Light transport matrices

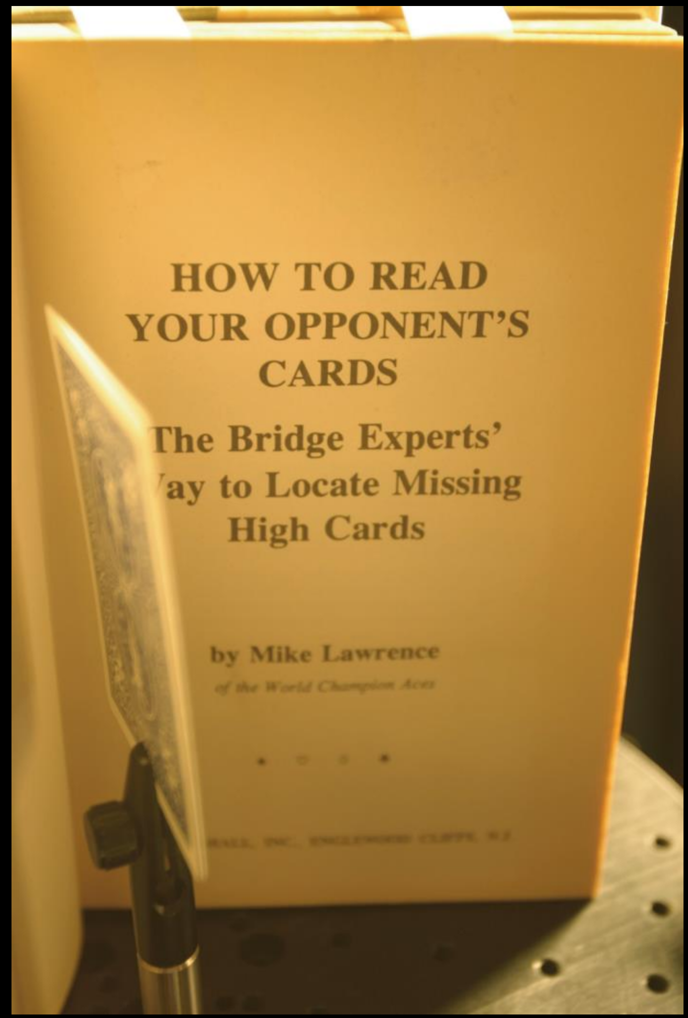

15-463, 15-663, 15-862
Computational Photography
Fall 2018, Lecture 18

http://graphics.cs.cmu.edu/courses/15-463

Course announcements

- Homework 5 has been posted.
  - Due on Friday November 9^{th}.
- Any problems with homework 4?
- No elevator pitch presentations for final projects.
- Extra office hours this week:
  - Monday 1:30-3:30 pm.
  - Tuesday noon-2:00 pm.
  - Friday's office hours will be held by Alankar.
- Great talk this Thursday: Eric Fossum, inventor of the CMOS sensor, will talk about quantum (i.e., photon-counting) CMOS sensors.

Overview of today’s lecture

- Leftover from last time: Generalized bas-relief ambiguity.
- The light transport matrix.
- Image-based relighting.
- Photometric stereo revisited.
- Optical computing using the light transport matrix.
- Dual photography.

Slide credits

These slides were directly adapted from:

- Matt O’Toole (CMU).

The light transport matrix

How do these three images relate to each other?


photo with lights 1 & 2 turned on


photo with light 1 turned on


photo with light 2 turned on

How do these three images relate to each other?

the superposition principle

photo with lights 1 & 2 turned on

=

photo with light 1 turned on

+

photo with light 2 turned on

photo taken under two light sources =
sum of photos taken under each source individually

the superposition principle

photo taken under two light sources =
sum of photos taken under each source individually

the superposition principle

why is the error not exactly zero?

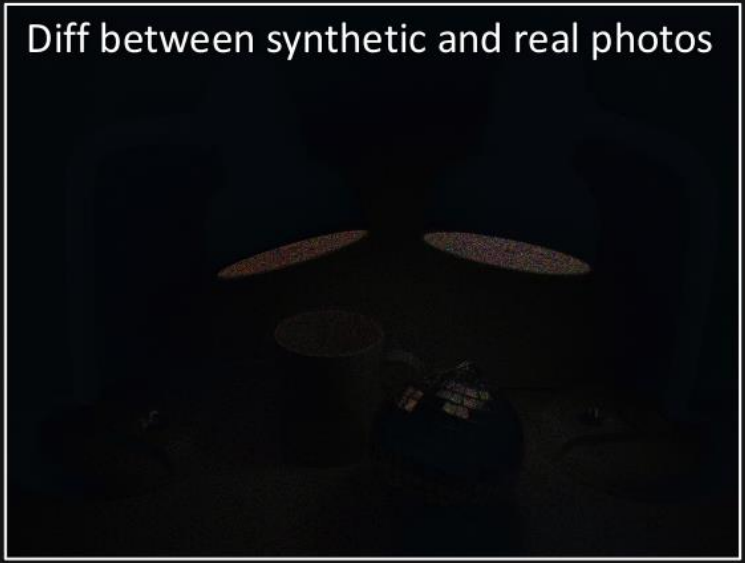

photo taken under two light sources =
sum of photos taken under each source individually

image-based relighting

=

image-based relighting

=

Weight 1

x 1

+

Weight 2

x 1

image-based relighting

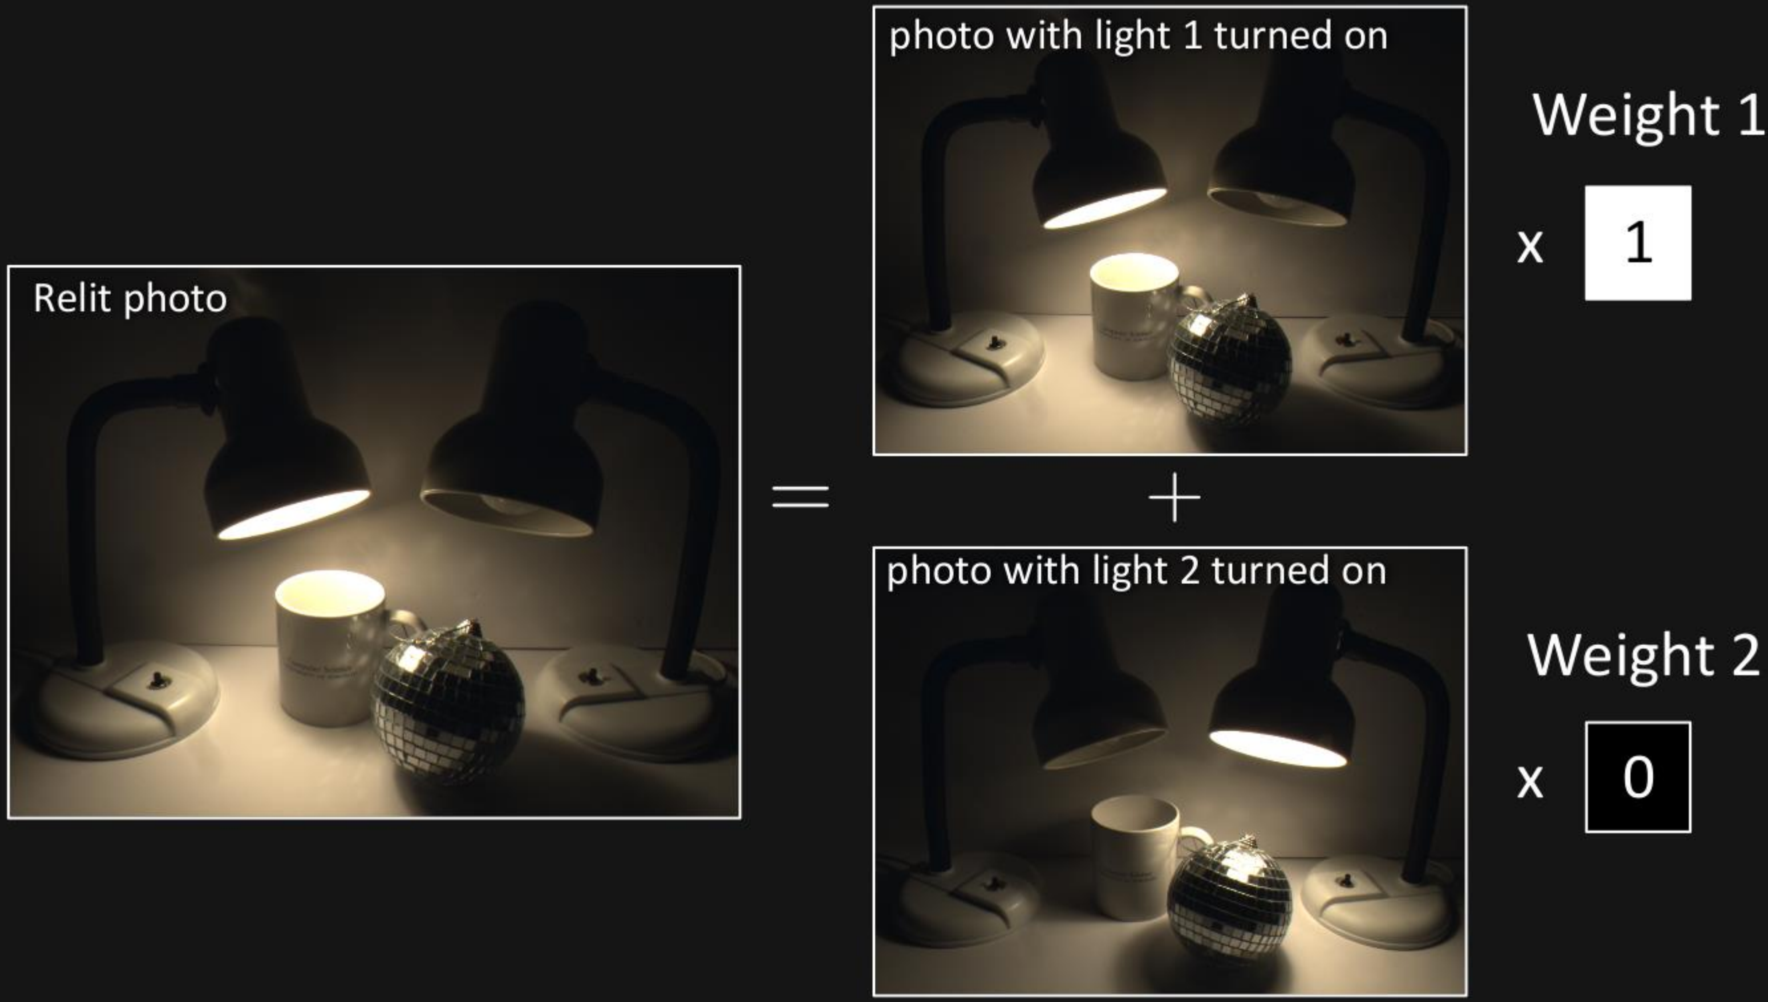

image-based relighting

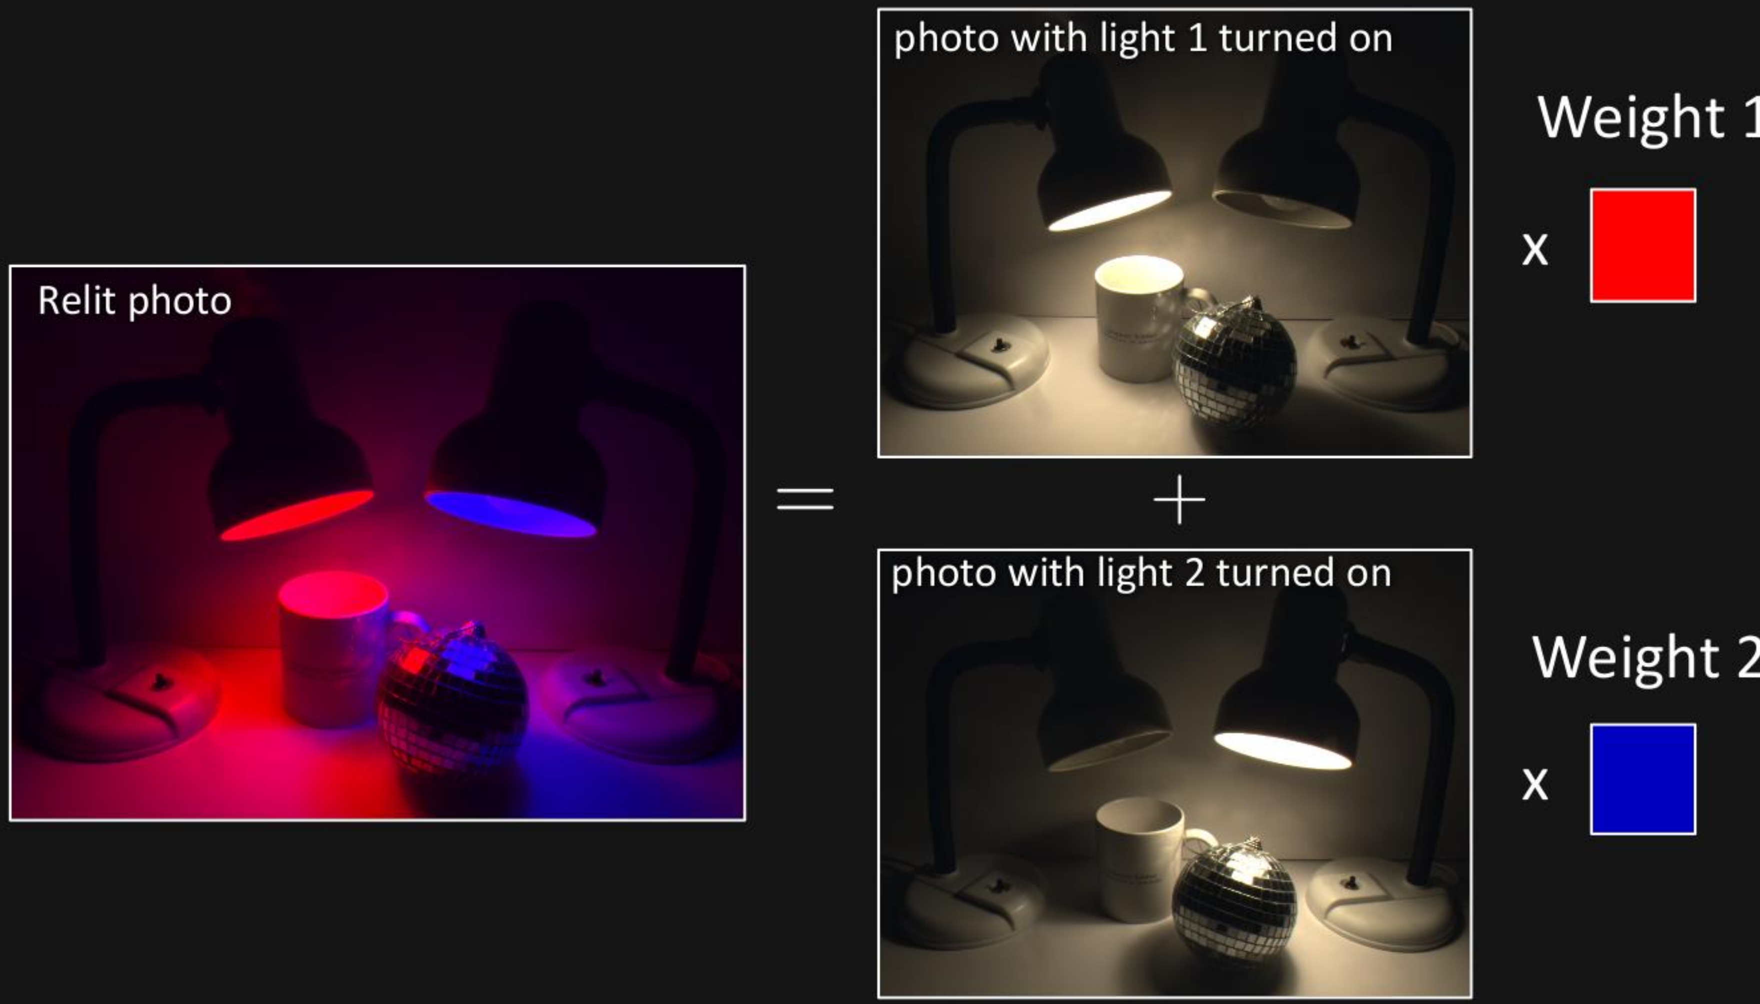


Relit photo $\mathbf{p}$
=
photo with light 1 turned on $\mathbf{T_1}$
Weight 1
x $\mathbf{l_1}$
+
photo with light 2 turned on $\mathbf{T_2}$
Weight 2
x $\mathbf{l_2}$

=

Weight 1 x $\mathbf{l_1}$ +

Weight 2 x $\mathbf{l_2}$

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

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 $\mathbf{l_2}$

n

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

n pixel values

Relit photo $\mathbf{p}$

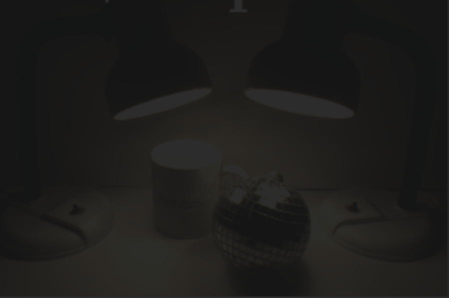

=

photo with light 1 turned on $\mathbf{T_1}$

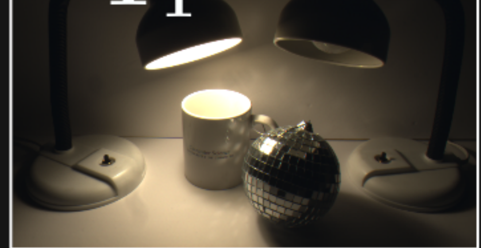

Weight 1 x $\mathbf{l_1}$ +

photo with light 2 turned on $\mathbf{T_2}$

Weight 2 x $\mathbf{l_2}$

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

n

n pixel values

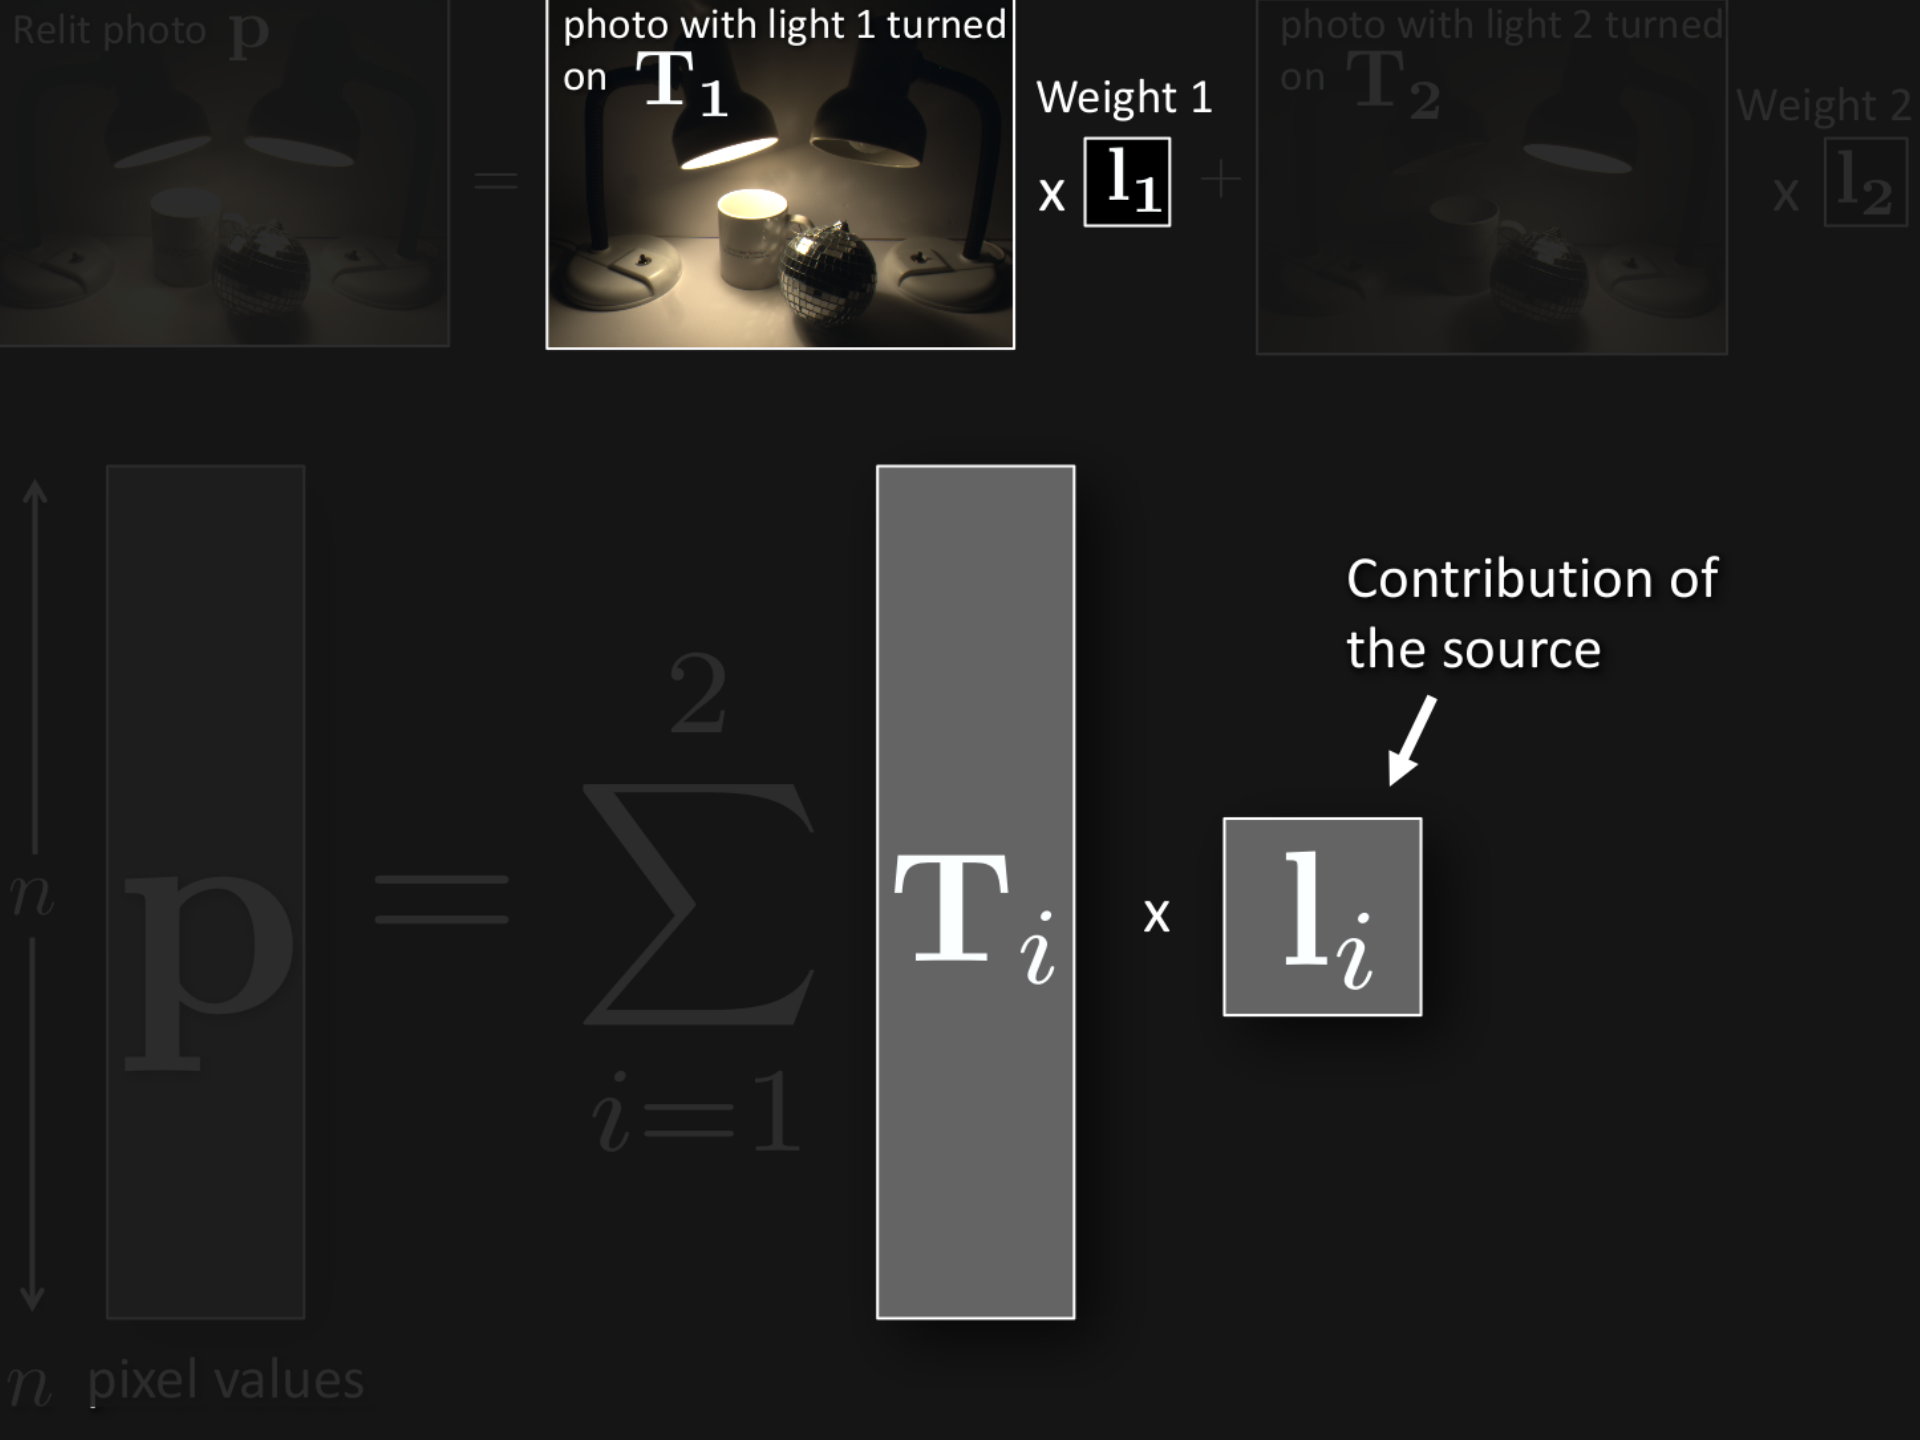

Relit photo p
=
photo with light 1 turned on $\mathbf{T_1}$
Weight 1
x $\mathbf{l_1}$
+
photo with light 2 turned on $\mathbf{T_2}$
Weight 2
x $\mathbf{l_2}$
n
$\mathbf{p} = \sum_{i=1}^{2} \mathbf{T}_i \times \mathbf{l}_i$
Contribution of the source
n pixel values

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 $\mathbf{l_2}$

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

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

n

n pixel values

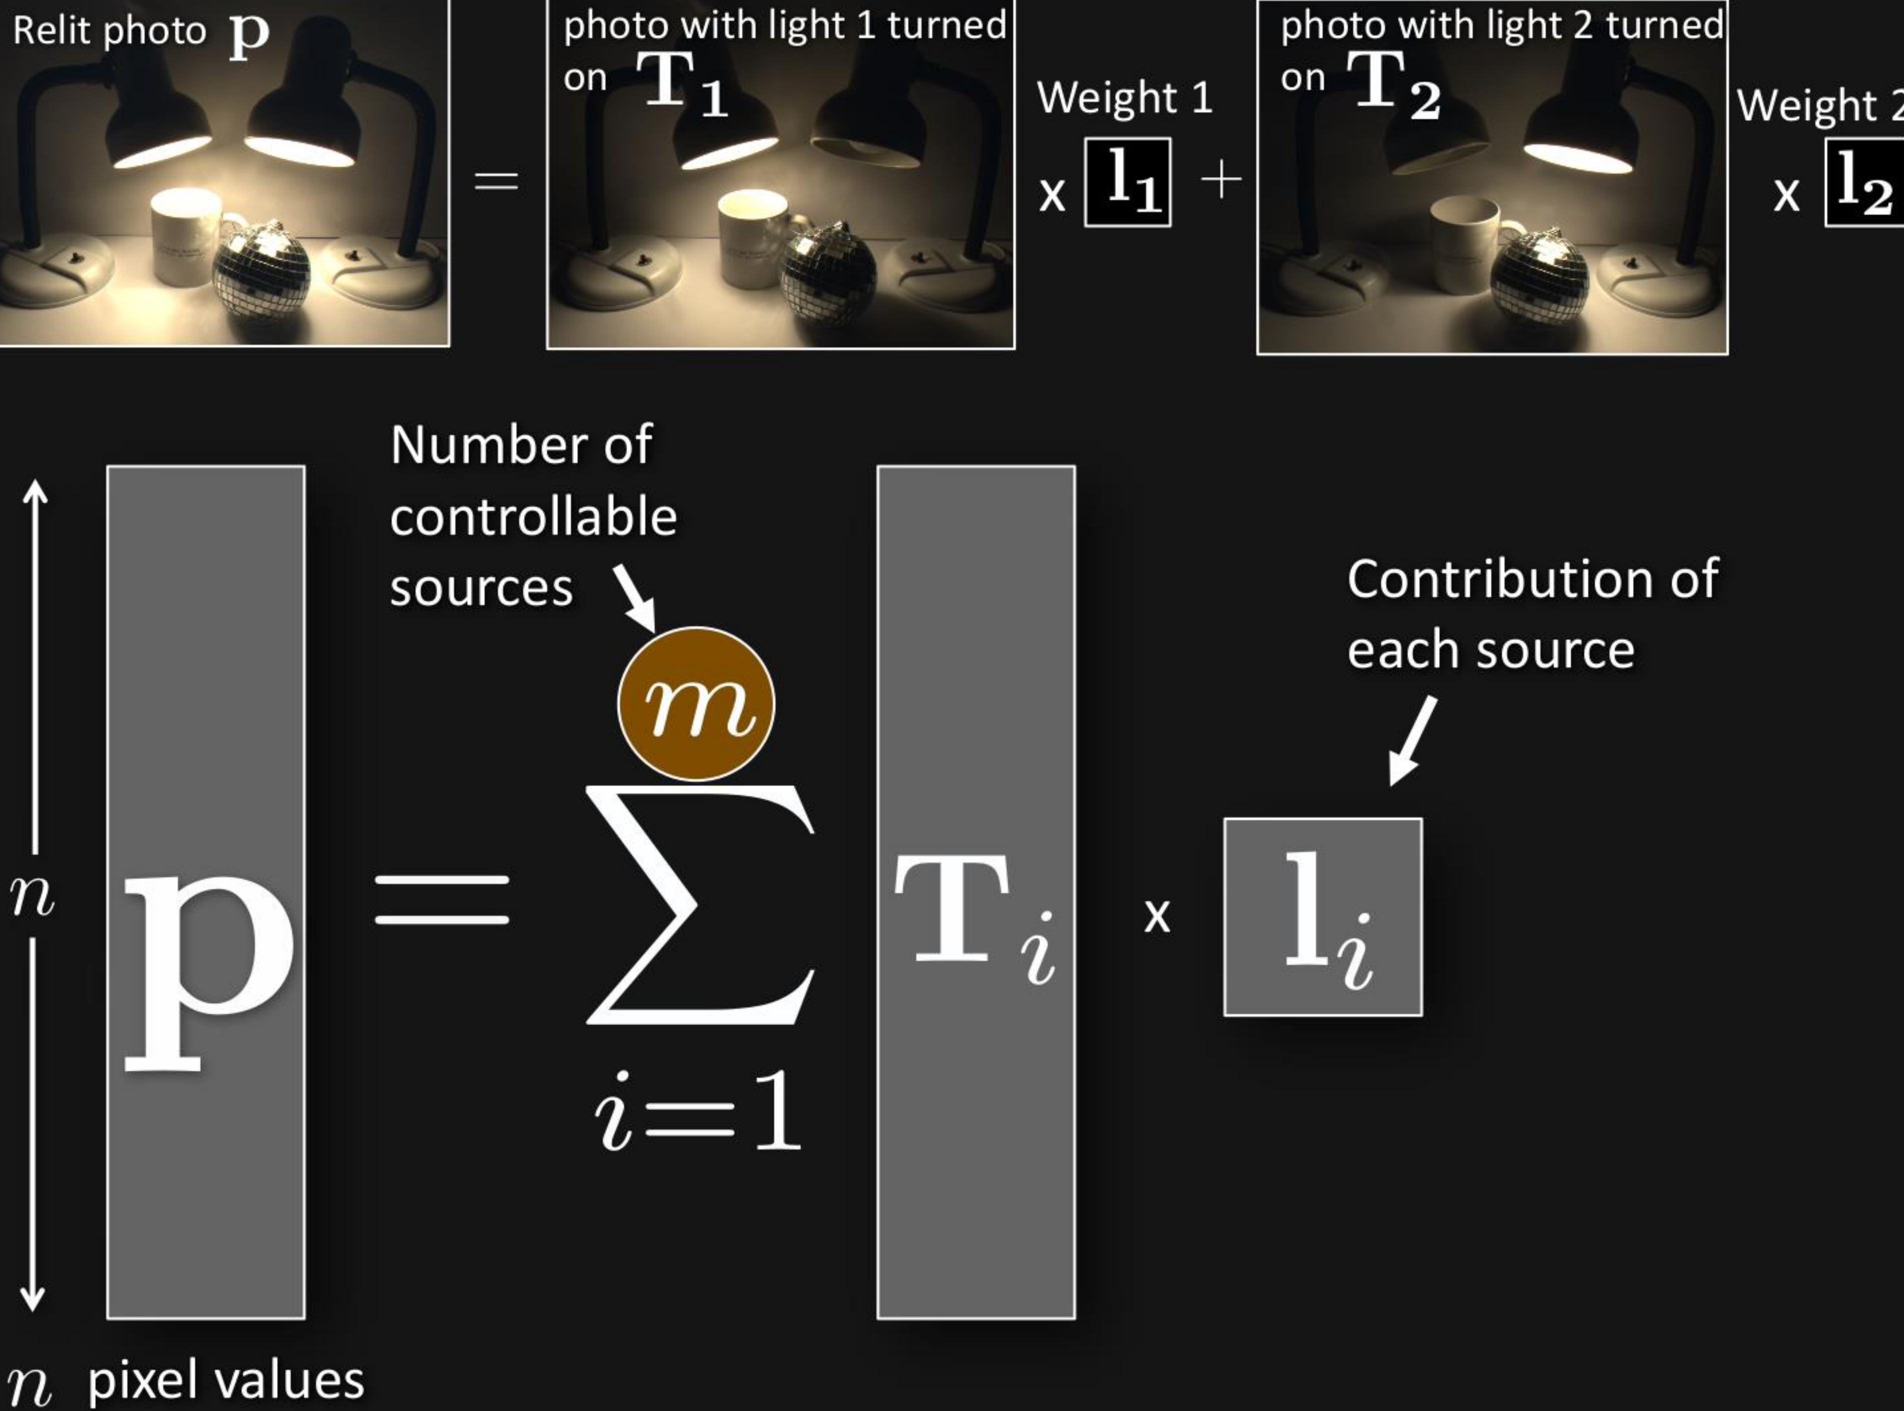
Relit photo $\mathbf{p}$
=
photo with light 1 turned on $\mathbf{T_1}$
Weight 1
x $\mathbf{l_1}$
+
photo with light 2 turned on $\mathbf{T_2}$
Weight 2
x $\mathbf{l_2}$
Number of controllable sources
Contribution of each source
$\mathbf{p} = \sum_{i=1}^{m} \mathbf{T}_i \times \mathbf{l}_i$
n
n pixel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

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 $\mathbf{l_2}$

light transport matrix

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

n pixel values

$n \times m$

m

Can you think of a case where we have a very large m?

Use a projector

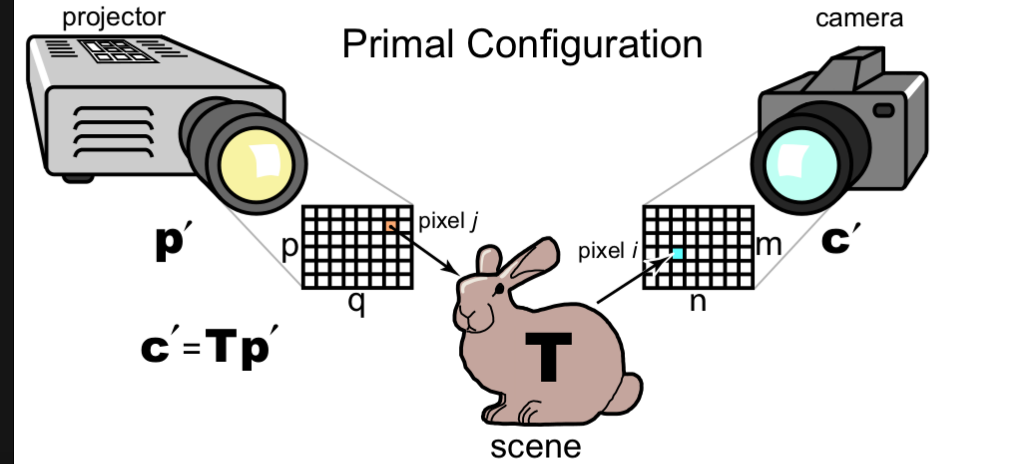

$$\underbrace{\mathbf{p}}_{n} = \underbrace{\mathbf{T}}_{n \times m} \; \underbrace{\mathbf{l}}_{m}$$

n pixel values

$n \times m$

What does each row and column of T correspond to here?

Use a projector

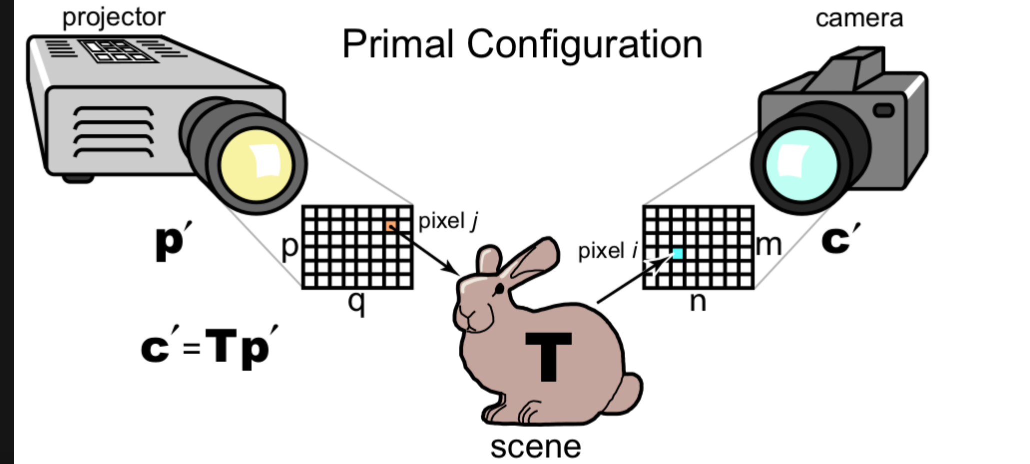

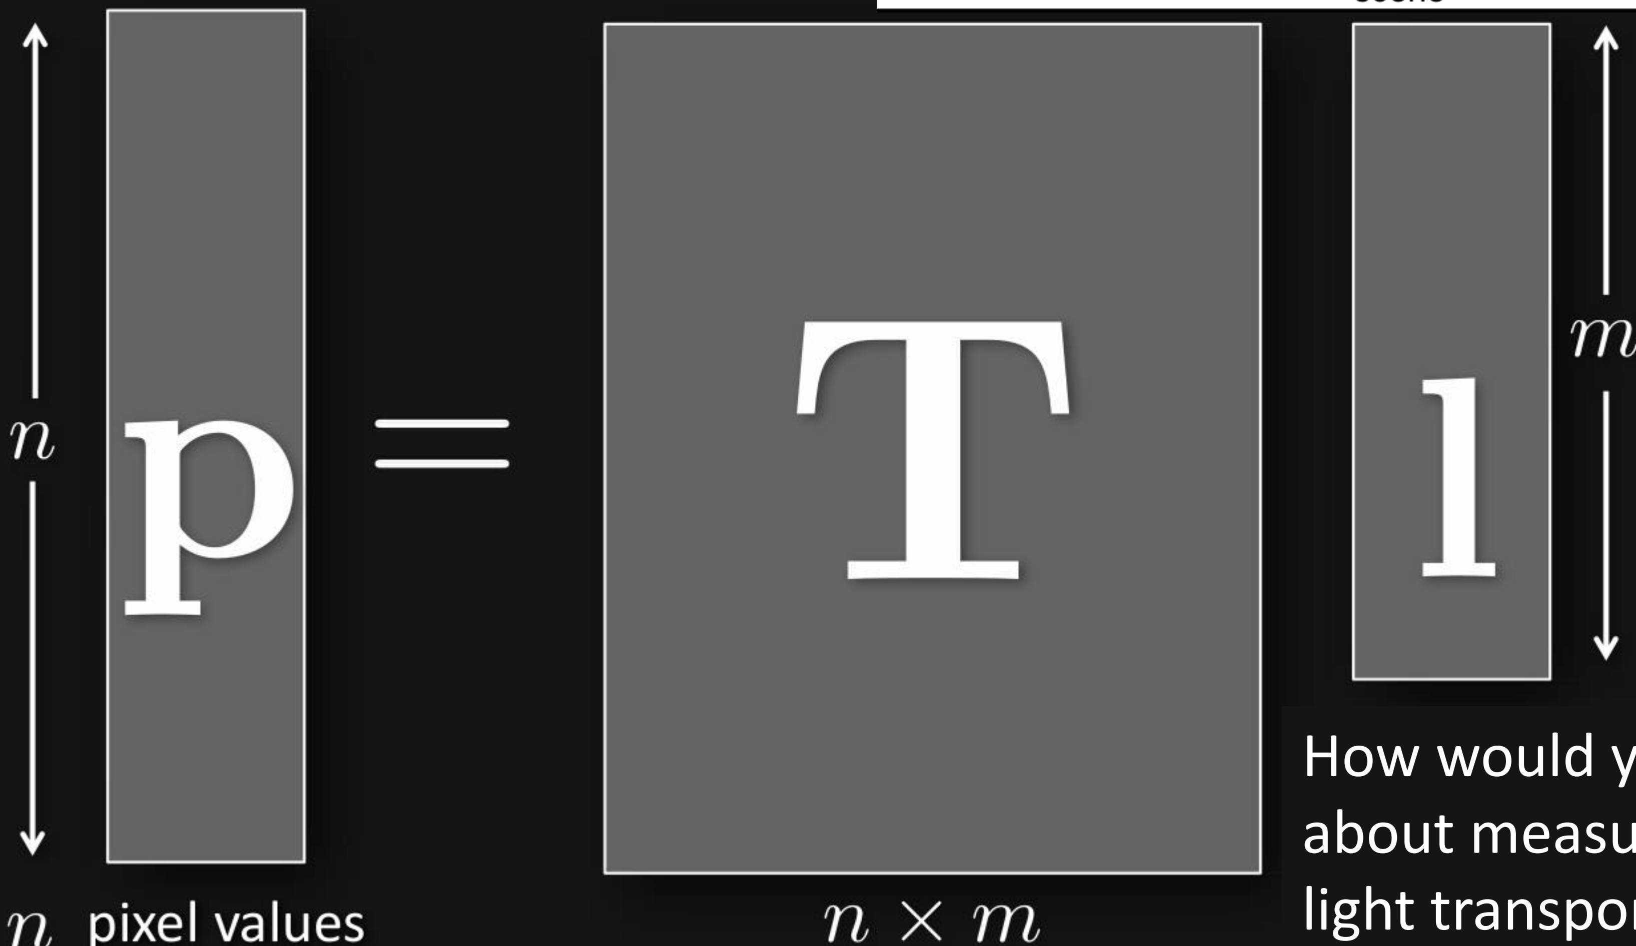

How would you go about measuring the light transport matrix?

Image-based relighting

Let’s say I have measured T.

- What does it mean to right-multiply it with some vector l?

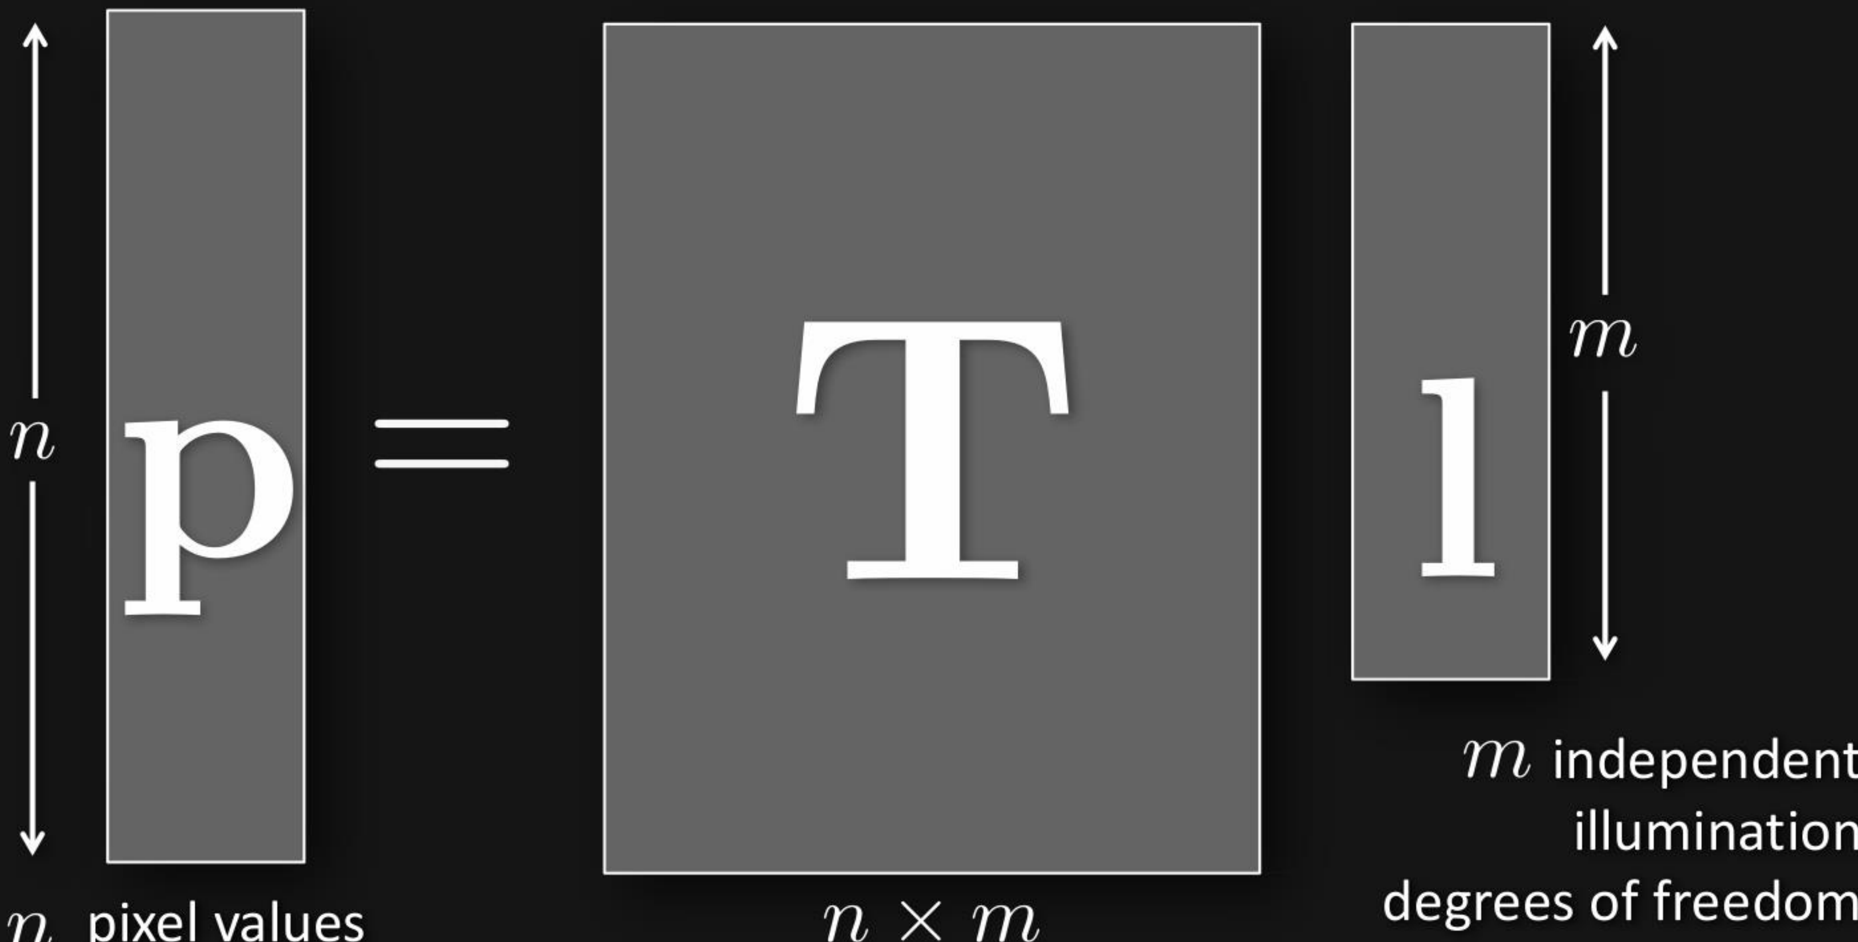

Image-based relighting: Use the measurements I already have of the scene (the pictures I took when measuring T) to simulate new illuminations of the scene.

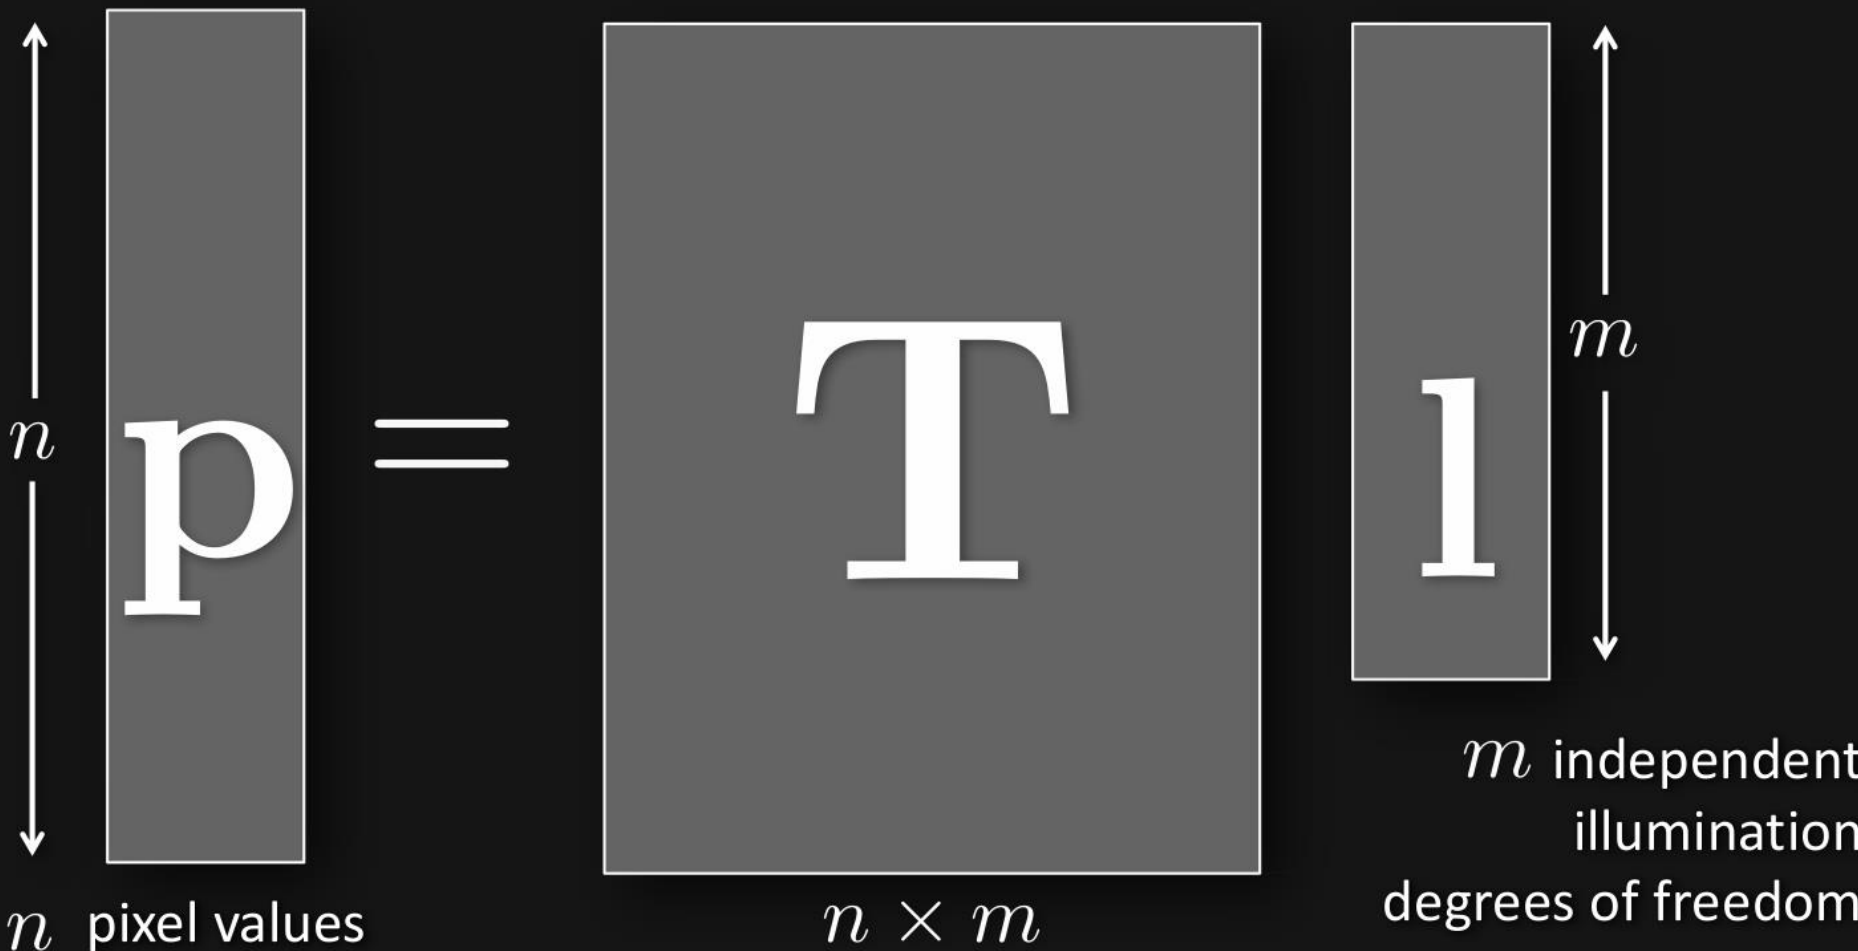

Acquiring the Reflectance Field [Debevec et al. 2000]

image-based rendering & relighting

Light stage 1.0

Reflectance field

Acquiring the Reflectance Field

image-based rendering & relighting

Debevec et al, SIG 2000

Great demonstration: https://www.youtube.com/watch?v=mkzLLz1tXds

Acquiring the Reflectance Field

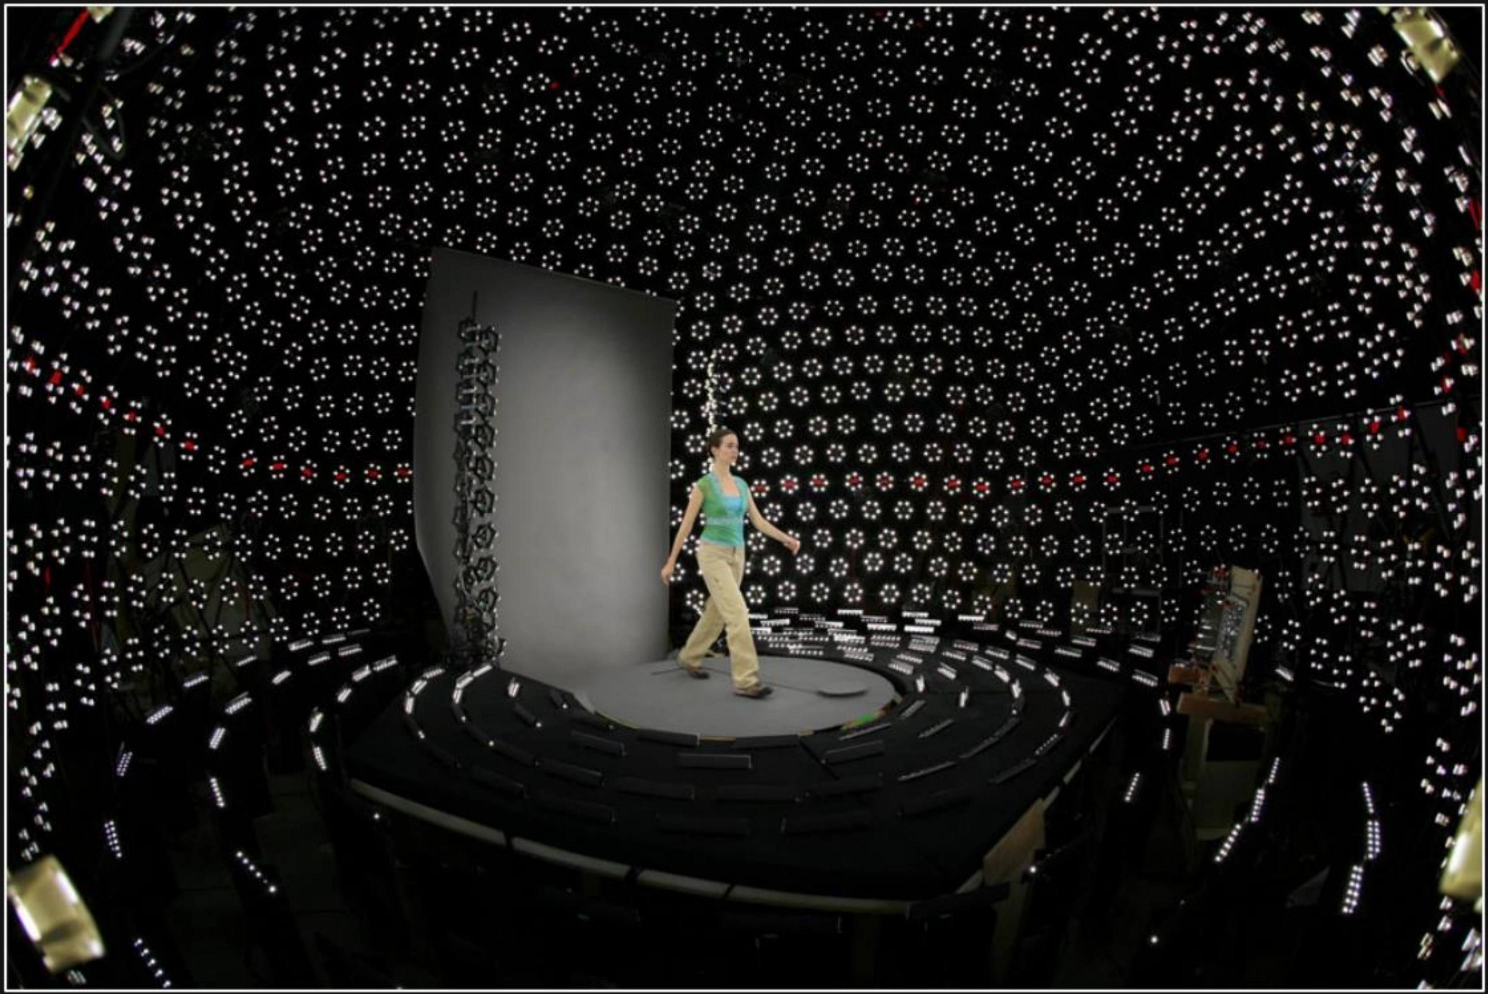

Light stage 6, Debevec et al., 2006

Photometric stereo revisited

the light transport matrix

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

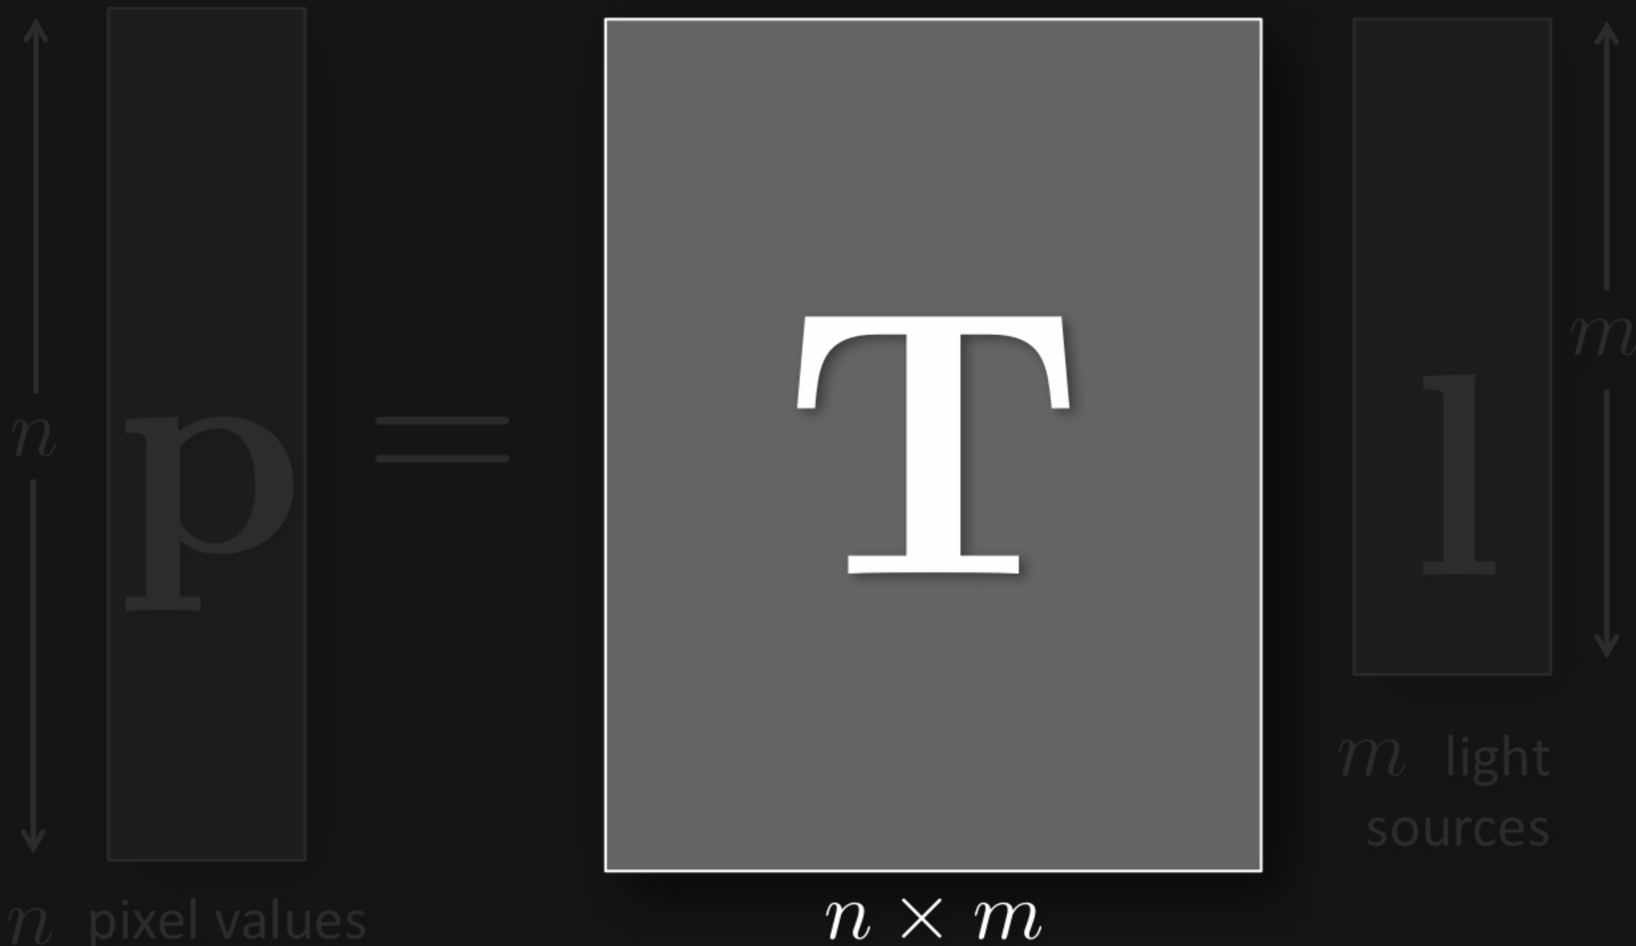

transport matrix is a function of scene geometry, reflectance, etc.

Photometric Stereo [Woodham, 1980]

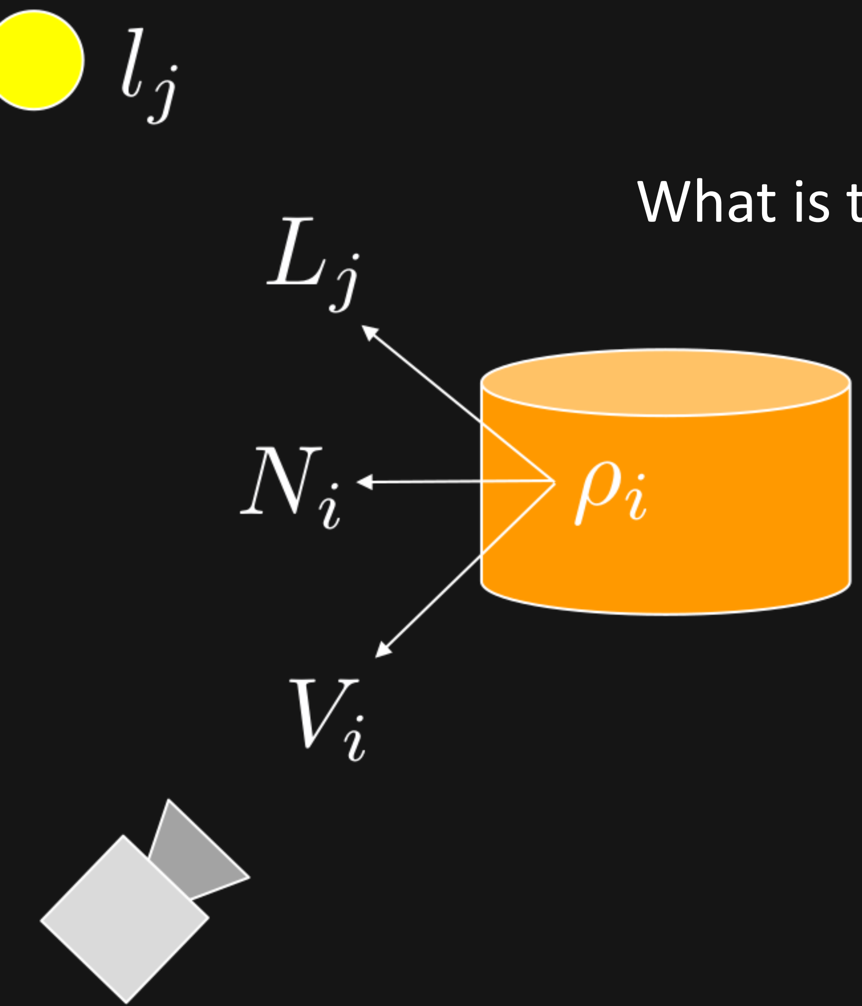

camera pixel i and light source j produce image intensity $\mathbf{t}_{ij}$

Diffuse reflections:

What is this?

$$\mathbf{t}_{ij} = \rho_i (N_i \cdot L_j) l_j$$

$$= \tilde{N}_i \cdot \tilde{L}_j$$

3x1 vector, unknown

3x1 vector, known

simplifying assumptions:
directional light source,
convex object

Photometric Stereo [Woodham, 1980]

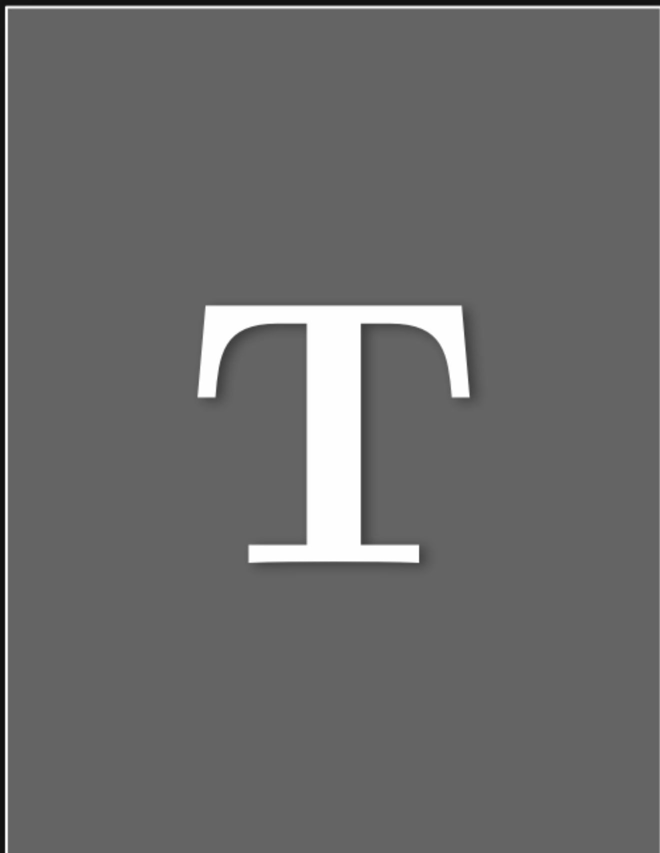

$n \times m$

Diffuse reflections:

$$\mathbf{t}_{ij} = \rho_i (N_i \cdot L_j) l_j$$

$$= \tilde{N}_i \cdot \tilde{L}_j$$

$\tilde{N}_i$: 3x1 vector, unknown

$\tilde{L}_j$: 3x1 vector, known

simplifying assumptions:
directional light source,
convex object

Photometric Stereo [Woodham, 1980]

$$\mathbf{T} = \tilde{N} \tilde{L}$$

$\mathbf{T}$: $n \times m$ (rank 3)

$\tilde{N}$: $n \times 3$

$\tilde{L}$: $3 \times m$

recover surface normals + albedo by decomposing transport matrix $\mathbf{T}$

Recovering Scene Geometry

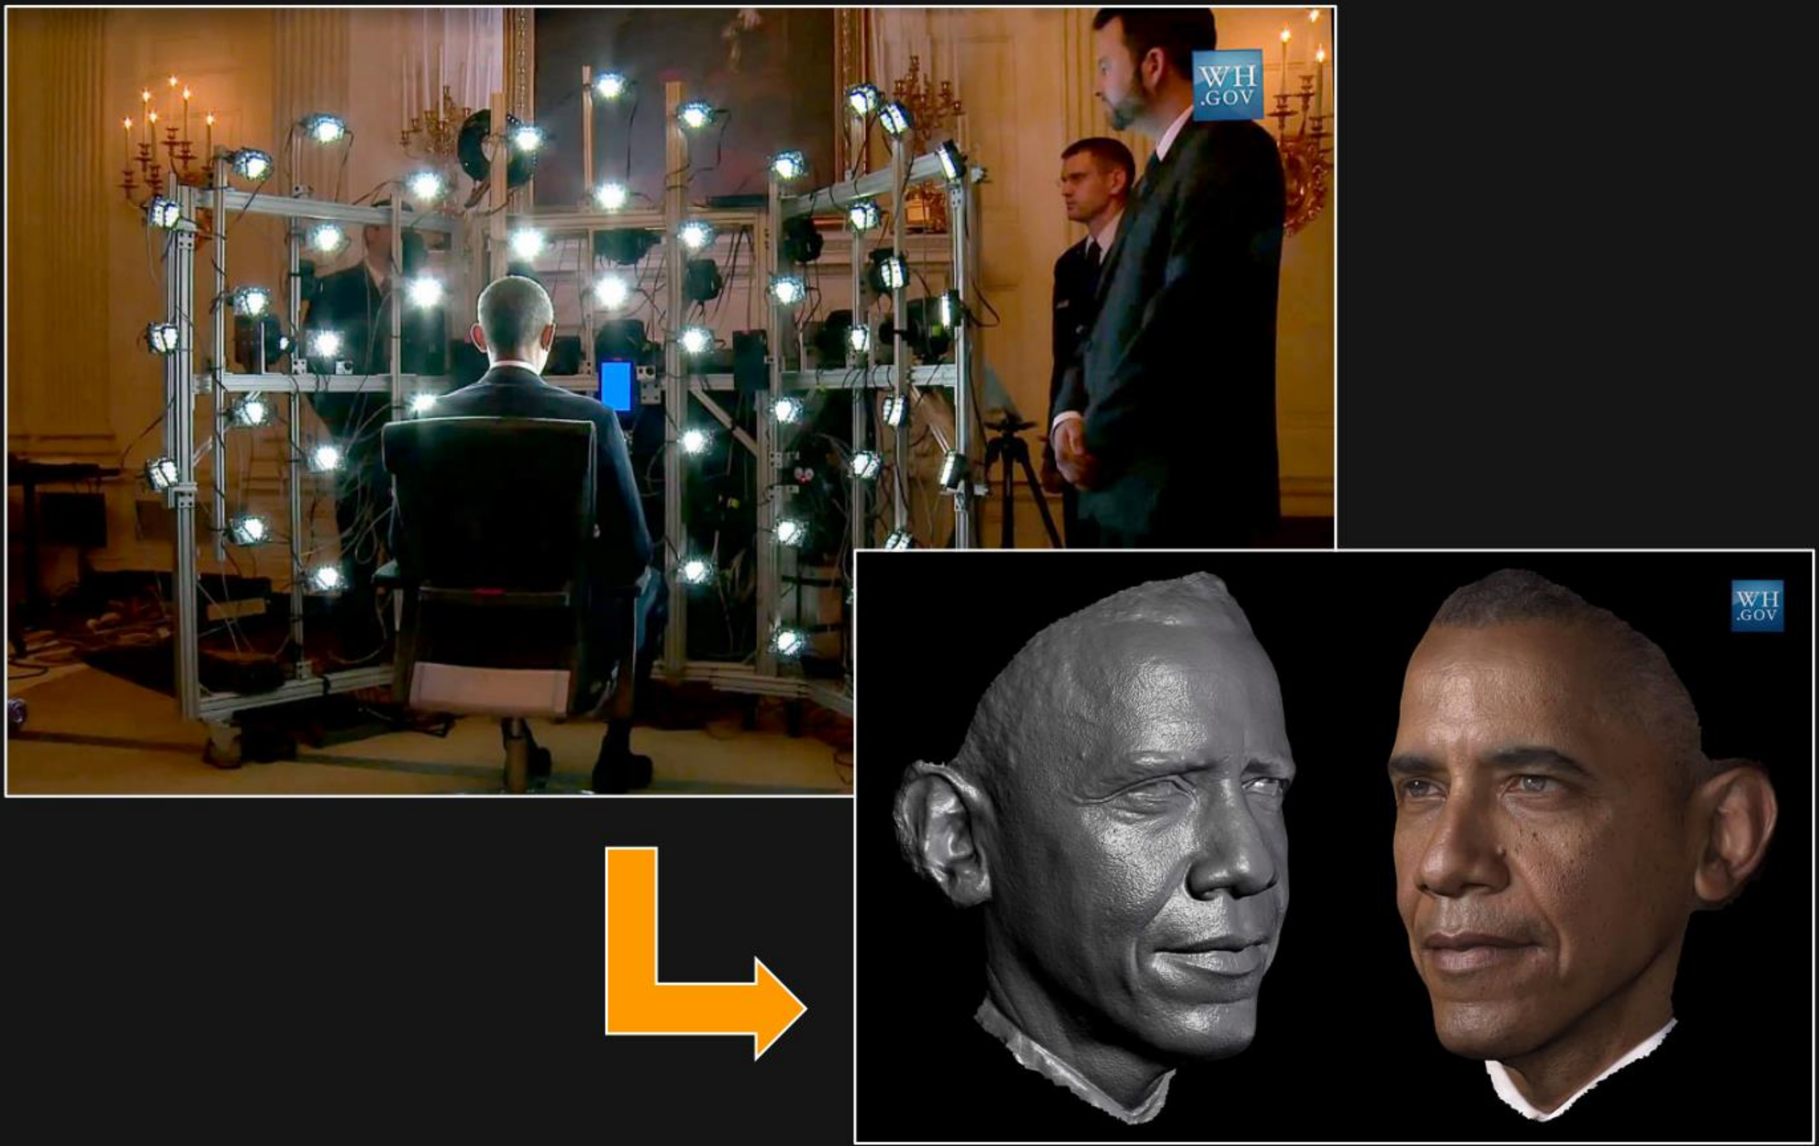

“Mobile” Light Stage, Debevec et al., 2014

https://www.youtube.com/watch? v=4GiLAOtjHNo

Optical computing using the light transport matrix

main difficulties

question: what are the challenges with analyzing $\mathbf{T}$?

$$\mathbf{p} = \mathbf{T}\mathbf{1}$$

main difficulties

question: what are the challenges with analyzing $\mathbf{T}$?

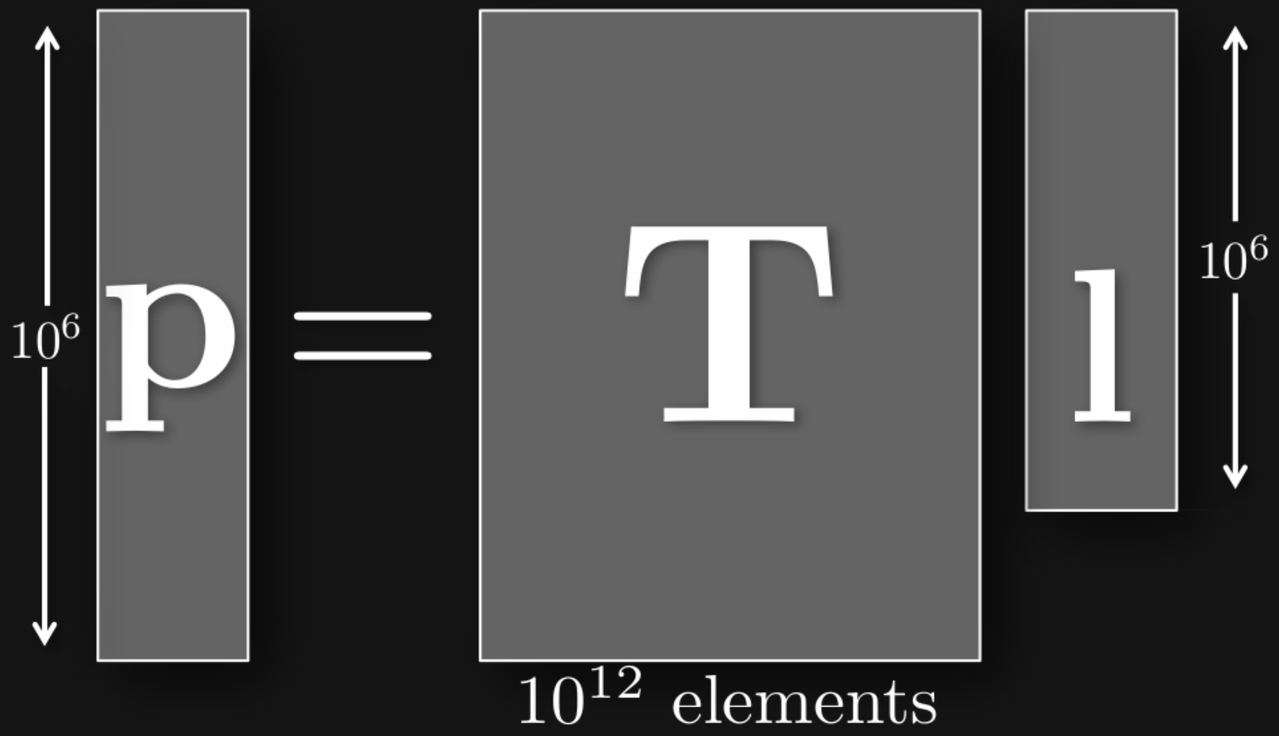

- matrix can be extremely large

main difficulties

question: what are the challenges with analyzing $\mathbf{T}$?

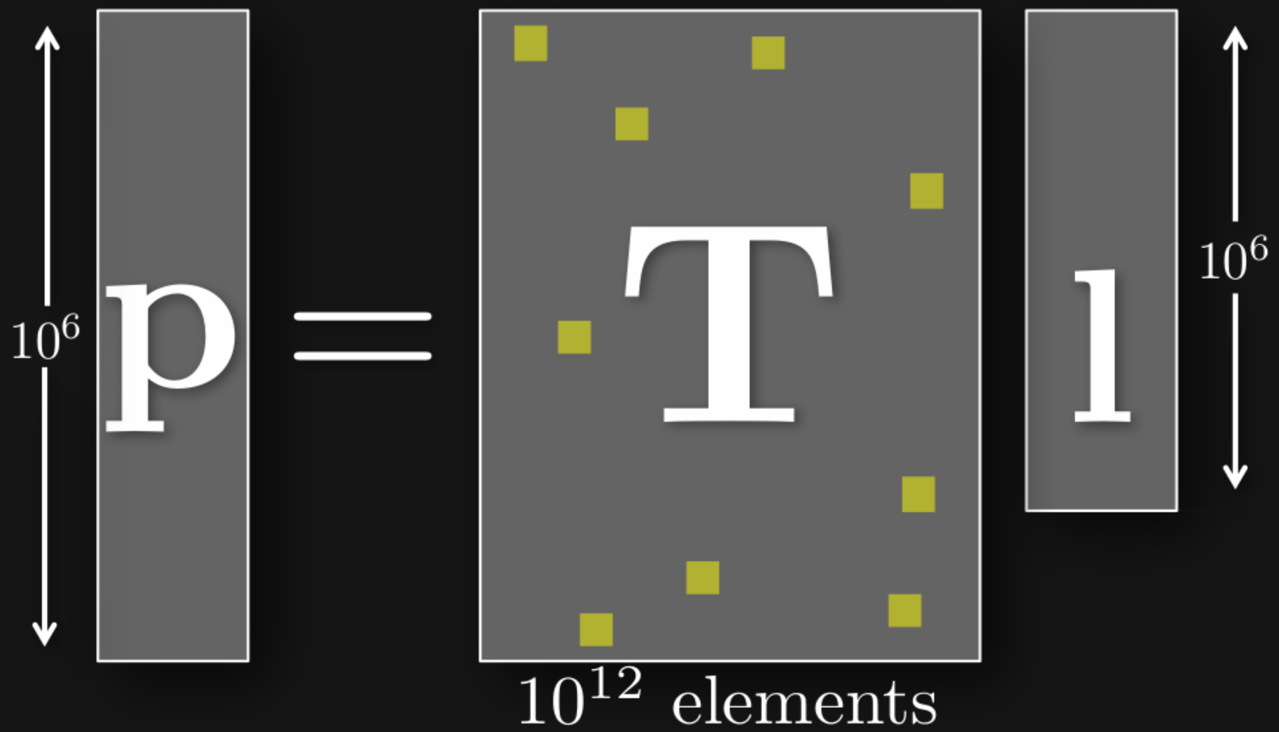

- matrix can be extremely large
- elements not directly accessible

main difficulties

question: what are the challenges with analyzing $\mathbf{T}$?

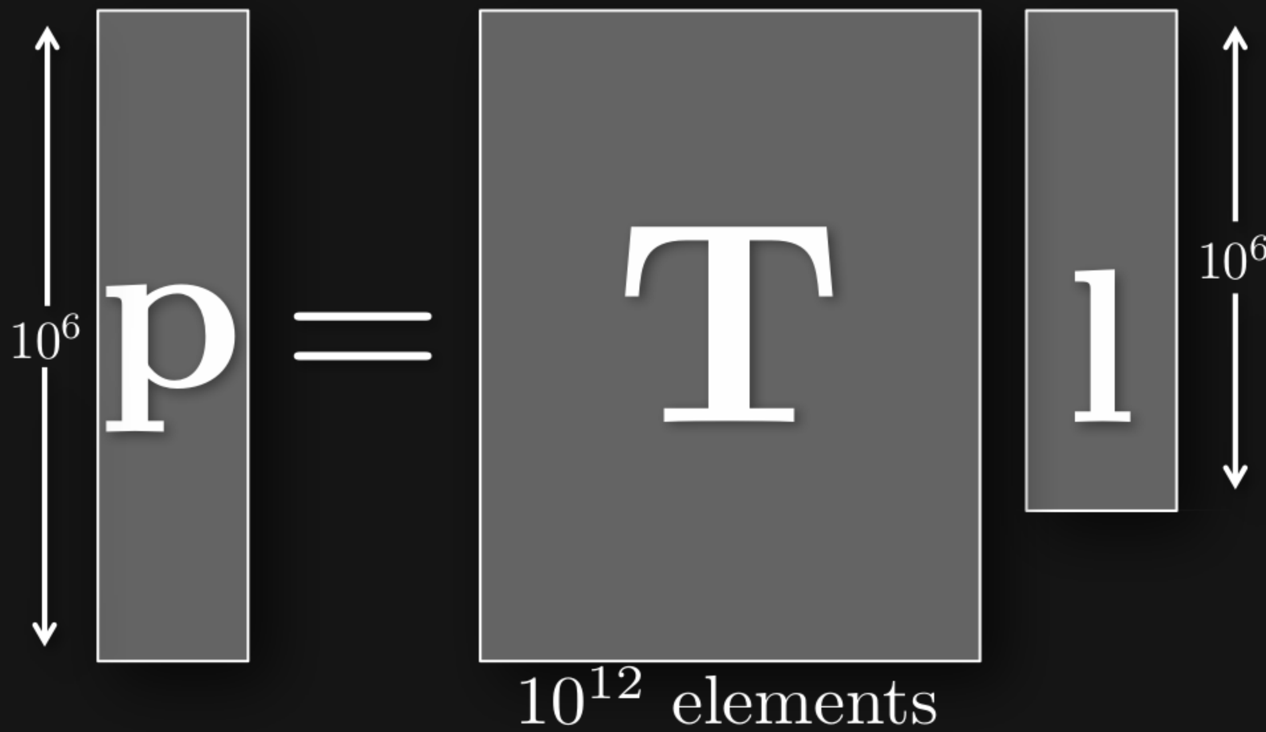

- matrix can be extremely large
- elements not directly accessible
- global structure poorly understood

computing with light

numerical algorithms implemented directly in optics

numerical domain

transport matrix

↓

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

↑ photo

↑ illumination vector

$\Longleftrightarrow$

optical domain

computing with light

numerical algorithms implemented directly in optics

numerical domain

transport matrix

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

photo

illumination vector

$\Longleftrightarrow$

optical domain

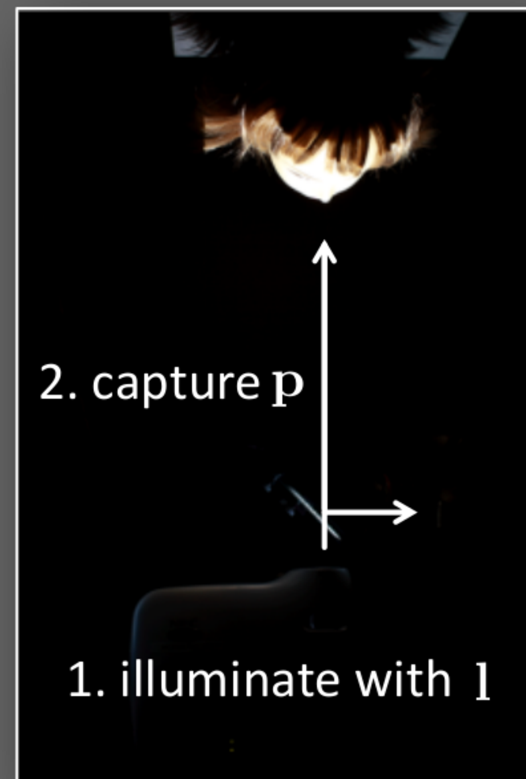

computing with light

numerical algorithms implemented 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

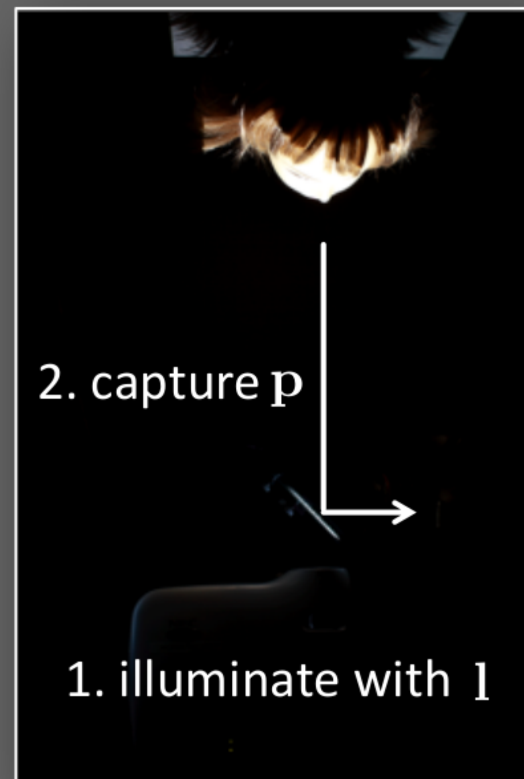

computing with light

numerical algorithms implemented directly in optics

numerical domain

```
function analyze(T)
...
for i = 1 to k {
    ...
    p_i = T l_i
    ...
    d_i = T r_i
    ...
}
...
return result
```

$\Longleftrightarrow$

optical domain

```
function analyze()
...
for i = 1 to k {
    ...
    project l_i, capture p_i
    ...
    project r_i, capture d_i
    ...
}
...
return result
```

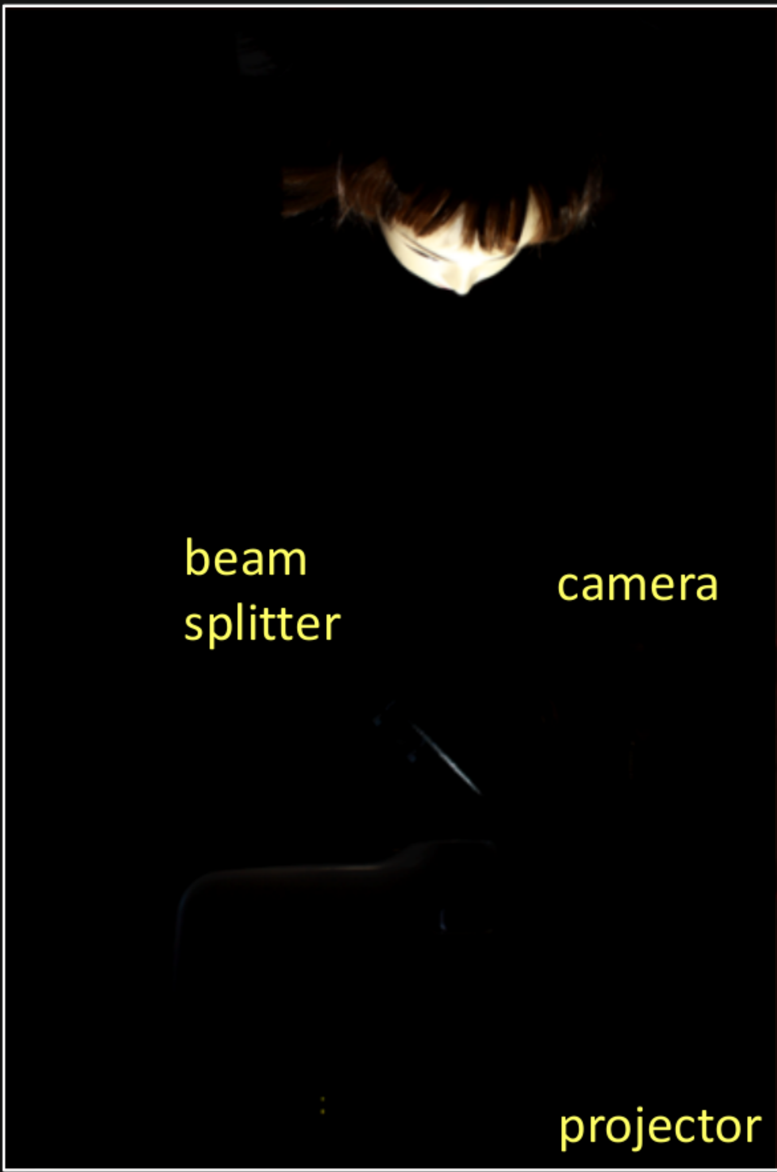

find an illumination pattern that
when projected onto scene,
we get the same photo back
(multiplied by a scalar)

project

capture

What do we call these patterns?

computing transport eigenvectors

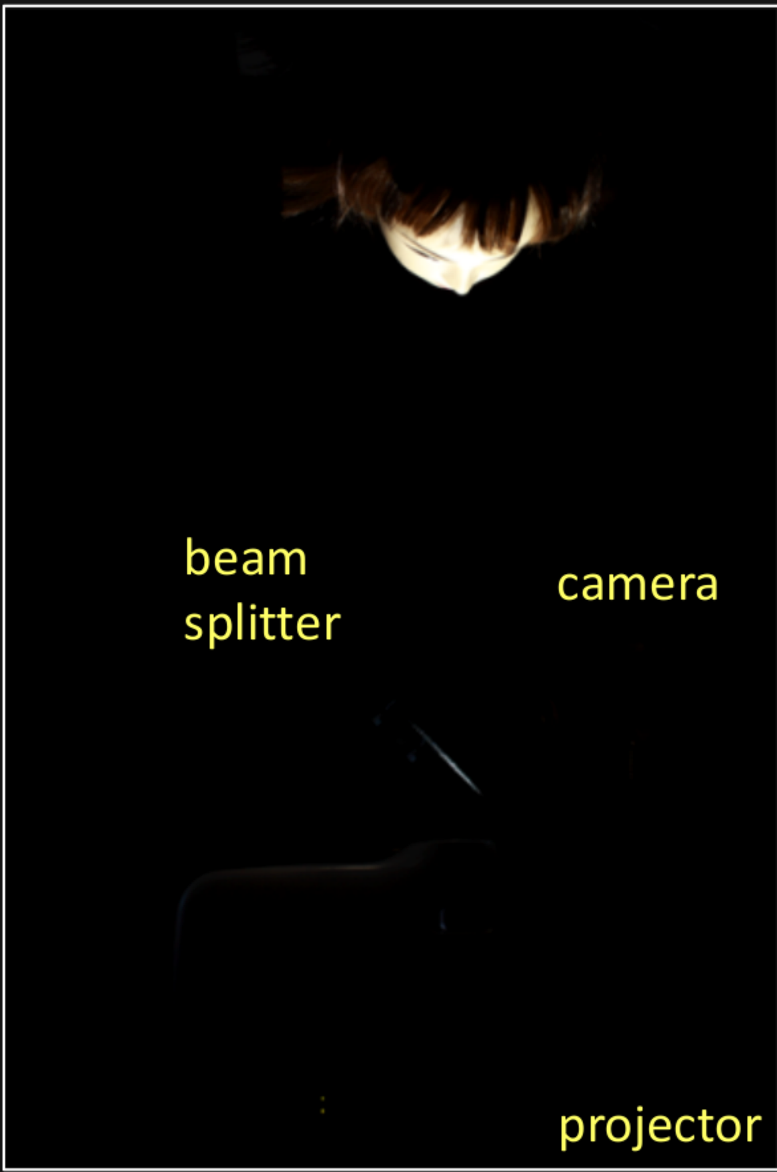

eigenvector of a square matrix T

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

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 $\mathbf{l}, \mathbf{T}\mathbf{l}, \mathbf{T}^2\mathbf{l}, \mathbf{T}^3\mathbf{l}, \ldots$

numerical domain

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

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

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

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

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

$\}$

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

$\Longleftrightarrow$

optical domain

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

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

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

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

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

$\}$

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

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 PowerIt($\mathbf{T}$)

$\mathbf{l}_1 =$ 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 domain

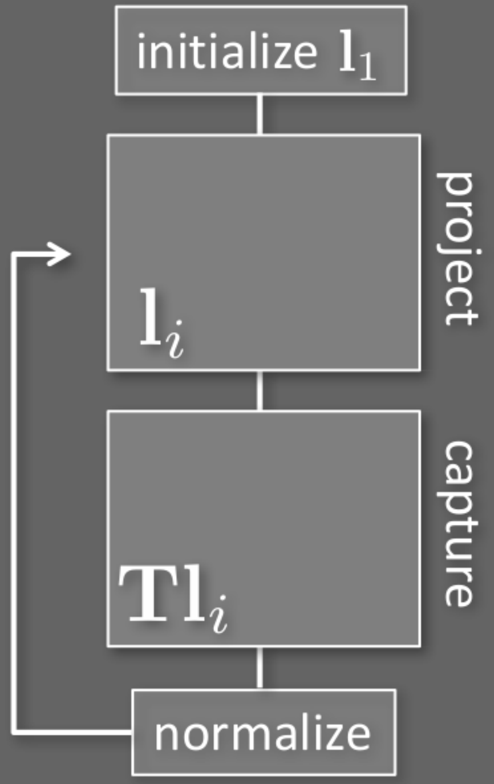

optical power iteration

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

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

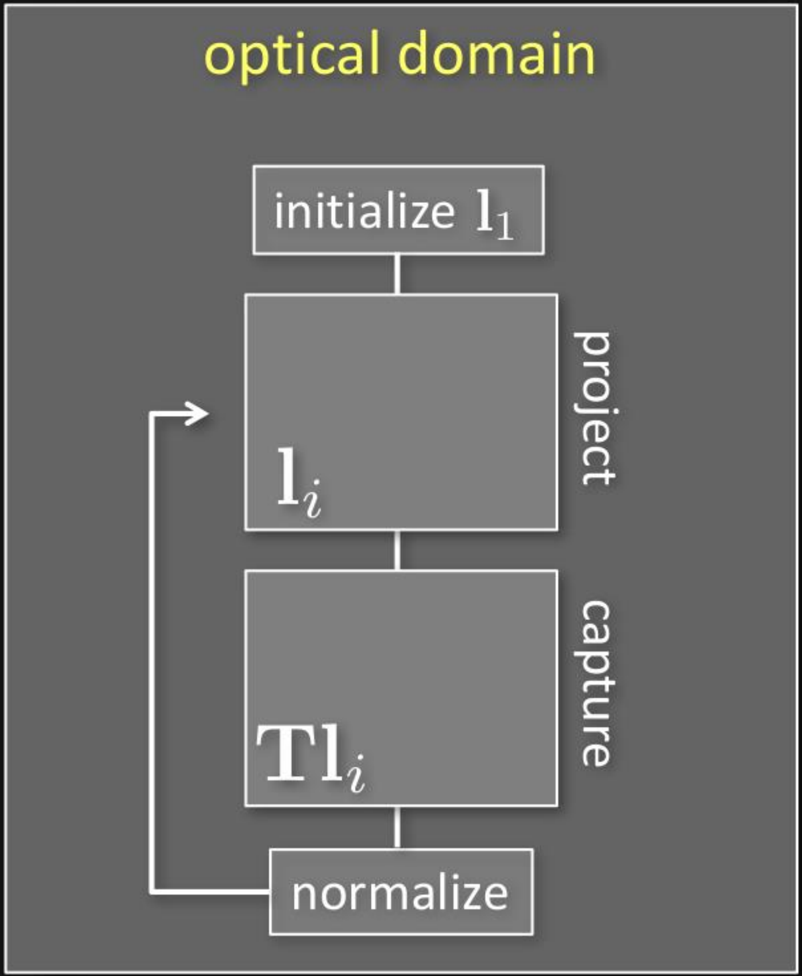

optical power iteration

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

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

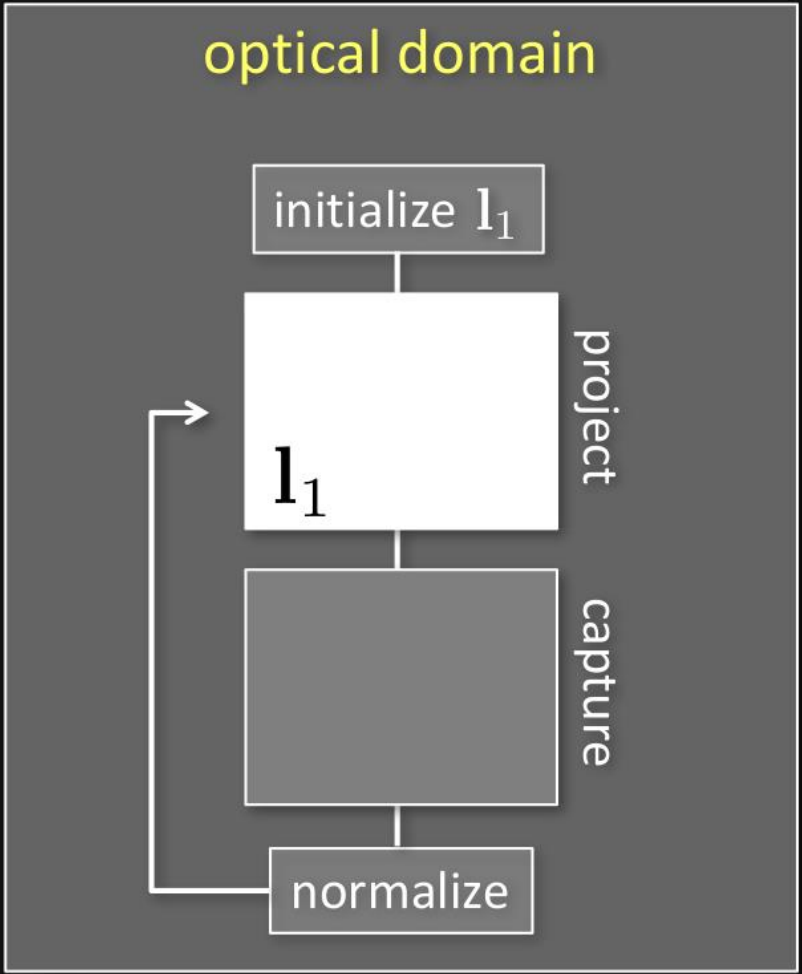

optical power iteration

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

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

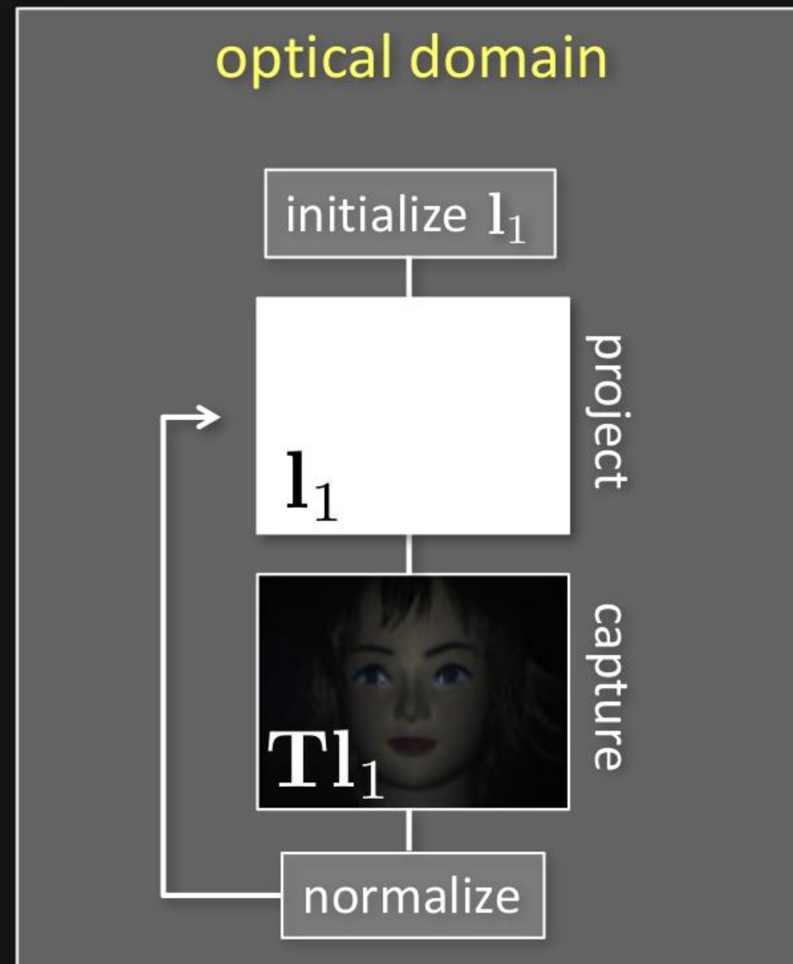

optical power iteration

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

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

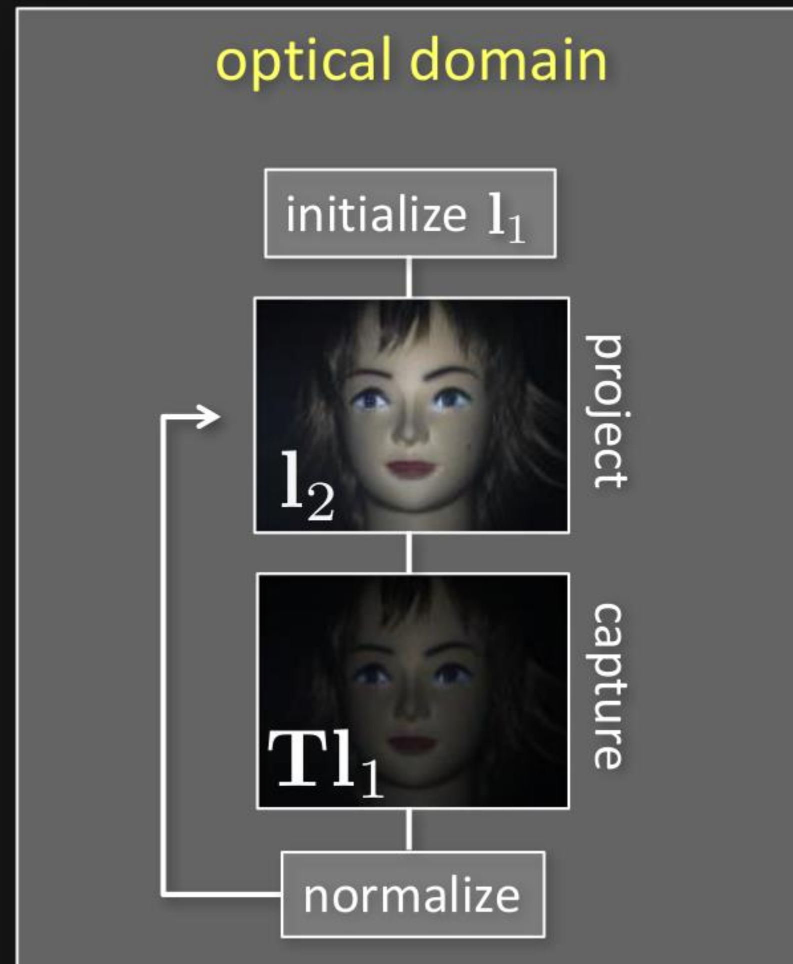

optical power iteration

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

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

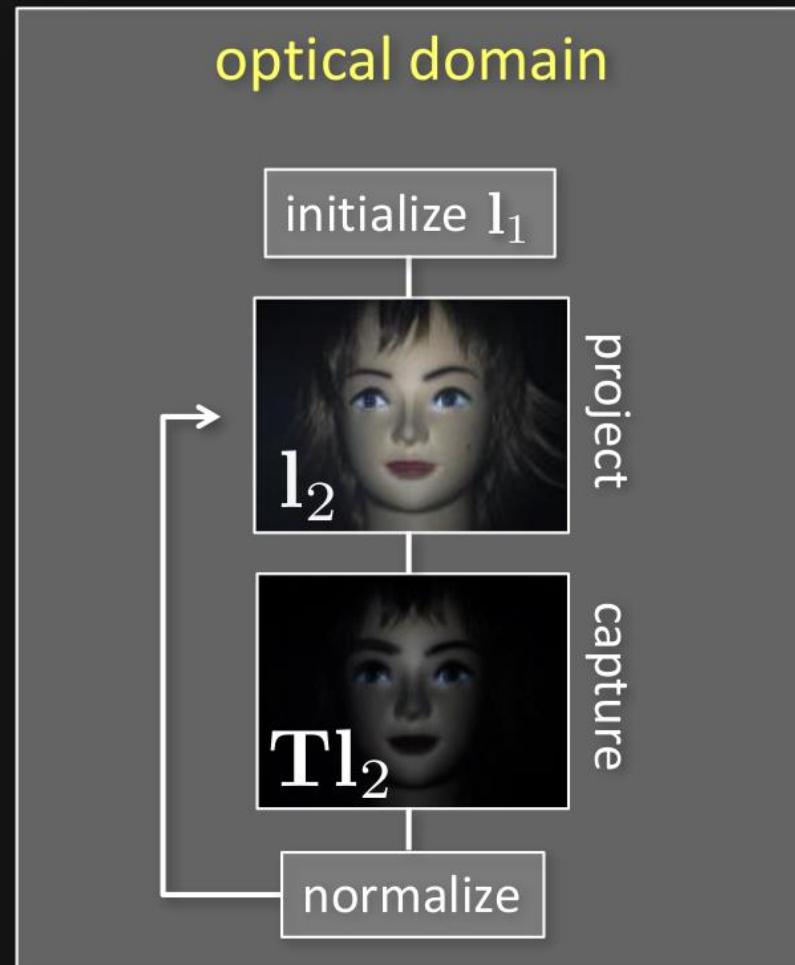

optical power iteration

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

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

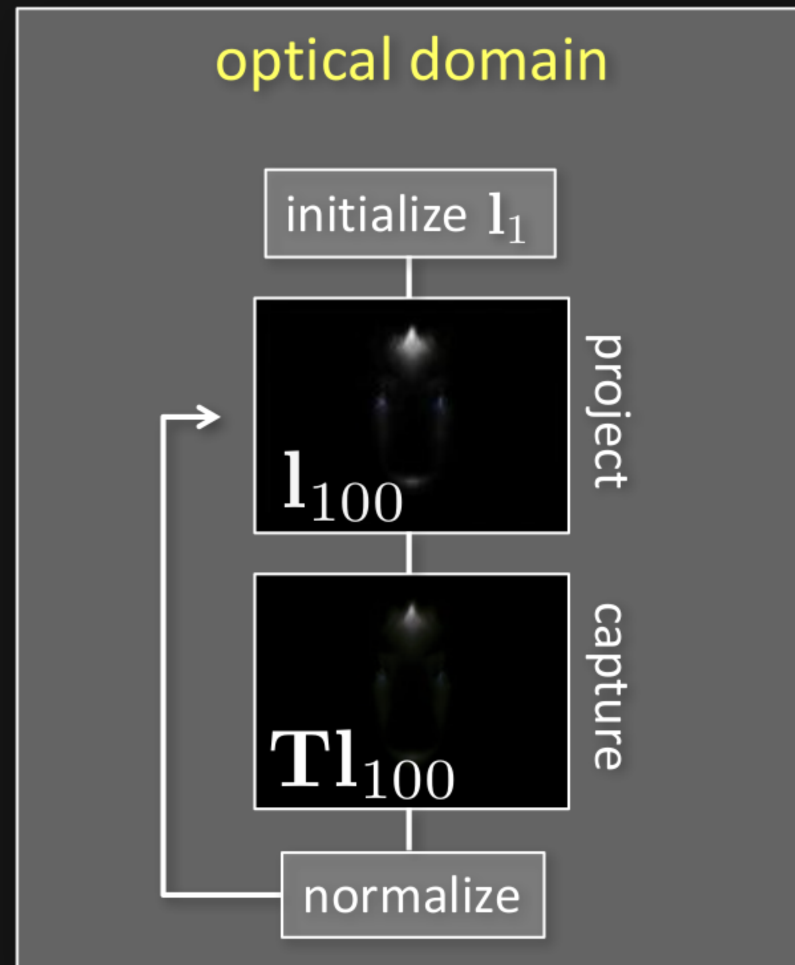

optical power iteration

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

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

optical domain

(approximate)
principal eigenvector

How would you measure the light transport matrix T?

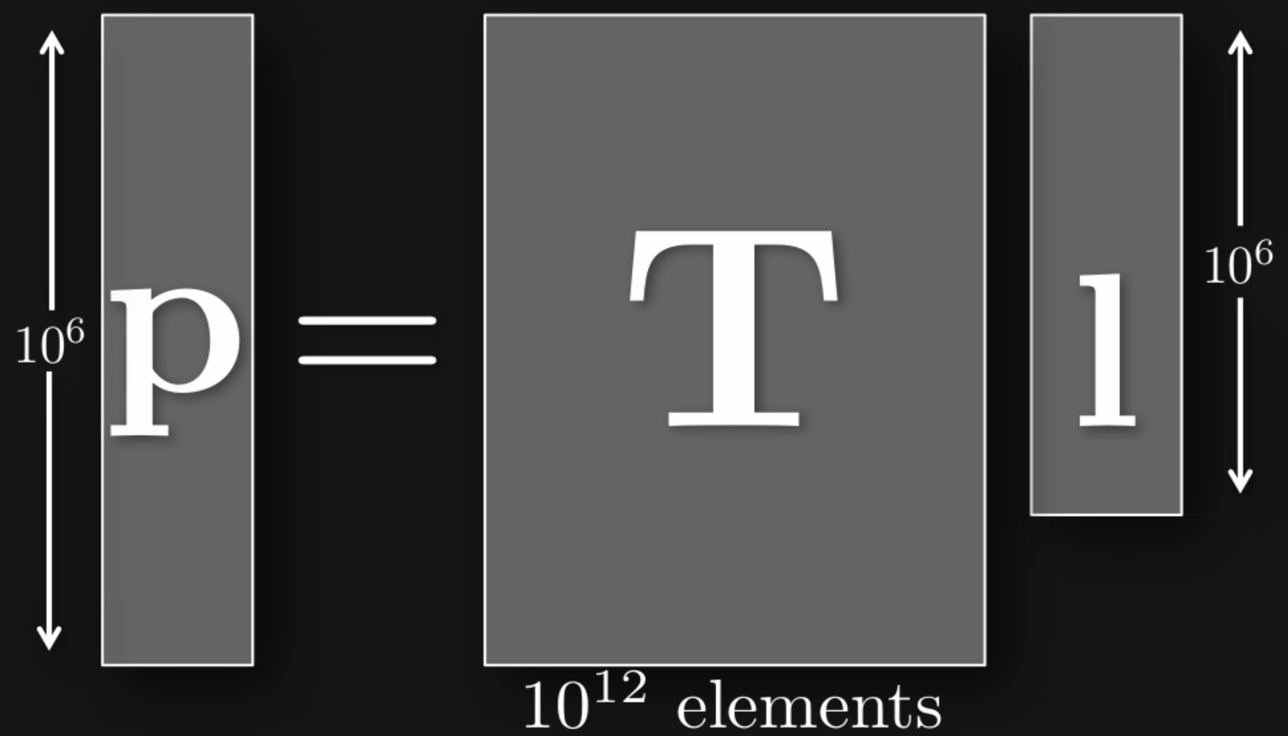

- matrix can be extremely large
- elements not directly accessible
- global structure poorly understood

How would you measure the light transport matrix T?

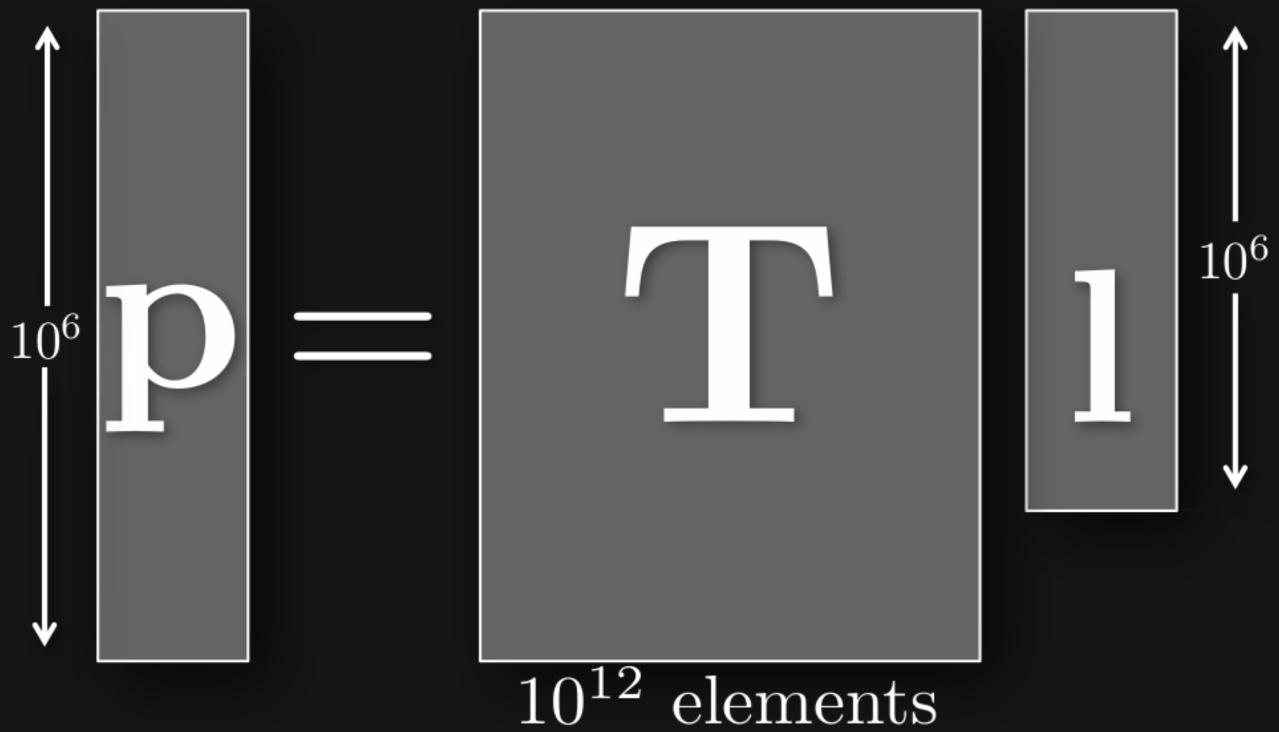

Exhaustive/naïve approach: turn on projector pixels one at a time and take a photo for each of them.

- What does each photo correspond to in T?

How would you measure the light transport matrix T?

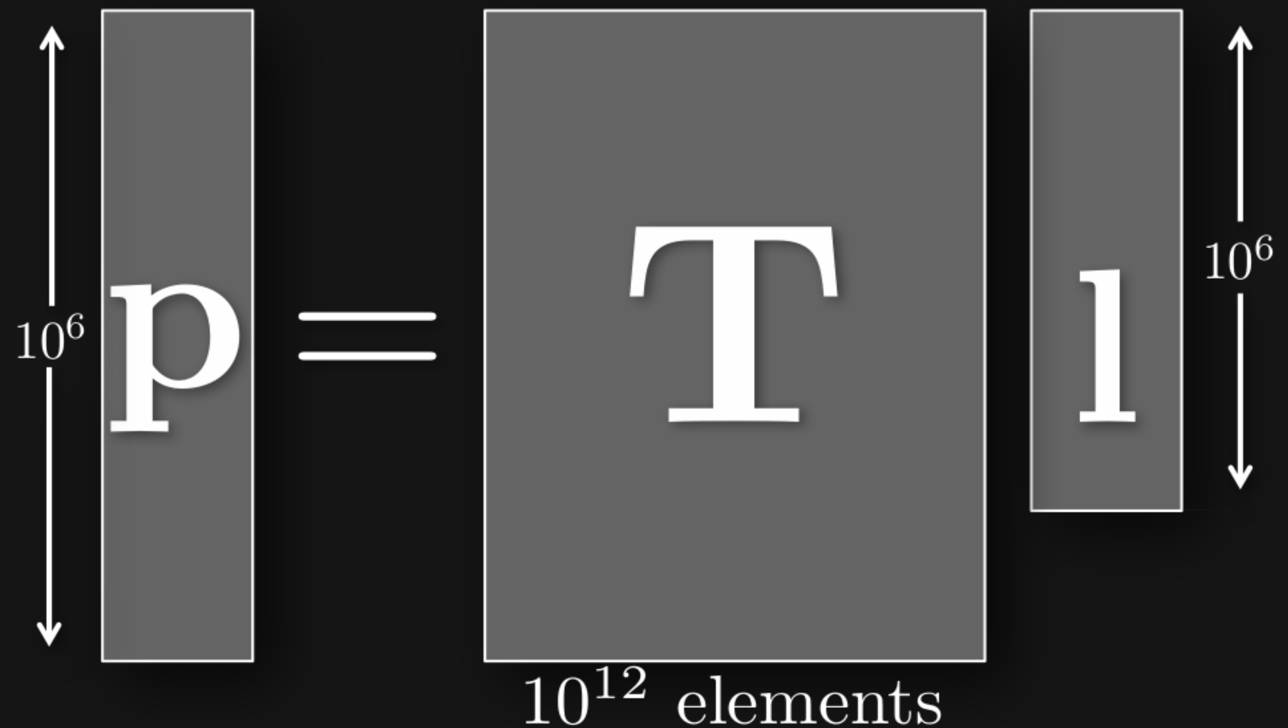

Exhaustive/naïve approach: turn on projector pixels one at a time and take a photo for each of them.

- How many photos do we need to capture?

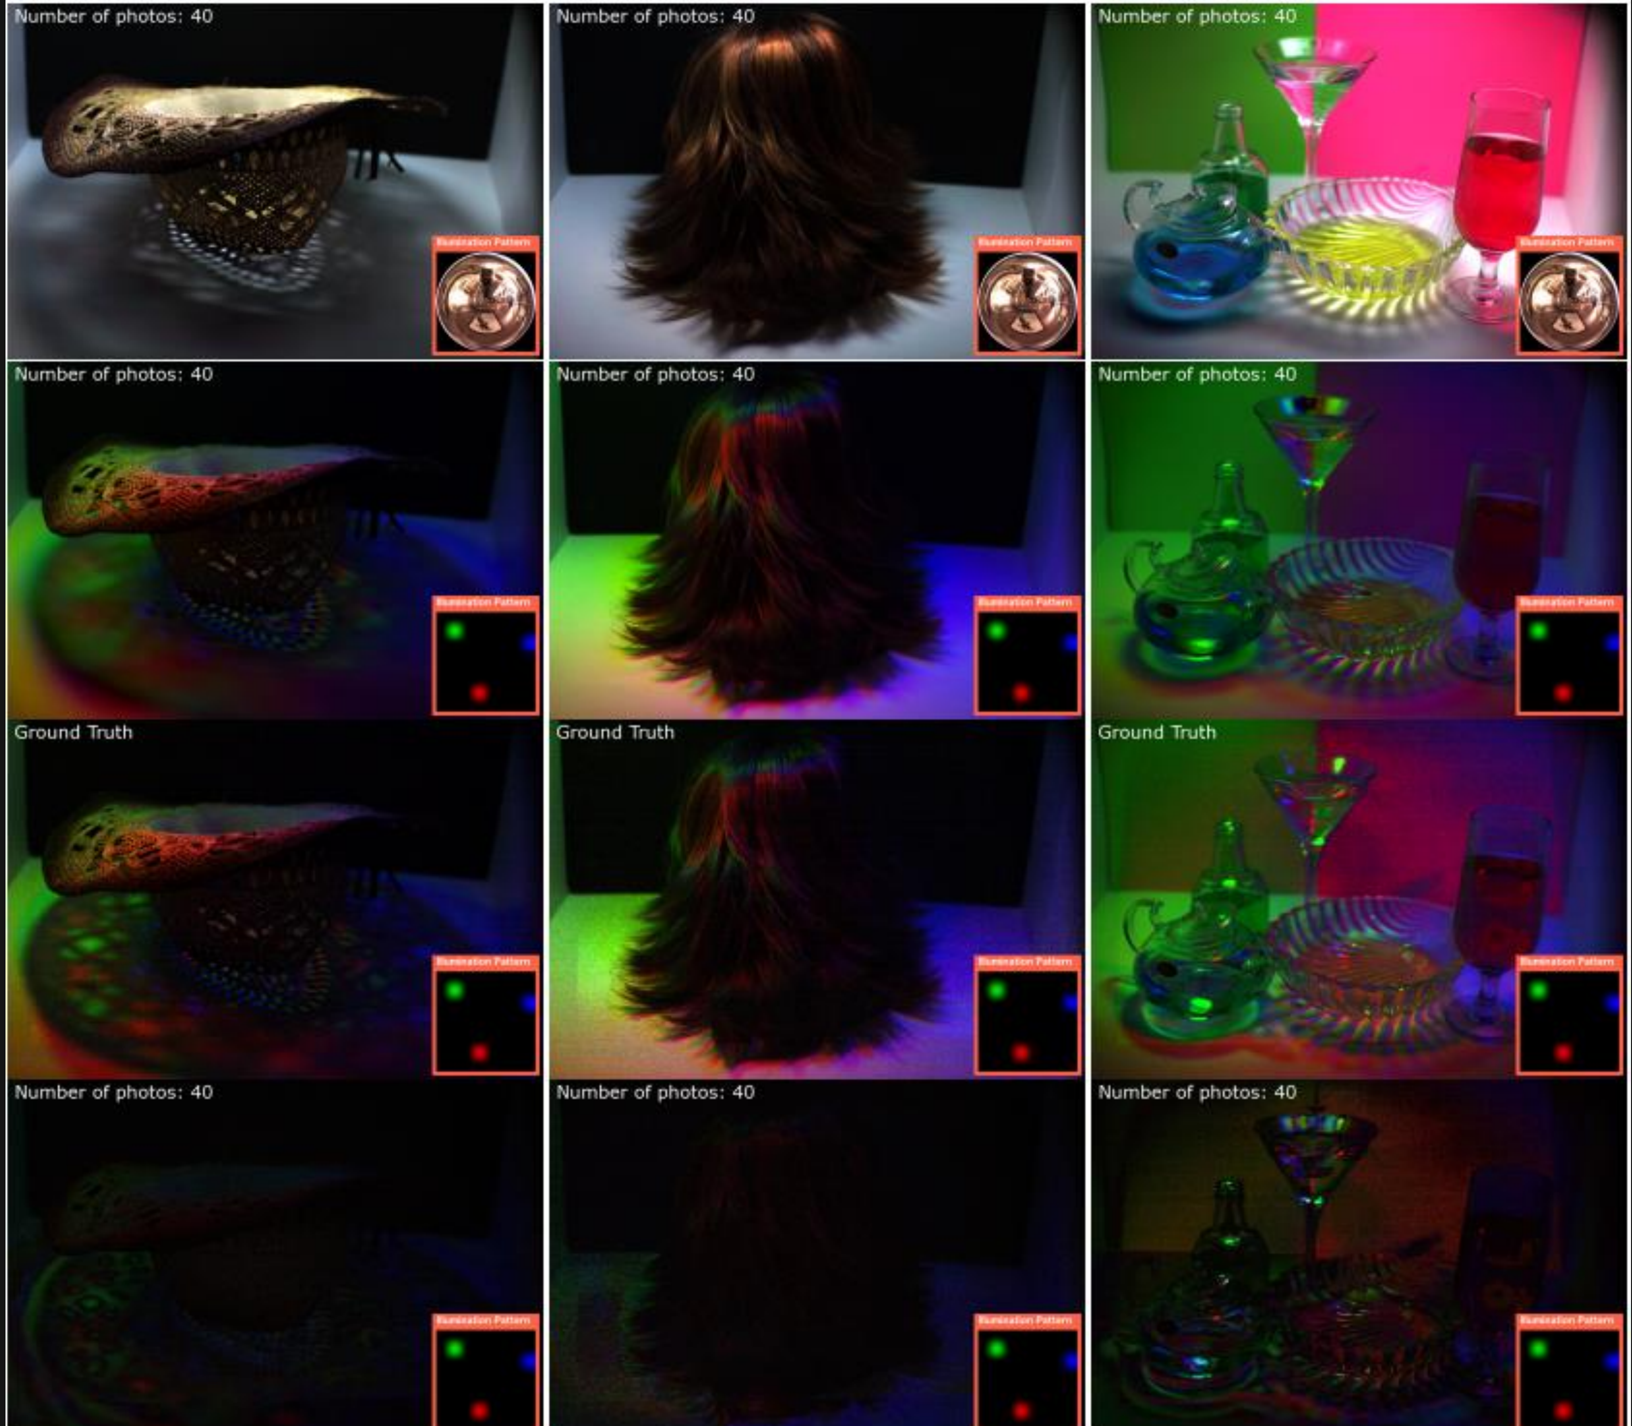
Number of photos: 40
Illumination Pattern
Number of photos: 40
Illumination Pattern
Number of photos: 40
Illumination Pattern
Number of photos: 40
Illumination Pattern
Number of photos: 40
Illumination Pattern
Number of photos: 40
Illumination Pattern
Ground Truth
Illumination Pattern
Ground Truth
Illumination Pattern
Ground Truth
Illumination Pattern
Number of photos: 40
Illumination Pattern
Number of photos: 40
Illumination Pattern
Number of photos: 40
Illumination Pattern

Inverse transport

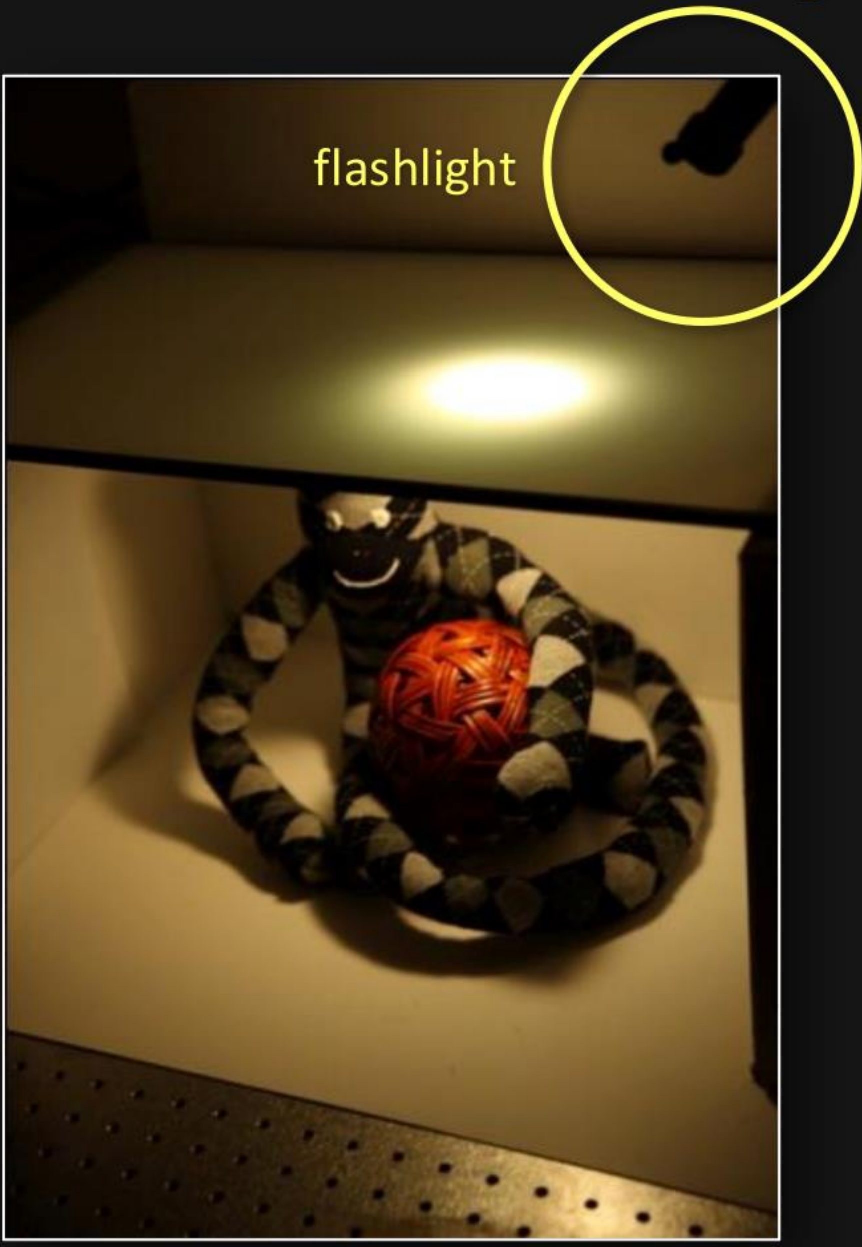
flashlight


diffuser

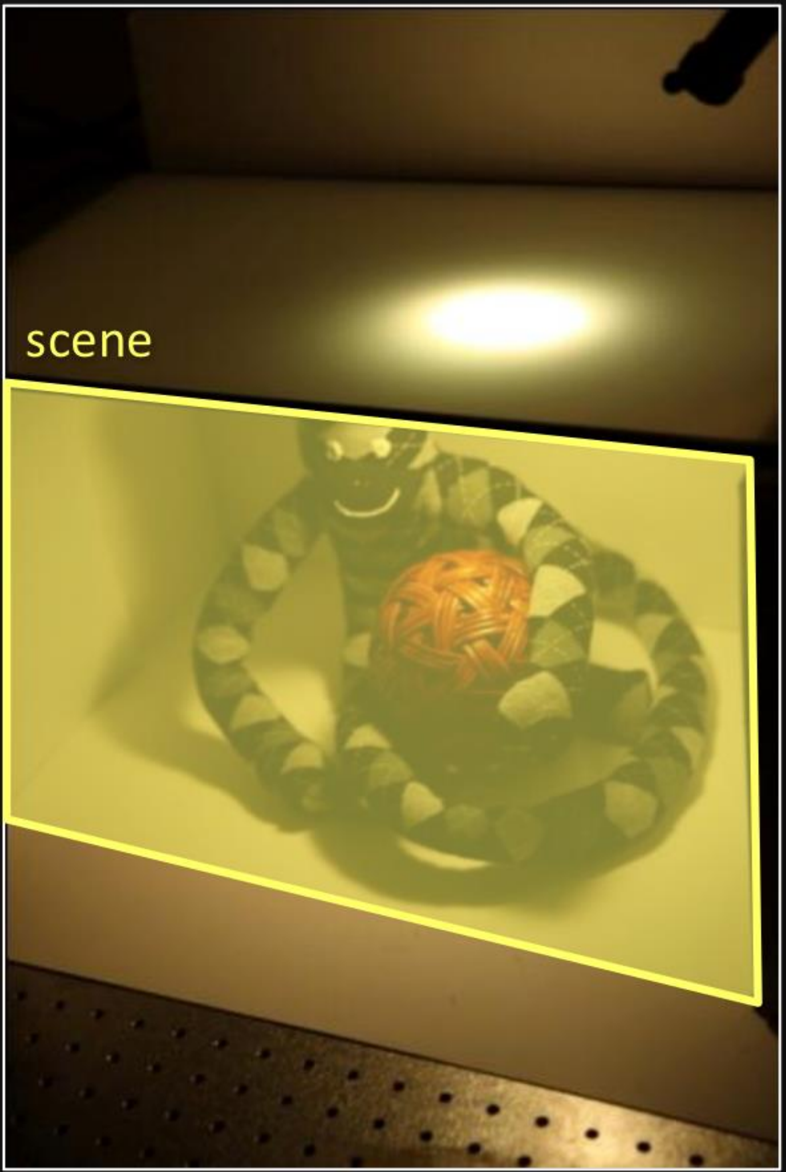
scene

input photo

How do you solve this problem if you know the light transport matrix T?

input photo

illumination

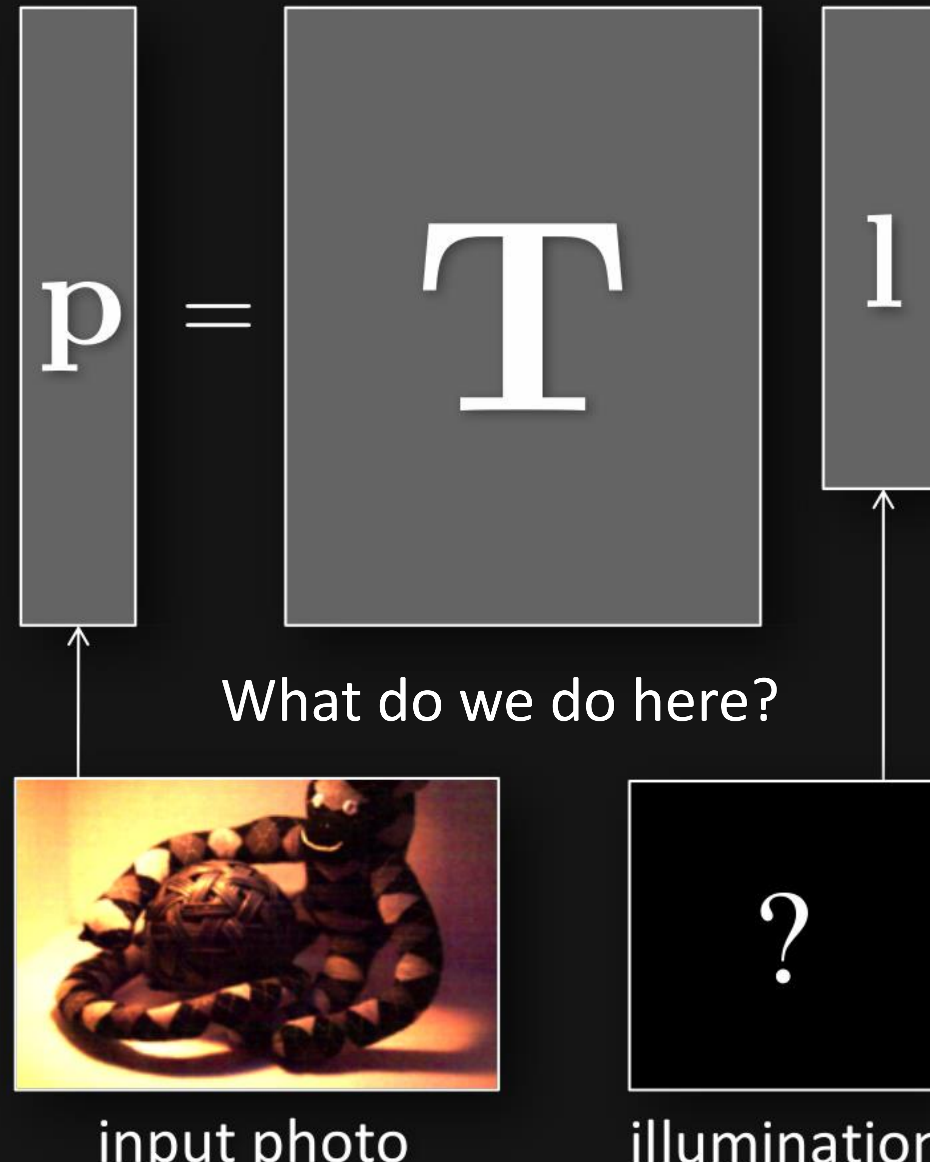

p
=
T
l
What do we do here?
?
input photo
illumination

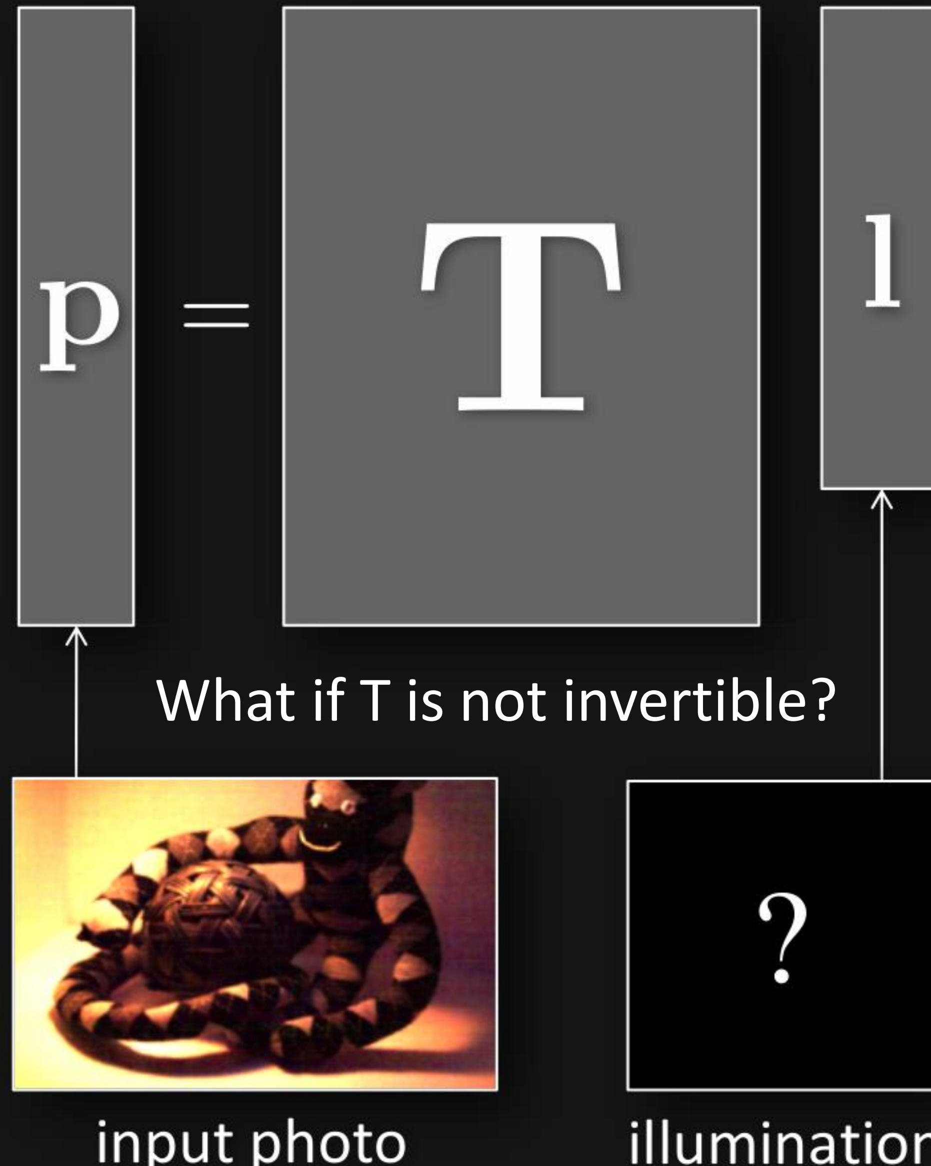
p
=
T
l
What if T is not invertible?
?
input photo
illumination

numerical goal

given photo $\mathbf{p}$, find illumination $\mathbf{l}$ that minimizes

$$\left\| \mathbf{T} \, \mathbf{l} - \mathbf{p} \right\|_2$$

How do you usually solve for l when T is large?

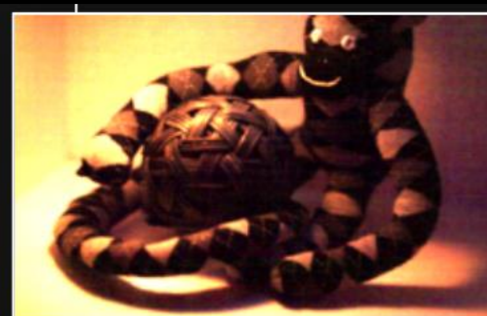

input photo

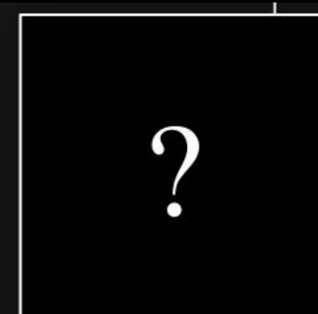

illumination

Reminder from lecture 10: Gradient descent

Given the loss function:

$$E(f) = \|Gf - v\|^2$$

Minimize by iteratively computing:

$$f^{i+1} = f^{i} - \eta^{i} r^{i}, \quad r^{i} = v - Af^{i}, \ \ \eta^{i} = \frac{\left(r^{i}\right)^{T} r^{i}}{(r^{i})^{T} A r^{i}} \qquad \text{for } i = 0, 1, \ldots, N$$

Is this cheaper than the pseudo-inverse approach?

- We never need to compute A, only its products with vectors f, r.
- Vectors f, r are images.
- Because A is the *Laplacian matrix*, these matrix-vector products can be efficiently computed using *convolutions* with the *Laplacian kernel*.

Gradient descent in this case

Given the loss function:

$$E(f) = \|Gf - v\|^2$$

Minimize by iteratively computing:

$$f^{i+1} = f^i - \eta^i r^i, \quad r^i = v - Af^i, \;\; \eta^i = \frac{\left(r^i\right)^T r^i}{(r^i)^T A r^i} \qquad \text{for } i = 0, 1, \ldots, N$$

Is this cheaper than the pseudo-inverse approach?

- We never need to compute A, only its products with vectors f, r.
- ~~Vectors f, r are images.~~ What are f, r in this case?
- ~~Because A is the *Laplacian matrix*, these matrix-vector products can be efficiently computed using *convolutions* with the *Laplacian kernel*.~~ How do we compute matrix-vector products efficiently in this case?

inverting light transport

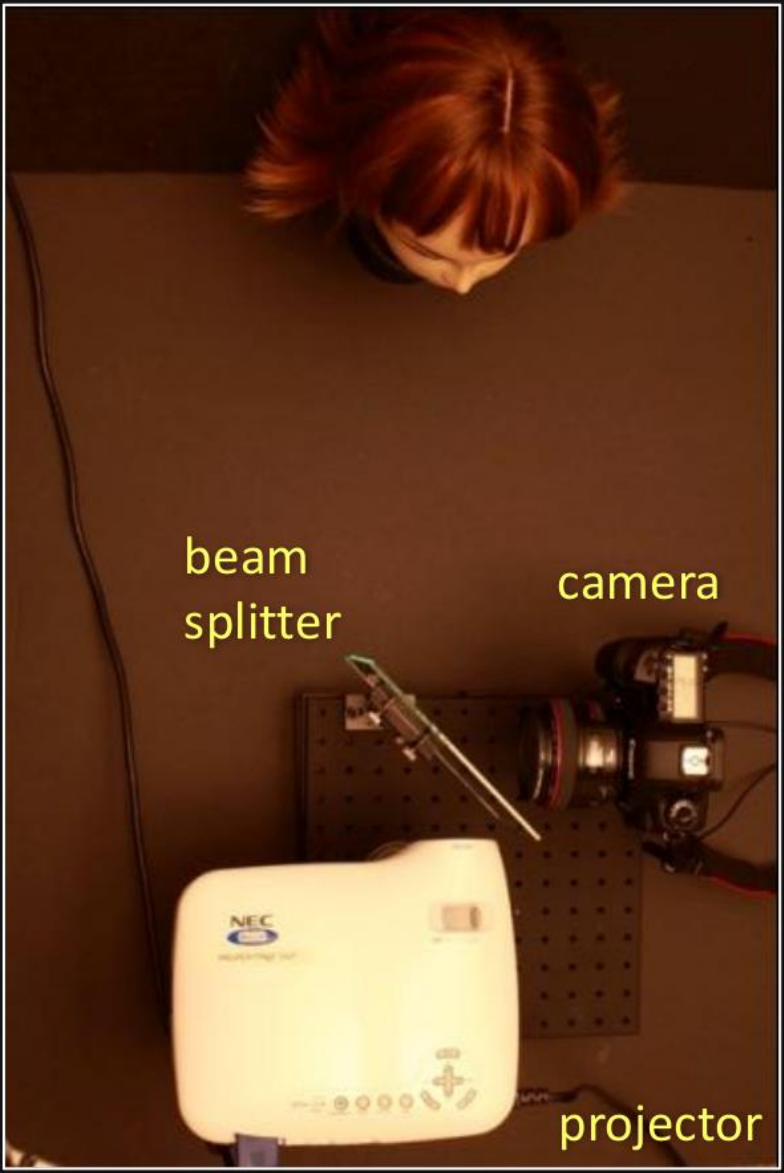

numerical goal

given photo $\mathbf{p}$, find illumination $\mathbf{l}$ that minimizes

$$\left\| \mathbf{T}\,\mathbf{l} - \mathbf{p} \right\|_2$$

remarks

- $\mathbf{T}$ low-rank or high-rank
- $\mathbf{T}$ unknown & not acquired
- illumination sequence will be specific to input photo

inverting light transport

input photo

actual illumination

Dual photography

How do the light transport matrices for these two scenes relate to each other?

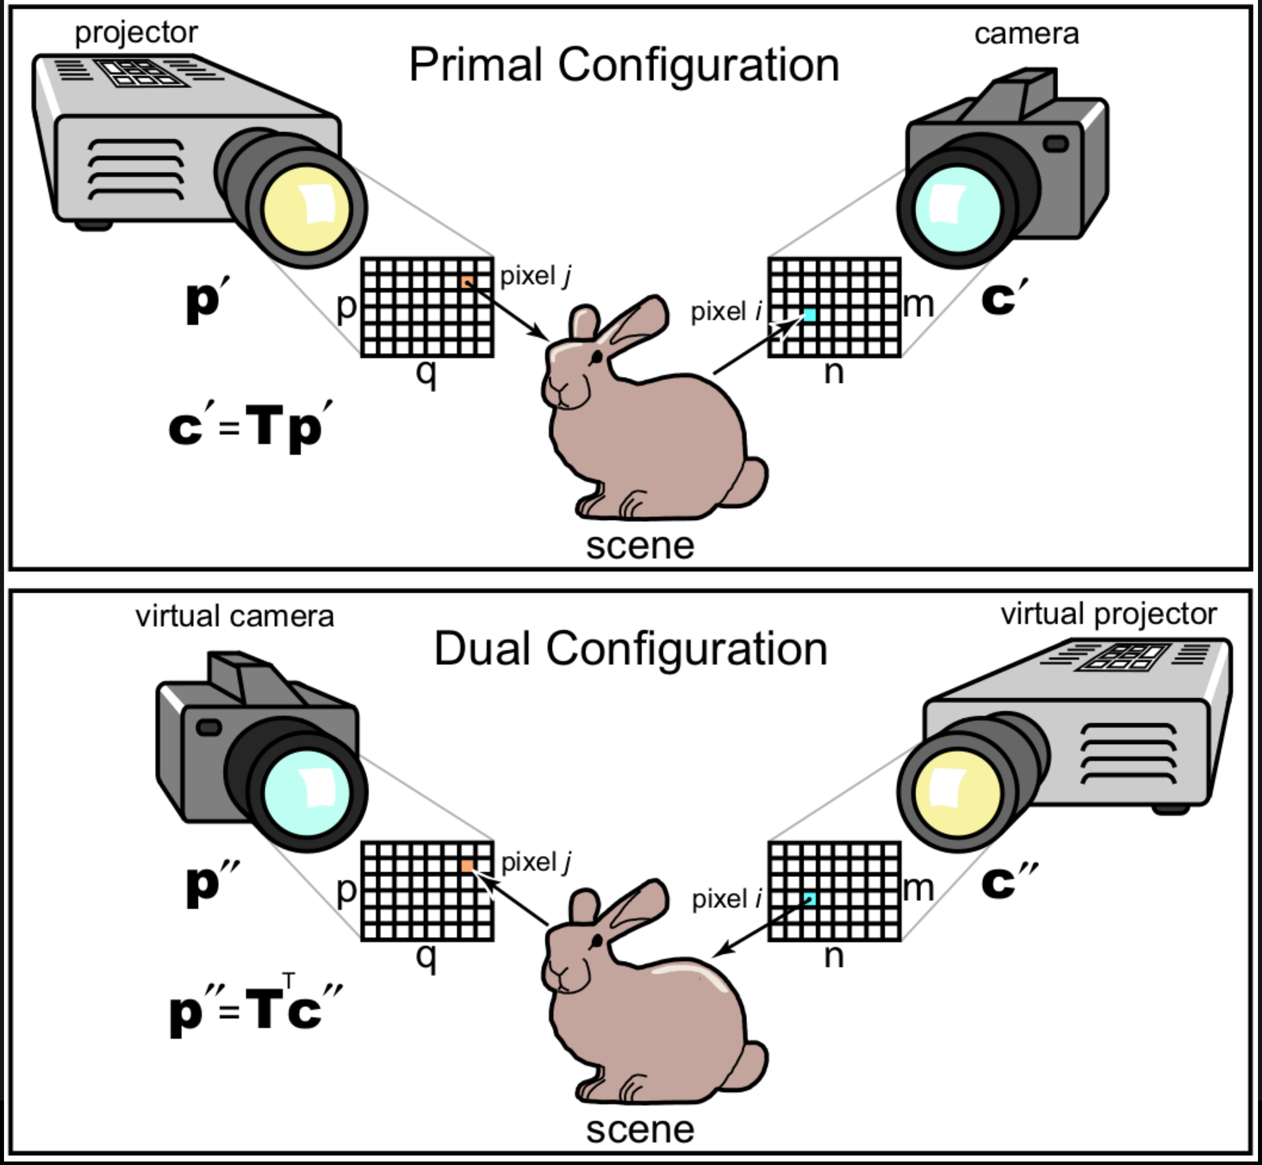

Helmholtz reciprocity: The two matrices are the transpose of each other.

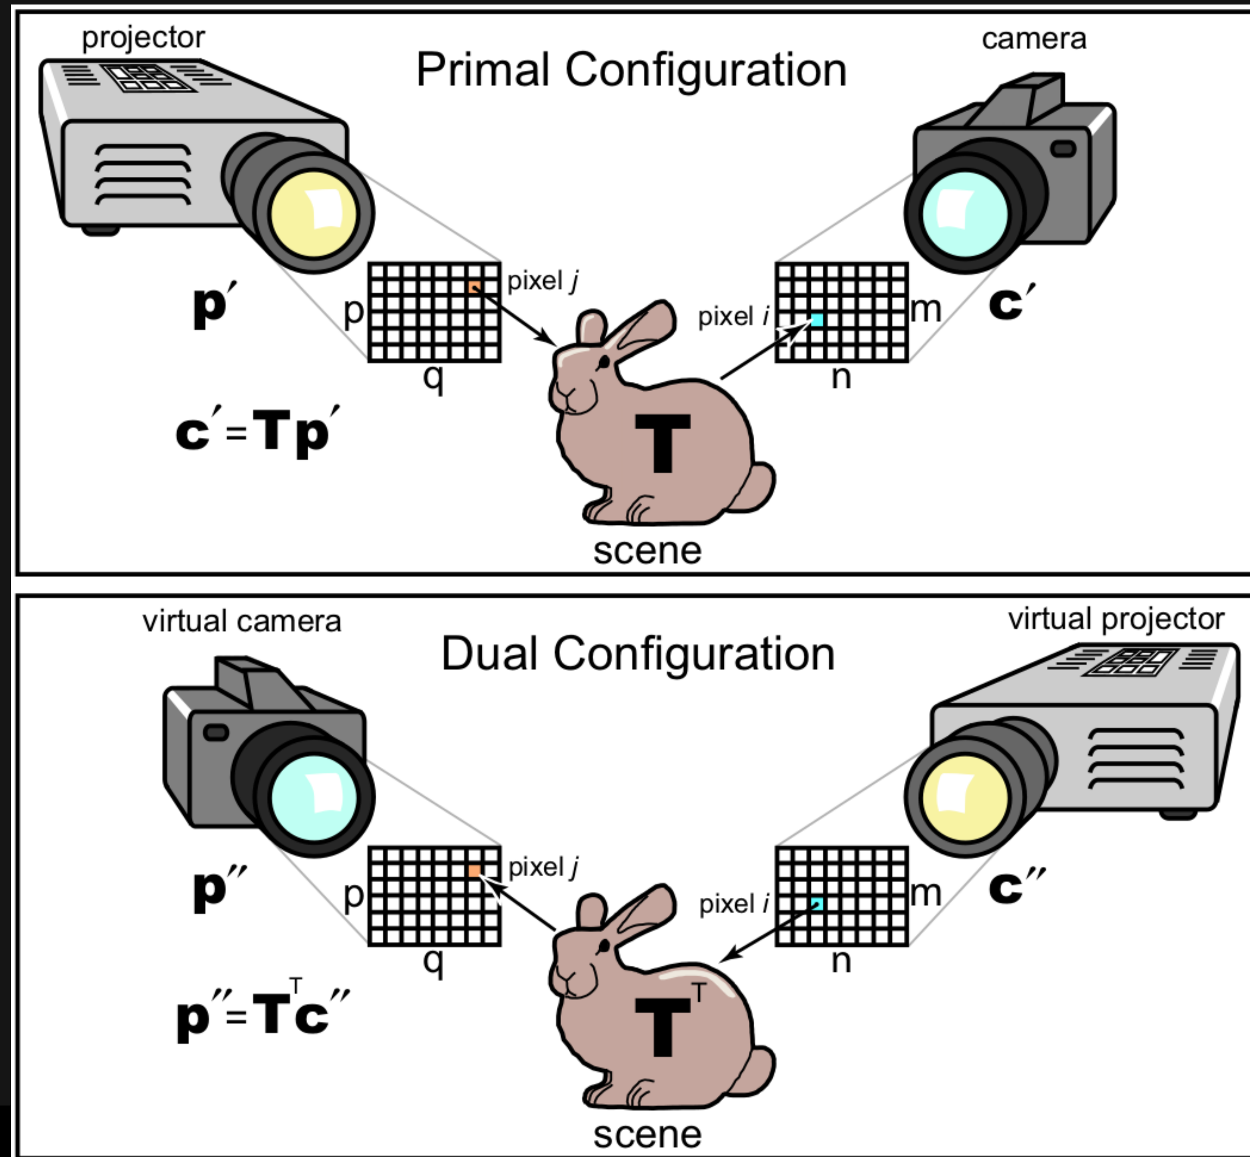

Great demonstration: https://www.youtube.com/watch?v=eV58Ko3iFuI

References

Basic reading:

- Sloan et al., “Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments,” SIGGRAPH 2002.
- Ng et al., “All-frequency shadows using non-linear wavelet lighting approximation,” SIGGRAPH 2003.
- Seitz et al., “A theory of inverse light transport,” ICCV 2005.

  These three papers all discuss the concept of light transport matrix in detail.
- Debevec et al., “Acquiring the reflectance field of a human face,” SIGGRAPH 2000.

  The paper on image-based relighting.
- Woodham et al., “Photometric stereo: A reflectance map technique for determining surface orientation from image intensity,” IUSIA 1979.

  The original photometric stereo paper.
- O’Toole and Kutulakos, “Optical computing for fast light transport analysis,” SIGGRAPH Asia 2010.

  The paper on eigenanalysis and optical computing using light transport matrices.
- Sen et al., “Dual photography,” SIGGRAPH 2005.

  The dual photography paper.

Additional reading:

- Peers et al., “Compressive light transport sensing,” TOG 2009.
- Wang et al., “Kernel Nyström method for light transport,” SIGGRAPH 2009.

  These two papers discuss alternative ways for efficient acquisition of the light transport matrix, using assumptions on its algebraic structure.