



Vision and color



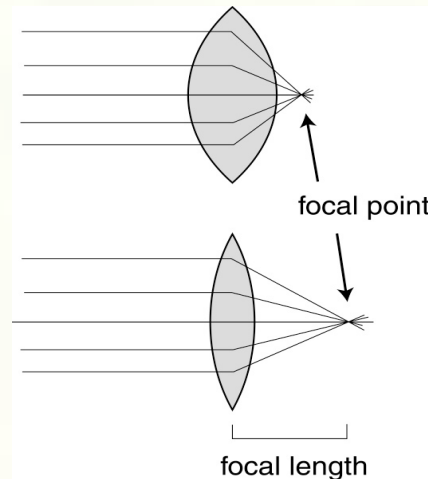
Reading

- Glassner, *Principles of Digital Image Synthesis*, pp. 5-32.
- Watt , Chapter 15.
- Brian Wandell. *Foundations of Vision*.
Sinauer Associates, Sunderland, MA, pp.
45-50 and 69-97, 1995.



Optics


- The human eye employs a lens to focus light.
- To quantify lens properties, we'll need some terms from *optics* (the study of sight and the behavior of light):
 - **Focal point** - the point where parallel rays converge when passing through a lens.
 - **Focal length** - the distance from the lens to the focal point.
 - **Diopter** - the reciprocal of the focal length, measured in meters.
 - Example: A lens with a “power” of 10D has a focal length of 0.1m.





Optics, cont'd

By tracing rays through a lens, we can generally tell where an object point will be focused to an image point:




This construction leads to the Gaussian lens formula:

$$\frac{1}{d_o} + \frac{1}{d_i} = \frac{1}{f}$$

Q: Given these three parameters, how does the human eye keep the world in focus?



Structure of the eye




Physiology of the human eye (Glassner, 1.1)

- The most important structural elements of the eye are:
 - **Cornea** - a clear coating over the front of the eye:
 - Protects eye against physical damage.
 - Provides initial focusing (40D).
 - **Iris** - Colored annulus with radial muscles.
 - **Pupil** - The hole whose size is controlled by the iris.



Structure of the eye, cont.




Physiology of the human eye (Glassner, 1.1)

- **Crystalline lens** - controls the focal distance:
 - Power ranges from 10 to 30D in a child.
 - Power and range reduces with age.
- **Ciliary body** - The muscles that compress the sides of the lens, controlling its power.

Q: As an object moves closer, do the ciliary muscles contract or relax to keep the object in focus?



Retina




Density of photoreceptors on the retina (Glassner, 1.4)

- **Retina** - a layer of photosensitive cells covering 200° on the back of the eye.
 - **Cones** - responsible for color perception.
 - **Rods** - Limited to intensity (but 10x more sensitive).
- **Fovea** - Small region (1 or 2°) at the center of the visual axis containing the highest density of cones (and no rods).



The human retina




Photomicrographs at increasing distances from the fovea. The large cells are cones; the small ones are rods. (Glassner , 1.5 and Wandell, 3.4).


- Photomicrographs at increasing distances from the fovea. The large cells are cones; the small ones are rods.



The human retina, cont'd



Photomicrograph of a cross-section of the retina near the fovea (Wandell, 5.1).




Light gathering by rods and cones (Wandell, 3.2)



Neuronal connections


- Even though the retina is very densely covered with photoreceptors, we have much more acuity in the fovea than in the periphery.



- In the periphery, the outputs of the photoreceptors are averaged together before being sent to the brain, decreasing the spatial resolution. As many as 1000 rods may converge to a single neuron.



Demonstrations of visual acuity



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



Blind spot demonstration (Glassner, 1.8)






Mach bands

- **Mach bands** were first discussed by Ernst Mach, an Austrian physicist.
- Appear when there are rapid variations in intensity, especially at C^0 intensity discontinuities:




- And at C^1 intensity discontinuities:





Mach bands, cont.

- **Possible cause:** lateral inhibition of nearby cells.




Lateral inhibition effect (Glassner, 1.25)

- **Q:** What image processing filter does this remind you of?



Higher Level Reasoning

- Many perceptual phenomena occur at a higher level in the brain




Edward H. Adelson

Checker Shadow Effect (Edward Adelson, 1995)




Higher Level Reasoning

- Many perceptual phenomena occur at a higher level in the brain



Edward H. Adelson




Checker Shadow Effect (Edward Adelson, 1995)



The radiant energy spectrum


- We can think of light as waves, instead of rays.
- Wave theory allows a nice arrangement of electromagnetic radiation (EMR) according to wavelength:





Emission spectra

- A light source can be characterized by an emission spectrum:



Emission spectra for daylight and a tungsten lightbulb (Wandell, 4.4)

The spectrum describes the energy at each wavelength.



What is color?

The eyes and brain turn an incoming emission spectrum into a discrete set of values.

The signal sent to our brain is somehow interpreted as *color*.

Color science asks some basic questions:


- When are two colors alike?
- How many pigments or primaries does it take to match another color?

One more question: why should we care?



Photopigments

- **Photopigments** are the chemicals in the rods and cones that react to light. Can respond to a single photon!
- Rods contain **rhodopsin**, which has peak sensitivity at about 500nm.




Rod sensitivity (Wandell ,4.6)

- Rods are active under low light levels, i.e., they are responsible for **scotopic** vision.




Univariance

- **Principle of univariance:** For any single photoreceptor, no information is transmitted describing the wavelength of the photon.



Measuring photoreceptor photocurrent (Wandell, 4.15)




Photocurrents measured for two light stimuli: 550nm (solid) and 659 nm (gray). The brightnesses of the stimuli are different, but the shape of the response is the same. (Wandell 4.17)



The color matching experiment

We can construct an experiment to see how to match a given test light using a set of lights called **primaries** with power control knobs.



The color matching experiment (Wandell, 4.10)

The primary spectra are $a(\lambda)$, $b(\lambda)$, $c(\lambda)$, ...

The power knob settings are A, B, C, ...



Rods and “color matching”

A rod responds to a spectrum through its spectral sensitivity function, $p(\lambda)$. The response to a test light, $t(\lambda)$, is simply:

$$P_t = \int t(\lambda)p(\lambda)d\lambda$$


How many primaries are needed to match the test light?

What does this tell us about rod color discrimination?



Cone photopigments

- Cones come in three varieties: L, M, and S.



Cone photopigment absorption (Glassner, 1.1)

- Cones are active under high light levels, i.e., they are responsible for **photopic** vision.



Cones and color matching

Color is perceived through the responses of the cones to light.

The response of each cone can be written simply as:

$$L_t = \int t(\lambda)l(\lambda)d\lambda$$

$$M_t = \int t(\lambda)m(\lambda)d\lambda$$

$$S_t = \int t(\lambda)s(\lambda)d\lambda$$

These are the only three numbers used to determine color.

Any pair of stimuli that result in the same three numbers will be indistinguishable.

How many primaries do you think we'll need to match t ?



Color matching

Let's assume that we need 3 primaries to perform the color matching experiment.

Consider three primaries, $a(\lambda)$, $b(\lambda)$, $c(\lambda)$, with three emissive power knobs, A , B , C .

The three knobs create spectra of the form:

$$e(\lambda) = Aa(\lambda) + Bb(\lambda) + Cc(\lambda)$$

What is the response of the l-cone?

$$\begin{aligned} L_{abc} &= \int e(\lambda)l(\lambda)d\lambda \\ &= \int [Aa(\lambda) + Bb(\lambda) + Cc(\lambda)]l(\lambda)d\lambda \\ &= \int Aa(\lambda)l(\lambda)d\lambda + \int Bb(\lambda)l(\lambda)d\lambda + \int Cc(\lambda)l(\lambda)d\lambda \\ &= A \int a(\lambda)l(\lambda)d\lambda + B \int b(\lambda)l(\lambda)d\lambda + C \int c(\lambda)l(\lambda)d\lambda \\ &= AL_a + BL_b + CL_c \end{aligned}$$

How about the m- and s-cones?



Color matching, cont'd

We end up with similar relations for all the cones:

$$L_{abc} = AL_a + BL_b + CL_c$$

$$M_{abc} = AM_a + BM_b + CM_c$$

$$S_{abc} = AS_a + BS_b + CS_c$$

We can re-write this as a matrix and equate to the test:

$$\begin{bmatrix} L_{abc} \\ M_{abc} \\ S_{abc} \end{bmatrix} = \begin{bmatrix} L_a & L_b & L_c \\ M_a & M_b & M_c \\ S_a & S_b & S_c \end{bmatrix} \begin{bmatrix} A \\ B \\ C \end{bmatrix} = \begin{bmatrix} L_t \\ M_t \\ S_t \end{bmatrix}$$

and then solve for the knob settings:

$$\begin{bmatrix} A \\ B \\ C \end{bmatrix} = \begin{bmatrix} L_a & L_b & L_c \\ M_a & M_b & M_c \\ S_a & S_b & S_c \end{bmatrix}^{-1} \begin{bmatrix} L_t \\ M_t \\ S_t \end{bmatrix}$$

In other words, we can choose the knob settings to cause the cones to react as we please!

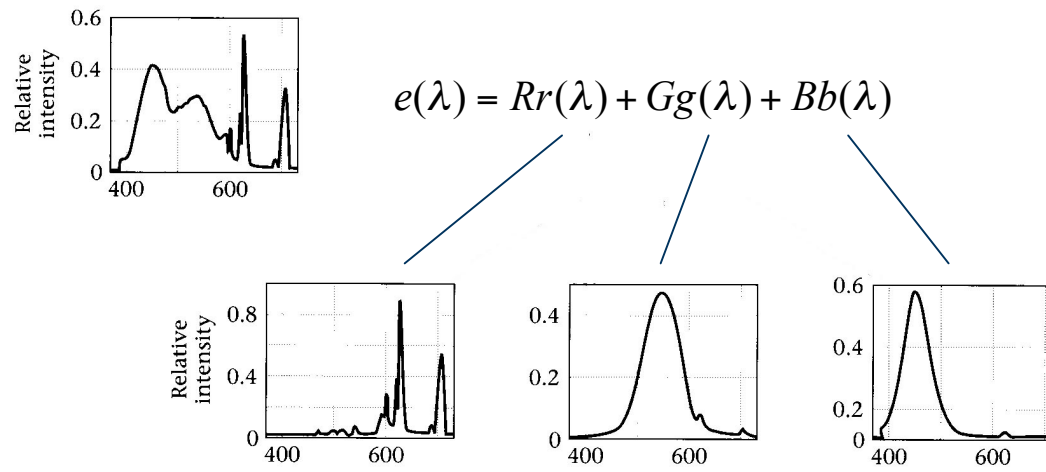
Well, one little “gotcha” – we may need to set the knob values to be negative.



Choosing Primaries

The primaries could be three color (monochromatic) lasers.

But, they can also be non-monochromatic, e.g., monitor phosphors:



Emission spectra for RGB monitor phosphors (Wandell B.3)



Color as linear projection

We can think of spectral functions in sampled form as n -dimensional vectors, where n is the number of samples.

$$t(\lambda) = \begin{bmatrix} t(\lambda_0) \\ t(\lambda_0 + \Delta_\lambda) \\ t(\lambda_0 + 2\Delta_\lambda) \\ \vdots \\ t(\lambda_0 + (n-1)\Delta_\lambda) \end{bmatrix}$$

In that case, computing the rod response is a projection from n dimensions to 1 dimension:

$$P_t = \begin{bmatrix} \cdots & p(\lambda) & \cdots \end{bmatrix} \begin{bmatrix} \vdots \\ t(\lambda) \\ \vdots \end{bmatrix}$$


Likewise, computing cone responses is a projection down to 3 dimensions:

$$\begin{bmatrix} L_t \\ M_t \\ S_t \end{bmatrix} = \begin{bmatrix} \cdots & l(\lambda) & \cdots \\ \cdots & m(\lambda) & \cdots \\ \cdots & s(\lambda) & \cdots \end{bmatrix} \begin{bmatrix} \vdots \\ t(\lambda) \\ \vdots \end{bmatrix}$$



Emission Spectrum is not Color

- Clearly, information is lost in this projection step...
- Different light sources can evoke exactly the same colors. Such lights are called **metamers**.




*A dim tungsten bulb and an RGB monitor set up to emit a metameric spectrum
(Wandell 4.11)*



Colored Surfaces

So far, we've discussed the colors of lights. How do *surfaces* acquire color?



Subtractive colour mixing (Wasserman 2.2)


A surface's **reflectance**, $\rho(\lambda)$, is its tendency to reflect incoming light across the spectrum.

Reflectance is combined “**subtractively**” with incoming light. Actually, the process is *multiplicative*:

$$I(\lambda) = \rho(\lambda)t(\lambda)$$



Subtractive Metamers



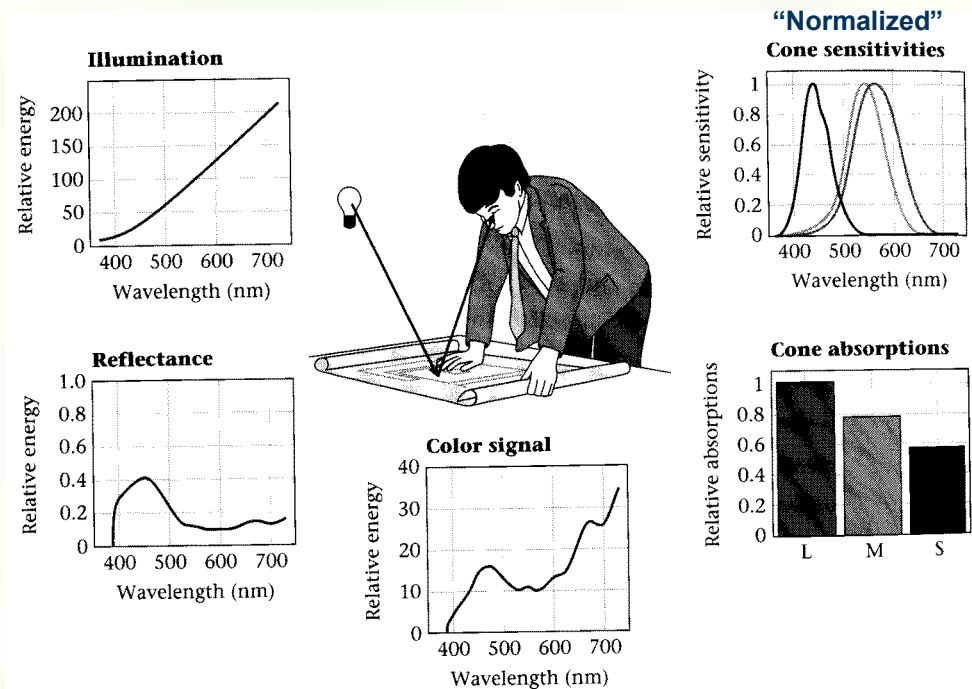
Surfaces that are metamers under only some lighting conditions (Wasserman 3.9)

Reflectance adds a whole new dimension of complexity to color perception.

The solid curve appears green indoors and out. The dashed curve looks green outdoors, but brown under incandescent light.



Illustration of Color Appearance



How light and reflectance become cone responses (Wandell, 9.2)



Lighting design

When deciding the kind of “feel” for an architectural space, the spectra of the light sources is critical.

Lighting design centers have displays with similar scenes under various lighting conditions.

For example:





What's wrong with RGB?


- RGB cannot represent some colors that we perceive! (using positive weights)
- In particular, pure hues other than R, G, or B cannot be represented.
- Reason: Adding two pure hues yields a de-saturated color (effectively has some “white” mixed in).
- Q: How can we fix this?

A: Pick “supersaturated” primaries.
That's what CIE color space does.

- Note that the supersaturated primaries cannot be perceived! So we've traded one problem for another...




Perceivable portion of CIE space





The CIE XYZ System

A standard created in 1931 by CIE, defined in terms of three color matching functions.



The XYZ color matching functions (Wasserman 3.8)

These functions are related to the cone responses roughly as:

$$\bar{x}(\lambda) \approx k_1 s(\lambda) + k_2 l(\lambda)$$

$$\bar{y}(\lambda) \approx k_3 m(\lambda)$$

$$\bar{z}(\lambda) \approx k_4 s(\lambda)$$



CIE Coordinates

Given an emission spectrum, we can use the CIE matching functions to obtain the X , Y and Z coordinates.

$$X = \int \bar{x}(\lambda)t(\lambda)d\lambda$$


$$Y = \int \bar{y}(\lambda)t(\lambda)d\lambda$$

$$Z = \int \bar{z}(\lambda)t(\lambda)d\lambda$$

Using the equations from the previous page, we can see that XYZ coordinates are closely related to LMS responses.



The CIE Colour Blob




Different views of the CIE color space (Foley II.1)



The CIE Chromaticity Diagram

- The CIE Chromaticity Diagram is a projection of the plane $X+Y+Z=1$.




The chromaticity diagram (a kind of slice through CIE space, Wasserman 3.7)

- **Dominant wavelengths** or **hues** go around the perimeter of the chromaticity diagram.
 - A color's dominant wavelength is where a line from white through that color intersects the perimeter.
 - Some colors, called *non-spectral* color's, don't have a dominant wavelength.



More About Chromaticity




- **Excitation purity or saturation** is measured in terms of a color's position on the line to its dominant wavelength.
- Complementary colors lie on opposite sides of white, and can be mixed to get white.



Gamuts

Not every output device can reproduce every color. A device's range of reproducible colors is called its **gamut**.



Gamuts of a few common output devices in CIE space (Foley, II.2)



Next class: Modern graphics HW

■ Topic:

How do we make the Z buffer algorithm run really fast?

■ Read:

- ["The GeForce 6 Series GPU Architecture"](#), Emmett Kilgariff and Randima Fernando, Chapter 30 in GPU Gems 2, edited by Matt Pharr, 2005.