Light

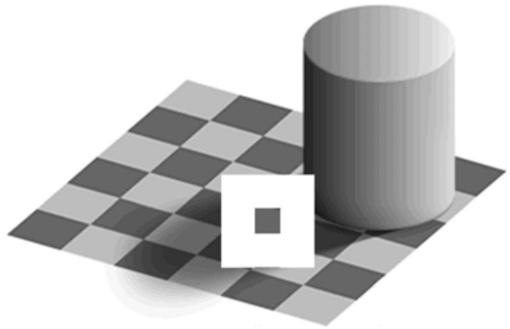

by Ted Adelson

Readings

- Szeliski, 2.2, 2.3.2
- Radiolab podcast on color

Properties of light

Today

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

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

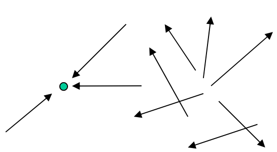

$$R(X, Y, Z, \theta, \phi, \lambda, t)$$

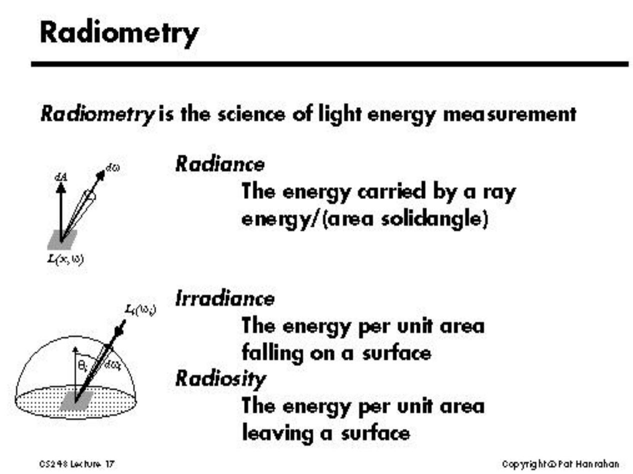

The light field

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

- Known as the **plenoptic function**
- If you know *R*, you can predict how the scene would appear from any viewpoint. How?

The **light field** $R(u, v, s, t)$ — t is *not* time (different from above t !)

- Assume radiance does not change along a ray
  – what does this assume about the world?
- Parameterize rays by intersection with two planes:

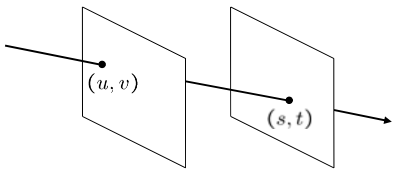

- Usually drop λ and time parameters
- How could you capture a light field?

Stanford light field gantry

Lytro

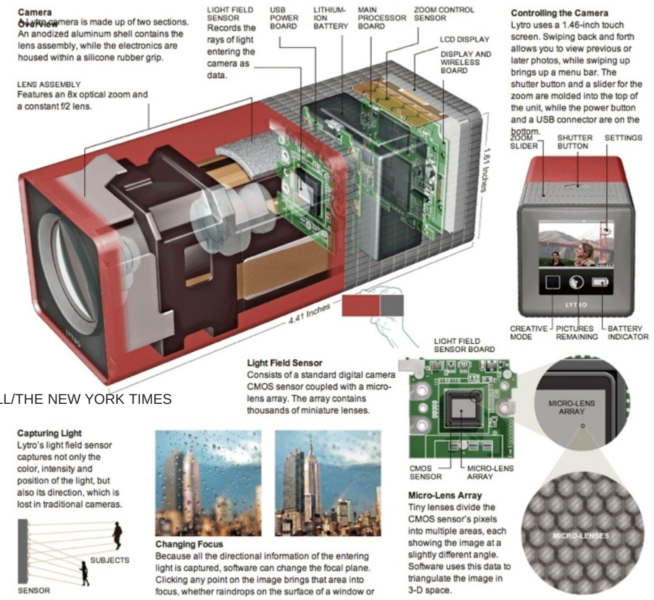

FRANK O'CONNELL/THE NEW YORK TIMES

Lightfield Camera

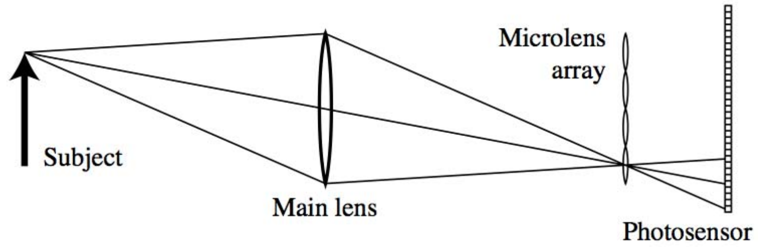

Figure 1: Conceptual schematic (not drawn to scale) of our camera, which is composed of a main lens, microlens array and a photosensor. The main lens focuses the subject onto the microlens array. The microlens array separates the converging rays into an image on the photosensor behind it.

From Ng et al., Stanford CSTR 2005-02

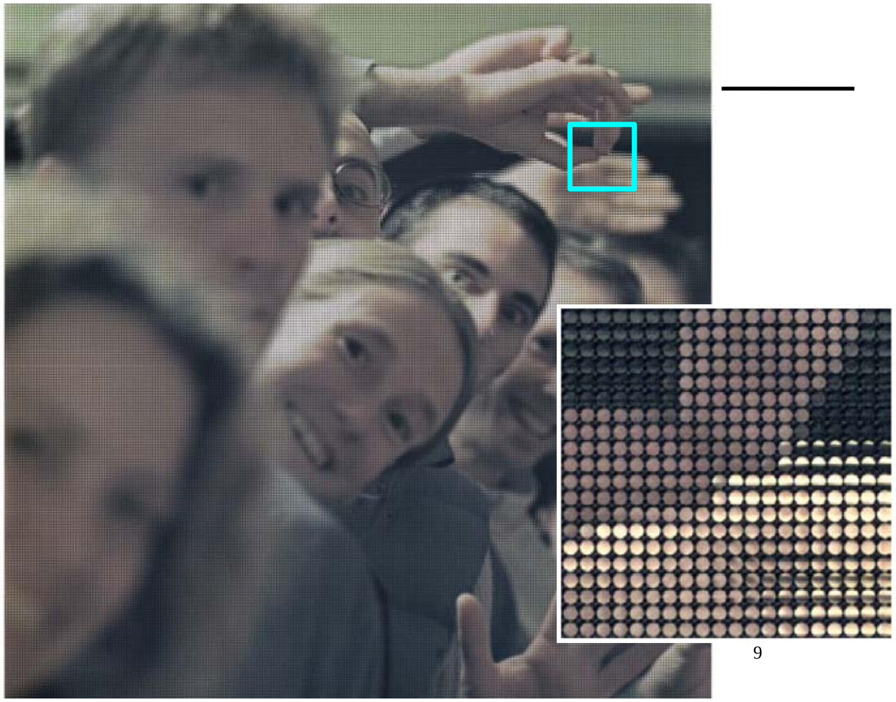

What is light?

Electromagnetic radiation (EMR) moving along rays in space

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

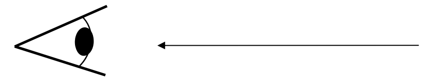

Perceiving light

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

The visible light spectrum

We “see” electromagnetic radiation in a range of wavelengths

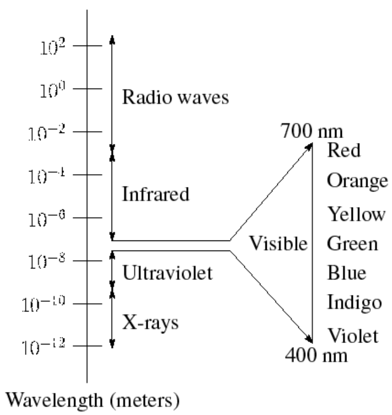

Light spectrum

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

- How much power (or energy) at each wavelength

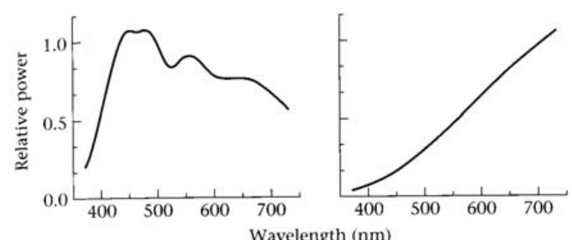

daylight tungsten bulb

Our visual system converts a light spectrum into “color”

- This is a rather complex transformation

The human visual system

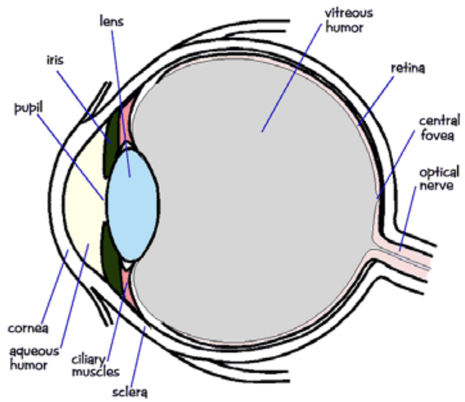

Color perception

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

Density of rods and cones

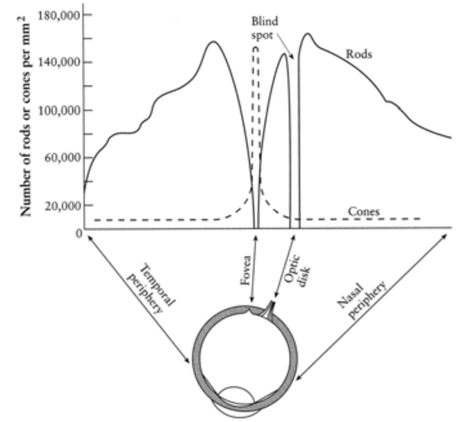

Rods and cones are *non-uniformly* distributed on the retina

- 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

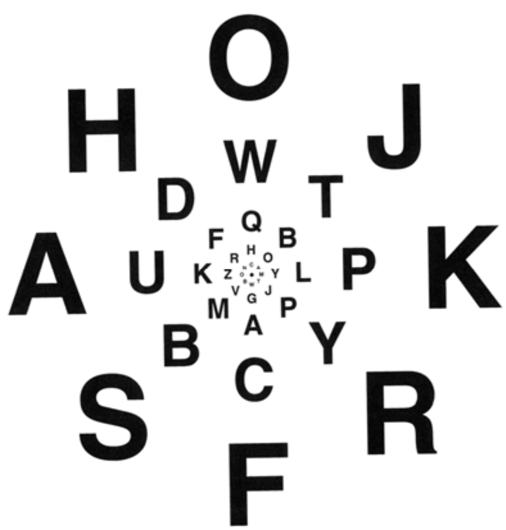

With one eye shut, at the right distance, all of these letters should appear equally legible (Glassner, 1.7).

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 seem lighter or darker depending on the surround:

- **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.

After images

Tired photoreceptors

- Send out negative response after a strong stimulus

http://www.michaelbach.de/ot/mot_adaptSpiral/index.html

Color perception

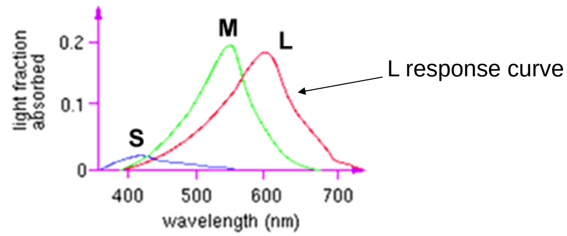

Three types of cones

- 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
  - varies from person to person (and with age)
- Colorblindness—deficiency in at least one type of cone

Color perception

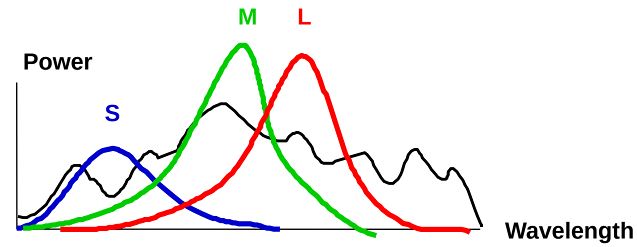

Rods and cones act as filters on the spectrum

- 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

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

Camera response function

Now how about the mapping f from radiance to pixels?

- It's also complex, but better understood
- This mapping f known as the film or camera *response function*

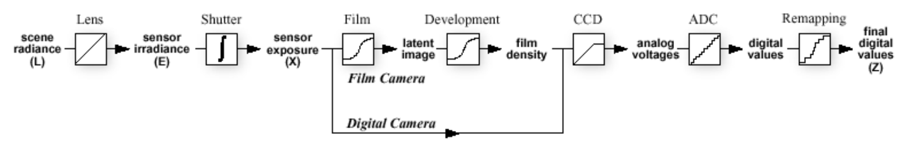

How can we recover radiance values given pixel values?

Why should we care?

- Useful if we want to estimate material properties
- Shape from shading requires radiance
- Enables creating high dynamic range images

What does the response function depend on?

$f(\text{shutter speed, aperture, film stock, digitizer, ...})$

Recovering the camera response

Method 1

- Carefully model every step in the pipeline
  - measure aperture, model film, digitizer, etc.
  - this is *really* hard to get right

Method 2

- Calibrate (estimate) the response function
  - Image several objects with known radiance
  - Measure the pixel values
  - Fit a function

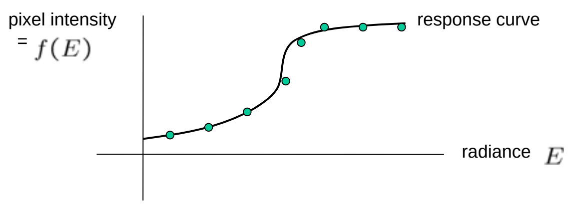

  - Find the inverse: f^{-1} maps pixel intensity to radiance

Recovering the camera response

Method 3

- Calibrate the response function from several images
  - Consider taking images with shutter speeds 1/1000, 1/100, 1/10, and 1
  - Q: What is the relationship between the radiance or pixel values in consecutive images?
  - A: 10 times as much radiance
  - Can use this to recover the camera response function

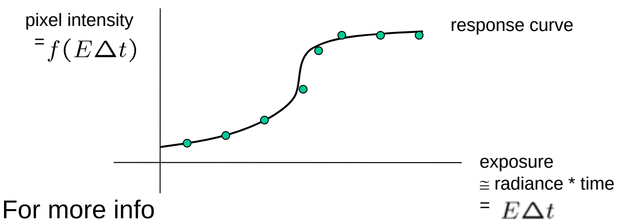

For more info

- P. E. Debevec and J. Malik. *Recovering High Dynamic Range Radiance Maps from Photographs*. In *SIGGRAPH 97*, August 1997

High dynamic range imaging

Figure 6: *Sixteen photographs of a church taken at 1-stop increments from 30 sec to $\frac{1}{1000}$ sec. The sun is directly behind the rightmost stained glass window, making it especially bright. The blue borders seen in some of the image margins are induced by the image registration process.*

Techniques

- Debevec: http://www.debevec.org/Research/HDR/
- Columbia: http://www.cs.columbia.edu/CAVE/software/rascal/rrgallery.php

Light transport

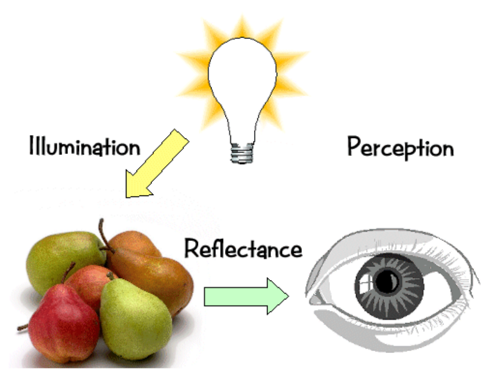

Light sources

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

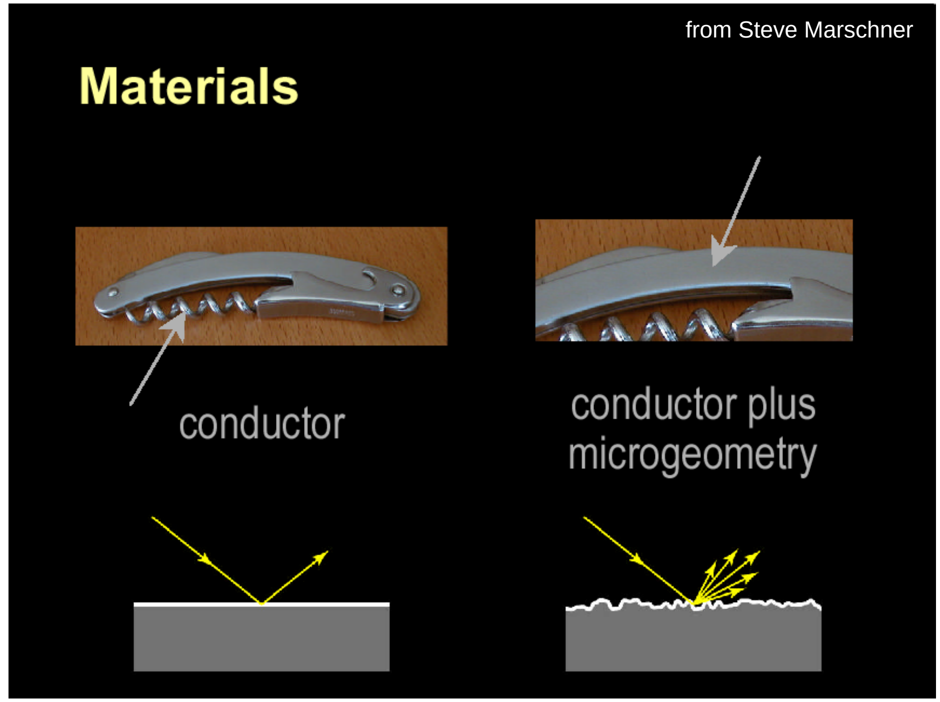

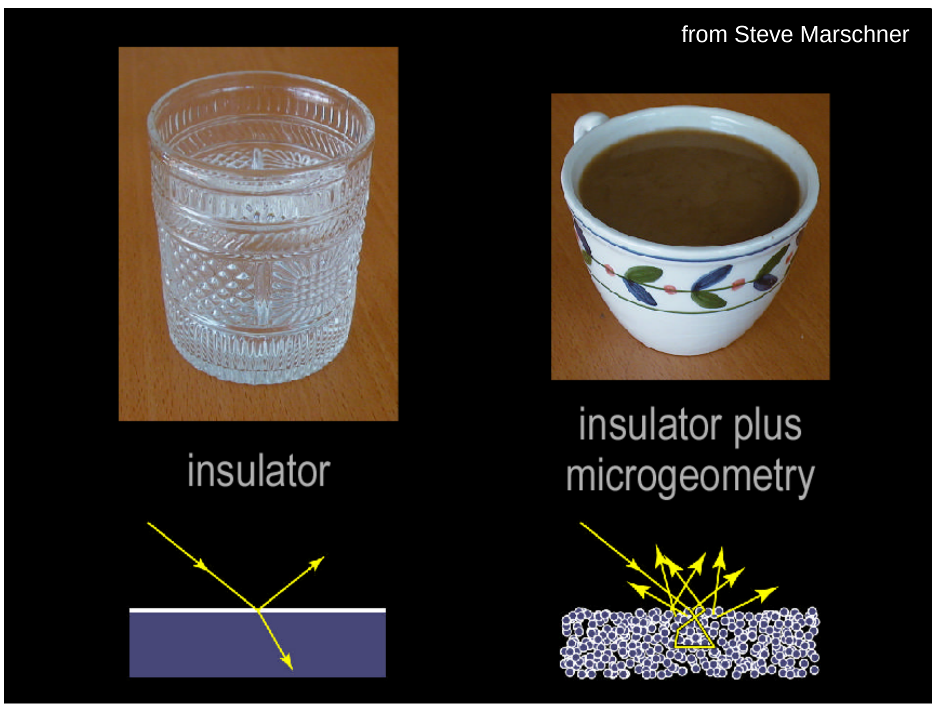

The interaction of light and matter

What happens when a light ray hits a point on 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”
- Some gets reflected
  - as we saw before, it could be reflected in multiple directions at once

Let’s consider the case of reflection in detail

- In the most general case, a single incoming ray could be reflected in all directions. How can we describe the amount of light reflected in each direction?

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?

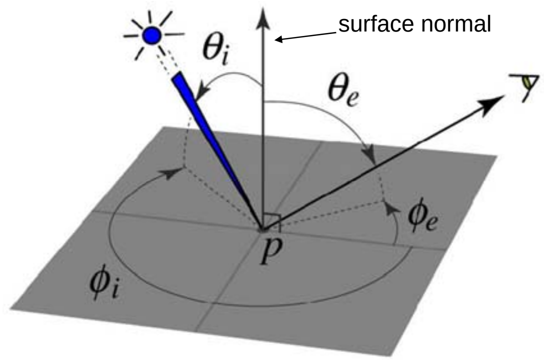

Answer given by the BRDF: $\rho(\theta_i, \phi_i, \theta_e, \phi_e)$

Diffuse reflection

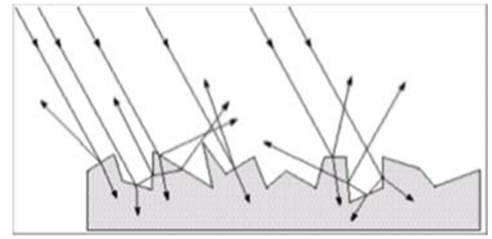

Diffuse reflection

- 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

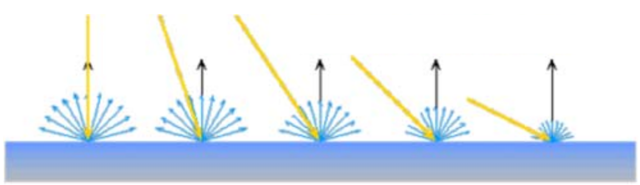

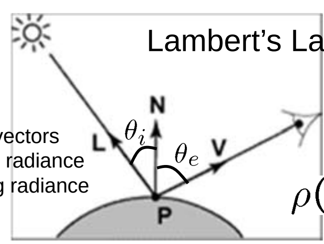

L, **N**, **V** unit vectors
I_e = outgoing radiance
I_i = incoming radiance

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

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**

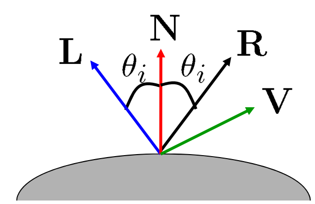

$$I_e = \begin{cases} I_i & \text{if } \mathbf{V} = \mathbf{R} \\ 0 & \text{otherwise} \end{cases}$$

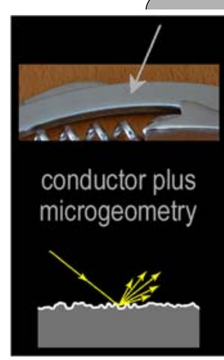

Near-perfect mirrors have a **highlight** around **R**

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

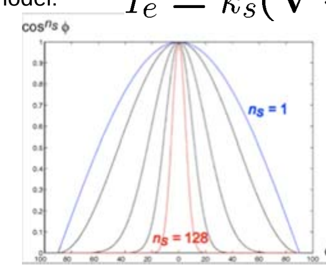

Specular reflection

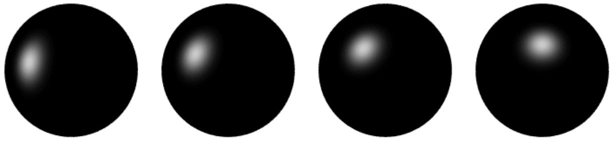

Moving the light source

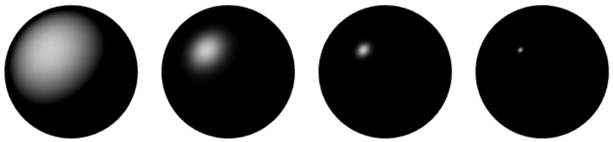

Changing n_s

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
$(x)_+ = \max(x, 0)$

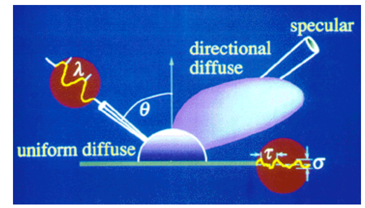

Measuring the BRDF

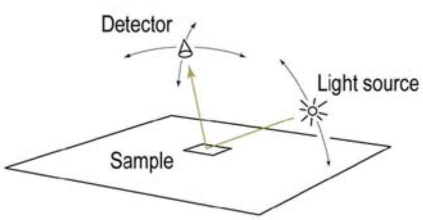

traditional

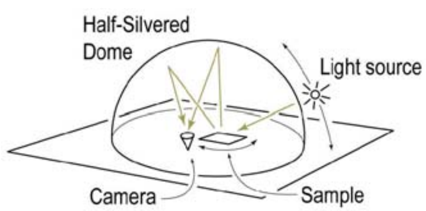

design by Greg Ward

Gonioreflectometer

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

Columbia-Utrecht Database

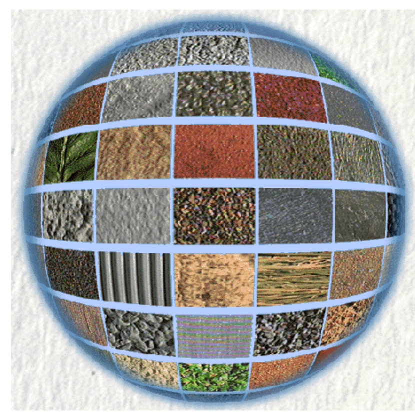

Captured BRDF models for a variety of materials

- http://www.cs.columbia.edu/CAVE/software/curet/index.php