


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

Computational Cameras: Redefining the Image

Achuta Kadambi (*University of California, Los Angeles*)

Revolutions in Imaging




750 BC


(The First Lens)

Ancient Assyria

Revolutions in Imaging




750 BC
(The First Lens)
Ancient Assyria




1816 AD
(The First Photo)
Joseph Niepce

“Revolutions” in Imaging



750 BC
(The First Lens)
Ancient Assyria



1816 AD
(The First Photo)
Joseph Niepce

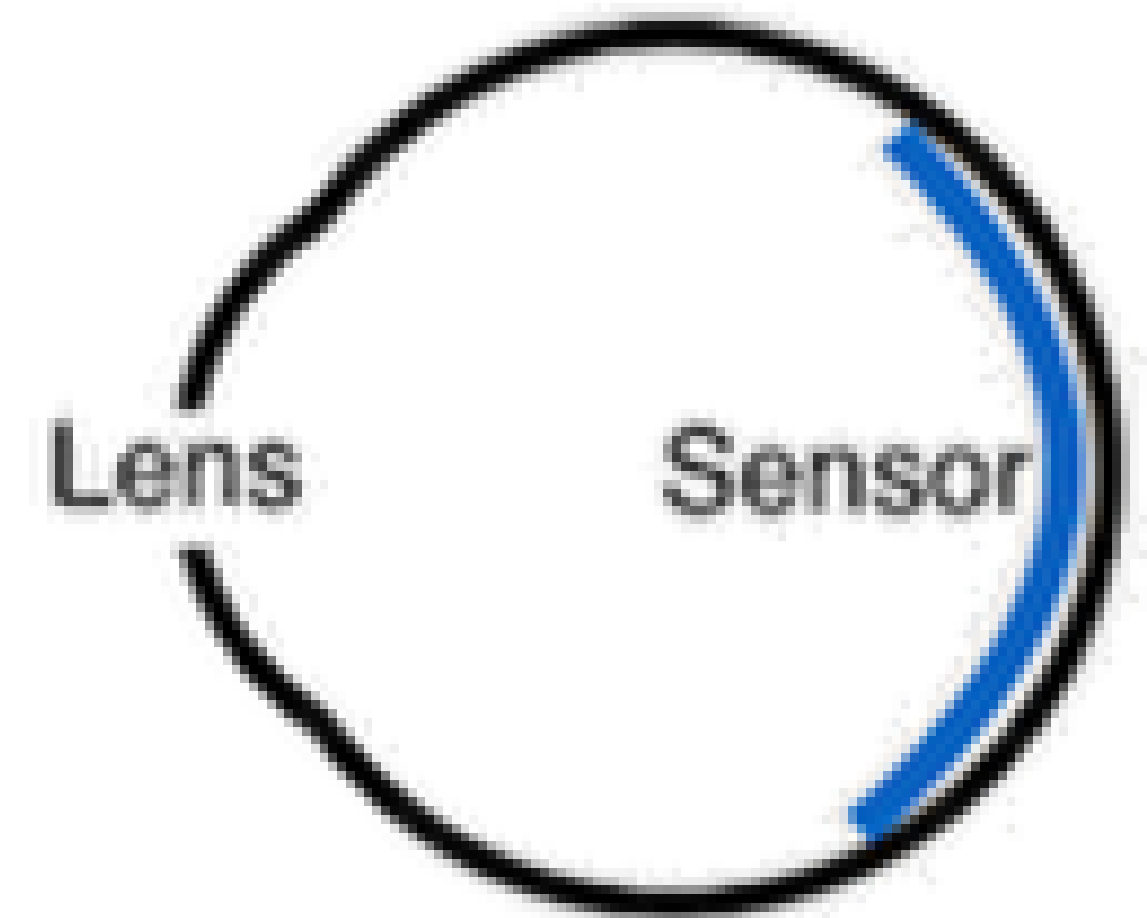
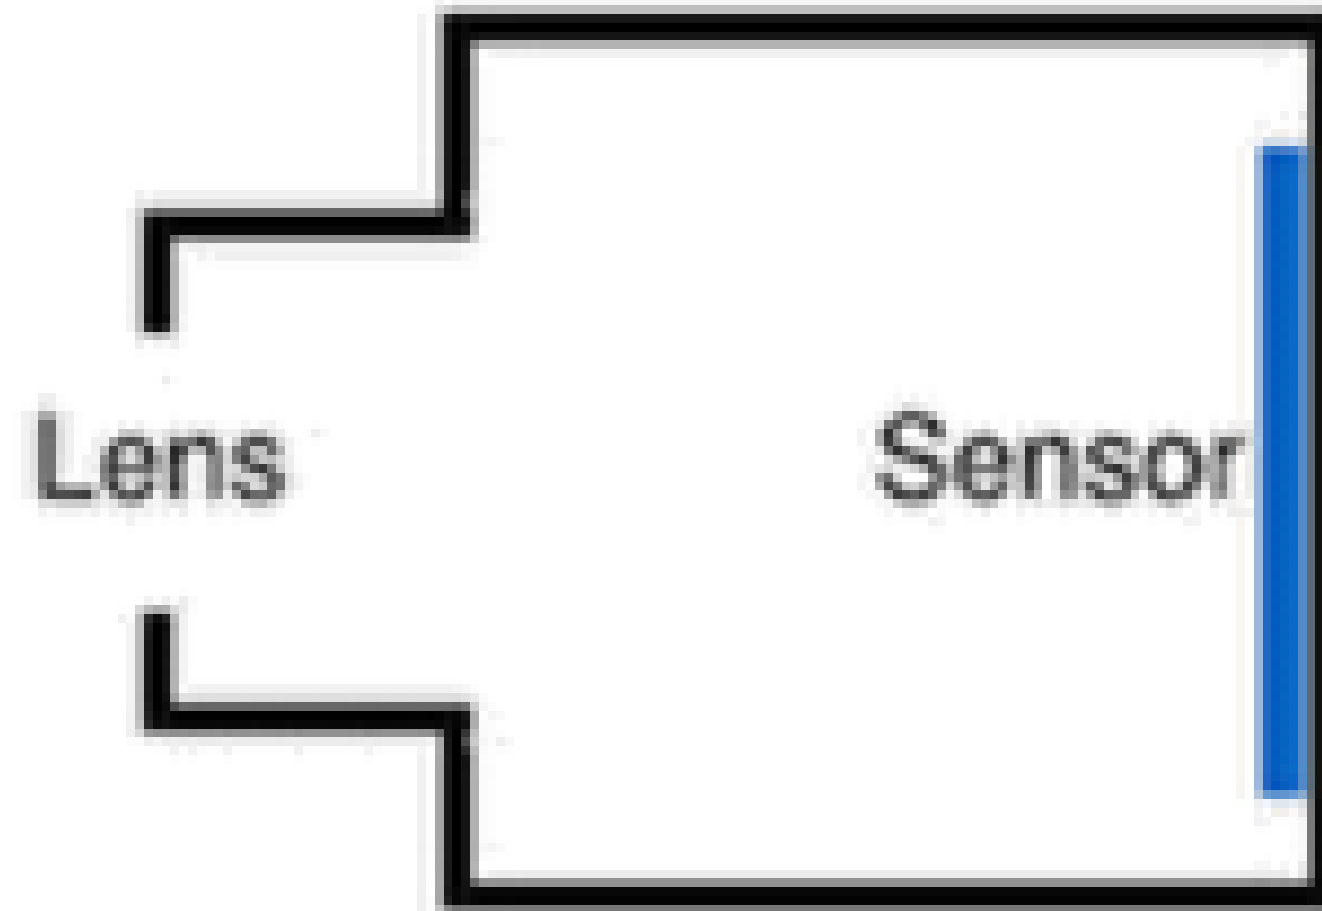


Is this really a revolution?
(Modern Photograph)

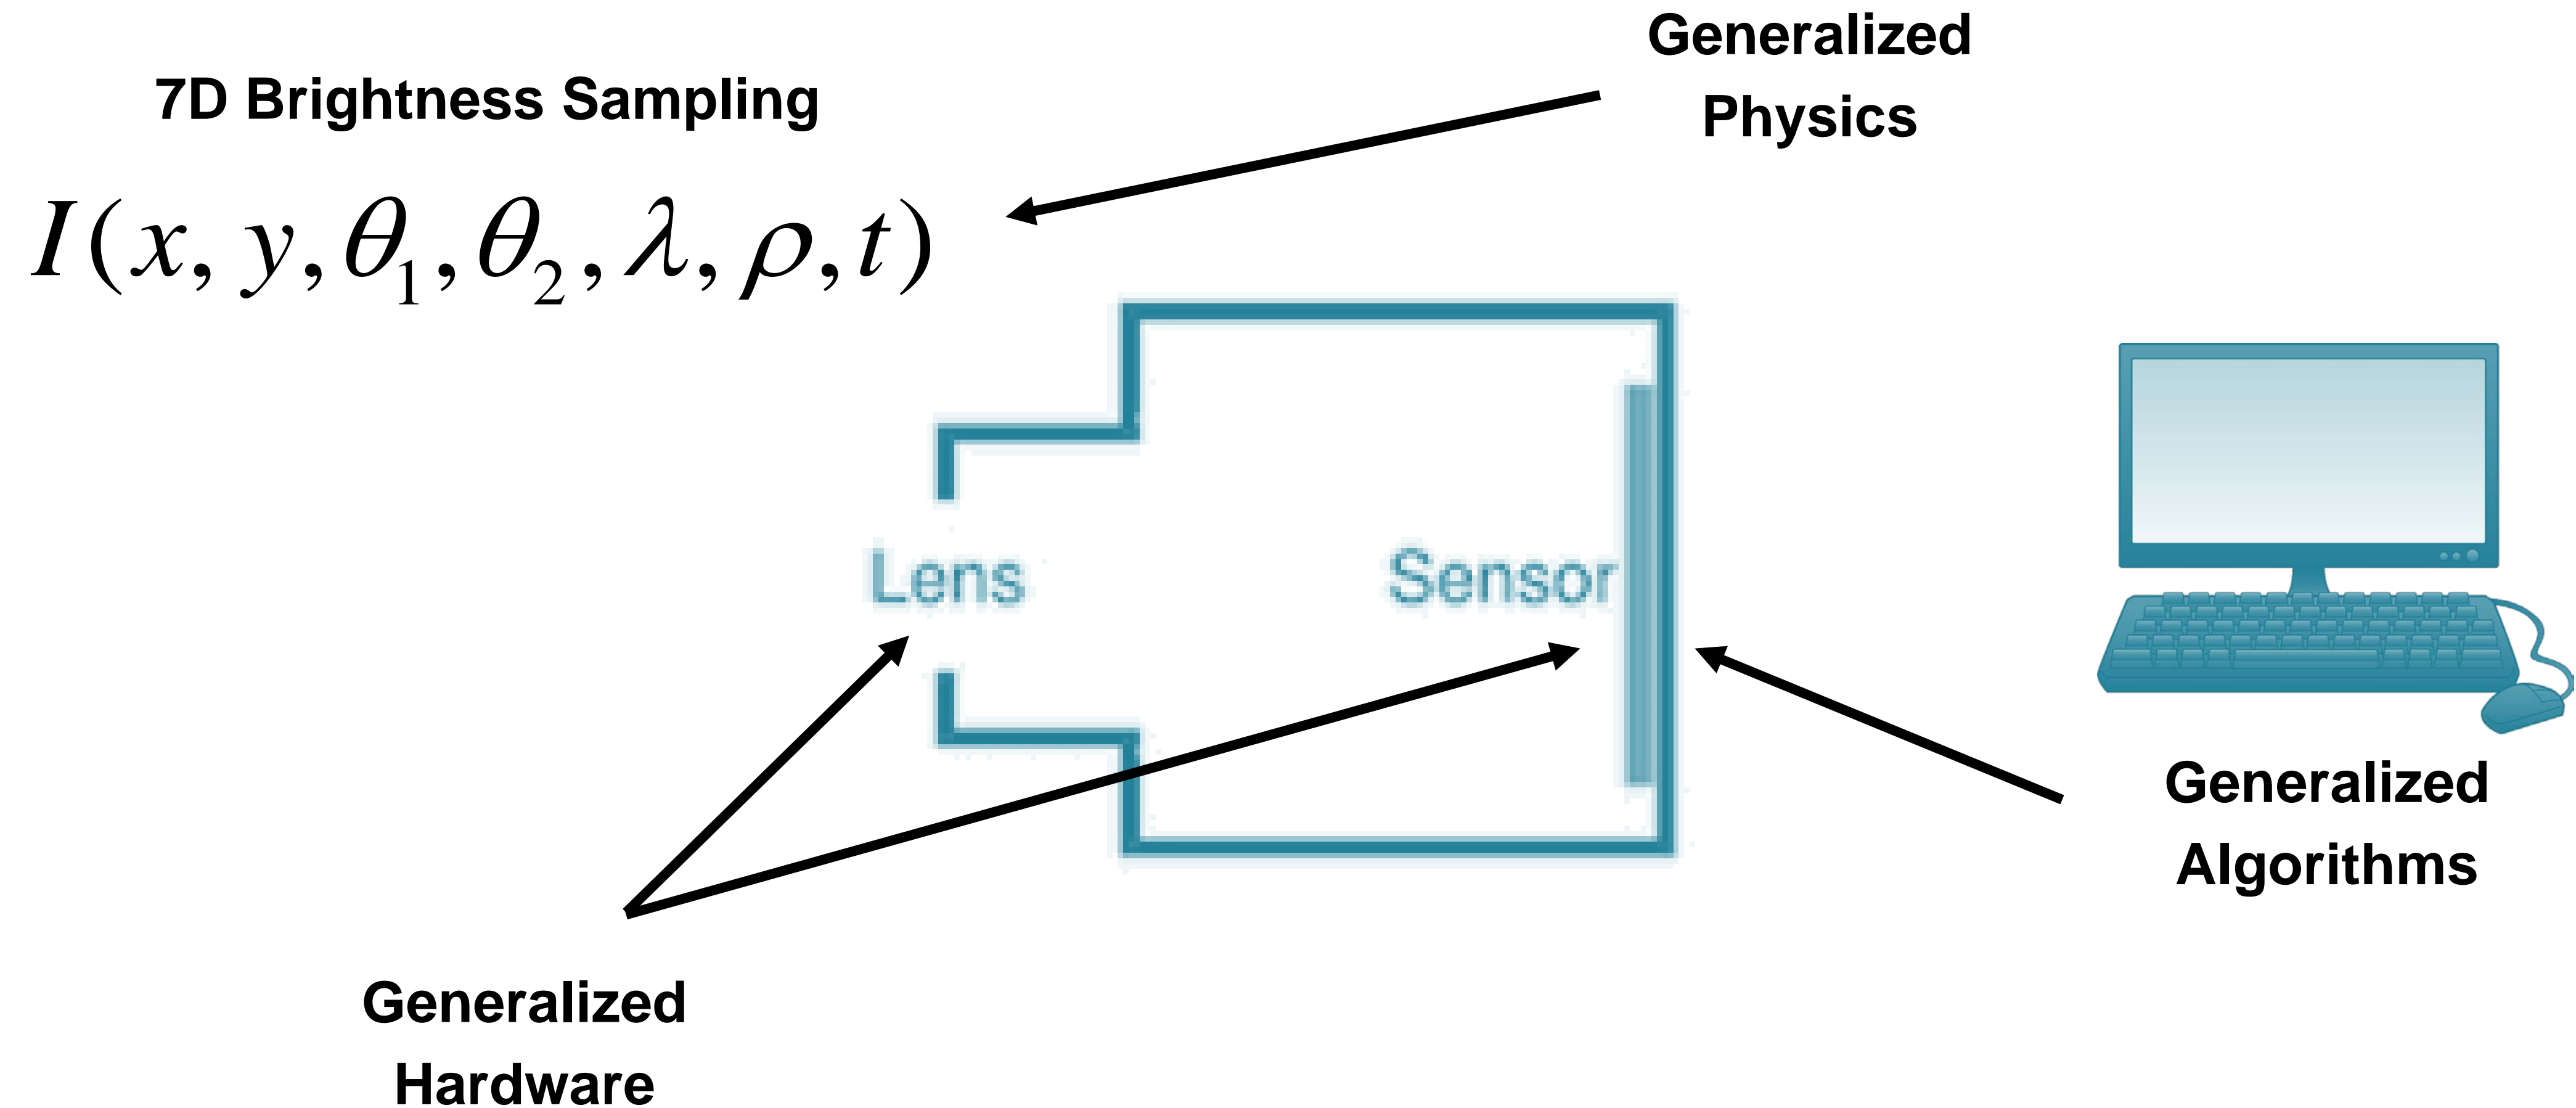
Ordinary Cameras are Boring

Capture of 2D Phenomena

$$I(x, y)$$



Computational Imaging Revolution




Physics, Hardware, and Computation

Broad Goal: Rethink existing limitations of imaging

Hardware


[IEEE'17]



[CVPR'16]



[ICCP'14]



[ToG'13]




[IJCV'17]




[ICCV'15]



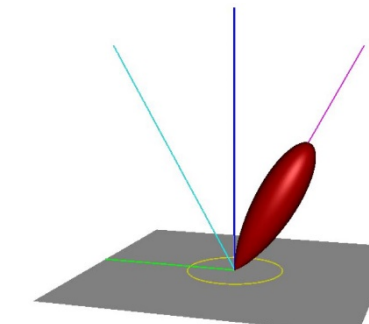
[Opt Lett'14]



[CVPR'15]



[ToG'16]




Physics

Computation

Physics, Hardware, and Computation

Hardware


[IEEE'17]



[CVPR'16]



[ICCP'14]



[ToG'13]




Hot Topic: 3D Imaging

[IJCV'17]



[ICCV'15]




[Opt Lett'14]



[CVPR'15]



[ToG'16]




Physics

Computation

3D Cameras are Ready to Disrupt

3D cameras capture the (x,y,z) position of light reflections



Helping cars navigate



Guiding medical robots



Scanning and printing objects





Microsoft Kinect v2





Microsoft Kinect v2






Multistriple Laser Scan



NextEngine 3D

\$3000 USD

Raster




Multistripe Laser Scan



NextEngine 3D

\$3000 USD

Raster




Bring in Physics

7D Brightness Sampling

$$I(x, y, \theta_1, \theta_2, \lambda, \rho, t)$$


Generalized
Physics




Generalized
Hardware

Generalized
Computation

Polarization of Light




Polarization of Light





Plane of Polarization


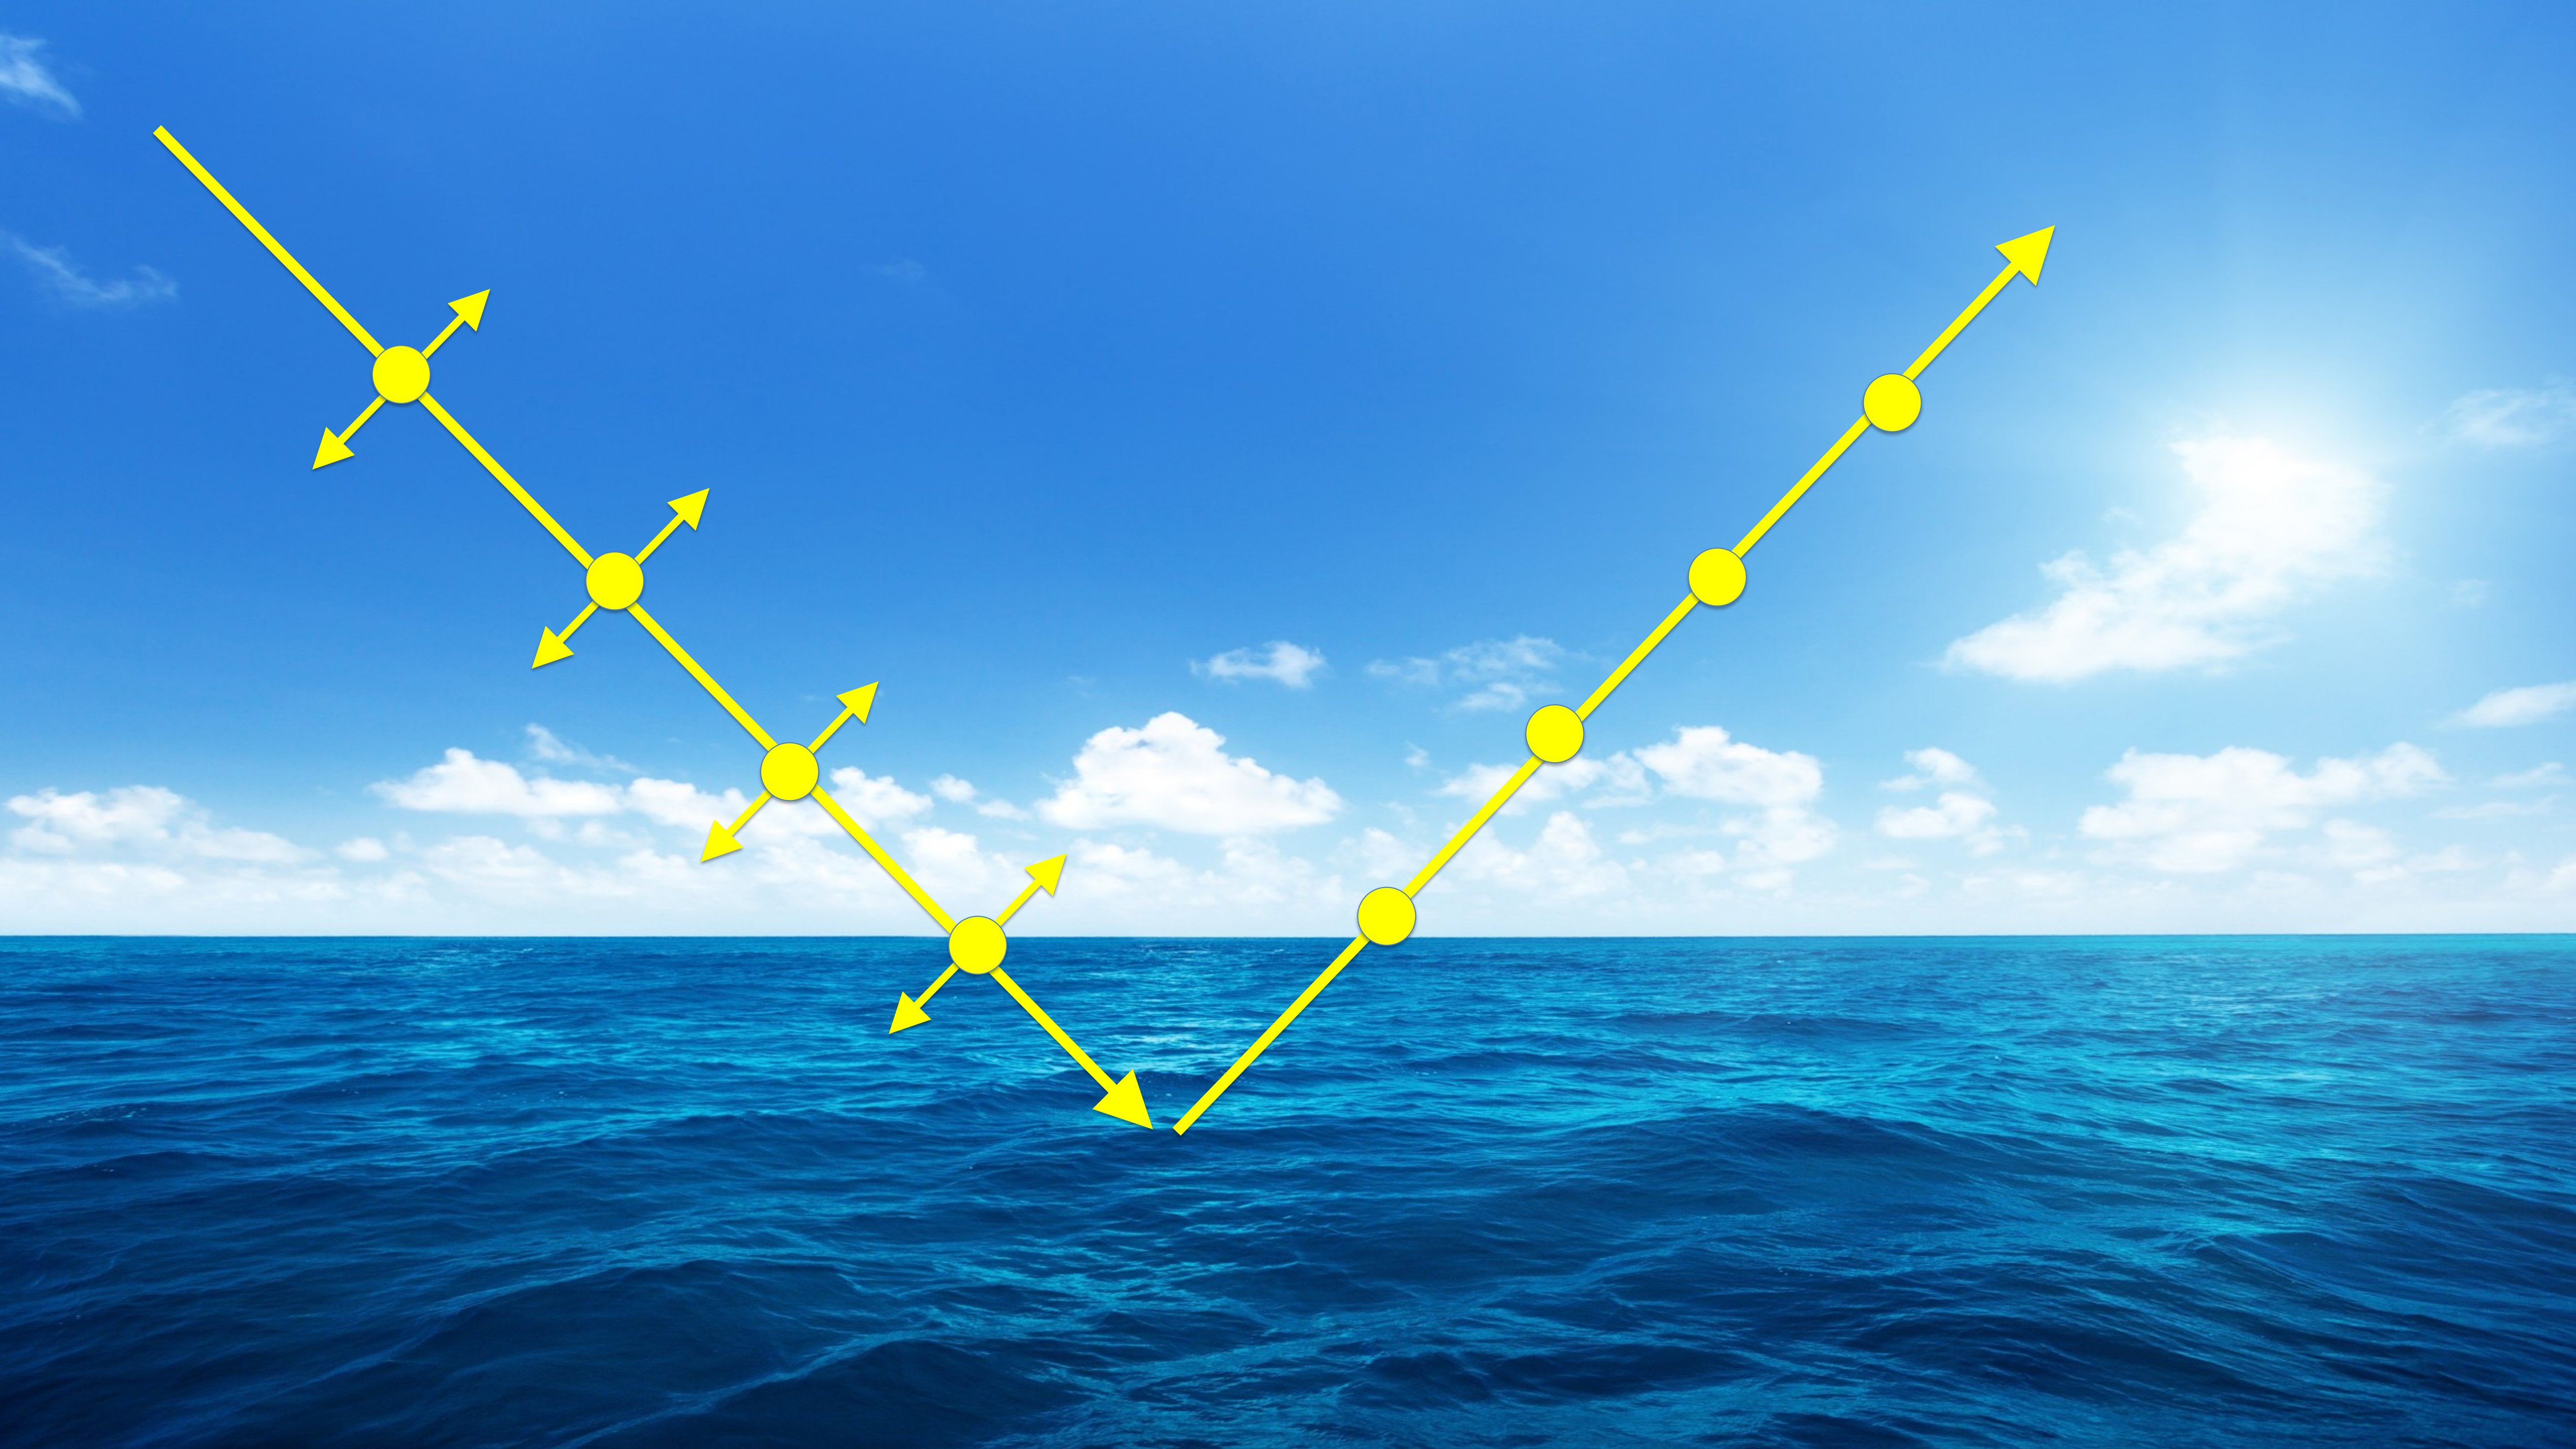
Polarization of Light



Plane of Polarization




Plane of Polarization







\$30 Polarizing Filter



Photos at Different Angles




Horizontal Filter Orientation



Vertical Filter Orientation


Colors map to 8-bit grayscale



Horizontal Filter Orientation




Vertical Filter Orientation



Difference Image


Colors map to 8-bit grayscale



Horizontal Filter Orientation



Vertical Filter Orientation

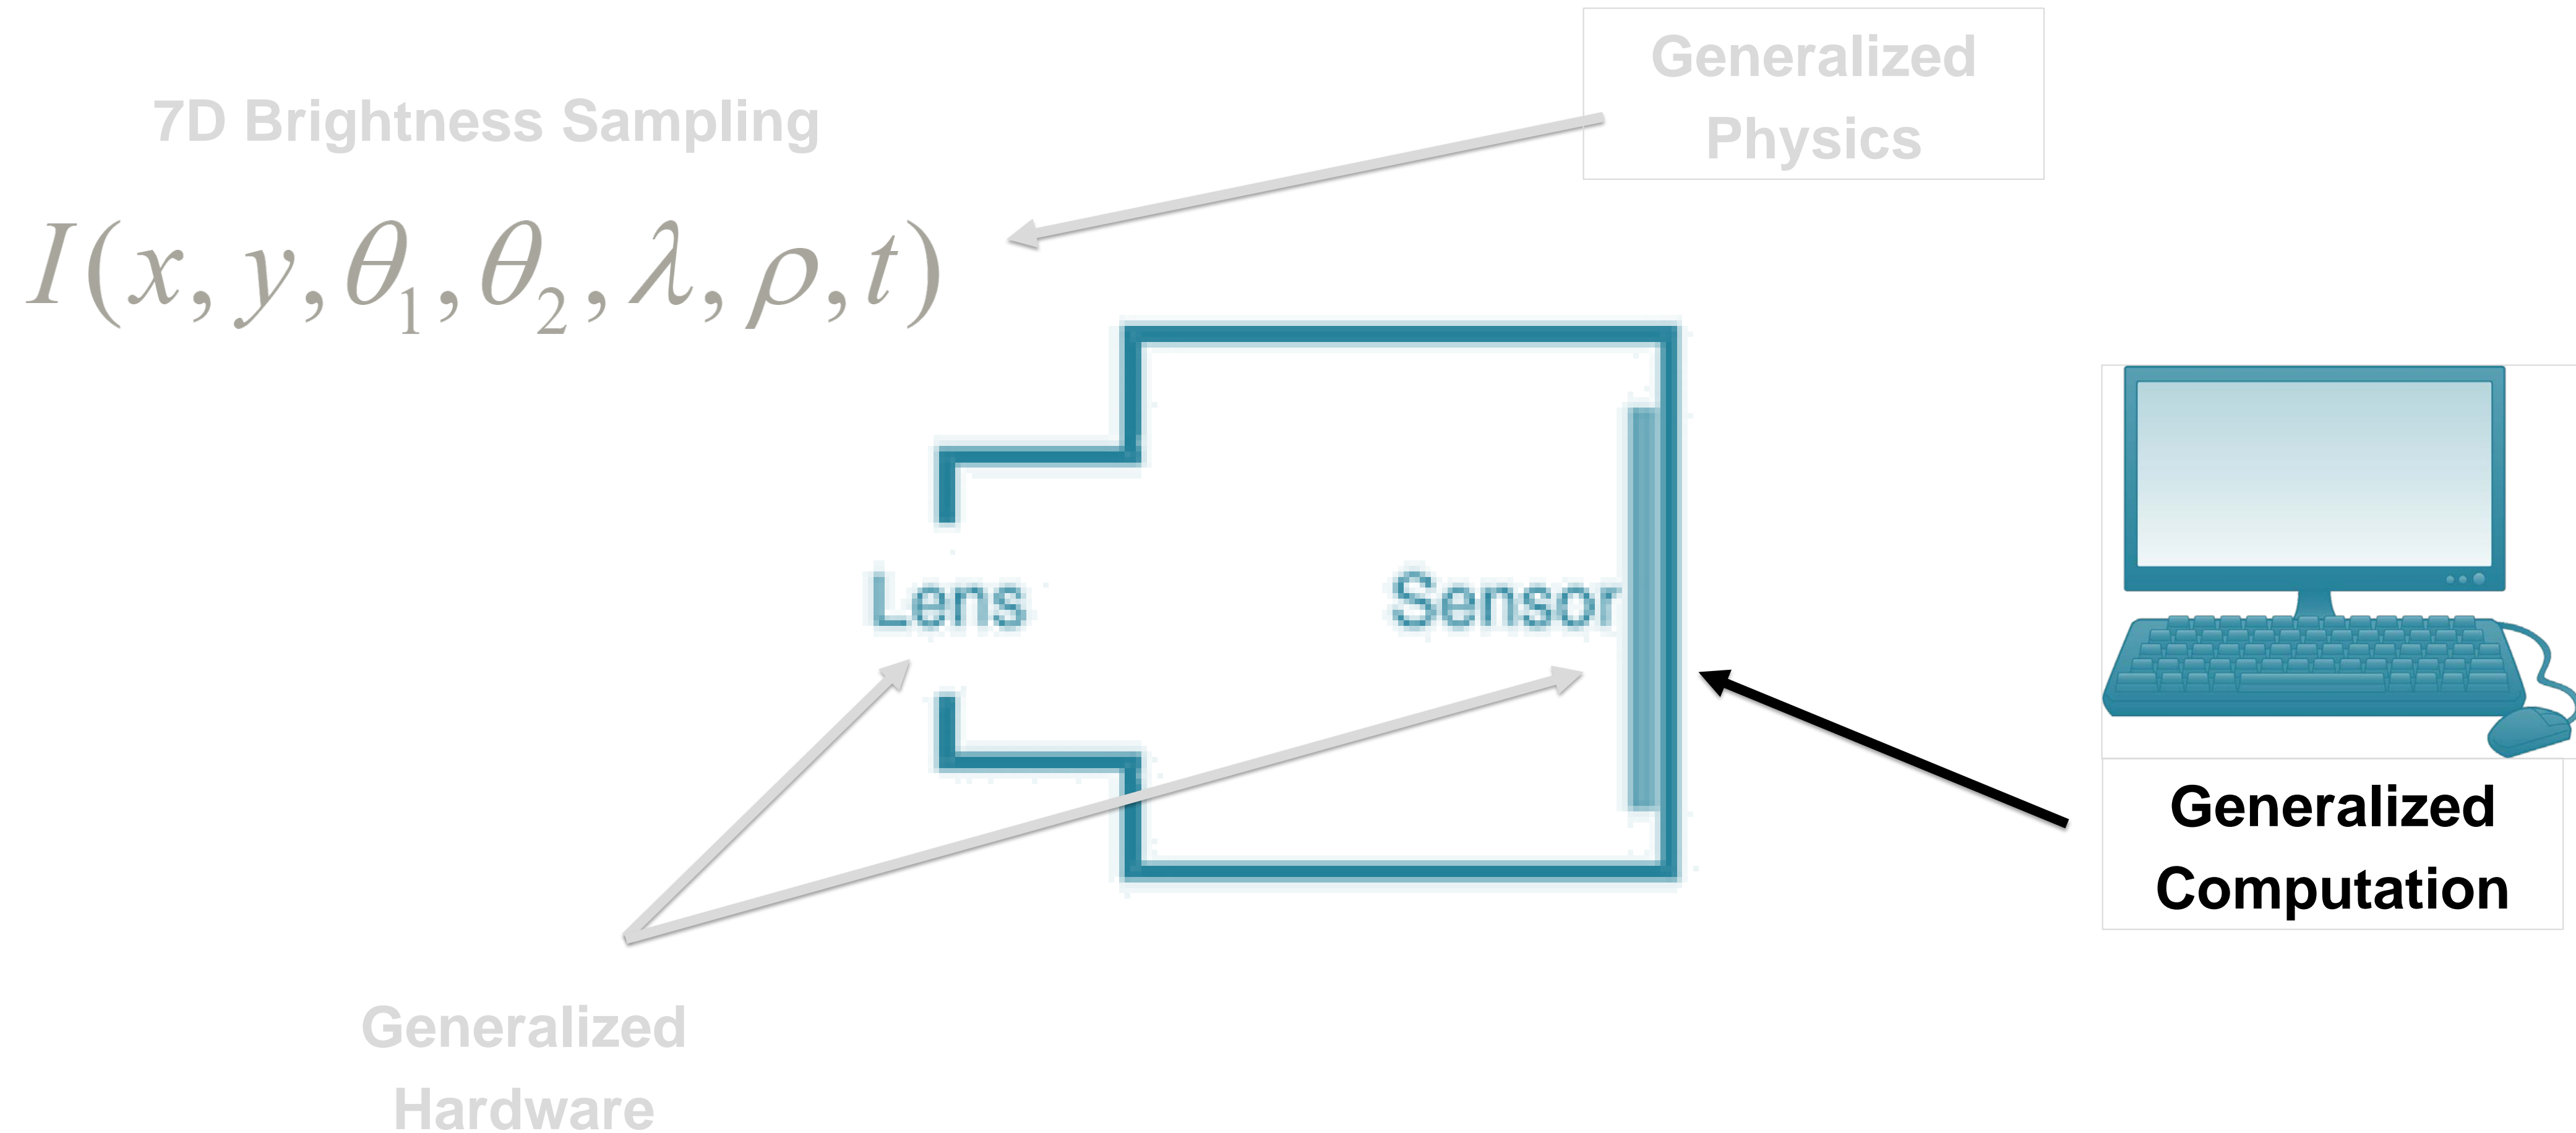


Difference Image
(Tonemapped to 8-bit)

We need to add computation!

Colors map to 8-bit grayscale

Final Ingredient: Bring in Computation



“Polarized 3D”




Computational Imaging
Using Polarization

“Polarized 3D”




Computational Imaging
Using Polarization




“Polarized 3D”



Computational Imaging
Using Polarization





Fresnel Electromagnetic Equations




Fresnel Electromagnetic Equations

Integrating normals to
obtain 3D shape




Fresnel Electromagnetic Equations

Integrating normals to
obtain 3D shape




Goal: Solve for zenith angle at
each position




Fresnel Electromagnetic Equations

Integrating normals to
obtain 3D shape




Goal: Solve for zenith angle at
each position




Fresnel Electromagnetic Equations

Integrating normals to
obtain 3D shape



Goal: Solve for zenith angle at
each position



Fresnel Electromagnetic Equations


Old Principle [Fresnel 1819]

$$r_{\perp} = \frac{\cos \theta_i - n \cos \theta_t}{\cos \theta_i + n \cos \theta_t}$$

$$r_{\parallel} = \frac{\cos \theta_i - n \cos \theta_t}{\cos \theta_t + n \cos \theta_i}$$

SfP crux: Solve for theta

Goal: Solve for zenith angle at each position



Fresnel Electromagnetic Equations

Old Principle [Fresnel 1819]


$$r_{\perp} = \frac{\cos \theta_i - n \cos \theta_t}{\cos \theta_i + n \cos \theta_t}$$

$$r_{\parallel} = \frac{\cos \theta_i - n \cos \theta_t}{\cos \theta_t + n \cos \theta_i}$$

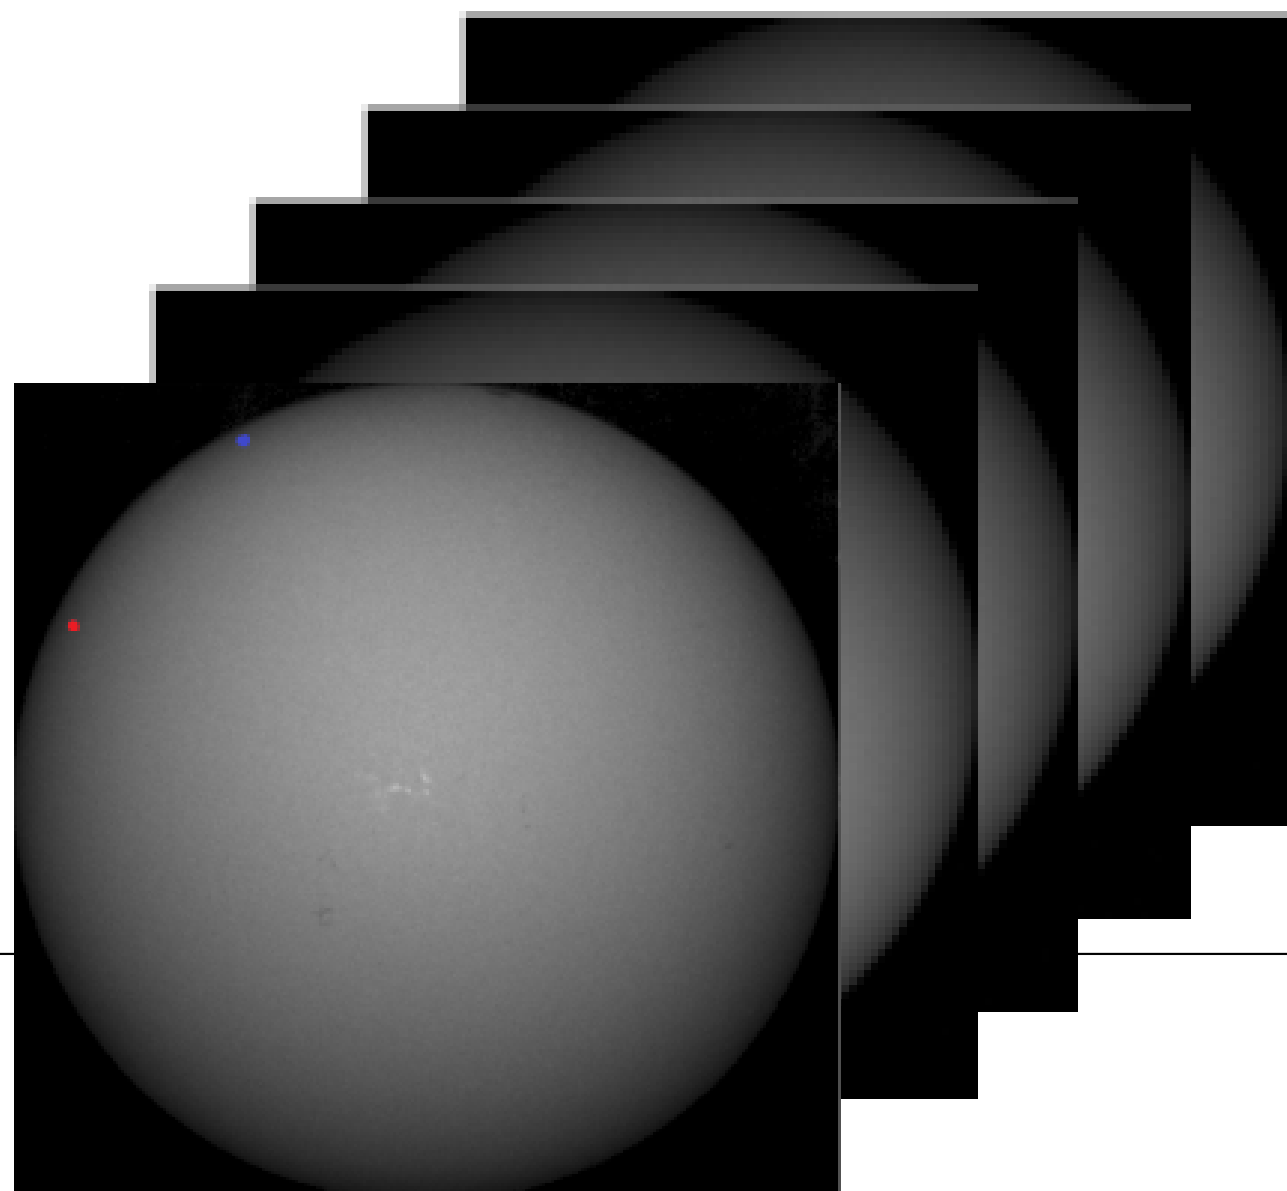
SfP crux: Solve for theta



⚠ Need to know refractive index

Goal: Solve for zenith angle at each position









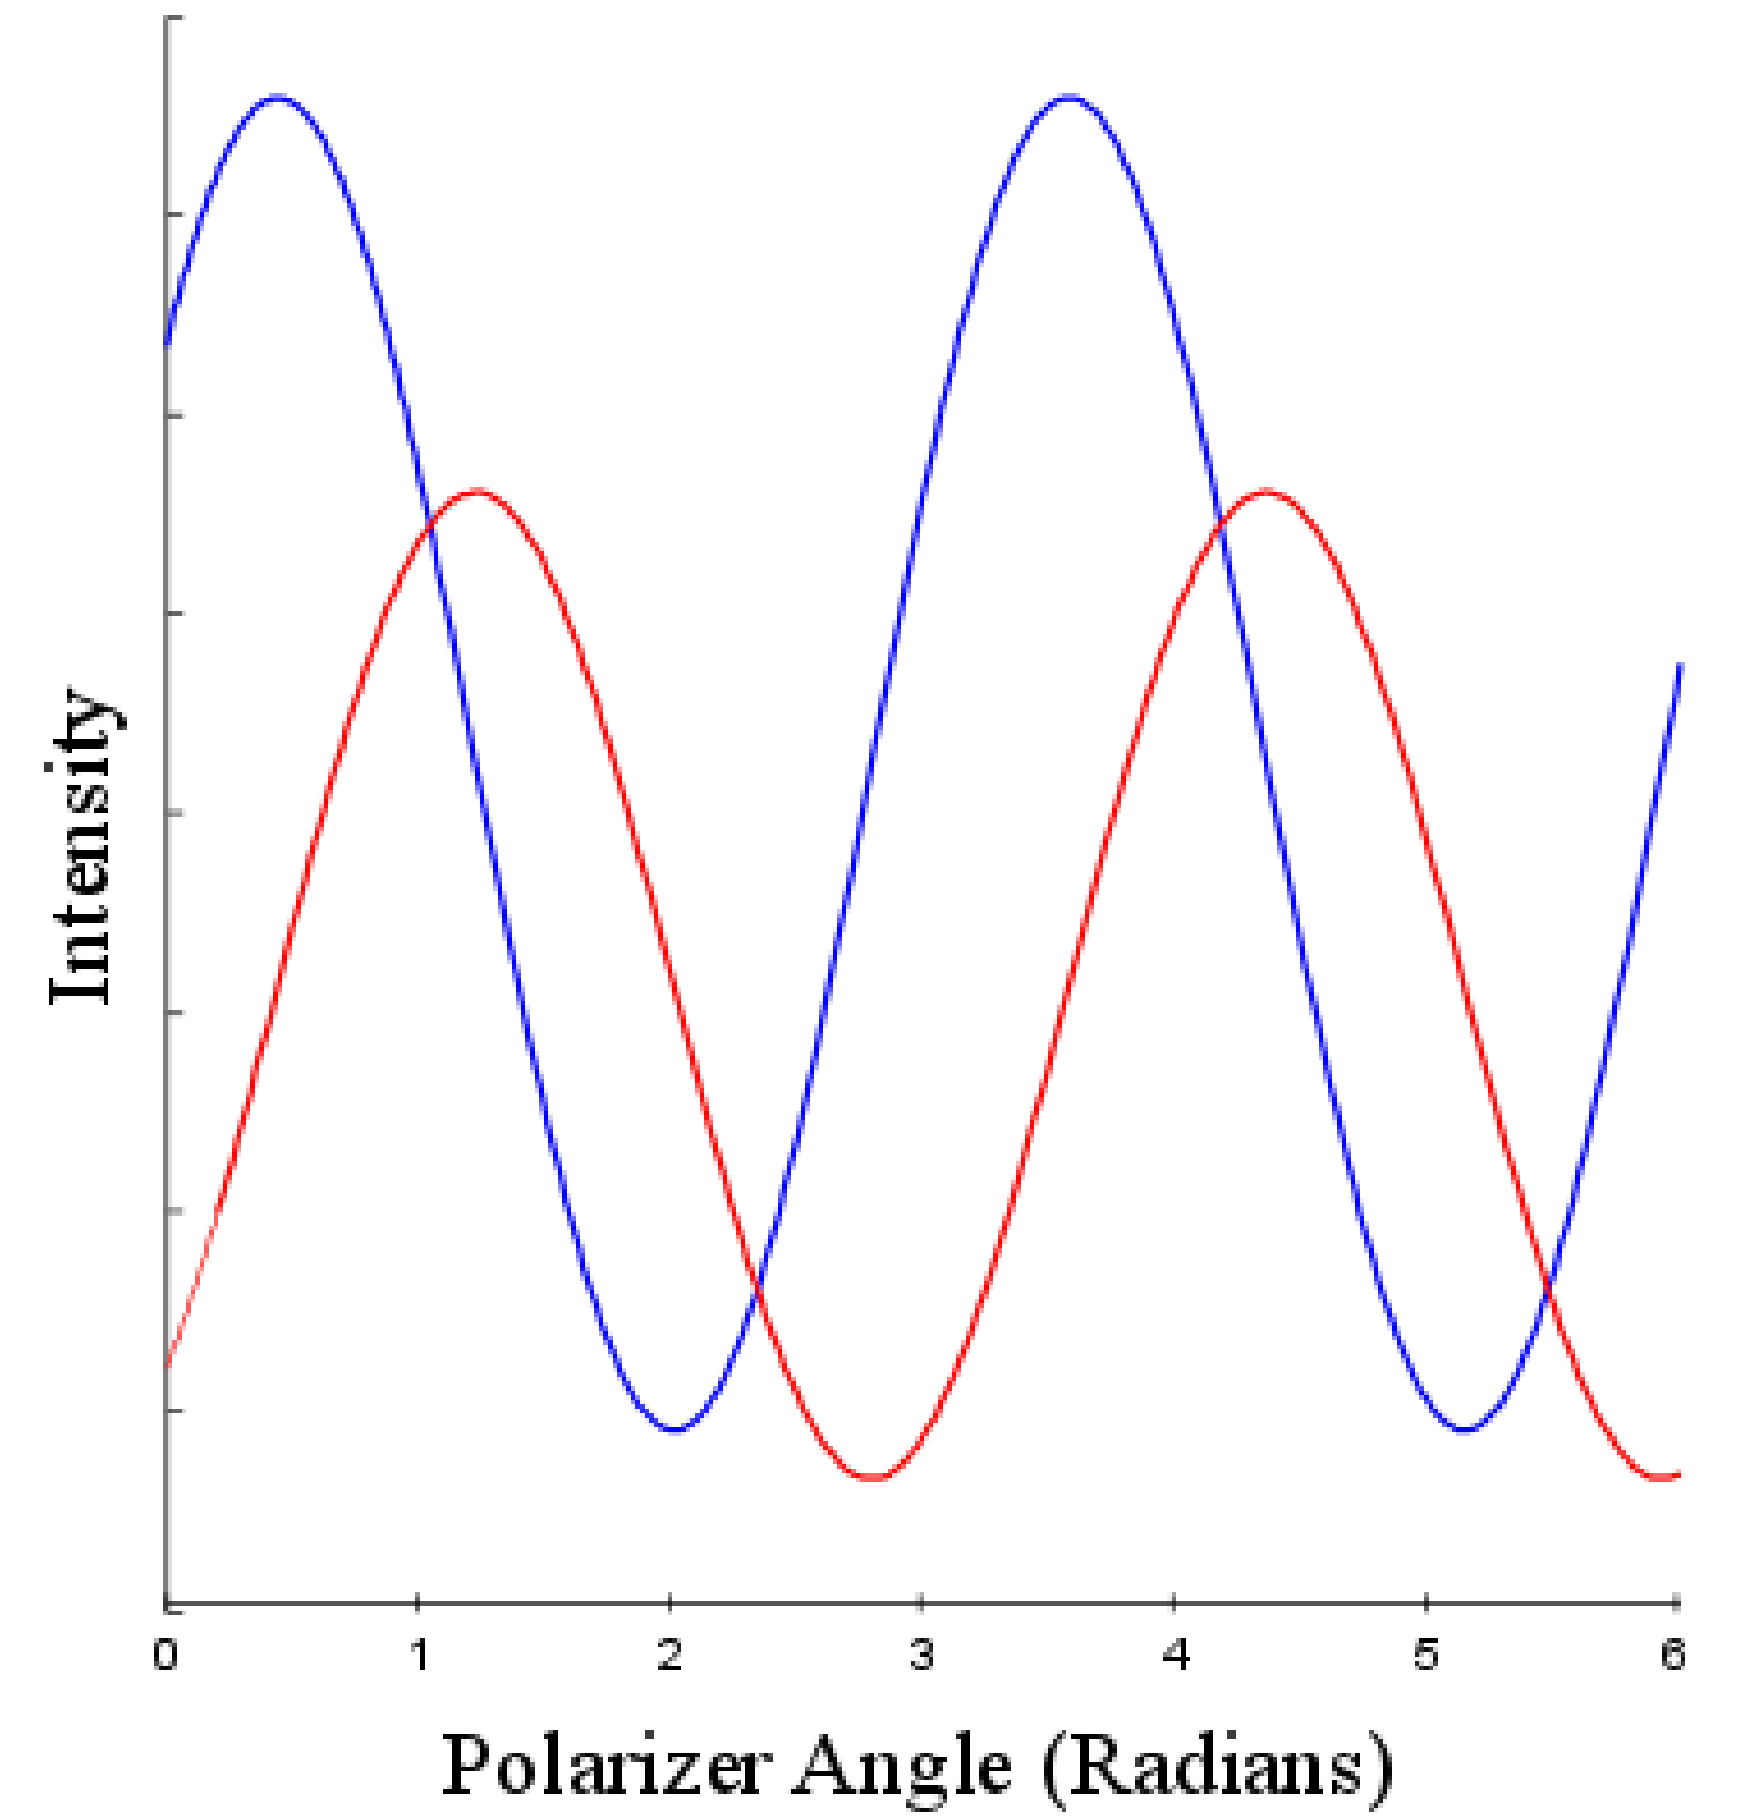
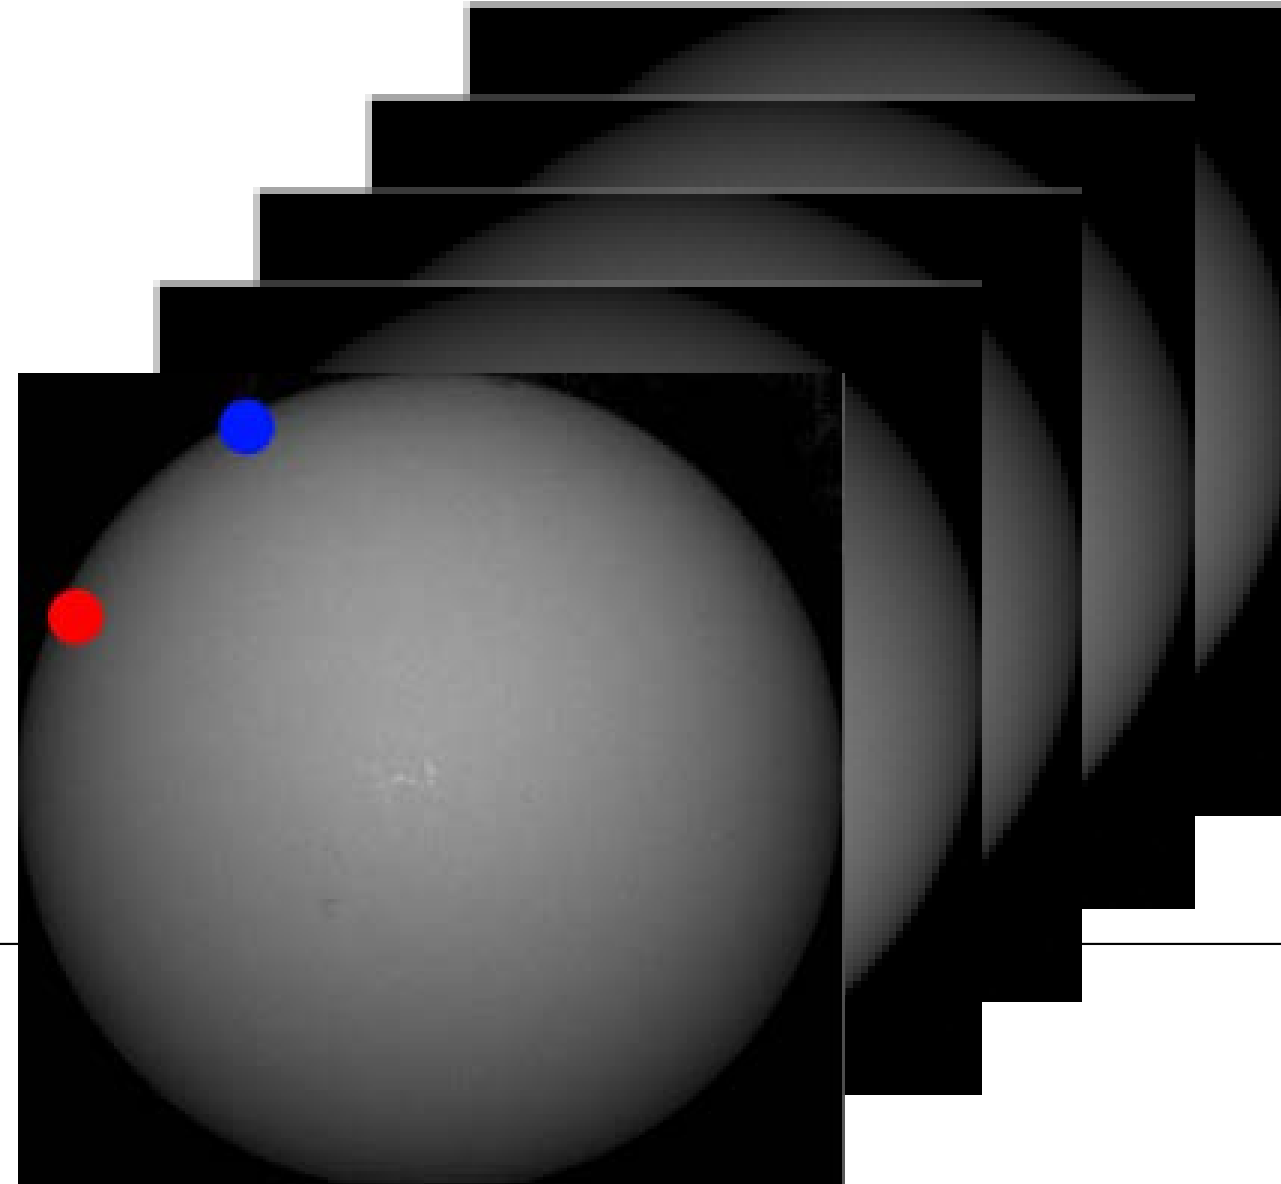




Image Formation Model



$$I(\phi_{\text{pol}}) = \frac{I_{\max} + I_{\min}}{2} + \frac{I_{\max} - I_{\min}}{2} \cos\left(2\left(\phi_{\text{pol}} - \phi\right)\right)$$

Goal: Solve for azimuth phase

Image Formation Model




$$I(\phi_{\text{pol}}) = \frac{I_{\max} + I_{\min}}{2} + \frac{I_{\max} - I_{\min}}{2} \cos\left(2\left(\phi_{\text{pol}} - \phi\right)\right)$$

Suppose $\exists \phi$ and $\phi' = \phi + \pi$

Goal: Solve for azimuth phase

Image Formation Model



$$I(\phi_{\text{pol}}) = \frac{I_{\max} + I_{\min}}{2} + \frac{I_{\max} - I_{\min}}{2} \cos\left(2(\phi_{\text{pol}} - \phi)\right)$$

Suppose $\exists \phi$ and $\phi' = \phi + \pi$

Goal: Solve for azimuth phase


Azimuthal Ambiguity problem with 2^P solutions

TLDR: EM analysis alone is underconstrained

1. **Refractive Index Needs to be Known**
2. **π Ambiguity in Surface Normal**

EM analysis alone is underconstrained

1. **Refractive Index Needs to be Known**
2. **π Ambiguity in Surface Normal**
3. **Low SNR for some geometries**
4. **Surface conforms to existing Fresnel models**



**State of the art using
EM analysis**


Are we barking up the right tree?

Emailing Sr. Researchers: "Why did you stop working on this problem?" (Rephrased)


Exemplary Response: "The polarization signal is subtle and too many physical constants unknown"

- A Senior Professor in Computer Vision

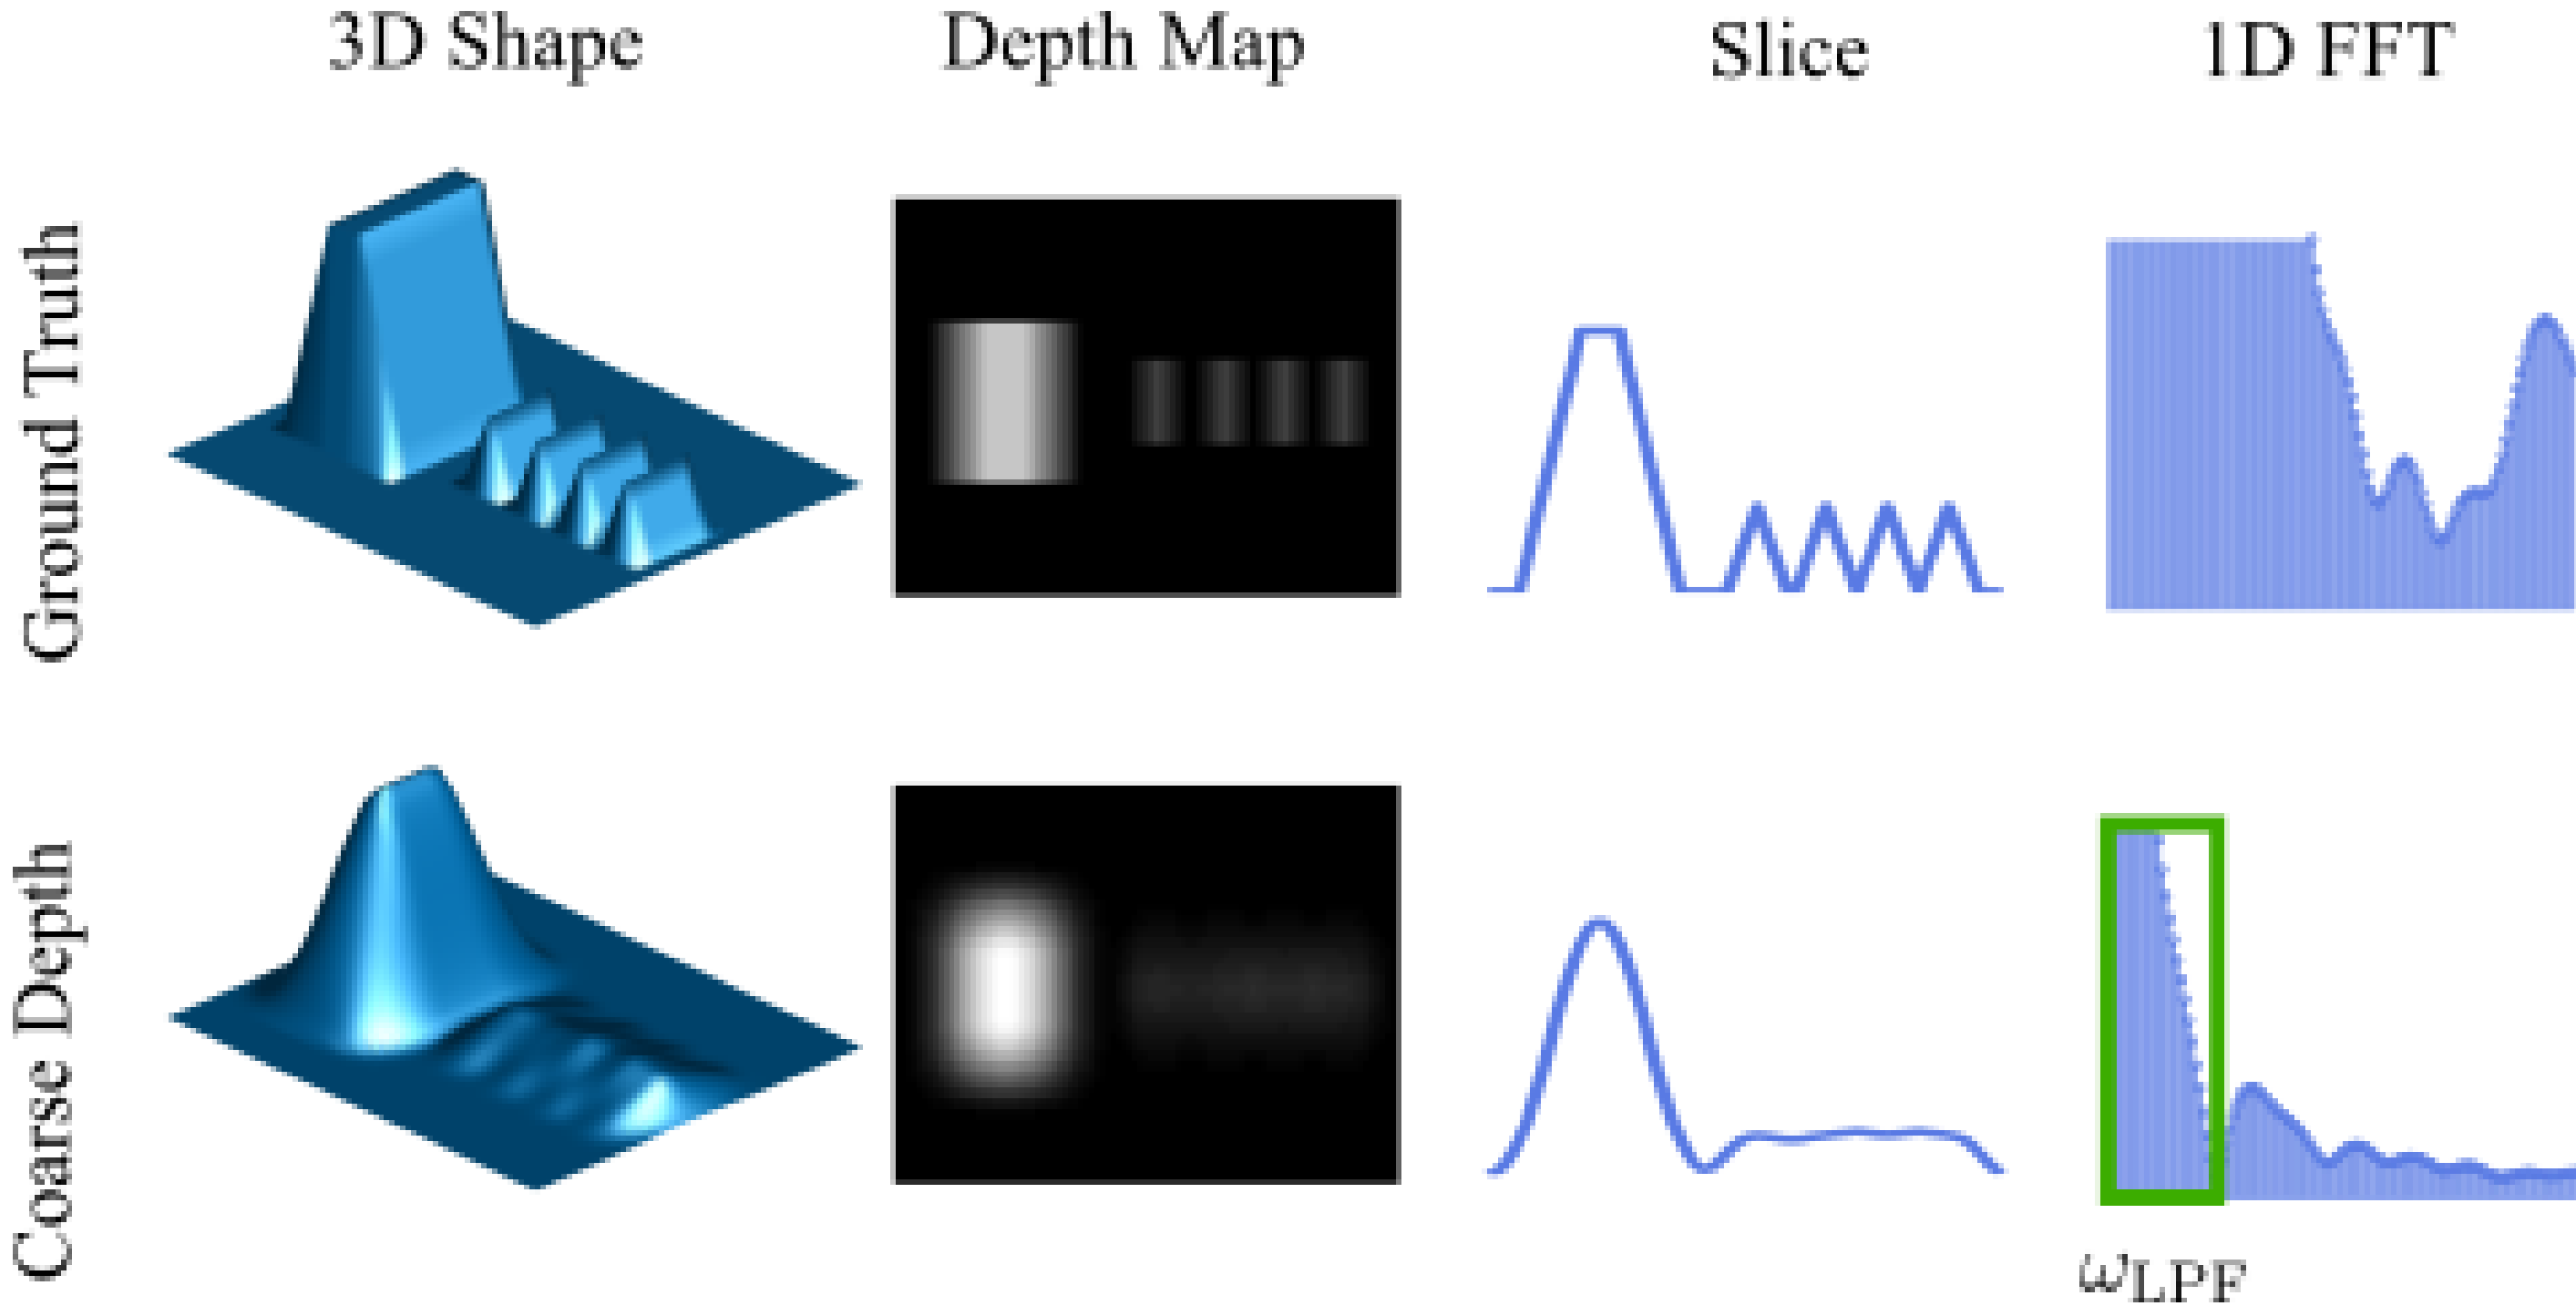
Frequency Analysis



Frequency Analysis

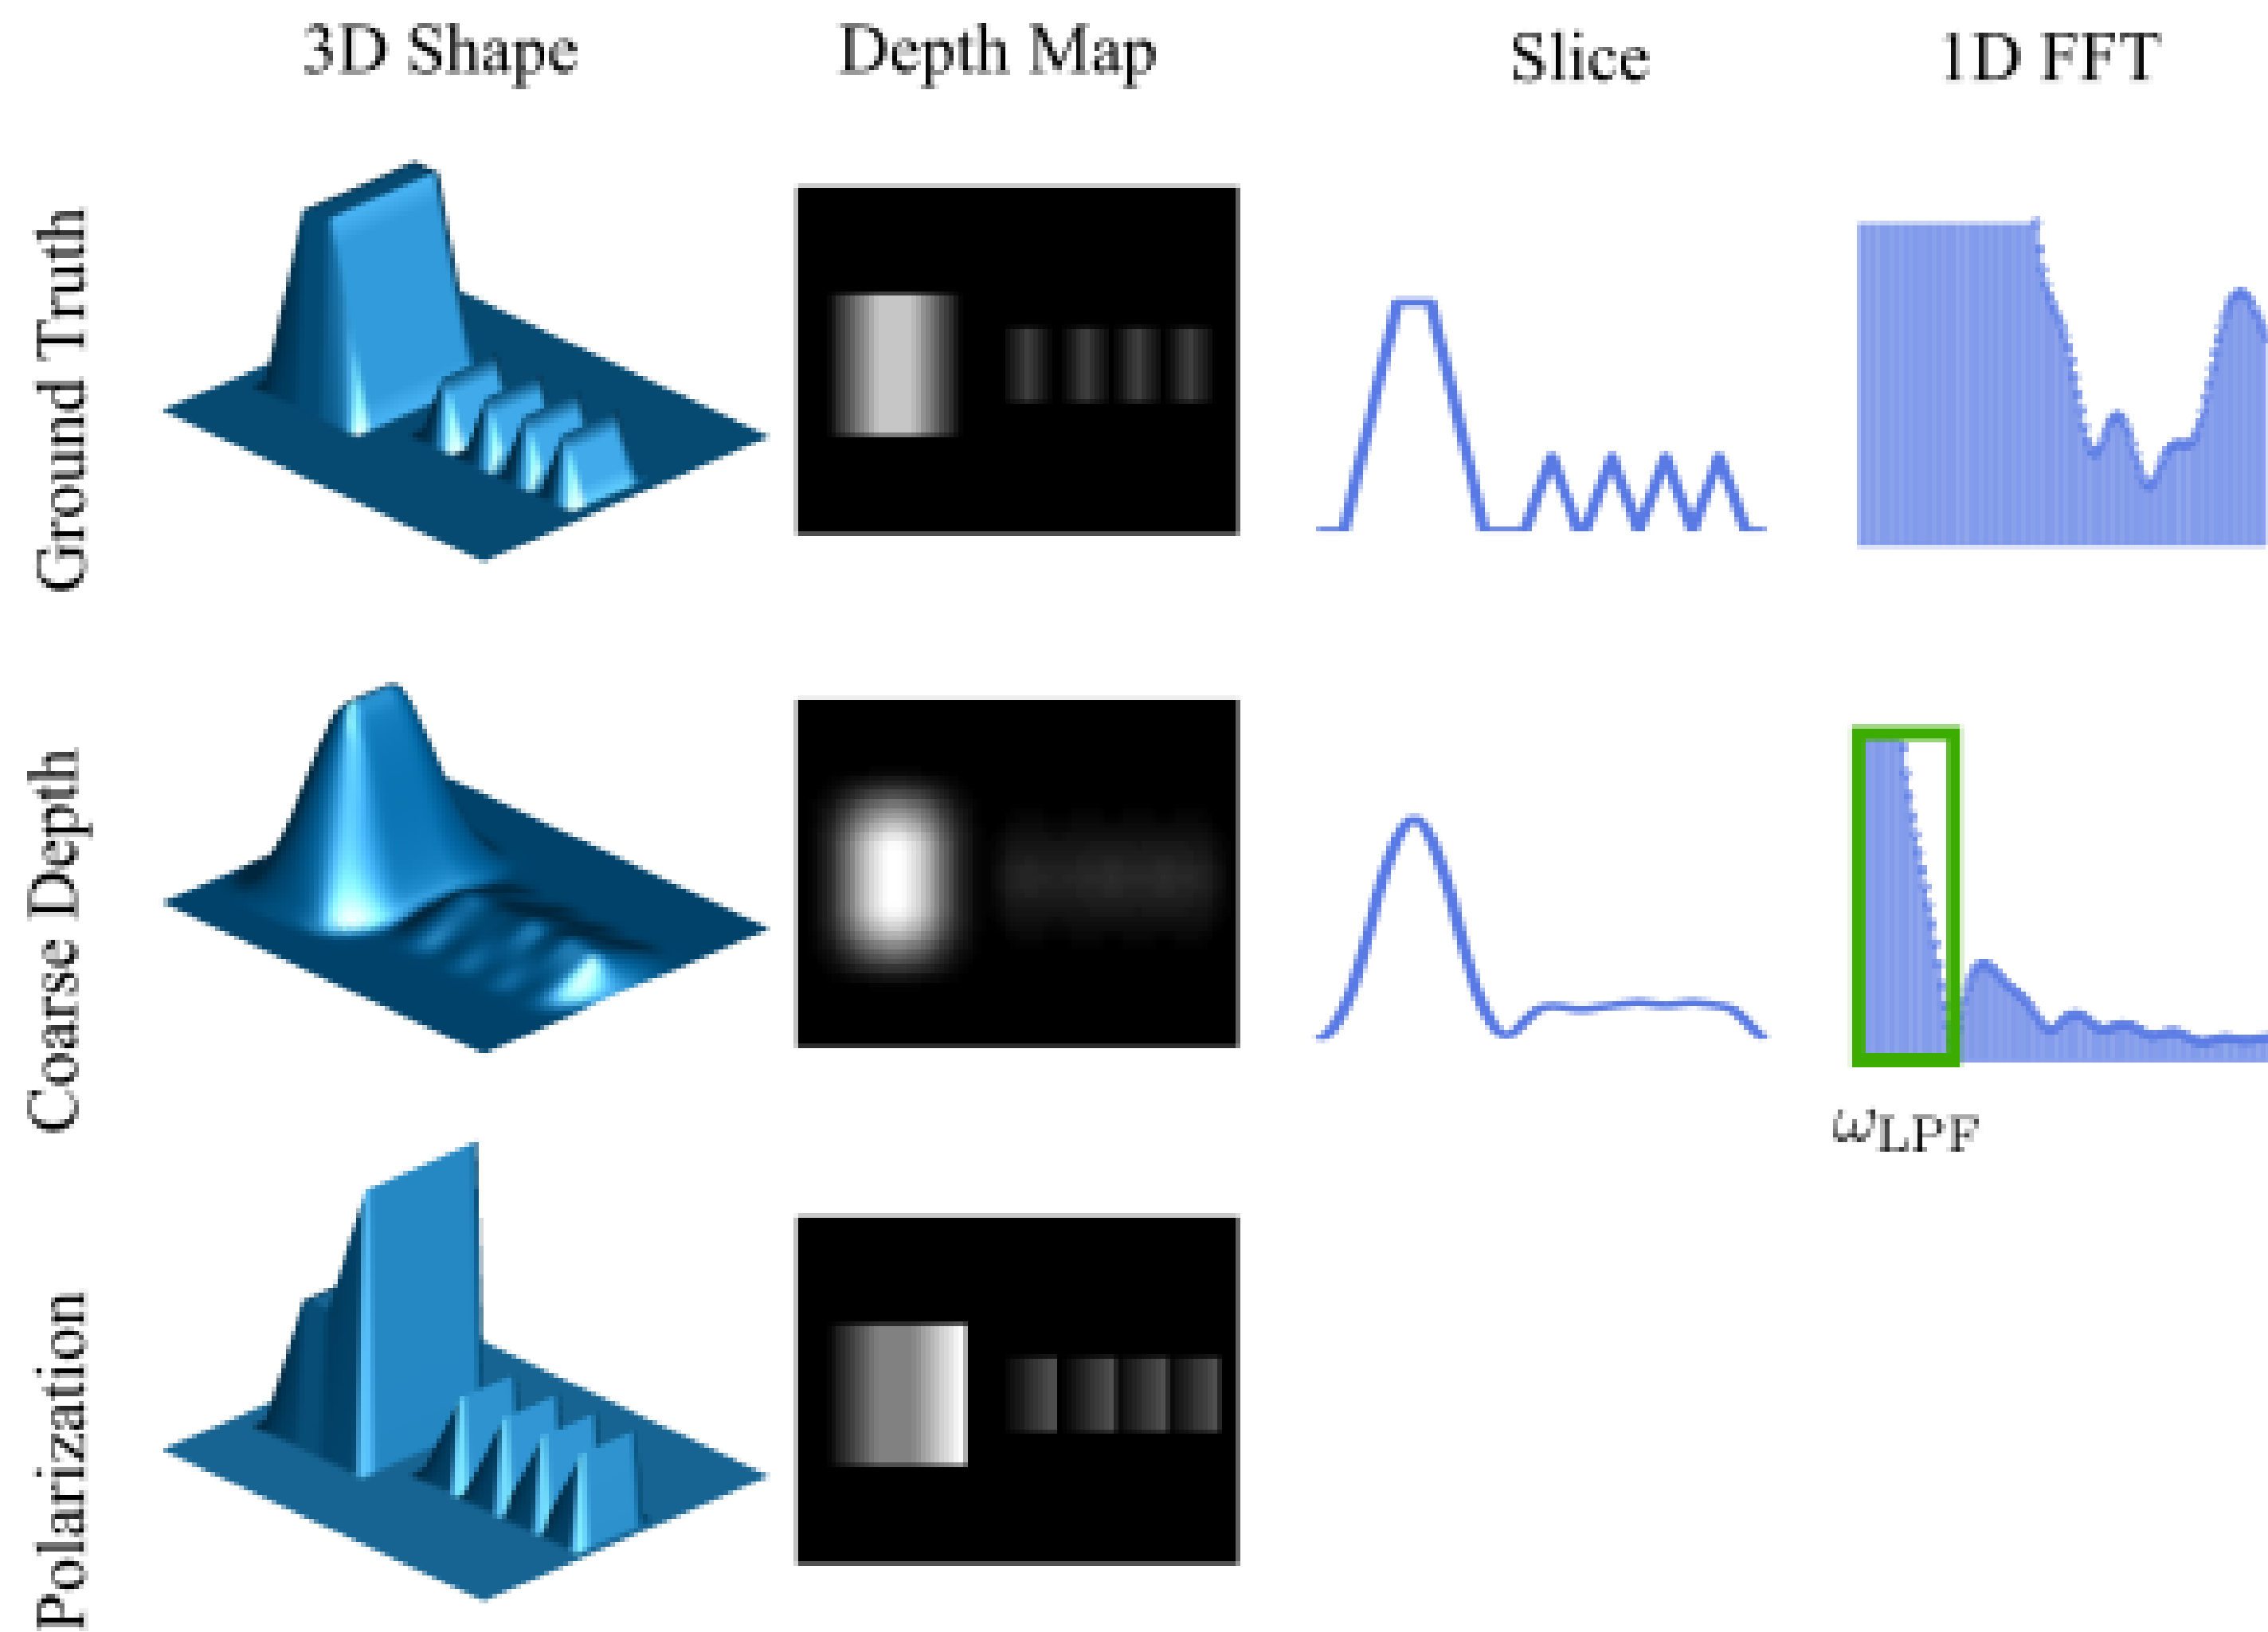


Frequency Analysis

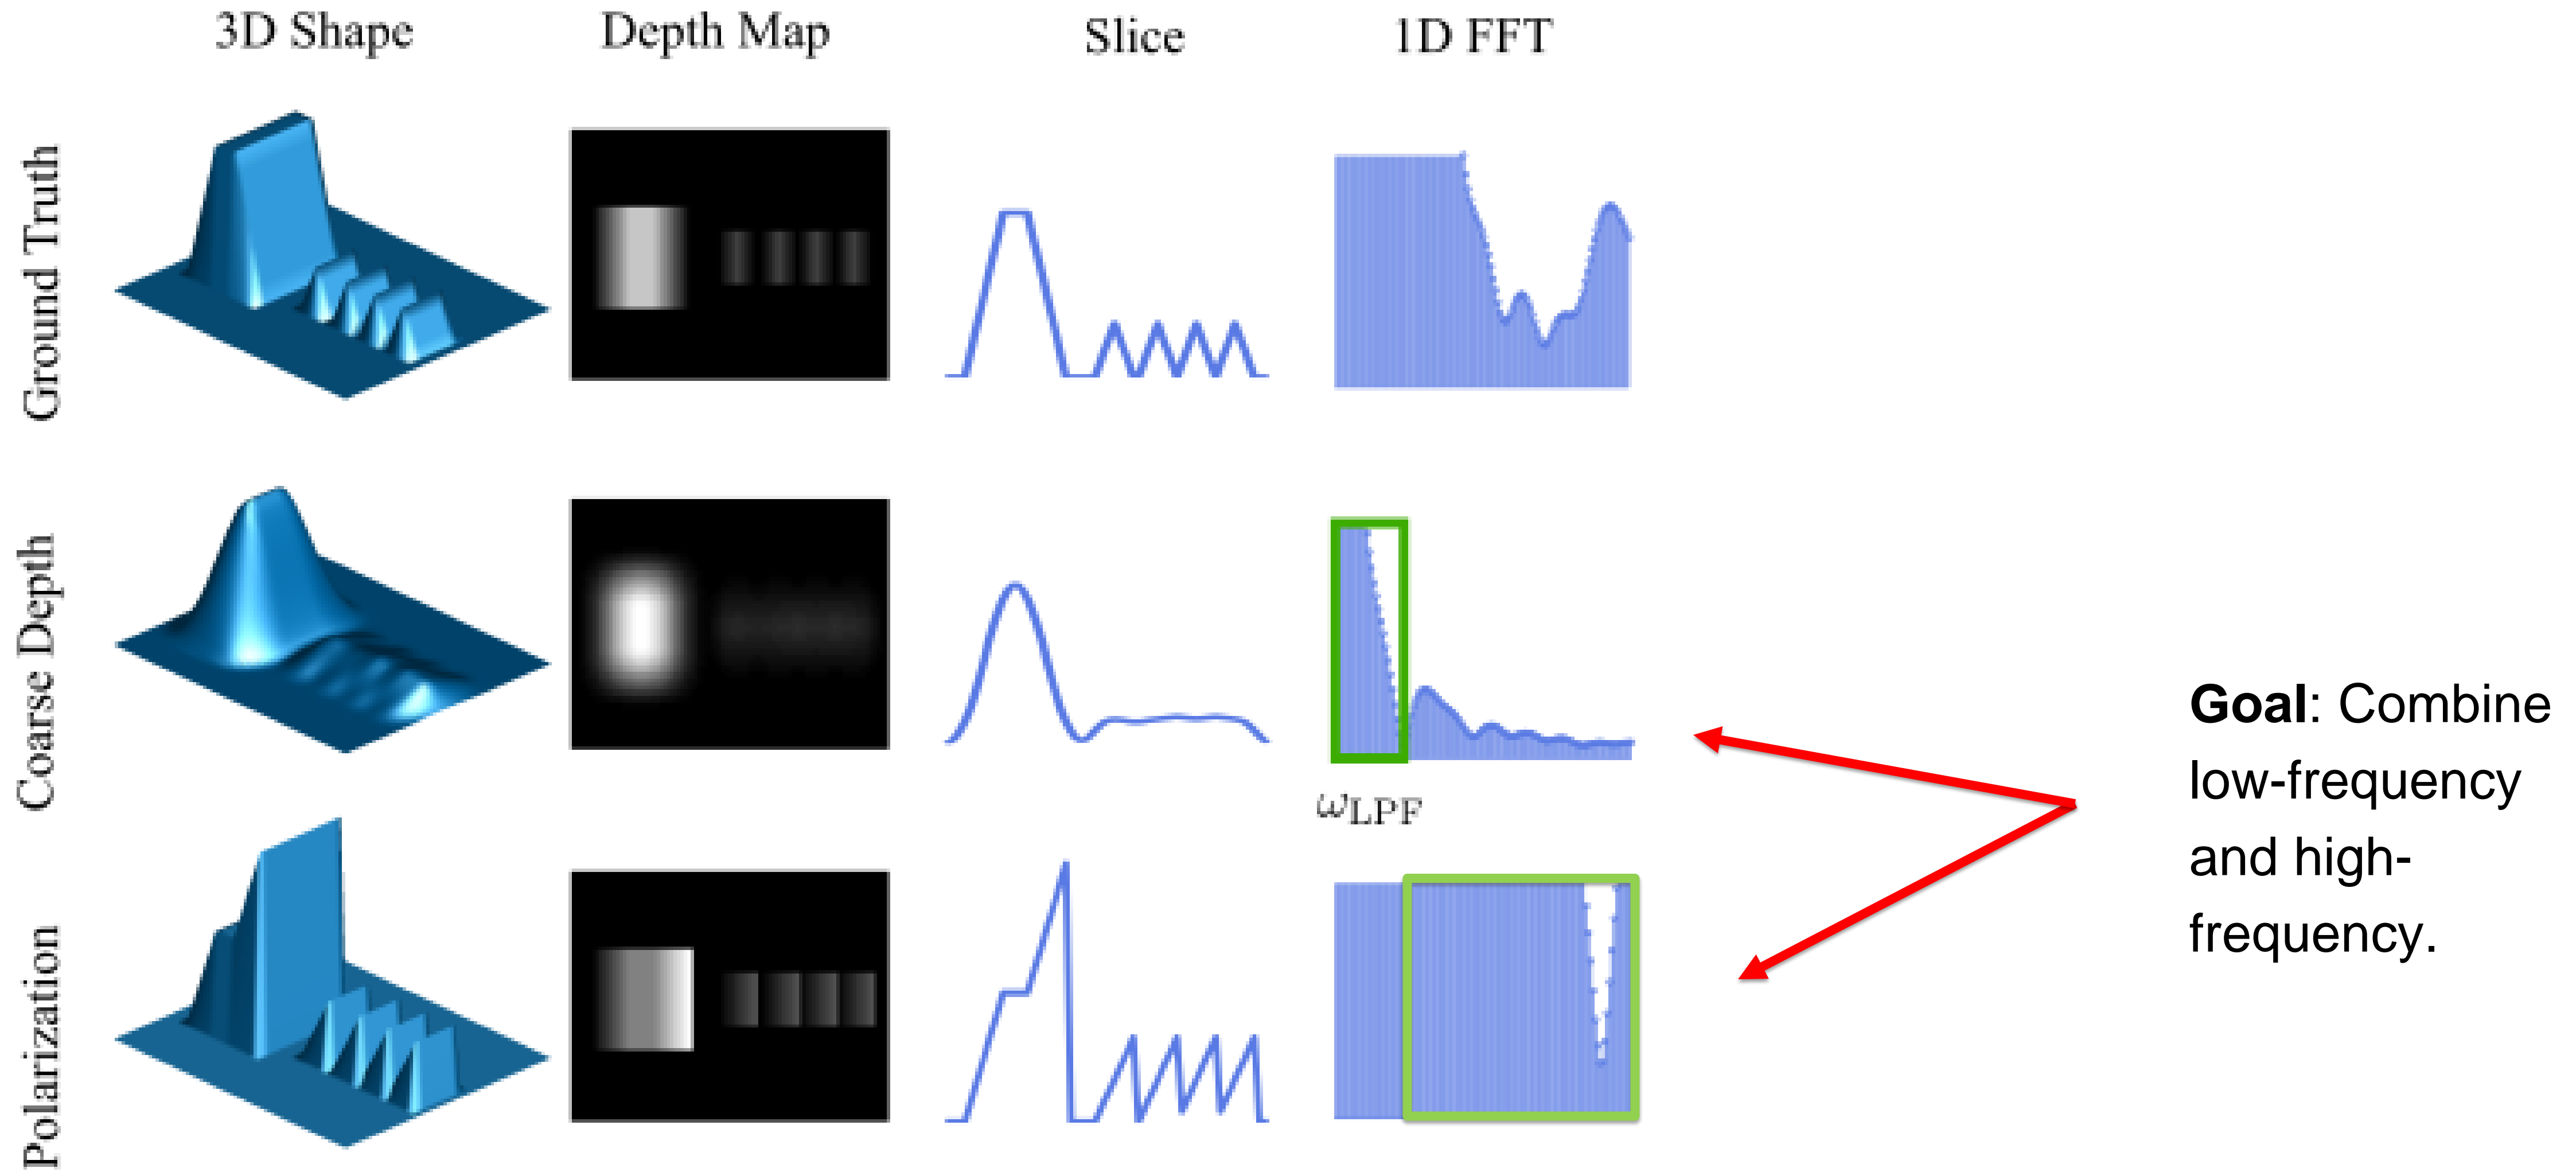


ω_{LPF}


Frequency Analysis



Frequency Analysis



Understanding our Input Data

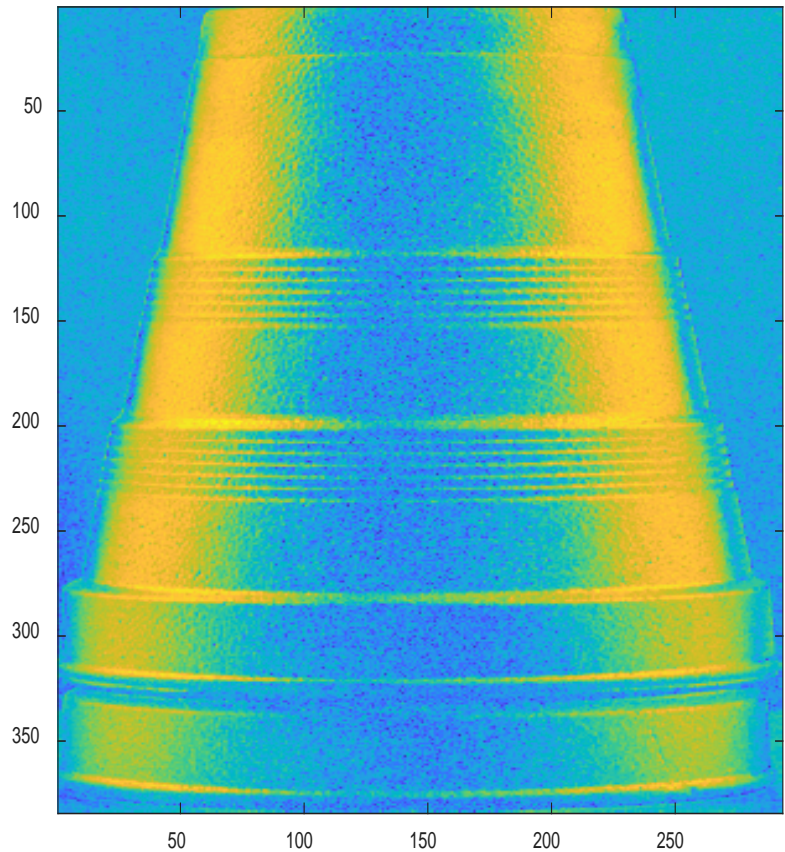


Coarse Depth Estimate




Obtained from Any
Depth Estimator

ϕ




Zenith Angle




Estimated from
Fresnel Equations

θ




Azimuth Angle





Gradient Domain Correction (Azimuth Only)

Azimuth Angle from Fresnel Eq.




Azimuth Angle of Deriv. Depth



Differentiate



Constrained Azimuth




Gradient Domain Correction

$$\hat{\mathcal{A}} = \arg \min_{\mathcal{A}} \left\| \mathbf{N}^{\text{depth}} - \mathcal{A} \left(\mathbf{N}^{\text{polar}} \right) \right\|_2^2 + \gamma \mathcal{I}_2 (\mathbf{A})$$



subject to $\mathbf{A} \in \{0, 1\}$,

Gradient Domain Correction (Azimuth Only)

Azimuth Angle from Fresnel Eq.



Azimuth Angle of Deriv. Depth




Differentiate



Gradient Domain Correction


Constrained Azimuth




$$\hat{\mathcal{A}} = \arg \min_{\mathcal{A}} \left\| \mathbf{N}^{\text{depth}} - \mathcal{A} \left(\mathbf{N}^{\text{polar}} \right) \right\|_2^2 + \gamma \mathcal{F}_2(\mathbf{A})$$

subject to $\mathbf{A} \in \{0, 1\}$,

Integrate the Surface (Naïve)



Corrected
Gradients




Integration
Algorithm



Output

Regularized Integrator



Corrected
Gradients




Integration
Algorithm

Output



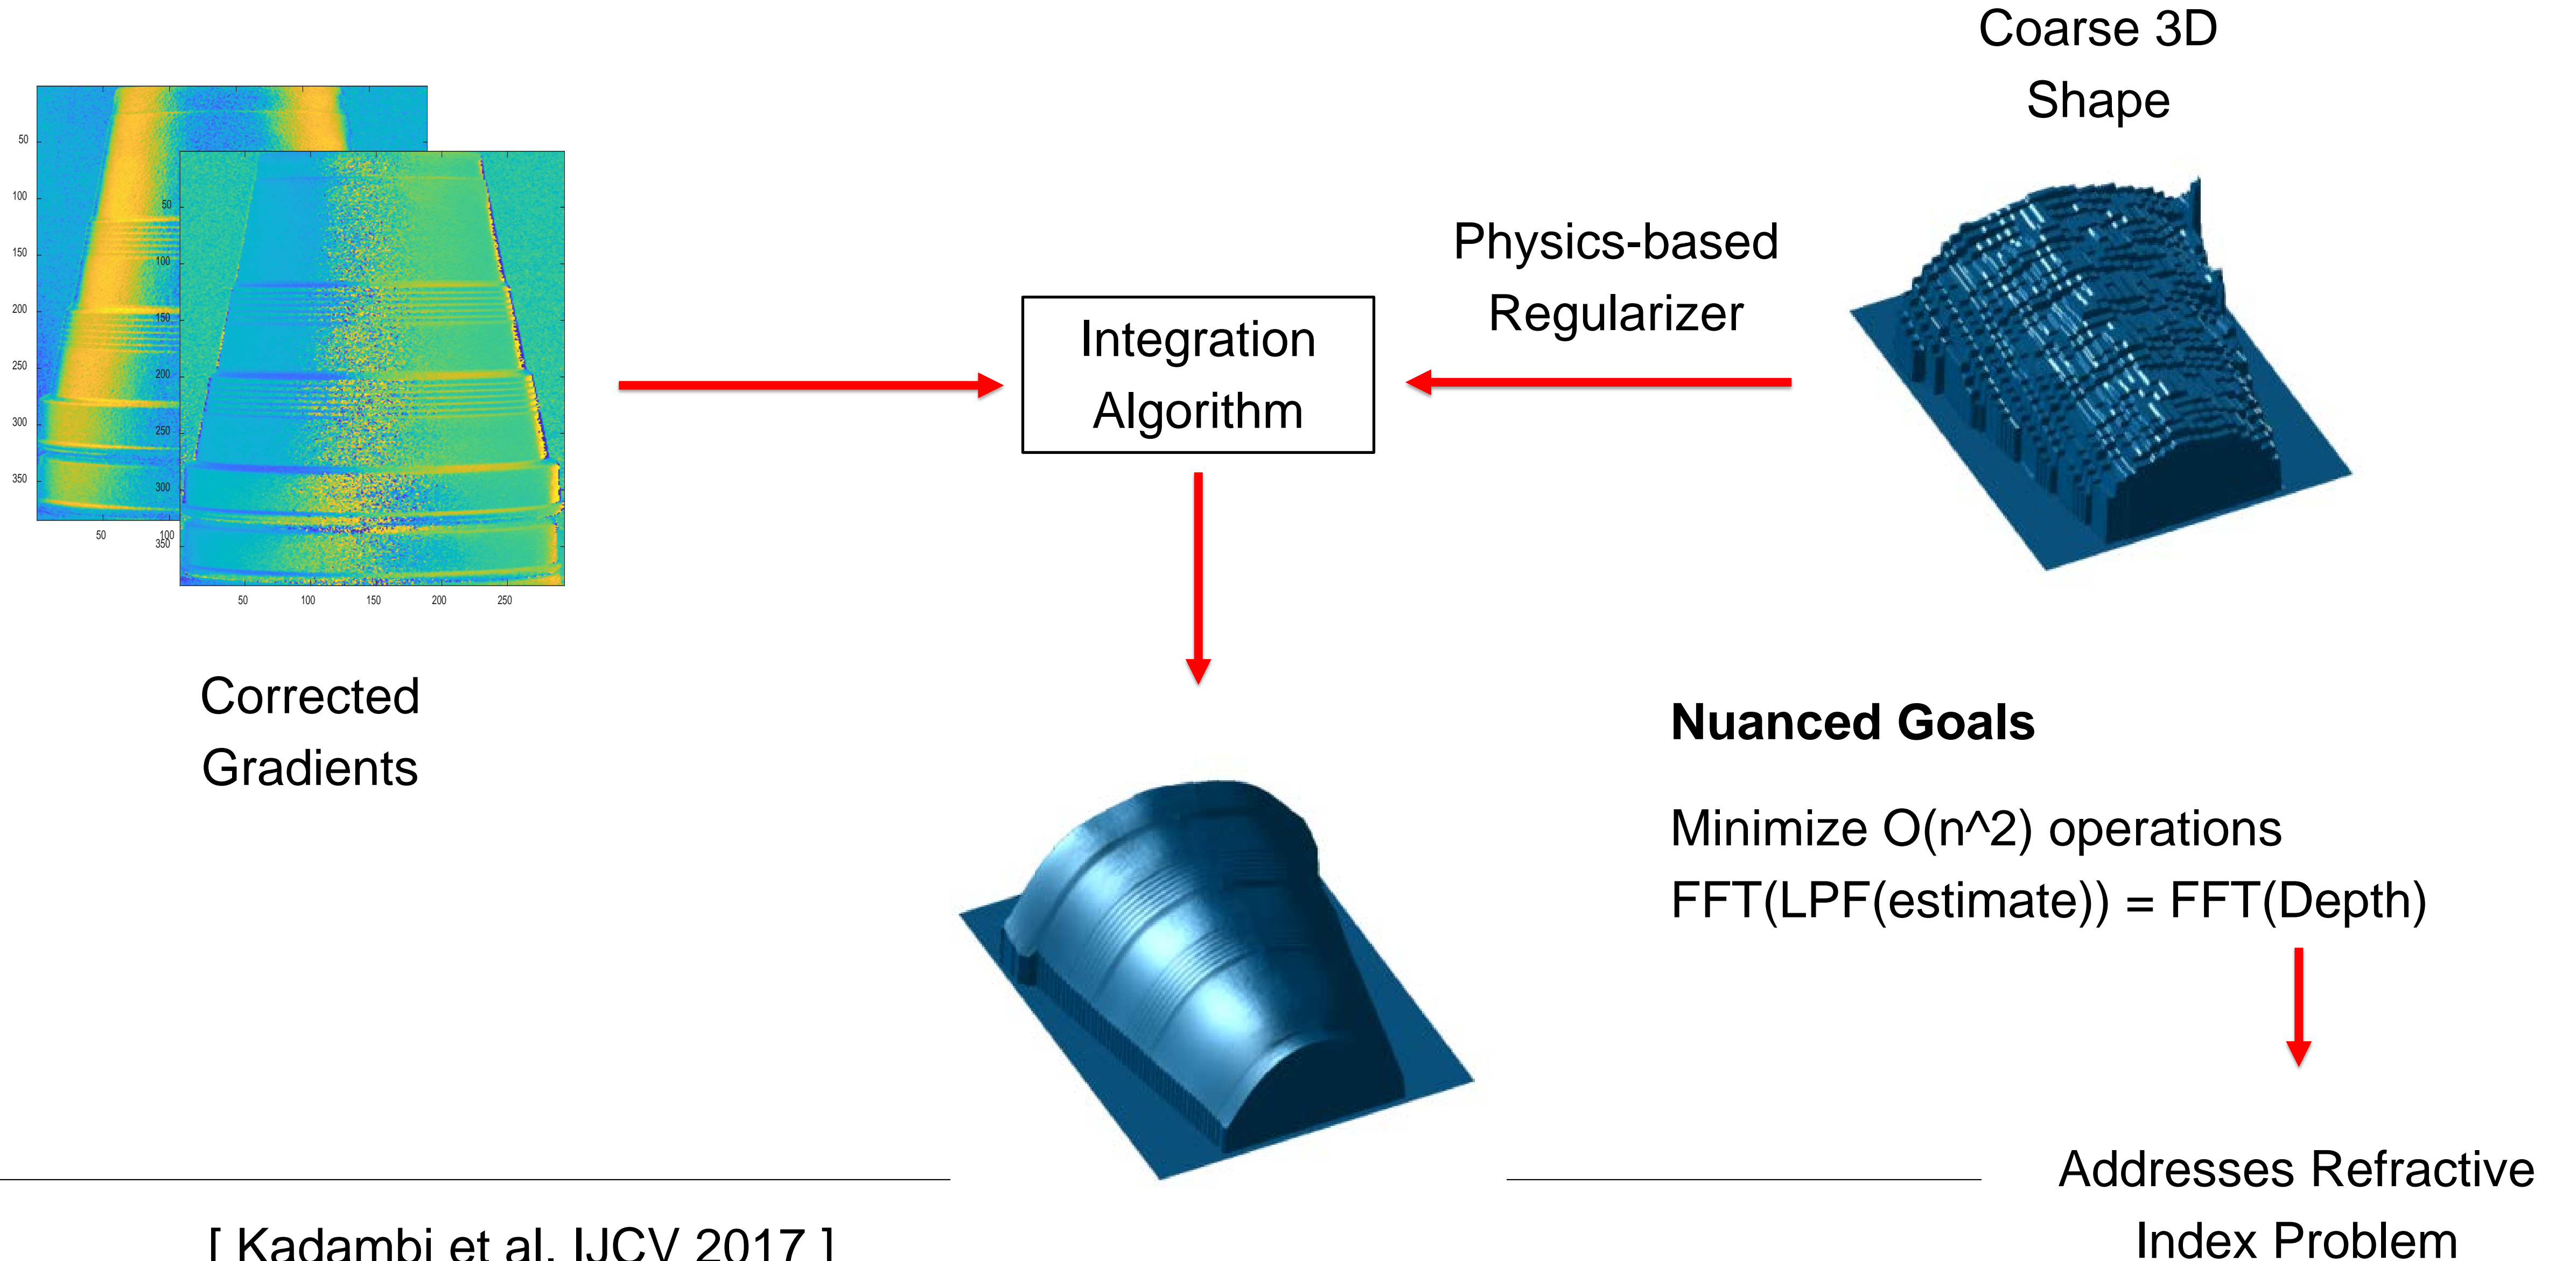
Regularizer

Coarse 3D
Shape



Naïve Regularizer Strategy: Penalize deviations in the surface normal integration with the coarse depth map

Regularized Integrator

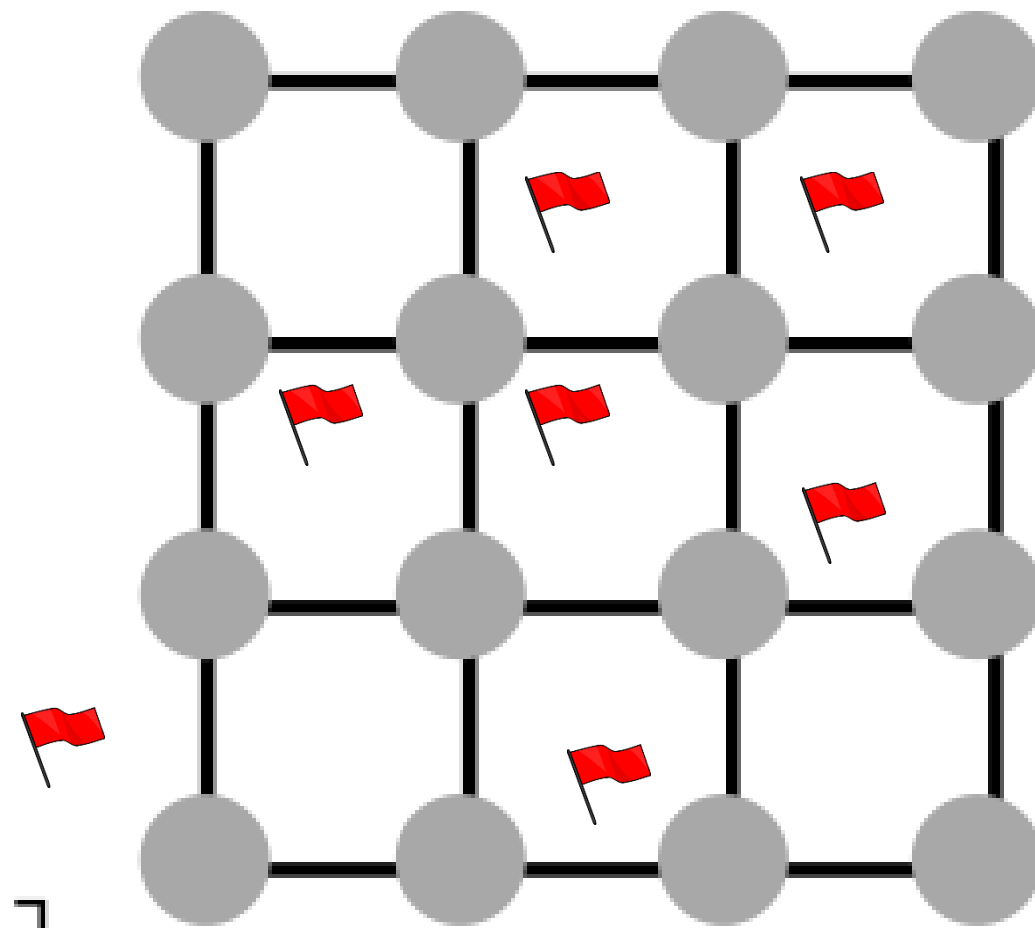


Integration Algorithm Posed as Graph Problem


Spanning Tree Integration

Kruskal's Algorithm: $O(E \log E)$


$$\begin{bmatrix} \lambda \mathbf{M} \odot \mathbf{I} \\ \nabla_S^2 \end{bmatrix} \text{VEC}(\hat{\mathbf{D}}) = \begin{bmatrix} \lambda \text{VEC}(\mathbf{M} \odot \mathbf{D}) \\ \nabla_S^T(\mathbf{N}^{\text{corr}}) \end{bmatrix}$$




Full Gradients




Minimum Spanning Tree



Ground Truth Surface




Full Integration



Spanning Tree Integration

Gradients can be flagged if the polarization signal is too small (e.g. DOLP)

Visual Debugging




Kinect




3 Polar Photos


Visual Debugging



Kinect




3 Polar Photos




Fresnel


Visual Debugging




Kinect



3 Polar Photos




Fresnel




Grad. Corr.


Visual Debugging




Kinect




3 Polar Photos



Fresnel




Grad. Corr.



Graph Integ.

Alternate Approaches



Microsoft Kinect ver. 2






Image-based Machine Learning [Laina16]




Shading Refinement [Wu14]



Industry Laser Scanner

Trained on 3D Database from NYU [Silberman12]




Polarization Approach


Physics, Hardware, and Computation

Hardware


[IEEE'17]



[CVPR'16]



[ICCP'14]



[ToG'13]




**Application: Relight
Consumer Photos**

[IJCV'17]



[ICCV'15]




[Opt Lett'14]



[CVPR'15]



[ToG'16]




Physics

Computation


Physics, Hardware, and Computation

Hardware


[IEEE'17]



[CVPR'16]



[ICCP'14]



[ToG'13]



**Application: Relight
Consumer Photos**

[IJCV'17]



[ICCV'15]




[Opt Lett'14]



[CVPR'15]



[ToG'16]

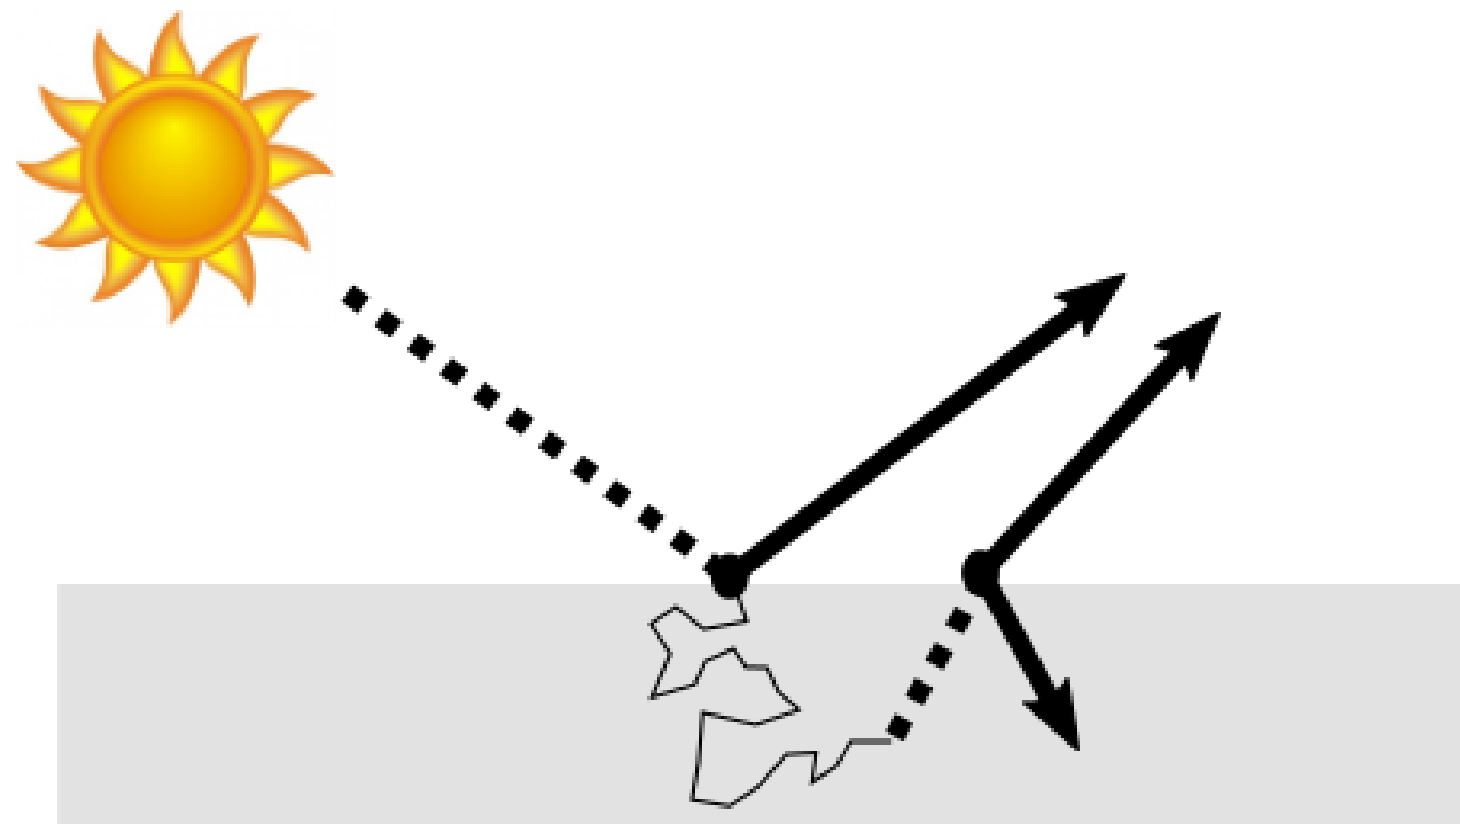


Physics

Computation

Provable Guarantees on Light Transport

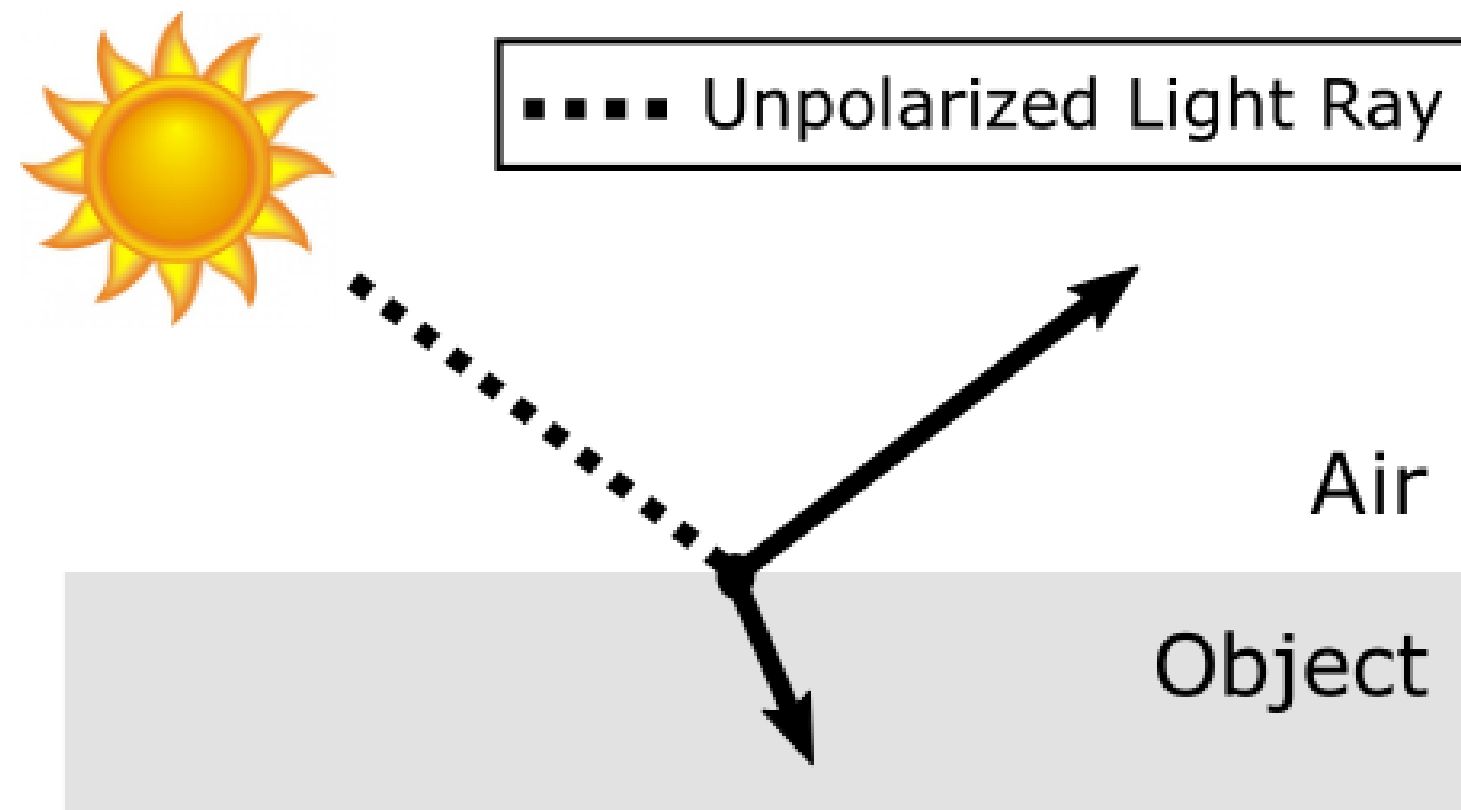
Proposition 1 *Stochastic ambiguities in the azimuthal angle are avoided when polarized reflections are either diffuse or specular dominant.*



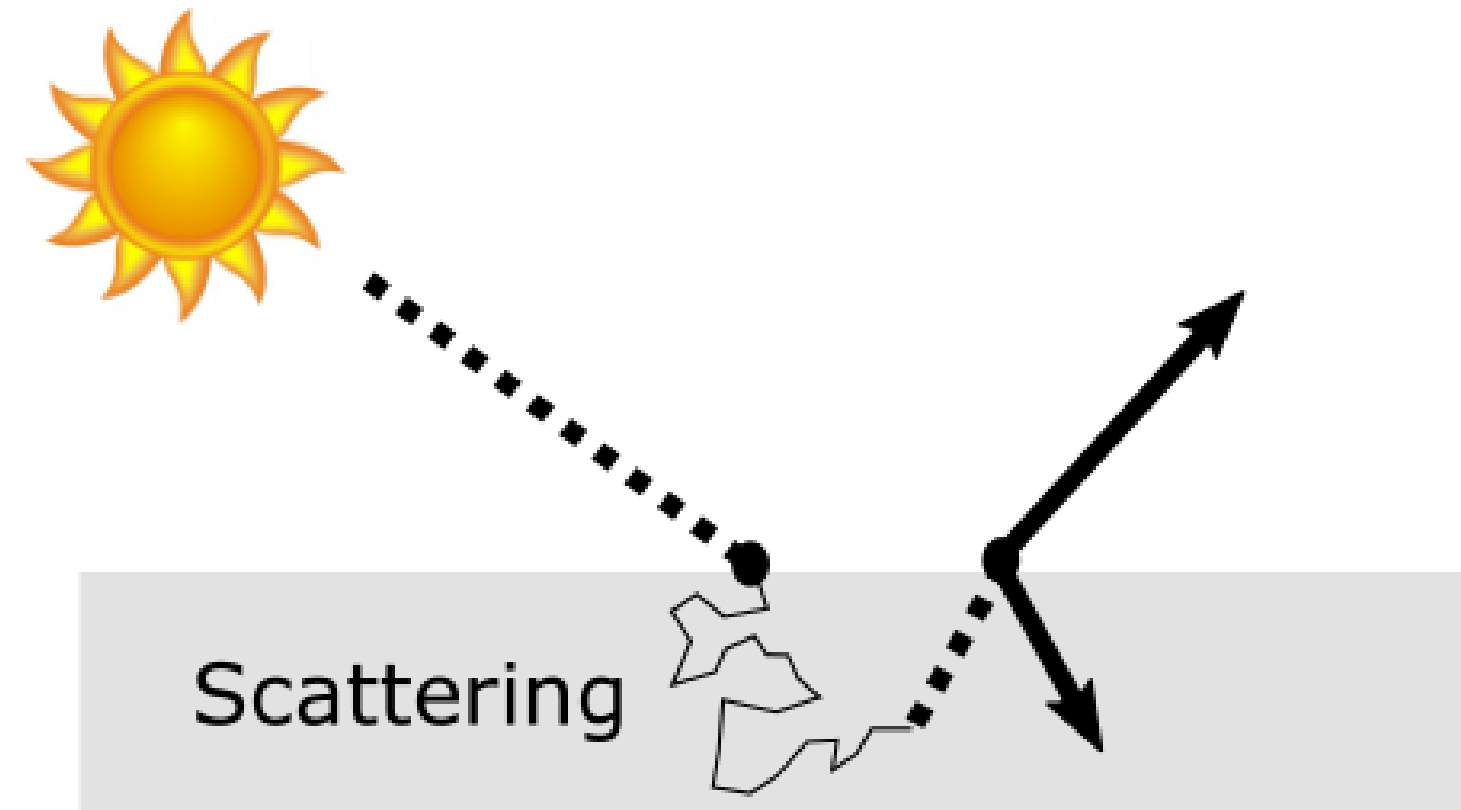
(c) Mixed Reflection

Proposition 2 *Assuming the conditions of Proposition 1 hold, perturbations in zenith angle due to mixed reflections can be corrected by applying the rotation operator $\hat{\mathcal{R}}$, as described in Eq. 10.*


Surface Assumption



(a) Specular Reflection



(b) Diffuse Reflection



(c) Mixed Reflection

Fresnel Equations
[Fresnel 1819]

$$\rho^{\text{spec}} = \frac{2n \tan \theta \sin \theta}{\tan^2 \theta \sin^2 \theta + |n^*|^2}$$

OR

$$\rho = \frac{(n - \frac{1}{n})^2 \sin^2 \theta}{2 + 2n^2 - (n + \frac{1}{n})^2 \sin^2 \theta + 4 \cos \theta \sqrt{n^2 - \sin^2 \theta}}$$

“Reversed Fresnel”
[Atkinson et al. 2006]

“Mixed Fresnel”
[Kadambi et al. 2017]

Contribution:
analyze a
mixture model


One of the most challenging objects to scan



Paints of different colors

Uncontrolled Ambient light
from window



Kinect





Kinect



Shading [Wu 14]



Polarized 3D




Passive Polarimetric Imaging


Varying Reflection Class




Passive Polarimetric Imaging



Kinect



Shading [Wu 14]







Polarized 3D

Varying Reflection Class




Passive Polarimetric Imaging




Many lingering challenges!

Limitations

(-) Requires Multiple Images



(-) Requires Coarse Depth



(-) Not real-time
(3 sec/frame)



Solutions

Multipolar
Camera

Depth Priors
from Machine
Learning

GPU or
physics-based
learning

Note: we have only scratched surface of just *linear* polarization

Birefringence, circular polarization, etc.

Future pathways to probe



**Fine structures for
Bioengineering**




Smartphone 3D Vision



Long-range 3D sensing

Physics, Hardware, and Computation


Hardware



[CVPR'16]



[ICCP'14]



[ToG'13]




[IJCV'17]



[ICCV'15]




[Opt Lett'14]



[CVPR'15]



[ToG'16]




Physics

Computation


Physics, Hardware, and Computation

Hardware


[IEEE'17]



[CVPR'16]



[ICCP'14]



[ToG'13]



Goal: Separate Light Bounces

[IJCV'17]



[ICCV'15]




[Opt Lett'14]



[CVPR'15]




[ToG'16]



Physics

Computation



car

car

person


bicycle

bicycle

bicycle

MIT Tech Review

Separating Multiple Bounces of Light




This pixel measures both blue and red paths. Can we split the light?

Ill-posed Problem

$$I = I_A + I_B$$

So the problem needs to be constrained


Previous Work: Smoothness in Space



Smoothness a Limiting Assumption

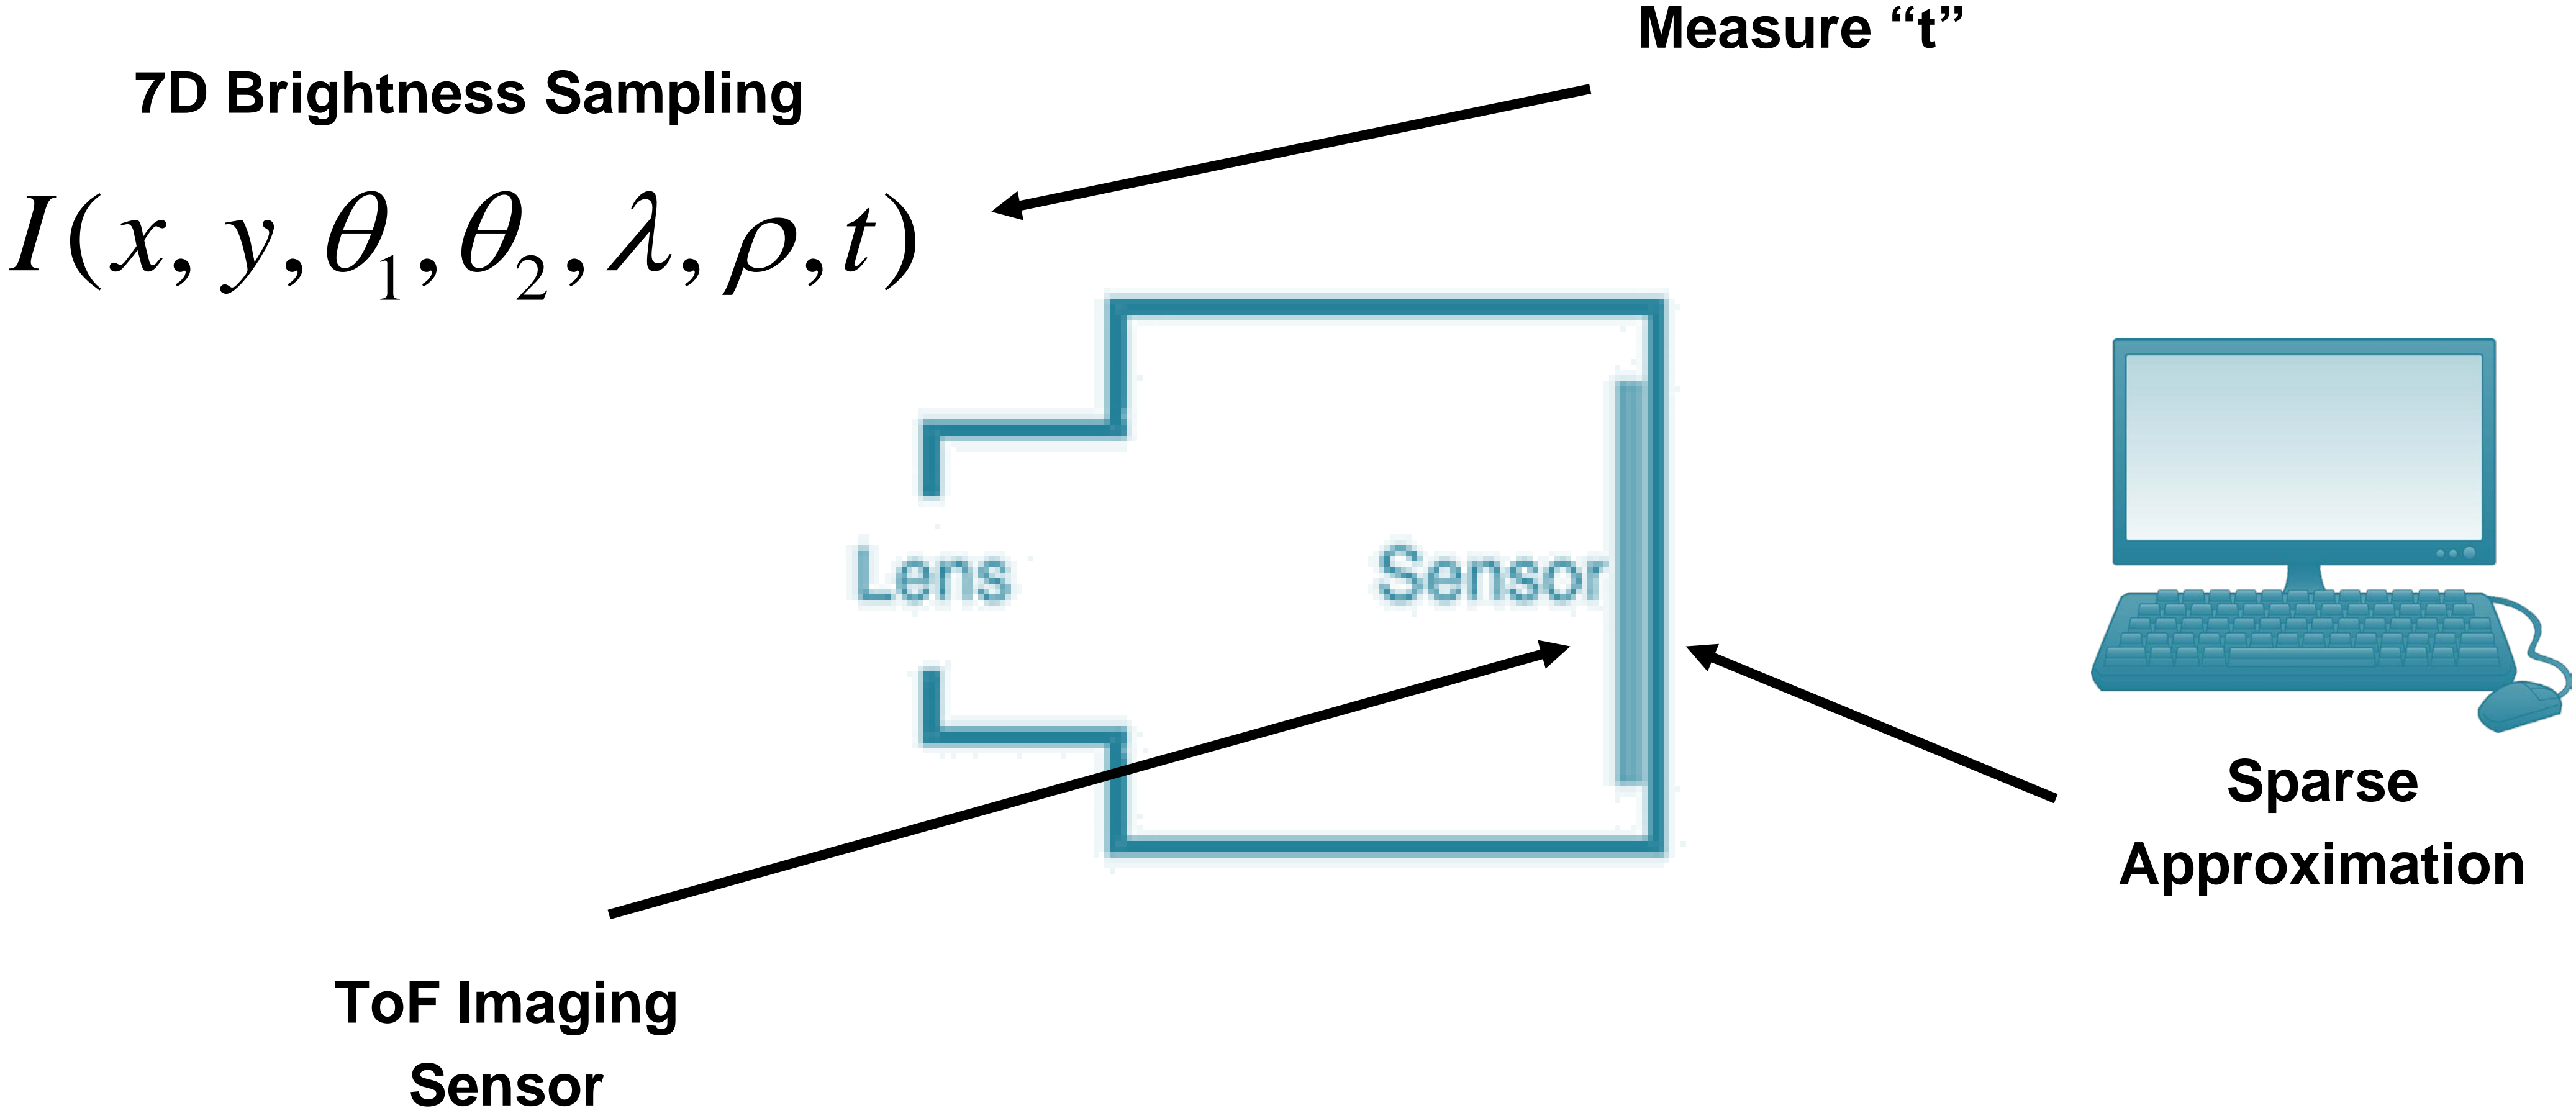
"If a laser beam hits a water drop... the lidar can think it's an object and slam on the brakes"

- Director of CMU autonomous driving lab



Multipath Interference for CW emitters

Generalizing the Camera to Systems Thinking



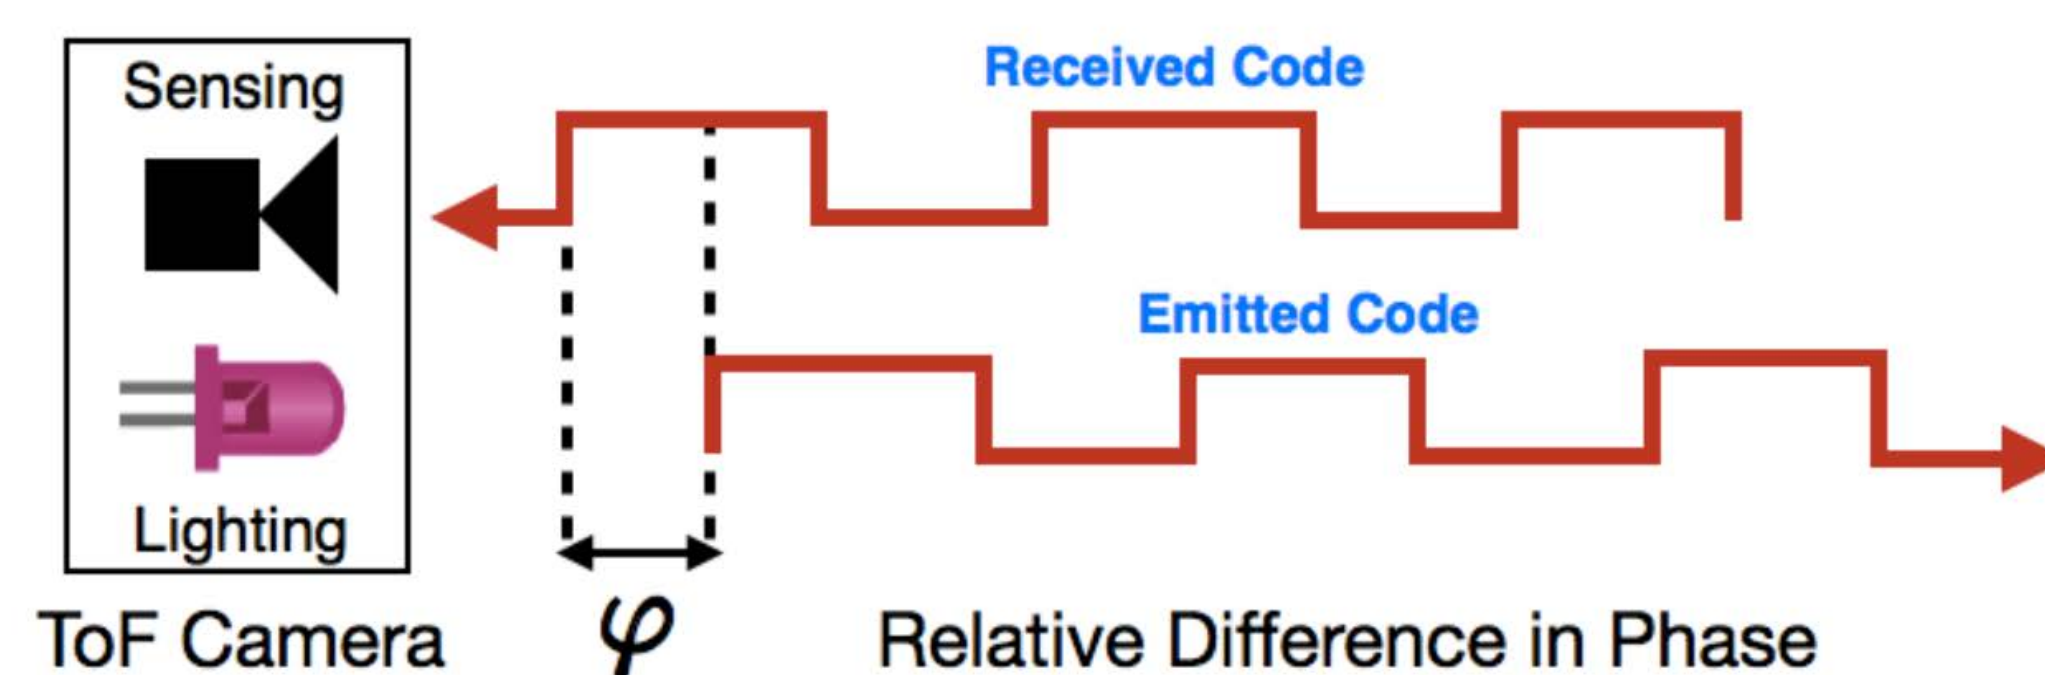
Overview of Time of Flight

$$d = vt \quad v = 3 * 10^8 \text{ m/s}$$



So, we need a camera to measure time delay.



Example: Microsoft Kinect (ver 2)




$$z = \frac{c\Phi}{2\pi f_\omega}, \quad c \approx 3 \times 10^8 \text{ m/s}$$



Recall: Strobing Pattern is MHz (nanosecond periods)


Interference




Measure one Phase that is in between (mod 2π)

Recall: Strobing Pattern is MHz (nanosecond periods)

Each Pixel Becomes a Linear Time-invariant System




Probing Kernel



Spike location due to optical path length


Temporal Response

=




Measured Signal

Each Pixel Becomes a Linear Time-invariant System




Probing Kernel

*




Spike location due to optical path length

=




Measured Signal




Probing Kernel

*




Temporal Response

=





Measured Signal

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?

Imaging device should be “spread spectrum” → term from telecommunications



Prior-based Orthogonal Matching Pursuit

We explore the matching pursuit class of problems which approximate the original l0 program:

$$\arg \min_{\vec{x}} \left\| \vec{y} - \mathbf{H}\vec{x} \right\|_2^2 + \lambda \left\| \vec{x} \right\|_0^2$$

We make two modifications here:

1. Non-negativity constraints (Bruckstein, Elad, and Zibulevsky, ISCCSP 2008).
2. Proximity constraints

Nonnegativity

- a) Consider only positive projections when searching for the next atom.
- b) When updating the residual use a solver to impose positivity on the coefficients.

Proximity Constraints

$$p(\vec{x}) \in \square(\vec{x}; \mu, \sigma)$$

$$\arg \max_{\vec{x}} p(\vec{x} | \vec{y}) \rightarrow \arg \max_{\vec{x}} \underbrace{p(\vec{y} | \vec{x})}_{\text{Likelihood}} \underbrace{p(\vec{x})}_{\text{Prior}}$$


Theory → Instantiation

Nanophotography: AM modulation with customizable signal encoding




Imaging at 8 Billion Effective FPS


Imaging at 8 Billion Effective FPS





Nanocamera




Imaging at 8 Billion Effective FPS



Nanocamera




Captured Data



[Kadambi et al. SIGGRAPH13]

Note non-smooth multipath (e.g. Specularities)

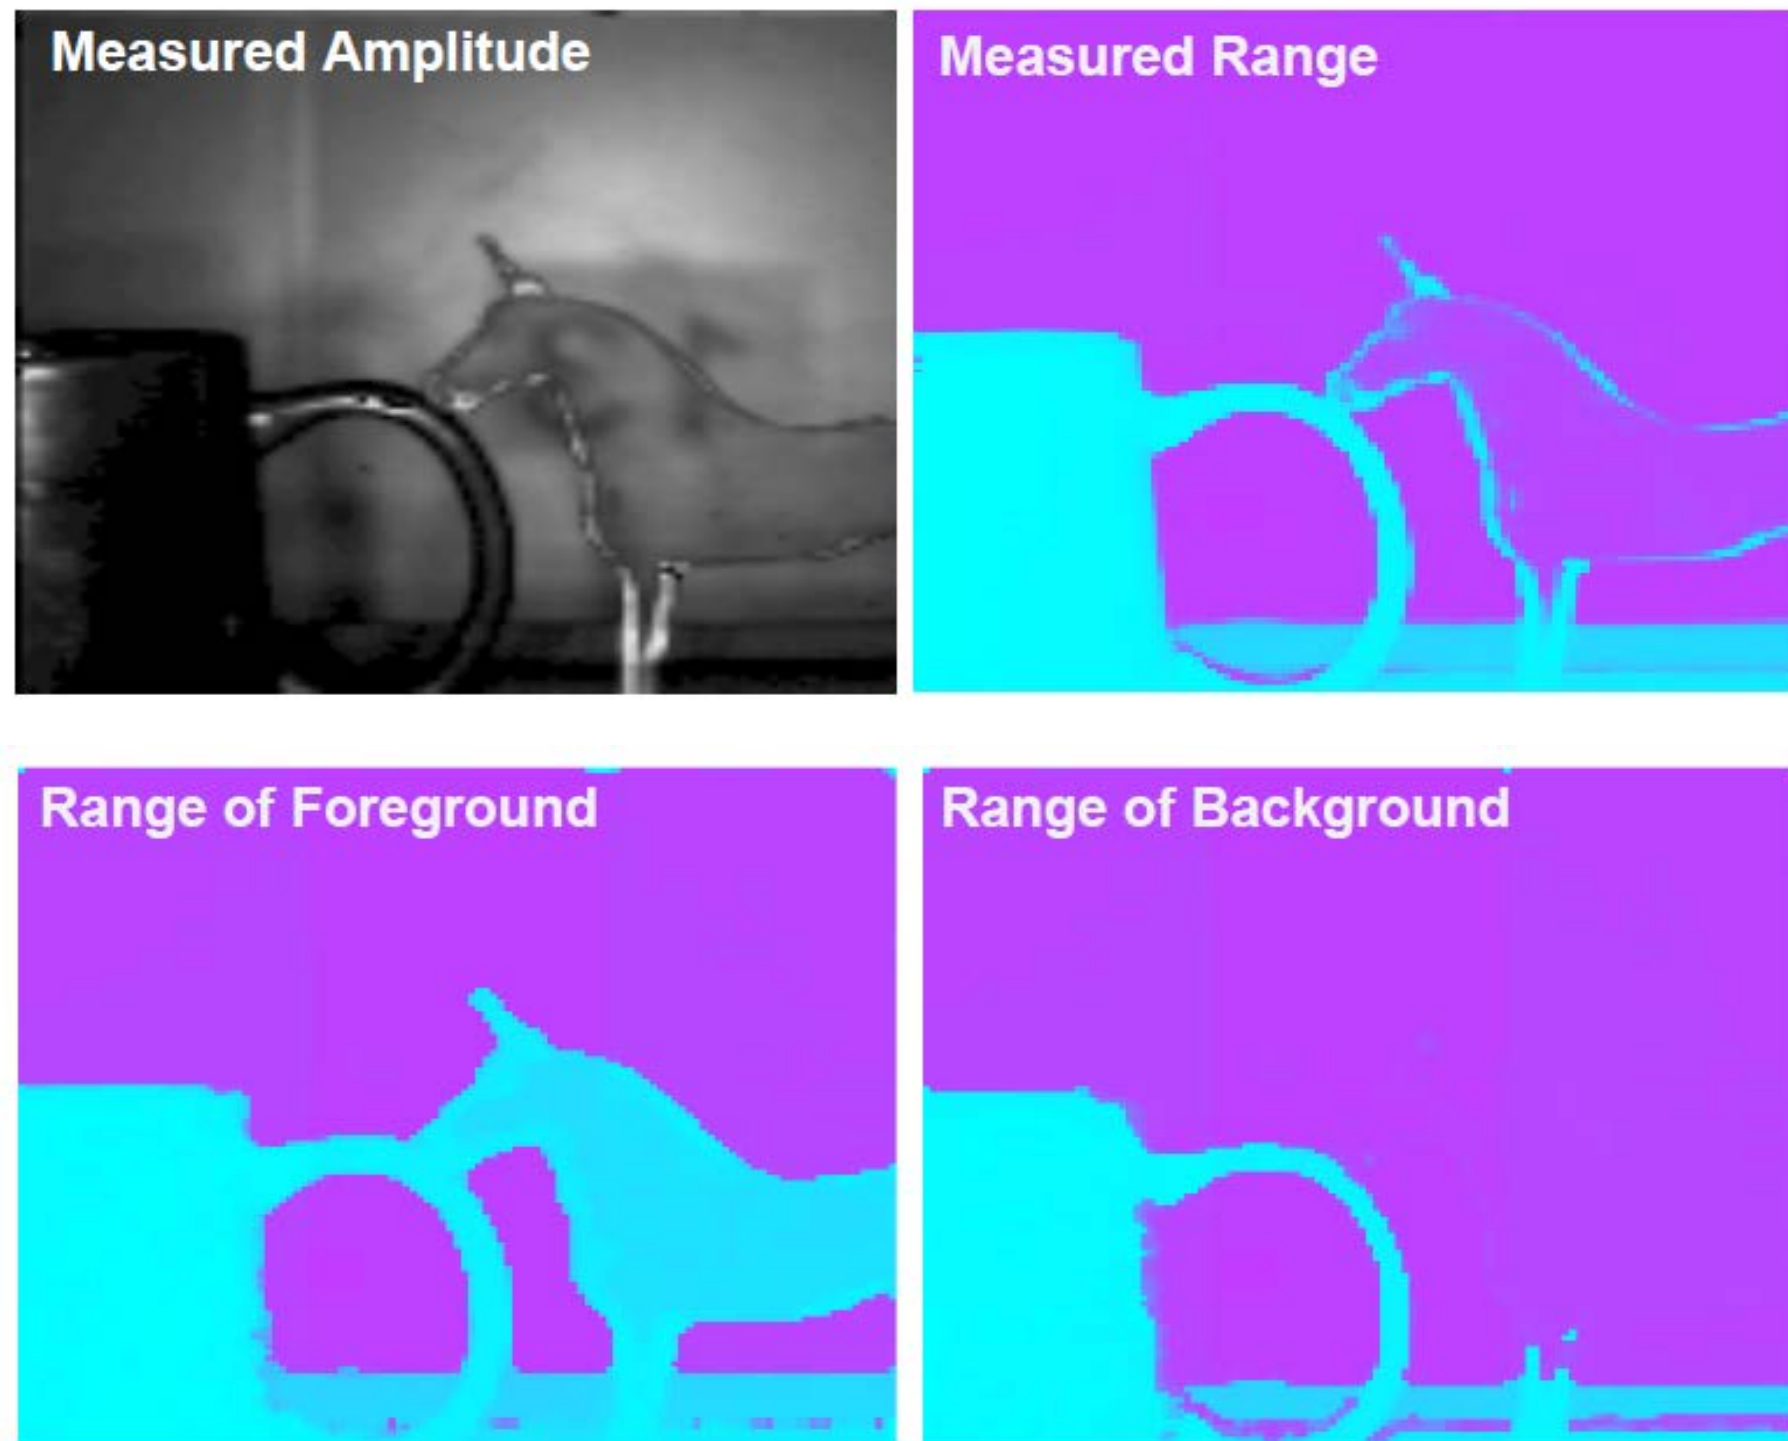
Tool Application: 3D Imaging in Multipath Environments




Contribution: Multi-bounce 3D imaging [Kadambi et al. SIGGRAPH Asia 2013]

Transient imaging tackling Edge Cases in Imaging

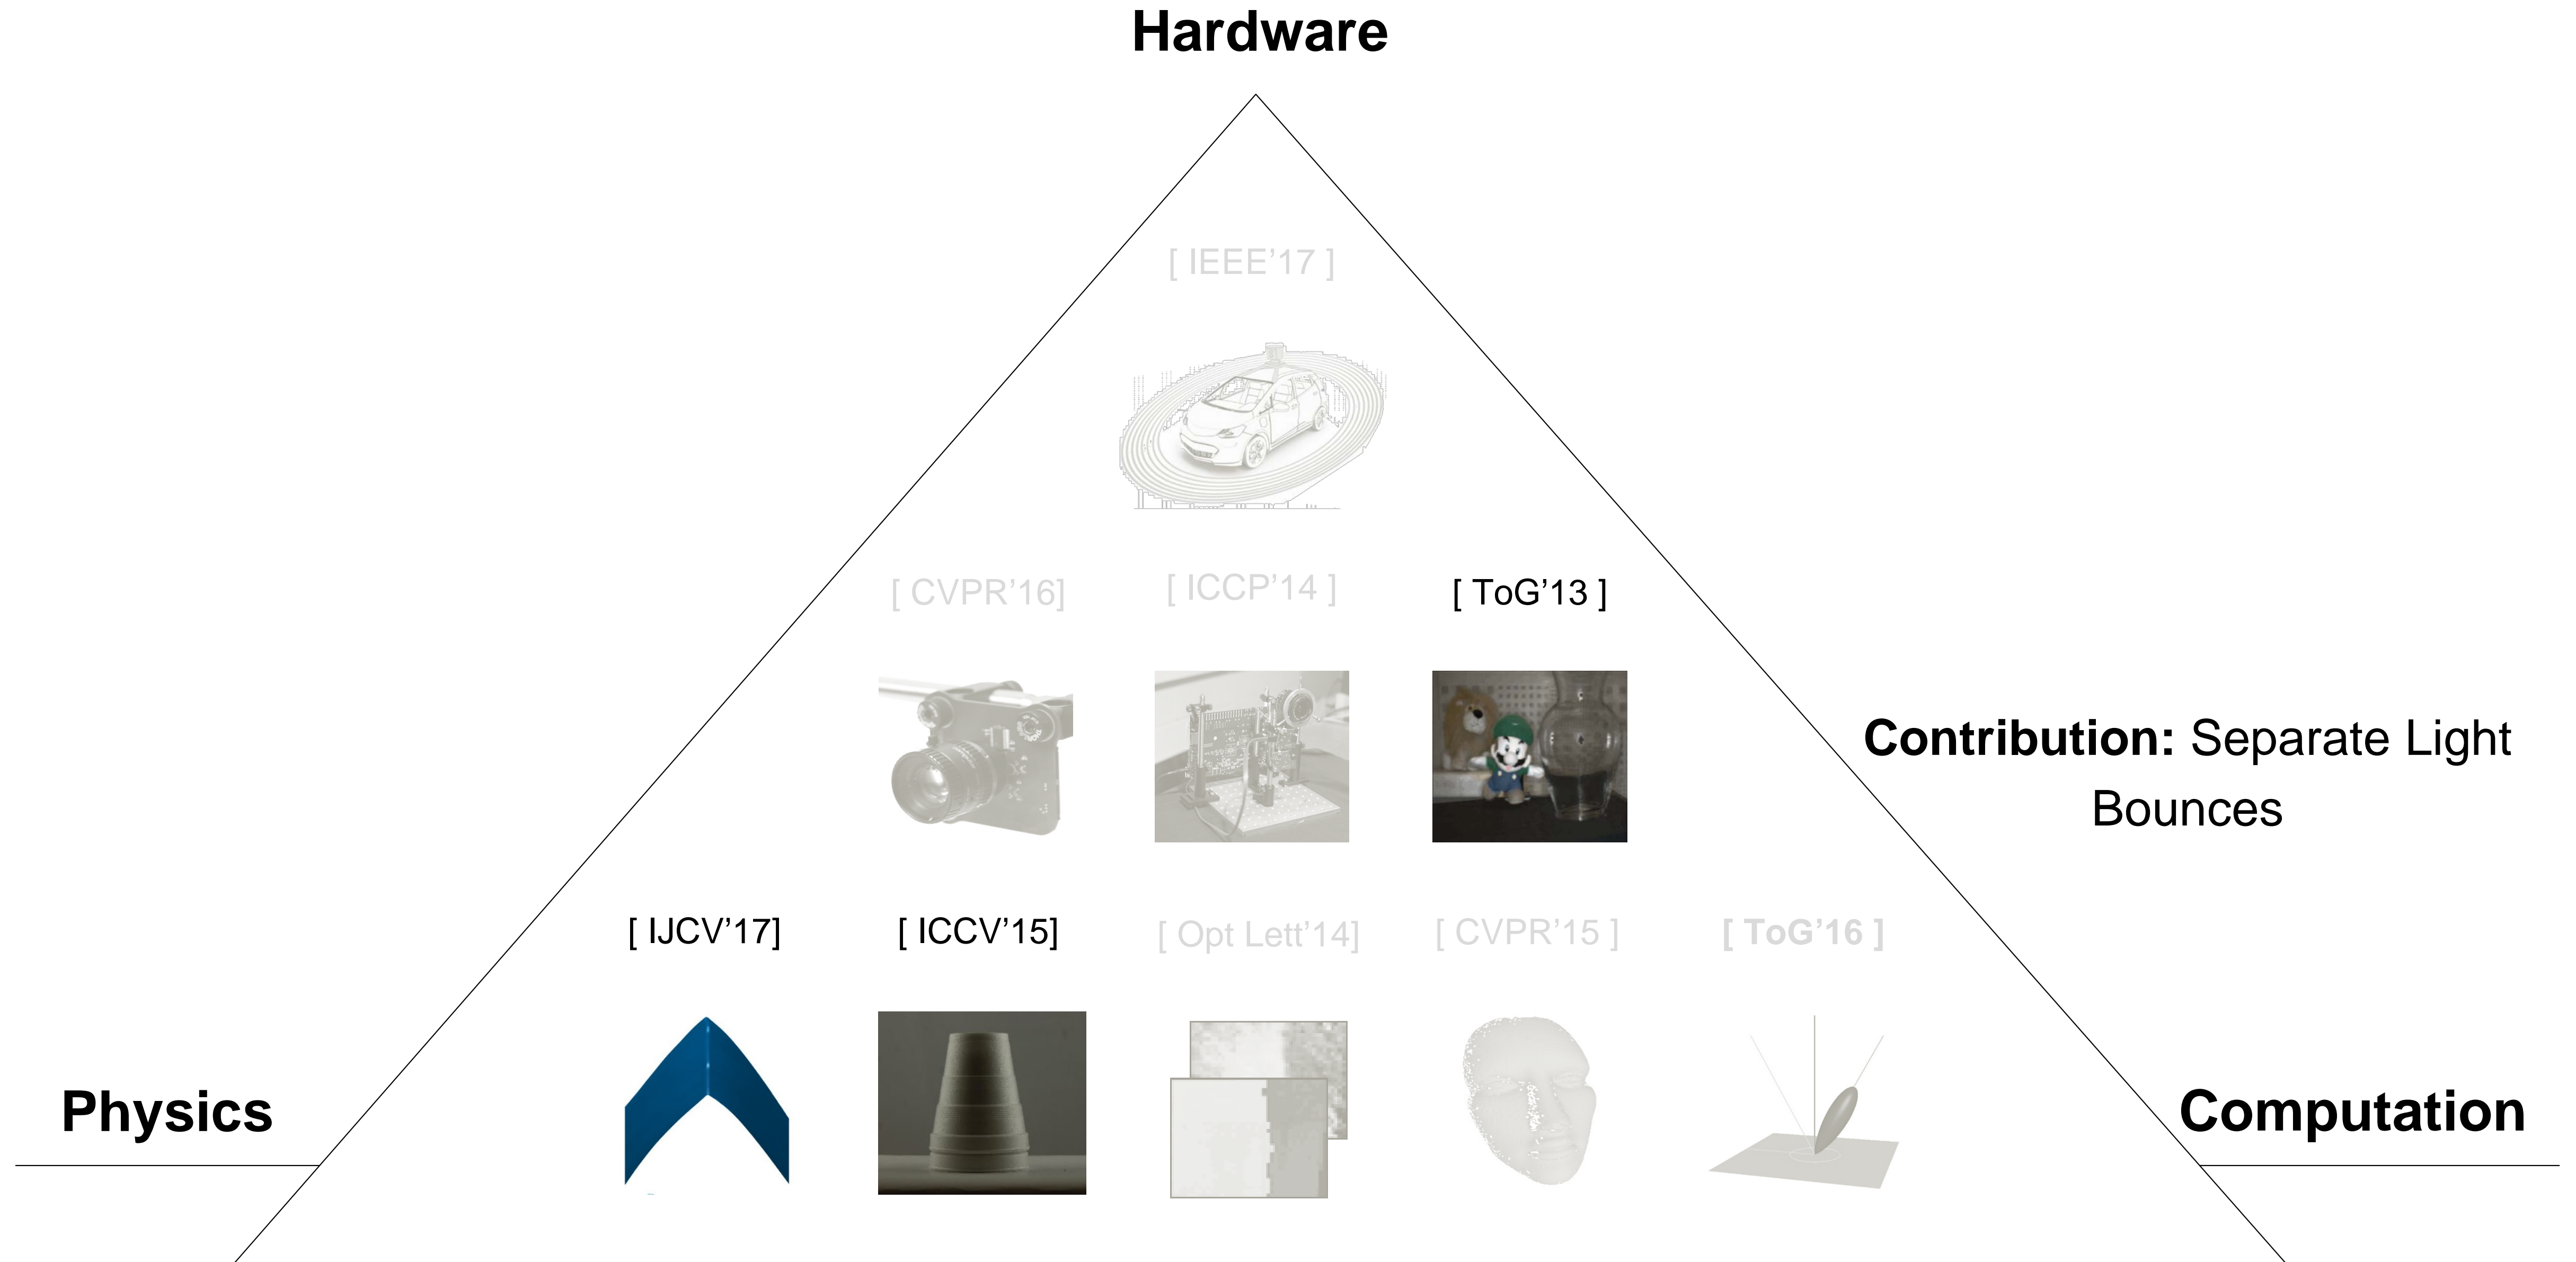
3D Shape of Translucent Objects




Seeing around Corners



Physics, Hardware, and Computation





Physics, Hardware, and Computation



Tool Application: Multi-Light Separation



A



Scene

Tool Application: Multi-Light Separation

Naïve Approach





Scene

$$I = I_A$$

Tool Application: Multi-Light Separation

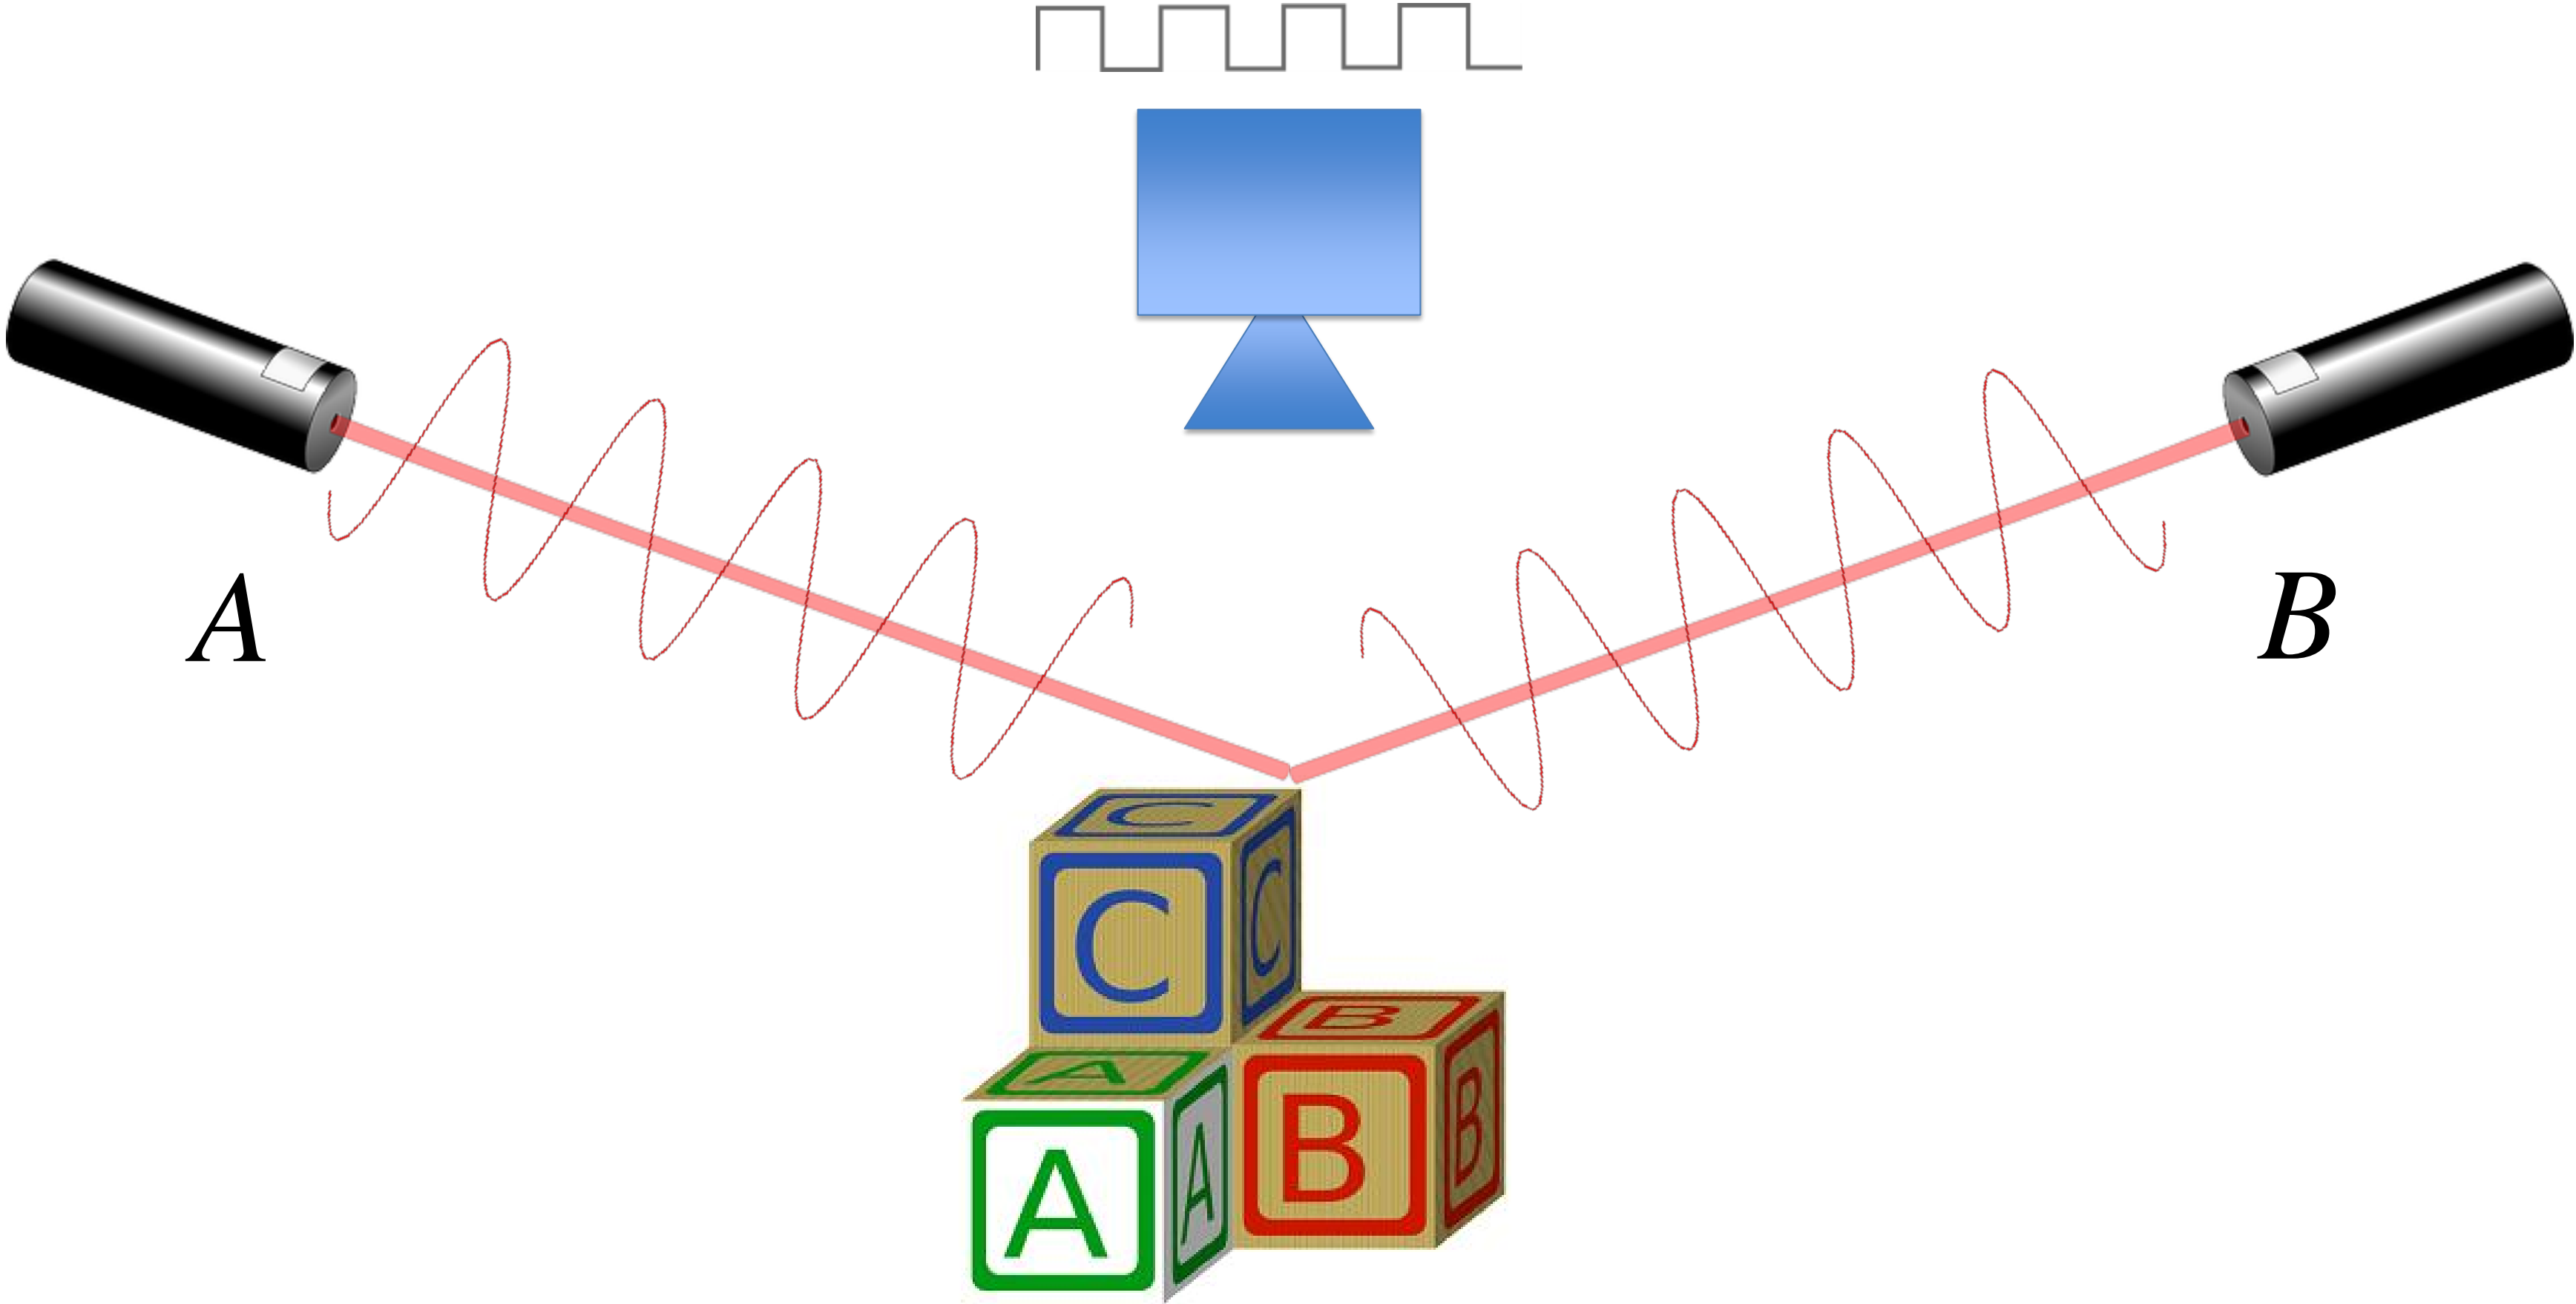
Naïve Approach



Scene

$$I = I_B$$

Tool Application: Multi-Light Separation




Scene

Sense of "Optimal Codes" in context of SNR and power bandwidth


Tool Application: Multi-Light Separation




Tool Application: Multi-Light Separation






Tool Application: Multi-Light Separation



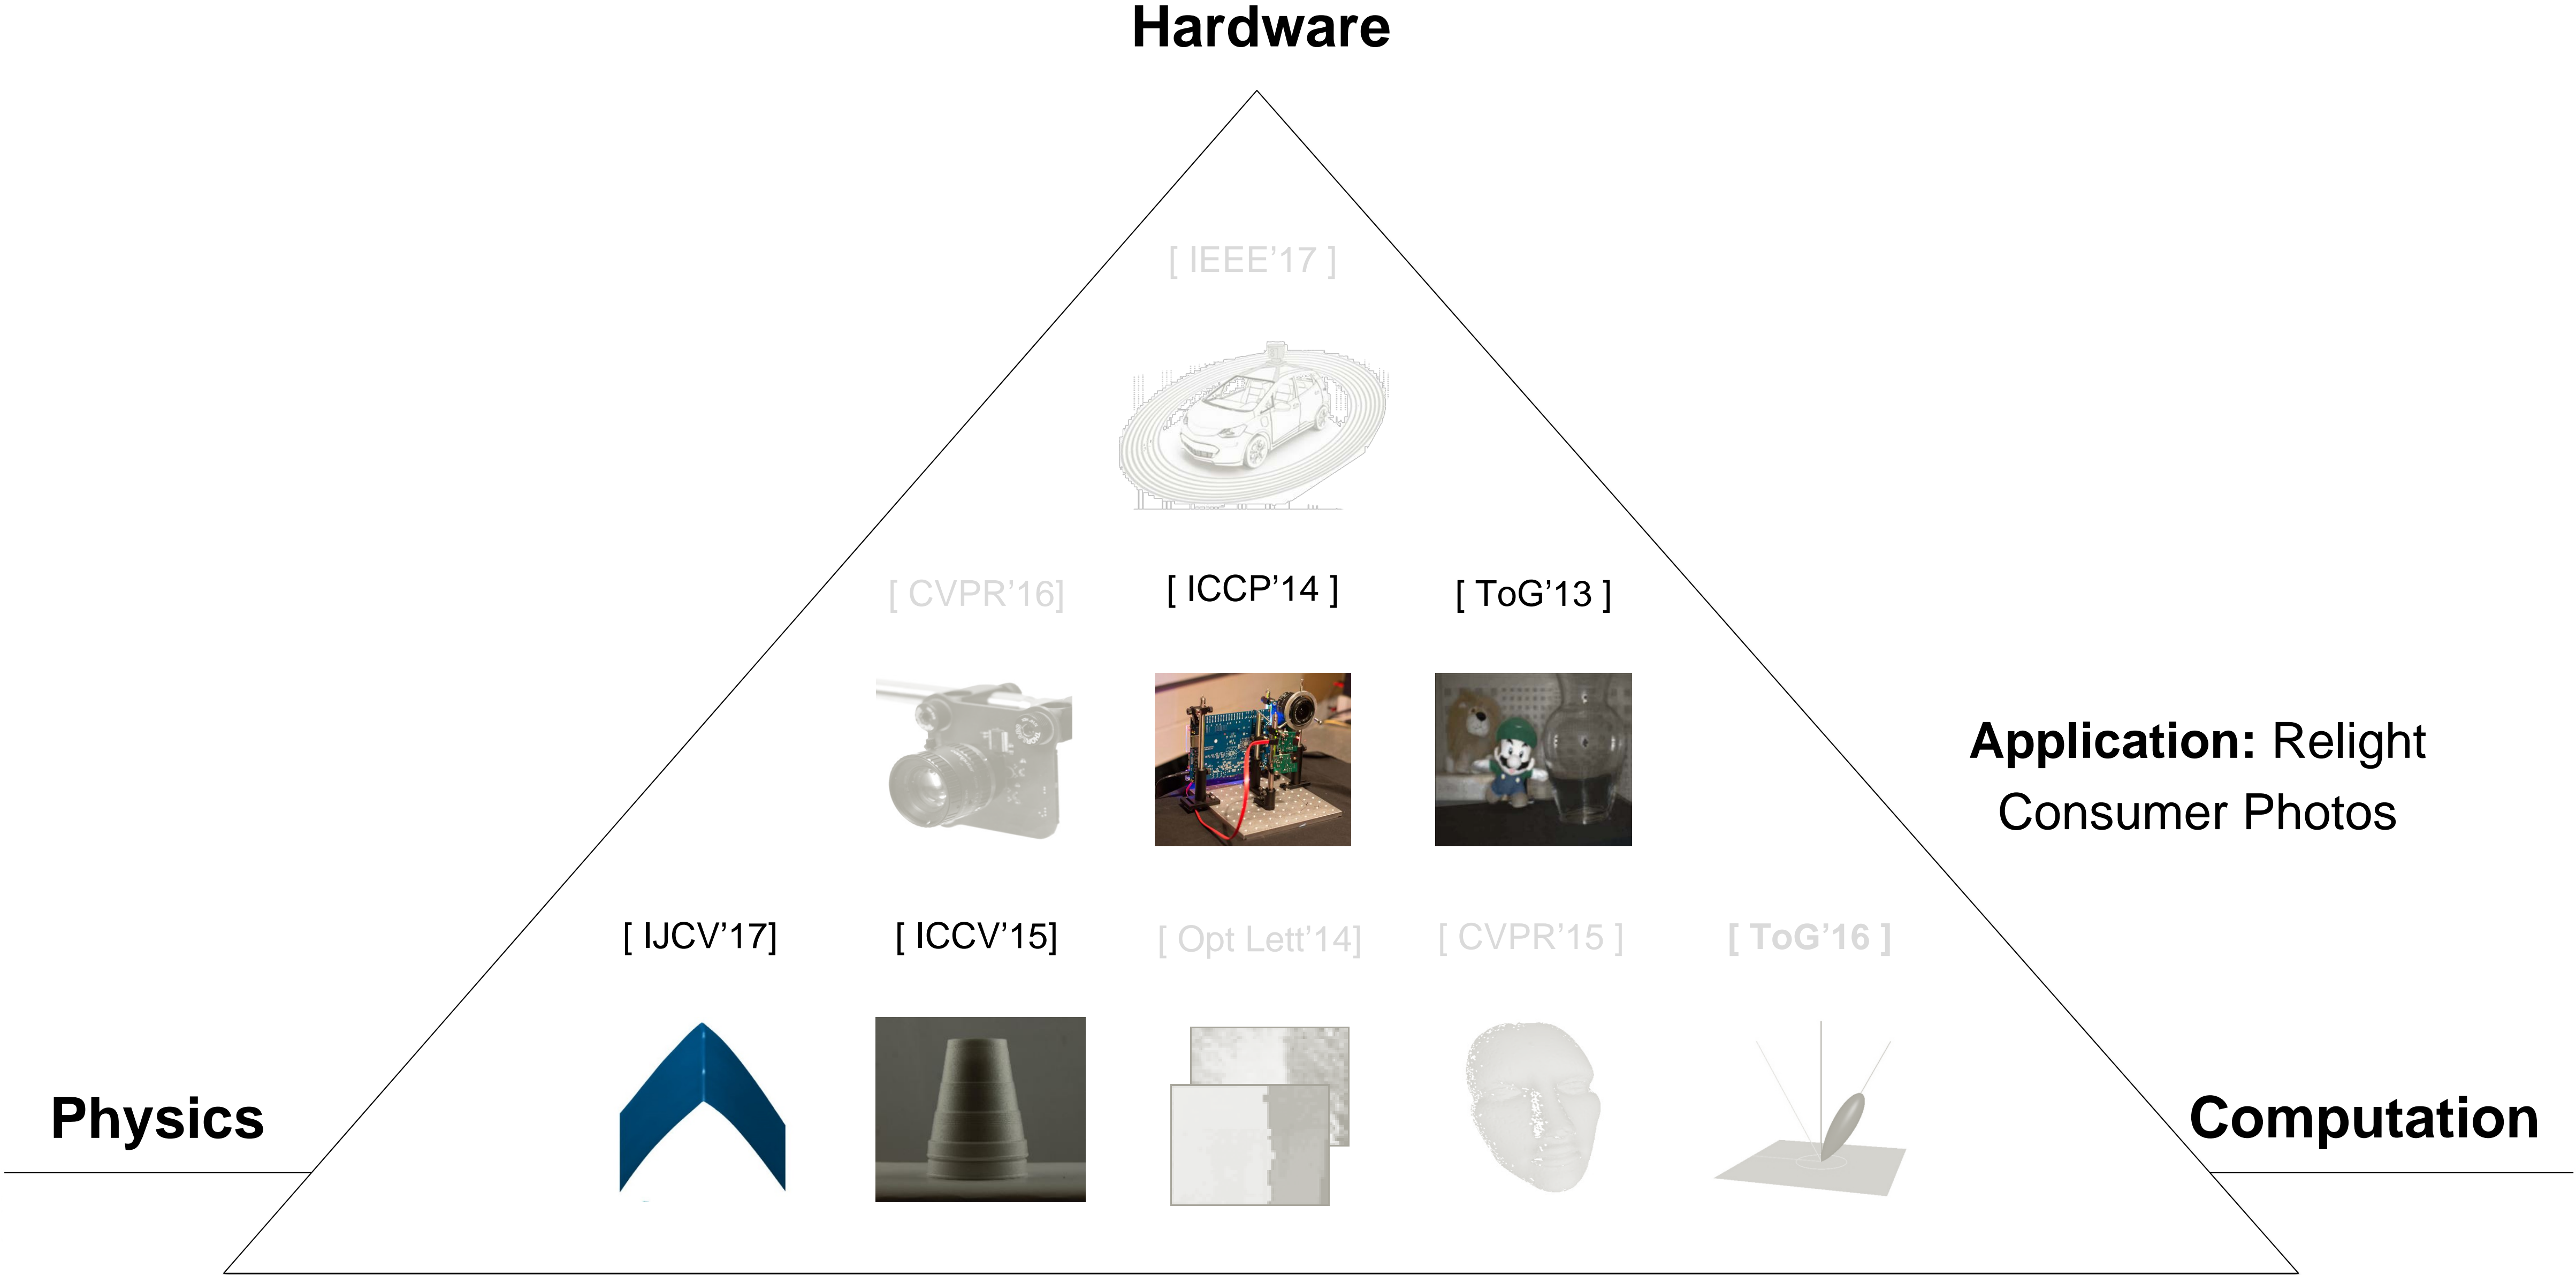
Tool Application: Multi-Light Separation



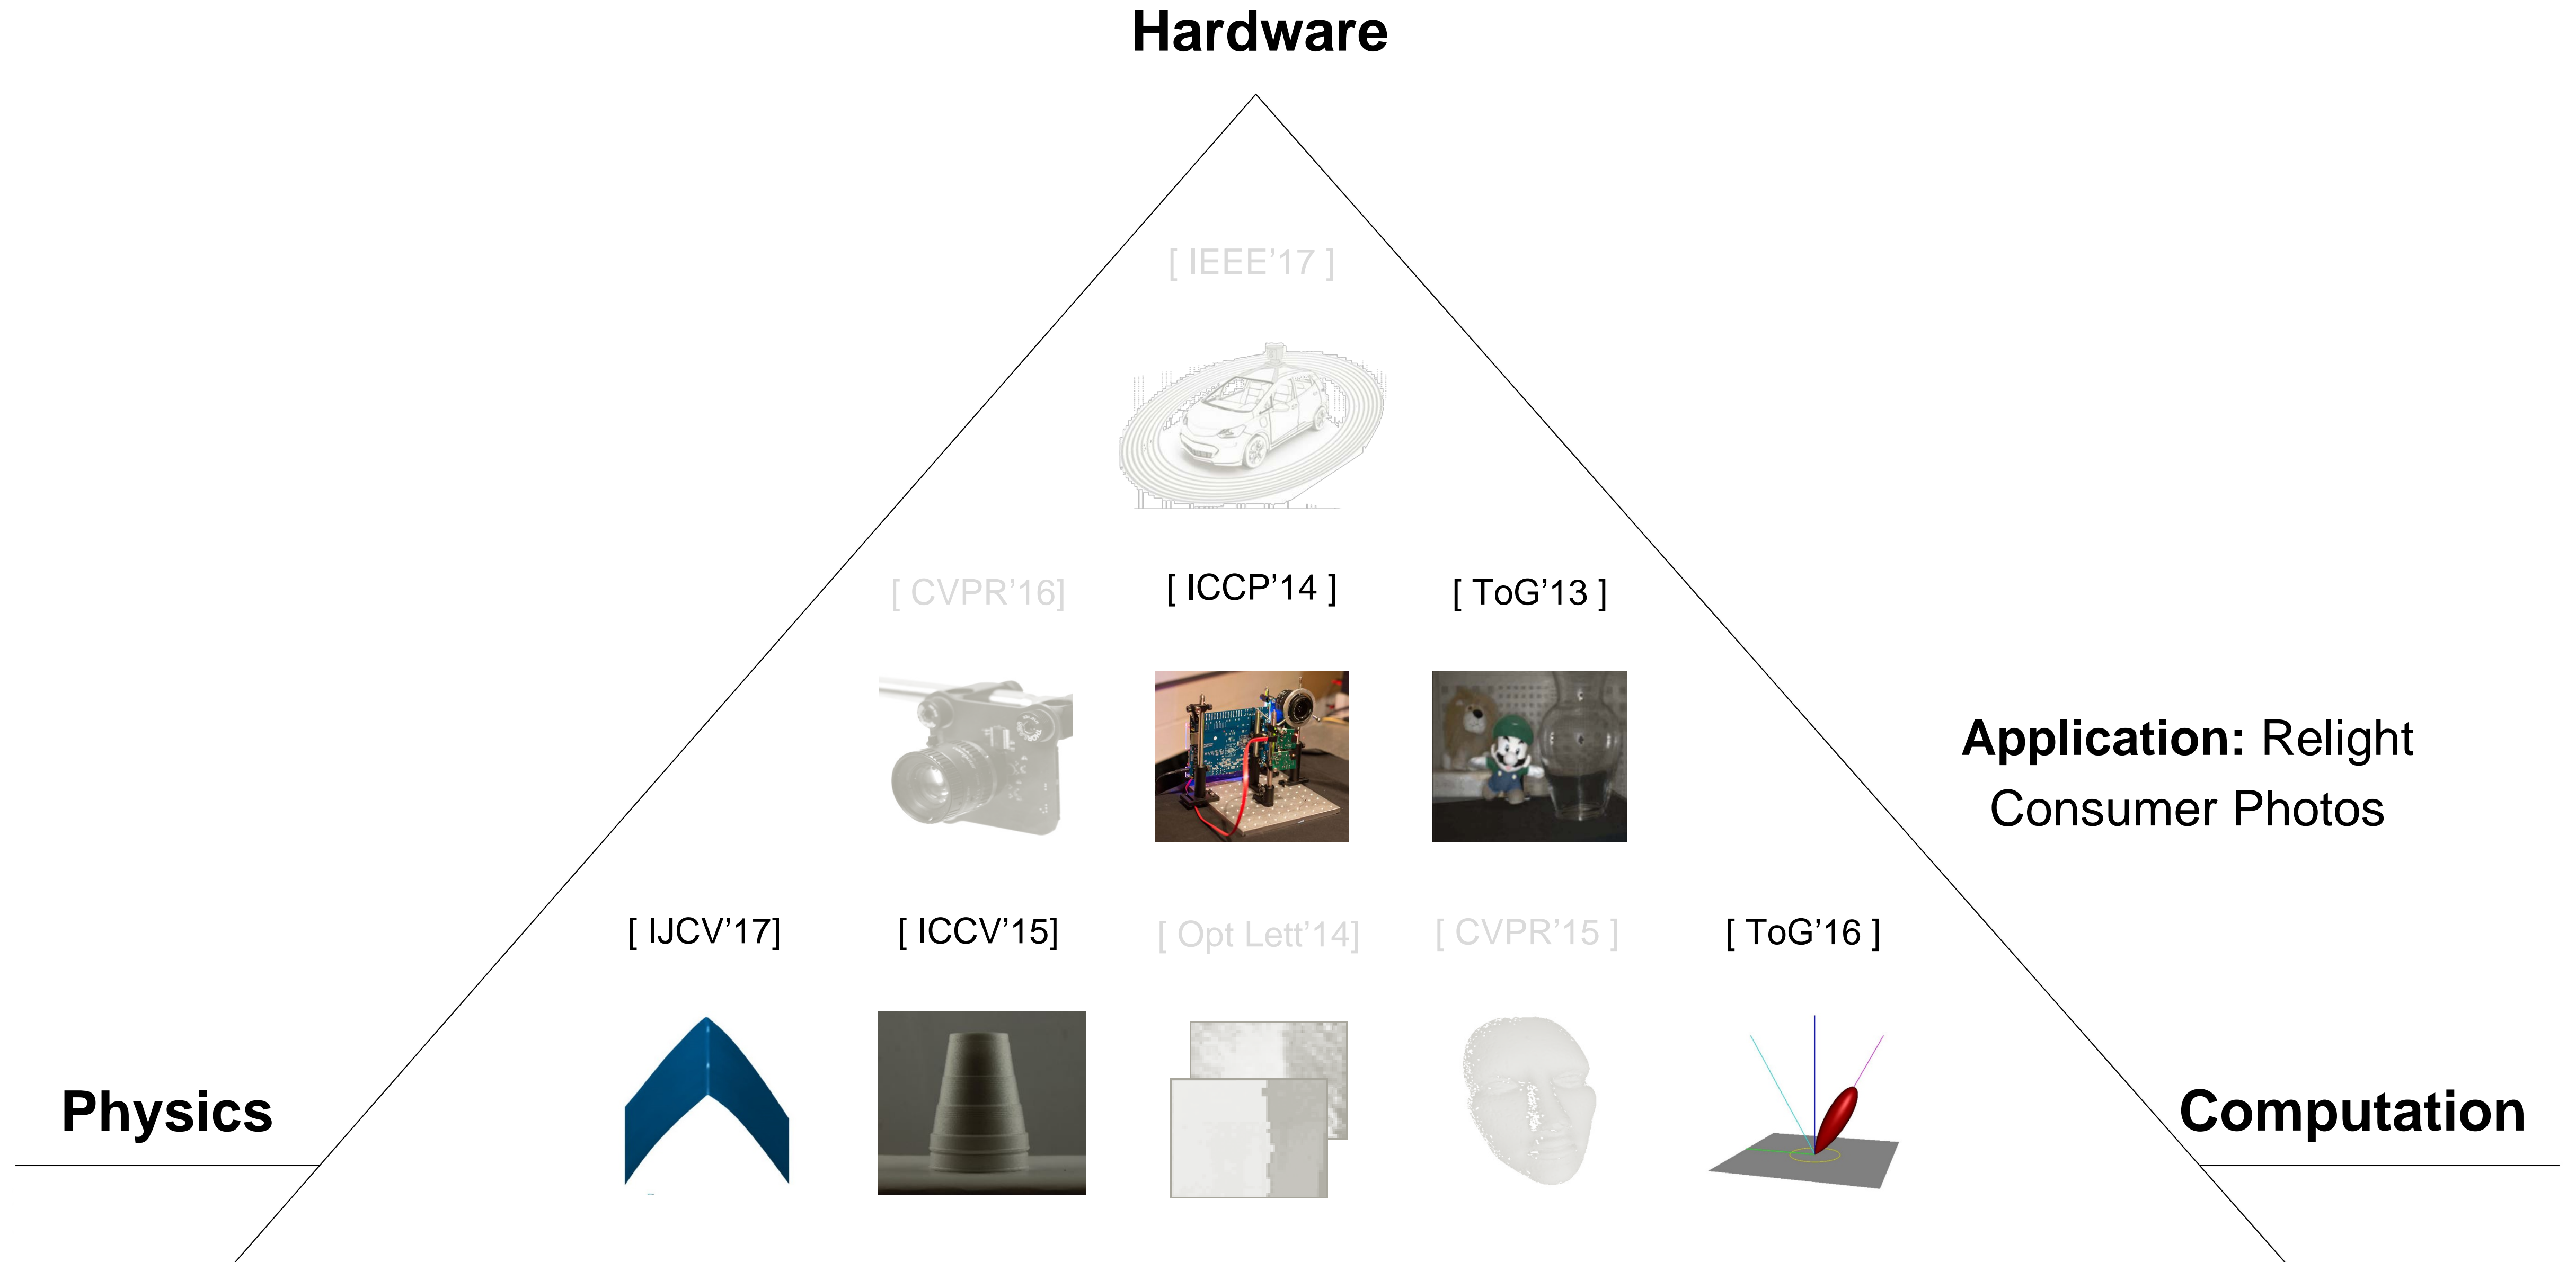
Tool Application: Multi-Light Separation



Physics, Hardware, and Computation



Physics, Hardware, and Computation





Scattering is not always discrete




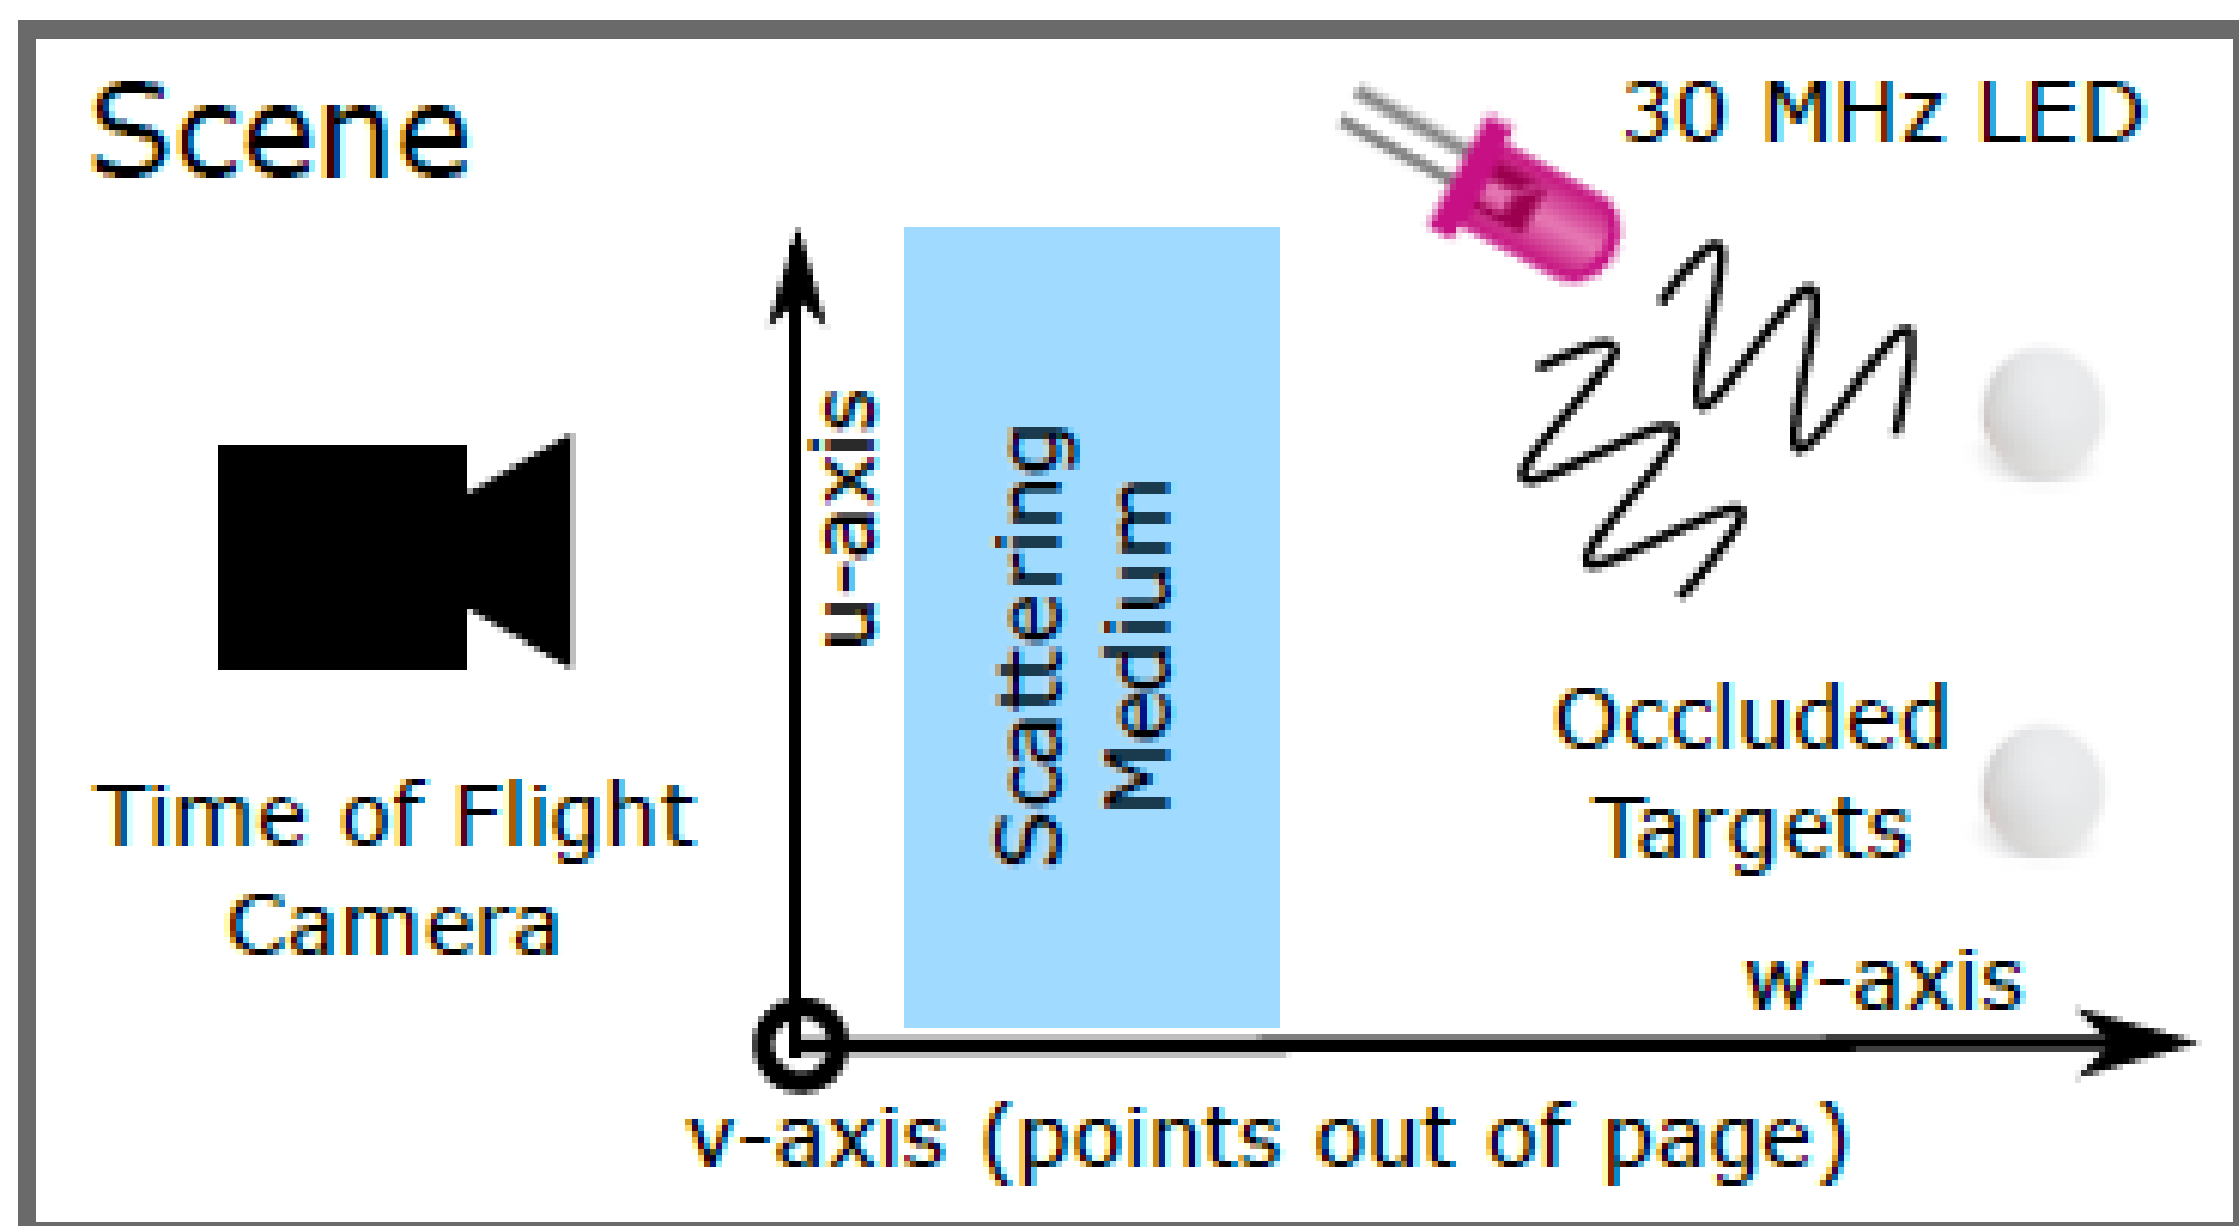
Exploit Space-time Diversity

What's behind the water?






(b) Highly Scattering




Physics, Hardware, and Computation

Hardware


[IEEE'17]



[CVPR'16]



[ICCP'14]




[ToG'13]




[IJCV'17]



[ICCV'15]




[Opt Lett'14]



[CVPR'15]



[ToG'16]



Physics


Computation

Physics, Hardware, and Computation

Hardware

When does the separation work?


[IEEE'17]



[CVPR'16]



[ICCP'14]



[ToG'13]



[IJCV'17]



[ICCV'15]




[Opt Lett'14]



[CVPR'15]





[ToG'16]



Physics

Computation

Intuition



A Guarantee on Resolvability of Light Paths

□ **Proposition** [Kadambi et al. CVPR16]

Two light-paths of equal amplitude can be separated if the difference in optical path length is greater than:

$$\Delta z > 0.6 \frac{c}{\Delta f_M}$$

Difference in optical path	Δz
Camera Frequency Bandwidth	Δf_M
Speed of Light	c




Figure 3. We show that the noiseless bound derived in *Proposition 1* is valid in typical shot noise limited scenarios.


* Valid in shot noise limited cases

Physics, Hardware, and Computation

Hardware

Contribution: separation works at high-frequencies


[IEEE'17]



[CVPR'16]



[ICCP'14]



[ToG'13]




[IJCV'17]



[ICCV'15]




[Opt Lett'14]



[CVPR'15]



[ToG'16]



Physics


Computation

Physics, Hardware, and Computation

Hardware

Goal: High-frequency 10 GHz Implementation


[IEEE'17]



[CVPR'16]



[ICCP'14]




[ToG'13]



[IJCV'17]



[ICCV'15]




[Opt Lett'14]



[CVPR'15]



[ToG'16]




Physics


Computation

New “MicroLIDAR” Project for Descattering


Study optical wave and electronic modulation jointly



CMOS Camera
[Kadambi et al. CVPR16]



Fiber optic **MicroLIDAR**
[Kadambi et al. IEEE17]



The MIT Press




Textbook: “Principles of Time-resolved Imaging” (Free in Fall’18)

Physics, Hardware, and Computation

Hardware

Goal: High-frequency 10 GHz Implementation


[IEEE'17]



[CVPR'16]



[ICCP'14]



[ToG'13]




[IJCV'17]



[ICCV'15]




[Opt Lett'14]



[CVPR'15]



[ToG'16]




Physics

Computation

Physics, Hardware, and Computation

Hardware


[IEEE'17]



[CVPR'16]



[ICCP'14]



[ToG'13]




**Application: Relight
Consumer Photos**


[IJCV'17]




[ICCV'15]




[Opt Lett'14]



[CVPR'15]



[ToG'16]




Physics

Computation

Health Applications

Hardware


[IEEE'17]



[CVPR'16]



[ICCP'14]



[ToG'13]




[IJCV'17]




[ICCV'15]




[Opt Lett'14]




[CVPR'15]



[ToG'16]







Health Applications



Physics

Computation

Canon Document Scanner into an X-ray imager?




The future of image processing is 7D?



Polarization for 3D



Transient Imaging



Space-Transient Imaging

7D Image Processing

Academic Questions:

How do multiple bounces relate to polarization?

*How can we overcome **diffraction**?*

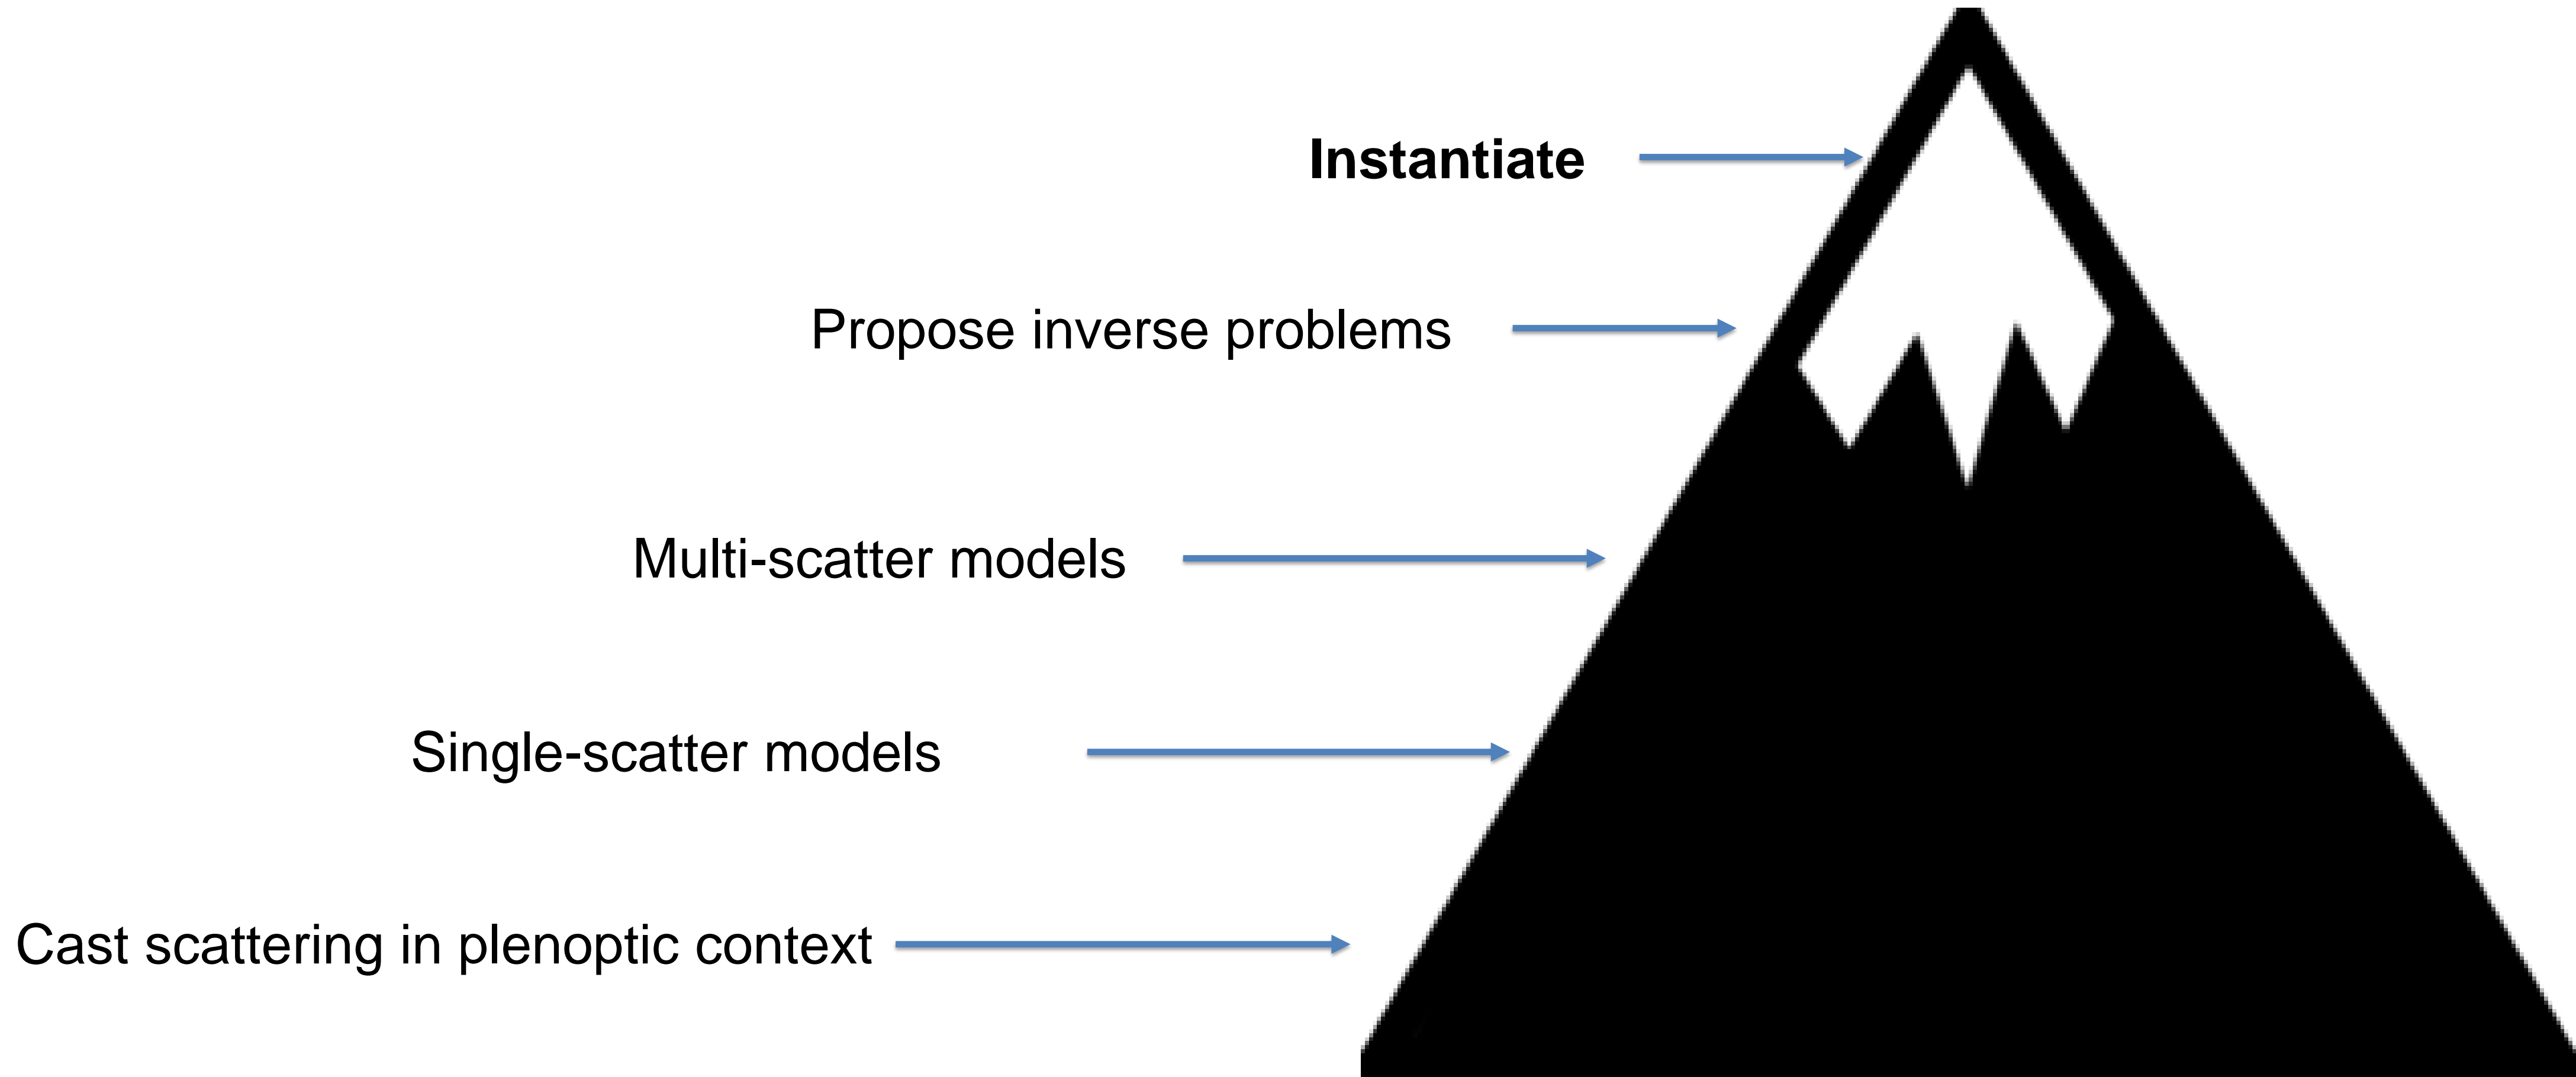
*How does **scatter** relate to transience?*

Provable Guarantees of Imaging Correctness?

Driving through Fog: Separate ballistic and scattered photons based on minute axial path delays or polarimetric signal



Driving through Fog



What are other possible mountains?

Future Direction: Physics-based machine learning

Goal: How can we unite physics and learning-based ideas?

Challenges to overcome:

- Need specialized datasets
- Need specialized CG rendering
- Need to rethink algorithms
- Need to justify approach




Image-based Machine Learning [Laina16]


← Trained on 3D Database from NYU [Silberman12]

Future Direction: Proving “correctness” of an imaging approach

New problems, new bounds – novel problem statements require new, provable bounds

Towards optimality – can we prove analytically that a computational camera is the best possible solution?

New systems, new simulators – e.g. puts pressure on SIGGRAPH and CVPR to evolve.



Demand Pressure for New Cameras




Drone Visual Systems



Evolution of Nanoscopes



Image Deeper through Tissue



Autonomous Driving



Mobile Photography



Sound + Light

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