Color

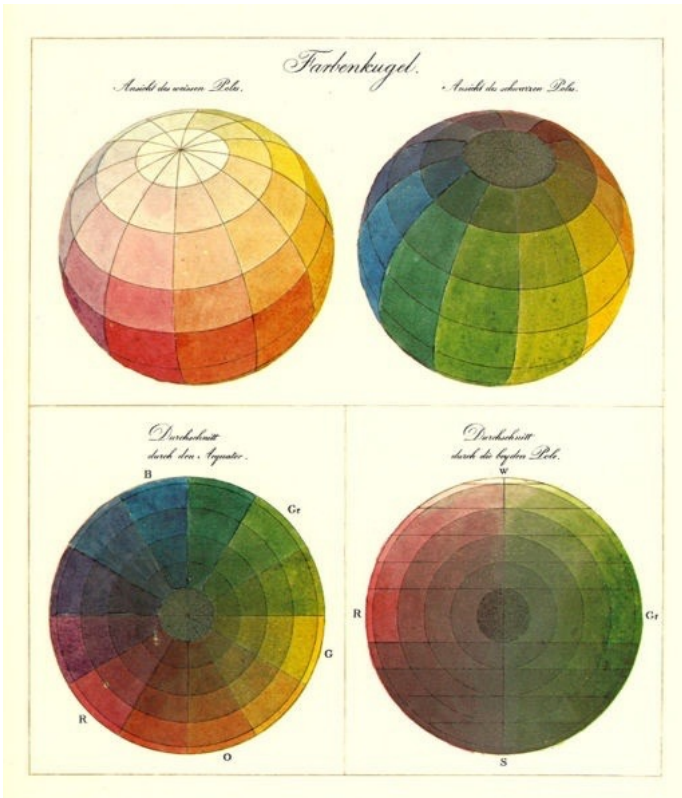

Phillip Otto Runge (1777-1810)

Outline

- Physical origin of color
- Spectra of sources and surfaces
- Physiology of color vision
- Quantifying color perception
- Color spaces
- Color constancy, white balance

What is color?

- Color is the result of interaction between physical light in the environment and our visual system

- "Color is a psychological property of our visual experiences when we look at objects and lights, *not* a physical property of those objects or lights"

  -- S. Palmer, *Vision Science: Photons to Phenomenology*

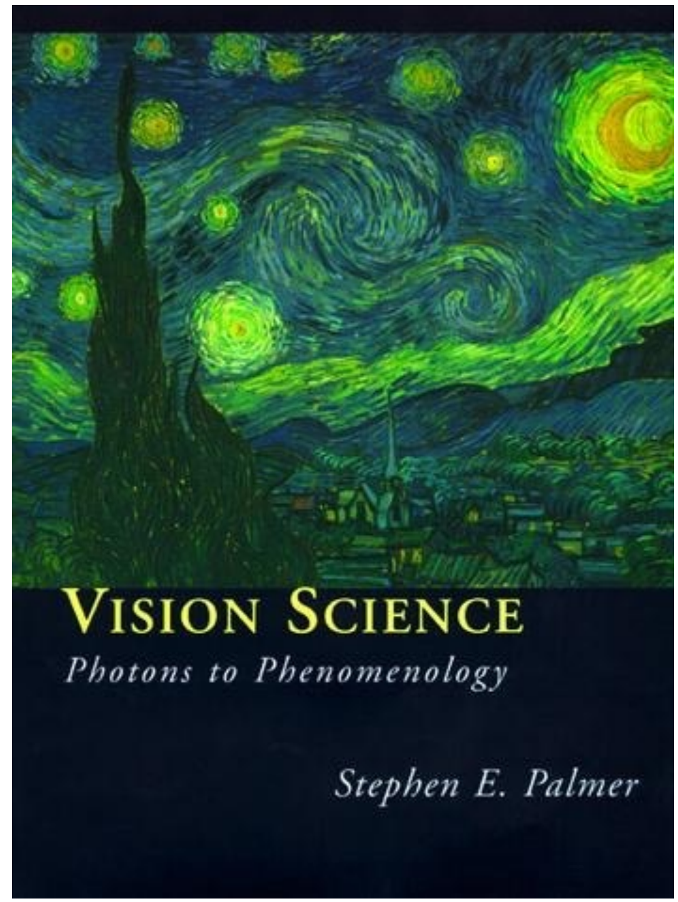

Causes of color

- The sensation of color is occurs in the brain
- You can get this sensation in a variety of ways
  - Dreaming
  - Hallucination
  - Pressure on the eyelids
- Most usual
  - Response of sensors in the eye to light with different amounts of energy at different wavelengths

XXXXX	BLUE	YELLOW
XXXXX	GREEN	BLUE
XXXXX	RED	GREEN
XXXXX	YELLOW	RED
XXXXX	BLUE	YELLOW
XXXXX	RED	GREEN
XXXXX	GREEN	BLUE
XXXXX	BLUE	YELLOW
XXXXX	YELLOW	RED
XXXXX	RED	GREEN

Stroop effect

- Color perception affects all sorts of other perceptual phenomena

- You name colors slowly IF the word is a different color name!

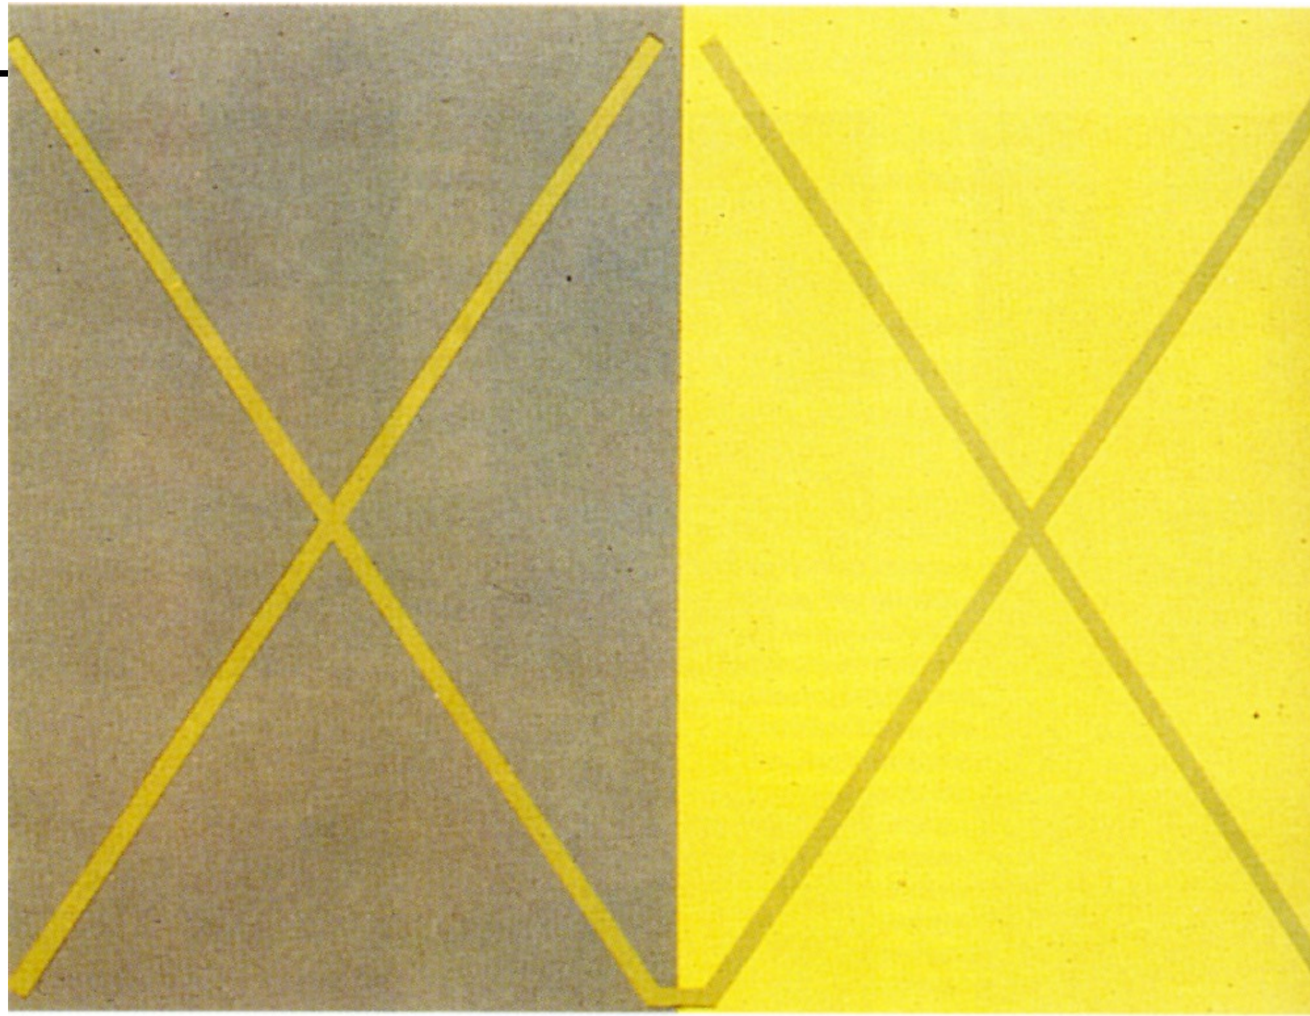

Your perception of color is relative to other colors

Different amounts of energy at different wavelengths

Light could be

- emitted with wavelengths absent (flourescent light vs. incandescent light)
- differentially reflected - e.g. paint on a surface
- differentially refracted - e.g. Newton's prism
- subject to wavelength dependent specular reflection (most metals).
- Flourescence -
  - invisible wavelengths absorbed and reemitted at visible wavelengths.
- Phosphorescence (ditto, energy, longer timescale)

Electromagnetic spectrum

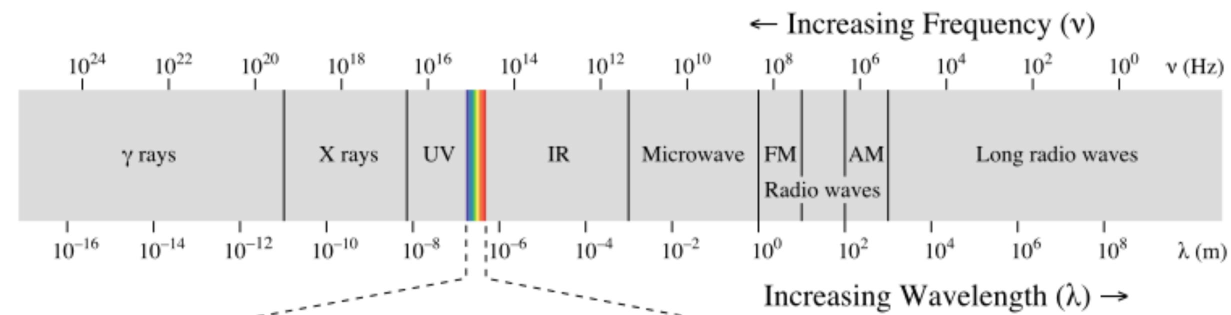

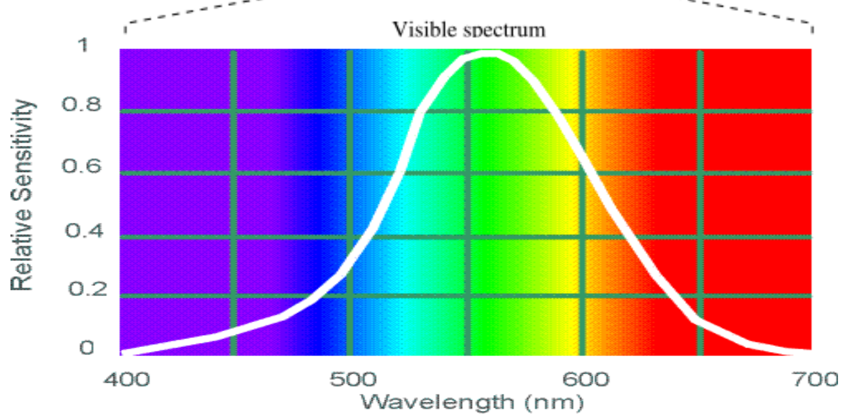

Human Luminance Sensitivity Function

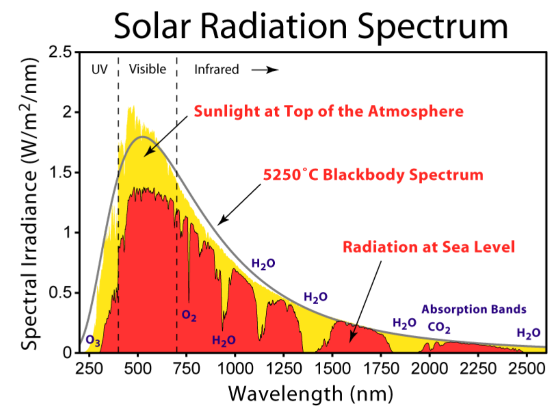

The Physics of Light

Any source of light can be completely described physically by its spectrum: the amount of energy emitted (per time unit) at each wavelength 400 - 700 nm.

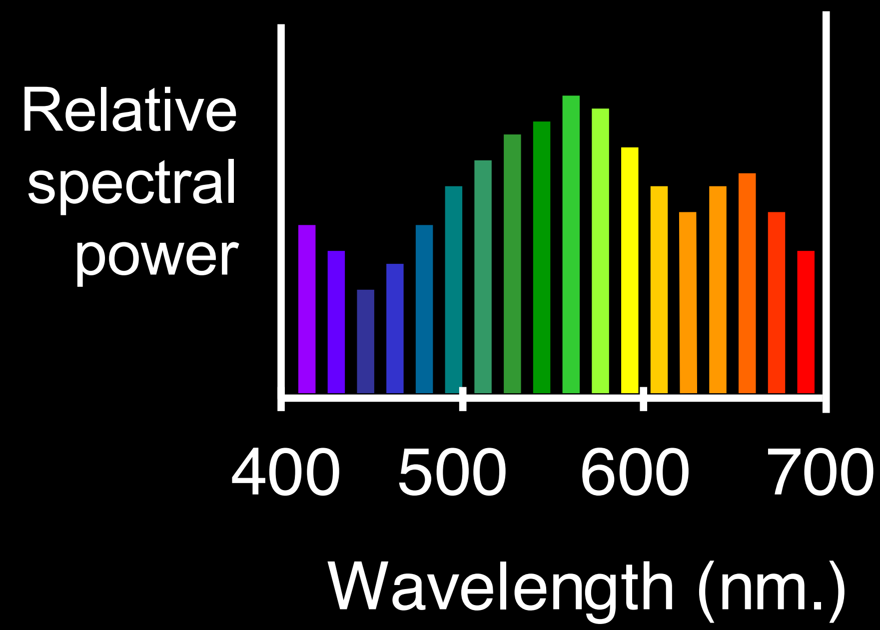

Spectra of Light Sources

Some examples of the spectra of light sources

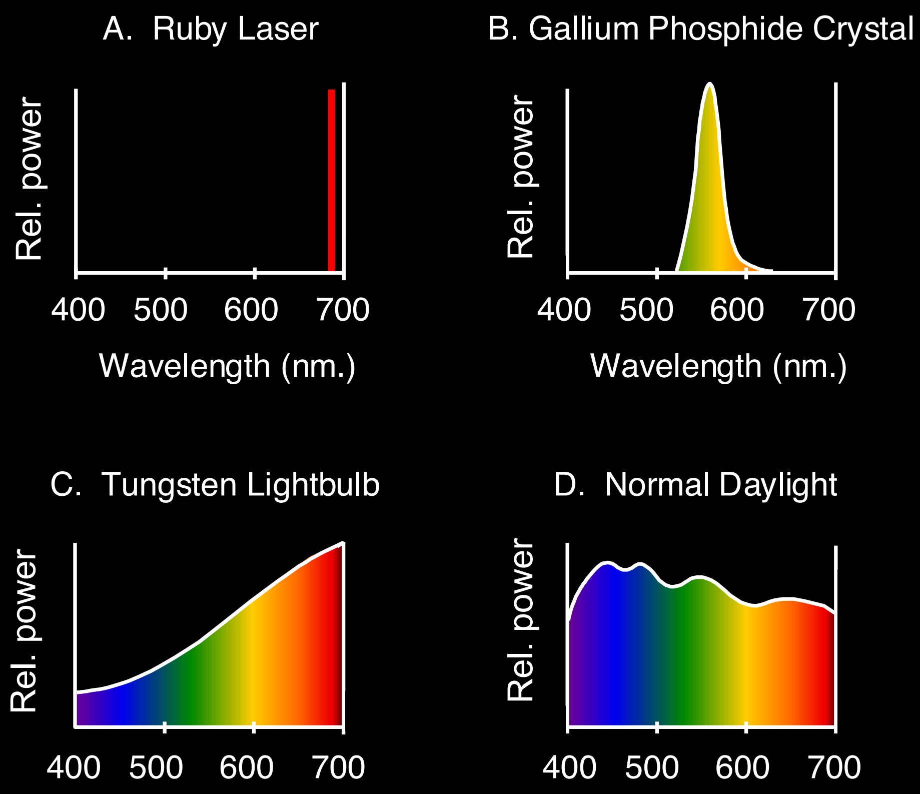

Spectra of light sources

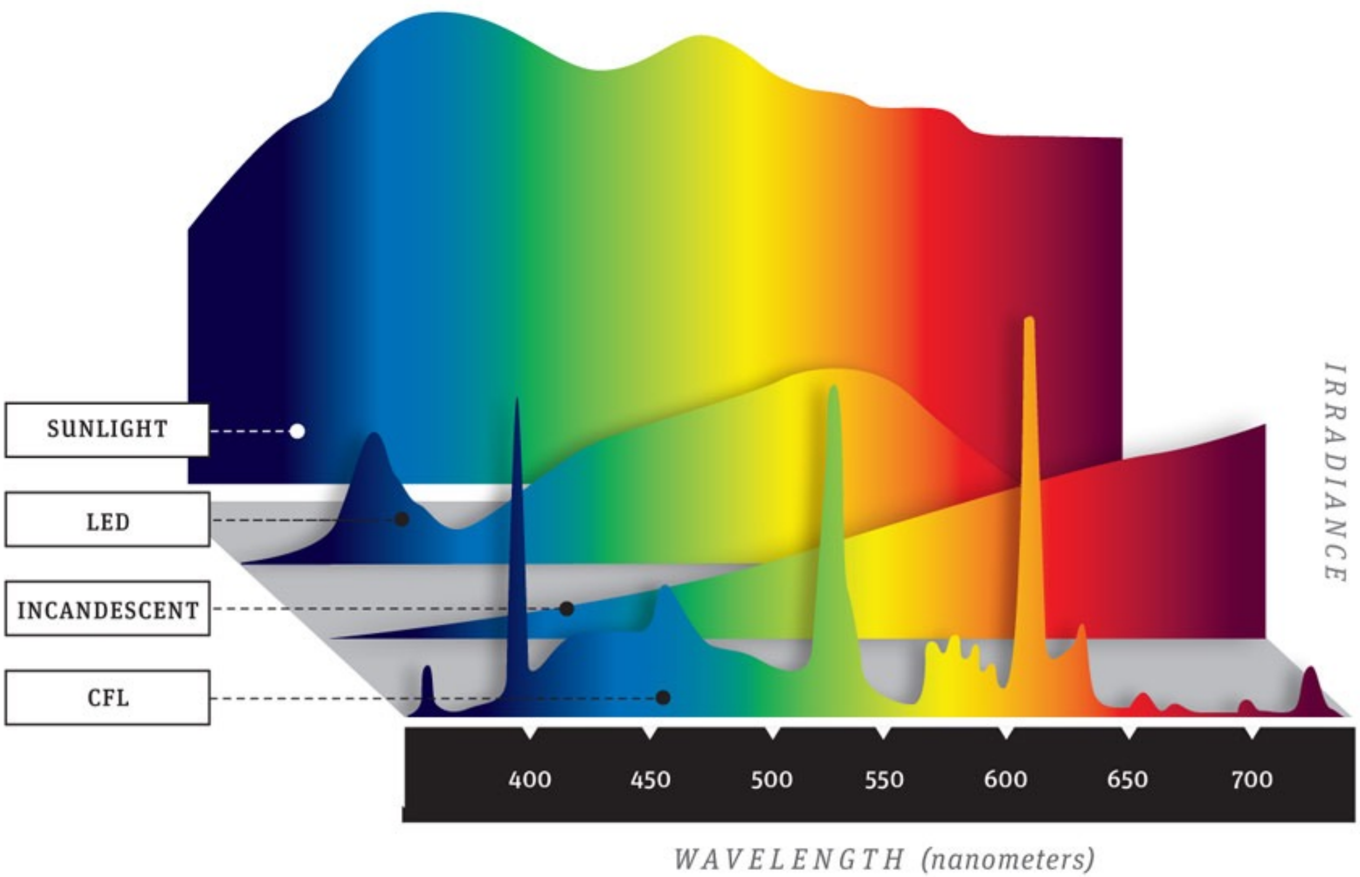

Source: Popular Mechanics

Reflectance Spectra of Surfaces

Some examples of the reflectance spectra of surfaces

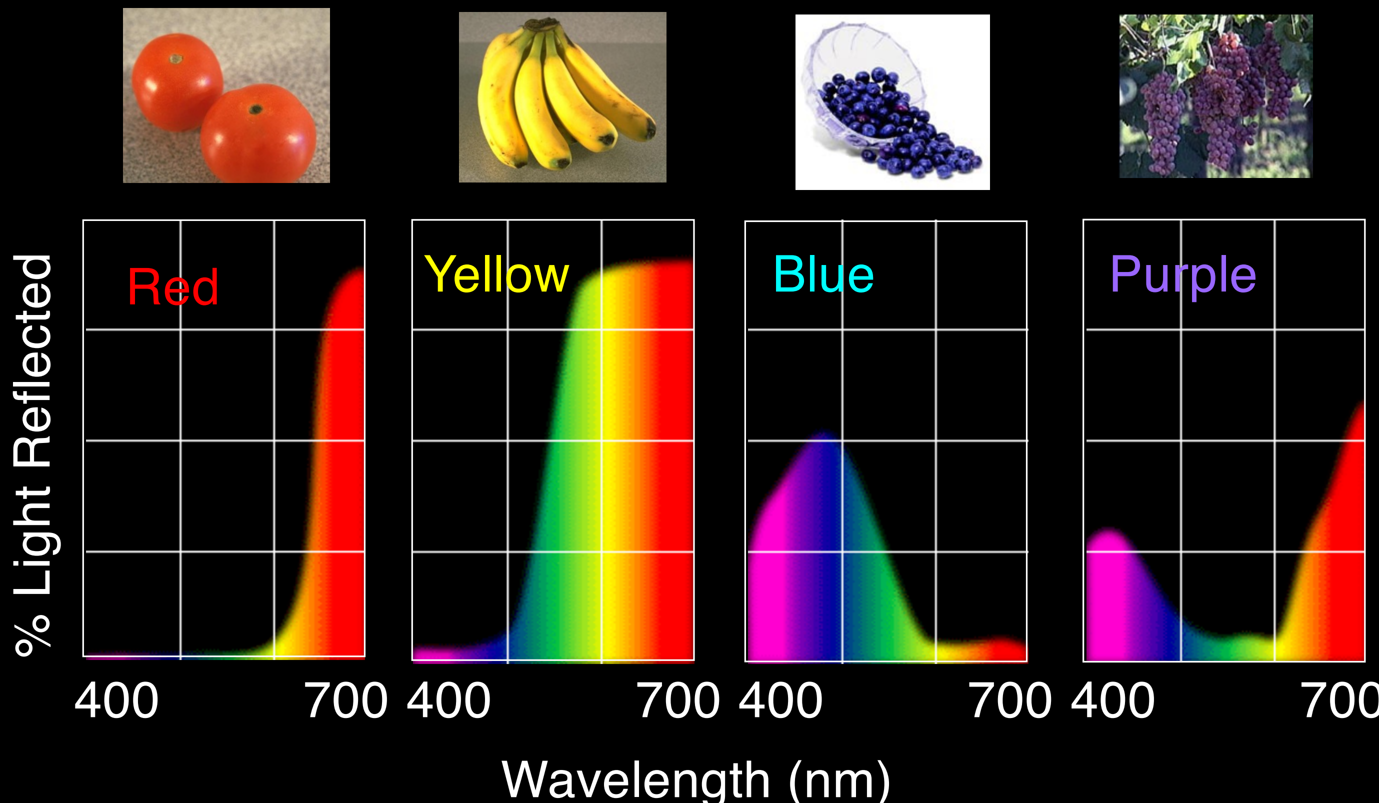

Interaction of light and surfaces

- Reflected color is the result of interaction of light source spectrum with surface reflectance

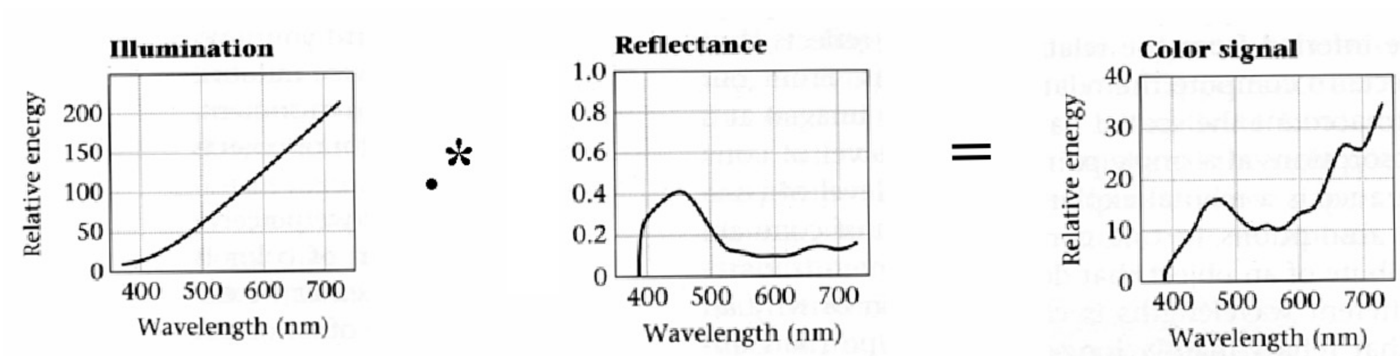

Interaction of light and surfaces

- What is the observed color of any surface under monochromatic light?

[Olafur Eliasson, *Room for one color*]

Outline

- Physical origin of color
- Spectra of sources and surfaces
- Physiology of color vision

The Eye

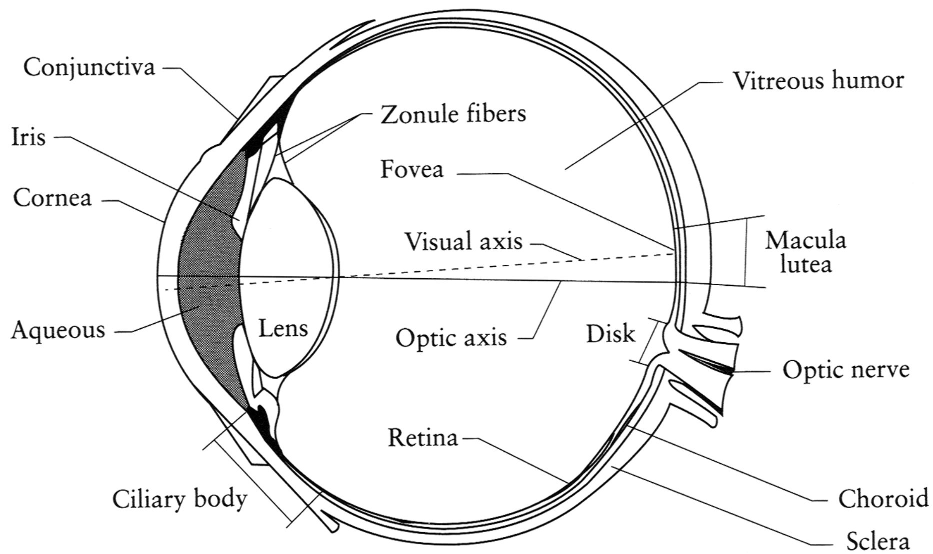

The human eye is a camera (sort of)

- **Lens** - changes shape by using ciliary muscles (to focus on objects at different distances)
- **Pupil** - the hole (aperture) whose size is controlled by the iris
- **Iris** - colored annulus with radial muscles
- **Retina** - photoreceptor cells

Slide by Steve Seitz

Rods and cones, fovea

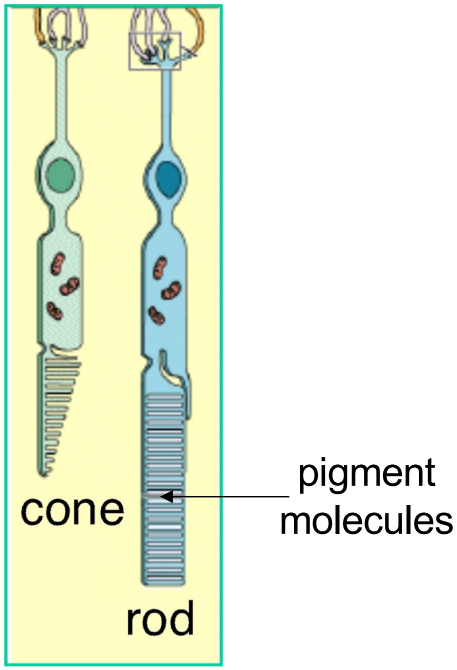

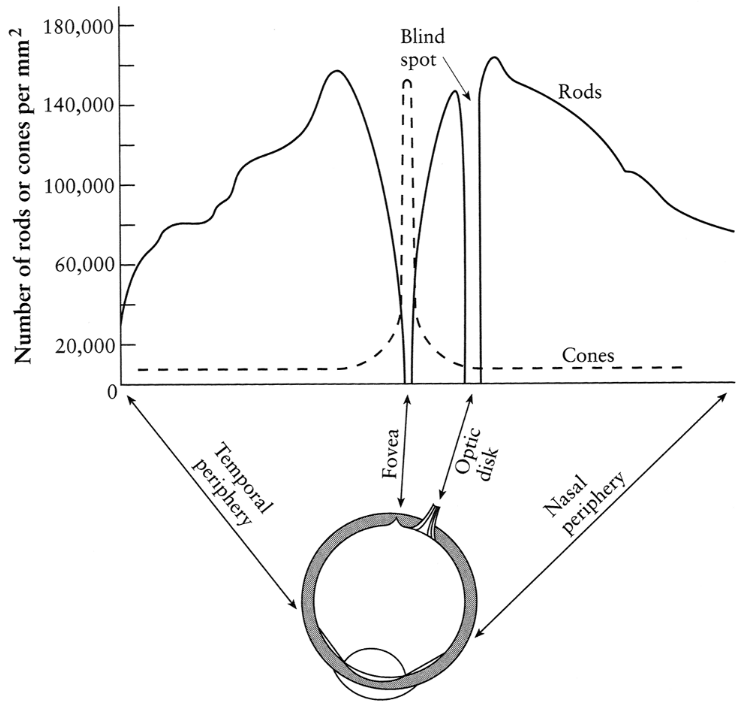

Rods are responsible for intensity, cones for color perception

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

- **Fovea** - Small region (1 or 2°) at the center of the visual field containing the highest density of cones – *and no rods*

Slide by Steve Seitz

Rod / cone sensitivity

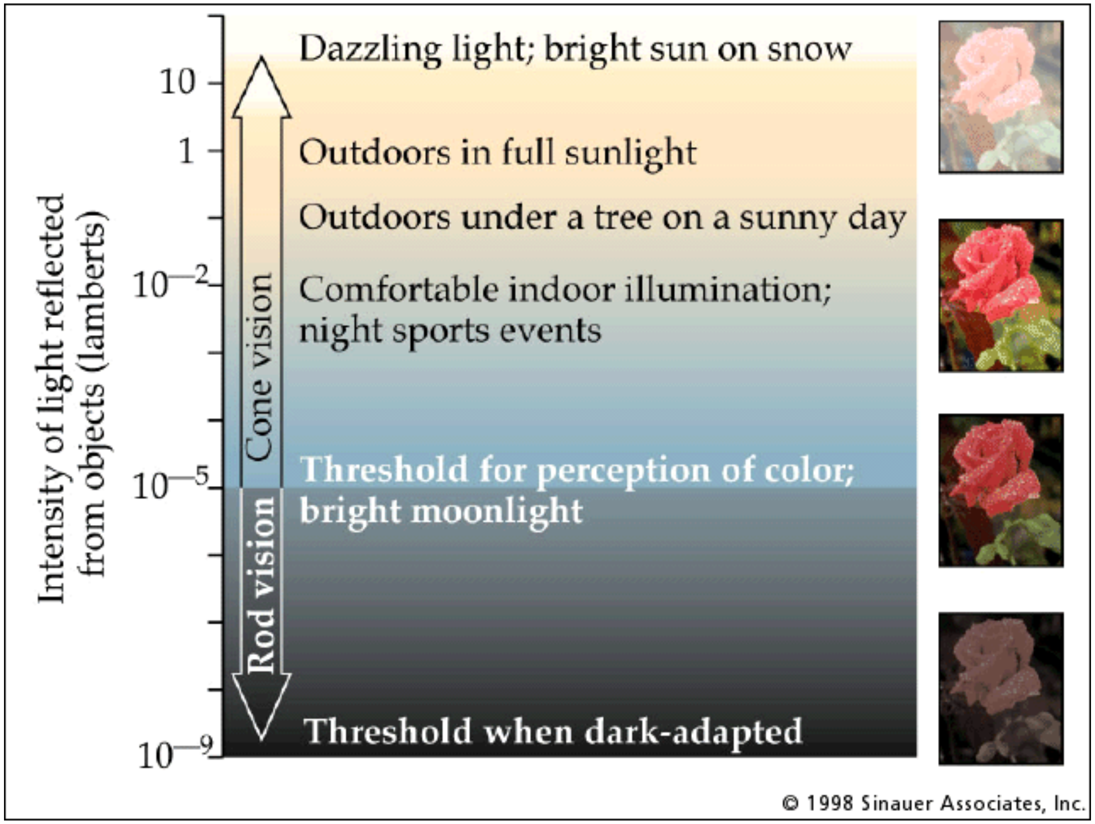

Why can't we read in the dark?

Slide by A. Efros

Physiology of Color Vision

Three kinds of cones:

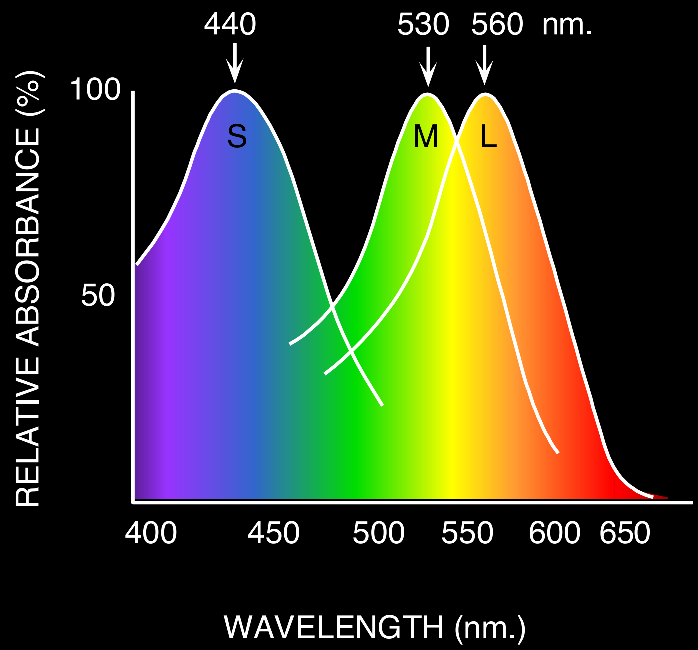

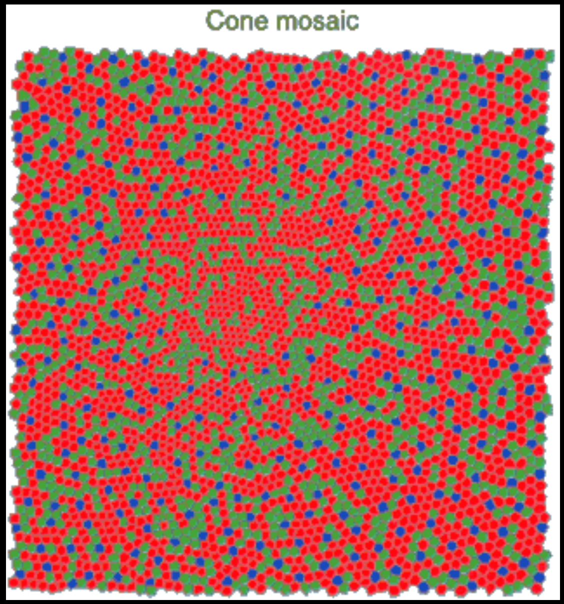

- Ratio of L to M to S cones: approx. 10:5:1
- Almost no S cones in the center of the fovea

Physiology of color vision: Fun facts

- Some animals have one (night animals), two (e.g., dogs), four (fish, birds), five (pigeons, some reptiles/amphibians), or even 12 (mantis shrimp) types of cone

http://en.wikipedia.org/wiki/Color_vision

Adapted from D. Hoiem

Color perception

- Color receptors act as filters on the spectrum
  - Let $t(\lambda)$ be the target spectrum and $c_1(\lambda)$, $c_2(\lambda)$, and $c_3(\lambda)$ be the response curves of the three types of cones
  - Then the total response of each type of cone is:

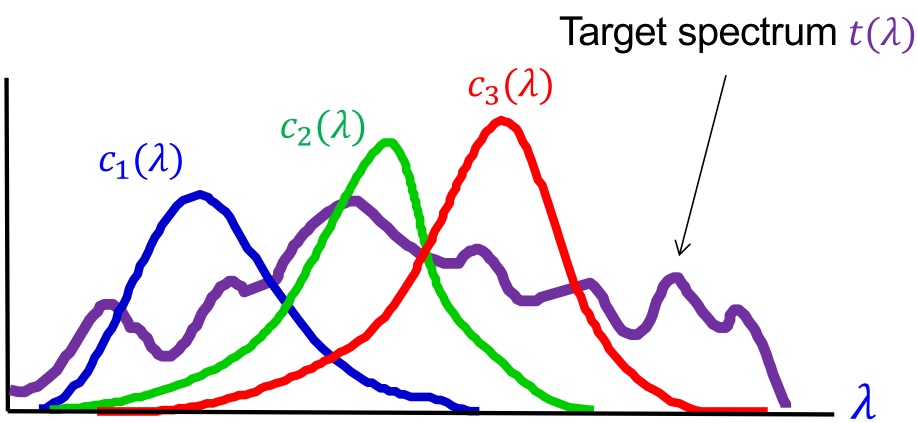

Color perception

- Color receptors act as filters on the spectrum
  - Let $t(\lambda)$ be the target spectrum and $c_1(\lambda)$, $c_2(\lambda)$, and $c_3(\lambda)$ be the response curves of the three types of cones
  - Then the total response of each type of cone is:

$$R_1 = \int_\lambda t(\lambda) c_1(\lambda) d\lambda$$

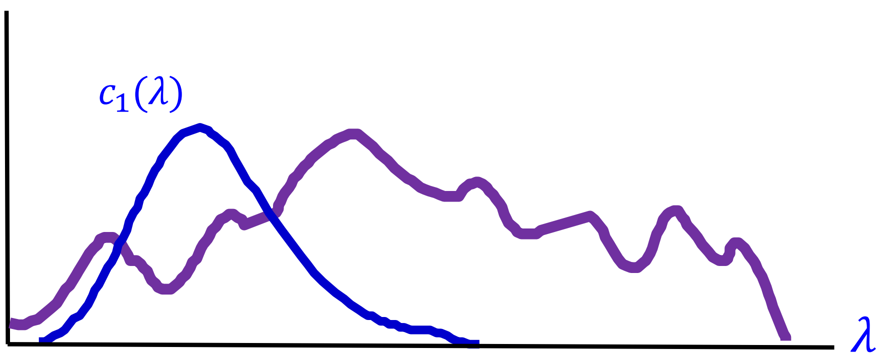

Color perception

- Color receptors act as filters on the spectrum
  - Let $t(\lambda)$ be the target spectrum and $c_1(\lambda)$, $c_2(\lambda)$, and $c_3(\lambda)$ be the response curves of the three types of cones
  - Then the total response of each type of cone is:

$$R_1 = \int_\lambda t(\lambda) c_1(\lambda) d\lambda \qquad R_2 = \int_\lambda t(\lambda) c_2(\lambda) d\lambda$$

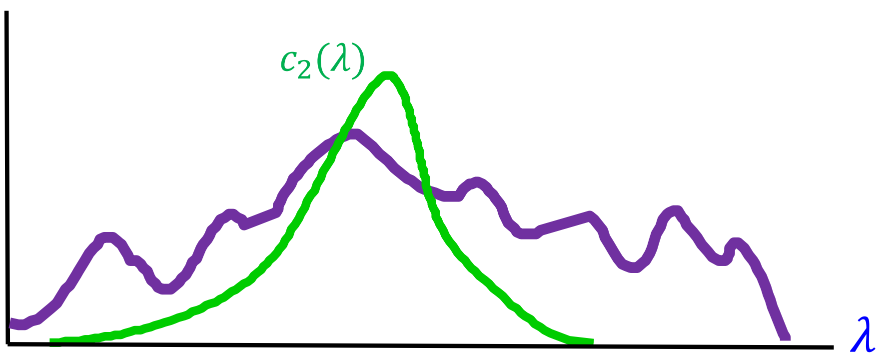

Color perception

- Color receptors act as filters on the spectrum
  - Let $t(\lambda)$ be the target spectrum and $c_1(\lambda)$, $c_2(\lambda)$, and $c_3(\lambda)$ be the response curves of the three types of cones
  - Then the total response of each type of cone is:

$$R_1 = \int_\lambda t(\lambda)c_1(\lambda)d\lambda \qquad R_2 = \int_\lambda t(\lambda)c_2(\lambda)d\lambda \qquad R_3 = \int_\lambda t(\lambda)c_3(\lambda)d\lambda$$

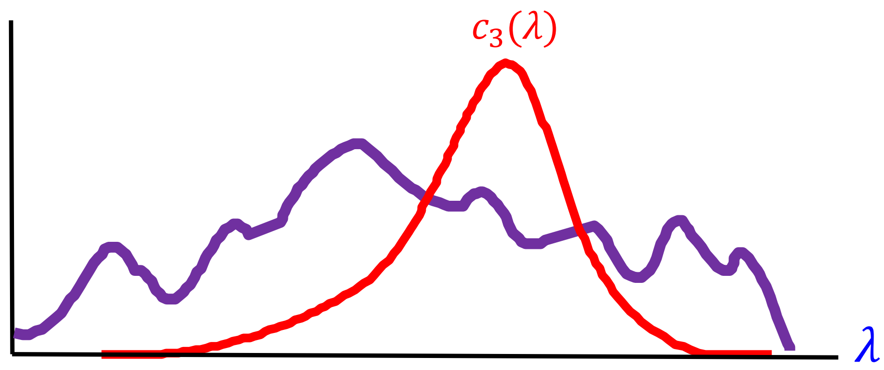

Color perception

- Color receptors act as filters on the spectrum
  - Each type of cone returns one number
- How can we represent an entire spectrum with three numbers?
  - We can't – most of the information is lost!

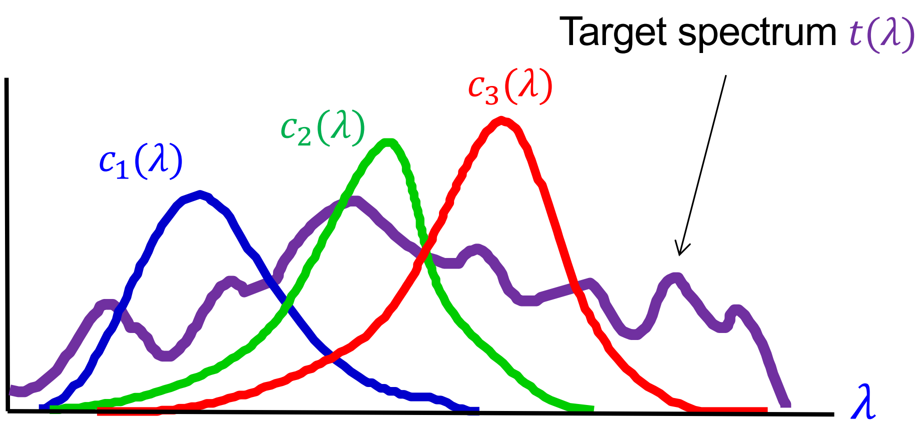

Metamers: Spectra that appear indistinguishable

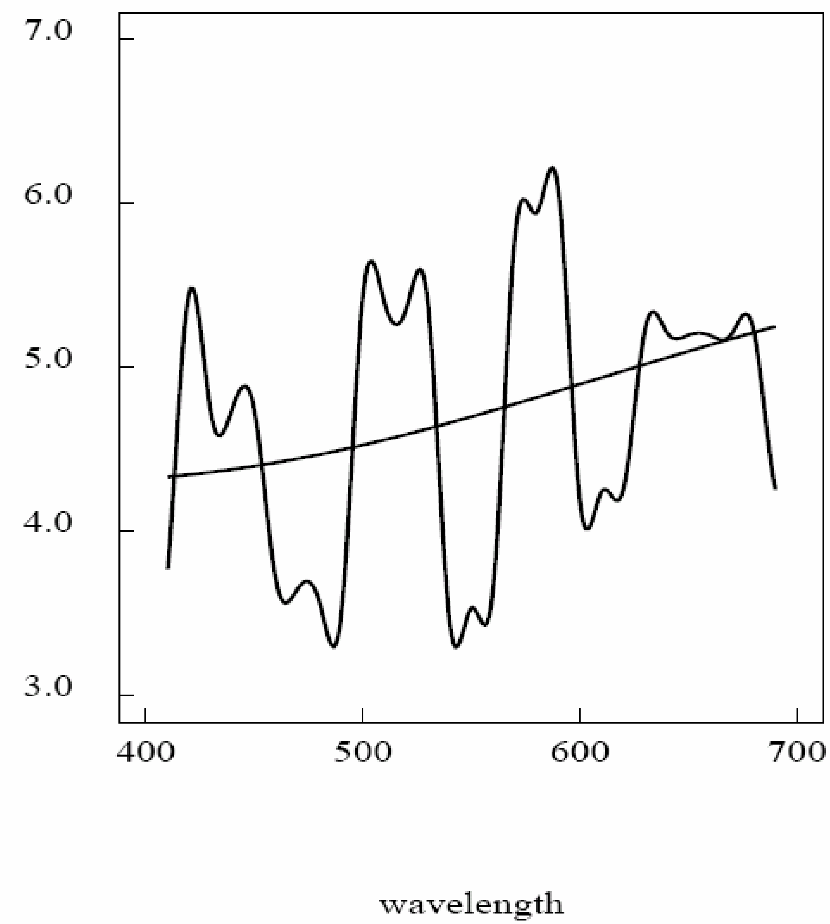

Color: Outline

- Physical origin of color
- Spectra of sources and surfaces
- Physiology of color vision
- Quantifying color perception

Quantifying color perception

- Spectral distributions go through a “black box” (human visual system) and are perceived as color
- The only way to quantify the “black box” is to perform a human study

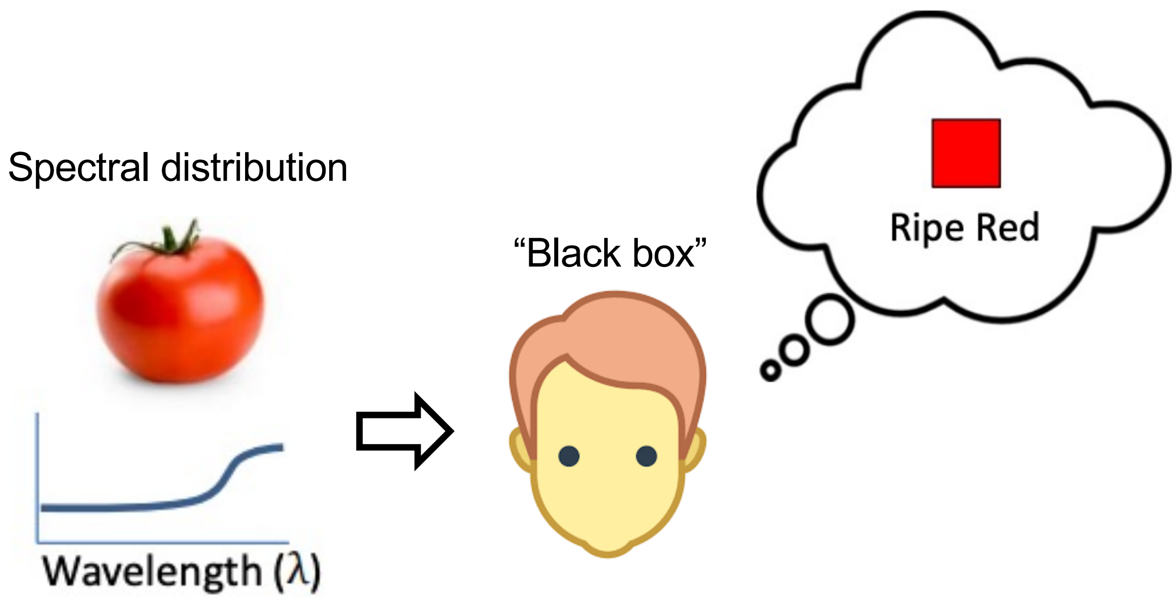

Source: M. Brown

Color matching experiments

- We would like to understand which spectra produce the same color sensation in people under similar viewing conditions

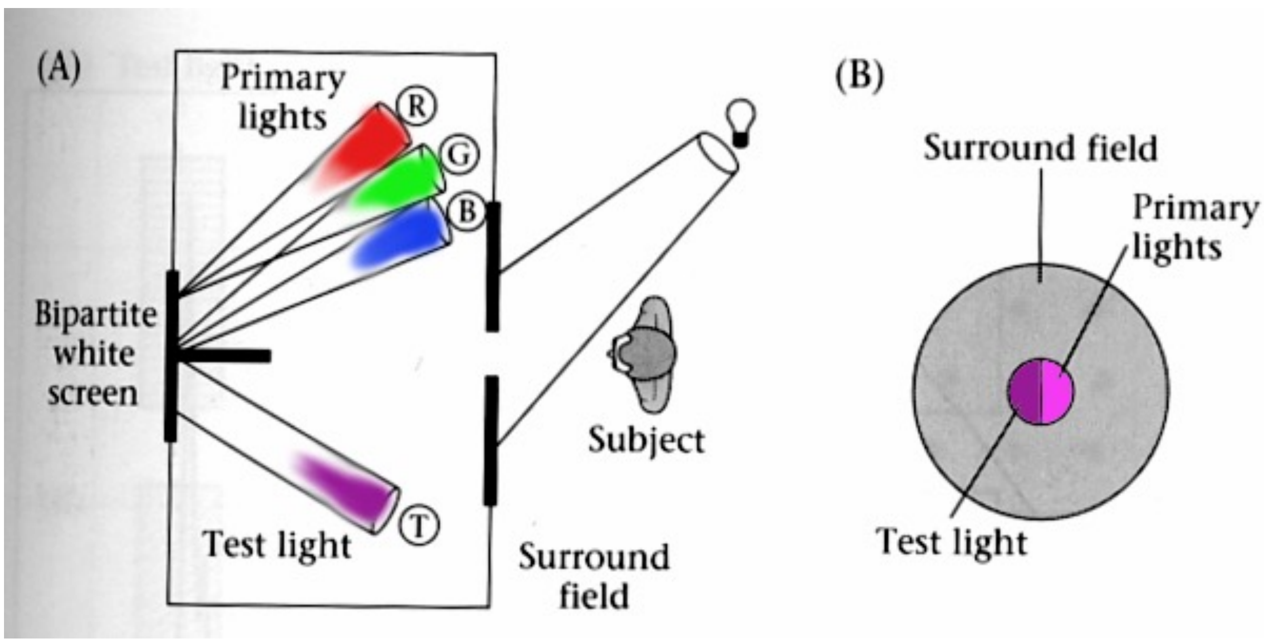

Wandell, Foundations of Vision, 1995

Color matching experiment 1

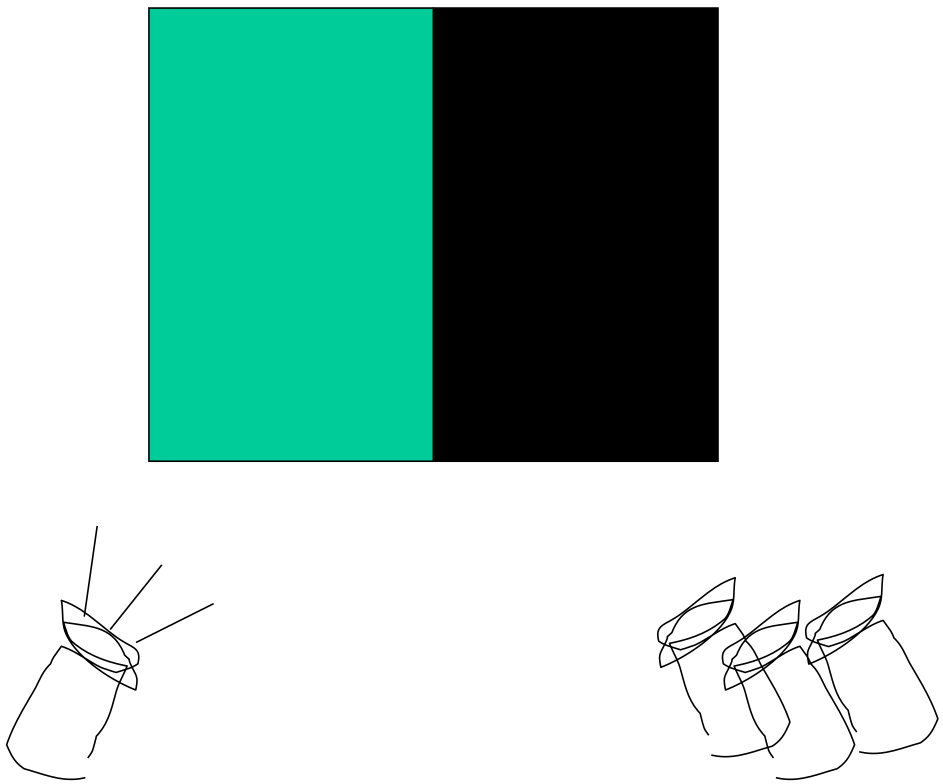

Source: W. Freeman

Color matching experiment 1

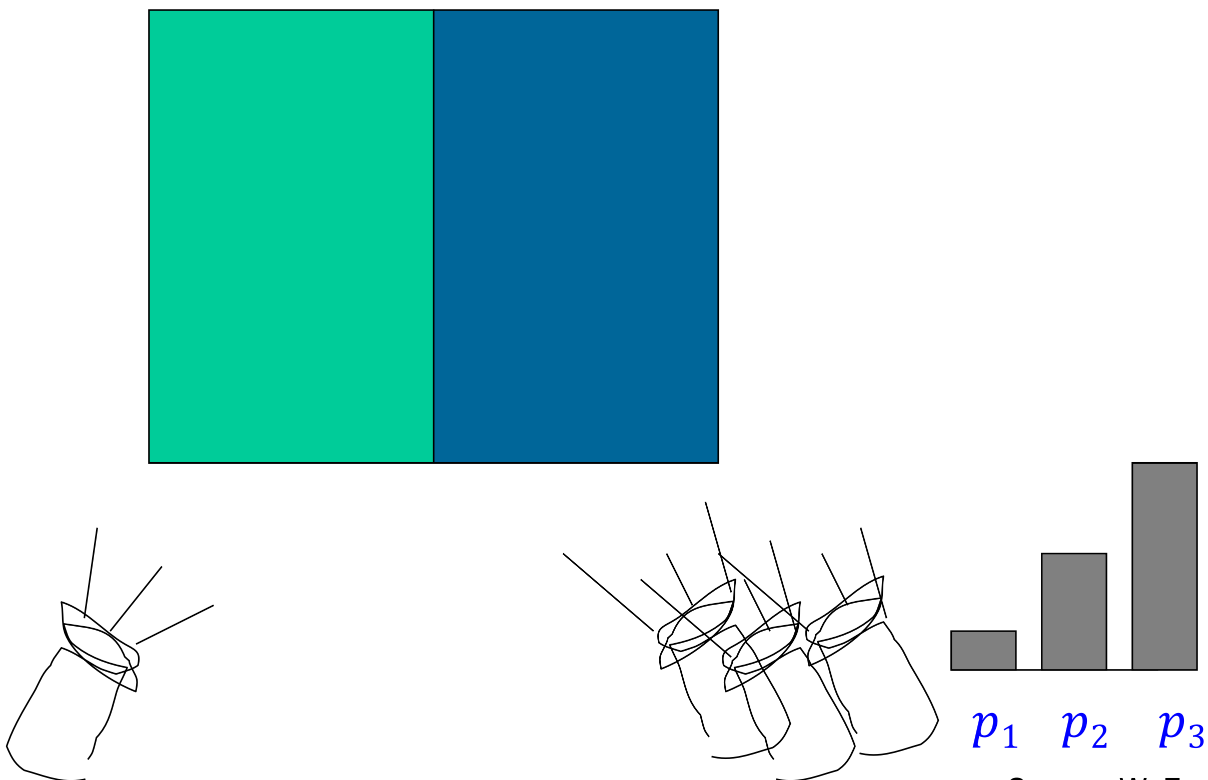

Source: W. Freeman

Color matching experiment 1

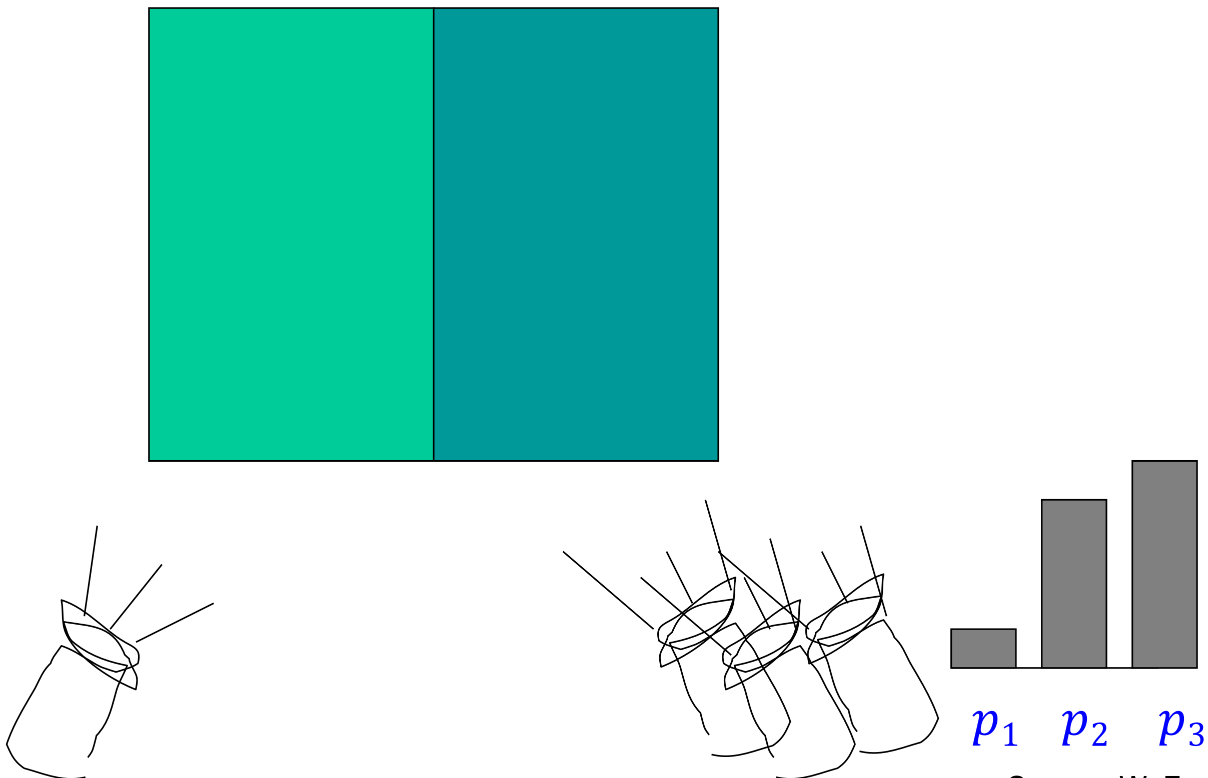

Source: W. Freeman

Color matching experiment 1

The primary color amounts needed for a match

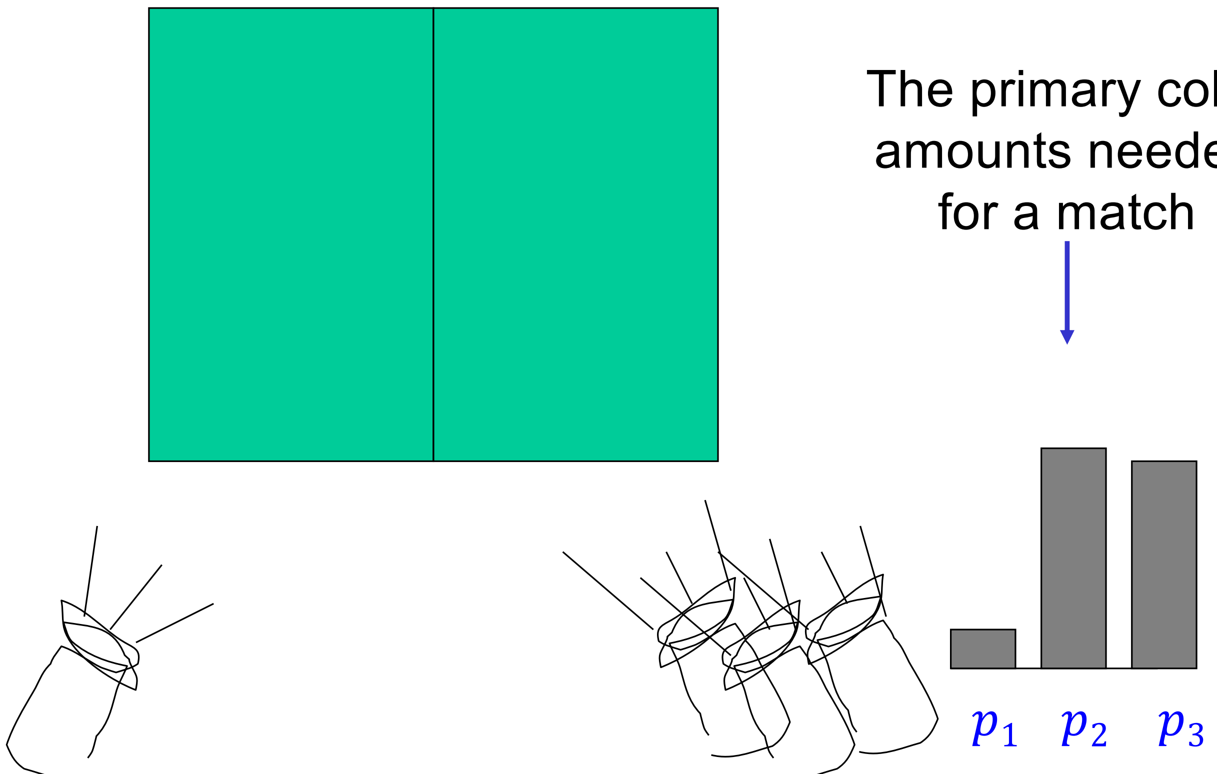

Source: W. Freeman

Color matching experiment 2

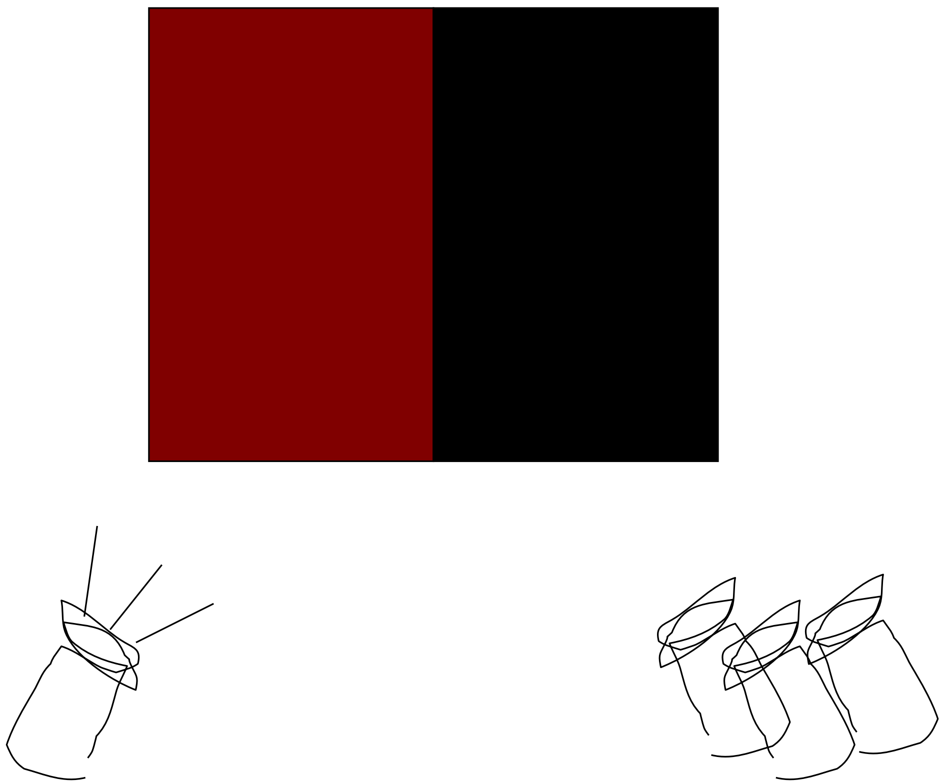

Source: W. Freeman

Color matching experiment 2

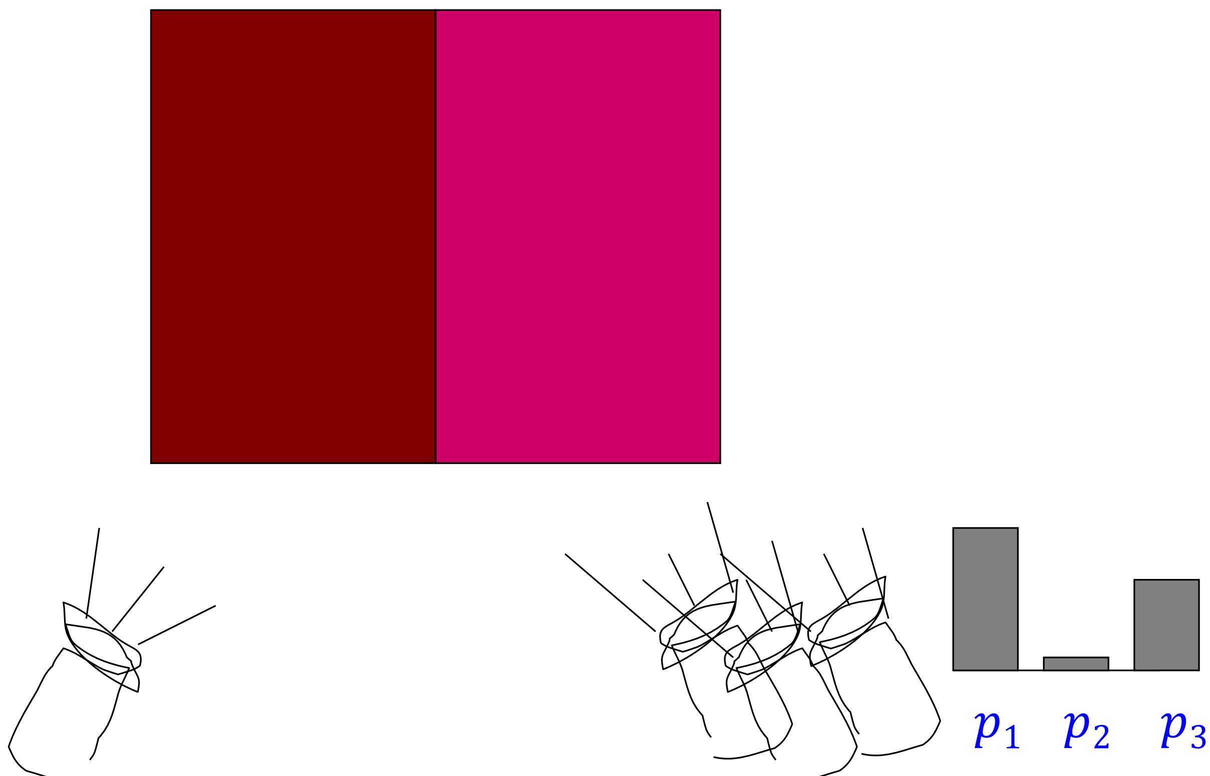

Source: W. Freeman

Color matching experiment 2

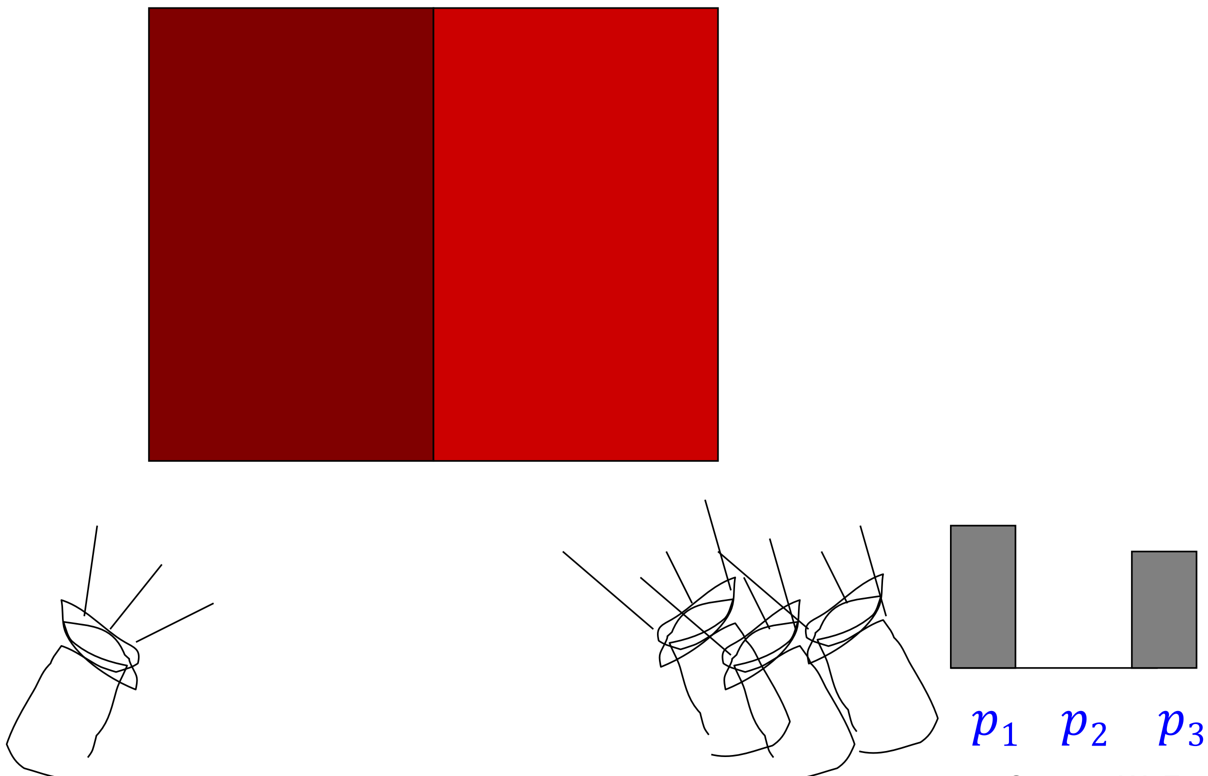

Source: W. Freeman

Color matching experiment 2

We say a “negative” amount of p_2 was needed to make the match, because we added it to the test color’s side.

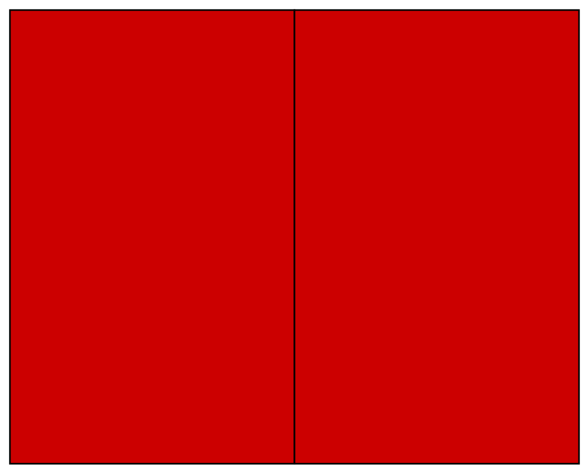

The primary color amounts needed for a match:

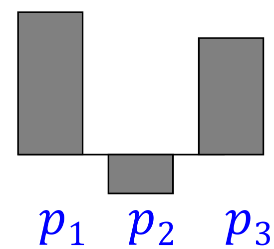

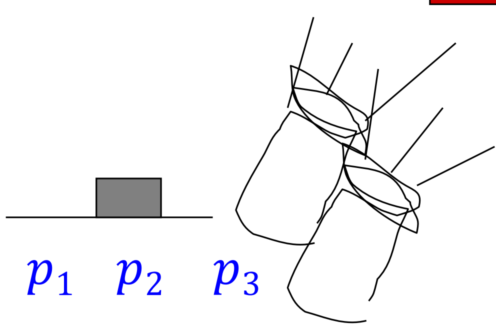

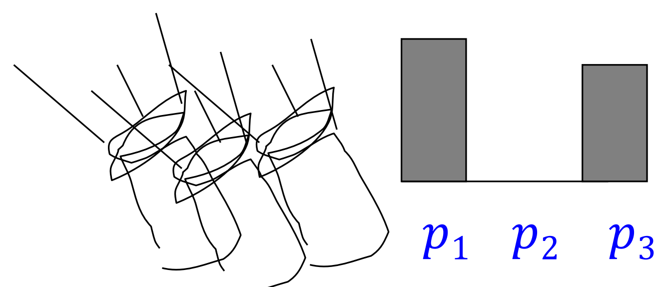

Source: W. Freeman

Empirical properties of color matching

- **Trichromacy:**
  - Most* people can match any given test light with three *independent* primaries
  - For the same light and primaries, most* people select the same weights
  - Thus, three numbers are sufficient for encoding color
  - This observation dates back to Thomas Young in the 18th century

Color matching is linear

- Let a_1, a_2, and a_3 be weights of primaries P_1, P_2, P_3 needed to match a test light A. We write $A = (a_1, a_2, a_3)$.
- Empirically, color matching obeys *Grassman's laws*:
  - If two lights can be matched with the same weights, then they match each other:
    - If $A = (a_1, a_2, a_3)$ and $B = (a_1, a_2, a_3)$, then $A = B$
  - If we mix two lights, then mixing the matches will match the result:
    - If $A = (a_1, a_2, a_3)$ and $B = (b_1, b_2, b_3)$, then

$$A + B = (a_1 + b_1, a_2 + b_2, a_3 + b_3)$$

  - If we scale the test light, then the matches get scaled by the same amount:

$$uA = (ua_1, ua_2, ua_3)$$

If's, and's and but's

- Some people do NOT match according to Grassman's laws
- Various causes
  - Missing a cone type -> dichromats
    - Usually L or M cone is missing
      - » "Red-green color blindness"
      - » About 8% of males; 0.5% of females
    - S cone missing
      - » rare
  - Missing two cone types -> monochromats
    - Rare
  - Missing all cone types -> monochromats
    - Rare
  - All cones present, but still doesn't see color -> central achromatopsia
    - Very rare

If’s, and’s and but’s

- Some people match according to Grassman’s laws
  - But use unusual sets of weights
  - Their cone sensitivities are different from that of others
  - Anomalous trichromats
- All this is tightly linked to genetics
  - Variant cone sensitivities <-> variants in genes encoding cone proteins
  - L, M cone genes are on x-chromosome, hence the gender link
  - Mild evidence suggests there may be females who can match with four primaries
- All this is tightly linked to mortality
  - Color deficient spider monkeys die young, on average
  - Color deficient humans may die younger, on average

Outline

- Physical origin of color
- Spectra of sources and surfaces
- Physiology of color vision
- Quantifying color perception
- Color spaces

Why specify color numerically?

- Accurate color reproduction is commercially valuable
  - e.g. Kodak yellow, painting a house.
- Of the order of 10 color names are widely recognized by English speakers - other languages have fewer/more, but not much more.
  - There's a great deal of structure to the way colour is spoken about (Berlin-Kay), but not much precision
- Color reproduction problems increased by prevalence of digital imaging - eg. digital libraries of art.
  - Choosing pixel values to reproduce/evoke experiences, e.g. an architectural model.
  - Consistency in user interfaces, monitor-printer consistency, monitor-lino consistency, etc.

Linear color spaces

- Fixing three **primaries** defines a linear color space in which the **coordinates** of a color are given by the weights of the primaries used to match it
- How can we find the coordinates of an arbitrary color signal?
- We need **color matching functions**, or amounts of each primary needed to match monochromatic sources at each wavelength

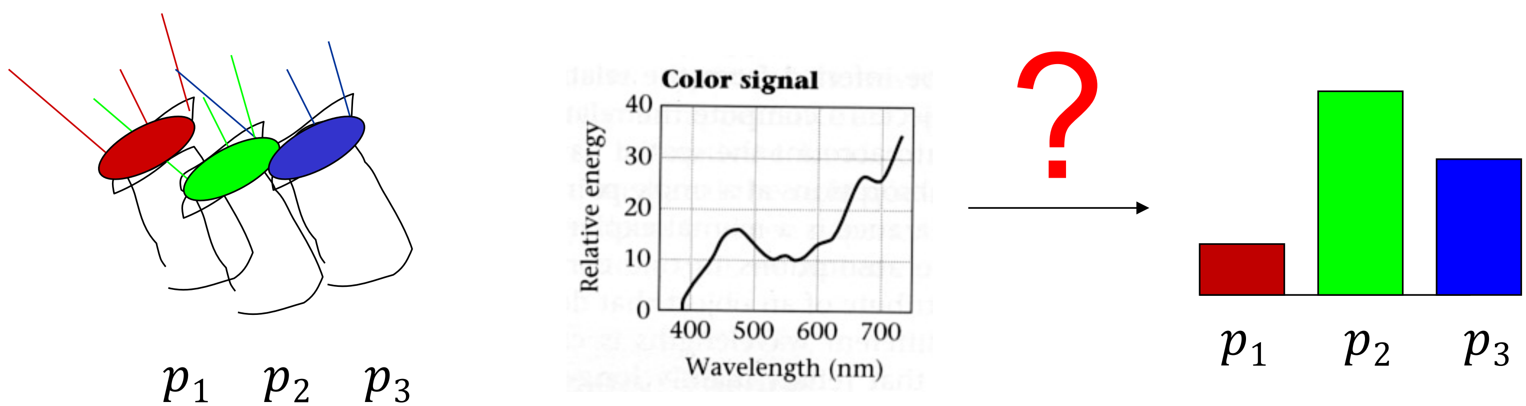

Linear color spaces

- Let $t(\lambda)$ be the spectrum of the target signal and $c_1(\lambda)$, $c_2(\lambda)$, and $c_3(\lambda)$ the color matching functions

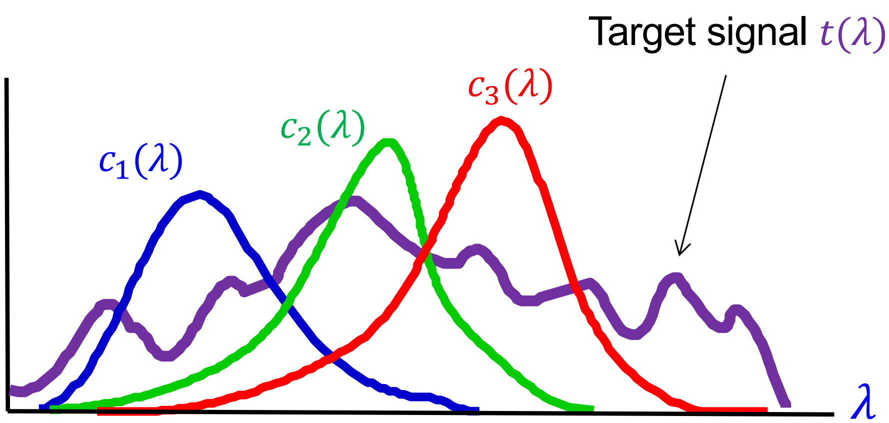

Linear color spaces

- Let $t(\lambda)$ be the spectrum of the target signal and $c_1(\lambda)$, $c_2(\lambda)$, and $c_3(\lambda)$ the color matching functions
- Then the weights of the primaries needed to match $t(\lambda)$ are:

$$w_1 = \int_{\lambda} t(\lambda) c_1(\lambda) d\lambda$$

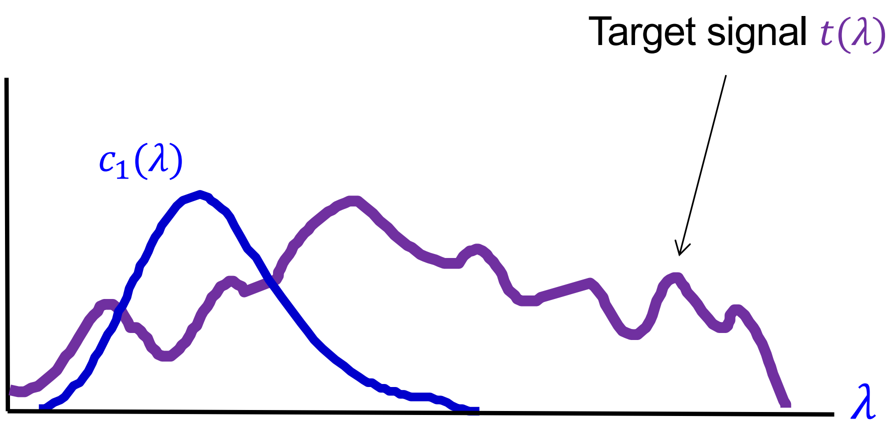

Linear color spaces

- Let $t(\lambda)$ be the spectrum of the target signal and $c_1(\lambda)$, $c_2(\lambda)$, and $c_3(\lambda)$ the color matching functions
- Then the weights of the primaries needed to match $t(\lambda)$ are:

$$w_1 = \int_\lambda t(\lambda)c_1(\lambda)d\lambda \qquad w_2 = \int_\lambda t(\lambda)c_2(\lambda)d\lambda$$

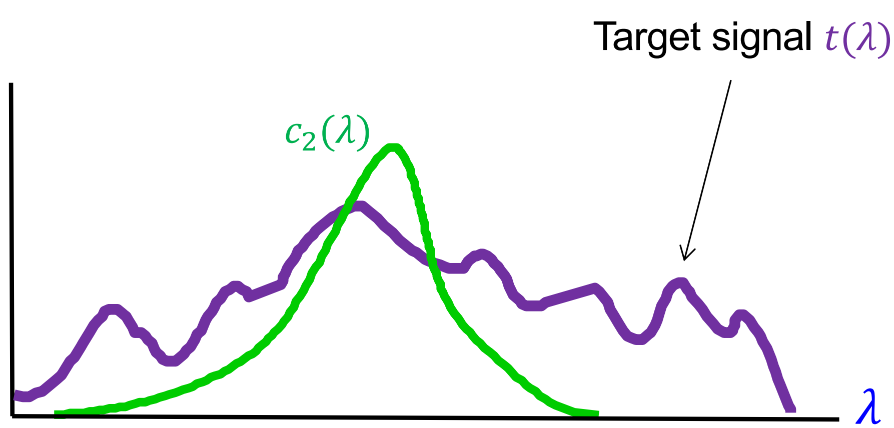

Linear color spaces

- Let $t(\lambda)$ be the spectrum of the target signal and $c_1(\lambda)$, $c_2(\lambda)$, and $c_3(\lambda)$ the color matching functions
- Then the weights of the primaries needed to match $t(\lambda)$ are:

$$w_1 = \int_\lambda t(\lambda)c_1(\lambda)d\lambda \qquad w_2 = \int_\lambda t(\lambda)c_2(\lambda)d\lambda \qquad w_3 = \int_\lambda t(\lambda)c_3(\lambda)d\lambda$$

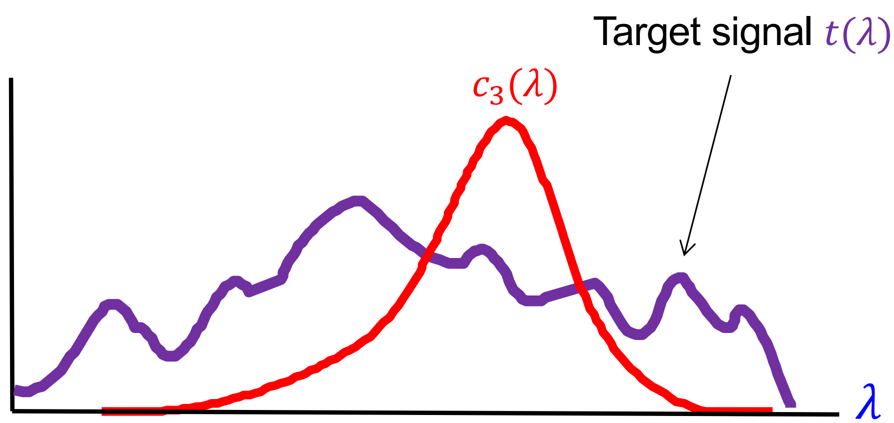

Matching functions act as filters on the target spectrum, like response curves of color receptors!

RGB color space

- Primaries are single-wavelength sources, matching functions for R and G have negative values for parts of the spectrum

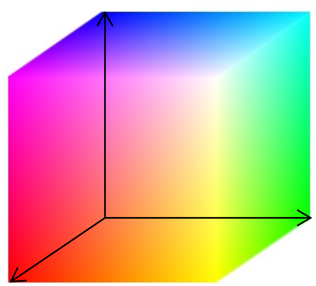

RGB primaries

$p_1 = 645.2$ nm

$p_2 = 525.3$ nm

$p_3 = 444.4$ nm

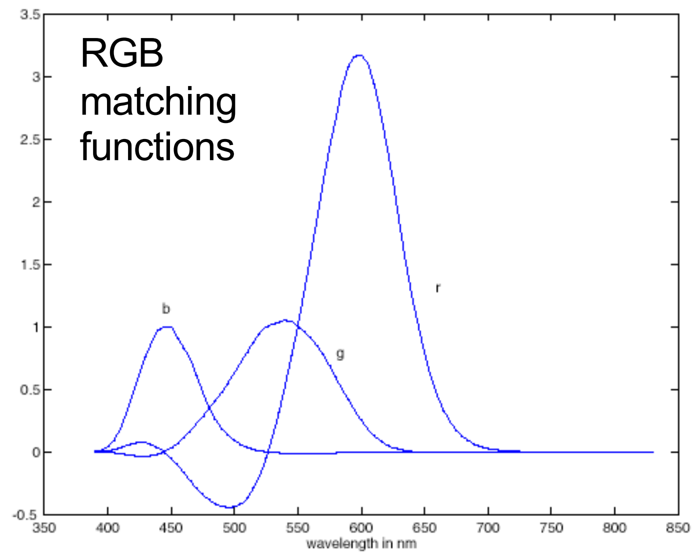

Comparison of RGB matching functions with best 3x3 linear transformation of cone responses

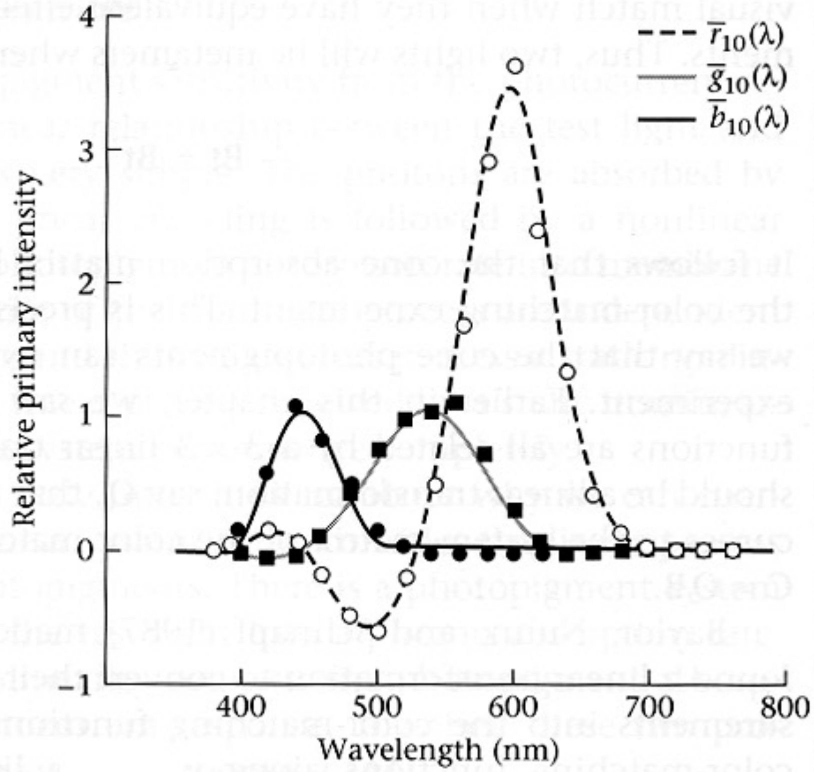

4.20 COMPARISON OF CONE PHOTOCURRENT RESPONSES AND THE COLOR-MATCHING FUNCTIONS. The cone photocurrent spectral responsivities are within a linear transformation of the color-matching functions, after a correction has been made for the optics and inert pigments in the eye. The smooth curves show the Stiles and Burch (1959) color-matching functions. The symbols show the matches predicted from the photocurrents of the three types of macaque cones. The predictions included a correction for absorption by the lens and other inert pigments in the eye. Source: Baylor, 1987.

Wandell, Foundations of Vision, 1995

Linear color spaces: CIE XYZ

- Primaries are *imaginary*, but matching functions are everywhere positive
- The Y parameter corresponds to brightness or *luminance*

Matching functions

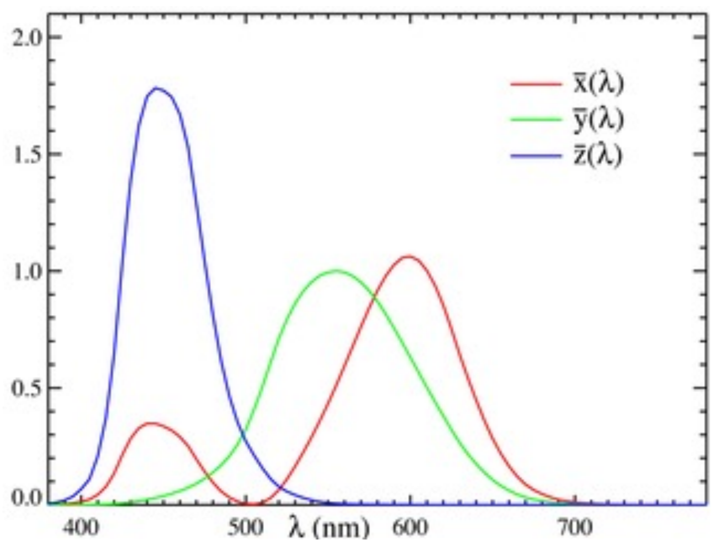

2D visualization:
$x = X/(X+Y+Z),\ y = Y/(X+Y+Z)$

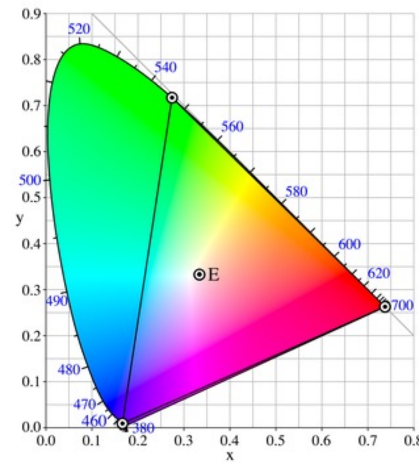

http://en.wikipedia.org/wiki/CIE_1931_color_space

Linear color spaces: CIE XYZ

- CIE XYZ is based on color matching experiments carried out in late 1920s by W. David Wright (Imperial College) and John Guild (National Physical Laboratory, London)
- The experiments used 17 “standard observers” (10 by Wright, 7 by Guild)

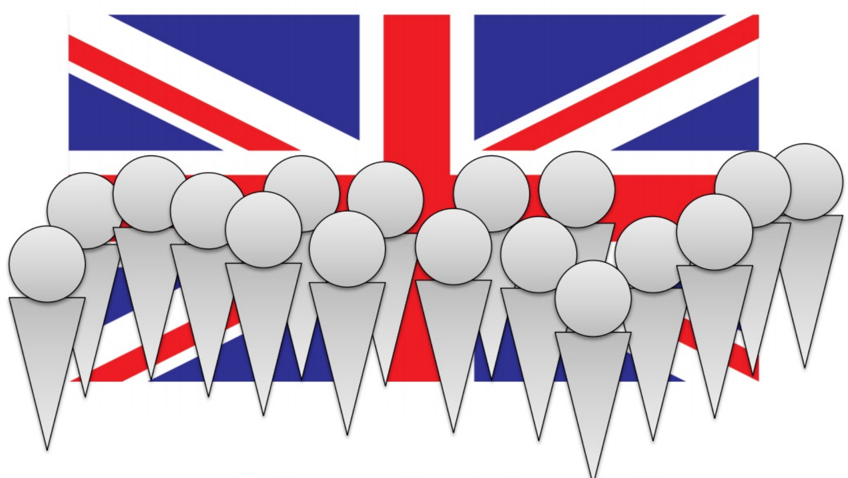

“The Standard Observers”

Source: M. Brown

Uniform color spaces

- Unfortunately, differences in x, y coordinates do not reflect perceptual color differences
- CIE $u'v'$ is a projective transform of x, y to make the ellipses more uniform

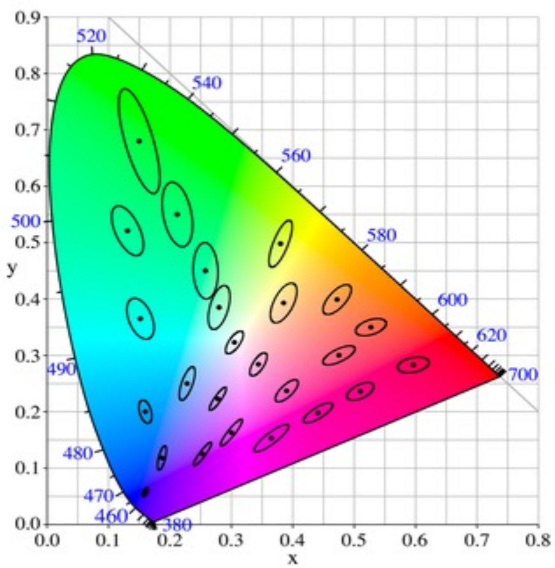

McAdam ellipses: Just noticeable differences in color

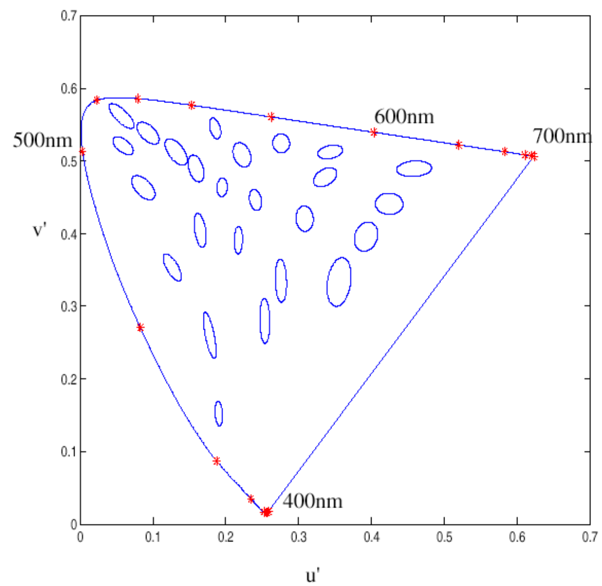

Nonlinear color spaces: HSV

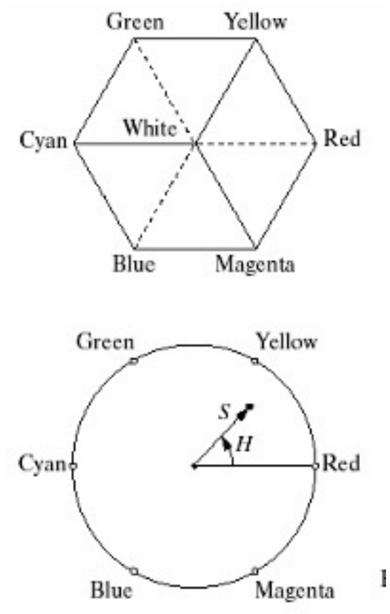

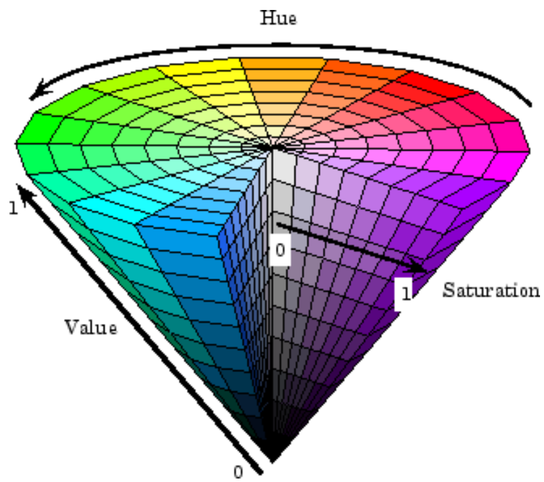

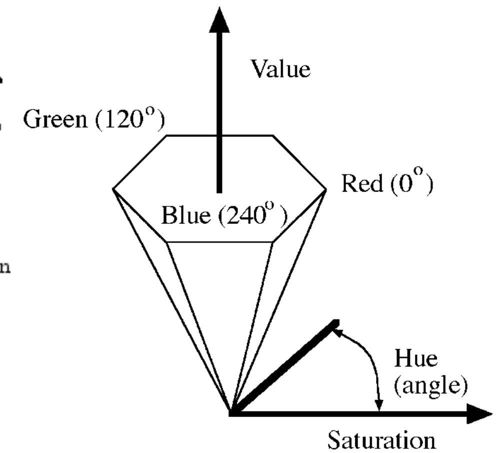

- Perceptually meaningful dimensions: Hue, Saturation, Value (Intensity)
- RGB cube on its vertex

Outline

- Physical origin of color
- Spectra of sources and surfaces
- Physiology of color vision
- Quantifying color perception
- Color spaces
- Color constancy, white balance

Color perception

- Color/lightness constancy
  - The ability of the human visual system to perceive the intrinsic reflectance properties of the surfaces despite changes in illumination conditions

J. S. Sargent, The Daughters of Edward D. Boit, 1882

Chromatic adaptation

- The visual system changes its sensitivity depending on the luminances prevailing in the visual field
  - The exact mechanism is still not fully understood
- Adapting to different brightness levels
  - Changing the size of the iris opening (i.e., the aperture) changes the amount of light that can enter the eye
  - Think of walking into a building from full sunshine
- Adapting to different color temperature
  - The receptive cells on the retina change their sensitivity
  - For example: if there is an increased amount of red light, the cells receptive to red decrease their sensitivity until the scene looks white again
  - We actually adapt better in brighter scenes: this is why candlelit scenes still look yellow

http://www.schorsch.com/kbase/glossary/adaptation.html

Checker shadow illusion

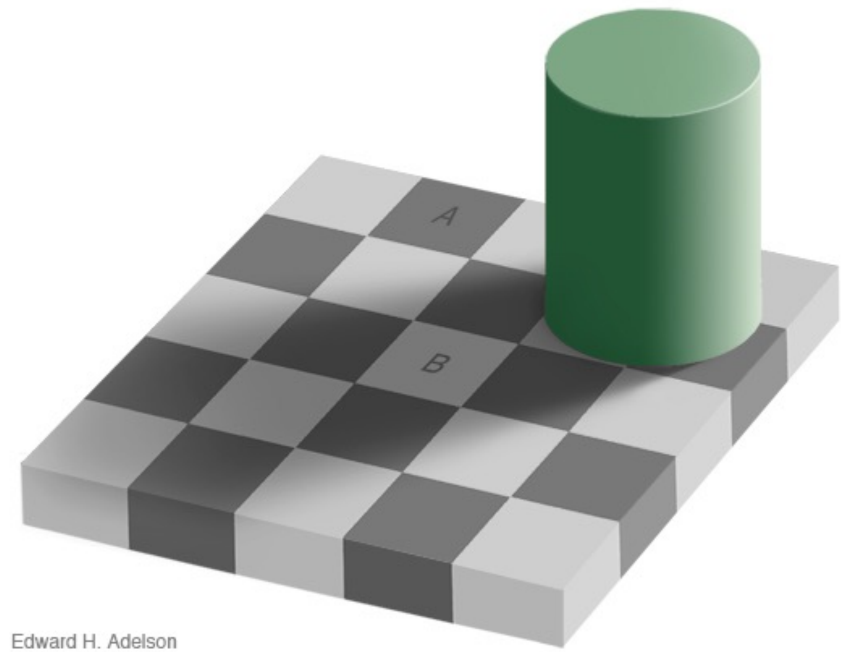

https://en.wikipedia.org/wiki/Checker_shadow_illusion

Checker shadow illusion

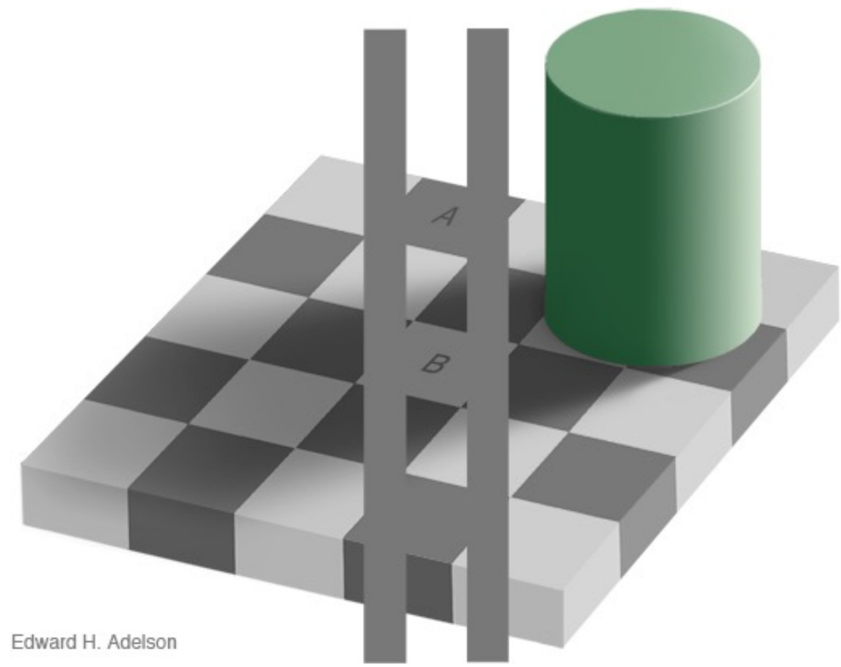

- Possible explanations: simultaneous contrast, reflectance vs. illumination edges

https://en.wikipedia.org/wiki/Checker_shadow_illusion

This strawberry cake has no red pixels!

https://www.digitaltrends.com/photography/non-red-strawberries/

Lightness Constancy

- Lightness constancy
  - how light is the surface, independent of the brightness of the illuminant
  - issues
    - spatial variation in illumination
    - absolute standard
  - Human lightness constancy is very good
- Assume
  - frontal 1D “Surface”
  - slowly varying illumination
  - quickly varying surface reflectance

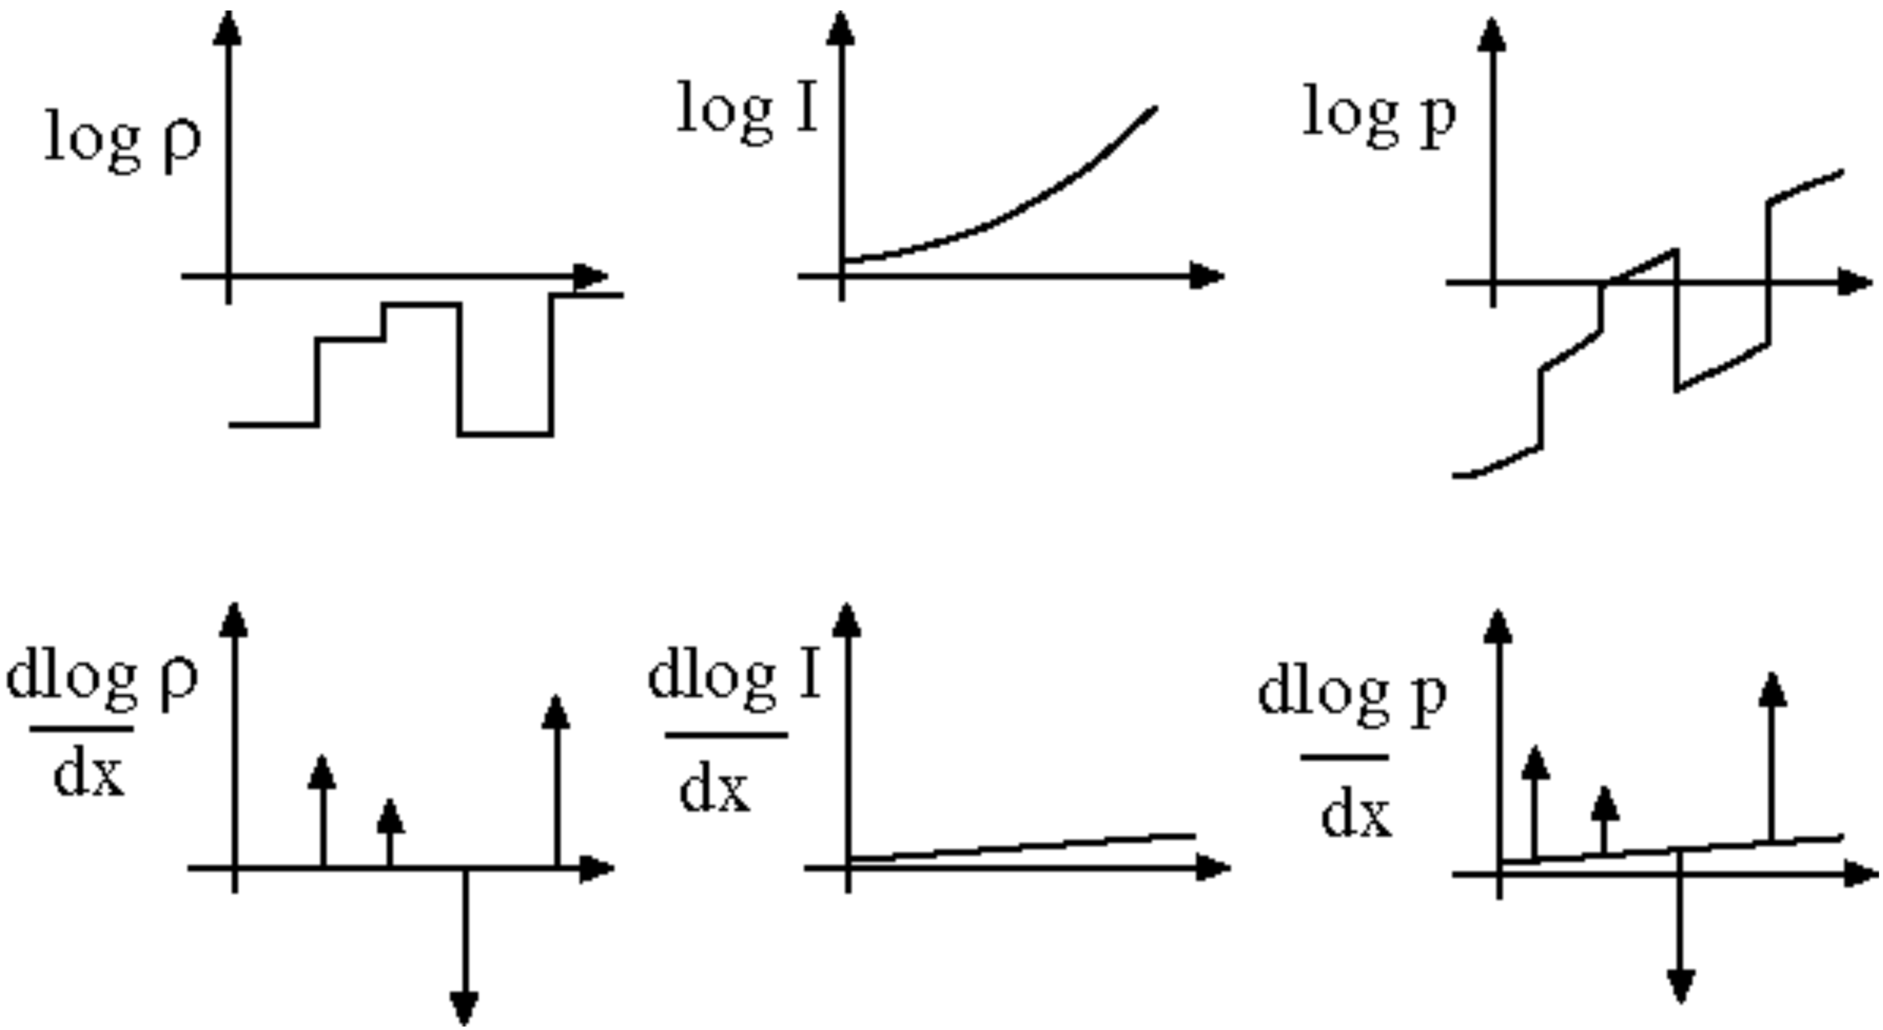

log ρ
log I
log p
dlog ρ / dx
dlog I / dx
dlog p / dx

Thresholded $\frac{d\log p}{dx}$

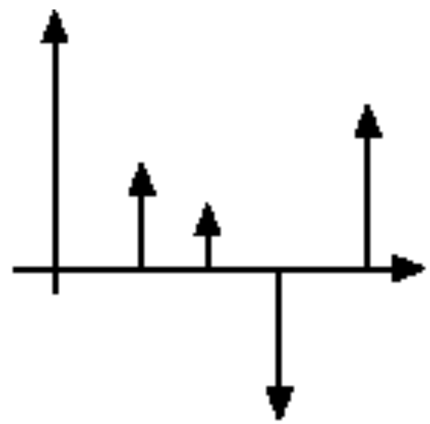

Integrate
This to get

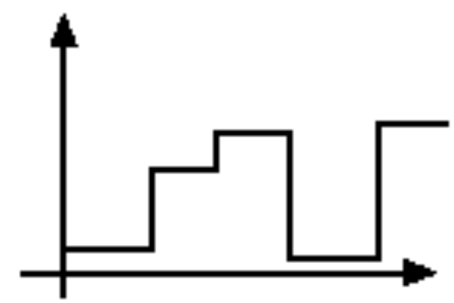

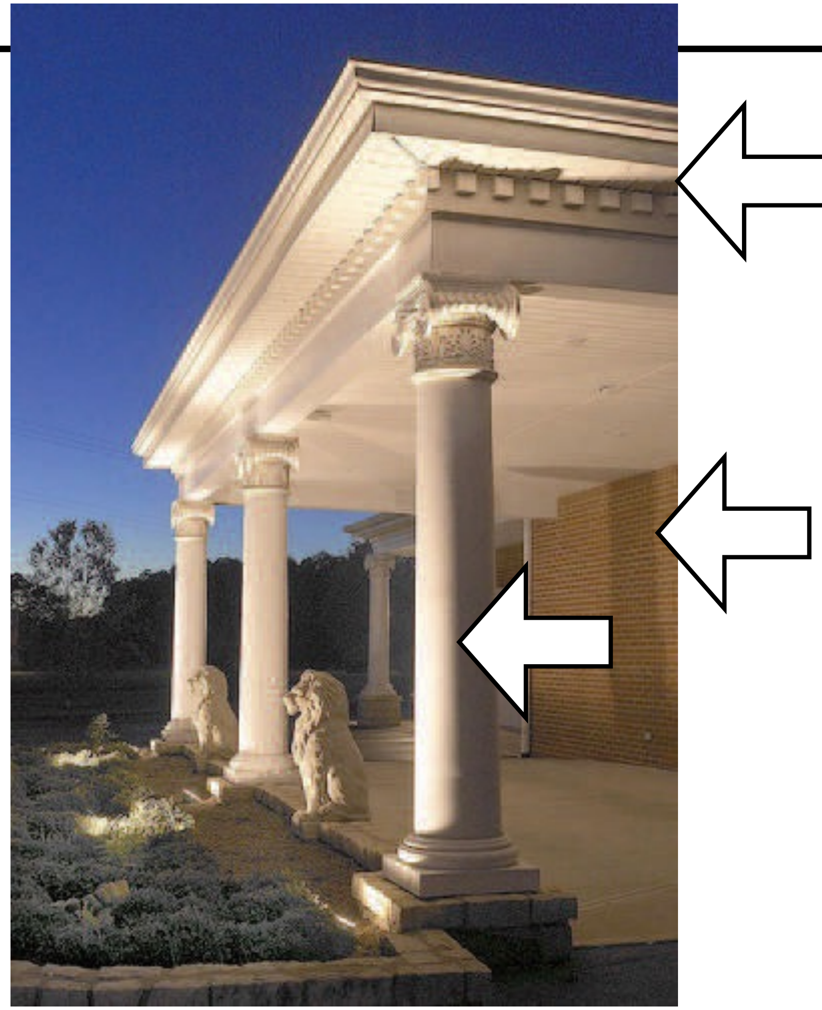

Stage lighting

White balance

- Analogous to color constancy mechanisms in human vision, cameras have mechanisms to adapt to the illumination in the environment so that neutral (white or gray) objects look neutral

Incorrect white balance

Correct white balance

http://www.cambridgeincolour.com/tutorials/white-balance.htm

White balance

- Film cameras:
  - Different types of film or different filters for different illumination conditions

- Digital cameras:
  - Automatic white balance
  - White balance settings corresponding to several common illuminants
  - Custom white balance using a reference object

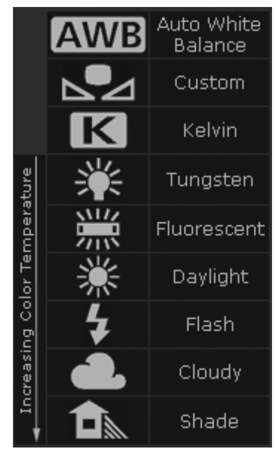

http://www.cambridgeincolour.com/tutorials/white-balance.htm

White balance

- Von Kries adaptation: Multiply each channel by a *gain factor*

- Best way: gray card
  - Take a picture of a neutral object (white or gray)
  - If the object is recorded as r_w, g_w, b_w use weights $1/r_w, 1/g_w, 1/b_w$

White balance

- Without gray cards: we need to “guess” which pixels correspond to white objects
- Gray world assumption
  - The image average $\bar{r}, \bar{g}, \bar{b}$ is gray
  - Use weights $1/\bar{r}, 1/\bar{g}, 1/\bar{b}$
- Brightest pixel assumption
  - Highlights usually have the color of the light source
  - Use weights inversely proportional to the values of the brightest pixels
- Gamut mapping
  - Gamut: convex hull of all pixel colors in an image
  - Find the transformation that matches the gamut of the image to the gamut of a “typical” image under white light
- Use image statistics, learning techniques

Is white balance solved?

New College, Oxford, September 2022

Is white balance solved?

Photographers in cities like San Francisco and Portland have been sharing apocalyptic images of red/orange skies as wildfire smoke literally blots out the sun. But many smartphone photographers trying to do the same thing have tried and failed over and over. It turns out Auto White Balance is ruining their shots.

https://petapixel.com/2020/09/10/smartphones-are-ruining-wildfire-sky-photos-with-auto-white-balance/

Is white balance solved?

- When there are several types of illuminants in the scene, different reference points will yield different results

Reference: moon

Reference: stone

http://www.cambridgeincolour.com/tutorials/white-balance.htm

Is white balance solved?

- When there are several types of illuminants in the scene, different reference points will yield different results
- Possible solution: spatially varying white balance

Input

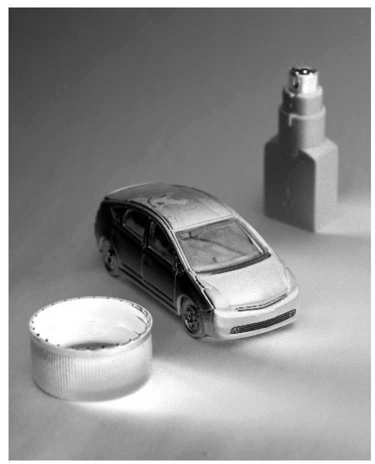

Alpha map

Output

E. Hsu, T. Mertens, S. Paris, S. Avidan, and F. Durand, Light Mixture Estimation for Spatially Varying White Balance, SIGGRAPH 2008

Uses of color in computer vision