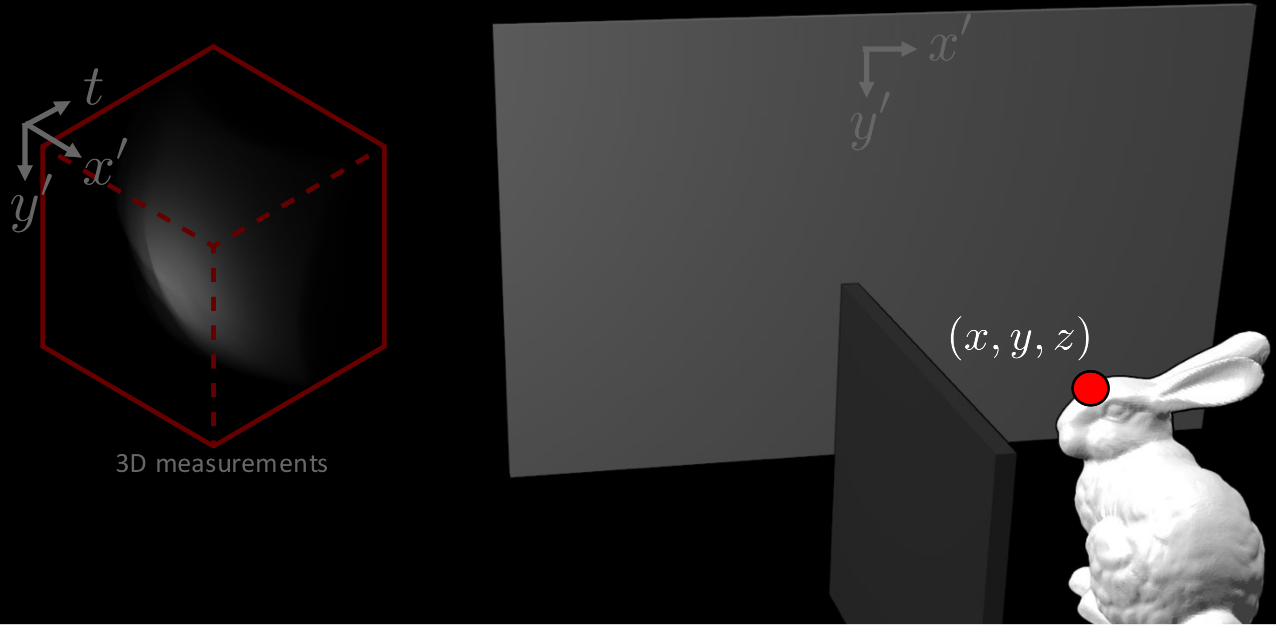

$$\underbrace{\tau(x', y', t)}_{\text{3D measurements}} = \iiint_\Omega \underbrace{\frac{1}{r_l^2 r^2}}_{\text{radiometric term}} \underbrace{\delta\left(r_l + r - tc\right)}_{\text{geometric term}} \cdot \underbrace{\rho(x, y, z)}_{\text{hidden 3D volume}} \, dx \; dy \; dz$$

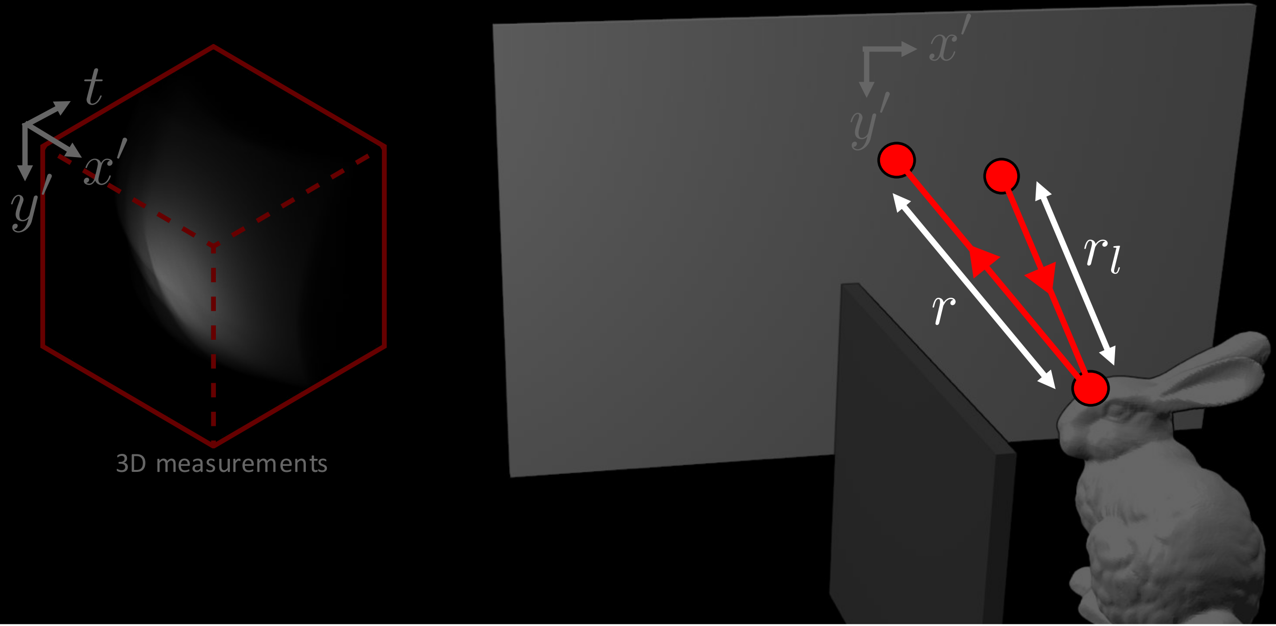

$$\underbrace{\tau(x',y',t)}_{\text{3D measurements}} = \iiint_\Omega \underbrace{\frac{1}{r_l^2 r^2}}_{\text{radiometric term}} \underbrace{\delta\left(r_l + r - tc\right)}_{\text{geometric term}} \cdot \underbrace{\rho(x,y,z)}_{\text{hidden 3D volume}} \, dx \; dy \; dz$$

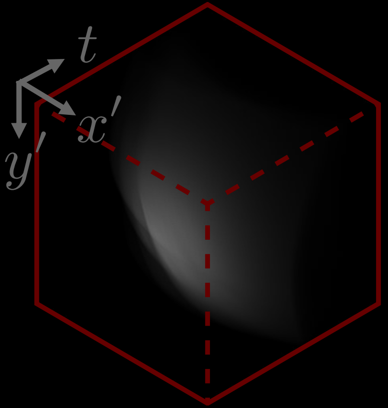

3D measurements

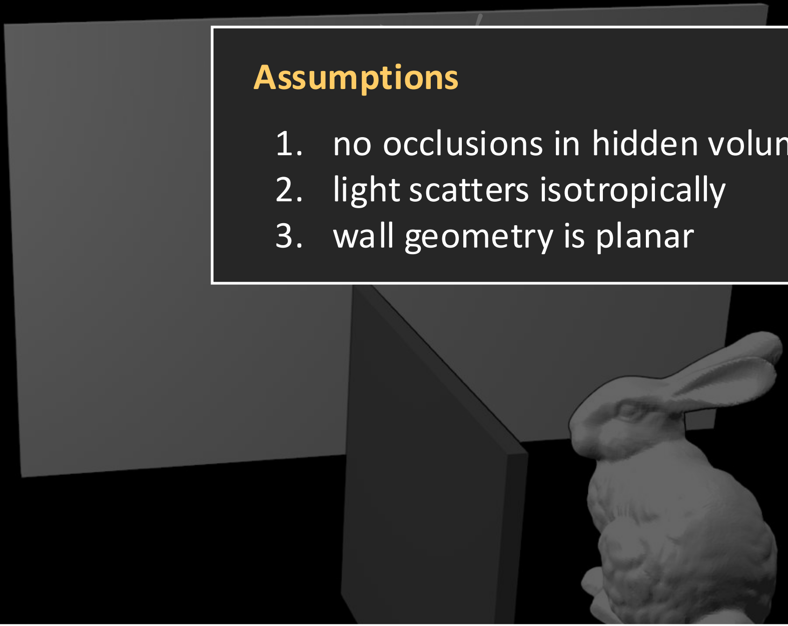

Assumptions

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

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

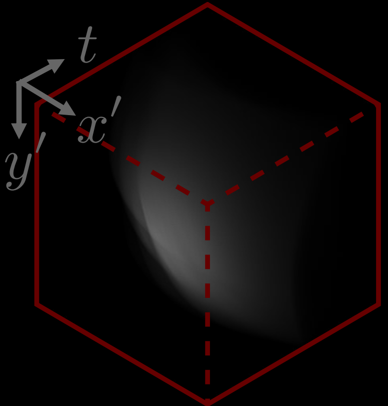

3D measurements

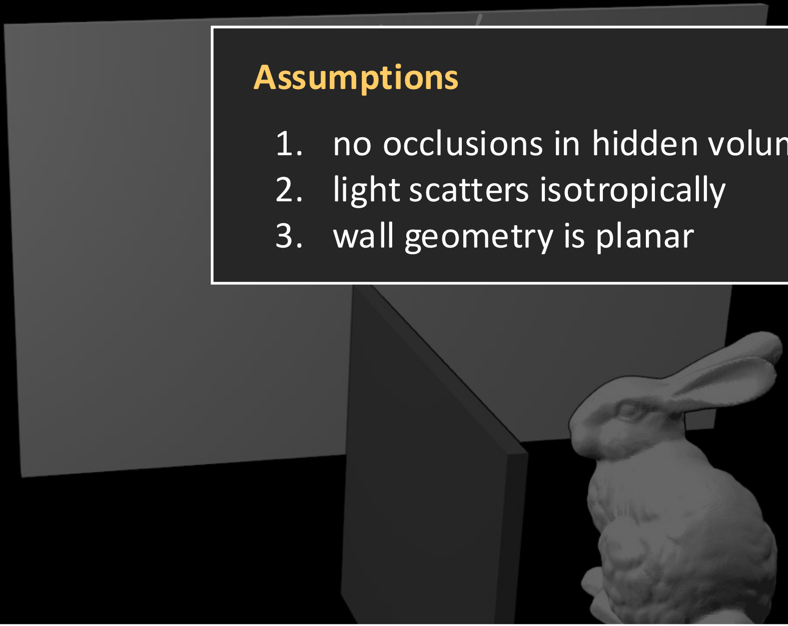

Assumptions

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

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

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

NLOS image formation mode:

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

measurements $n^3 \times 1$

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

unknown volume $n^3 \times 1$

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.

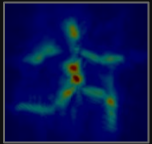

Iterative Inversion [Gupta 12, Heide 13]

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 $z = \sqrt{u}$, $t = 2\sqrt{v}/c$)

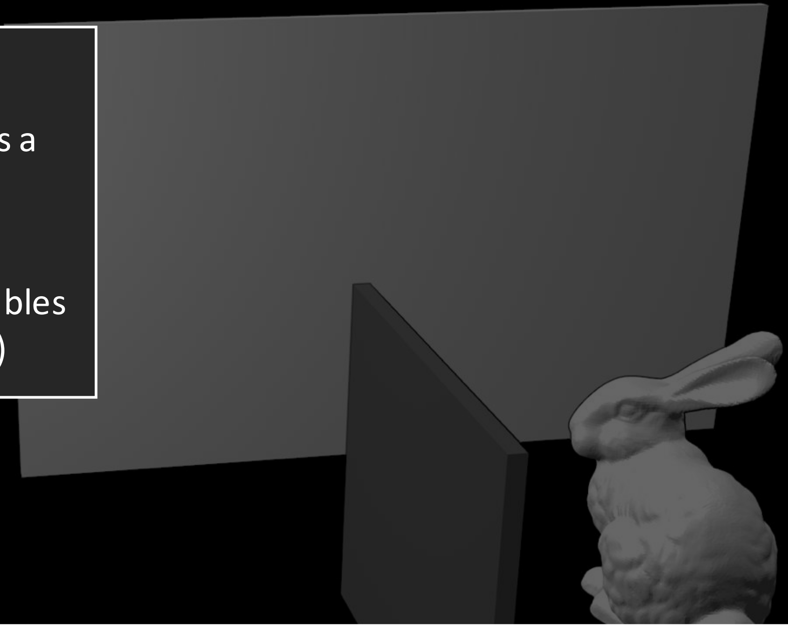

$$\tau(x', y', t) = \iiint_\Omega \frac{1}{r_l^2 r^2} \delta\left(r_l + r - tc\right) \cdot \rho(x, y, z)\ dx\ dy\ dz$$

Our approach

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

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

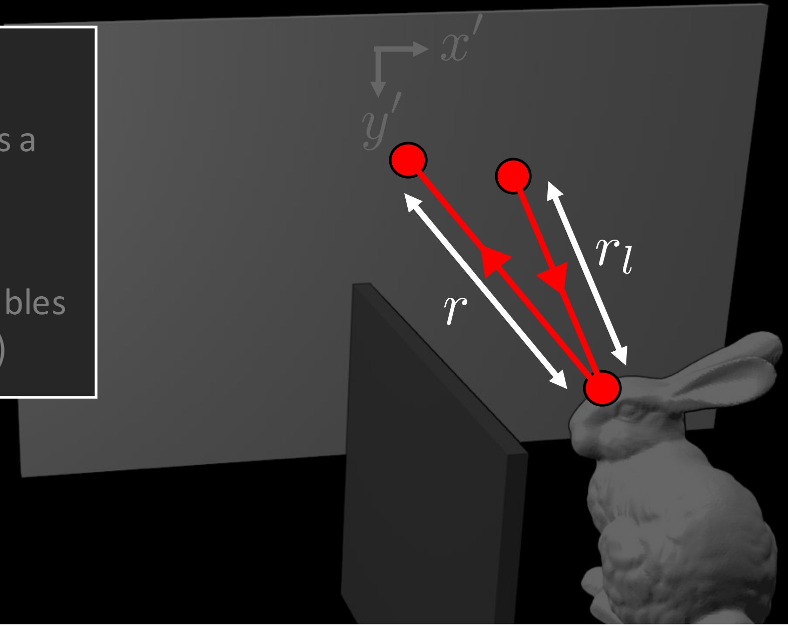

$$\tau(x', y', t) = \iiint_{\Omega} \frac{1}{r_l^2 r^2} \delta\left(r_l + r - tc\right) \cdot \rho(x, y, z)\; dx\; dy\; dz$$

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

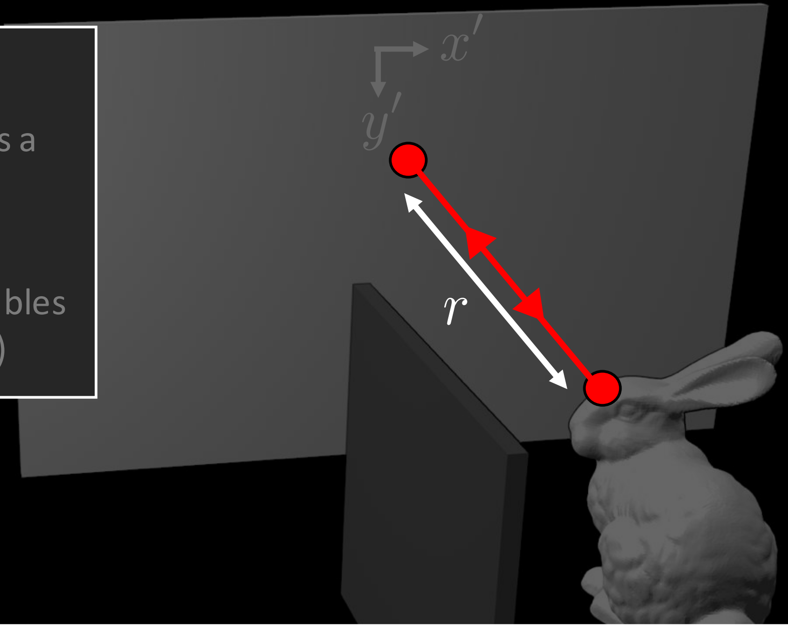

$$\tau(x', y', t) = \iiint_\Omega \frac{1}{r^4} \delta\left(2r - tc\right) \cdot \rho(x, y, z)\; dx\; dy\; dz$$

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

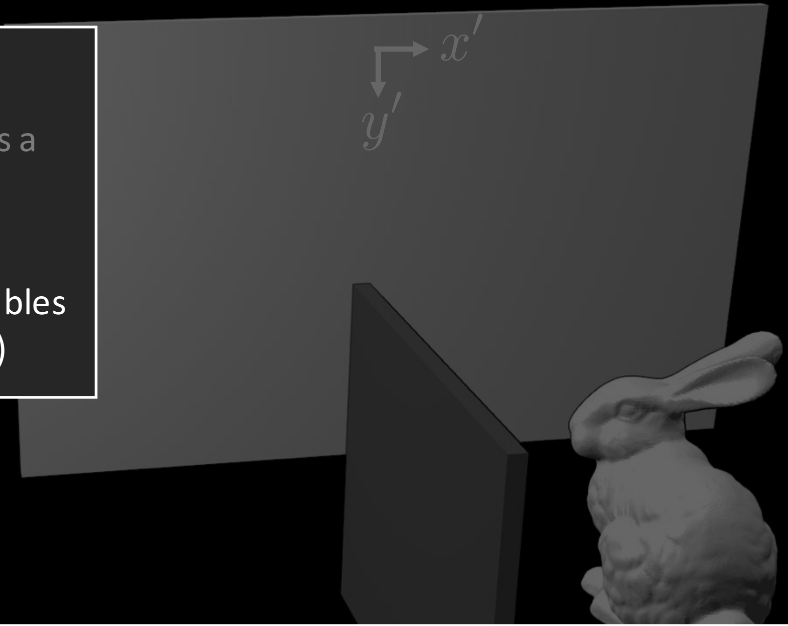

$$\tau(x', y', t) = \iiint_\Omega \frac{1}{r^4} \delta\left(2r - tc\right) \cdot \rho(x, y, z) \; dx \; dy \; dz$$

Our approach

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

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

3D measurements

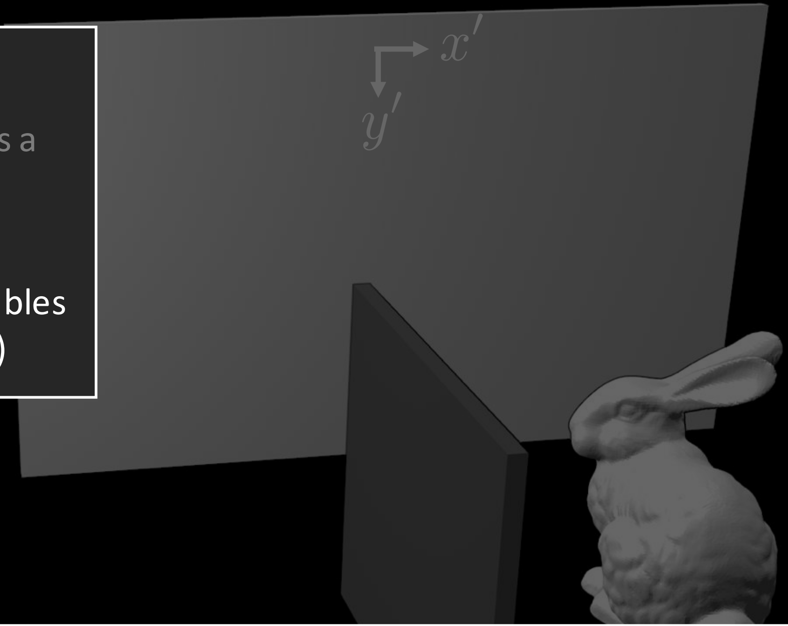

$$v^{3/2}\tau(x', y', \frac{2}{c}\sqrt{v}) = \iiint_\Omega \frac{1}{2\sqrt{u}} \delta\left((x'-x)^2 + (y'-y)^2 + u - v\right) \cdot \rho(x, y, \sqrt{u}) dx dy du$$

Our approach

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

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

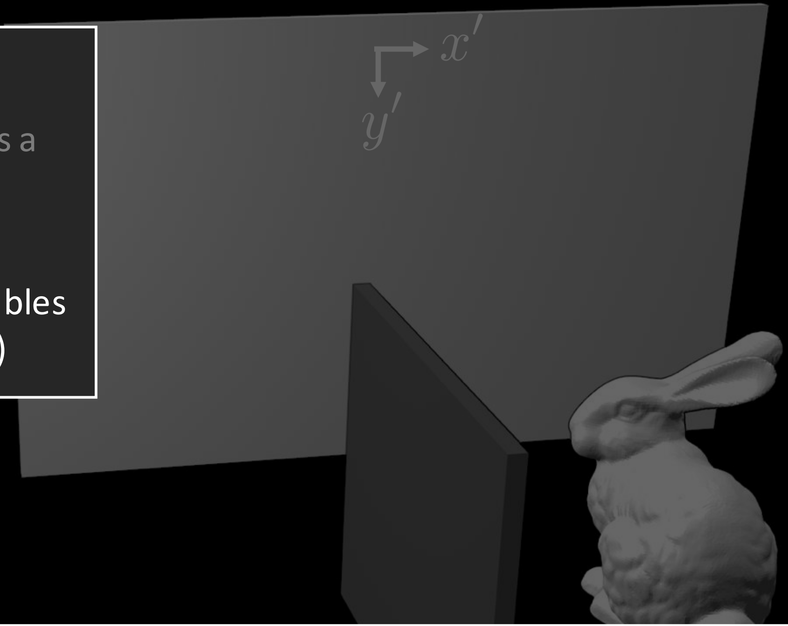

$$v^{3/2}\tau(x', y', \tfrac{2}{c}\sqrt{v}) = \iiint_\Omega \frac{1}{2\sqrt{u}}\delta\left((x'-x)^2 + (y'-y)^2 + u - v\right) \cdot \rho(x, y, \sqrt{u})\,dx\,dy\,du$$

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

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.

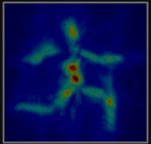

Iterative Inversion [Gupta 12, Heide 13]

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 $n^3 \times 1$

Backpropagation [Velten 12, Buttafava 15]

Flops: $O(n^5)$

Memory: $O(n^3)$

Runtime: Approx. 10 min.

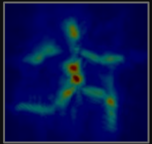

Iterative Inversion [Gupta 12, Heide 13]

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 $n \times n \times n$

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

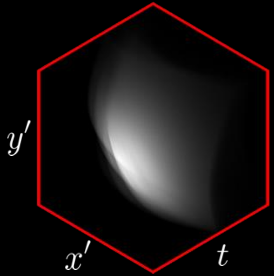

O'Toole et al., Nature 2018

measurements

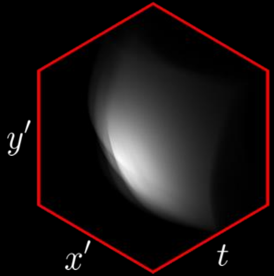

Step 1: resample and attenuate along t -axis

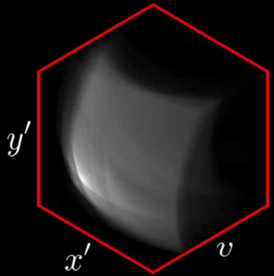

O'Toole et al., Nature 2018

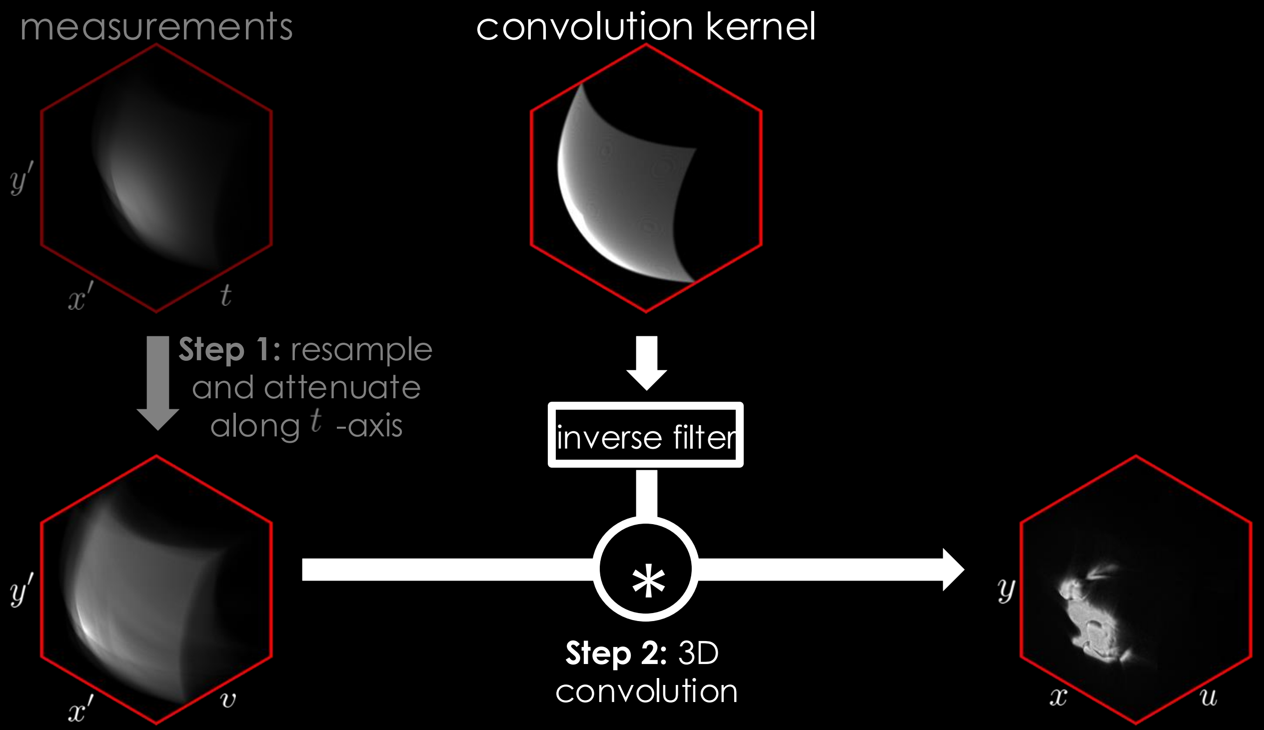

O'Toole et al., Nature 2018

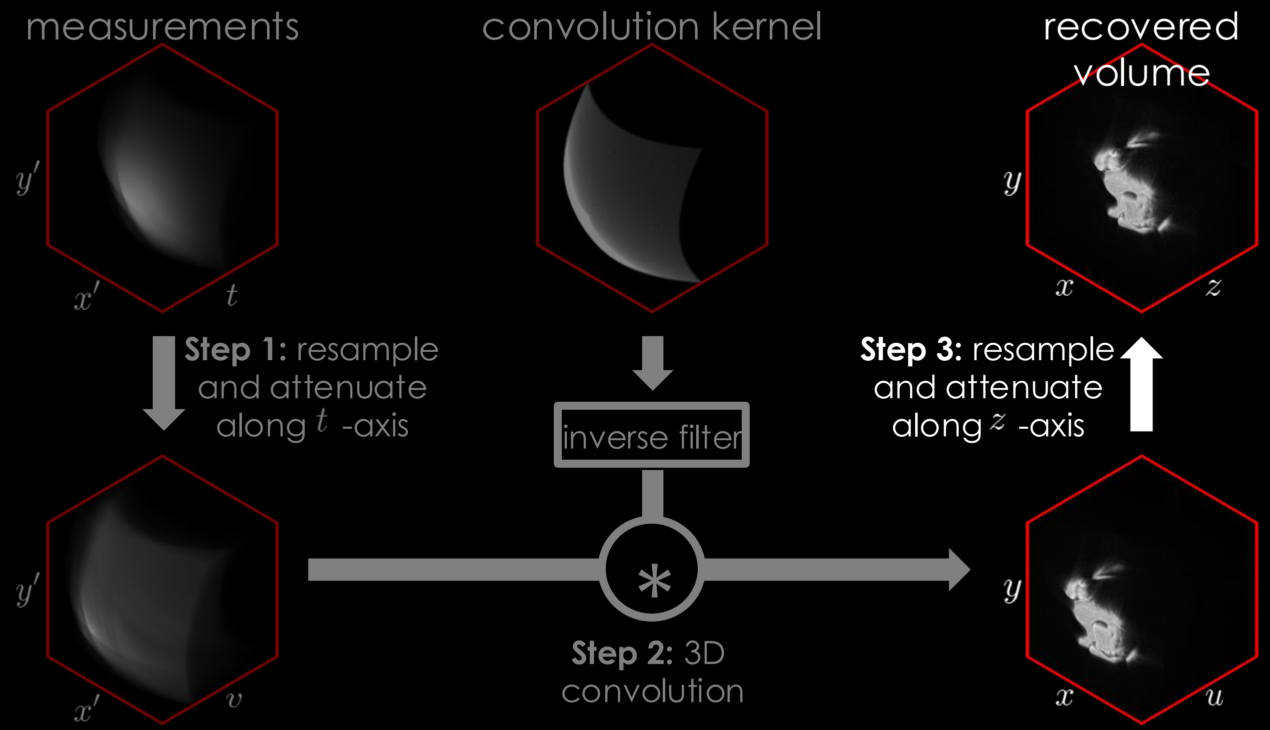

O'Toole et al., Nature 2018

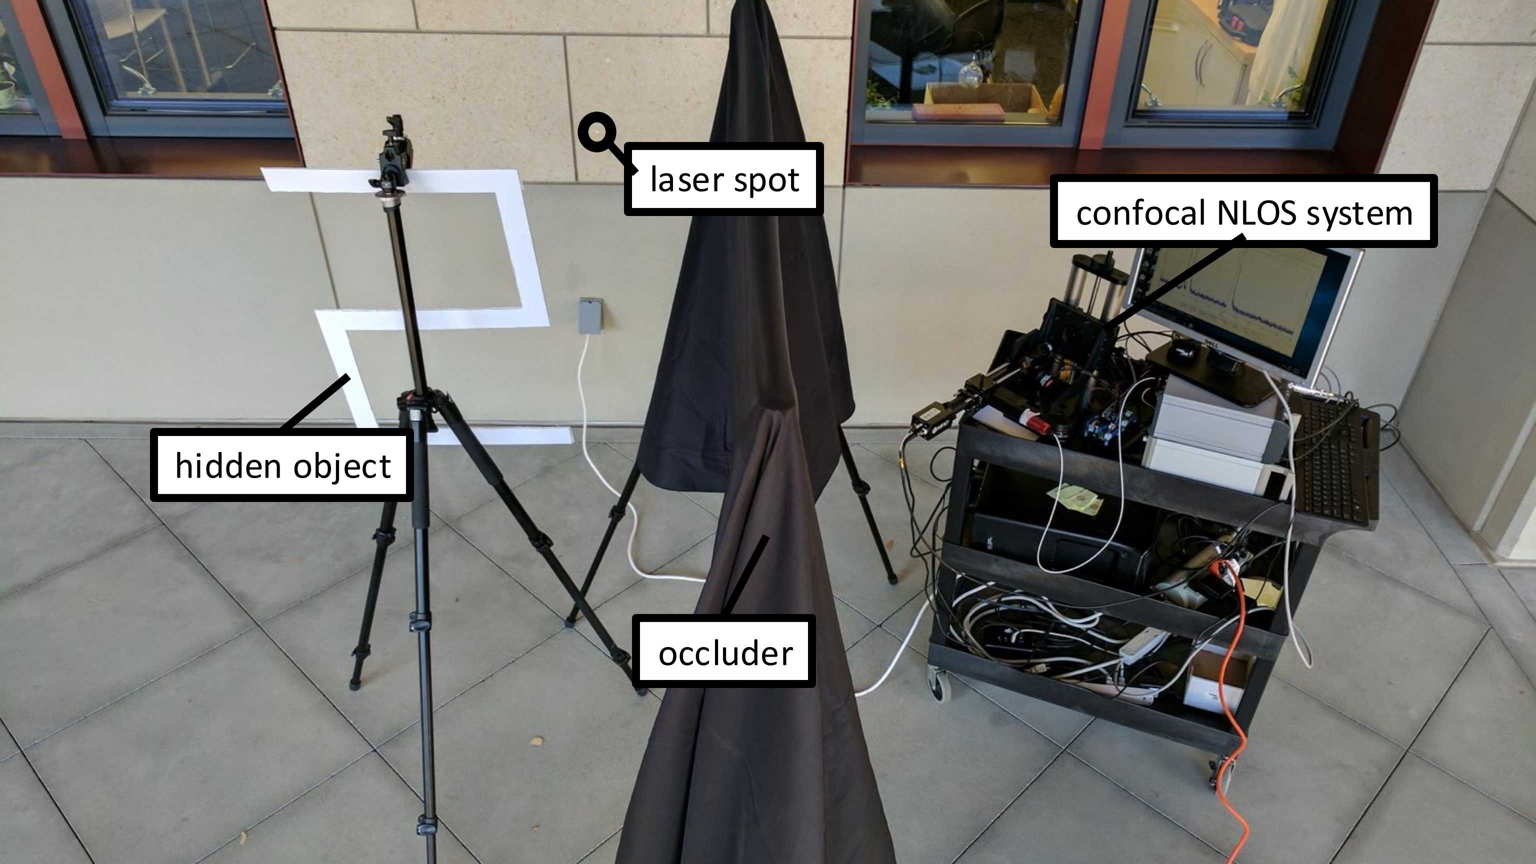
laser spot
confocal NLOS system
hidden object
occluder


1.25 meters

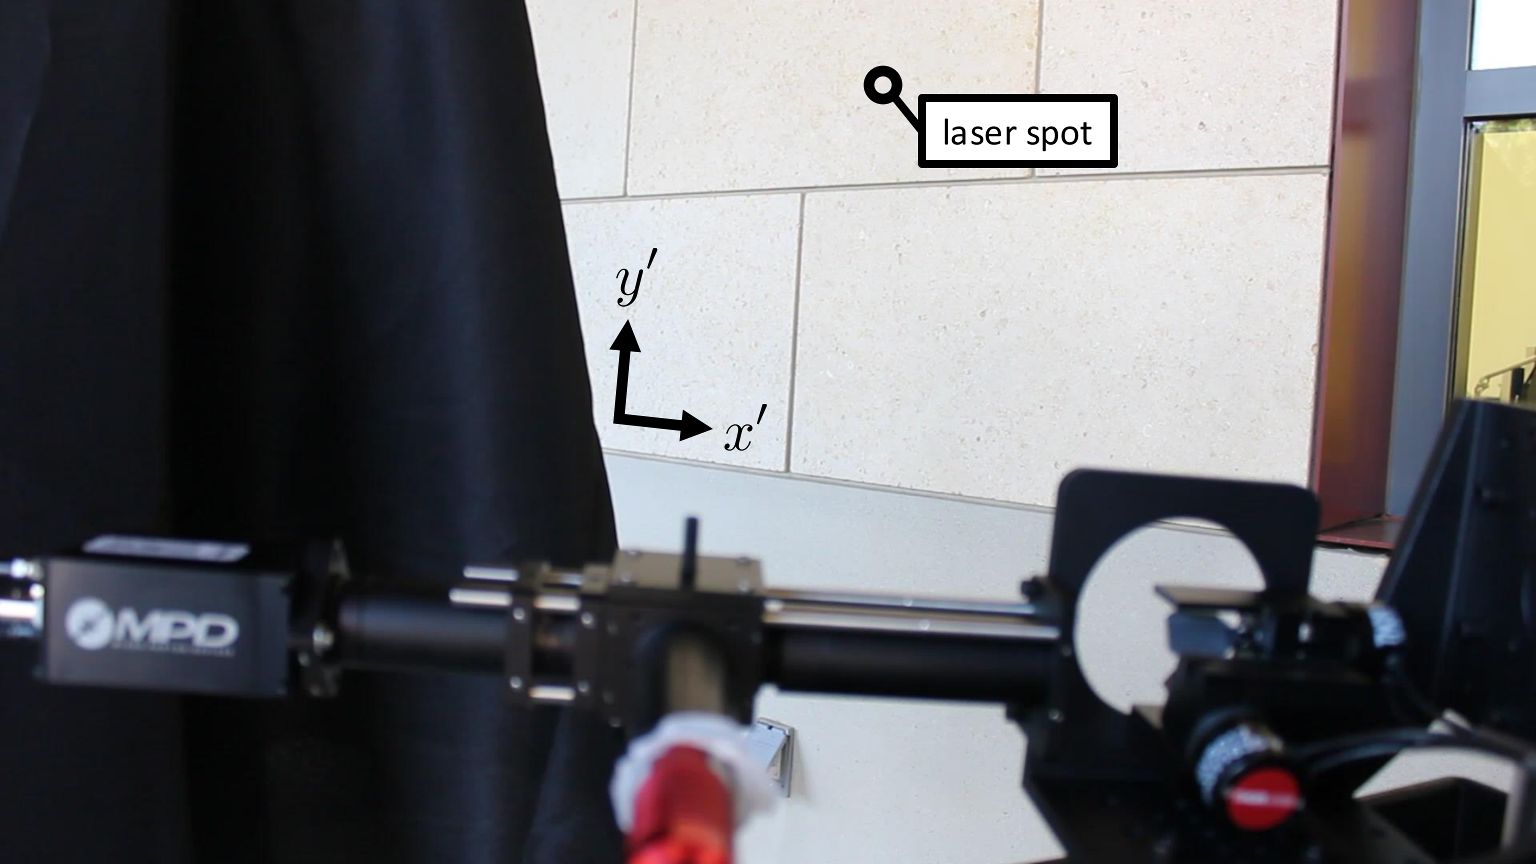

laser spot
y'
x'
MPD

measurements

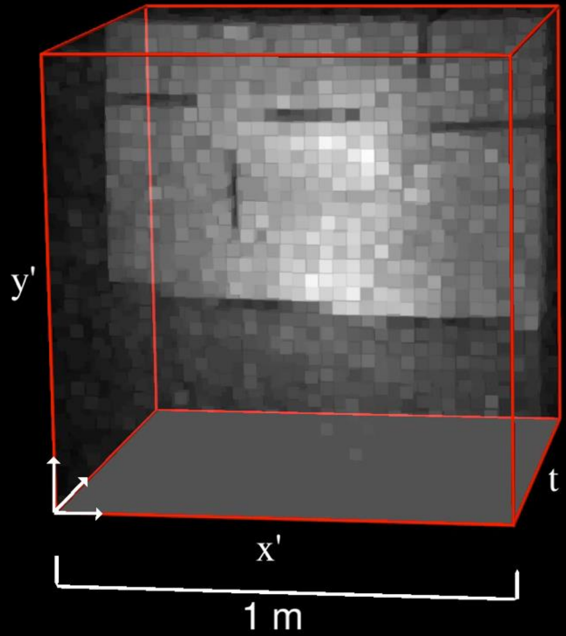

Maximum Intensity Projection

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.

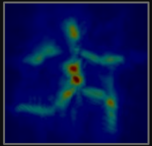

3D Deconvolution (with Light-Cone Transform) [O'Toole et al. 2018]

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

Memory: $O(n^3)$

Runtime: $<$ 1 second

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

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

Memory: $O(n^5)$

Runtime: $>$ 1 hour

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)
- neural rendering for propagating light

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

wa

hidden
object

confocal sampling

image formation model

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

wavefield

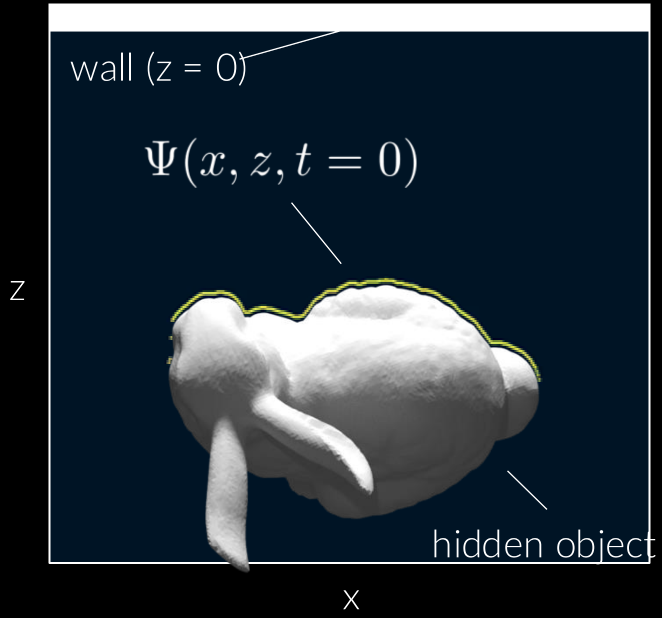

$$\Psi(x, z = 0, t)$$

confocal measurements

general solution (time reversal)

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

wavefield

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

confocal measurements

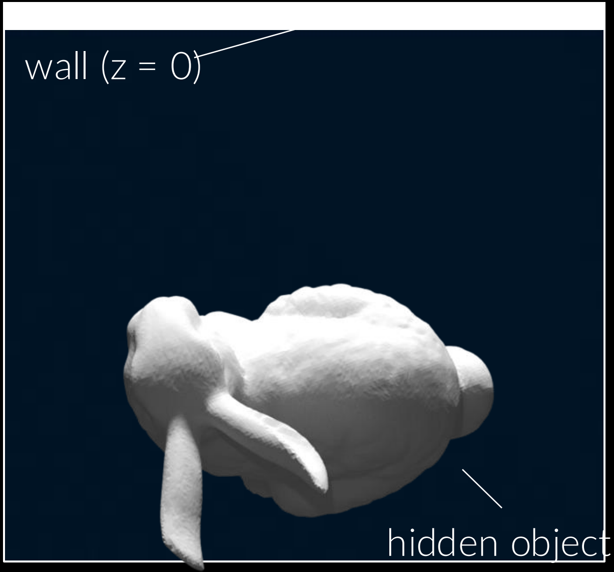

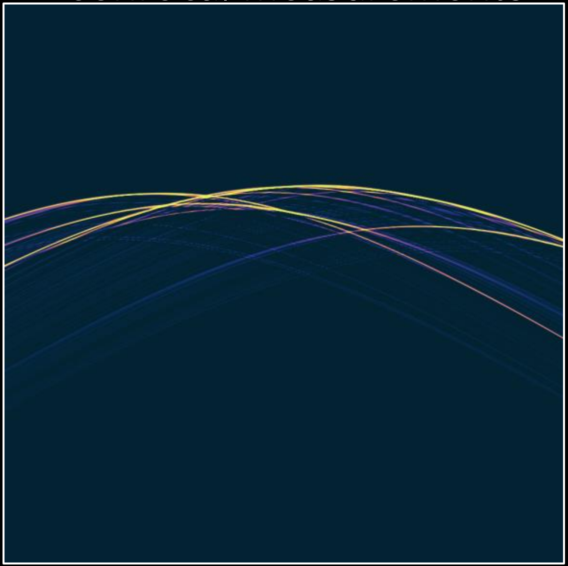

general solution (time reversal)

finite-difference time-domain method

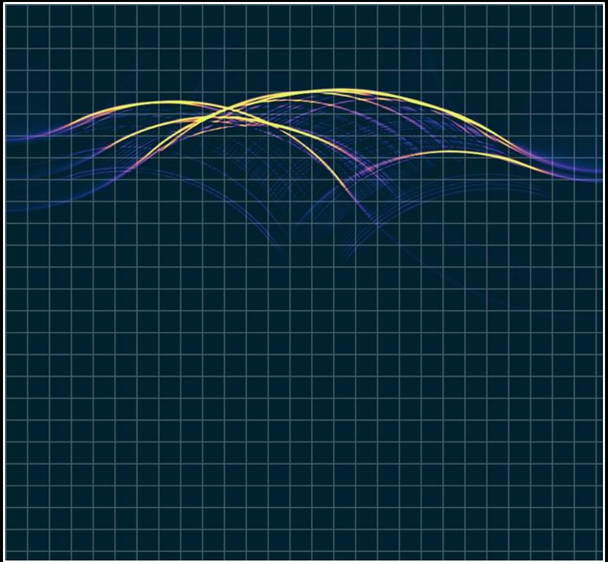

1. approximate wave equation with finite differences

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

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

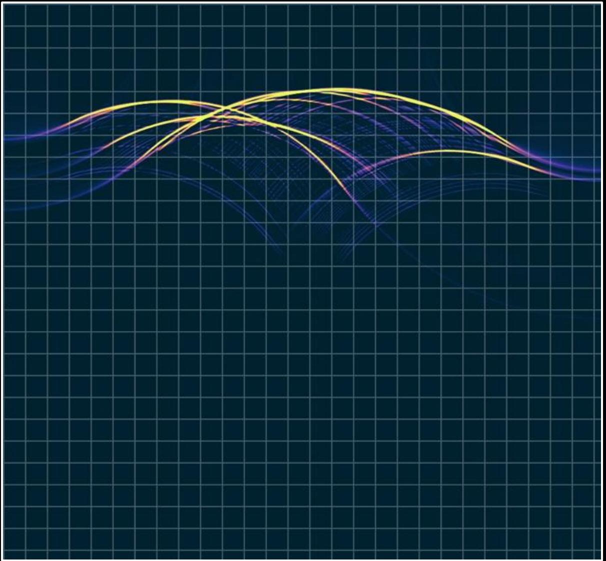

1. approximate wave equation with finite differences

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

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

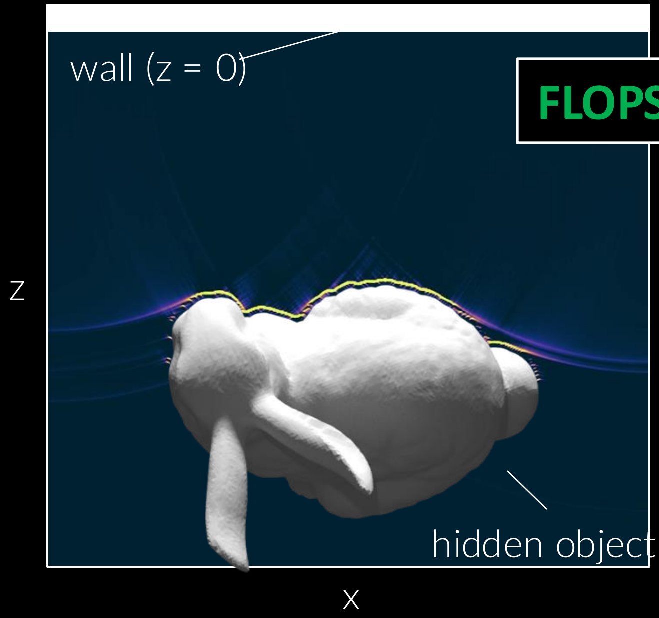

FLOPS: $O(n^3 \log n)$

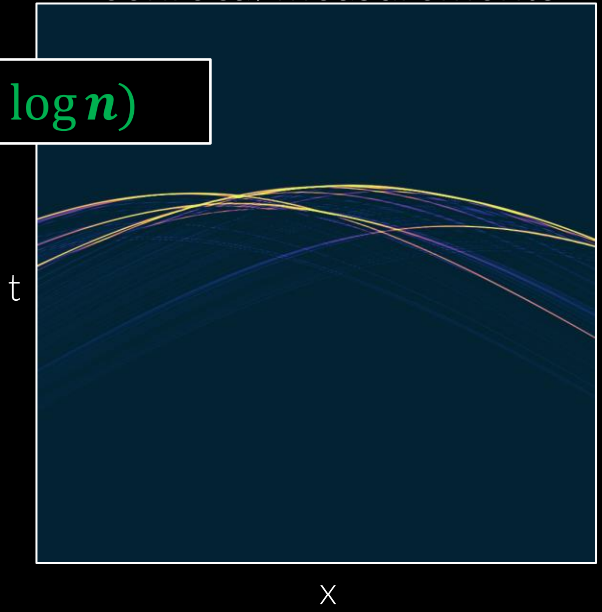

hardware prototype

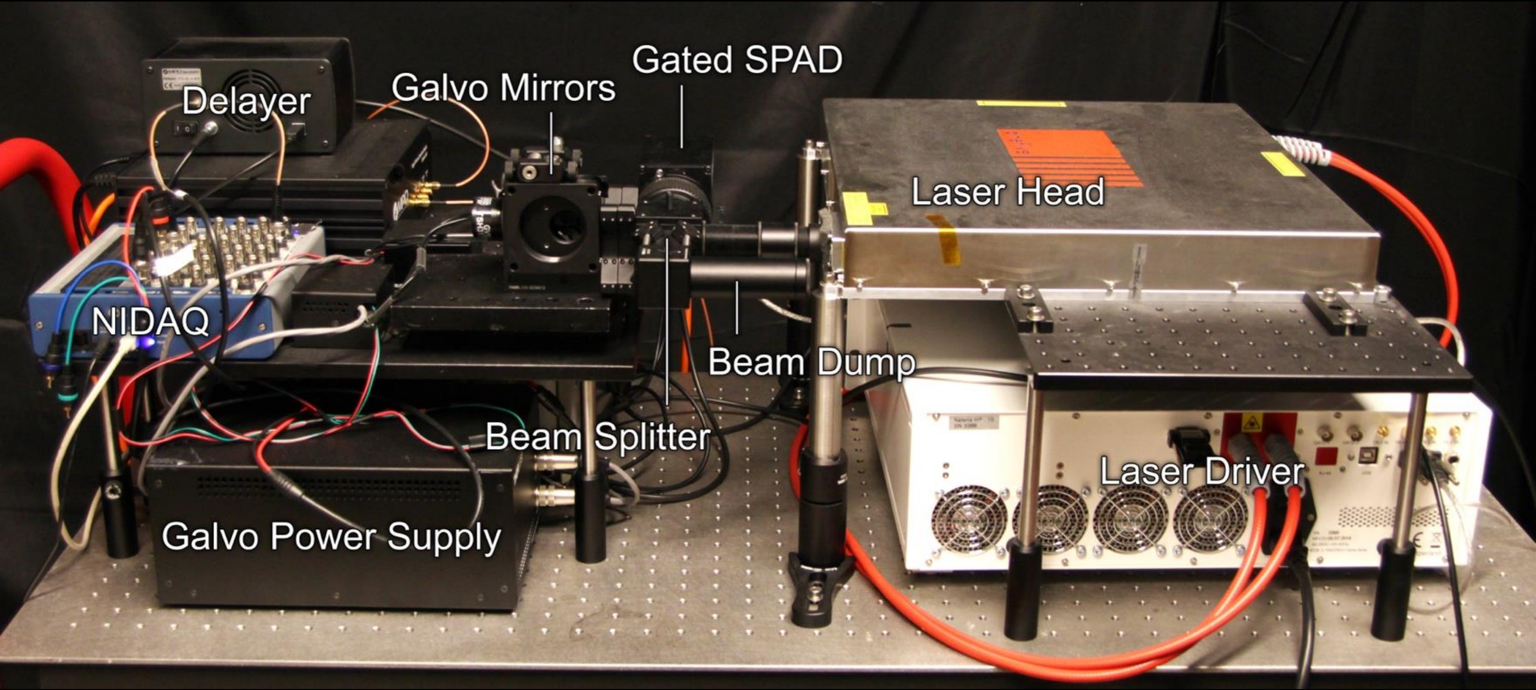

hardware prototype

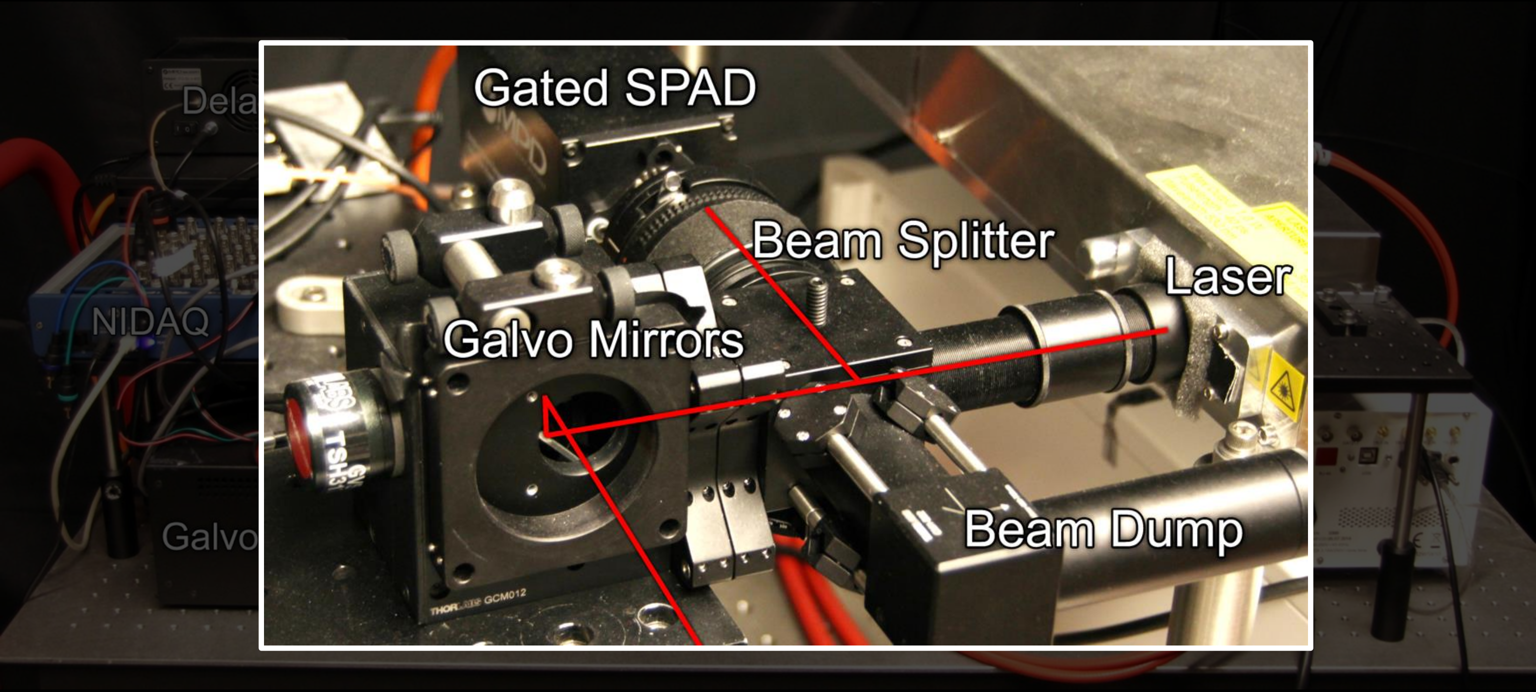

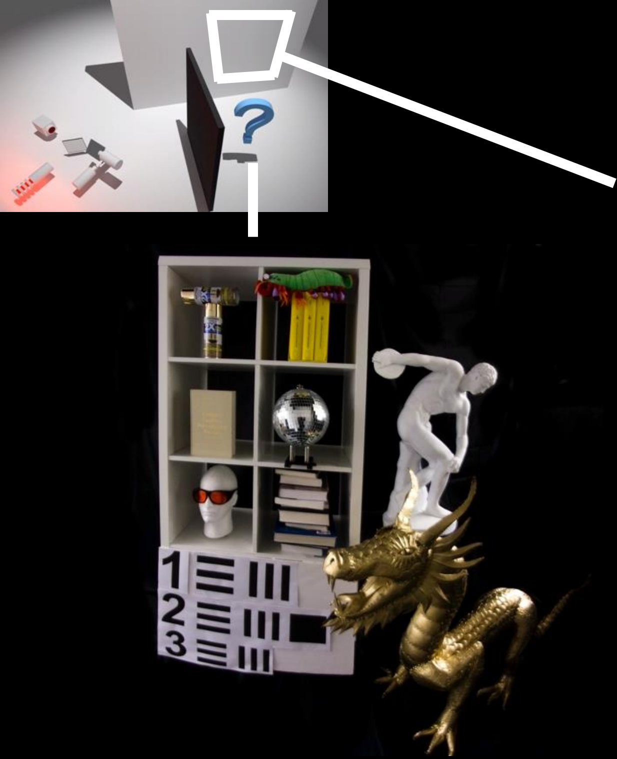

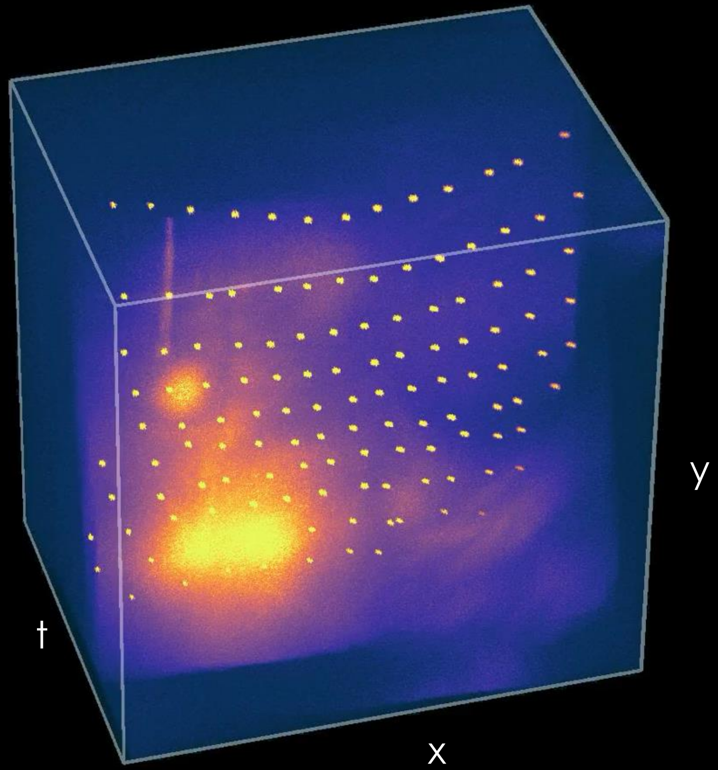

Captured

Lindell et al., SIGGRAPH 201

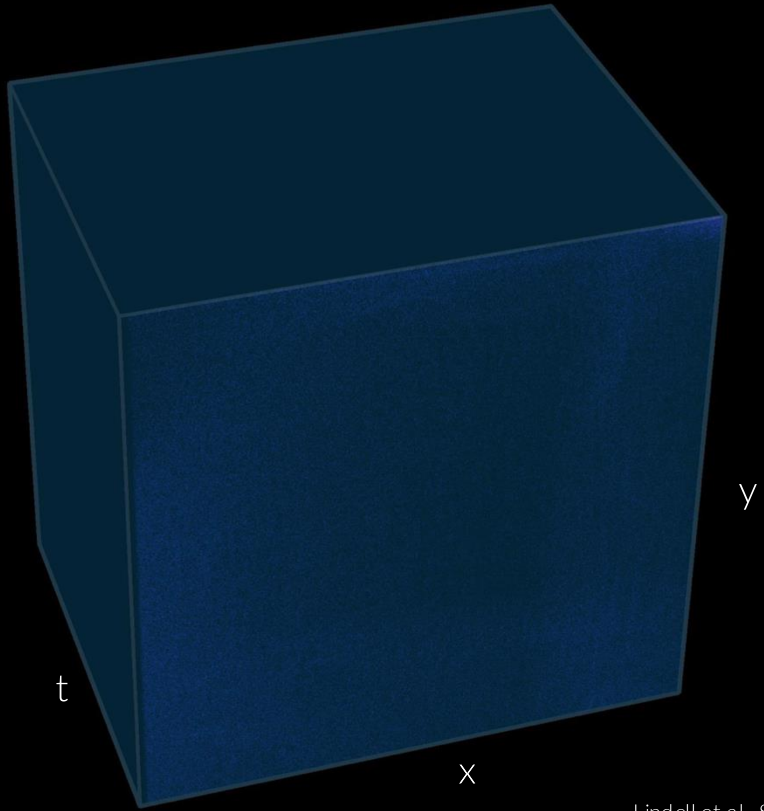

Lindell et al., SIGGRAPH 2019

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)

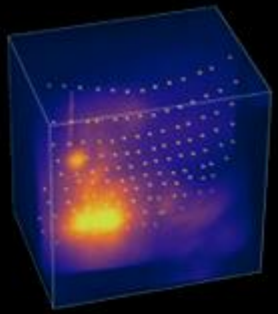

$\mathcal{F}$

Resamp

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

f-k Migration

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

$$\underbrace{\Psi(x, y, z, t)}_{\text{wavefield}} = \iiint \underbrace{\bar{\Phi}(k_x, k_y, f)}_{\text{Fourier transform of measurements}} e^{2\pi i(k_x x + k_y y + k_z z - ft)} dk_x dk_y df$$

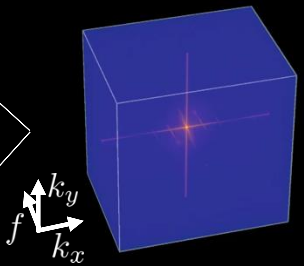

Set t=0 to get migrated solution

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

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

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

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

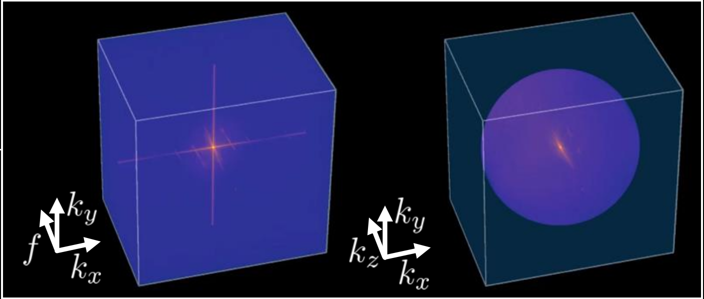

Use dispersion relation[1] to perform substitution of variables

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

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

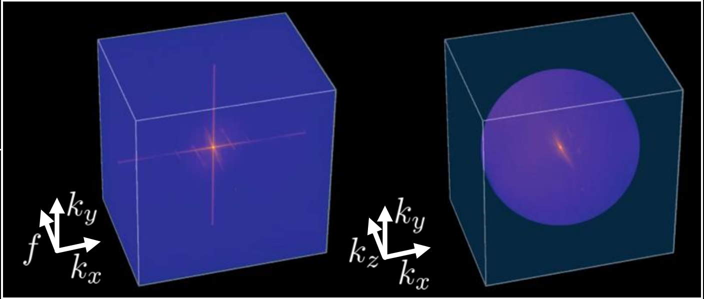

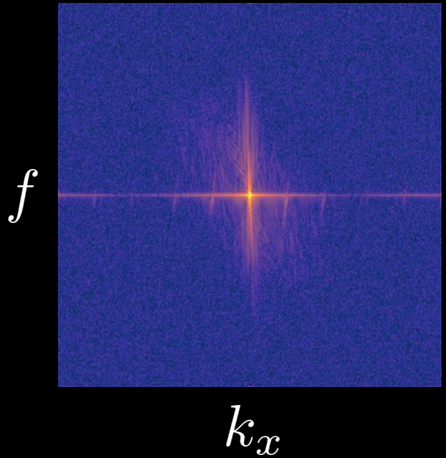

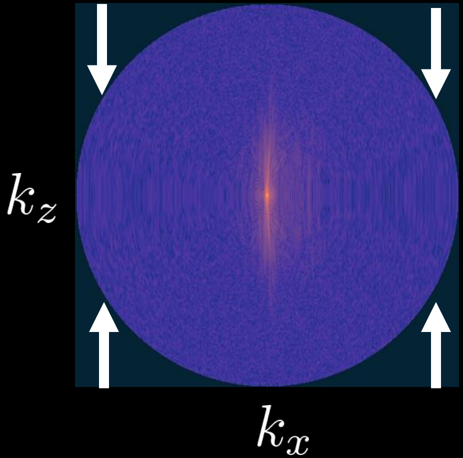

Use dispersion relation[1] to perform substitution of variables

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

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

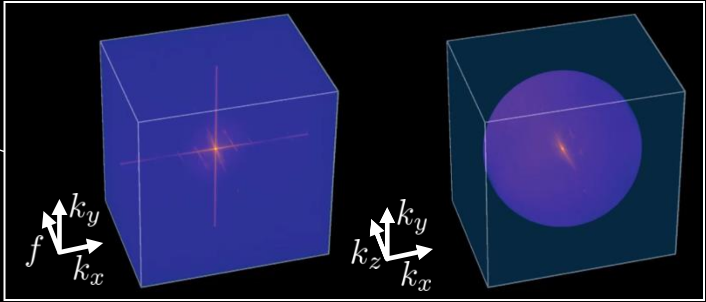

The migrated solution is an inverse Fourier Transform!

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

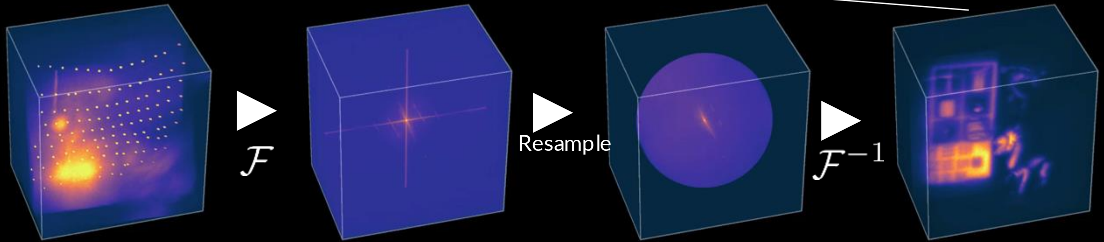

f–k Migration

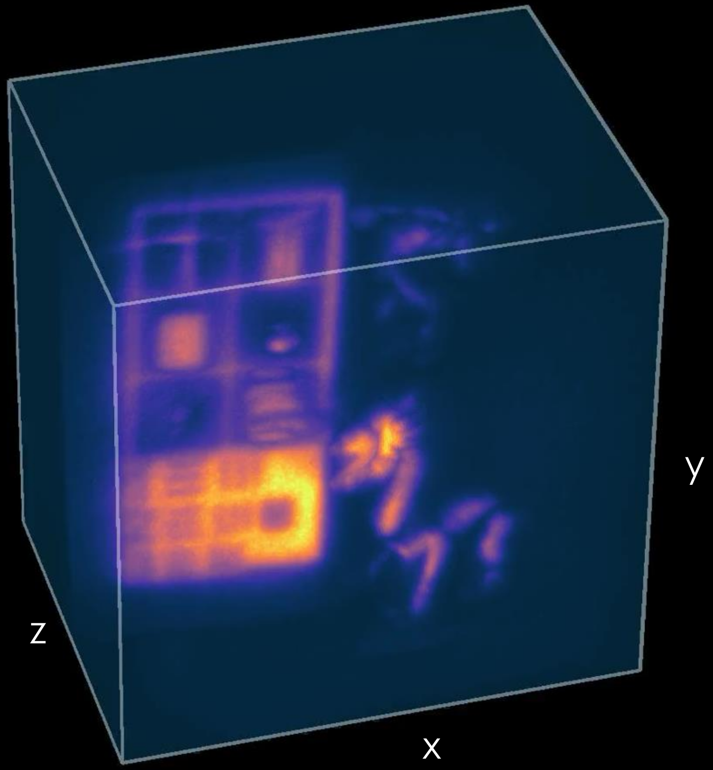

Dimensions: 2 x 2 m
Exposure: 180 min
Reconstruction time: ~90 sec

Lindell et al., SIGGRAPH 201

Reconstruction Comparison

dimensions: 2 m x 2 m x 1.5 m

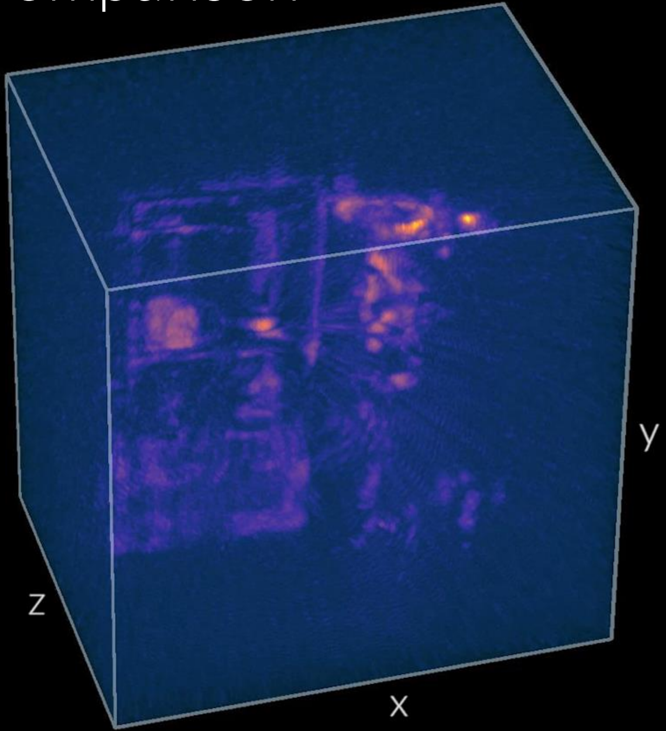

Filtered Backprojection

hardware prototype

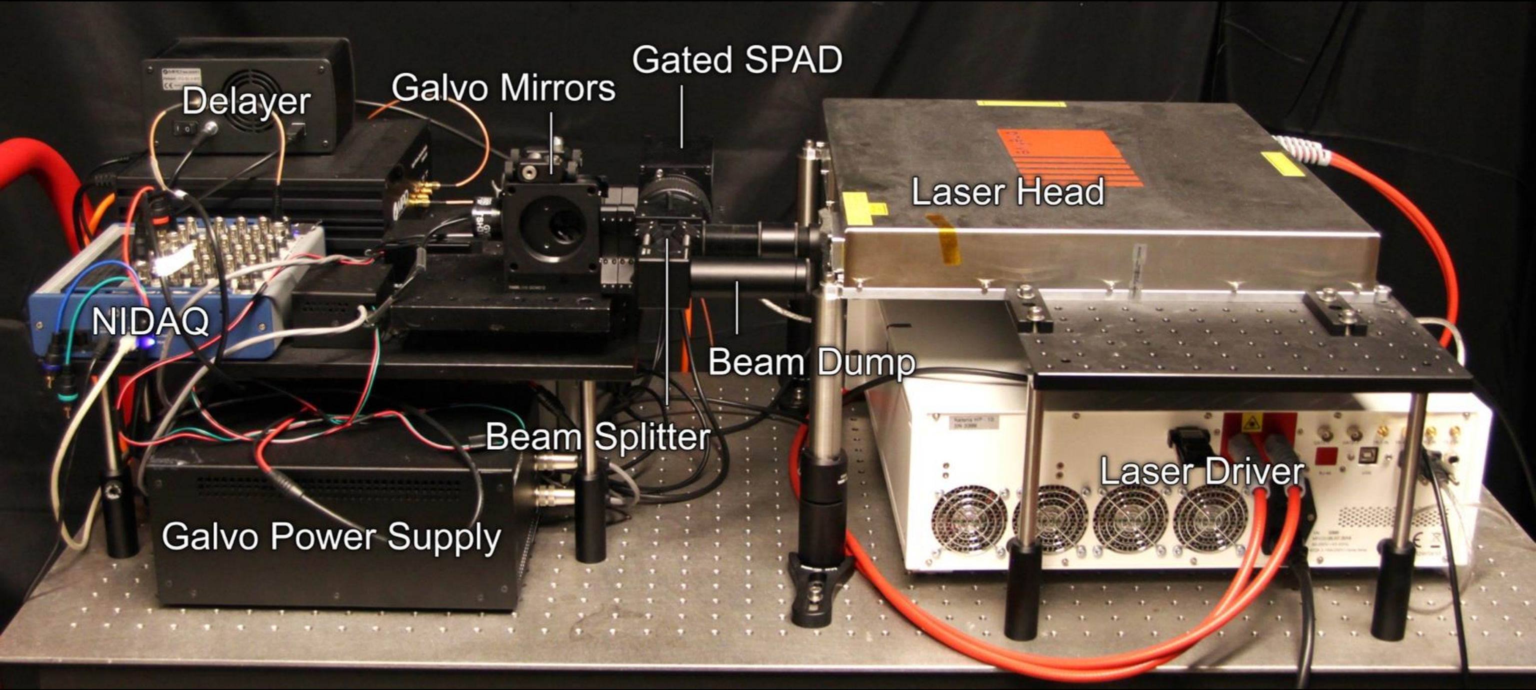

hardware prototype

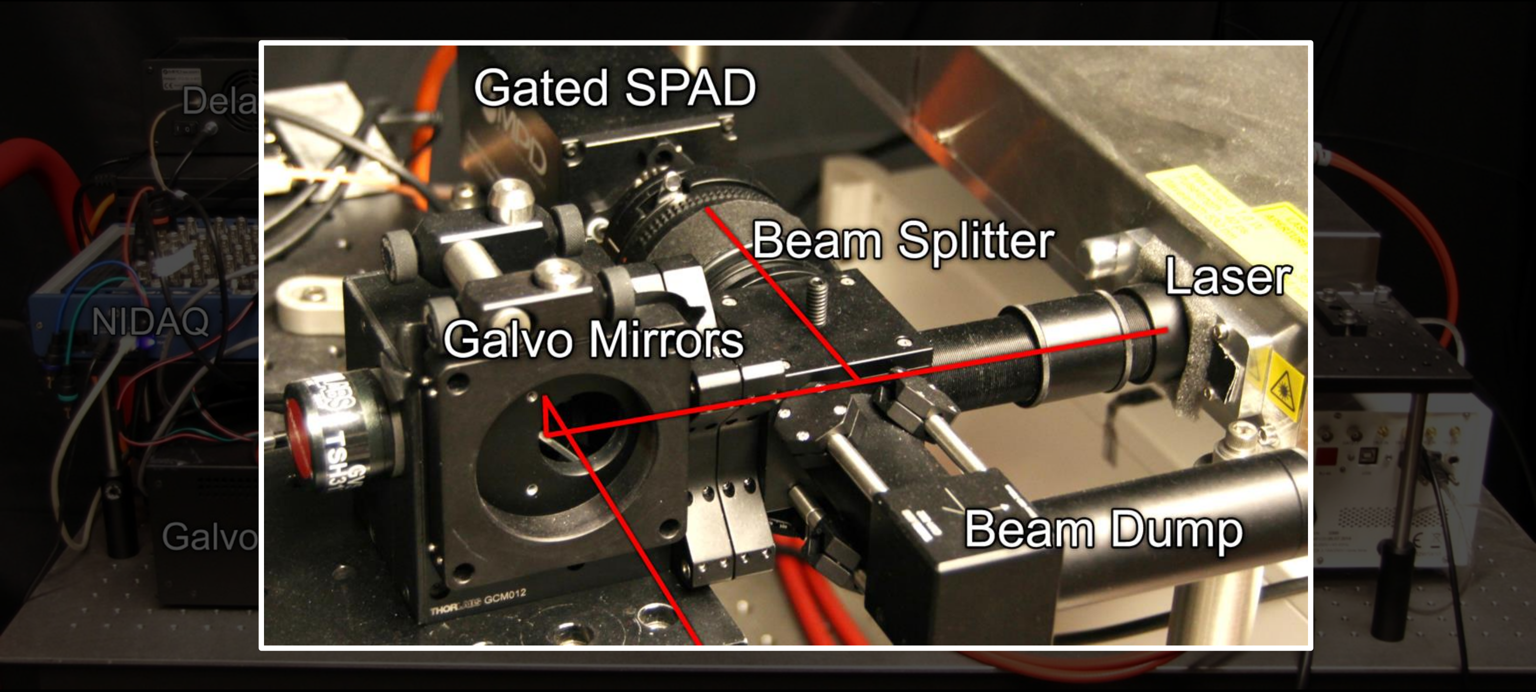

real-time scanning

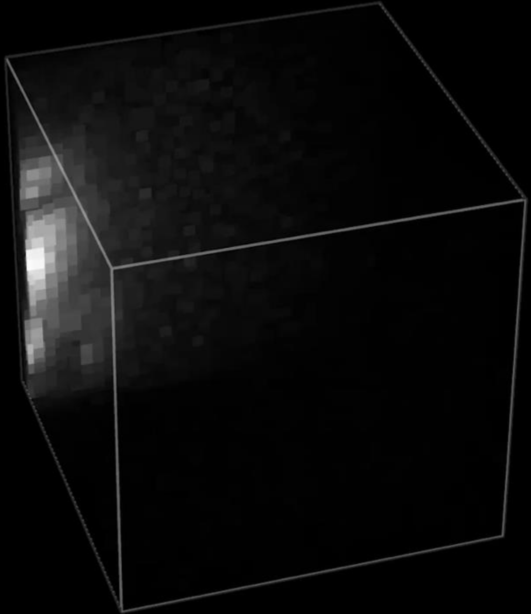

Framerate: 4 Hz
Resolution: 32 x 32
Dimensions: 2 m x 2 m x 2 m
Reconstruction time: ~1 s per frame

Lindell et al., SIGGRAPH 2019

Outlook

Directional Light-Cone Transform

hidden scene

Recovered surface

[Young et al., CVPR 2020]

Outlook

Keyhole NLOS Imaging

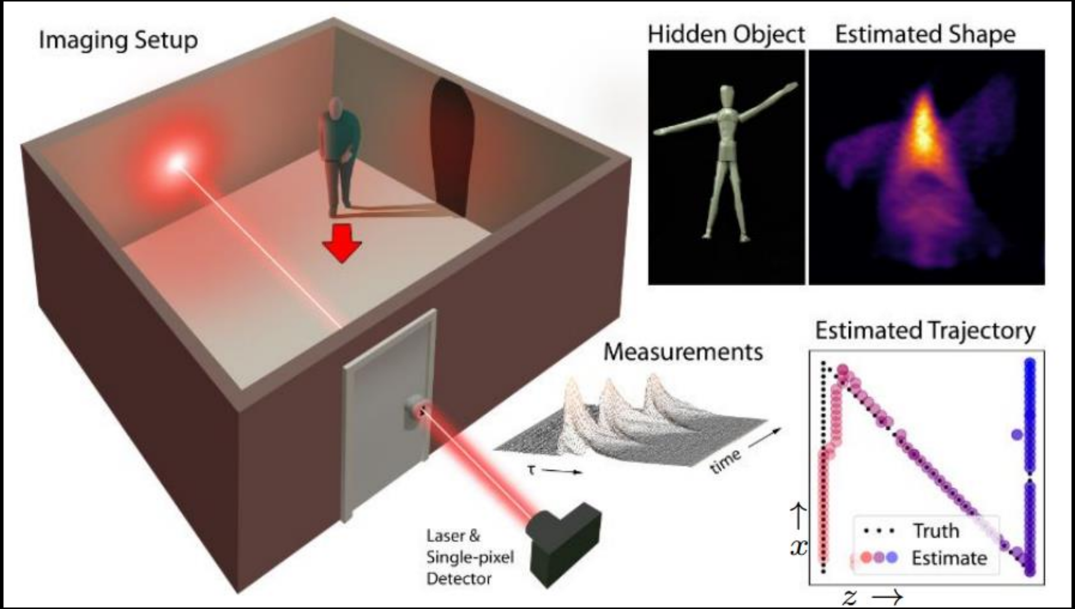

[Metzler et al., IEEE TCI 2021]

overview

- Time-resolved imaging
- Single-photon avalanche diodes (SPADs)
- Single-photon lidar
- Non-line-of-sight imaging
- neural rendering for propagating light

relativity of simultaneity

relativity of simultaneity

relativity of simultaneity

relativity of simultaneity

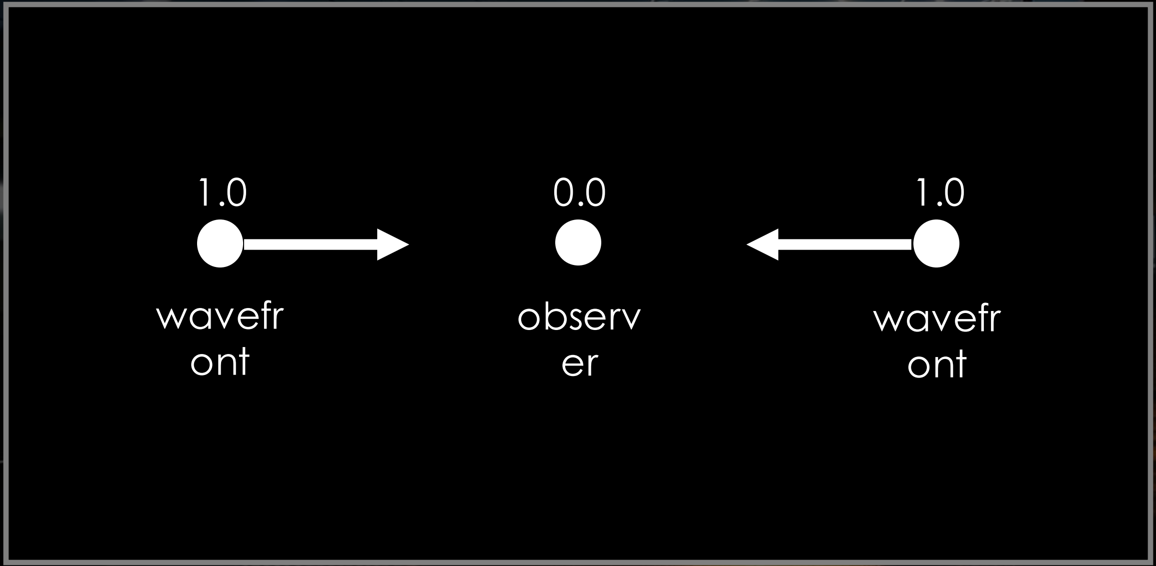

relativity of simultaneity

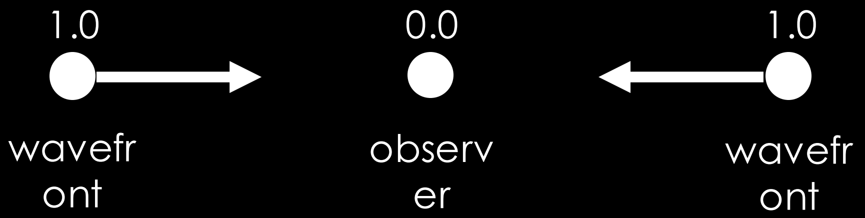

appear to arrive simultaneously

relativity of simultaneity

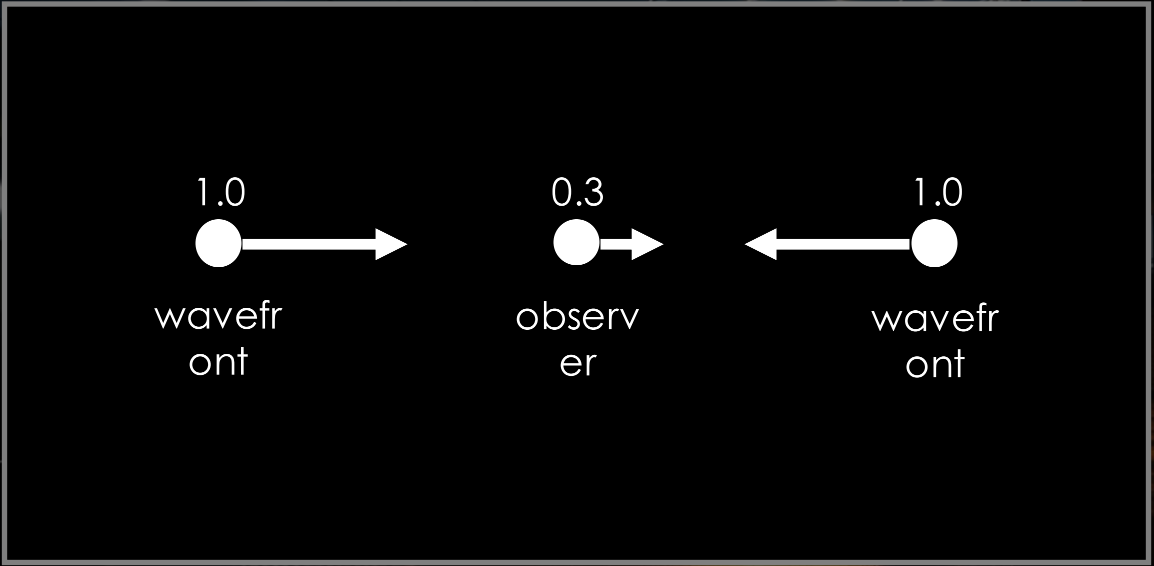

relativity of simultaneity

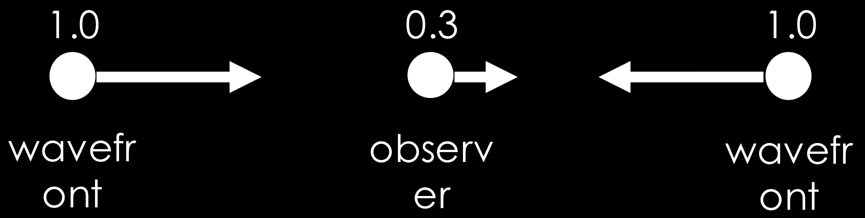

wavefront to the right closes faster—right event happens first!

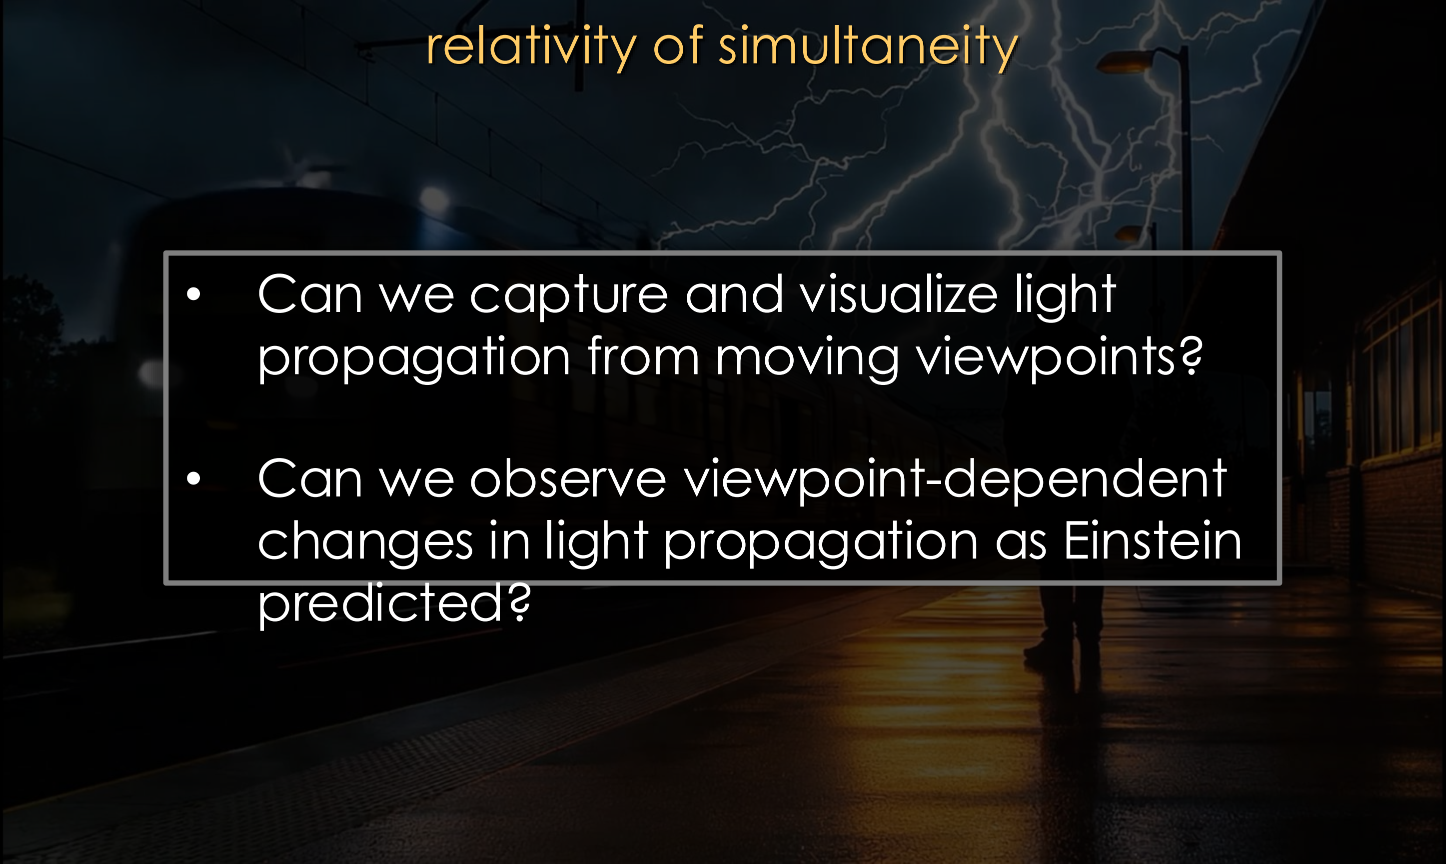

relativity of simultaneity

- Can we capture and visualize light propagation from moving viewpoints?
- Can we observe viewpoint-dependent changes in light propagation as Einstein predicted?

"transient" videography at 250 billion frames per second

[Malik et al. '2

“transient” videography at 250 billion frames per second

pulsed
laser

[Malik et al. ‘2

"transient" videography at 250 billion frames per second

pulsed laser

single-photon avalanche diode (SPAD)

[Malik et al. '2

“transient” videography at 250 billion frames per second

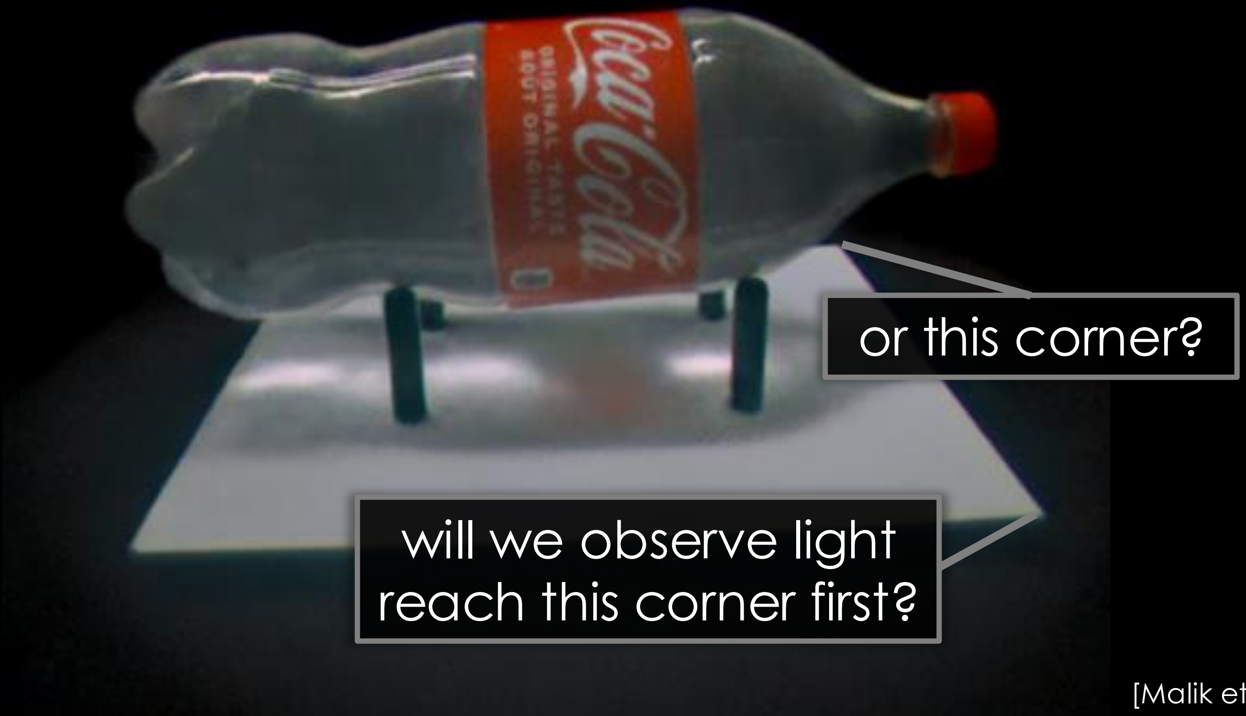

[Malik et al. ‘2

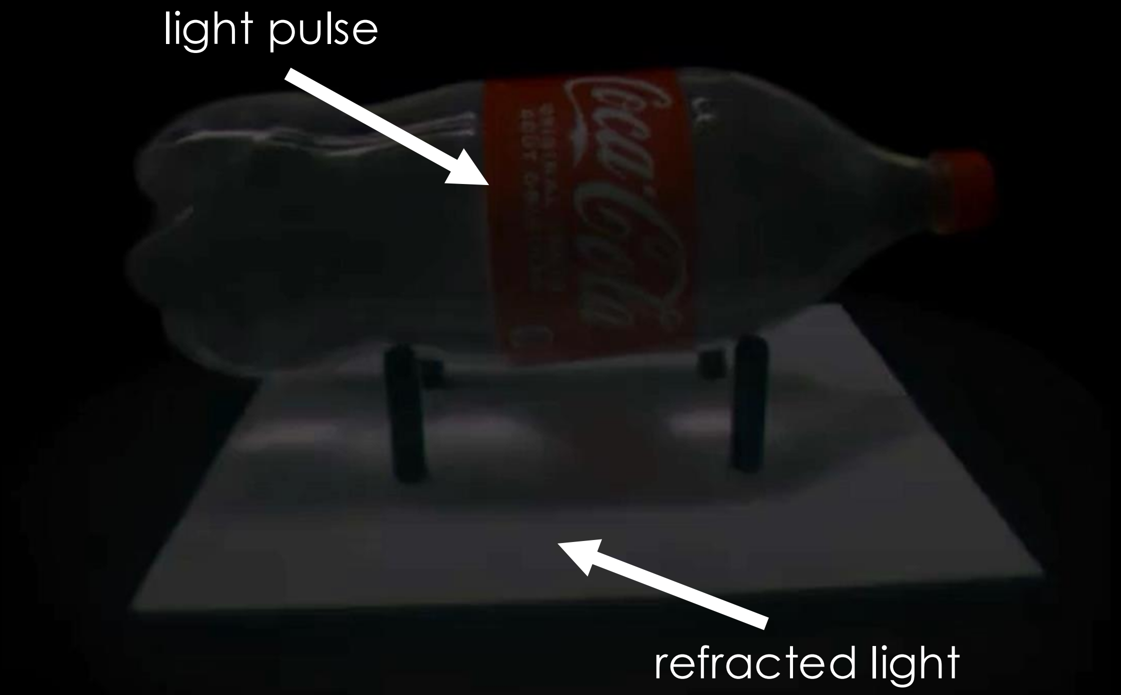

[Malik et al. '2

light focused at ca

[Malik et al. '2

[Malik et al. '2

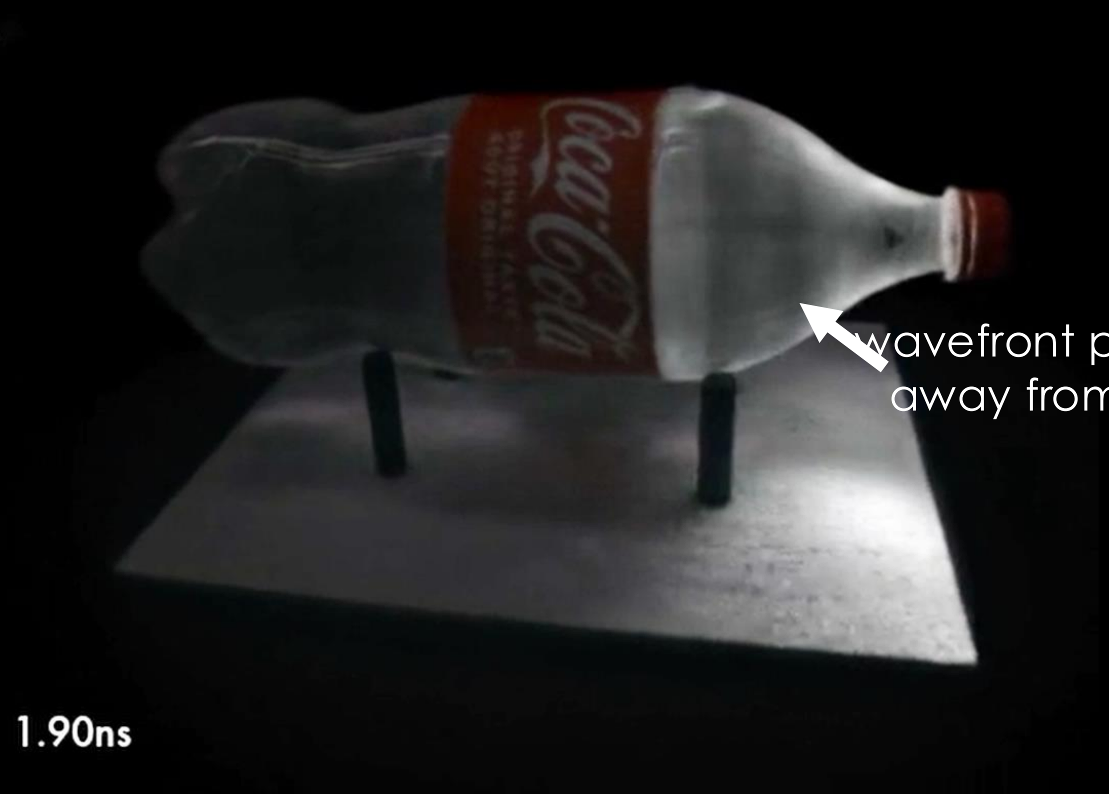

[Malik et al. ‘2

[Malik et al. ‘2

[Malik et al. ‘2

[Malik et al. ‘2

shortest path length from bottle to camera

[Malik et al. ‘2

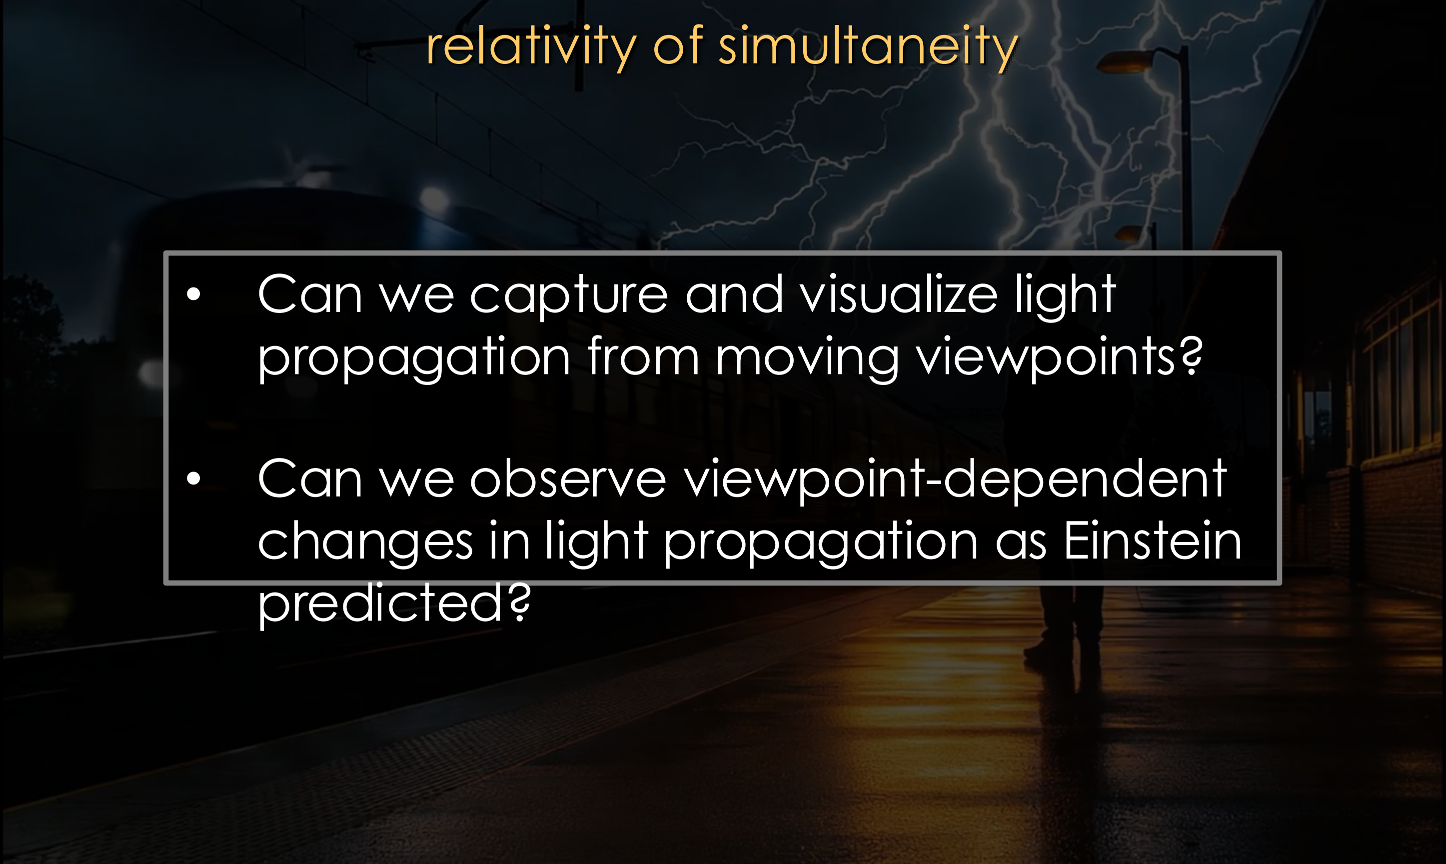

relativity of simultaneity

- Can we capture and visualize light propagation from moving viewpoints?
- Can we observe viewpoint-dependent changes in light propagation as Einstein predicted?

relativity of simultaneity

Yes! (and no)

- we can’t move at relativistic speeds

relativity of simultaneity

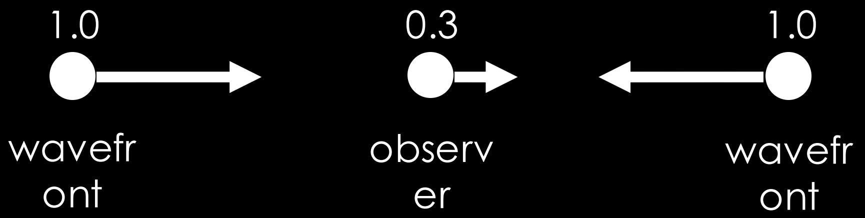

does the right wavefront appear to travel 1.3c?

relativity of simultaneity

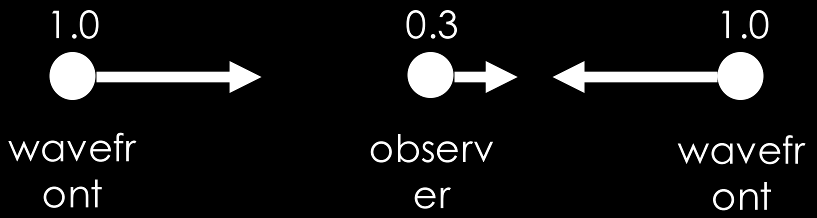

No! space contracts so that the speed of light is maintained
(Lorentz contraction)

relativity of simultaneity

Yes! (and no)

- we can’t move at relativistic speeds
- we can capture and visualize propagation of light, including viewpoint-dependent effects

related work

related work

transient imaging

streak cameras

[Velten '13, Gao '14, …]

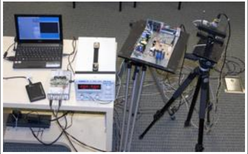

time-of-flight cameras

[Heide '13, O'Toole '14, …]

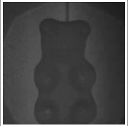

interferometry

[Gkioulekas '15, …]

SPADs

[O'Toole '17, Lindell, '18, …]

related work

transient imaging

streak
cameras
[Velten '13, Gao '14,
…]

time-of-flight cameras
[Heide '13, O'Toole '14, …]

interferome
try
[Gkioulekas '15,
…]

SPADs
[O'Toole '17, Lindell, '18,
…]

- limited to single-viewpoint capture

related work

transient imaging

streak cameras

[Velten '13, Gao '14, …]

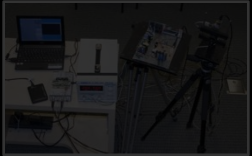

time-of-flight cameras

[Heide '13, O'Toole '14, …]

interferometry

[Gkioulekas '15, …]

SPADs

[O'Toole '17, Lindell, '18, …]

- limited to single-viewpoint capture

NeRFs & video novel view synthesis

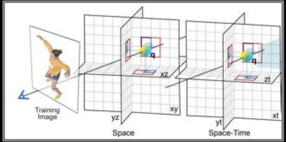

[Li '22]
[Cao & Johnson '23]
[Fridovich-Keil '23]
[Wang '23]

related work

transient imaging

streak
cameras
[Velten '13, Gao '14,
…]

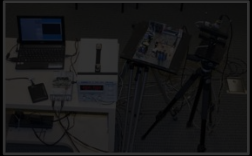

time-of-flight cameras
[Heide '13, O'Toole '14, …]

interferome
try
[Gkioulekas '15,
…]

SPADs
[O'Toole '17, Lindell, '18,
…]

- limited to single-viewpoint capture

NeRFs & video novel view synthesis

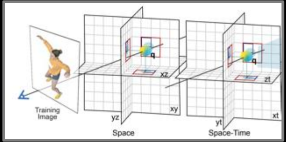

[Li '22]
[Cao & Johnson '23]
[Fridovich-Keil '23]
[Wang '23]

- do not account for finite speed of light

related work

transient imaging

streak cameras
[Velten '13, Gao '14, …]

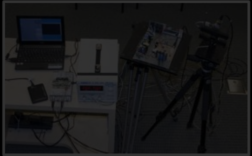

time-of-flight cameras
[Heide '13, O'Toole '14, …]

interferometry
[Gkioulekas '15, …]

SPADs
[O'Toole '17, Lindell, '18, …]

- limited to single-viewpoint capture

NeRFs & video novel view synthesis

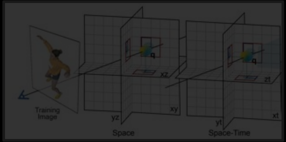

[Li '22]
[Cao & Johnson '23]
[Fridovich-Keil '23]
[Wang '23]

- do not account for speed of light

Transient NeRF

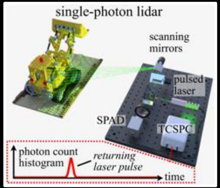

[Malik '23]

related work

transient imaging

streak cameras
[Velten '13, Gao '14, …]

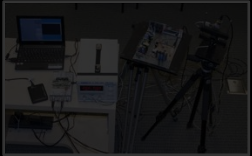

time-of-flight cameras
[Heide '13, O'Toole '14, …]

interferometry
[Gkioulekas '15, …]

SPADs
[O'Toole '17, Lindell, '18, …]

- limited to single-viewpoint capture

NeRFs & video novel view synthesis

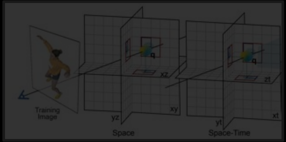

[Li '22]
[Cao & Johnson '23]
[Fridovich-Keil '23]
[Wang '23]

- do not account for speed of light

Transient NeRF

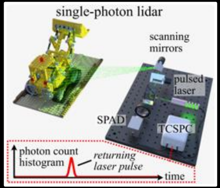

[Malik '23]

- only models direct reflections

related work

transient imaging

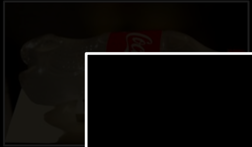

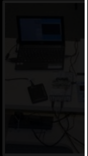

contributions

[Gkioulekas '15, …]

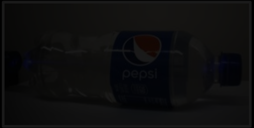

SPADs

[O'Toole '17, Lindell, '18, …]

- limited to single-viewpoint capture

video novel view synthesis

[Malik '23]

- only models direct reflections

related work

contributions

- new method for novel view synthesis of light propagation videos

transient imaging

video novel view synthesis

[Gkioulekas '15, ...]

SPADs

[O'Toole '17, Lindell, '18, ...]

[Malik '23]

- limited to single-viewpoint capture

- only models direct reflections

related work

transient imaging

video novel view synthesis

contributions

- new method for novel view synthesis of light propagation videos
- introduce a new neural field parameterization to make this possible

[Gkioulekas '15, …]

SPADs

[O'Toole '17, Lindell, '18, …]

[Malik '23]

- limited to single-viewpoint capture
- only models direct reflections

related work

transient imaging

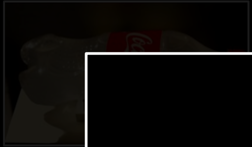

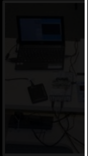

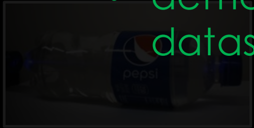

SPADs

[O'Toole '17, Lindell, '18, …]

- limited to single-viewpoint capture

video novel view synthesis

[Malik '23]

- only models direct reflections

contributions

- new method for novel view synthesis of light propagation videos
- introduce a new neural field parameterization to make this possible
- demonstrate on a new multi-viewpoint transient dataset

method

method

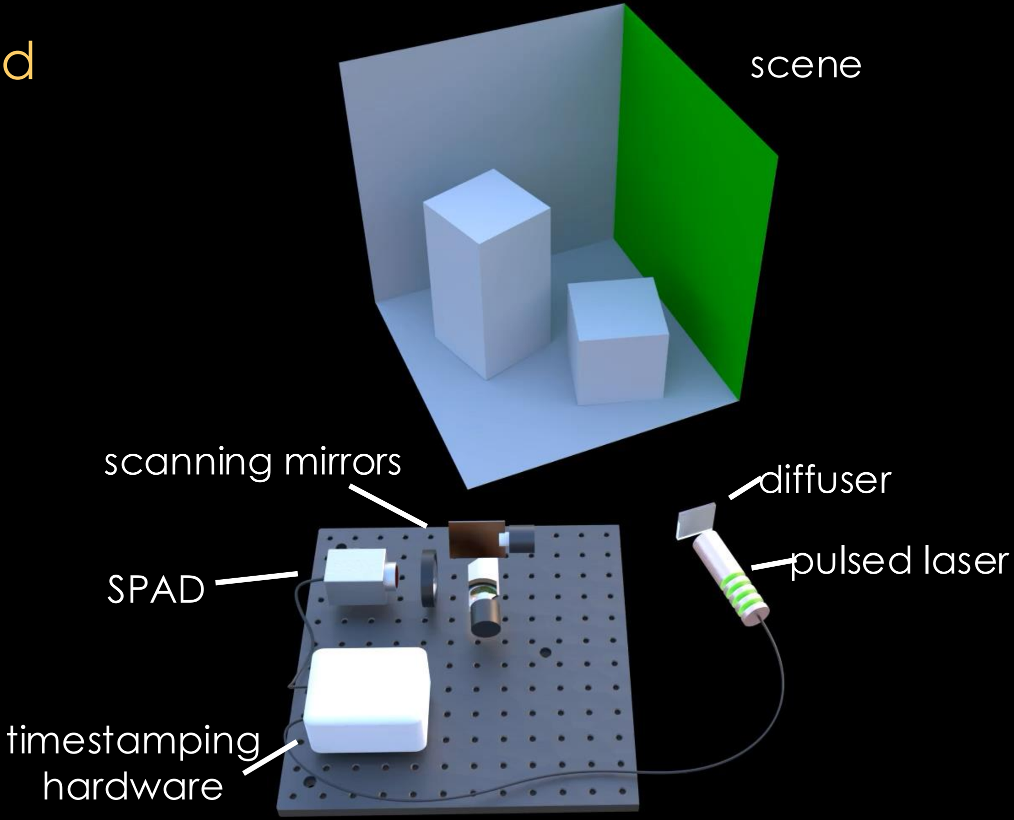
scene
scanning mirrors
diffuser
pulsed laser
SPAD
timestamping
hardware

method

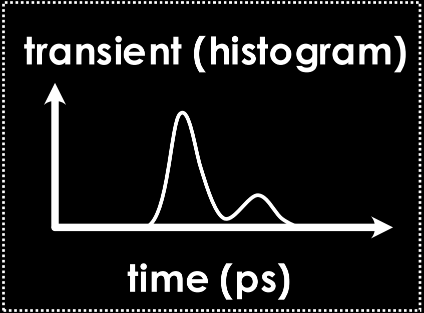
transient (histogram)
time (ps)

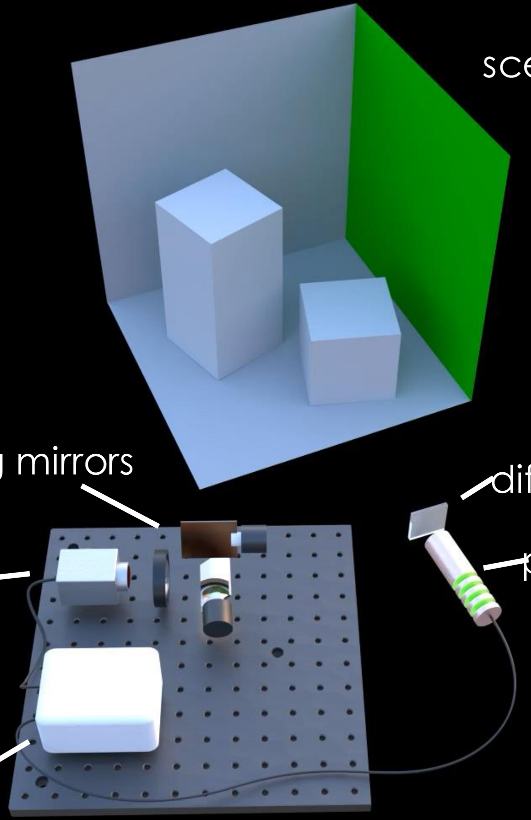
scene
scanning mirrors
diffuser
pulsed laser
SPAD
timestamping hardware

method

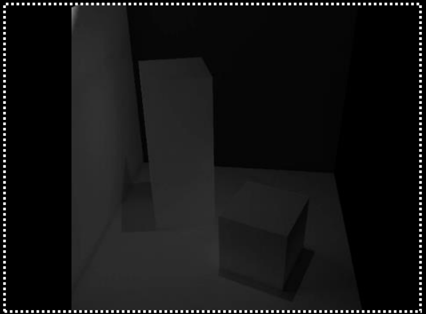

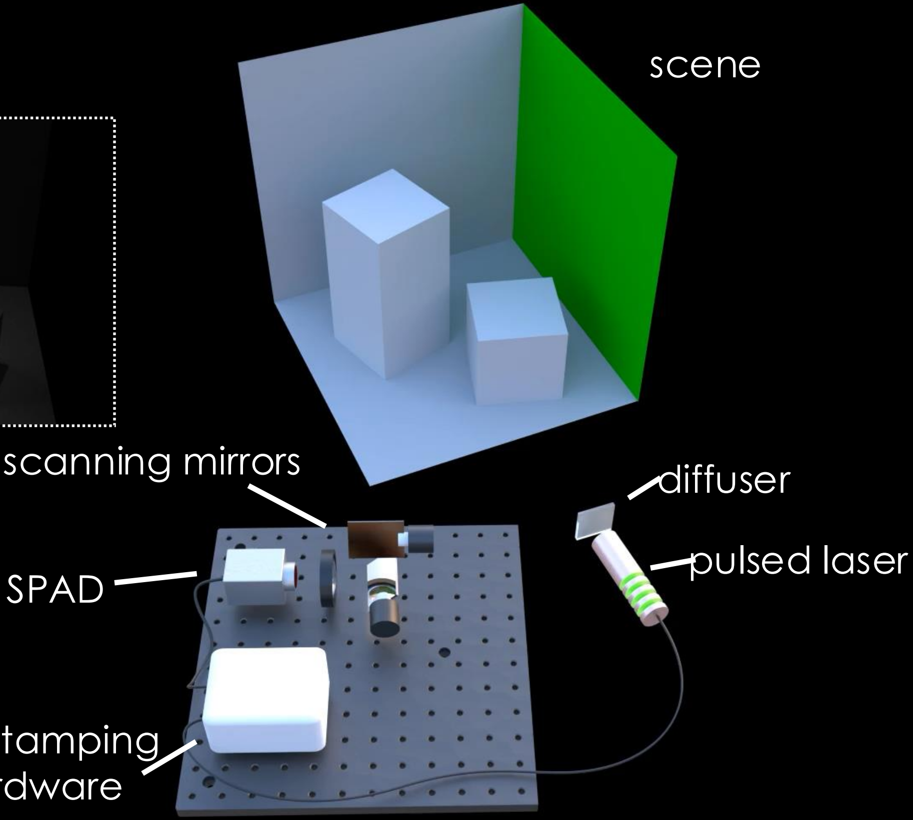
scene
scanning mirrors
diffuser
pulsed laser
SPAD
timestamping
hardware

method

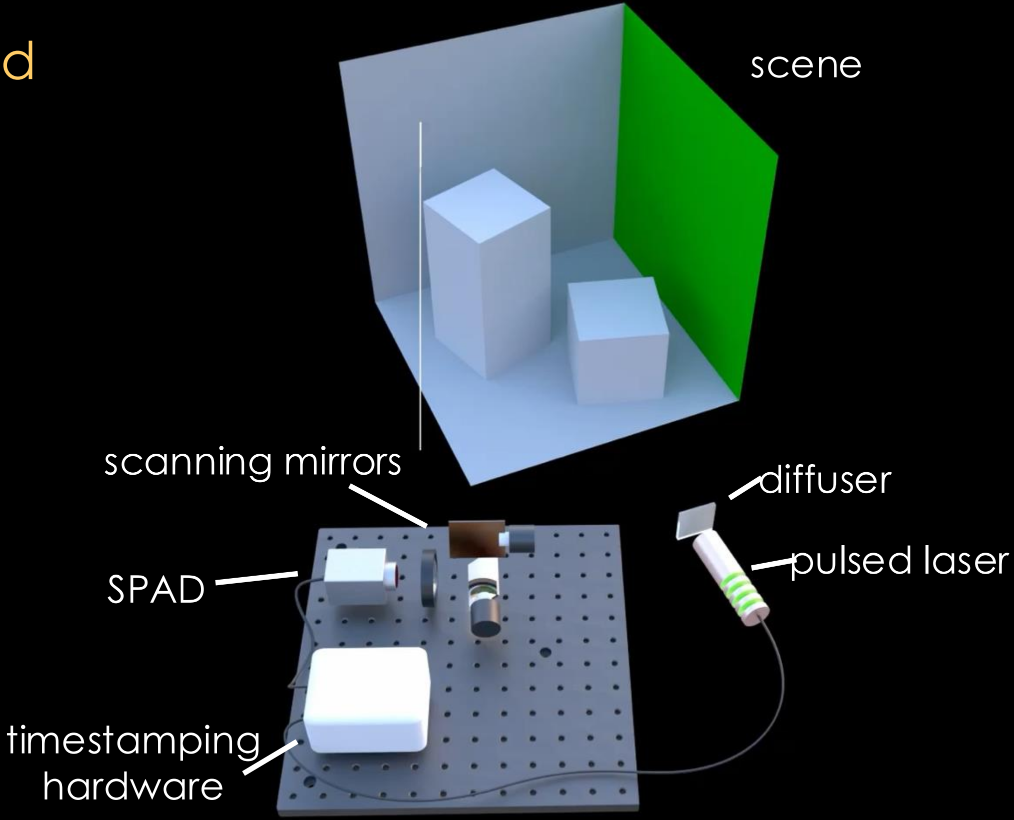
scene
scanning mirrors
diffuser
pulsed laser
SPAD
timestamping
hardware

method

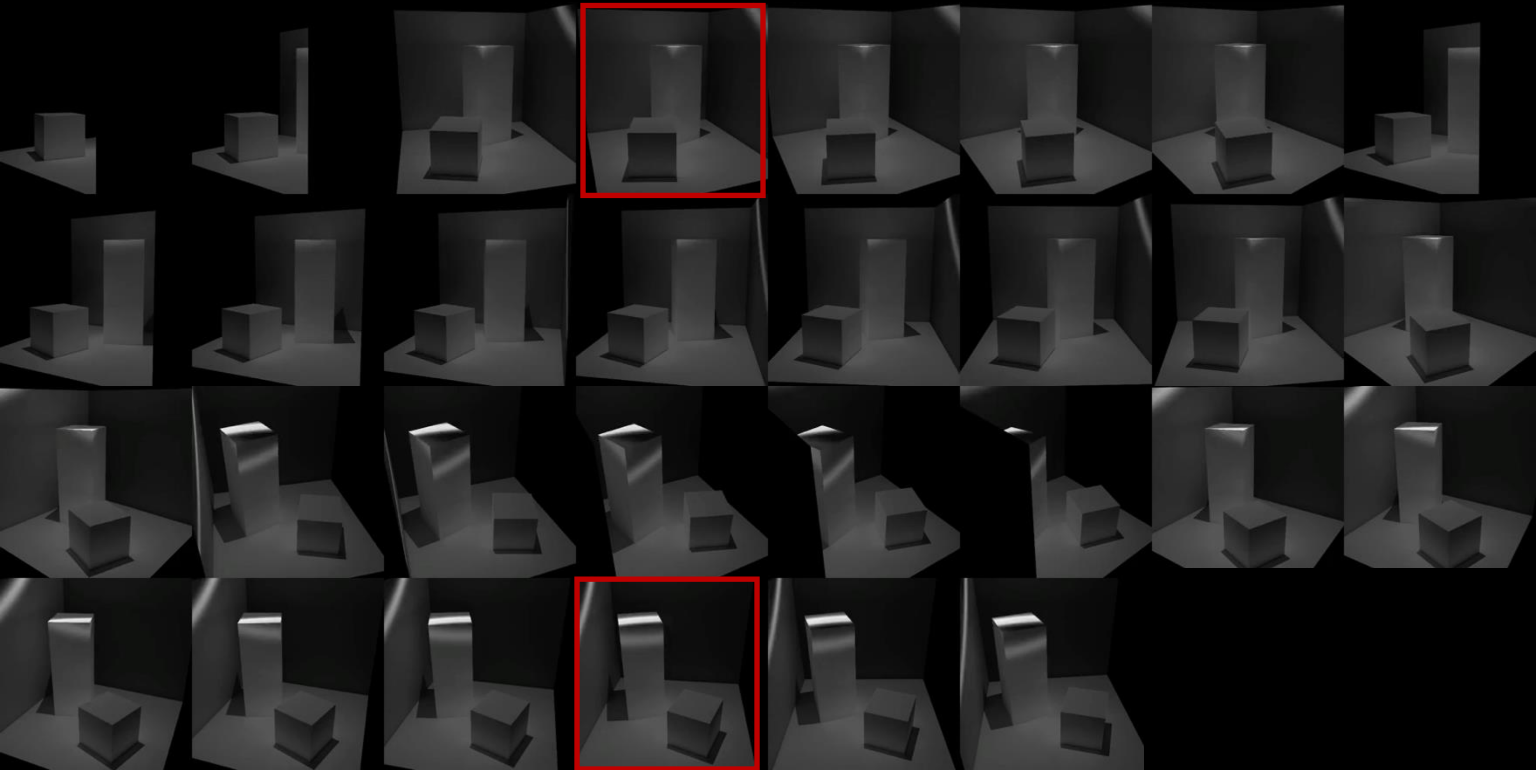

method

transient light transport is viewpoint dependent!

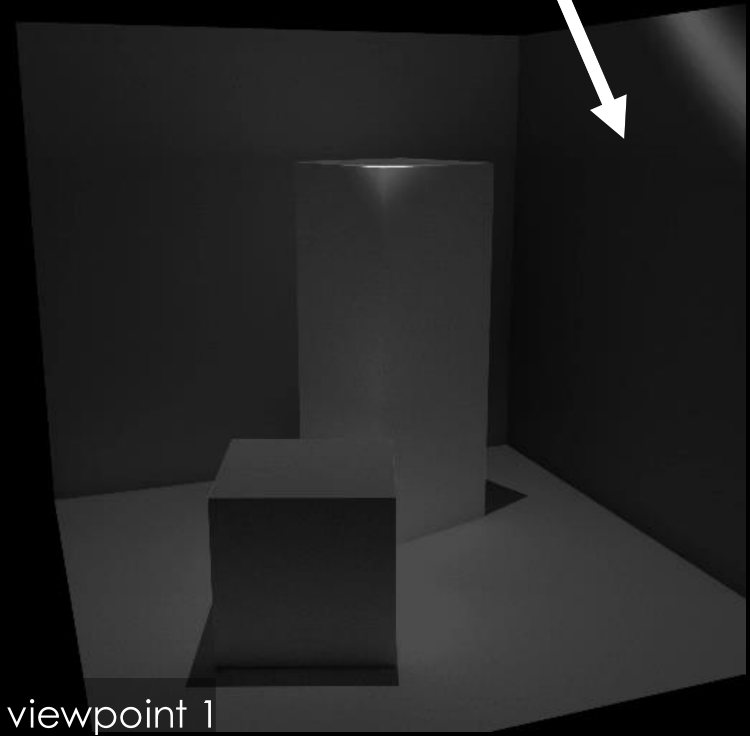

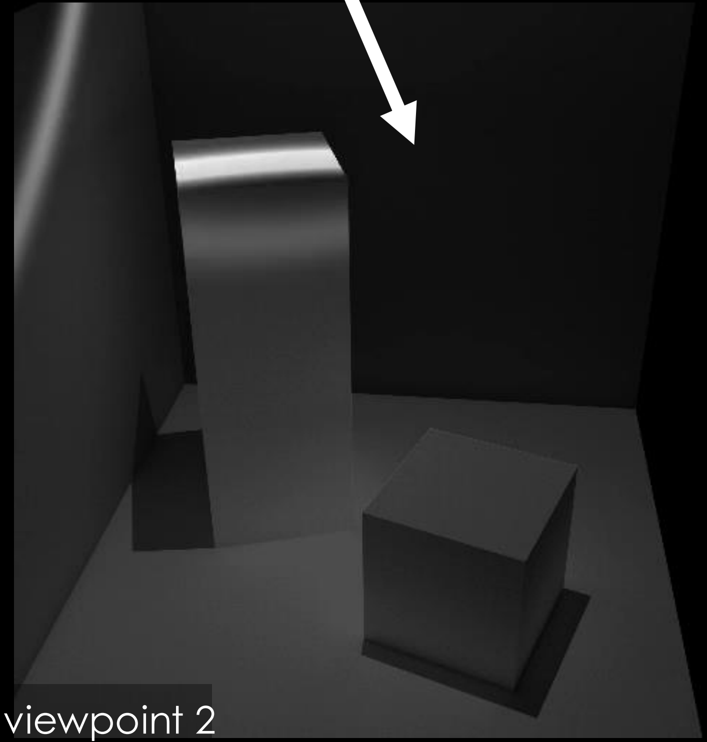

method

- every camera viewpoint has its own temporal reference frame!

method

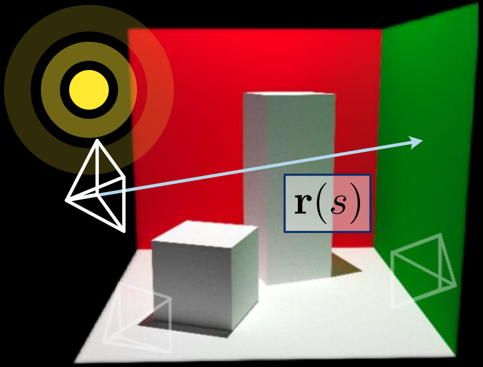

scene

method

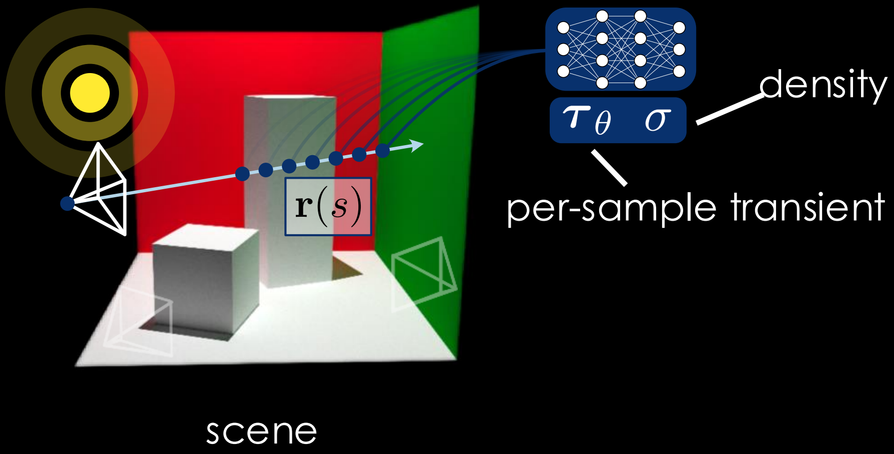

$\mathbf{r}(s)$
τ_θ σ
density
per-sample transient
scene

method

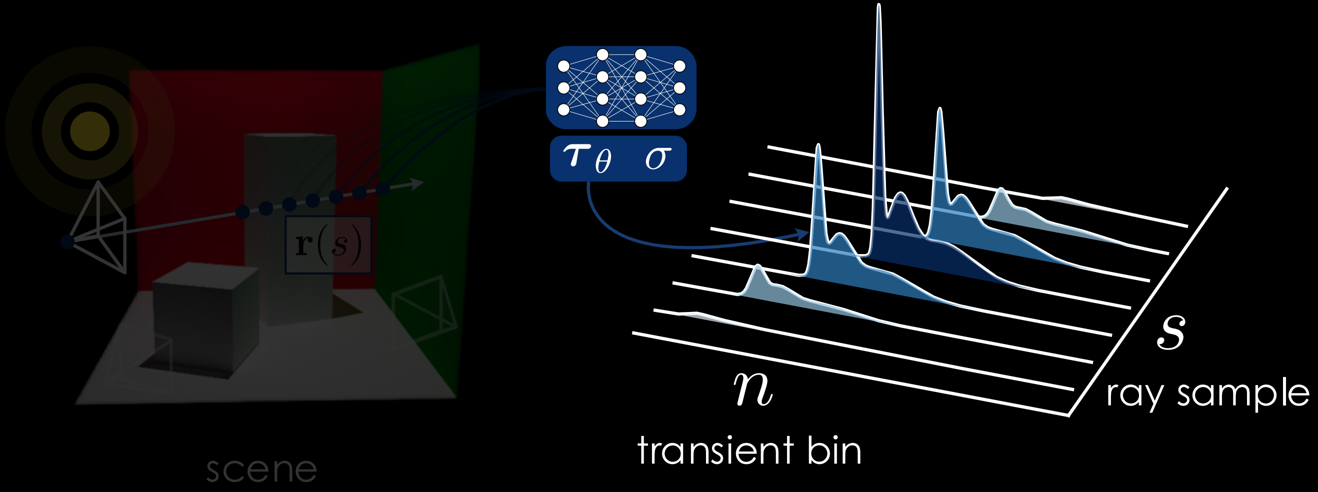

method

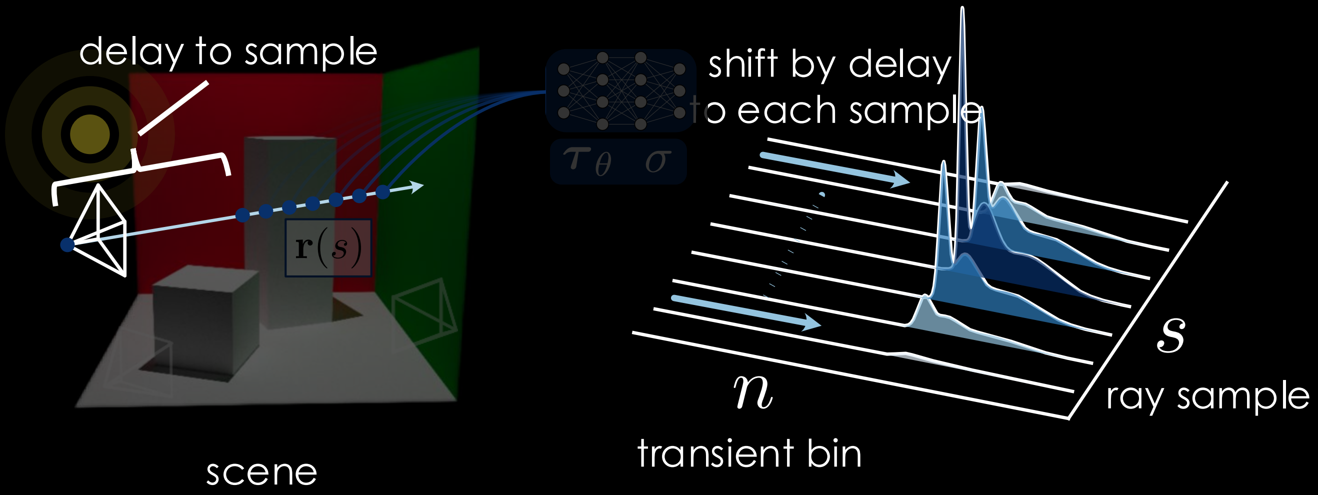

delay to sample
$\mathbf{r}(s)$
τ_θ
σ
shift by delay
to each sample
s
ray sample
n
transient bin
scene

method

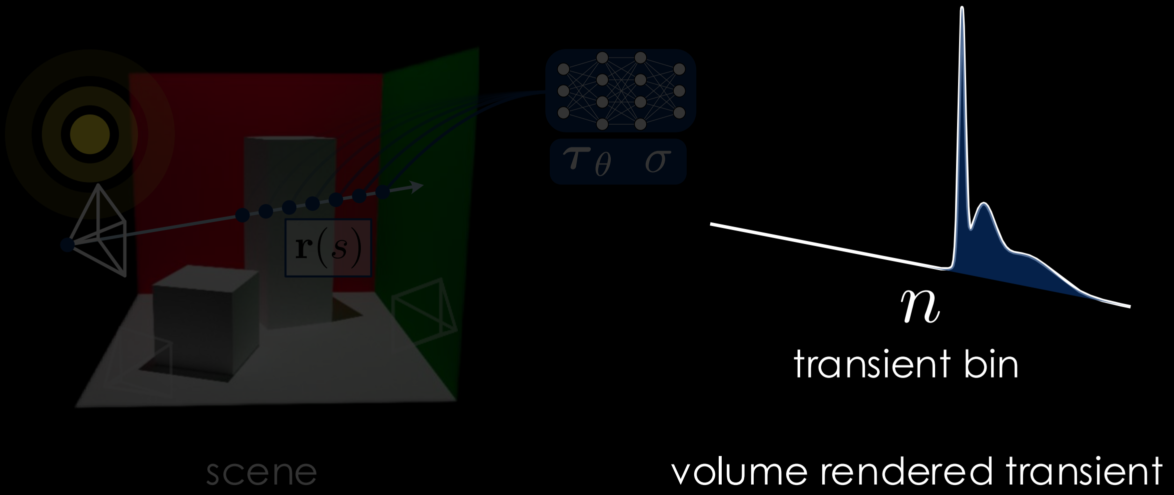

τ_θ σ
$\mathbf{r}(s)$
n
transient bin
scene
volume rendered transient

method

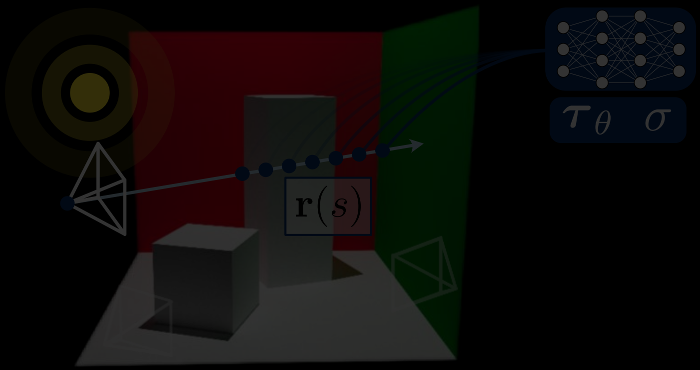

scene

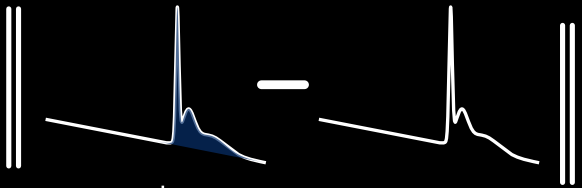

loss function

method

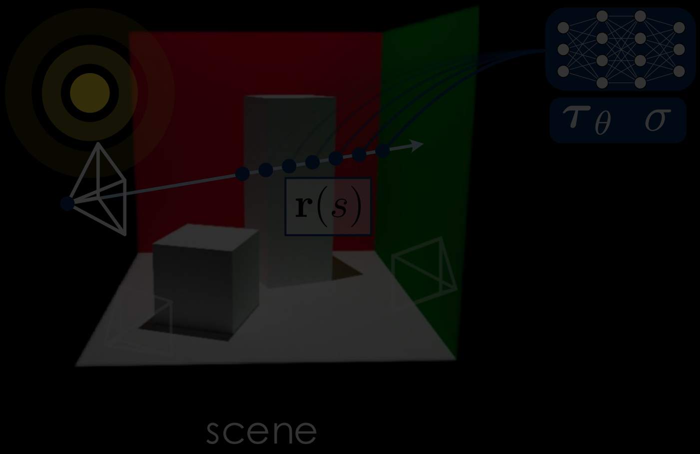

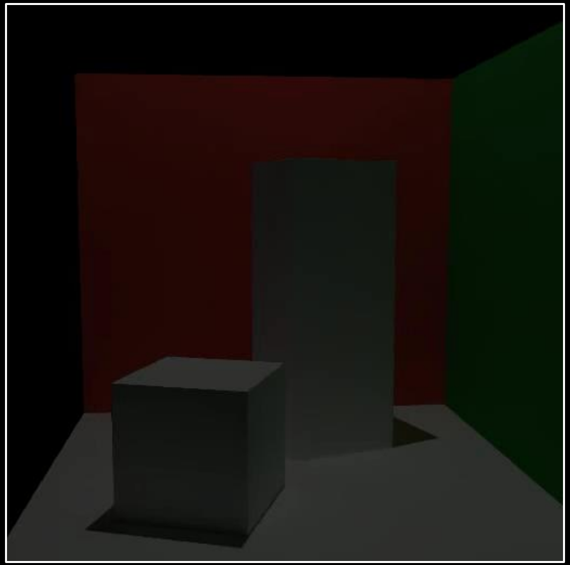

- render from given viewpoint
- linearly increase time

hardware prototype

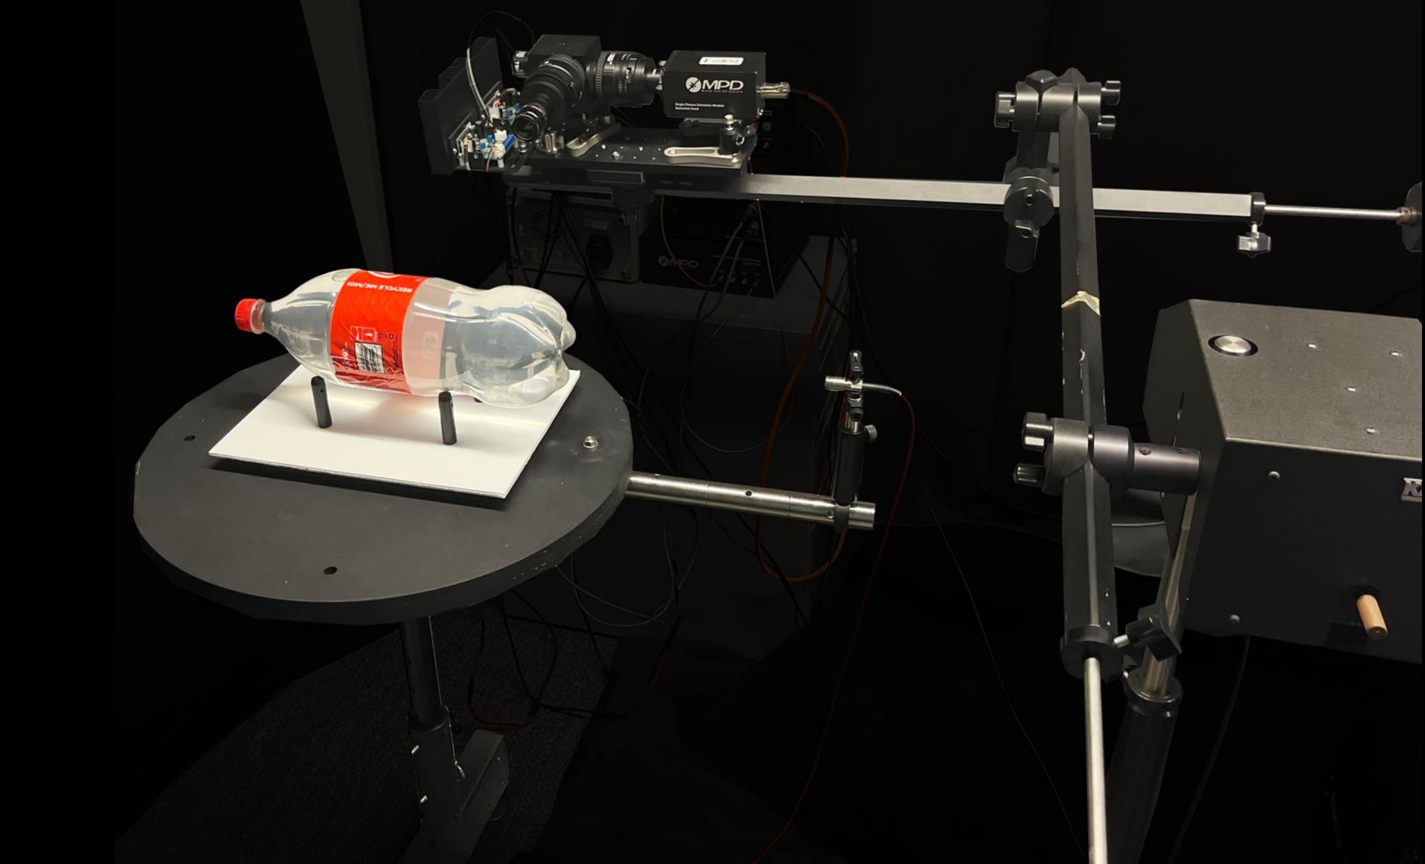
MPD

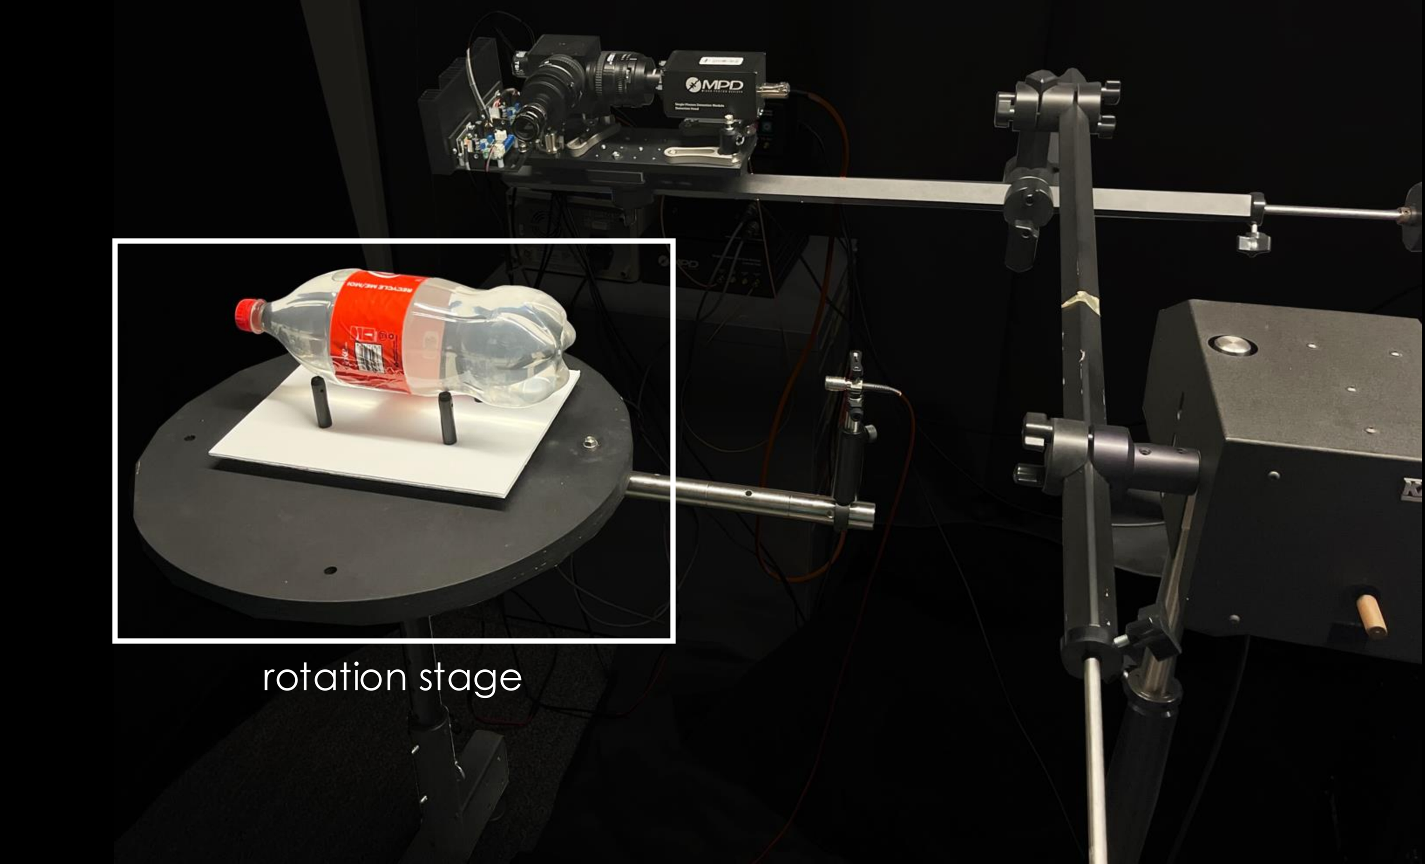
MPD
rotation stage

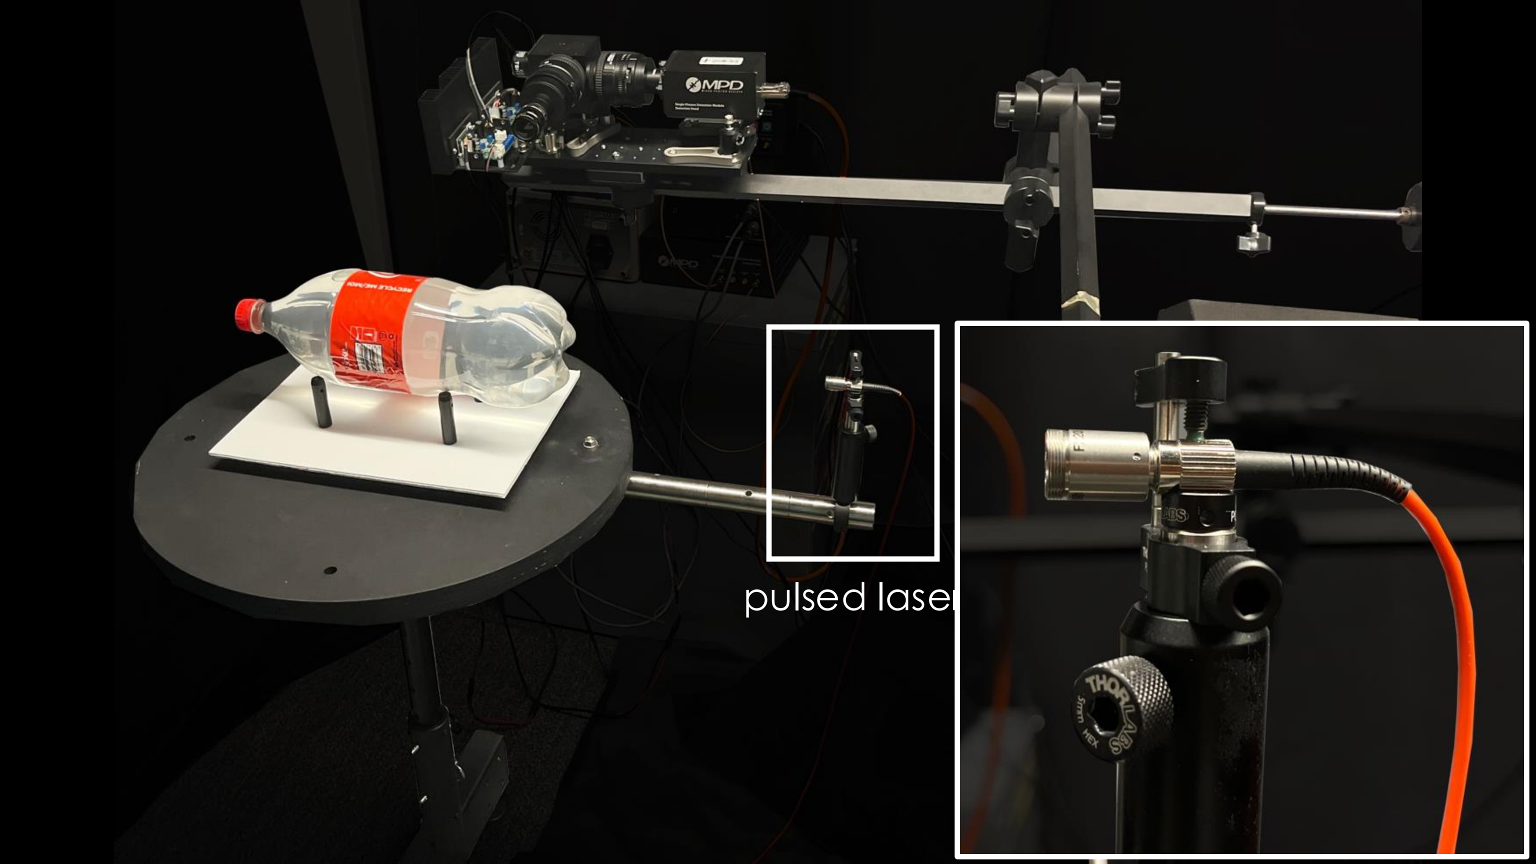

MPD
pulsed laser

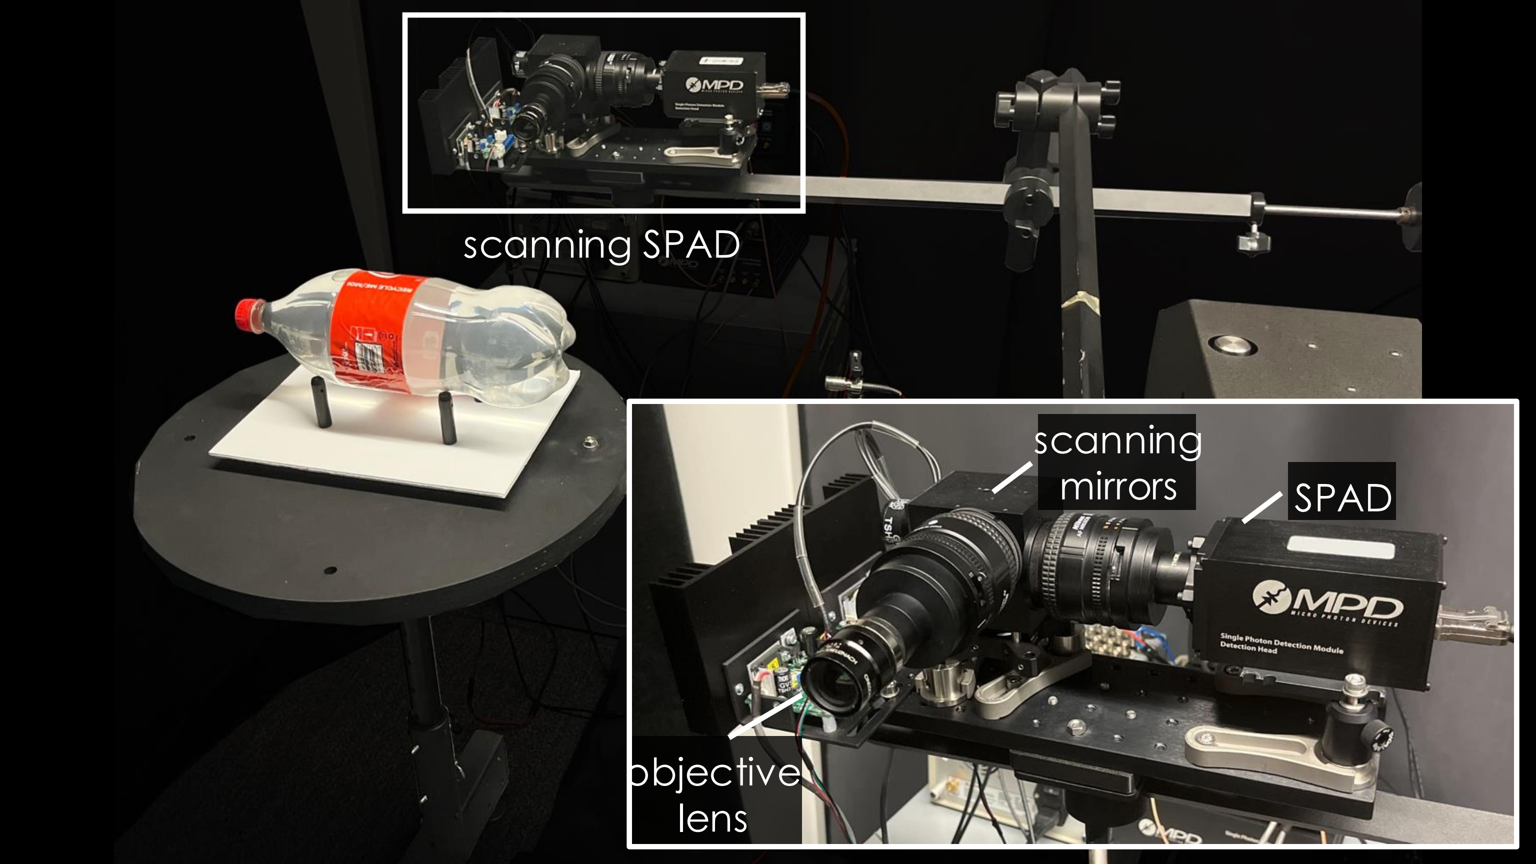

scanning SPAD
scanning mirrors
SPAD
objective lens
MPD
MICRO PHOTON DEVICES
Single Photon Detection Module
Detection Head

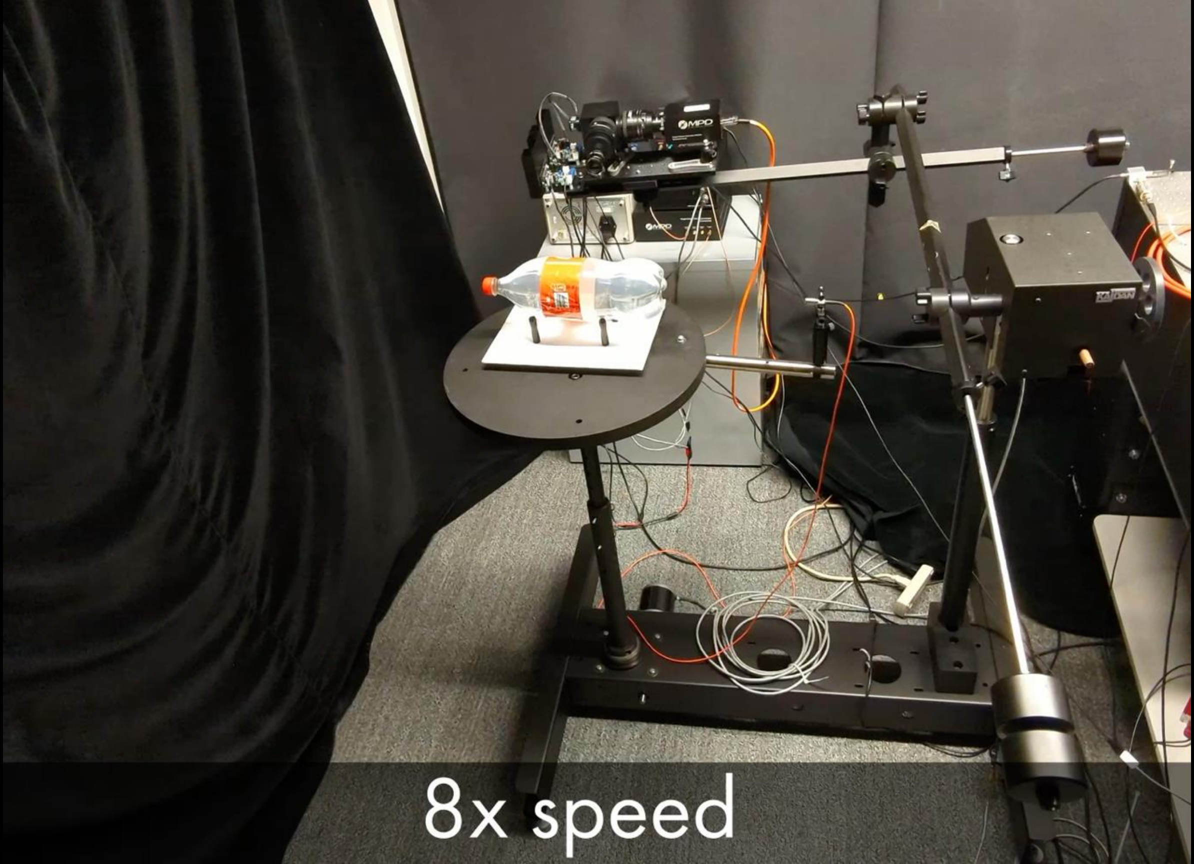
MPD
8x speed

results

simulated scene (ground truth)

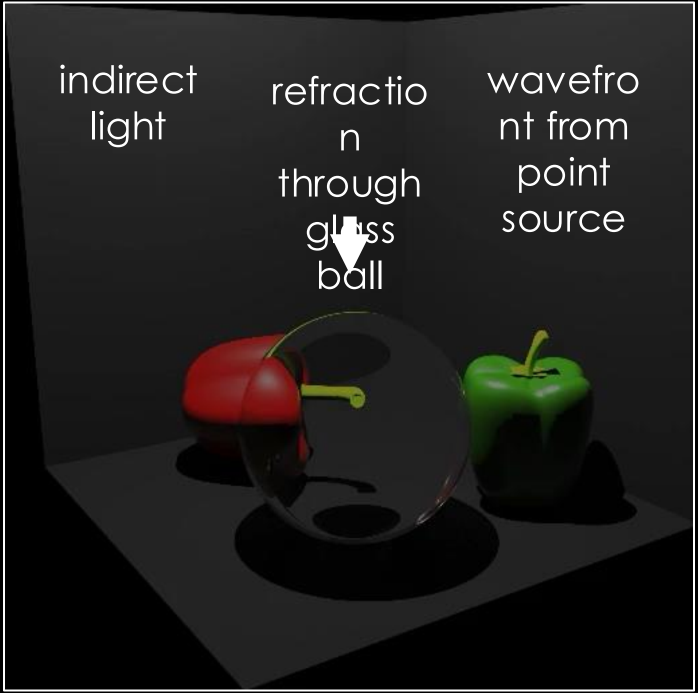

simulated scene (ground truth)

integrated image

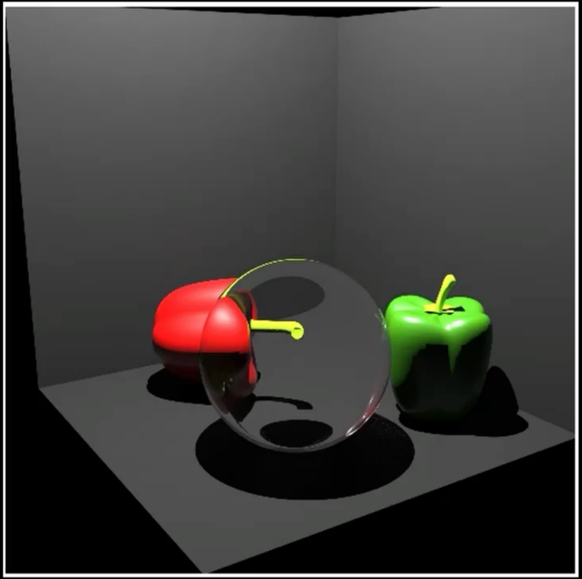

transient

moving camera (novel viewpoints)

Transient NeRF

K-Planes

proposed

moving camera (novel viewpoints —captured data)

color image

transient video (proposed)

[Malik et al. '2

moving camera (novel viewpoints —captured data)

color image

transient video (proposed)

[Malik et al. '2

moving camera (novel viewpoints —captured data)

color image

transient video (proposed)

[Malik et al. '2

moving camera (novel viewpoints—captured data)

color image

transient video (proposed)

[Malik et al. '2

moving camera (novel viewpoints—captured data)

color image

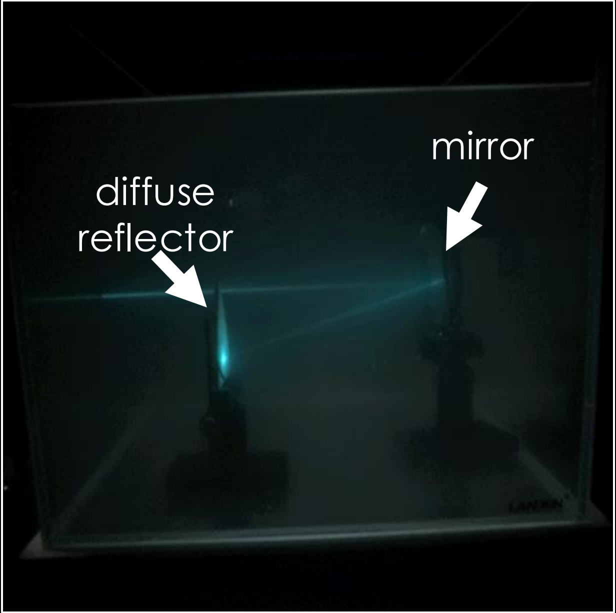

transient video (proposed)

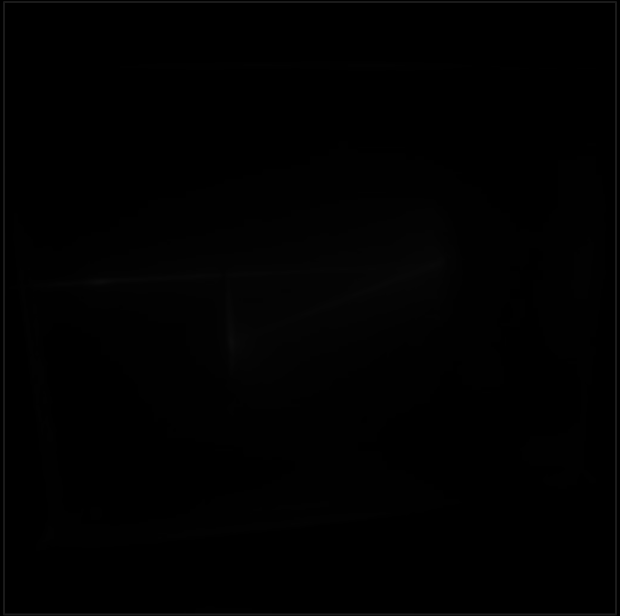

[Malik et al. '2

moving camera (novel viewpoints—captured data)

color image

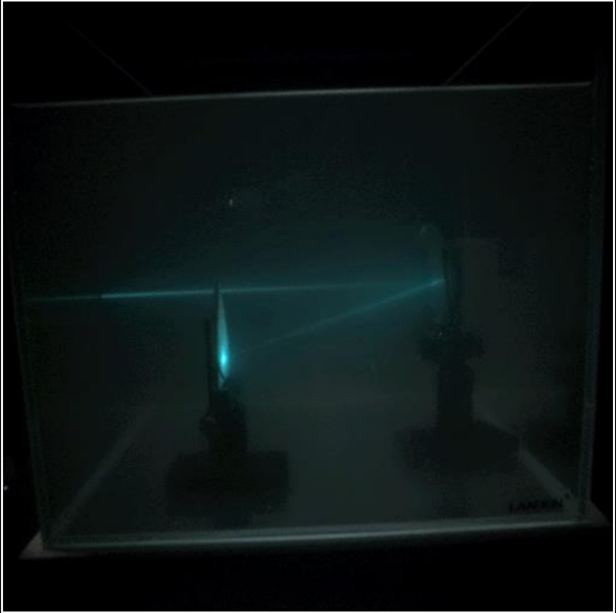

transient video (proposed)

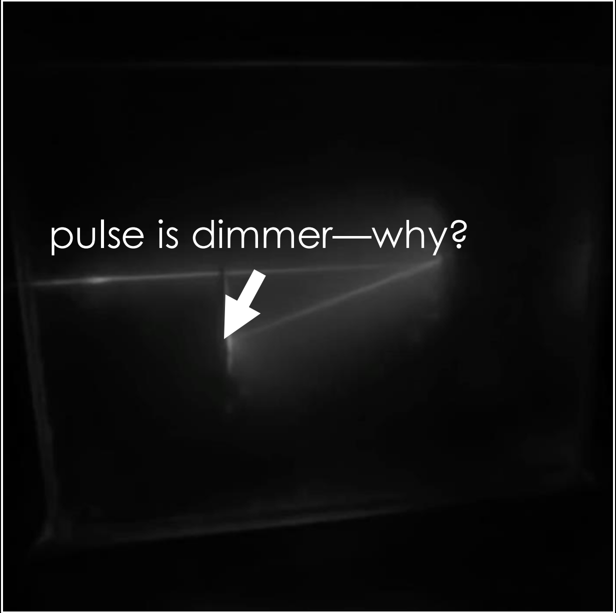

[Malik et al. '2

moving camera (novel viewpoints—captured data)

color image

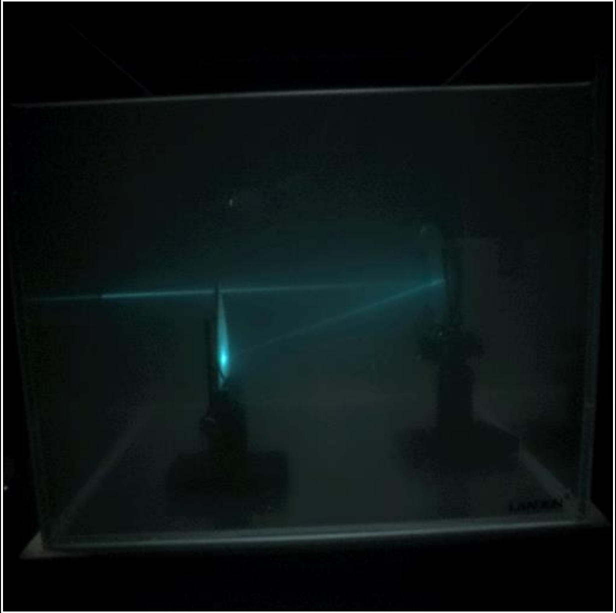

transient video (proposed)

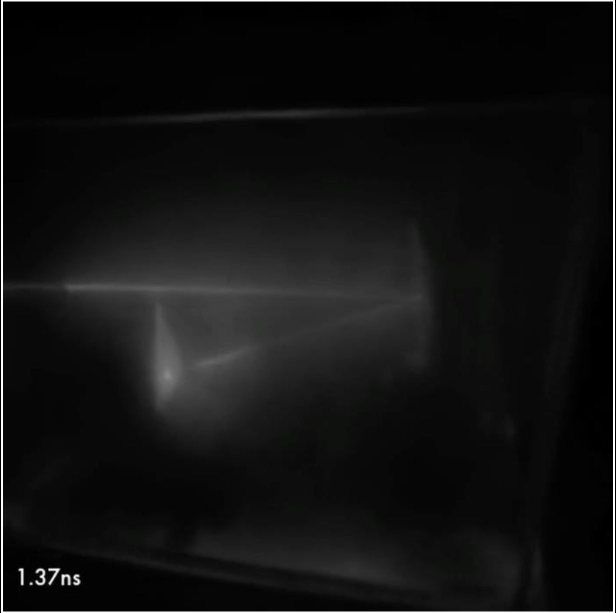

[Malik et al. '2

moving camera (novel viewpoints—captured data)

color image

transient video (proposed)

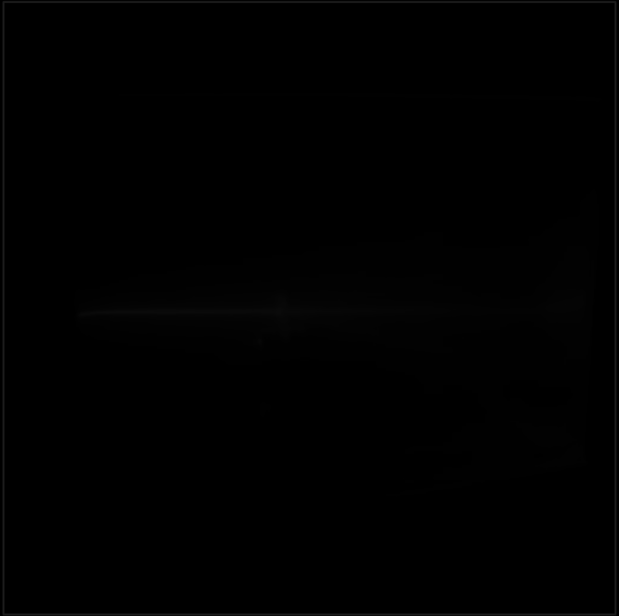

[Malik et al. '2

moving camera (novel viewpoints—captured scene)

color image

transient video (proposed)

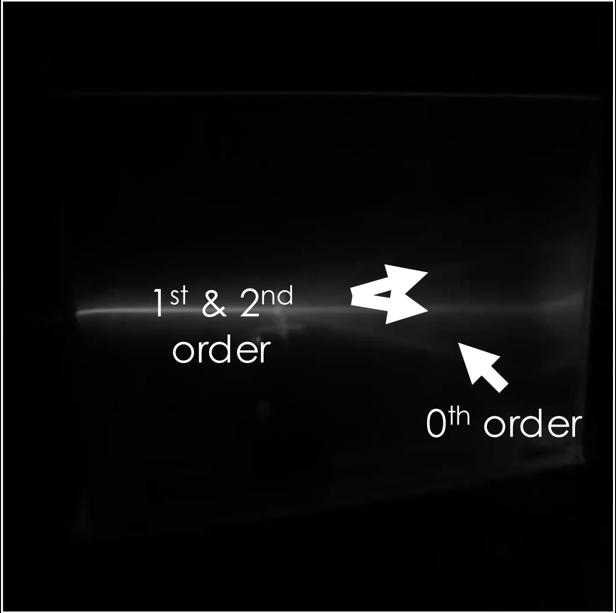

[Malik et al. '2

moving camera (novel viewpoints—captured scene)

color image

transient video (proposed)

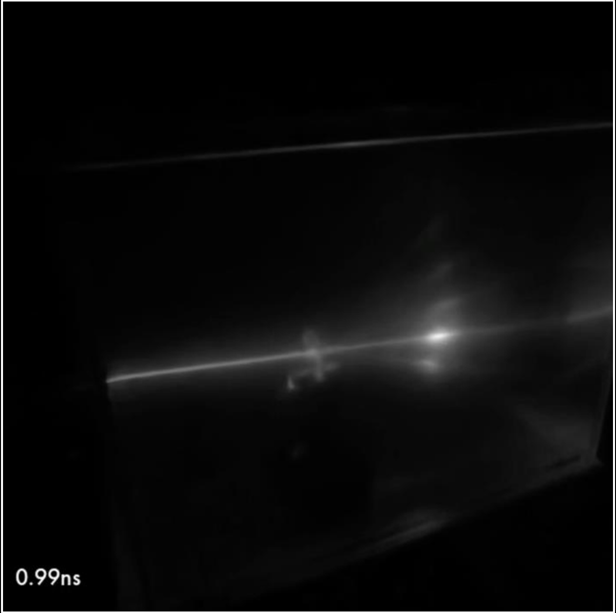

[Malik et al. '2

captured scene (novel viewpoints)

color image

transient video

[Malik et al. ‘2

captured scene (novel viewpoints)

color image

transient video

[Malik et al. '2

captured scene (novel viewpoints)

color image

transient video

[Malik et al. '2

time-warping (Velten et al. 2012)

no warping

depth-based warping

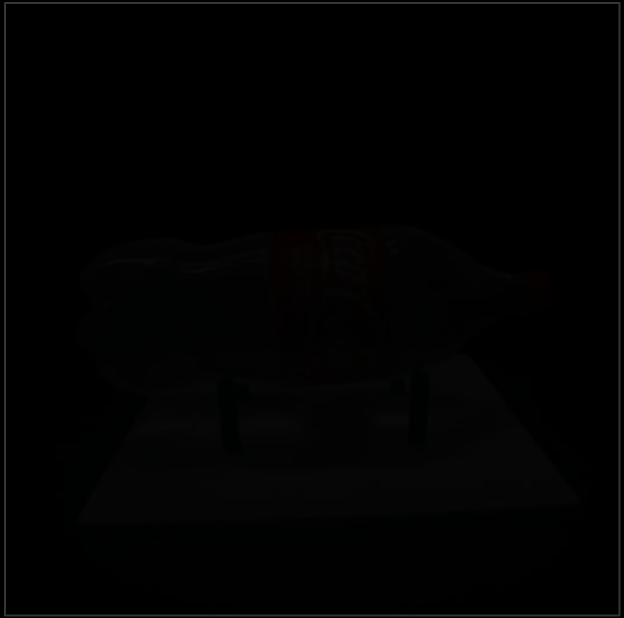

time-warping (Velten et al. 2012)

no warping

depth-based warping

discussion

future work

- Can we render & incorporate relativistic effects?

future work

- Can we render & incorporate relativistic effects?
- this work focuses primarily on visualization, but…
  - can we use multiply-scattered light to recover 3D geometry?
  - what about material or reflectance properties?
  - imaging through scattering media?

future work

- 3D visualization of light transport
  - artistic, educational, and scientific visualizations
  - visualize non-linear or quantum optical effects?
  - imaging fast & slow


Frame rate: 10.0Hz
Intensity tone mapped
Elapsed time: 24s + 500ms

[Wei et al. '23]

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!

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).
- Wei, Mian, et al. "Passive Ultra-Wideband Single-Photon Imaging." Proceedings of the IEEE/CVF International Conference on Computer Vision. 2023.
- Malik, Anagh, et al. "Flying With Photons: Rendering Novel Views of Propagating Light." arXiv preprint arXiv:2404.06493 (2024).