Photometric Image Formation

Introduction to Computer Vision I

CSE 152A

Lecture 3

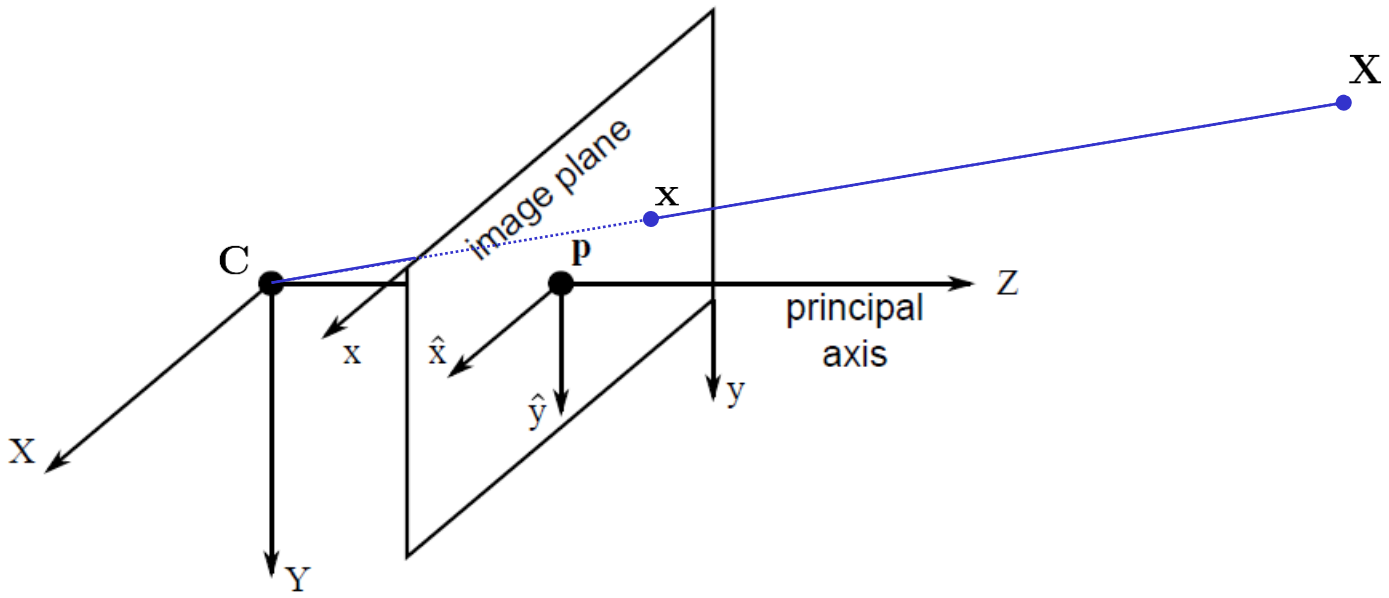
Announcements

- Assignment 0 is due Oct 7, 11:59 PM
- Assignment 1 will be released Oct 7
  - Due Oct 21, 11:59 PM

Image Formation: Outline

- Geometric image formation
  - How do 3D world points project to 2D image points?
- Photometric image formation
  - What color is the projected point?

Geometric image formation



The projective camera

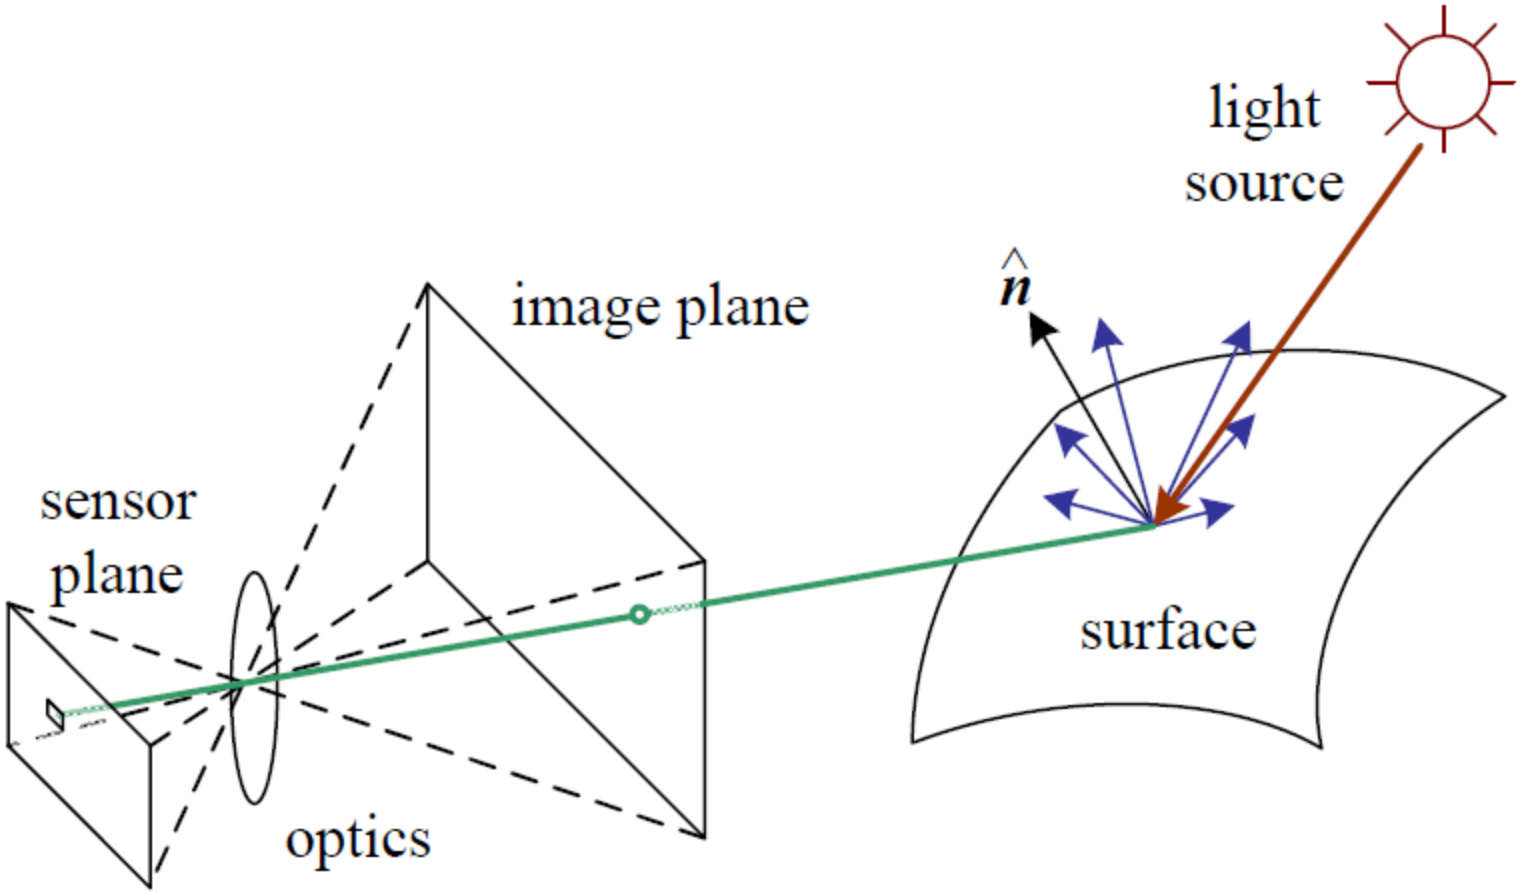
- Extrinsic parameters: Since the camera coordinate frame may not align with the world coordinate frame, there is a 3D Euclidean transformation from world coordinates to camera coordinates
- Intrinsic parameters: Since the scene units (e.g., cm) differ from the pixel coordinate frame units (i.e., pixels) and origin (i.e., upper left pixel), there is a 2D affine transformation comprised of focal length in x and y directions, skew (which is 0 in real cameras), and principal point, all in terms of pixel dimensions

$$\begin{bmatrix} x \\ y \\ w \end{bmatrix} = \begin{bmatrix} \alpha_x & 0 & x_0 \\ 0 & \alpha_y & y_0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} r_{11} & r_{12} & r_{13} & t_X \\ r_{21} & r_{22} & r_{23} & t_Y \\ r_{31} & r_{32} & r_{33} & t_Z \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \\ T \end{bmatrix}$$

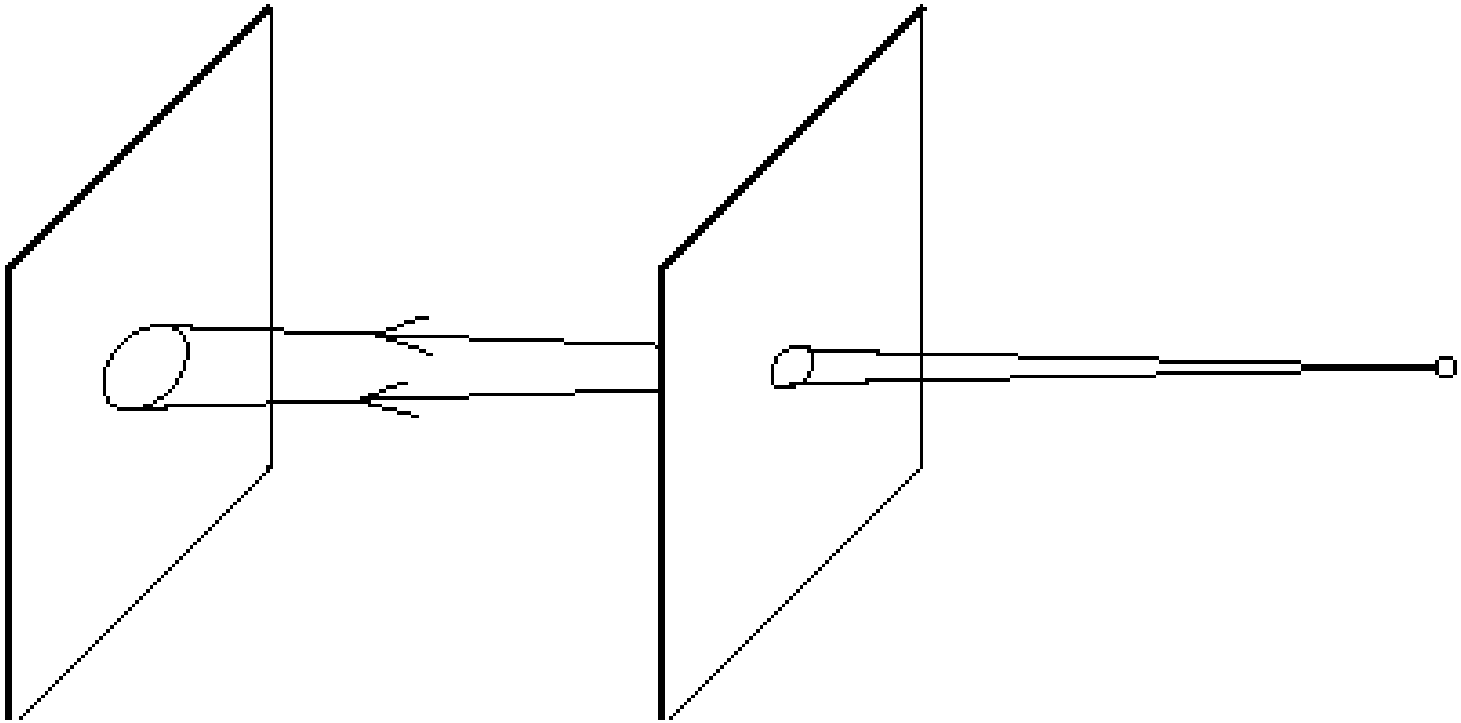
Intrinsic parameters

Extrinsic parameters

Photometric image formation



Beyond the pinhole Camera
Getting more light – Bigger Aperture



Pinhole Camera Images with Variable Aperture

2 mm

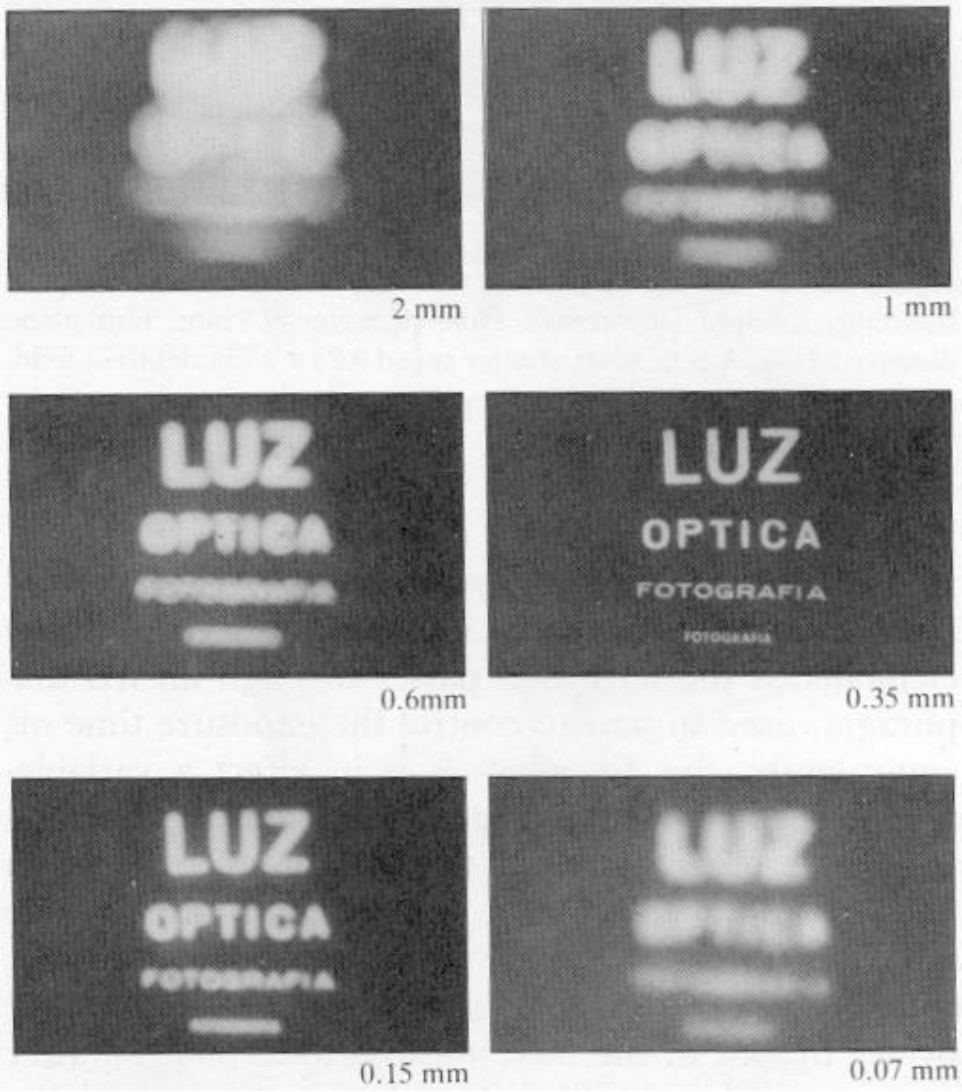
1mm

.6 mm

.35 mm

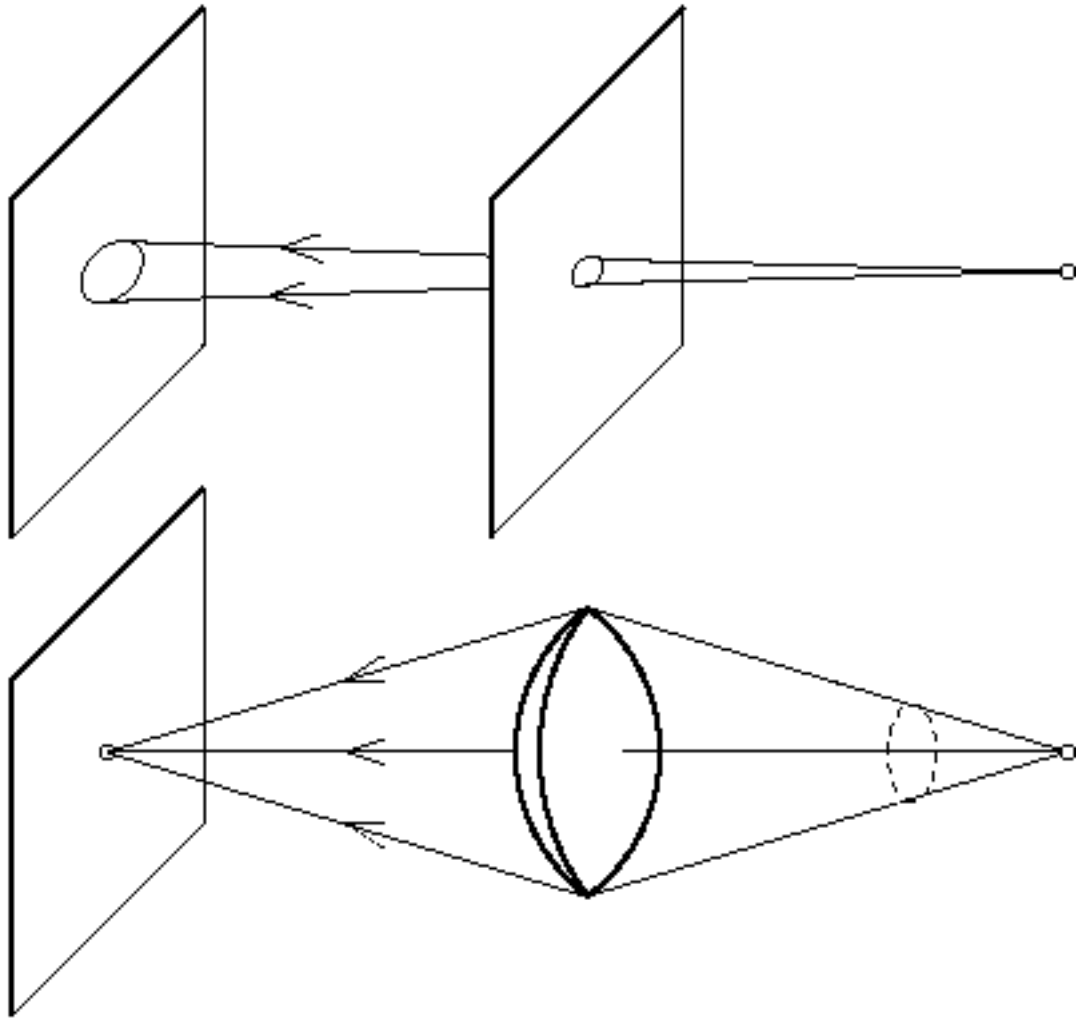
.15 mm

.07 mm

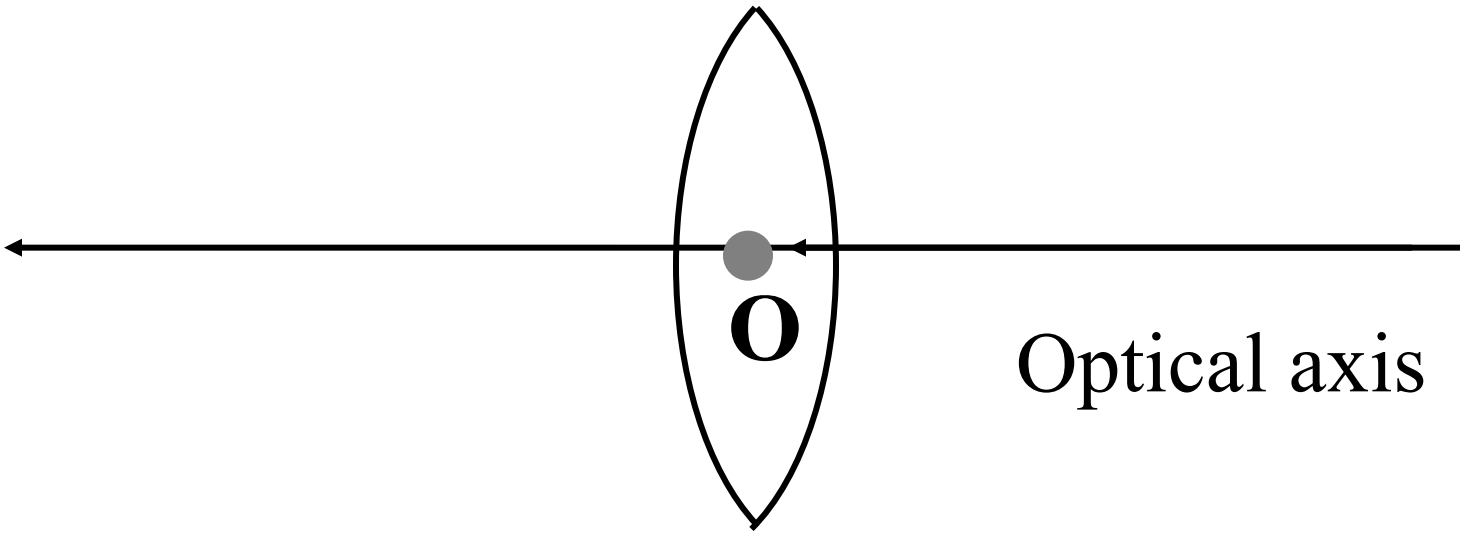


The reason for lenses

We need light, but big pinholes cause blur.

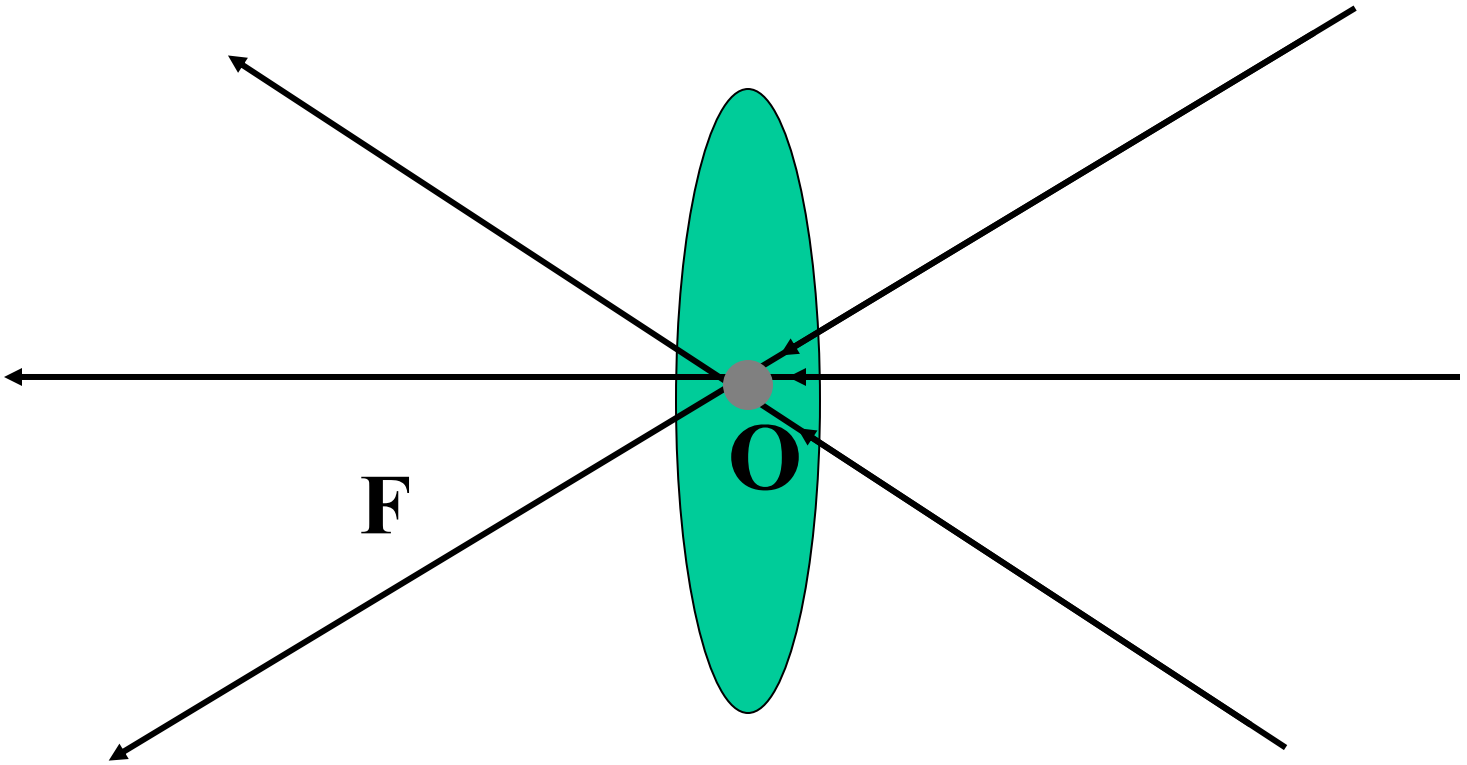


Thin Lens



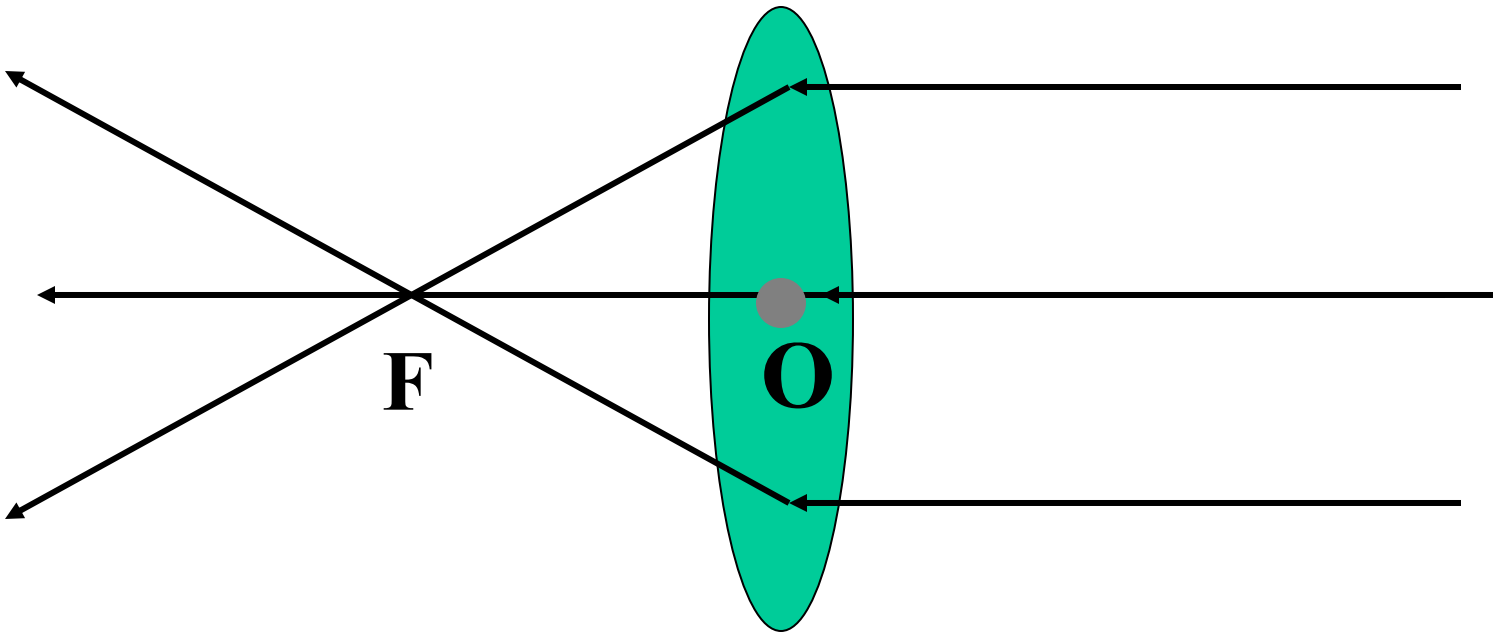
- Rotationally symmetric about optical axis
- Spherical interfaces

Thin Lens: Center



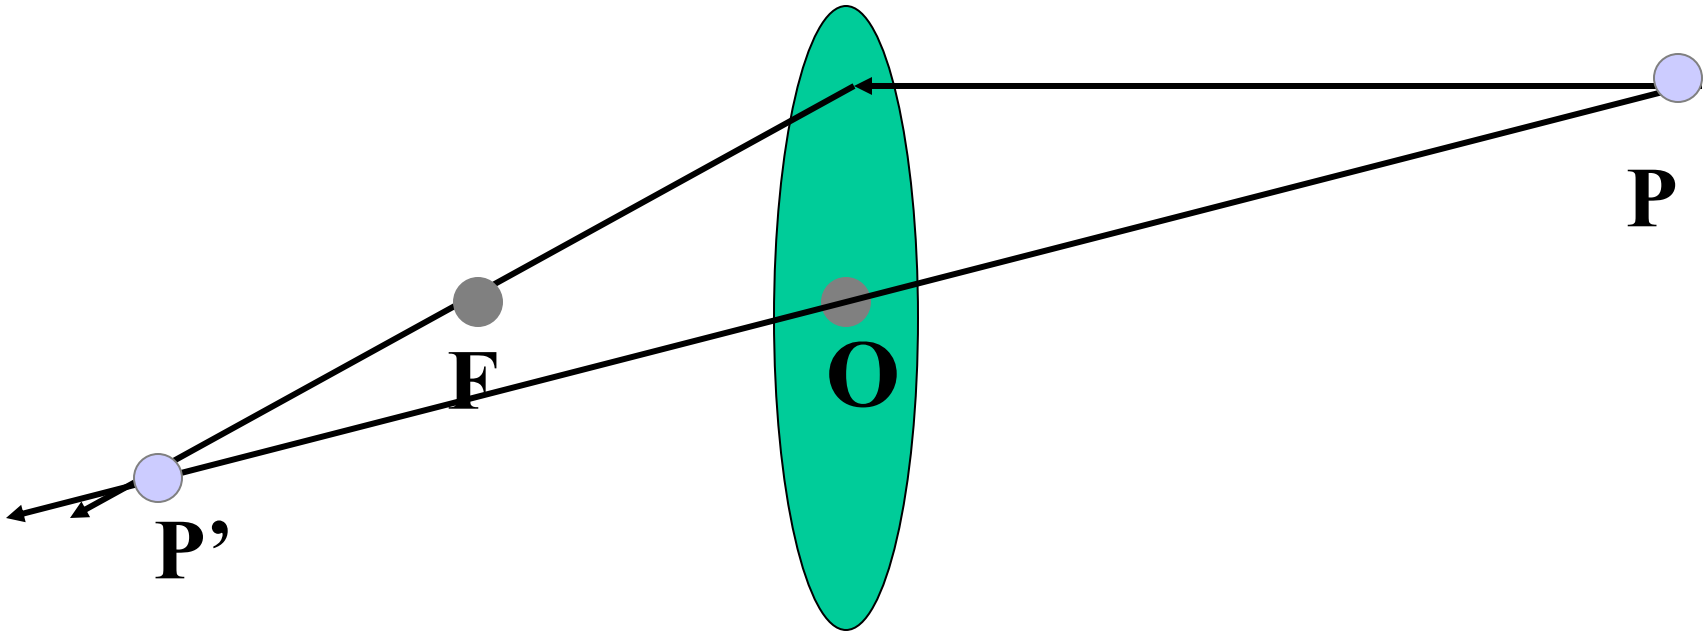
- All rays that enter lens along line pointing at **O** emerge in same direction

Thin Lens: Focus



Parallel lines pass through the focus F

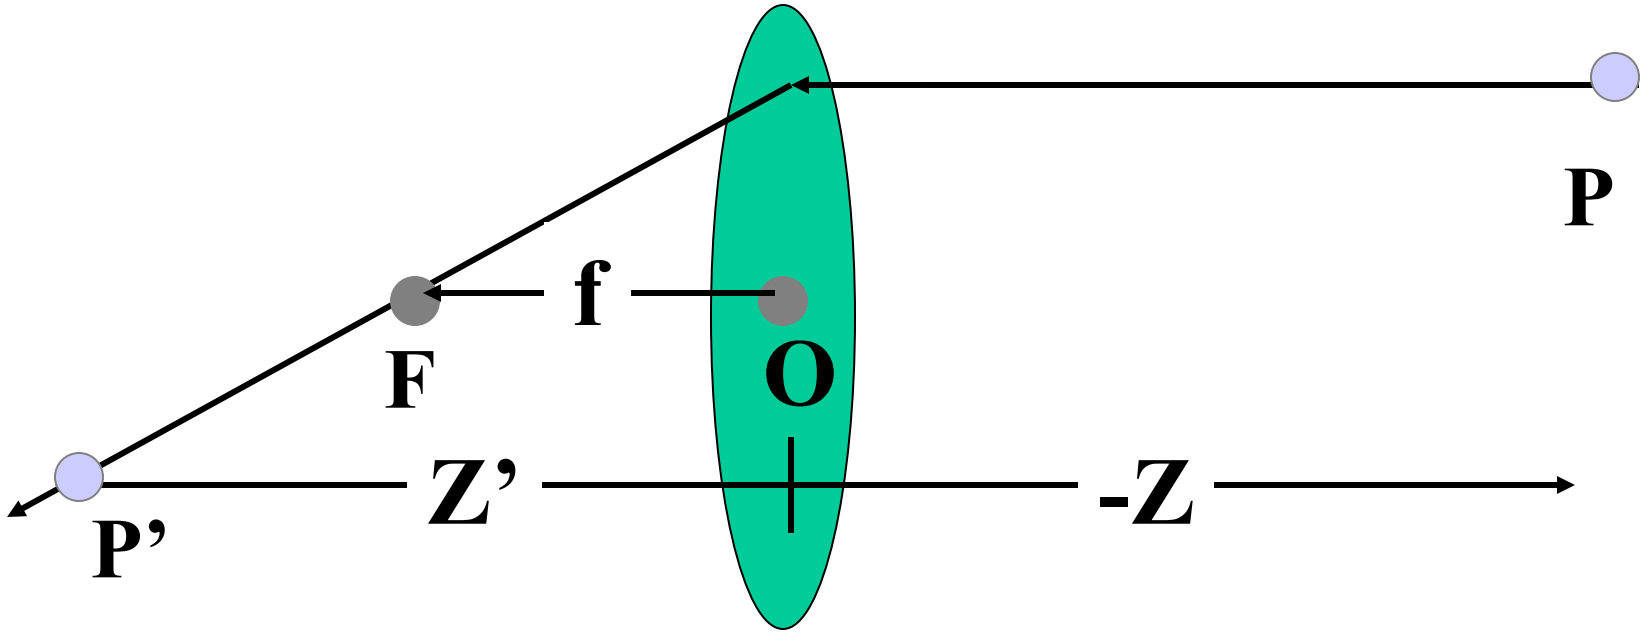
Thin Lens: Image of Point



All rays passing through lens and starting at **P** converge upon **P'**

So light gather capability of lens is given the area of the lens and all the rays focus on P' instead of become blurred like a pinhole

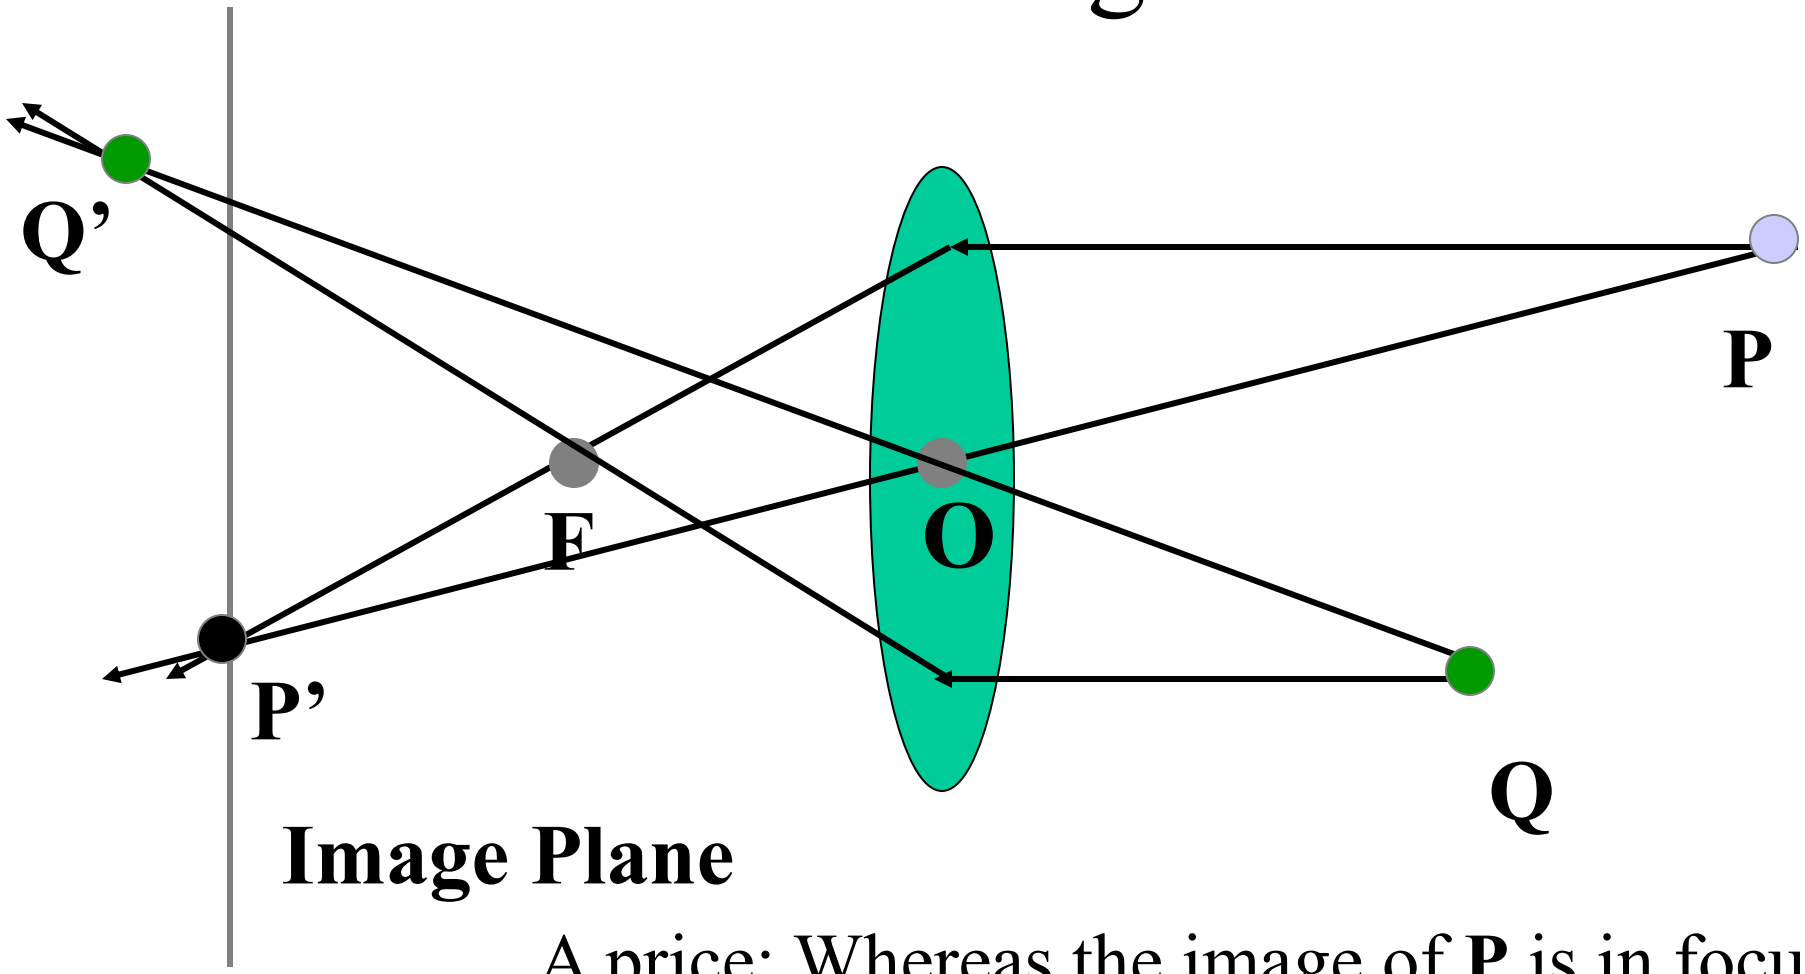
Thin Lens: Image of Point



$$\frac{1}{z'} - \frac{1}{z} = \frac{1}{f}$$

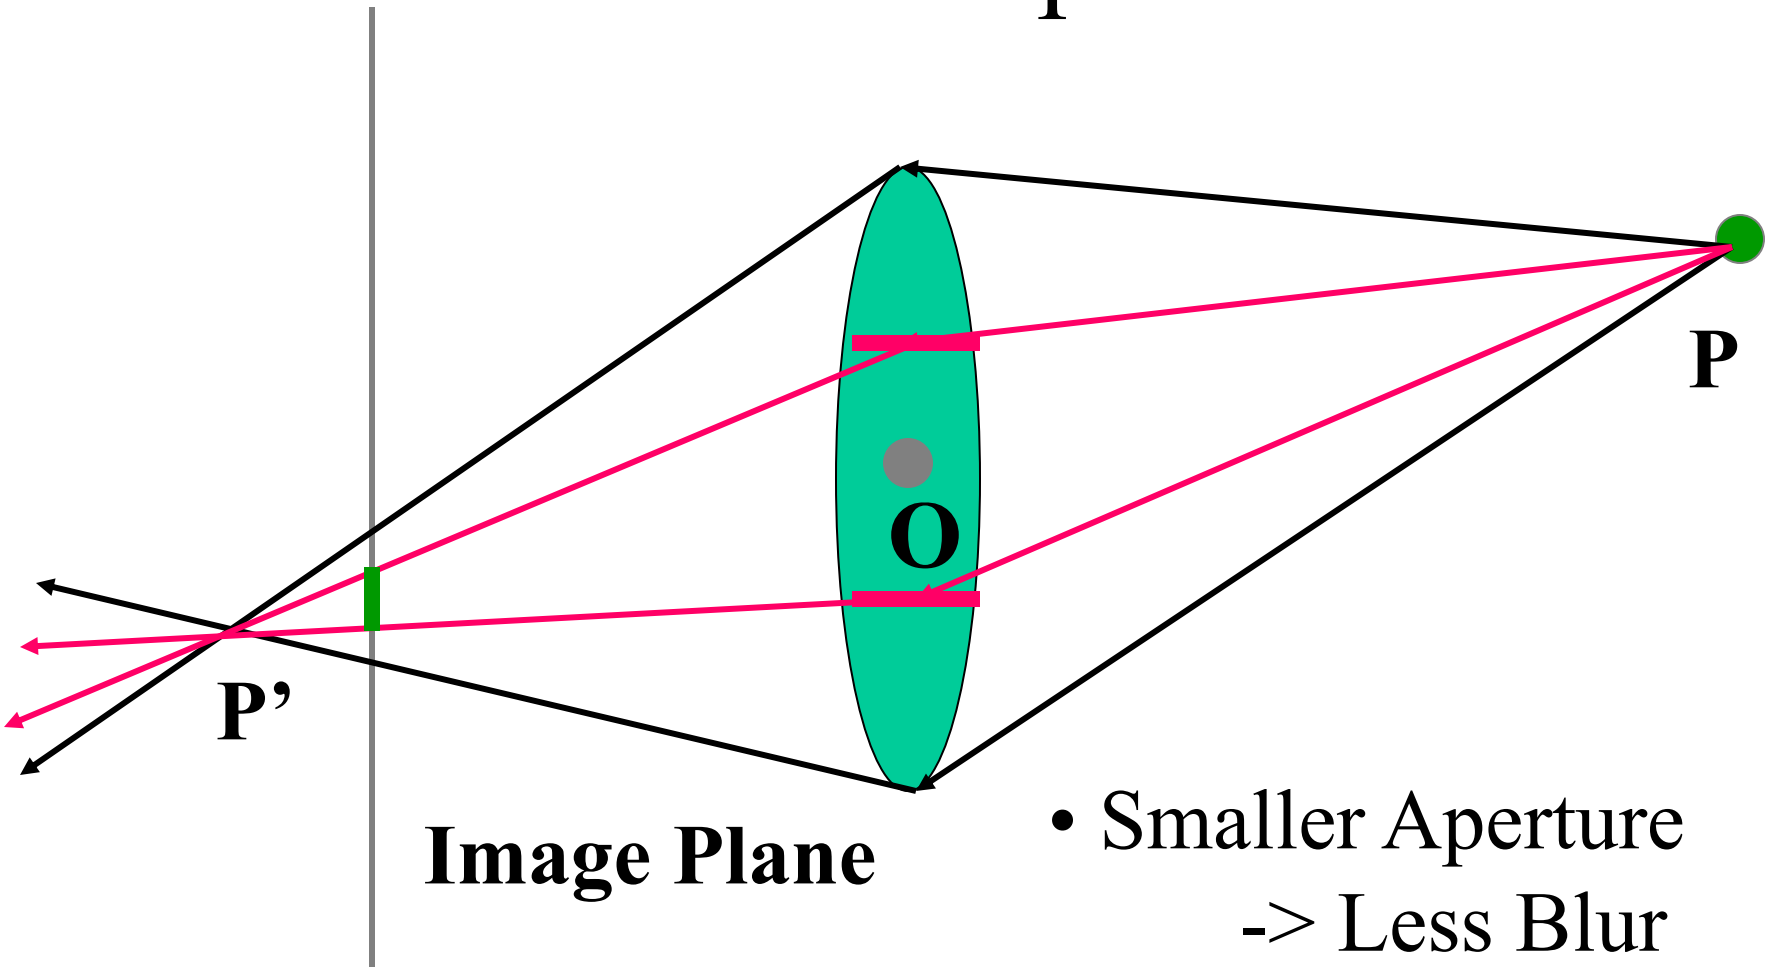
Relation between depth of Point (-Z) and the depth where it focuses (Z')

Thin Lens: Image Plane



A price: Whereas the image of **P** is in focus, the image of **Q** is not

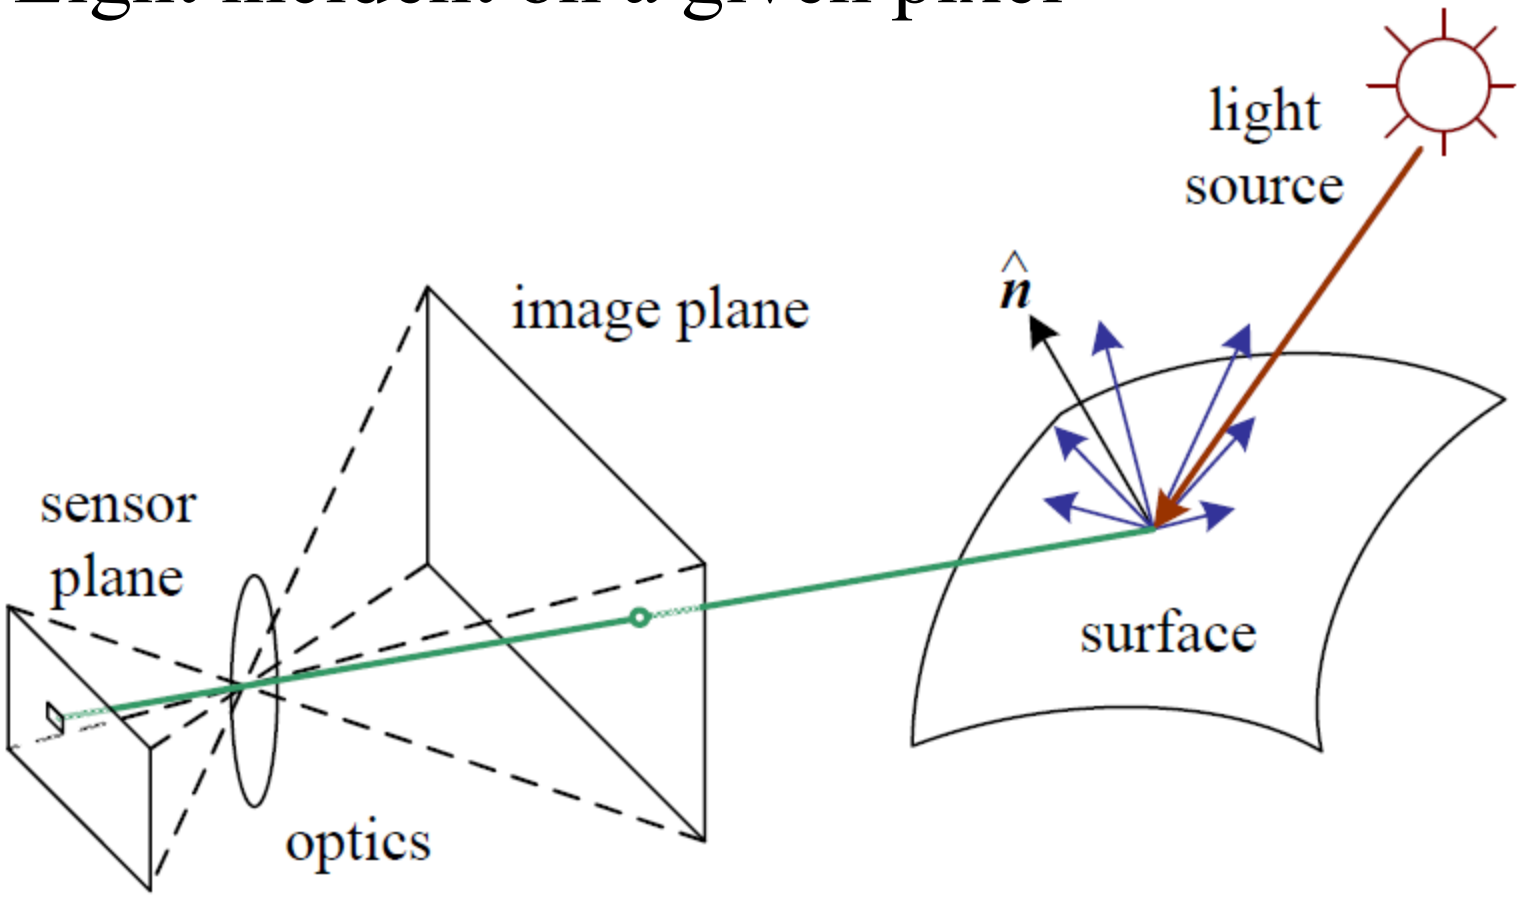
Thin Lens: Aperture



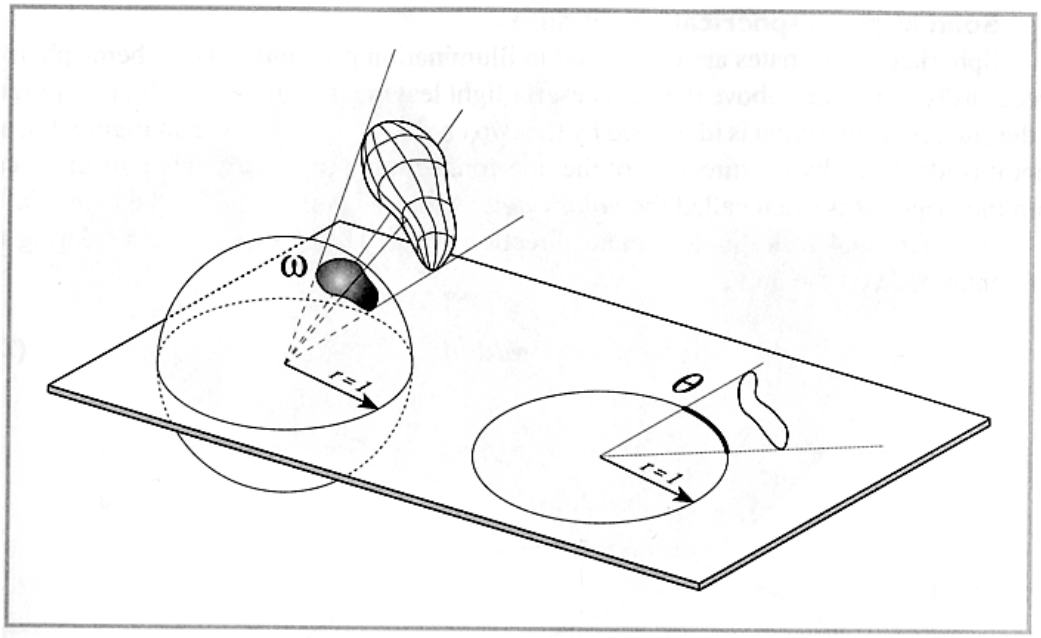
- Smaller Aperture
  -> Less Blur
- Pinhole -> No Blur

Photometric image formation

- Light incident on a given pixel



Measuring Angle



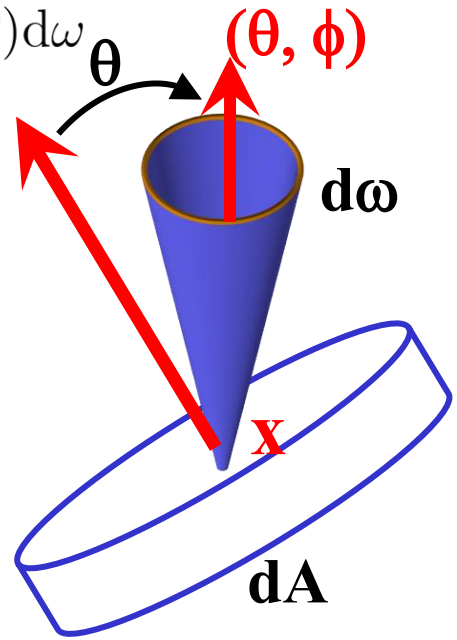
- The **solid angle** subtended by an object from a point P is the area of the projection of the object onto the unit sphere centered at P
- Definition is analogous to projected angle in 2D
- Measured in *steradians*, sr
- If I am at P and I look out, the solid angle tells me how much of my view is filled with an object

Radiance

- Power traveling at some point in a specified direction, per unit area perpendicular to the direction of travel, per unit solid angle
  - Units: watts per square meter per steradian, $W/m^2/sr = W\ m^{-2}\ sr^{-1}$

$$L(\mathbf{X}, \theta, \phi) = \frac{P}{(\mathrm{d}A \cos\theta)\mathrm{d}\omega}$$

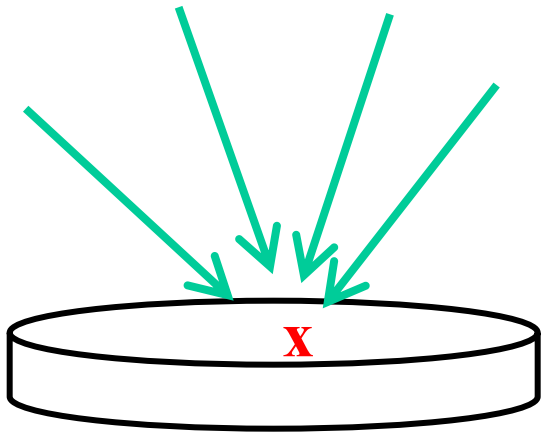
radiance in direction different from surface normal, use spherical coordinates θ, ϕ



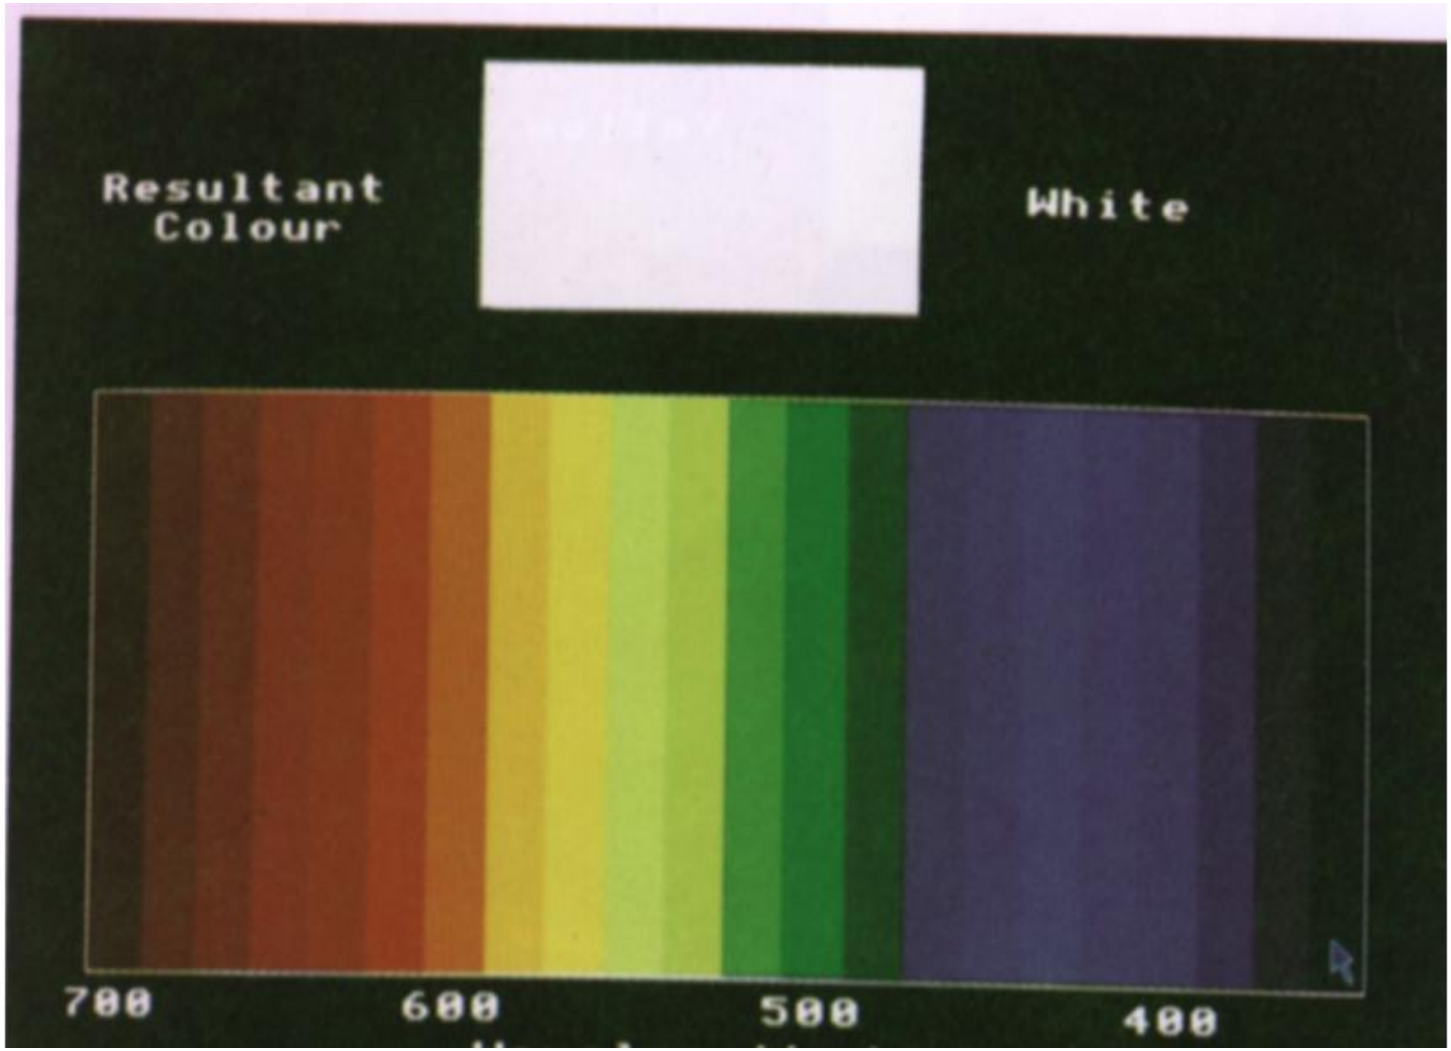
Irradiance

- Total power arriving at the surface (from all incoming angles)
  - Units: power per unit area, $W/m^2 = W\ m^{-2}$

$$E(\mathbf{X}) = \int_{\text{hemisphere}} L(\mathbf{X}, \theta, \phi) \cos\theta \mathrm{d}\omega$$



Visible Light Spectrum



Camera sensor

- Measured pixel intensity is a function of irradiance E integrated over
  - Pixel's area (x,y)
  - range of wavelengths λ
  - some period of time t

spatial response of pixel

spectral response of pixel

$$I = \int_t \int_\lambda \int_y \int_x E(x, y, \lambda, t) s(x, y) q(\lambda) \, \mathrm{d}x \, \mathrm{d}y \, \mathrm{d}\lambda \, \mathrm{d}t$$

- Ideally, the camera response function R is linear to the radiance, but it may not be

$$I = R\left(\int_t \int_\lambda \int_y \int_x E(x, y, \lambda, t) s(x, y) q(\lambda) \, \mathrm{d}x \, \mathrm{d}y \, \mathrm{d}\lambda \, \mathrm{d}t \right)$$

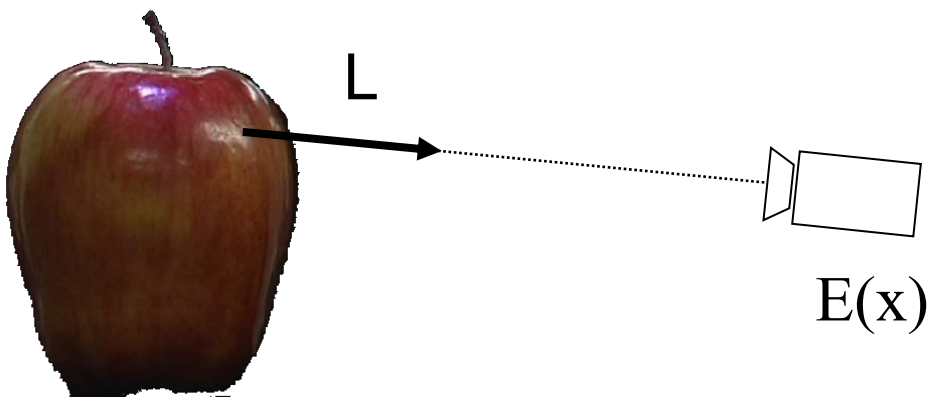
Image irradiance is proportional scene radiance

For a camera with a thin lens, it can be shown that

$$E(x) = k_L L$$

where

- E(x) is the image irradiance at point x
- L is the radiance coming from a scene point projecting to image point x
- k_L is a proportionality constant that may depend on the lens and may be a function of x

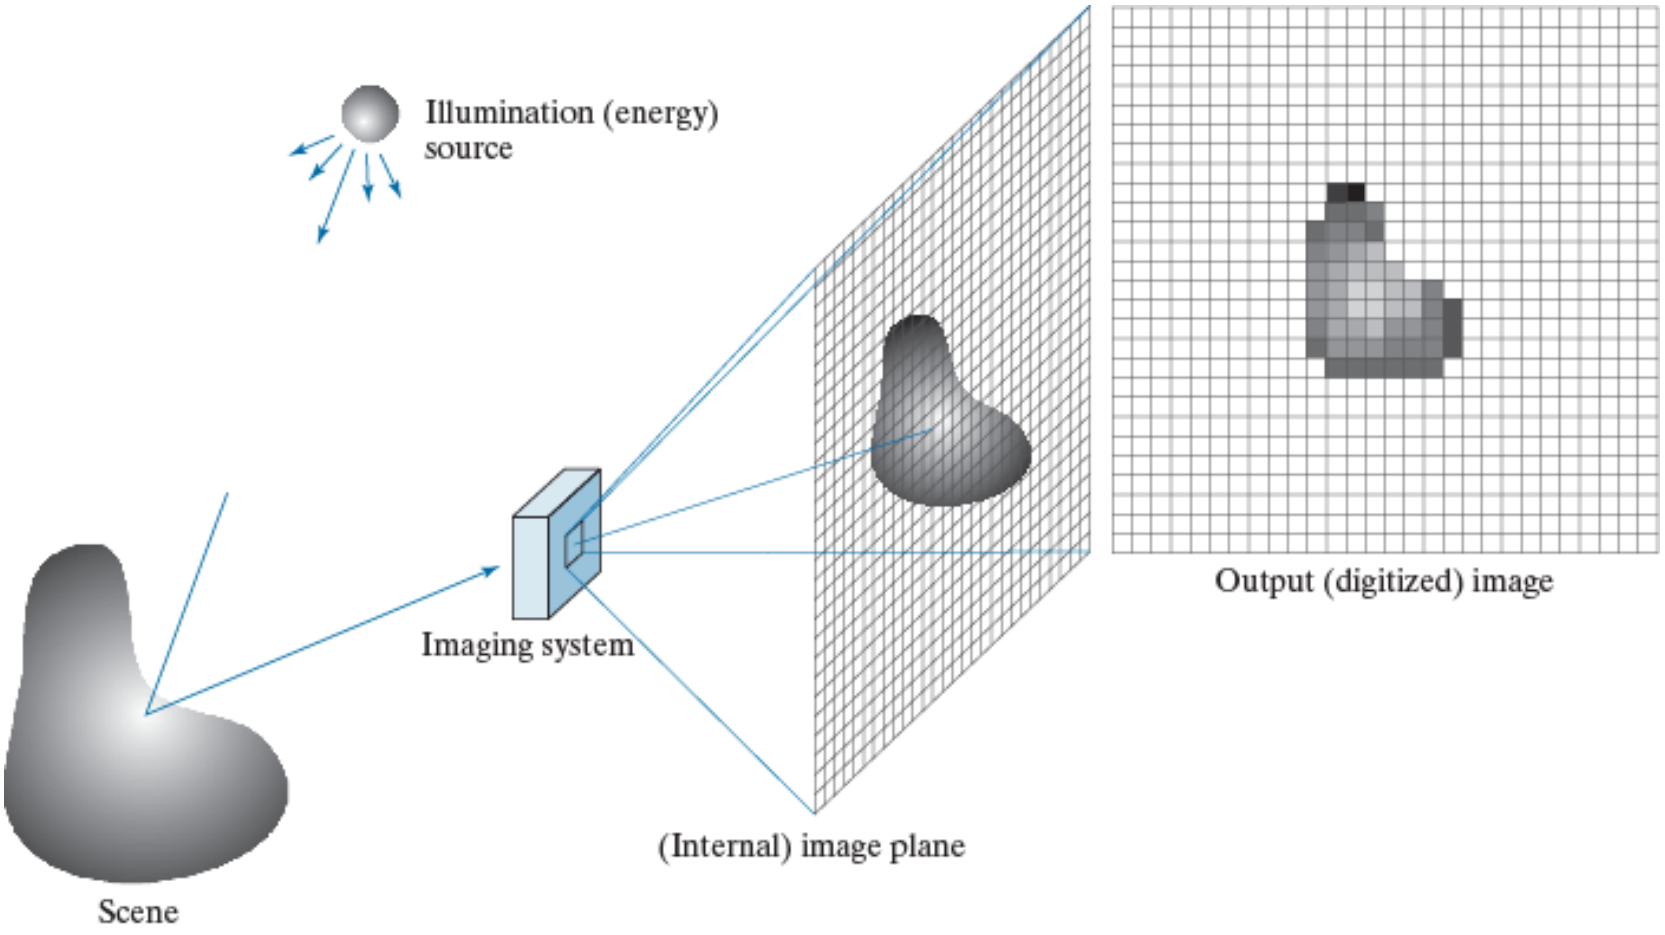


Combined with linear sensor model, we have

$$I = k_c k_L L$$

In other words, the measured pixel intensity is proportional to the radiance

Image acquisition



Color Cameras

Eye:

Three types of Cones

Cameras:

1. Filter wheel
2. Prism (with 3 sensors)
3. Filter mosaic

… and X3

Filter wheel

Rotate multiple filters in front of lens

Allows more than 3 color bands



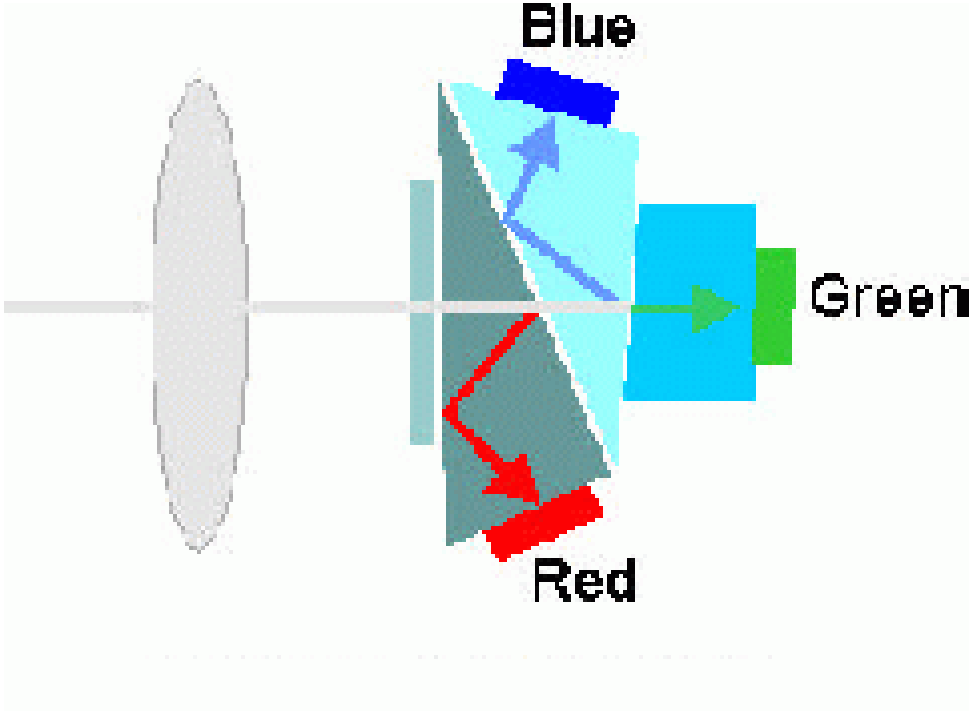
Only suitable for static scenes

Prism color camera

Separate light in 3 beams using dichroic prism

Requires 3 sensors & precise alignment

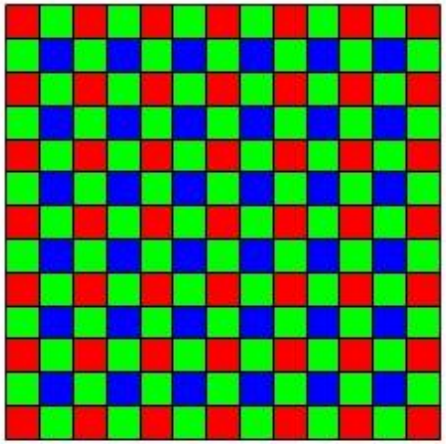
Good color separation



Filter mosaic



Coat filter directly on sensor



Bayer filter

Color filter array (CFA)

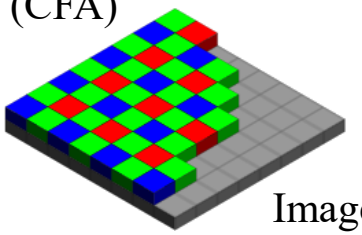
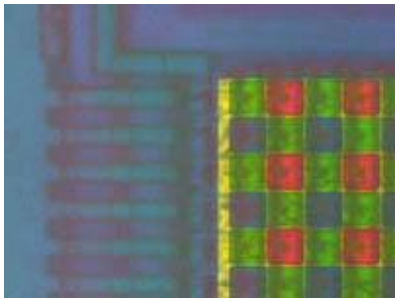


Image sensor



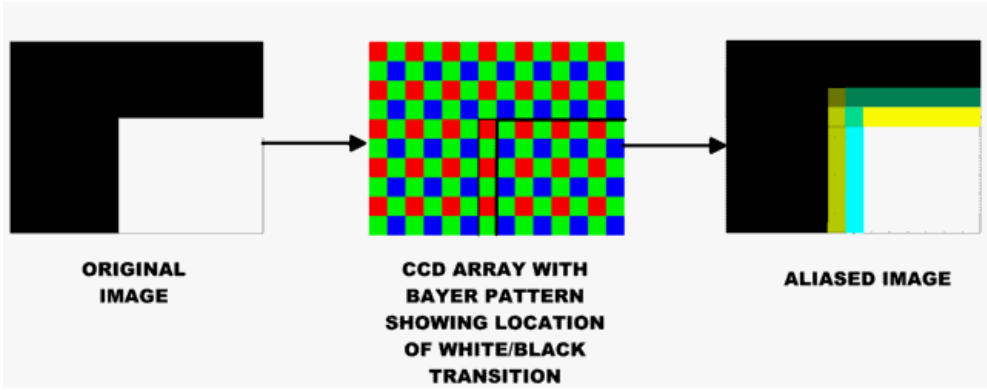
Demosaicing (obtain full color & full resolution image)

G	R	G	R
B	G	B	G
G	R	G	R
B	G	B	G

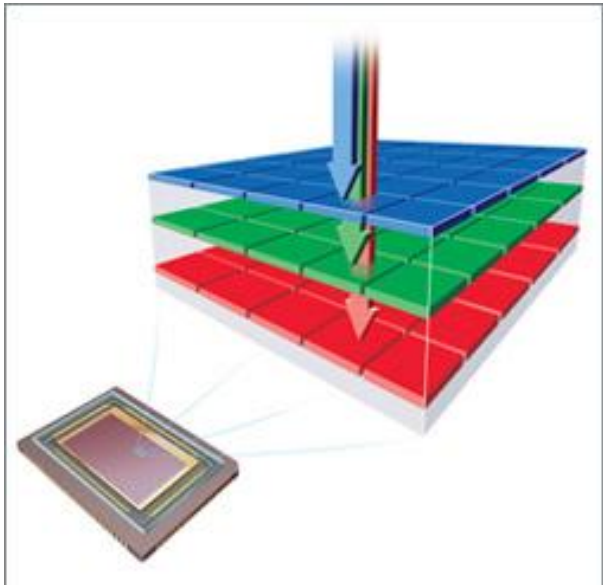
CFA

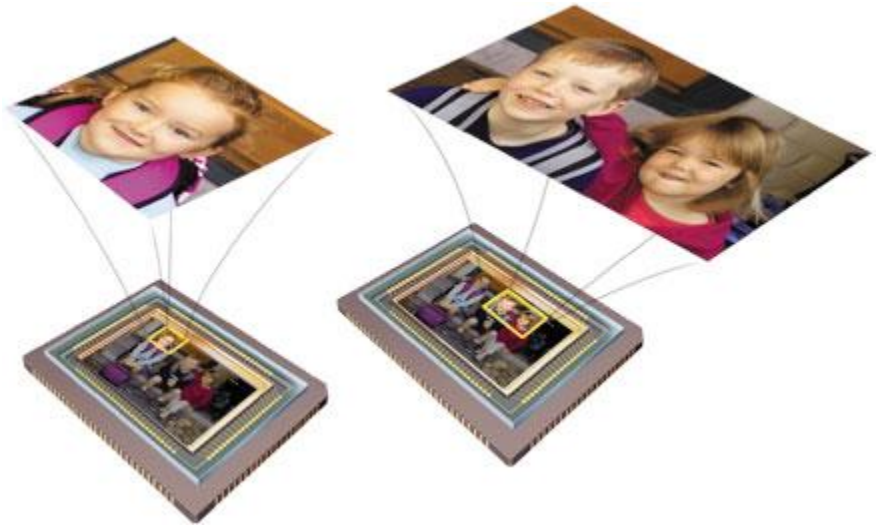
rGb	Rgb	rGb	Rgb
rgB	rGb	rgB	rGb
rGb	Rgb	rGb	Rgb
rgB	rGb	rgB	rGb

Interpolated (lower case) pixel values

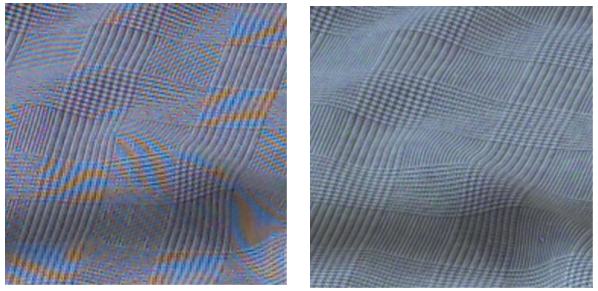


Color CMOS sensor
Foveon's X3

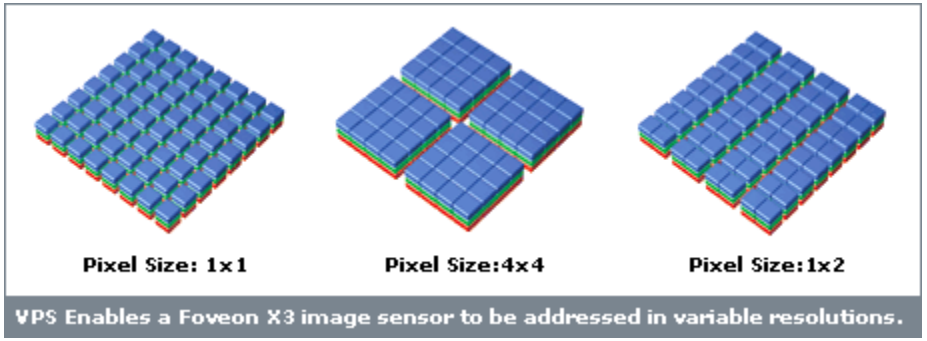




better image quality



smarter pixels



Light at surfaces

Many effects when light strikes a surface -- could be:

- Reflected
  - Mirror
- Transmitted
  - Skin, glass
- Scattered
  - Milk
- Travel along the surface and leave at some other point
- Absorbed

We will assume:

- All the light leaving a point is due to that arriving at that point
- Surfaces don't fluoresce
  - e.g., scorpions, detergents
- Surfaces don't emit light (i.e., are cool)

Light at surfaces











BRDF

- Bi-directional Reflectance Distribution Function

$$\rho(\theta_{in}, \phi_{in}; \theta_{out}, \phi_{out})$$

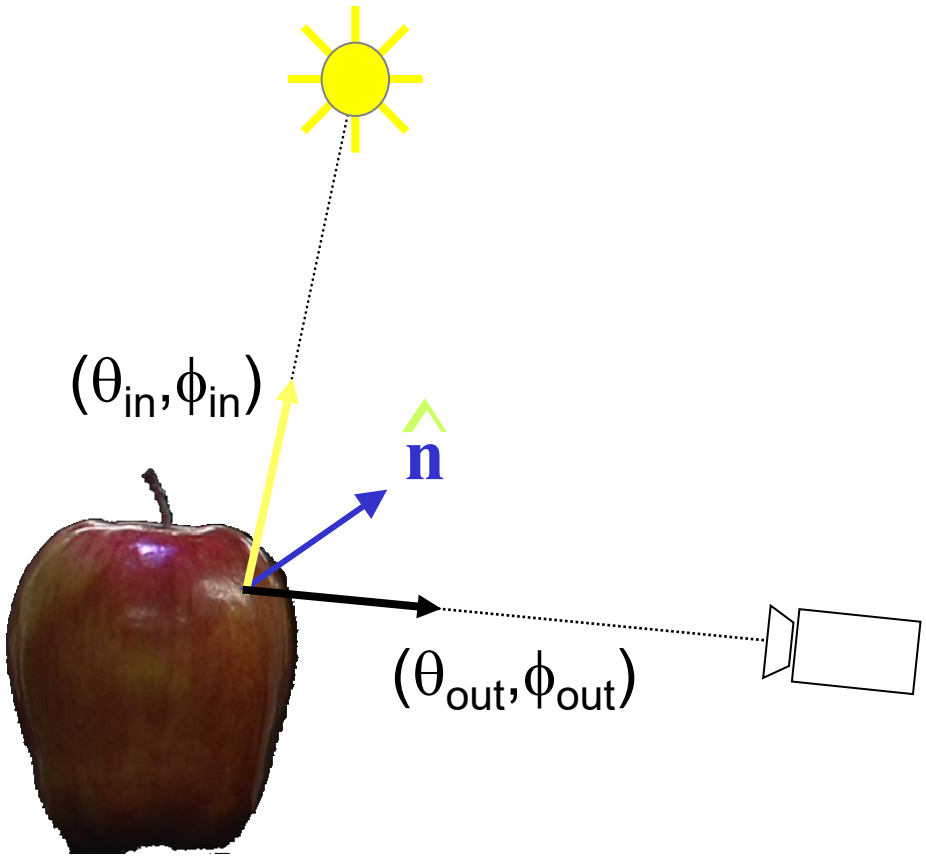
- Function of
  - Incoming light direction:

$$\theta_{in}, \phi_{in}$$

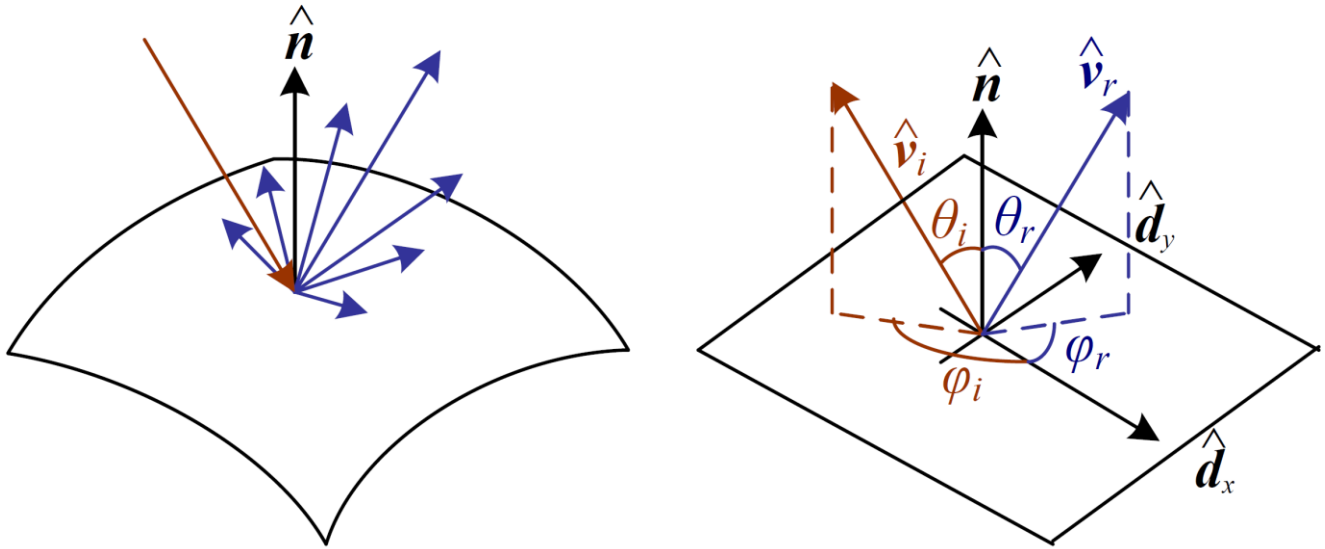
  - Outgoing light direction:

$$\theta_{out}, \phi_{out}$$

- Ratio of emitted radiance to incident irradiance



Lighting, reflectance, and shading

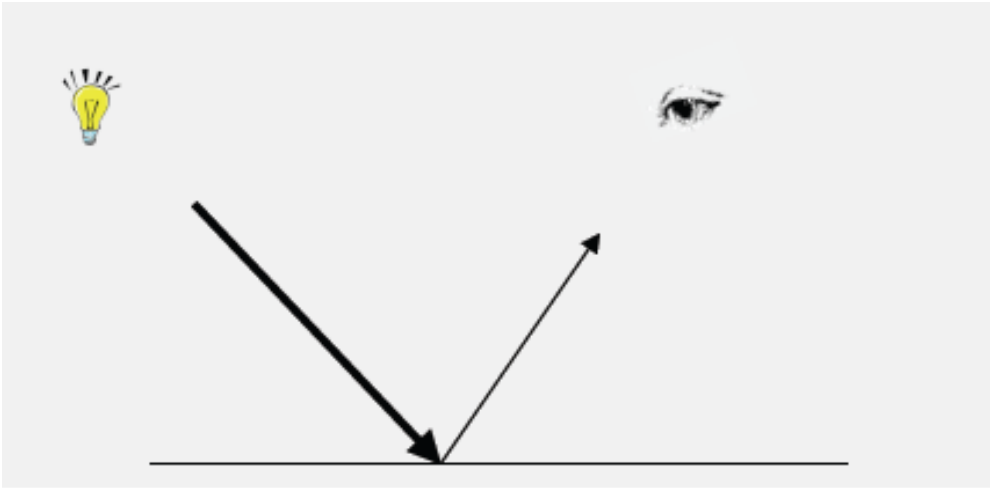


BRDF $f_r(\theta_i, \phi_i, \theta_r, \phi_r; \lambda)$

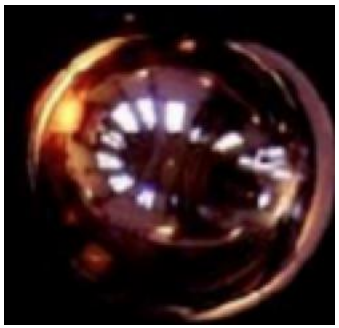
Specular reflection

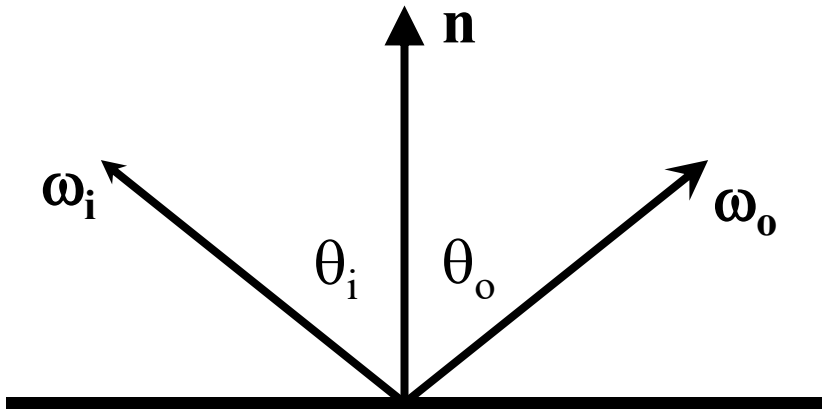
- Ideal specular reflection is mirror reflection
  - Perfectly smooth surface
  - Incoming light ray is bounced in single direction
  - Angle of incidence equals angle of reflection





Specular Reflection: Smooth Surface





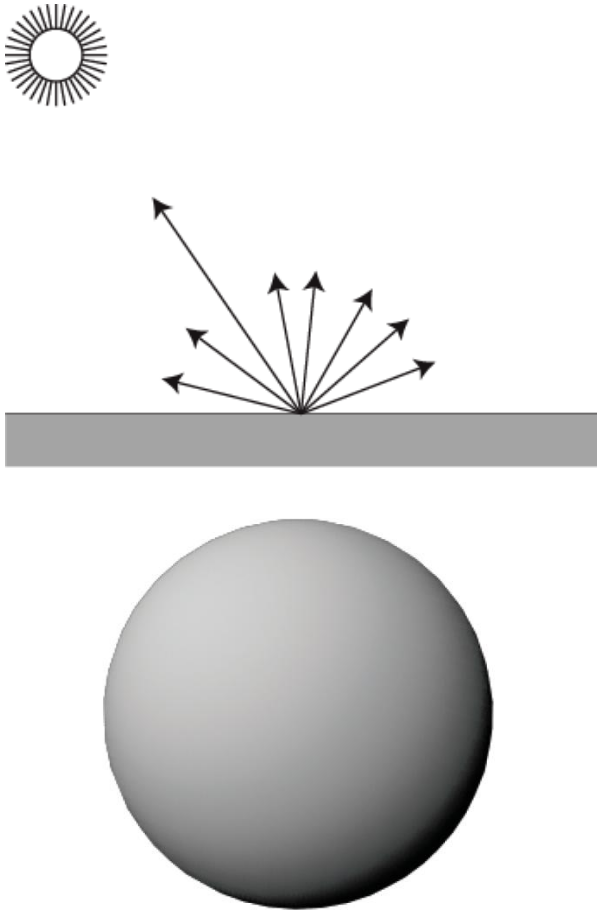
- **n, ω_i, ω_o** are coplanar
- $\theta_i = \theta_o$

$$\boldsymbol{\omega}_o = 2(\boldsymbol{\omega}_i{}^\top \mathbf{n})\mathbf{n} - \boldsymbol{\omega}_i$$

Speculum – Latin for "Mirror"

Diffuse surface

- Ideal diffuse material reflects light equally in all directions
- View-independent
- Matte, not shiny materials
  - Paper
  - Unfinished wood
  - Unpolished stone



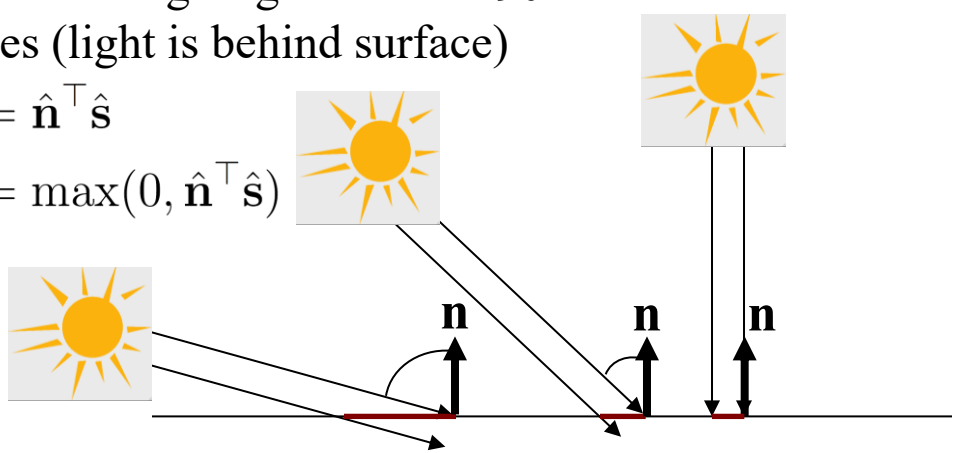
Diffuse reflection

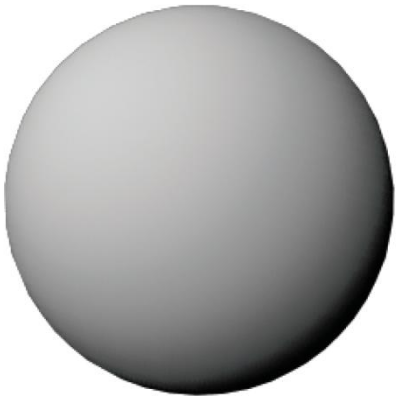
- Beam of parallel rays shining on a surface
  - Area covered by beam varies with the angle between the beam and the normal
  - The larger the area, the less incident light per area
  - Incident light per unit area is proportional to the cosine of the angle between the normal and the light rays
- Object darkens as normal turns away from light
- Lambert's cosine law (Johann Heinrich Lambert, 1760)
- Diffuse surfaces are also called Lambertian surfaces

Do not allow angles greater than 90 degrees (light is behind surface)

$$\cos\theta = \hat{\mathbf{n}}^\top \hat{\mathbf{s}}$$

$$\cos^+\theta = \max(0, \hat{\mathbf{n}}^\top \hat{\mathbf{s}})$$





Glossy surface

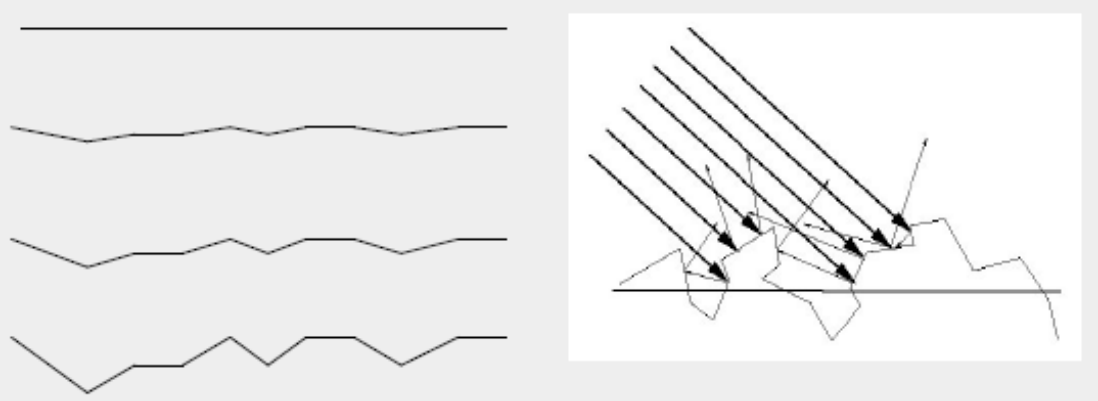
- Assume surface composed of small mirrors with random orientation (micro-facets)
- Smooth surfaces
  - Micro-facet normals close to surface normal
  - Sharp highlights
- Rough surfaces
  - Micro-facet normals vary strongly
  - Blurry highlight

Polished

Smooth

Rough

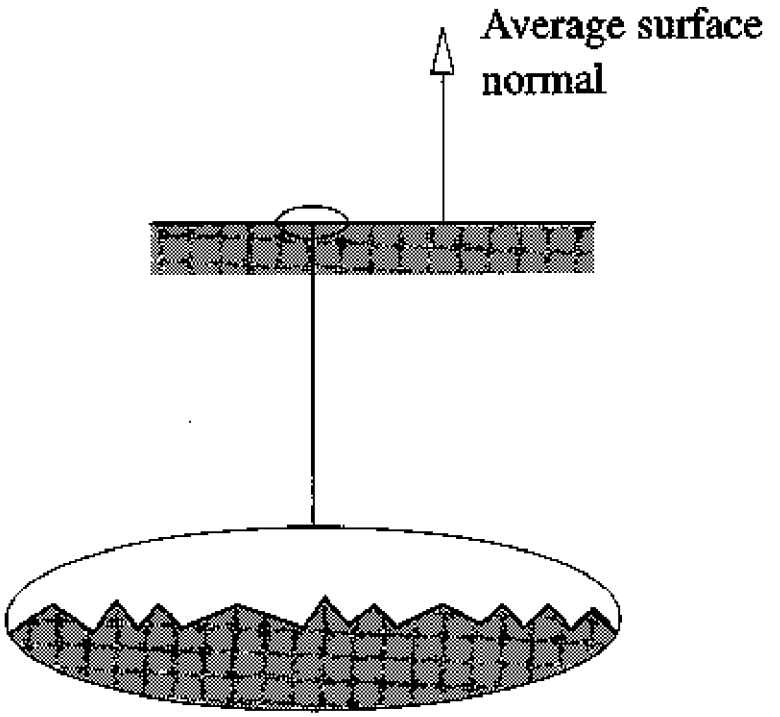
Very rough



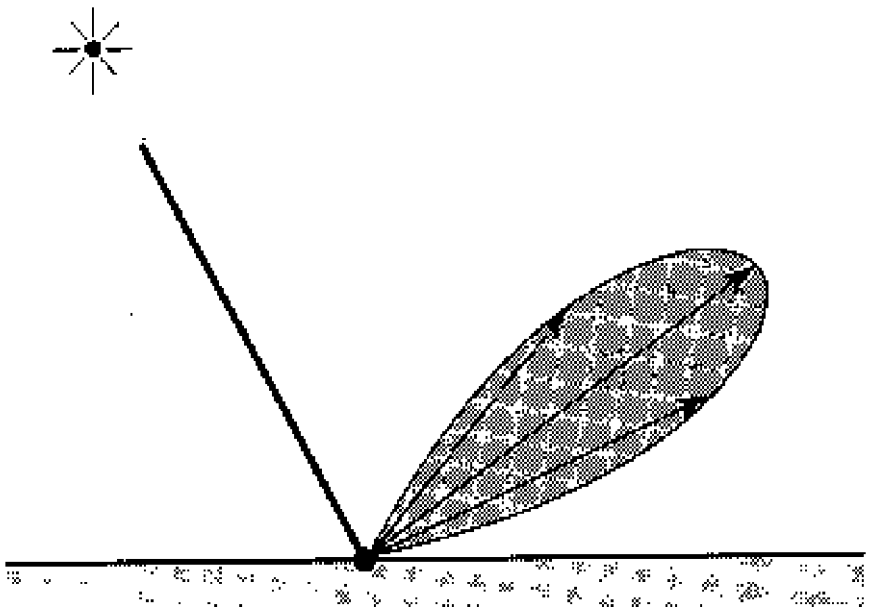
Glossy reflection

- Expect most light to be reflected in mirror direction
- Because of micro-facets, some light is reflected slightly off ideal reflection direction
- Reflection
  - Brightest when view vector is aligned with reflection
  - Decreases as angle between view vector and reflection direction increases

Phong reflectance model



Symmetric V-shaped
grooves – 'microfacets'



Phong Lobe
(Lobe illustrates brightness in a direction)

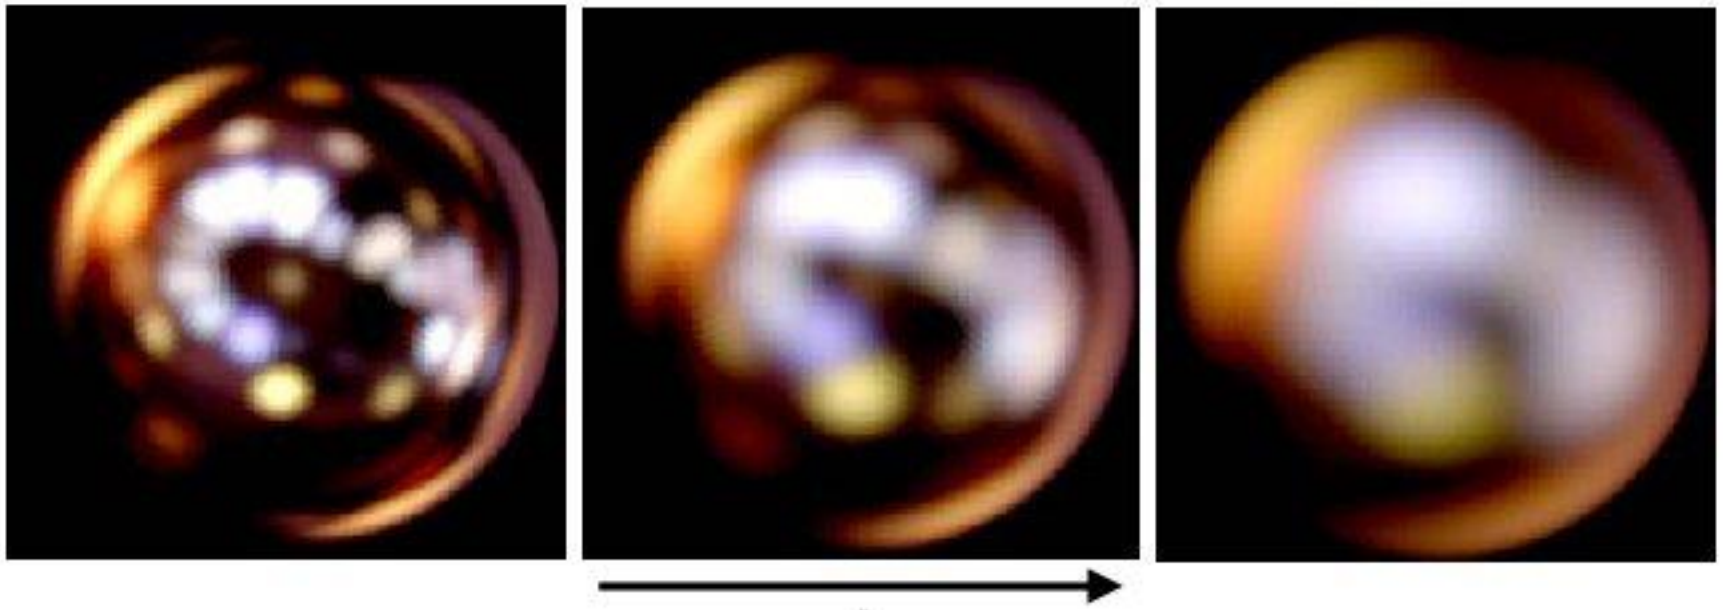
Phong Model



Mirror



Diffuse



General BRDF



Example: velvet

Portrait of Sir Thomas Morre, Hans Holbein the Younger, 1527

Light

- Special light sources
  - Point sources
  - Distant point sources
  - Area sources

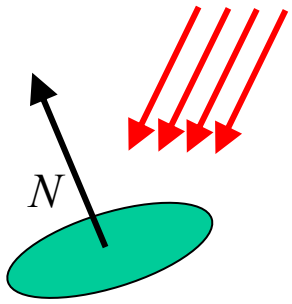
Point light source

- Similar to light bulbs
- An infinitesimally small point that radiates light equally in all directions
  - Light vector varies across receiving surface
  - Intensity drops off proportionally to the inverse square of the distance from the light
    - Reason for inverse square falloff:
      Surface area of sphere $A = 4\pi r^2$

Light from a distant source

- Note, if light is very far away, then view light as coming from a direction in 3D
- Directional light source
  - Light rays are parallel
  - Direction and intensity are the same everywhere
  - As if the source were infinitely far away

$$\rho_d(x)\left(N(x)^T S(x)\right)$$

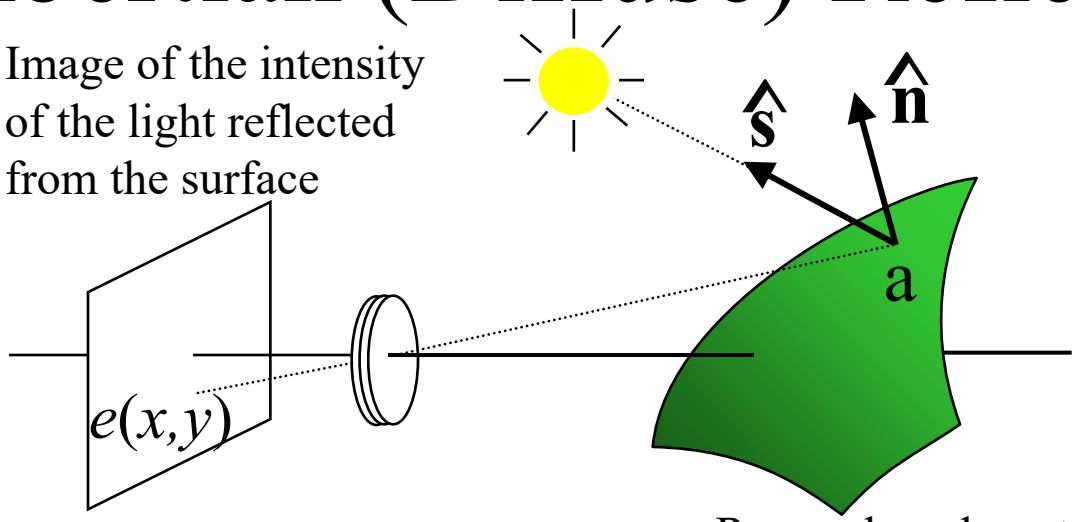


Remember, do not allow angles greater than 90 degrees (light is behind surface)

$$\cos(\theta) = \hat{\mathbf{n}}^\top \hat{\mathbf{s}}$$

$$\cos^+(\theta) = \max(0, \hat{\mathbf{n}}^\top \hat{\mathbf{s}})$$

Lambertian (Diffuse) Reflection



Remember, do not allow angles greater than 90 degrees (light is behind surface)

$$e(x,y) = a(x,y)s_0 \cos\theta$$

$$e(x,y) = a(x,y)s_0 \hat{\mathbf{n}}(x,y)^\top \hat{\mathbf{s}}(x,y), \text{ where } \cos\theta = \hat{\mathbf{n}}(x,y)^\top \hat{\mathbf{s}}(x,y)$$

$$\cos\theta = \hat{\mathbf{n}}^\top \hat{\mathbf{s}}$$

$$e(x,y) = a(x,y)s_0 \max(0, \hat{\mathbf{n}}(x,y)^\top \hat{\mathbf{s}}(x,y))$$

$$\cos^+\theta = \max(0, \hat{\mathbf{n}}^\top \hat{\mathbf{s}})$$

$$e(x,y) = a(x,y) \max(0, \hat{\mathbf{n}}(x,y)^\top \mathbf{s}(x,y)), \text{ where } \mathbf{s}(x,y) = s_0 \hat{\mathbf{s}}(x,y)$$

where

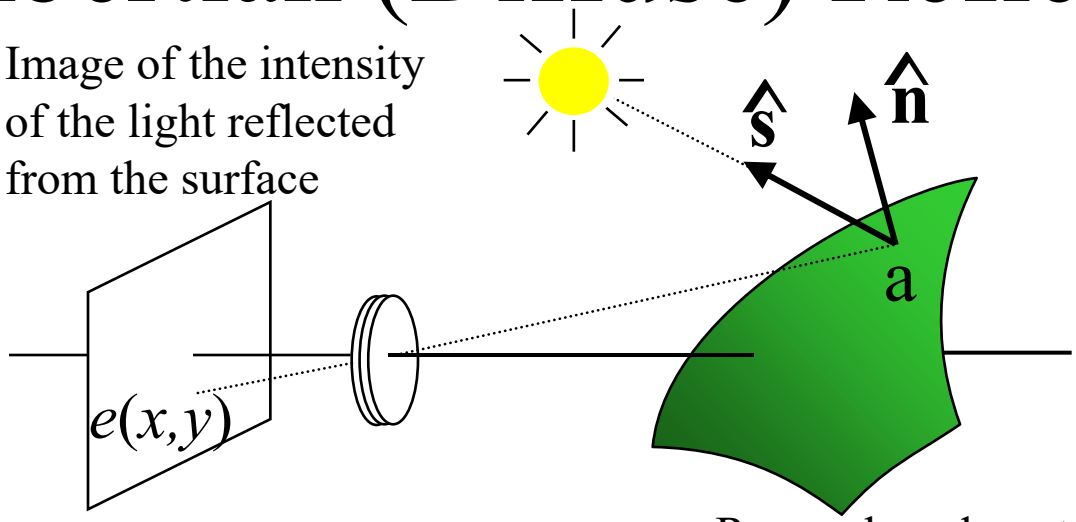
$a(x,y)$ is the albedo of the surface facet imaged by pixel

$\hat{\mathbf{n}}(x,y)$ is the unit normal of the surface facet imaged by pixel

s_0 is the intensity of the light source

$\hat{\mathbf{s}}(x,y)$ is the unit direction to the light source from the surface facet imaged by pixel

Lambertian (Diffuse) Reflection



Remember, do not allow angles greater than 90 degrees (light is behind surface)

$$e(x,y) = a(x,y)s_0 \cos\theta$$

$$e(x,y) = a(x,y)s_0 \hat{\mathbf{n}}(x,y)^\top \hat{\mathbf{s}}(x,y), \text{ where } \cos\theta = \hat{\mathbf{n}}(x,y)^\top \hat{\mathbf{s}}(x,y)$$

$$\cos\theta = \hat{\mathbf{n}}^\top \hat{\mathbf{s}}$$

$$e(x,y) = a(x,y)s_0 \max(0, \hat{\mathbf{n}}(x,y)^\top \hat{\mathbf{s}}(x,y))$$

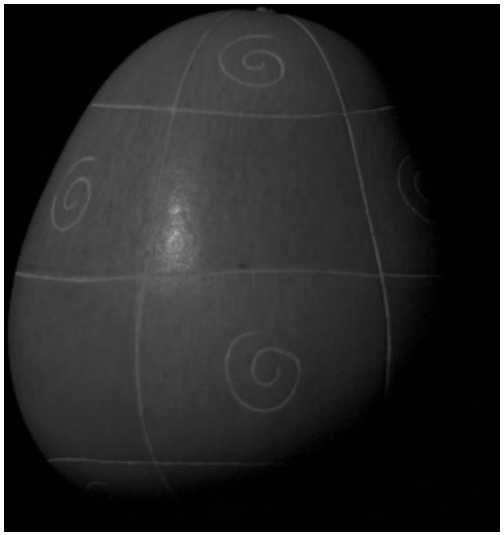
$$\cos^+\theta = \max(0, \hat{\mathbf{n}}^\top \hat{\mathbf{s}})$$

$$e(x,y) = a(x,y) \max(0, \hat{\mathbf{n}}(x,y)^\top \mathbf{s}(x,y)), \text{ where } \mathbf{s}(x,y) = s_0 \hat{\mathbf{s}}(x,y)$$

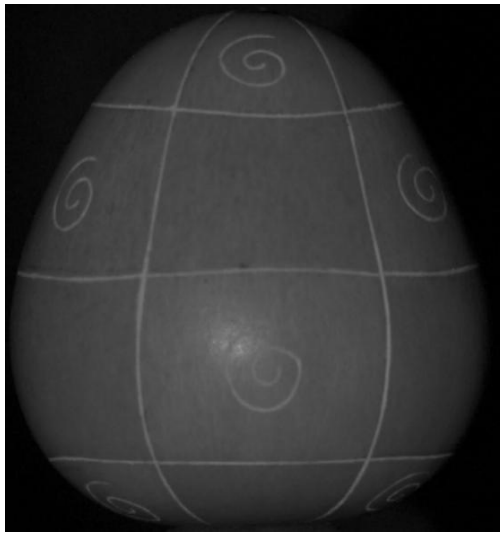
- For color (instead of grayscale)
  - Surface color and light color are RGB
  - Need to compute RGB values of reflected color separately

Multiple Light Sources

- Light is additive
- Integrate over all light sources
- For example, given two light sources
  - Render image 1 using light source 1
  - Render image 2 using light source 2
  - Final image = image 1 + image 2



+



=

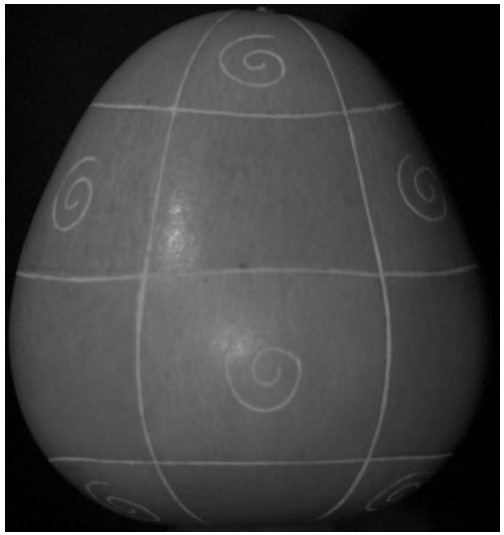
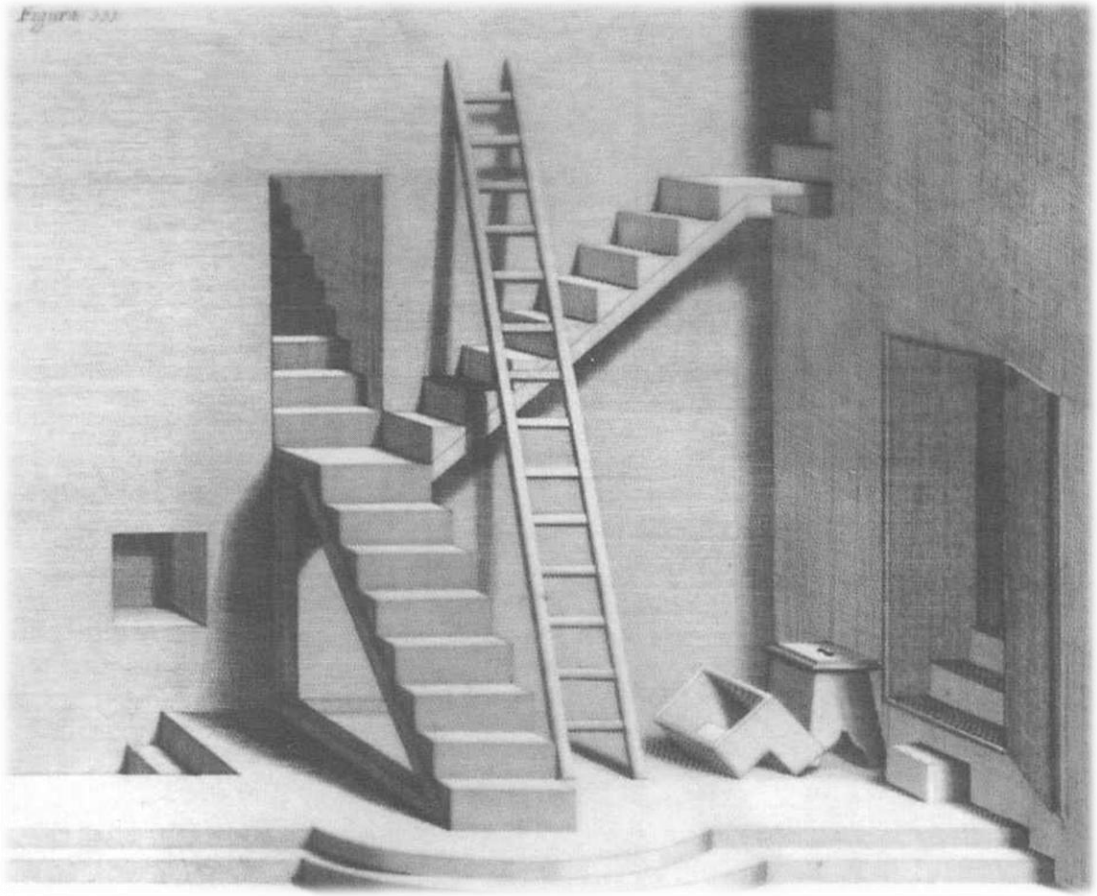


image 1 image 2 final image

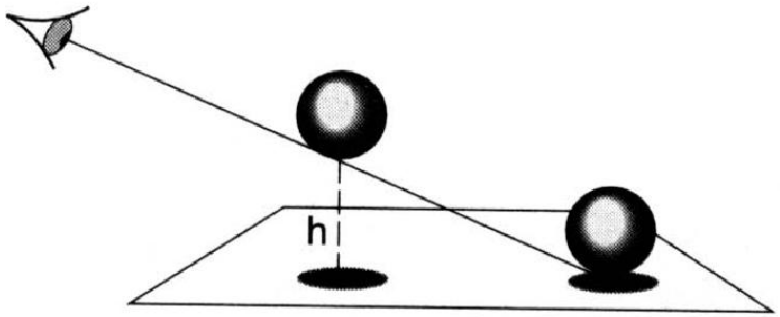
Shadows

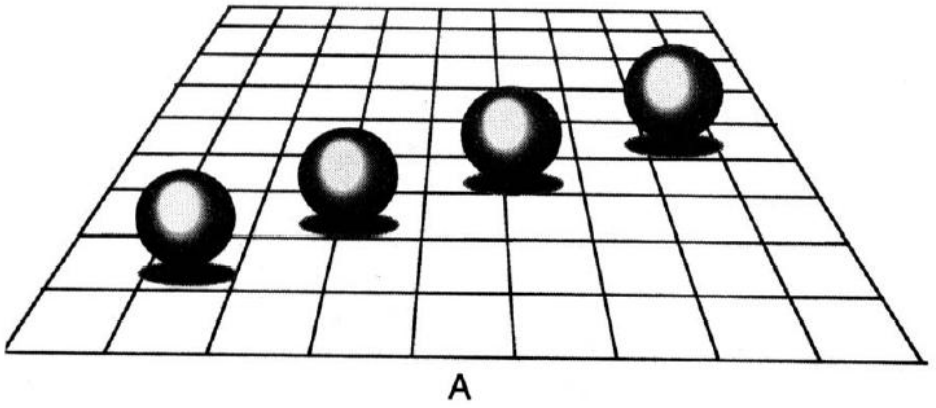
- Give additional cues on scene lighting

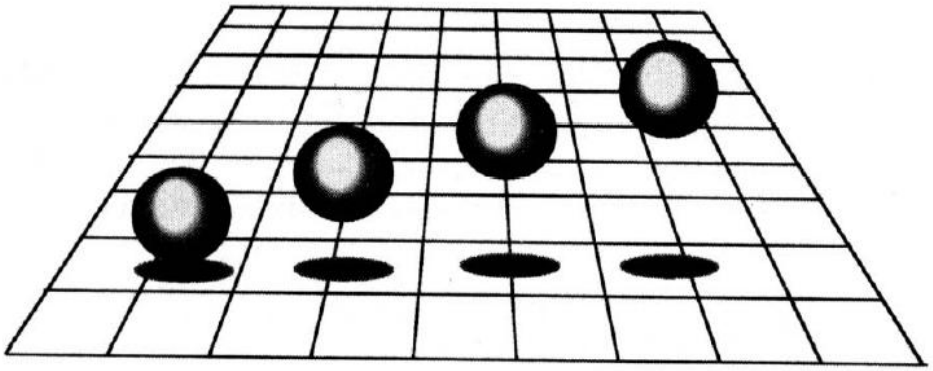


Shadows

- Contact points
- Depth cues

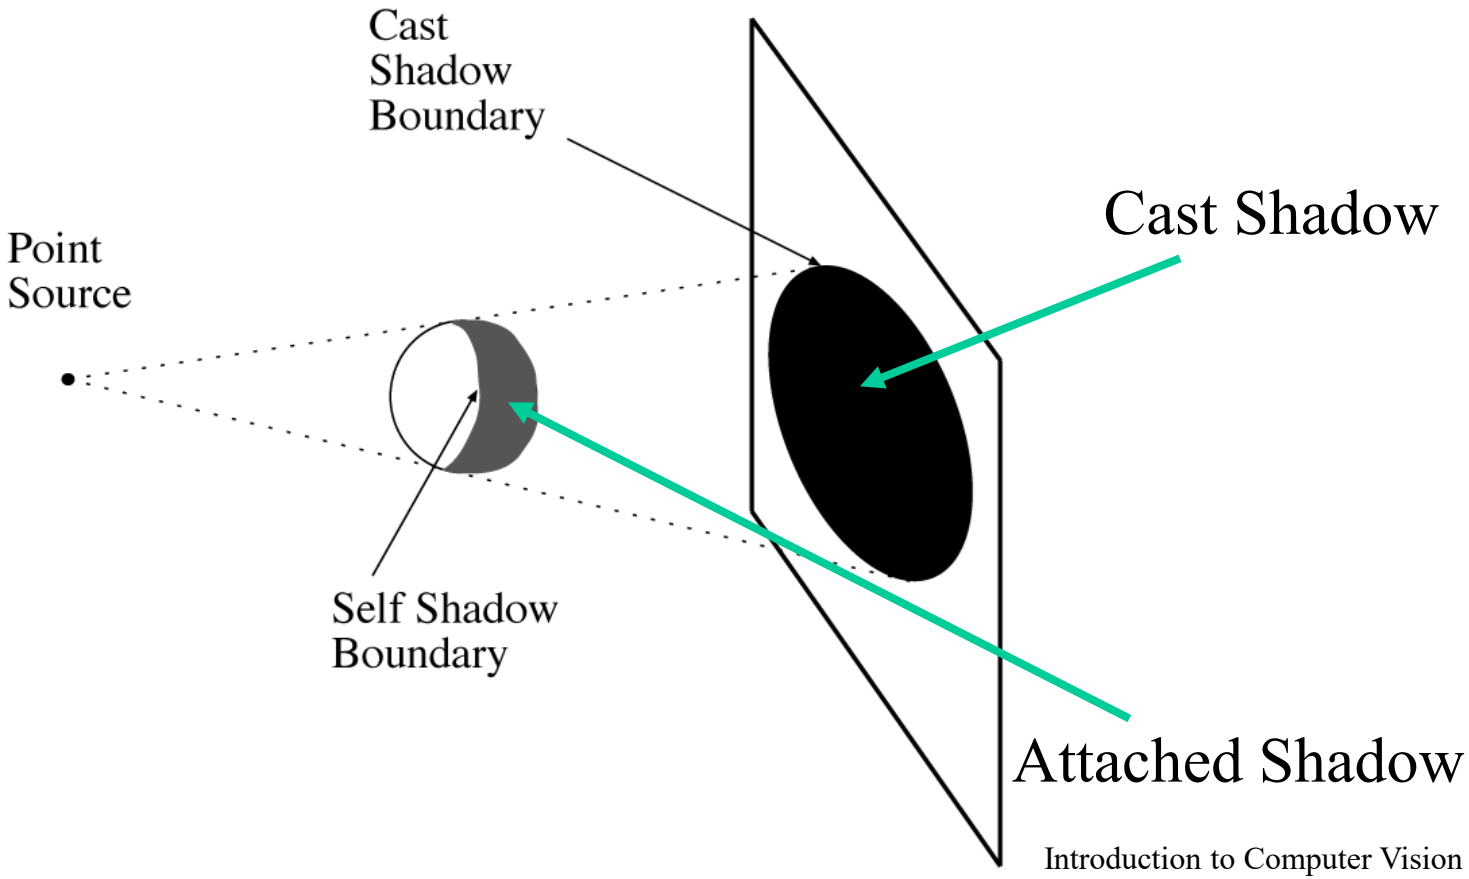






Shadows cast by a point source

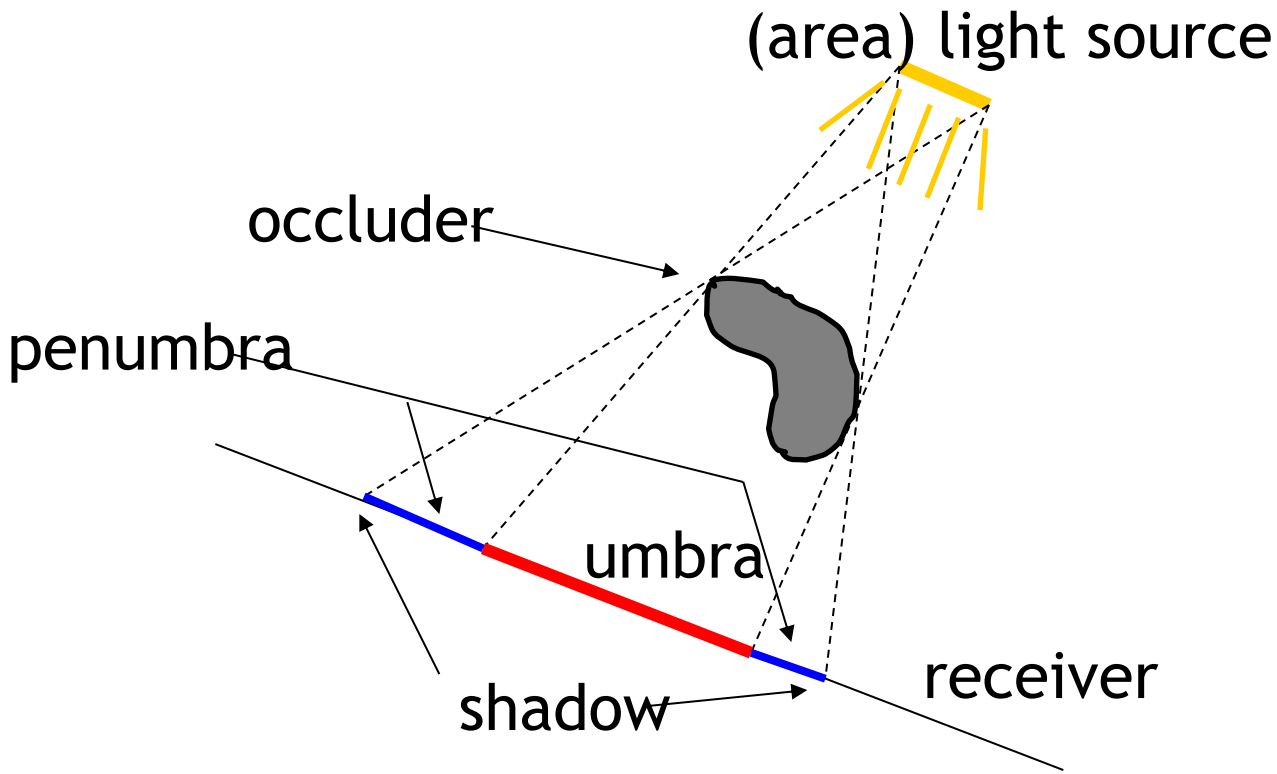
- A point that cannot see the source is in shadow
- For point sources, two types of shadows: cast shadows & attached shadows



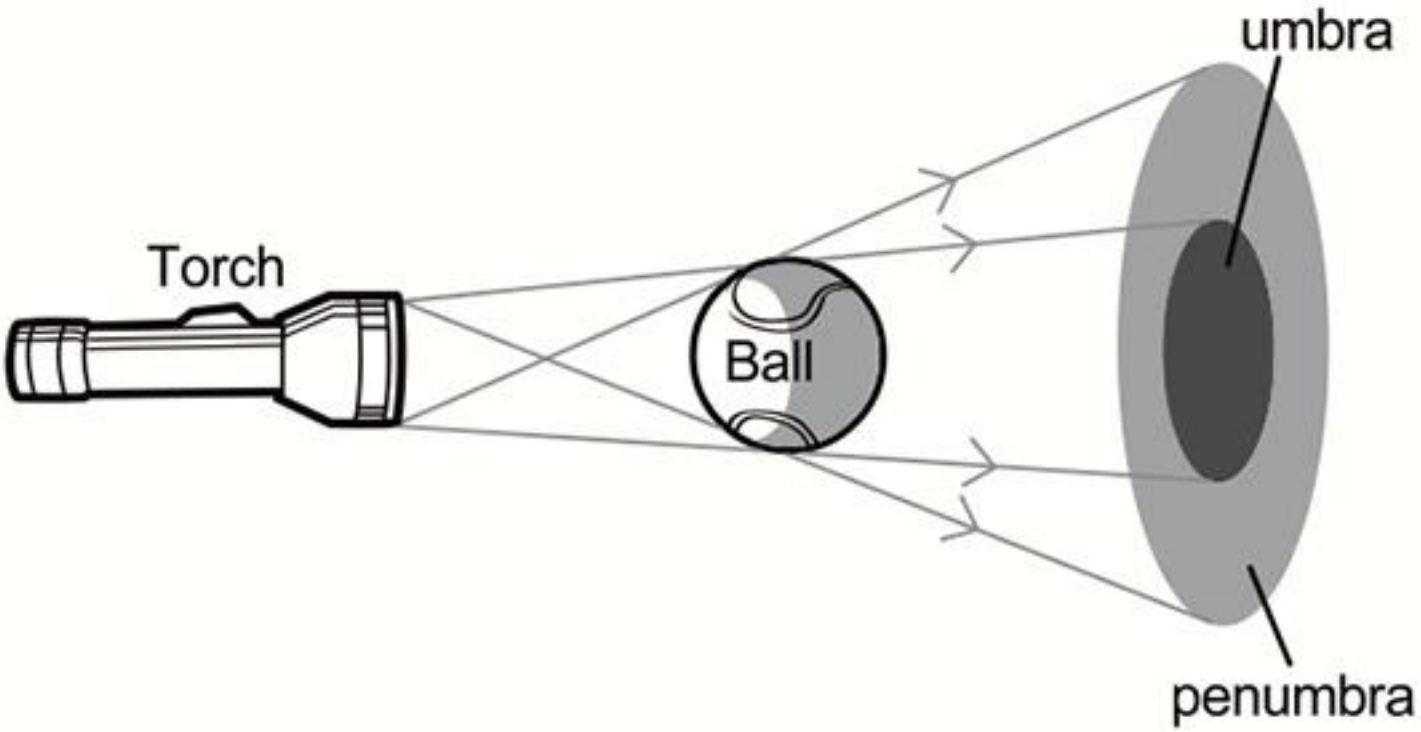
Terminology

Umbra: fully shadowed region

Penumbra: partially shadowed region

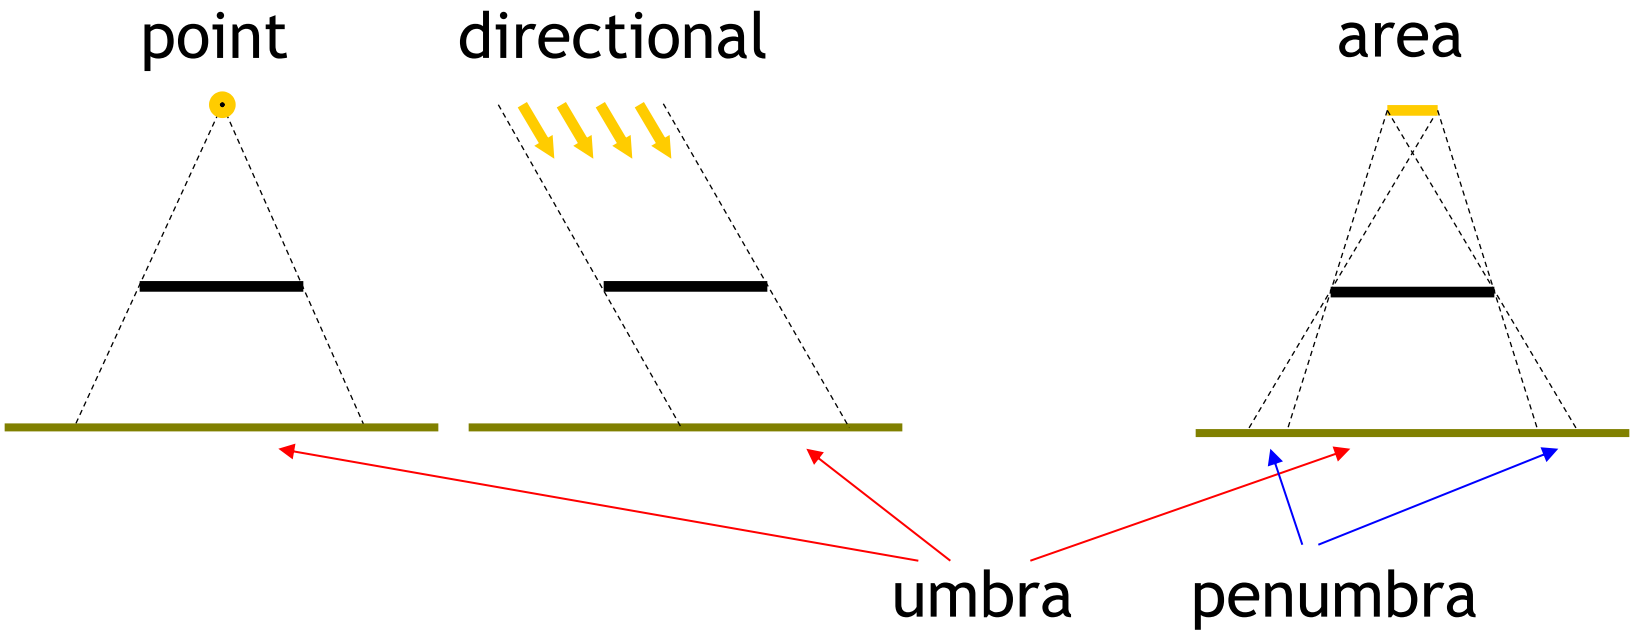


Penumbra and Umbra



Hard and soft shadows

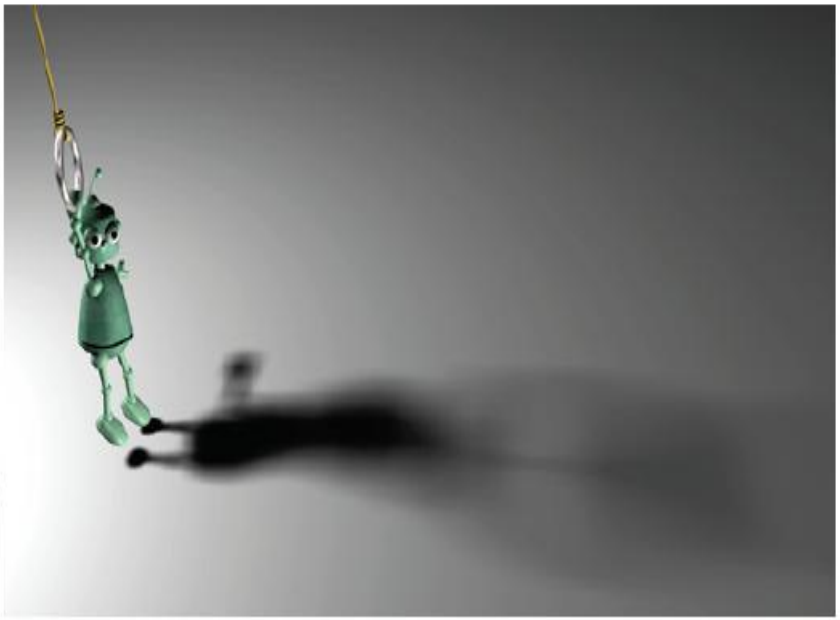
- Point and directional lights lead to hard shadows, no penumbra
- Area light sources lead to soft shadows, with penumbra



Hard and soft shadows



Hard shadow from point light source



Soft shadow from area light source

Next Lecture

- Photometric Stereo