CS6670: Computer Vision

Noah Snavely

Lecture 21: Light, reflectance and photometric stereo

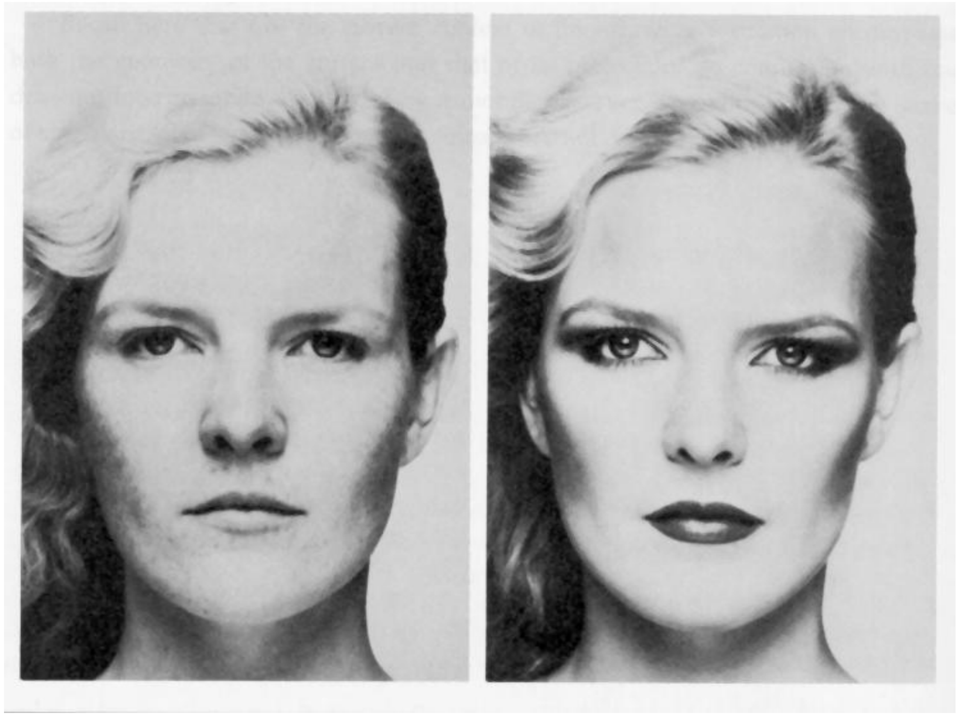

Announcements

- Final projects
  - Midterm reports due November 24 (next Tuesday) by 11:59pm (upload to CMS)
  - State the problem
  - Describe progress so far, any problems that have come up

What is light?

Electromagnetic radiation (EMR) moving along rays in space

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

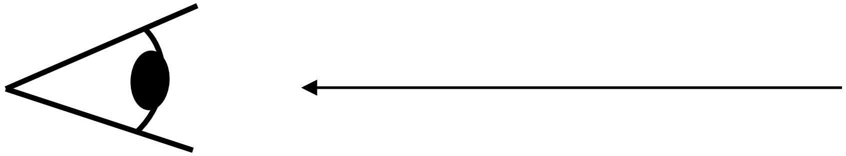

Perceiving light

- 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

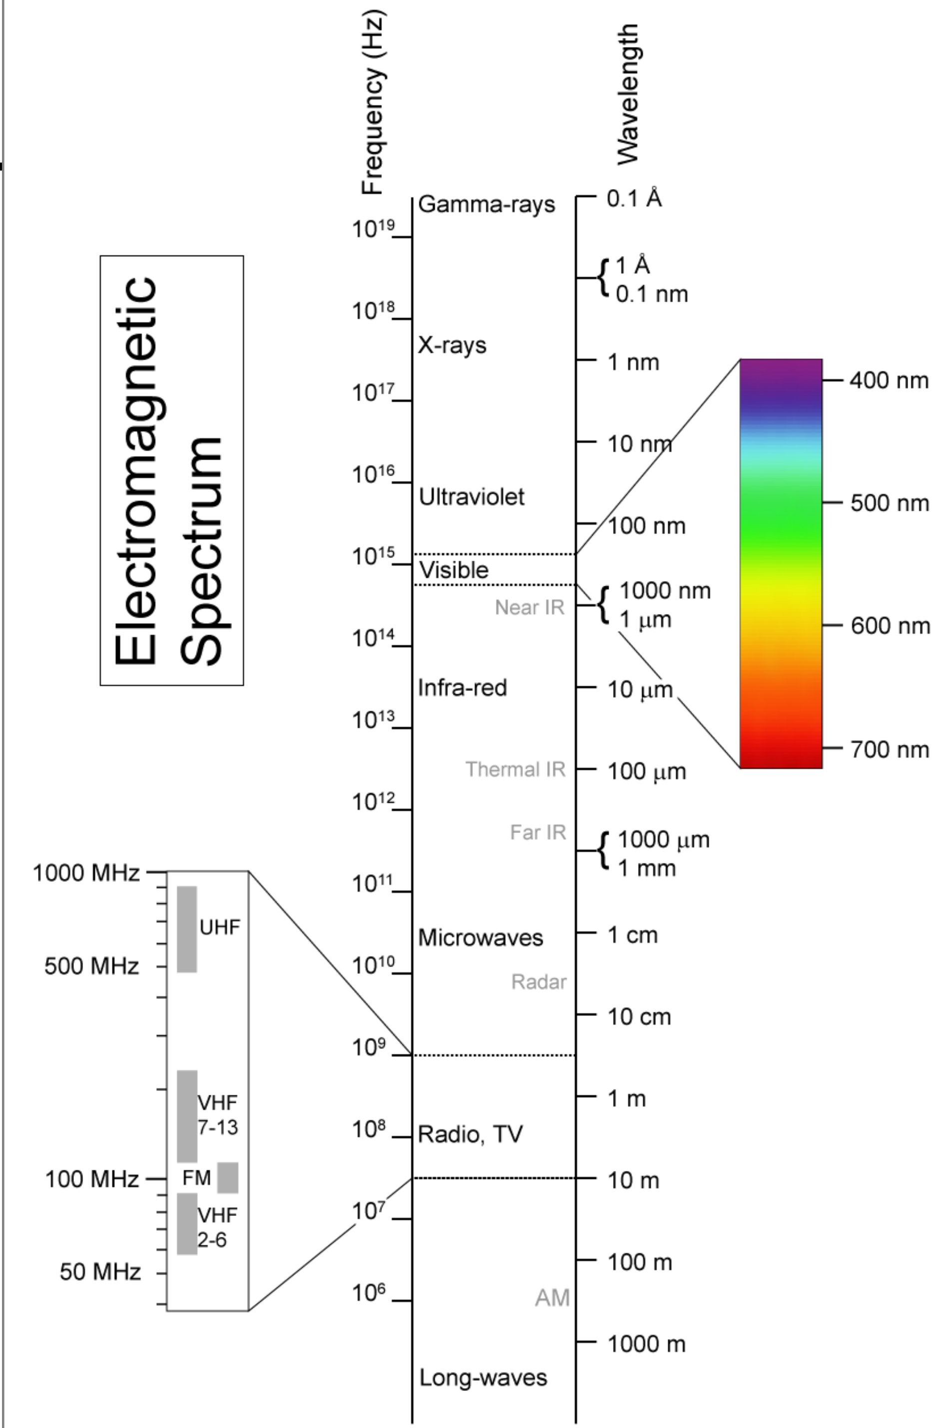

Light spectrum

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

- How much power (or energy) at each wavelength

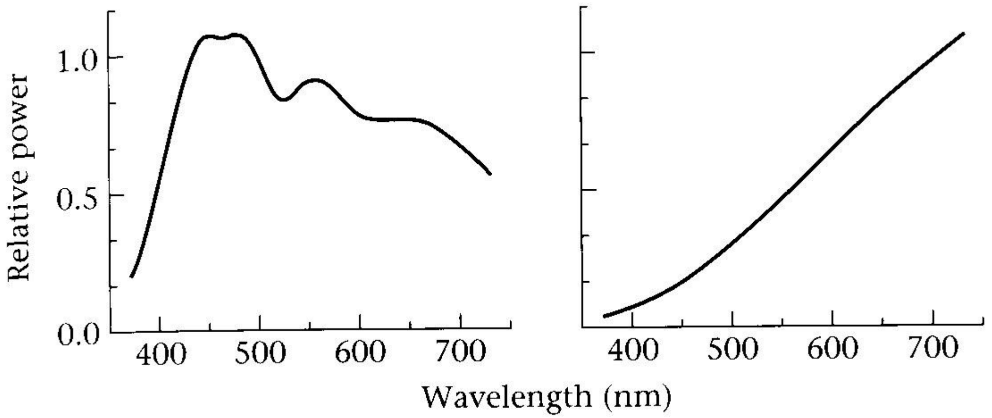

daylight

tungsten bulb

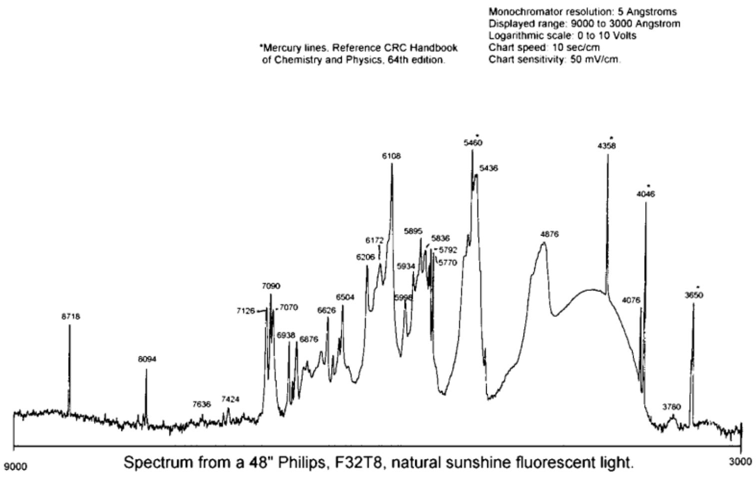

fluorescent light

Our visual system converts a light spectrum into “color”

- This is a rather complex transformation

The human visual system

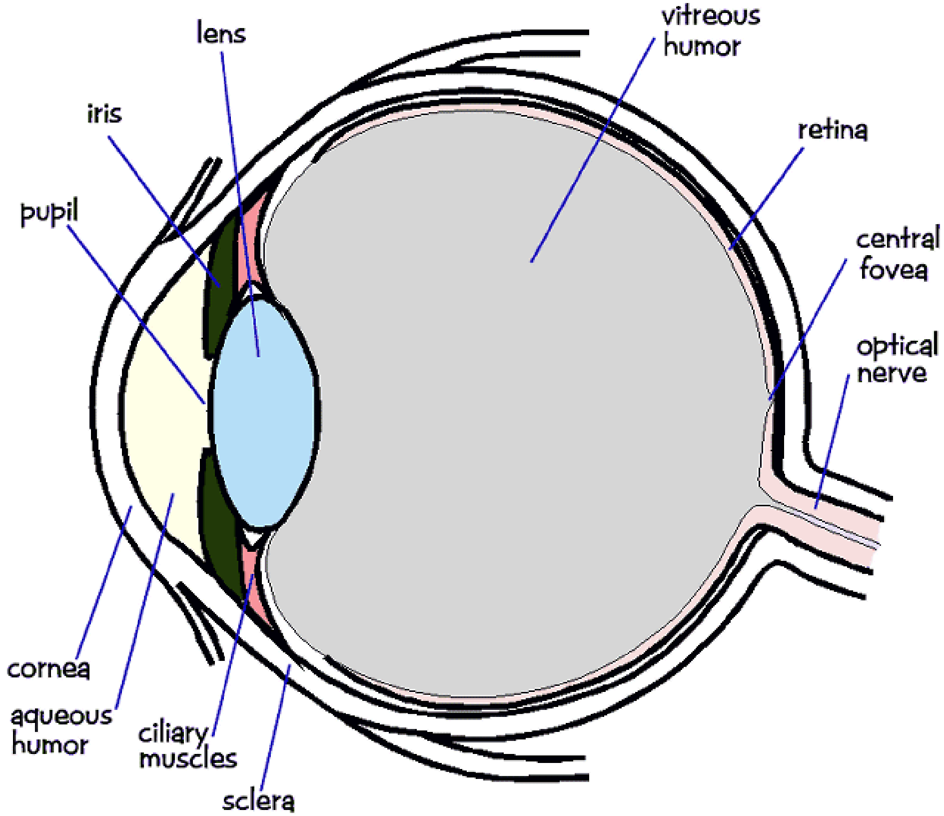

Color perception

- 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

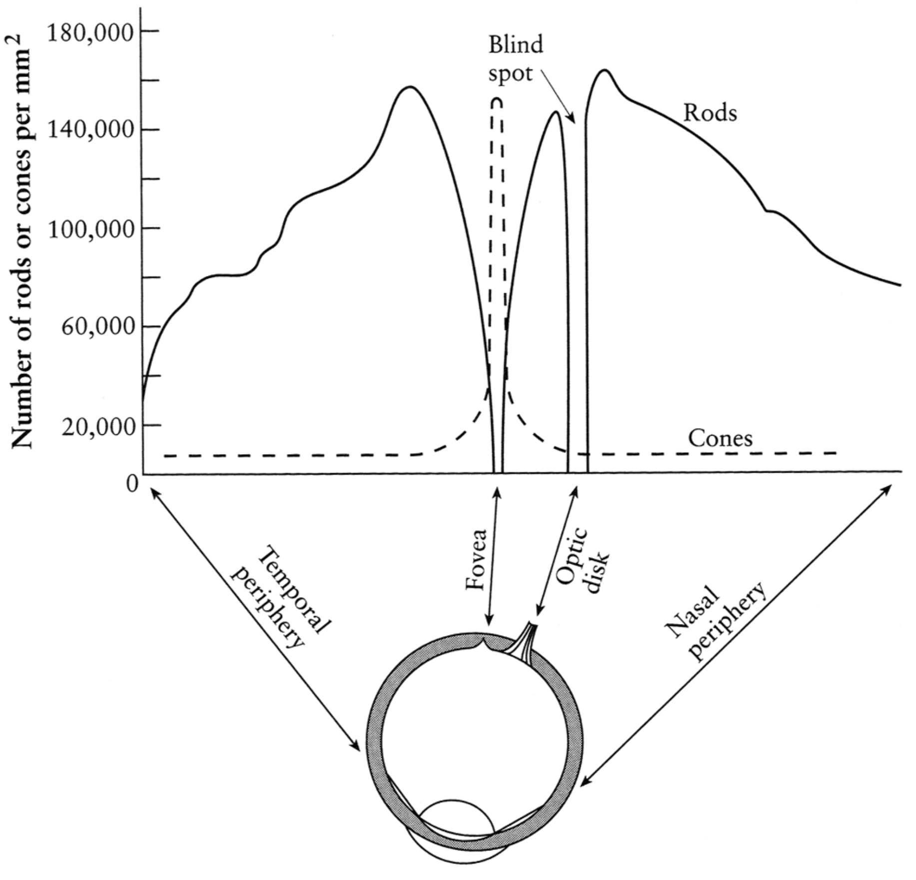

Rods and cones are *non-uniformly* distributed on the retina

- 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

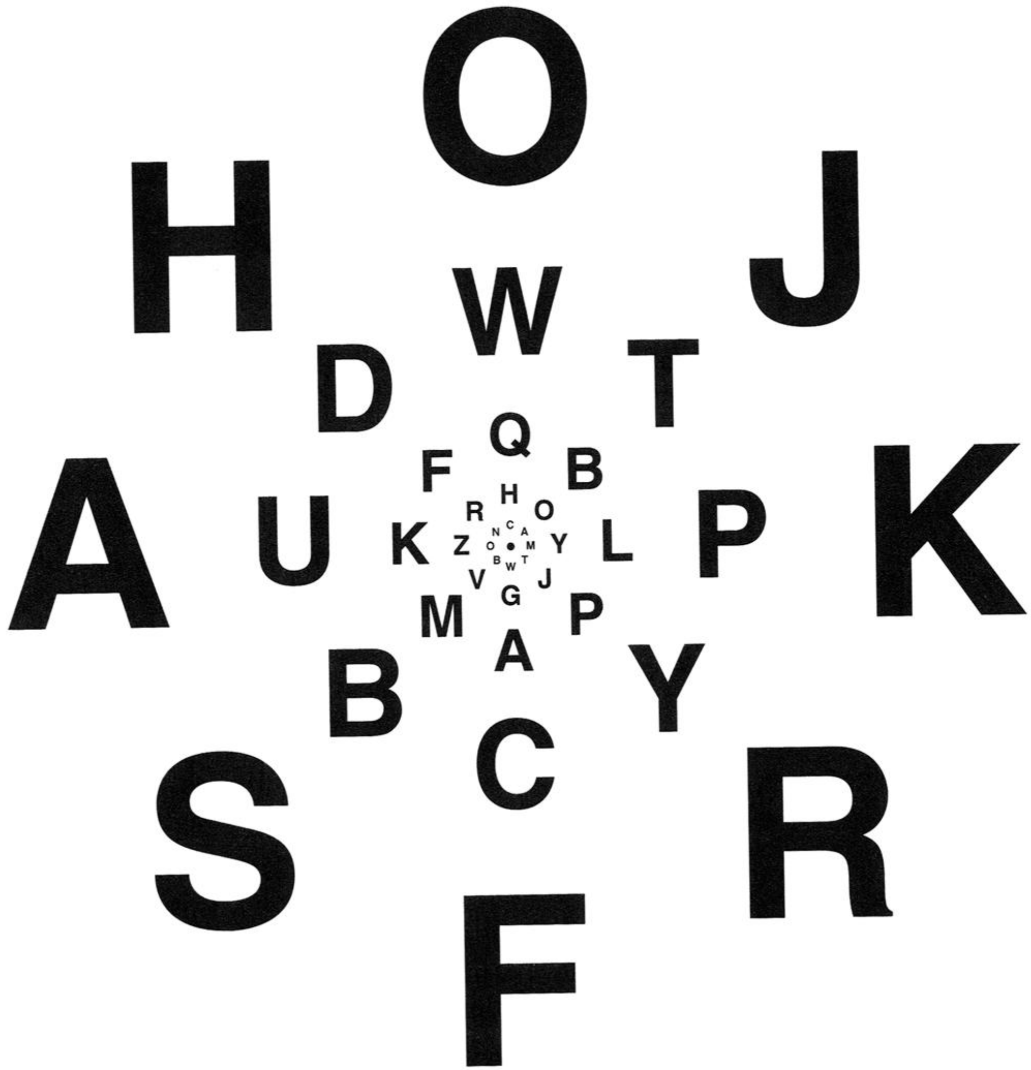

With one eye shut, at the right distance, all of these letters should appear equally legible (Glassner, 1.7).

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.

“Approximate brightness constancy, a similar effect, makes us tend to see objects in terms of their *reflecting power rather than the amount of light they actually reflect.* Thus we can almost always identify *a* piece of white paper as white even though it is placed in shadow where it actually reflects much less light to the eye than a piece of gray paper in full illumination.”

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.

Color perception

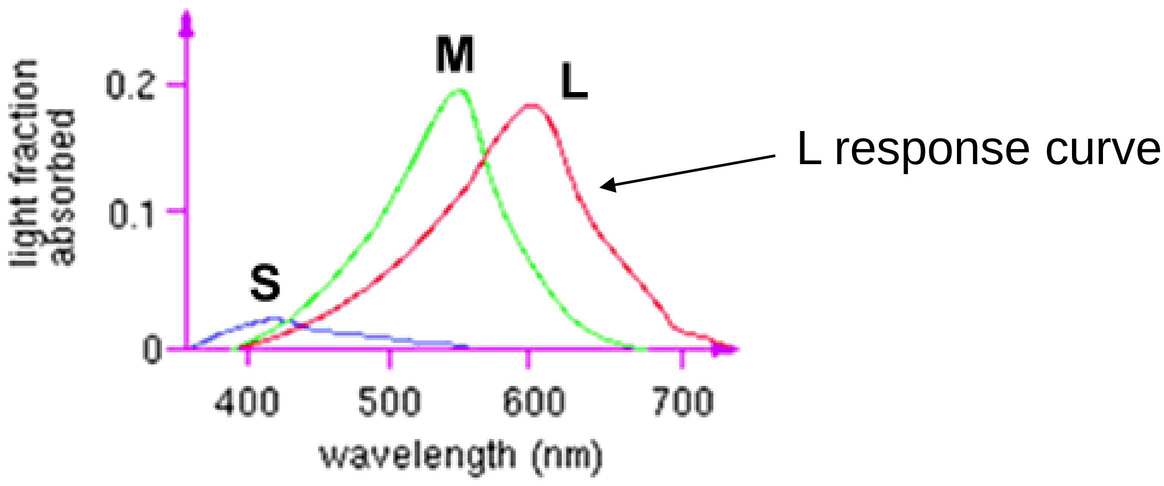

Three types of cones

- 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
- Colorblindness—deficiency in at least one type of cone

Color perception

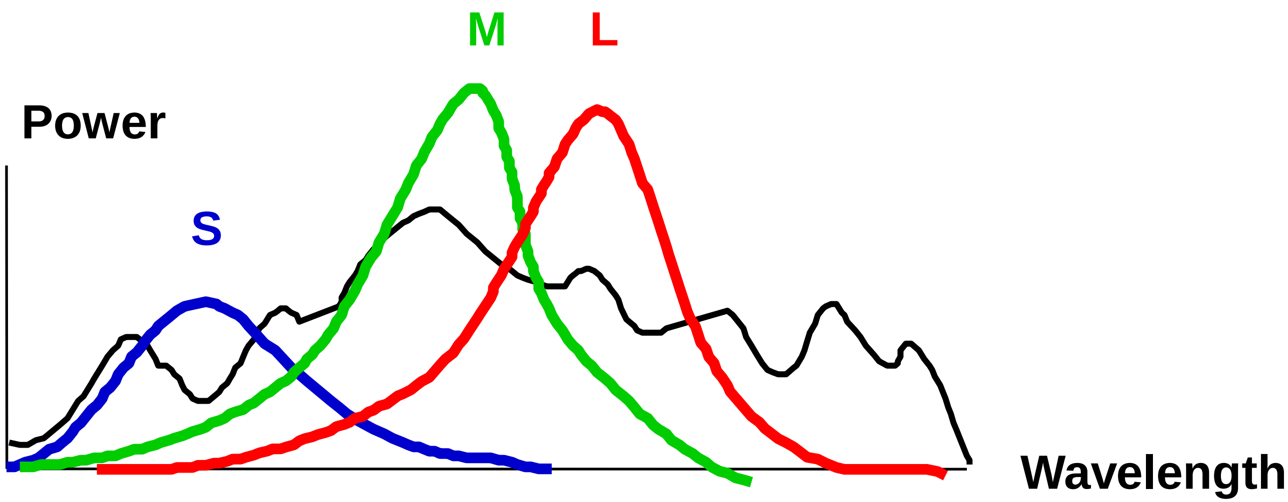

Rods and cones act as filters on the spectrum

- 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

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
- Afterimages

The same is true for cameras

- But we have tools to correct for these effects
  - Coming soon: Computational Photography lecture

Light transport

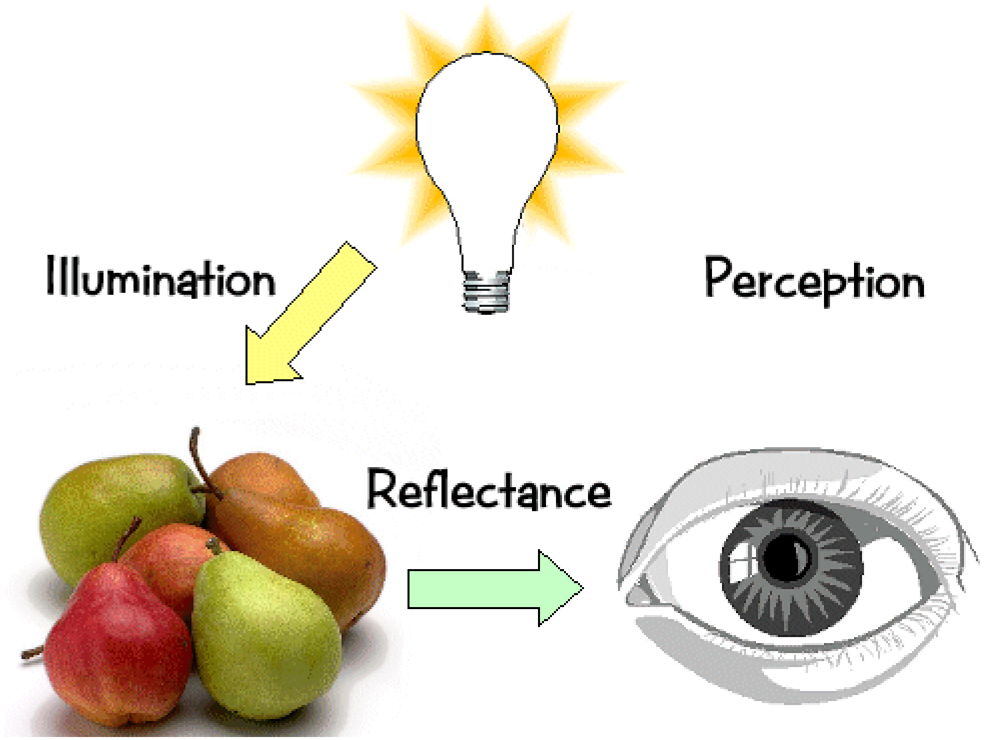

Light sources

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?

Reflectance spectrum (albedo)

To a first approximation, surfaces absorb some wavelengths of light and reflect others

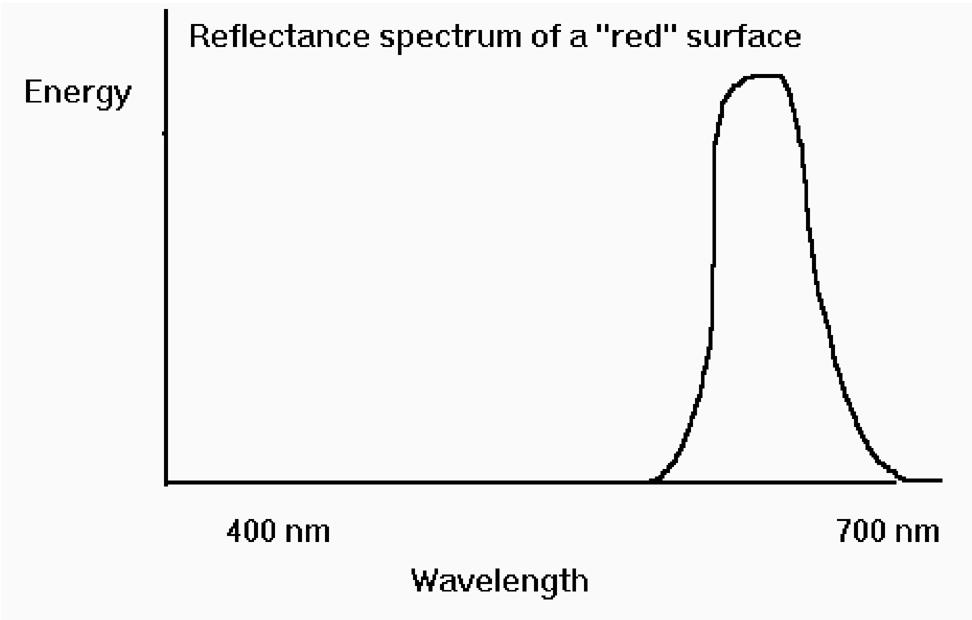

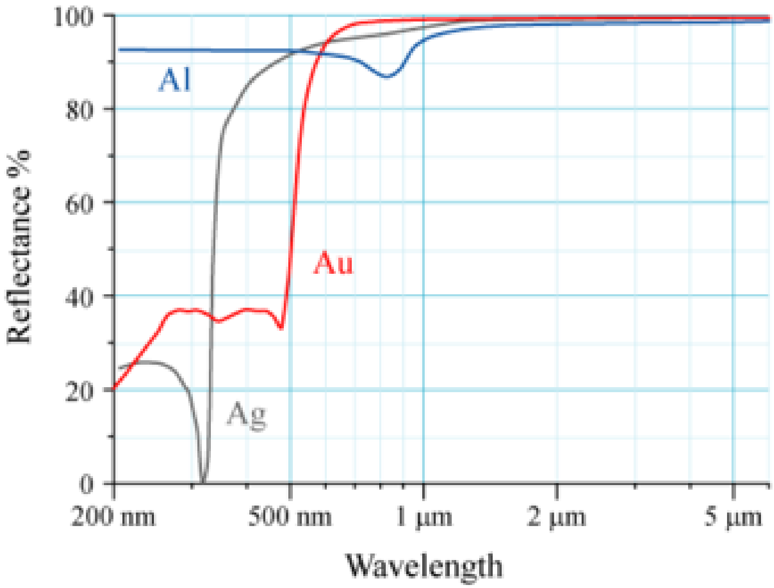

These spectra are multiplied by the spectra of the incoming light

Specular reflection/ transmission

conductor

insulator

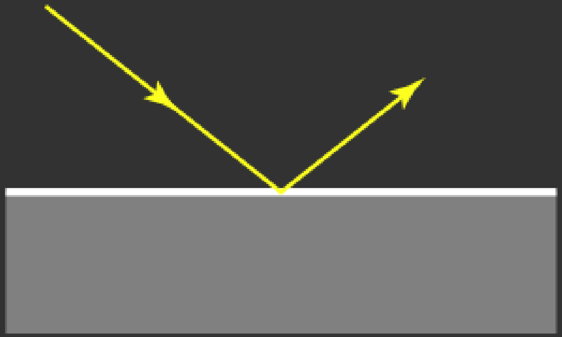

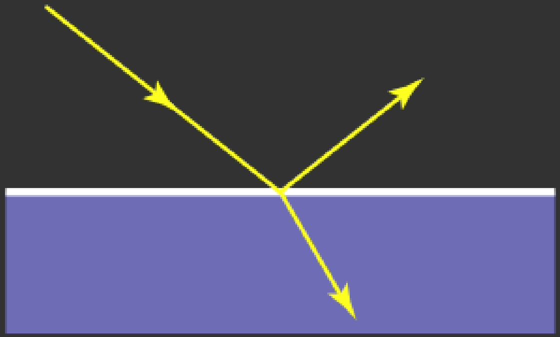

from Steve Marschner

Non-smooth-surfaced materials

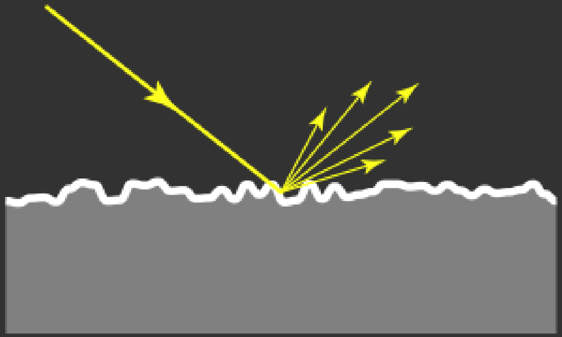

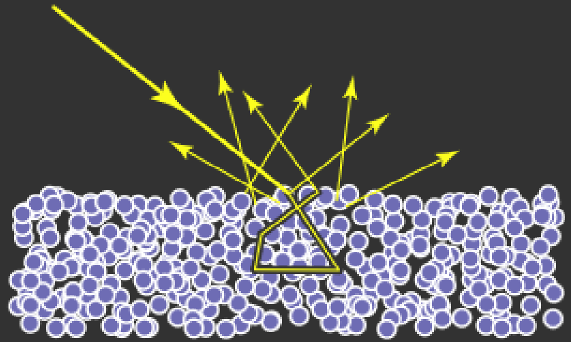

from Steve Marschner

Direct

Global

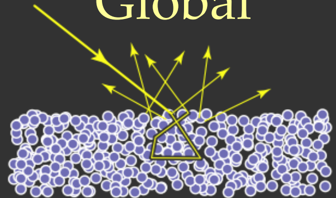

Classic reflection behavior

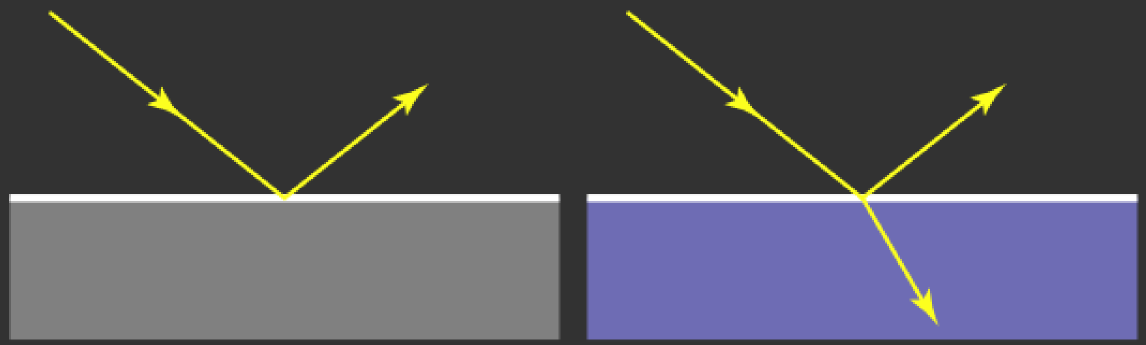

ideal specular (Fresnel)

rough specular Lambertian

from Steve Marschner

What happens when a light ray hits an object?

Some of the light gets absorbed

- converted to other forms of energy (e.g., heat)

Some gets transmitted through the object

- possibly bent, through “refraction”
- a transmitted ray could possible bounce back

Some gets reflected

- as we saw before, it could be reflected in multiple directions (possibly all directions) at once

Let’s consider the case of reflection in detail

The BRDF

The Bidirectional Reflection Distribution Function

- Given an incoming ray (θ_i, ϕ_i) and outgoing ray (θ_e, ϕ_e) what proportion of the incoming light is reflected along outgoing ray?

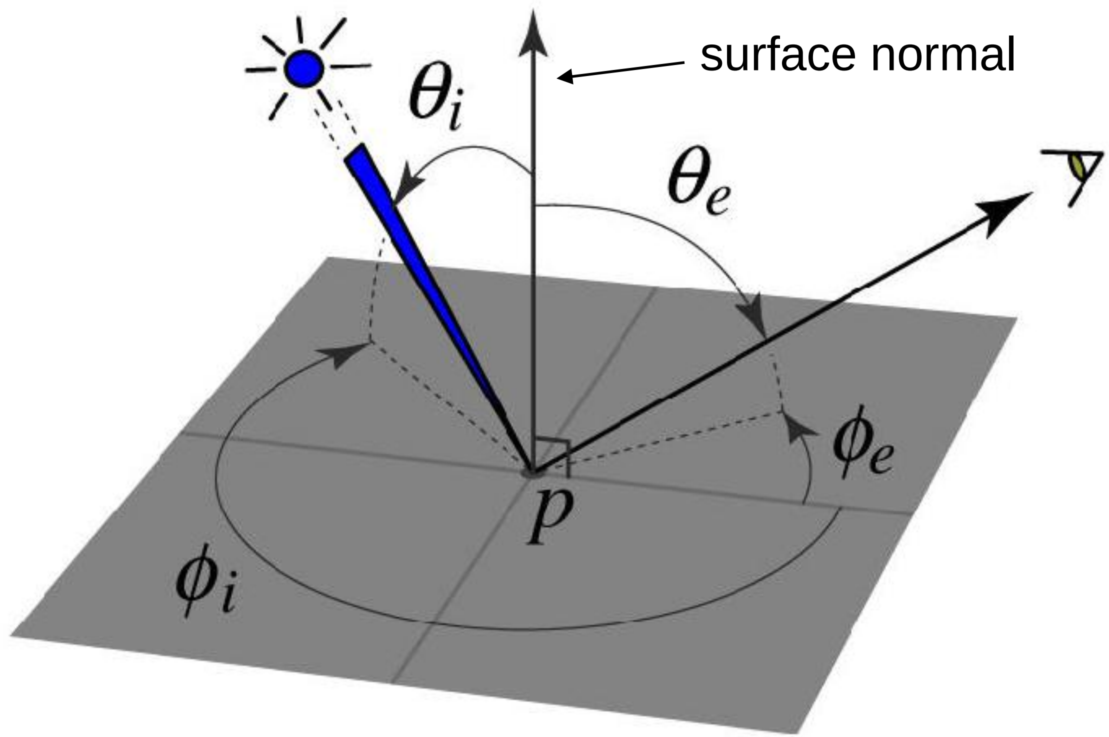

Answer given by the BRDF: $\rho(\theta_i, \phi_i, \theta_e, \phi_e)$

Constraints on the BRDF

Energy conservation

- Quantity of outgoing light ≤ quantity of incident light
  - integral of BRDF ≤ 1

Helmholtz reciprocity

- reversing the path of light produces the same reflectance

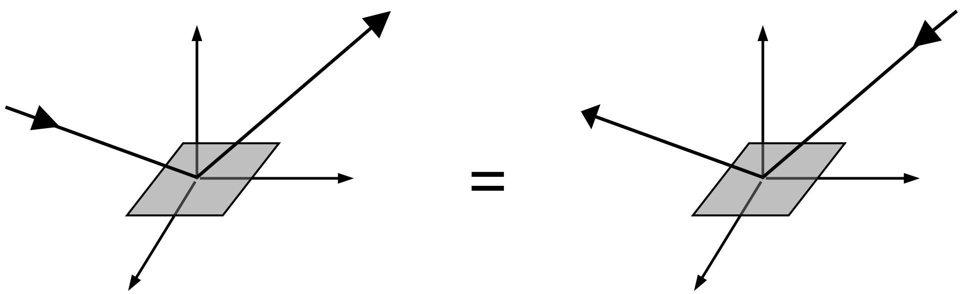

Diffuse reflection

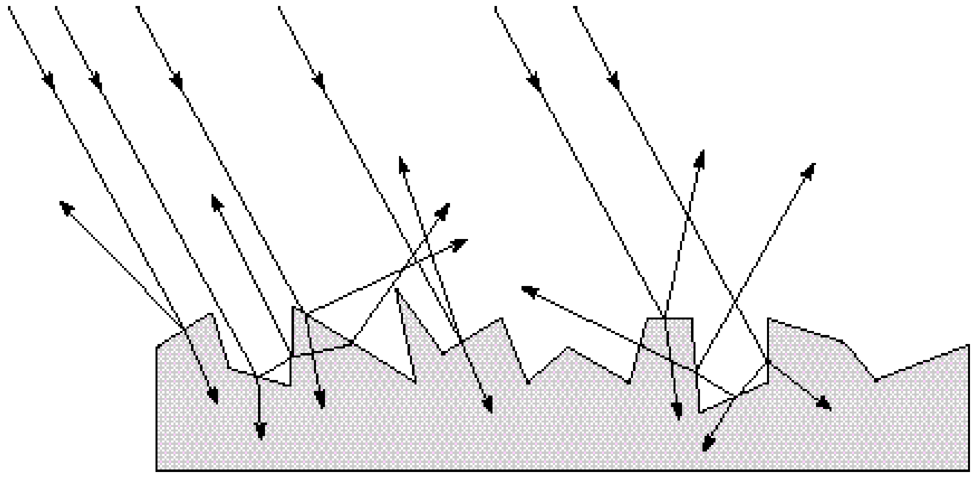

Diffuse reflection

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

Diffuse reflection

Diffuse reflection governed by **Lambert's law**

- Viewed brightness does not depend on viewing direction
- Brightness *does* depend on direction of illumination
- This is the model most often used in computer vision

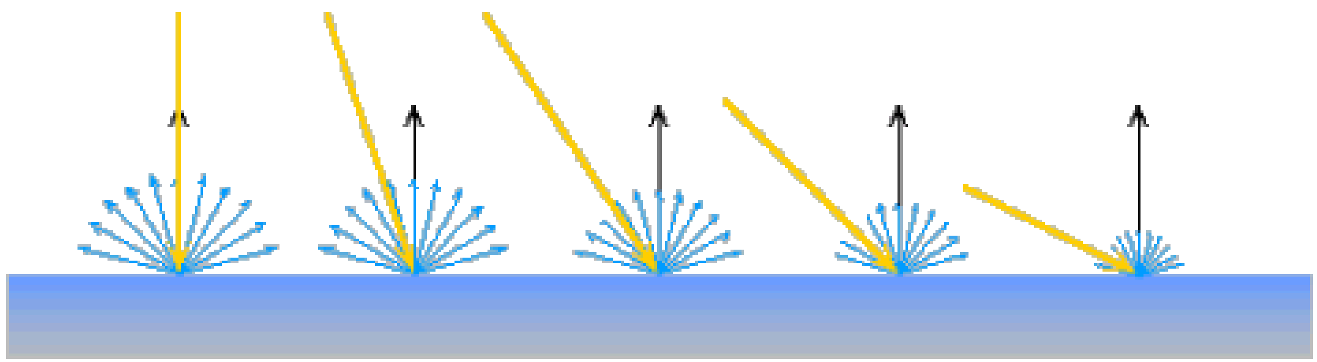

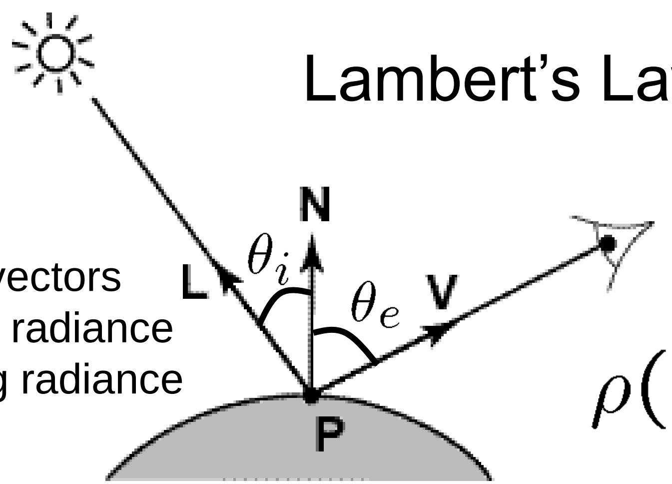

L, **N**, **V** unit vectors
I_e = outgoing radiance
I_i = incoming radiance

Lambert's Law: $I_e = k_d \mathbf{N} \cdot \mathbf{L} I_i$

k_d is called **albedo**

BRDF for **Lambertian surface**

$$\rho(\theta_i, \phi_i, \theta_e, \phi_e) = k_d cos\theta_i$$

Specular reflection

For a perfect mirror, light is reflected about **N**

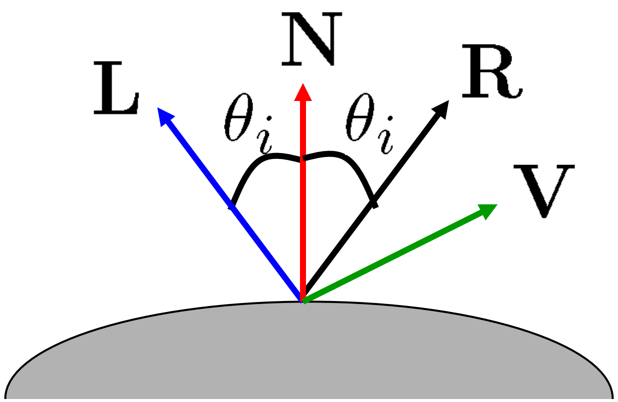

$$I_e = \begin{cases} I_i & \text{if } \mathbf{V} = \mathbf{R} \\ 0 & \text{otherwise} \end{cases}$$

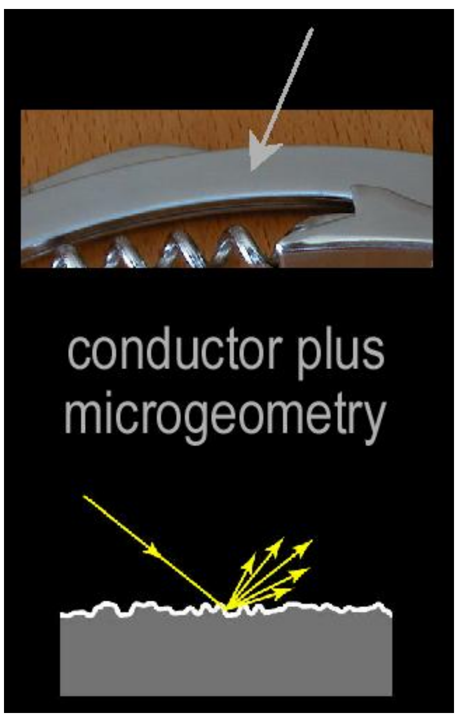

Near-perfect mirrors have a **highlight** around **R**

- common model: $I_e = k_s(\mathbf{V} \cdot \mathbf{R})^{n_s} I_i$

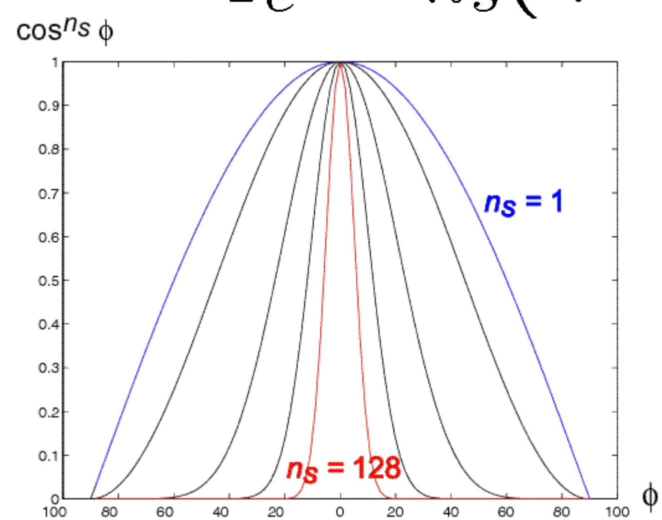

Specular reflection

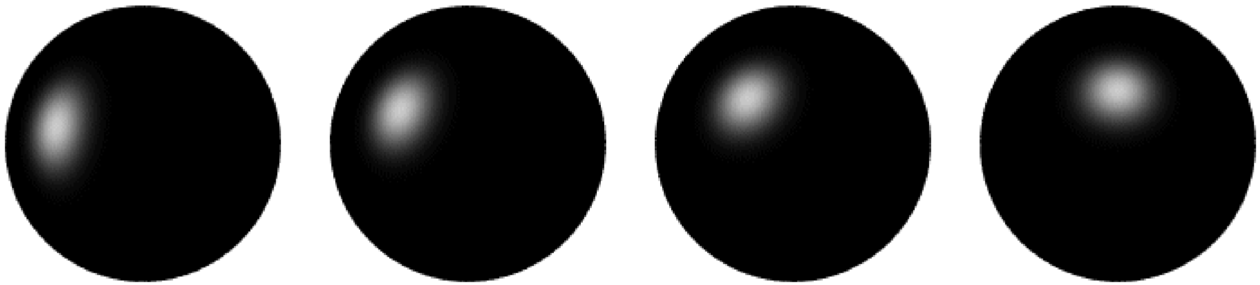

Moving the light source

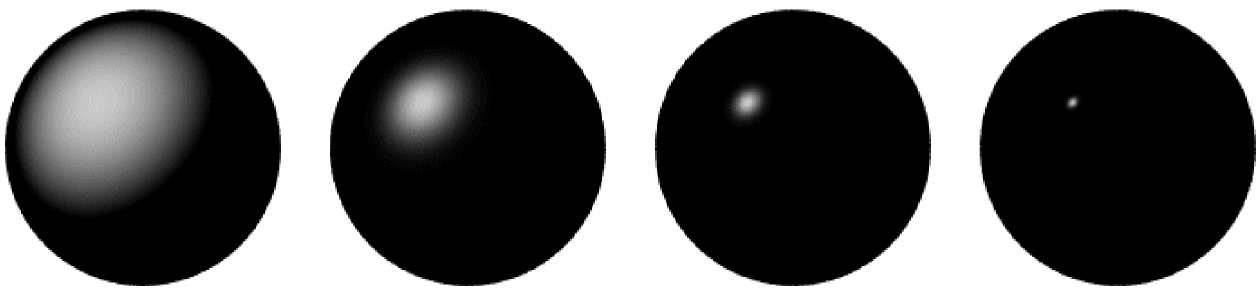

Changing n_s

Phong illumination model

Phong approximation of surface reflectance

- Assume reflectance is modeled by three components
  - Diffuse term
  - Specular term
  - Ambient term (to compensate for inter-reflected light)

$$I_e = k_a I_a + I_i \left[k_d (\mathbf{N} \cdot \mathbf{L})_+ + k_s (\mathbf{V} \cdot \mathbf{R})_+^{n_s} \right]$$

L, **N**, **V** unit vectors
I_e = outgoing radiance
I_i = incoming radiance
I_a = ambient light
k_a = ambient light reflectance factor
$(x)_+ = \max(x, 0)$

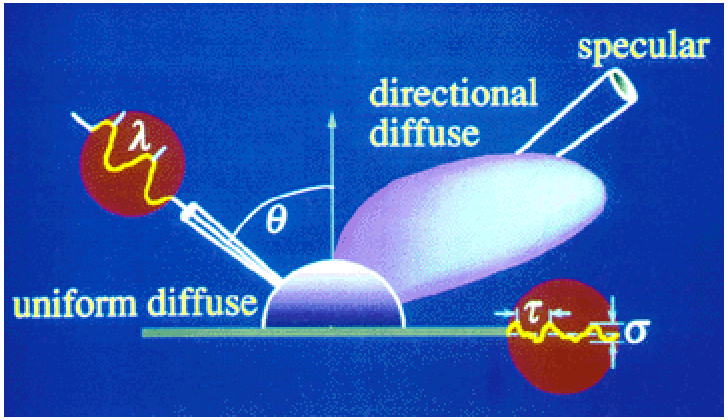

BRDF models

Phenomenological

- Phong [75]
- Ward [92]
- Lafortune et al. [97]
- Ashikhmin et al. [00]

Physical

- Cook-Torrance [81]
- Dichromatic [Shafer 85]
- He et al. [91]

Here we're listing only some well-known examples

Measuring the BRDF

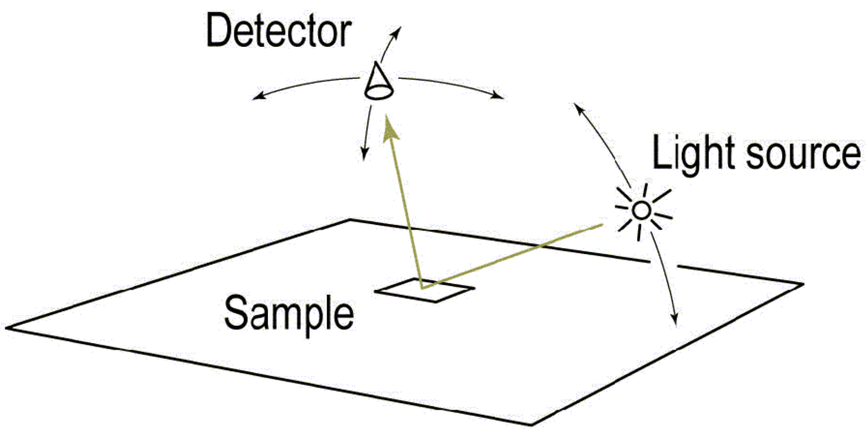

traditional

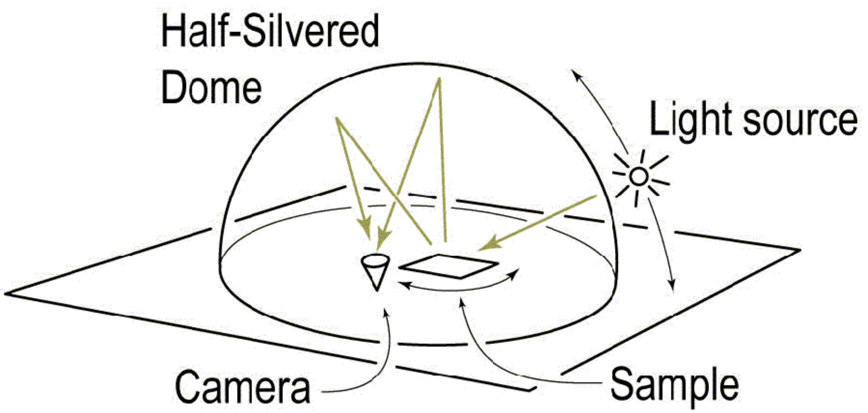

design by Greg Ward

Gonioreflectometer

- Device for capturing the BRDF by moving a camera + light source
- Need careful control of illumination, environment

BRDF databases

- MERL (Matusik et al.): 100 isotropic, 4 nonisotropic, dense

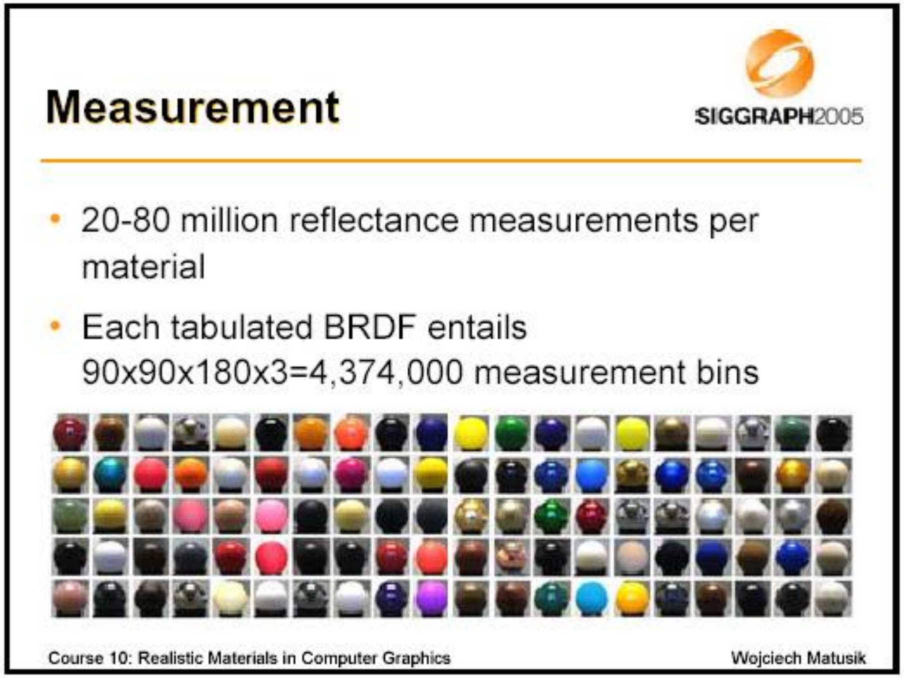

- CURET (Columbia-Utrect): 60 samples, more sparsely sampled, but also bidirectional texure functions (BTF)

Questions?

- 3-minute break

Photometric Stereo

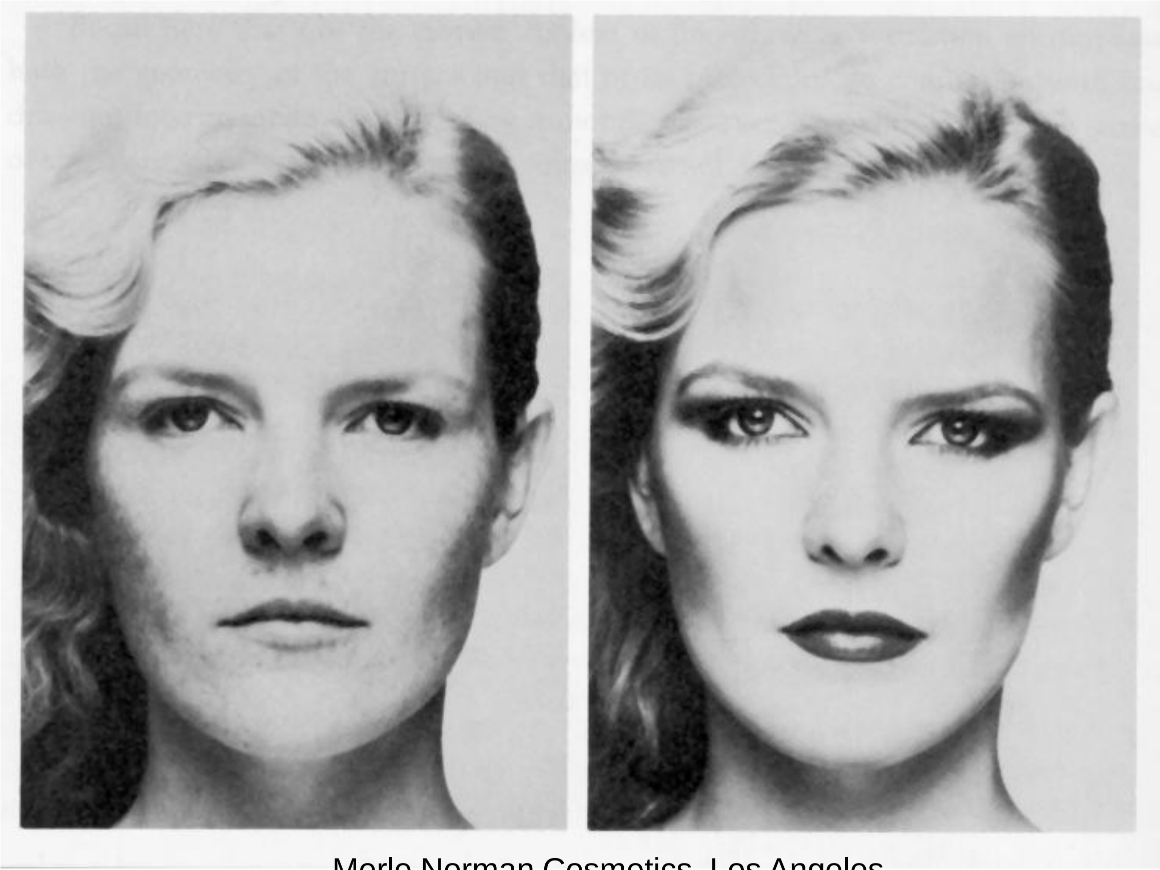

Merle Norman Cosmetics, Los Angeles

Readings

- R. Woodham, *Photometric Method for Determining Surface Orientation from Multiple Images*. Optical Engineering 19(1)139-144 (1980). (PDF)

Diffuse reflection

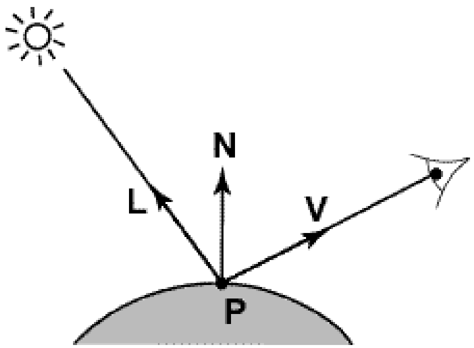

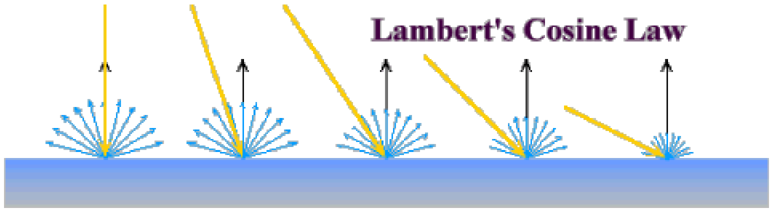

$$R_e = k_d \mathbf{N} \cdot \mathbf{L} R_i$$

image intensity of **P** → $I = k_d \mathbf{N} \cdot \mathbf{L}$

Simplifying assumptions

- $I = R_e$: camera response function f is the identity function:
  - can always achieve this in practice by solving for f and applying f^{-1} to each pixel in the image
- $R_i = 1$: light source intensity is 1
  - can achieve this by dividing each pixel in the image by R_i

Shape from shading

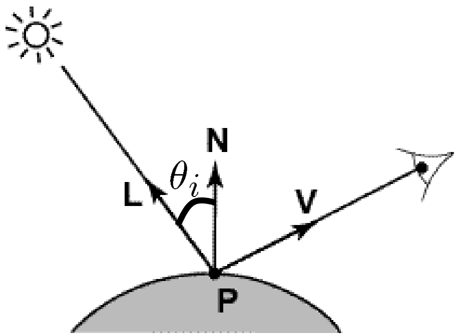

Suppose $k_d = 1$

$$\begin{aligned} I &= k_d \mathbf{N} \cdot \mathbf{L} \\ &= \mathbf{N} \cdot \mathbf{L} \\ &= cos\ \theta_i \end{aligned}$$

You can directly measure angle between normal and light source

- Not quite enough information to compute surface shape
- But can be if you add some additional info, for example
  - assume a few of the normals are known (e.g., along silhouette)
  - constraints on neighboring normals—“integrability”
  - smoothness
- Hard to get it to work well in practice
  - plus, how many real objects have constant albedo?

Photometric stereo

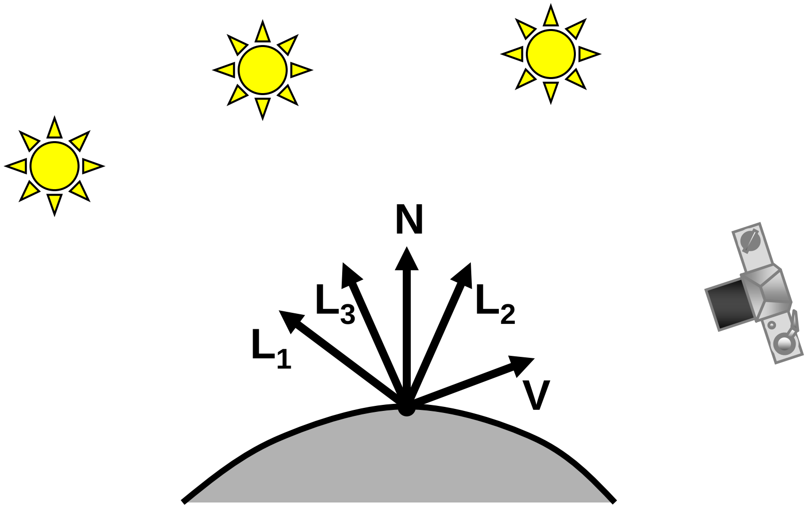

$$I_1 = k_d \mathbf{N} \cdot \mathbf{L_1}$$
$$I_2 = k_d \mathbf{N} \cdot \mathbf{L_2}$$
$$I_3 = k_d \mathbf{N} \cdot \mathbf{L_3}$$

Can write this as a matrix equation:

$$\begin{bmatrix} I_1 \\ I_2 \\ I_3 \end{bmatrix} = k_d \begin{bmatrix} \mathbf{L_1}^T \\ \mathbf{L_2}^T \\ \mathbf{L_3}^T \end{bmatrix} \mathrm{N}$$

Solving the equations

$$\underbrace{\begin{bmatrix} I_1 \\ I_2 \\ I_3 \end{bmatrix}}_{\substack{\mathbf{I} \\ 3 \times 1}} = \underbrace{\begin{bmatrix} \mathbf{L_1}^T \\ \mathbf{L_2}^T \\ \mathbf{L_3}^T \end{bmatrix}}_{\substack{\mathbf{L} \\ 3 \times 3}} \underbrace{k_d \mathbf{N}}_{\substack{\mathbf{G} \\ 3 \times 1}}$$

$$\mathbf{G} = \mathbf{L}^{-1}\mathbf{I}$$

$$k_d = \|\mathbf{G}\|$$

$$\mathbf{N} = \frac{1}{k_d}\mathbf{G}$$

More than three lights

Get better results by using more lights

$$\begin{bmatrix} I_1 \\ \vdots \\ I_n \end{bmatrix} = \begin{bmatrix} \mathbf{L_1} \\ \vdots \\ \mathbf{L_n} \end{bmatrix} k_d \mathbf{N}$$

Least squares solution:

$$\begin{aligned} \mathbf{I} &= \mathbf{LG} \\ \mathbf{L^T I} &= \mathbf{L^T L G} \\ \mathbf{G} &= (\mathbf{L^T L})^{-1} (\mathbf{L^T I}) \end{aligned}$$

Solve for N, k_d as before

What's the size of $\mathbf{L^T L}$?