Light and color

370: Intro to Computer Vision

Subhransu Maji

Feb 4 & 6, 2025

College of INFORMATION AND COMPUTER SCIENCES

UMASS AMHERST

XXXXXX	GREEN	GREEN
XXXXXX	BLUE	BLUE
XXXXXX	YELLOW	YELLOW
XXXXXX	PURPLE	PURPLE
XXXXXX	ORANGE	ORANGE
XXXXXX	RED	RED
XXXXXX	WHITE	WHITE
XXXXXX	PURPLE	PURPLE
XXXXXX	ORANGE	ORANGE
XXXXXX	BLUE	BLUE
XXXXXX	RED	RED
XXXXXX	GREEN	GREEN
XXXXXX	WHITE	WHITE
XXXXXX	YELLOW	YELLOW
XXXXXX	PURPLE	PURPLE
XXXXXX	RED	RED
XXXXXX	GREEN	GREEN
XXXXXX	BLUE	BLUE

Color is affected by context

Joseph Albers, Interaction of Colors

Color is affected by context

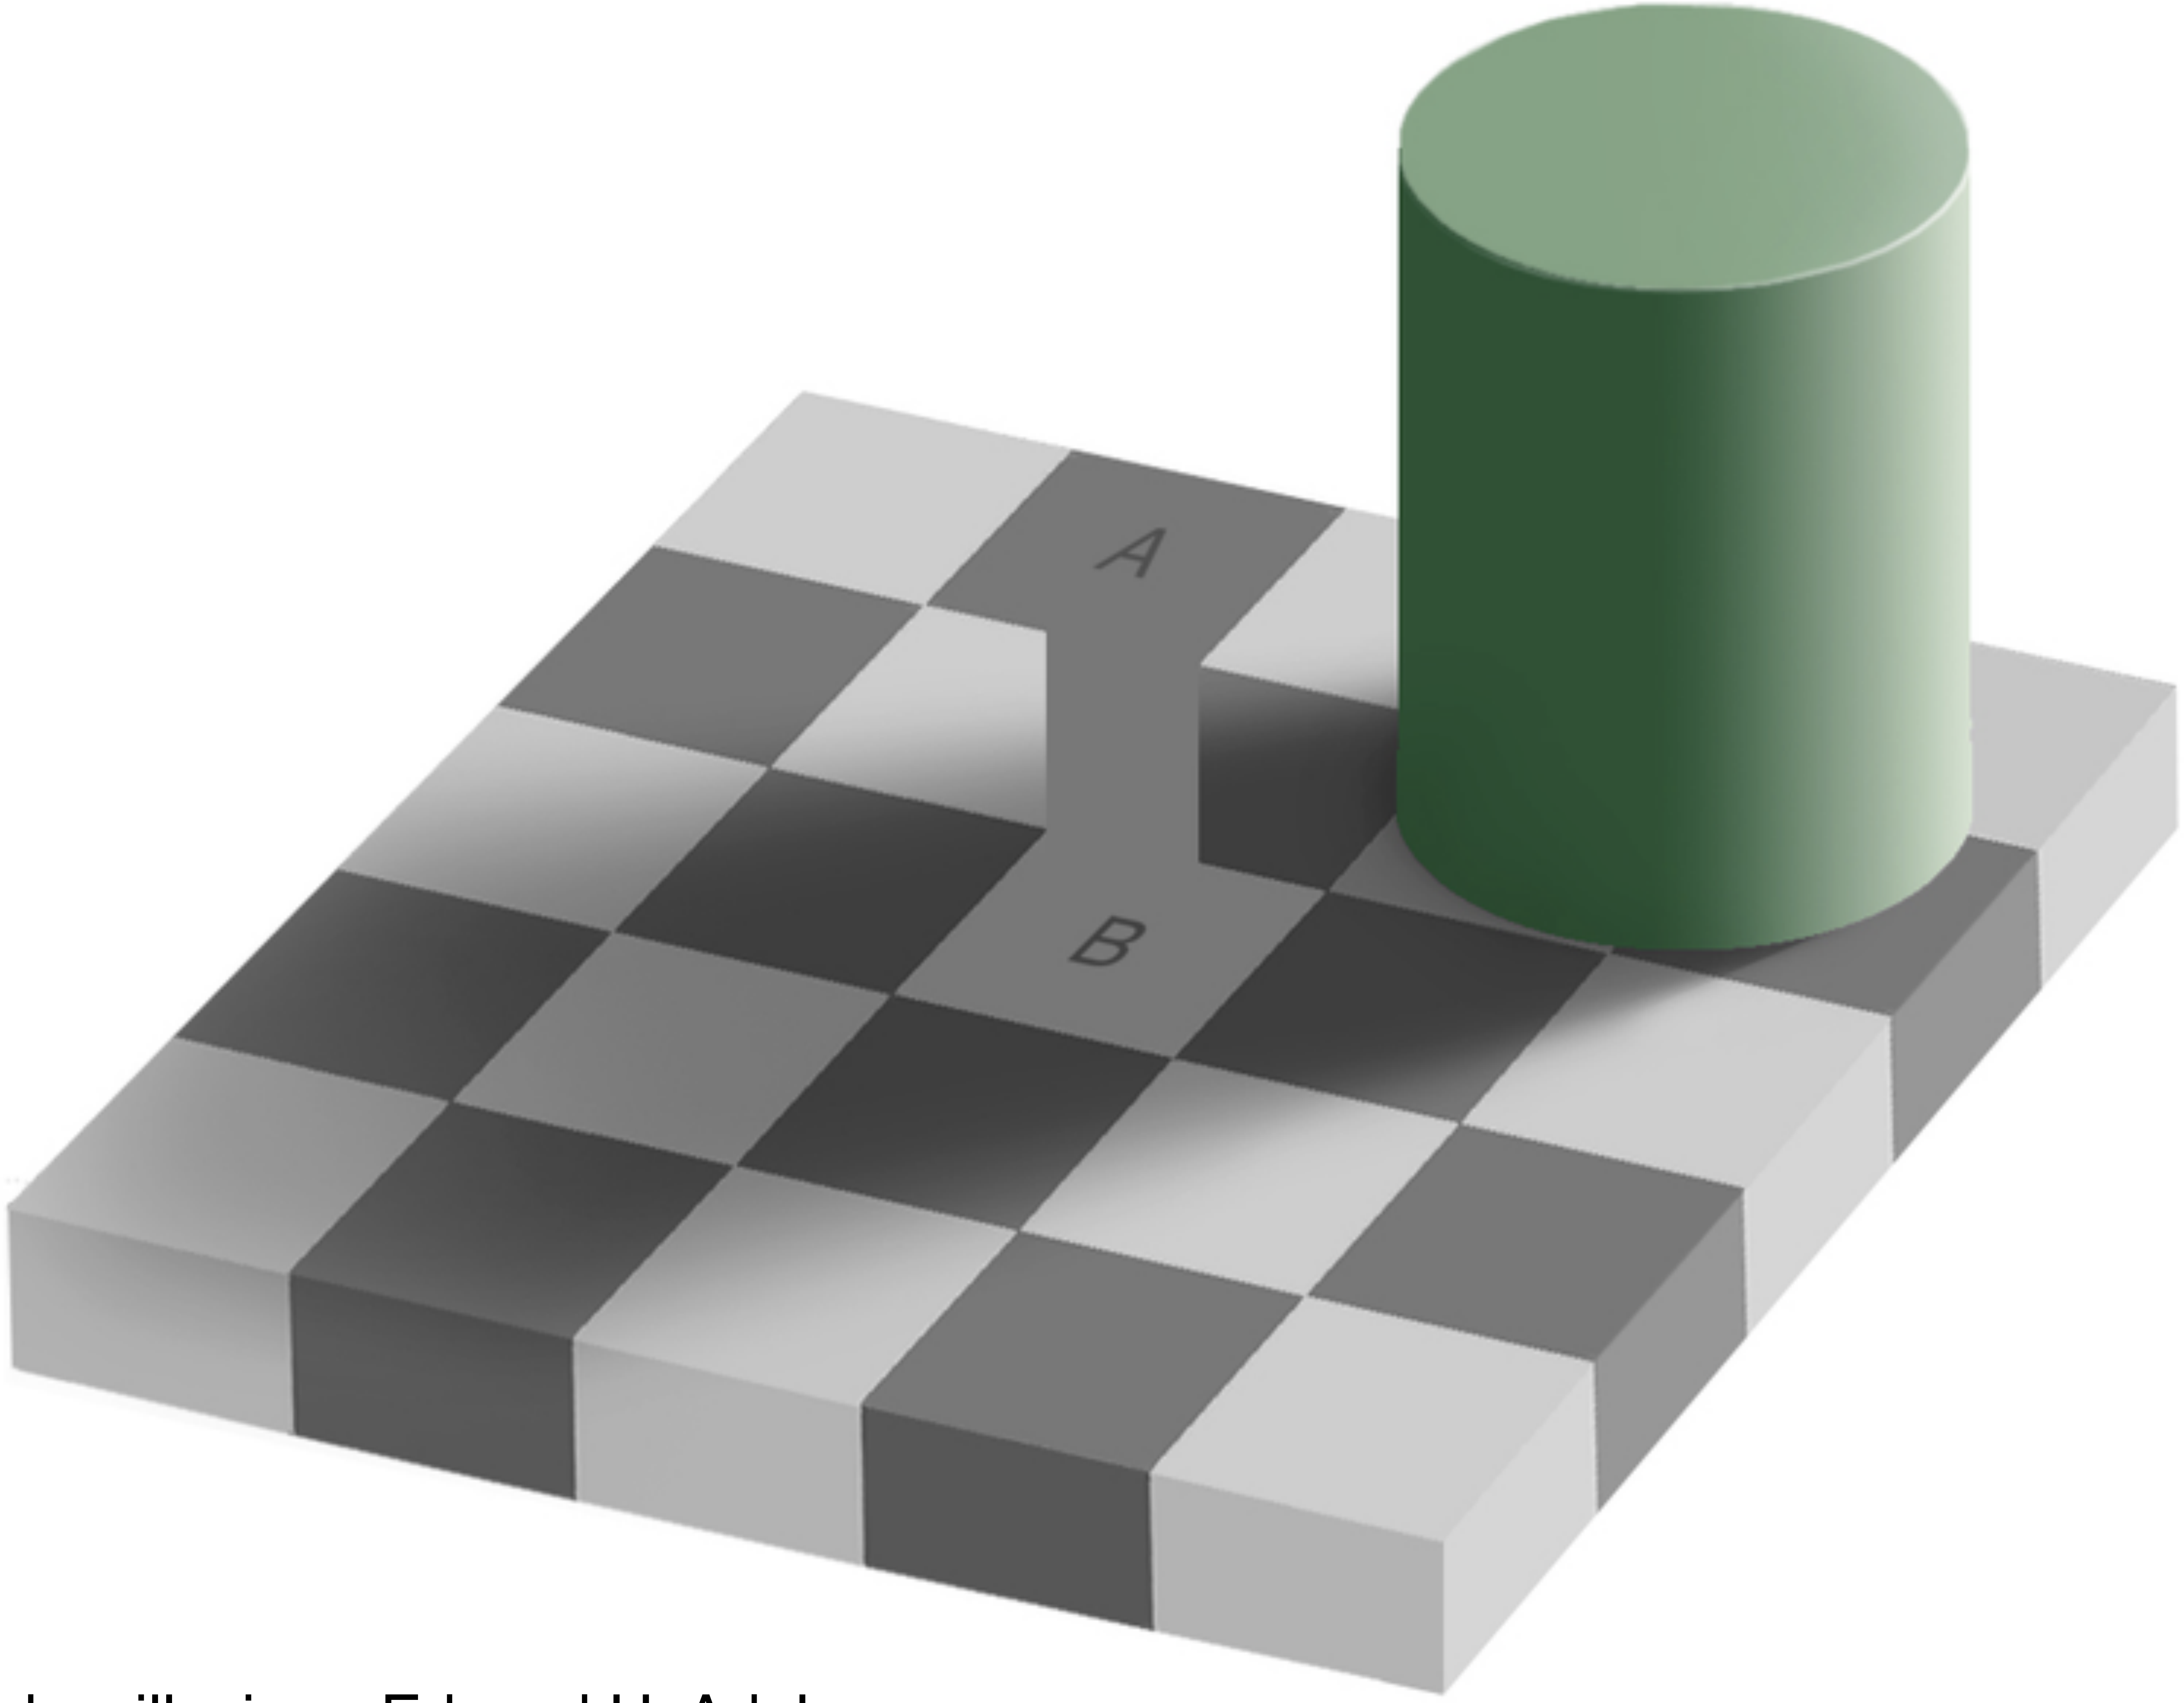

Checker shadow illusion - Edward H. Adelson

What color is this dress?

white and gold
or
blue and black

Let’s use class questions

Enroll in Class Question Before Tomorrow

We will use class question for low-stakes in-class quizzes. Please follow these instructions to sign up before the class tomorrow. Class code **RSTLG**

If you already have a Class Question account, skip to step 2. If you are new, start at step 1.

1. Go to classquestion.com/students and click "Click here to register". This link will allow you to register for the site.
2. Once you have registered, go to classquestion.com/students and sign in.
3. Click "Add Class" at the bottom. Enter the Class Code for this class: **RSTLG** and then click "Add Class".
4. Your class will be added to the dropdown menu at the top. You can now click the "Sign In" button to log into your class!

Overview

Spectral basis of light

Color perception in the human eye

Tristimulus theory and color spaces

Interaction of light and surfaces

Color phenomena

What is color?

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

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

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

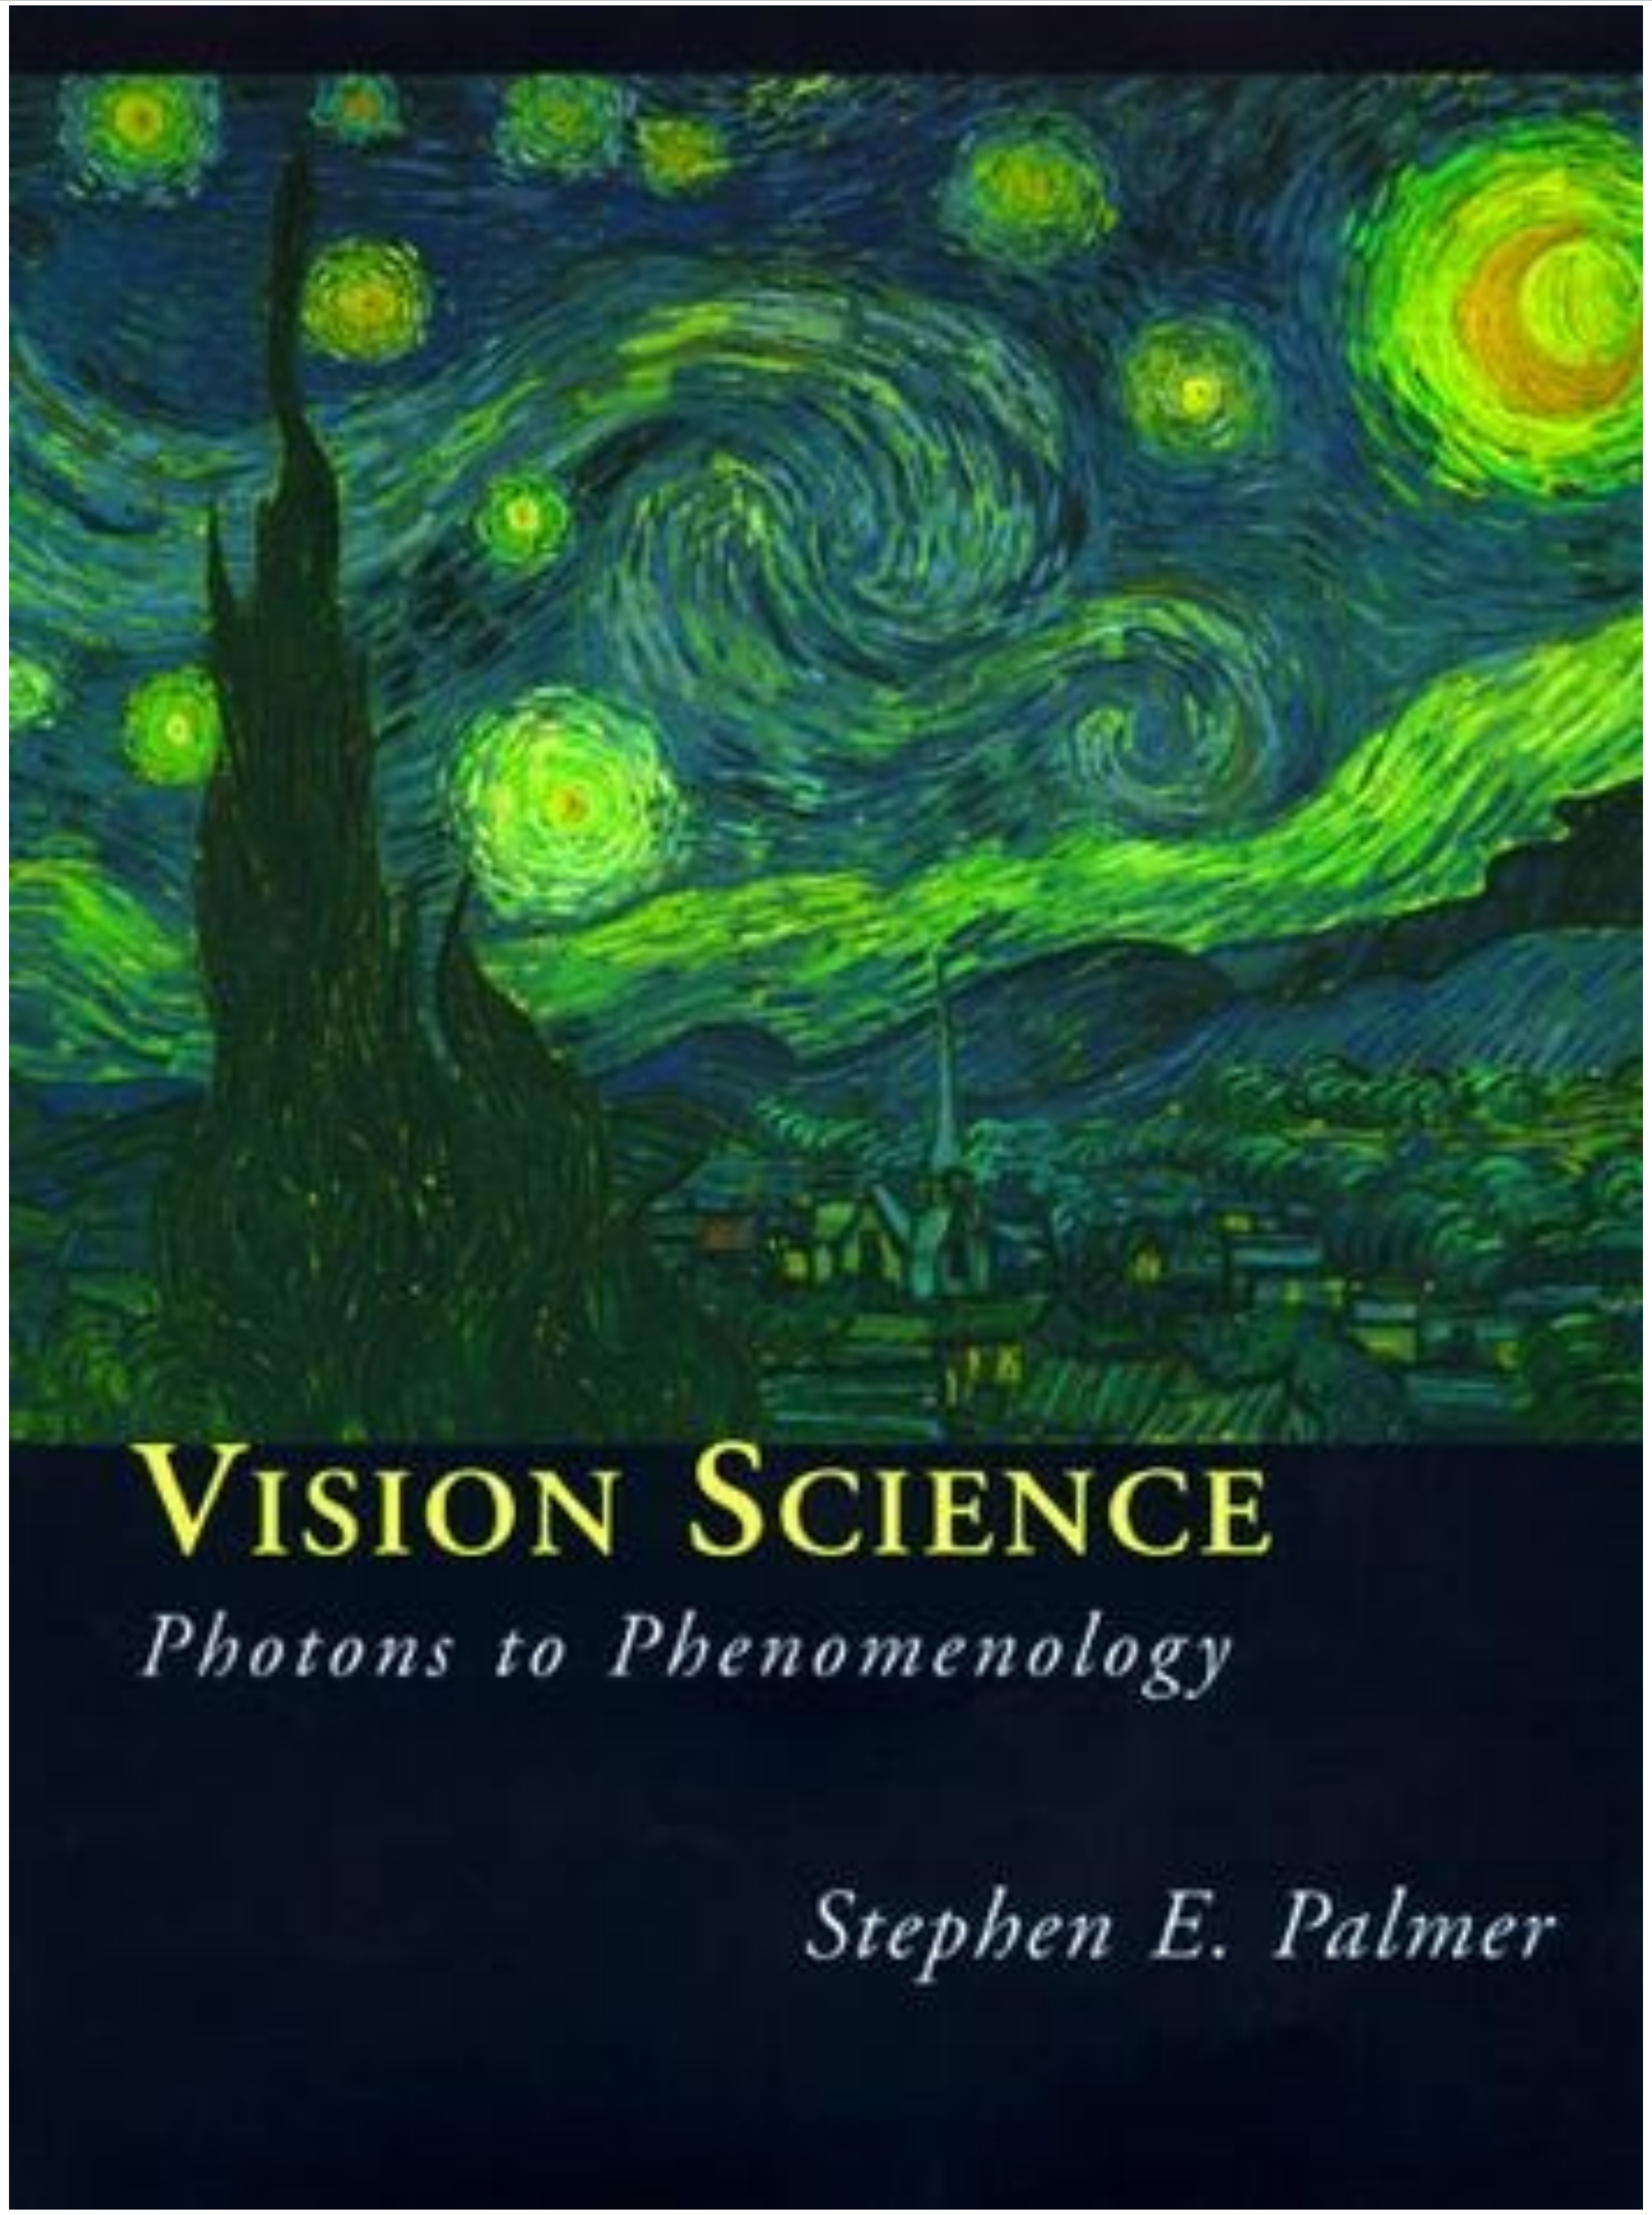

Newton’s theory of light

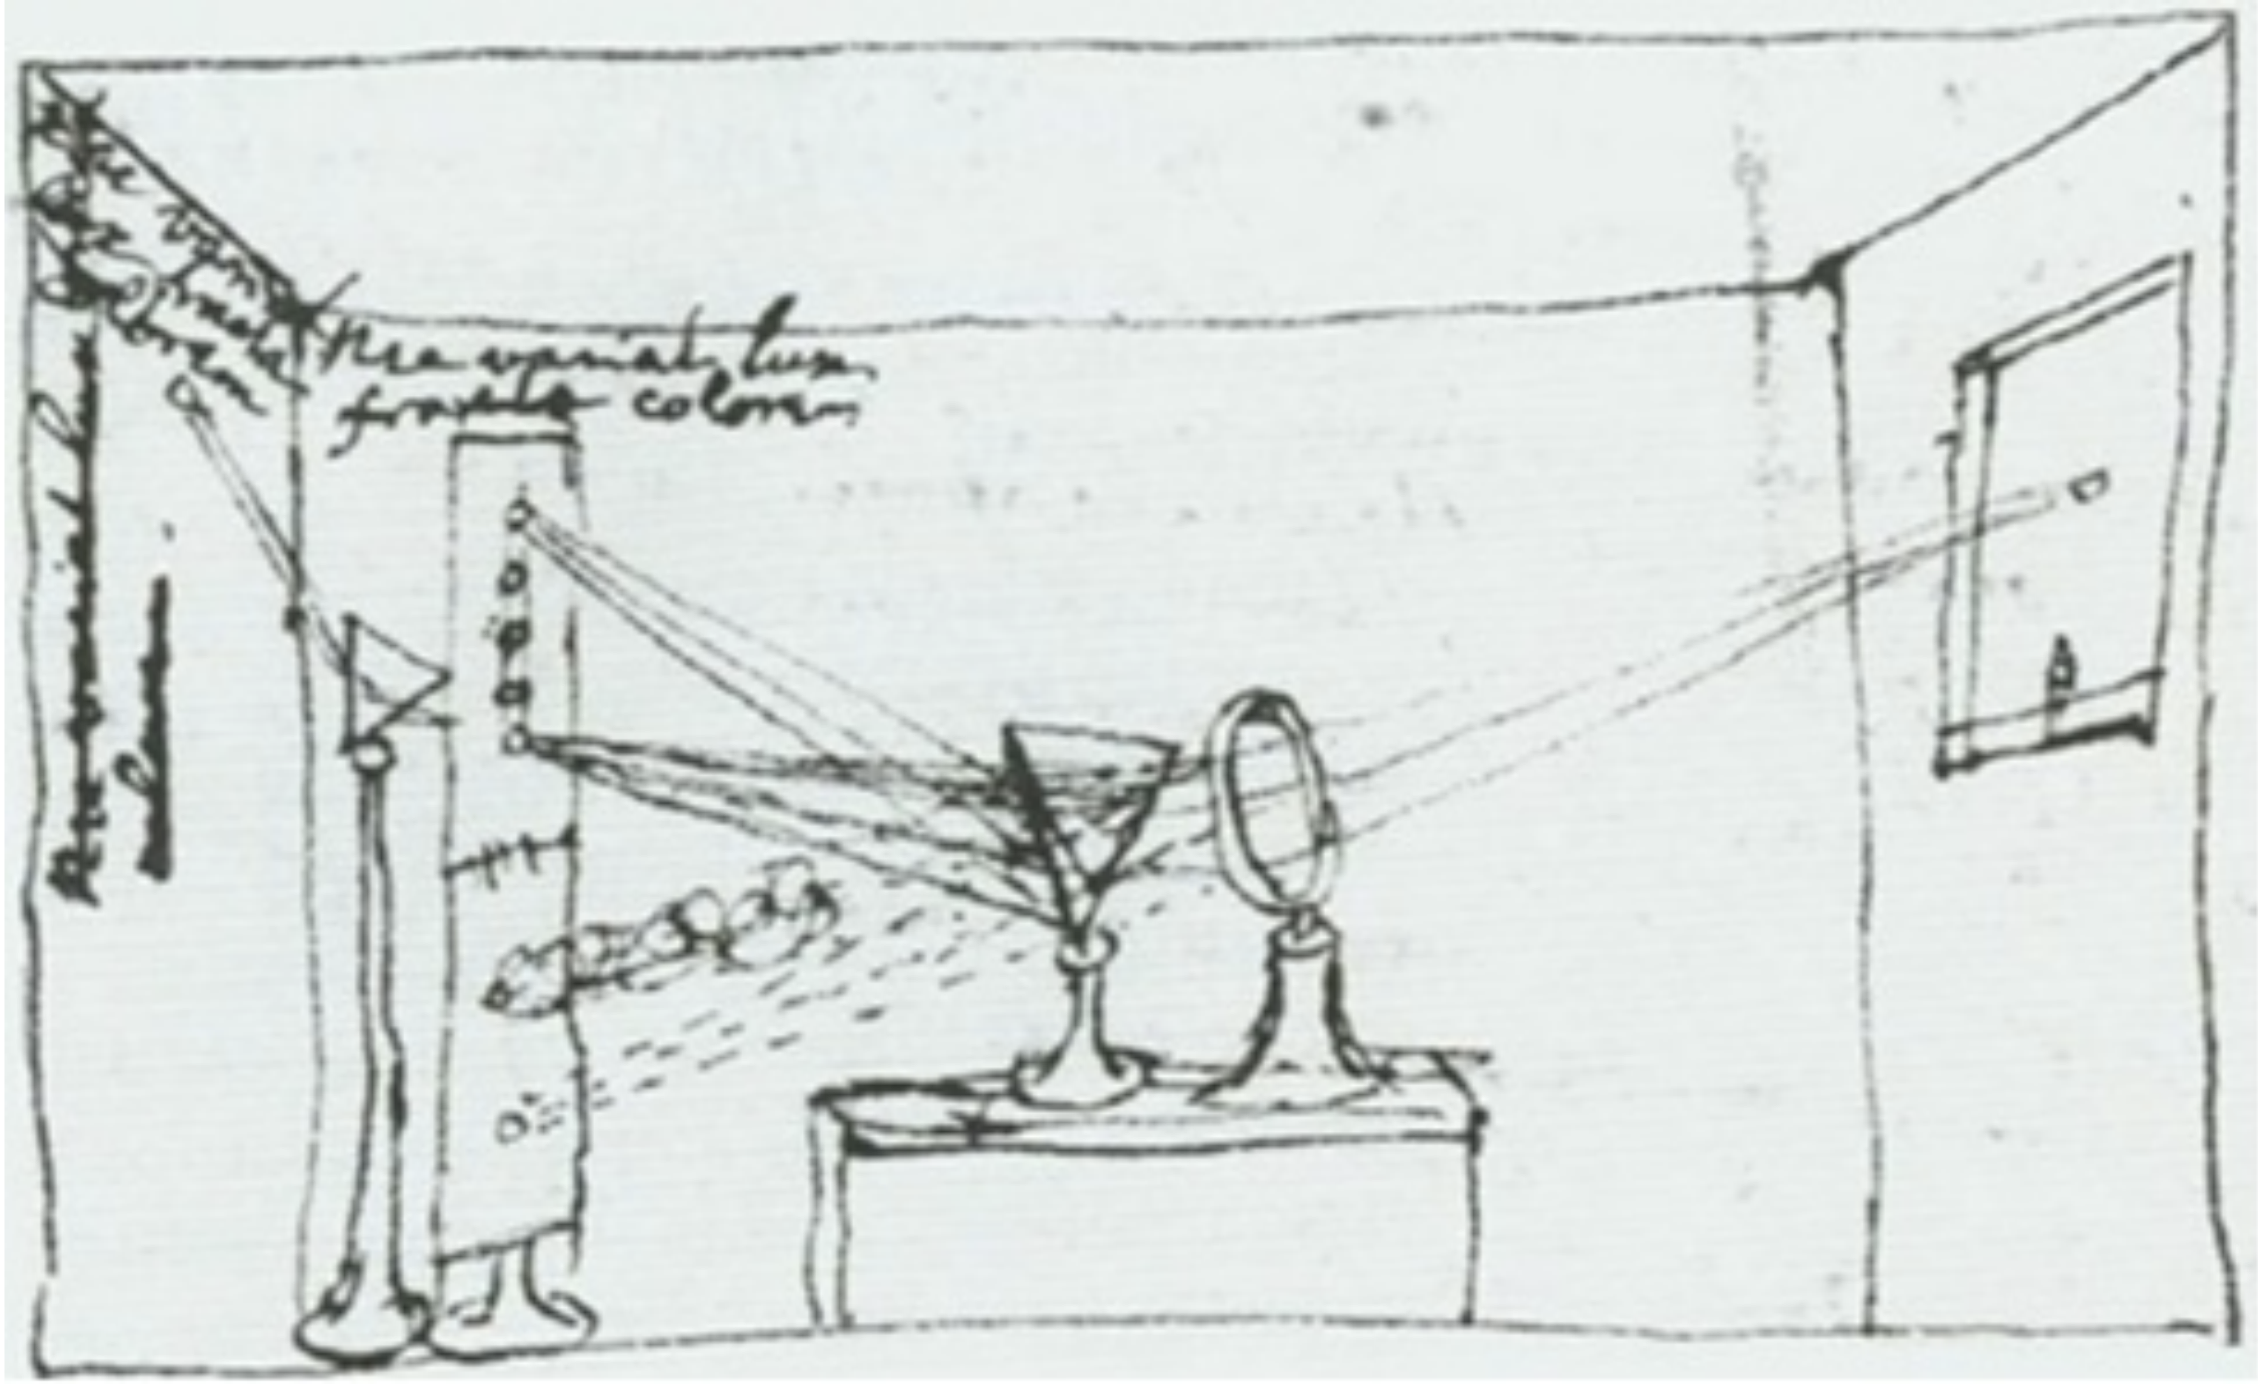

Newton's sketch of his crucial experiment in which light from the sun is refracted through a prism. One color is then refracted through a second prism to show that it undergoes no further change. Light is then shown to be composed of the colors refracted in the second prisms.

Image credit: **Warden and Fellows**

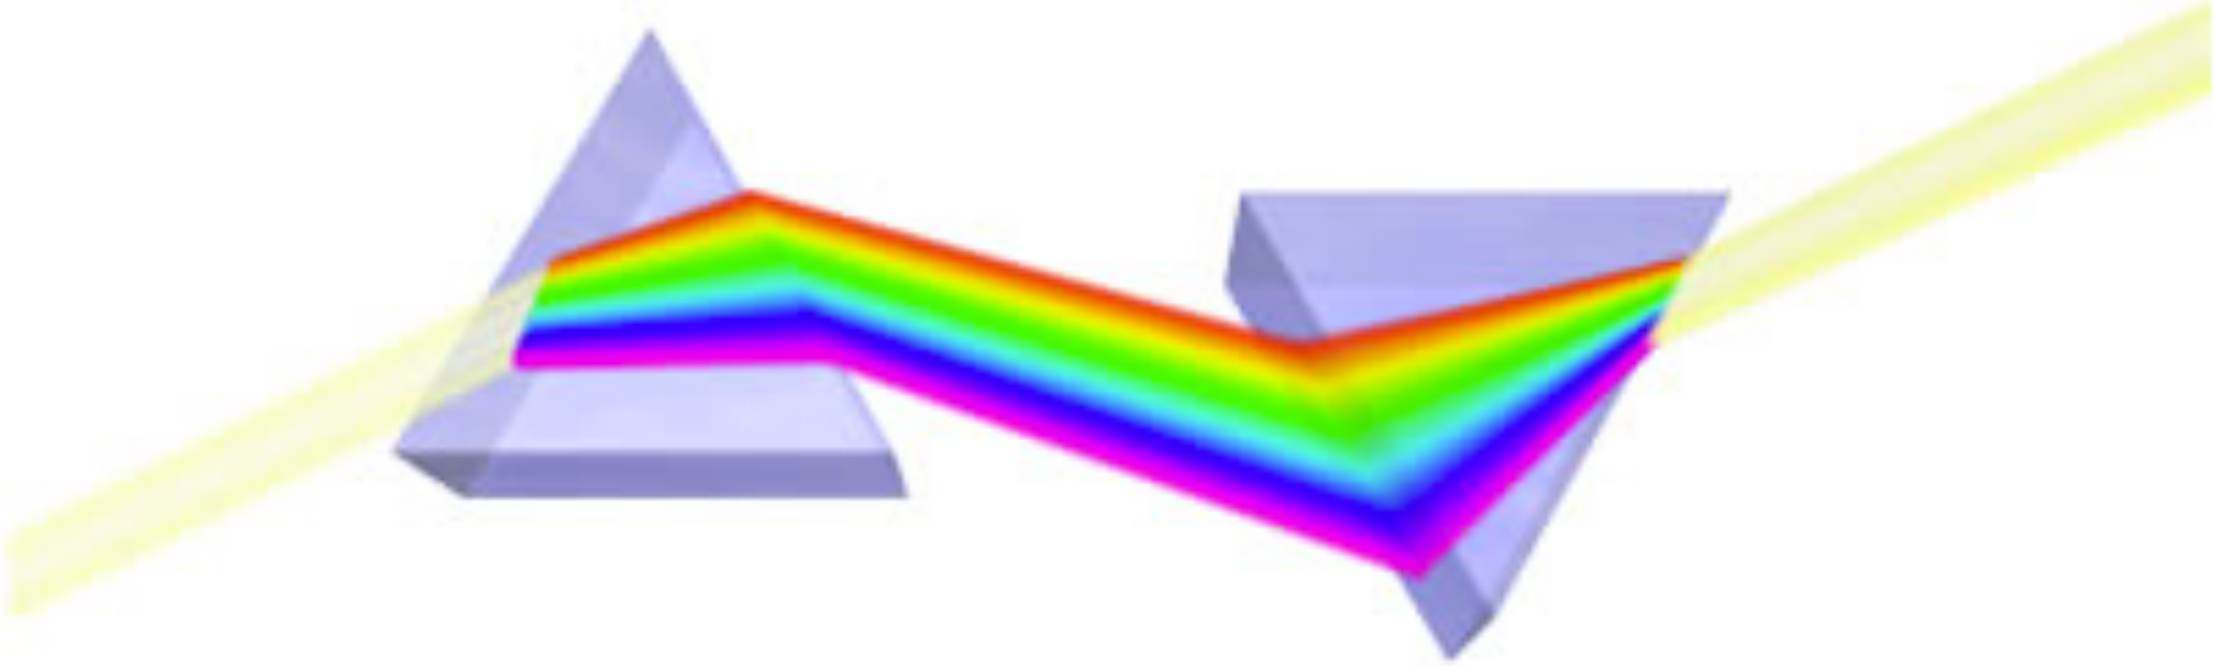

The electromagnetic spectrum

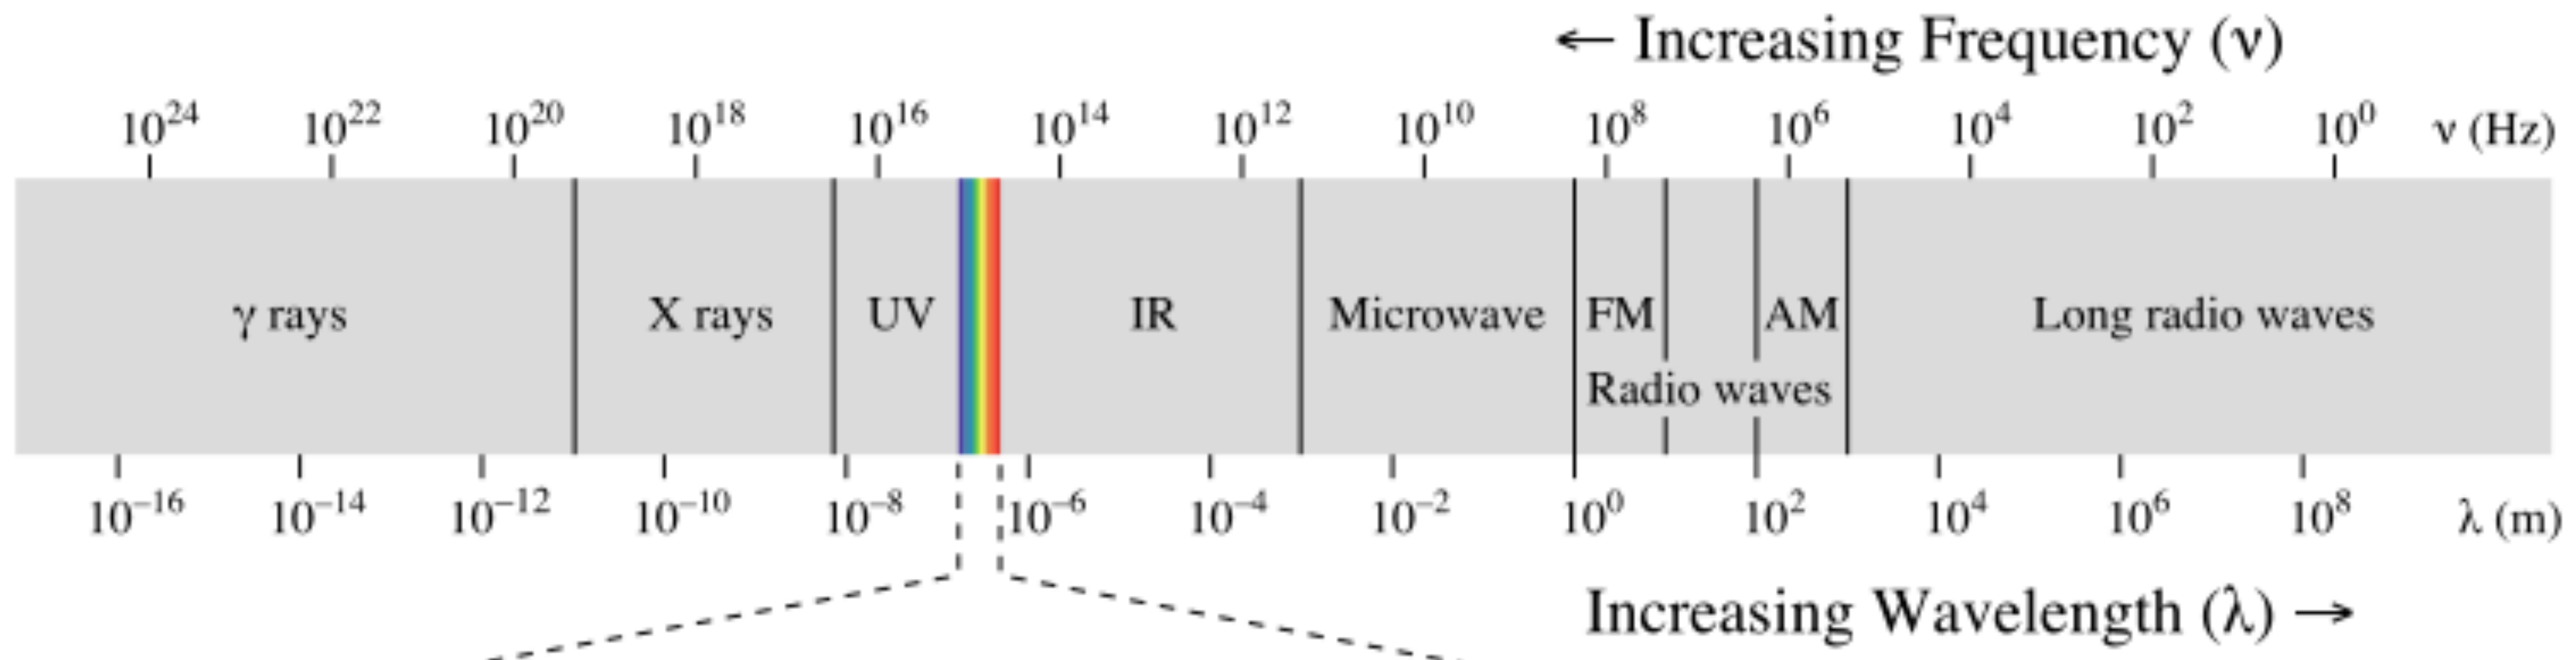

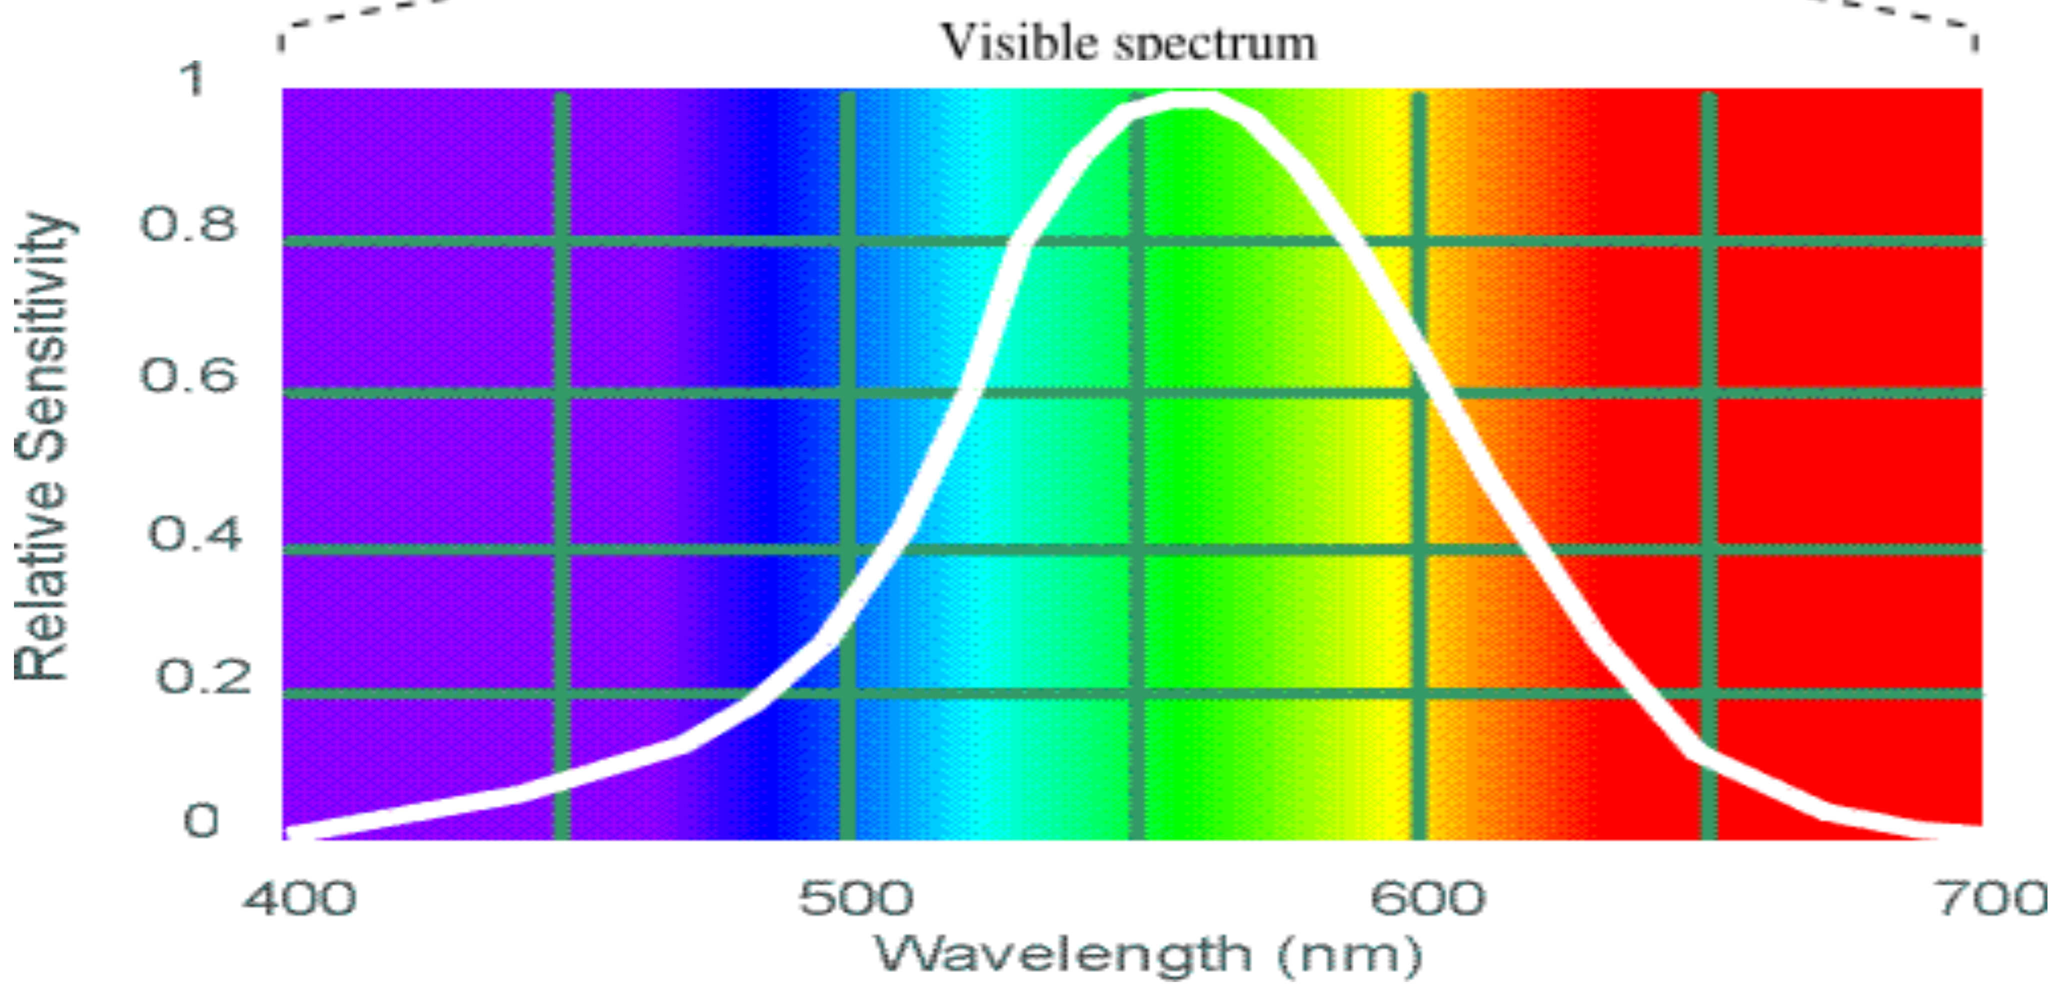

Human Luminance Sensitivity Function

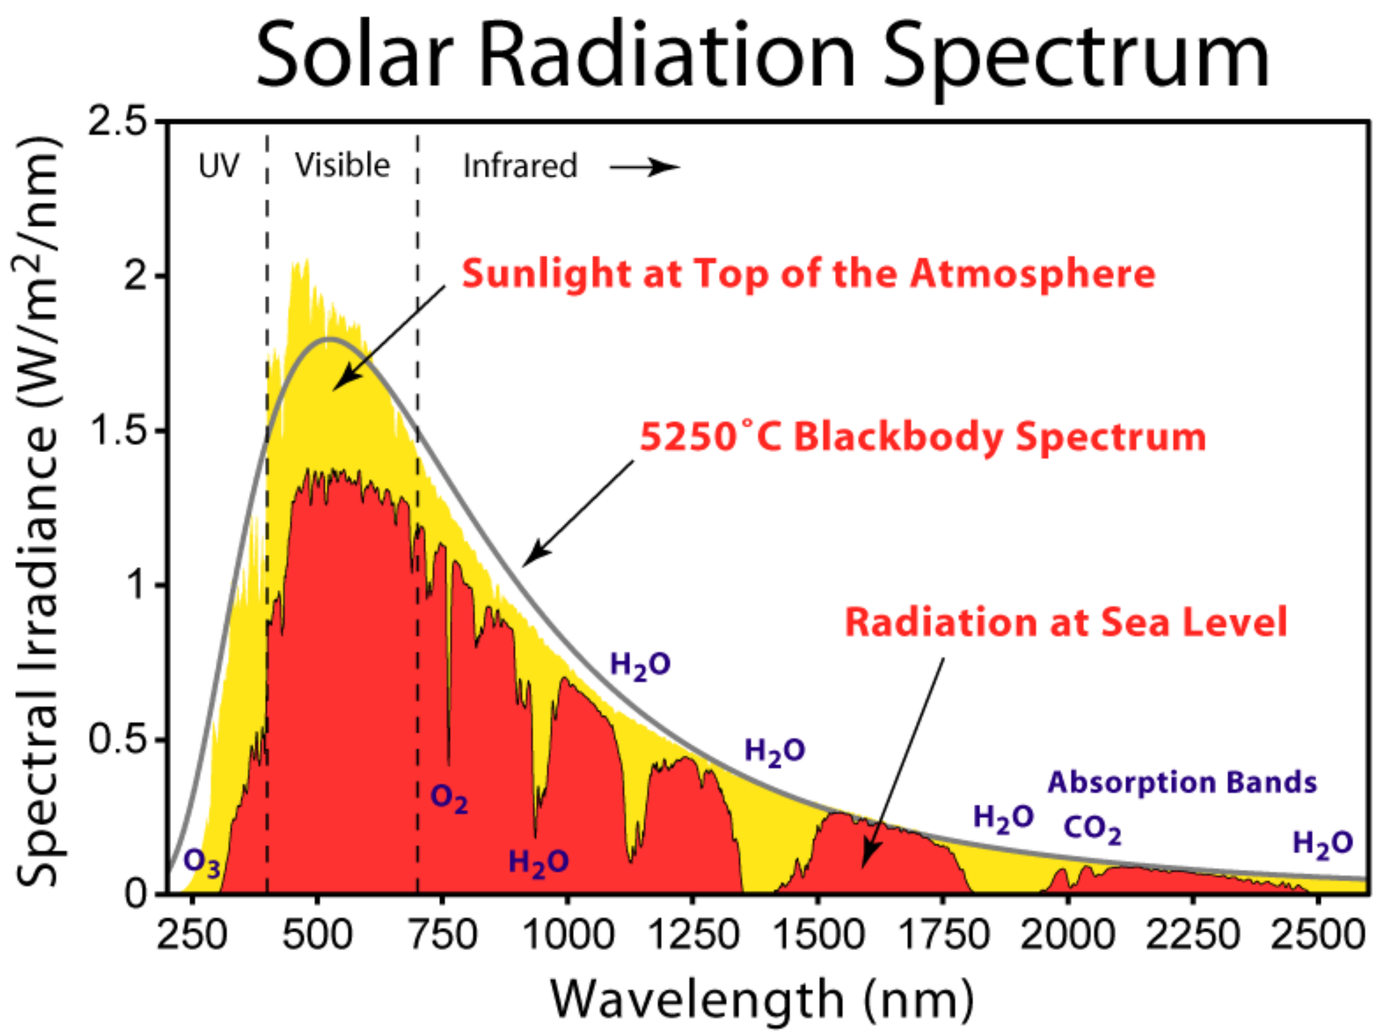

The Physics of Light

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

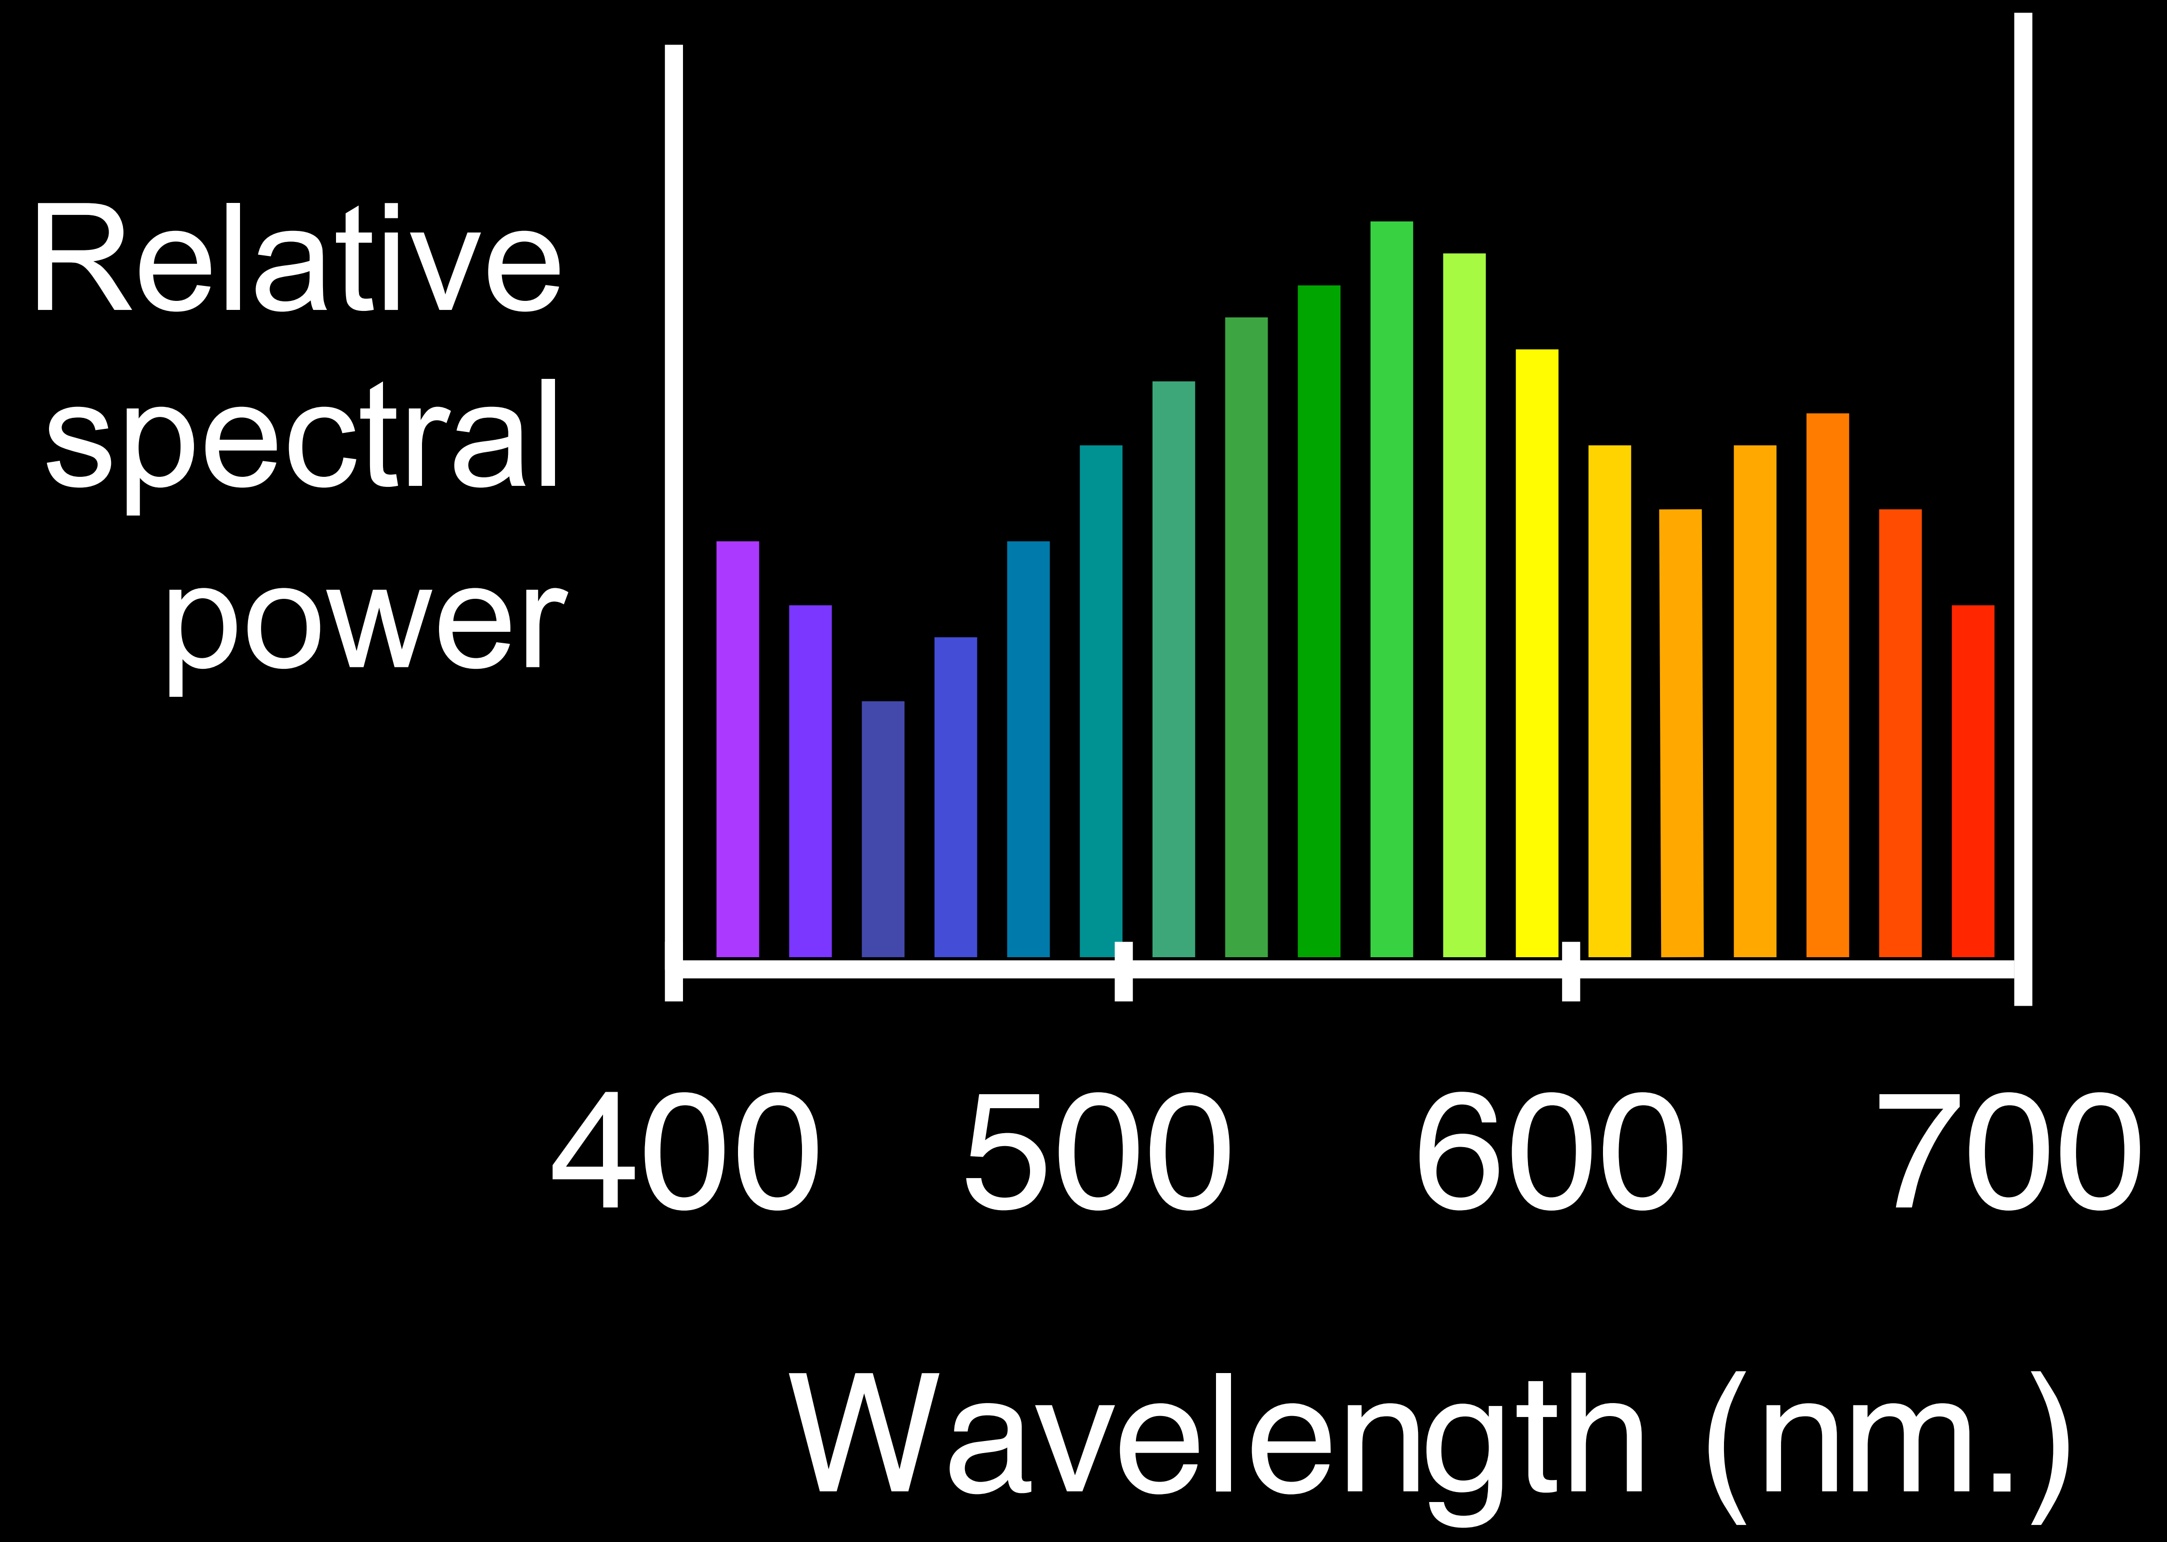

Spectra of Light Sources

Some examples of the spectra of light sources

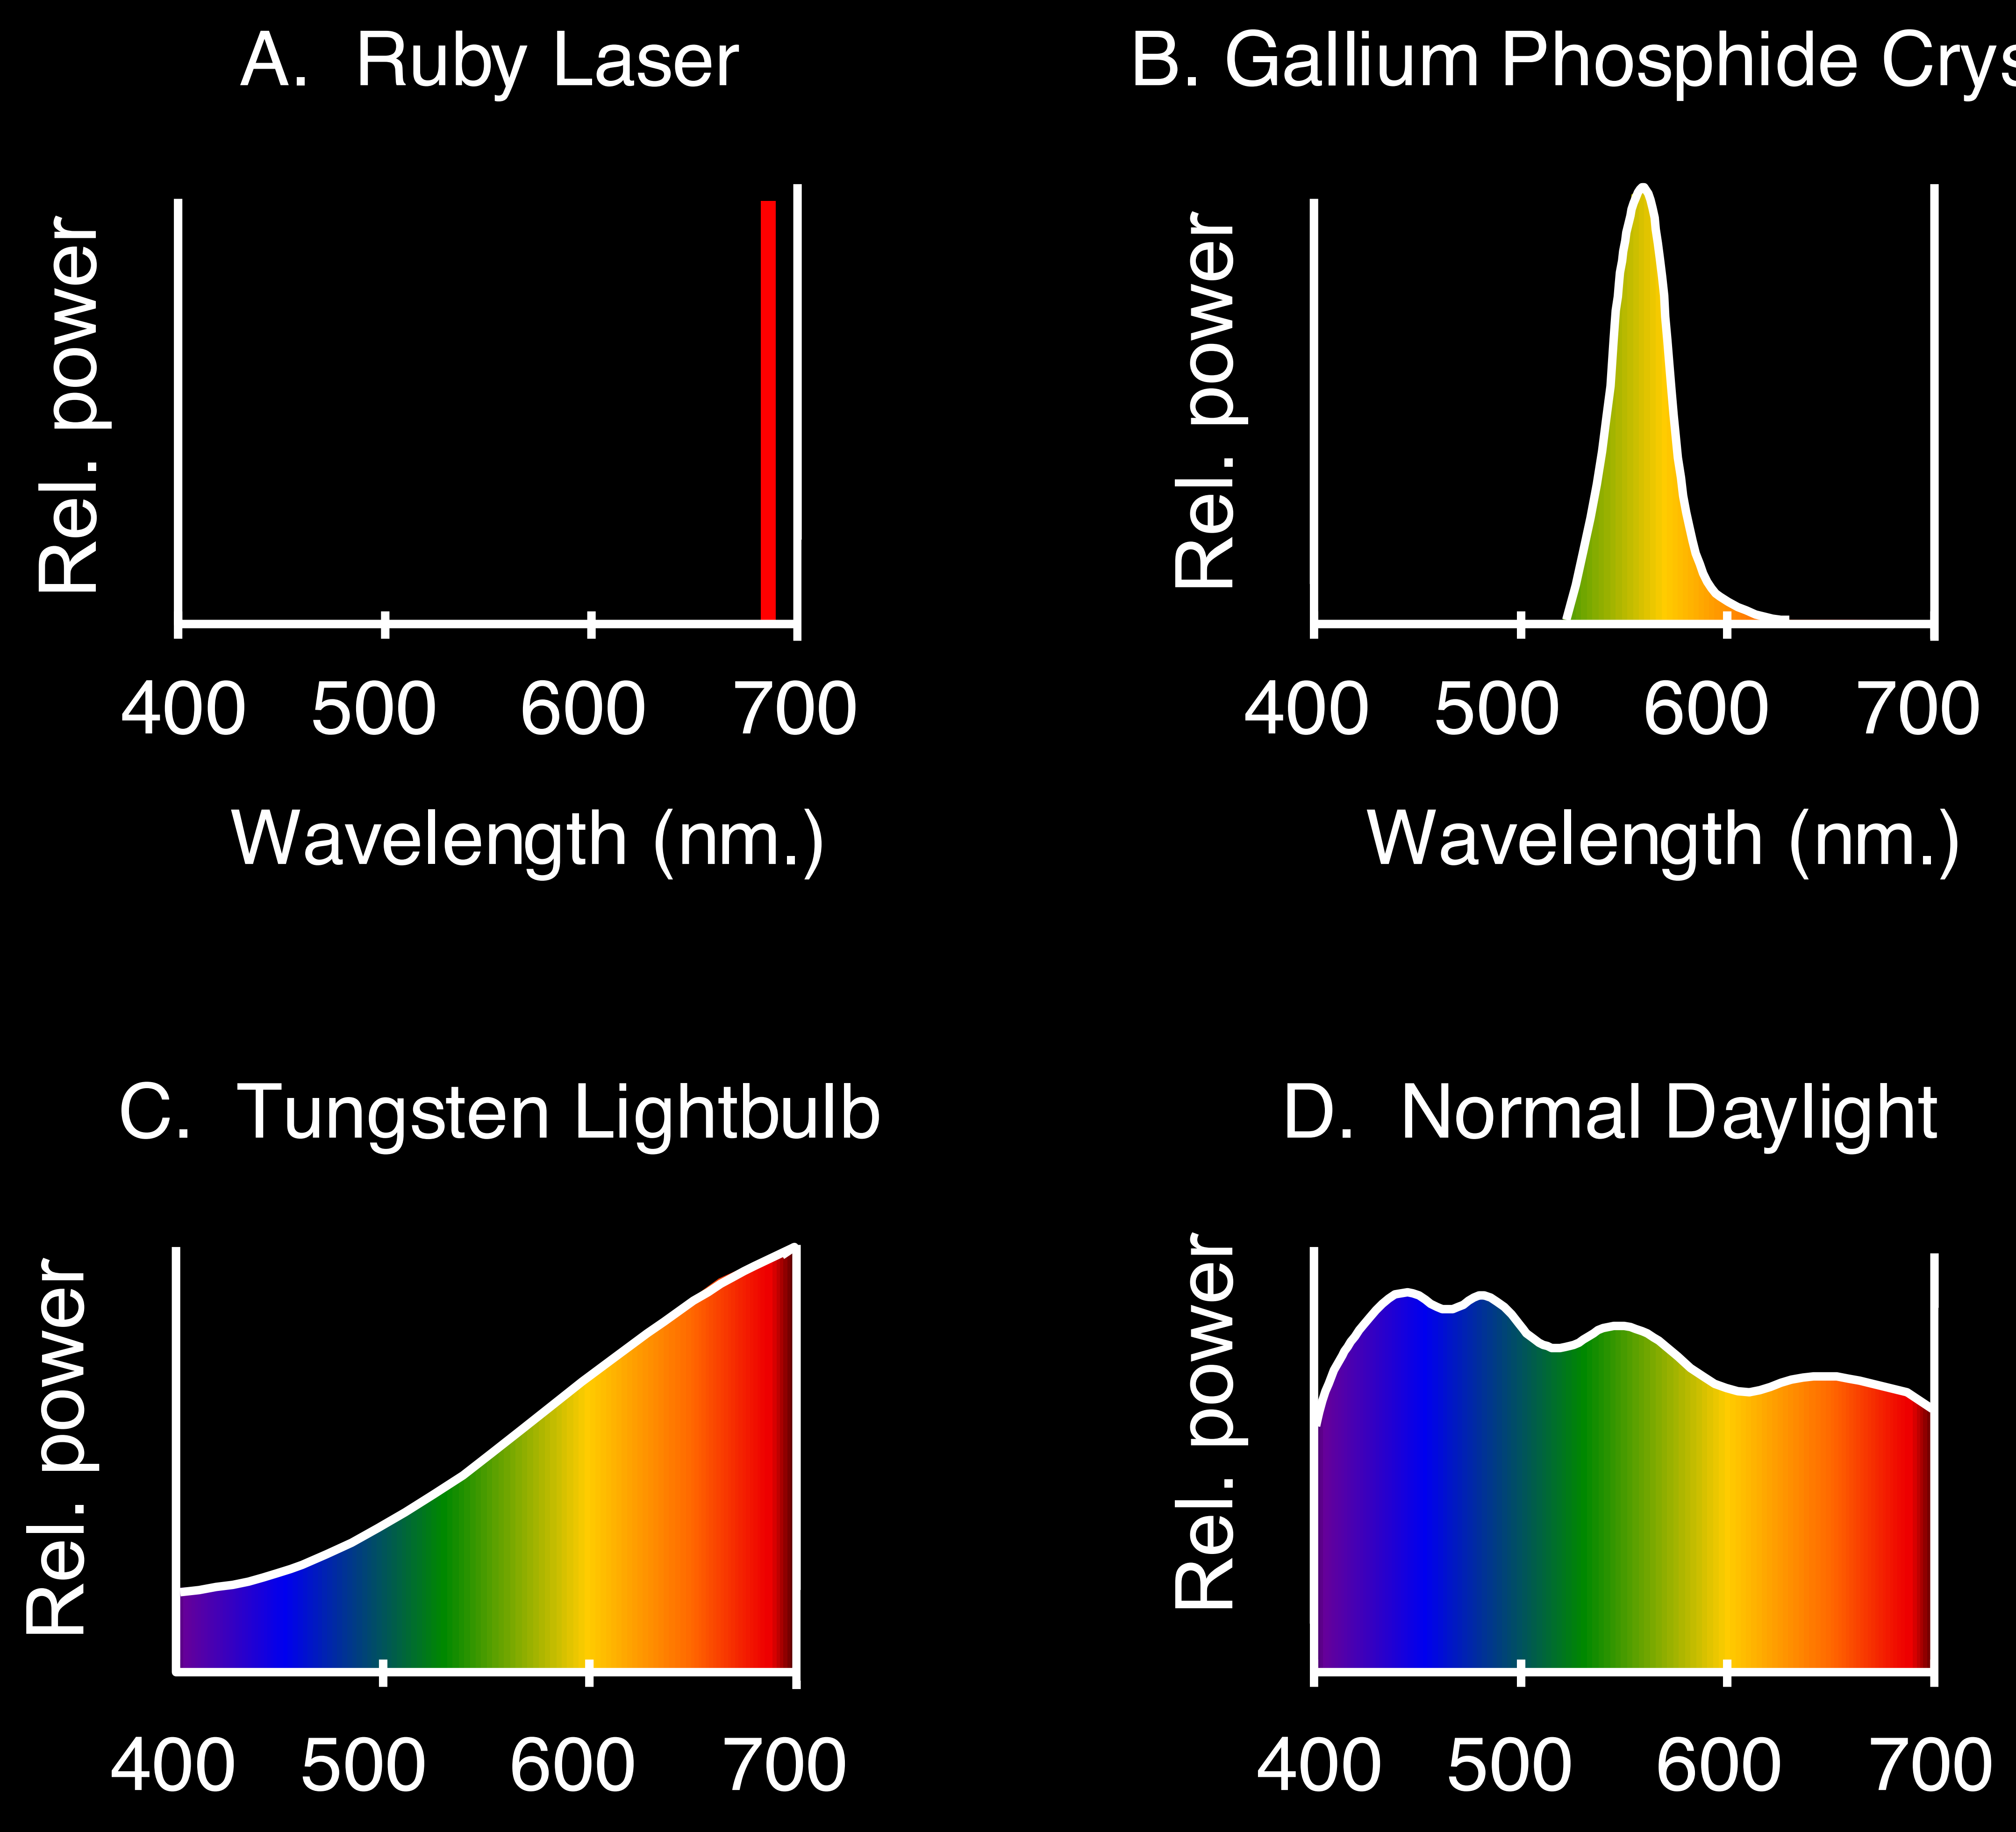

Reflectance Spectra of Surfaces

Some examples of the reflectance spectra of surfaces

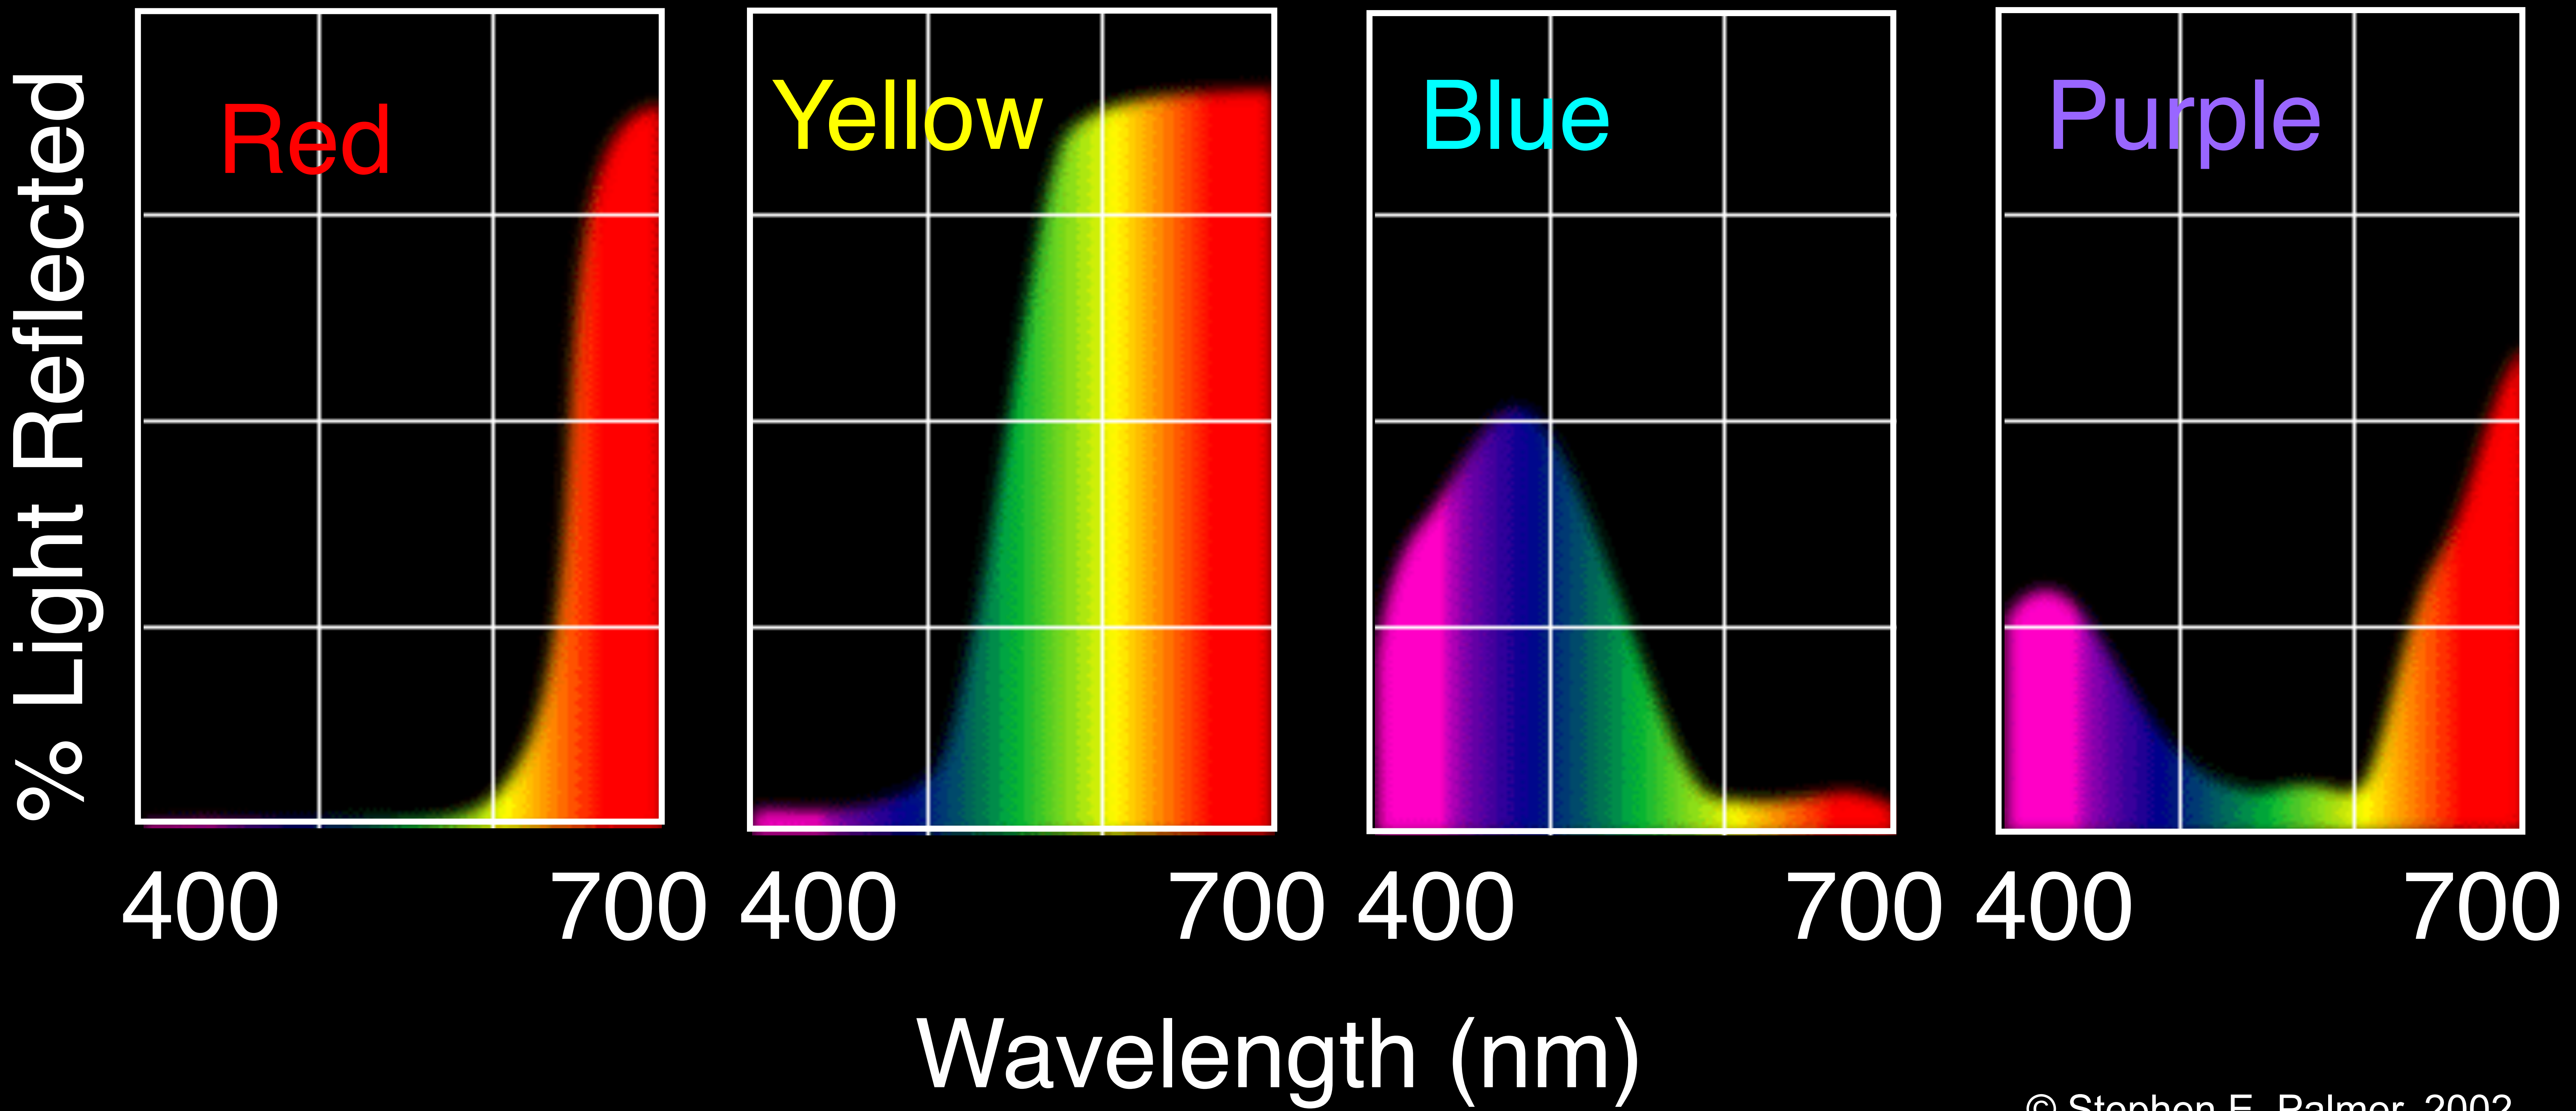

The eye

The human eye is a sophisticated camera!

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

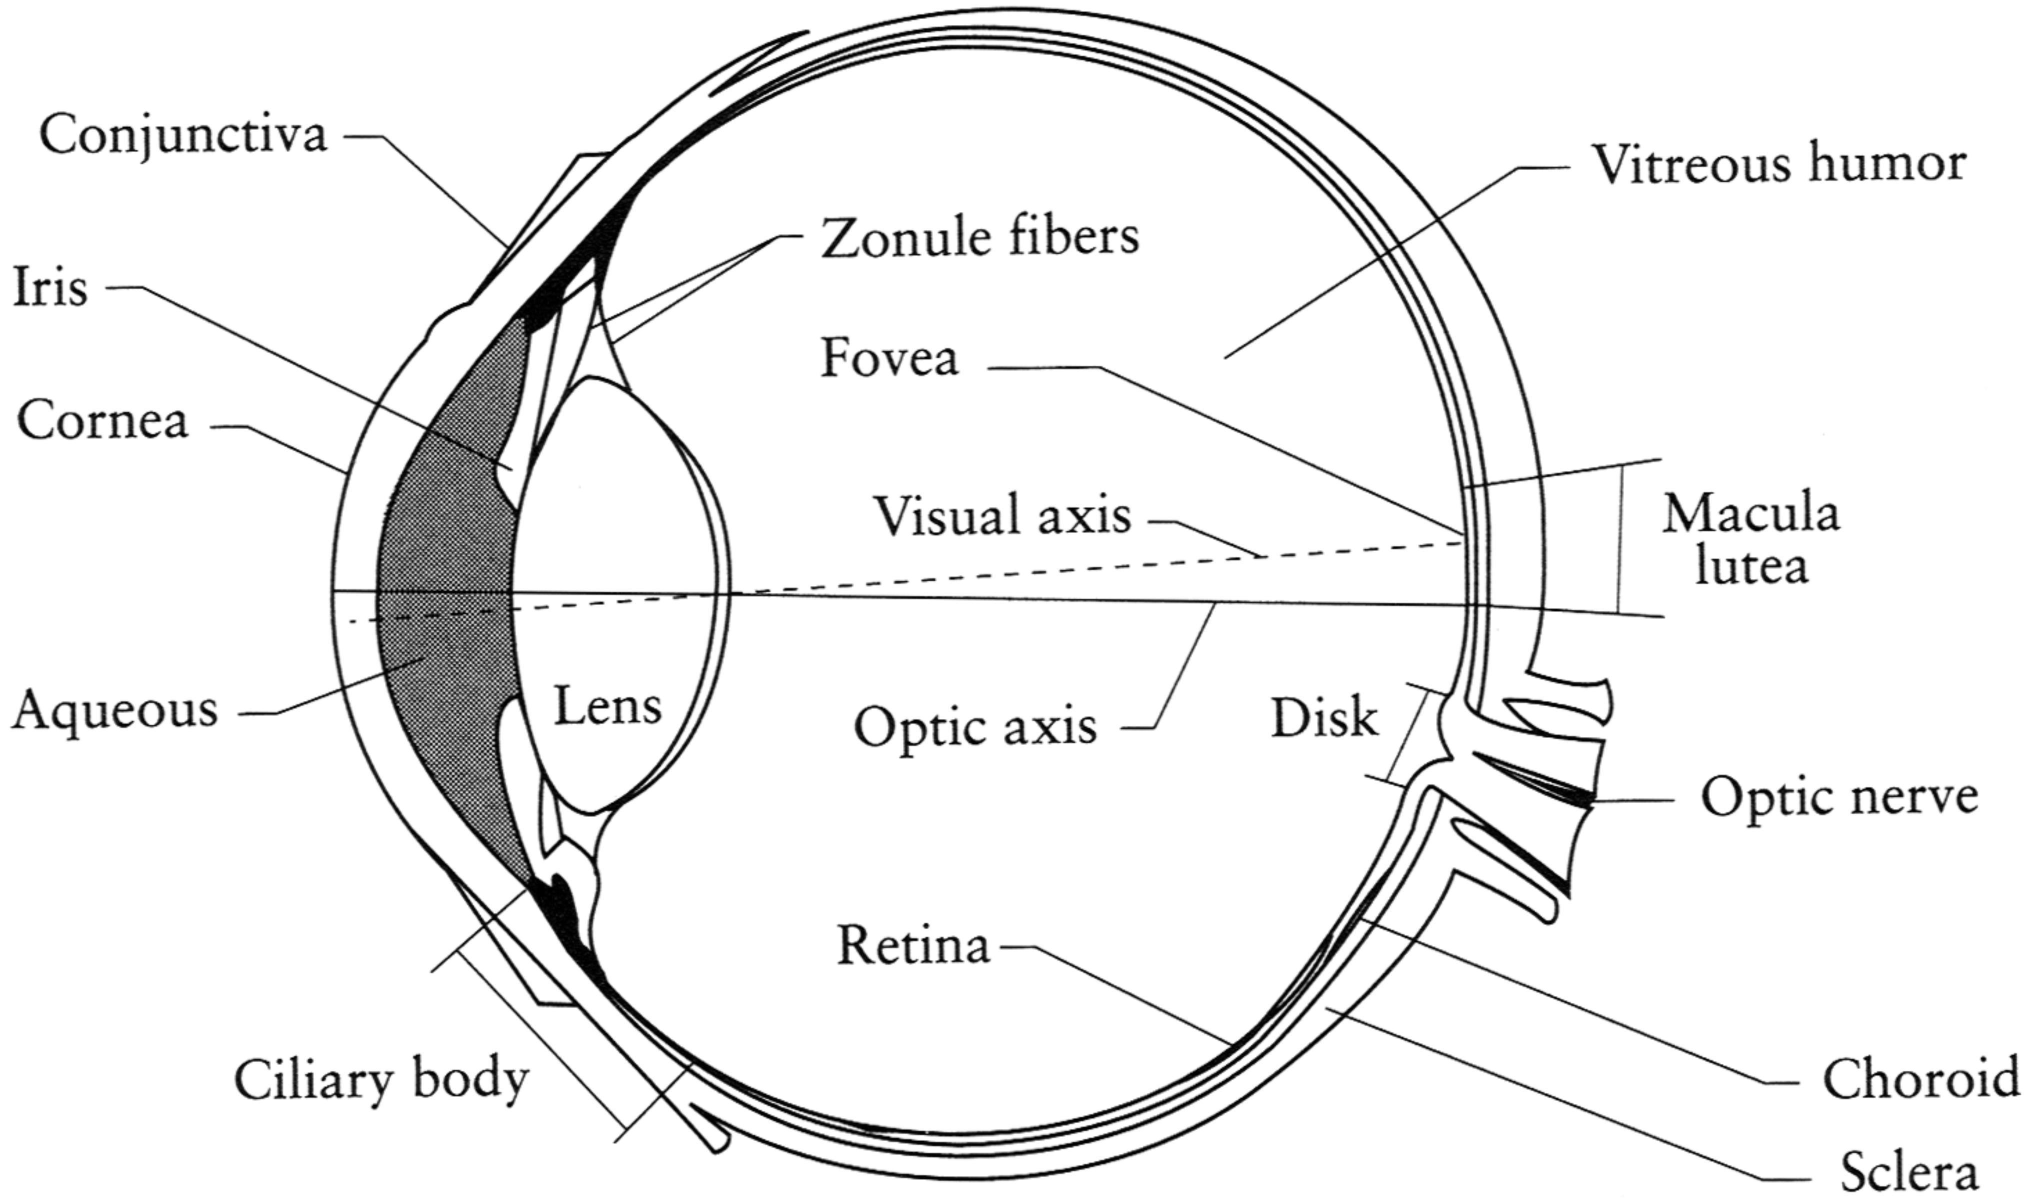

Rods and cones, fovea

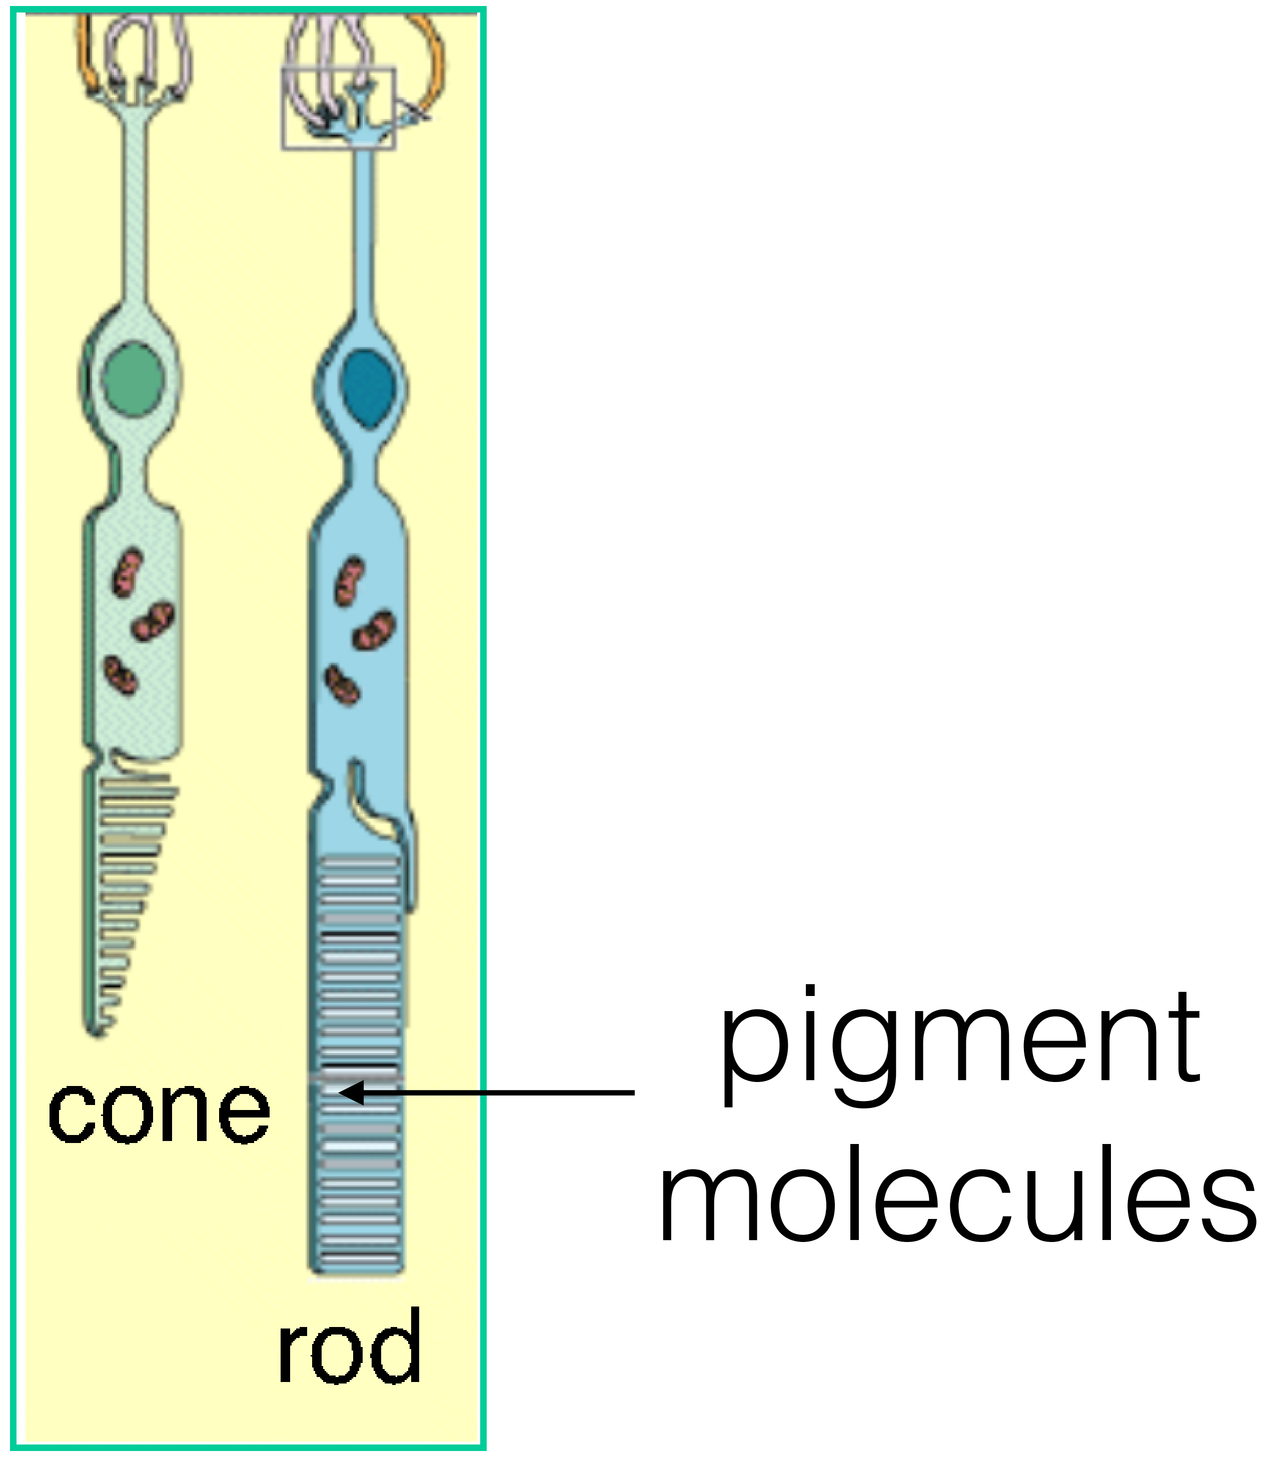

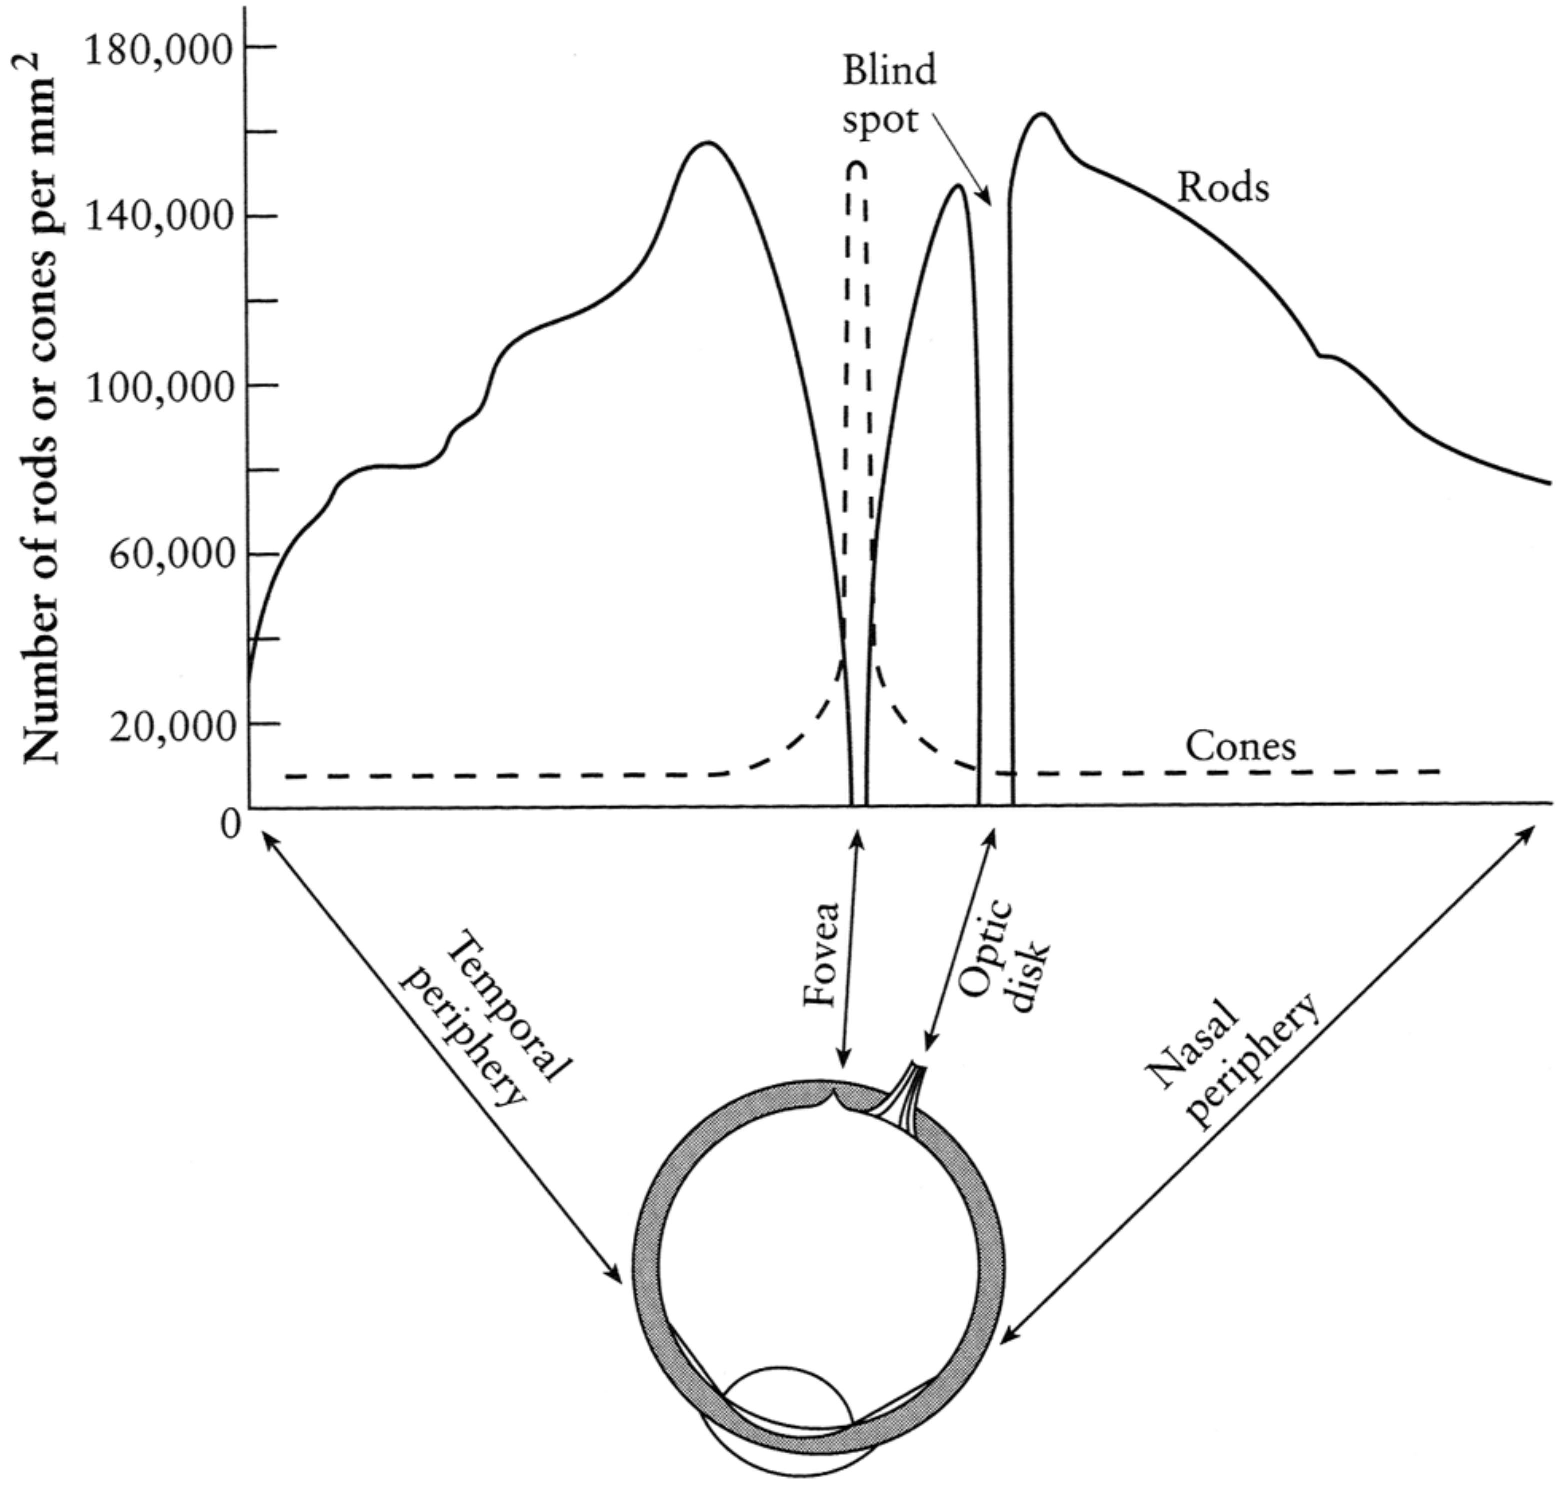

Rods are responsible for intensity, cones for color perception

Rods and cones are non-uniformly distributed on the retina

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

There are about 5 million cones and 100 million rods in each eye

Demonstration of visual acuity

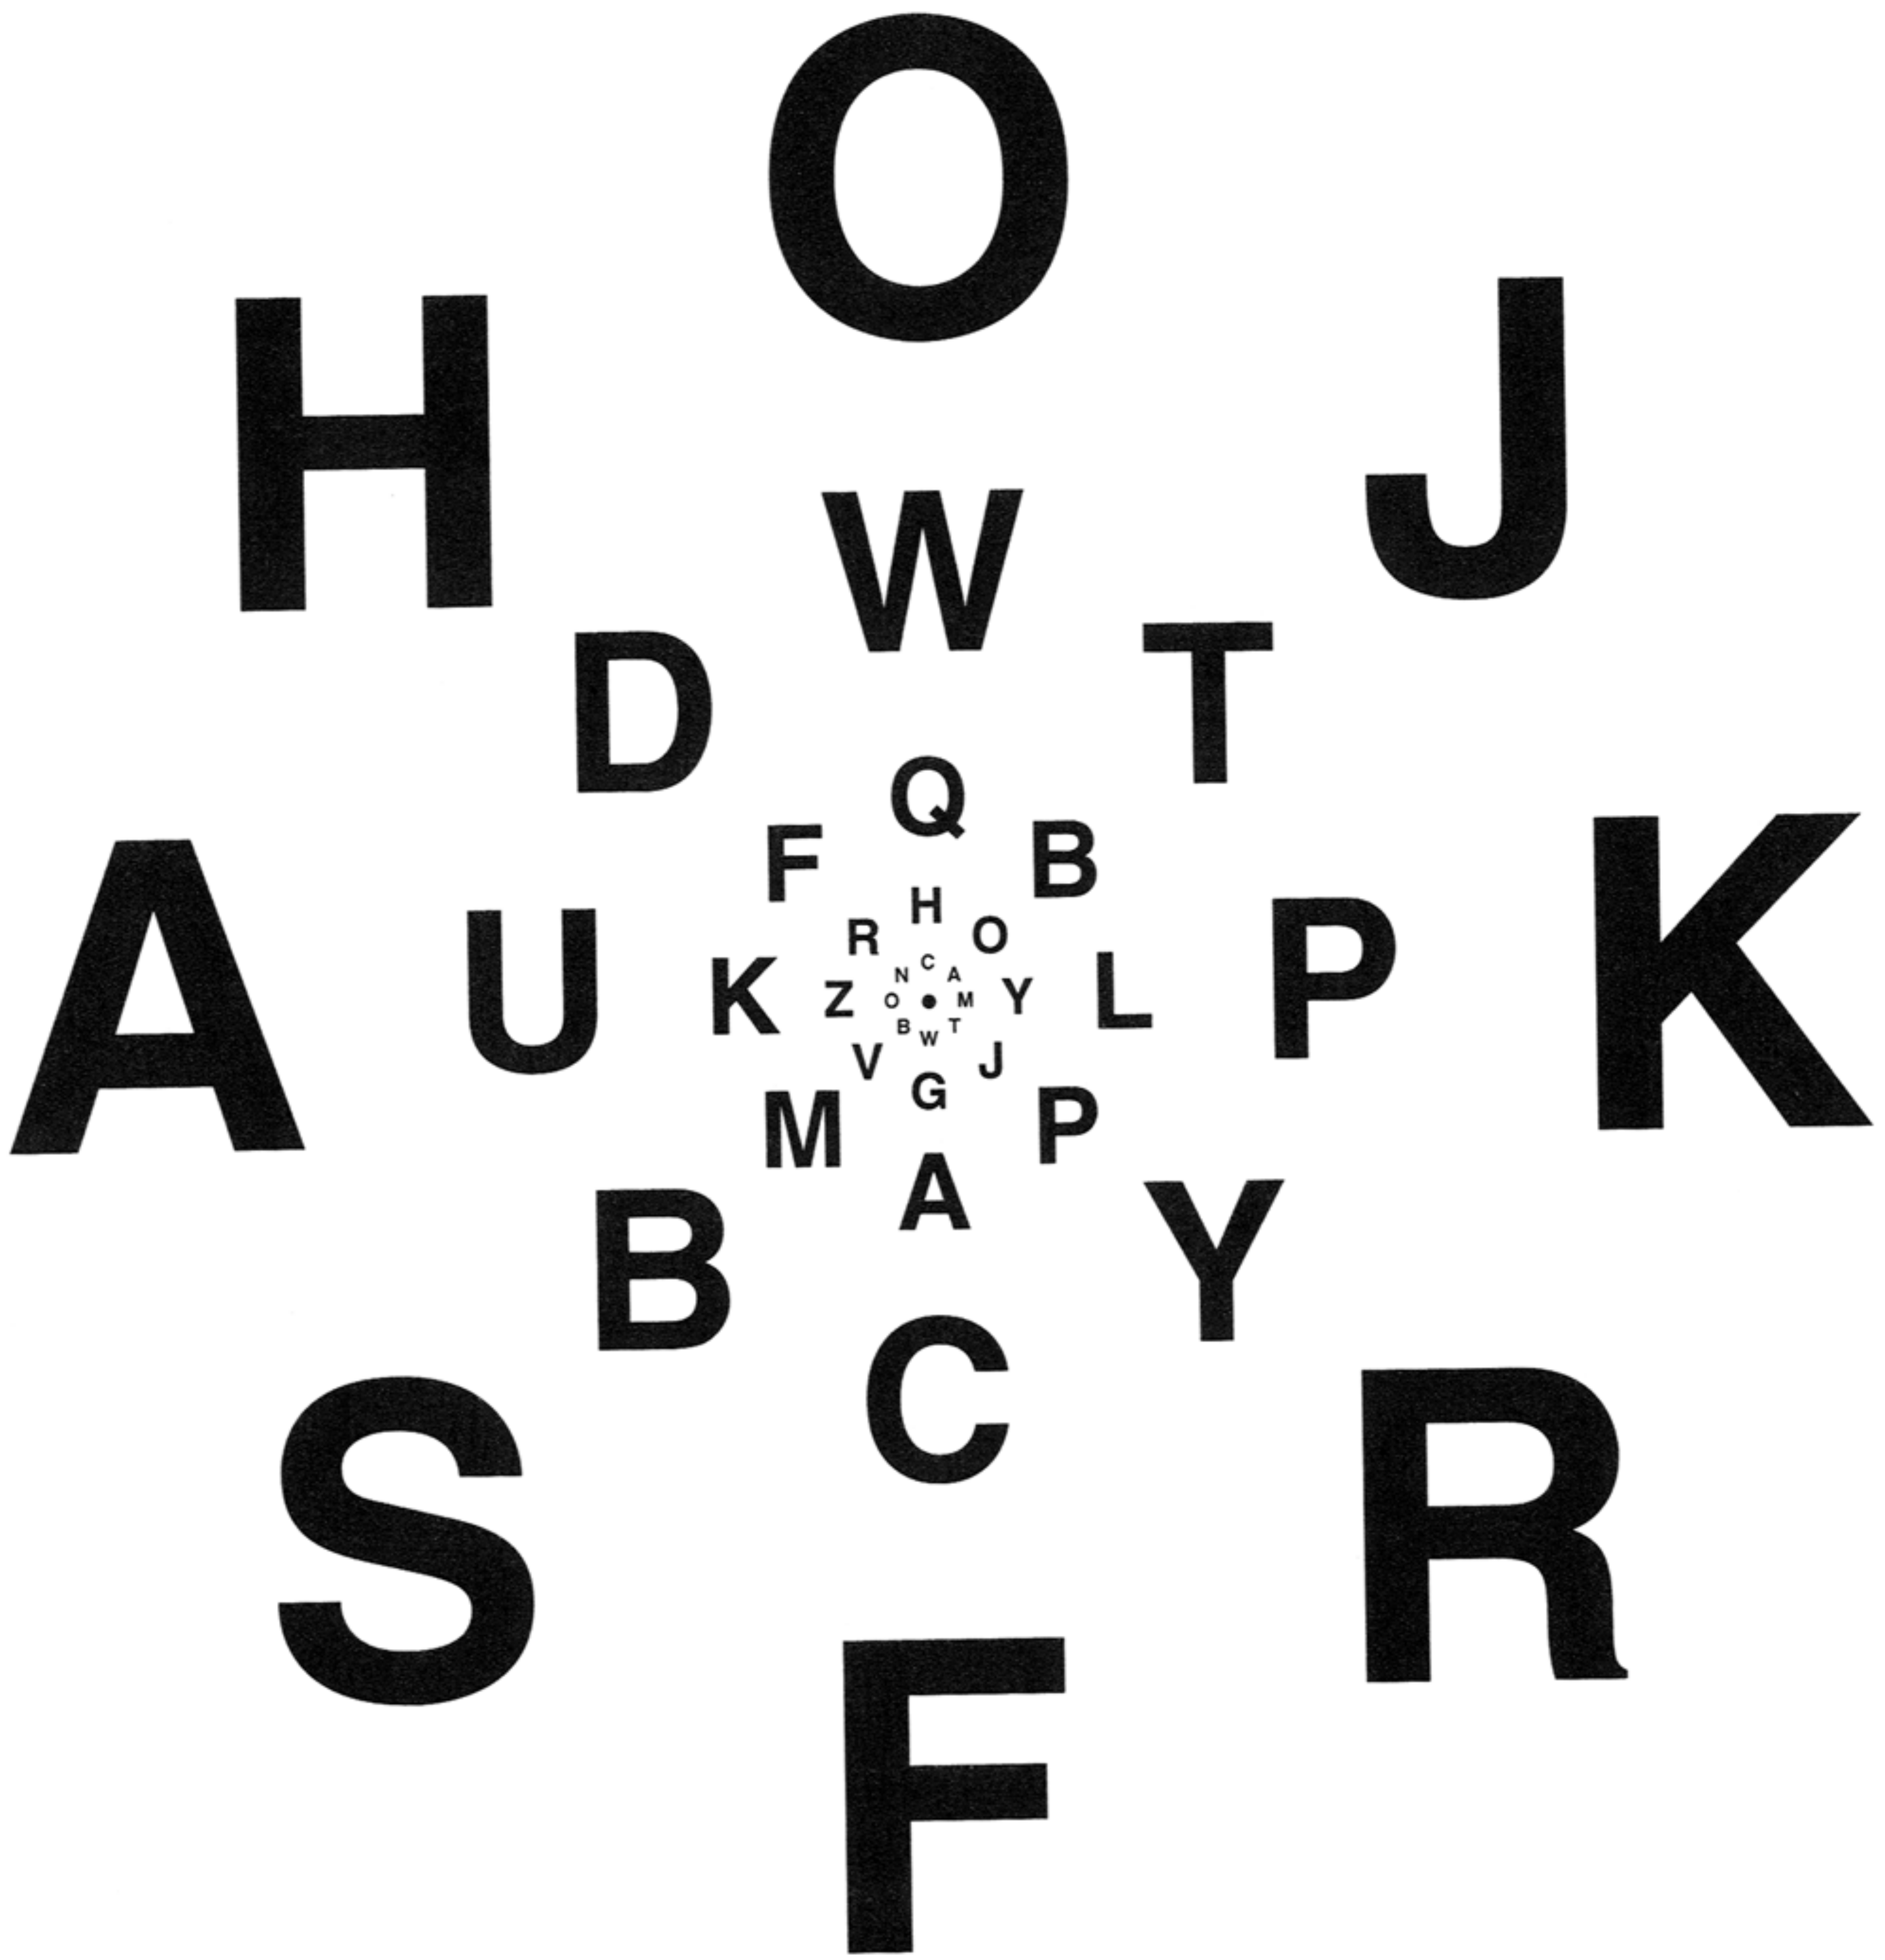

With one eye shut, at the right distance, all of these letters should appear equally legible (Glassner, 1.7).

Slide by Steve Seitz

Blind spot

With left eye shut, look at the cross on the left. At the right distance, the circle on the right should disappear (Glassner, 1.8).

Slide by Steve Seitz

Rod/cone sensitivity

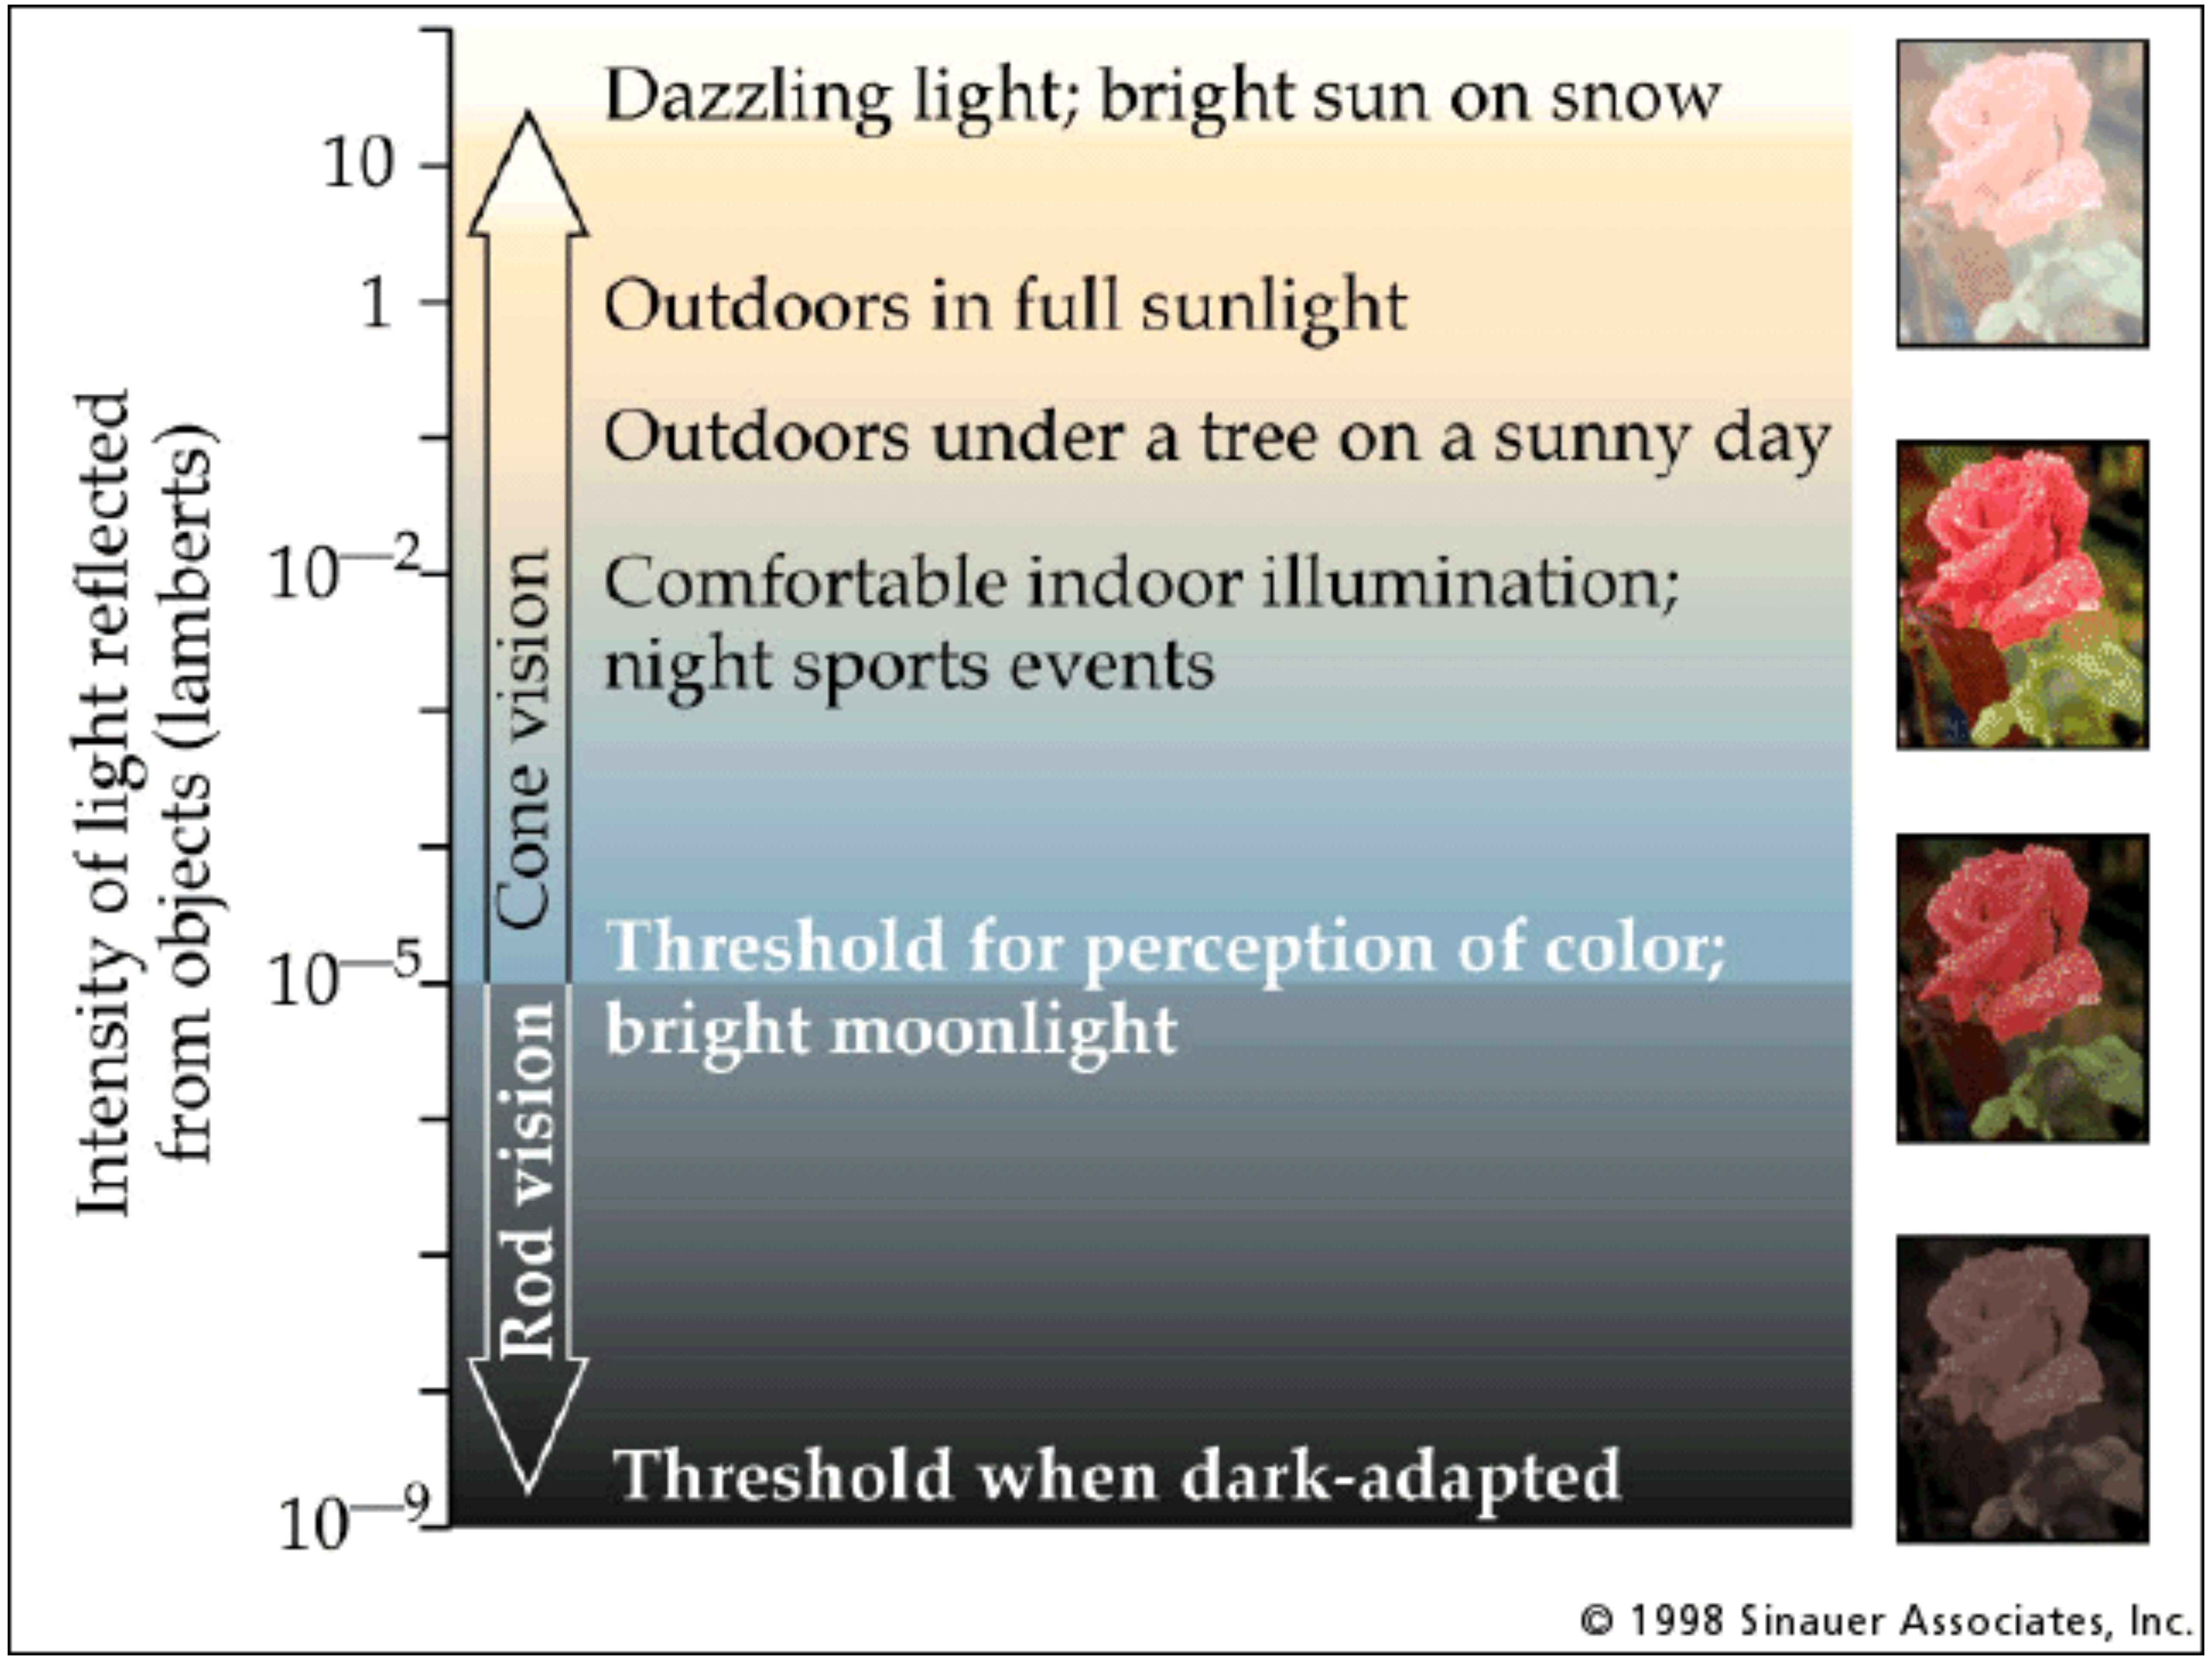

Why can't we read in the dark?

Slide by A. Efros

Physiology of Color Vision

Three kinds of cones:

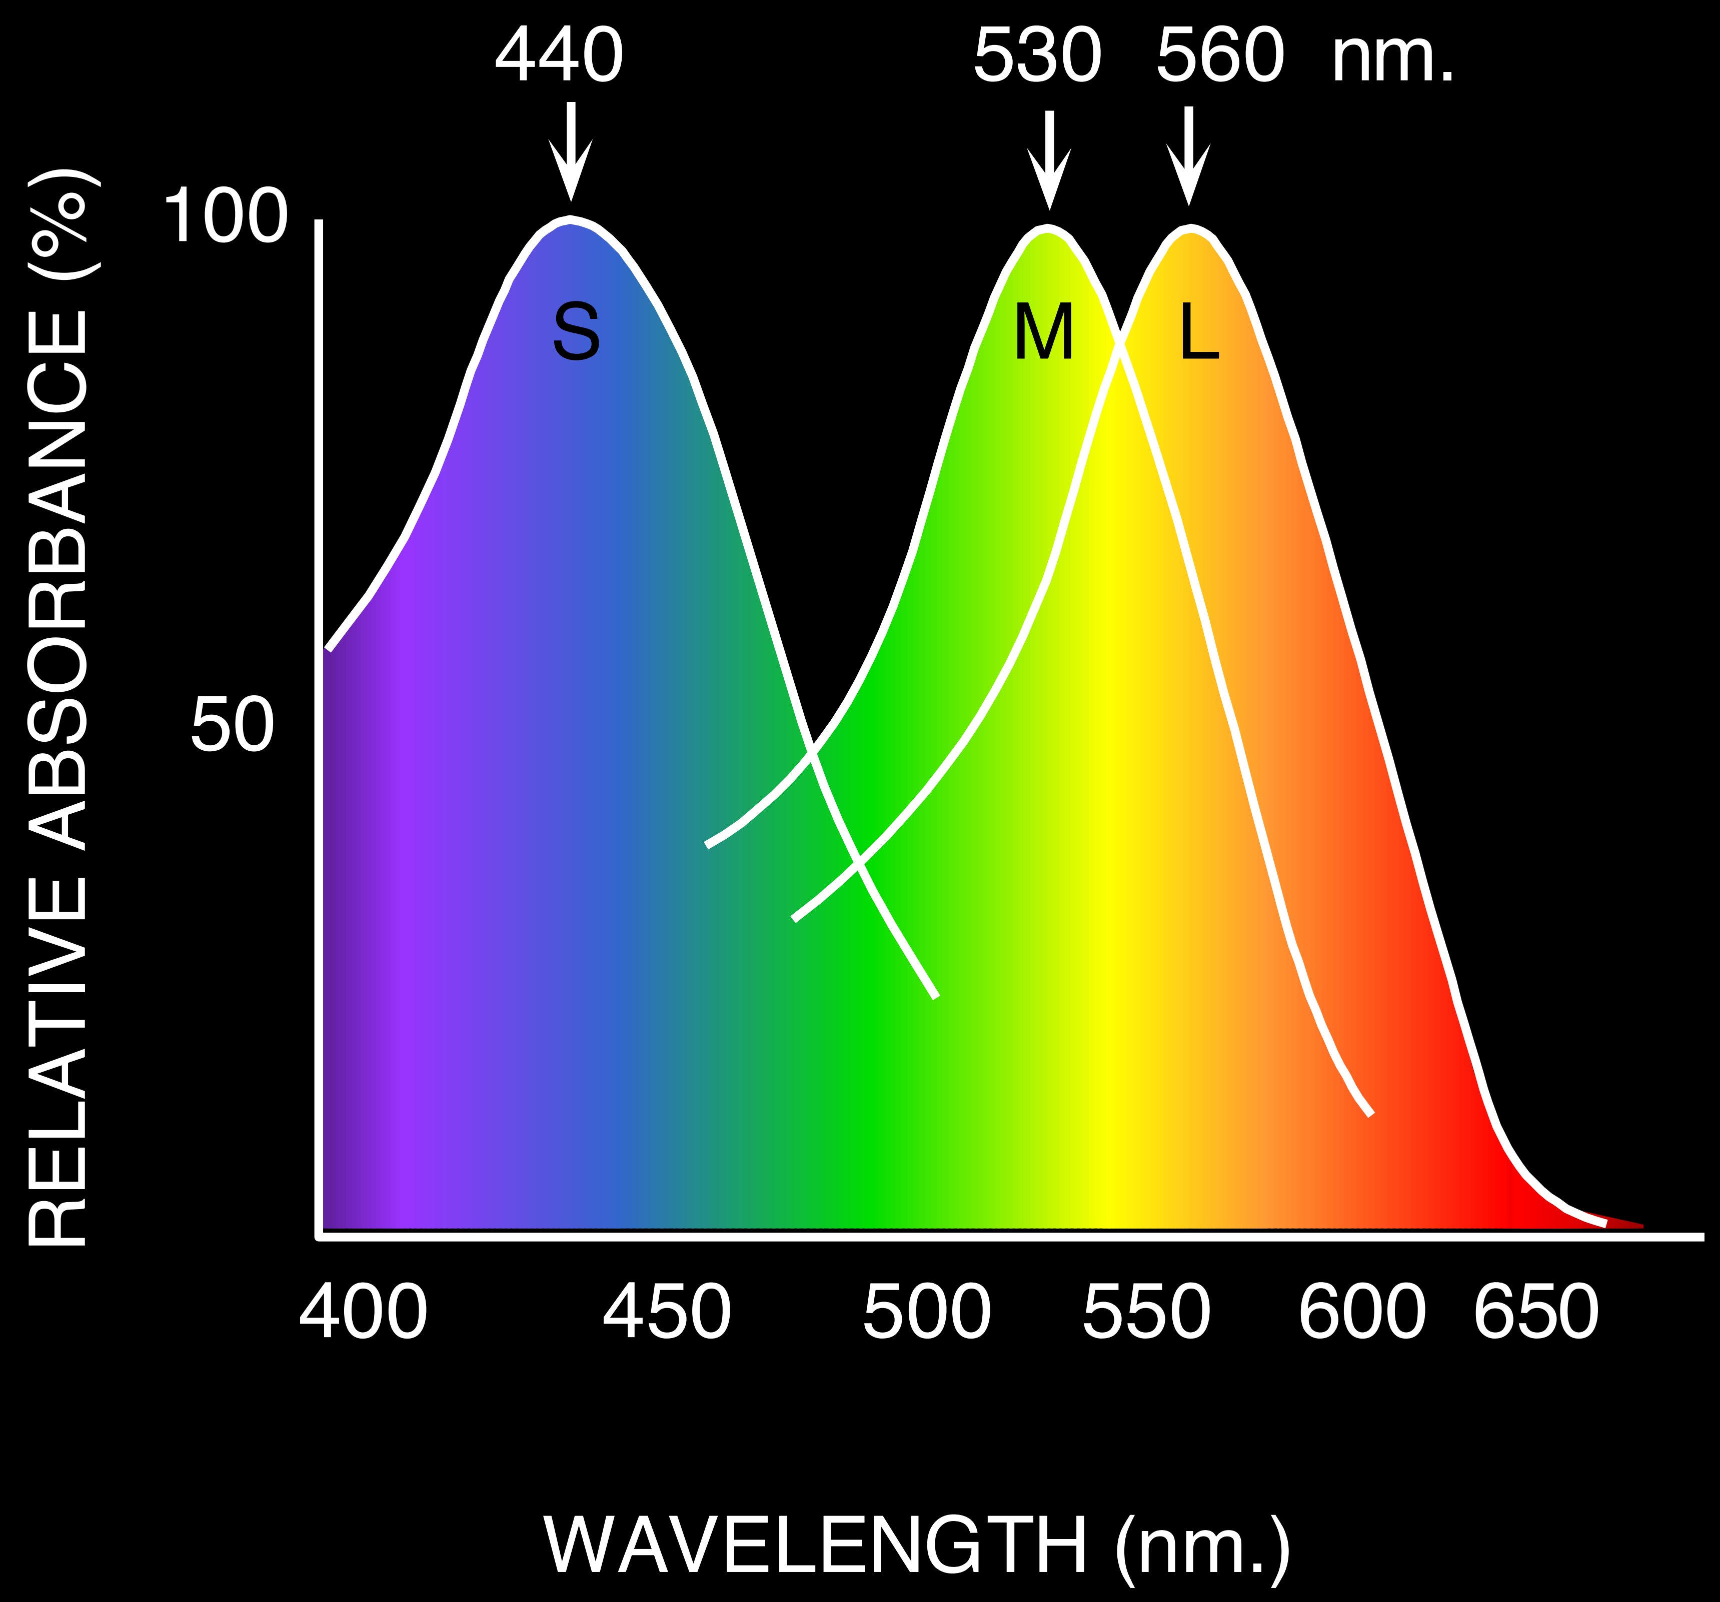

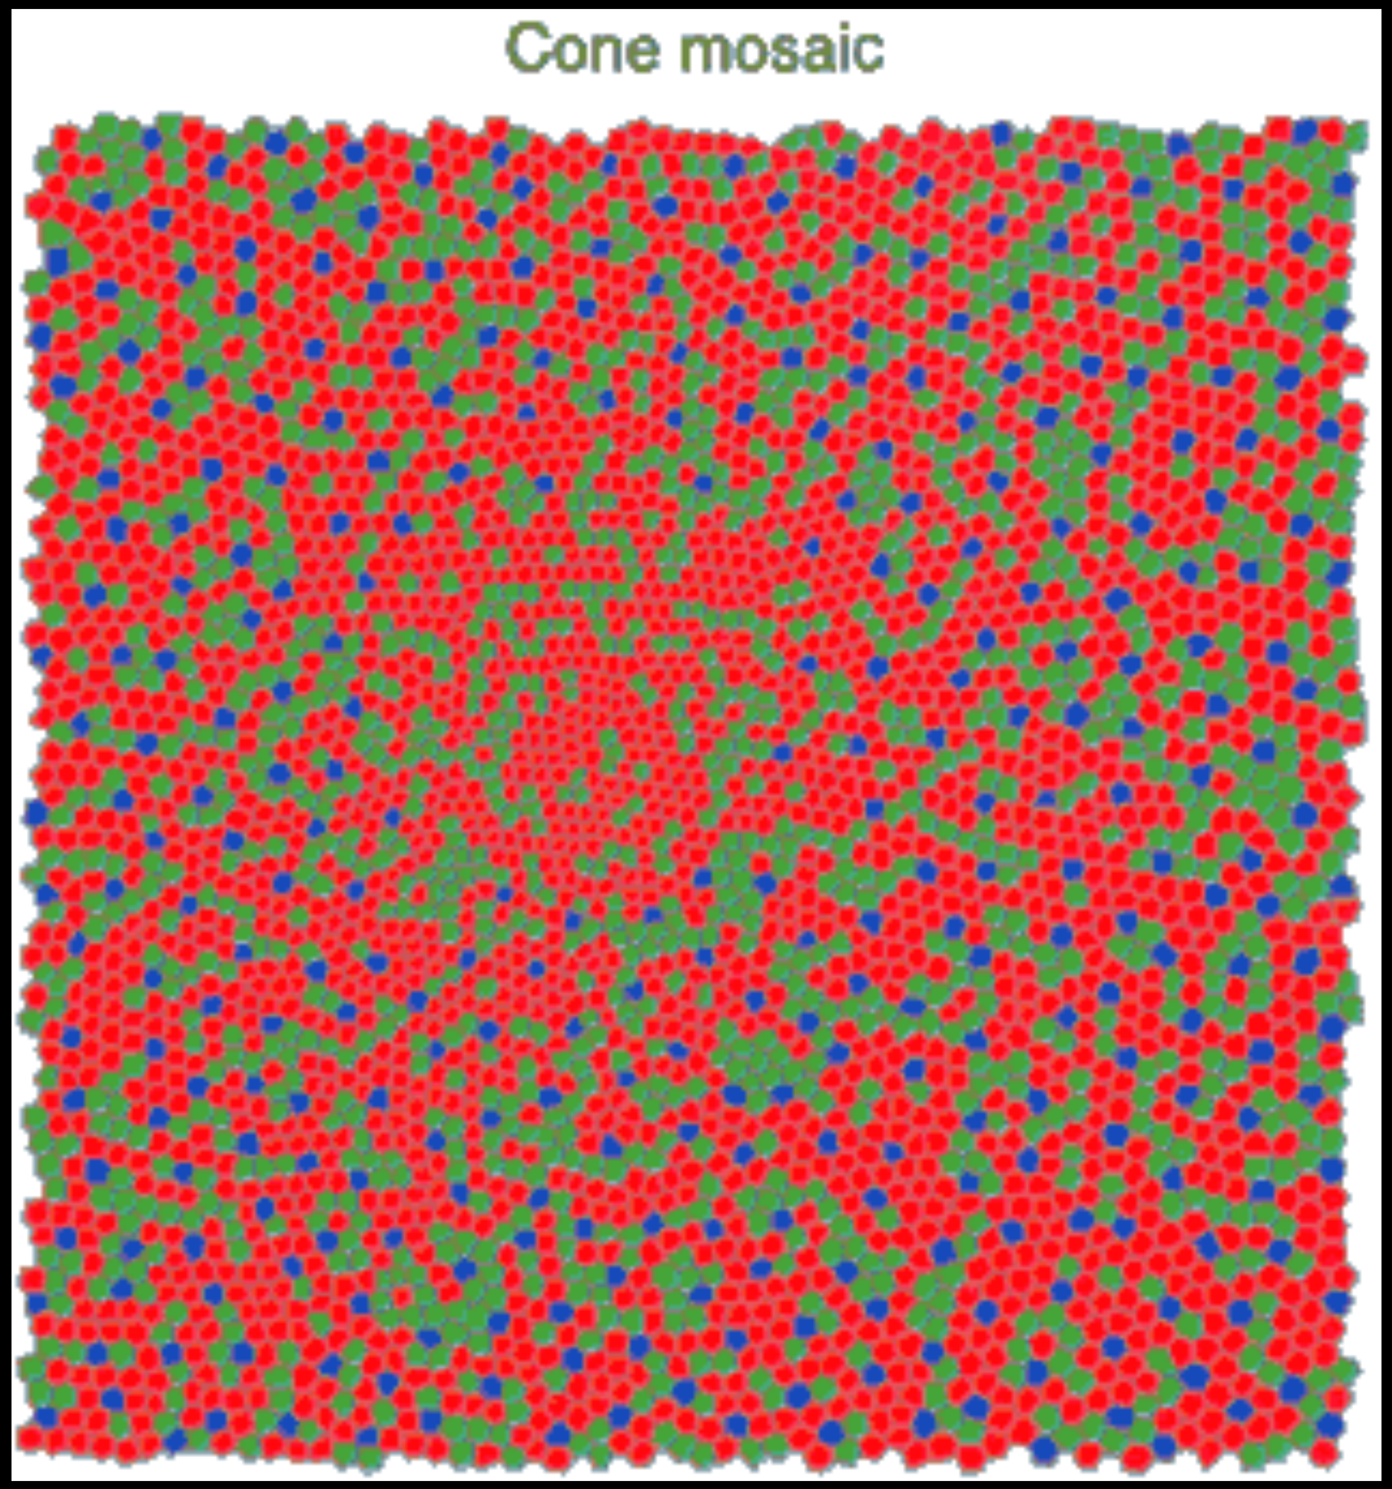

- Ratio of L to M to S cones: approx. 10:5:1
- Almost no S cones in the center of the fovea

Color perception

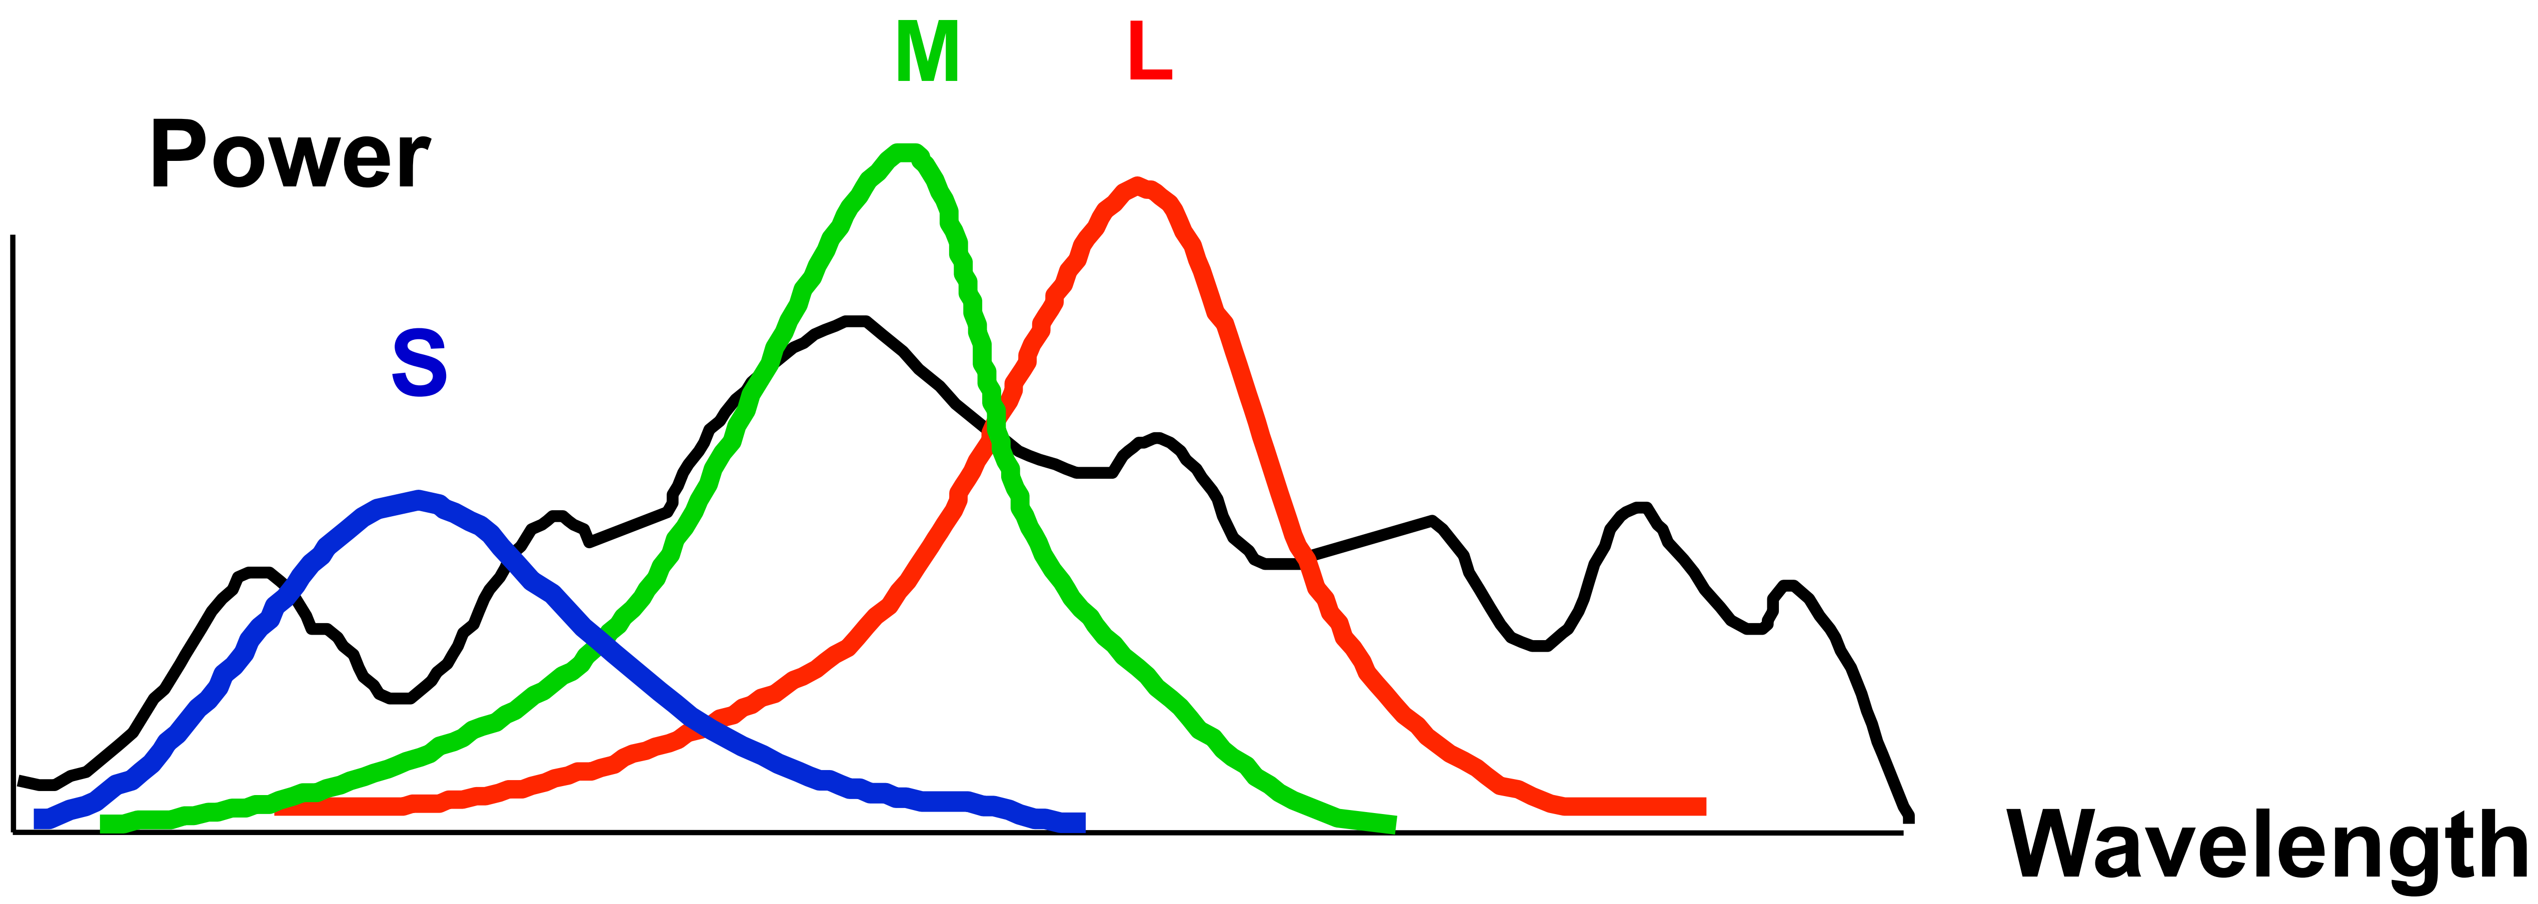

Rods and cones act as filters on the spectrum

To get the output of a filter, multiply its response curve by the spectrum, integrate over all wavelengths. Thus, each cone yields one number.

Example: $L = [0.1, 0.3, 0.2, 0.1, 0.1, 0.0, 0.0, 0.0, 0.0]$

$S = [0.0, 0.1, 0.1, 0.2, 0.2, 0.3, 0.1, 0.0, 0.0]$

$$R_S = \sum_i L(i) \times S(i) = 0.09$$

Color perception

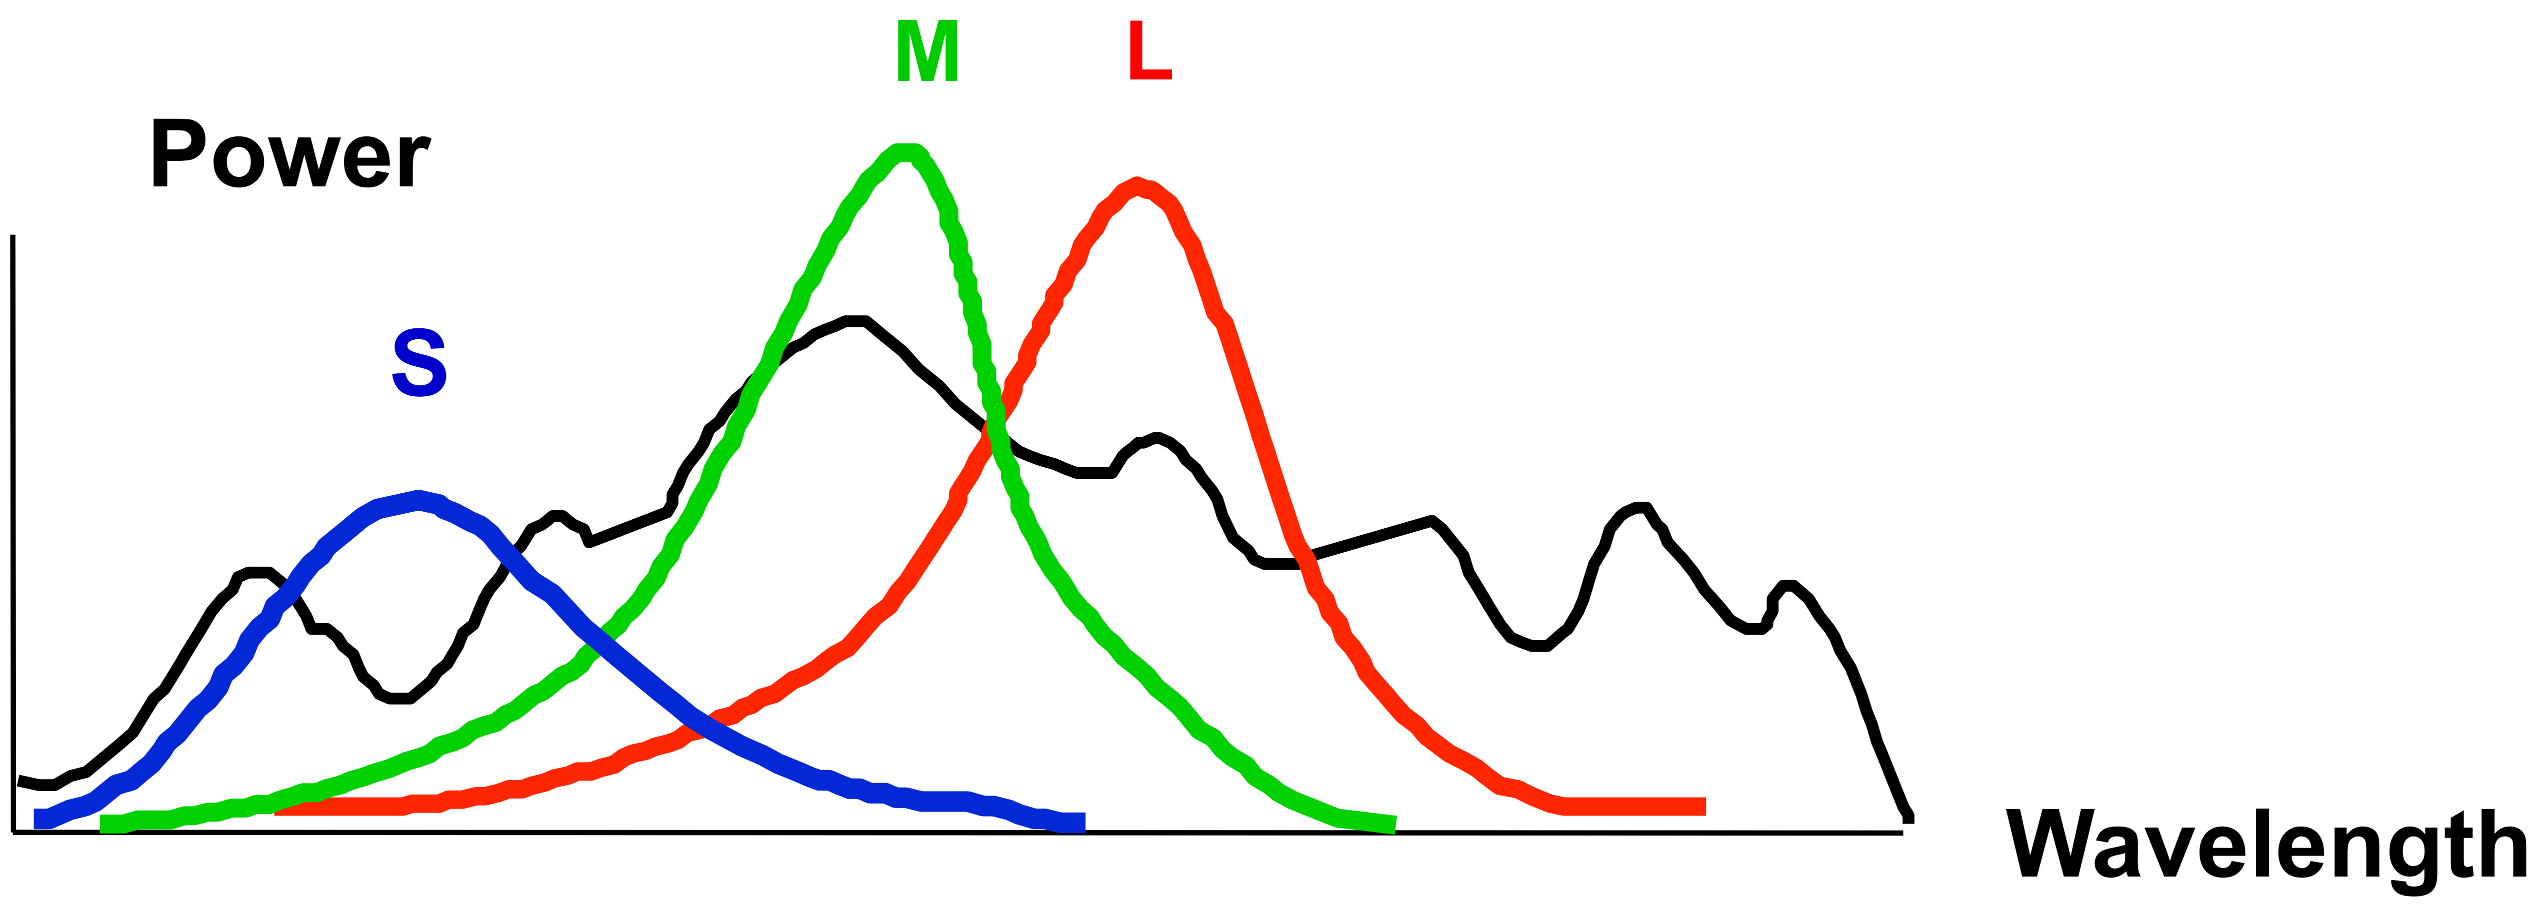

Rods and cones act as filters on the spectrum

To get the output of a filter, multiply its response curve by the spectrum, integrate over all wavelengths. Thus, each cone yields one number.

How can we represent an entire spectrum with 3 numbers?

We can’t! A lot of the information is lost

- As a result, two different spectra may appear indistinguishable.
- Such spectra are known as metamers

Spectra of some real-world surfaces

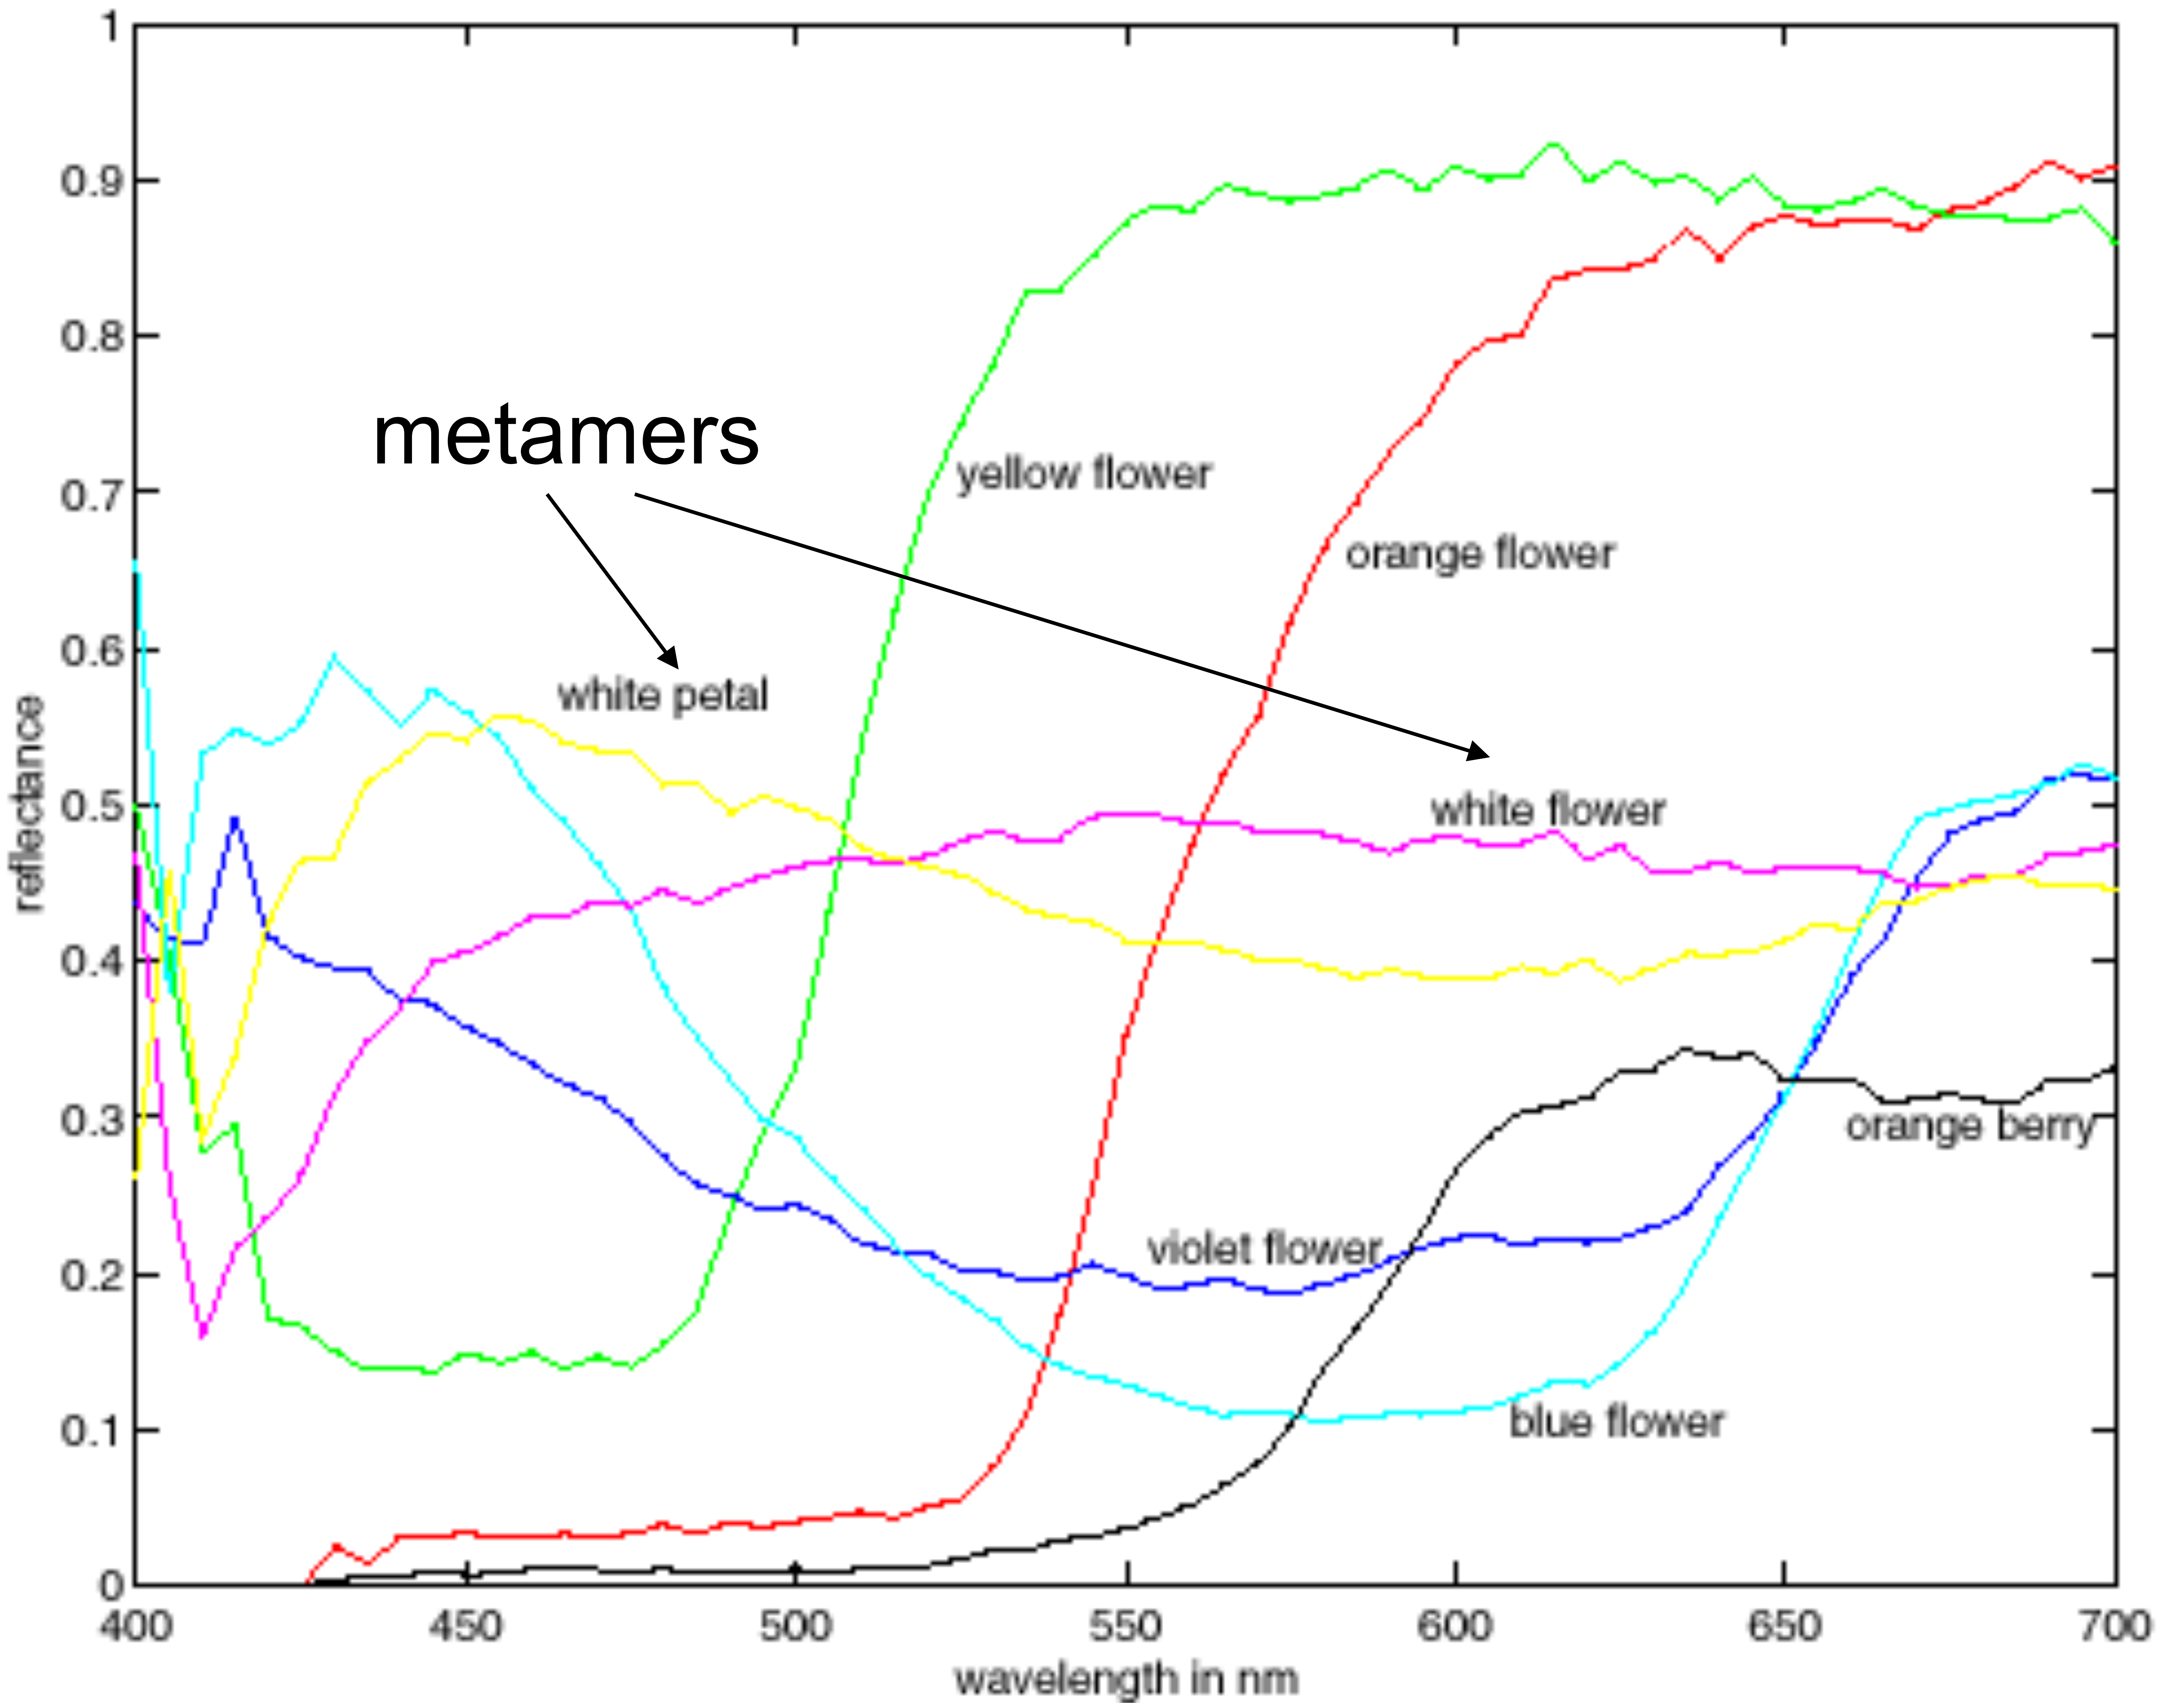

Physiology of color vision: fun facts

"M" and "L" pigments are encoded on the X-chromosome

- That's why men are more likely to be color blind
- "L" gene has high variation, so some women may be *tetra-chromatic*

Color blindness

- **Red-green color blindness** — mutation in L or M photoreceptors; difficulty in discriminating red and green hues
- **Blue-yellow color blindness** — mutation in S photoreceptors; difficulty in discriminating bluish and greenish hues, yellowish and reddish hues

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

More information @ http://en.wikipedia.org/wiki/Color_vision

Mantis Shrimp: Extraordinary Eyes

Homo sapiens

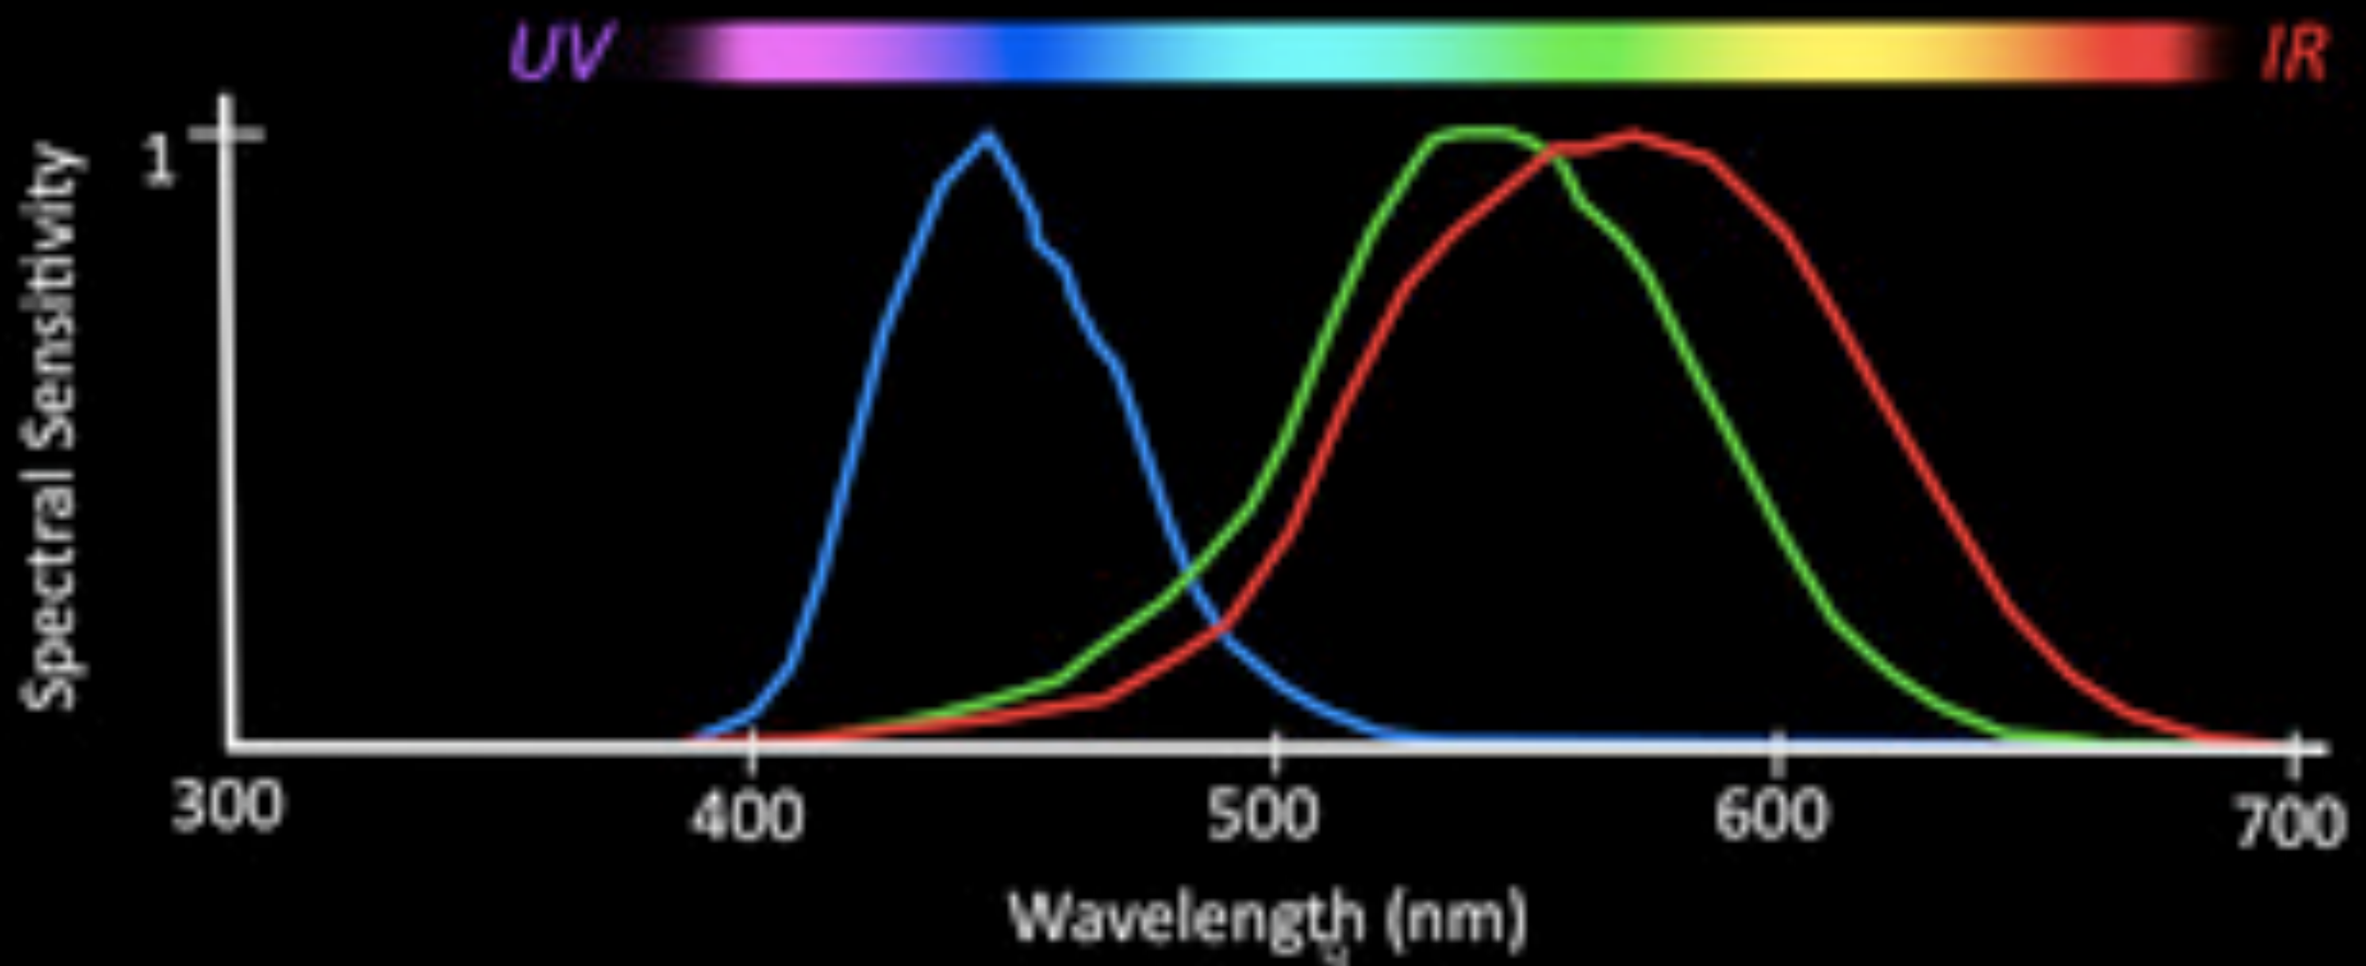

Neogonodactylus oestedii

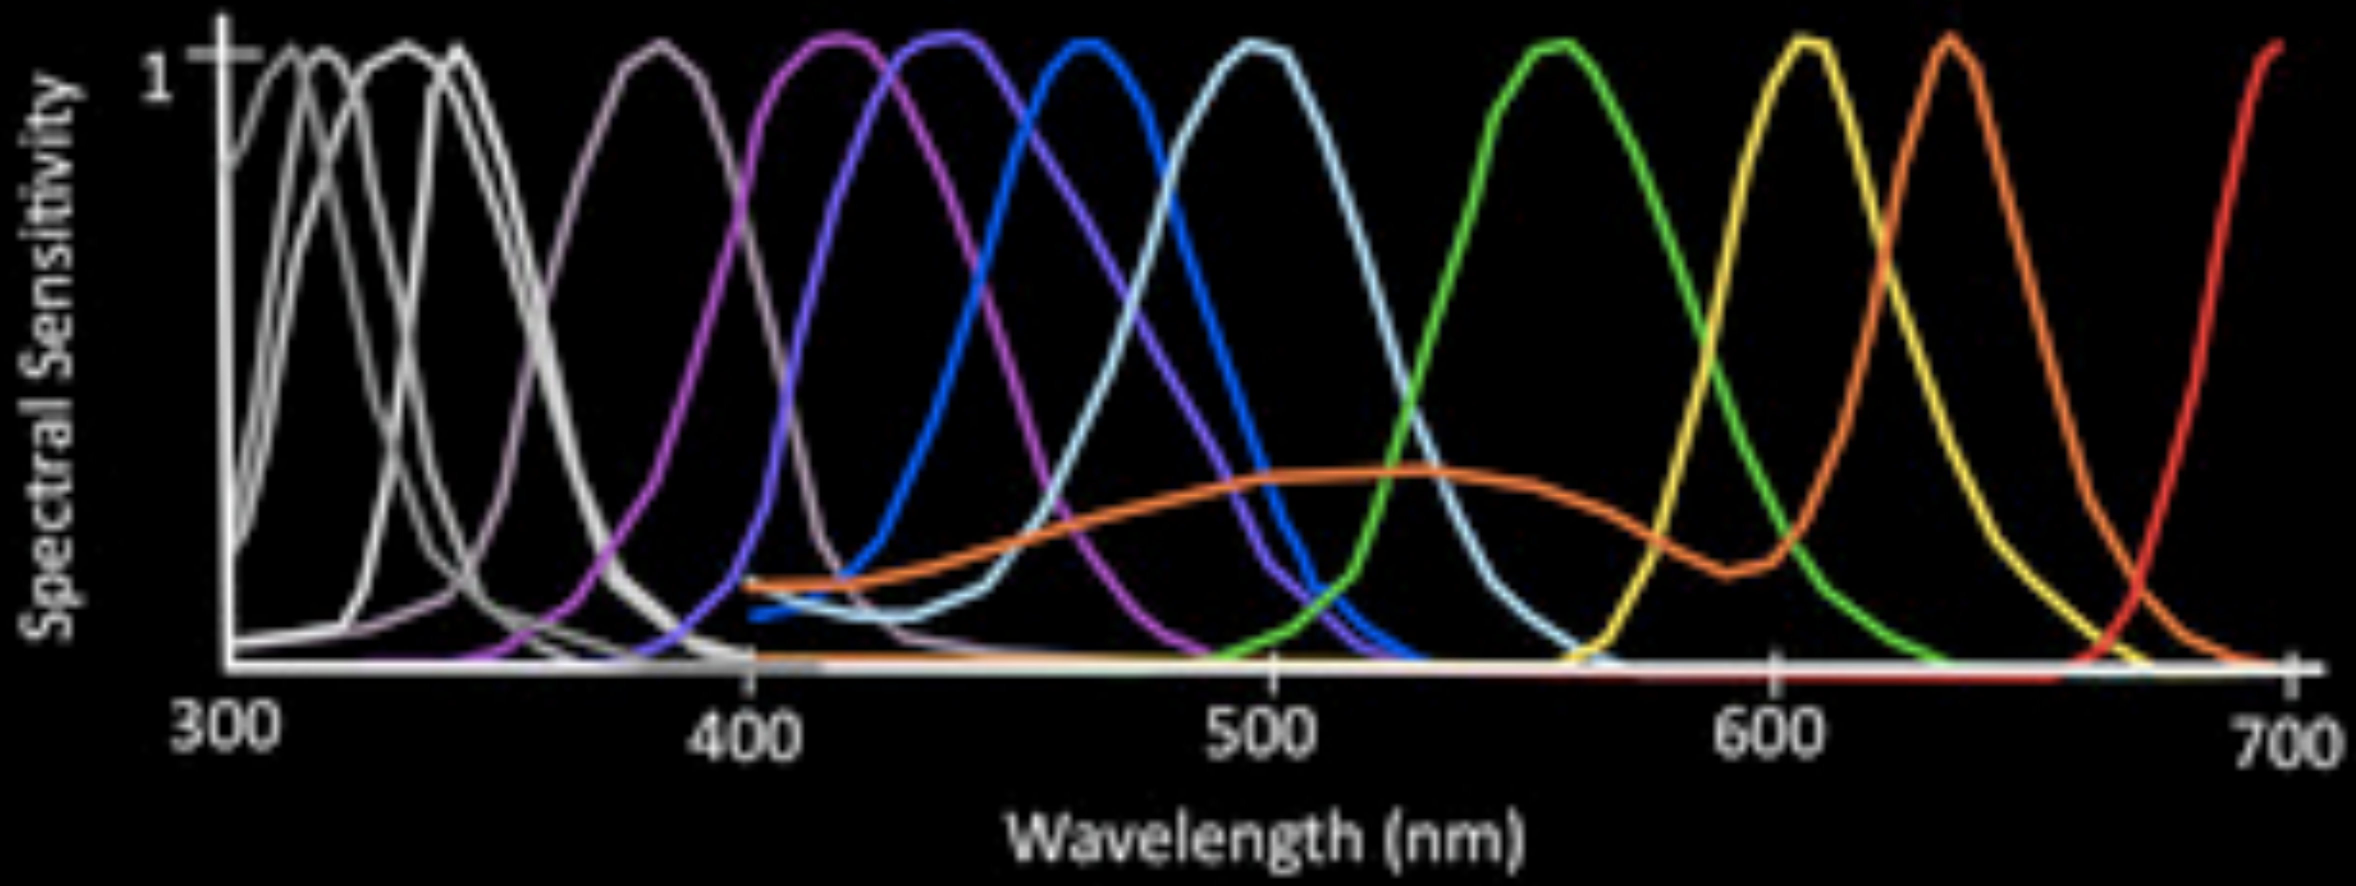

Marshall *et al.*, 2007; Marshall and Oberwinkler, 1999

How insects see

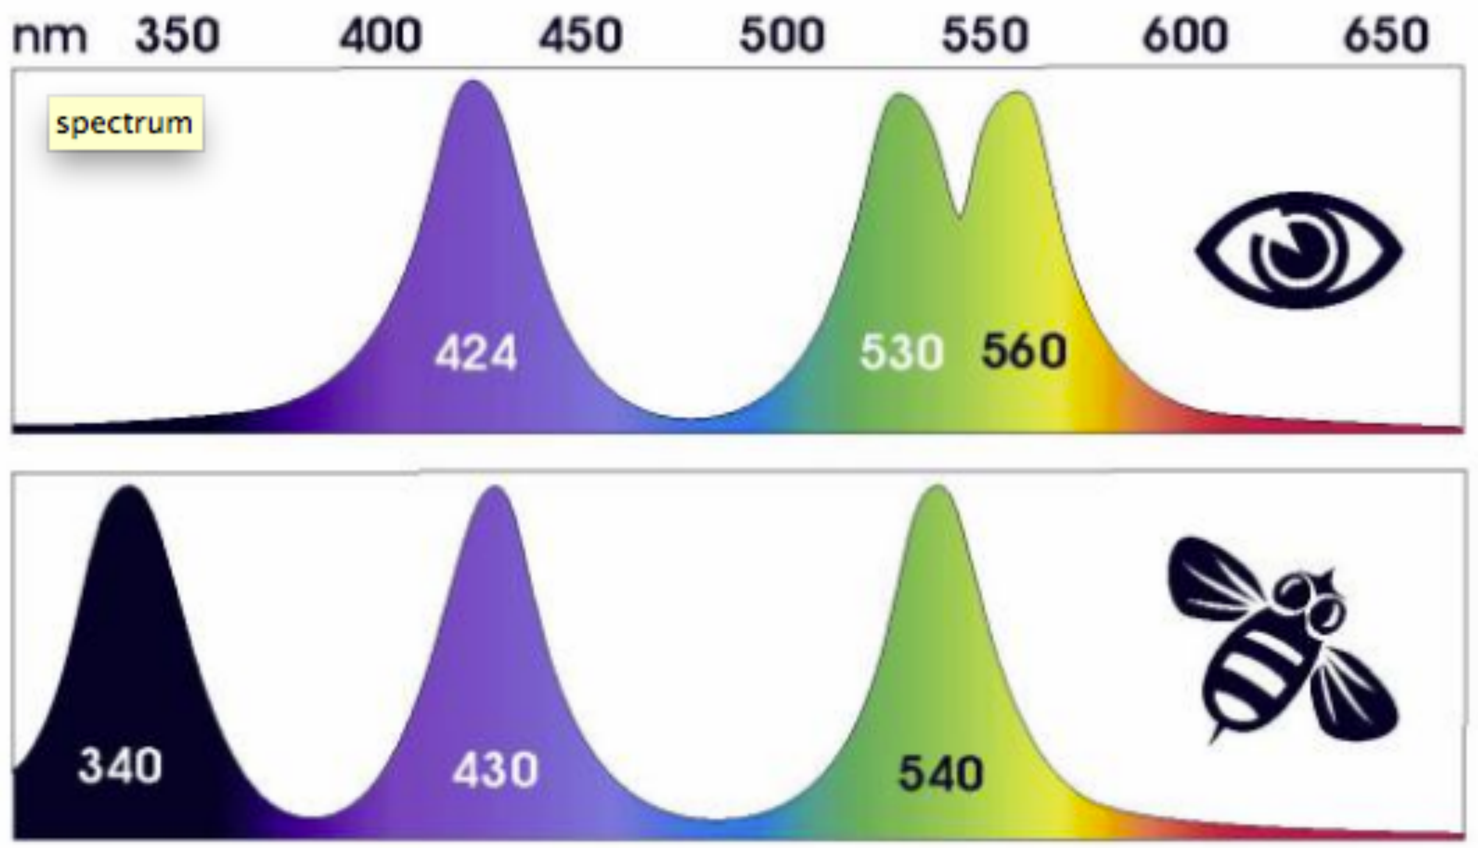

http://fieldguidetohummingbirds.wordpress.com/2008/11/11/do-we-see-what-bees-see/

visible light image

simulated bee vision

http://photographyoftheinvisibleworld.blogspot.de/

Standardizing color experience

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

Color matching experiments

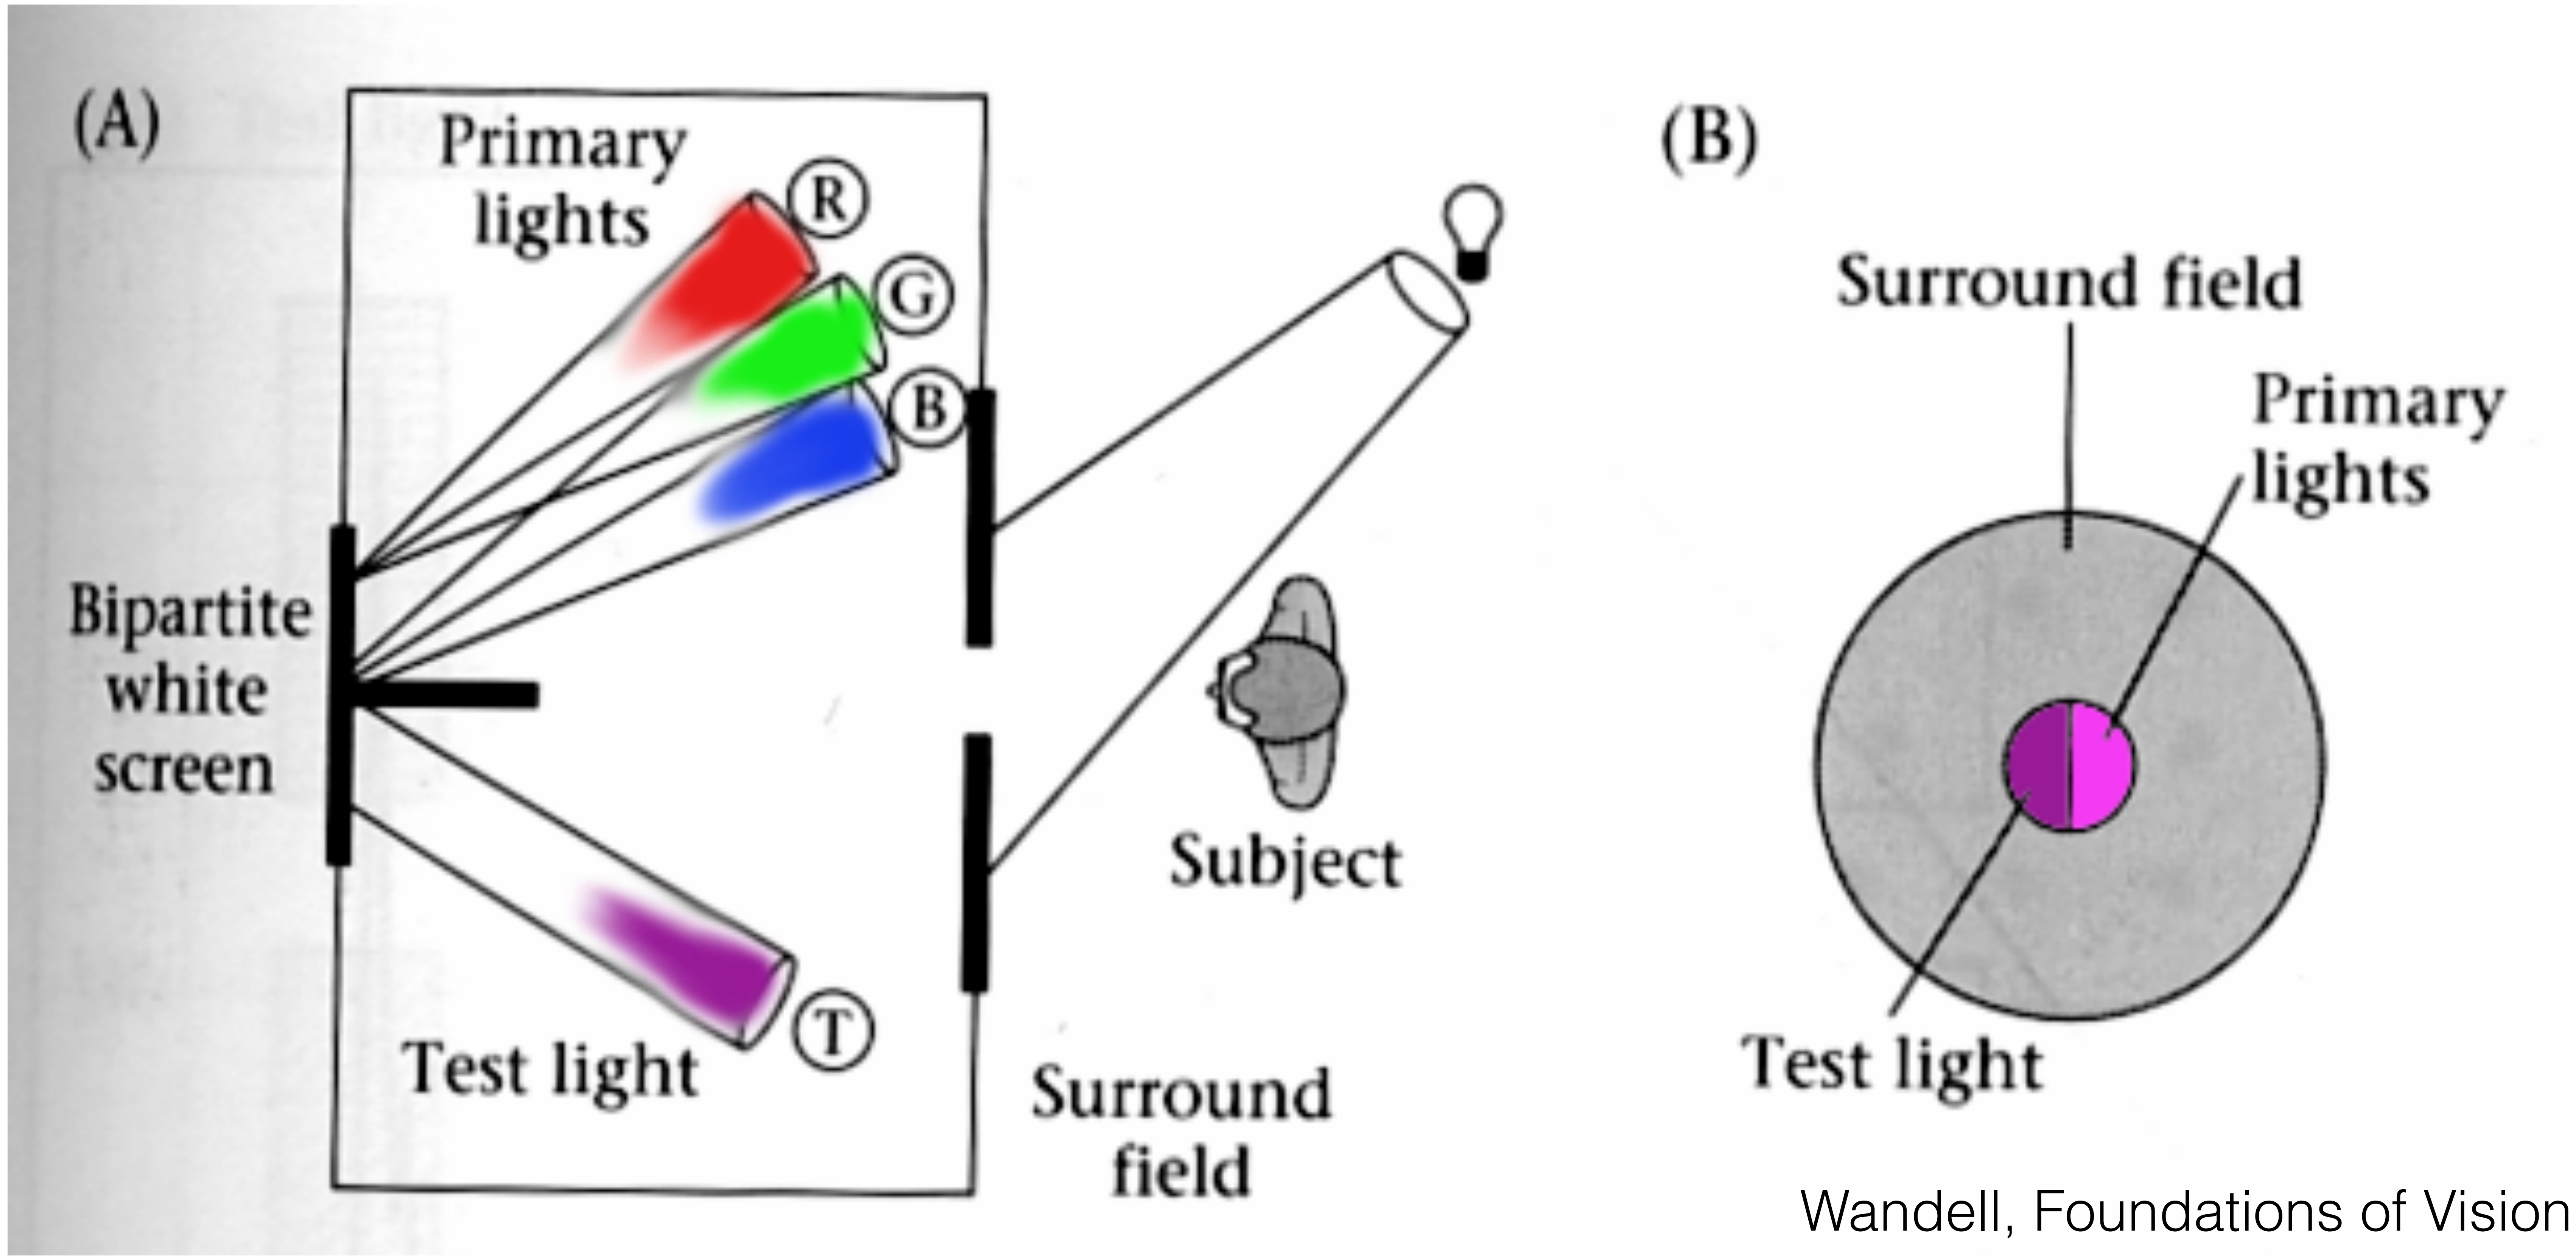

Wandell, Foundations of Vision, 1995

Color matching experiment 1

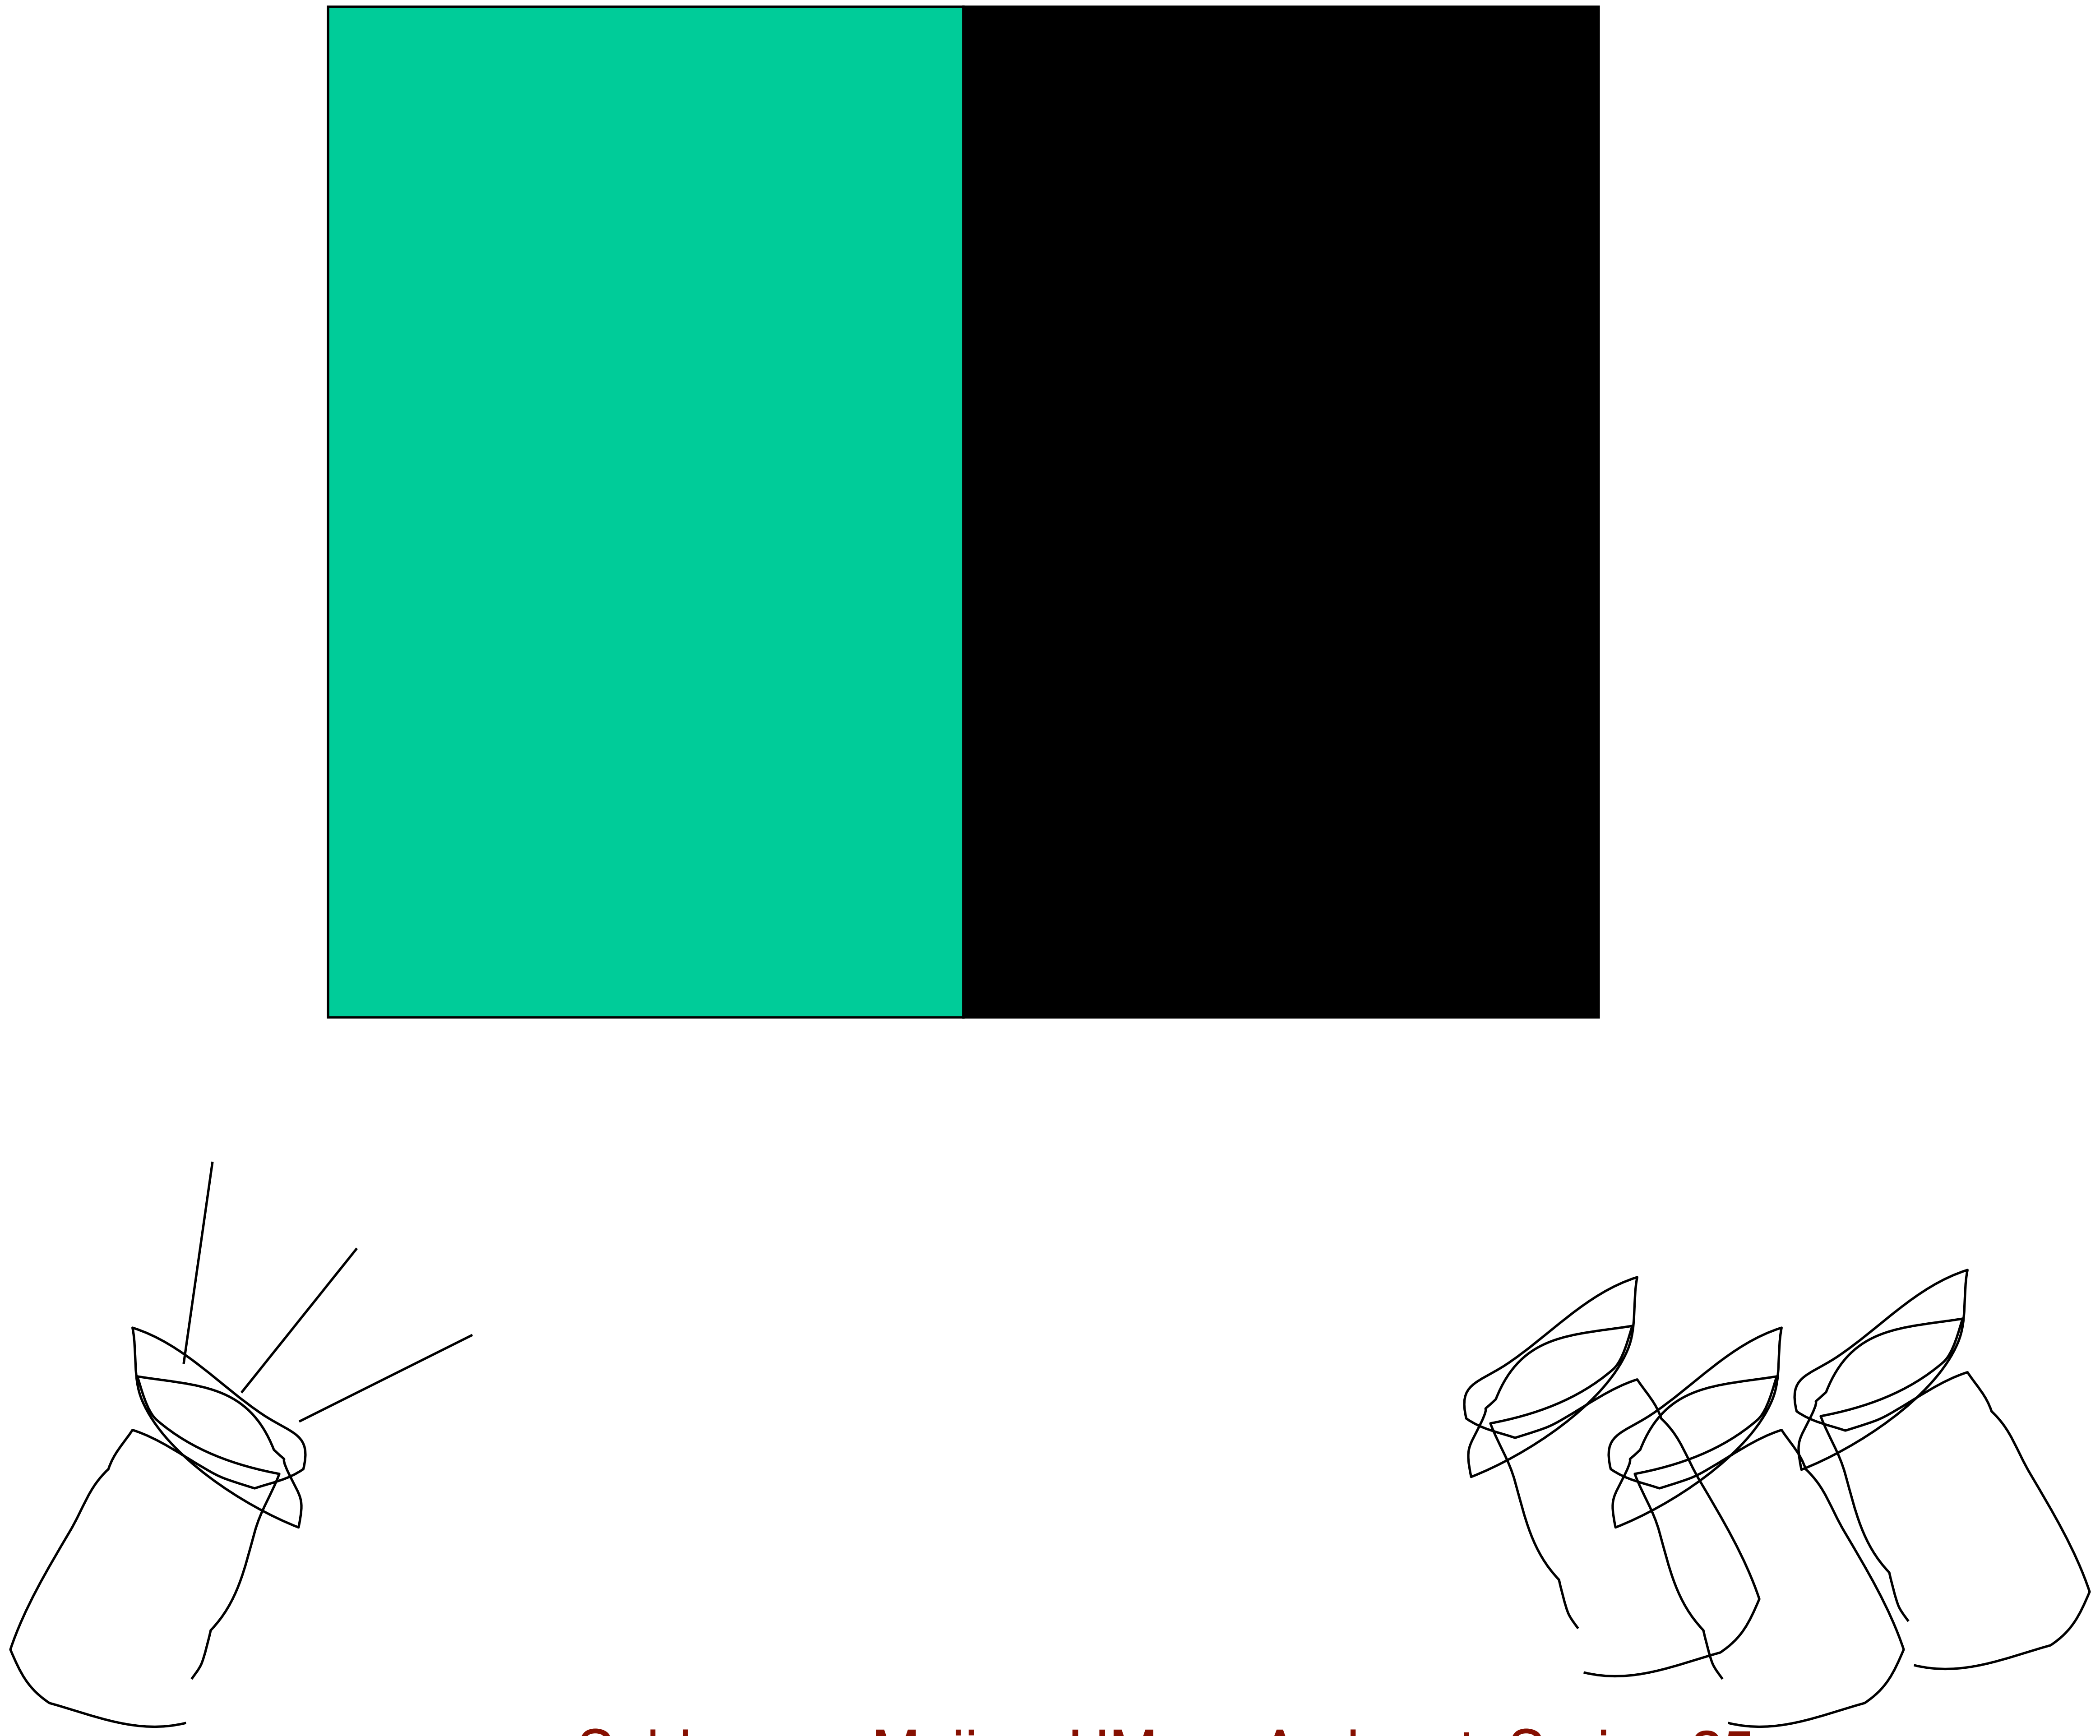

Source: W. Freeman

Color matching experiment 1

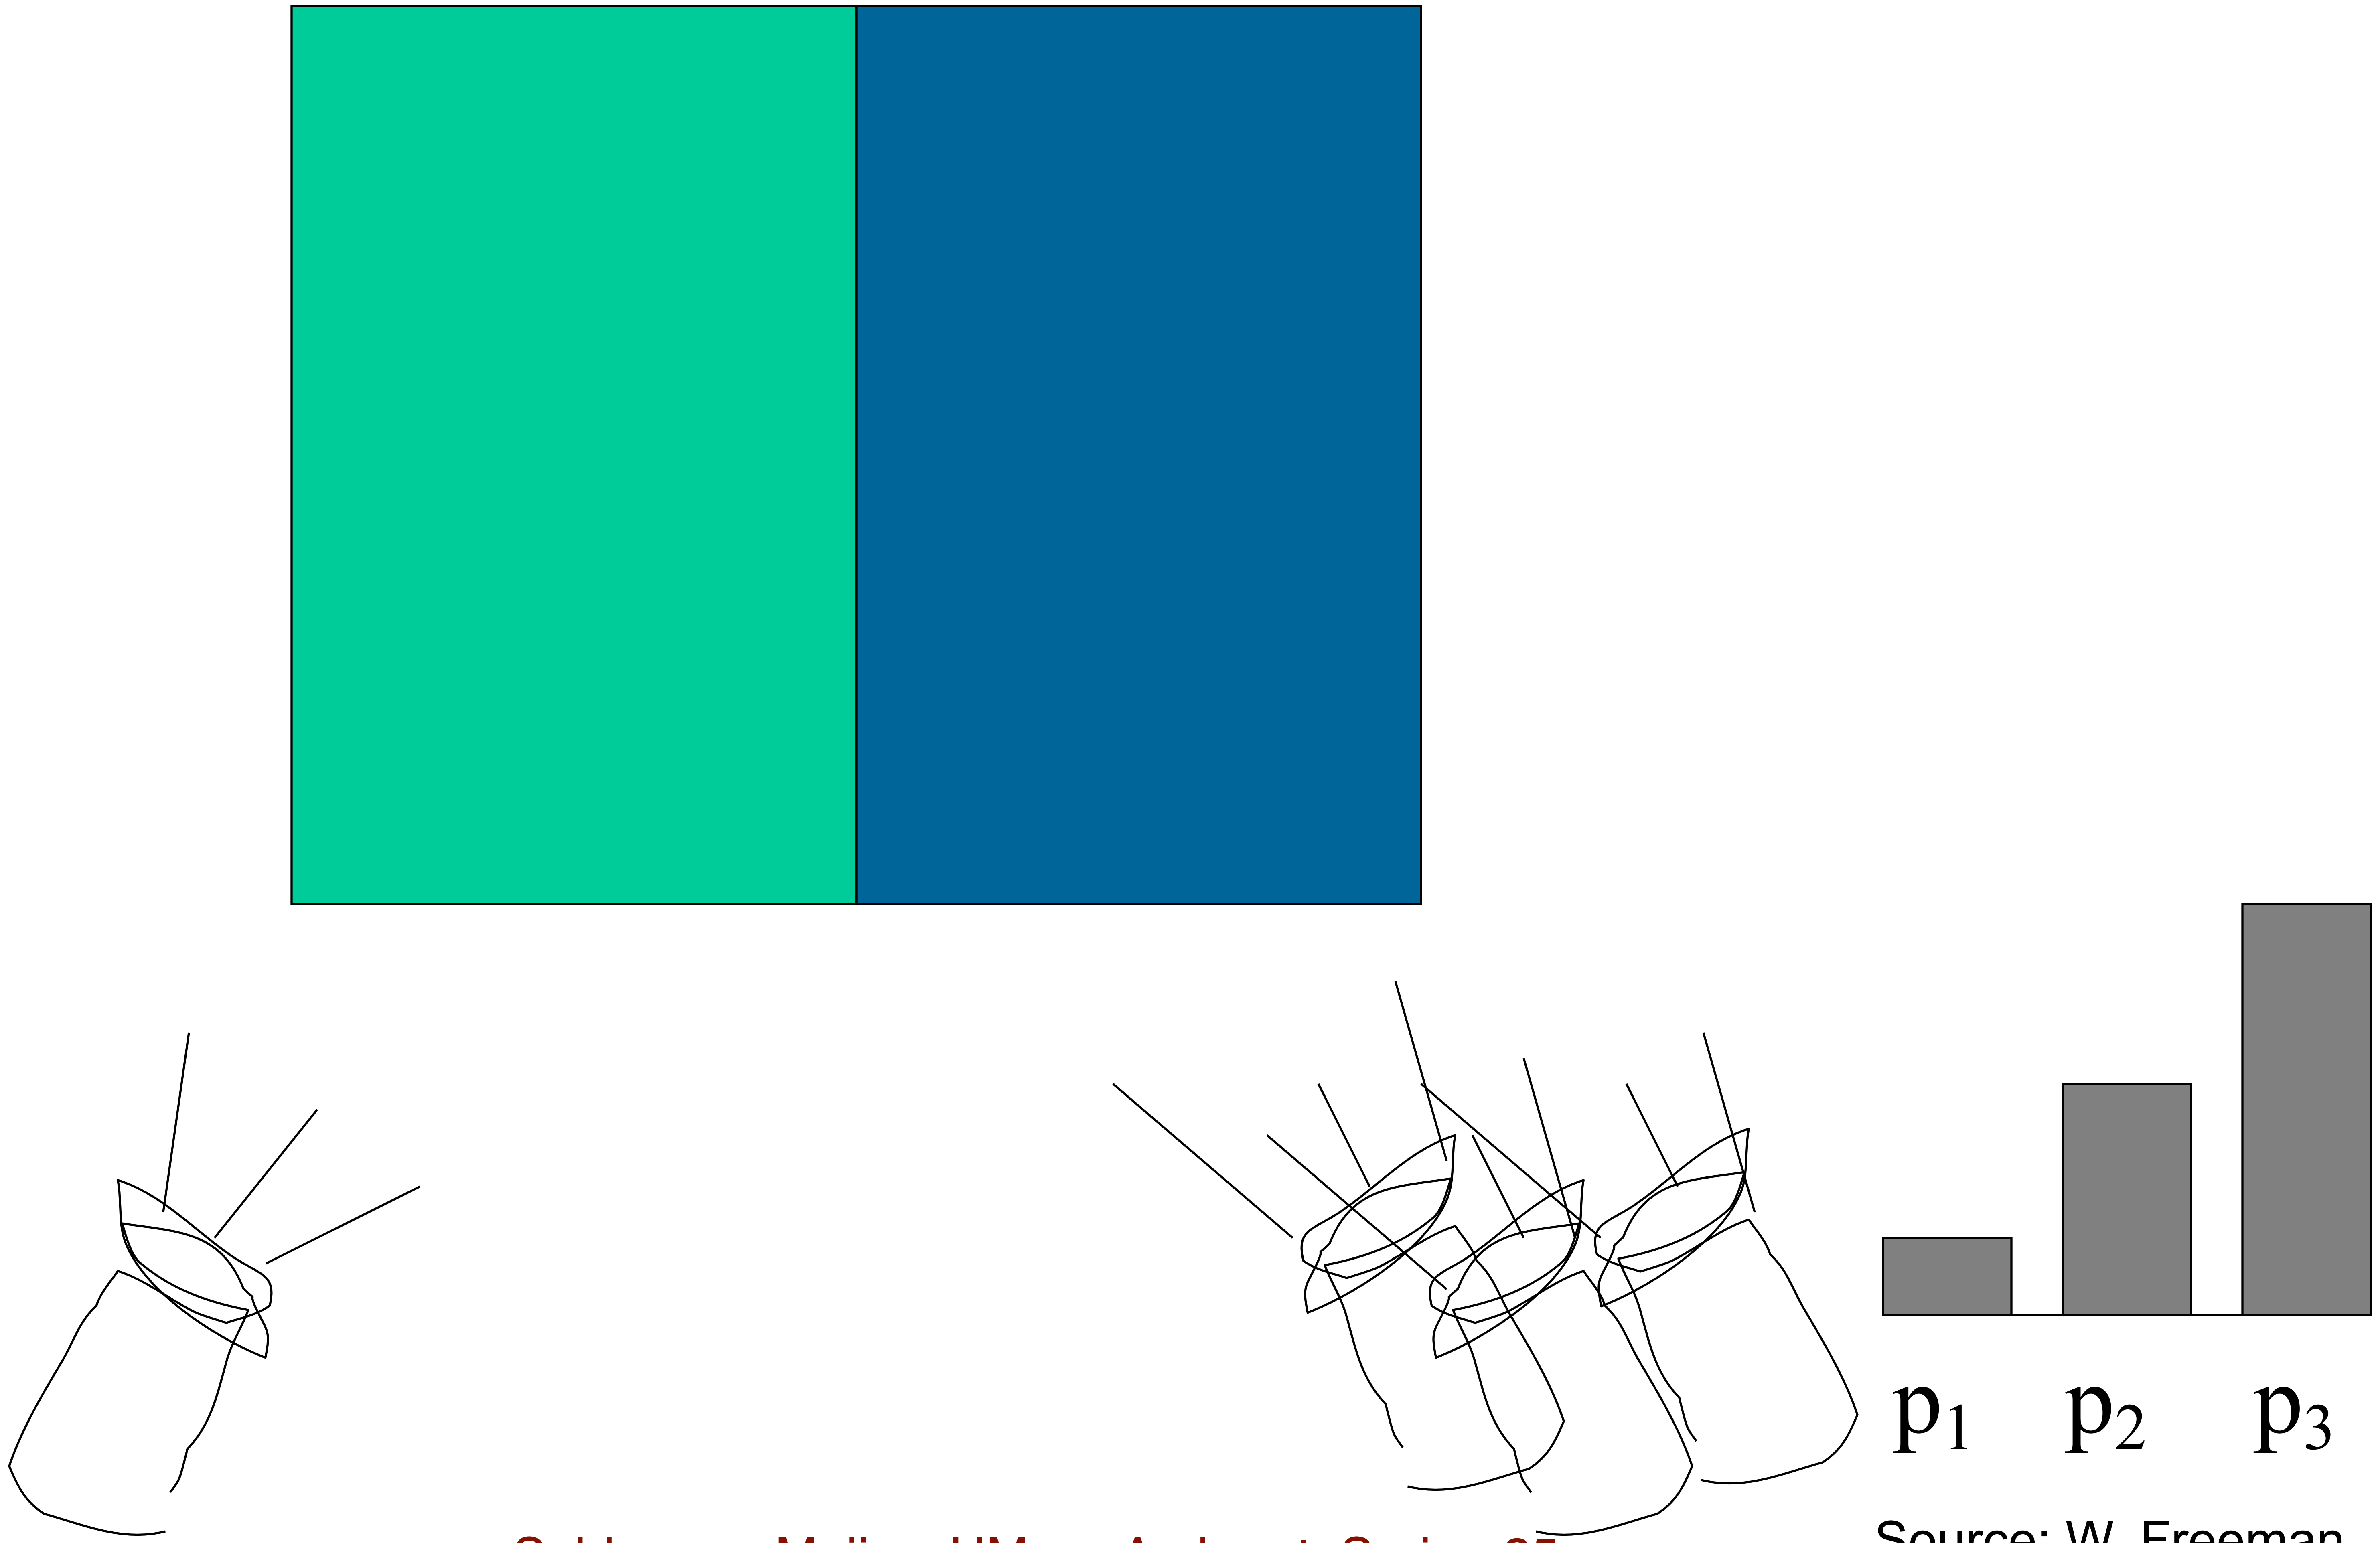

Source: W. Freeman

Color matching experiment 1

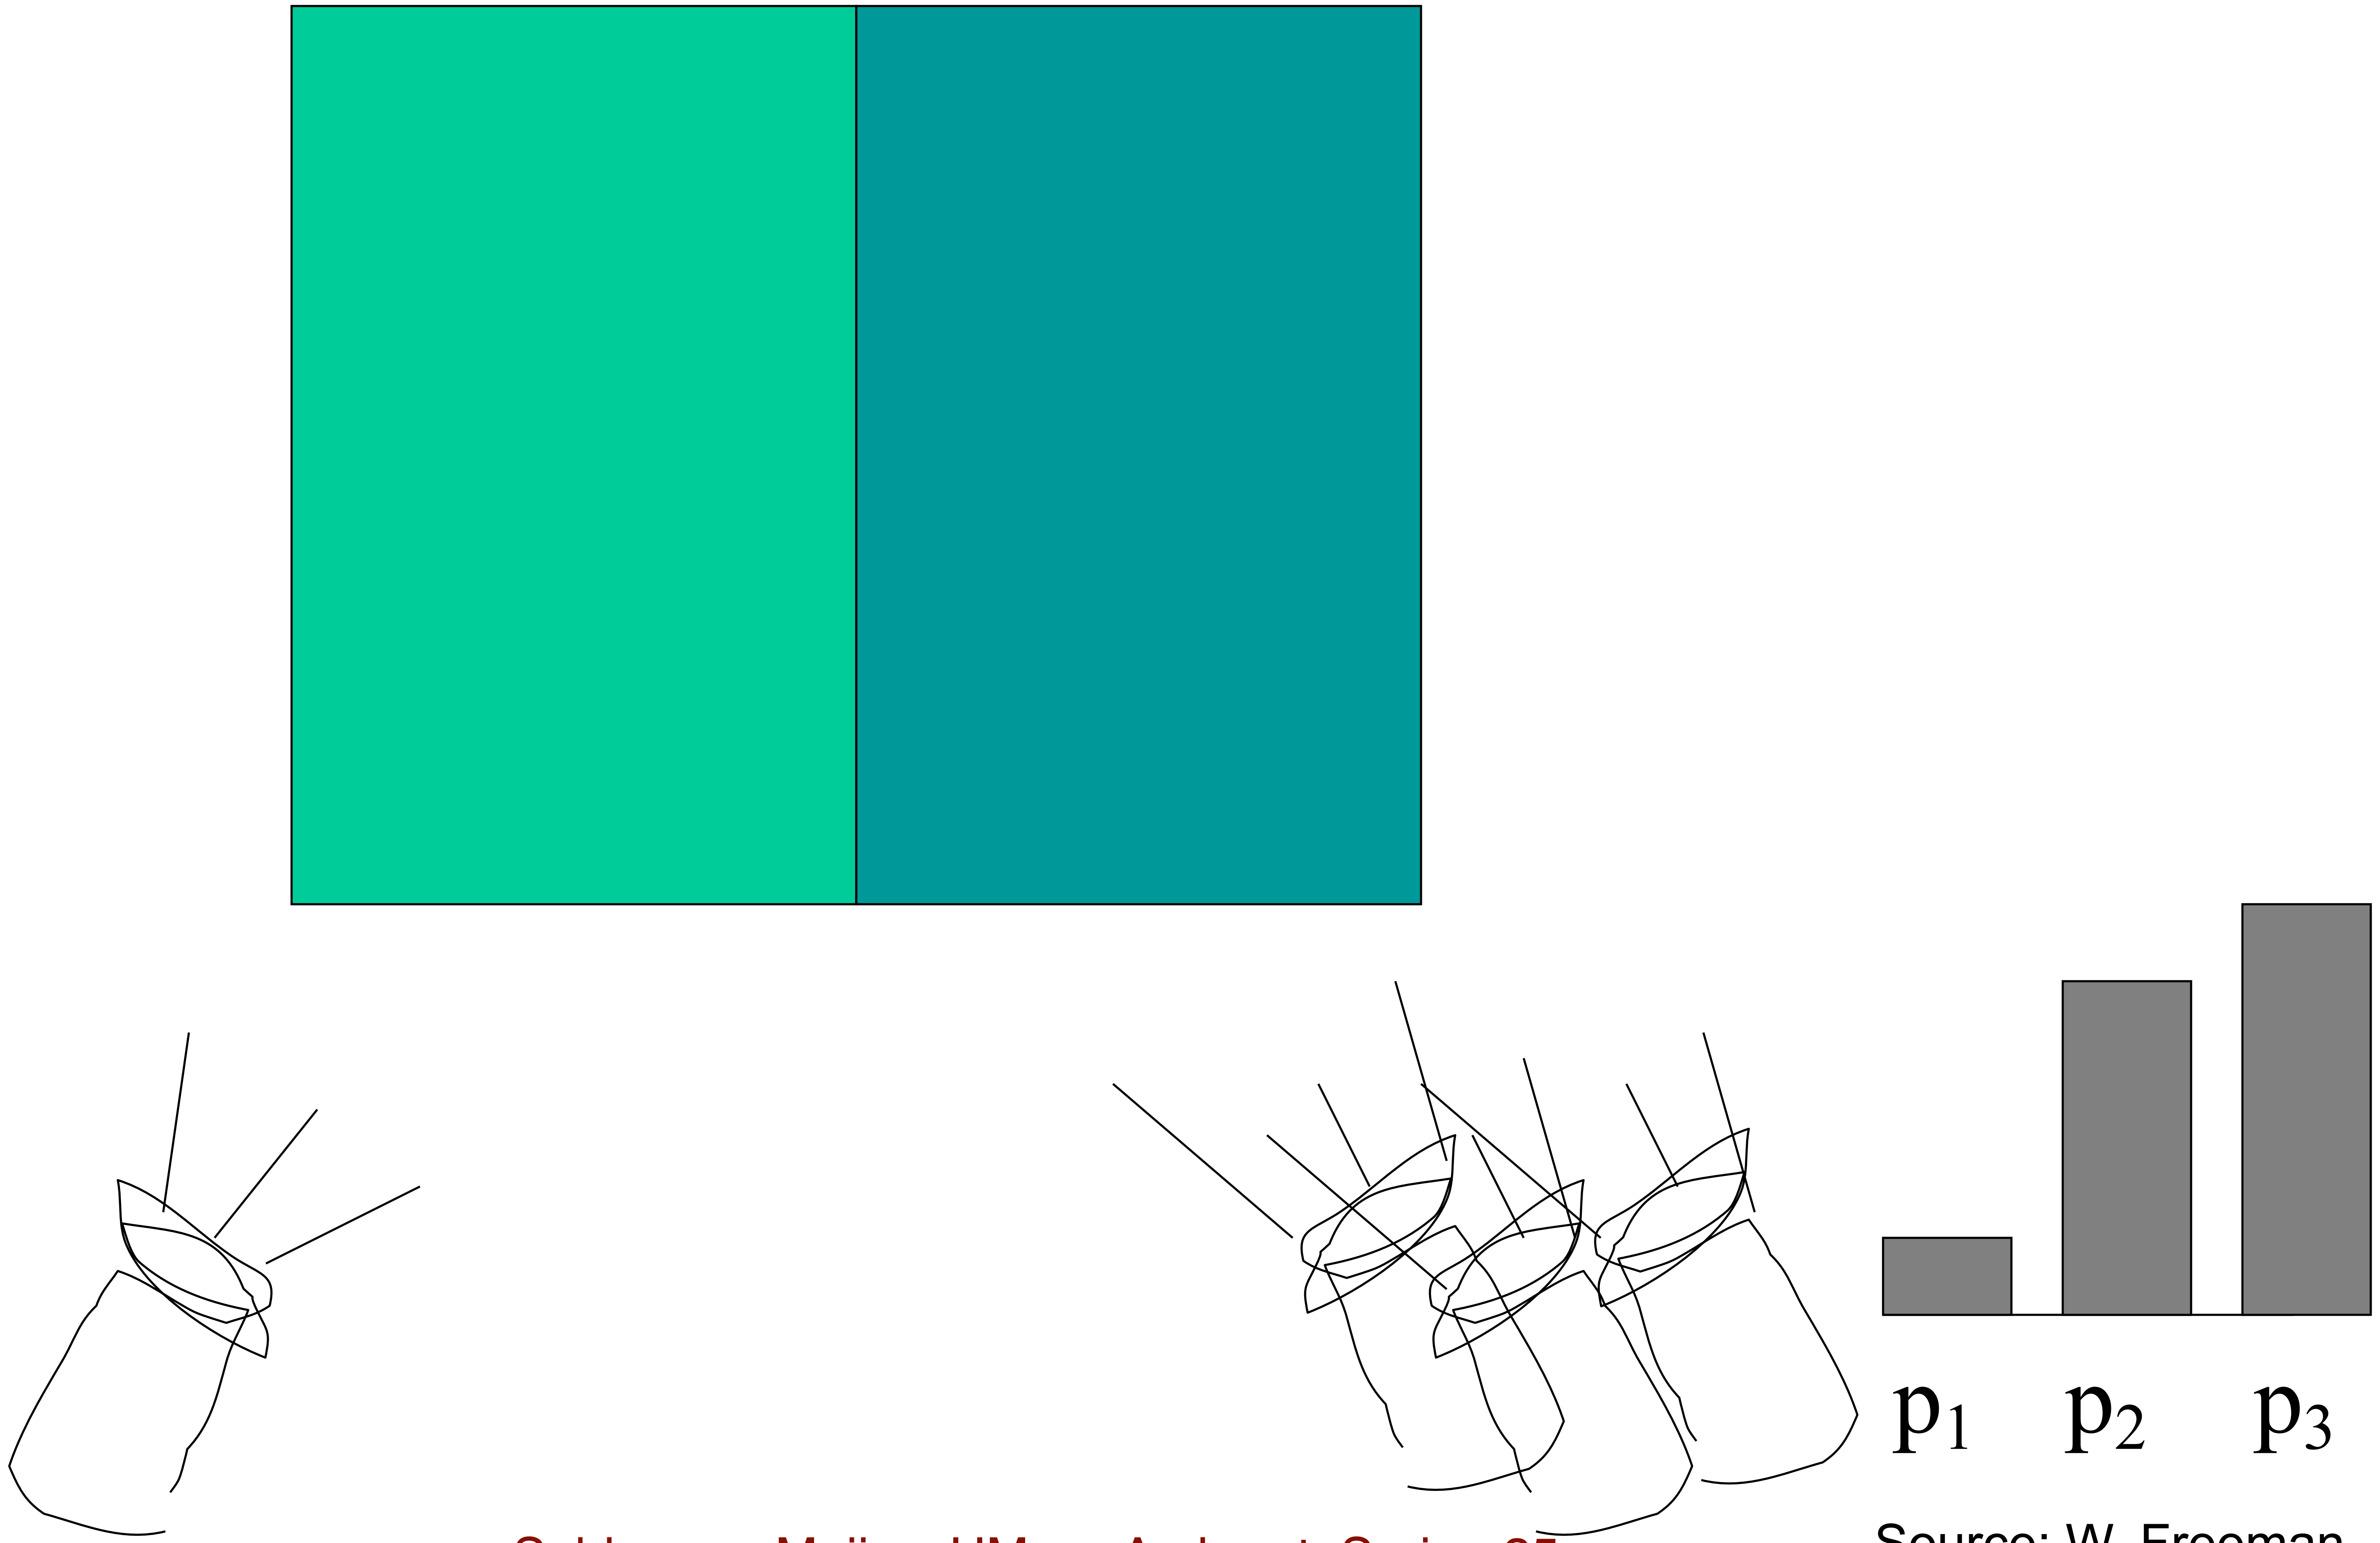

Source: W. Freeman

Color matching experiment 1

The primary color amounts needed for a match

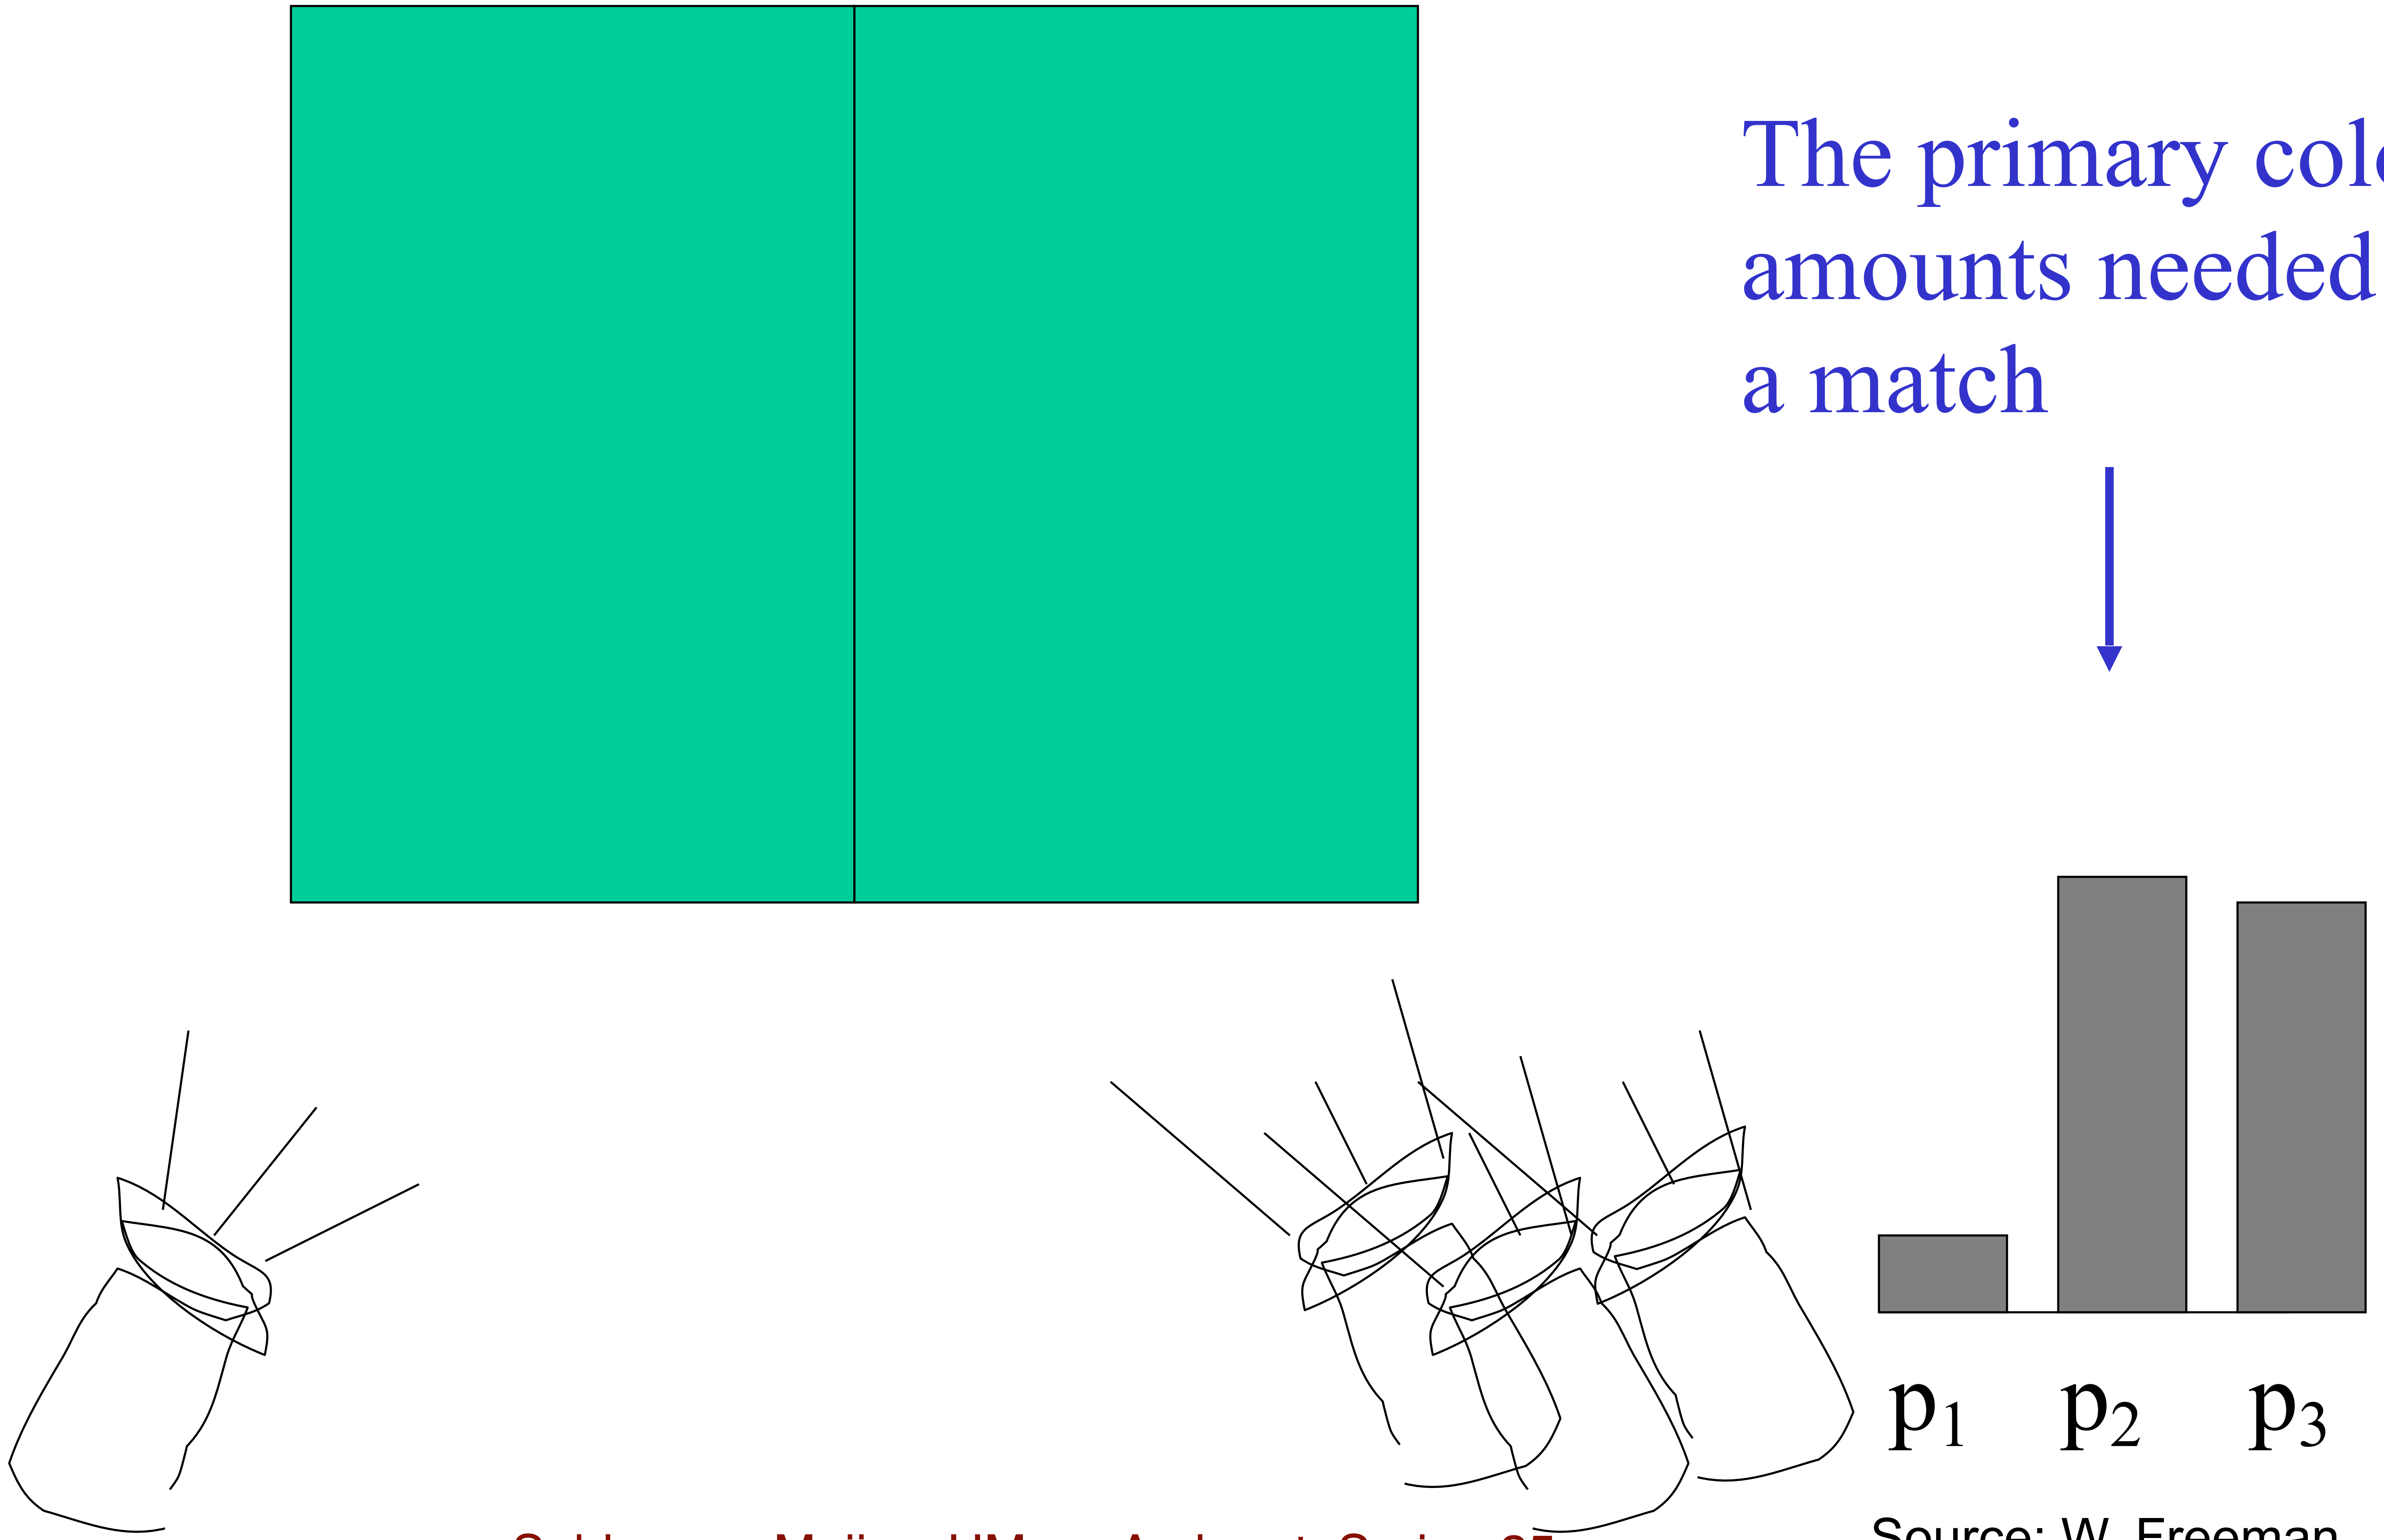

Source: W. Freeman

Color matching experiment 2

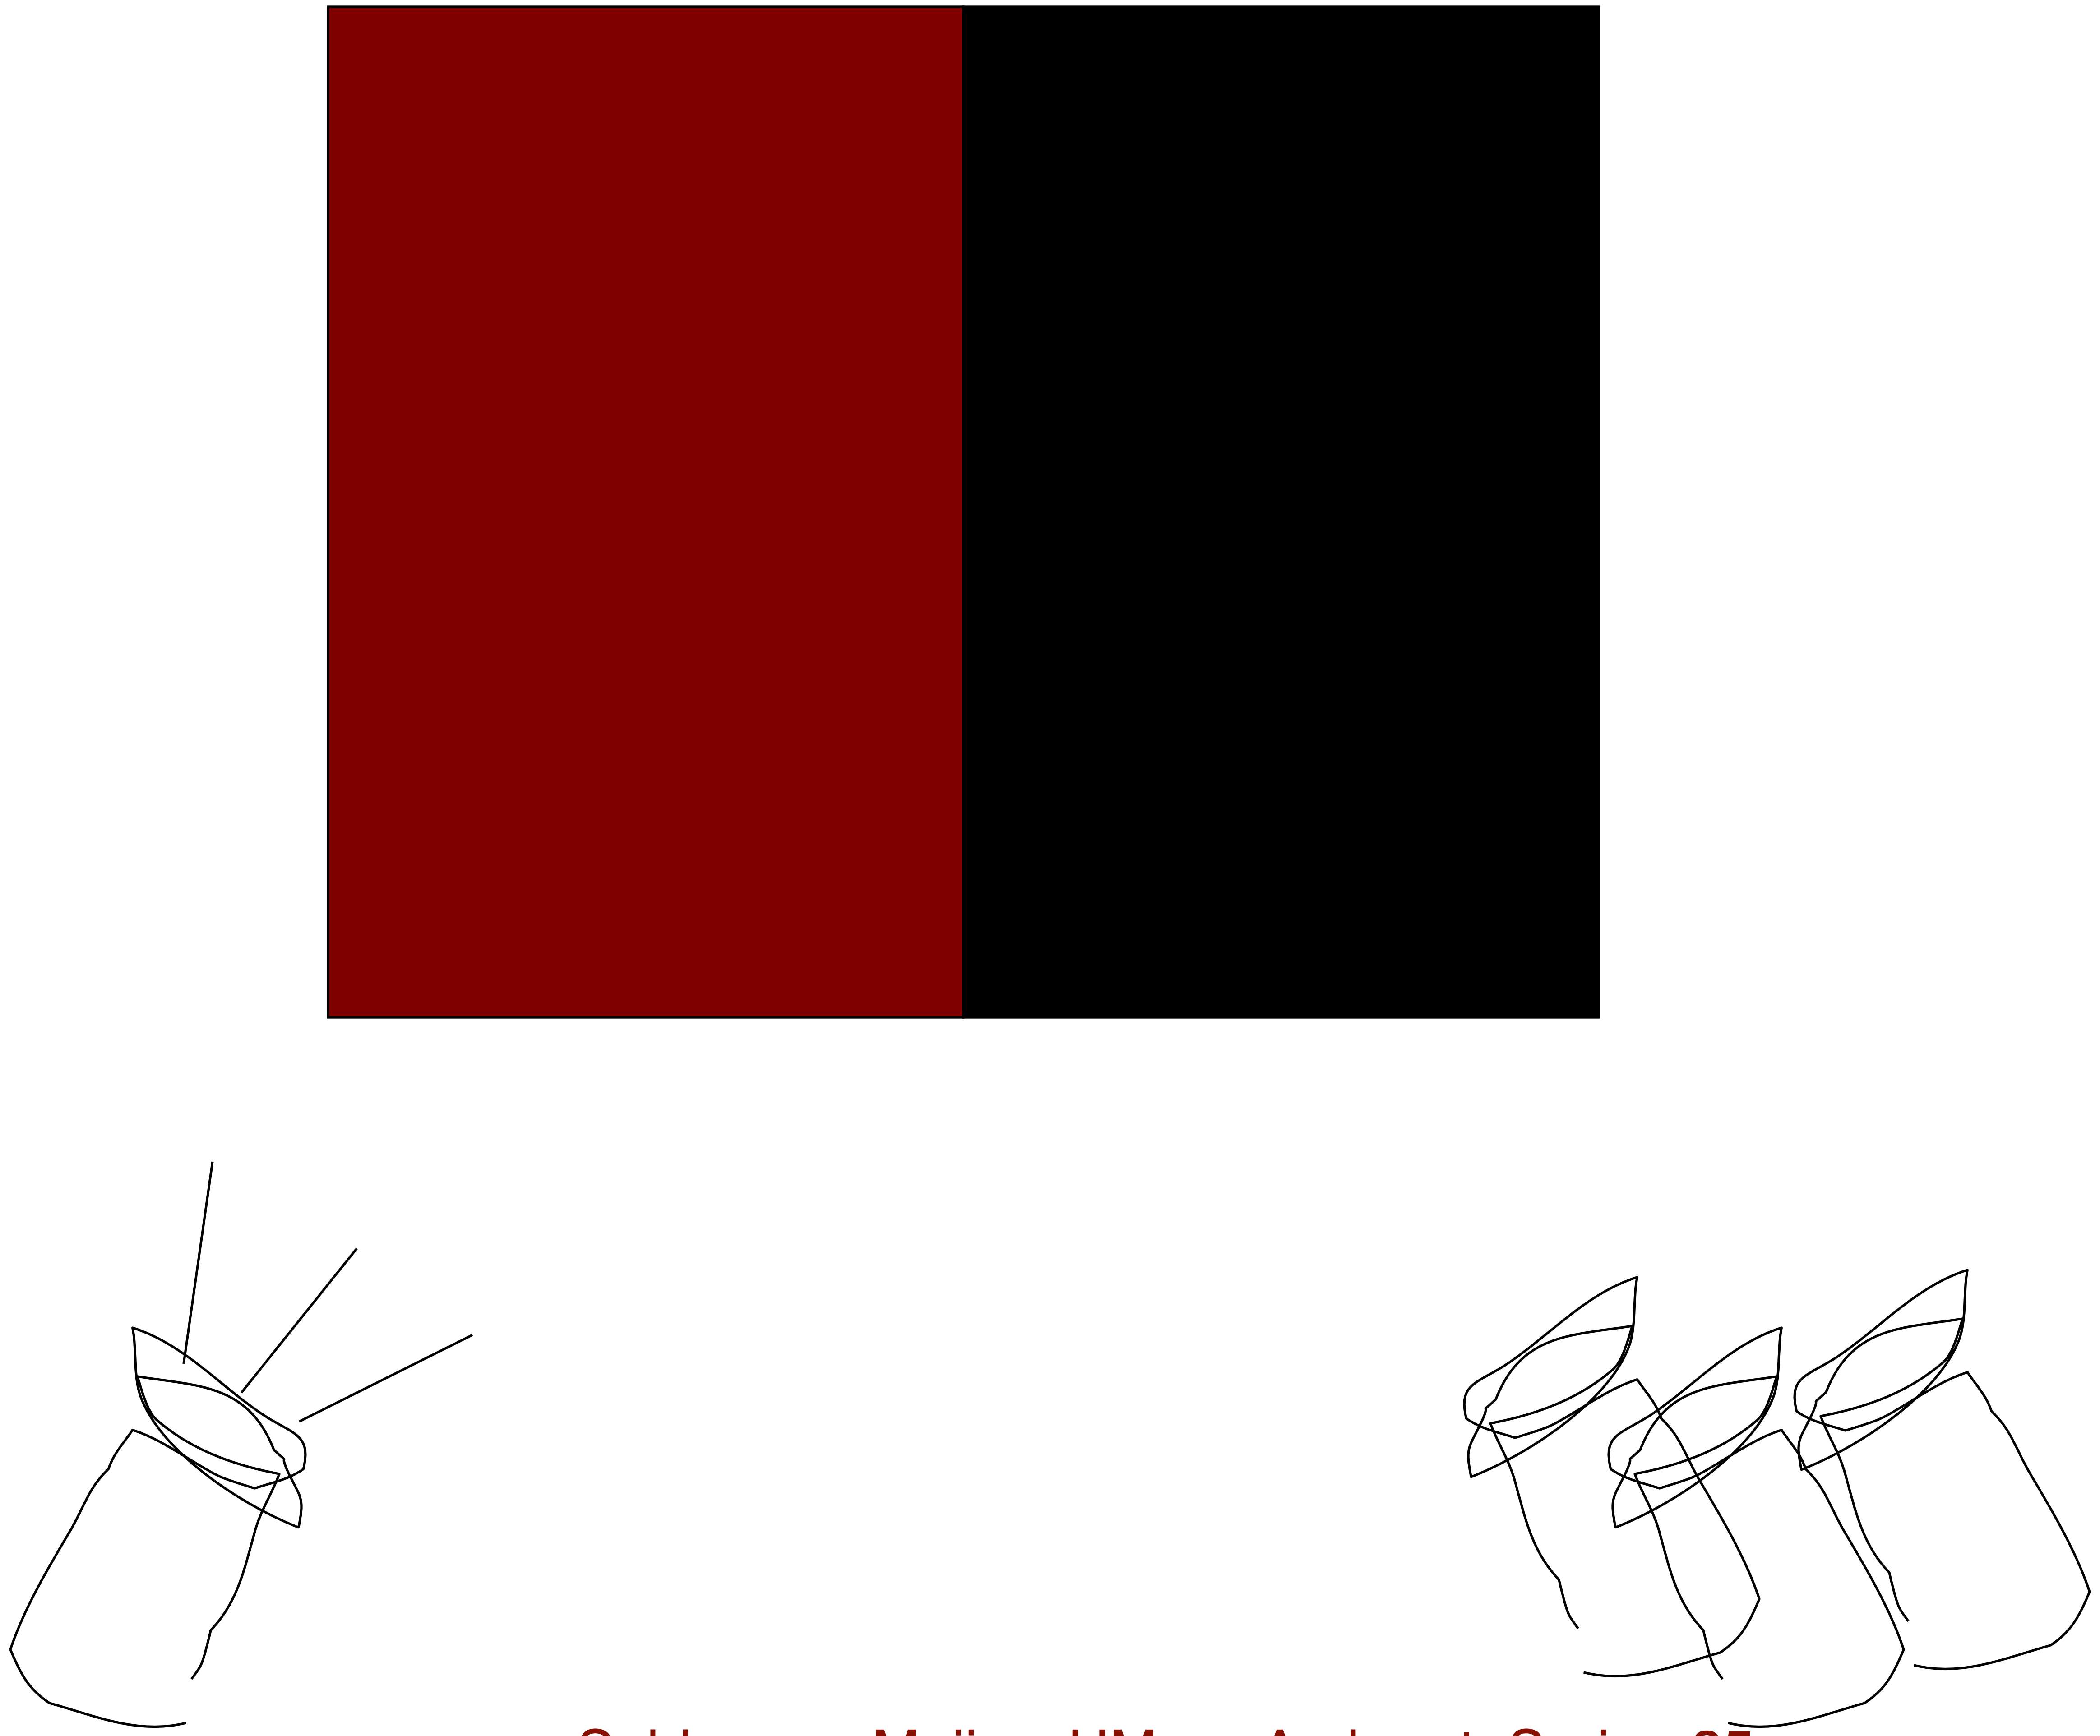

Source: W. Freeman

Color matching experiment 2

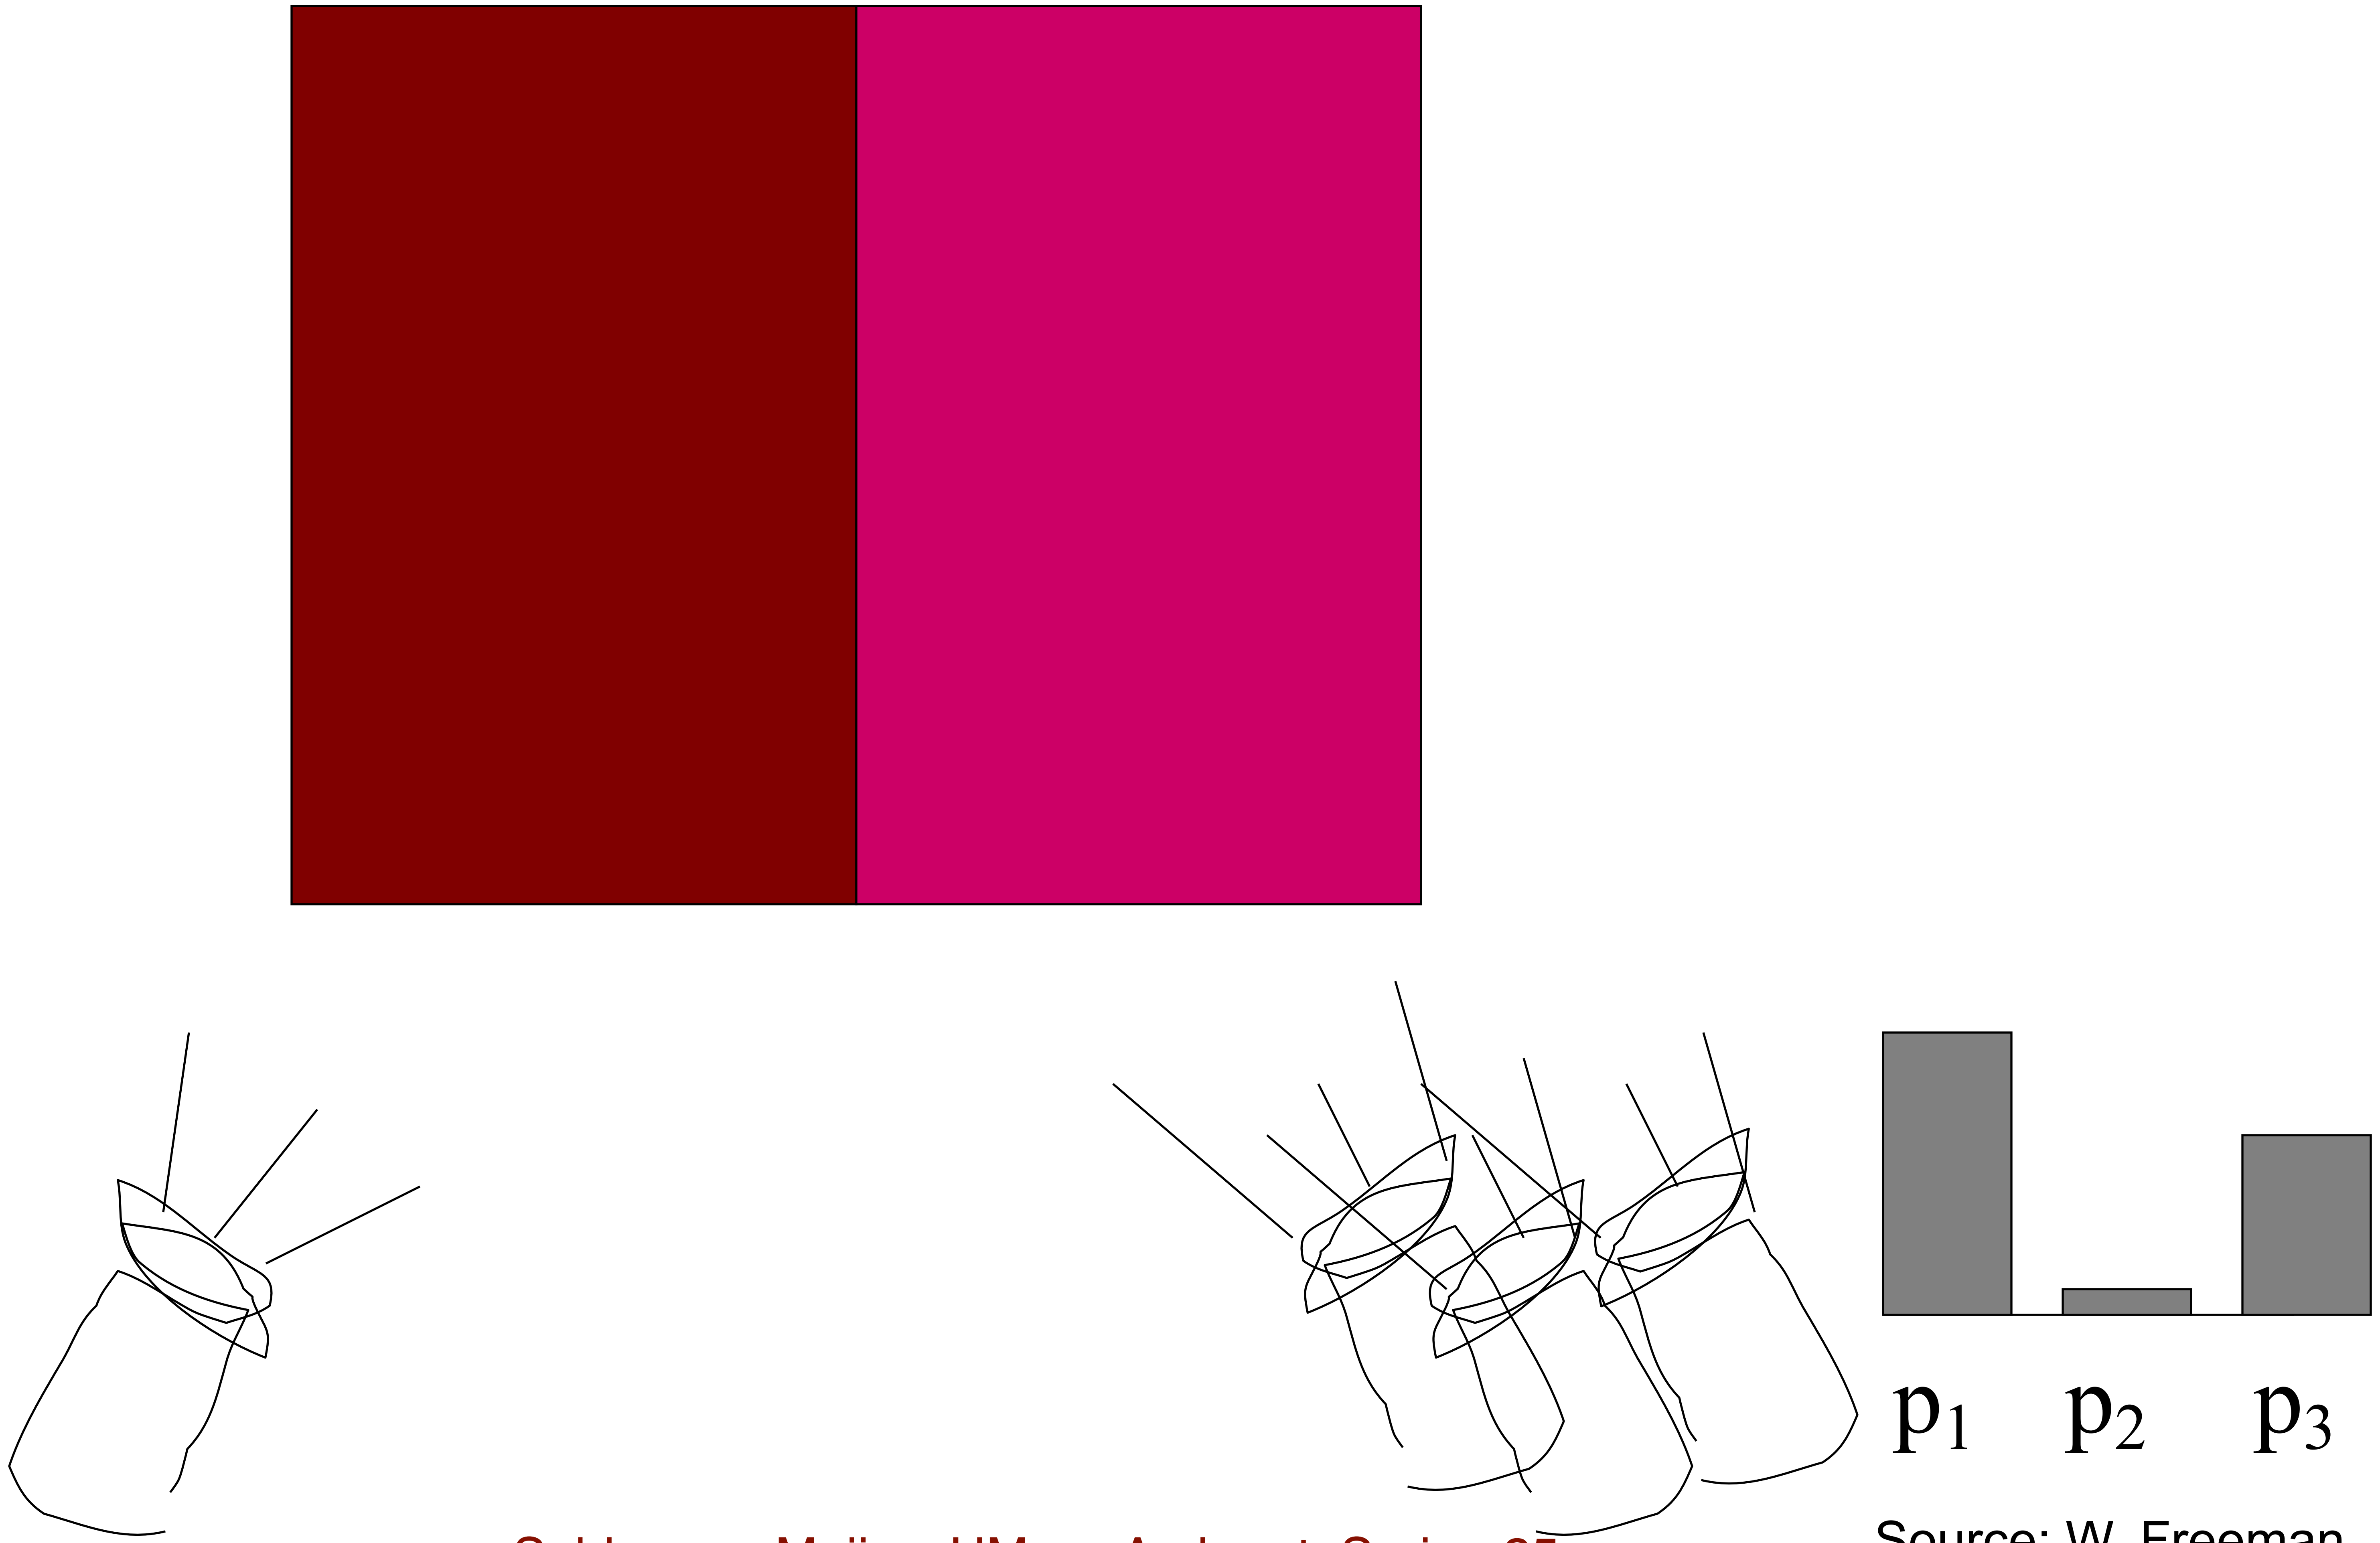

Source: W. Freeman

Color matching experiment 2

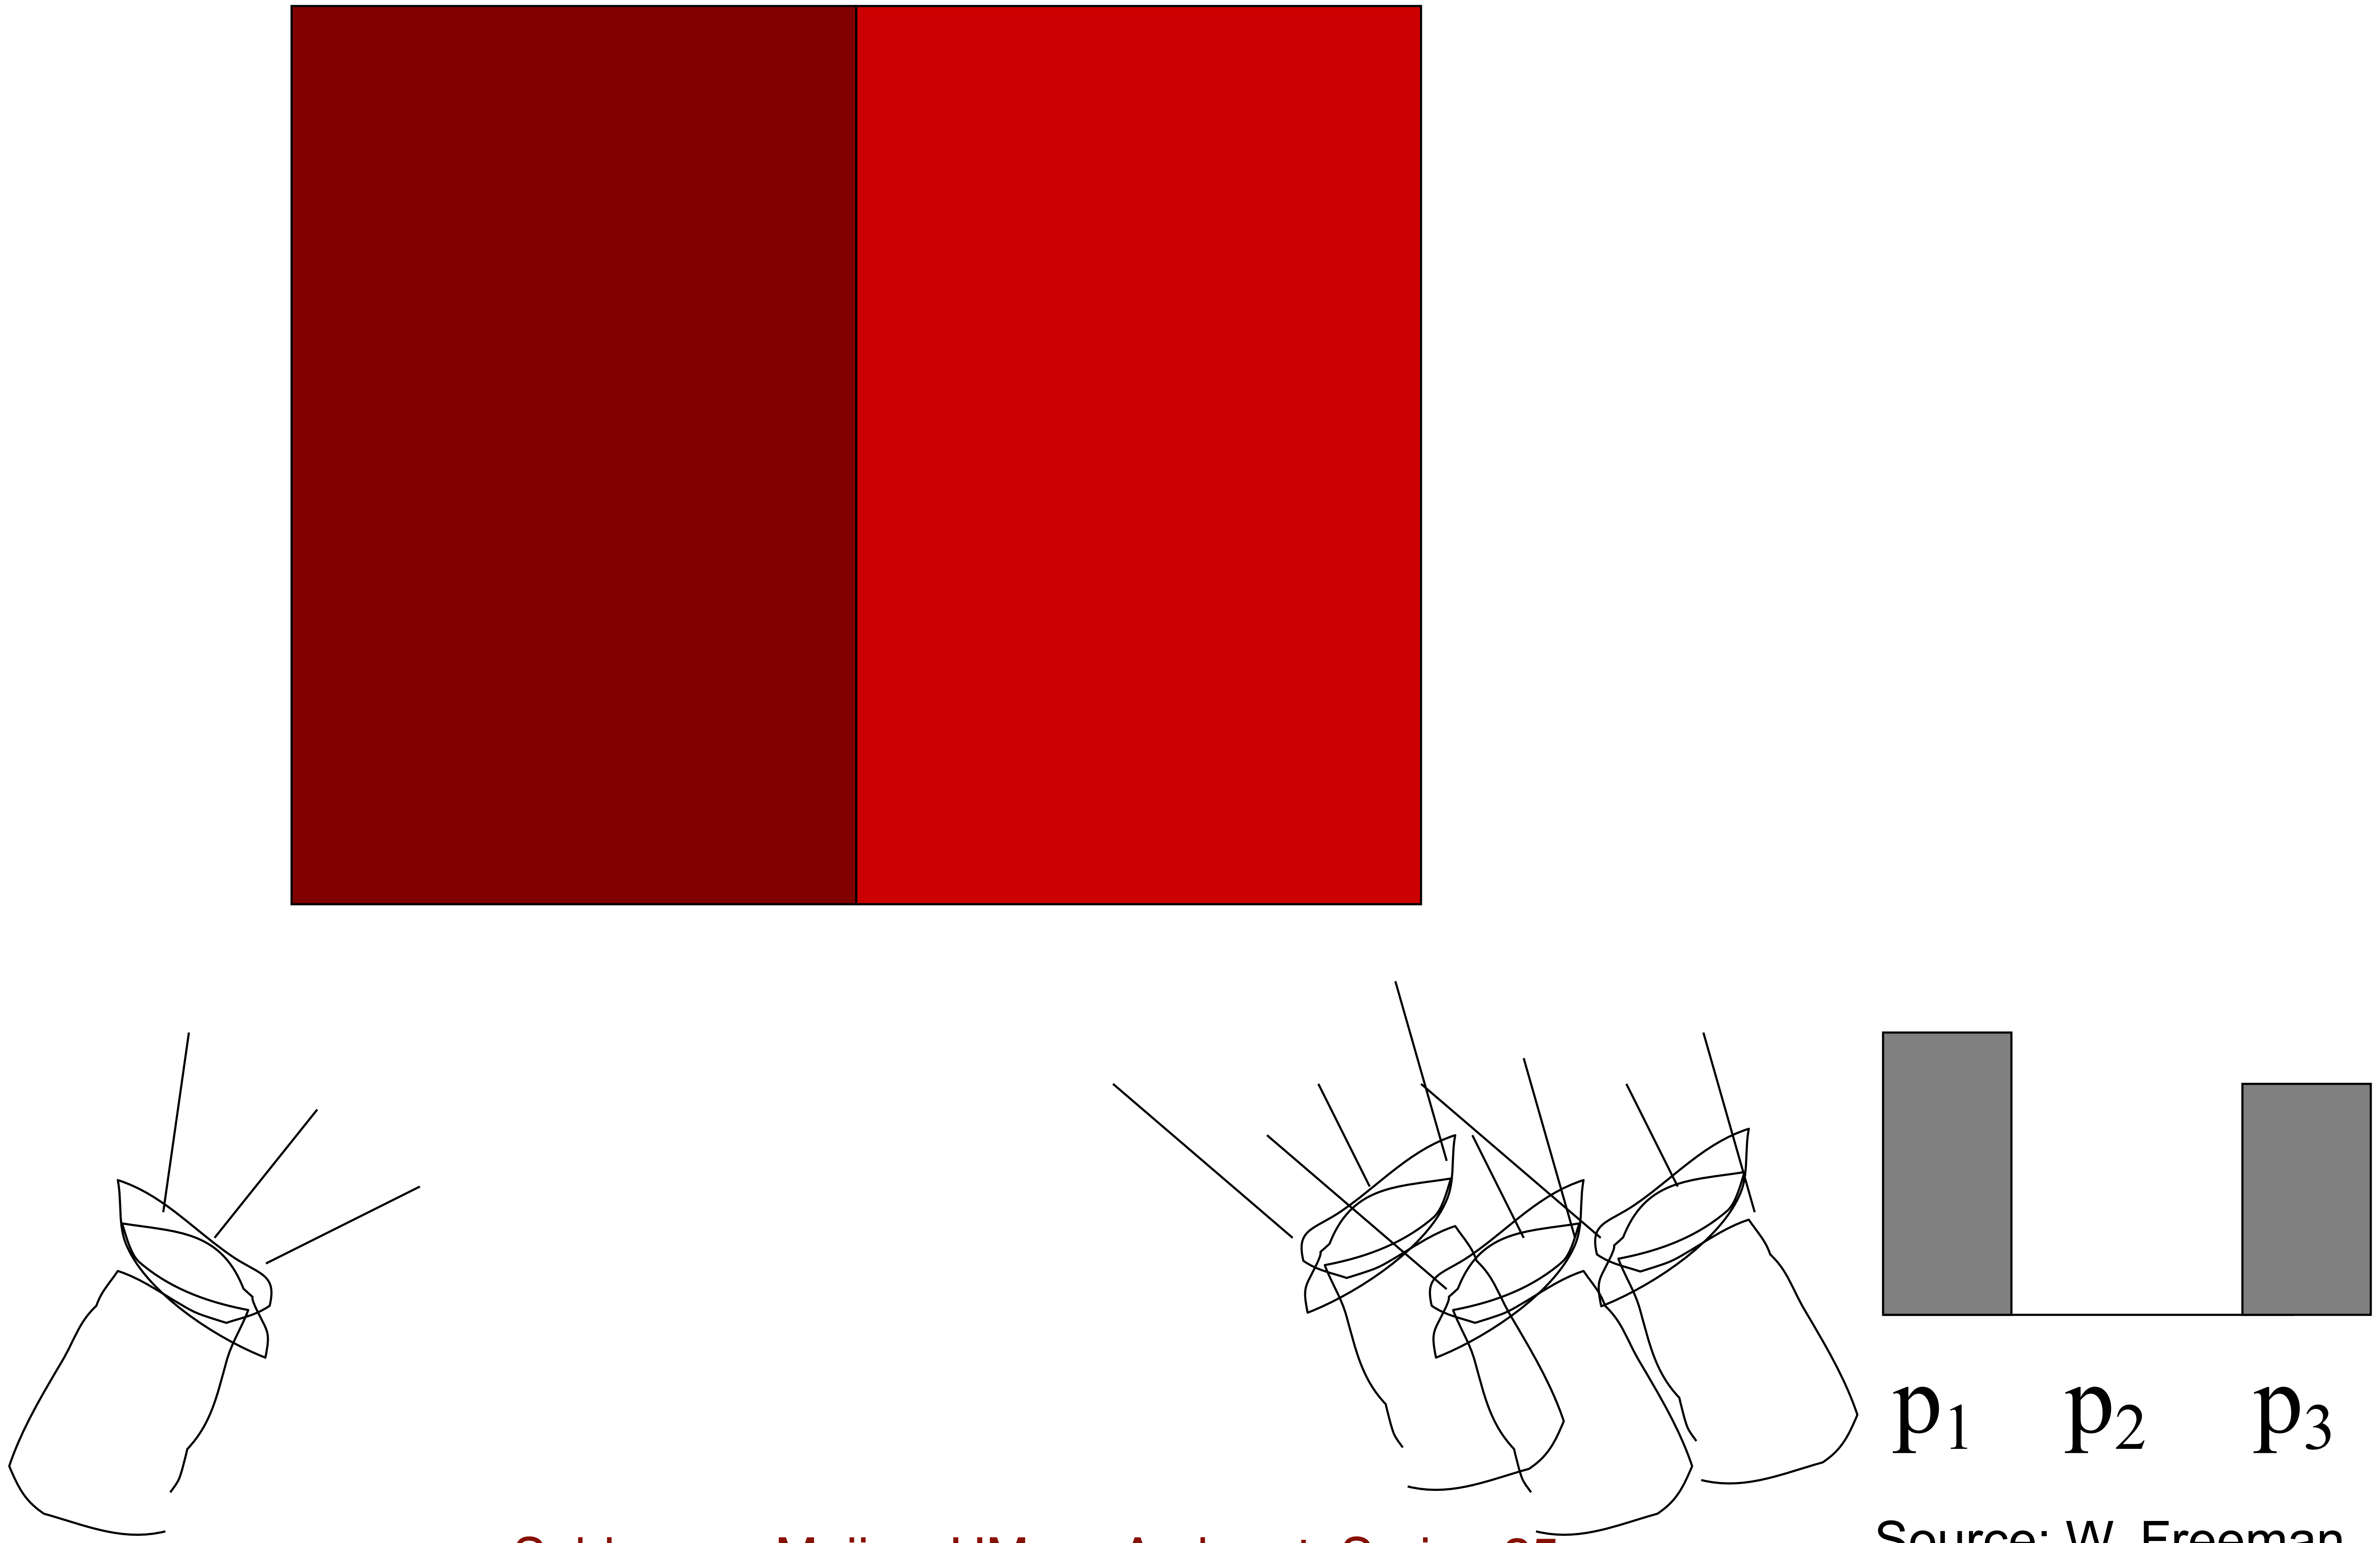

Source: W. Freeman

Color matching experiment 2

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

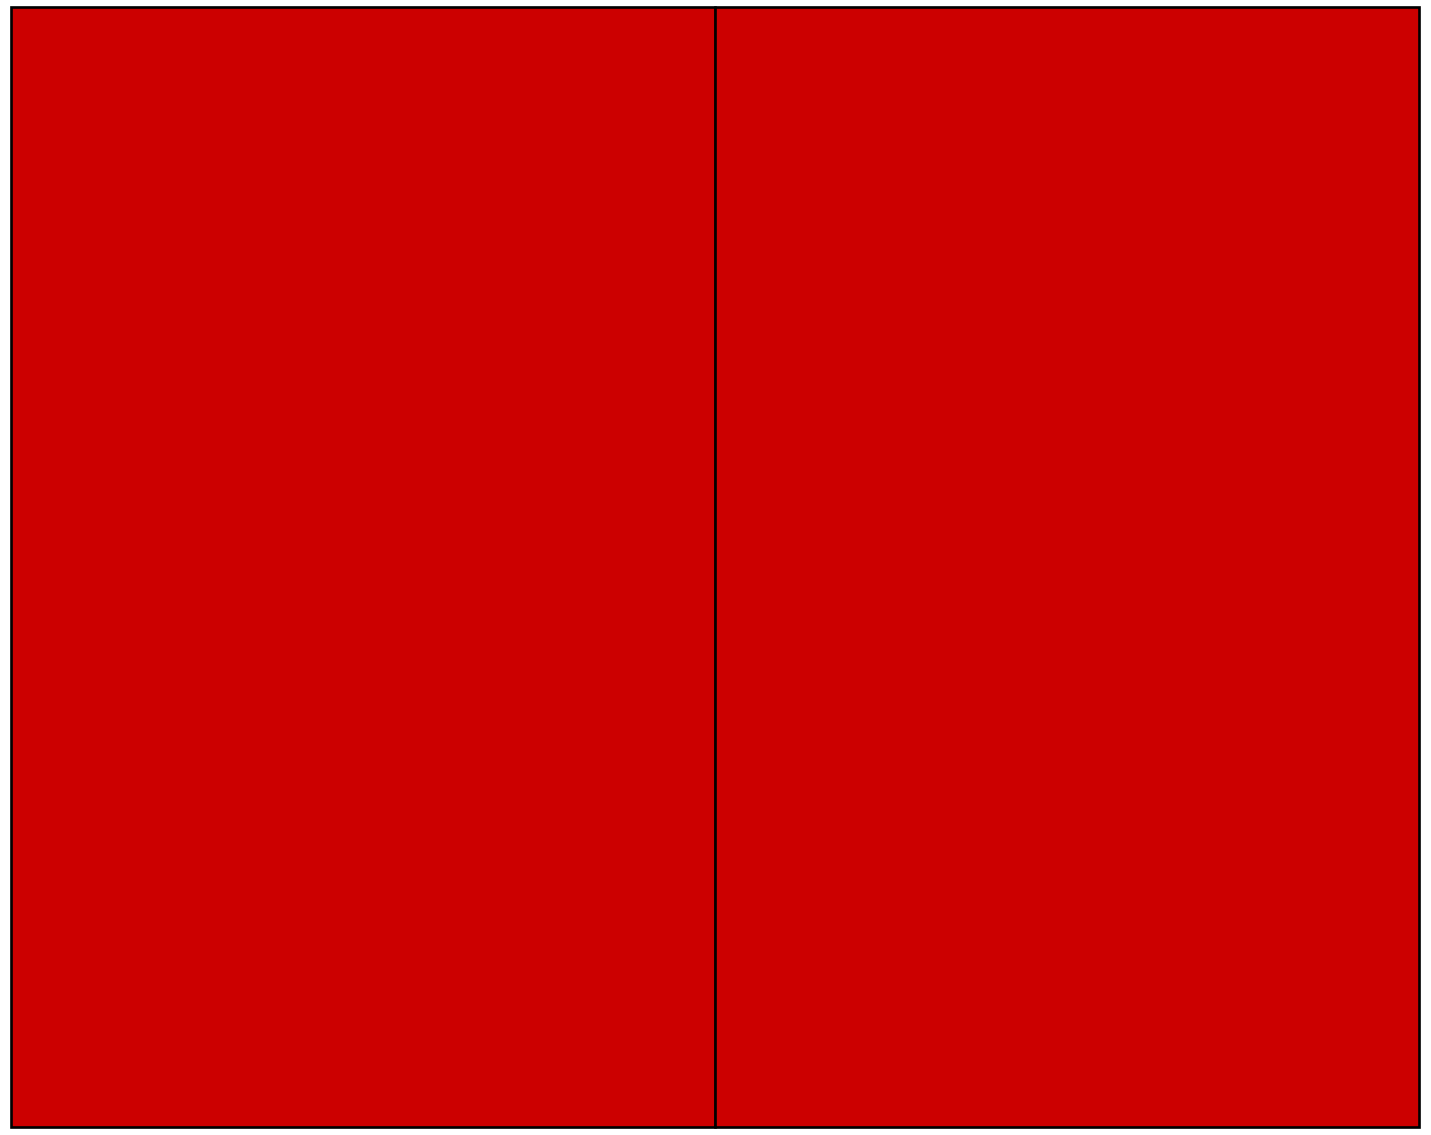

The primary color amounts needed for a match:

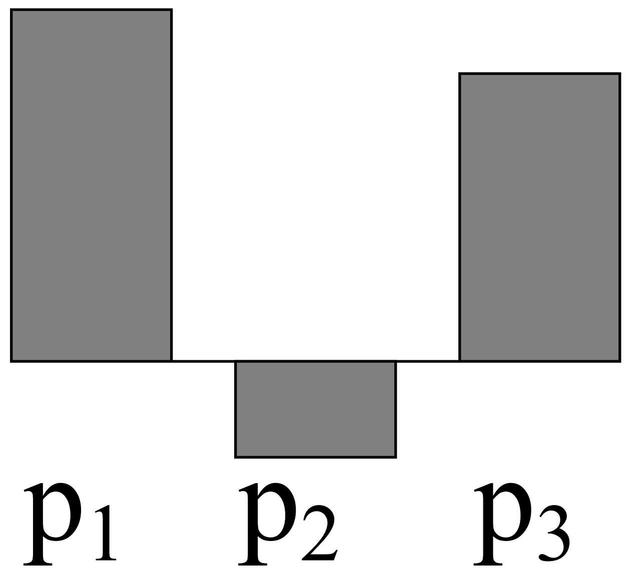

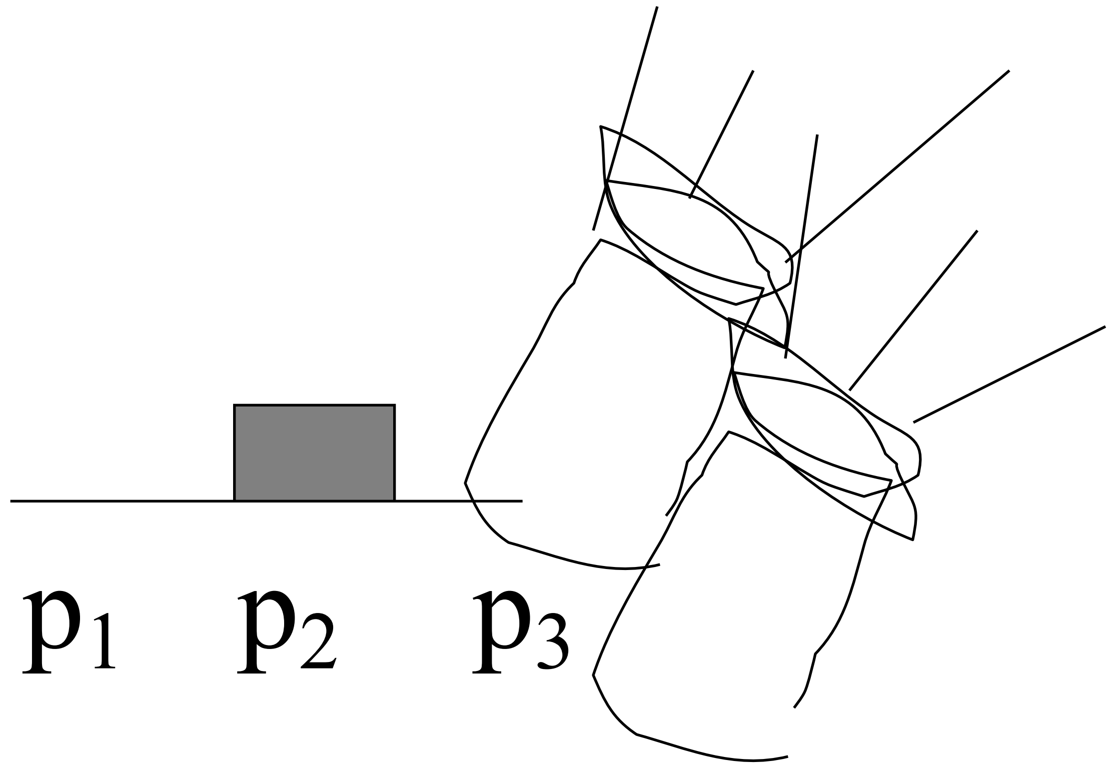

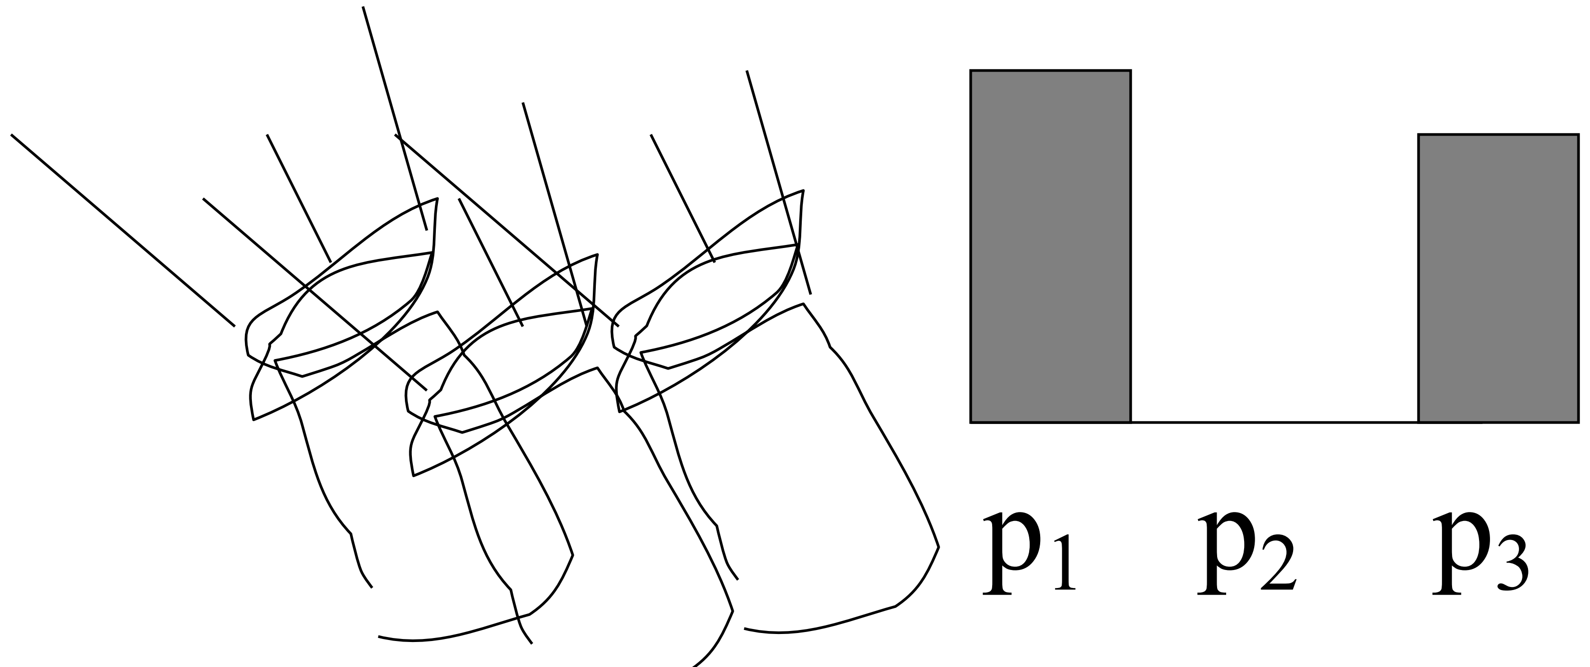

Source: W. Freeman

Trichromacy

In color matching experiments, most people can match any given light with three primaries

- Primaries must be *independent*

For the same light and same primaries, most people select the same weights

- Exception: color blindness

Trichromatic color theory

- Three numbers seem to be sufficient for encoding color
- Dates back to 18th century (Thomas Young)

Grassman's Laws (1853)

Color matching appears to be 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

How to compute the weights of the primaries to match any spectral signal?

Given: a choice of three primaries and a target color signal

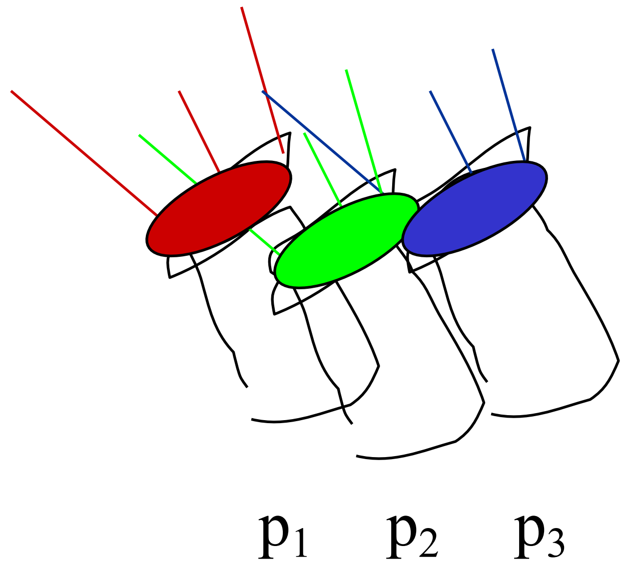

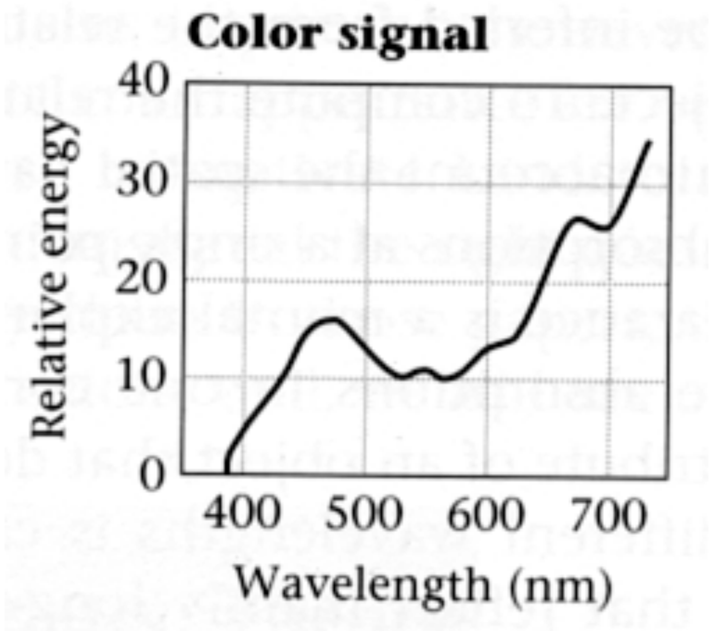

?

Find: weights of the primaries needed to match the color signal

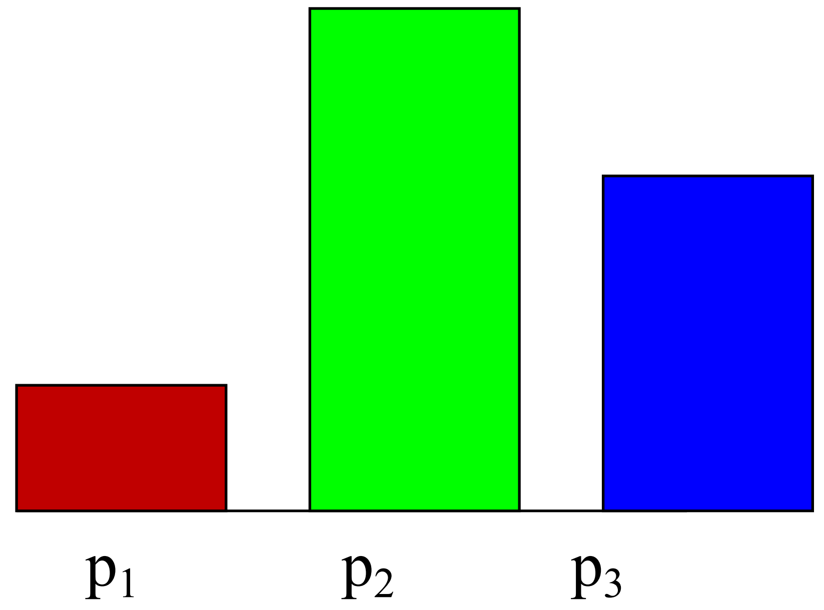

Color matching function: primary color

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.

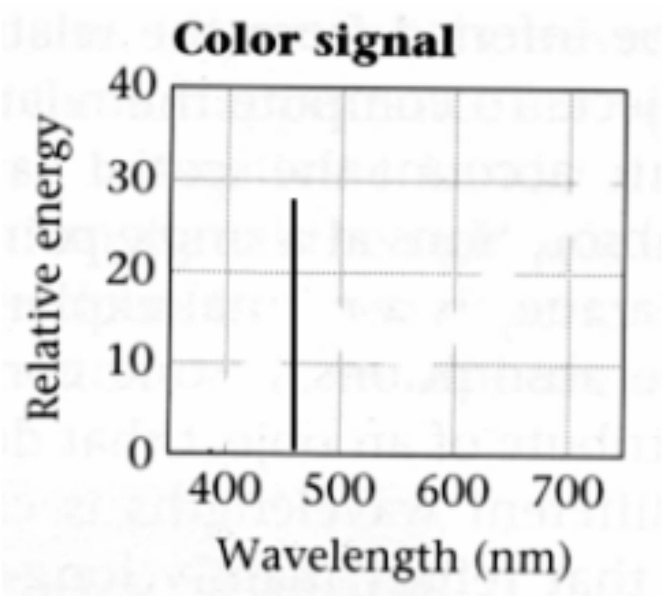

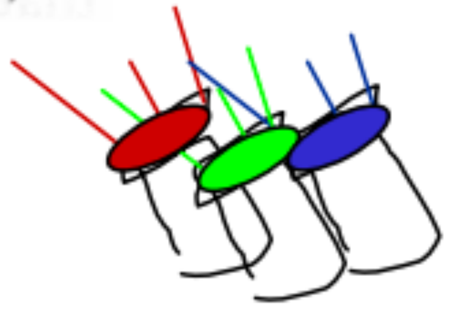

And any spectral signal can be thought of as a linear combination of very many monochromatic lights, with the linear coefficient given by the spectral power at each wavelength.

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

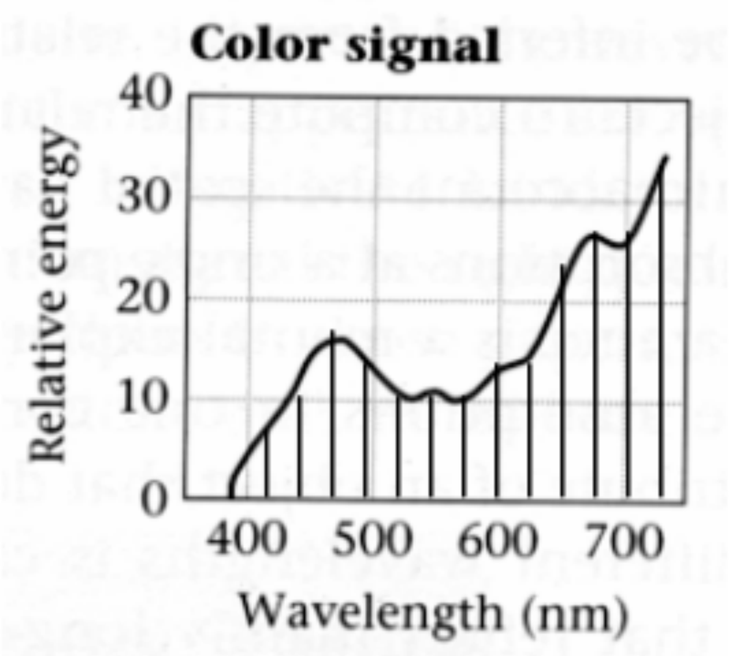

Source: W. Freeman

Color matching functions: any color

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

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

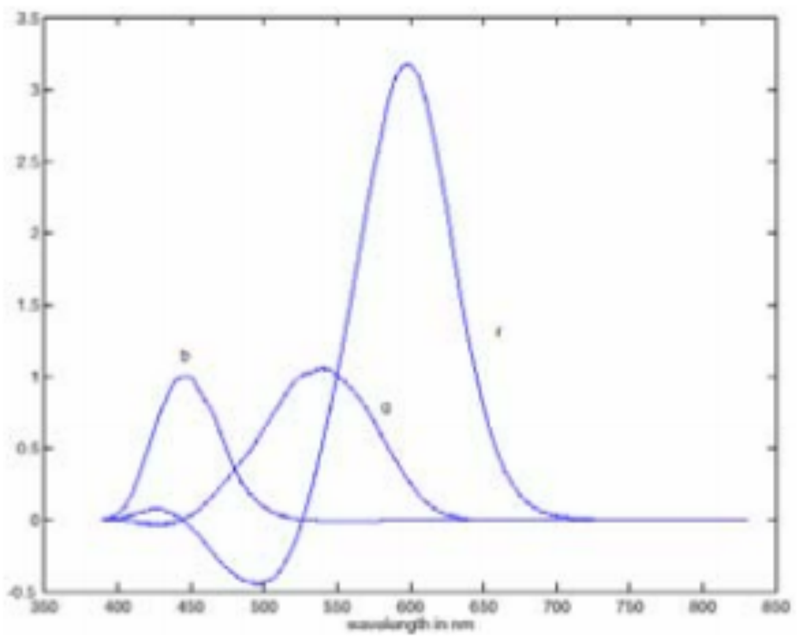

Let the new spectral signal be described by the vector *t*.

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

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

RGB space

Primaries are monochromatic lights (for monitors, they correspond to the three types of phosphors)

Subtractive matching required for some wavelengths

RGB primaries

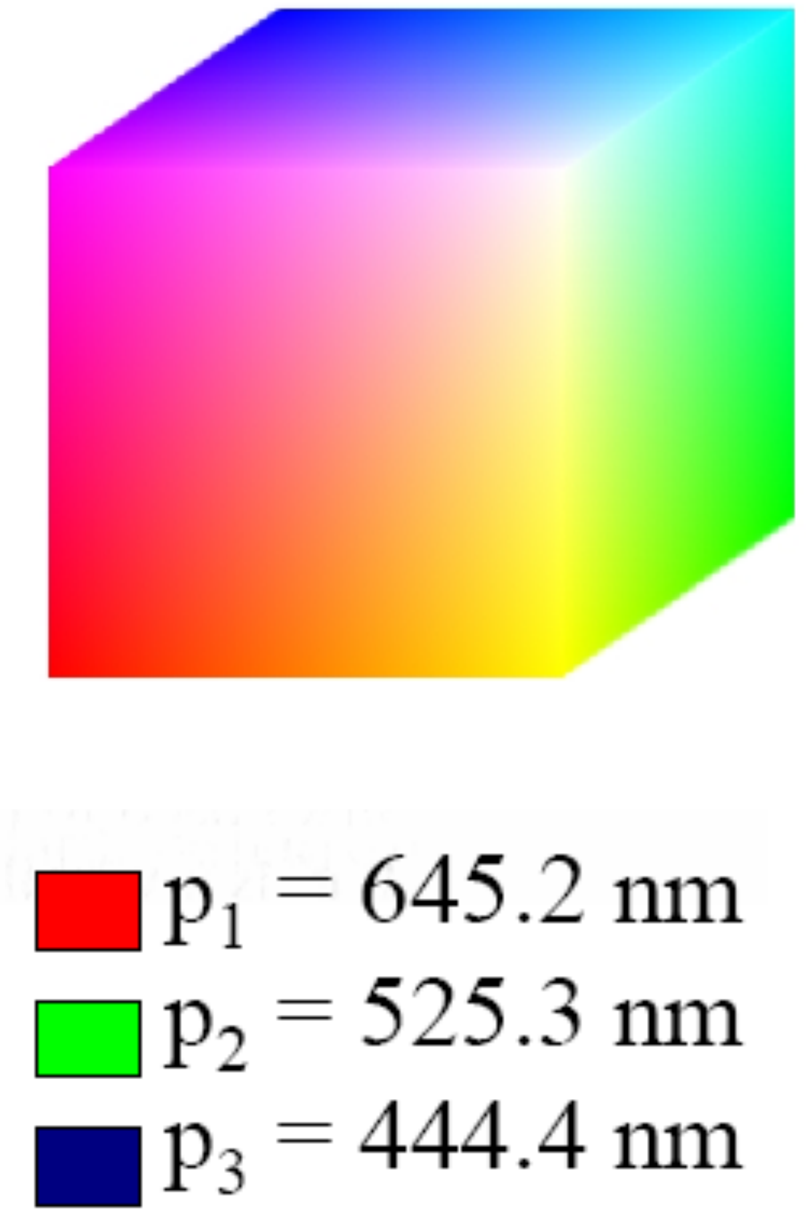

RGB matching functions

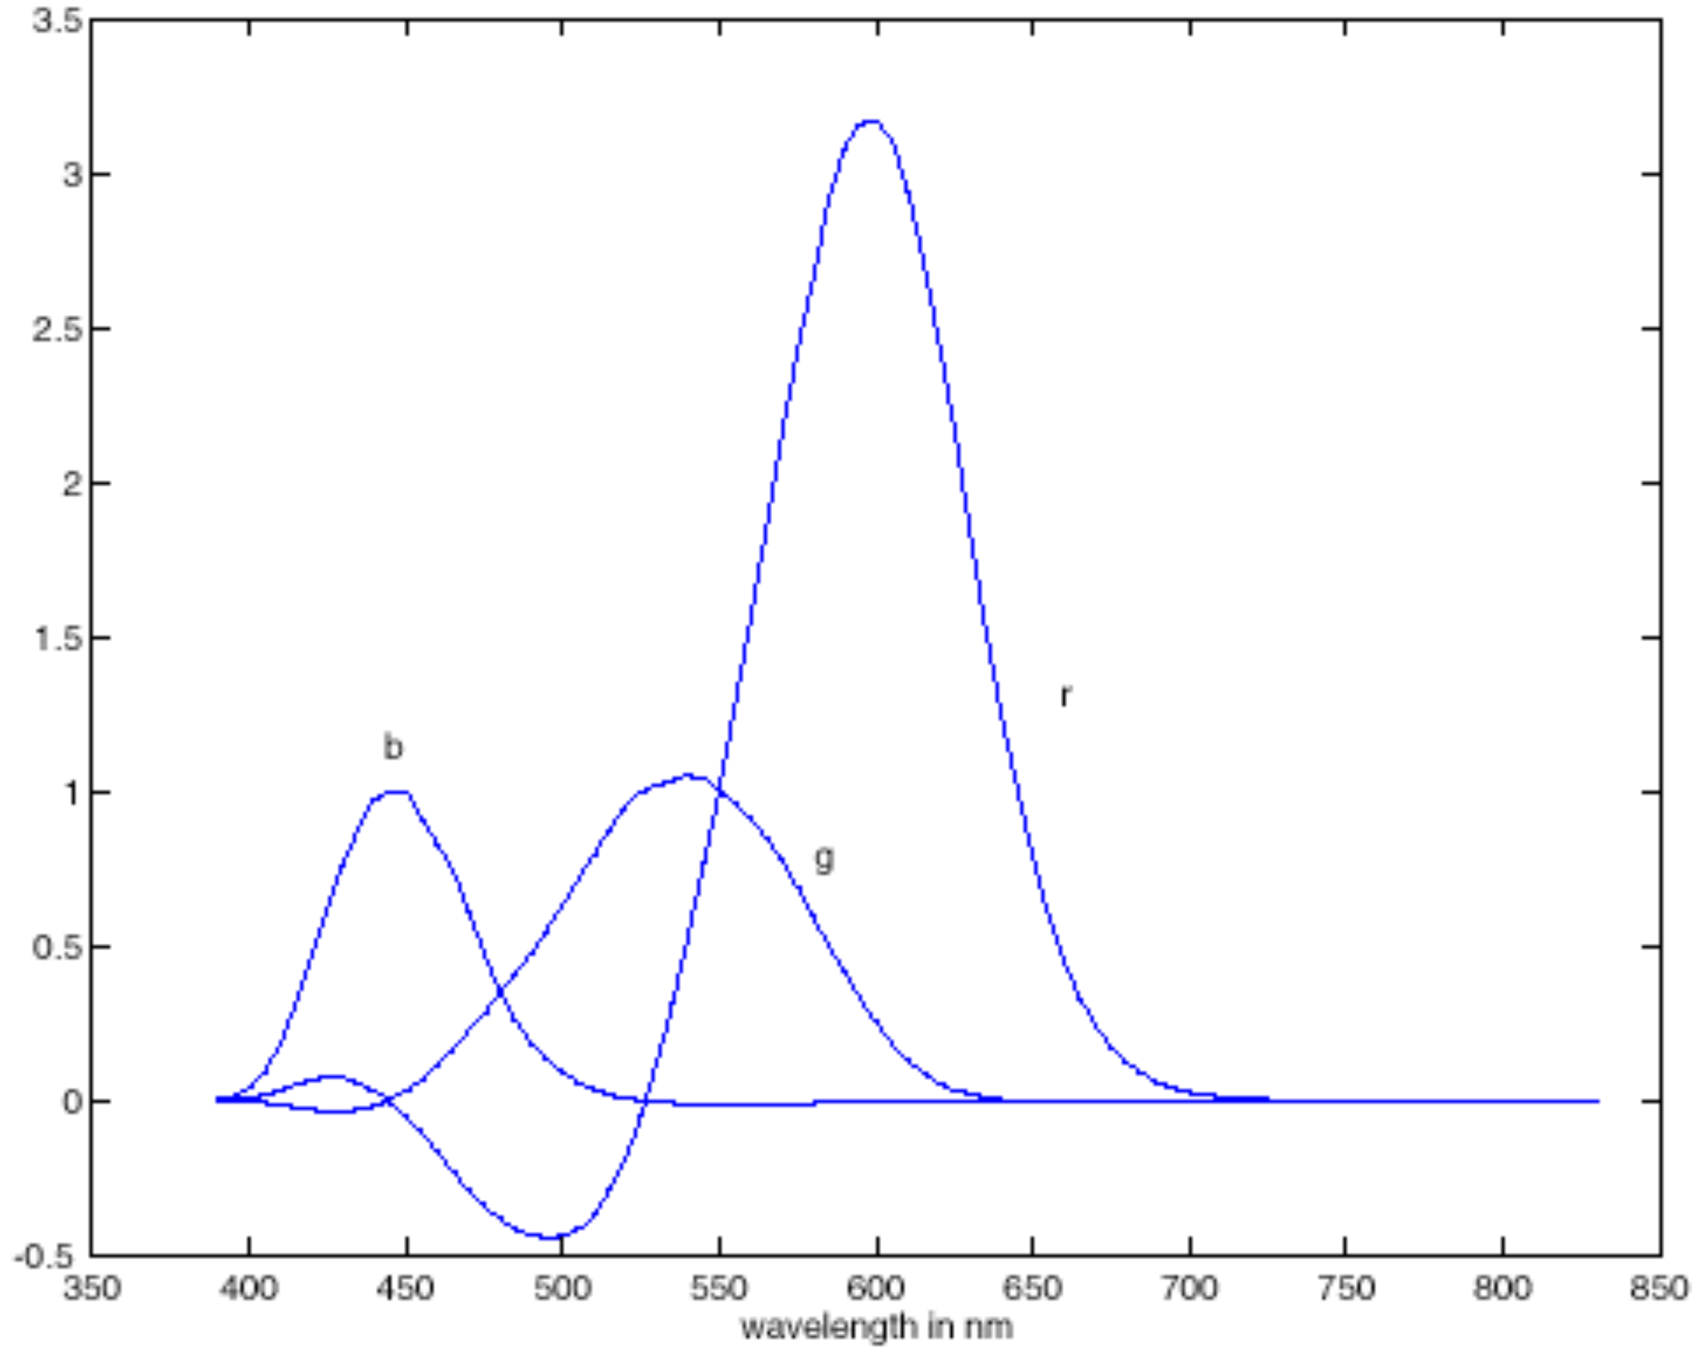

Color displays

Given a color image (say 8 bits for R, G, B)

- Turn on 3 subpixels with power proportional to RGB values

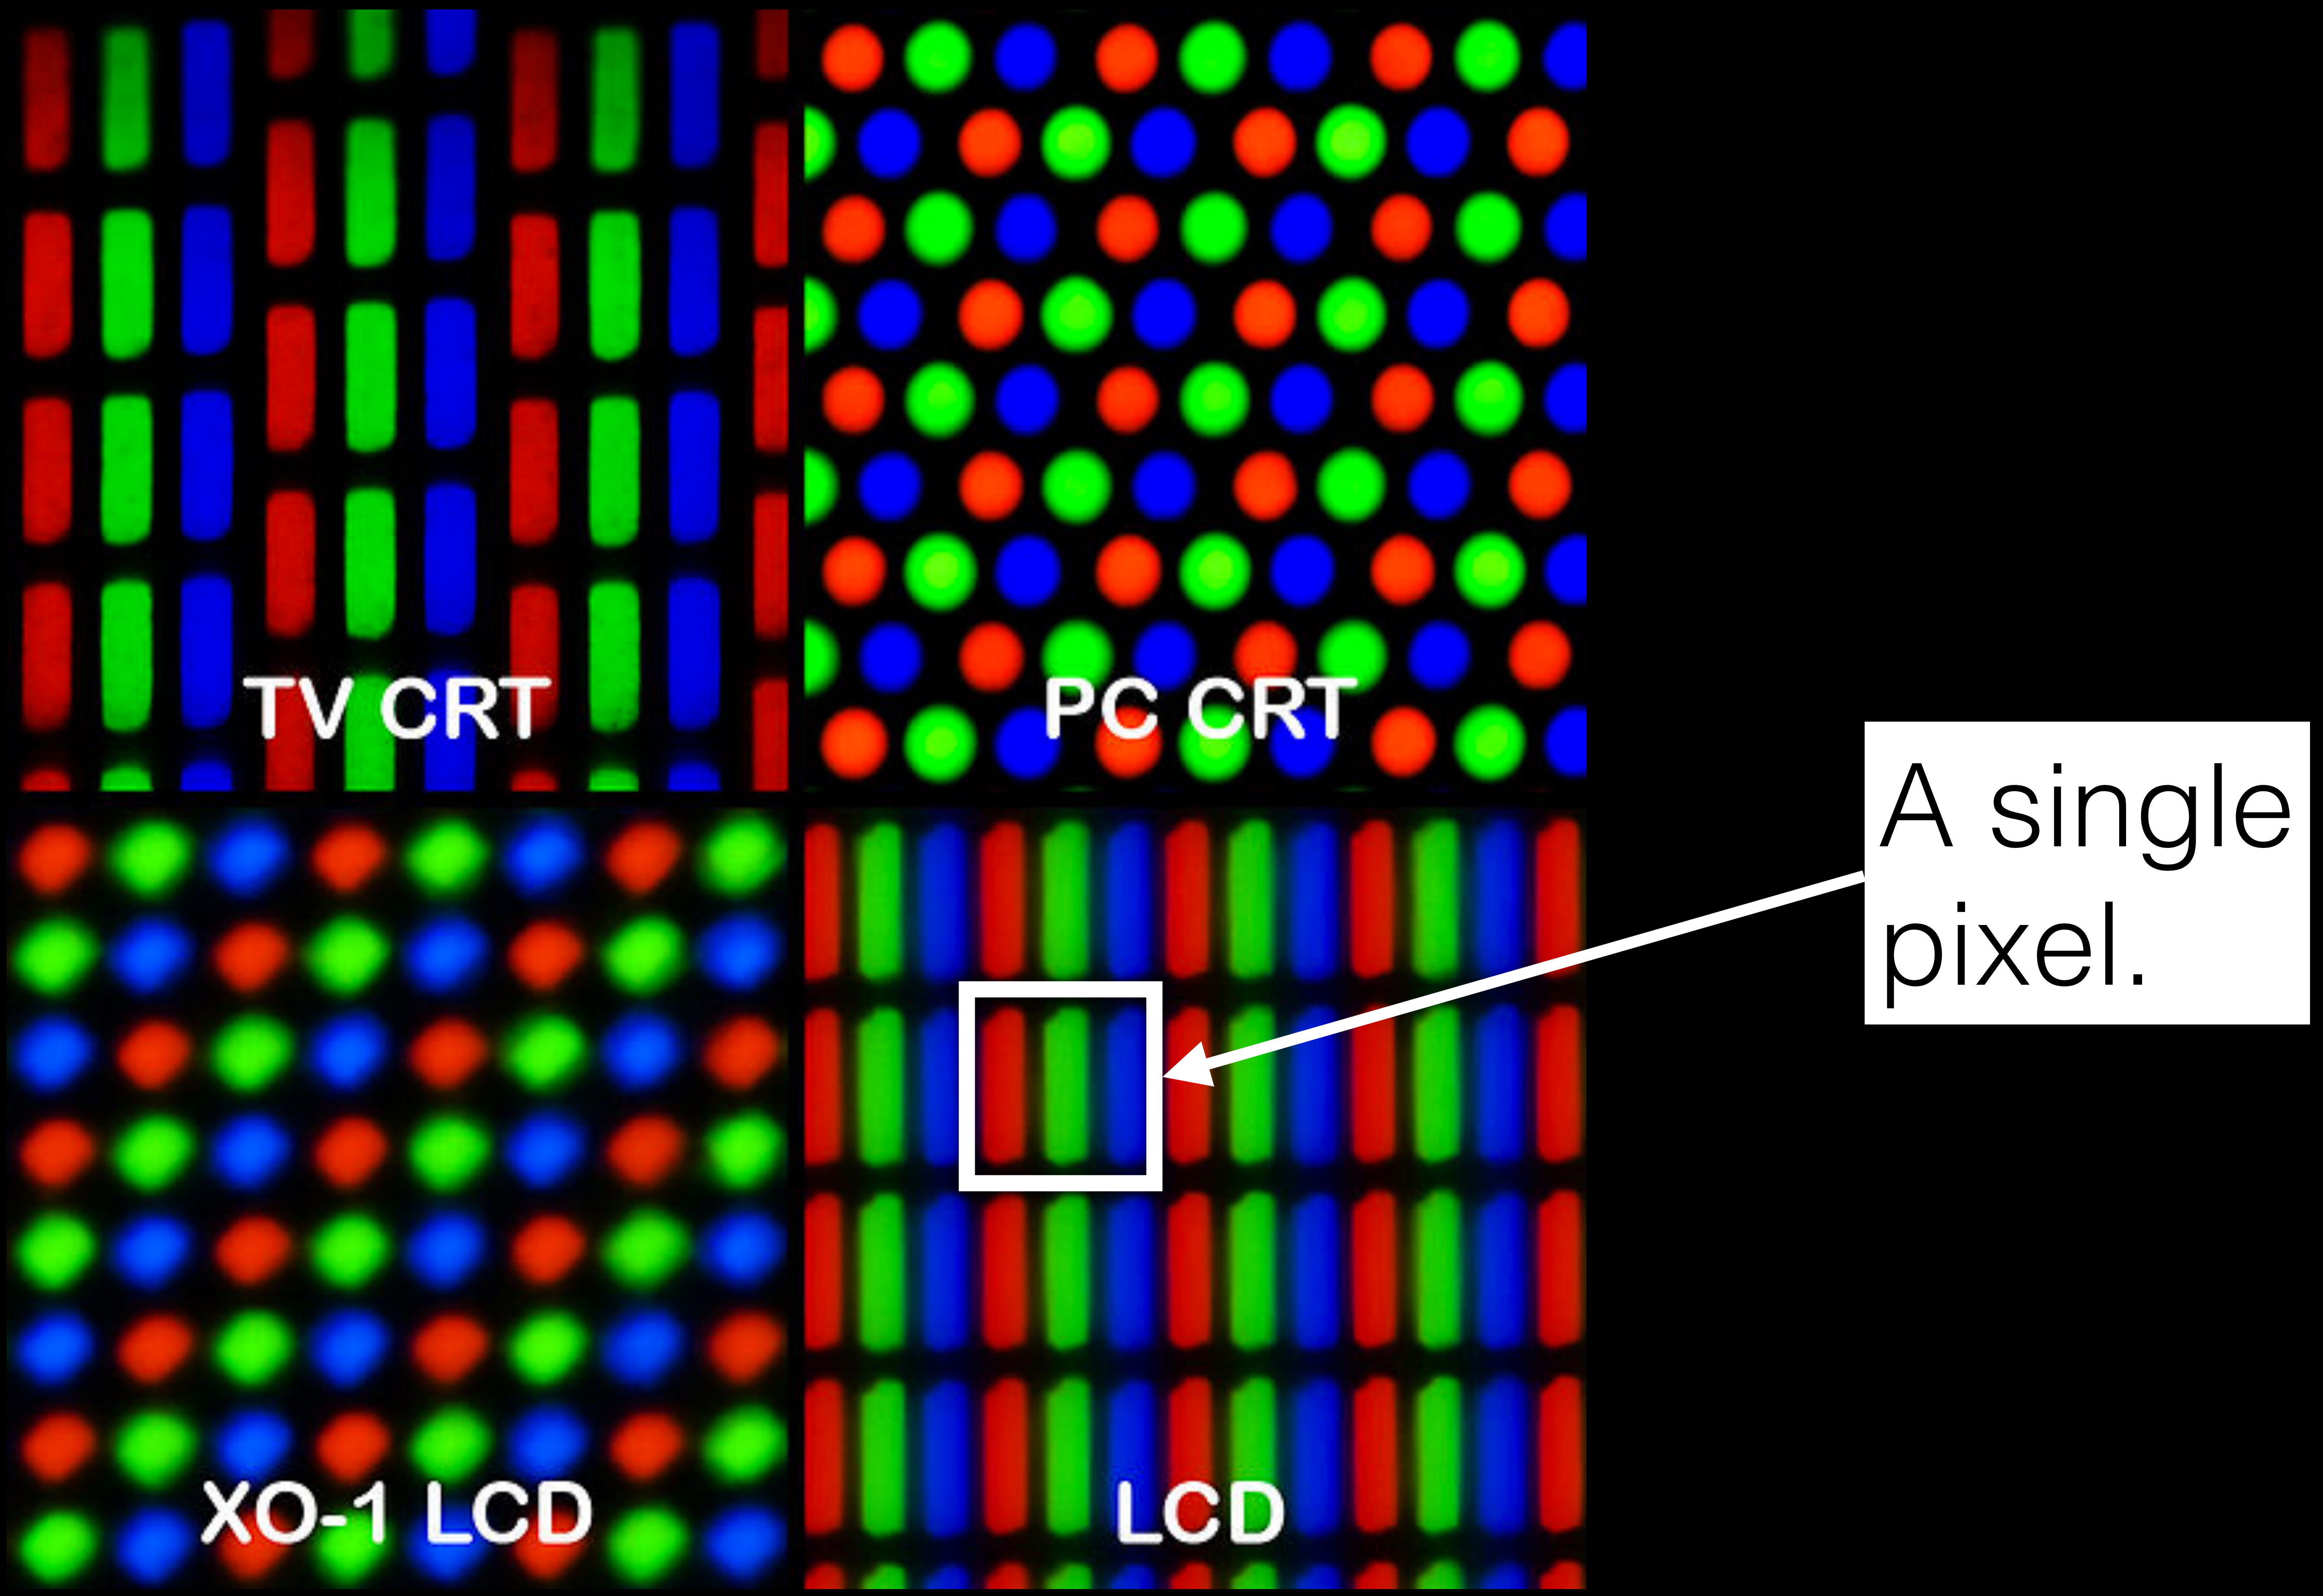

https://en.wikipedia.org/wiki/File:Pixel_geometry_01_Pengo.jpg

“White” text on color display

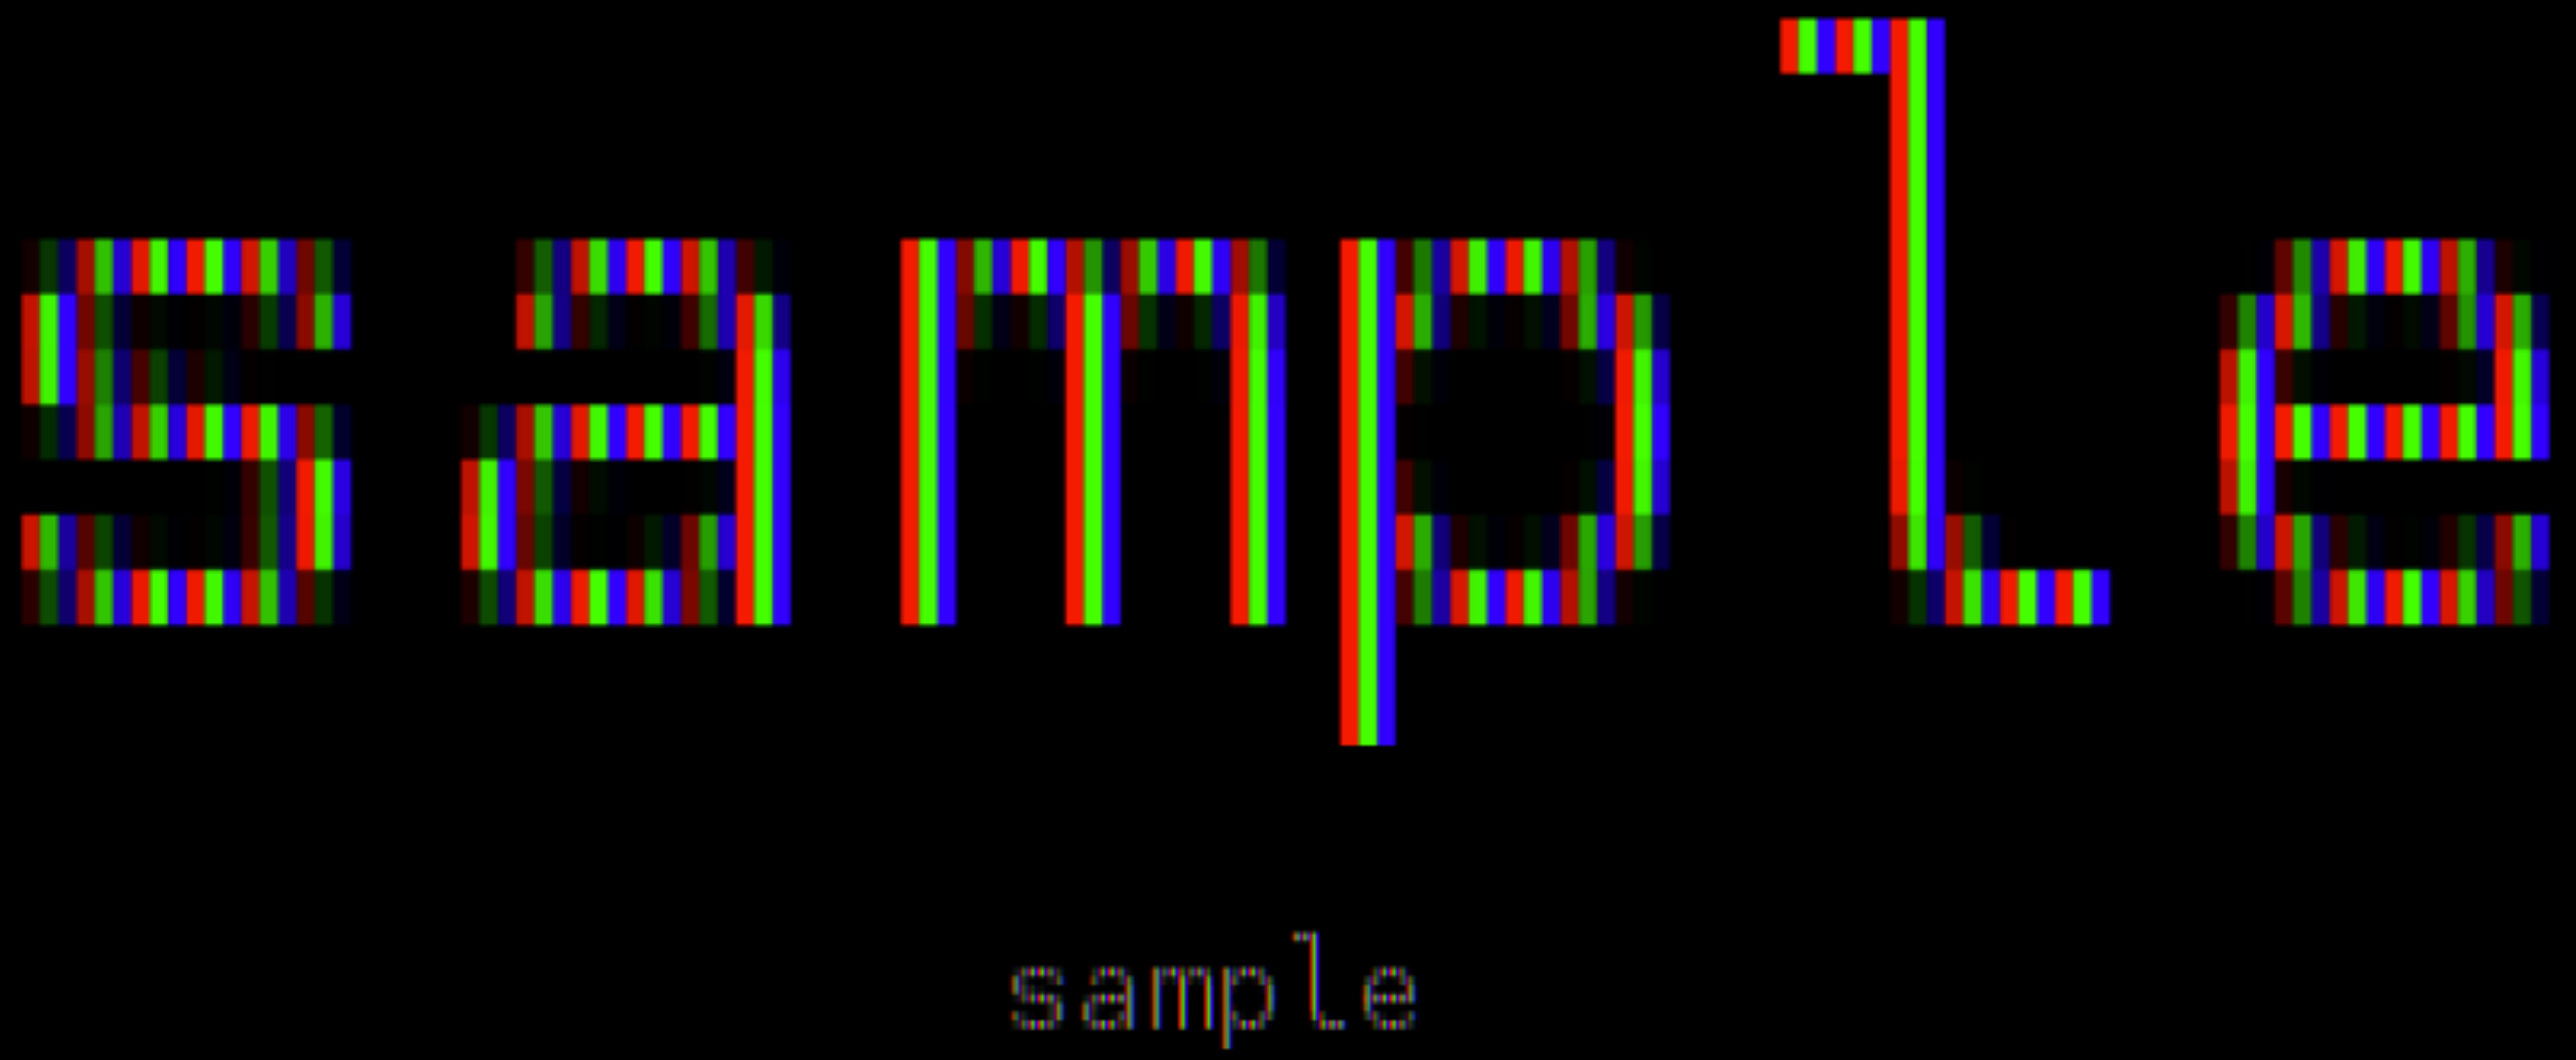

http://en.wikipedia.org/wiki/Subpixel_rendering

Uniform color spaces

Unfortunately, differences in x,y coordinates do not reflect perceptual color differences

CIE u’v’ is a transform of x,y to make the ellipses more uniform

xyz

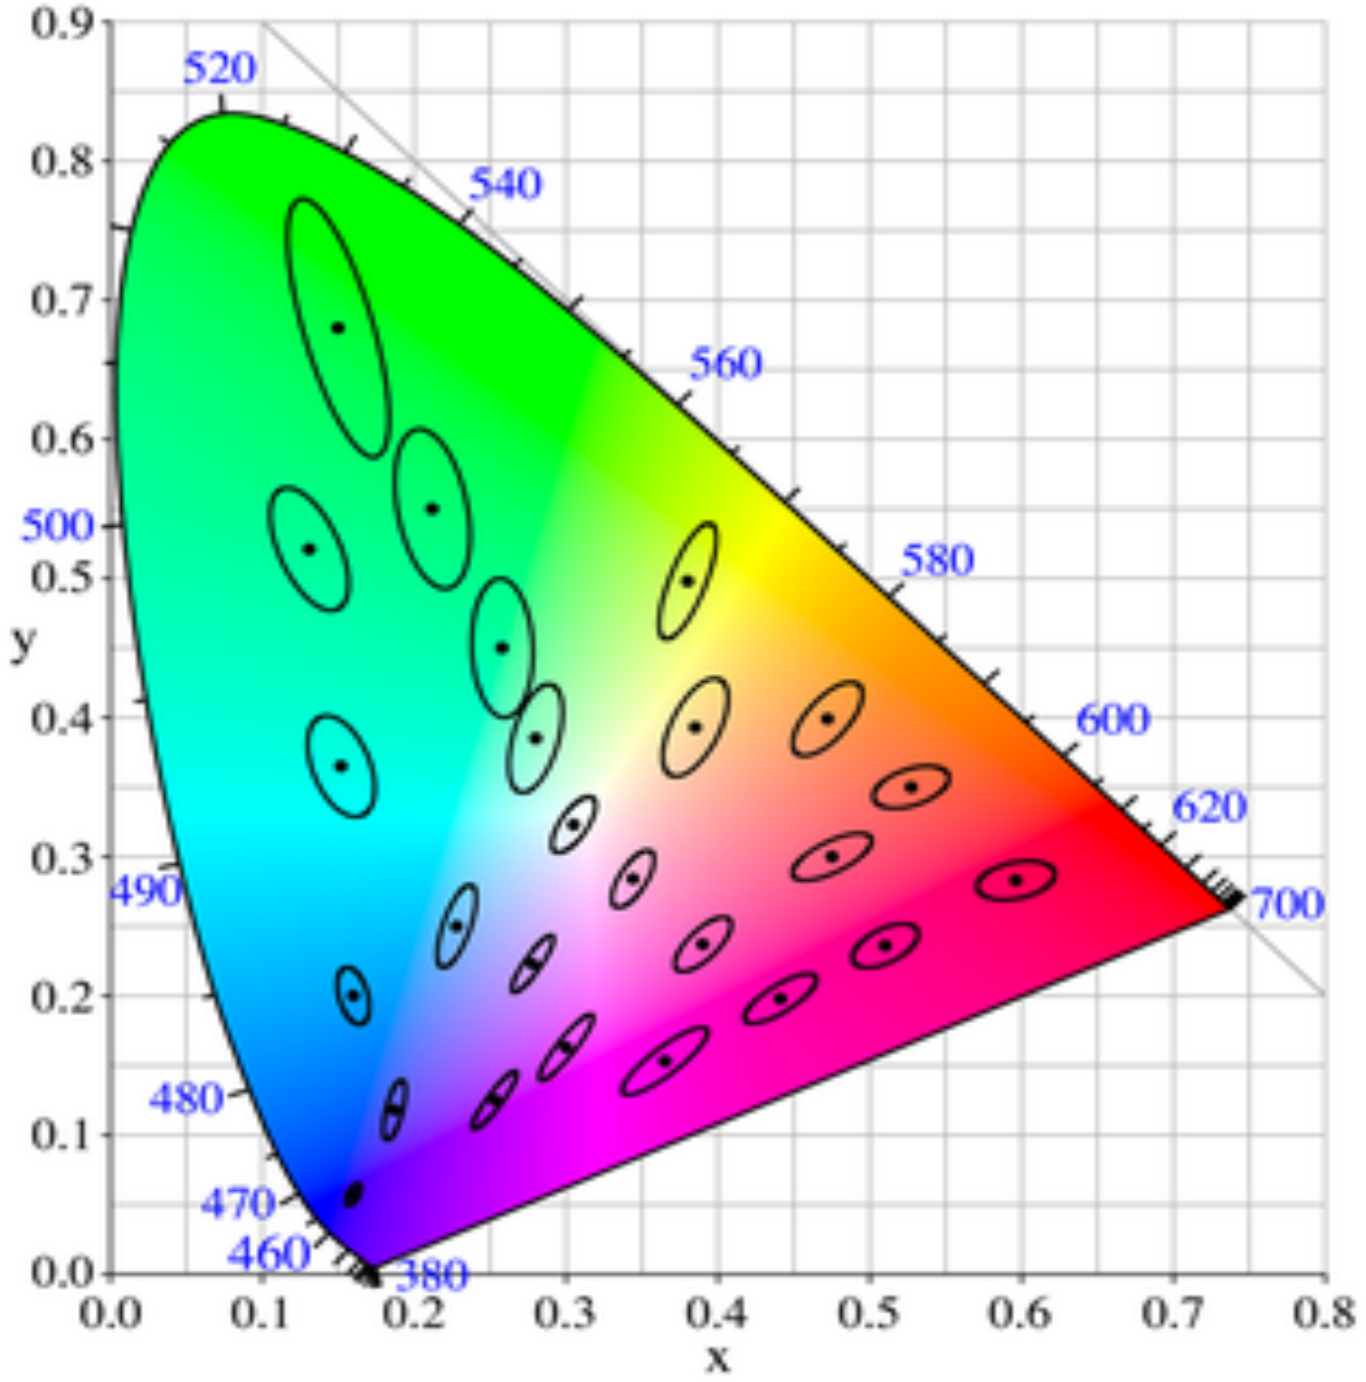

McAdam ellipses: Just noticeable differences in color

lu’v’

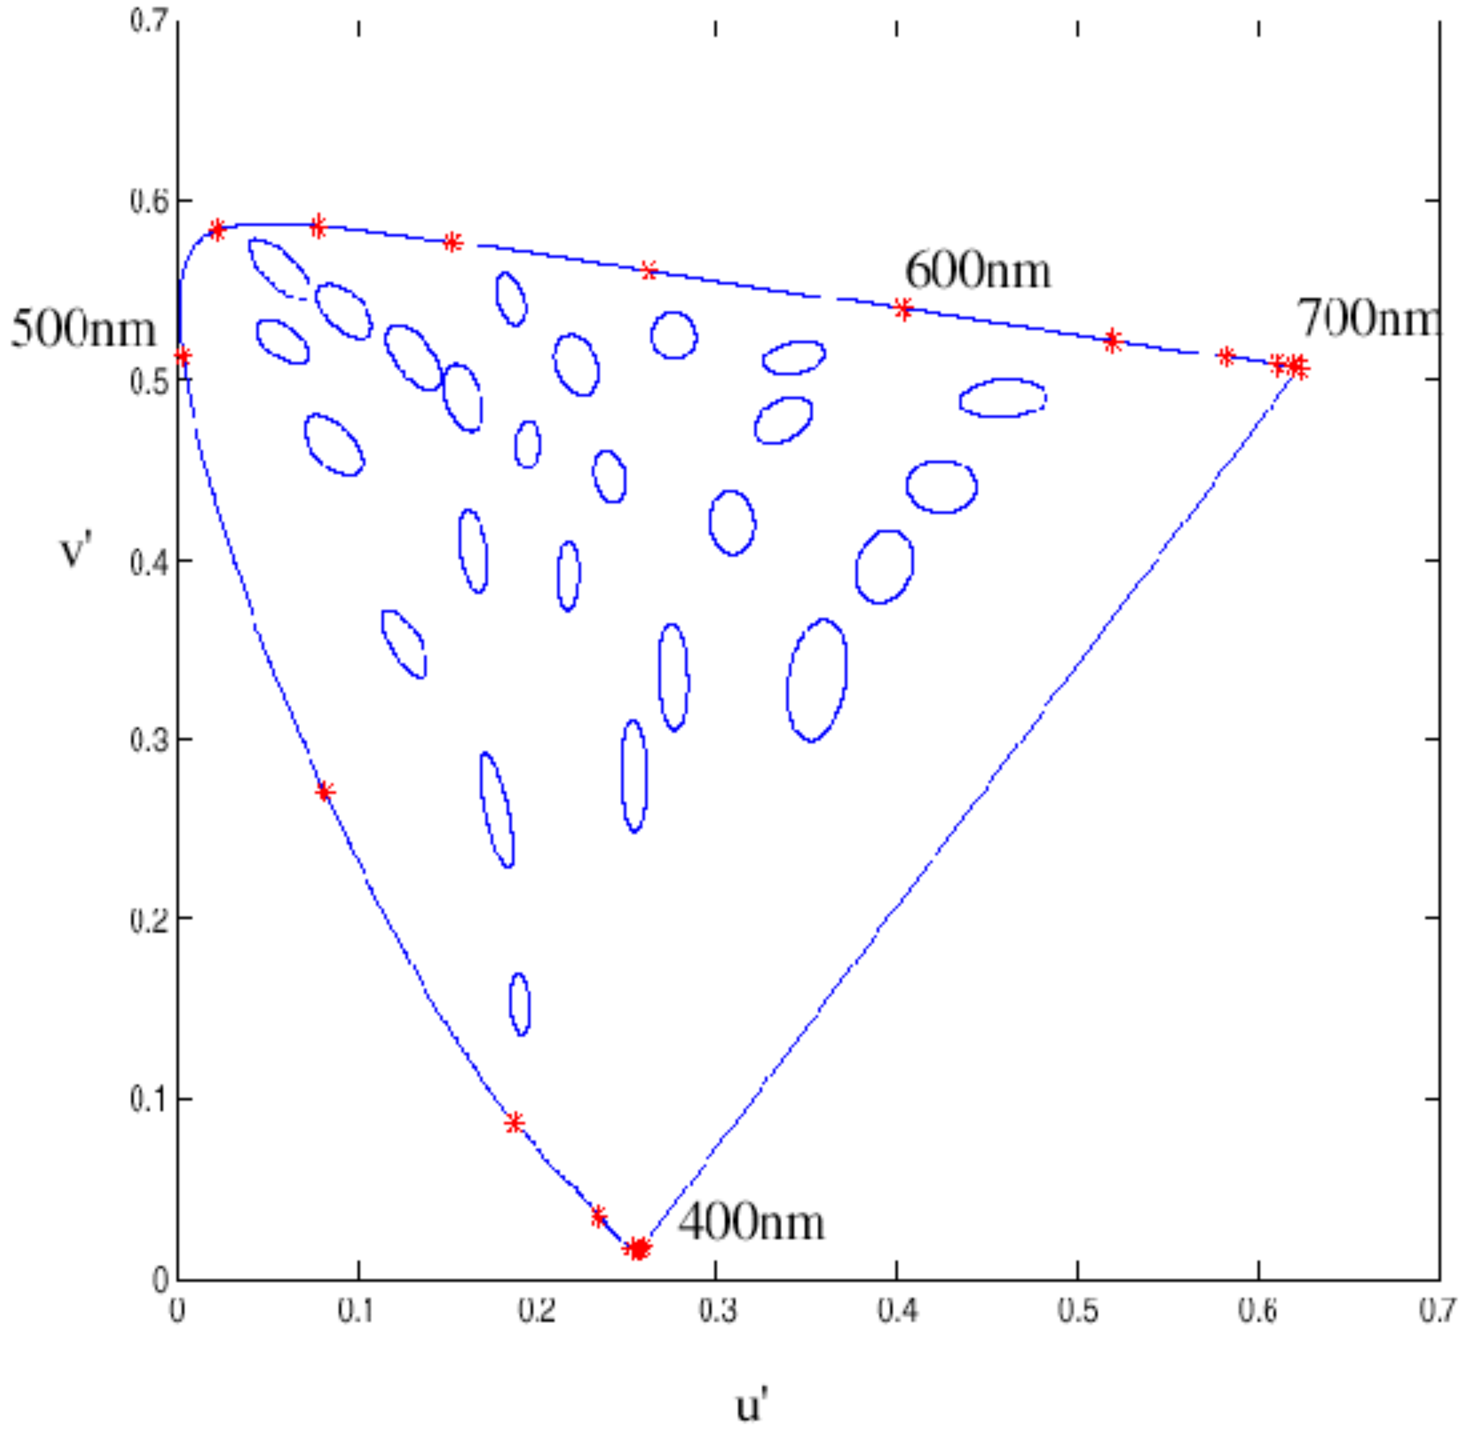

Nonlinear color spaces: HSV

Perceptually meaningful dimensions: Hue, Saturation, Value (Intensity)

RGB cube on its vertex

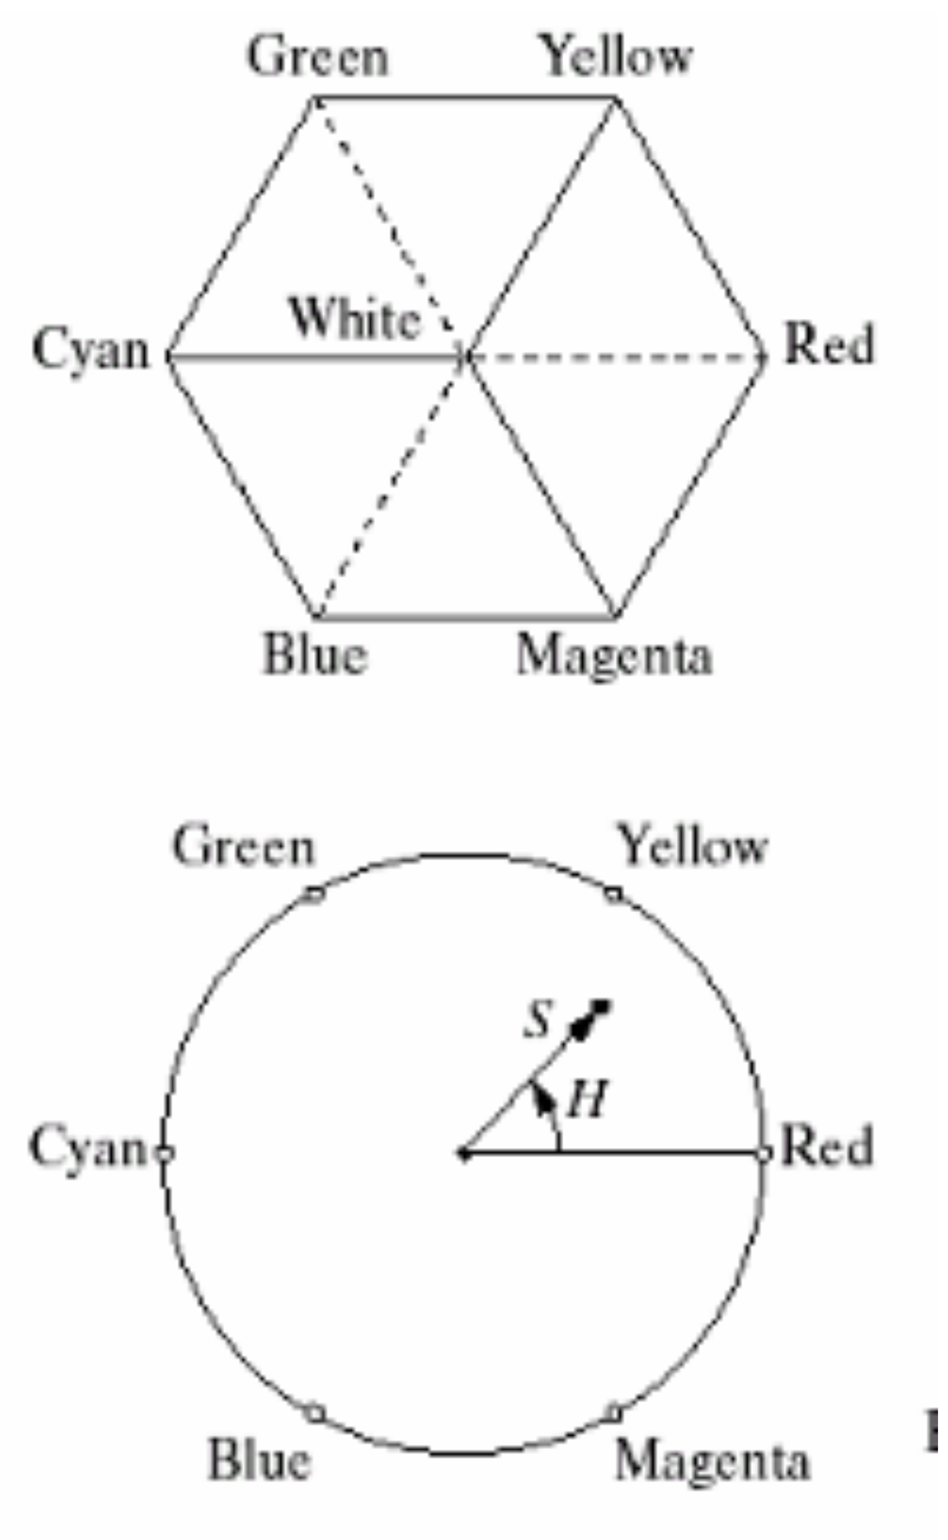

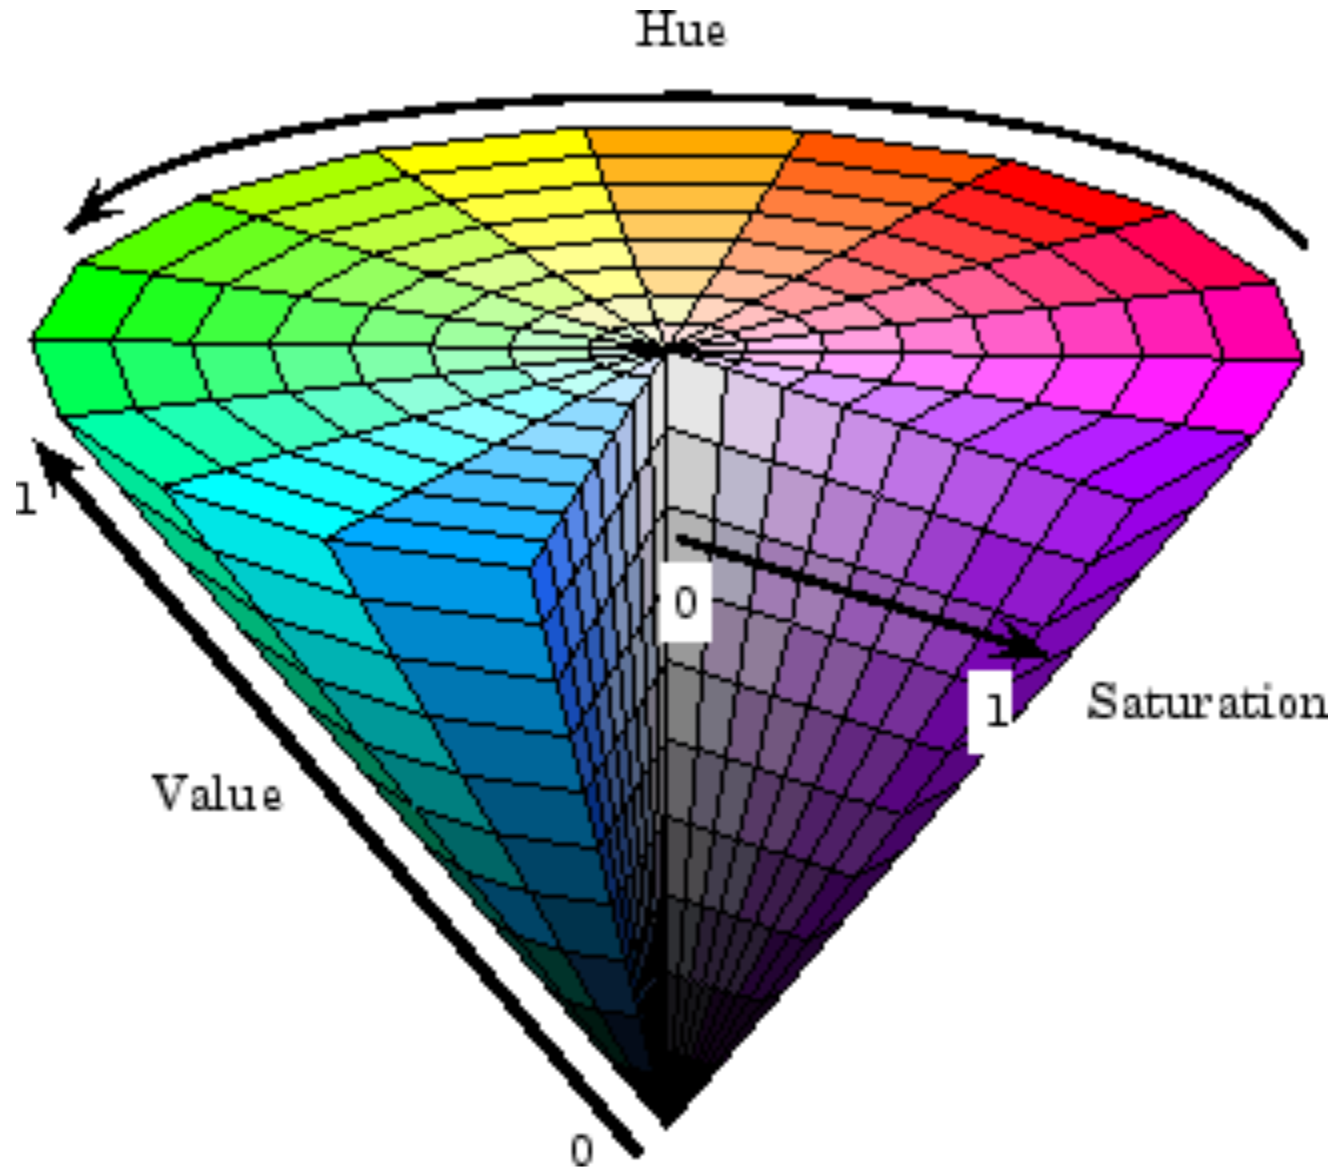

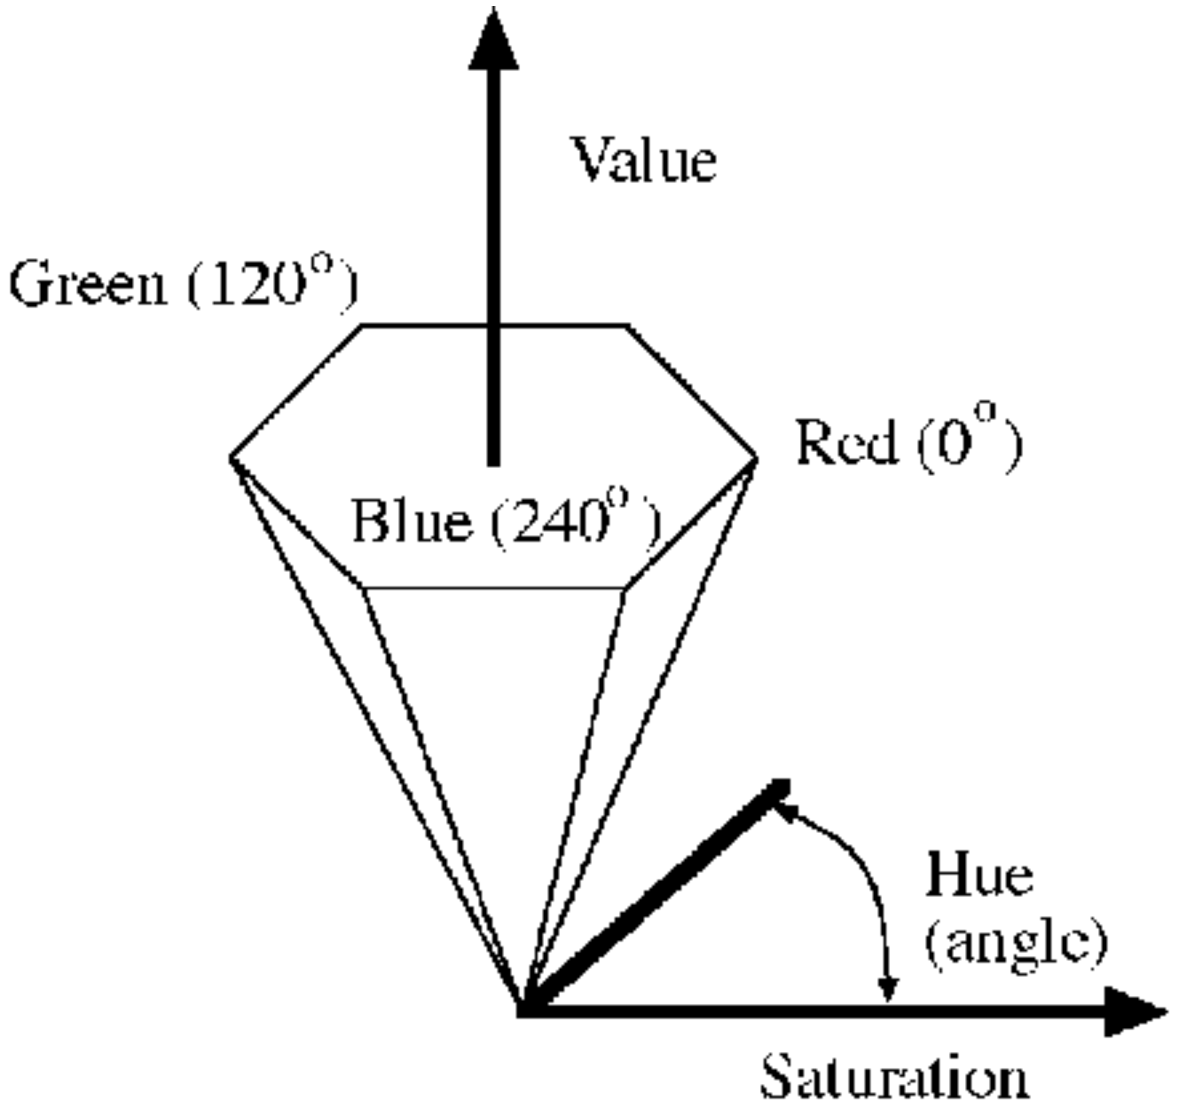

Some early attempts in color spaces

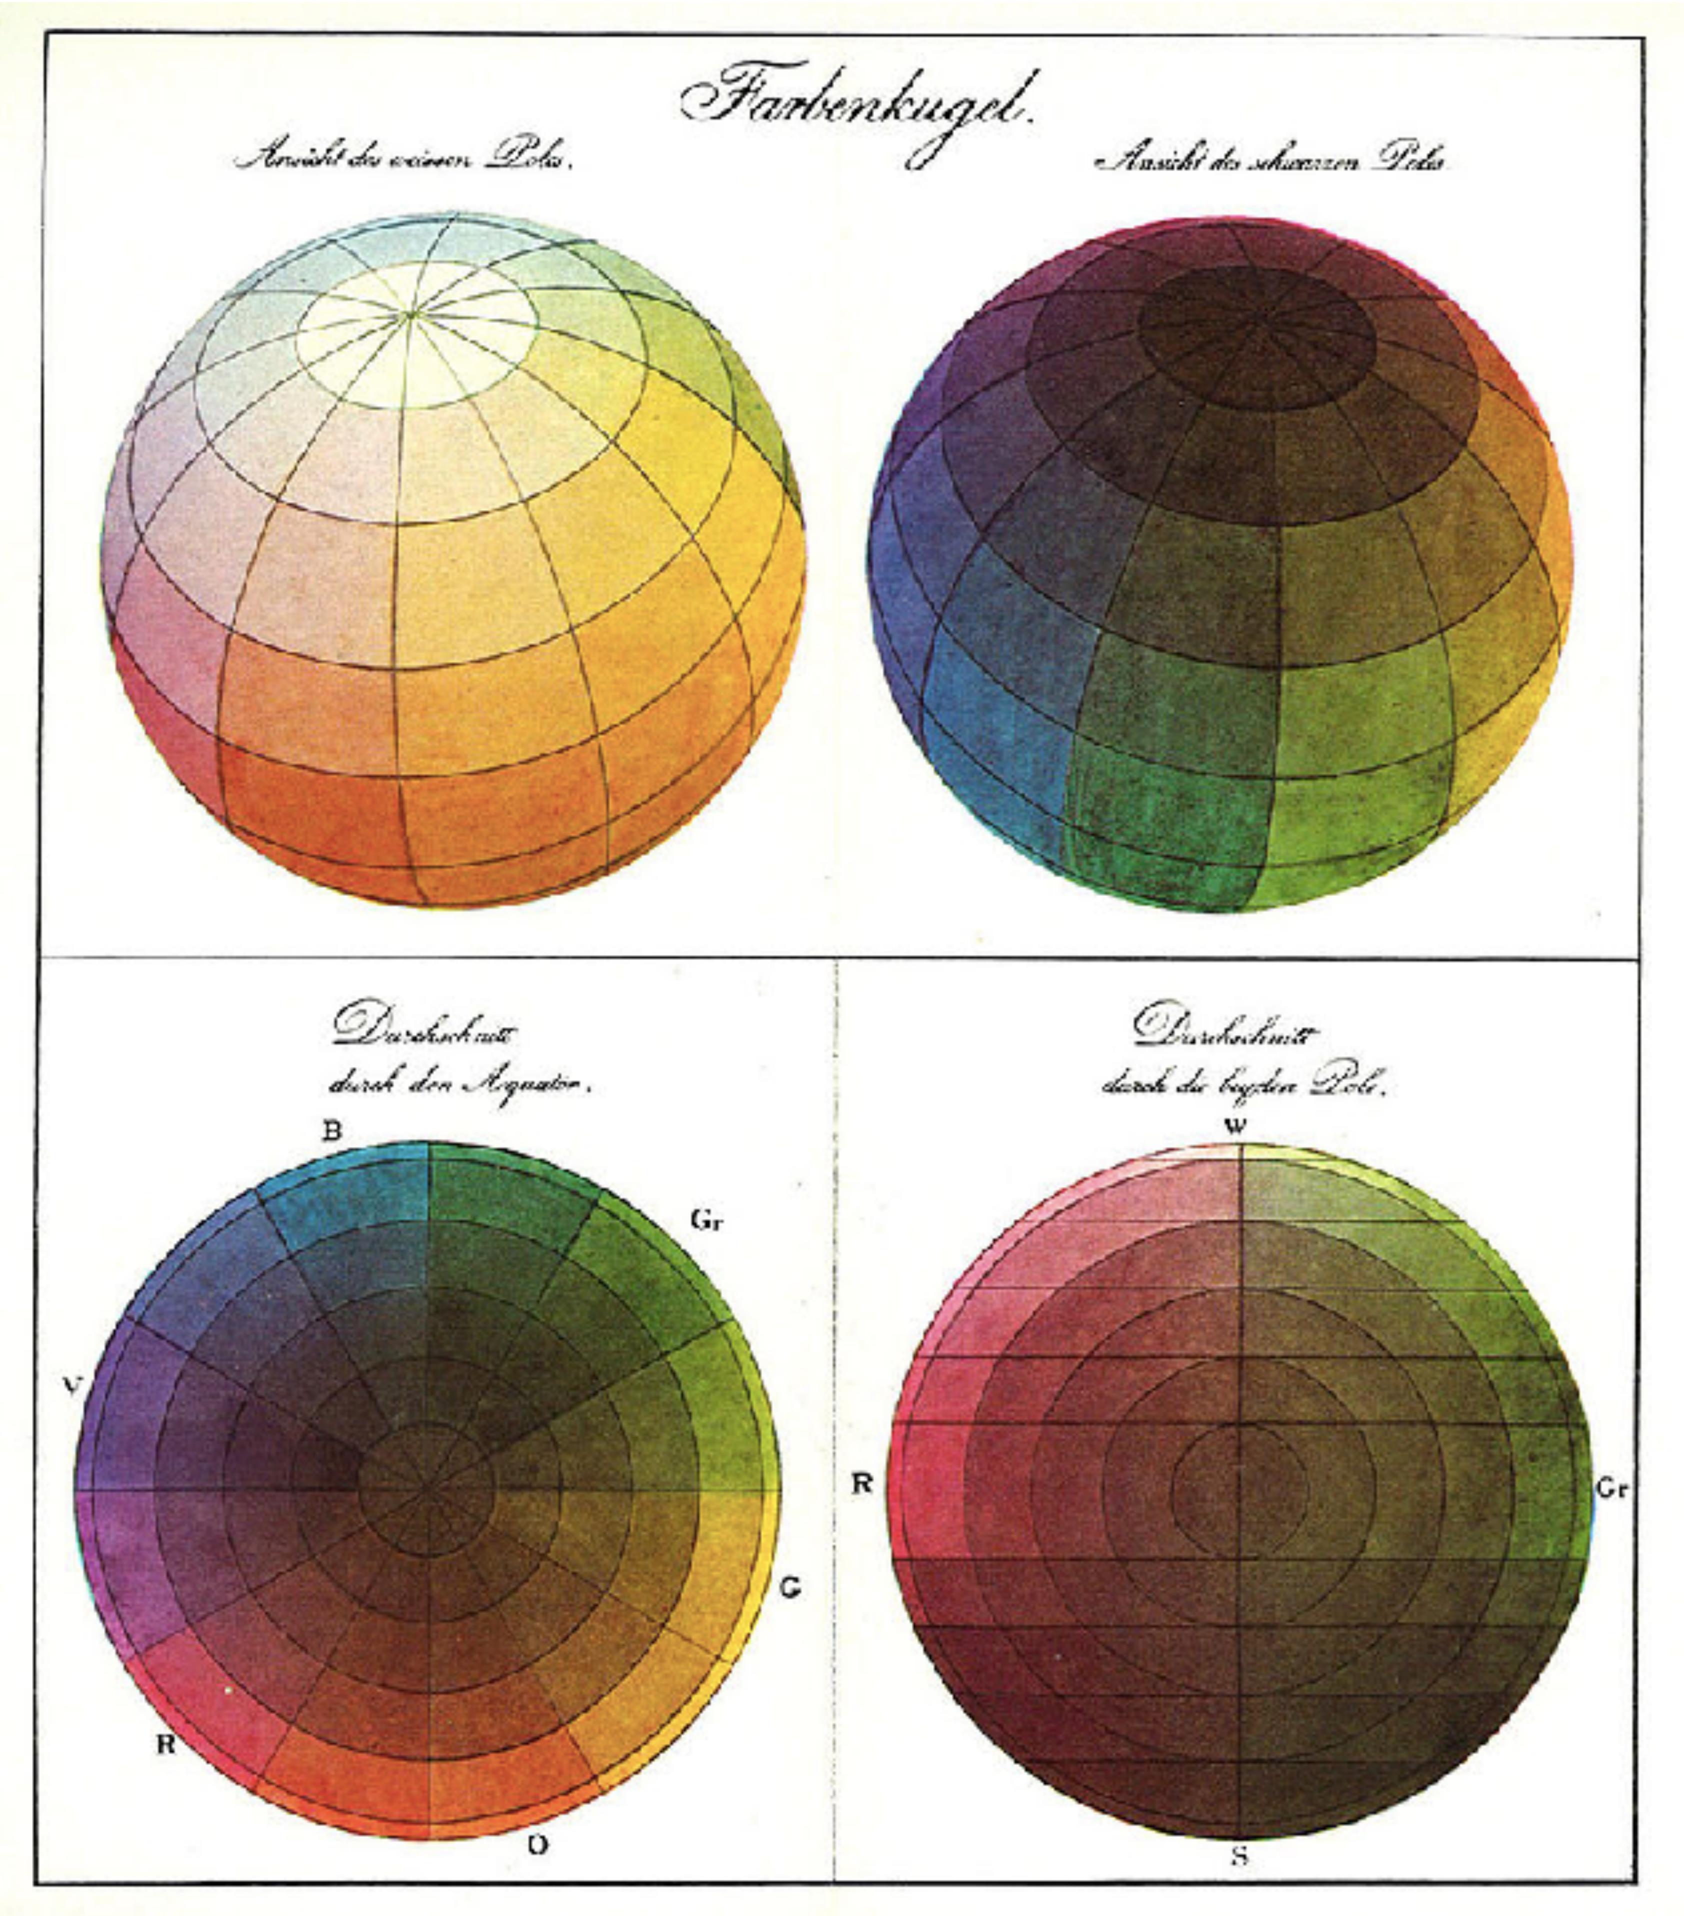

Philipp Otto Runge's Farbenkugel (color sphere), 1810

A BALANCED COLOR SPHERE

Munsell's balanced color sphere, 1900, from A Color Notation, 1905

Color constancy

The ability of the human visual system to perceive color relatively constant despite changes in illumination conditions

We perceive the same color both in shadow and sunlight

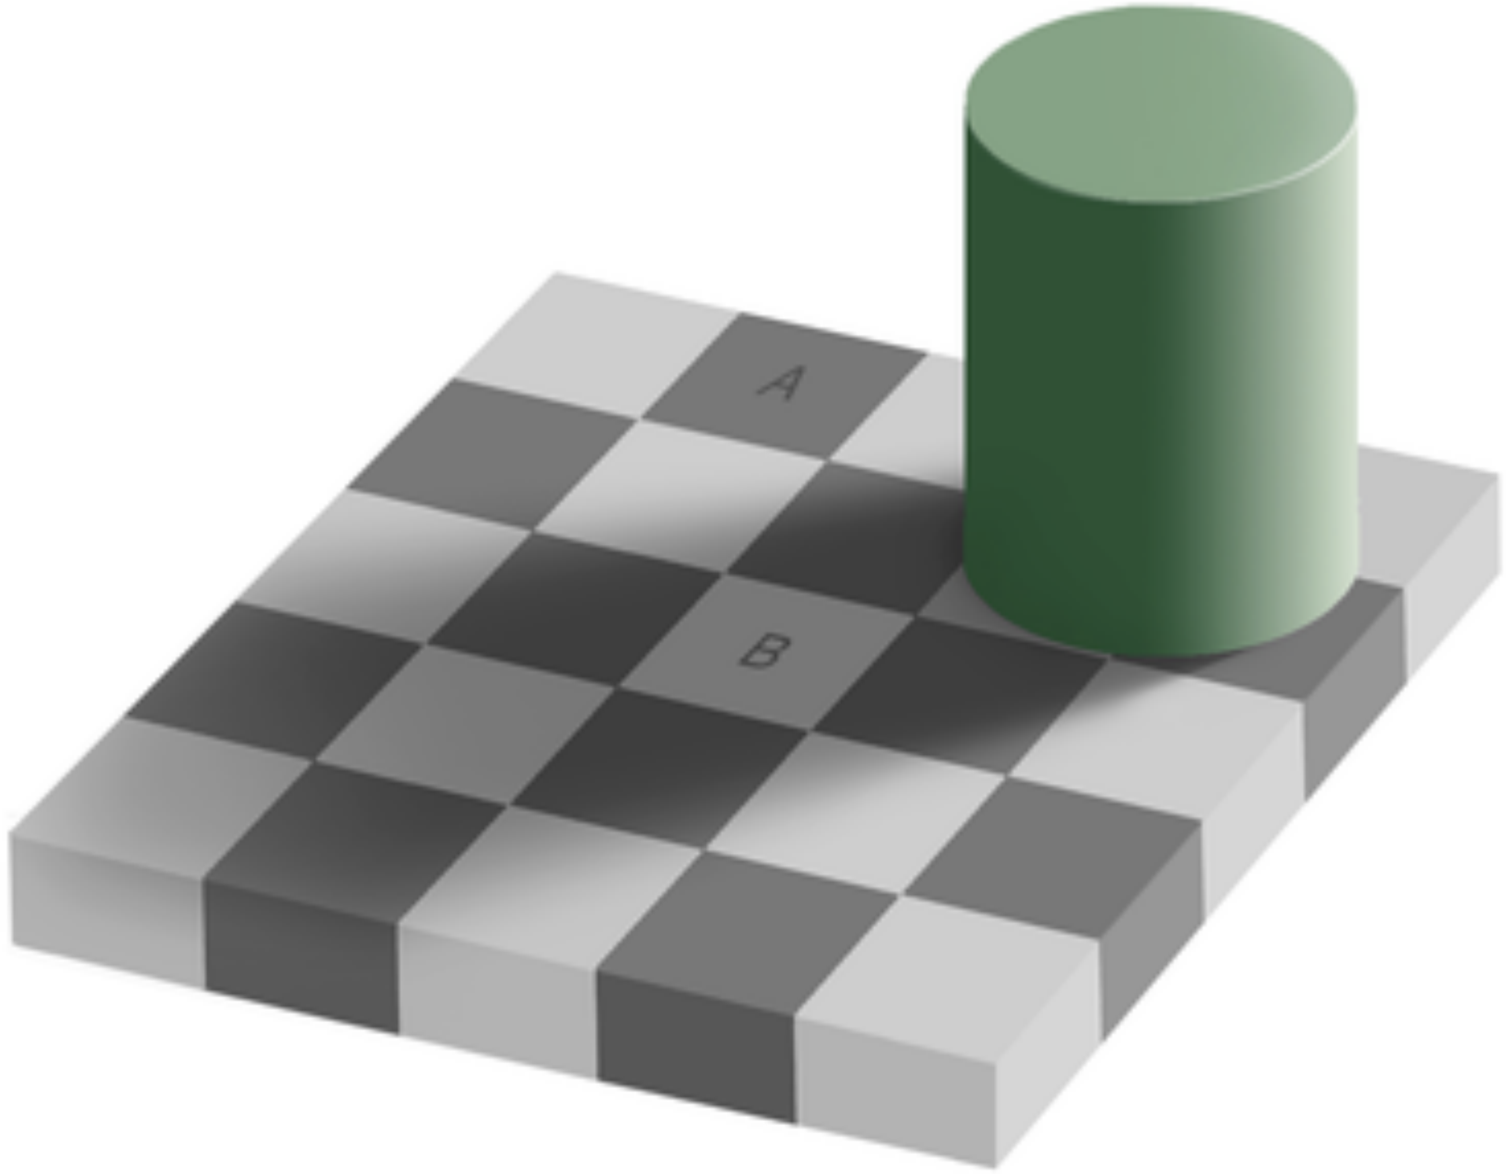

Color constancy causes A and B to look different although the pixel values are the same

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

Interaction of light and surfaces

Reflected color is the result of interaction between the light source spectrum and the reflection surface reflectance

Rough approximation due to other phenomenon

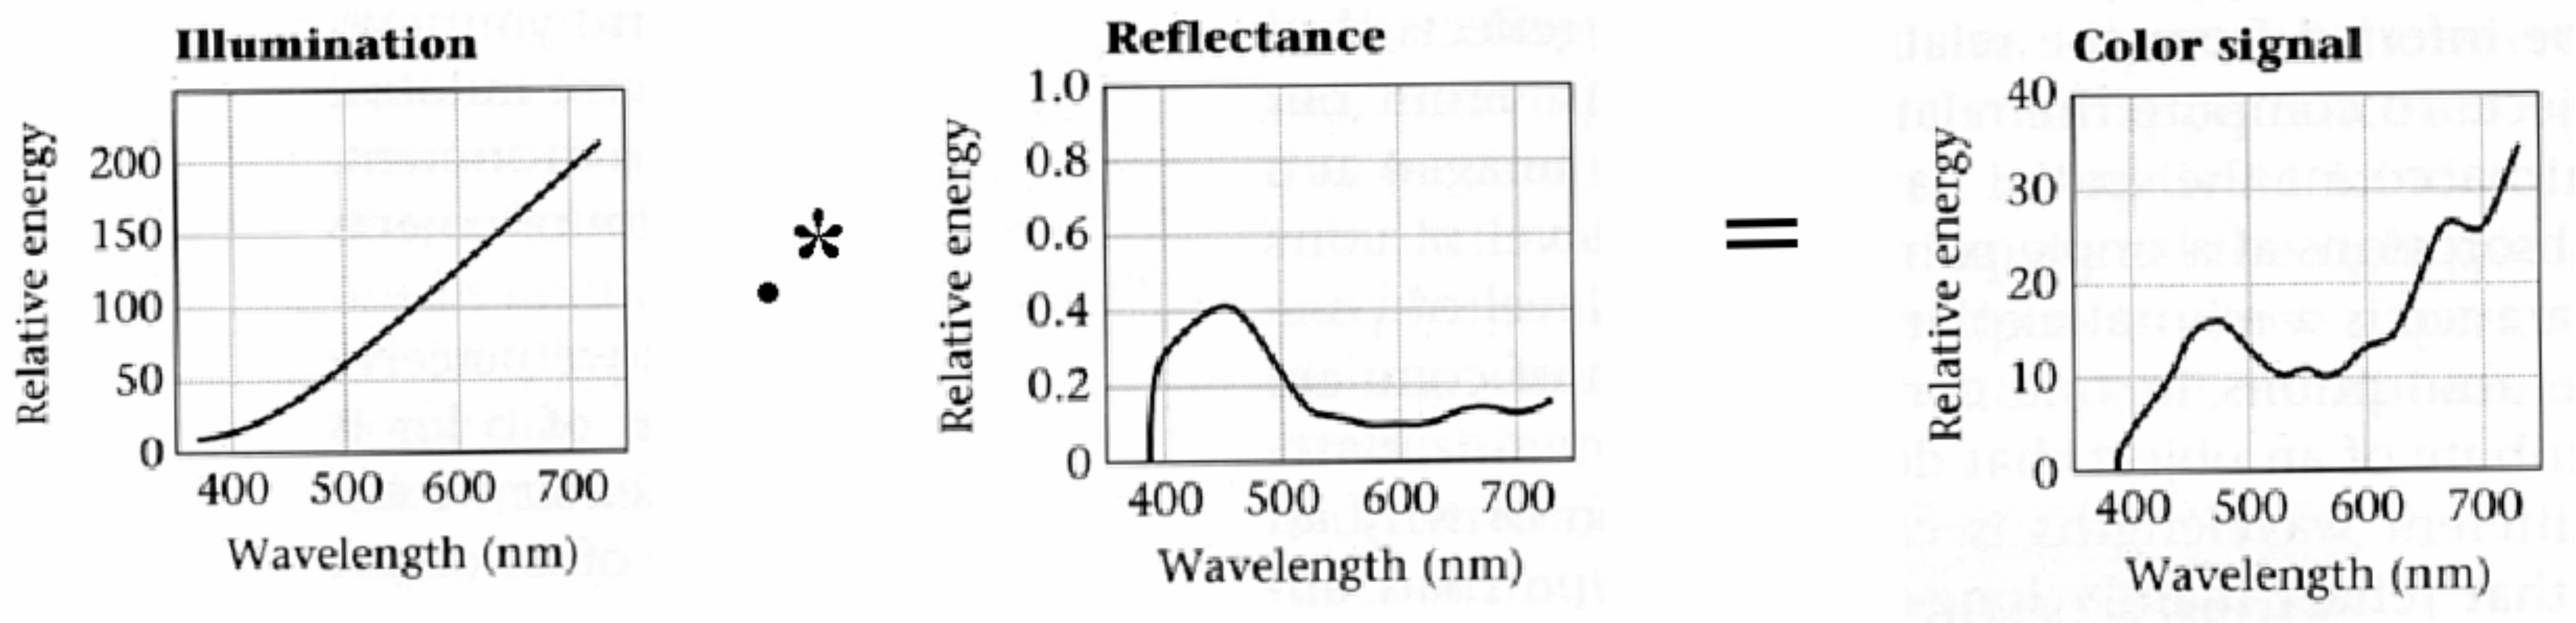

Interaction of light and surfaces

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

Interaction of light and surfaces

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

Room for one color, Olafur Eliasson

Recap: interaction of light and surfaces

Reflected color is the result of interaction between the light source spectrum and the reflection surface reflectance

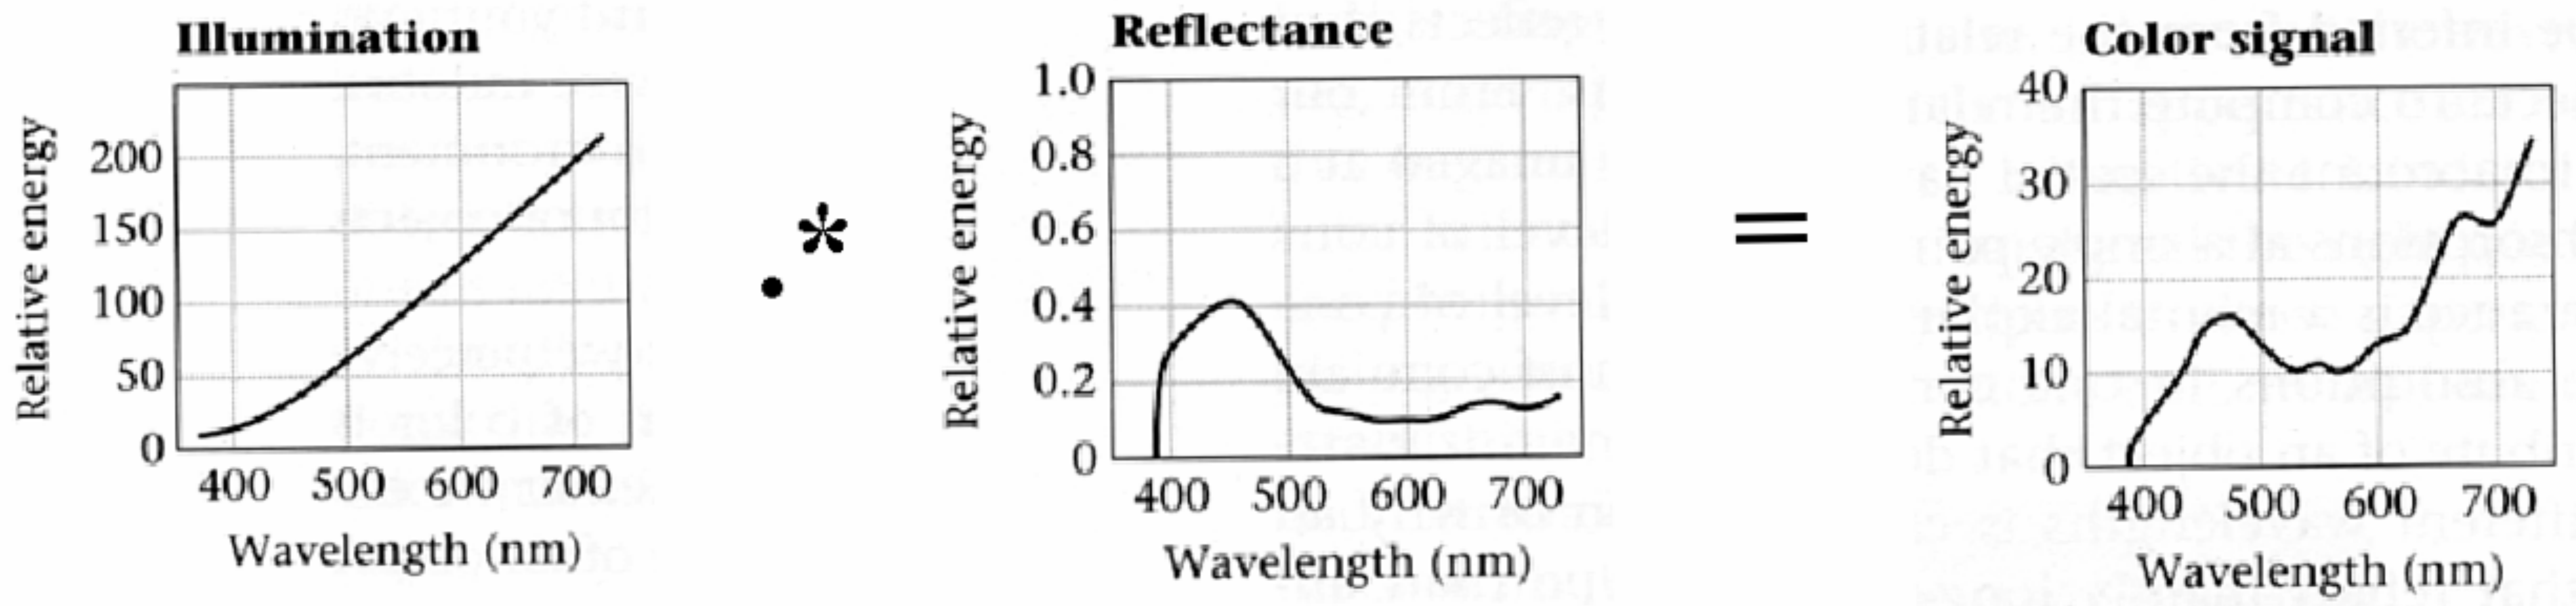

Color constancy

white and gold
or
blue and black

light is blue, white is tinted, blue and gold unchanged

light is yellow, black reflects yellow, blue unchanged

Color constancy

http://xkcd.com/1492/

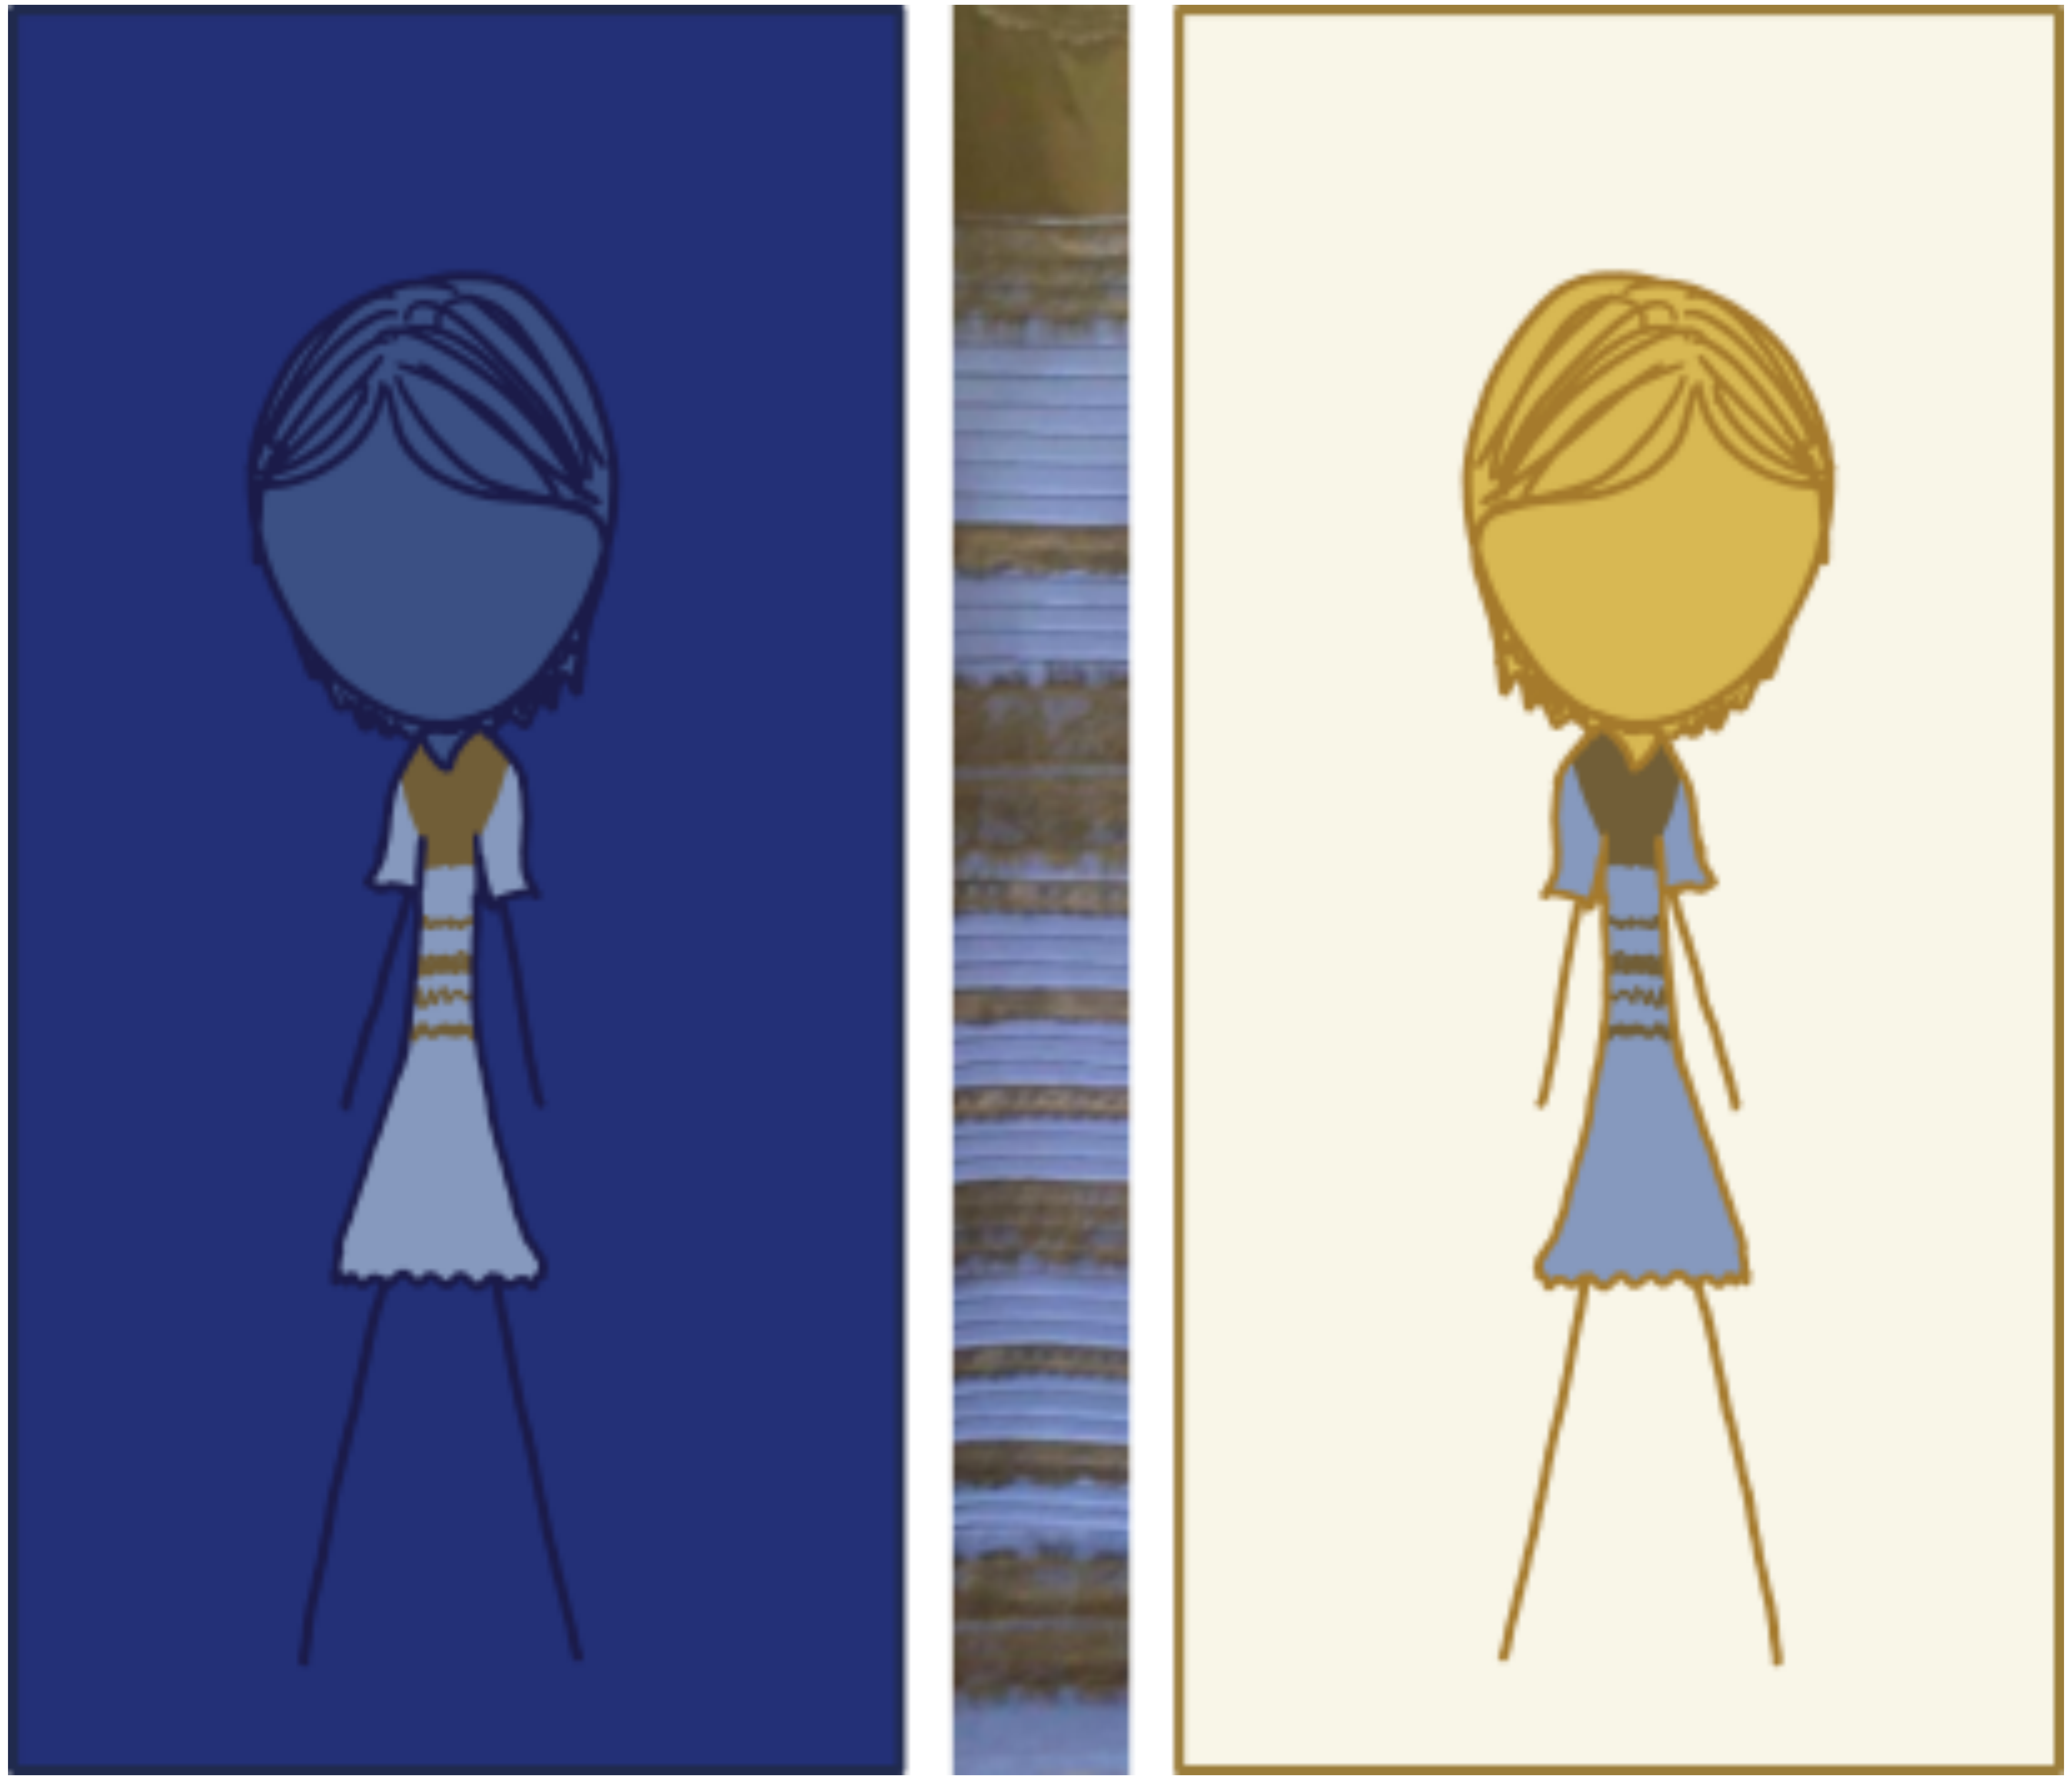

Color constancy

http://xkcd.com/1492/

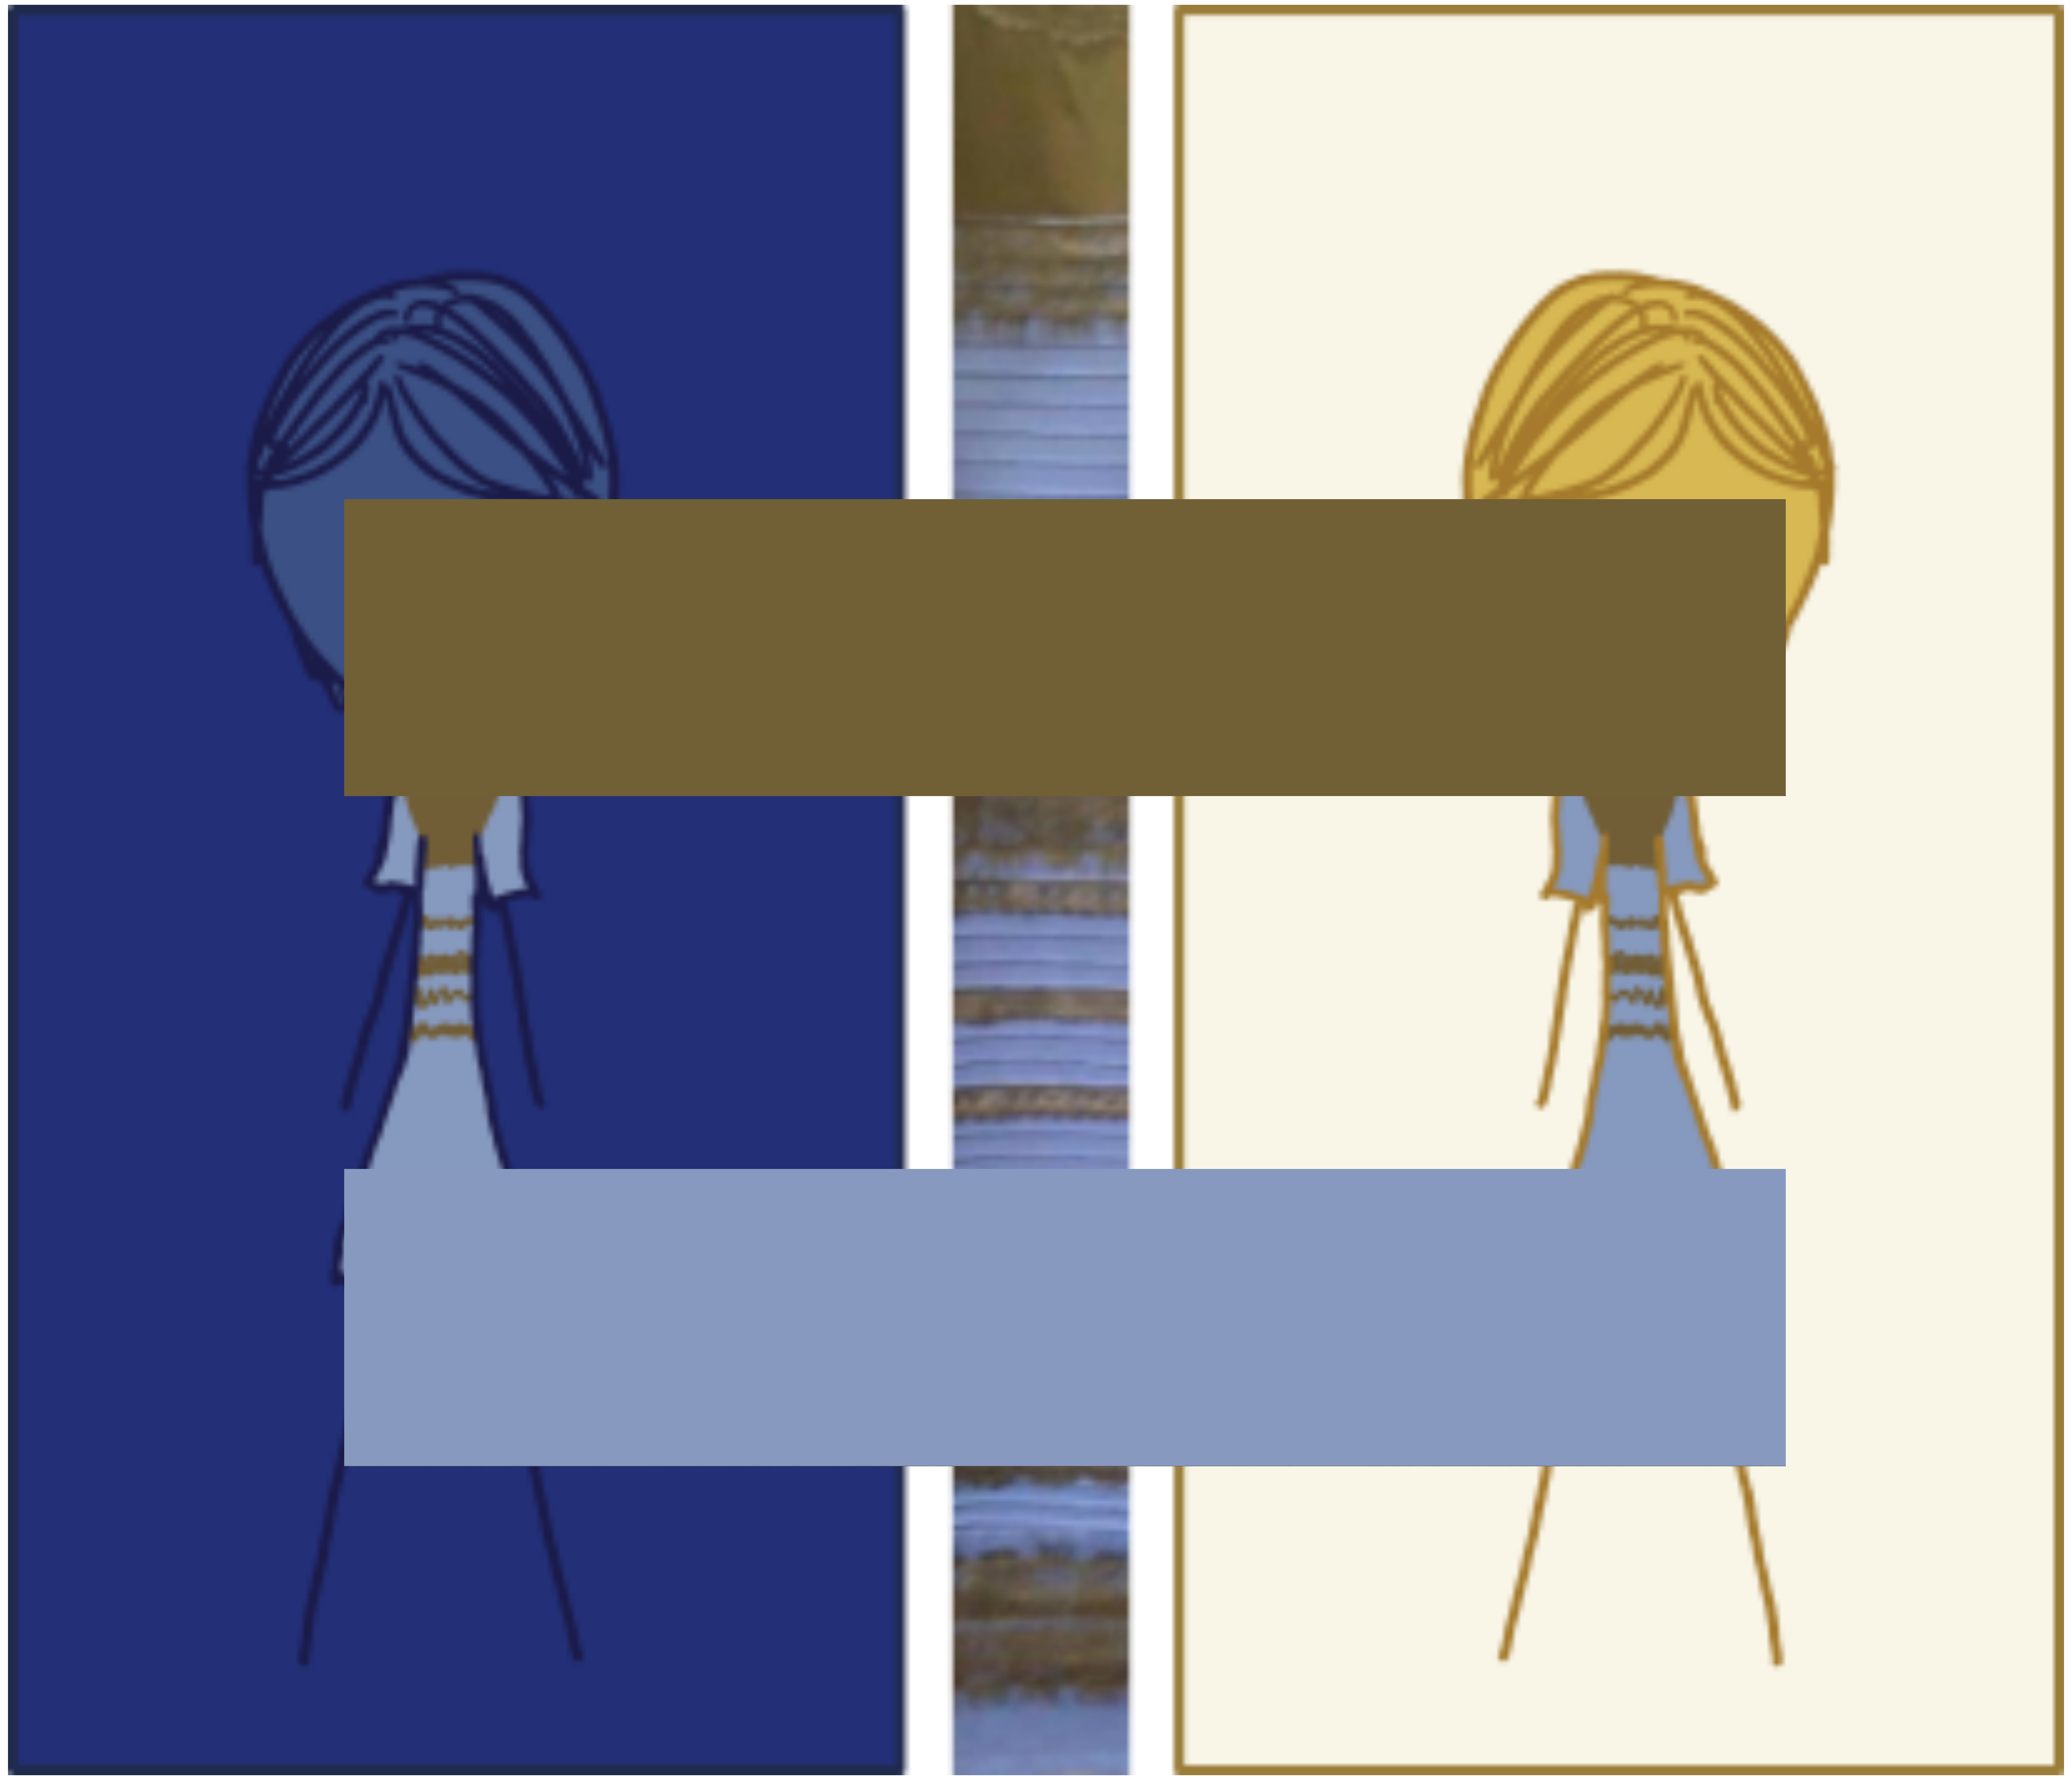

Chromatic adaptation

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

White balance

When looking at a picture on screen or print, our eyes are adapted 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

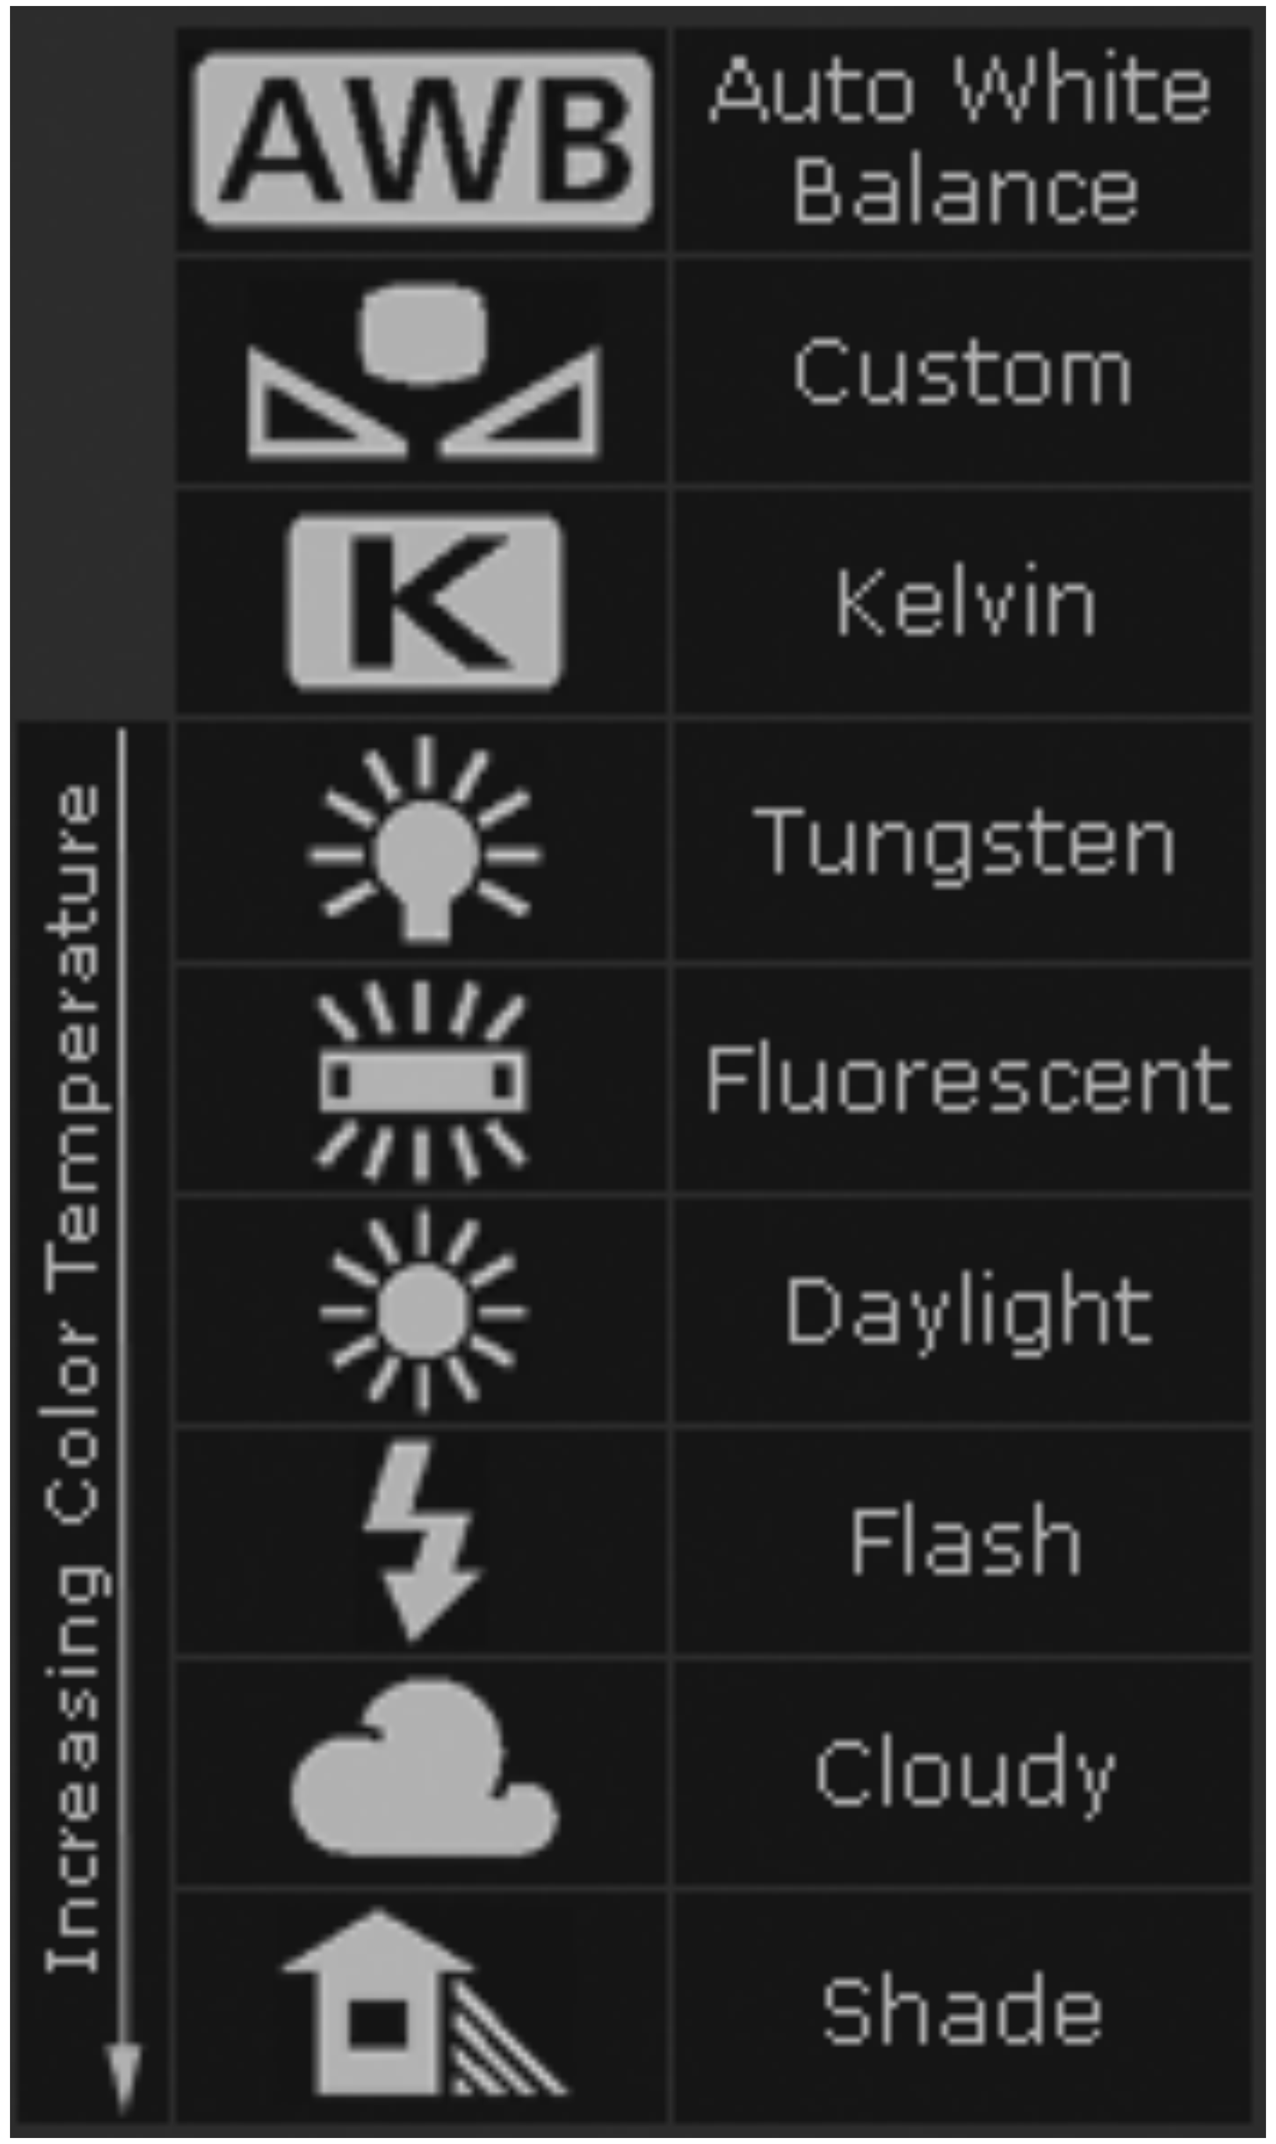

http://www.cambridgeincolour.com/tutorials/white-balance.htm

White balance

Von Kries adaptation

- Multiply each channel by a gain factor

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 $1/r_w$, $1/g_w$, $1/b_w$

Source: L. Lazebnik

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 natural image statistics

Source: L. Lazebnik

Failure cases

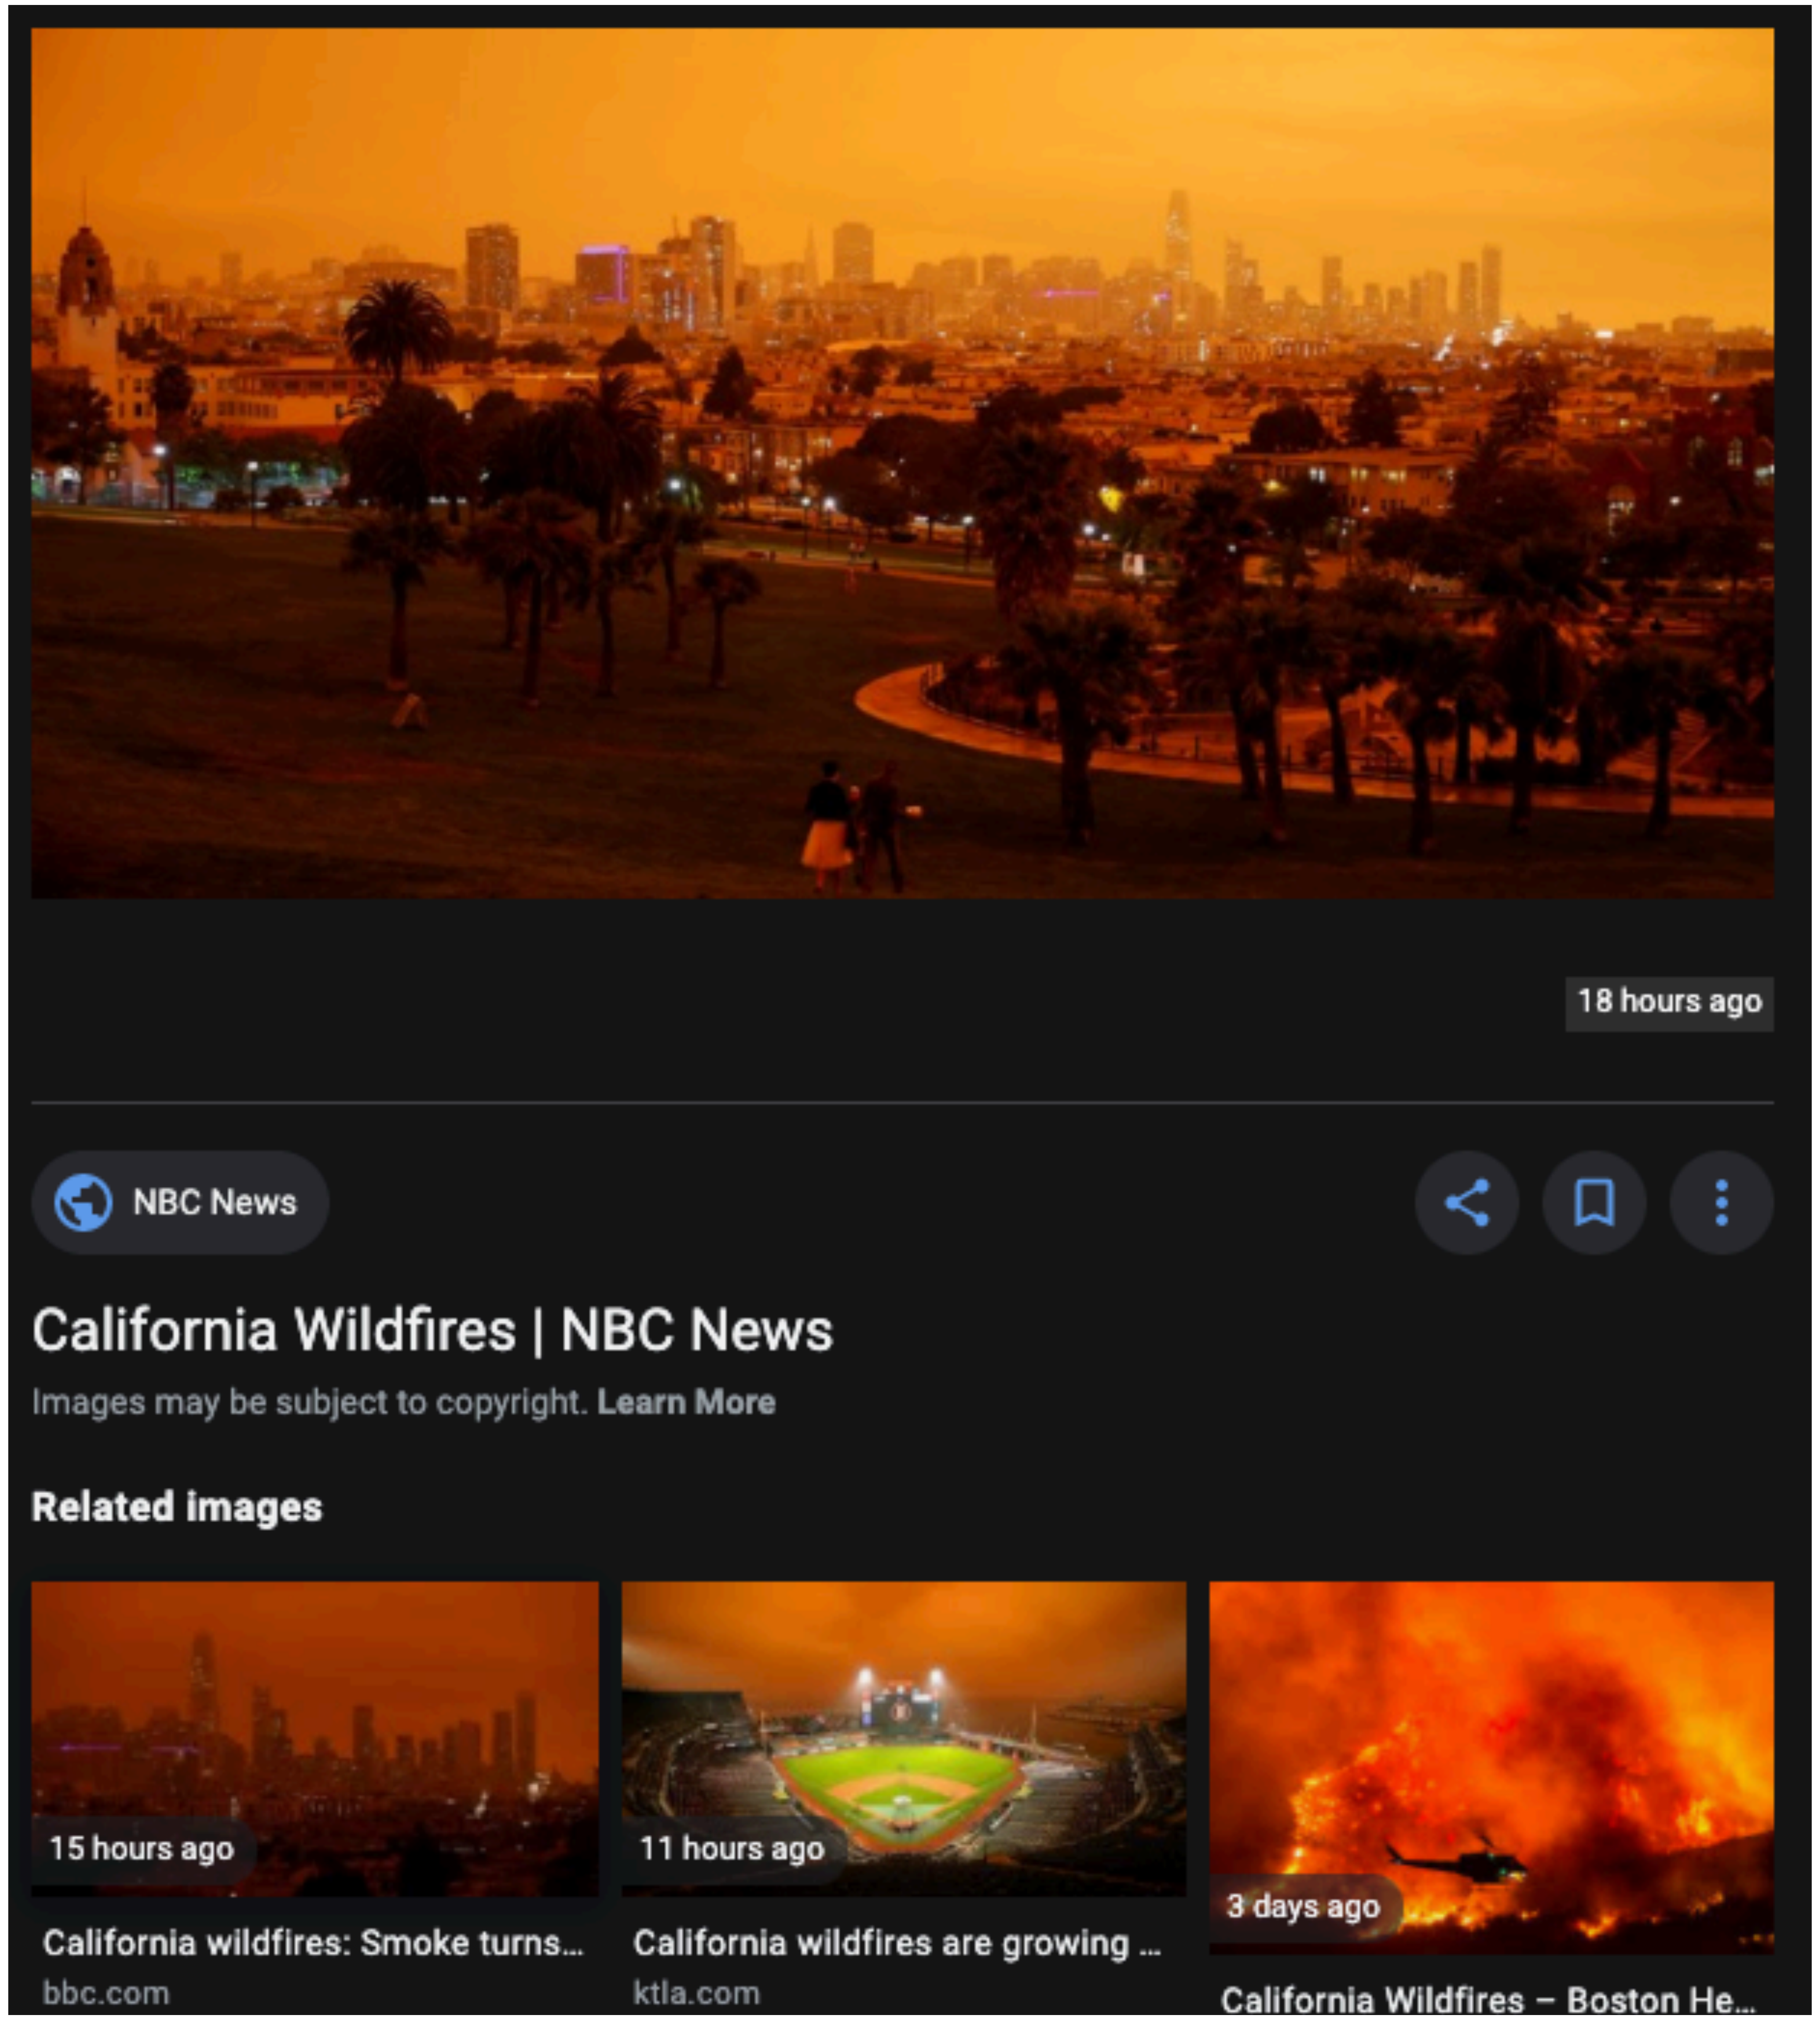

Color and language

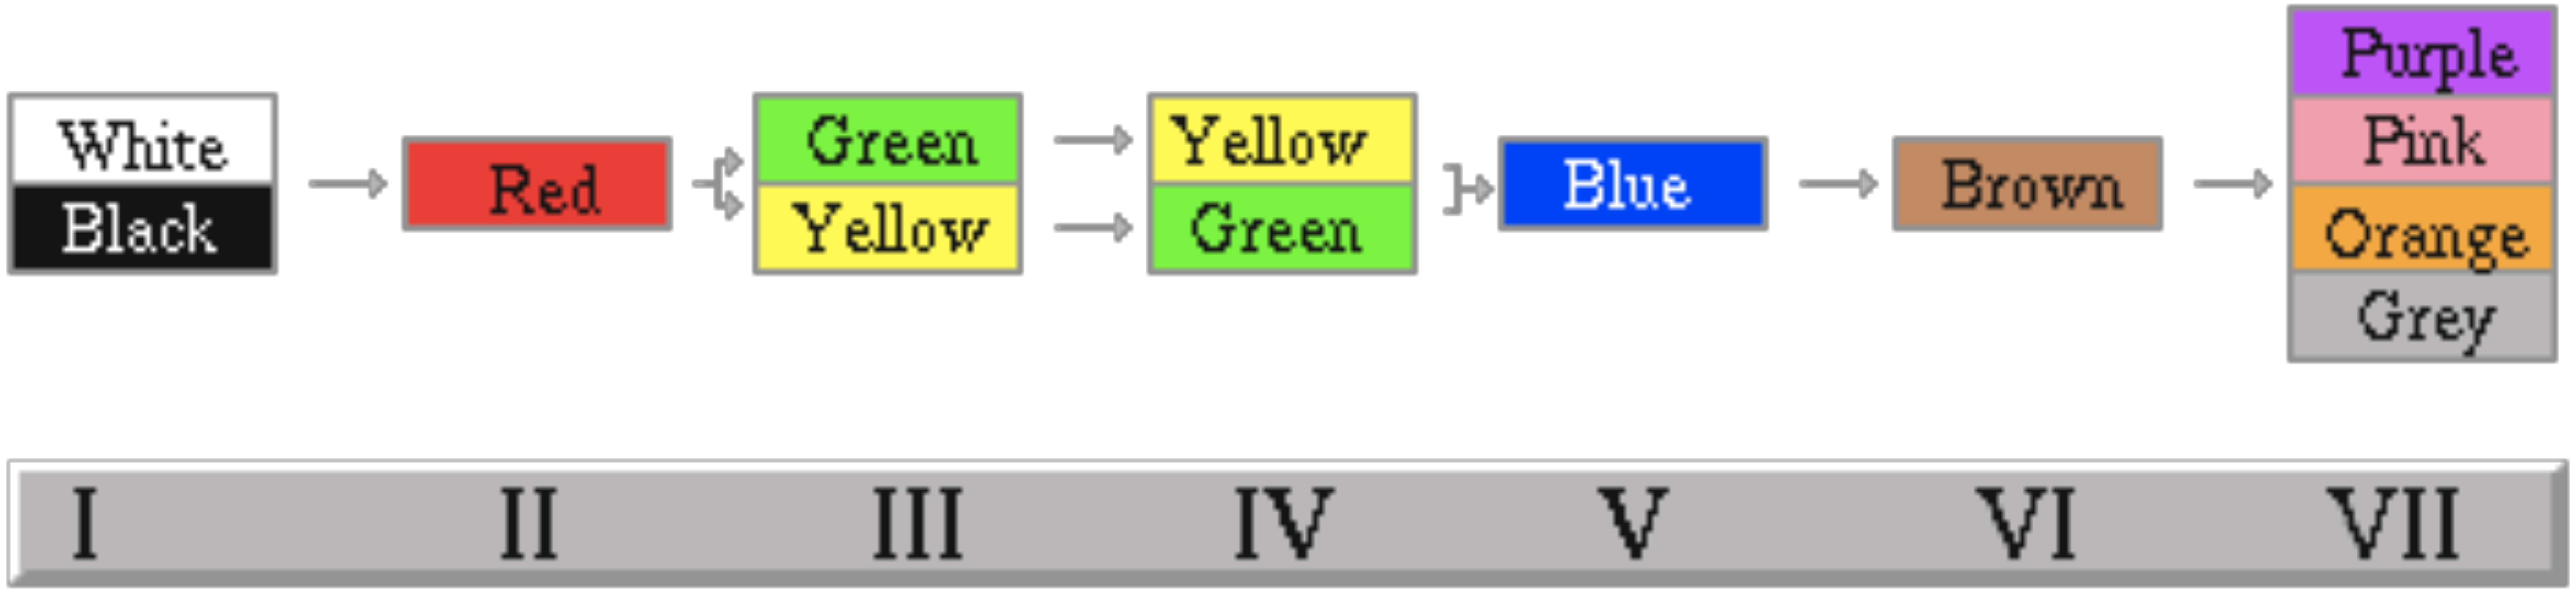

Evolution of color terms across ~20 diverse languages

B. Berlin and P. Kay, Basic Color Terms: Their Universality and Evolution (1969)

Grayscale to color

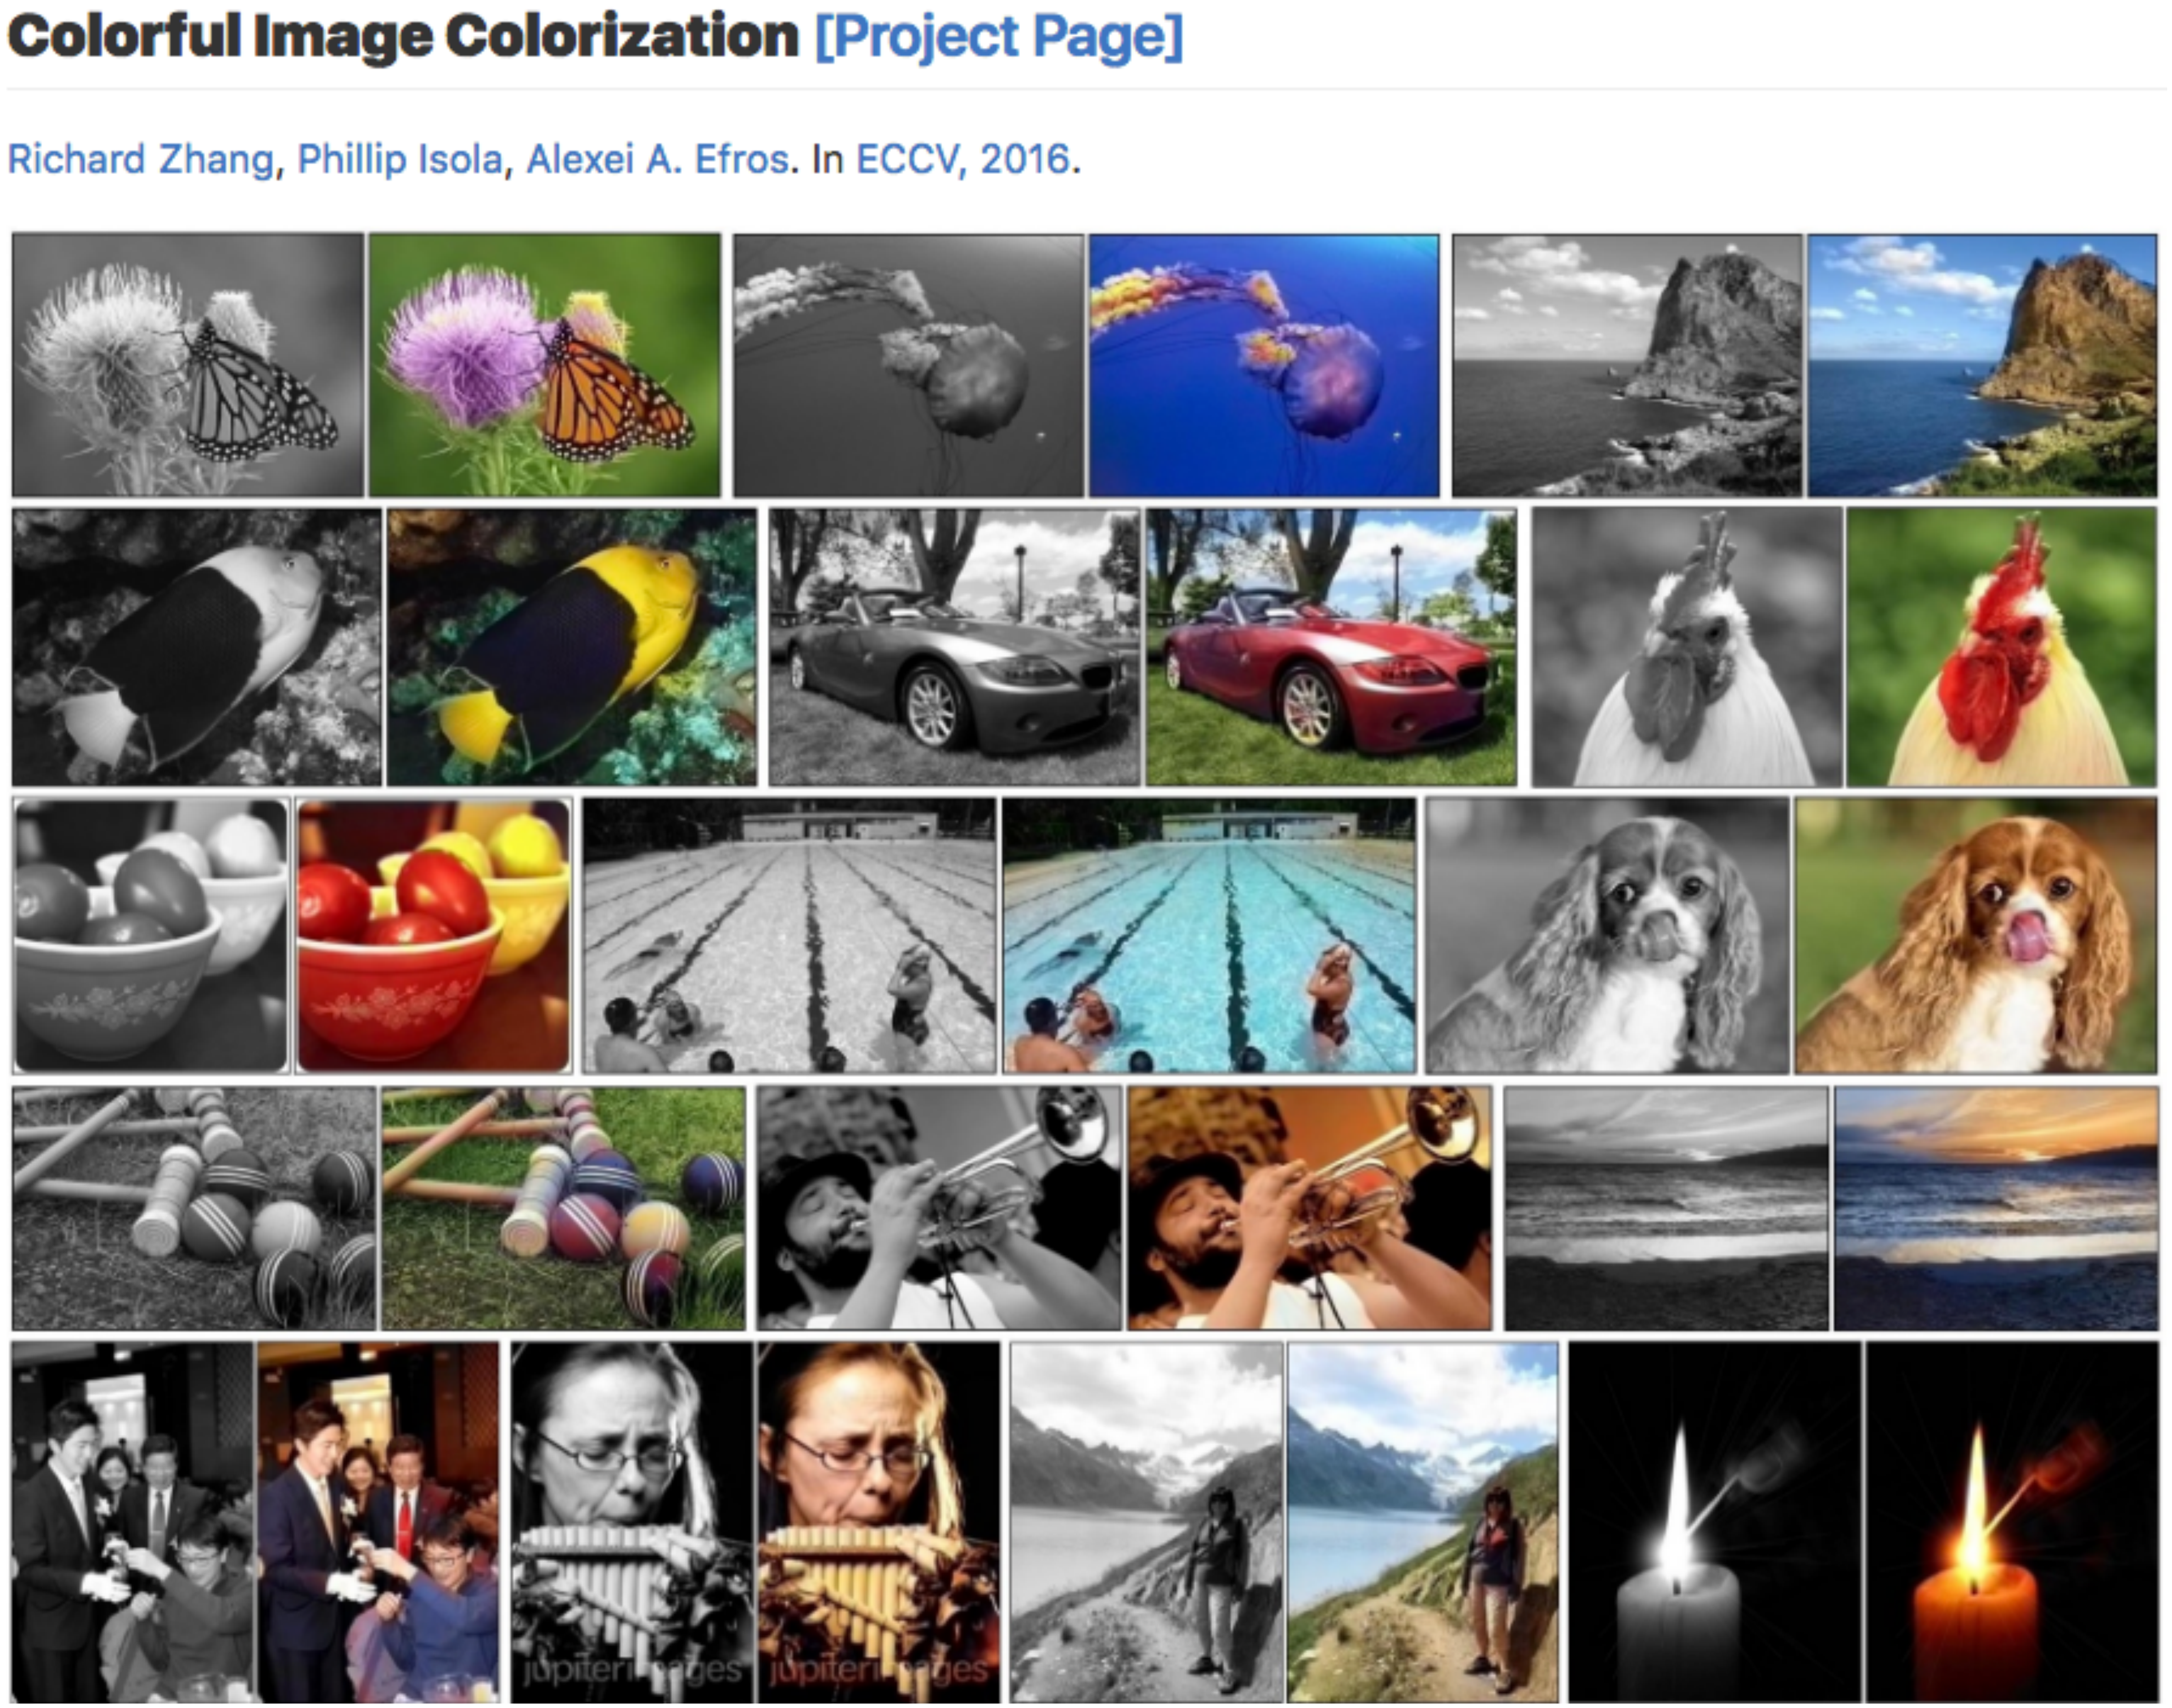

https://github.com/richzhang/colorization

Further readings and thoughts …

Color matching applet

- http://graphics.stanford.edu/courses/cs178/applets/colormatching.html

B. Berlin and P. Kay, Basic Color Terms: Their Universality and Evolution (1969)

- It is a book. The library has some copies.

D.A. Forsyth, A novel algorithm for color constancy

- Gamut based approach
- http://luthuli.cs.uiuc.edu/~daf/papers/colorconst.pdf

Recent work on grayscale images to color using CNNs

- Turns out that this is a good way to learn representations without supervision!

Watch this talk …

https://www.ted.com/talks/beau_lotto_optical_illusions_show_how_we_see/up-next