Course Schedule

http://imagingav.media.mit.edu

13:30 – 13:50	Course Introduction *Ramesh Raskar (MIT)*
13:50 – 15:00	Existing Sensors and Their Limits ***Guy Satat (MIT)***, *Achuta Kadambi (UCLA)*
15:00 – 15:10	Break
15:10 – 15:50	Emerging 3D Sensors *Achuta Kadambi (UCLA)*
15:50 – 16:30	Imaging in Bad Weather *Guy Satat (MIT)*
16:30 – 16:40	Break
16:40 – 17:20	Deep Learning-based Computational Imaging *Jan Kautz (NVIDIA)*
17:20 – 17:30	Conclusion and Open Problems

Course Schedule

http://imagingav.media.mit.edu

13:30 – 13:50	Course Introduction *Ramesh Raskar (MIT)*
13:50 – 15:00	Existing Sensors and Their Limits *Guy Satat (MIT),* ***Achuta Kadambi (UCLA)***
15:00 – 15:10	Break
15:10 – 15:50	Emerging 3D Sensors *Achuta Kadambi (UCLA)*
15:50 – 16:30	Imaging in Bad Weather *Guy Satat (MIT)*
16:30 – 16:40	Break
16:40 – 17:20	Deep Learning-based Computational Imaging *Jan Kautz (NVIDIA)*
17:20 – 17:30	Conclusion and Open Problems

A Dive into LIDAR

Achuta Kadambi *(MIT/UCLA)*


3 Colors


Photo credit: Rabbi Yair Hoffman

Photo credit: Conde Nast

Photo credit: David Matthews Jr

Teaching Machines to See the World

Teaching Machines to See is Hard

PC: MIT Tech Review


Spinning
Imager
Photo credit: Xavier Harding

Time of Flight

Time it takes for an object, particle, or wave to travel a distance through a medium.

LIDAR

Police Speed Gun

Microsoft Kinect

1638: First Time of Flight Camera?

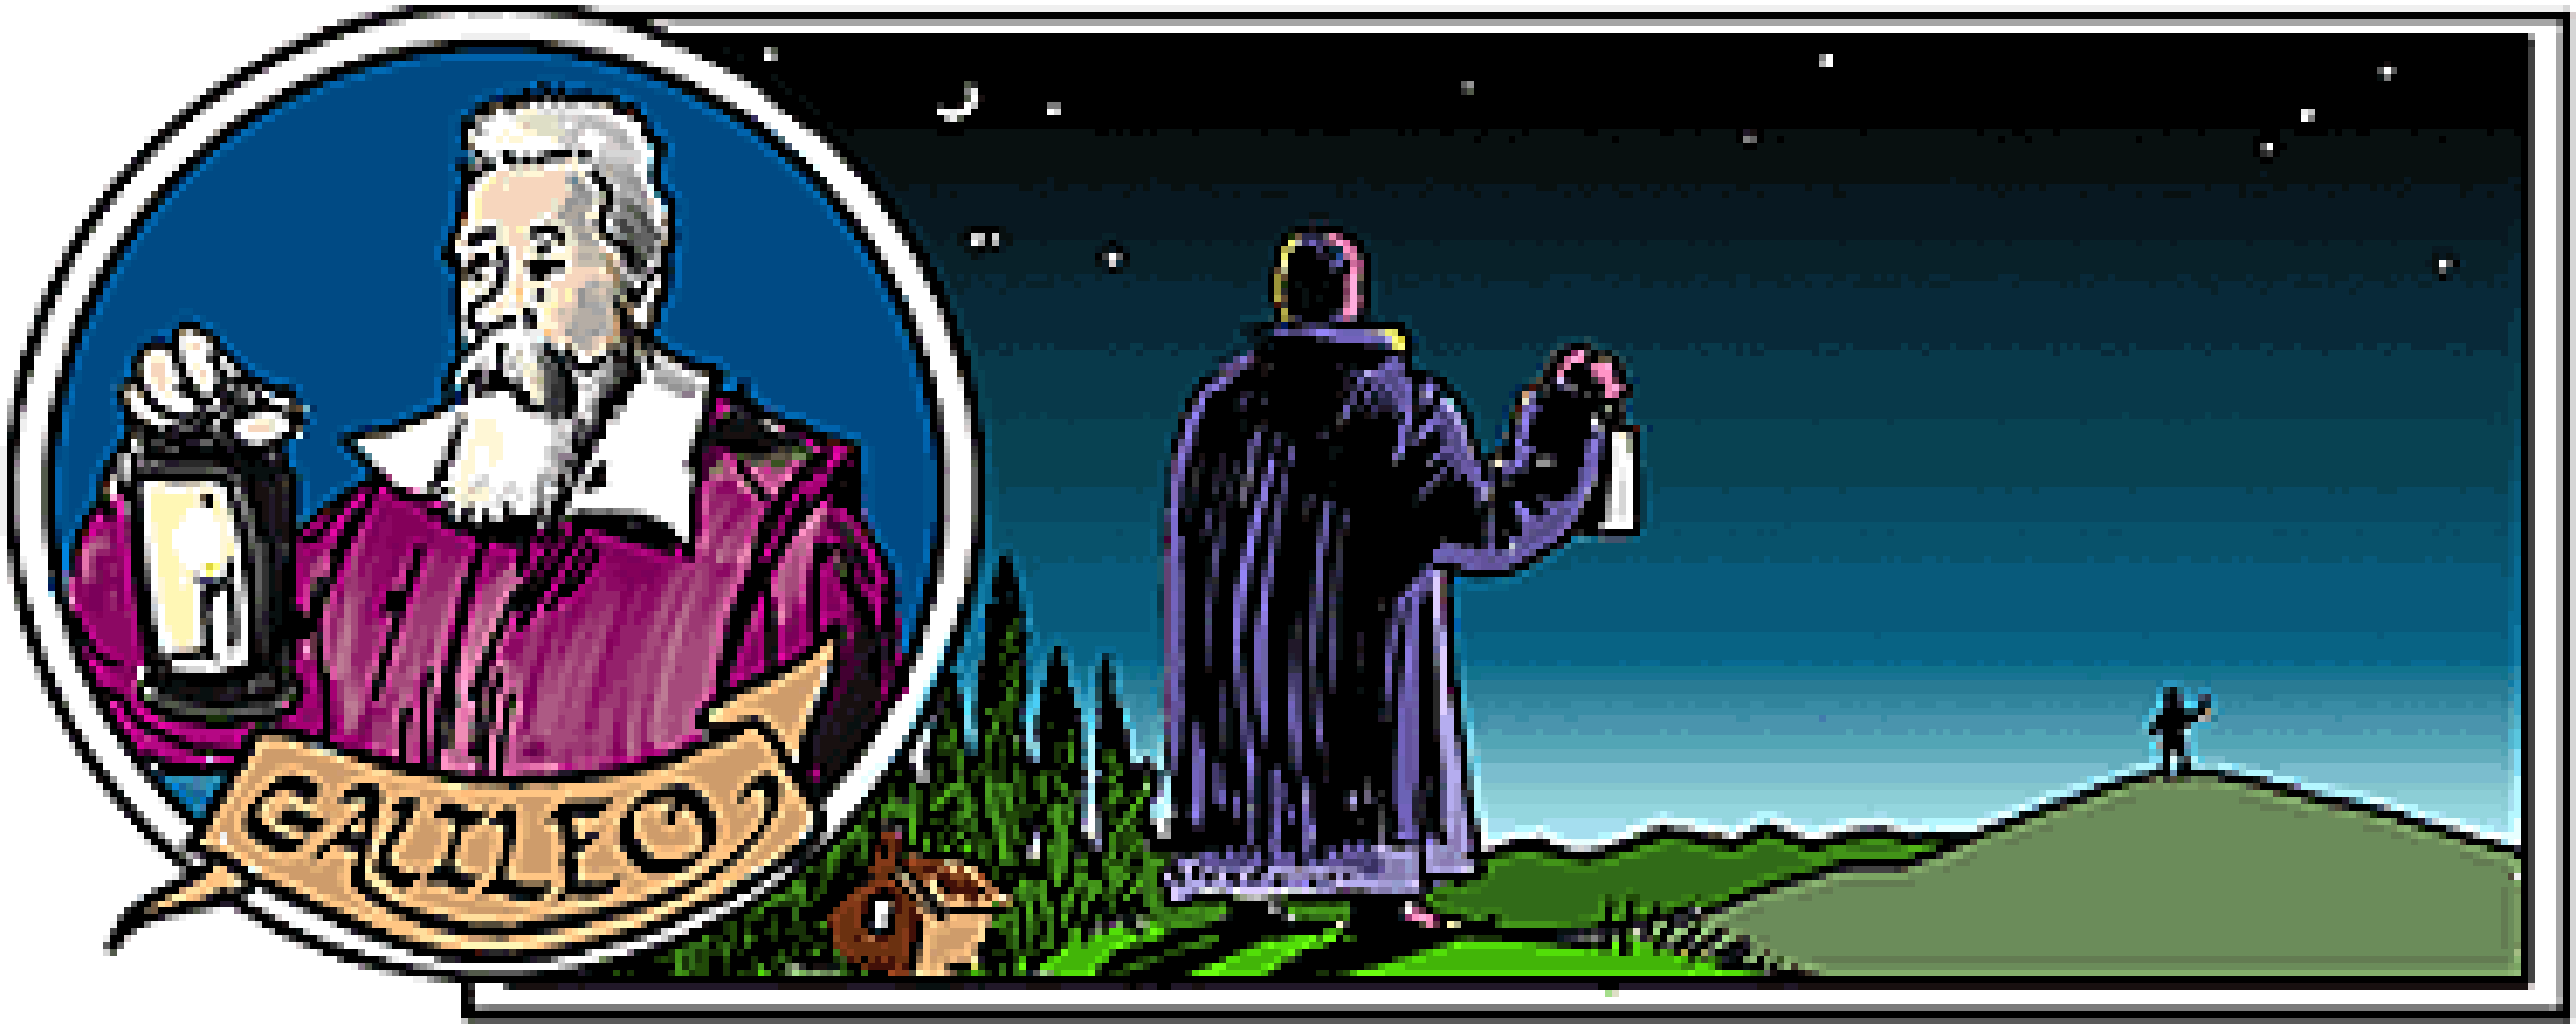

Graphic courtesy: phys.colorado.edu

Time of Flight in your Living Room (MS Kinect)

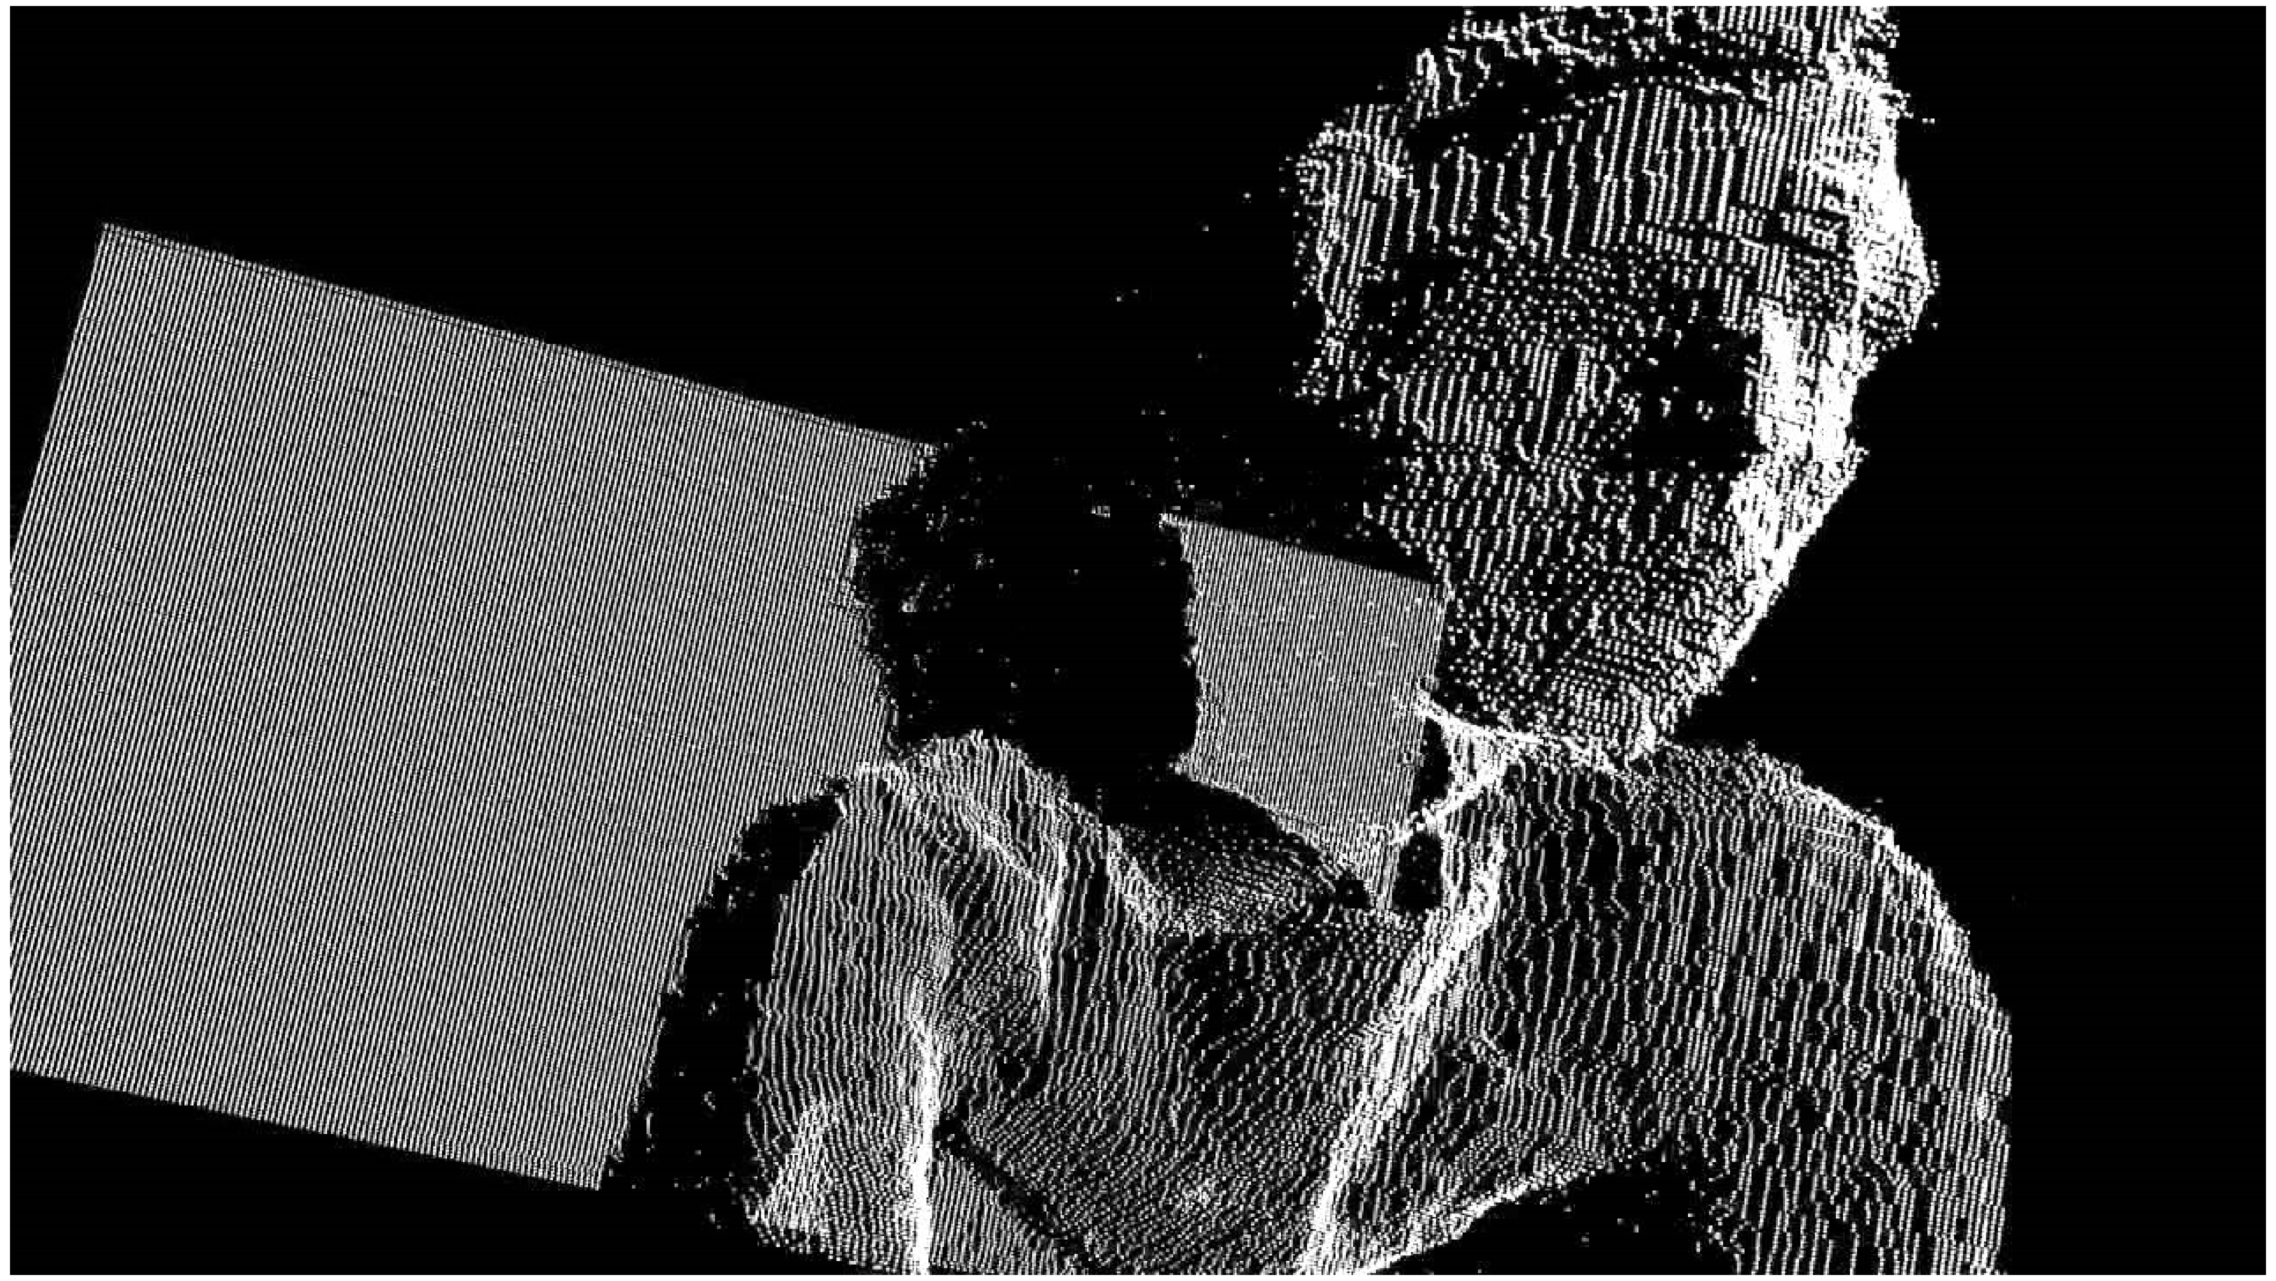

Time of Flight on your Car (Velodyne LIDAR)

Silicon Valley spinoff from Velodyne Acoustics

Breakthrough ~2005: Unlike other LIDAR
it rotates and obtains a wide-FOV

Spins ~15 Hz (e.g. ~500 rpm)

Exemplary specs:

32 photon detectors
16 lasers

How does LIDAR work today?

Sending photons → target

$$d = vt$$

Delving Deeper into the LIDAR Taxonomy

Impulse

Continuous-Wave

Impulse 3D Imaging

LIDAR Imaging System

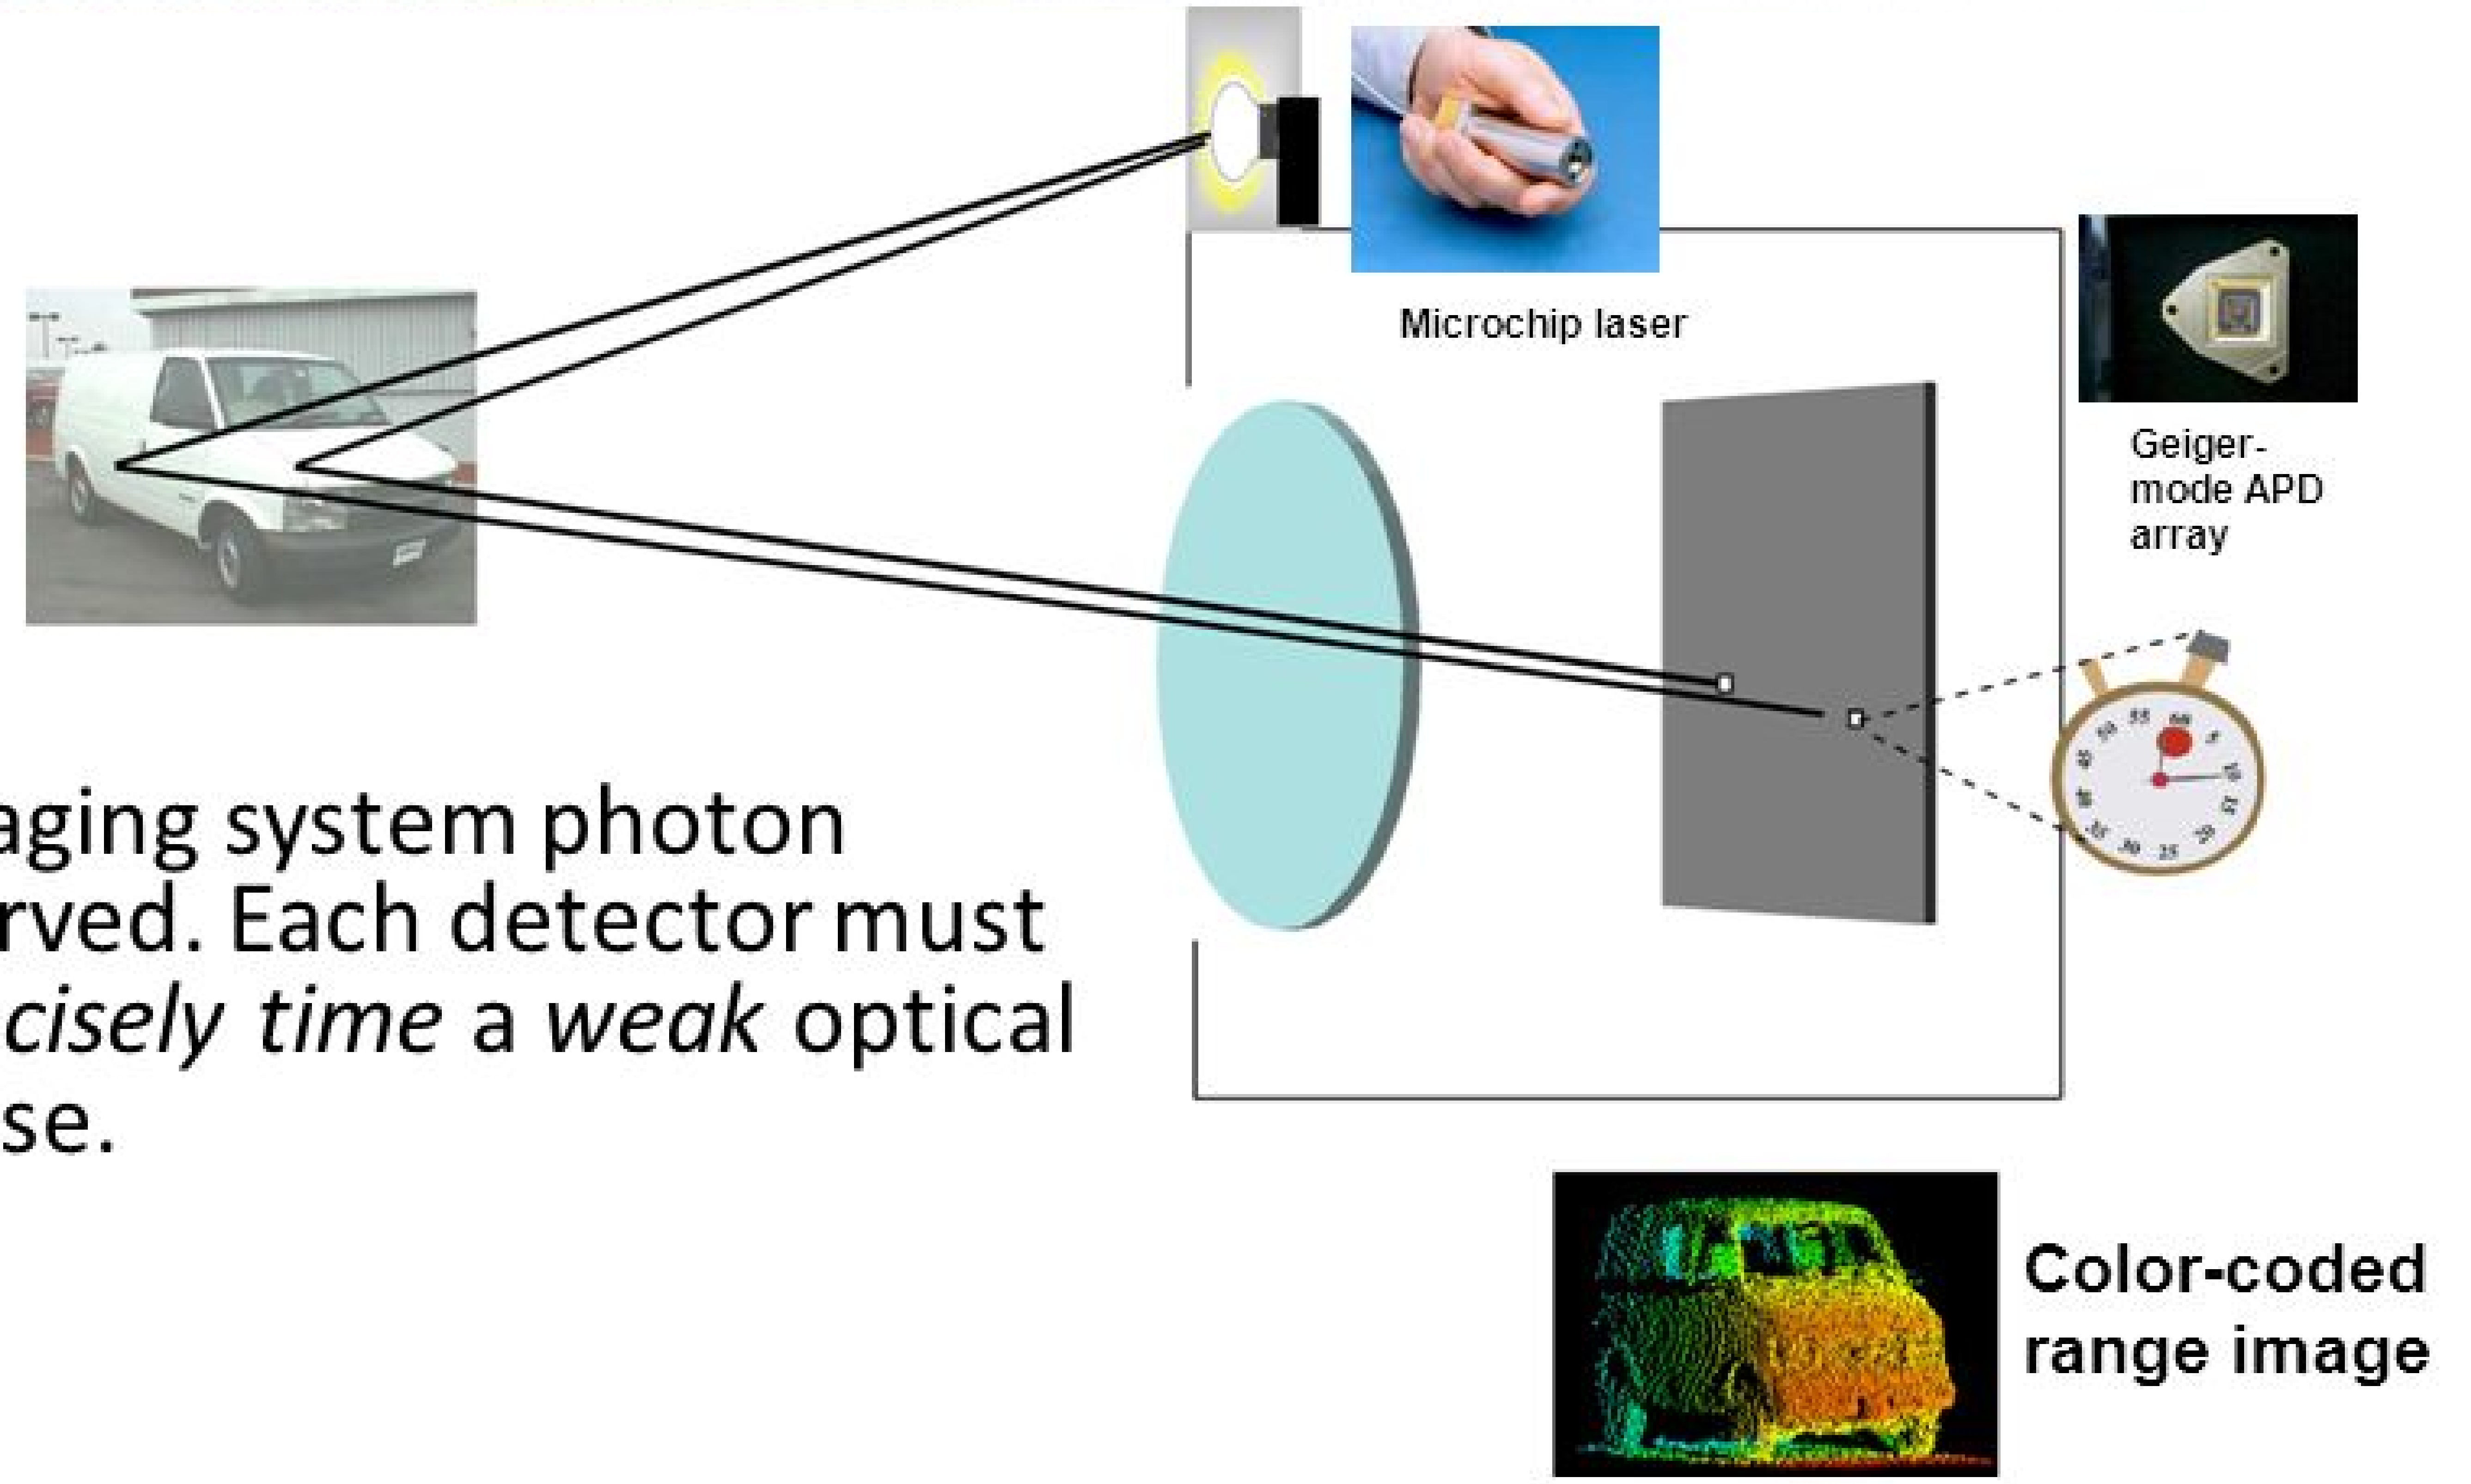

- Imaging system photon starved. Each detector must *precisely time* a *weak* optical pulse.

Achuta Kadambi
MIT/UCLA

The Avalanche Photodiode

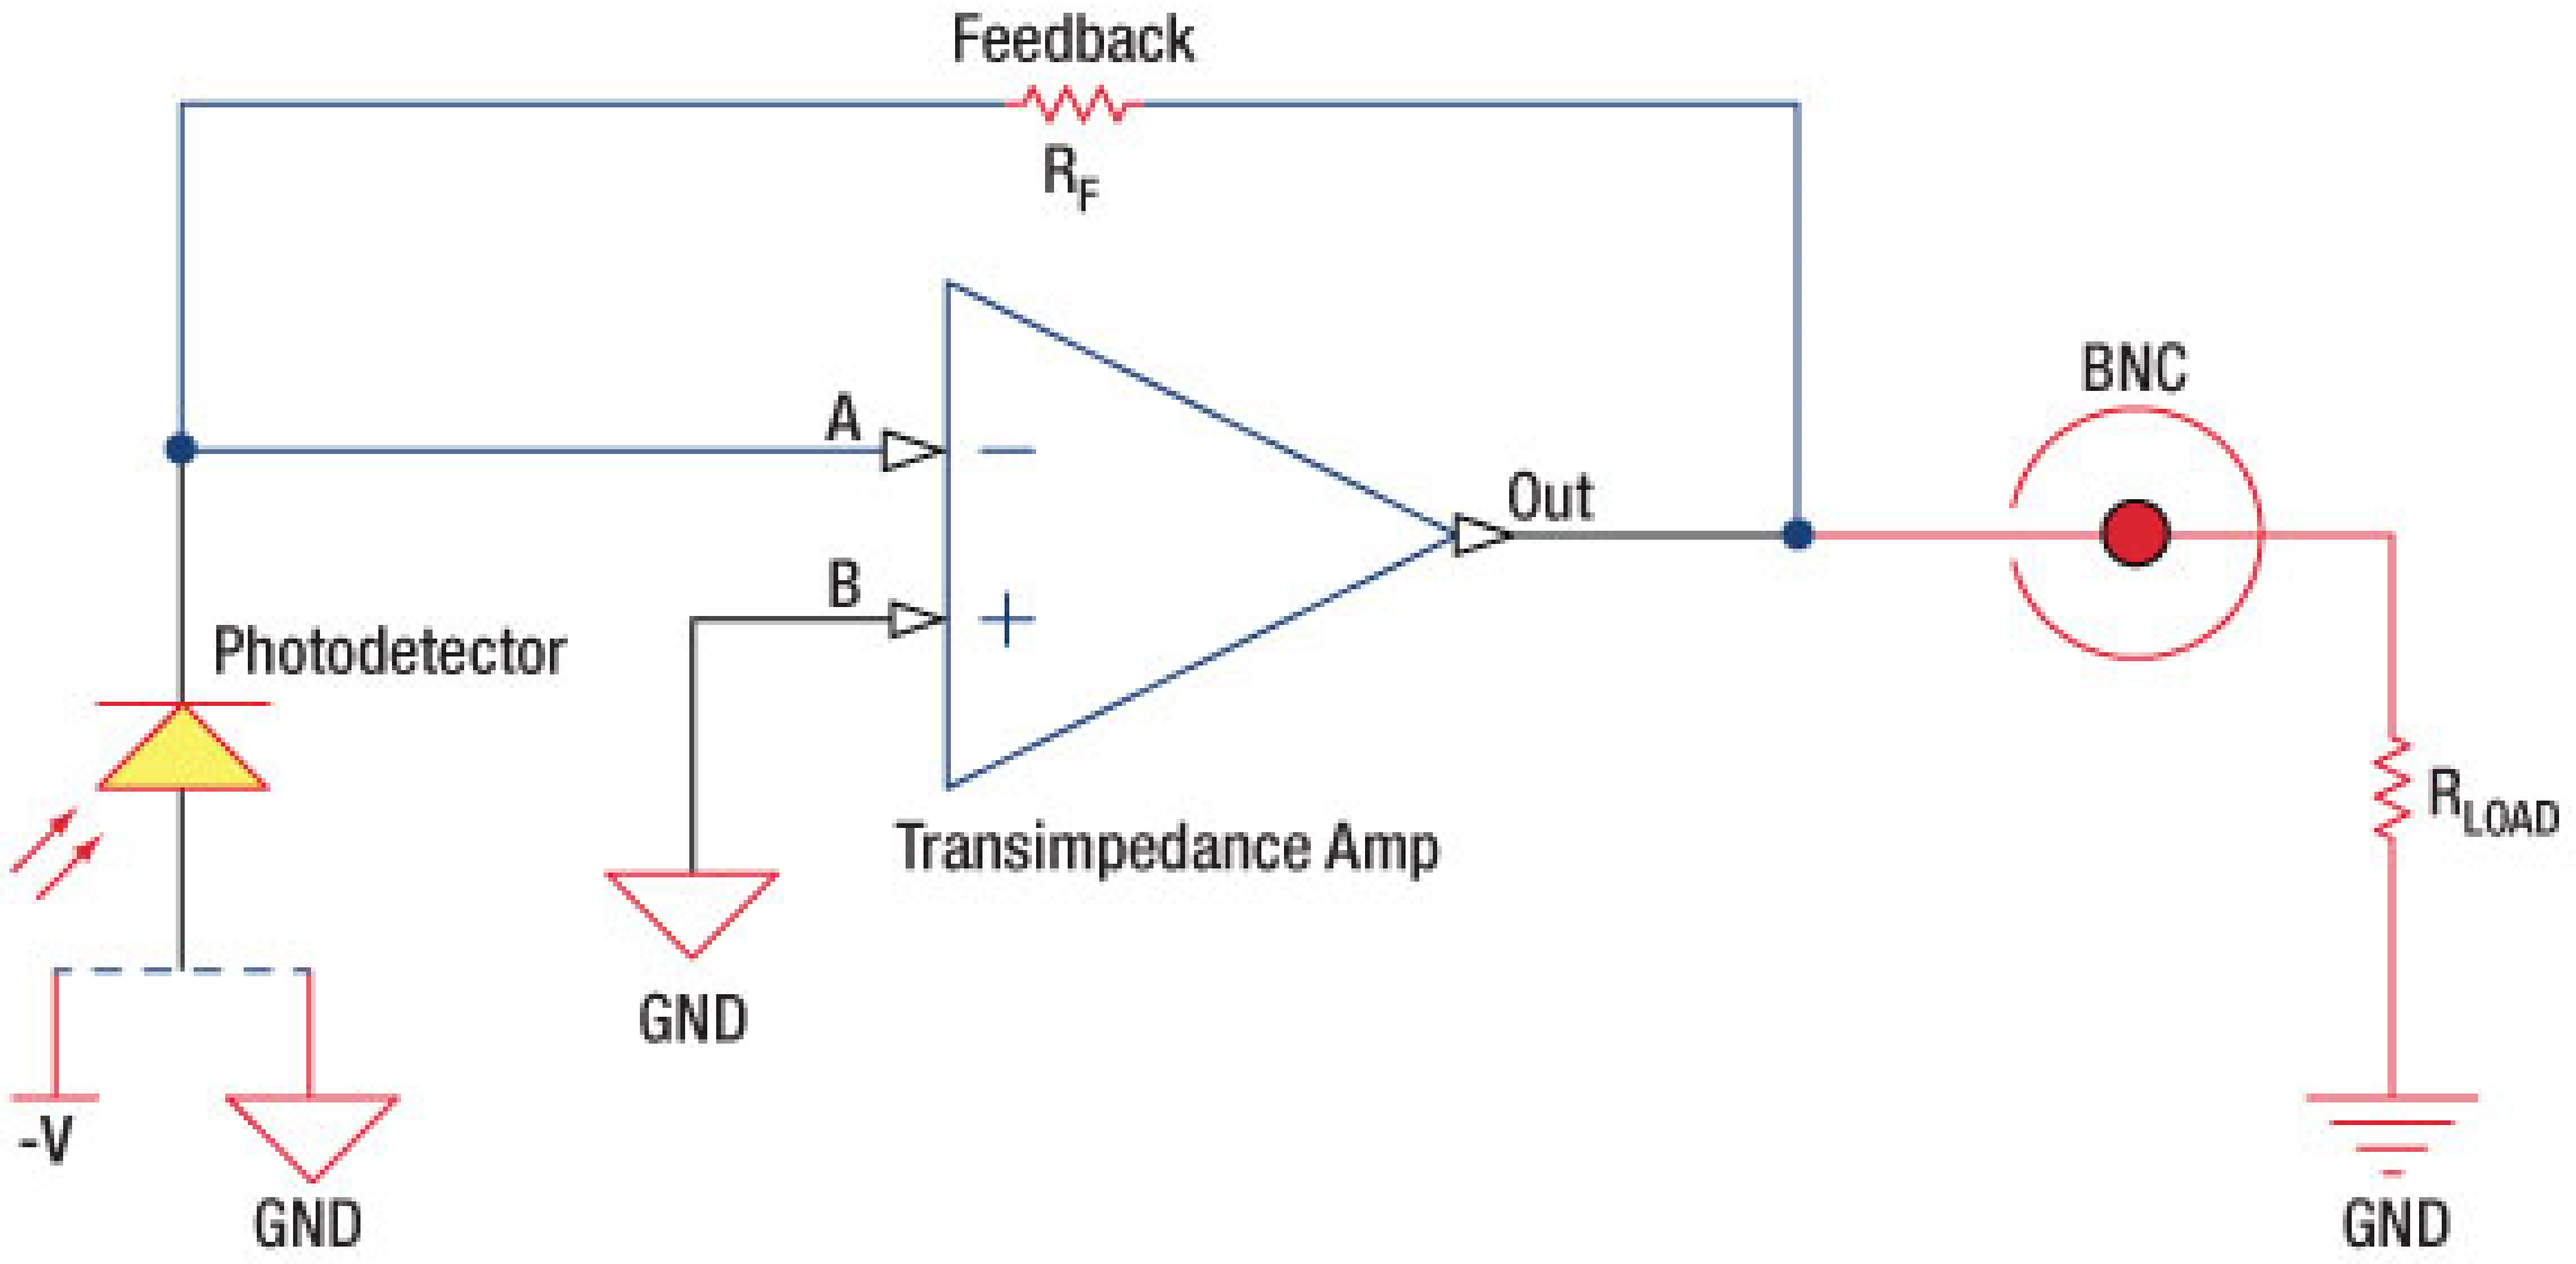

Thor Labs APD

Key Ingredient : Reverse Bias

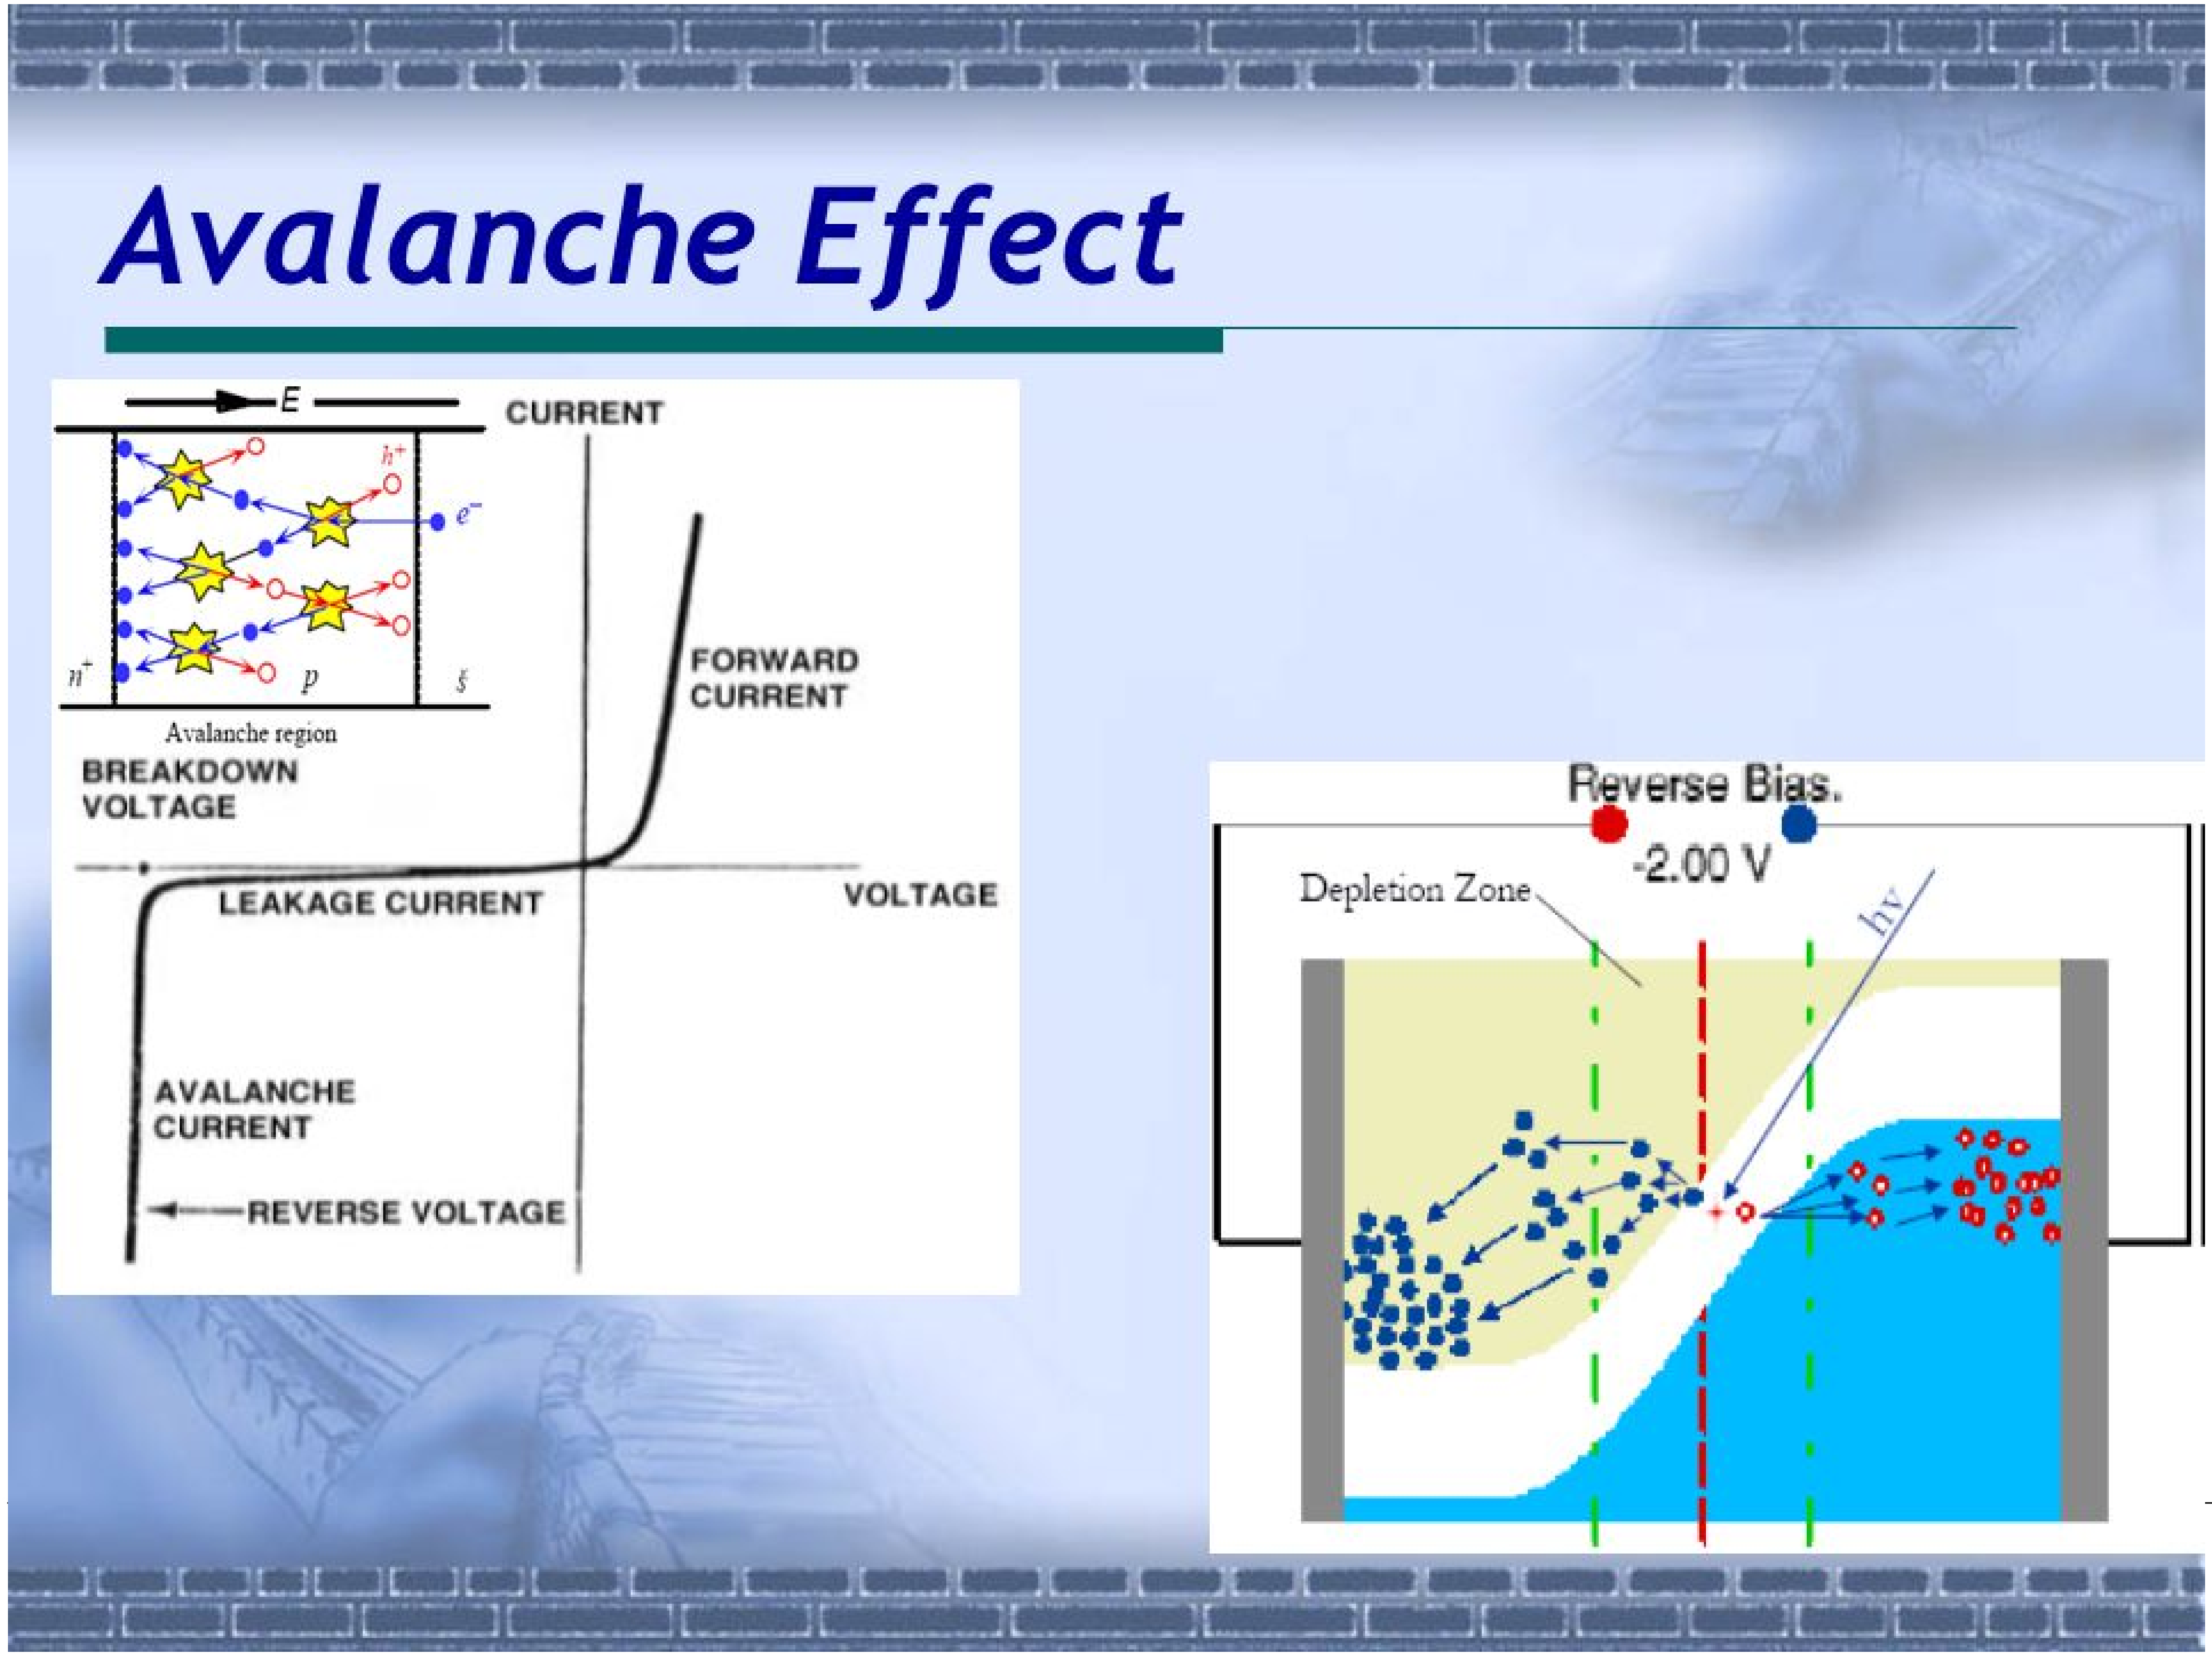

Slide from Jie Zhang

Remaining Complexity is in the Alignment and Scanning

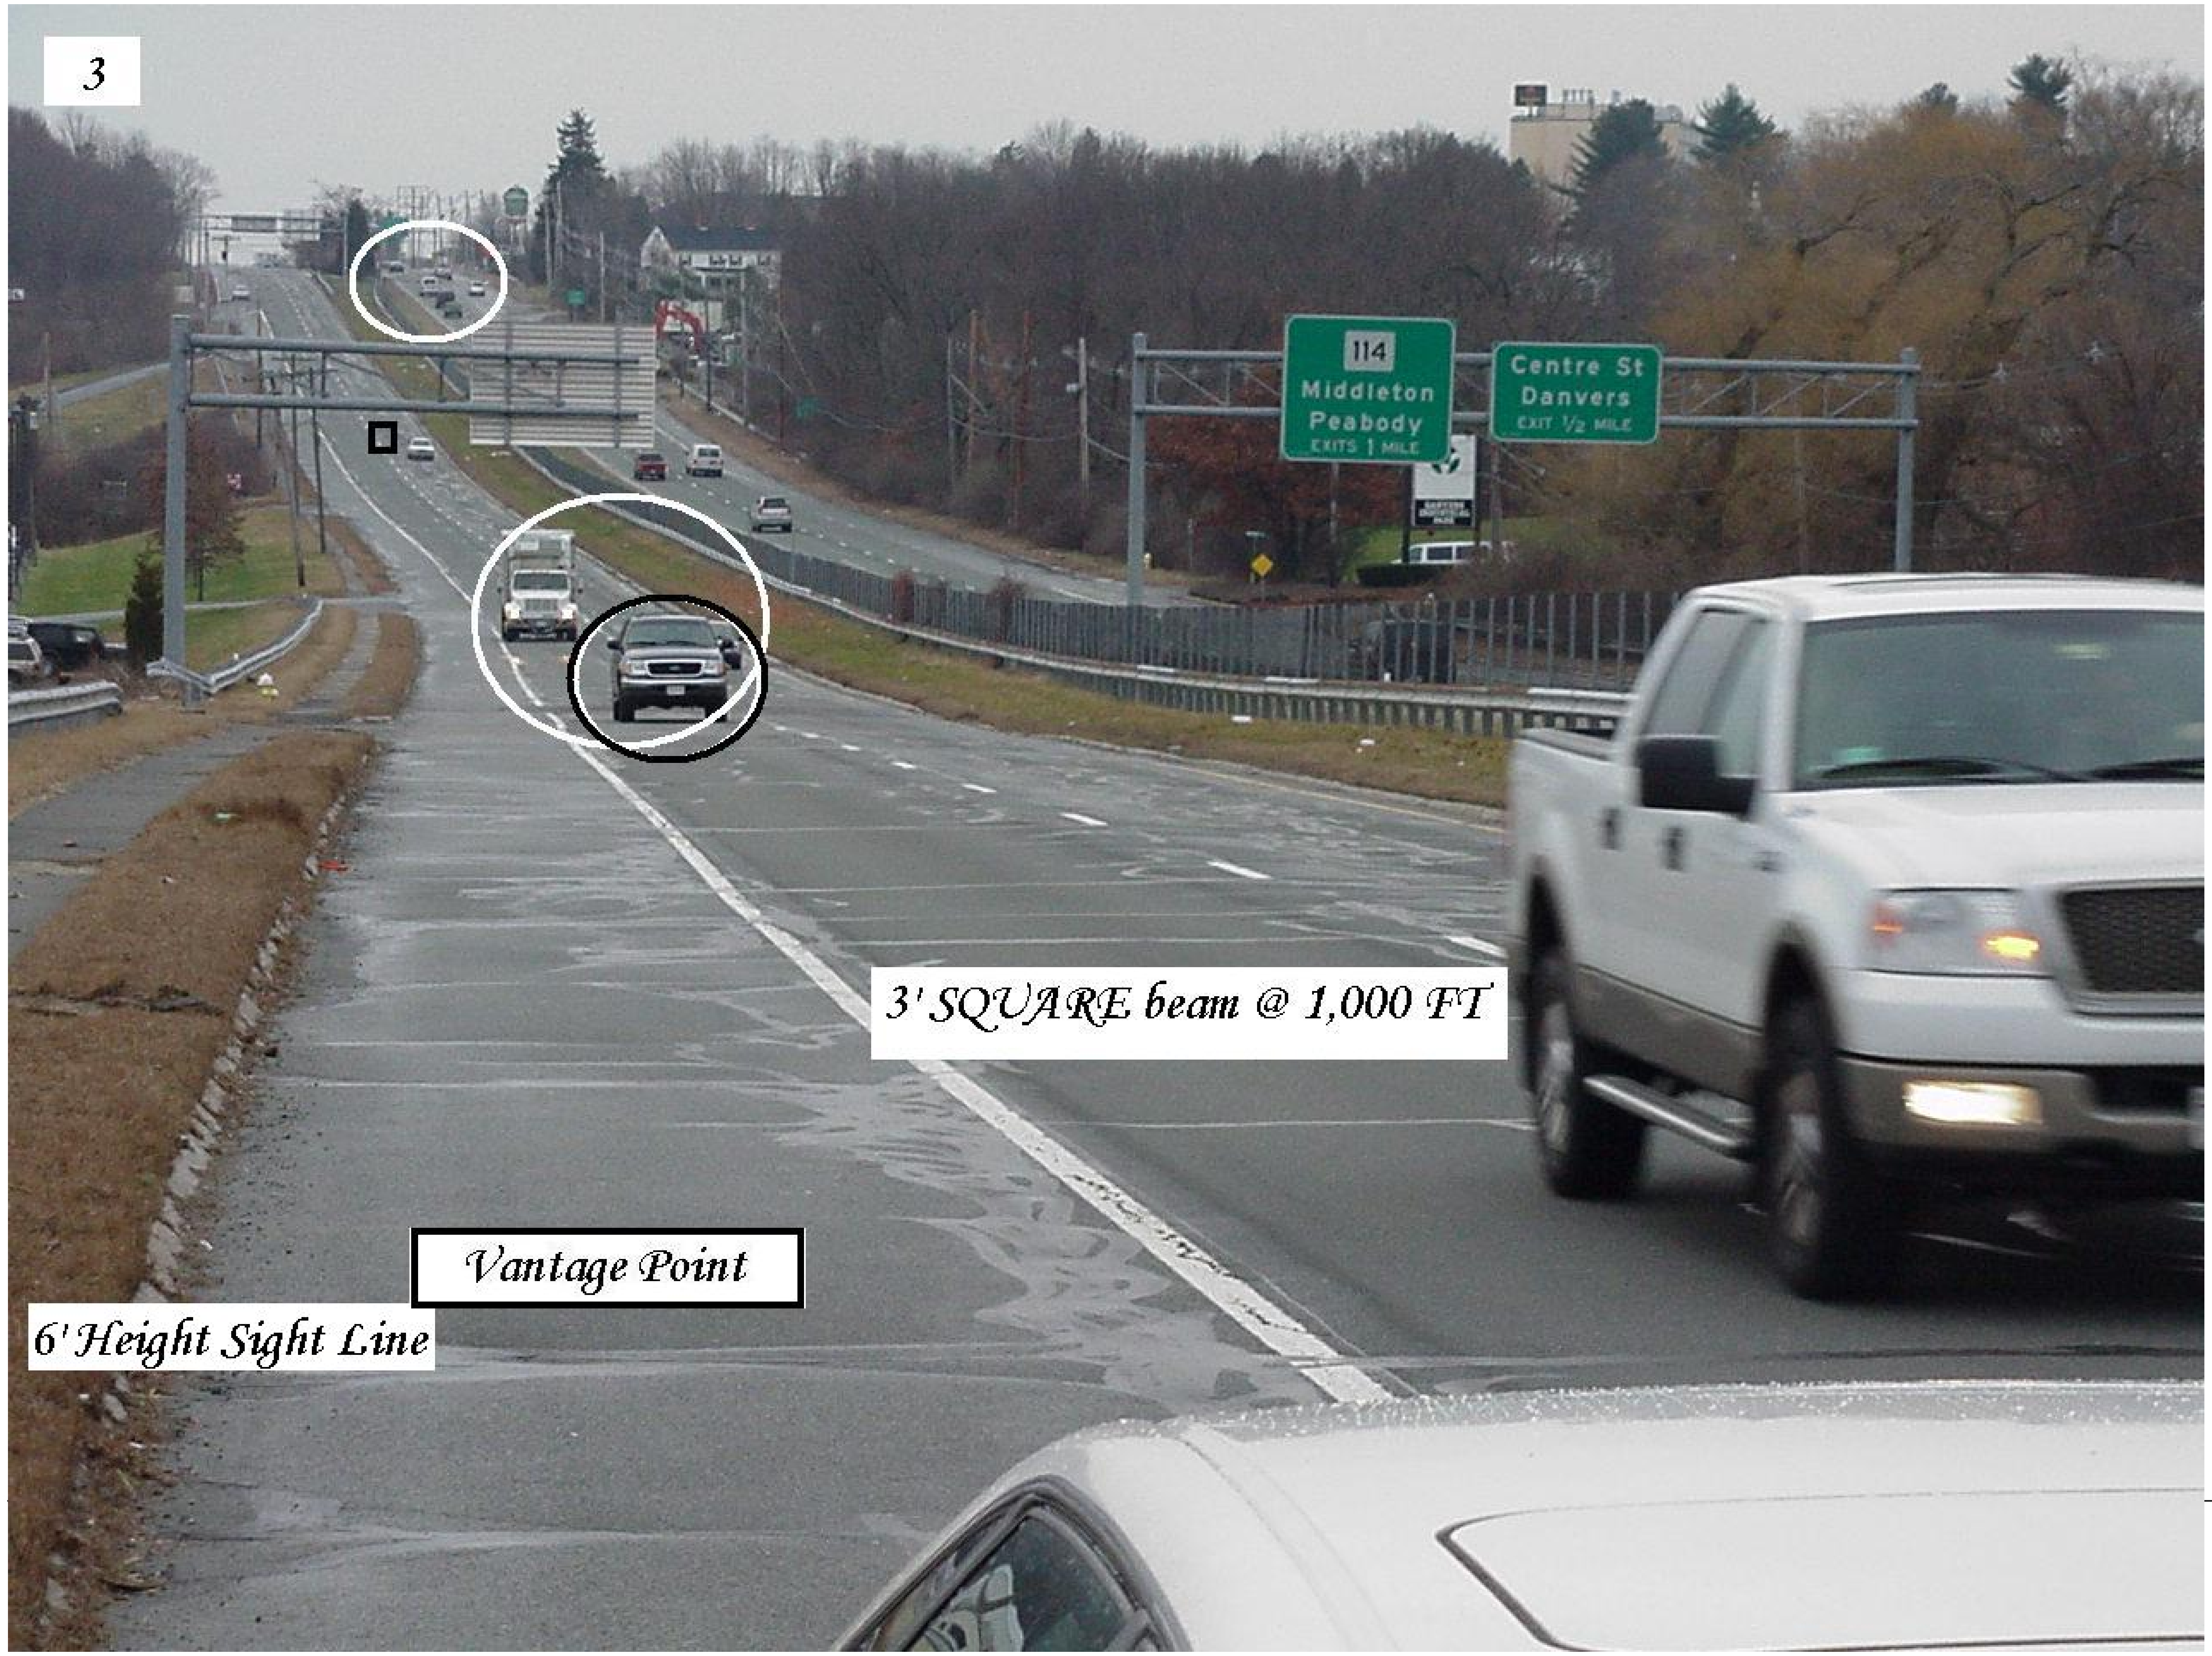

To go from a Police LIDAR to something more, we need to narrow the beam, rotate, align, etc.

This is Velodyne's contribution

Cost has come down drastically (VLP-16 is $4K)

Continuous-Wave 3D Imaging

$$d = vt \qquad v = 3 * 10^8 \, m/s$$

So, we need a camera to measure time delay.

Example: Microsoft Kinect (2013)

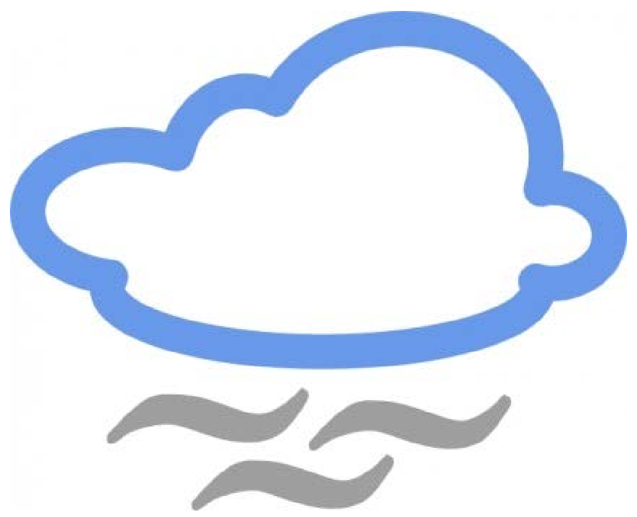

Rain

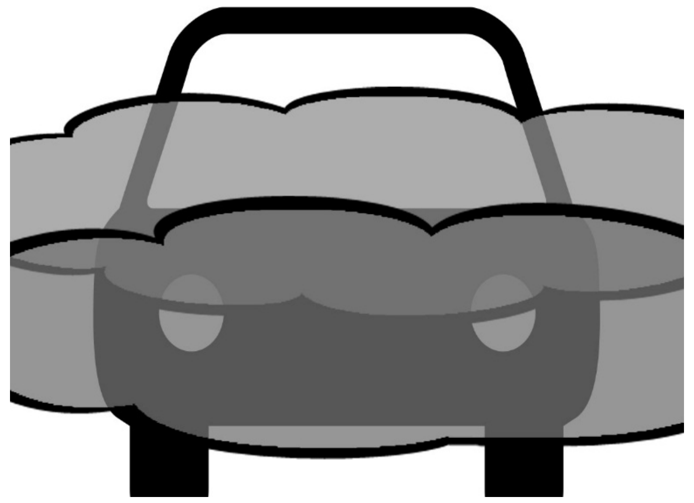

Fog

Unusual Objects

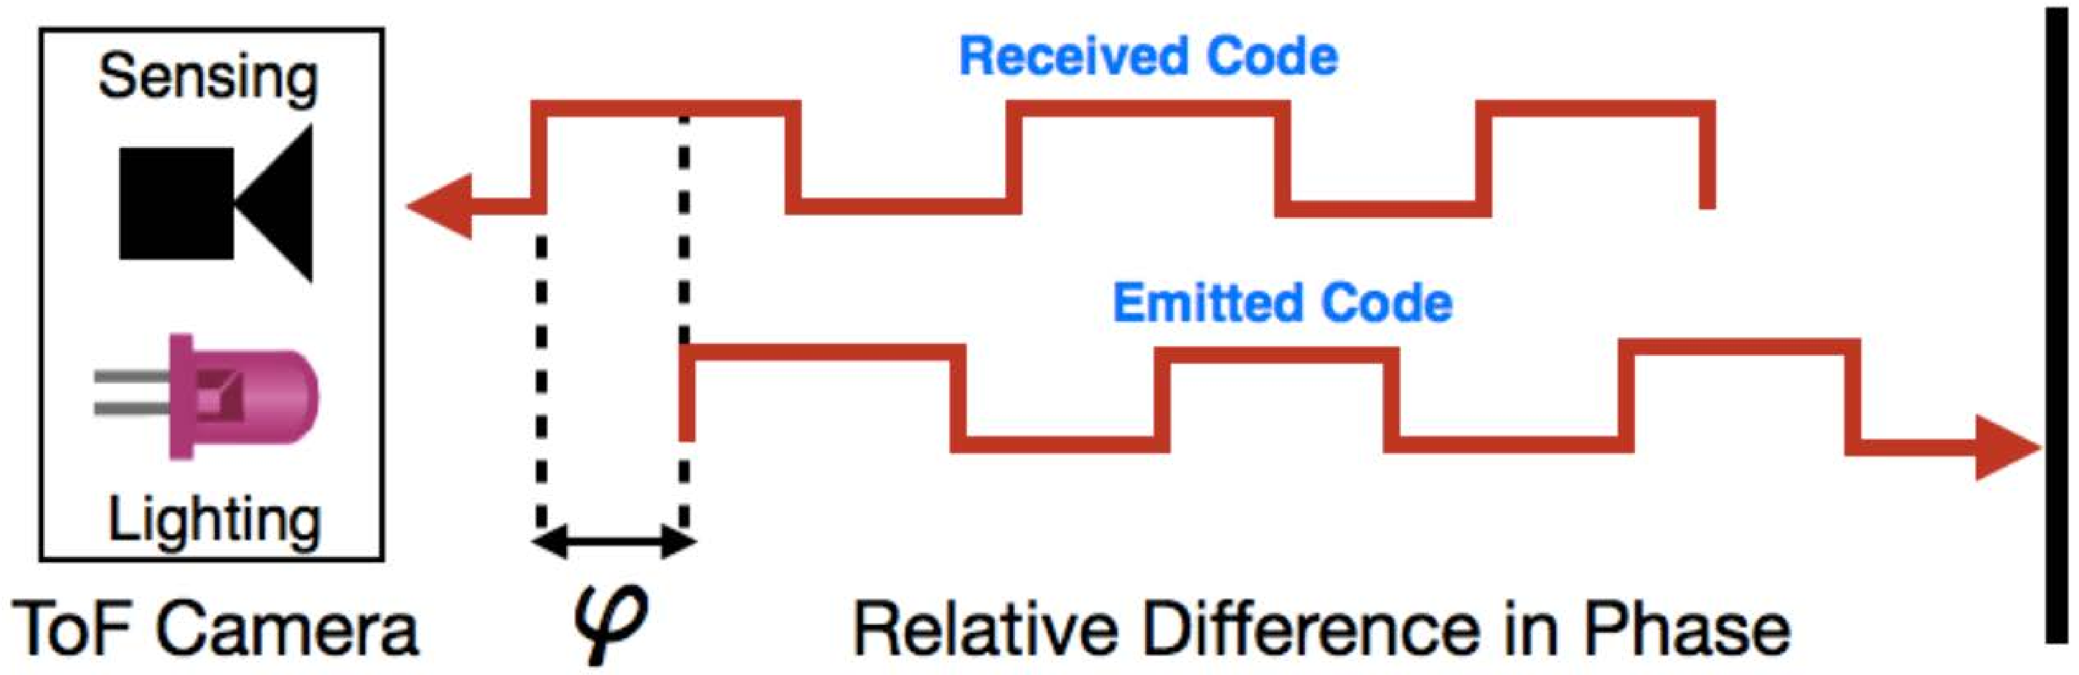

$$z = \frac{c\Phi}{2\pi f_\omega}, \quad c \approx 3\times10^8 \,\mathrm{m/s}$$

Two Dominant Types of CW Architectures

FMCW

AMCW

We continuously modulate a laser

FMCW Architecture

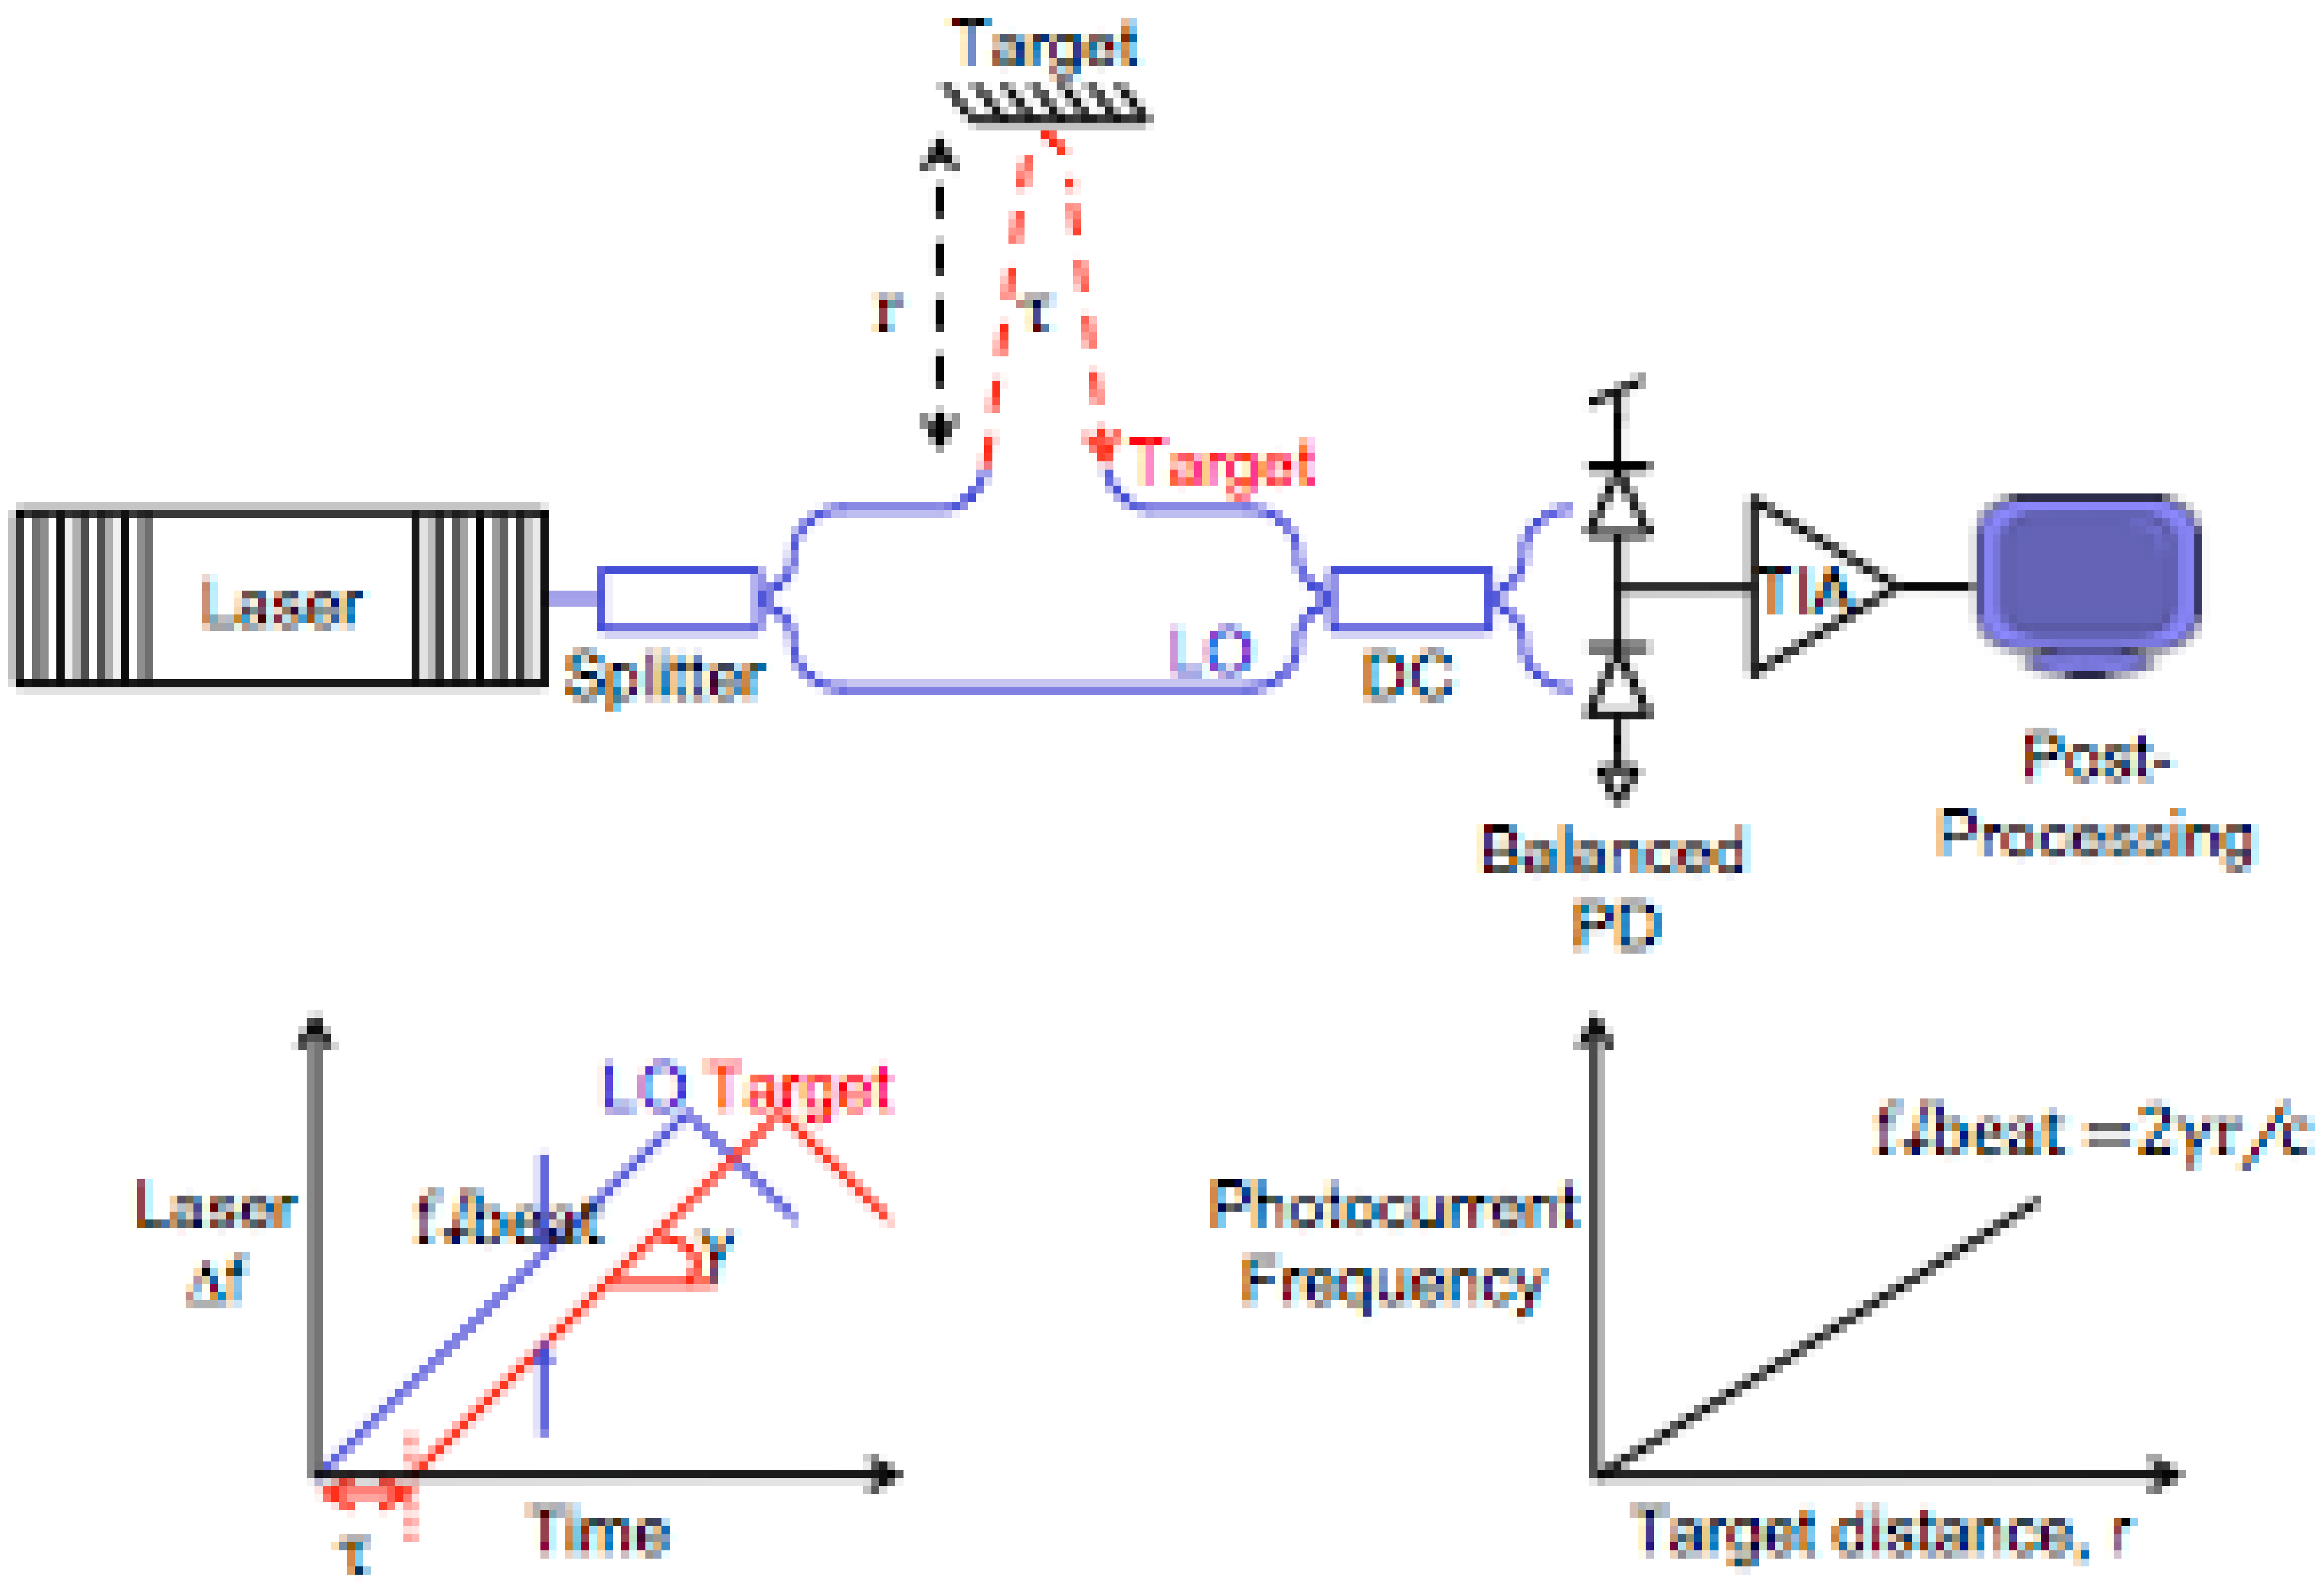

Figure from Ming Wu Lab, UC Berkeley

AMCW Architecture

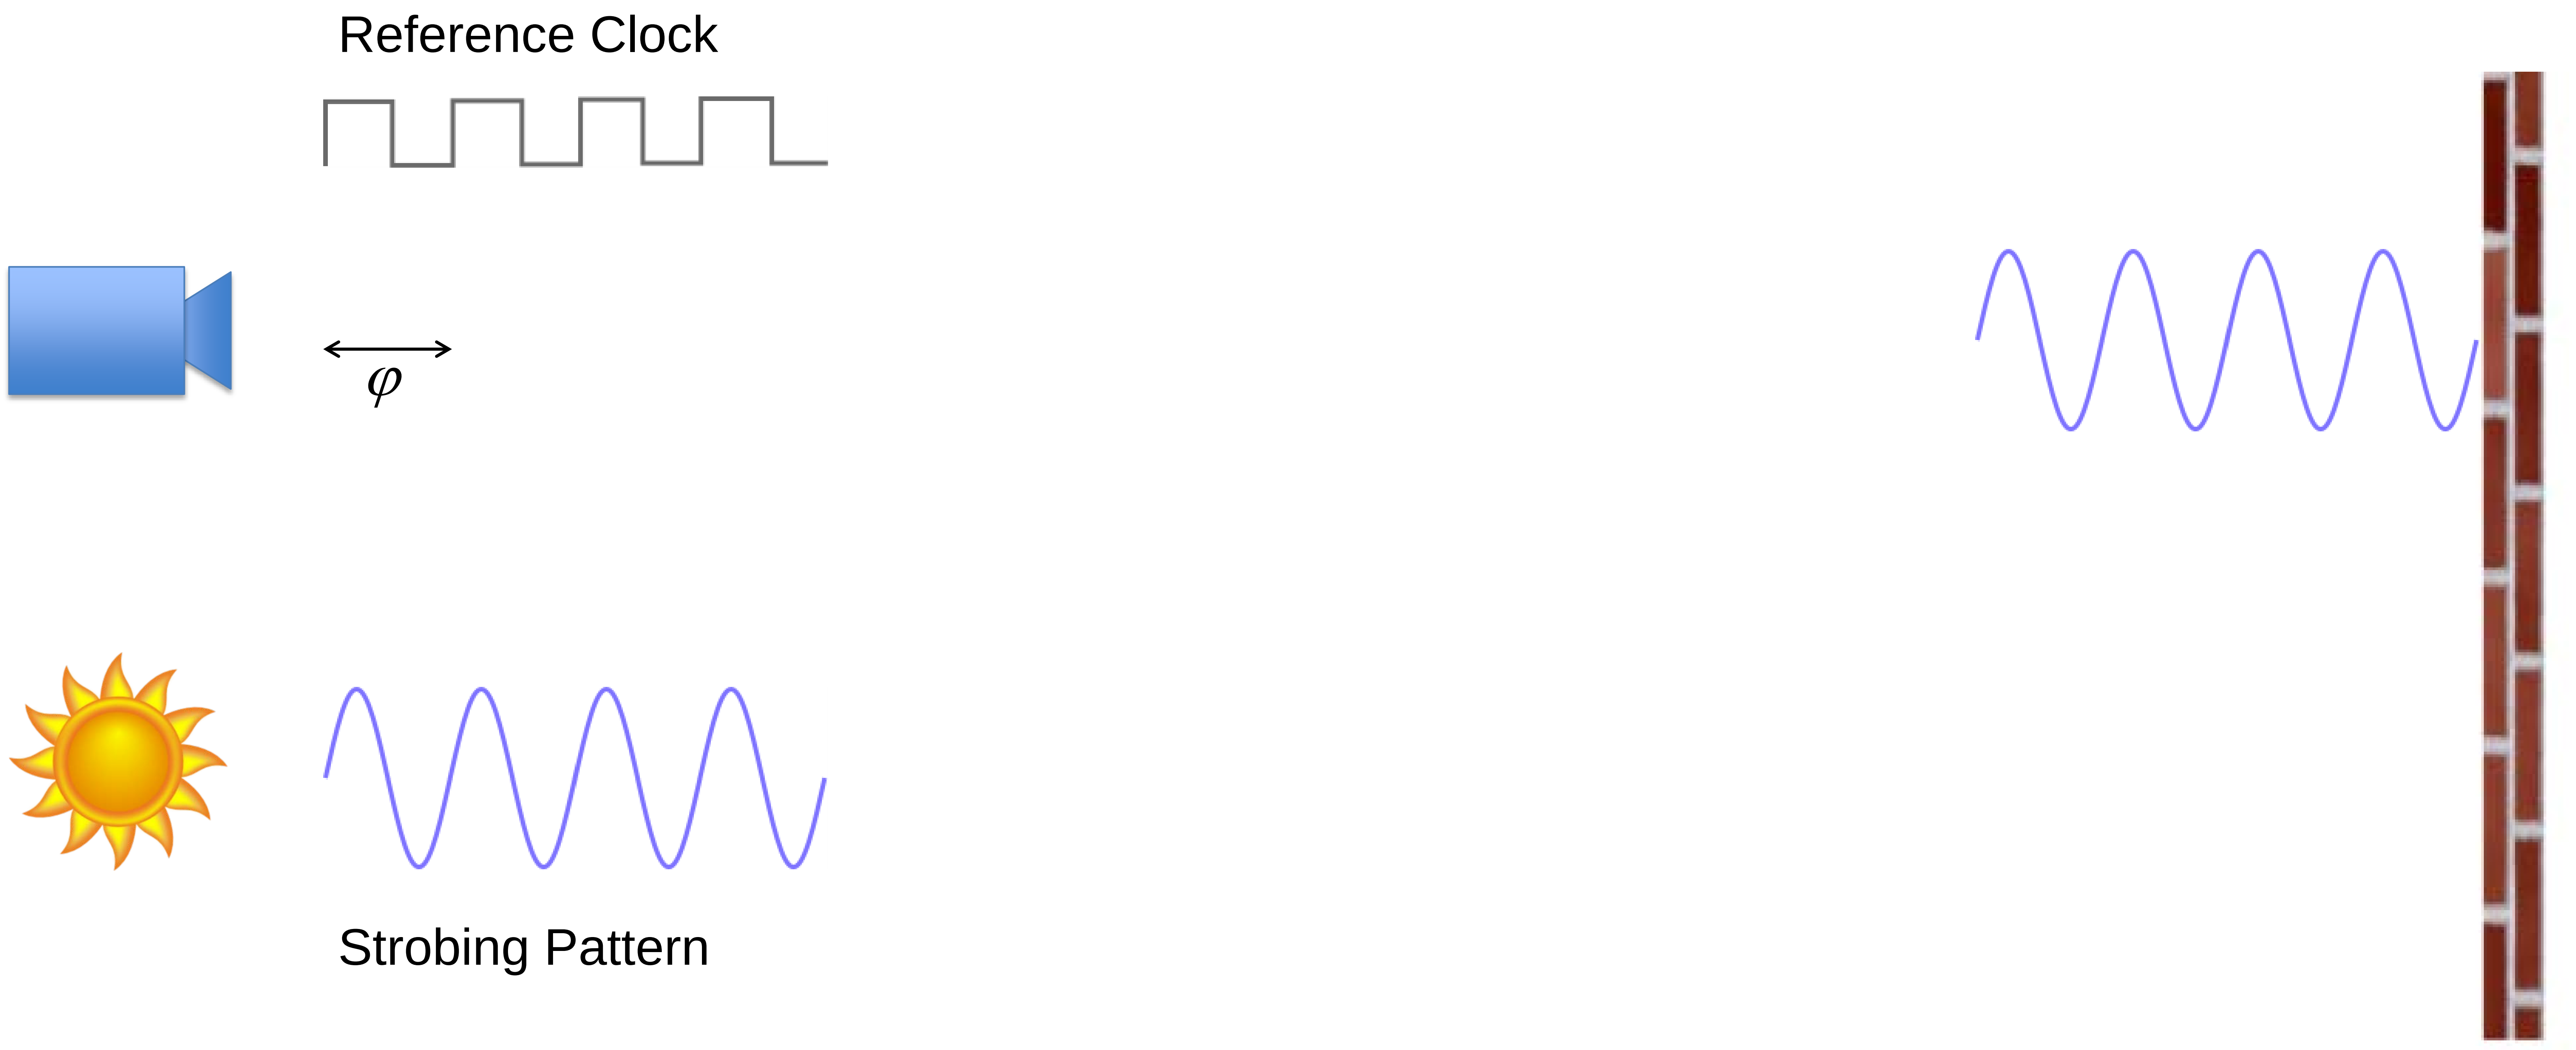

Recall: Strobing Pattern is MHz (nanosecond periods)

Achuta Kadambi
MIT/UCLA

Interference

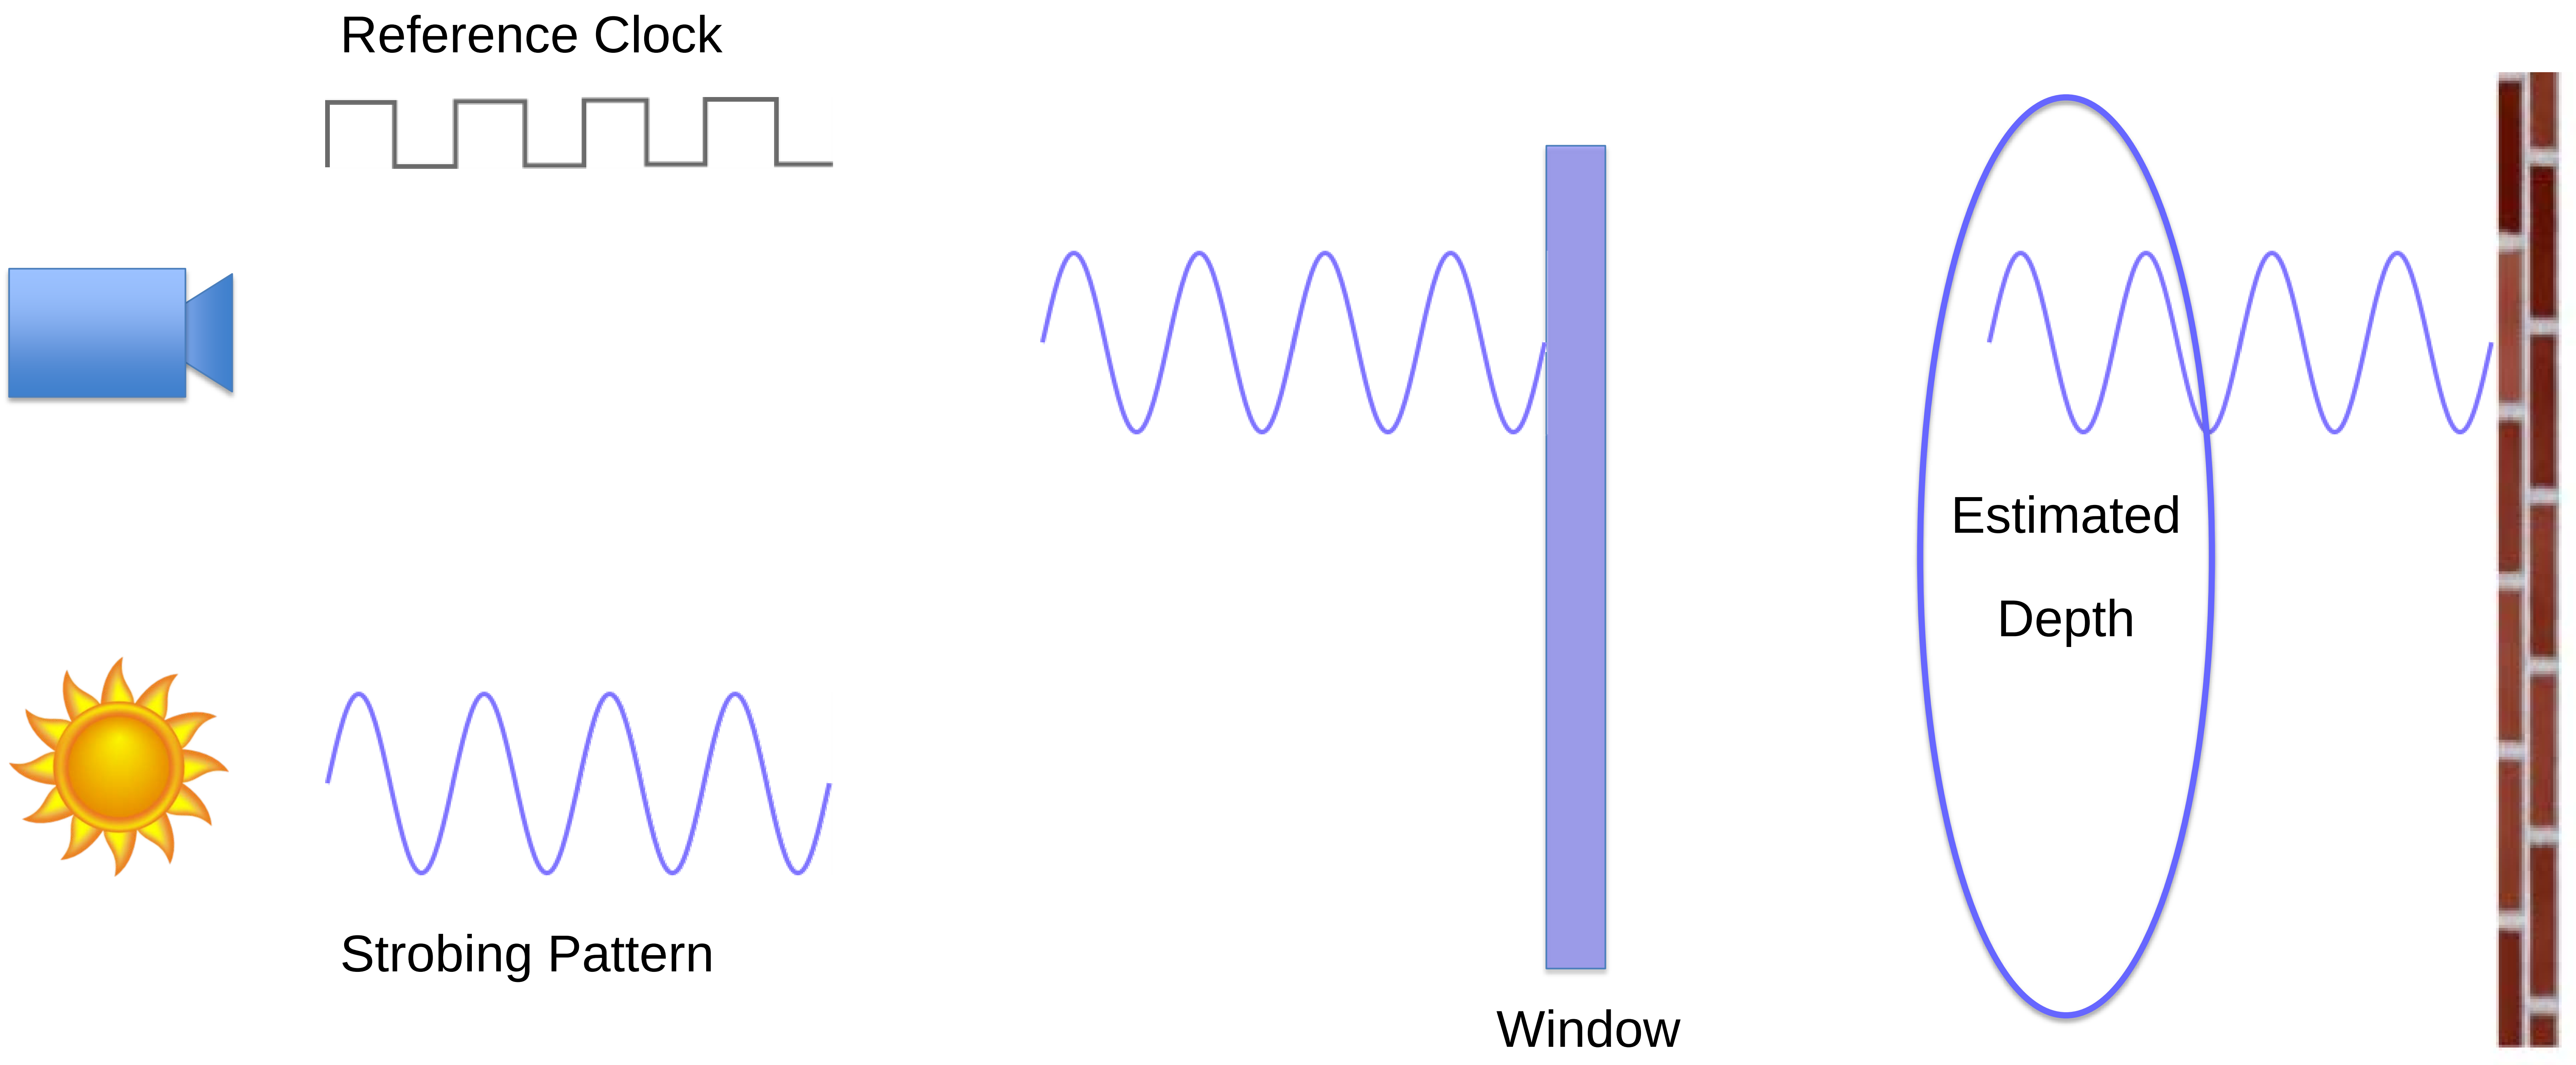

Measure one Phase that is in between (mod 2 \pi)

Recall: Strobing Pattern is MHz (nanosecond periods)

How the Phase ToF (Kinect) Works

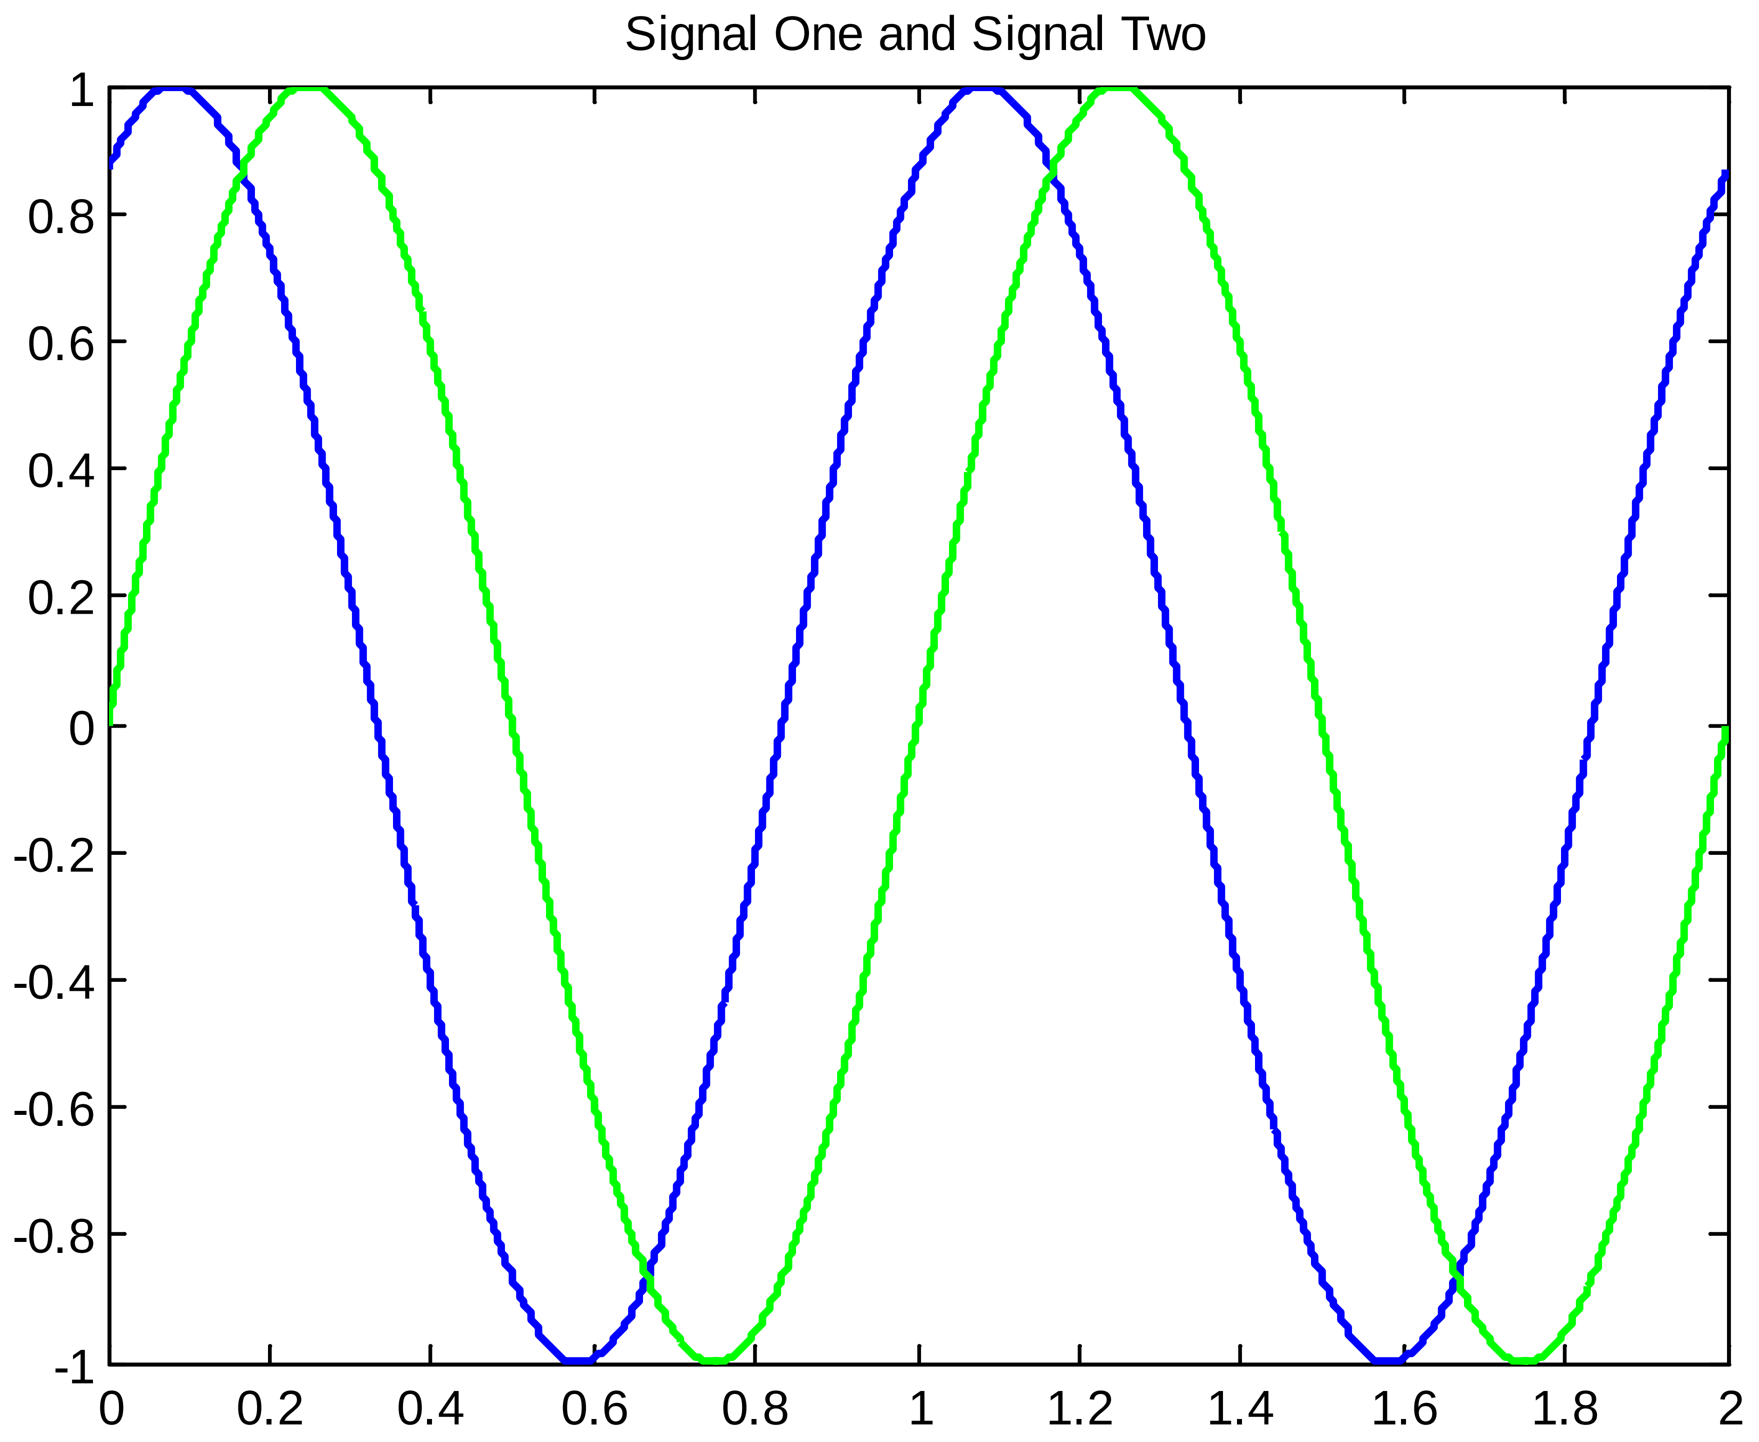

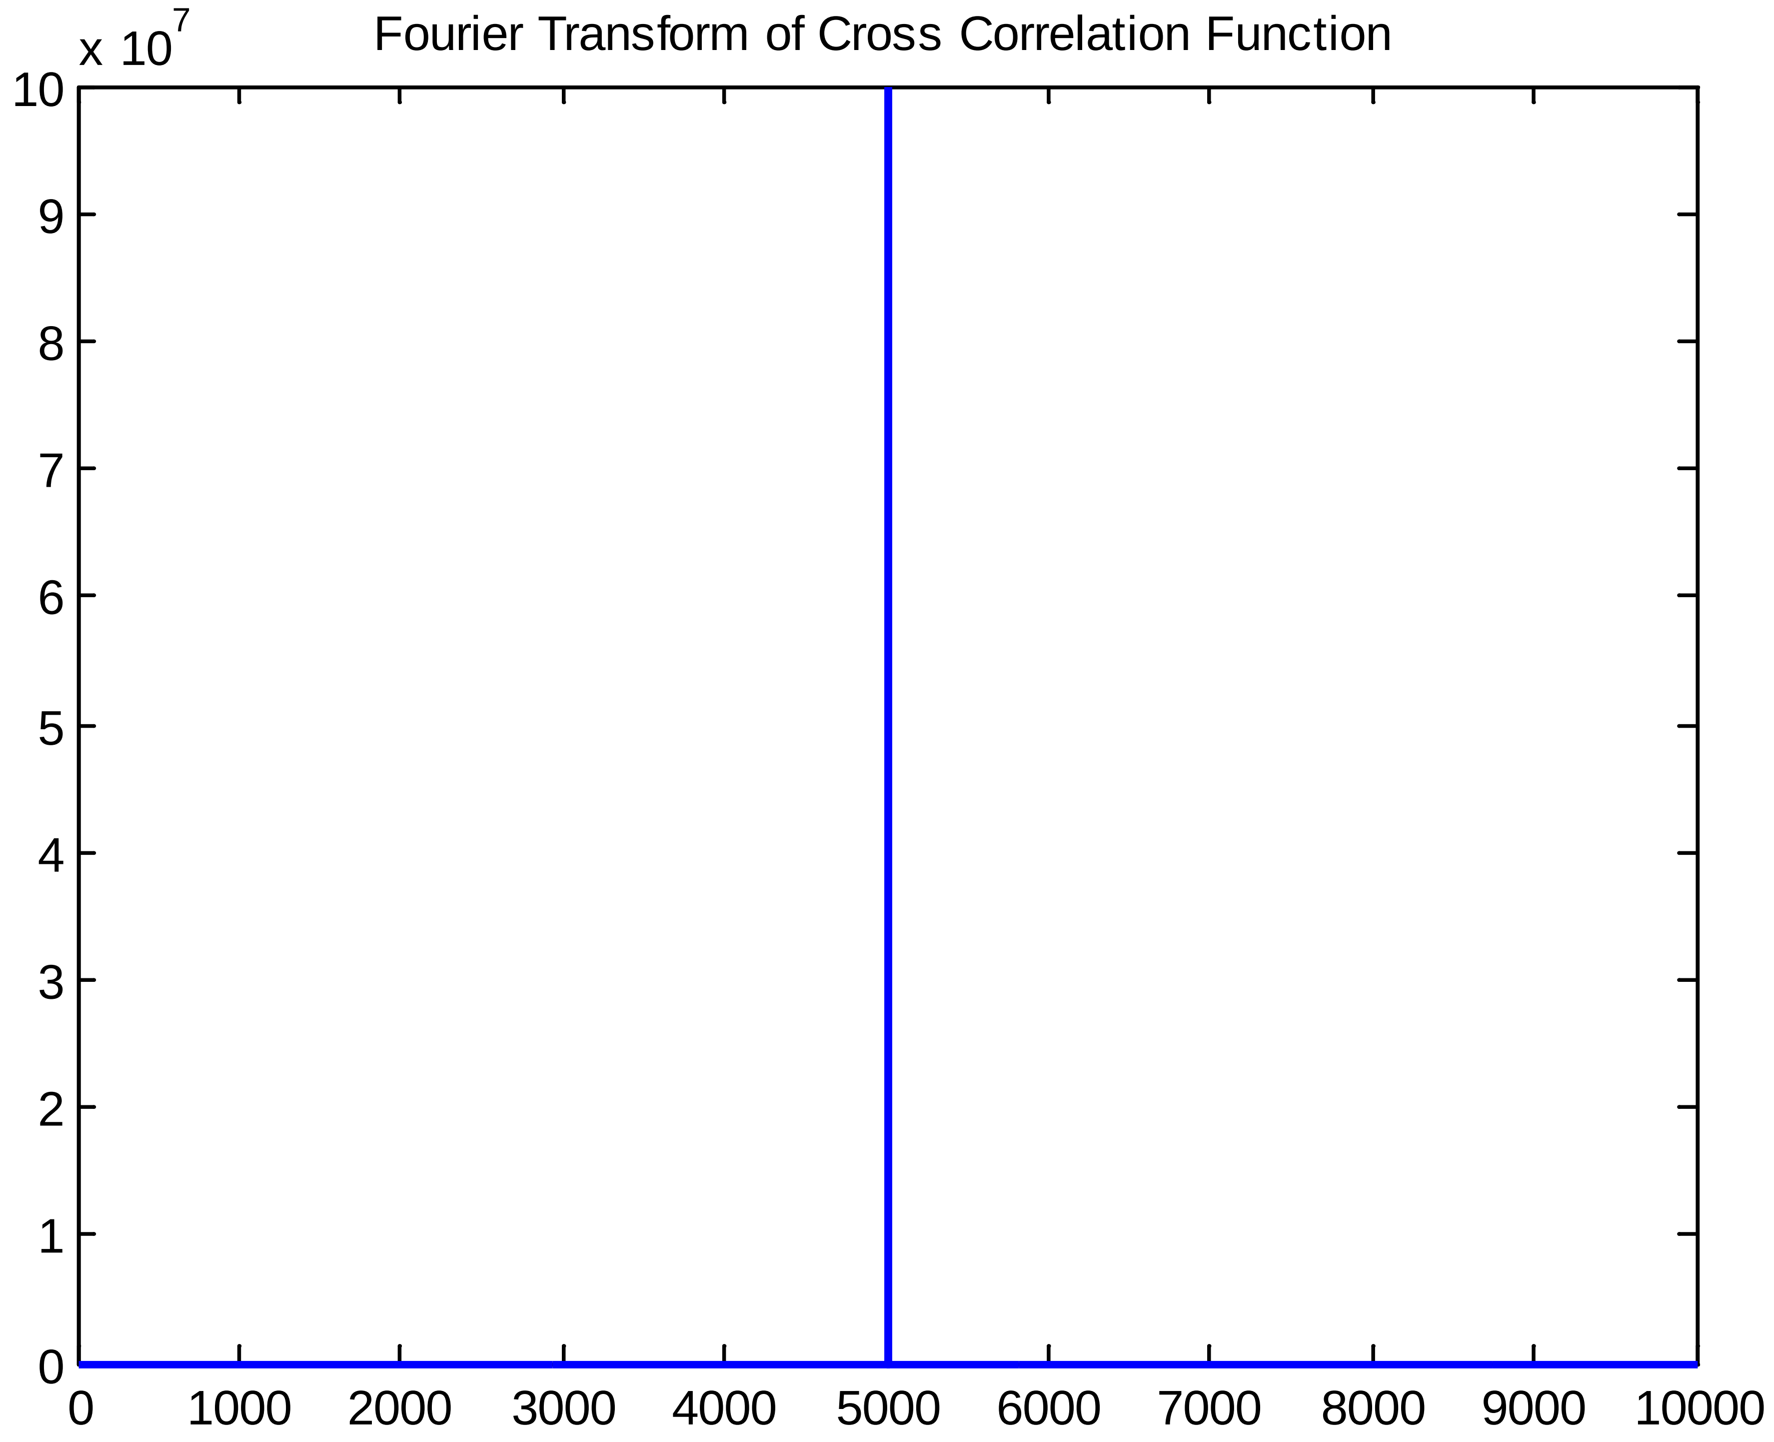


```
dt = 0.0001;
N = 20000;
t = 0:dt:(N-1)*dt;

phi1 = pi/3; phi2 = 0;
s1 = sin(2*pi*1*t + phi1);
s2 = sin(2*pi*1*t+ phi2);

xcFFT = conj( fft(s1) ) .*
fft(s2);

[~, fundamental_idx] =
max(xcFFT);

phase_difference = angle(
xcFFT(fundamental_idx) );
```

Recall: Strobing Pattern is MHz (nanosecond periods)

Each Pixel Becomes a Linear Time-invariant System

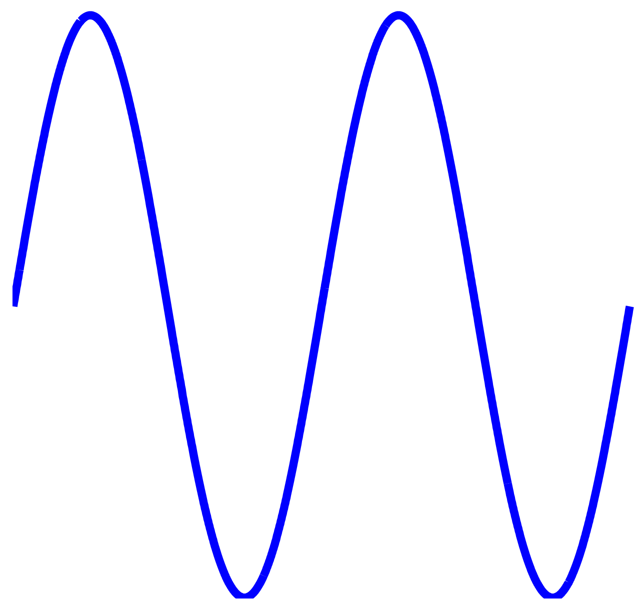

Probing Kernel

Each Pixel Becomes a Linear Time-invariant System

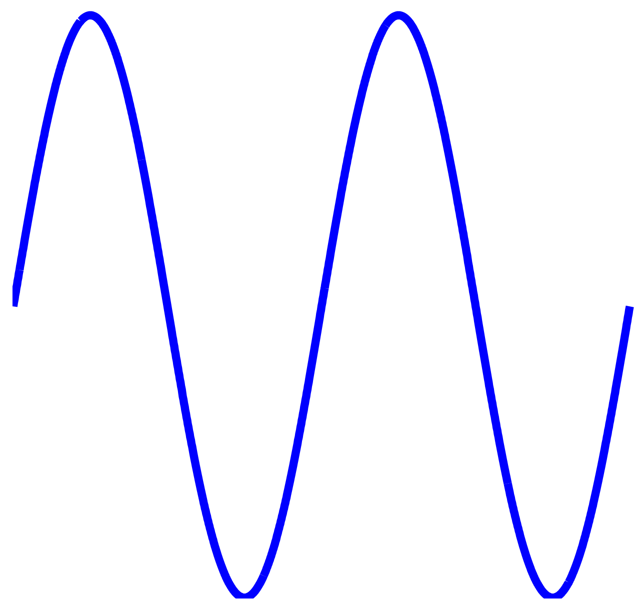

*

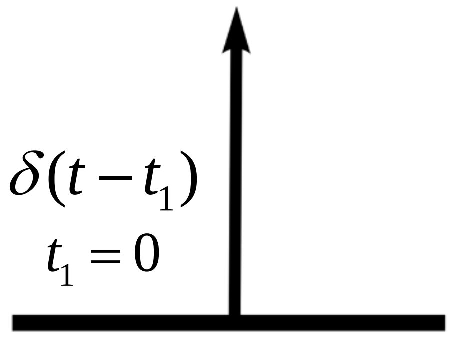

$\delta(t - t_1)$

$t_1 = 0$

Spike location due to optical path length

Probing Kernel

Temporal Response

Each Pixel Becomes a Linear Time-invariant System

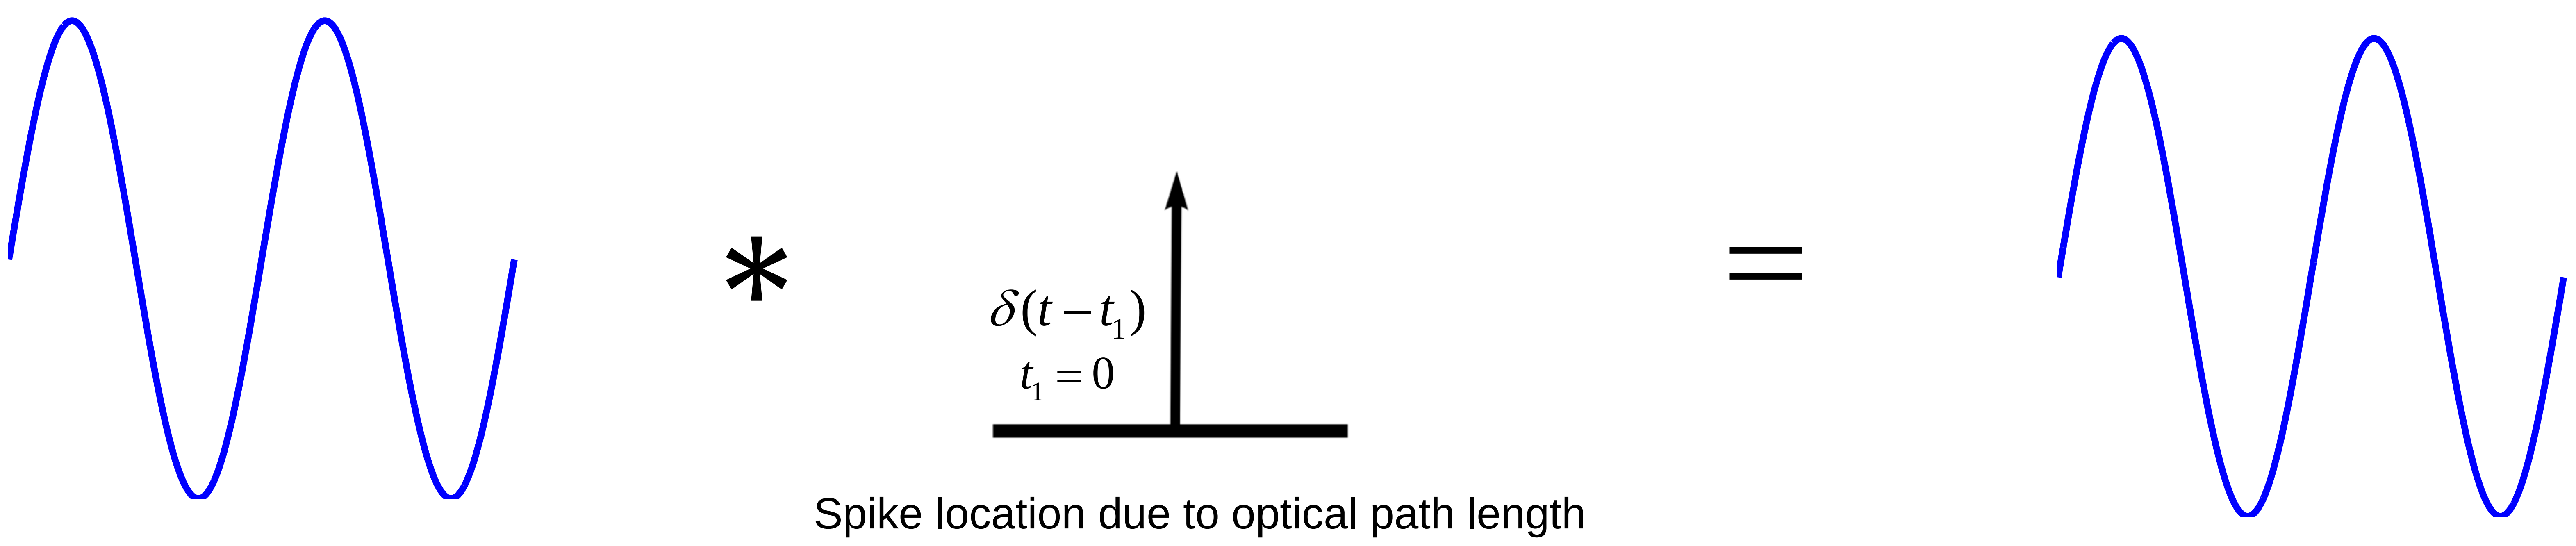

Probing Kernel

Temporal Response

Measured Signal

Each Pixel Becomes a Linear Time-invariant System

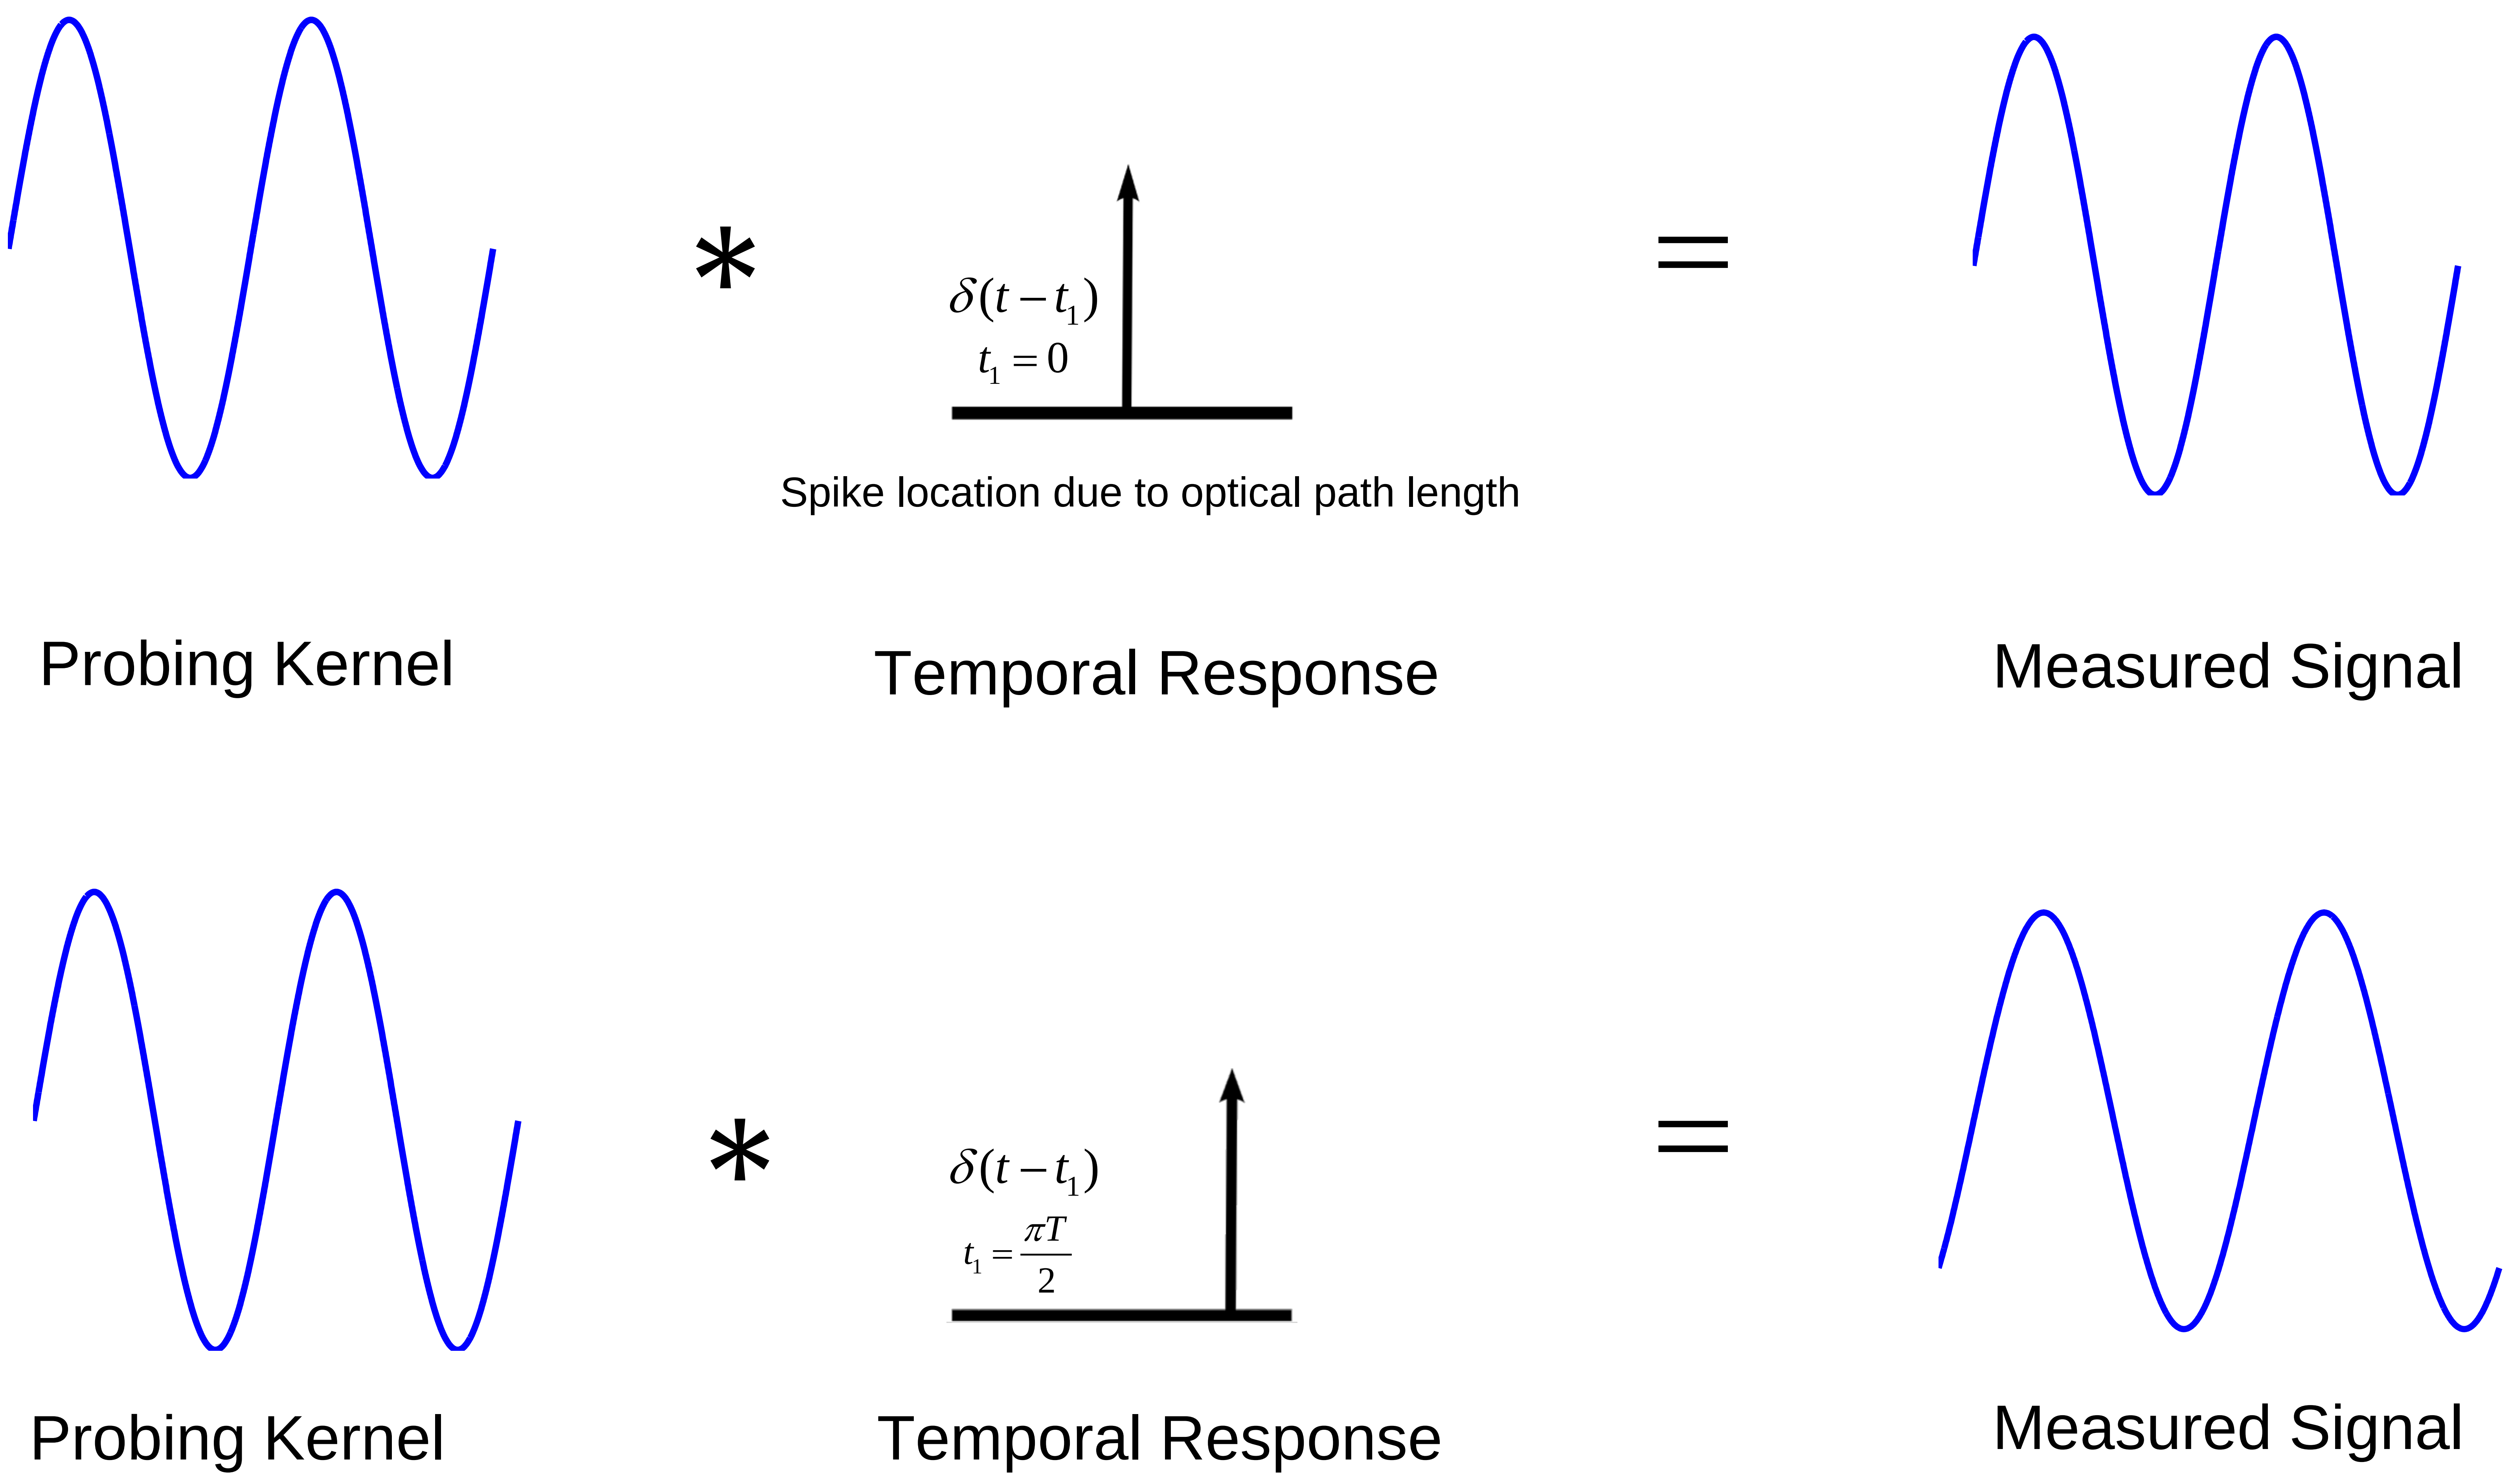

Multipath Interference

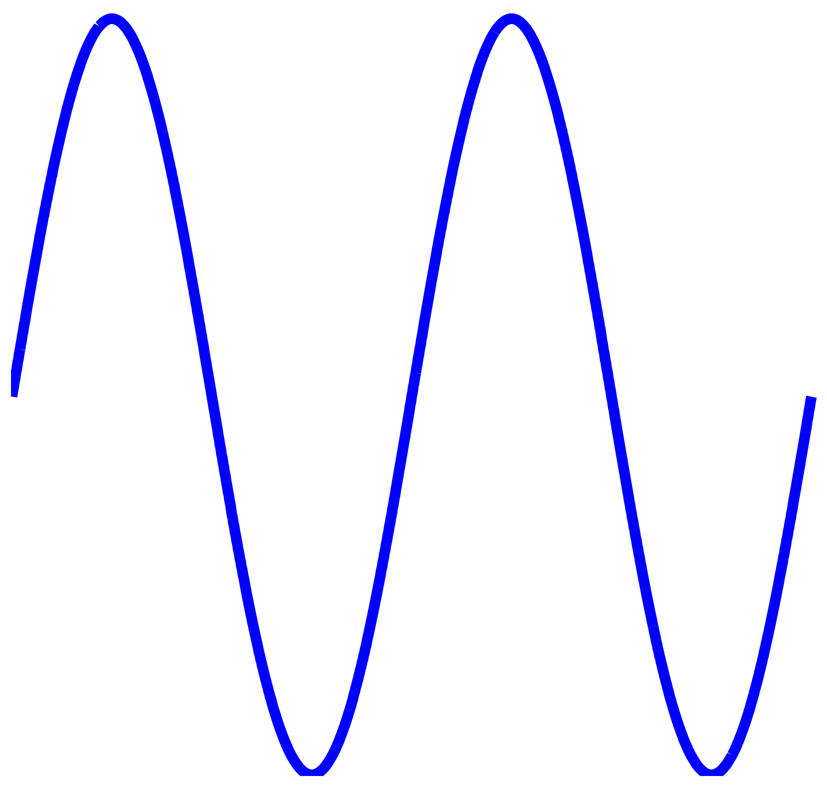

Probing Kernel

Multipath Interference

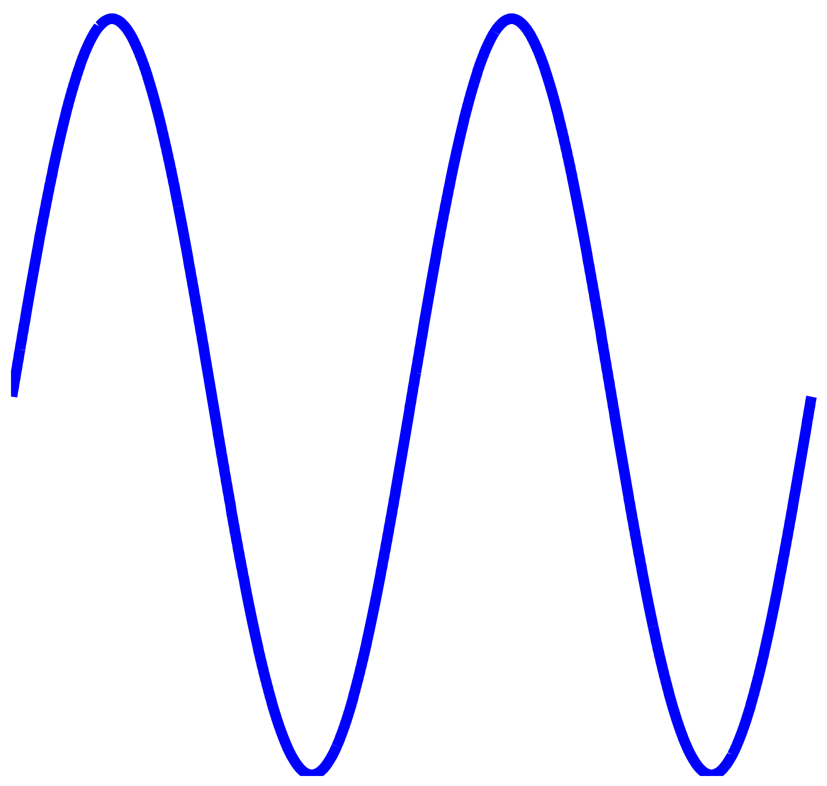

Probing Kernel

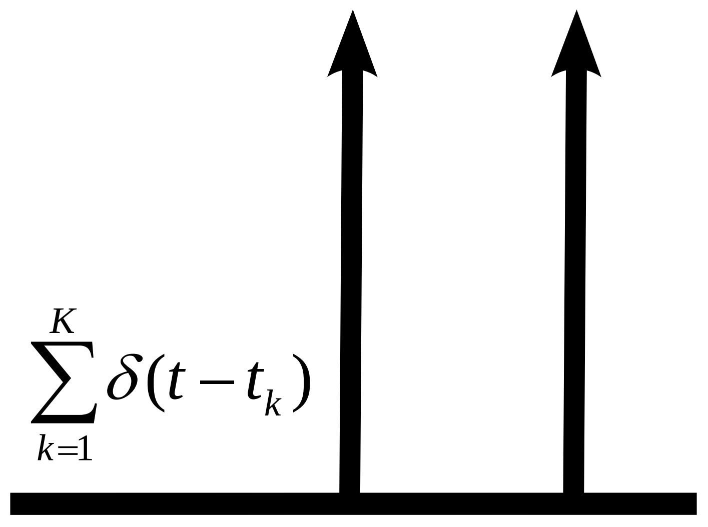

Temporal Response

Multipath Interference

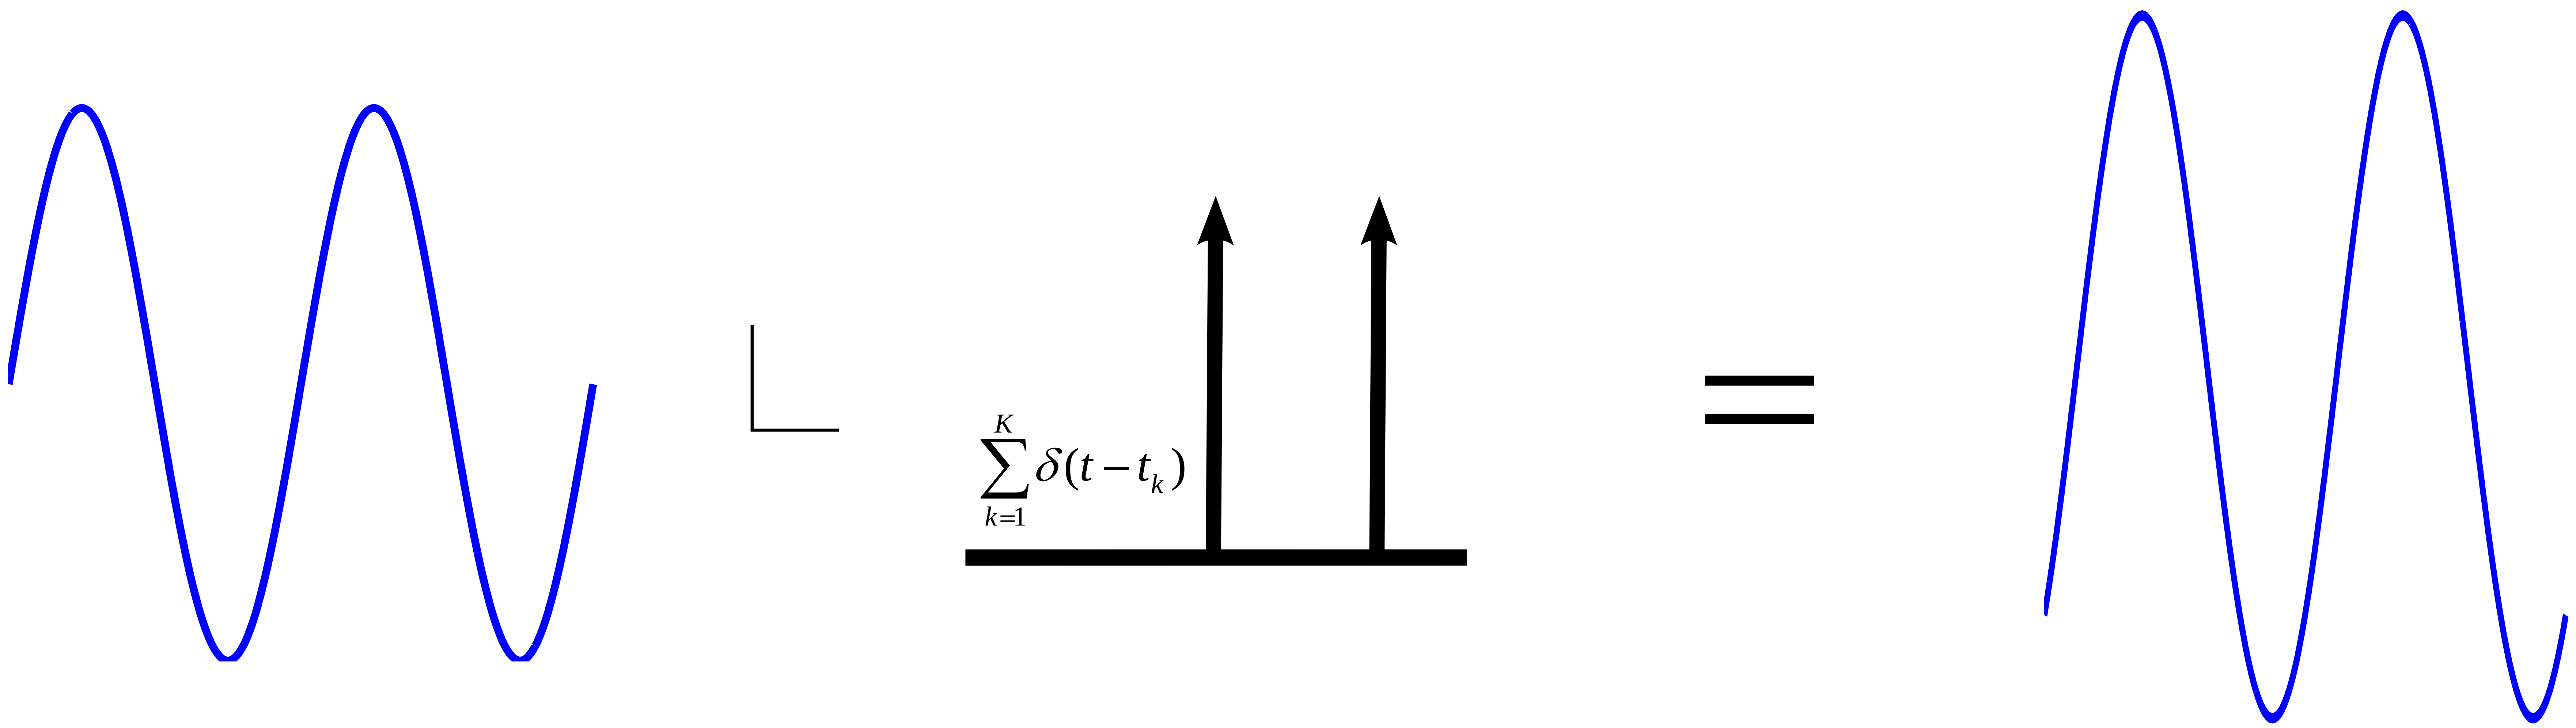

Probing Kernel | Temporal Response | Measured Signal

Problem of **unicity**: sum of two sines at same frequency but different phase is a single sine wave at a mixed phase

Multipath Interference

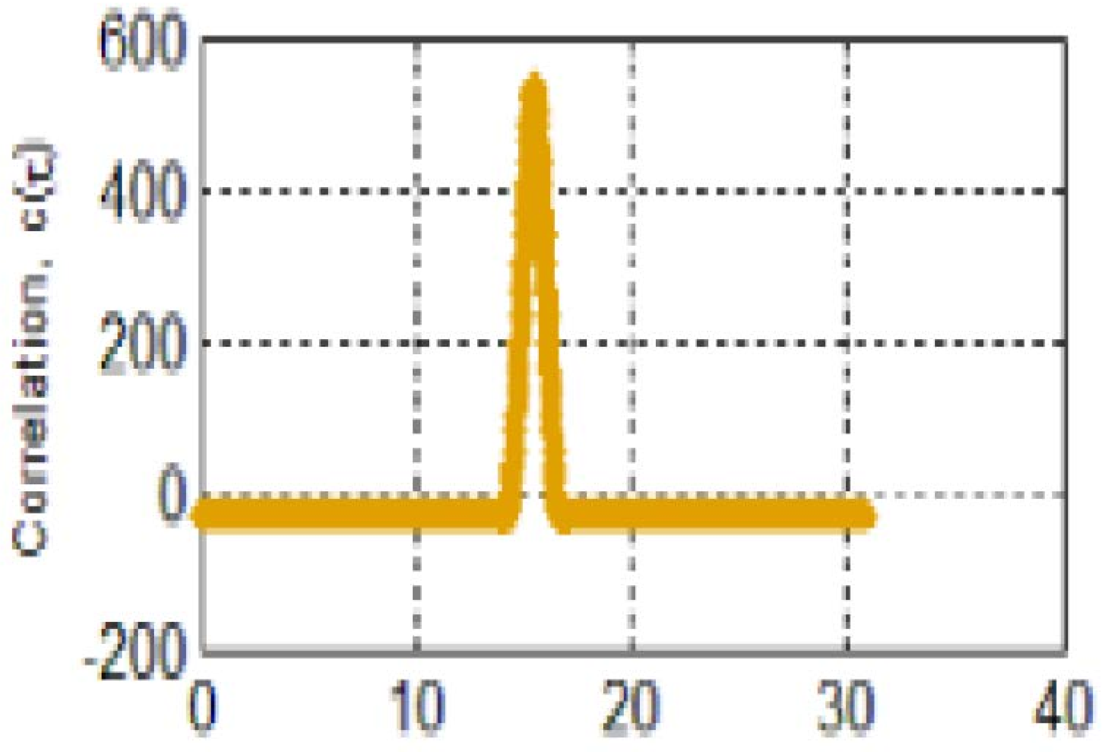

Probing Kernel

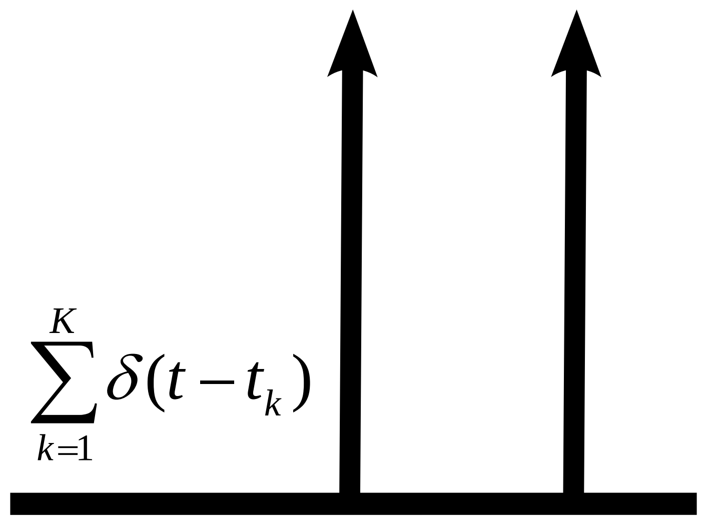

Temporal Response

=

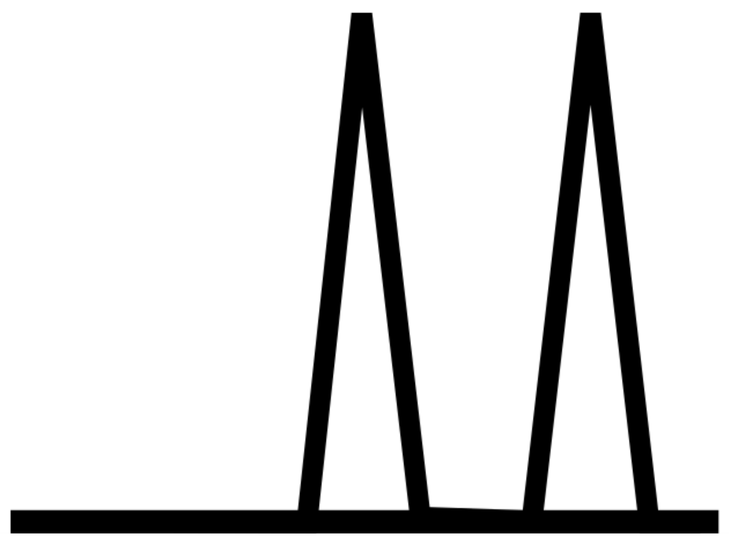

Measured Signal

This is a cartoon

What AM modulation is optimal?

Camera becomes comms device, repurposing theory from classic EECS fields (radar, telecom)

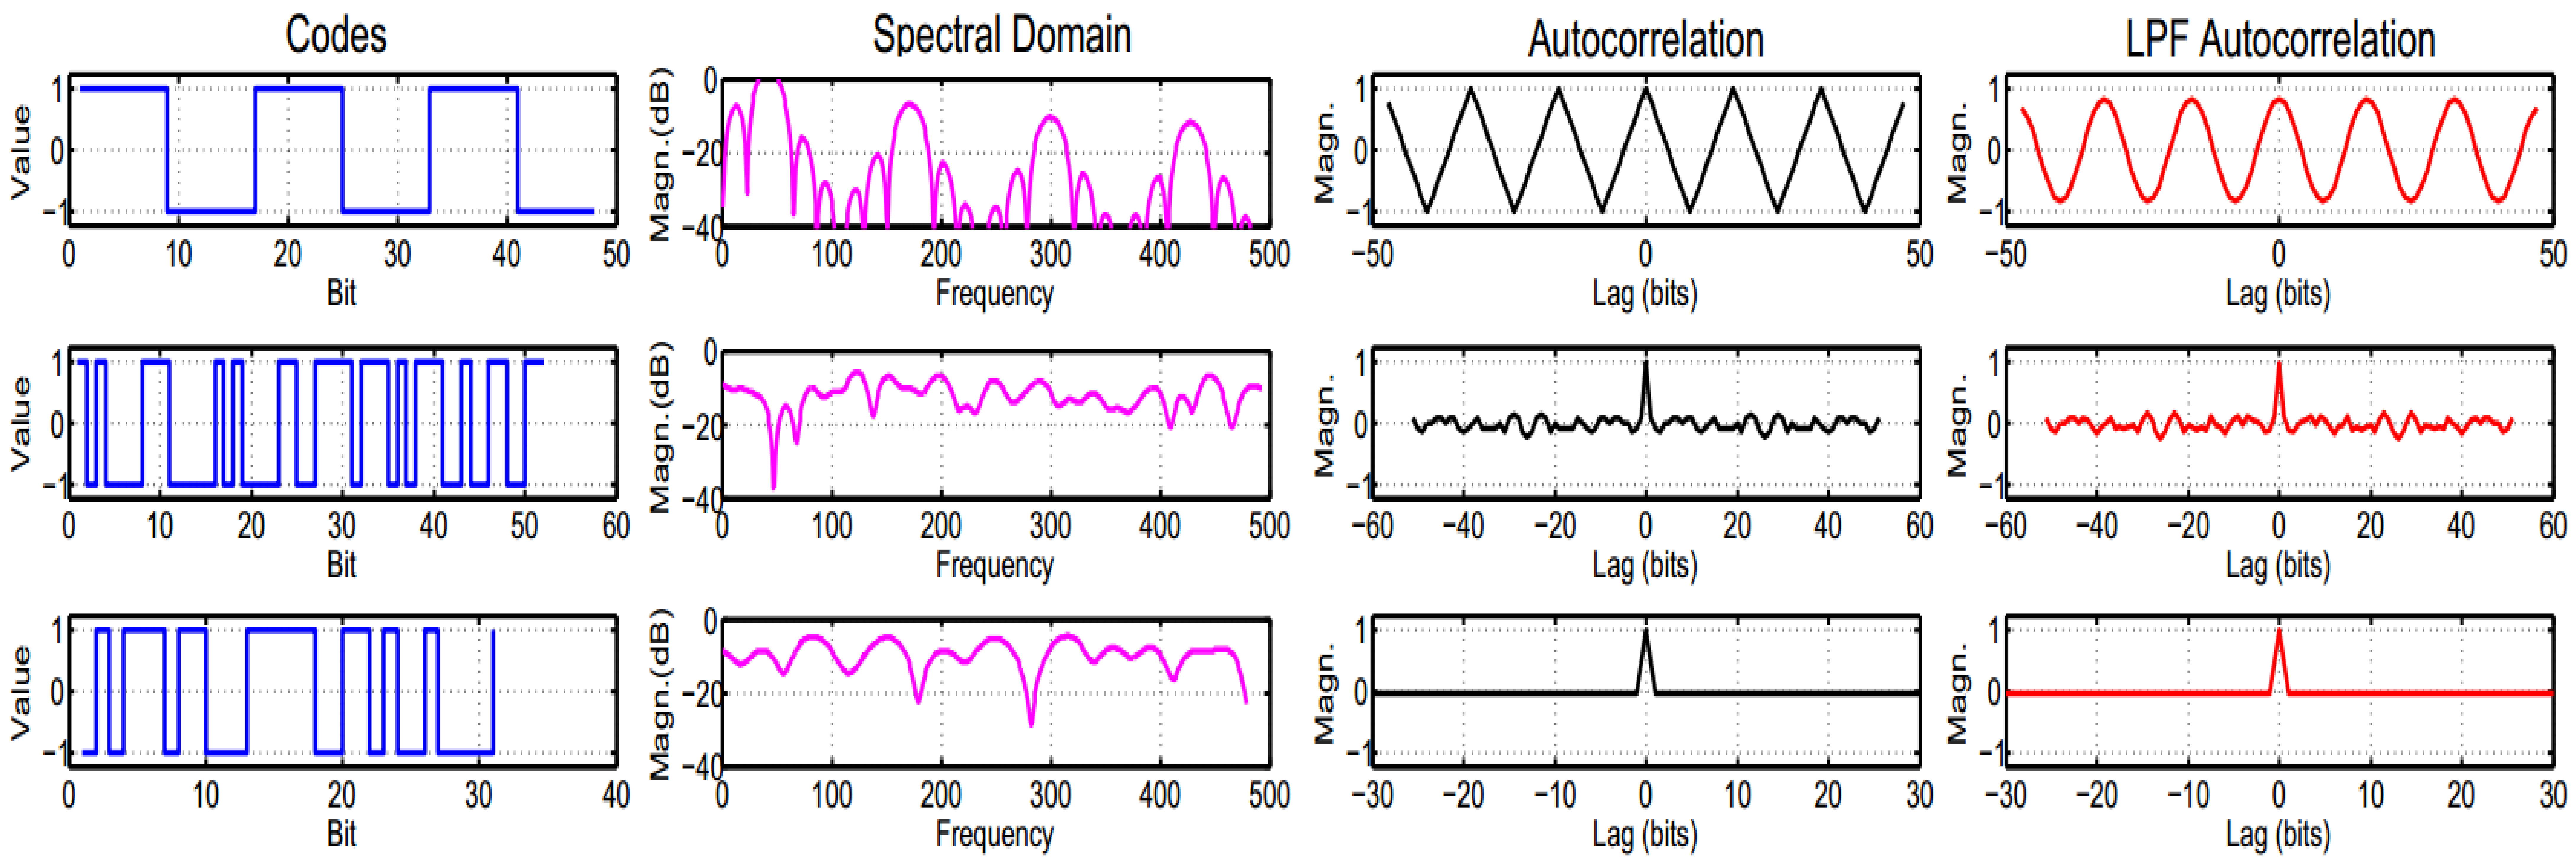

Design → Implementation

Nanophotography: AM modulation with customizable signal encoding

Siggraph Asia 2013

Empirical Validation of Resolution

Position of camera, geometry of checkered pattern known.

Validation of 100 picoseconds for singlepath

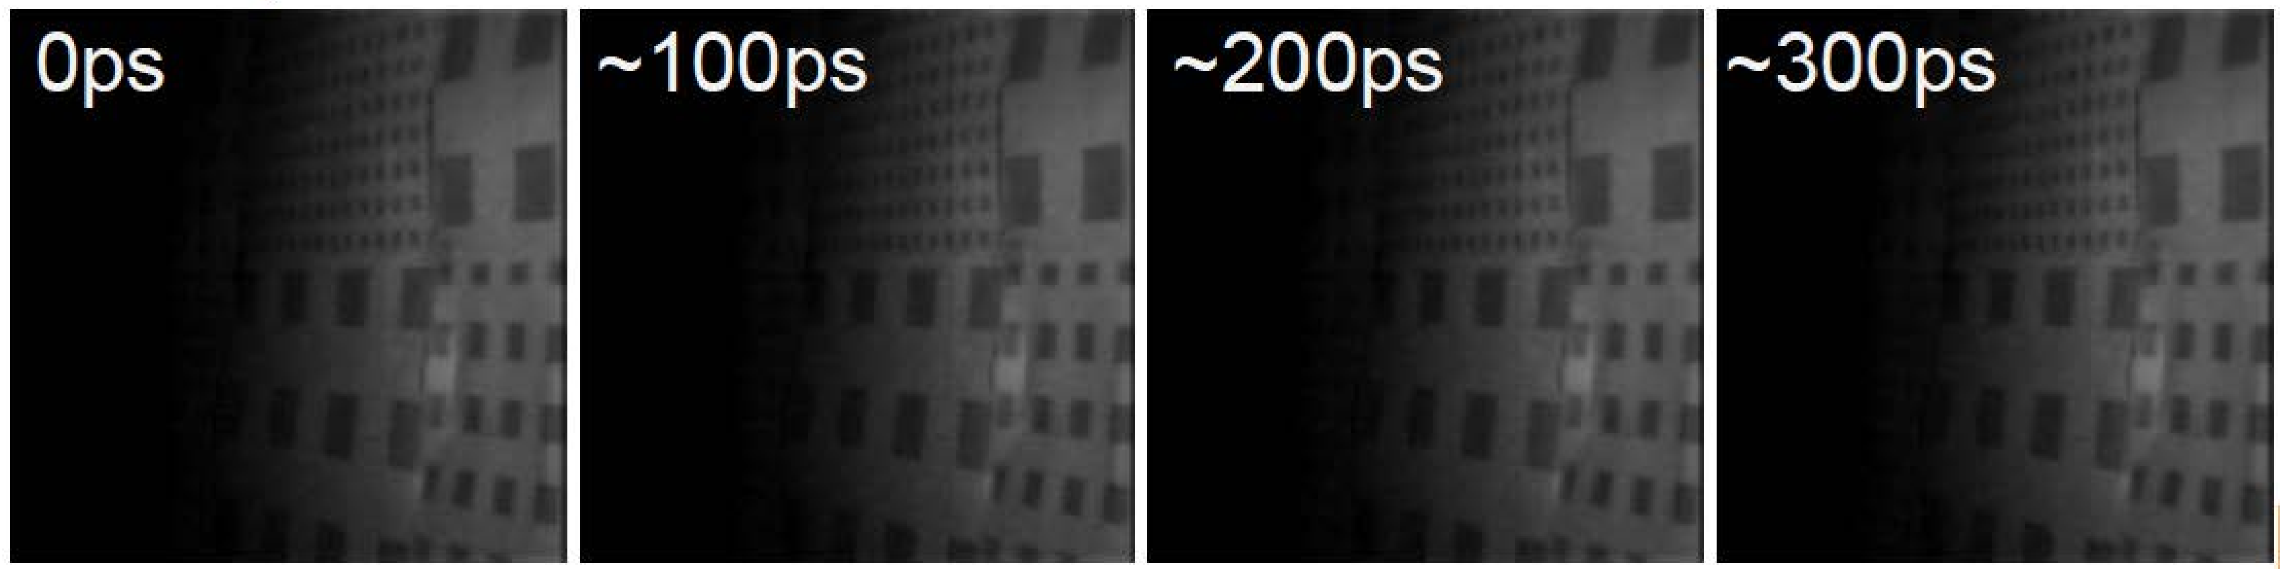

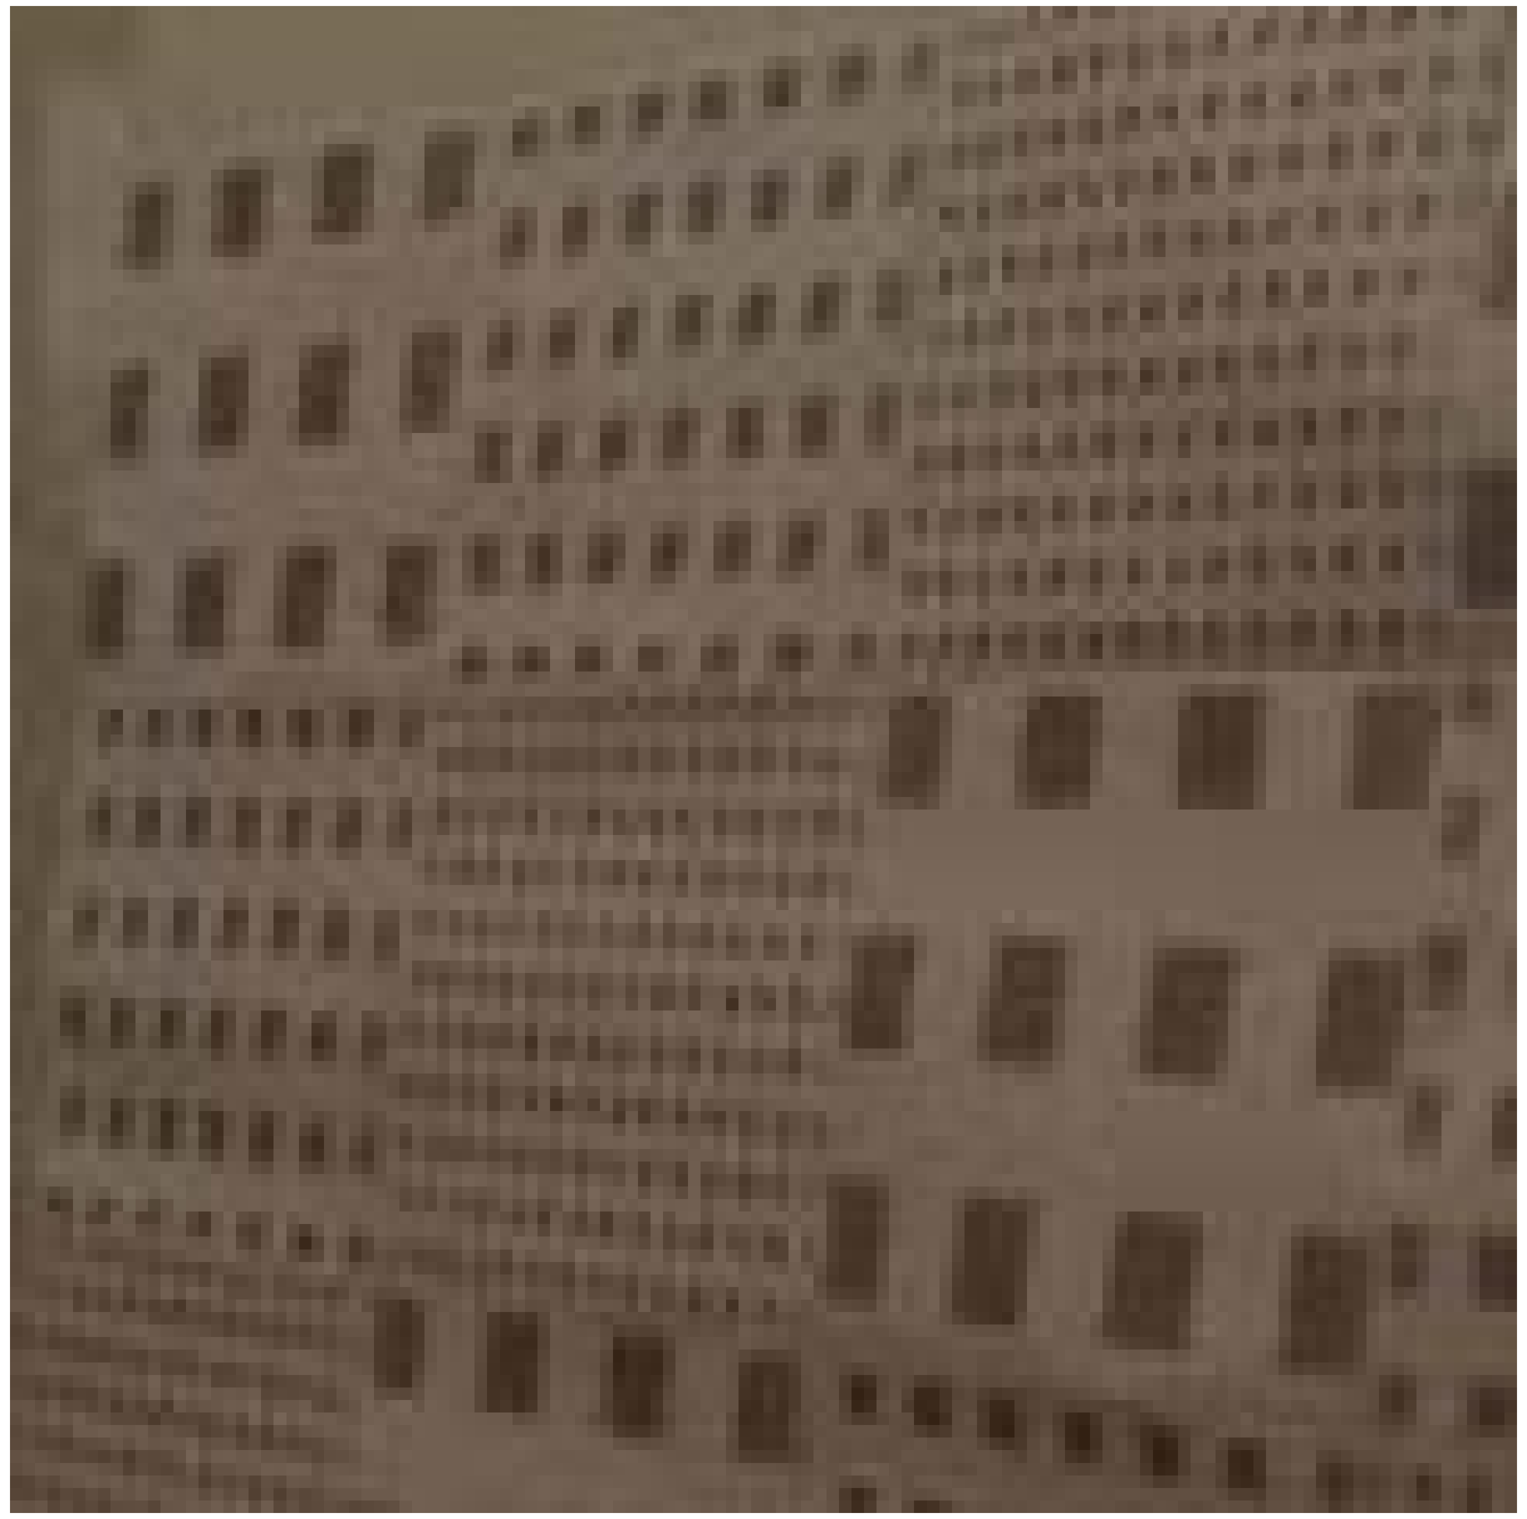

Siggraph Asia 2013

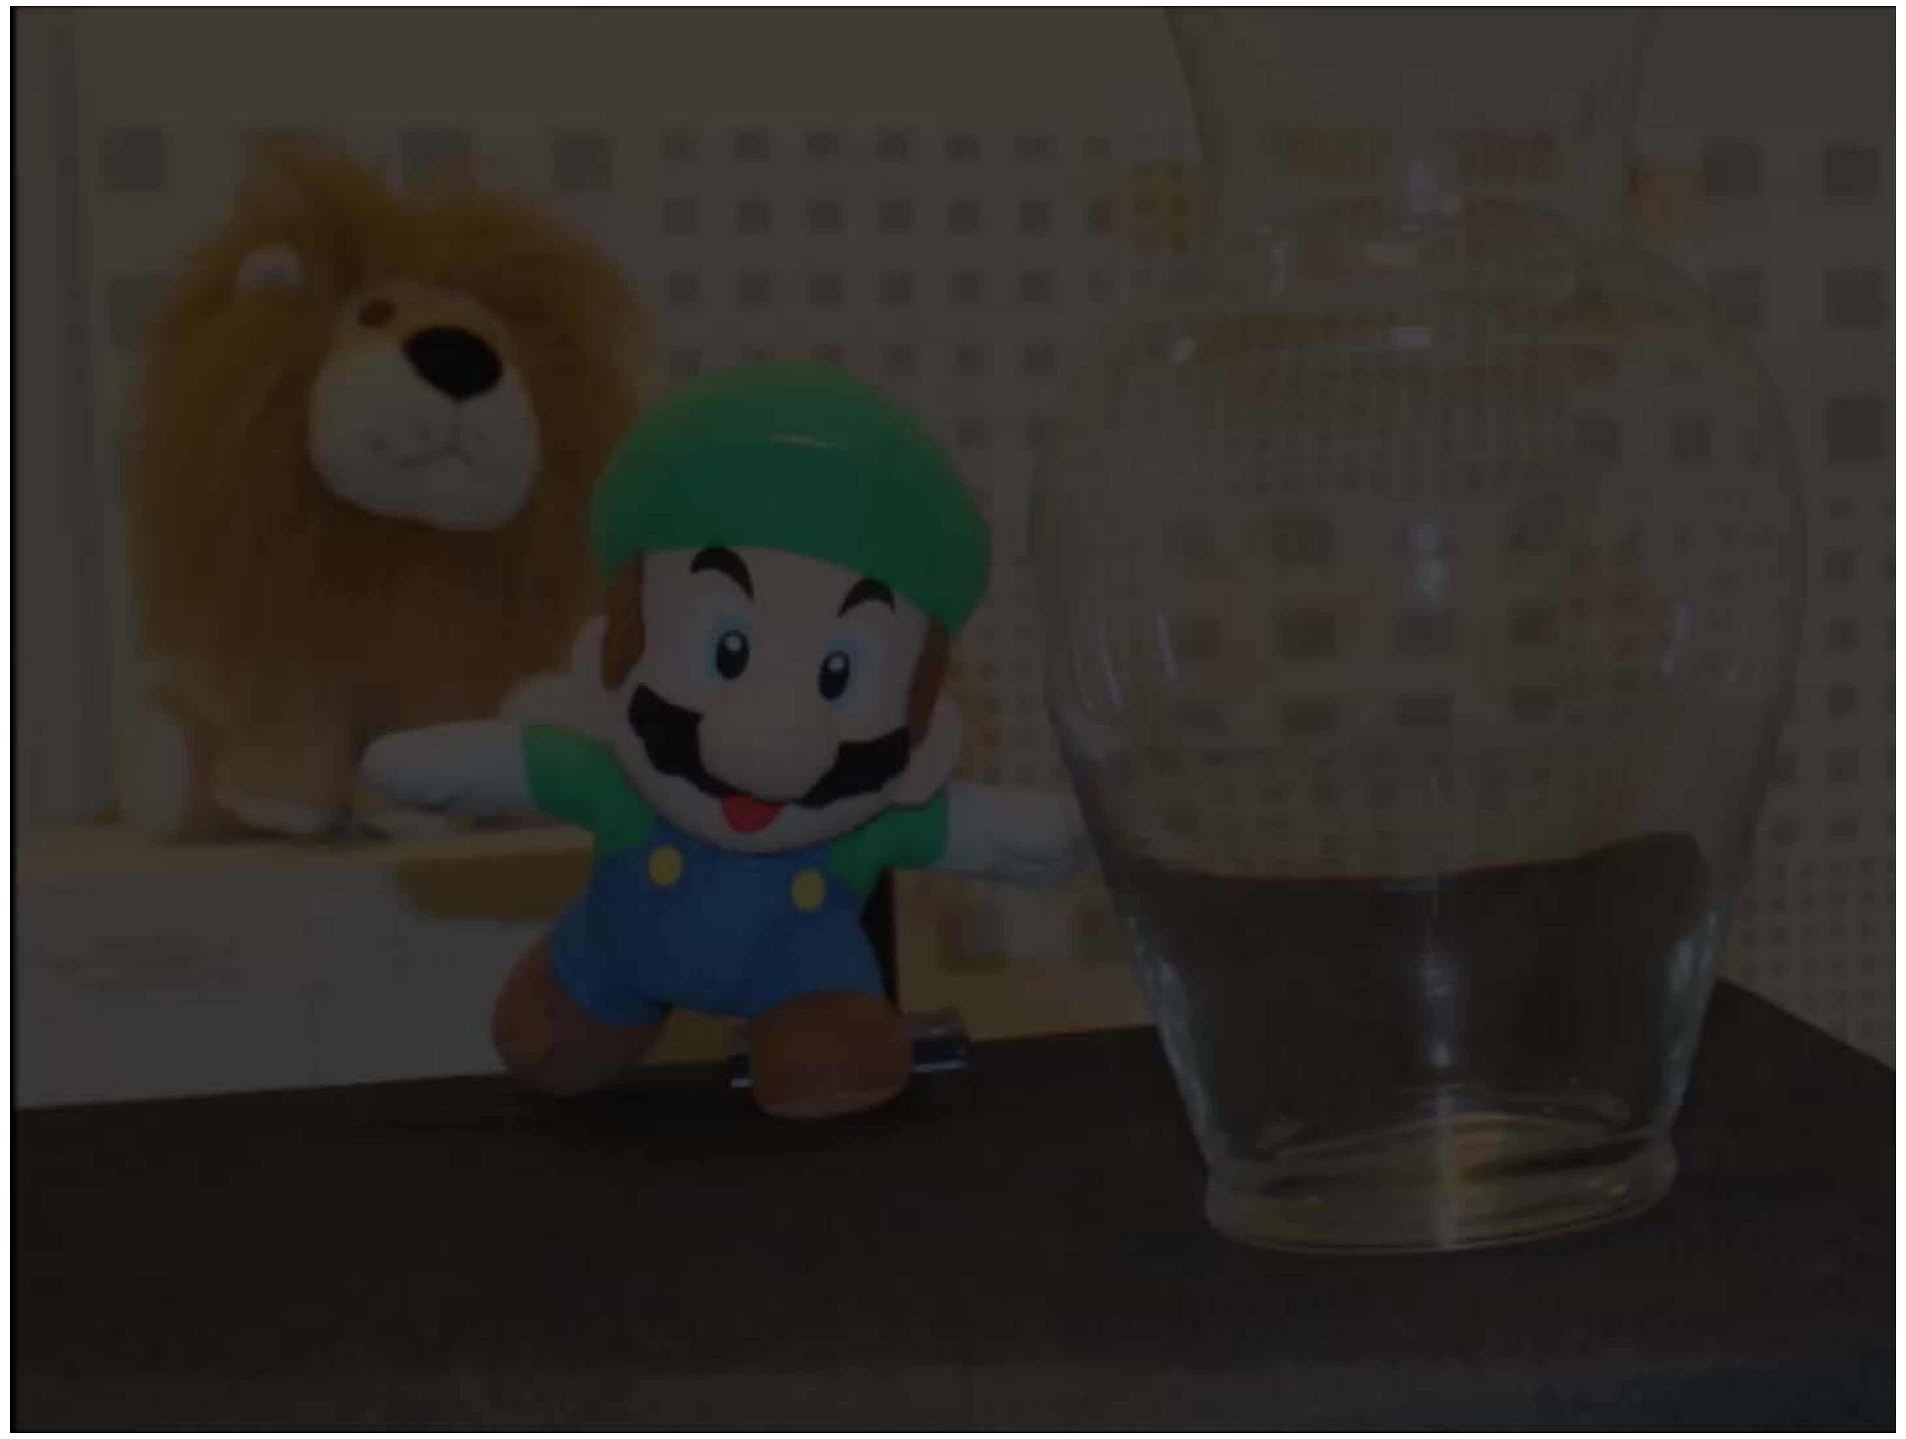

Siggraph Asia 2013

Why is multipath separation important?

Error Suppression

Scene

Measured Phase Image

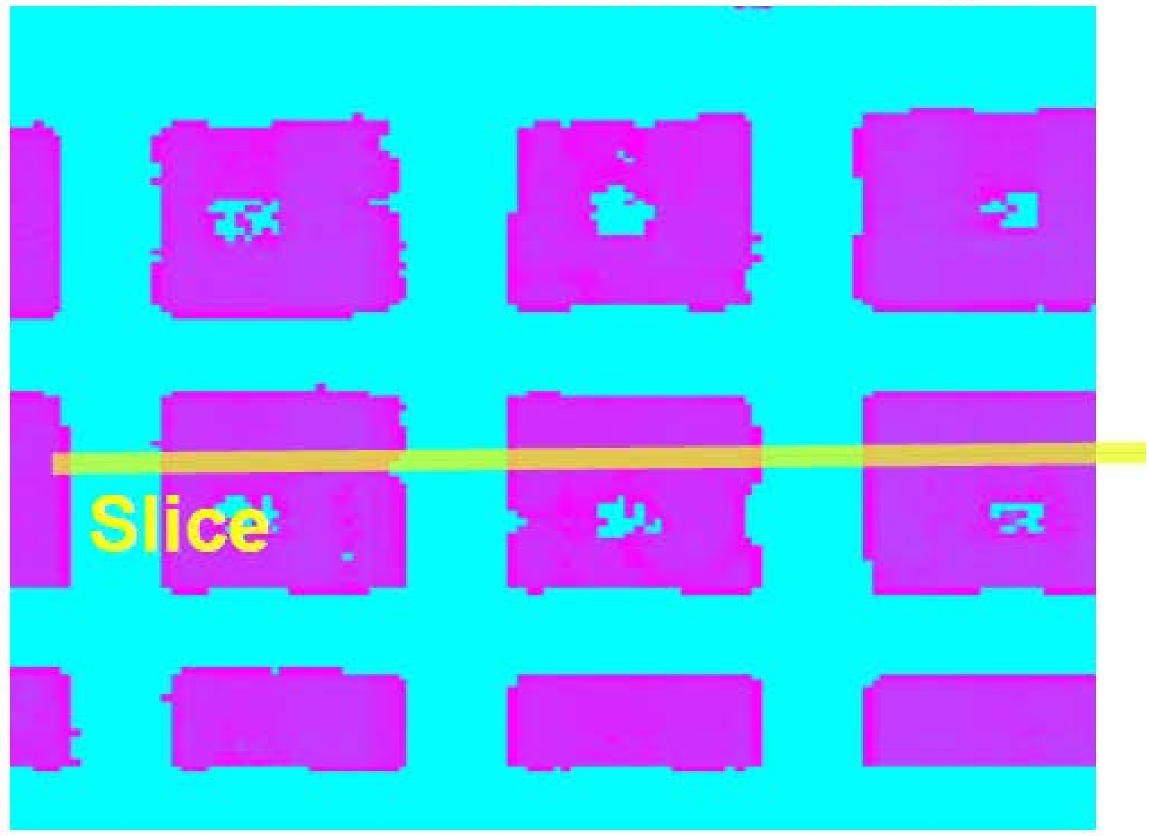

Error Suppression

Scene

Measured Phase Image

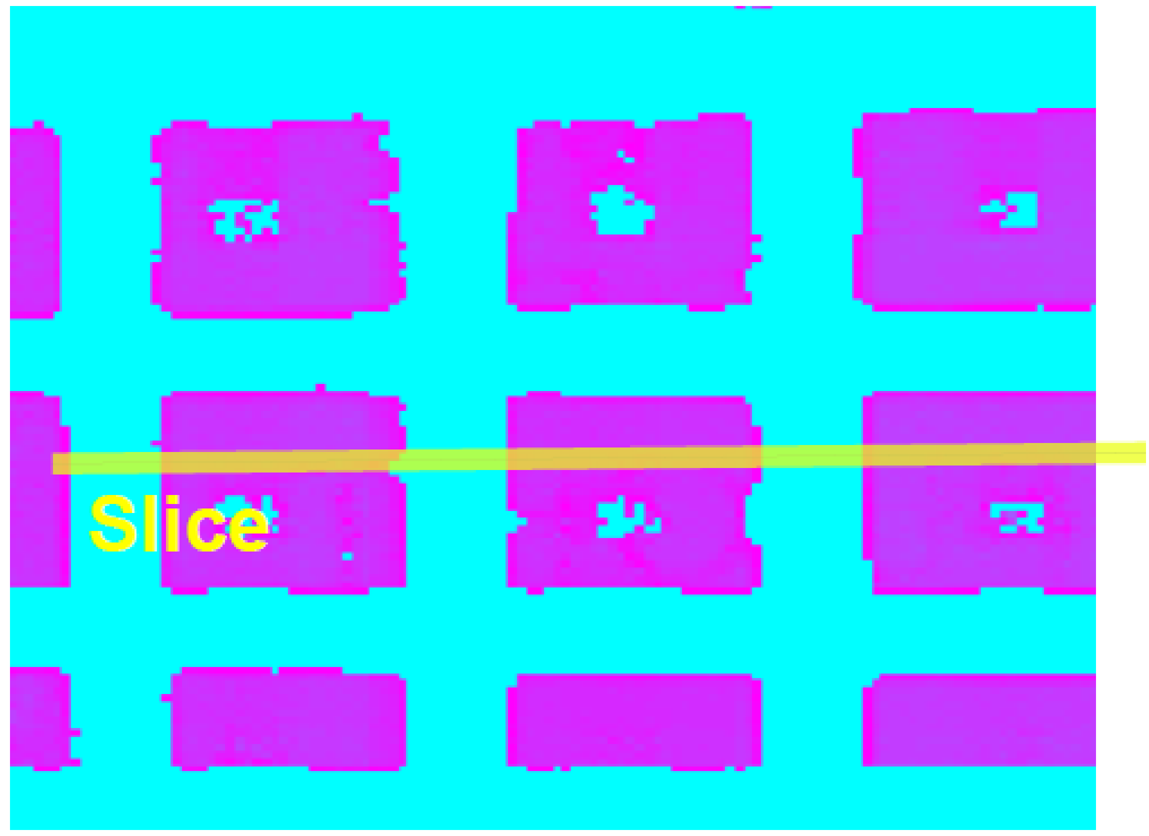

Measured Phase Slice

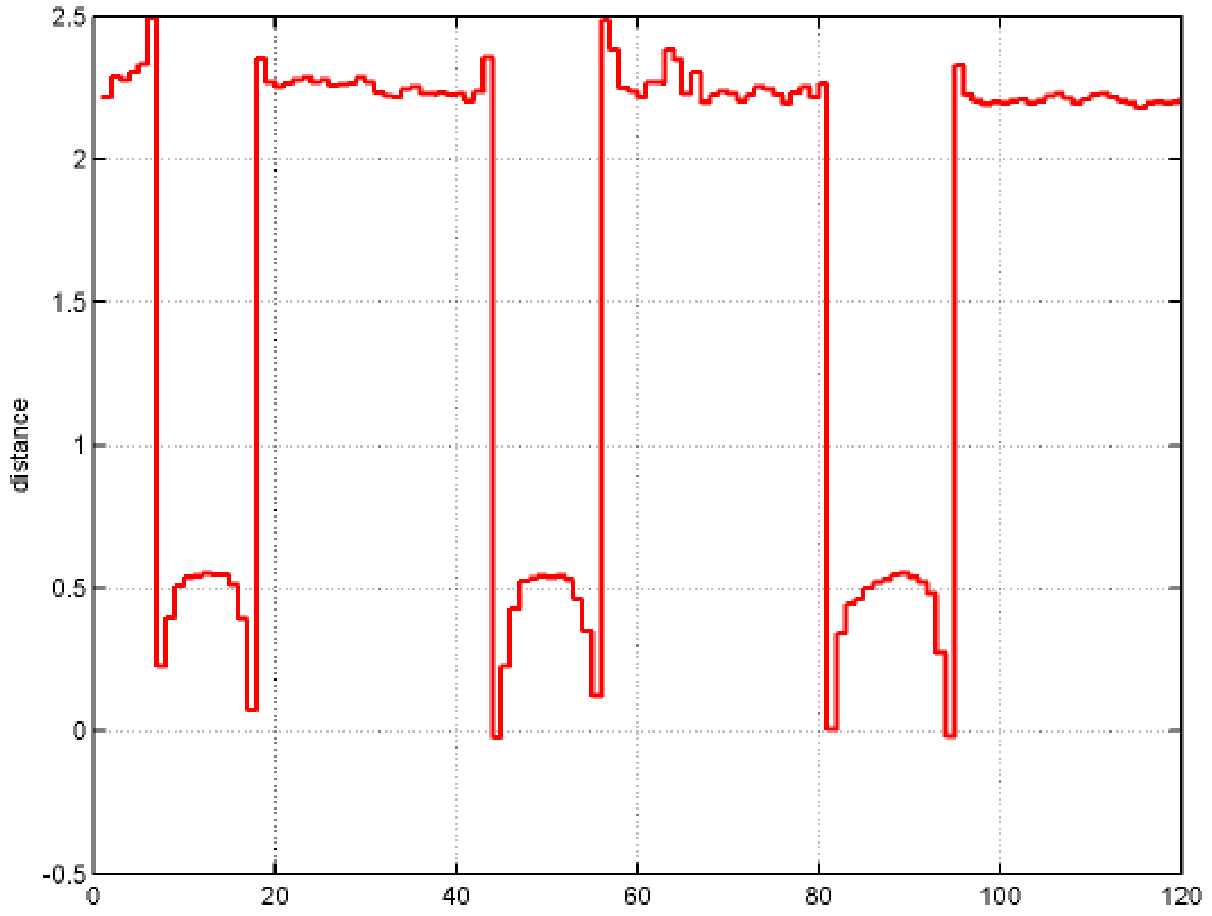

Error Suppression

Scene

Measured Phase Image

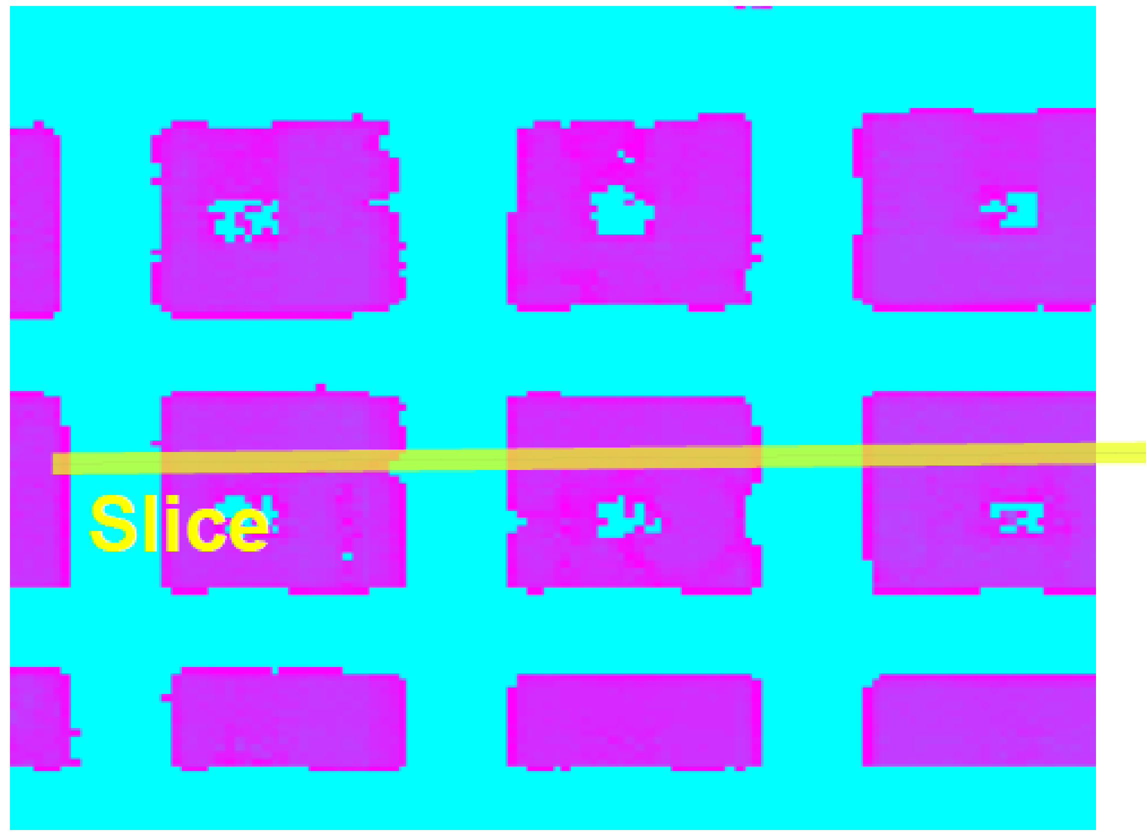

Measured Phase Slice

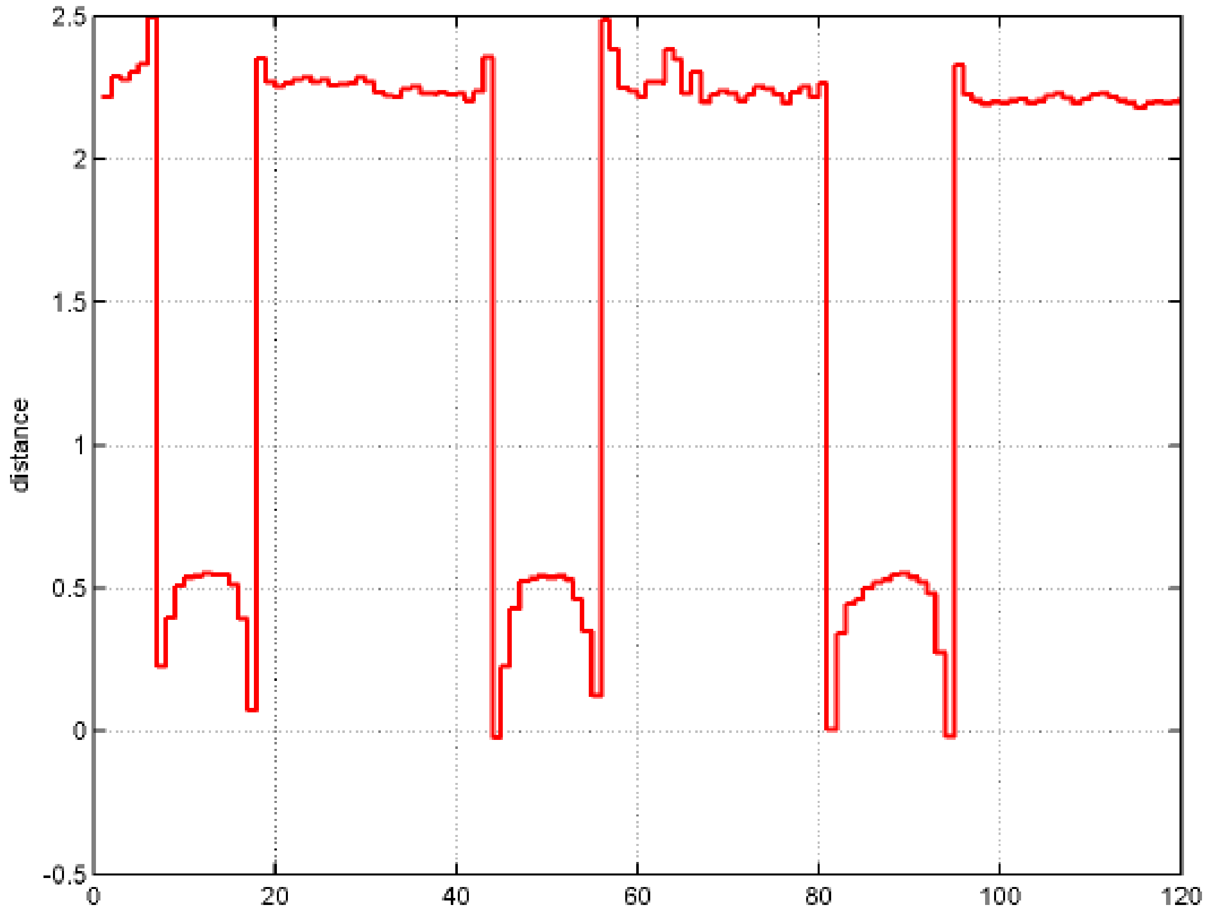

Deconvolved Phase Slice

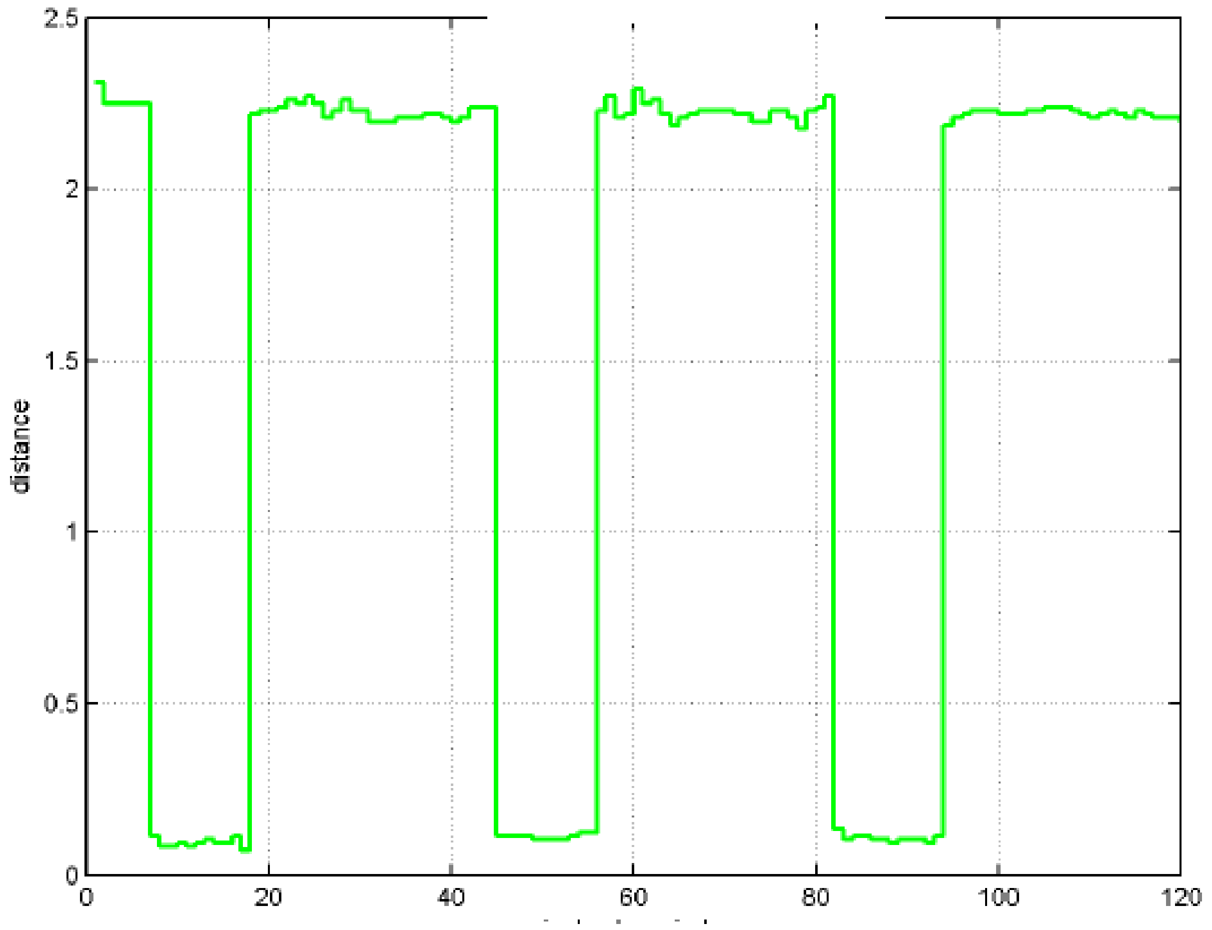

Achuta Kadambi
MIT/UCLA

Transparent 3D

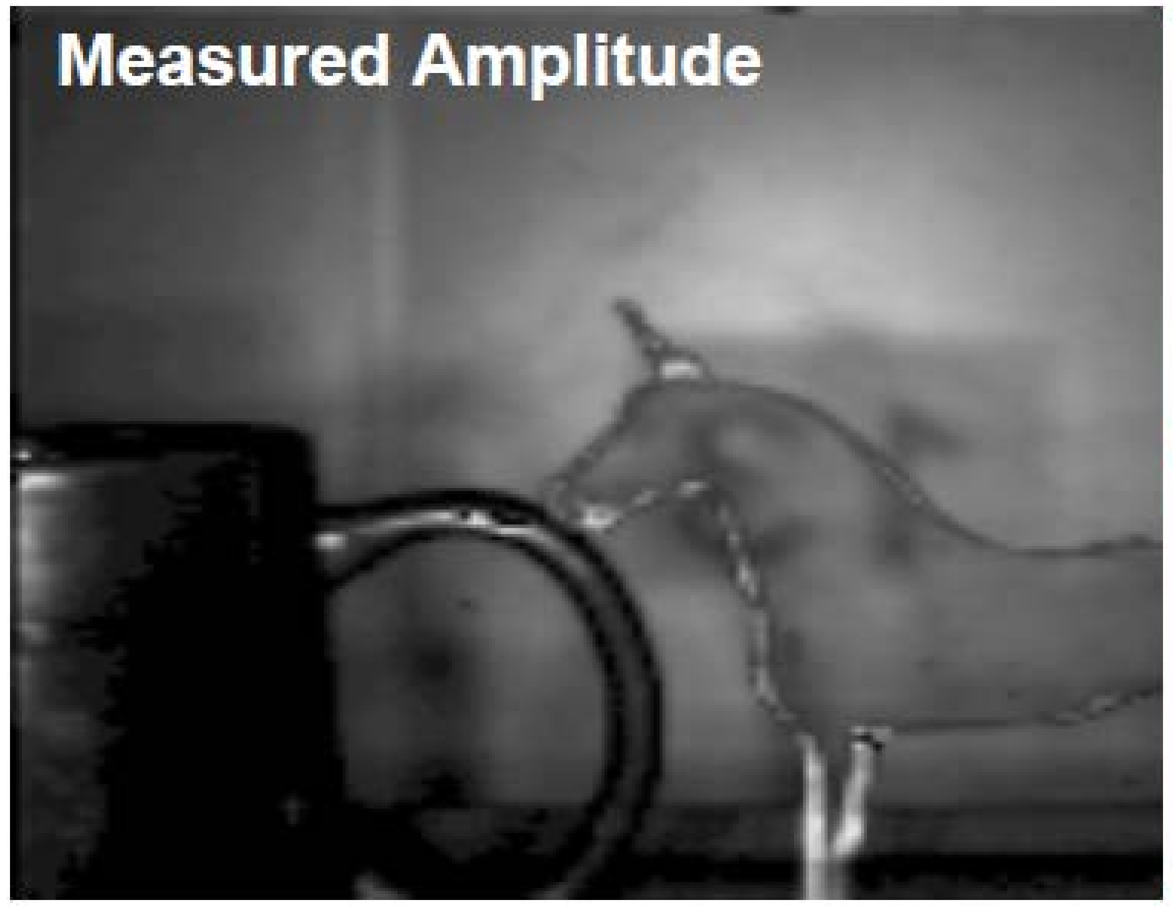

Transparent 3D

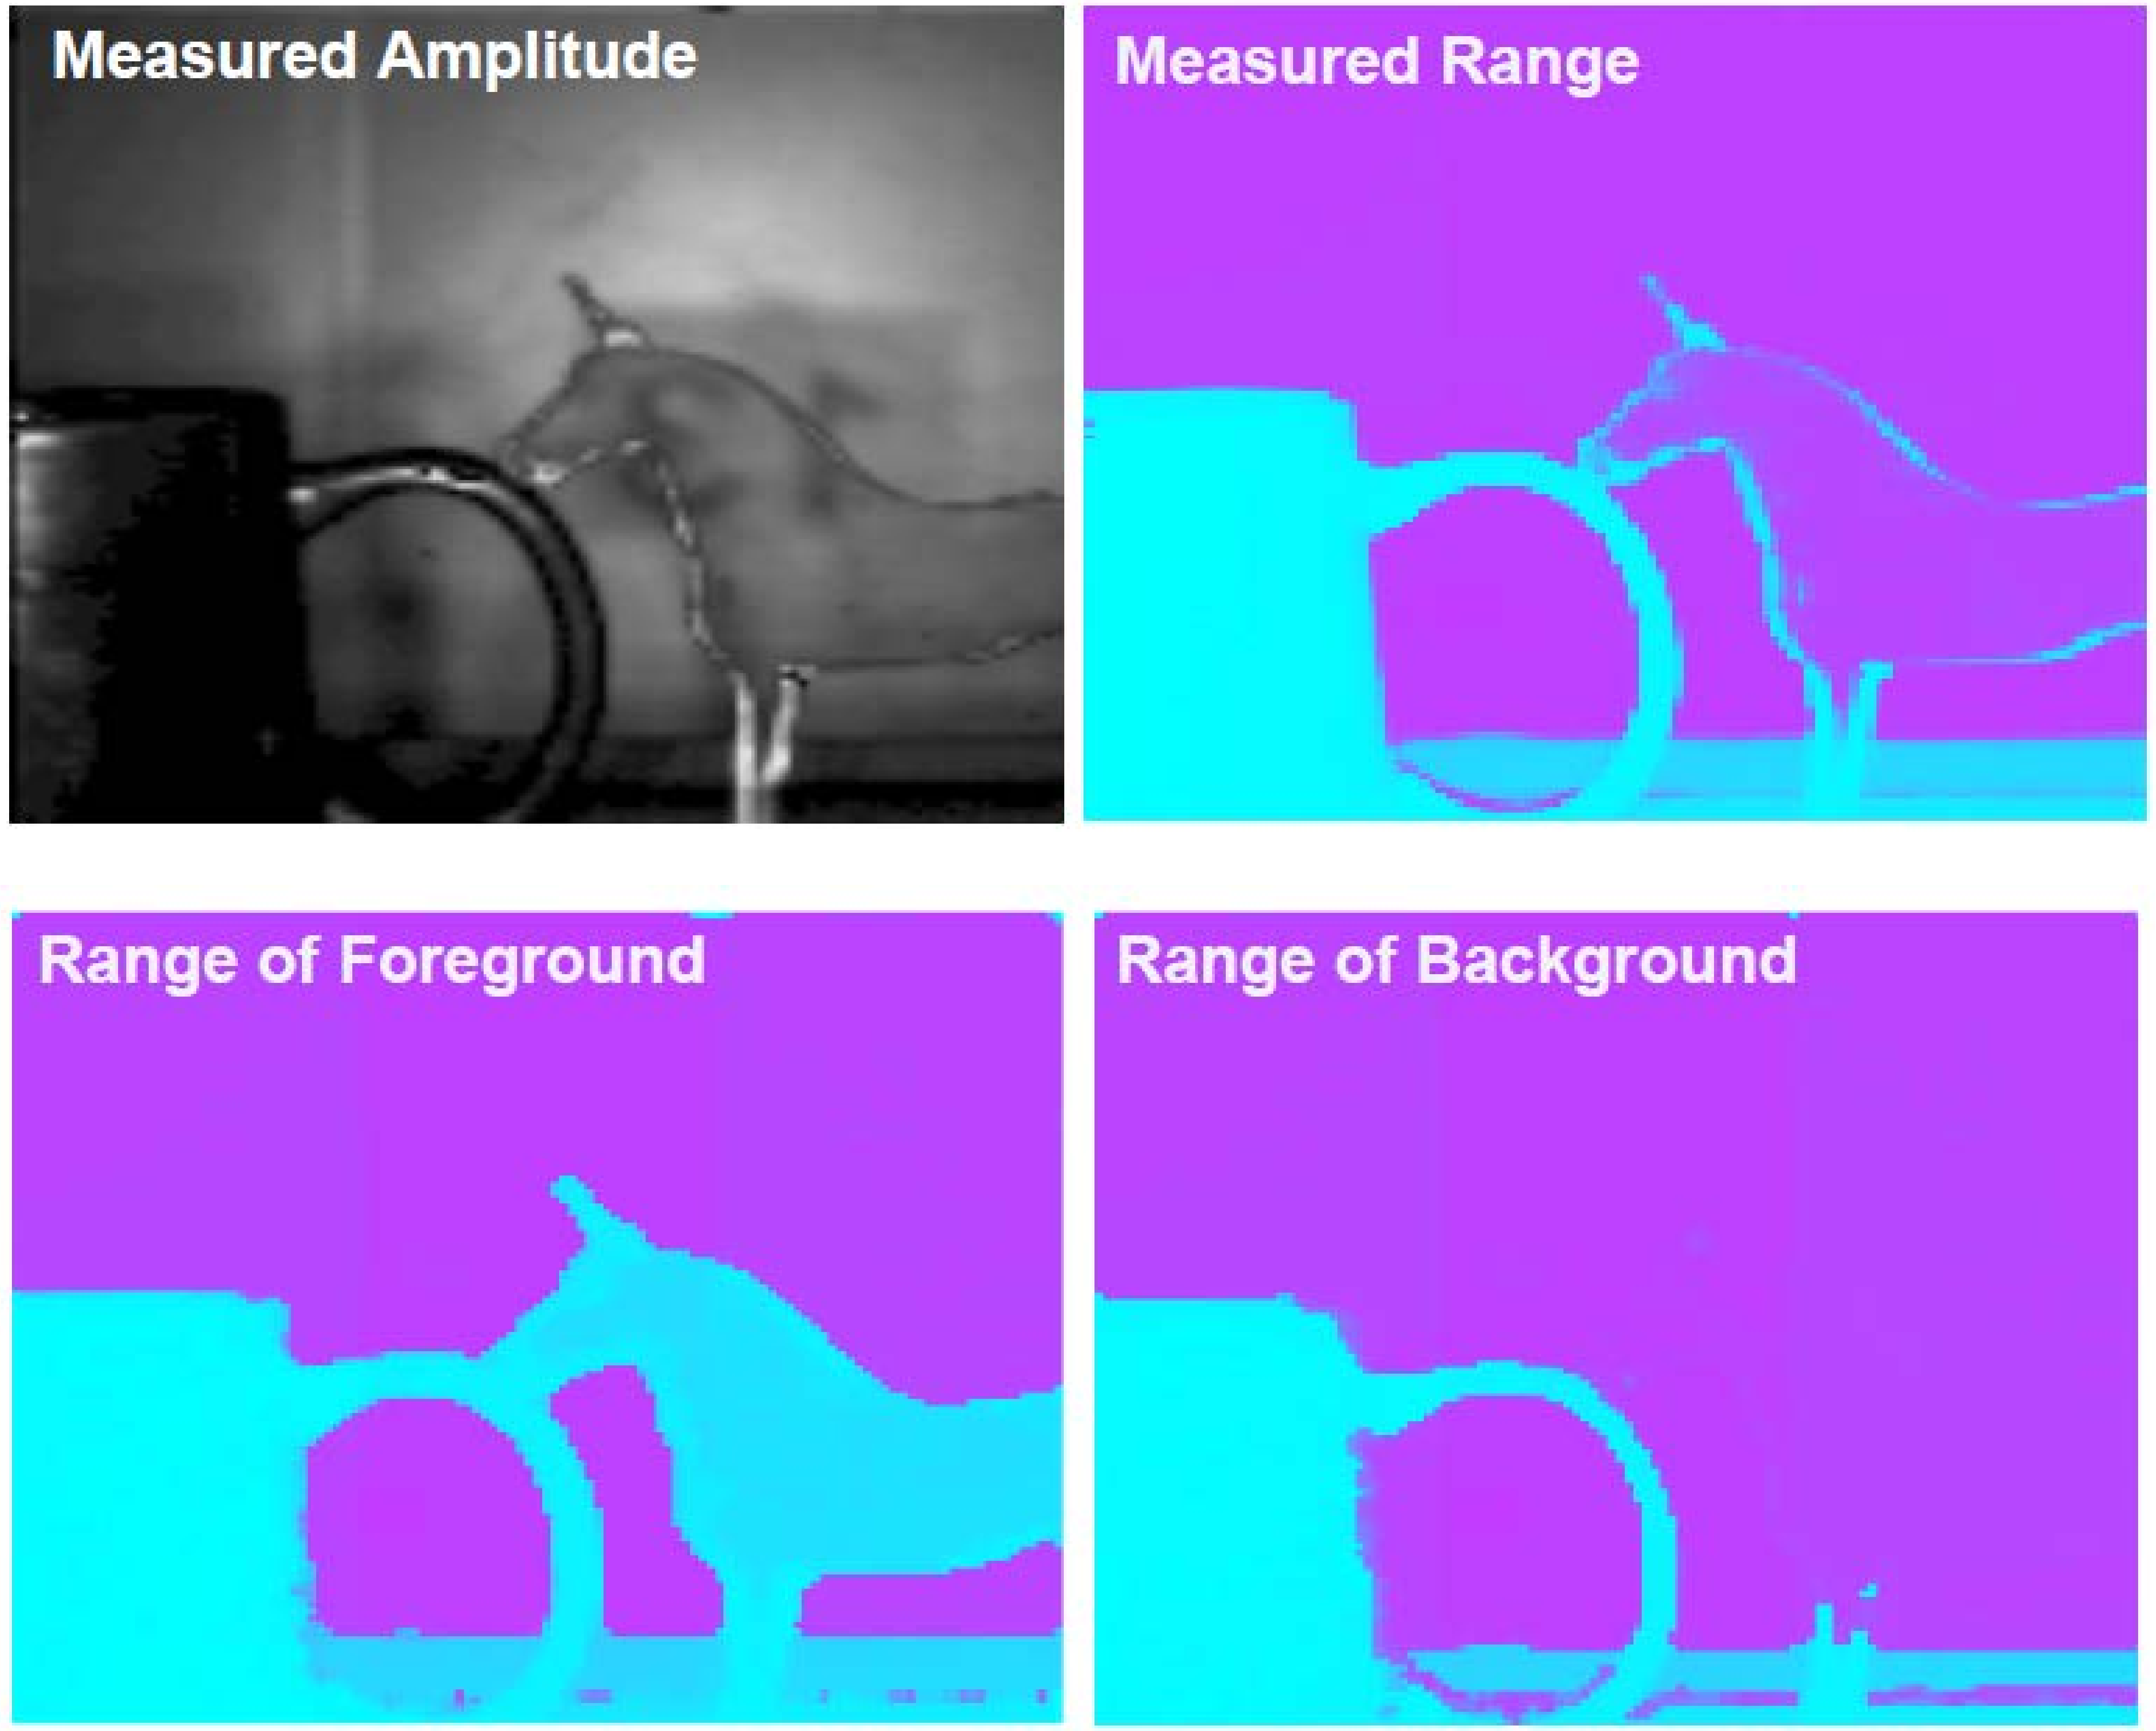

Scientific Applications

Image: a quantum dot on a slide [Bhandari2015]

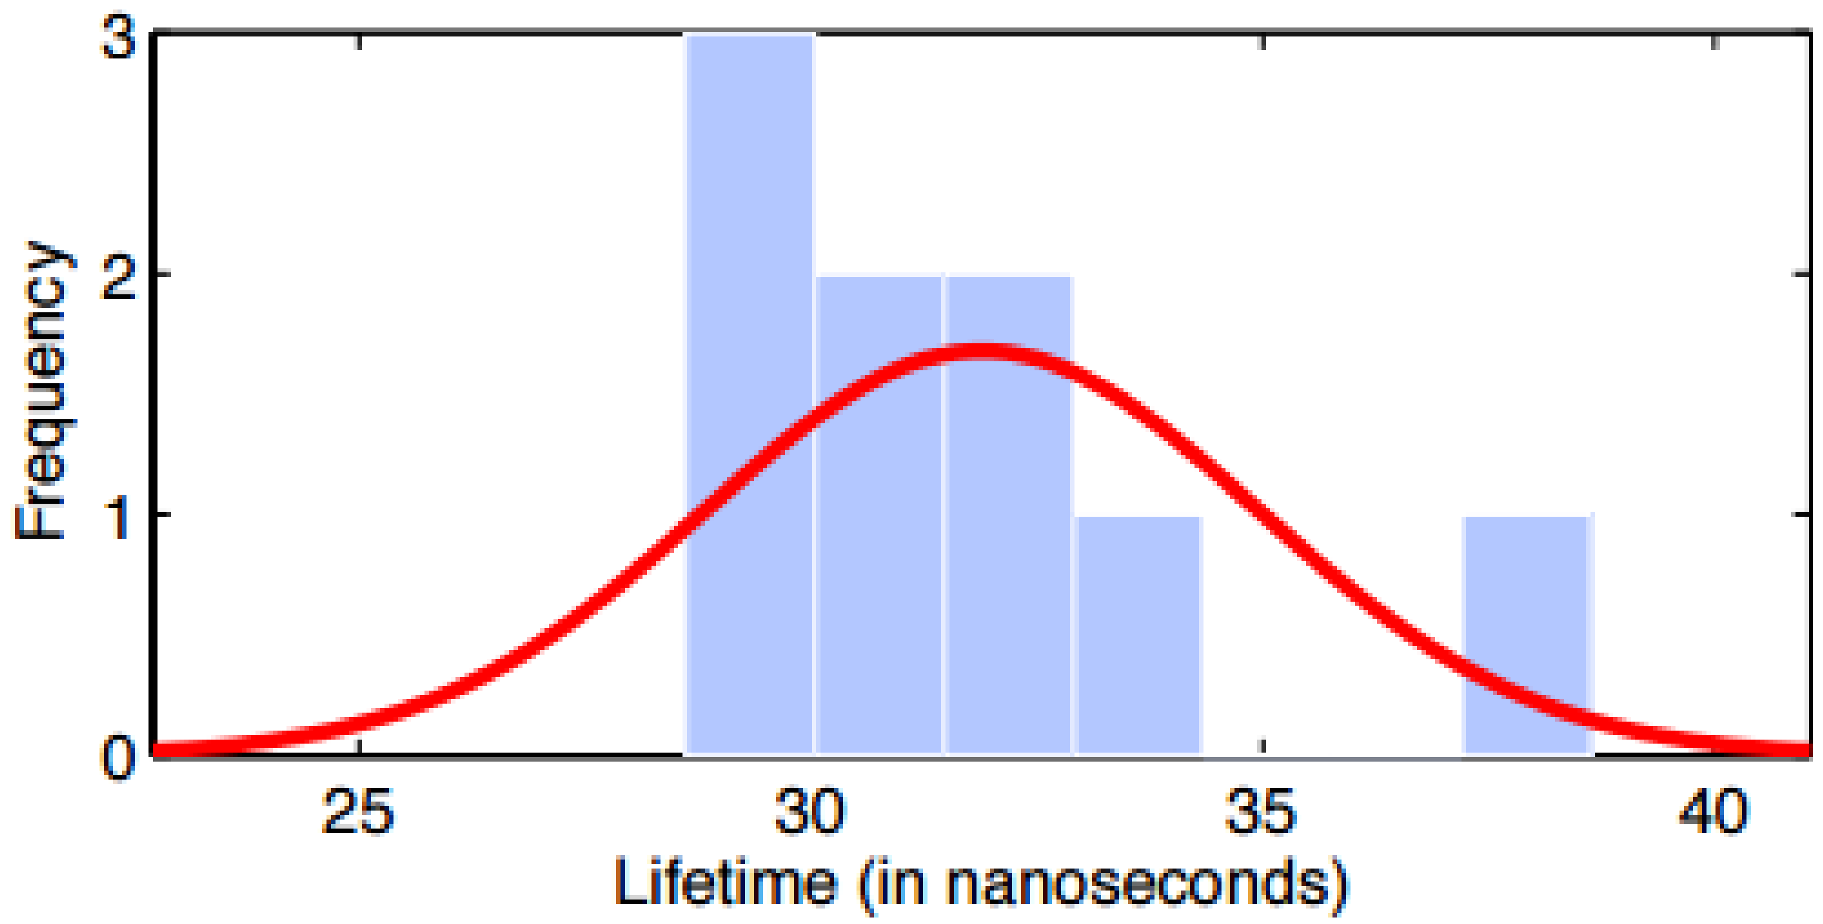

Light emitted from fluorophore is time delayed

Can be reflection-separated from background signal

Lead author of fluor. work is Ayush Bhandari, MIT who has several lines in this area.

Future of LIDAR might be both impulse and CW

Photons Wavefronts

Open Problem: “Curse of Light Speed”

Hard to get quality 3D with time of flight alone

1ns → light travels 1 foot

How does LIDAR work today?

Hard problem .. Light moves fast!

$$d = vt$$

How does LIDAR work today?

LIDAR is hard → light is fast

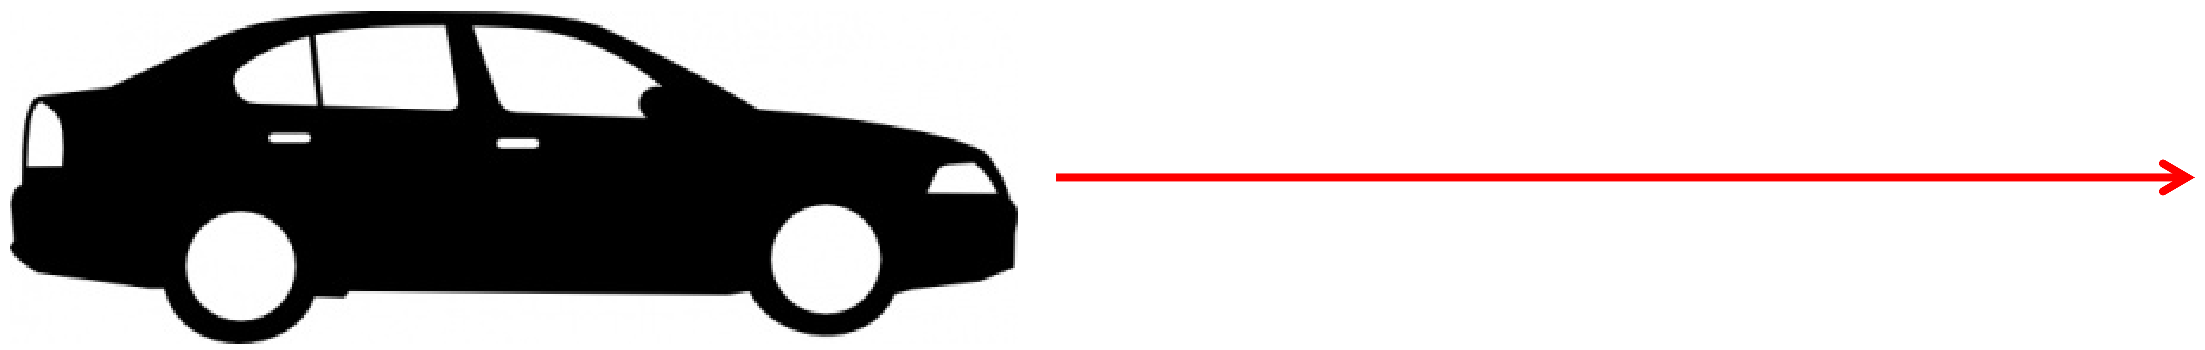

$$d = vt$$

5mm tag

How does LIDAR work today?

LIDAR is hard → light is fast

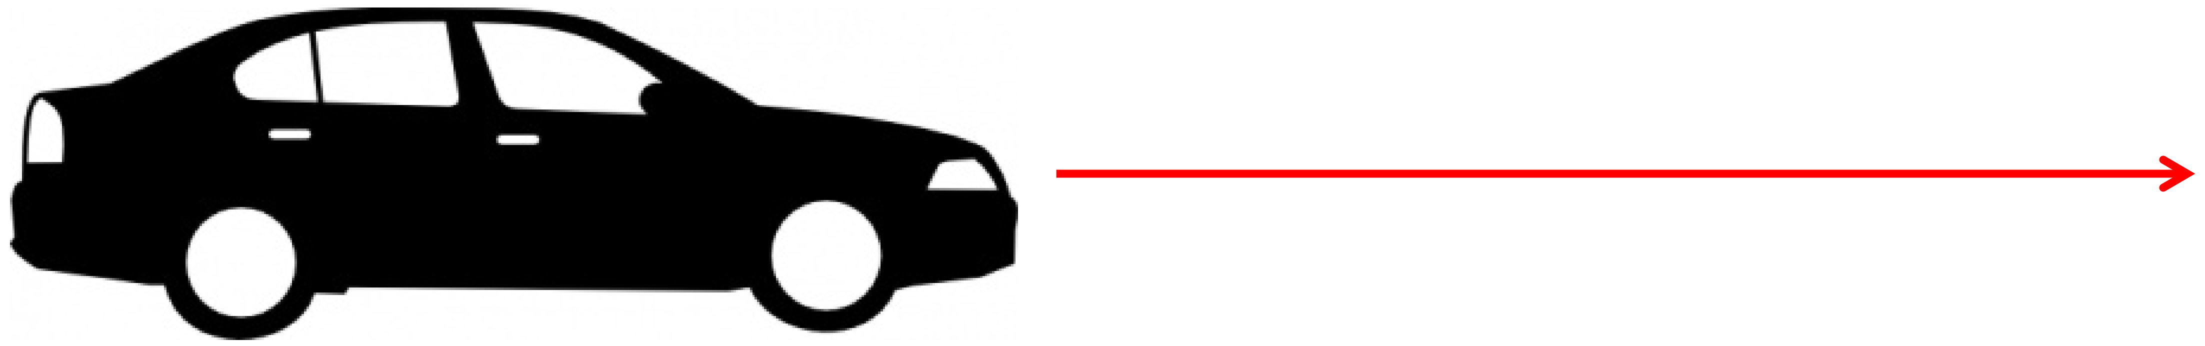

$$d = vt$$

5mm tag

MIT Proposal: Heterodyne LIDAR at 2 micrometers

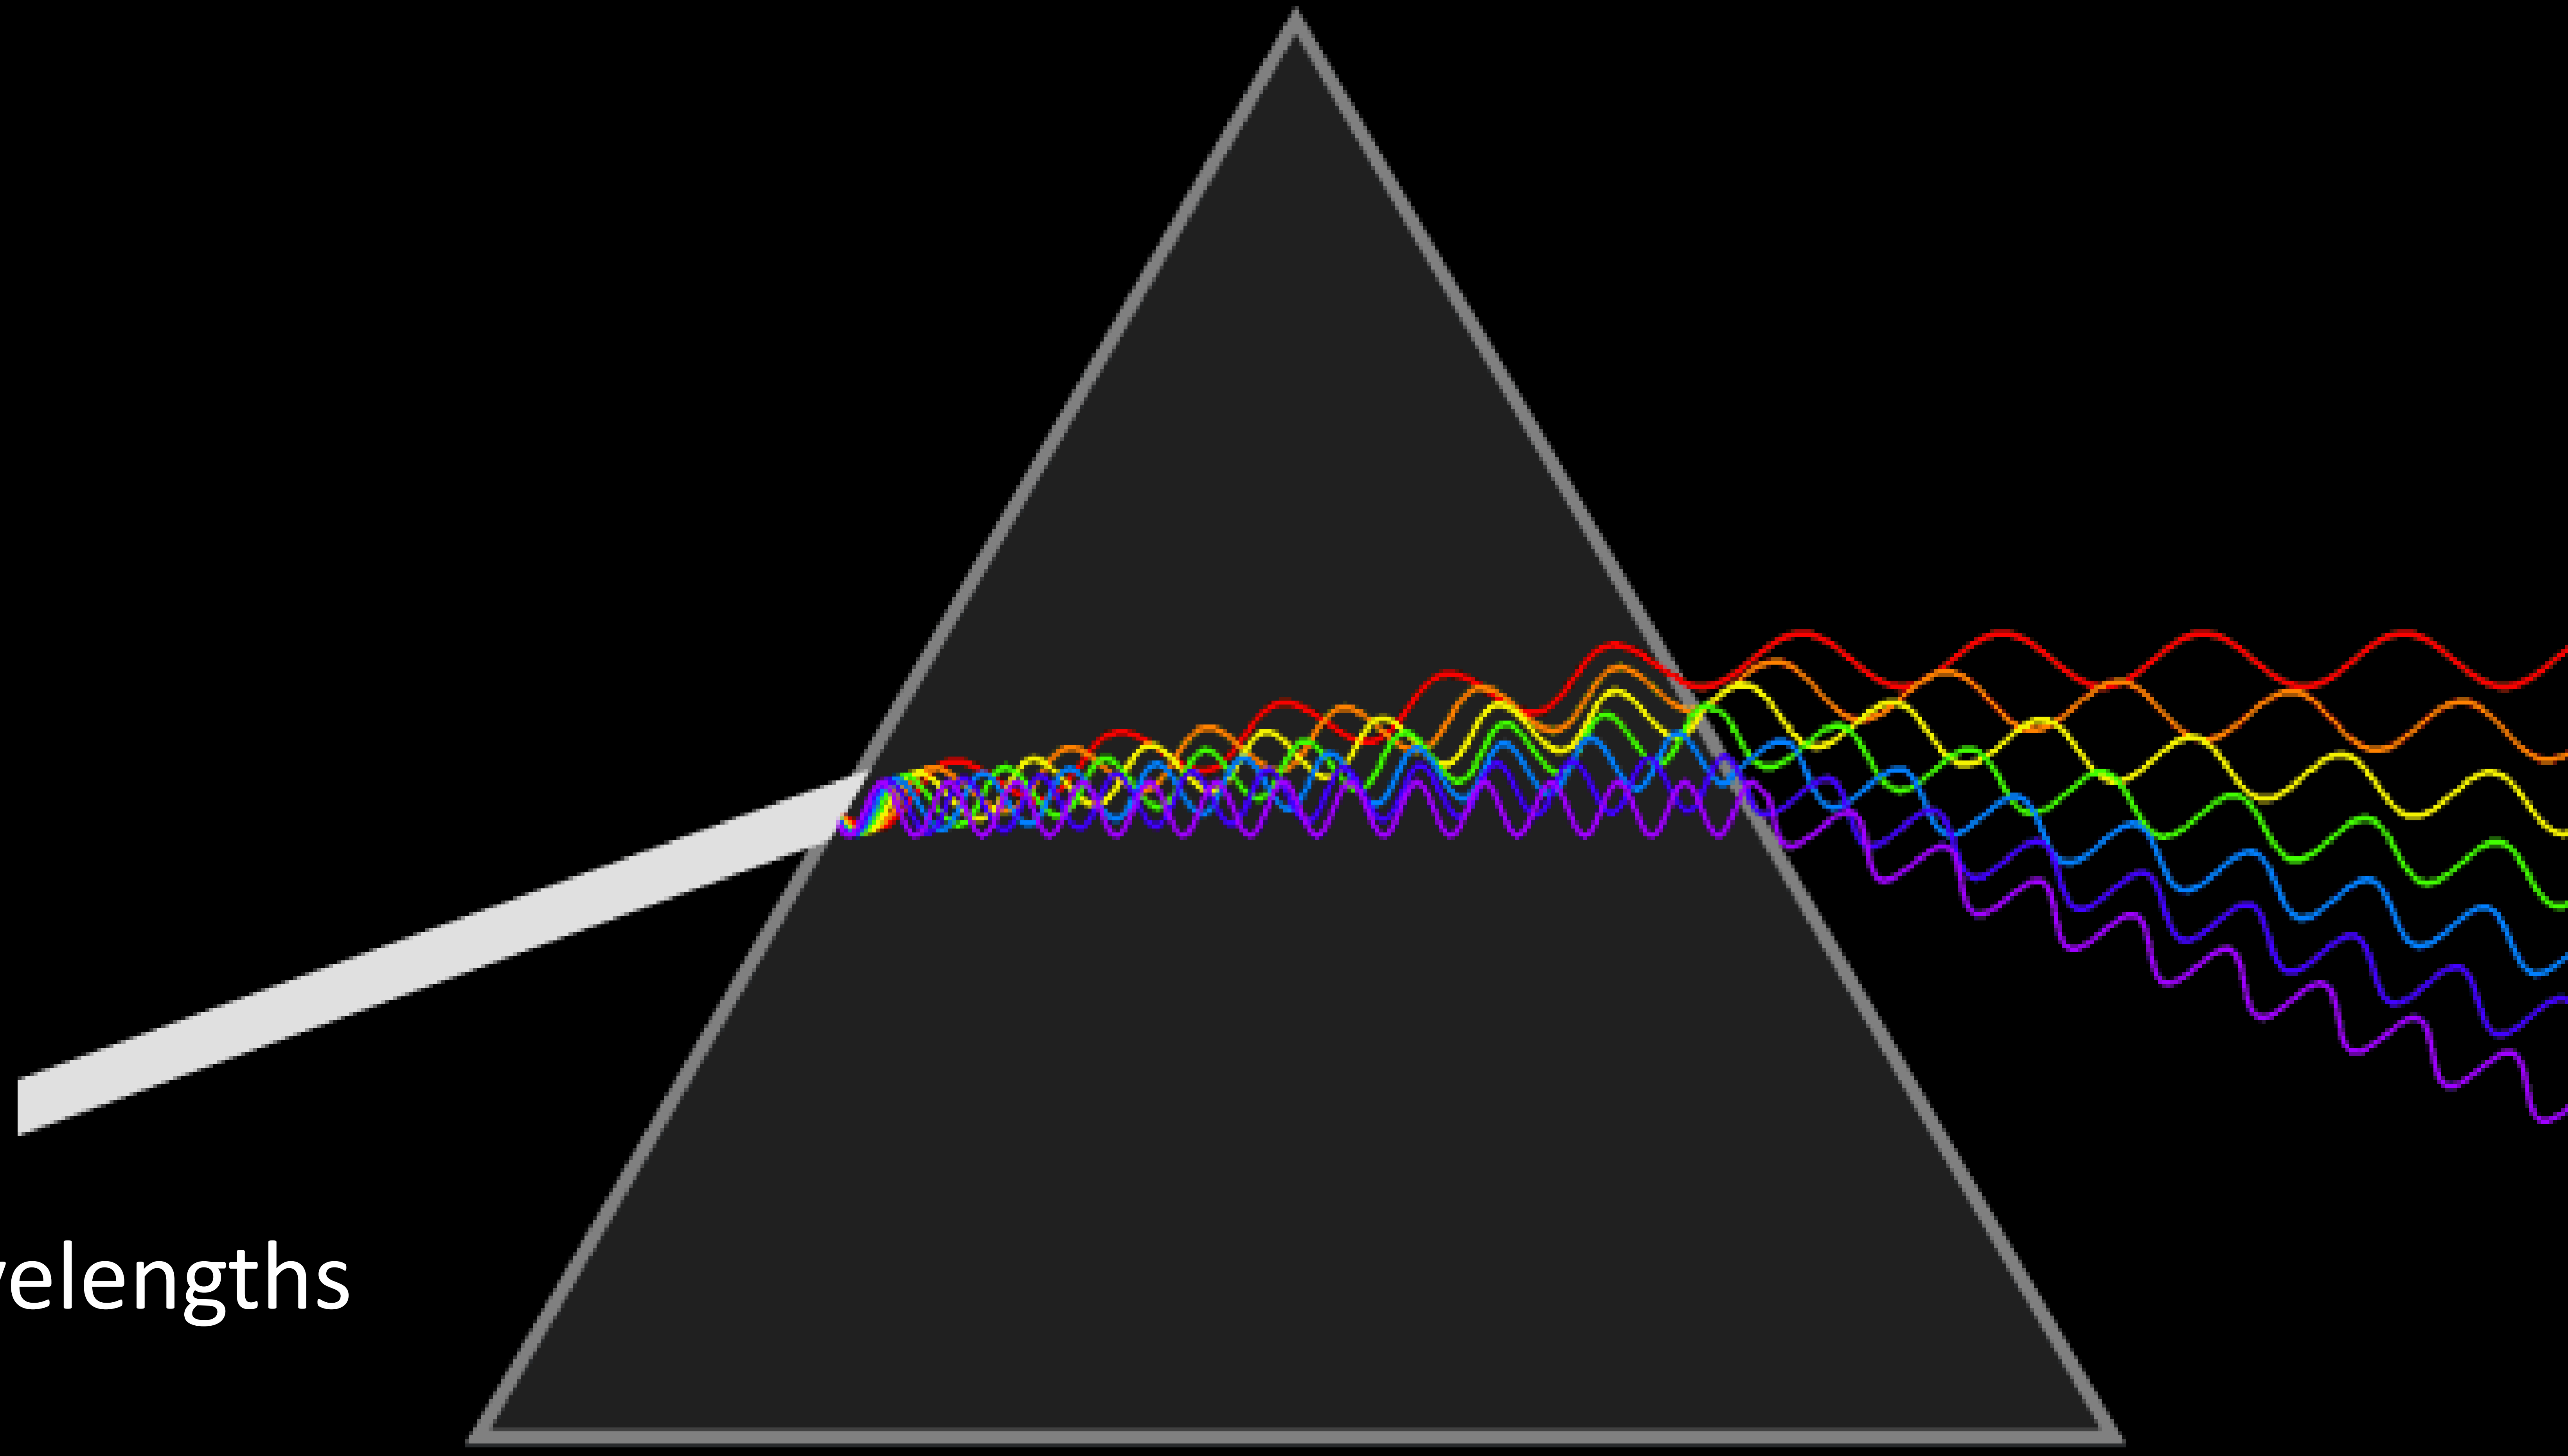

Exploits light at different wavelengths

Photo credit: Wikipedia

New MIT LIDAR

Photons

$$d = vt$$

Photons and Wavefronts

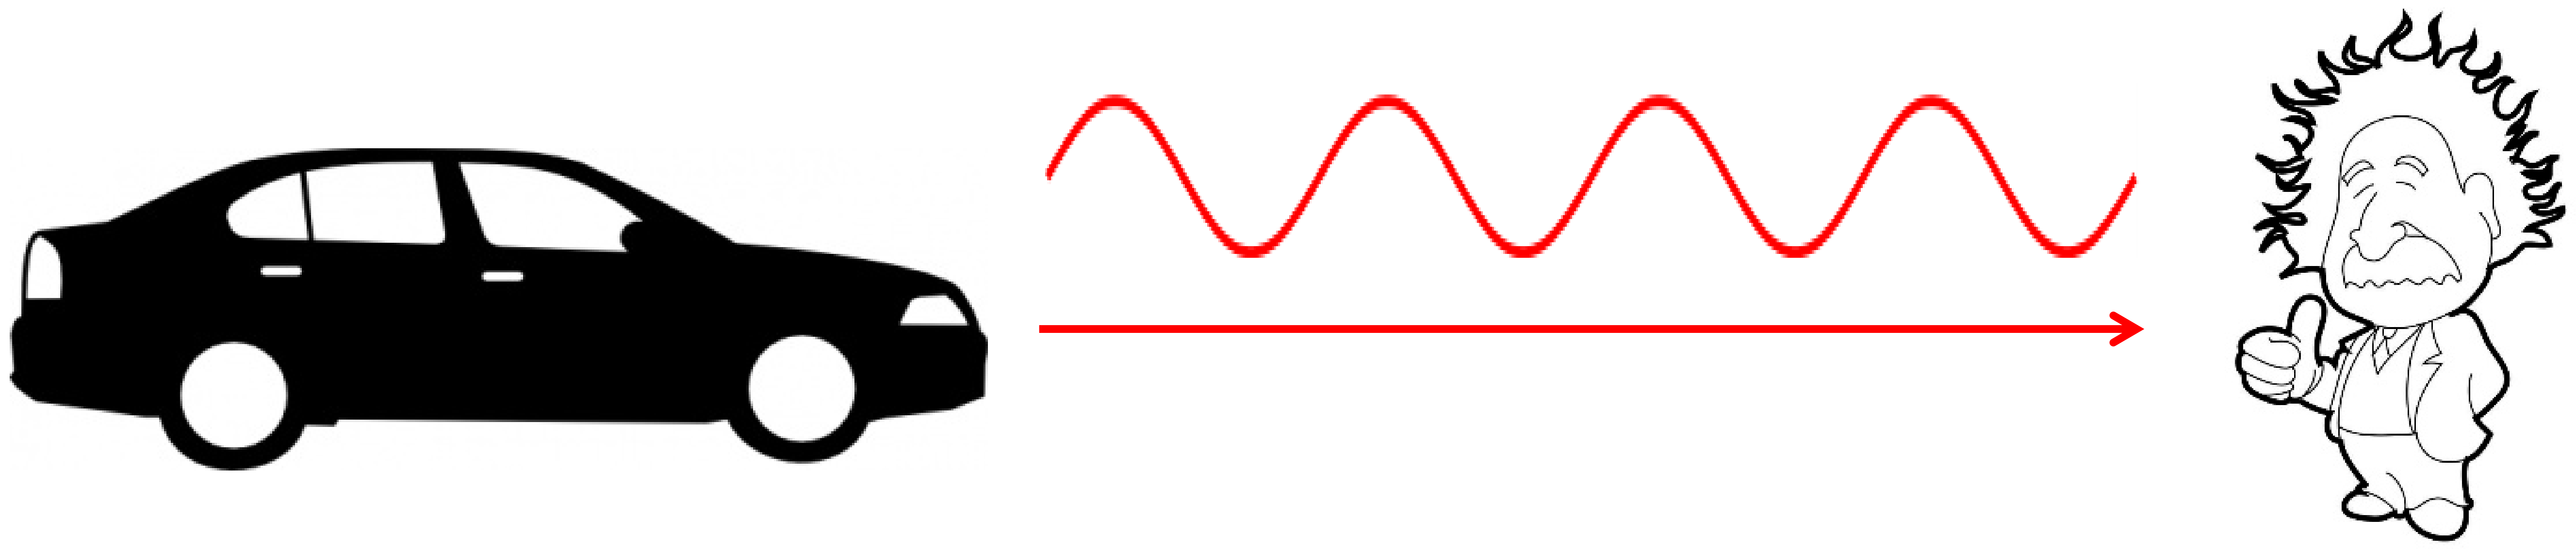

New MIT LIDAR

Photons

Photons and Wavefronts

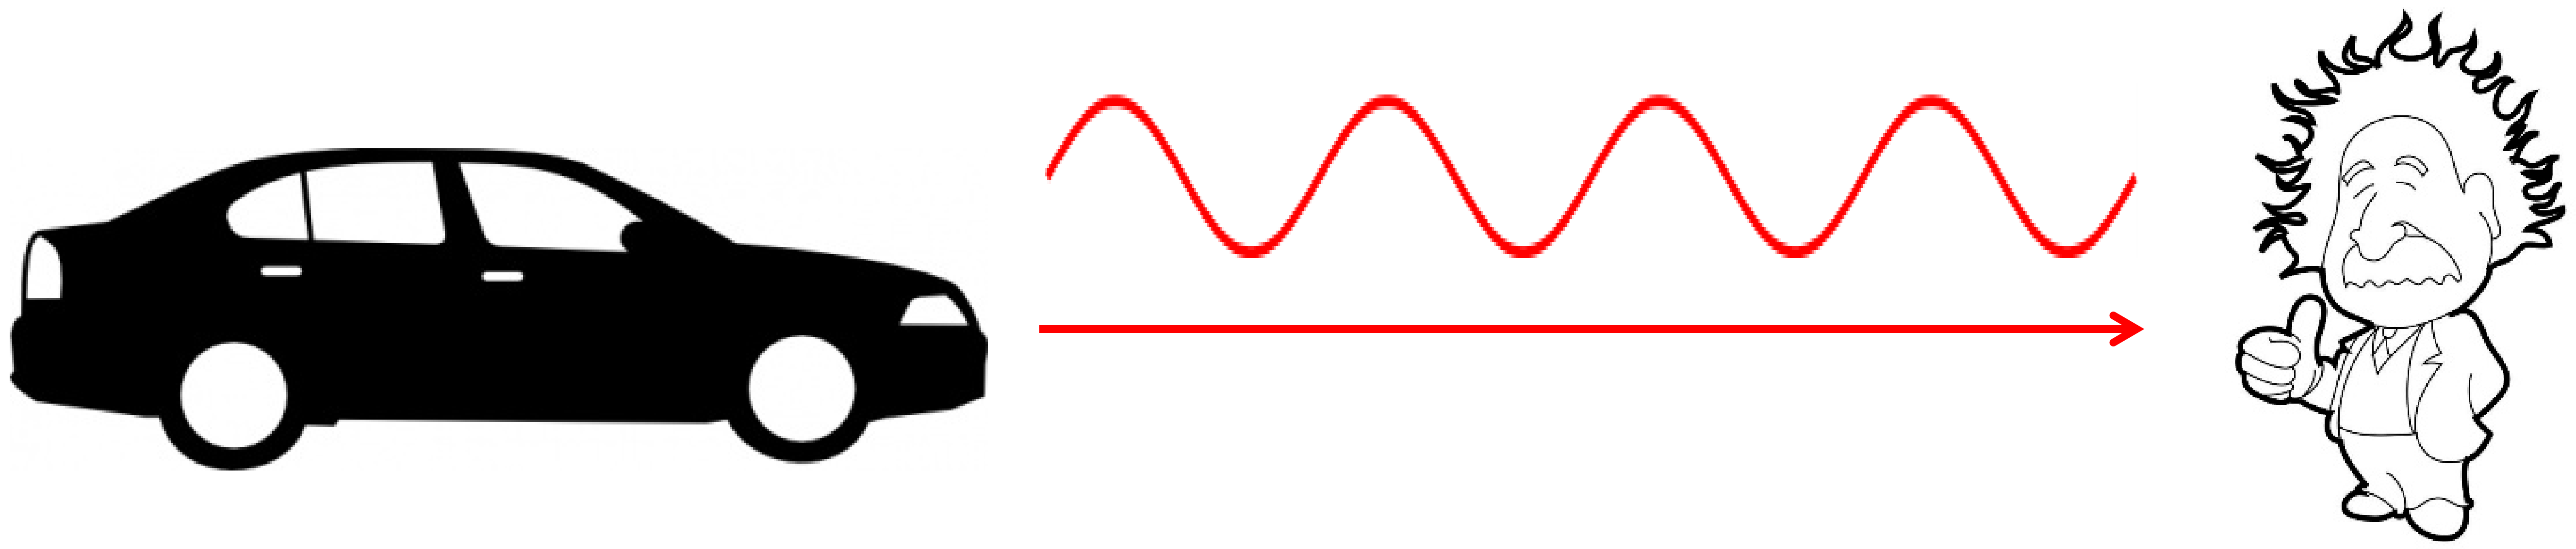

2 micron Depth Resolution at Long Range
(100X improvement)

Photons x Wavefronts

Photons

$$E_{\text{reflected}}\left(t, z = z_{\text{imager}}\right) = \underbrace{E_0\, e^{j(\omega t - \omega \frac{z_{\text{ToF}}}{c})}}_{\text{optical term}} \times \left(e^{jn_0}\, \underbrace{e^{j\eta \cos(\Omega t - \Omega \frac{z_{\text{ToF}}}{c})}}_{\text{modulation term}} + 1\right)$$

Wavefront

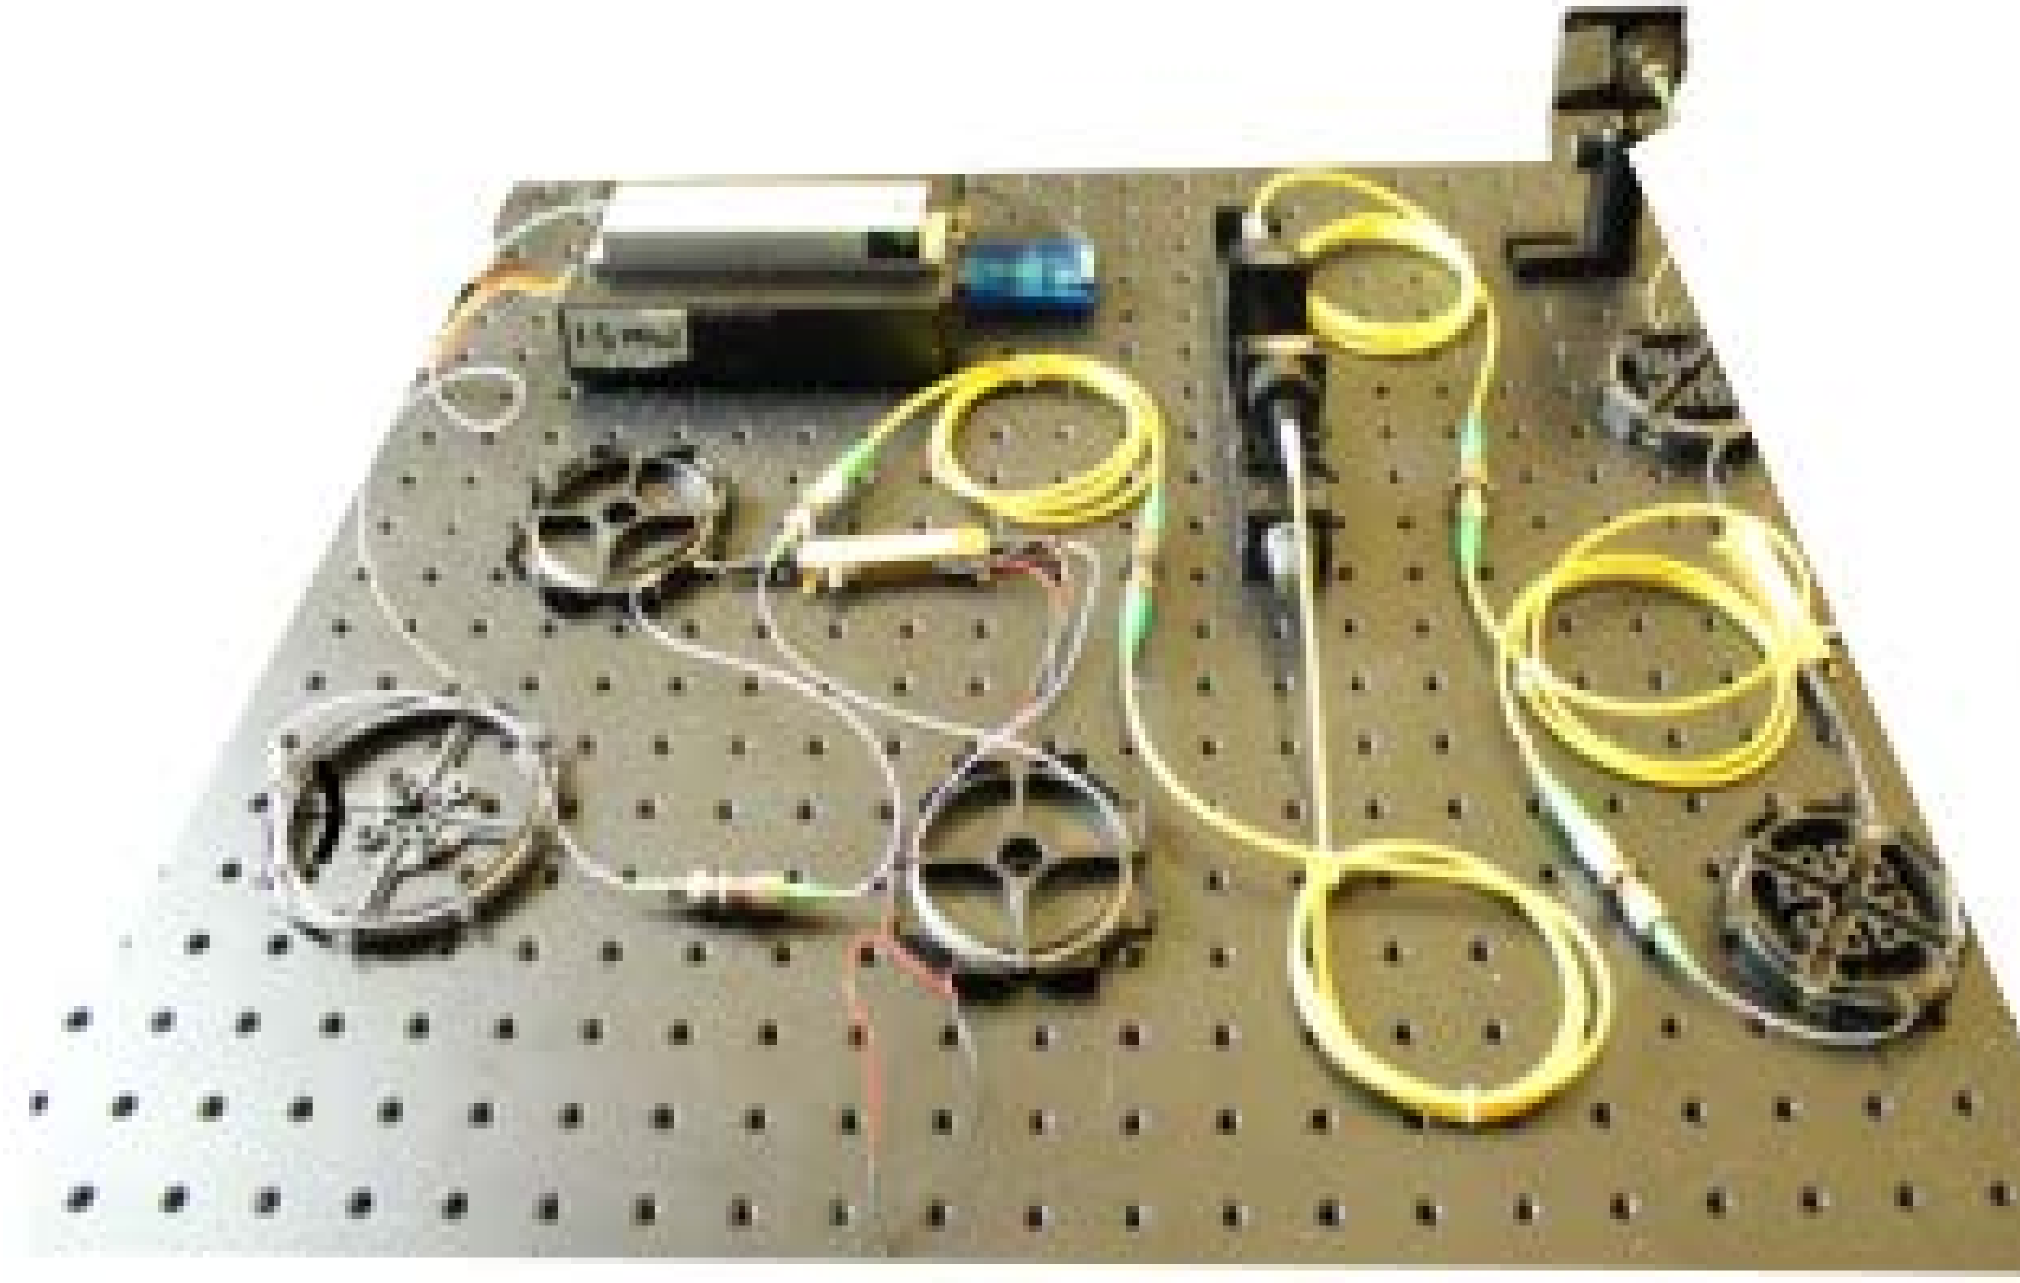

Portable, GHz LIDAR

Towards 500 meter range

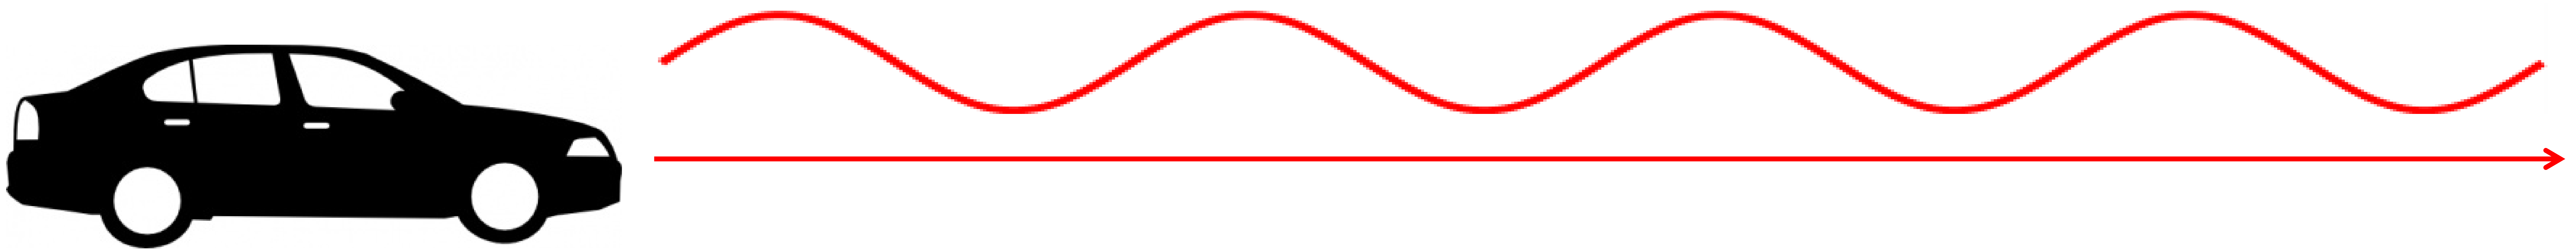

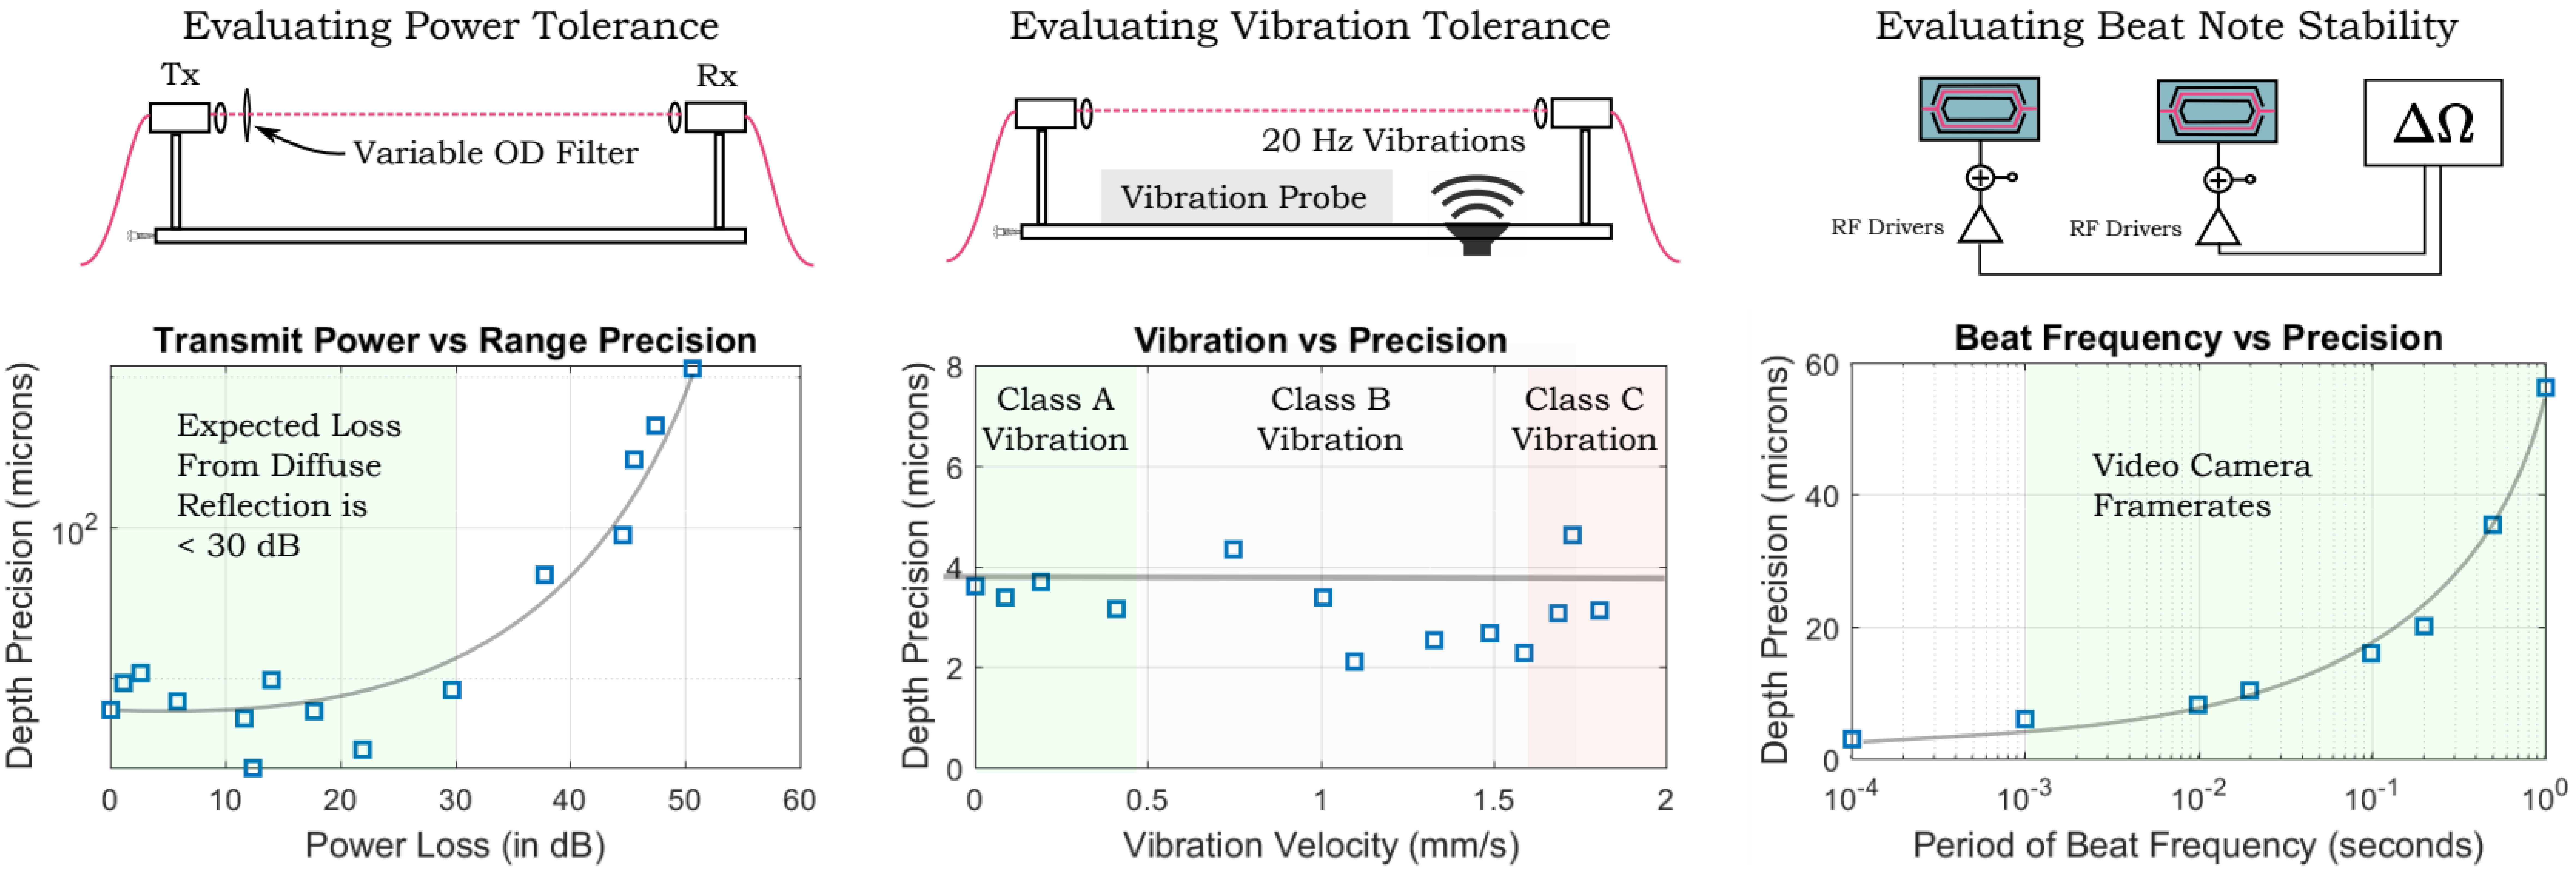

The LIDAR Arms Race

Quanergy
(tech unicorn 1.5 billion valuation)

Strobe (acq. By GM)

Waymo/Google

The LIDAR Arms Race

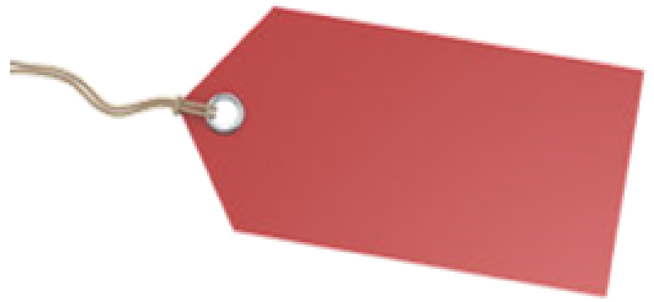

Existing technologies tag light at mm precision

Strobe (acq. By GM)

Quanergy
(tech unicorn 1.5 billion valuation)

Waymo/Google

Next Level → Light Tagging at Micrometers

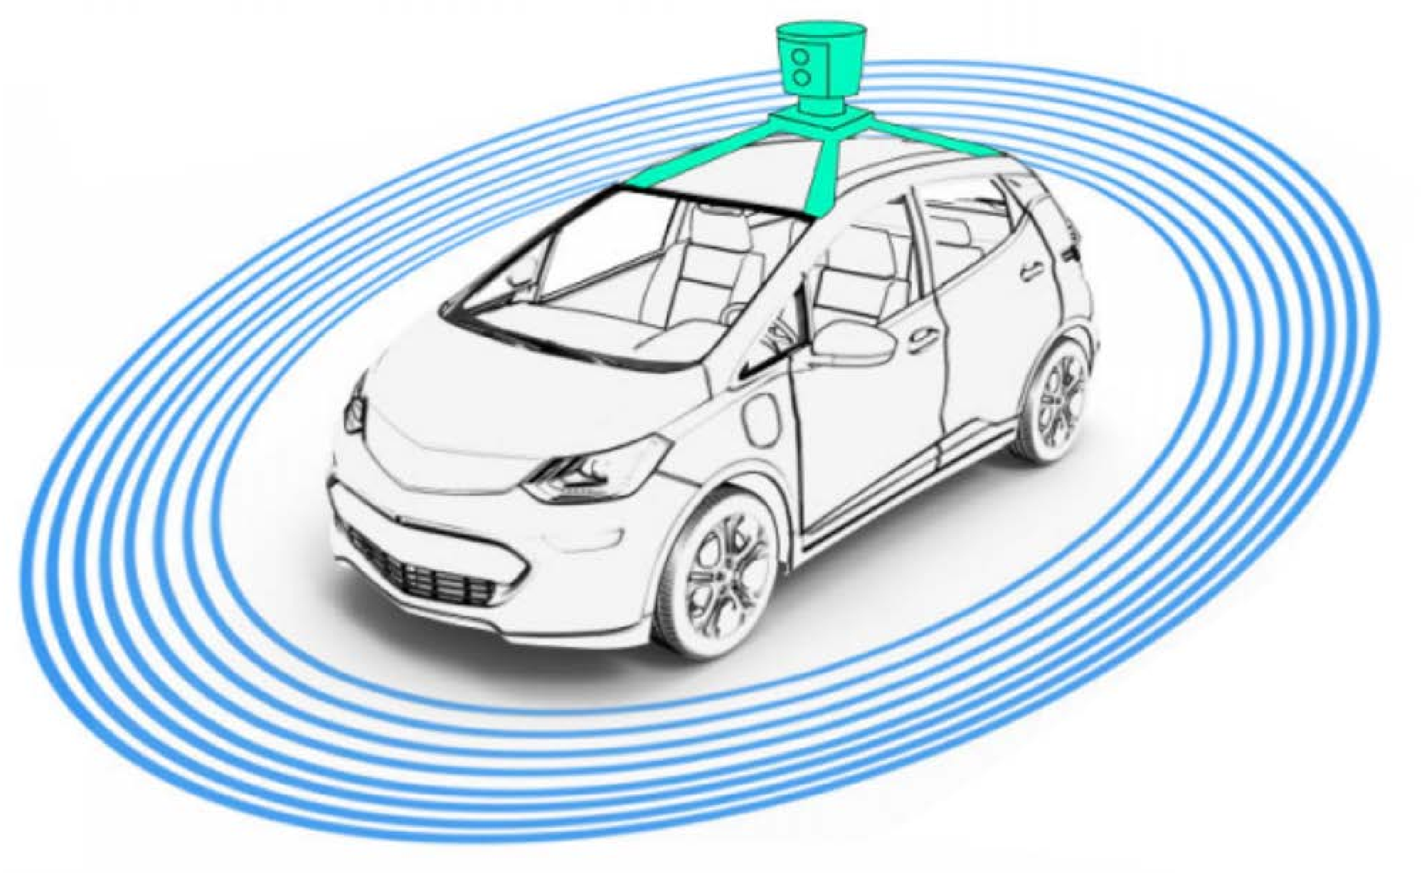

Vehicles
(Driving Through Fog)

Airplanes
Small 3D Features

Imaging thru Tissue

Time of Flight @ CVPR This Year

Time-resolved Light Transport Decomposition for Thermal Photometric Stereo

Kenichiro Tanaka Nobuhiro Ikeya Tsuyoshi Takatani
Hiroyuki Kubo Takuya Funatomi Yasuhiro Mukaigawa
Nara Institute of Science and Technology (NAIST), Japan
{ktanaka, ikeya.nobuhiro.ii6, takatani.tsuyoshi.to2, hkubo, funatomi, mukaigawa}@is.naist.jp

Paper exploits time-resolved statistics of temperature to understand materials

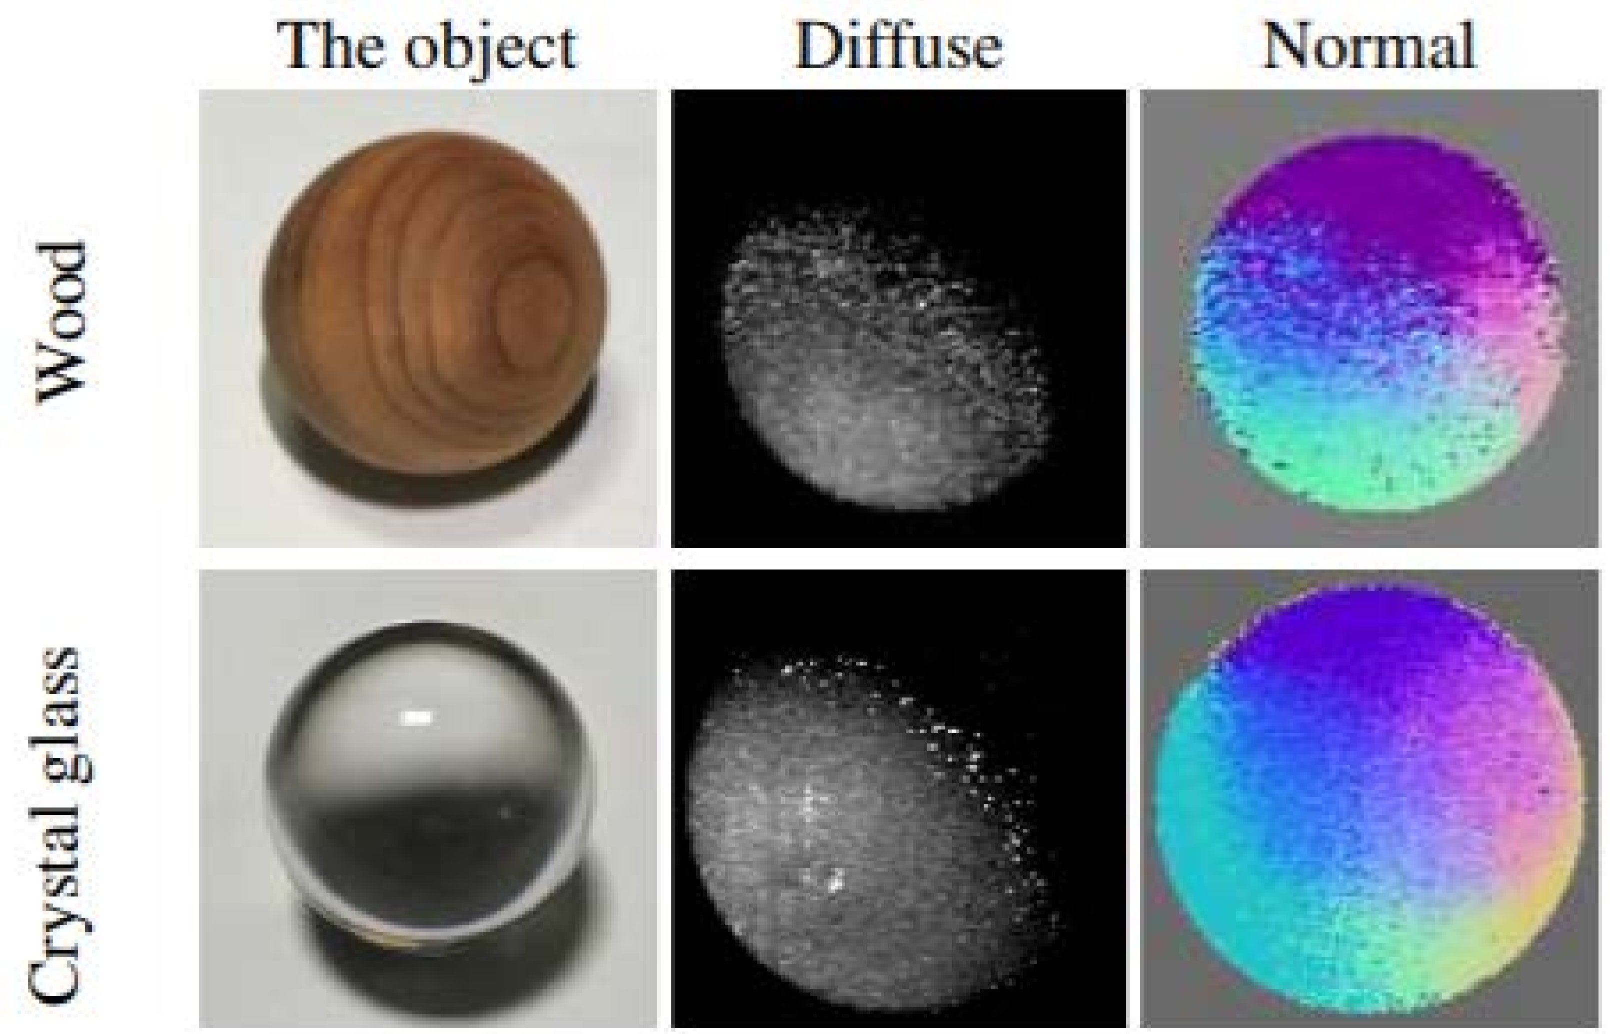


Autonomous
cars Driving
Through Fog?


As if it were a
sunny day

Course Schedule

13:30 – 13:50	Course Introduction *Ramesh Raskar (MIT)*
13:50 – 15:00	Existing Sensors and Their Limits *Guy Satat (MIT), Achuta Kadambi (UCLA)*
15:00 – 15:10	Break
15:10 – 15:50	Emerging 3D Sensors *Achuta Kadambi (UCLA)*
15:50 – 16:30	Imaging in Bad Weather *Guy Satat (MIT)*
16:30 – 16:40	Break
16:40 – 17:20	Deep Learning-based Computational Imaging *Jan Kautz (NVIDIA)*
17:20 – 17:30	Conclusion and Open Problems

Course Schedule

http://imagingav.media.mit.edu

13:30 – 13:50	Course Introduction *Ramesh Raskar (MIT)*
13:50 – 15:00	Existing Sensors and Their Limits *Guy Satat (MIT), Achuta Kadambi (UCLA)*
15:00 – 15:10	Break
15:10 – 15:50	Emerging 3D Sensors *Achuta Kadambi (UCLA)*
15:50 – 16:30	Imaging in Bad Weather *Guy Satat (MIT)*
16:30 – 16:40	Break
16:40 – 17:20	Deep Learning-based Computational Imaging *Jan Kautz (NVIDIA)*
17:20 – 17:30	Conclusion and Open Problems

End with some joke

How does a Self-Driving Car See Today?

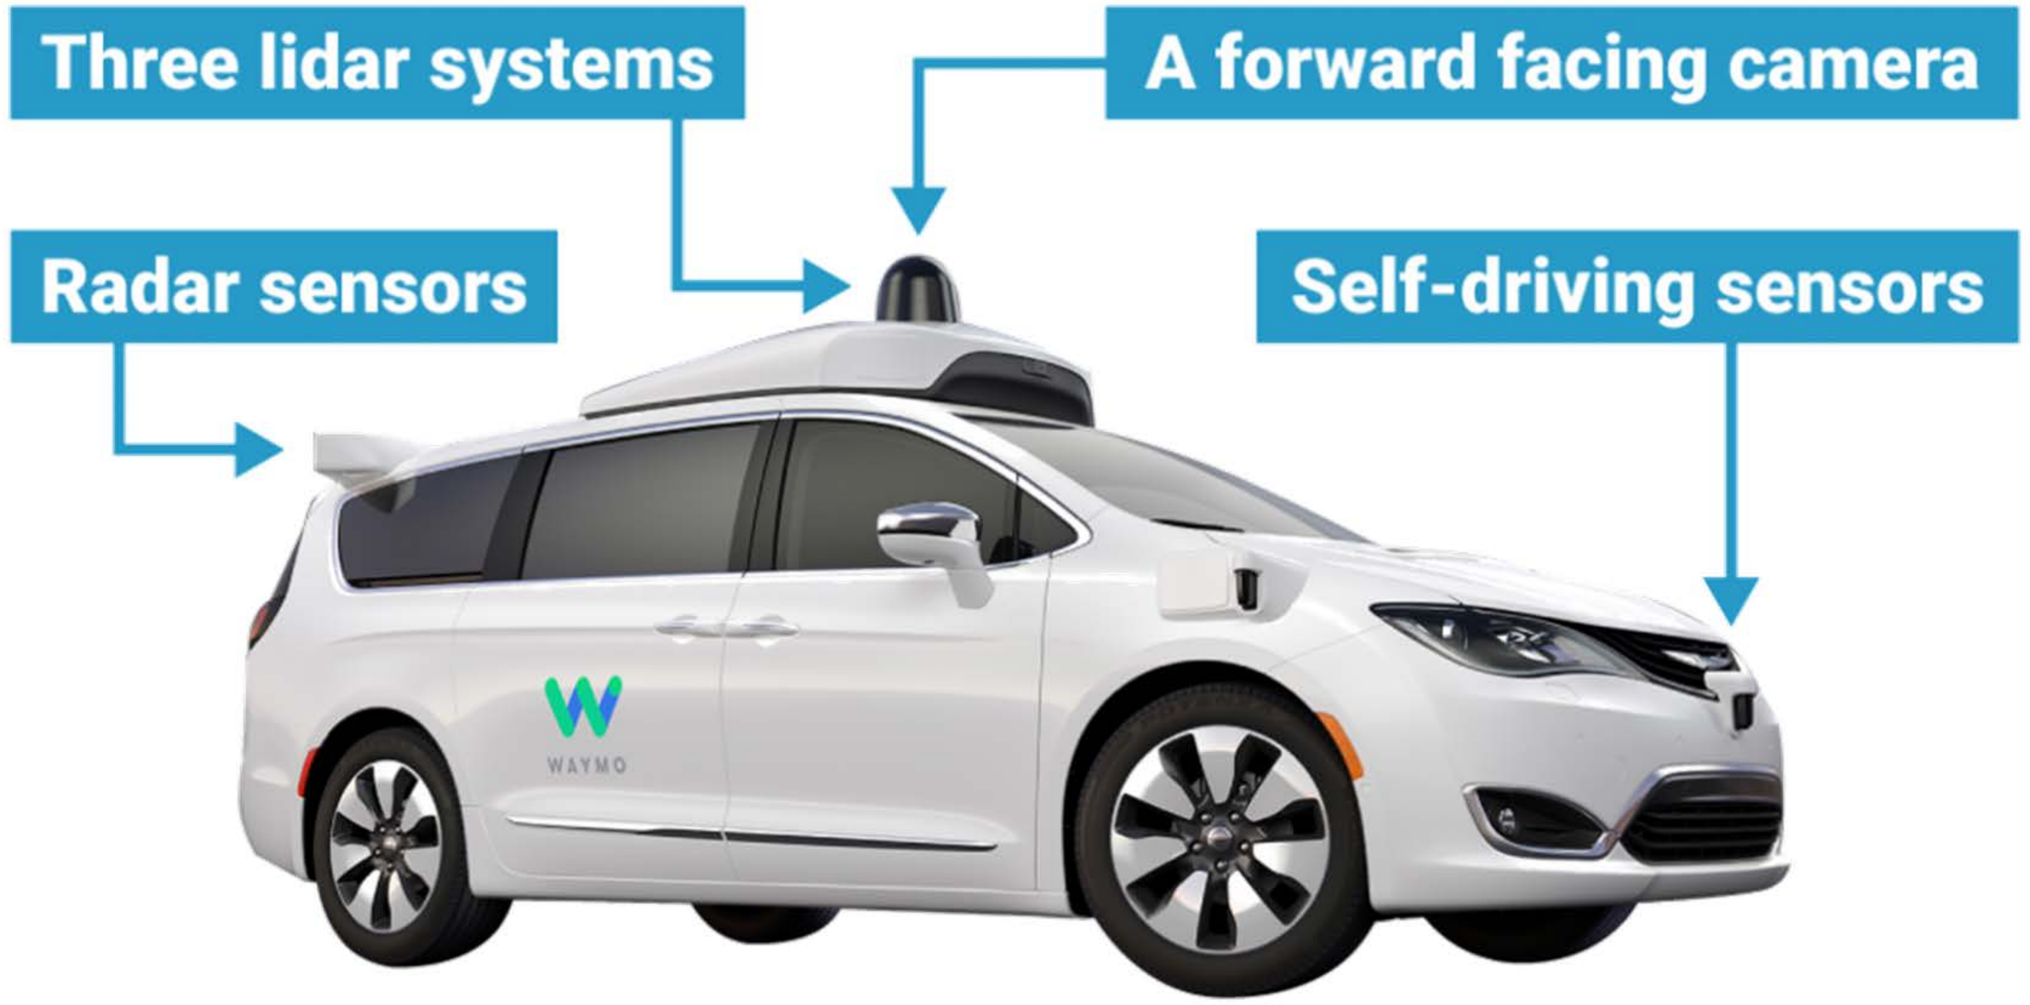

Waymo (Google)


16 Colors
?
? ? ? ?
? ? ? ?
? ? ? ?

Beyond Vision – Camera 2025

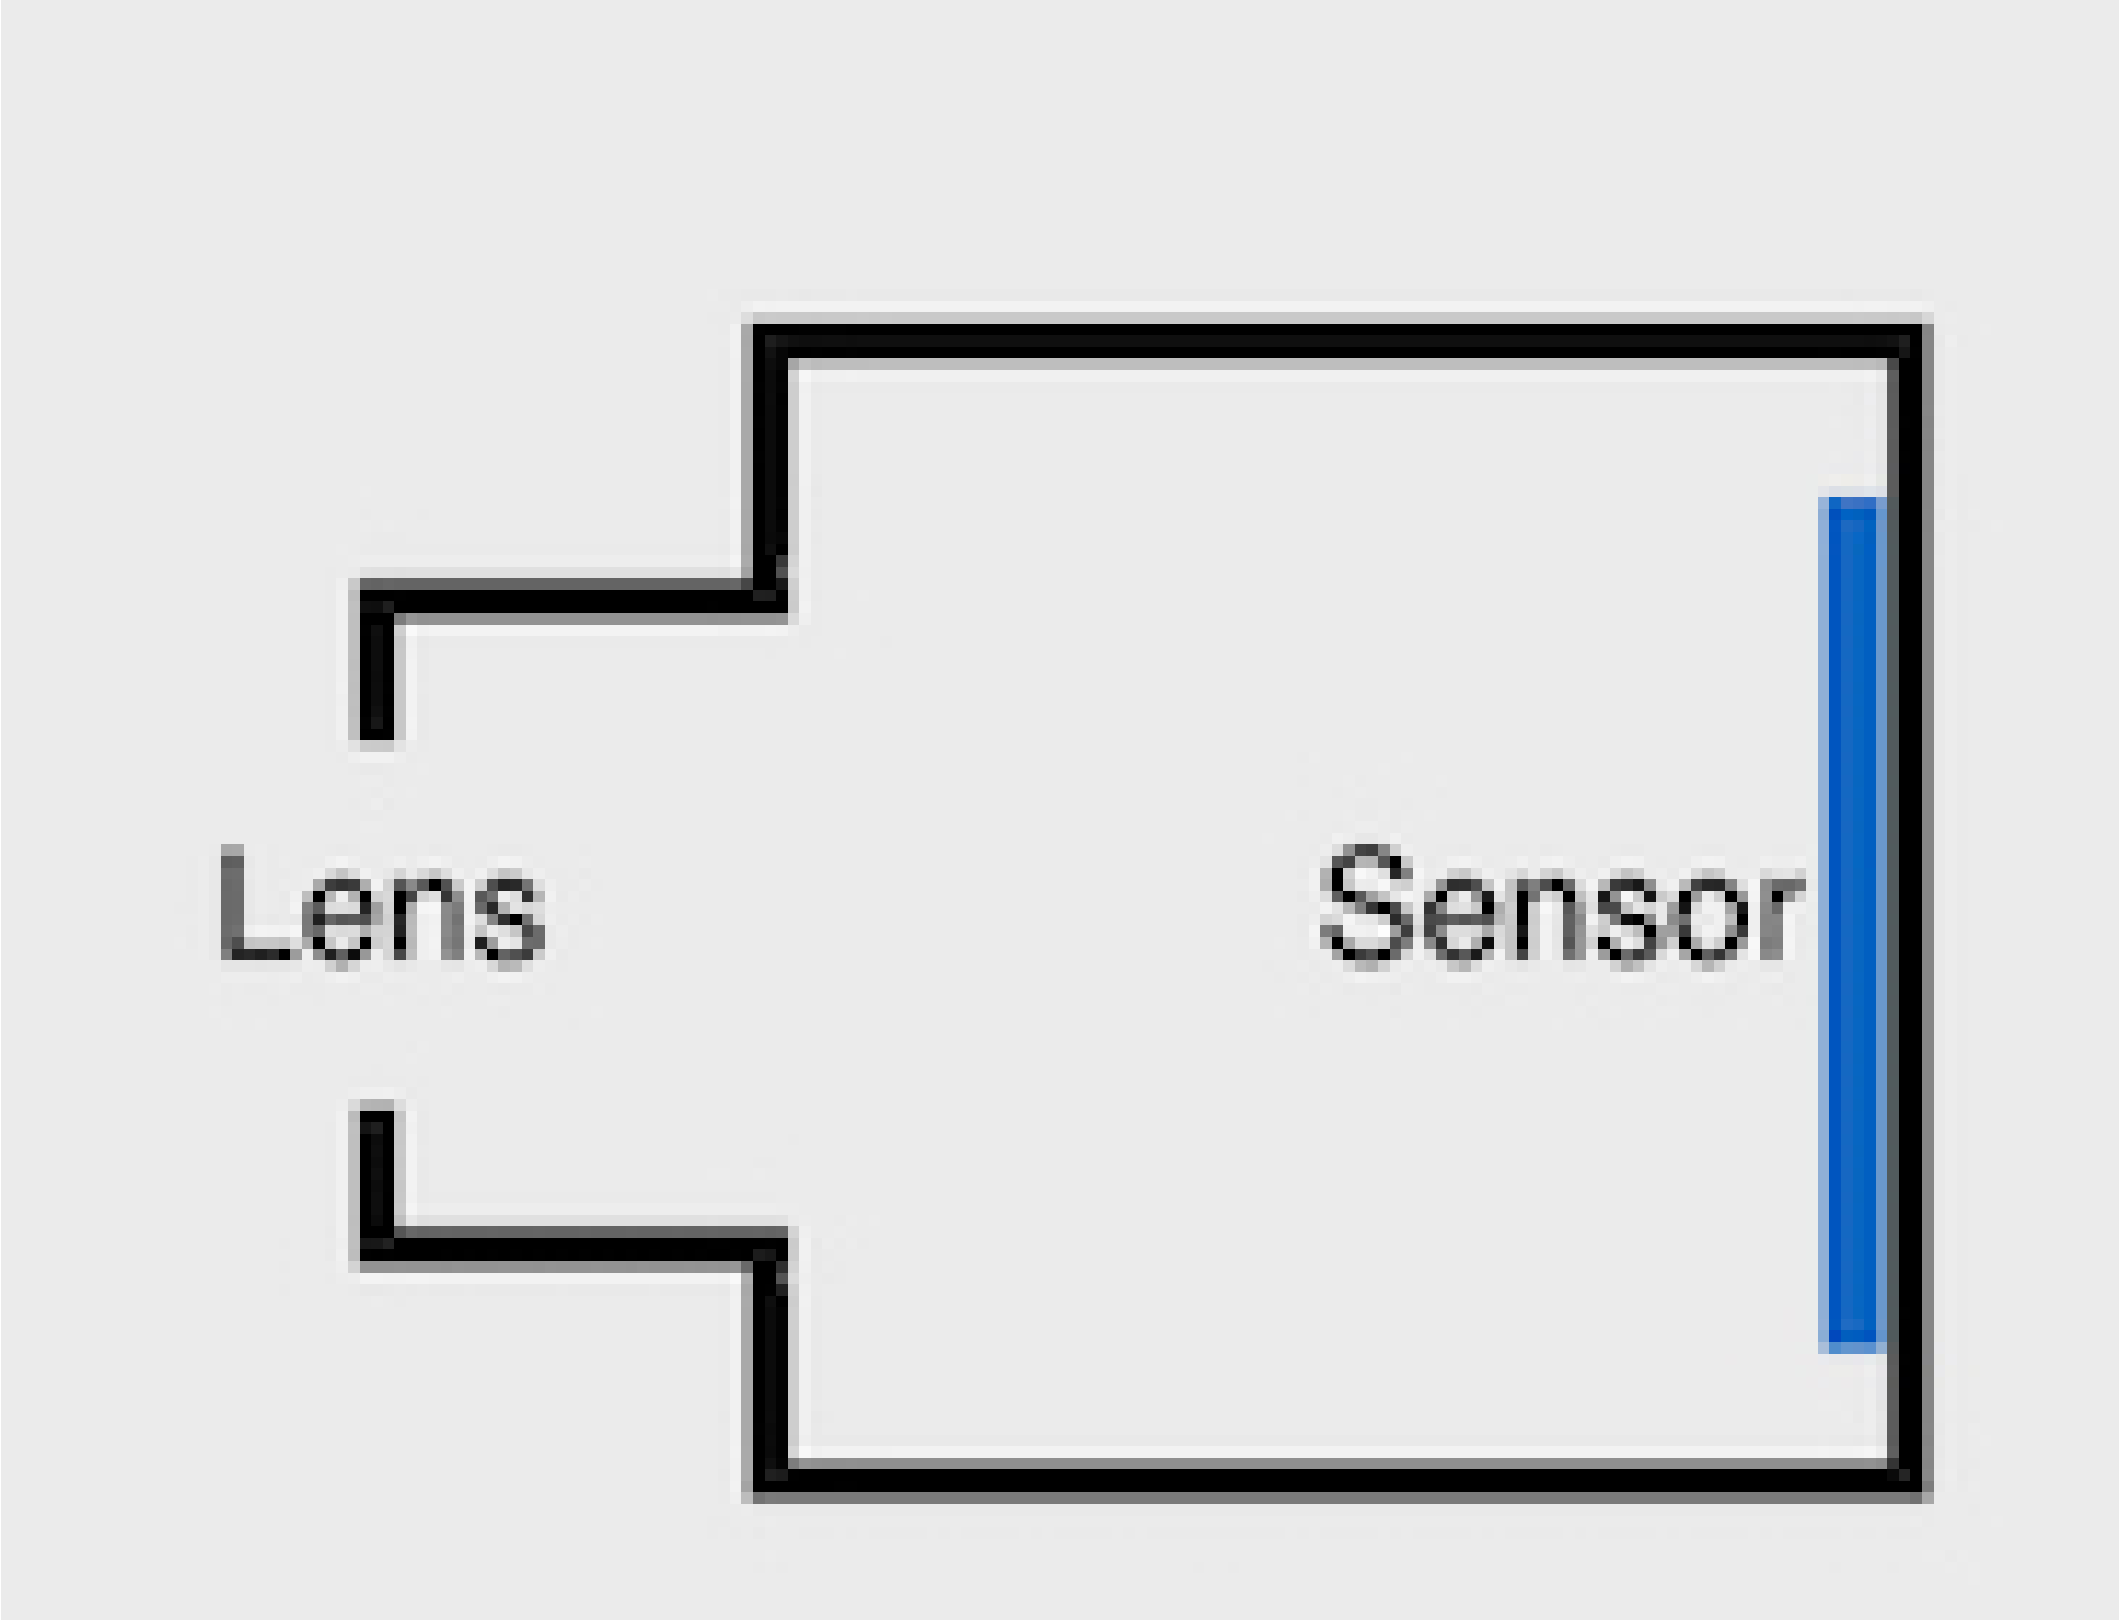

Graphic courtesy Jamie Schiel – Inspired by Laura Waller

Beyond Vision – Camera 2025

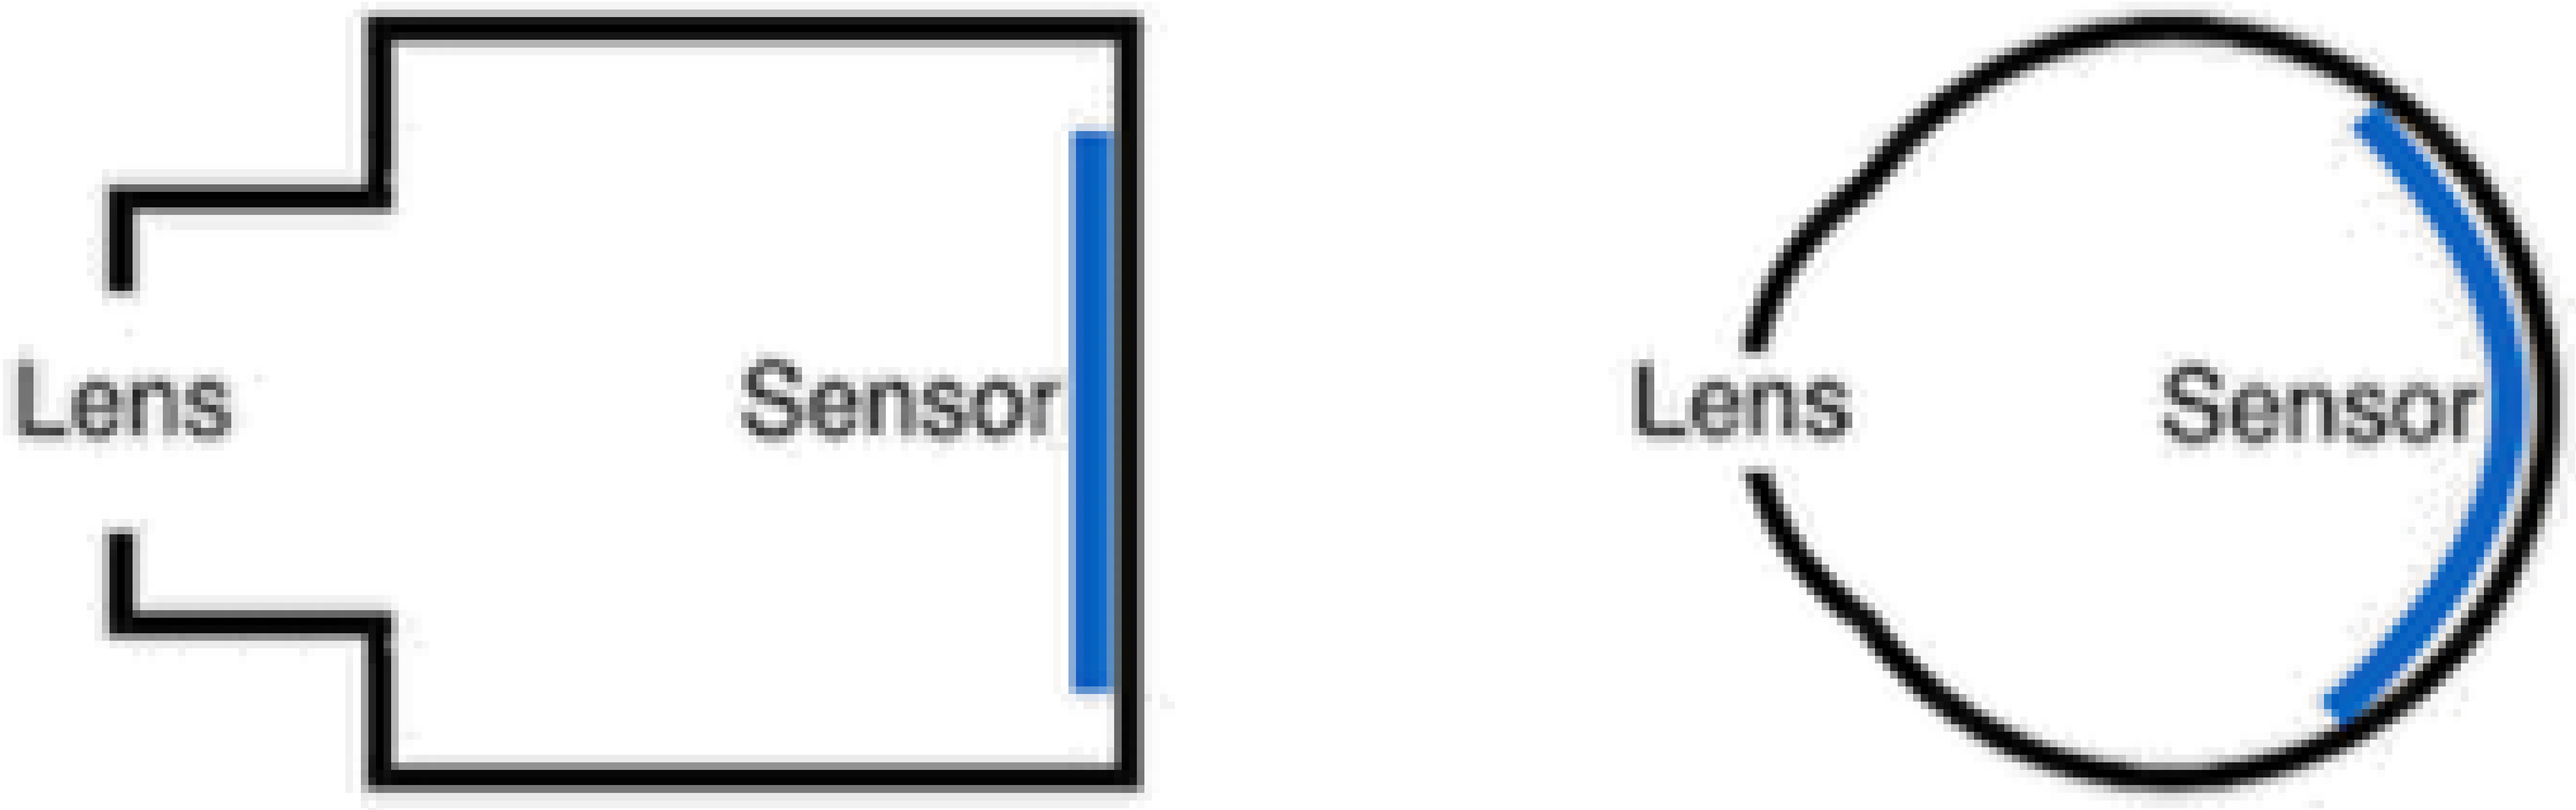

Graphic courtesy Jamie Schiel – Inspired by Laura Waller

Why should the camera mimic the human eye?

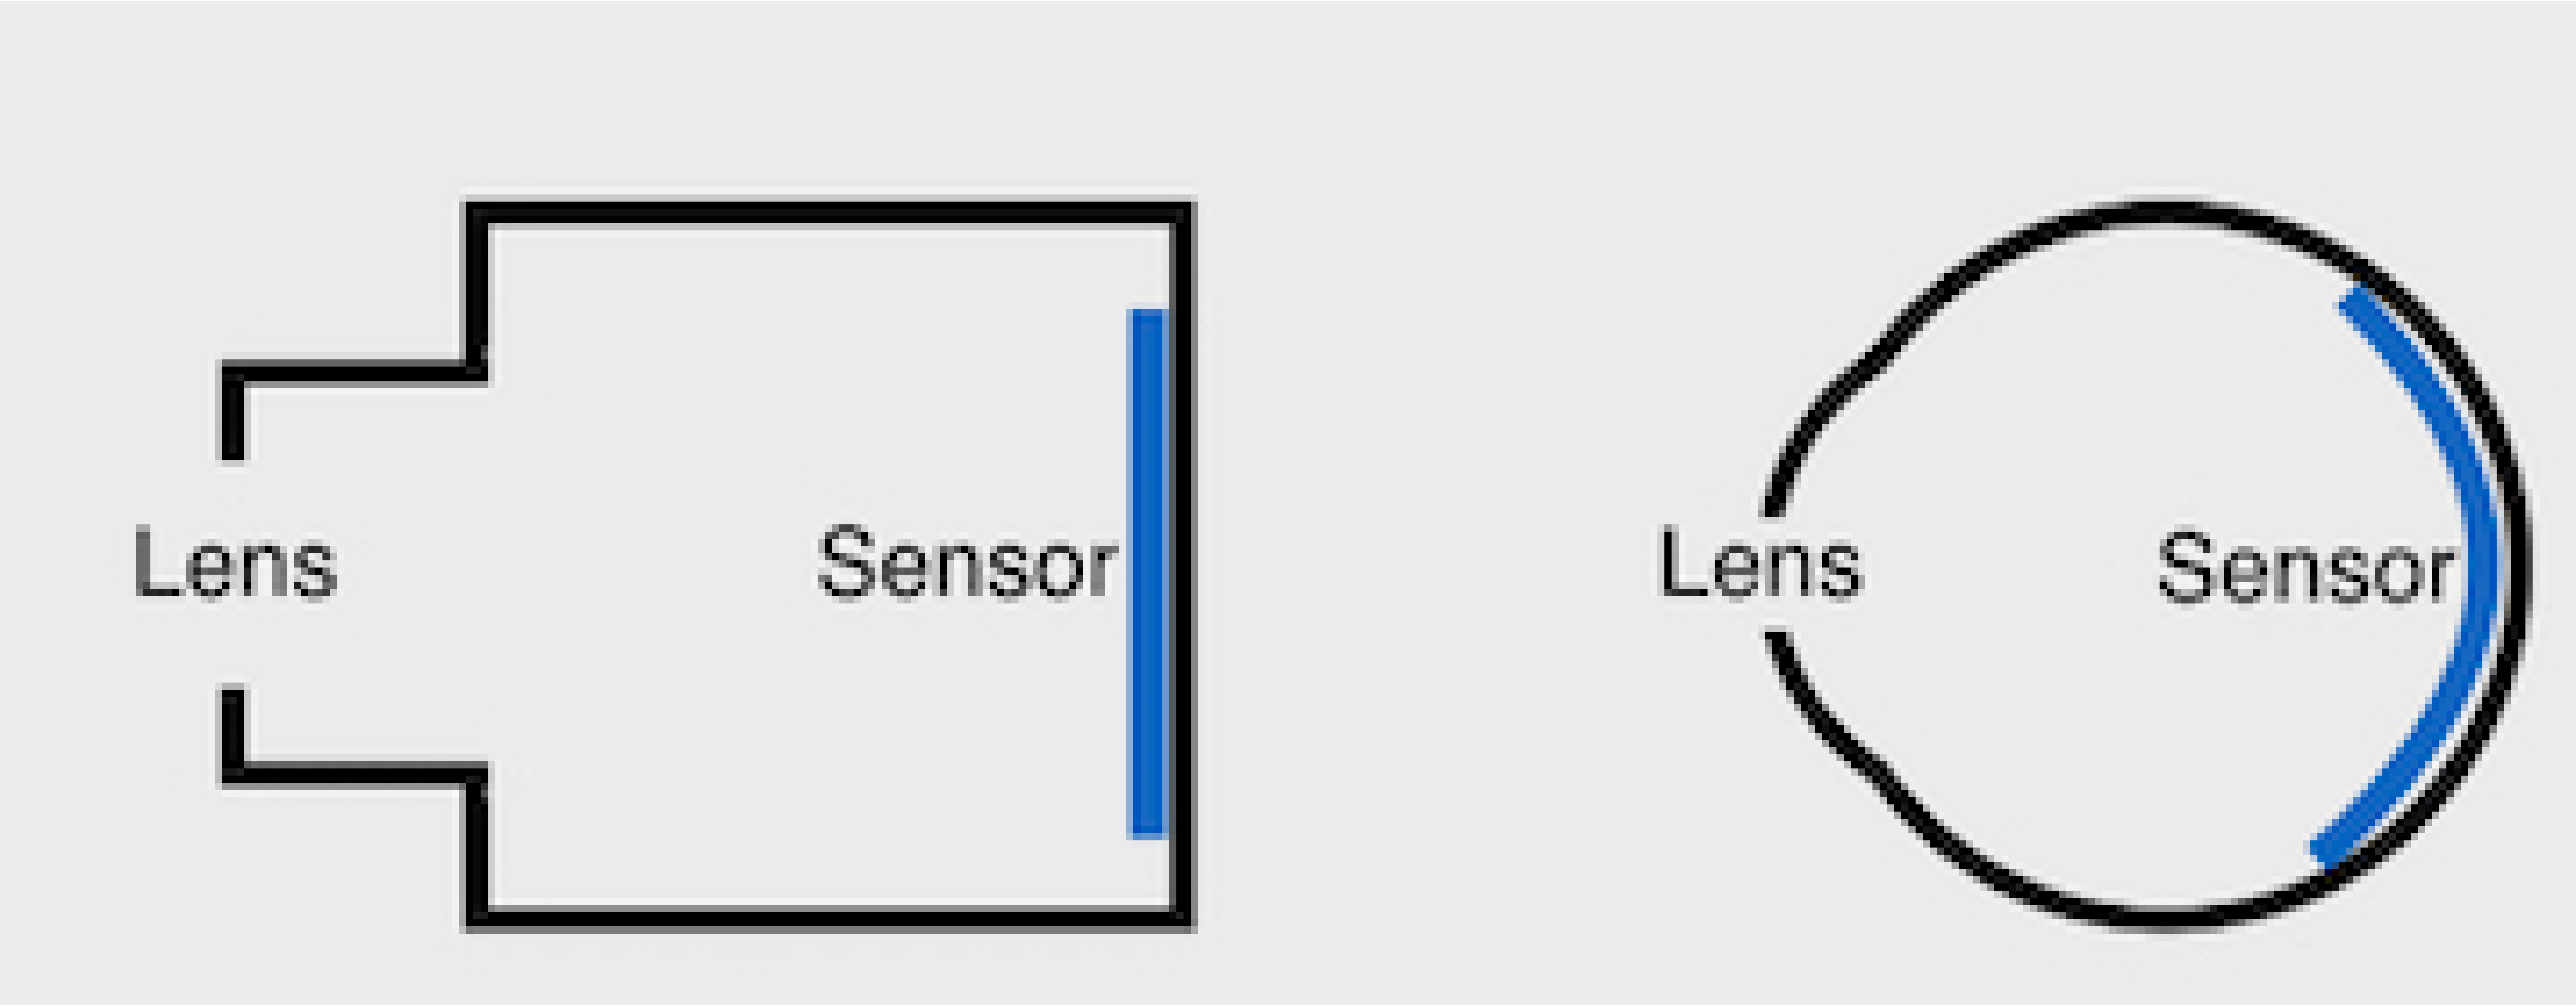

Graphic courtesy Jamie Schiel – Inspired by Laura Waller