Specular surfaces

- Another important class of surfaces is specular (mirror-like).
  - specular surfaces reflect a significant amount of energy in the specular (mirror) direction
  - produces “highlights”
- Two related cases
  - a perfect mirror
  - a fuzzy mirror
- Typically there is a diffuse (Lambertian) component as well (effects are additive)

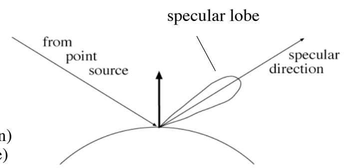

Computing reflection (specular) direction

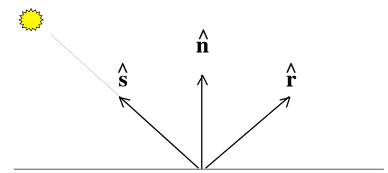

Computing reflection (specular) direction

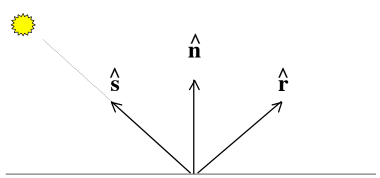

$$\hat{\mathbf{s}} + \hat{\mathbf{r}} = k\hat{\mathbf{n}} \qquad \text{and} \qquad \hat{\mathbf{n}} \bullet \hat{\mathbf{s}} = \hat{\mathbf{n}} \bullet \hat{\mathbf{r}}$$

$$\hat{\mathbf{n}} \bullet \hat{\mathbf{s}} + \hat{\mathbf{n}} \bullet \hat{\mathbf{r}} = k \;\Rightarrow\; k = \mathbf{2}\hat{\mathbf{n}} \bullet \hat{\mathbf{s}}$$

$$So \;\; \hat{\mathbf{r}} = \mathbf{2}(\hat{\mathbf{n}} \bullet \hat{\mathbf{s}})\hat{\mathbf{n}} - \hat{\mathbf{s}}$$

Phong’s model of specularities

- There are very few cases where the exact shape of the specular lobe matters.
- Typically:
  - very, very small --- mirror
  - small -- blurry mirror
  - bigger -- see only light sources as “specula
  - very big -- faint specularities
- Phong’s model
  - reflected energy falls off with

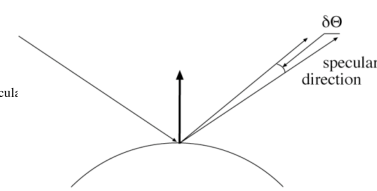

$$\cos^n(\delta\vartheta)$$

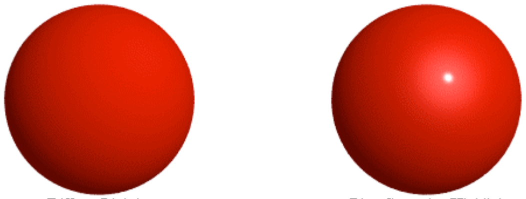

from
http://www.geocities.com/SiliconValley/Horizon/6933/shading.html

More about sources

- Exitance of a source is
  - the internally generated power radiated per unit area on the radiating surface
- A source will have both
  - radiosity, because it reflects
  - exitance, because it emits

Radiosity leaving = Exitance + Radiosity due to incoming light

More about sources

What about multiple lights?

What about real, nearby, point sources?

What about extended sources?

Standard nearby point source model (Lambertian reflection)

$$\rho_d(x)\left(\frac{N(x)\bullet S(x)}{r(x)^2}\right)$$

- N is the illuminated surface normal
- rho is diffuse albedo
- S is source vector - a vector from x to the source, whose length is the intensity term
  - this works because a dot-product is basically a cosine
- r(x) is distance from surface point to source --- term occurs because source “looks smaller” as we move away--or, alternatively, its energy is spread out over a larger surface.

Standard distant point source model

- If the source is far away, this formula reduces to the same as the Lambertian formula from before:

then $r(x) \cong R$

and $S(x) \cong S$

if we let $S_d = \left(S / R^2 \right)$

we get $\rho_d(x)\left(N(x) \bullet S_d\right)$

Line sources

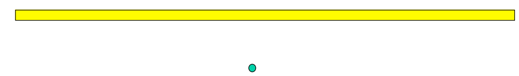

Radiosity due to line source varies with inverse distance,
if the source is long enough (derivation is through integration of the contributions along the line)

General extended sources

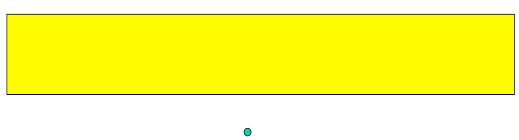

Can be handled by doing the integration (we won't)

What if the source is large relative to the distance to it?

How about the hemisphere of the sky?

Shadows cast by a point source

- A point that can't see the source is in shadow
- For point sources, the geometry is simple

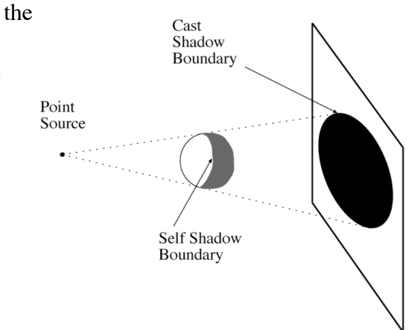

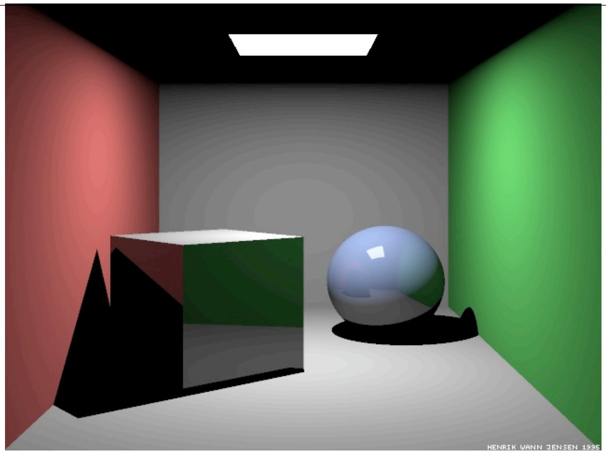

Ray-traced Cornell box, due to Henrik Jensen, http://www.gk.dtu.dk/~hwj

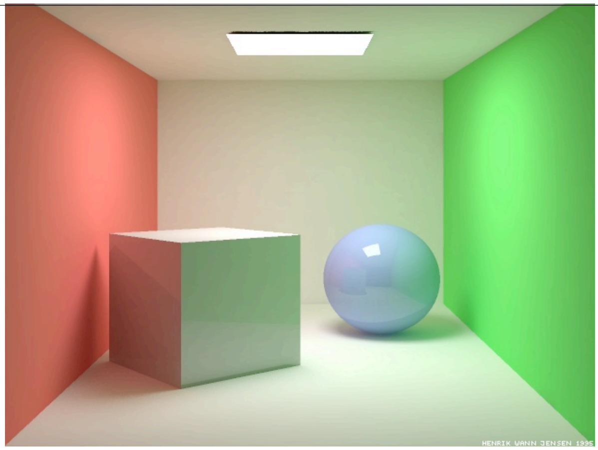

Radiosity Cornell box, due to Henrik Jensen, http://www.gk.dtu.dk/~hwj, rendered with ray tracer

Flat shading

- Compute shading value inside polygon using interpolate
- Flat shading
  - Use polygon normal for Lambertian shading
  - Advantages:
    - fast -- one shading value per polygon
  - Disadvantages:
    - inaccurate -- looks blocky

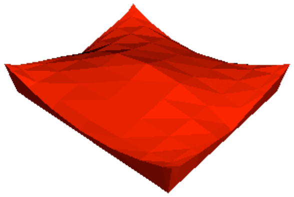

Figure from http://freespace.virgin.net/hugo.elias/graphics/x_polygo.htm

Gouraud (Interpolated) Shading

- Use normal at each vertex of polygon
  - known either from the "mesh" construction, OR, average normals where polygons meet.
- Shade these, and linearly interpolate
- Advantages
  - fast (interpolation can even be put into scan line algorithm)
  - much smoother
- Disadvantages:
  - specularities get lost

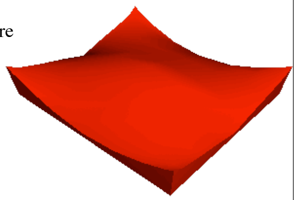

Phong Shading

- Interpolate normals instead of pixel values
  - Shade using normal estimate for each point
  - Advantage
    - high quality, narrow specularities
  - Disadvantage
    - more expensive than Gouraud

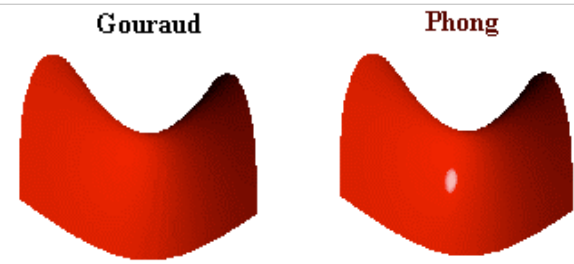

from
http://www.geocities.com/SiliconValley/Horizon/6933/shading.html

What about the color of the light?

So far, we have not dealt with the color of the light--the implicit assumption being that it is "white" and thus does not change the color of anything.

This is often not the case!

Naïve (but common) model

Consider the color of the light to be specified by its (R,G,B)--technically the color of a perfect uniform reflector (white surface).

Similarly, now specify the albedo as a triple--one for each channel. The color of a Lambertian surface is then:

$$(R,G,B) = (\rho_R S_R,\ \ \rho_G S_G,\ \ \rho_B S_B)(\mathbf{n} \bullet \mathbf{s})$$

Naïve (but common) model

Naïve because we assume that the red part of the light does not interact with green or blue albedos, etc.

(Referred to as the diagonal model)

What about specular surfaces?

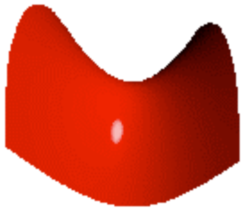

Specular surfaces

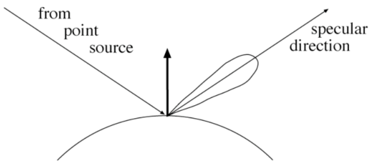

- Important point: The specular part of the reflected light usually carries the color of the **light**
- Technically, this is the case for dielectrics--plastics, paints, glass.
- Important exception is metals (e.g. gold, copper)

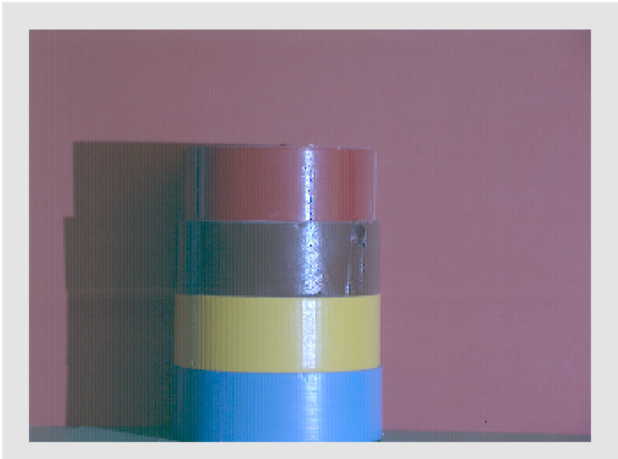

Example: Dielectrics

- Examples: Paints, plastics
- Reasonably well approximated by a specular part and a Lambertian body part.

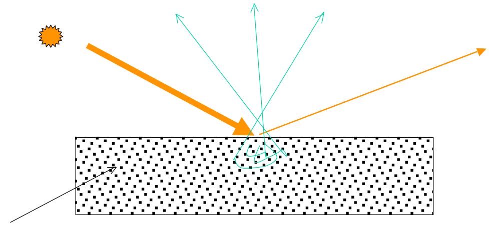

Non conductive matrix with scattering particles of the order of the wavelength of light---note: the same general process explains why the sky is blue.

The colors of the rainbow

- Light is electromagnetic radiation, occurring at different wavelengths (or photon energies)
- The radiation around us is a mix of these
- Visible portion is about 400 to 700 nm
- Certain applications may require modeling some UV also.
- Light is specified by its spectrum recording how much power is at each wavelength.

Sunlight

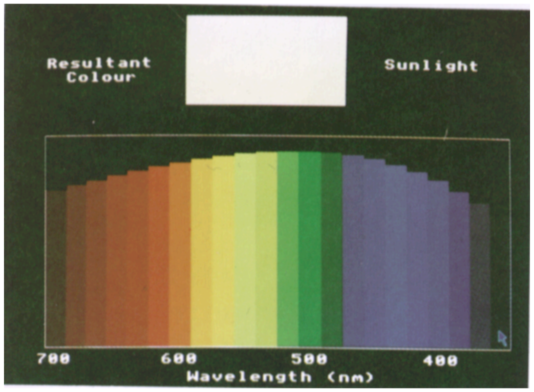

Two disparate source spectra

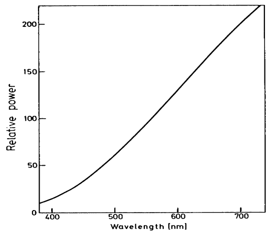

Fig. 4.1. Wavelength composition of light from a tungsten-filament lamp [typified by CIE ILL A (Sect. 4.6)]. Relative spectral power distribution curve. Color temperature: 2856 K

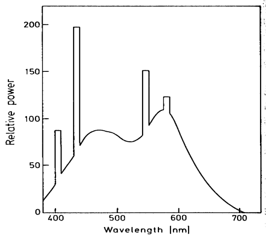

Fig. 4.2. Wavelength composition of light from a daylight fluorescent lamp. Typical relative spectral power distribution curve. Correlated color temperature: 6000 K. (Based on data of Jerome reported in [Ref. 3.14, p. 37])

Energy spectra of 20 other common lights

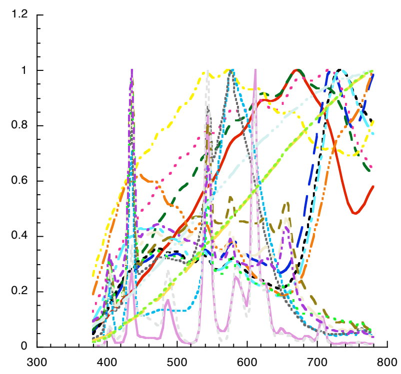

Absorption spectra: real pigments

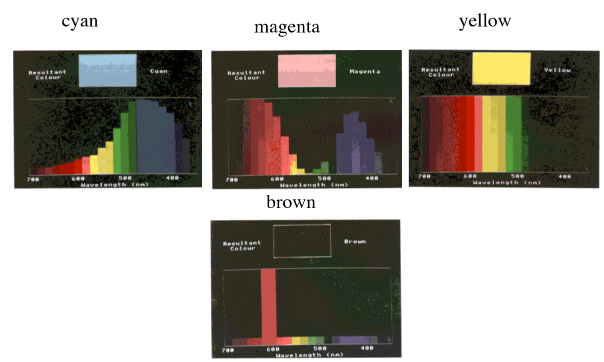

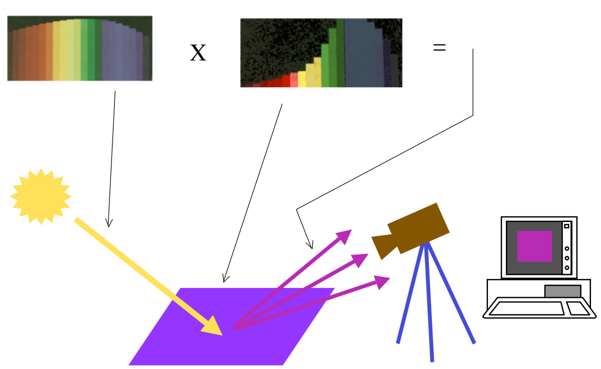

Sensors

Sensors (including those in your eyes) have a varied sensitivity over wavelength

Different variations lead to different kinds of sensor responses (“colors” in a naïve sense)

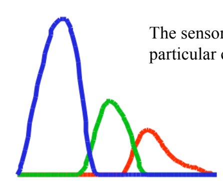

The sensor sensitivity spectra for a particular color camera (they vary a lot).

<--- wavelength --->

Image Formation (Spectral)

(**R**,**G**,**B**) = $\int_{380}^{780}$  * 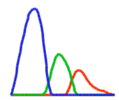 dλ

More formally,

The response of an image capture system to a light signal $L(\lambda)$ associated with a given pixels is modeled by

$$\rho^{(k)} = \int L(\lambda) R^{(k)}(\lambda) d\lambda$$

where $R^{(k)}(\lambda)$ is the sensor response function for the k^{th} channel.

Note the usual case of three channels

$$(R, G, B) = \left(\rho^{(1)}, \rho^{(2)}, \rho^{(3)}\right)$$

Discrete Version

Often we represent functions by vectors. For example, a spectra might be represented by 101 samples in the range of 380 to 780 nm in steps of 4nm.

Then $L(\lambda)$ becomes the vector **L**, $R^{(k)}(\lambda)$ becomes the vector $\mathbf{R^k}$, and the response is given by a dot product:

$$\rho^{(k)} = \mathbf{L} \bullet \mathbf{R^{(k)}}$$

Sensor/light interaction example

R=(0,0,1,3,7,4,2,0,0,0)

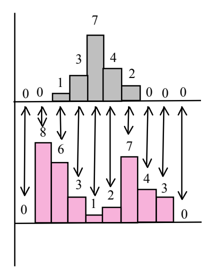

L=(0,8,6,3,1,2,7,4,3,0)

Multiply lined up pairs of numbers and then sum up

Sensor/light interaction example

R=(0,0,1,3,7,4,2,0,0,0)

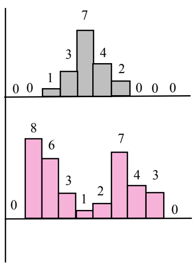

L=(0,8,6,3,1,2,7,4,3,0)

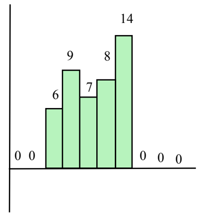

L .* R=
(0*0, 0*8, 1*6, 3*3, 7*1, 4*2, 2*7, 0*4, 0*3, 0*0)
= (0, 0, 6, 9, 7, 8, 14, 0, 0, 0)

L • P = 0 + 0 + 6 + 9 + 7 + 8 + 14
= 44

Image Formation (Spectral)

- Note that by this model, light capture is linear.
- Formally this means

?

Image Formation (Spectral)

- Note that image formation loses spectral information
- This means that two quite different spectra can map into the same color

One tricky bit

Electronic capture (e.g. “CCD”) is linear, but typically the circuitry will put the sensor responses through a non-linear mapping (e.g. approximate square root).

This is because display is usually either non-linear due to physics (CRT) or by design (to be like a CRT). This is better because there is less relative noise where humans will notice it.

(A bit more on this later).

Causes of color

- The sensation of color is caused by the brain.
- One way to get it is through a **response** of the eye to the presence/absence of light at various wavelengths.
- Dreaming, hallucination, etc.
- Pressure on the eyelids

Trichromaticity

Empirical fact--colors can be approximately described/matched by three quantities (assuming normal color vision).

Need to reconcile this observation with the spectral characterization of light

Color receptors

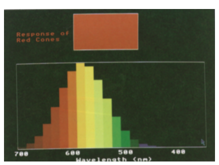

"Long" cone

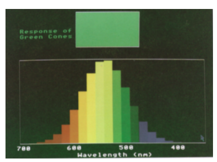

"Medium" cone

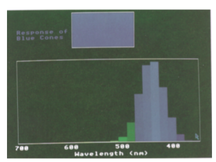

"Short" cone

Some understanding results from an analogy with camera sensors

Directly determining the camera like sensitivity response is hard!

Colour Reproduction

Motivates specifying color numerically (there are other reasons to do this also)

General (man in the street) observation--color reproduction *sort of* works.

Specifying Colour

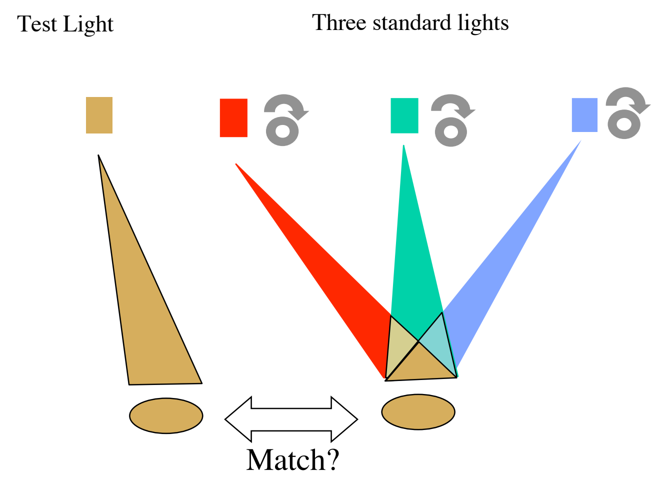

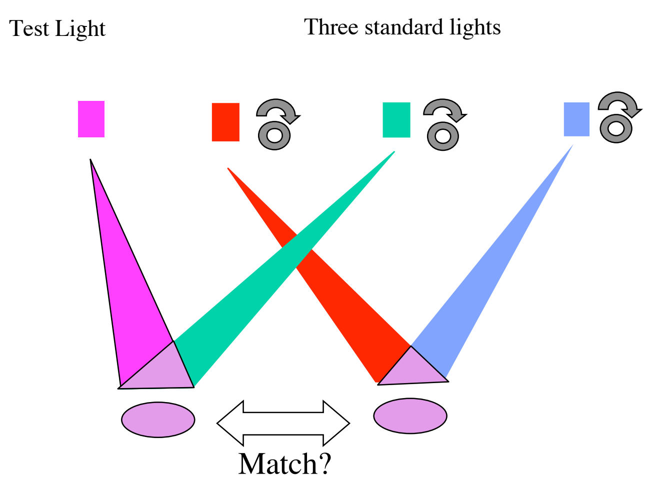

Trichromacy

Experimental fact about people (with “normal” colour vision)---matching works (for reasonable lights), provided that we are sometimes allowed negative values.

Our “knob” positions correspond to (X,Y,Z) in the standard colorimetry system.

Technical detail: (X,Y,Z) are actually arranged to be **positive** by a linear transformation, but these “knob” positions **cannot** correspond to any **physical** light.

Specifying Colour

(50,150,75)

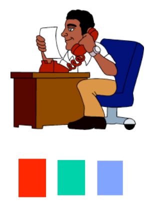

(50,150,75)

Specifying Colour

We don't want to do a matching experiment every time we want to use a new color!

Grassman's Contribution

Colour matching is linear

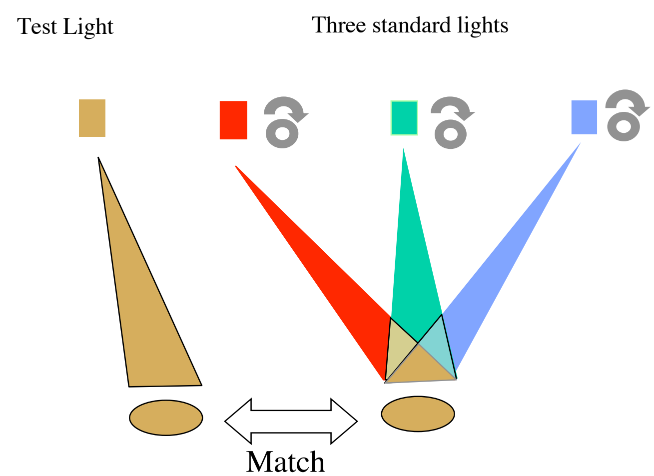

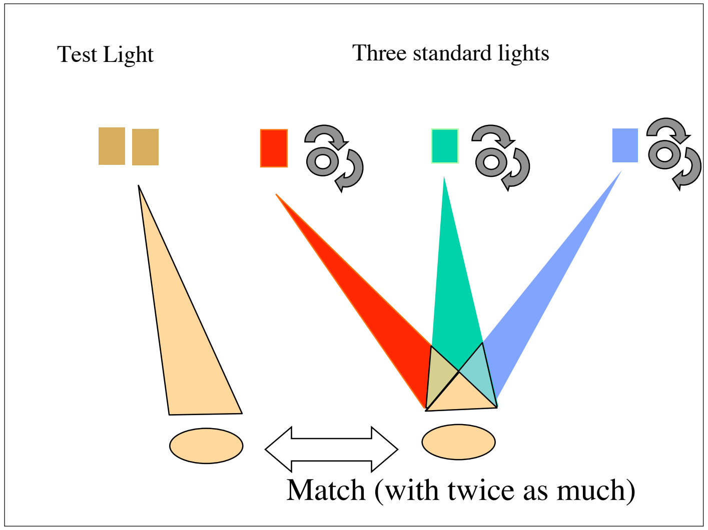

Matching is Linear (Part 1)

C1 is matched with (X1,Y1,Z1)

C = a*C1

C is matched with a * (X1, Y1, Z1)

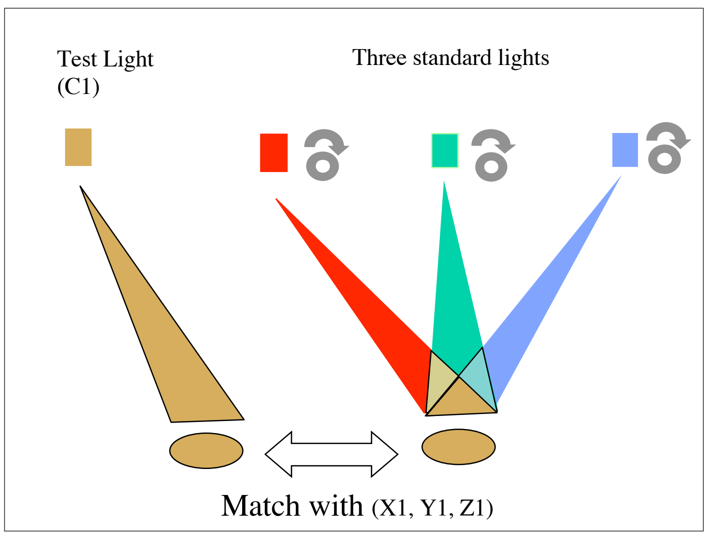

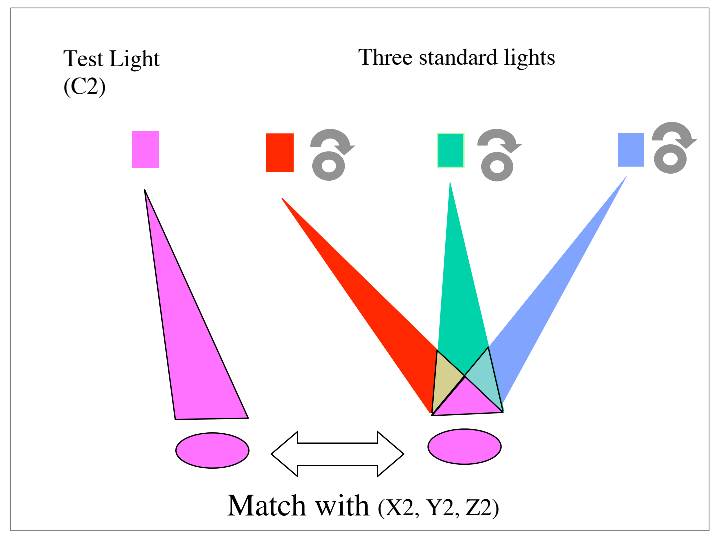

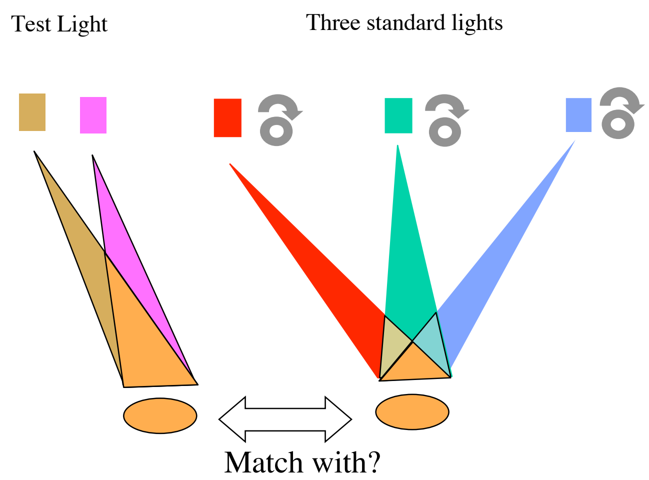

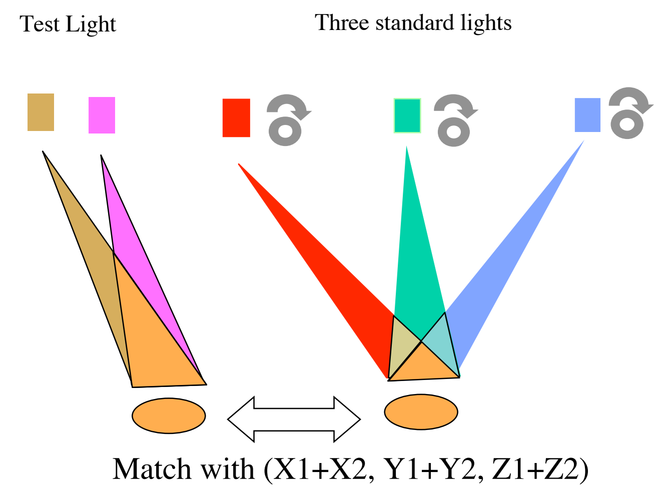

Matching is Linear (formal)

C = a*C1 + b*C2

C1 is matched with (X1,Y1,Z1)
C2 is matched with (X2,Y2,Z2)

C is matched by
a*(X1,Y1,Z1) + b*(X2,Y2,Z2)

Specifying Color

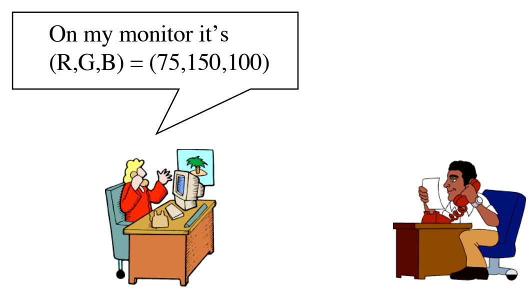

Specifying Colour

But what is (R,G,B)?

Specifying Colour

R matches (X_r, Y_r, Z_r)

G matches (X_g, Y_g, Z_g)

B matches (X_b, Y_b, Z_b)

Specifying Colour

Then by
(R,G,B)=(75,150,100)
you mean (X,Y,Z),
where …..

$$X = 75^* X_r + 150^* X_g + 100^* X_b$$

$$Y = 75^* Y_r + 150^*Y_g + 100^* Y_b$$

$$Z = 75^* Z_r + 150^*Z_g + 100^* Z_b$$

(No need to match--just compute!)

Specifying Colour

… , now that we have specified the colour, I can print it!

$$\begin{vmatrix} X \\ Y \\ Z \end{vmatrix} = \begin{vmatrix} X_r & X_g & X_b \\ Y_r & Y_g & Y_b \\ Z_r & Z_g & Z_b \end{vmatrix} \begin{vmatrix} 75 \\ 100 \\ 150 \end{vmatrix}$$

$$\begin{vmatrix} X \\ Y \\ Z \end{vmatrix} = \begin{vmatrix} X_r & X_g & X_b \\ Y_r & Y_g & Y_b \\ Z_r & Z_g & Z_b \end{vmatrix} \begin{vmatrix} R \\ G \\ B \end{vmatrix}$$

$$\begin{vmatrix} X \\ Y \\ Z \end{vmatrix} = M \begin{vmatrix} R \\ G \\ B \end{vmatrix}$$

Colour Reproduction
(Monitors & Projectors)

$$\to \begin{vmatrix} X \\ Y \\ Z \end{vmatrix}_{apple}$$

Find (R,G,B)

$$\begin{vmatrix} X \\ Y \\ Z \end{vmatrix}_{apple} = M \begin{vmatrix} R \\ G \\ B \end{vmatrix}_{apple}$$

$$\begin{vmatrix} R \\ G \\ B \end{vmatrix}_{apple} = M^{-1} \begin{vmatrix} X \\ Y \\ Z \end{vmatrix}_{apple}$$

XYZ color space

XYZ color space is a linear transformation of the matches to standard lights.

The transformation is used to ensure that all color co-ordinates are positive

This means that XYZ corresponds to matches of fictitious (physically impossible) lights.

The gamut of all colors

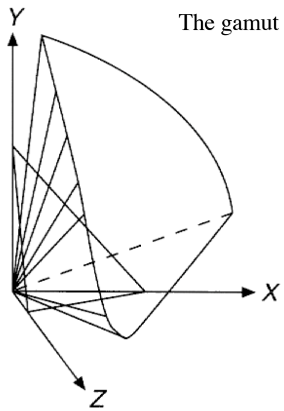

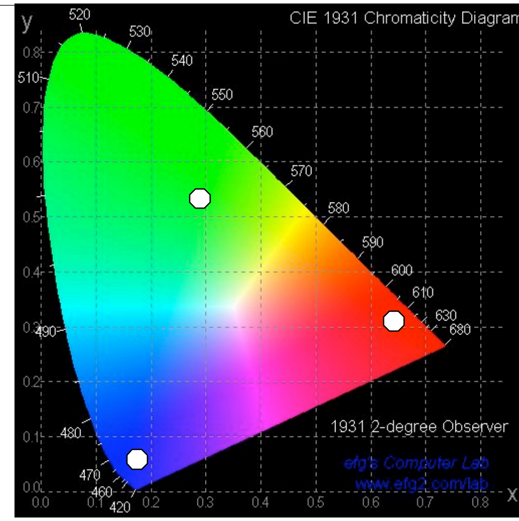

Available from efg2.com

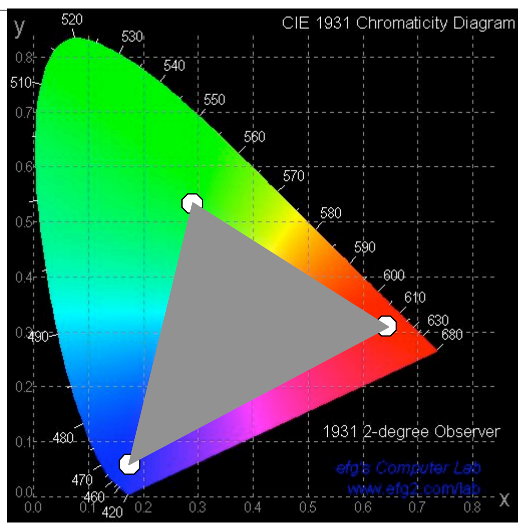

Available from efg2.com

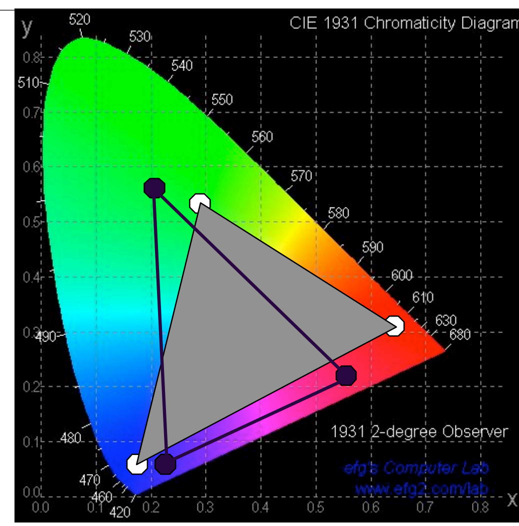

Available from efg2.com

Qualitative features of CIE x, y

- Linearity implies that colors obtainable by mixing lights with colors A, B lie on line segment with endpoints at A and B
- Monochromatic colours (spectral colors) run along the “Spectral Locus”

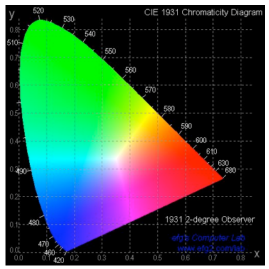

Qualitative features of CIE x, y

- Why the funny shape?

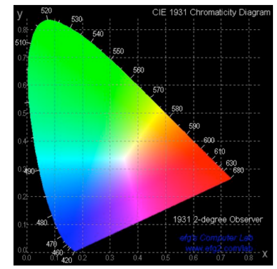

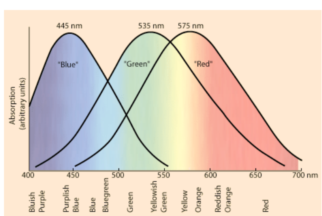

One measurement of human cone absorption

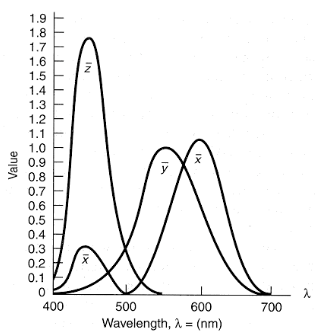

XYZ response curves

Matching is only for “aperture” color

- When color is viewed in the context of other colors numerous effects occur which complicate the characterization of color (simultaneous contrast, color constancy, etc)
- Other complications include chromatic aberration in the eye and different spatial resolution for different colors (these are linked)

Colour Reproduction

Key point--color reproduction is based on “metamerism”

Metameric match--colors which match, despite different spectra.

Duplicating spectra would work, but for practical reasons, we duplicate the match.

For reflective surfaces, e.g prints, this means that the match can change if the illumination changes.

The quest for uniform colour spaces

- Definition of uniform: equal (small!) steps give the same perceived color changes.
- XYZ is not uniform!
- Uniformity only applies to small differences. There is no theory for numerically deciding if yellow is perceptually closer to green or red.

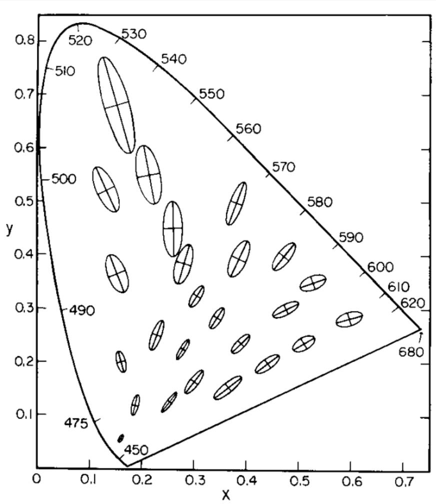

MacAdam Ellipses (scaled by a factor of 10) on CIE x, y

Mixing pigments in CIE

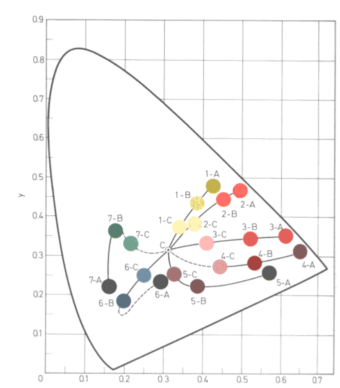

Color matching is linear, but combining pigments is not necessarily linear like mixing light .

Shading values for colored surfaces

- Simplest:
  - Use appropriate shading model in 3 channels, instead of one
  - Implies red albedo, green albedo, blue albedo, etc.
  - Works because the shading model is independent of wavelength.
  - Can lead to somewhat inaccurate colour reproduction in some cases - particularly coloured light on coloured surfaces
- Better
  - Use appropriate shading model at many different wavelength samples - 7 is usually enough
  - Estimate receptor response in eye using sum over wavelength
  - Set up pixel value to generate that receptor response

Monitor Gamma

A typical image encoding is **NOT** linear. Often a gamma correction is included. This leads to no end of confusion.

A "gamma" corrected image is ready to drive a CRT monitor, and has advantages that quantization (8 bits) errors are *roughly* uniformly distributed--that fact that this works is a convenient accident.

Monitor Gamma

Due to the physics involved, CRT monitor brightness is proportional to voltage^(2.5)

This is further hacked to give the "standard" gamma of 2.2

So, if an image looks good on a CRT, it is likely to be non-linear by pow(1/2.2)

LCD--more linear, but then hardware/software can be hacked to be like CRT

Confusing? Yes!

Recursive ray tracing

H&B, page 597

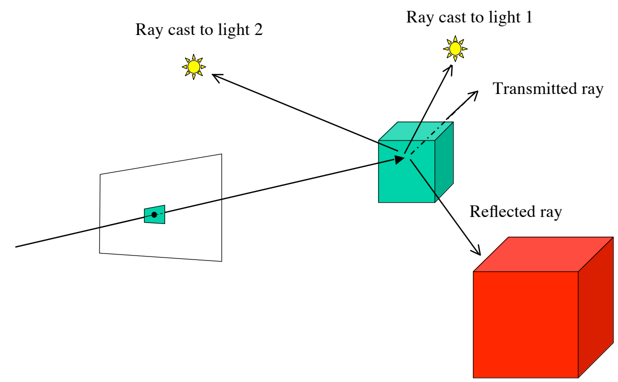

Recursive ray tracing rendering algorithm

- Cast ray from pinhole (projection center) through pixel, determine nearest intersection
- Compute components by casting rays
  - to sources = shadow ray (diffuse and for specular lobe)
  - along reflected direction = reflected ray
  - along transmitted dir = refracted ray
- Determine each component and add them up with contribution from ambient illumination.
- To determine some of the components, the ray tracer must be called **recursively**.

Recursive ray tracing rendering (cont)

- Recursion needs to stop at some point!
- Contributions die down after multiple bounces---there is no such thing as a perfect reflector---so we either set mirror reflections to be less than 100% (even if the user asks for 100%), or simply include an attenuation factor for each new ray.
- Can also model absorption due to light traveling in medium
  - Usually ignored in air, but depends on the application
  - Translucent absorption is exponential in depth

$$I = I_0 e^{-\alpha d}$$

- Recursion is stopped when contributions are too small
  - need to track the cumulative effect
  - common to also limit the depth explicitly

Mechanics

- Primary issue is intersection computations.
  - E.g. sphere, triangle.
- Polygon (should feel familiar!)
- Find point on plane of polygon and then determine if it is inside
  - One way is to make an argument with angles
  - Another way---thinking of the polygon as a surface of a polyhedra---is to check if the point is on the inside side of each of the other planes of the polyhedra.
- Sphere, relatively simple algebra.

Poly details

Review! (recall picking)

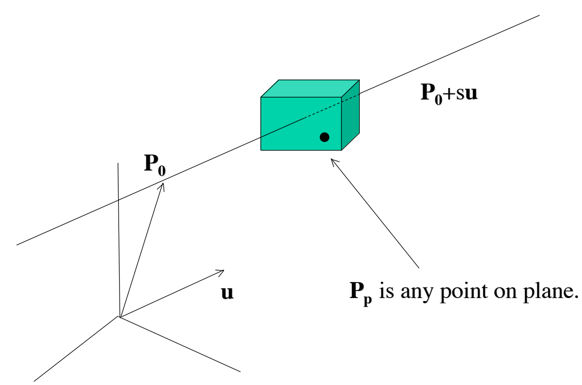

Poly details

Review! (recall picking)

To find the intersection of the ray and the plane, solve:

$$\left(\mathbf{P_0} + s\mathbf{u} - \mathbf{P_p}\right) \bullet \mathbf{n} = 0$$

Once you have the point of intersection, $\mathbf{P_i}$, test that it is inside by testing against all other faces.

$$\left(\mathbf{P_i} - \mathbf{P_p}\right) \bullet \mathbf{n} < 0$$

Note that **n** and $\mathbf{P_p}$ are now from those *other* faces.

Sphere details (H&B, 602)

May be helpful for grad version of A6.

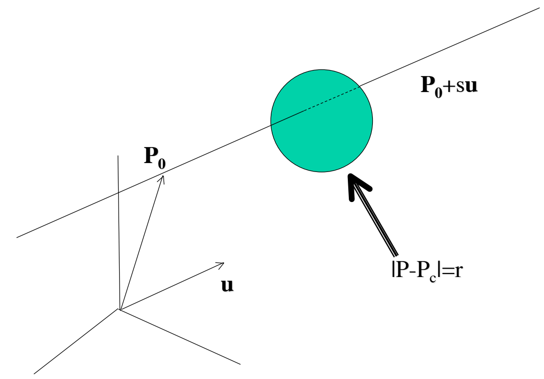

Sphere details (H&B, 602)

May be helpful for grad version of A6.

$$\left|\mathbf{P_0} + s\mathbf{u} - \mathbf{P_c}\right| = r$$

$$\left|\Delta\mathbf{P} + s\mathbf{u}\right| = r$$

$$\left(\Delta\mathbf{P} + s\mathbf{u}\right) \bullet \left(\Delta\mathbf{P} + s\mathbf{u}\right) = r^2$$

$$\Delta\mathbf{P} \bullet \Delta\mathbf{P} - r^2 + 2s\Delta\mathbf{P} \bullet \mathbf{u} + s^2\mathbf{u} \bullet \mathbf{u} = 0$$

The last expression is easily solved using the quadratic equation. If the discriminant is negative (complex solutions), then the ray does not intersect the sphere.

Sphere details (H&B, 602)

May be helpful for grad version of A6.

Recall that if: $as^2 + bs + c = 0$

The “discriminant” is: $b^2 - 4ac$

The solution is: $s = \dfrac{-b \pm \sqrt{b^2 - 4ac}}{2a}$

Note that in the book, **u** is a unit vector, so **u•u**=1, thus a=1, and b has a factor of 2 that is removed by dividing by 2a=2, to get equation 10-71.

Review!

Reflection Details

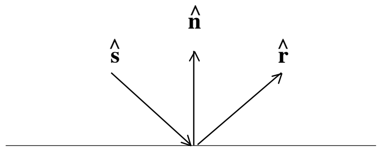

$\hat{\mathbf{s}} + \hat{\mathbf{r}} = k\hat{\mathbf{n}}$ and $\hat{\mathbf{n}} \bullet \hat{\mathbf{s}} = \hat{\mathbf{n}} \bullet \hat{\mathbf{r}}$

$$\hat{\mathbf{n}} \bullet \hat{\mathbf{s}} + \hat{\mathbf{n}} \bullet \hat{\mathbf{r}} = k \;\Rightarrow\; k = \mathbf{2}\hat{\mathbf{n}} \bullet \hat{\mathbf{s}}$$

So $\hat{\mathbf{r}} = \mathbf{2}(\hat{\mathbf{n}} \bullet \hat{\mathbf{s}})\hat{\mathbf{n}} - \hat{\mathbf{s}}$

Refraction Details

Index of refraction, n, is the ratio of speed of light in a vacuum, to speed of light in medium.

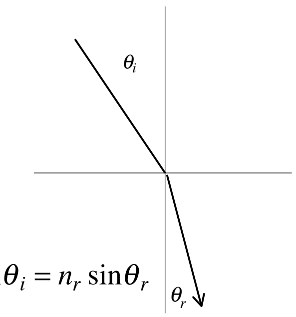

Typical values:

air:	1.00 (nearly)
water:	1.33
glass:	1.45-1.6
diamond:	2.2

$$n_i \sin\theta_i = n_r \sin\theta_r$$

The indices of refraction for the two media, and the incident angle, θ_i, yield the refracted angle θ_r. (Also need planarity).