Light, Camera and Shading

CS 543 / ECE 549 – Saurabh Gupta

Spring 2020, UIUC

http://saurabhg.web.illinois.edu/teaching/ece549/sp2020/

Many slides adapted from S. Seitz, L. Lazebnik, D. Hoiem, D. Forsyth

Recap

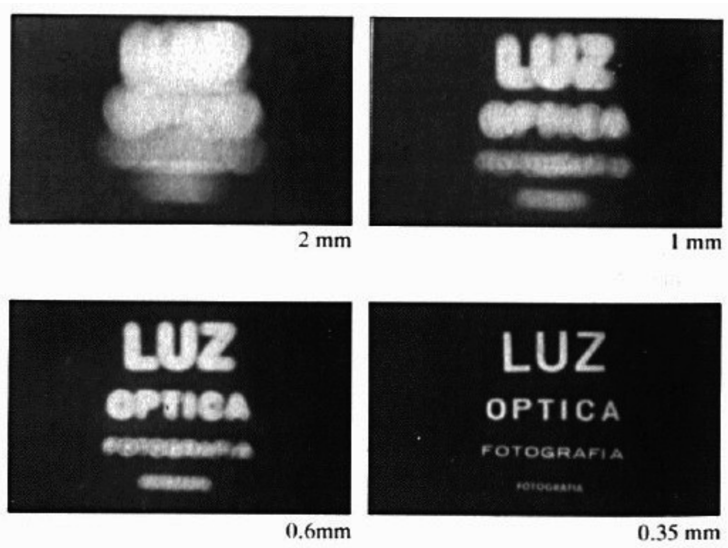

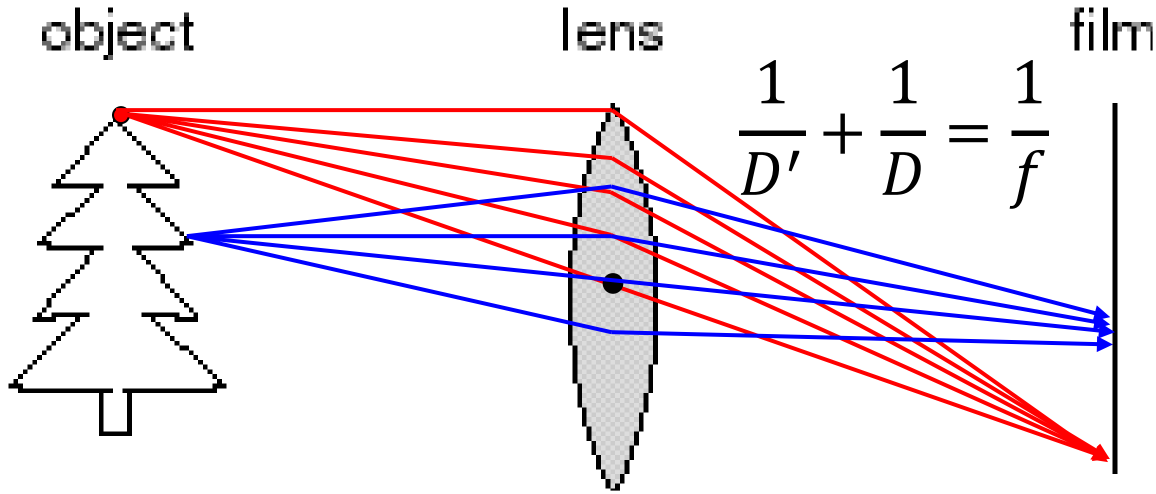

$$\frac{1}{D'} + \frac{1}{D} = \frac{1}{f}$$

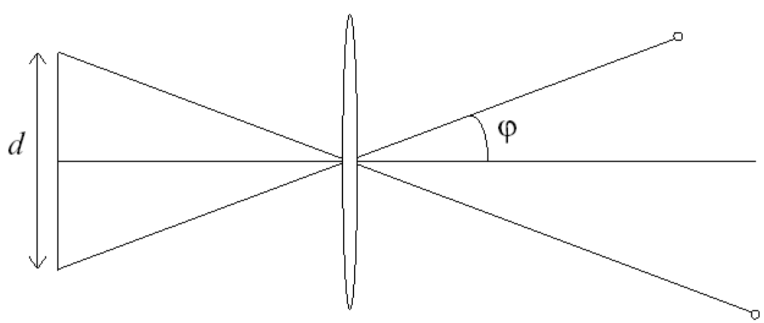

$$\varphi = tan^{-1}\left(\frac{d}{2f}\right)$$

Recap

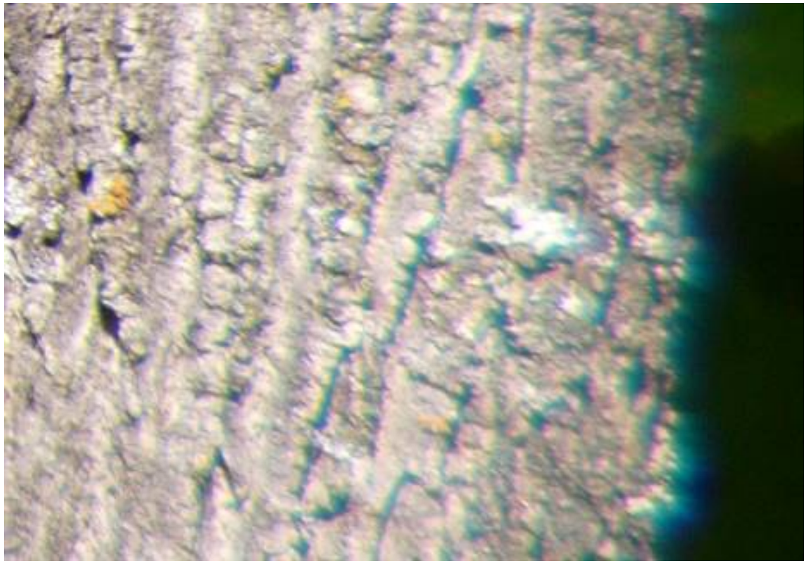

Recap

$$E = \left[\frac{\pi}{4} \left(\frac{d}{f} \right)^2 \cos^4 \alpha \right] L$$

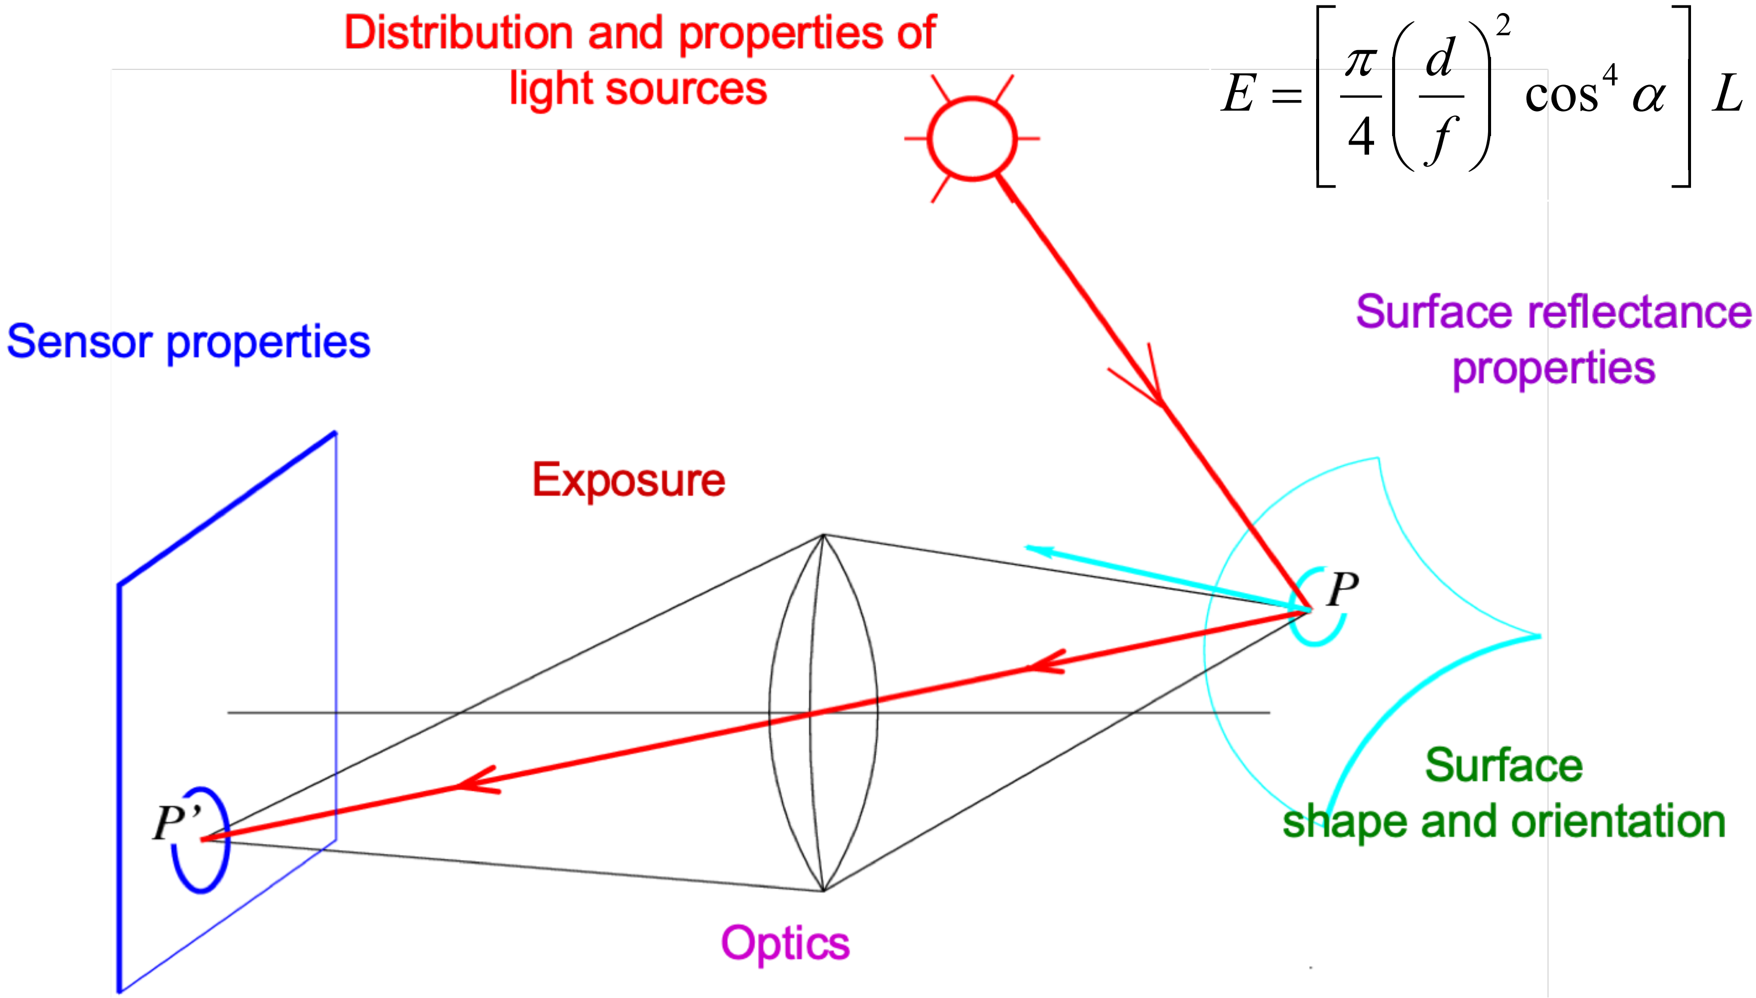

Recap

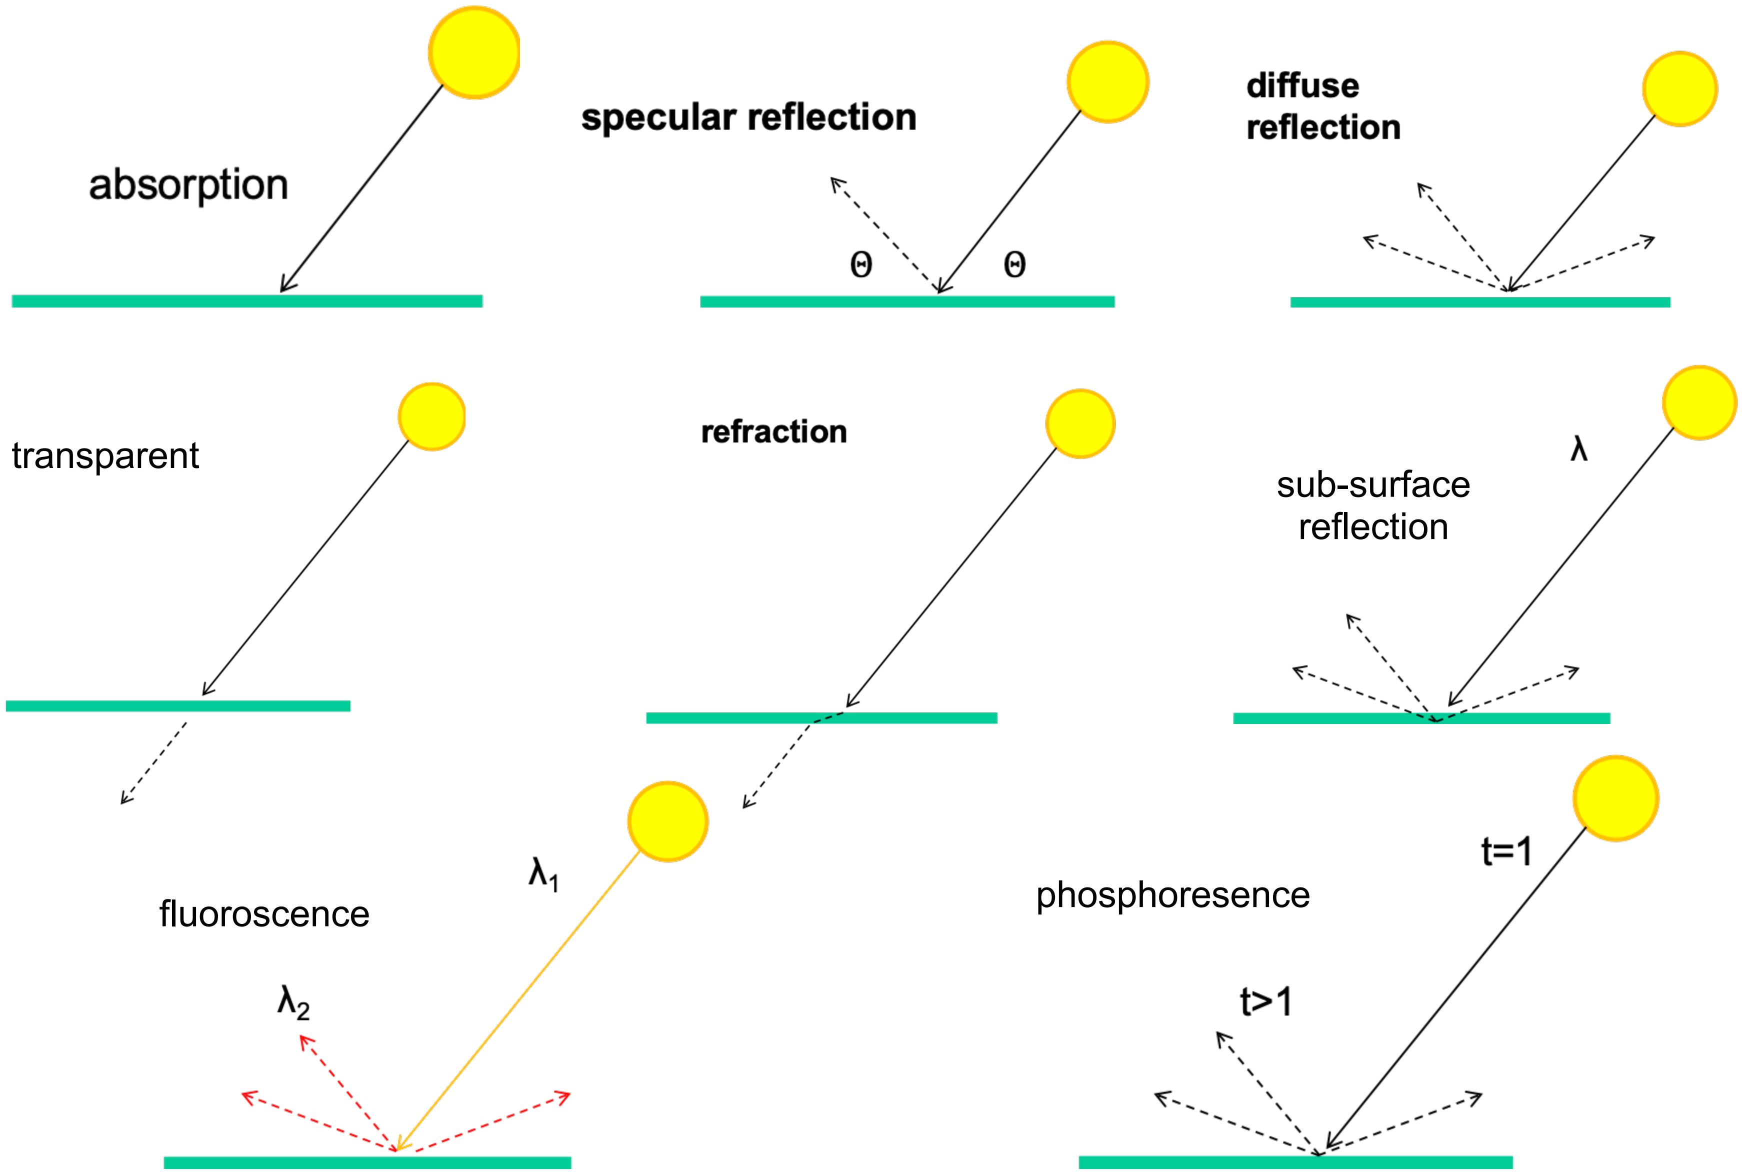
absorption
specular reflection
Θ
Θ
diffuse
reflection
transparent
refraction
sub-surface
reflection
λ
fluoroscence
λ1
λ2
phosphoresence
t=1
t>1

Overview

- Cameras with lenses
  - Depth of field
  - Field of view
  - Lens aberrations
- Brightness of a pixel
  - Small taste of radiometry
  - In-camera transformation of light
  - Reflectance properties of surfaces
  - Lambertian reflection model
  - Shape from shading
- Color

Most surfaces have both

Specularity = spot where specular reflection dominates (typically reflects light source)

Photo: northcountryhardwoodfloors.com

Typically, specular component is small

Slide from D. Hoiem

Specular reflection

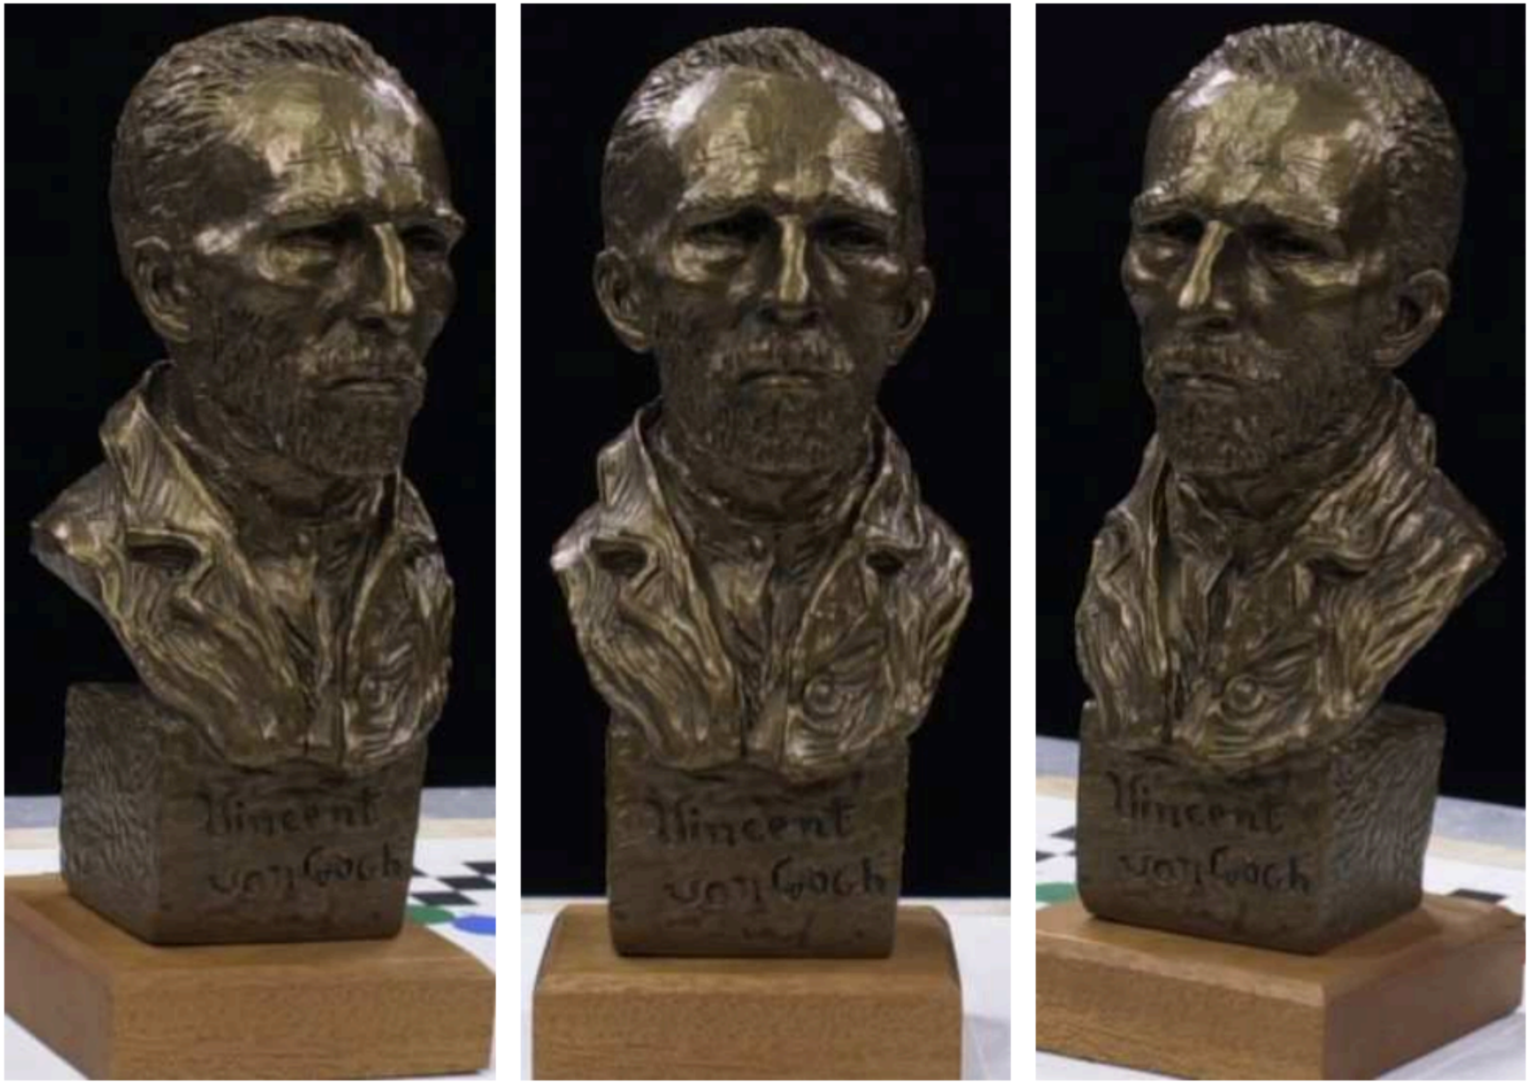

Picture source

Slide from L. Lazebnik

When light hits a typical surface

- Some light is absorbed

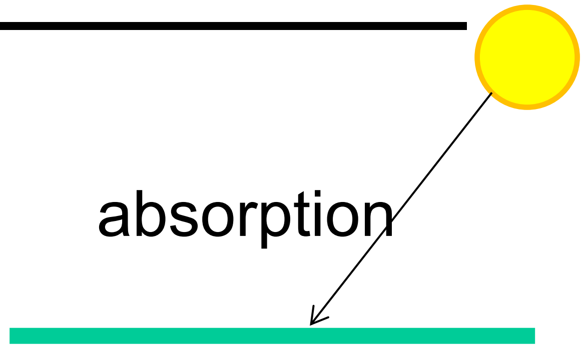

- Some light is reflected diffusely
  - Independent of viewing direction

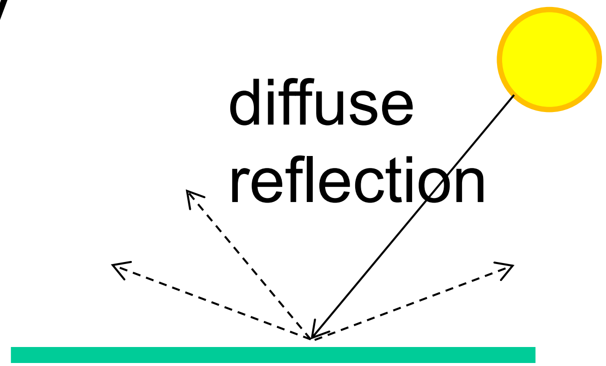

- Some light is reflected specularly
  - Light bounces off (like a mirror), depends on viewing direction

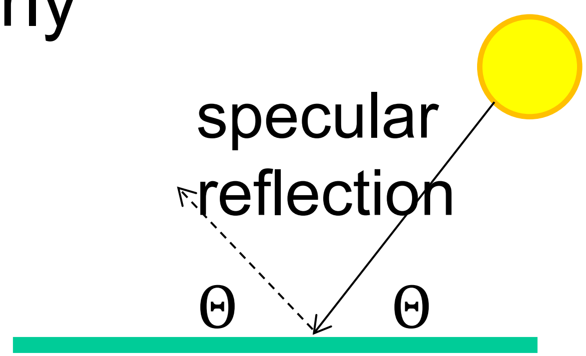

Slide from D. Hoiem

Bidirectional Reflectance Distribution Function (BRDF)

- How bright a surface appears when viewed from one direction when light falls on it from another
- Definition: ratio of the radiance in the emitted direction to irradiance in the incident direction

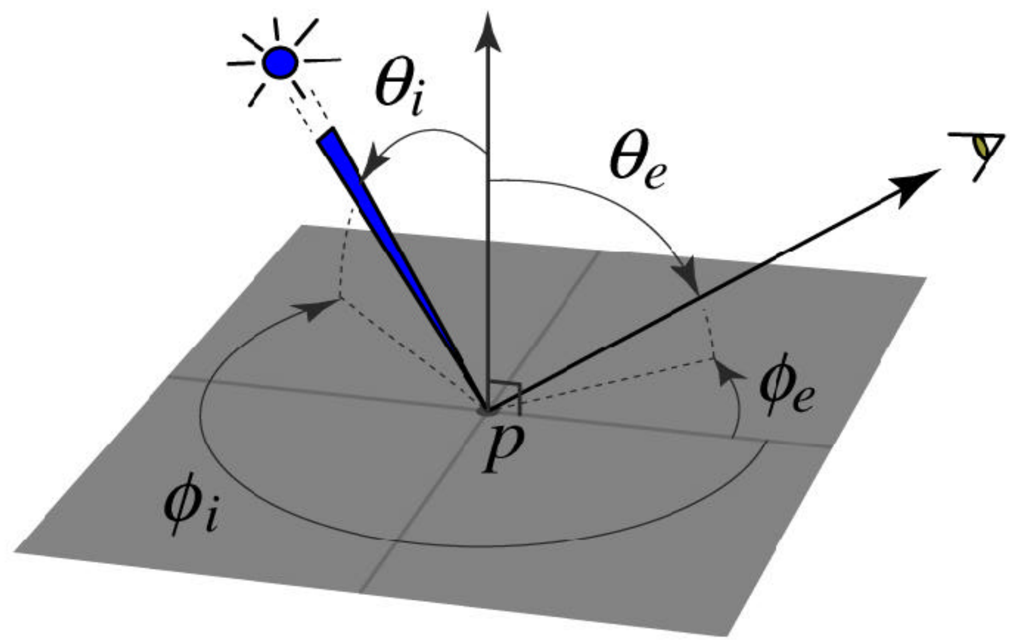

Source: Steve Seitz

Lambertian reflectance model

Some light is absorbed (function of albedo ρ)

Remaining light is scattered, *equally* in all directions.

Examples: soft cloth, concrete, matte paints

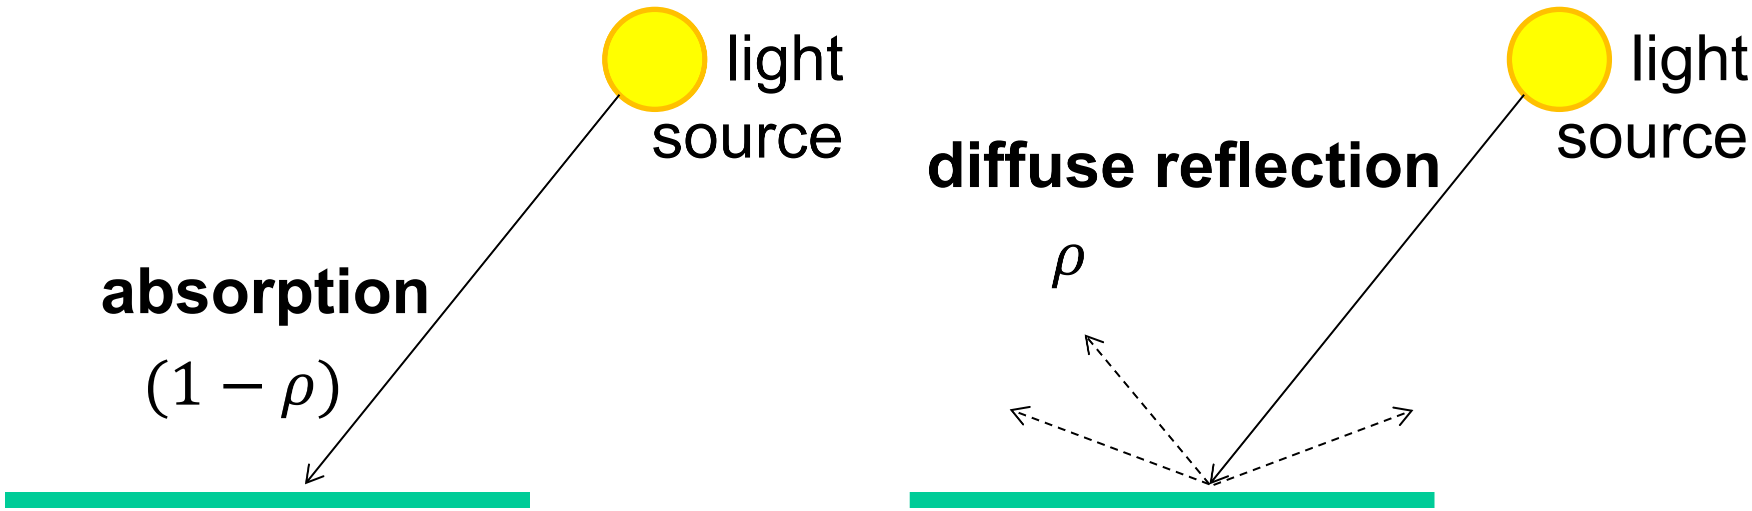

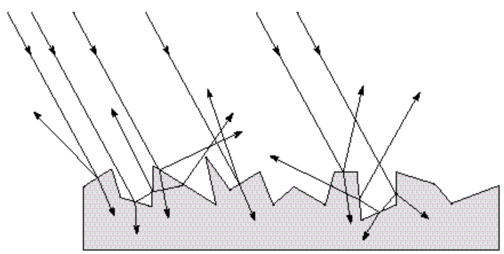

Slide from D. Hoiem

Intensity and Surface Orientation

Intensity depends on illumination angle because less light comes in at oblique angles.

ρ = albedo

$\boldsymbol{S}$ = directional source

$\boldsymbol{N}$ = surface normal

I = reflected intensity

$$I(x) = \rho(x)(\boldsymbol{S} \cdot \boldsymbol{N}(x))$$

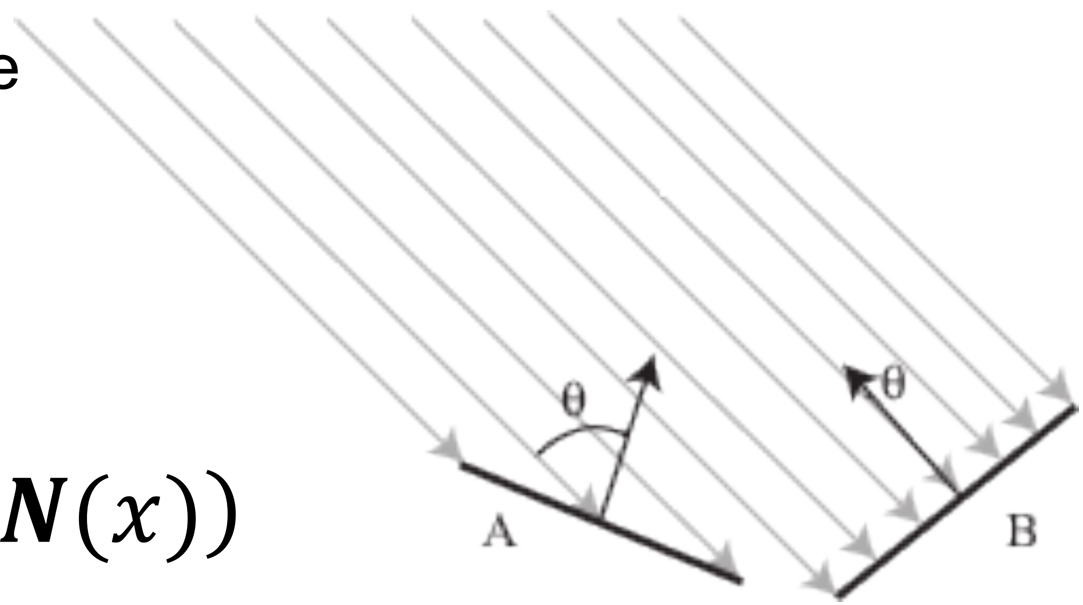

Slide: Forsyth

Photometric stereo (shape from shading)

- Can we reconstruct the shape of an object based on shading cues?

Slide from L. Lazebnik

Photometric stereo

Assume:

- A Lambertian object
- A *local shading model* (each point on a surface receives light only from sources visible at that point)
- A set of *known* light source directions
- A set of pictures of an object, obtained in exactly the same camera/object configuration but using different sources
- Orthographic projection

Goal: reconstruct object shape and albedo

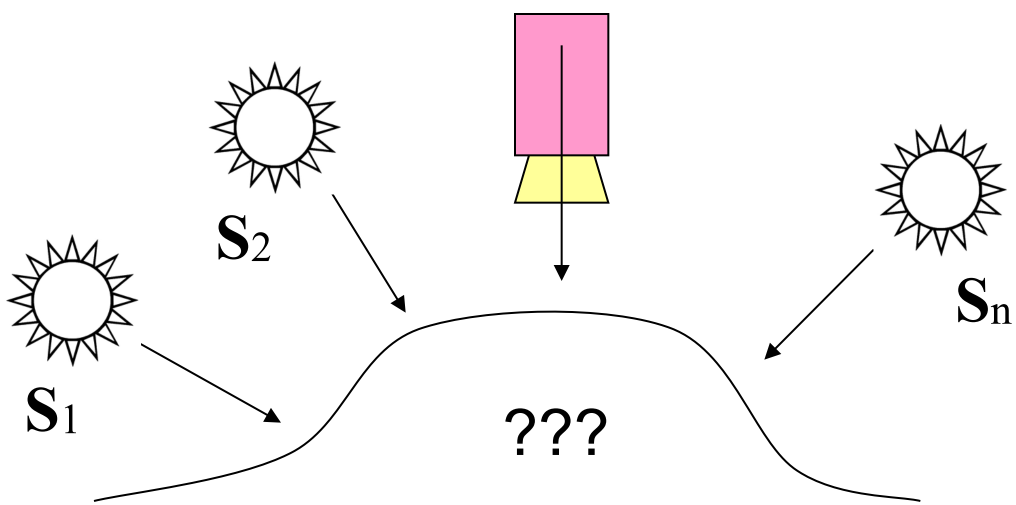

Example 1

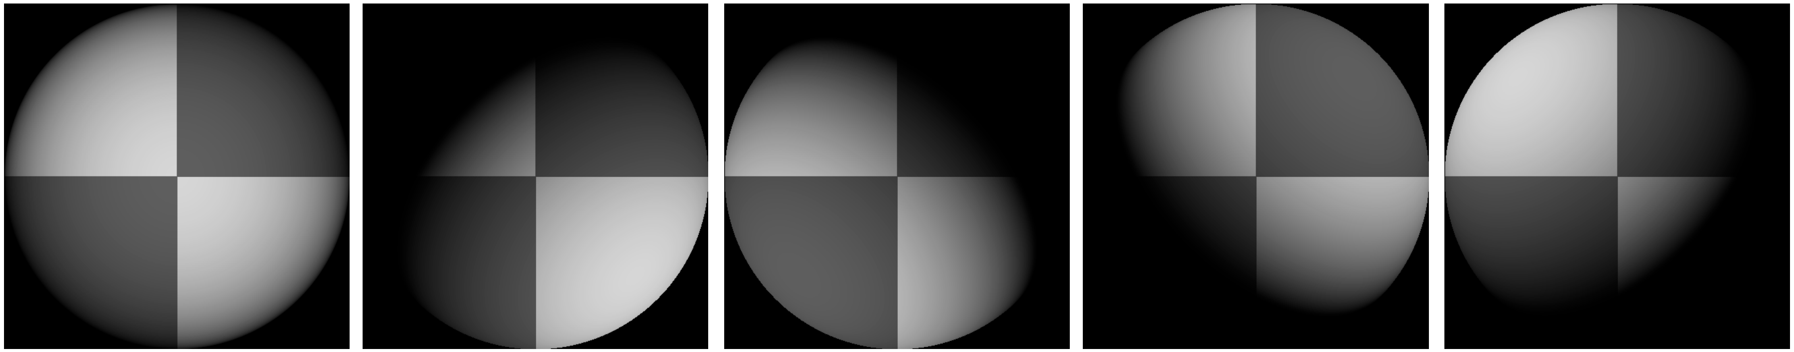

Recovered albedo

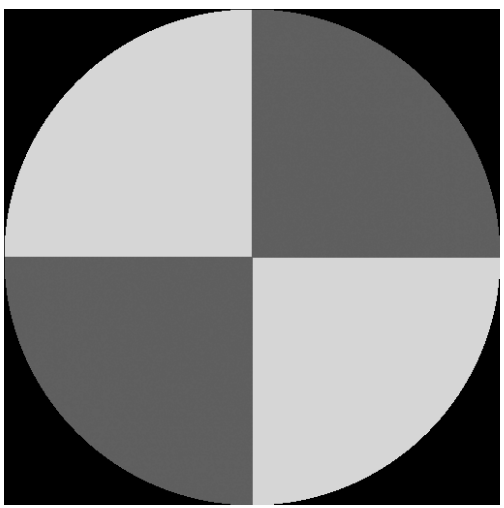

Recovered normal field

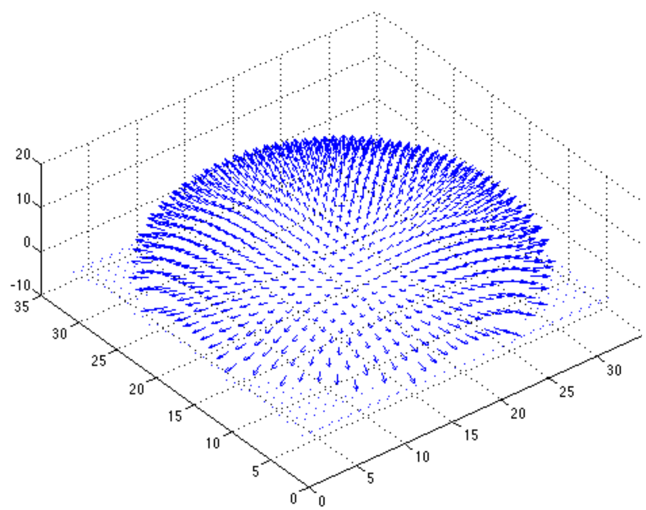

Recovered surface model

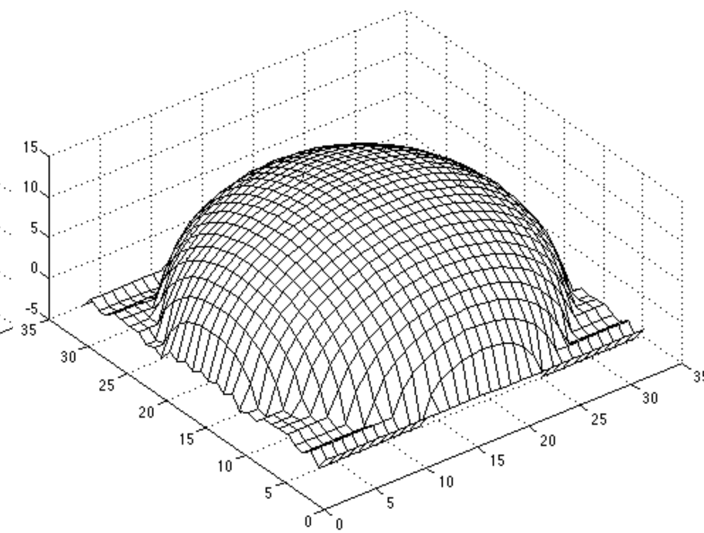

Slide from L. Lazebnik

F&P 2nd ed., sec. 2.2.4

Example 2

Input

Recovered albedo

Recovered normal field

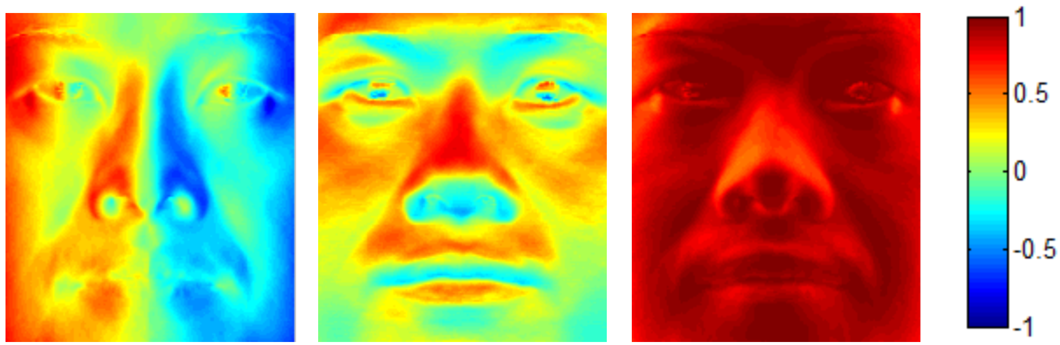

x y z

Recovered surface model

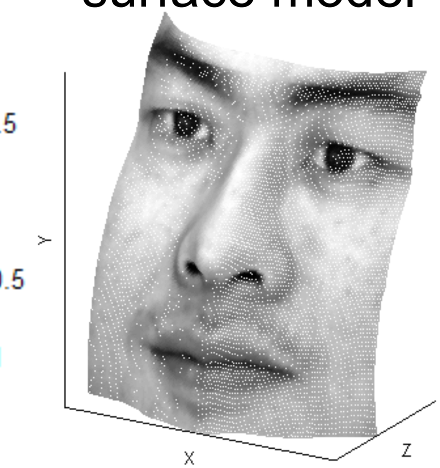

Slide from L. Lazebnik

Image model

- **Known:** source vectors $\mathbf{S}_j$ and pixel values $I_j(x,y)$
- **Unknown:** surface normal $\mathbf{N}(x,y)$ and albedo $\rho(x,y)$

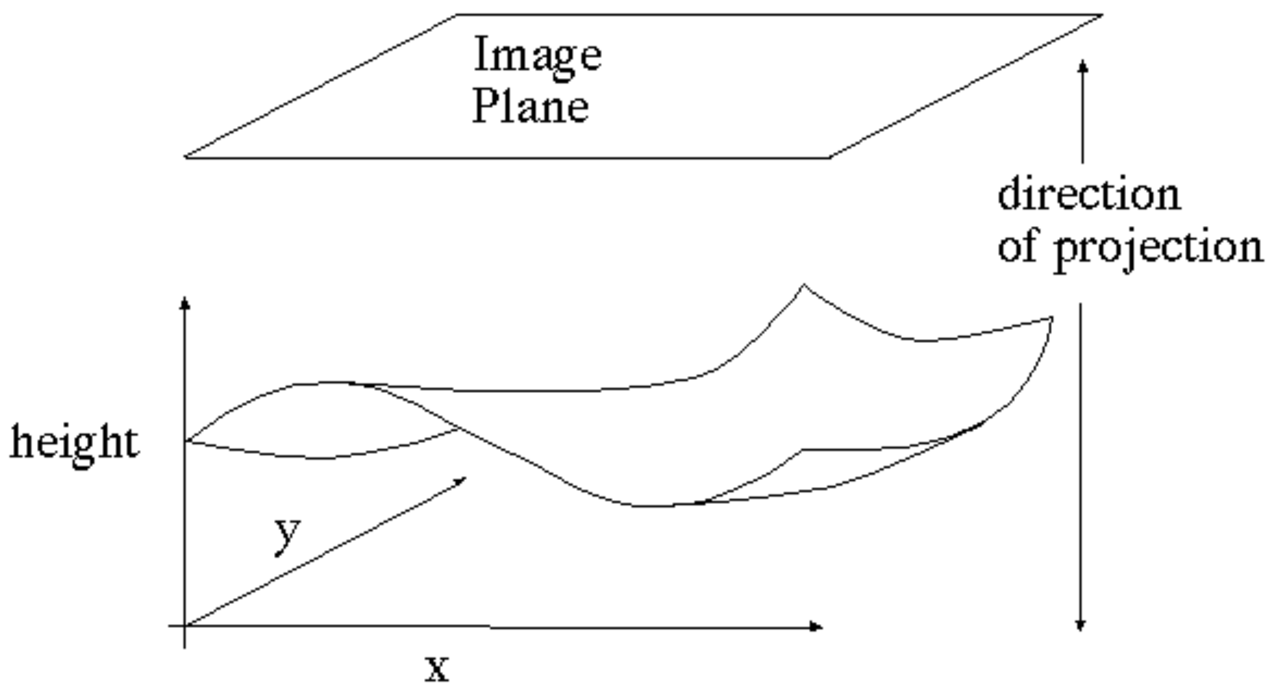

Slide from L. Lazebnik

F&P 2nd ed., sec. 2.2.4

Image model

- **Known:** source vectors $\mathbf{S}_j$ and pixel values $I_j(x,y)$
- **Unknown:** surface normal $\mathbf{N}(x,y)$ and albedo $\rho(x,y)$
- Assume that the response function of the camera is a linear scaling by a factor of k
- Lambert's law:

$$
\begin{aligned}
I_j(x,y) &= k\,\rho(x,y)\big(\mathbf{N}(x,y)\cdot\mathbf{S}_j\big) \\
&= \big(\rho(x,y)\mathbf{N}(x,y)\big)\cdot(k\mathbf{S}_j) \\
&= \mathbf{g}(x,y)\cdot\mathbf{V}_j
\end{aligned}
$$

Slide from L. Lazebnik

F&P 2nd ed., sec. 2.2.4

Least squares problem

- For each pixel, set up a linear system:

$$\begin{bmatrix} I_1(x,y) \\ I_2(x,y) \\ \vdots \\ I_n(x,y) \end{bmatrix} = \begin{bmatrix} \mathbf{V}_1^T \\ \mathbf{V}_2^T \\ \vdots \\ \mathbf{V}_n^T \end{bmatrix} \mathbf{g}(x,y)$$

(n × 1) known — (n × 3) known — (3 × 1) unknown

- Obtain least-squares solution for $\mathbf{g}(x,y)$ (which we defined as $\mathbf{N}(x,y)\ \rho(x,y)$)
- Since $\mathbf{N}(x,y)$ is the unit normal, $\rho(x,y)$ is given by the magnitude of $\mathbf{g}(x,y)$
- Finally, $\mathbf{N}(x,y) = \mathbf{g}(x,y)\ /\ \rho(x,y)$

Slide from L. Lazebnik

F&P 2nd ed., sec. 2.2.4

Synthetic example

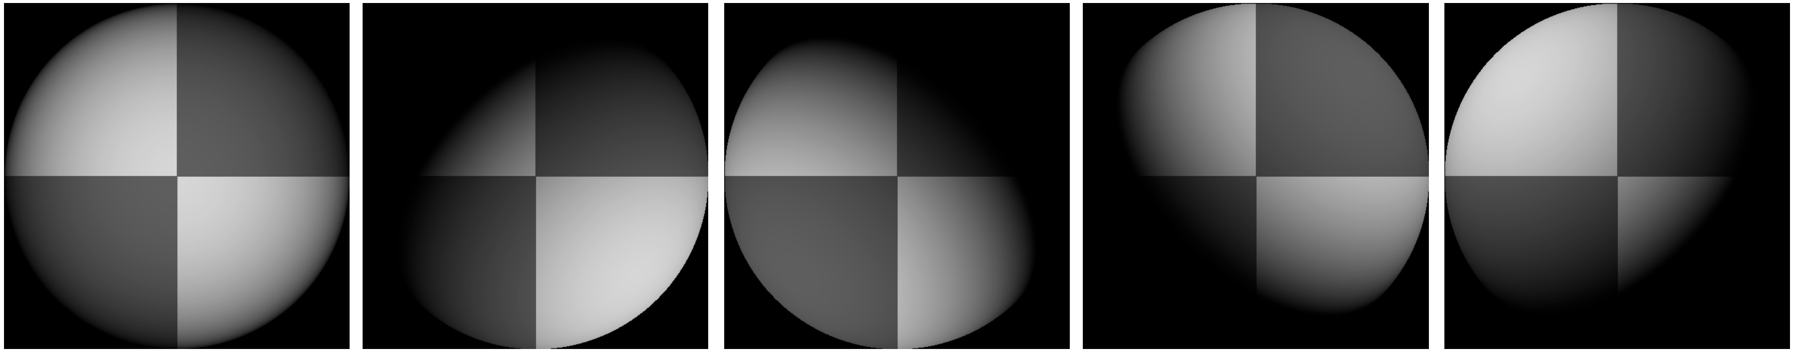

Recovered albedo

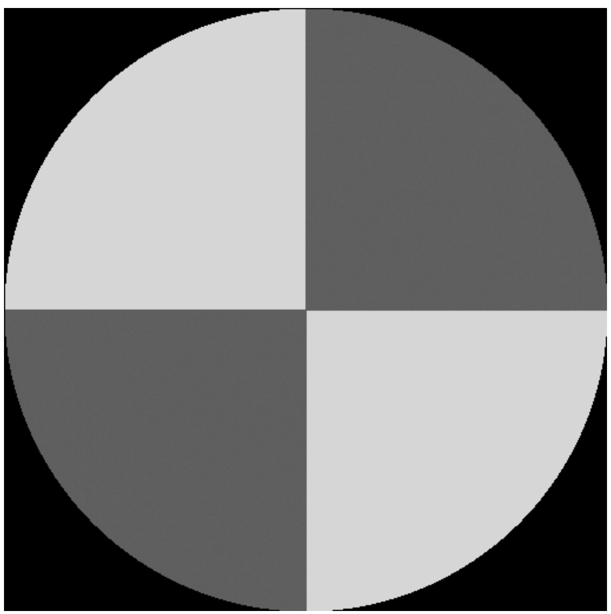

Recovered normal field

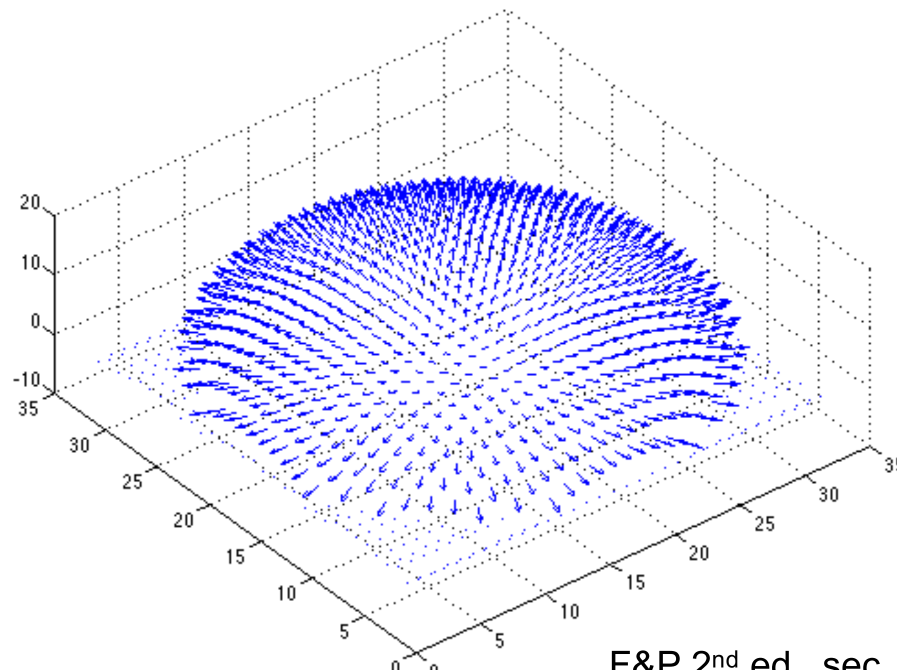

Slide from L. Lazebnik

F&P 2nd ed., sec. 2.2.4

Recovering a surface from normals

Recall the surface is written as

$$(x, y, f(x, y))$$

This means the normal has the form:

$$\mathbf{N}(x,y) = \frac{1}{\sqrt{f_x^2 + f_y^2 + 1}} \begin{pmatrix} f_x \\ f_y \\ 1 \end{pmatrix}$$

If we write the estimated vector *g* as

$$\mathbf{g}(x,y) = \begin{pmatrix} g_1(x,y) \\ g_2(x,y) \\ g_3(x,y) \end{pmatrix}$$

Then we obtain values for the partial derivatives of the surface:

$$f_x(x,y) = g_1(x,y) / g_3(x,y)$$

$$f_y(x,y) = g_2(x,y) / g_3(x,y)$$

Slide from L. Lazebnik

F&P 2nd ed., sec. 2.2.4

Recovering a surface from normals

We can now recover the surface height at any point by integration along some path, e.g.

$$f(x,y) = \int_0^x f_x(s,0)\,ds + \int_0^y f_y(x,t)\,dt + C$$

(for robustness, should take integrals over many different paths and average the results)

Integrability: for the surface f to exist, the mixed second partial derivatives must be equal:

$$\frac{\partial}{\partial y}(g_1(x,y)/g_3(x,y)) = \frac{\partial}{\partial x}(g_2(x,y)/g_3(x,y))$$

(in practice, they should at least be similar)

Surface recovered by integration

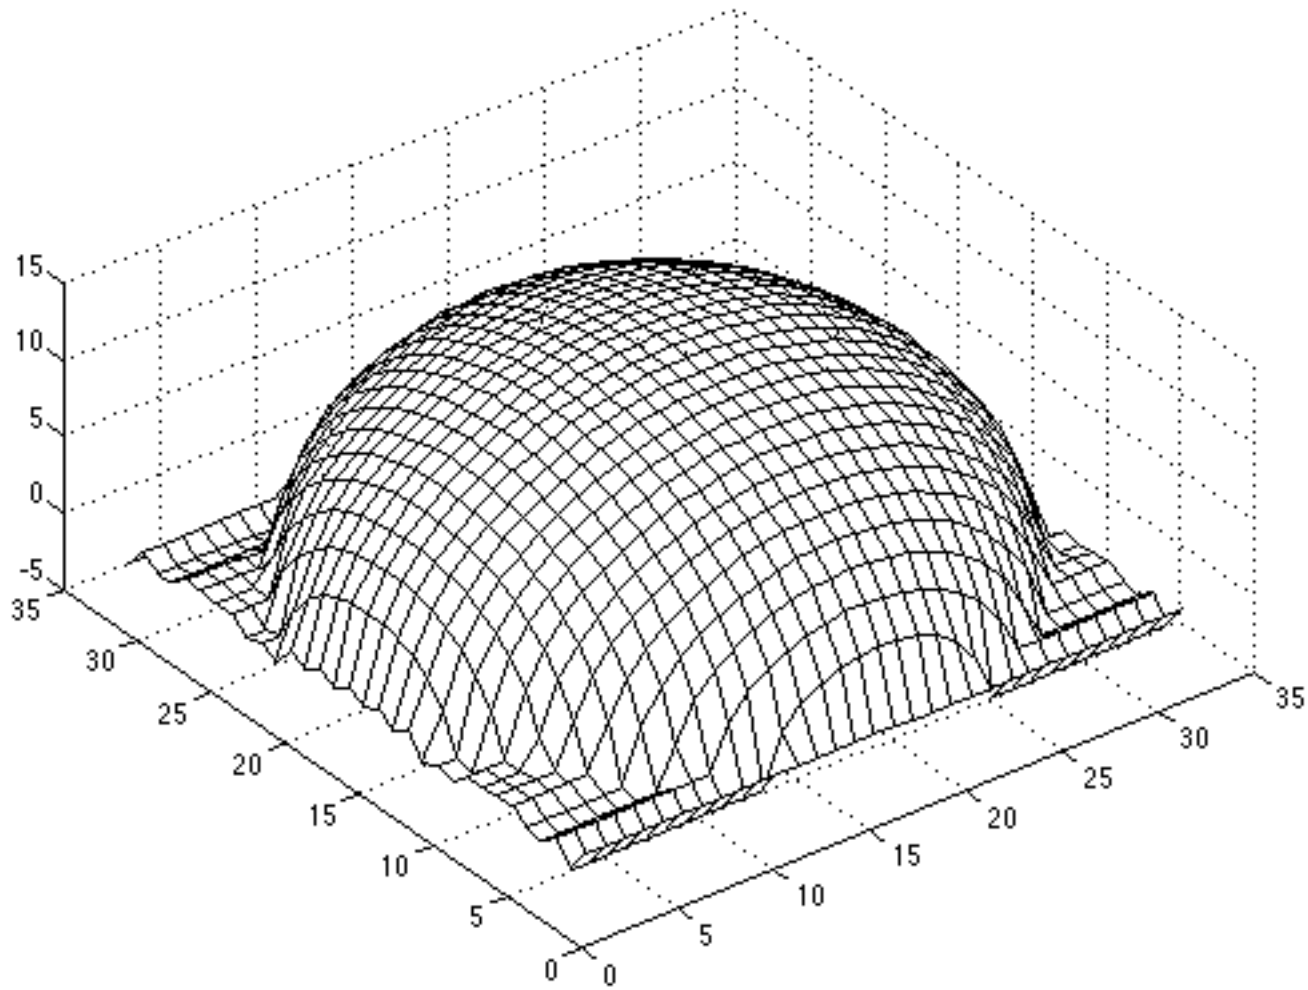

Slide from L. Lazebnik

F&P 2nd ed., sec. 2.2.4

Limitations

- Orthographic camera model
- Simplistic reflectance and lighting model
- No shadows
- No interreflections
- No missing data
- Integration is tricky

Slide from L. Lazebnik

Finding the direction of the light source

$$I(x,y) = \mathbf{N}(x,y) \cdot \mathbf{S}(x,y)$$

Full 3D case:

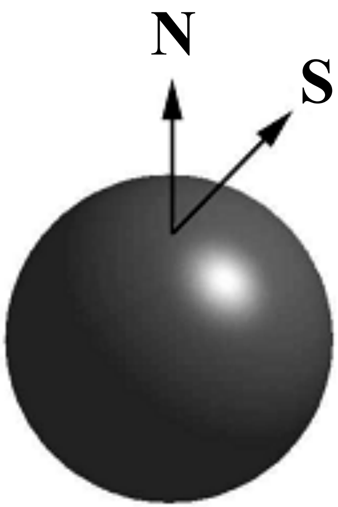

$$\begin{pmatrix} N_x(x_1,y_1) & N_y(x_1,y_1) & N_z(x_1,y_1) \\ N_x(x_2,y_2) & N_y(x_2,y_2) & N_z(x_2,y_2) \\ \vdots & \vdots & \vdots \\ N_x(x_n,y_n) & N_y(x_n,y_n) & N_z(x_n,y_n) \end{pmatrix} \begin{pmatrix} S_x \\ S_y \\ S_z \end{pmatrix} = \begin{pmatrix} I(x_1,y_1) \\ I(x_2,y_2) \\ \vdots \\ I(x_n,y_n) \end{pmatrix}$$

P. Nillius and J.-O. Eklundh, "Automatic estimation of the projected light source direction," CVPR 2001

Slide by L. Lazebnik

Finding the direction of the light source

Consider points on the *occluding contour*:

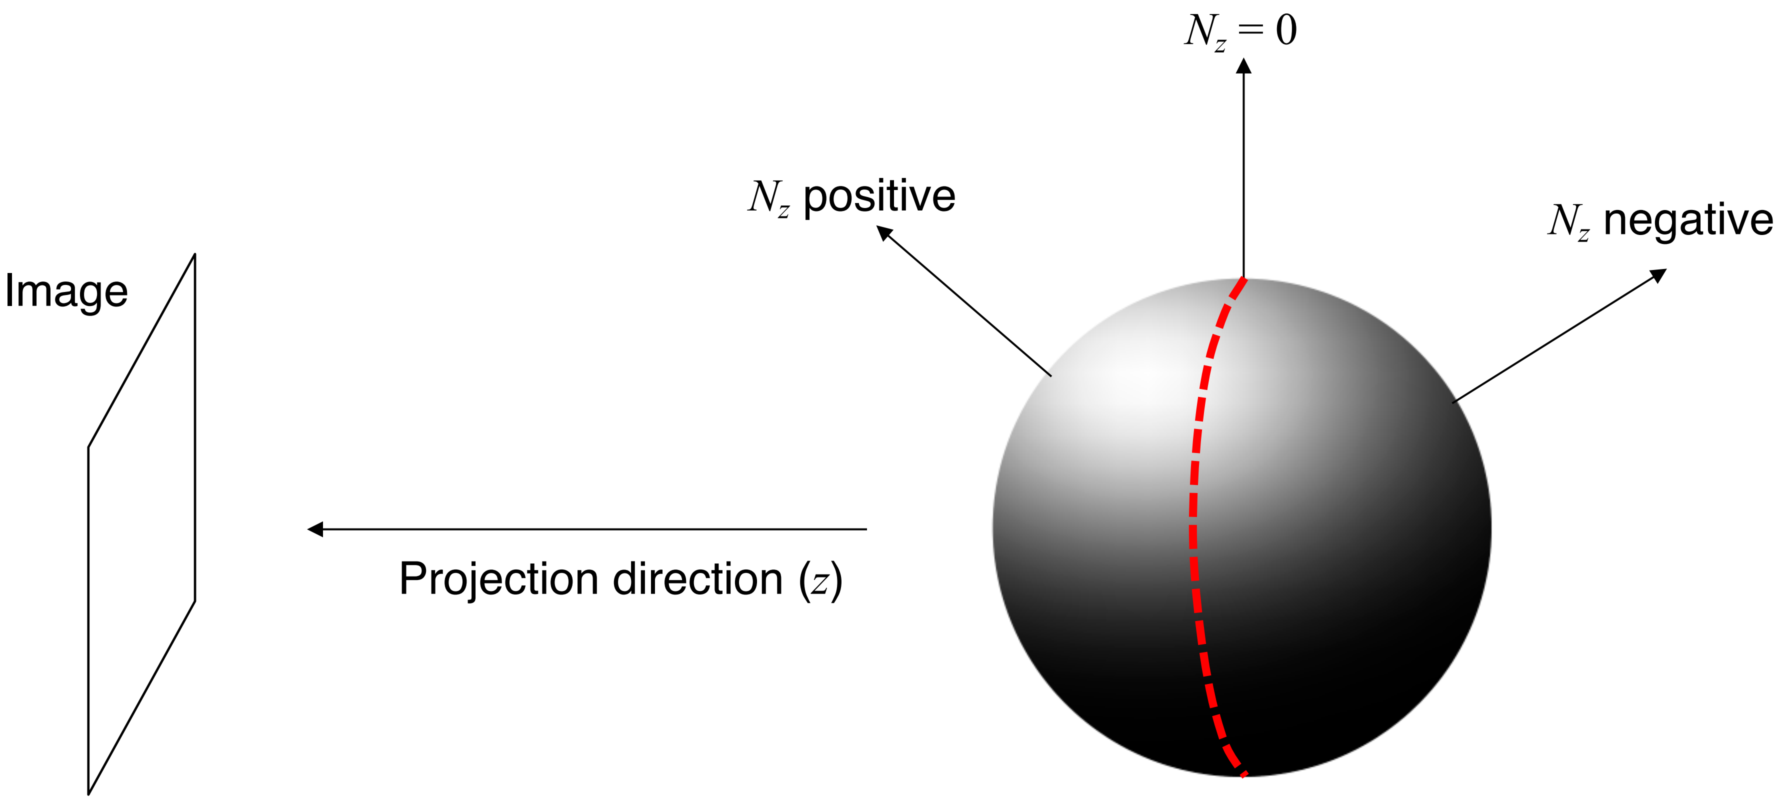

P. Nillius and J.-O. Eklundh, "Automatic estimation of the projected light source direction," CVPR 2001

Slide by L. Lazebnik

Finding the direction of the light source

$$I(x,y) = \mathbf{N}(x,y) \cdot \mathbf{S}(x,y)$$

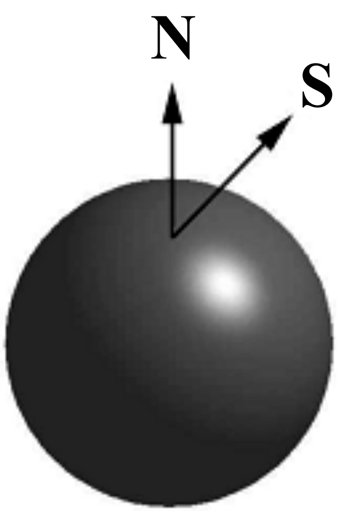

Full 3D case:

$$\begin{pmatrix} N_x(x_1,y_1) & N_y(x_1,y_1) & N_z(x_1,y_1) \\ N_x(x_2,y_2) & N_y(x_2,y_2) & N_z(x_2,y_2) \\ \vdots & \vdots & \vdots \\ N_x(x_n,y_n) & N_y(x_n,y_n) & N_z(x_n,y_n) \end{pmatrix} \begin{pmatrix} S_x \\ S_y \\ S_z \end{pmatrix} = \begin{pmatrix} I(x_1,y_1) \\ I(x_2,y_2) \\ \vdots \\ I(x_n,y_n) \end{pmatrix}$$

For points on the *occluding contour,* $N_z = 0$:

$$\begin{pmatrix} N_x(x_1,y_1) & N_y(x_1,y_1) \\ N_x(x_2,y_2) & N_y(x_2,y_2) \\ \vdots & \vdots \\ N_x(x_n,y_n) & N_y(x_n,y_n) \end{pmatrix} \begin{pmatrix} S_x \\ S_y \end{pmatrix} = \begin{pmatrix} I(x_1,y_1) \\ I(x_2,y_2) \\ \vdots \\ I(x_n,y_n) \end{pmatrix}$$

P. Nillius and J.-O. Eklundh, "Automatic estimation of the projected light source direction," CVPR 2001

Slide by L. Lazebnik

Finding the direction of the light source

P. Nillius and J.-O. Eklundh, “Automatic estimation of the projected light source direction,” CVPR 2001

Slide by L. Lazebnik

Application: Detecting composite photos

Fake photo

Real photo

M. K. Johnson and H. Farid, Exposing Digital Forgeries by Detecting Inconsistencies in Lighting, ACM Multimedia and Security Workshop, 2005.

Slide by L. Lazebnik

Overview

- Cameras with lenses
  - Depth of field
  - Field of view
  - Lens aberrations
- Brightness of a pixel
  - Small taste of radiometry
  - In-camera transformation of light
  - Reflectance properties of surfaces
  - Lambertian reflection model
  - Shape from shading
- Color