Light, Camera and Shading

CS 543 / ECE 549 – Saurabh Gupta

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

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

Building a Real Camera

Home-made pinhole camera

http://www.debevec.org/Pinhole/

Slide by A. Efros

Shrinking the aperture

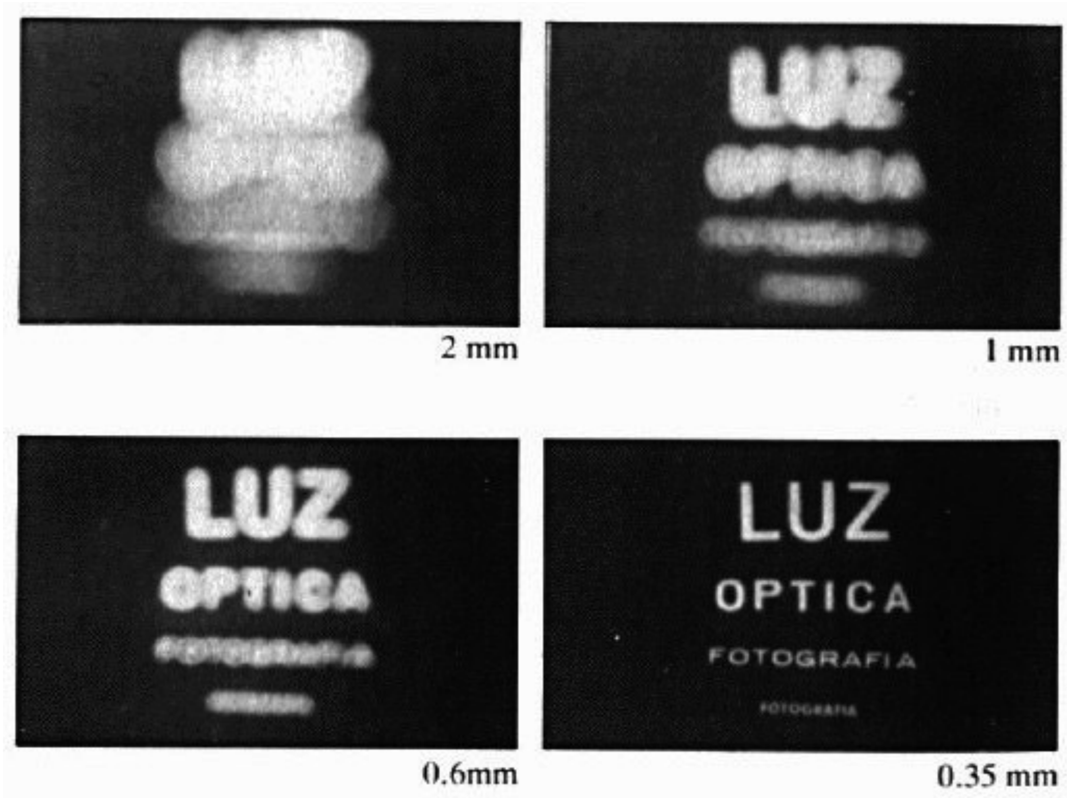

Why not make the aperture as small as possible?

- Less light gets through
- Diffraction effects…

Slide by Steve Seitz

Shrinking the aperture

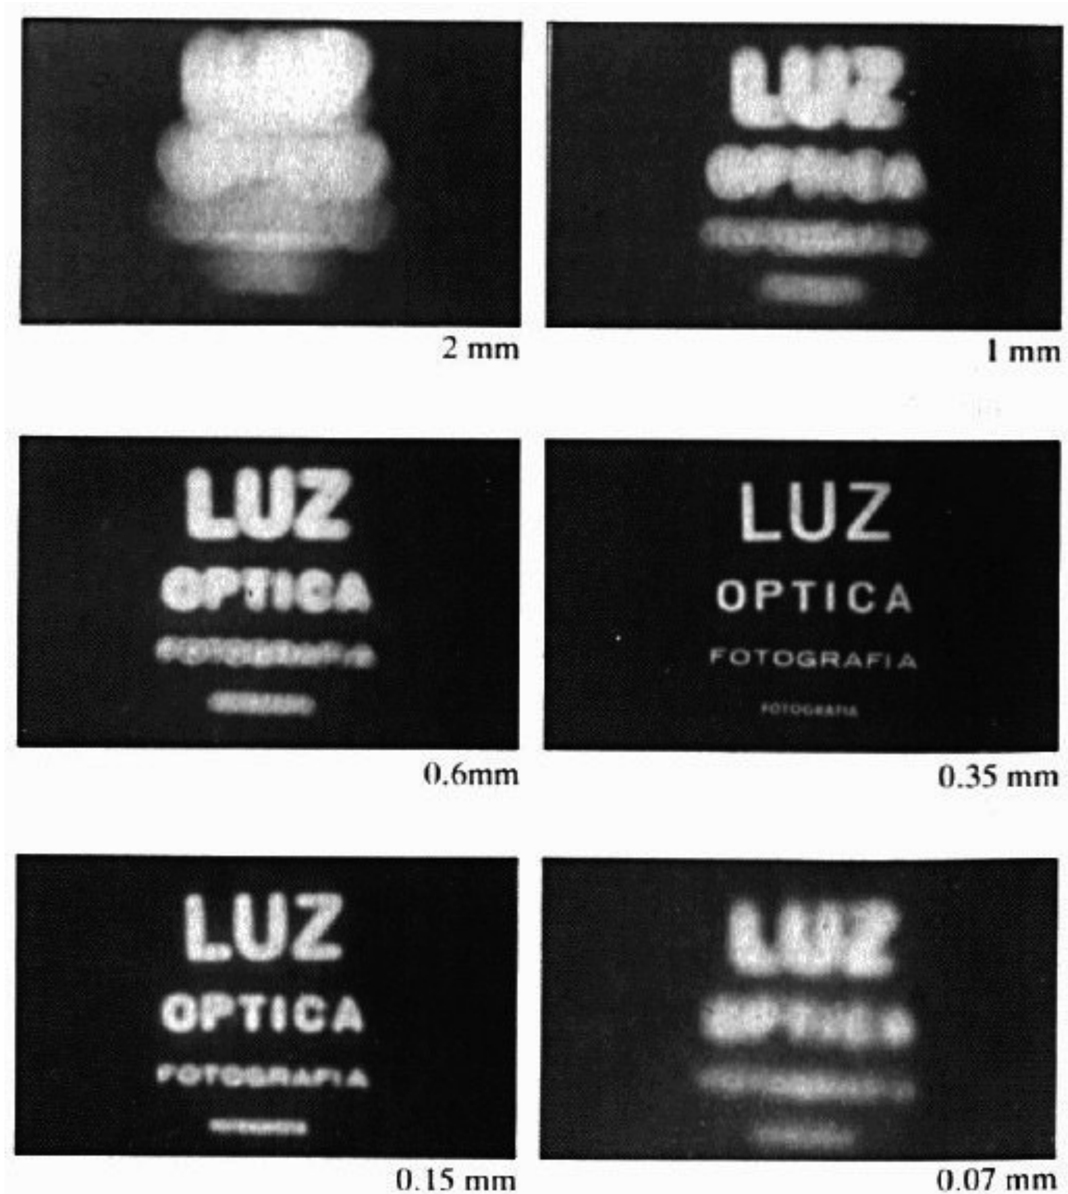

Slide by Steve Seitz

Adding a lens

Adding a lens

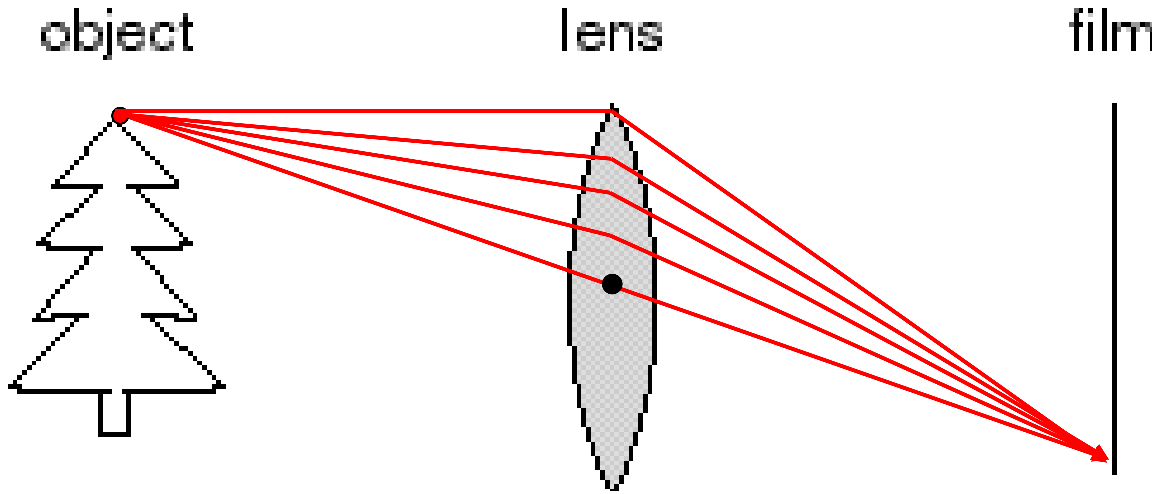

A lens focuses light onto the film

- Thin lens model:
  - Rays passing through the center are not deviated (pinhole projection model still holds)

Slide by Steve Seitz

Adding a lens

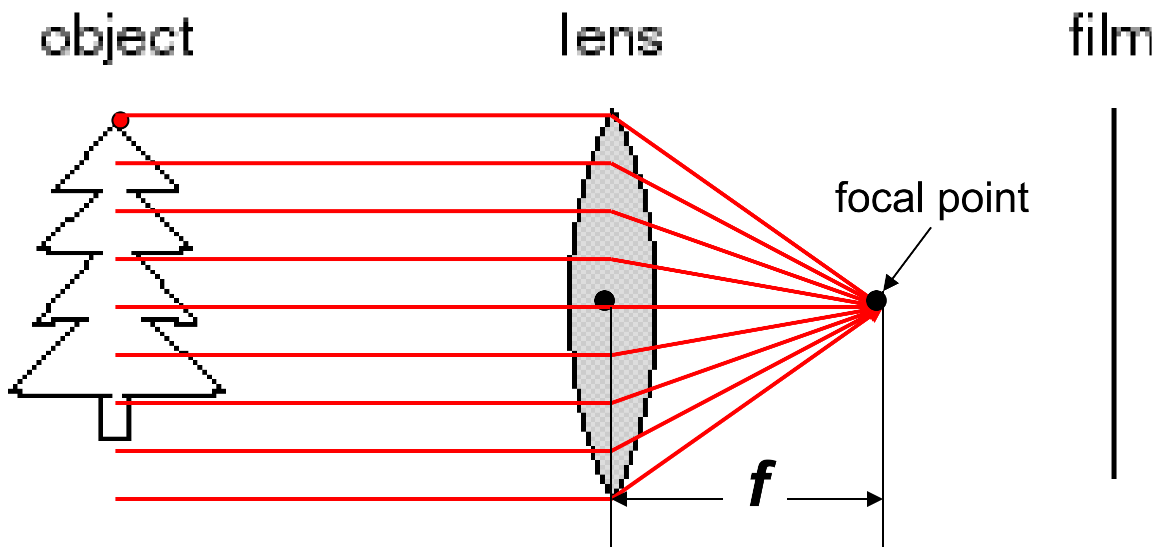

A lens focuses light onto the film

- Thin lens model:
  - Rays passing through the center are not deviated (pinhole projection model still holds)
  - All rays parallel to the optical axis pass through the *focal point*
  - All parallel rays converge to points on the *focal plane*

Slide by Steve Seitz

Thin lens formula

- Where does the lens focus the rays coming from a given point in the scene?

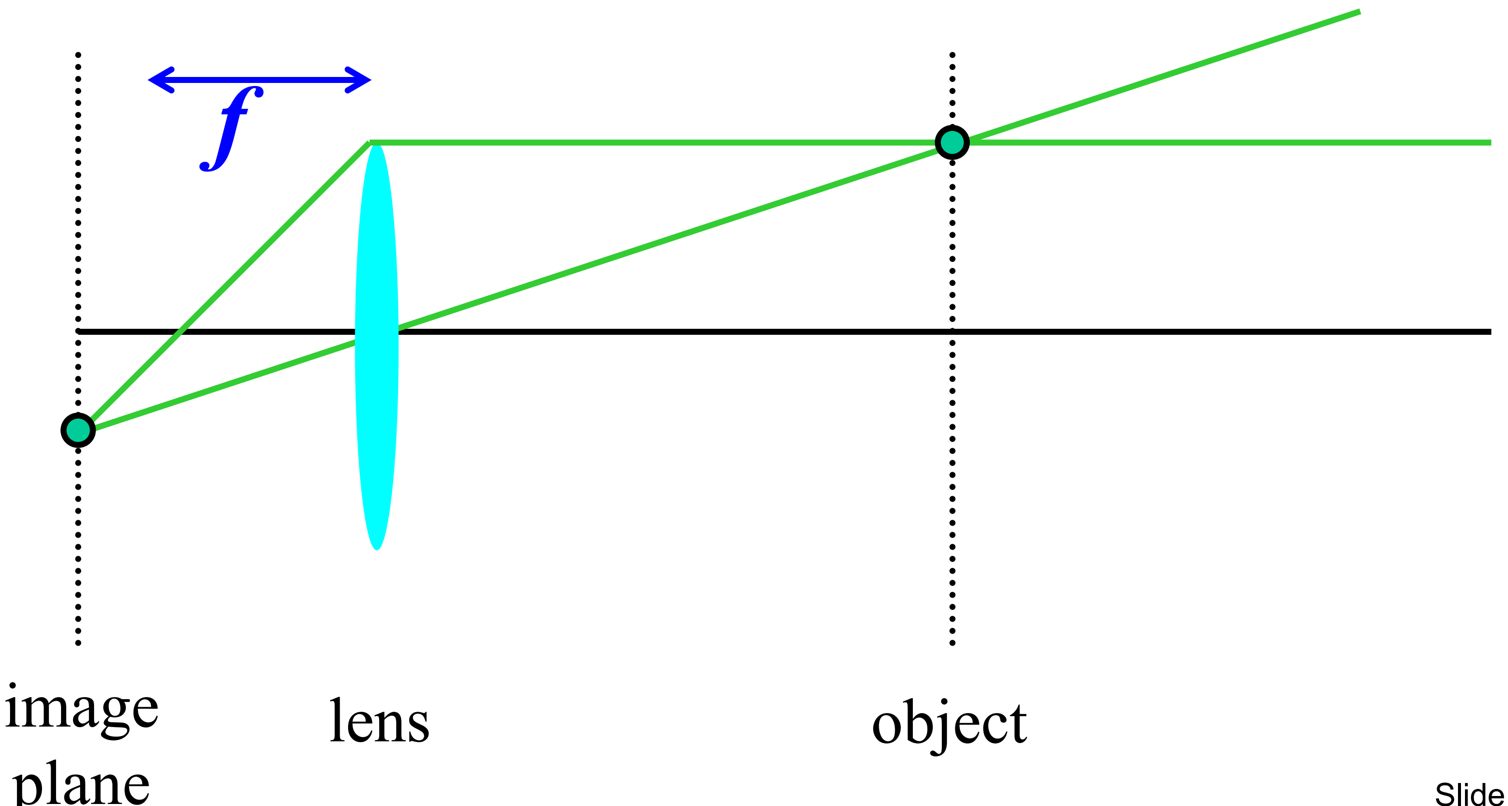

Slide by Frédo Durand

Thin lens formula

- What is the relation between the focal length (***f***), the distance of the object from the optical center (***D***), and the distance at which the object will be in focus (***D′***)?

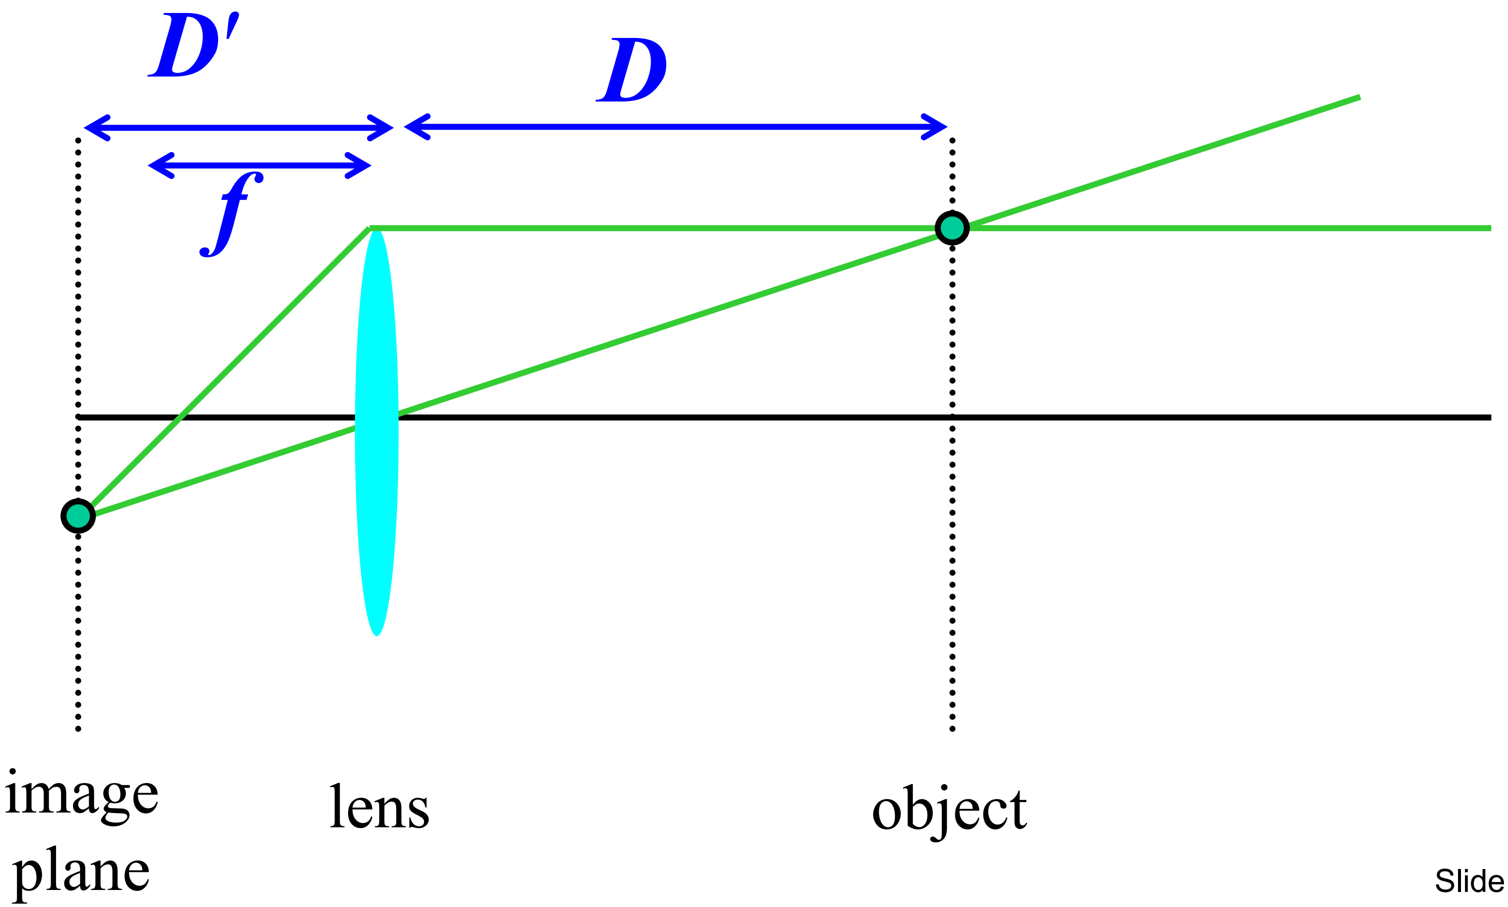

Slide by Frédo Durand

Thin lens formula

Similar triangles everywhere!

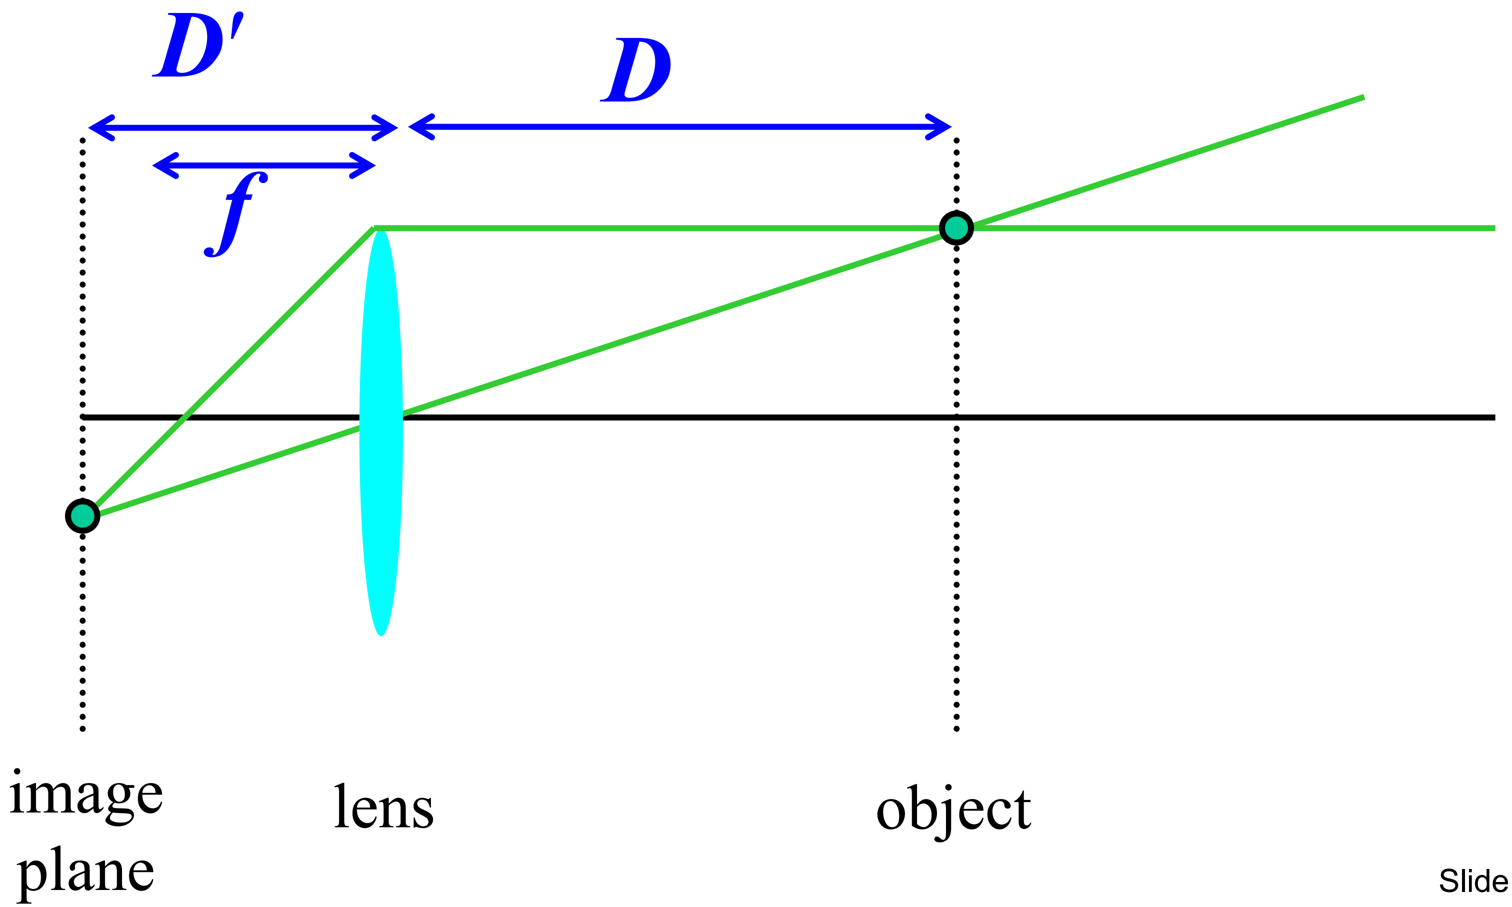

Slide by Frédo Durand

Thin lens formula

Similar triangles everywhere!

$$y'/y = D'/D$$

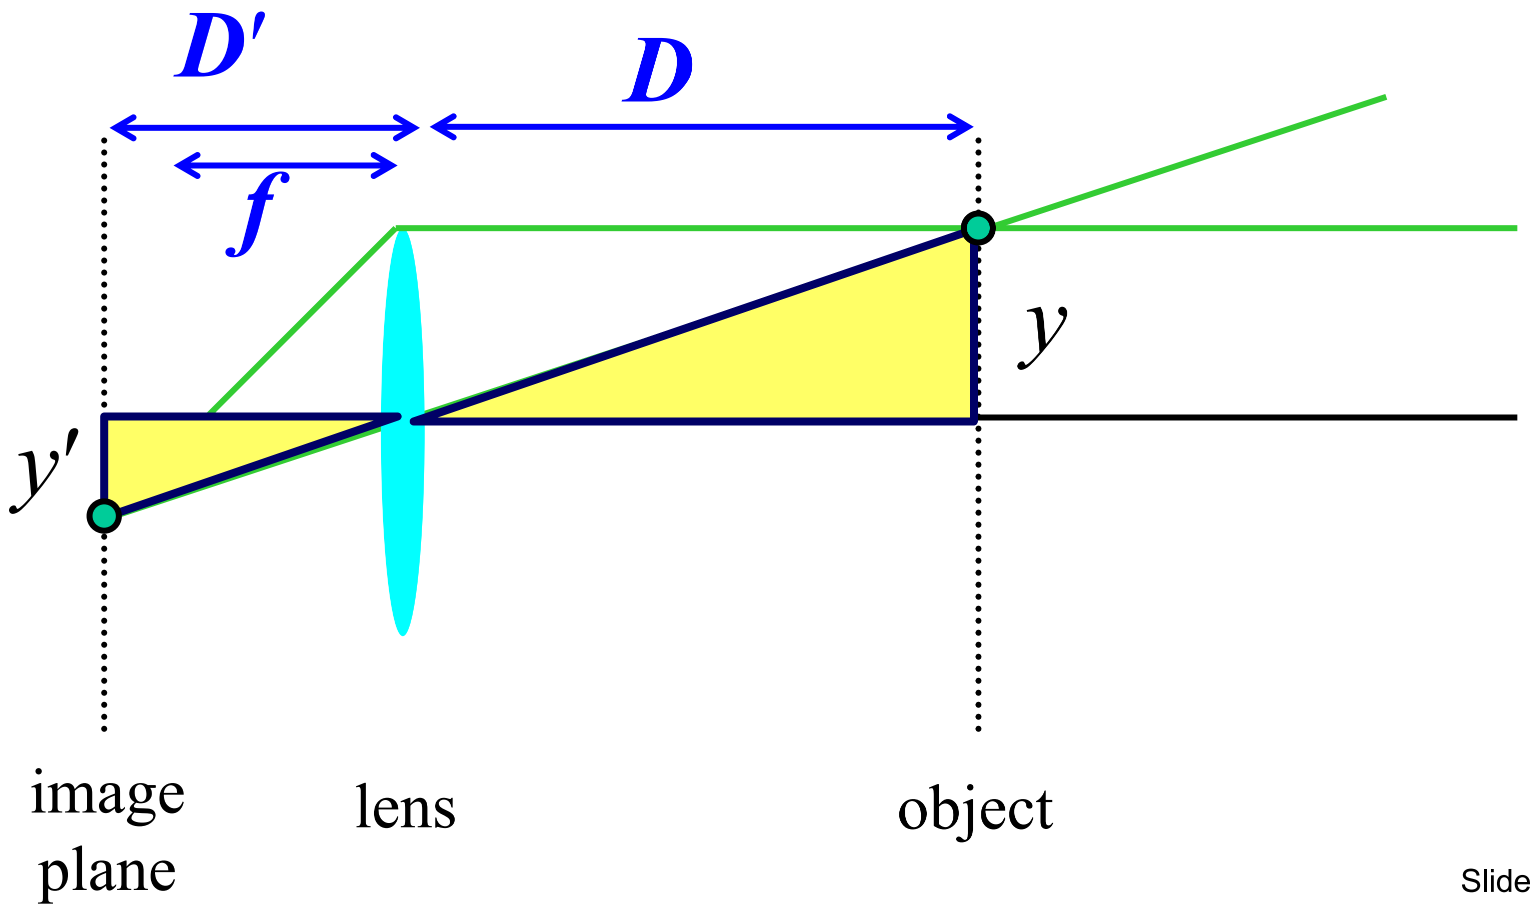

Slide by Frédo Durand

Thin lens formula

Similar triangles everywhere!

$$y'/y = D'/D$$

$$y'/y = (D'-f)/f$$

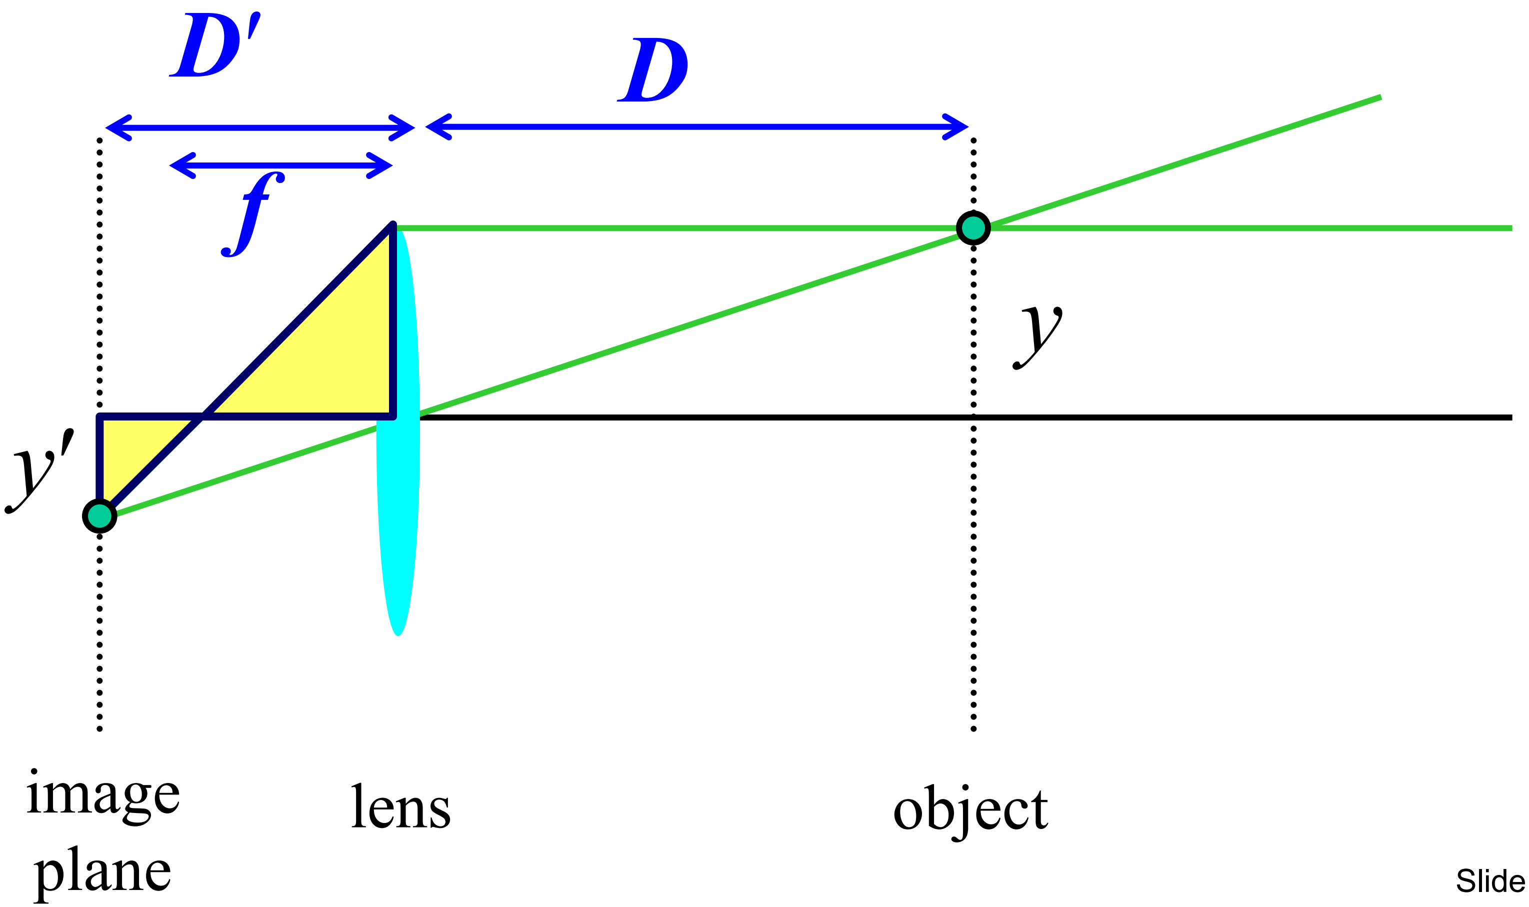

Slide by Frédo Durand

Thin lens formula

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

Any point satisfying the thin lens equation is in focus.

What happens when D is very large?

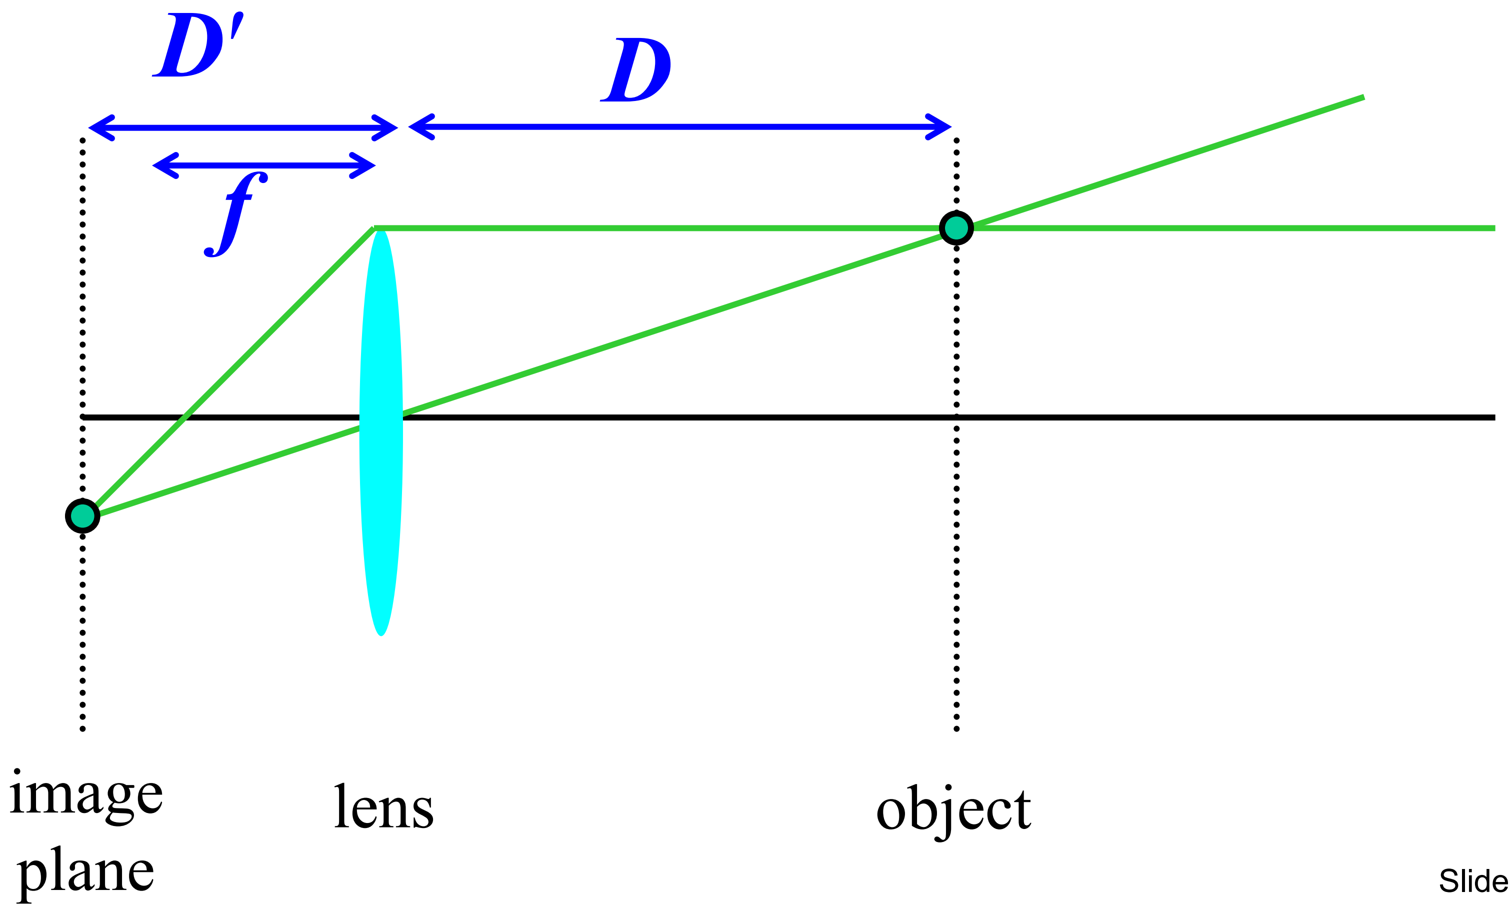

Slide by Frédo Durand

Depth of Field

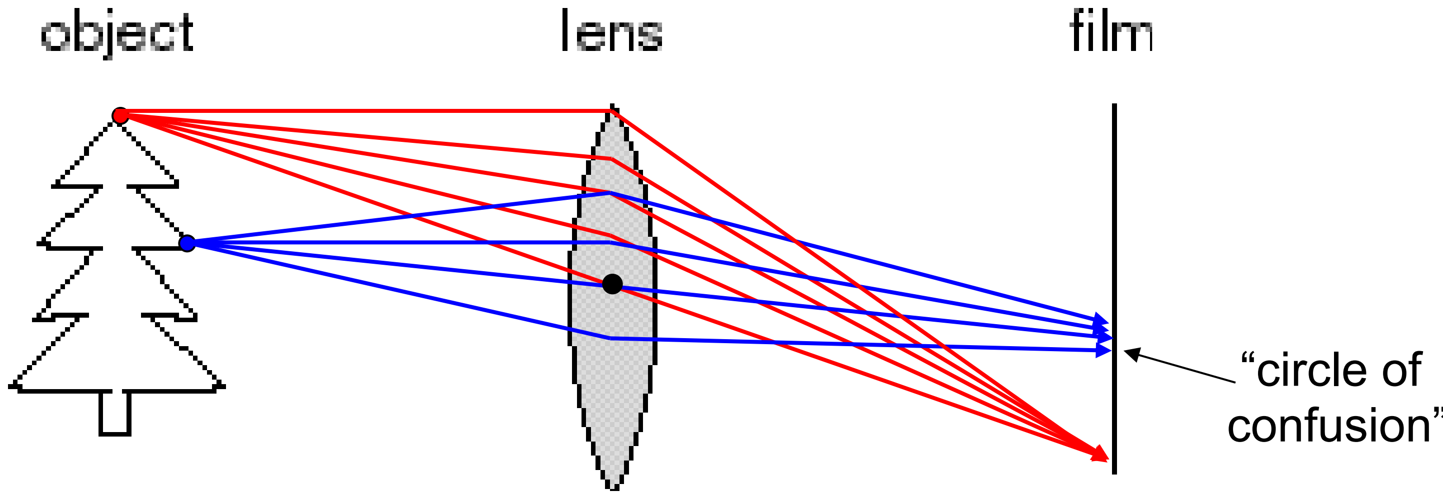

For a fixed focal length, there is a specific distance at which objects are “in focus”

- Other points project to a “circle of confusion” in the image

Depth of Field

http://www.cambridgeincolour.com/tutorials/depth-of-field.htm

Slide by A. Efros

Controlling depth of field

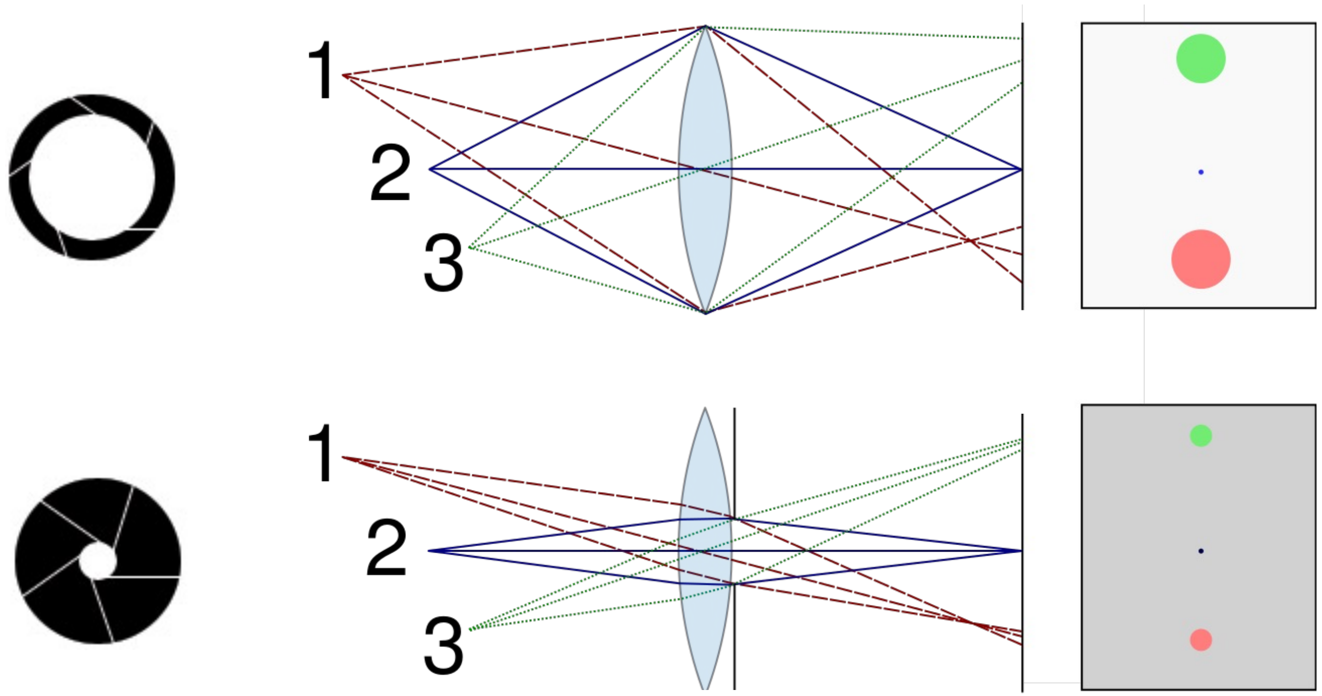

Changing the aperture size affects depth of field

- A smaller *aperture* increases the range in which the object is approximately in focus
- But small aperture reduces amount of light – need to increase *exposure*

http://en.wikipedia.org/wiki/File:Depth_of_field_illustration.svg

Varying the aperture

Large aperture = small DOF

Small aperture = large DOF

Slide by A. Efros

Field of View

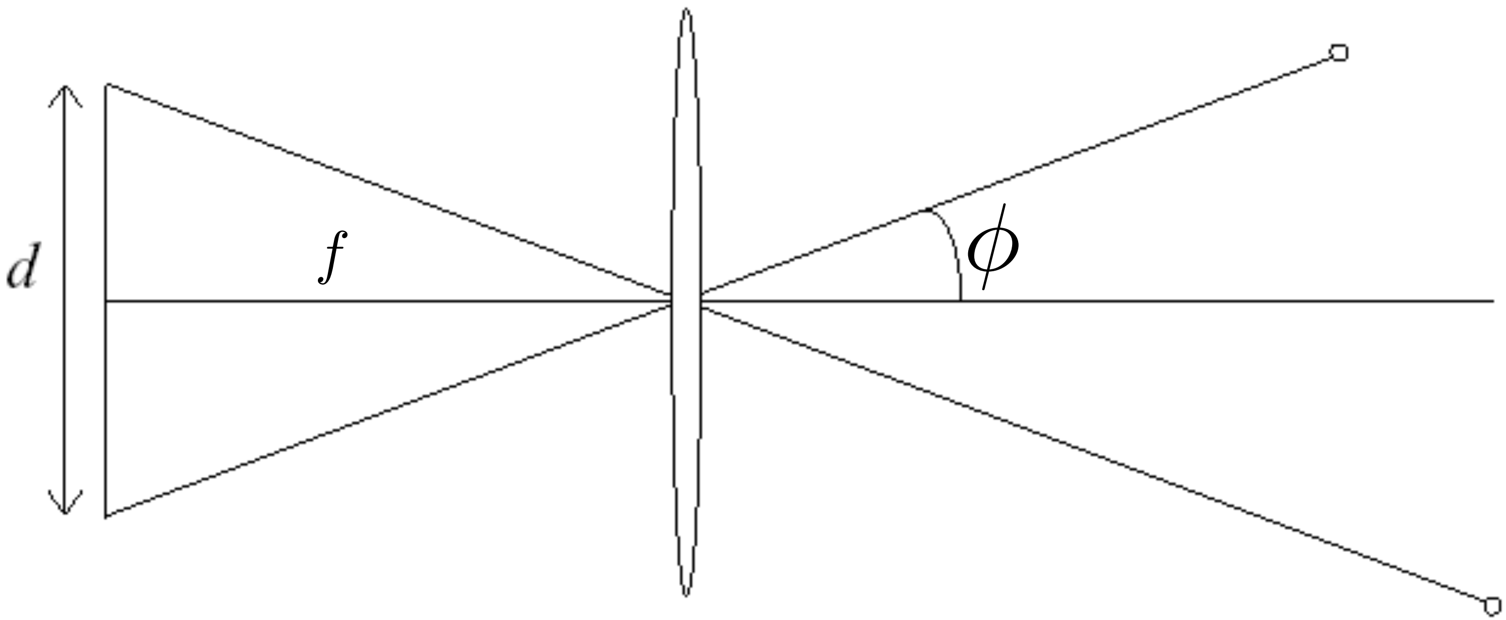

FOV depends on focal length and size of the camera retina

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

Larger focal length = smaller FOV

Slide by A. Efros

Field of View

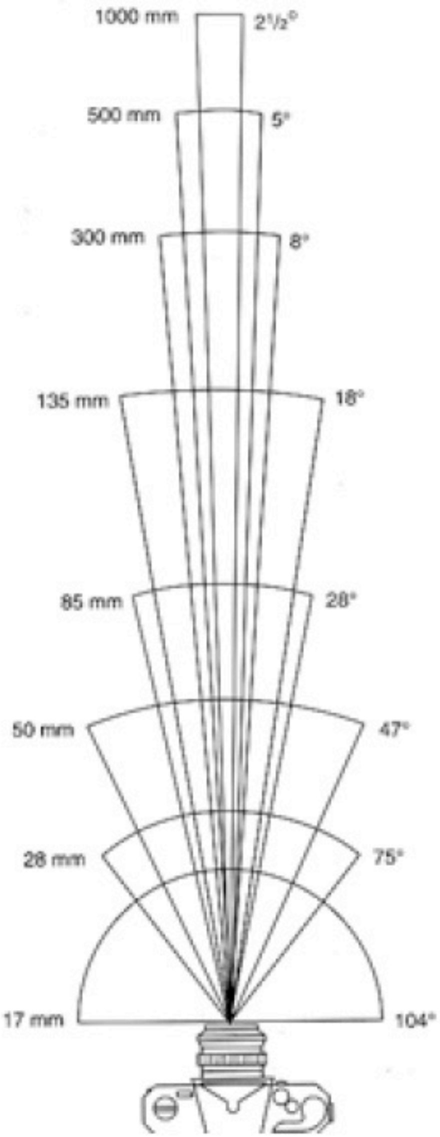

17mm

28mm

50mm

85mm

Slide by A. Efros

Field of View

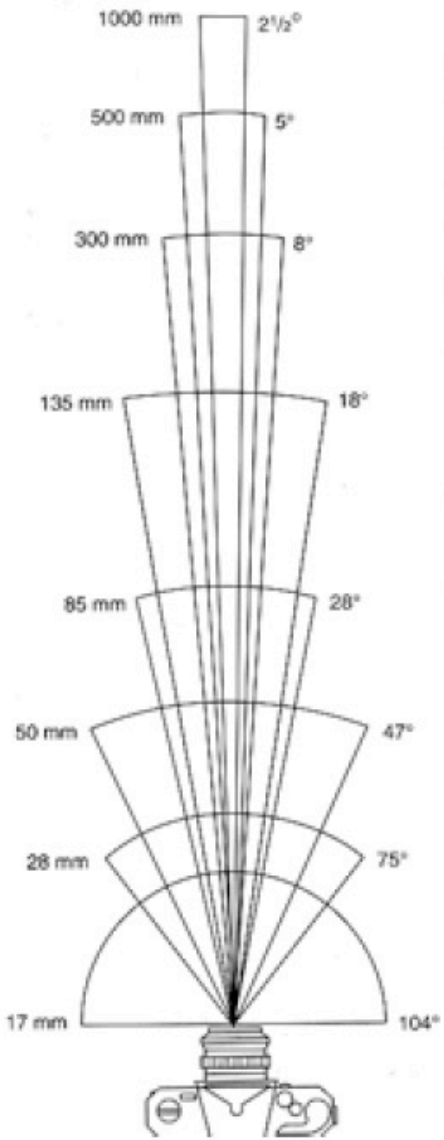

135mm

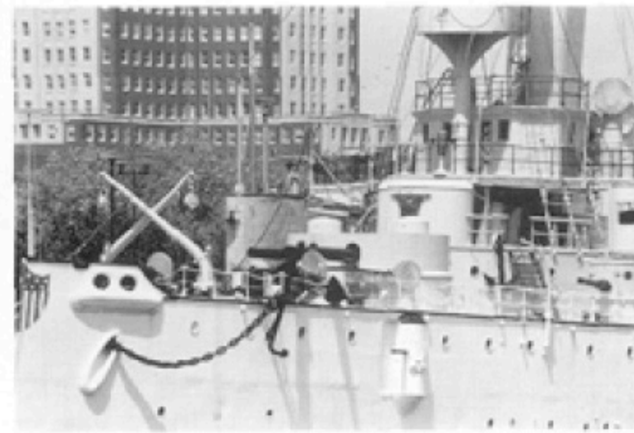

300mm

Slide by A. Efros

Field of View / Focal Length

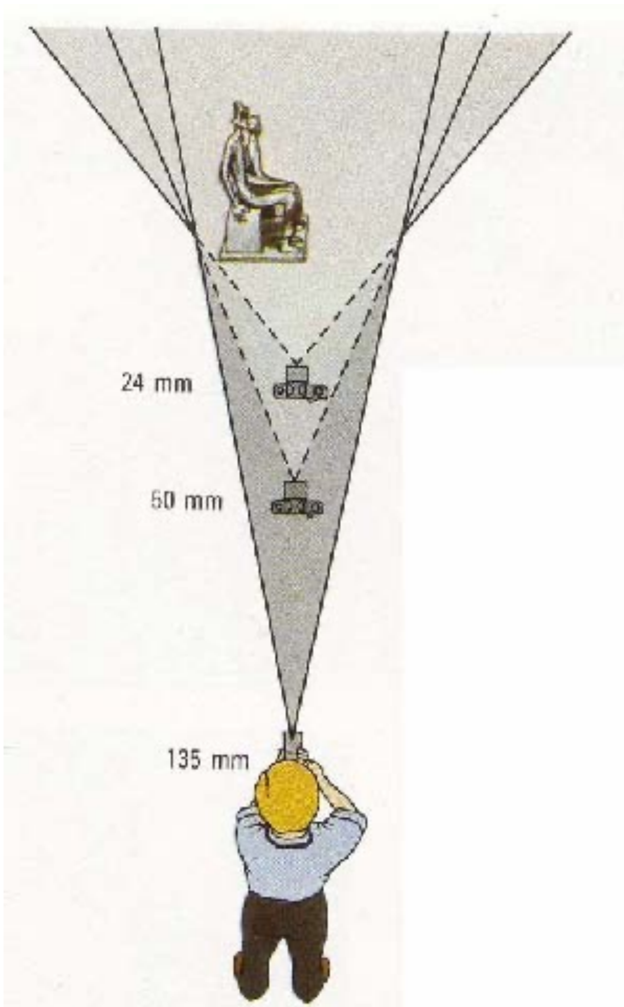

Large FOV, small *f*
Camera close to car

Small FOV, large *f*
Camera far from the car

Sources: A. Efros, F. Durand

Same effect for faces

wide-angle

standard

telephoto

Source: F. Durand

Approximating an orthographic camera

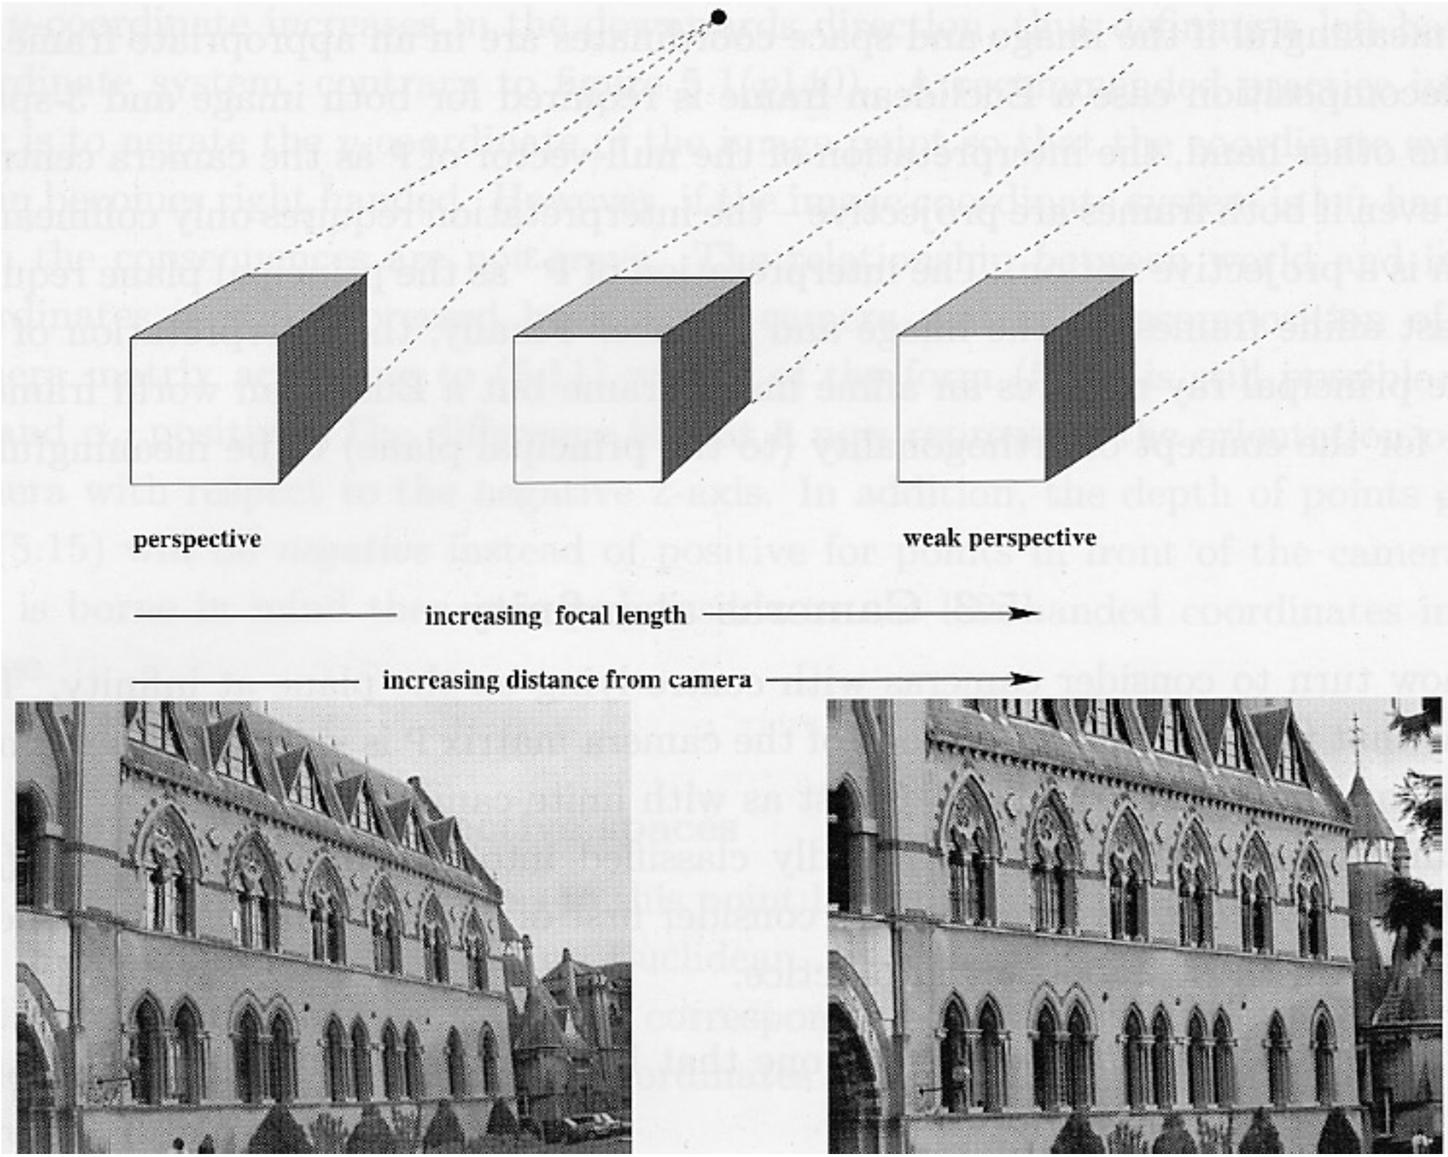

Source: Hartley & Zisserman

The dolly zoom

- Continuously adjusting the focal length while the camera moves away from (or towards) the subject

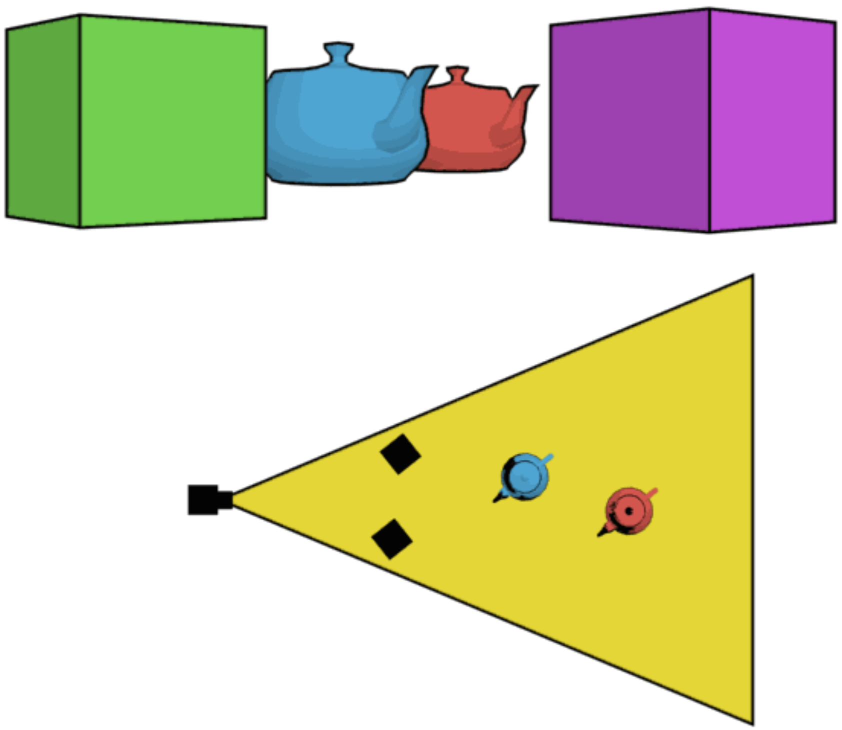

http://en.wikipedia.org/wiki/Dolly_zoom

The dolly zoom

- Continuously adjusting the focal length while the camera moves away from (or towards) the subject
- "The Vertigo shot"

[Example of dolly zoom from *Goodfellas*](#) (YouTube)

[Example of dolly zoom from *La Haine*](#) (YouTube)

Real lenses

Lens flaws: Vignetting

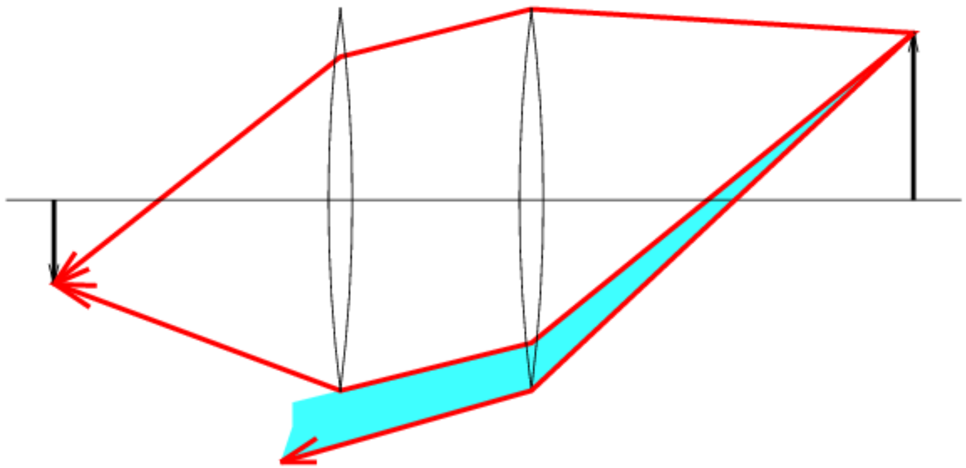

Radial Distortion

- Caused by imperfect lenses
- Deviations are most noticeable near the edge of the lens

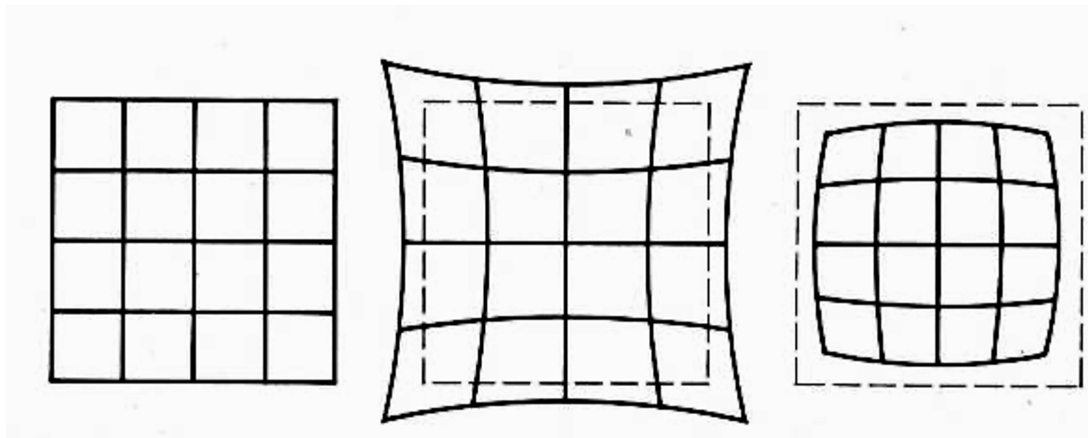

No distortion Pin cushion Barrel

Lens flaws: Spherical aberration

Spherical lenses don't focus light perfectly

- Rays farther from the optical axis focus closer

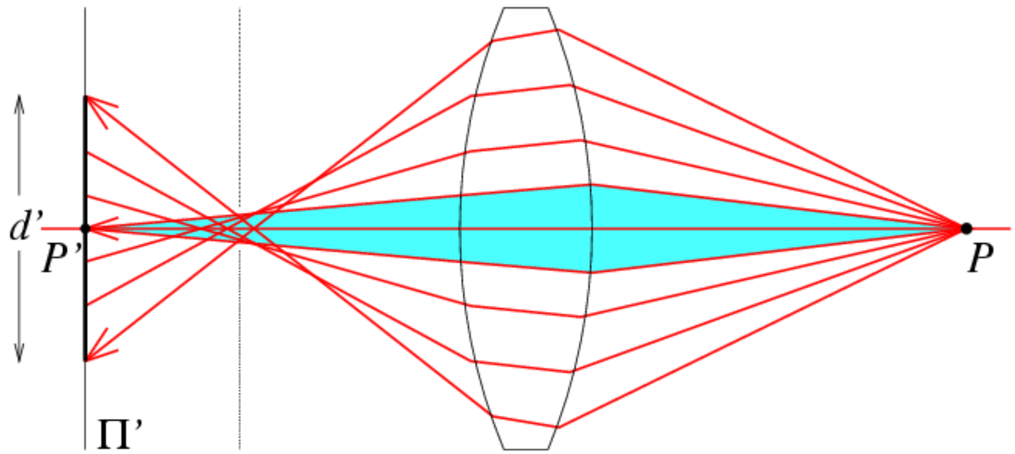

Lens Flaws: Chromatic Aberration

Lens has different refractive indices for different wavelengths: causes color fringing

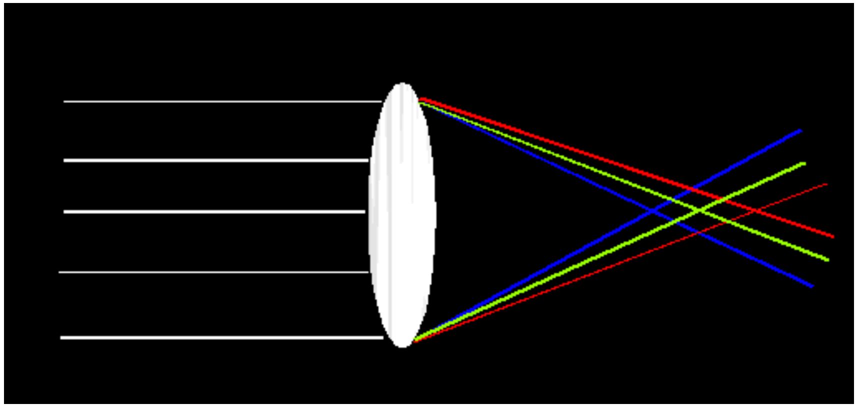

Near Lens Center

Near Lens Outer Edge

Lens Flaws: Chromatic Aberration

Researchers tried teaching a network about objects by forcing it to assemble jigsaws.

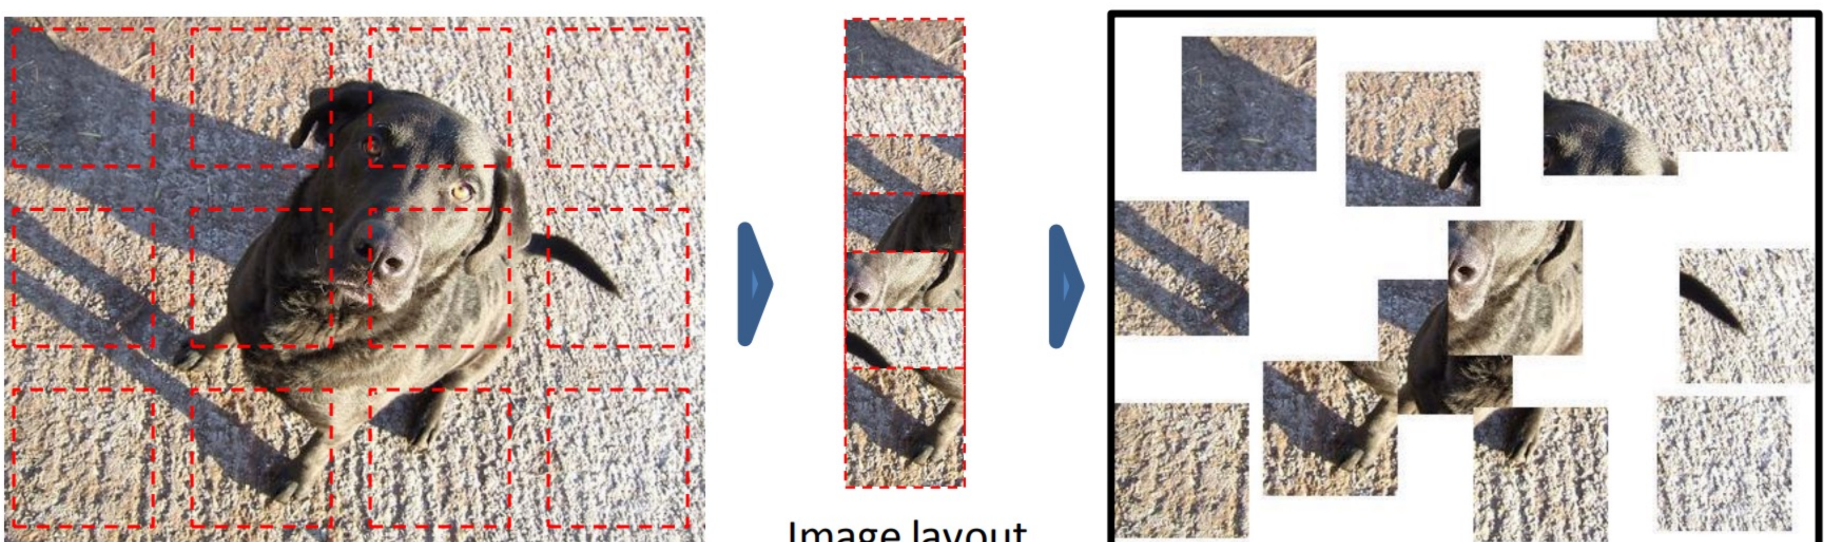

Initial layout, with sampled patches in red

Image layout is discarded

We can recover image layout automatically

Slide Credit: C. Doersch

Rolling Shutter

Rolling Shutter: pixels read in sequence

Can get global reading, but $$$

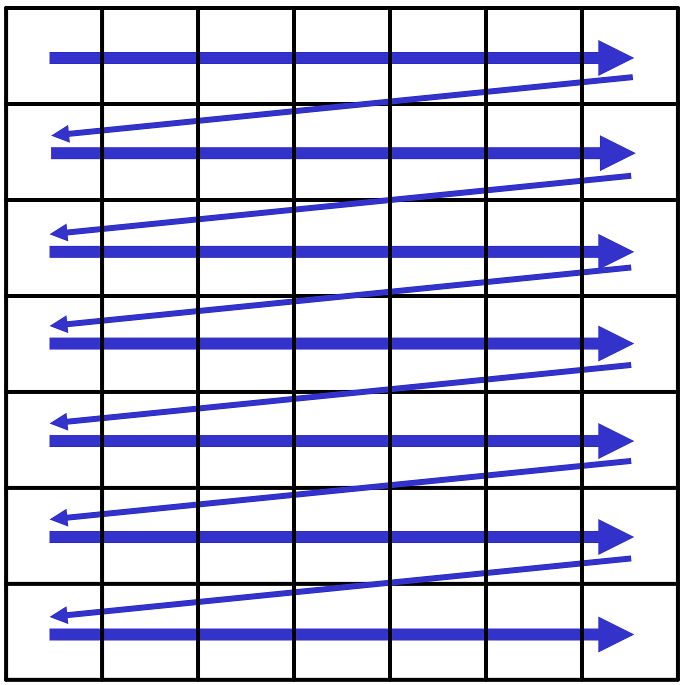

Overview

- 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

Image formation

What determines the brightness of an image pixel?

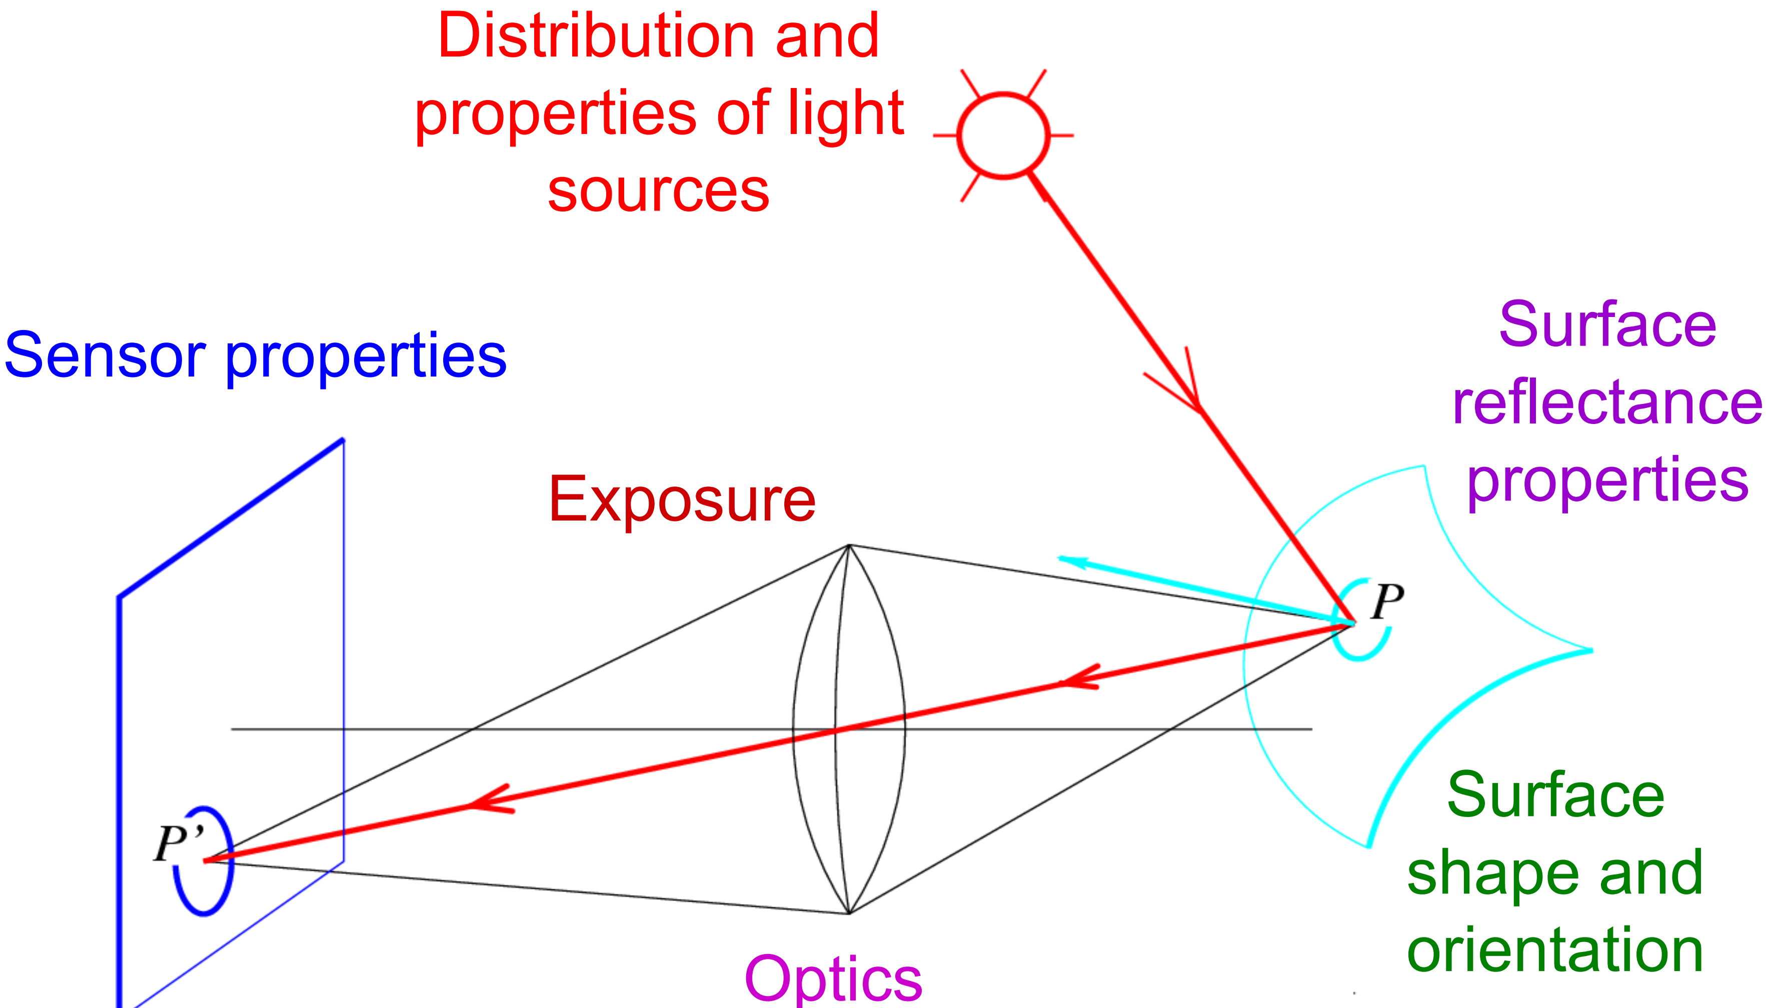

Fundamental radiometric relation

L: *Radiance* emitted from *P* toward *P*'

- Energy carried by a ray (Watts per sq. meter per steradian)

E: *Irradiance* falling on *P*' from the lens

- Energy arriving at a surface (Watts per sq. meter)

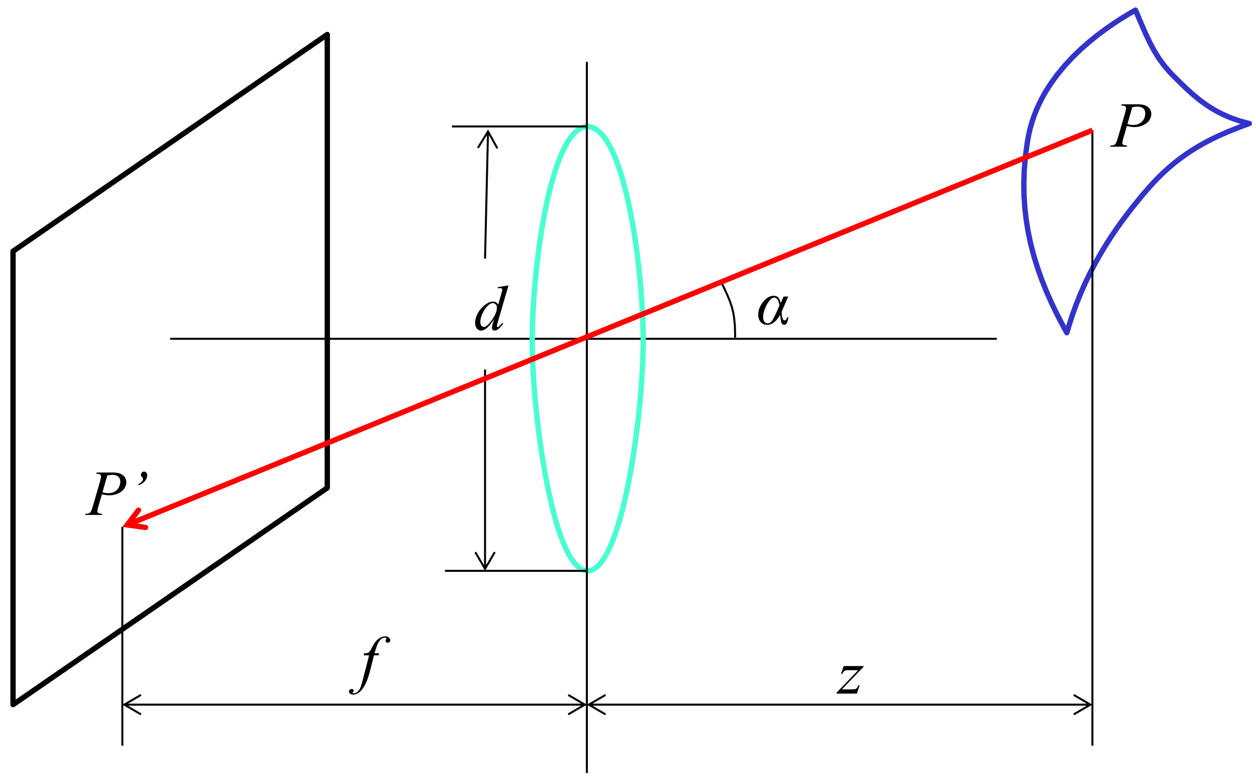

What is the relationship between *E* and *L*?

Szeliski 2.2.3

Fundamental radiometric relation

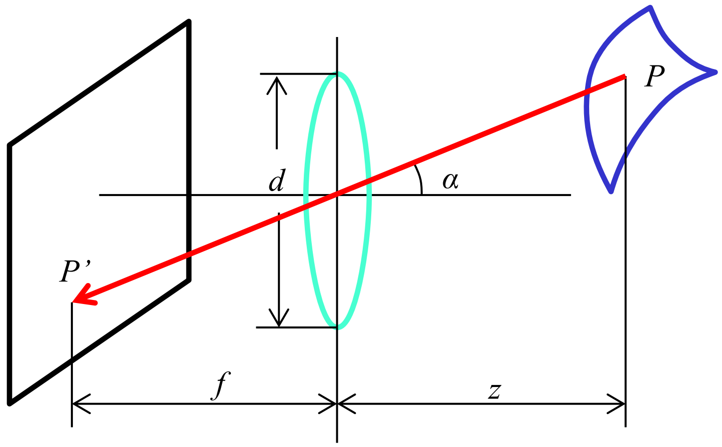

$$E = \left[\frac{\pi}{4} \left(\frac{d}{f} \right)^2 \cos^4 \alpha \right] L$$

- Image irradiance is linearly related to scene radiance
- Irradiance is proportional to the area of the lens and inversely proportional to the squared distance between the lens and the image plane
- The irradiance falls off as the angle between the viewing ray and the optical axis increases

Relation between Image Irradiance E and Scene Radiance L

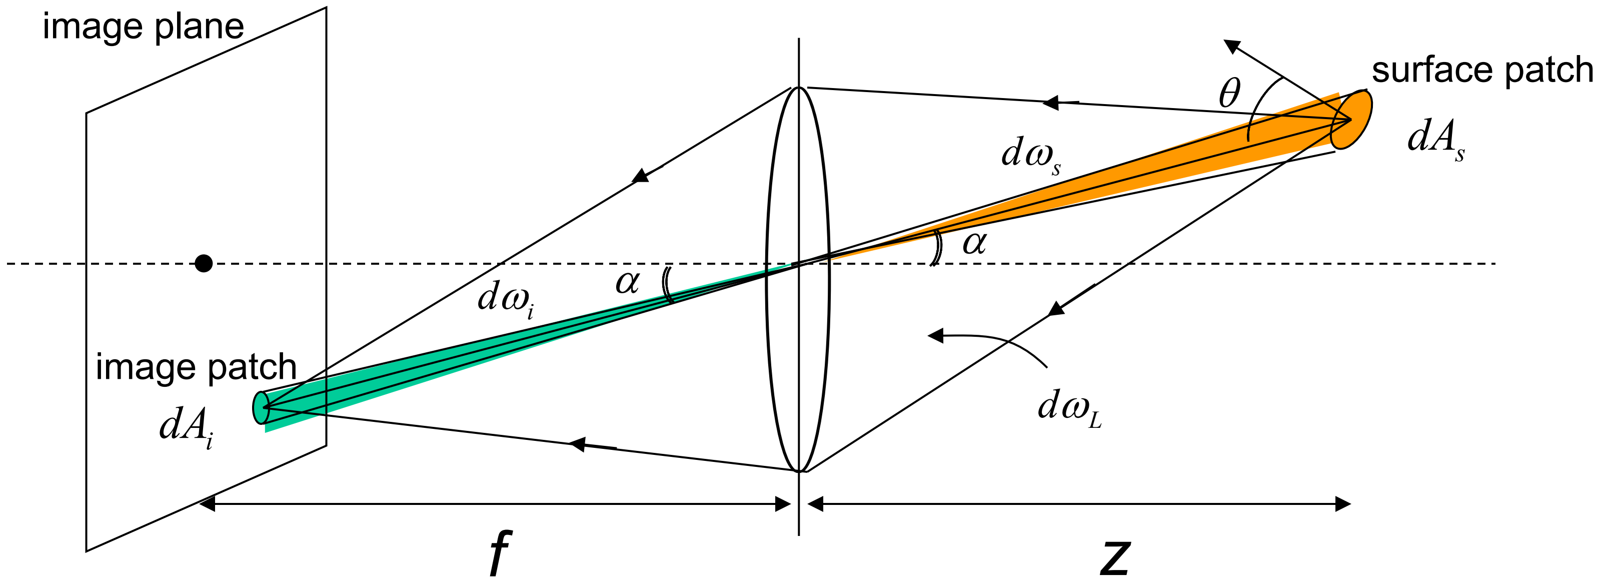

- Solid angles of the double cone (orange and green):

$$d\omega_i = d\omega_s \qquad \frac{dA_i \cos\alpha}{(f/\cos\alpha)^2} = \frac{dA_s \cos\theta}{(z/\cos\alpha)^2}$$

$$\boxed{\frac{dA_s}{dA_i} = \frac{\cos\alpha}{\cos\theta}\left(\frac{z}{f}\right)^2} \rightarrow (1)$$

- Solid angle subtended by lens:

$$\boxed{d\omega_L = \frac{\pi d^2}{4}\frac{\cos\alpha}{(z/\cos\alpha)^2}} \rightarrow (2)$$

Slide from S Narasimhan.

Relation between Image Irradiance E and Scene Radiance L

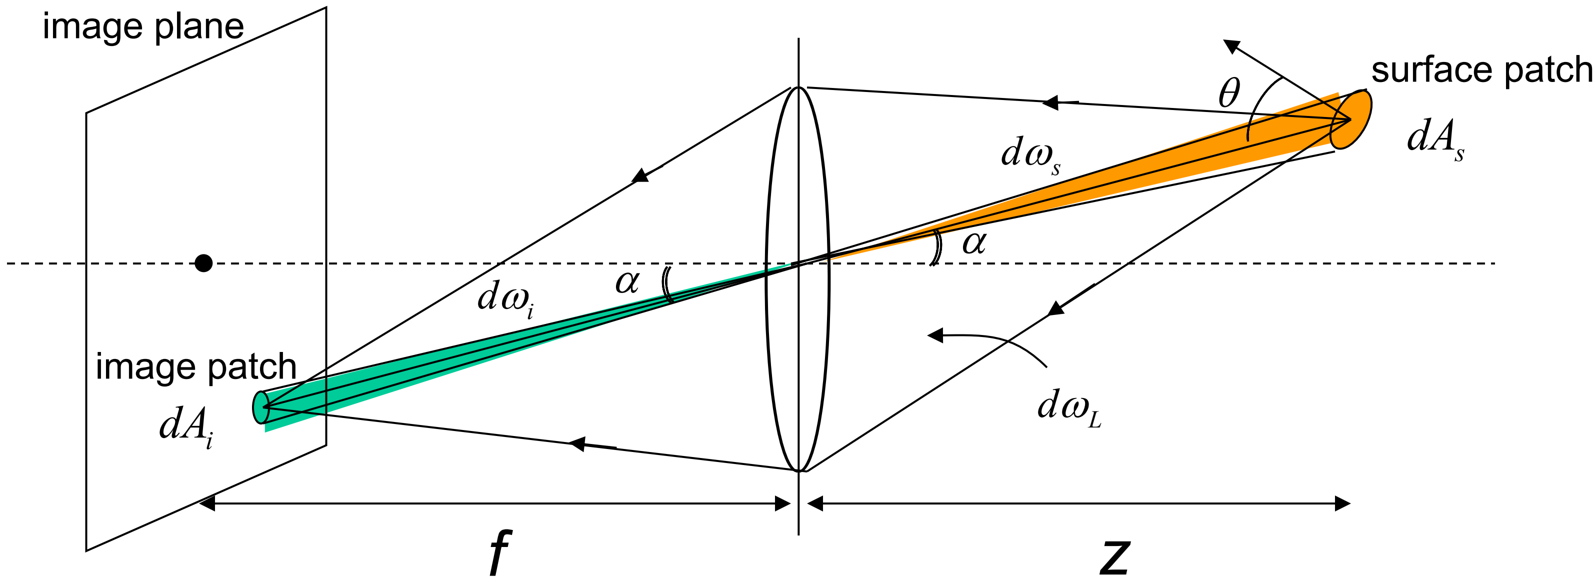

- Flux received by lens from dA_s = Flux projected onto image dA_i

$$\boxed{L\,(dA_s \cos\theta)\,d\omega_L \;=\; E\,dA_i} \longrightarrow (3)$$

- From (1), (2), and (3):

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

- Image irradiance is proportional to Scene Radiance!
- Small field of view → Effects of 4th power of cosine are small.

Slide from
S Narasimhan.

From light rays to pixel values

$$X = E \cdot \Delta t$$

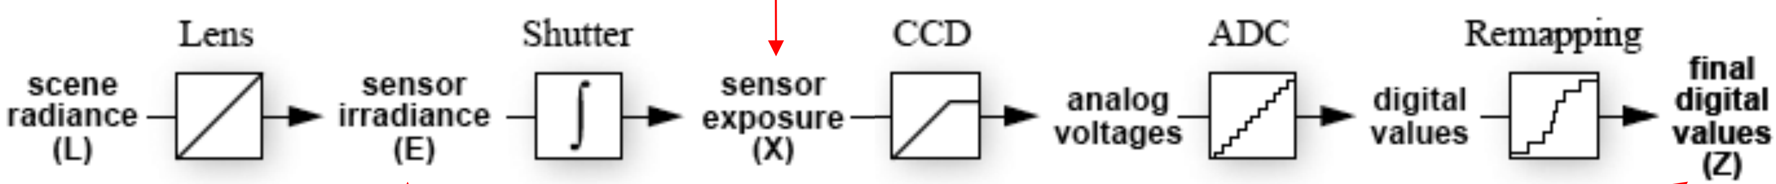

$$E = \left[\frac{\pi}{4} \left(\frac{d}{f} \right)^2 \cos^4 \alpha \right] L$$

$$Z = f(E \cdot \Delta t)$$

- Camera response function: the mapping f from irradiance to pixel values
  - Useful if we want to estimate material properties
  - Enables us to create *high dynamic range (HDR) images*
  - Classic reference: P. E. Debevec and J. Malik, *Recovering High Dynamic Range Radiance Maps from Photographs*, *SIGGRAPH 97*

Basic models of reflection

Specular: light bounces off at the incident angle

- E.g., mirror

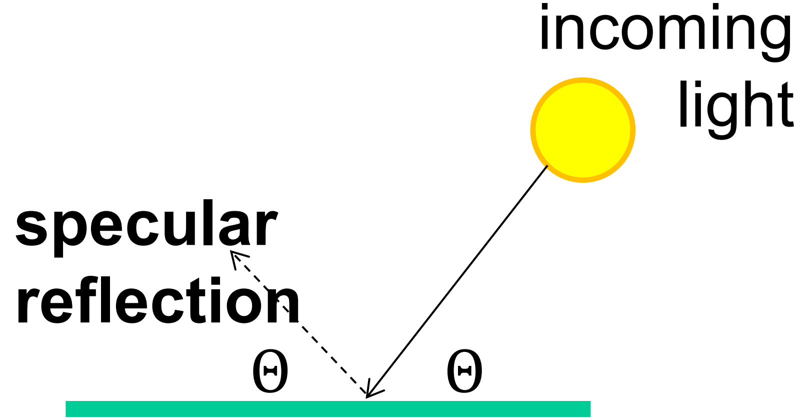

Diffuse: light scatters in all directions

- E.g., brick, cloth, rough wood

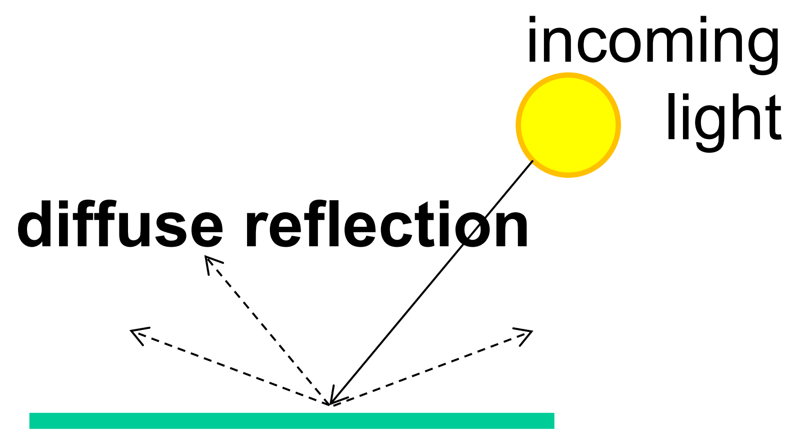

Slide from D. Hoiem

Other possible effects

transparency

light source

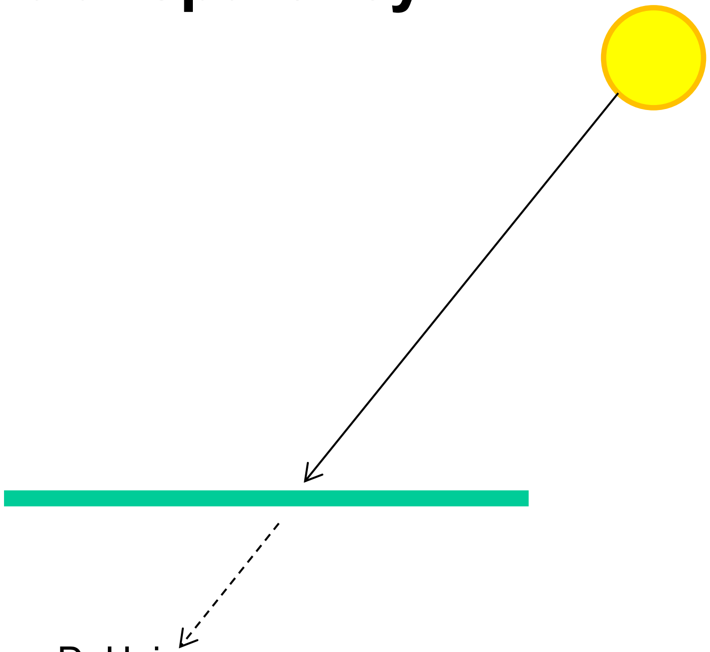

refraction

light source

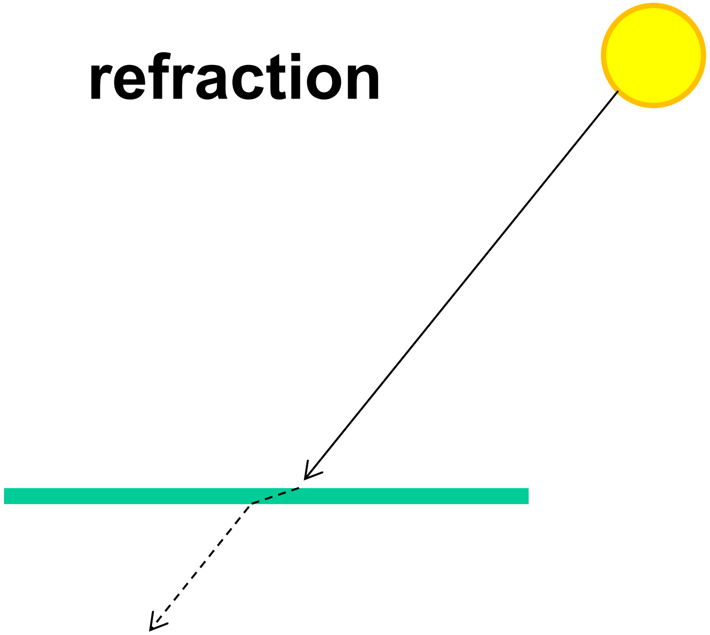

Slide from D. Hoiem

Other possible effects

subsurface scattering

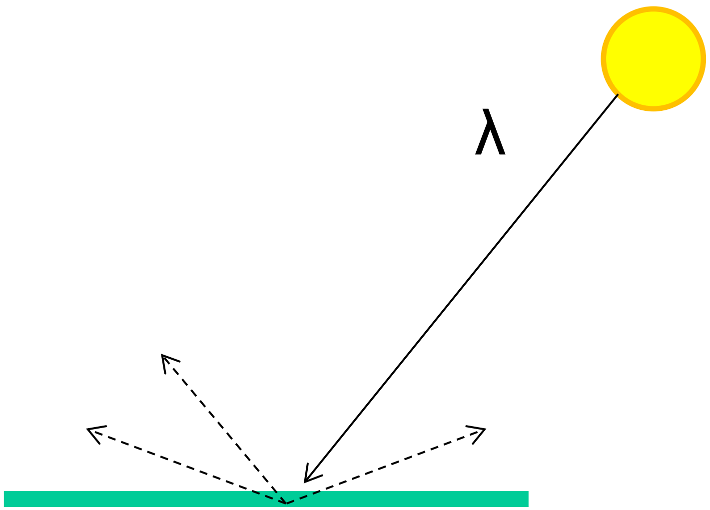

Slide from D. Hoiem

https://en.wikipedia.org/wiki/Subsurface_scattering#/media/File:Skin_Subsurface_Scattering.jpg

Other possible effects

fluorescence

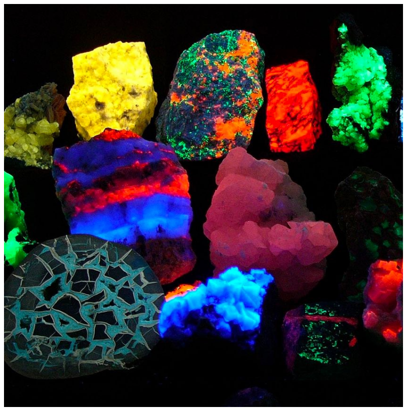

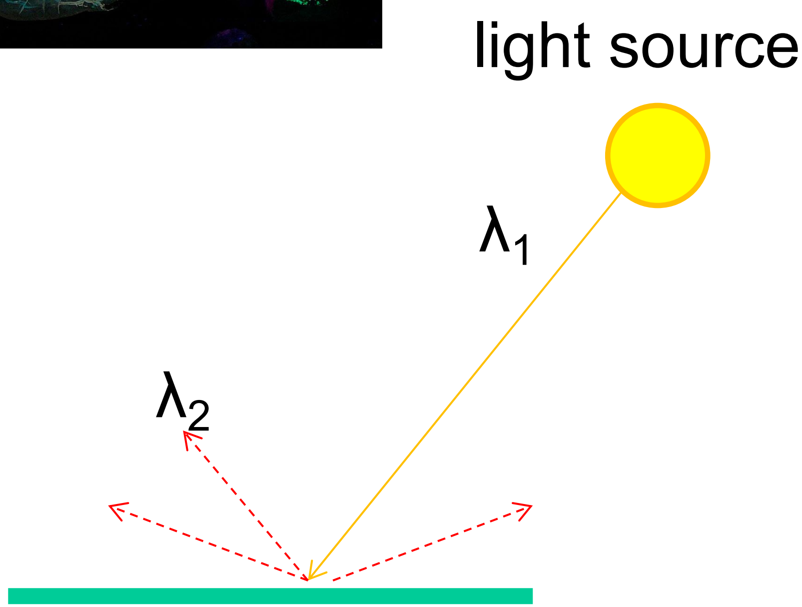

phosphorescence

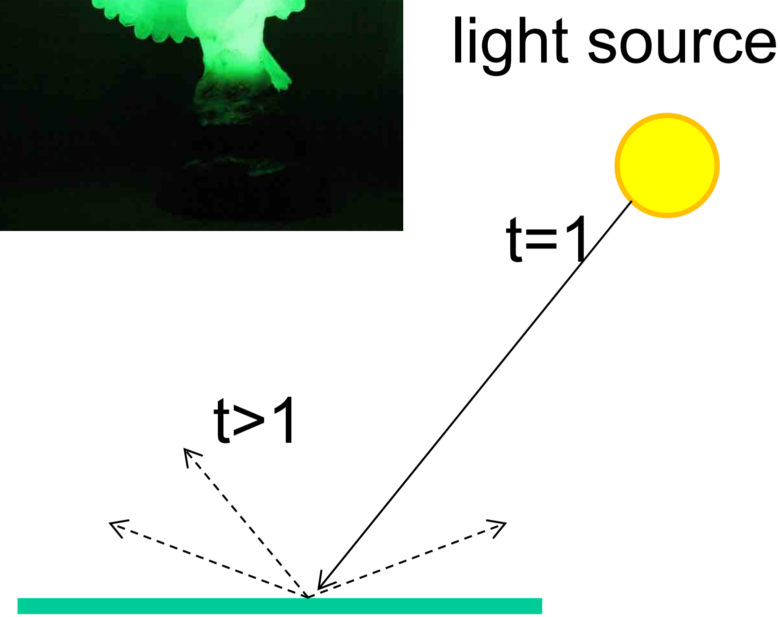

Overview

- 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

Lambertian reflectance model

Some light is absorbed (function of albedo ρ)

Remaining light is scattered (diffuse reflection)

Examples: soft cloth, concrete, matte paints

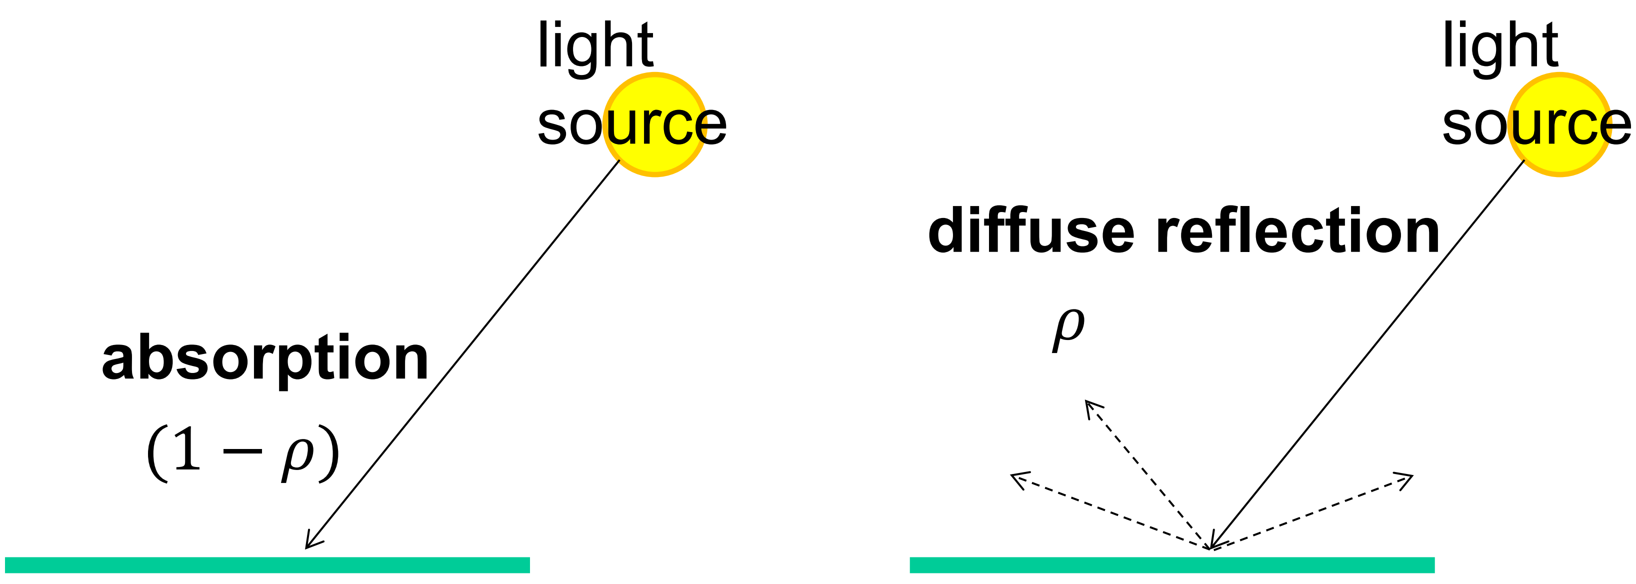

Slide from D. Hoiem

Specular Reflection

Reflected direction depends on light orientation and surface normal

- E.g., mirrors are fully specular

Flickr, by suzysputnik

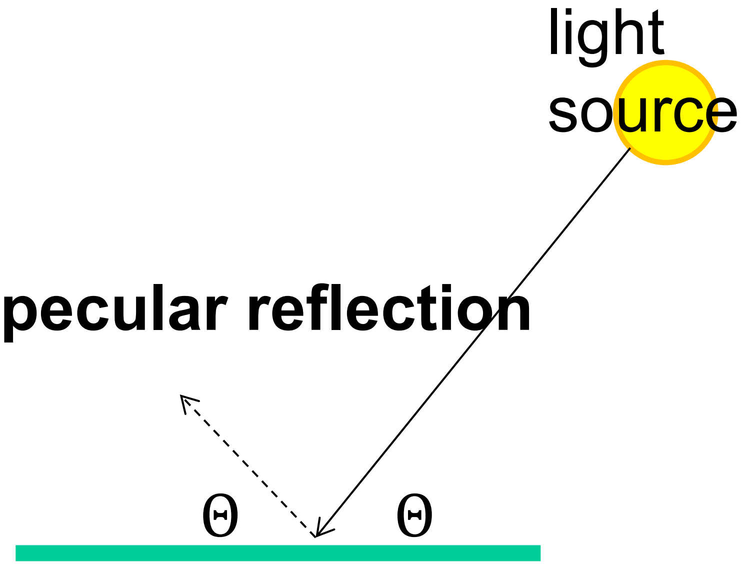

Flickr, by piratejohnny

Most surfaces have both specular and diffuse components

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

Photo: northcountryhardwoodfloors.com

Typically, specular component is small

Slide from D. Hoiem

BRDF: Bidirectional Reflectance Distribution Function

- Model of local reflection that tells 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

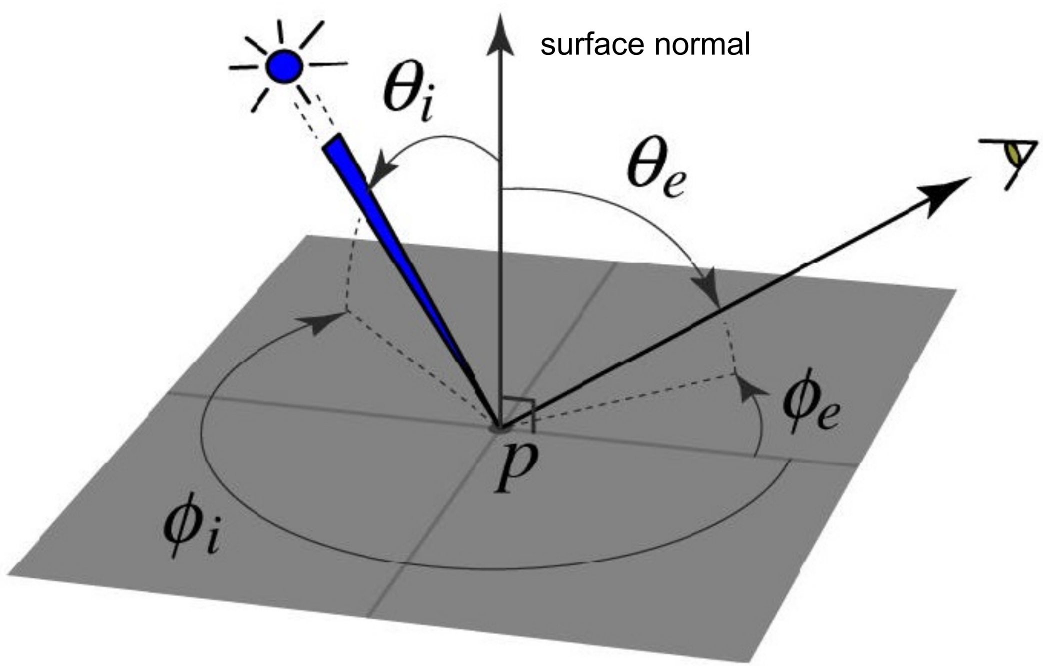

Slide credit: S. Savarese

BRDFs can be incredibly complicated…








Diffuse reflection

- Light is reflected equally in all directions
  - Dull, matte surfaces like chalk or latex paint
  - Microfacets scatter incoming light randomly

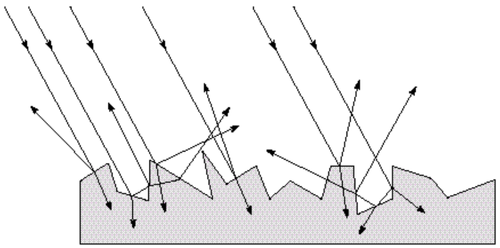

- Brightness of the surface depends on the incidence of illumination

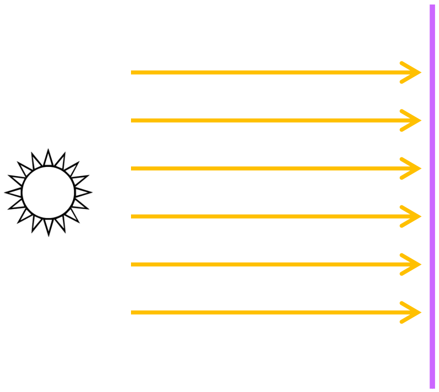

brighter

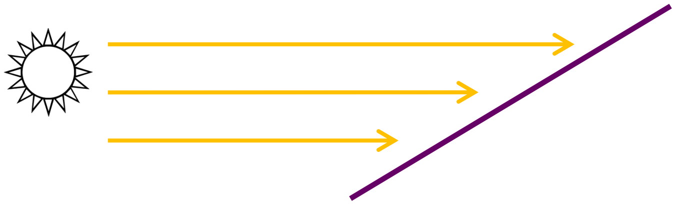

darker

Diffuse reflection: Lambert's law

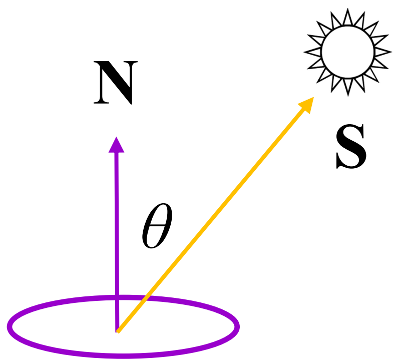

$$B = \rho \, \mathbf{N} \cdot \mathbf{S}$$
$$= \rho \left\| \mathbf{S} \right\| \cos\theta$$

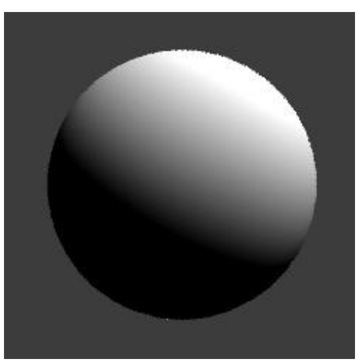

- *B*: radiosity (total power leaving the surface per unit area)
- ρ: albedo (fraction of incident irradiance reflected by the surface)
- *N*: unit normal
- *S*: source vector (magnitude proportional to intensity of the source)

Slide from L. Lazebnik

Specular reflection

- Radiation arriving along a source direction leaves along the **specular direction** (source direction reflected about normal)
- On real surfaces, energy usually goes into a lobe of directions
- **Phong model**: reflected energy falls of with $\cos^n(\delta\theta)$

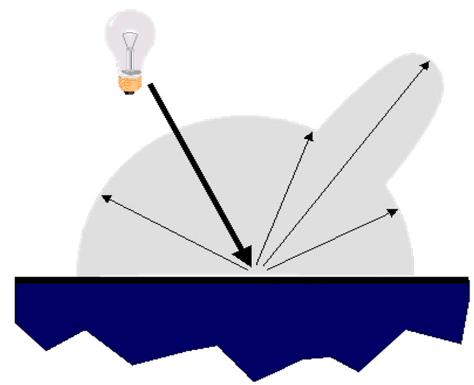

Moving the light source

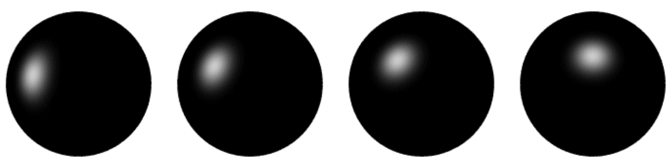

Changing the exponent

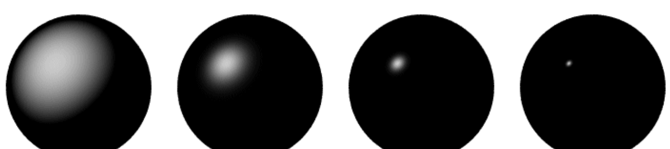

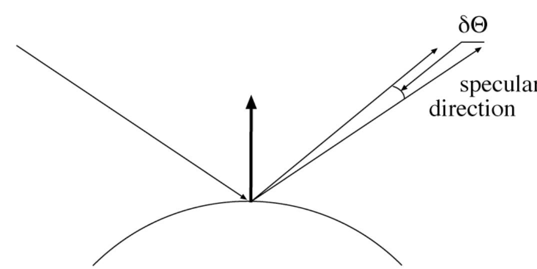

Slide from L. Lazebnik

Specular reflection

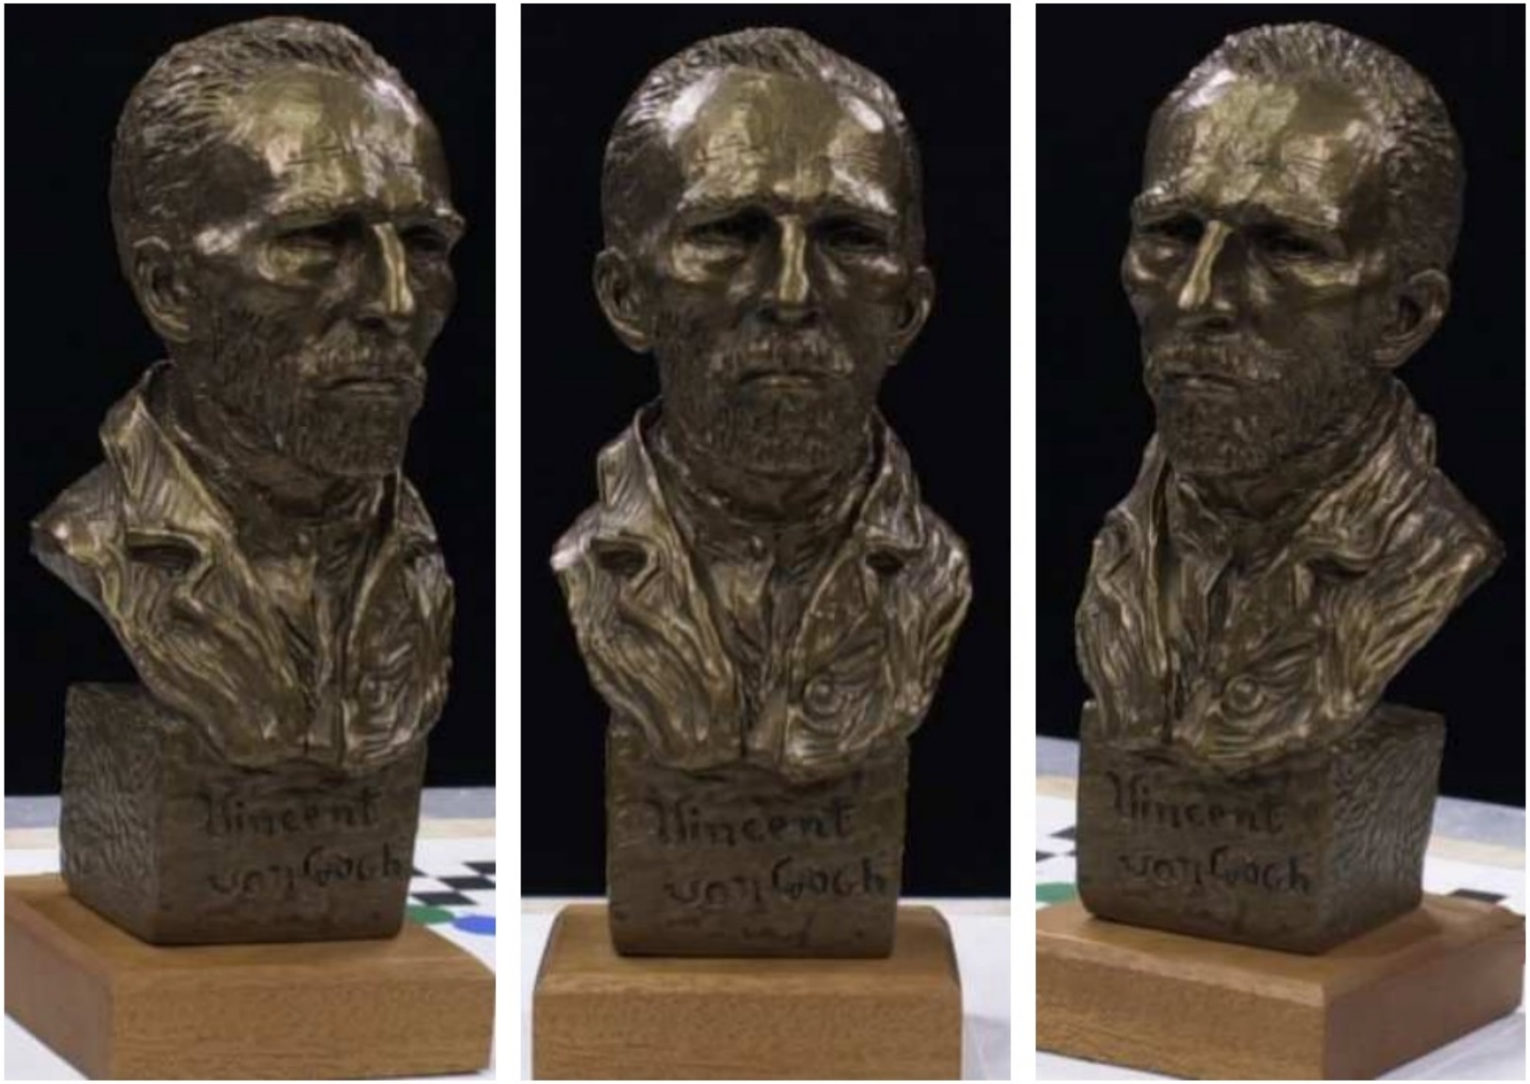

Picture source

Slide from L. Lazebnik

Lambertian + Specular Model

- I(x) = Ambient Term + Diffuse Term + Specular Term

Photometric stereo (shape from shading)

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

Luca della Robbia,
Cantoria, 1438

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

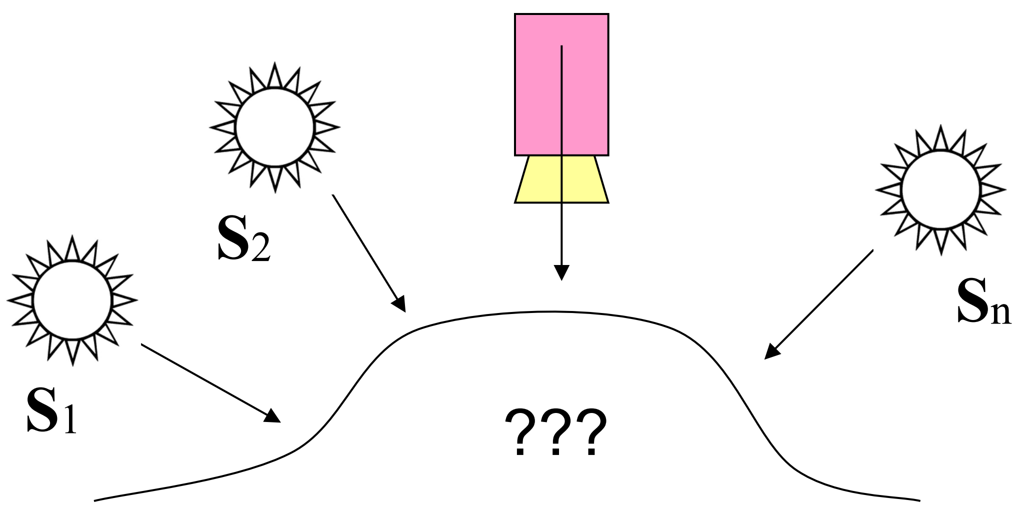

Slide from L. Lazebnik

F&P 2nd ed., sec. 2.2.4

Example 1

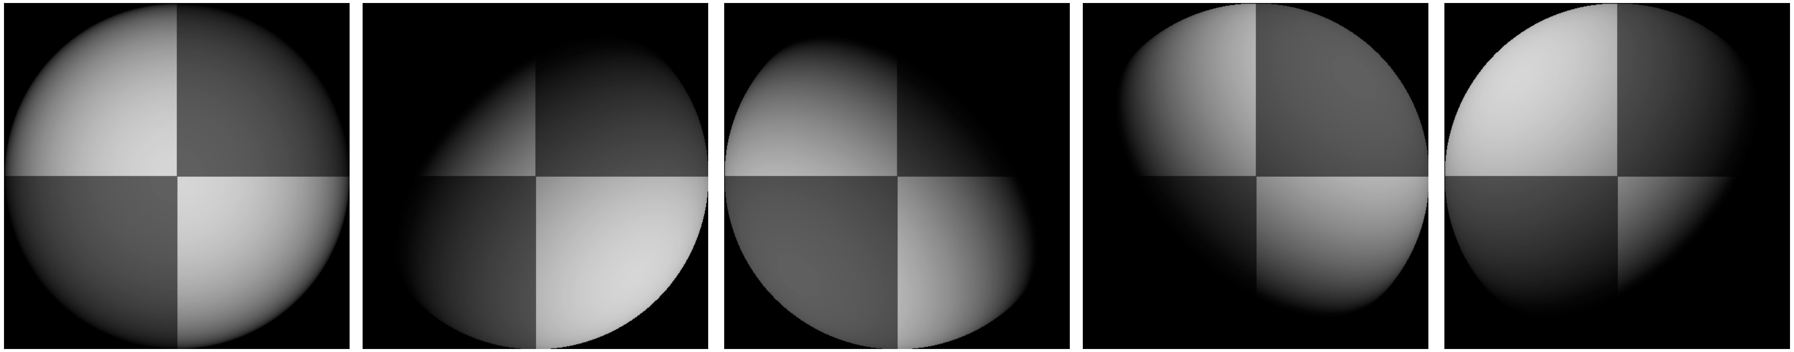

Recovered albedo

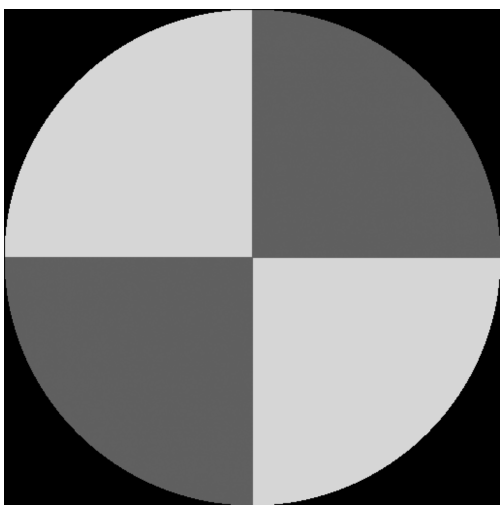

Recovered normal field

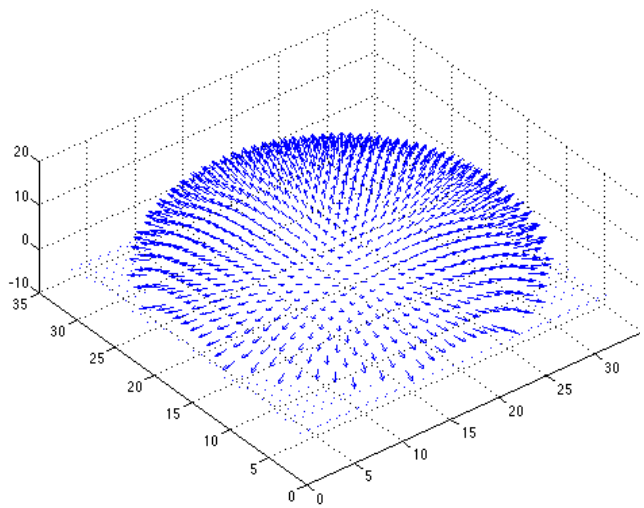

Recovered surface model

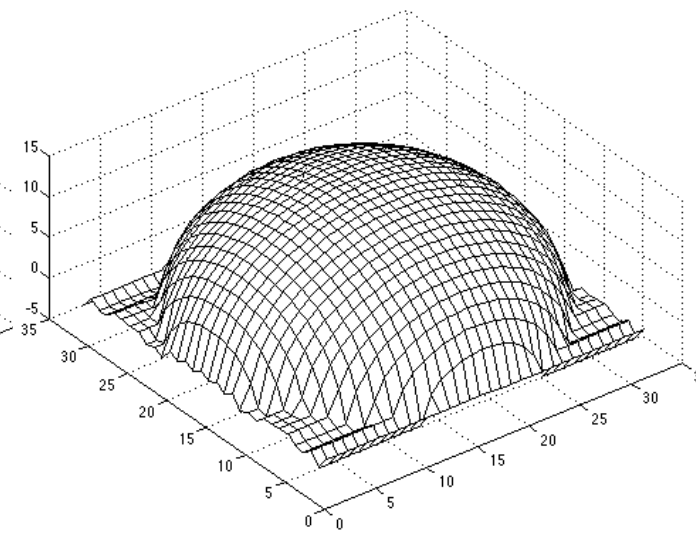

Slide from L. Lazebnik

F&P 2nd ed., sec. 2.2.4

Example 2

Input

Recovered albedo

Recovered normal field

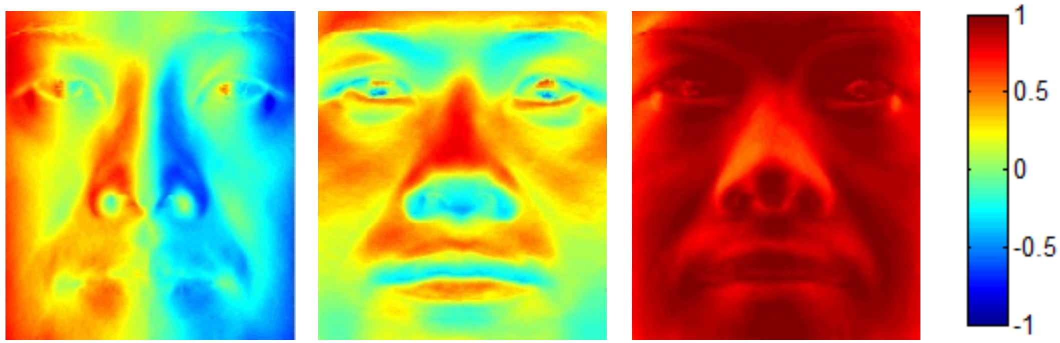

x y z

Recovered surface model

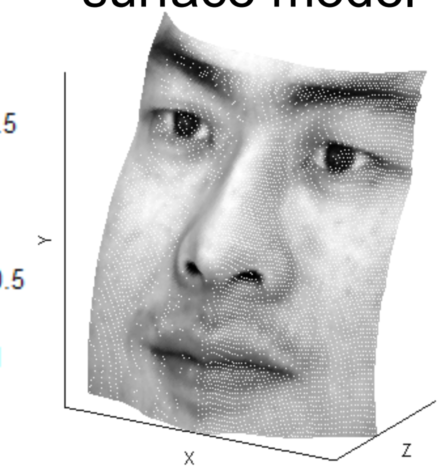

Slide from L. Lazebnik

Image model

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

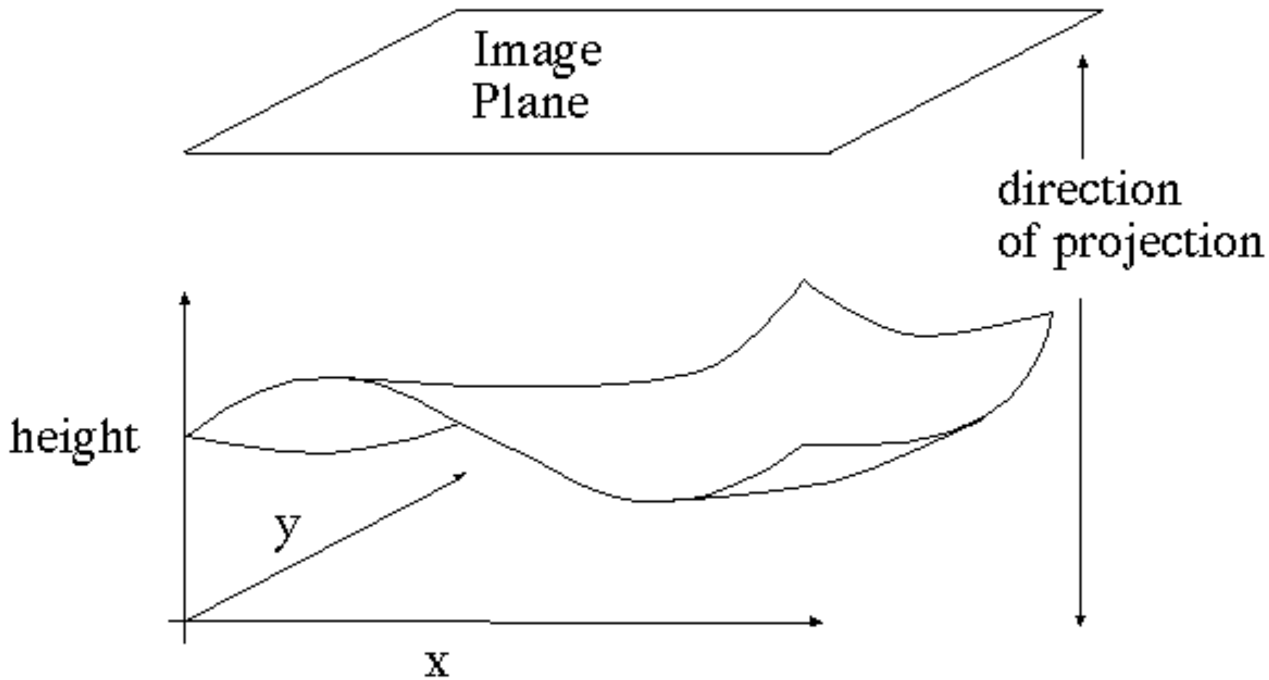

Slide from L. Lazebnik

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

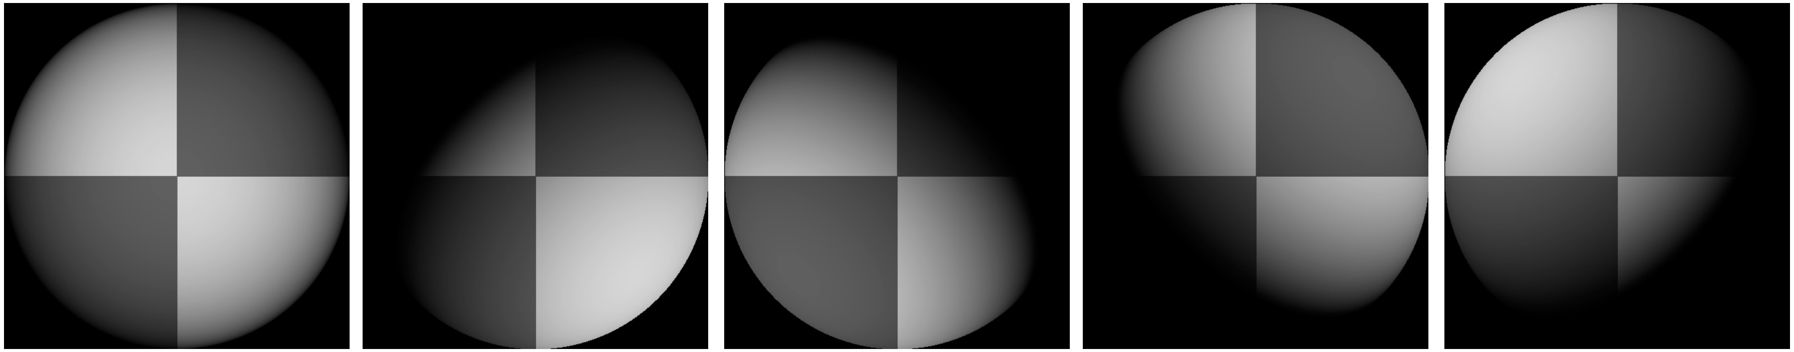

Recovered albedo

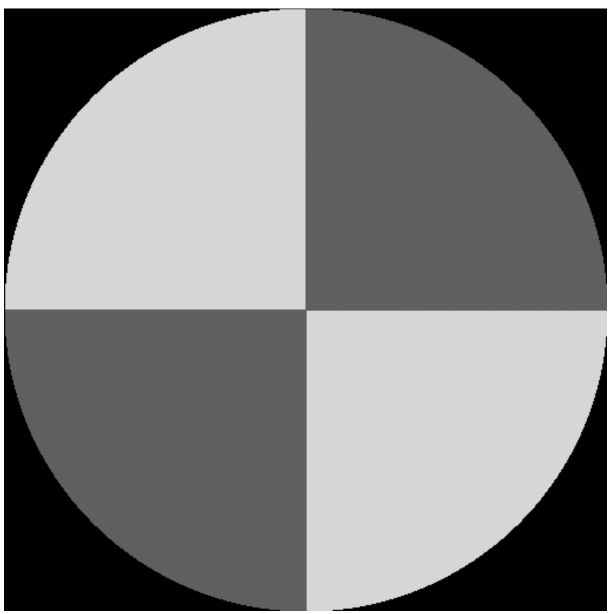

Recovered normal field

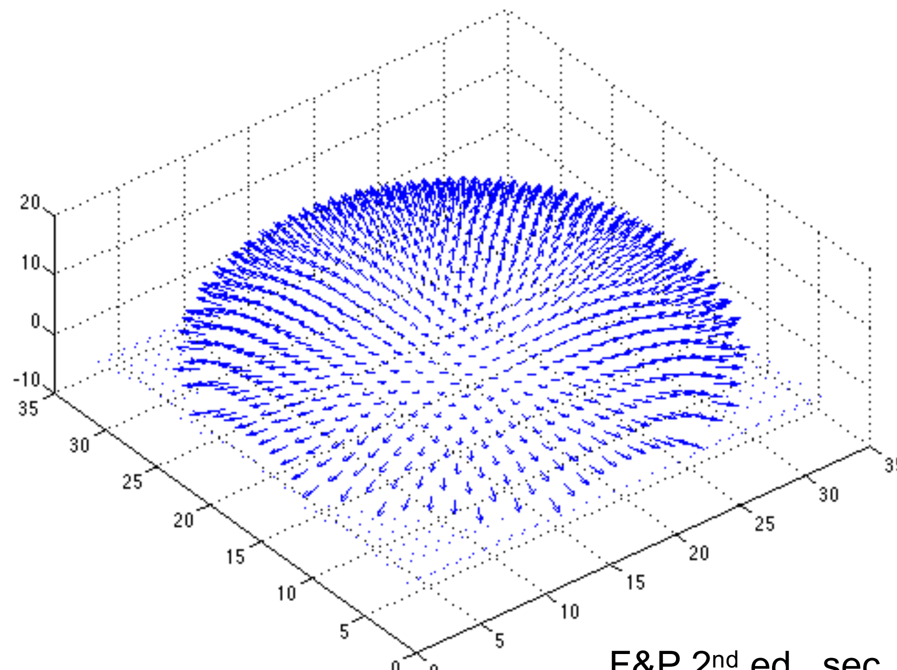

F&P 2nd ed., sec. 2.2.4

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:

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

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

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)

Slide from L. Lazebnik

F&P 2nd ed., sec. 2.2.4

Surface recovered by integration

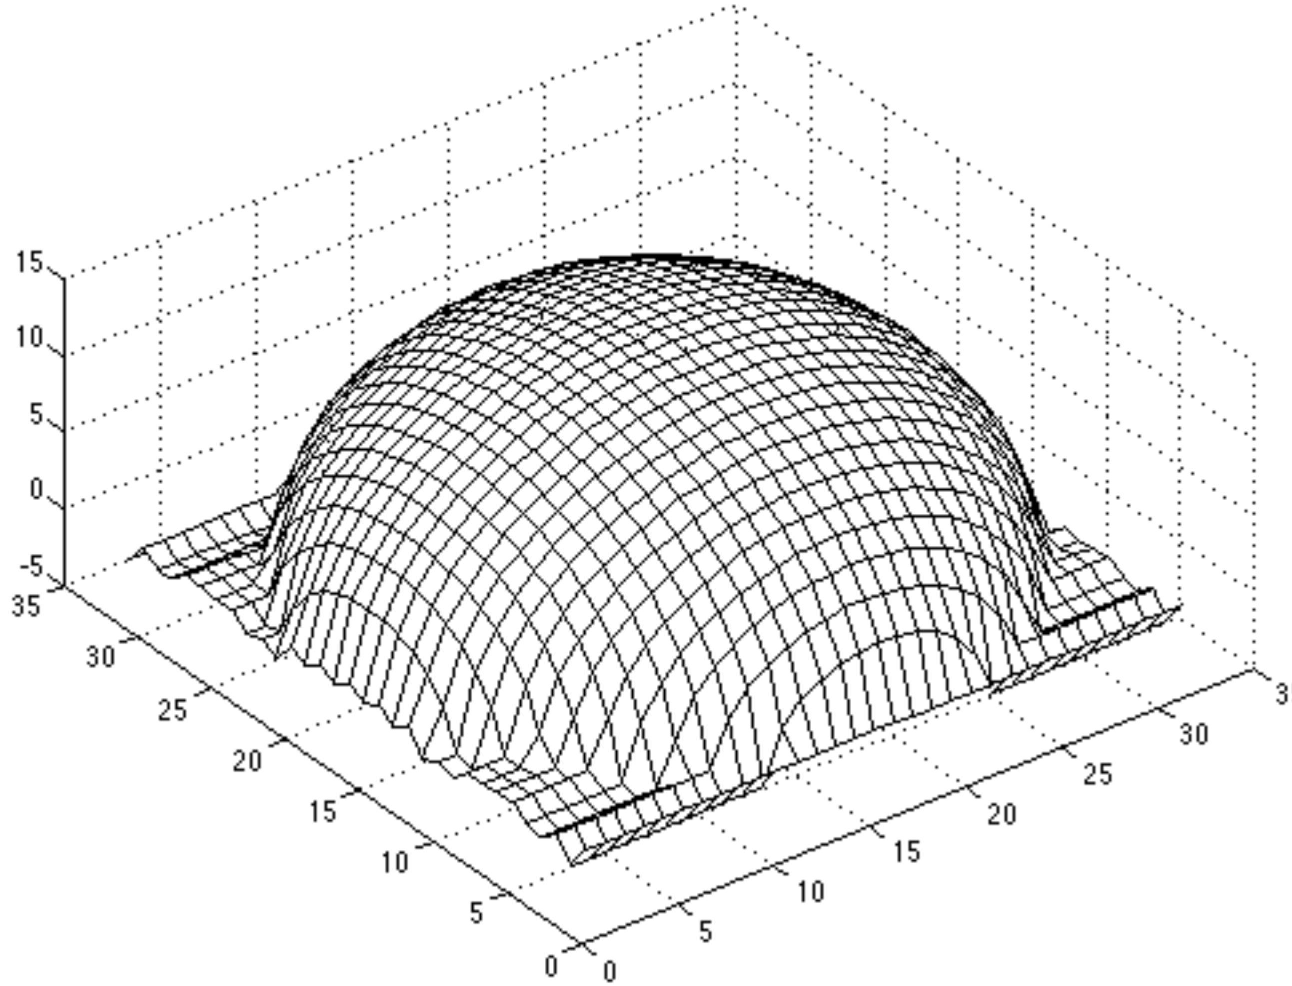

Slide from L. Lazebnik

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

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

Full 3D case:

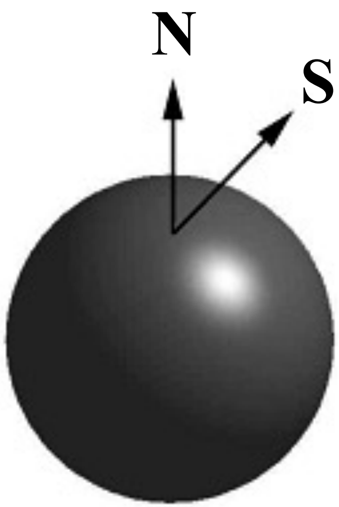

$$\begin{pmatrix} N_x(x_1,y_1) & N_y(x_1,y_1) & N_z(x_1,y_1) \\ N_x(x_2,y_2) & N_y(x_2,y_2) & N_z(x_2,y_2) \\ \vdots & \vdots & \vdots \\ N_x(x_n,y_n) & N_y(x_n,y_n) & N_z(x_n,y_n) \end{pmatrix} \begin{pmatrix} S_x \\ S_y \\ S_z \end{pmatrix} = \begin{pmatrix} I(x_1,y_1) \\ I(x_2,y_2) \\ \vdots \\ I(x_n,y_n)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

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

Finding the direction of the light source

Consider points on the *occluding contour*:

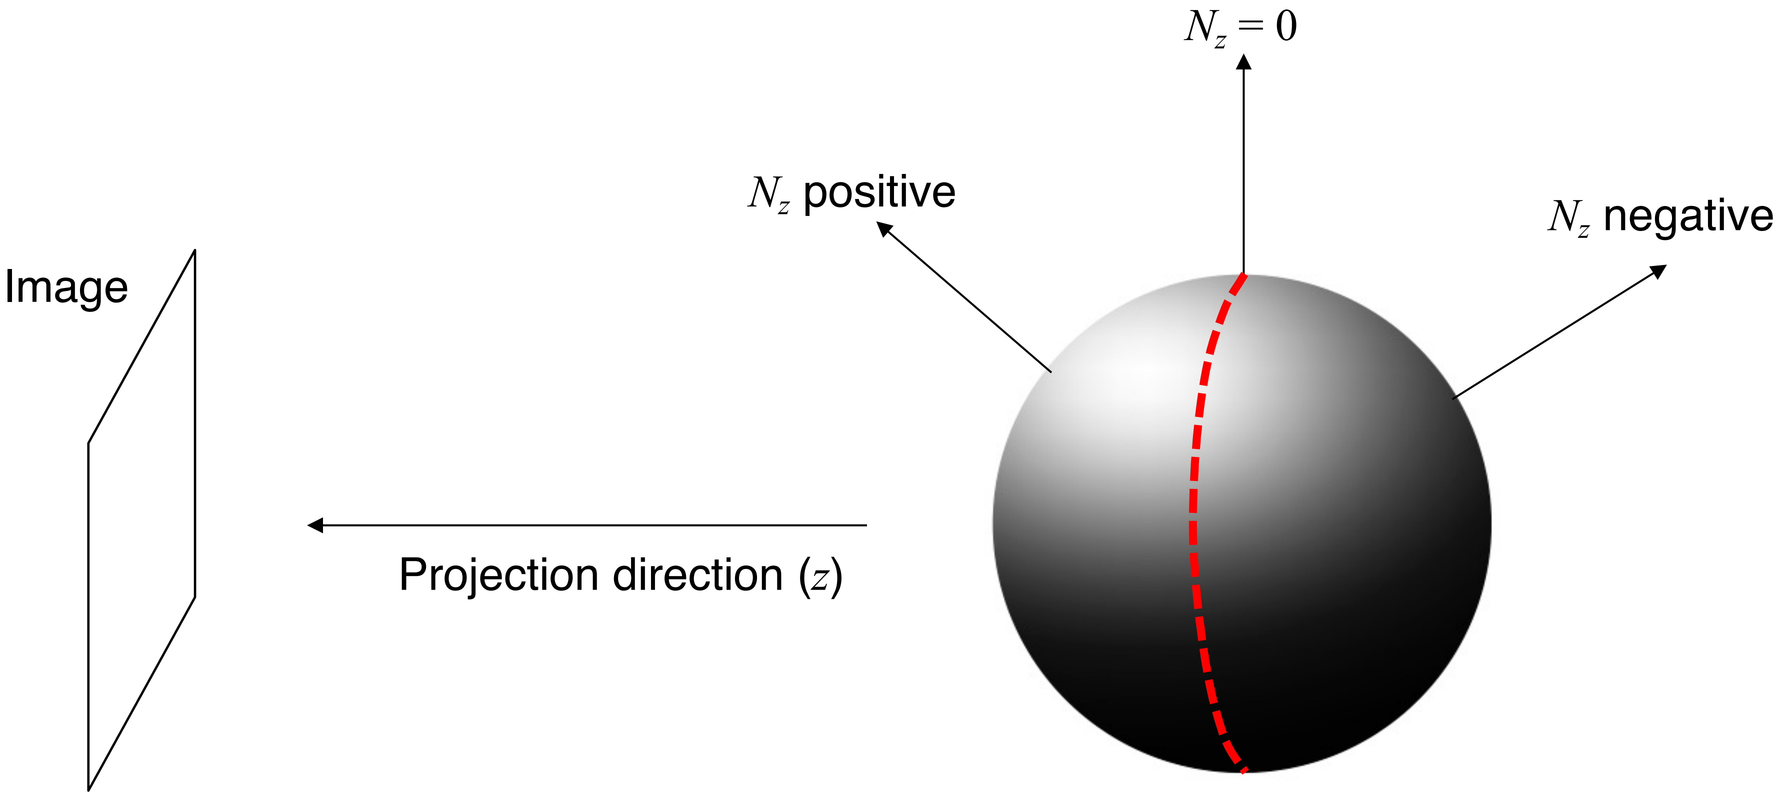

P. Nillius and J.-O. Eklundh, “Automatic estimation of the projected light source direction,” CVPR 2001

Finding the direction of the light source

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

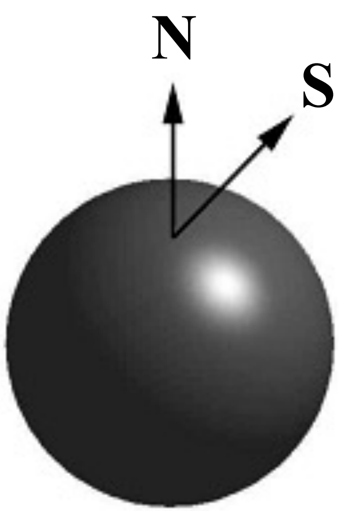

Full 3D case:

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

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

Finding the direction of the light source

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

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.