Color and color constancy

6.869, MIT
Bill Freeman
Antonio Torralba

Feb. 22, 2011

Why does a visual system need color?

http://www.hobbylinc.com/gr/pll/pll5019.jpg

Why does a visual system need color?
(an incomplete list…)

Why does a visual system need color? (an incomplete list…)

- To tell what food is edible.

Why does a visual system need color? (an incomplete list…)

- To tell what food is edible.
- To distinguish material changes from shading changes.

Why does a visual system need color? (an incomplete list…)

- To tell what food is edible.
- To distinguish material changes from shading changes.
- To group parts of one object together in a scene.

Why does a visual system need color? (an incomplete list…)

- To tell what food is edible.
- To distinguish material changes from shading changes.
- To group parts of one object together in a scene.
- To find people's skin.

Why does a visual system need color? (an incomplete list…)

- To tell what food is edible.
- To distinguish material changes from shading changes.
- To group parts of one object together in a scene.
- To find people’s skin.
- Check whether a person’s appearance looks normal/healthy.

Lecture outline

- Color physics.
- Color perception.

Lecture outline

- Color physics.
- Color perception.

Color

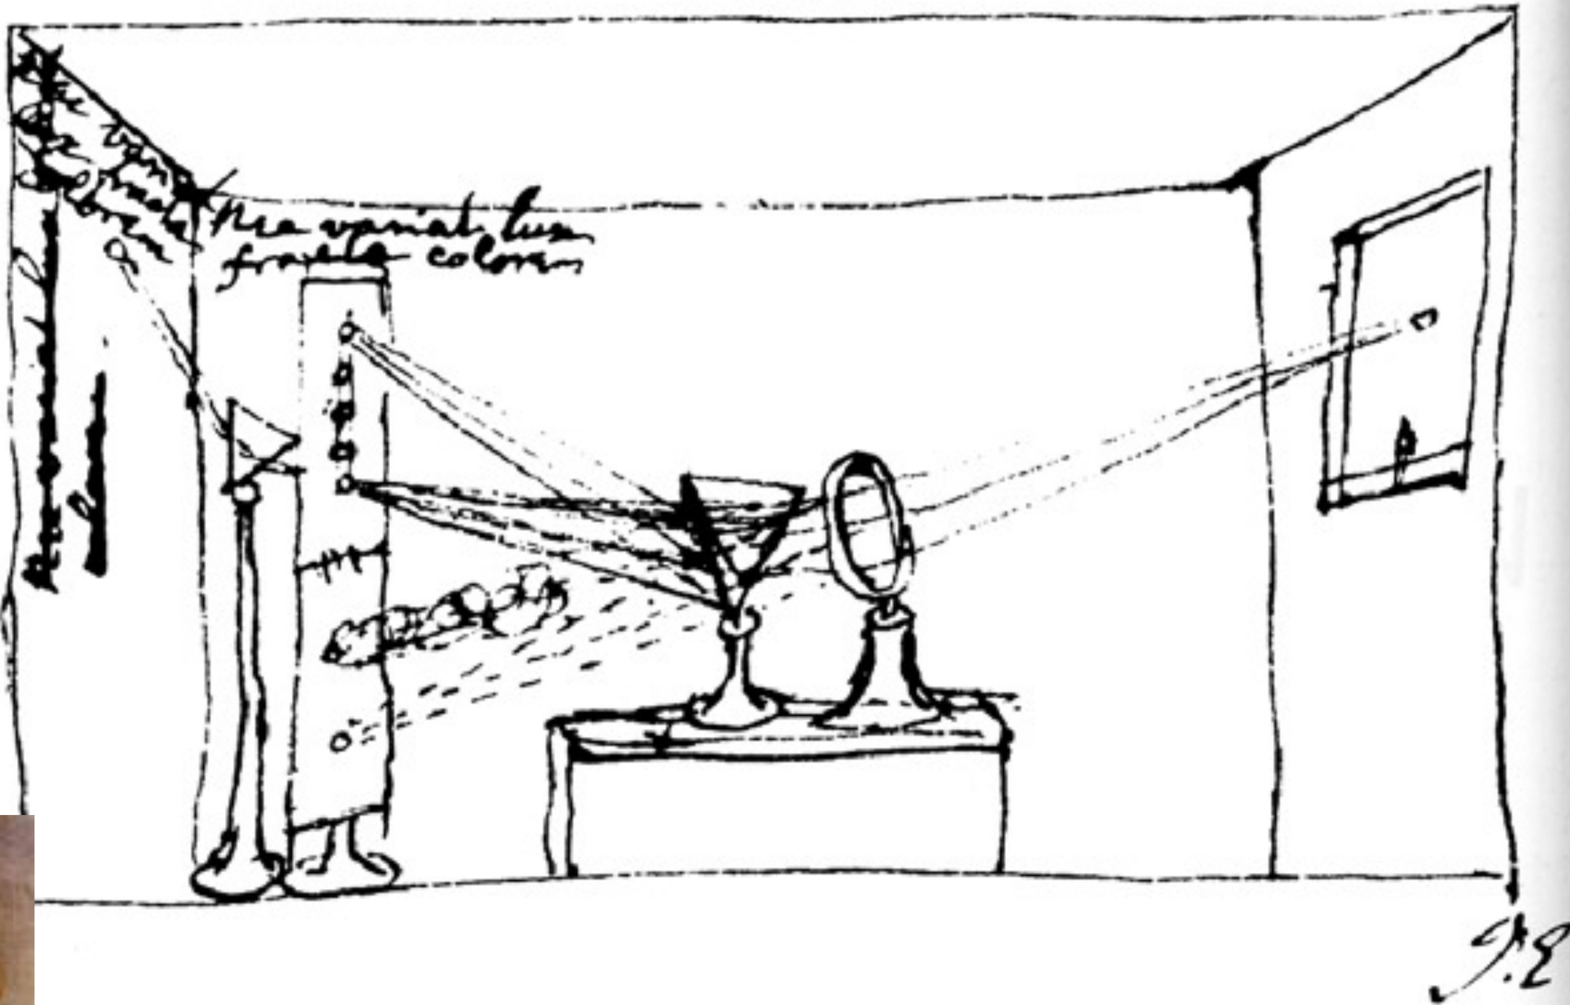

www.popularpersons.org

4.1 NEWTON'S SUMMARY DRAWING of his experiments with light. Using a point source of light and a prism, Newton separated sunlight into its fundamental components. By reconverging the rays, he also showed that the decomposition is reversible.

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

Spectral colors

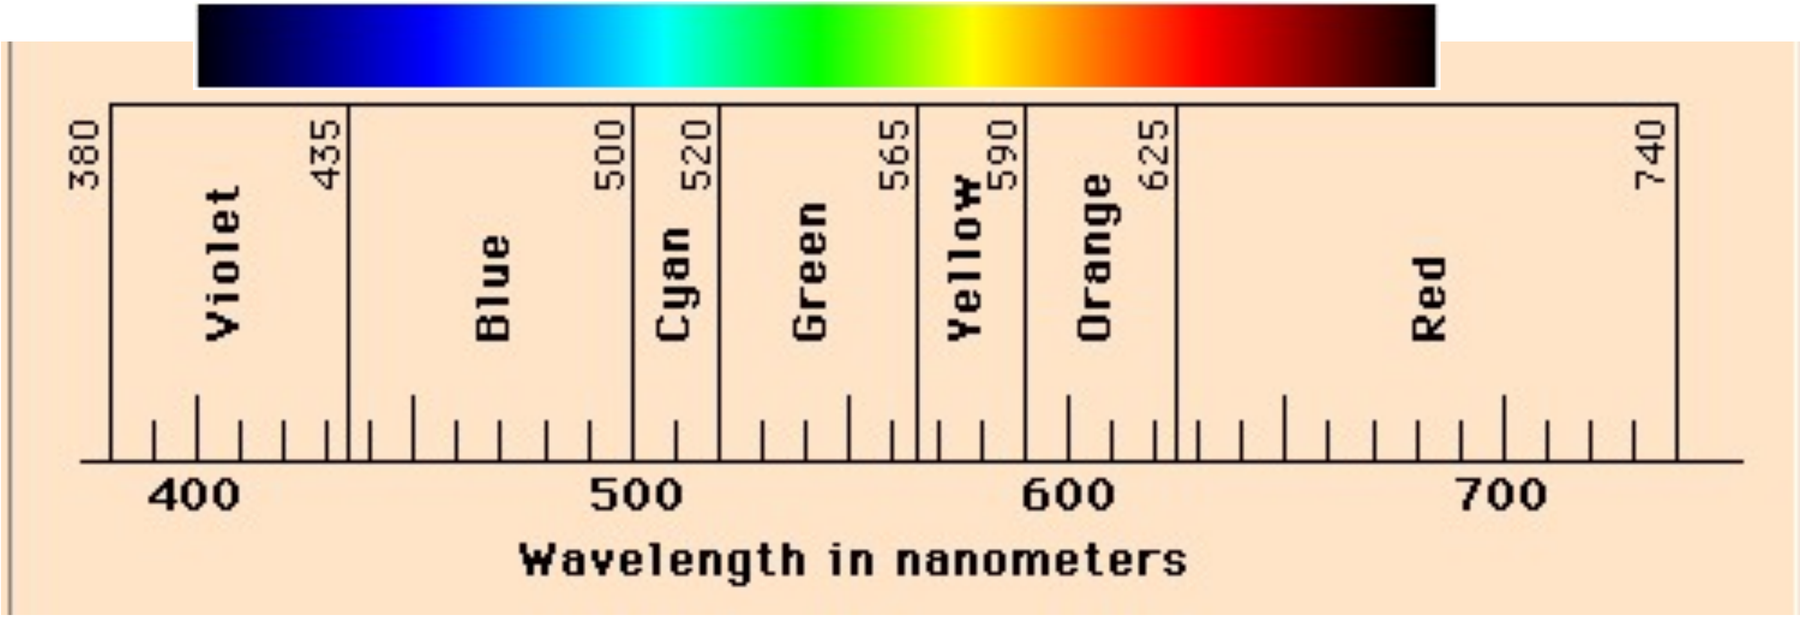

http://hyperphysics.phy-astr.gsu.edu/hbase/vision/specol.html#c2

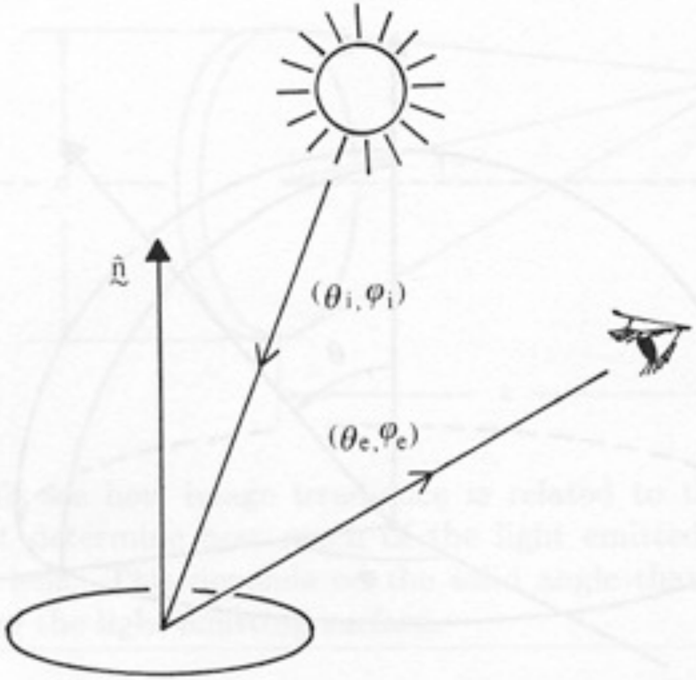

Horn, 1986

Figure 10-7. The bidirectional reflectance distribution function is the ratio of the radiance of the surface patch as viewed from the direction (θ_e, ϕ_e) to the irradiance resulting from illumination from the direction (θ_i, ϕ_i).

Radiometry for color

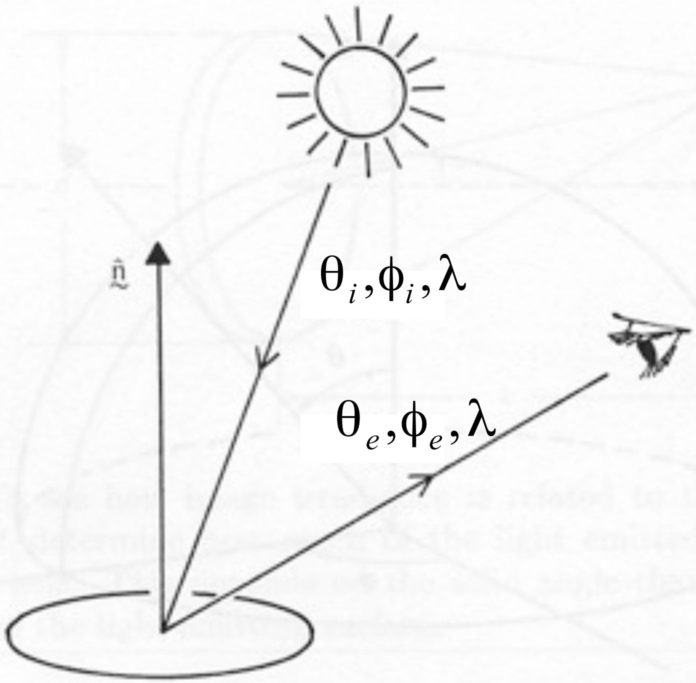

Horn, 1986

Figure 10-7. The bidirectional reflectance distribution function is the ratio of the radiance of the surface patch as viewed from the direction (θ_e, ϕ_e) to the irradiance resulting from illumination from the direction (θ_i, ϕ_i).

Radiometry for color

Spectral radiance: power in a specified direction, per unit area, per unit solid angle, per unit wavelength

$$BRDF = f(\theta_i, \phi_i, \theta_e, \phi_e, \lambda) = \frac{L(\theta_e, \phi_e, \lambda)}{E(\theta_i, \phi_i, \lambda)}$$

Spectral irradiance: incident power per unit area, per unit wavelength

Simplified rendering models: BRDF ➔ reflectance

For diffuse reflections, we replace the BRDF calculation with a wavelength-by-wavelength scalar multiplier

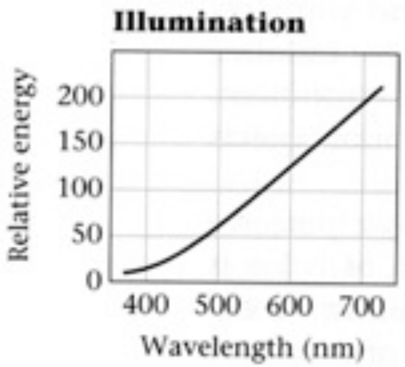

.*

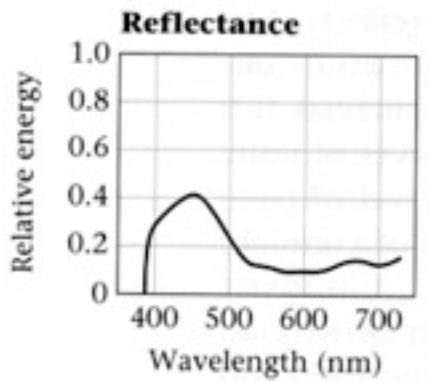

=

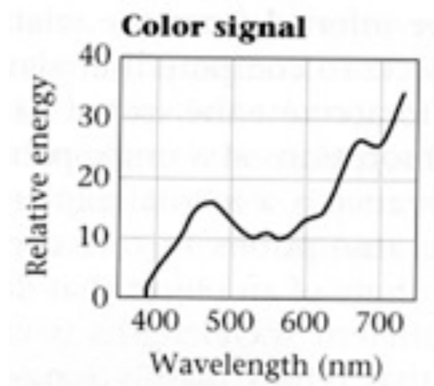

Foundations of Vision, by Brian Wandell, Sinauer Assoc., 1995

Simplified rendering models: transmittance

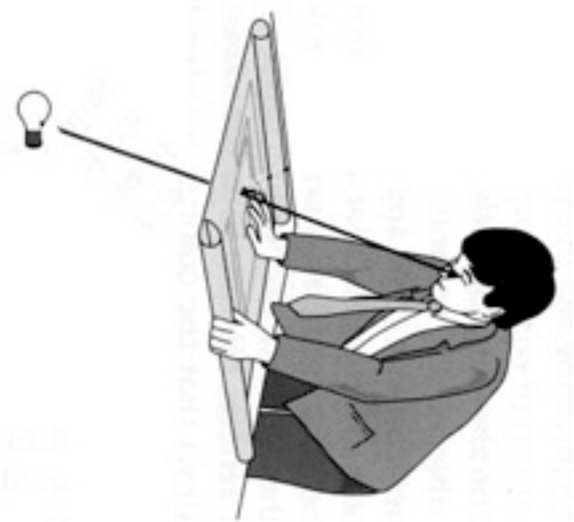

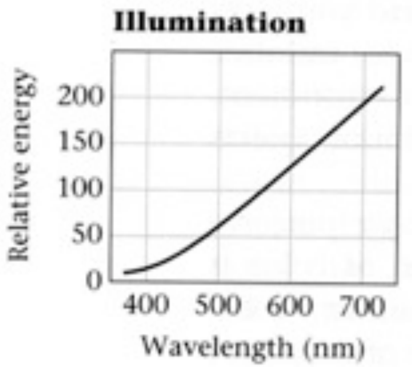

.*

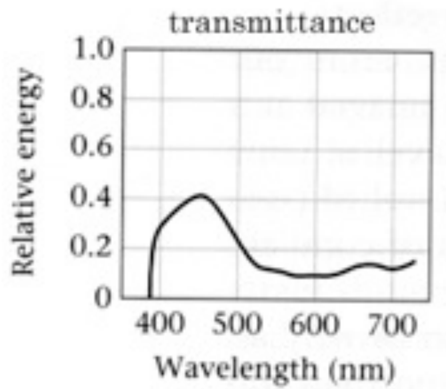

=

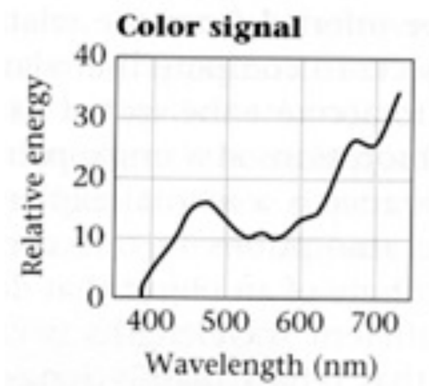

Foundations of Vision, by Brian Wandell, Sinauer Assoc., 1995

Two illumination spectra

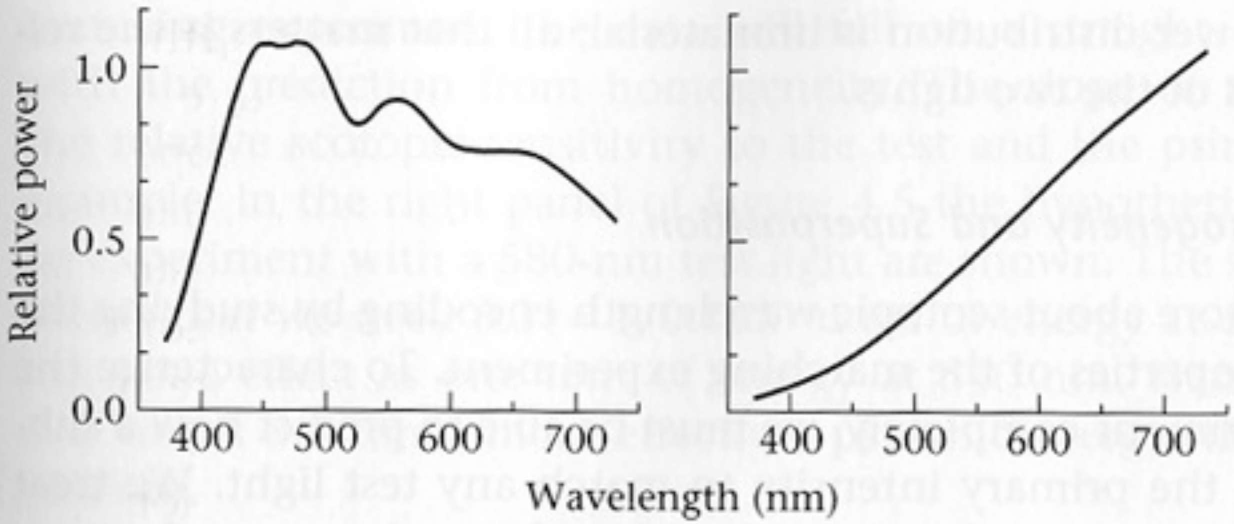

4.4 THE SPECTRAL POWER DISTRIBUTION of two important light sources are shown: (left) blue skylight and (right) a tungsten bulb.

Blue sky

Tungsten light bulb

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

Some reflectance spectra

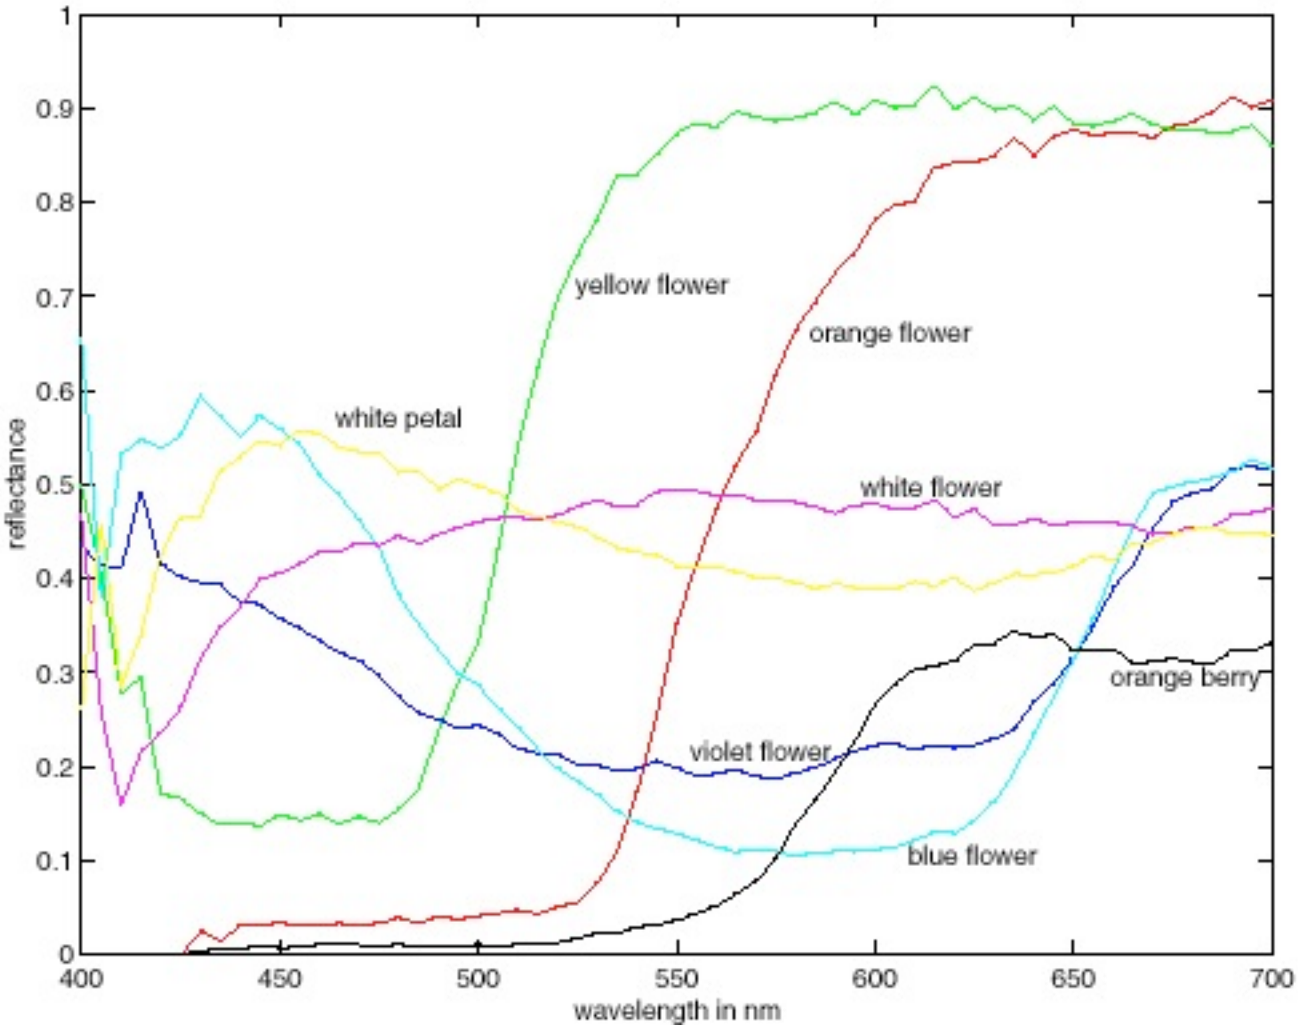

Spectral albedoes for several different leaves, with color names attached. Notice that different colours typically have different spectral albedo, but that different spectral albedoes may result in the same perceived color (compare the two whites). Spectral albedoes are typically quite smooth functions. Measurements by E.Koivisto.

Forsyth, 2002

Color names for cartoon spectra

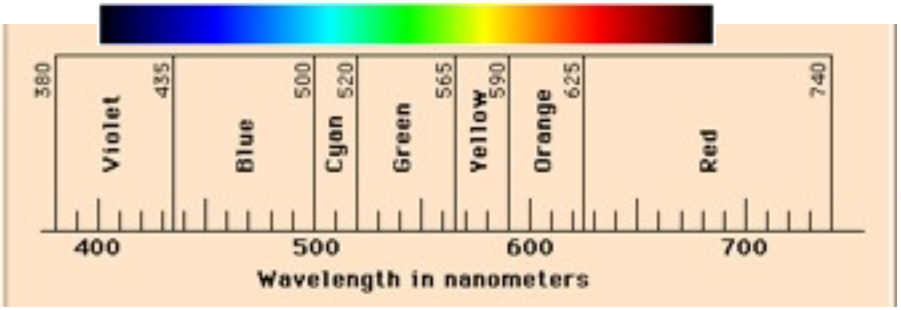

Color names for cartoon spectra

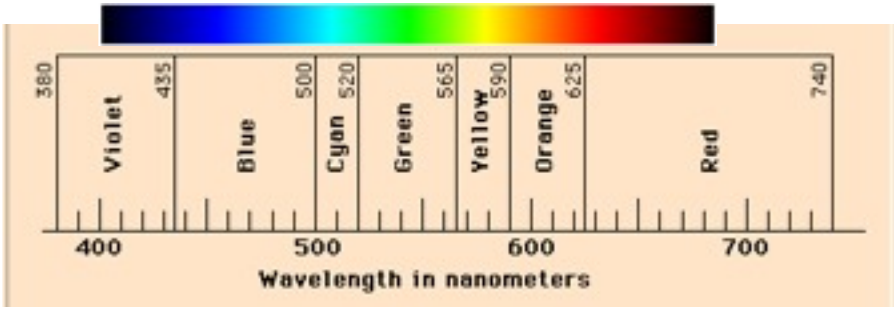

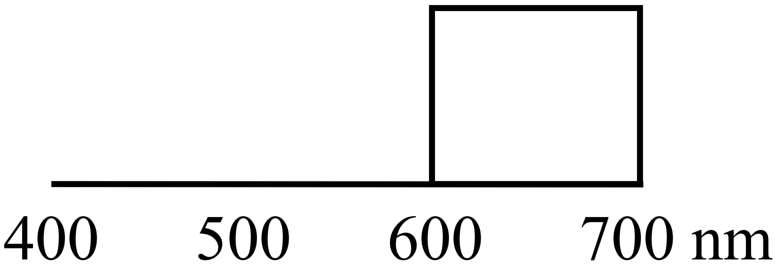

Color names for cartoon spectra

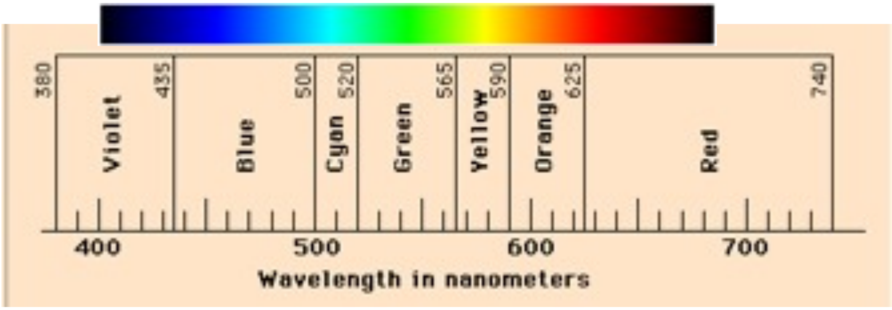

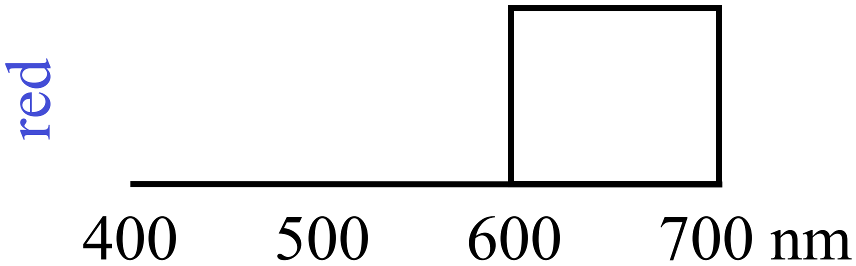

Color names for cartoon spectra

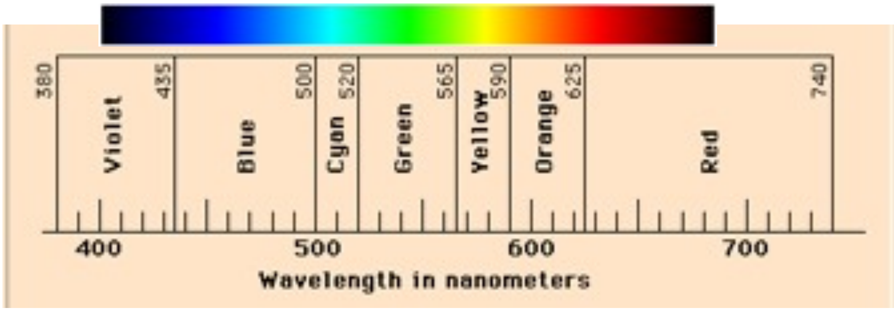

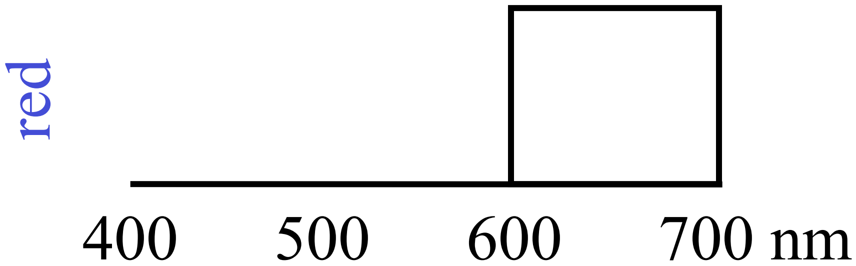

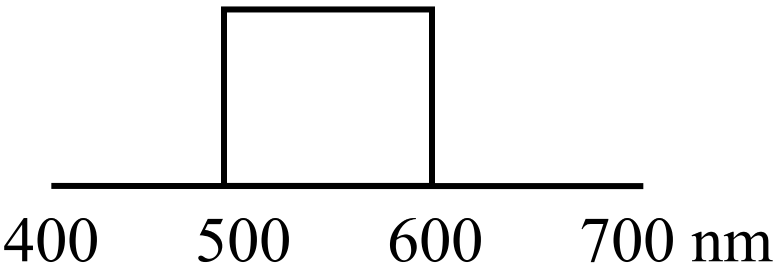

Color names for cartoon spectra

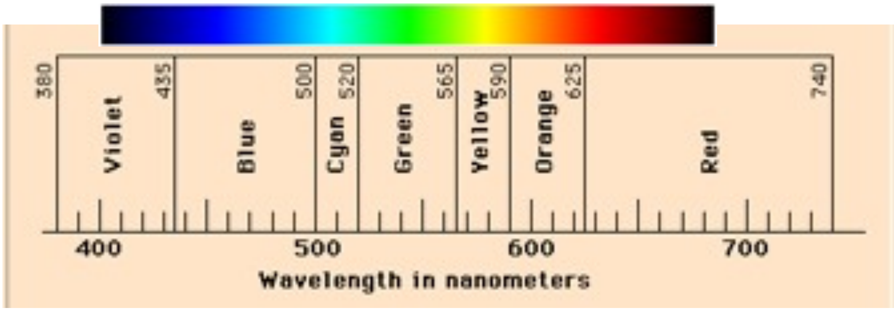

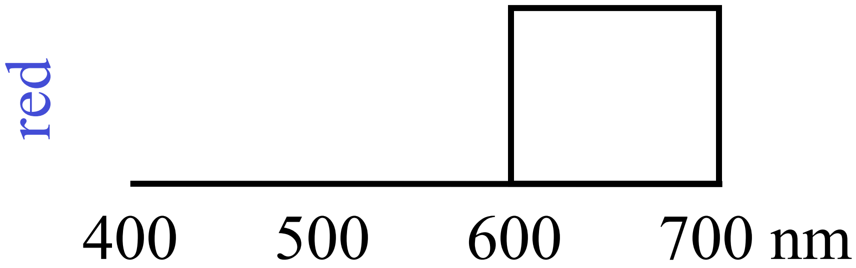

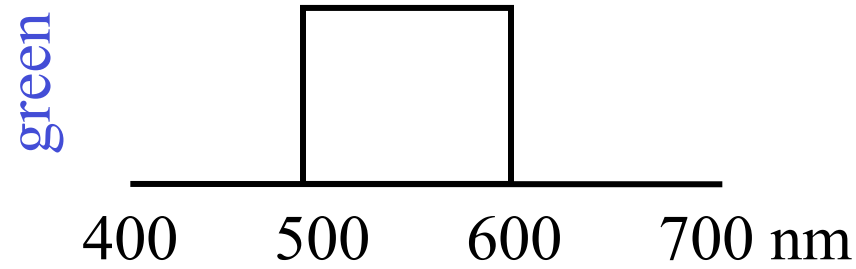

Color names for cartoon spectra

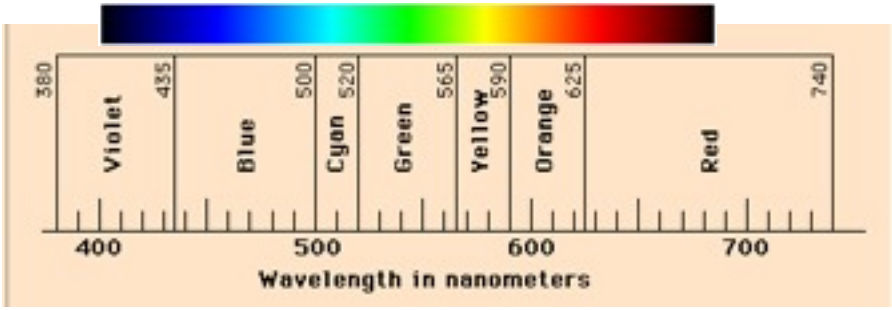

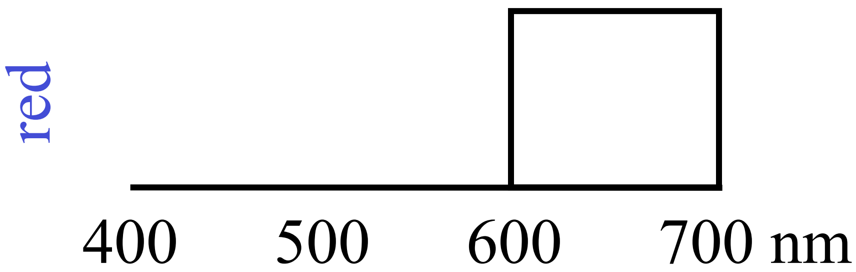

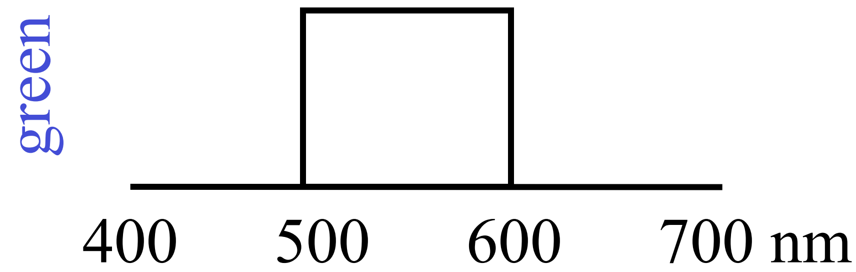

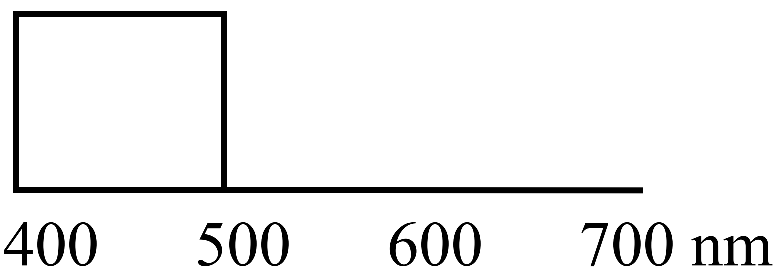

Color names for cartoon spectra

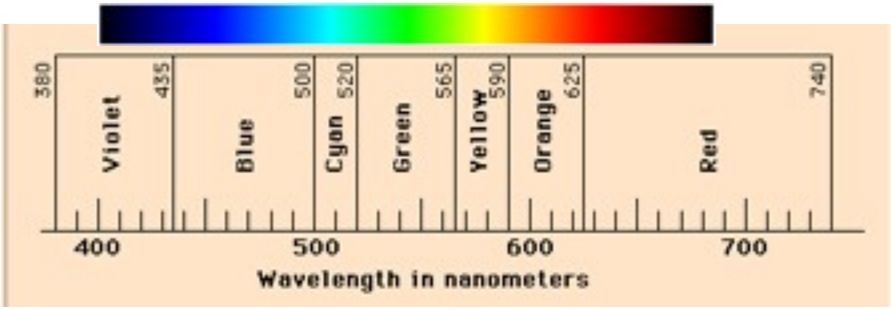

red

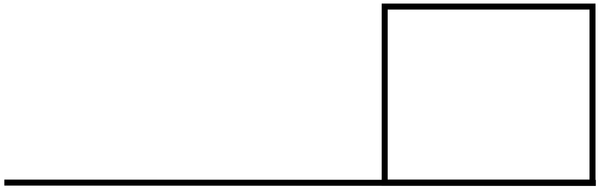

400 500 600 700 nm

green

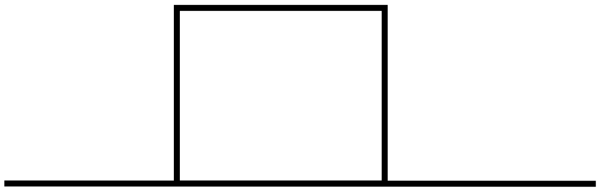

400 500 600 700 nm

blue

400 500 600 700 nm

Color names for cartoon spectra

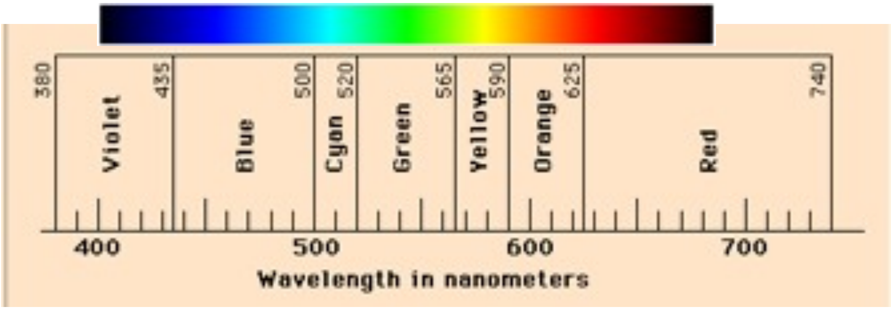

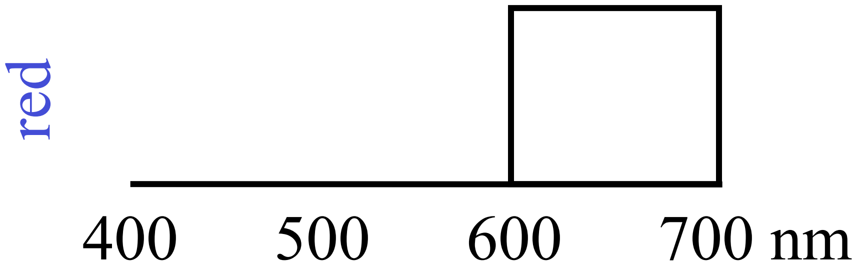

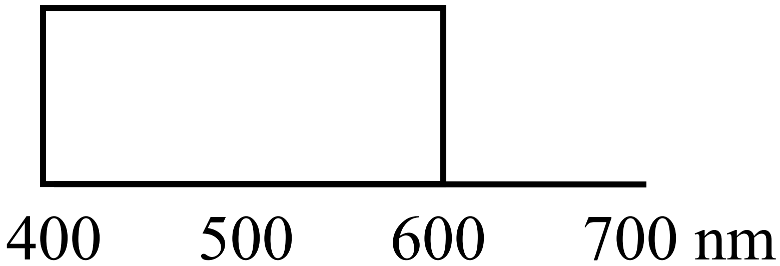

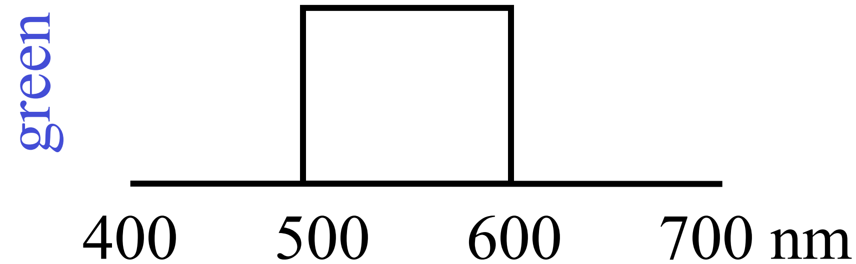

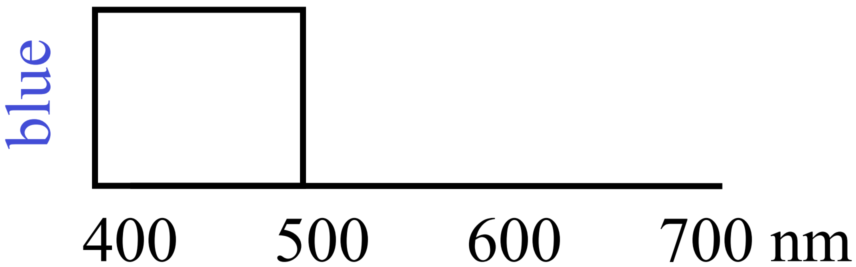

Color names for cartoon spectra

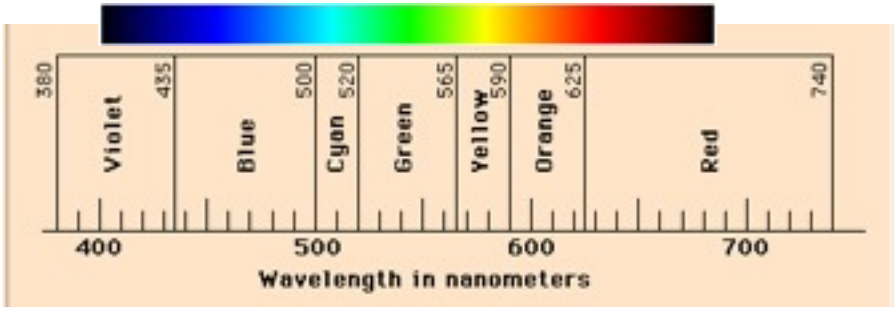

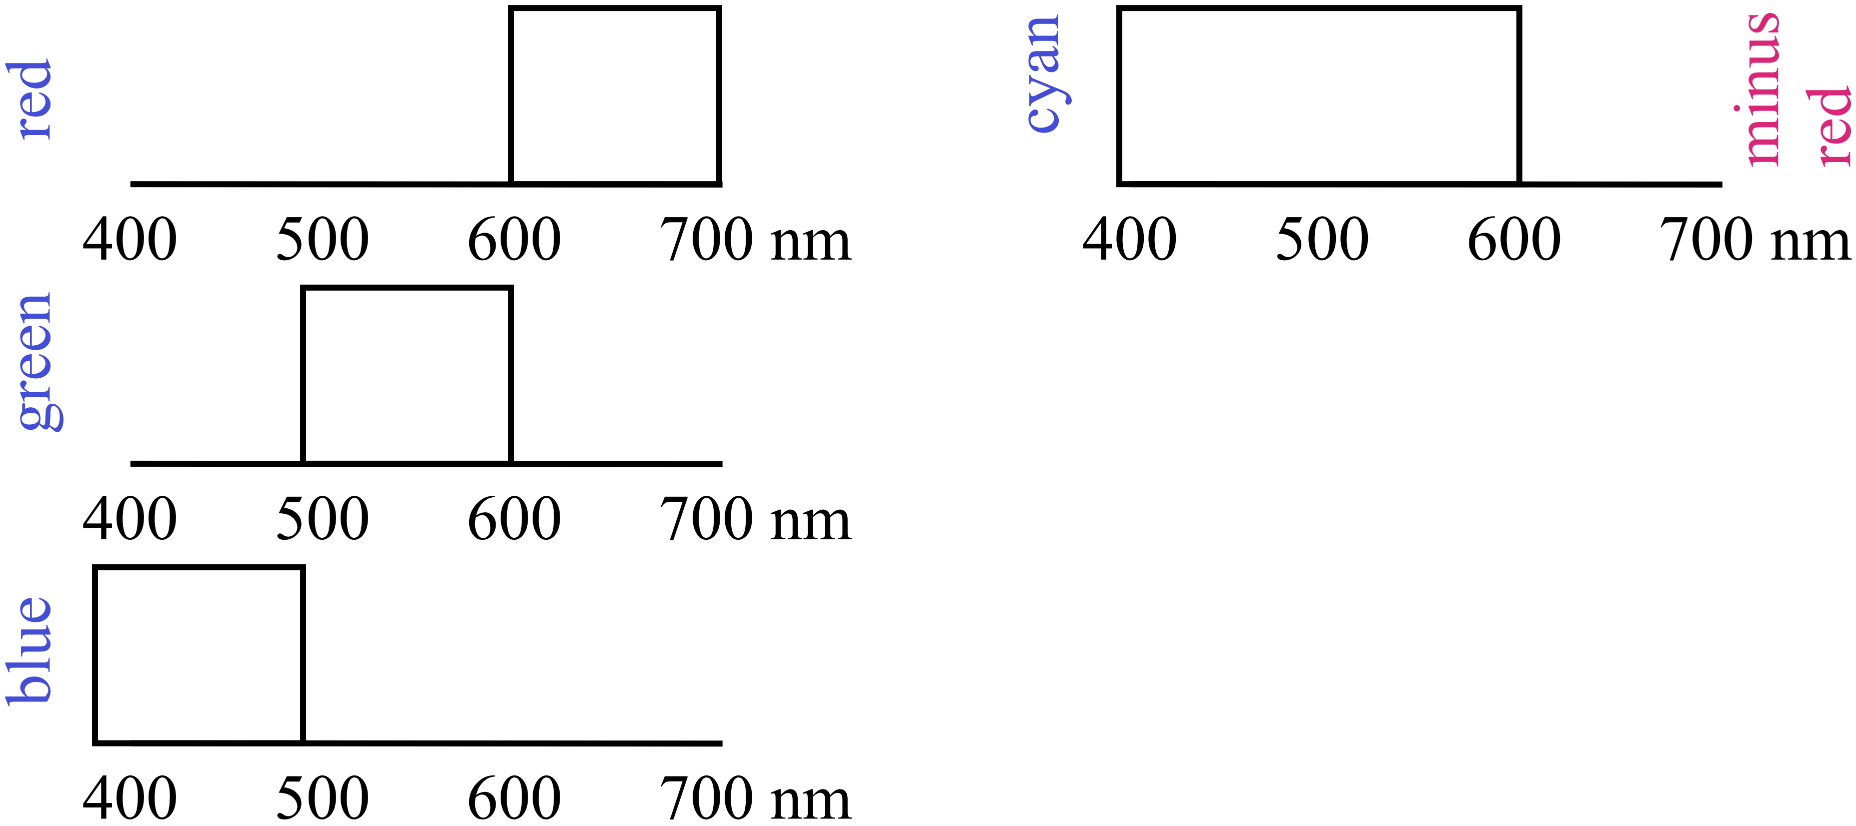

Color names for cartoon spectra

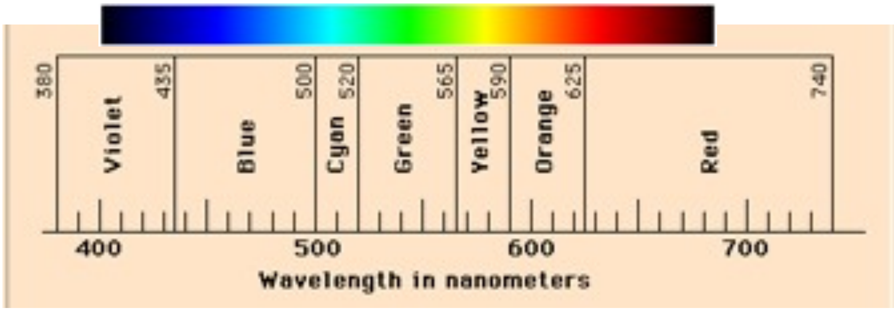

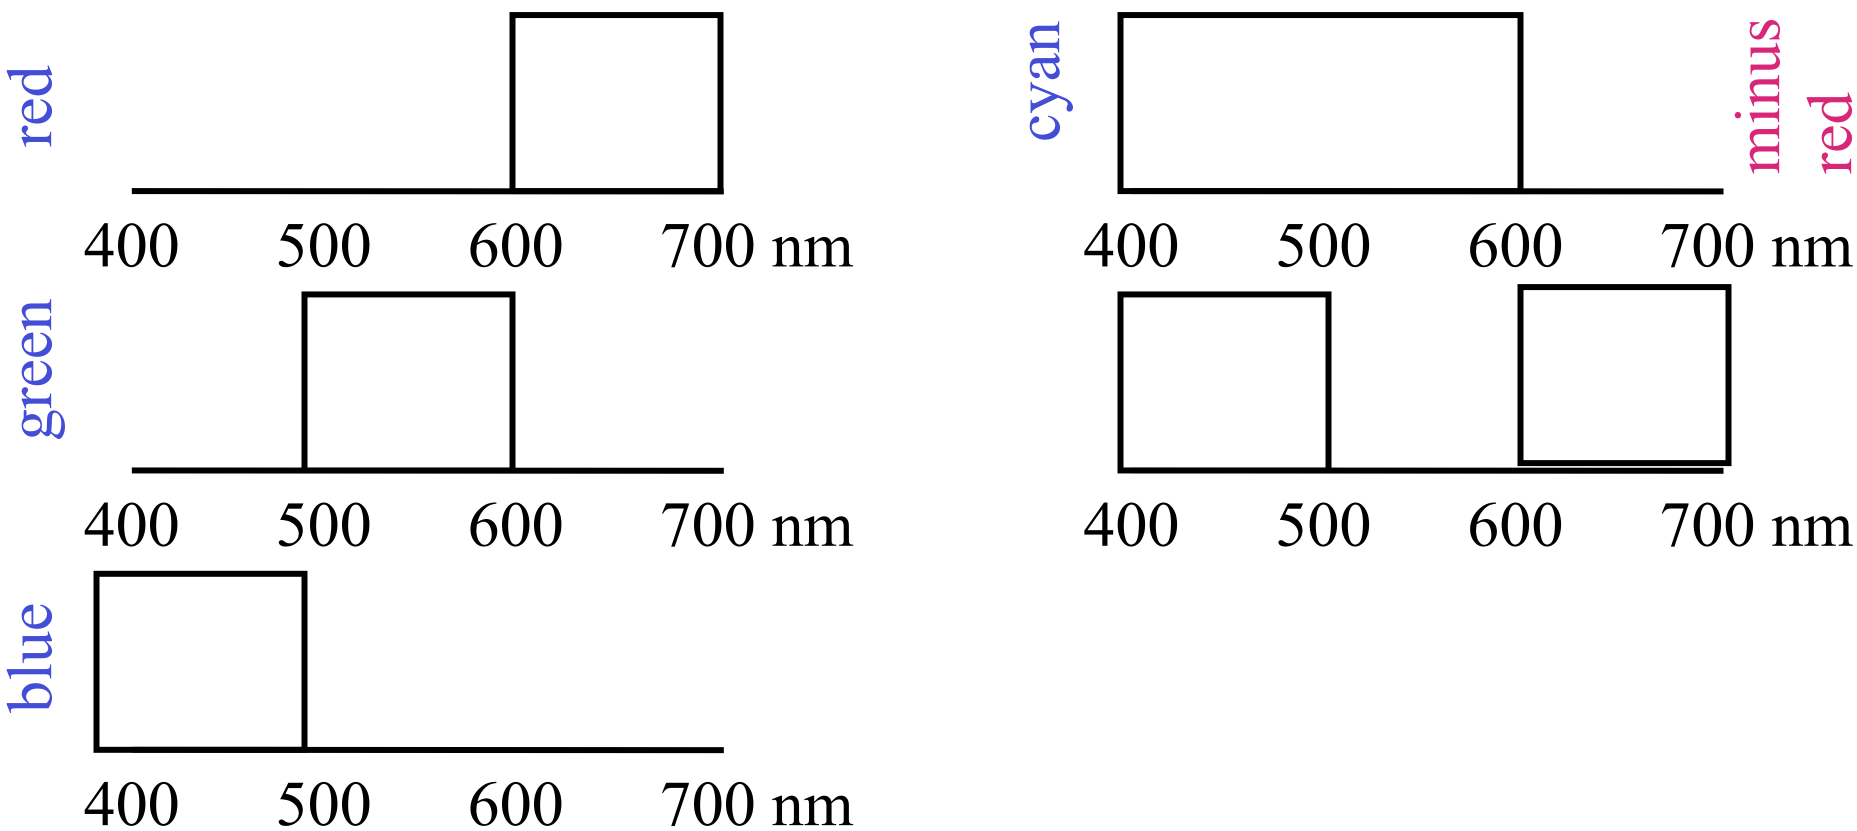

Color names for cartoon spectra

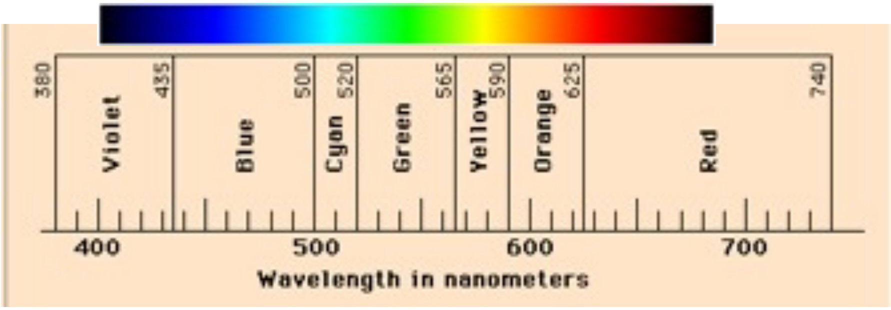

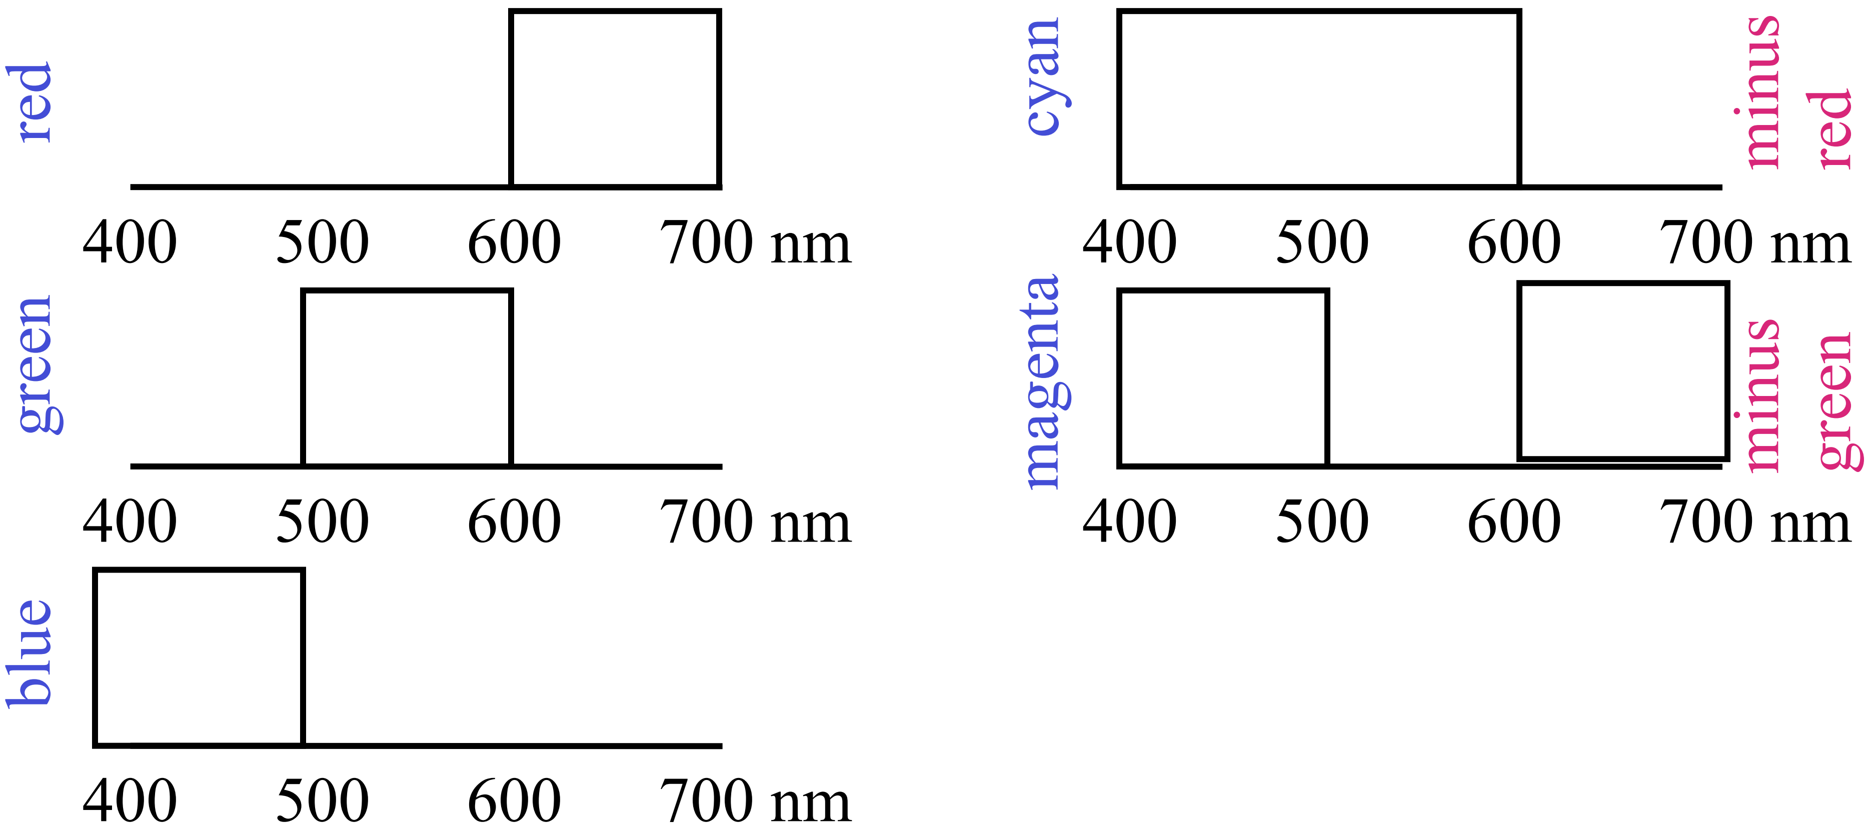

Color names for cartoon spectra

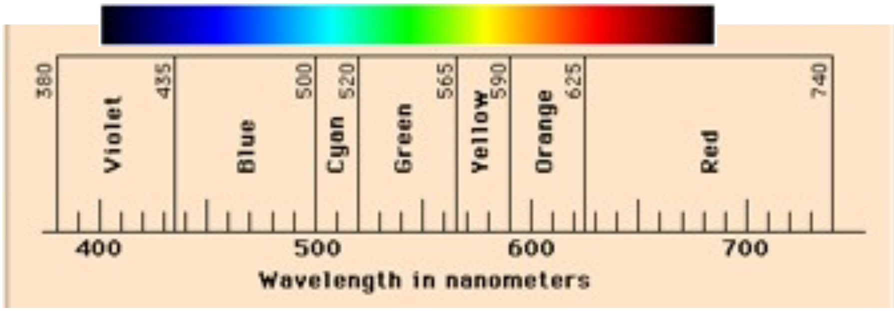

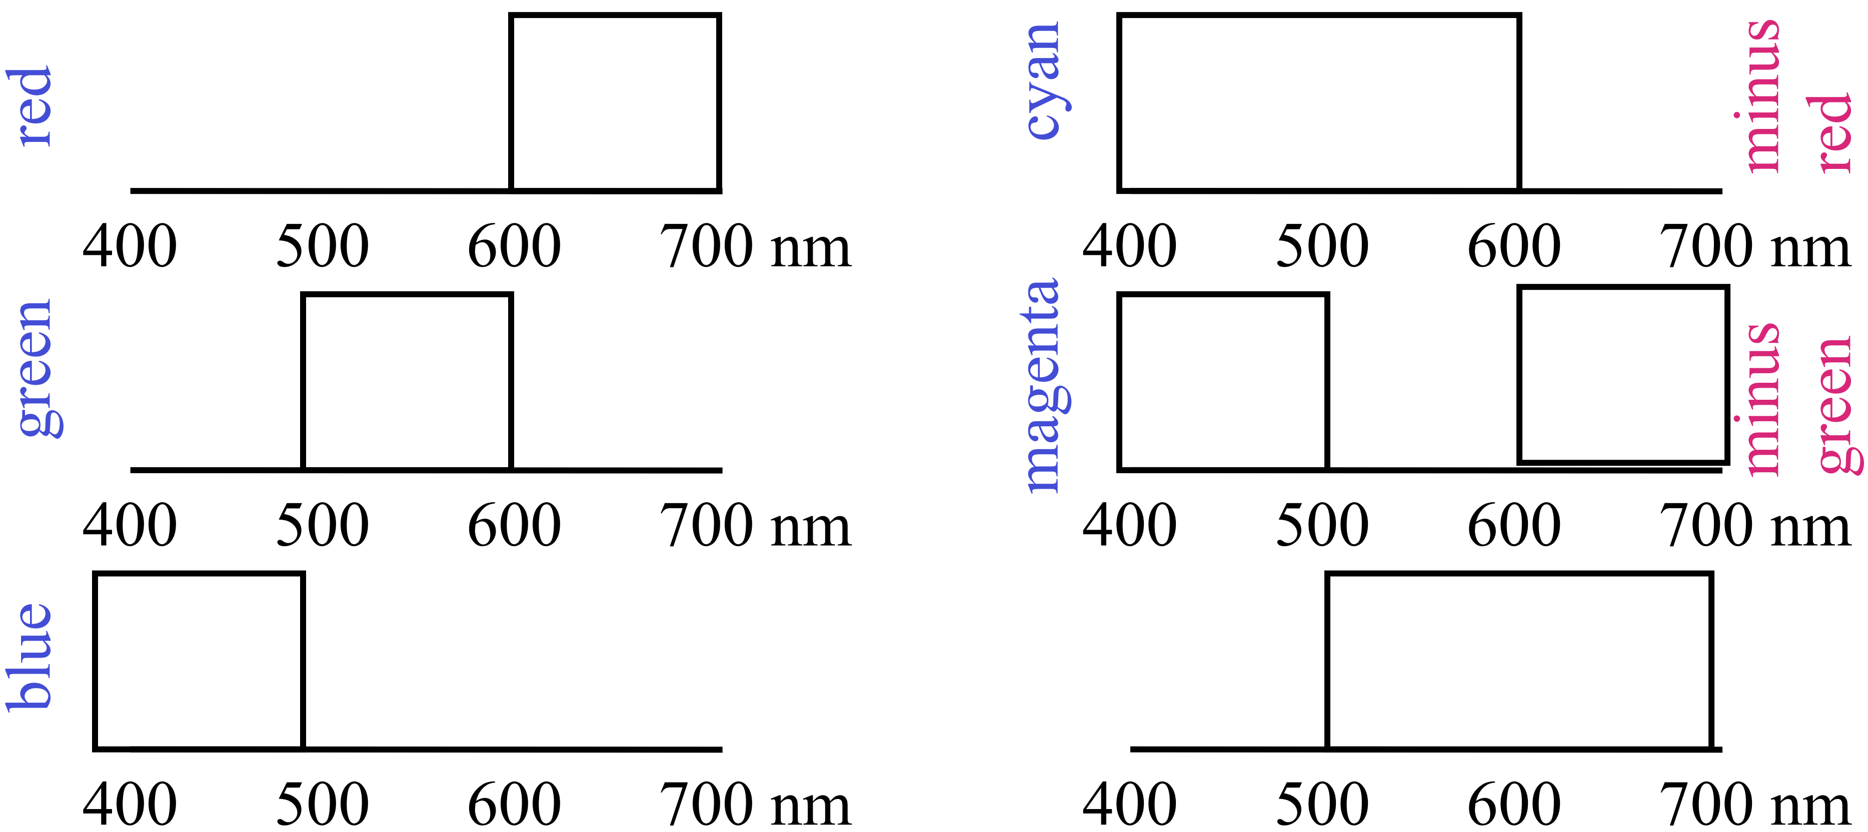

Color names for cartoon spectra

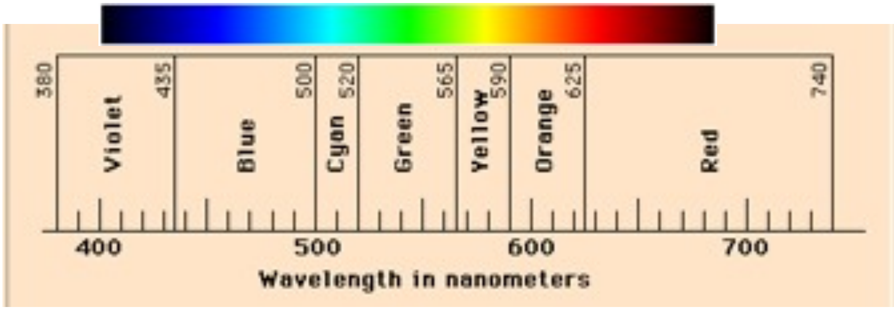

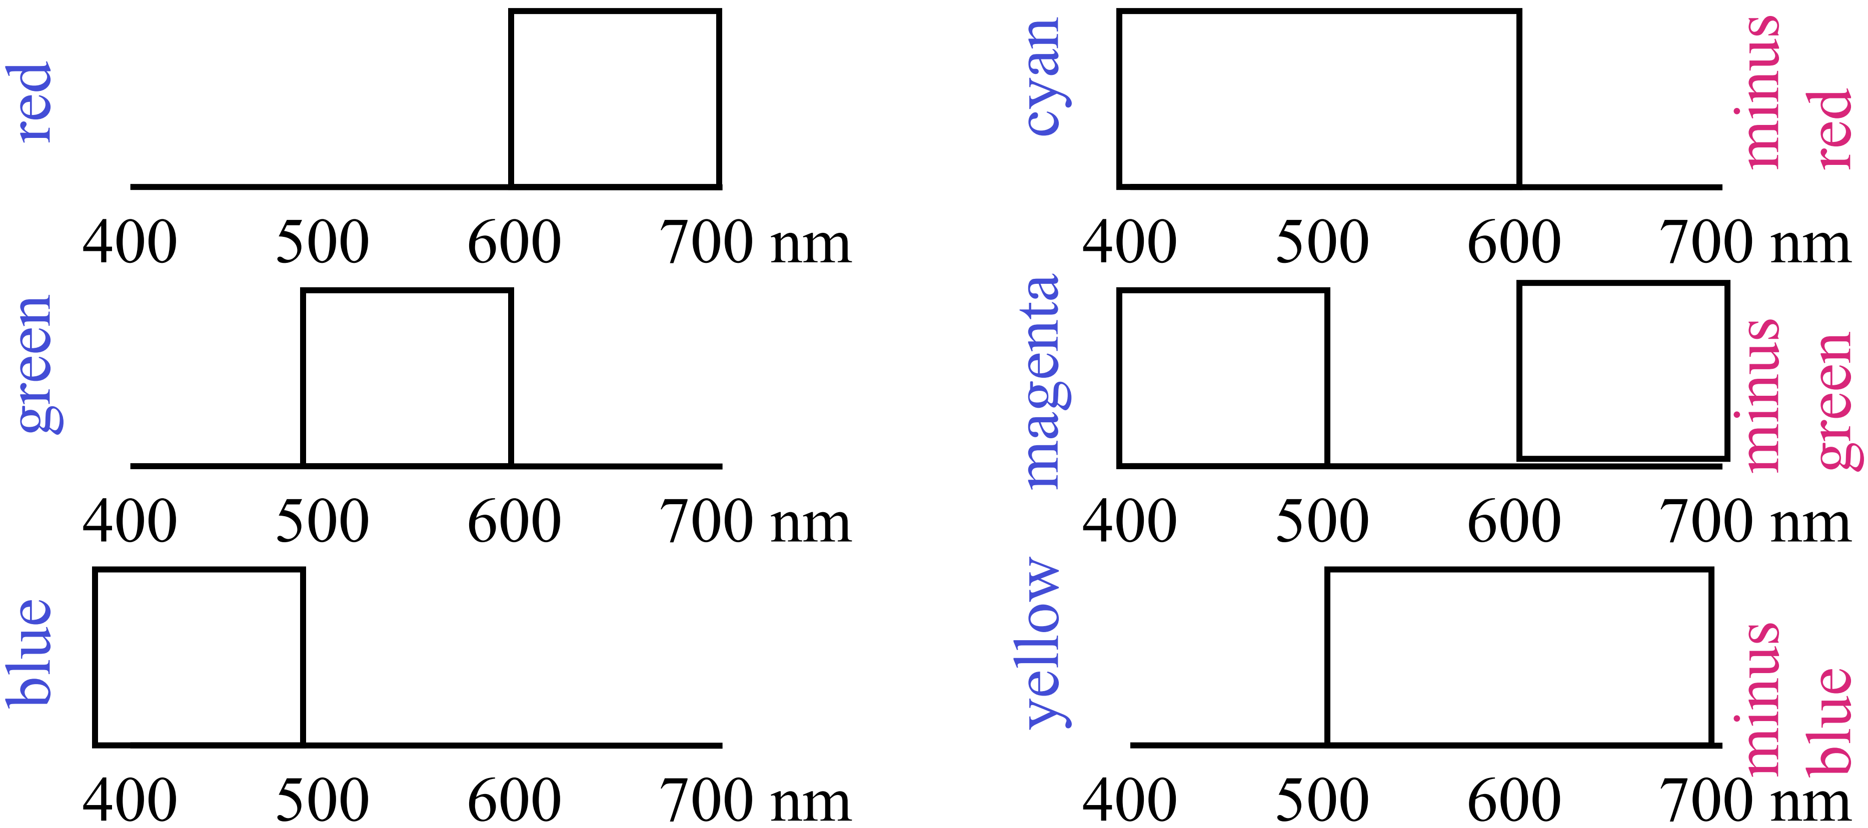

Additive color mixing

Additive color mixing

When colors combine by *adding* the color spectra. Example color displays that follow this mixing rule: CRT phosphors, multiple projectors aimed at a screen, Polachrome slide film.

Additive color mixing

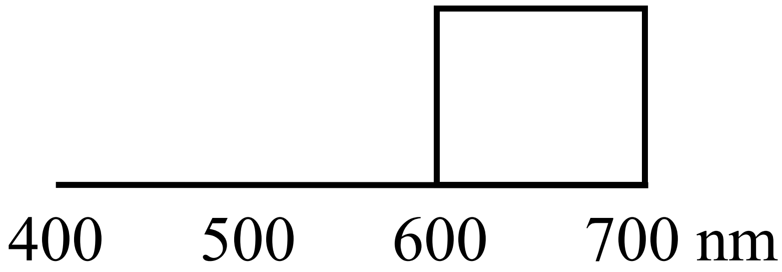

When colors combine by *adding* the color spectra. Example color displays that follow this mixing rule: CRT phosphors, multiple projectors aimed at a screen, Polachrome slide film.

Additive color mixing

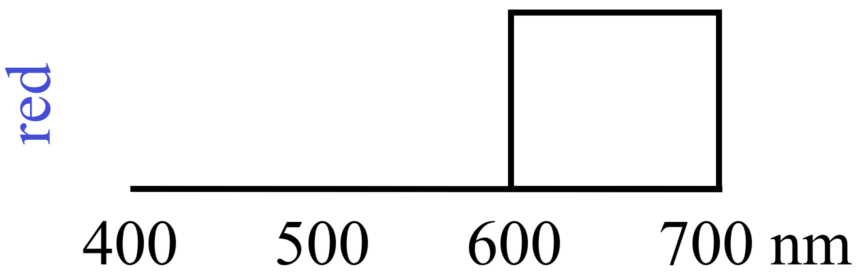

When colors combine by *adding* the color spectra. Example color displays that follow this mixing rule: CRT phosphors, multiple projectors aimed at a screen, Polachrome slide film.

Additive color mixing

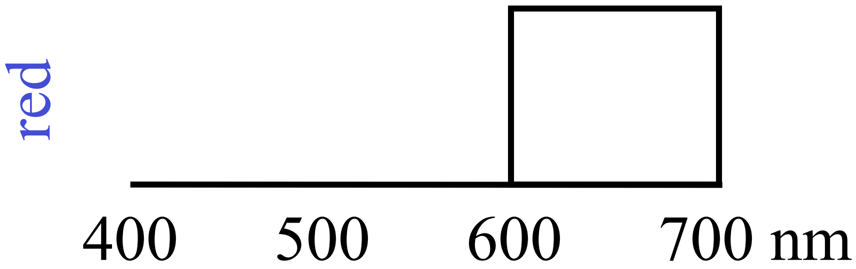

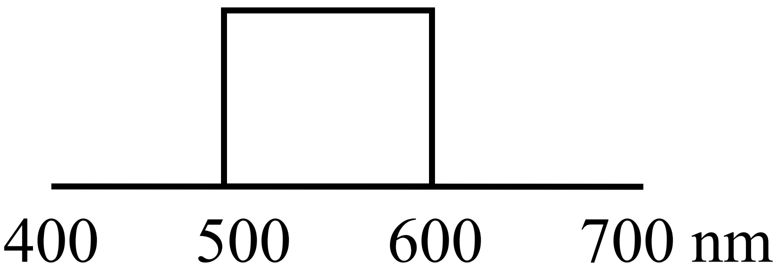

When colors combine by *adding* the color spectra. Example color displays that follow this mixing rule: CRT phosphors, multiple projectors aimed at a screen, Polachrome slide film.

Additive color mixing

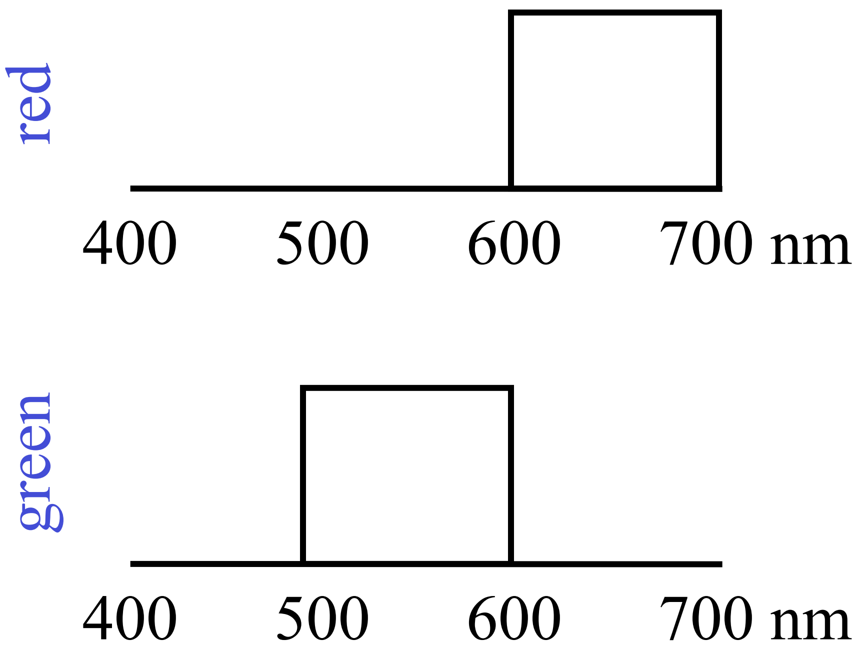

When colors combine by *adding* the color spectra. Example color displays that follow this mixing rule: CRT phosphors, multiple projectors aimed at a screen, Polachrome slide film.

Additive color mixing

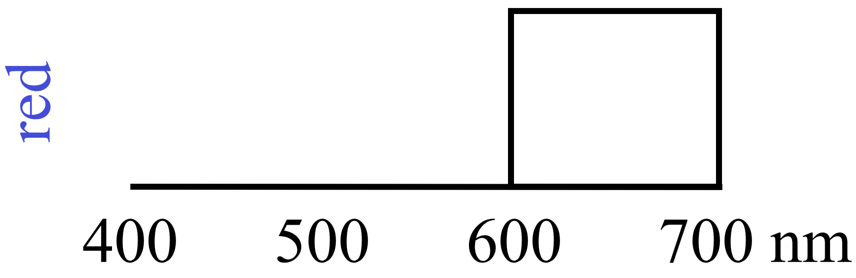

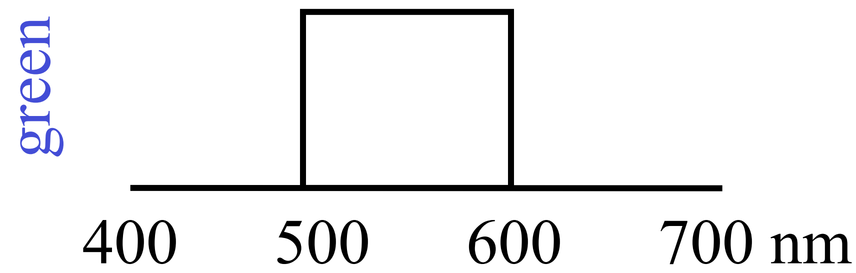

When colors combine by *adding* the color spectra. Example color displays that follow this mixing rule: CRT phosphors, multiple projectors aimed at a screen, Polachrome slide film.

Red and green make…

Additive color mixing

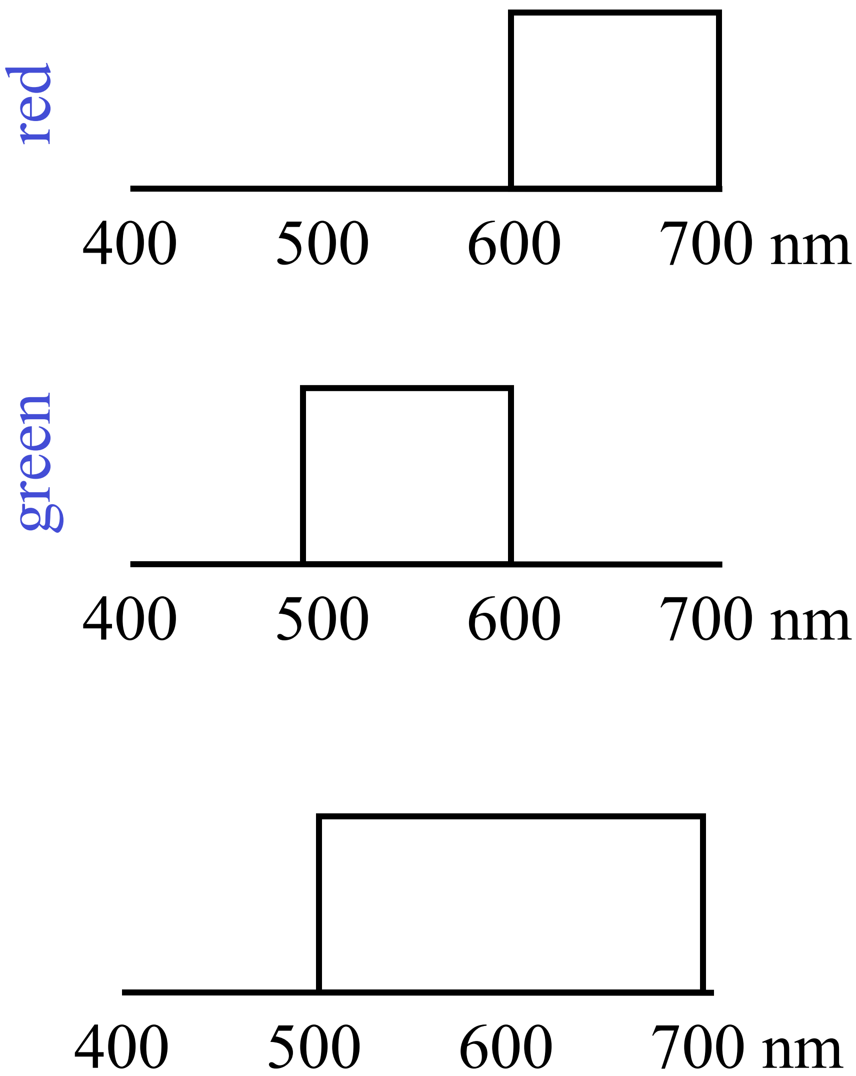

When colors combine by *adding* the color spectra. Example color displays that follow this mixing rule: CRT phosphors, multiple projectors aimed at a screen, Polachrome slide film.

Red and green make…

Additive color mixing

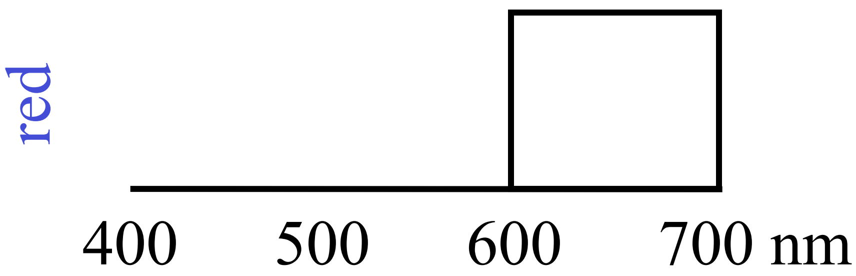

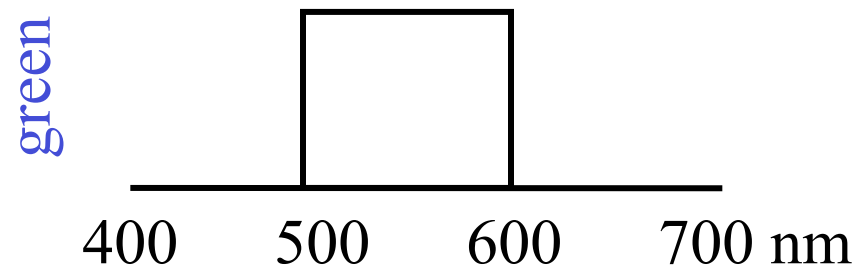

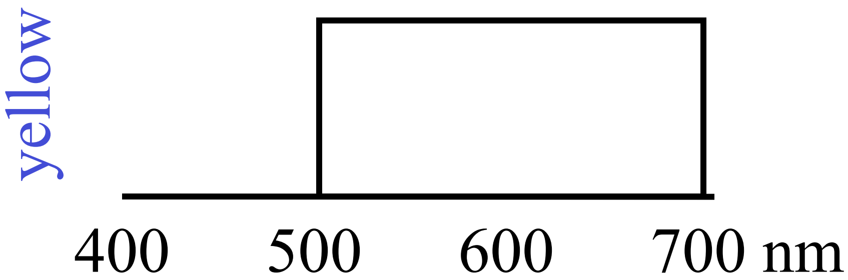

When colors combine by *adding* the color spectra. Example color displays that follow this mixing rule: CRT phosphors, multiple projectors aimed at a screen, Polachrome slide film.

Red and green make…

Yellow!

Subtractive color mixing

Subtractive color mixing

When colors combine by *multiplying* the color spectra. Examples that follow this mixing rule: most photographic films, paint, cascaded optical filters, crayons.

Subtractive color mixing

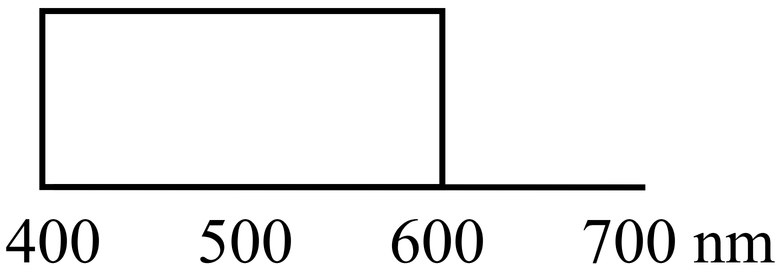

When colors combine by *multiplying* the color spectra. Examples that follow this mixing rule: most photographic films, paint, cascaded optical filters, crayons.

Subtractive color mixing

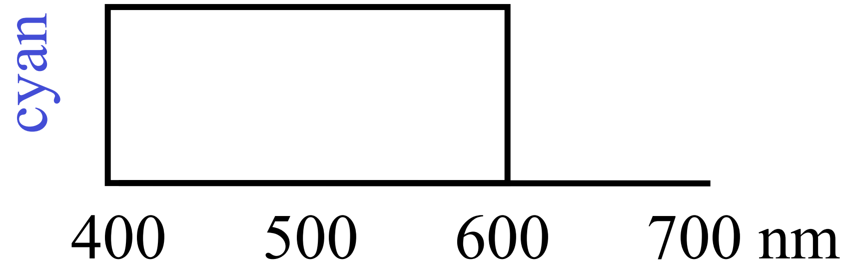

When colors combine by *multiplying* the color spectra. Examples that follow this mixing rule: most photographic films, paint, cascaded optical filters, crayons.

Subtractive color mixing

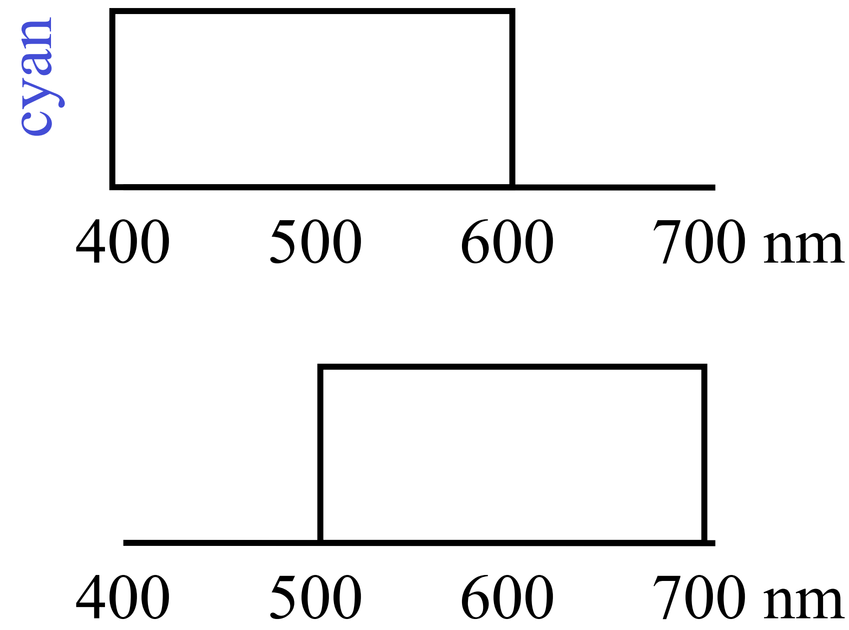

When colors combine by *multiplying* the color spectra. Examples that follow this mixing rule: most photographic films, paint, cascaded optical filters, crayons.

Subtractive color mixing

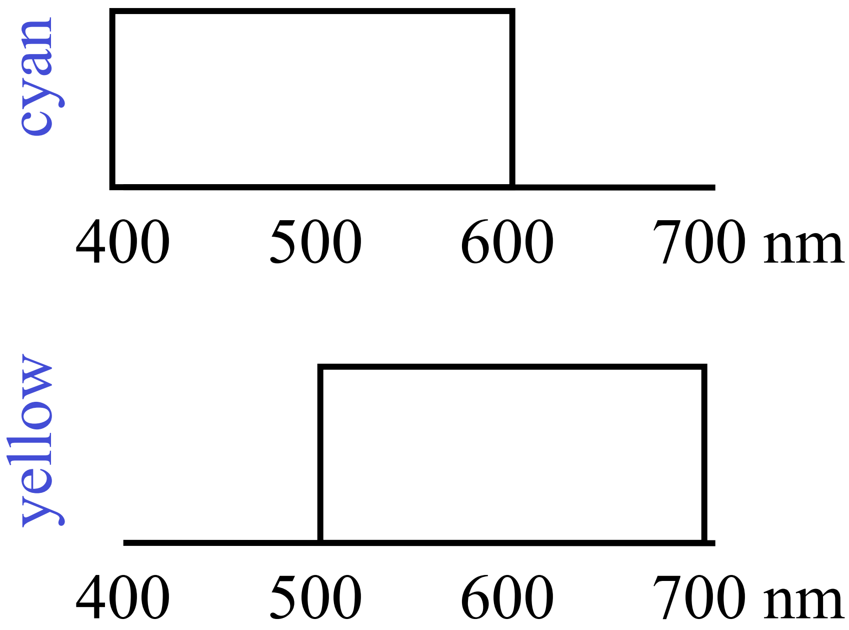

When colors combine by *multiplying* the color spectra. Examples that follow this mixing rule: most photographic films, paint, cascaded optical filters, crayons.

Subtractive color mixing

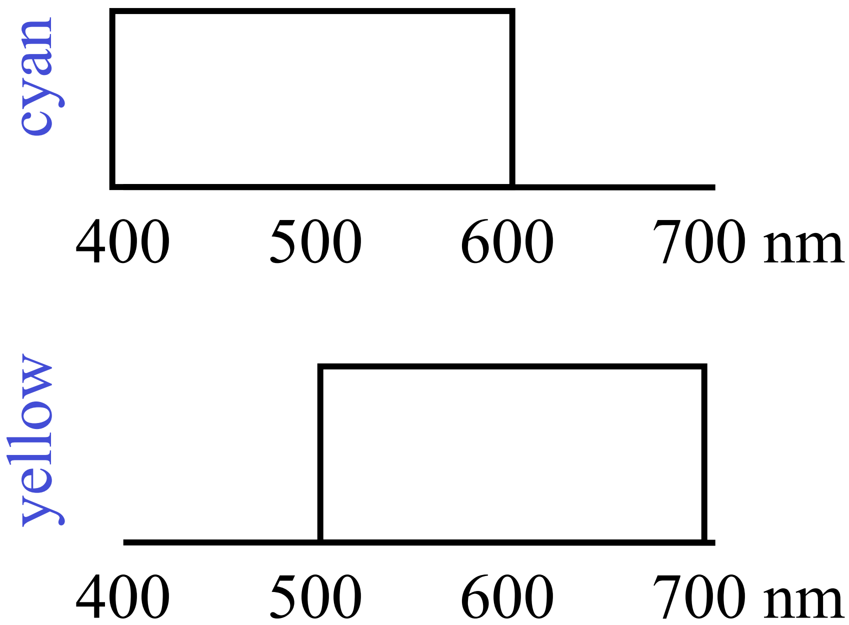

When colors combine by *multiplying* the color spectra. Examples that follow this mixing rule: most photographic films, paint, cascaded optical filters, crayons.

Cyan and yellow (in crayons, called “blue” and yellow) make…

Subtractive color mixing

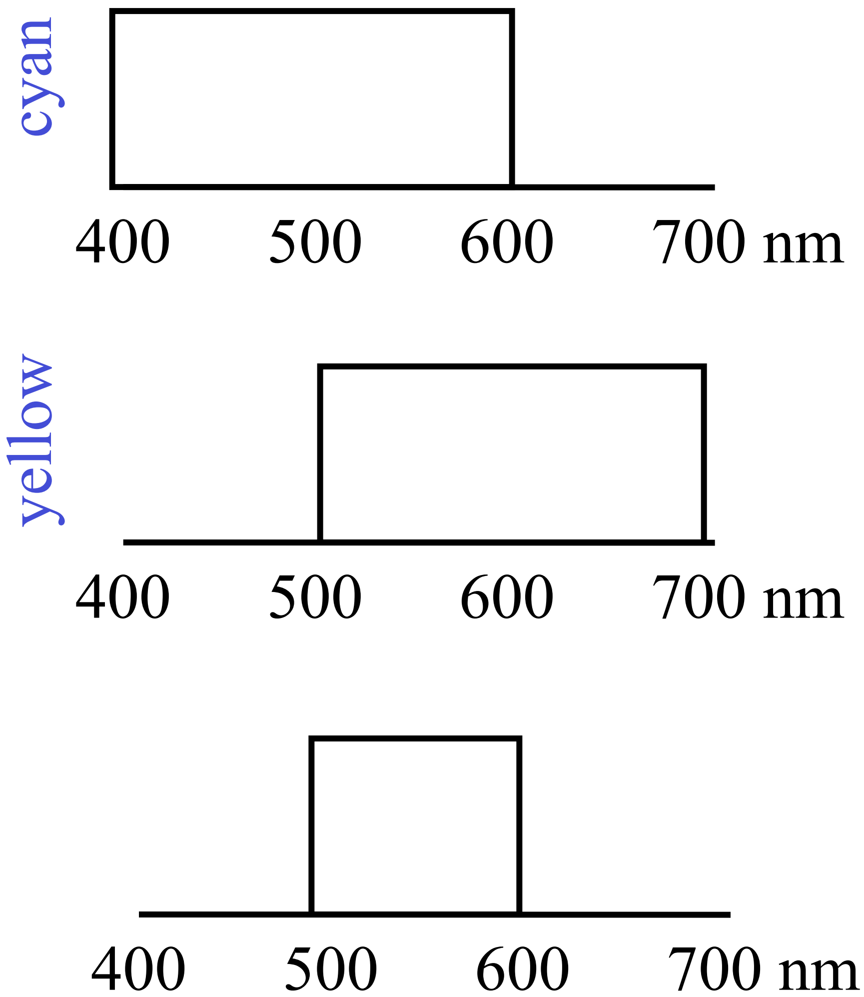

When colors combine by *multiplying* the color spectra. Examples that follow this mixing rule: most photographic films, paint, cascaded optical filters, crayons.

Cyan and yellow (in crayons, called “blue” and yellow) make…

Subtractive color mixing

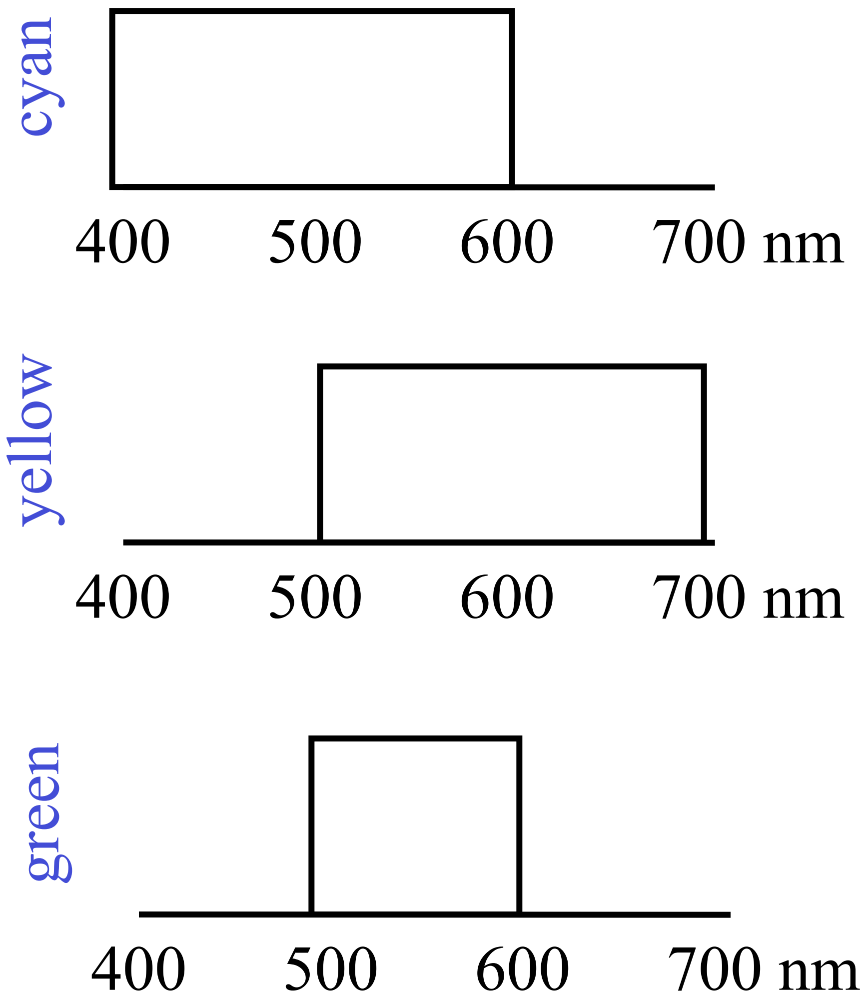

When colors combine by *multiplying* the color spectra. Examples that follow this mixing rule: most photographic films, paint, cascaded optical filters, crayons.

Cyan and yellow (in crayons, called "blue" and yellow) make…

Green!

Overhead projector demo

- Subtractive color mixing

Low-dimensional models for color spectra

$$\begin{pmatrix} \vdots \\ a(\lambda) \\ \vdots \end{pmatrix} \approx \begin{pmatrix} \vdots & \vdots & \vdots \\ a_1(\lambda) & a_2(\lambda) & a_3(\lambda) \\ \vdots & \vdots & \vdots \end{pmatrix} \begin{pmatrix} \omega_1 \\ \omega_2 \\ \omega_3 \end{pmatrix}$$

How to find a linear model for color spectra:

--form a matrix, D, of measured spectra, 1 spectrum per column.

--[u, s, v] = svd(D) satisfies D = u*s*v‘

--the first n columns of u give the best (least-squares optimal) n-dimensional linear bases for the data, D:

$$D \approx u(:,1:n) * s(1:n,1:n) * v(1:n,:)'$$

Macbeth Color Checker

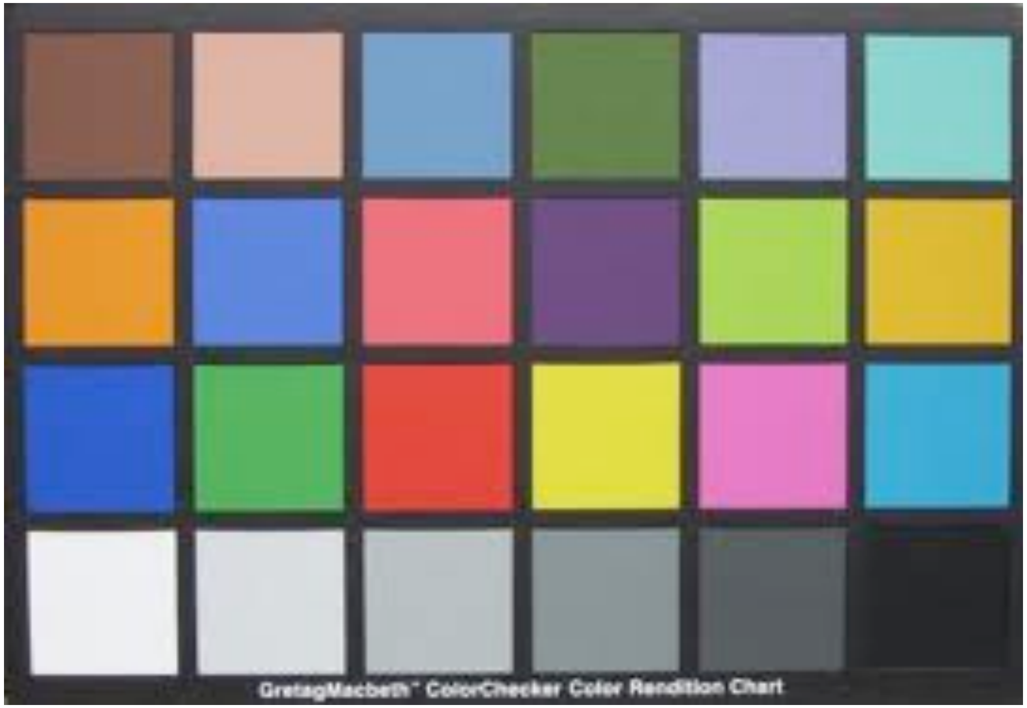

http://www.flickr.com/photos/erikaneola/2231647456/

My Macbeth Color Checker Tattoo

I think I have all the other color checker photos beat...

Yes, the tattoo is real.
No, it is not a rubik's cube.

THIS PHOTOGRAPH IS COPYRIGHT 2007 THE X-RITE CORPORATION!

A photograph from this session can be viewed on the X-Rite Website: www.xrite.com/top_munsell.aspx

Basis functions for Macbeth color checker

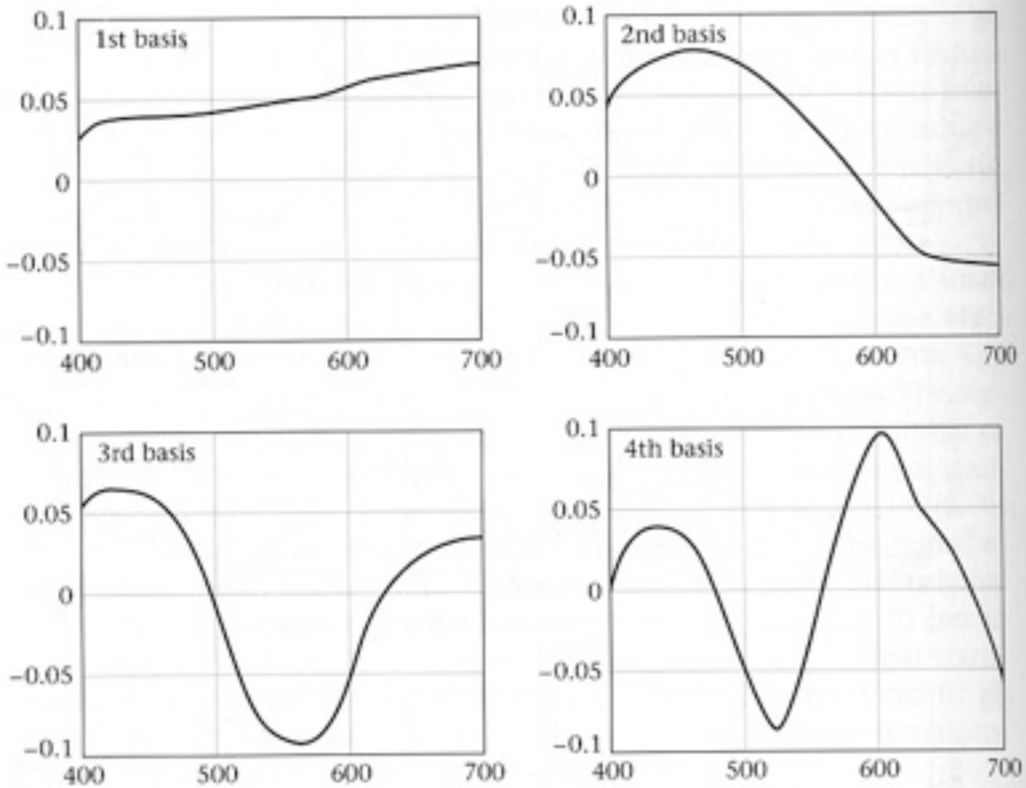

9.9 BASIS FUNCTIONS OF THE LINEAR MODEL FOR THE MACBETH COLORCHECKER. The surface-reflectance functions in the collection vary smoothly with wavelength, as do the basis functions. The first basis function is all positive and explains the most variance in the surface-reflectance functions. The basis functions are ordered in terms of their relative significance for reducing the error in the linear-model approximation to the surfaces.

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

Fitting color spectra with low-dimensional linear models

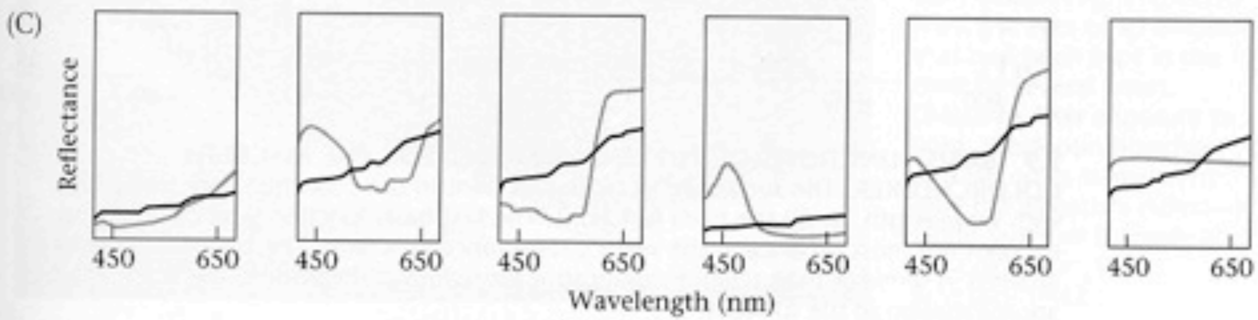

n = 1

9.8 A LINEAR MODEL TO APPROXIMATE THE SURFACE REFLECTANCES IN THE MACBETH COLORCHECKER. The panels in each row of this figure show the surface-reflectance functions of six colored surfaces (shaded lines) and the approximation to these functions using a linear model (solid lines). The approximations using linear models with (A) three, (B) two, and (C) one dimension are shown.

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

Fitting color spectra with low-dimensional linear models

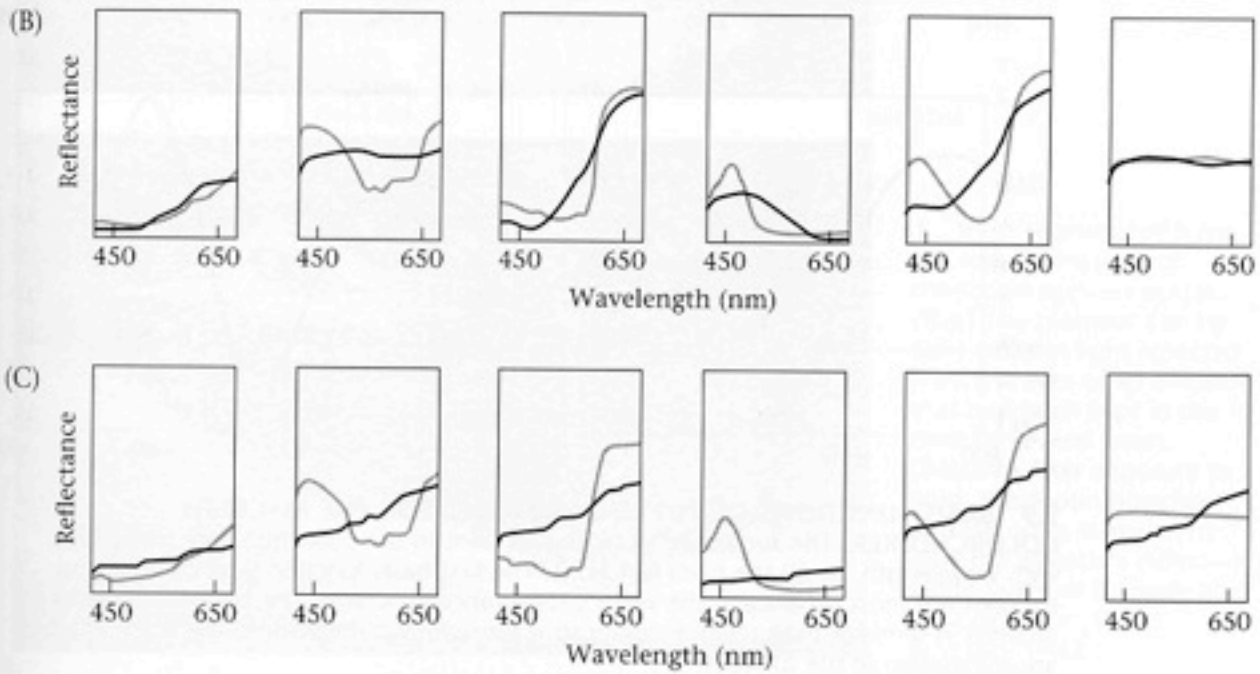

n = 2

n = 1

9.8 A LINEAR MODEL TO APPROXIMATE THE SURFACE REFLECTANCES IN THE MACBETH COLORCHECKER. The panels in each row of this figure show the surface-reflectance functions of six colored surfaces (shaded lines) and the approximation to these functions using a linear model (solid lines). The approximations using linear models with (A) three, (B) two, and (C) one dimension are shown.

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

Fitting color spectra with low-dimensional linear models

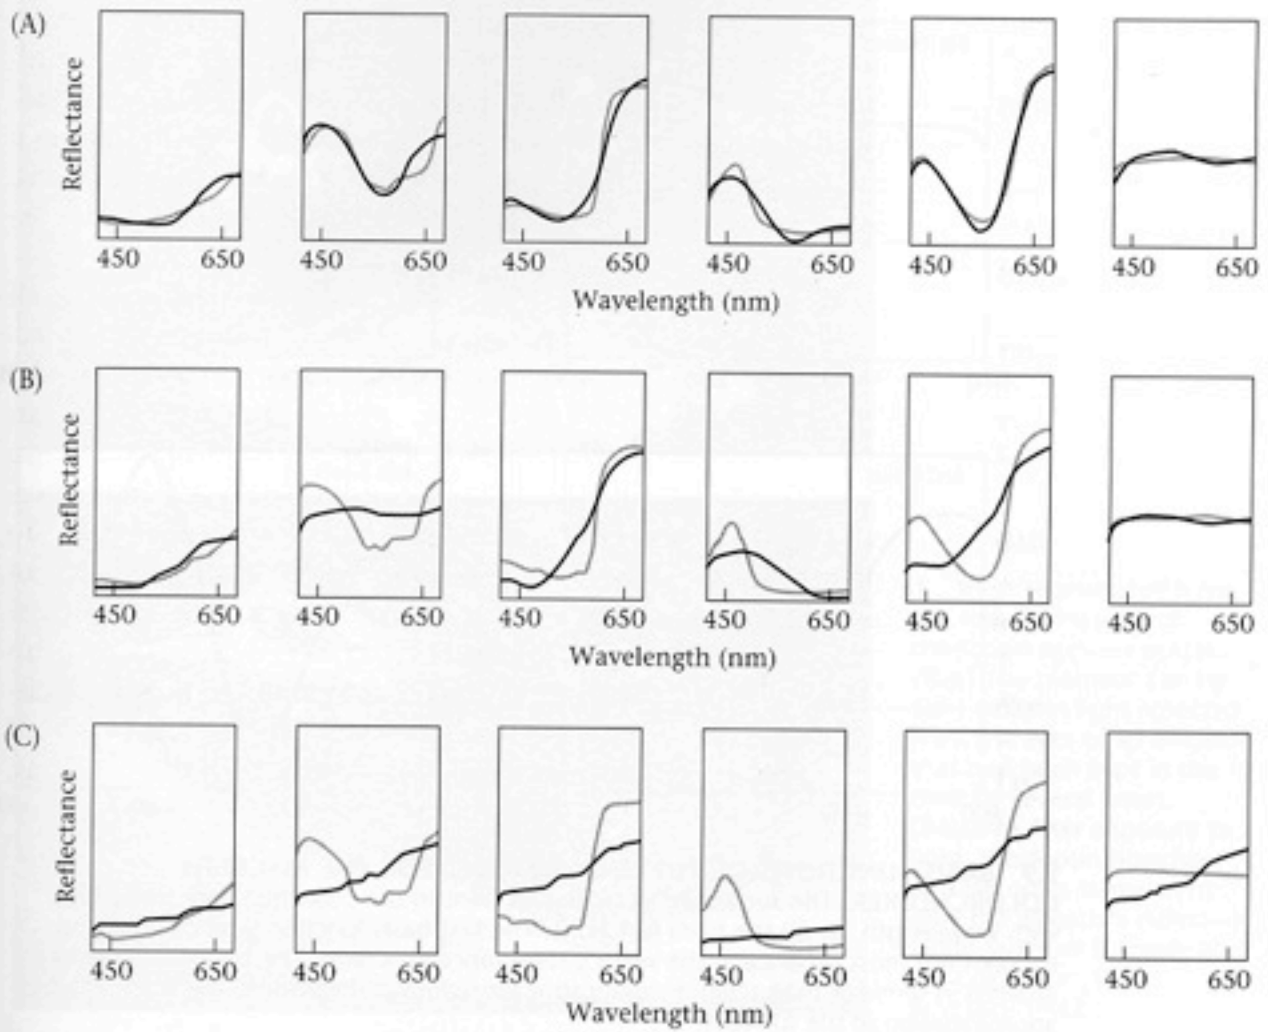

n = 3

n = 2

n = 1

9.8 A LINEAR MODEL TO APPROXIMATE THE SURFACE REFLECTANCES IN THE MACBETH COLORCHECKER. The panels in each row of this figure show the surface-reflectance functions of six colored surfaces (shaded lines) and the approximation to these functions using a linear model (solid lines). The approximations using linear models with (A) three, (B) two, and (C) one dimension are shown.

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

Lecture outline

- Color physics.
- Color perception.

Lecture outline

- Color physics.
- Color perception
  - part 1: assume perceived color only depends on light spectrum.
  - part 2: the more general case.

The assumption for color perception, part 1

The assumption for color perception, part 1

- We know color appearance really depends on:

The assumption for color perception, part 1

- We know color appearance really depends on:
  - The illumination

The assumption for color perception, part 1

- We know color appearance really depends on:
  - The illumination
  - Your eye’s adaptation level

The assumption for color perception, part 1

- We know color appearance really depends on:
  - The illumination
  - Your eye's adaptation level
  - The colors and scene interpretation surrounding the observed color.

The assumption for color perception, part 1

- We know color appearance really depends on:
  - The illumination
  - Your eye’s adaptation level
  - The colors and scene interpretation surrounding the observed color.

The assumption for color perception, part 1

- We know color appearance really depends on:
  - The illumination
  - Your eye's adaptation level
  - The colors and scene interpretation surrounding the observed color.

- But for now we will assume that <u>the spectrum of the light arriving at your eye completely determines the perceived color</u>.

Color standards are important in industry

Address http://www.ams.usda.gov/fv/ppbweb/PPBfilecodes/105a15.htm

AMS USDA

Fruit and Vegetable Programs

AMS USDA SEARCH

Processed Products Standards and Quality Certification

Visual Aids and Inspection Aids Approved For Use in Ascertaining Grades of Processed Fruits and Vegetables (Photo)

- Frozen Red Tart Cherries
- Orange Juice (Processed)
- Canned Tomatoes
- Frozen French Fried Potatoes
- Tomato Products
- Maple Syrup
- Honey
- Frozen Lima Beans
- Canned Mushrooms
- Peanut Butter
- Canned Pimientos
- Frozen Peas
- Canned Clingstone Peaches
- Headspace Gauge
- Canned Applesauce
- Canned Freestone Peaches
- Canned Ripe Olives

Return to: Processed Products Bran

New Page 1 - Microsoft Internet Explorer

Address http://www.ams.usda.gov/fv/ppbweb/PPBfilecodes/visual%20aid%20photo.htm

Image of Inspecti

UNITED STATES DEPARTMENT OF AGRICULTURE

COLOR STANDARDS

for

FROZEN FRENCH FRIED POTATOES

FOURTH EDITION, 1988

MUNSELL COLOR
BALTIMORE, MARYLAND
64-1

Color trademarks

CURRENTLY REGISTERED COLOR TRADEMARKS

http://blog.patents-tms.com/?p=52

A color trademark is a non-conventional trademark where at least one color is used to identify the commercial origin of a product or service. A color trademark must meet the same requirements of a conventional trademark. Thus, the color trademark must either be inherently distinctive or have acquired secondary meaning. To be inherently distinctive, the color must be arbitrarily or suggestively applied to a product or service. In contrast, to acquire secondary meaning, consumers must associate the color used on goods or services as originating from a single source. Below is a selection of some currently registered color trademarks in the U.S. Trademark Office:

MARK/COLOR(S)/OWNER:

BANK OF AMERICA 500
blue, red & grey
Bank of America Corporation

NATIONAL CAR RENTAL
green
NCR Affiliate Servicer, Inc.

FORD
blue
Ford Motor Company

VISTEON
orange
Ford Motor Company

76
red & blue
ConocoPhillips Company

VW
silver, metallic blue, black and white
Volkswagen Aktiengesellschaft Corp.

THE HOME DEPOT
orange
Homer TLC, Inc.

HONDA
red
Honda Motor Co., Ltd.

M MARATHON
brown, orange, yellow
Marathon Oil Company

M MARATHON
gray, black & white
Marathon Oil Company

COSTCO
red
Costco Wholesale Membership, Inc.

TEENAGE MUTANT NINJA TURTLES MUTANTS & MONSTERS
red, green, yellow, black, grey and white
Mirage Studios, Inc.

TARGET

What we need from a color measurement system

- Given a color, how do you assign a number to it?
- Given an input power spectrum, what is its numerical color value, and how do we control our printing/projection/cooking system to match it?

What’s the machinery in the eye?

Eye Photoreceptor responses

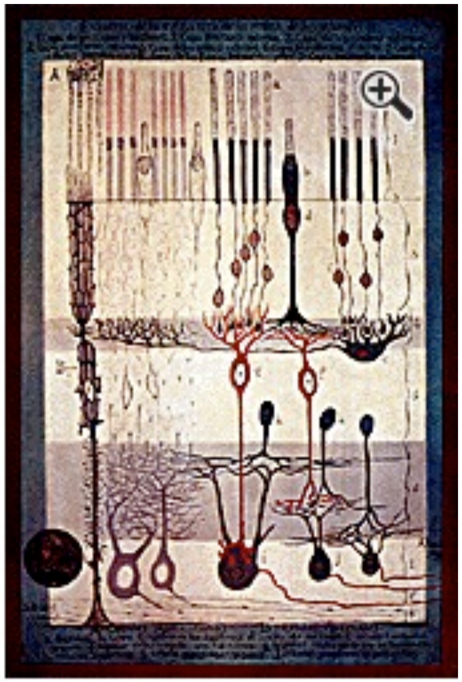

Instituto Cajal. CSIC. Madrid.

The intricate layers and connections of nerve cells in the retina were drawn by the famed Spanish anatomist Santiago Ramón y Cajal around 1900. Rod and cone cells are at the top. Optic nerve fibers leading to the brain may be seen at bottom right.

(Where do you think the light comes in?)

Human Photoreceptors

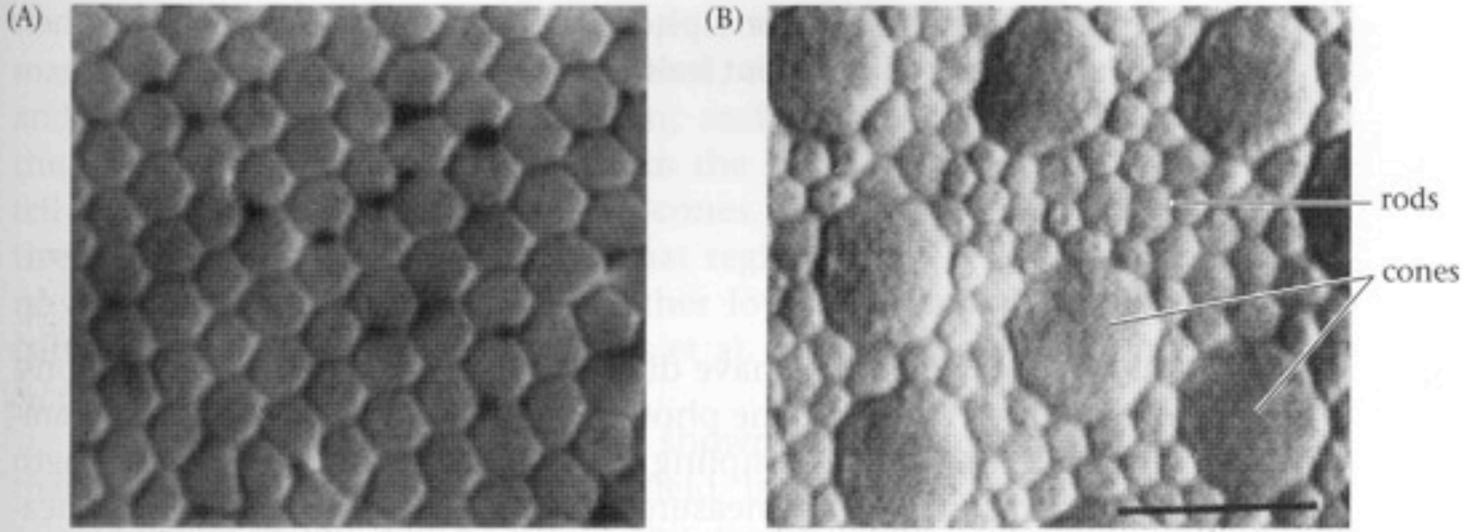

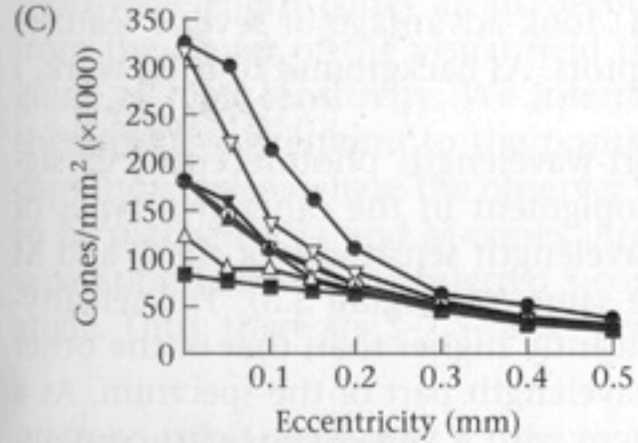

3.4 THE SPATIAL MOSAIC OF THE HUMAN CONES. Cross sections of the human retina at the level of the inner segments showing (A) cones in the fovea, and (B) cones in the periphery. Note the size difference (scale bar = 10 μm), and that, as the separation between cones grows, the rod receptors fill in the spaces. (C) Cone density plotted as a function of distance from the center of the fovea for seven human retinas; cone density decreases with distance from the fovea. Source: Curcio et al., 1990.

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

Human eye photoreceptor spectral sensitivities

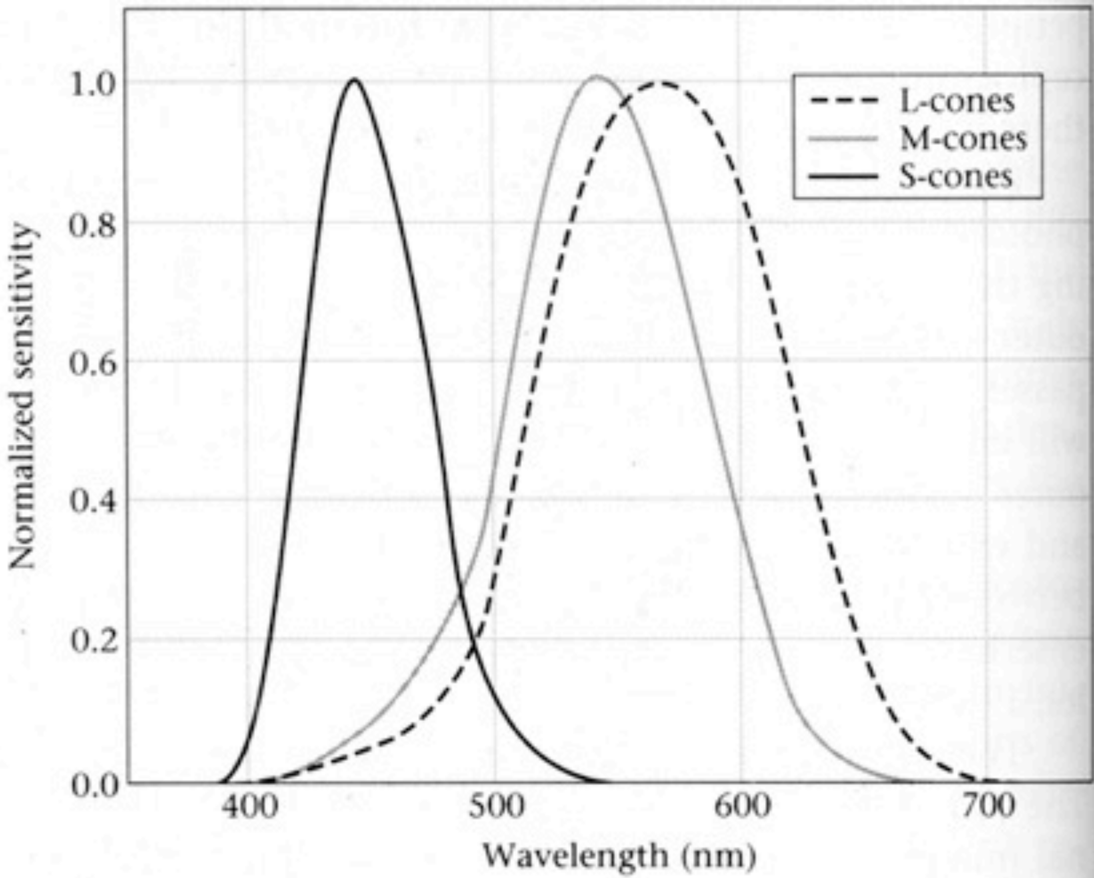

3.3 SPECTRAL SENSITIVITIES OF THE L-, M-, AND S-CONES in the human eye. The measurements are based on a light source at the cornea, so that the wavelength loss due to the cornea, lens, and other inert pigments of the eye plays a role in determining the sensitivity. Source: Stockman and MacLeod, 1993.

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

Color matching experiment

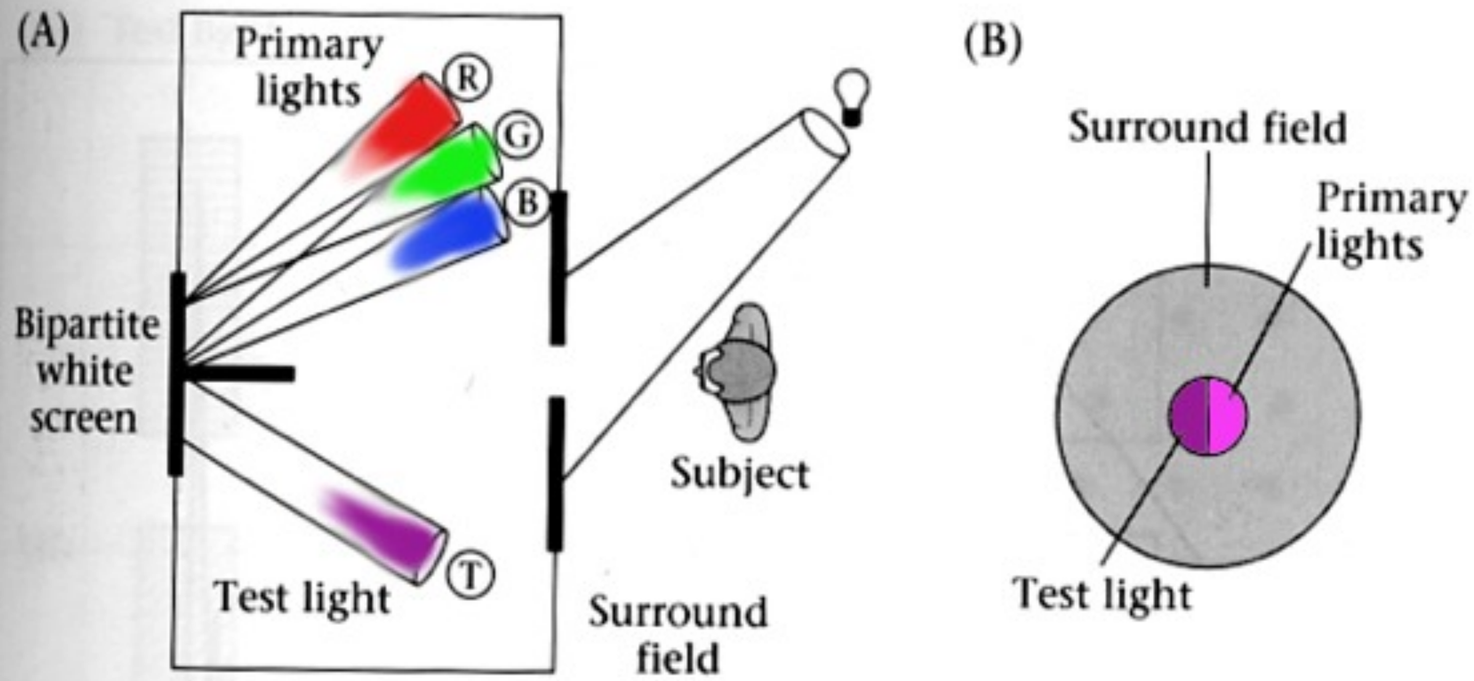

4.10 THE COLOR-MATCHING EXPERIMENT. The observer views a bipartite field and adjusts the intensities of the three primary lights to match the appearance of the test light. (A) A top view of the experimental apparatus. (B) The appearance of the stimuli to the observer. After Judd and Wyszecki, 1975.

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

Color matching experiment 1

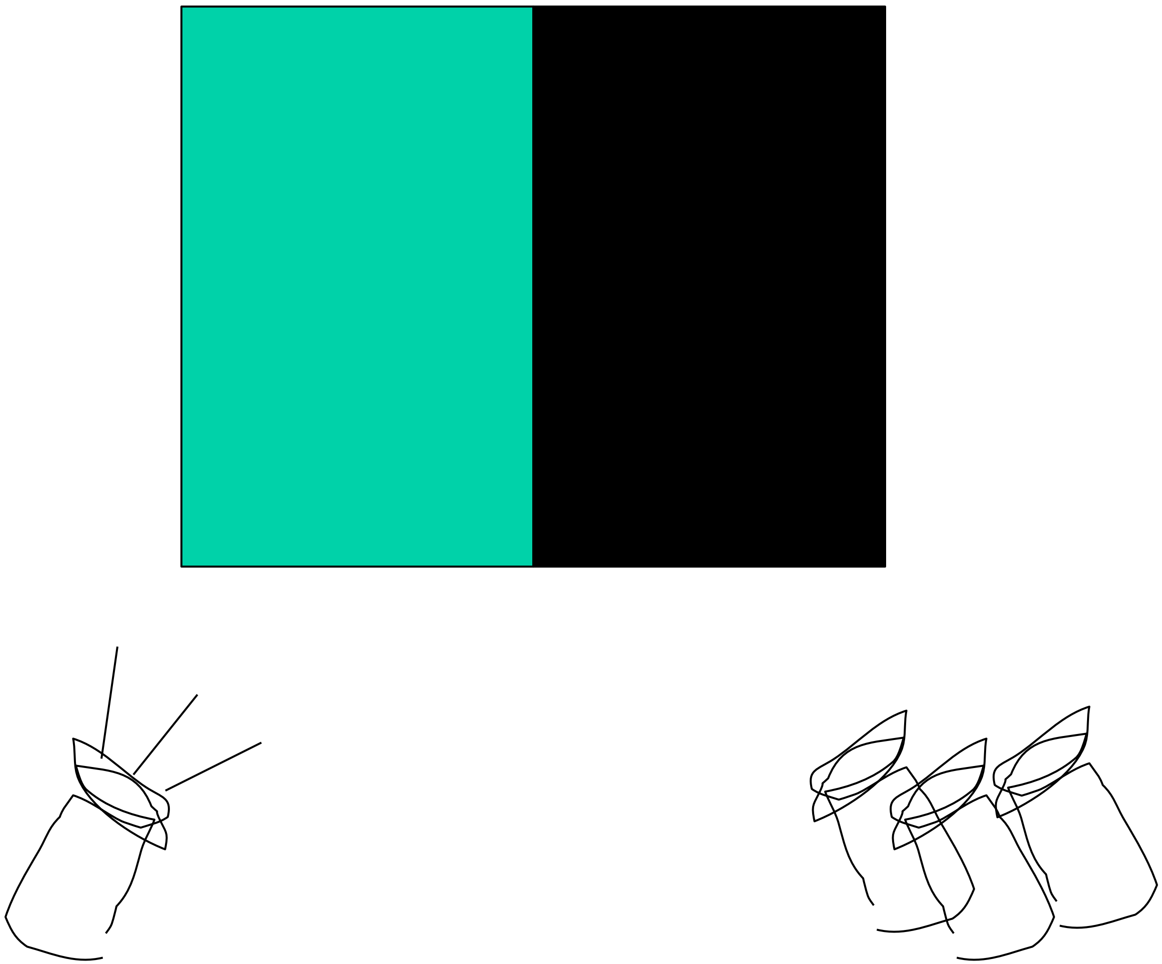

Color matching experiment 1

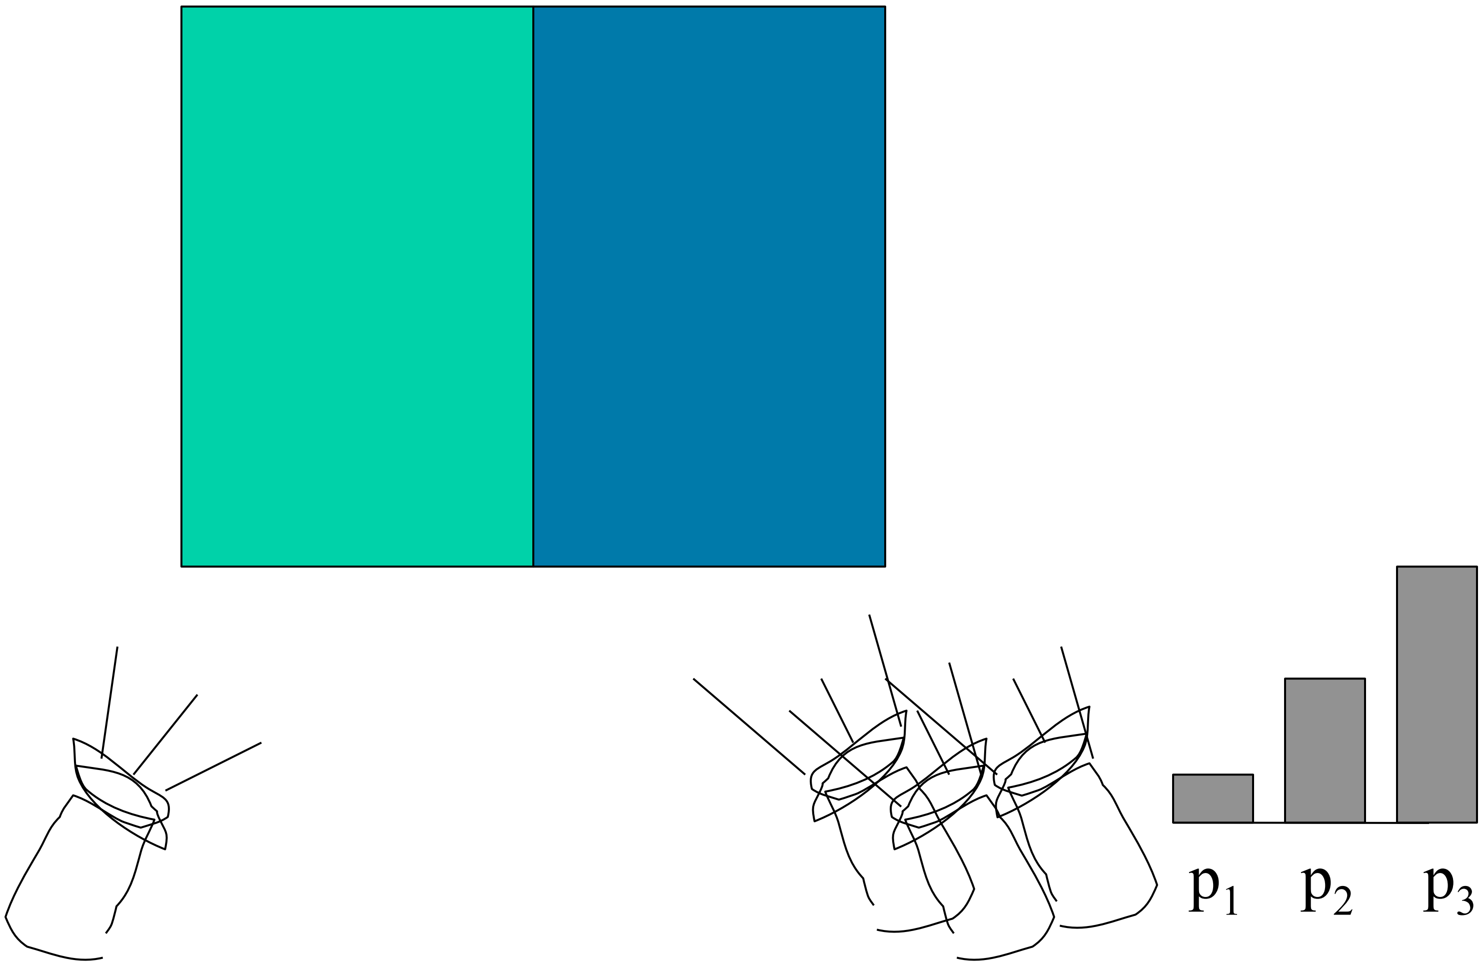

Color matching experiment 1

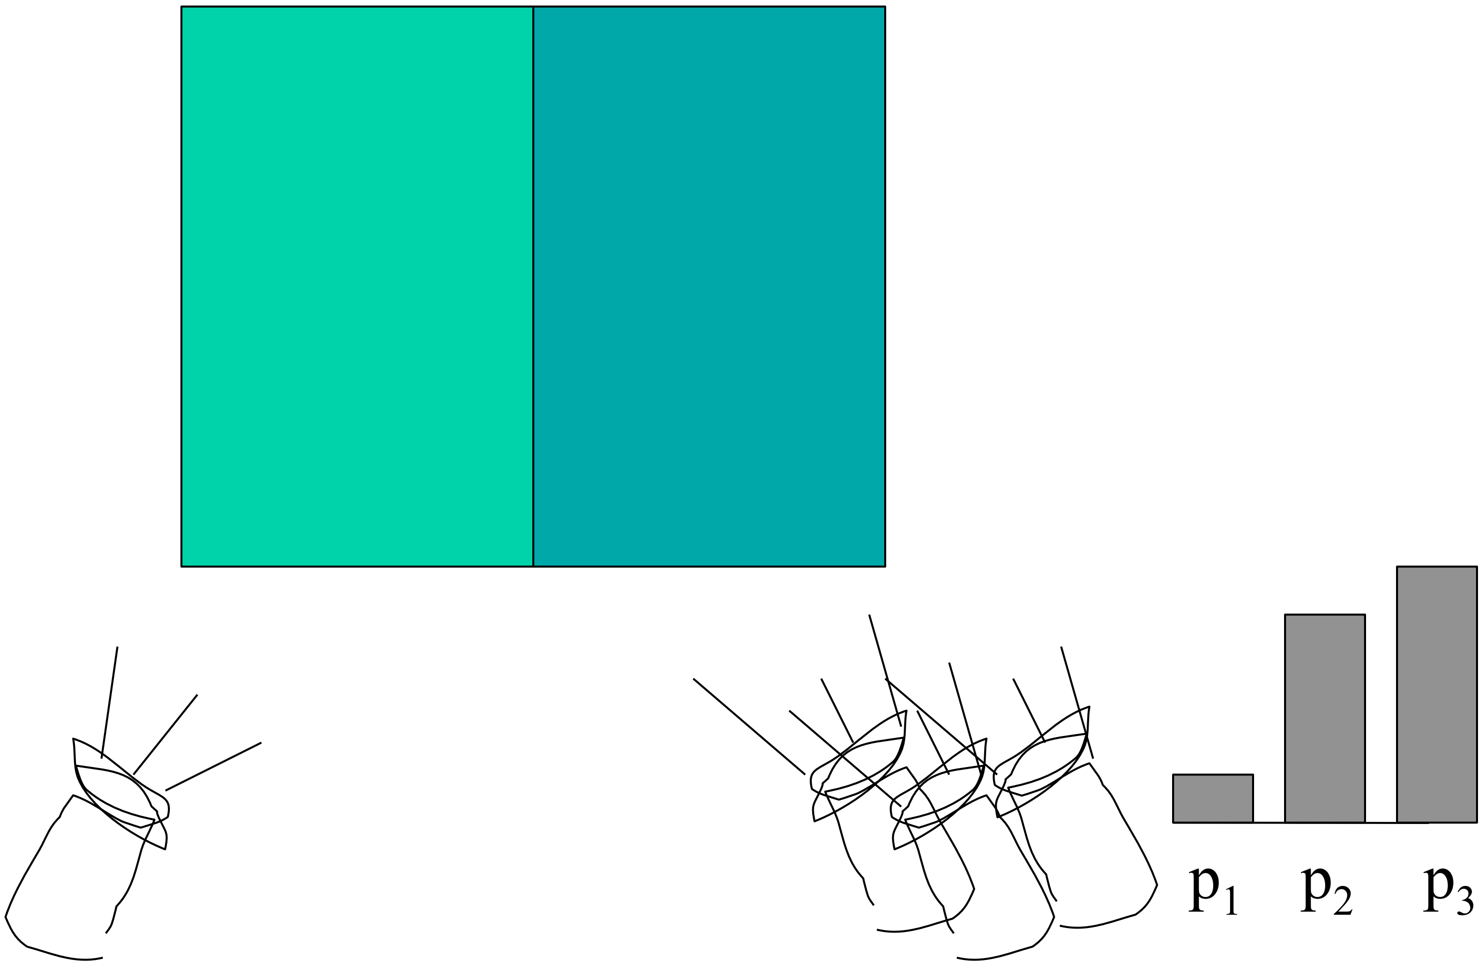

Color matching experiment 1

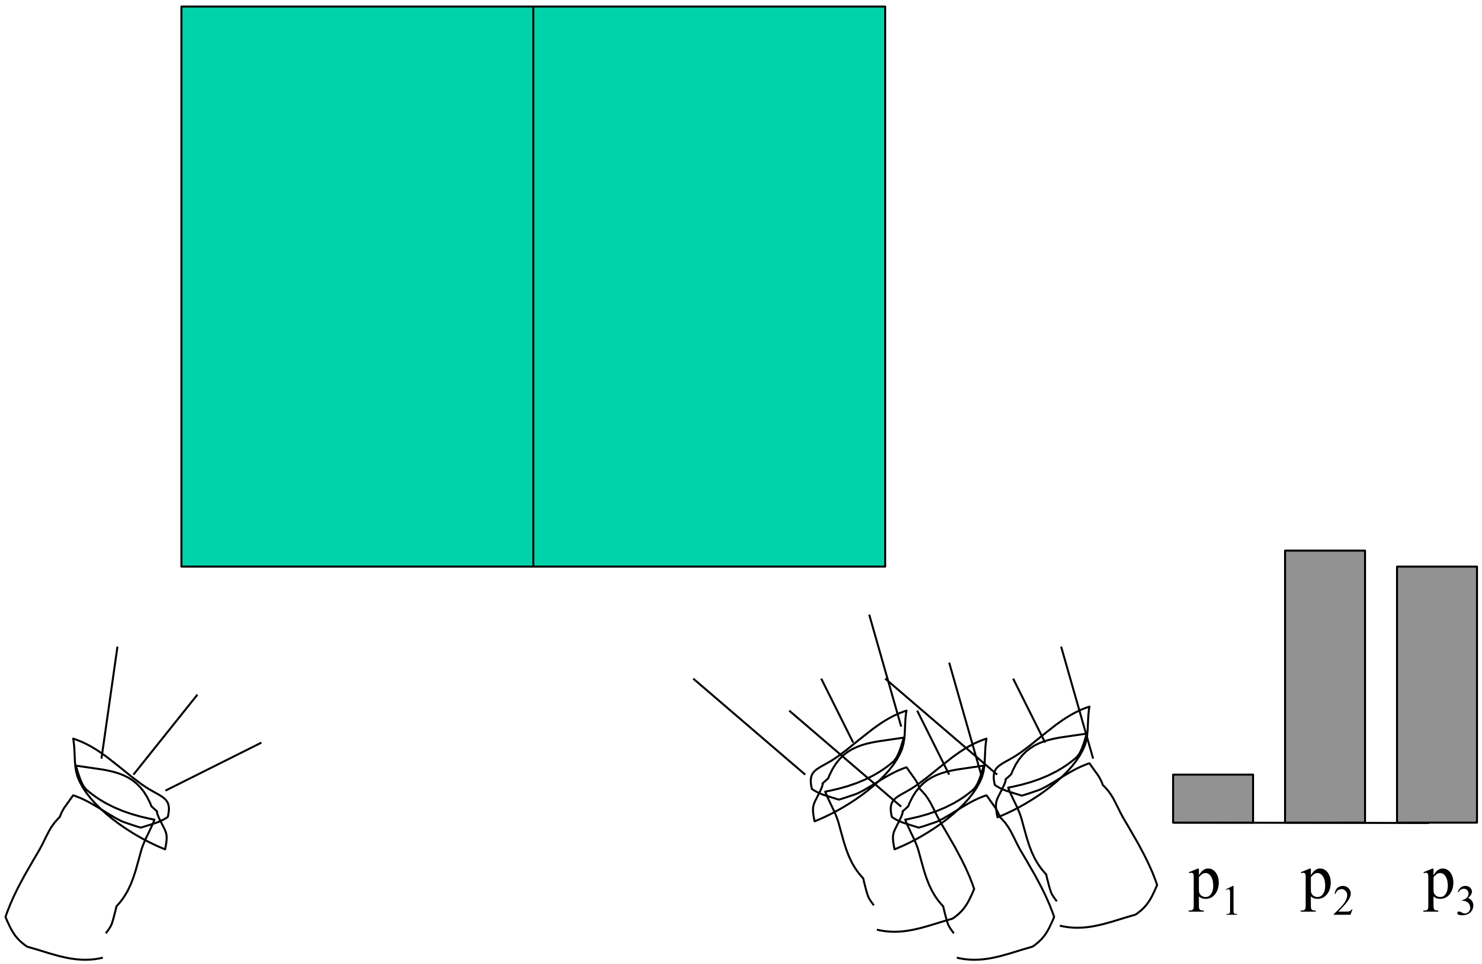

Color matching experiment 1

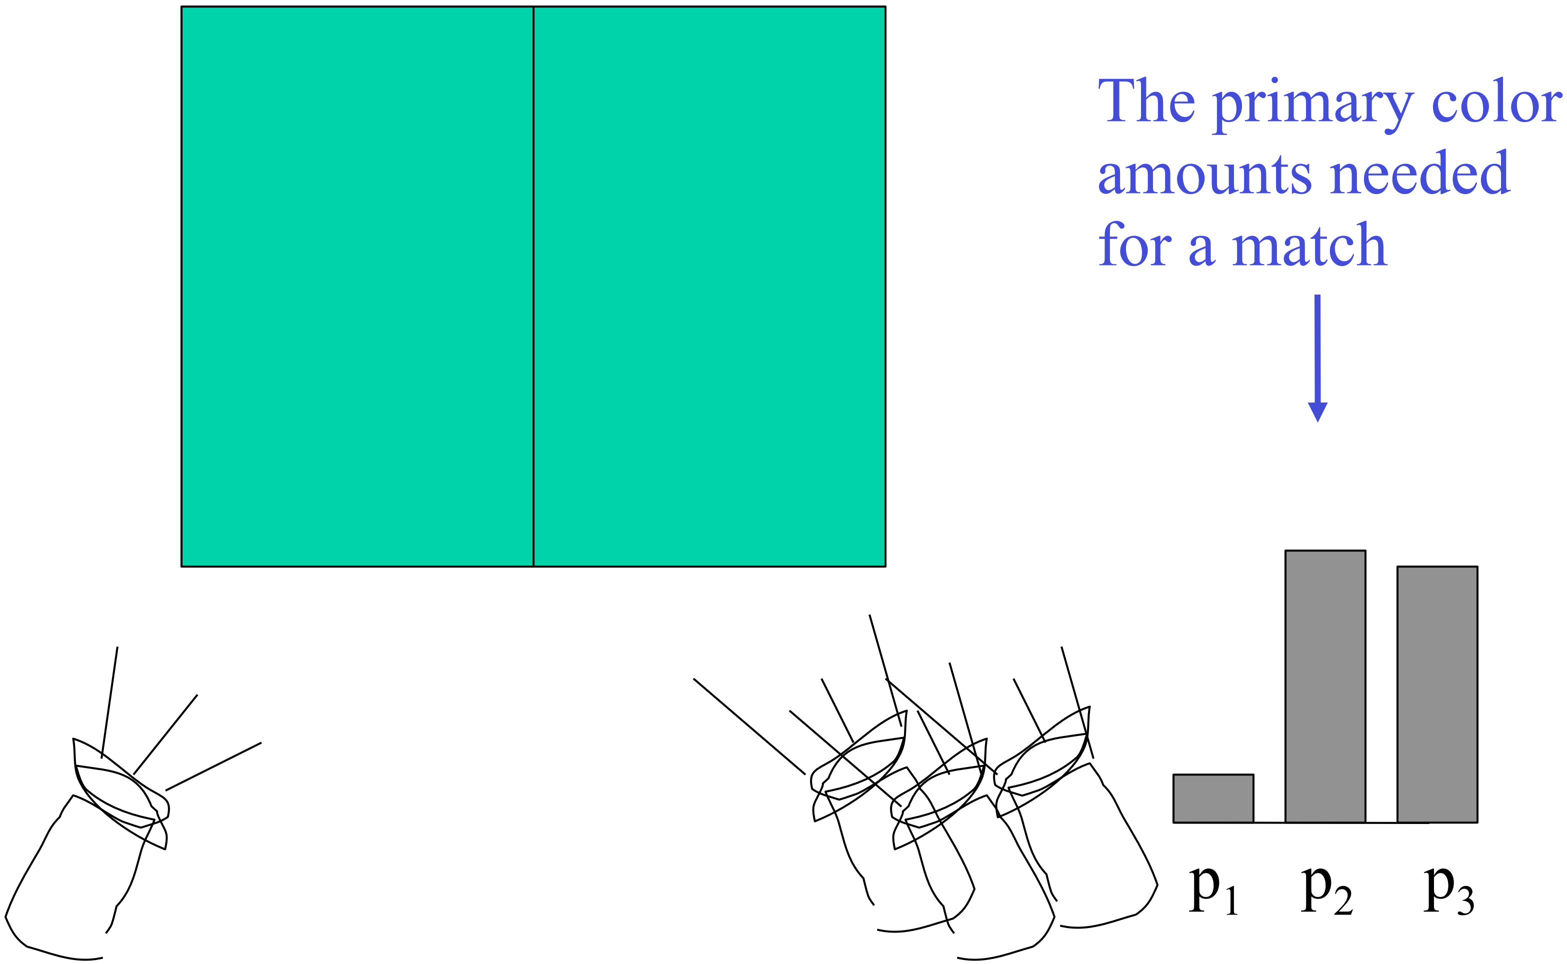

Color matching experiment 2

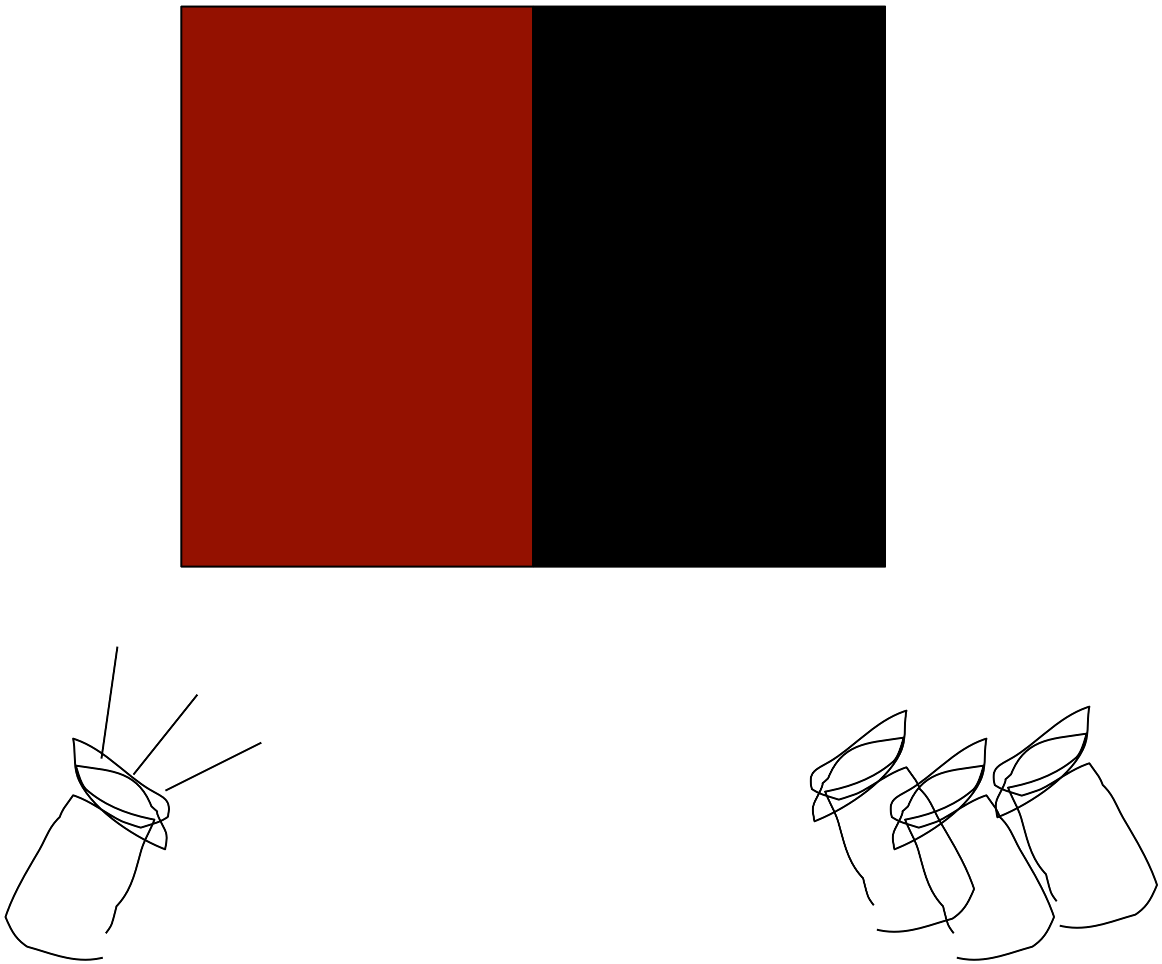

Color matching experiment 2

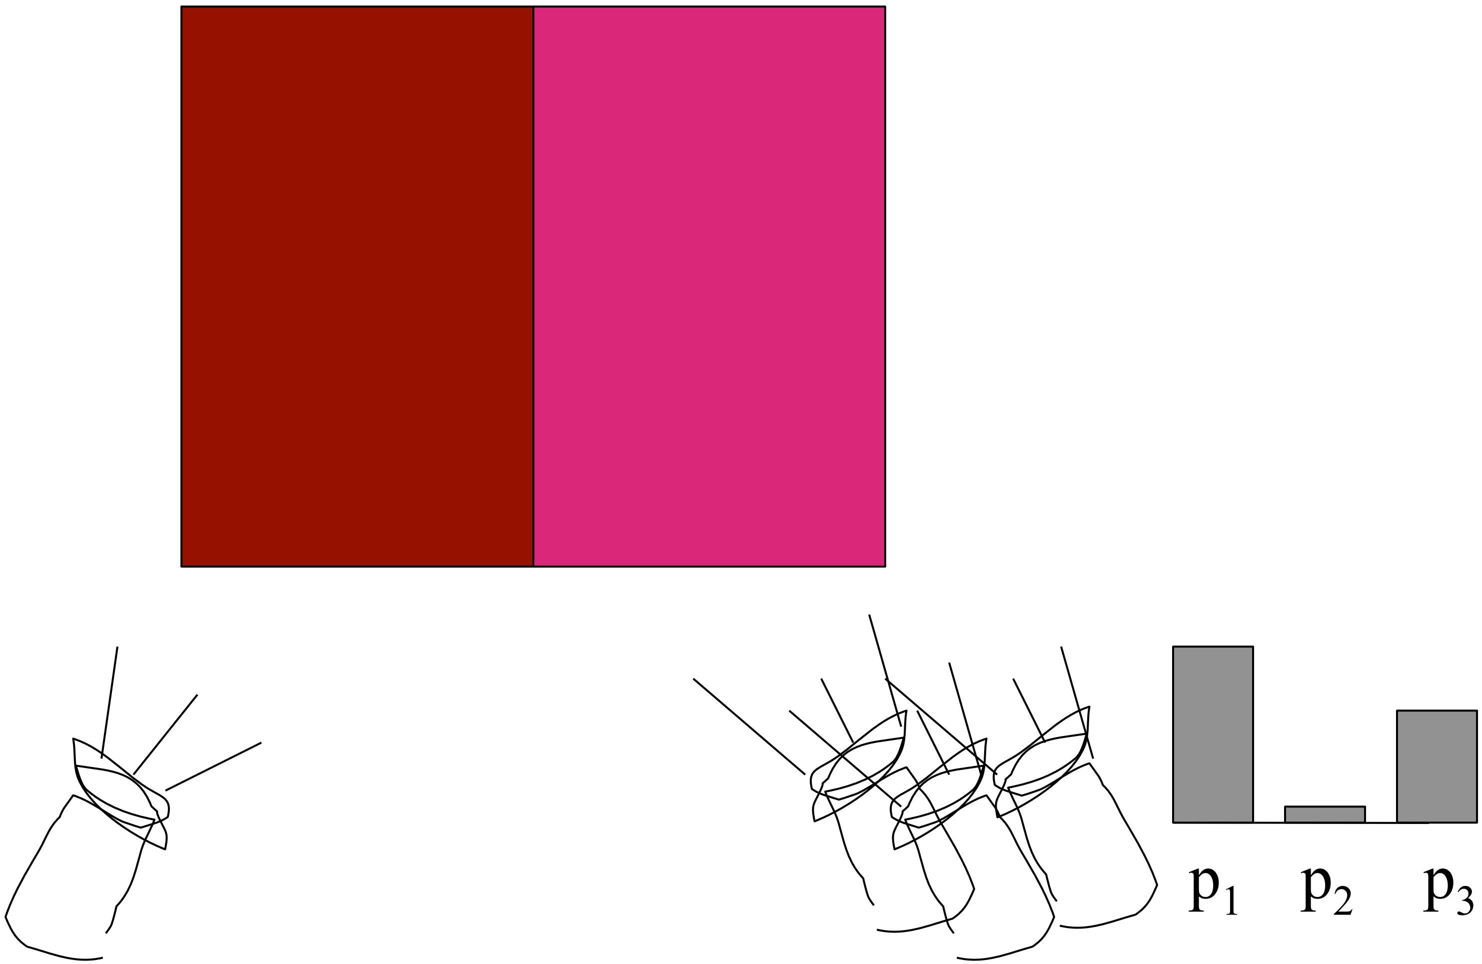

Color matching experiment 2

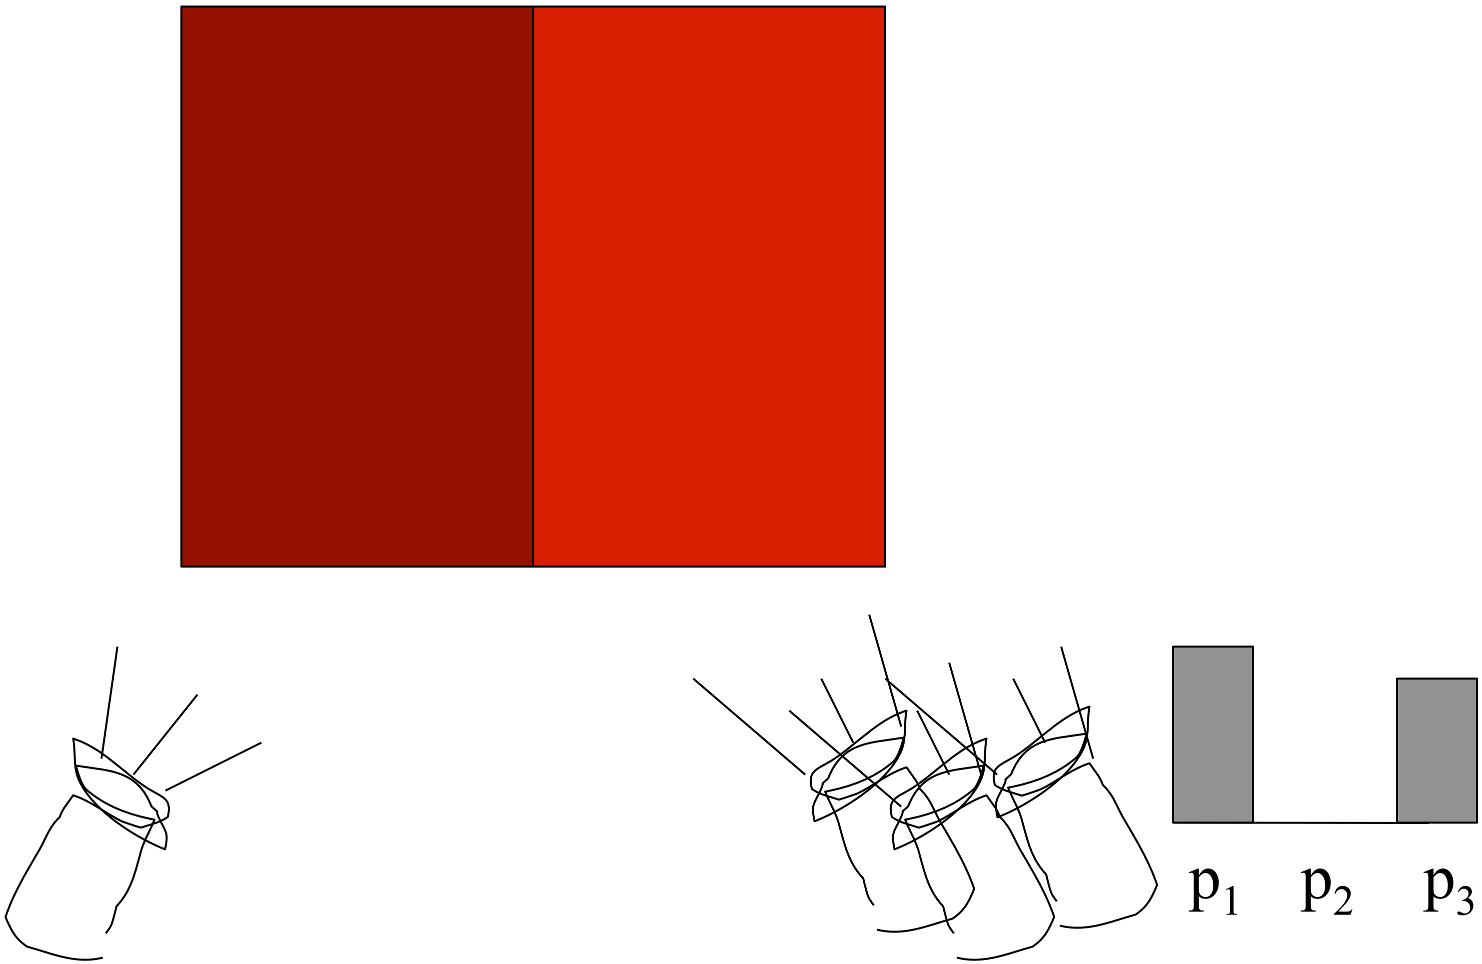

Color matching experiment 2

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

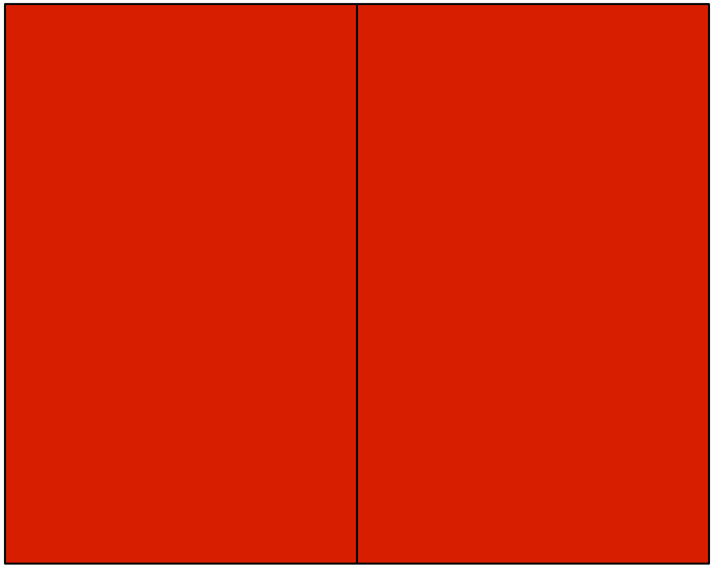

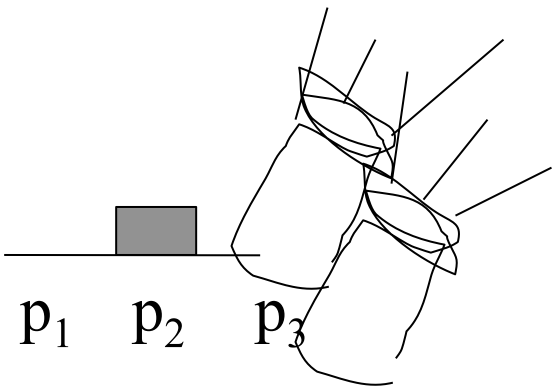

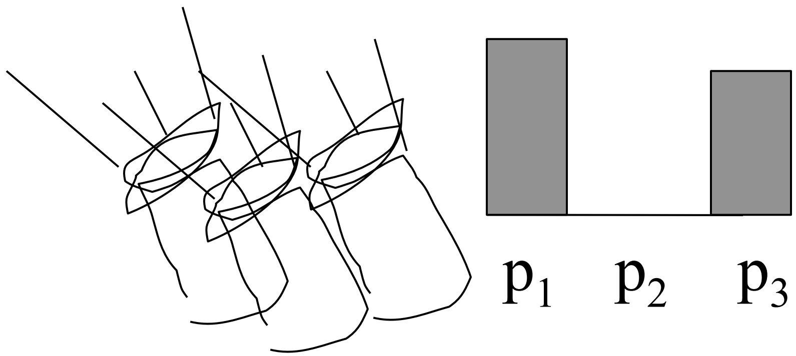

Color matching experiment 2

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

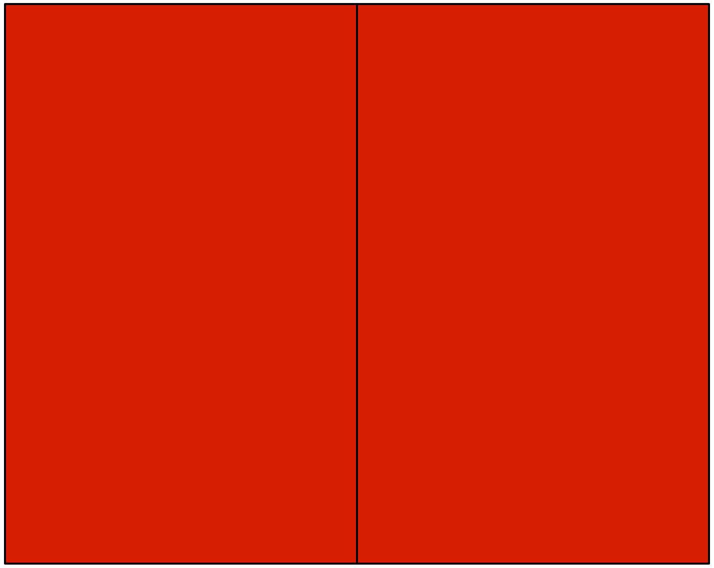

The primary color amounts needed for a match:

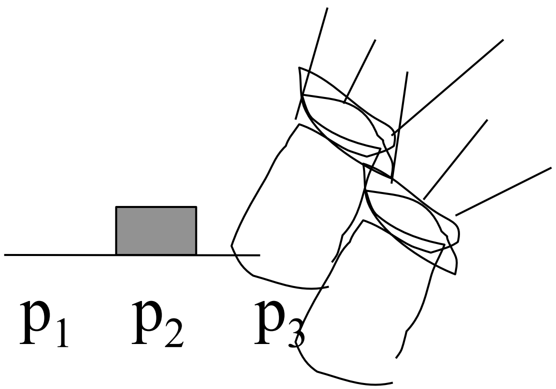

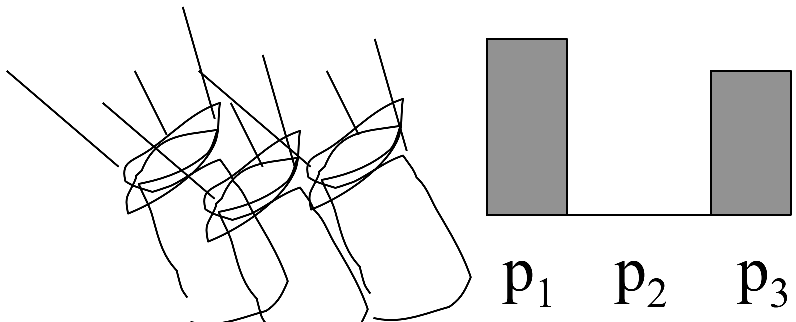

Color matching experiment 2

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

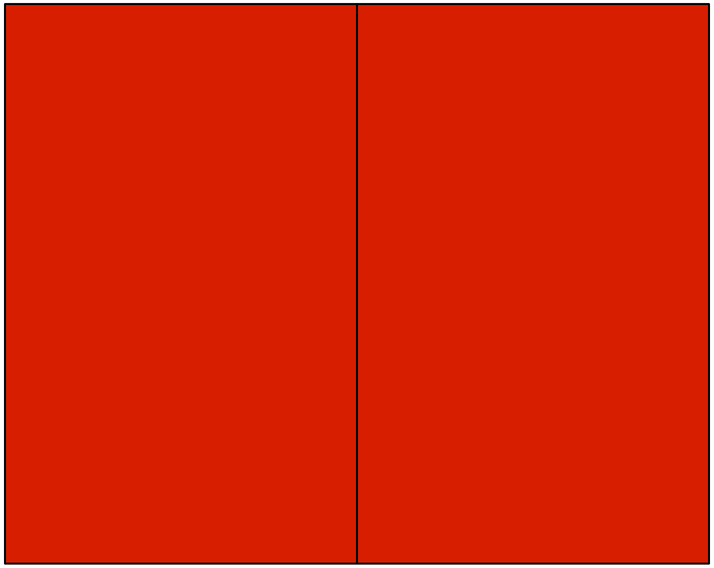

The primary color amounts needed for a match:

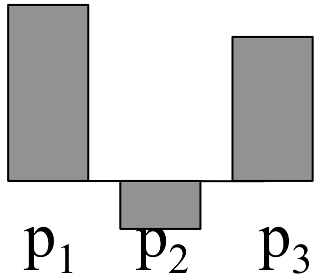

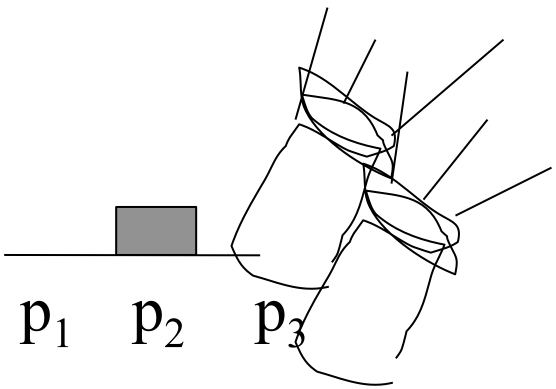

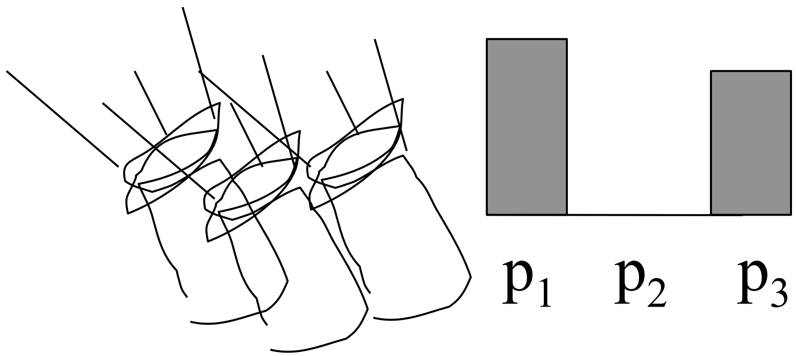

Color matching superposition (Grassman's laws)

If A_1 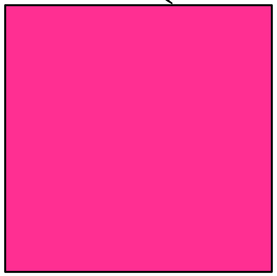 matches B_1 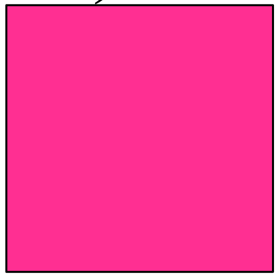

and A_2 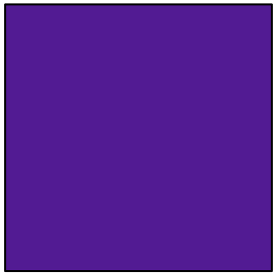 matches B_2 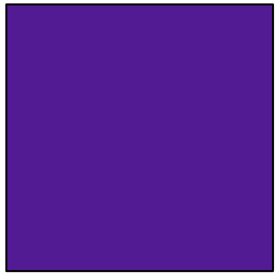

then $A_1 + A_2$ 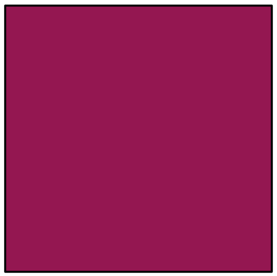 matches $B_1 + B_2$ 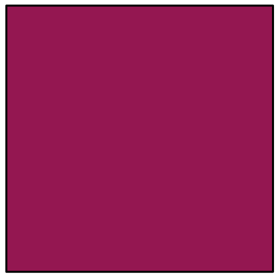

To measure a color

1. Choose a set of 3 primary colors (three power spectra).
2. Determine how much of each primary needs to be added to a probe signal to match the test light.

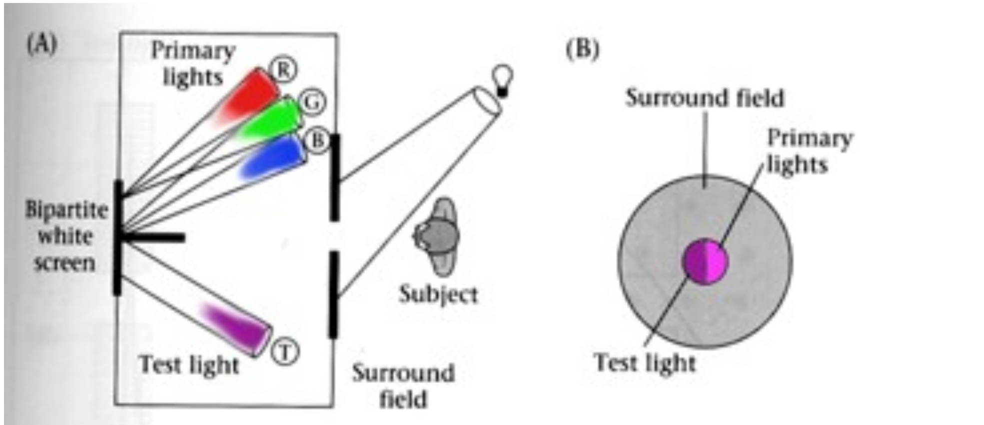

4.10 THE COLOR-MATCHING EXPERIMENT. The observer views a bipartite field and adjusts the intensities of the three primary lights to match the appearance of the test light. (A) A top view of the experimental apparatus. (B) The appearance of the stimuli to the observer. After Judd and Wyszecki, 1975.

To measure a color

1. Choose a set of 3 primary colors (three power spectra).
2. Determine how much of each primary needs to be added to a probe signal to match the test light.

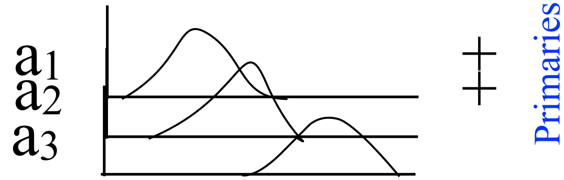

weighted sum of primaries

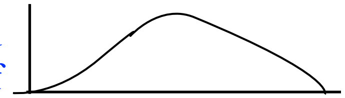

project

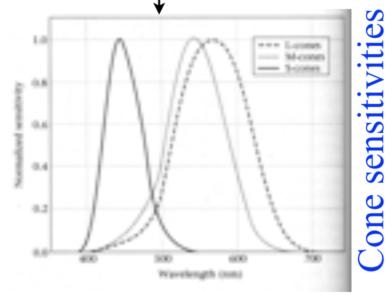

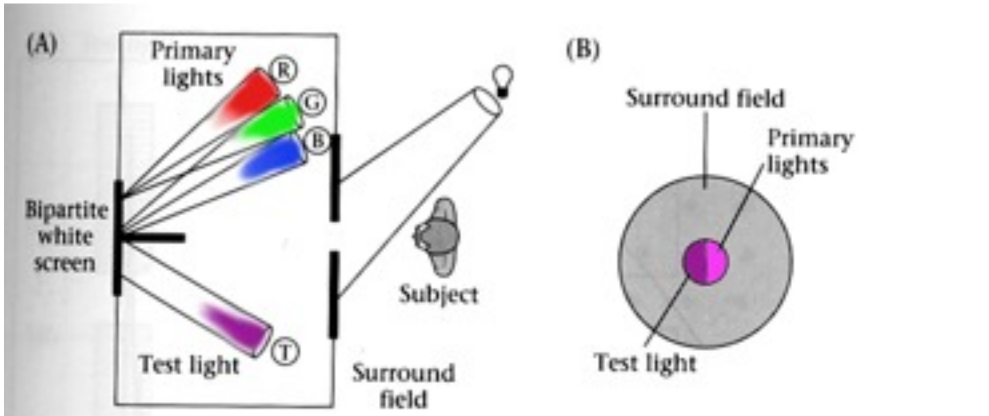

4.10 THE COLOR-MATCHING EXPERIMENT. The observer views a bipartite field and adjusts the intensities of the three primary lights to match the appearance of the test light. (A) A top view of the experimental apparatus. (B) The appearance of the stimuli to the observer. After Judd and Wyszecki, 1975.

test light

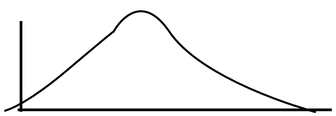

project

Cone sensitivities

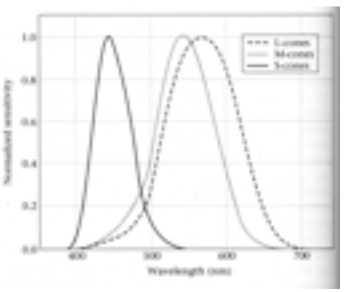

L, M, S responses

Cone response curves as basis vectors in a 3-d subspace of light power spectra

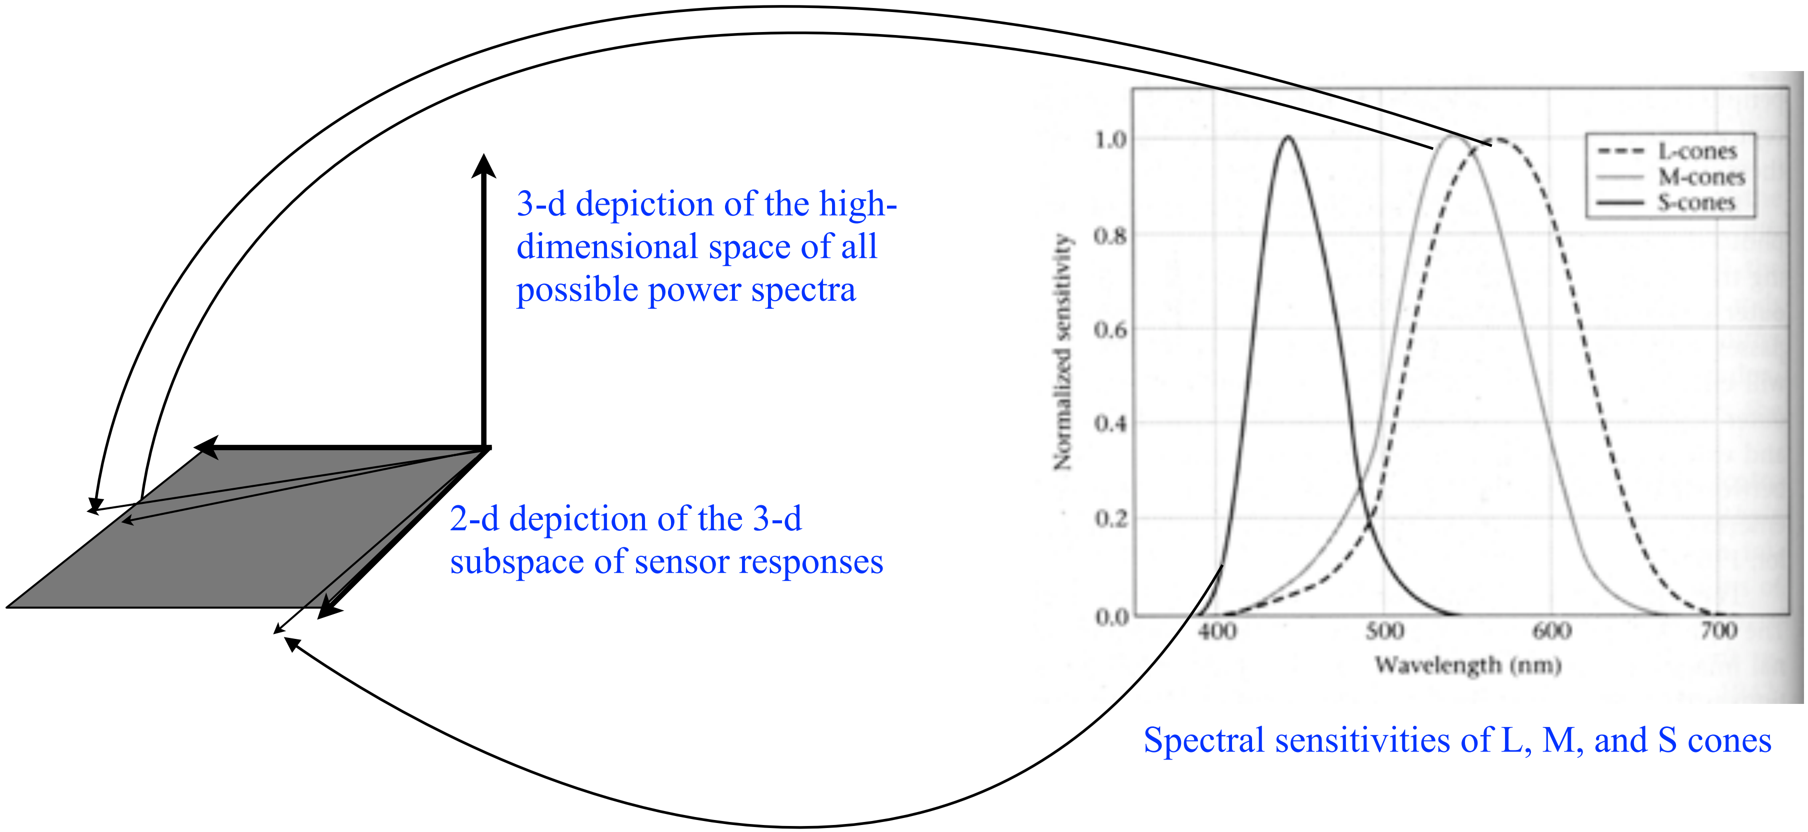

What we need from a color measurement system

- Given a color, how do you assign a number to it?
- Given an input power spectrum, what is its numerical color value, and how do we control our printing/projection/cooking system to match it?

"Color matching functions" let us find other basis vectors for the eye response subspace of light power spectra

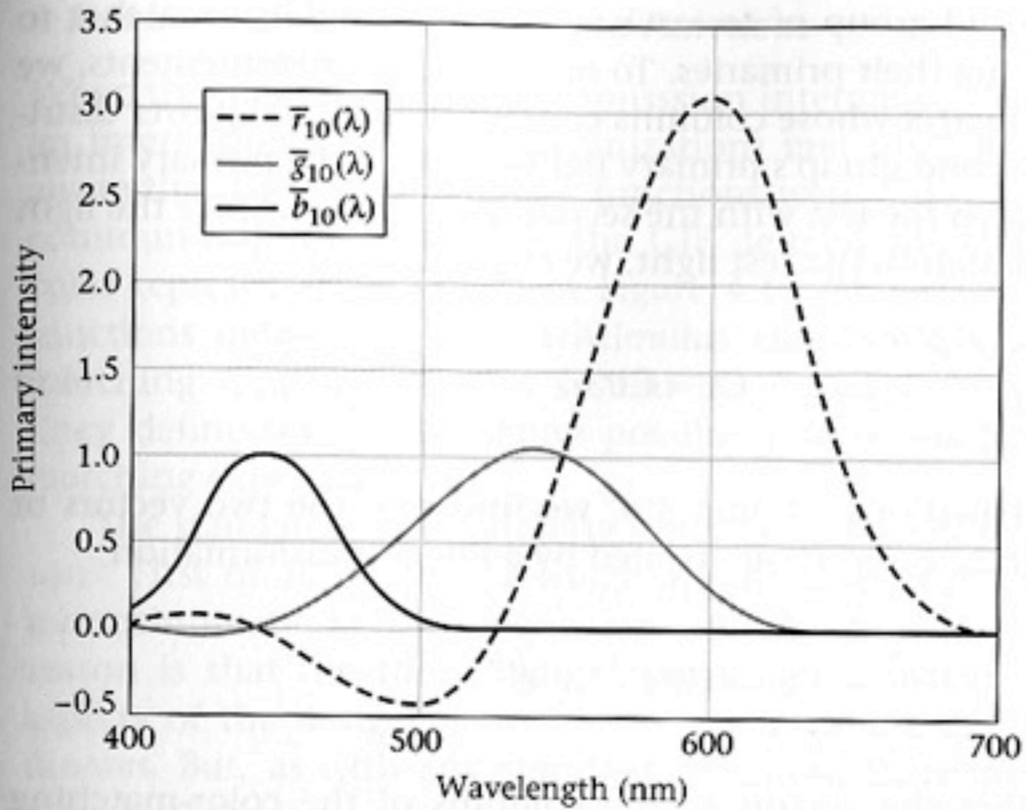

p_1 = 645.2 nm

p_2 = 525.3 nm

p_3 = 444.4 nm

4.13 THE COLOR-MATCHING FUNCTIONS ARE THE ROWS OF THE COLOR-MATCHING SYSTEM MATRIX. The functions measured by Stiles and Burch (1959) using a 10-degree bipartite field and primary lights at the wavelengths 645.2 nm, 525.3 nm, and 444.4 nm with unit radiant power are shown. The three functions in this figure are called $\bar{r}_{10}(\lambda)$, $\bar{g}_{10}(\lambda)$, and $\bar{b}_{10}(\lambda)$.

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

"Color matching functions" let us find other basis vectors for the eye response subspace of light power spectra

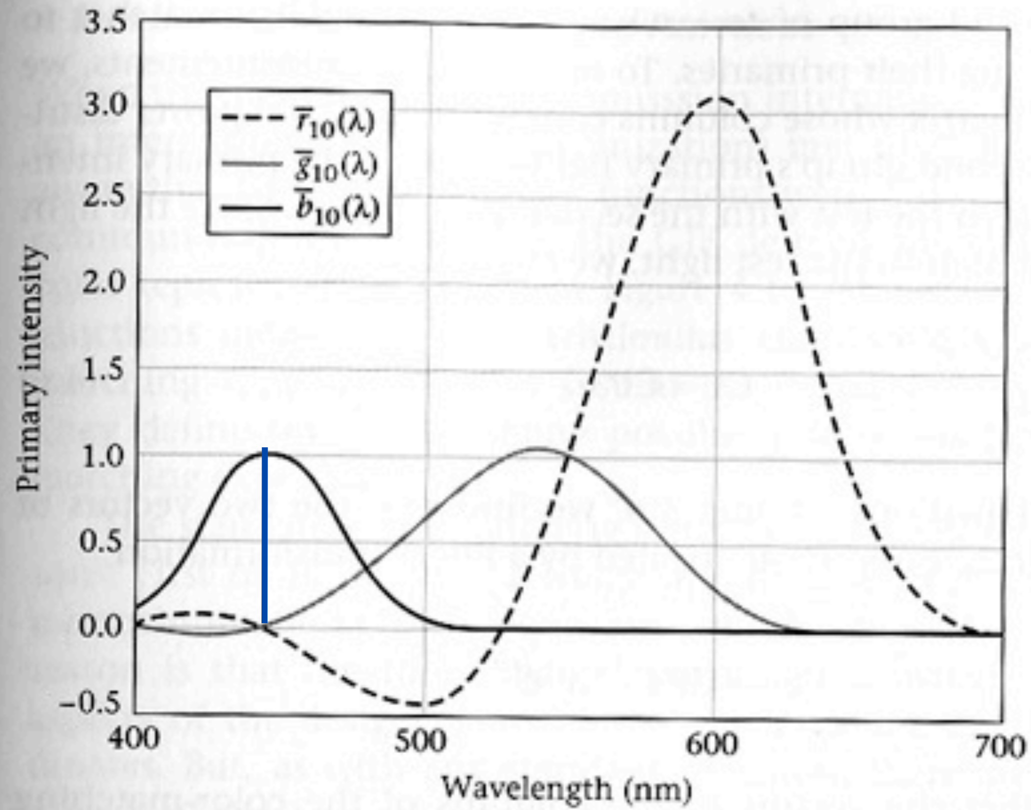

p_1 = 645.2 nm

p_2 = 525.3 nm

p_3 = 444.4 nm

4.13 THE COLOR-MATCHING FUNCTIONS ARE THE ROWS OF THE COLOR-MATCHING SYSTEM MATRIX. The functions measured by Stiles and Burch (1959) using a 10-degree bipartite field and primary lights at the wavelengths 645.2 nm, 525.3 nm, and 444.4 nm with unit radiant power are shown. The three functions in this figure are called $\bar{r}_{10}(\lambda)$, $\bar{g}_{10}(\lambda)$, and $\bar{b}_{10}(\lambda)$.

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

"Color matching functions" let us find other basis vectors for the eye response subspace of light power spectra

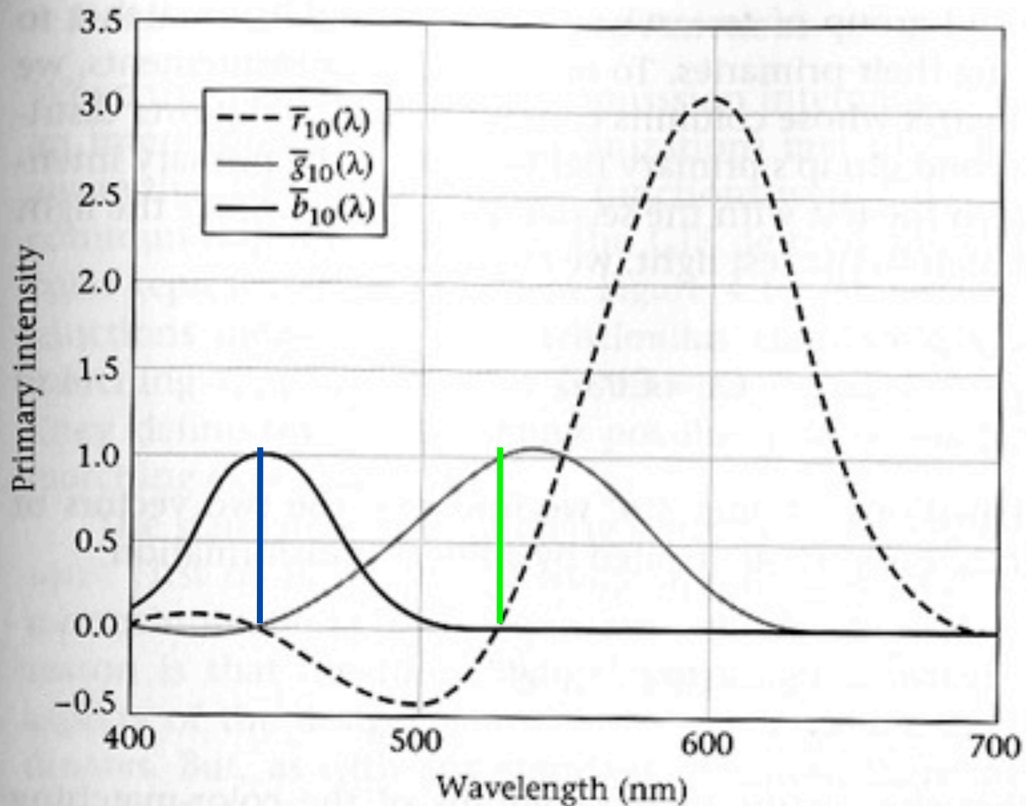

p_1 = 645.2 nm

p_2 = 525.3 nm

p_3 = 444.4 nm

4.13 THE COLOR-MATCHING FUNCTIONS ARE THE ROWS OF THE COLOR-MATCHING SYSTEM MATRIX. The functions measured by Stiles and Burch (1959) using a 10-degree bipartite field and primary lights at the wavelengths 645.2 nm, 525.3 nm, and 444.4 nm with unit radiant power are shown. The three functions in this figure are called $\bar{r}_{10}(\lambda)$, $\bar{g}_{10}(\lambda)$, and $\bar{b}_{10}(\lambda)$.

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

“Color matching functions” let us find other basis vectors for the eye response subspace of light power spectra

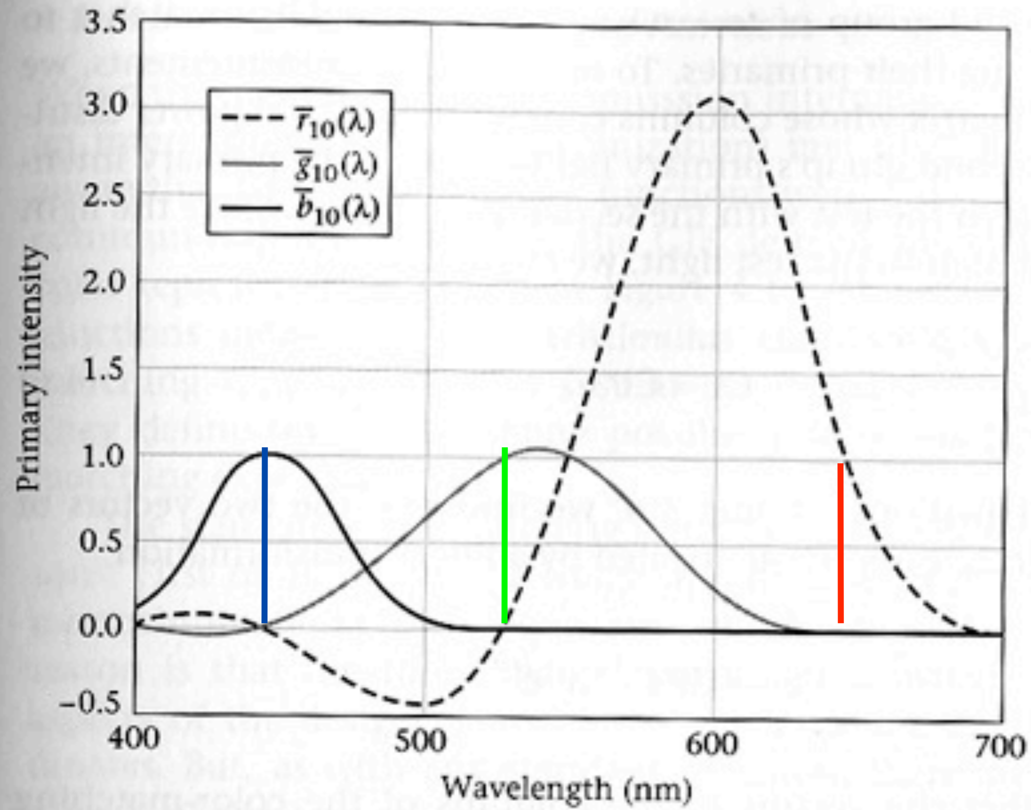

$p_1 = 645.2$ nm

$p_2 = 525.3$ nm

$p_3 = 444.4$ nm

4.13 THE COLOR-MATCHING FUNCTIONS ARE THE ROWS OF THE COLOR-MATCHING SYSTEM MATRIX. The functions measured by Stiles and Burch (1959) using a 10-degree bipartite field and primary lights at the wavelengths 645.2 nm, 525.3 nm, and 444.4 nm with unit radiant power are shown. The three functions in this figure are called $\bar{r}_{10}(\lambda)$, $\bar{g}_{10}(\lambda)$, and $\bar{b}_{10}(\lambda)$.

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

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

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

We know that a monochromatic light of λ_i wavelength will be matched by the amounts

of each primary.

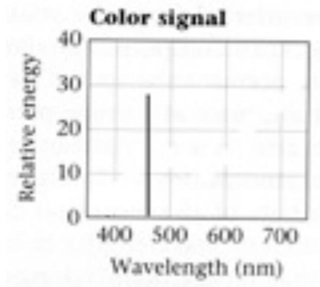

Using the color matching functions to predict the primary match to a new spectral signal

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

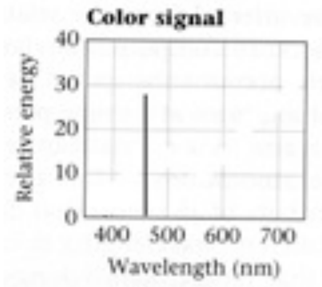

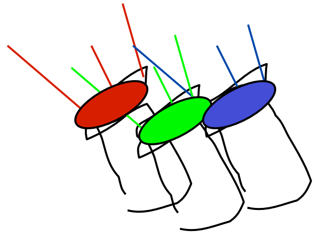

Using the color matching functions to predict the primary match to a new spectral signal

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

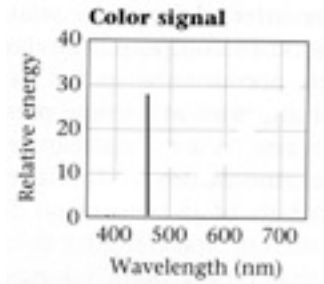

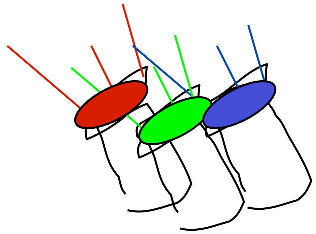

And any spectral signal can be thought of as a linear combination of very many monochromatic lights, with the linear coefficient given by the spectral power at each wavelength.

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

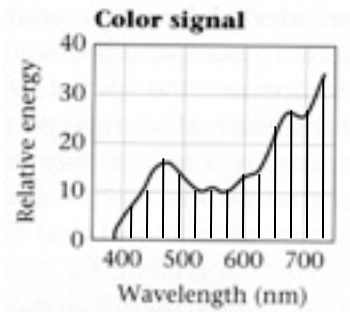

Using the color matching functions to predict the primary match to a new spectral signal

Using the 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

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

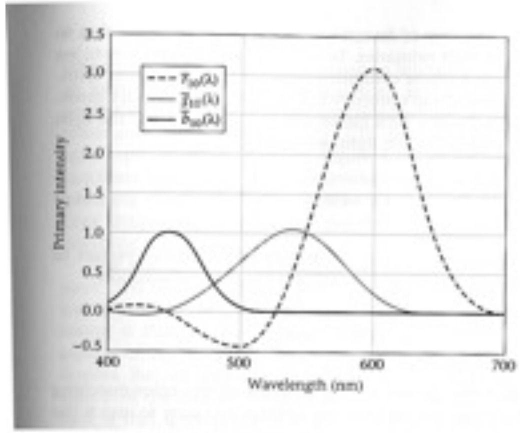

Using the color matching functions to predict the primary match to a new spectral signal

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

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

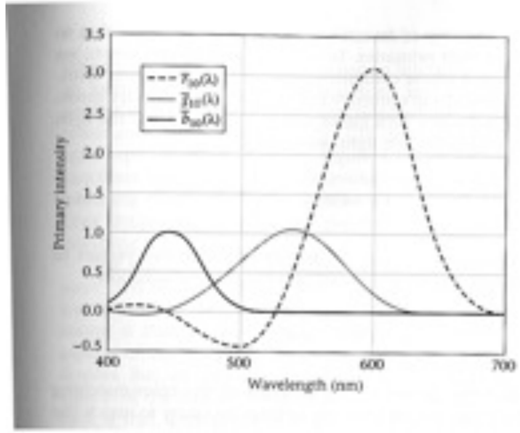

Let the new spectral signal be described by the vector t.

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

Using the 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

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

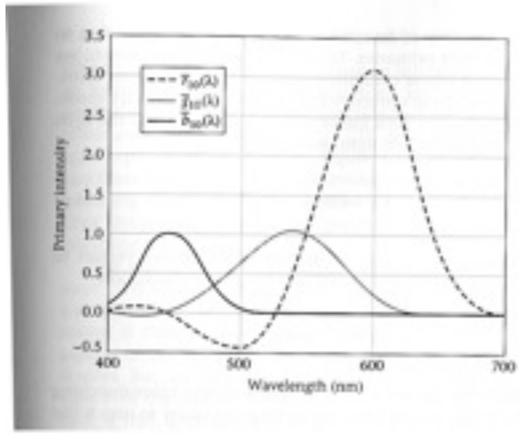

Let the new spectral signal be described by the vector t.

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

Then the amounts of each primary needed to match t are:

Using the 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

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

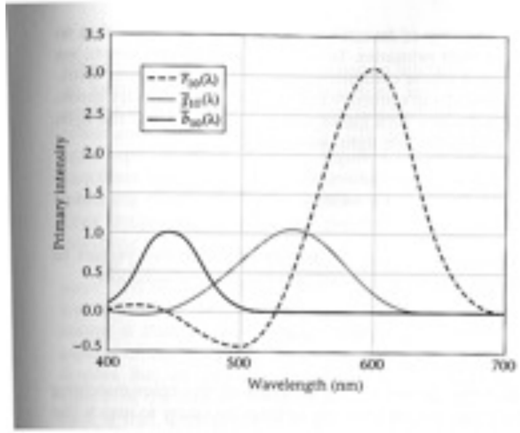

Let the new spectral signal be described by the vector t.

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

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

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.

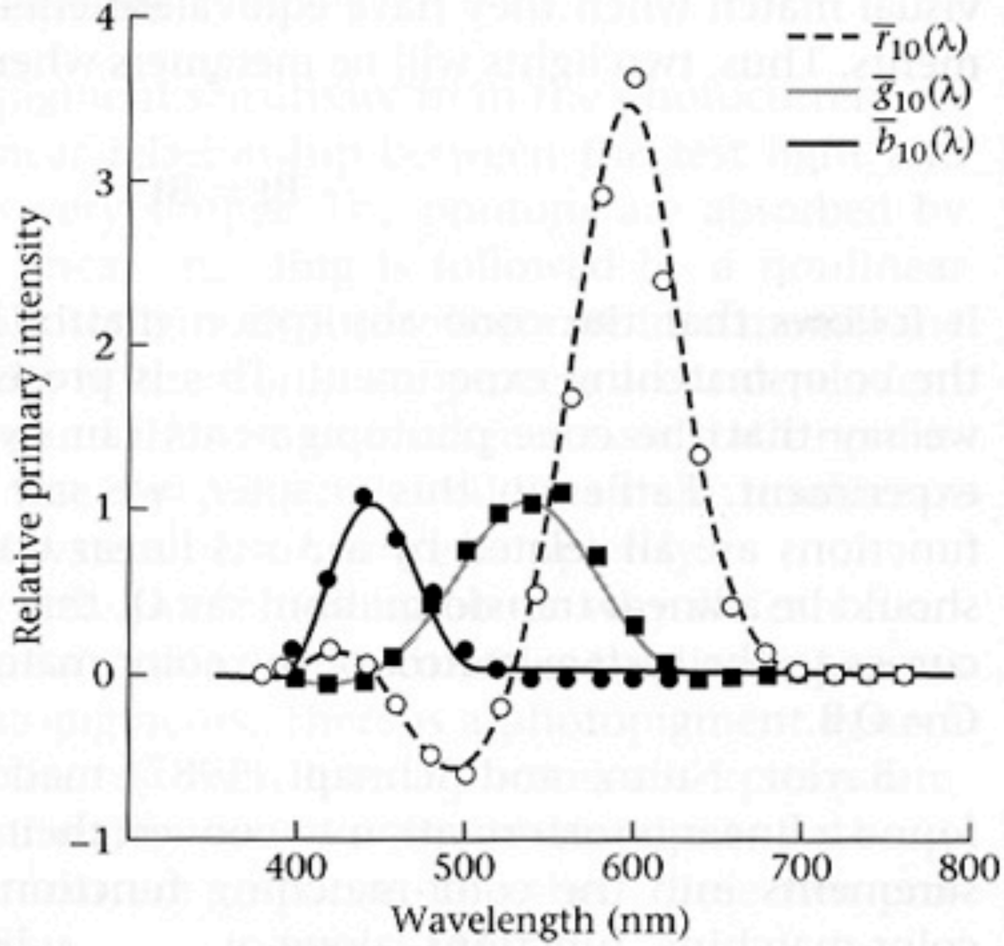

Foundations of Vision, by Brian Wandell, Sinauer Assoc., 1995

CIE XYZ color space

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

CIE XYZ color space

- Commission Internationale d'Eclairage, 1931 (International Commission on Illumination).

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

CIE XYZ color space

- Commission Internationale d’Eclairage, 1931 (International Commission on Illumination).
- “…as with any standards decision, there are some irratating aspects of the XYZ color-matching functions as well…no set of physically realizable primary lights that by direct measurement will yield the color matching functions.”

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

CIE XYZ color space

- Commission Internationale d'Eclairage, 1931 (International Commission on Illumination).
- "…as with any standards decision, there are some irratating aspects of the XYZ color-matching functions as well…no set of physically realizable primary lights that by direct measurement will yield the color matching functions."
- "Although they have served quite well as a technical standard, and are understood by the mandarins of vision science, they have served quite poorly as tools for explaining the discipline to new students and colleagues outside the field."

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

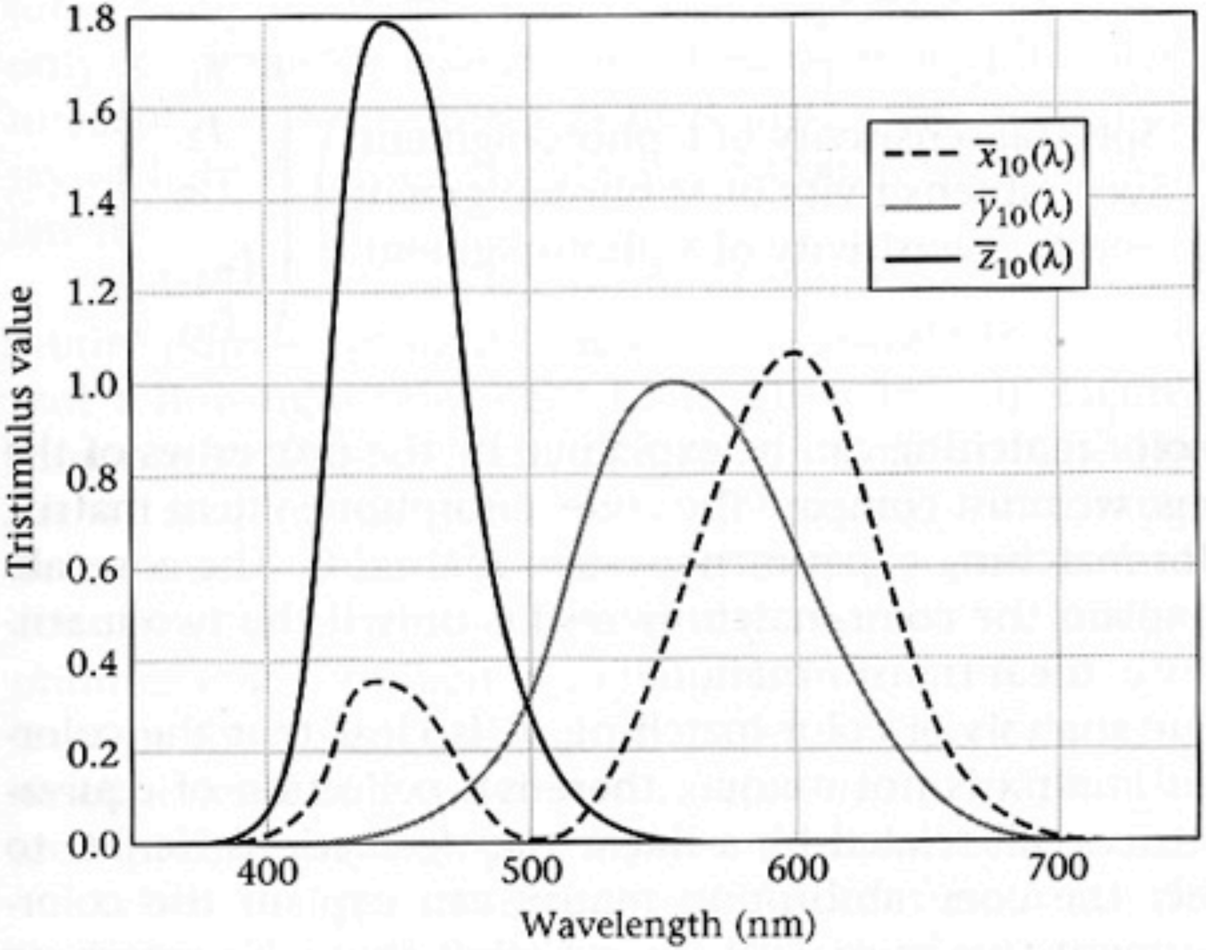

4.14 THE XYZ STANDARD COLOR-MATCHING FUNCTIONS. In 1931 the CIE standardized a set of color-matching functions for image interchange. These color-matching functions are called $\bar{x}(\lambda)$, $\bar{y}(\lambda)$, and $\bar{z}(\lambda)$. Industrial applications commonly describe the color properties of a light source using the three primary intensities needed to match the light source that can be computed from the XYZ color-matching functions.

CIE XYZ: Color matching functions are positive everywhere, but primaries are "imaginary" (require adding light to the test color's side in a color matching experiment). Usually compute x, y, where $x=X/(X+Y+Z)$

$y=Y/(X+Y+Z)$

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

Color metamerism: different spectra looking the same color

Two spectra, t and s, perceptually match when

$$C\vec{t} = C\vec{s}$$

where C are the color matching functions for some set of primaries.

Color metamerism: different spectra looking the same color

Two spectra, t and s, perceptually match when

$$C\vec{t} = C\vec{s}$$

where C are the color matching functions for some set of primaries.

Graphically,

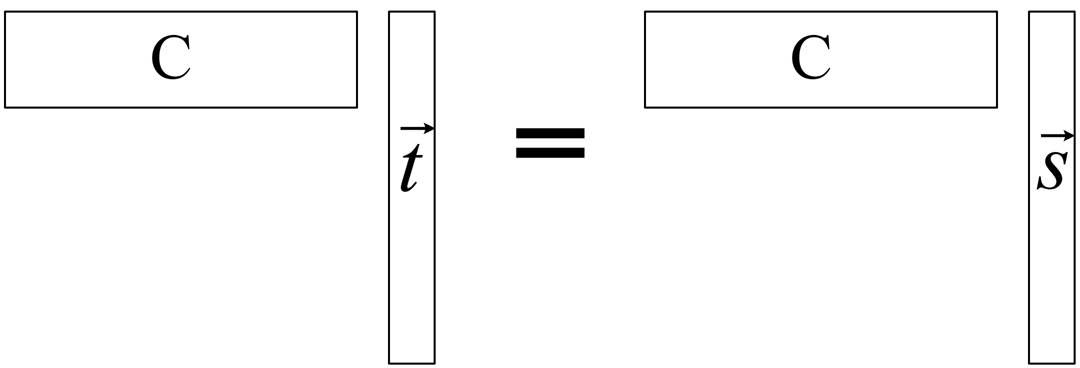

Metameric lights

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

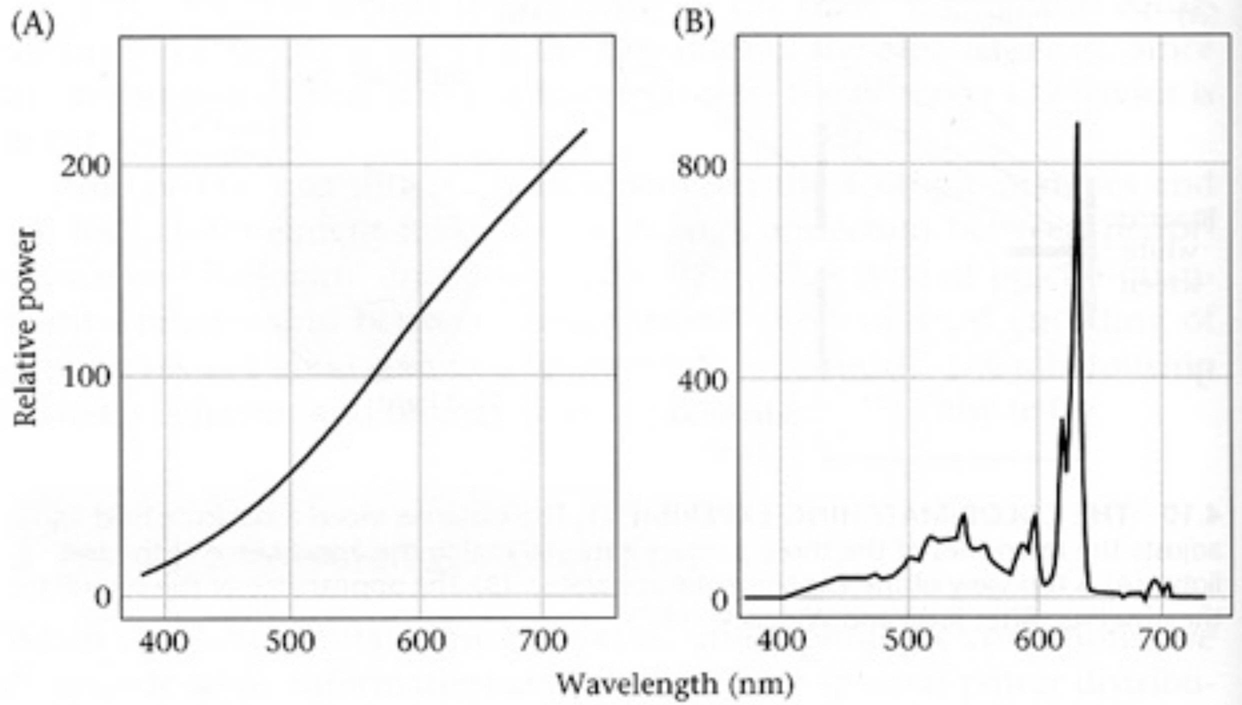

4.11 METAMERIC LIGHTS. Two lights with these spectral power distributions appear identical to most observers and are called metamers. (A) An approximation to the spectral power distribution of a tungsten bulb. (B) The spectral power distribution of light emitted from a conventional television monitor whose three phosphor intensities were set to match the light in panel A in appearance.

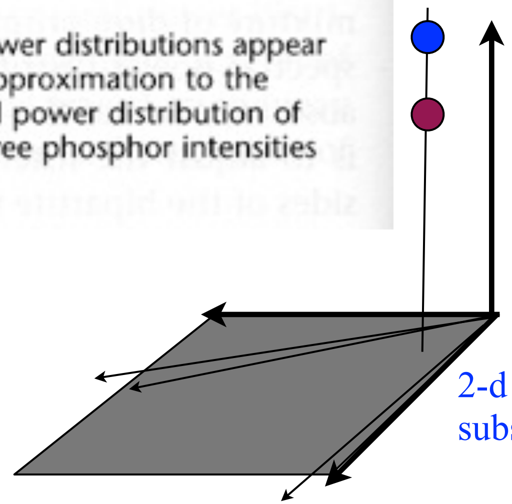

Concepts in color measurement

- What are colors?
  - Arise from power spectrum of light.
- How represent colors:
  - Pick primaries
  - Measure color matching functions (CMF’s)
  - Matrix mult power spectrum by CMF’s to find color as the 3 primary color values.
- How share color descriptions between people?
  - Standardize on a few sets of primaries.
  - Translate colors between systems of primaries.

Lecture outline

- Color physics.
- Color perception
  - part 1: assume perceived color only depends on light spectrum.
  - part 2: the more general case.

Color constancy demo

- We assumed that the spectrum impinging on your eye determines the object color. That’s often true, but not always.   Here’s a counter-example…

OSA Illumination Estimation Contest

http://color.psych.upenn.edu/osacontest2011/

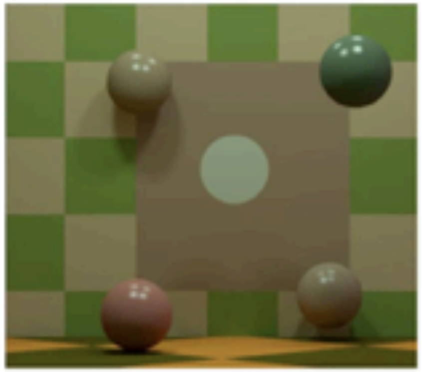

[Optical Society of America](#)

[Fall Vision Meeting](#)

Spectrum Recovery Competition, 2011

[David H. Brainard](#) and [Alex R. Wade](#)

- [The Contest](#)
- [The Prize](#)
- [Entering and The Rules](#)
- [Image Generation](#)
- [Calibration Image](#)
- [Wavelength Sampling](#)
- [Data File Format](#)
- [Illuminant Spectral Power Distributions](#)
- [Surface Reflectance Functions](#)
- [Cone Coordinates](#)
- [Error Measure](#)
- [Sample Program](#)
- [Image Data](#)
- [Frequently Asked Questions (FAQ)](#)
- [References](#)

http://color.psych.upenn.edu/osacontest2011/#images

Visited · Getting Started · Latest Headlines · Google · Apple · Yahoo! · YouTube · Wikipedia · News · Popular · Google Maps · Elsevier Editorial Sy...

nload zip archive containing all of the image data as well as the .mat files described above (~50 MB).

ach image, we provide the LMS cone coordinates in a MATLAB .mat file. Each of these images contains a N by M by 3 matrix called theImage. T
planes provide the L, M, and S cone coordinates at each pixel respectively. We also provide an JPEG rendering of each image. These are simply f
lization and should not be used as actual image data for the contest - these images were scaled and/or tone-mapped by hand to produce reasonable
es for display. The archive also contains the calibration image and its illuminant spectrum. For fun, the jpeg images are shown below.

Calibration Image:

Image 1:

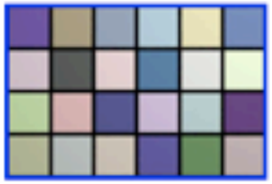

Image 2:

Image 3:

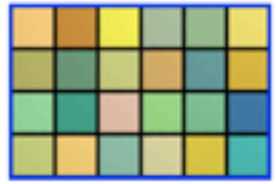

Image 4:

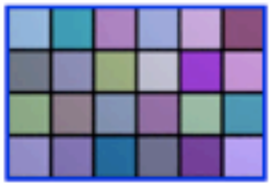

Image 5:

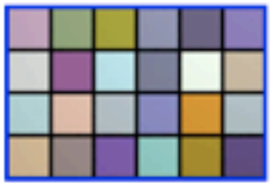

Image 6:

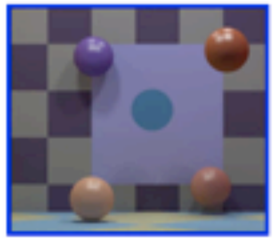

Image 7:

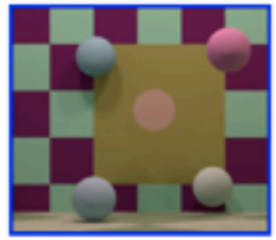

Image 8:

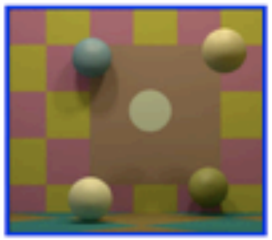

Image 9:

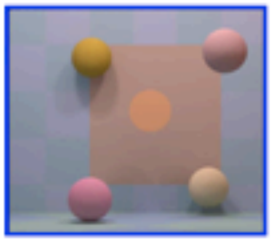

Image 10:

Rendering equation for jth observation

$$\begin{pmatrix} L_j \\ M_j \\ S_j \end{pmatrix} = \mathbf{E}^T(\mathbf{A}\vec{x}_j^s \;.*\; \mathbf{B}\vec{x}^i)$$

$$\begin{pmatrix} L_j \\ M_j \\ S_j \end{pmatrix} =$$

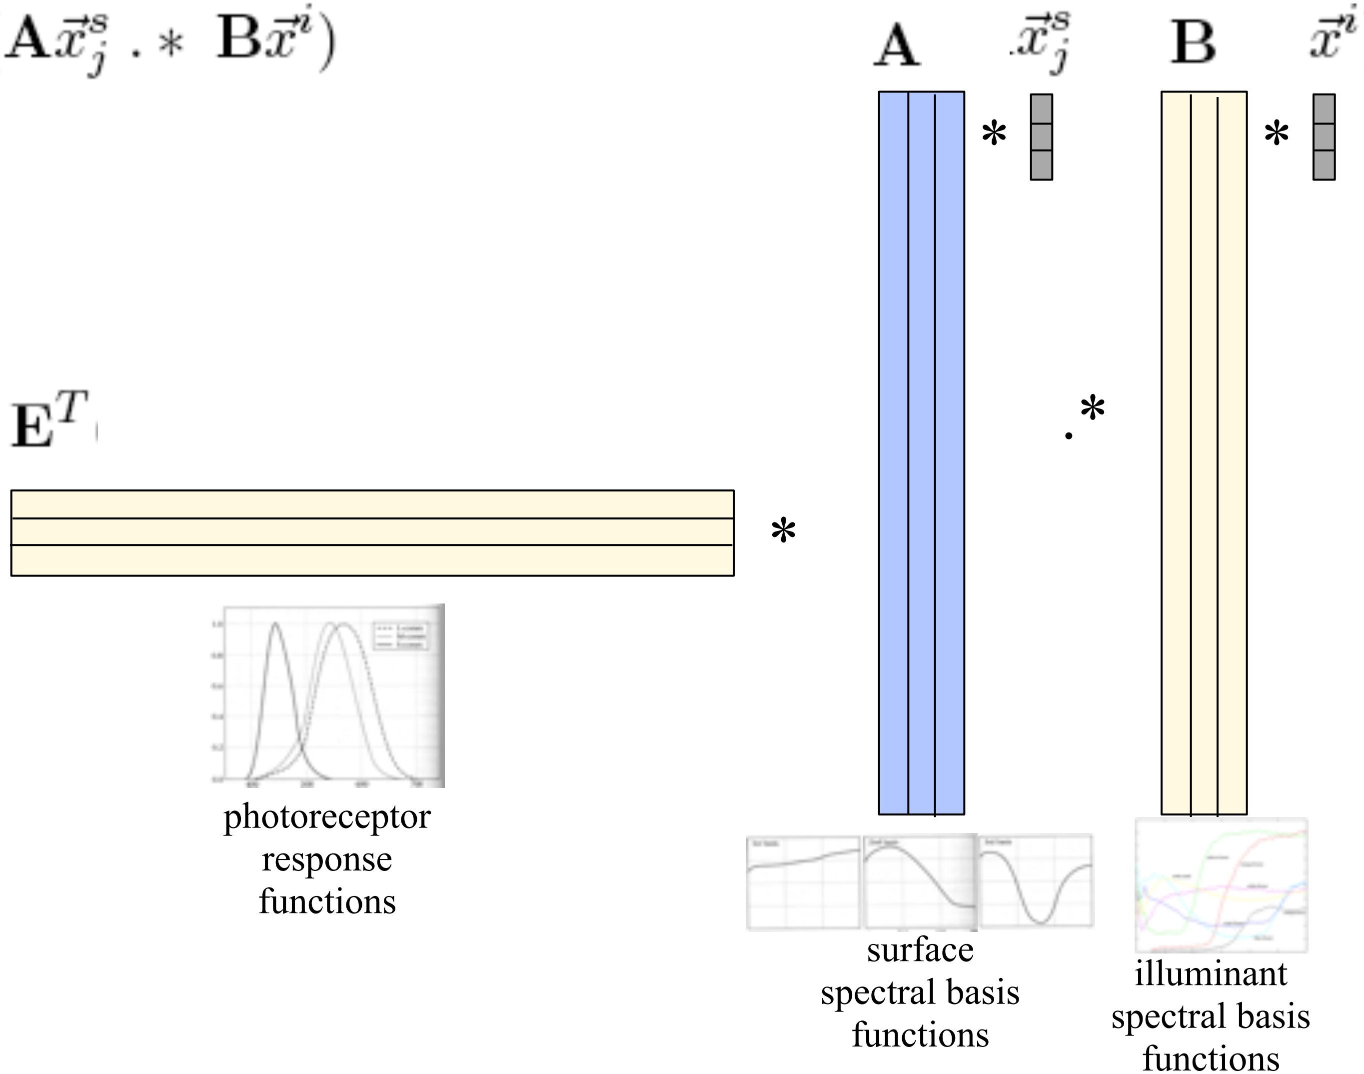

Color constancy solution 1: find white in the scene

Let the kth patch be the white one, with surface coefficients assumed to be $\vec{x}^W$
Then we can solve for the illuminant coefficient, $\vec{x}^i$

$$\begin{pmatrix} L_k \\ M_k \\ S_k \end{pmatrix} = \underbrace{\mathbf{E}^T(\mathbf{A}\vec{x}^W \; .* \; \mathbf{B}\vec{x}^i)}_{\text{a 3x3 matrix}}$$

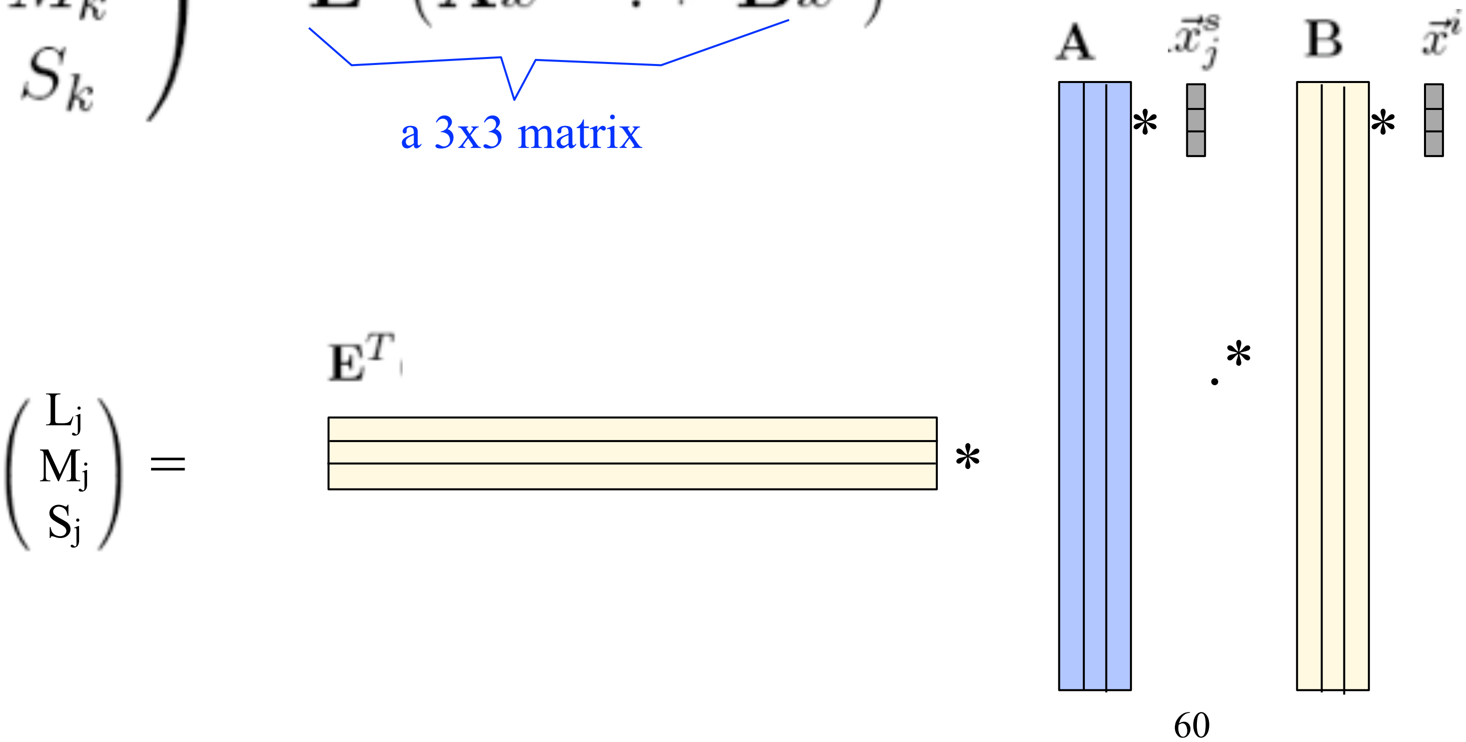

Color constancy solution 2: assume scene colors average to grey

$$\frac{1}{N}\sum_j \begin{pmatrix} L_j \\ M_j \\ S_j \end{pmatrix} = \mathbf{E}^T(\mathbf{A}\frac{1}{N}\sum_j \vec{x}^s_j \;.*\; \mathbf{B}\vec{x}^i)$$

$$= \underbrace{\mathbf{E}^T(\mathbf{A}\vec{x}^G \;.*\; \mathbf{B}}_{\text{a 3x3 matrix}}\vec{x}^i)$$

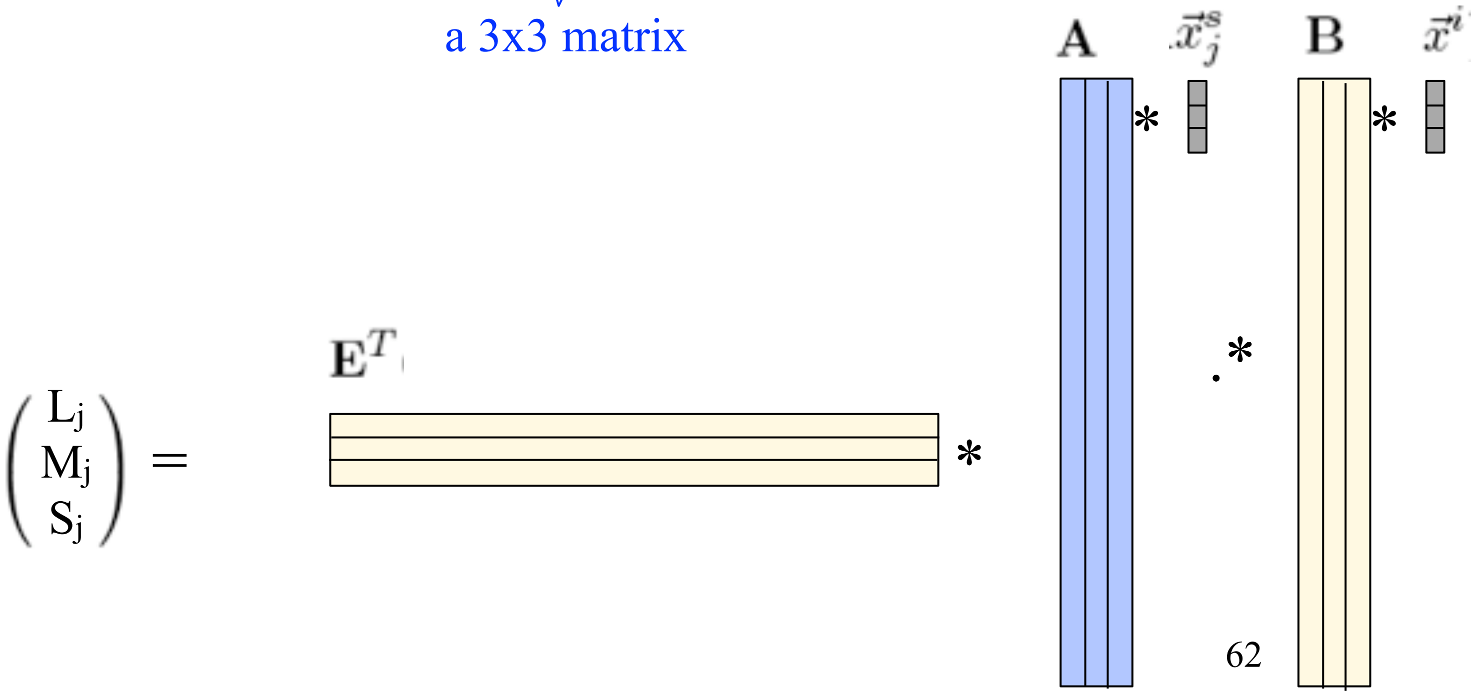

an image that violates both assumptions

http://1.bp.blogspot.com/_vsiS4vPB35s/S9FiRzKmyEI/AAAAAAAAAUc/TSb5RVWDM9Q/s1600/NATURE-GreenForest_1024x768.jpeg

Bayesian approach

Bayes rule

$$P(\vec{x}|\vec{y}) = kP(\vec{y}|\vec{x})P(\vec{x})$$

Likelihood

$$P(\vec{y}_j|\vec{x}^i, \vec{x}^s_j) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp \frac{-|\vec{y}_j - \vec{f}(\vec{x}^i, \vec{x}^s_j)|^2}{2\sigma^2}$$

$$\vec{y}_j = \begin{pmatrix} L_j \\ M_j \\ S_j \end{pmatrix}$$

$$\vec{f}(\vec{x}^i, \vec{x}^s_j) = \mathbf{E}^T(\mathbf{A}\vec{x}^s_j \ .* \ \mathbf{B}\vec{x}^i)$$

Posterior

$$P(\vec{x}|\vec{y}) = P(\vec{x}^i)\prod_j P(\vec{y}_j|\vec{x}^i, \vec{x}^s_j)P(\vec{x}^s_j)$$

Likelihood term for a b = 1 problem

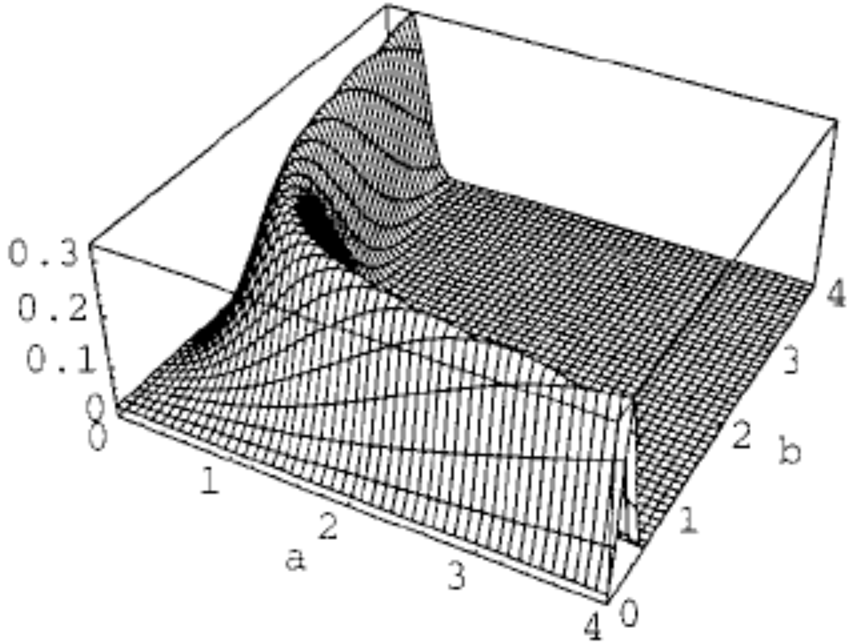

$$\begin{array}{c} L_j \\ M_j \\ S_j \end{array} = \mathbf{E}^T * \mathbf{A} * \vec{x}_j^s \; .* \; \mathbf{B} * \vec{x}^i$$

Bayesian approach: priors on surfaces and illuminants

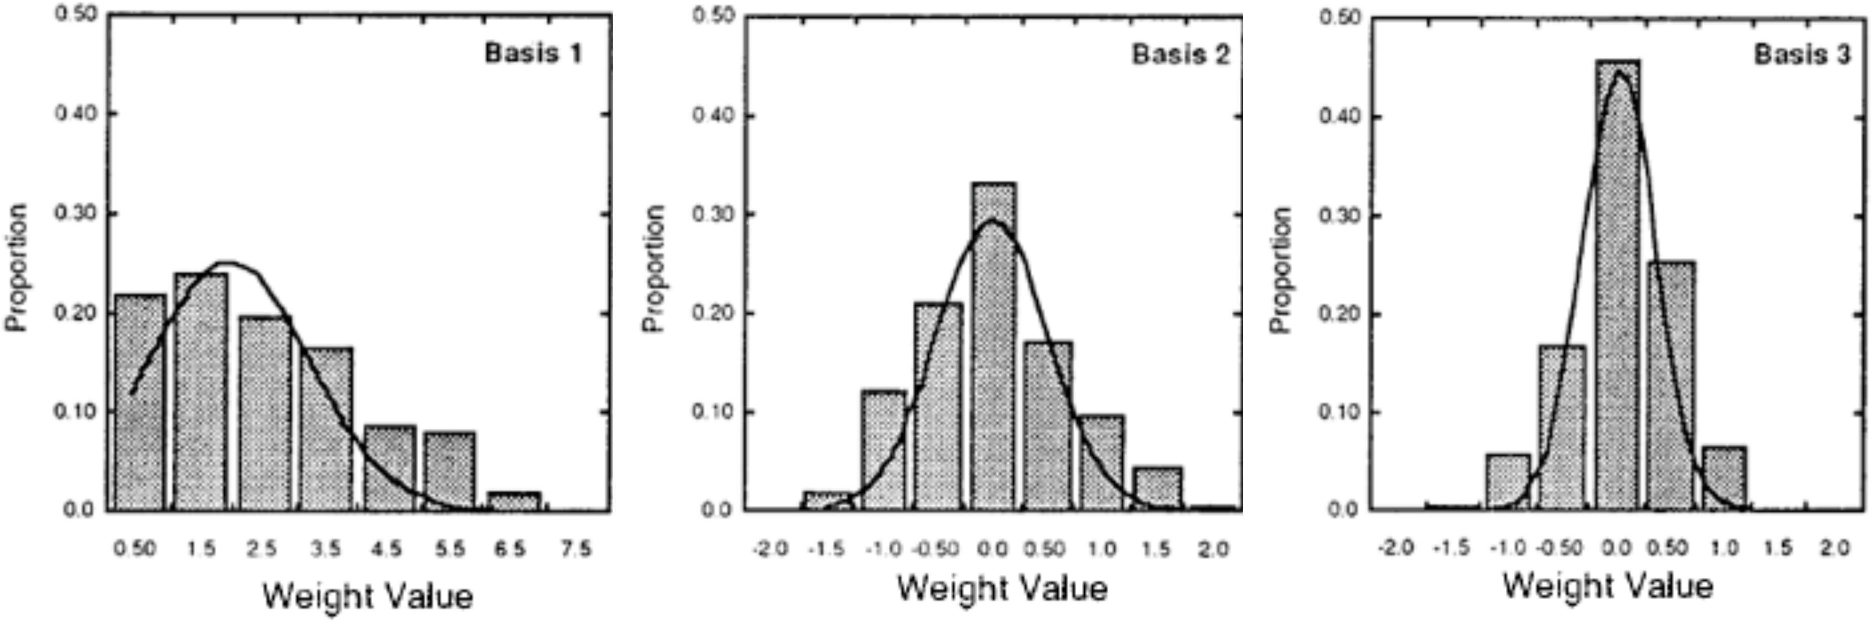

Distribution of surface weights. The histograms show the distribution of linear model weights derived from the measurements of Kelly *et al.*[68] and Nickerson.[69] Each histogram corresponds to one basis vector. The solid curves show the fit of a truncated trivariate normal distribution to the weights.

Picking a single best x

From the supplementary notes for this lecture:

with the *loss function*, which specifies the penalty for guessing wrong. Let $\hat{\vec{x}}$ be your estimate of the parameters, $\vec{x}$. Then $L(\hat{\vec{x}}, \vec{x})$ is the loss incurred by guessing $\hat{\vec{x}}$ when the true value was $\vec{x}$. With the posterior probability, we can calculate the expected loss, $\bar{L}(\hat{\vec{x}}, \vec{x})$

$$\bar{L}(\hat{\vec{x}}, \vec{x}) = \int_{\vec{x}} L(\hat{\vec{x}}, \vec{x}) P(\vec{x}|\vec{y}) \tag{6.20}$$

We often use a loss function which is only a function of $\hat{\vec{x}} - \vec{x}$.

Two loss functions (left), and the (minus) expected losses for the 1=ab problem

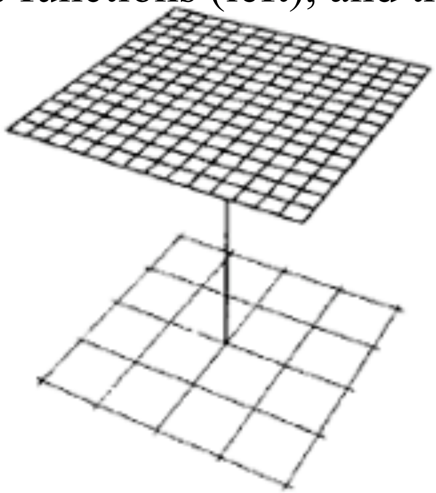

(a) MAP loss function

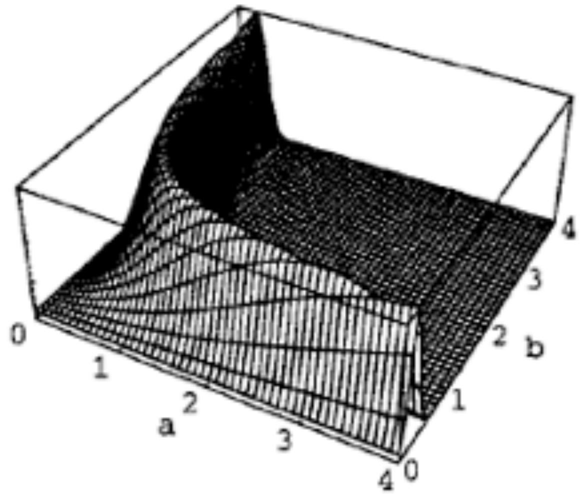

(d) (minus) MAP expected loss

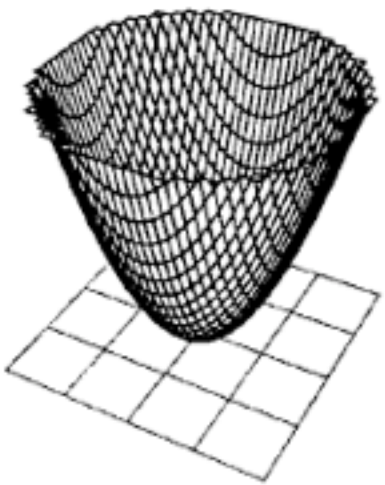

(b) MMSE loss function

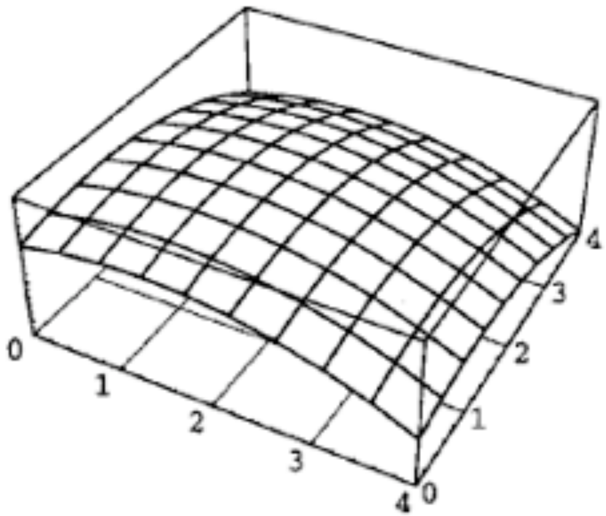

(e) (minus) MMSE expected loss

MAP estimate of illumination spectrum

- Start from some illuminant candidate.
- Find the surface colors that would best explain the observed data.
  - Evaluate the corresponding likelihoold and prior probability terms.
- Move to another illuminant choice.

MMSE estimate of illumination spectrum

For the MMSE estimate, we will use a Monte Carlo method (averaging many different trials). We will take many random draws of candidate illuminant spectra, nd the corresponding surface colors that would explain the observed image data, and then check how probable that set of surface colors would be. We'll use that probability as a weight to form a weighted average of the sampled illumination spectra, which will be the MMSE estimate.

appendix

- supplemental slides about the CIE color space, and spatial resolution and color.

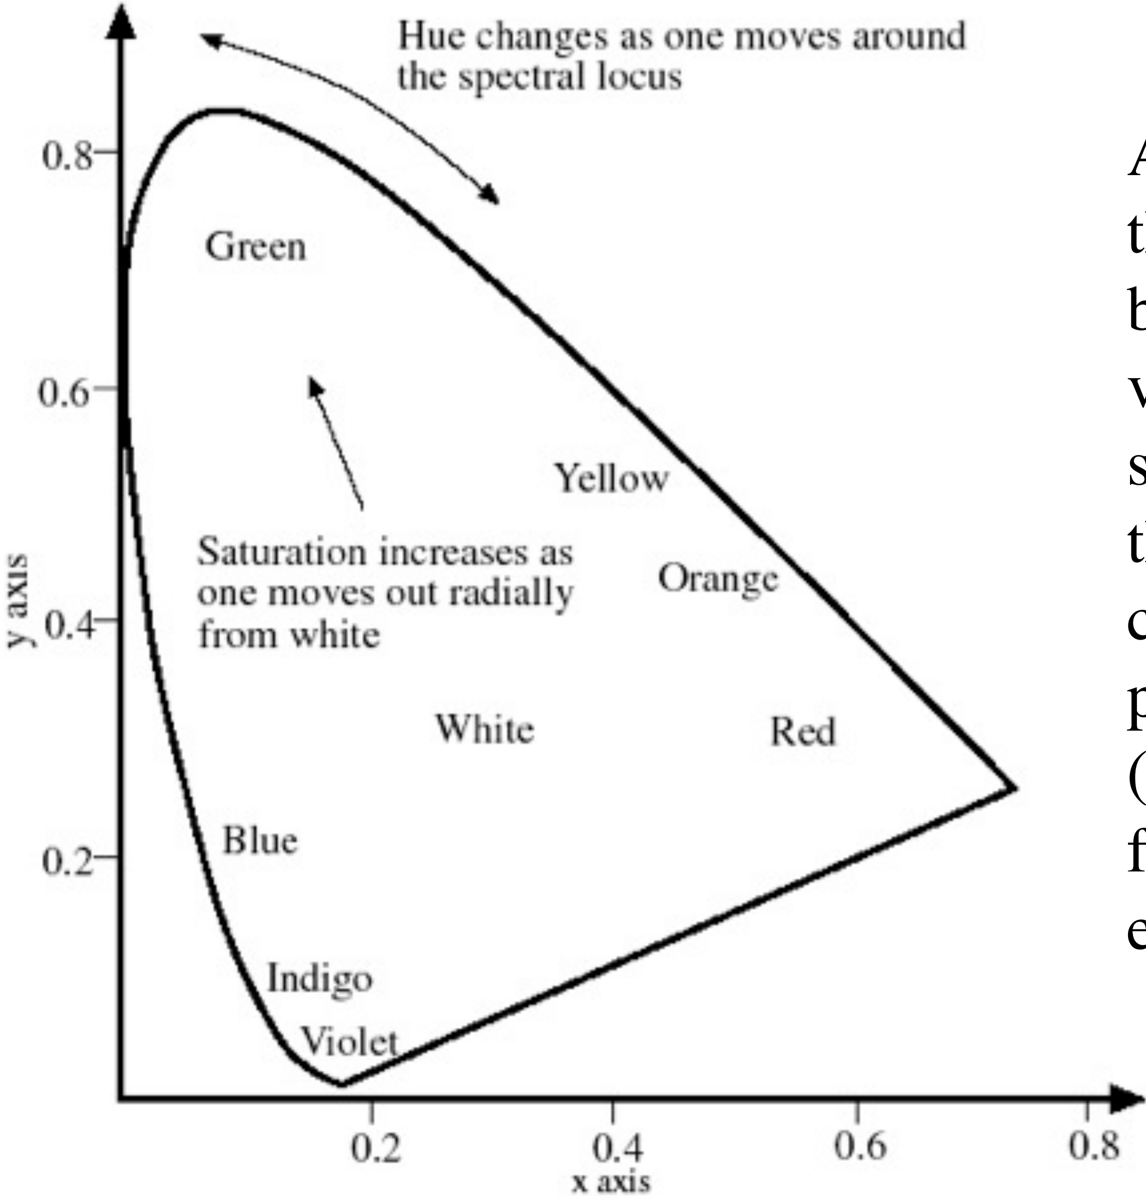

A qualitative rendering of the CIE (x,y) space. The blobby region represents visible colors. There are sets of (x, y) coordinates that don't represent real colors, because the primaries are not real lights (so that the color matching functions could be positive everywhere).

Forsyth & Ponce

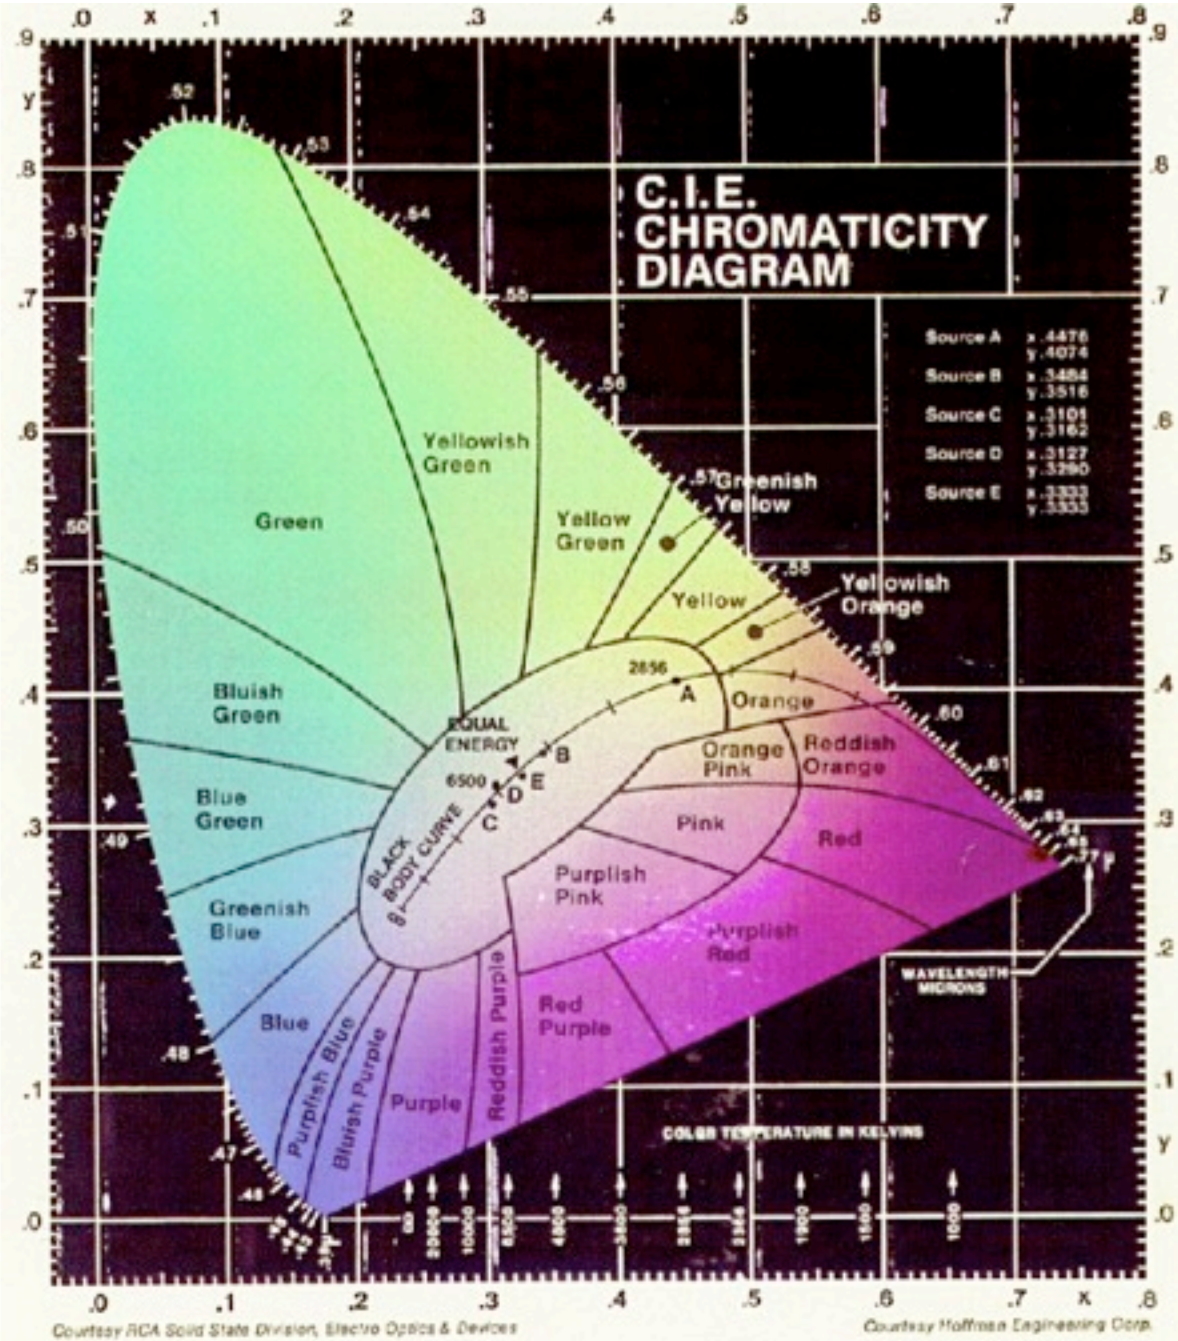
C.I.E.
CHROMATICITY
DIAGRAM
Source A x .4476 y .4074
Source B x .3484 y .3516
Source C x .3101 y .3162
Source D x .3127 y .3290
Source E x .3333 y .3333
Green
Yellowish Green
Yellow Green
Greenish Yellow
Yellow
Yellowish Orange
Orange
Orange Pink
Reddish Orange
Pink
Red
Purplish Pink
Purplish Red
Red Purple
Reddish Purple
Purple
Bluish Purple
Purplish Blue
Blue
Greenish Blue
Blue Green
Bluish Green
EQUAL ENERGY
BLACK BODY CURVE
6500
2856
A
B
C
D
E
WAVELENGTH MICRONS
COLOR TEMPERATURE IN KELVINS
Courtesy RCA Solid State Division, Electro Optics & Devices
Courtesy Hoffman Engineering Corp.

Pure wavelength in chromaticity diagram

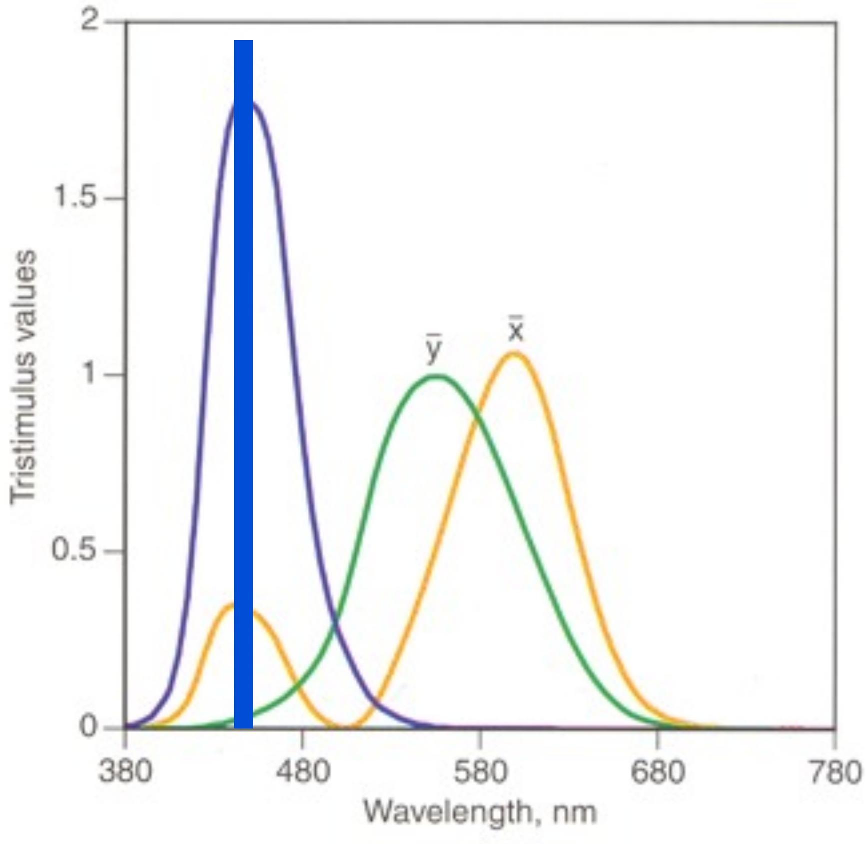

The 1931 standard observer, as it is usually shown.

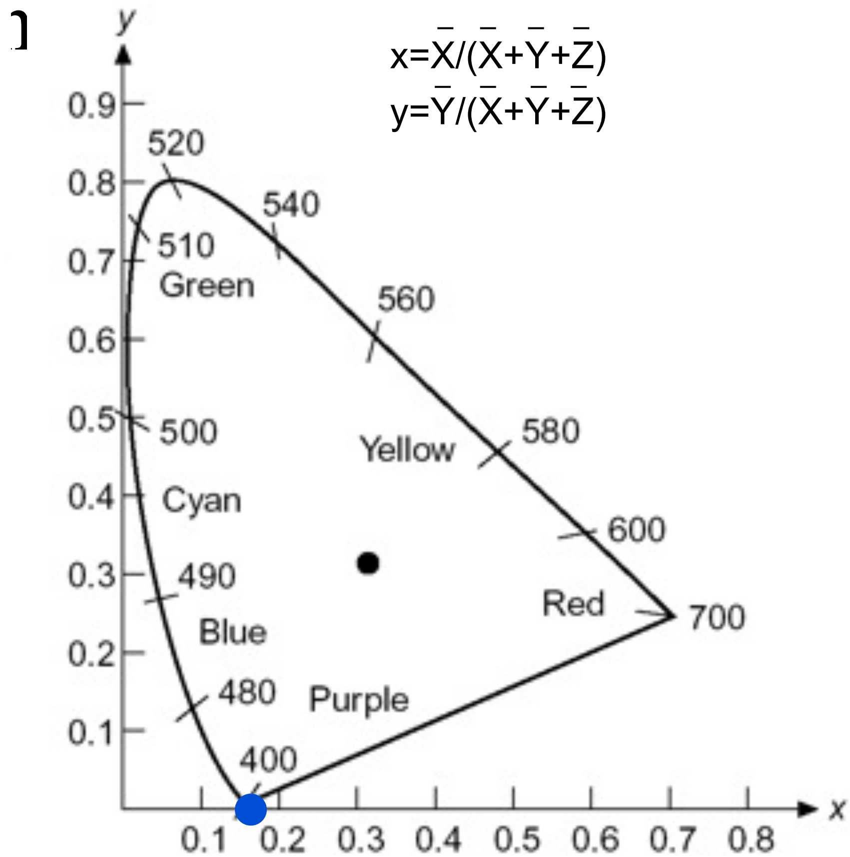

$x=\bar{X}/(\bar{X}+\bar{Y}+\bar{Z})$

$y=\bar{Y}/(\bar{X}+\bar{Y}+\bar{Z})$

Pure wavelength in chromaticity diagram

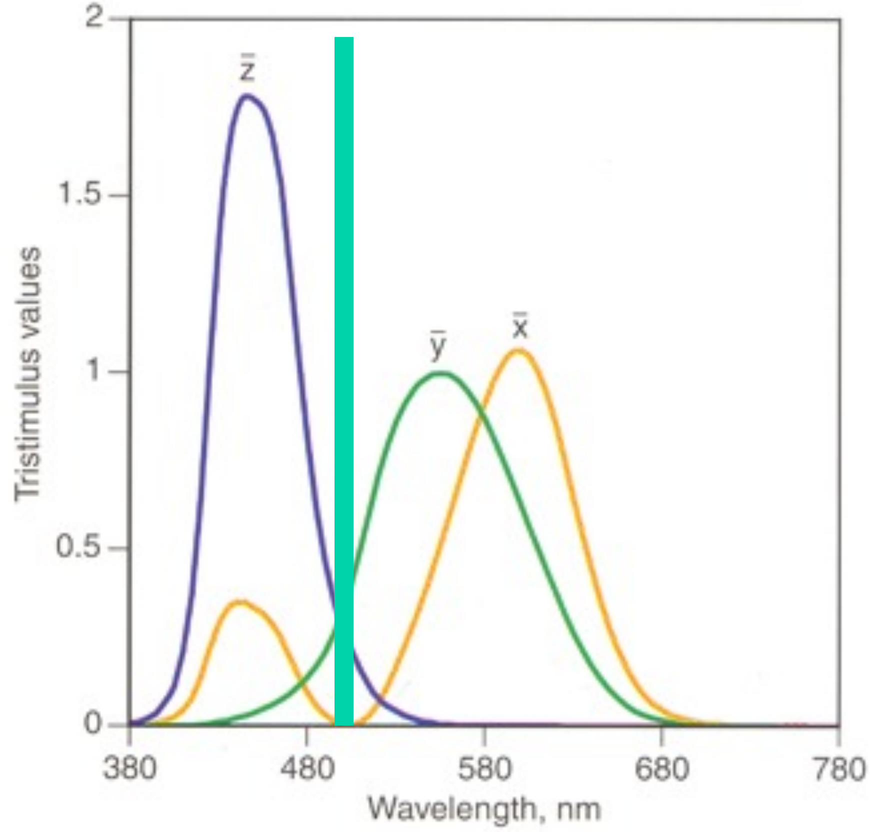

The 1931 standard observer, as it is usually shown.

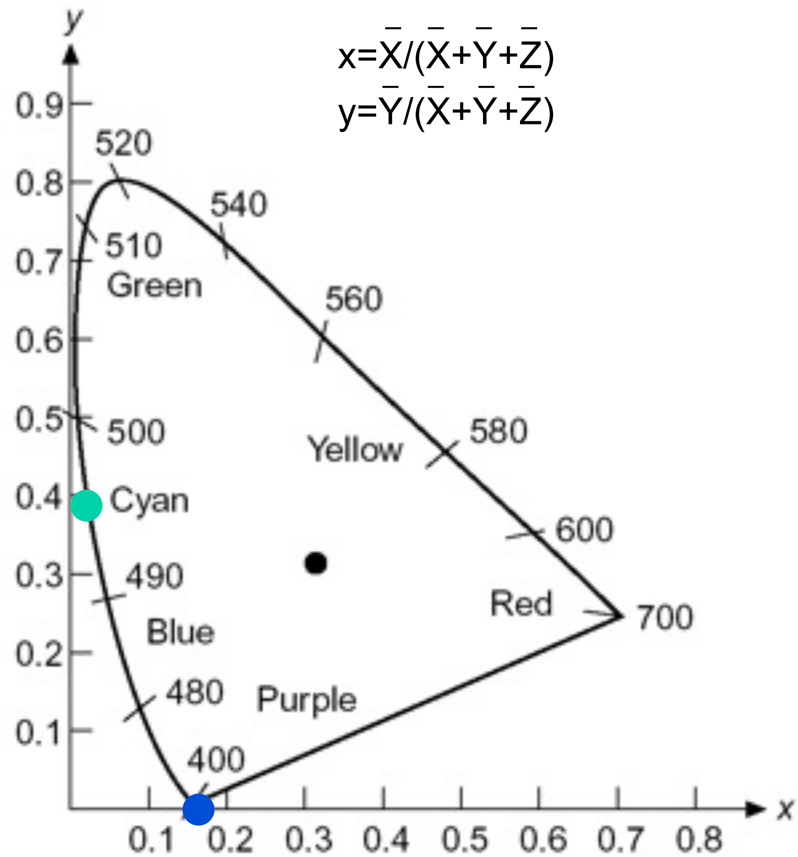

$x=\bar{X}/(\bar{X}+\bar{Y}+\bar{Z})$
$y=\bar{Y}/(\bar{X}+\bar{Y}+\bar{Z})$

Pure wavelength in chromaticity diagram

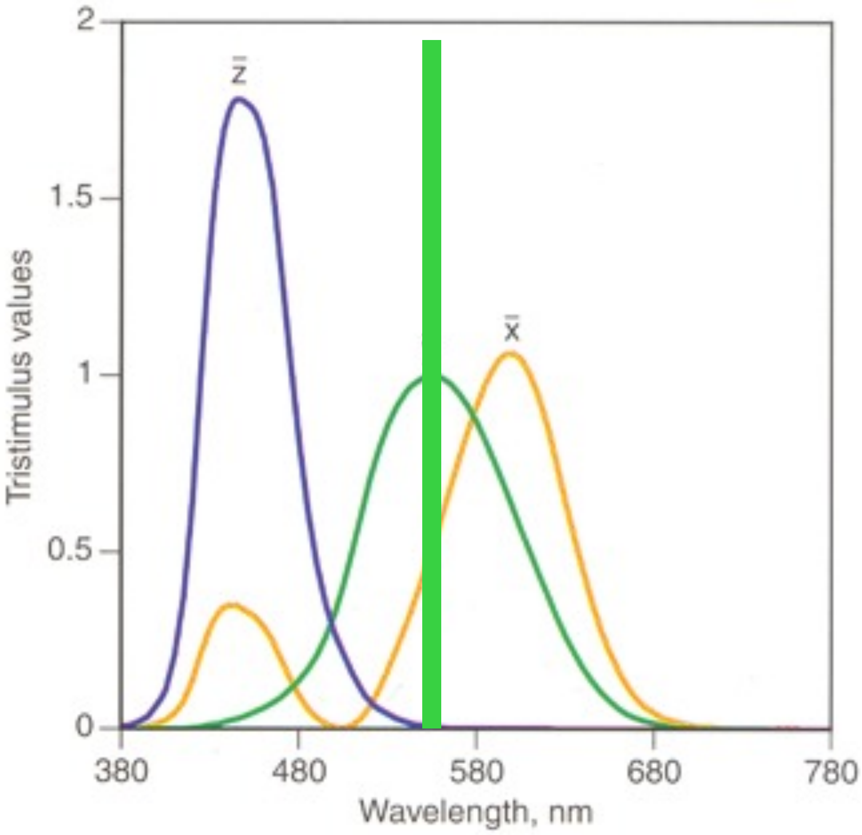

The 1931 standard observer, as it is usually shown.

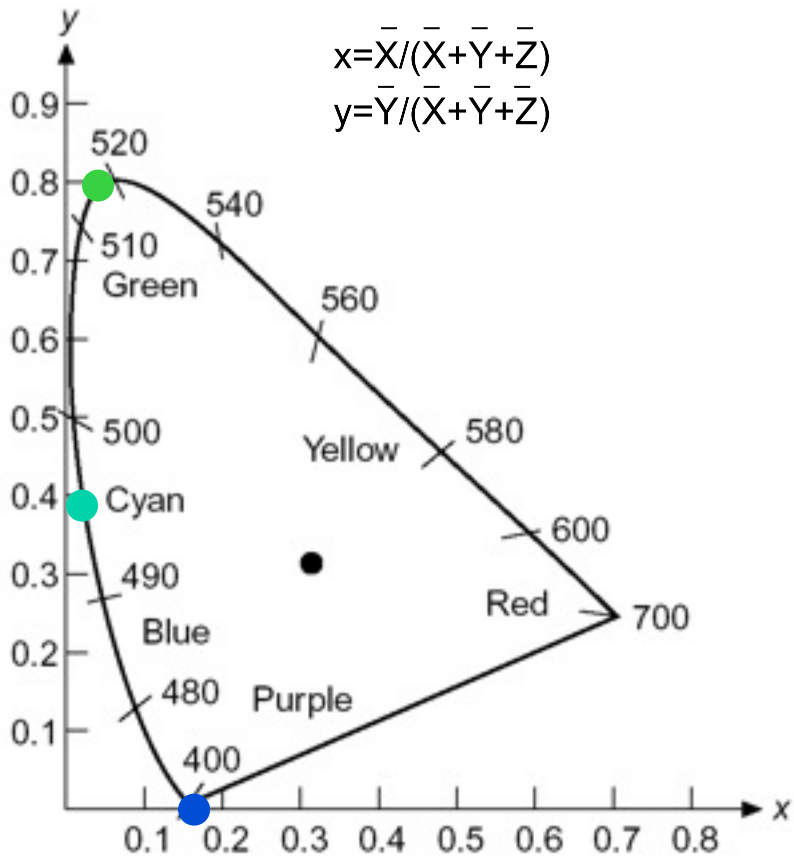

$x = \bar{X}/(\bar{X}+\bar{Y}+\bar{Z})$

$y = \bar{Y}/(\bar{X}+\bar{Y}+\bar{Z})$

Pure wavelength in chromaticity diagram

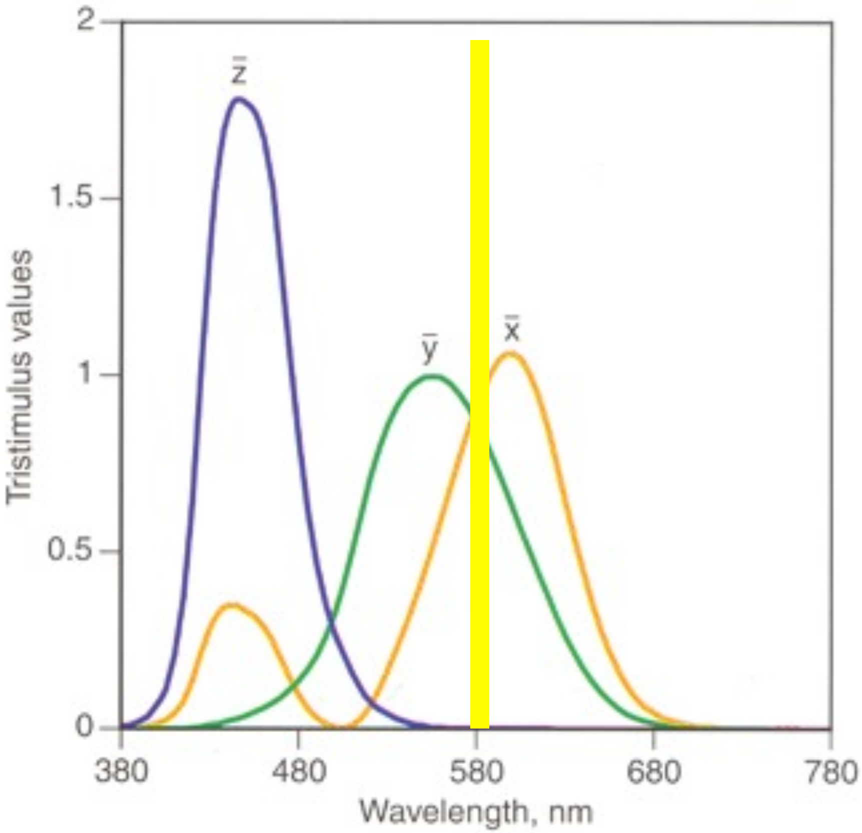

The 1931 standard observer, as it is usually shown.

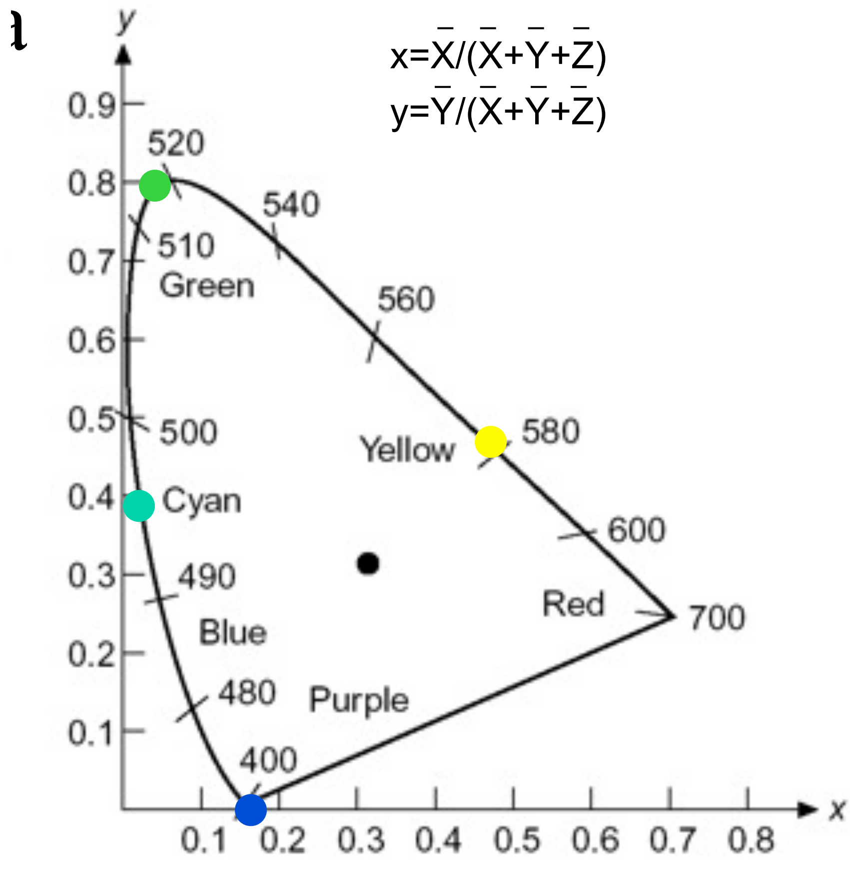

Pure wavelength in chromaticity diagram

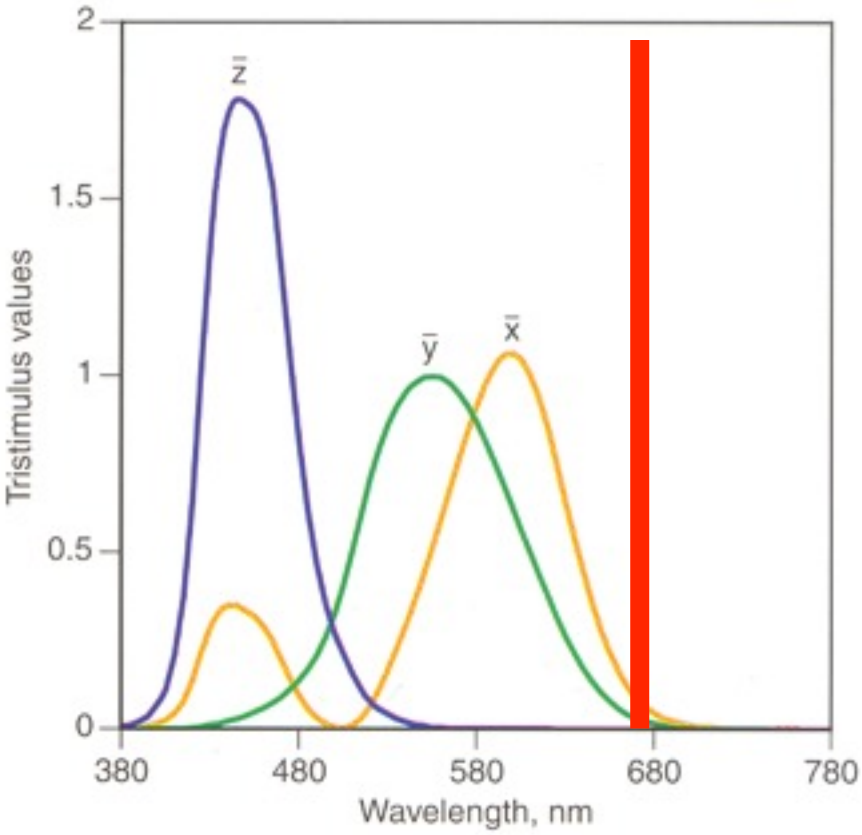

The 1931 standard observer, as it is usually shown.

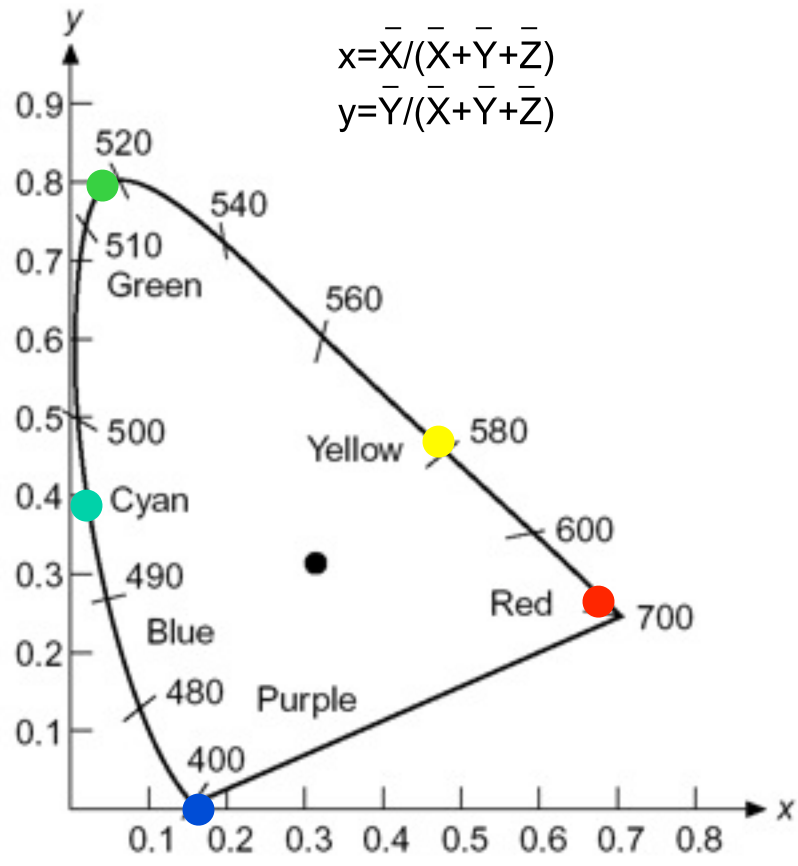

$x=\bar{X}/(\bar{X}+\bar{Y}+\bar{Z})$

$y=\bar{Y}/(\bar{X}+\bar{Y}+\bar{Z})$

CIE chromaticity diagram

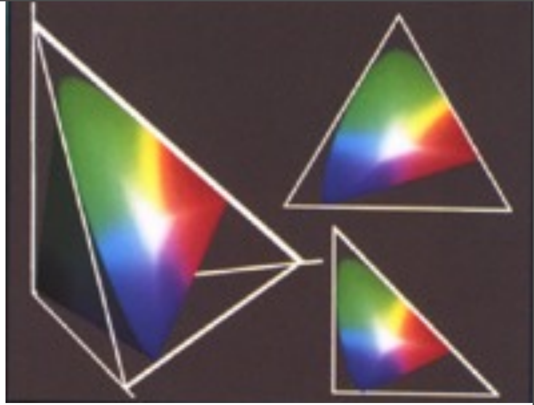

- Spectrally pure colors lie along boundary
- Weird shape comes from shape of matching curves
  and restriction to positive stimuli
- Note that some hues do not correspond to a pure spectrum (purple-violet)
- Standard white light (approximates sunlight) at C

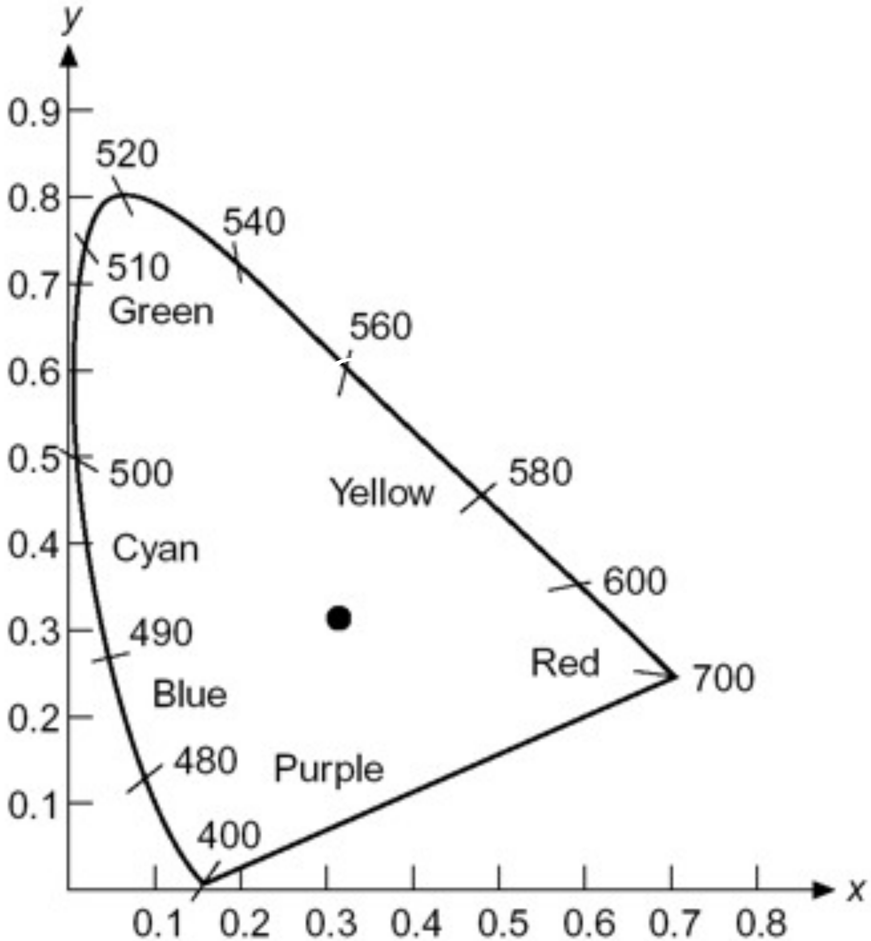

CIE color space

$$\begin{pmatrix} R \\ G \\ B \end{pmatrix} = \begin{pmatrix} 3.24 & -1.54 & -0.50 \\ -0.97 & 1.88 & 0.04 \\ 0.06 & -0.20 & 1.06 \end{pmatrix} \begin{pmatrix} X \\ Y \\ Z \end{pmatrix}$$

$$\begin{pmatrix} X \\ Y \\ Z \end{pmatrix} = \begin{pmatrix} 0.41 & 0.36 & 0.18 \\ 0.21 & 0.72 & 0.07 \\ 0.02 & 0.12 & 0.95 \end{pmatrix} \begin{pmatrix} R \\ G \\ B \end{pmatrix}$$

- Can think of X, Y , Z as coordinates
- Linear transform from typical RGB or LMS
- Always positive (because physical spectrum is positive and matching curves are positives)
- Note that many points in XYZ do not correspond to visible

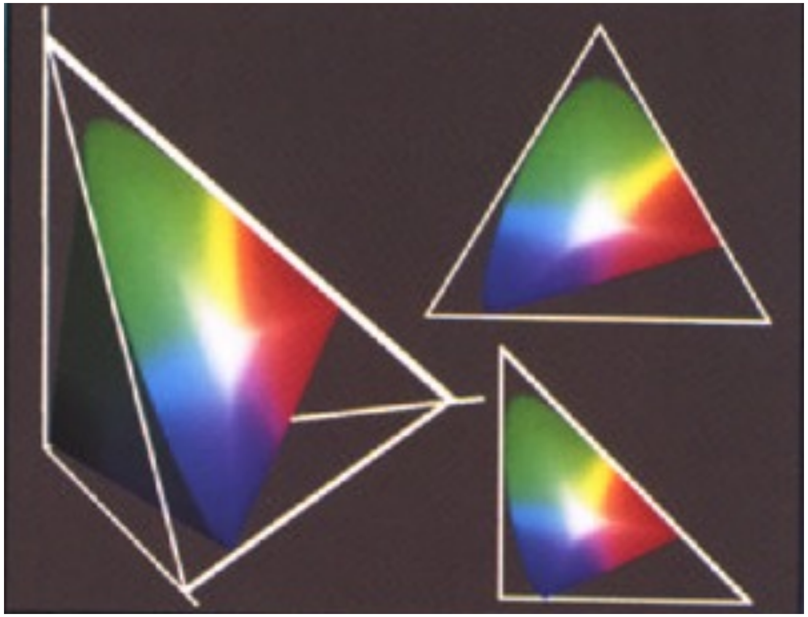

Another psychophysical fact: luminance and chrominance channels in the brain

From W. E. Glenn, in Digital Images and Human Vision, MIT Press, edited by Watson, 1993

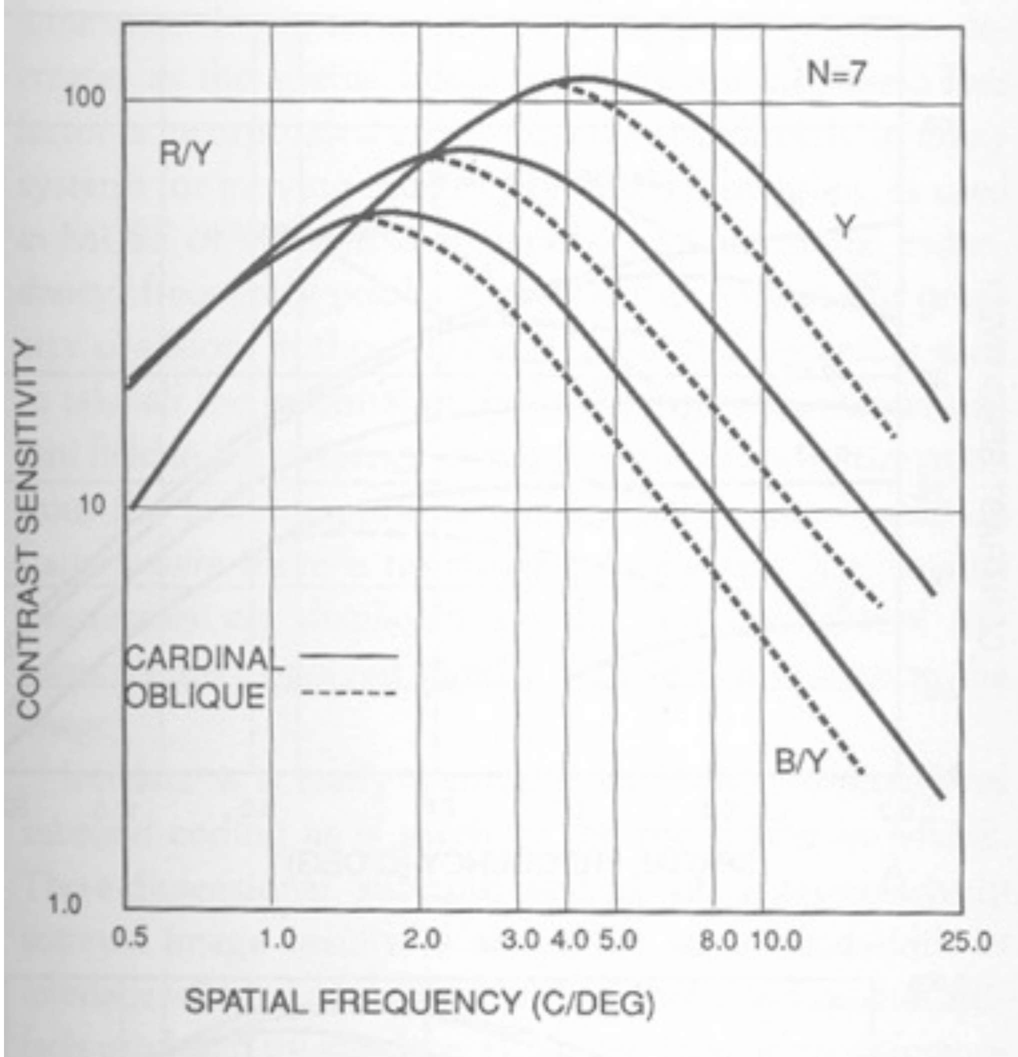

Figure 6.1
Contrast sensitivity threshold functions for static luminance gratings (Y) and isoluminance chromaticity gratings (R/Y,B/Y) averaged over seven observers.

NTSC color components: Y, I, Q

$$\begin{pmatrix} Y \\ I \\ Q \end{pmatrix} = \begin{pmatrix} 0.299 & 0.587 & 0.114 \\ 0.596 & -0.274 & -0.322 \\ 0.211 & -0.523 & 0.312 \end{pmatrix} \begin{pmatrix} R \\ G \\ B \end{pmatrix}$$

NTSC - RGB

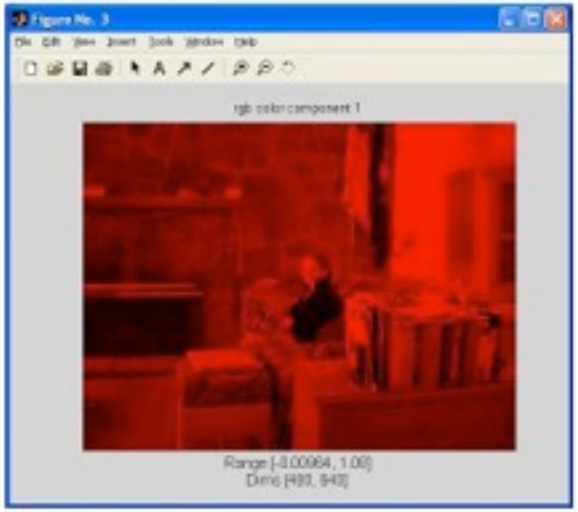
Figure No. 3
rgb color component 1
Dims [480, 640]

R

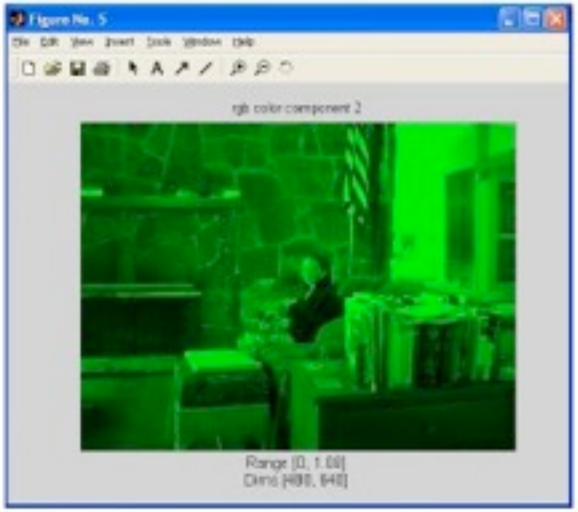
Figure No. 5
rgb color component 2
Range [0, 1.08]
Dims [480, 640]

G

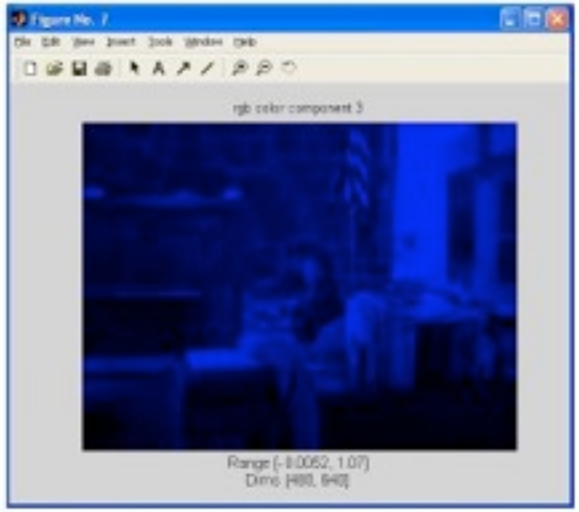
Figure No. 7
rgb color component 3
Dims [480, 640]

B

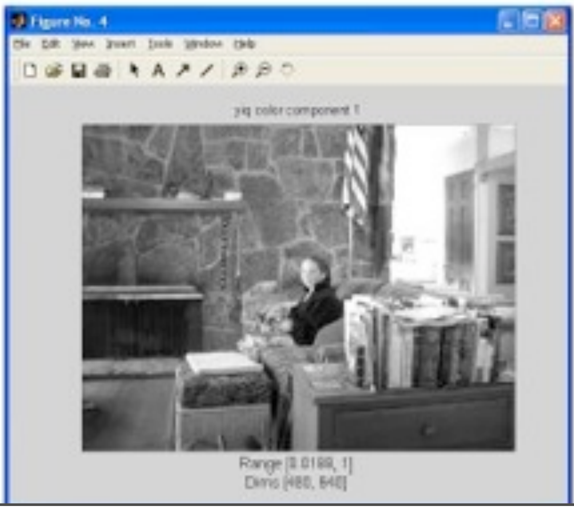
Figure No. 4
yiq color component 1
Dims [480, 640]

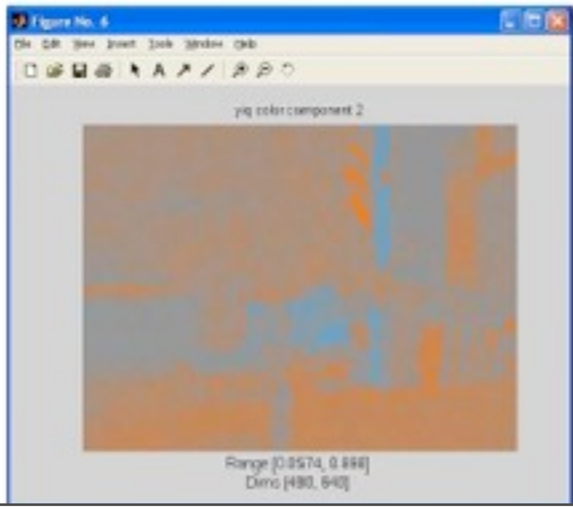
Figure No. 6
yiq color component 2
Dims [480, 640]

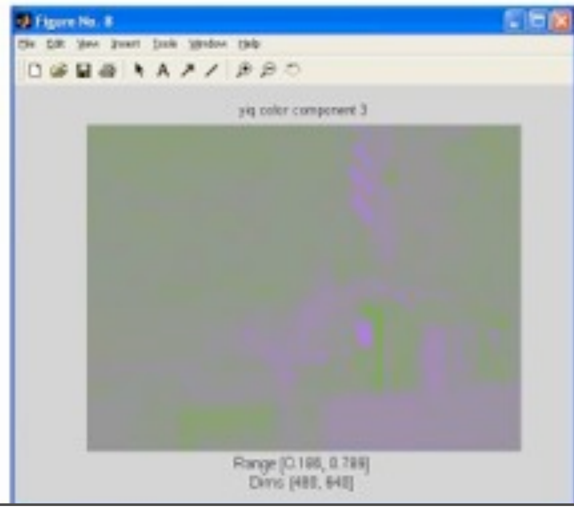
Figure No. 8
yiq color component 3
Dims [480, 640]

Spatial resolution and color

original

R

G

B

Blurring the G component

original

processed

R

G

B

Blurring the G component

original

processed

R

G

B

Blurring the R component

original

processed

R

G

B

Blurring the R component

original

processed

R

G

B

Blurring the B component

original

R

G

B

Blurring the B component

original

processed

R

G

B

From W. E. Glenn, in Digital Images and Human Vision, MIT Press, edited by Watson, 1993

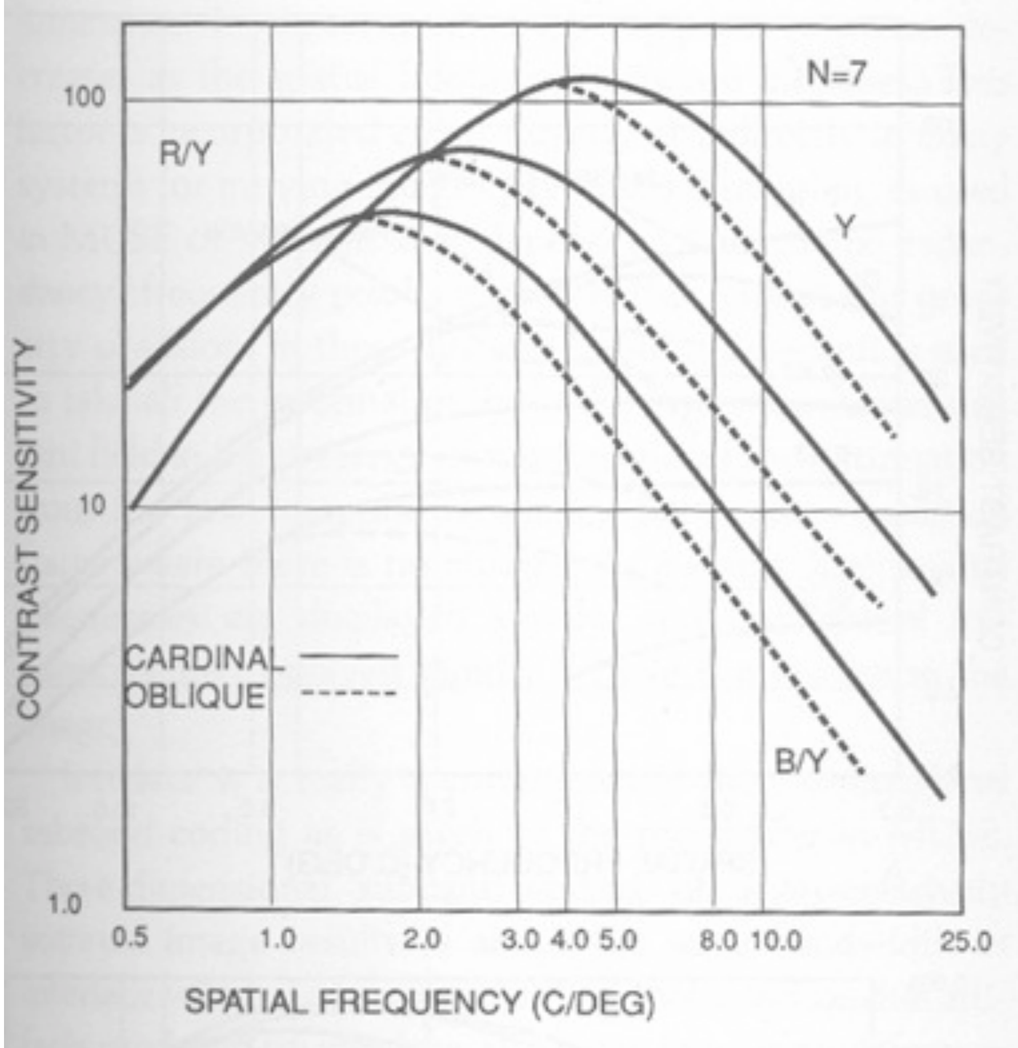

Figure 6.1
Contrast sensitivity threshold functions for static luminance gratings (Y) and isoluminance chromaticity gratings (R/Y,B/Y) averaged over seven observers.

Lab color components

L

a

b

A rotation of the color coordinates into directions that are more perceptually meaningful:
L: luminance,
a: red-green,
b: blue-yellow

Blurring the L Lab component

original

L

a

b

Blurring the L Lab component

original

processed

L

a

b

Blurring the a Lab component

original

L

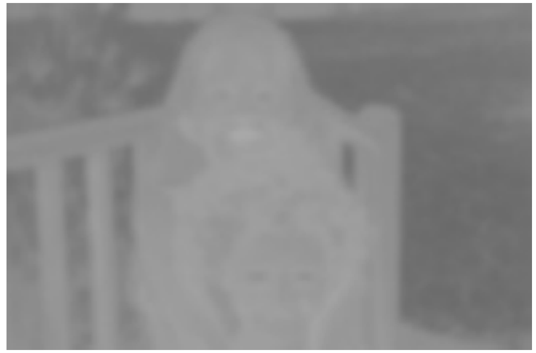

a

b

Blurring the a Lab component

original

processed

L

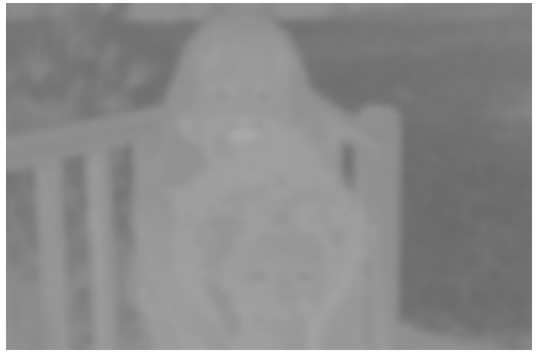

a

b

Blurring the b Lab component

original

L

a

b

Blurring the b Lab component

original

processed

L

a

b

Class project idea 2: time-lapse photography temporal color filtering

- Some colors change slowly over time and we can't easily perceive those long-term changes.
- Take photographs over time of imagery you want to analyze, and include a color calibration card in the scene.
- From the measurements over the card, you can pull out the illumination spectrum for each photo, and show each image as if they were all taken under the same illumination.
- Then color differences between images should correspond to true surface color changes. Temporally filter the color-balanced time-lapse imagery to accentuate the color changes of your subject over time. This will give you a color magnifying glass to exaggerate color changes over time.



Selected Bibliography

Vision Science
by Stephen E. Palmer
MIT Press; ISBN: 0262161834
760 pages (May 7, 1999)

Billmeyer and Saltzman's Principles of Color Technology, 3rd Edition
by Roy S. Berns, Fred W. Billmeyer, Max Saltzman
Wiley-Interscience; ISBN: 047119459X
304 pages 3 edition (March 31, 2000)

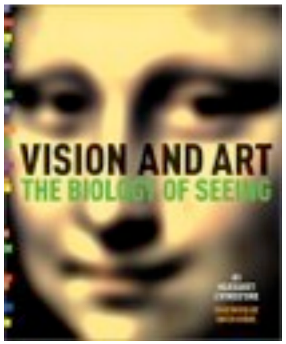

Vision and Art : The Biology of Seeing
by Margaret Livingstone, David H. Hubel
Harry N Abrams; ISBN: 0810904063
208 pages (May 2002)

Selected Bibliography

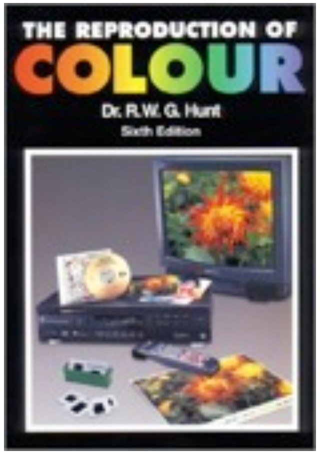

The Reproduction of Color
by R. W. G. Hunt
Fountain Press, 1995

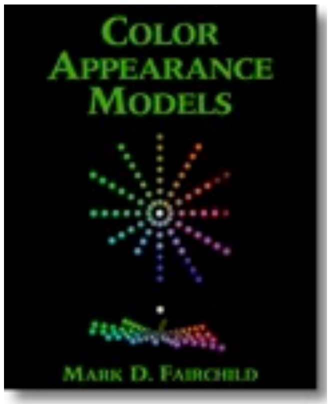

Color Appearance Models
by Mark Fairchild
Addison Wesley, 1998

Other color references

- Reading:
  - Chapter 6, Forsyth & Ponce
  - Chapter 4 of Wandell, Foundations of Vision, Sinauer, 1995 has a good treatment of this.