COMPUTATIONAL ASPECTS OF DIGITAL PHOTOGRAPHY

Light & Color (continued)

Wojciech Jarosz

wojciech.k.jarosz@dartmouth.edu

Dartmouth

Administrivia

Assignment 2 available now

- back to programming
- due next Wednesday

Last time…

Light & Color

- Physics background
- Color perception & measurement
- Color reproduction
- Color spaces

What is light?

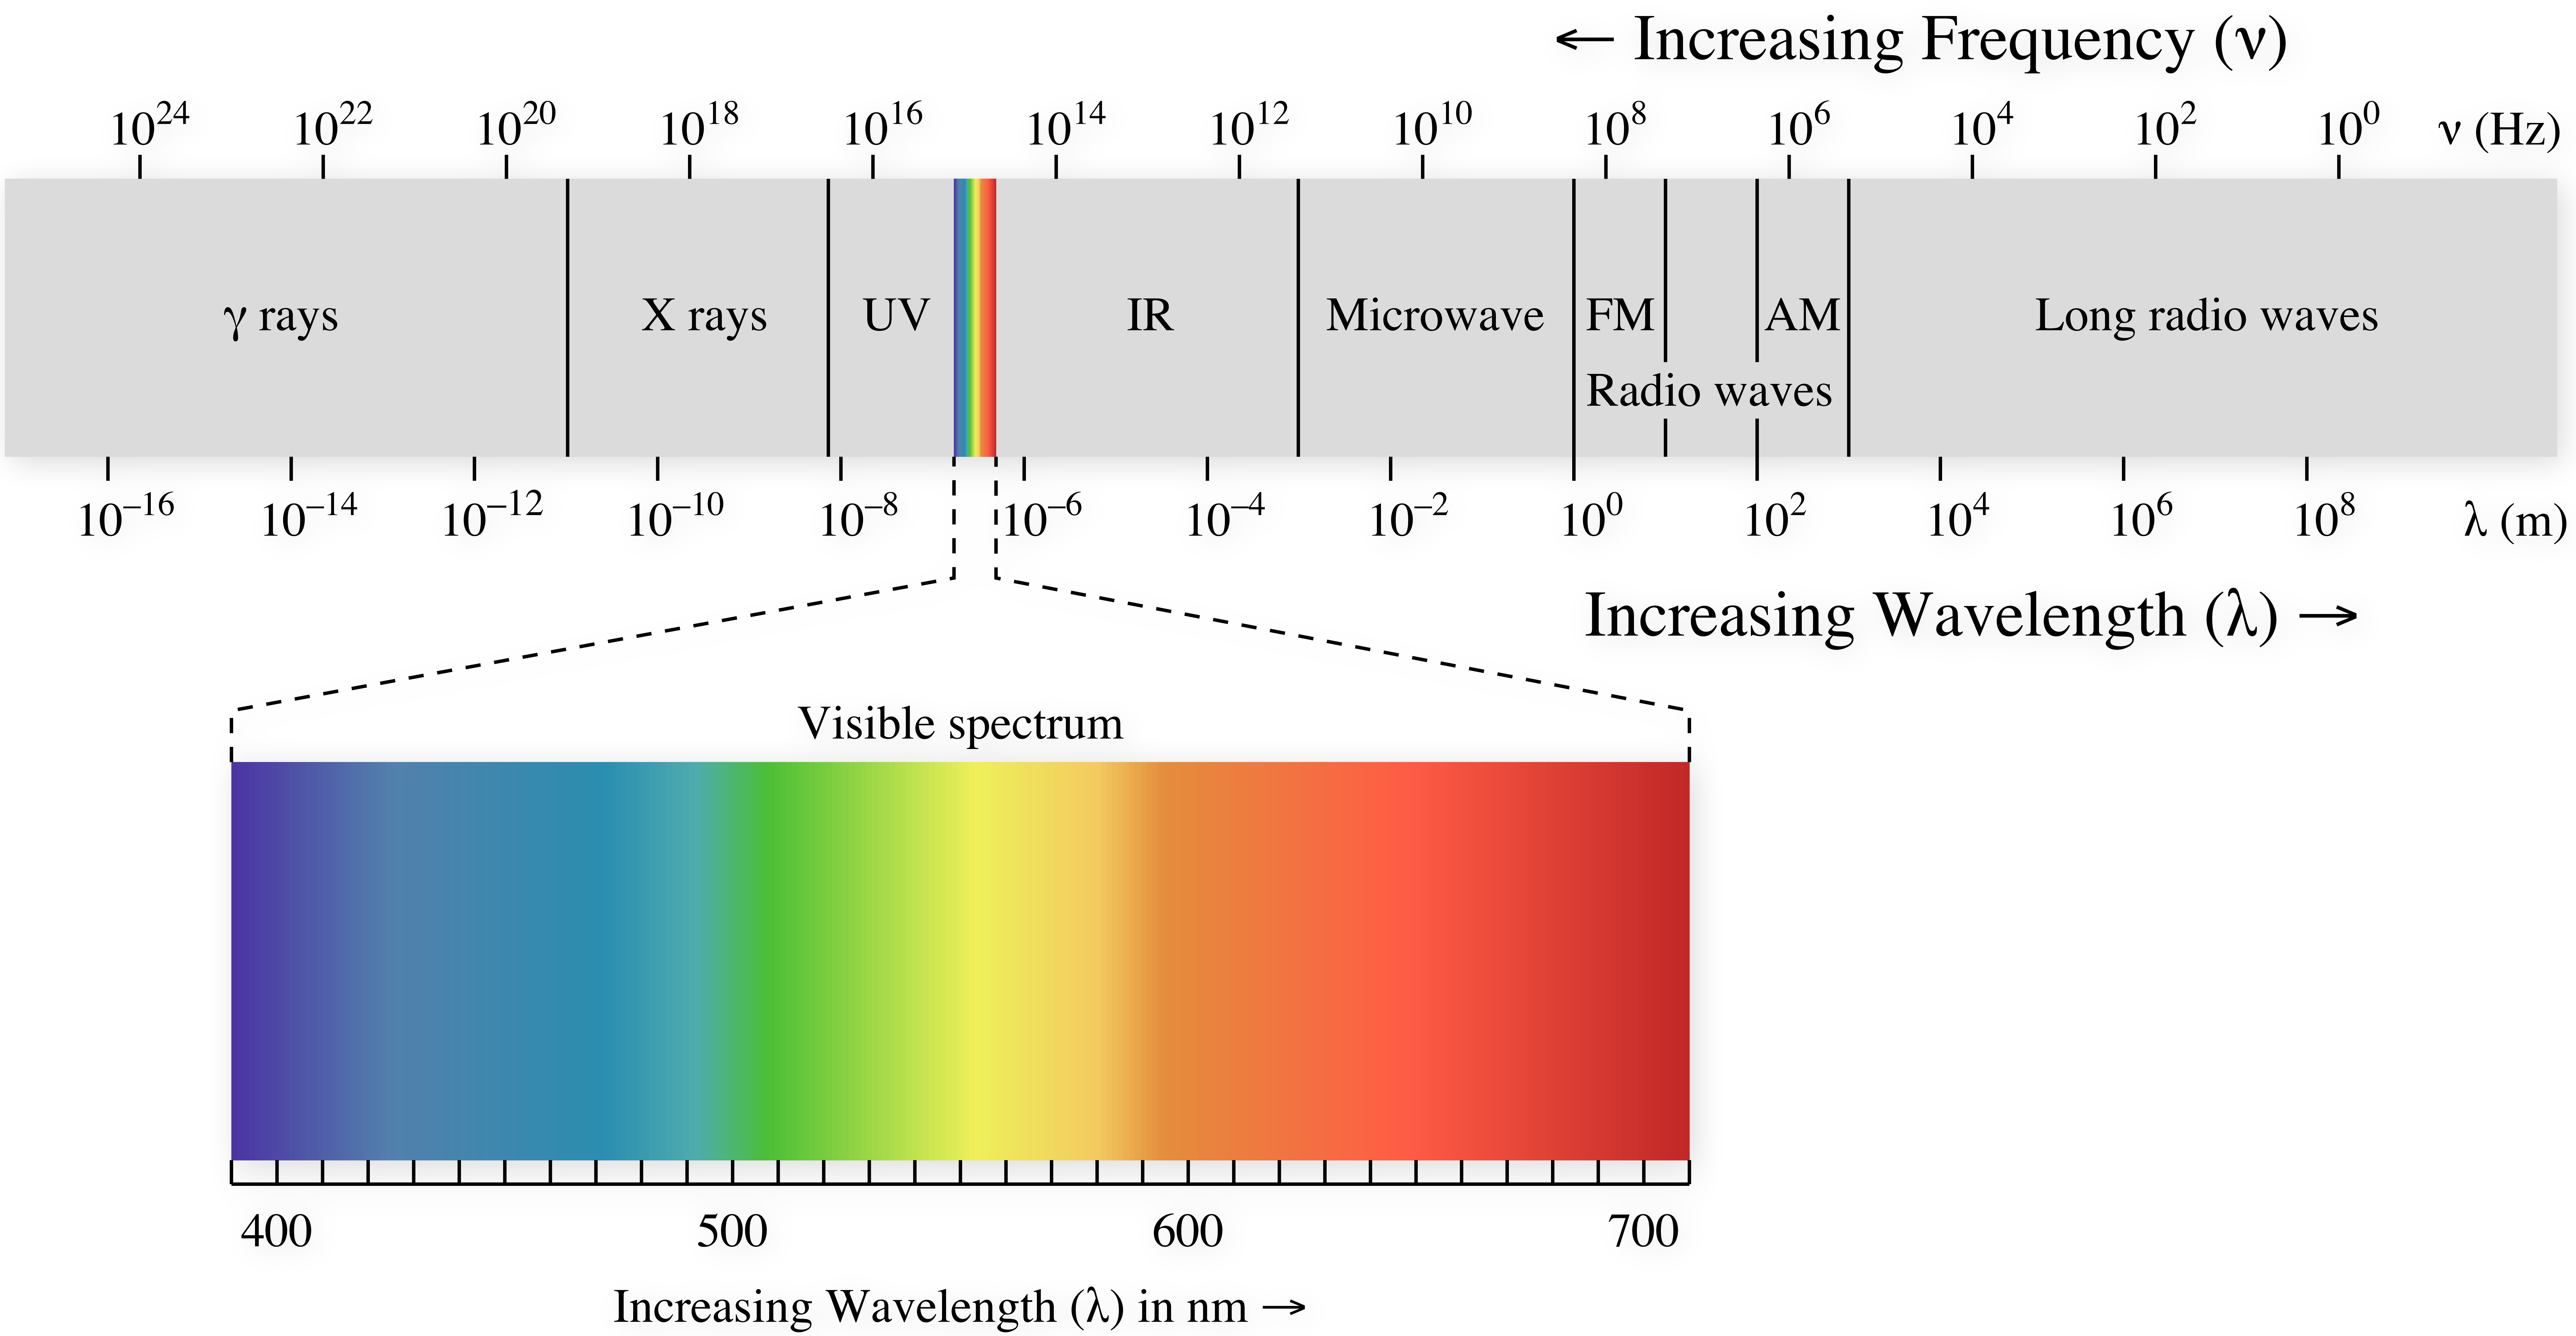

Spectral distribution function (SPD)

Light can be a mixture of many wavelengths

- SPD: intensity as a function of wavelength over enter spectrum

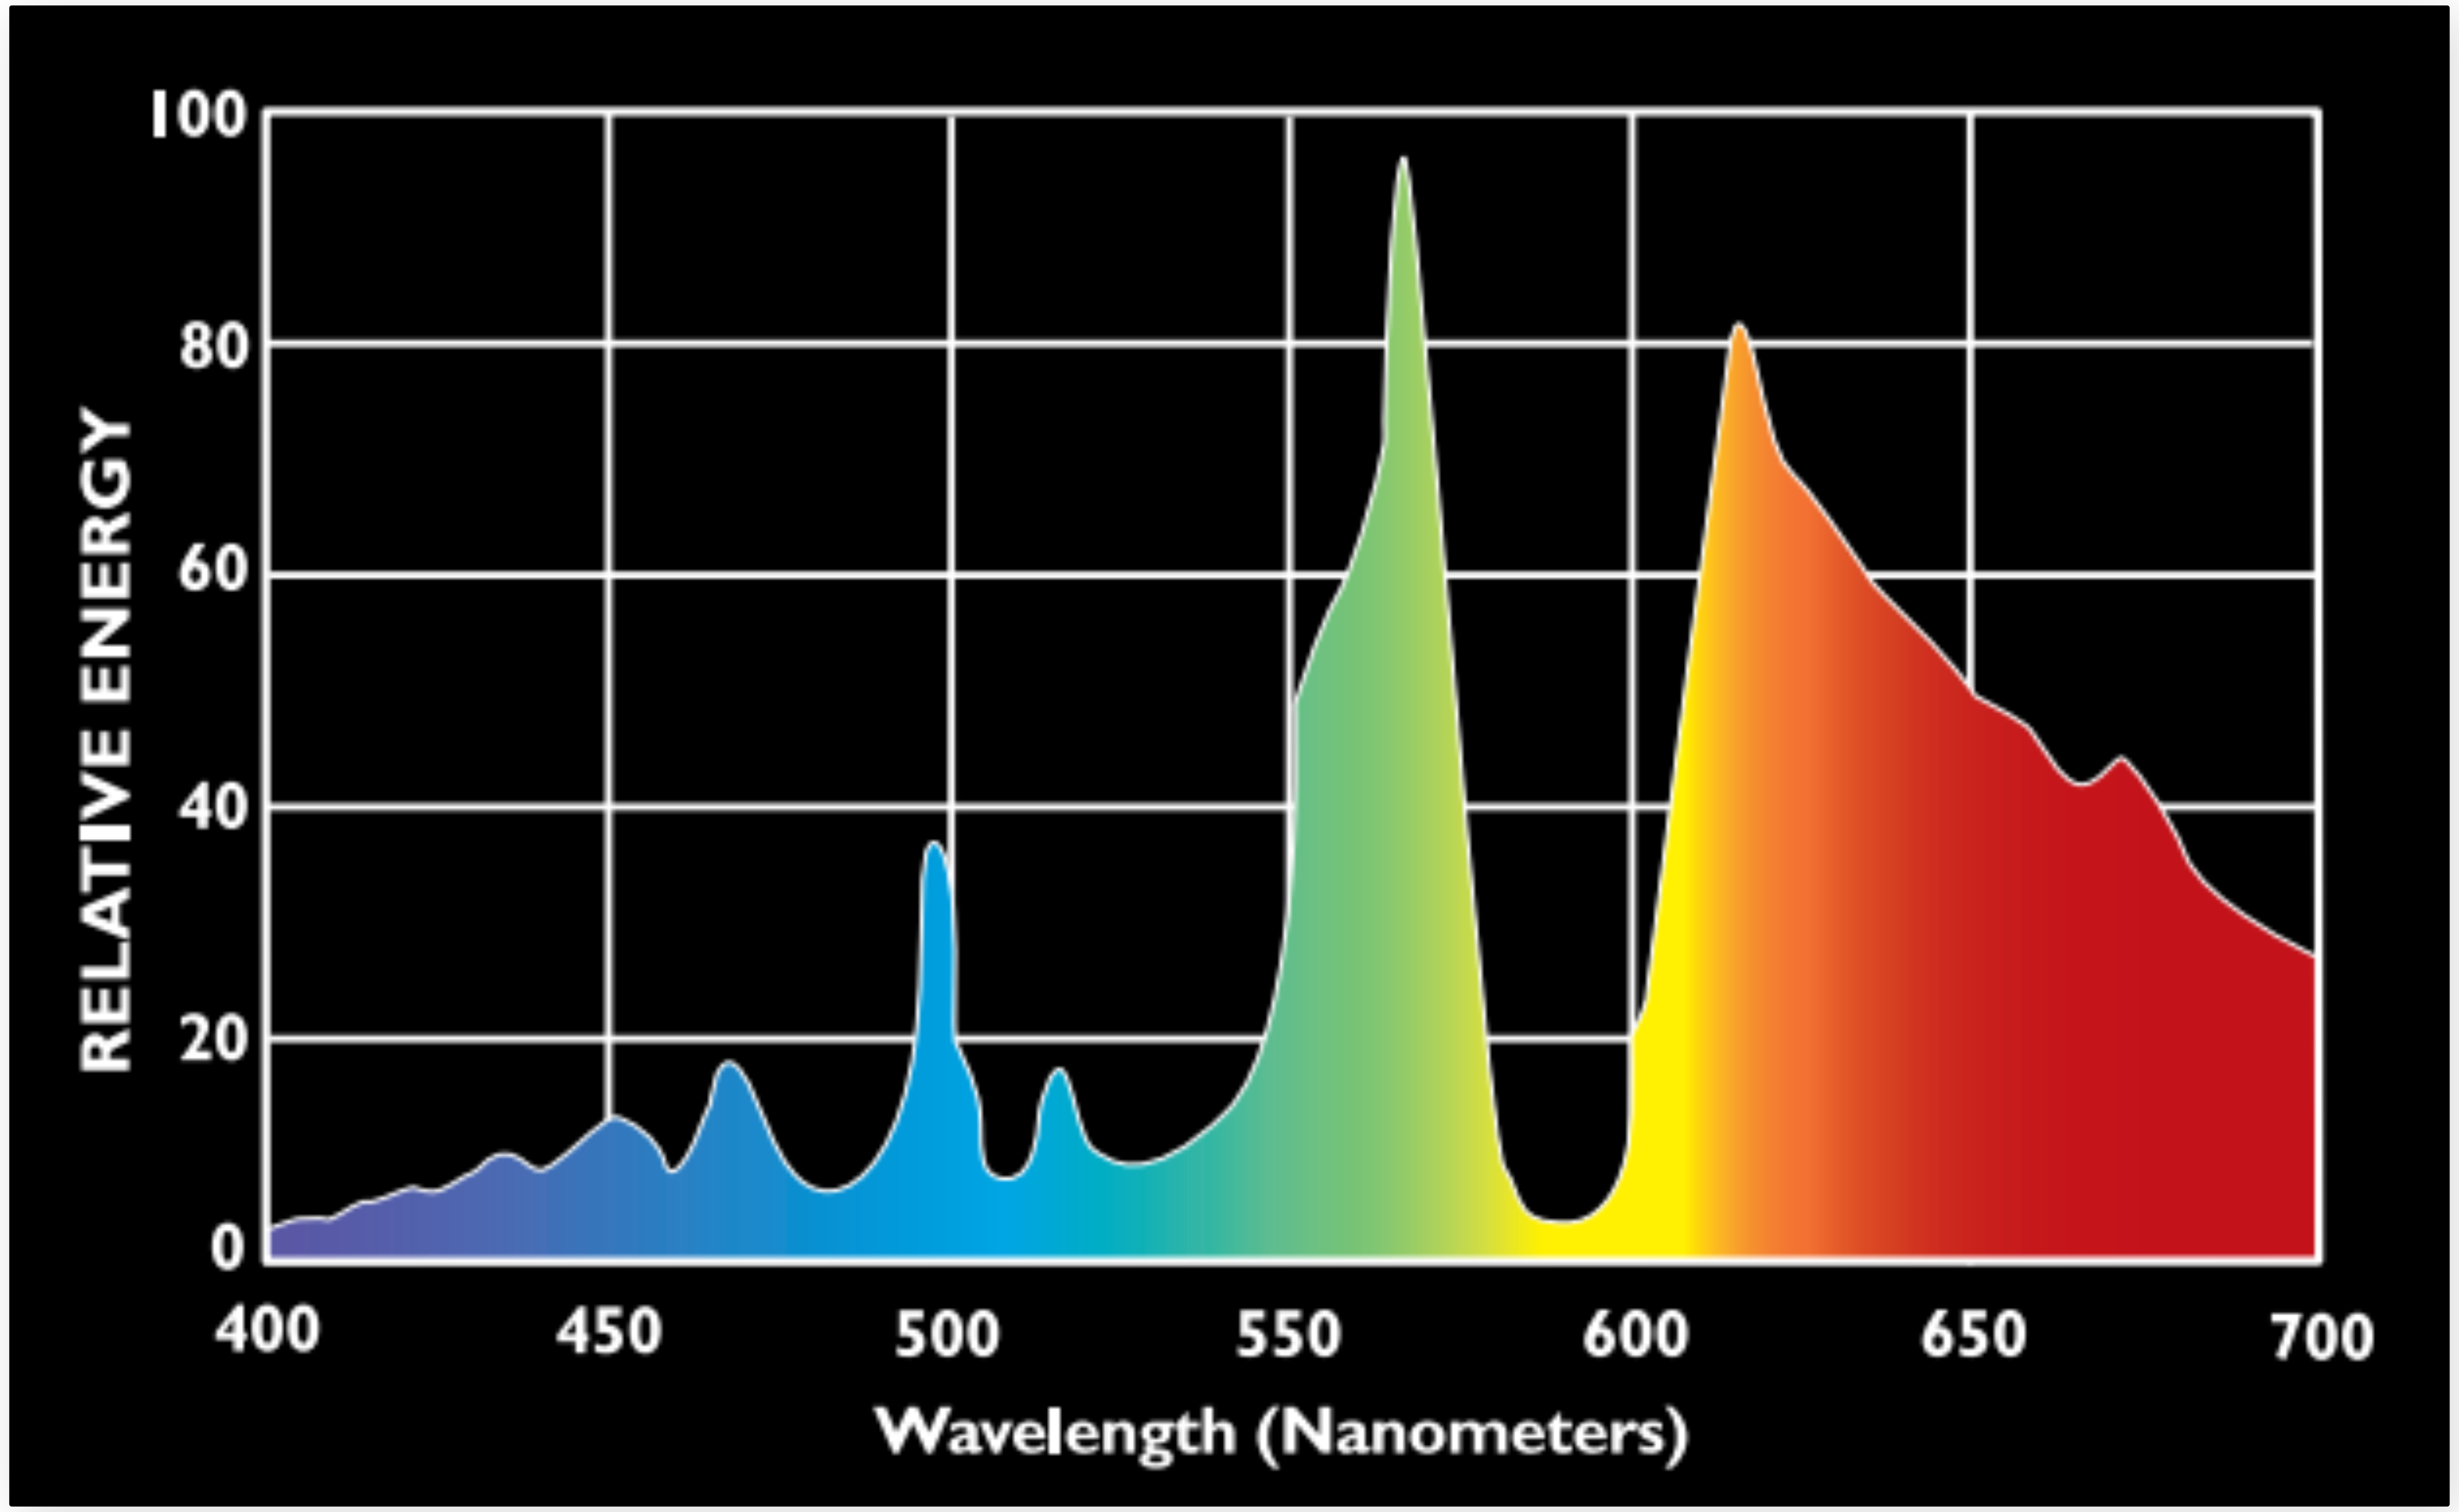

Light-matter interaction

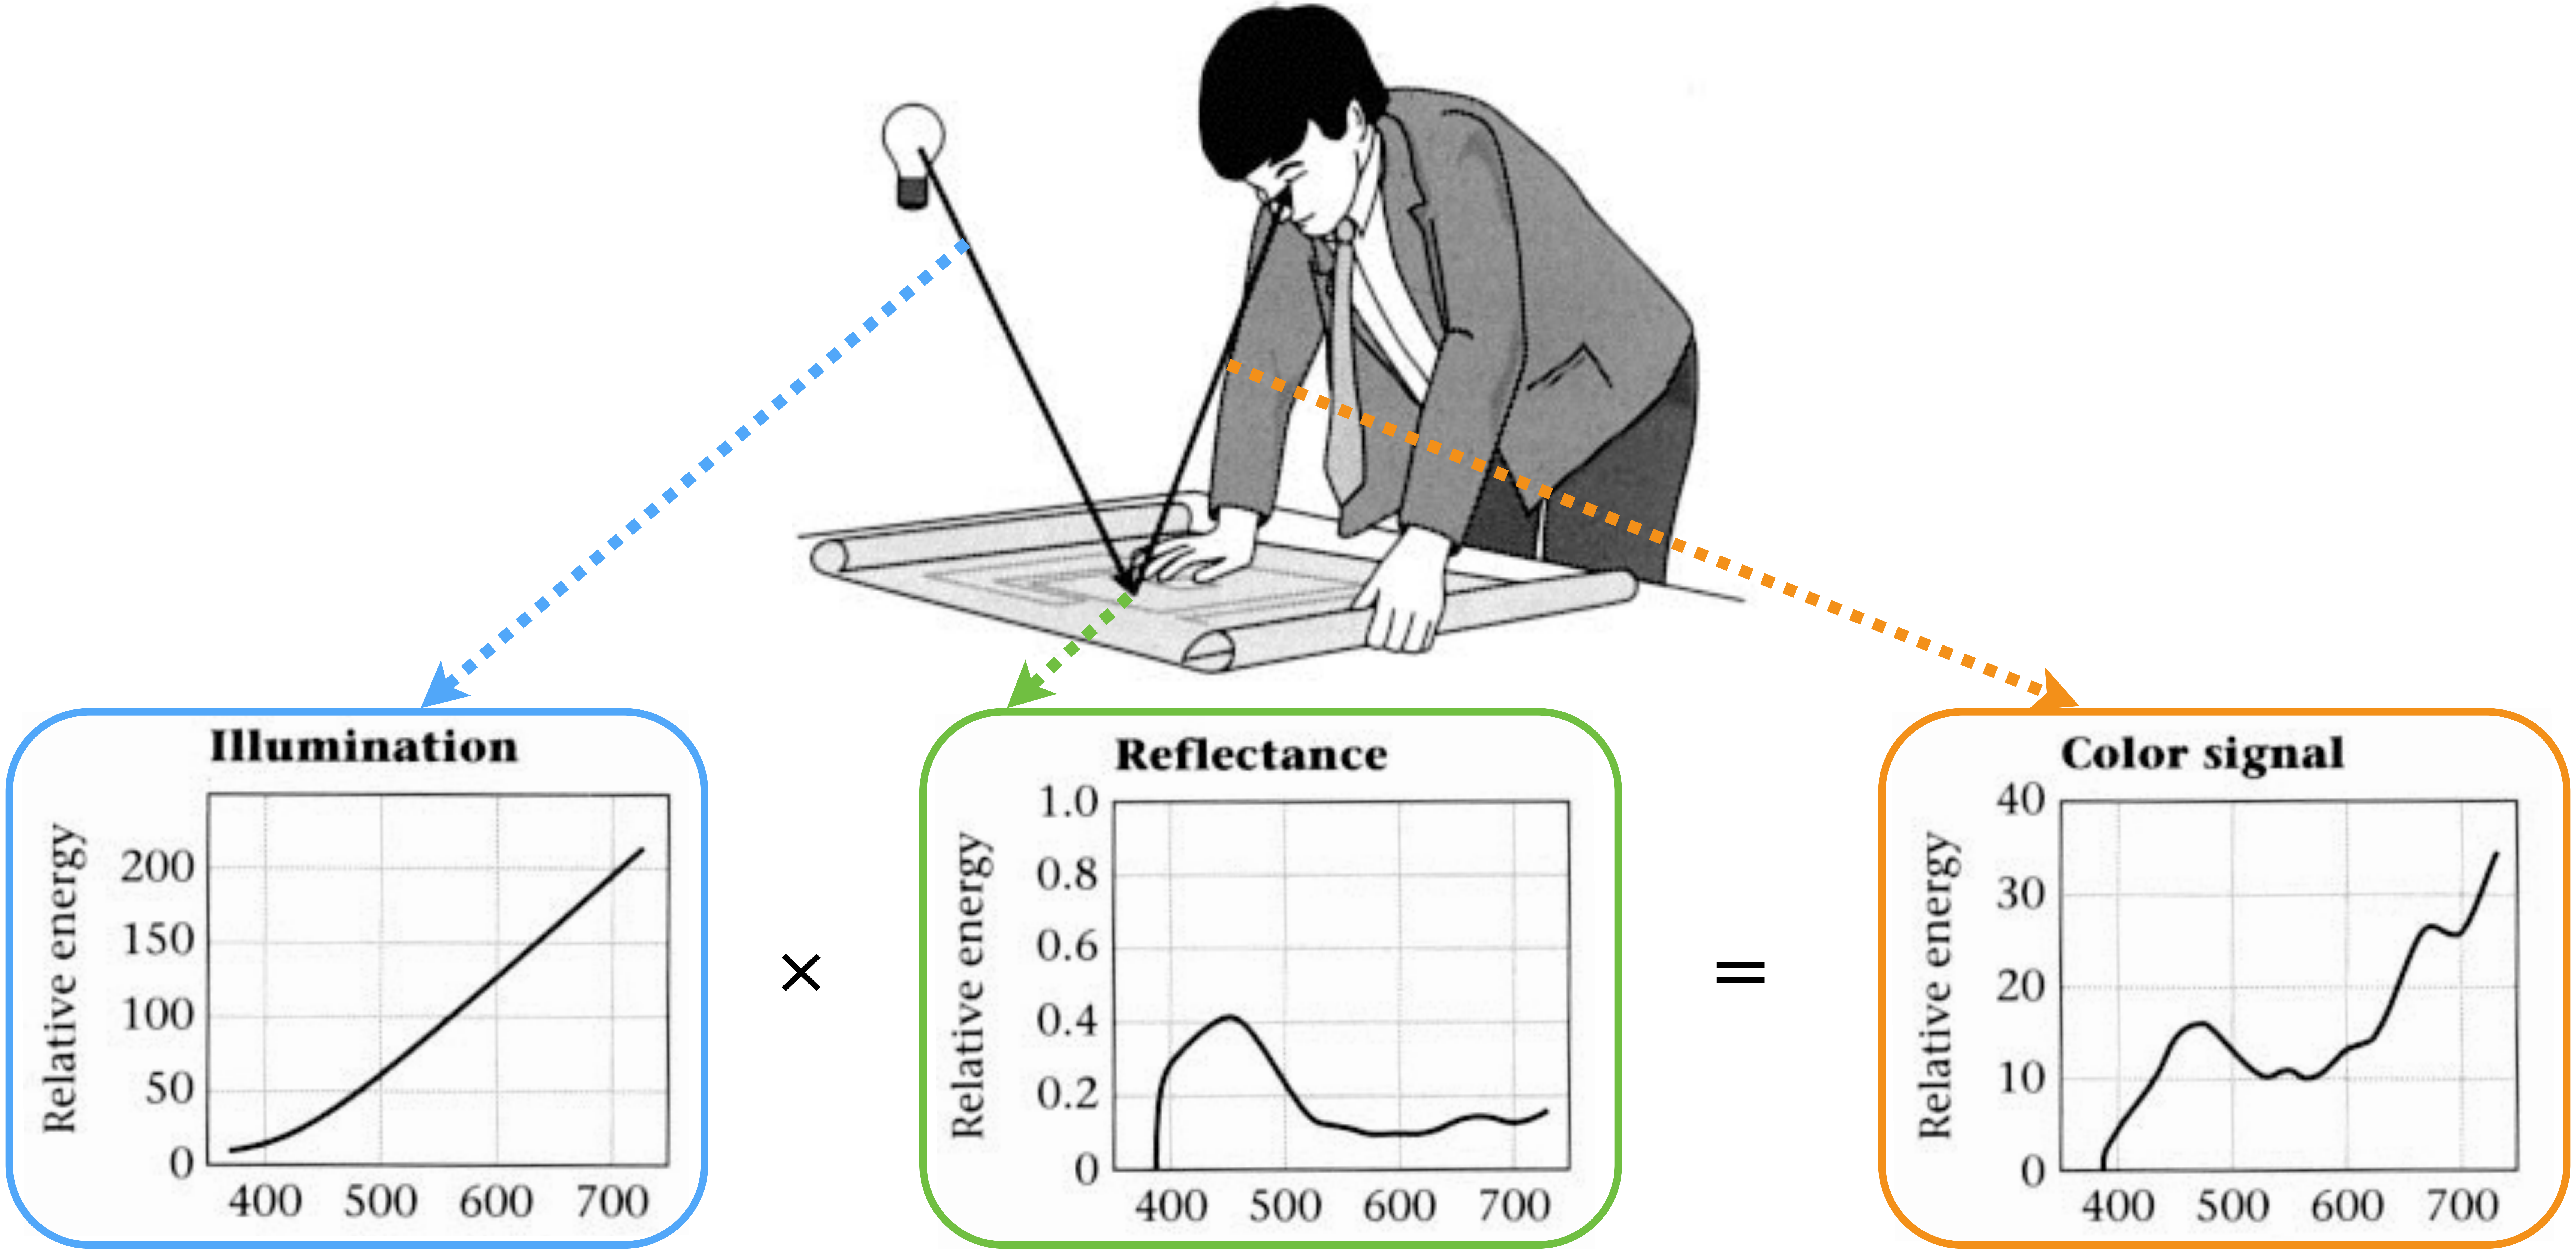

Physical light to perceptual color

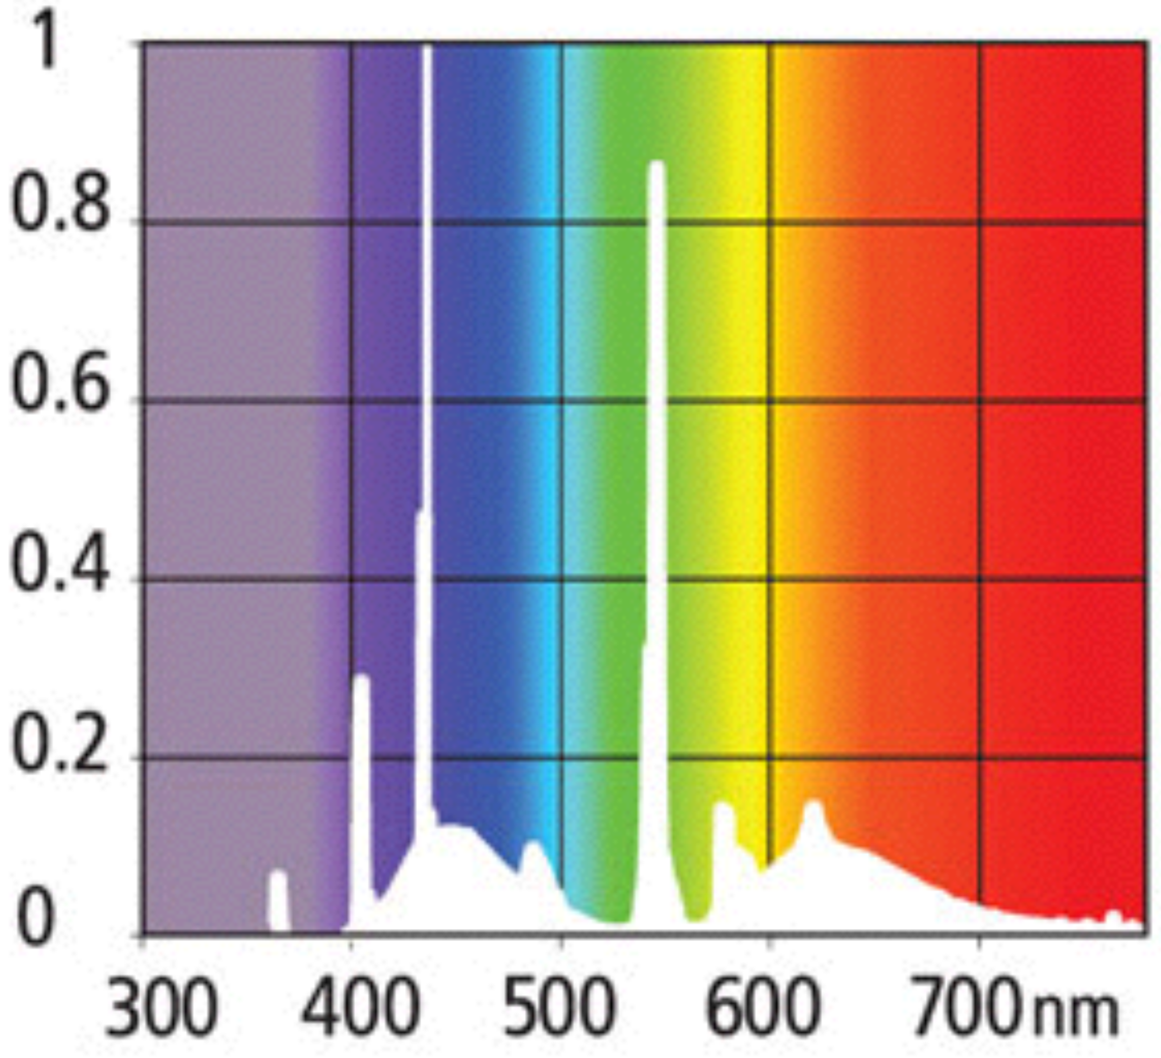

Physical

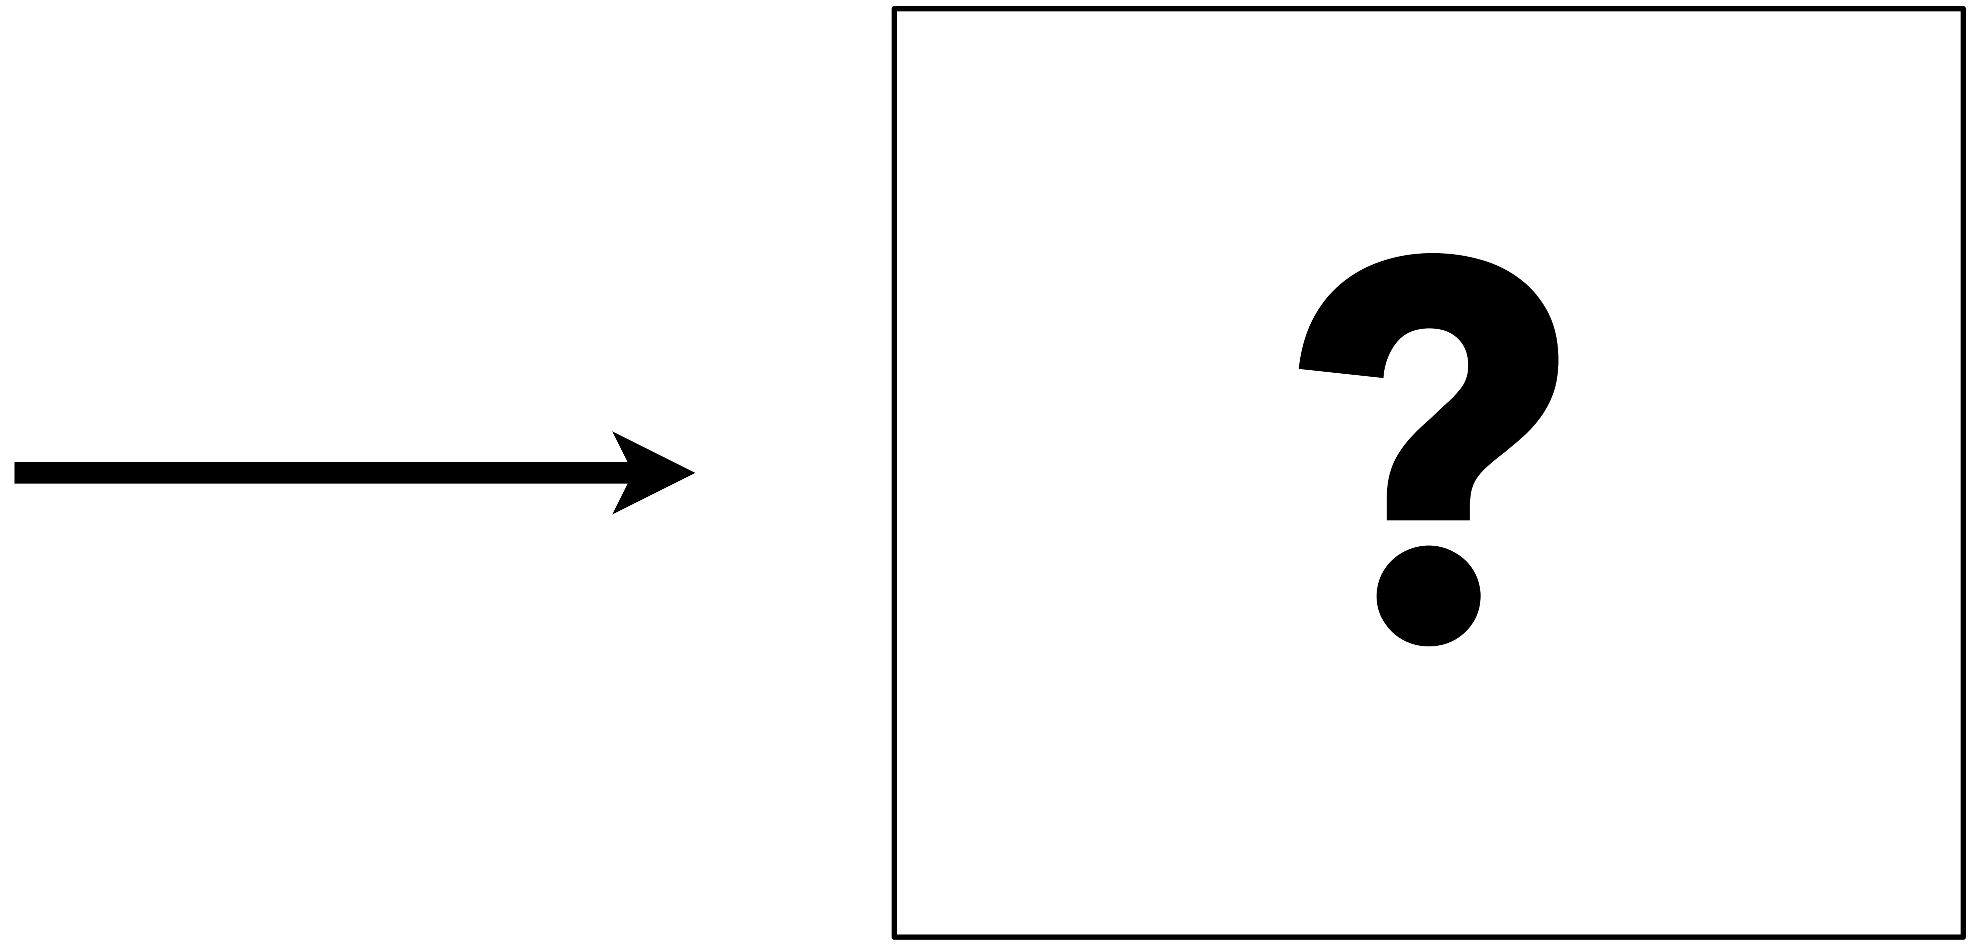

Perceptual

The eye as a measurement device

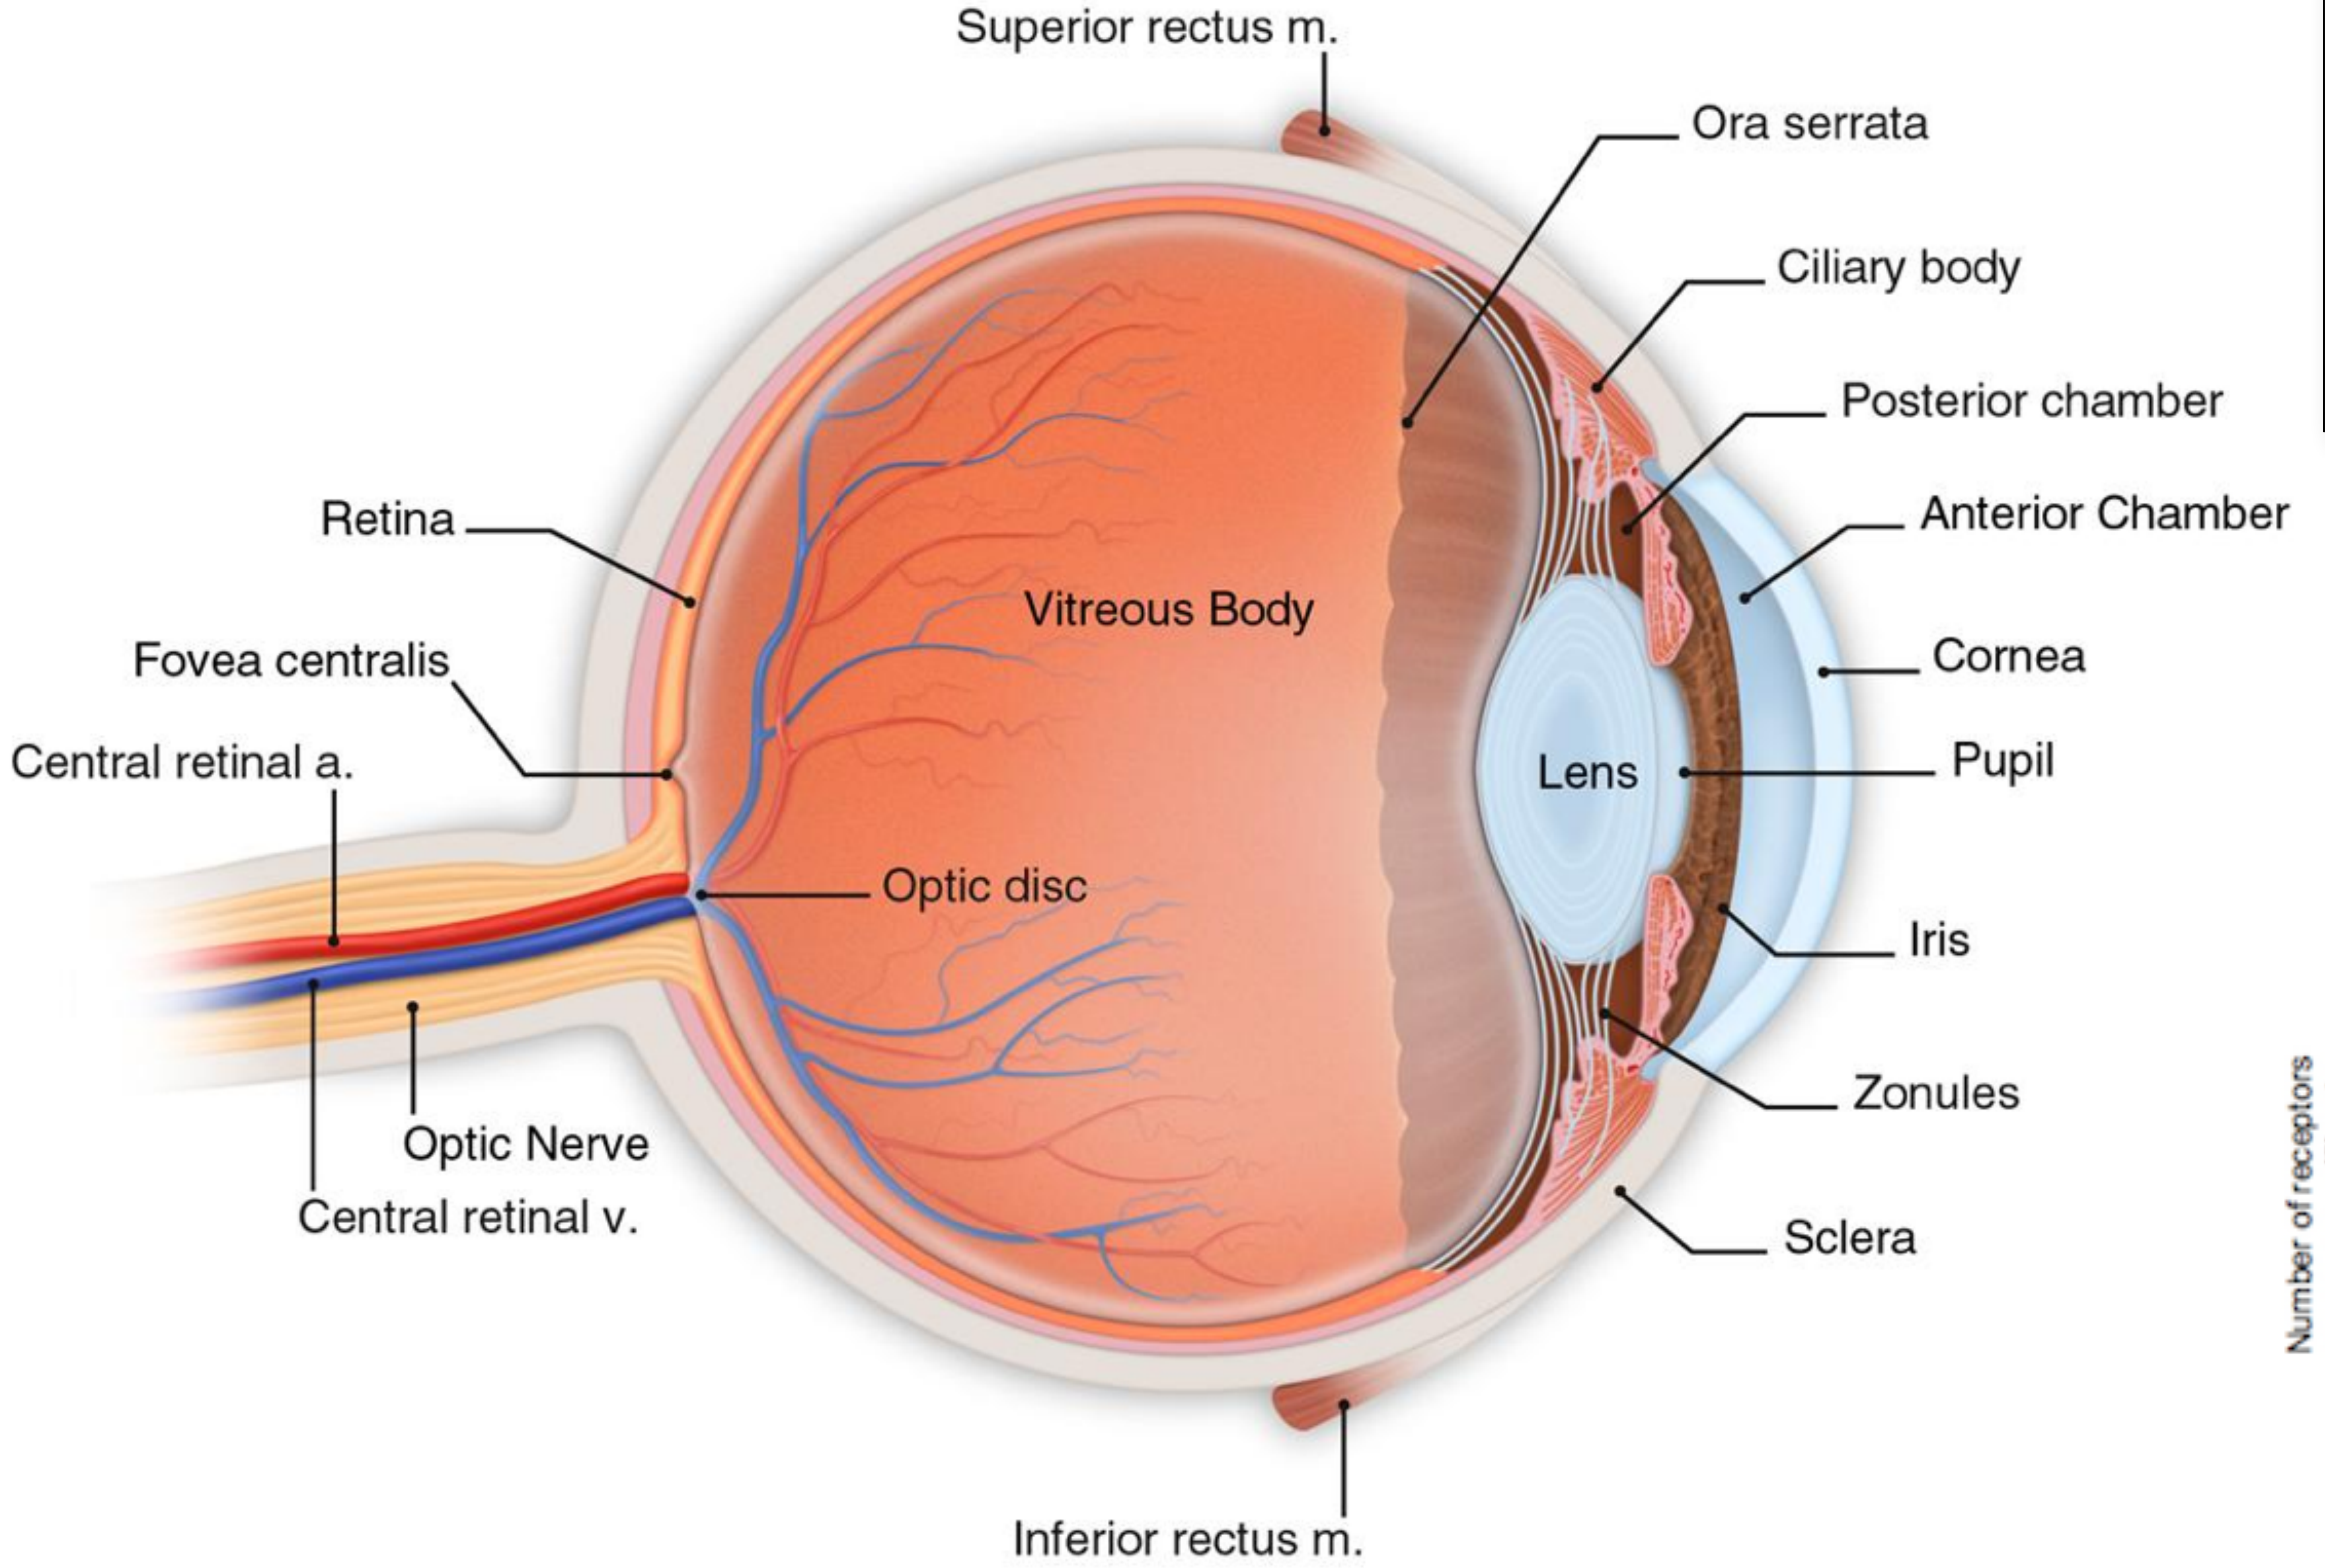

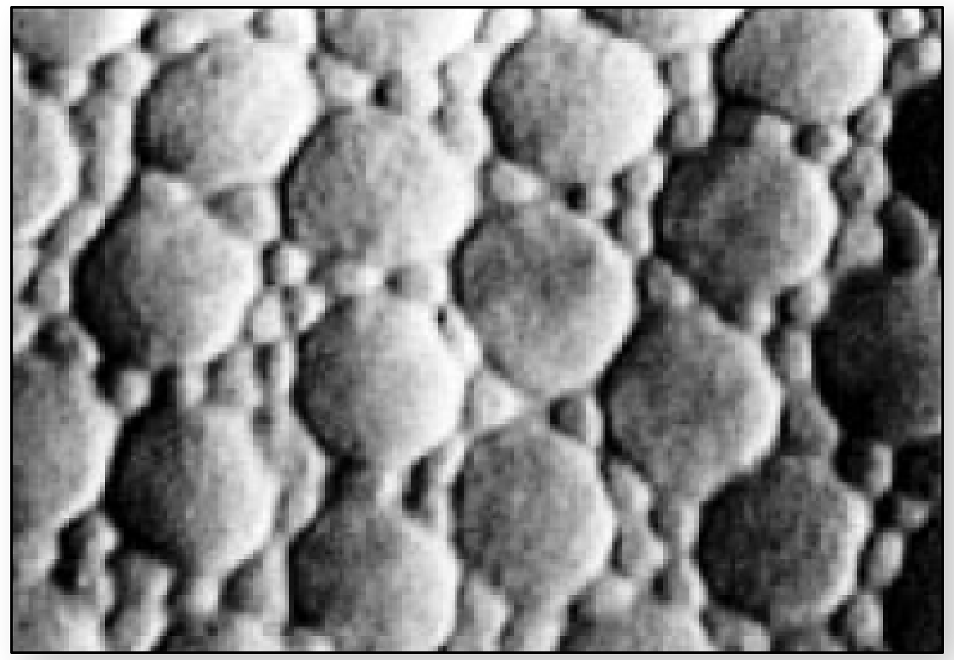

near fovea

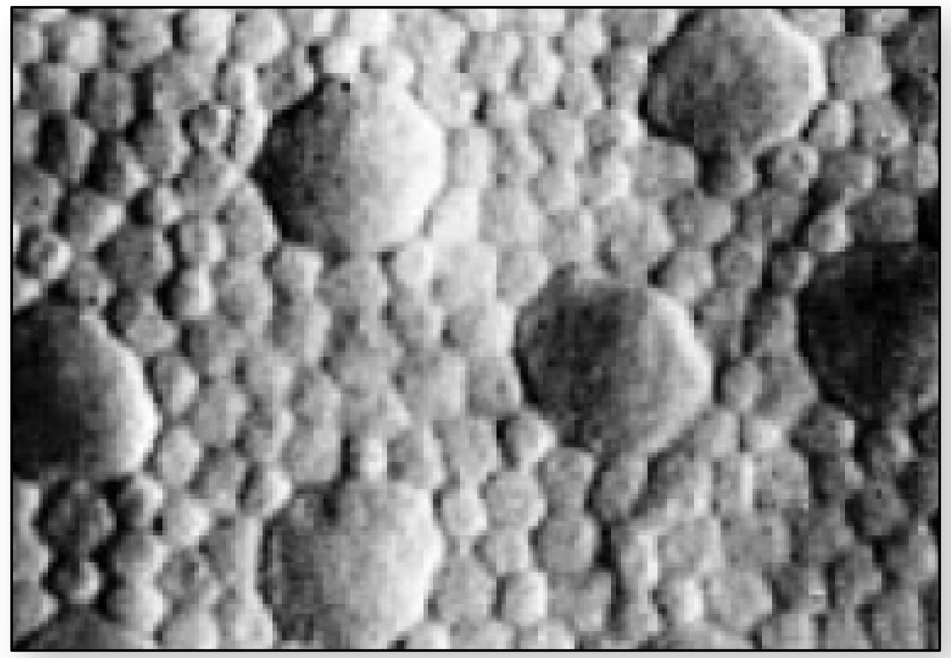

away from fovea

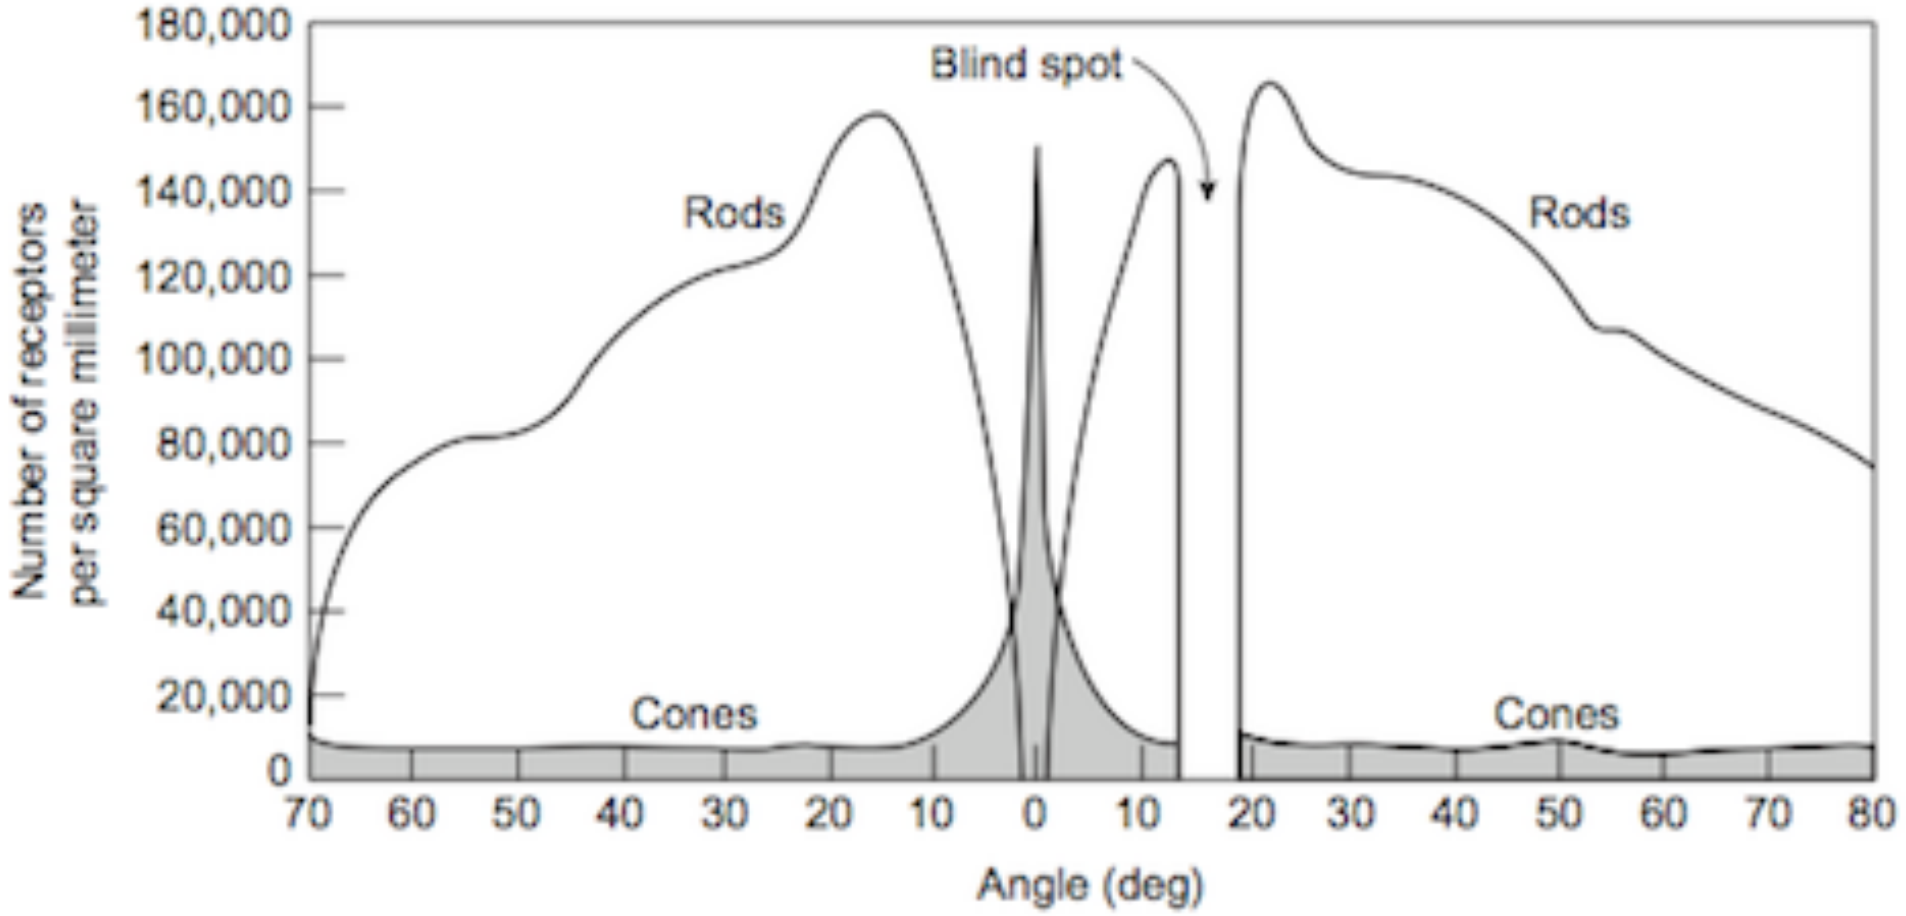

A simple light detector model

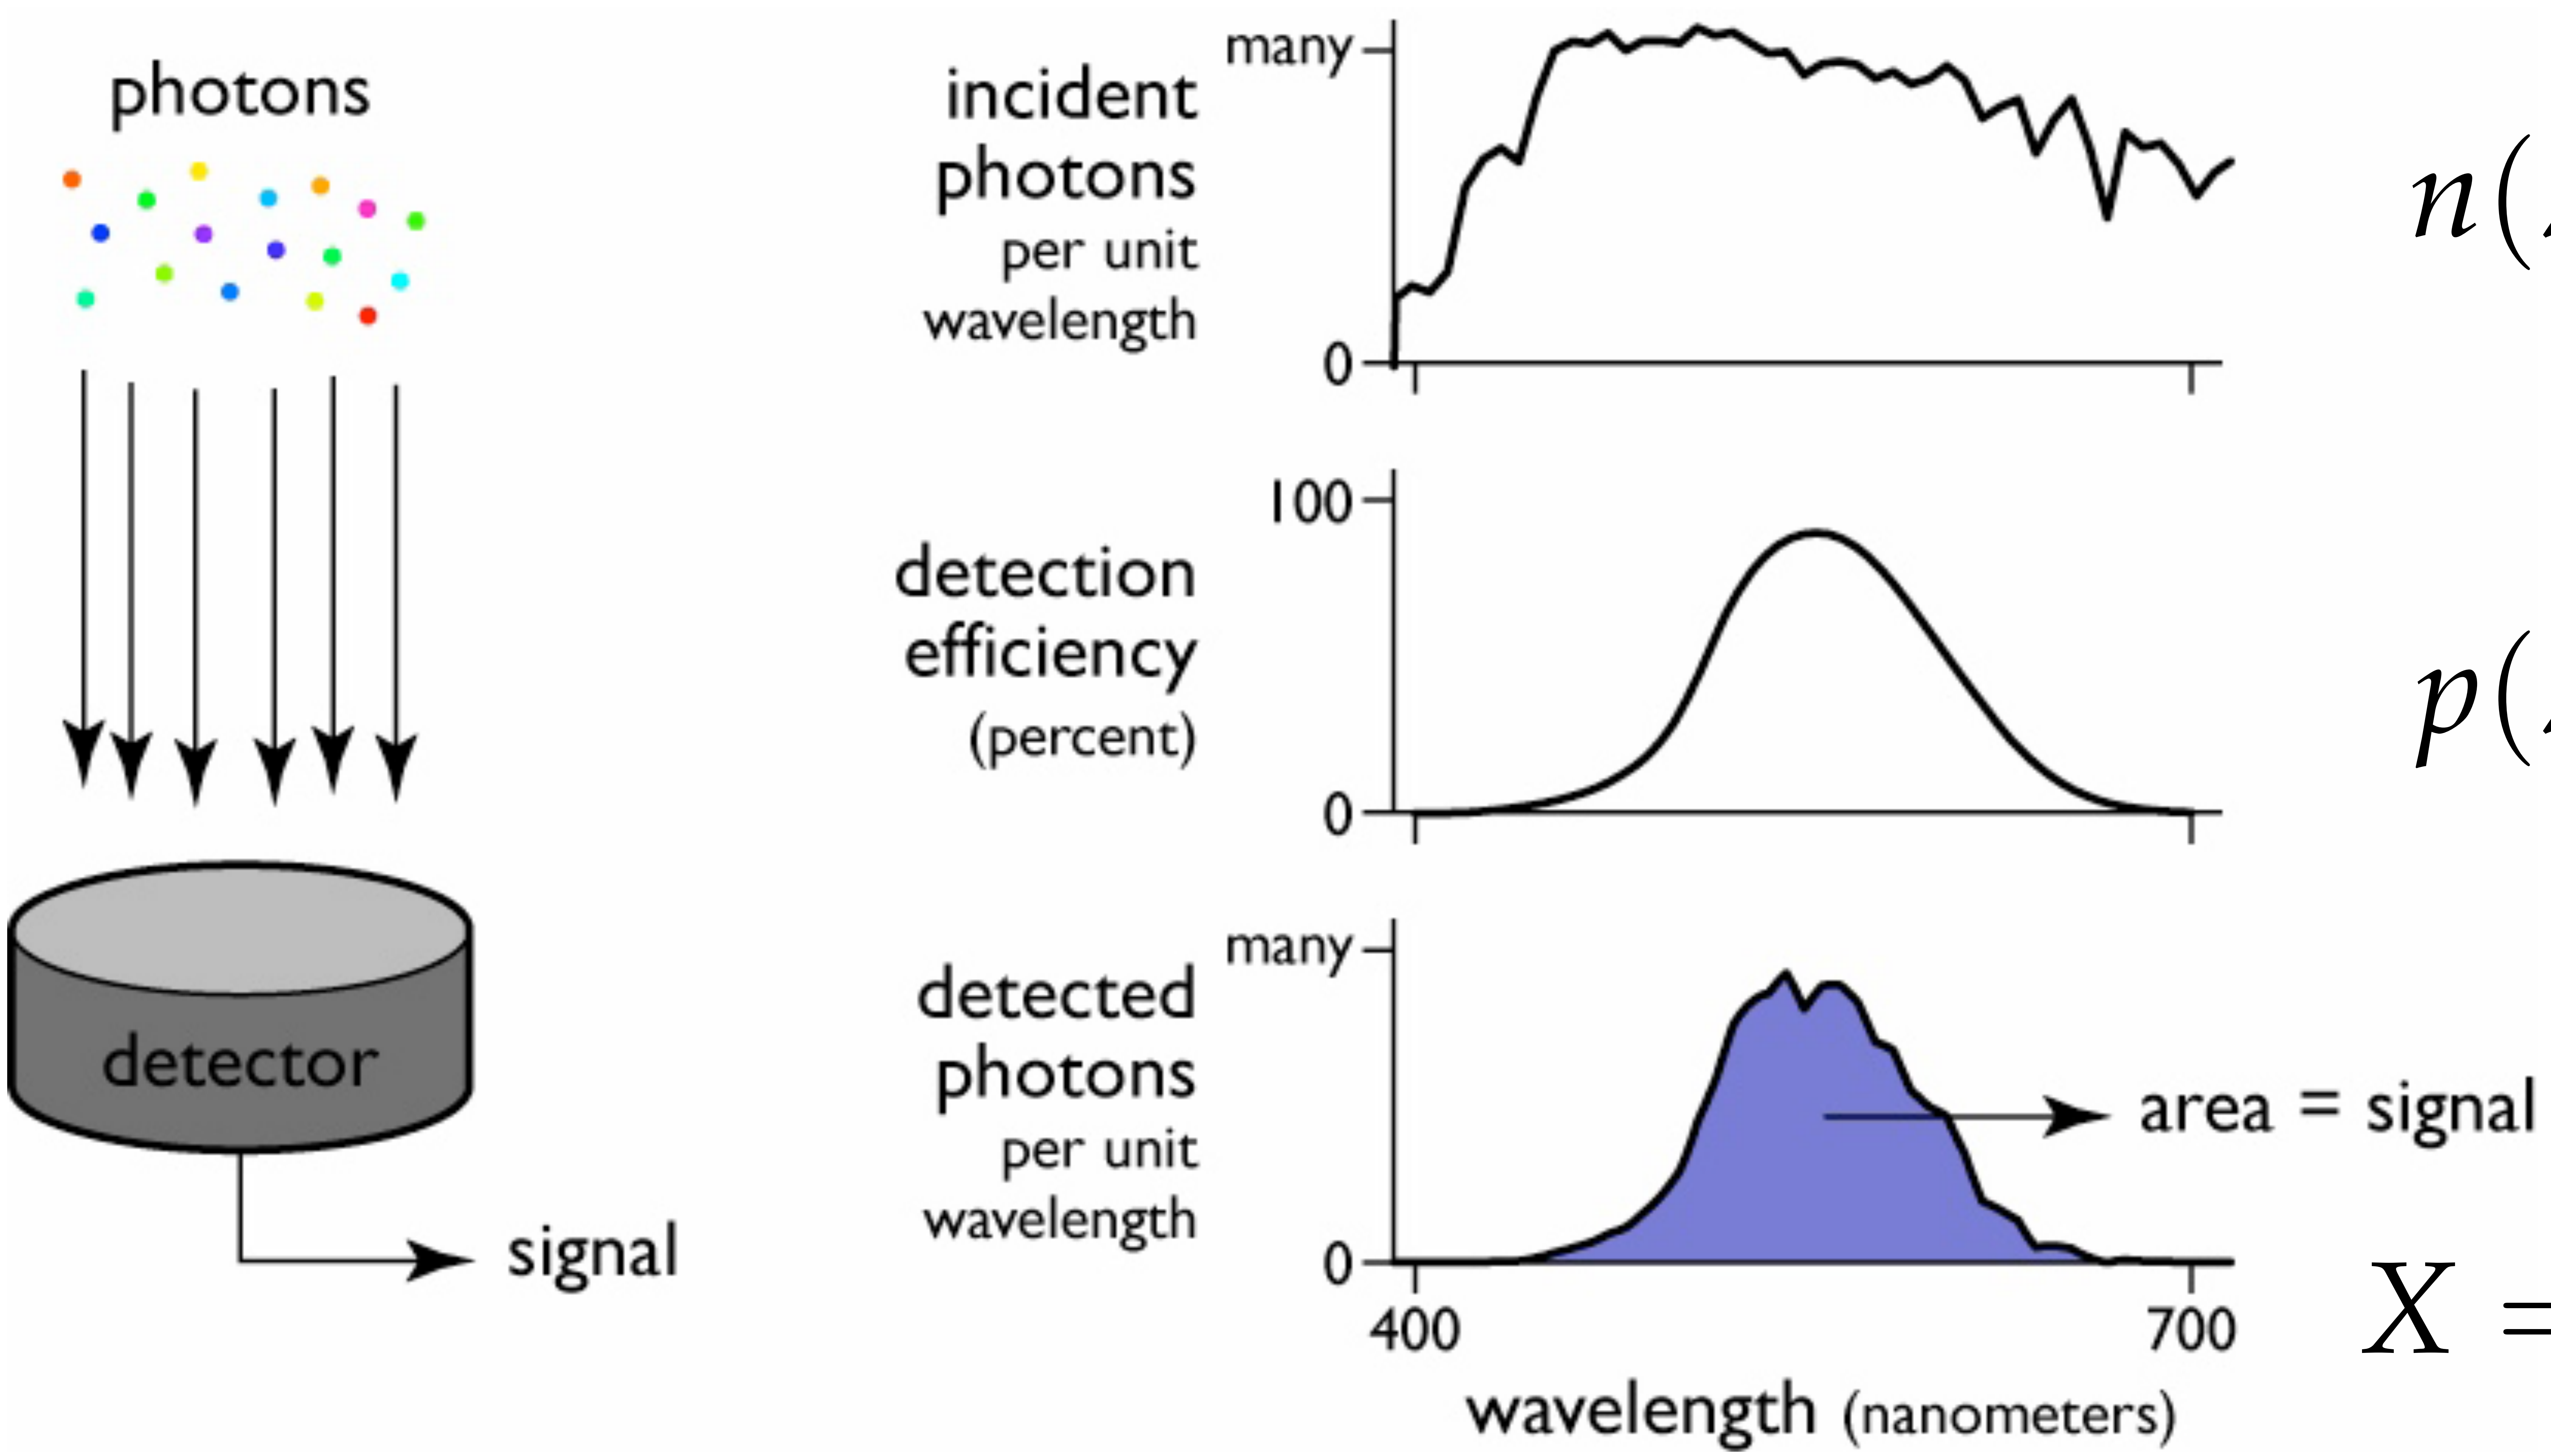

$n(\lambda)$

$p(\lambda)$

$$X = \int n(\lambda)p(\lambda)\,\mathrm{d}\lambda$$

Light detection math

$$X = \int s(\lambda)\, r(\lambda)\, \mathrm{d}\lambda$$

measured signal

input spectrum

detector's sensitivity

$$S = \int r_S(\lambda)\, s(\lambda)\, \mathrm{d}\lambda$$

$$M = \int r_M(\lambda)\, s(\lambda)\, \mathrm{d}\lambda$$

$$L = \int r_L(\lambda)\, s(\lambda)\, \mathrm{d}\lambda$$

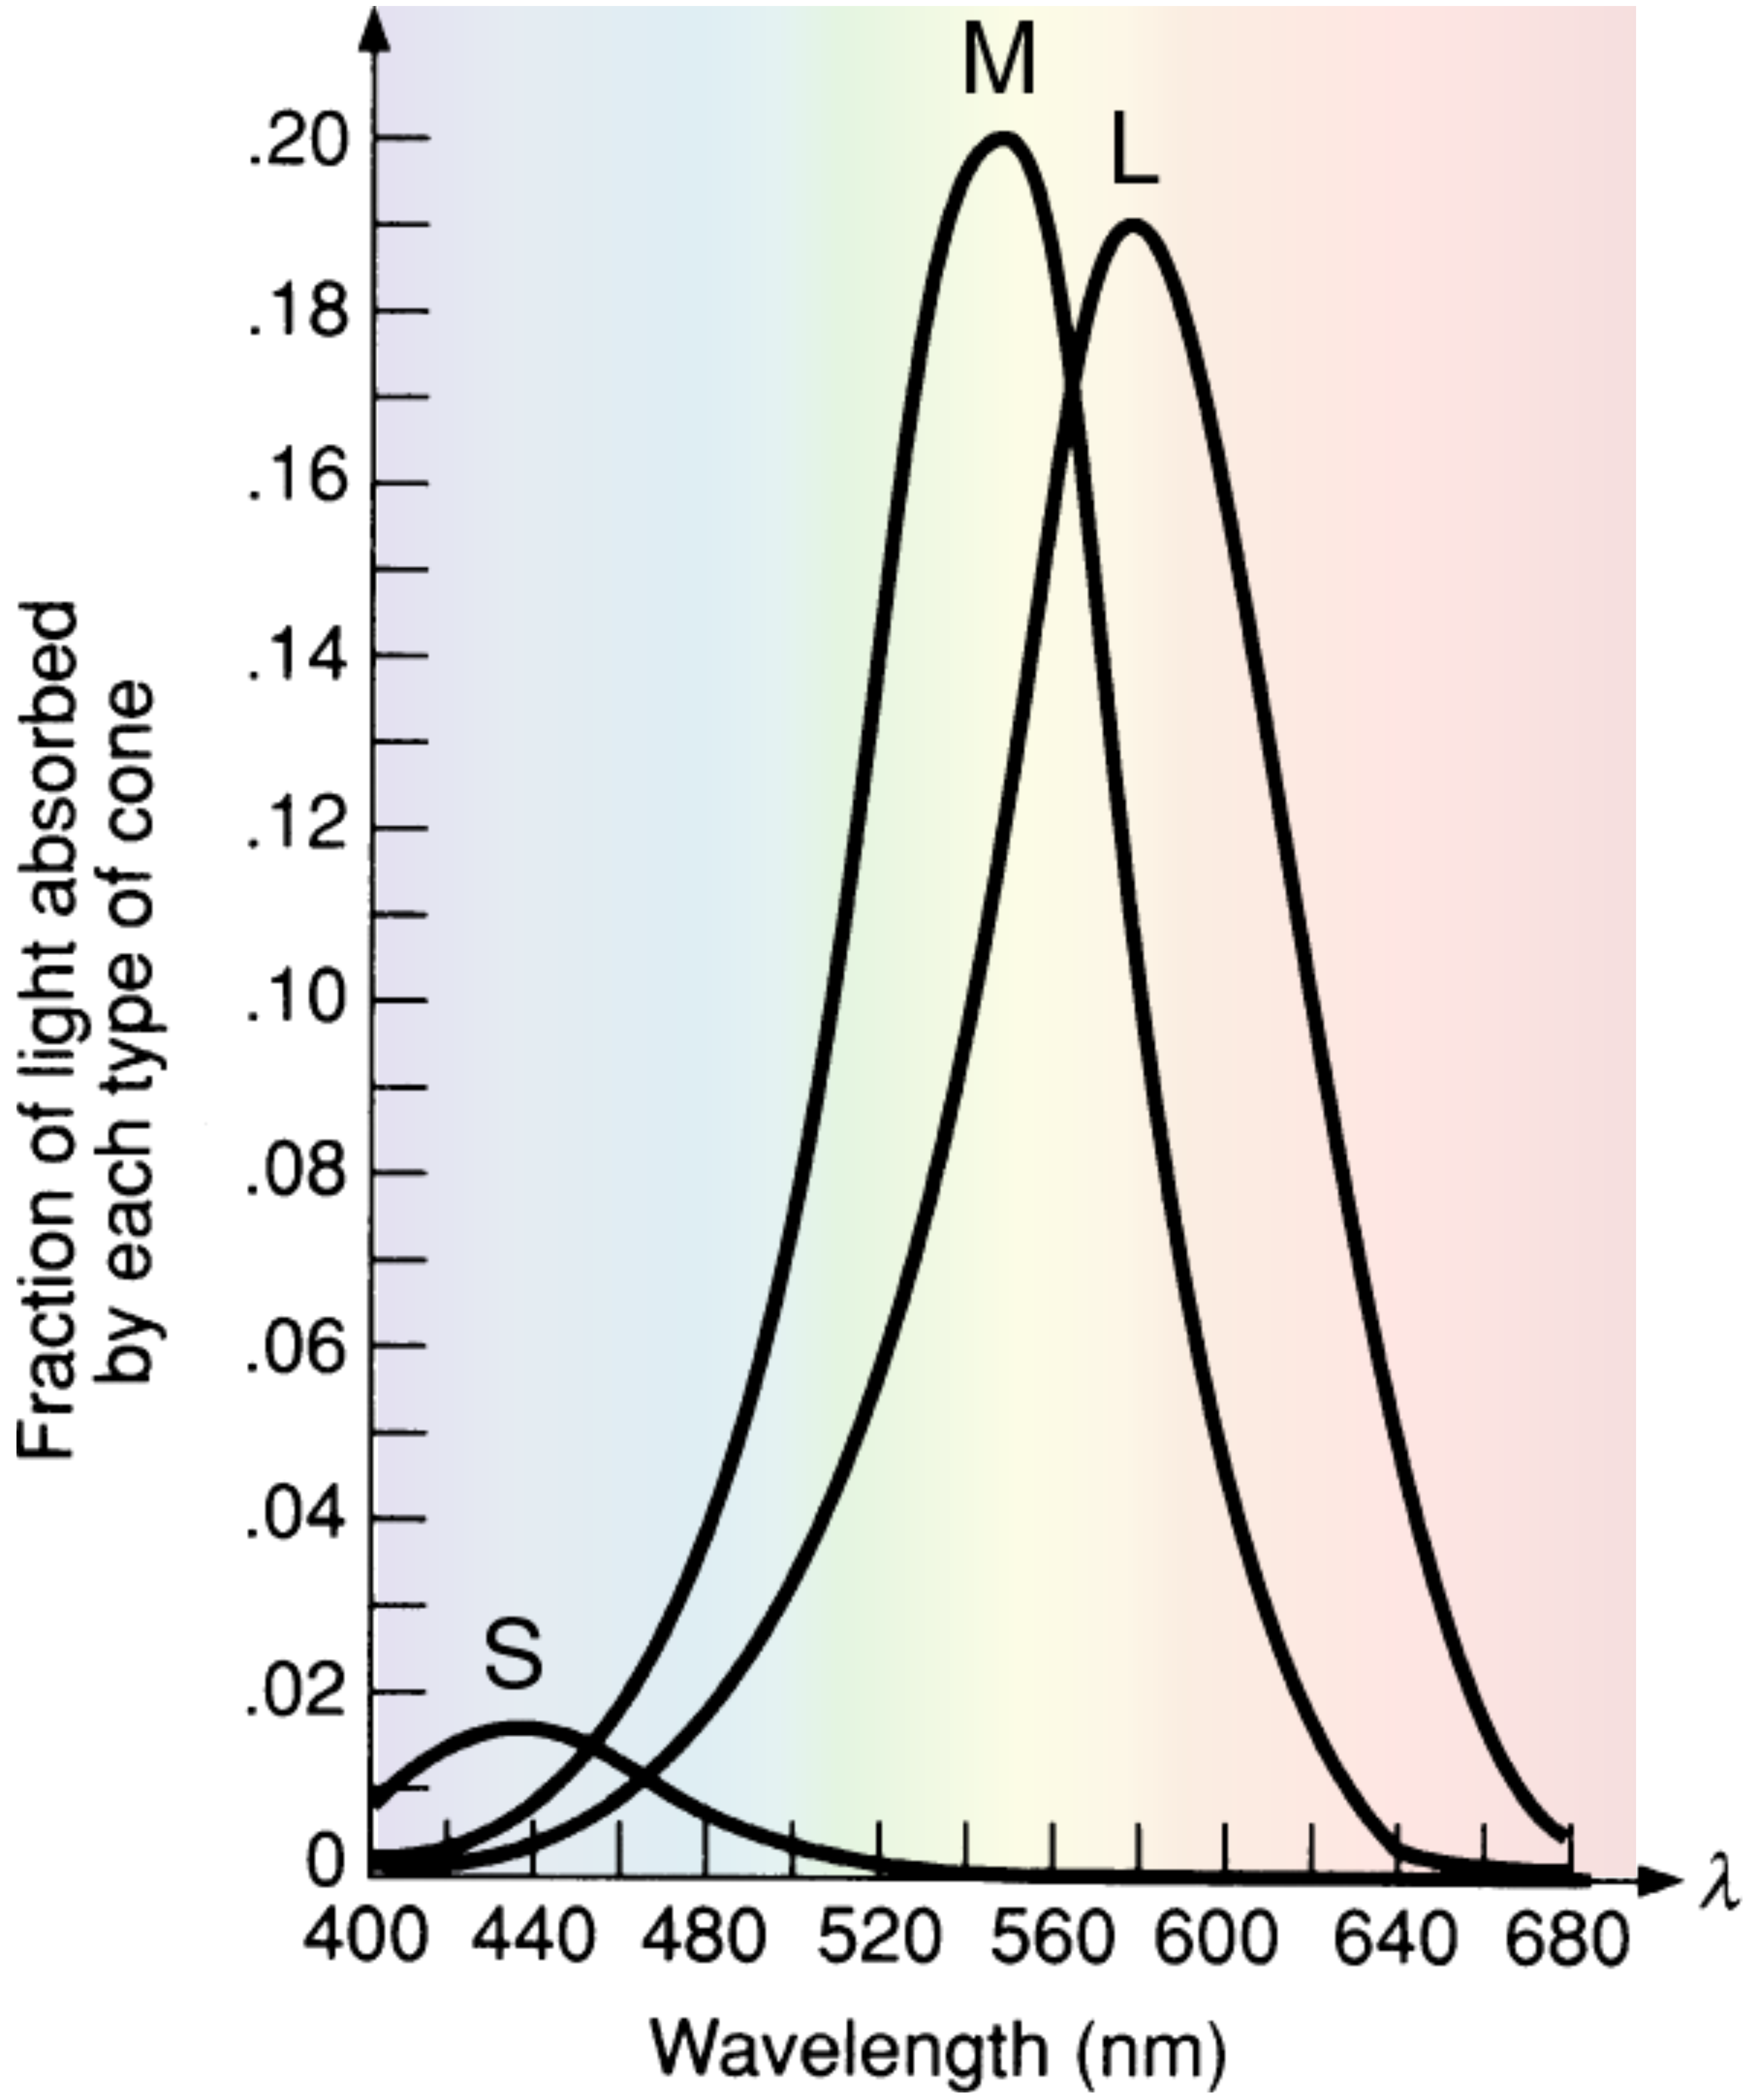

Stimulus
(arbitrary spectrum)

Start with infinite number of values (one per wavelength)

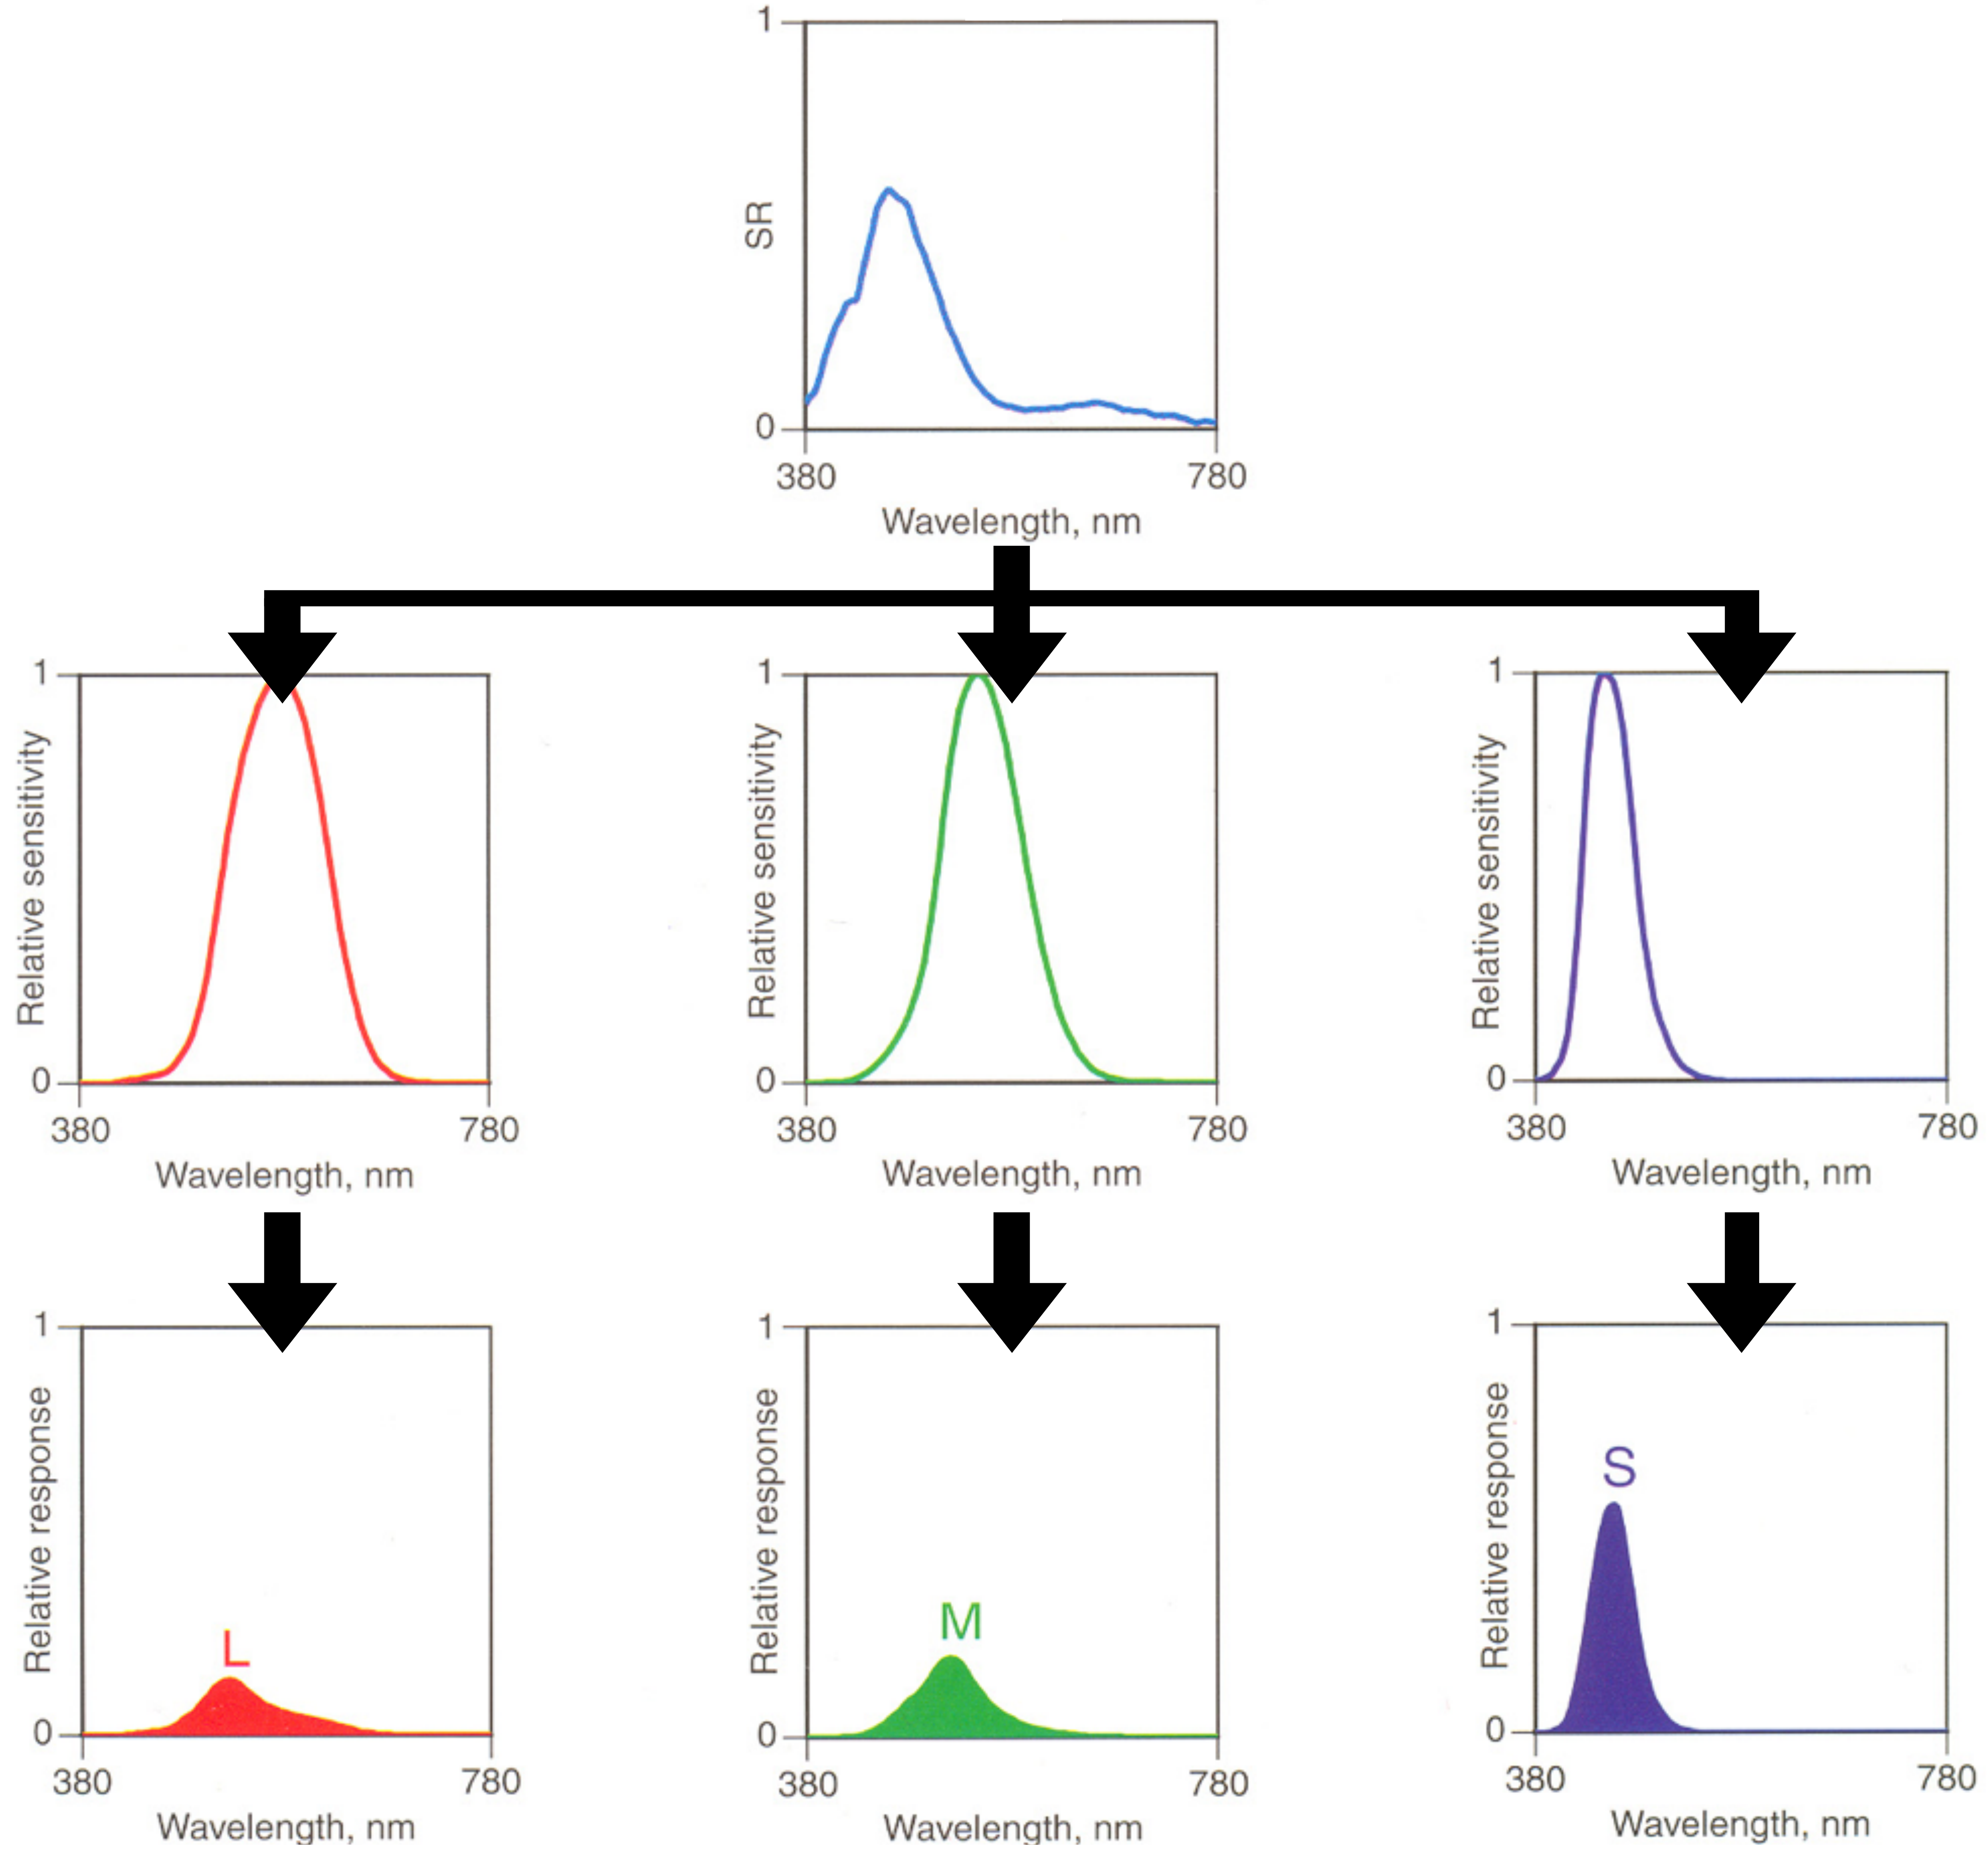

Response curves

Multiply

Integrate

1 number

1 number

1 number

End up with 3 values (one per cone type)

Linear algebra interpretation

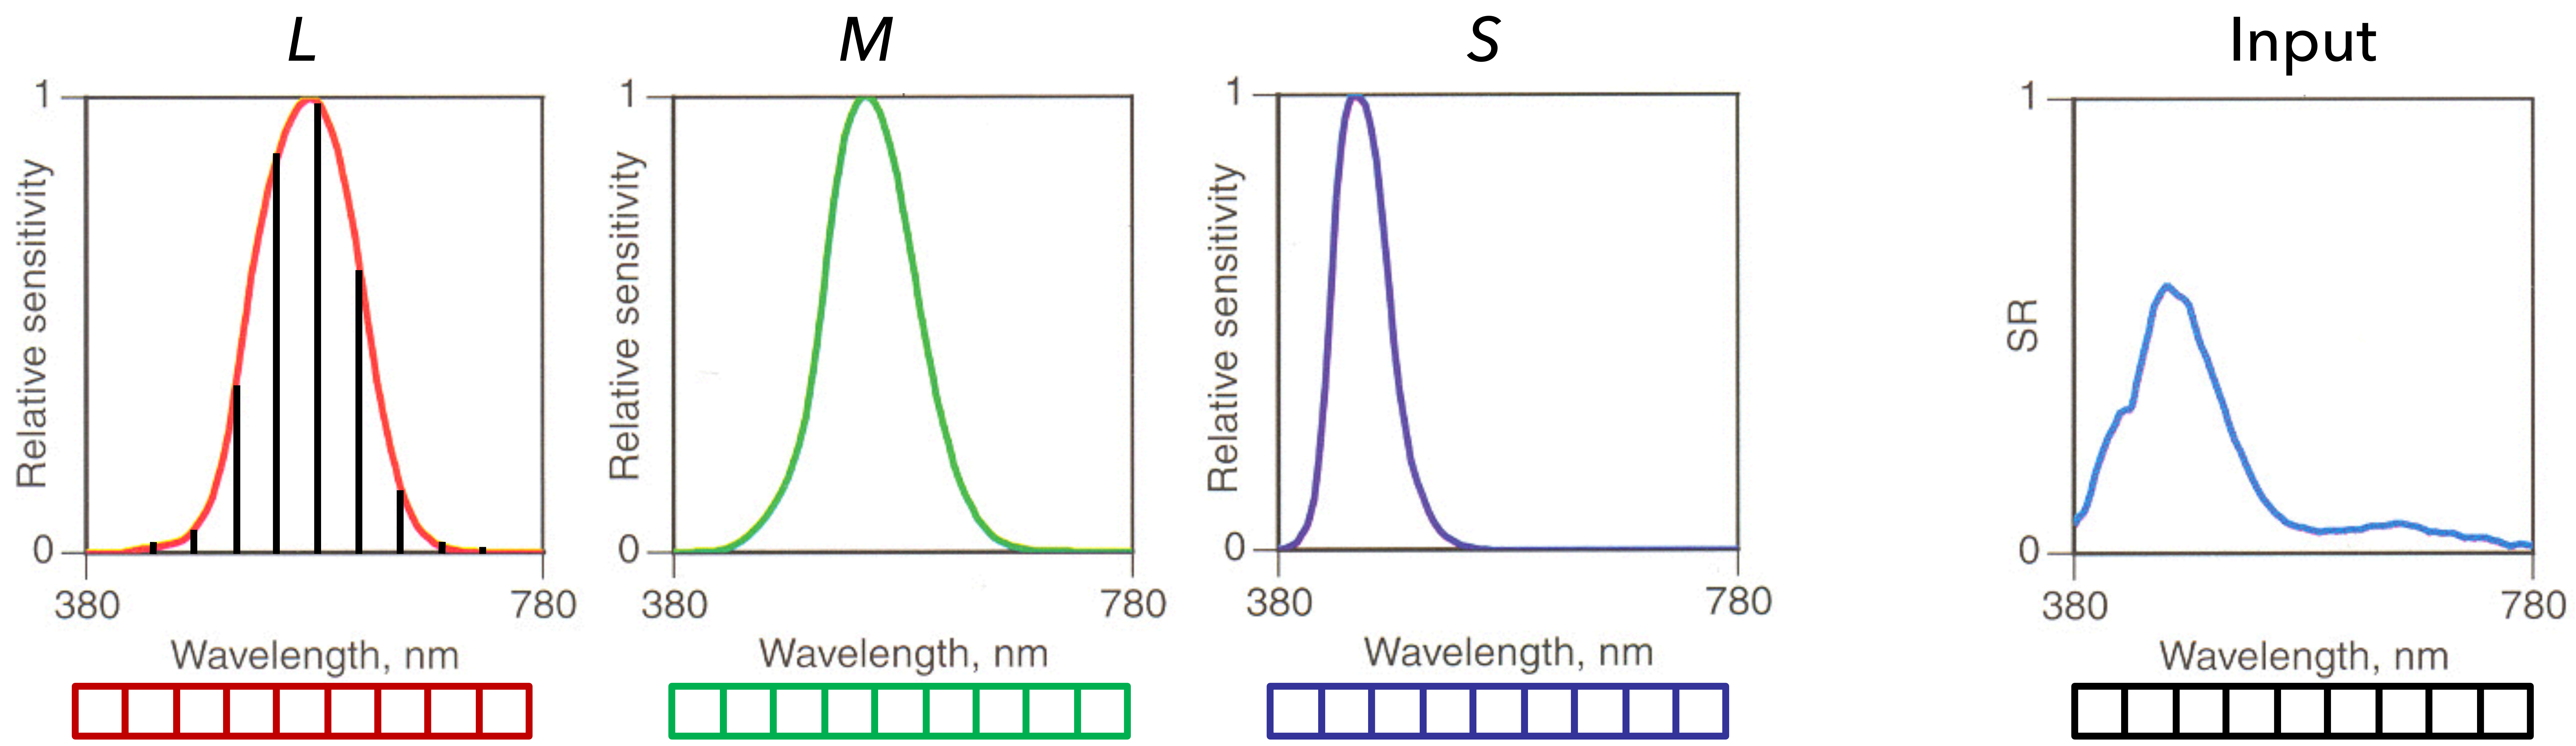

Linear algebra interpretation

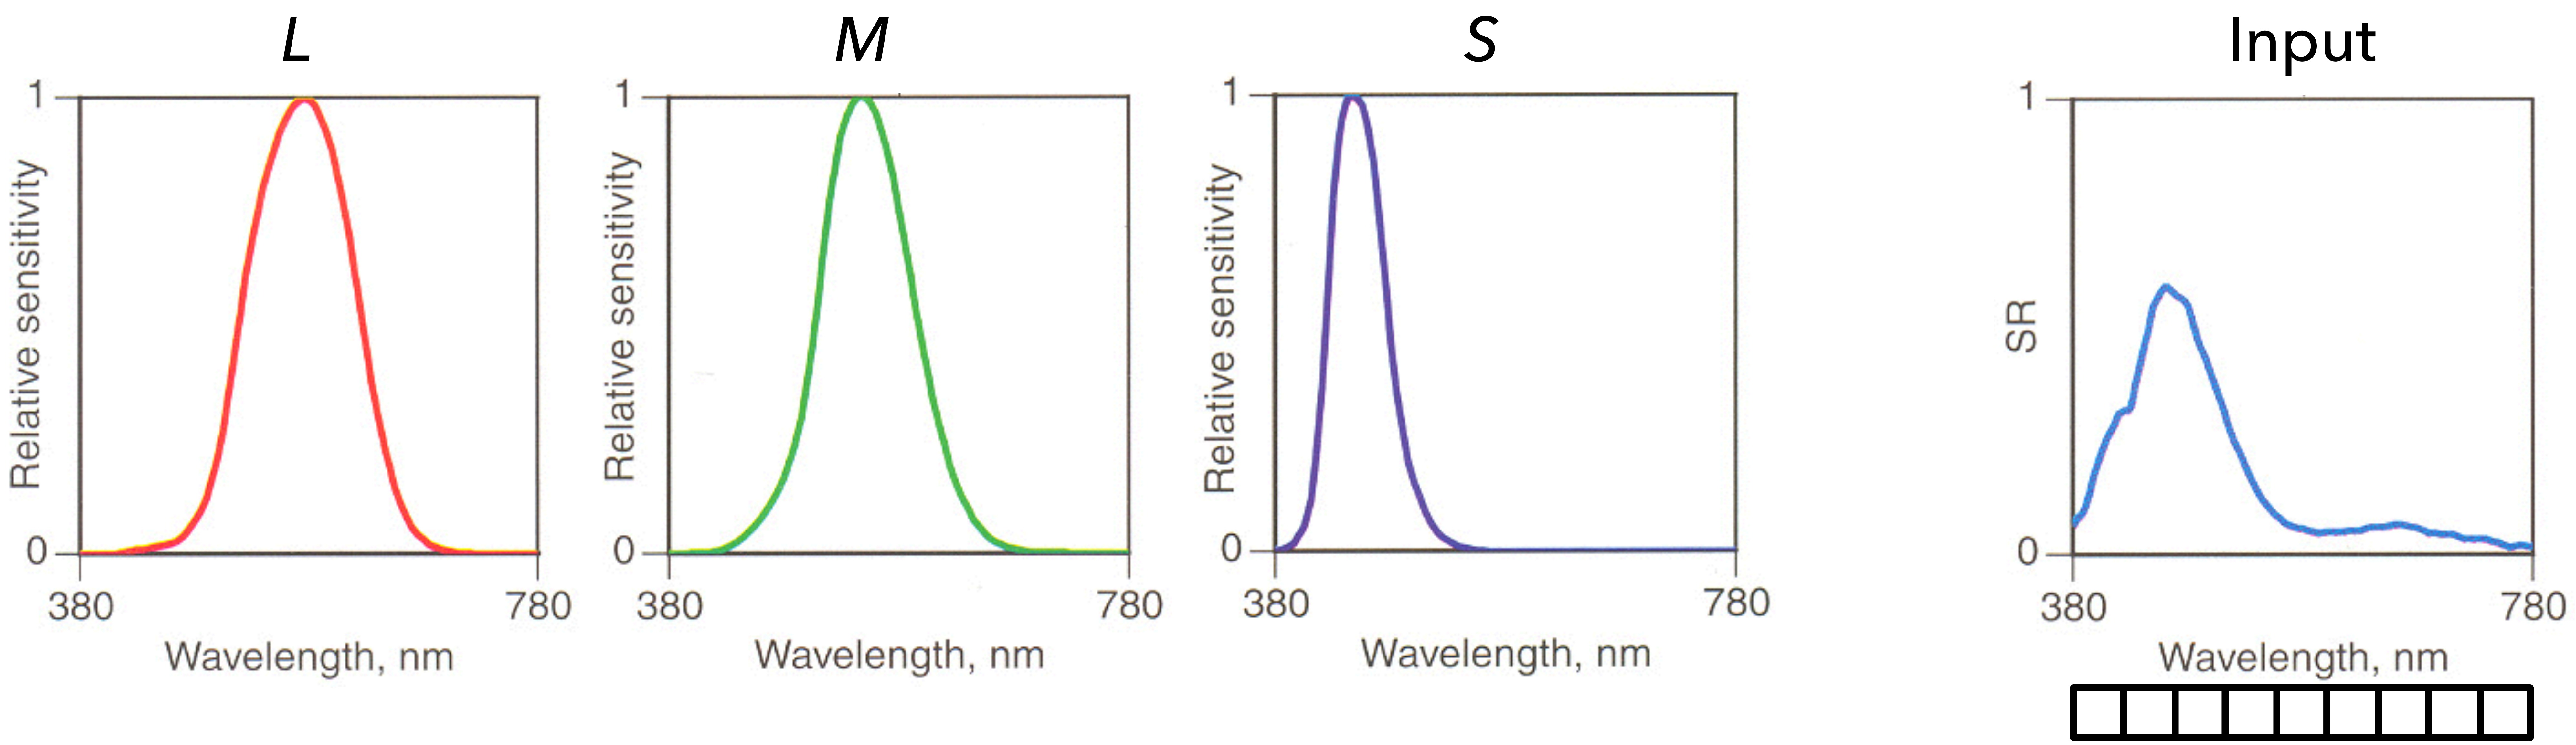

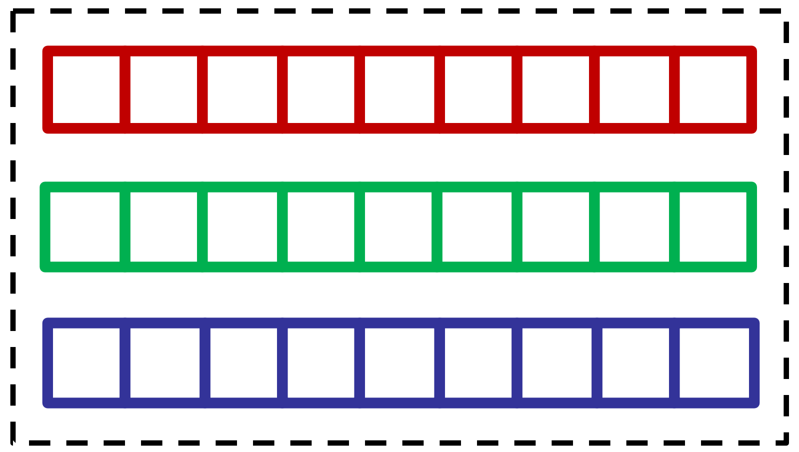

cone sensitivities

Linear algebra interpretation

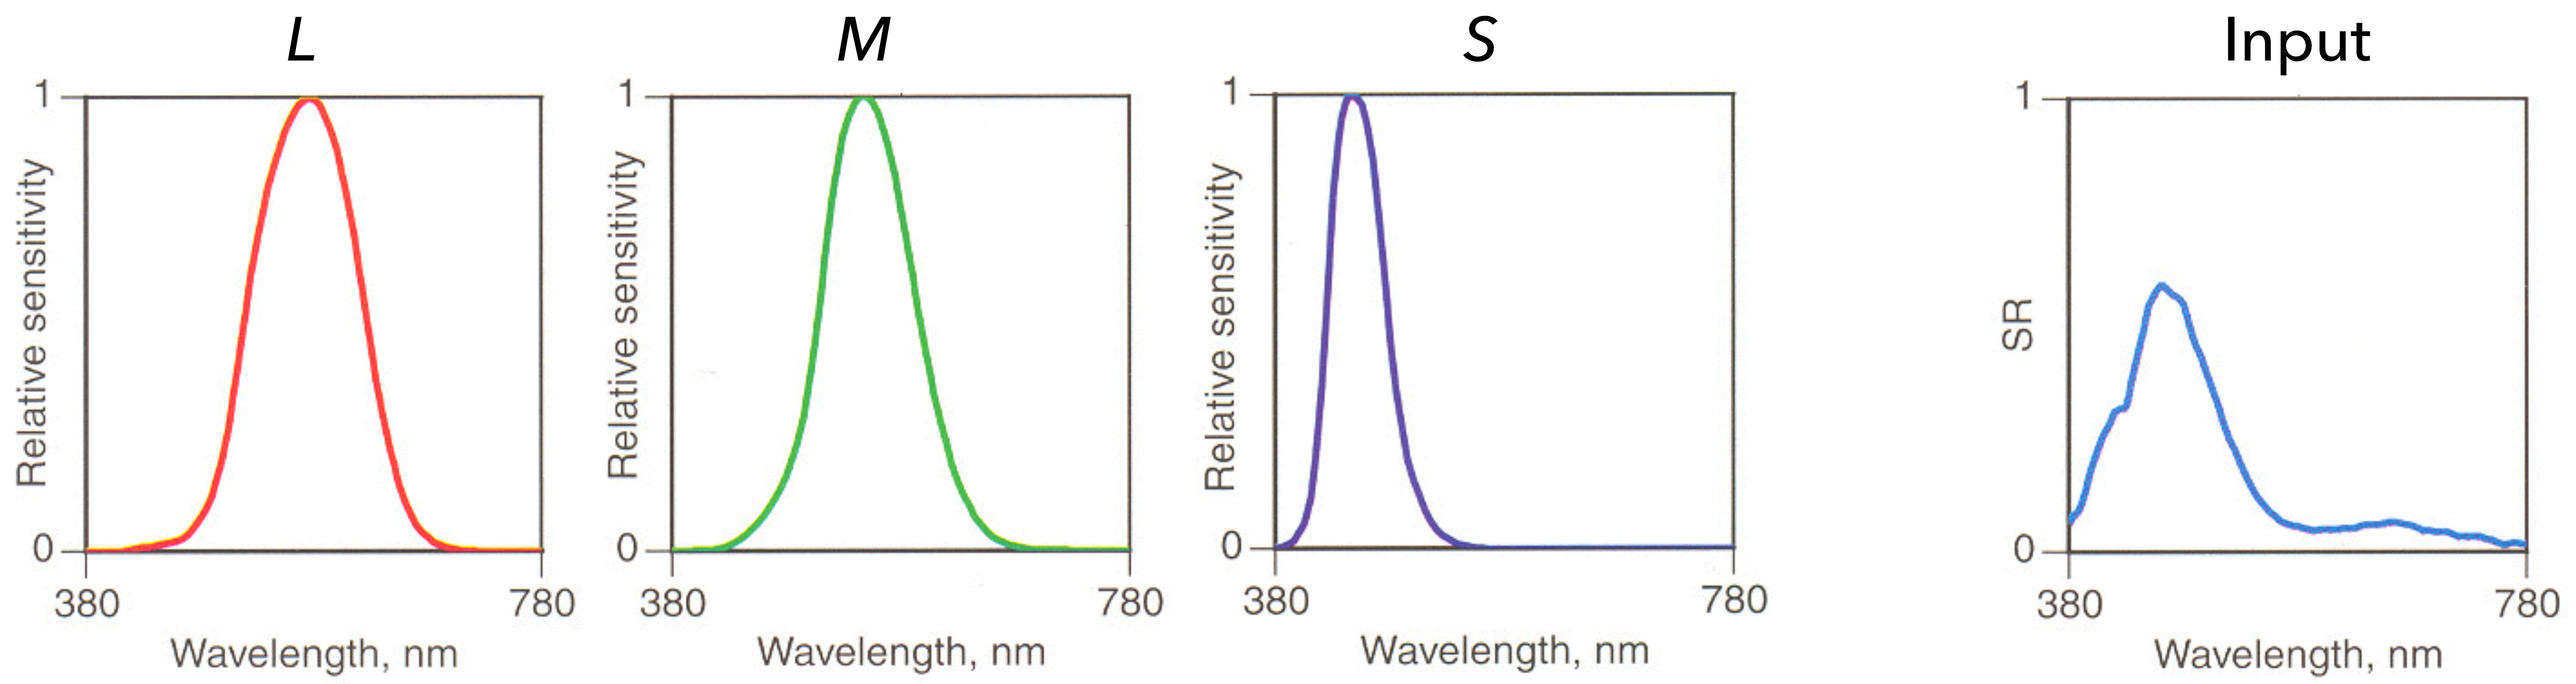

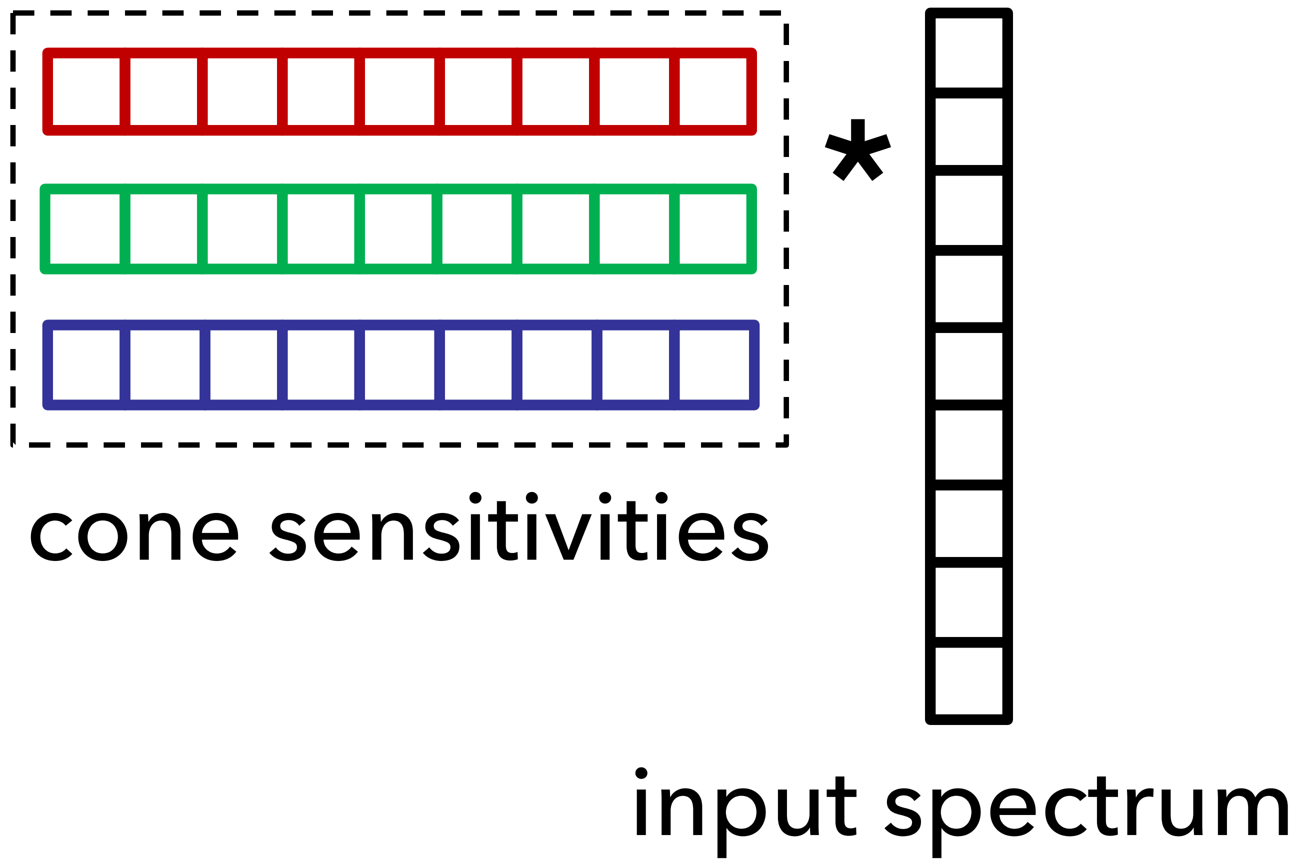

Linear algebra interpretation

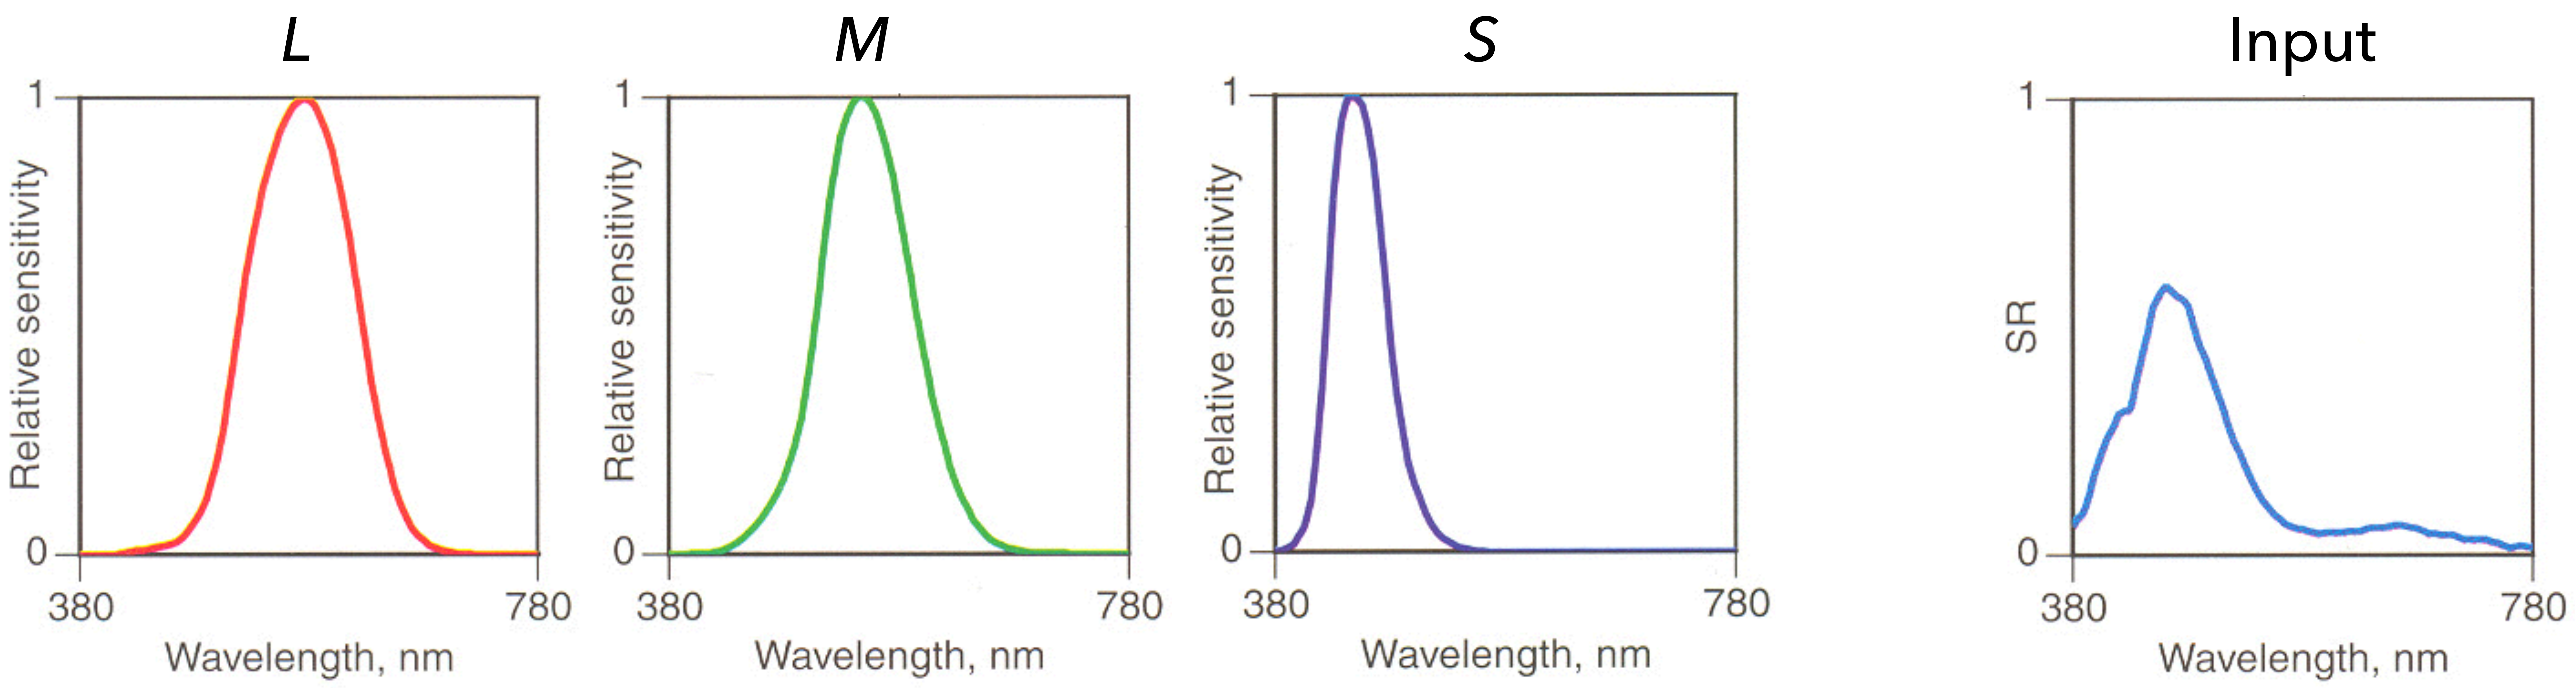

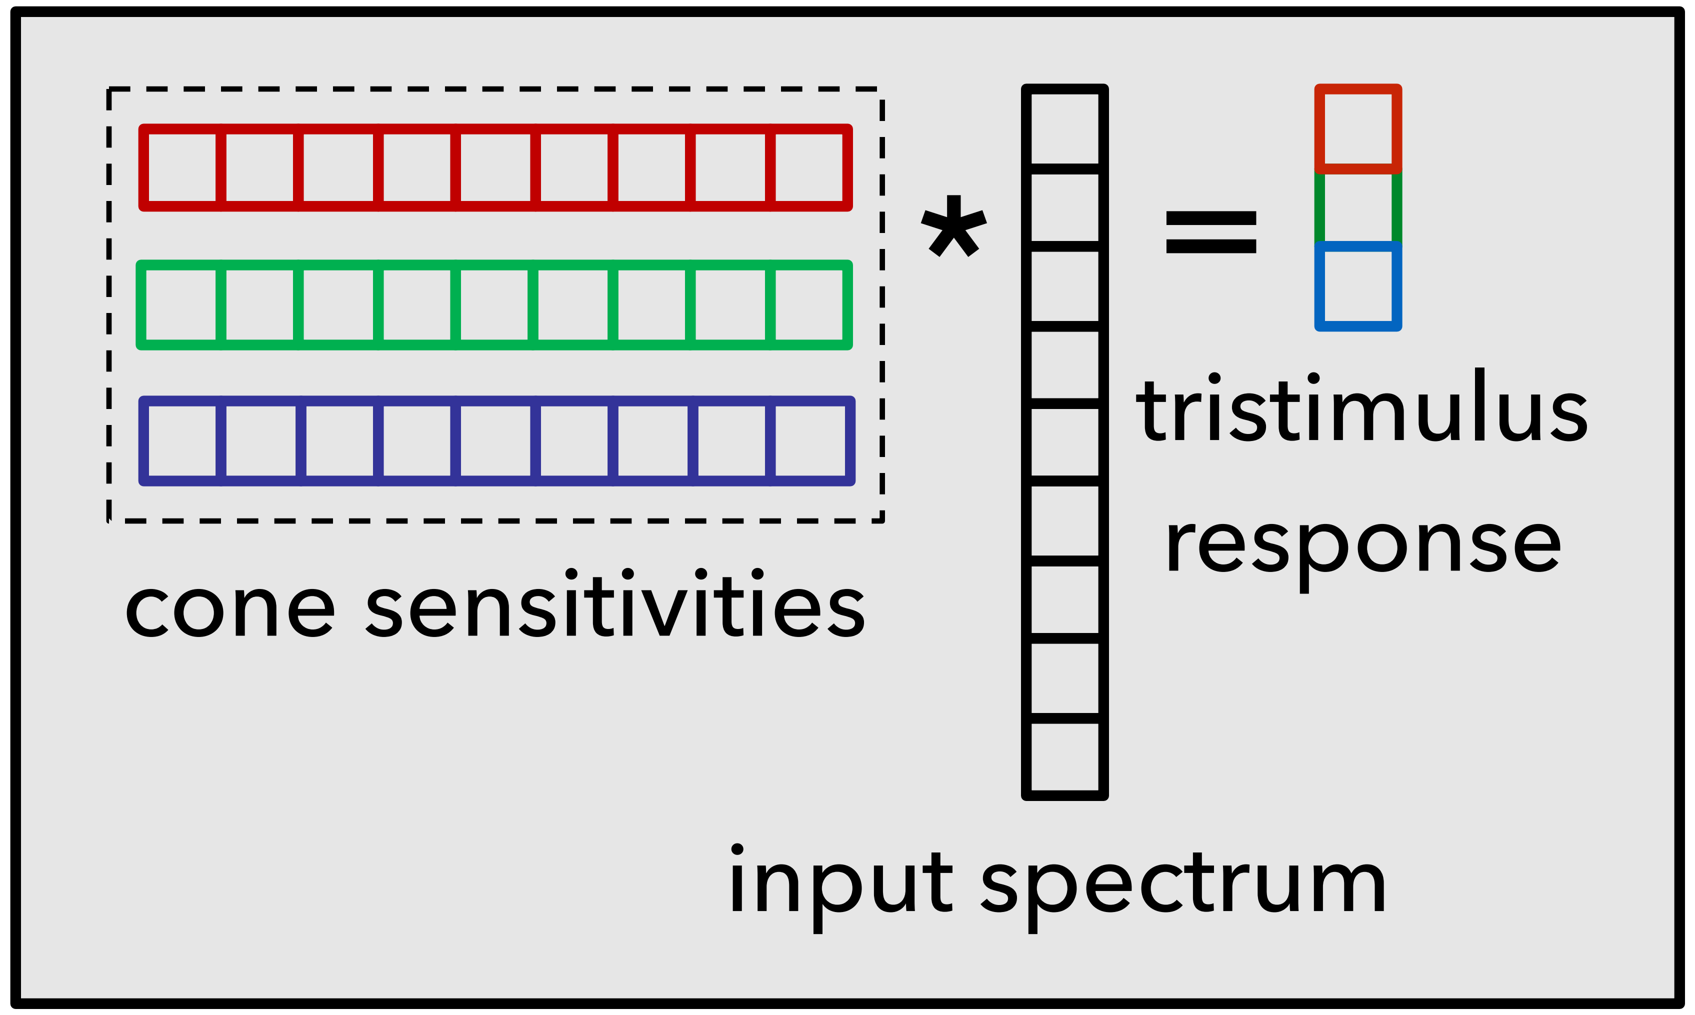

Tristimulus response is a matrix-vector multiplication

Cone responses to a spectrum s

Integral notation:

$$S = \int r_S(\lambda)\, s(\lambda)\, \mathrm{d}\lambda = r_S \cdot s$$

$$M = \int r_M(\lambda)\, s(\lambda)\, \mathrm{d}\lambda = r_M \cdot s$$

$$L = \int r_L(\lambda)\, s(\lambda)\, \mathrm{d}\lambda = r_L \cdot s$$

Matrix notation:

$$\begin{bmatrix} S \\ M \\ L \end{bmatrix} = \begin{bmatrix} \text{—} & r_S & \text{—} \\ \text{—} & r_M & \text{—} \\ \text{—} & r_L & \text{—} \end{bmatrix} \begin{bmatrix} | \\ s \\ | \end{bmatrix}$$

r_S, r_M and r_L are N-dimensional vectors, where $N = \infty$

Physical light to perceptual color

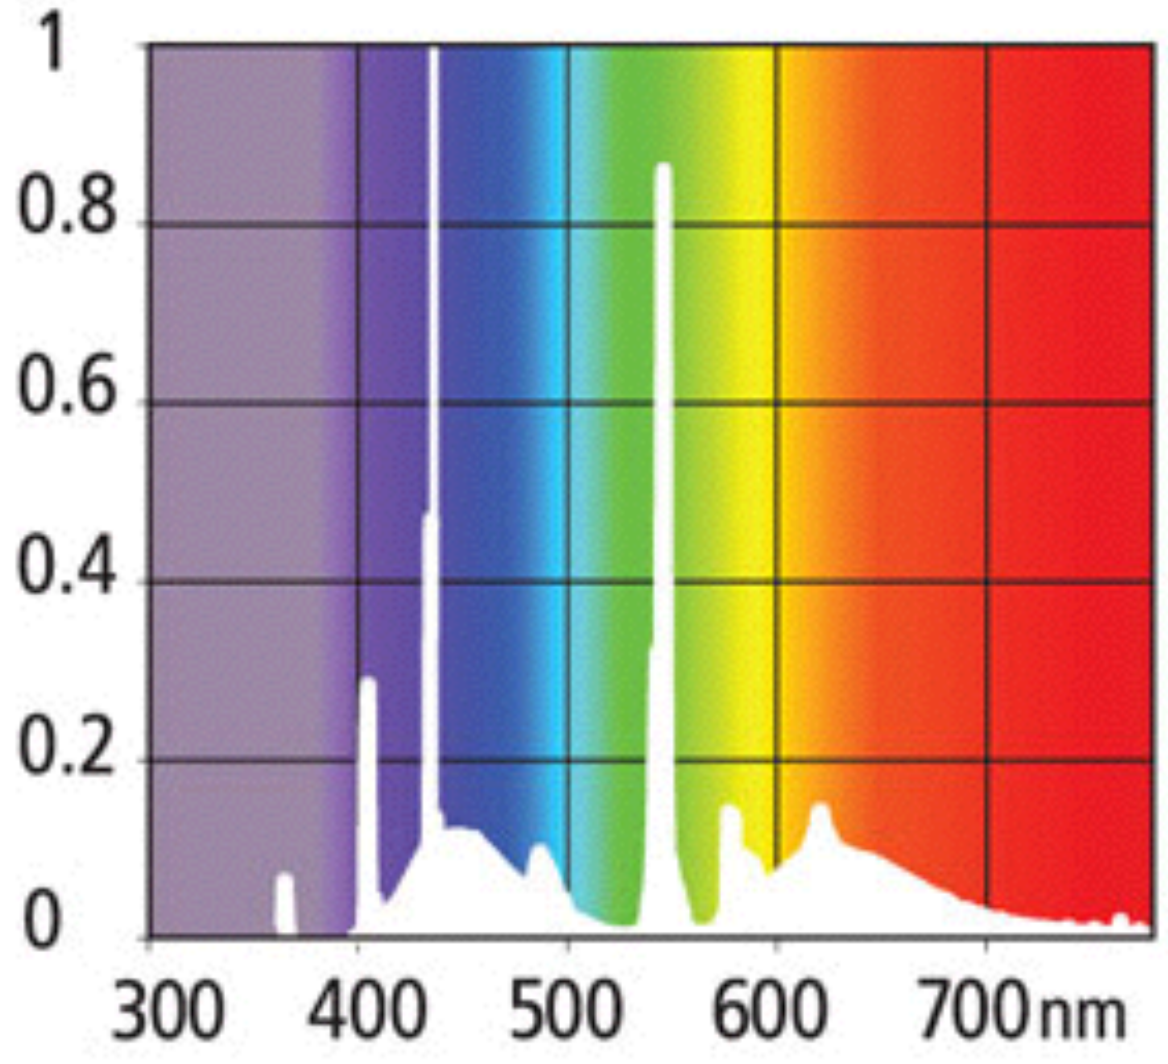

Physical

$$\begin{bmatrix} S \\ M \\ L \end{bmatrix} = \begin{bmatrix} — & r_S & — \\ — & r_M & — \\ — & r_L & — \end{bmatrix} \begin{bmatrix} | \\ s \\ | \end{bmatrix}$$

Perceptual

Basic fact of colorimetry

Take a spectrum (which is an infinity of numbers)

Eye produces three numbers (a projection to 3D)

This throws away a lot of information!

- many spectra can produce same S, M, L tristimulus values!
- *metamers*
- affected by illuminant

Warning: tricky thing with color

Cone responses overlap & are not orthogonal!

Basis functions for analysis

- eyes, cameras

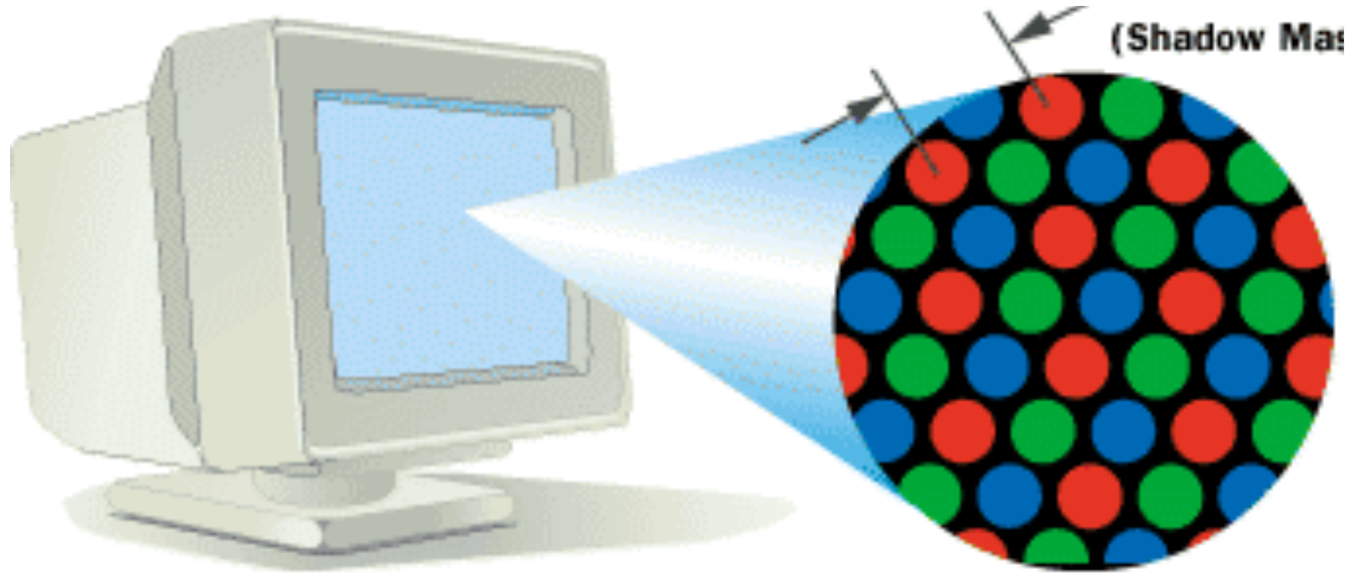

are different than for synthesis

- lights, monitors

The RGB in your camera is different than the RGB in your monitor!

Color reproduction (the right way)

We want to compute the combination of R, G, B that will project to the same visual response as s

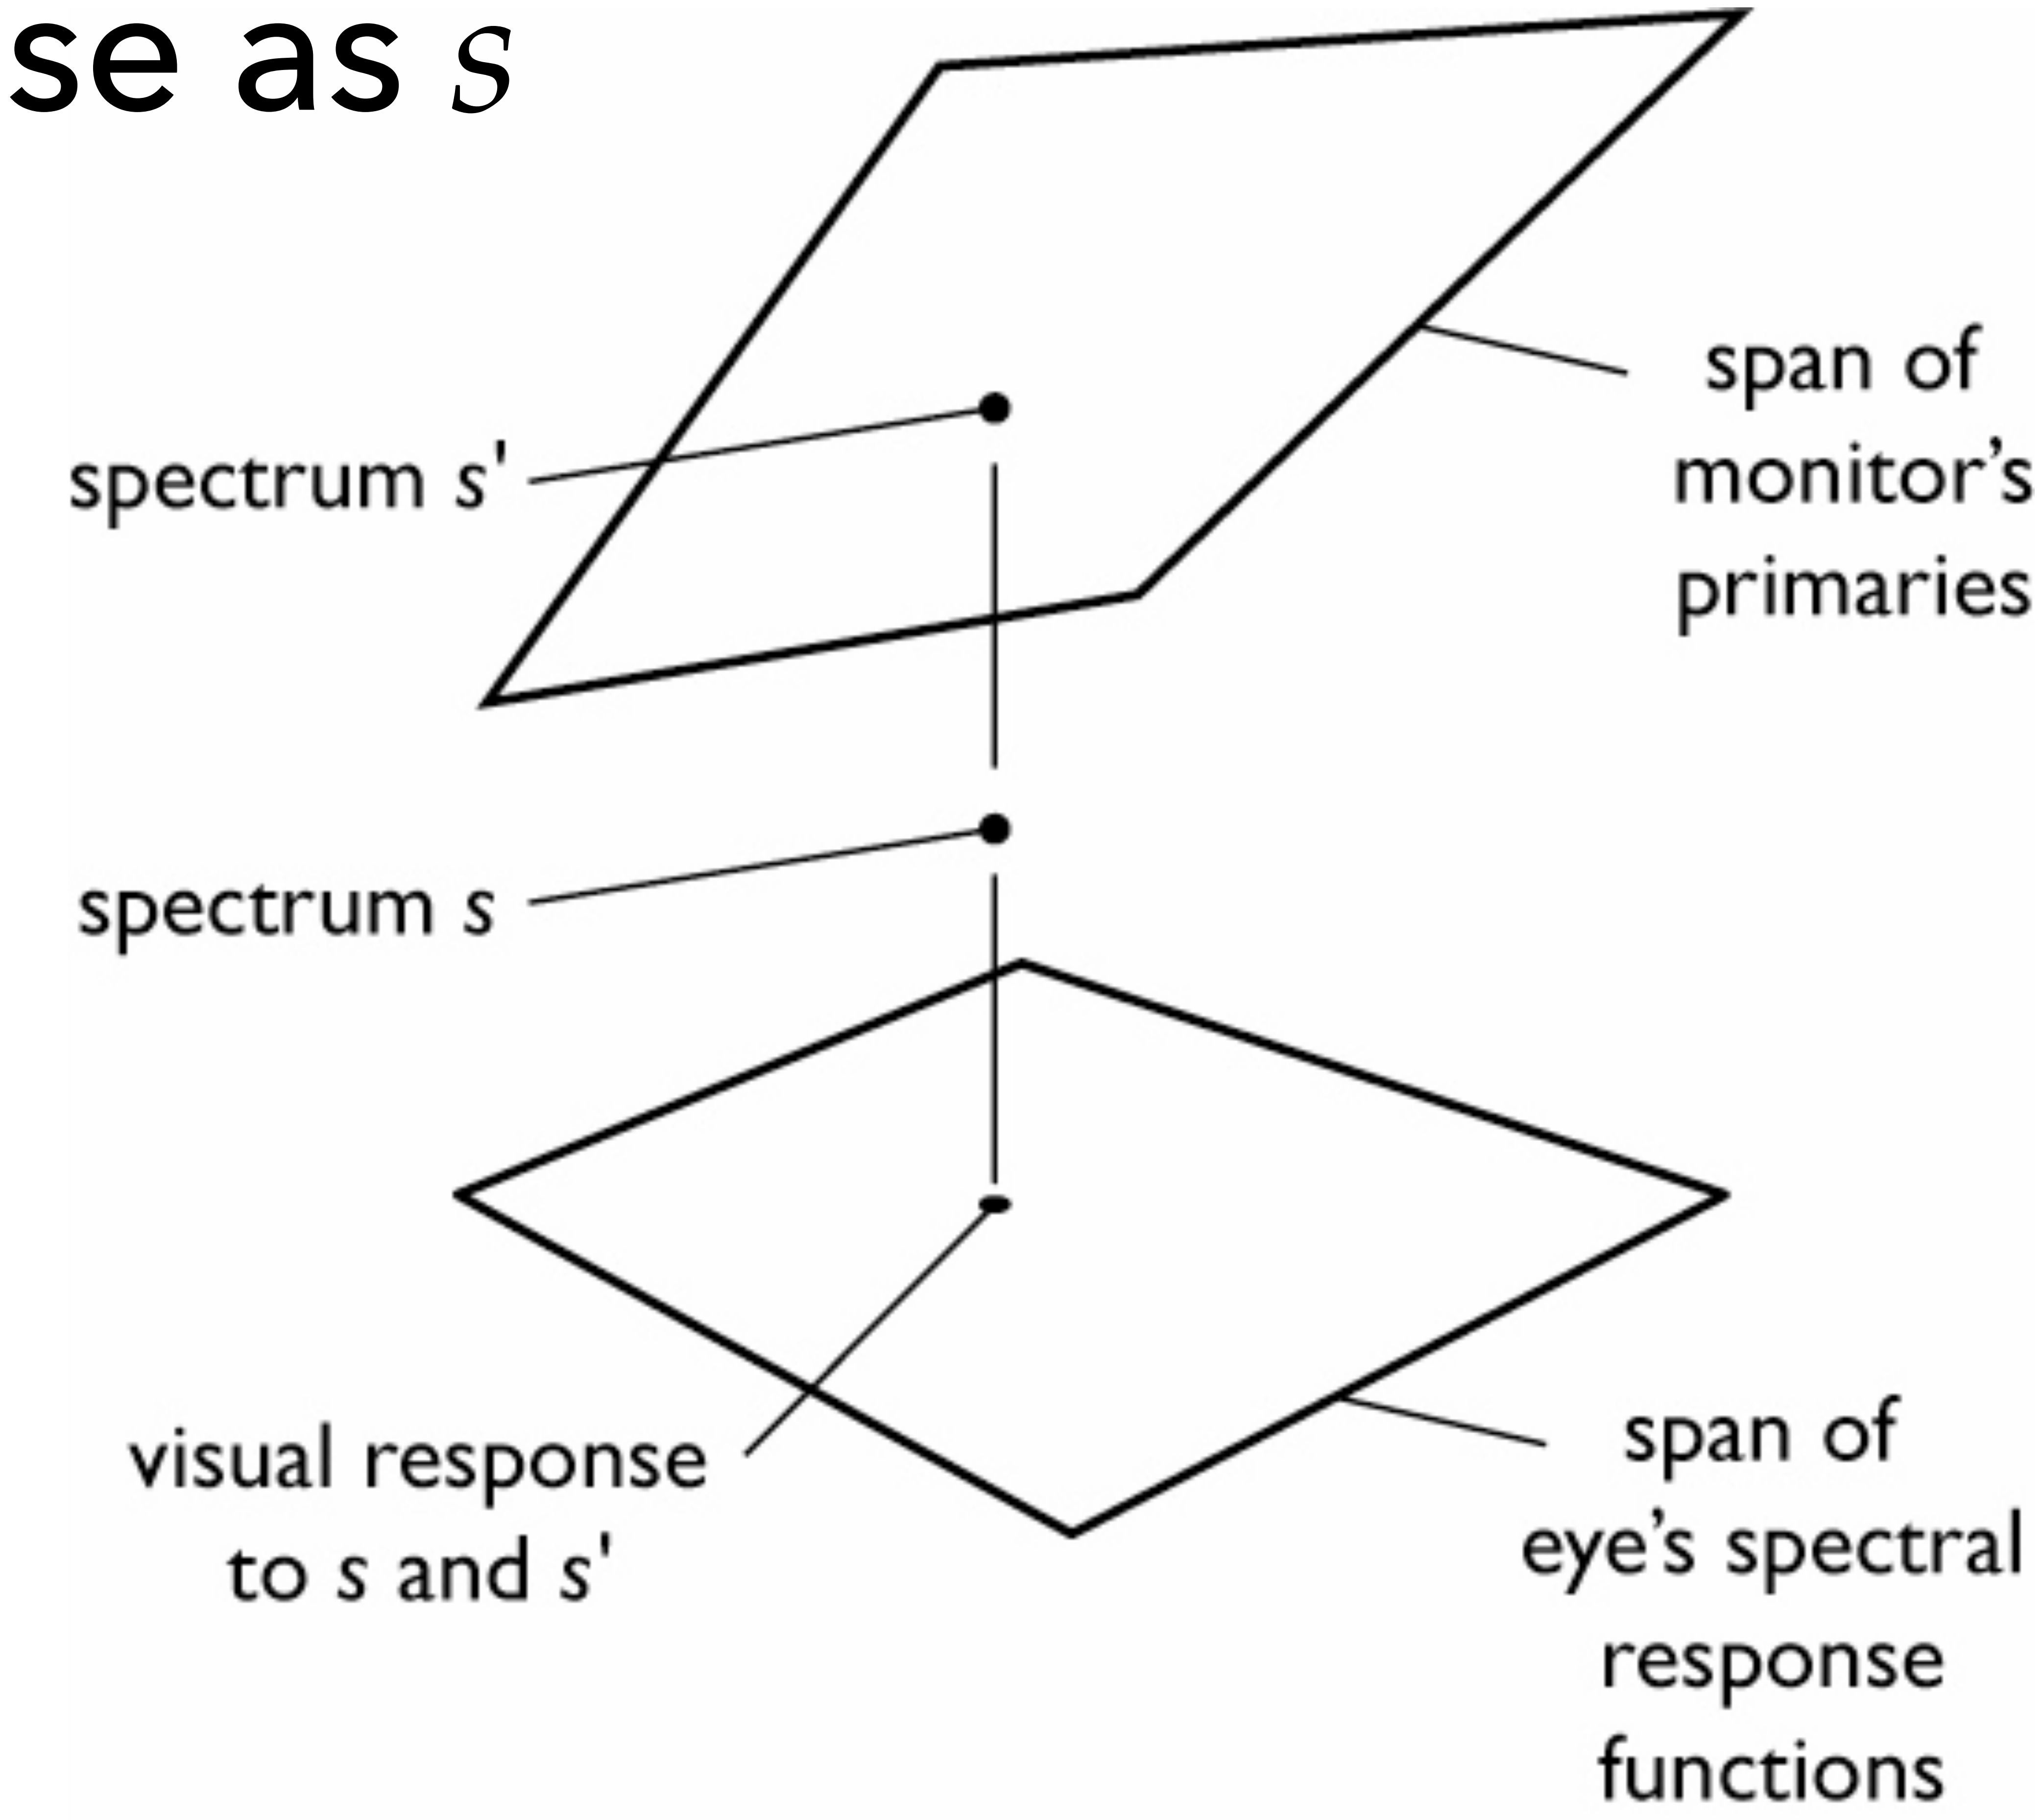

Color reproduction as linear algebra

What color do we see when we look at the display?

- Feed C to display

$$C$$

Color reproduction as linear algebra

What color do we see when we look at the display?

- Feed C to display
- Display produces s_a

$$M_{RGB}\, C$$

Color reproduction as linear algebra

What color do we see when we look at the display?

- Feed C to display
- Display produces s_a
- Eye looks at s_a and produces E

$$E = M_{SML}\, M_{RGB}\, C$$

$$\begin{bmatrix} S \\ M \\ L \end{bmatrix} = \begin{bmatrix} r_S \cdot s_R & r_S \cdot s_G & r_S \cdot s_B \\ r_M \cdot s_R & r_M \cdot s_G & r_M \cdot s_B \\ r_L \cdot s_R & r_L \cdot s_G & r_L \cdot s_B \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix}$$

Color reproduction as linear algebra

Goal of reproduction: visual response to s and s_a is the same:

$$M_{SML}\, s = M_{SML}\, s_a$$

Substitute in expression for s_a,

$$M_{SML}\, s = M_{SML}\, M_{RGB}\, C$$

$$C = \underbrace{(M_{SML}\, M_{RGB})^{-1} M_{SML}}_{\text{color matching matrix for RGB}}\, s$$

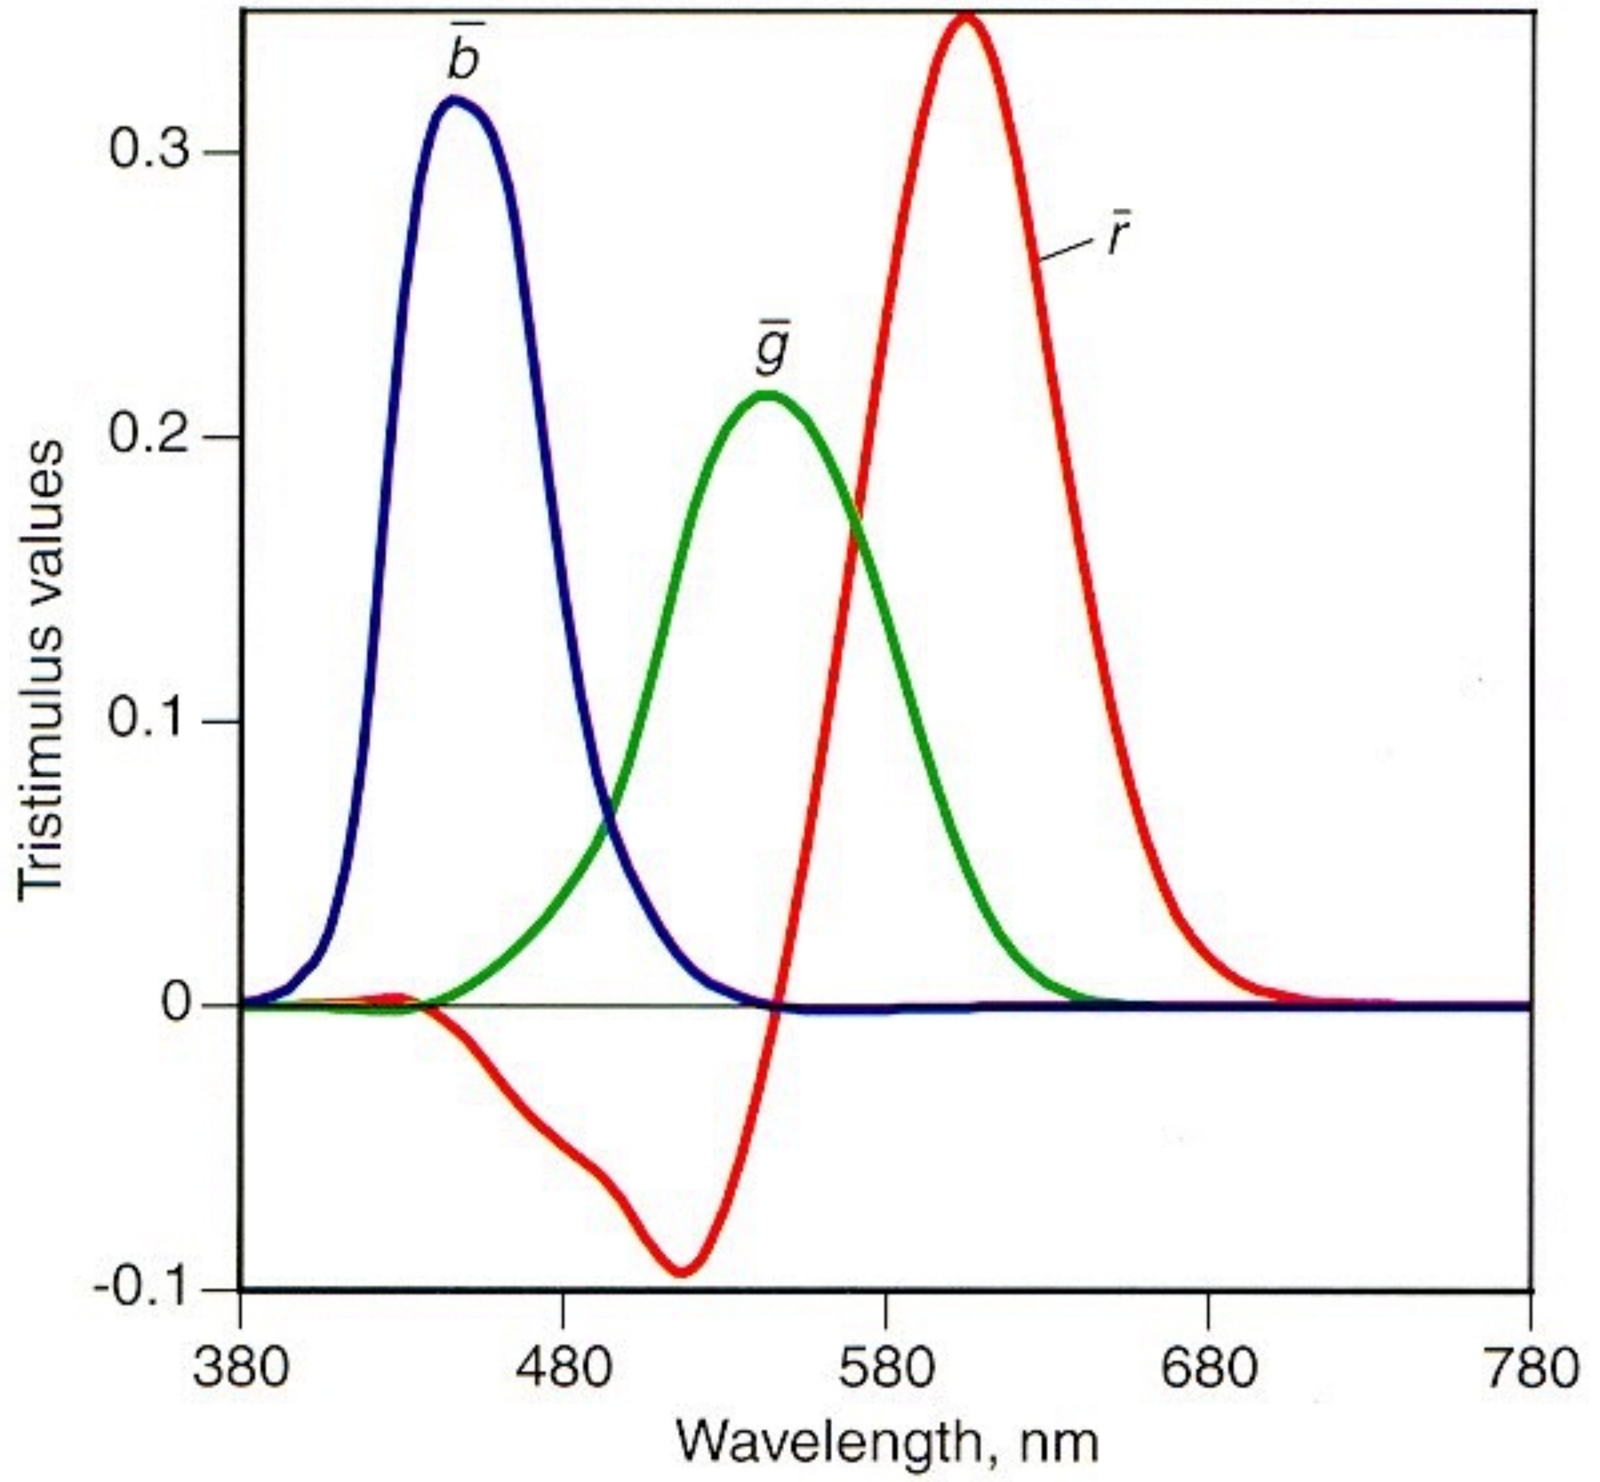

These curves are the color-matching functions for the 1931 standard observer, The average results of 17 color-normal observers having matched each wavelength of the equal-energy spectrum with primaries of 435.8 nm, 546.1 nm, and 700 nm.

Meaning of these curves/rows

Monochromatic wavelength λ can be reproduced with:

$b(\lambda)$ amount of the 435.8nm primary,

$+ g(\lambda)$ amount of the 546.1 primary,

$+ r(\lambda)$ amount of the 700 nm primary

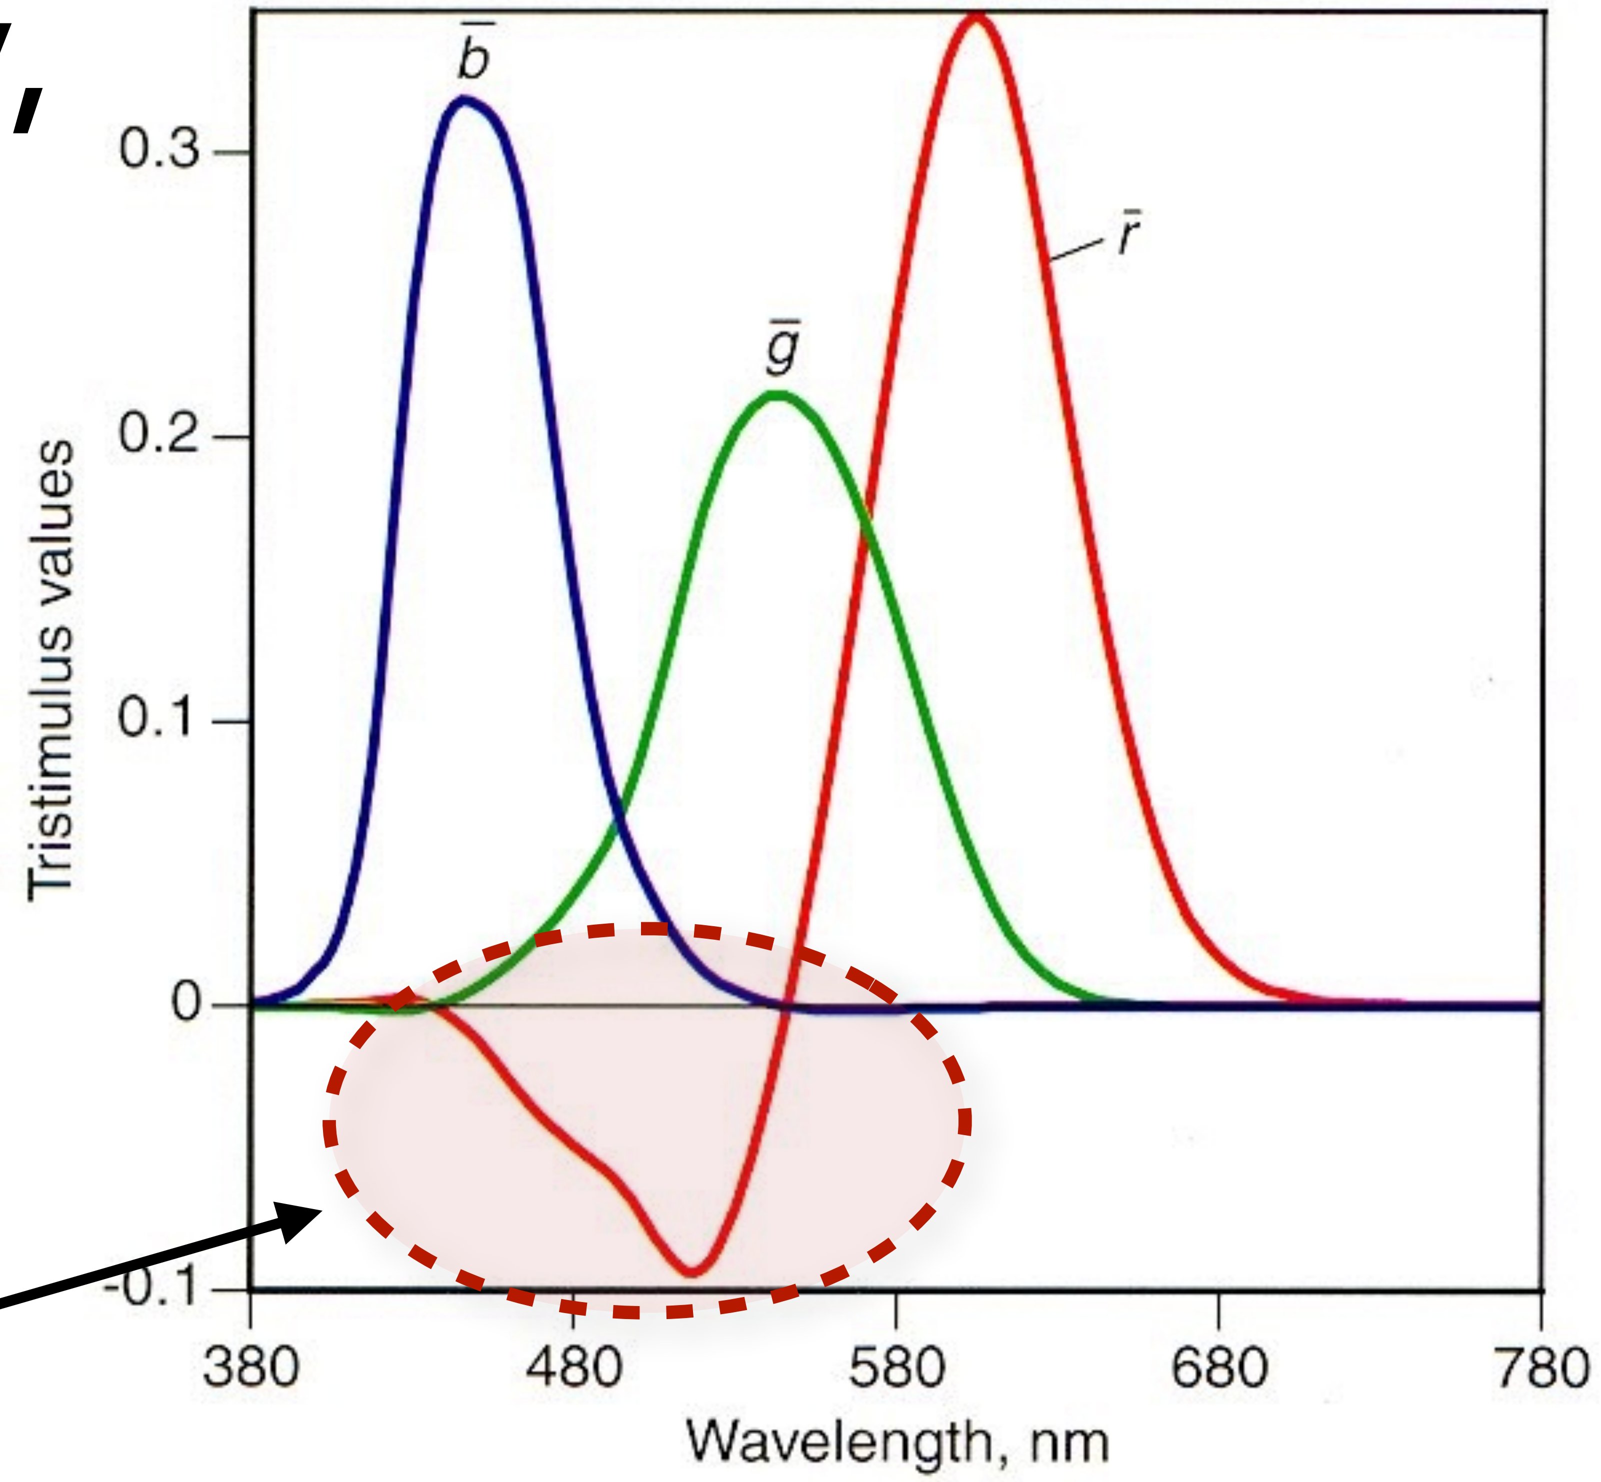

These curves are the color-matching functions for the 1931 standard observer, The average results of 17 color-normal observers having matched each wavelength of the equal-energy spectrum with primaries of 435.8 nm, 546.1 nm, and 700 nm.

Negative light required?

CIE XYZ color space

Linear algebra to the rescue!

Purely positive basis functions

Linear transformation of CIE RGB

Non-physical primaries

CIE XYZ ←- CIE RGB

$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = \frac{1}{0.17697} \begin{bmatrix} 0.49 & 0.31 & 0.20 \\ 0.17697 & 0.81240 & 0.01063 \\ 0.00 & 0.01 & 0.99 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix}$$

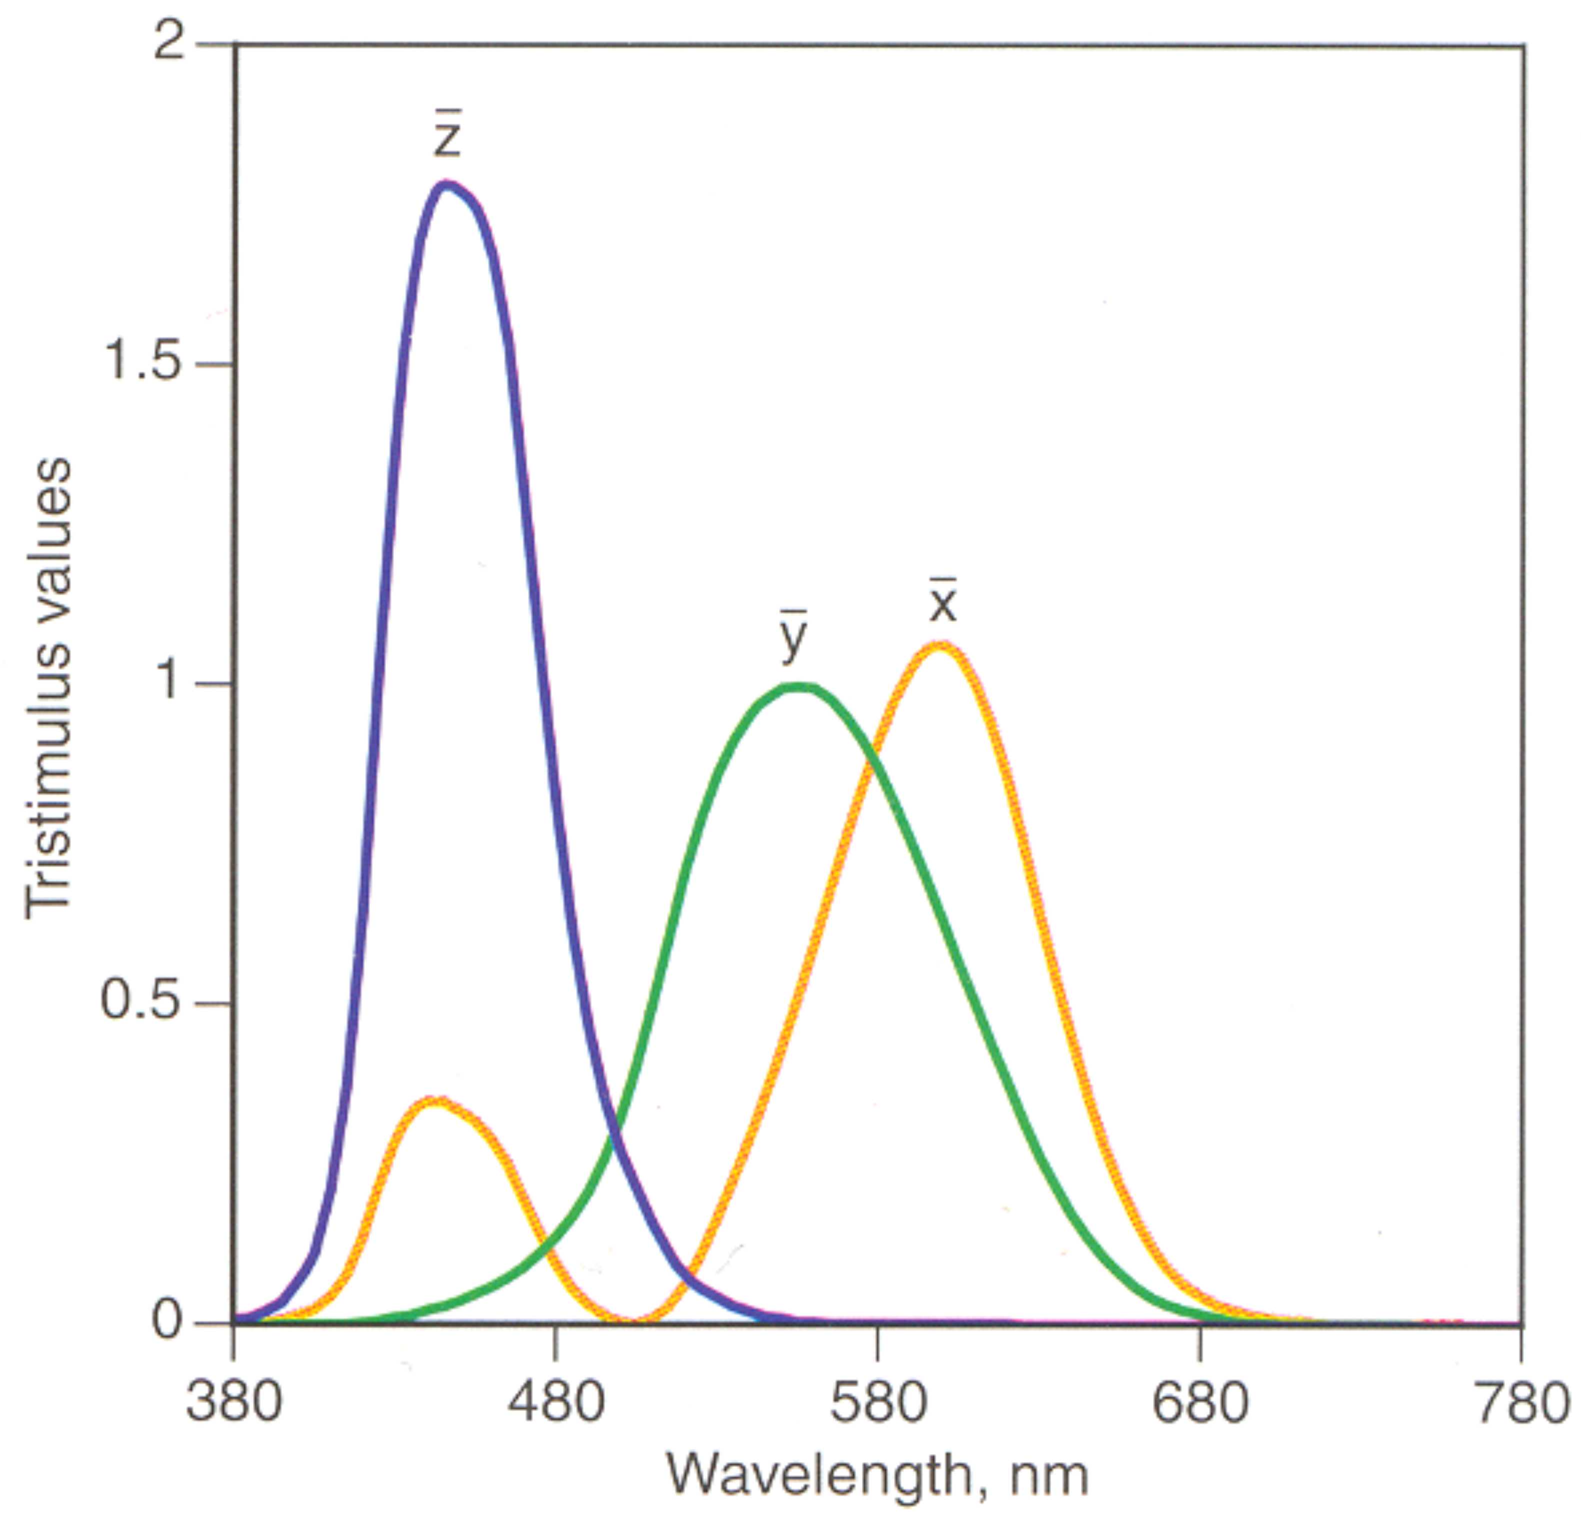

CIE XYZ color cone

3D spaces can be hard to visualize

Chrominance is our notion of color, as opposed to brightness/luminance

Recall that our eyes correct for multiplicative scale factors

- discount light intensity

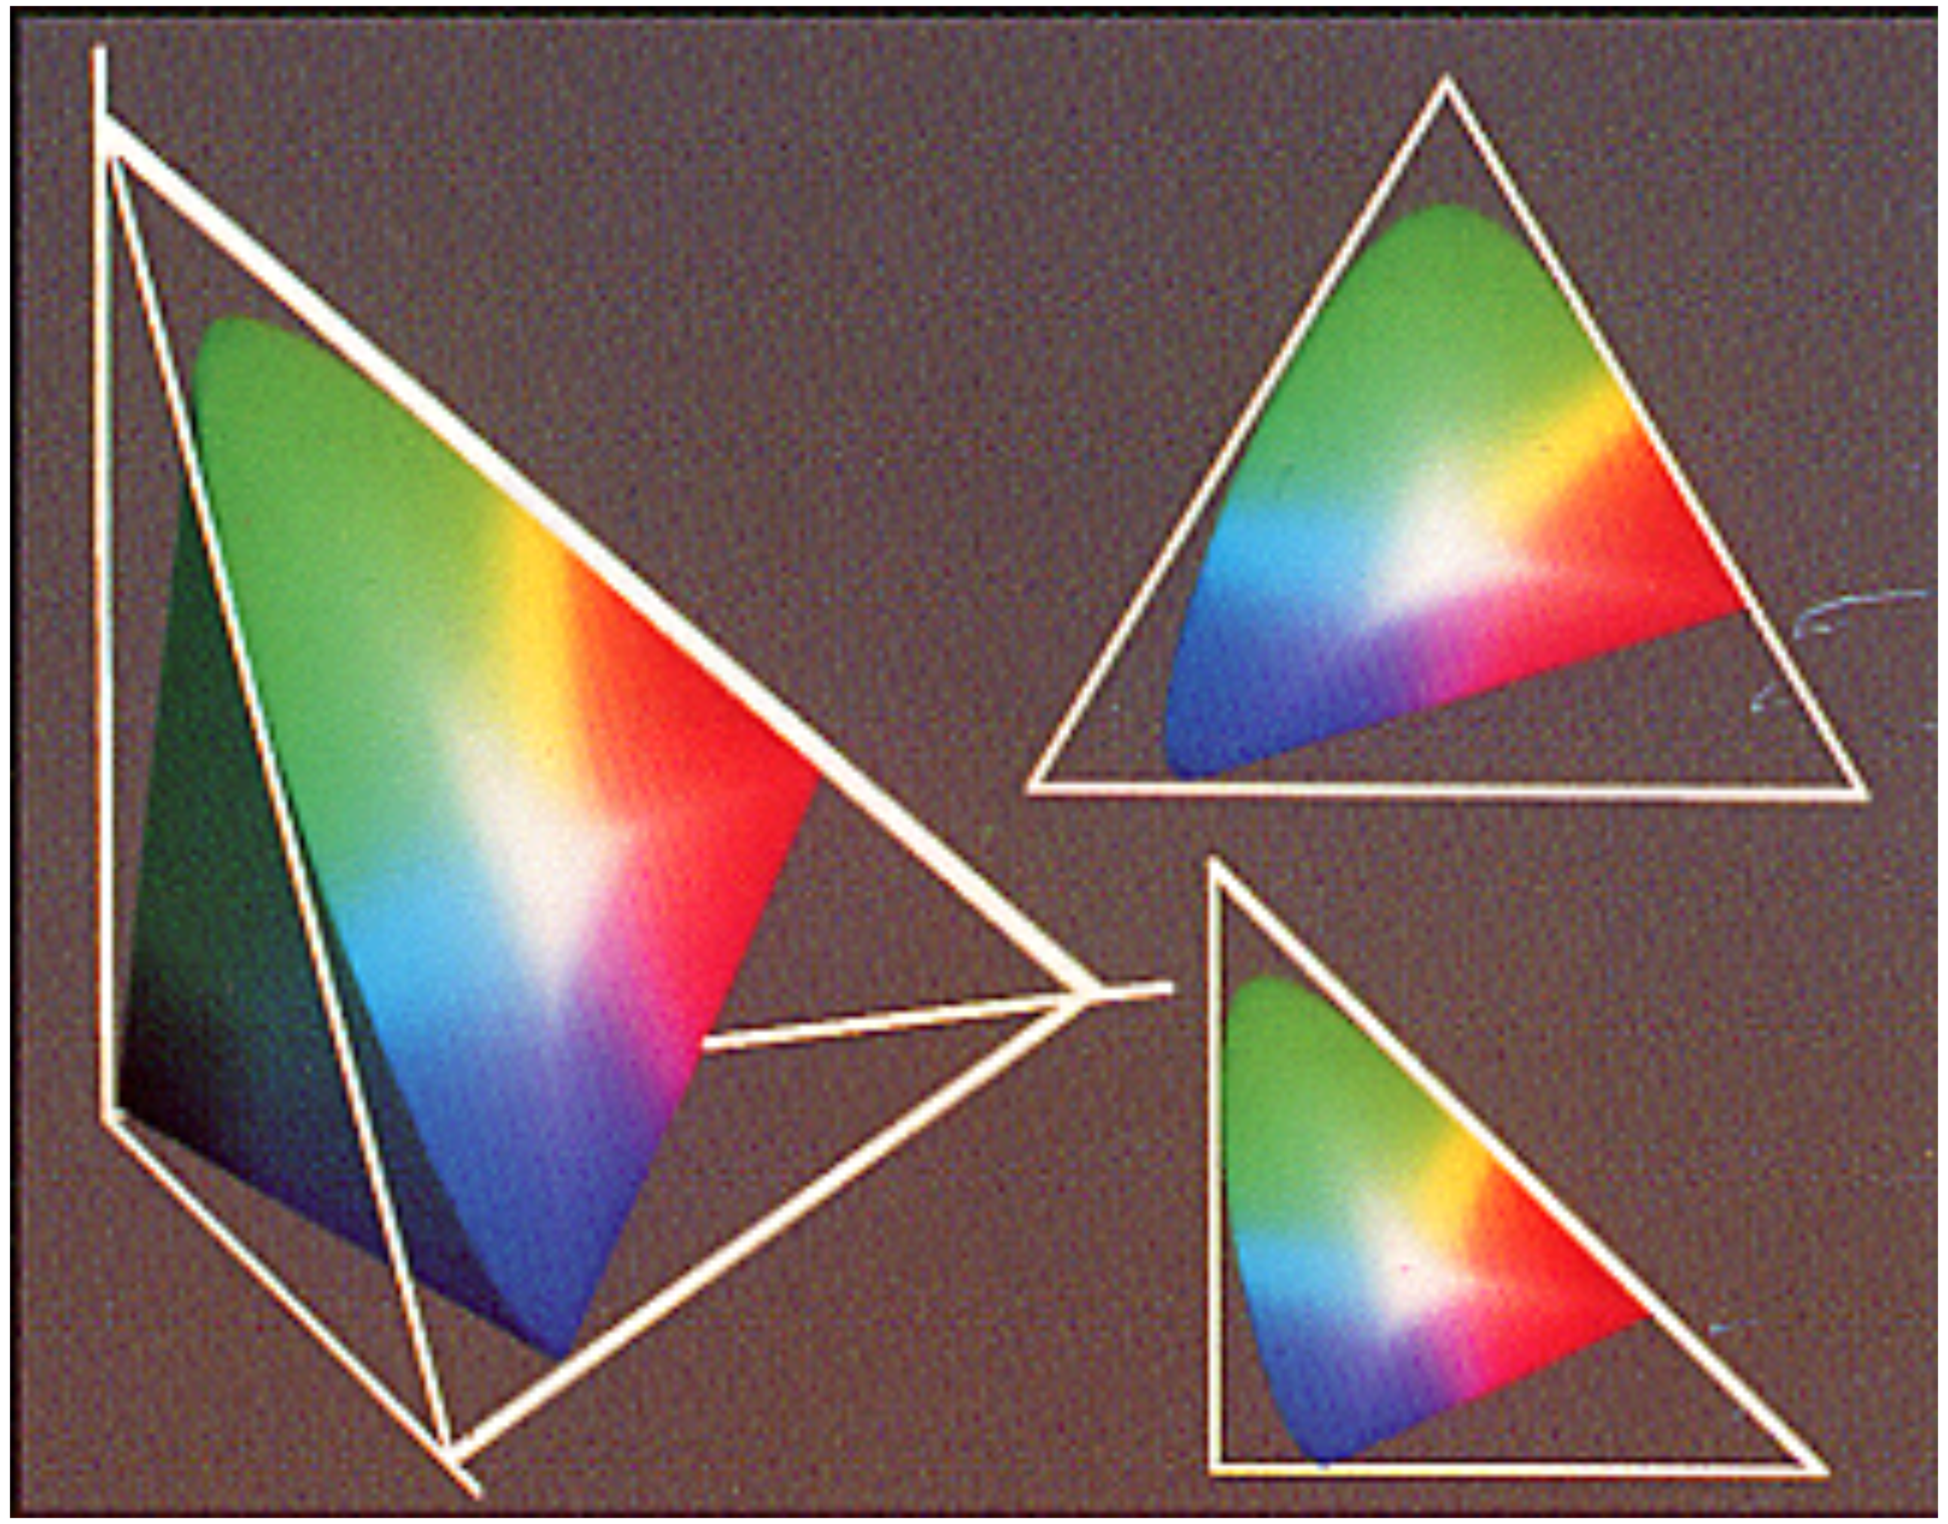

Chromaticity Diagram

The CIE xyY Color Space

Chromaticity (x,y) can be derived by normalizing the XYZ color components:

$$x = \frac{X}{X+Y+Z} \qquad y = \frac{Y}{X+Y+Z}$$

- (x,y) characterize color
- Y characterizes *brightness*

Combining xy with Y allows us to represent any color

Plotting on xy plane allows us to see all colors of a single brightness

CIE Chromaticity Chart

Spectral colors along curved boundary

Linear combination of two colors: line connecting two points

Linear combination of 3 colors span a triangle (color gamut)

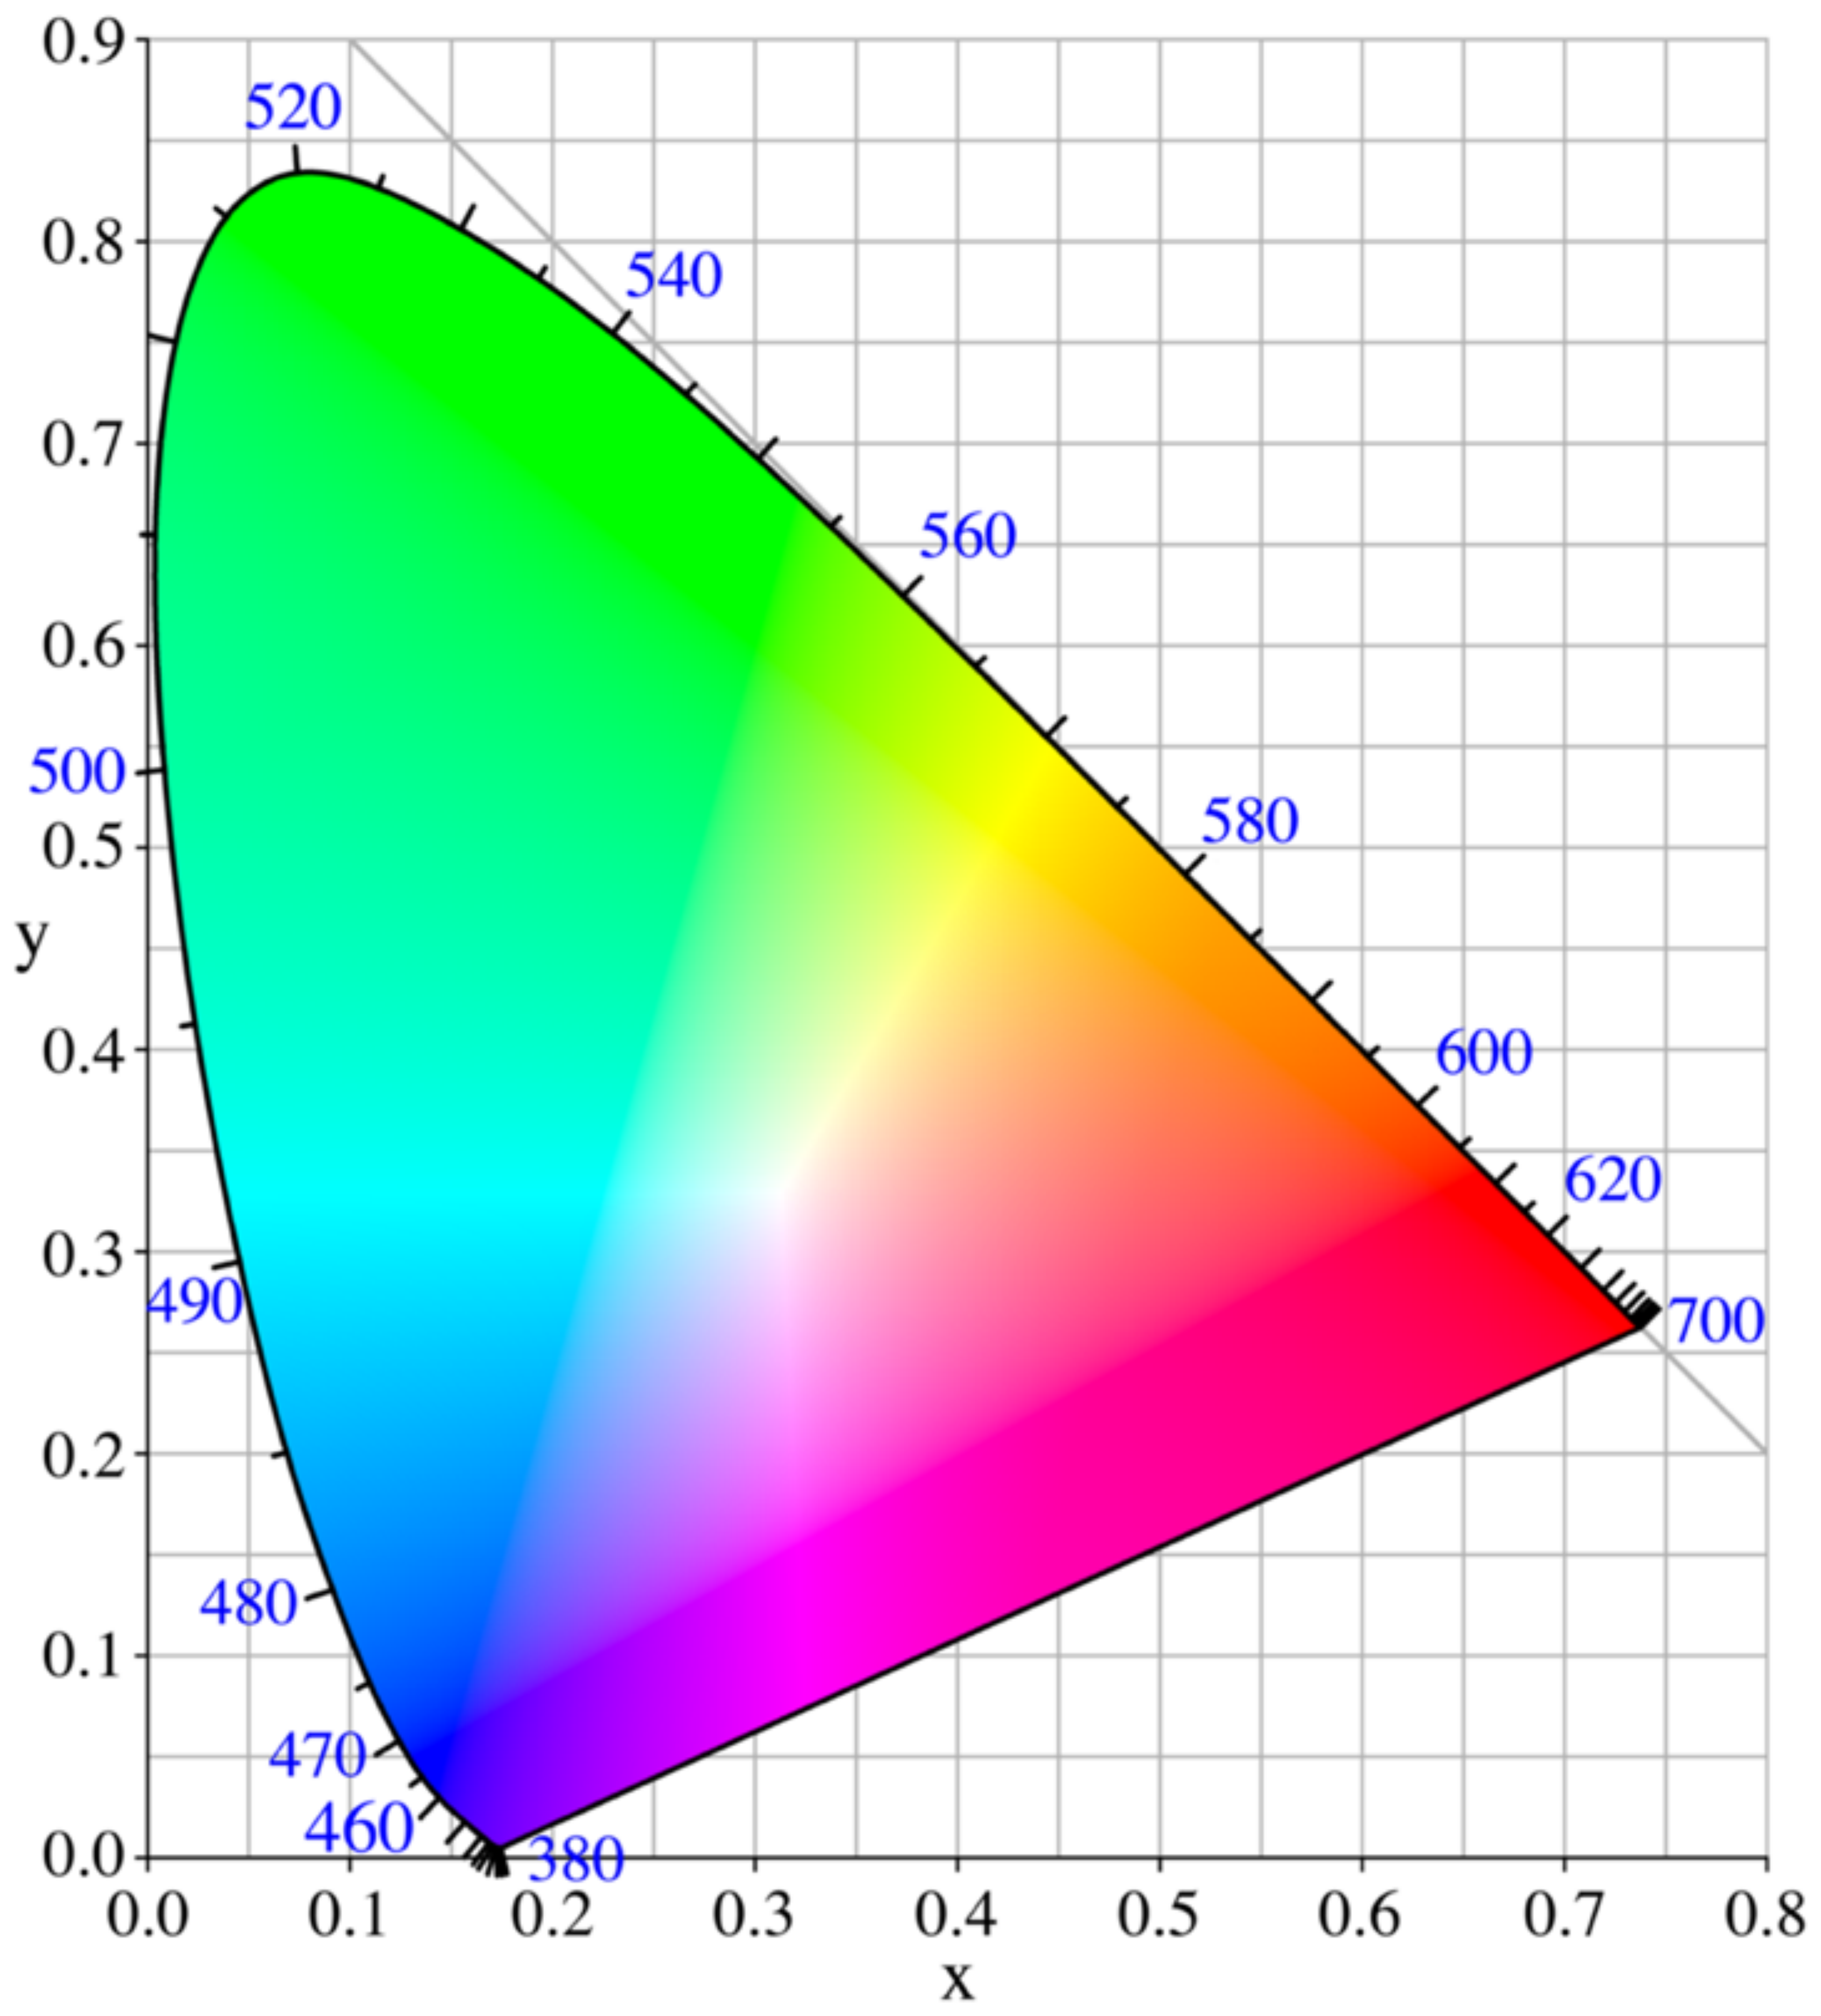

CIE RGB Color Space

Color primaries at:
435.8, 546.1, 700.0 nm

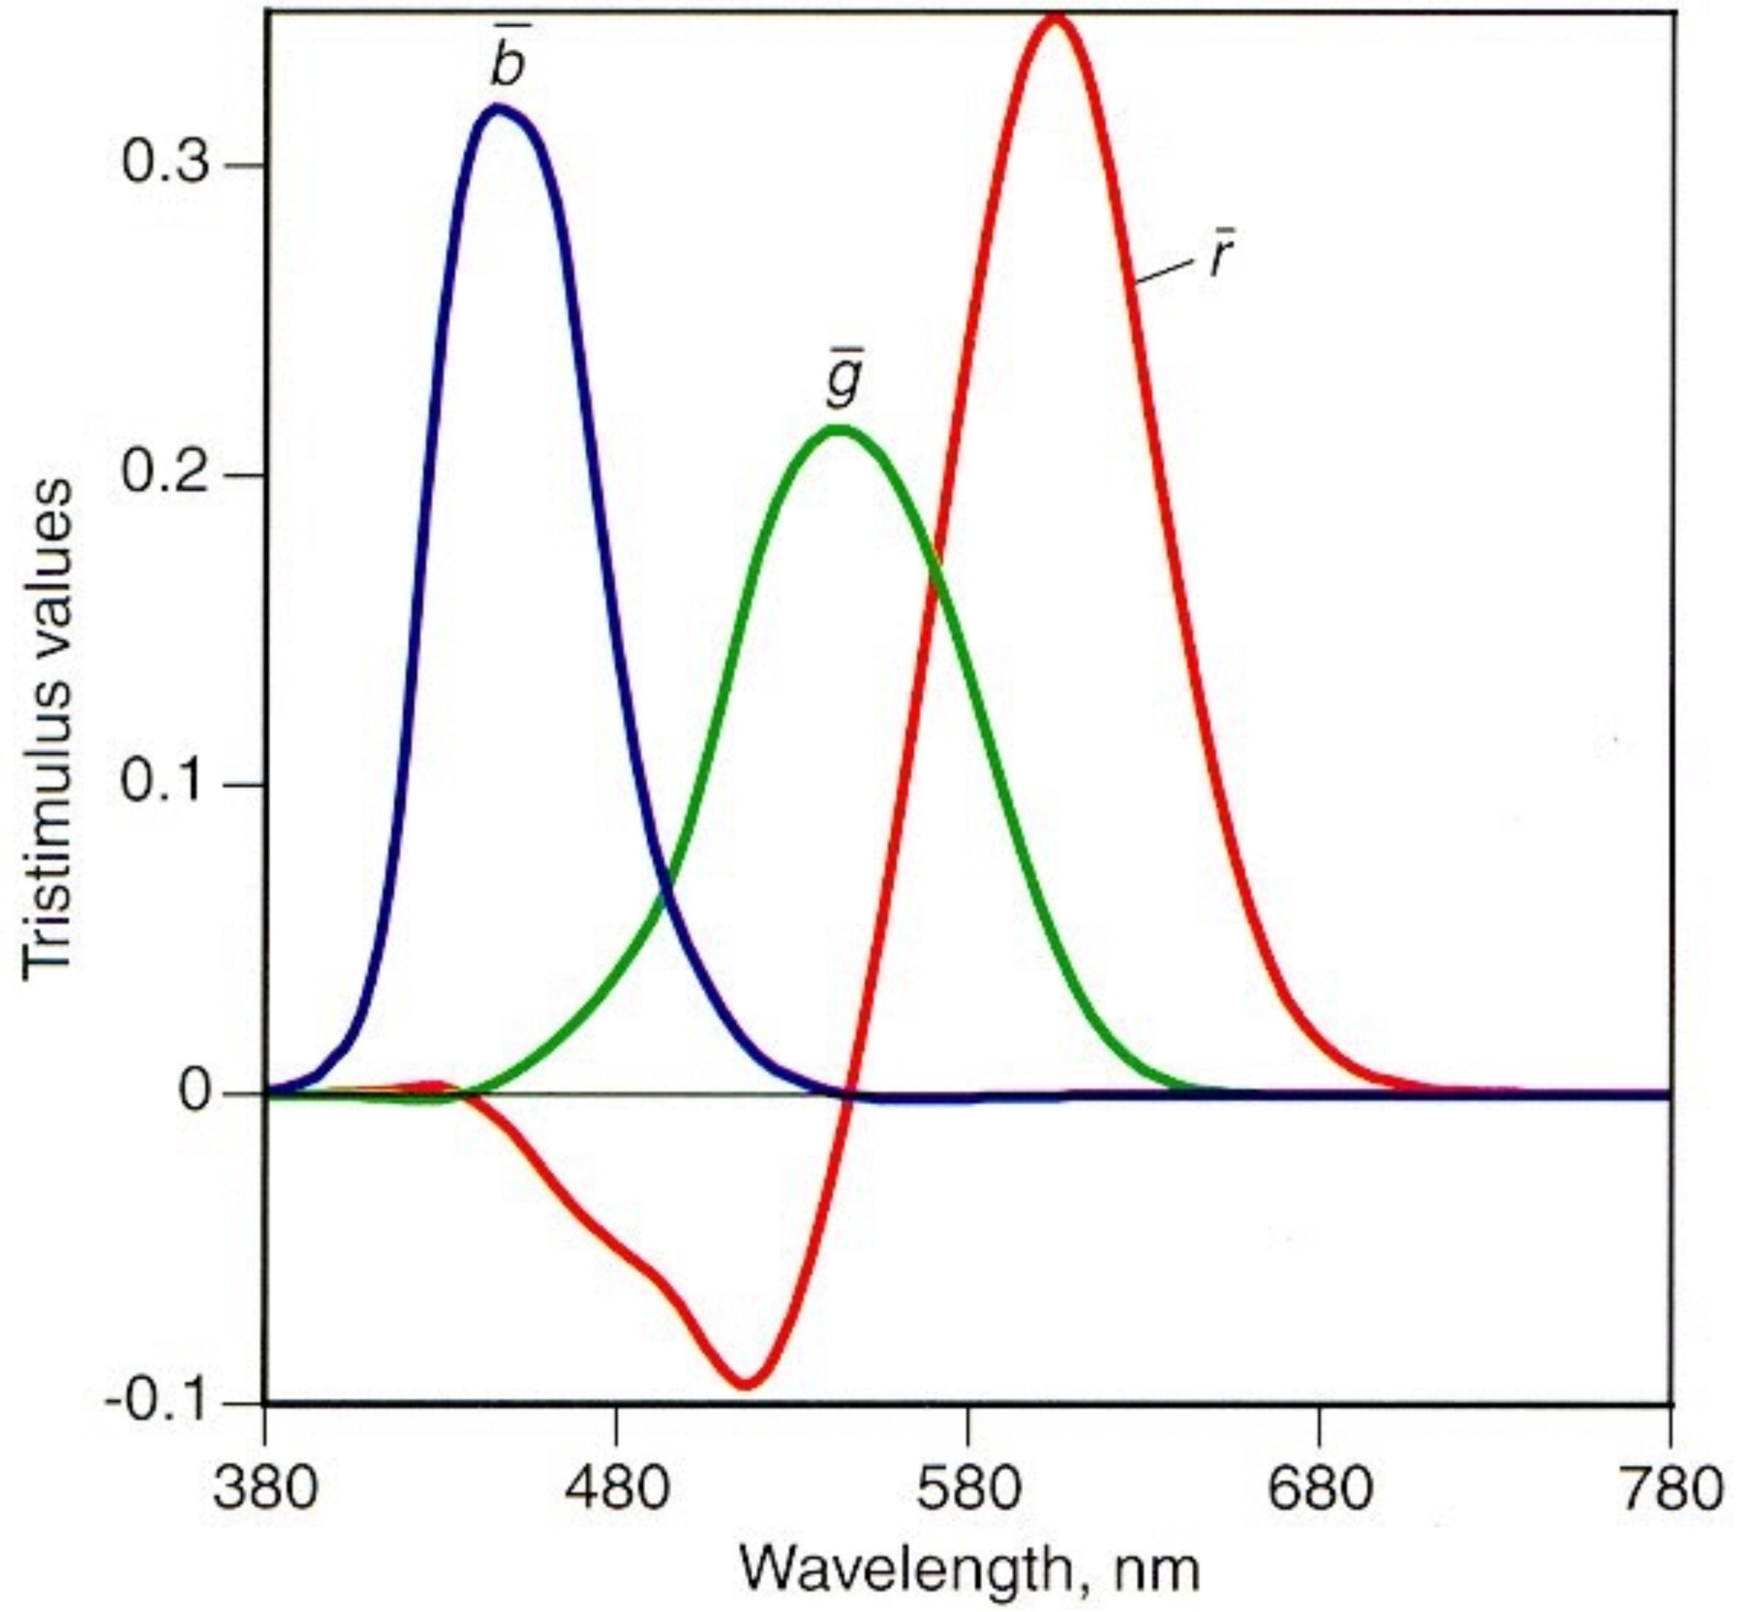

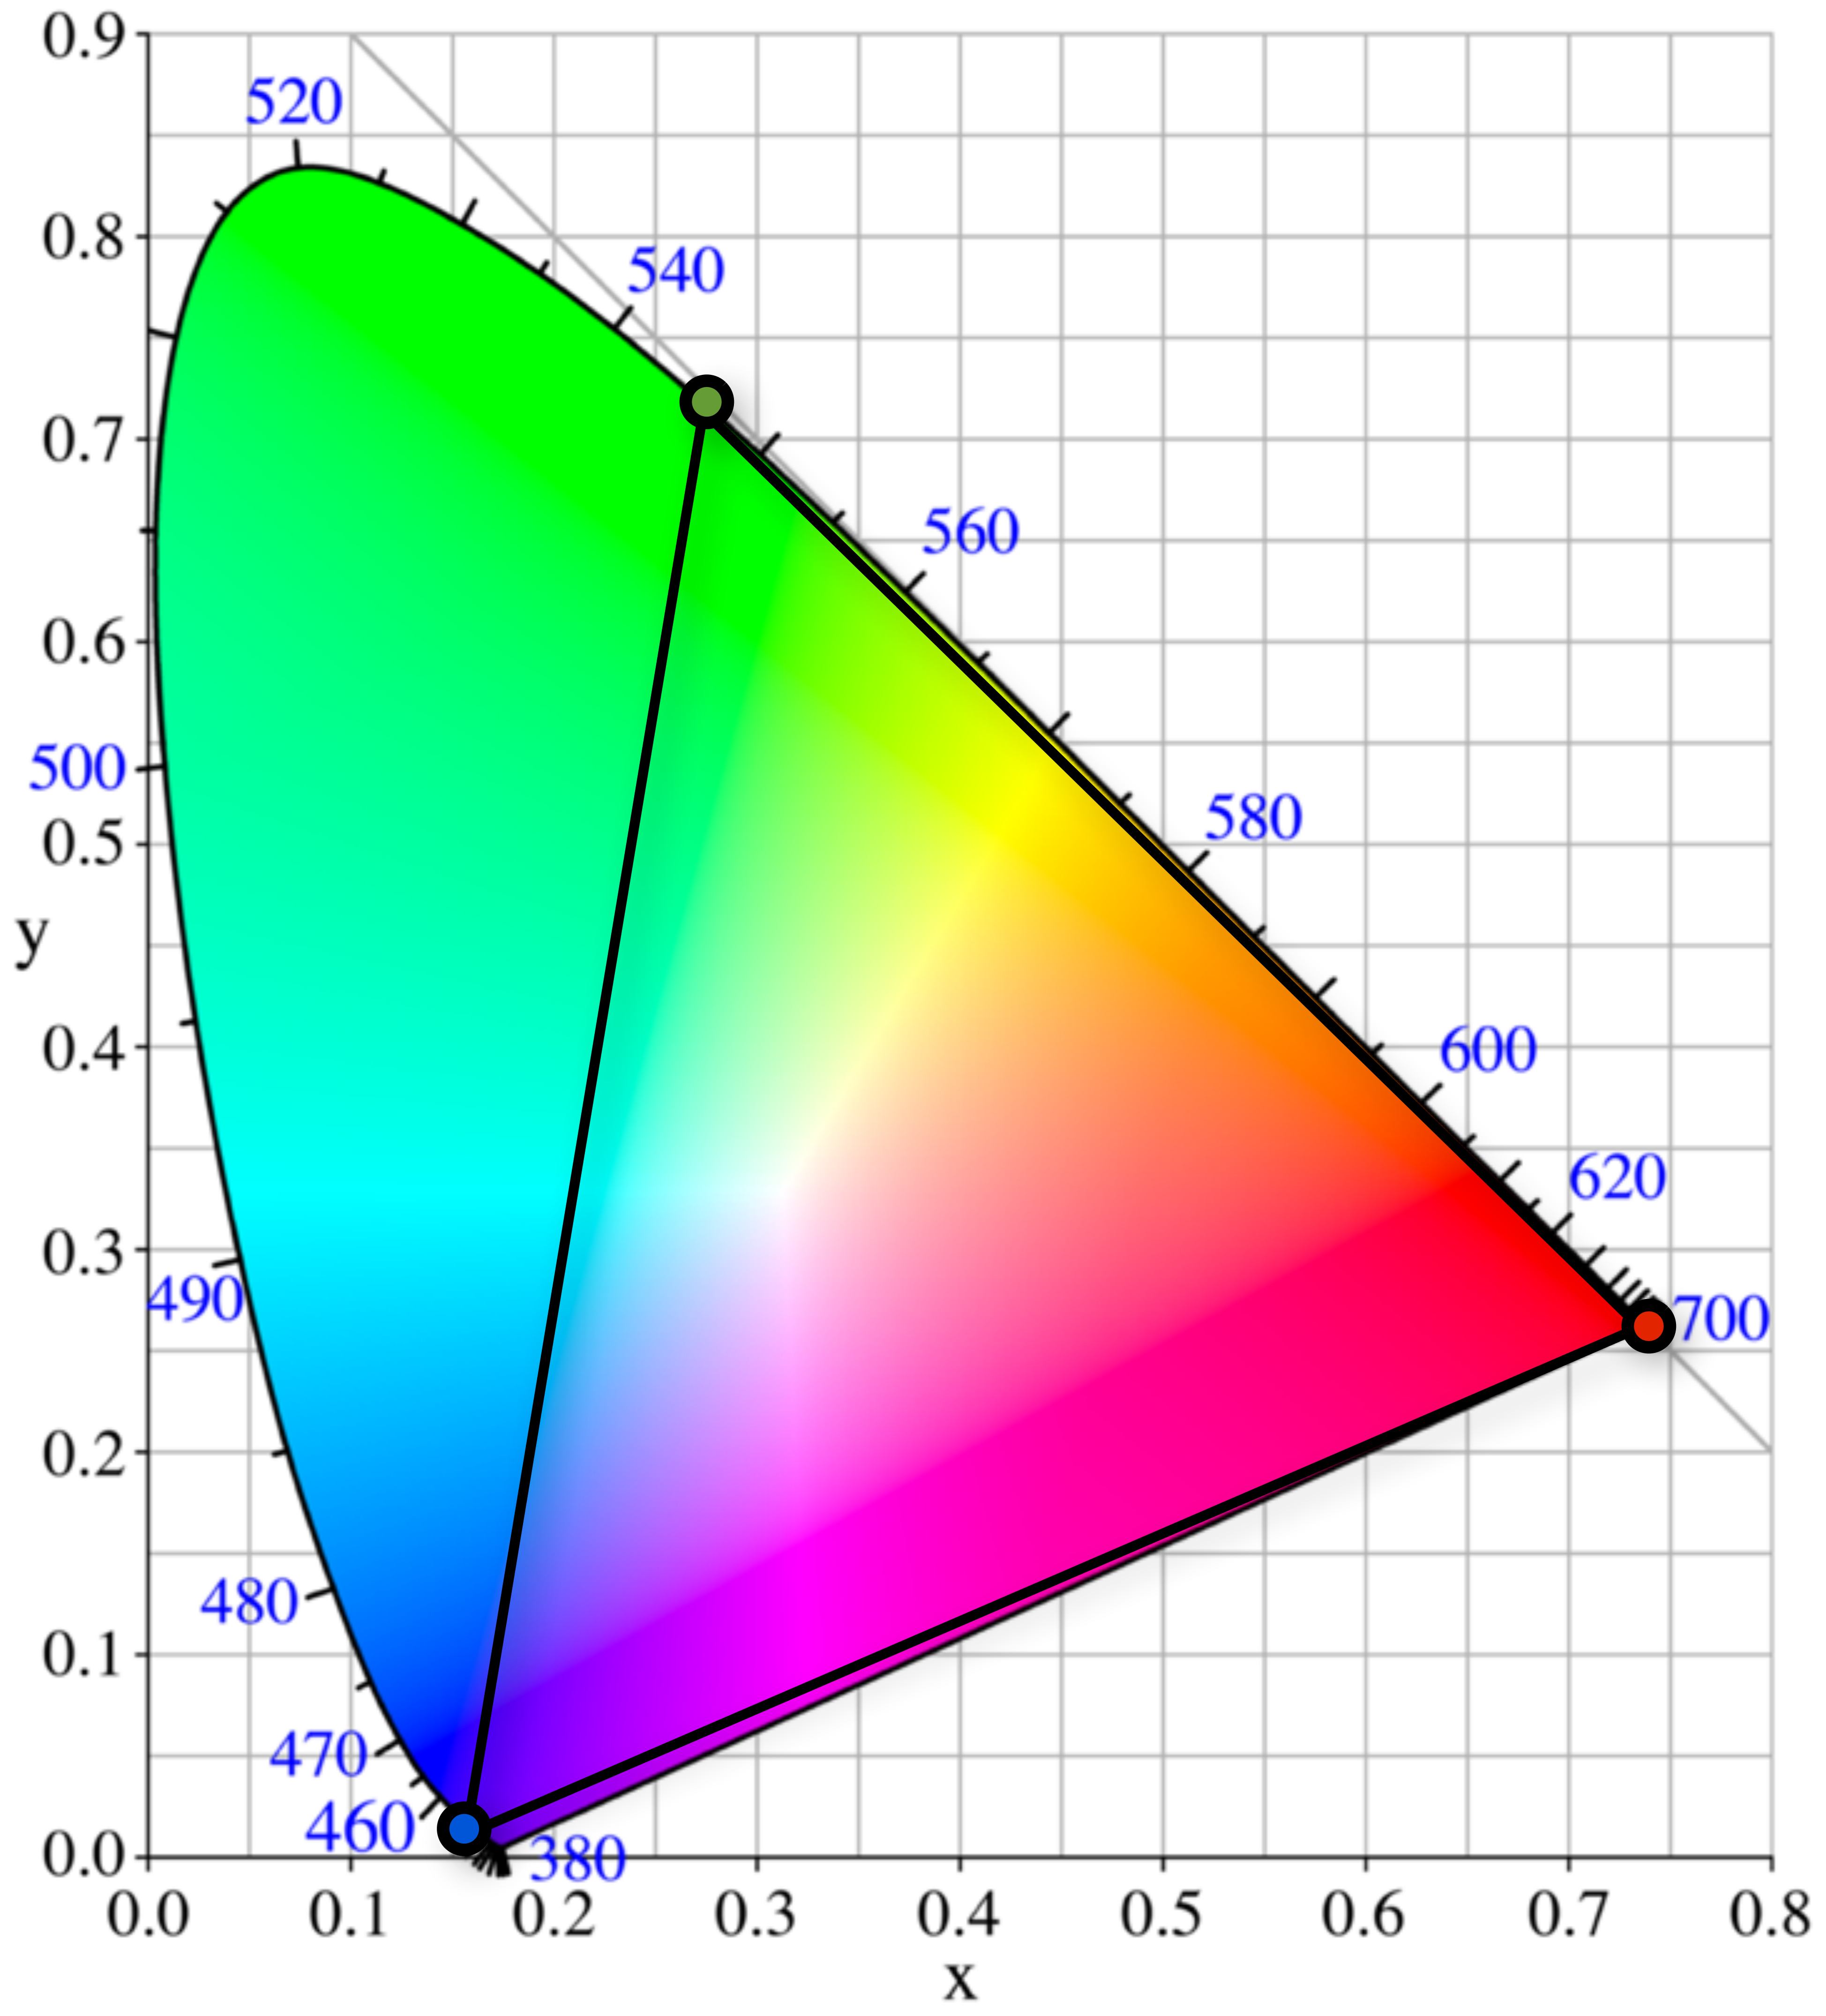

Color Gamuts

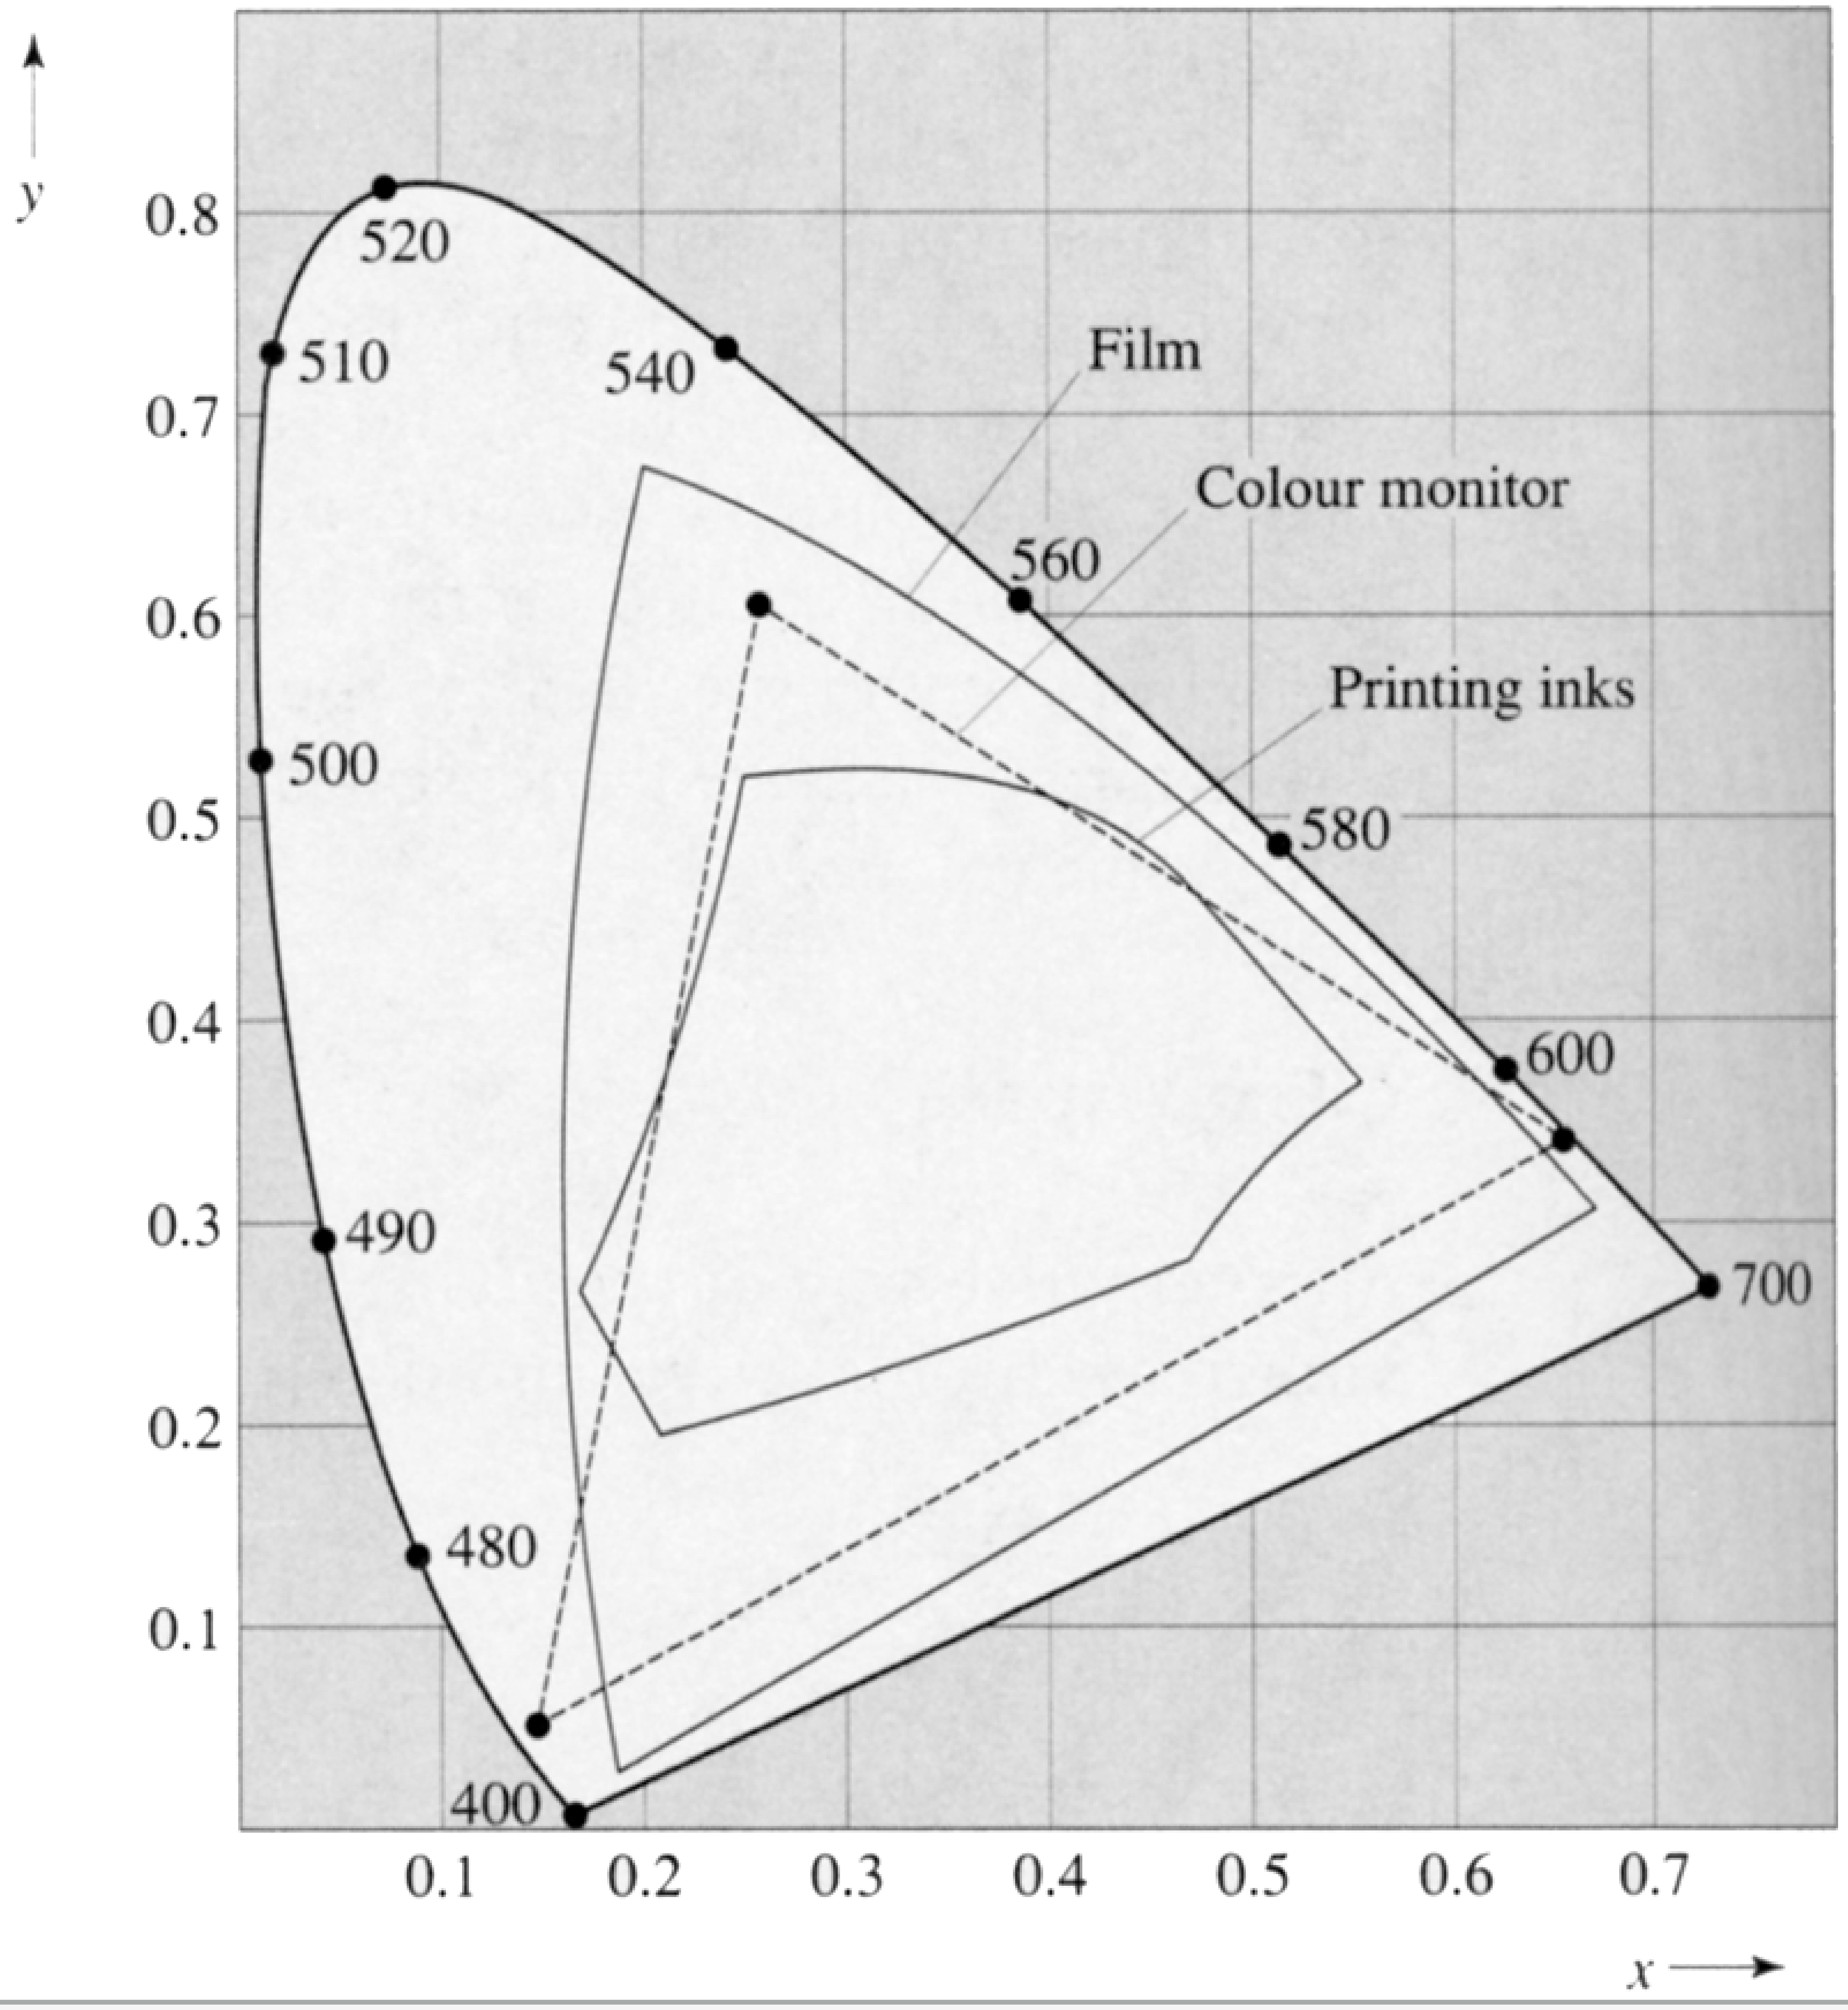

CIE Chromaticity Chart Features

White Point

Dominant wavelength

Inverse color

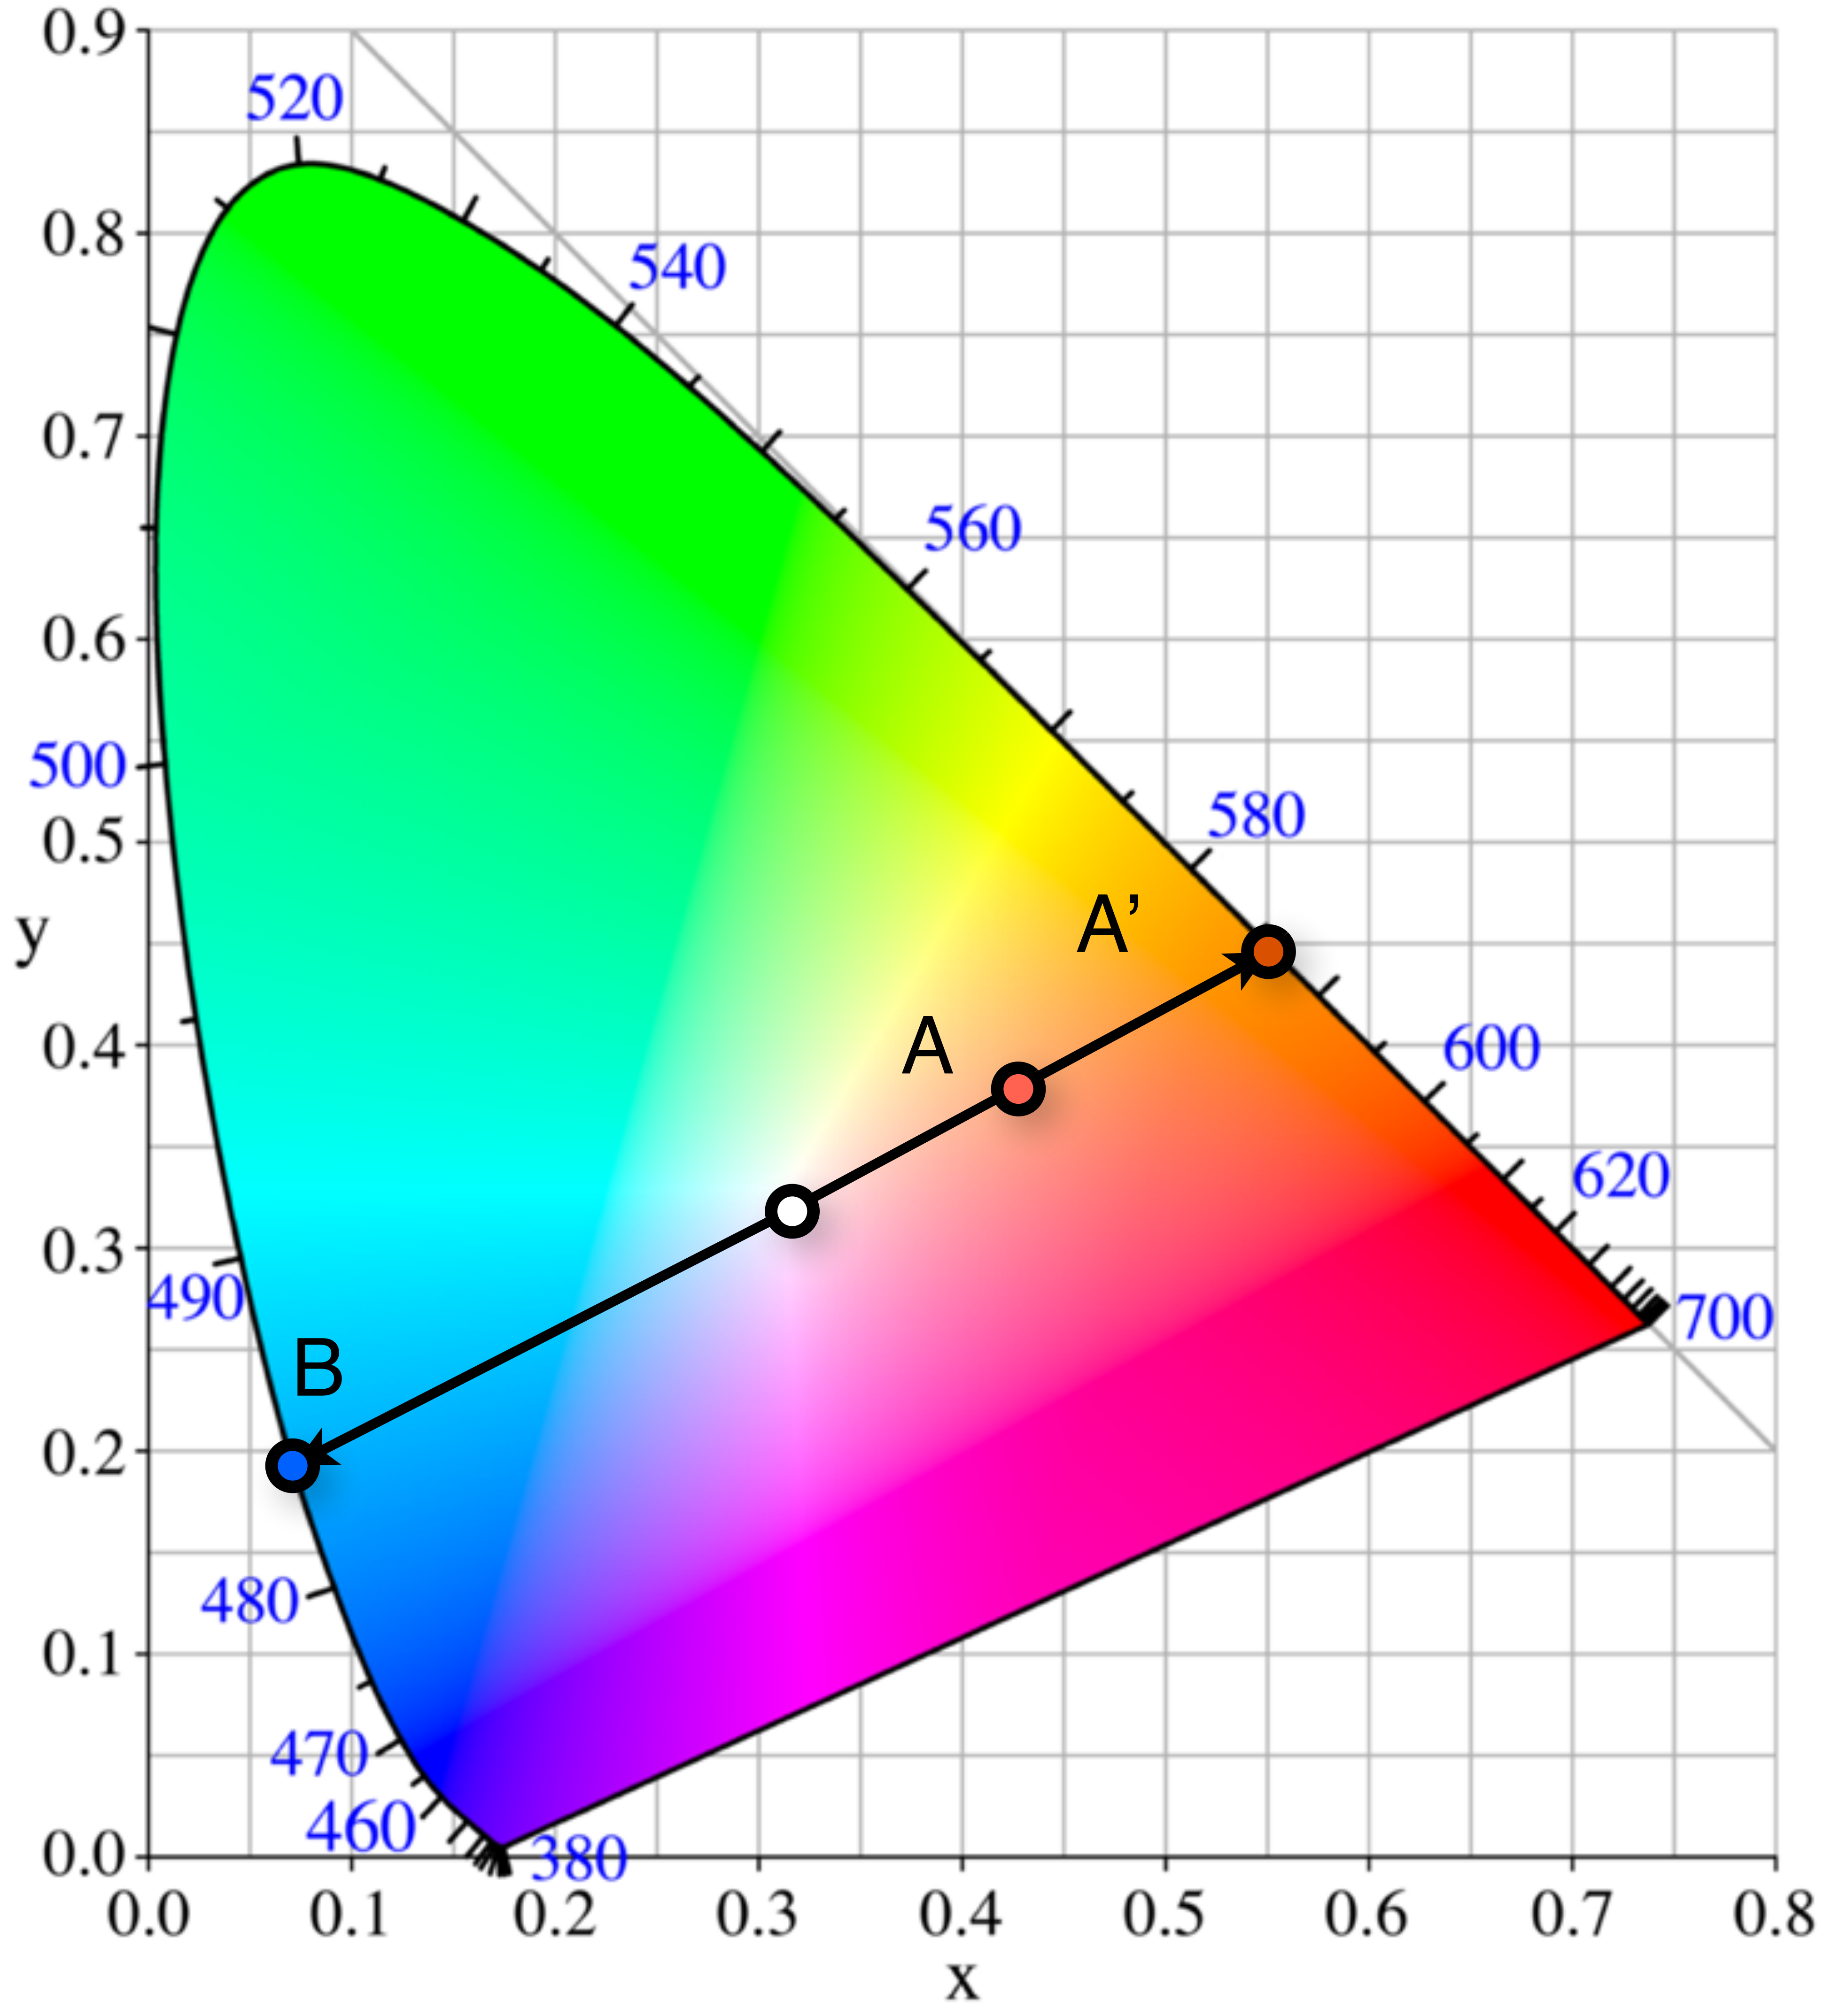

Perceptually-Uniform Color Spaces

All these color spaces so far are perceptually non-uniform:

- two colors close together in space are not necessarily visually similar
- two colors far apart are not necessarily very different!

Measuring "perceptual distance" in color spaces important for many industries

Experiments by MacAdams

MacAdams Color Ellipses

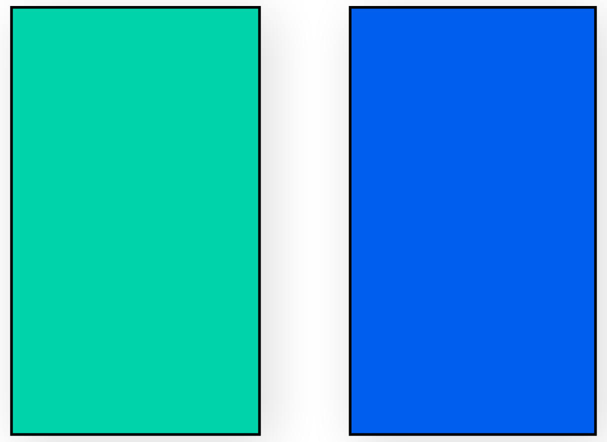

Test patches

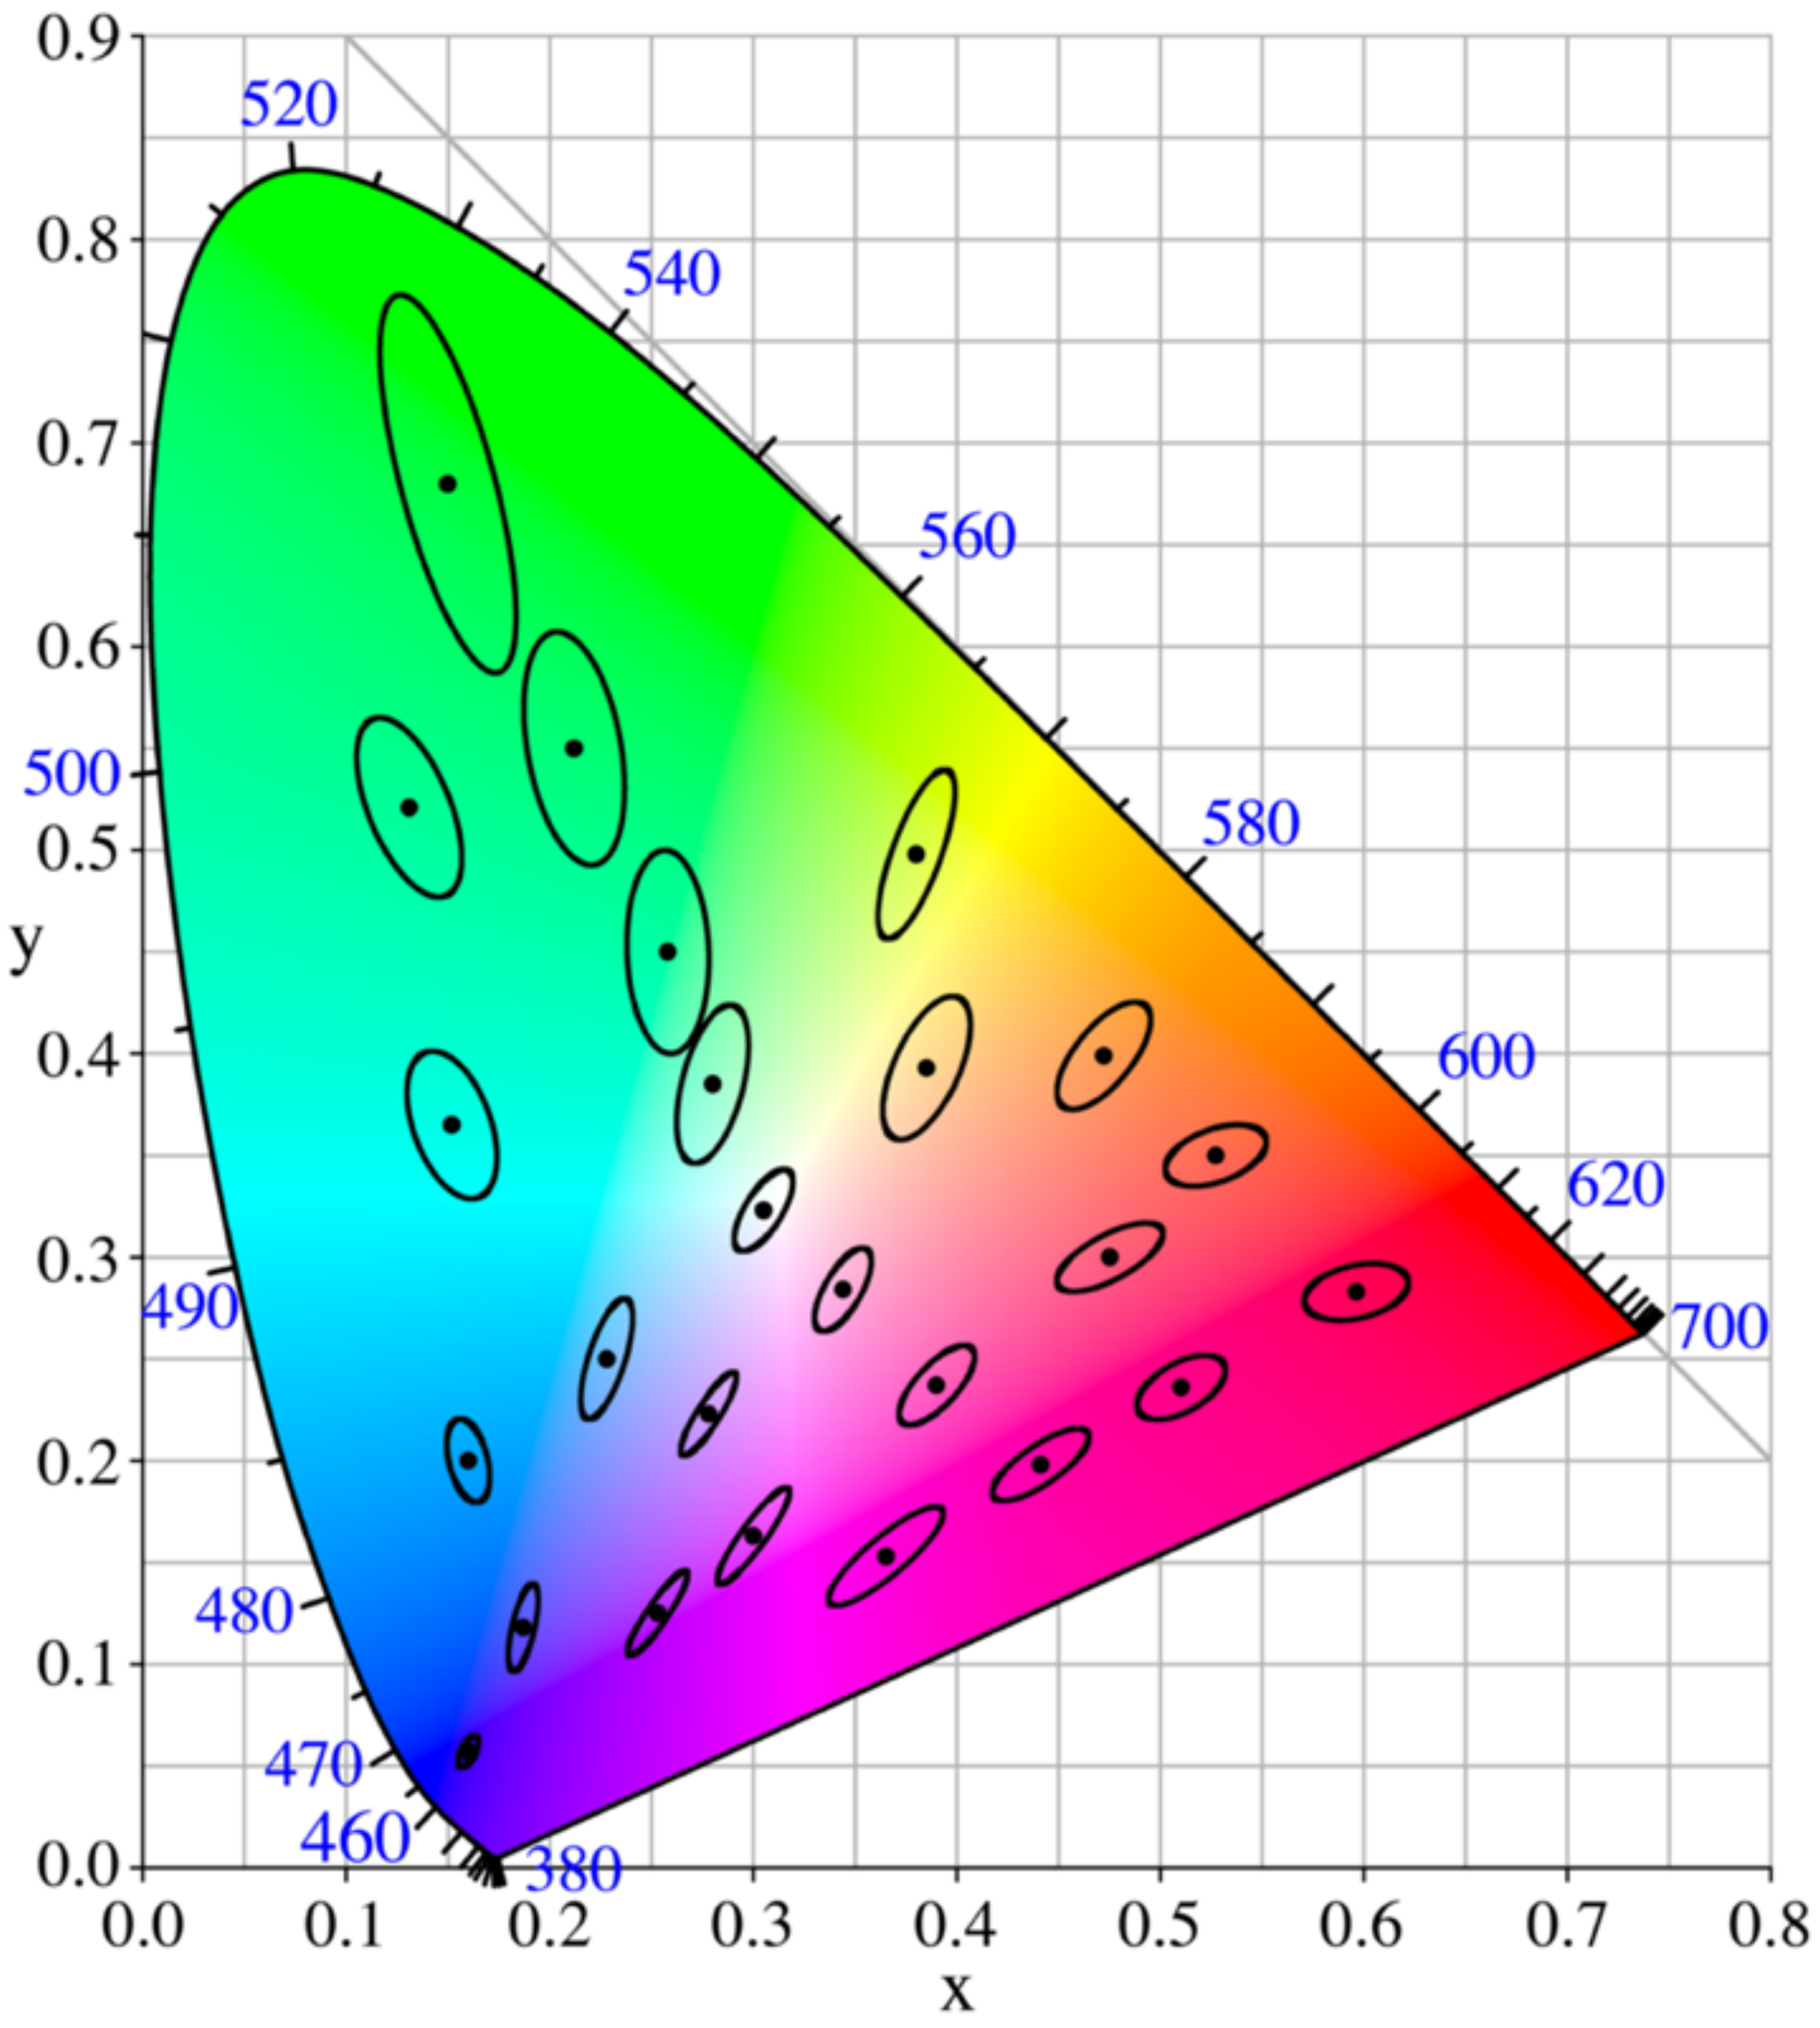

CIELab and CIELuv Color Spaces

Two attempts to make a perceptually-uniform color space

MacAdams ellipses become nearly (but not perfectly) circular

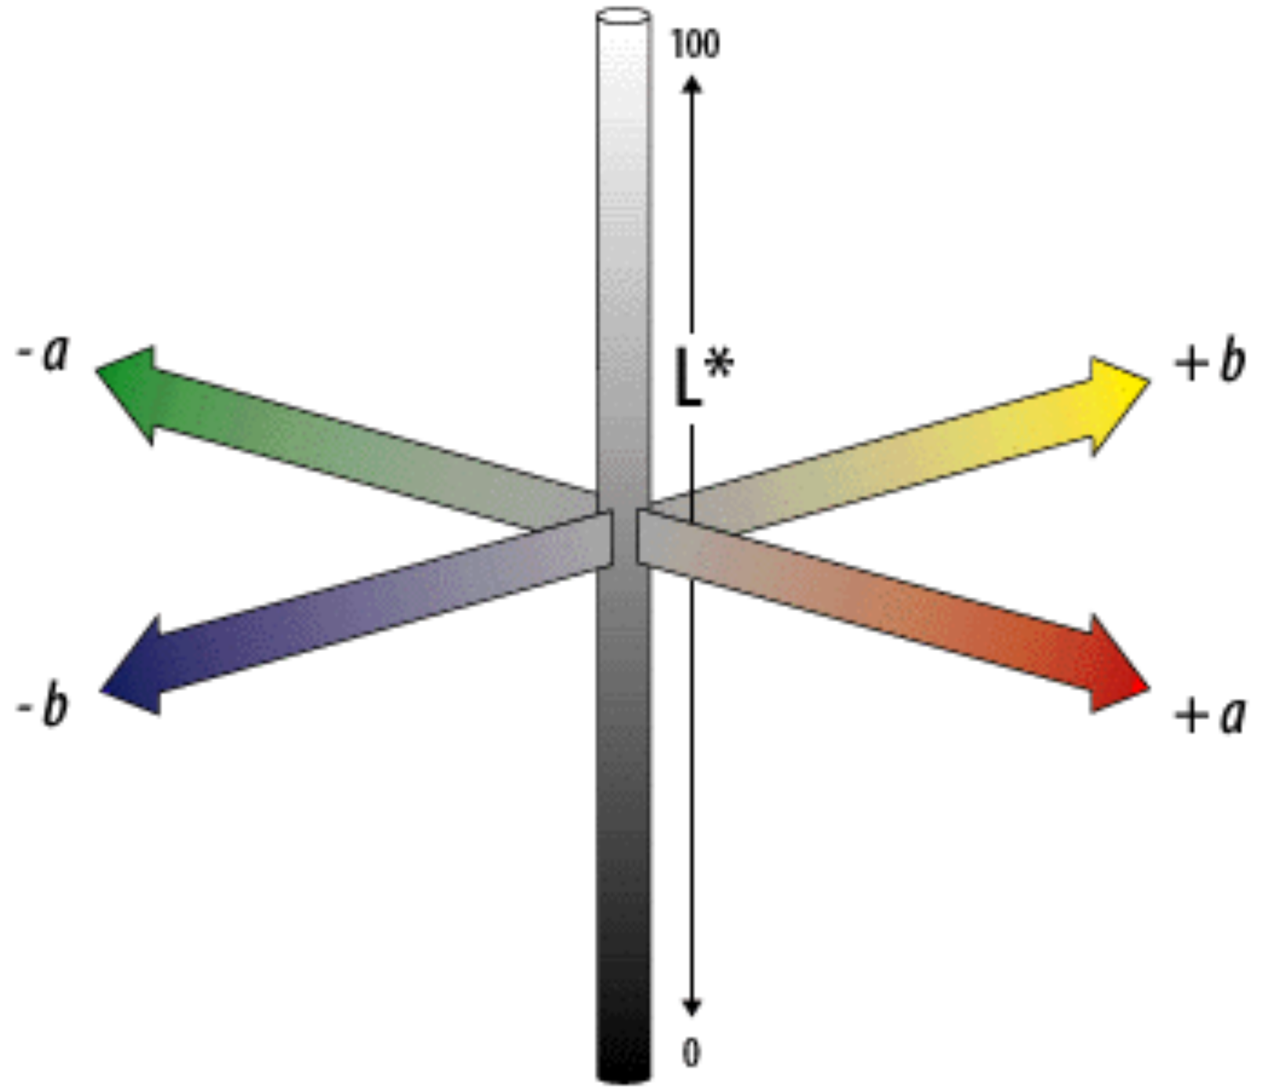

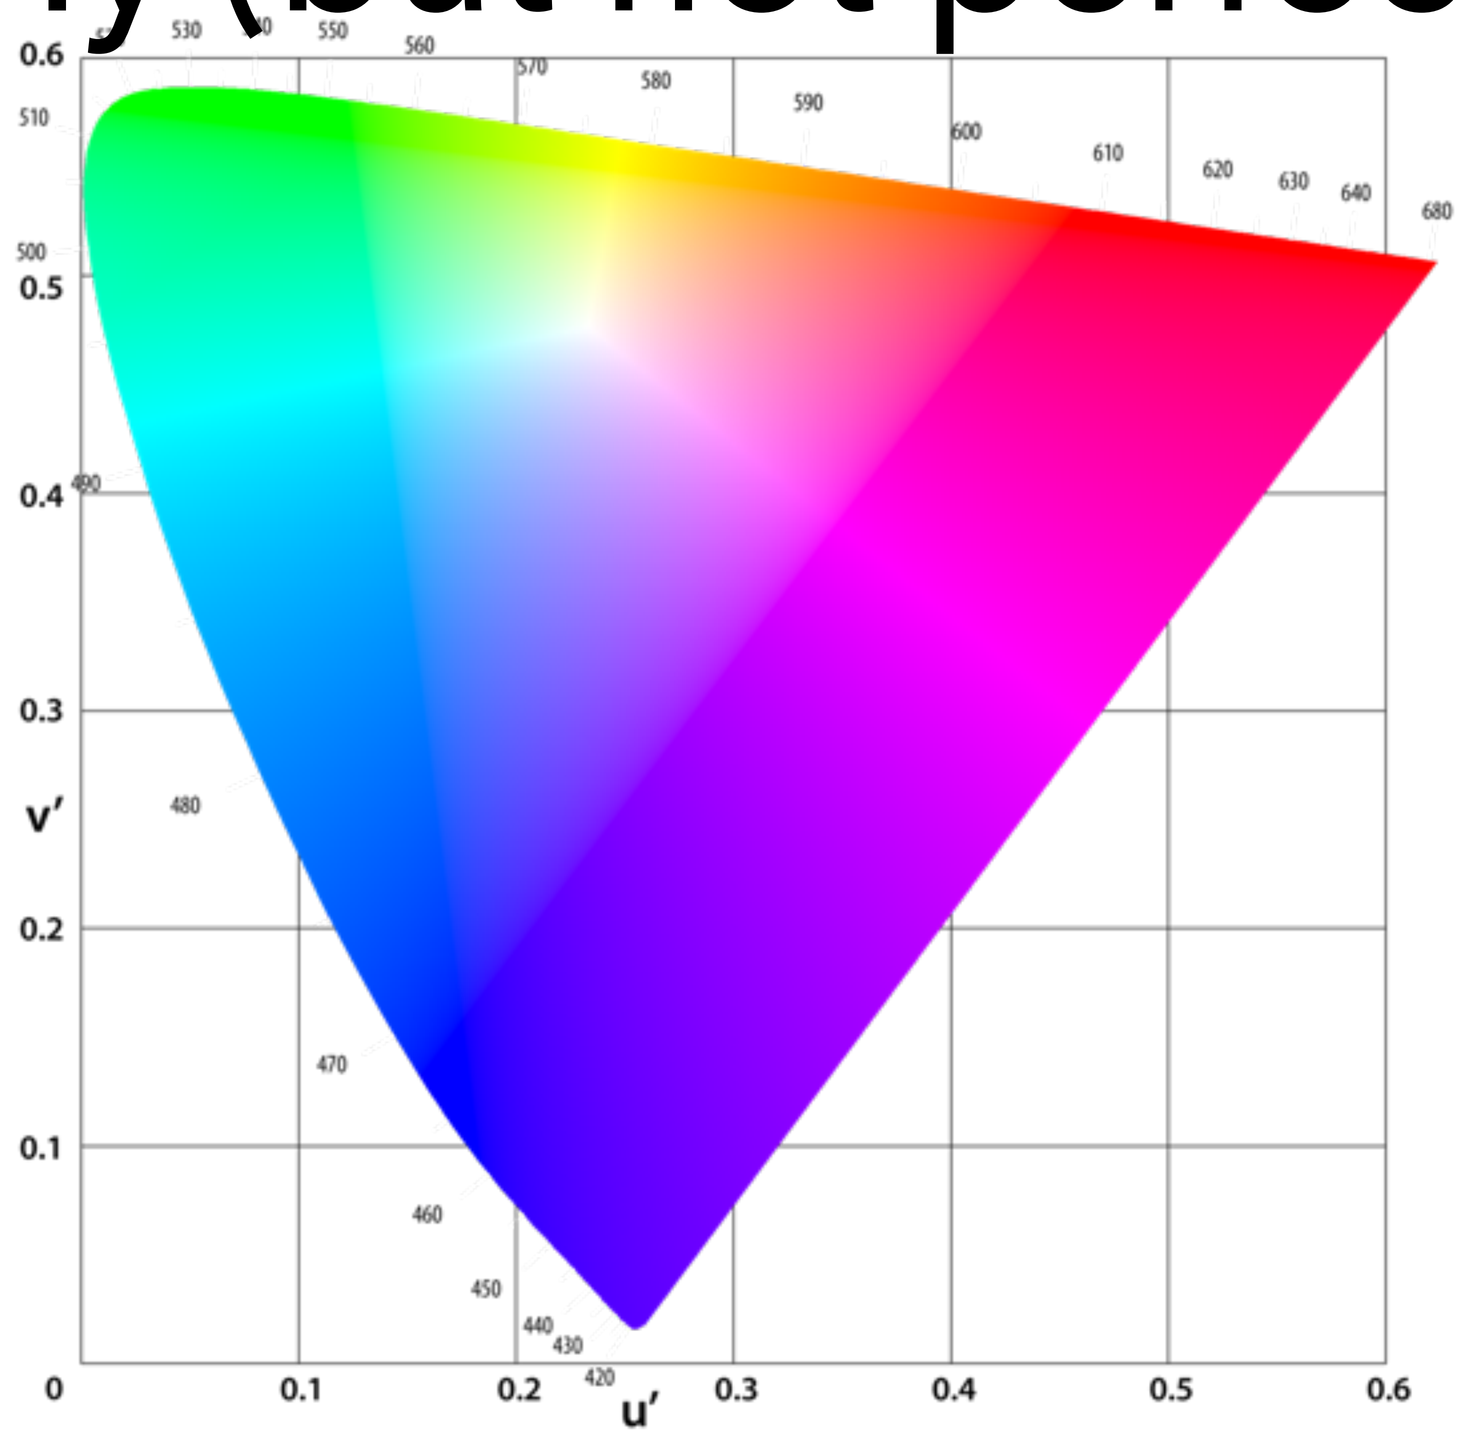

Higher-level color perception

Higher-level color perception

Color perception is much more complicated than response of SML cones…

Visual pathway

- A lot happens after the cones
- But: cone responses are input to further processing

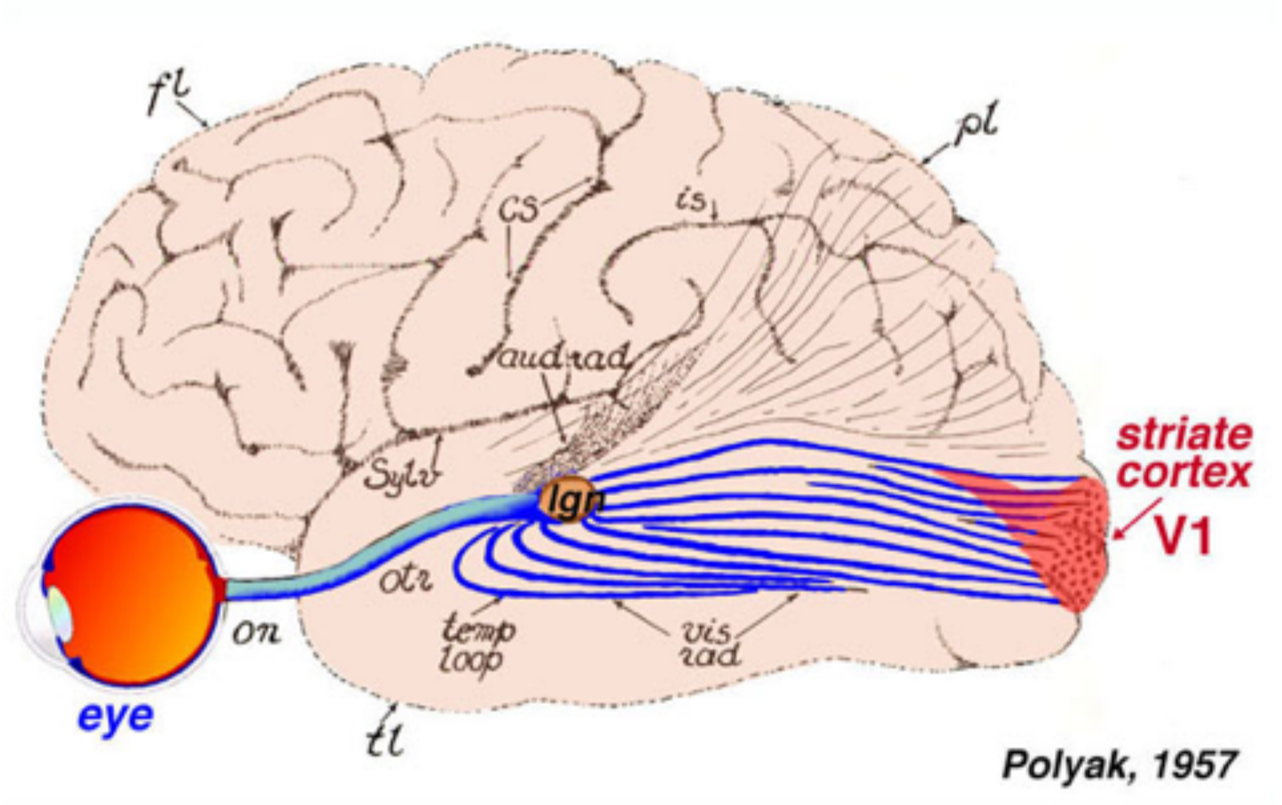

Color constancy

Also known as chromatic adaptation

Color of object is perceived as the same even under varying illumination

For example:

- A white sheet of paper under green illumination is still perceived as white, even though the reflected light is green! The human brain infers the white color from the context, which is “green-ish“ too because of the green illumination.

Color constancy

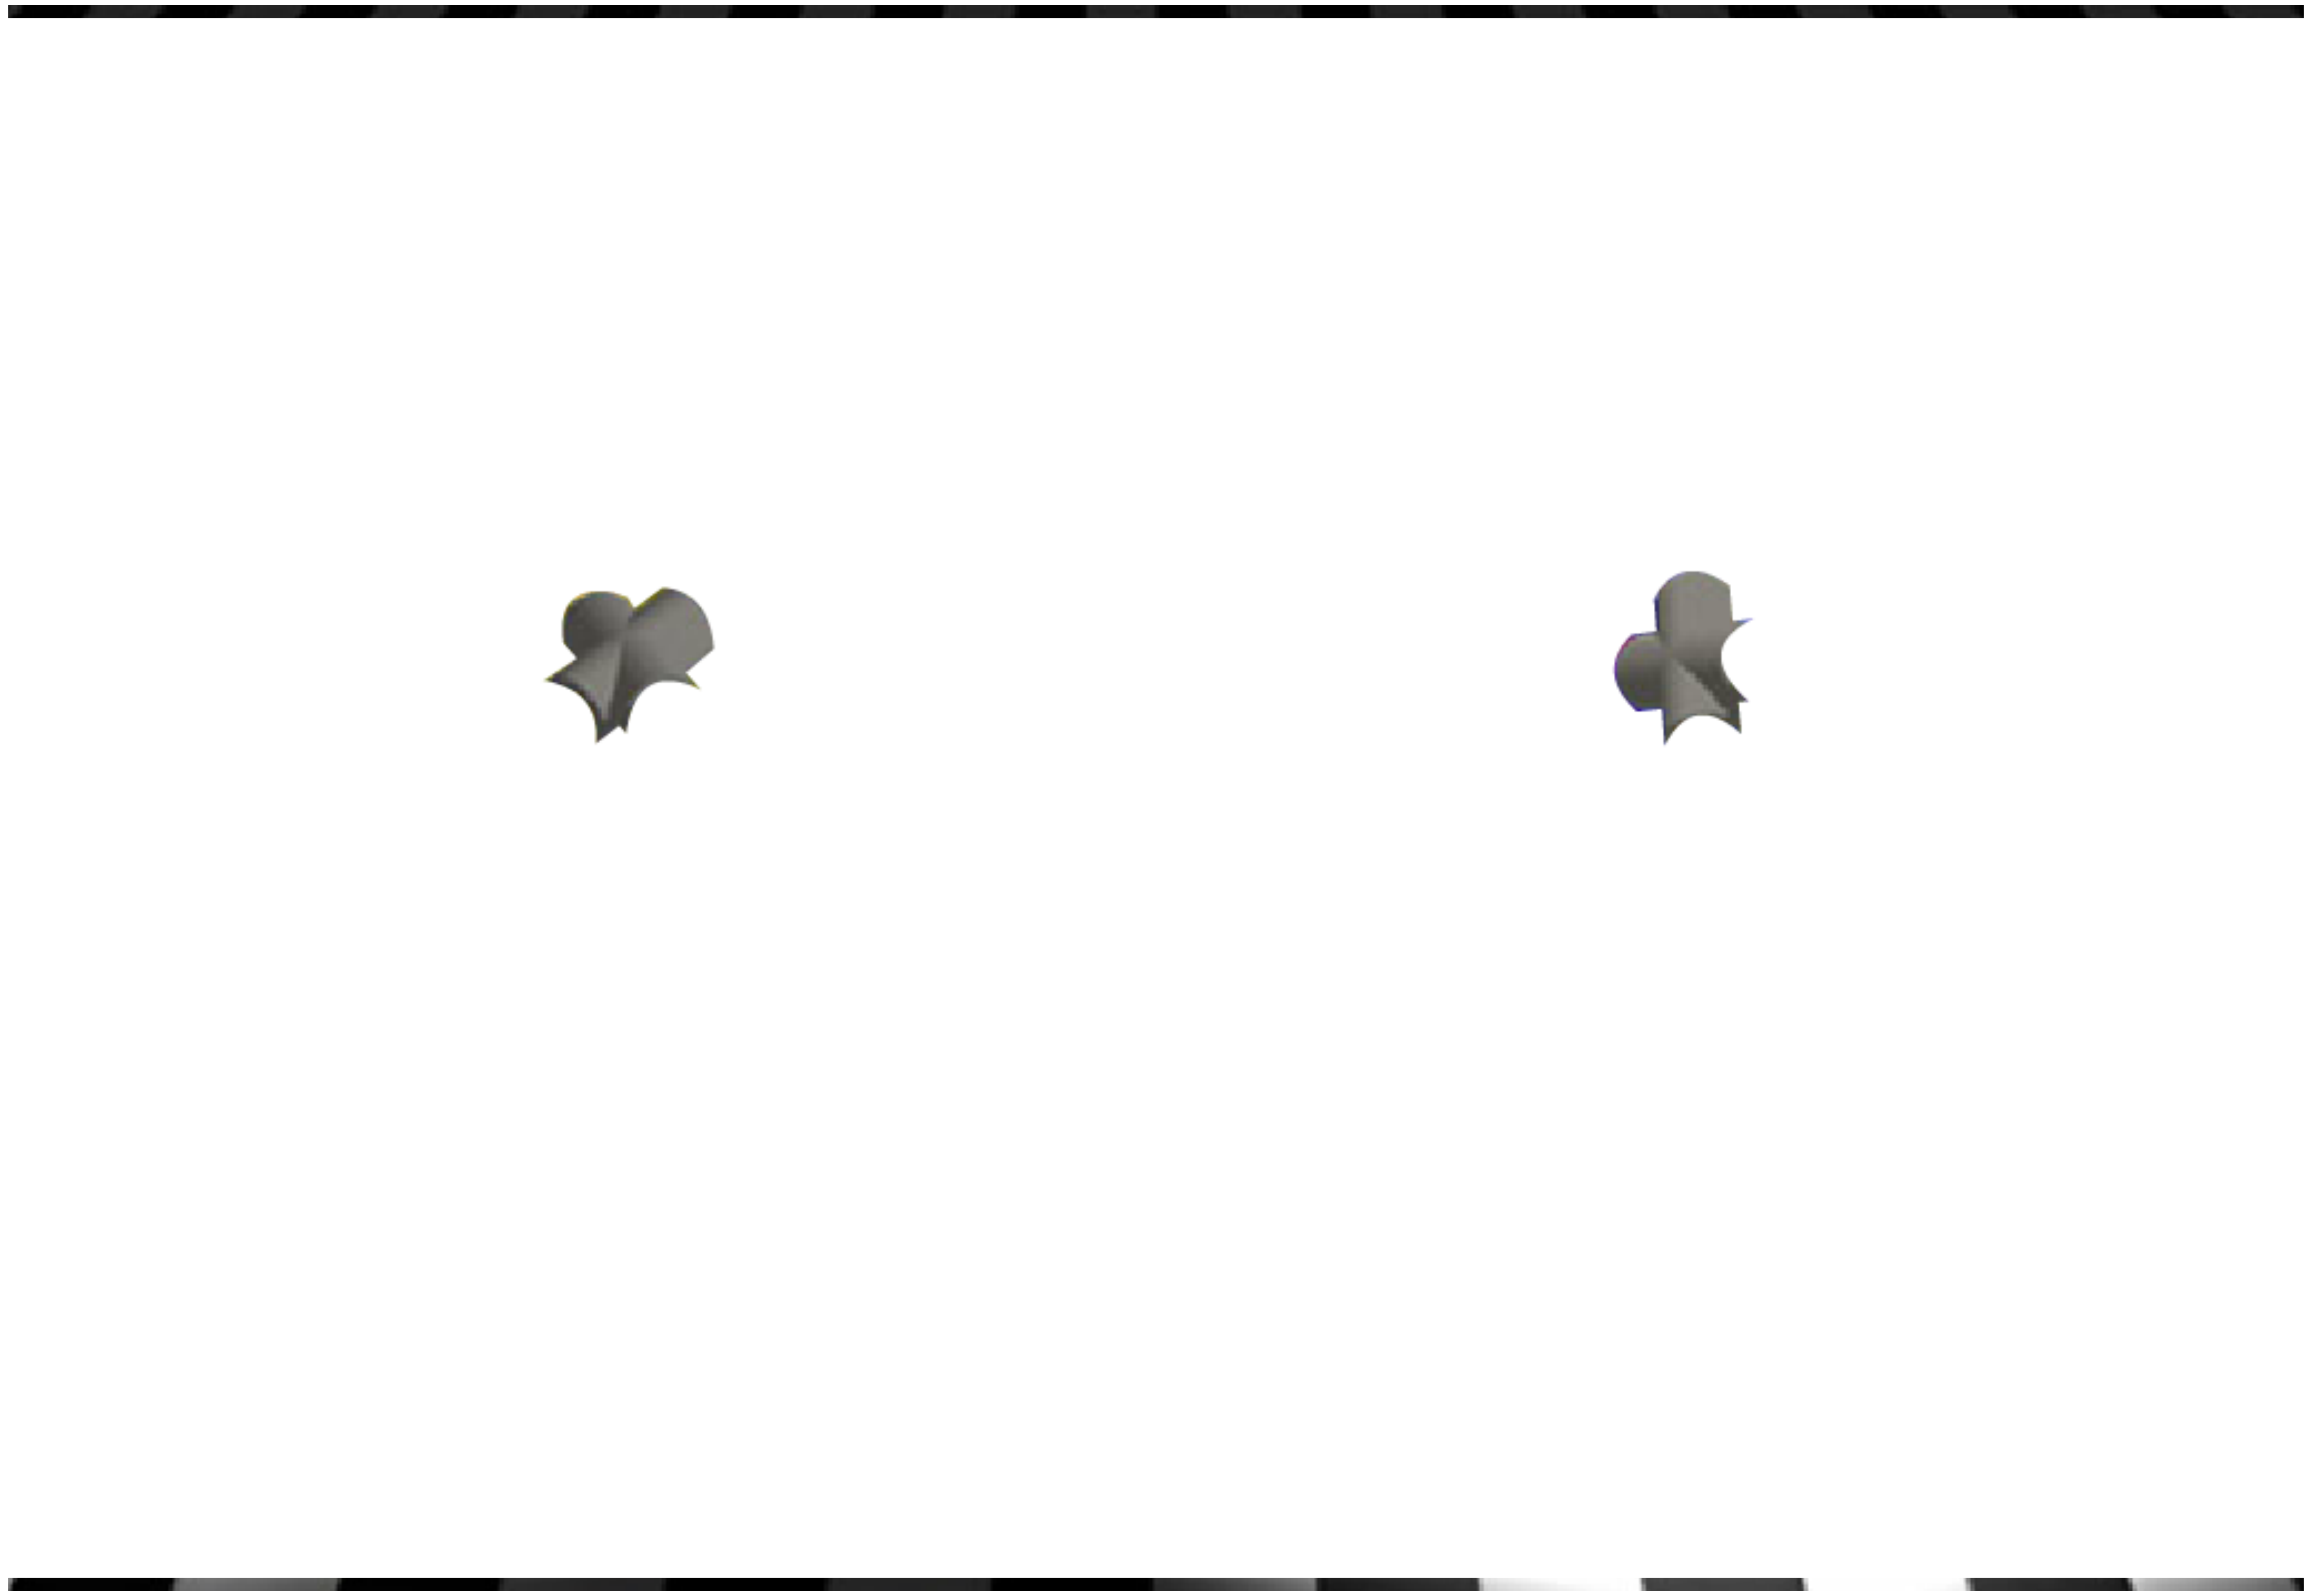

blue and black?

or

white and gold?

Color constancy failure

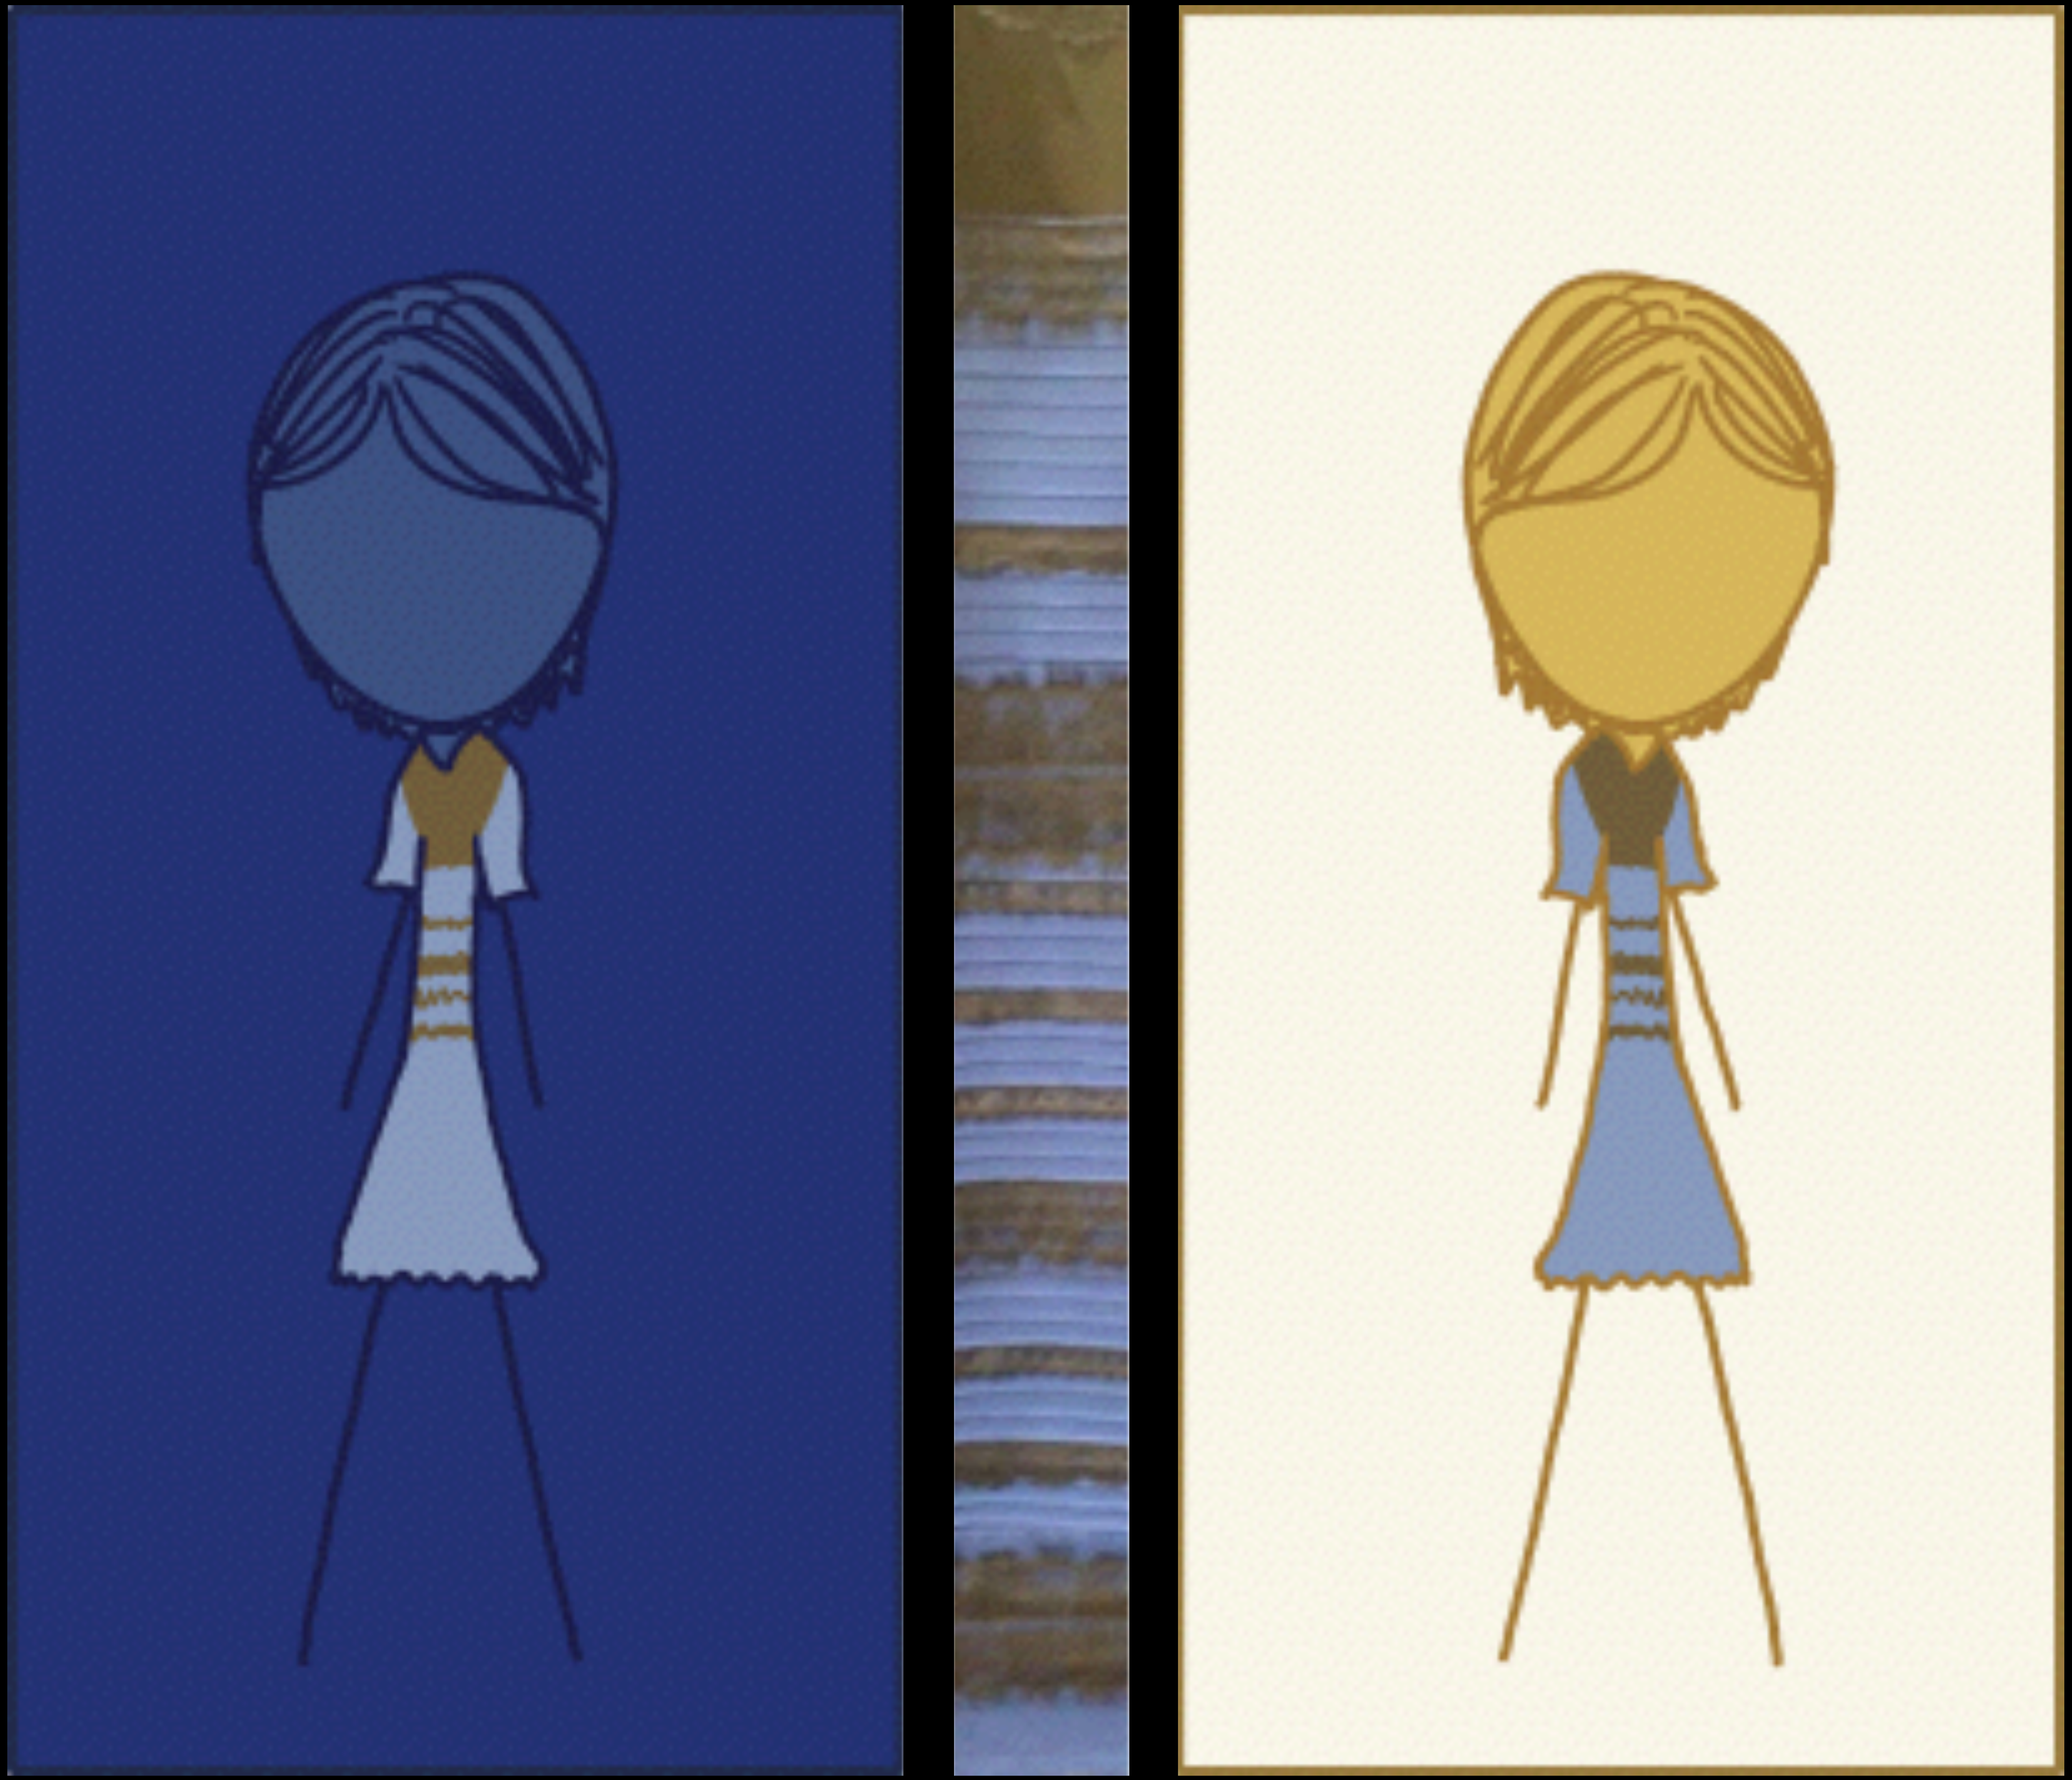

http://xkcd.com/1492/

Hering's opponent process theory (1874)

After sensing by cones, colors are encoded as red versus green, blue versus yellow, and black versus white

Physiological evidence found in the 1950s

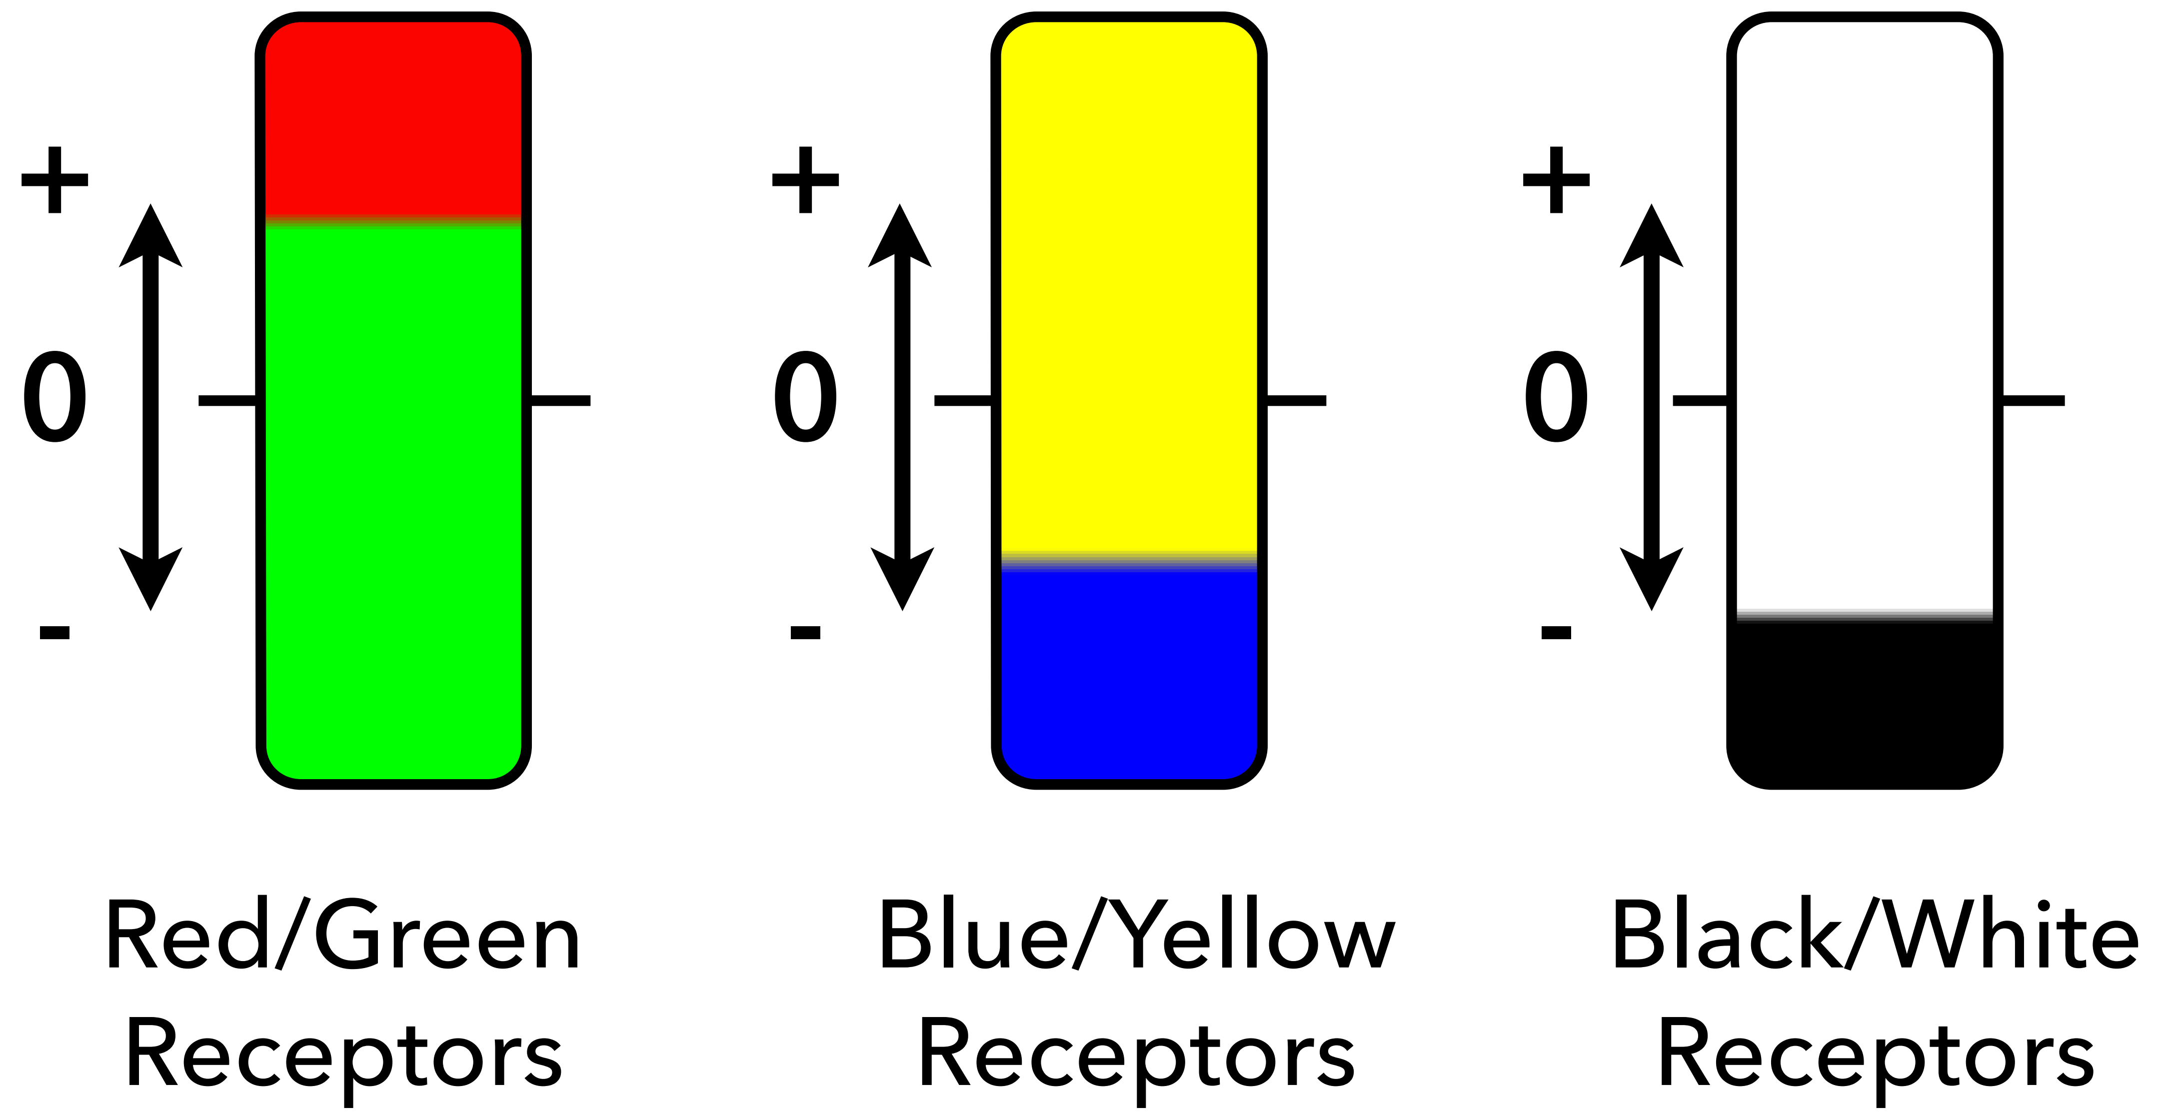

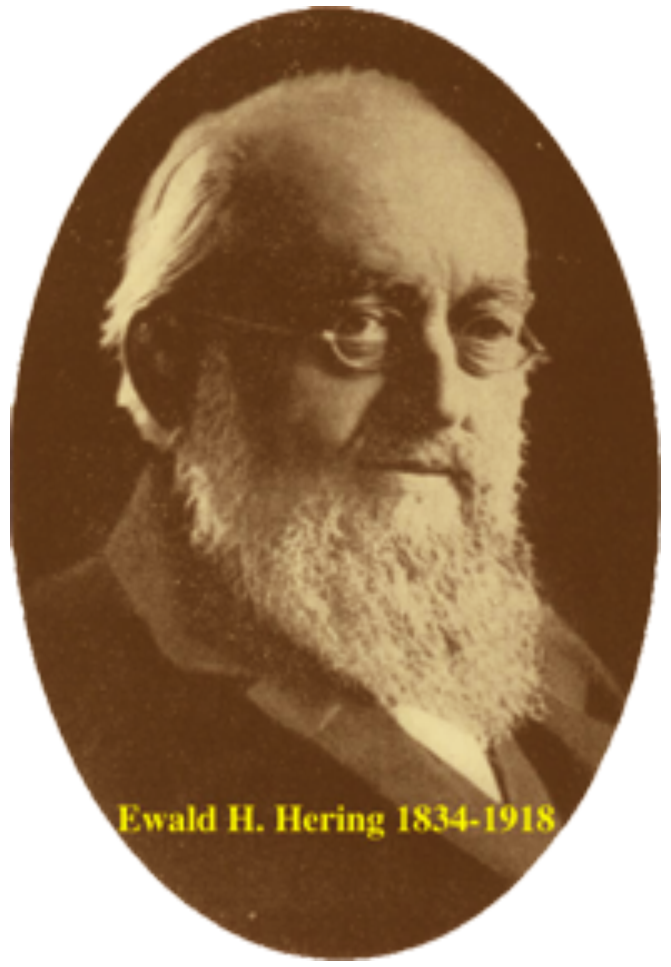

Dual process theory

Inputs are LMS cone responses

Output has a different parameterization:

- Light-dark
- Blue-yellow
- Red-green

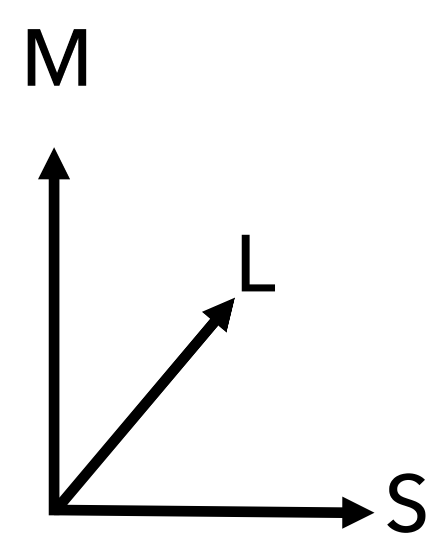

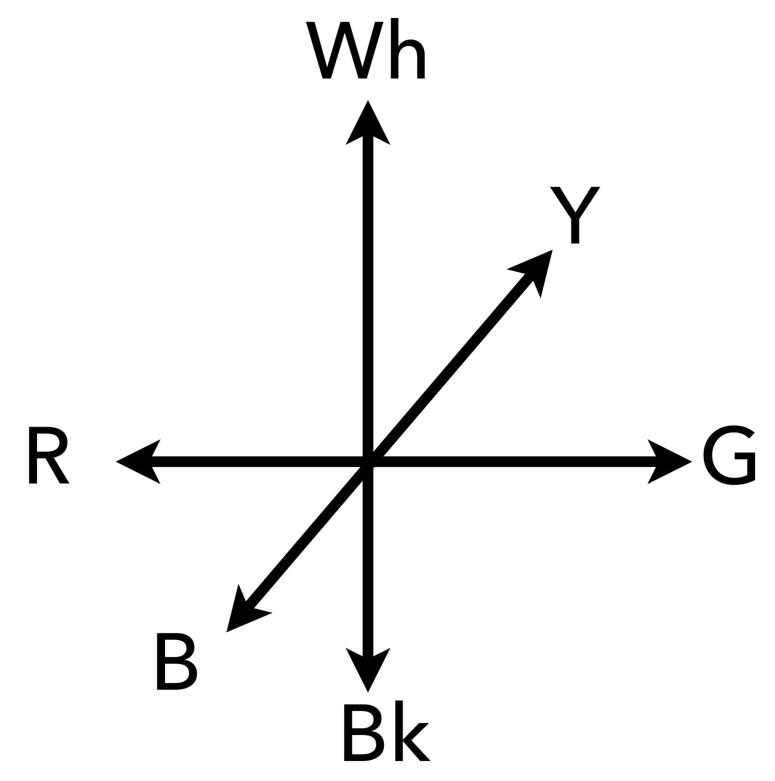

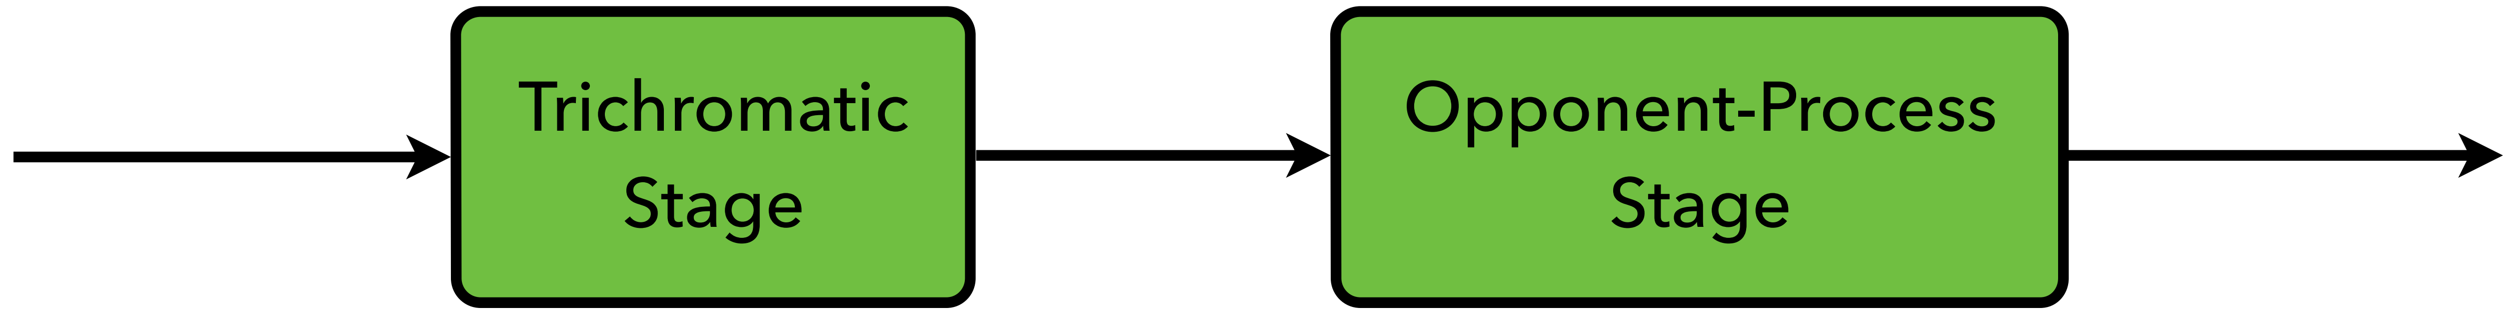

Color opponents wiring

Sums for brightness

Differences for color opponents

At the end, it's just a 3x3 matrix compared to LMS

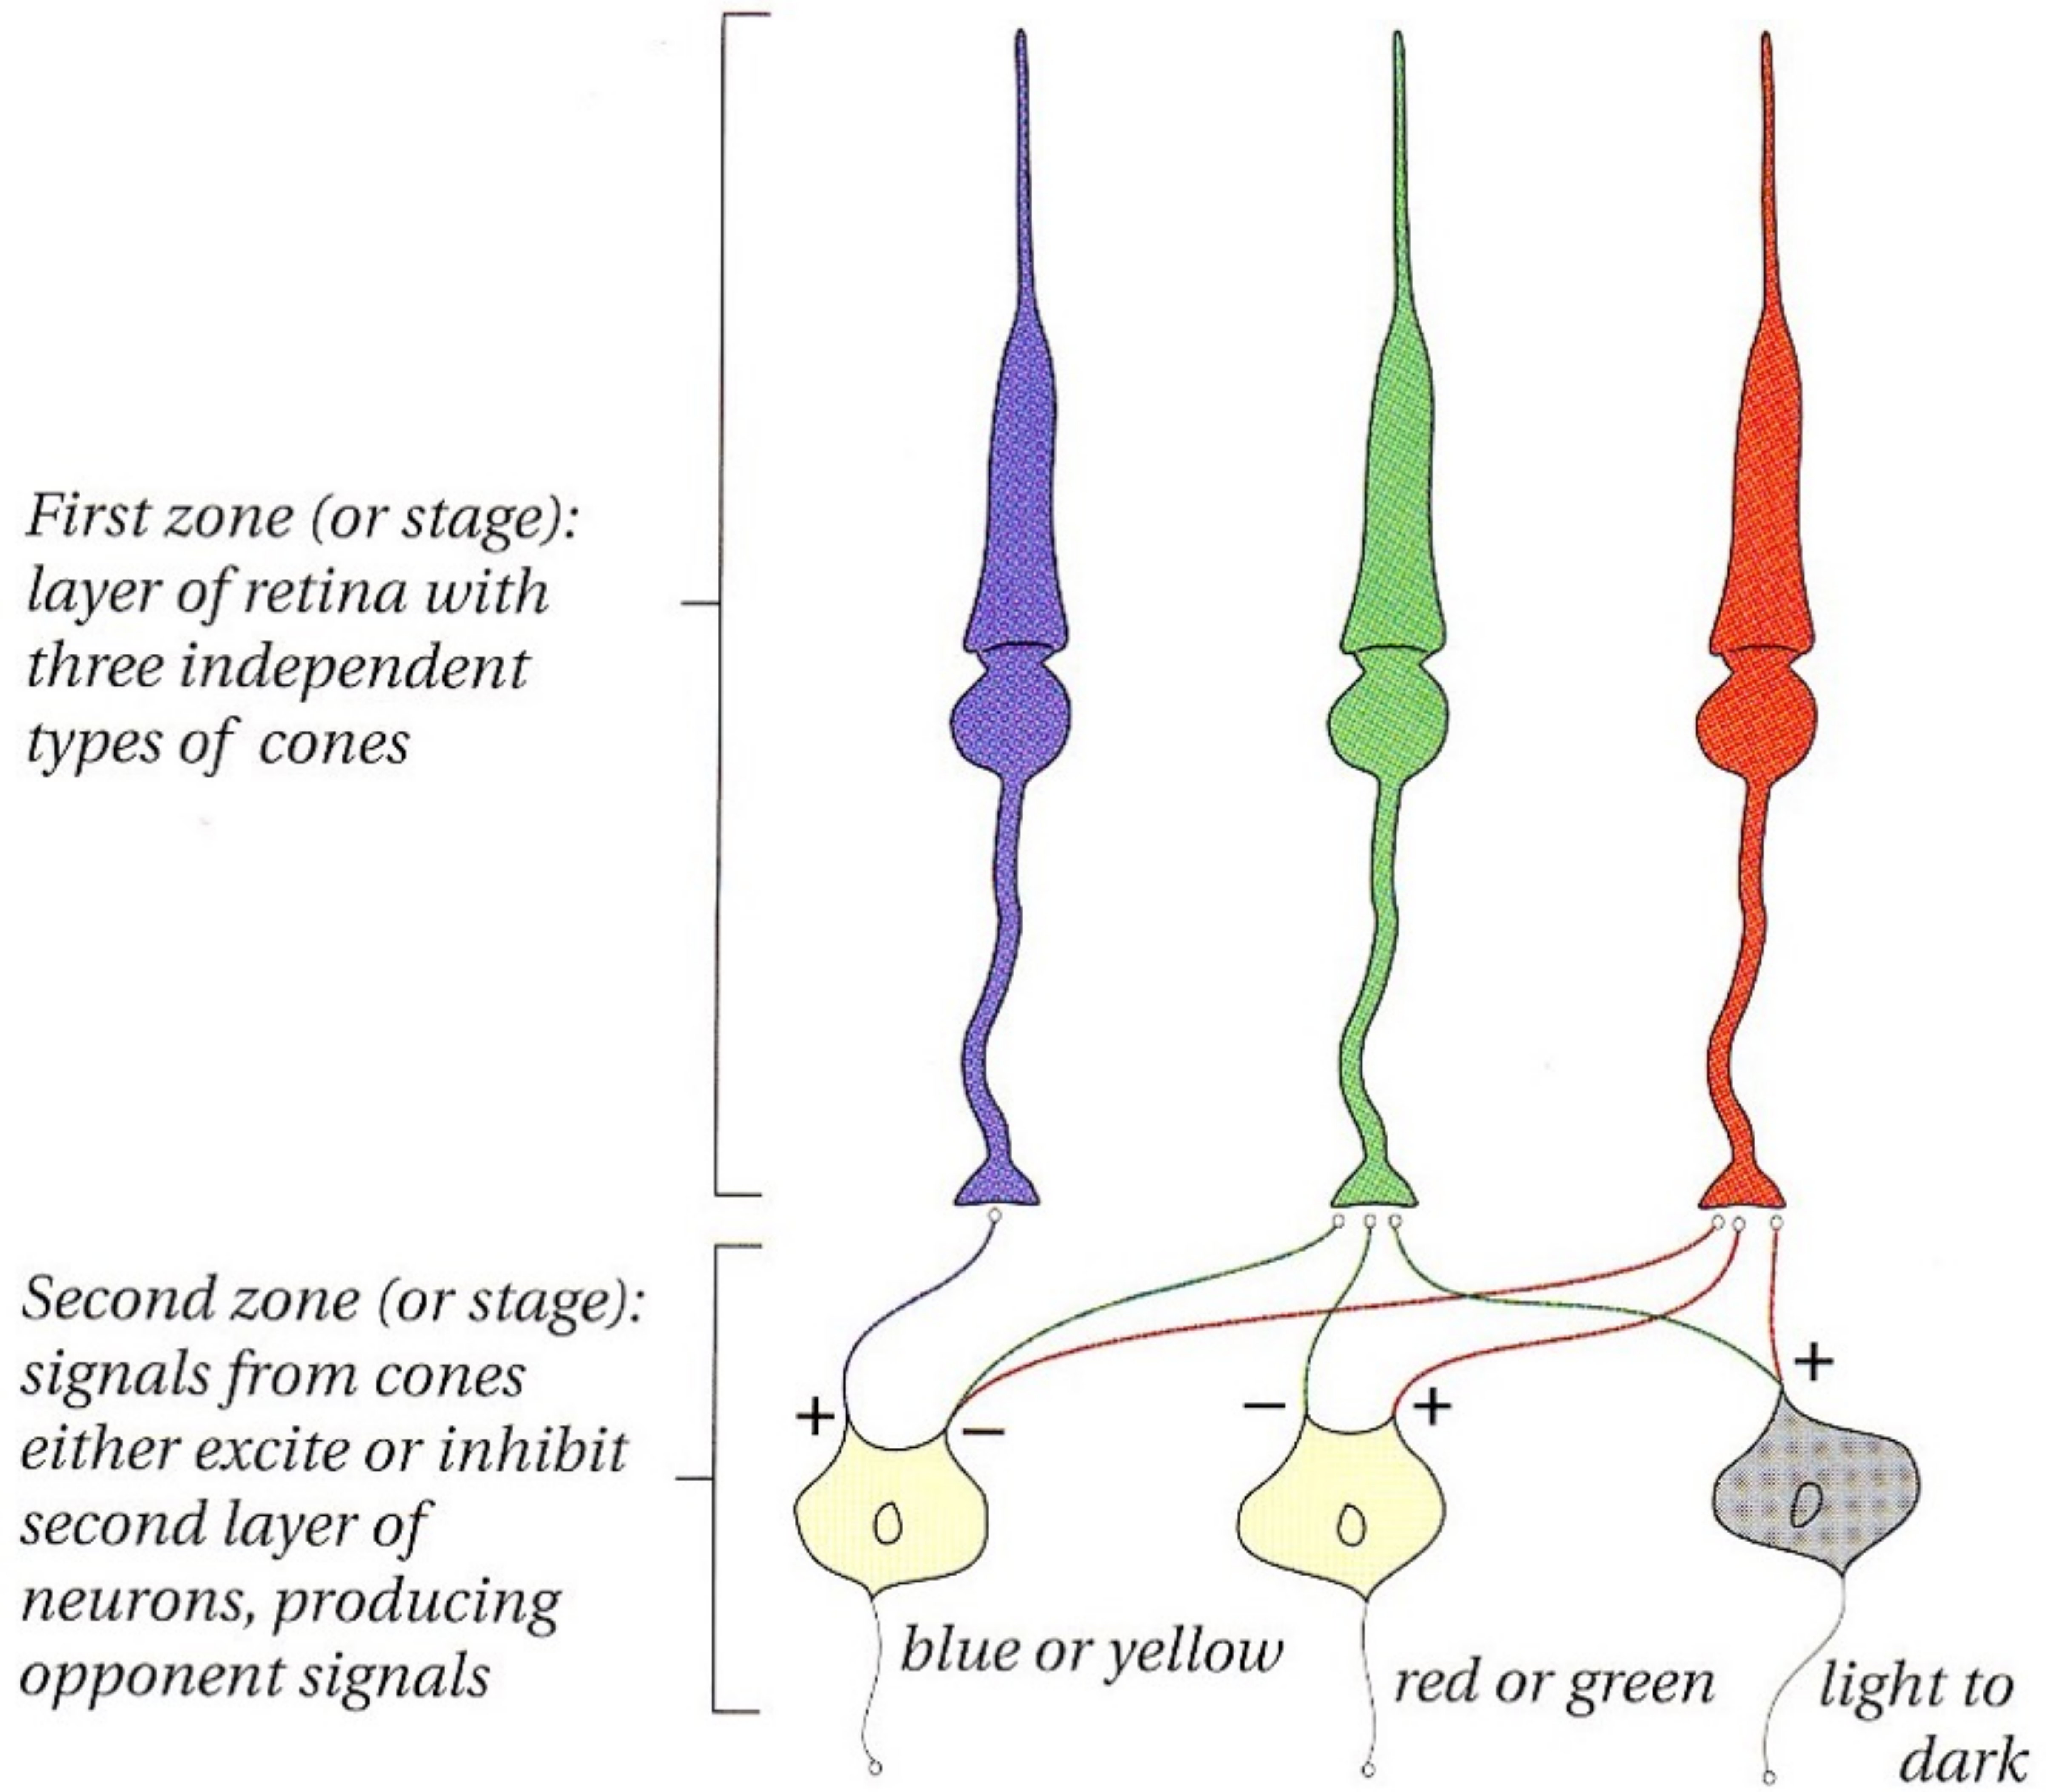

Opponent Colors

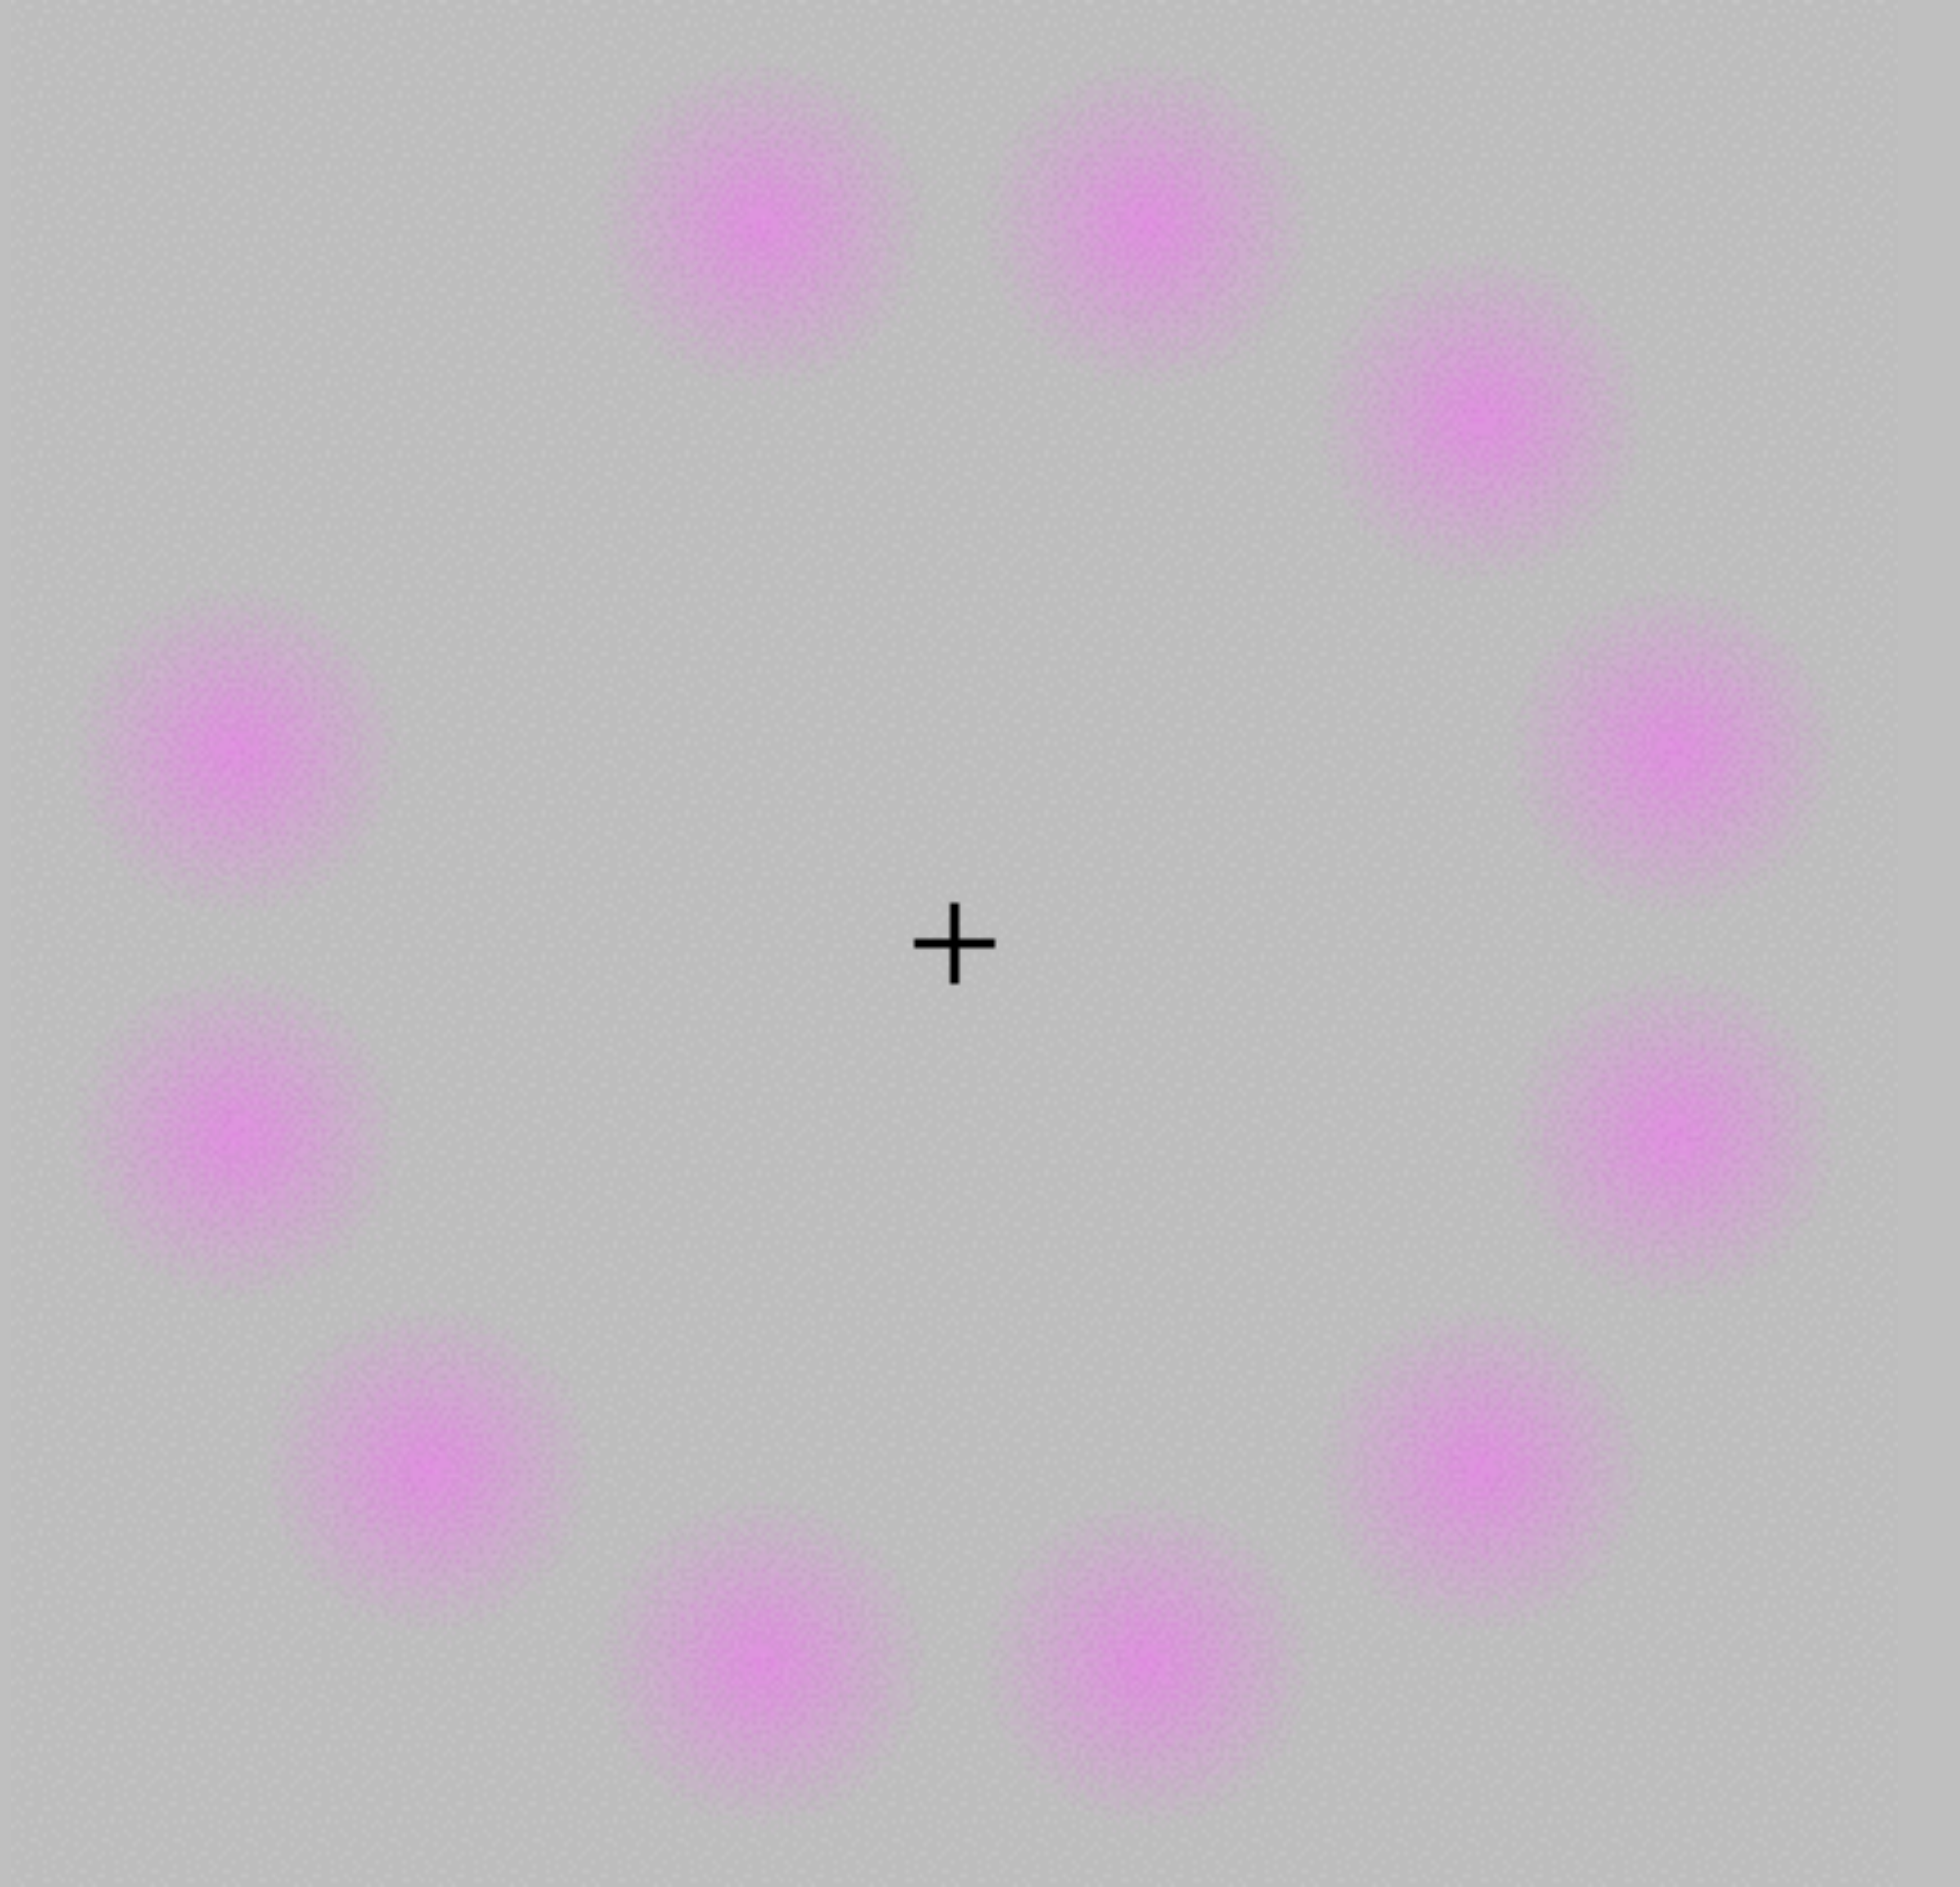

Fixate on the + for >30s. Blinking is OK.

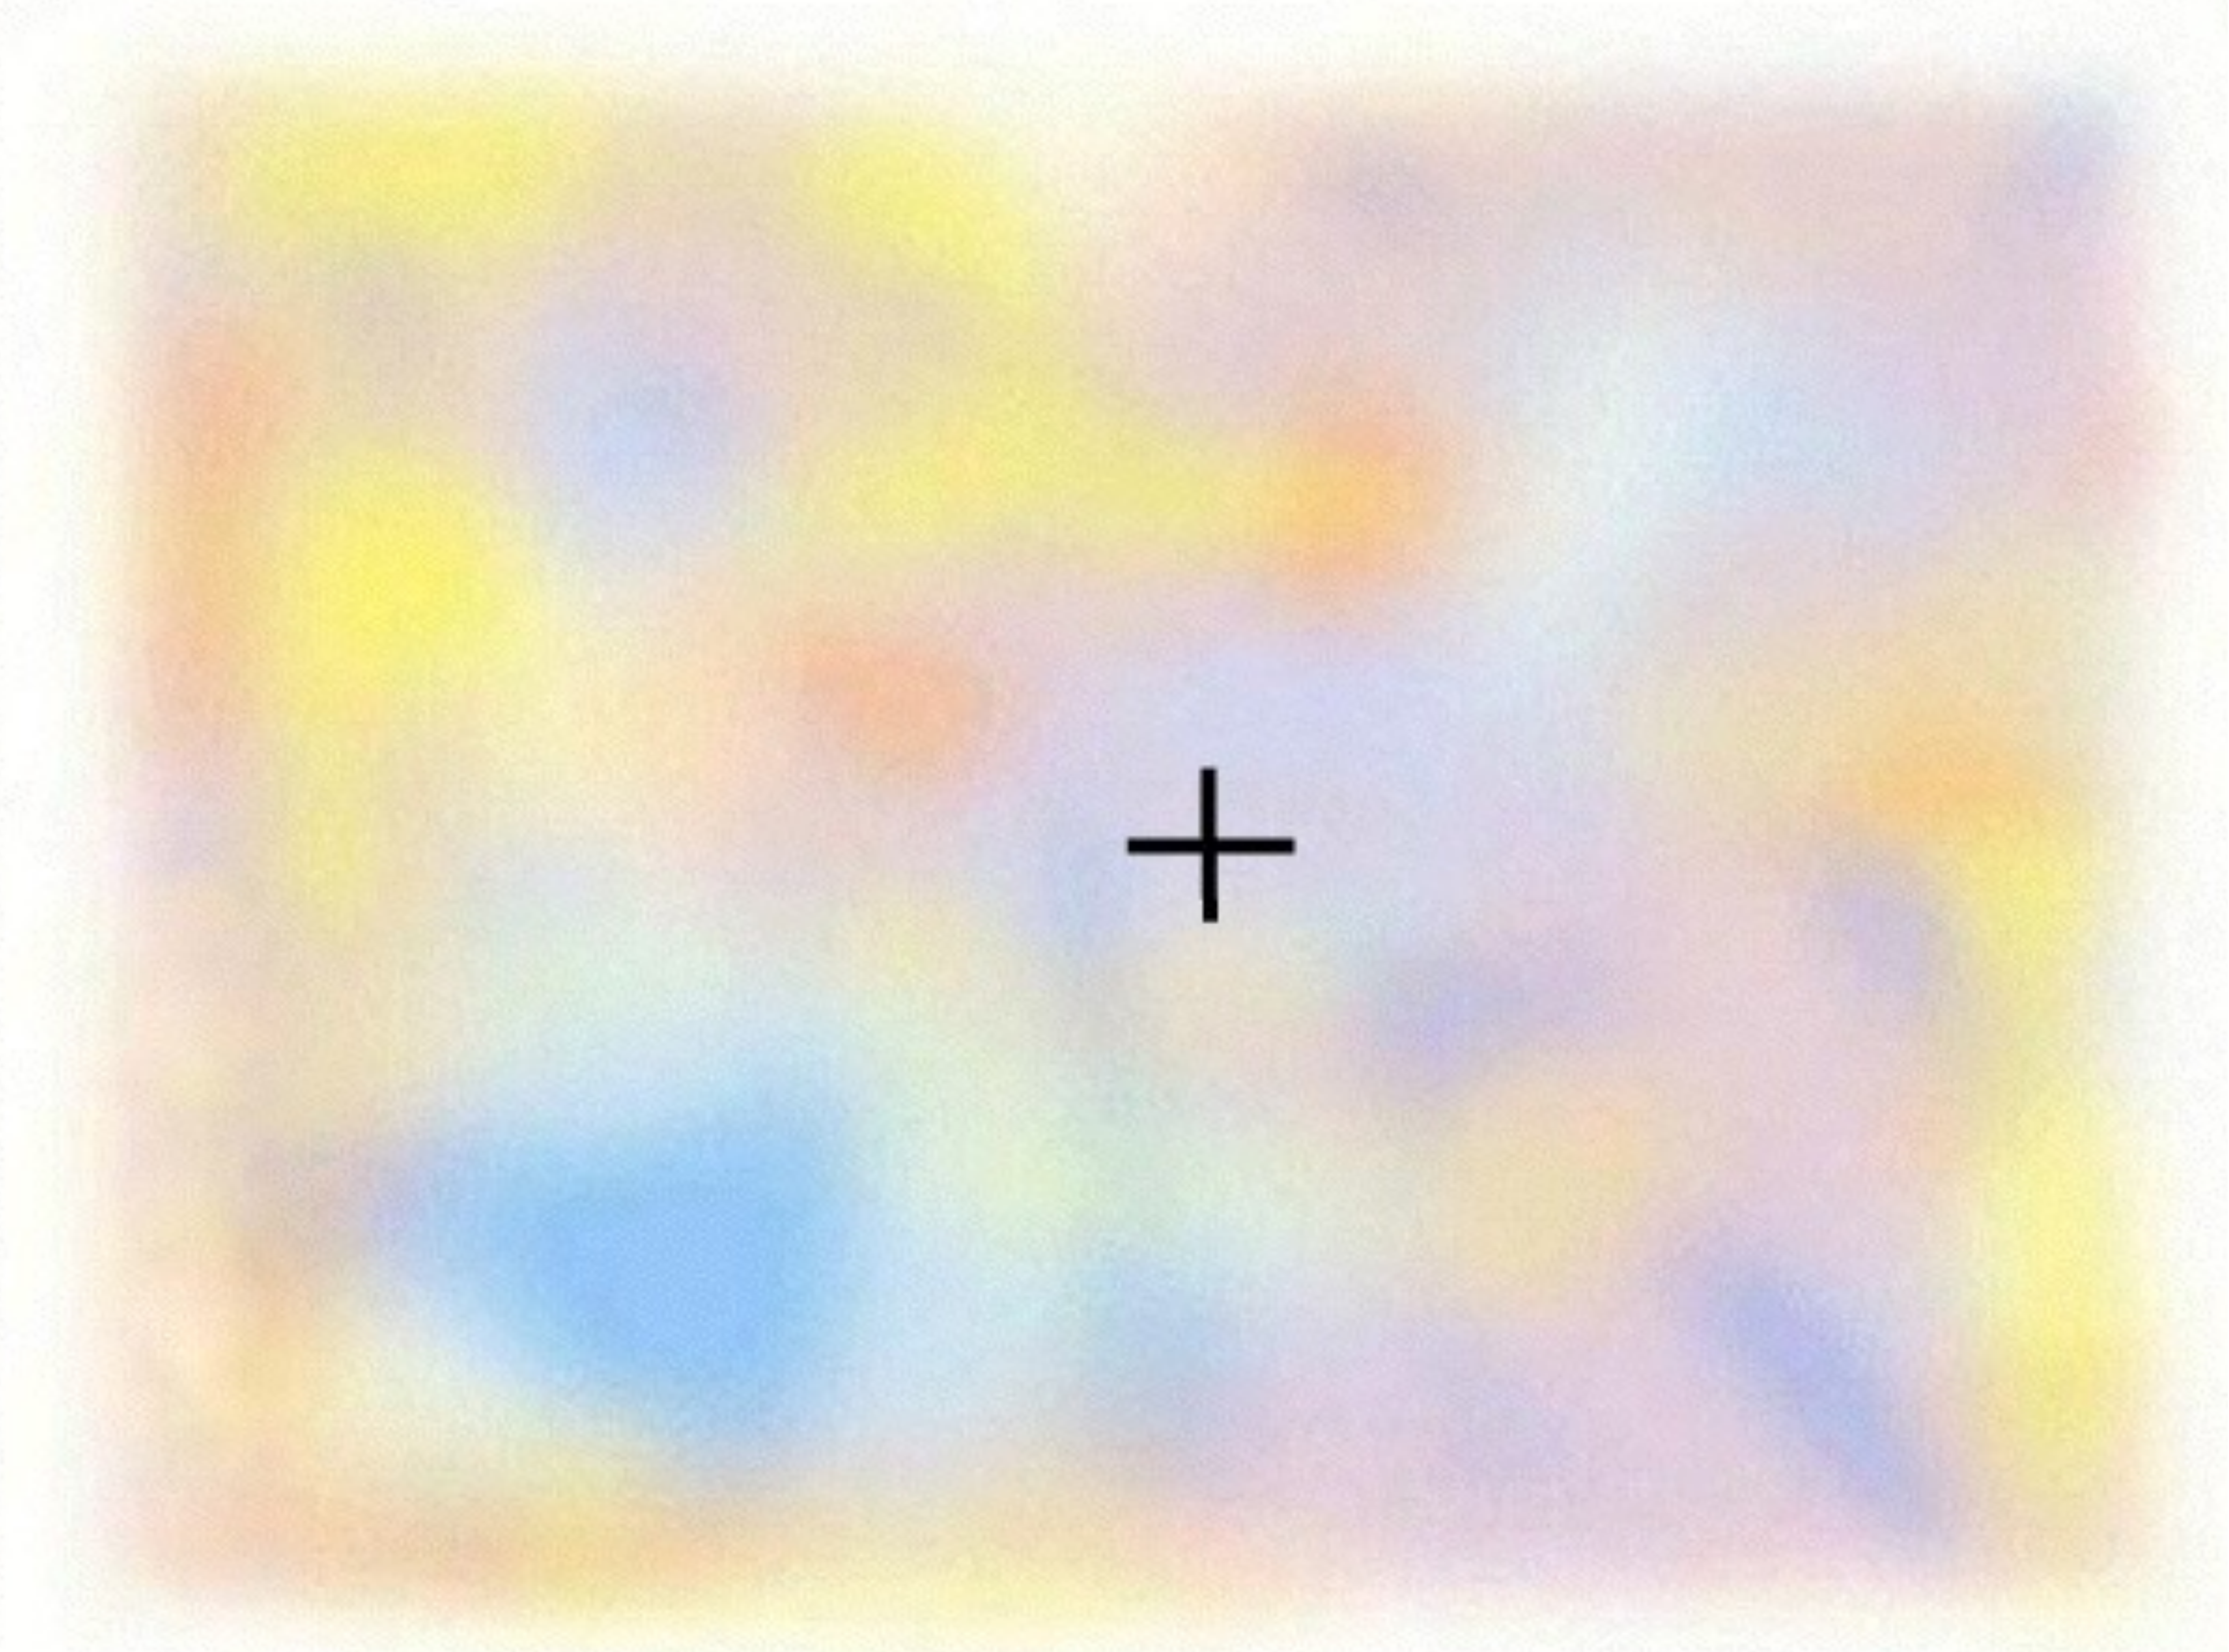

When you eventually move your eyes the image will suddenly reappear or become more vibrant if it didn't disappear.


© Wojciech Jarosz

Opponent color spaces

Luminance, red-green, blue-yellow

CIELab

YUV

$$\begin{bmatrix} Y' \\ U \\ V \end{bmatrix} = \begin{bmatrix} 0.299 & 0.587 & 0.114 \\ -0.14713 & -0.28886 & 0.436 \\ 0.615 & -0.51499 & -0.10001 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix}$$

$$\begin{bmatrix} R \\ G \\ B \end{bmatrix} = \begin{bmatrix} 1 & 0 & 1.13983 \\ 1 & -0.39465 & -0.58060 \\ 1 & 2.03211 & 0 \end{bmatrix} \begin{bmatrix} Y' \\ U \\ V \end{bmatrix}$$

YcrCb

- used a lot in image/video compression

$$\begin{aligned} Y' &= 0 + (0.299 \cdot R'_D) + (0.587 \cdot G'_D) + (0.114 \cdot B'_D) \\ C_B &= 128 - (0.168736 \cdot R'_D) - (0.331264 \cdot G'_D) + (0.5 \cdot B'_D) \\ C_R &= 128 + (0.5 \cdot R'_D) - (0.418688 \cdot G'_D) - (0.081312 \cdot B'_D) \end{aligned}$$

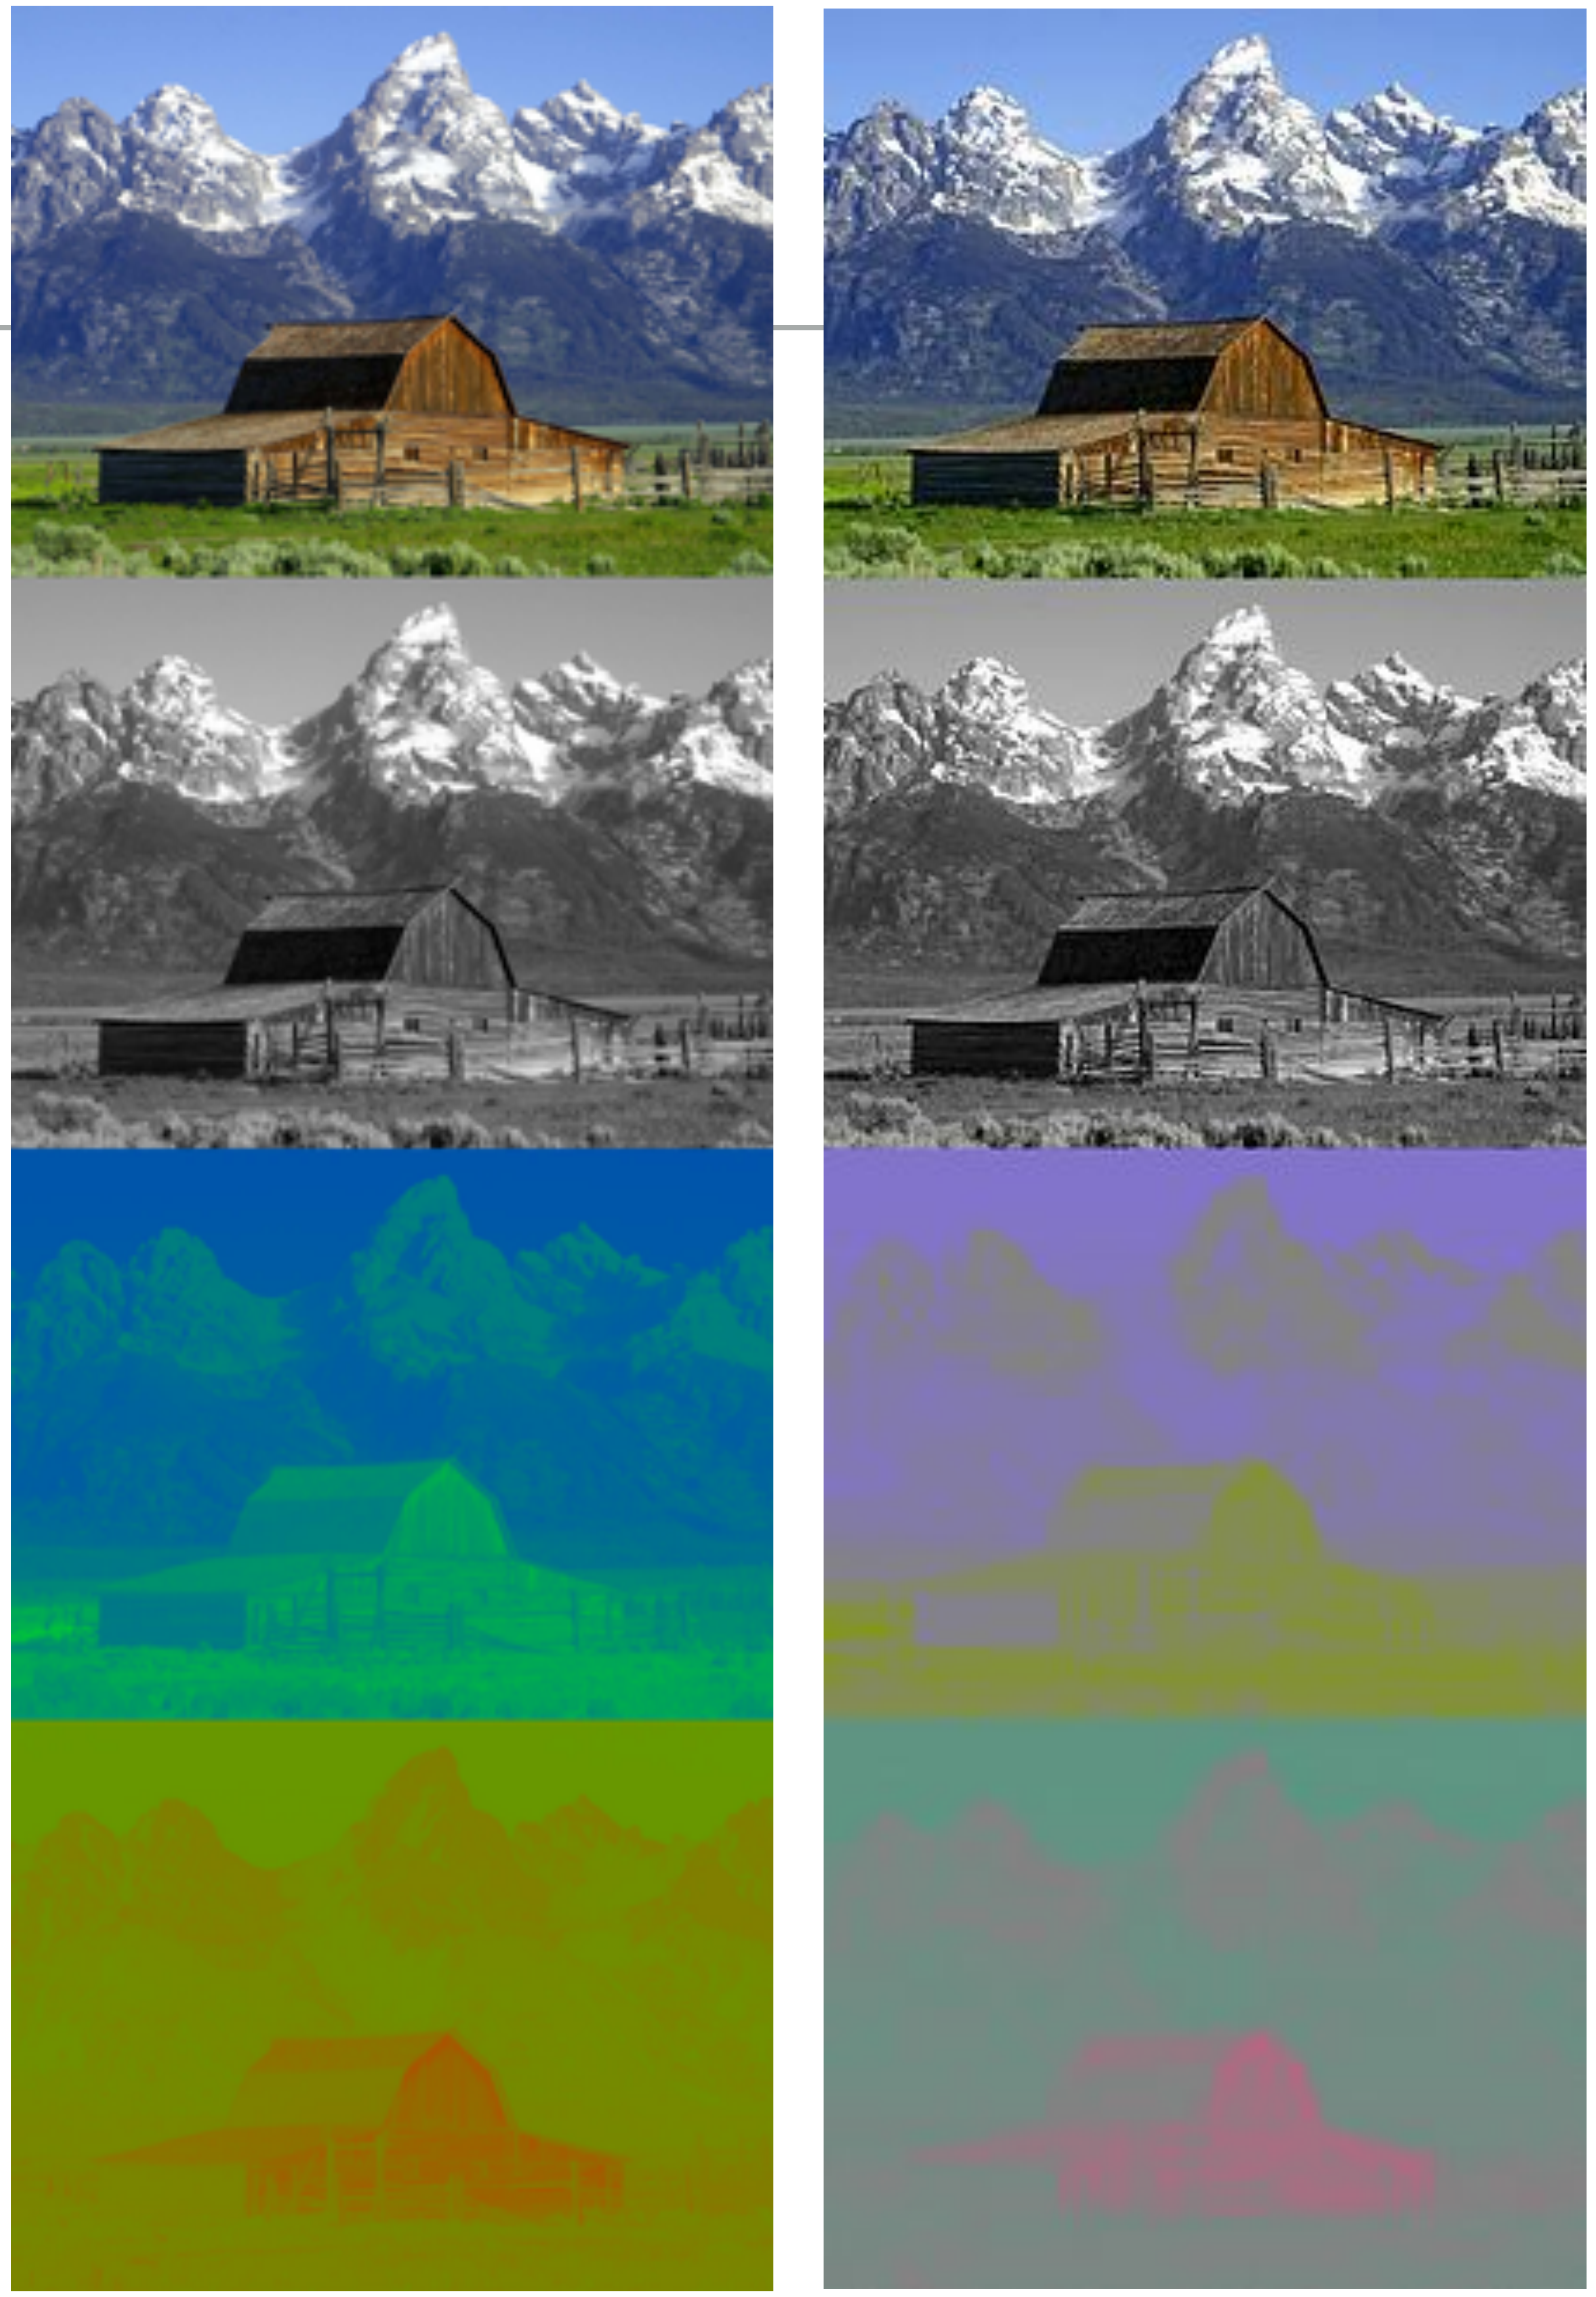

YUV YCrCb

Programming Assignment 2

Gamma

Colorspaces

- color2gray
- luminance–chrominance separation
- luminance-only brightness/contrast
- RGB⇆YUV

Spanish castle illusion

Grayworld whitebalance

Histograms

- autolevels
- visualize RGB histogram
- histogram equalization & histogram matching


© Wojciech Jarosz




© Wojciech Jarosz

© Wojciech Jarosz
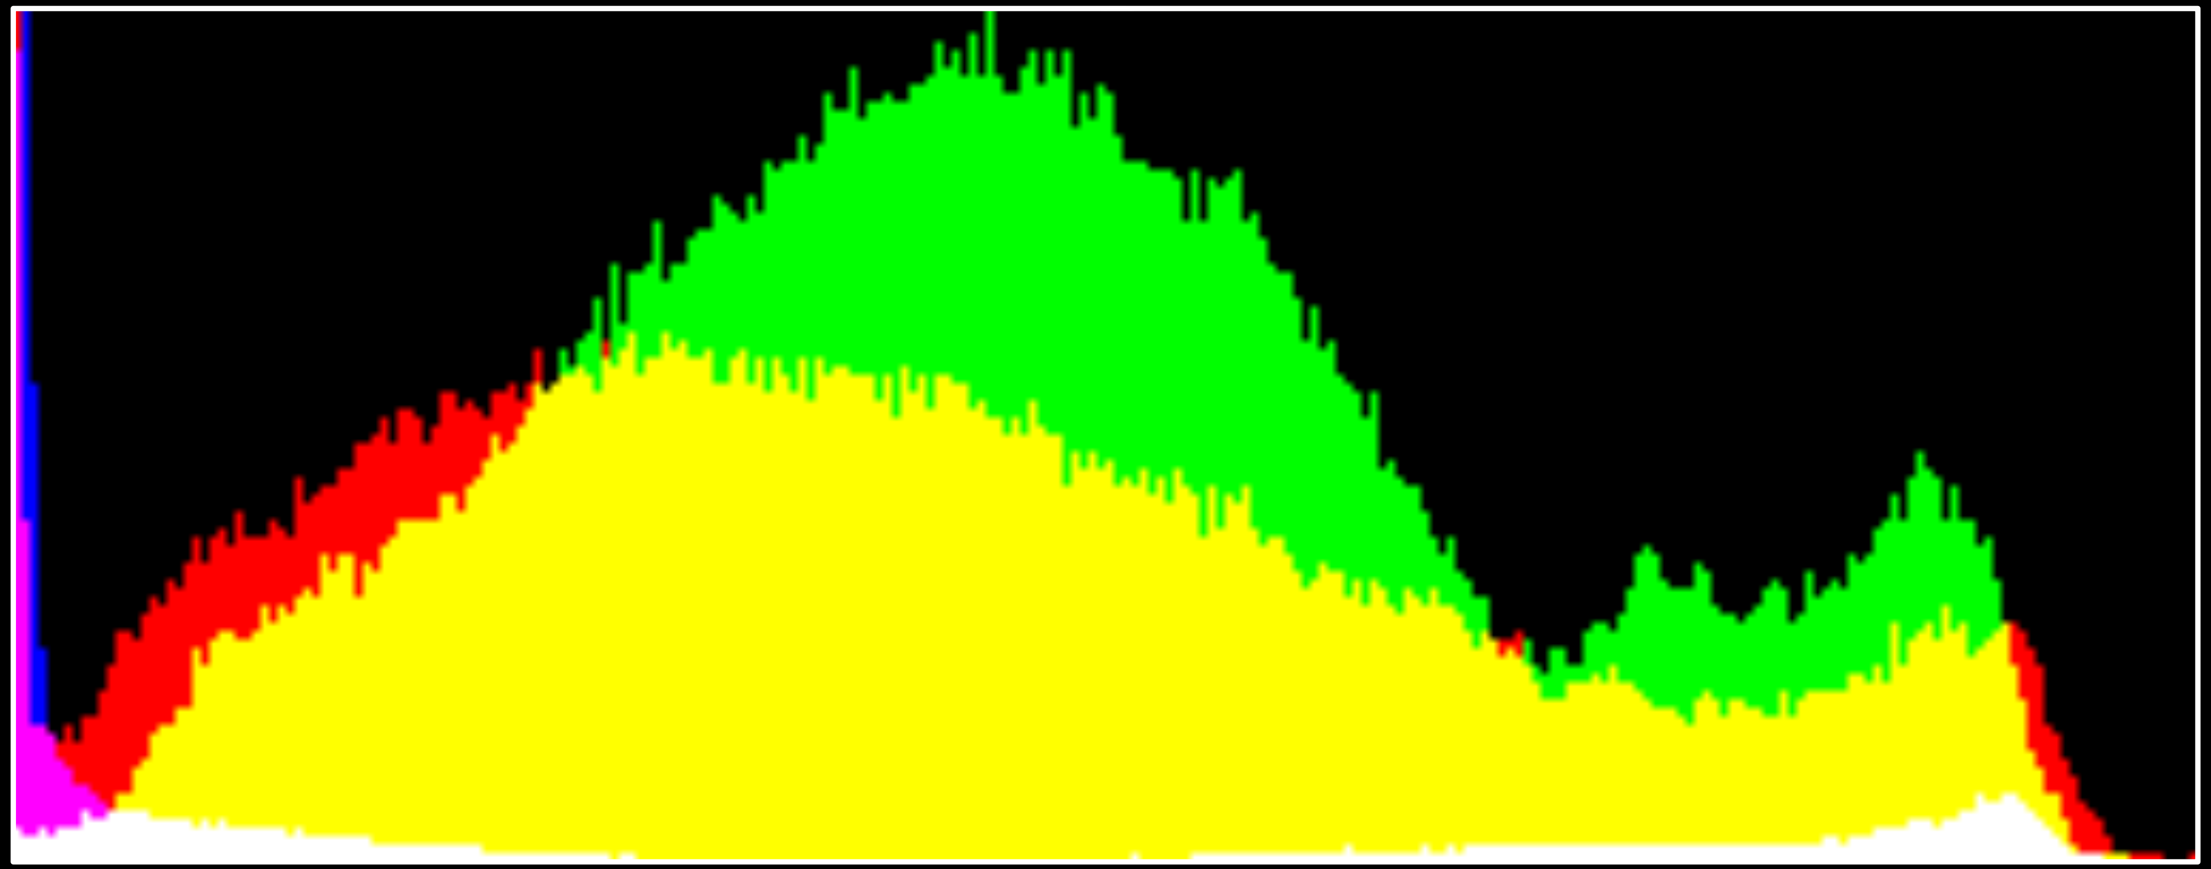

Slide credits

Frédo Durand

Steve Marschner

Matthias Zwicker