CS5670: Computer Vision

Noah Snavely

Light & Perception

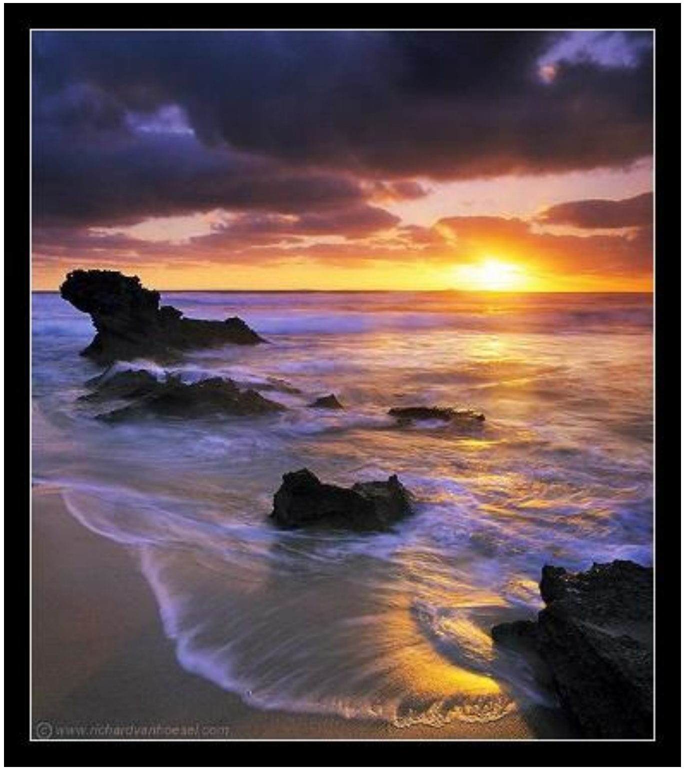

Announcements

- Project 3
  - code due this Thursday, 3/29, by 11:59pm
  - artifact due Friday, 3/30, by 11:59pm

Can we determine shape from lighting?

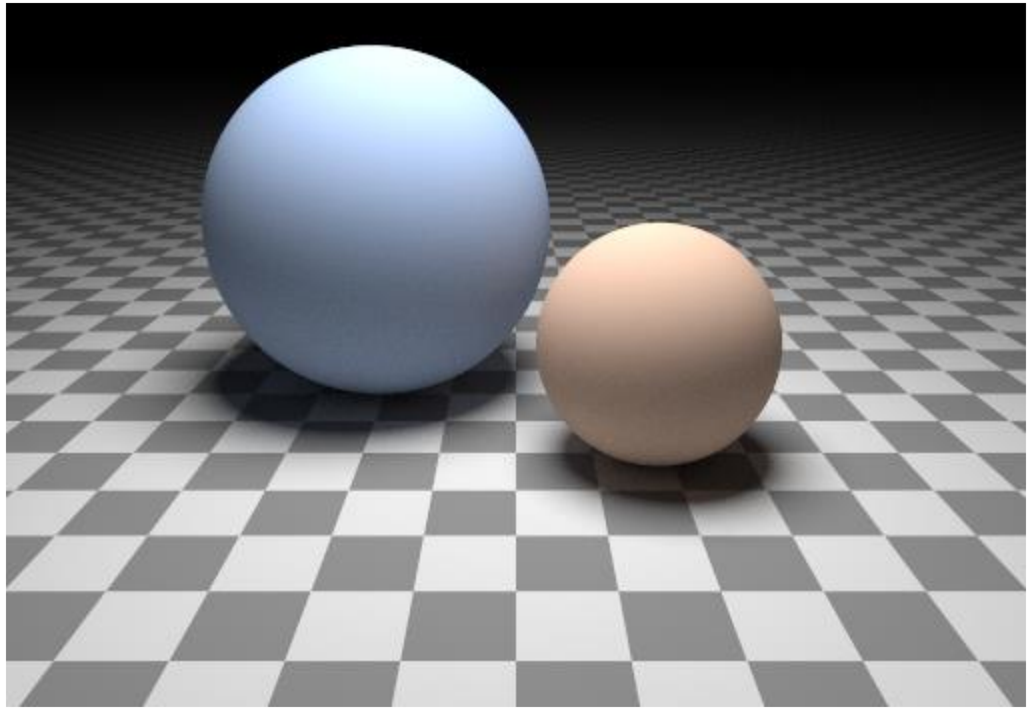

- Are these spheres?
  - Or just flat discs painted with varying albedo?
  - There is ambiguity between *shading* and *reflectance*
  - But still, as humans we can understand the shapes of these objects

What we know: Stereo

Key Idea: use feature motion to understand shape

Next: Photometric Stereo

Key Idea: use pixel brightness to understand shape

Next: Photometric Stereo

Key Idea: use pixel brightness to understand shape

Photometric Stereo

What results can you get?

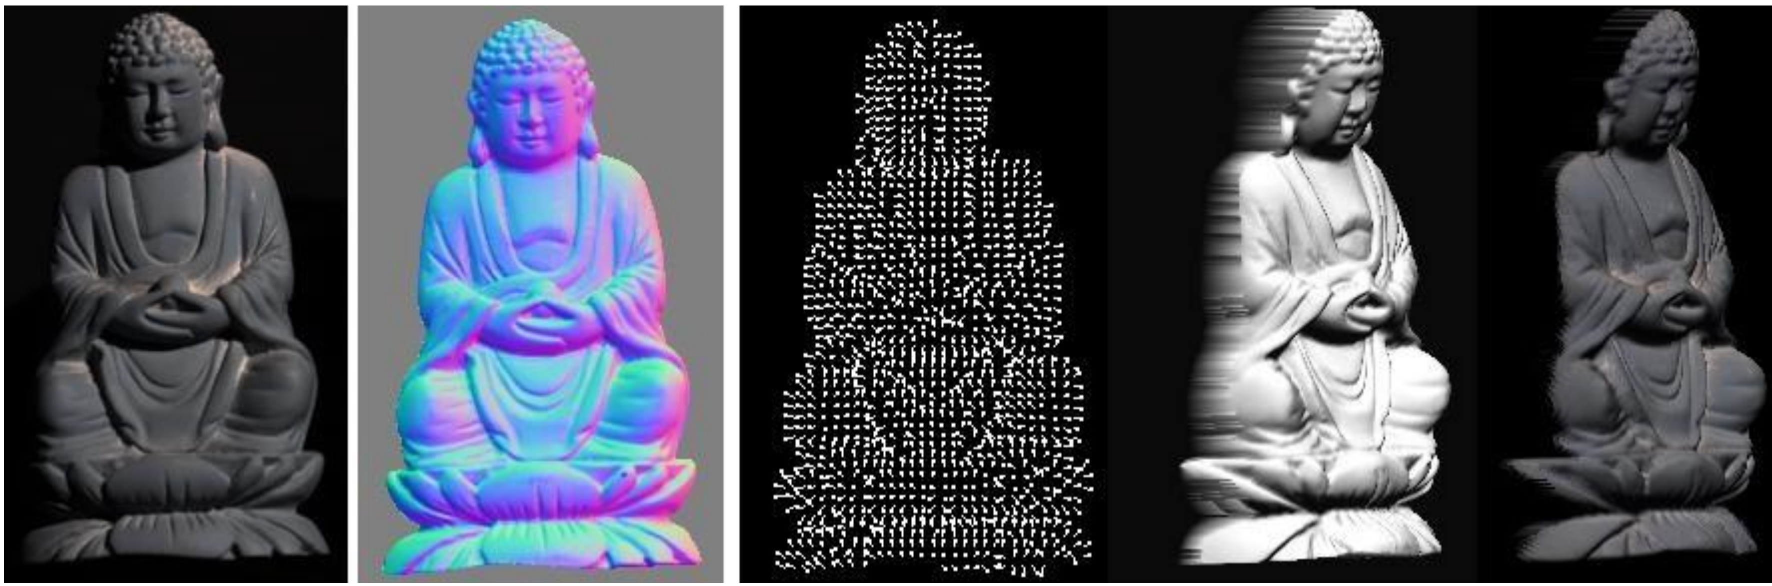

Input (1 of 12) | Normals (RGB colormap) | Normals (vectors) | Shaded 3D rendering | Textured 3D rendering

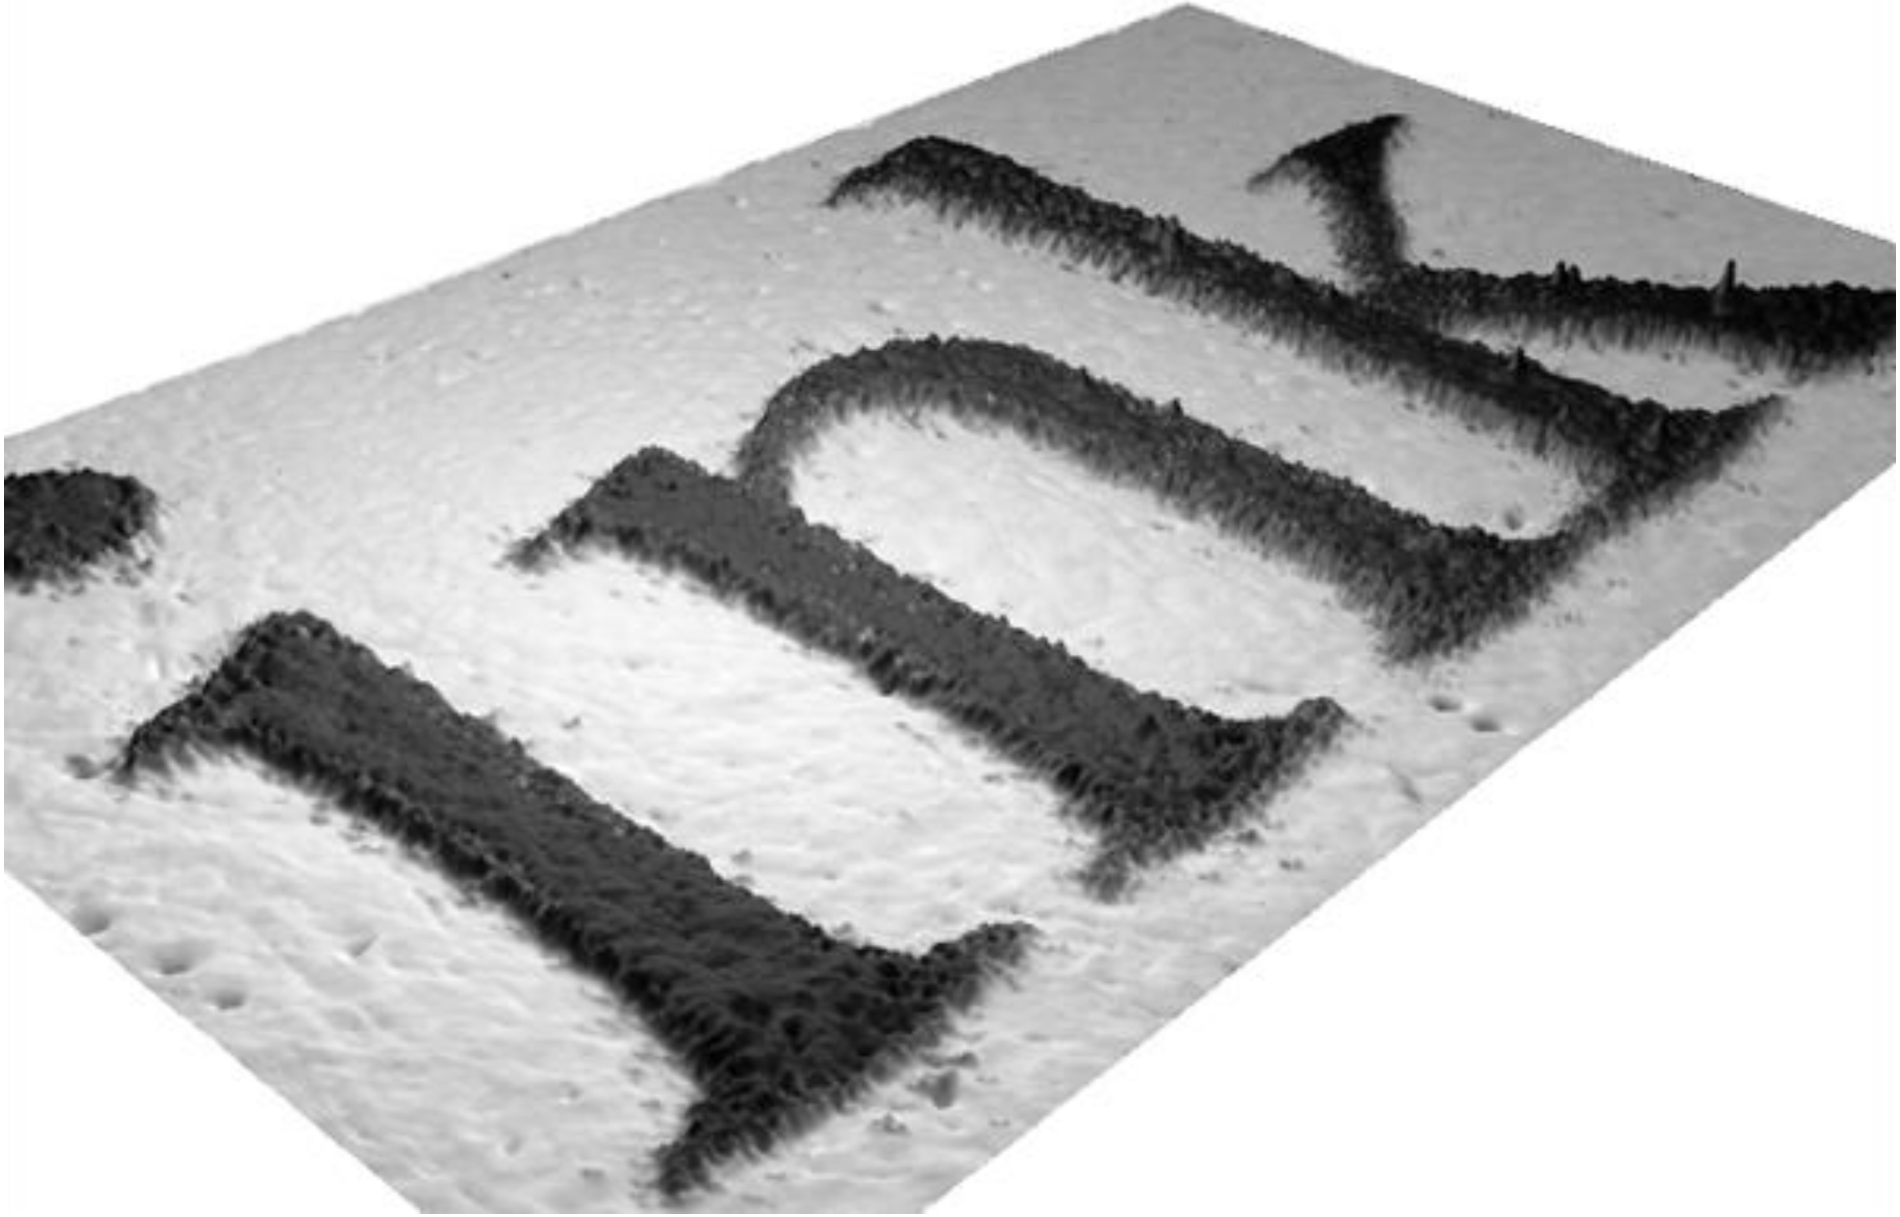

Light

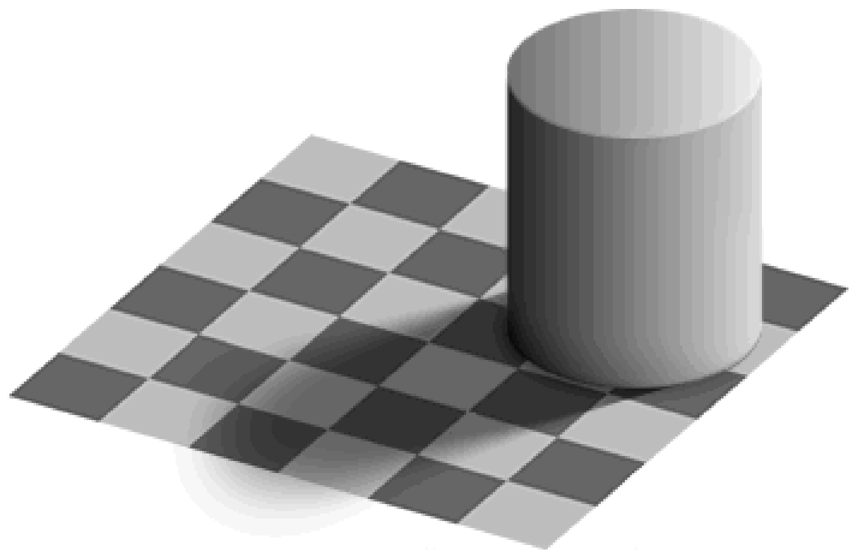

by Ted Adelson

Readings

- Szeliski, 2.2, 2.3.2

Light

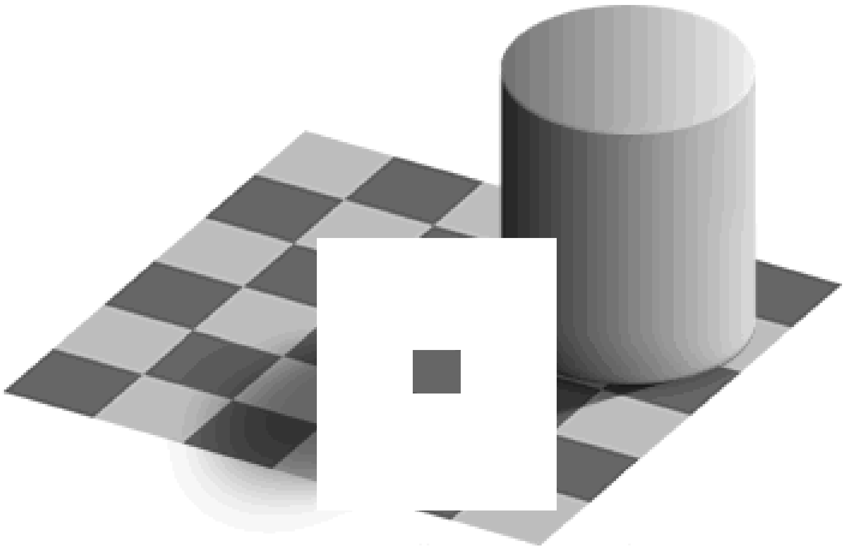

by Ted Adelson

Readings

- Szeliski, 2.2, 2.3.2

Properties of light

Today

- What is light?
- How do we measure it?
- How does light propagate?
- How does light interact with matter?

Radiometry

What determines the brightness of a pixel?

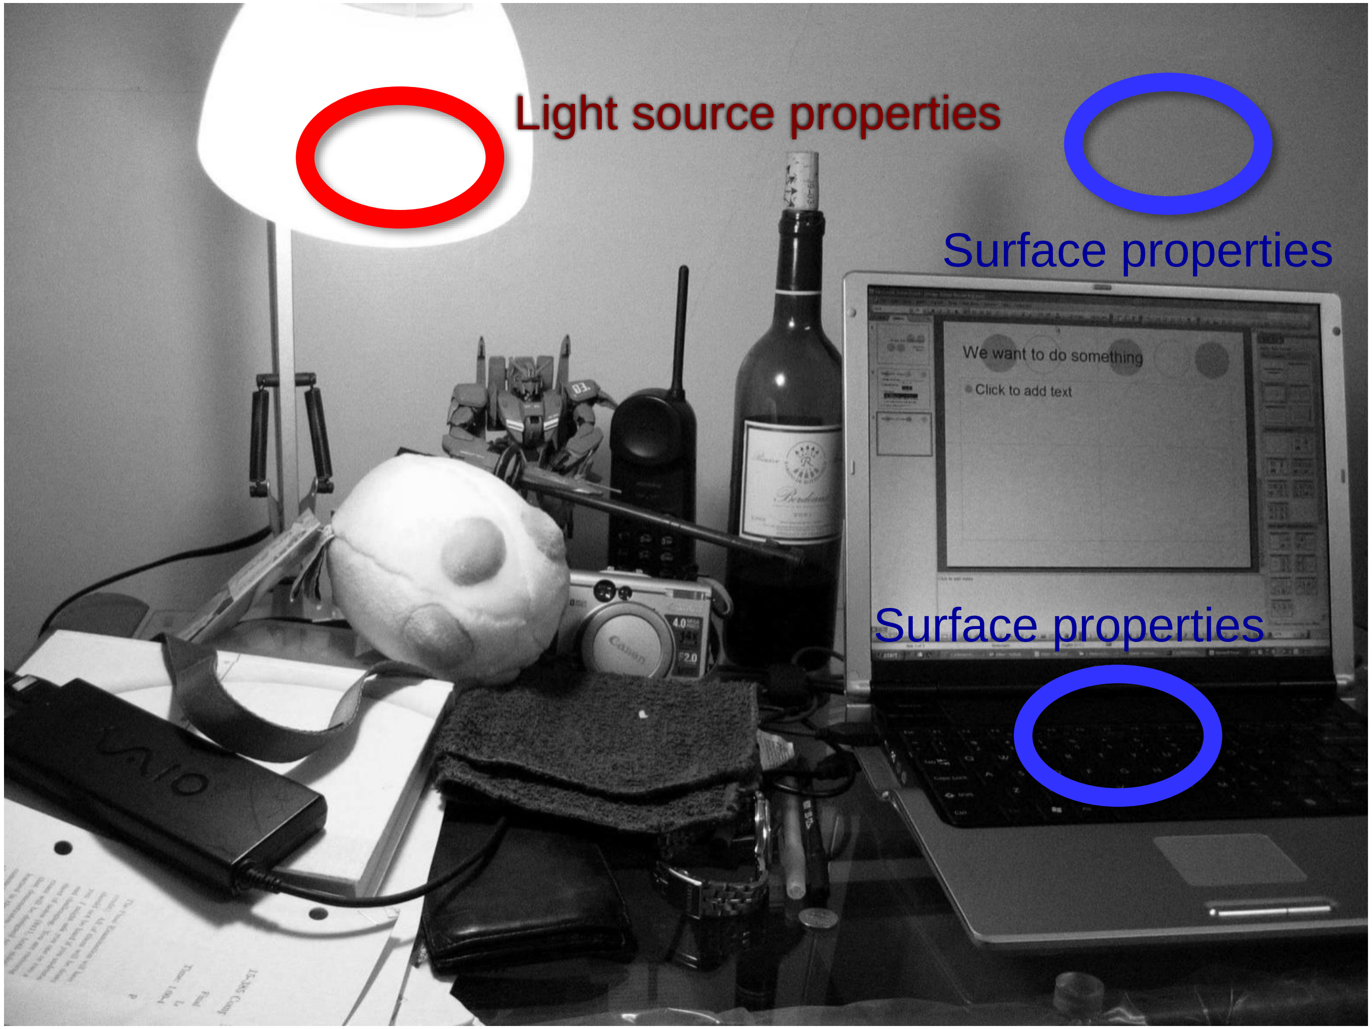

Radiometry

What determines the brightness of a pixel?

Radiometry

What determines the brightness of a pixel?

Radiometry

What determines the brightness of an image pixel?

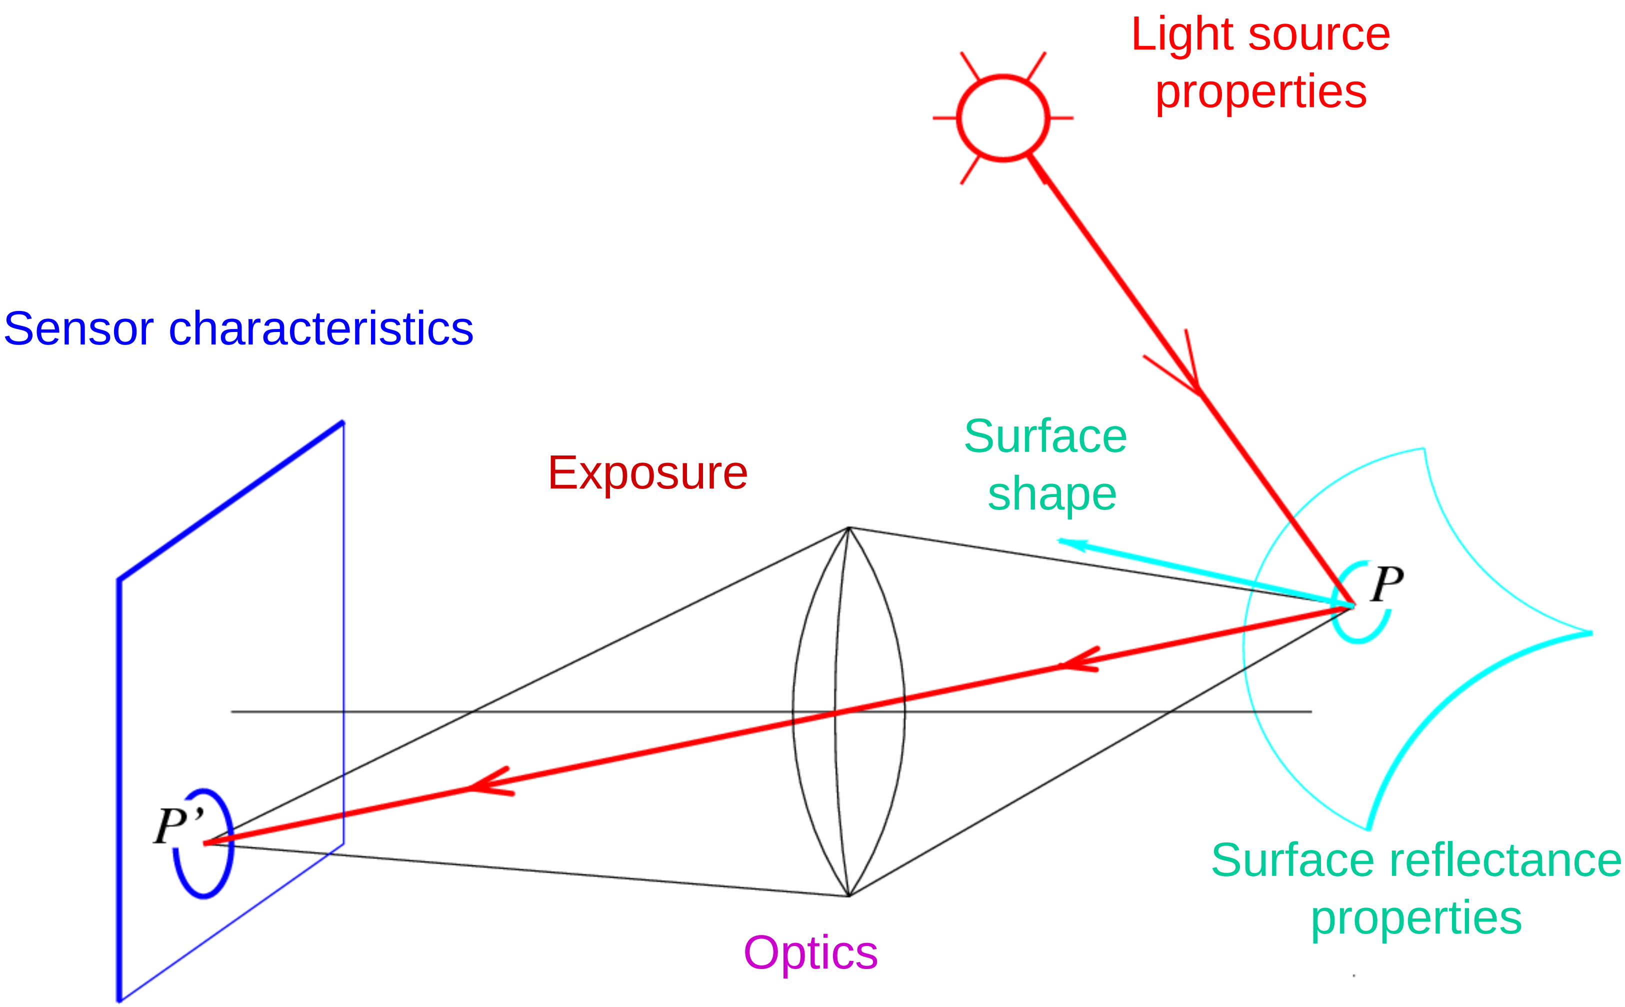

Slide by L. Fei-Fei

What is light?

Electromagnetic radiation (EMR) moving along rays in space

- R(λ) is EMR, measured in units of power (watts)
  - λ is wavelength

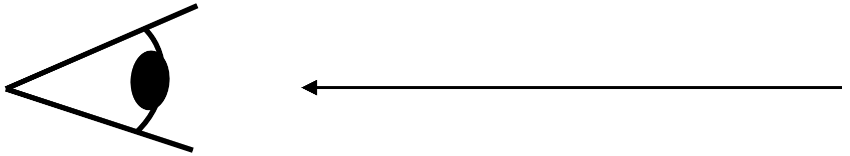

Light field

- We can describe all of the light in the scene by specifying the radiation (or **"radiance"** along all light rays) arriving at every point in space and from every direction

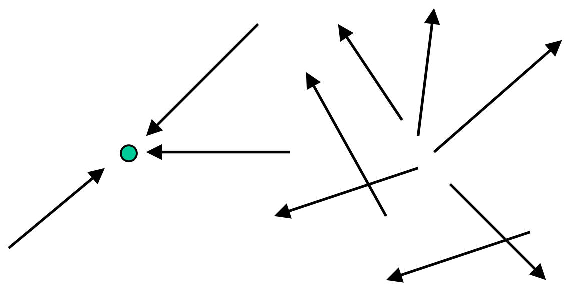

$$R(X, Y, Z, \theta, \phi, \lambda, t)$$

Color perception

Electromagnetic radiation (EMR) moving along rays in space

- R(λ) is EMR, measured in units of power (watts)
  - λ is wavelength

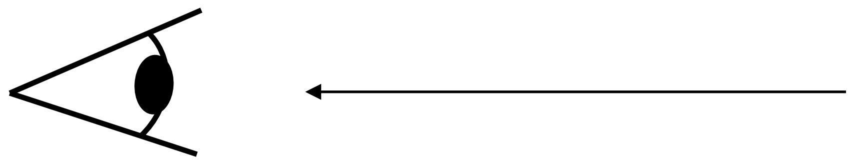

Perceiving light

- How do we convert radiation into “color”?
- What part of the spectrum do we see?

Visible light

We "see" electromagnetic radiation in a range of wavelengths

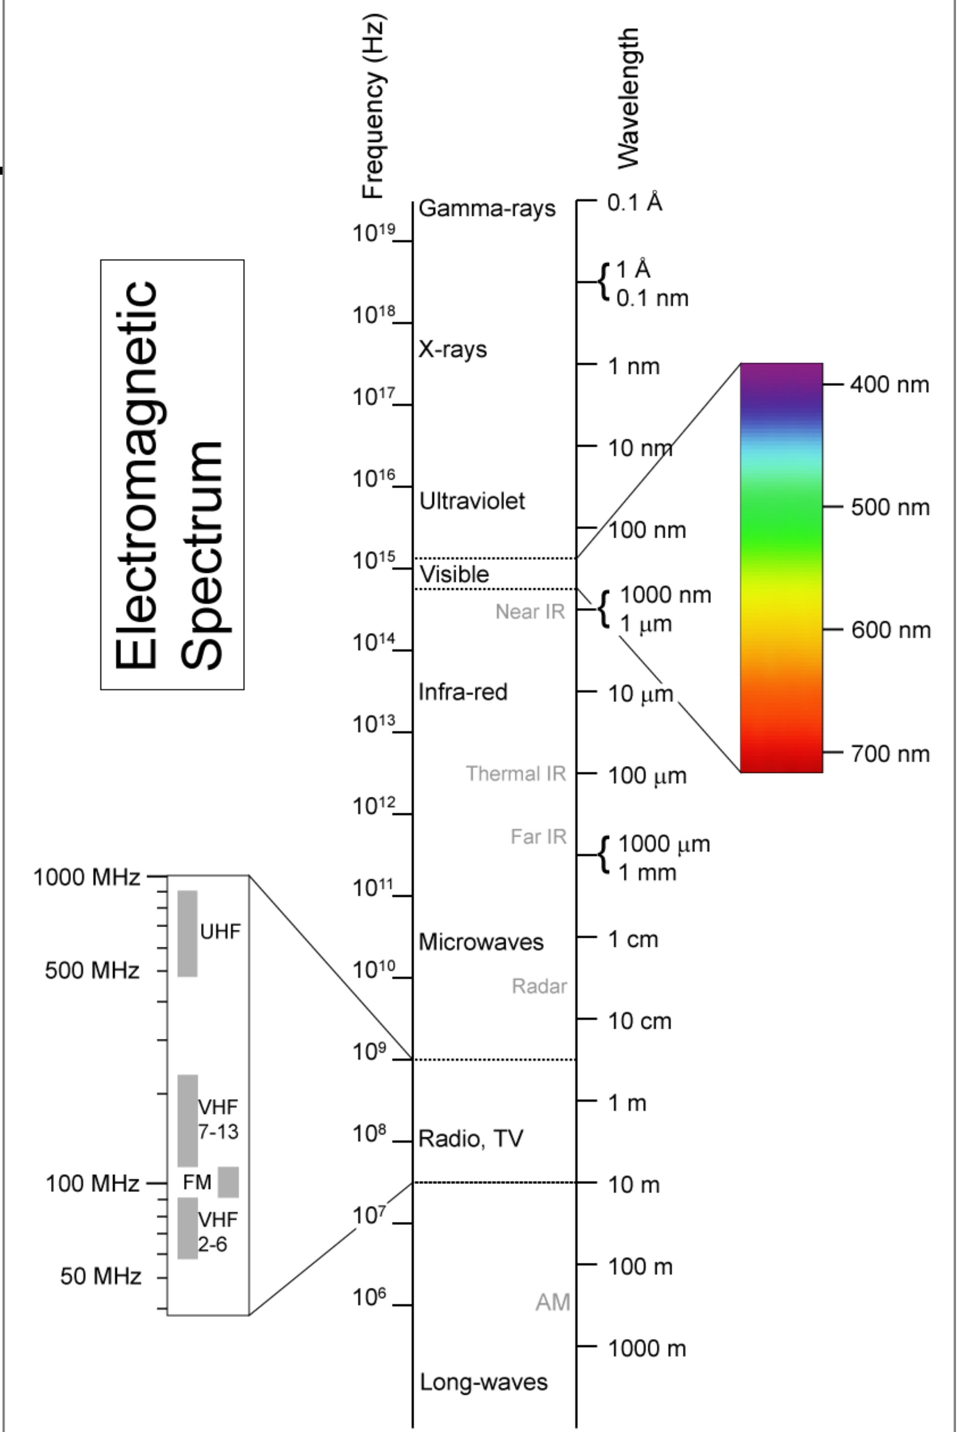

Light spectrum

The appearance of light depends on its power **spectrum**

- How much power (or energy) at each wavelength

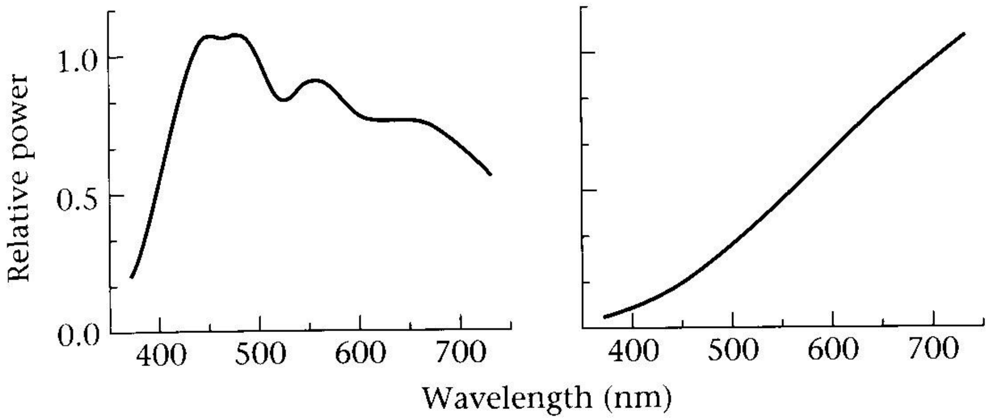

daylight

tungsten bulb

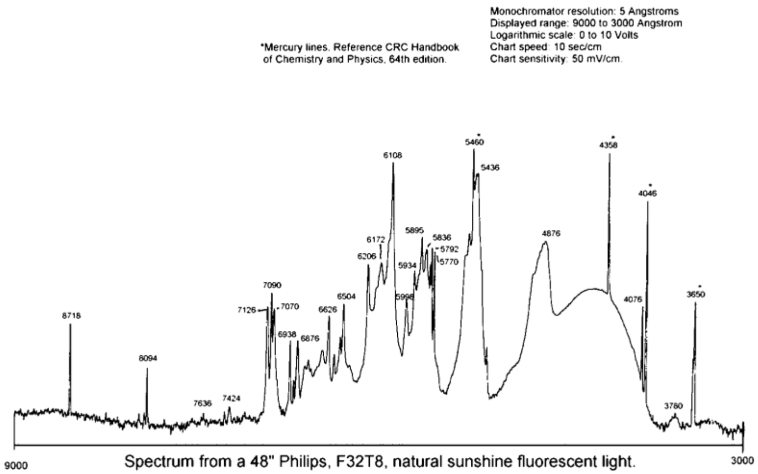

fluorescent bulb

Our visual system converts a light spectrum into “color”

- This is a rather complex transformation

The human visual system

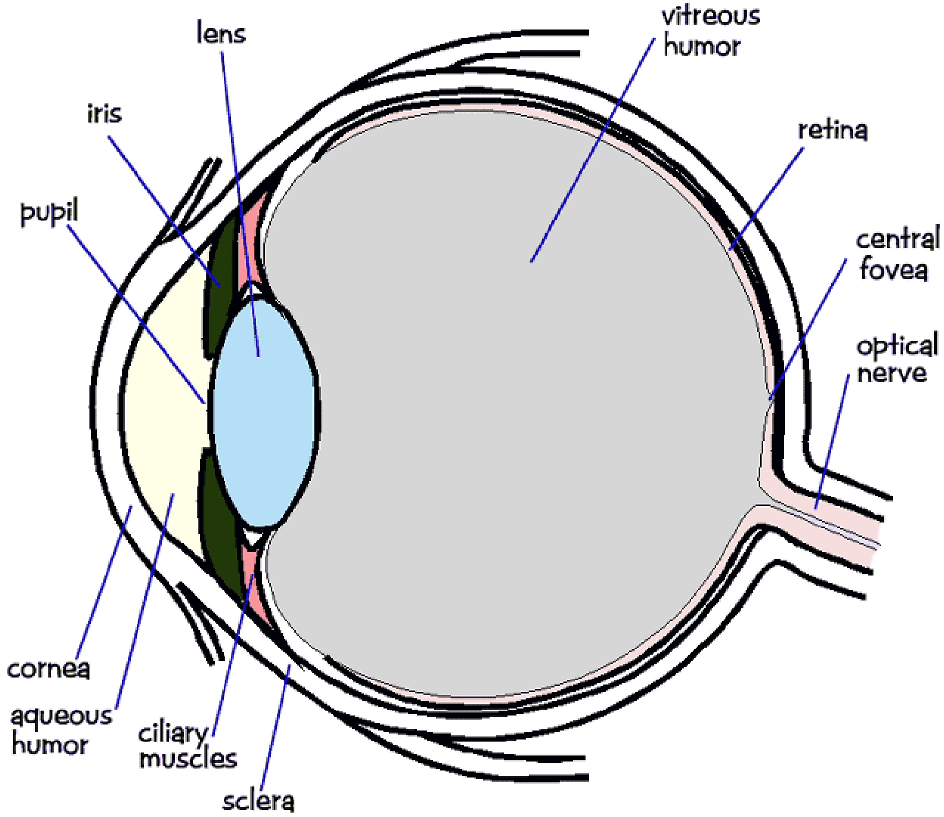

Color perception

- Light hits the retina, which contains photosensitive cells
  - rods and cones
- These cells convert the spectrum into a few discrete values

Density of rods and cones

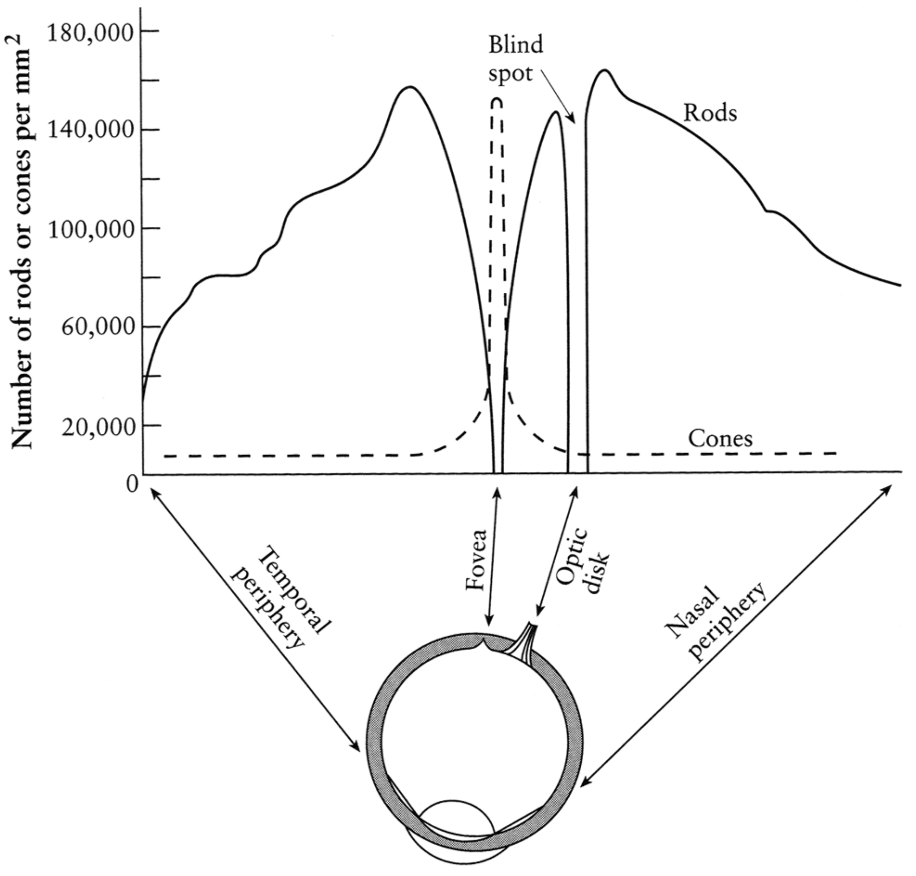

Rods and cones are *non-uniformly* distributed on the retina

- Rods responsible for intensity, cones responsible for color
- **Fovea** - Small region (1 or 2°) at the center of the visual field containing the highest density of cones (and no rods).
- Less visual acuity in the periphery—many rods wired to the same neuron

Demonstrations of visual acuity

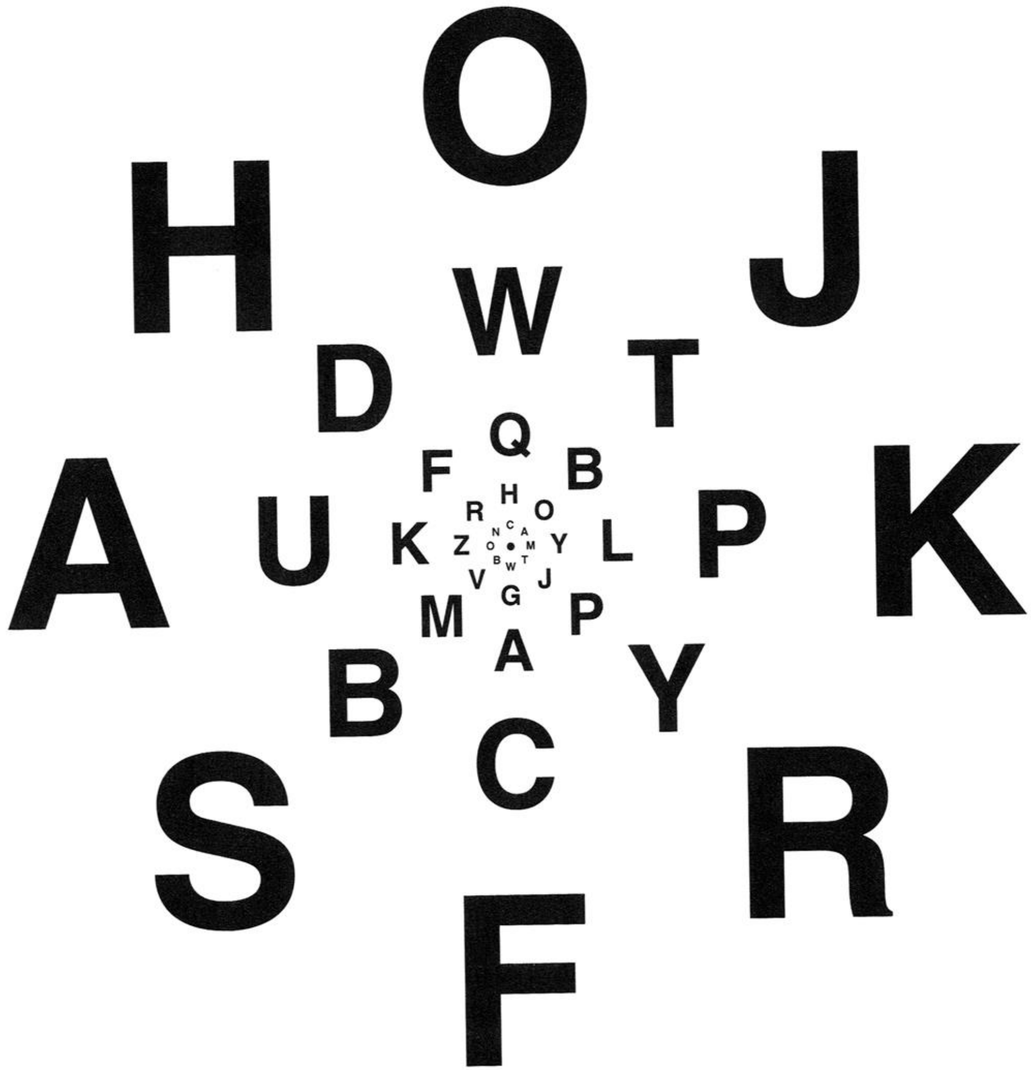

With one eye shut, at the right distance, all of these letters should appear equally legible (Glassner, 1.7).

Demonstrations of visual acuity

With left eye shut, look at the cross on the left. At the right distance, the circle on the right should disappear (Glassner, 1.8).

Brightness contrast and constancy

The apparent brightness depends on the surrounding region

- **brightness contrast**: a constant colored region seems lighter or darker depending on the surrounding intensity:

  - http://www.sandlotscience.com/Contrast/Checker_Board_2.htm

- **brightness constancy**: a surface looks the same under widely varying lighting conditions.

Light response is nonlinear

Our visual system has a large *dynamic range*

- We can resolve both light and dark things at the same time
- One mechanism for achieving this is that we sense light intensity on a *logarithmic scale*
  - an exponential intensity ramp will be seen as a linear ramp
- Another mechanism is *adaptation*
  - rods and cones adapt to be more sensitive in low light, less sensitive in bright light.

Visual dynamic range

Background	Luminance (candelas per square meter)
Horizon sky	
Moonless overcast night	0.00003
Moonless clear night	0.0003
Moonlit overcast night	0.003
Moonlit clear night	0.03
Deep twilight	0.3
Twilight	3
Very dark day	30
Overcast day	300
Clear day	3,000
Day with sunlit clouds	30,000
Daylight fog	
Dull	300–1,000
Typical	1,000–3,000
Bright	3,000–16,000
Ground	
Overcast day	30–100
Sunny day	300
Snow in full sunlight	16,000

FIGURE 1.13
Luminance of everyday backgrounds. *Source:* Data from Rea, ed., *Lighting Handbook 1984 Reference and Application*, fig. 3-44, p. 3-24.

A piece of white paper can be 1,000,000,000 times brighter in outdoor sunlight than in a moonless night.

BUT in a given lighting condition, light ranges over only about two orders of magnitude.

Visual dynamic range

If we were sensitive to this whole range all the time, we wouldn’t be able to discriminate lightness levels in a typical scene.

The visual system solves this problem by restricting the ‘dynamic range’ of its response to match the current overall or ‘ambient’ light level.

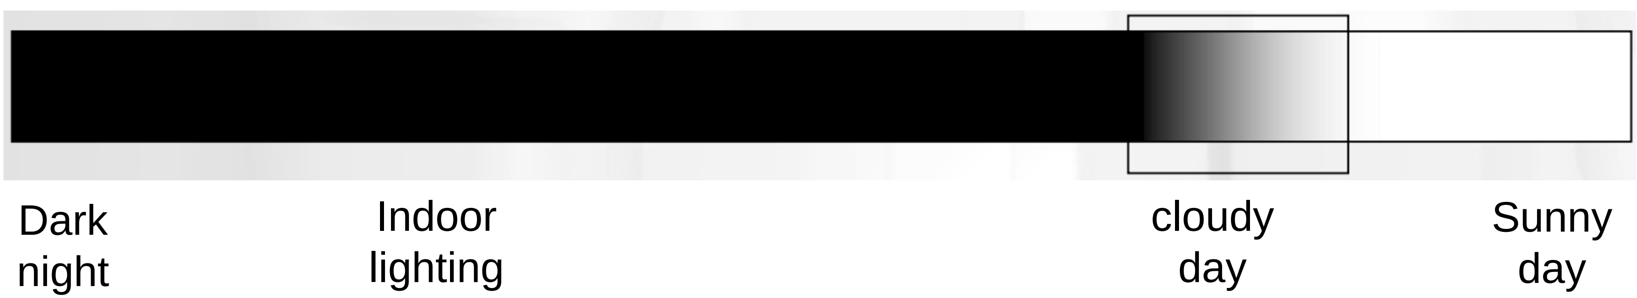

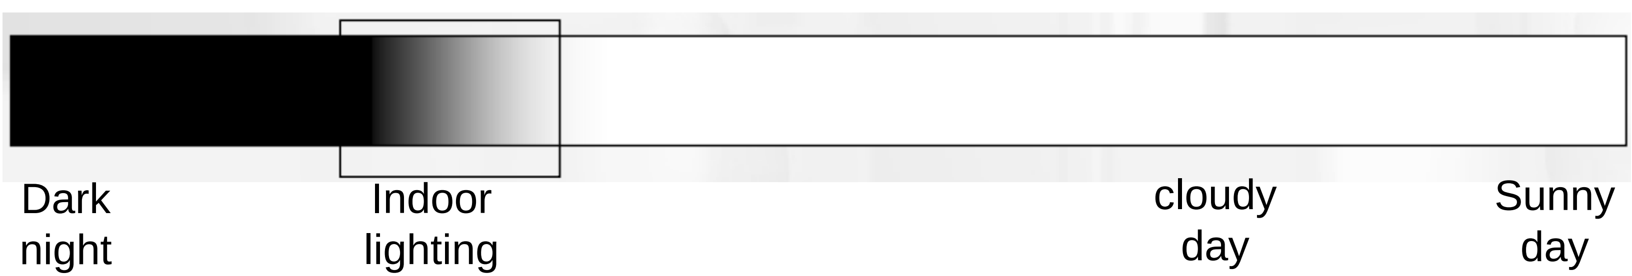

Color perception

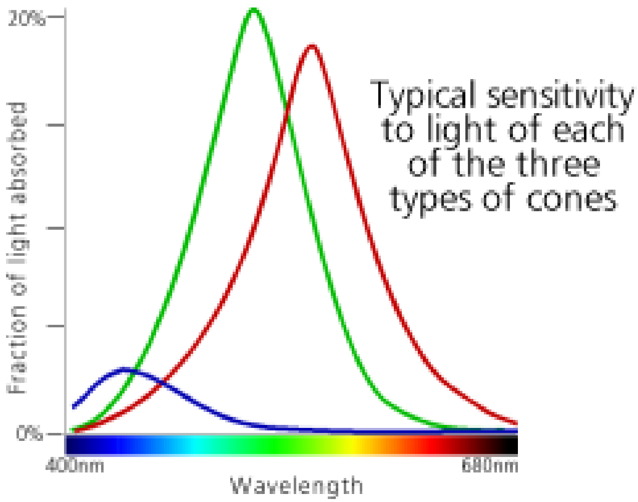

L response curve

Three types of cones

- Each is sensitive in a different region of the spectrum
  - but regions overlap
  - Short (S) corresponds to blue
  - Medium (M) corresponds to green
  - Long (L) corresponds to red
- Different sensitivities: we are more sensitive to green than red
  - varies from person to person (and with age)
- Colorblindness—deficiency in at least one type of cone

Color perception

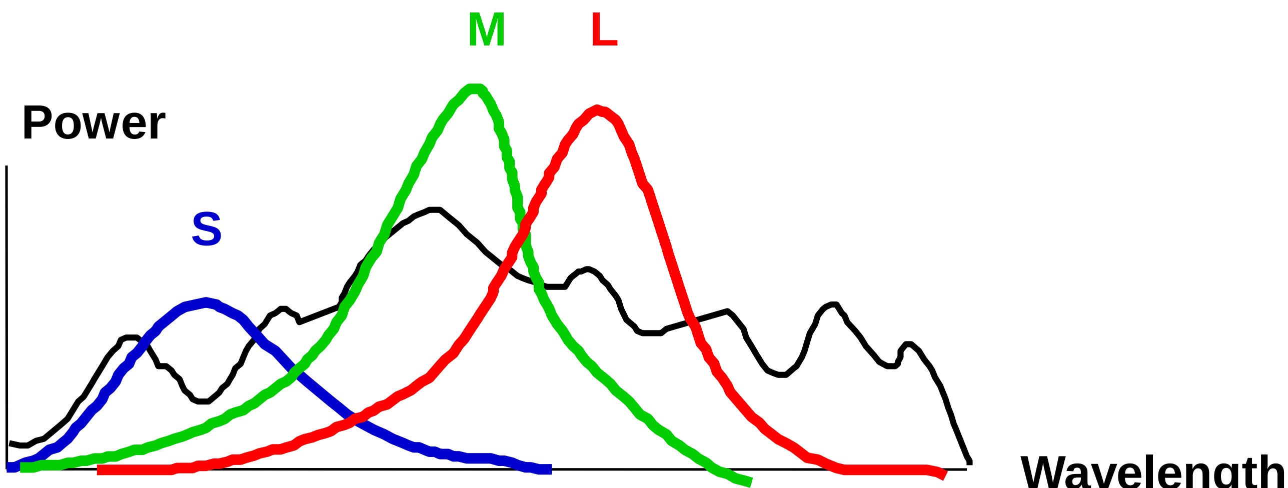

Rods and cones act as filters on the spectrum

- To get the output of a filter, multiply its response curve by the spectrum, integrate over all wavelengths
  - Each cone yields one number
- Q: How can we represent an entire spectrum with 3 numbers?
- A: We can't! Most of the information is lost.
  - As a result, two different spectra may appear indistinguishable
    - » such spectra are known as **metamers**
    - » http://www.cs.brown.edu/exploratories/freeSoftware/repository/edu/brown/cs/exploratories/applets/spectrum/metamers_guide.html

What kind of bulb is it?

http://www.chemistryland.com/CHM107Lab/Exp7/Spectroscope/Spectroscope.html

Perception summary

The mapping from radiance to perceived color is quite complex!

- We throw away most of the data
- We apply a logarithm
- Brightness affected by pupil size
- Brightness contrast and constancy effects

The same is true for cameras

- But we have tools to correct for these effects
  - (Computational Photography)

Light transport

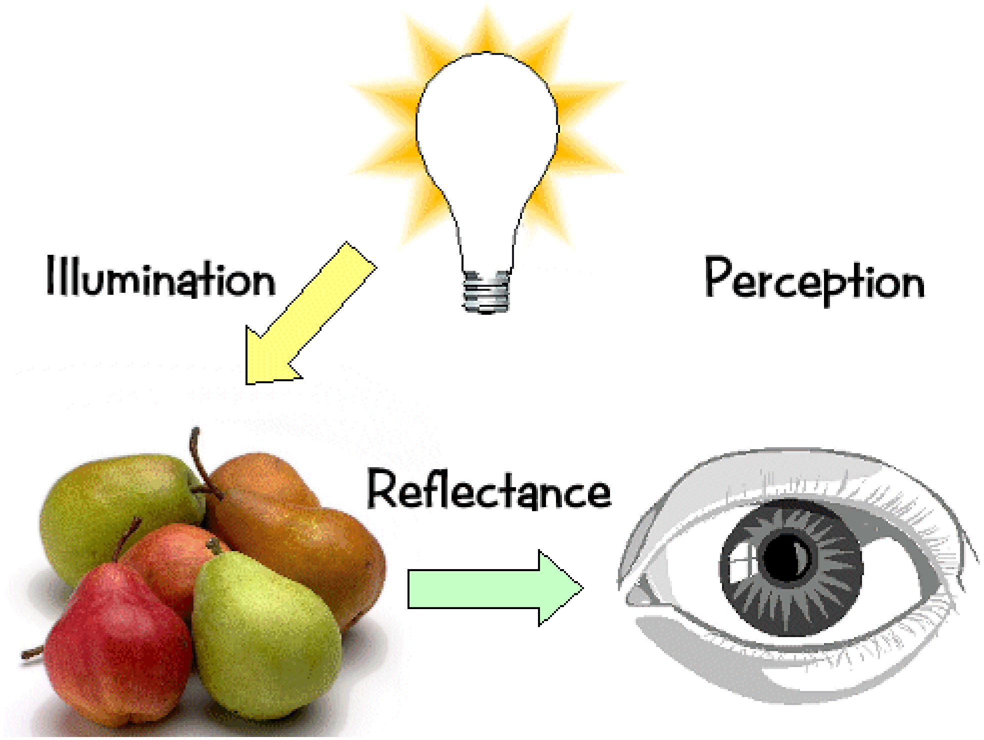

Light sources

Basic types

- point source
- directional source
  - a point source that is infinitely far away
- area source
  - a union of point sources

More generally

- a light field can describe *any* distribution of light sources

What happens when light hits an object?

Modeling Image Formation

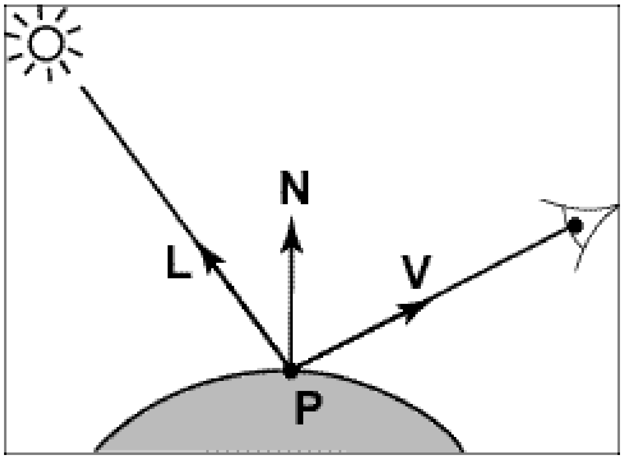

We need to reason about:

- How light interacts with the scene
- How a pixel value is related to light energy in the world

Track a “ray” of light all the way from light source to the sensor

Directional Lighting

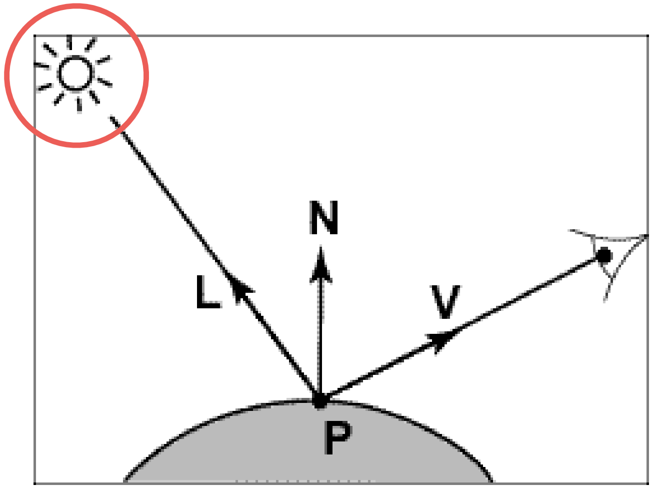

- Key property: all rays are parallel
- Equivalent to an infinitely distant point source

Lambertian Reflectance

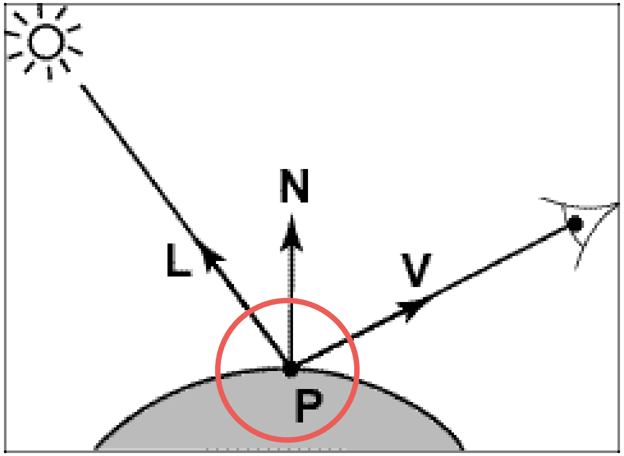

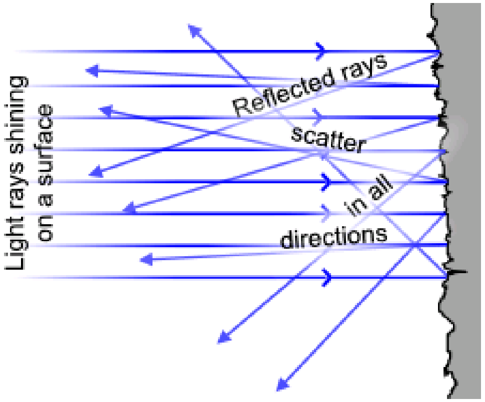

$$I = N \cdot L$$

Image intensity $=$ Surface normal $\cdot$ Light direction

Image intensity $\propto$ cos(angle between N and L)

Materials - Three Forms

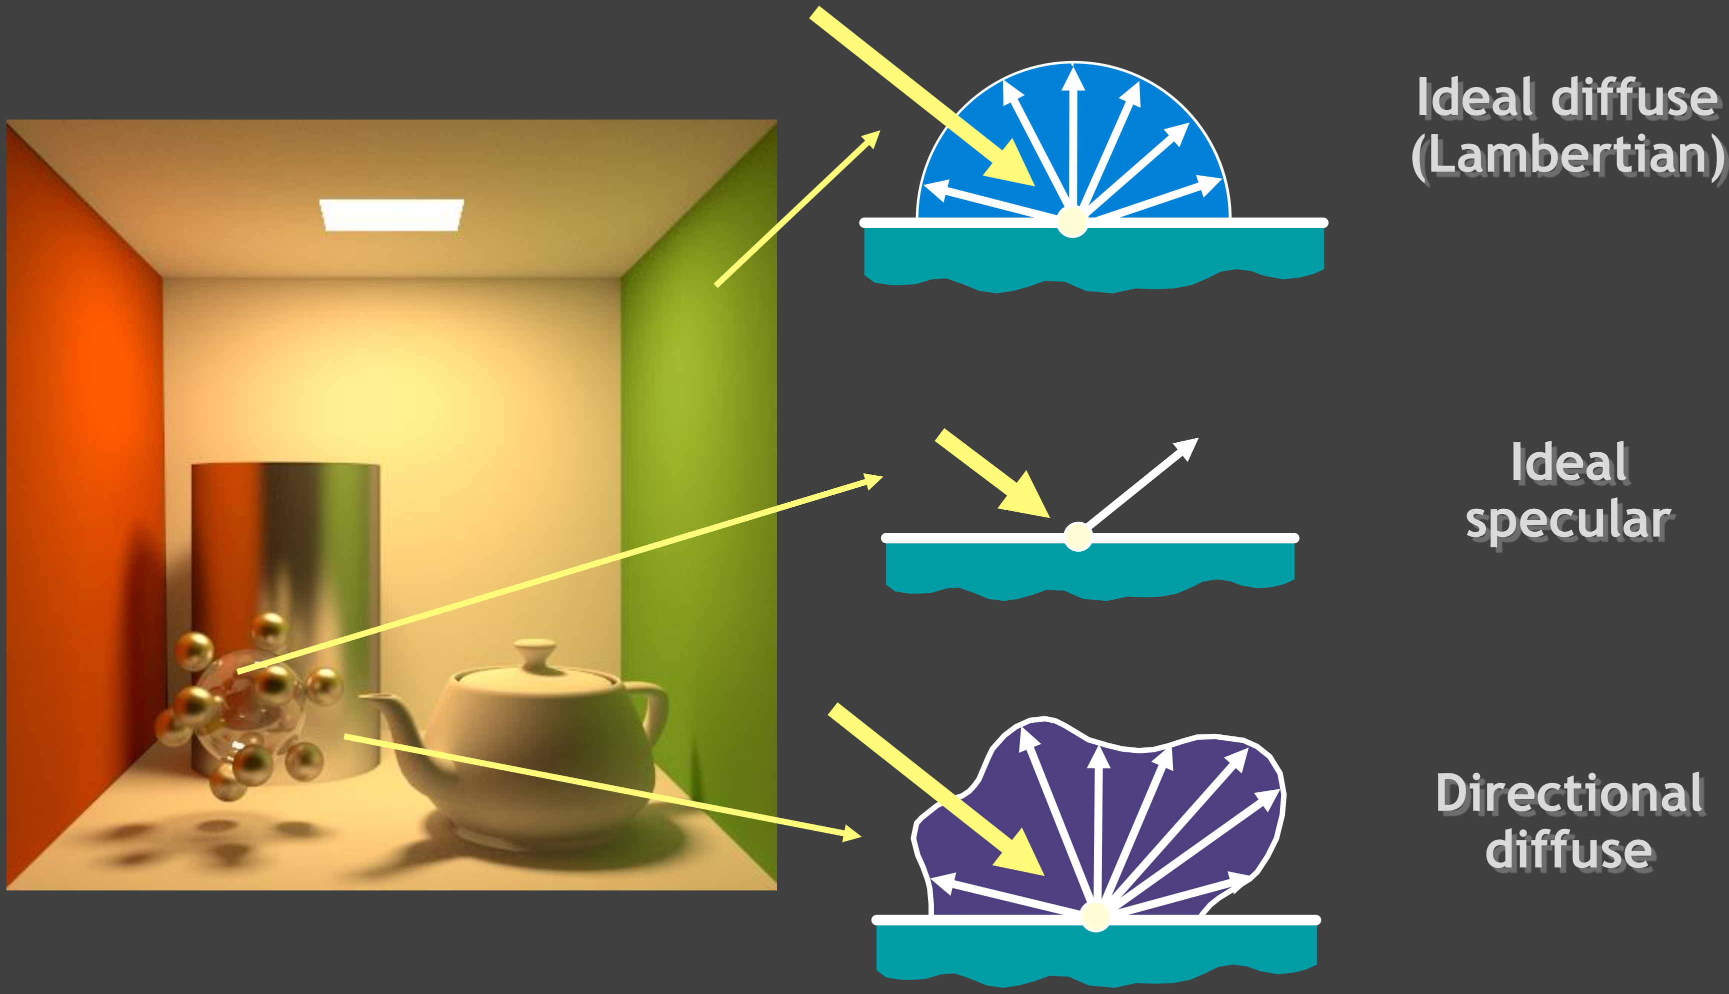

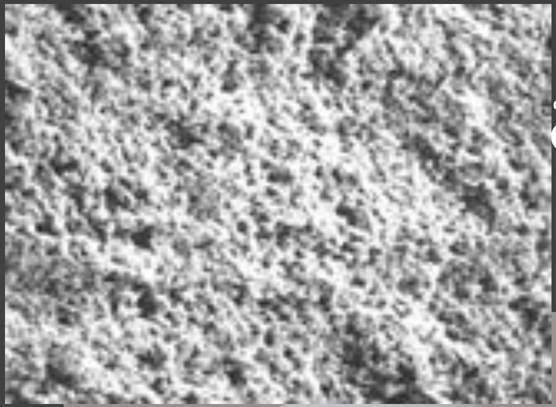

Ideal diffuse (Lambertian)

Ideal specular

Directional diffuse

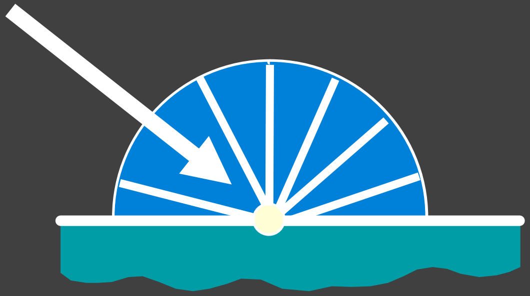

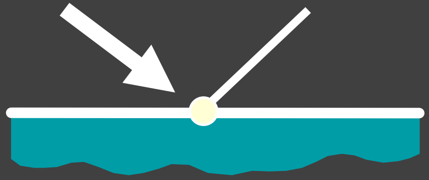

Ideal Diffuse Reflection

- Characteristic of multiple scattering materials
- An idealization but reasonable for matte surfaces

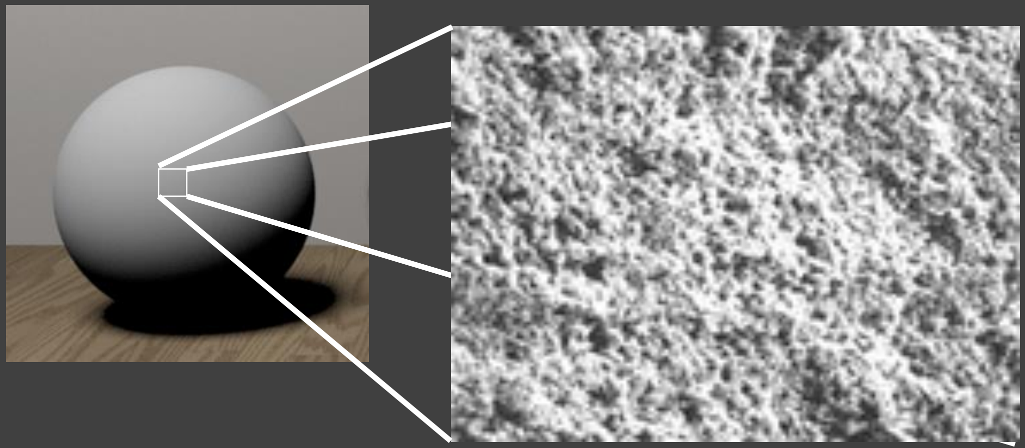

Lambertian Reflectance

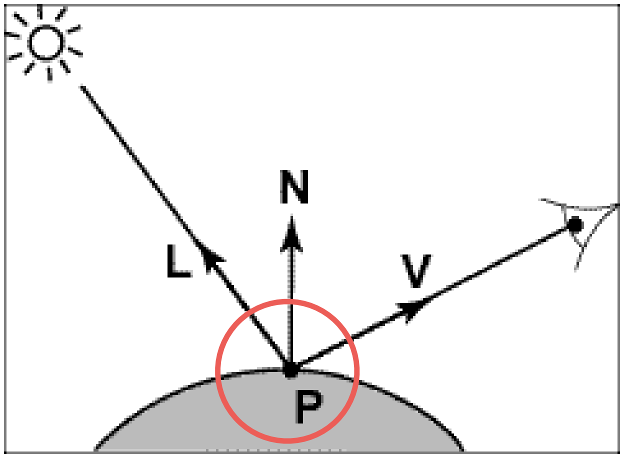

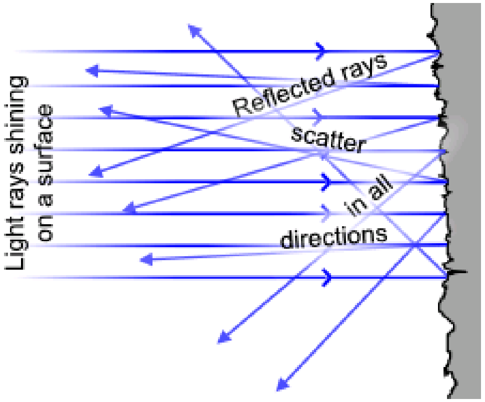

1. Reflected energy is proportional to cosine of angle between L and N **(incoming)**

2. Measured intensity is viewpoint-independent **(outgoing)**

Lambertian Reflectance: Incoming

1. Reflected energy is proportional to cosine of angle between L and N

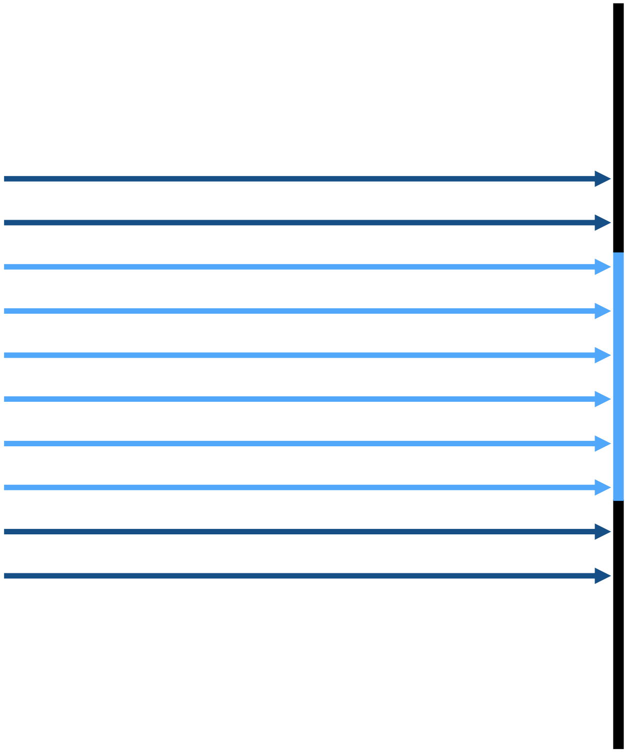

Lambertian Reflectance: Incoming

1. Reflected energy is proportional to cosine of angle between L and N

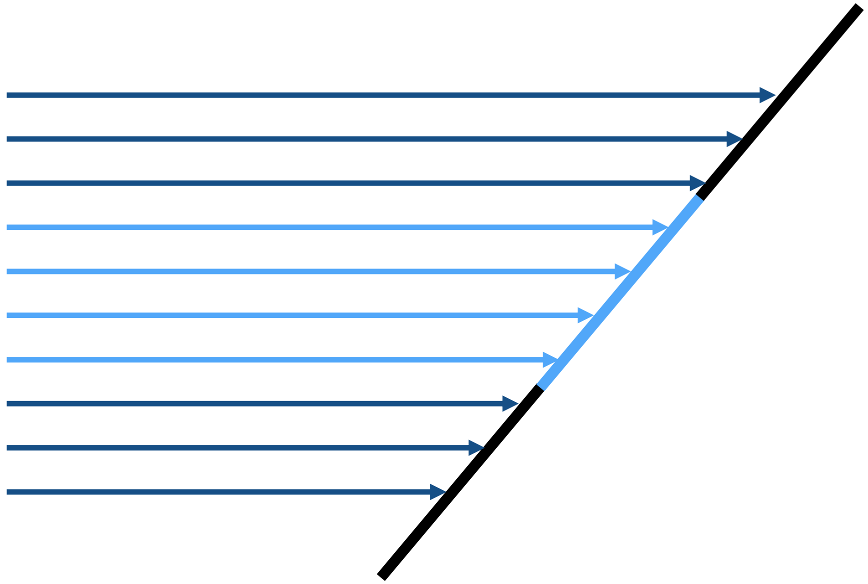

Lambertian Reflectance: Incoming

1. Reflected energy is proportional to cosine of angle between L and N

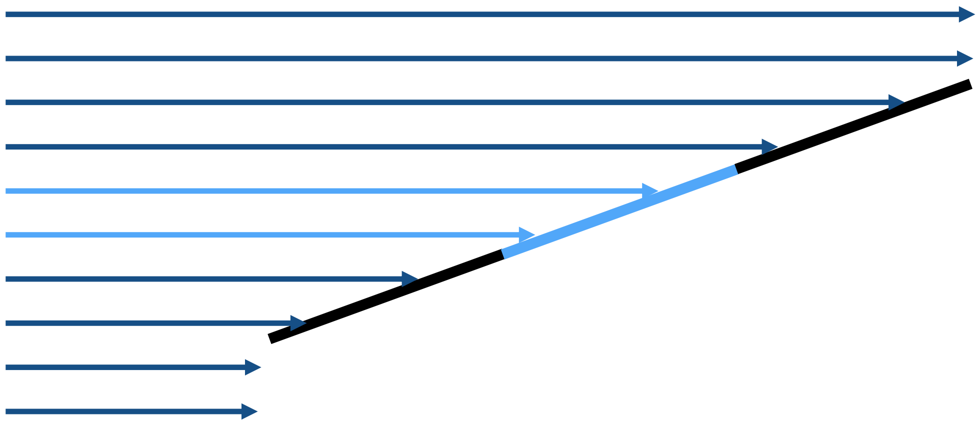

Light hitting surface is proportional to the **cosine**

Lambertian Reflectance: Outgoing

1. Radiance (what we see) is viewpoint-independent

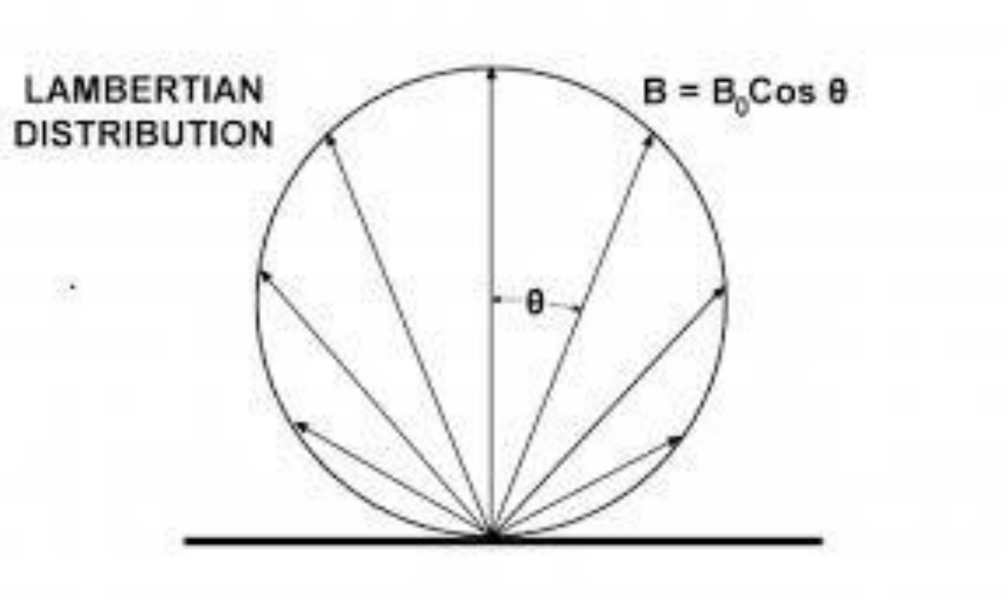

Lambertian Reflectance: Outgoing

1. Radiance (what the eye sees) is viewpoint-independent

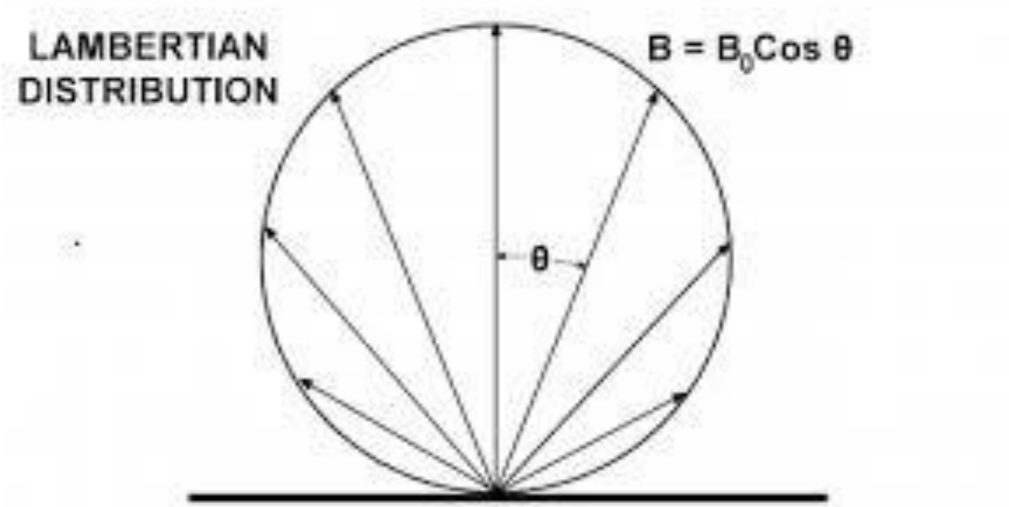

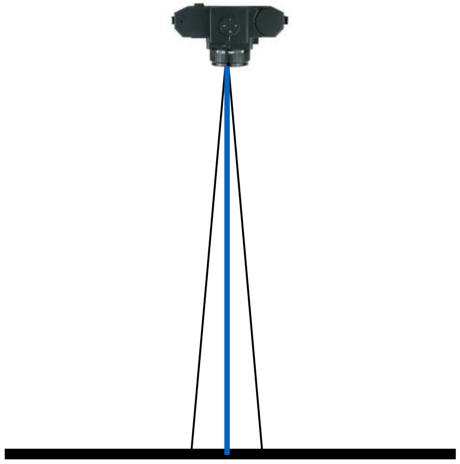

Lambertian Reflectance: Outgoing

1. Measured intensity is viewpoint-independent

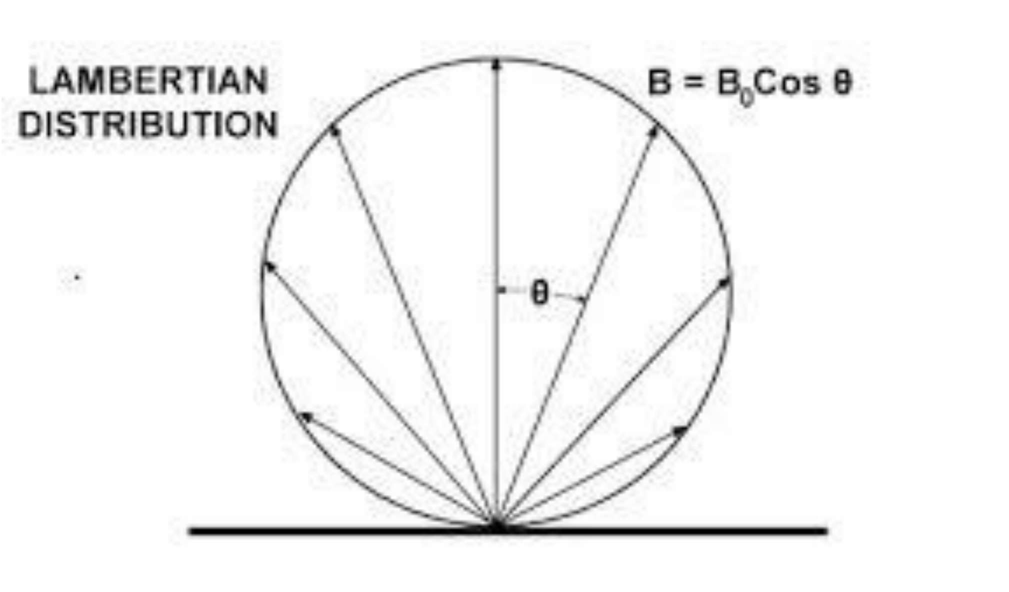

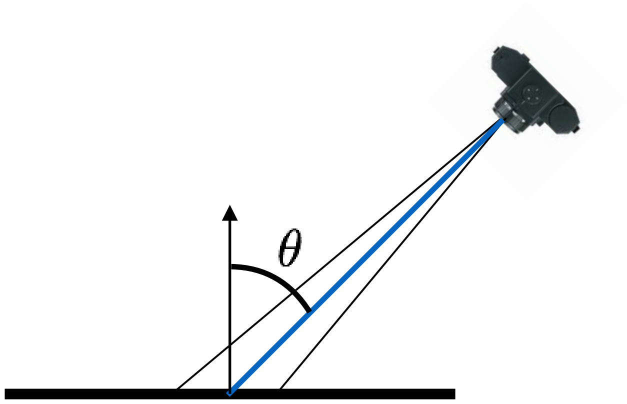

Lambertian Reflectance: Outgoing

1. Measured intensity is viewpoint-independent

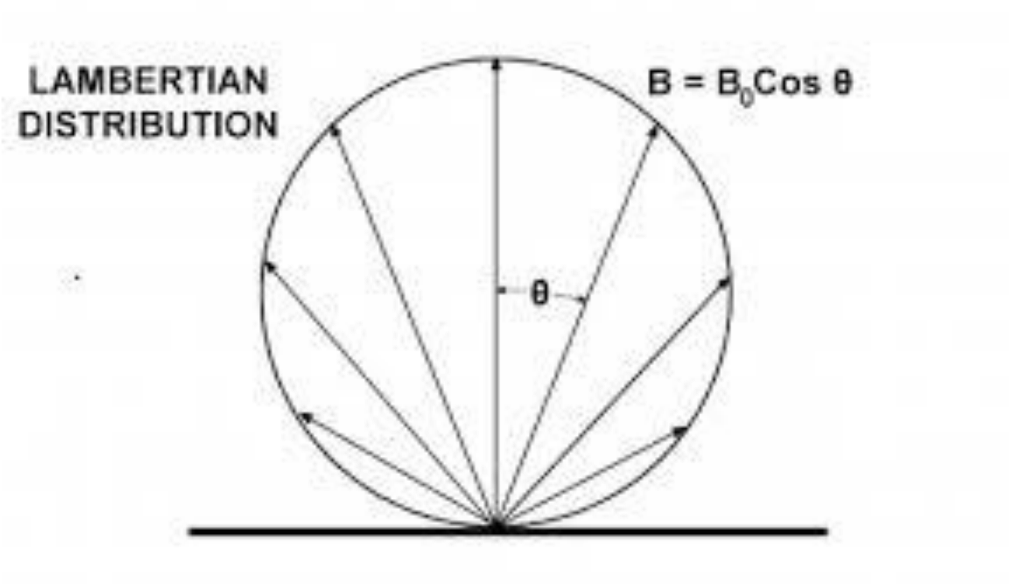

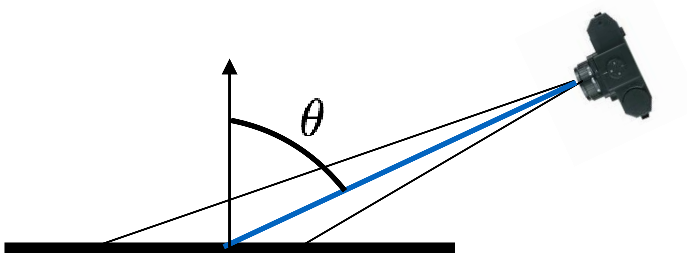

A cos (θ)

Radiance
(what eye sees)

$$\propto B_0 \cos(\theta) \frac{1}{\cos(\theta)}$$

Image Formation Model: Final

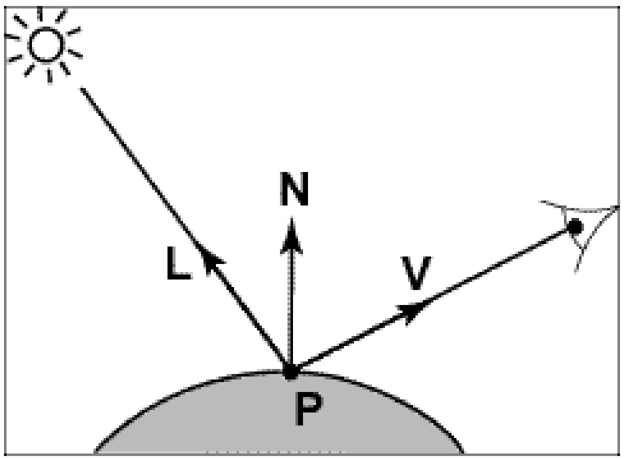

$$I = k_d \mathbf{N} \cdot \mathbf{L}$$

1. Diffuse albedo: what fraction of incoming light is reflected?
   - Introduce scale factor k_d
2. Light intensity: how much light is arriving?
   - Compensate with camera exposure (global scale factor)
3. Camera response function
   - Assume pixel value is linearly proportional to incoming energy (perform radiometric calibration if not)

A Single Image: Shape from Shading

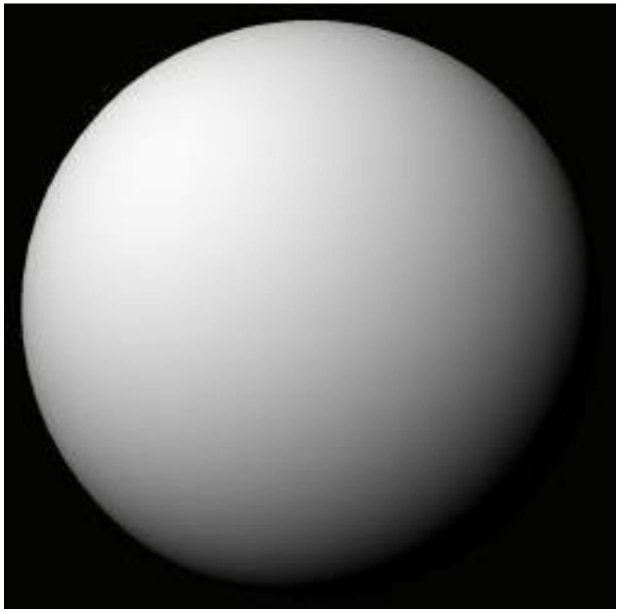

$$I = k_d \mathbf{N} \cdot \mathbf{L}$$

Assume k_d is 1 for now.

What can we measure from one image?

- $\cos^{-1}(I)$ is the angle between N and L
- Add assumptions:
  - Constant albedo
  - A few known normals (e.g. silhouettes)
  - Smoothness of normals

In practice, SFS doesn't work very well:
assumptions are too restrictive,
too much ambiguity in nontrivial scenes.

Application: Detecting composite photos

Real photo

Fake photo

Questions?