Color

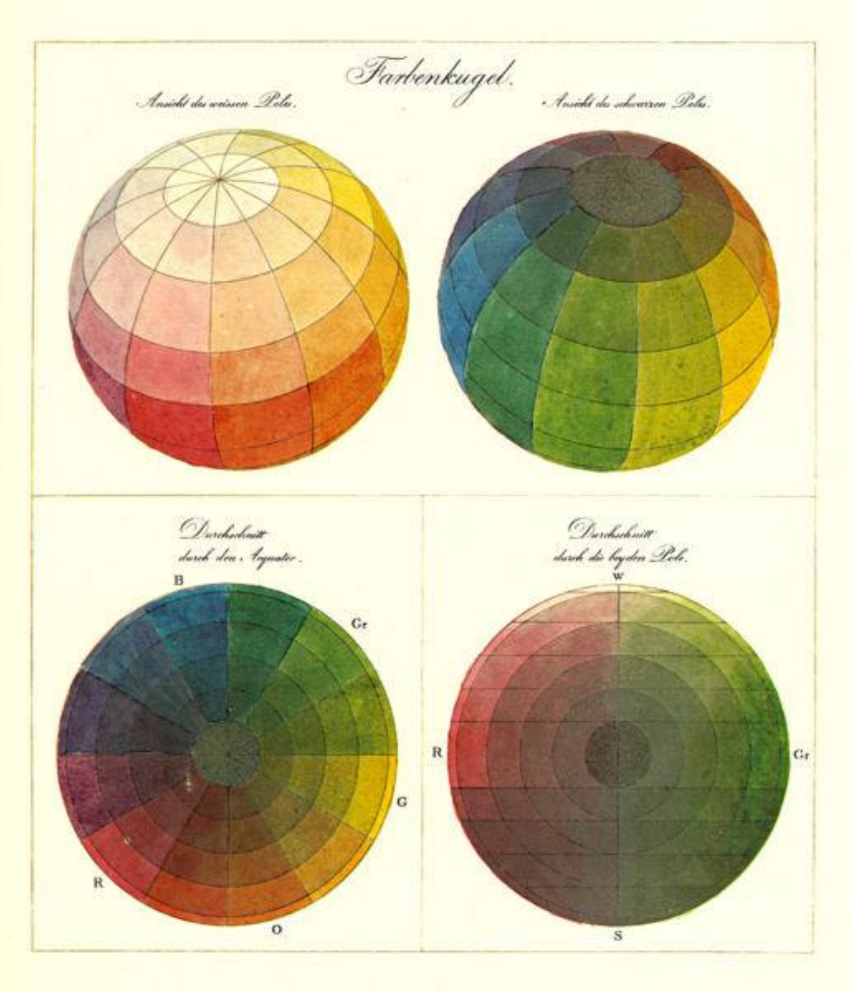

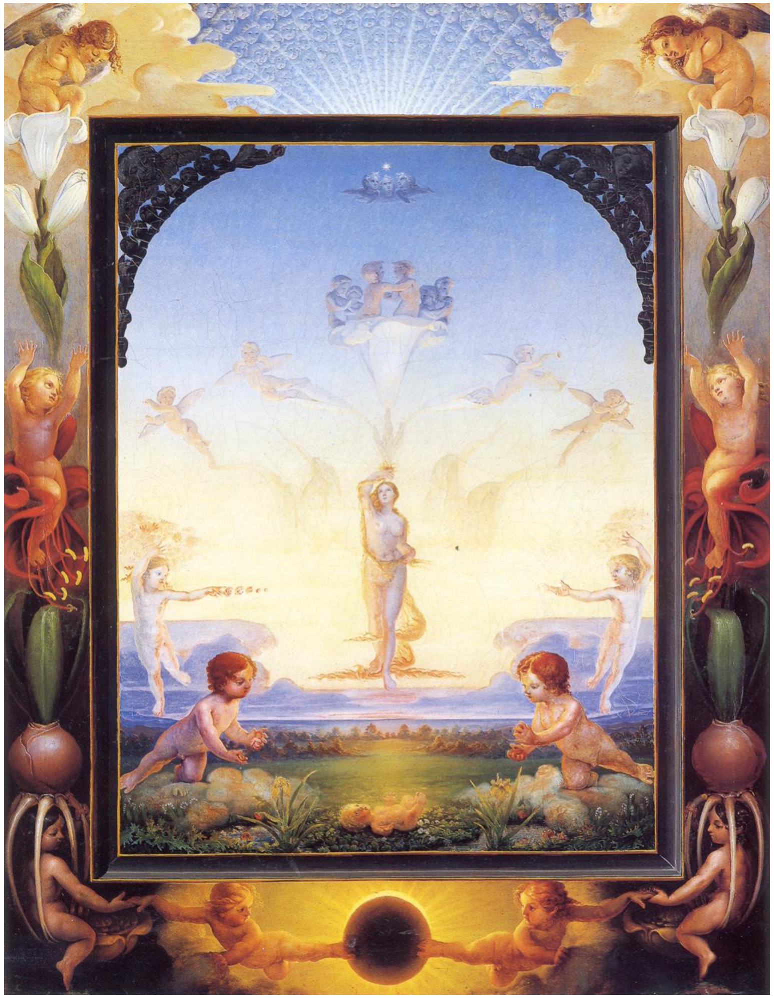

Phillip Otto Runge (1777-1810)

Review: Radiometry

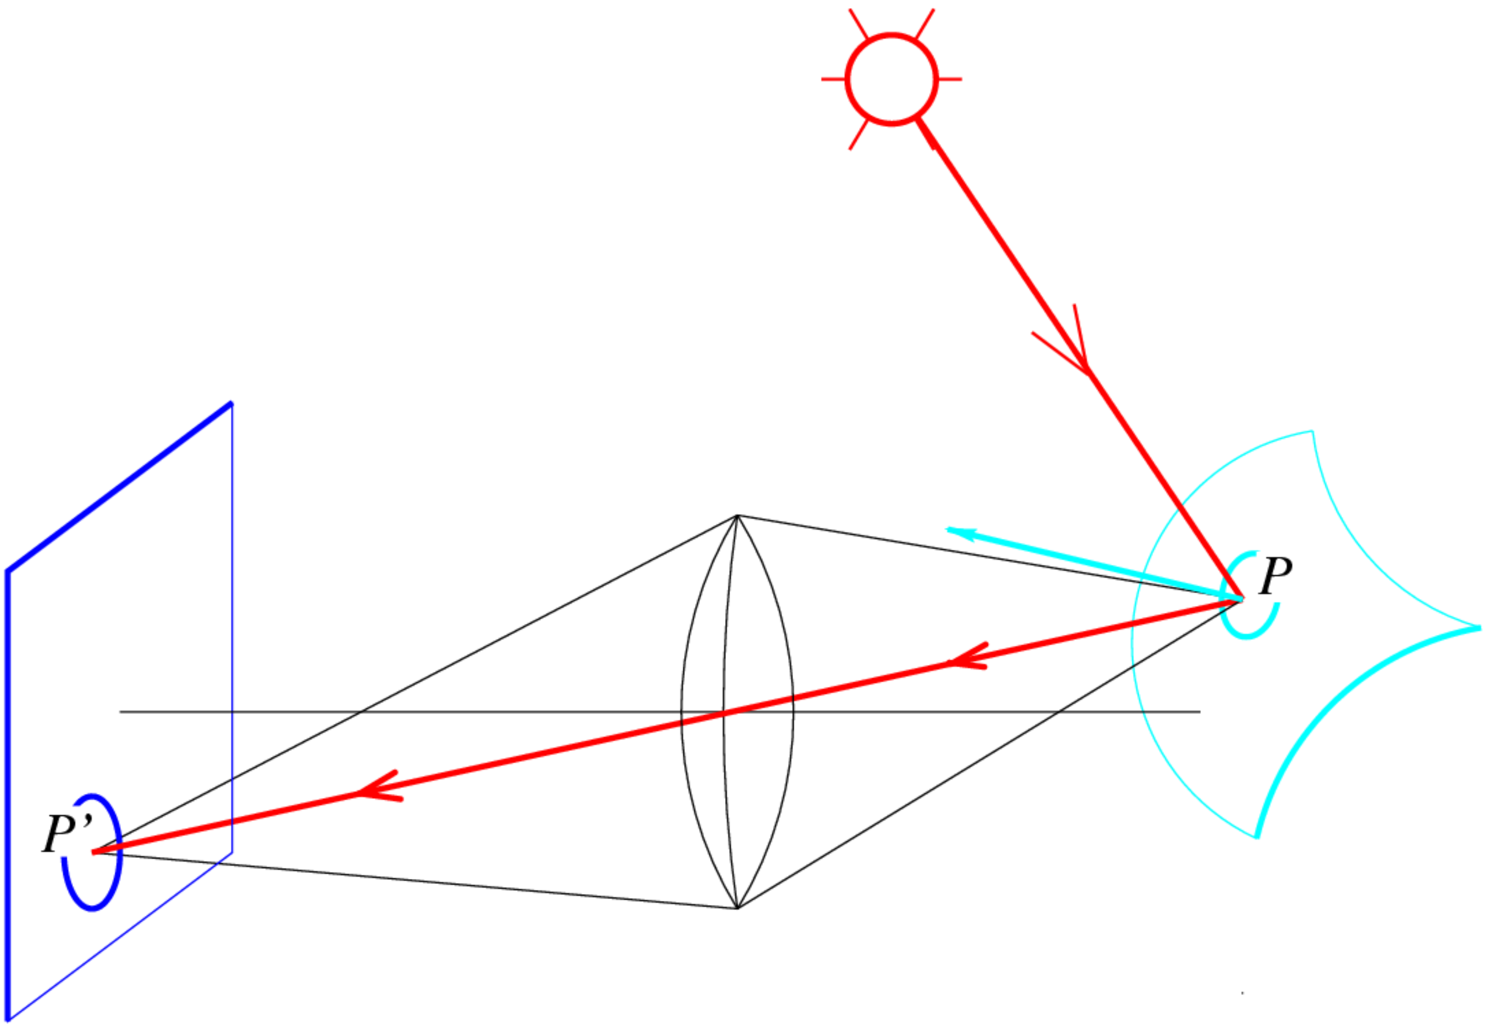

Slide by L. Fei-Fei

Review: Shape from shading

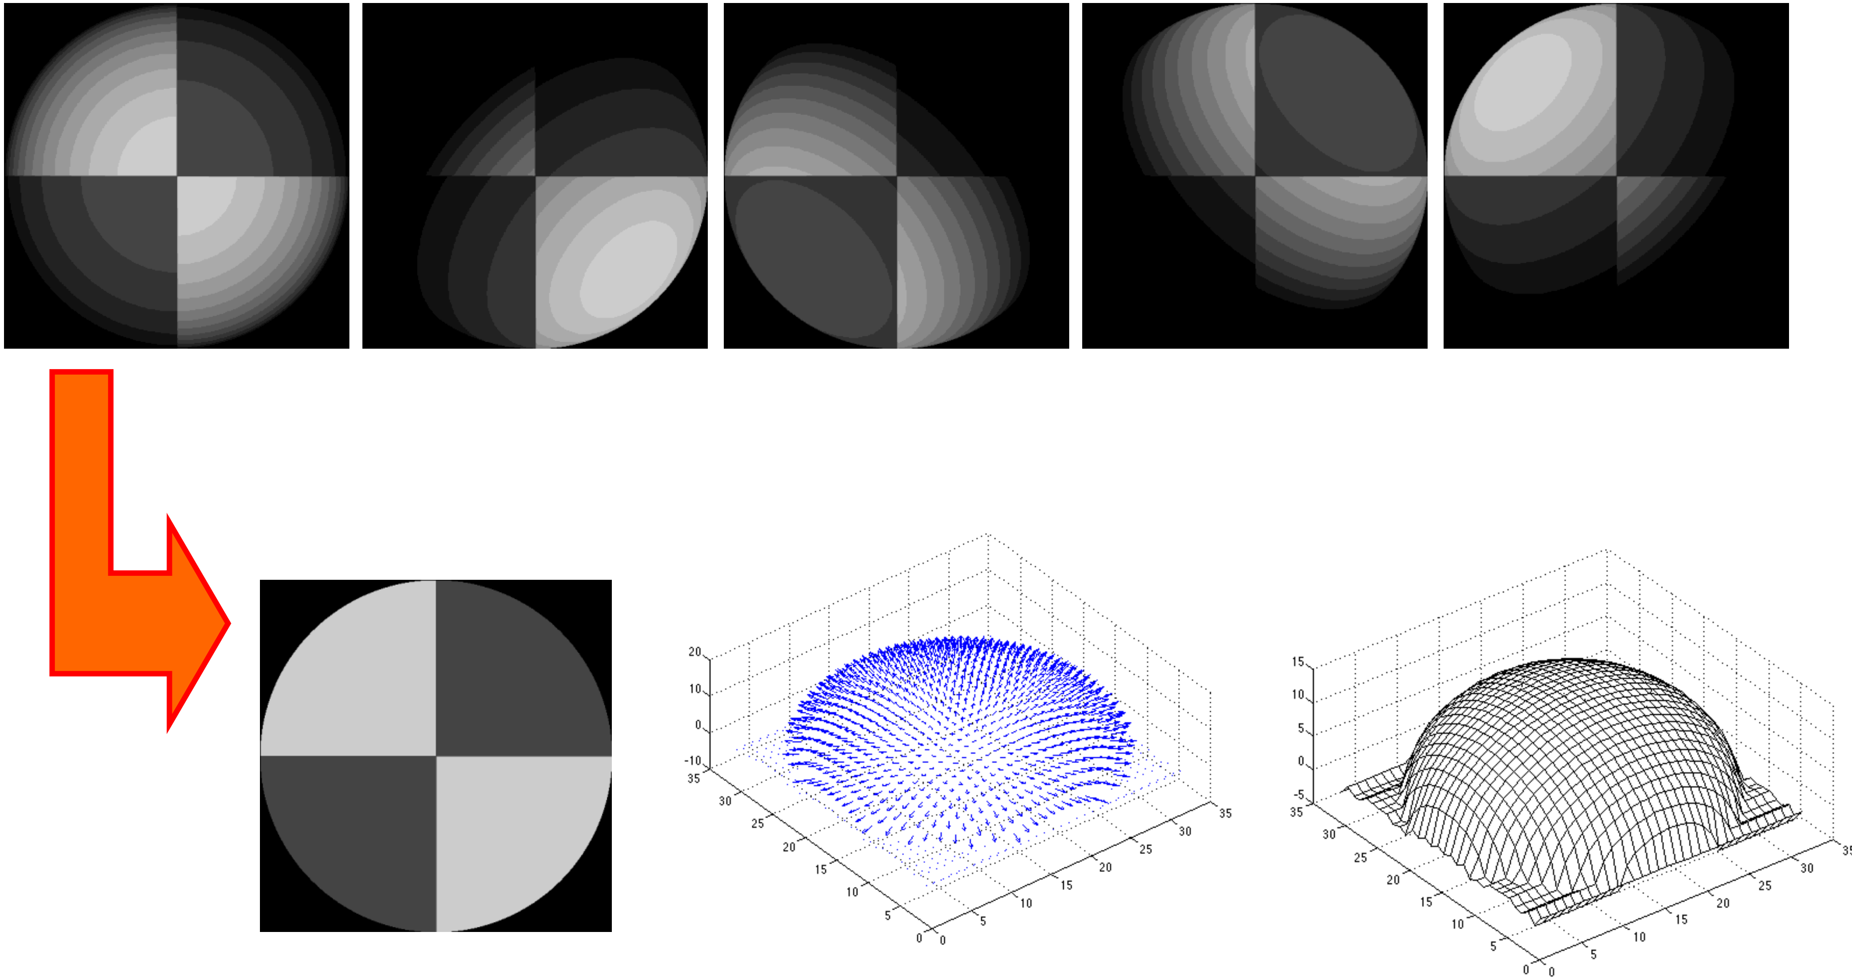

Forsyth & Ponce, Sec. 5.4

Today: Color

- The nature of color
- Color processing in the human visual system
- Color spaces
- Adaptation and constancy
- White balance
- Uses of color in computer vision

What is color?

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

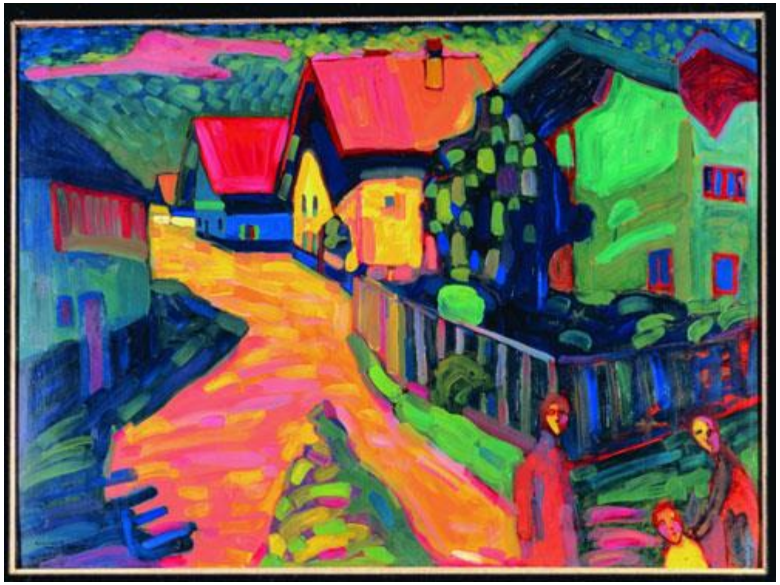

Wassily Kandinsky (1866-1944), Murnau Street with Women, 1908

Electromagnetic spectrum

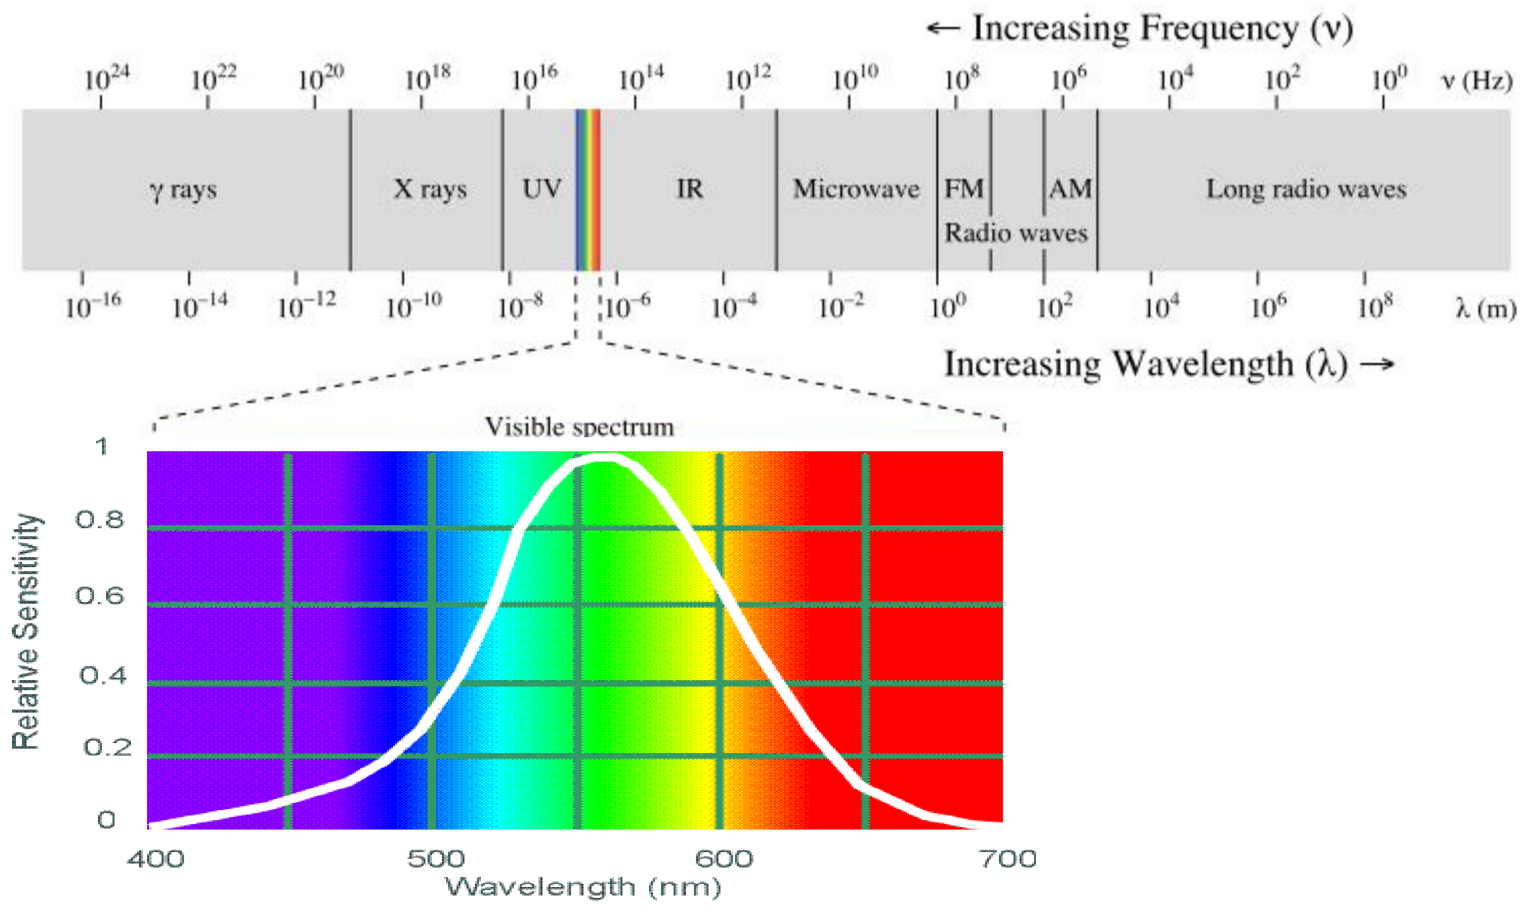

Human Luminance Sensitivity Function

Why do we see light at these wavelengths?

Because that's where the sun radiates electromagnetic energy

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.

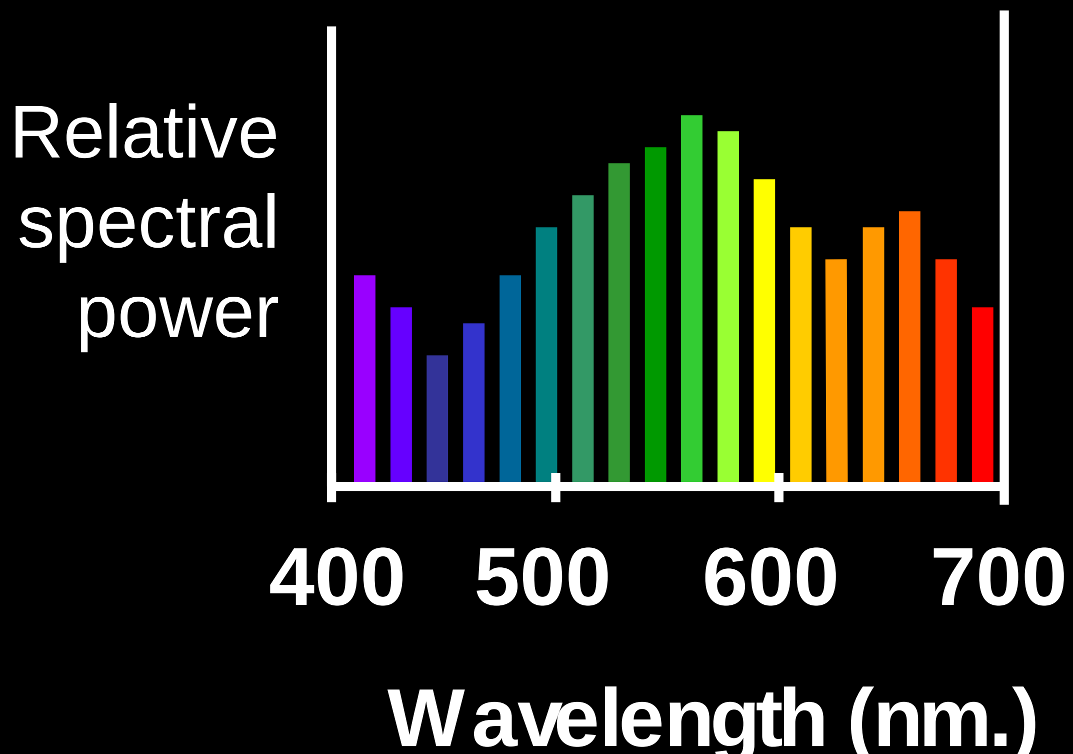

The Physics of Light

Some examples of the spectra of light sources

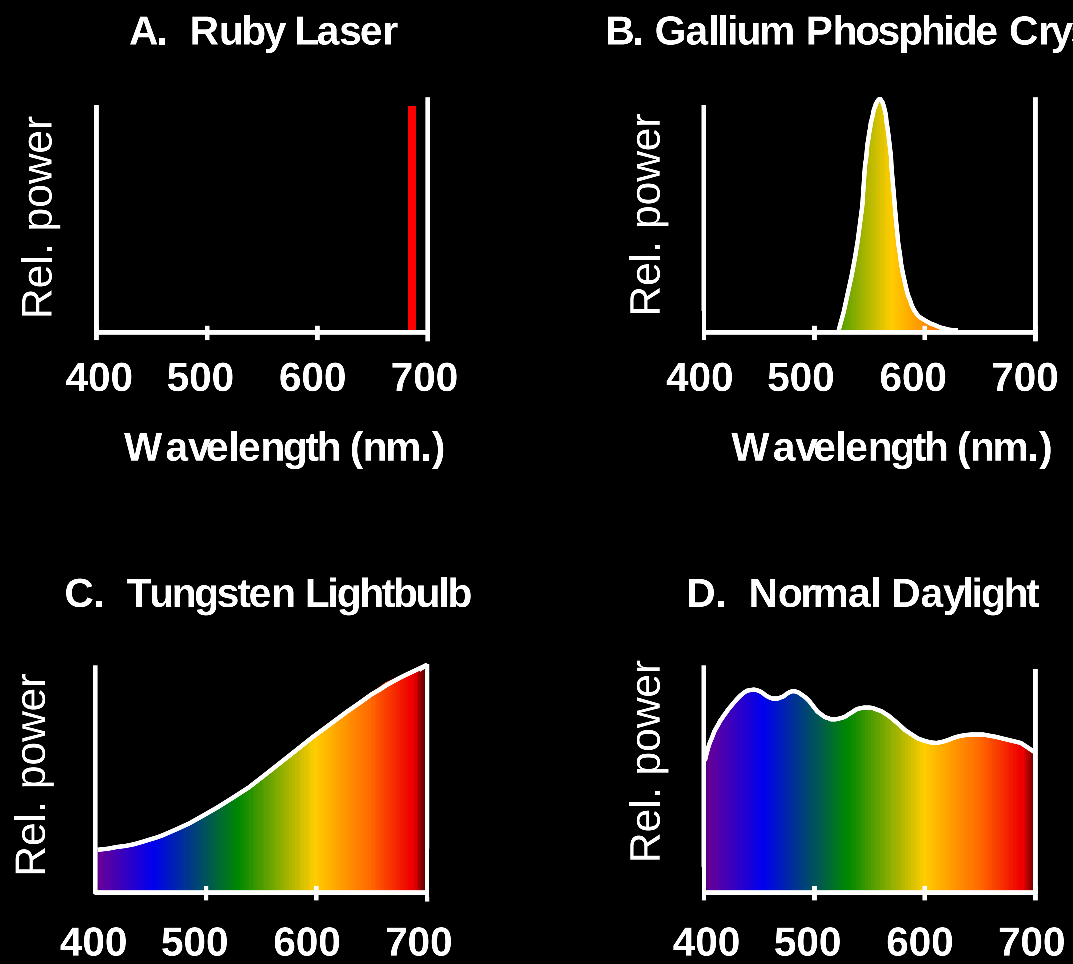

The Physics of Light

Some examples of the <u>reflectance</u> spectra of <u>surfaces</u>

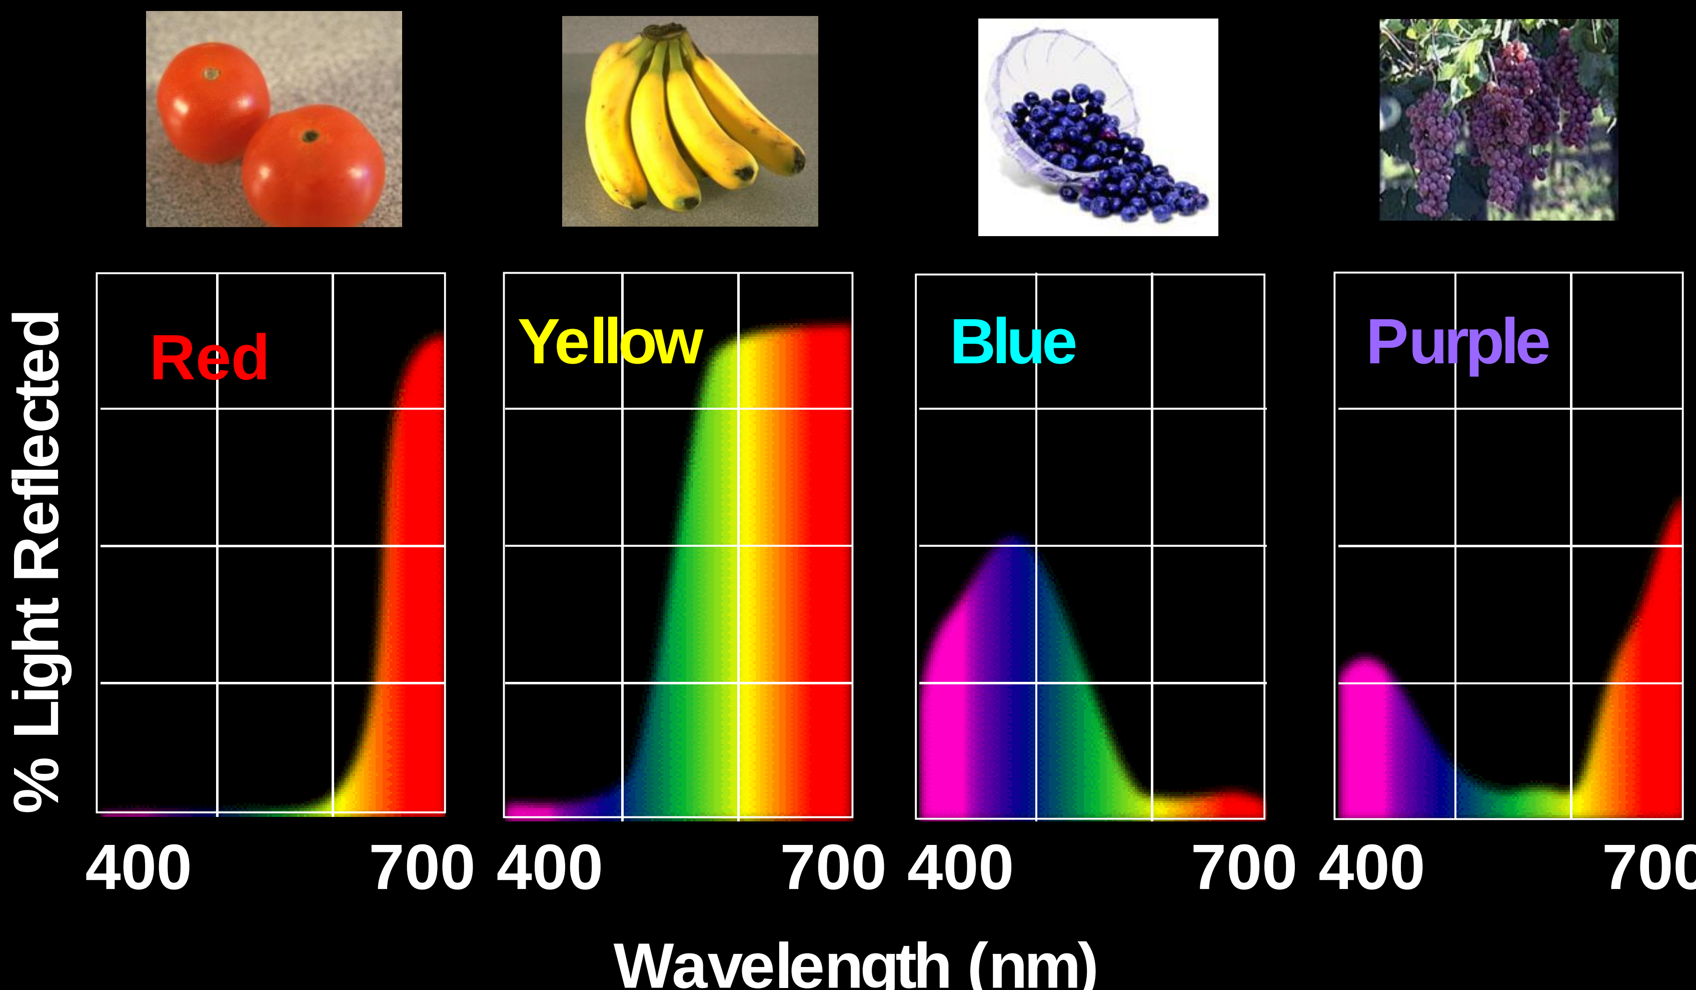

Interaction of light and surfaces

- Observed color is the result of interaction of light source spectrum with surface reflectance
- Spectral radiometry
  - All definitions and units are now "per unit wavelength"
  - All terms are now "spectral"

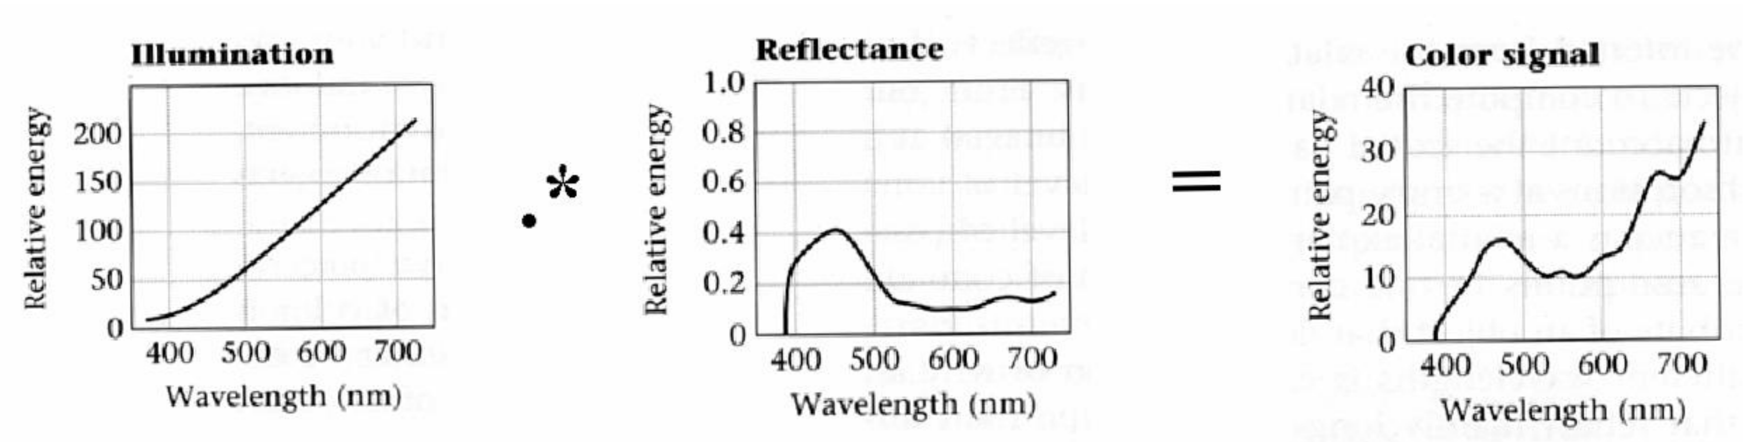

The Eye

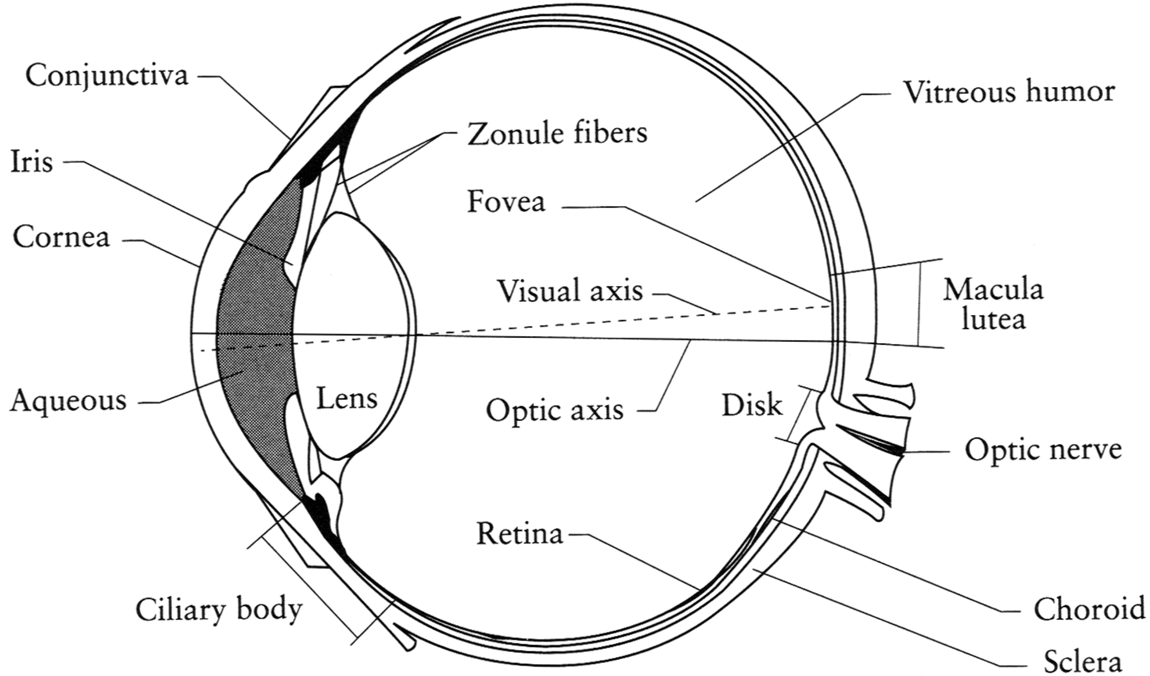

The human eye is a camera!

- **Iris** - colored annulus with radial muscles
- **Pupil** - the hole (aperture) whose size is controlled by the iris
- **Lens** - changes shape by using ciliary muscles (to focus on objects at different distances)
- What's the "film"?
  – photoreceptor cells (rods and cones) in the **retina**

Slide by Steve Seitz

Density of rods and cones

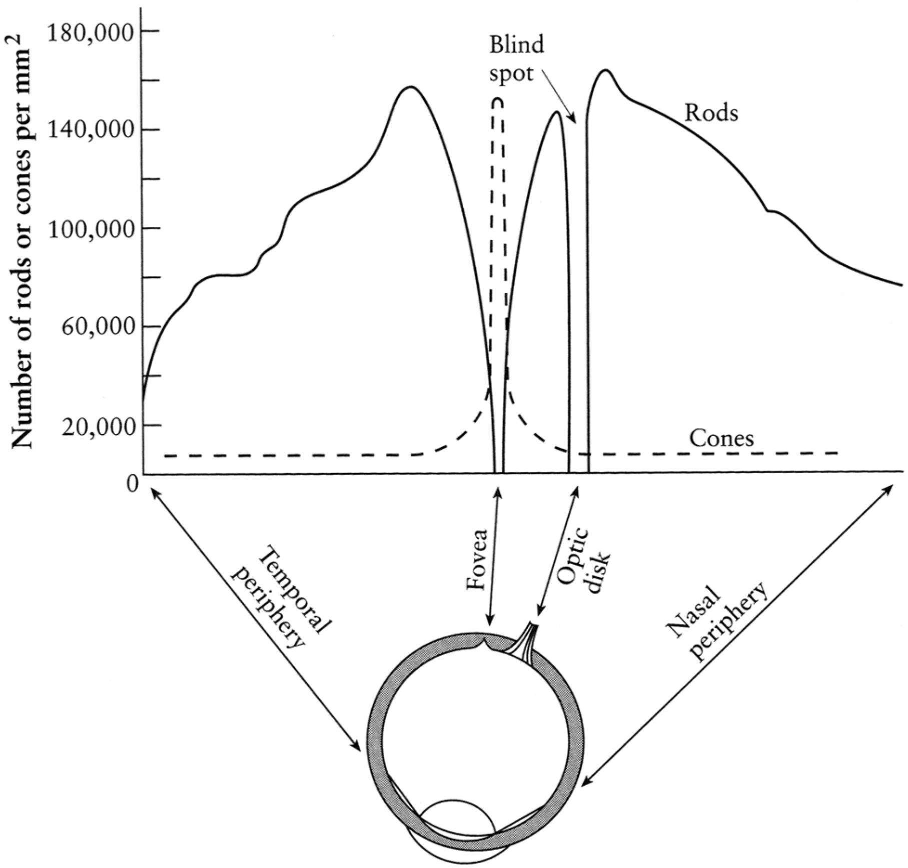

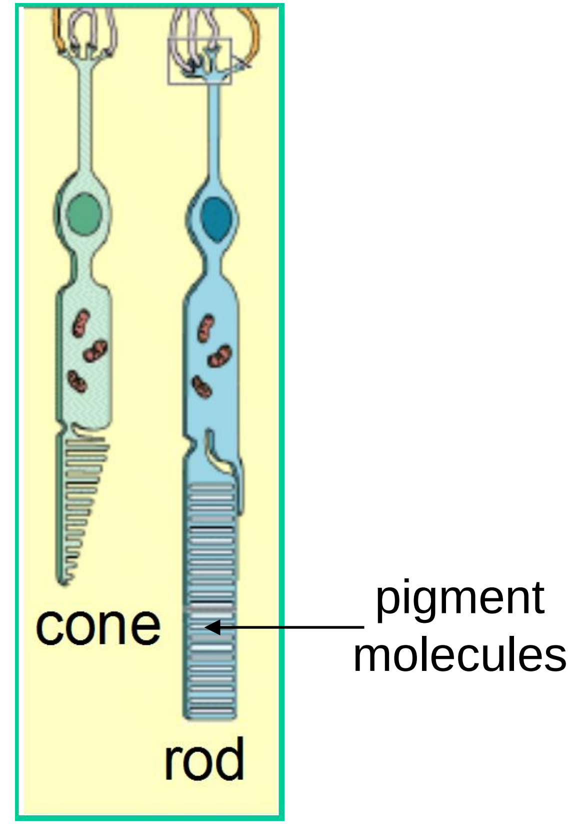

Rods and cones are *non-uniformly* distributed on the retina

- Rods responsible for intensity, cones responsible for color
- **Fovea** - Small region (1 or 2°) at the center of the visual field containing the highest density of cones (and no rods).
- Less visual acuity in the periphery—many rods wired to the same neuron

Rod / Cone sensitivity

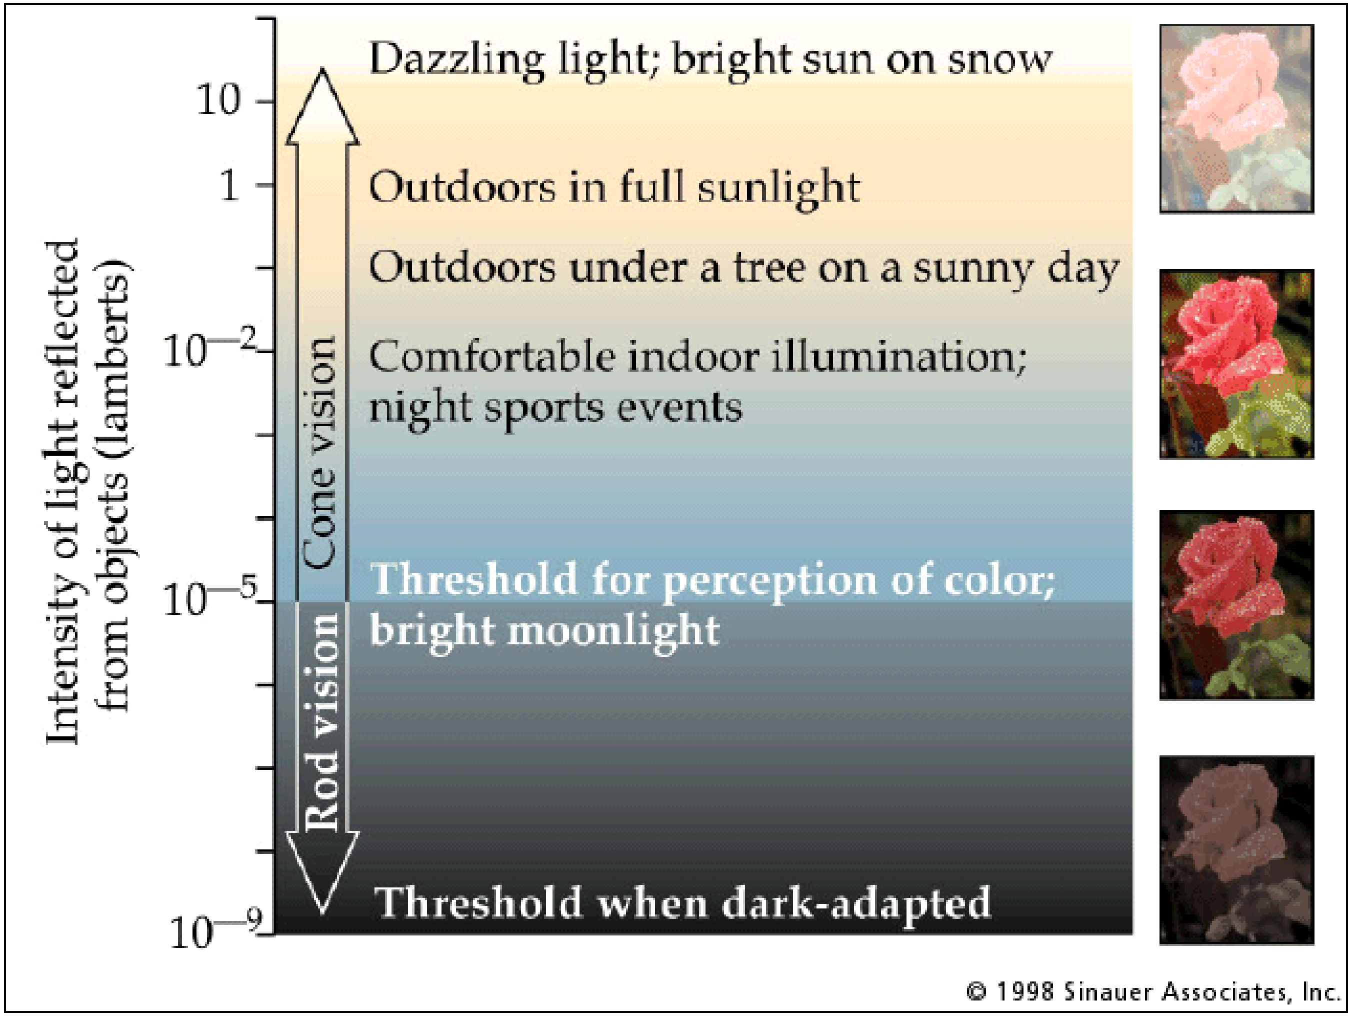

Why can't we read in the dark?

Slide by A. Efros

Physiology of Color Vision

Three kinds of cones:

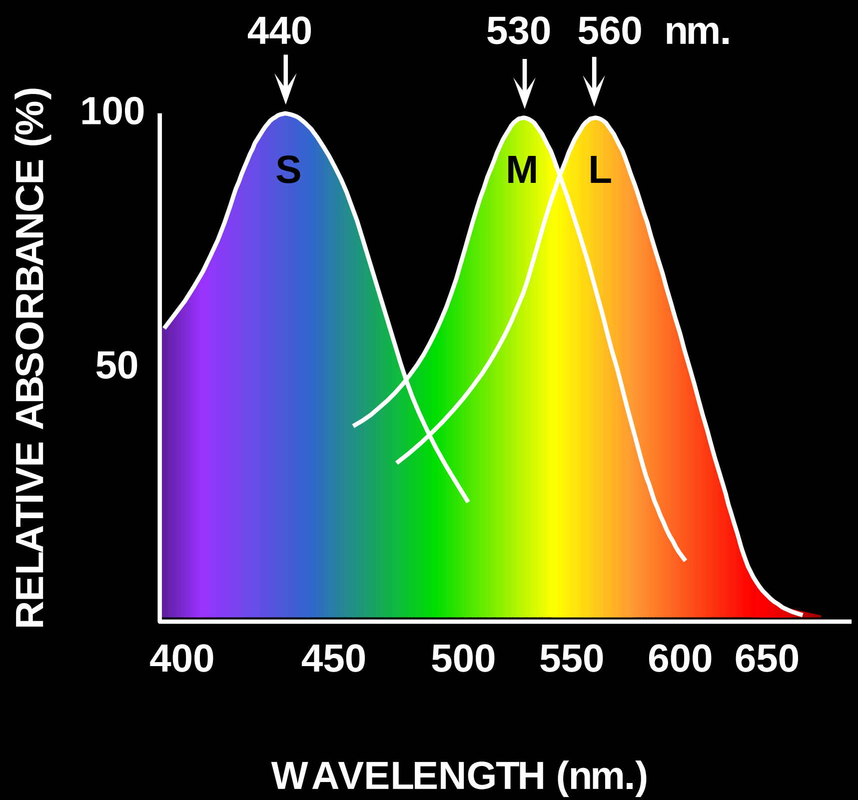

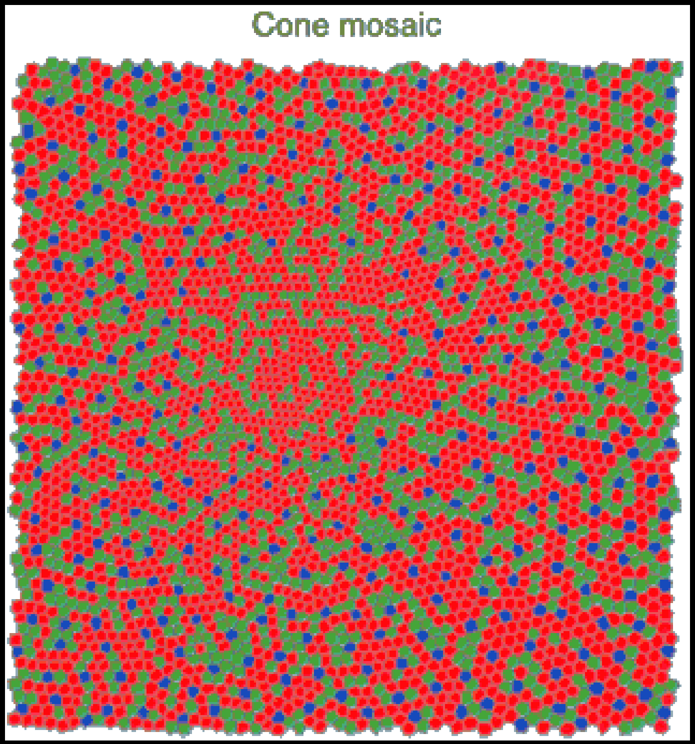

- Ratio of L to M to S cones: approx. 10:5:1
- Almost no S cones in the center of the fovea

Color interpolation in human visual system

Brewster's colors: evidence of interpolation from spatially offset color samples

Scale relative to human photoreceptor size: each line covers about 7 photoreceptors

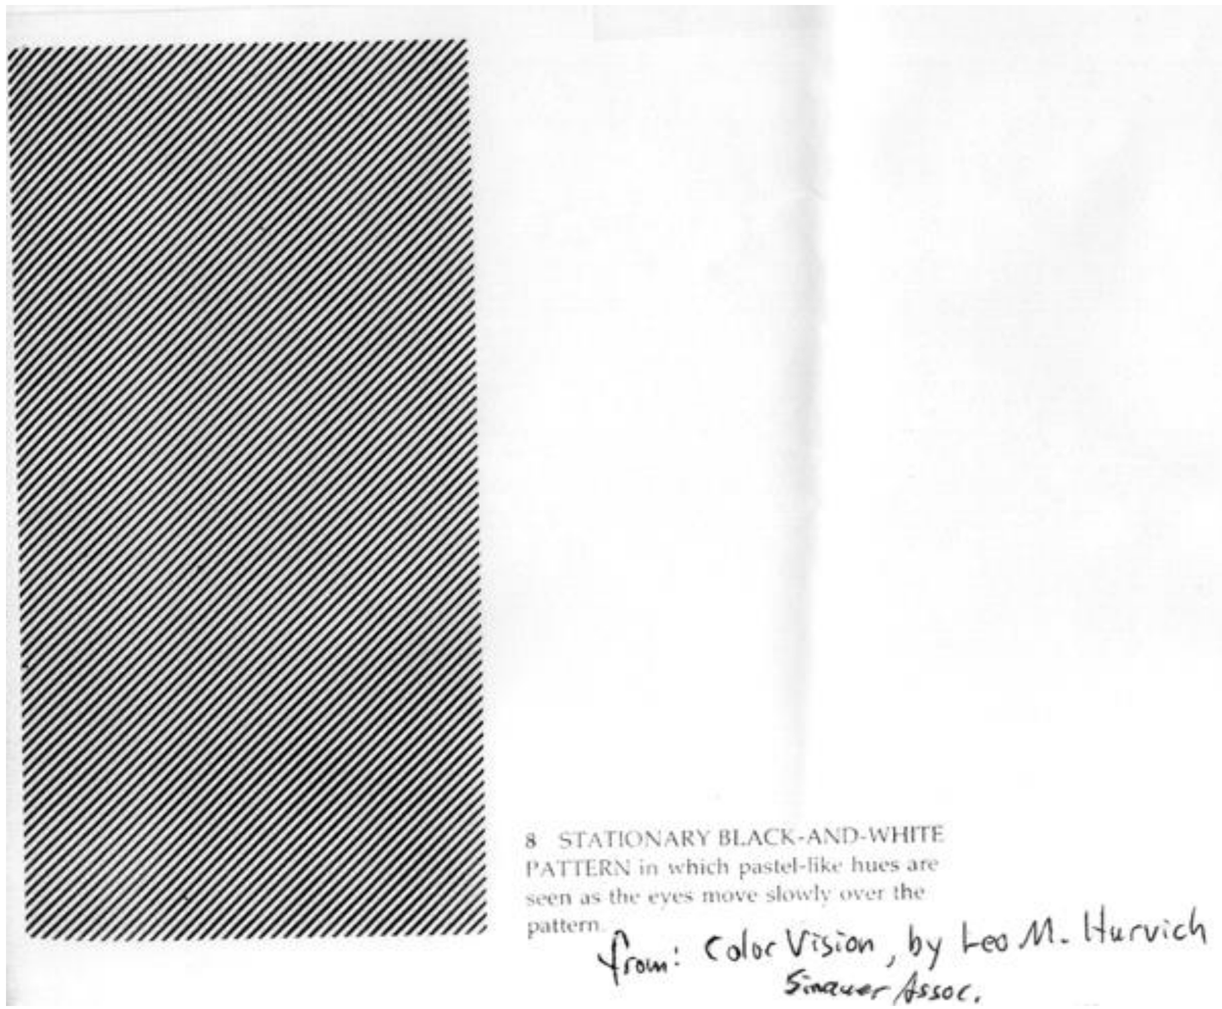

Source: F. Durand

Color perception

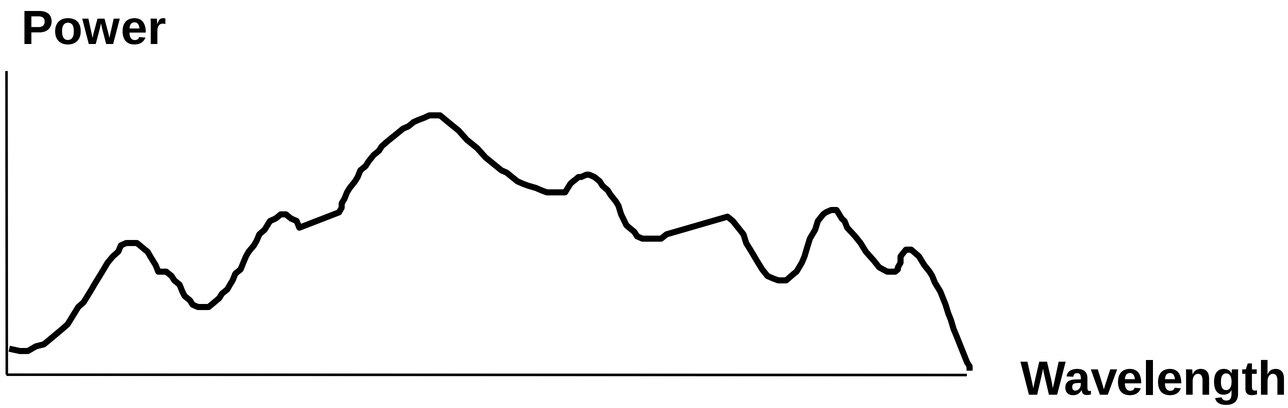

Rods and cones act as filters on the spectrum

- To get the output of a filter, multiply its response curve by the spectrum, integrate over all wavelengths
  - Each cone yields one number
- Q: How can we represent an entire spectrum with 3 numbers?

Color perception

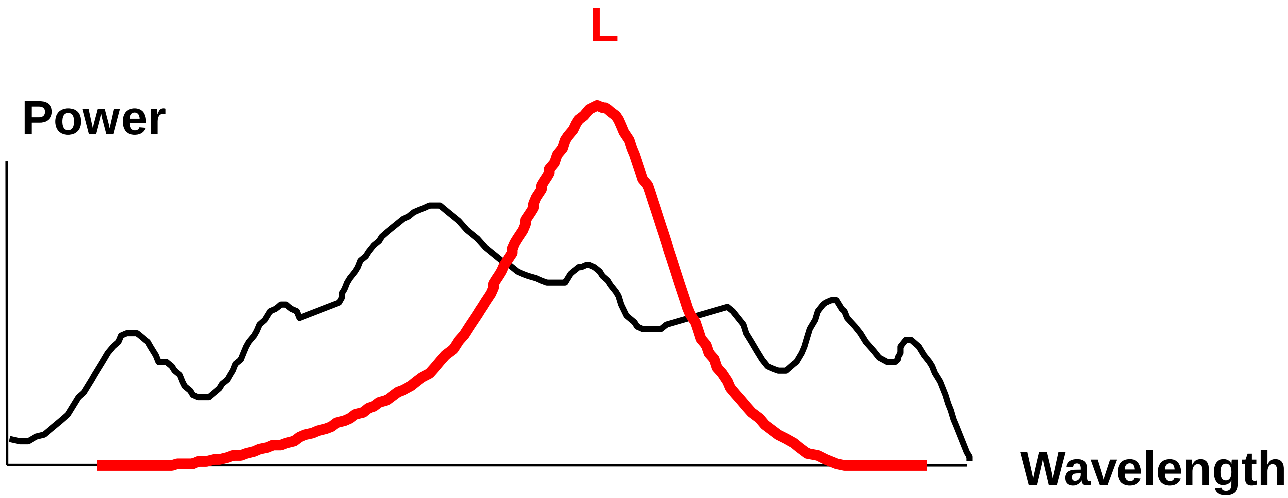

Rods and cones act as filters on the spectrum

- To get the output of a filter, multiply its response curve by the spectrum, integrate over all wavelengths
  - Each cone yields one number
- Q: How can we represent an entire spectrum with 3 numbers?

Slide by Steve Seitz

Color perception

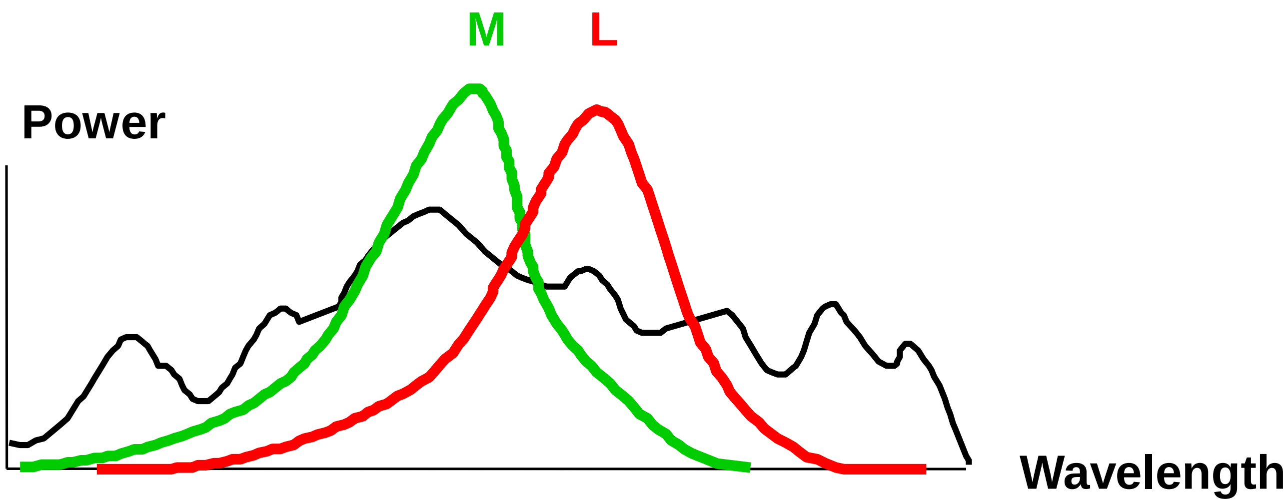

Rods and cones act as filters on the spectrum

- To get the output of a filter, multiply its response curve by the spectrum, integrate over all wavelengths
  - Each cone yields one number
- Q: How can we represent an entire spectrum with 3 numbers?

Slide by Steve Seitz

Color perception

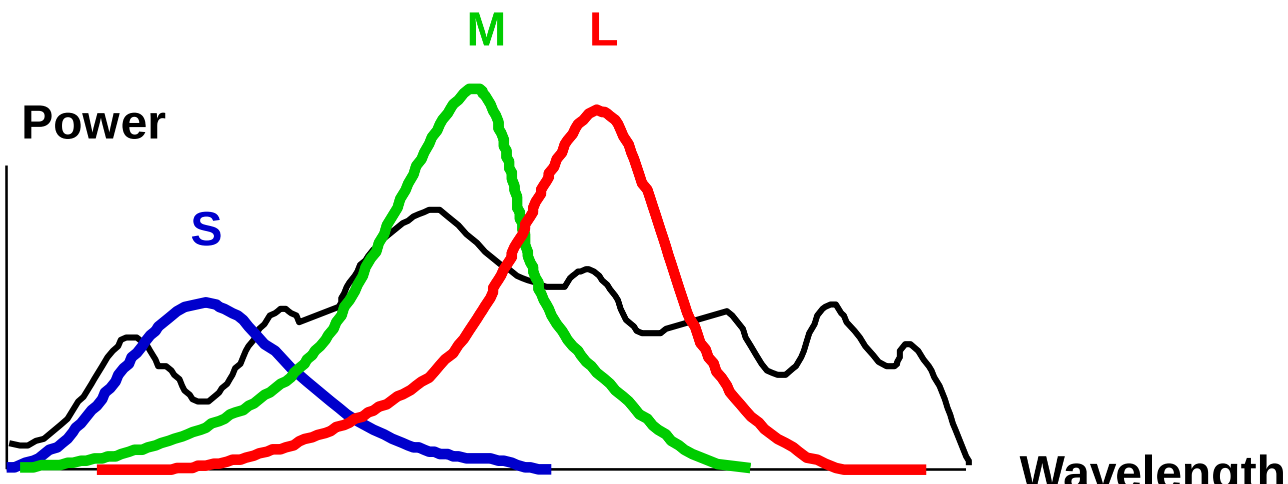

Rods and cones act as filters on the spectrum

- To get the output of a filter, multiply its response curve by the spectrum, integrate over all wavelengths
  - Each cone yields one number
- Q: How can we represent an entire spectrum with 3 numbers?

Slide by Steve Seitz

Color perception

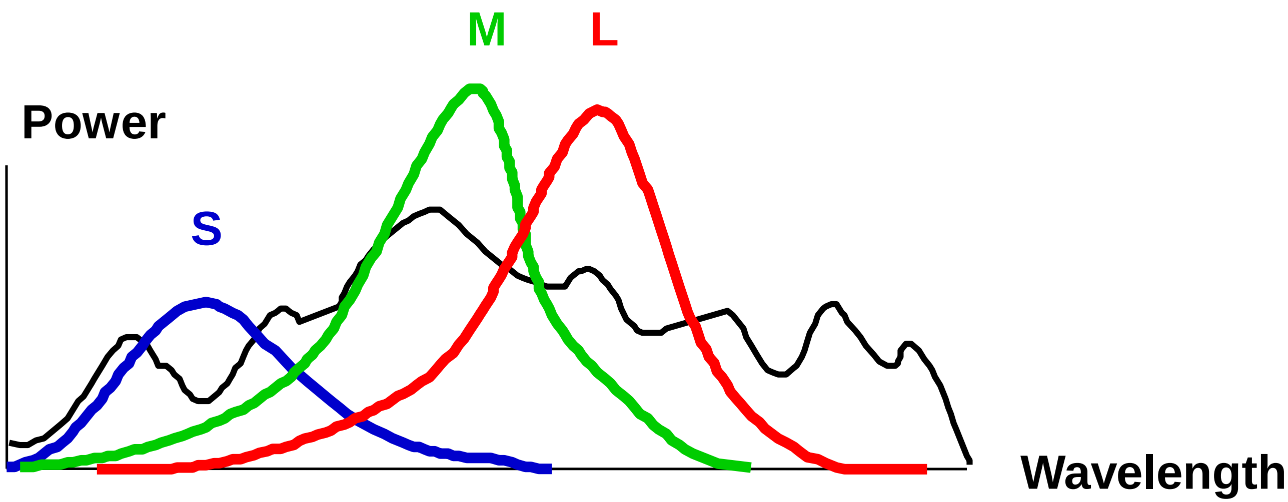

Rods and cones act as filters on the spectrum

- To get the output of a filter, multiply its response curve by the spectrum, integrate over all wavelengths
  - Each cone yields one number
- Q: How can we represent an entire spectrum with 3 numbers?
- A: We can't! Most of the information is lost.
  - As a result, two different spectra may appear indistinguishable
    - » such spectra are known as **metamers**

Slide by Steve Seitz

Spectra of some real-world surfaces

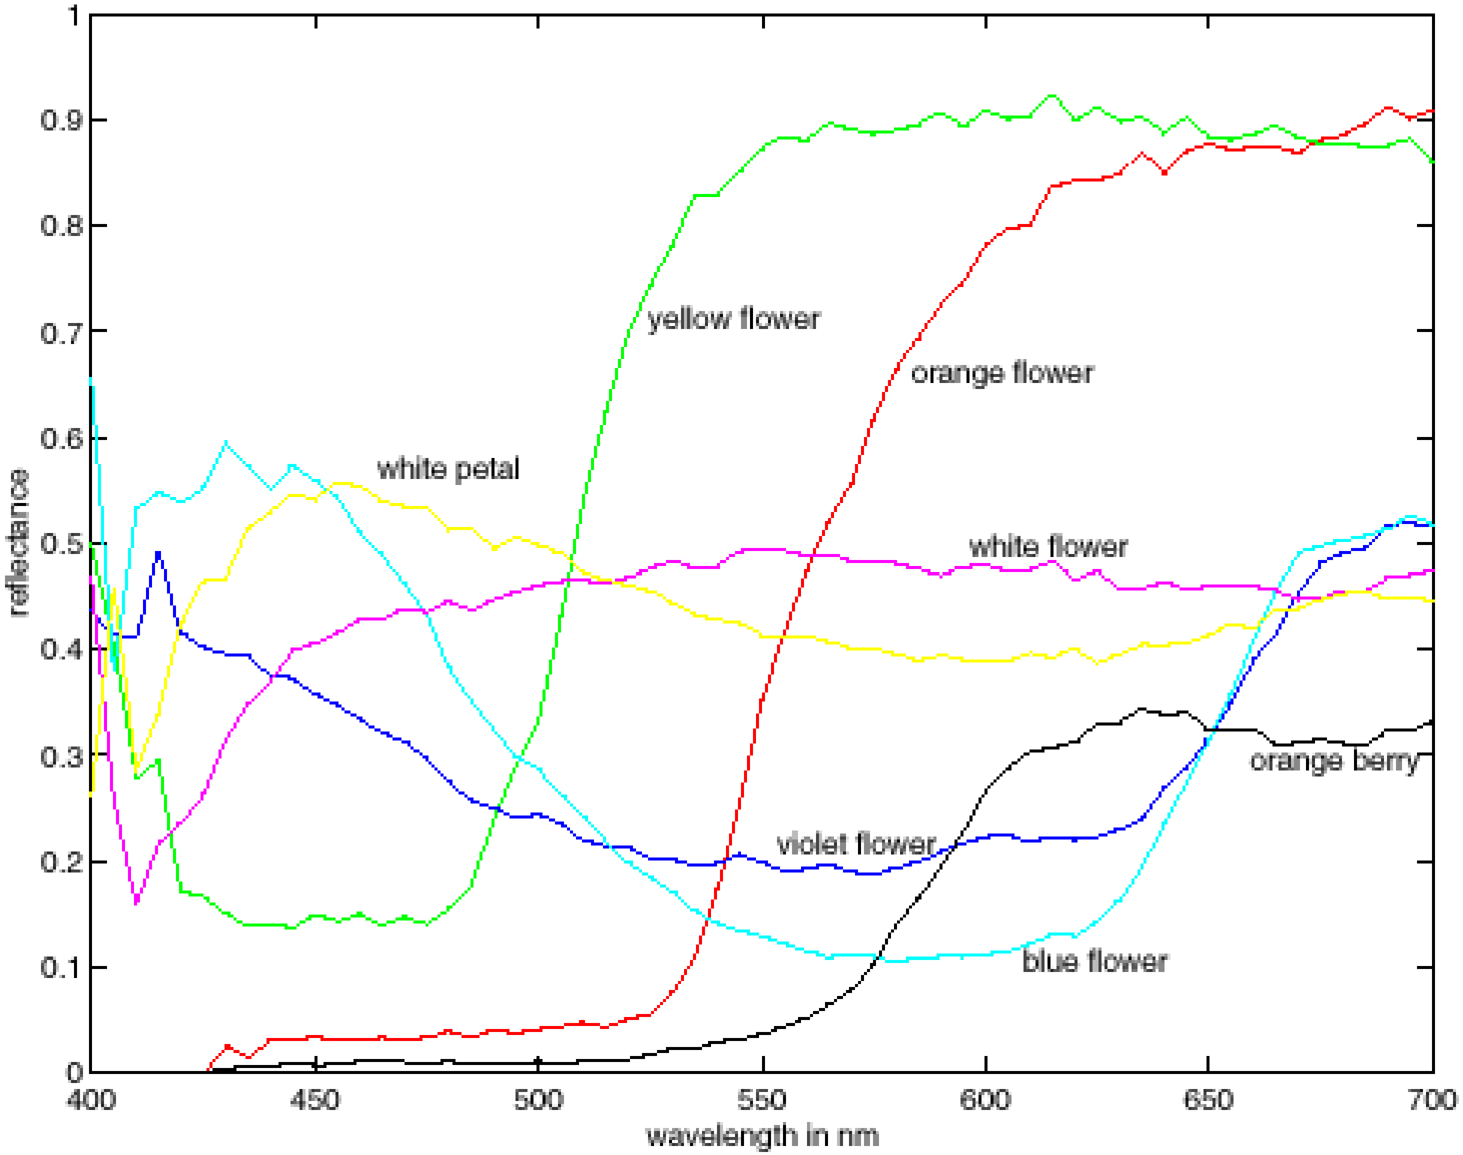

Spectra of some real-world surfaces

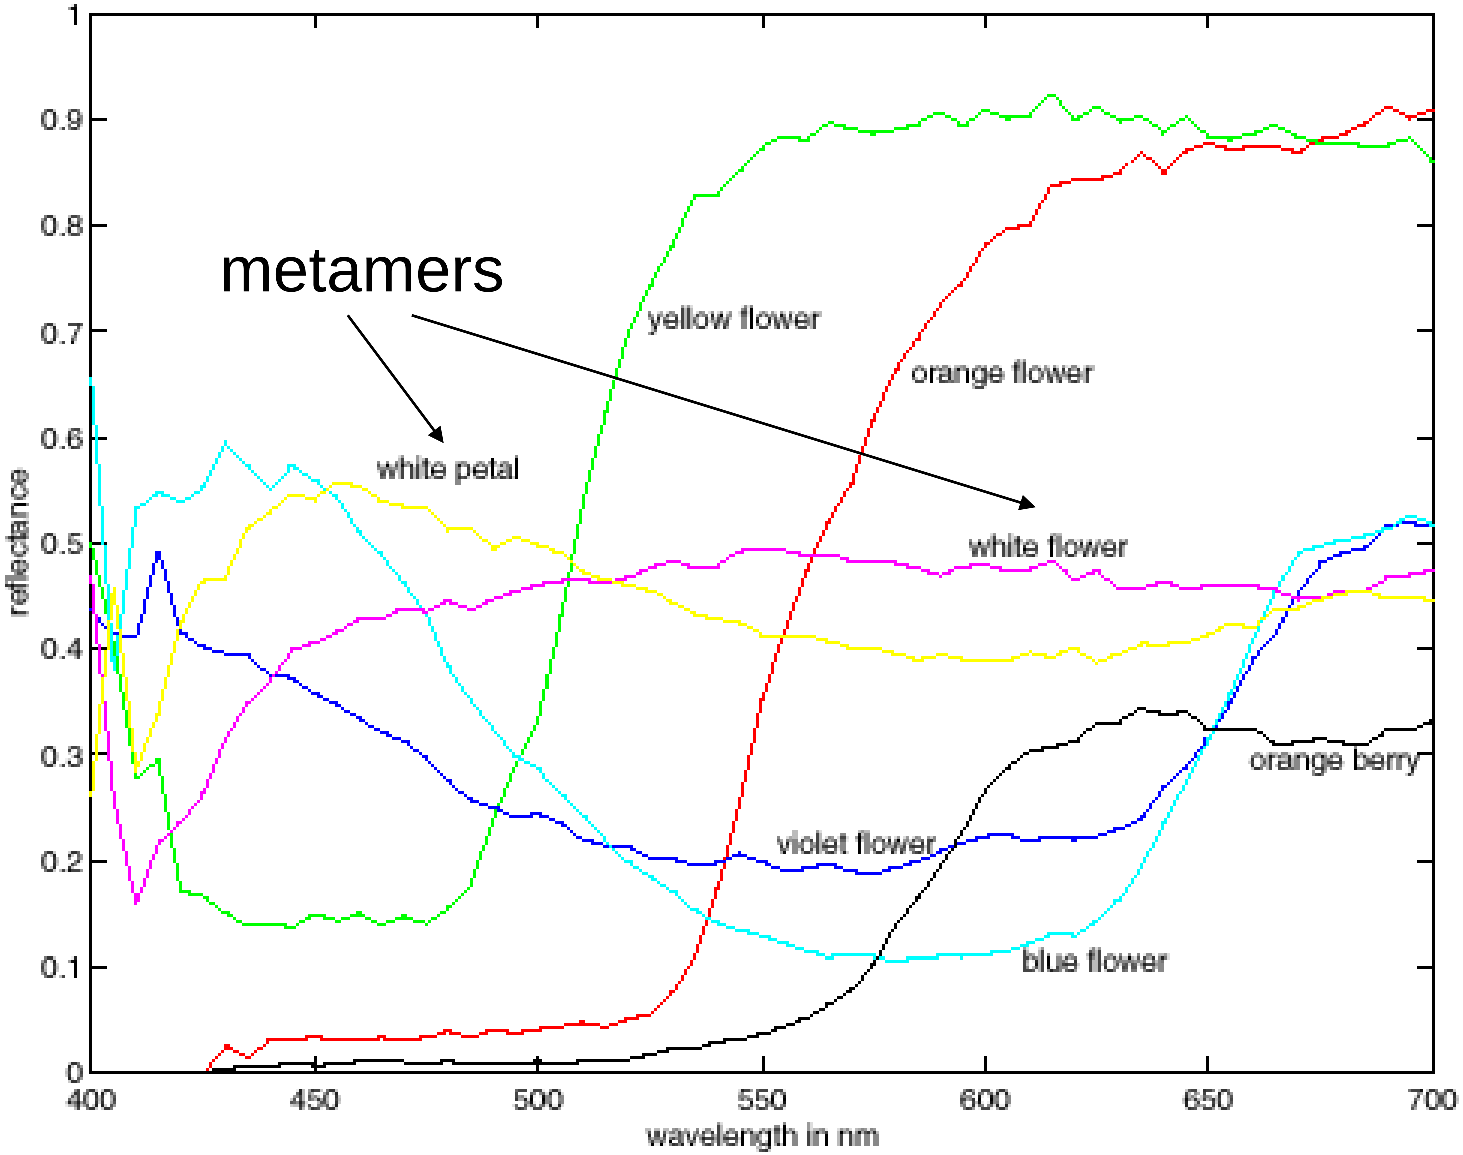

Metamers

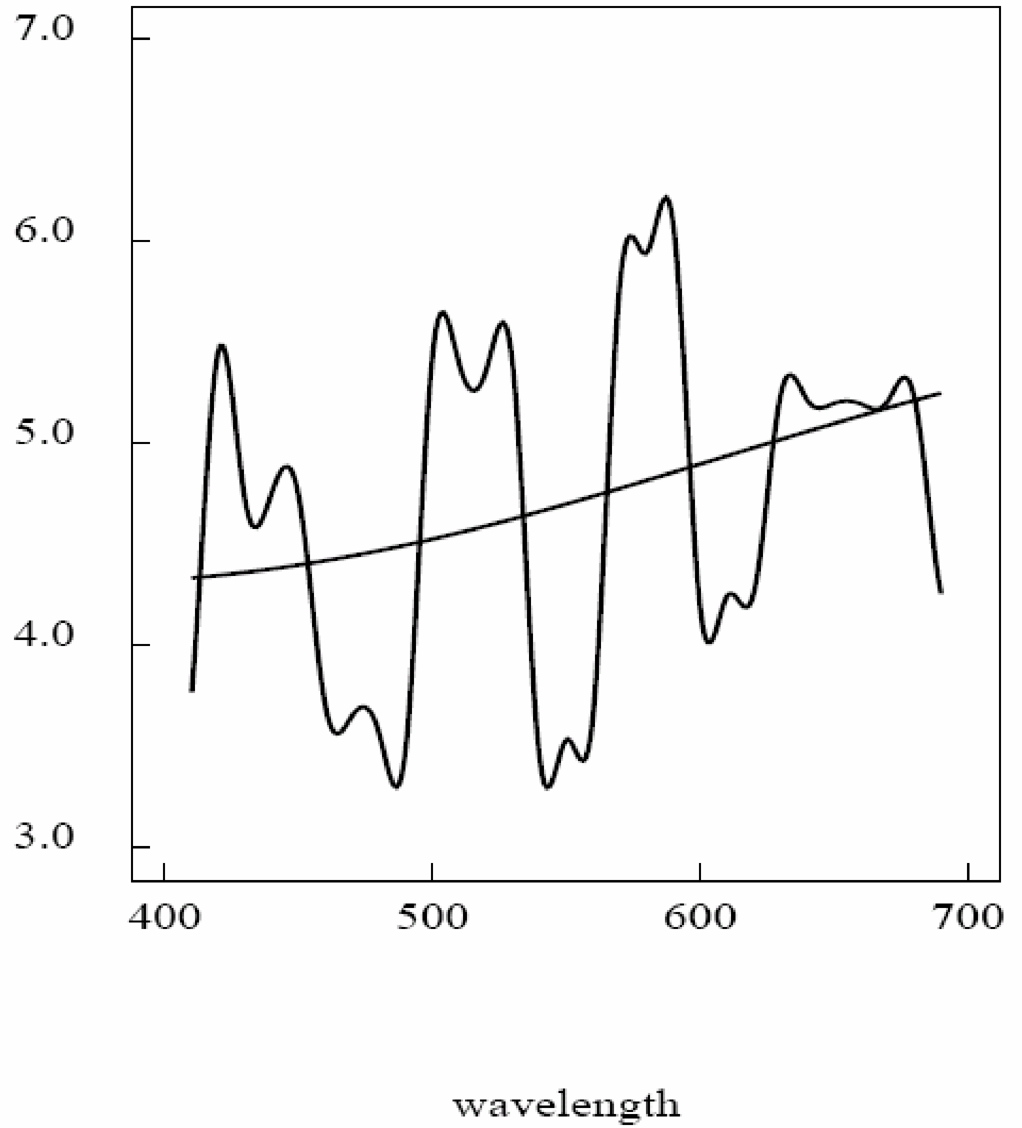

Standardizing color experience

- We would like to understand which spectra produce the same color sensation from people under similar viewing conditions
- Color matching experiments

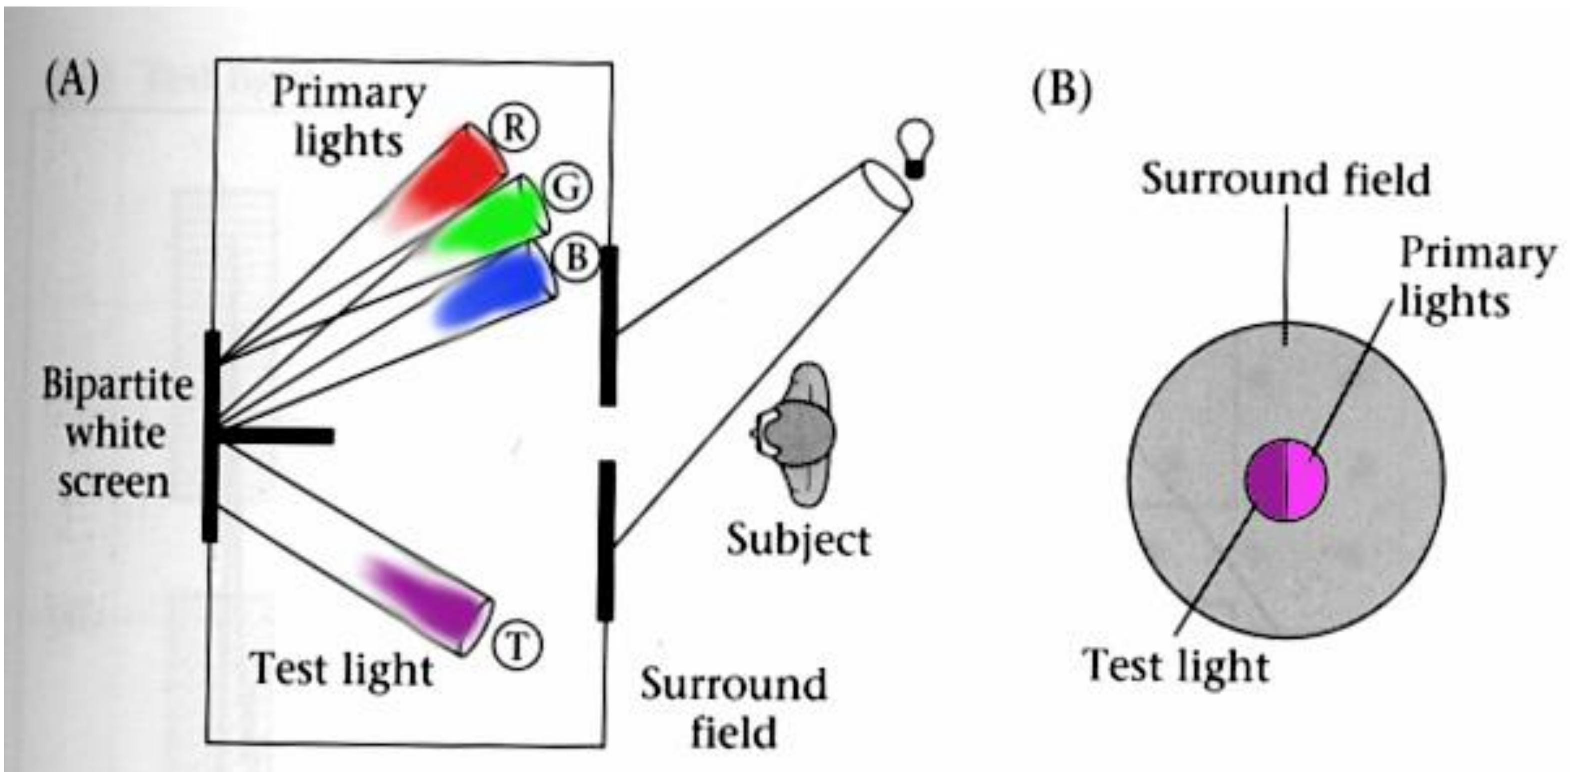

Foundations of Vision, by Brian Wandell, Sinauer Assoc., 1995

Color matching experiment 1

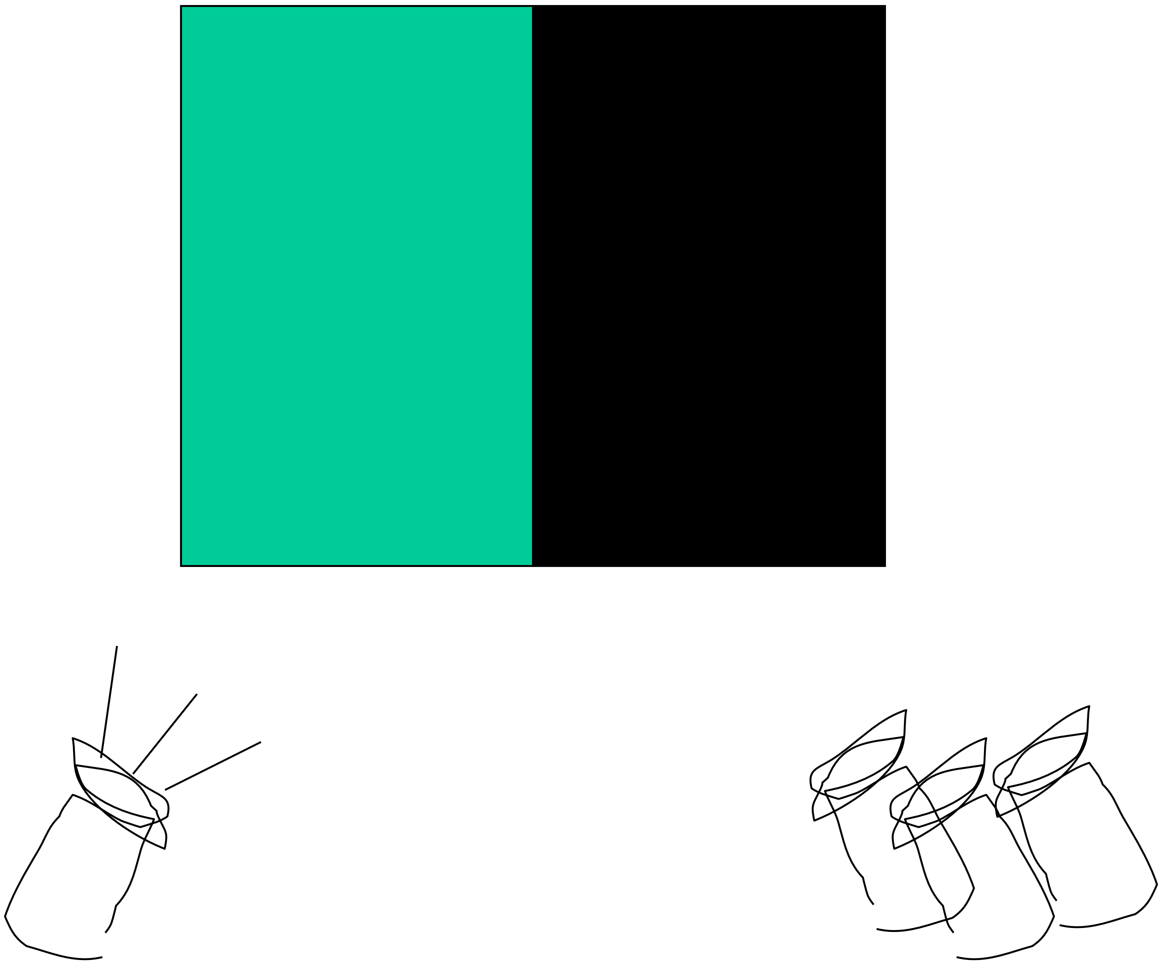

Source: W. Freeman

Color matching experiment 1

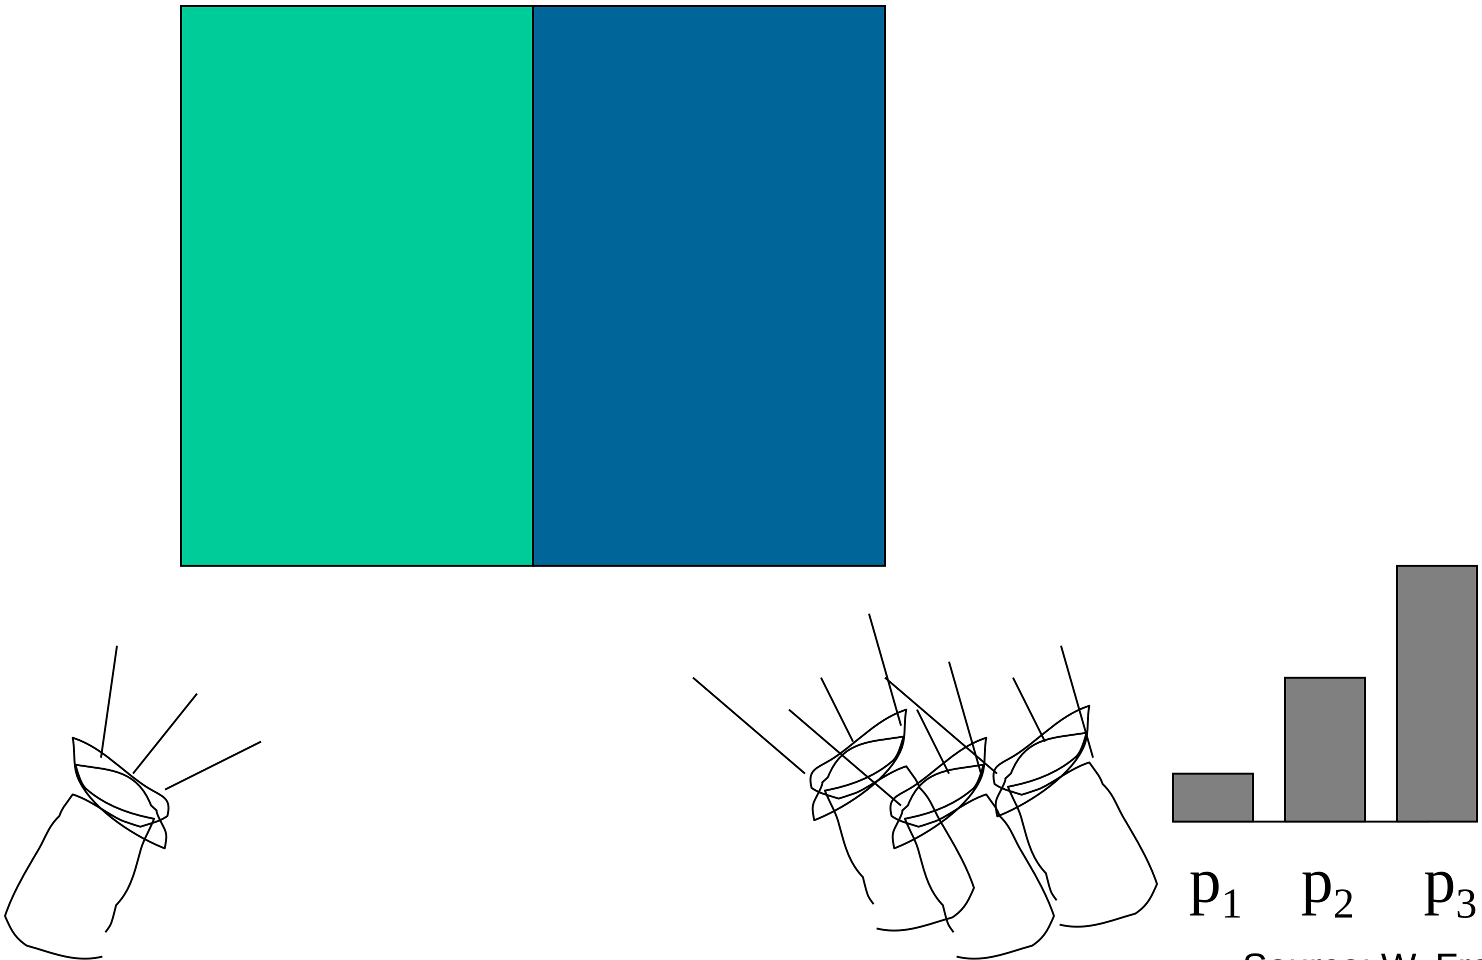

Source: W. Freeman

Color matching experiment 1

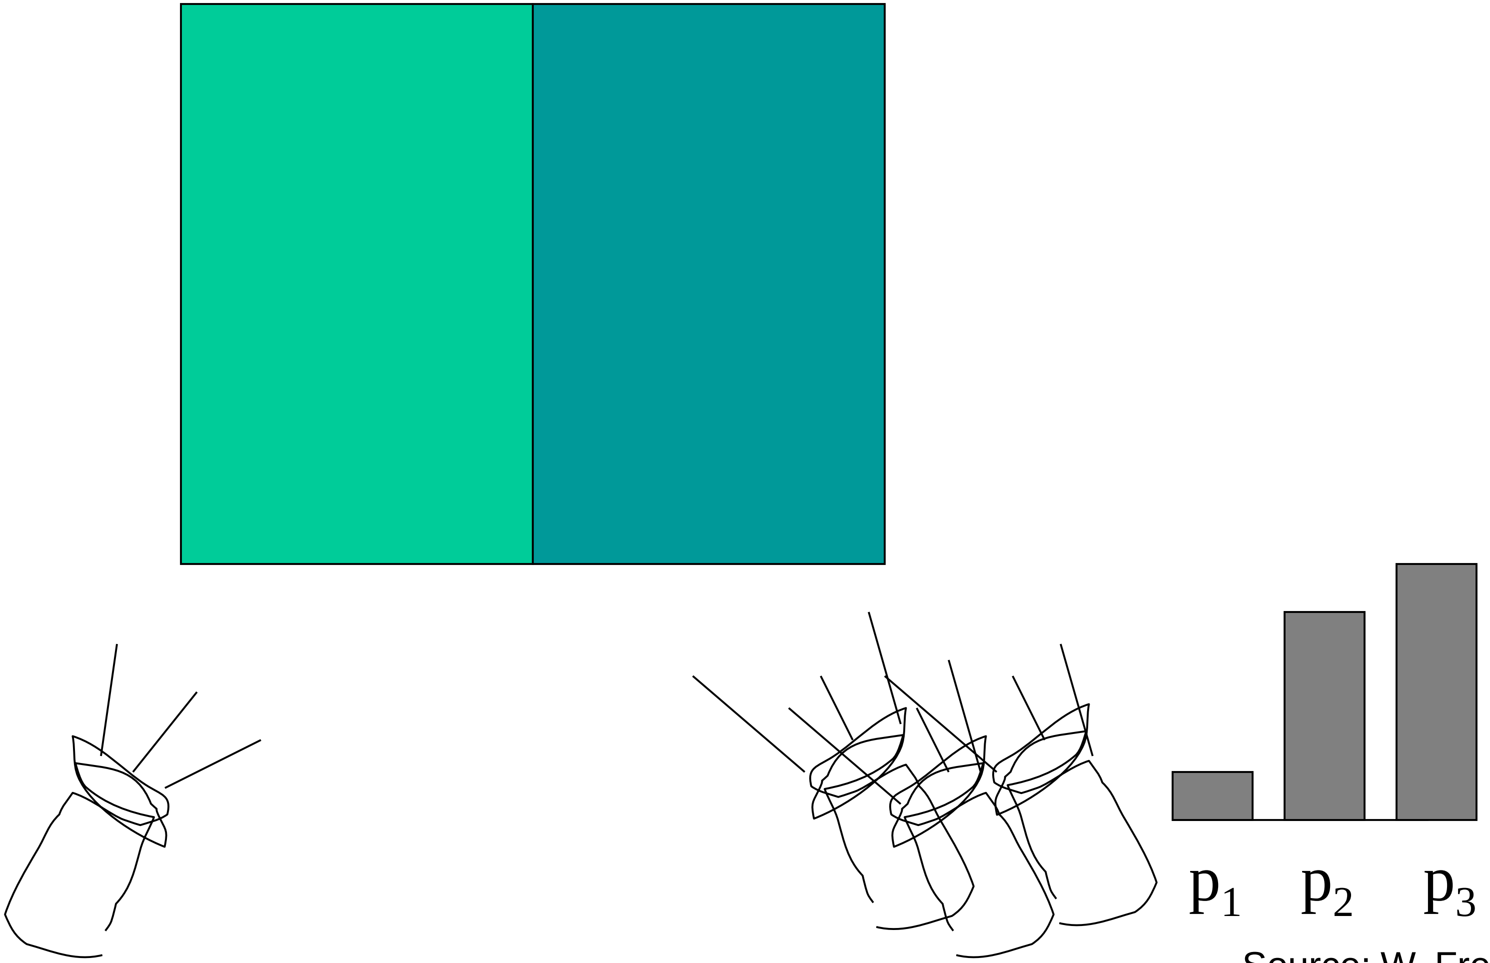

Source: W. Freeman

Color matching experiment 1

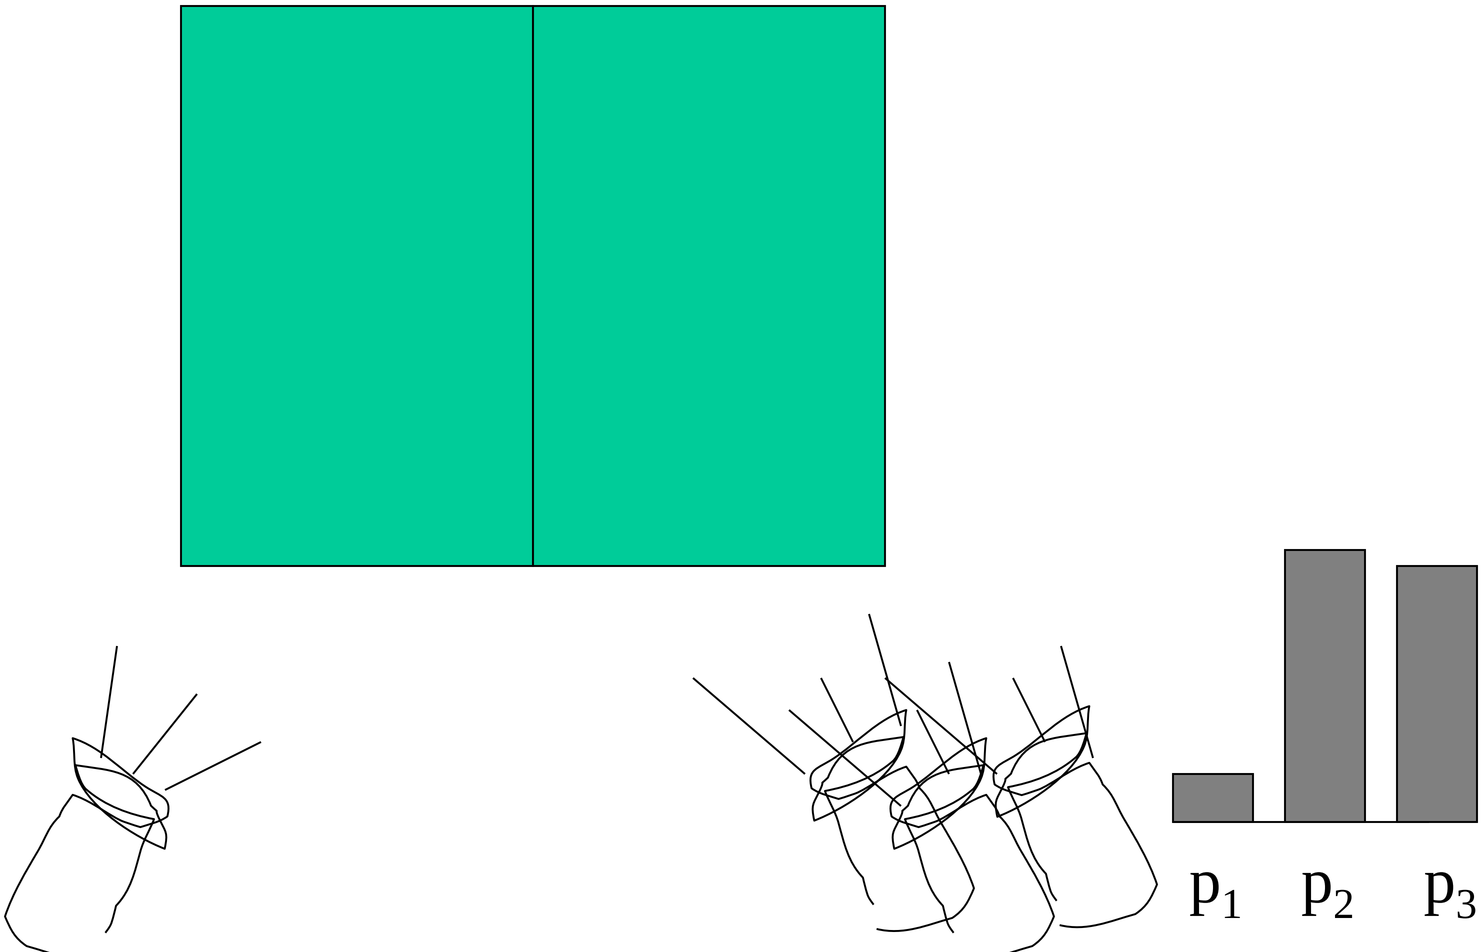

Source: W. Freeman

Color matching experiment 1

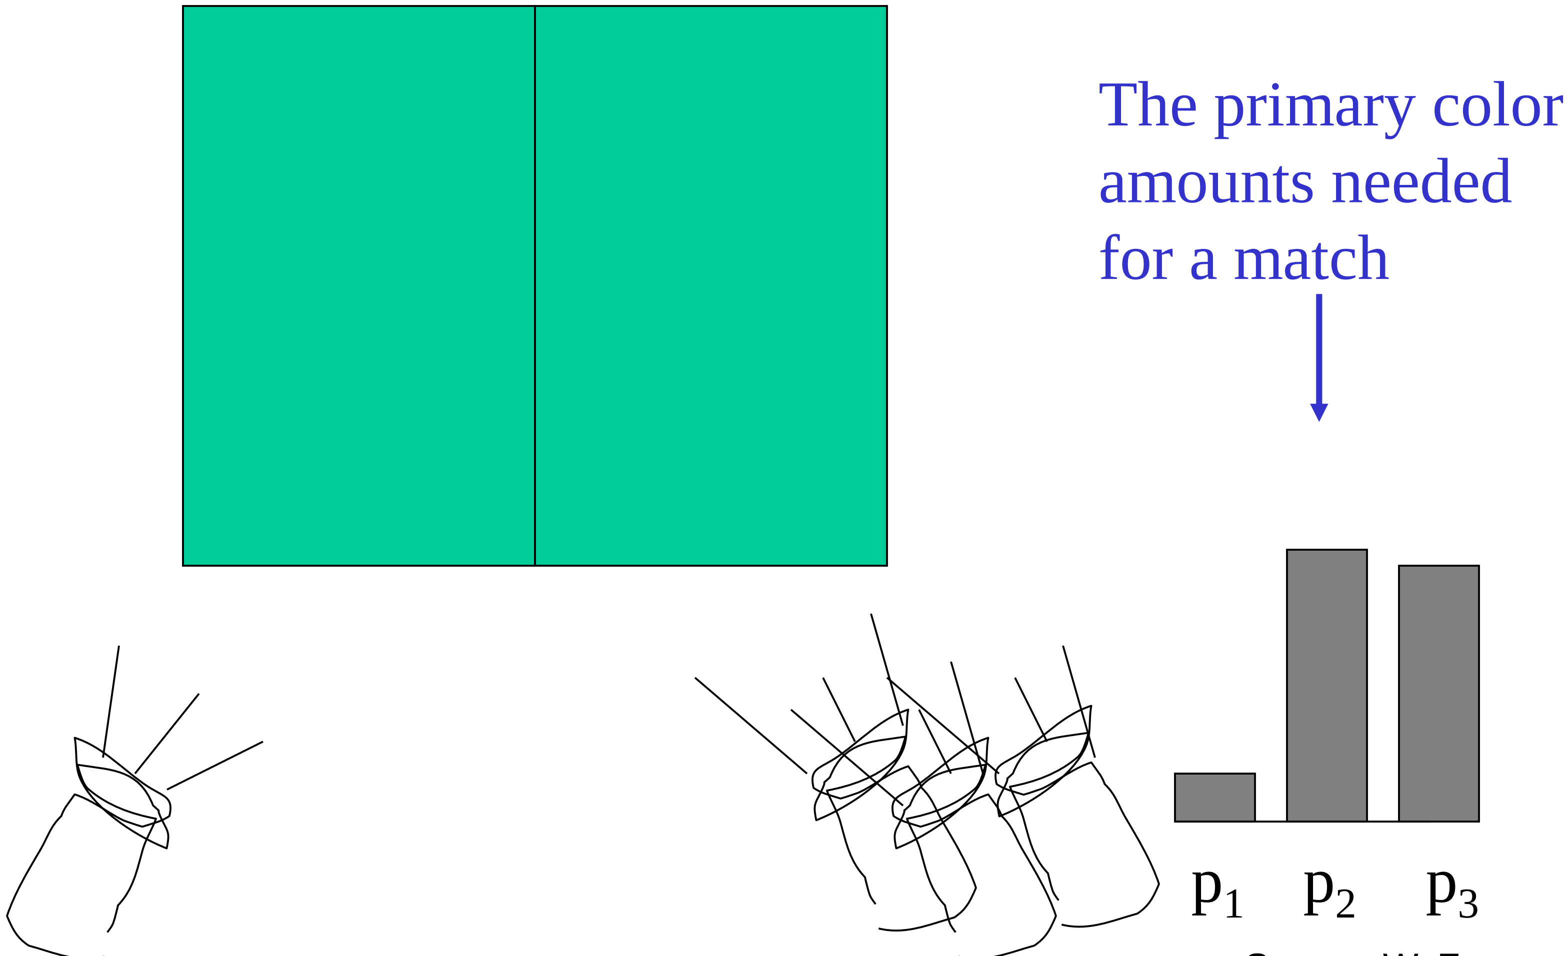

Source: W. Freeman

Color matching experiment 2

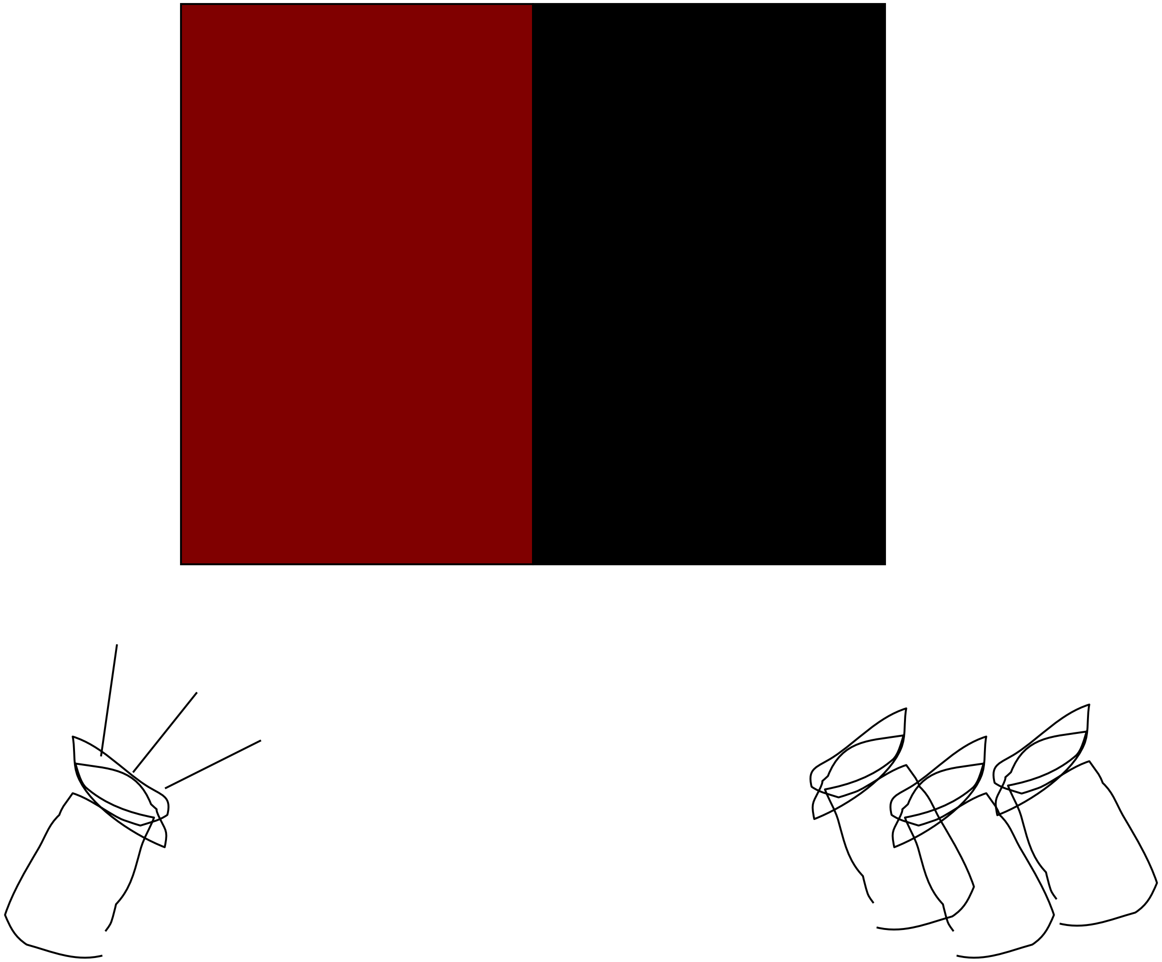

Source: W. Freeman

Color matching experiment 2

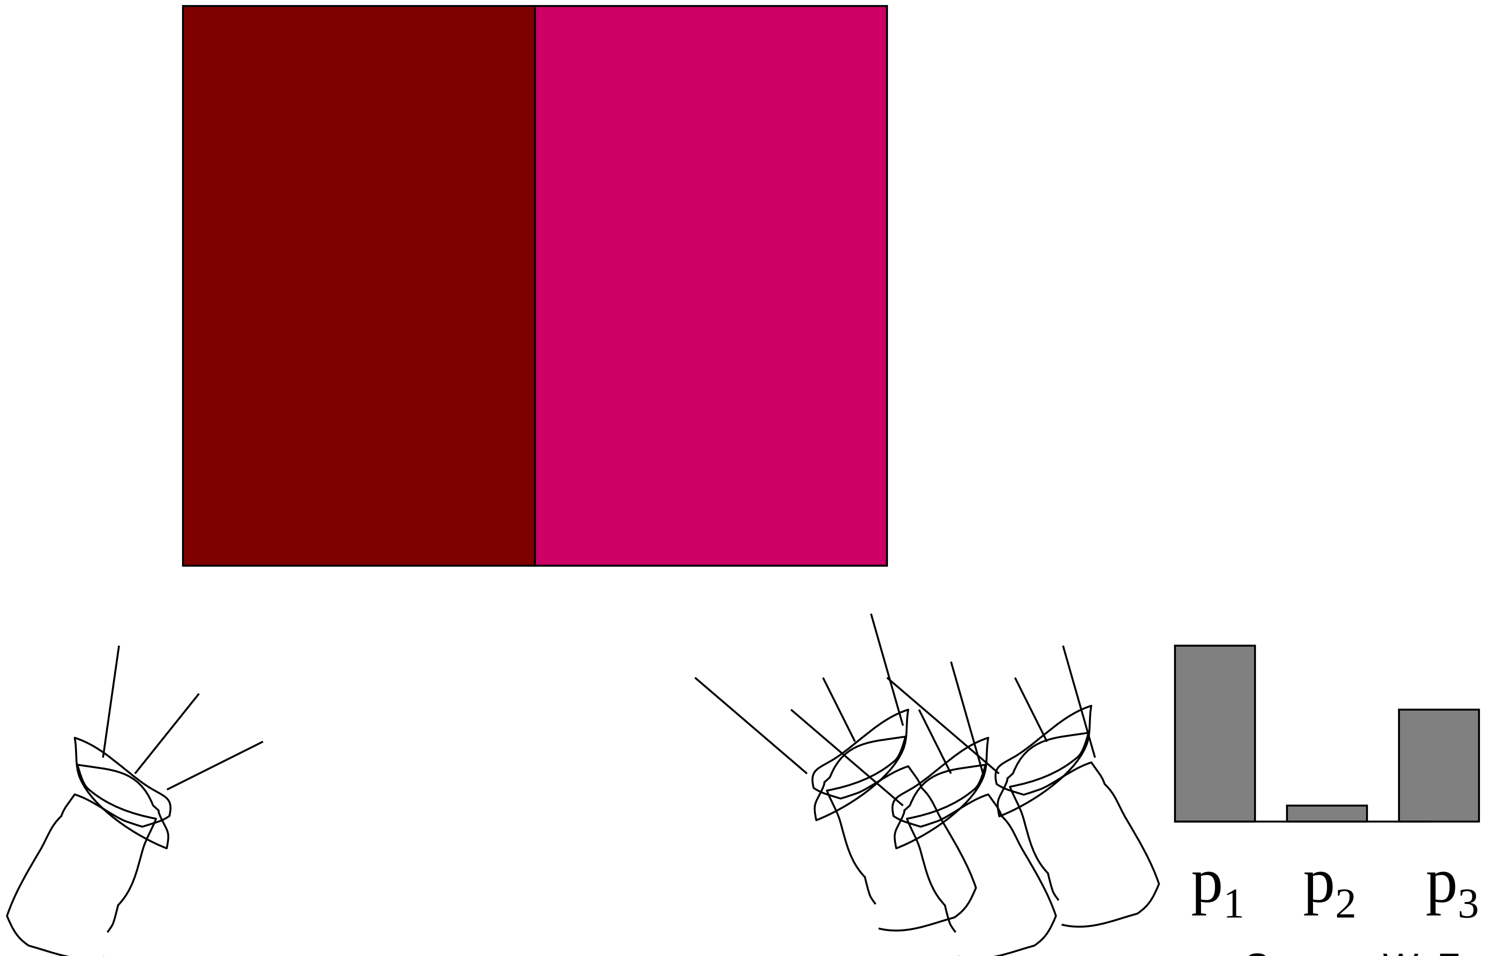

Source: W. Freeman

Color matching experiment 2

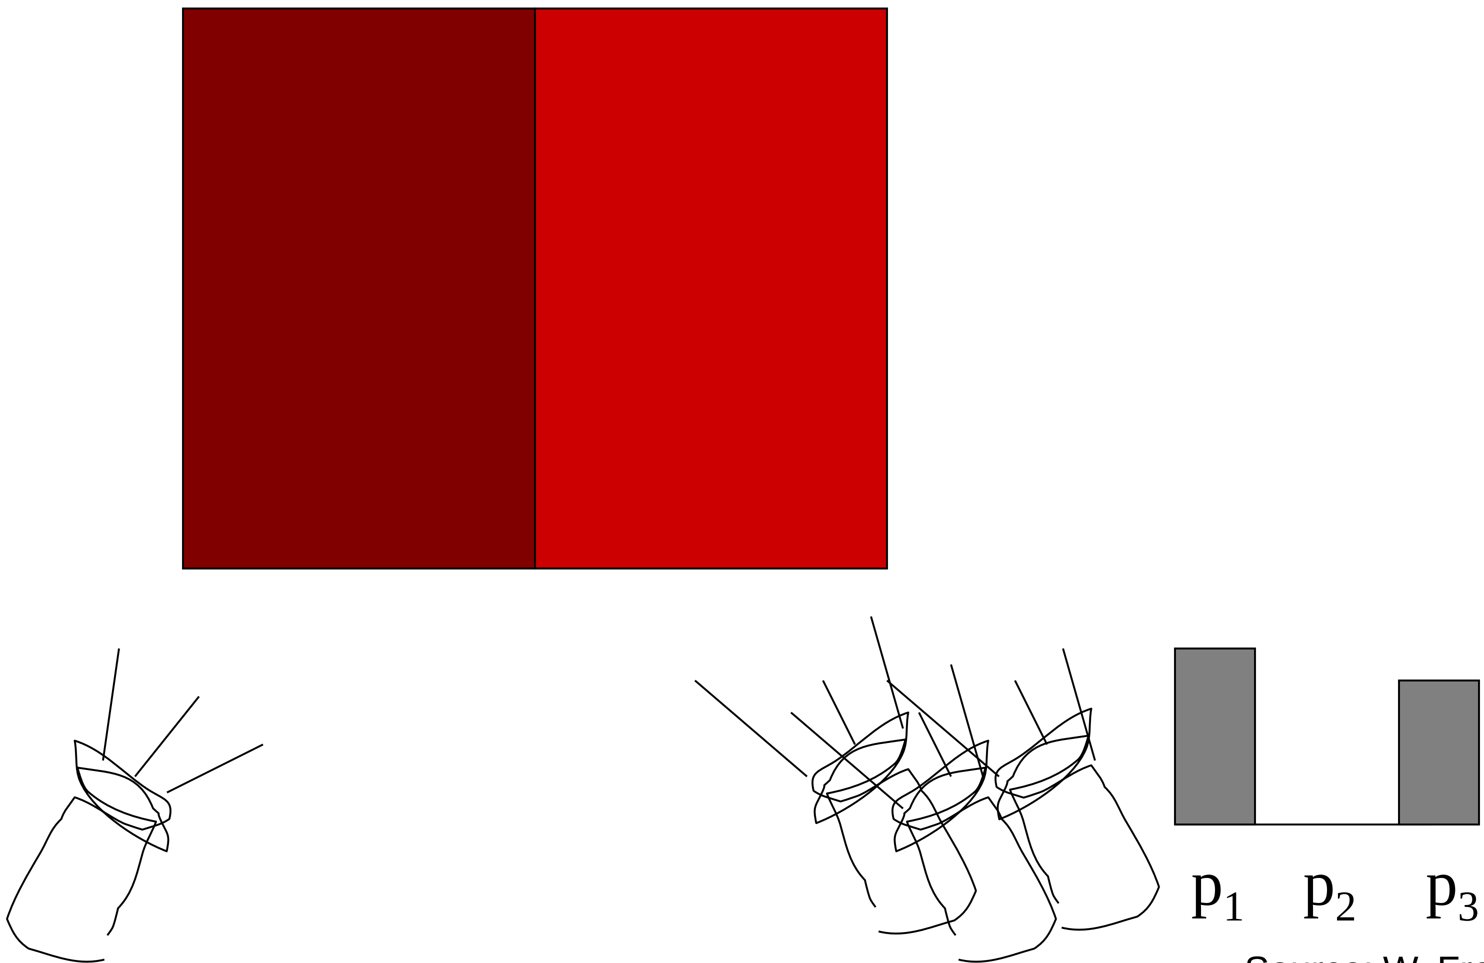

Source: W. Freeman

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.

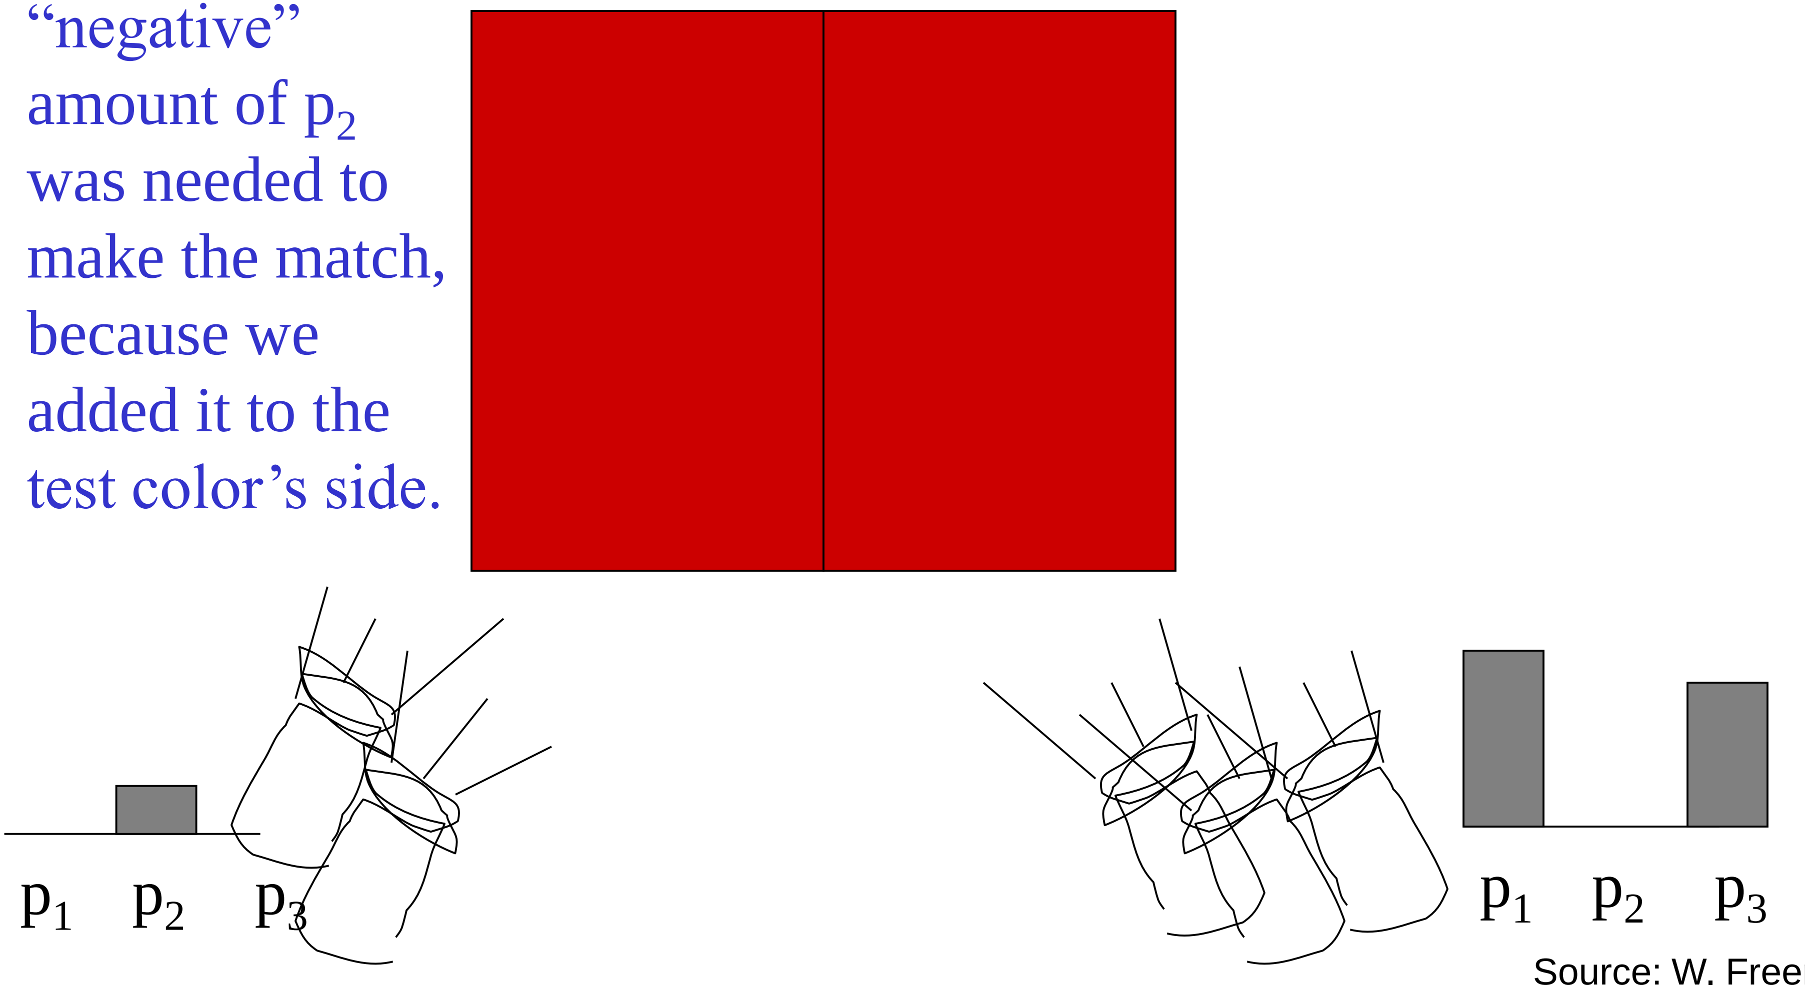

Source: W. Freeman

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:

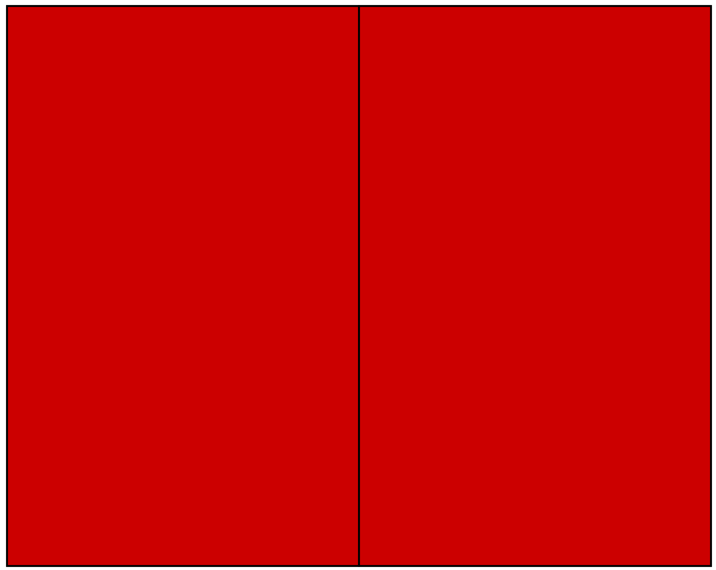

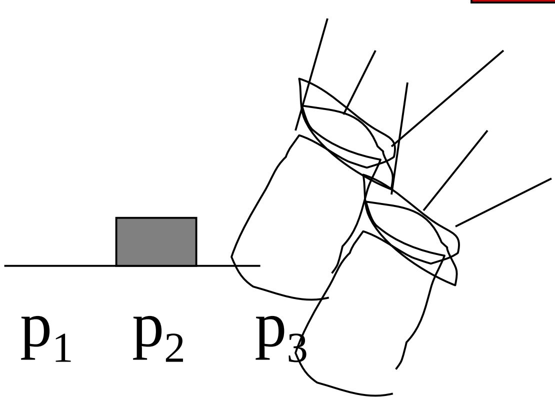

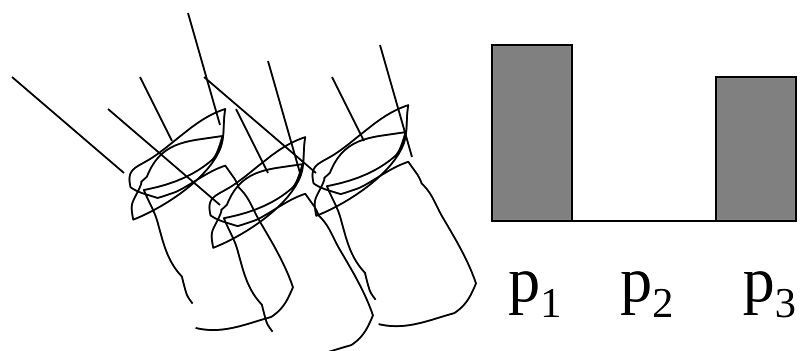

Source: W. Freeman

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.

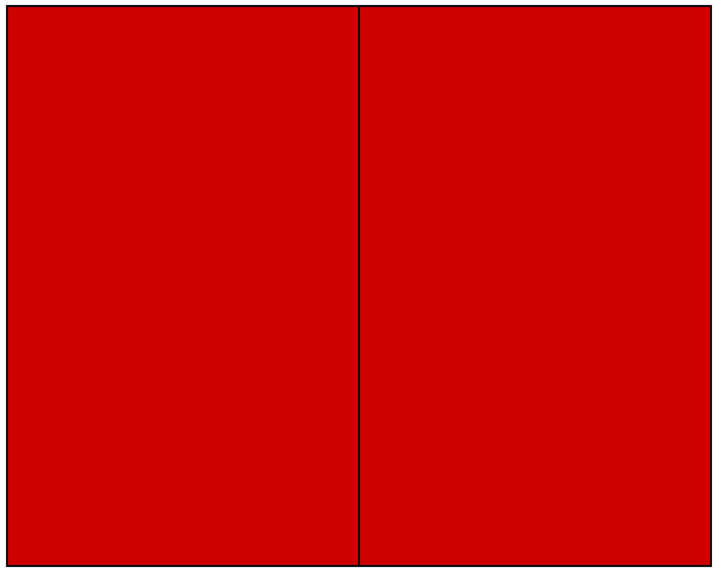

The primary color amounts needed for a match:

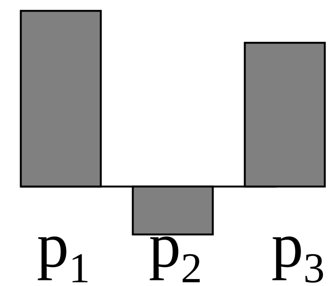

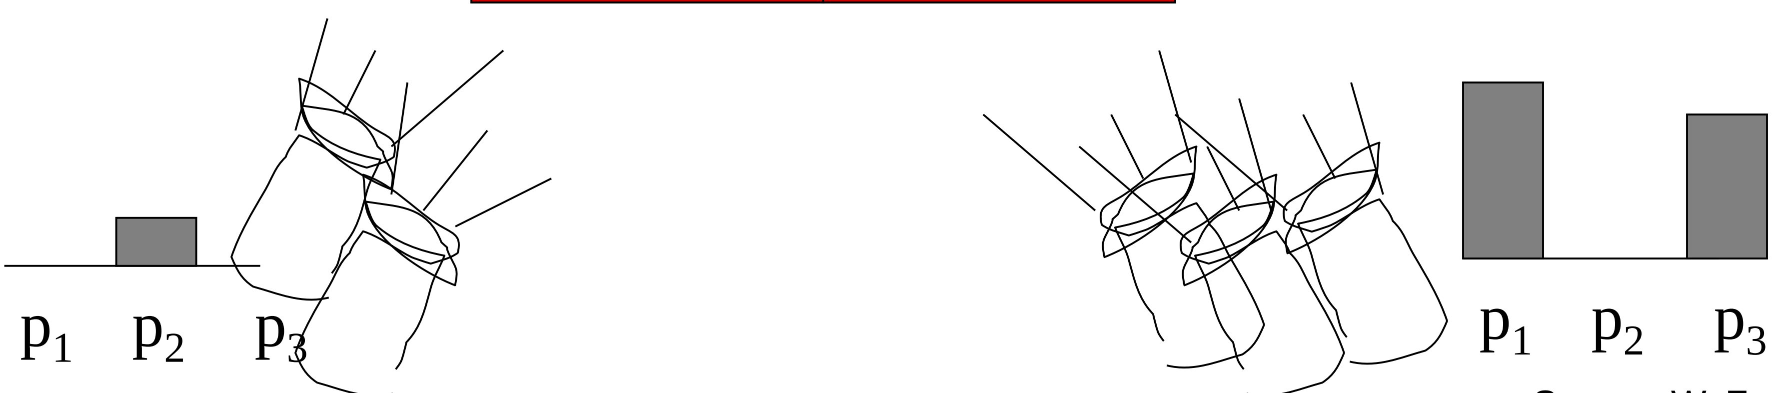

Source: W. Freeman

Trichromacy

- Three numbers seem to be sufficient for encoding color
- In color matching experiments, most people can match any given light with three primaries
  - Exception: color blindness
- For the same light and same primaries, most people select the same weights
- Trichromatic color theory dates back to 18th century (Thomas Young)

Grassman's Laws

- Additive matching is linear

Grassman’s Laws

- Additive matching is linear
- 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 $A = B$.

Grassman's Laws

- Additive matching is linear
- 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 mix two test lights, then mixing the matches will match the result:
  - 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 $A+B = (u_1+v_1) P_1 + (u_2+v_2) P_2 + (u_3+v_3) P_3$.

Grassman's Laws

- Additive matching is linear
- 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 mix two test lights, then mixing the matches will match the result:
  - 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 $A+B = (u_1+v_1) P_1 + (u_2+v_2) P_2 + (u_3+v_3) P_3$.
- 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 $kA = (ku_1) P_1 + (ku_2) P_2 + (ku_3) P_3$.

Linear color spaces

- Defined by a choice of three primaries
- The coordinates of a color are given by the weights of the primaries used to match it
- *Matching functions*: weights required to match single-wavelength light sources

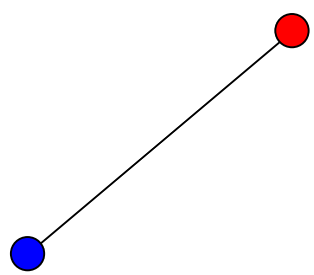

mixing two lights produces colors that lie along a straight line along them in color space

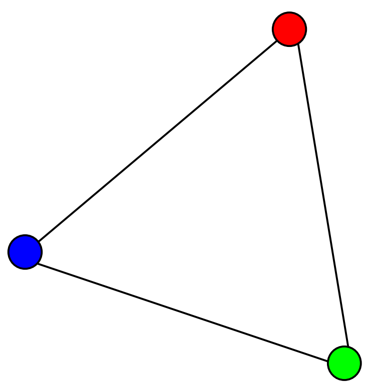

mixing three lights produces colors that lie within the triangle they define in color space

How to compute the color match for any color signal for any set of primary colors

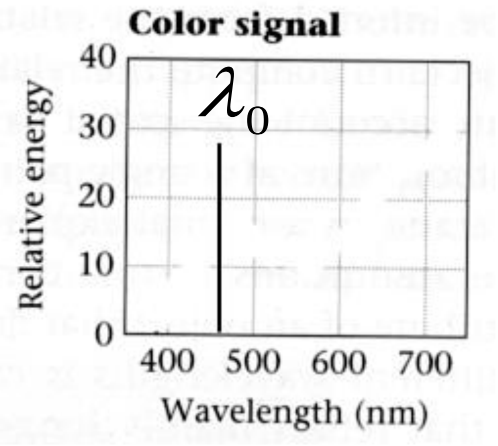

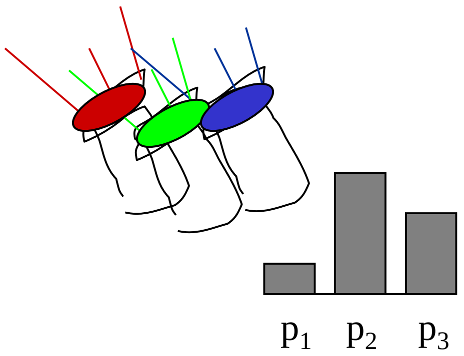

Source: W. Freeman

How to compute the color match for any color signal for any set of primary colors

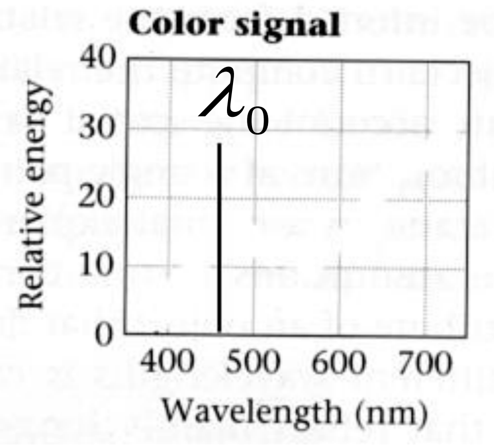

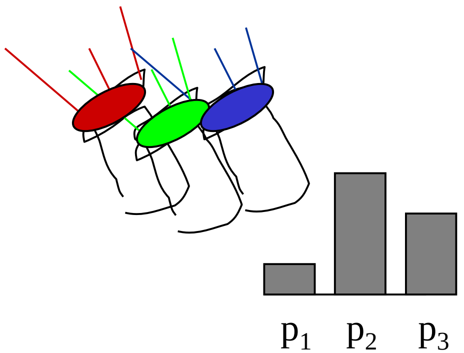

- Pick a set of primaries, $p_1(\lambda), p_2(\lambda), p_3(\lambda)$

Source: W. Freeman

How to compute the color match for any color signal for any set of primary colors

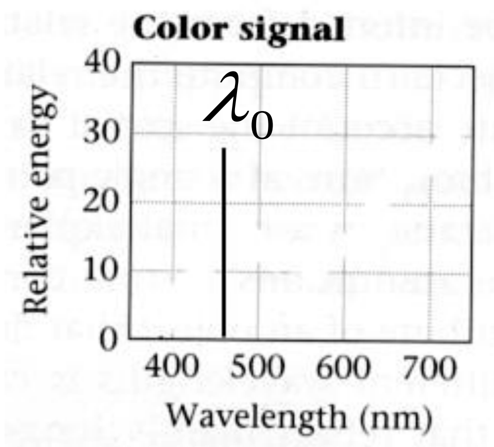

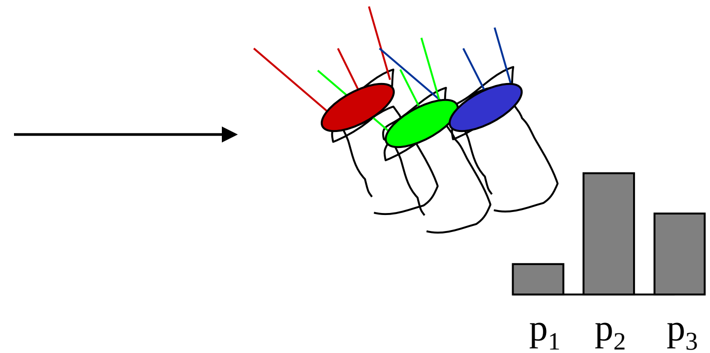

- Pick a set of primaries, $p_1(\lambda), p_2(\lambda), p_3(\lambda)$
- Measure the amount of each primary, $c_1(\lambda_0), c_2(\lambda_0), c_3(\lambda_0)$ needed to match a monochromatic light, $t(\lambda_0)$ at each spectral wavelength λ_0 (pick some spectral step size). These are the color matching functions.

Source: W. Freeman

Using color matching functions to predict the matches for a new spectral signal

We know that a monochromatic light λ_i of wavelength will be matched by the amounts $c_1(\lambda_i), c_2(\lambda_i), c_3(\lambda_i)$ of each primary.

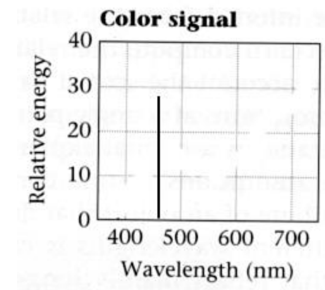

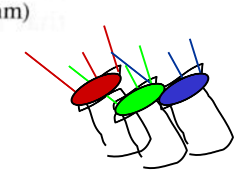

And any spectral signal can be thought of as a linear combination of very many monochromatic lights, with the linear coefficient given by the spectral power at each wavelength.

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

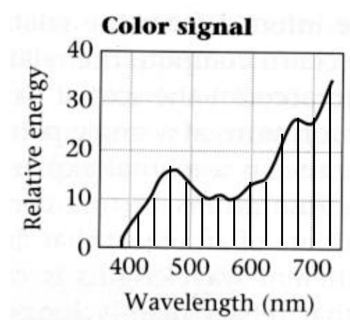

Source: W. Freeman

Using color matching functions to predict the primary match to a new spectral signal

Store the color matching functions in the rows of the matrix, *C*

$$C = \begin{pmatrix} c_1(\lambda_1) & \cdots & c_1(\lambda_N) \\ c_2(\lambda_1) & \cdots & c_2(\lambda_N) \\ c_3(\lambda_1) & \cdots & c_3(\lambda_N) \end{pmatrix}$$

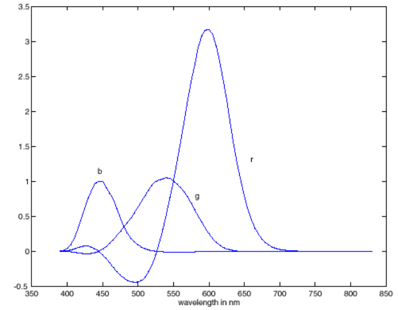

Let the new spectral signal be described by the vector *t*.

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

Then the amounts of each primary needed to match t are: $\vec{e} = C\vec{t}$

The components e_1, e_2, e_3 describe the color of *t*. If you have some other spectral signal, *s*, and *s* matches *t* perceptually, then e_1, e_2, e_3, will also match *s* (by Grassman's Laws)

Source: W. Freeman

Linear color spaces: RGB

- Primaries are monochromatic lights (for monitors, they correspond to the three types of phosphors)
- *Subtractive matching* required for some wavelengths

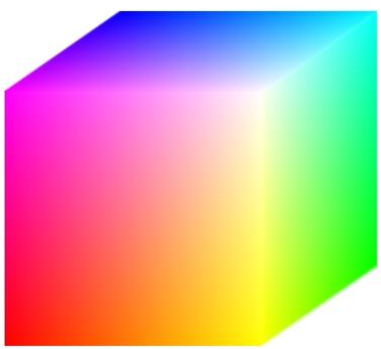

$p_1 = 645.2$ nm
$p_2 = 525.3$ nm
$p_3 = 444.4$ nm

RGB matching functions

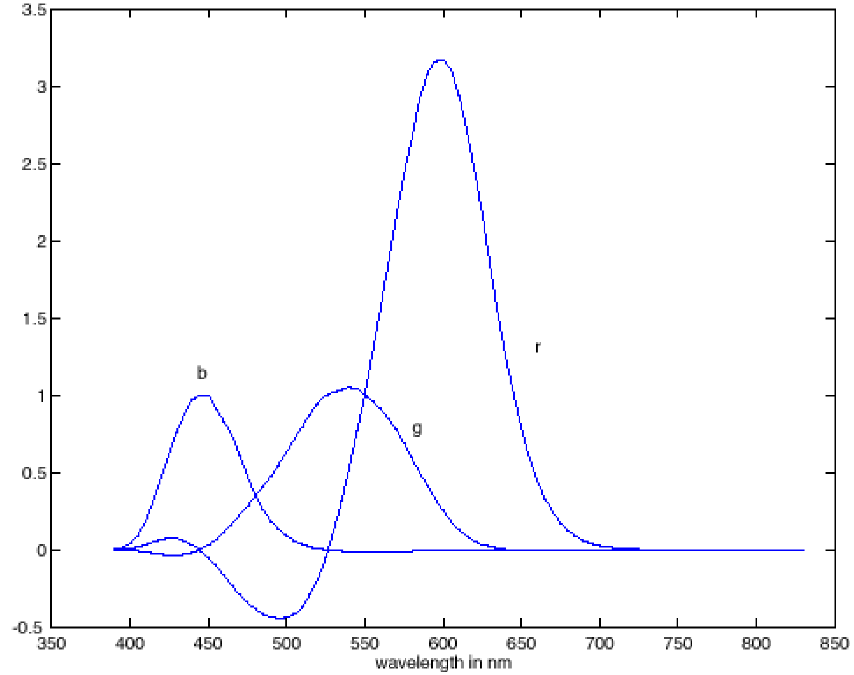

Comparison of color matching functions with best 3x3 transformation of cone responses

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.

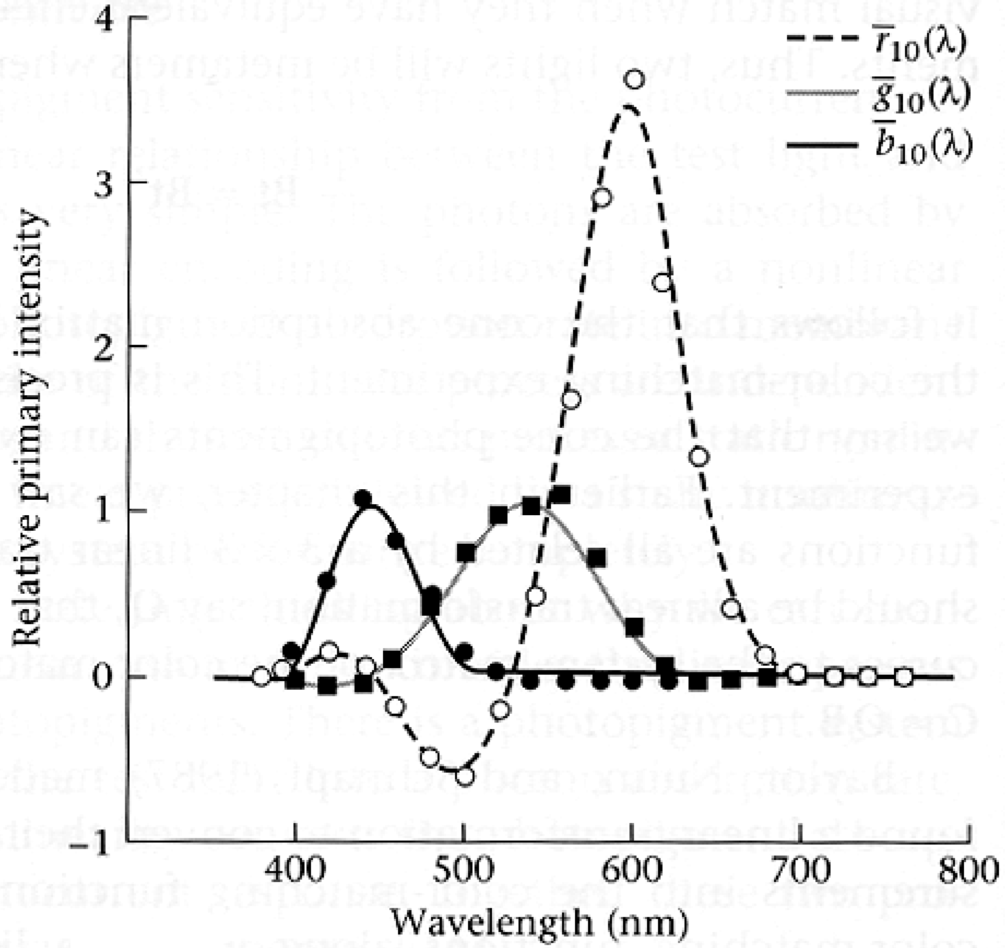

Foundations of Vision, by Brian Wandell, Sinauer Assoc., 1995

Linear color spaces: CIE XYZ

- Established in 1931 by the [International Commission on Illumination]
- Primaries are imaginary, but matching functions are everywhere positive
- 2D visualization: draw (x,y), where $x = X/(X+Y+Z)$, $y = Y/(X+Y+Z)$

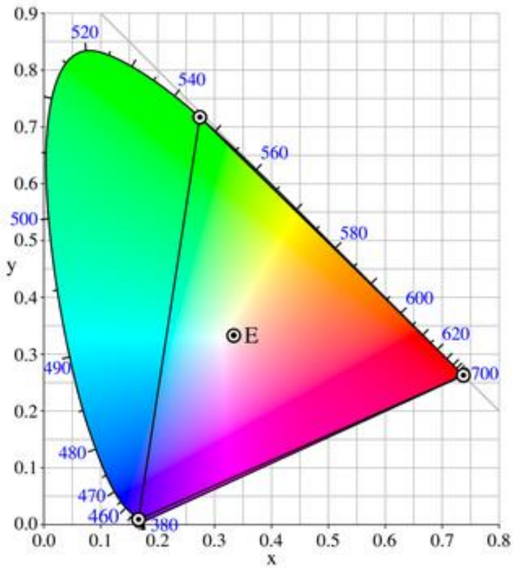

Matching functions

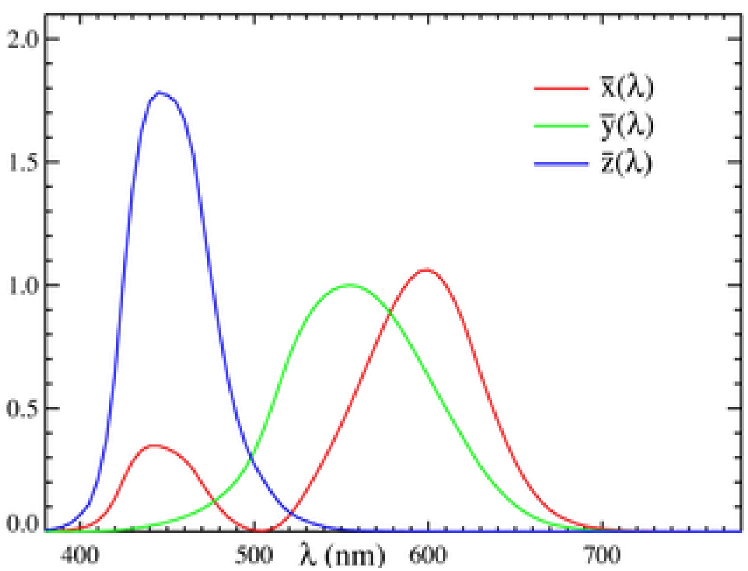

Uniform color spaces

- 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

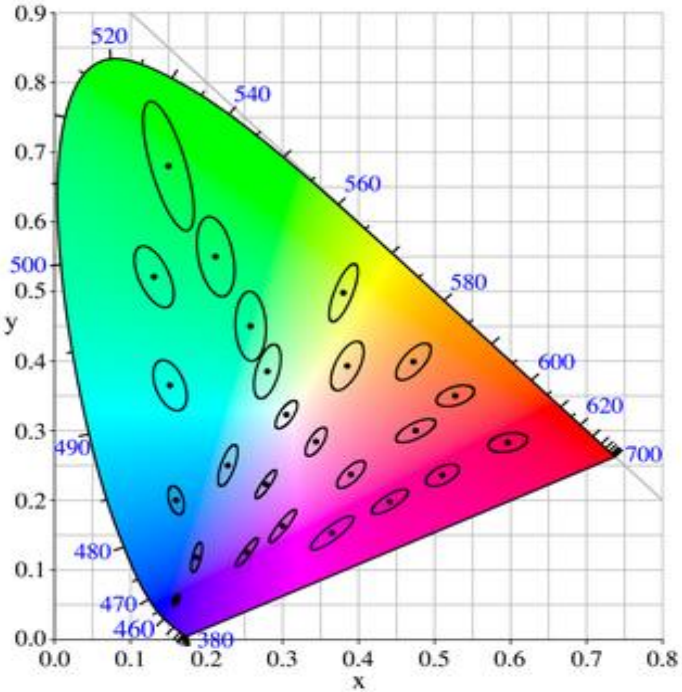

McAdam ellipses: Just noticeable differences in color

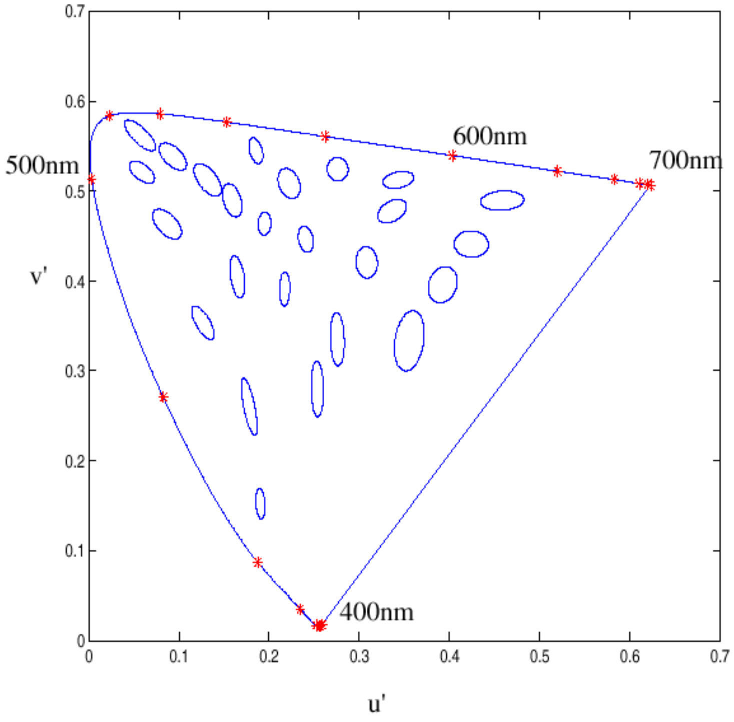

Uniform color spaces

- 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
- Next generation: CIE L*a*b* (Koenderink: “an awful mix of magical numbers and arbitrary functions that somehow ‘fit’ the eye measure”)

Nonlinear color spaces: HSV

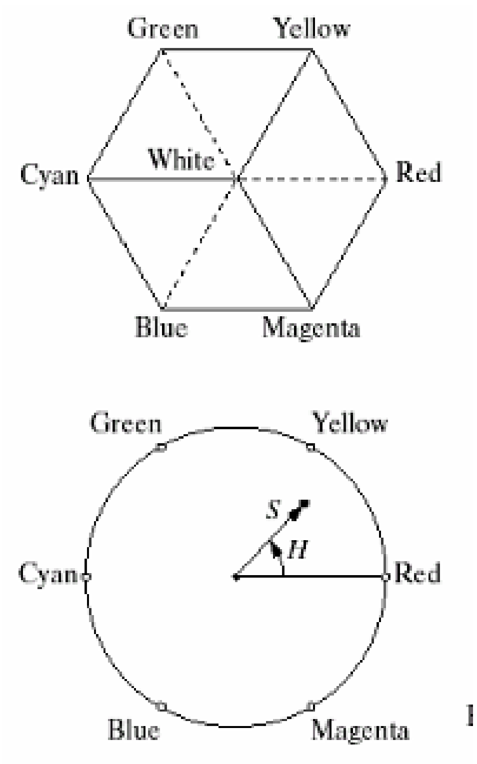

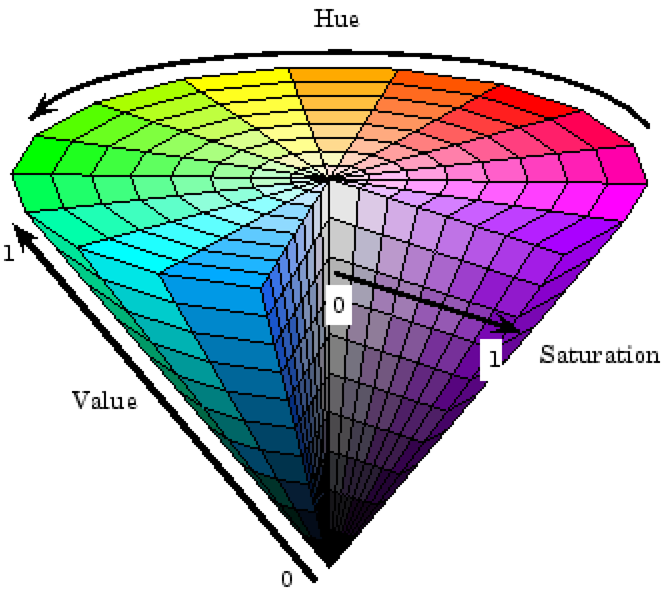

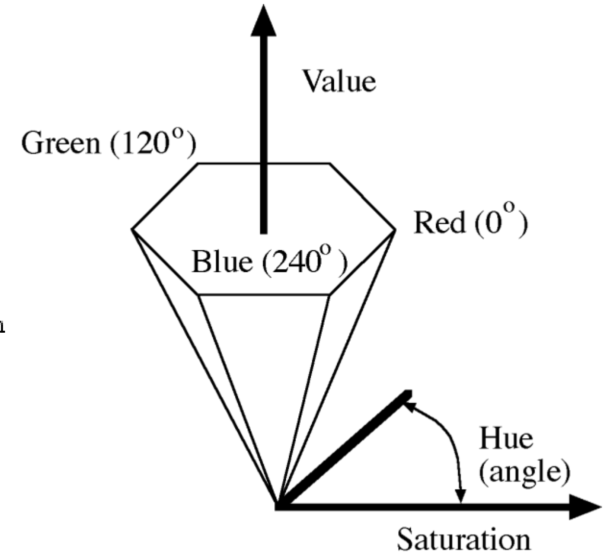

- Perceptually meaningful dimensions: Hue, Saturation, Value (Intensity)
- RGB cube on its vertex

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

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
- Instantaneous effects
  - Simultaneous contrast: background color affects perceived color of the target
  - Mach bands

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
- Instantaneous effects
  - Simultaneous contrast: background color affects perceived color of the target
  - Mach bands
- Gradual effects
  - Light/dark adaptation
  - Chromatic adaptation
  - Afterimages

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
- Instantaneous effects
  - Simultaneous contrast: background color affects perceived color of the target
  - Mach bands
- Gradual effects
  - Light/dark adaptation
  - Chromatic adaptation
  - Afterimages
- High-level effects
  - Color naming

Lightness constancy

White in light and in shadow

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

Slide by F. Durand

Lightness constancy

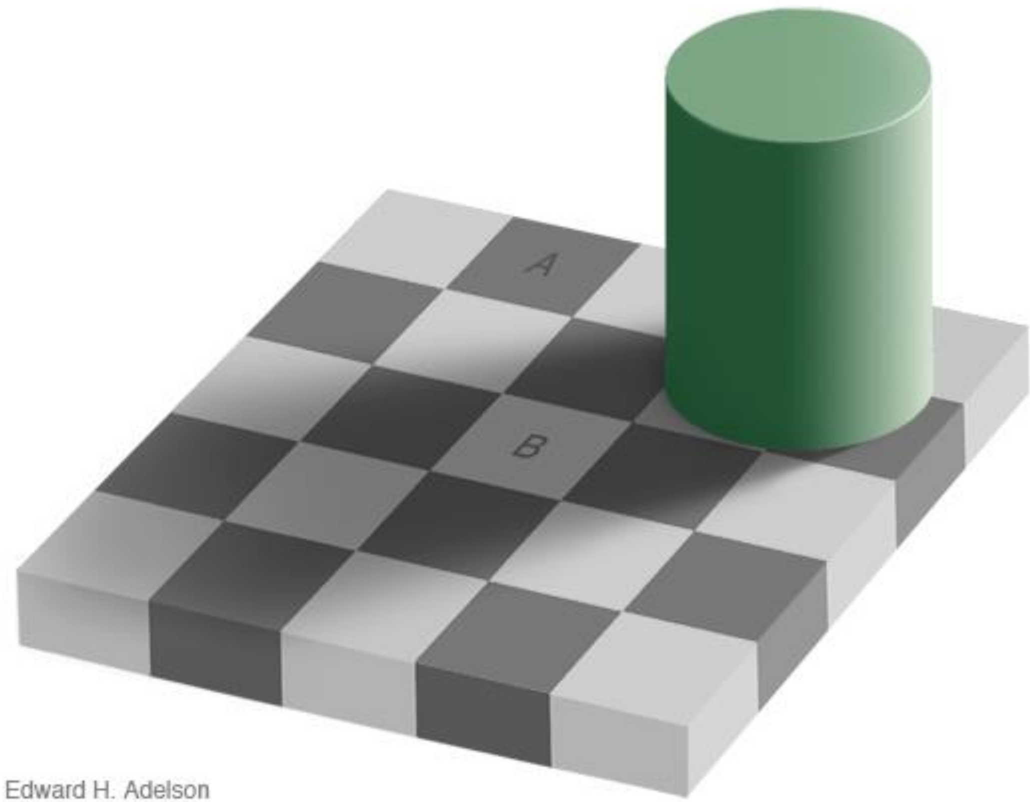

http://web.mit.edu/persci/people/adelson/checkershadow_illusion.html

Lightness constancy

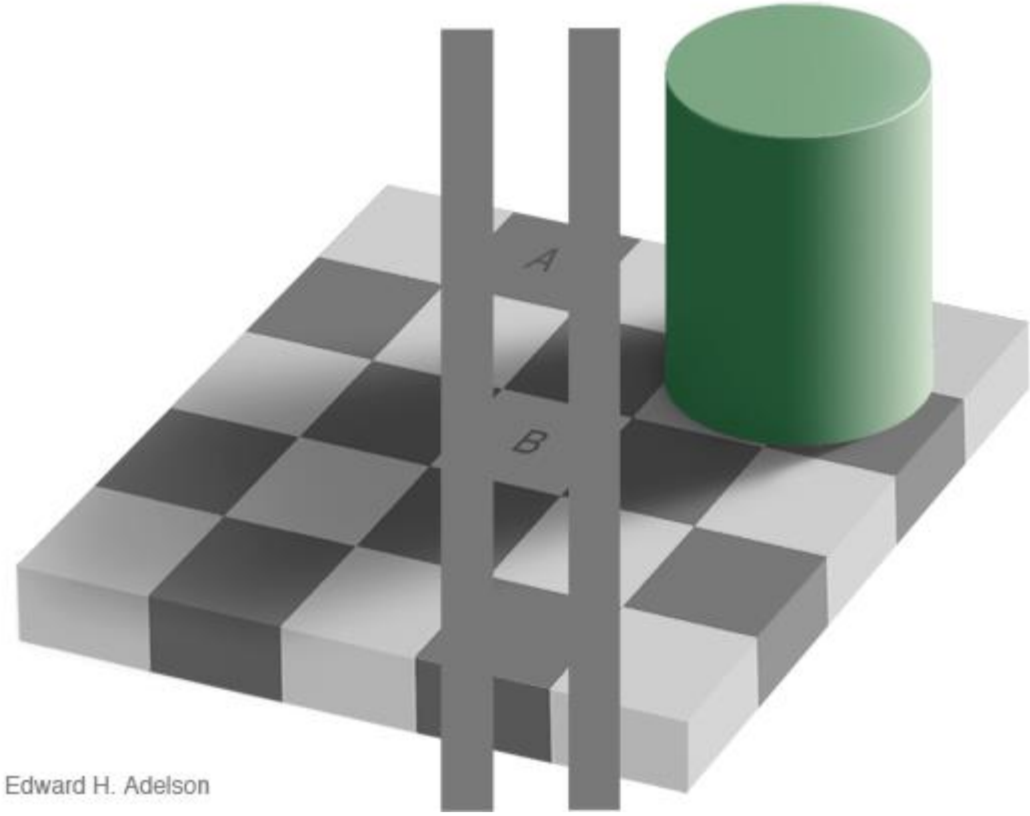

http://web.mit.edu/persci/people/adelson/checkershadow_illusion.html

Lightness constancy

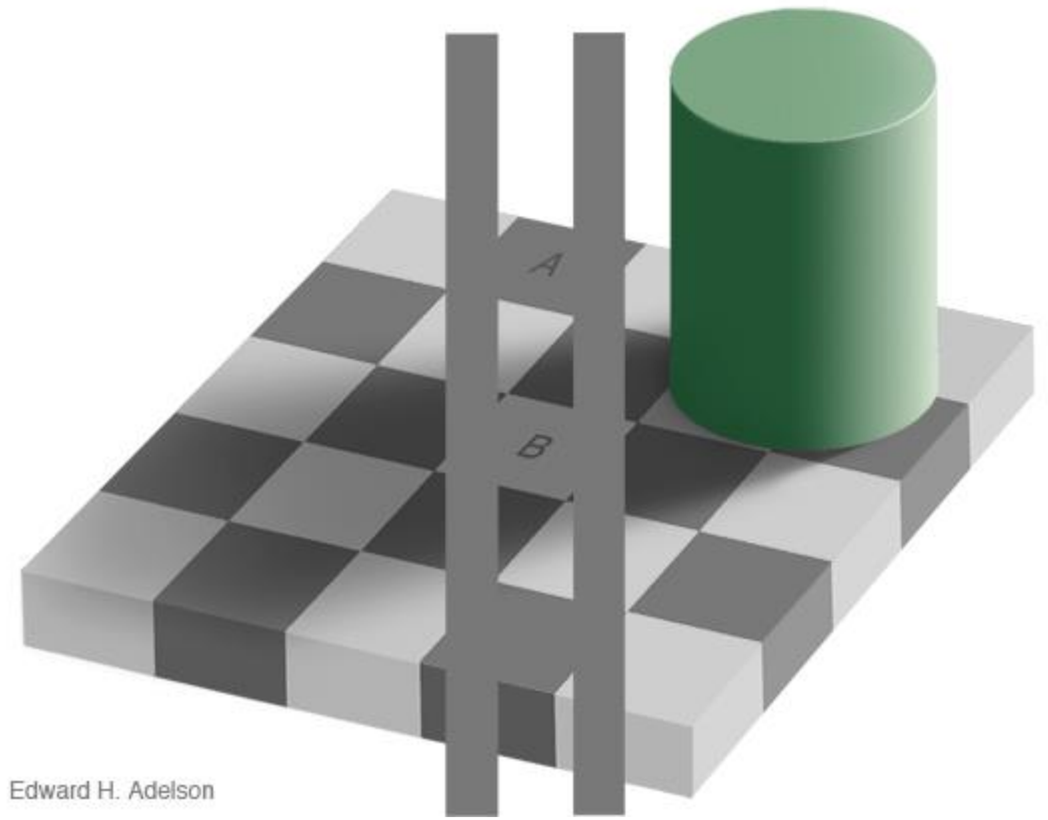

- Possible explanations
  - Simultaneous contrast
  - Reflectance edges vs. illumination edges

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

Simultaneous contrast/Mach bands

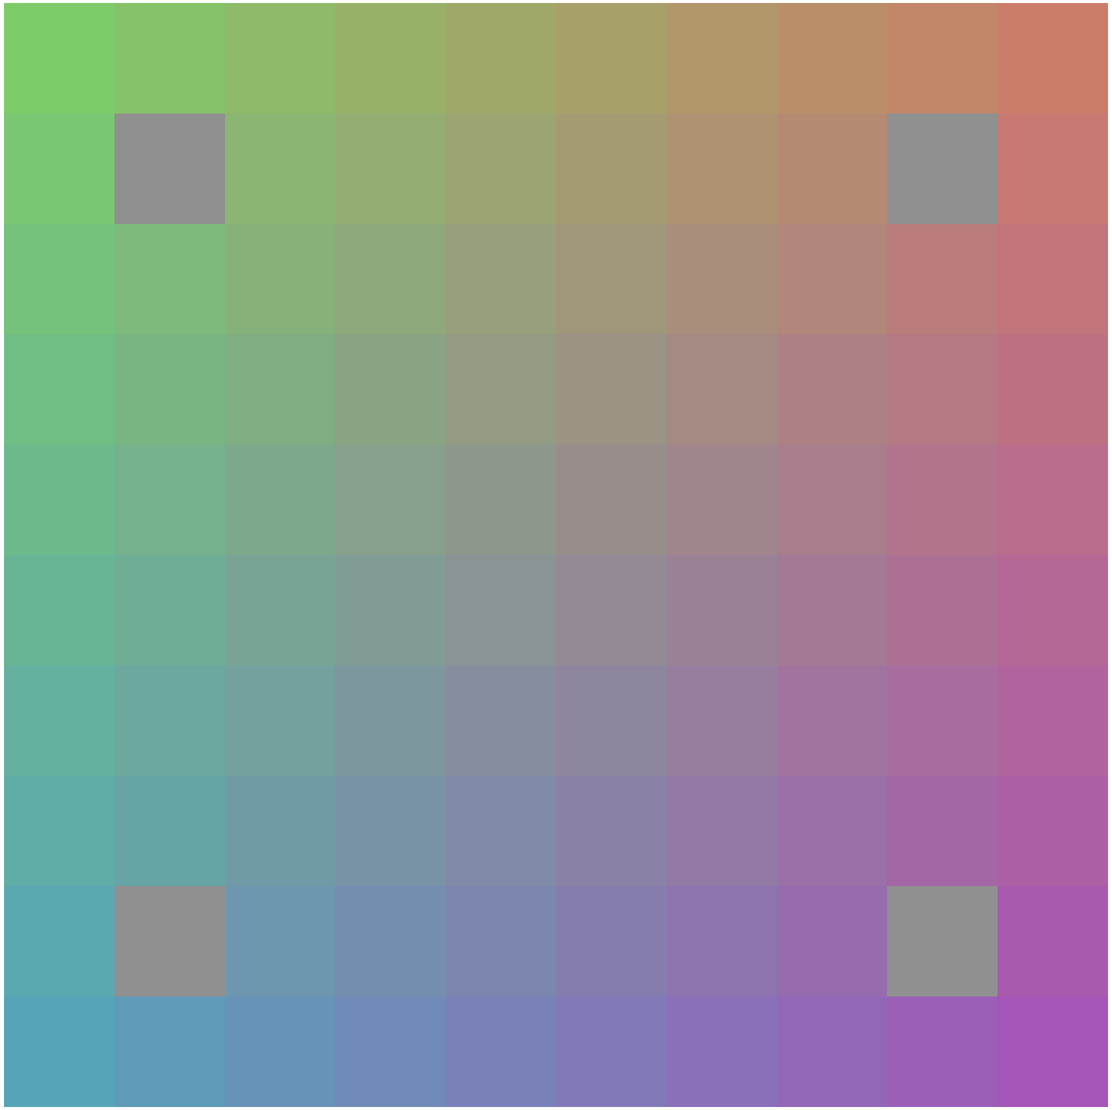

Source: D. Forsyth

Chromatic adaptation

- The visual system changes its sensitivity depending on the luminances prevailing in the visual field
  - The exact mechanism is poorly 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

Chromatic adaptation

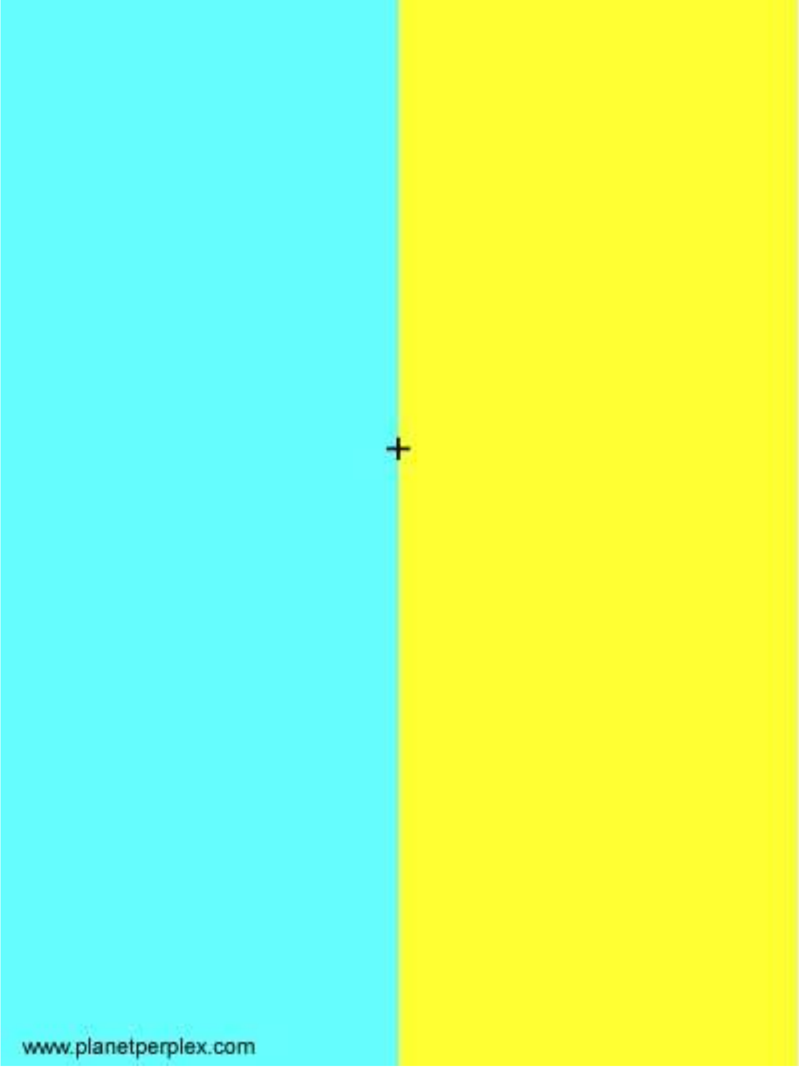

Name that color

Green

Red

Blue

Yellow

Green

Red

Orange

Purple

Brown

Yellow

Black

Useful reference

Stephen E. Palmer, **Vision Science: Photons to Phenomenology**, MIT Press, 1999

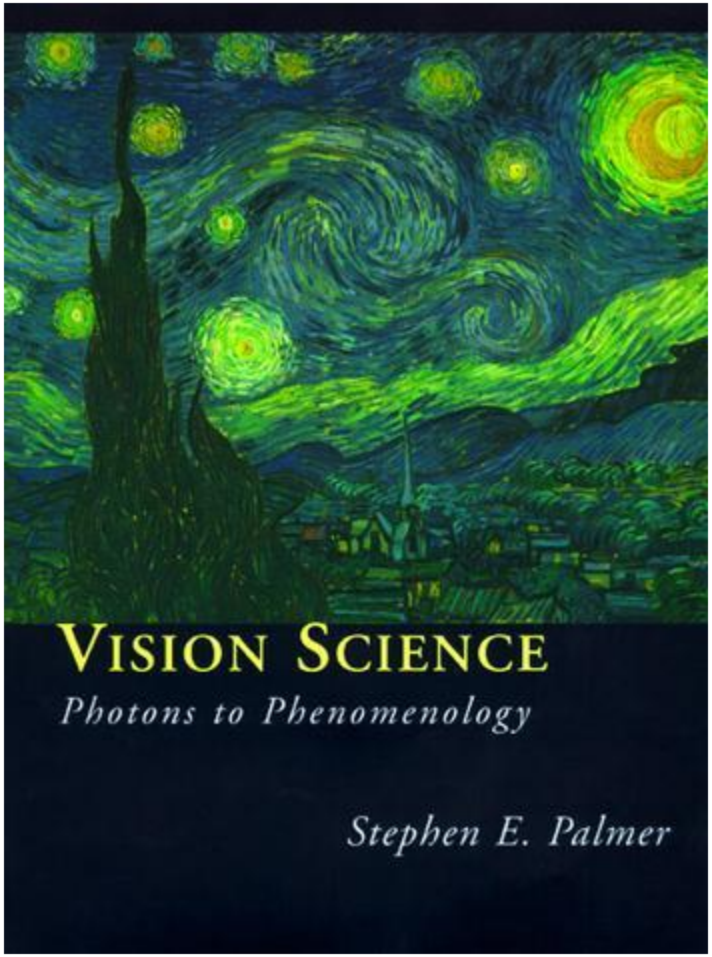

White balance

- When looking at a picture on screen or print, we adapt to the illuminant of the room, not to that of the scene in the picture
- When the white balance is not correct, the picture will have an unnatural color “cast”

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

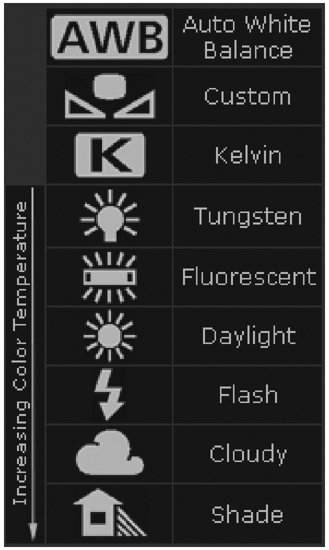

http://www.cambridgeincolour.com/tutorials/white-balance.htm

White balance can be tricky…

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

Reference: moon

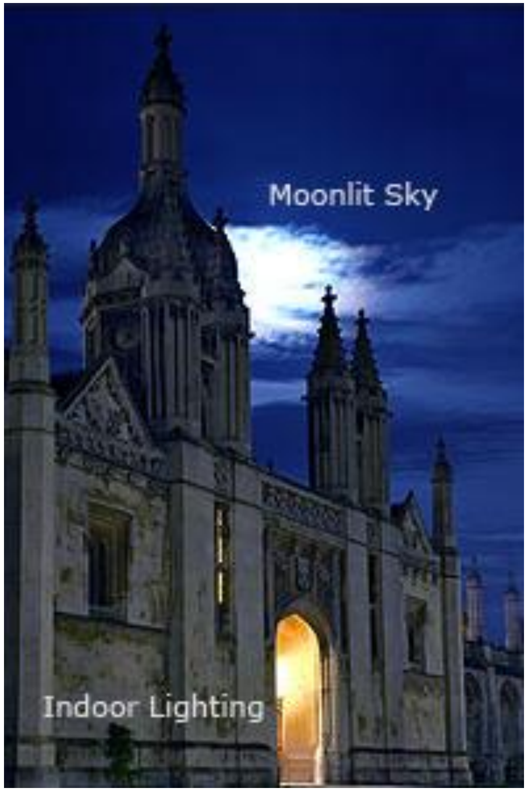

Reference: stone

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

White balance

- Von Kries adaptation
  - Multiply each channel by a gain factor
  - Note that the light source could have a more complex effect, corresponding to an arbitrary 3x3 matrix
- Best way: gray card
  - Take a picture of a neutral object (white or gray)
  - Deduce the weight of each channel
    - If the object is recoded 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

White balance

- Without gray cards: we need to “guess” which pixels correspond to white objects
- Gray world assumption
  - The image average r_{ave}, g_{ave}, b_{ave} is gray
  - Use weights $1/r_{ave}$, $1/g_{ave}$, $1/b_{ave}$

White balance

- Without gray cards: we need to “guess” which pixels correspond to white objects
- Gray world assumption
  - The image average r_{ave}, g_{ave}, b_{ave} is gray
  - Use weights $1/r_{ave}$, $1/g_{ave}$, $1/b_{ave}$
- Brightest pixel assumption
  - Highlights usually have the color of the light source
  - Use weights inversely proportional to the values of the brightest pixels

White balance

- Without gray cards: we need to “guess” which pixels correspond to white objects
- Gray world assumption
  - The image average r_{ave}, g_{ave}, b_{ave} is gray
  - Use weights $1/r_{ave}$, $1/g_{ave}$, $1/b_{ave}$
- 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

White balance

- Without gray cards: we need to “guess” which pixels correspond to white objects
- Gray world assumption
  - The image average r_{ave}, g_{ave}, b_{ave} is gray
  - Use weights $1/r_{ave}$, $1/g_{ave}$, $1/b_{ave}$
- 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

White balance by recognition

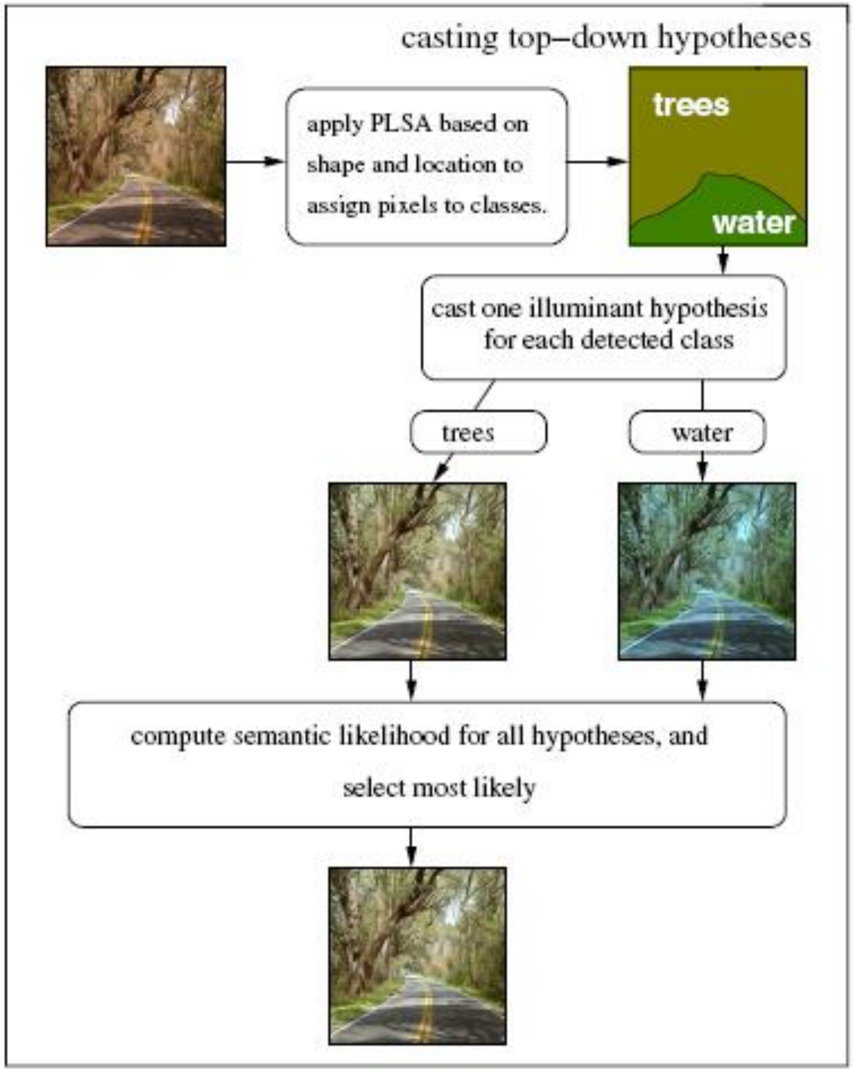

- Key idea: For each of the semantic classes present in the image, compute the illuminant that transforms the pixels assigned to that class so that the average color of that class matches the average color of the same class in a database of "typical" images

J. Van de Weijer, C. Schmid and J. Verbeek, [Using High-Level Visual Information for Color Constancy](), ICCV 2007.

Uses of color in computer vision

Color histograms for indexing and retrieval

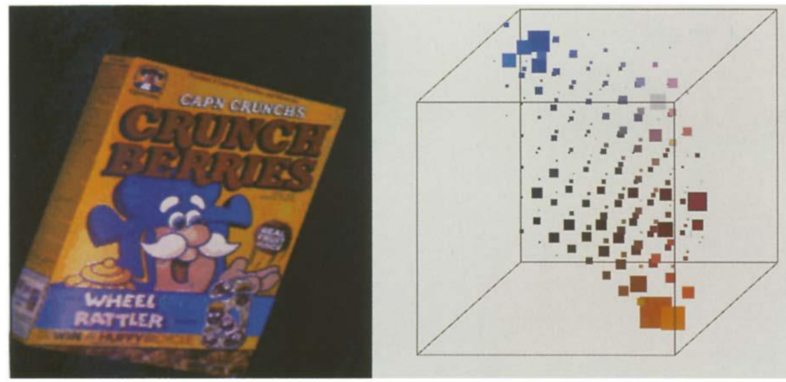

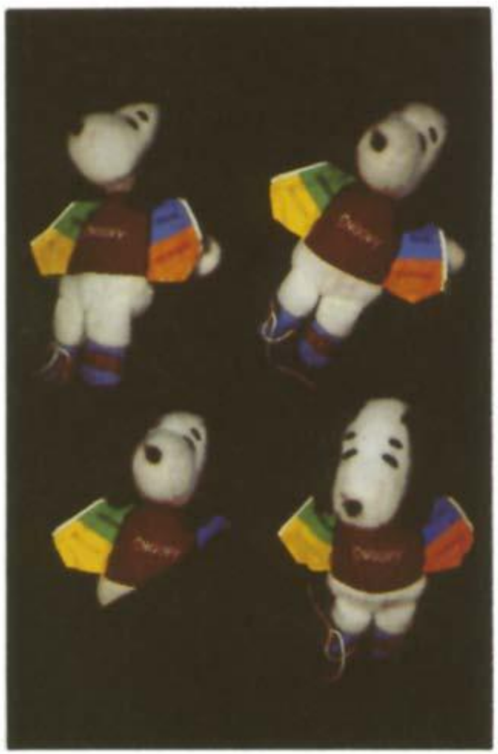

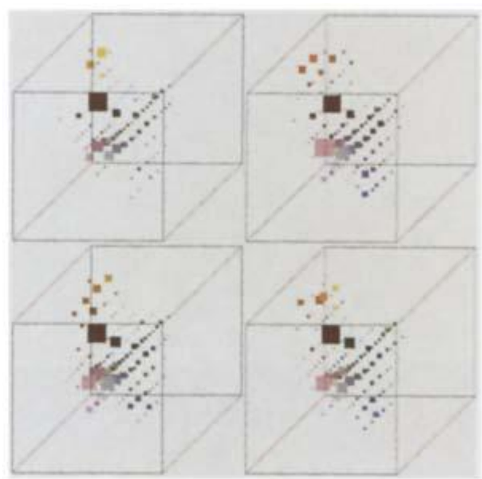

Swain and Ballard, [Color Indexing](), IJCV 1991.

Uses of color in computer vision

Skin detection

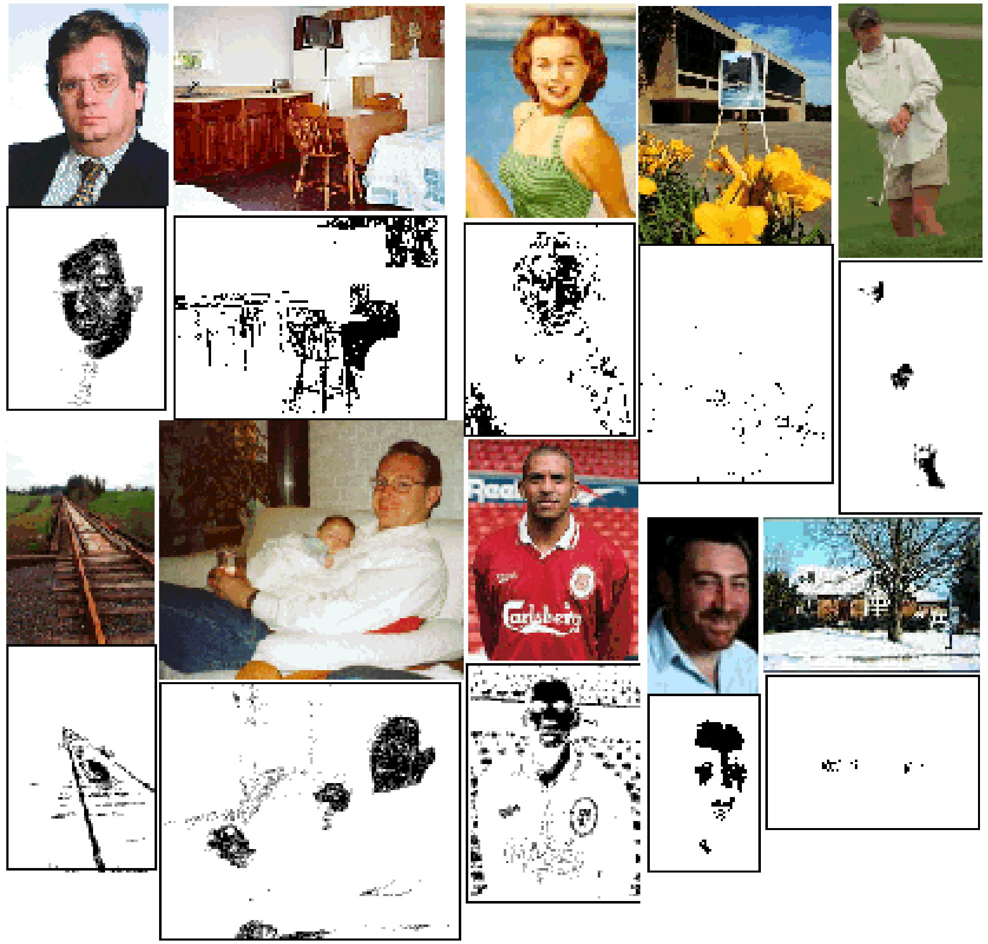

M. Jones and J. Rehg, [Statistical Color Models with Application to Skin Detection](), IJCV 2002.

Uses of color in computer vision

Image segmentation and retrieval

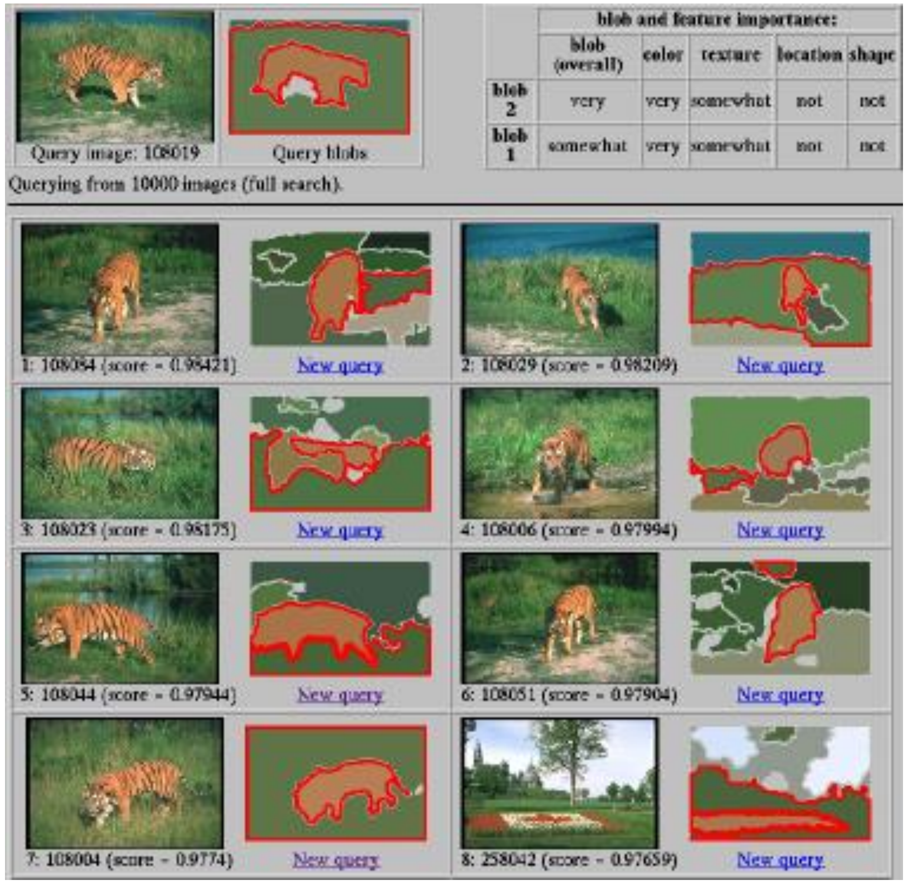

C. Carson, S. Belongie, H. Greenspan, and Ji. Malik, Blobworld: Image segmentation using Expectation-Maximization and its application to image querying, ICVIS 1999.

Uses of color in computer vision

Robot soccer

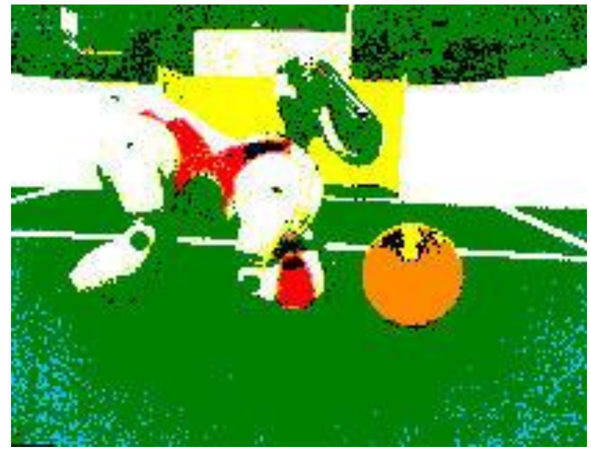

M. Sridharan and P. Stone, Towards Eliminating Manual Color Calibration at RoboCup. RoboCup-2005: Robot Soccer World Cup IX, Springer Verlag, 2006

Source: K. Grauman

Uses of color in computer vision

Building appearance models for tracking

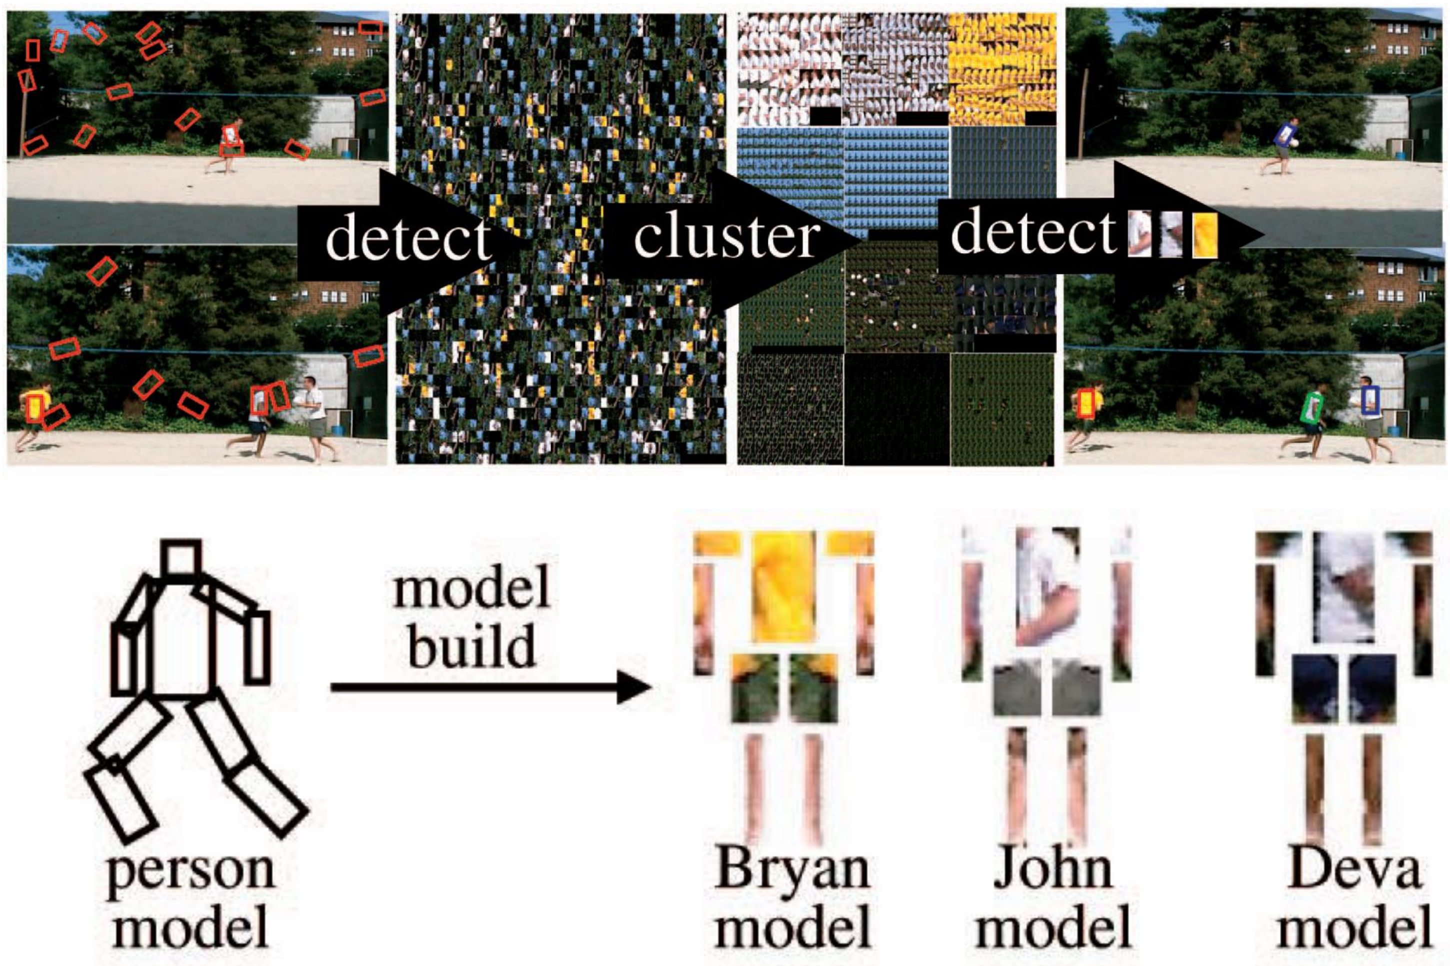

D. Ramanan, D. Forsyth, and A. Zisserman. [Tracking People by Learning their Appearance](). PAMI 2007.

Uses of color in computer vision

Judging visual realism

J.-F. Lalonde and A. Efros. Using Color Compatibility for Assessing Image Realism. ICCV 2007.

Next time: Linear filtering and edge detection