Photometric stereo

15-463, 15-663, 15-862
Computational Photography
Fall 2020, Lecture 18

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

Course announcements

- Homework assignment 5 is due on November 16th.
    - Large bonus component.
    - Make sure to start early, as photometric stereo data acquisition is tricky.

- Final project logistics:
    - Final project presentations scheduled for December 17th.
    - I will email those of you needing equipment separately for pickups, starting tomorrow.

Overview of today's lecture

- Light sources.
- Some notes about radiometry.
- Photometric stereo.
- Uncalibrated photometric stereo.
- Generalized bas-relief ambiguity.

Slide credits

Many of these slides were adapted from:

- Srinivasa Narasimhan (16-385, Spring 2014).
- Todd Zickler (Harvard University).
- Steven Gortler (Harvard University).
- Kayvon Fatahalian (Stanford University; CMU 15-462, Fall 2015).

Light sources

“Physics-based” computer vision (a.k.a “inverse optics”)

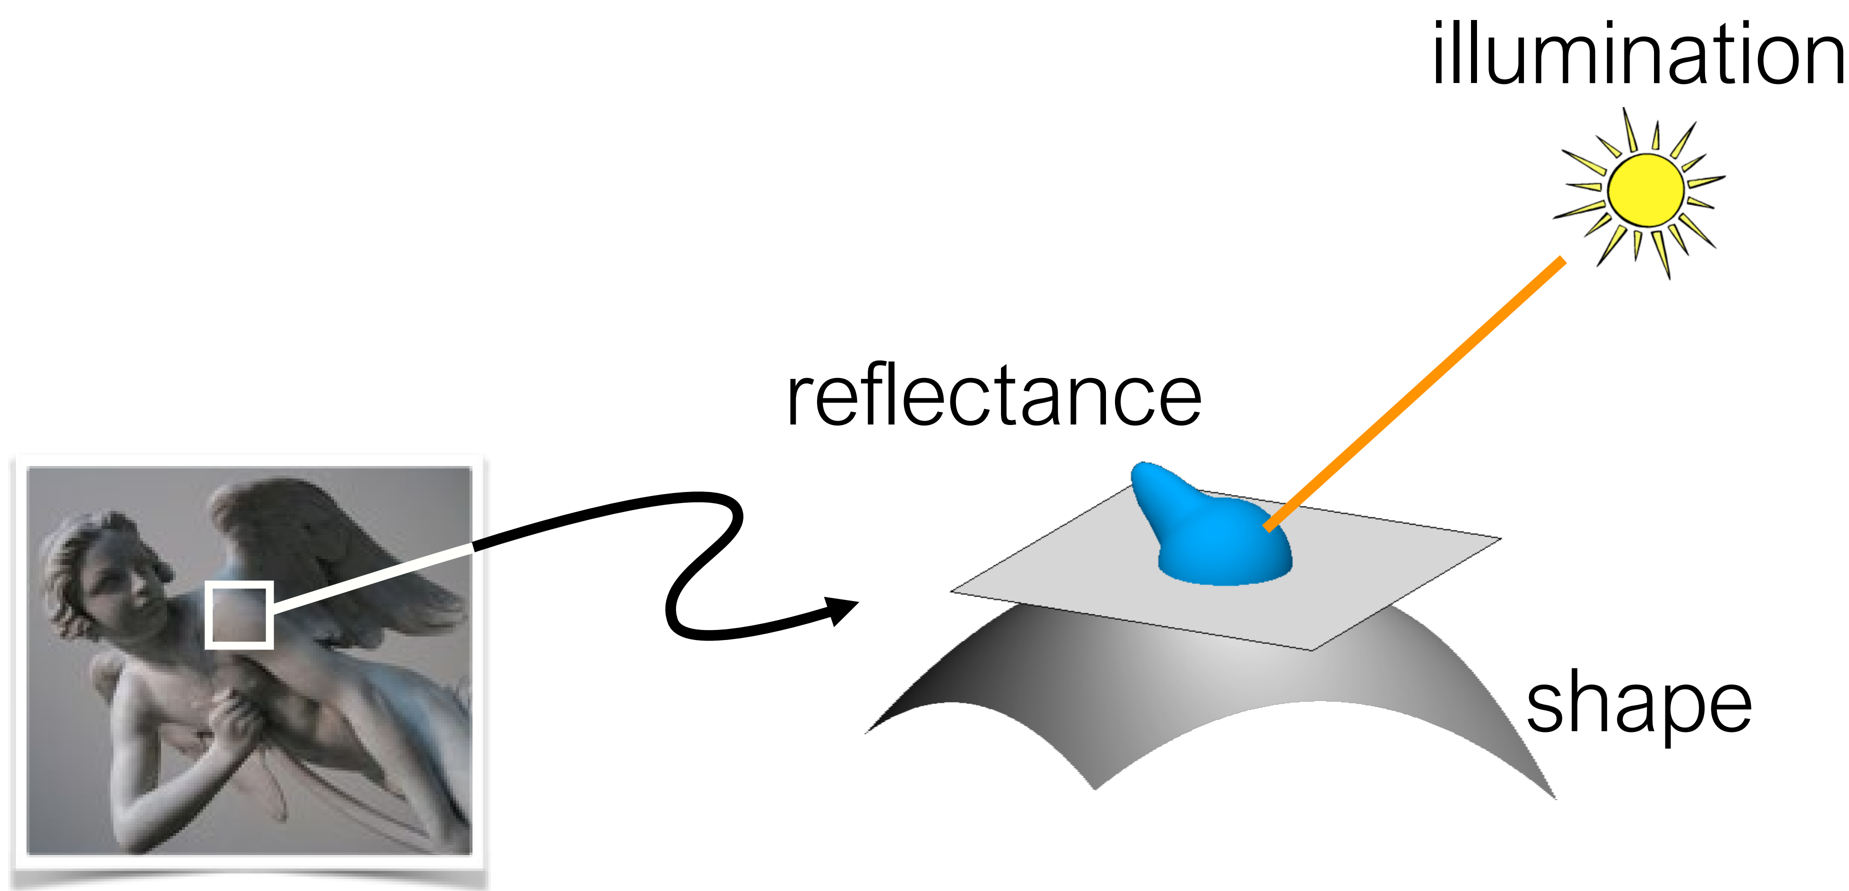

$\mathbf{I} \Longrightarrow$ shape, illumination, reflectance

Lighting models: Plenoptic function

- Radiance as a function of position and direction
- Radiance as a function of position, direction, and time
- Spectral radiance as a function of position, direction, time and wavelength

$$L(x, \omega, t, \lambda)$$

5D

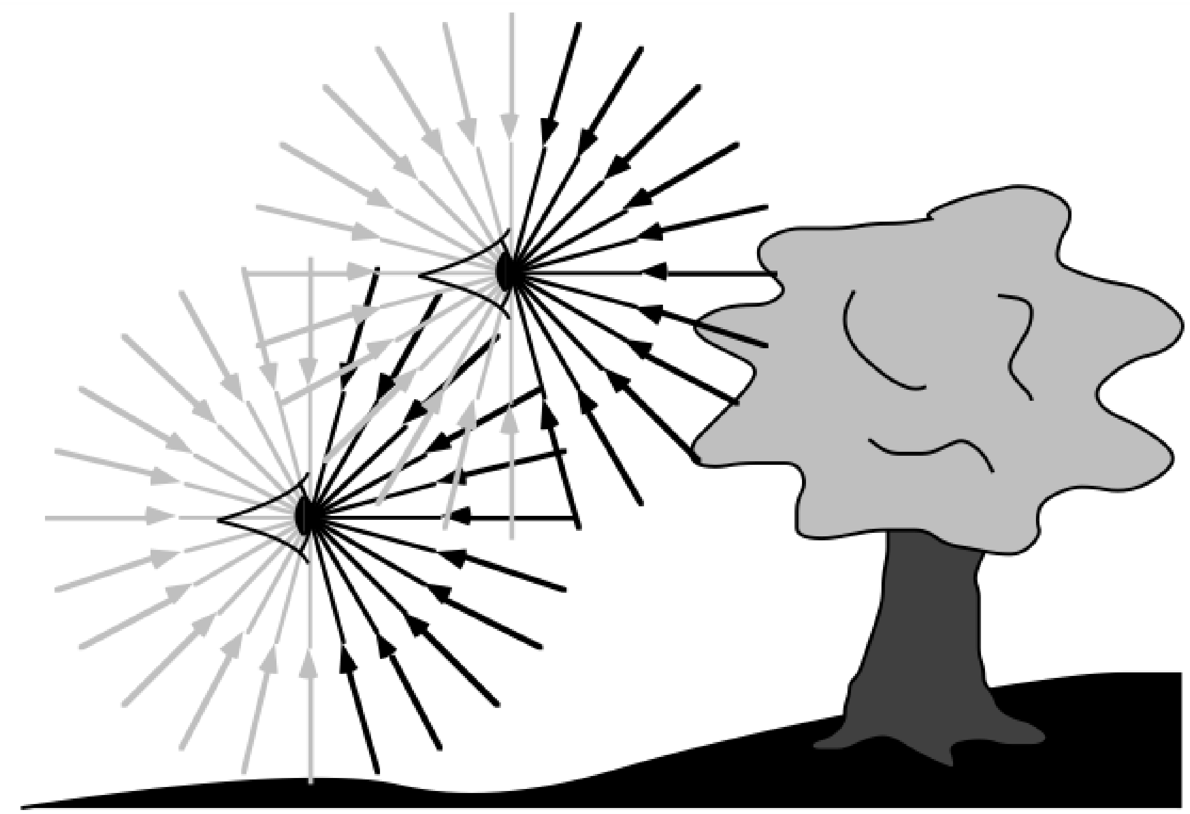

Fig.1.3
The plenoptic function describes the information available to an observer at any point in space and time. Shown here are two schematic eyes-which one should consider to have punctate pupils-gathering pencils of light rays. A real observer cannot see the light rays coming from behind, but the plenoptic function does include these rays.

[Adelson and Bergen, 1991]

Lighting models: far-field (or directional) approximation

- Assume that, over the observed region of interest, all source of incoming flux are relatively far away

$$L(x, \omega, t, \lambda) \longrightarrow L(\omega, t, \lambda)$$

$$L(x, \omega) \longrightarrow L(\omega)$$

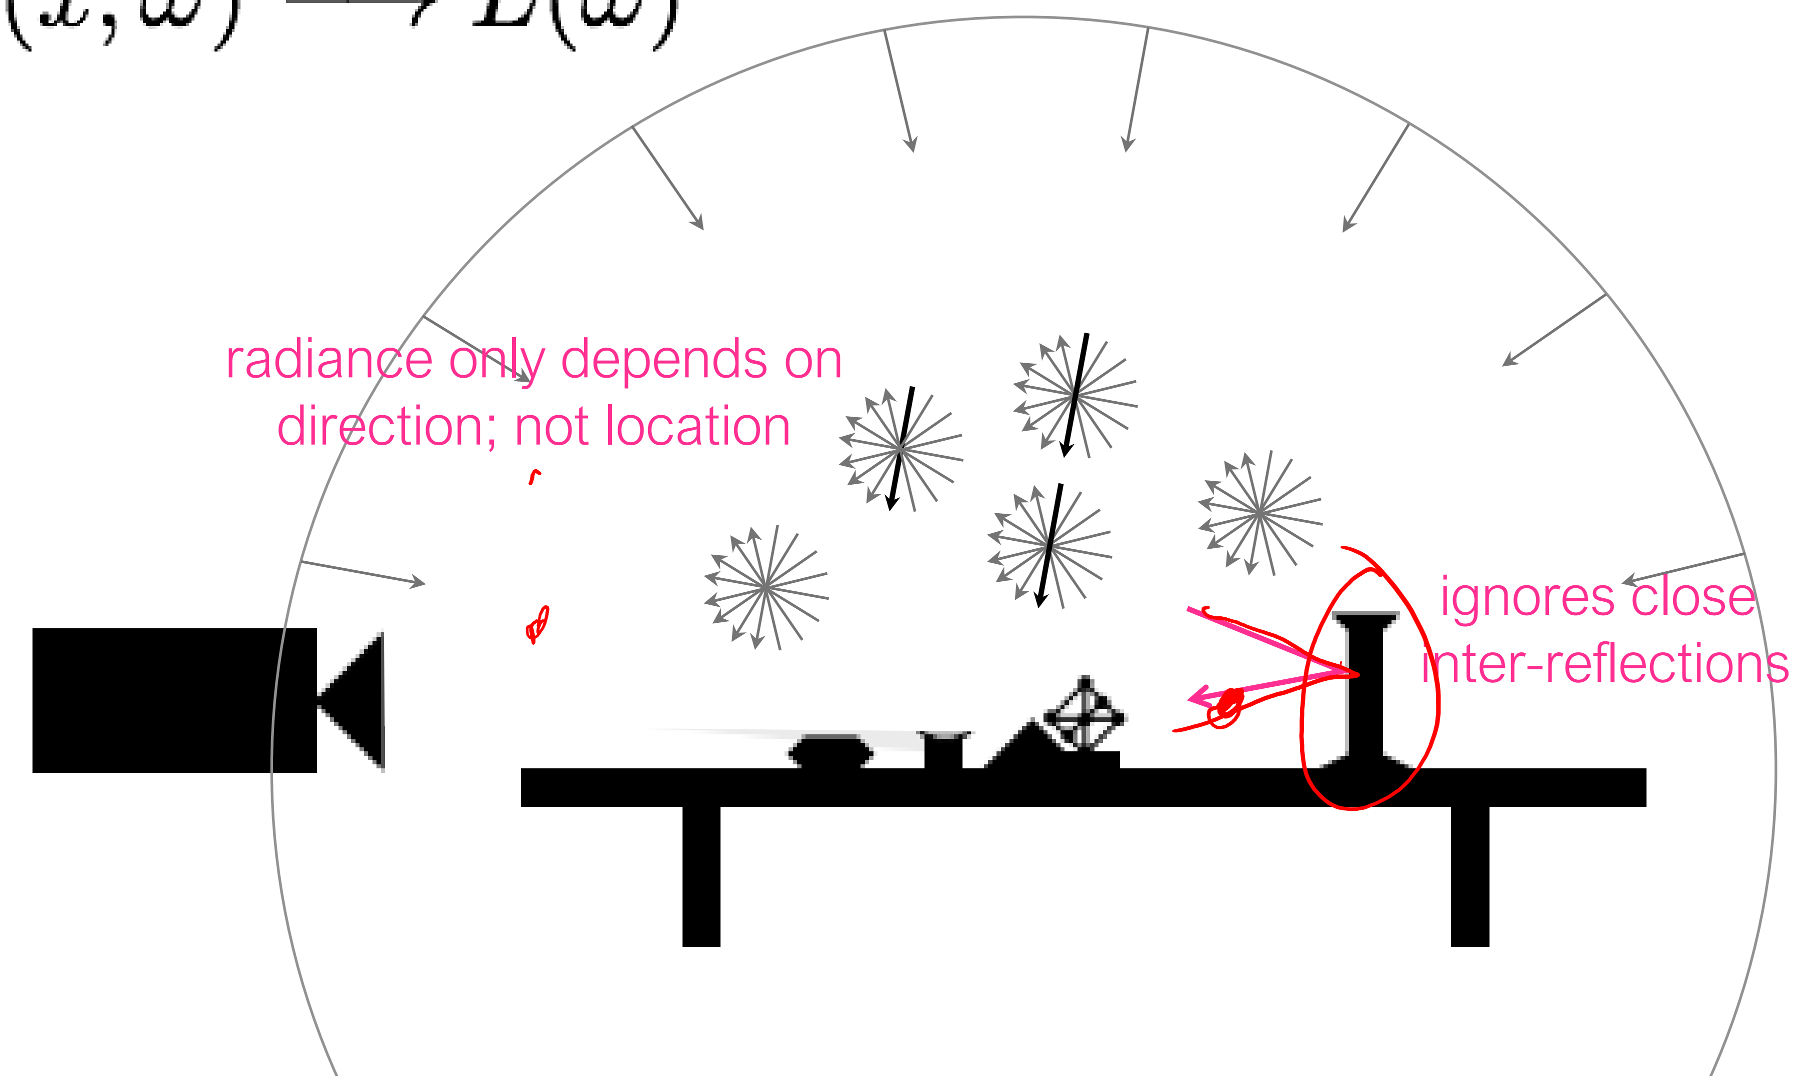

[Debevec, 1998]

Application: augmented reality

$L(\omega)$

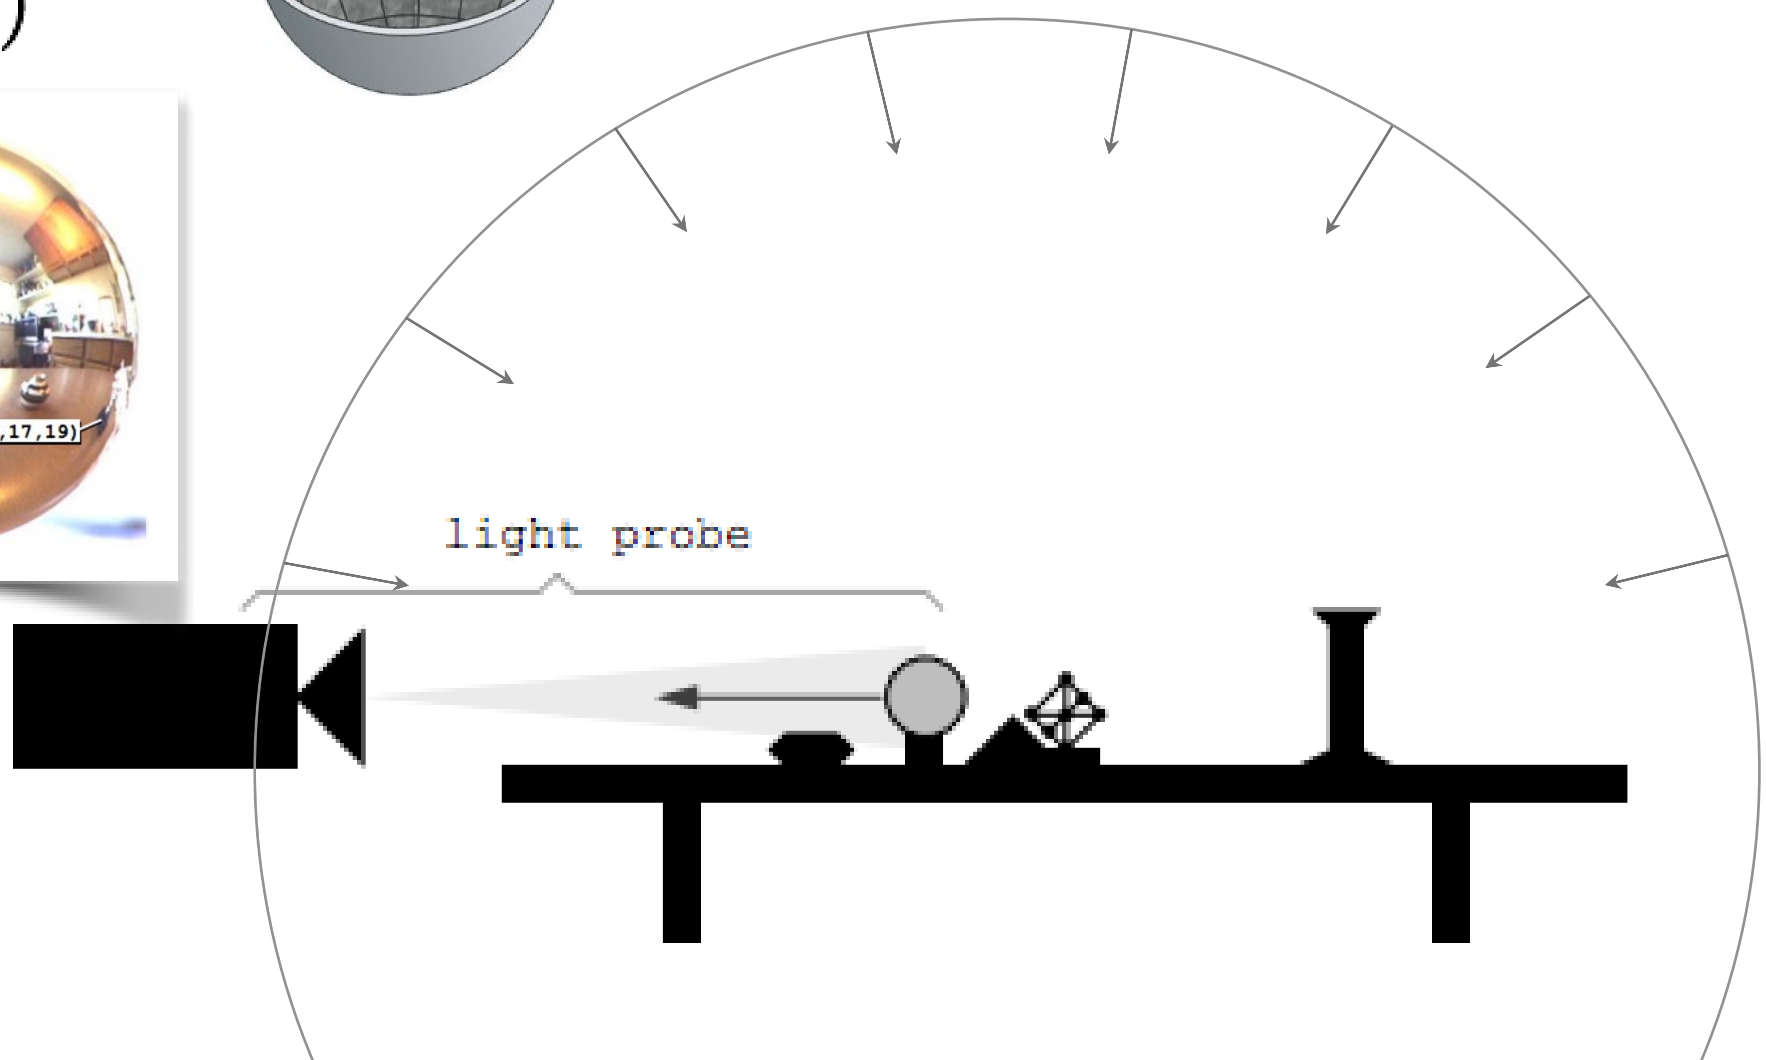

[Debevec, 1998]

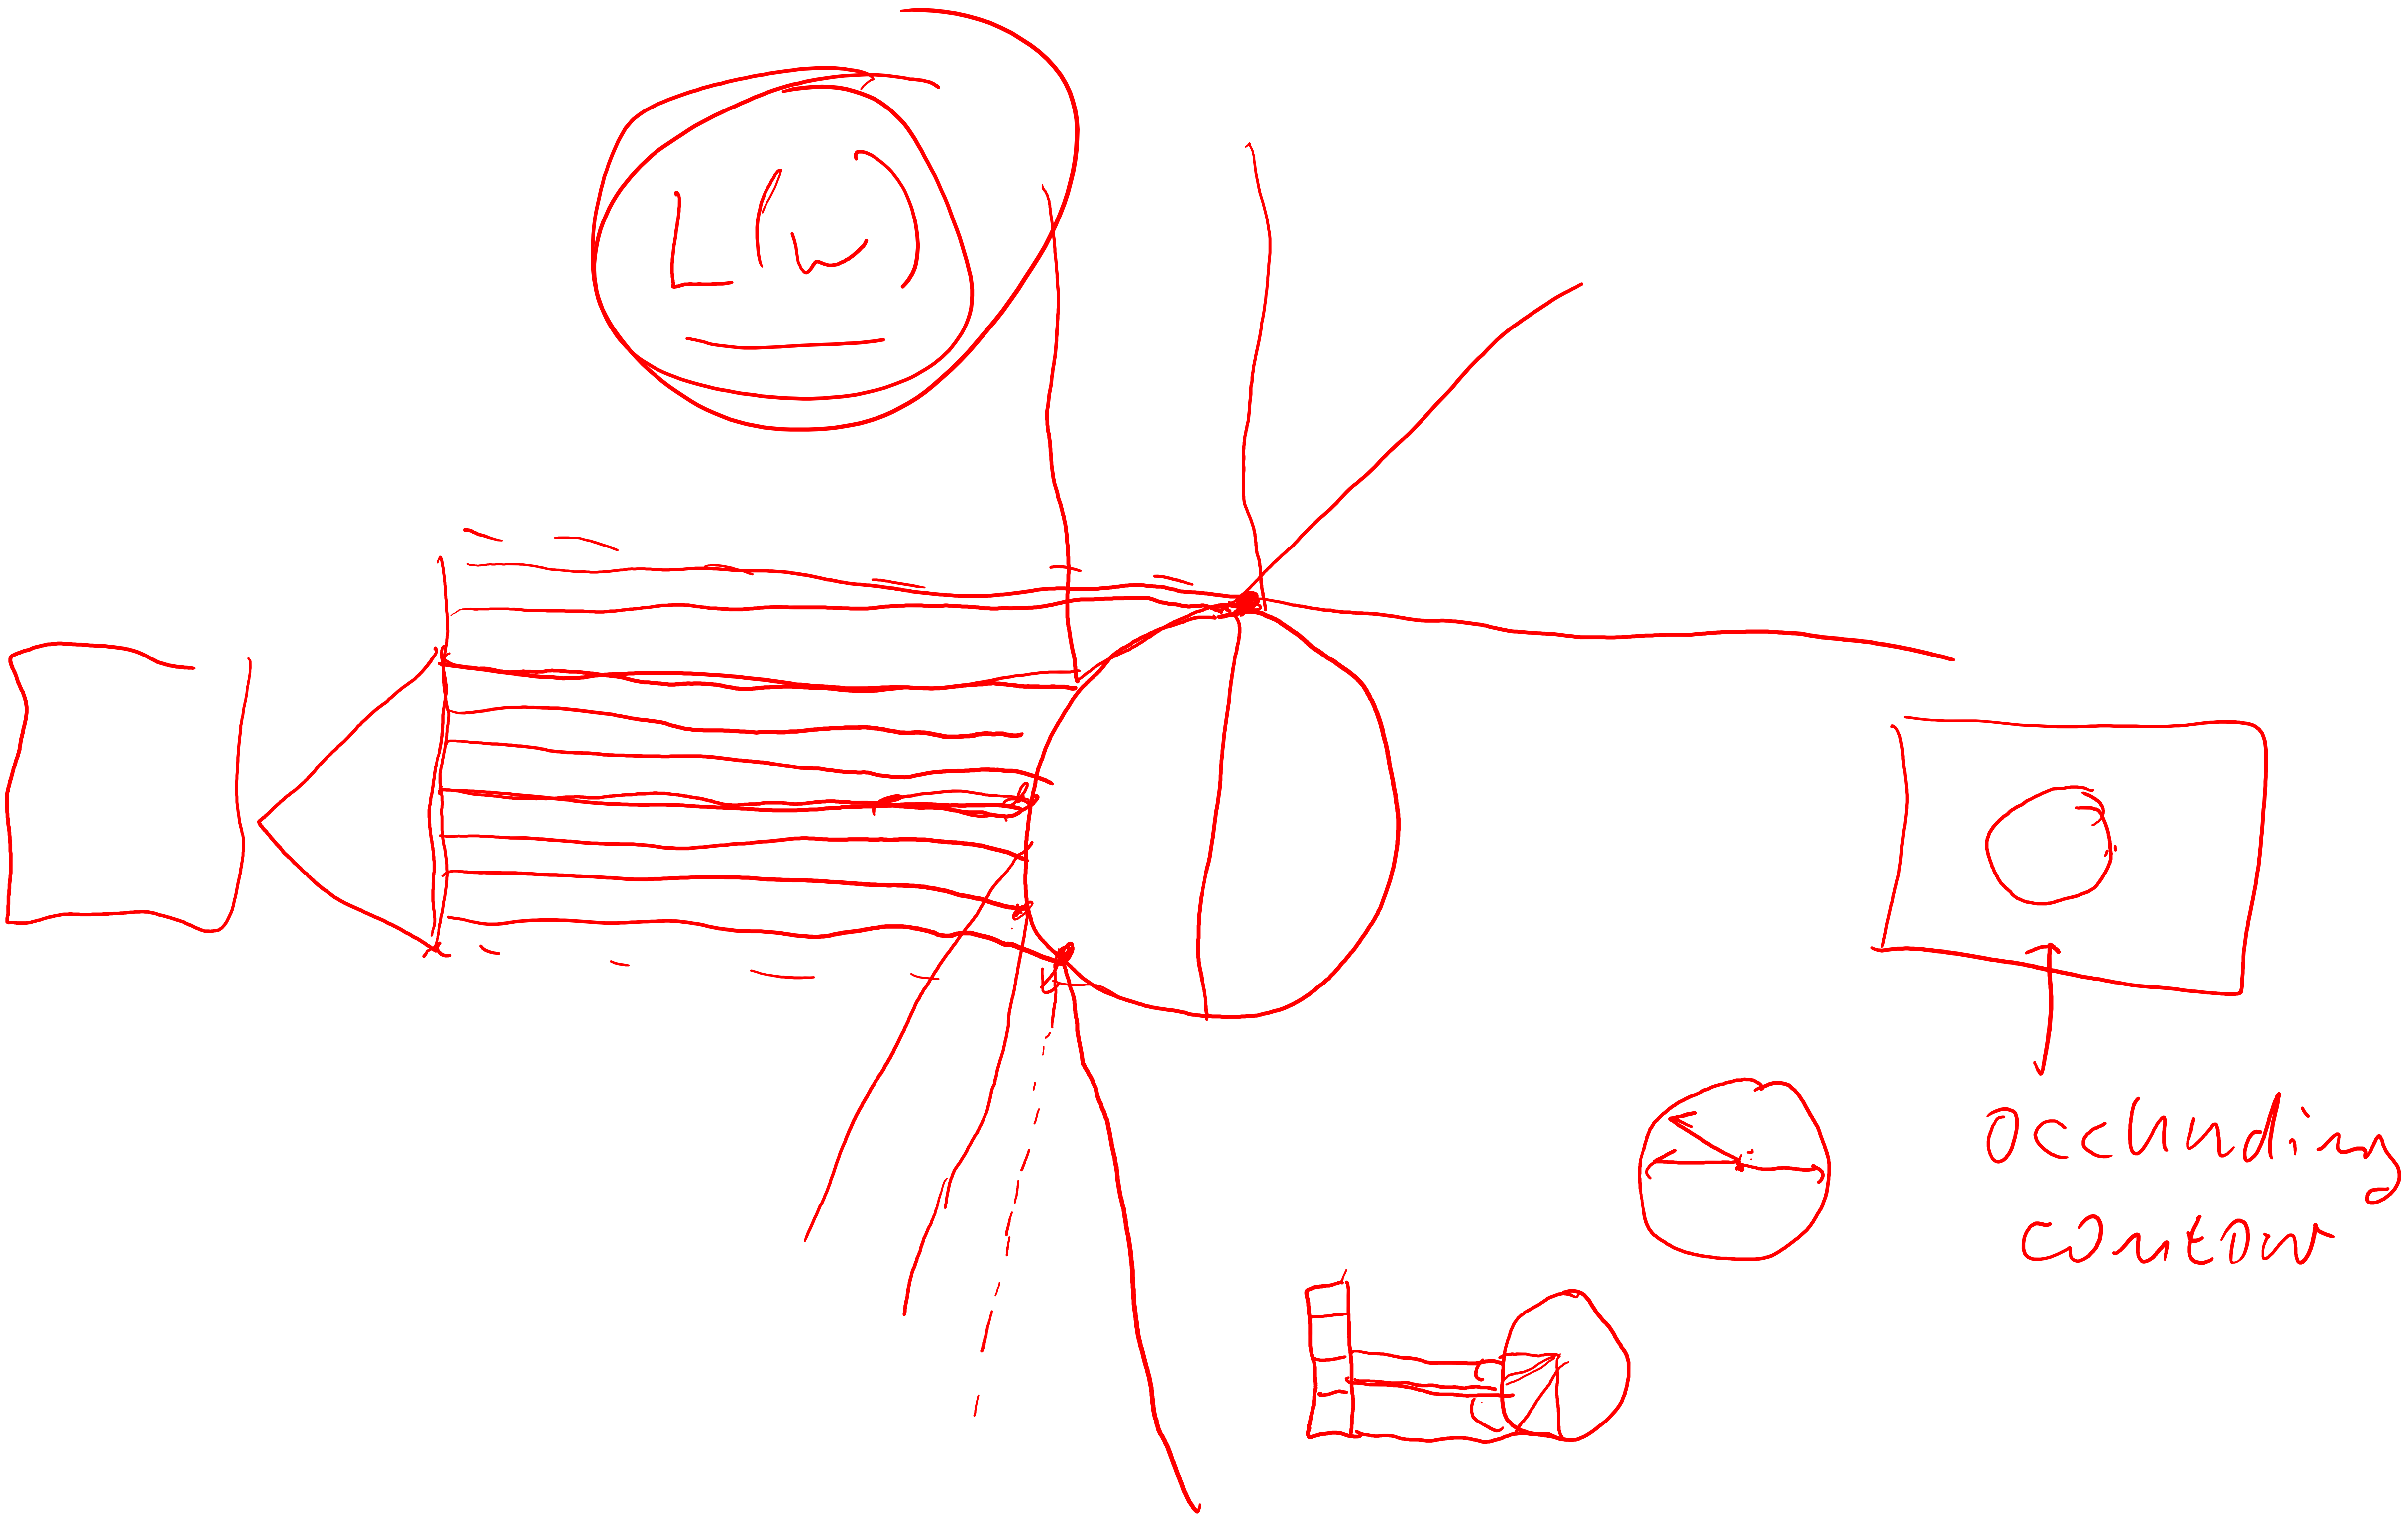
Lw
occluding
contour

Application: augmented reality

(a) Background photograph

(b) Camera calibration grid and light probe

(g) Final result with differential rendering

[Debevec, 1998]

Application: augmented reality

http://gl.ict.usc.edu/LightStages/

Application: augmented reality

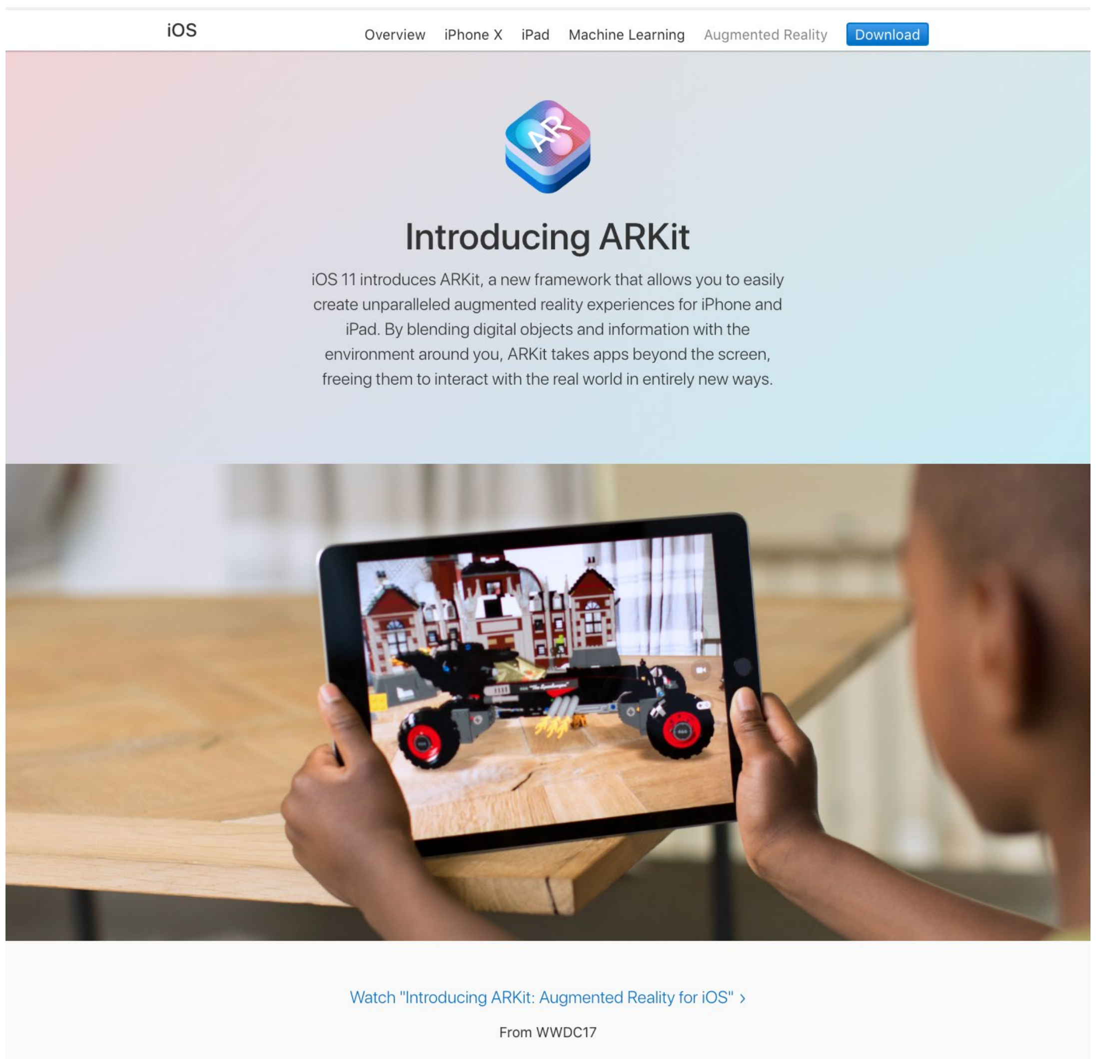

[https://developer.apple.com/arkit/]

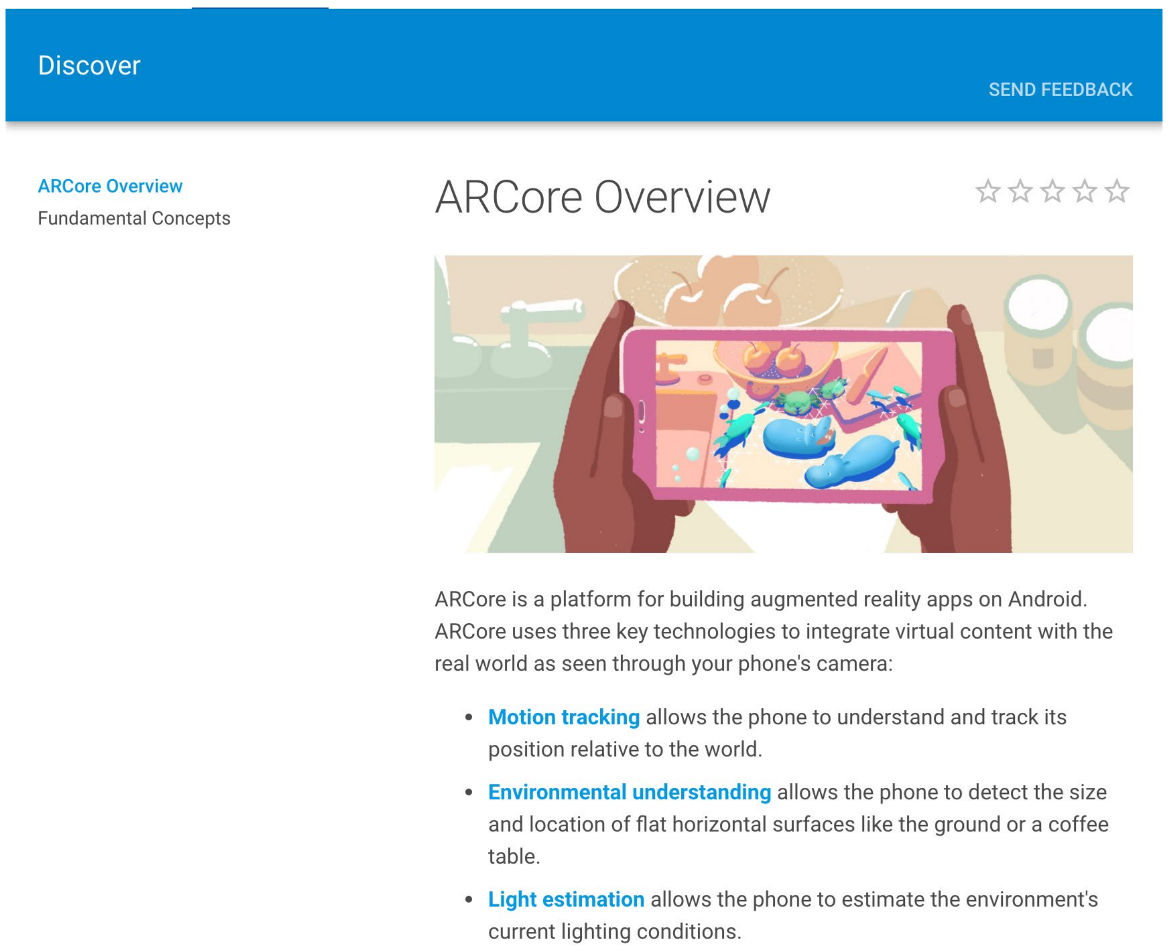

[https://developers.google.com/ar/]

Lighting models: far-field approximation

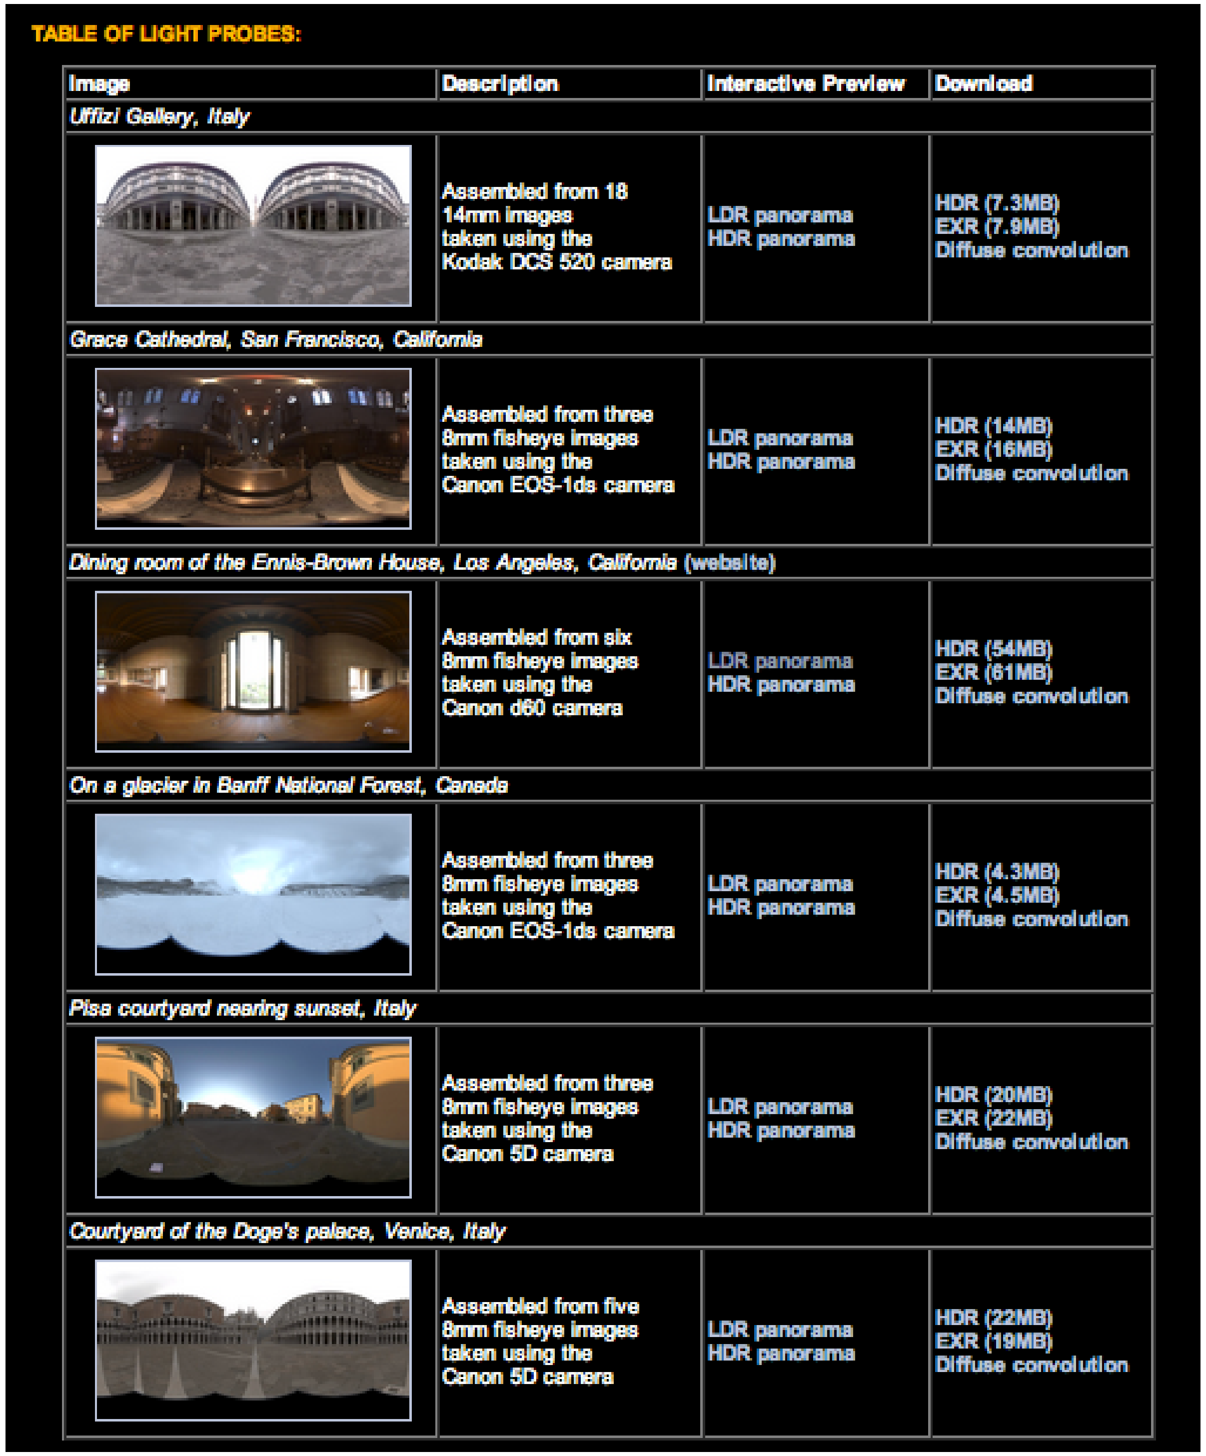

TABLE OF LIGHT PROBES:

Image	Description	Interactive Preview	Download
Uffizi Gallery, Italy			
	Assembled from 18 14mm images taken using the Kodak DCS 520 camera	LDR panorama HDR panorama	HDR (7.3MB) EXR (7.9MB) Diffuse convolution
Grace Cathedral, San Francisco, California			
	Assembled from three 8mm fisheye images taken using the Canon EOS-1ds camera	LDR panorama HDR panorama	HDR (14MB) EXR (16MB) Diffuse convolution
Dining room of the Ennis-Brown House, Los Angeles, California (website)			
	Assembled from six 8mm fisheye images taken using the Canon d60 camera	LDR panorama HDR panorama	HDR (54MB) EXR (61MB) Diffuse convolution
On a glacier in Banff National Forest, Canada			
	Assembled from three 8mm fisheye images taken using the Canon EOS-1ds camera	LDR panorama HDR panorama	HDR (4.3MB) EXR (4.5MB) Diffuse convolution
Pisa courtyard nearing sunset, Italy			
	Assembled from three 8mm fisheye images taken using the Canon 5D camera	LDR panorama HDR panorama	HDR (20MB) EXR (22MB) Diffuse convolution
Courtyard of the Doge's palace, Venice, Italy			
	Assembled from five 8mm fisheye images taken using the Canon 5D camera	LDR panorama HDR panorama	HDR (22MB) EXR (19MB) Diffuse convolution

- One can download far-field lighting environments that have been captured by others

  [http://gl.ict.usc.edu/Data/HighResProbes/]

- A number of apps and software exist to help you capture capture your own environments using a light probe

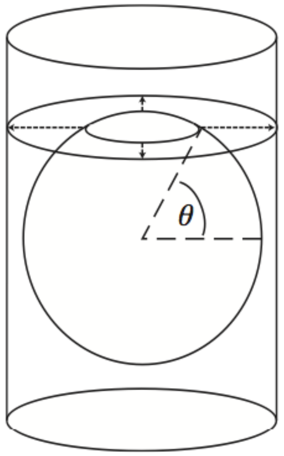

Figure 6. To produce the equal-area cylindrical projection of a spherical map, one projects each point on the surface of the sphere horizontally outward onto the cylinder, and then unwraps the cylinder to obtain a rectangular "panoramic" map.

Application: inferring outdoor illumination

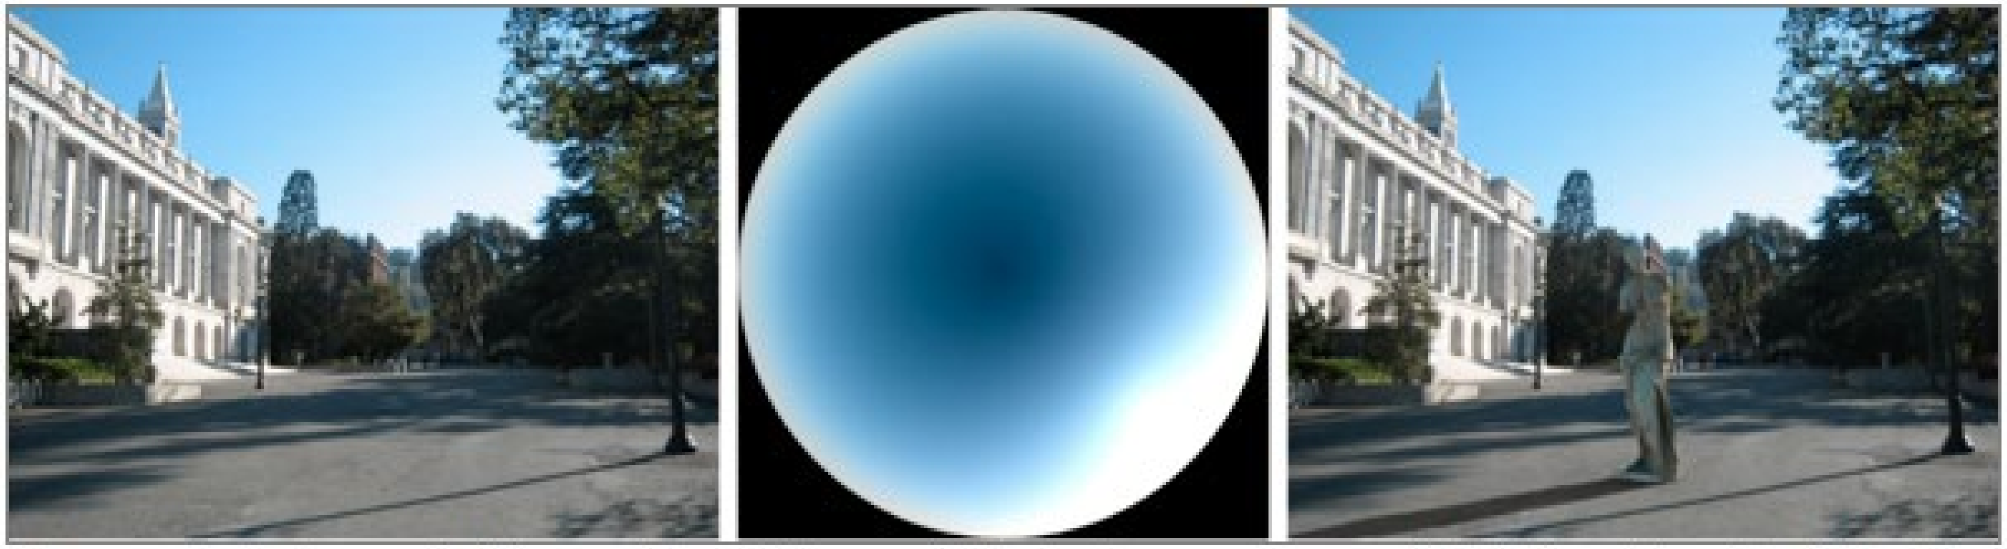

From a single image (left), we estimate the most likely sky appearance (middle) and insert a 3-D object (right). Illumination estimation was done entirely automatically.

[Lalonde et al., 2009]

A further simplification: Low-frequency illumination

$$L(\omega) = \sum_i a_i Y_i(\omega)$$

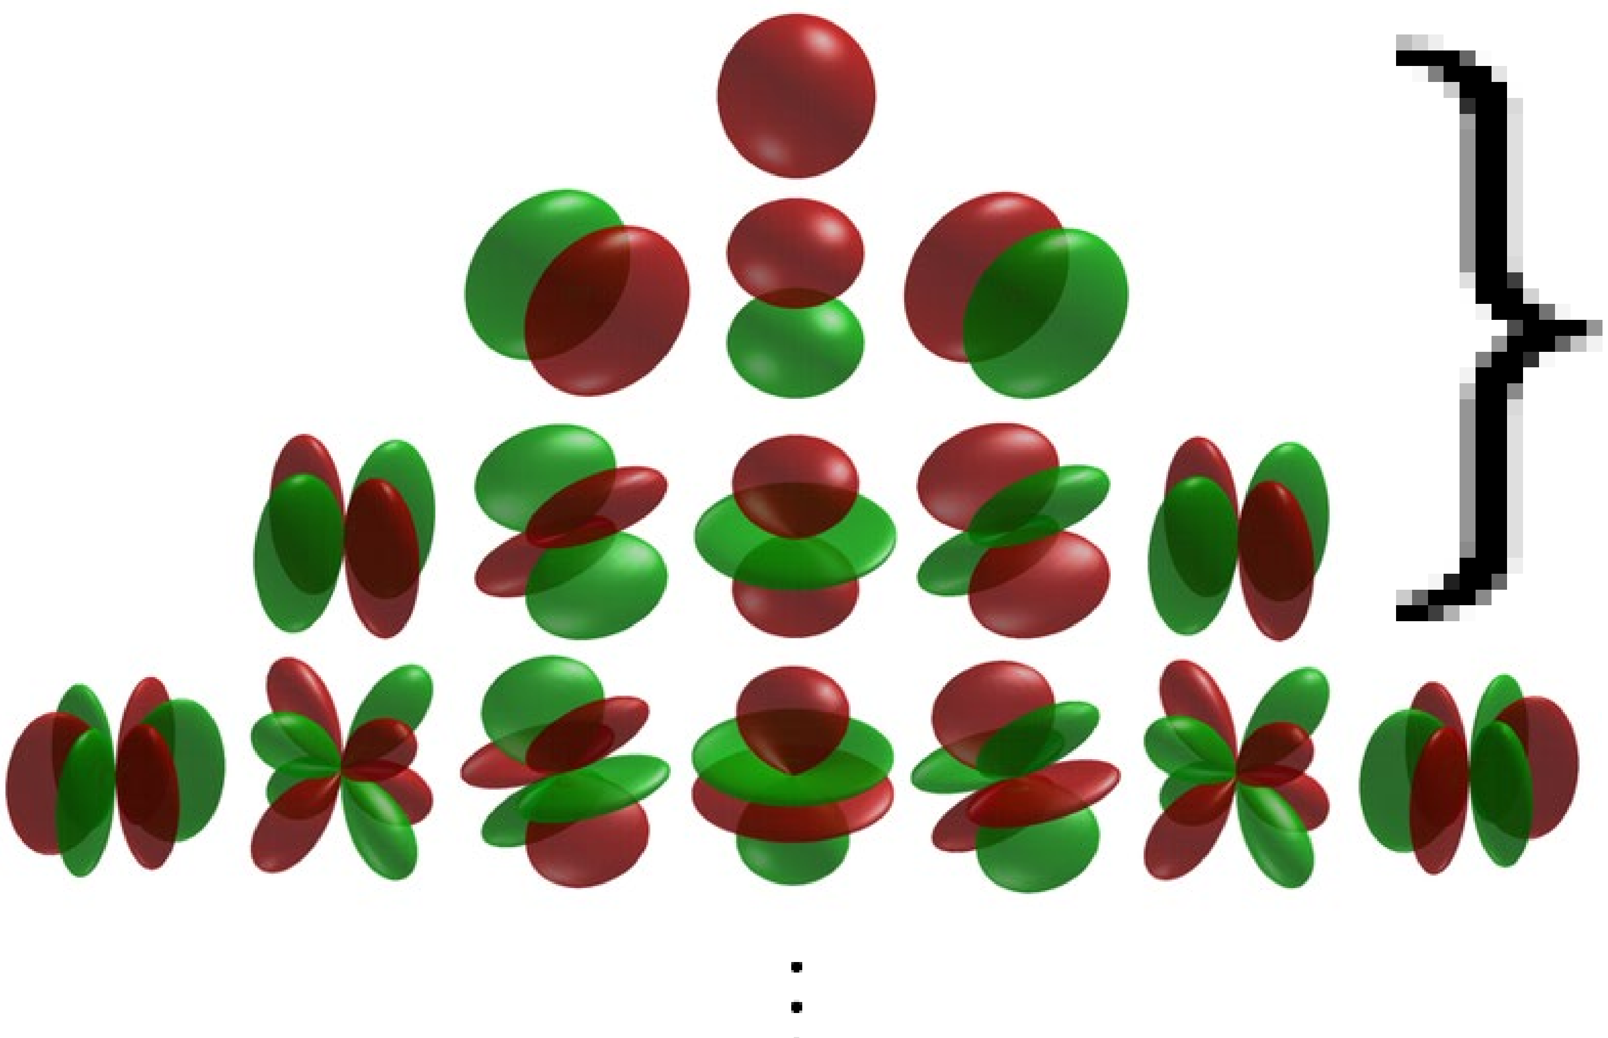

First nine basis functions are sufficient for re-creating Lambertian appearance

[Ramamoorthi and Hanrahan, 2001; Basri and Jacobs, 2003]

Low-frequency illumination

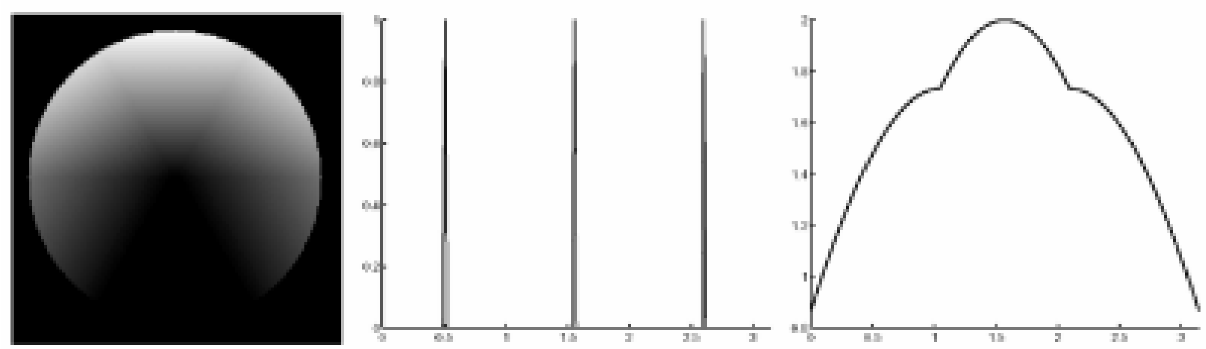

Fig. 2. On the left, a white sphere illuminated by three directional (distant point) sources of light. All the lights are parallel to the image plane, one source illuminates the sphere from above and the two others illuminate the sphere from diagonal directions. In the middle, a cross-section of the lighting function with three peaks corresponding to the three light sources. On the right, a cross-section indicating how the sphere reflects light. We will make precise the intuition that the material acts as a low-pass filtering, smoothing the light as it reflects it.

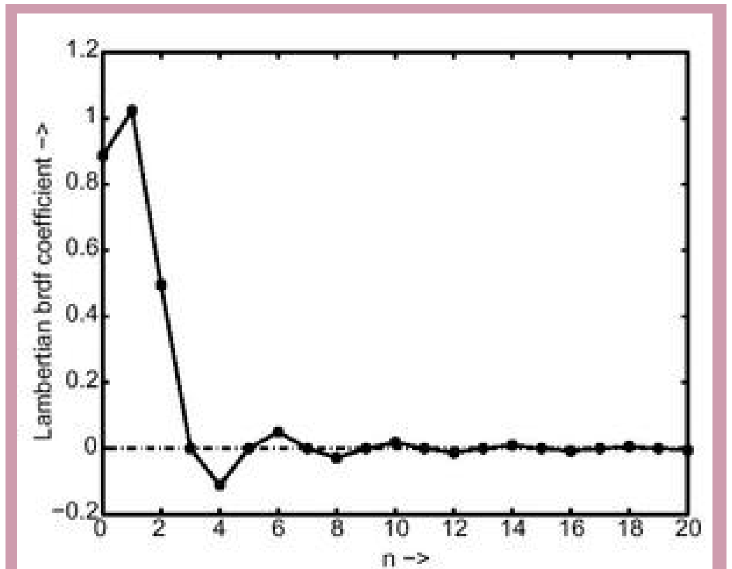

Figure 3. Plot of spherical harmonic terms in Lambertian BRDF filter.

[Ramamoorthi and Hanrahan, 2001; Basri and Jacobs, 2003]

Low-frequency illumination

$$L(\omega) = \sum_i a_i Y_i(\omega)$$

$\Downarrow$ Truncate to first 9 terms

$$\vec{\ell} = (\ell_1, \ldots, \ell_9)$$

[Ramamoorthi and Hanrahan, 2001; Basri and Jacobs, 2003]

Application: Trivial rendering

Capture light probe

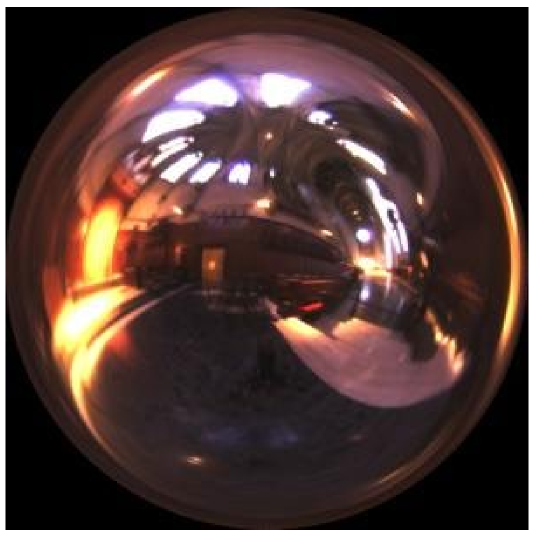

Low-pass filter (truncate to first nine SHs)

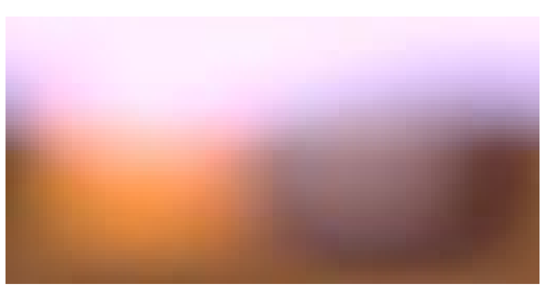

Rendering a (convex) diffuse object in this environment simply requires a lookup based on the surface normal at each pixel

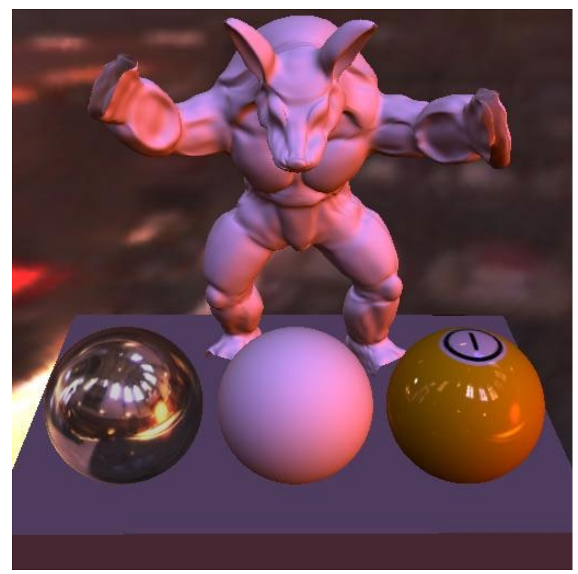

White-out: Snow and Overcast Skies

CAN'T perceive the shape of the snow covered terrain!

CAN perceive shape in regions
lit by the street lamp!!

WHY?

Diffuse Reflection from Uniform Sky

$$L^{surface}(\theta_r,\phi_r) \quad = \quad \int_{-\pi}^{\pi}\int_{0}^{\pi/2} L^{src}(\theta_i,\phi_i)\, f(\theta_i,\phi_i;\theta_r,\phi_r)\cos\theta_i\ \sin\theta_i\, d\theta_i\, d\phi_i$$

- Assume Lambertian Surface with Albedo = 1 (no absorption)

$$f(\theta_i,\phi_i;\theta_r,\phi_r) = \frac{1}{\pi}$$

- Assume Sky radiance is constant

$$L^{src}(\theta_i,\phi_i) = L^{sky}$$

- Substituting in above Equation:

$$\boxed{L^{surface}(\theta_r,\phi_r) \quad = \quad L^{sky}}$$

Radiance of any patch is the same as Sky radiance !! (white-out condition)

Even simpler: Directional lighting

- Assume that, over the observed region of interest, all source of incoming flux is from one direction

$$L(x, \omega, t, \lambda) \longrightarrow L(x, t, \lambda) \longrightarrow s(t, \lambda)\delta(\omega = \omega_o(t))$$

$$L(x, \omega) \longrightarrow L(\omega) \longrightarrow \underline{s}\delta(\omega = \underline{\omega_o})$$

- Convenient representation

$$\vec{\boldsymbol{\ell}} = (\ell_x, \ell_y, \ell_z)$$

"light direction" $\underline{\hat{\boldsymbol{\ell}}} = \dfrac{\vec{\boldsymbol{\ell}}}{||\vec{\boldsymbol{\ell}}||}$ $\equiv \omega_o$

"light strength" $||\vec{\boldsymbol{\ell}}||$ $\equiv s$

Simple shading

ASSUMPTION 1:
LAMBERTIAN

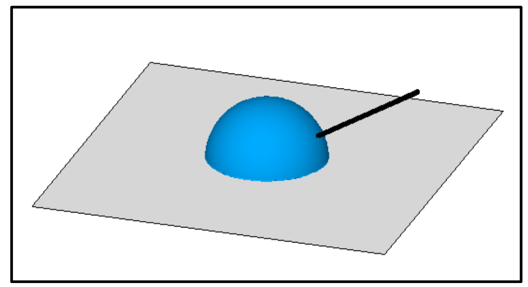

ASSUMPTION 2:
DIRECTIONAL LIGHTING

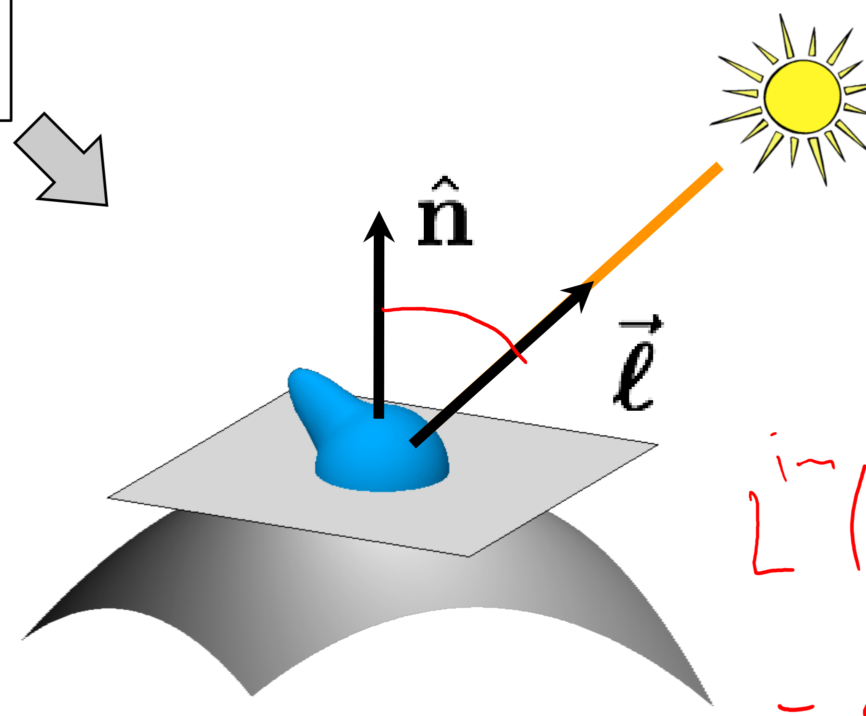

$$L^{\text{out}}(\hat{\omega}) = \int_{\Omega_{\text{in}}} \underbrace{f(\hat{\omega}_{\text{in}}, \hat{\omega}_{\text{out}})}\,\underbrace{L^{\text{in}}(\hat{\omega}_{\text{in}})}\,\underbrace{\cos\theta_{\text{in}}}\, d\hat{\omega}_{\text{in}}$$

$$\boxed{I = a\hat{\mathbf{n}}^{\top}\vec{\boldsymbol{\ell}}}$$

$$L^{in}(\omega_{in})$$

$$= \underbrace{s\,\delta(\hat{l} - \omega_{in})}$$

“N-dot-l” shading

ASSUMPTION 1:
LAMBERTIAN

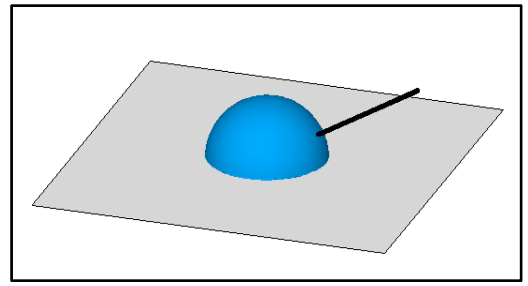

ASSUMPTION 2:
DIRECTIONAL LIGHTING

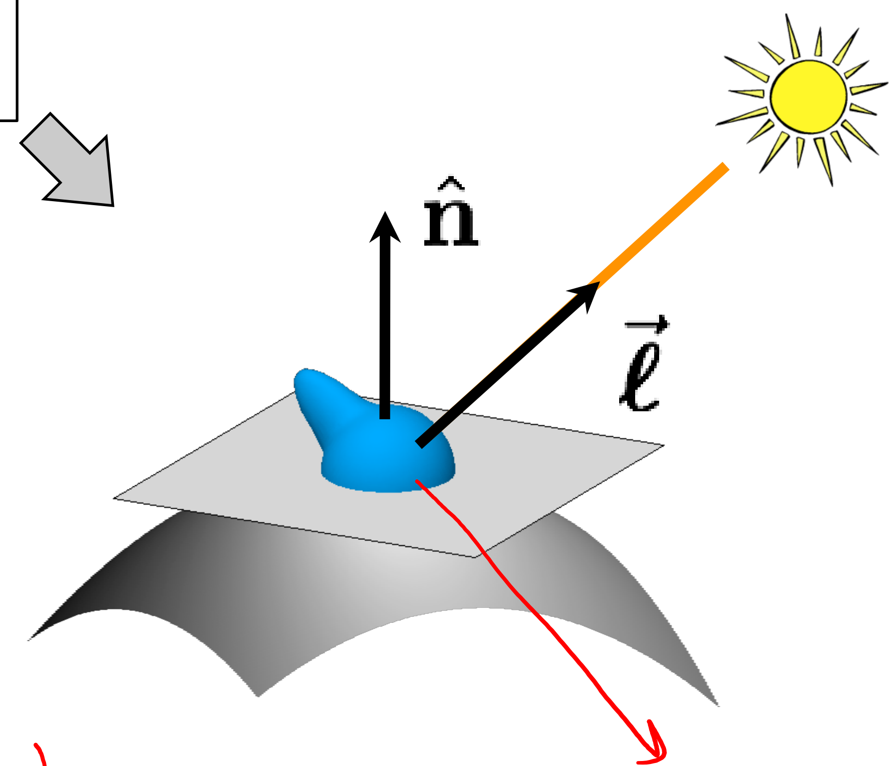

$$L^{\text{out}}(\hat{\boldsymbol{\omega}}) = \int_{\Omega_{\text{in}}} f(\hat{\boldsymbol{\omega}}_{\text{in}}, \hat{\boldsymbol{\omega}}_{\text{out}}) L^{\text{in}}(\hat{\boldsymbol{\omega}}_{\text{in}}) \cos\theta_{\text{in}} d\hat{\boldsymbol{\omega}}_{\text{in}}$$

$$\boxed{I = a\underbrace{\hat{\mathbf{n}}^\top \vec{\boldsymbol{\ell}}}}$$

$a \cdot \max(n^\top l, 0)$

An ideal point light source

$$L(\boldsymbol{x}, \boldsymbol{\omega}) = \frac{s}{||\boldsymbol{x} - \boldsymbol{x}_o||^2} \; \delta\left(\boldsymbol{\omega} = \frac{\boldsymbol{x} - \boldsymbol{x}_o}{||\boldsymbol{x} - \boldsymbol{x}_o||}\right)$$

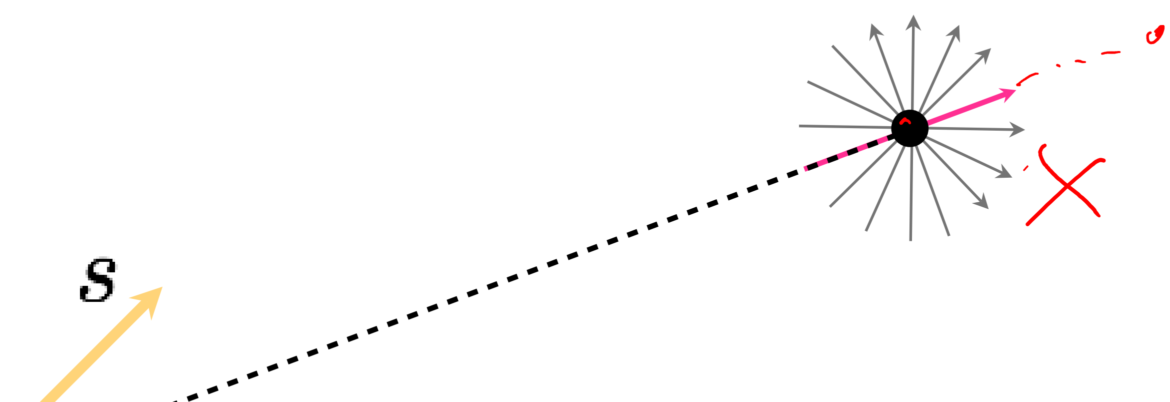

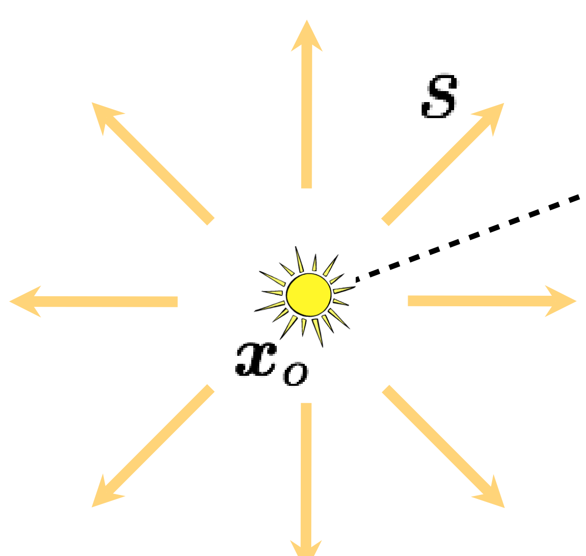

Think of this as a spatially-varying directional source where

1. the direction is away from x_o
2. the strength is proportional to 1/(distance)^2

Summary of some useful lighting models

- plenoptic function (function on 5D domain)
- far-field illumination (function on 2D domain)
- low-frequency far-field illumination (nine numbers)
- directional lighting (three numbers = direction and strength)
- point source (four numbers = location and strength)

Some notes about radiometry

What about color?

Spectral radiance

- Distribution of radiance as a function of wavelength.
- All of the phenomena we described in this lecture can be extended to take into account color, by considering separate radiance and BRDF functions independently for each wavelength.

retinal color

$$\mathbf{c}(\ell(\lambda)) = (c_s, c_m, c_l)$$

$$c_s = \int k_s(\lambda)\ell(\lambda)d\lambda$$

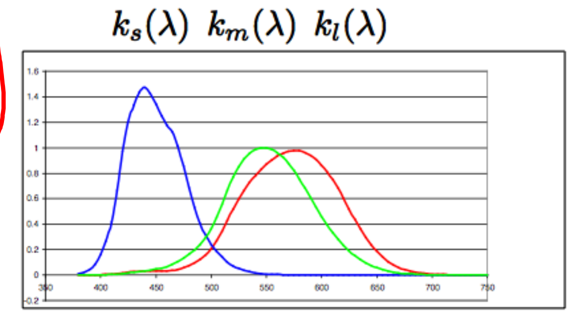

LMS senstivity functions

$$L^o(\omega_o, \lambda_o) = \int\int f(\omega_o, \omega_i, \lambda_o, \lambda_i) L^i(\omega_i, \lambda_i) \langle \omega_i, \hat{n} \rangle d\omega_i d\lambda_i$$

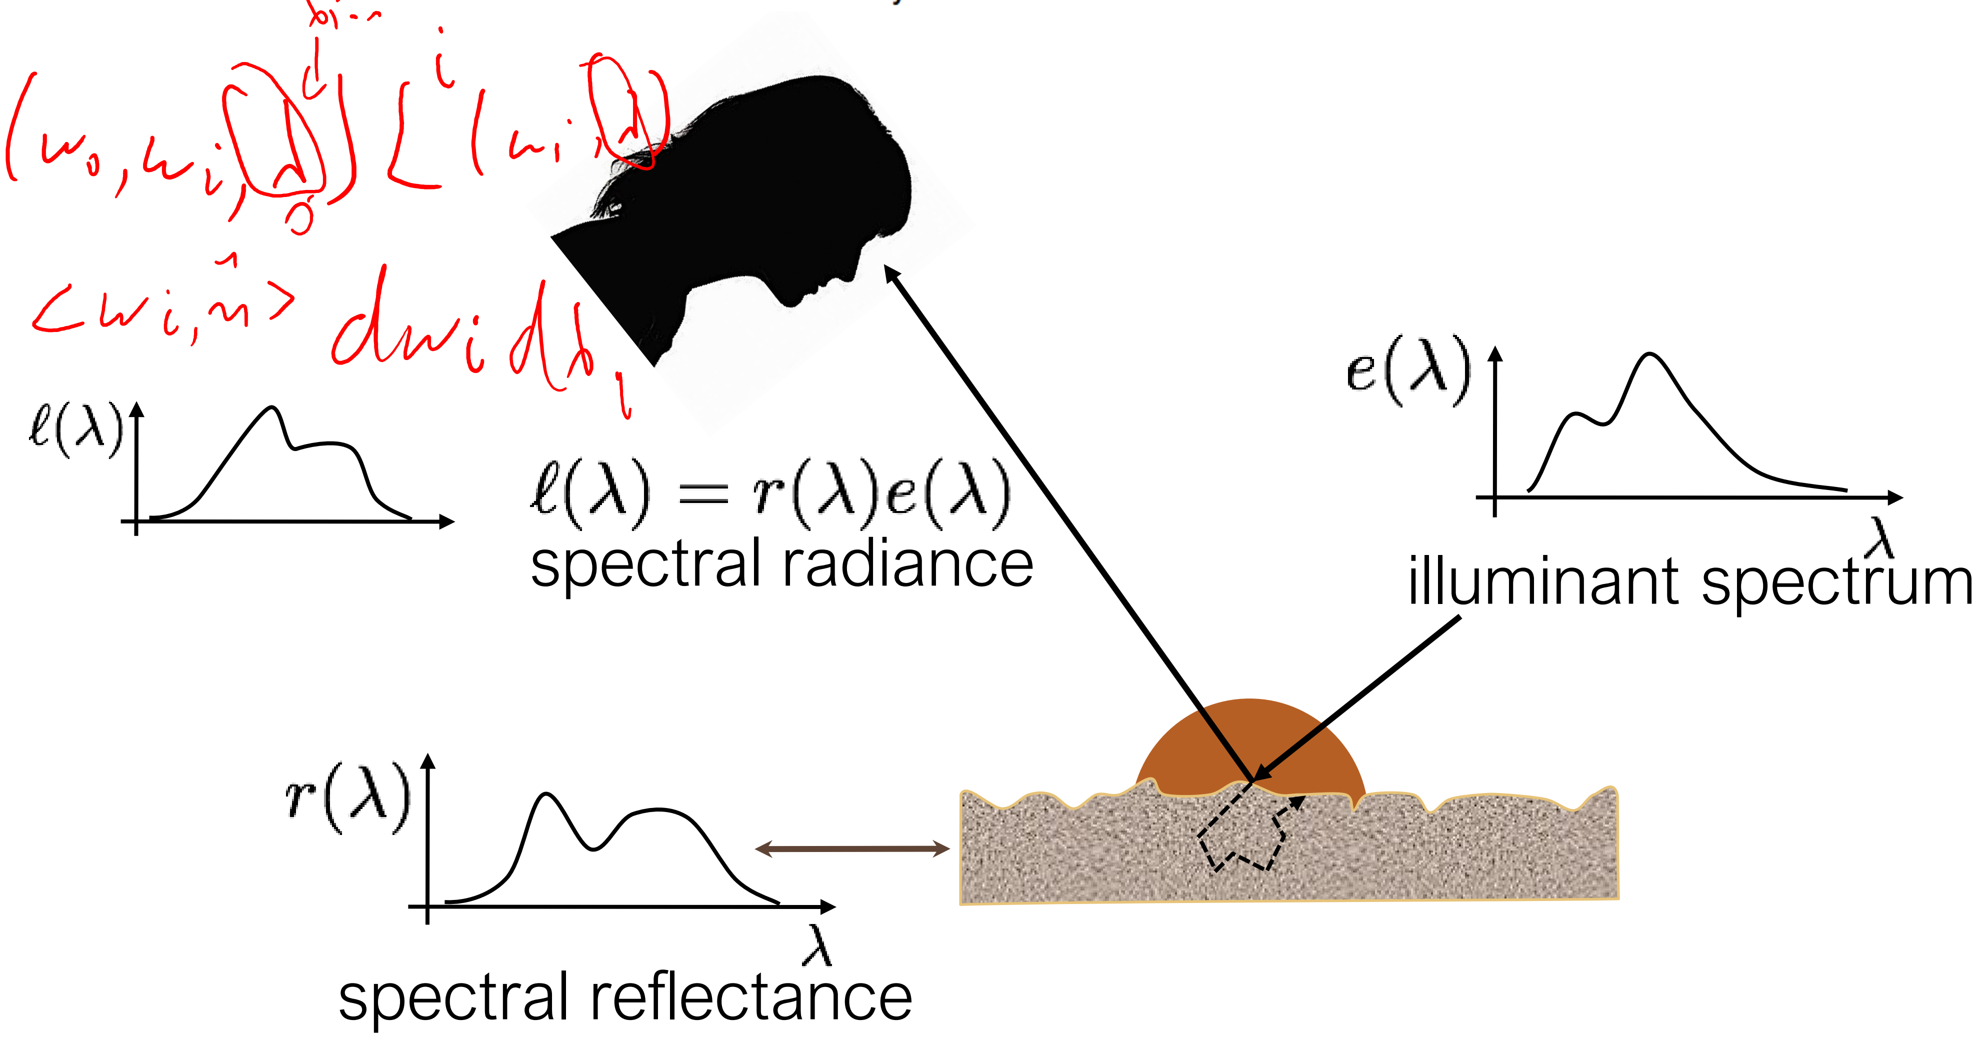

Spectral radiance

- Distribution of radiance as a function of wavelength.
- All of the phenomena we described in this lecture can be extended to take into account color, by considering separate radiance and BRDF functions independently for each wavelength.

Does this view of color ignore any important phenomena?

Spectral radiance

- Distribution of radiance as a function of wavelength.
- All of the phenomena we described in this lecture can be extended to take into account color, by considering separate radiance and BRDF functions independently for each wavelength.

Does this view of color ignore any important phenomena?

- Things like fluorescence and any other phenomena where light changes color.

Spectral Sensitivity Function (SSF)

- *Any* light sensor (digital or not) has different sensitivity to different wavelengths.
- This is described by the sensor's *spectral sensitivity function* .
- When measuring light of a some SPD $\Phi(\lambda)$, the sensor produces a *scalar* $f(\lambda)$ response:

$$\text{sensor response} \longrightarrow R = \int_\lambda \underbrace{\Phi(\lambda)}_{\text{light SPD}} \underbrace{f(\lambda)}_{\text{sensor SSF}} d\lambda$$

Weighted combination of light's SPD: light contributes more at wavelengths where the sensor has higher sensitivity.

The spectral sensitivity function converts radiometric units (radiance, irradiance) defined per wavelength, to radiometric quantities where wavelength has been averaged out.

Radiometry versus photometry

- **All radiometric quantities have equivalents in photometry**
- **Photometry: accounts for response of human visual system to electromagnetic radiation**
- **Luminance (Y) is photometric quantity that corresponds to radiance: integrate radiance over all wavelengths, weight by eye's luminous efficacy curve, e.g.:**

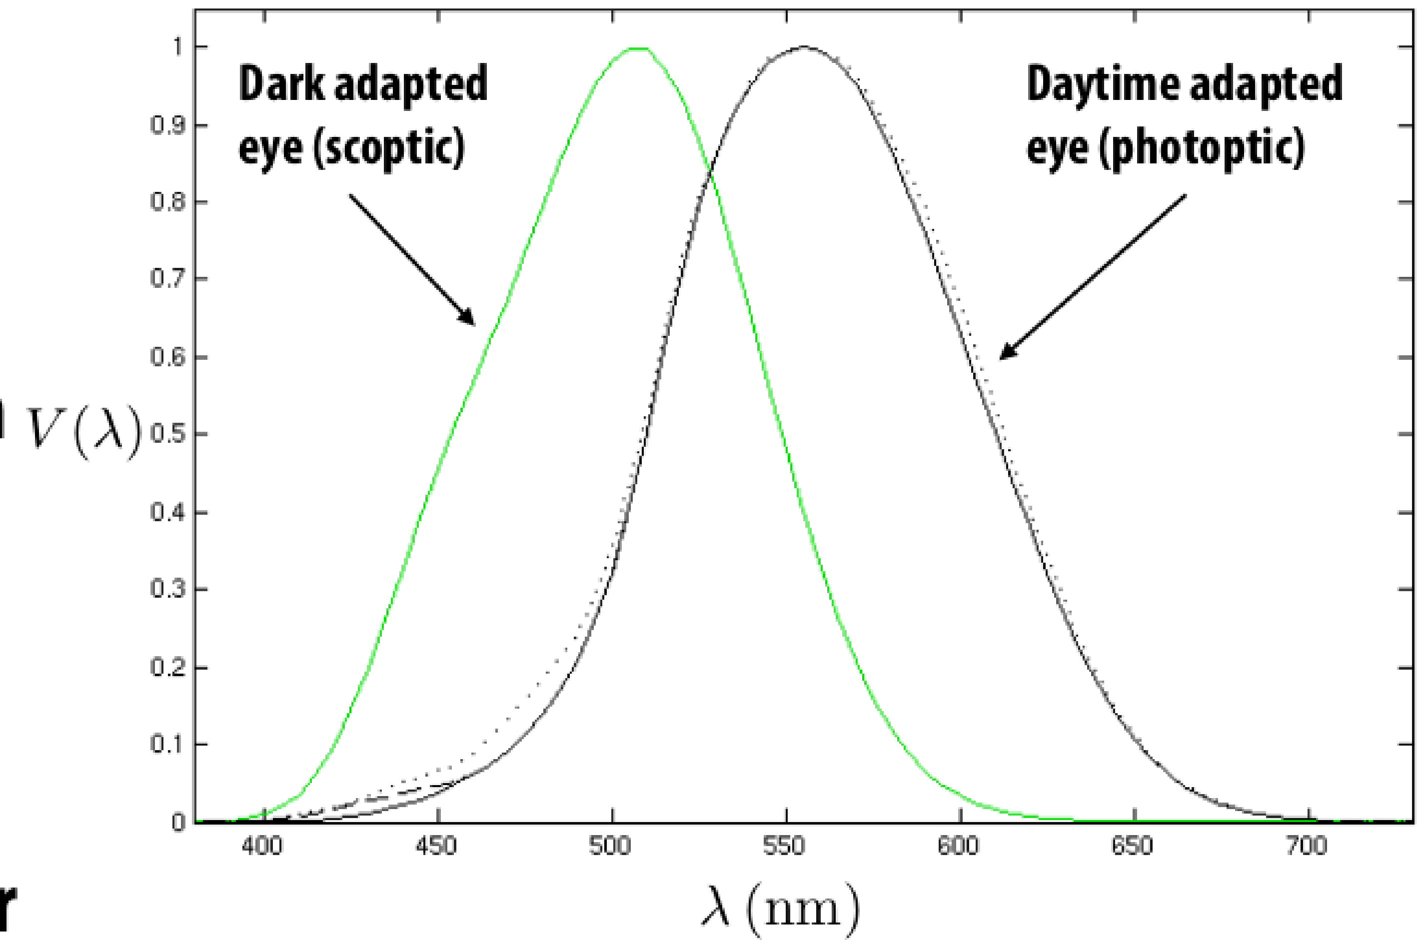

$$Y(\mathrm{p}, \omega) = \int_0^\infty L(\mathrm{p}, \omega, \lambda)\, V(\lambda)\, \mathrm{d}\lambda$$

Radiometry versus photometry

Physics	Radiometry	Photometry
Energy	Radiant Energy	Luminous Energy
Flux (Power)	Radiant Power	Luminous Power
Flux Density	Irradiance (incoming) Radiosity (outgoing)	Illuminance (incoming) Luminosity (outgoing)
Angular Flux Density	Radiance	Luminance
Intensity	Radiant Intensity	Luminous Intensity

Radiometry versus photometry

Photometry	MKS	CGS	British
Luminous Energy	Talbot	Talbot	Talbot
Luminous Power	Lumen	Lumen	Lumen
Illuminance Luminosity	Lux	Phot	Footcandle
Luminance	Nit, Apostlib, Blondel	Stilb Lambert	Footlambert
Luminous Intensity	Candela	Candela	Candela

Modern LED light

Input power: 11 W

Output: 815 lumens

(~ 80 lumens / Watt)

Incandescent bulbs:

~15 lumens / Watt)

Quiz 1: Measurement of a sensor using a thin lens

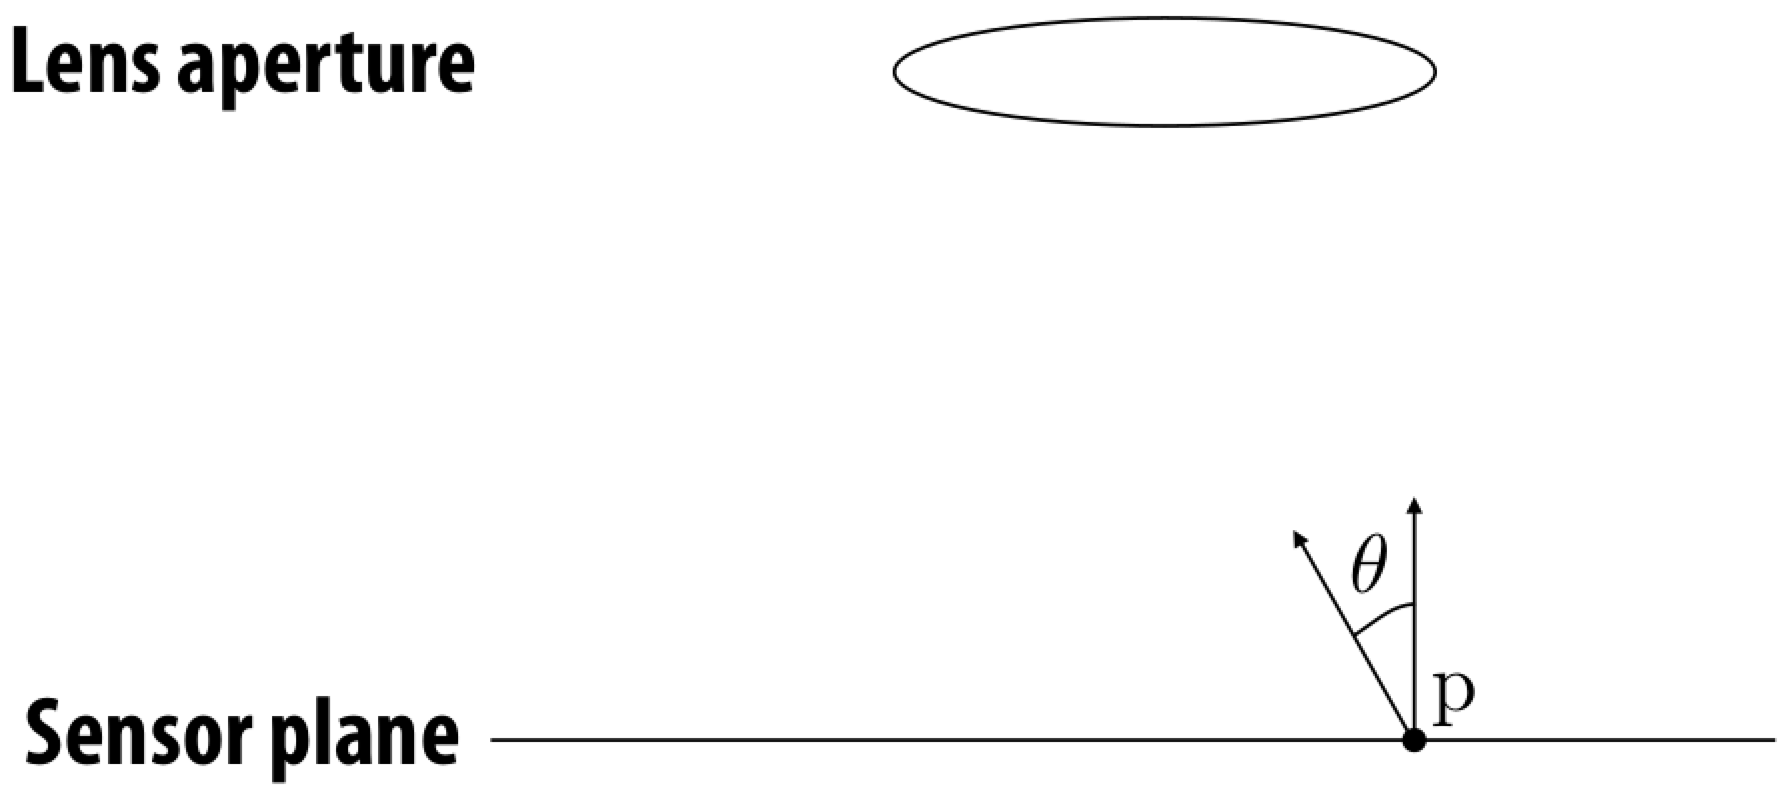

What integral should we write for the power measured by infinitesimal pixel p?

$$\int L(\omega, p)\, d\omega$$

Quiz 1: Measurement of a sensor using a thin lens

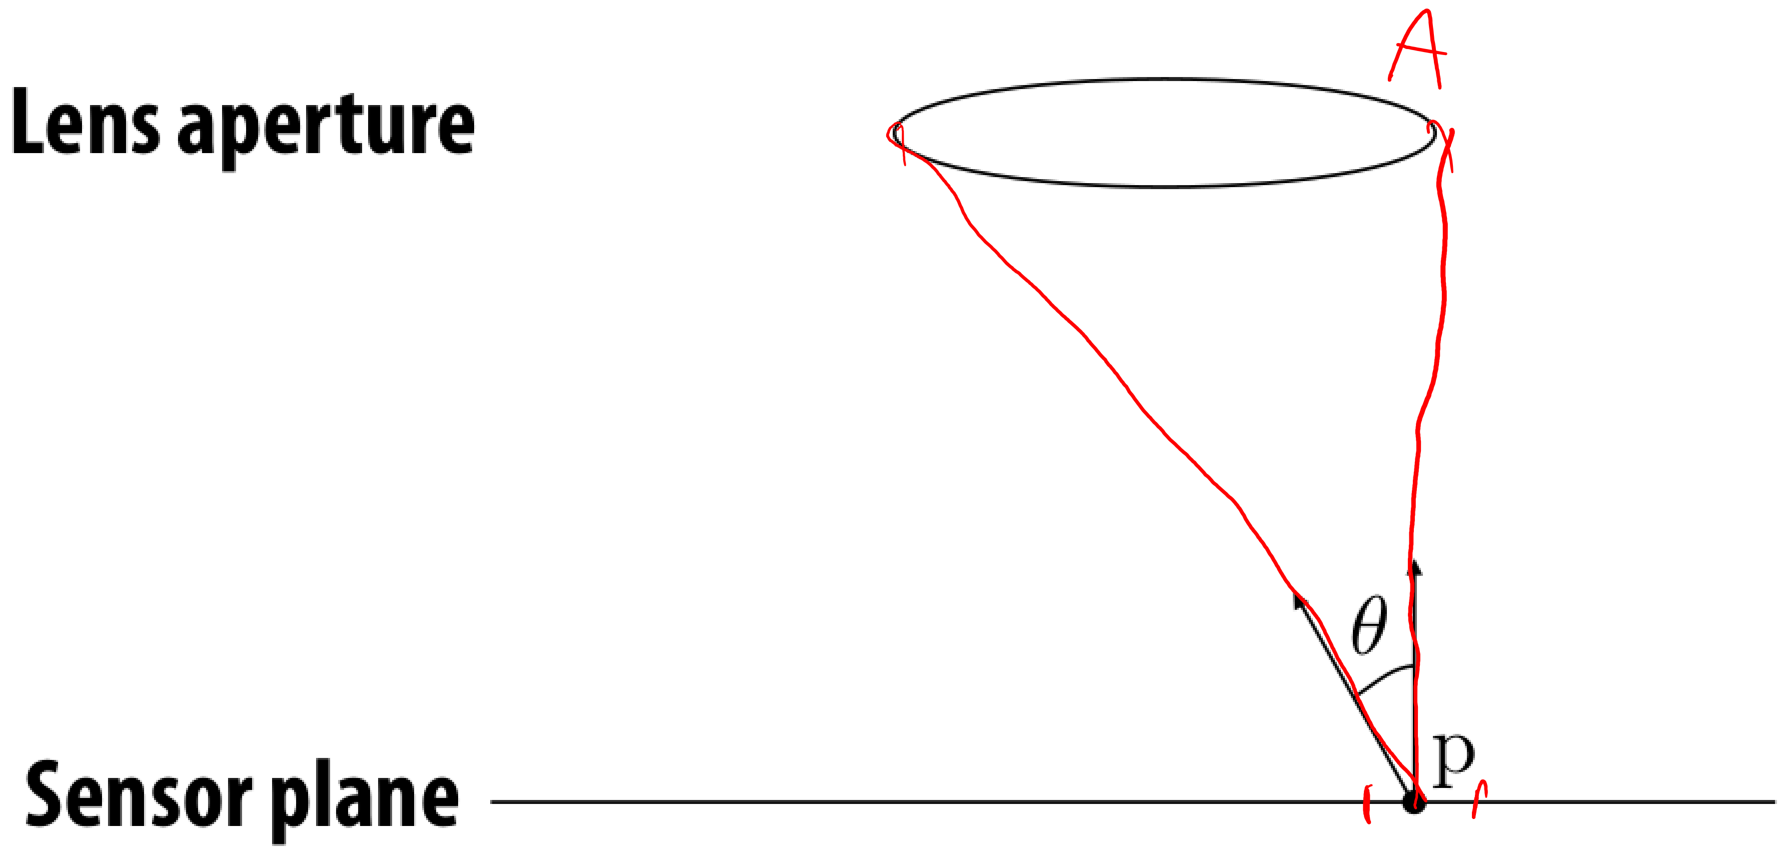

What integral should we write for the power measured by infinitesimal pixel p?

$$E(\mathrm{p}, t) = \int_{H^2} L_i(\mathrm{p}, \omega', t) \cos\theta \, \mathrm{d}\omega'$$

Can I transform this integral over the hemisphere to an integral over the aperture area?

Quiz 1: Measurement of a sensor using a thin lens

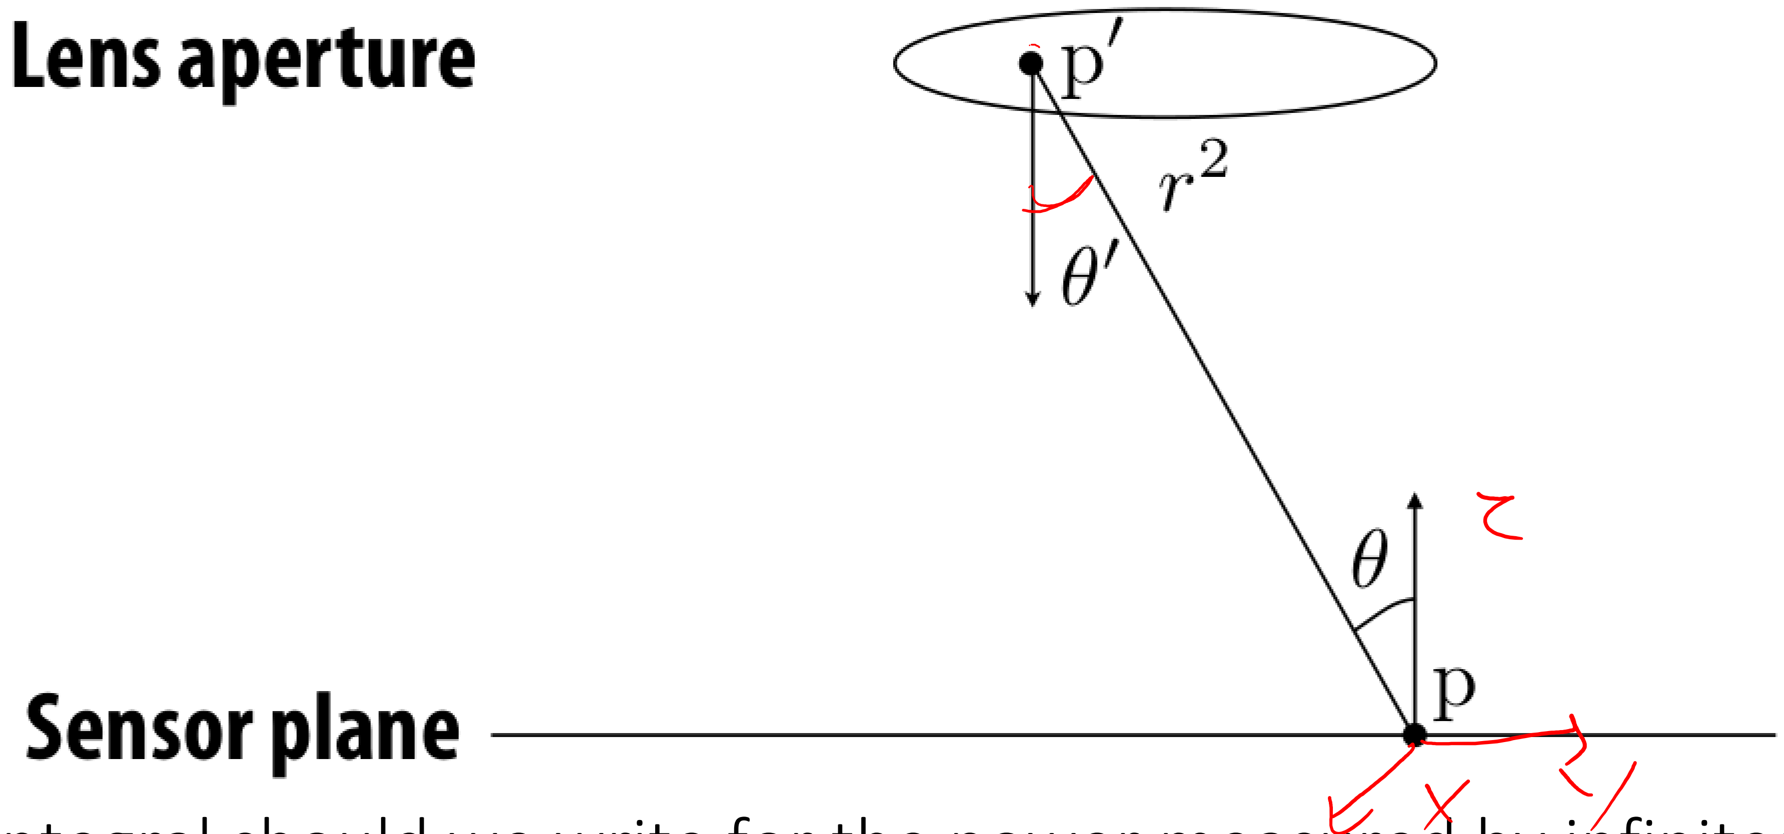

What integral should we write for the power measured by infinitesimal pixel p?

$$E(\mathrm{p}, t) = \int_{H^2} L_i(\mathrm{p}, \omega', t) \cos\theta \, \mathrm{d}\omega'$$

$\sin\theta \, d\theta \, d\varphi$

$w = \begin{pmatrix} \sin\theta \sin\varphi \\ \sin\theta \cos\varphi \\ \cos\theta \end{pmatrix}$

Can I transform this integral over the hemisphere to an integral over the aperture area?

$$E(\mathrm{p}, t) = \int_{A} L(\mathrm{p}' \to \mathrm{p}, t) \frac{\cos\theta \cos\theta'}{||\mathrm{p}' - \mathrm{p}||^2} \, \mathrm{d}A'$$

Transform integral over solid angle to integral over lens aperture

Quiz 1: Measurement of a sensor using a thin lens

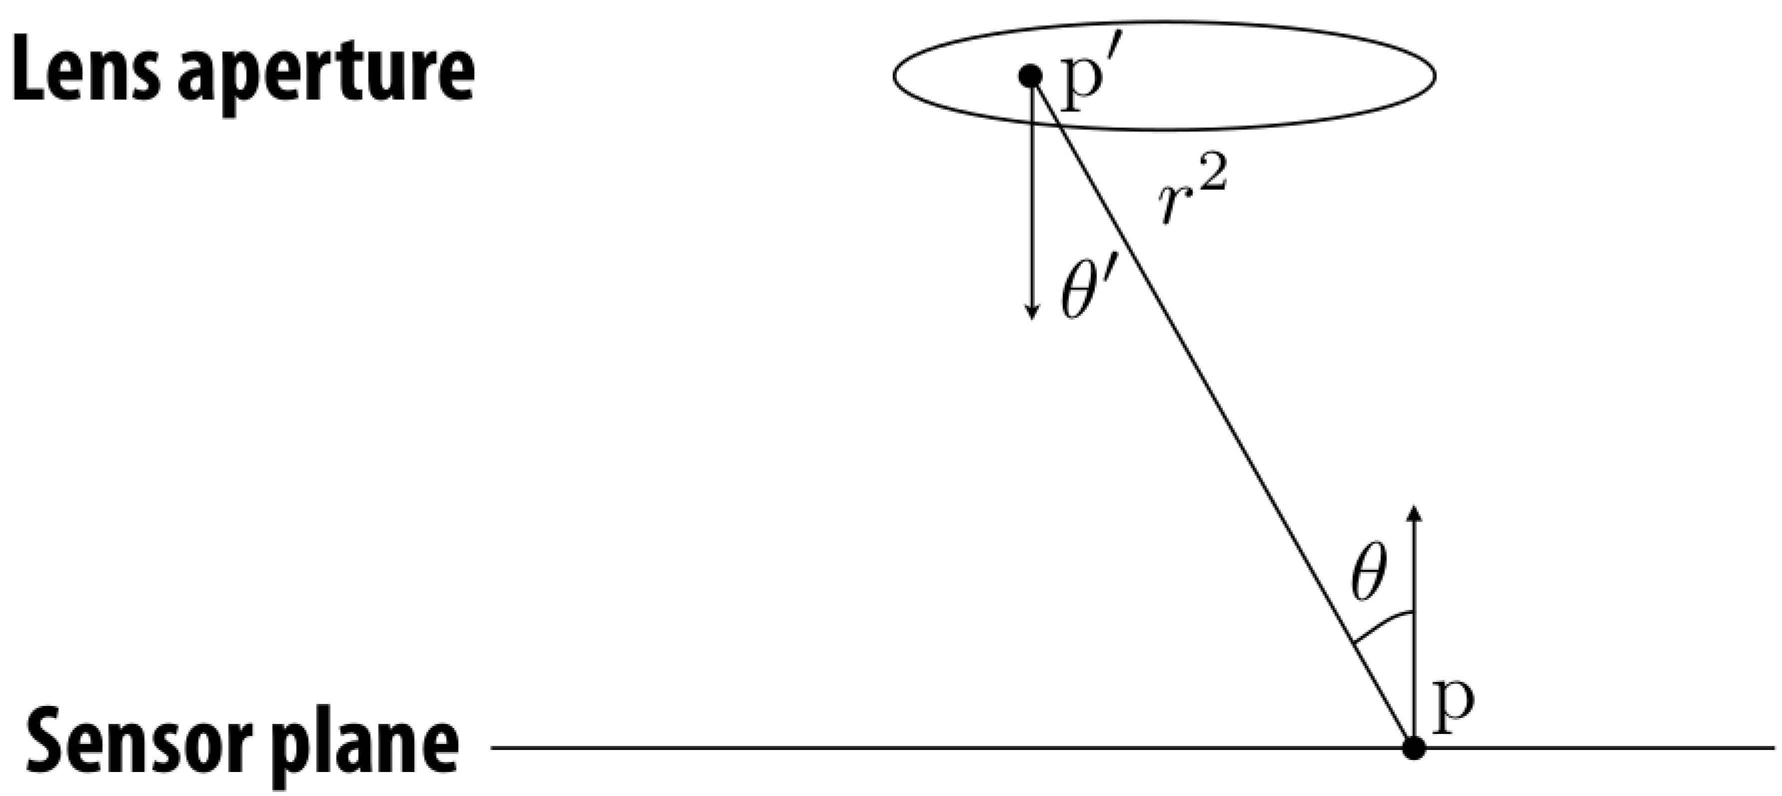

$$E(\mathrm{p}, t) = \int_A L(\mathrm{p}' \to \mathrm{p}, t) \frac{\cos\theta \cos\theta'}{||\mathrm{p}' - \mathrm{p}||^2} \, \mathrm{d}A'$$

Transform integral over solid angle to integral over lens aperture

$$= \int_A L(\mathrm{p}' \to \mathrm{p}, t) \frac{\cos^2\theta}{||\mathrm{p}' - \mathrm{p}||^2} \, \mathrm{d}A'$$

Assume aperture and film plane are parallel: $\theta = \theta'$

Can I write the denominator in a more convenient form?

Quiz 1: Measurement of a sensor using a thin lens

Lens aperture

$$||\mathrm{p}' - \mathrm{p}|| = \frac{d}{\cos\theta}$$

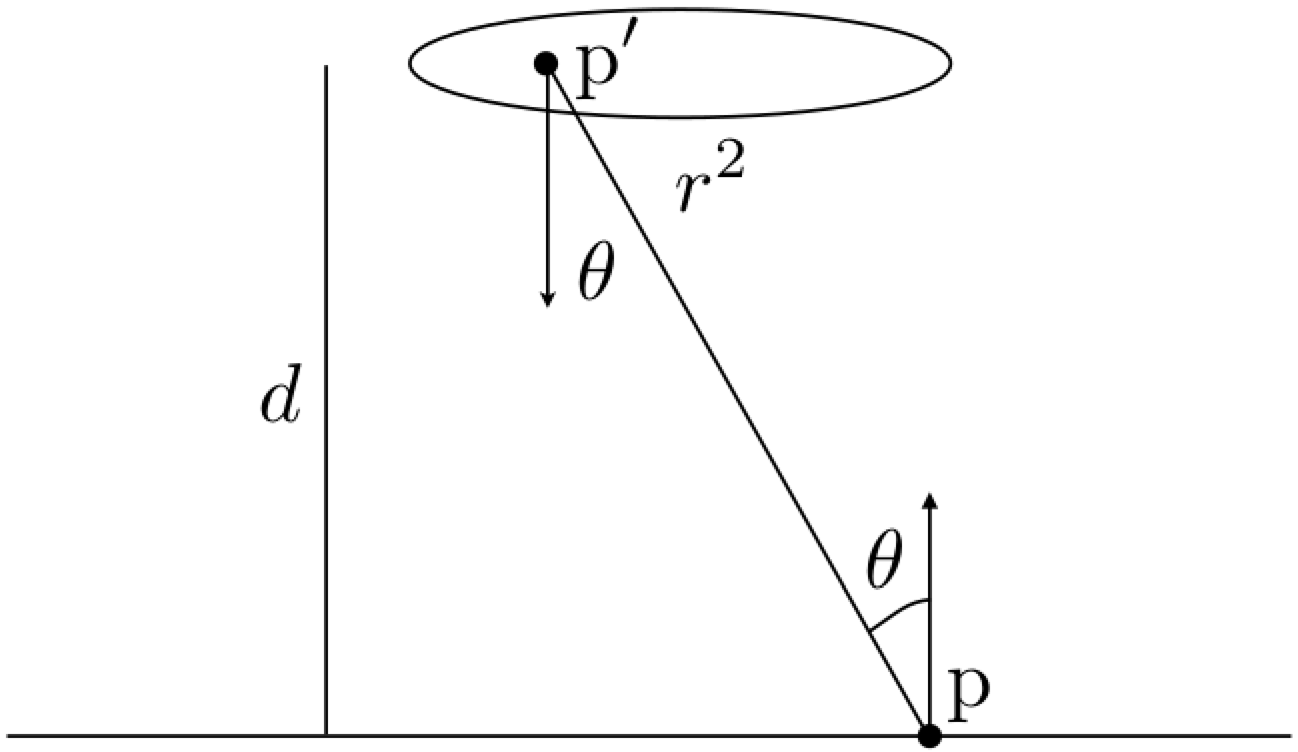

Sensor plane

$$E(\mathrm{p}, t) = \int_A L(\mathrm{p}' \to \mathrm{p}, t) \frac{\cos^2\theta}{||\mathrm{p}' - \mathrm{p}||^2} \, \mathrm{d}A'$$

$$= \frac{1}{d^2} \int_A L(\mathrm{p}' \to \mathrm{p}, t) \cos^4\theta \, \mathrm{d}A'$$

What does this say about the image I am capturing?

Vignetting

Fancy word for: pixels far off the center receive less light

white wall under uniform light

more interesting example of vignetting

Four types of vignetting:

- Mechanical: light rays blocked by hoods, filters, and other objects.
- Lens: similar, but light rays blocked by lens elements.
- Natural: due to radiometric laws (“cosine fourth falloff”).
- Pixel: angle-dependent sensitivity of photodiodes.

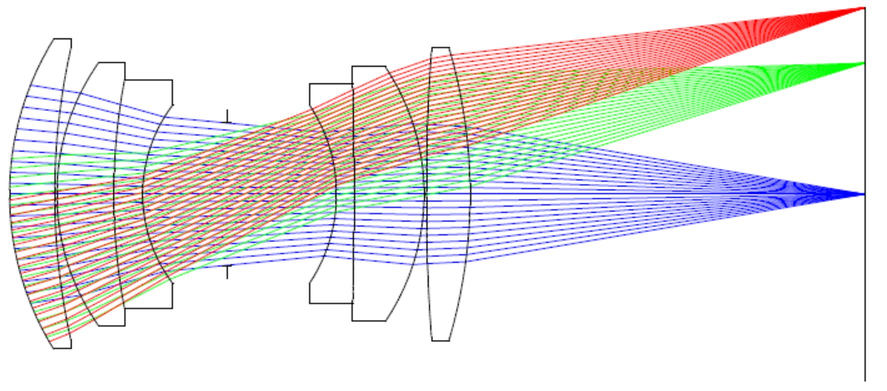

Quiz 2: BRDF of the moon

What BRDF does the moon have?

Quiz 2: BRDF of the moon

What BRDF does the moon have?

- Can it be diffuse?

Quiz 2: BRDF of the moon

What BRDF does the moon have?

- Can it be diffuse?

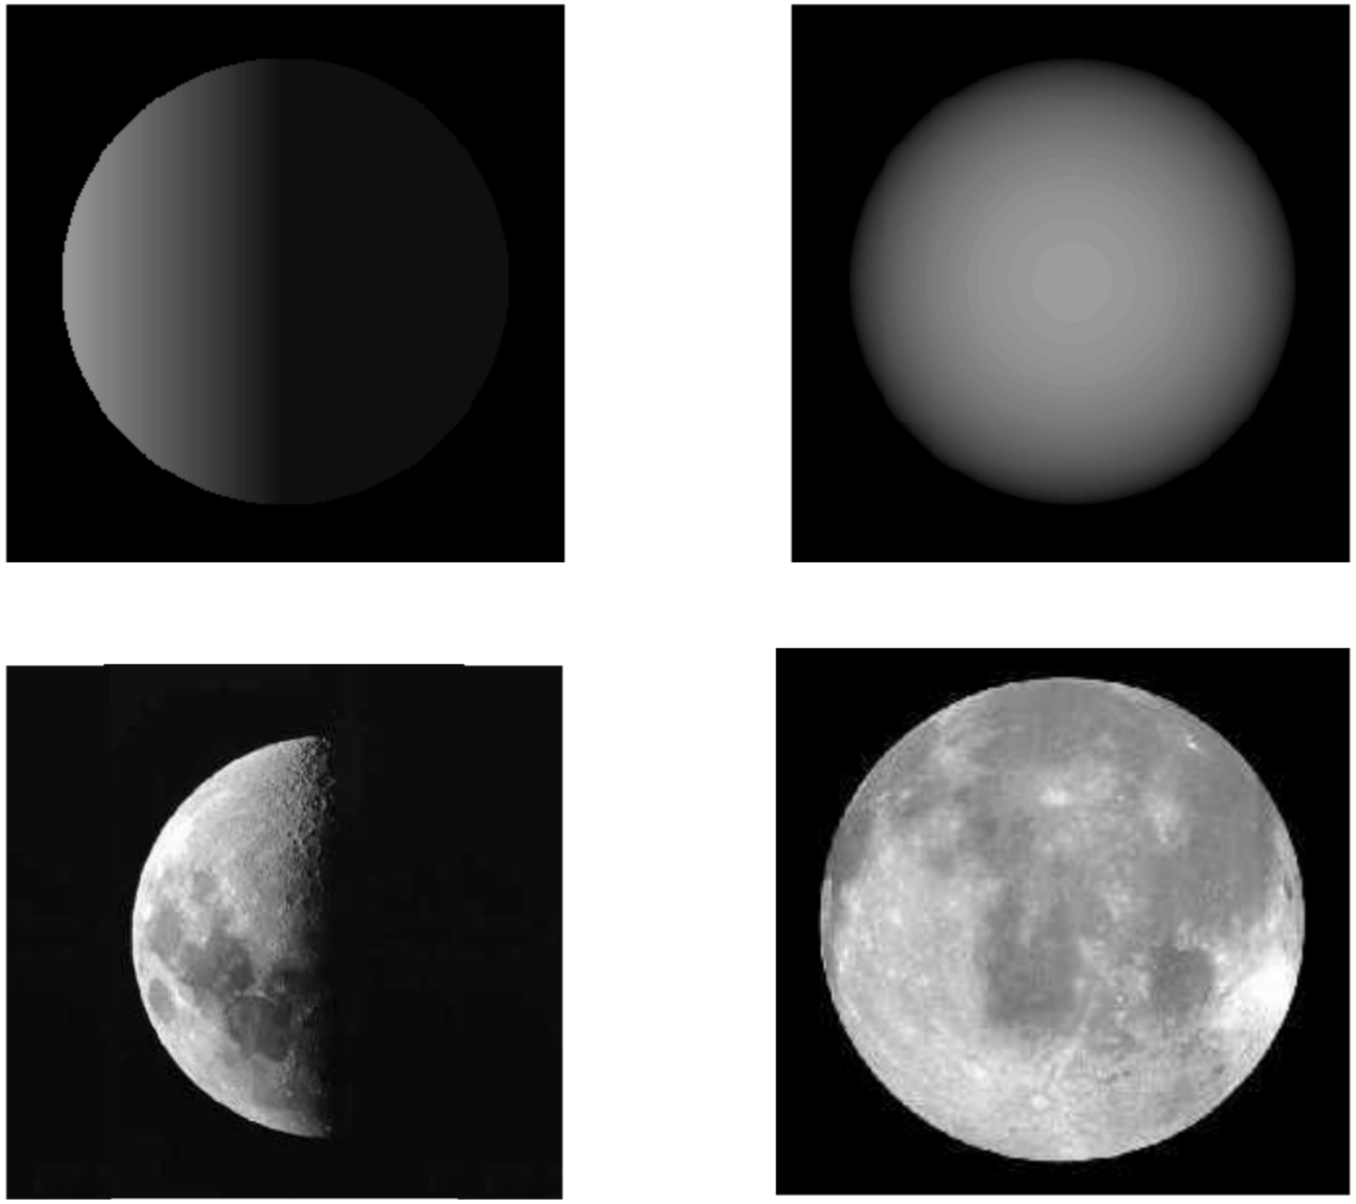

$\rightarrow n \cdot l$

Even though the moon appears matte, its edges remain bright.

Rough diffuse appearance

Surface Roughness Causes Flat Appearance

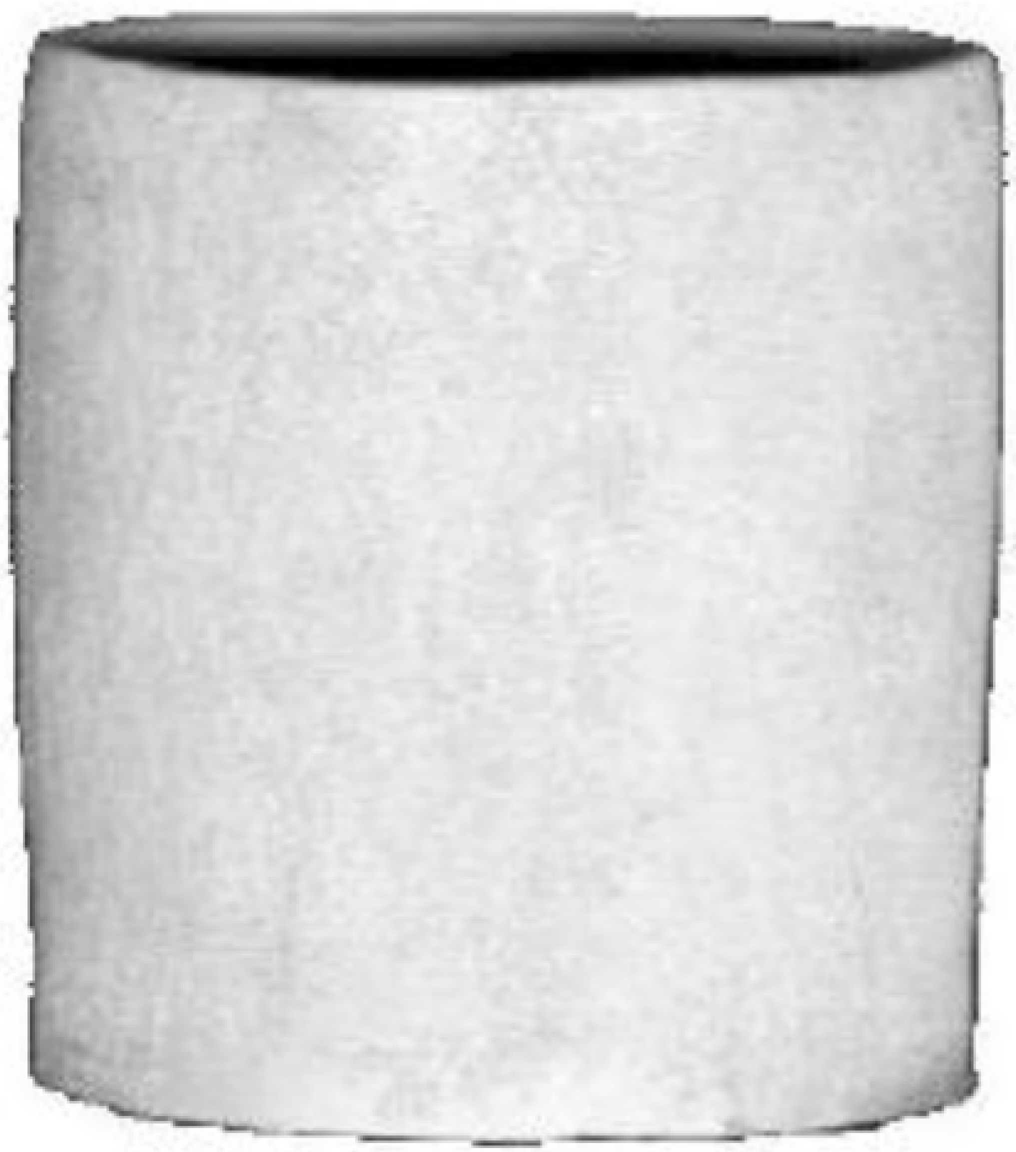

Actual Vase

Lambertian Vase

Five important equations/integrals to remember

Flux measured by a sensor of area X and directional receptivity W:

$$\Phi(W,X)=\int_X\int_W L(\hat{\boldsymbol{\omega}},x)\cos\theta d\boldsymbol{\omega} dA$$

Reflectance equation:

$$L^{\text{out}}(\hat{\boldsymbol{\omega}})=\int_{\Omega_{\text{in}}} f(\hat{\boldsymbol{\omega}}_{\text{in}},\hat{\boldsymbol{\omega}}_{\text{out}})L^{\text{in}}(\hat{\boldsymbol{\omega}}_{\text{in}})\cos\theta_{\text{in}}d\hat{\boldsymbol{\omega}}_{\text{in}}$$

Radiance under directional lighting and Lambertian BRDF ("n-dot-l shading"):

$$L^{\text{out}}=a\hat{\mathbf{n}}^\top\vec{\boldsymbol{\ell}}$$

Conversion of a (hemi)-spherical integral to a surface integral:

$$\int_{H^2} L_i(\mathrm{p},\omega',t)\ \cos\theta\,\mathrm{d}\omega'\ =\int_A L(\mathrm{p}'\to\mathrm{p},t)\frac{\cos\theta\cos\theta'}{||\mathrm{p}'-\mathrm{p}||^2}\,\mathrm{d}A'$$

Computing (hemi)-spherical integrals:

$$d\omega=\frac{dA}{r^2}=\sin\theta\,d\theta\,d\phi$$

and

$$\int d\omega=\int_0^{\pi}\int_0^{2\pi}\sin\theta\,d\theta\,d\phi$$

Photometric stereo

Image Intensity and 3D Geometry

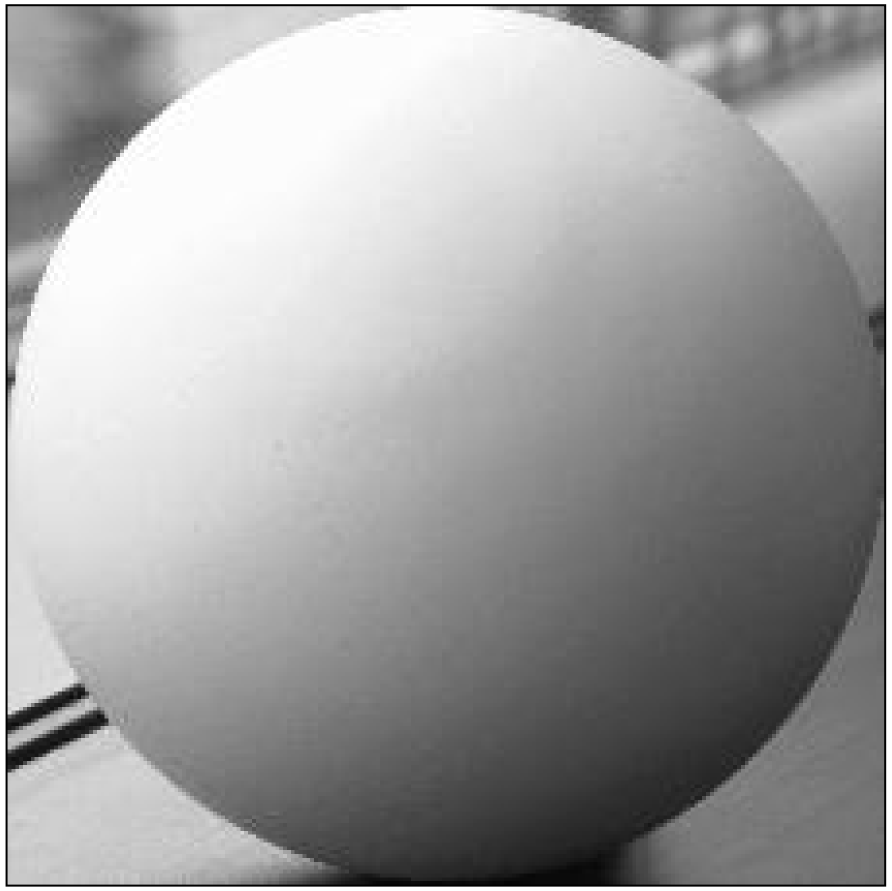

- *Shading* as a cue for shape reconstruction
- What is the relation between intensity and shape?

"N-dot-l" shading

ASSUMPTION 1:
LAMBERTIAN

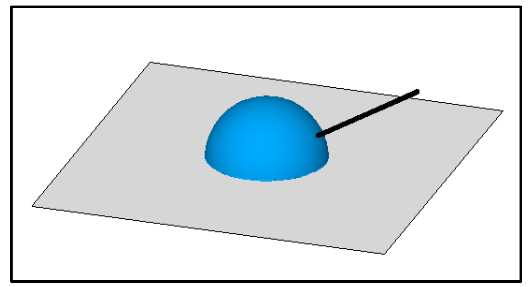

ASSUMPTION 2:
DIRECTIONAL LIGHTING

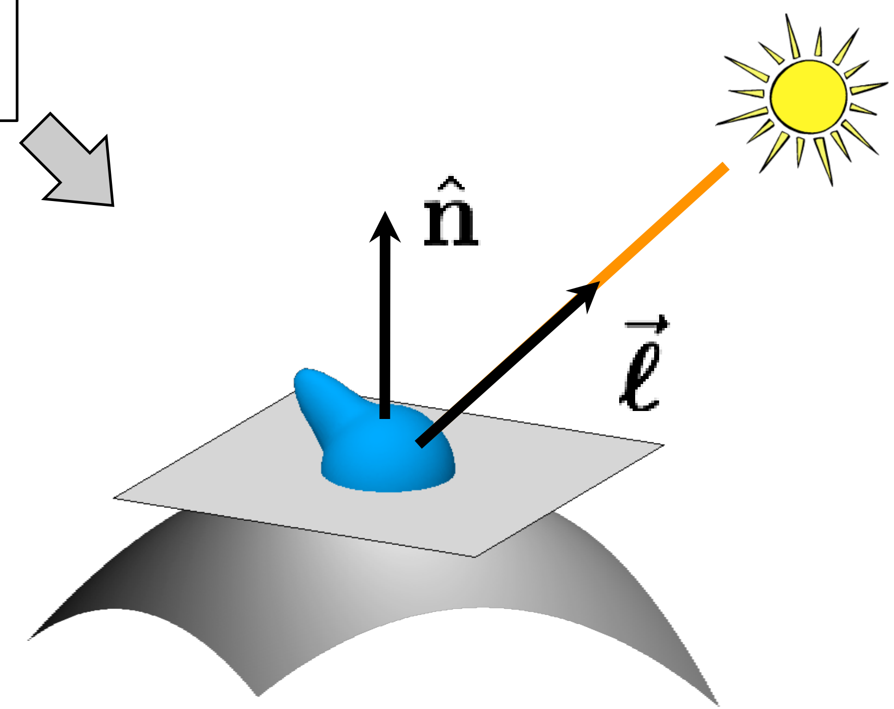

$$L^{\text{out}}(\hat{\boldsymbol{\omega}}) = \int_{\Omega_{\text{in}}} f(\hat{\boldsymbol{\omega}}_{\text{in}}, \hat{\boldsymbol{\omega}}_{\text{out}}) L^{\text{in}}(\hat{\boldsymbol{\omega}}_{\text{in}}) \cos\theta_{\text{in}} d\hat{\boldsymbol{\omega}}_{\text{in}}$$

$$\boxed{I = a\hat{\mathbf{n}}^\top \vec{\ell}}$$

Why do we call these normal "shape"?

Surfaces and normals

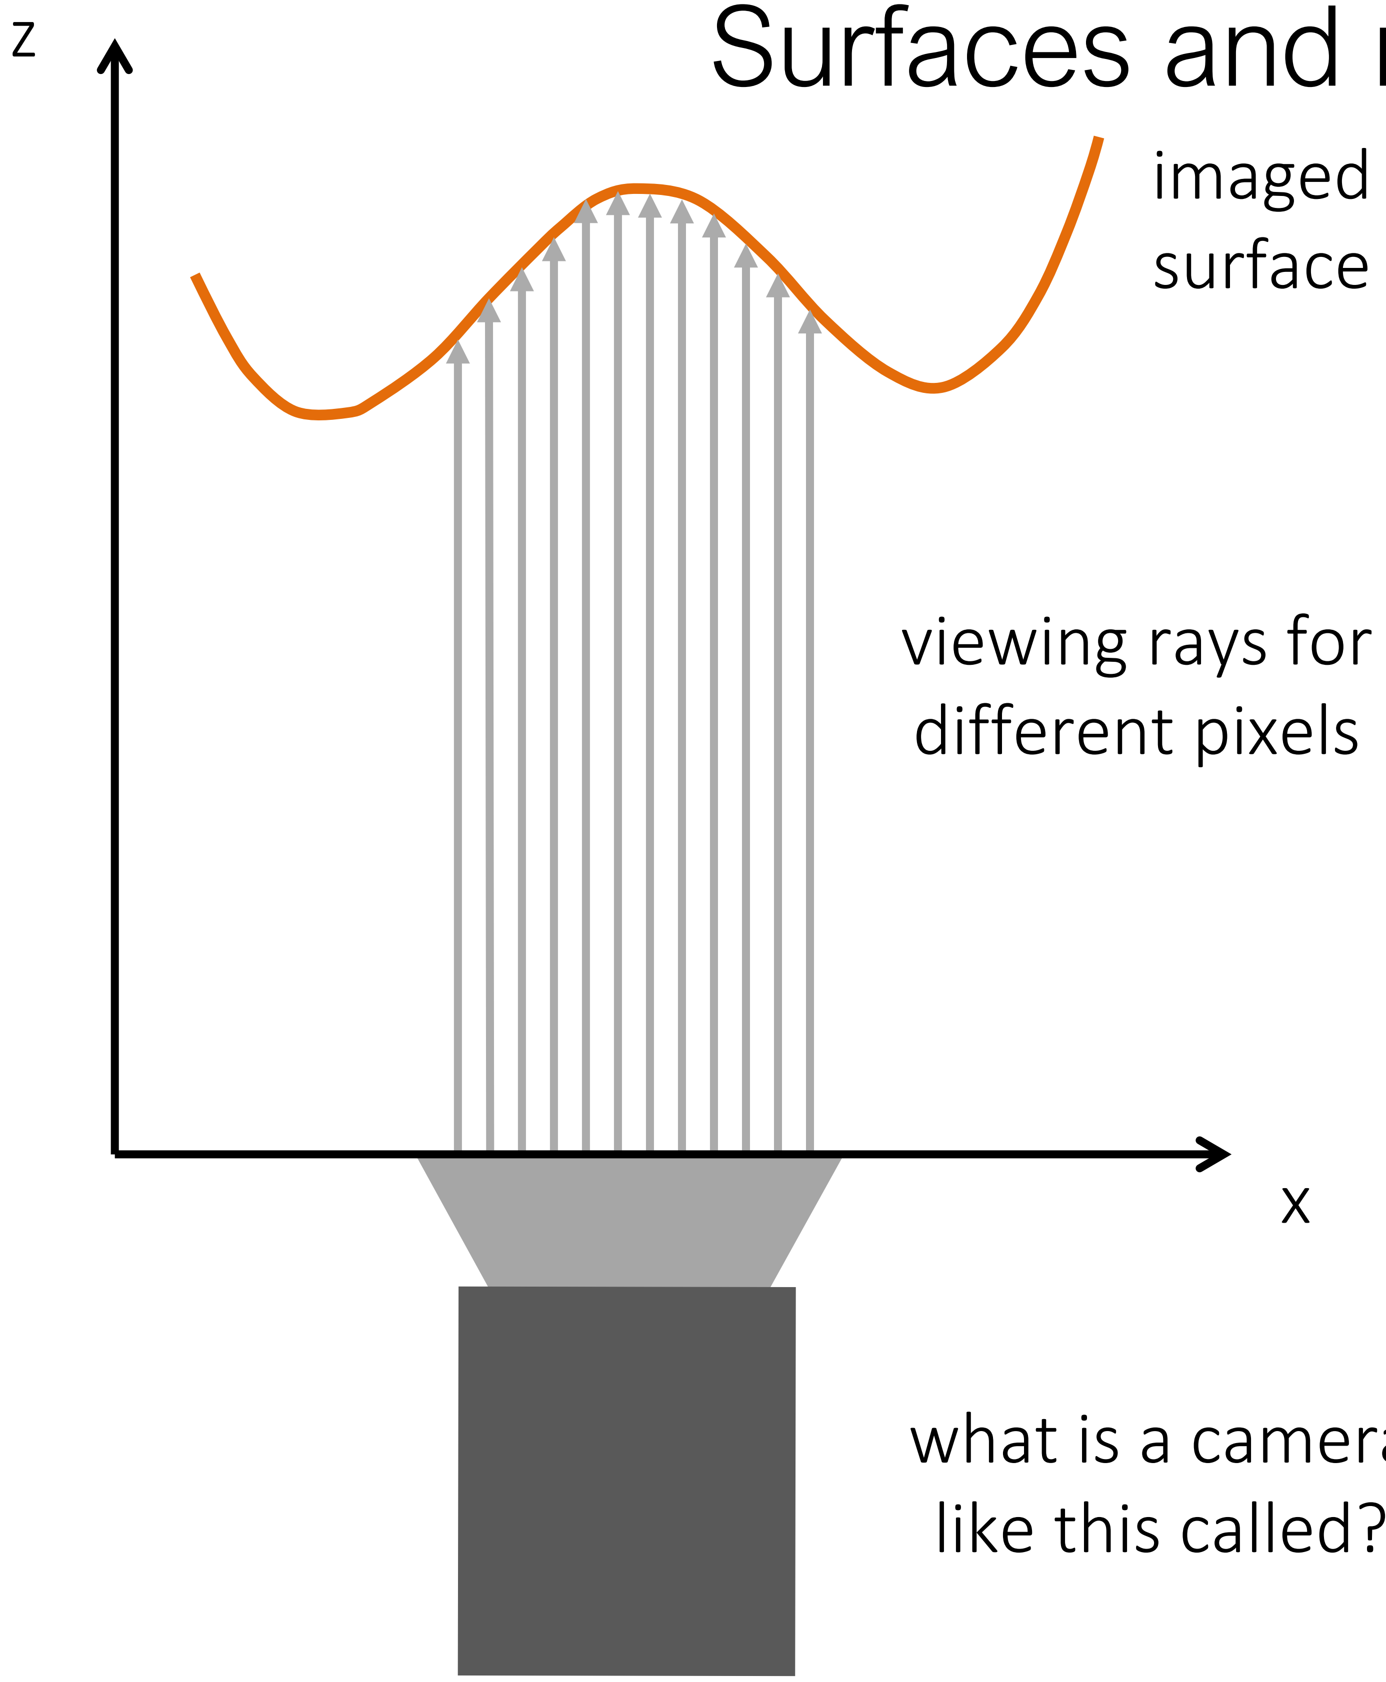

what is a camera
like this called?

Surfaces and normals

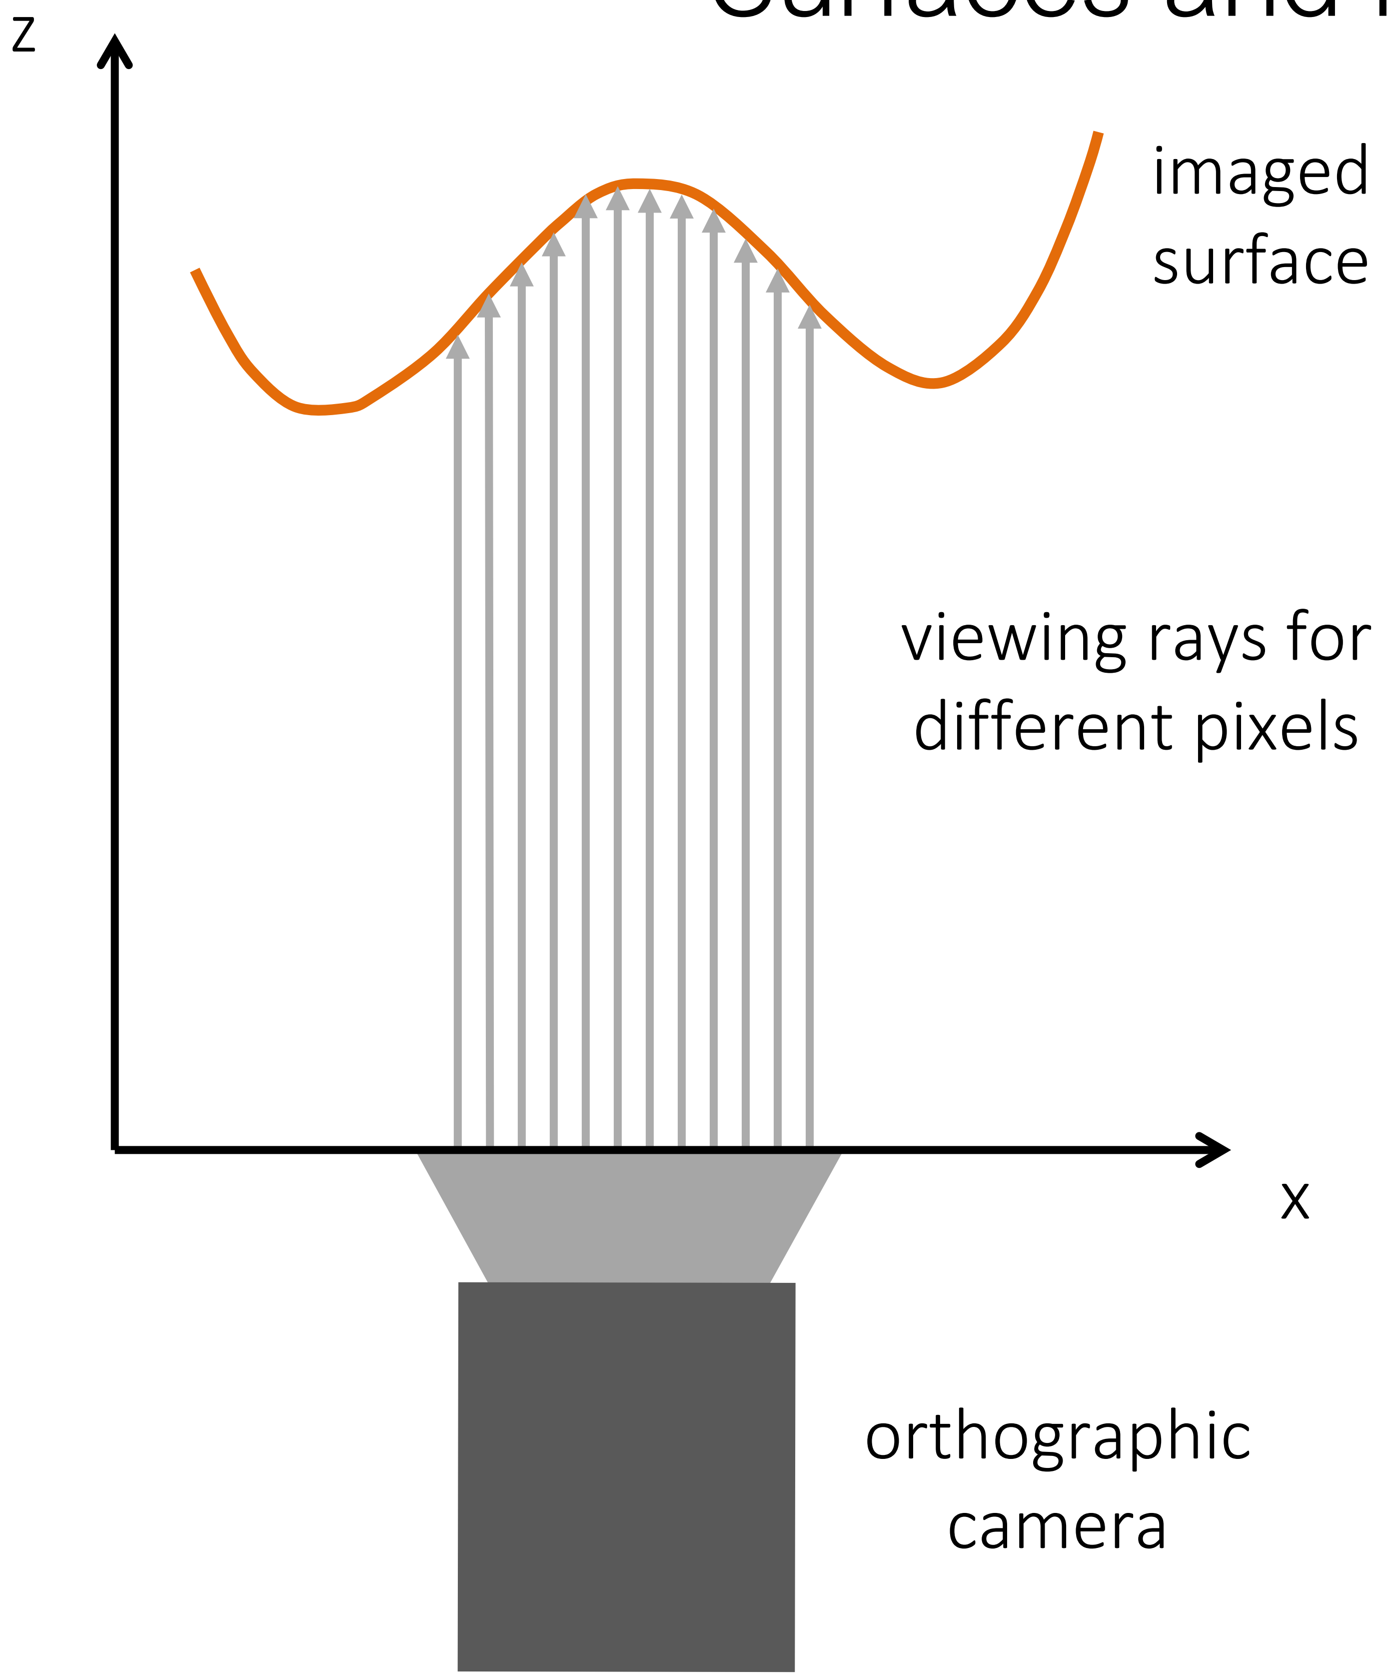

Surface representation as a depth field (also known as Monge surface):

$$\underset{\uparrow\ \text{depth at each pixel}}{z} = f\underbrace{(x, y)}_{\text{pixel coordinates on image place}}$$

How does surface normal relate to this representation?

Surfaces and normals

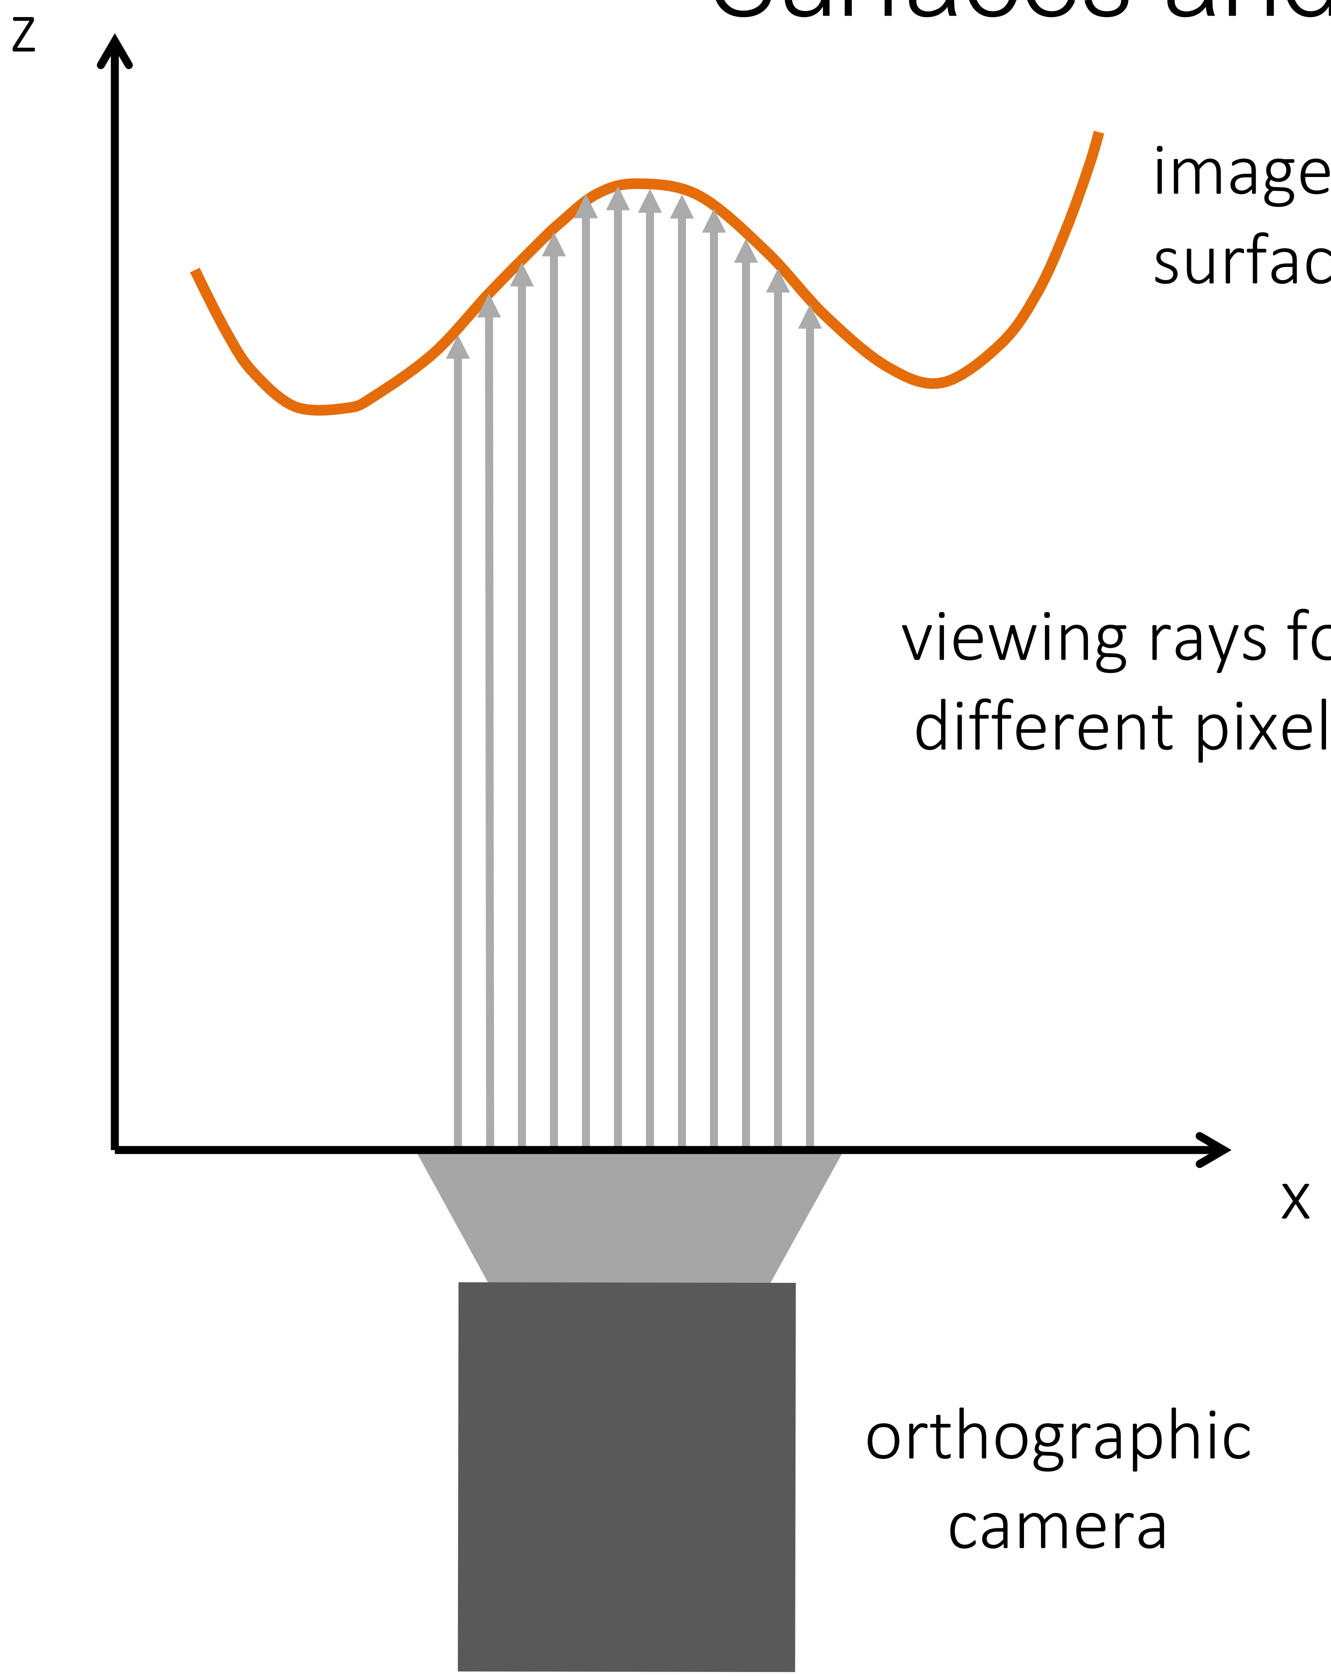

Surface representation as a depth image (also known as Monge surface):

$$z = f(\underbrace{x, y}_{\text{pixel coordinates on image place}})$$

depth at each pixel

Unnormalized normal:

$$\tilde{n}(x, y) = \left(\frac{df}{dx}, \frac{df}{dy}, -1\right)$$

Actual normal:

$$n(x, y) = \tilde{n}(x, y) / \|\tilde{n}(x, y)\|$$

Normals are scaled spatial derivatives of depth image!

Shape from a Single Image?

- Given a single image of an object with known surface reflectance taken under a known light source, can we recover the shape of the object?

$$\underline{I} = \vec{n} \cdot \underline{L}$$

Human Perception

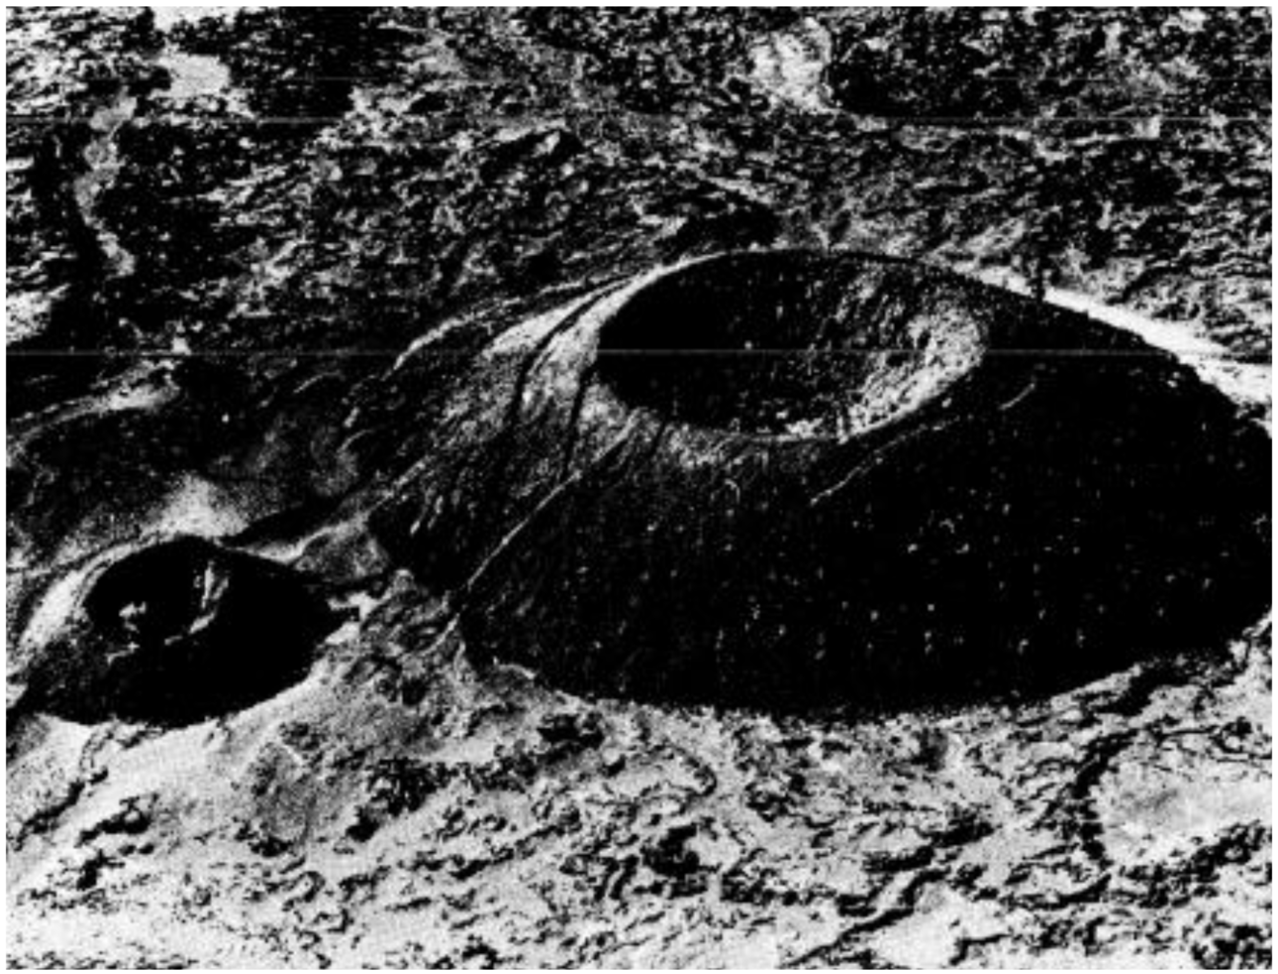

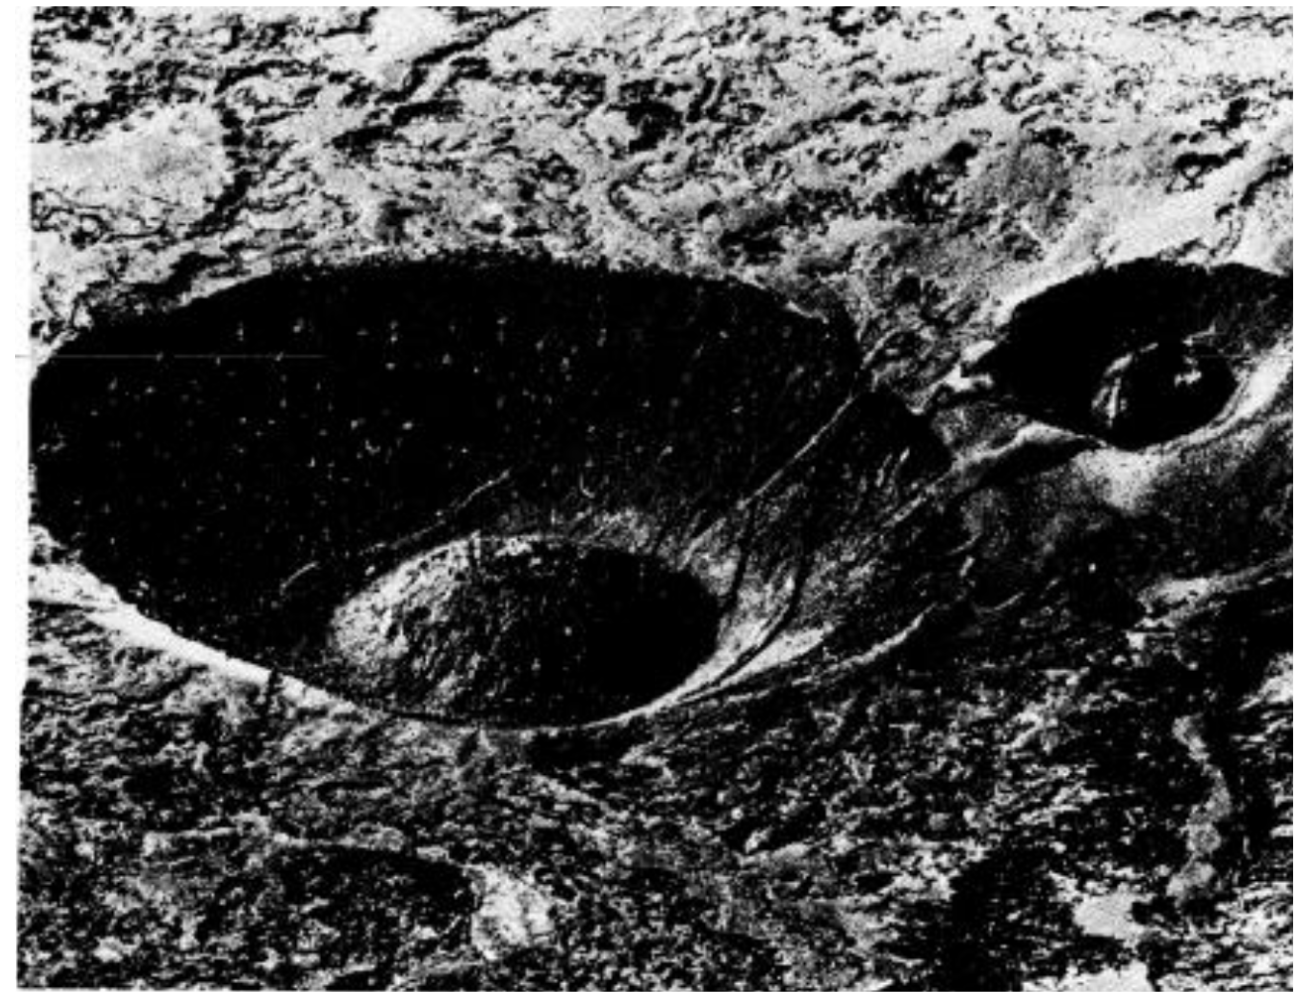

Examples of the classic bump/dent stimuli used to test lighting assumptions when judging shape from shading, with shading orientations (a) 0° and (b) 180° from the vertical.

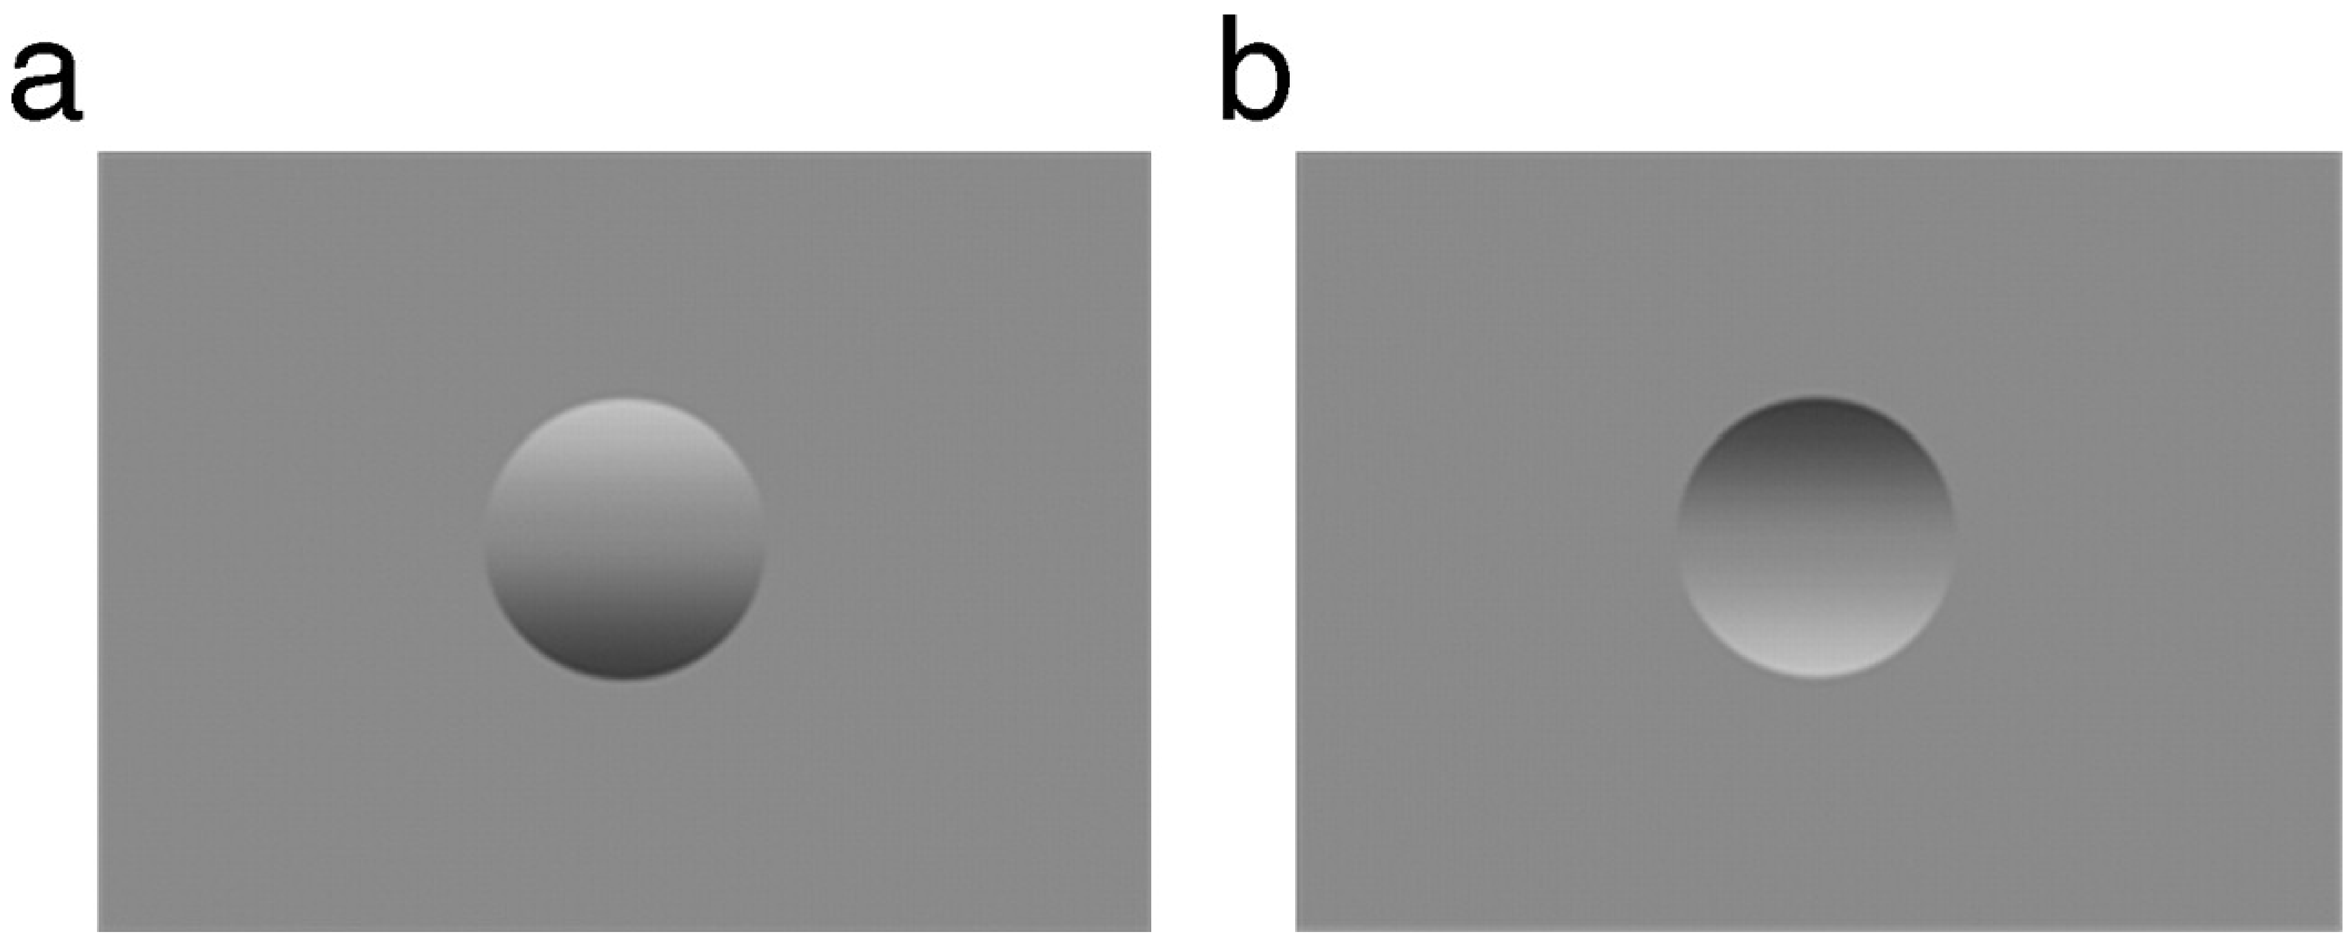

Thomas R et al. J Vis 2010;10:6

Human Perception

- Our brain often perceives shape from shading.
- Mostly, it makes many assumptions to do so.
- For example:

  Light is coming from above (sun).

  Biased by occluding contours.

by V. Ramachandran

Single-lighting is ambiguous

ASSUMPTION 1:
LAMBERTIAN

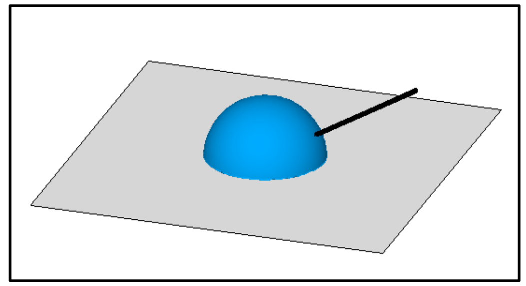

ASSUMPTION 2:
DIRECTIONAL LIGHTING

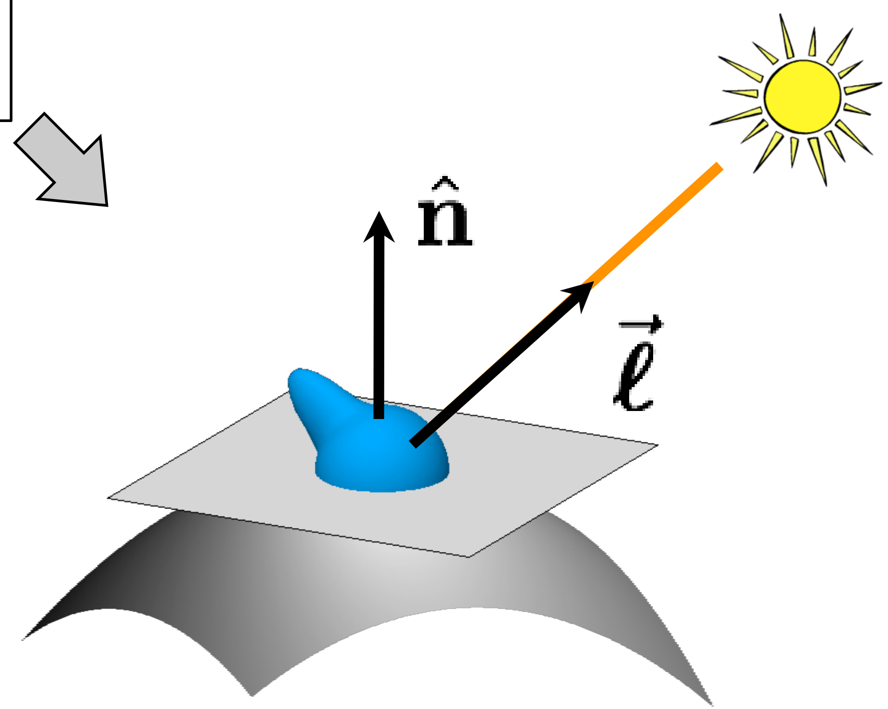

$$L^{\mathrm{out}}(\hat{\boldsymbol{\omega}}) = \int_{\Omega_{\mathrm{in}}} f(\hat{\boldsymbol{\omega}}_{\mathrm{in}}, \hat{\boldsymbol{\omega}}_{\mathrm{out}}) L^{\mathrm{in}}(\hat{\boldsymbol{\omega}}_{\mathrm{in}}) \cos\theta_{\mathrm{in}} d\hat{\boldsymbol{\omega}}_{\mathrm{in}}$$

$$I = a\hat{\mathbf{n}}^{\top}\vec{\boldsymbol{\ell}}$$

Lambertian photometric stereo

$$
\begin{aligned}
I_1 &= a\hat{\boldsymbol{n}}^\top \vec{\boldsymbol{\ell}}_1 \\
I_2 &= a\hat{\boldsymbol{n}}^\top \vec{\boldsymbol{\ell}}_2 \\
&\vdots \\
I_N &= a\hat{\boldsymbol{n}}^\top \vec{\boldsymbol{\ell}}_N
\end{aligned}
$$

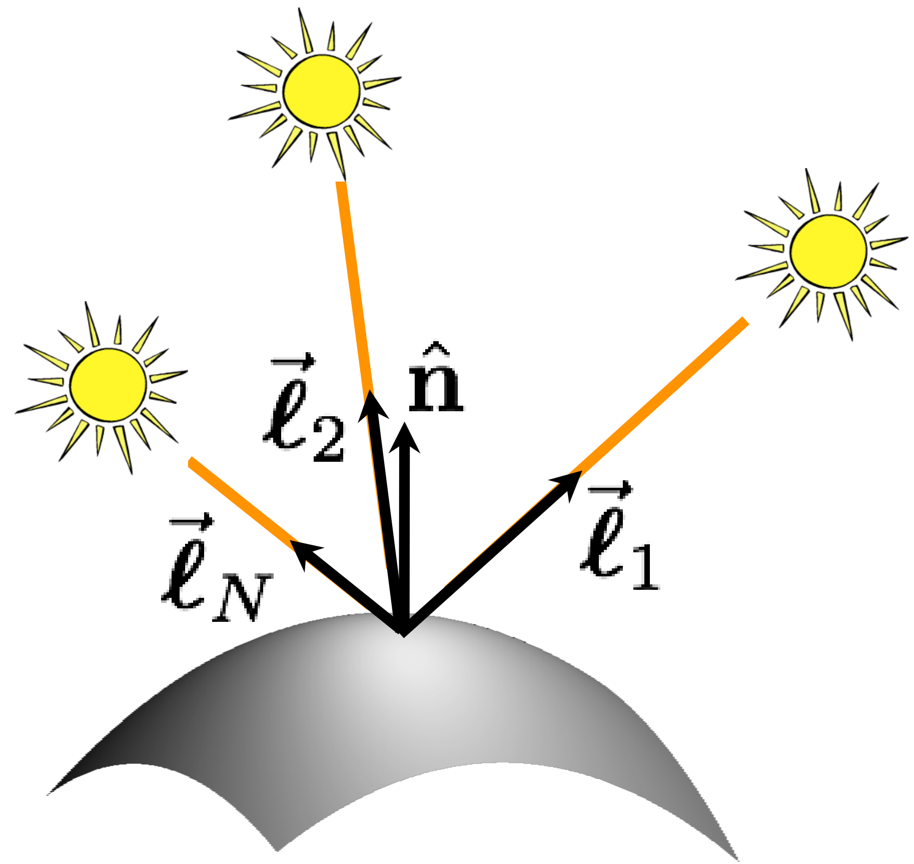

Assumption: We know the lighting directions.

Lambertian photometric stereo

$$
\begin{aligned}
I_1 &= a\hat{\boldsymbol{n}}^\top \vec{\boldsymbol{\ell}}_1 \\
I_2 &= a\hat{\boldsymbol{n}}^\top \vec{\boldsymbol{\ell}}_2 \\
&\vdots \\
I_N &= a\hat{\boldsymbol{n}}^\top \vec{\boldsymbol{\ell}}_N
\end{aligned}
$$

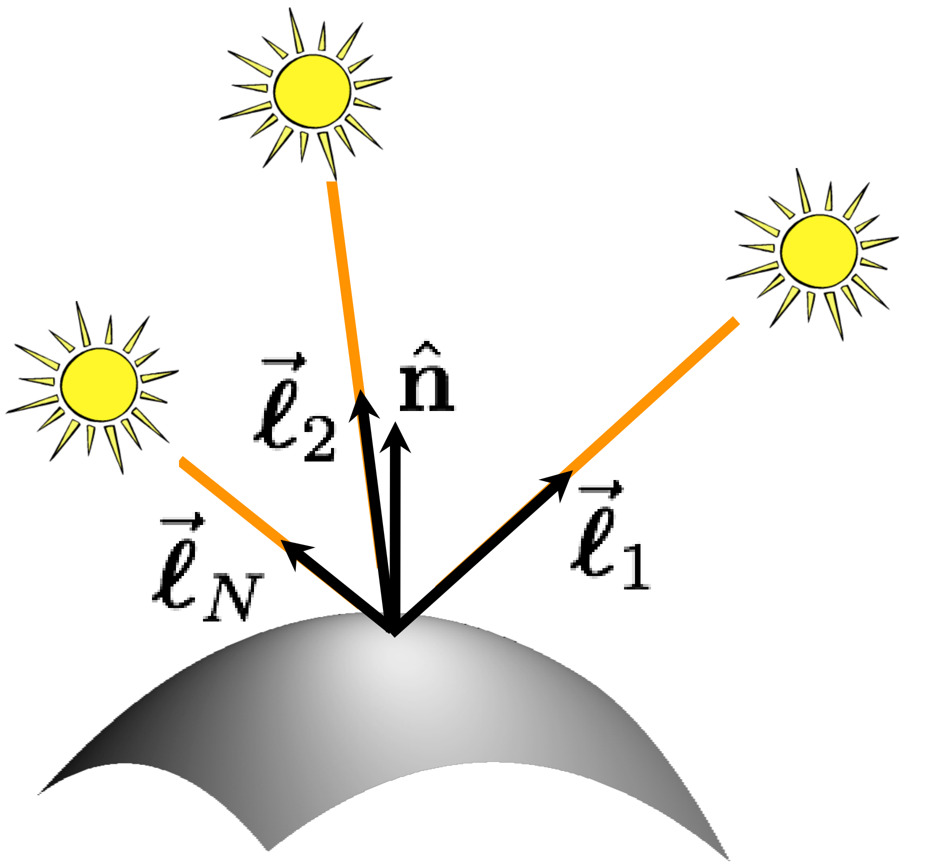

define “pseudo-normal” $\vec{\boldsymbol{b}} \triangleq a\hat{\boldsymbol{n}}$

solve linear system for pseudo-normal

$$
\begin{bmatrix} I_1 \\ I_2 \\ \vdots \\ I_N \end{bmatrix} = \begin{bmatrix} \vec{\boldsymbol{\ell}}_1^\top \\ \vec{\boldsymbol{\ell}}_2^\top \\ \vdots \\ \vec{\boldsymbol{\ell}}_N^\top \end{bmatrix} \begin{bmatrix} \vec{\boldsymbol{b}} \end{bmatrix}
$$

What are the dimensions of these matrices?

Lambertian photometric stereo

$$I_1 = a\hat{\boldsymbol{n}}^\top \vec{\boldsymbol{\ell}}_1$$
$$I_2 = a\hat{\boldsymbol{n}}^\top \vec{\boldsymbol{\ell}}_2$$
$$\vdots$$
$$I_N = a\hat{\boldsymbol{n}}^\top \vec{\boldsymbol{\ell}}_N$$

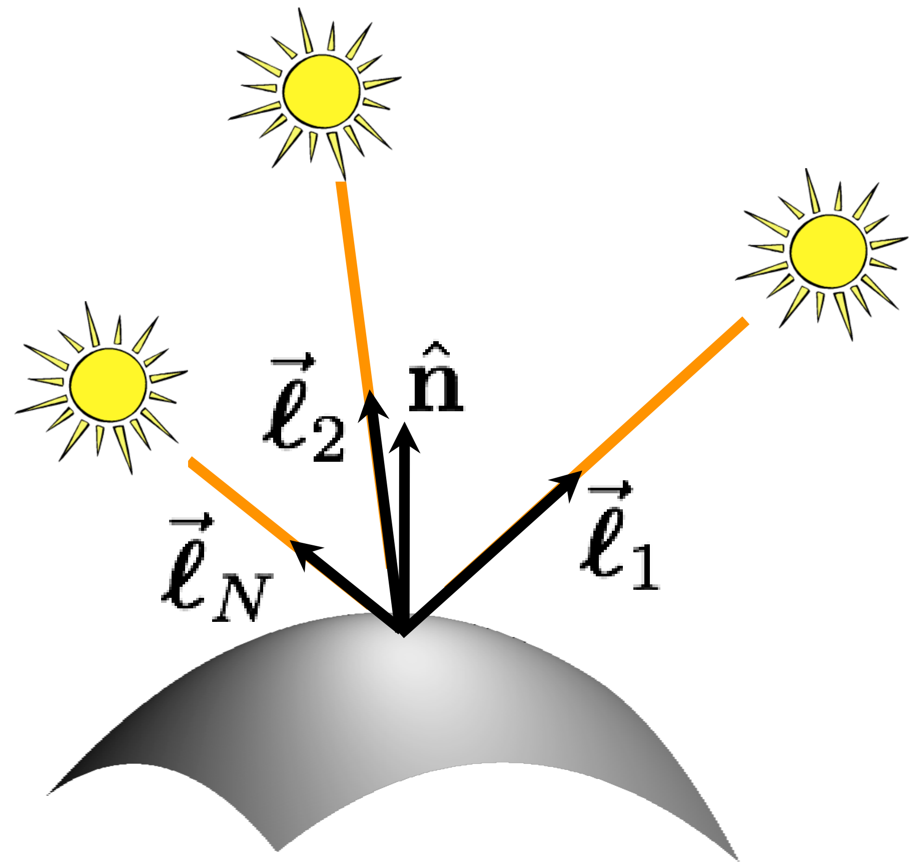

define "pseudo-normal" $\vec{\boldsymbol{b}} \triangleq a\hat{\boldsymbol{n}}$

solve linear system for pseudo-normal

$$\begin{bmatrix} I_1 \\ I_2 \\ \vdots \\ I_N \end{bmatrix}_{N\times 1} = \begin{bmatrix} \vec{\boldsymbol{\ell}}_1^\top \\ \vec{\boldsymbol{\ell}}_2^\top \\ \vdots \\ \vec{\boldsymbol{\ell}}_N^\top \end{bmatrix}_{N\times 3} \begin{bmatrix} \vec{\boldsymbol{b}} \end{bmatrix}_{3\times 1}$$

What are the knowns and unknowns?

Lambertian photometric stereo

$$
\begin{aligned}
I_1 &= a\hat{\boldsymbol{n}}^\top \vec{\boldsymbol{\ell}}_1 \\
I_2 &= a\hat{\boldsymbol{n}}^\top \vec{\boldsymbol{\ell}}_2 \\
&\vdots \\
I_N &= a\hat{\boldsymbol{n}}^\top \vec{\boldsymbol{\ell}}_N
\end{aligned}
$$

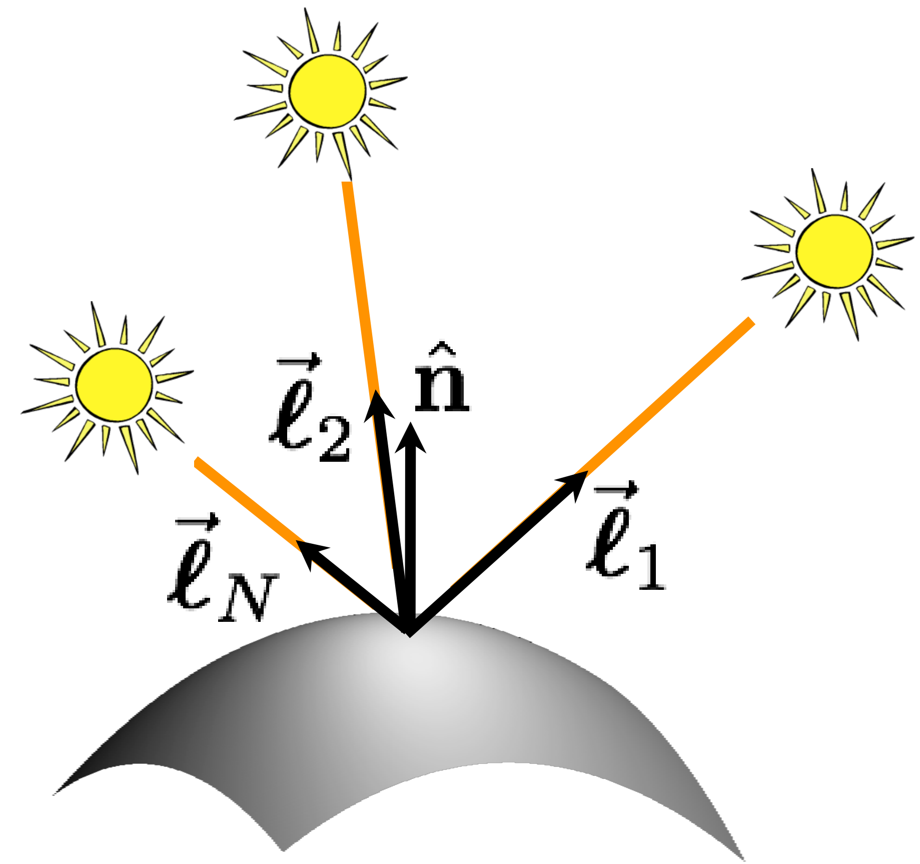

define "pseudo-normal" $\vec{\boldsymbol{b}} \triangleq a\hat{\boldsymbol{n}}$

solve linear system for pseudo-normal

$$
\begin{bmatrix} I_1 \\ I_2 \\ \vdots \\ I_N \end{bmatrix}_{N\times 1} = \begin{bmatrix} \vec{\boldsymbol{\ell}}_1^\top \\ \vec{\boldsymbol{\ell}}_2^\top \\ \vdots \\ \vec{\boldsymbol{\ell}}_N^\top \end{bmatrix}_{N\times 3} \begin{bmatrix} \vec{\boldsymbol{b}} \end{bmatrix}_{3\times 1}
$$

How many lights do I need for unique solution?

Lambertian photometric stereo

$$
\begin{aligned}
I_1 &= a\hat{\boldsymbol{n}}^\top \vec{\boldsymbol{\ell}}_1 \\
I_2 &= a\hat{\boldsymbol{n}}^\top \vec{\boldsymbol{\ell}}_2 \\
&\vdots \\
I_N &= a\hat{\boldsymbol{n}}^\top \vec{\boldsymbol{\ell}}_N
\end{aligned}
$$

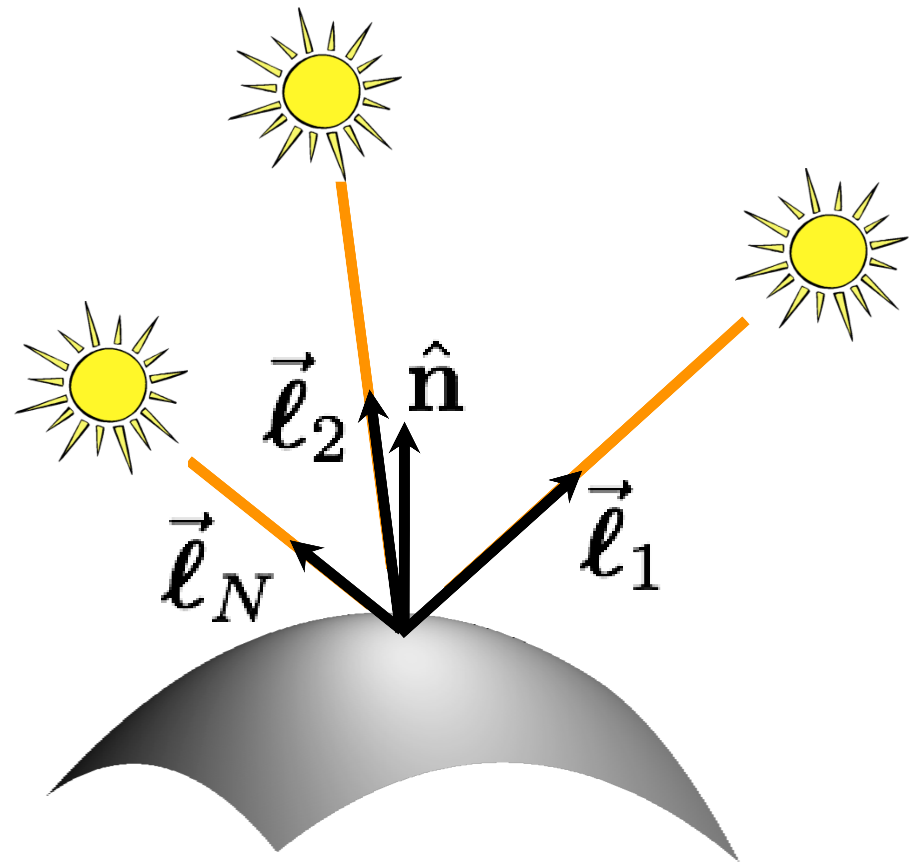

define “pseudo-normal” $\vec{\boldsymbol{b}} \triangleq a\hat{\boldsymbol{n}}$

solve linear system for pseudo-normal

$$
\begin{bmatrix} I_1 \\ I_2 \\ \vdots \\ I_N \end{bmatrix}_{N\times 1} = \begin{bmatrix} \vec{\boldsymbol{\ell}}_1^\top \\ \vec{\boldsymbol{\ell}}_2^\top \\ \vdots \\ \vec{\boldsymbol{\ell}}_N^\top \end{bmatrix}_{N\times 3} \begin{bmatrix} \vec{\boldsymbol{b}} \end{bmatrix}_{3\times 1}
$$

once system is solved, b gives normal direction and albedo

How do we solve this system?

Solving the Equation with three lights

$$\underbrace{\begin{bmatrix} I_1 \\ I_2 \\ I_2 \end{bmatrix}}_{\underset{3\times1}{\mathbf{I}}} = \underbrace{\begin{bmatrix} \mathbf{s}_1^T \\ \mathbf{s}_2^T \\ \mathbf{s}_3^T \end{bmatrix}}_{\underset{3\times3}{\mathbf{S}}} \underbrace{\rho\mathbf{n}}_{\underset{3\times1}{\tilde{\mathbf{n}}}}$$

$$\tilde{\mathbf{n}} = \mathbf{S}^{-1}\mathbf{I}$$

inverse

$$\rho = |\tilde{\mathbf{n}}|$$

$$\mathbf{n} = \frac{\tilde{\mathbf{n}}}{|\tilde{\mathbf{n}}|} = \frac{\tilde{\mathbf{n}}}{\rho}$$

Is there any reason to use more than three lights?

More than Three Light Sources

- Get better SNR by using more lights

$$\begin{bmatrix} I_1 \\ \vdots \\ I_N \end{bmatrix} = \begin{bmatrix} \mathbf{s}_1^T \\ \vdots \\ \mathbf{s}_N^T \end{bmatrix} \rho \mathbf{n}$$

- Least squares solution:

$$\mathbf{I} = \mathbf{S}\tilde{\mathbf{n}} \qquad \longleftarrow \qquad N \times 1 = \underline{(N \times 3)}(3 \times 1)$$

$$\mathbf{S}^T\mathbf{I} = \mathbf{S}^T\mathbf{S}\tilde{\mathbf{n}}$$

$$\tilde{\mathbf{n}} = \boxed{\left(\mathbf{S}^T\mathbf{S}\right)^{-1}\mathbf{S}^T}\mathbf{I}$$

Moore-Penrose pseudo inverse

- Solve for $\rho, \mathbf{n}$ as before

Computing light source directions

- Trick: place a chrome sphere in the scene

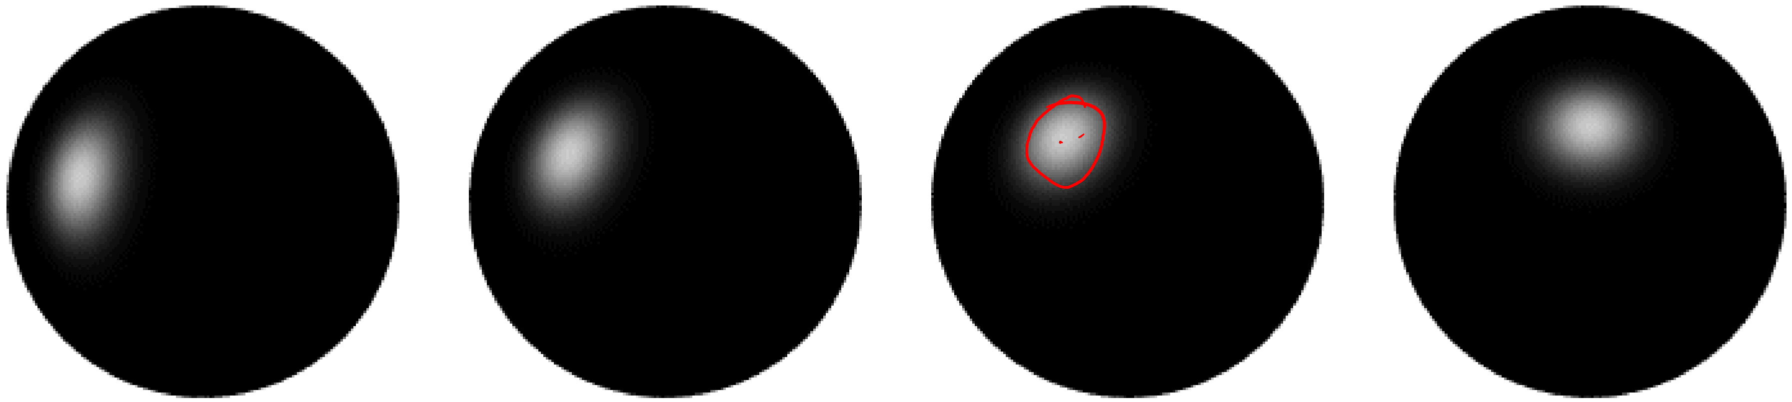

  – the location of the highlight tells you the source direction

Limitations

- Big problems
  - Doesn’t work for shiny things, semi-translucent things
  - Shadows, inter-reflections

- Smaller problems
  - Camera and lights have to be distant
  - Calibration requirements
    - measure light source directions, intensities
    - camera response function

Depth from normals

- Solving the linear system per-pixel gives us an estimated surface normal for each pixel

Input photo

Estimated normals

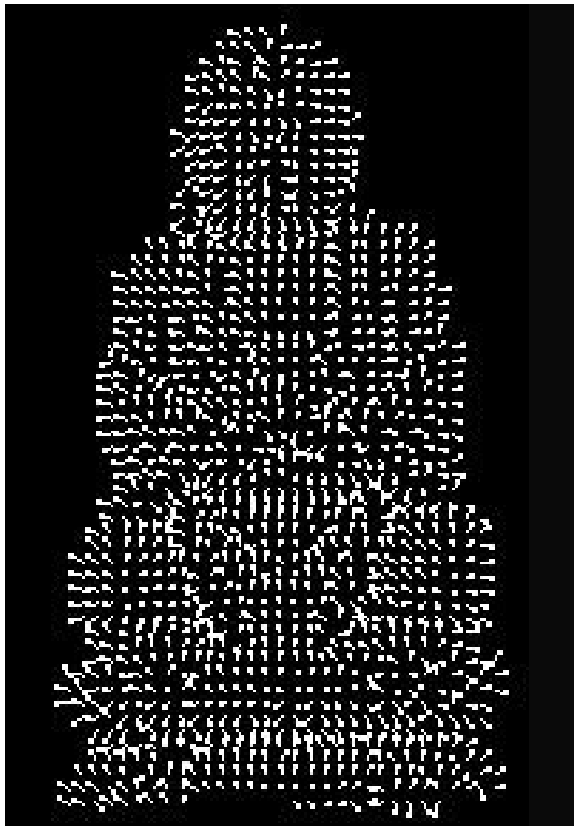

Estimated normals (needle diagram)

- How can we compute depth from normals?
  - Normals are like the “derivative” of the true depth

Surfaces and normals

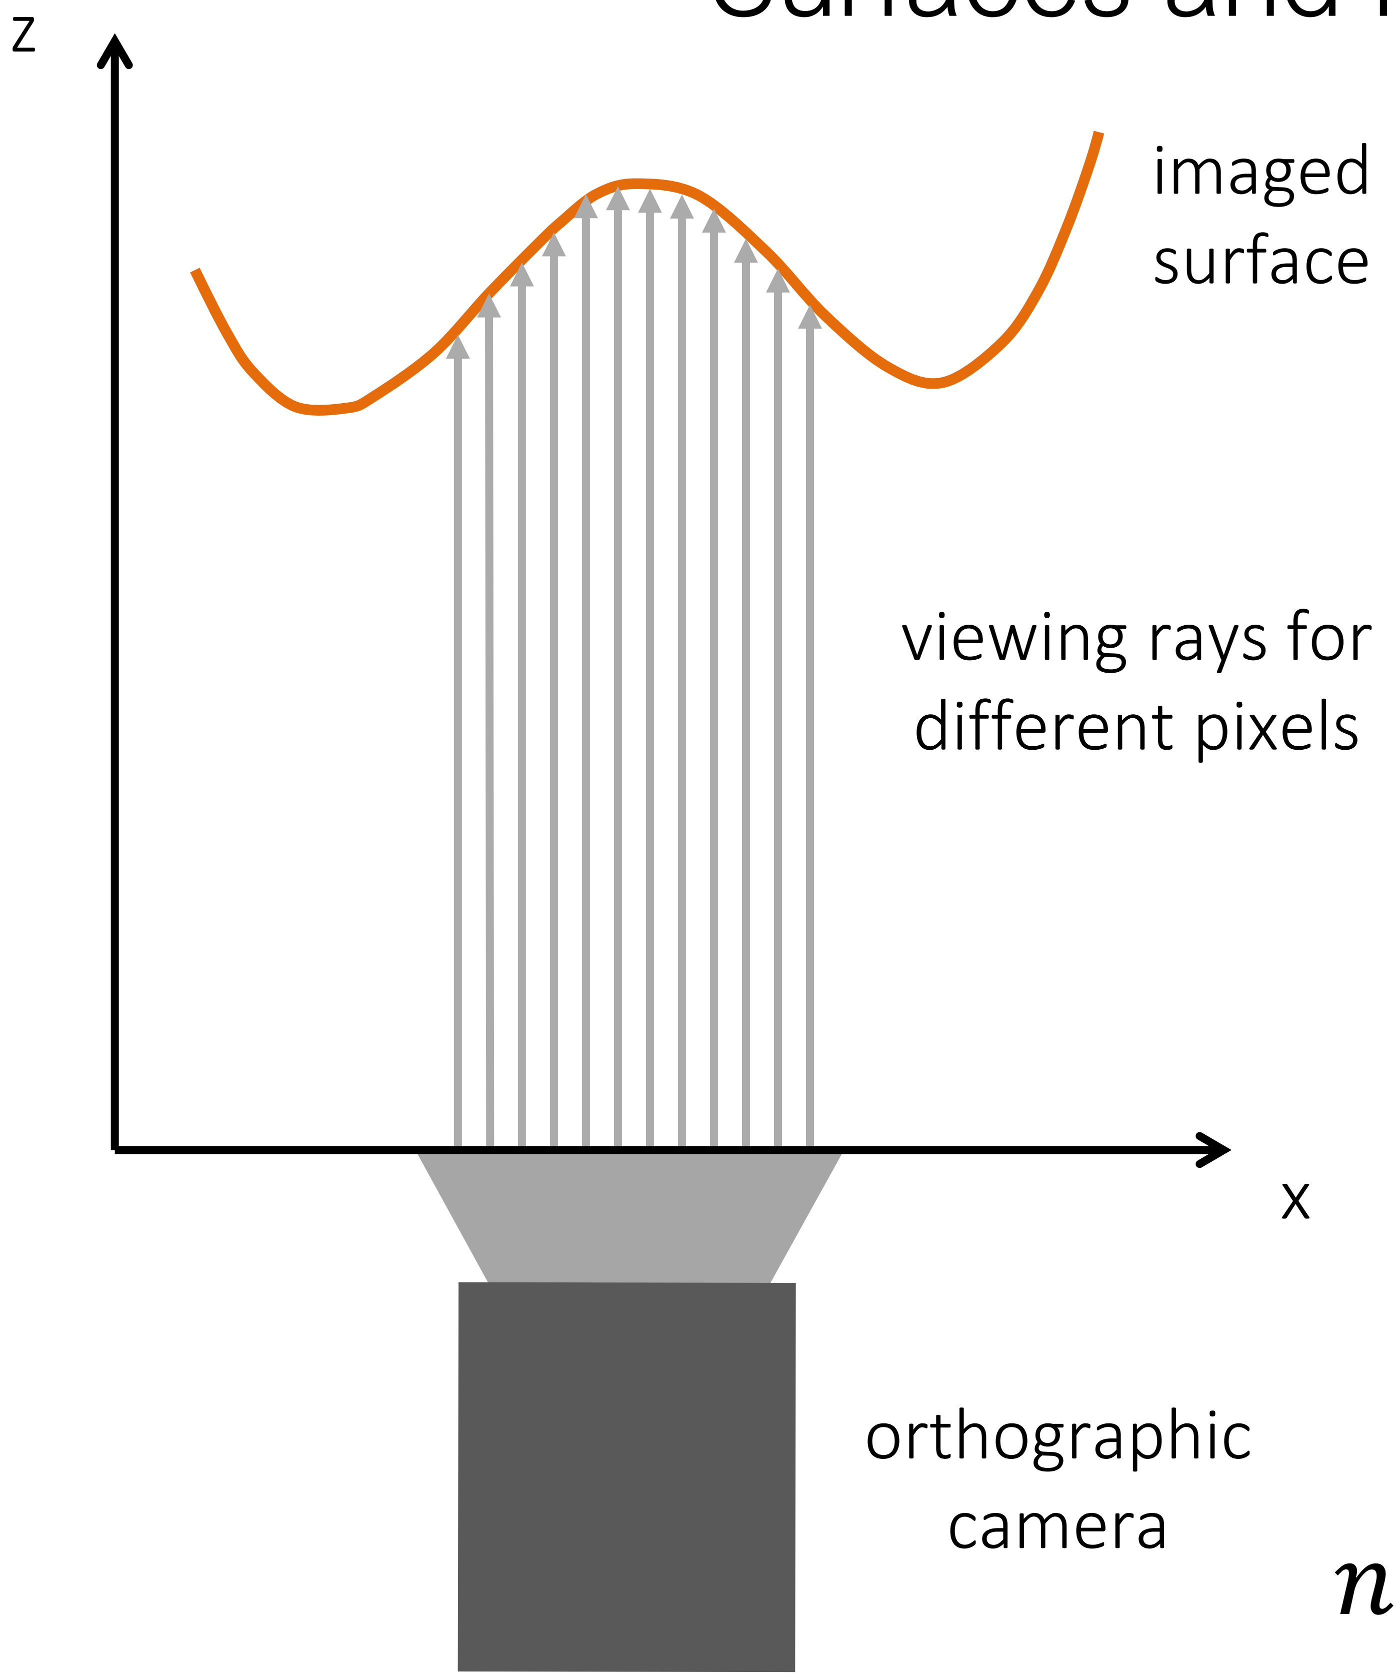

Surface representation as a depth image (also known as Monge surface):

$$\underset{\text{depth at each pixel}}{z} = f\underbrace{(x, y)}_{\text{pixel coordinates in image space}}$$

Unnormalized normal:

$$\tilde{n}(x, y) = \left(\frac{df}{dx}, \frac{df}{dy}, -1\right)$$

Actual normal:

$$n(x, y) = \tilde{n}(x, y) / \|\tilde{n}(x, y)\|$$

Normals are scaled spatial derivatives of depth image!

Depth from normals

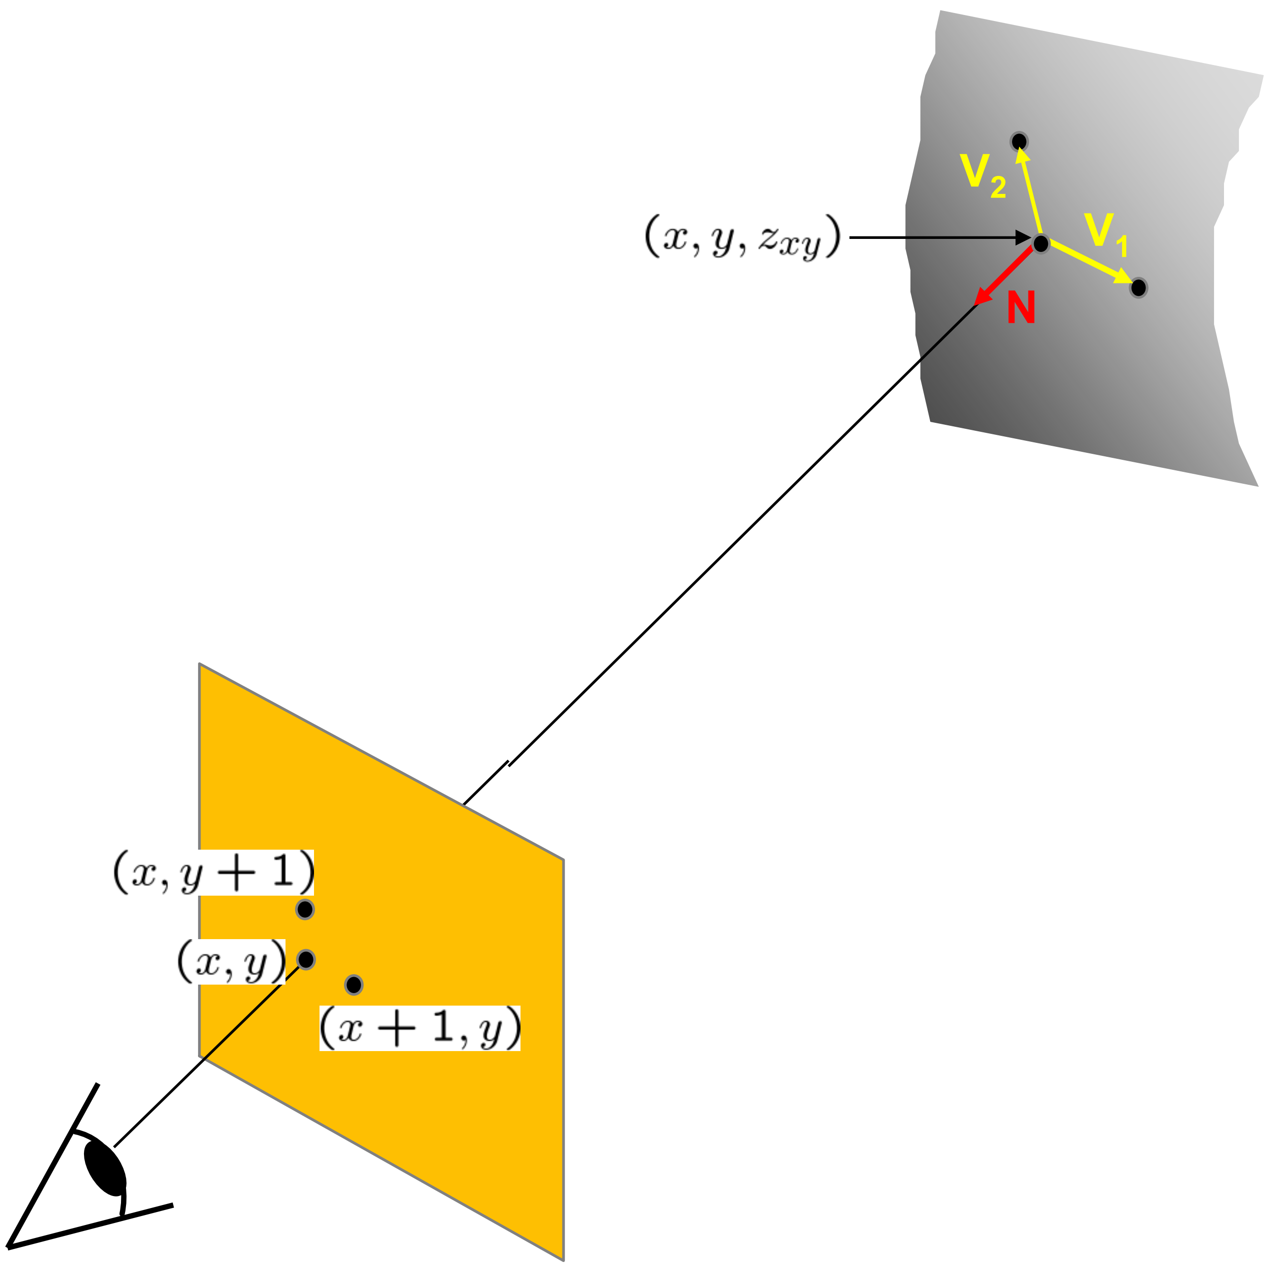

Use vector field integration techniques as in gradient-domain image processing.

Results

1. Estimate light source directions
2. Compute surface normals
3. Compute albedo values
4. Estimate depth from surface normals
5. Relight the object (with original texture and uniform albedo)

Results: Lambertian Sphere

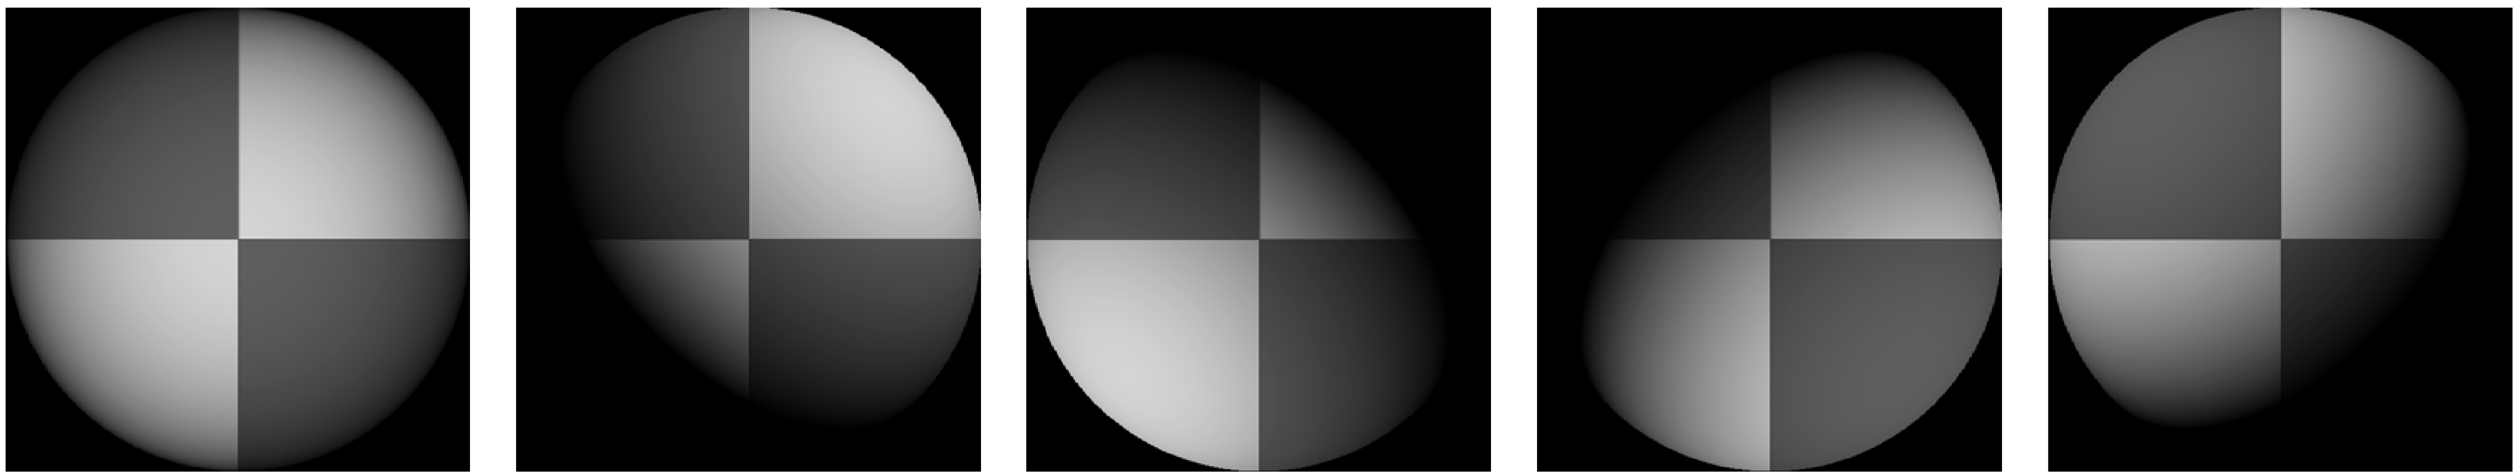

Input Images

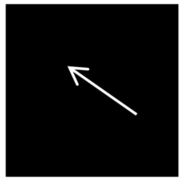

Needles are projections of surface normals on image plane

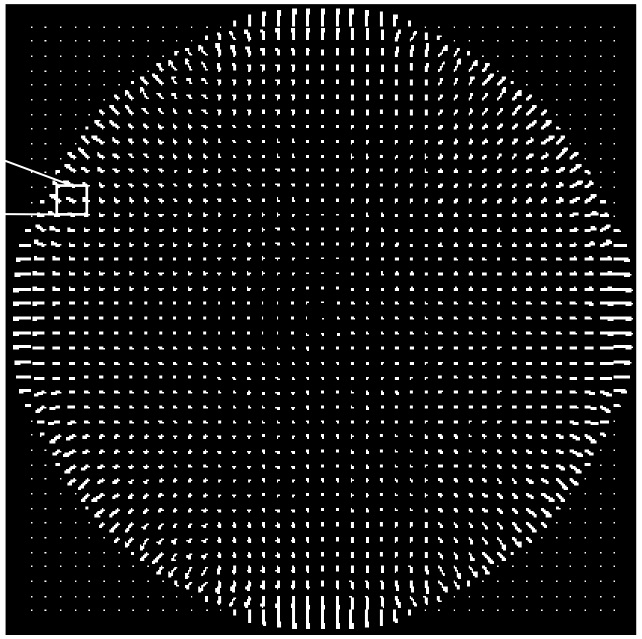

Estimated Surface Normals

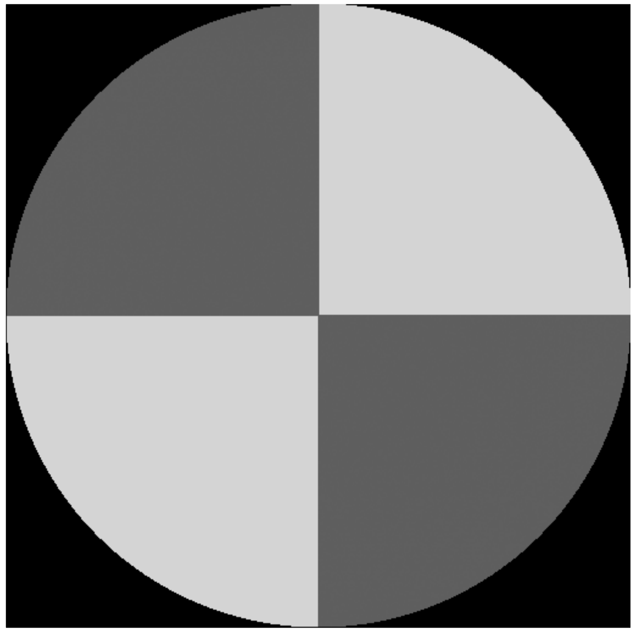

Estimated Albedo

Lambertain Mask

Results – Albedo and Surface Normal

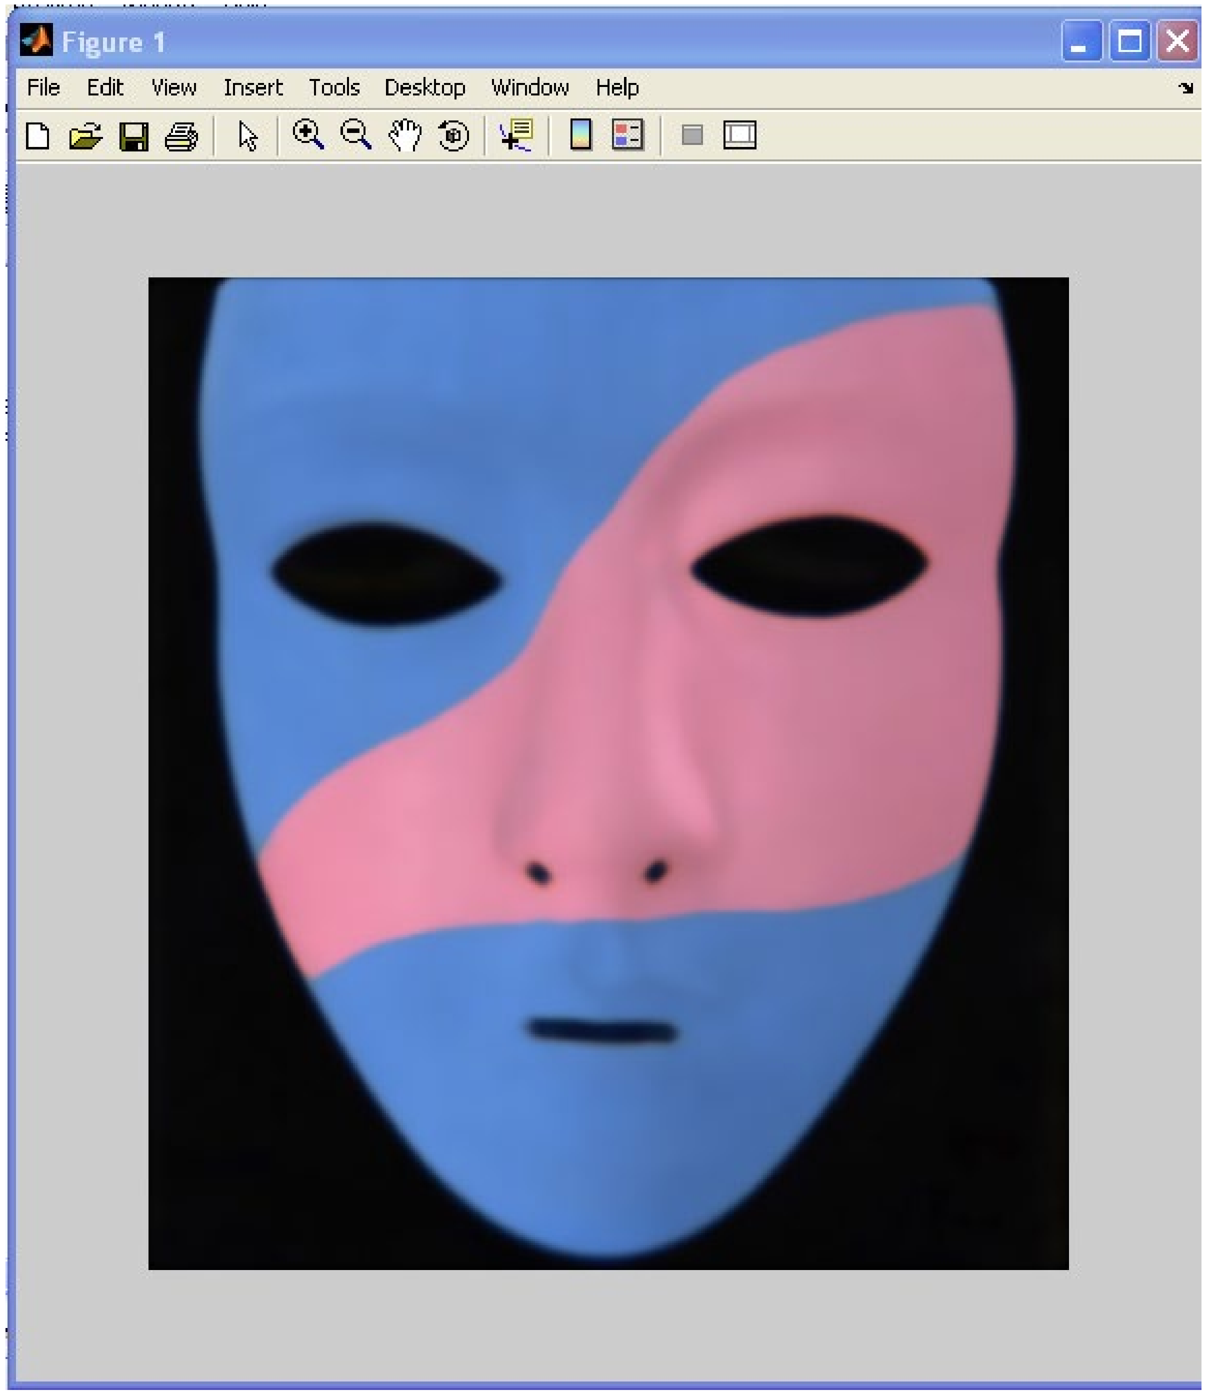

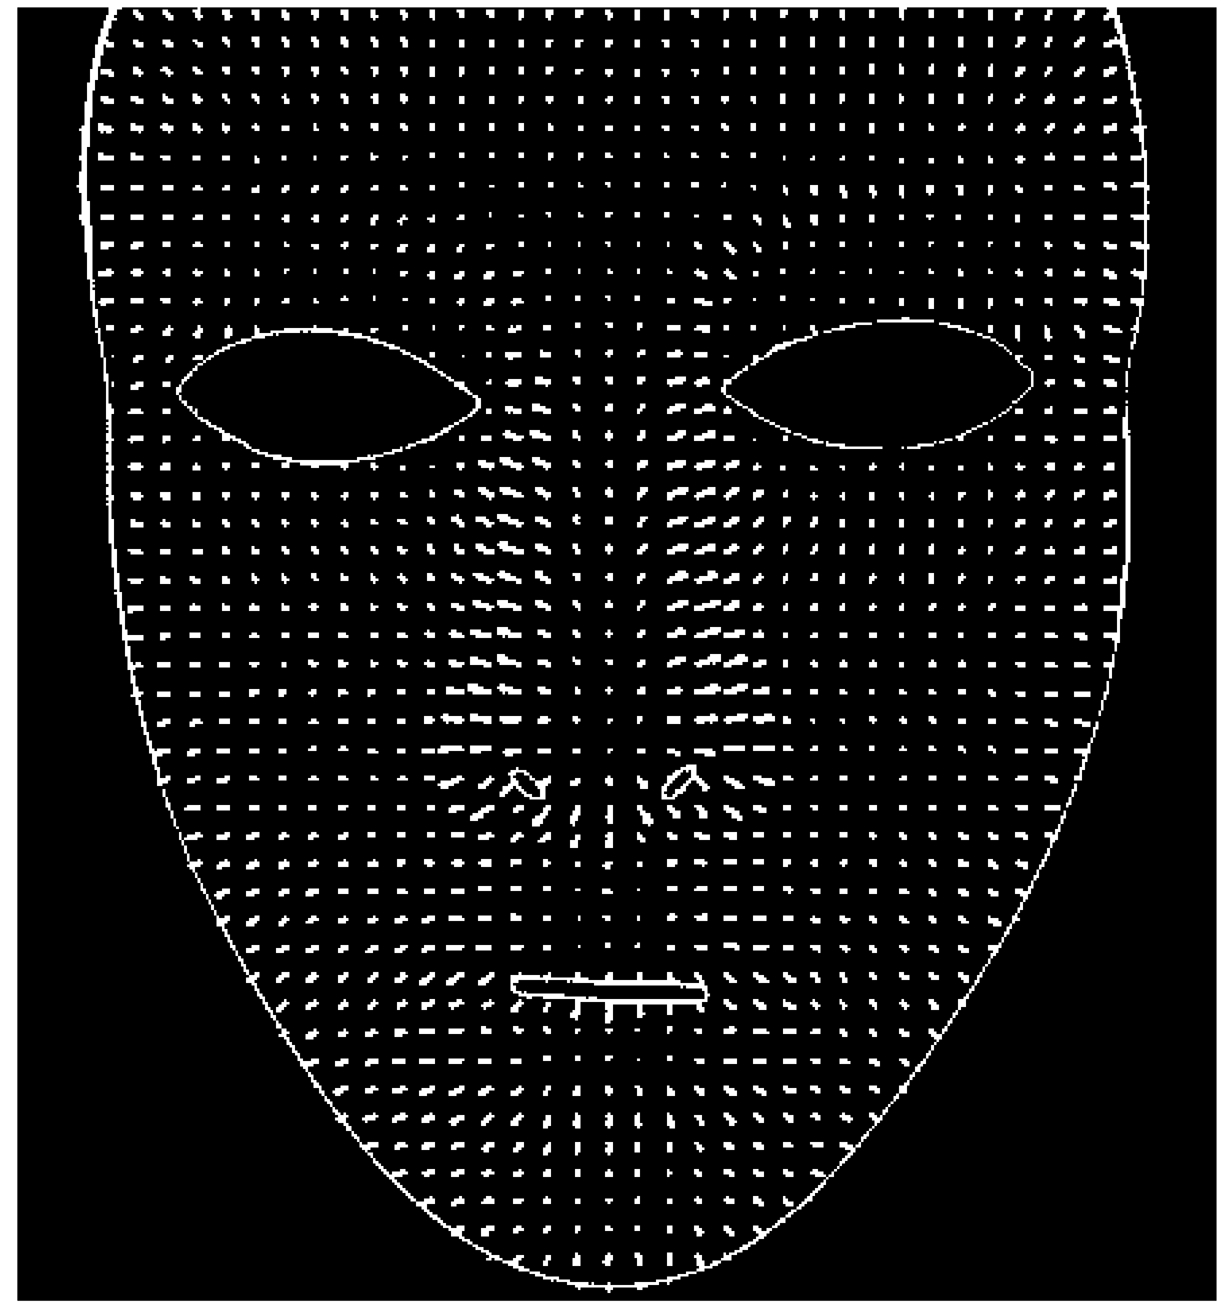

Results – Shape of Mask

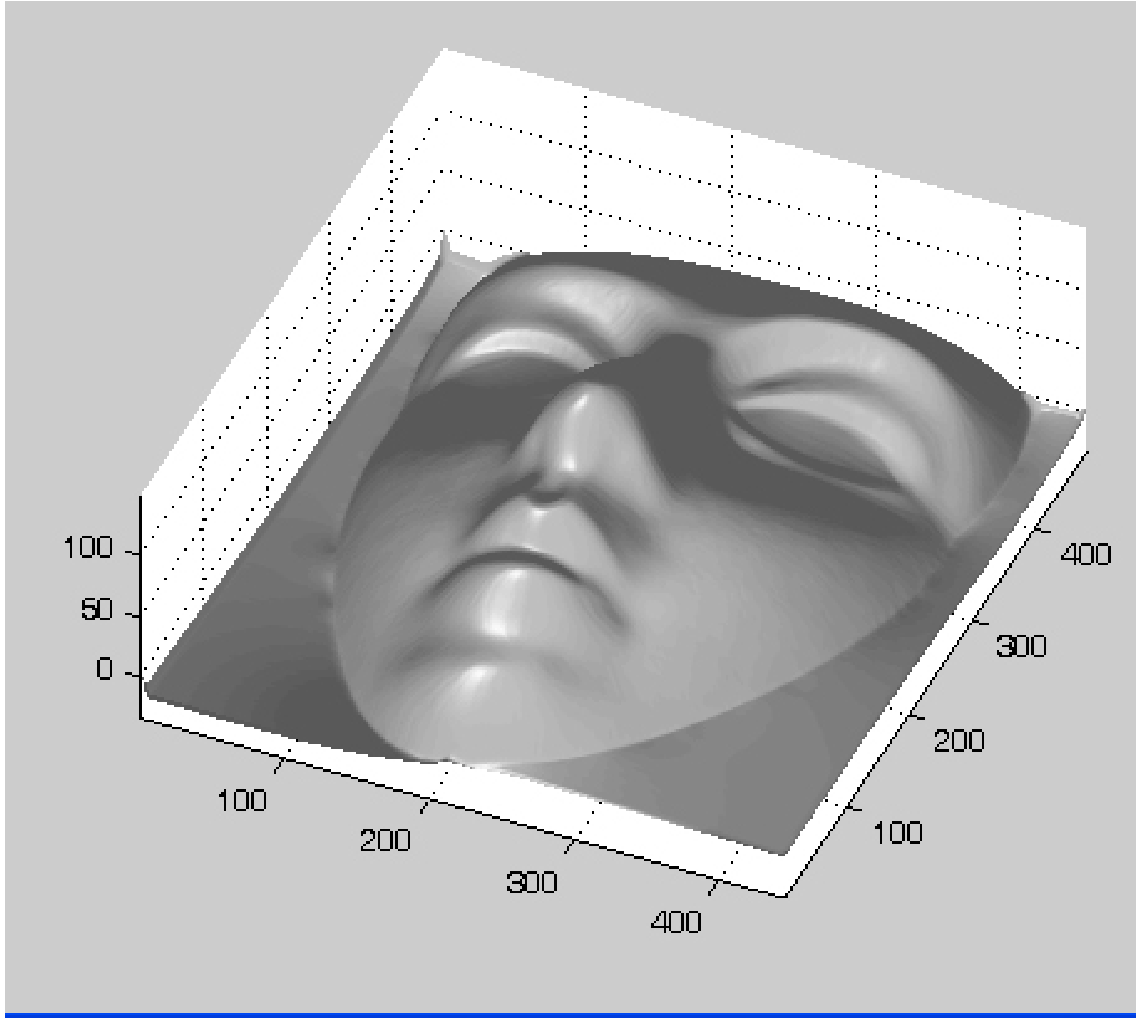

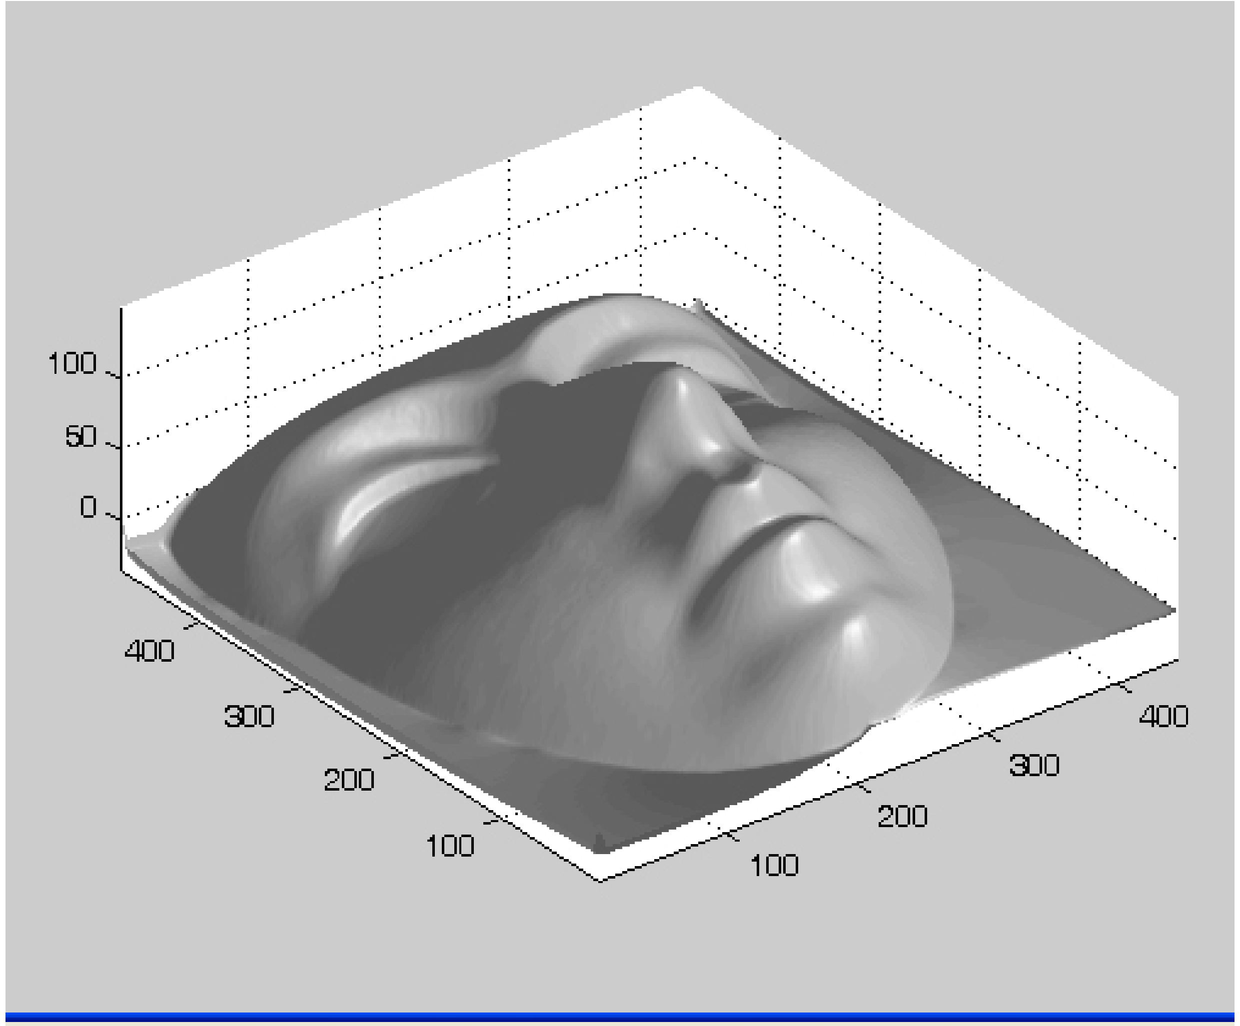

Results: Lambertian Toy





Non-idealities: interreflections

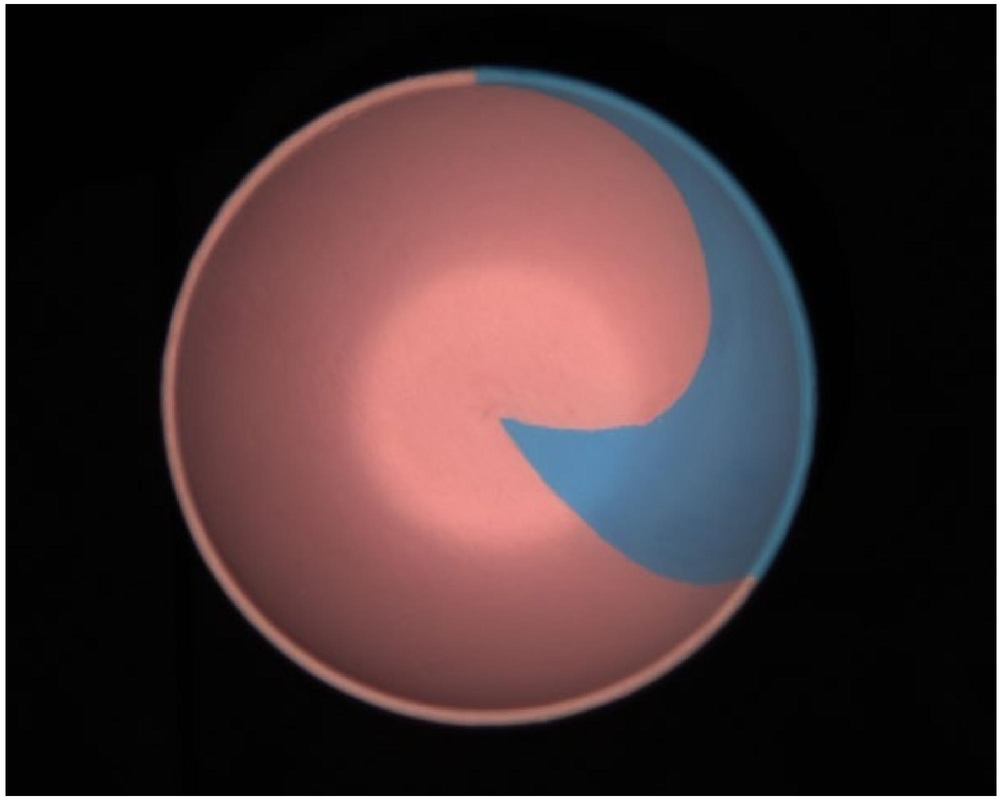

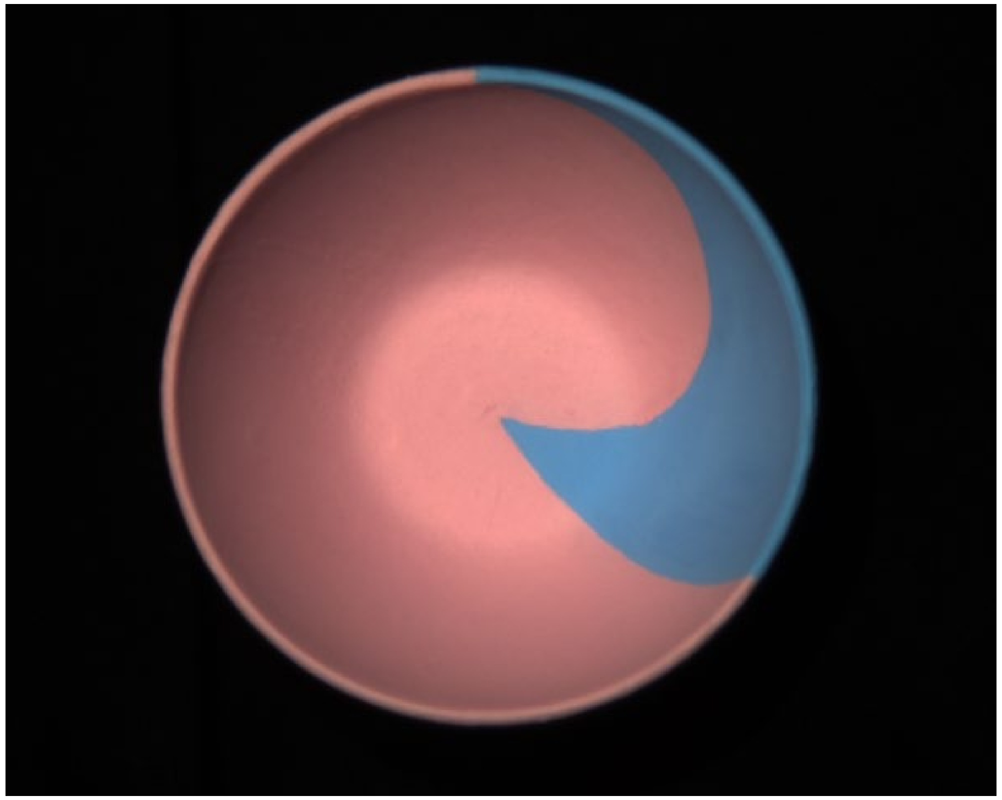

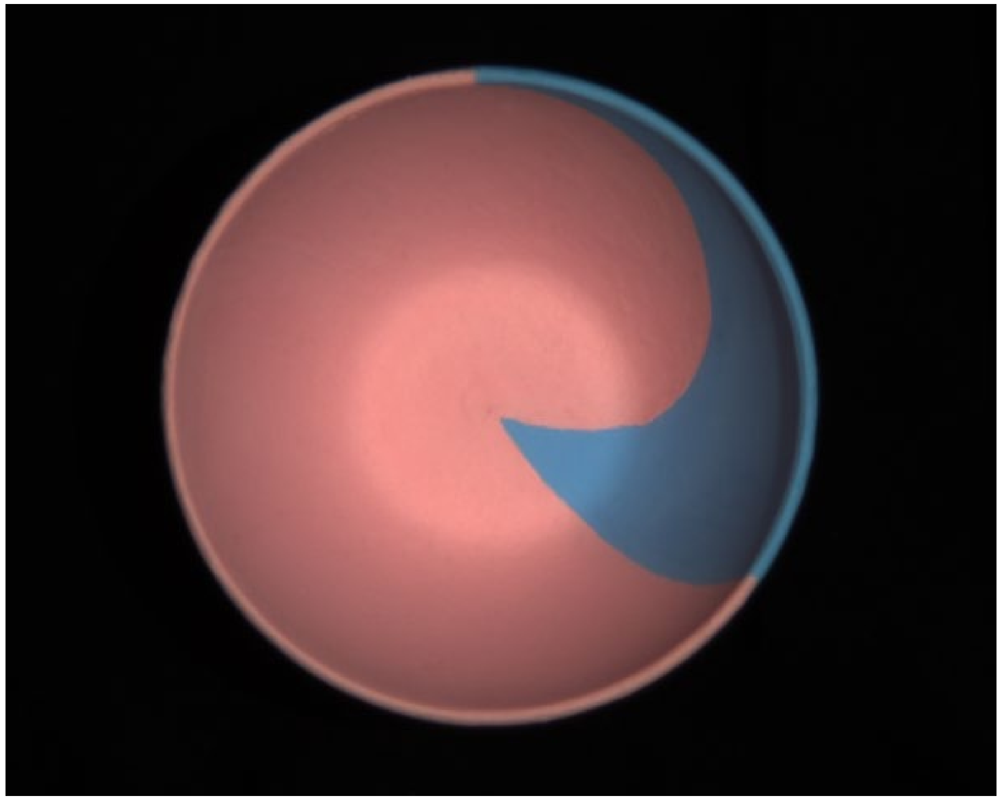

Non-idealities: interreflections

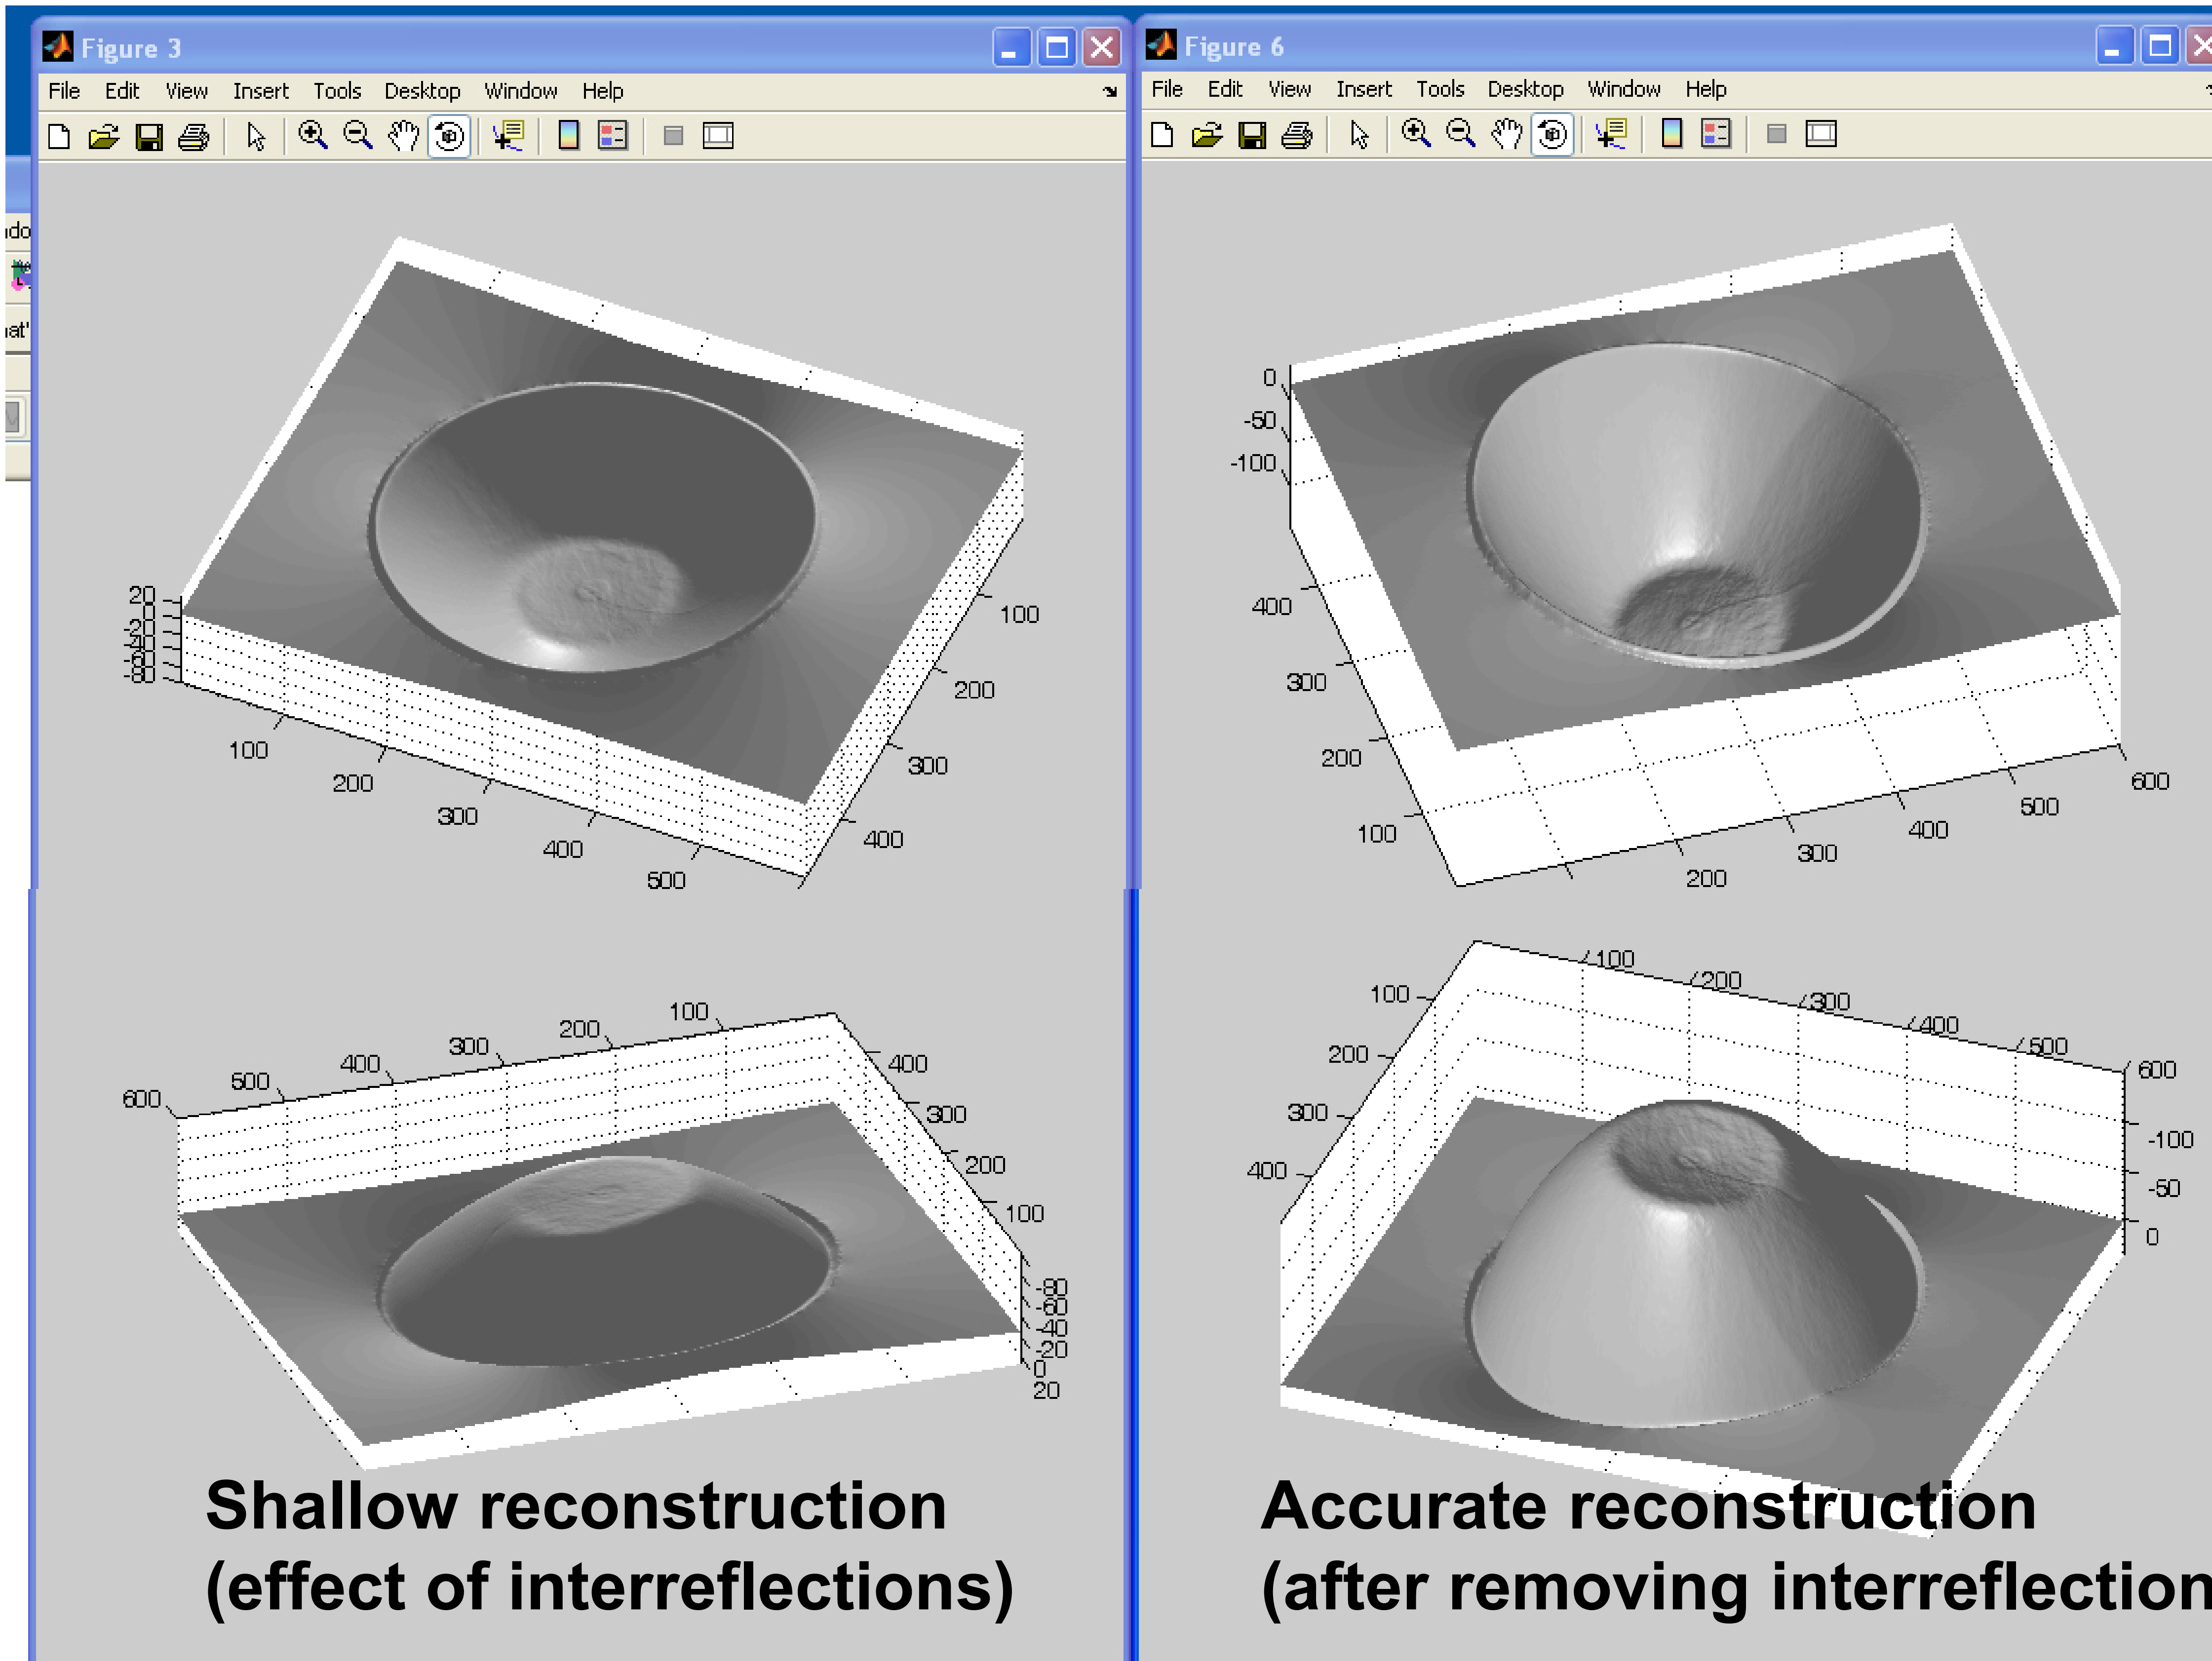

Shallow reconstruction (effect of interreflections)

Accurate reconstruction (after removing interreflections)

What if the light directions are unknown?

Uncalibrated photometric stereo

What if the light directions are unknown?

$$I_1 = a\hat{\boldsymbol{n}}^\top \vec{\boldsymbol{\ell}}_1$$
$$I_2 = a\hat{\boldsymbol{n}}^\top \vec{\boldsymbol{\ell}}_2$$
$$\vdots$$
$$I_N = a\hat{\boldsymbol{n}}^\top \vec{\boldsymbol{\ell}}_N$$

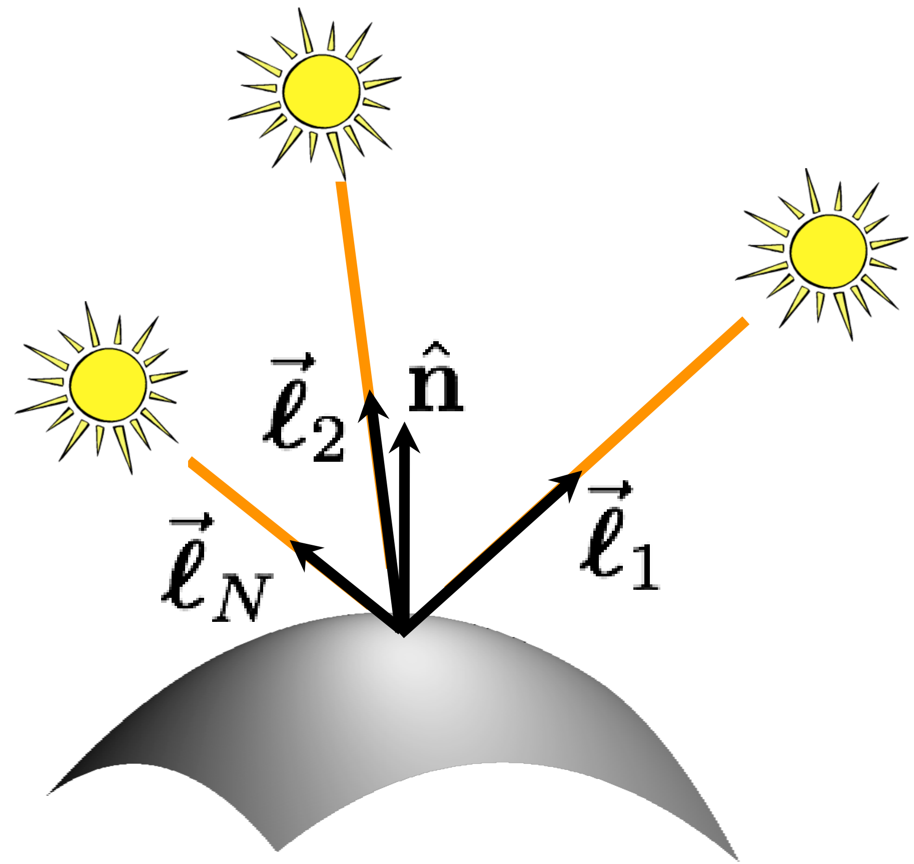

define "pseudo-normal" $\vec{\boldsymbol{b}} \triangleq a\hat{\boldsymbol{n}}$

solve linear system for pseudo-normal

$$\begin{bmatrix} I_1 \\ I_2 \\ \vdots \\ I_N \end{bmatrix}_{N\times 1} = \begin{bmatrix} \vec{\boldsymbol{\ell}}_1^\top \\ \vec{\boldsymbol{\ell}}_2^\top \\ \vdots \\ \vec{\boldsymbol{\ell}}_N^\top \end{bmatrix}_{N\times 3} \begin{bmatrix} \vec{\boldsymbol{b}} \end{bmatrix}_{3\times 1}$$

What if the light directions are unknown?

$$I_1 = a\hat{\boldsymbol{n}}^\top \vec{\boldsymbol{\ell}}_1$$
$$I_2 = a\hat{\boldsymbol{n}}^\top \vec{\boldsymbol{\ell}}_2$$
$$\vdots$$
$$I_N = a\hat{\boldsymbol{n}}^\top \vec{\boldsymbol{\ell}}_N$$

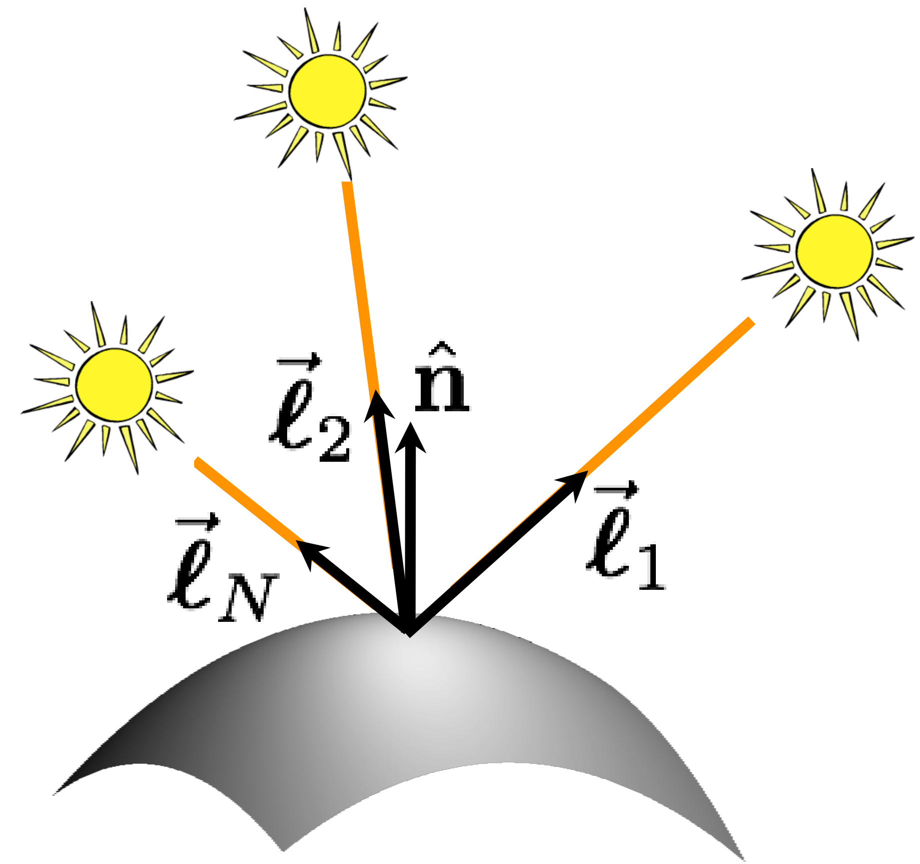

define "pseudo-normal" $\vec{\boldsymbol{b}} \triangleq a\hat{\boldsymbol{n}}$

solve linear system for pseudo-normal at each image pixel

$$\begin{bmatrix} I_1 \\ I_2 \\ \vdots \\ I_N \end{bmatrix}_{N\times M} = \begin{bmatrix} \vec{\boldsymbol{\ell}}_1^\top \\ \vec{\boldsymbol{\ell}}_2^\top \\ \vdots \\ \vec{\boldsymbol{\ell}}_N^\top \end{bmatrix}_{N\times 3} \begin{bmatrix} B \end{bmatrix}_{3\times M}$$

M: number of pixels

What if the light directions are unknown?

$$
\begin{aligned}
I_1 &= a\hat{\boldsymbol{n}}^\top \vec{\boldsymbol{\ell}}_1 \\
I_2 &= a\hat{\boldsymbol{n}}^\top \vec{\boldsymbol{\ell}}_2 \\
&\vdots \\
I_N &= a\hat{\boldsymbol{n}}^\top \vec{\boldsymbol{\ell}}_N
\end{aligned}
$$

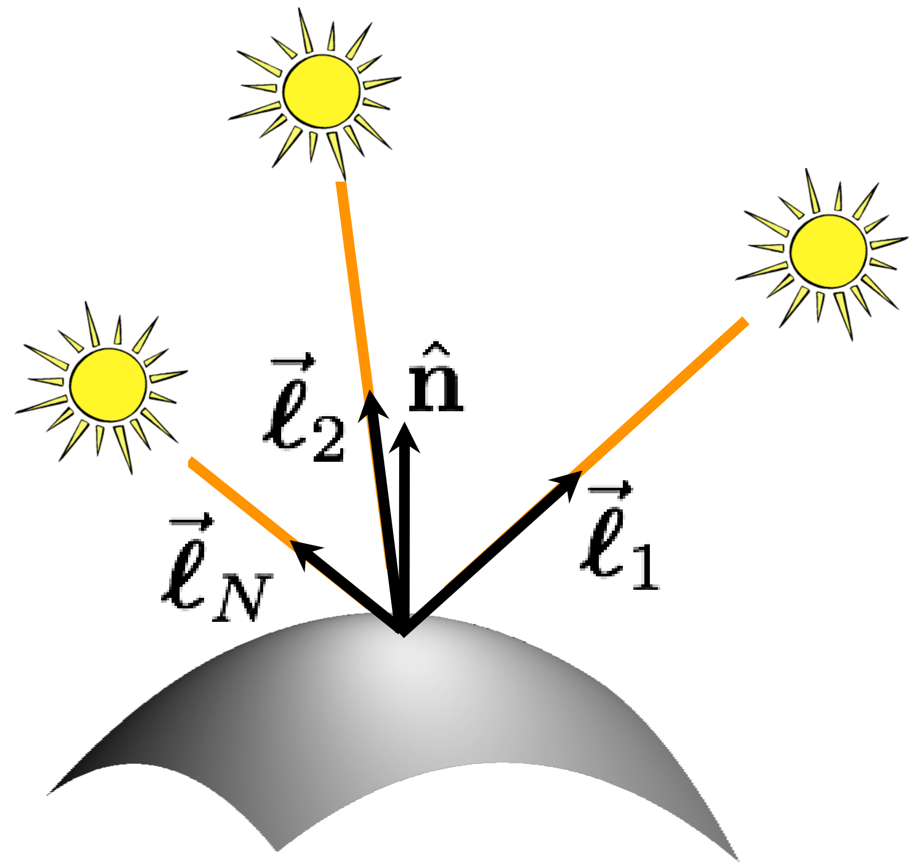

define "pseudo-normal" $\vec{\boldsymbol{b}} \triangleq a\hat{\boldsymbol{n}}$

solve linear system for pseudo-normal at each image pixel

$$
\begin{bmatrix} I_1 \\ I_2 \\ \vdots \\ I_N \end{bmatrix}_{N\times M} = \begin{bmatrix} \vec{\ell}_1^\top \\ \vec{\ell}_2^\top \\ \vdots \\ \vec{\ell}_N^\top \end{bmatrix}_{N\times 3} \begin{bmatrix} B \end{bmatrix}_{3\times M}
$$

How do we solve this system without knowing light matrix L?

Factorizing the measurement matrix

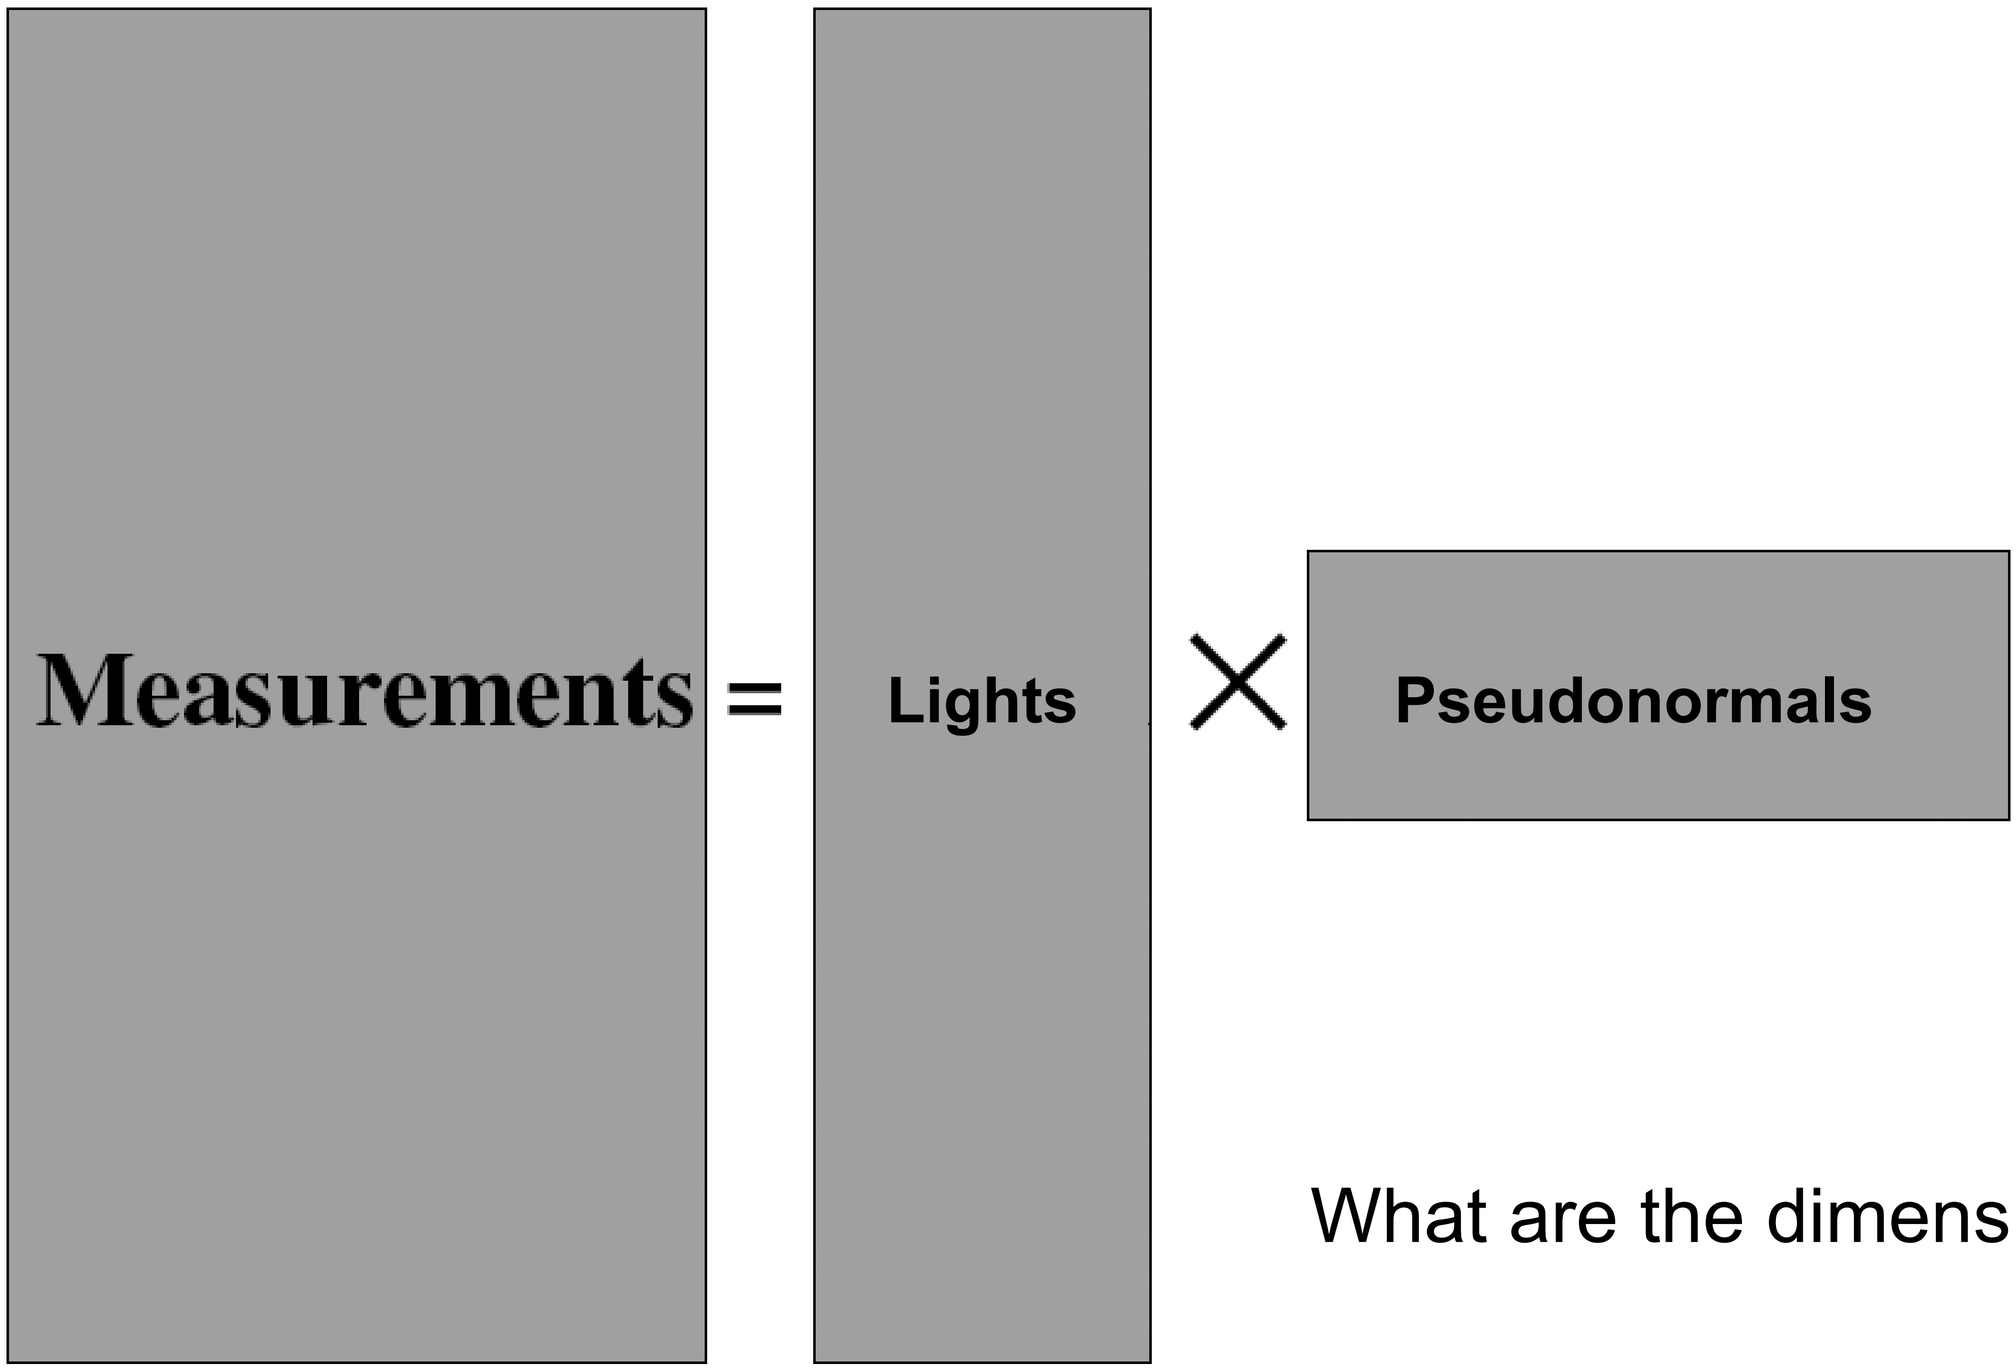

What are the dimensions?

Factorizing the measurement matrix

- Singular value decomposition:

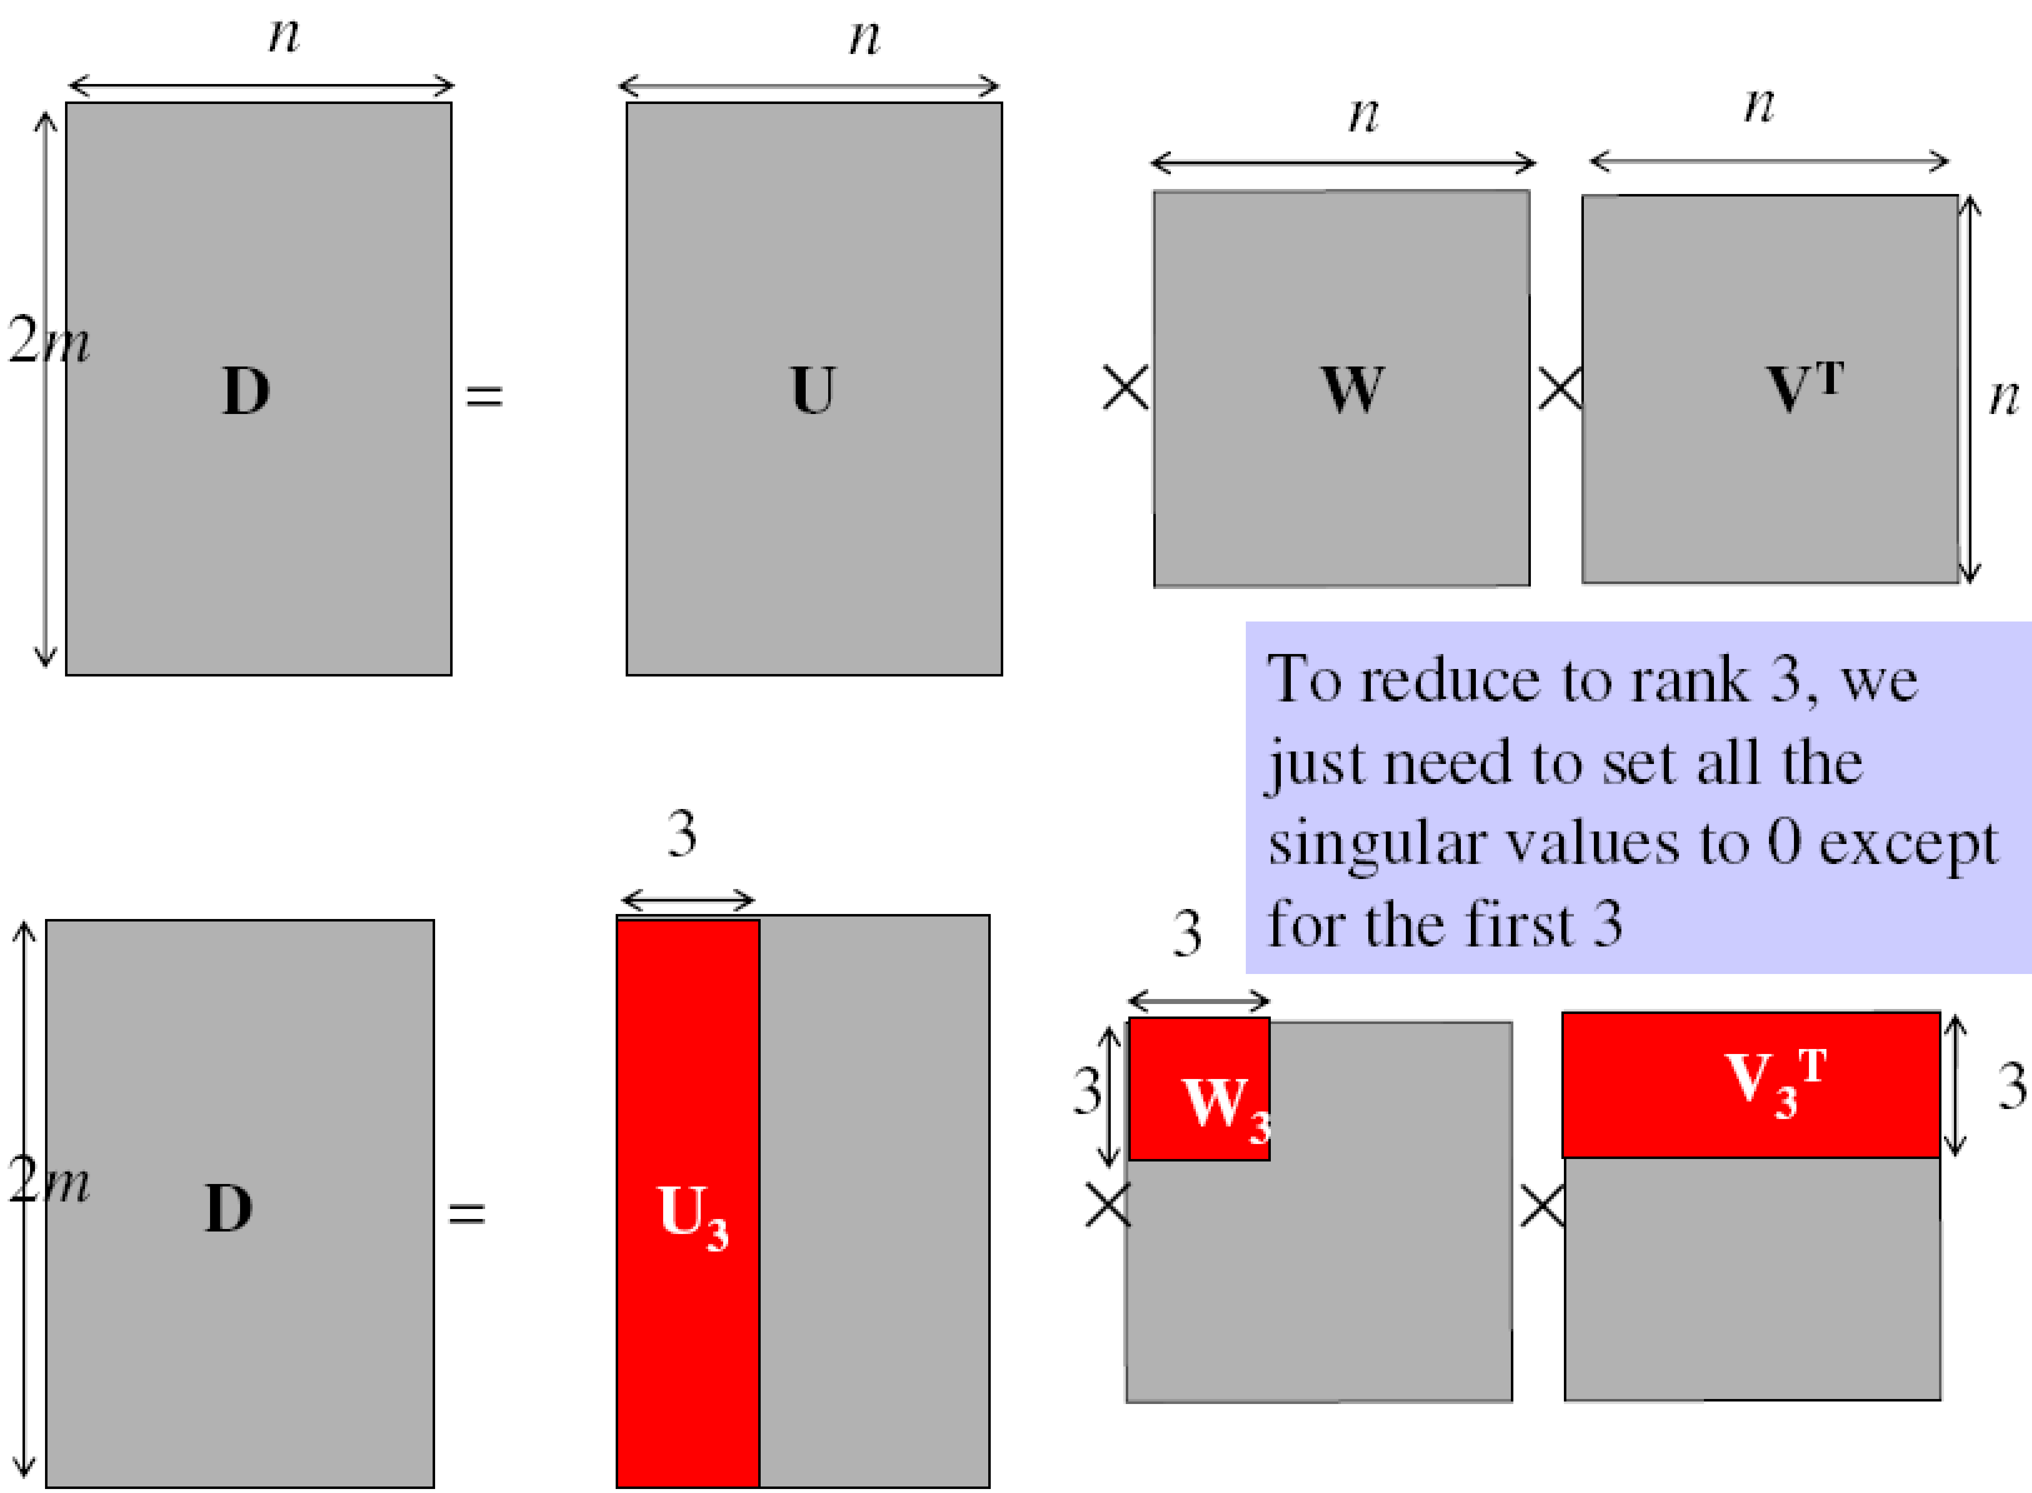

To reduce to rank 3, we just need to set all the singular values to 0 except for the first 3

This decomposition minimizes **$|I\text{-}LB|^2$**

Are the results unique?

Are the results unique?

We can insert any 3x3 matrix Q in the decomposition and get the same images:

$$\mathbf{I} = \mathbf{L}\,\mathbf{B} = (\mathbf{L}\,\mathbf{Q}^{-1})\,(\mathbf{Q}\,\mathbf{B})$$

Are the results unique?

We can insert any 3x3 matrix Q in the decomposition and get the same images:

$$\mathbf{I} = \mathbf{L}\,\mathbf{B} = (\mathbf{L}\,\mathbf{Q}^{-1})\,(\mathbf{Q}\,\mathbf{B})$$

Can we use any assumptions to remove some of these 9 degrees of freedom?

Generalized bas-relief ambiguity

Enforcing integrability

What does the matrix **B** correspond to?

Enforcing integrability

What does the matrix **B** correspond to?

- Surface representation as a depth image (also known as Monge surface):

$$\underset{\substack{\uparrow \\ \text{depth at each pixel}}}{z} = f(\underbrace{x, y}_{\text{pixel coordinates in image space}})$$

- Unnormalized normal:

$$\tilde{n}(x, y) = \left(\frac{df}{dx}, \frac{df}{dy}, -1\right)$$

- Actual normal:

$$n(x, y) = \tilde{n}(x, y) / \|\tilde{n}(x, y)\|$$

- Pseudo-normal:

$$b(x, y) = a(x, y) n(x, y)$$

- Rearrange into 3xN matrix **B**.

Enforcing integrability

What does the integrability constraint correspond to?

Enforcing integrability

What does the integrability constraint correspond to?

- Differentiation order should not matter:

$$\frac{d}{dy}\frac{df(x,y)}{dx} = \frac{d}{dx}\frac{df(x,y)}{dy}$$

- Can you think of a way to express the above using pseudo-normals **b**?

Enforcing integrability

What does the integrability constraint correspond to?

- Differentiation order should not matter:

$$\frac{d}{dy}\frac{df(x,y)}{dx} = \frac{d}{dx}\frac{df(x,y)}{dy}$$

- Can you think of a way to express the above using pseudo-normals **b**?

$$\frac{d}{dy}\frac{b_1(x,y)}{b_3(x,y)} = \frac{d}{dx}\frac{b_2(x,y)}{b_3(x,y)}$$

Enforcing integrability

What does the integrability constraint correspond to?

- Differentiation order should not matter:

$$\frac{d}{dy}\frac{df(x,y)}{dx} = \frac{d}{dx}\frac{df(x,y)}{dy}$$

- Can you think of a way to express the above using pseudo-normals **b**?

$$\frac{d}{dy}\frac{b_1(x,y)}{b_3(x,y)} = \frac{d}{dx}\frac{b_2(x,y)}{b_3(x,y)}$$

- Simplify to:

$$b_3(x,y)\frac{db_1(x,y)}{dy} - b_1(x,y)\frac{db_3(x,y)}{dy} = b_2(x,y)\frac{db_1(x,y)}{dx} - b_1(x,y)\frac{db_2(x,y)}{dx}$$

Enforcing integrability

What does the integrability constraint correspond to?

- Differentiation order should not matter:

$$\frac{d}{dy}\frac{df(x,y)}{dx} = \frac{d}{dx}\frac{df(x,y)}{dy}$$

- Can you think of a way to express the above using pseudo-normals **b**?

$$\frac{d}{dy}\frac{b_1(x,y)}{b_3(x,y)} = \frac{d}{dx}\frac{b_2(x,y)}{b_3(x,y)}$$

- Simplify to:

$$b_3(x,y)\frac{db_1(x,y)}{dy} - b_1(x,y)\frac{db_3(x,y)}{dy} = b_2(x,y)\frac{db_1(x,y)}{dx} - b_1(x,y)\frac{db_2(x,y)}{dx}$$

- If $\mathbf{B_e}$ is the pseudo-normal matrix we get from SVD, then find the 3x3 transform **D** such that $\mathbf{B=D \cdot B_e}$ is the closest to satisfying integrability in the least-squares sense.

Enforcing integrability

Does enforcing integrability remove all ambiguities?

Generalized Bas-relief ambiguity

If **B** is integrable, then:

- **B'=G^{-T}·B** is also integrable for all **G** of the form ($\lambda \neq 0$)

$$G = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ \mu & \nu & \lambda \end{bmatrix}$$

- Combined with transformed lights **S'=G·S**, the transformed pseudonormals produce the same images as the original pseudonormals.

- This ambiguity cannot be removed using shadows.

- This ambiguity *can* be removed using interreflections or additional assumptions.

This ambiguity is known as the generalized bas-relief ambiguity.

Generalized Bas-relief ambiguity

When μ = ν = 0, **G** is equivalent to the transformation employed by relief sculptures.

When μ = ν = 0 and $\lambda = +\text{-}1$, top/down ambiguity.

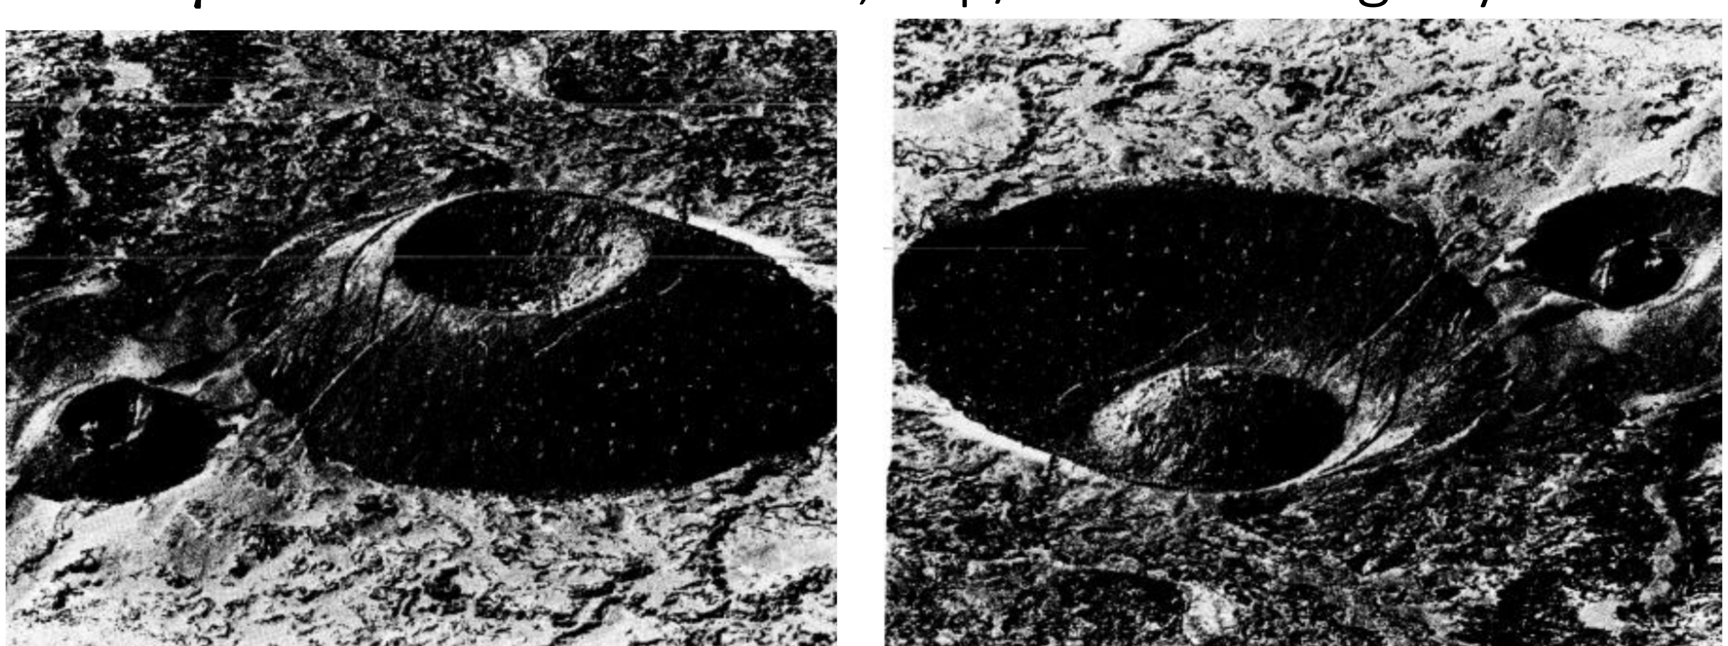

Otherwise, includes shearing.

What assumptions have we made for all this?

What assumptions have we made for all this?

- Lambertian BRDF
- Directional lighting
- No interreflections or scattering

Shape independent of BRDF via reciprocity: "Helmholtz Stereopsis"

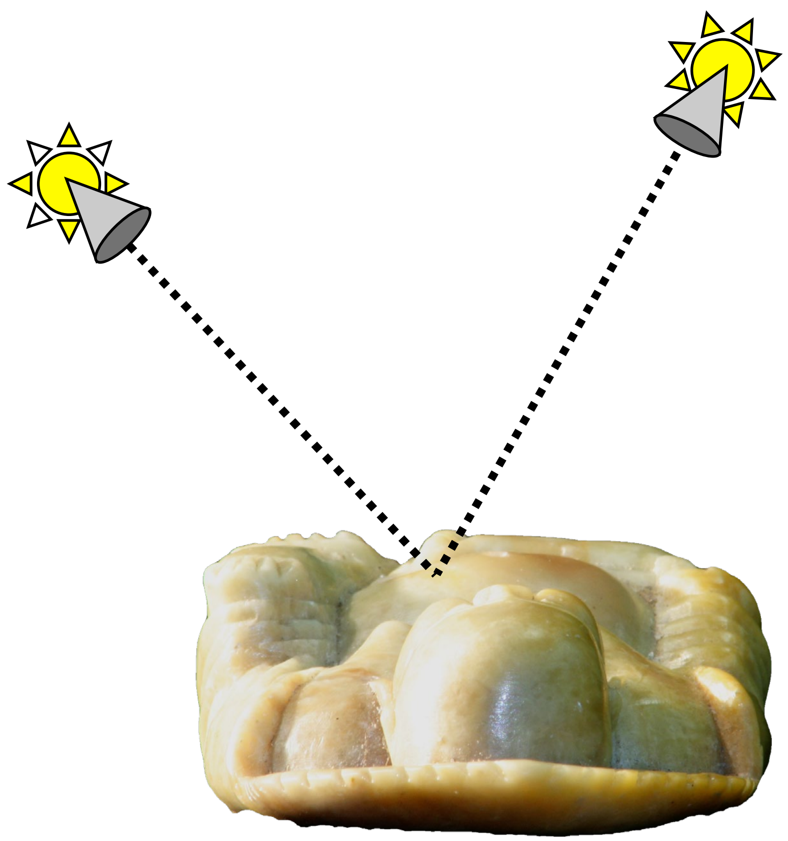

$$I = f(\text{shape}, \cancel{\text{illumination}}, \cancel{\text{reflectance}})$$

$$f^{-1} =$$

[Zickler et al., 2002]

References

Basic reading:

- Szeliski, Section 2.2.
- Gortler, Chapter 21.
  This book by Steven Gortler has a great *introduction* to radiometry, reflectance, and their use for image formation.

Additional reading:

- Oren and Nayar, “Generalization of the Lambertian model and implications for machine vision,” IJCV 1995.
  The paper introducing the most common model for rough diffuse reflectance.
- Debevec, “Rendering Synthetic Objects into Real Scenes,” SIGGRAPH 1998.
  The paper that introduced the notion of the environment map, the use of chrome spheres for measuring such maps, and the idea that they can be used for easy rendering.
- Lalonde et al., “Estimating the Natural Illumination Conditions from a Single Outdoor Image,” IJCV 2012.
  A paper on estimating outdoors environment maps from just one image.
- Basri and Jacobs, “Lambertian reflectance and linear subspaces,” ICCV 2001.
- Ramamoorthi and Hanrahan, “A signal-processing framework for inverse rendering,” SIGGRAPH 2001.
- Sloan et al., “Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments,” SIGGRAPH 2002.
  Three papers describing the use of spherical harmonics to model low-frequency illumination, as well as the low-pass filtering effect of Lambertian reflectance on illumination.
- Zhang et al., “Shape-from-shading: a survey,” PAMI 1999.
  A review of perceptual and computational aspects of shape from shading.
- Woodham, “Photometric method for determining surface orientation from multiple images,” Optical Engineering 1980.
  The paper that introduced photometric stereo.
- Yuille and Snow, “Shape and albedo from multiple images using integrability,” CVPR 1997.
- Belhumeur et al., “The bas-relief ambiguity,” IJCV 1999.
- Papadhimitri and Favaro, “A new perspective on uncalibrated photometric stereo,” CVPR 2013.
  Three papers discussing uncalibrated photometric stereo. The first paper shows that, when the lighting directions are not known, by assuming integrability, one can reduce unknowns to the bas-relief ambiguity. The second paper discusses the bas-relief ambiguity in a more general context. The third paper shows that, if instead of an orthographic camera one uses a perspective camera, this is further reduced to just a scale ambiguity.
- Alldrin et al., “Resolving the generalized bas-relief ambiguity by entropy minimization,” CVPR 2007.
  A popular technique for resolving the bas-relief ambiguity in uncalibrated photometric stereo.
- Zickler et al., “Helmholtz stereopsis: Exploiting reciprocity for surface reconstruction,” IJCV 2002.
  A method for photometric stereo reconstruction under arbitrary BRDF.
- Nayar et al., “Shape from interreflections,” IJCV 1991.
- Chandraker et al., “Reflections on the generalized bas-relief ambiguity,” CVPR 2005.
  Two papers discussing how one can perform photometric stereo (calibrated or otherwise) in the presence of strong interreflections.
- Frankot and Chellappa, “A method for enforcing integrability in shape from shading algorithms,” PAMI 1988.
- Agrawal et al., “What is the range of surface reconstructions from a gradient field?,” ECCV 2006.
  Two papers discussing how one can integrate a normal field to reconstruct a surface.