C280, Computer Vision

Prof. Trevor Darrell
trevor@eecs.berkeley.edu

Lecture 3: Color

Color

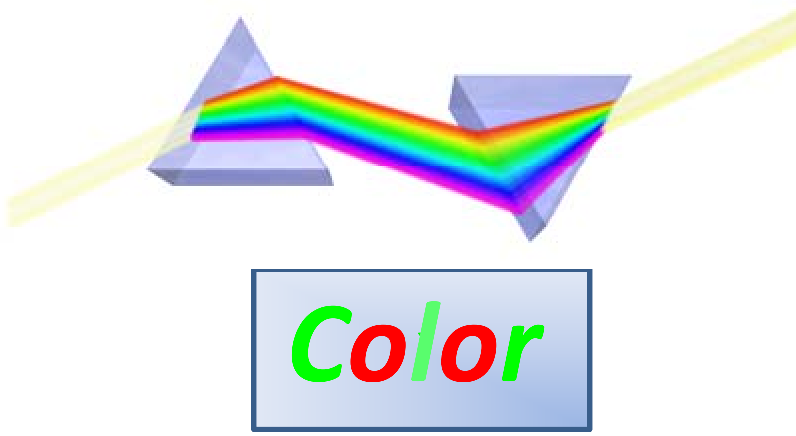

Readings:

- Forsyth and Ponce, Chapter 6
- Szeliski, 2.3.2

Last time

- Image formation affected by geometry, photometry, and optics.
- Projection equations express how world points mapped to 2d image.
- Homogenous coordinates allow linear system for projection equations.
- Lenses make pinhole model practical
- Photometry models: Lambertian, BRDF
- Digital imagers, Bayer demosaicing

Parameters (focal length, aperture, lens diameter, sensor sampling…) strongly affect image obtained.

K. Grauman

Slide Credits

- Kristen Grauman: 3-48, 50-75, 79-86
- Bob Woodham: 49, 87-90
- and others, indirectly (Steve Palmer, Brian Wandell, etc!)

Today: Color

- Measuring color
  - Spectral power distributions
  - Color mixing
  - Color matching experiments
  - Color spaces
    - Uniform color spaces

- Perception of color
  - Human photoreceptors
  - Environmental effects, adaptation

- Using color in machine vision systems

Color and light

- **Color of light** arriving at camera depends on
  – Spectral reflectance of the surface light is leaving
  – Spectral radiance of light falling on that patch

- **Color perceived** depends on
  – Physics of light
  – Visual system receptors
  – Brain processing, environment

Color and light

White light: composed of about equal energy in all wavelengths of the visible spectrum

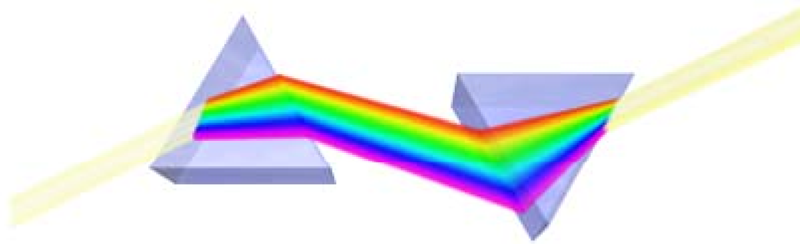

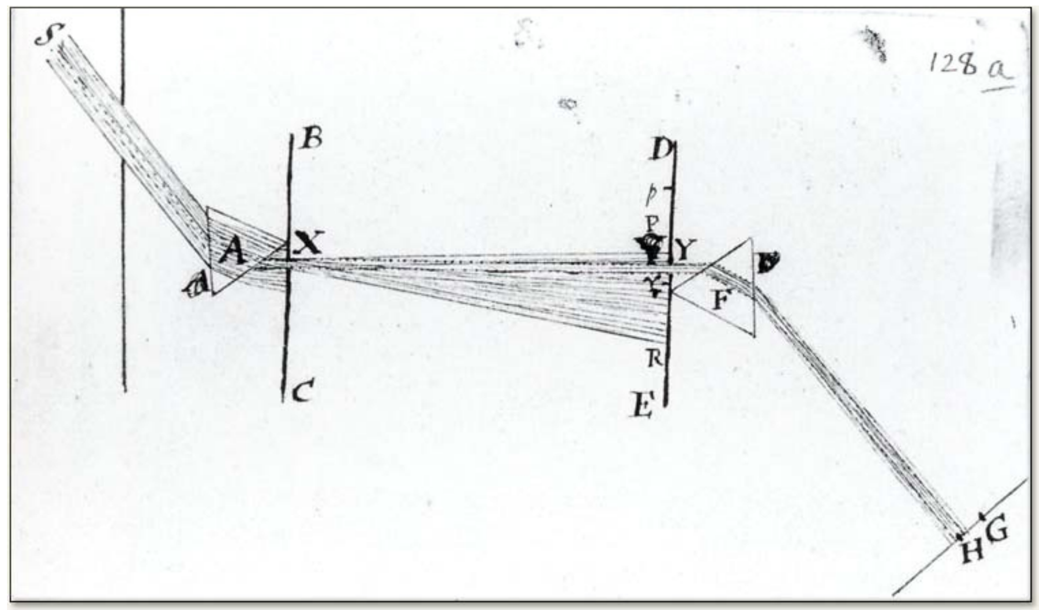

Newton 1665

Image from http://micro.magnet.fsu.edu/

Electromagnetic spectrum

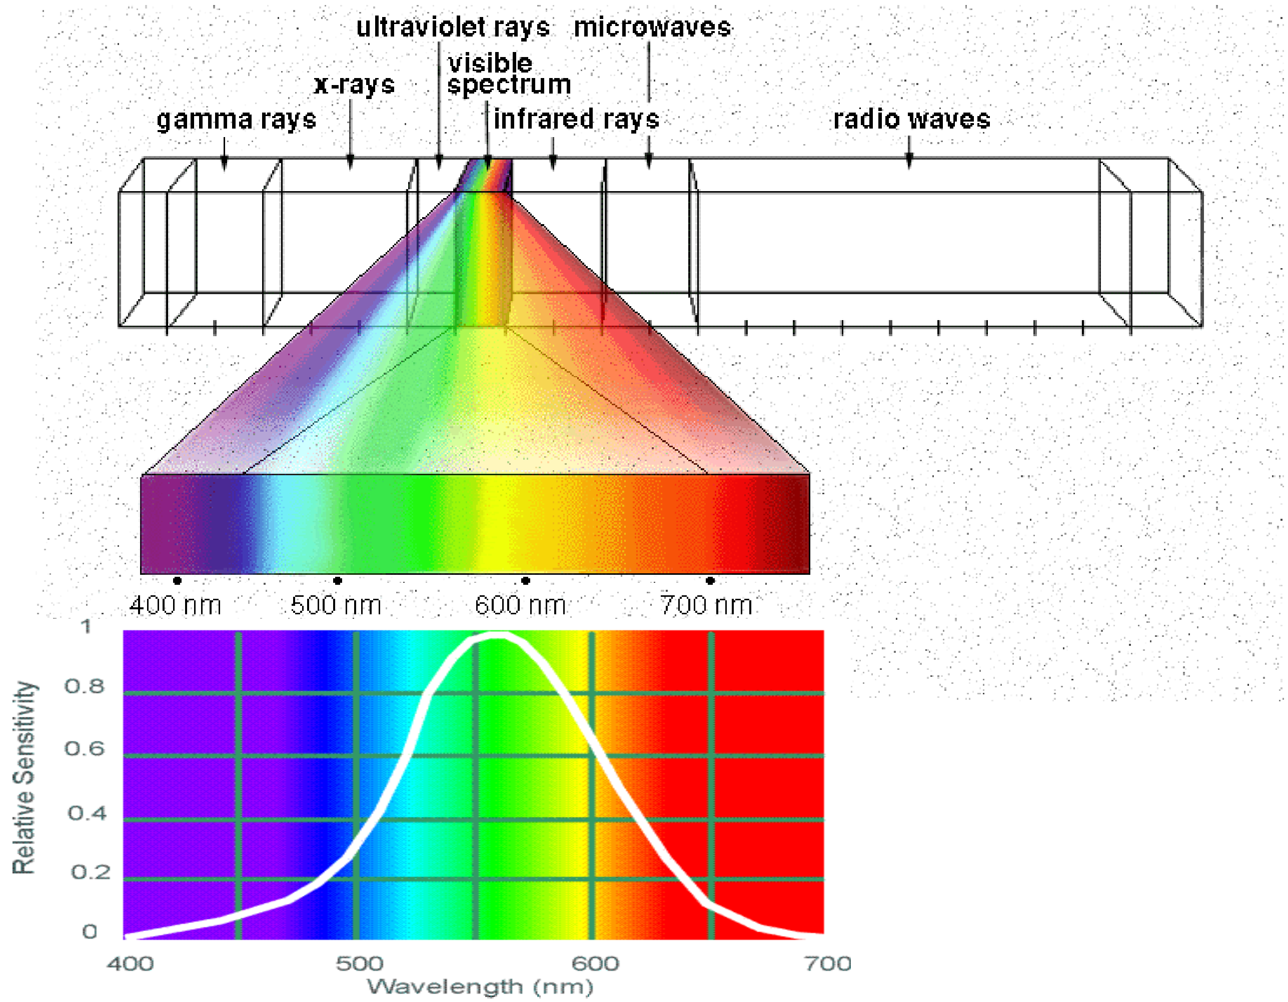

Human Luminance Sensitivity Function

Image credit: nasa.gov

Measuring spectra

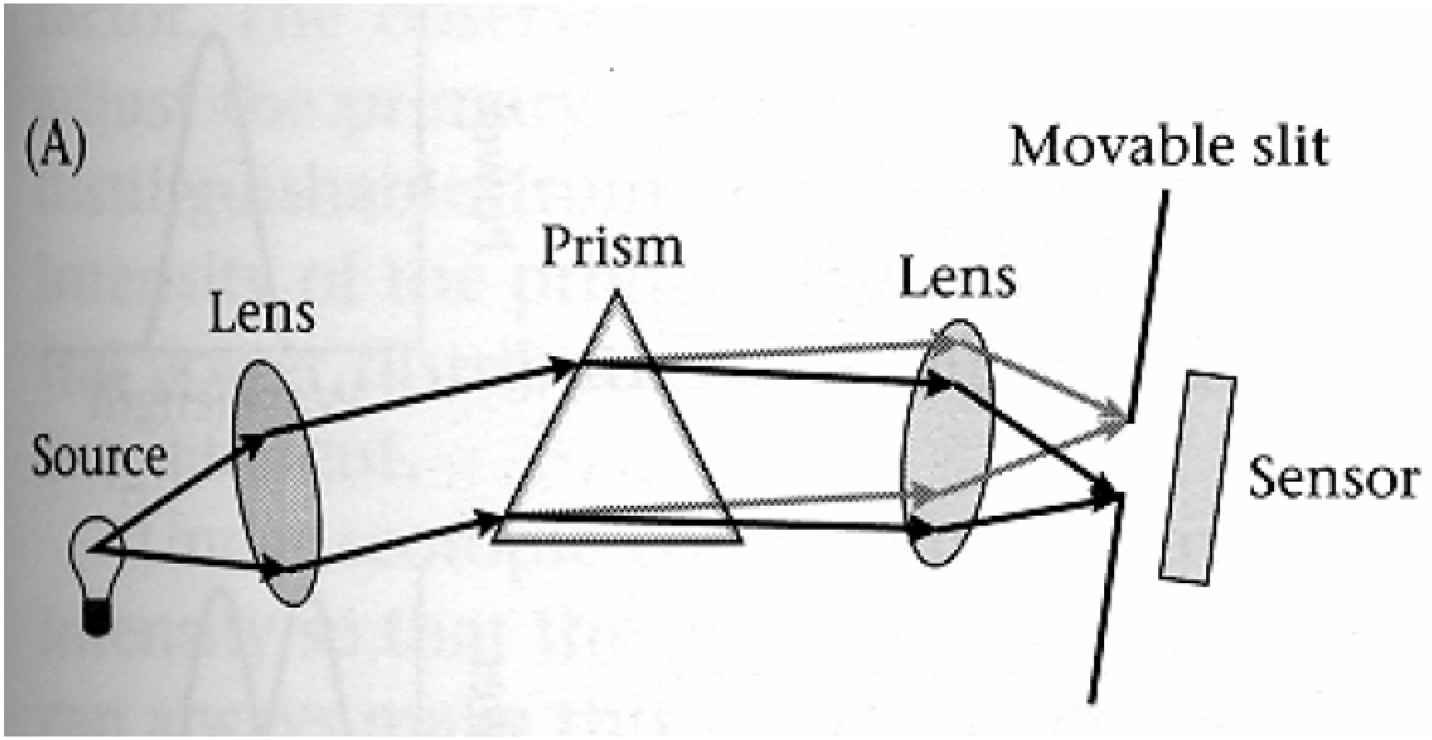

Spectroradiometer: separate input light into its different wavelengths, and measure the energy at each.

Foundations of Vision, B. Wandell

Spectral power distribution

- The power per unit area at each wavelength of a radiant object

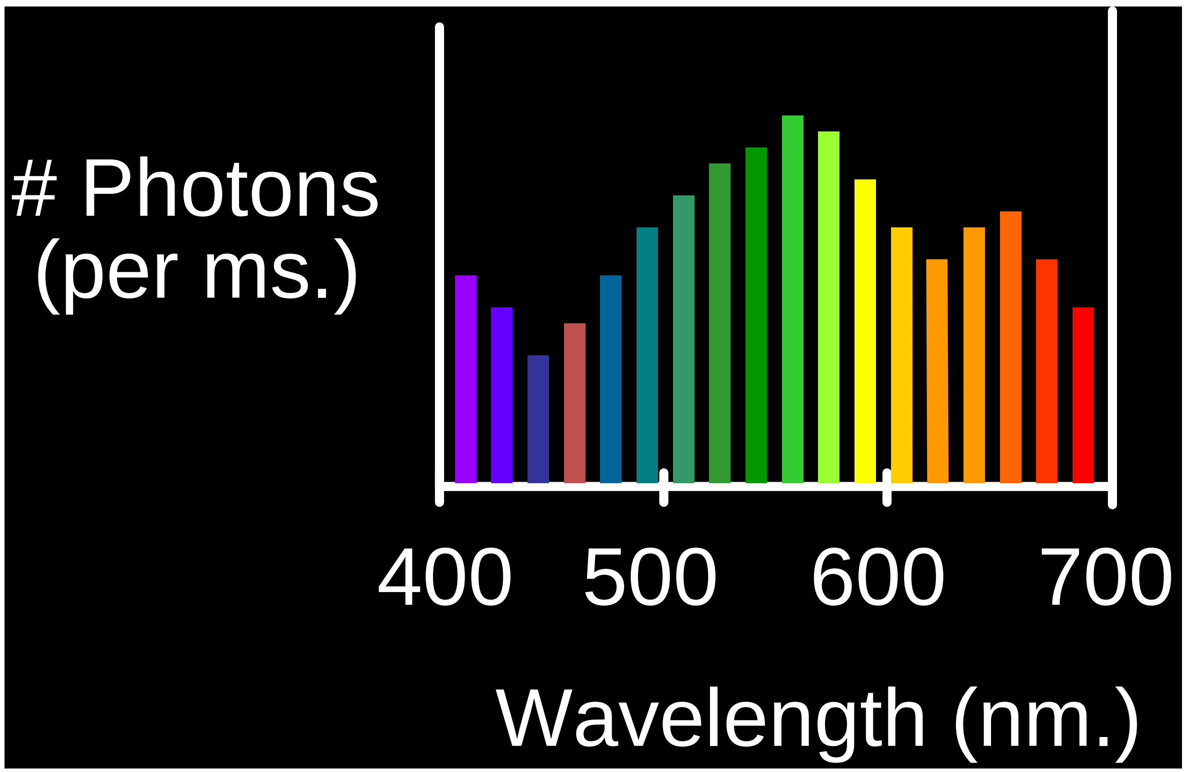

Spectral power distributions

Some examples of the spectra of light sources

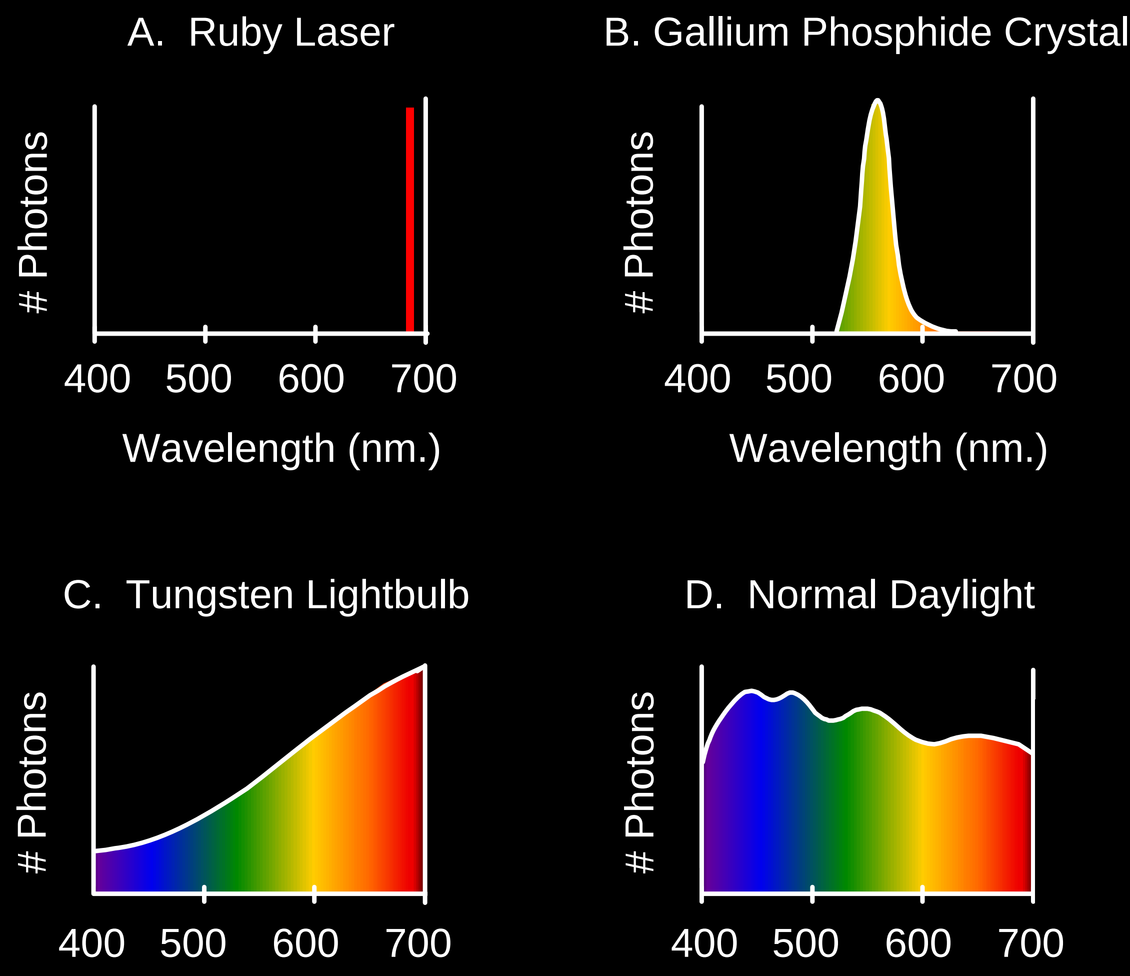

The color viewed is also affected by the surface's spectral reflectance properties.

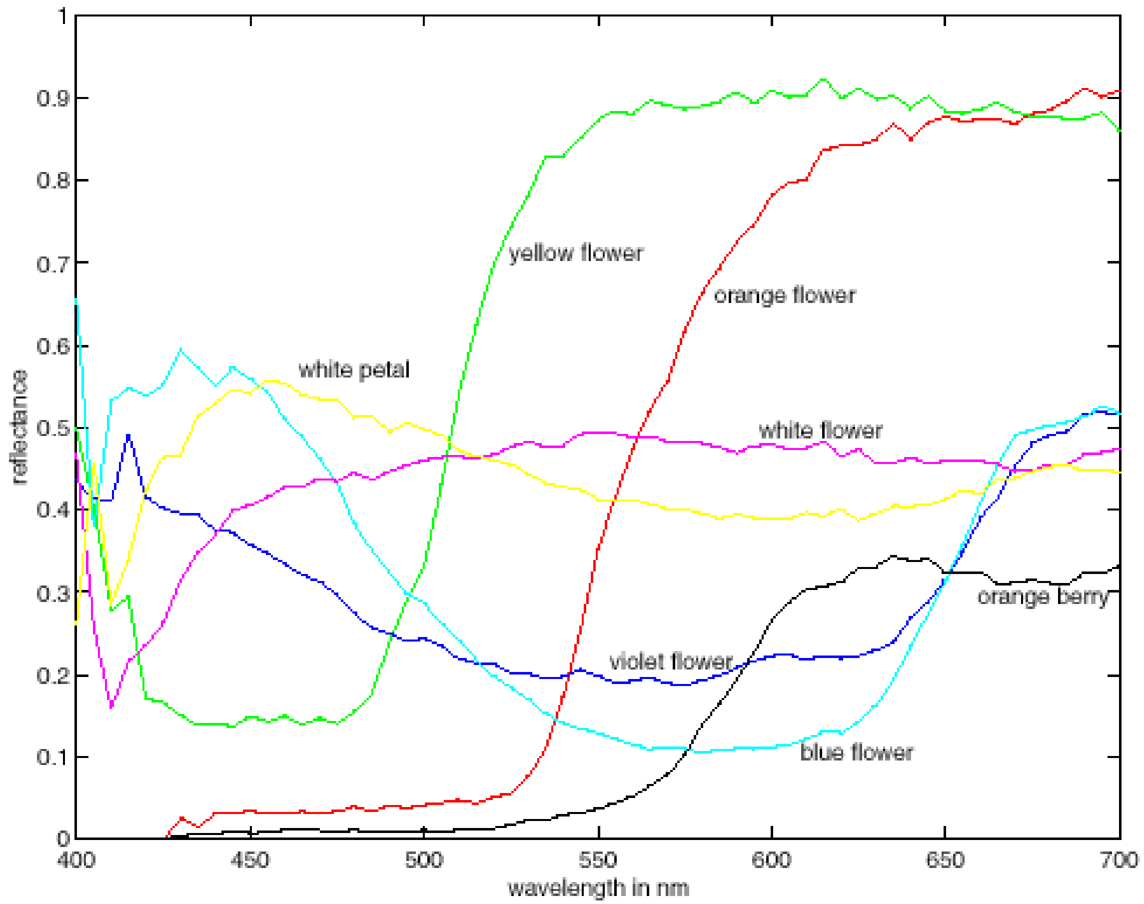

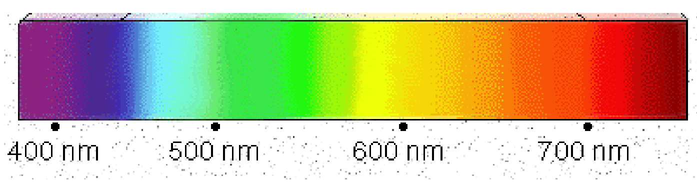

Spectral reflectances for some natural objects: how much of each wavelength is reflected for that surface

Forsyth & Ponce, measurements by E. Koivisto

Surface reflectance spectra

Some examples of the reflectance spectra of surfaces

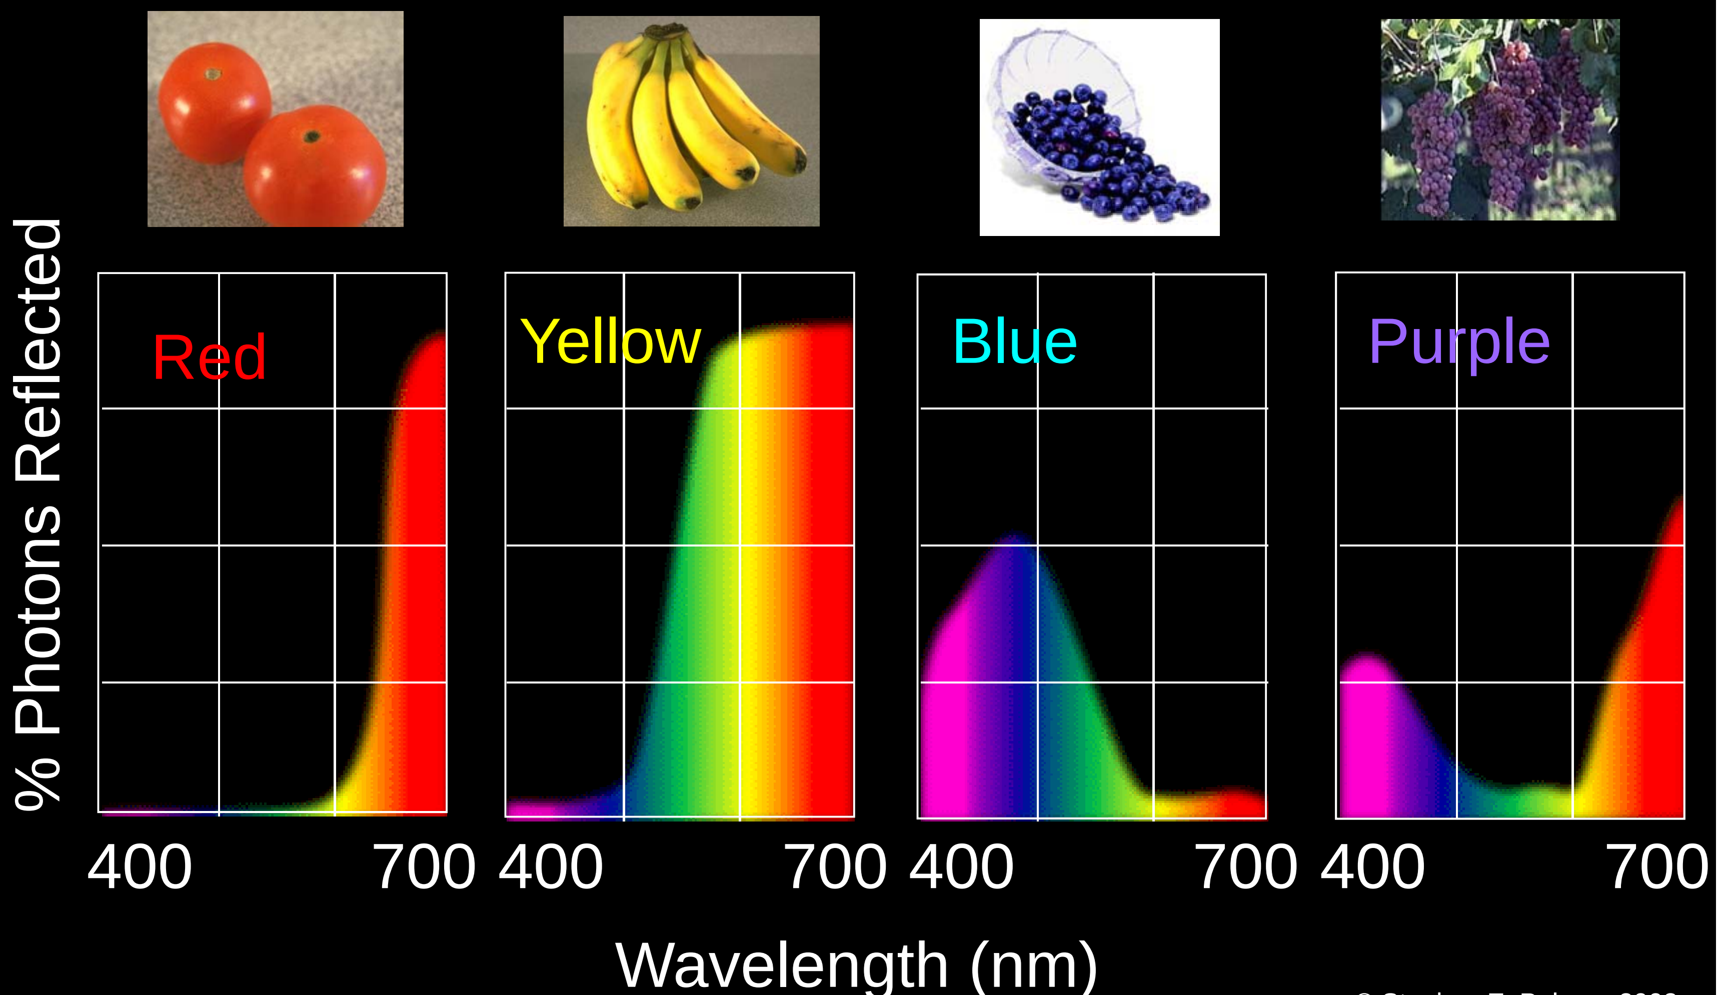

The Psychophysical Correspondence

There is no simple functional description for the perceived color of all lights under all viewing conditions, but ……

A helpful constraint:

Consider only physical spectra with normal distributions

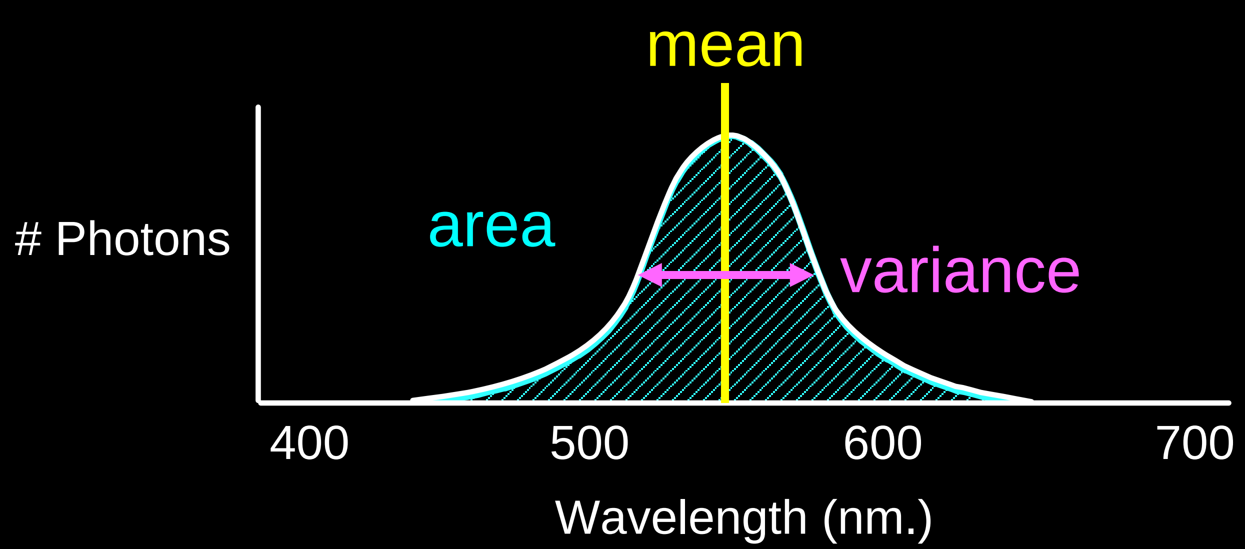

The Psychophysical Correspondence

Mean ⟷ Hue

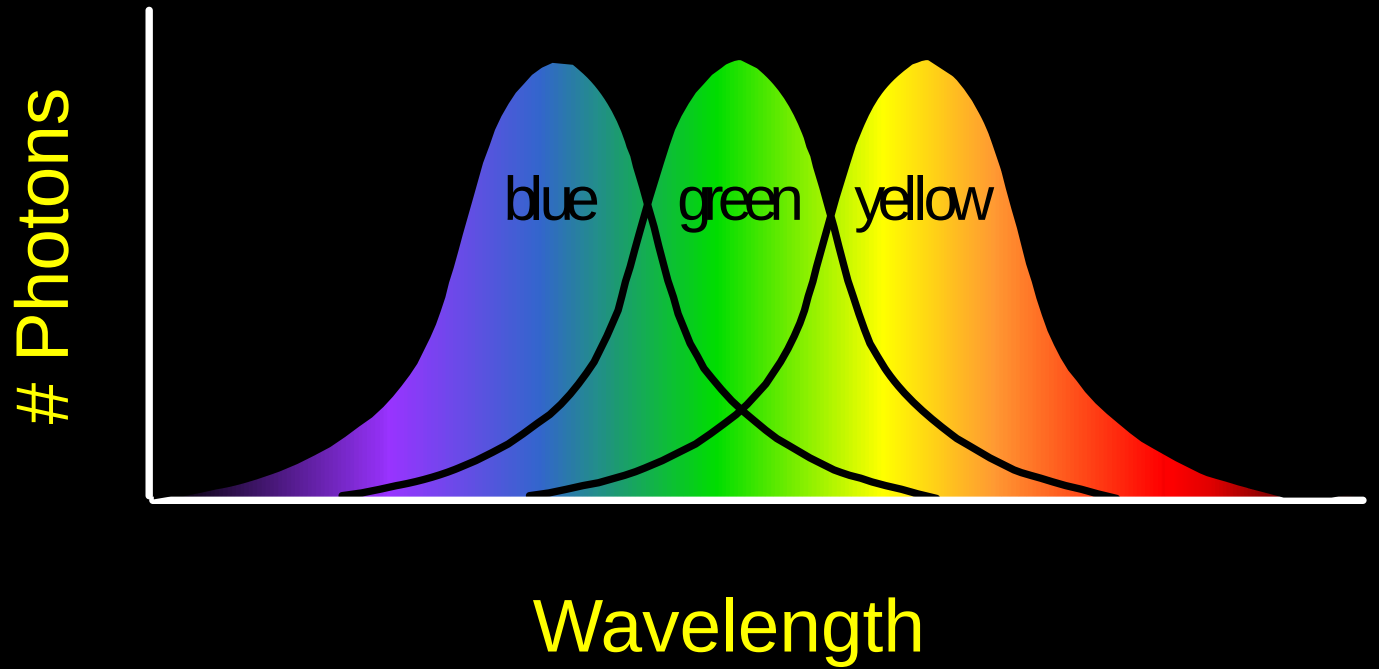

The Psychophysical Correspondence

Variance ⟷ Saturation

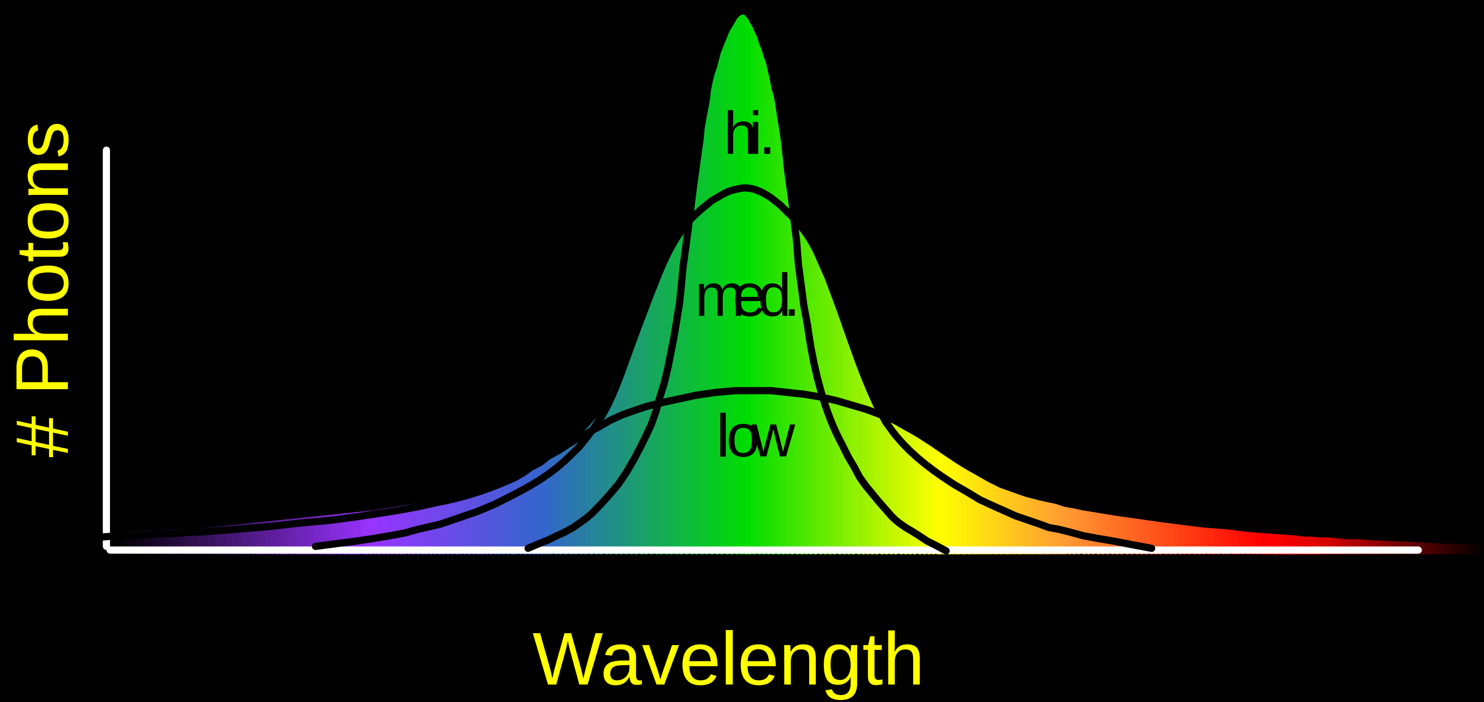

The Psychophysical Correspondence

Area ⟷ Brightness

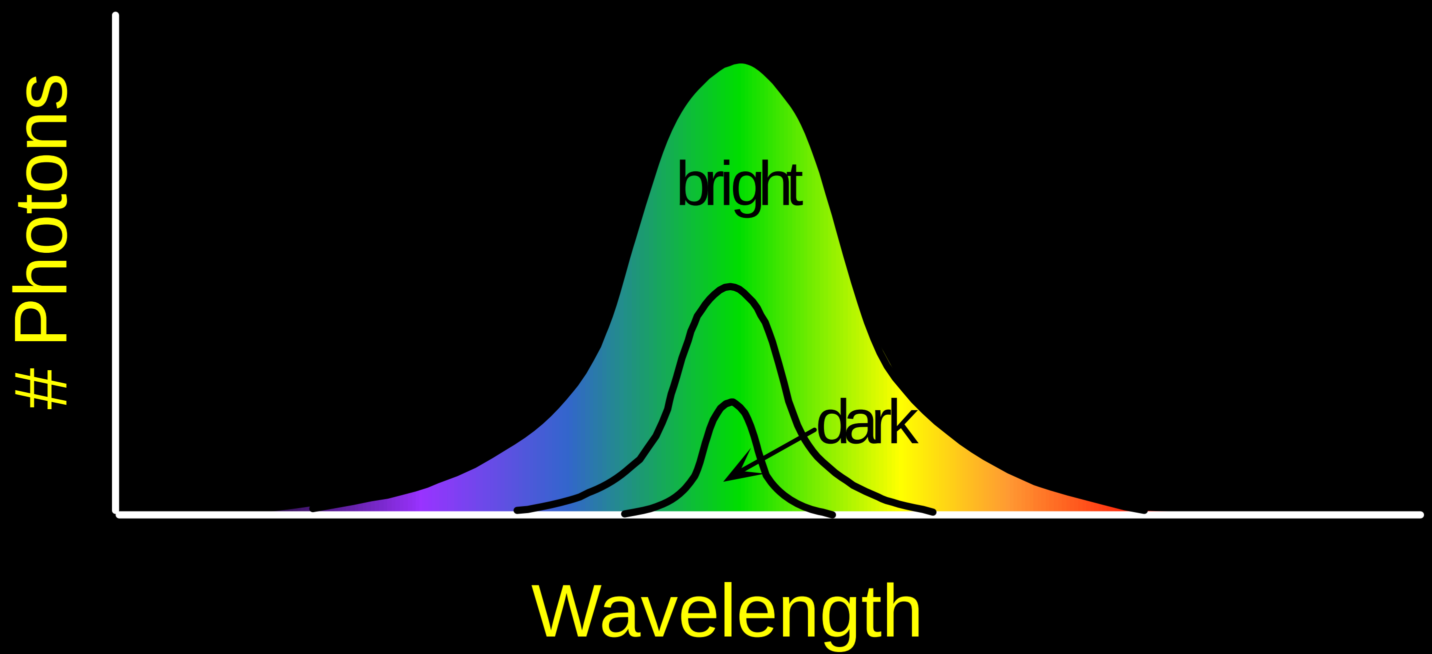

Color mixing

Cartoon spectra for color names:

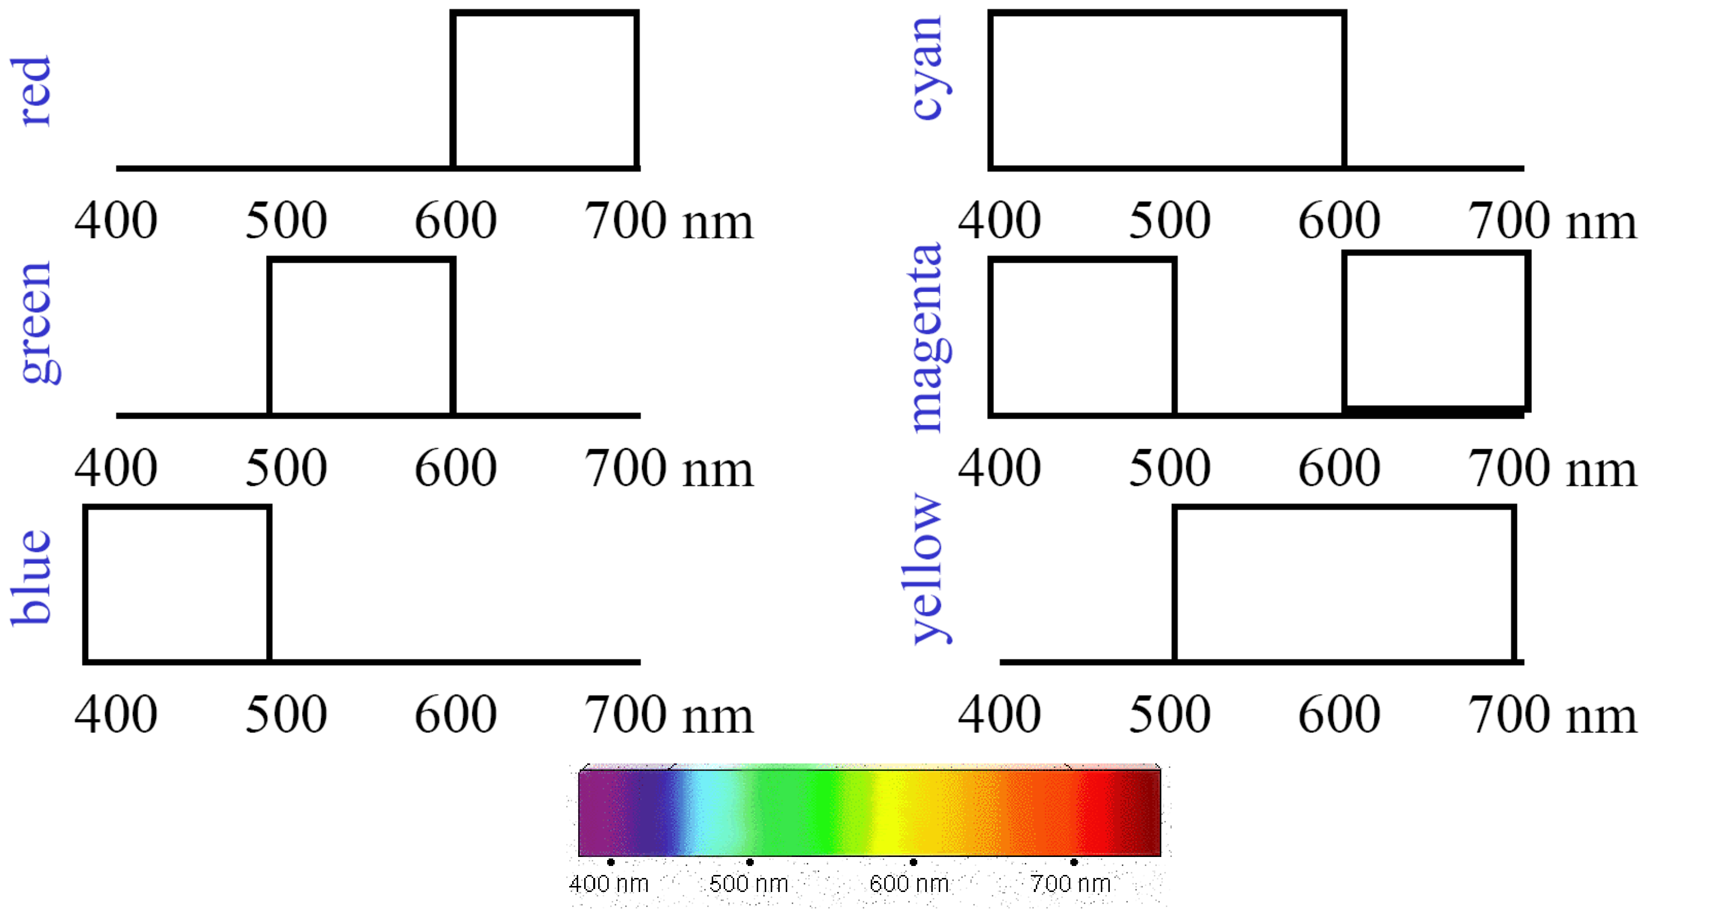

Source: W. Freeman

Additive color mixing

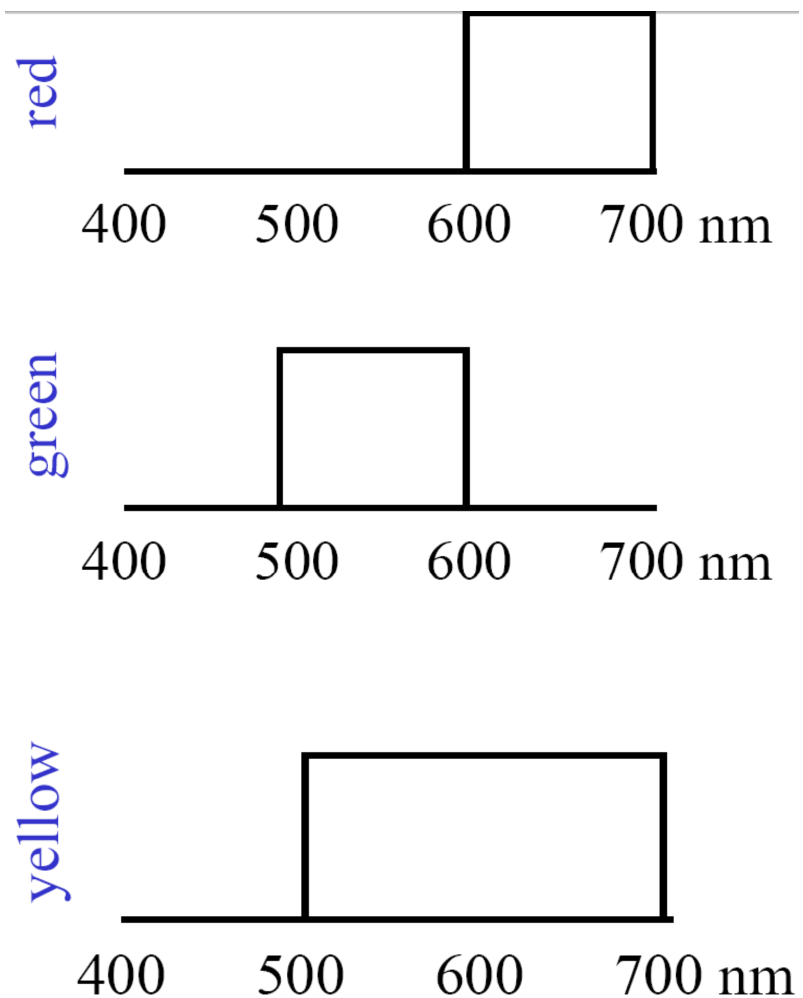

Colors combine by *adding* color spectra

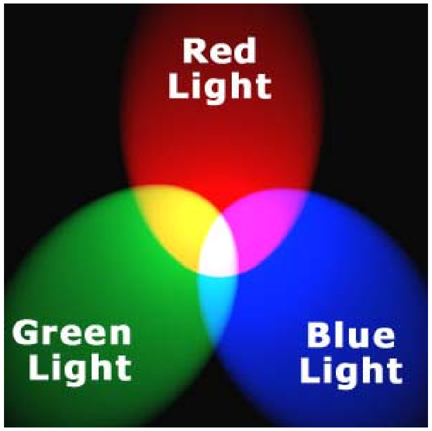

Light *adds* to black.

Source: W. Freeman

Examples of additive color systems

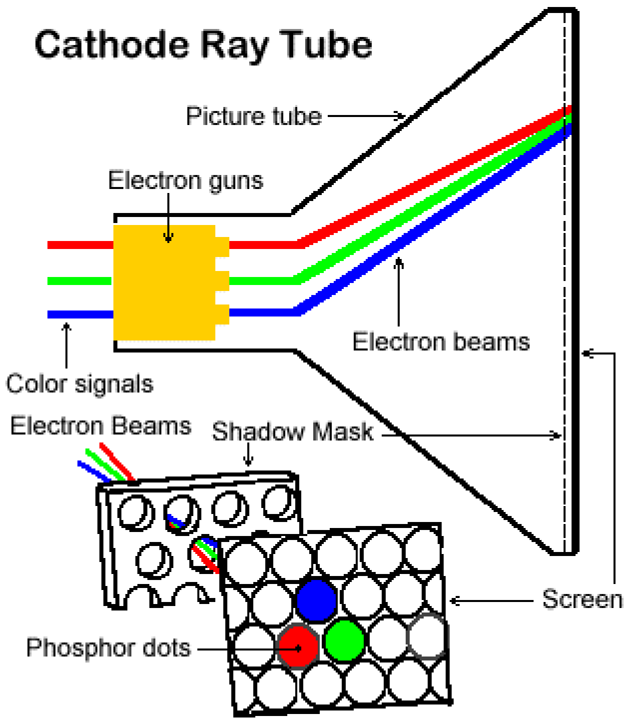

CRT phosphors

multiple projectors

http://www.jegsworks.com

http://www.crtprojectors.co.uk/

Superposition

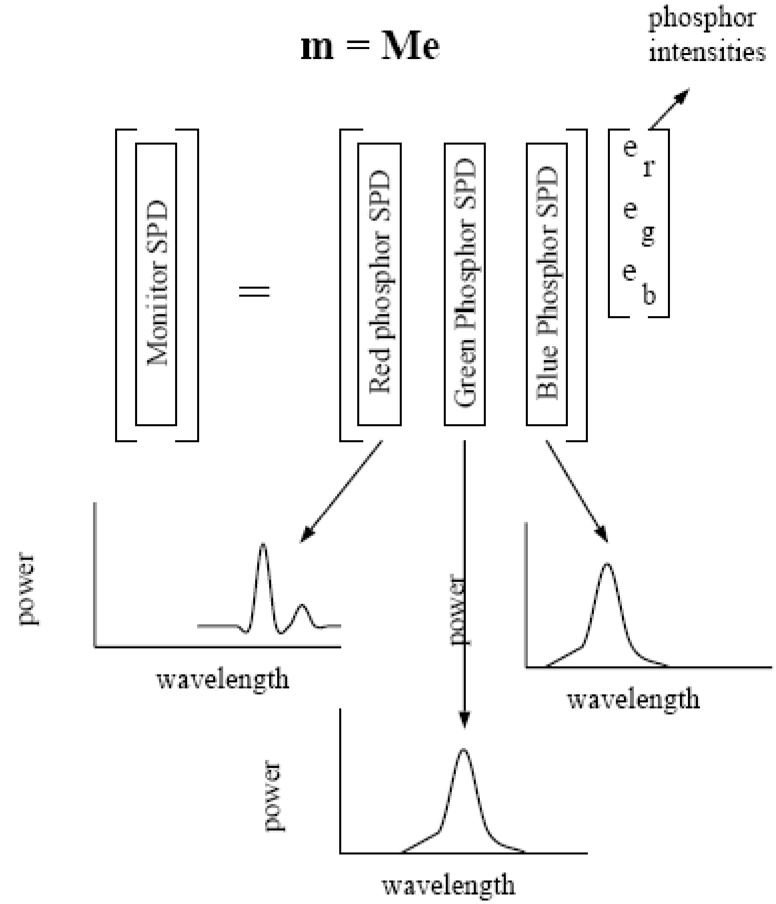

- Additive mixing:

  The spectral power distribution of the mixture is the sum of the spectral power distributions of the components.

Figure from B. Wandell, 1996

Subtractive color mixing

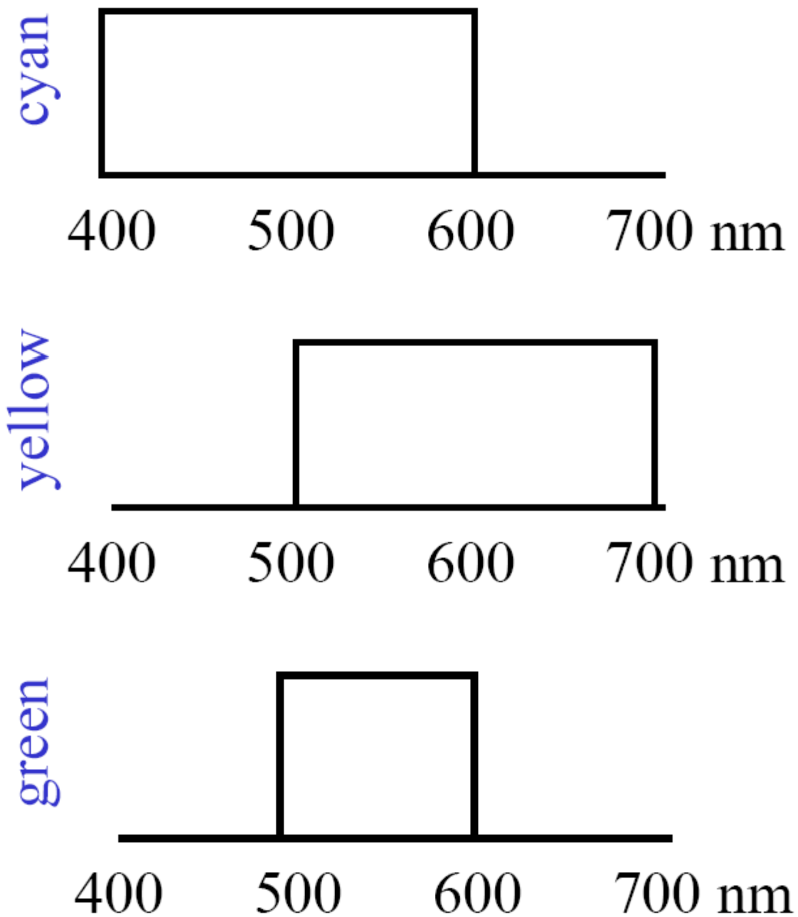

Colors combine by *multiplying* color spectra.

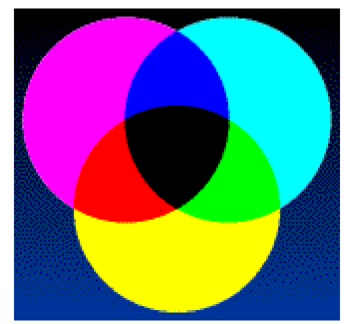

Pigments *remove* color from incident light (white).

Source: W. Freeman

Examples of subtractive color systems

- Printing on paper
- Crayons
- Most photographic film

Today: Color

- Measuring color
  - Spectral power distributions
  - Color mixing
  - Color matching experiments
  - Color spaces
    - Uniform color spaces

- Perception of color
  - Human photoreceptors
  - Environmental effects, adaptation

- Using color in machine vision systems

Why specify color *numerically*?

- Accurate color reproduction is commercially valuable
  - Many products are identified by color (“golden” arches)
- Few color *names* are widely recognized by English speakers
  - 11: black, blue, brown, grey, green, orange, pink, purple, red, white, and yellow.
  - Other languages have fewer/more.
  - Common to disagree on appropriate color names.
- Color reproduction problems increased by prevalence of digital imaging – e.g. digital libraries of art.
  - How to ensure that everyone perceives the same color?
  - **What spectral radiances *produce the same response* from people under simple viewing conditions?**

Color matching experiments

- Goal: find out what spectral radiances produce same response in human observers

Color matching experiments

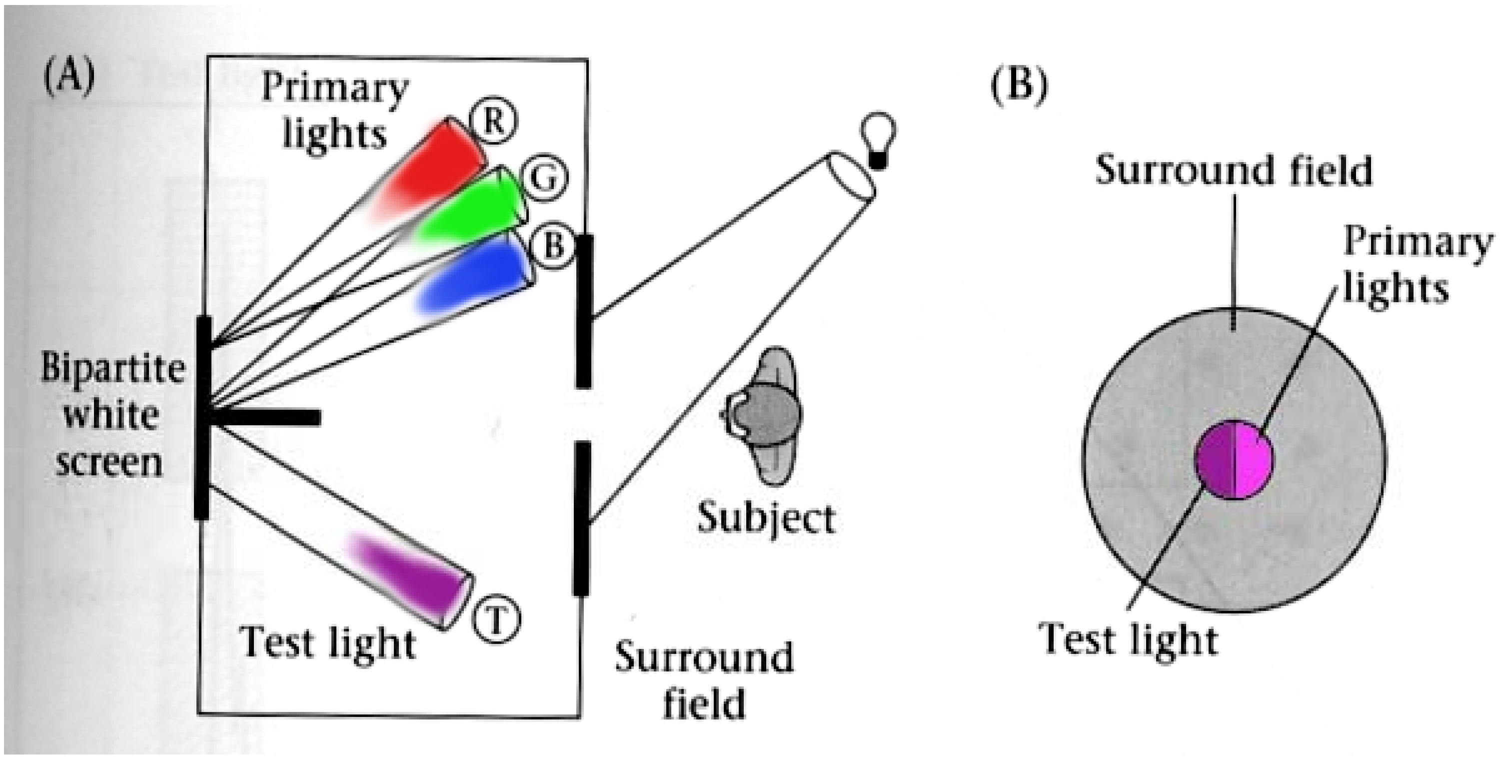

Observer adjusts weight (intensity) for primary lights (fixed SPD's) to match appearance of test light.

Foundations of Vision, by Brian Wandell, Sinauer Assoc., 1995

After Judd & Wyszecki.

Color matching experiments

- Goal: find out what spectral radiances produce same response in human observers
- Assumption: simple viewing conditions, where we say test light alone affects perception
  - Ignoring additional factors for now like adaptation, complex surrounding scenes, etc.

Color matching experiment 1

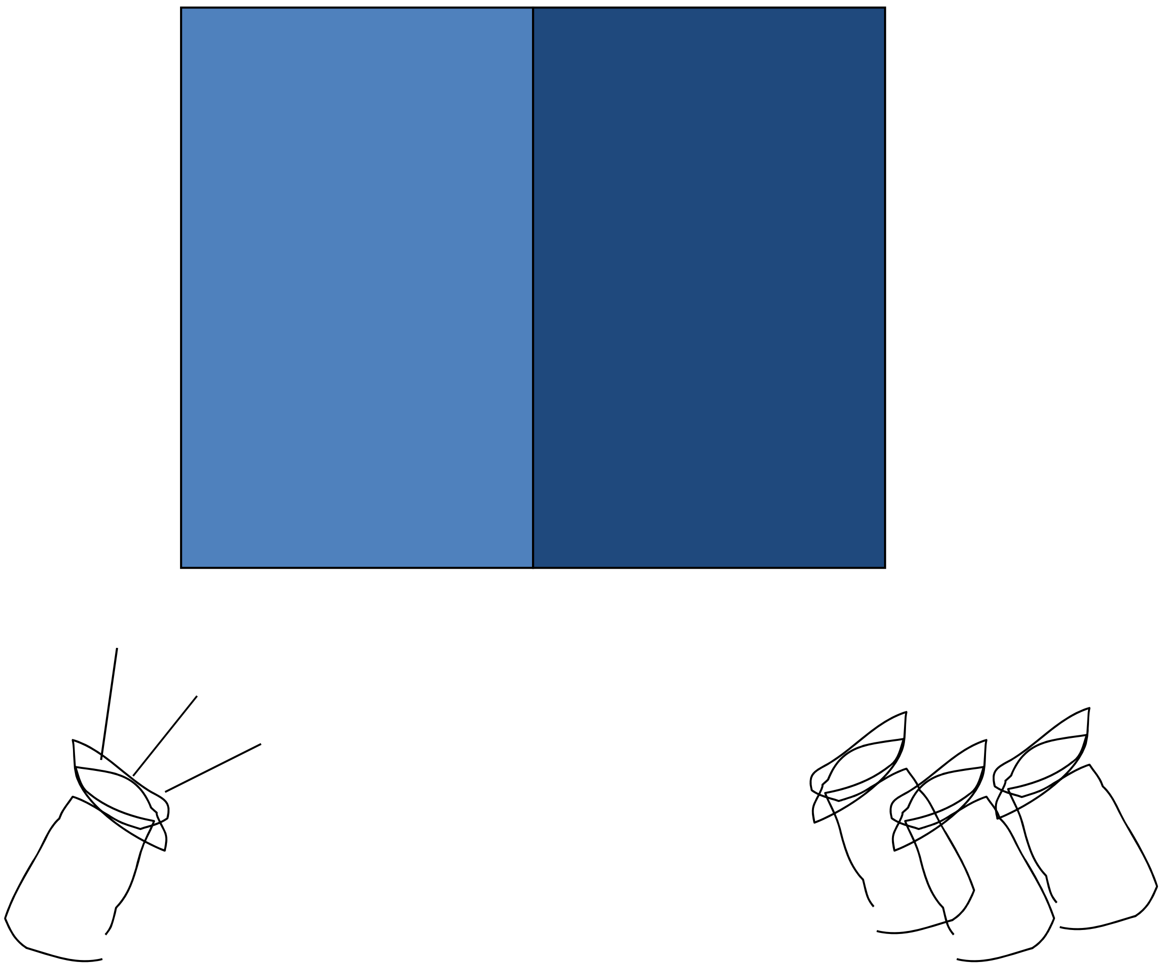

Slide credit: W. Freeman

Color matching experiment 1

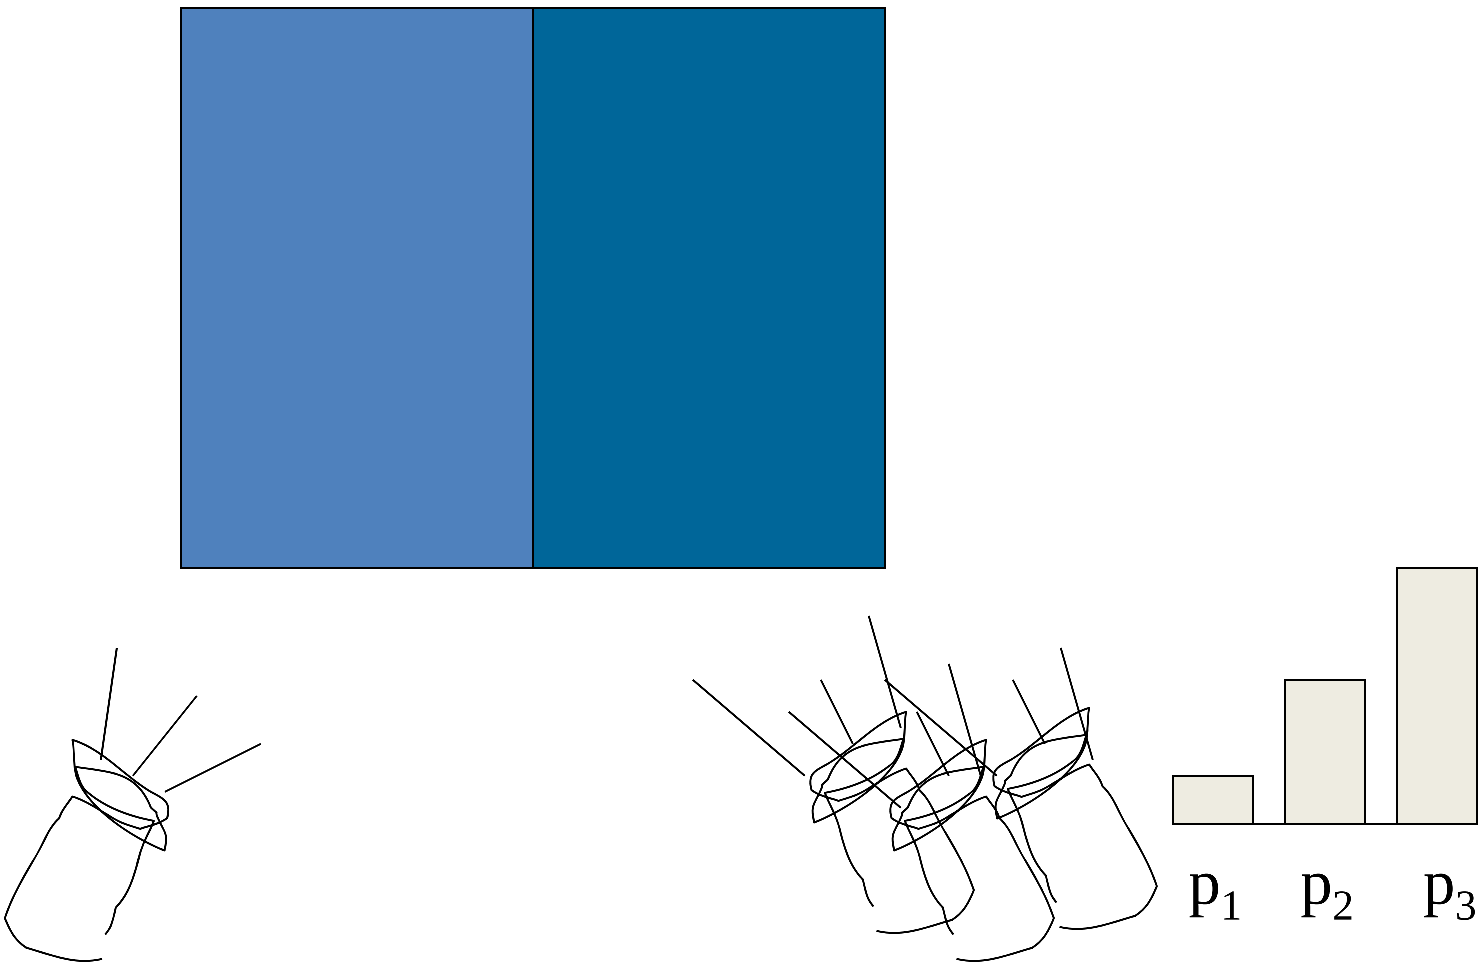

Slide credit: W. Freeman

Color matching experiment 1

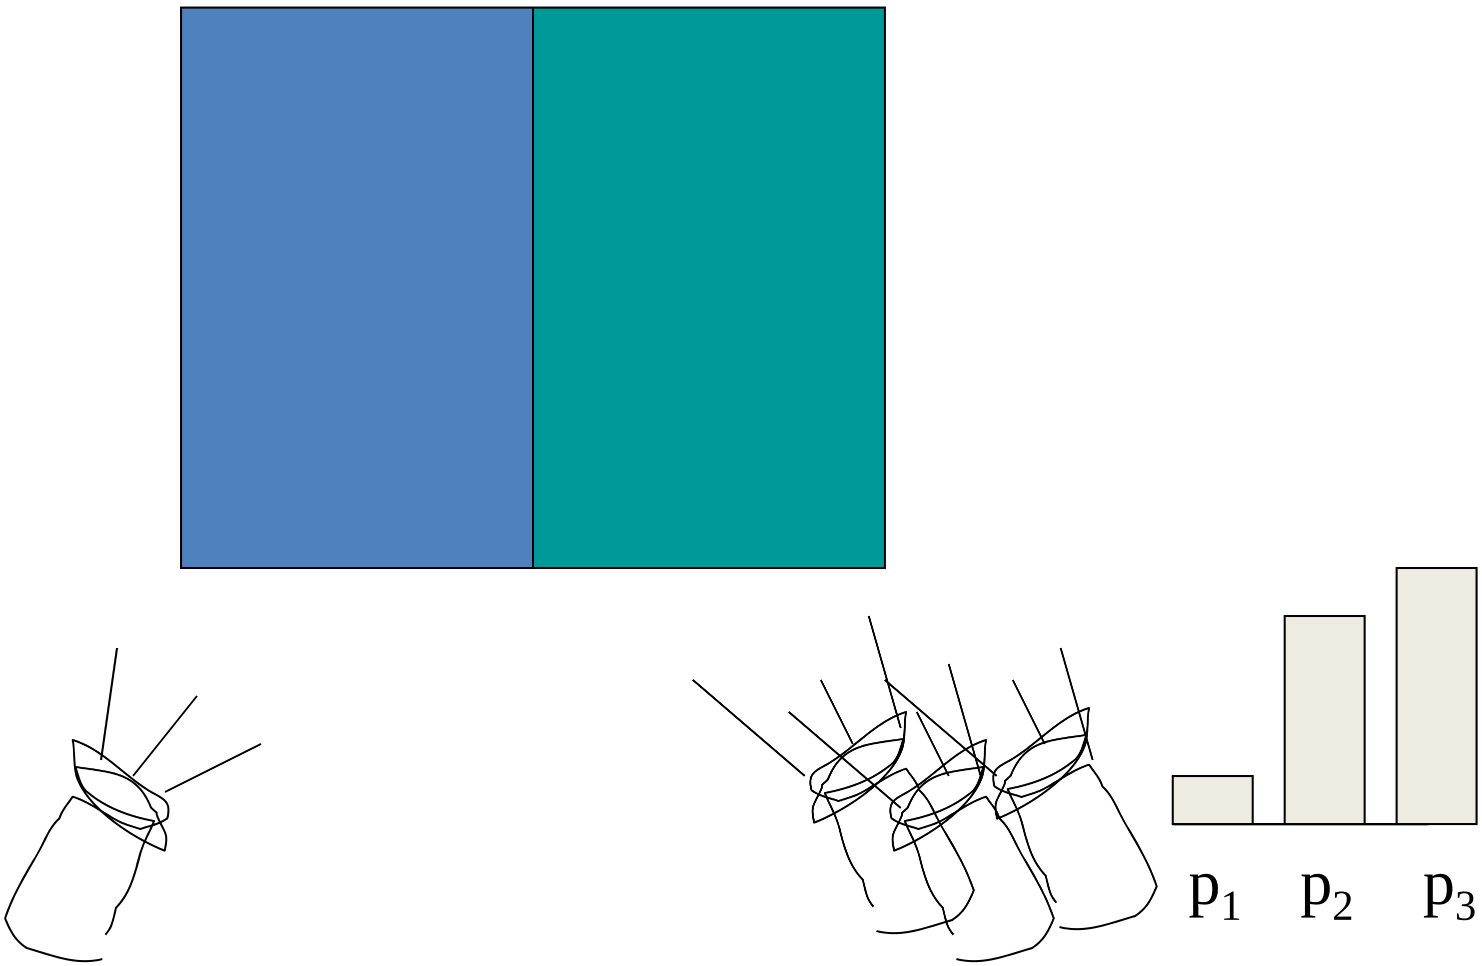

Slide credit: W. Freeman

Color matching experiment 1

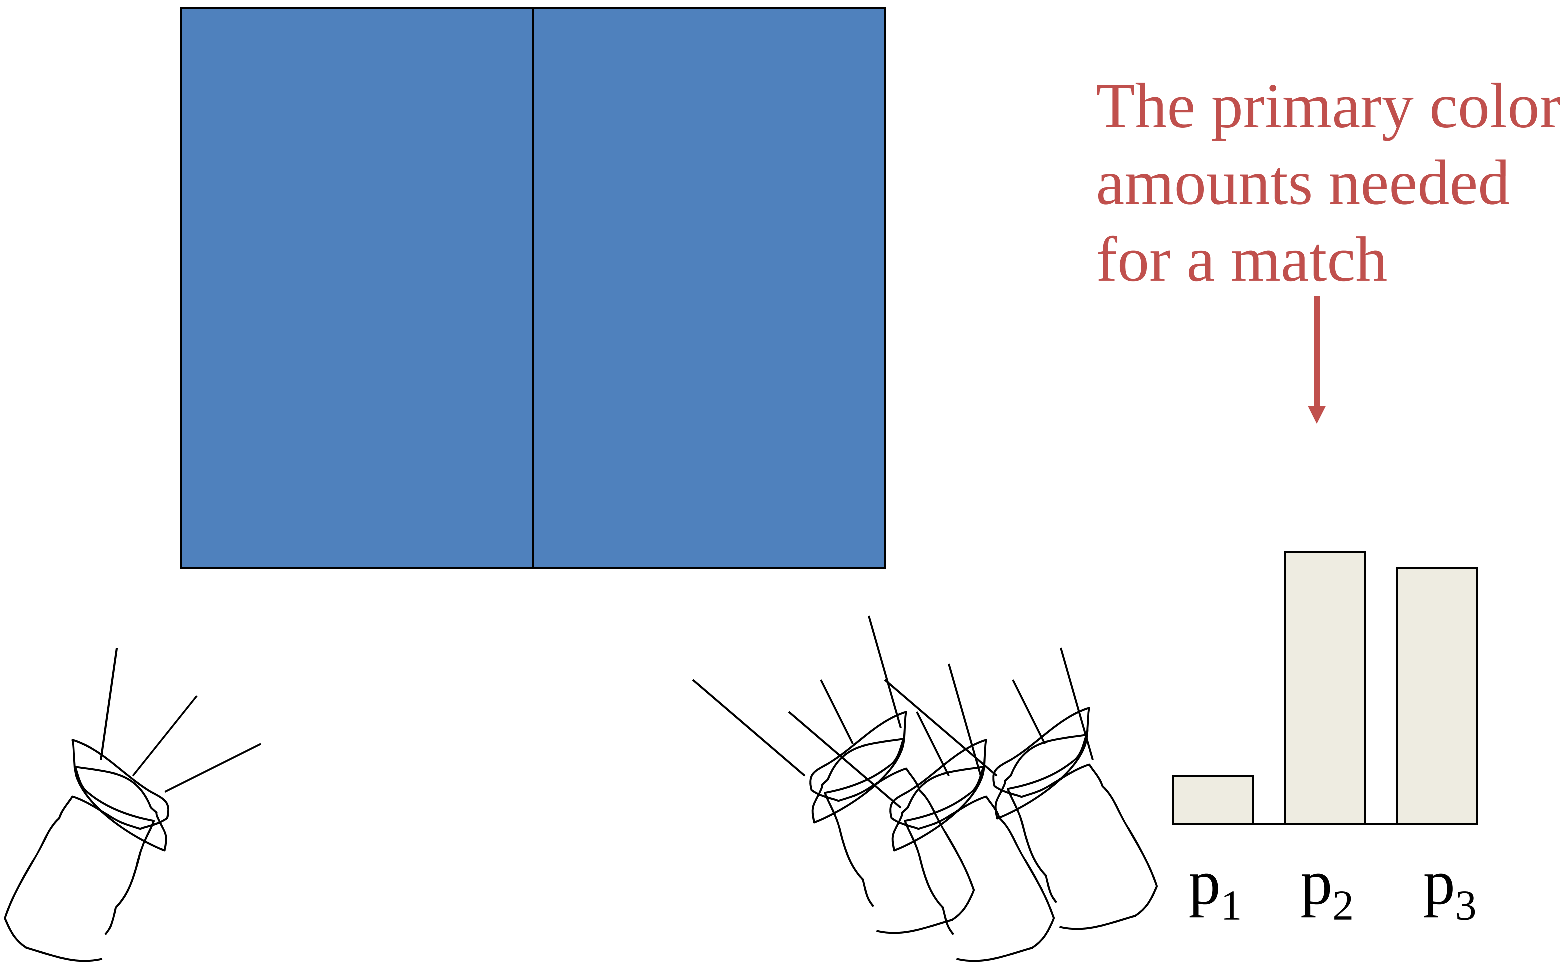

Slide credit: W. Freeman

Color matching experiment 2

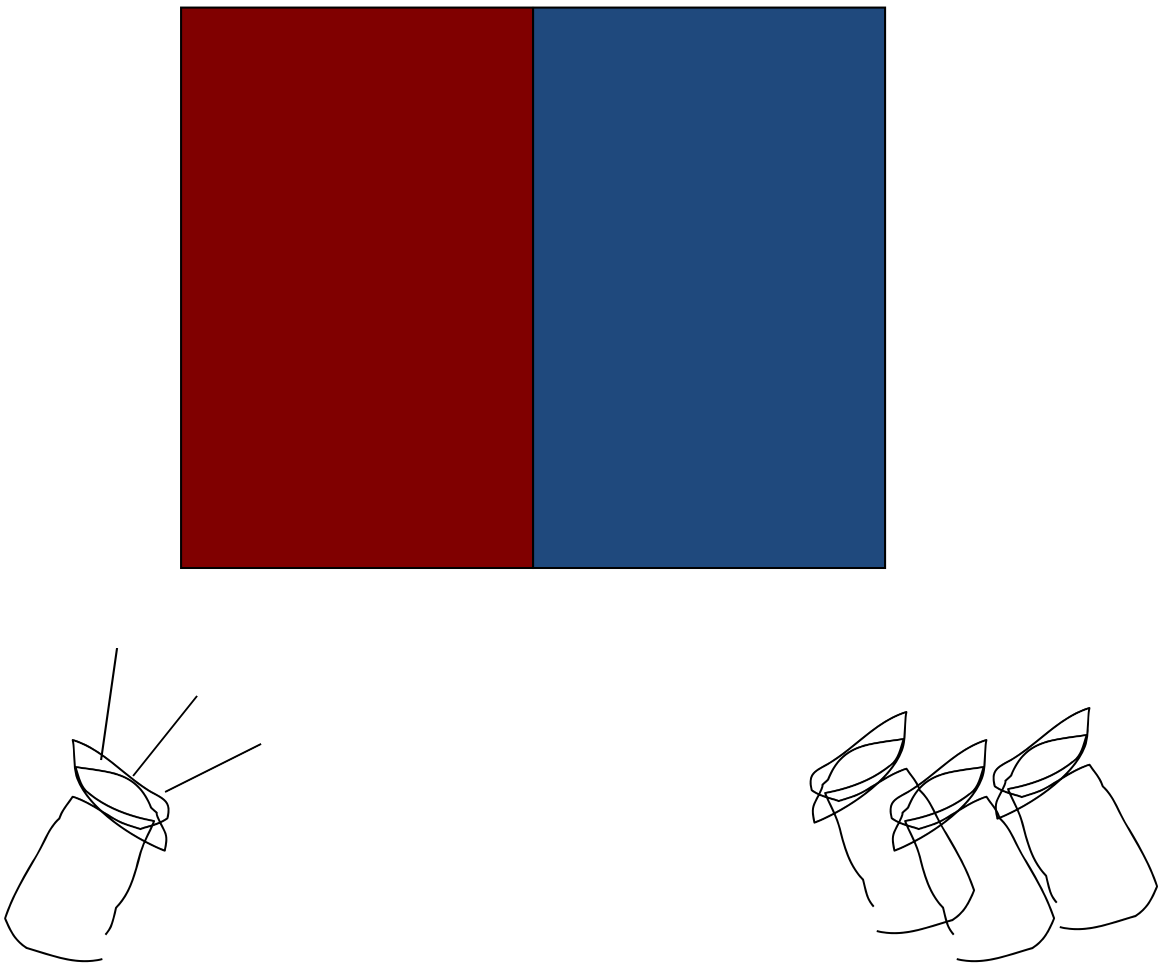

Slide credit: W. Freeman

Color matching experiment 2

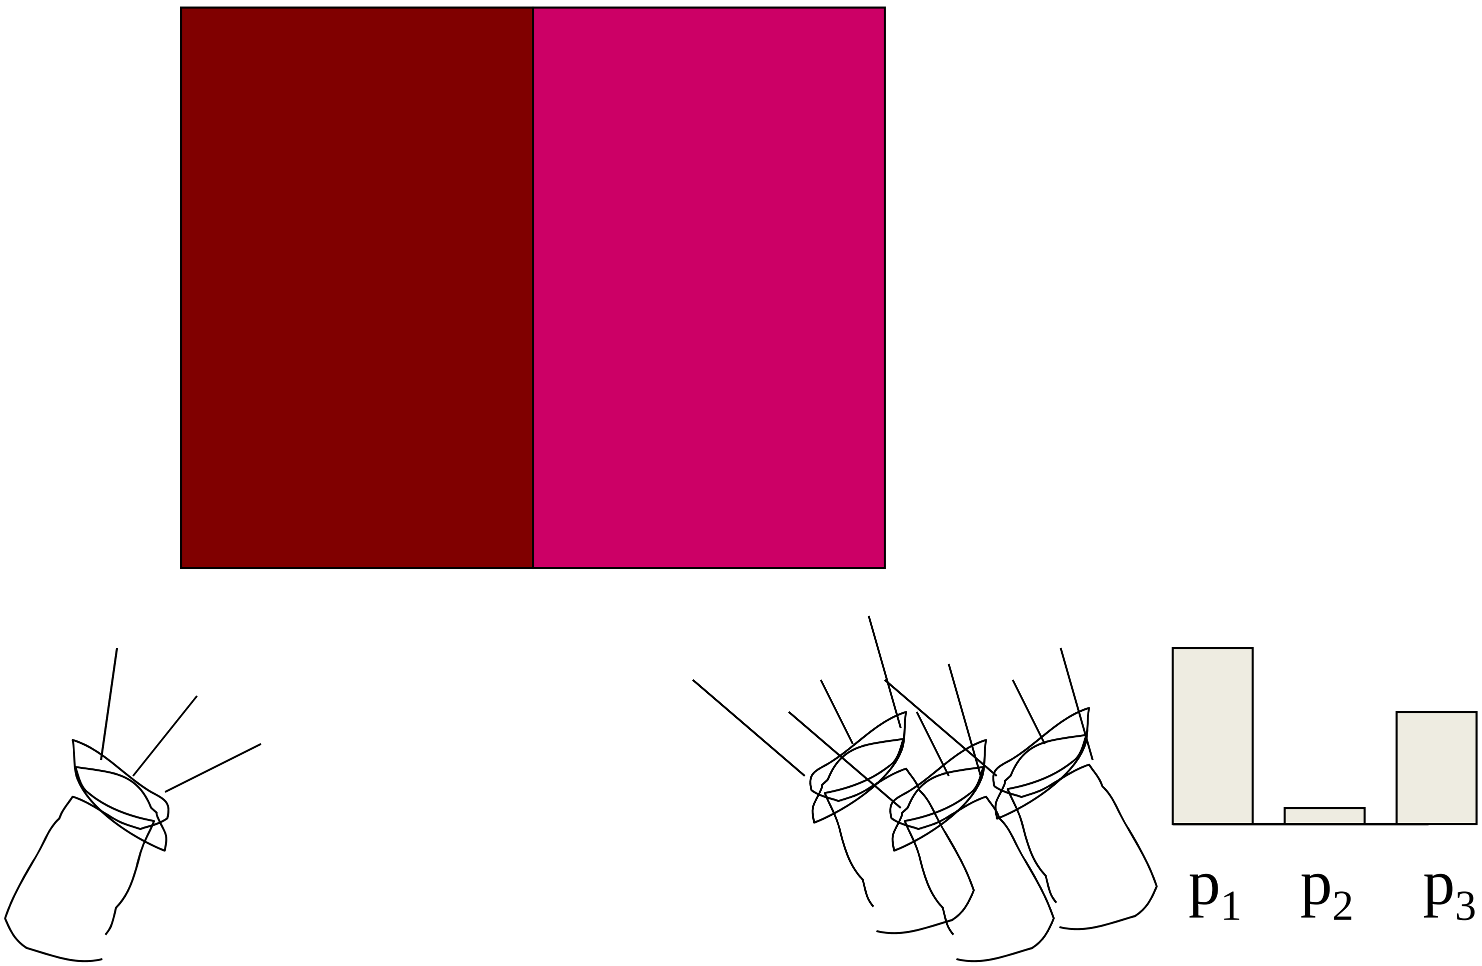

Slide credit: W. Freeman

Color matching experiment 2

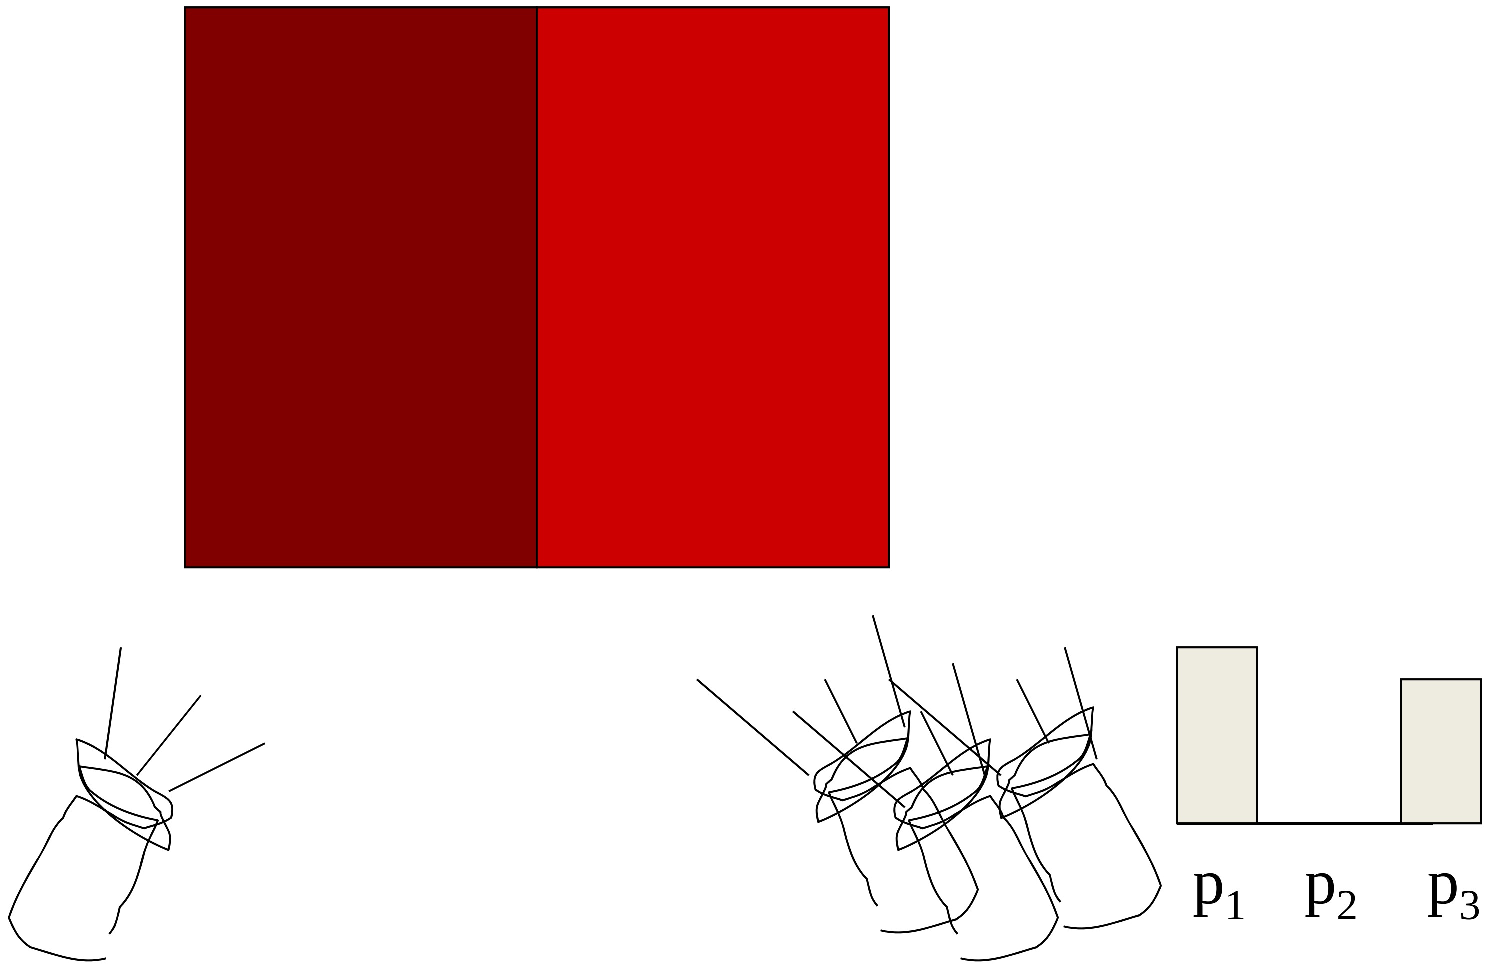

Slide credit: W. Freeman

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.

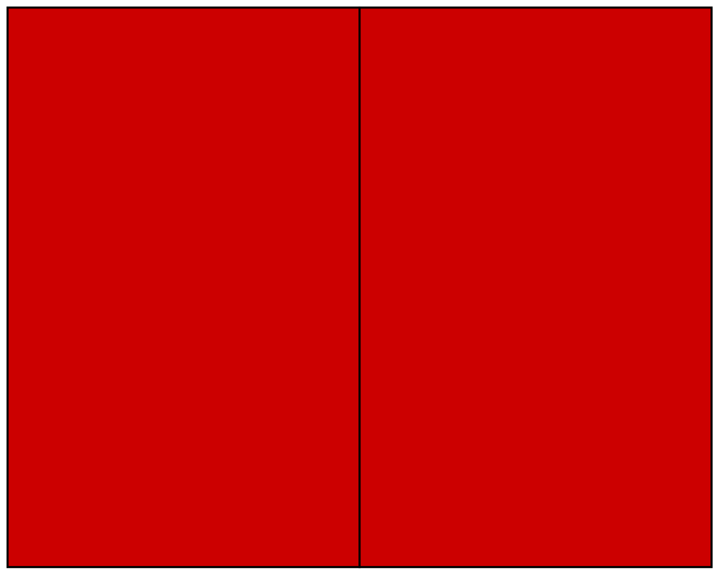

The primary color amounts needed for a match:

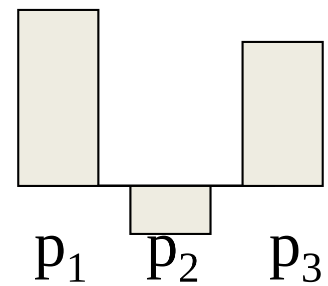

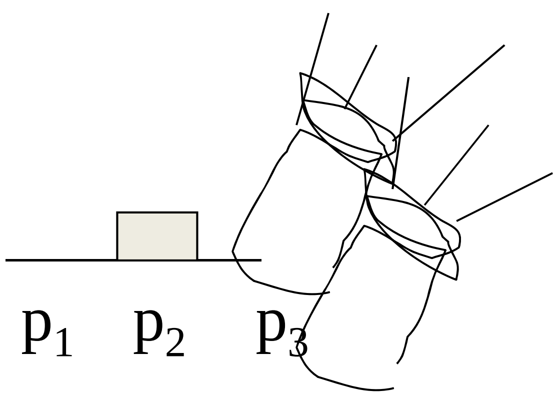

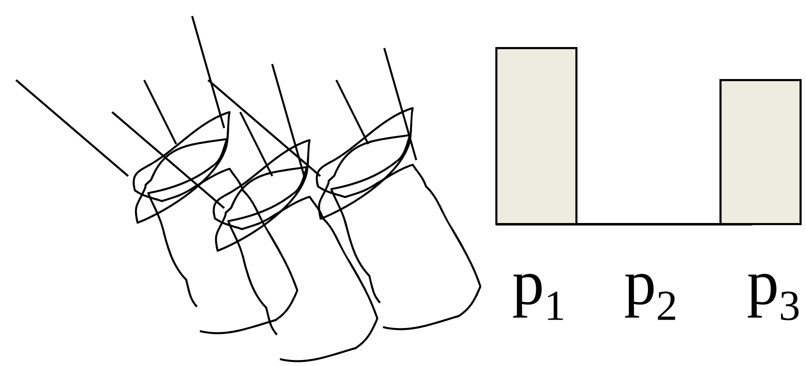

Color matching

- What must we require of the primary lights chosen?
- How are three numbers enough to represent entire spectrum?

Metamers

- If observer says a mixture is a match → receptor excitations of both stimuli must be equal
- But lights forming a *perceptual* match still may be *physically* different
  - Match light: must be combination of primaries
  - Test light: any light
- **Metamers**: pairs of lights that match perceptually but not physically

Metameric spectral power distributions

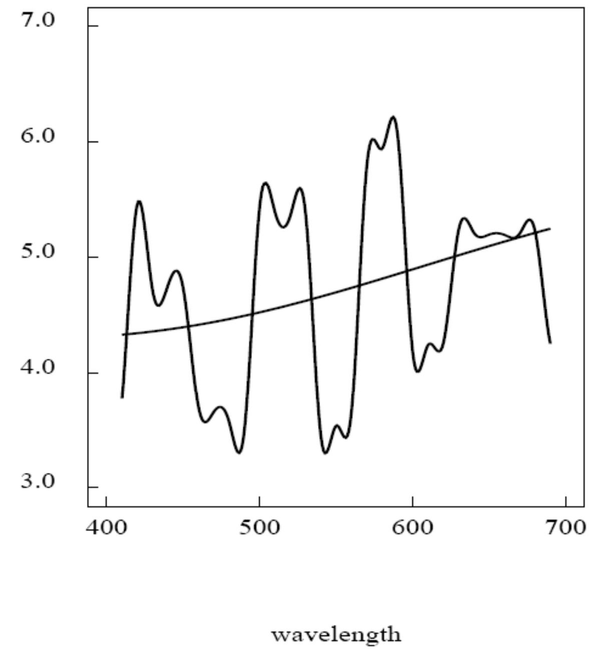

Fig from B. Wandell, 1996

Grassman's laws

- If two test lights can be matched with the same set of weights, then they match each other:
  - Suppose $A = u_1 P_1 + u_2 P_2 + u_3 P_3$ and $B = u_1 P_1 + u_2 P_2 + u_3 P_3$. Then $A = B$.
- If we scale the test light, then the matches get scaled by the same amount:
  - Suppose $A = u_1 P_1 + u_2 P_2 + u_3 P_3$.
    Then $kA = (ku_1) P_1 + (ku_2) P_2 + (ku_3) P_3$.
- If we mix two test lights, then mixing the matches will match the result (superposition):
  - Suppose $A = u_1 P_1 + u_2 P_2 + u_3 P_3$ and $B = v_1 P_1 + v_2 P_2 + v_3 P_3$. Then $A+B = (u_1+v_1) P_1 + (u_2+v_2) P_2 + (u_3+v_3) P_3$.

Here "=" means "matches".

Computing color matches

- How do we compute the weights that will yield a perceptual match for any test light using a given set of primaries?
  1. Select primaries
  2. Estimate their *color matching functions*: observer matches series of monochromatic lights, one at each wavelength
  3. Multiply matching functions and test light

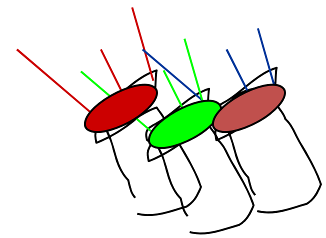

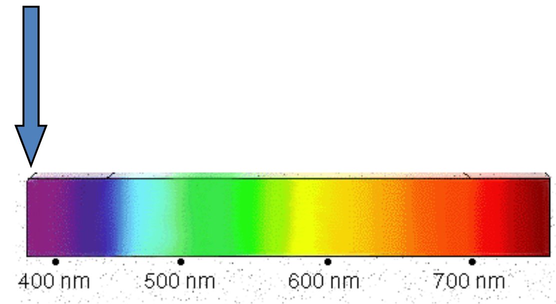

$$\mathbf{C} = \begin{pmatrix} c_1(\lambda_1) \\ c_2(\lambda_1) \\ c_3(\lambda_1) \end{pmatrix}$$

Computing color matches

Color matching functions for a particular set of primaries

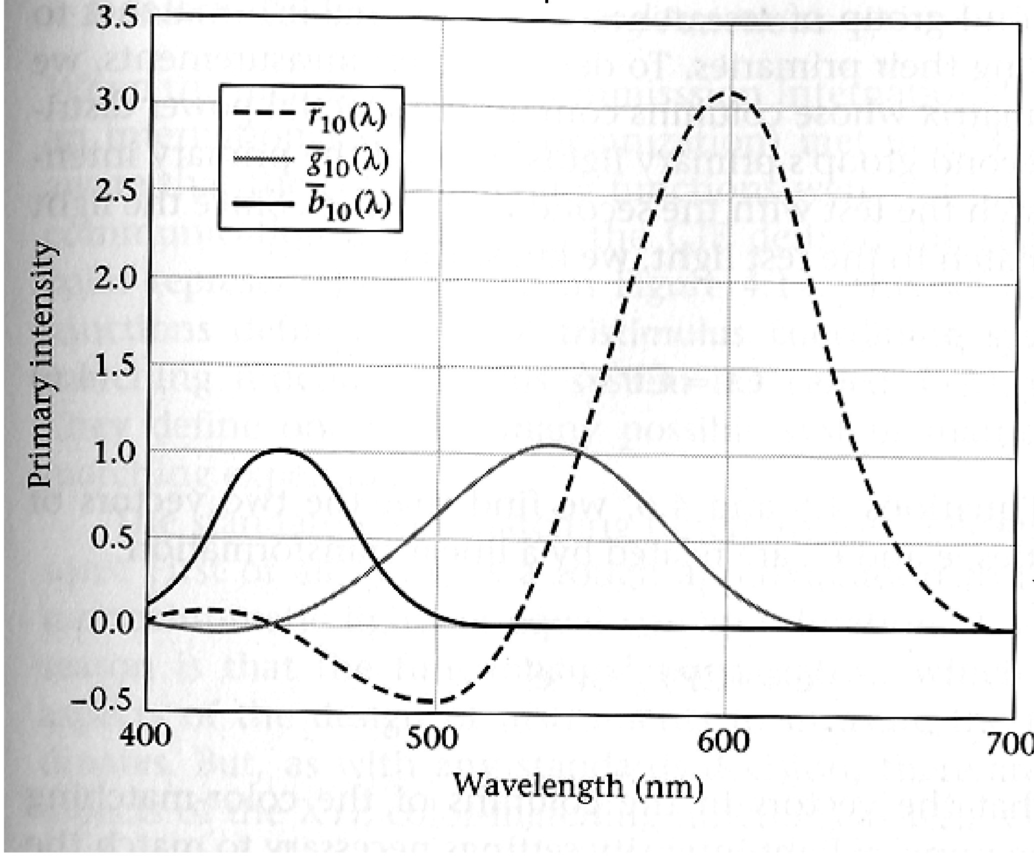

$p_1 = 645.2$ nm

$p_2 = 525.3$ nm

$p_3 = 444.4$ nm

Rows of matrix C

Foundations of Vision, by Brian Wandell, Sinauer Assoc., 1995

Slide credit: W. Freeman

Computing color matches

$$\lambda_i \quad \xleftrightarrow{\text{matches}} \quad c_1(\lambda_i), c_2(\lambda_i), c_3(\lambda_i)$$

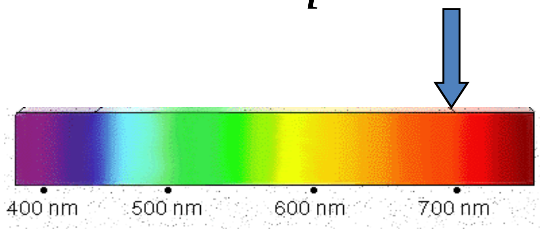

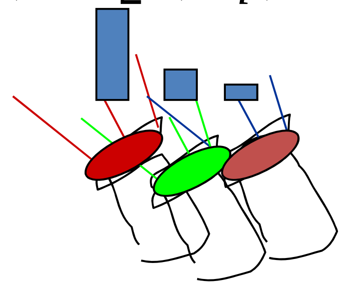

Now have matching functions for all monochromatic light sources, so we know how to match a unit of each wavelength.

Arbitrary new spectral signal is a linear combination of the monochromatic sources.

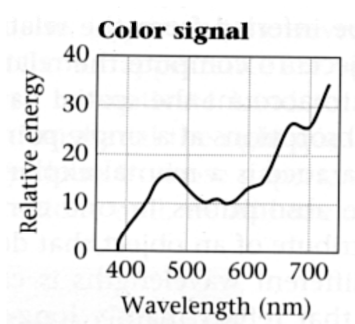

$$t = \begin{pmatrix} t(\lambda_1) \\ \dots \\ t(\lambda_N) \end{pmatrix}$$

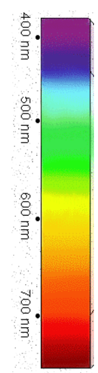

Computing color matches

So, given any set of primaries and their associated matching functions (C), we can compute weights (e) needed on each primary to give a perceptual match to any test light t (spectral signal).

$$\mathbf{e} = \mathbf{C}\mathbf{t}$$

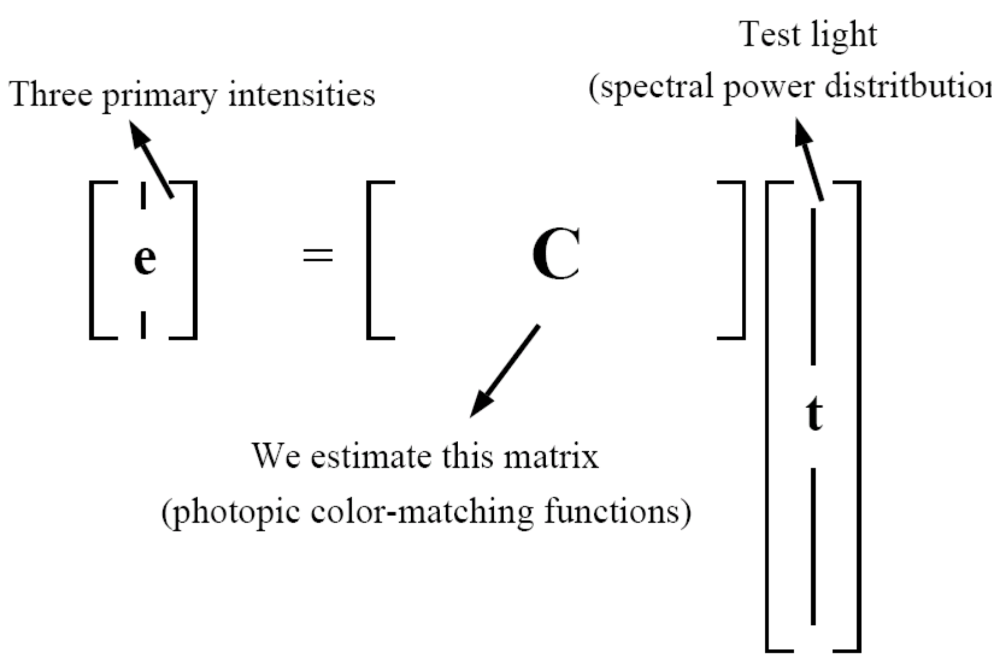

Fig from B. Wandell, 1996

Computing color matches

- Why is computing the color match for any color signal for a given set of primaries useful?
  - Want to paint a carton of Kodak film with the Kodak yellow color.
  - Want to match skin color of a person in a photograph printed on an ink jet printer to their true skin color.
  - Want the colors in the world, on a monitor, and in a print format to all look the same.

Adapted from W. Freeman

Image credit: pbs.org

Today: Color

- Measuring color
  - Spectral power distributions
  - Color mixing
  - Color matching experiments
  - Color spaces
    - Uniform color spaces
- Perception of color
  - Human photoreceptors
  - Environmental effects, adaptation
- Using color in machine vision systems

Standard color spaces

- Use a common set of primaries/color matching functions
- Linear color space examples
  - RGB
  - CIE XYZ
- Non-linear color space
  - HSV

RGB color space

- Single wavelength primaries
- Good for devices (e.g., phosphors for monitor), but not for perception

RGB color matching functions

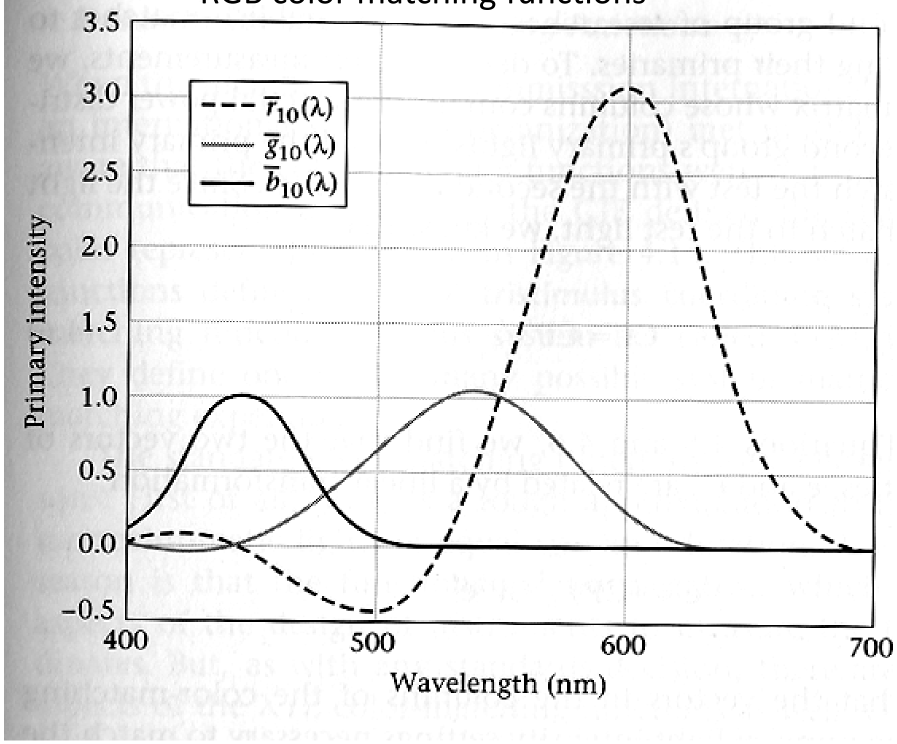

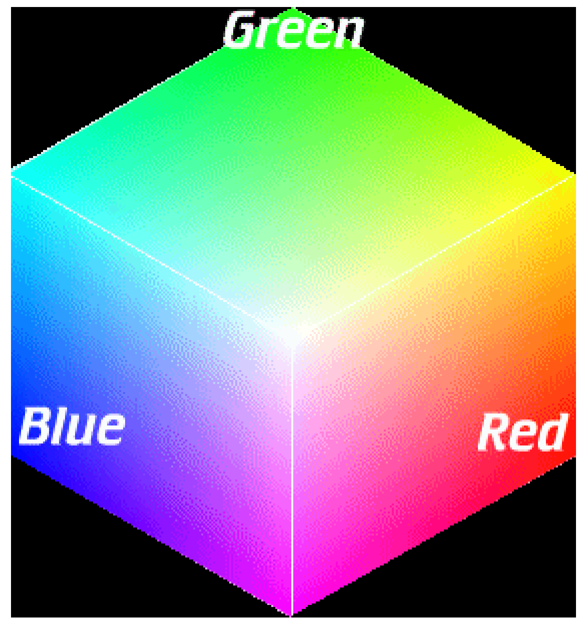

CIE XYZ color space

- Established by the commission international d'eclairage (CIE), 1931
- Usually projected to display:

  (x,y) = (X/(X+Y+Z), Y/(X+Y+Z))

CIE XYZ Color matching functions

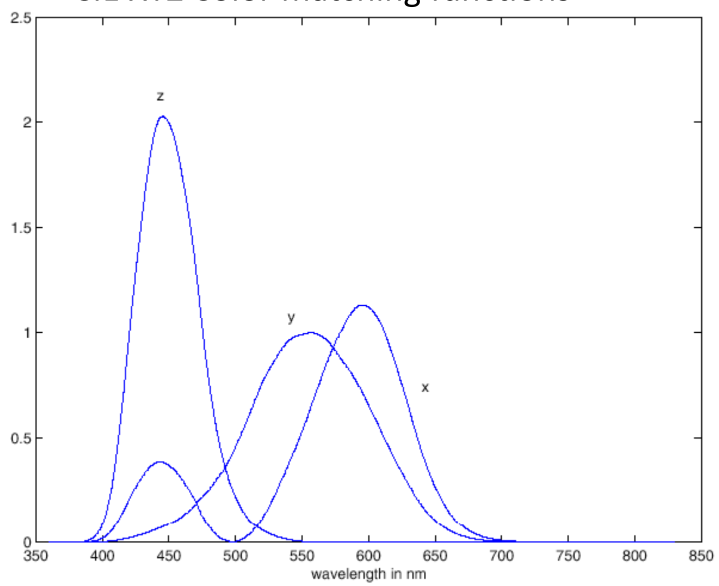

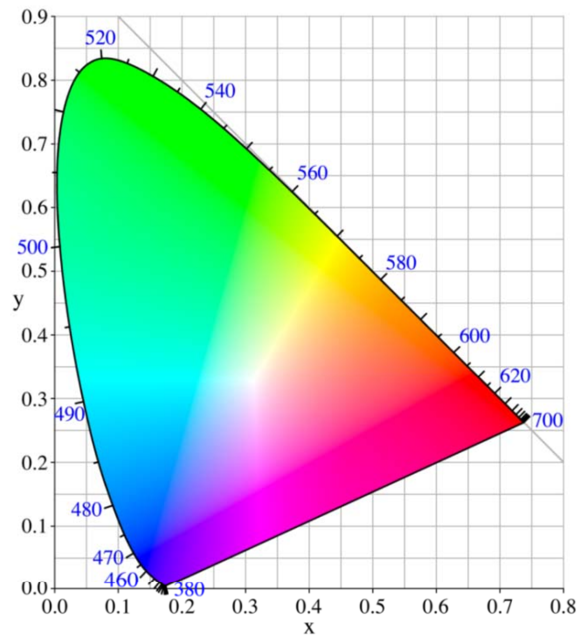

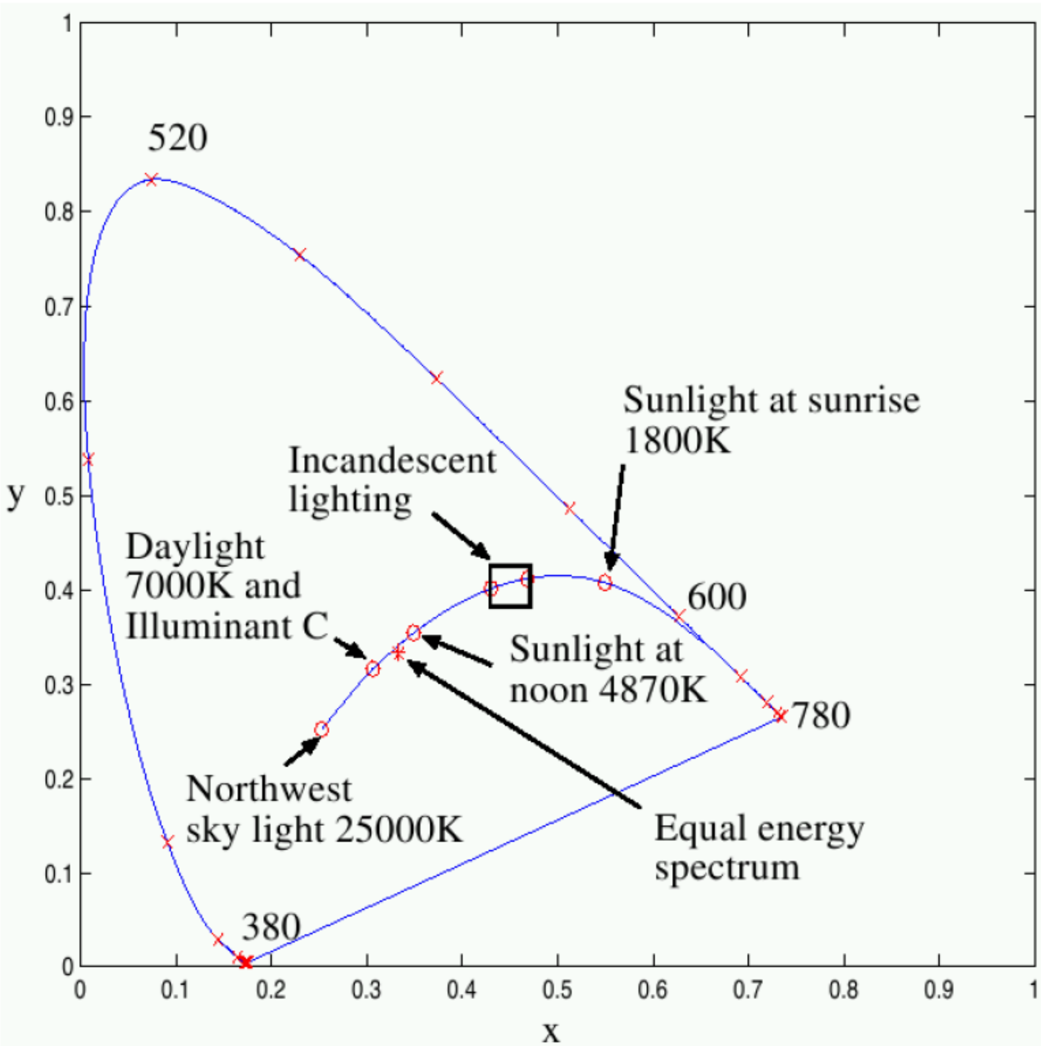

Forsyth & Ponce Figure 6.9

The figure shows:

- The spectral locus (i.e., the colours of monochromatic lights)
- The black-body locus (i.e., the colors of heated black-bodies)
- Some common illuminants

HSV color space

- **H**ue, **S**aturation, **V**alue (Brightness)
- Nonlinear – reflects topology of colors by coding **hue** as an angle
- Matlab: `hsv2rgb`, `rgb2hsv`.

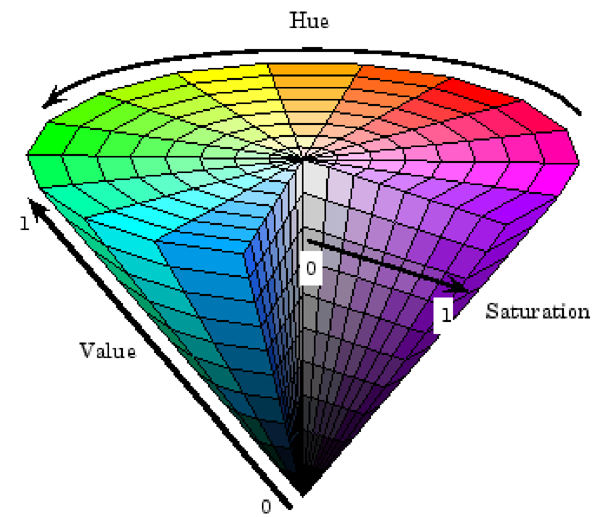

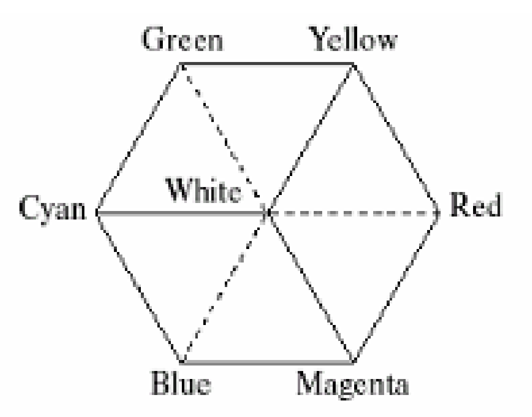

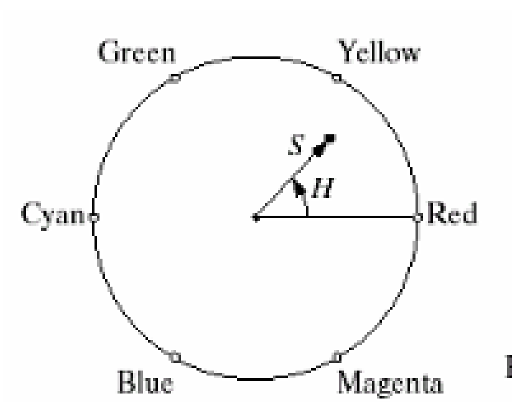

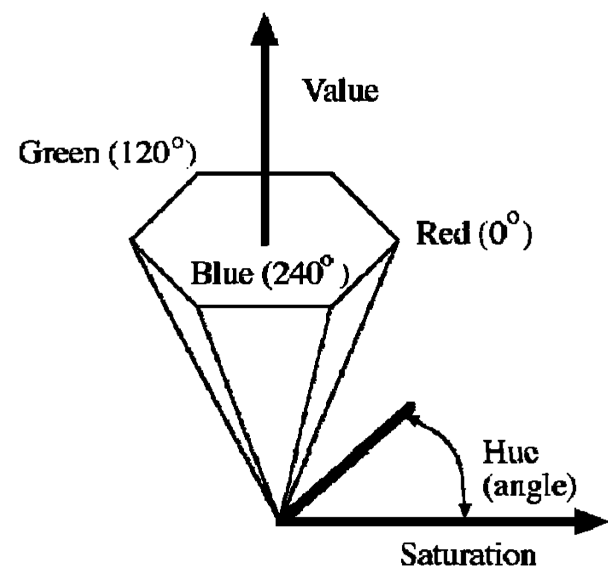

Image from mathworks.com

Distances in color space

- Are distances between points in a color space perceptually meaningful?

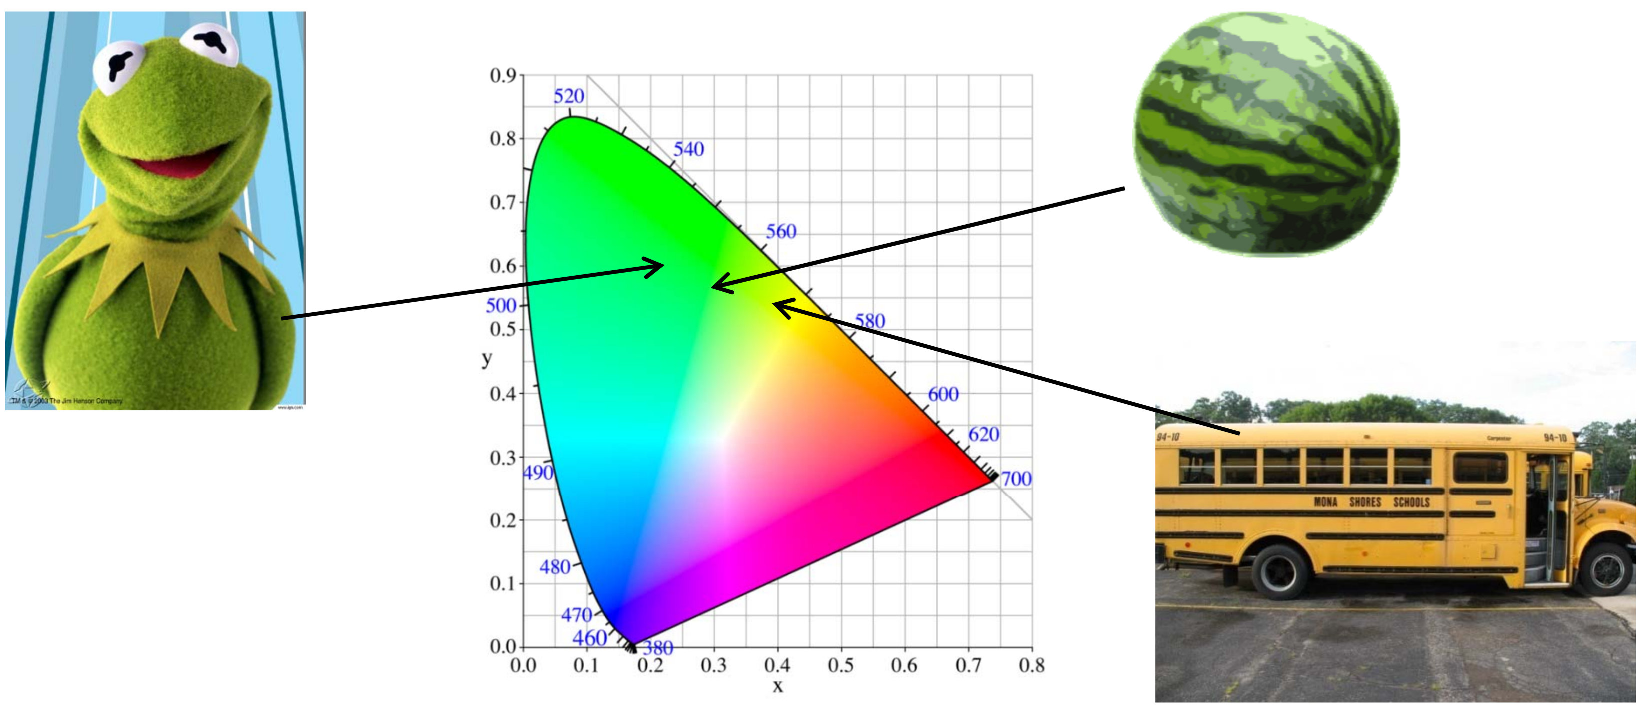

Distances in color space

- Not necessarily: CIE XYZ is **not** a *uniform* color space, so magnitude of differences in coordinates are poor indicator of color "distance".

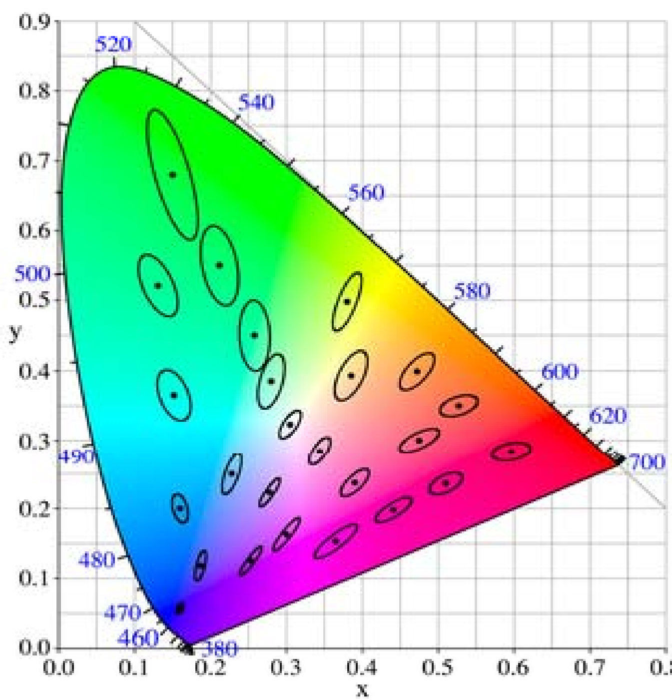

McAdam ellipses:
Just noticeable differences in color

Uniform color spaces

- Attempt to correct this limitation by remapping color space so that just-noticeable differences are contained by circles → distances more perceptually meaningful.

- Examples:
  – CIE u’v’
  – CIE Lab

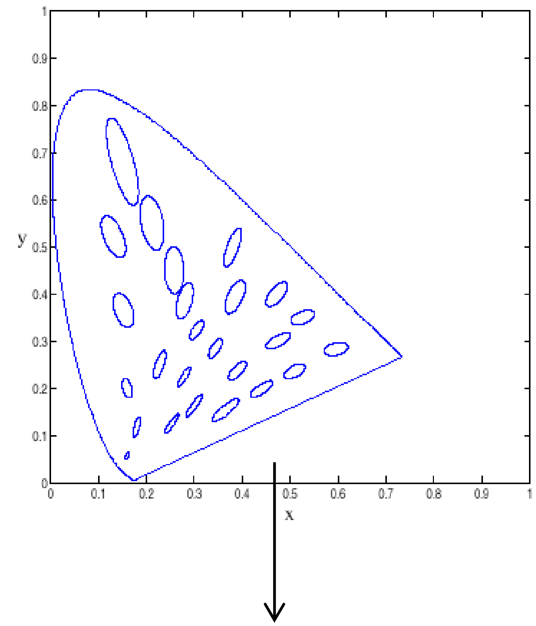

CIE XYZ

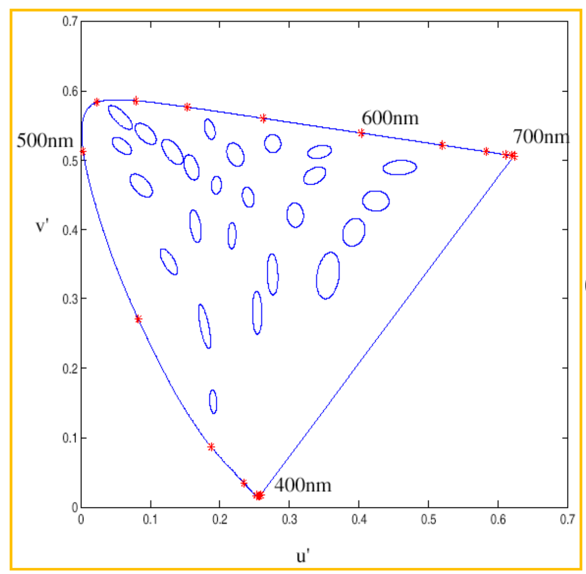

CIE u’v’

Today: Color

- Measuring color
  - Spectral power distributions
  - Color mixing
  - Color matching experiments
  - Color spaces
    - Uniform color spaces

- Perception of color
  - Human photoreceptors
  - Environmental effects, adaptation

- Using color in machine vision systems

Color

- **Color of light** arriving at camera depends on
  - Spectral reflectance of the surface light is leaving
  - Spectral radiance of light falling on that patch
- **Color perceived** depends on
  - Physics of light
  - Visual system receptors
  - Brain processing, environment

Human photoreceptors

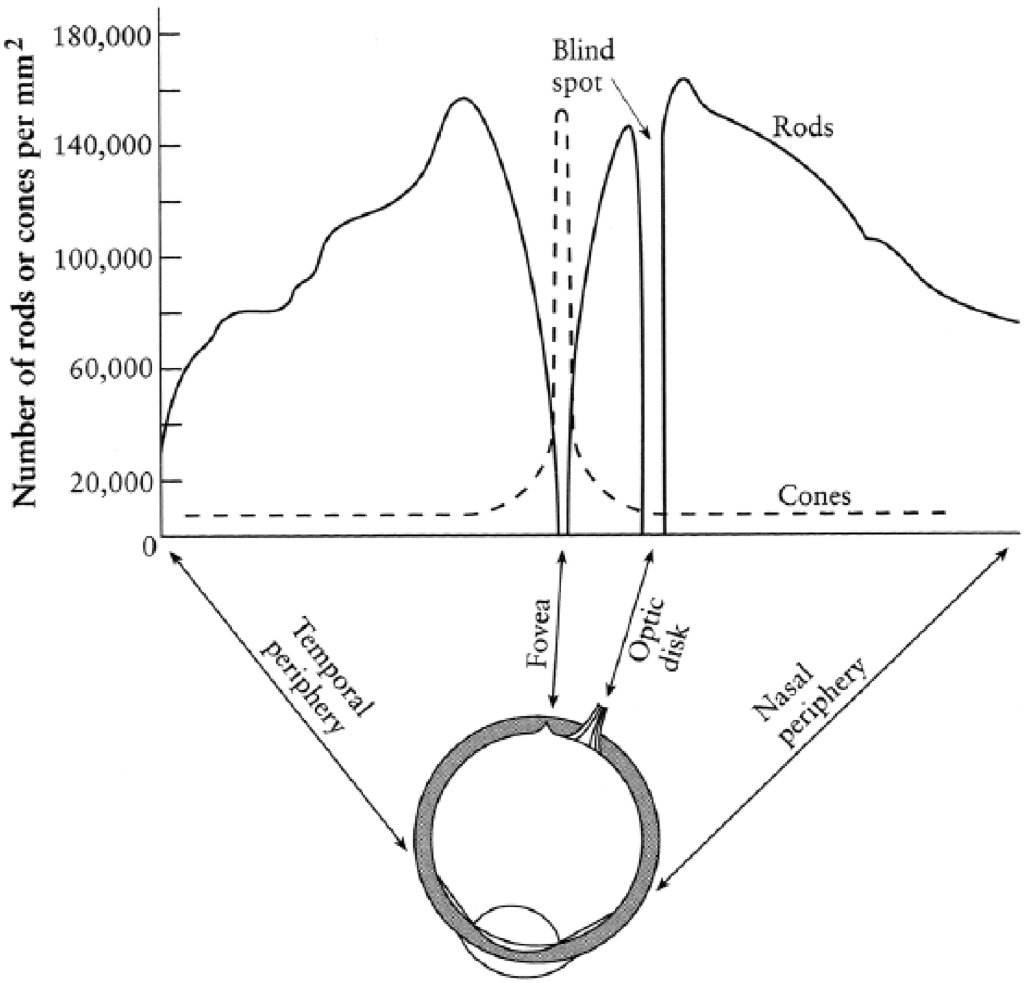

-Rods responsible for intensity
-Cones responsible for color
-Fovea: small region (1 or 2°) at the center of the visual field containing the highest density of cones (and no rods).
– Less visual acuity in the periphery

Adapted from Seitz, Duygulu

Two types of light-sensitive receptors

Cones
- cone-shaped
- less sensitive
- operate in high light
- color vision

Rods
- rod-shaped
- highly sensitive
- operate at night
- gray-scale vision

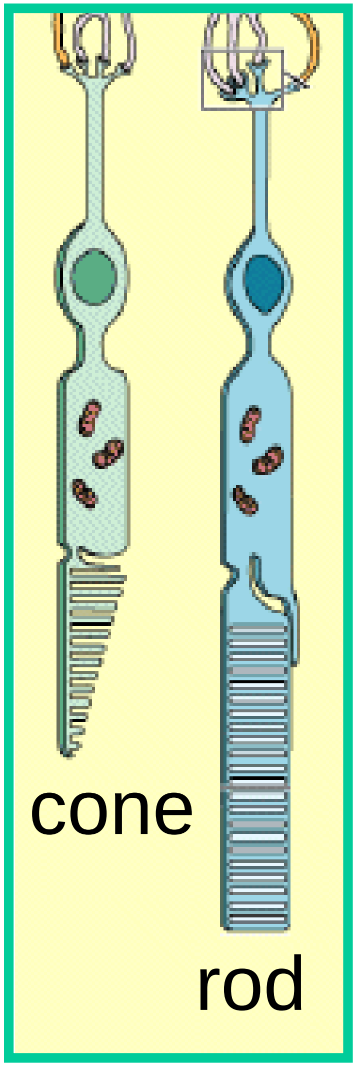

Slide credit: Alyosha Efros

Human photoreceptors

- React only to some wavelengths, with different sensitivity (light fraction absorbed)
- Brain fuses responses from local neighborhood of several cones for perceived color
- Sensitivities vary from person to person, and with age
- Color blindness: deficiency in at least one type of cone

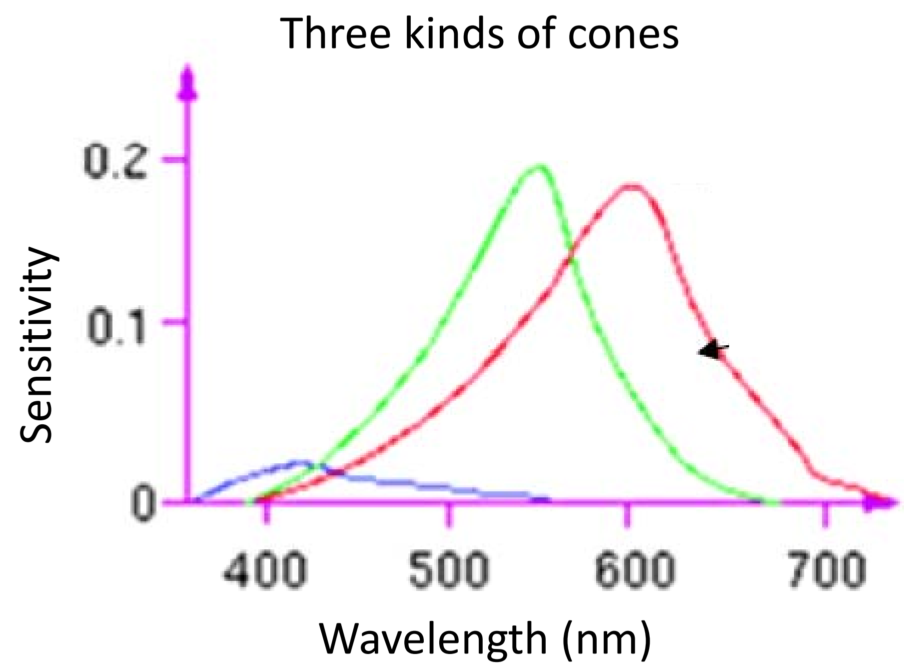

Human photoreceptors

Possible evolutionary pressure for developing receptors for different wavelengths in primates

Osorio & Vorobyev, 1996

Trichromacy

- Experimental facts:
  - Three primaries will work for most people if we allow subtractive matching; “trichromatic” nature of the human visual system
  - Most people make the *same* matches for a given set of primaries (i.e., select the same mixtures)

Environmental effects & adaptation

- *Chromatic adaptation*: we adapt to a particular illuminant
- *Assimilation, contrast effects, chromatic induction*: nearby colors affect what is perceived; receptor excitations interact across image and time
- *Afterimages*

Color matching ~= color appearance

Physics of light ~= perception of light

Chromatic adaptation

- If the visual system is exposed to a certain illuminant for a while, color system starts to adapt / skew.

Chromatic adaptation

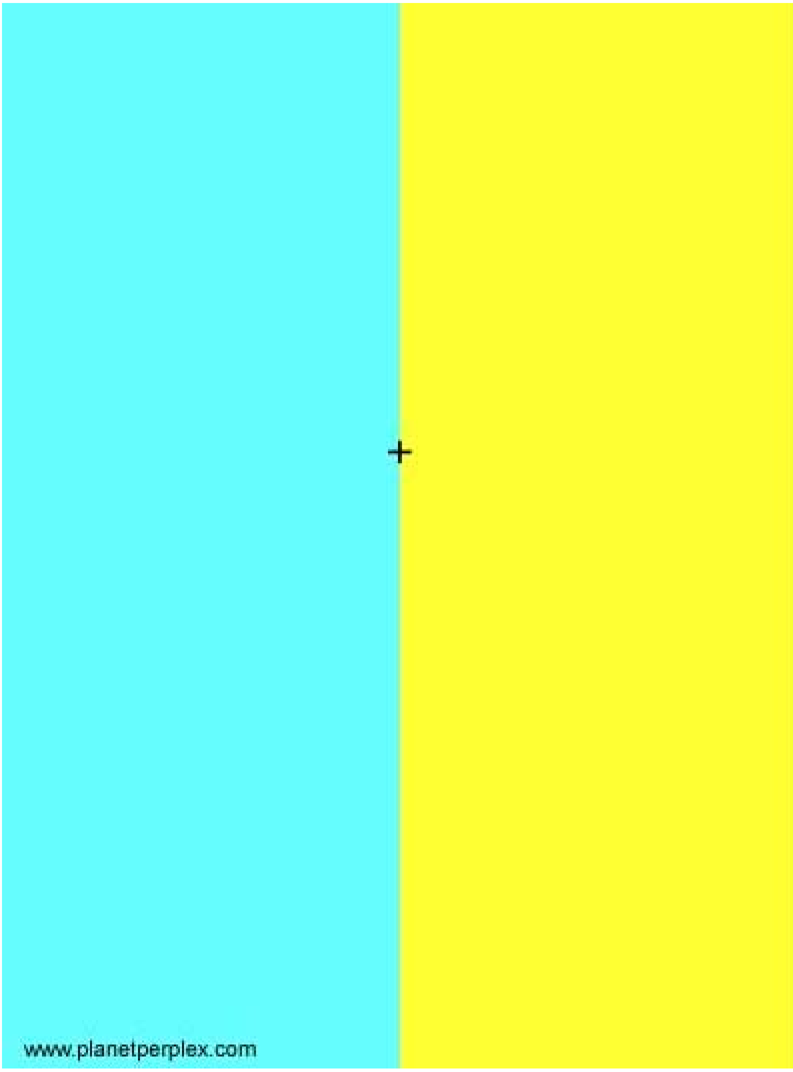

http://www.planetperplex.com/en/color_illusions.html

Brightness perception

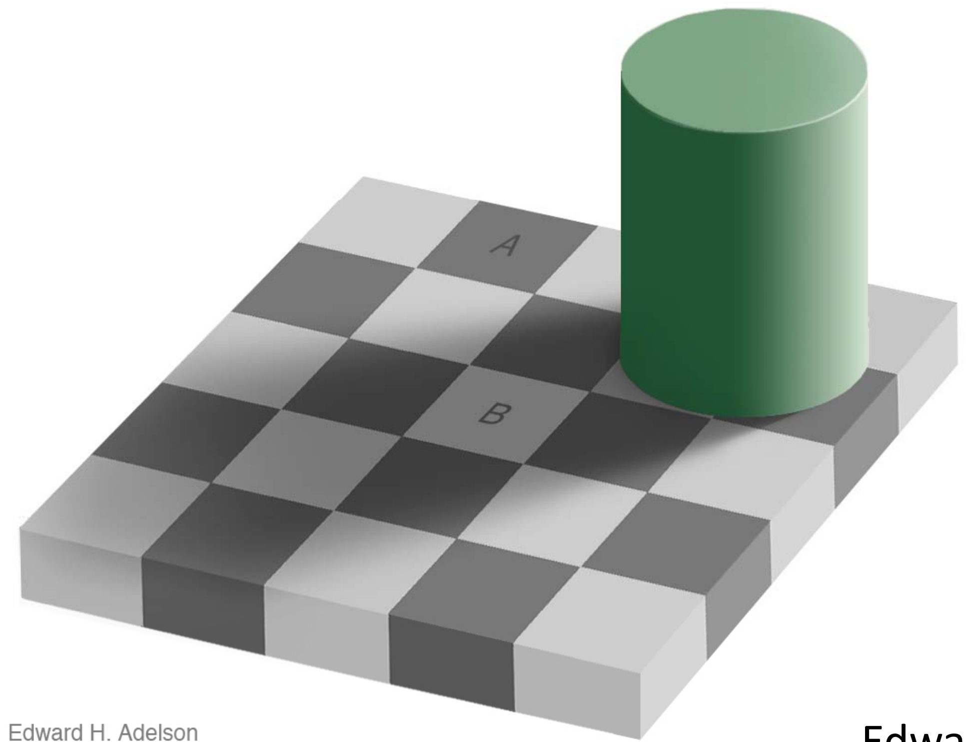

Edward Adelson

http://web.mit.edu/persci/people/adelson/illusions_demos.html

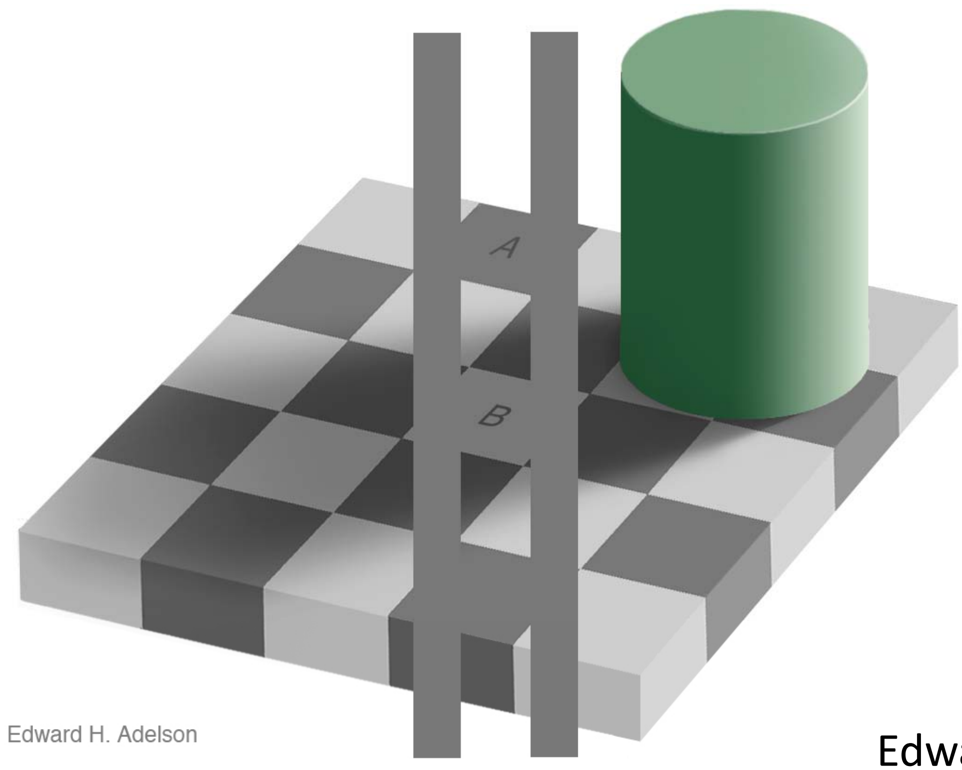

Edward Adelson

http://web.mit.edu/persci/people/adelson/illusions_demos.html

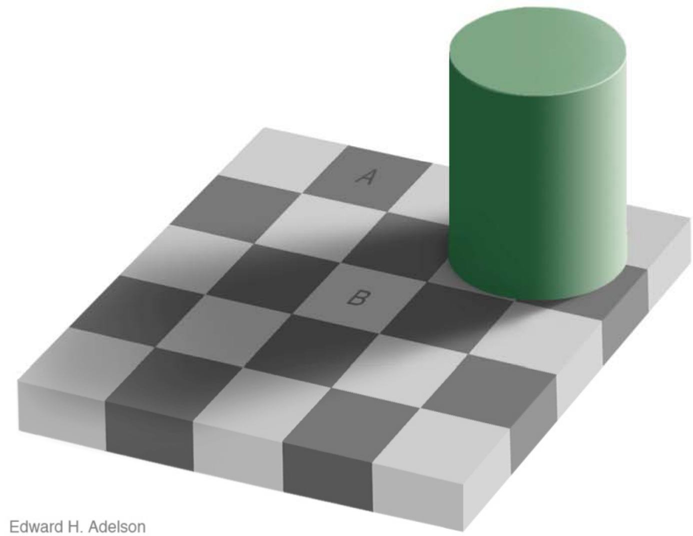

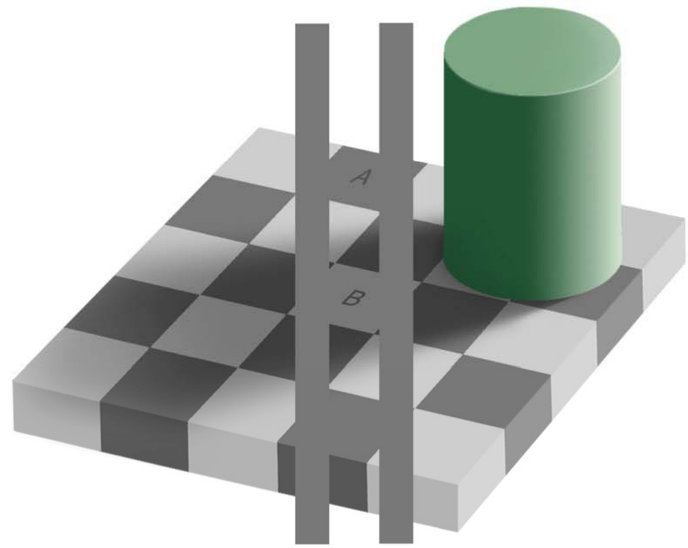

Edward Adelson

http://web.mit.edu/persci/people/adelson/illusions_demos.html

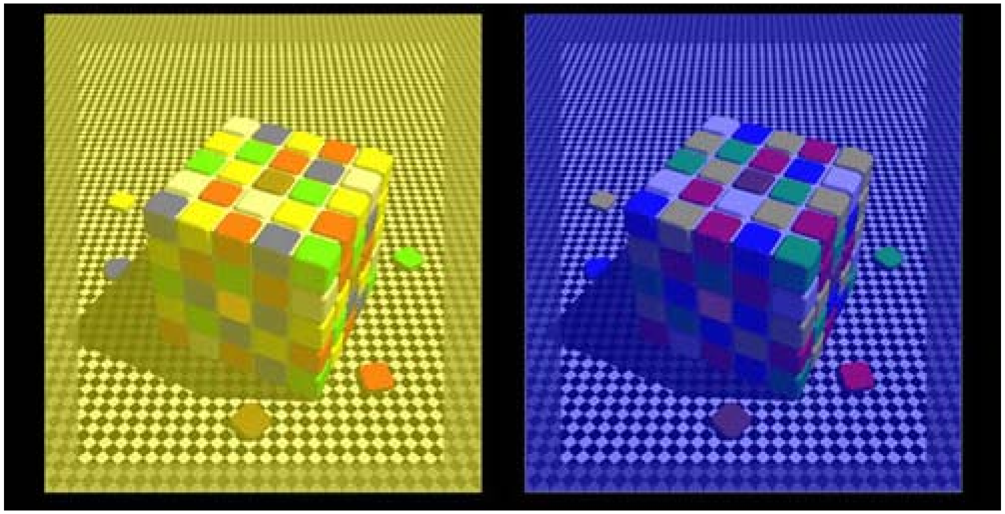

Look at blue squares

Look at yellow squares

-
- http://www.lottolab.org/articles/illusionsoflight.asp

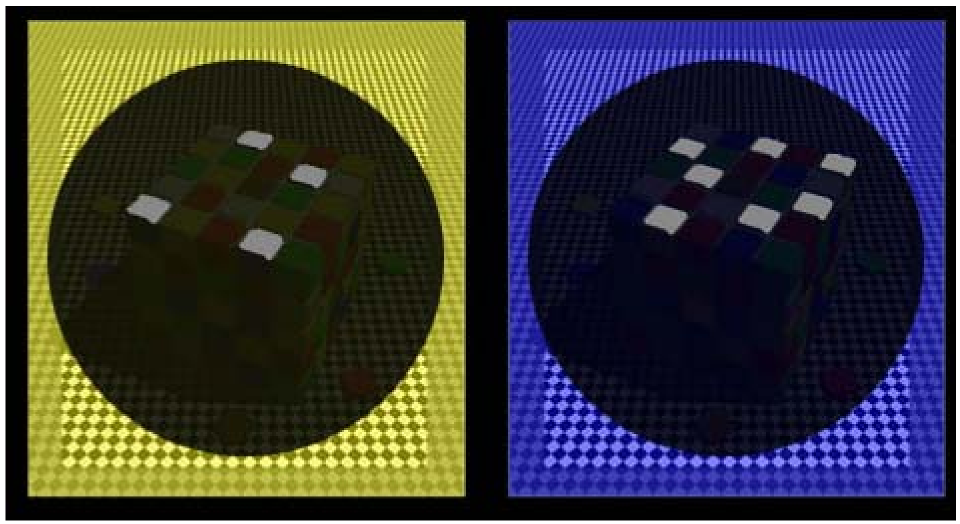

- Content © 2008 R.Beau Lotto
- http://www.lottolab.org/articles/illusionsoflight.asp

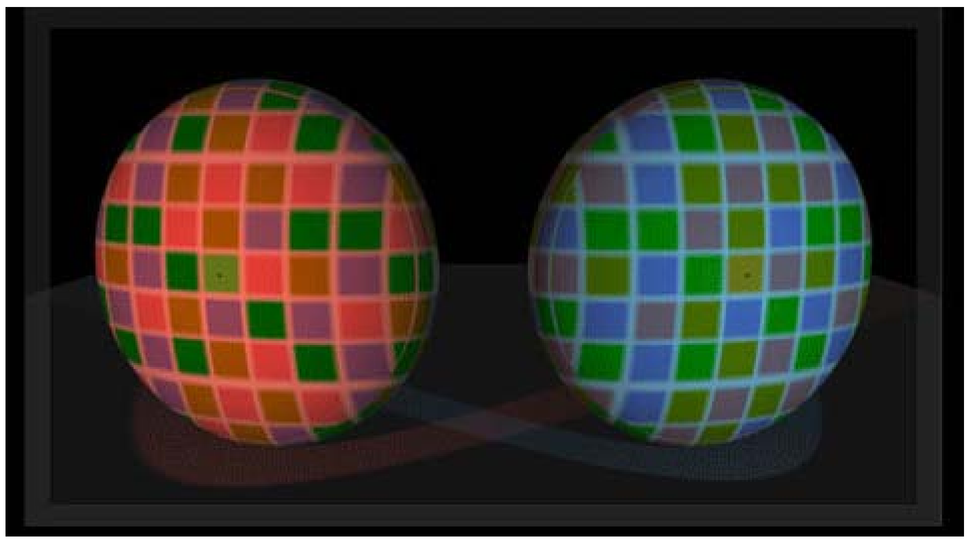

- Content © 2008 R.Beau Lotto
- http://www.lottolab.org/articles/illusionsoflight.asp

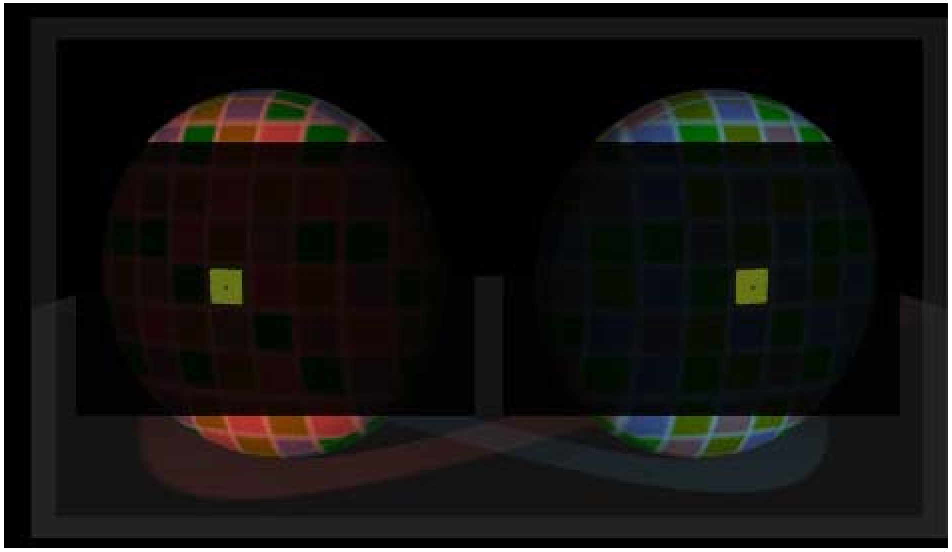

- Content © 2008 R.Beau Lotto
- http://www.lottolab.org/articles/illusionsoflight.asp

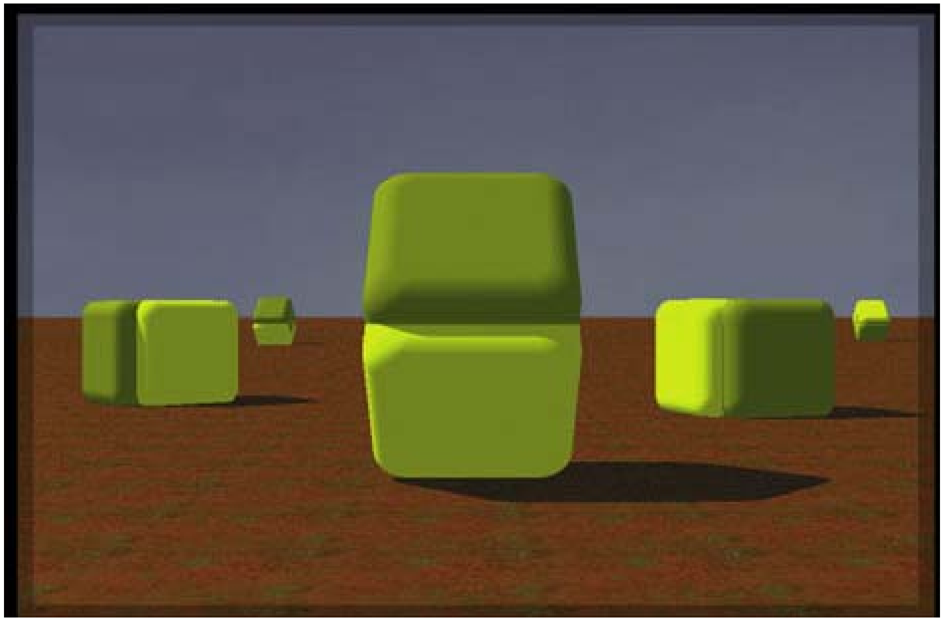

- Content © 2008 R.Beau Lotto
- http://www.lottolab.org/articles/illusionsoflight.asp

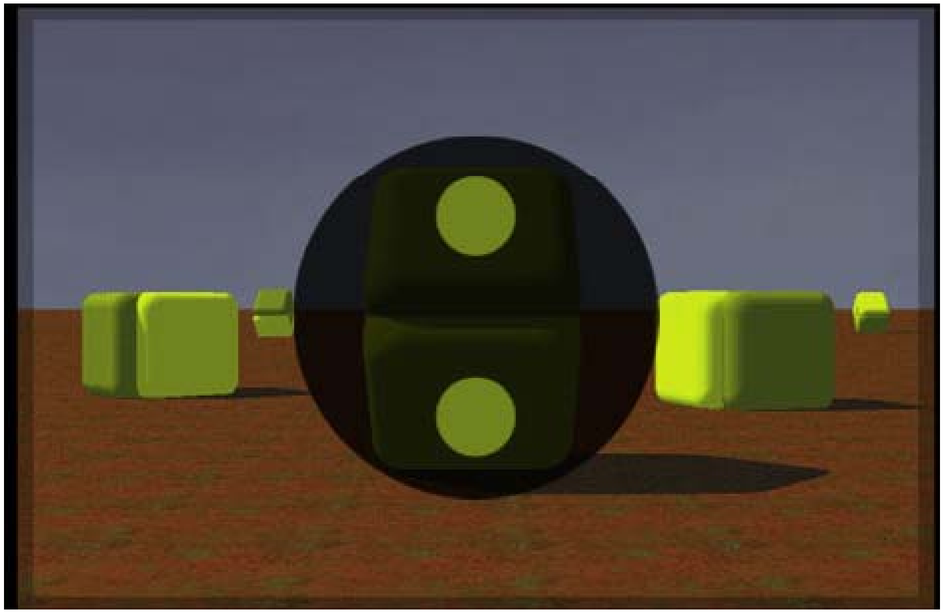

- Content © 2008 R.Beau Lotto
- http://www.lottolab.org/articles/illusionsoflight.asp

Contrast effects

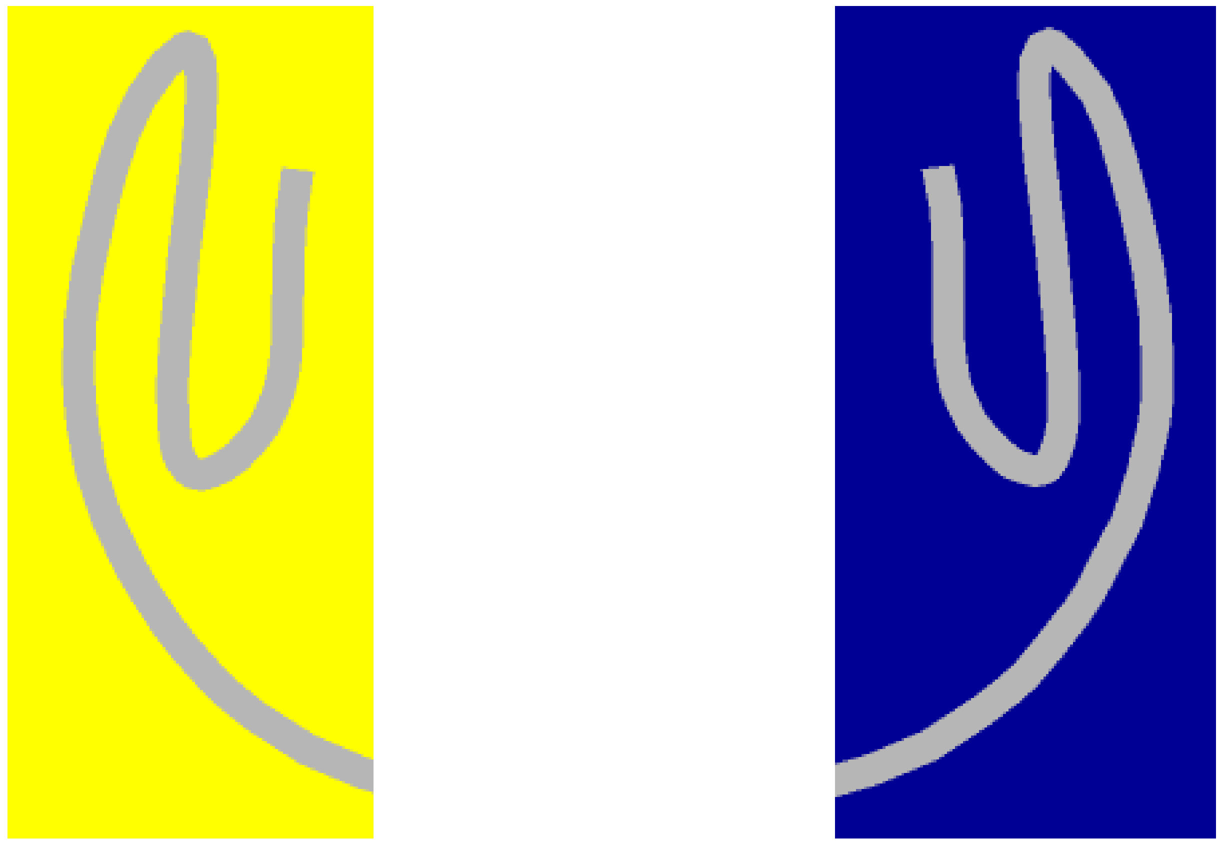

After images

- Tired photoreceptors send out negative response after a strong stimulus

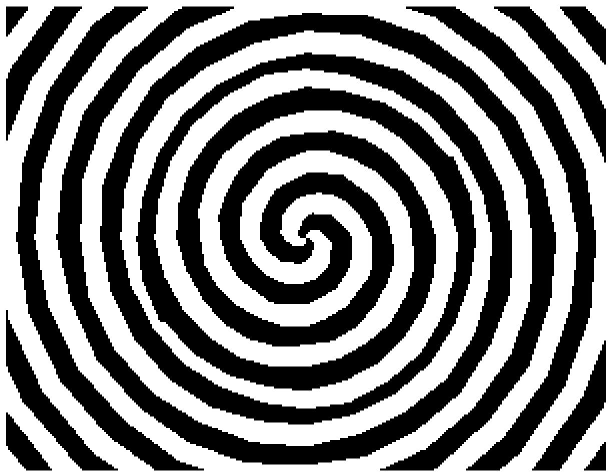

http://www.sandlotscience.com/Aftereffects/Andrus_Spiral.htm

Source: Steve Seitz

Name that color

Blue Red Green Cyan
Magenta Black Pink
Yellow Orange Violet
Brown Purple Cyan
Indigo Red Green Blue

High level interactions affect perception and processing.

Yellow Text on a Blue Background

- Is fairly easy to read unless the text is isoluminant with the background colour. As the luminance of the background becomes the same as the luminance of the text, it is very difficult to make out what the text says. So much so, that at this point I can write just about anything I want here and hardly anyone would want to put in the effort to see what it was I had written.

http://people.cs.ubc.ca/~tmm/courses/cpsc533c-03-spr/0203.rodmcfarland.ppt

Other isoluminance effects

- Stereoscopic depth is not detectable with isoluminant colours
- Isoluminance in animation makes it appear to be slower than the same animation in black-and-white
- Shape and form are best shown using luminance:

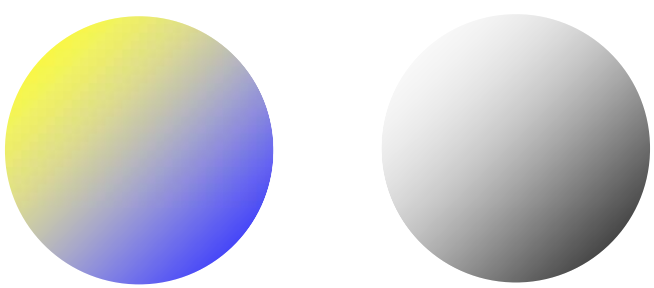

http://people.cs.ubc.ca/~tmm/courses/cpsc533c-03-spr/0203.rodmcfarland.ppt

“Two pathways”

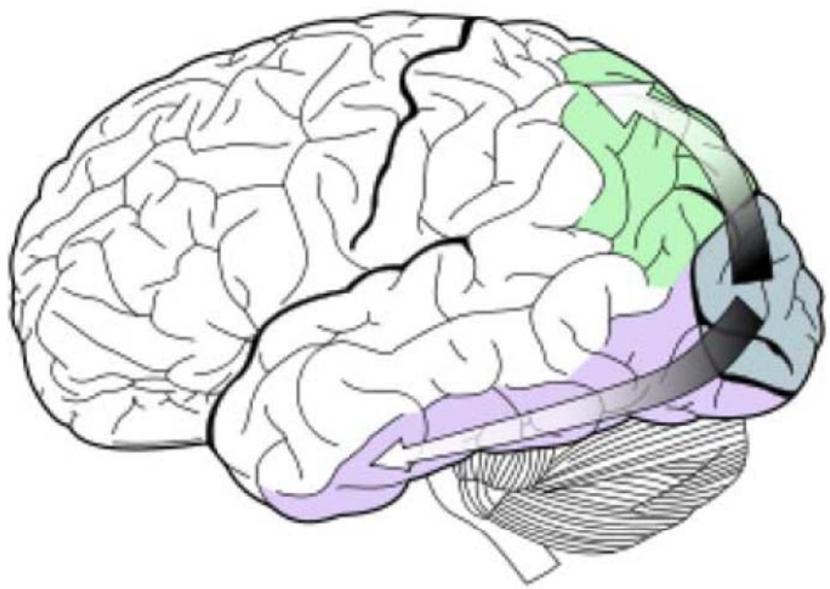

- Visual pathways in the human brain. The ventral stream(purple) is important in color recognition. The dorsal stream(green) is also shown. They originate from a common source in the visual cortex.
- Cliché: “what” vs “where”

en.wikipedia.org/wiki/Color_vision

Today: Color

- Measuring color
  - Spectral power distributions
  - Color mixing
  - Color matching experiments
  - Color spaces
    - Uniform color spaces

- Perception of color
  - Human photoreceptors
  - Environmental effects, adaptation

- Using color in machine vision systems

Color as a low-level cue for CBIR

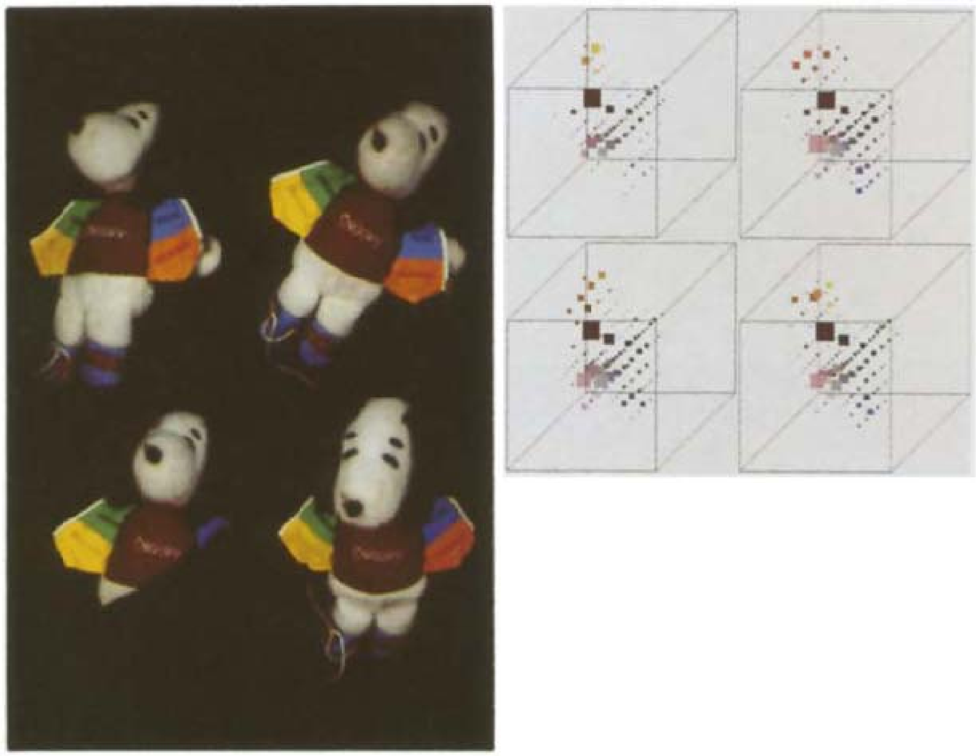

Swain and Ballard, Color Indexing, IJCV 1991

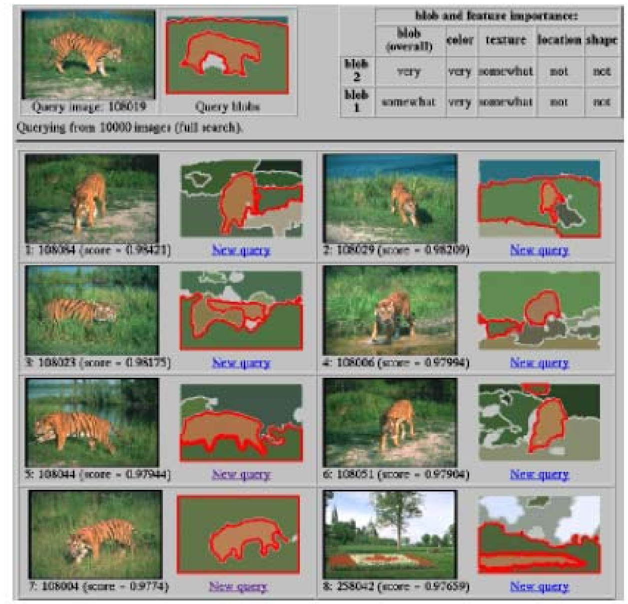

Blobworld system
Carson et al, 1999

Color as a low-level cue for CBIR

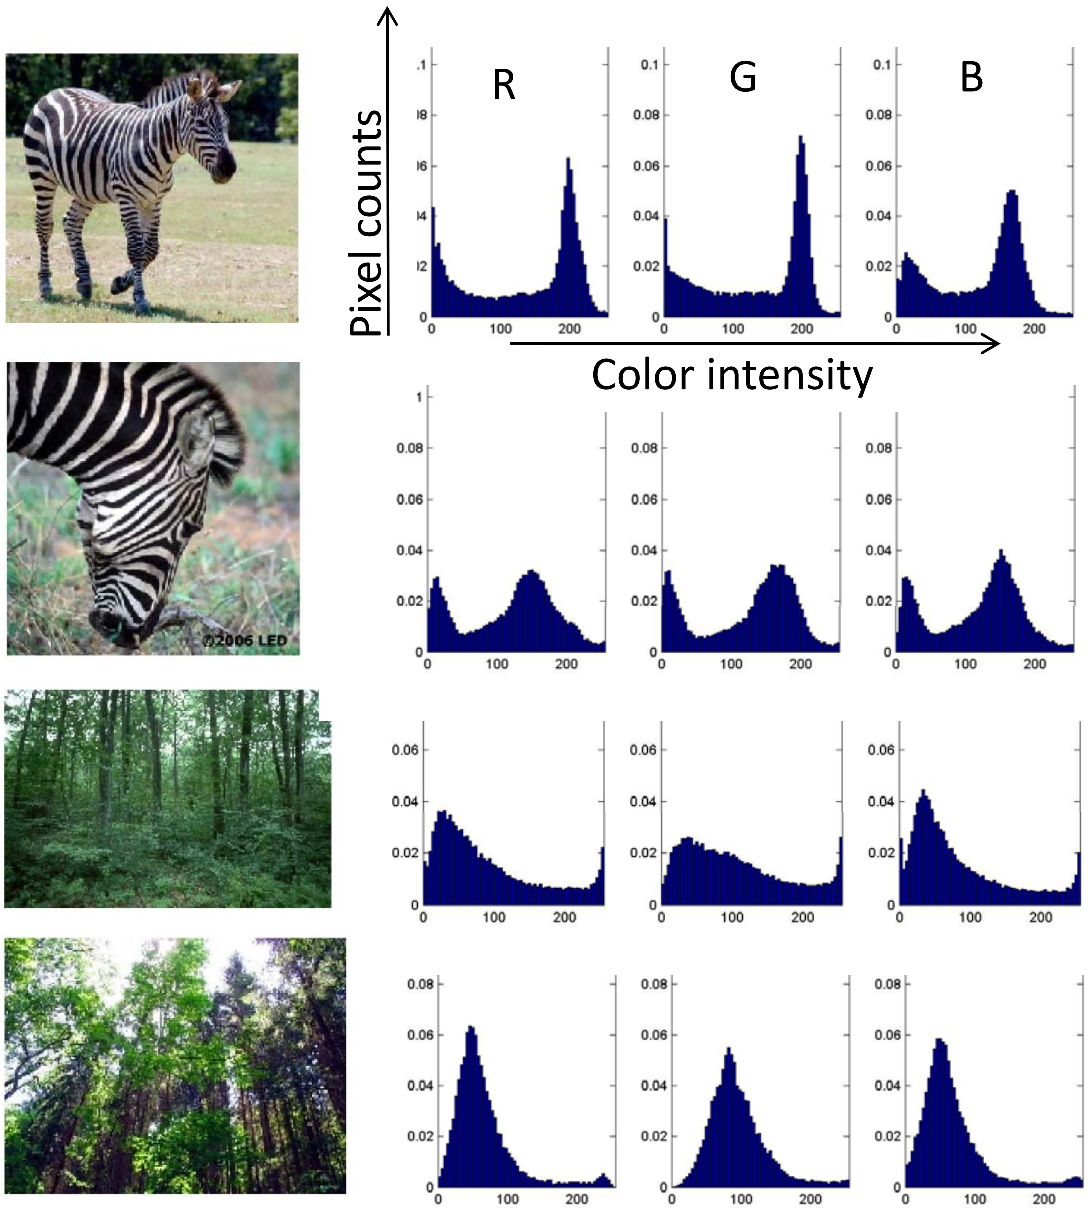

- Color histograms: Use distribution of colors to describe image
- No spatial info – invariant to translation, rotation, scale

Color as a low-level cue for CBIR

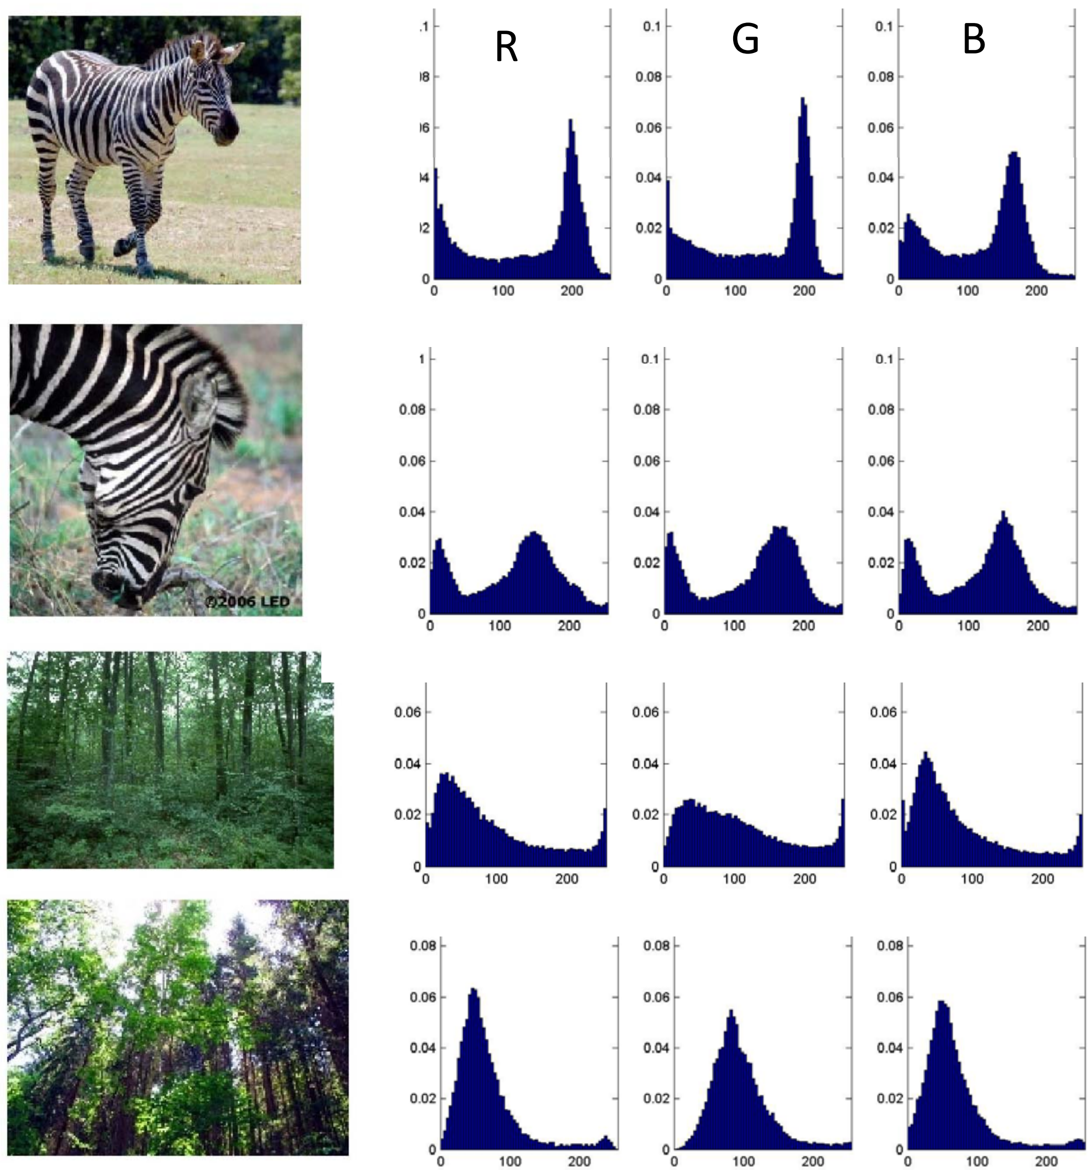

Given two histogram vectors, sum the minimum counts per bin:

$$I(x, y) = \sum_{i=1}^{n} \min(x_i, y_i)$$

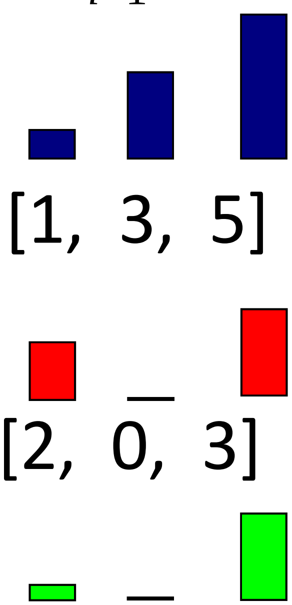

x = [1, 3, 5]

y = [2, 0, 3]

[1, 0, 3]

$\sum_i \min(x_i, y_i) = 4$

Color-based image retrieval

- Given collection (database) of images:
  - Extract and store one color histogram per image
- Given new query image:
  - Extract its color histogram
  - For each database image:
    - Compute intersection between query histogram and database histogram
  - Sort intersection values (highest score = most similar)
  - Rank database items relative to query based on this sorted order

Color-based image retrieval

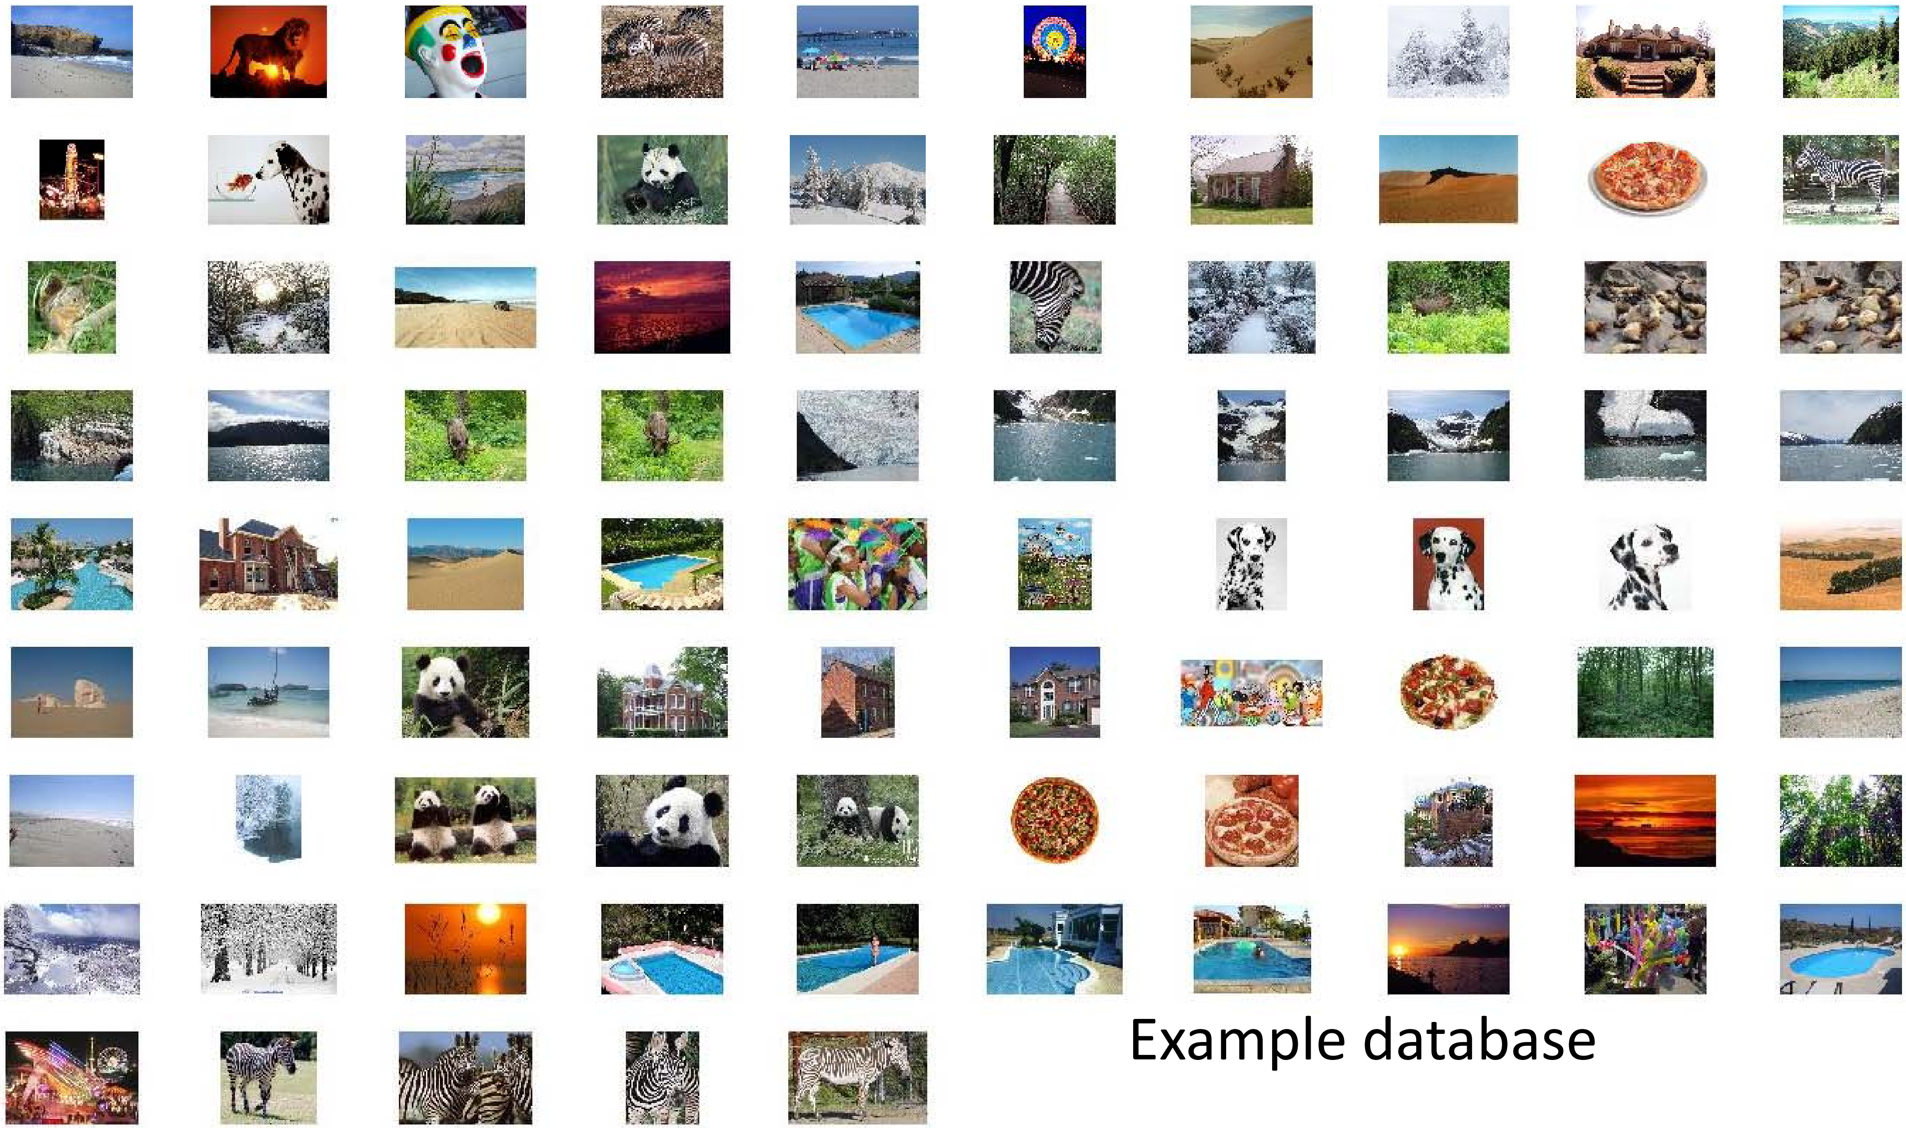

Example database

Color-based image retrieval

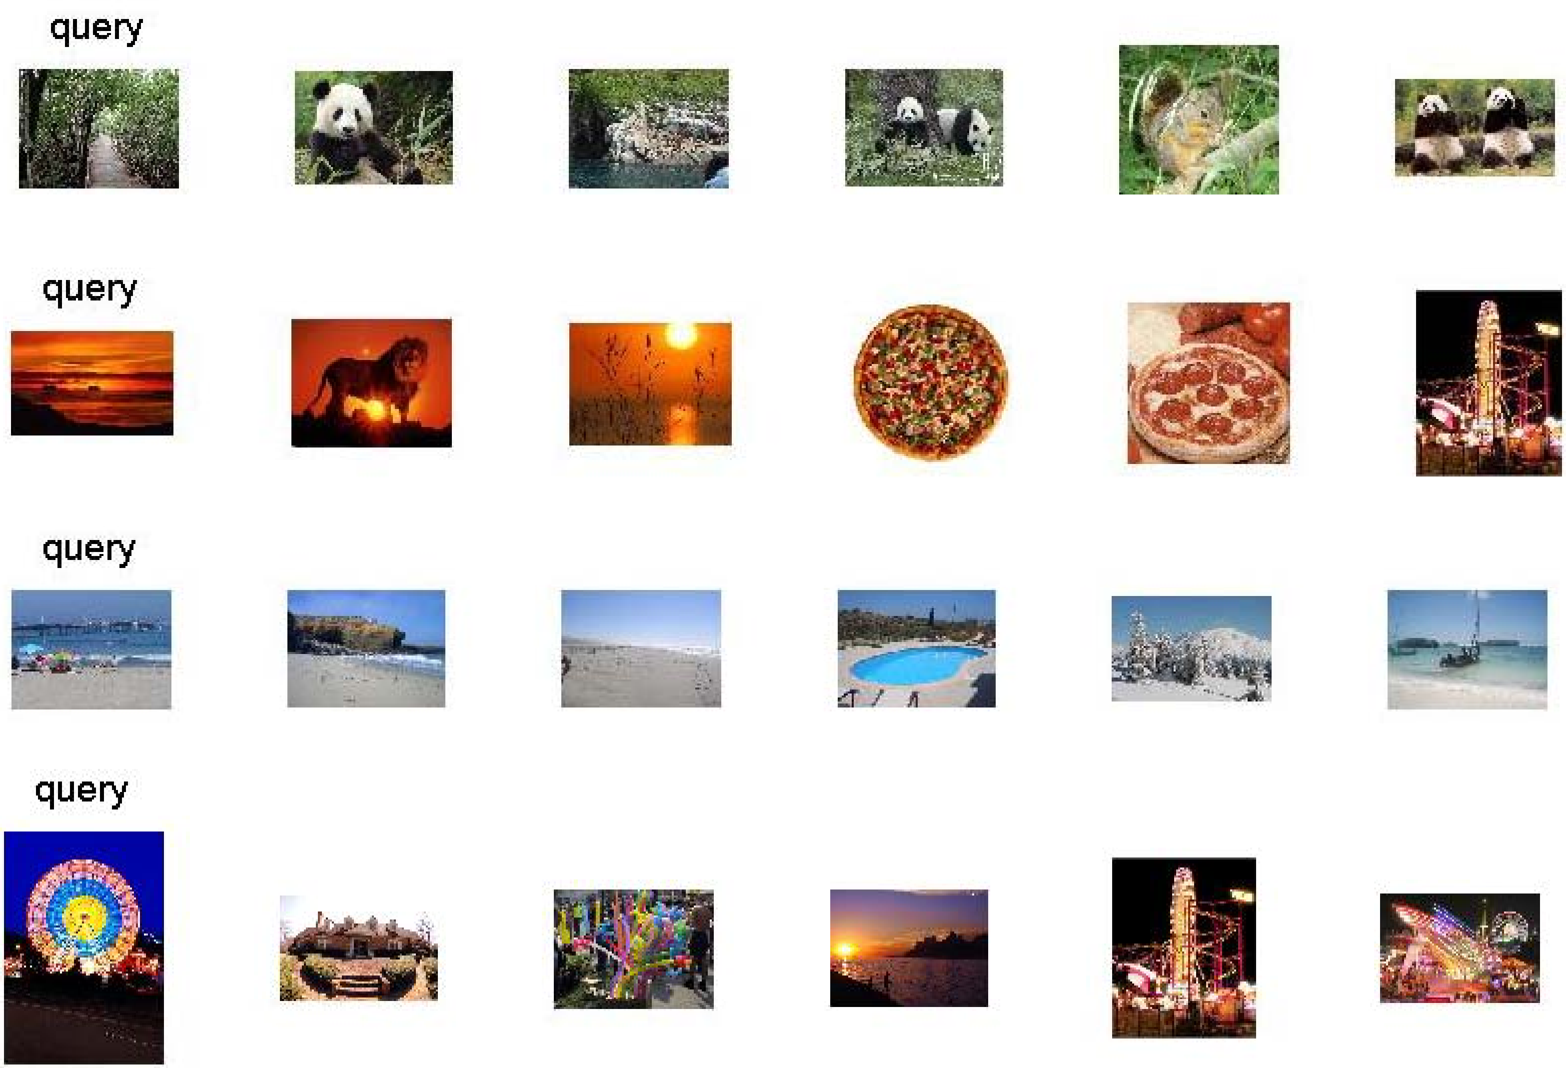

Example retrievals

Color-based image retrieval

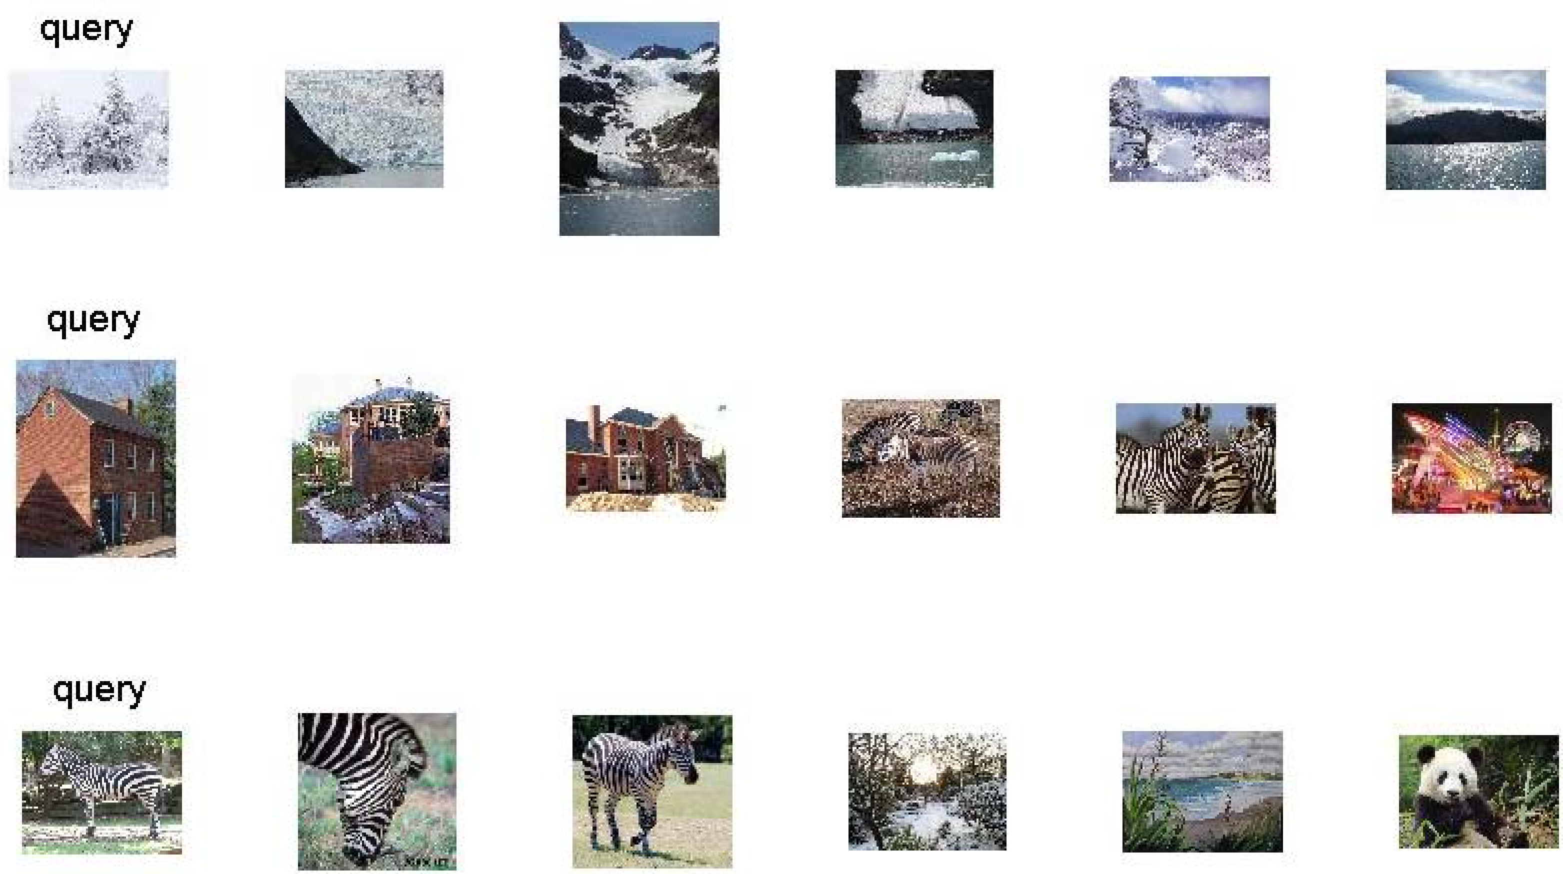

Example retrievals

Color-based skin detection

M. Jones and J. Rehg, Statistical Color Models with Application to Skin Detection, IJCV 2002.

Color-based segmentation for robot soccer

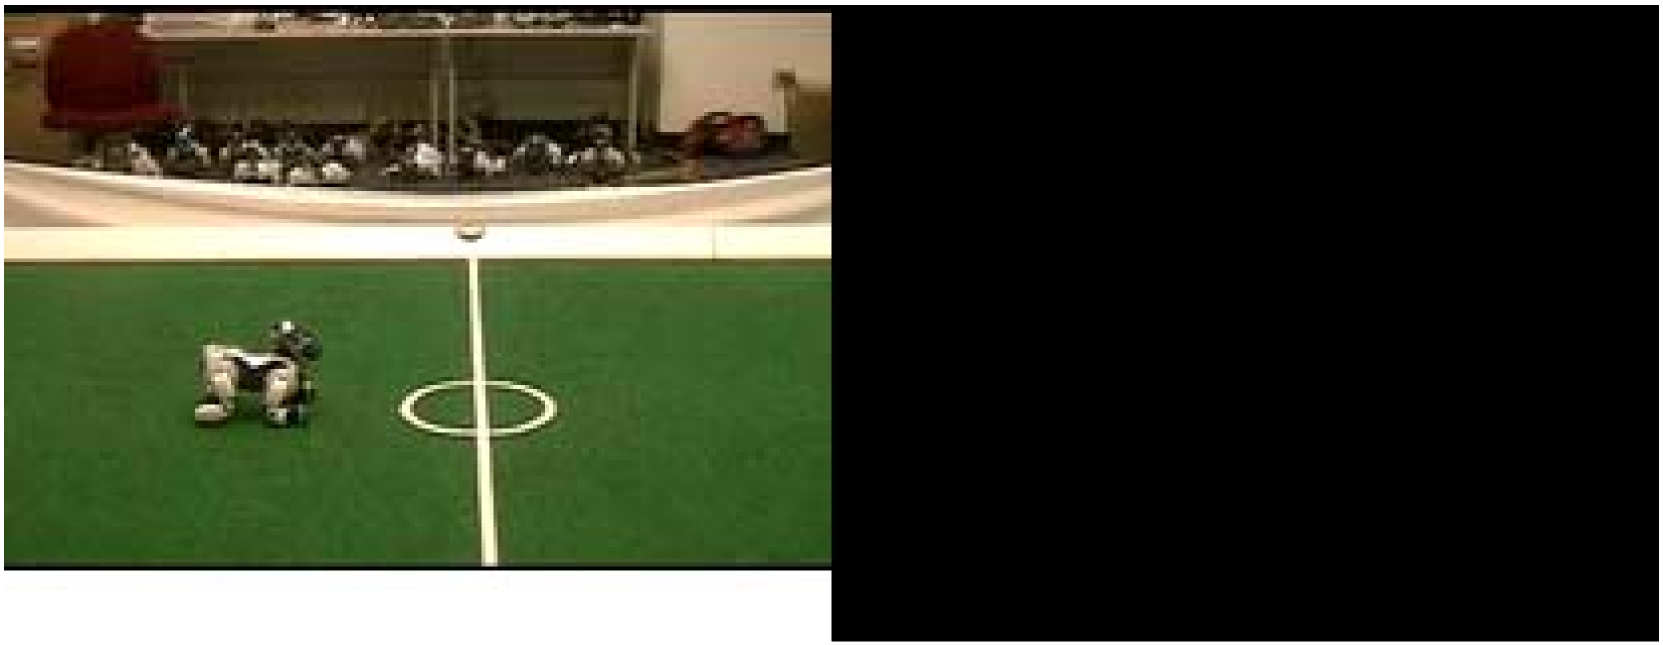

Towards Eliminating Manual Color Calibration at RoboCup. Mohan Sridharan and Peter Stone. RoboCup-2005: Robot Soccer World Cup IX, Springer Verlag, 2006

http://www.cs.utexas.edu/users/AustinVilla/?p=research/auto_vis

Color-based appearance models for body tracking

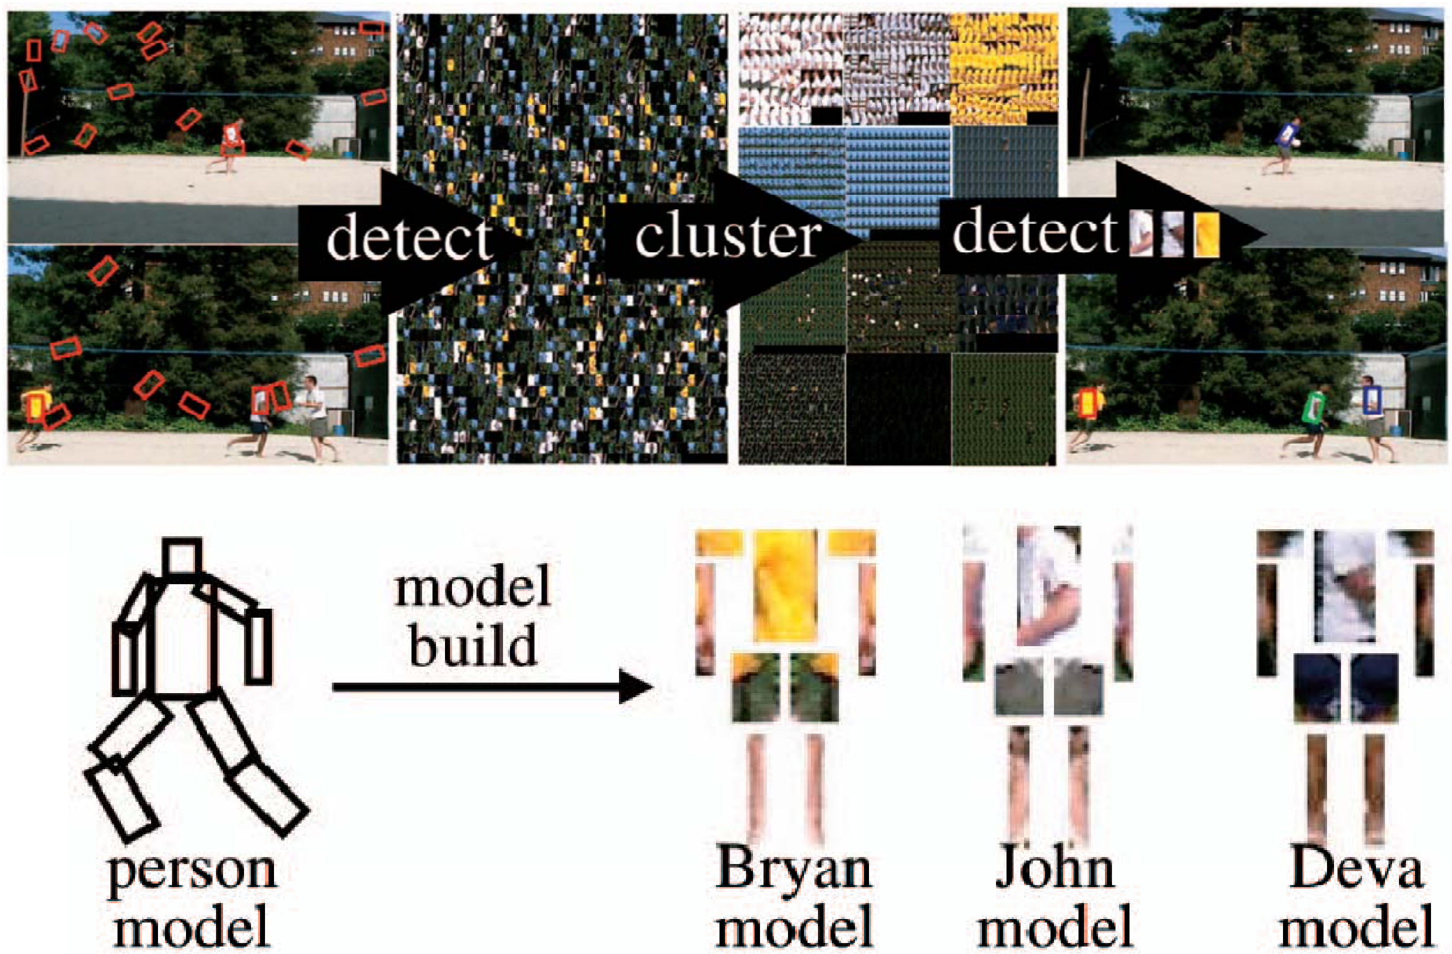

D. Ramanan, D. Forsyth, and A. Zisserman. Tracking People by Learning their Appearance. PAMI 2007.

Slide credit: L. Lazebnik

Viewing Colored Objects

- Assume diffuse (Lambertian) plus specular model
- Diffuse component
  - colour of reflected light depends on both illuminant and surface
- Specular component
  - specularities on dielectric (non-metalic) objects take the colour of the light
  - specularities on metals have colour of the metal

Finding Specularities

- Assume we are dealing with dielectrics
  — specularly reflected light is the same colour as the source
- Reflected light has two components
  — diffuse
  — specular
  — and we see a weighted sum of these two
- Specularities produce a characteristic “dog leg” in the histogram of receptor responses
  — in a patch of diffuse surface, we see a colour multiplied by different scaling constants (surface orientation)
  — in the specular patch, a new colour is added; a “dog leg” results

Skewed-T

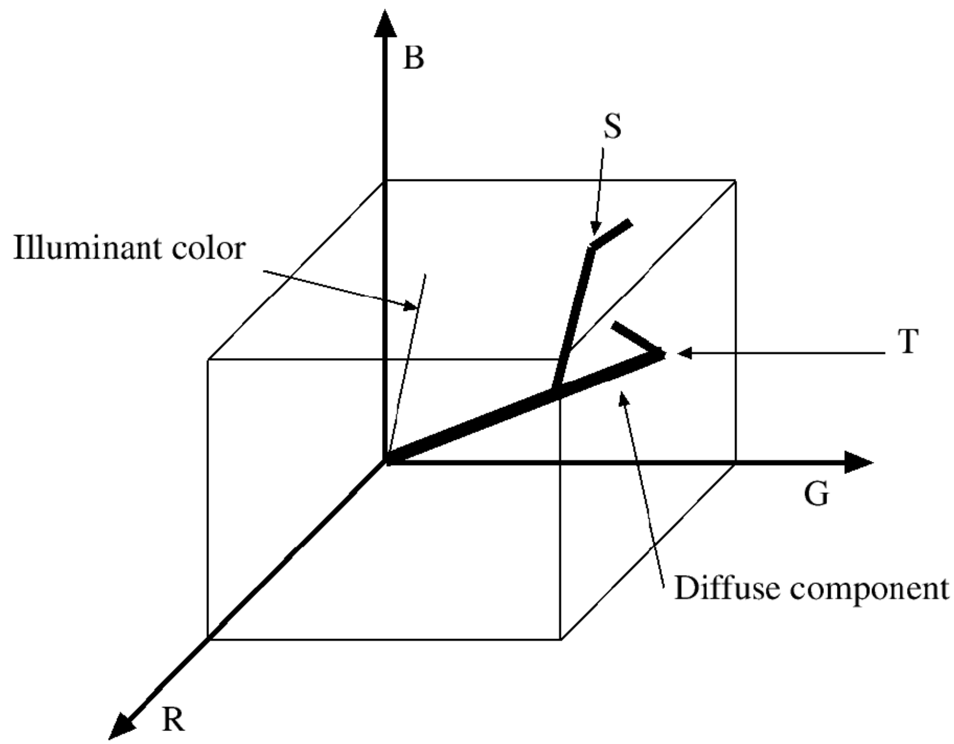

Forsyth & Ponce Figure 6.20

Skewed-T

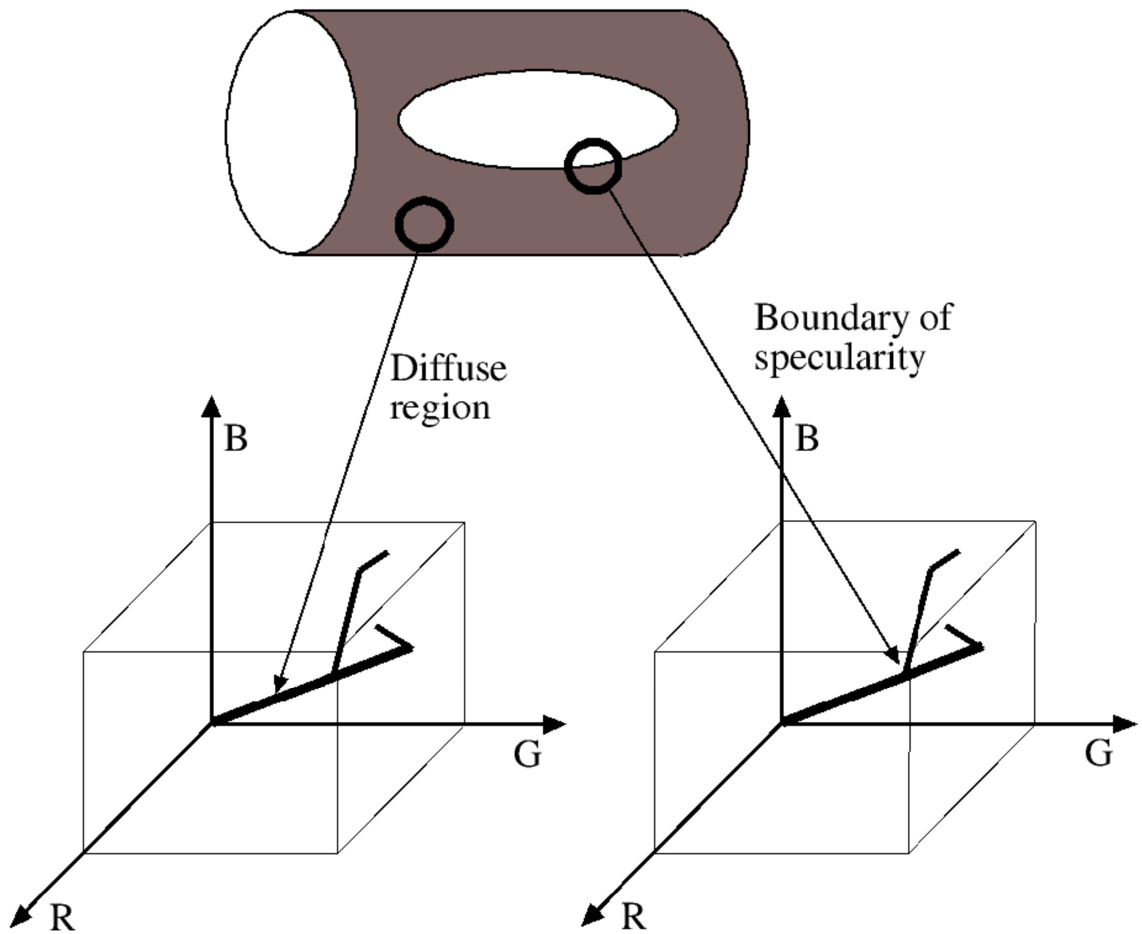

Forsyth & Ponce Figure 6.21

Slide Credits

- Kristen Grauman: 3-48, 50-75, 79-86
- Bob Woodham: 49, 87-90
- and others, indirectly (Steve Palmer, Brian Wandell, etc!)

Today: Color

- Measuring color
  - Spectral power distributions
  - Color mixing
  - Color matching experiments
  - Color spaces
    - Uniform color spaces

- Perception of color
  - Human photoreceptors
  - Environmental effects, adaptation

- Using color in machine vision systems

Next time

- ***Image Filtering***
  - Readings: Szel., Ch. 3.1-3.3
- Pset 0 due tomorrow, 5pm
  - download from http://www.box.net/shared/pvkynapd34
  - submit by email if desired
- Pset 1 available online on Monday
  - via bSpace and some mirror…