Special Topics:
Computational Imaging

http://16385.courses.cs.cmu.edu/

16-385 Computer Vision
Fall 2020, Lecture 27

computational imaging

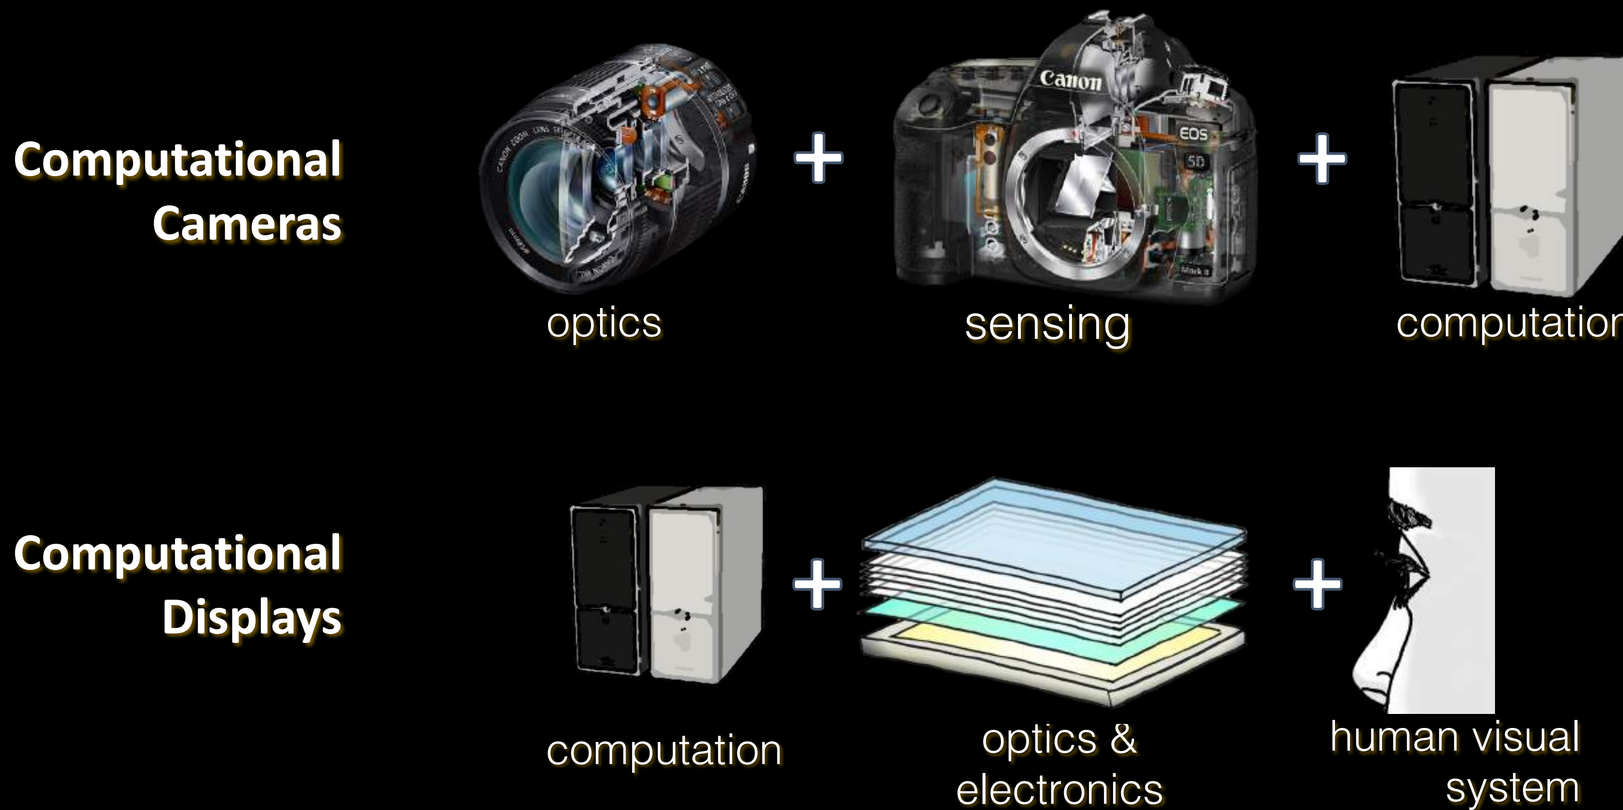

computational imaging

Computational Cameras

optics

+

sensing

+

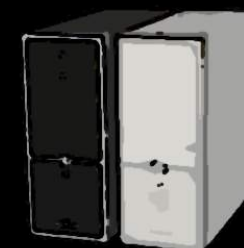

computation

Computational Displays

computation

+

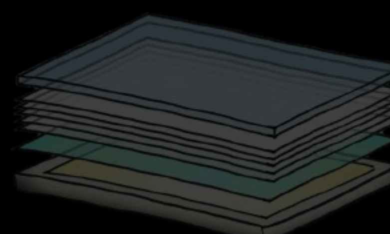

optics & electronics

+

human visual system

computational imaging

Computational Cameras

HDR Imaging [Debevec, Nayar, ...]

Super-resolution [Baker, ...]

Light Fields [Levoy, ...]

Computational Displays

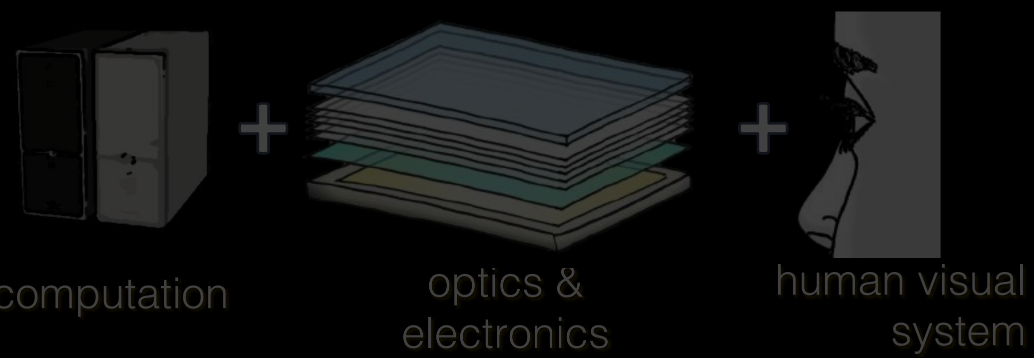

computational imaging

Computational Cameras

HDR Imaging [Debevec, Nayar, ...]

Super-resolution [Baker, ...]

Light Fields [Levoy, ...]

Computational Displays

HDR Display [Seetzen, ...]

Super-resolution [Hirsch, Heide, ...]

Light Fields [Wetzstein, ...]

computational imaging

Computational Cameras

+

Computational Displays

computational imaging

Computational Cameras

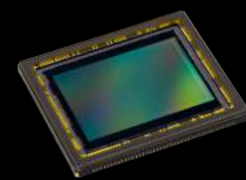

+

Computational Displays

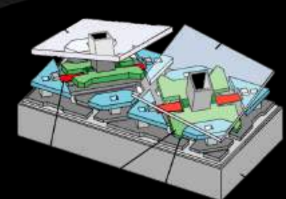

light transport in a general scene

light transport in a general scene

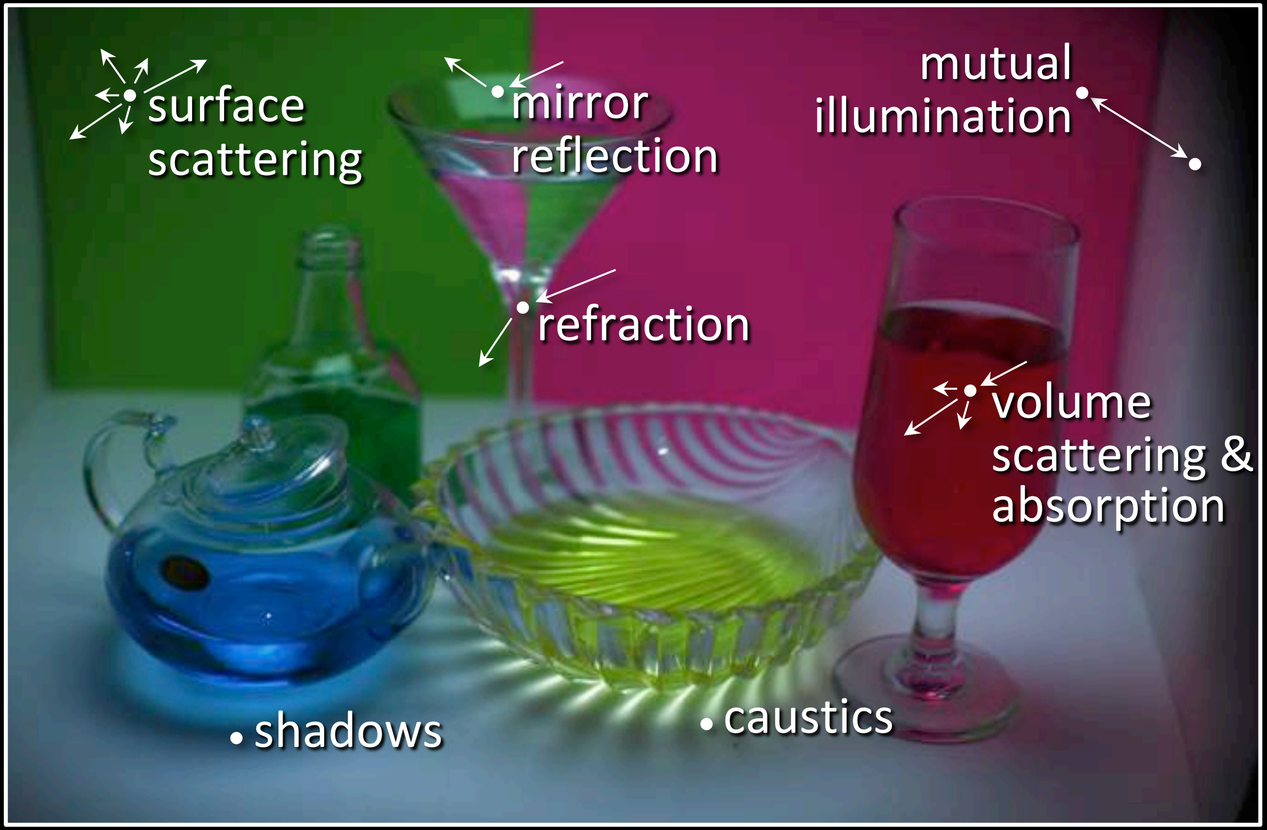

light transport in a general scene

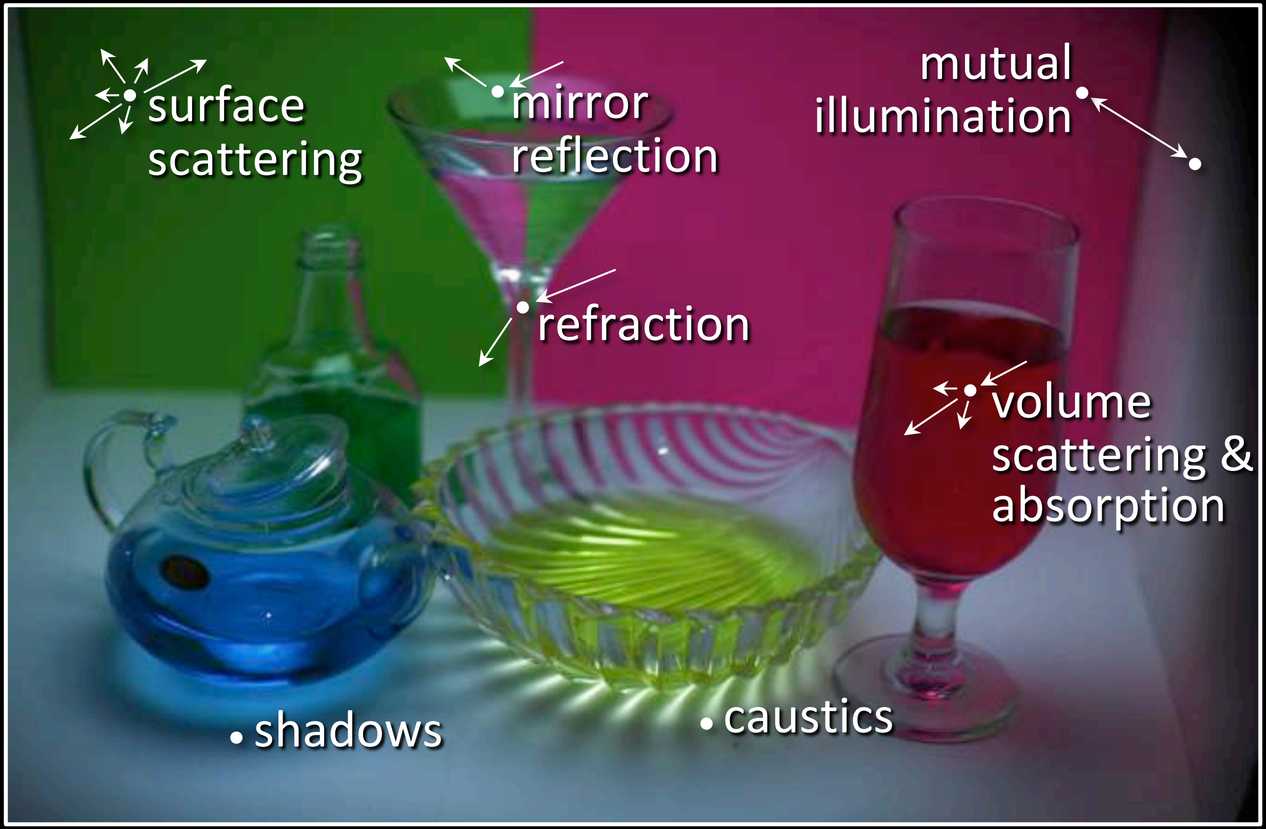

computational light transport

computational light transport involves using controllable light sources & cameras to sample, acquire or analyze a scene's transport function

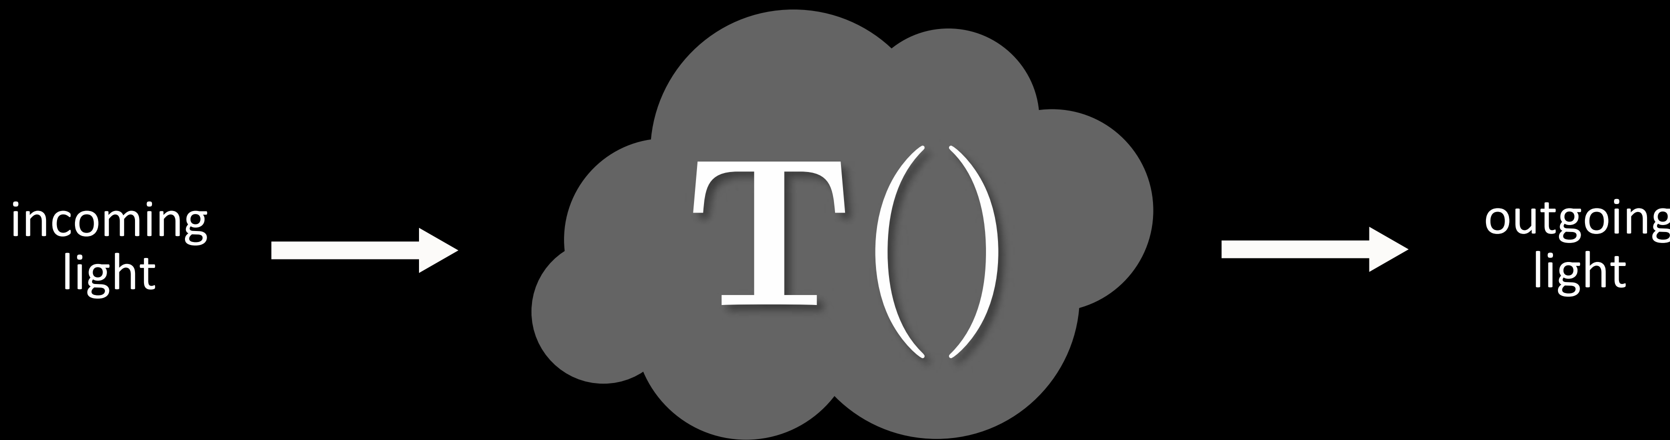

3D imaging for autonomous cars

LIDAR (light detection and ranging)
Velodyne VLS-128

3D imaging for autonomous cars

LIDAR (light detection and ranging
Velodyne VLS-128

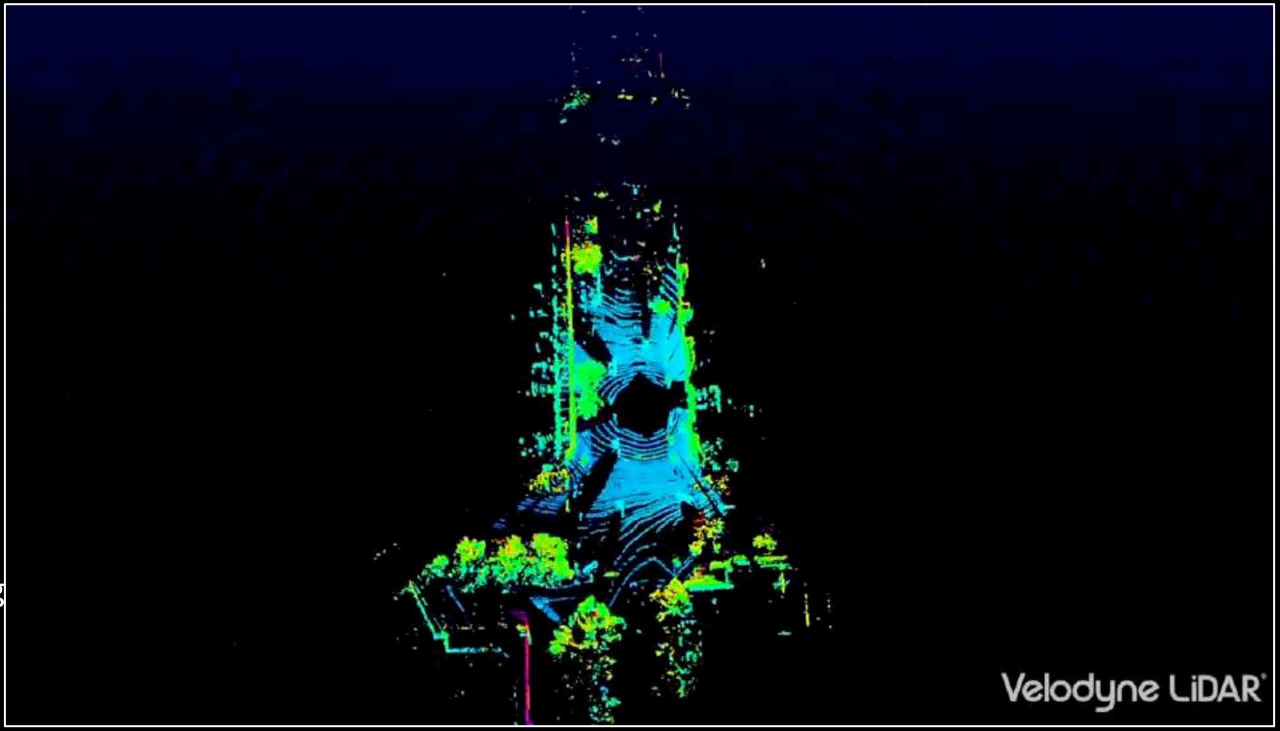

3D imaging for smartphones

iPhone X

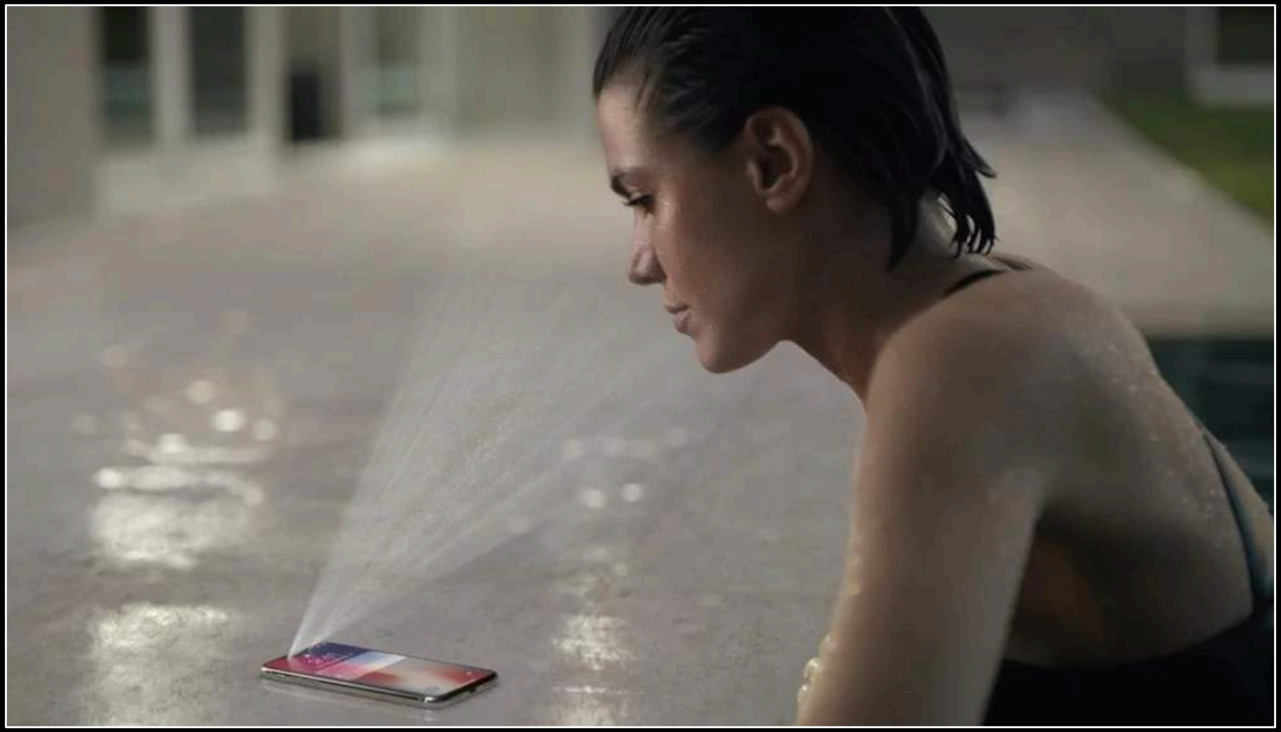

3D imaging for smartphones

iPhone X

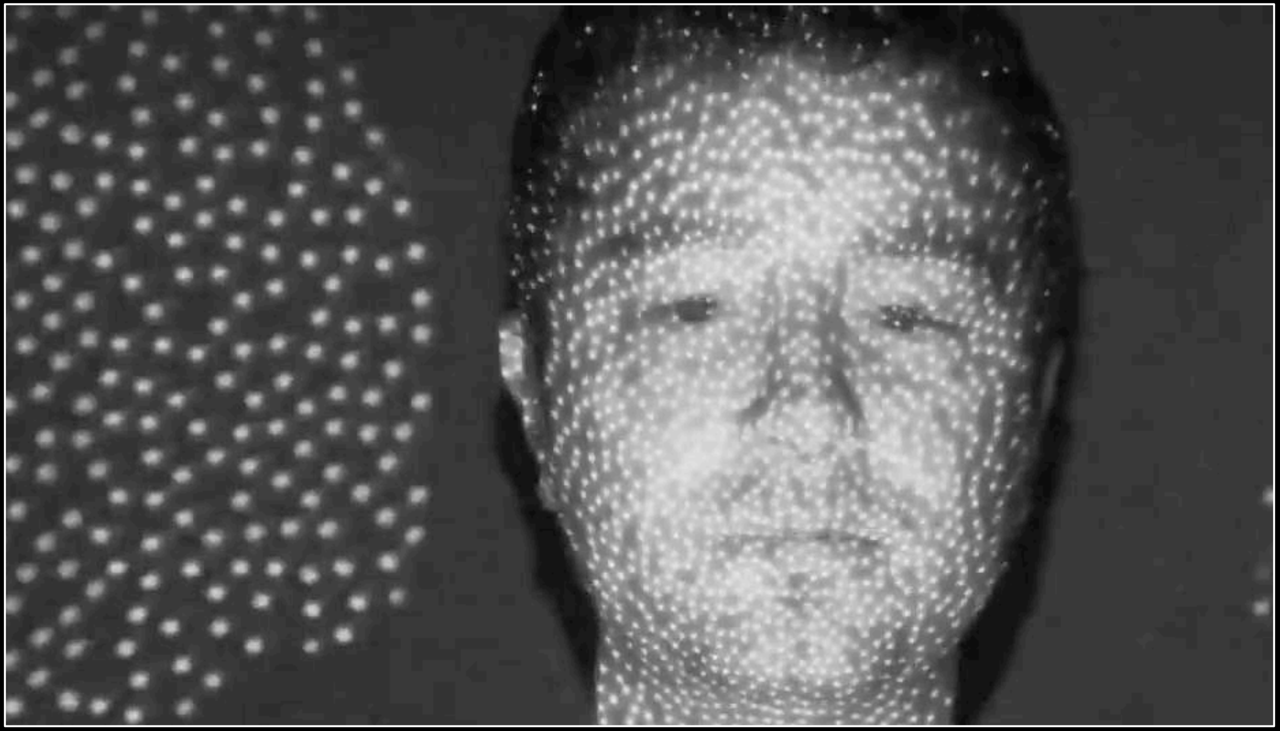

Paul Debevec’s light stage 6

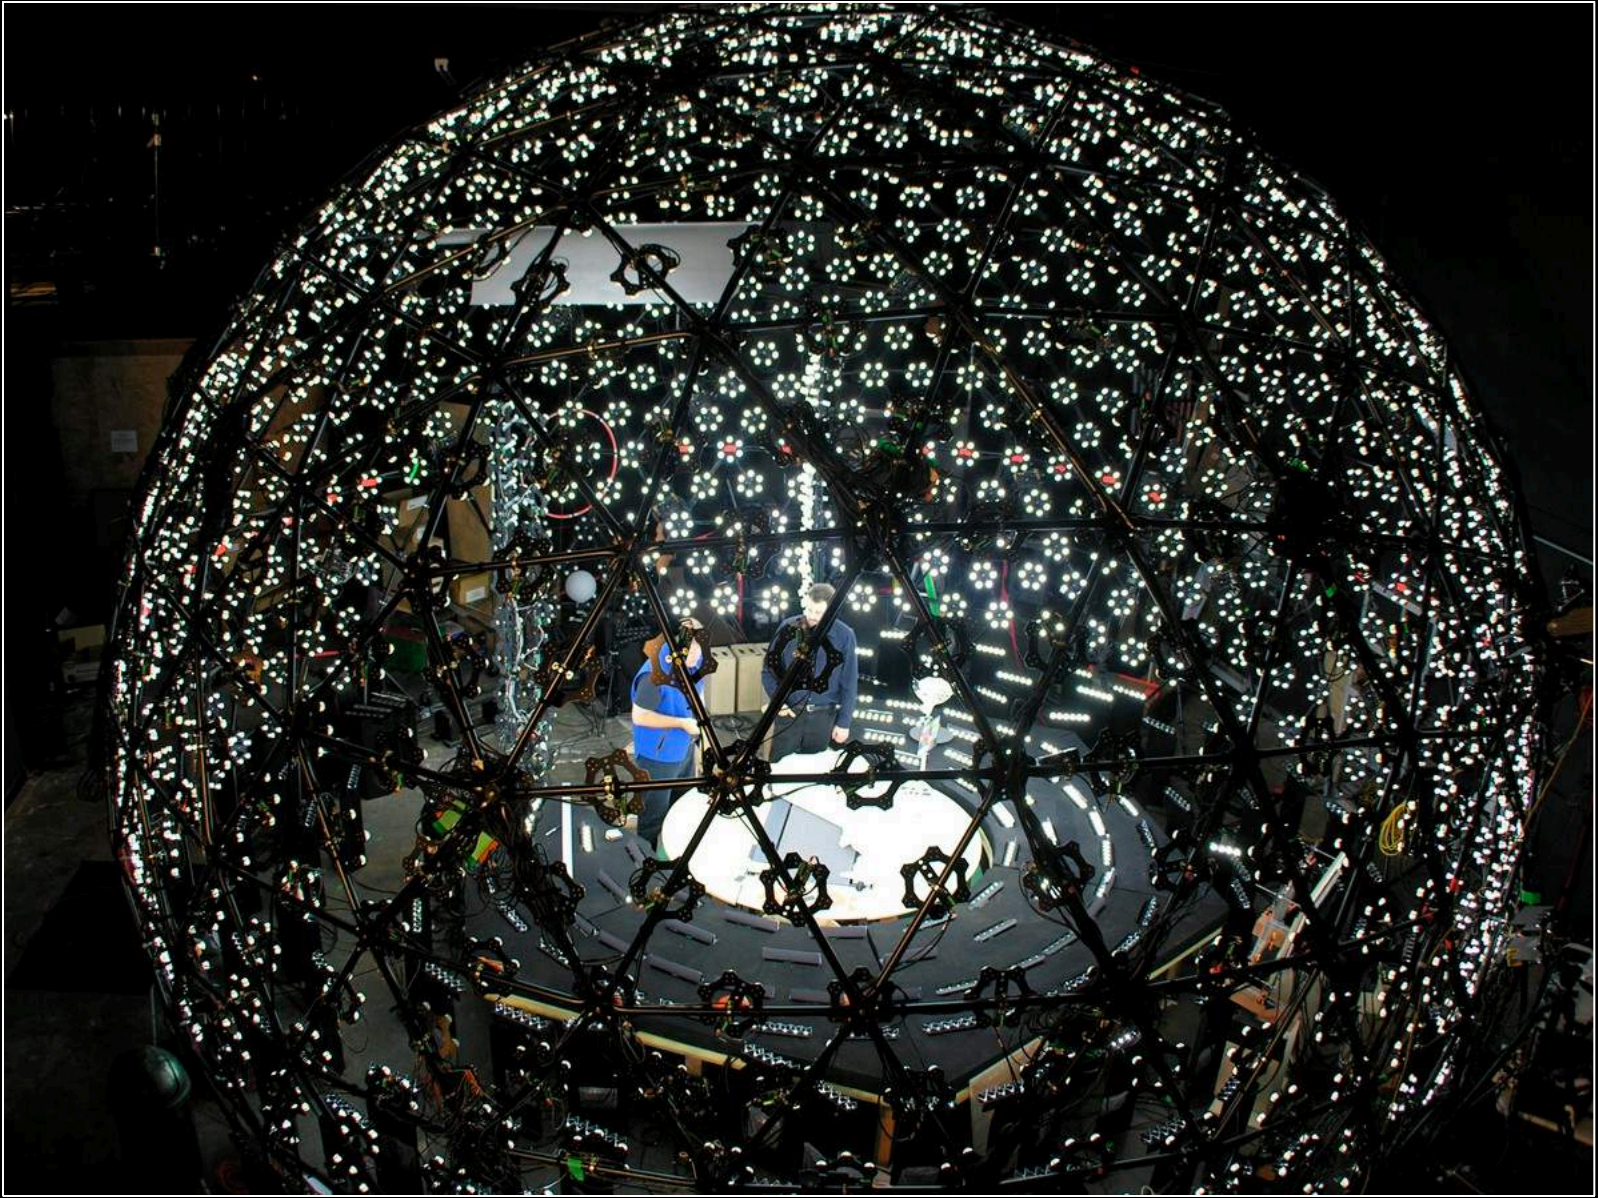

https://www.fxguide.com/fxfeatured/light_stage_6/

Paul Debevec's light stage 6

https://www.fxguide.com/fxfeatured/light_stage_6/

mobile light stage

https://news.usc.edu/71893/usc-digital-technology-creates-3-d-portraits-of-obama/

aperture correlation microscope (source: Zeiss)

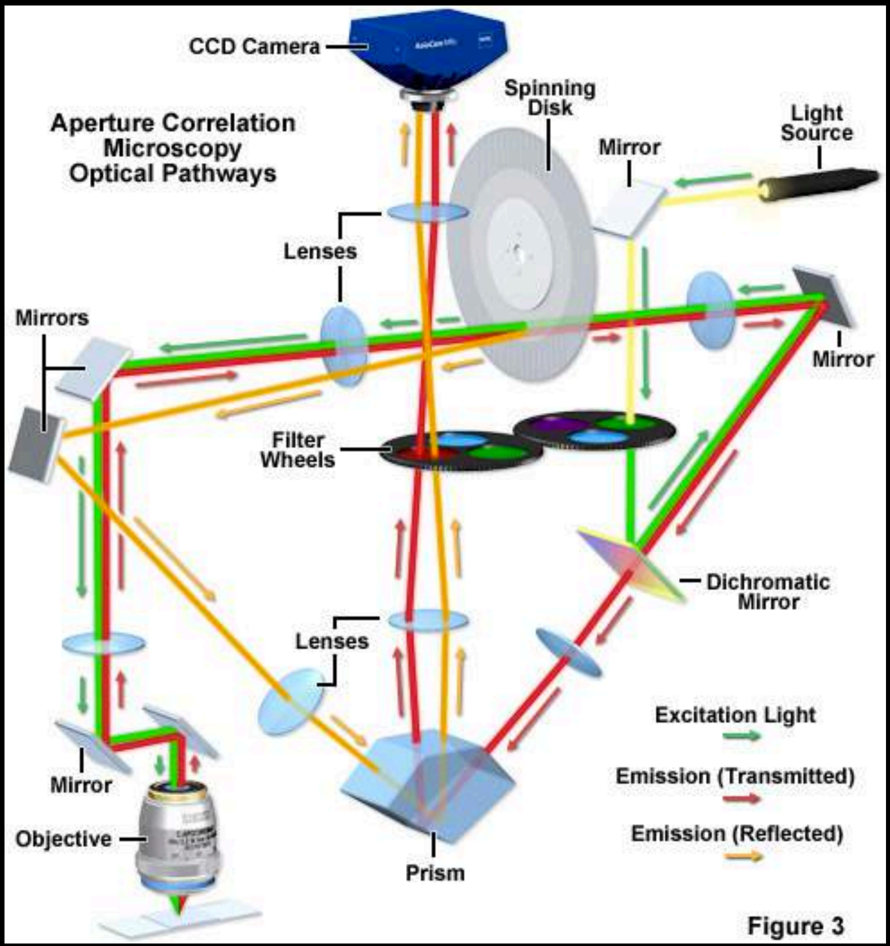

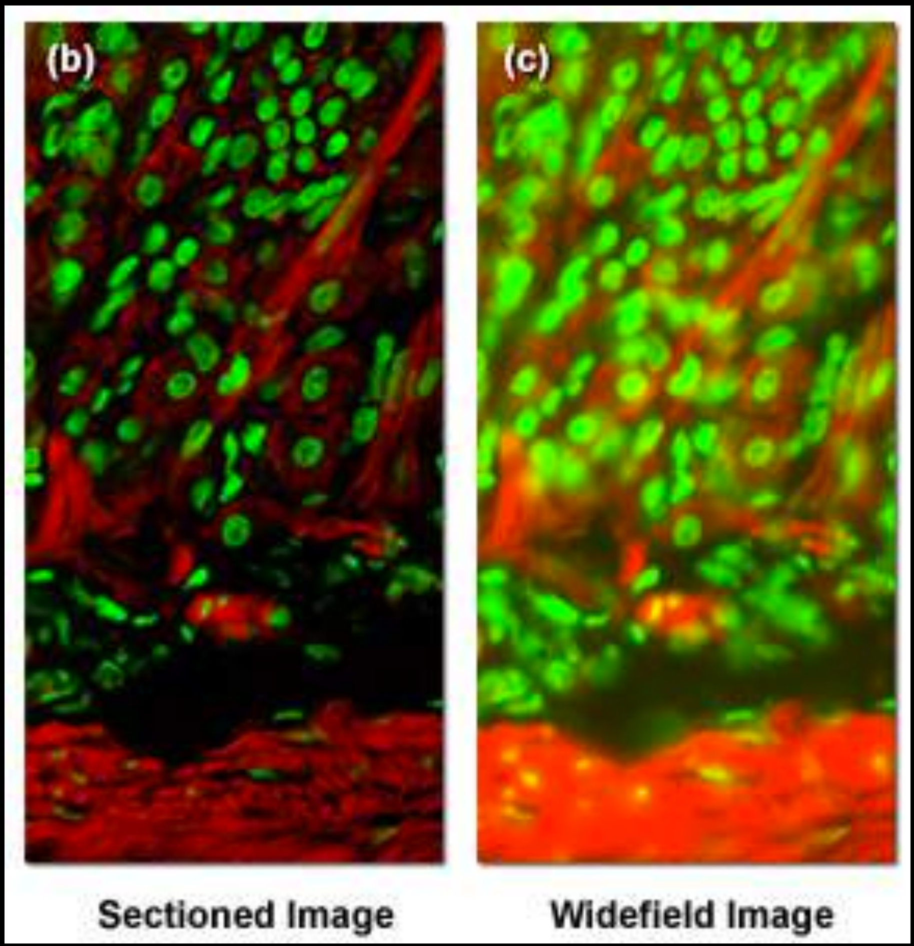

overview

1. **the light transport matrix:** a general model for the transfer of radiant energy

2. example transport matrices for real scenes

3. challenges associated with analyzing transport matrices

4. optical algorithms to analyze light transport

overview

1. **the light transport matrix:** a general model for the transfer of radiant energy

2. example transport matrices for real scenes

3. challenges associated with analyzing transport matrices

4. optical algorithms to analyze light transport

modeling light transport

how do we model light transport between one light source and one sensor?

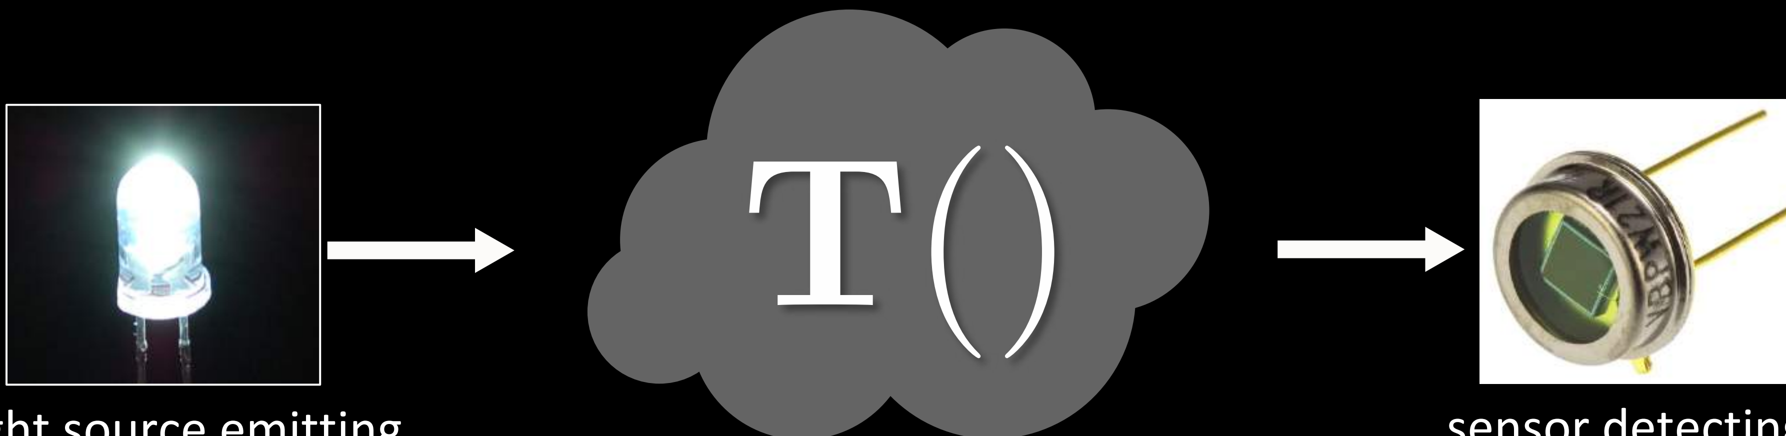

light source emitting radiant energy **x** (measured in Joules)

sensor detecting radiant energy **y** (measured in Joules)

modeling light transport

how do we model light transport between one light source and one sensor?

light source emitting radiant energy **x** (measured in Joules)

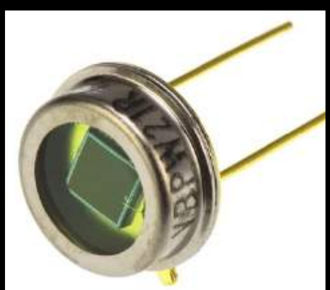

sensor detecting radiant energy **y** (measured in Joules)

observations:

- radiant energy is always non-negative, i.e., can't detect "negative" energy"
- the transport function T() is homogeneous of degree 1, i.e., the function satisfies T(s*x) = s T(x)

property 1: homogeneity of degree 1

weight

1

x

=

scene light by light source at 100% intensity

observations:

- radiant energy is always non-negative, i.e., can't detect "negative" energy"
- the transport function T() is homogeneous of degree 1, i.e., the function satisfies T(s*x) = s T(x)

property 1: homogeneity of degree 1

weight

2

x

=

scene light by light source at 200% intensity

observations:

- radiant energy is always non-negative, i.e., can't detect "negative" energy"
- the transport function T() is homogeneous of degree 1, i.e., the function satisfies T(s*x) = s T(x)

property 1: homogeneity of degree 1

weight

0.5

x

=

scene light by light source at 50% intensity
(fails for saturated pixels)

observations:

- radiant energy is always non-negative, i.e., can't detect "negative" energy"
- the transport function T() is homogeneous of degree 1, i.e., the function satisfies T(s*x) = s T(x)

example when homogeneity of degree 1 condition does **not** hold true

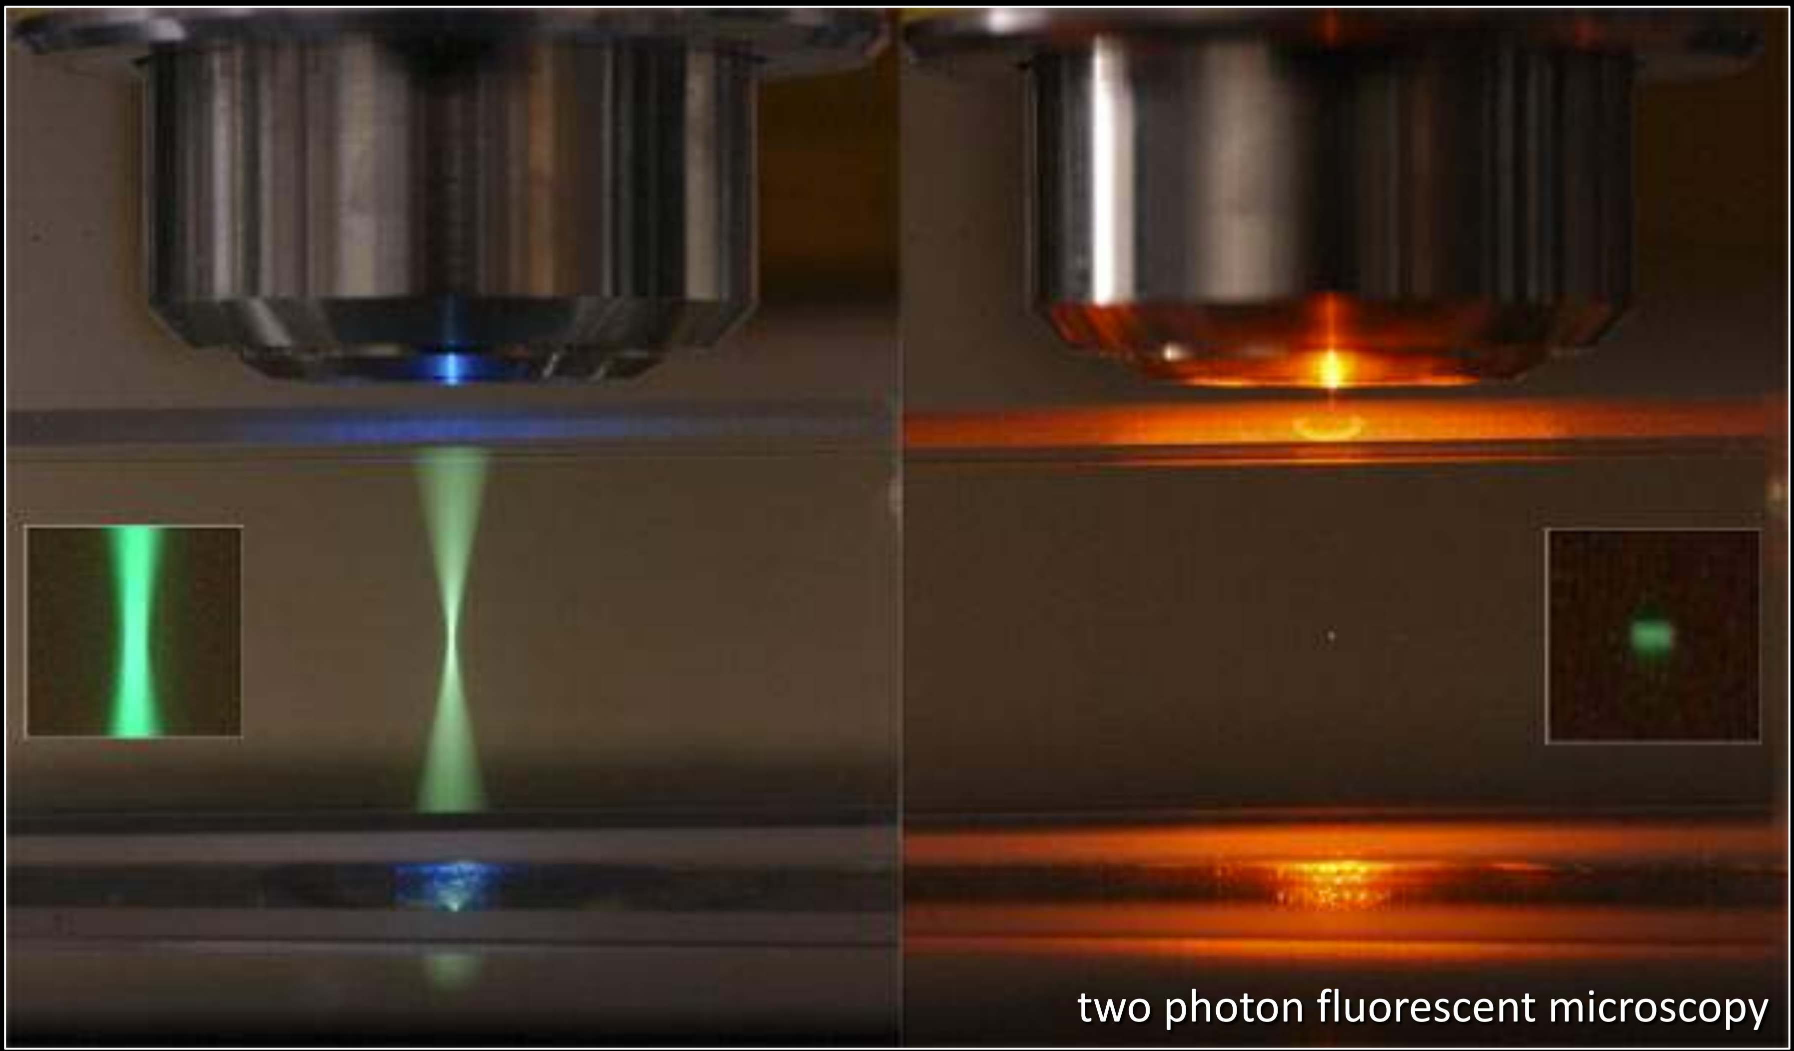

computational light transport

multiple light sources

sensor

observation:

- measurement under two light sources equals the sum of measurements taken under each source individually, i.e., T(x_1, x_2) = T(x_1, 0) + T(0, x_2)

property 2: additivity

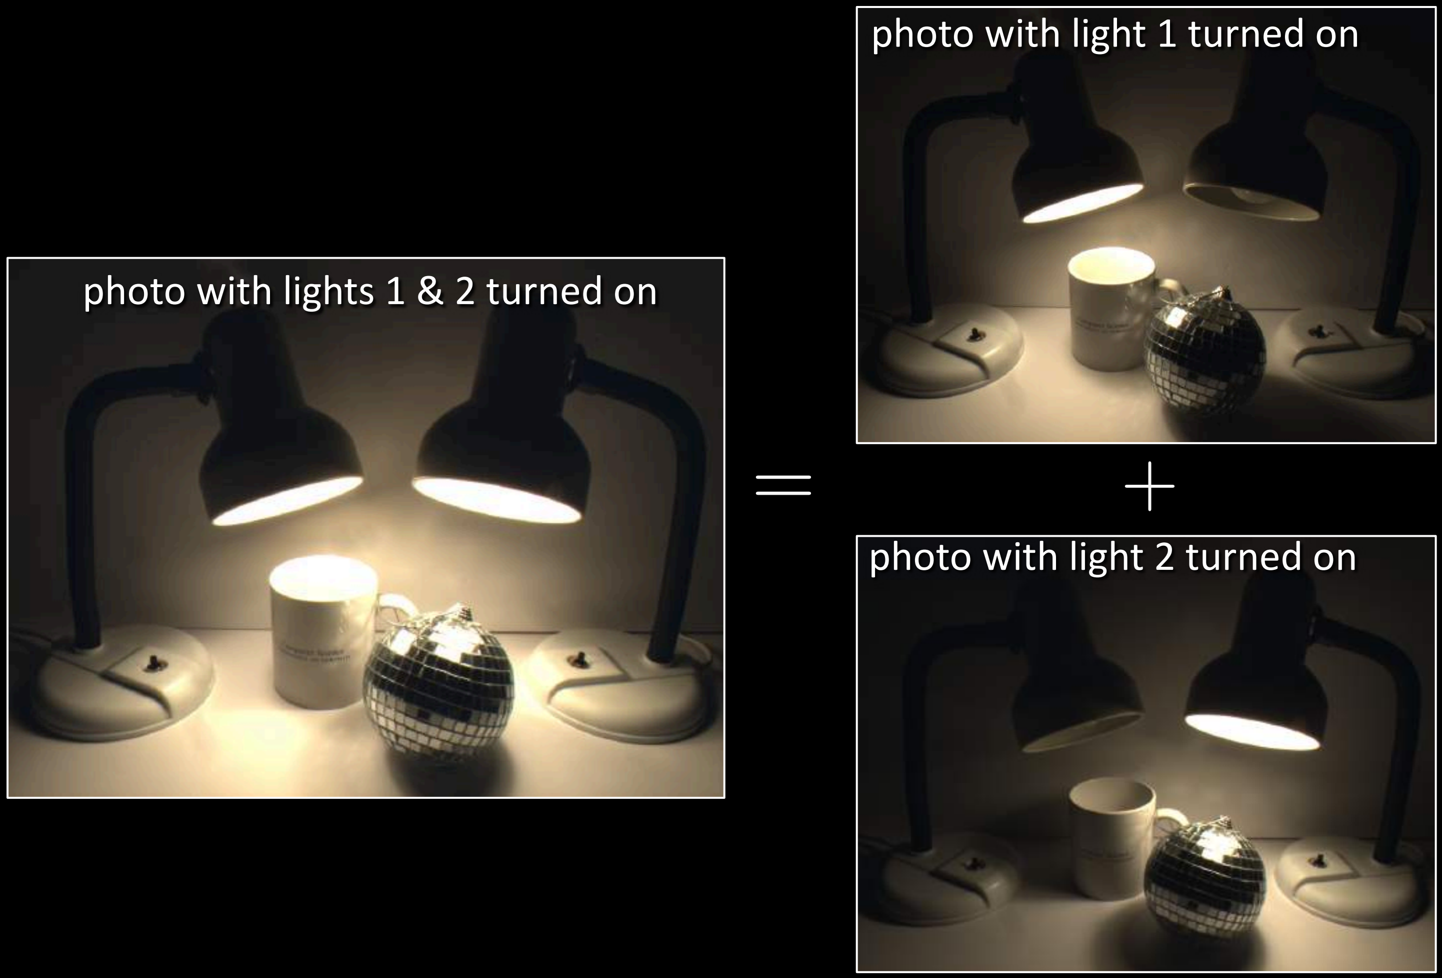

observation:

- measurement under two light sources equals the sum of measurements taken under each source individually, i.e., $T(x_1, x_2) = T(x_1, 0) + T(0, x_2)$

property 2: additivity

observation:

- measurement under two light sources equals the sum of measurements taken under each source individually, i.e., $T(x_1, x_2) = T(x_1, 0) + T(0, x_2)$

property 2: additivity

observation:

- measurement under two light sources equals the sum of measurements taken under each source individually, i.e., T(x_1, x_2) = T(x_1, 0) + T(0, x_2)

image-based relighting

=

image-based relighting

=

Weight 1

x 1

+

Weight 2

x 1

image-based relighting

=

Weight 1

x 1

+

Weight 2

x 0

image-based relighting

=

Weight 1

x

+

Weight 2

x

Relit photo $\mathbf{p}$
=
photo with light 1 turned on $\mathbf{T_1}$
Weight 1
x $\mathbf{l_1}$
+
photo with light 2 turned on $\mathbf{T_2}$
Weight 2
x $\mathbf{l_2}$

=

Weight 1

x $\mathbf{l_1}$ +

Weight 2

x $\mathbf{l_2}$

$$\mathbf{p} = \sum_{i=1}^{2} \mathbf{T}_i \times \mathbf{l}_i$$

=

Weight 1

x $\mathbf{l_1}$ +

Weight 2

x $\mathbf{l_2}$

$$n \quad \mathbf{p} = \sum_{i=1}^{2} \mathbf{T}_i \times \mathbf{l}_i$$

n pixel values

=

Weight 1

x $\mathbf{l_1}$ +

Weight 2

x $\mathbf{l_2}$

$$\mathbf{p} = \sum_{i=1}^{2} \mathbf{T}_i \times \mathbf{l}_i$$

n

n pixel values

=

Weight 1

x $\mathbf{l}_1$

+

Weight 2

x $\mathbf{l}_2$

$$n \quad \mathbf{p} = \sum_{i=1}^{2} \mathbf{T}_i \times \mathbf{l}_i$$

Contribution of the source

n pixel values

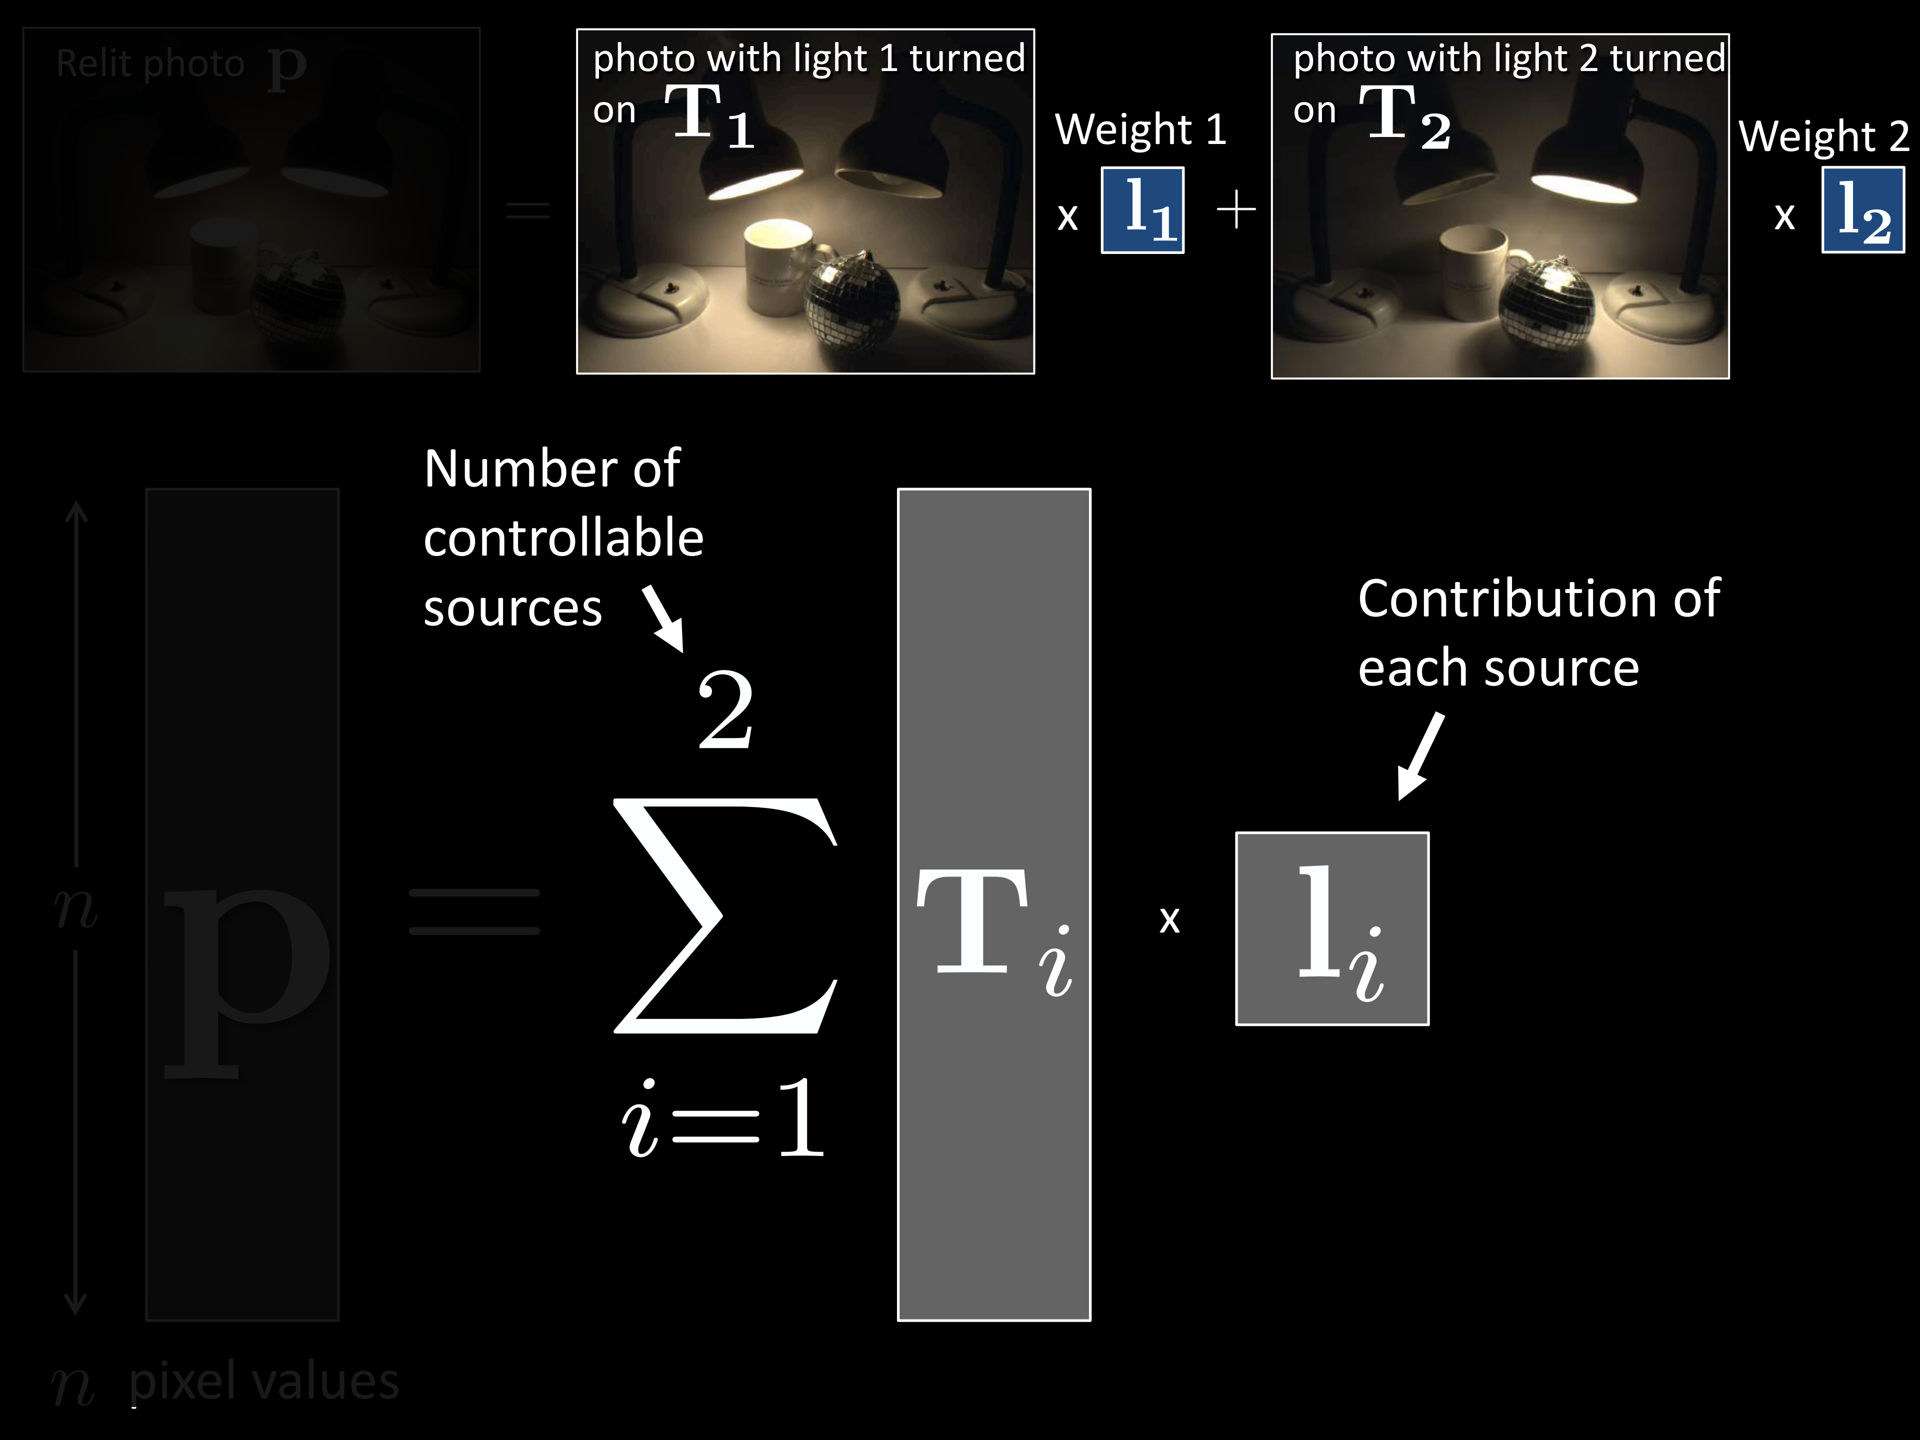

Relit photo p
photo with light 1 turned on $\mathbf{T_1}$
Weight 1
x $\mathbf{l_1}$
+
photo with light 2 turned on $\mathbf{T_2}$
Weight 2
x $\mathbf{l_2}$
Number of controllable sources
Contribution of each source
$\mathbf{p} = \sum_{i=1}^{2} \mathbf{T}_i \times \mathbf{l}_i$
n
n pixel values

$=$

Weight 1

x $\mathbf{l_1}$ $+$

Weight 2

x $\mathbf{l_2}$

Number of controllable sources

Contribution of each source

$$\mathbf{p} = \sum_{i=1}^{2} \mathbf{T}_i \times \mathbf{l}_i$$

n

n pixel values

$=$

Weight 1

x $\mathbf{l_1}$ $+$

Weight 2

x $\mathbf{l_2}$

Number of controllable sources

Contribution of each source

$$n \quad \mathbf{p} = \sum_{i=1}^{2} \mathbf{T}_i \times \mathbf{l}_i$$

n pixel values

$=$

Weight 1 x $\mathbf{l_1}$ $+$

Weight 2 x $\mathbf{l_2}$

Number of controllable sources

Contribution of each source

$$\mathbf{p} = \sum_{i=1}^{m} \mathbf{T}_i \times \mathbf{l}_i$$

n

n pixel values

$=$

Weight 1

x $\mathbf{l_1}$ $+$

Weight 2

x $\mathbf{l_2}$

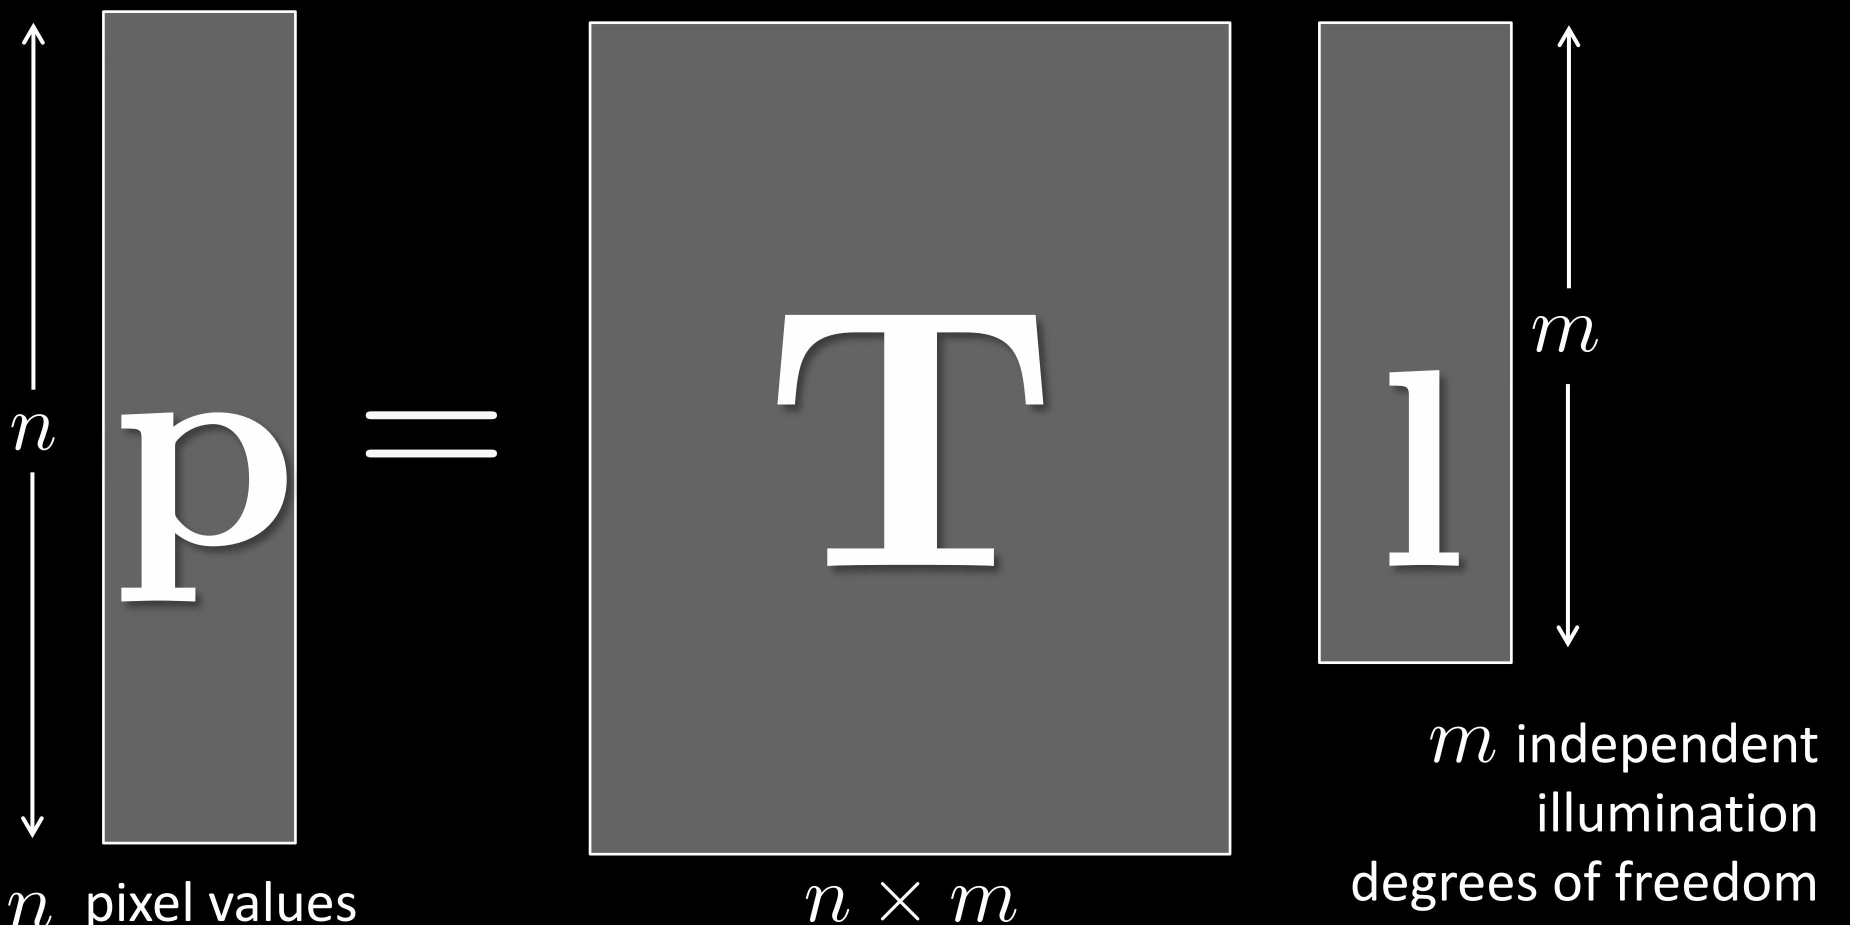

modeling light transport with color

light source emitting
radiant energy **x**
(measured in Joules)

with specific wavelength w

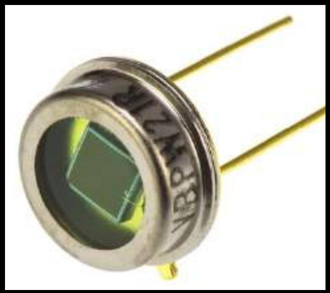

sensor detecting
radiant energy **y**
(measured in Joules)

with specific wavelength w'

bispectral BRDF / transport matrix

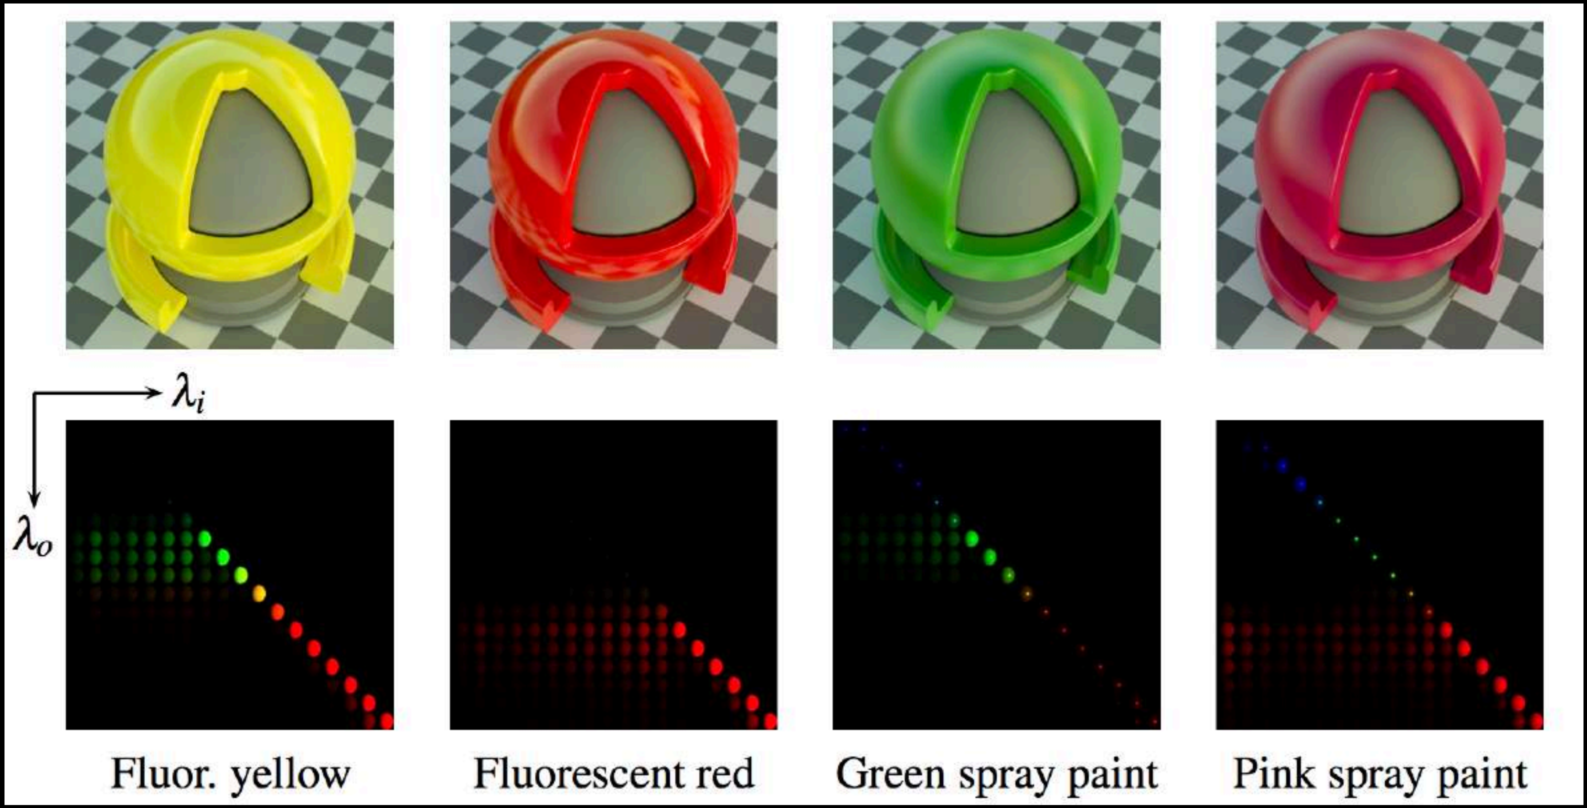

"Acquisition and Analysis of Bispectral Bidirectional Reflectance and Reradiation Distribution Functions", Hullin et al. 2010

modeling light transport **with polarization**

light source emitting
radiant energy **x**
(measured in Joules)

with specific polarization state

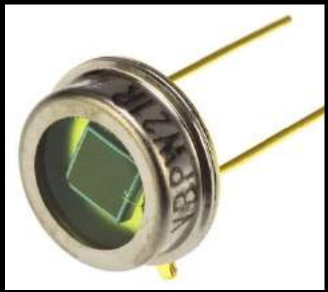

sensor detecting
radiant energy **y**
(measured in Joules)

with specific polarization state

modeling light transport **with polarization**

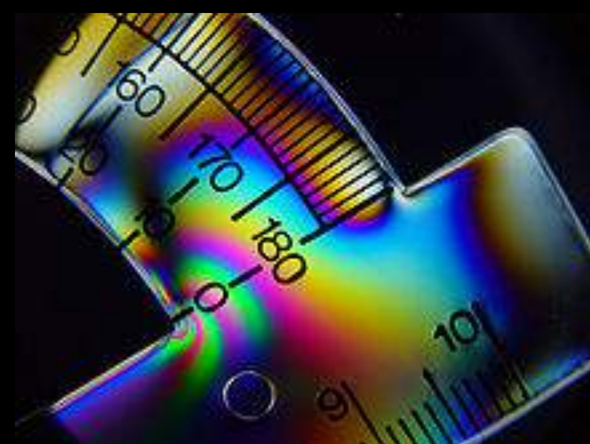

modeling light transport **with time**

light source emitting radiant energy **x** (measured in Joules)

at time t

$T()$

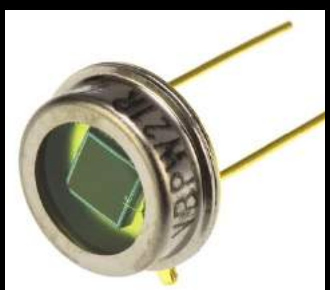

sensor detecting radiant energy **y** (measured in Joules)

at time t'

modeling light transport **with time**

[Raskar et al. 2011]

the light transport matrix

Sloan et al 02, Ng et al 03, Seitz et al 05, Sen et al 05, …

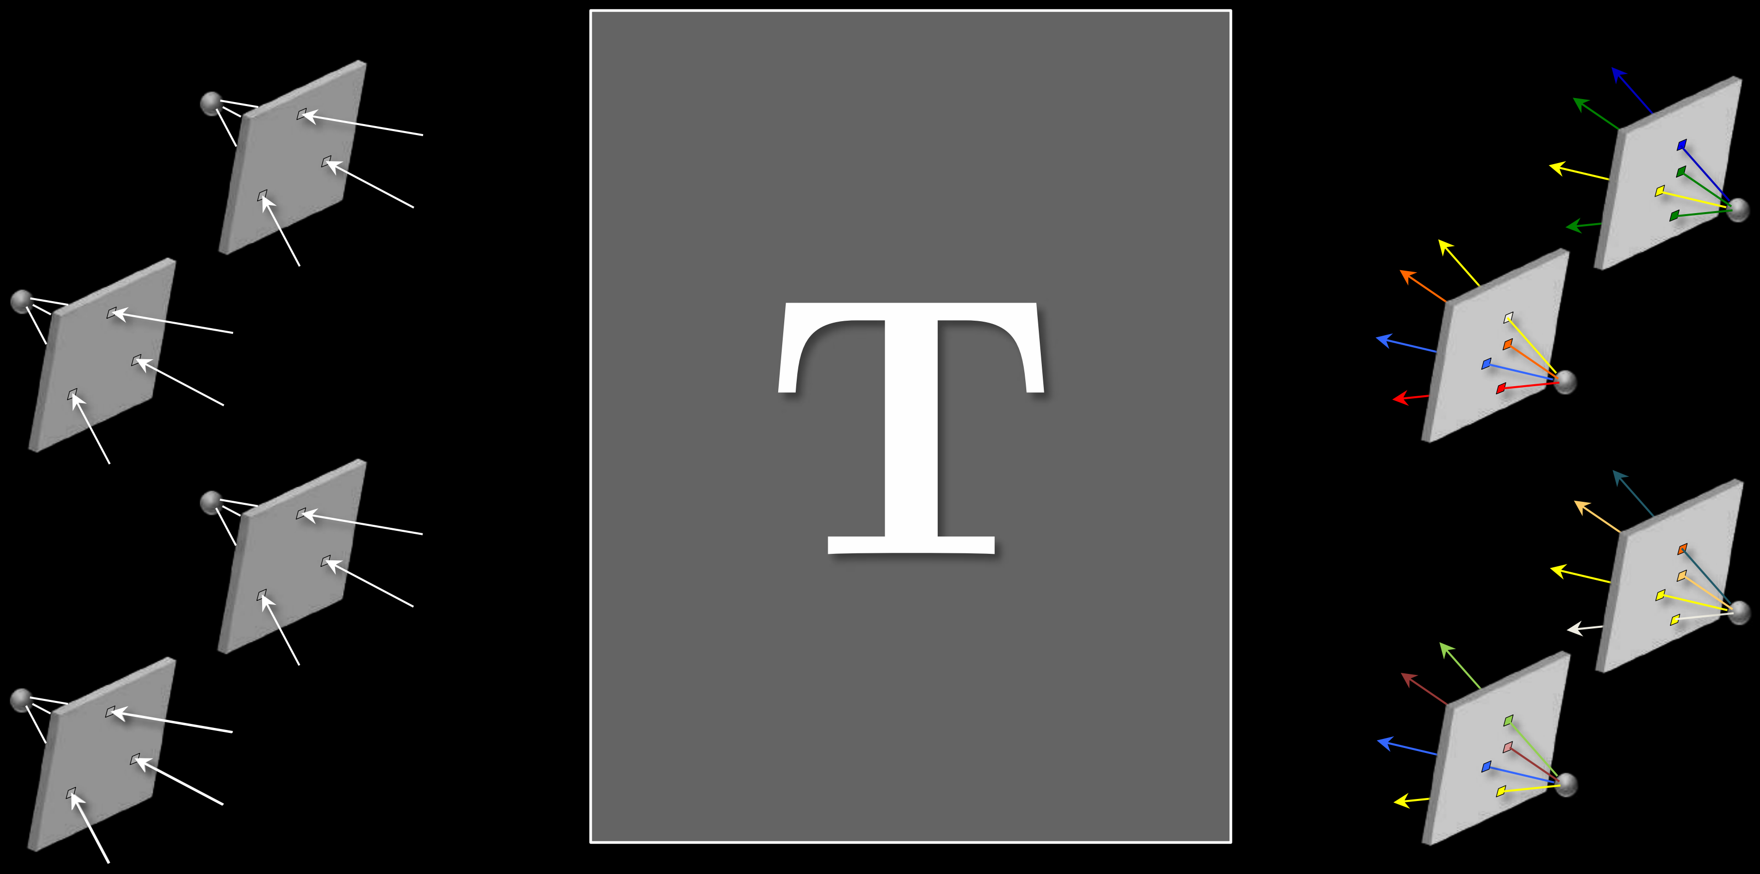

transport matrix represents the set of photos under all possible (controllable) lighting conditions

overview

1. **the light transport matrix:** a general model for the transfer of radiant energy

2. example transport matrices for real scenes

3. challenges associated with analyzing transport matrices

4. optical algorithms to analyze light transport

convex scene, diffuse reflectance, projector

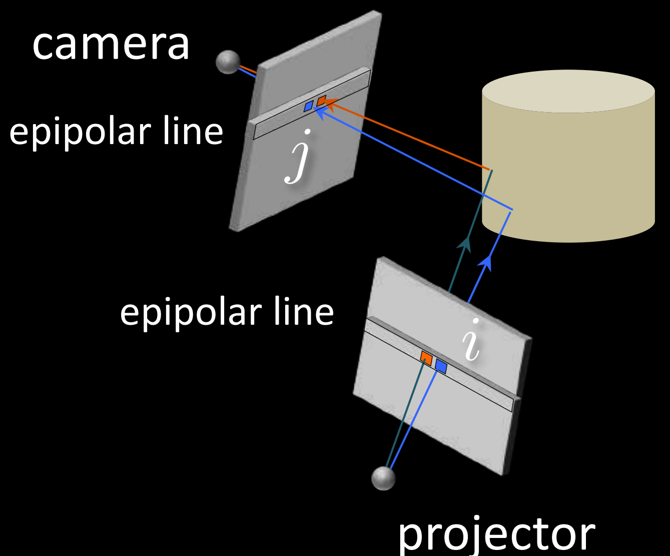

camera view

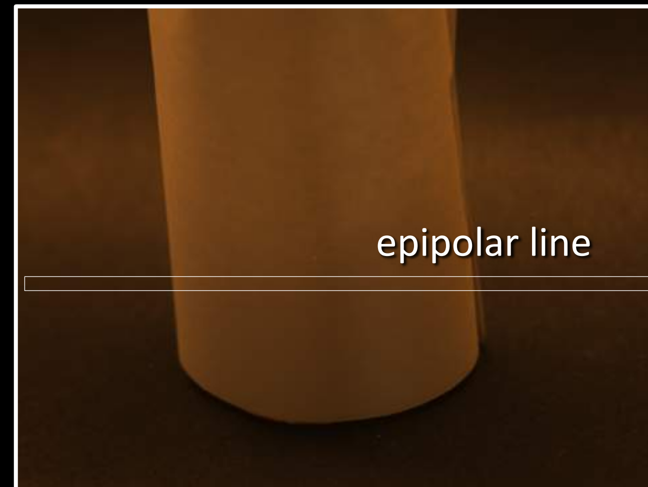

1 projector pixel on
ambient illumination

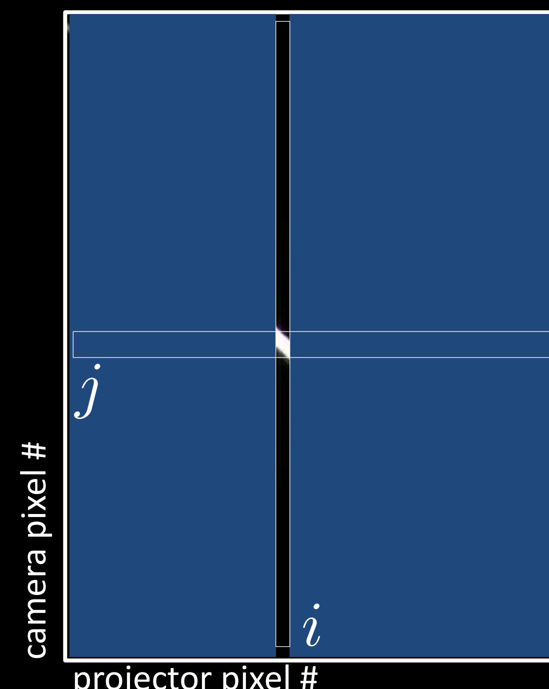

$\mathbf{T}$ (1 epipolar line)

$\mathbf{T} \Leftrightarrow$ stereo disparity map

acquiring $\mathbf{T} \Leftrightarrow$ structured-light 3D scanning

$\mathbf{T}$ sparse, high rank

convex scene, diffuse reflectance, point sources

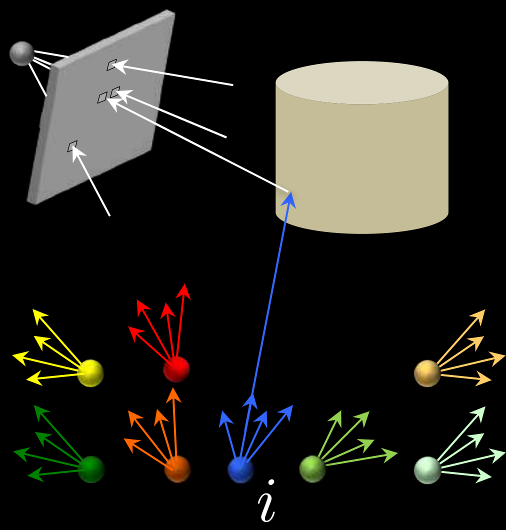

1 point source

ambient illumination

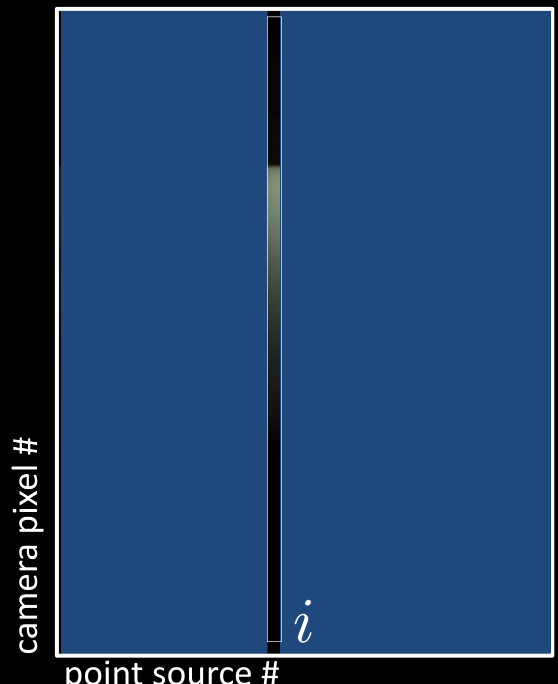

$\mathbf{T}$ (for 1 image row)

no shadows $\Rightarrow$

$rank(\mathbf{T}) = 3$ [Shashua, PhD 92]

attached shadows $\Rightarrow$

$rank(\mathbf{T}) \approx 9$ [Basri & Jacobs, PAMI 01]

analyzing $\mathbf{T} \Leftrightarrow$ photometric stereo

convex scene, specular reflectance, point sources

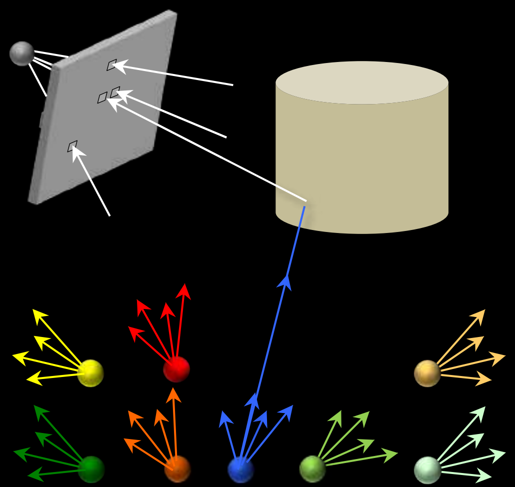

1 point source illumination

ambient illumination

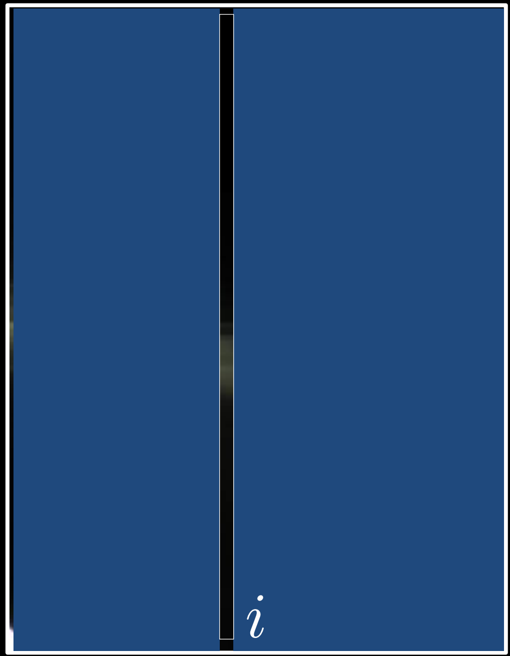

$\mathbf{T}$ (for 1 image row)

specular reflectance $\Rightarrow$
$\mathbf{T}$ can become full rank
[Ramamoorthi & Hanrahan, SIG 01]

analyzing $\mathbf{T} \Leftrightarrow$
shape-from-specularities
[Sanderson et al, PAMI 89]

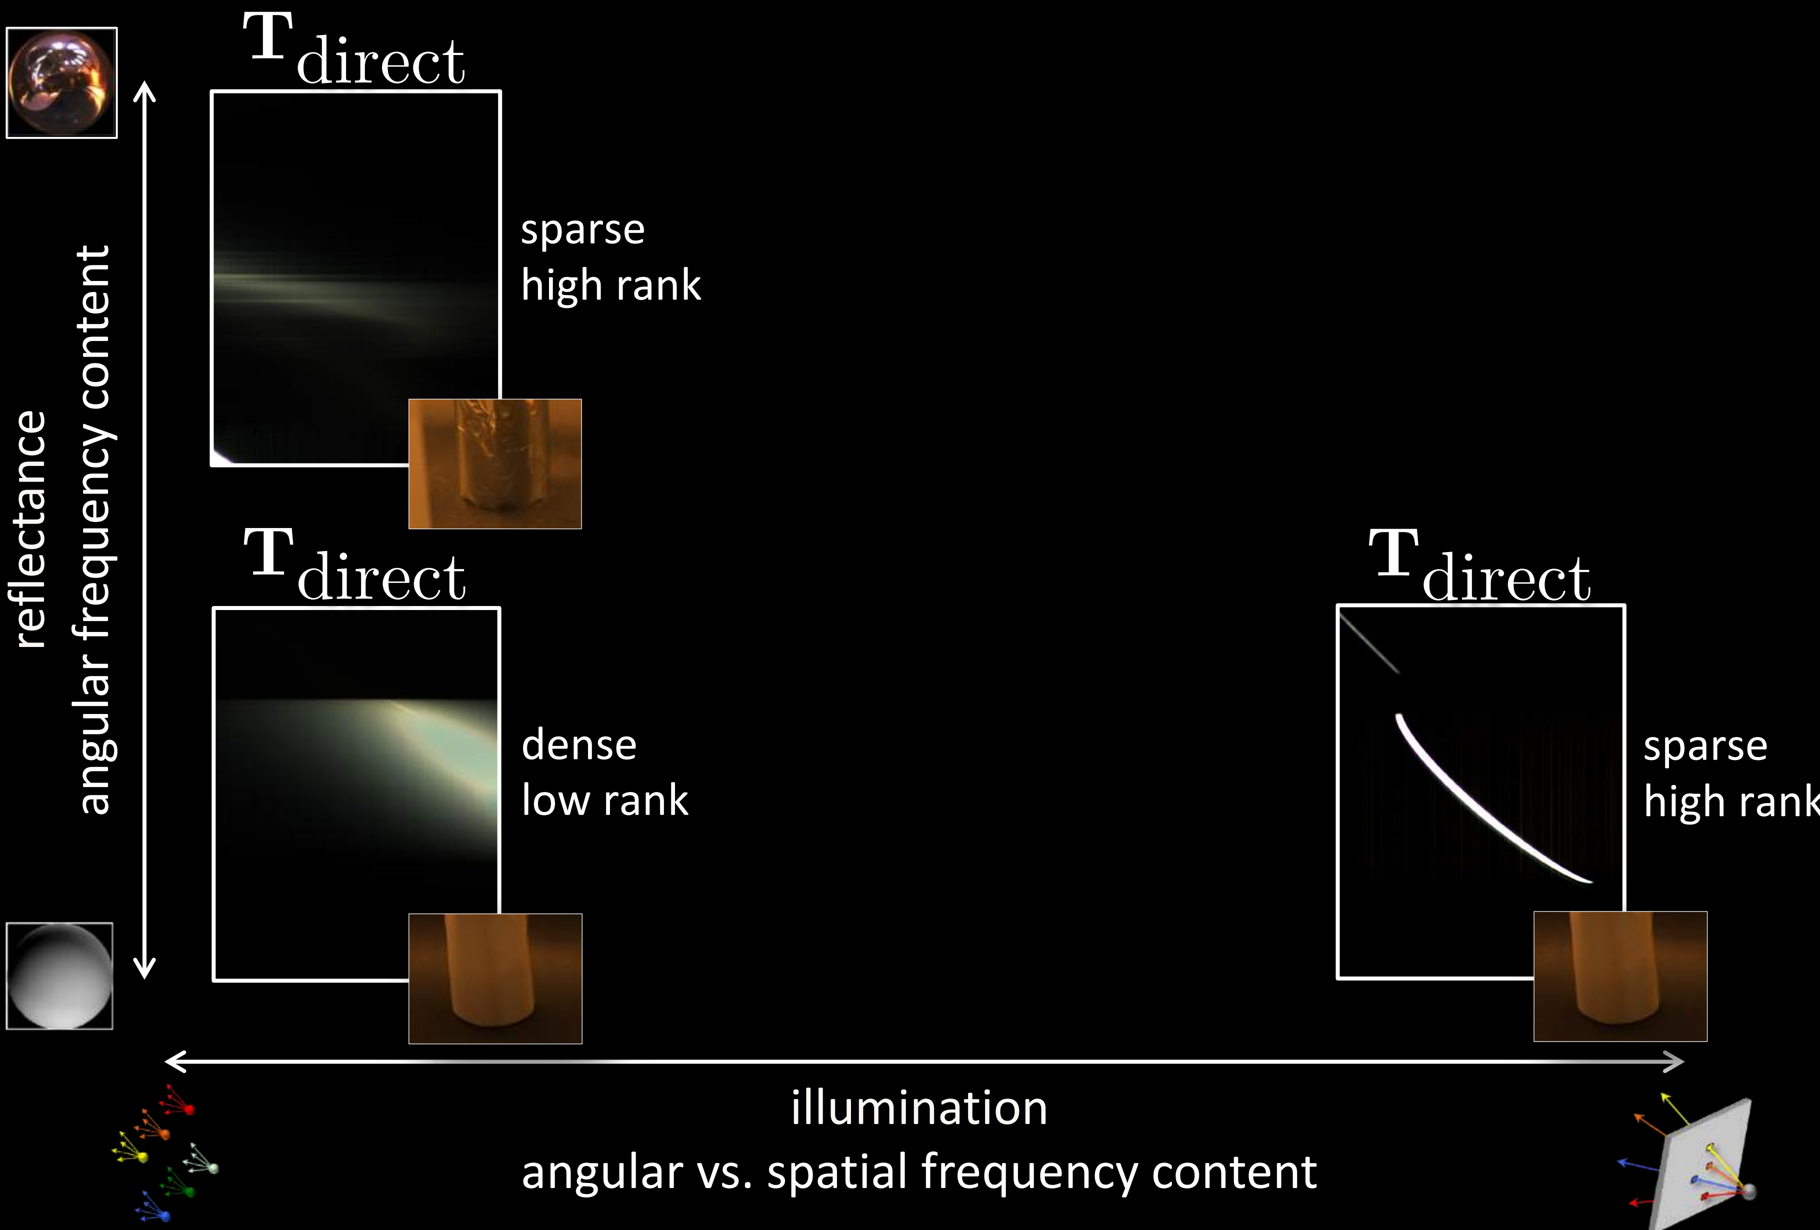

$\mathbf{T}_{\text{direct}}$
sparse
high rank
$\mathbf{T}_{\text{direct}}$
dense
low rank
$\mathbf{T}_{\text{direct}}$
sparse
high rank
reflectance
angular frequency content
illumination
angular vs. spatial frequency content

convex scene, diffuse reflectance, projector

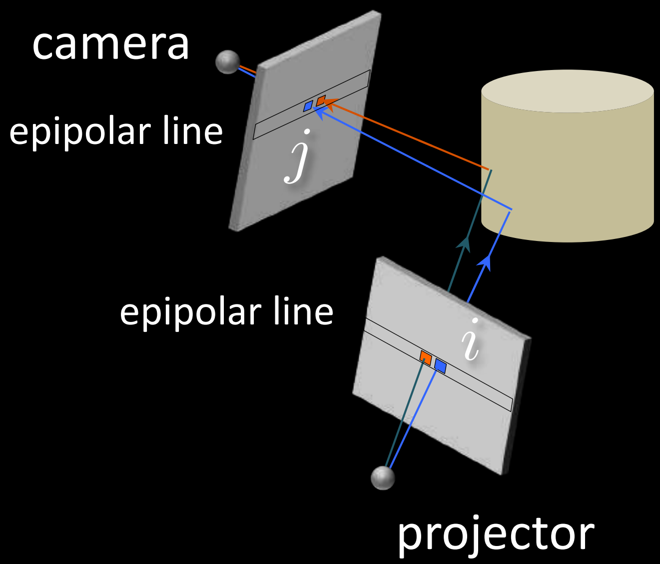

camera view

ambient illumination

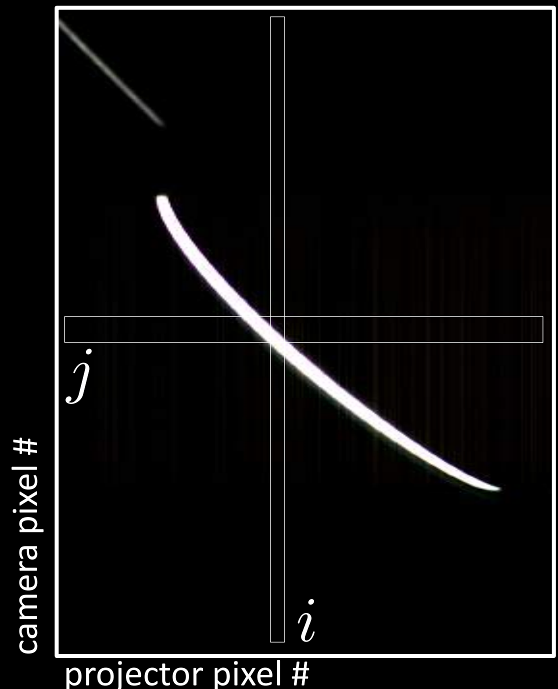

$\mathbf{T}$ (1 epipolar line)

convex scene, translucency, projector

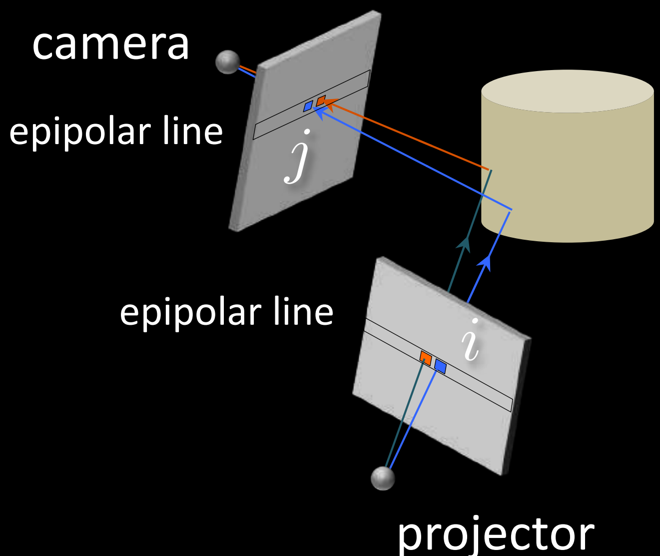

camera view

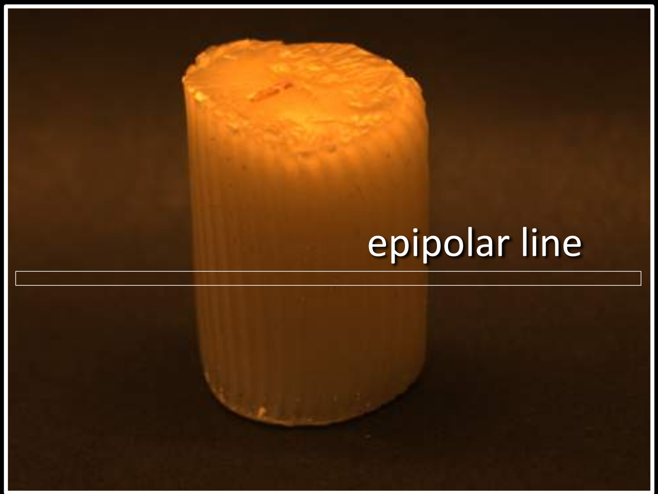

ambient illumination

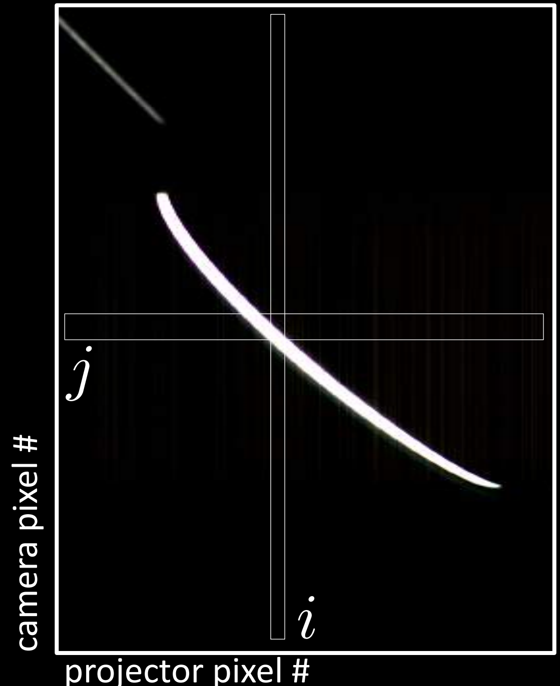

$\mathbf{T}$ (1 epipolar line)

convex scene, translucency, projector

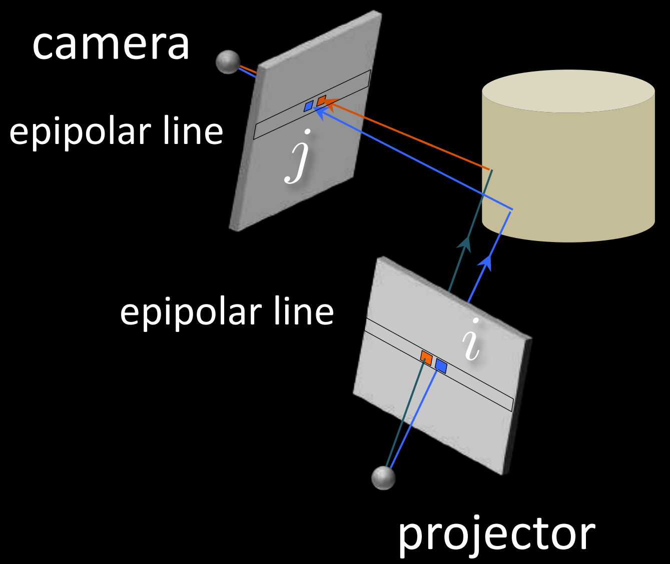

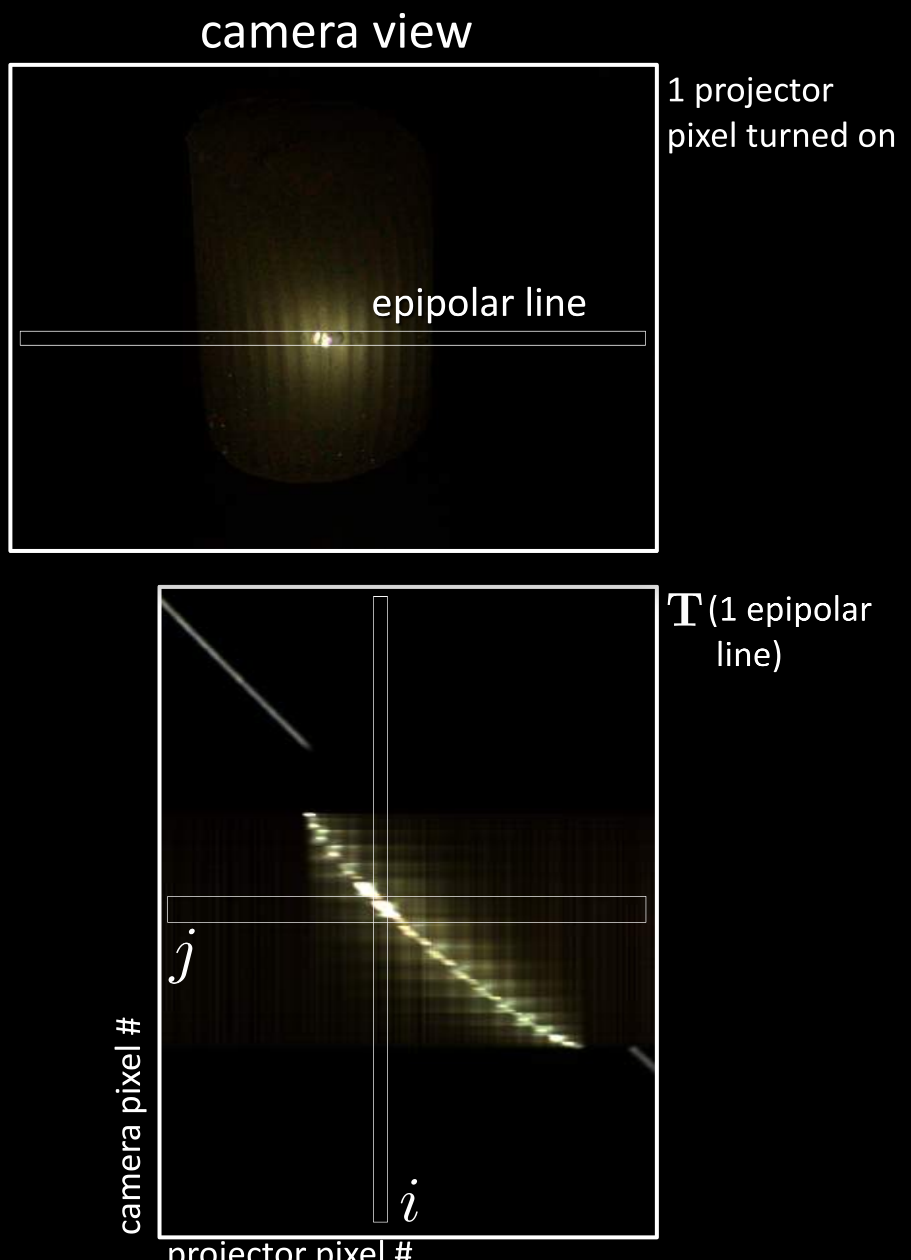

$$\mathbf{T} = \mathbf{T}_{\text{direct}} + \mathbf{T}_{\text{indirect}}$$

recovering $\mathbf{T}_{\text{direct}} \Leftrightarrow$ structured-light 3D scanning

$\mathbf{T}$ less sparse, usually high rank

convex scene, translucency, projector

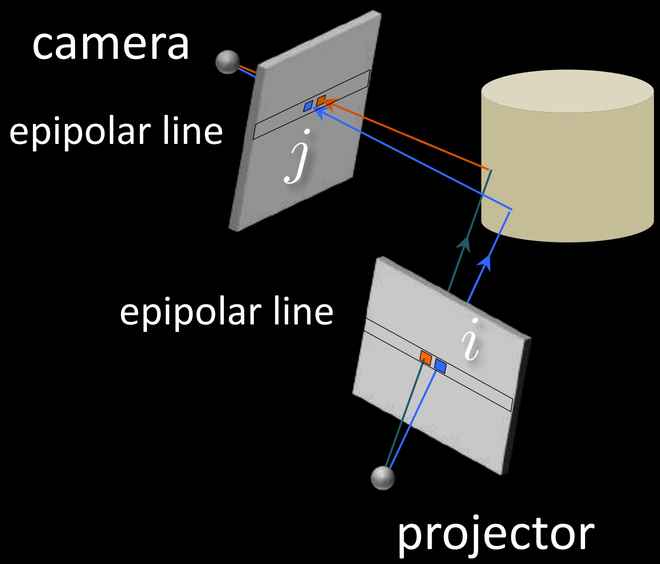

camera view

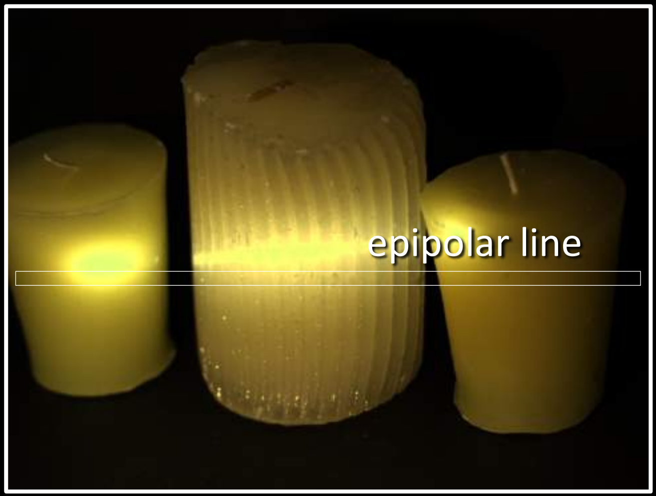

1 projector pixel turned on

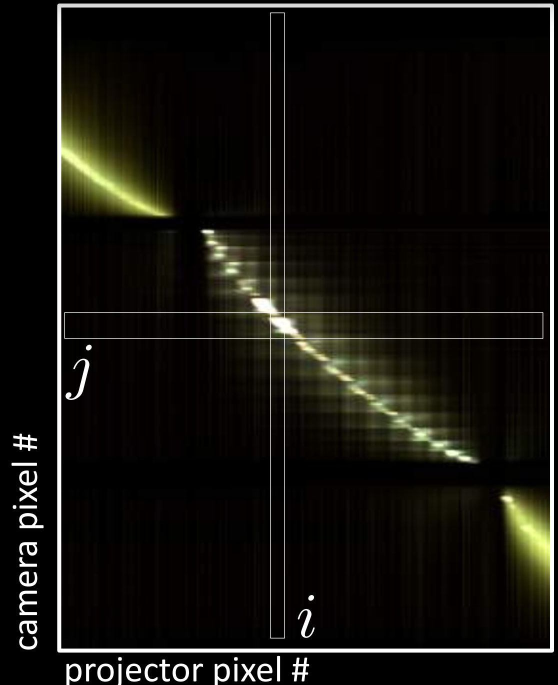

$\mathbf{T}$ (1 epipolar line)

$$\mathbf{T} = \mathbf{T}_{\text{direct}} + \mathbf{T}_{\text{indirect}}$$

recovering $\mathbf{T}_{\text{direct}} \Leftrightarrow$ structured-light 3D scanning

$\mathbf{T}$ less sparse, usually high rank

general scene, projector

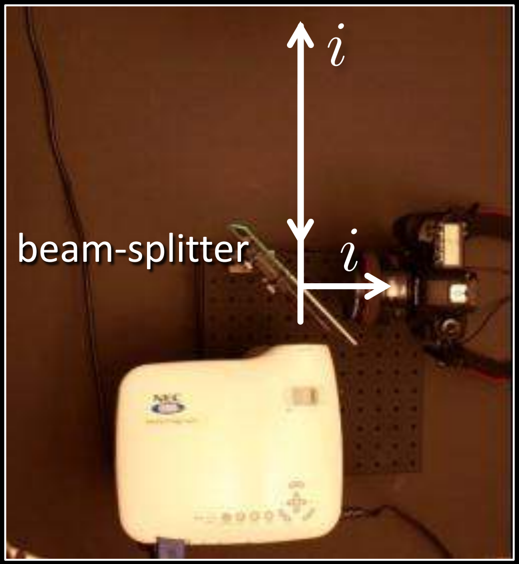

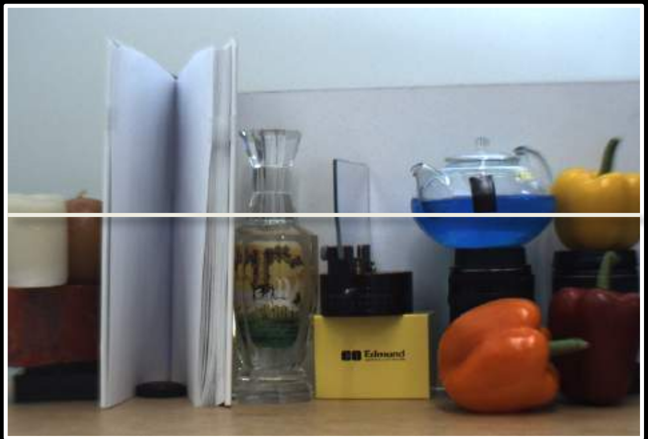

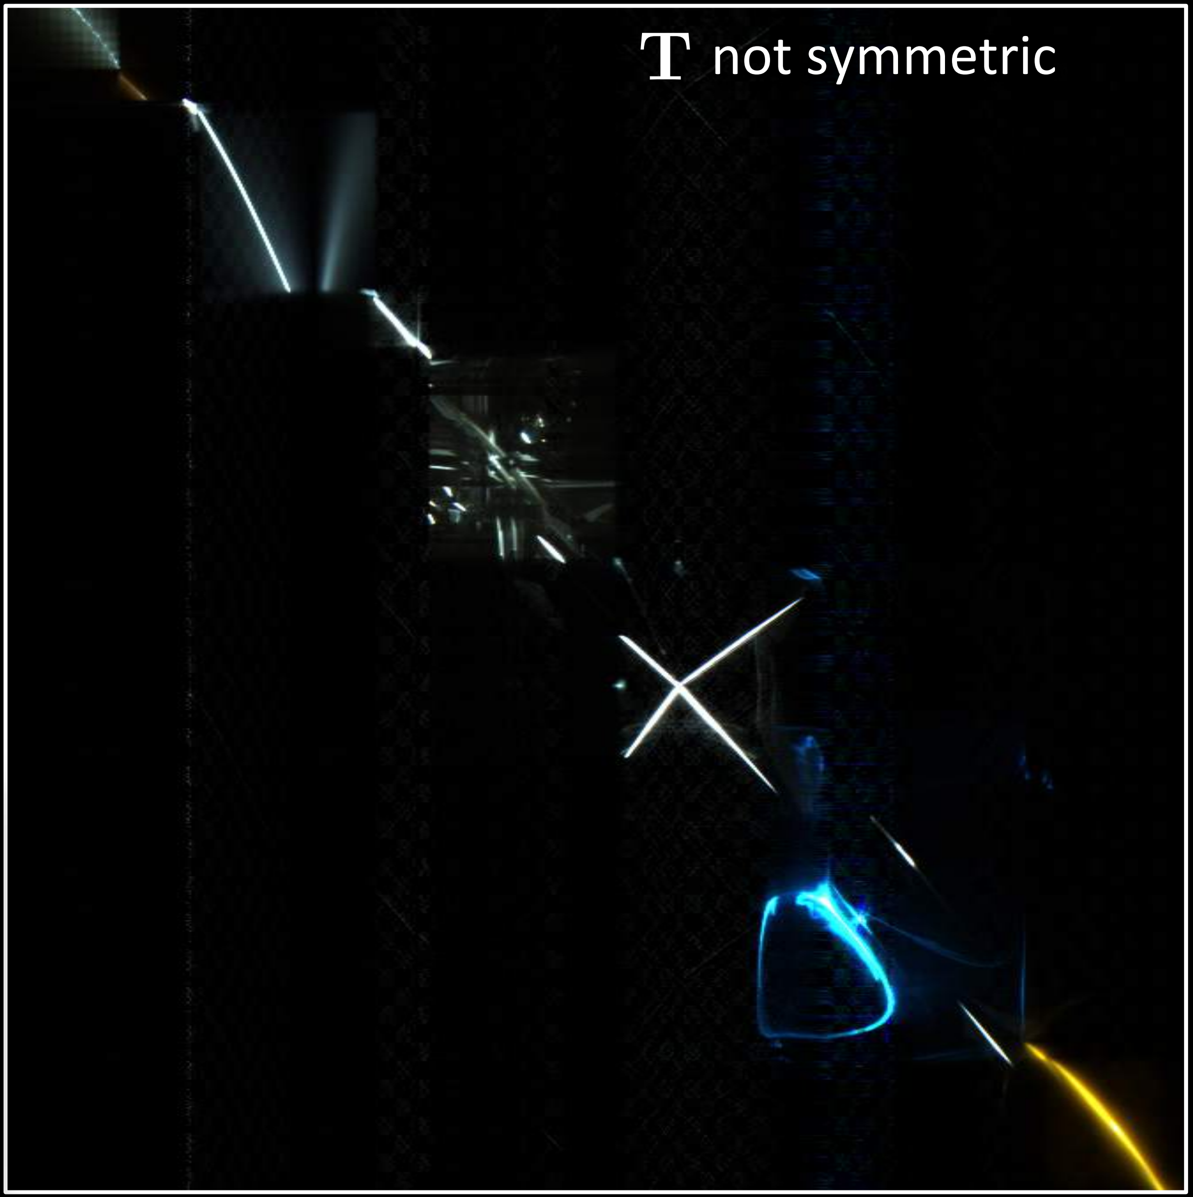

general scene, coaxial projector & camera

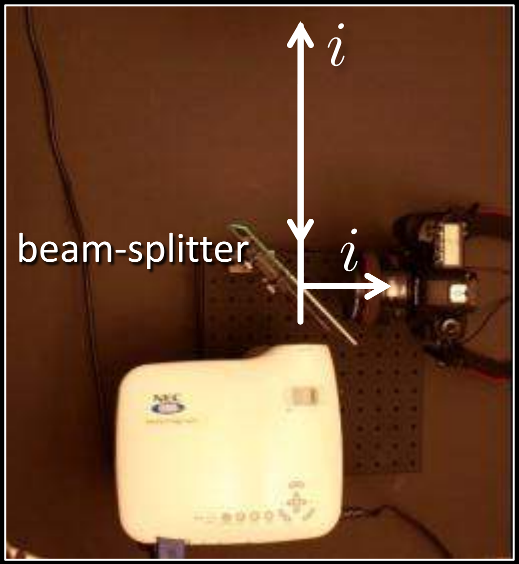

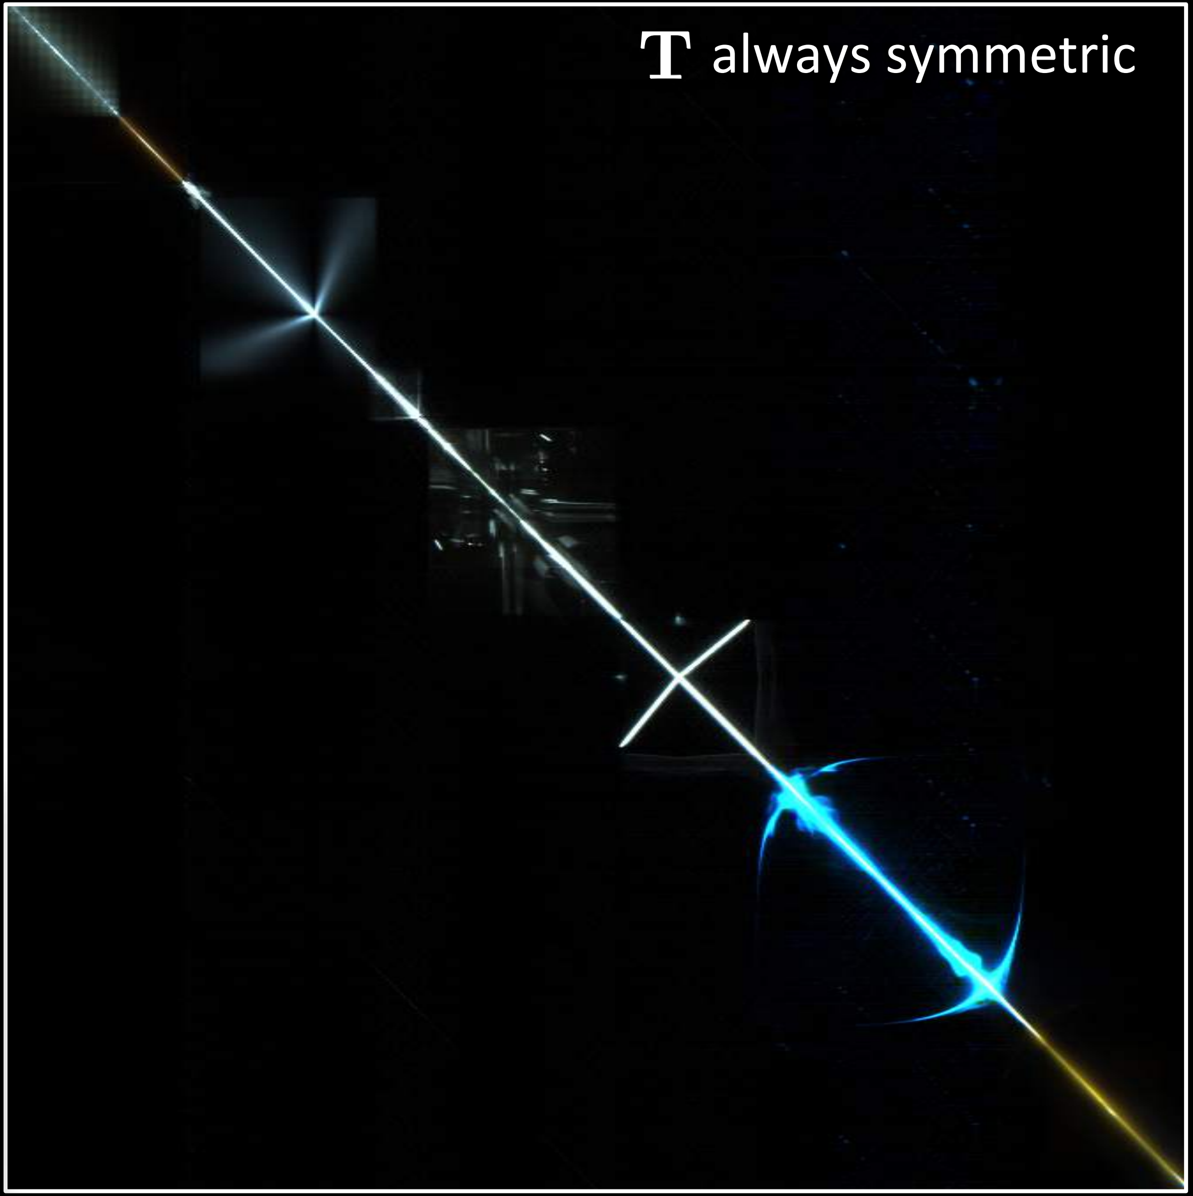

general scene, array of point sources

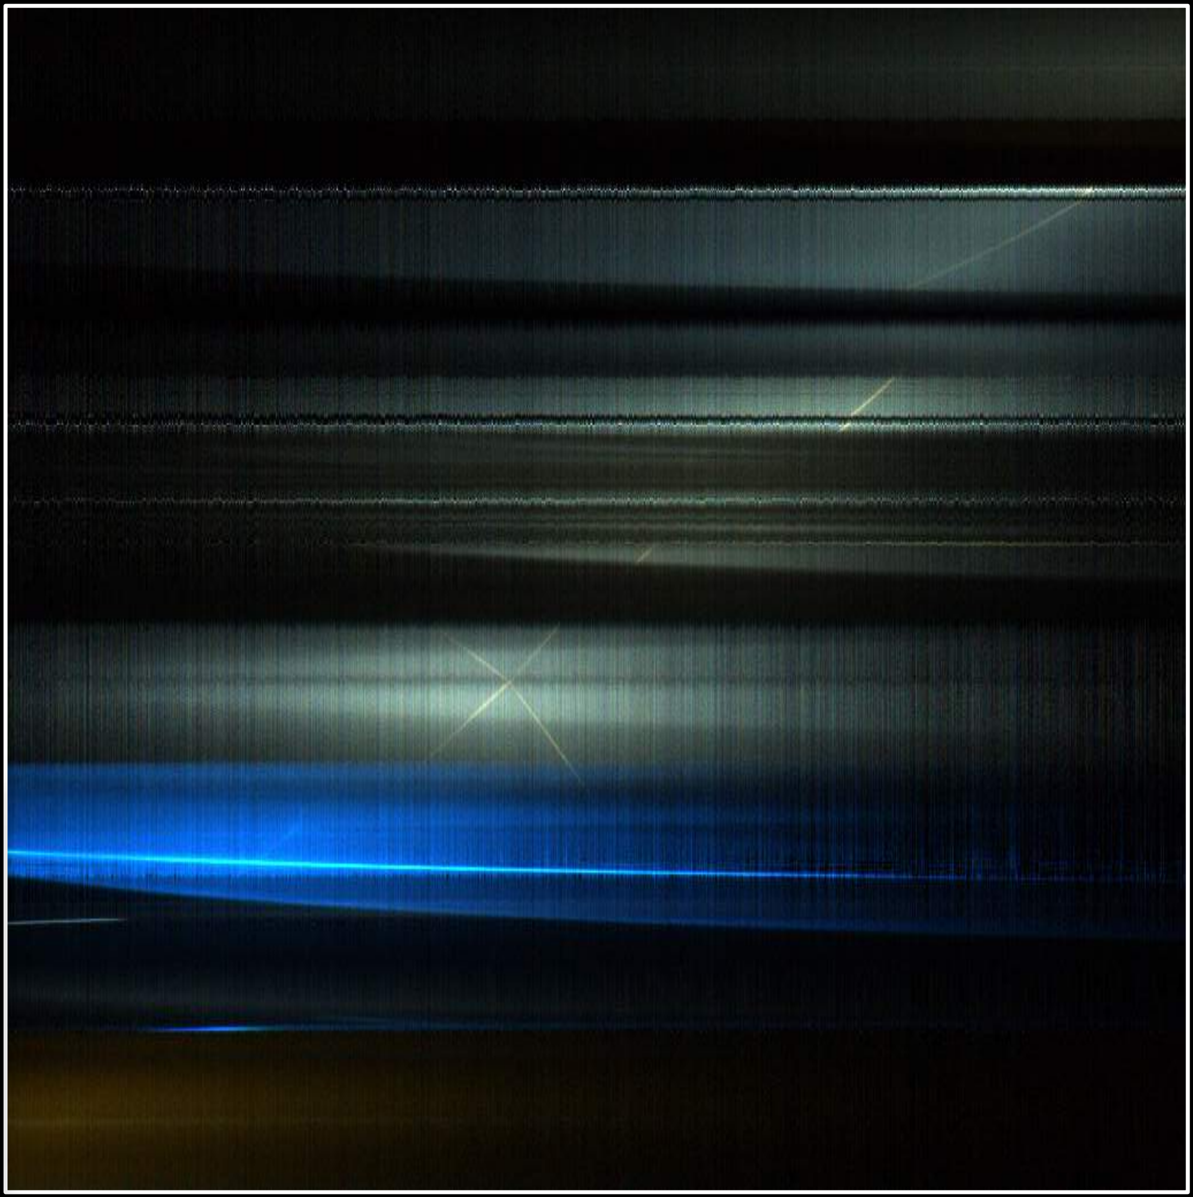

overview

1. **the light transport matrix:** a general model for the transfer of radiant energy

2. example transport matrices for real scenes

3. challenges associated with analyzing transport matrices

4. optical algorithms to analyze light transport

basic matrix properties

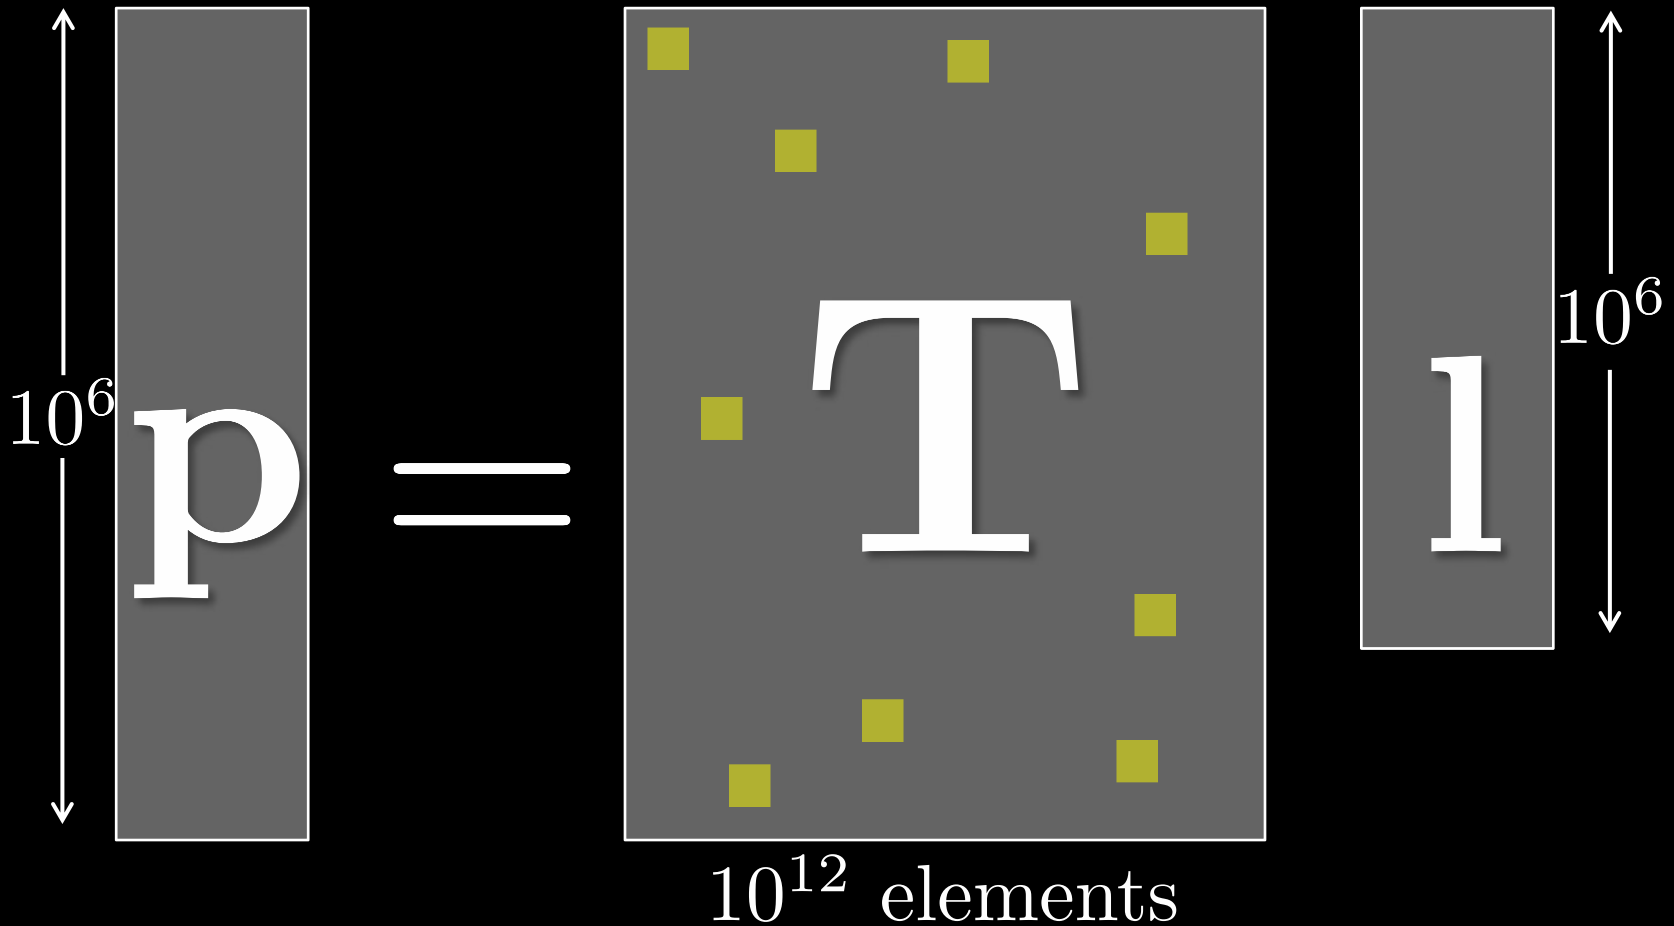

unknown & extremely large

no random access to its elements

relation to scene geometry & reflectance can be complex

overview

1. **the light transport matrix:** a general model for the transfer of radiant energy

2. example transport matrices for real scenes

3. challenges associated with analyzing transport matrices

4. optical algorithms to analyze light transport

Computing with Light

Key idea: analyze the transport matrix by implementing iterative numerical algorithms directly in optics

numerical domain

transport matrix

$$\mathbf{p} = \mathbf{T}\,\mathbf{l}$$

photo

illumination pattern

$\Longleftrightarrow$

optical domain

Computing with Light

Key idea: analyze the transport matrix by implementing iterative numerical algorithms directly in optics

numerical domain

transport matrix

$$\mathbf{p} = \mathbf{T}\,\mathbf{l}$$

photo

illumination pattern

$\Longleftrightarrow$

optical domain

Computing with Light

Key idea: analyze the transport matrix by implementing iterative numerical algorithms directly in optics

numerical domain

function analyze($\mathbf{T}$)

. . .

for $i = 1$ to k {

. . .

$\mathbf{p}_i = \mathbf{T}\mathbf{l}_i$

. . .

$\mathbf{d}_i = \mathbf{T}\mathbf{r}_i$

. . .

}

. . .

return result

$\Longleftrightarrow$

optical domain

Computing with Light

Key idea: analyze the transport matrix by implementing iterative numerical algorithms directly in optics

numerical domain

$\textbf{function}$ analyze($\mathbf{T}$)

$\ldots$

$\textbf{for}\ i = 1 \text{ to } k\ \{$

$\quad\ldots$

$\quad\mathbf{p}_i = \mathbf{T}\mathbf{l}_i$

$\quad\ldots$

$\quad\mathbf{d}_i = \mathbf{T}\mathbf{r}_i$

$\quad\ldots$

$\}$

$\ldots$

$\textbf{return}$ result

$\Longleftrightarrow$

optical domain

$\textbf{function}$ analyze()

$\ldots$

$\textbf{for}\ i = 1 \text{ to } k\ \{$

$\quad\ldots$

$\quad$project $\mathbf{l}_i$, capture $\mathbf{p}_i$

$\quad\ldots$

$\quad$project $\mathbf{r}_i$, capture $\mathbf{d}_i$

$\quad\ldots$

$\}$

$\ldots$

$\textbf{return}$ result

Computing Transport Eigenvectors

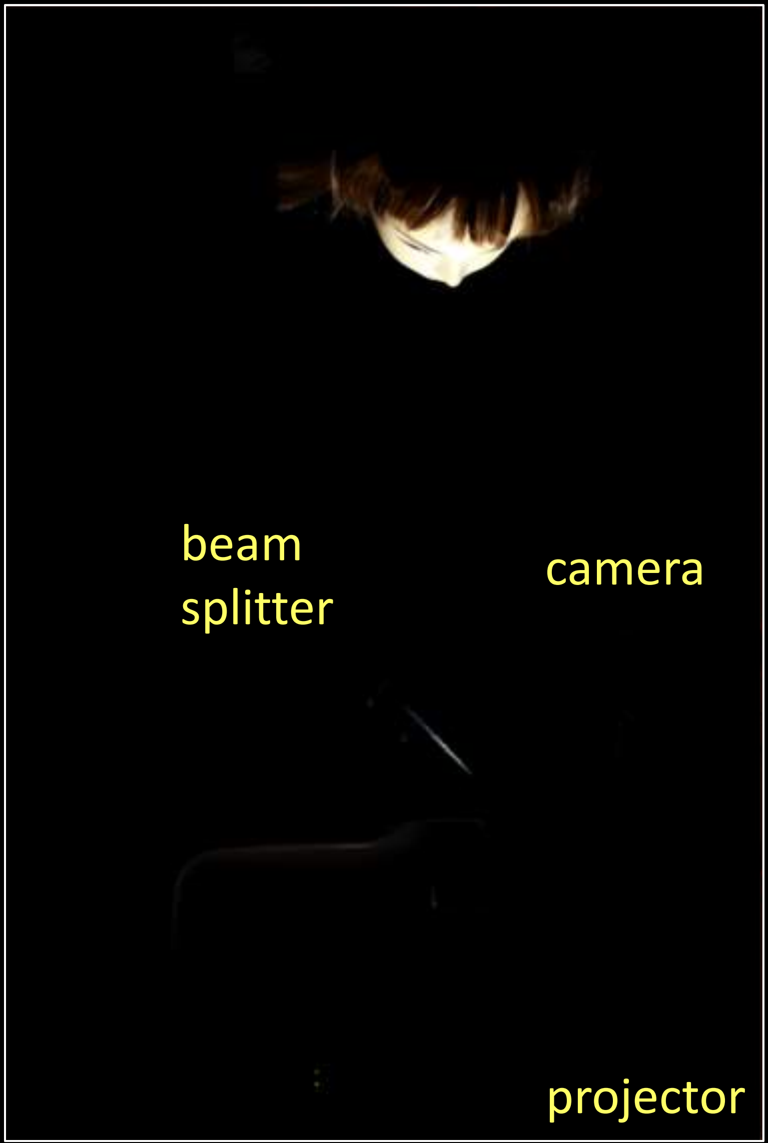

Eigenvector of a square matrix T

when projected onto scene,
we get the same photo back
(multiplied by a scalar)

project

capture

Numerical goal

find $\mathbf{l}, \lambda$ such that

$$\mathbf{T}\mathbf{l} = \lambda\mathbf{l}$$

and λ is maximal

Optical Power Iteration

Goal: find principal eigenvector of $\mathbf{T}$

Observation: it is a fixed point of the sequence $\mathbf{l}, \mathbf{T}\mathbf{l}, \mathbf{T}^2\mathbf{l}, \mathbf{T}^3\mathbf{l}, \ldots$

numerical domain

function $\mathrm{PowerIt}(\mathbf{T})$

$\mathbf{l}_1 = \text{initial vector}$

for $i = 1$ to k {

$\mathbf{p}_i = \mathbf{T}\mathbf{l}_i$

$\mathbf{l}_{i+1} = \mathbf{p}_i / \|\mathbf{p}_i\|_2$

}

return $\mathbf{l}_{i+1}$

Properties

- linear convergence [Trefethen and Bau 1997]
- eigenvalues must be distinct
- $\mathbf{l}_1$ cannot be orthogonal to principal eigenvector

Optical Power Iteration

Goal: find principal eigenvector of $\mathbf{T}$

Observation: it is a fixed point of the sequence $\mathbf{l}, \mathbf{Tl}, \mathbf{T}^2\mathbf{l}, \mathbf{T}^3\mathbf{l}, \ldots$

numerical domain

$\textbf{function}\ \mathrm{PowerIt}(\mathbf{T})$

$\mathbf{l}_1 = \text{initial vector}$

$\textbf{for}\ i = 1\ \text{to}\ k\ \{$

$\quad \mathbf{p}_i = \mathbf{T}\mathbf{l}_i$

$\quad \mathbf{l}_{i+1} = \mathbf{p}_i / \|\mathbf{p}_i\|_2$

$\}$

$\textbf{return}\ \mathbf{l}_{i+1}$

$\Longleftrightarrow$

optical domain

$\textbf{function}\ \mathrm{PowerIt}()$

$\mathbf{l}_1 = \text{initial vector}$

$\textbf{for}\ i = 1\ \text{to}\ k\ \{$

$\quad \text{project}\ \mathbf{l}_i,\ \text{capture}\ \mathbf{p}_i$

$\quad \mathbf{l}_{i+1} = \mathbf{p}_i / \|\mathbf{p}_i\|_2$

$\}$

$\textbf{return}\ \mathbf{l}_{i+1}$

Optical Power Iteration

Goal: find principal eigenvector of $\mathbf{T}$

Observation: it is a fixed point of the sequence $\mathbf{l}, \mathbf{T}\mathbf{l}, \mathbf{T}^2\mathbf{l}, \mathbf{T}^3\mathbf{l}, \ldots$

numerical domain

function $\text{PowerIt}(\mathbf{T})$

$\mathbf{l}_1 = \text{initial vector}$

for $i = 1$ to k {

$\quad \mathbf{p}_i = \mathbf{T}\mathbf{l}_i$

$\quad \mathbf{l}_{i+1} = \mathbf{p}_i / \|\mathbf{p}_i\|_2$

}

return $\mathbf{l}_{i+1}$

$\Longleftrightarrow$

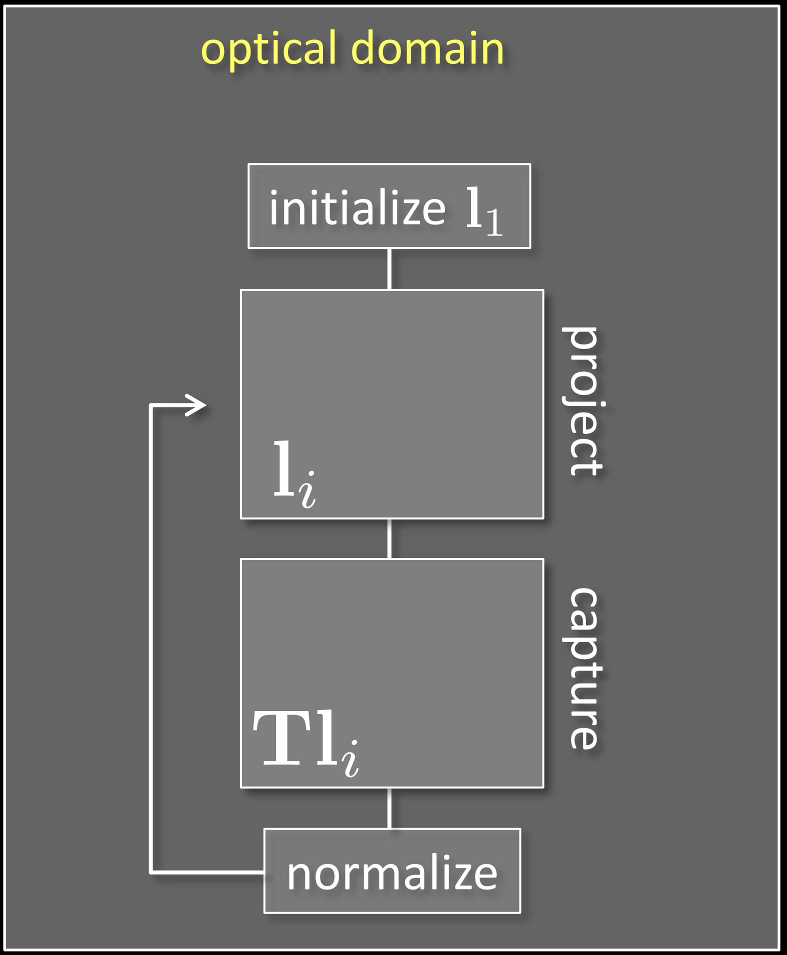

Optical Power Iteration

Goal: find principal eigenvector of $\mathbf{T}$

Observation: it is a fixed point of the sequence $\mathbf{l}, \mathbf{T}\mathbf{l}, \mathbf{T}^2\mathbf{l}, \mathbf{T}^3\mathbf{l}, \ldots$

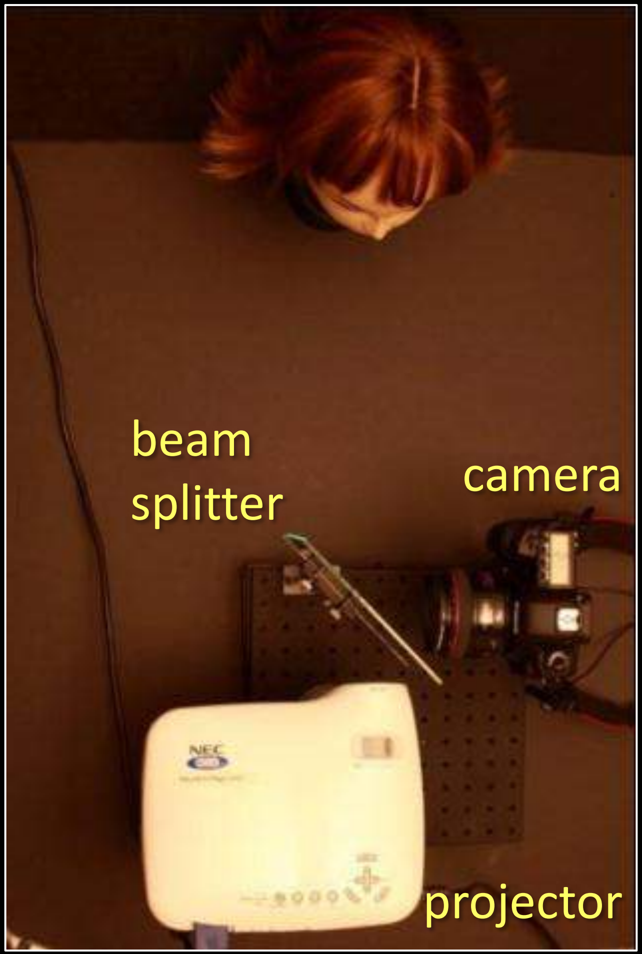

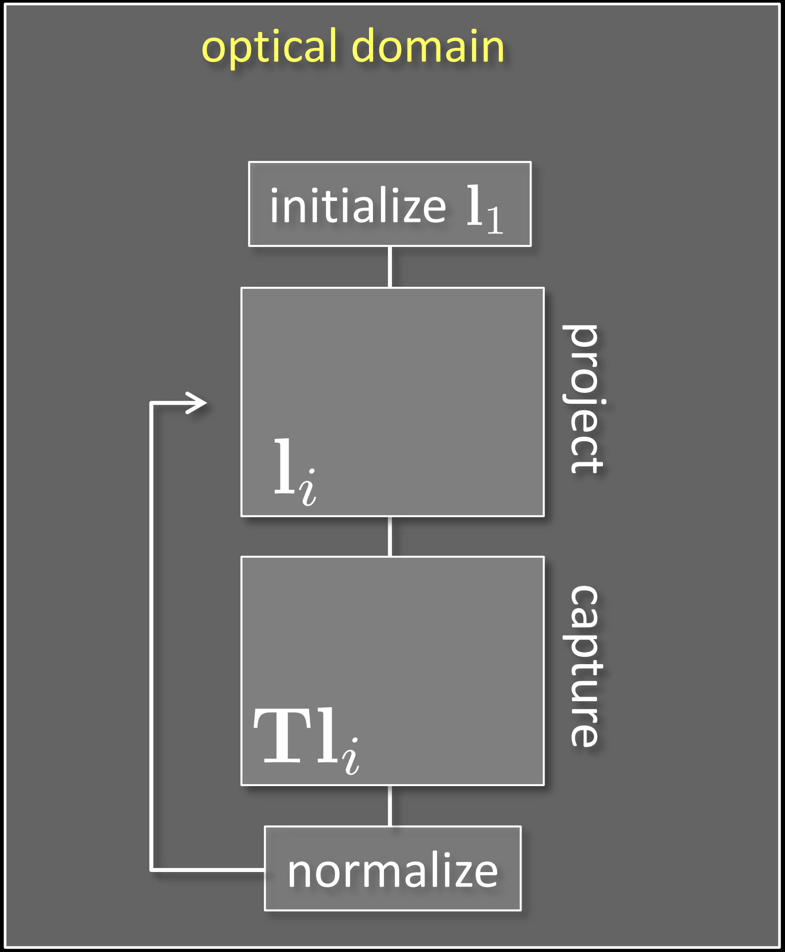

Optical Power Iteration

Goal: find principal eigenvector of $\mathbf{T}$

Observation: it is a fixed point of the sequence $\mathbf{l}, \mathbf{T}\mathbf{l}, \mathbf{T}^2\mathbf{l}, \mathbf{T}^3\mathbf{l}, \ldots$

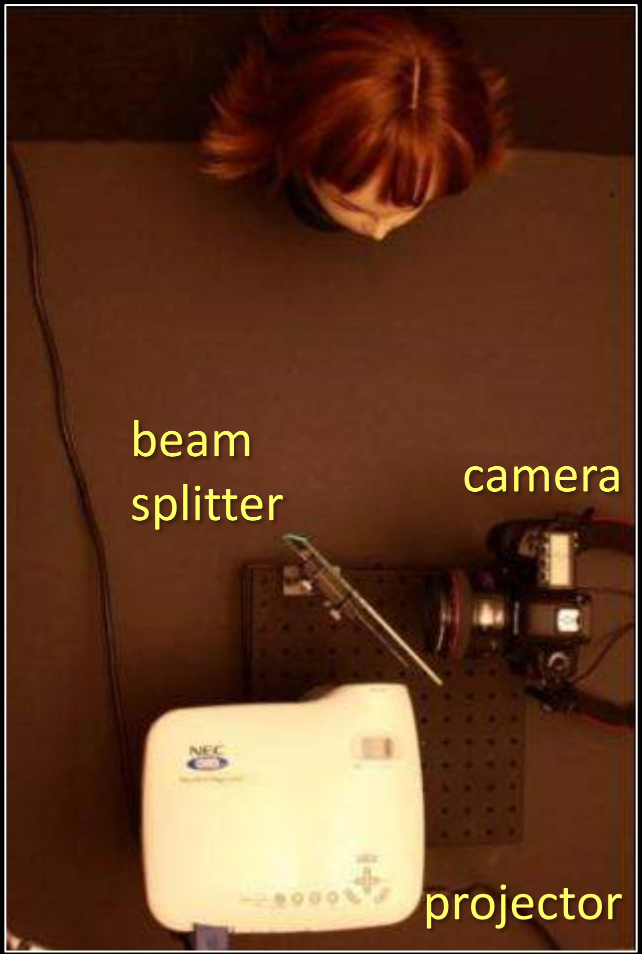

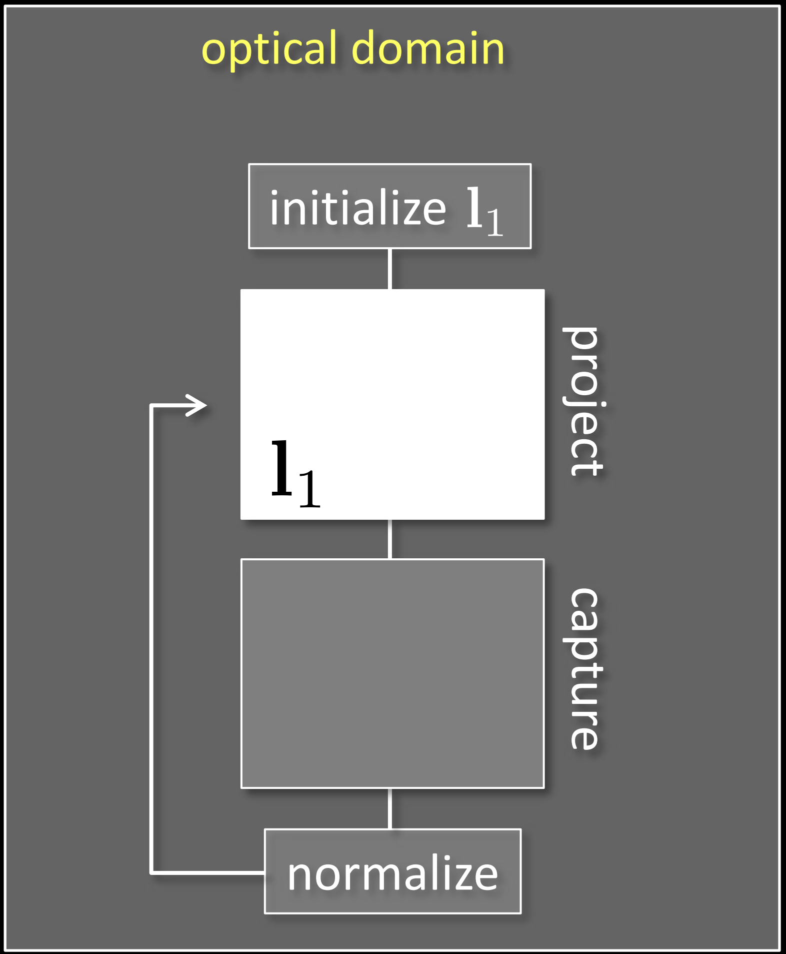

Optical Power Iteration

Goal: find principal eigenvector of $\mathbf{T}$

Observation: it is a fixed point of the sequence $\mathbf{l}, \mathbf{T}\mathbf{l}, \mathbf{T}^2\mathbf{l}, \mathbf{T}^3\mathbf{l}, \ldots$

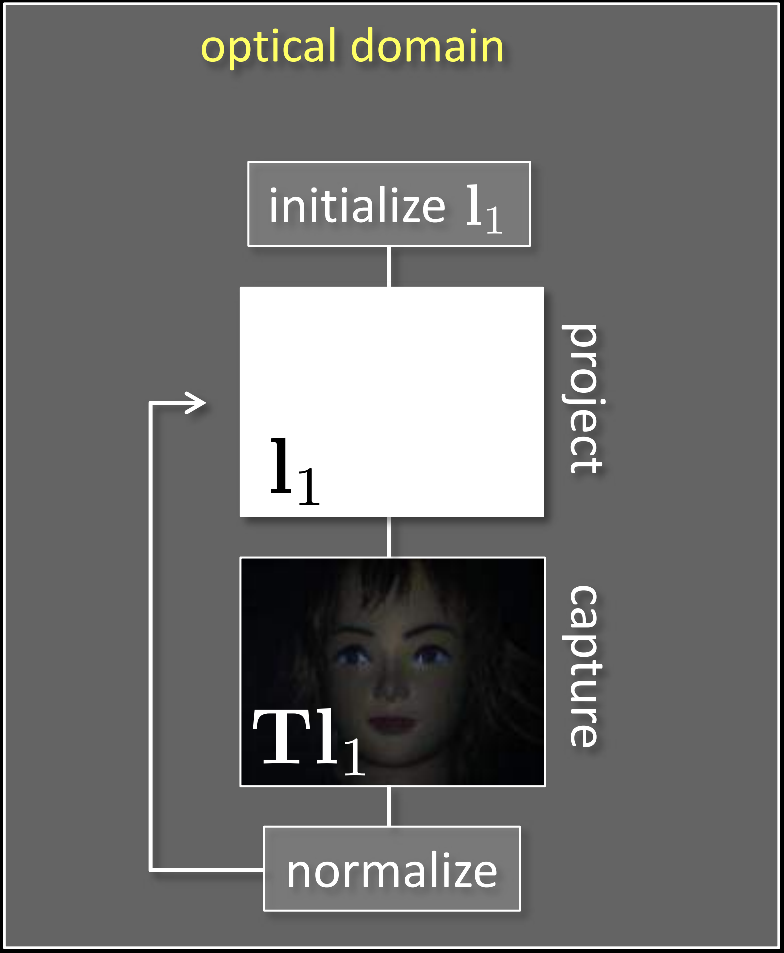

Optical Power Iteration

Goal: find principal eigenvector of $\mathbf{T}$

Observation: it is a fixed point of the sequence $\mathbf{l}, \mathbf{T}\mathbf{l}, \mathbf{T}^2\mathbf{l}, \mathbf{T}^3\mathbf{l}, \ldots$

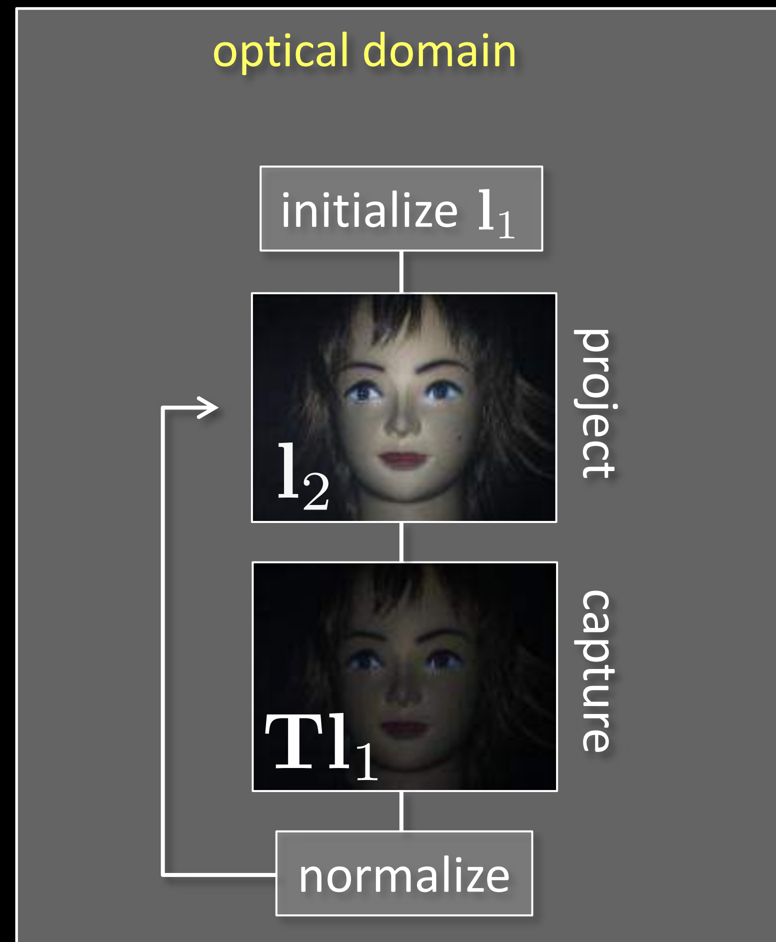

Optical Power Iteration

Goal: find principal eigenvector of $\mathbf{T}$

Observation: it is a fixed point of the sequence $\mathbf{l}, \mathbf{T}\mathbf{l}, \mathbf{T}^2\mathbf{l}, \mathbf{T}^3\mathbf{l}, \ldots$

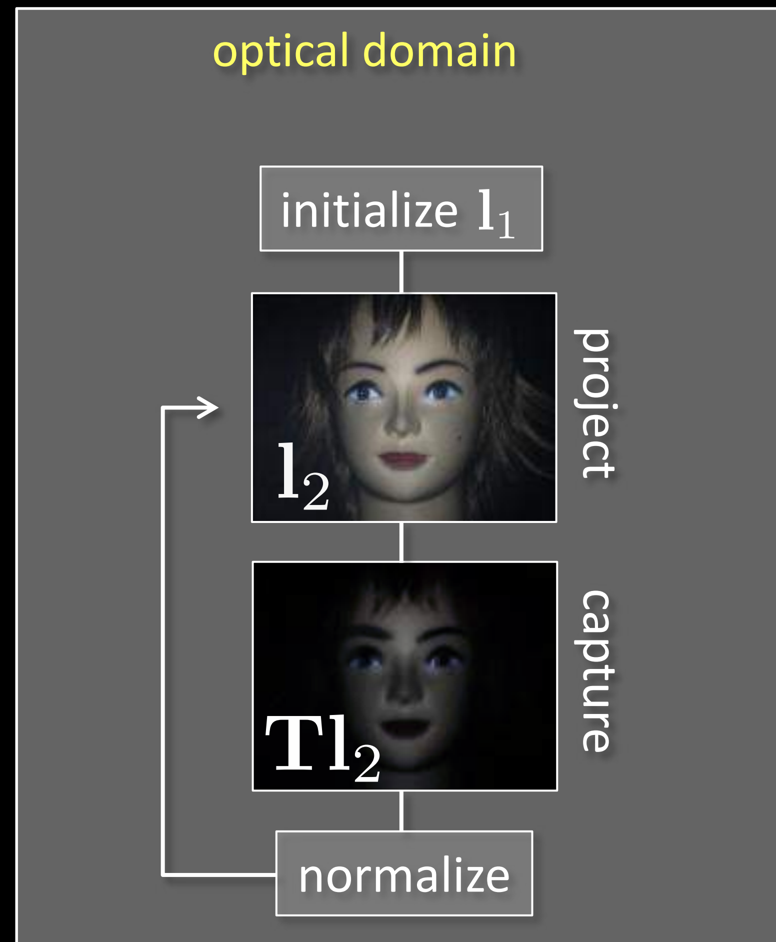

Optical Power Iteration

Goal: find principal eigenvector of $\mathbf{T}$

Observation: it is a fixed point of the sequence $\mathbf{l}, \mathbf{T}\mathbf{l}, \mathbf{T}^2\mathbf{l}, \mathbf{T}^3\mathbf{l}, \ldots$

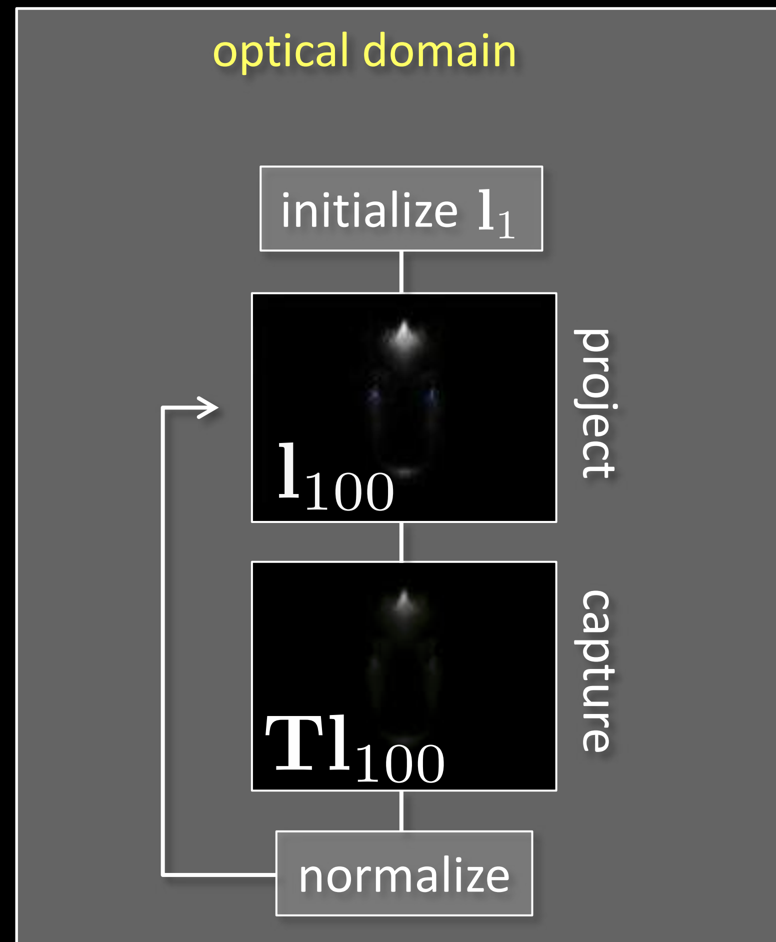

Optical Power Iteration

Goal: find principal eigenvector of $\mathbf{T}$

Observation: it is a fixed point of the sequence $\mathbf{l}, \mathbf{T}\mathbf{l}, \mathbf{T}^2\mathbf{l}, \mathbf{T}^3\mathbf{l}, \ldots$

optical domain

(approximate)
principal eigenvector

Rank-k Transport Approximation

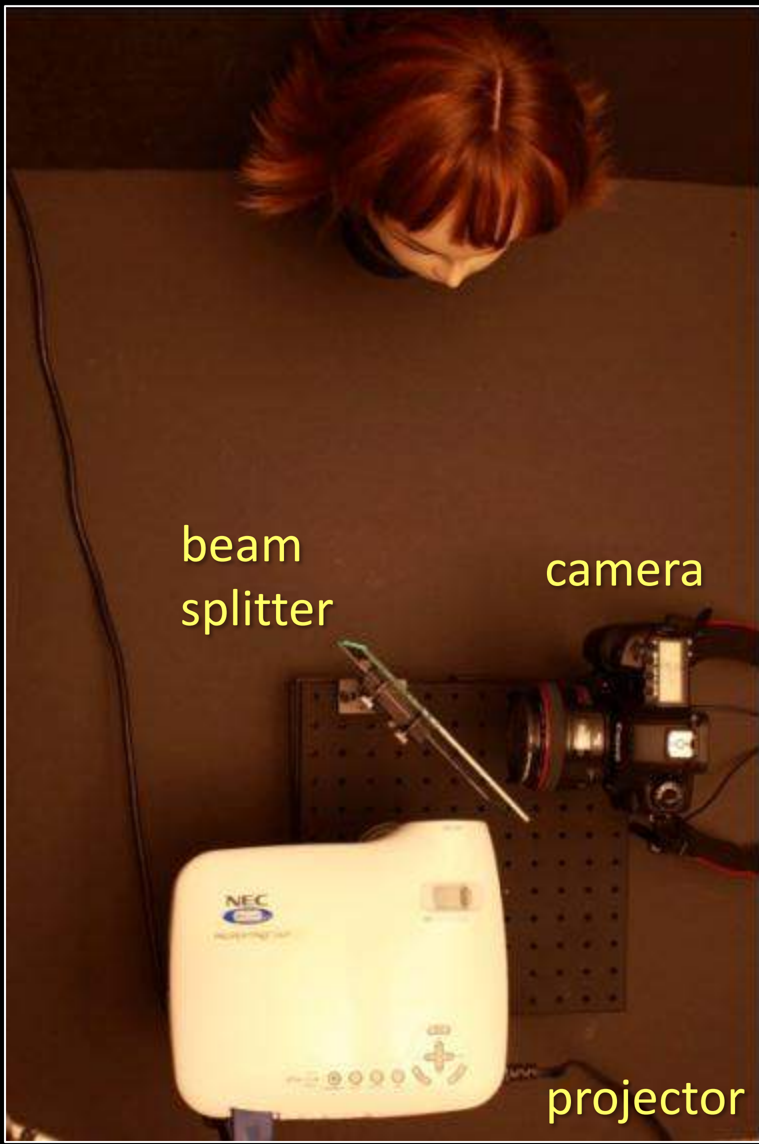

Numerical goal [Simon and Zha 2000]

find matrices $\mathbf{P}_{n \times k}, \mathbf{L}_{k \times m}$

that minimize

$$\left\| \mathbf{T} - \mathbf{P}\mathbf{L} \right\|_F$$

k

Symmetric $\mathbf{T}$

- 1 camera, 1 projector
- 2 k photos for rank- k approx.

Nonsymmetric $\mathbf{T}$

- 2 cameras, 2 projectors
- 4 k photos for rank- k approx.

Results: Optical Arnoldi

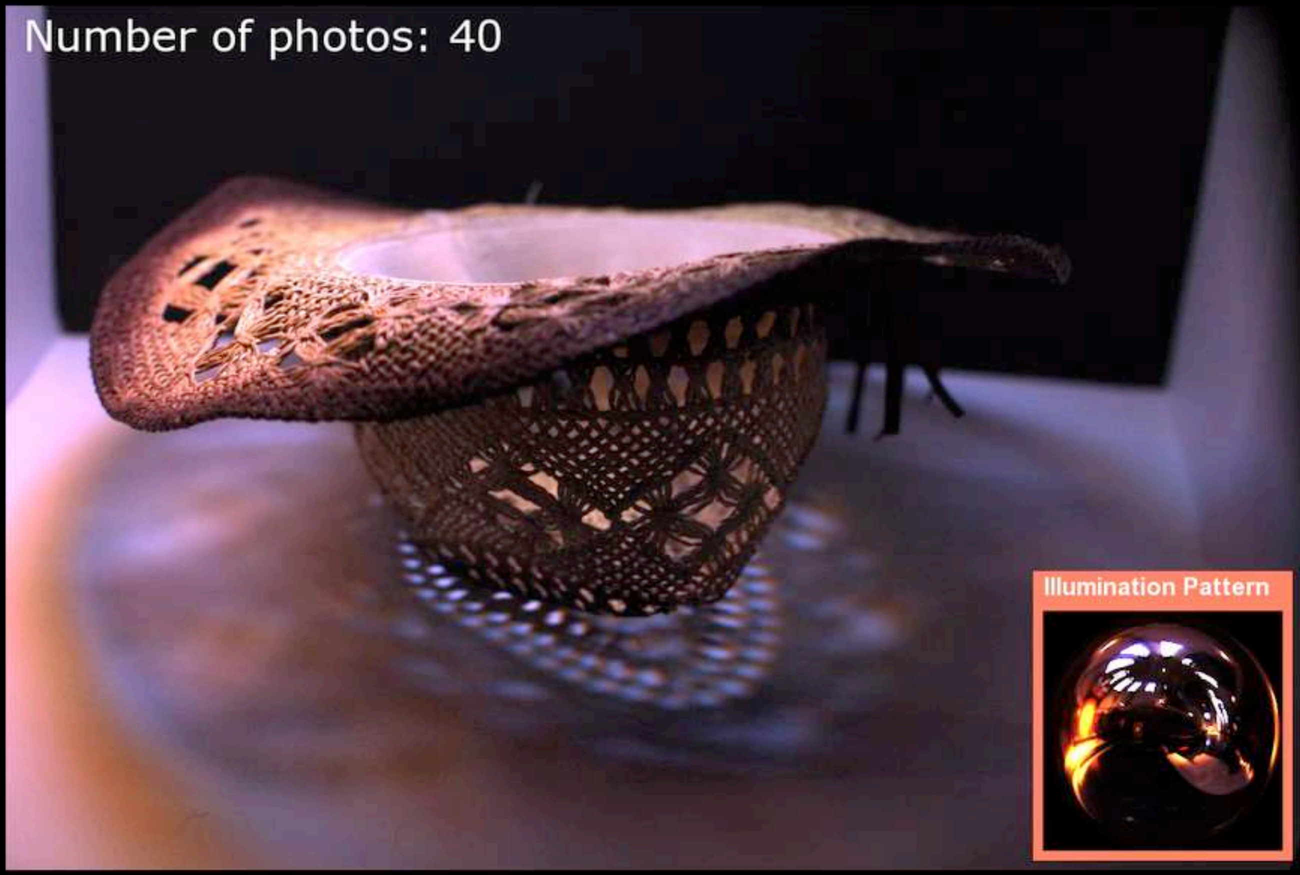

concluding remarks

- the light transport matrix is a general model for describing the transfer or radiant energy
- the entries of a transport matrix describes all possible observations one can make of a scene
- transport matrix is often too large to measure directly in practice
- numerical algorithms can be partially or fully implemented in the optical domain