Computational Aspects of Digital Photography

Light & Color

Wojciech Jarosz

wojciech.k.jarosz@dartmouth.edu

Dartmouth

Administritivia

Assignment 1 due tomorrow

- Please bring your pinhole cameras to class on Thursday for “show and tell” :-)

Assignment 2 available soon

- back to programming
- due next Wednesday

Fill out poll to select paper for presentations (before class Thursday)

- http://goo.gl/forms/Yfwp3ee5sW

Today

Light & Color

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

What is light?

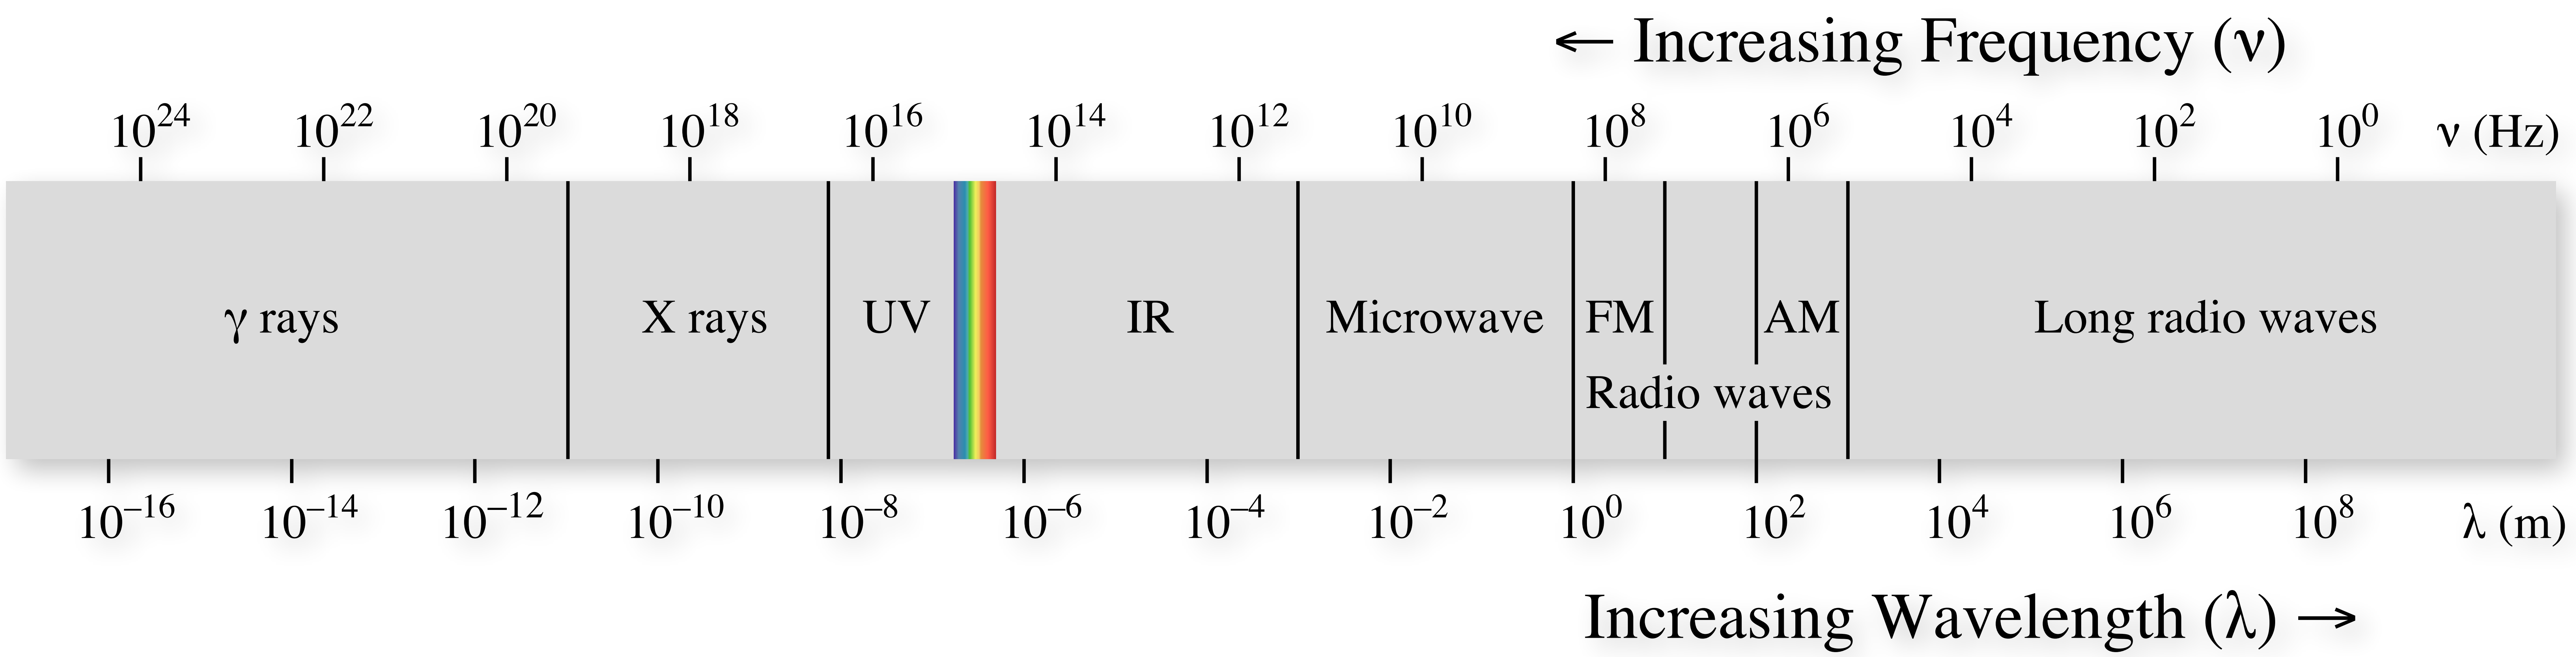

A form of electromagnetic (EM) radiation

- like x-rays, microwaves, radio waves, etc
- characterized by wavelength
- amplitude determines intensity

We perceive a limited section of the spectrum as "visible" light

What is light?

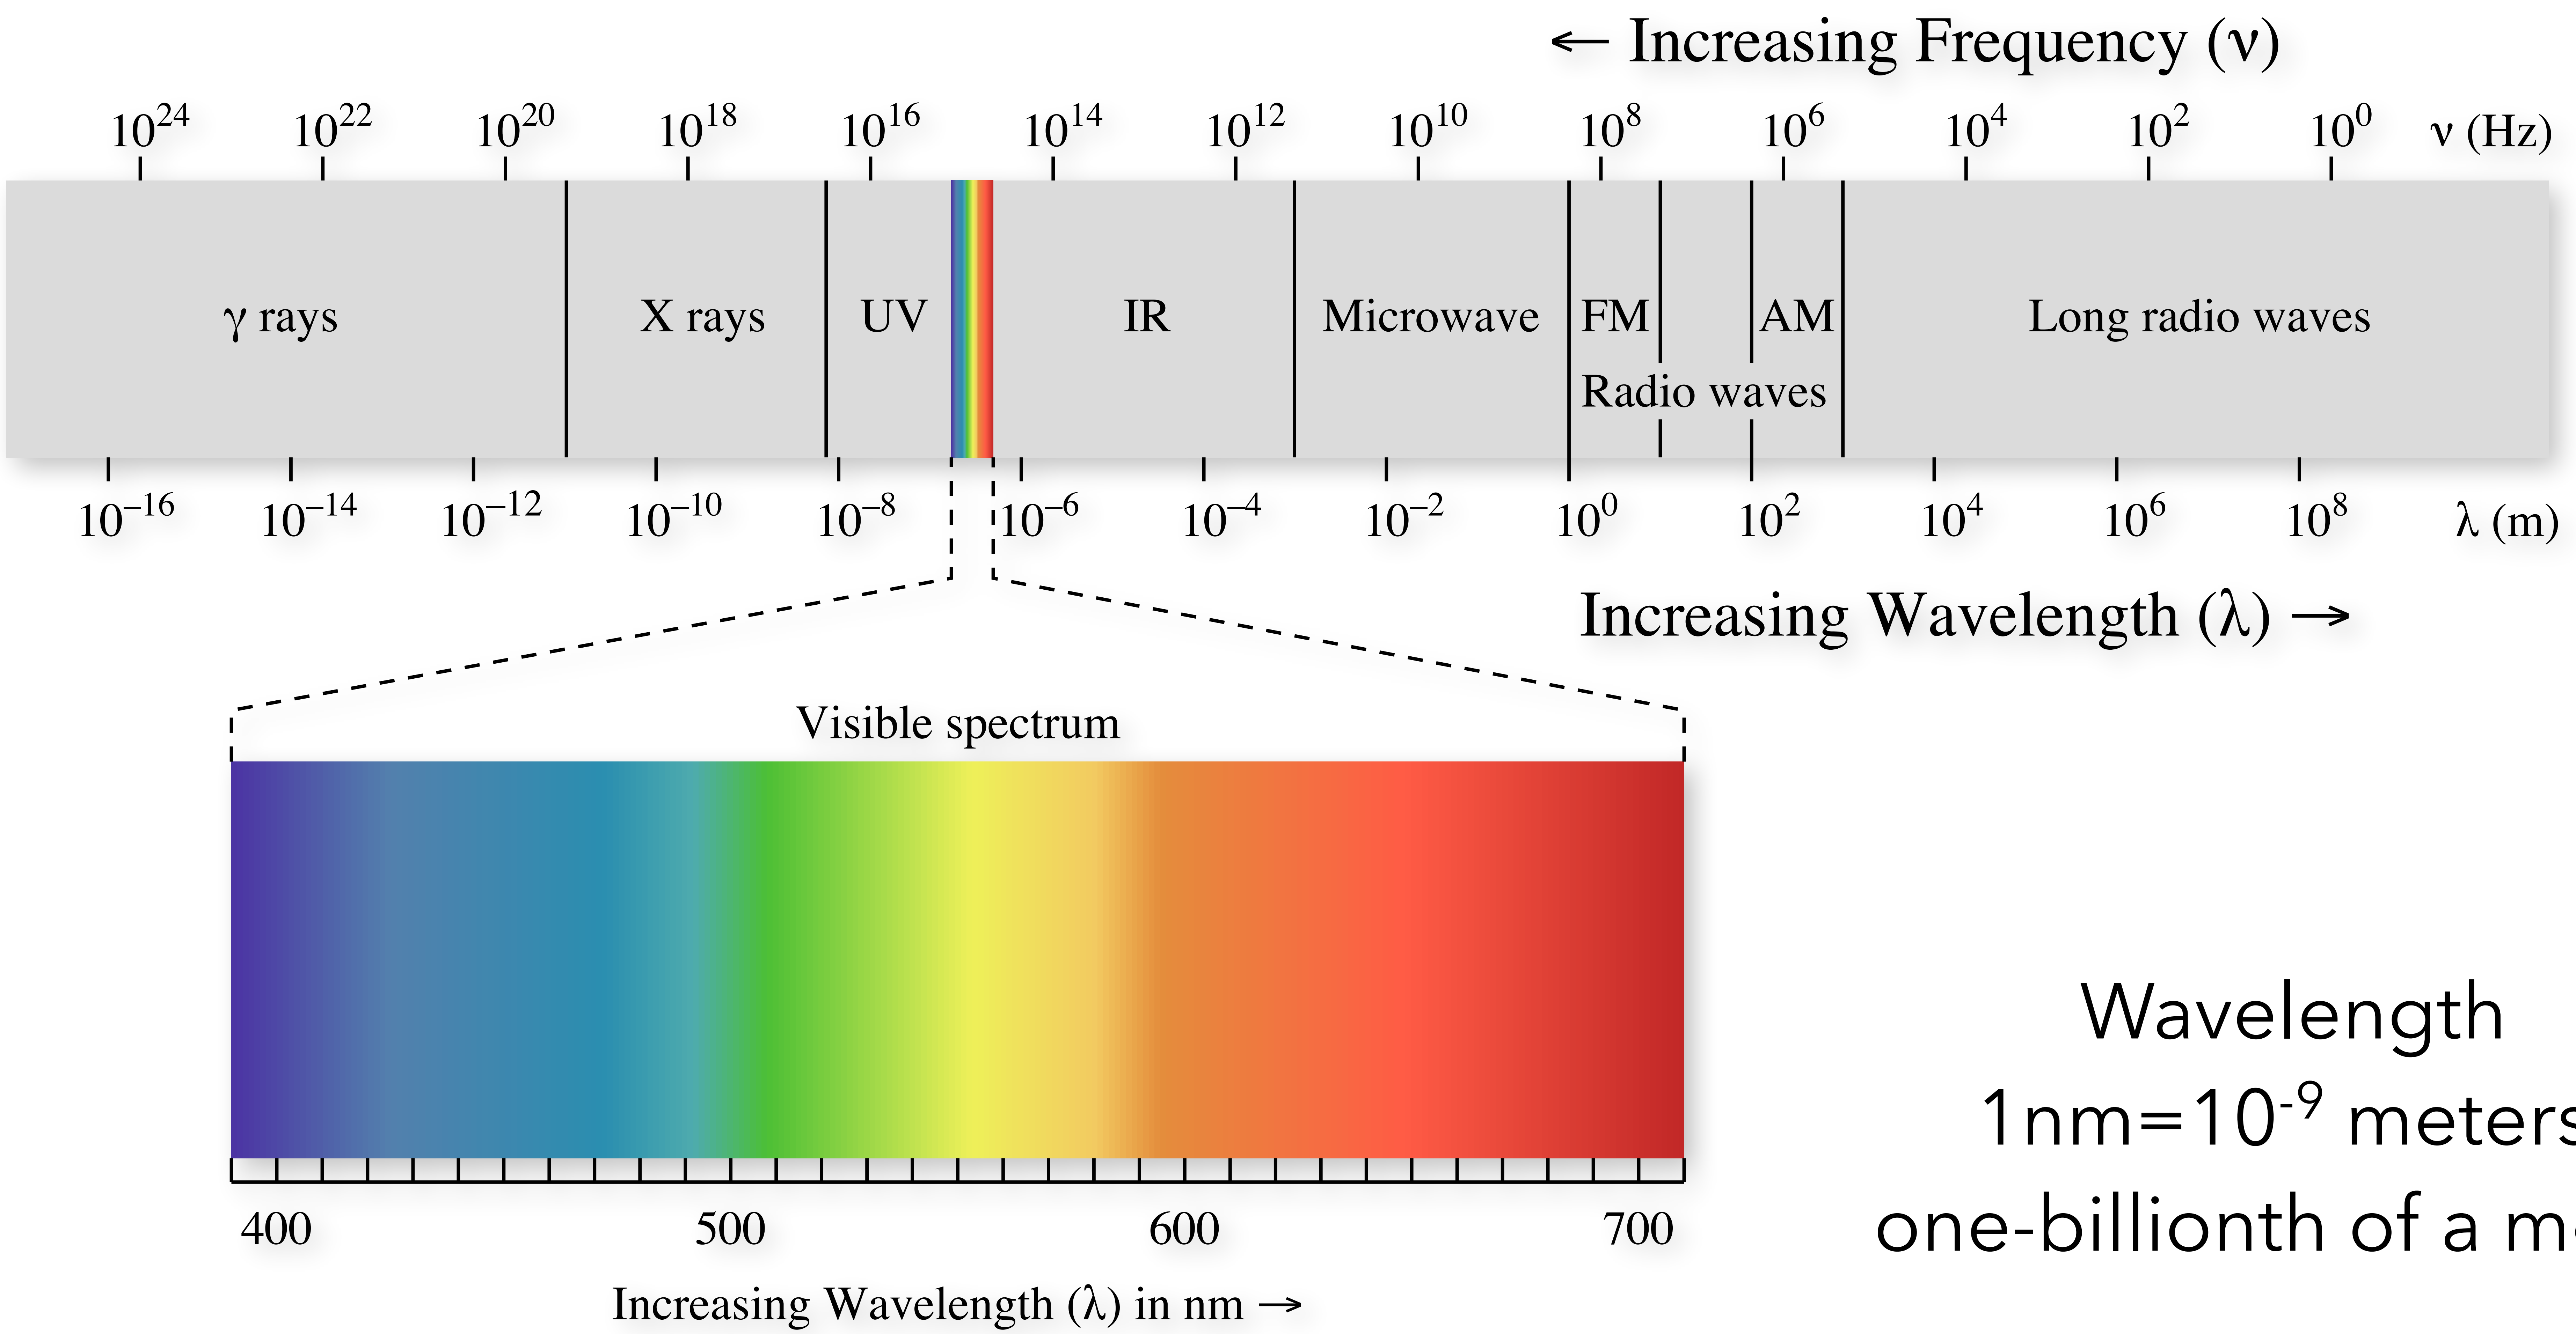

Wavelength

1nm=10^{-9} meters,

one-billionth of a meter

speed of light = wavelength * frequency

Light transport: Geometric optics

Simplified model

- http://en.wikipedia.org/wiki/Geometrical_optics

Roughly speaking

- Light is transported along straight rays
- When light interacts with material, it may be reflected or refracted

Can model most effects that are important for our daily experience

Light transport: Geometric optics

Rays carry a spectrum of electromagnetic energy

- An "energy distribution"

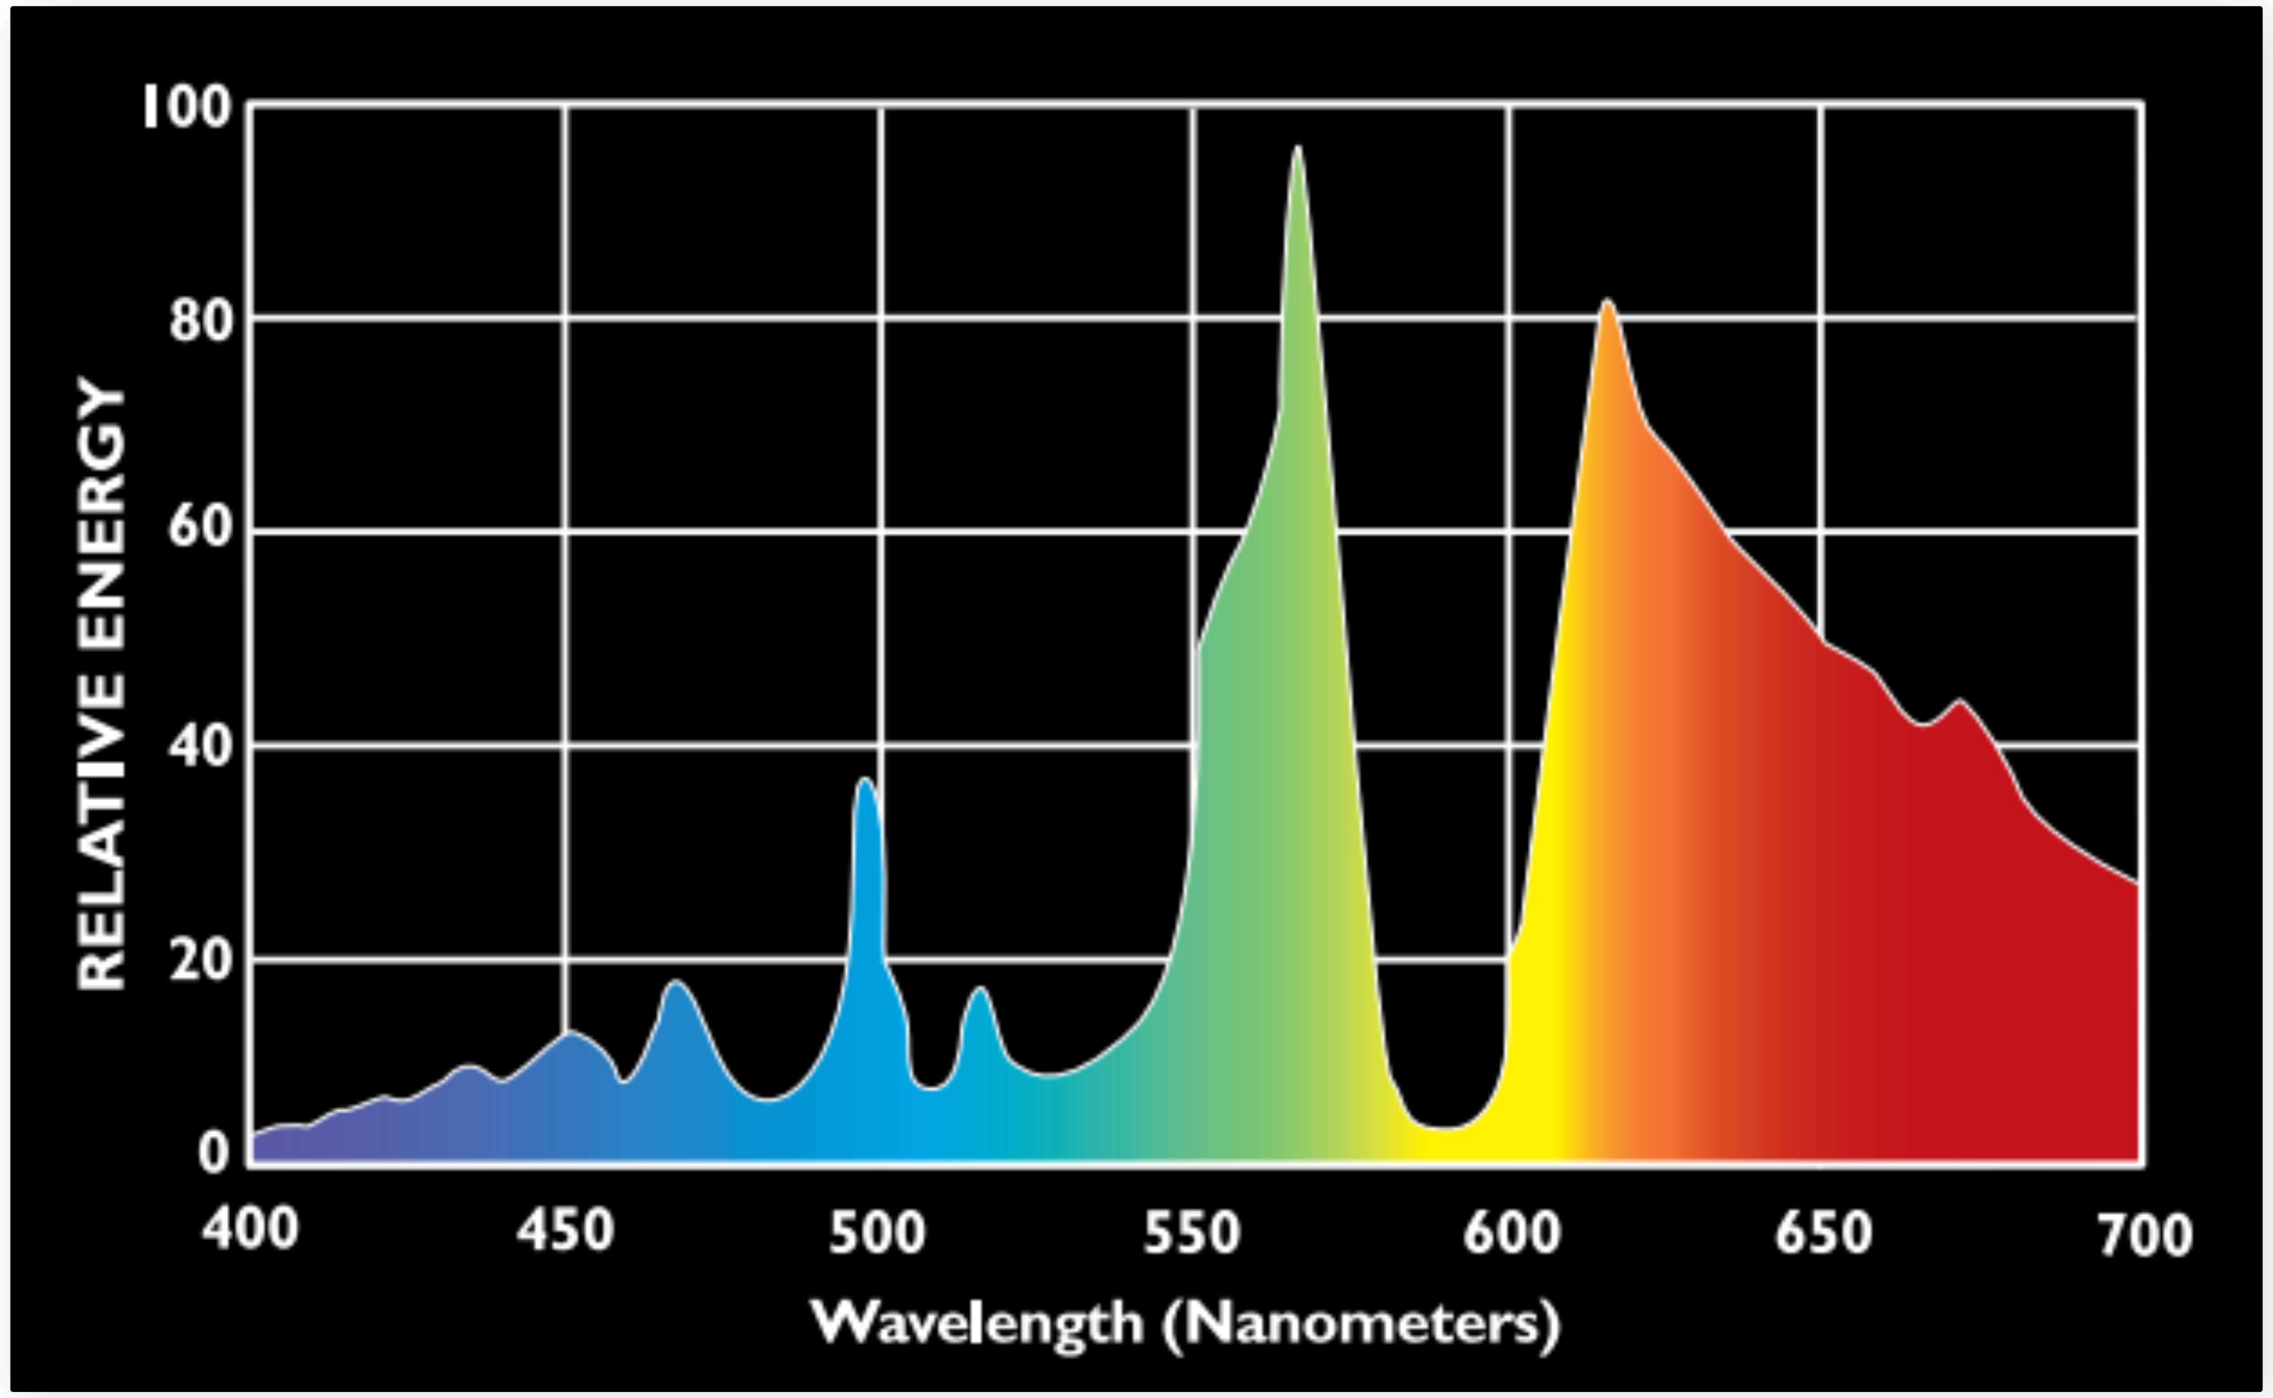

Spectral distribution of light

Light can be a mixture of many wavelengths

- each with some intensity
- represented by continuous function
  - $s(\lambda)$ = intensity at wavelength λ
- **spectral power distribution (SPD)**: intensity as a function of wavelength over enter spectrum

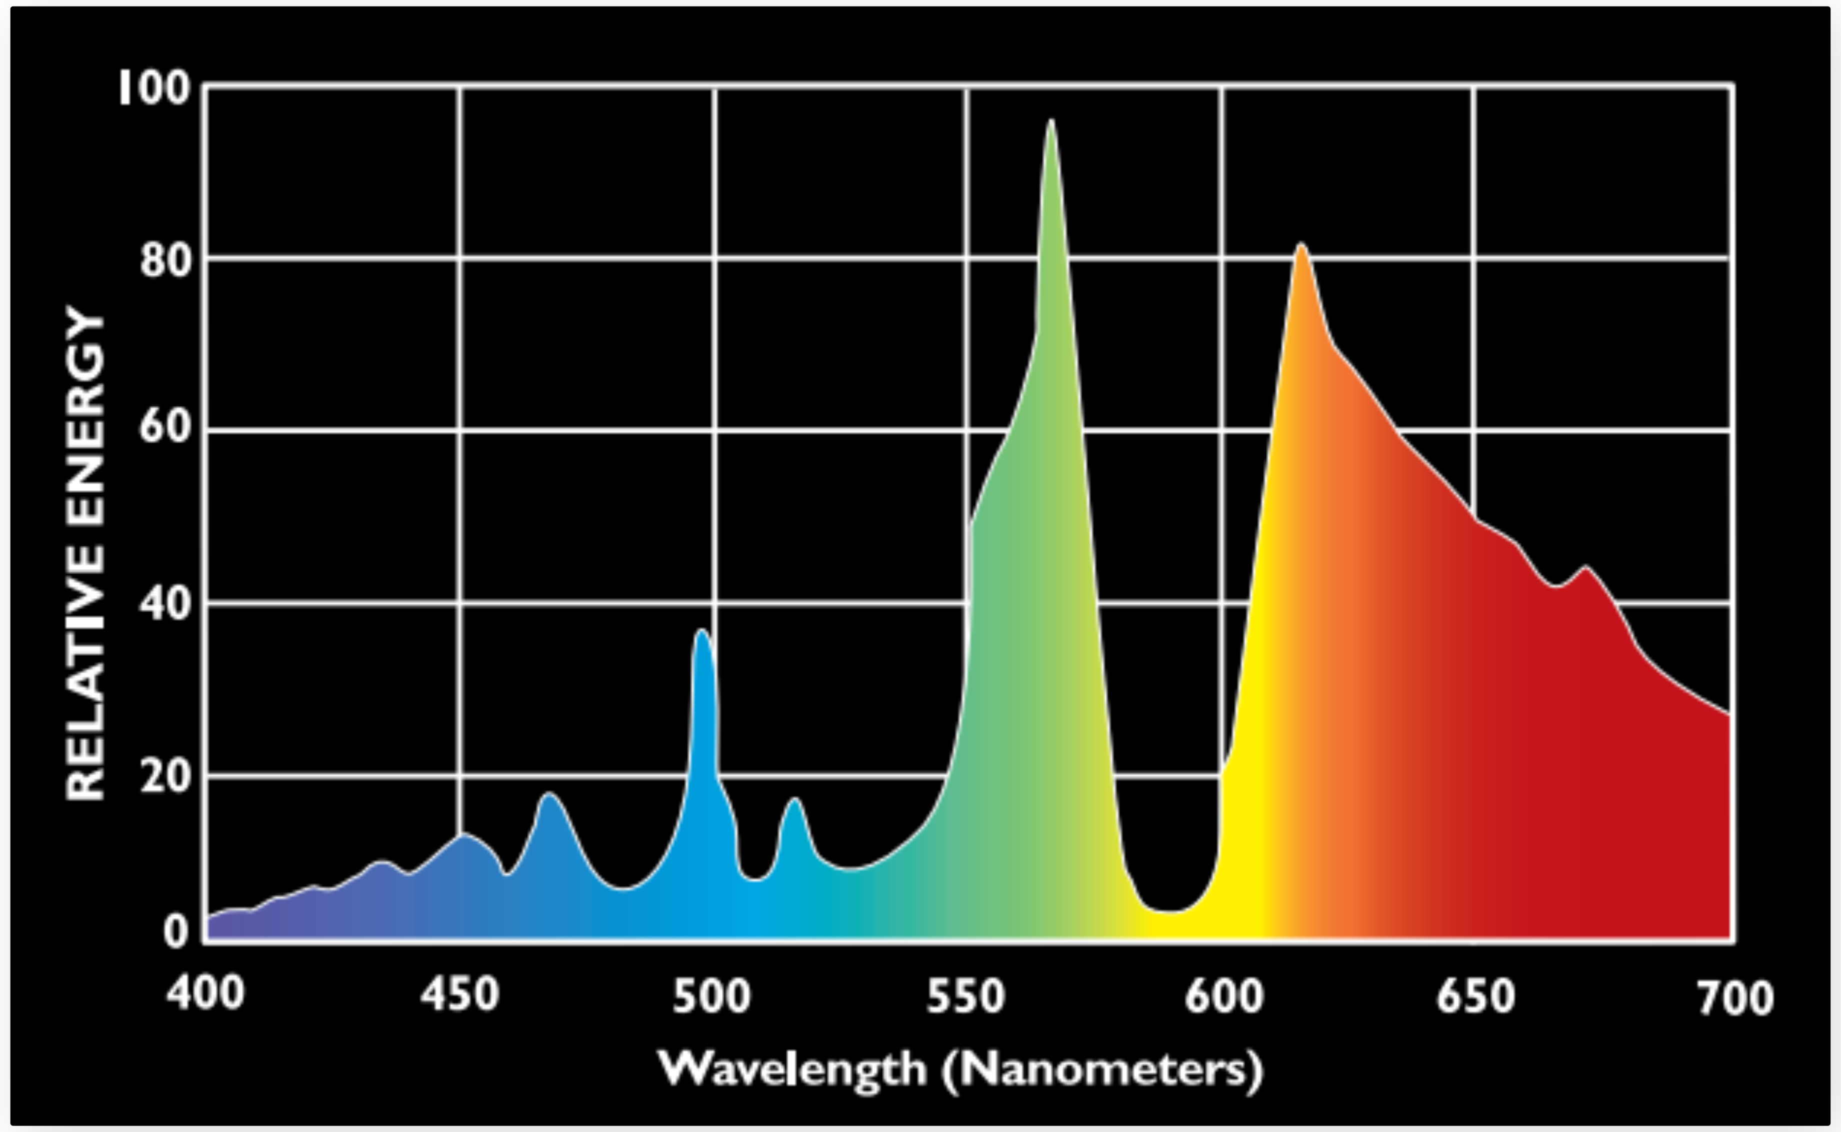

We perceive these distributions as colors

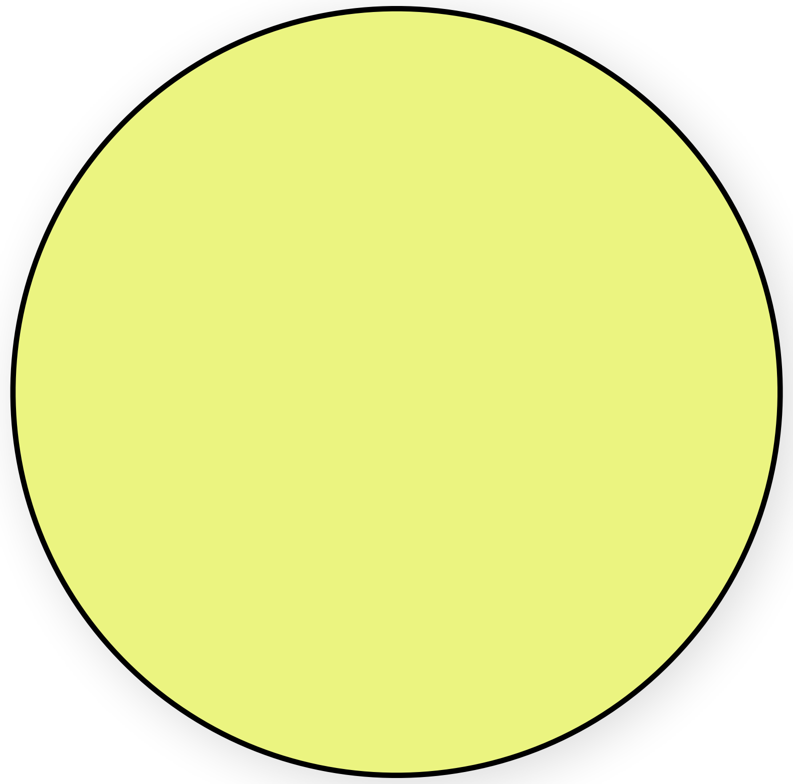

Light-matter interaction

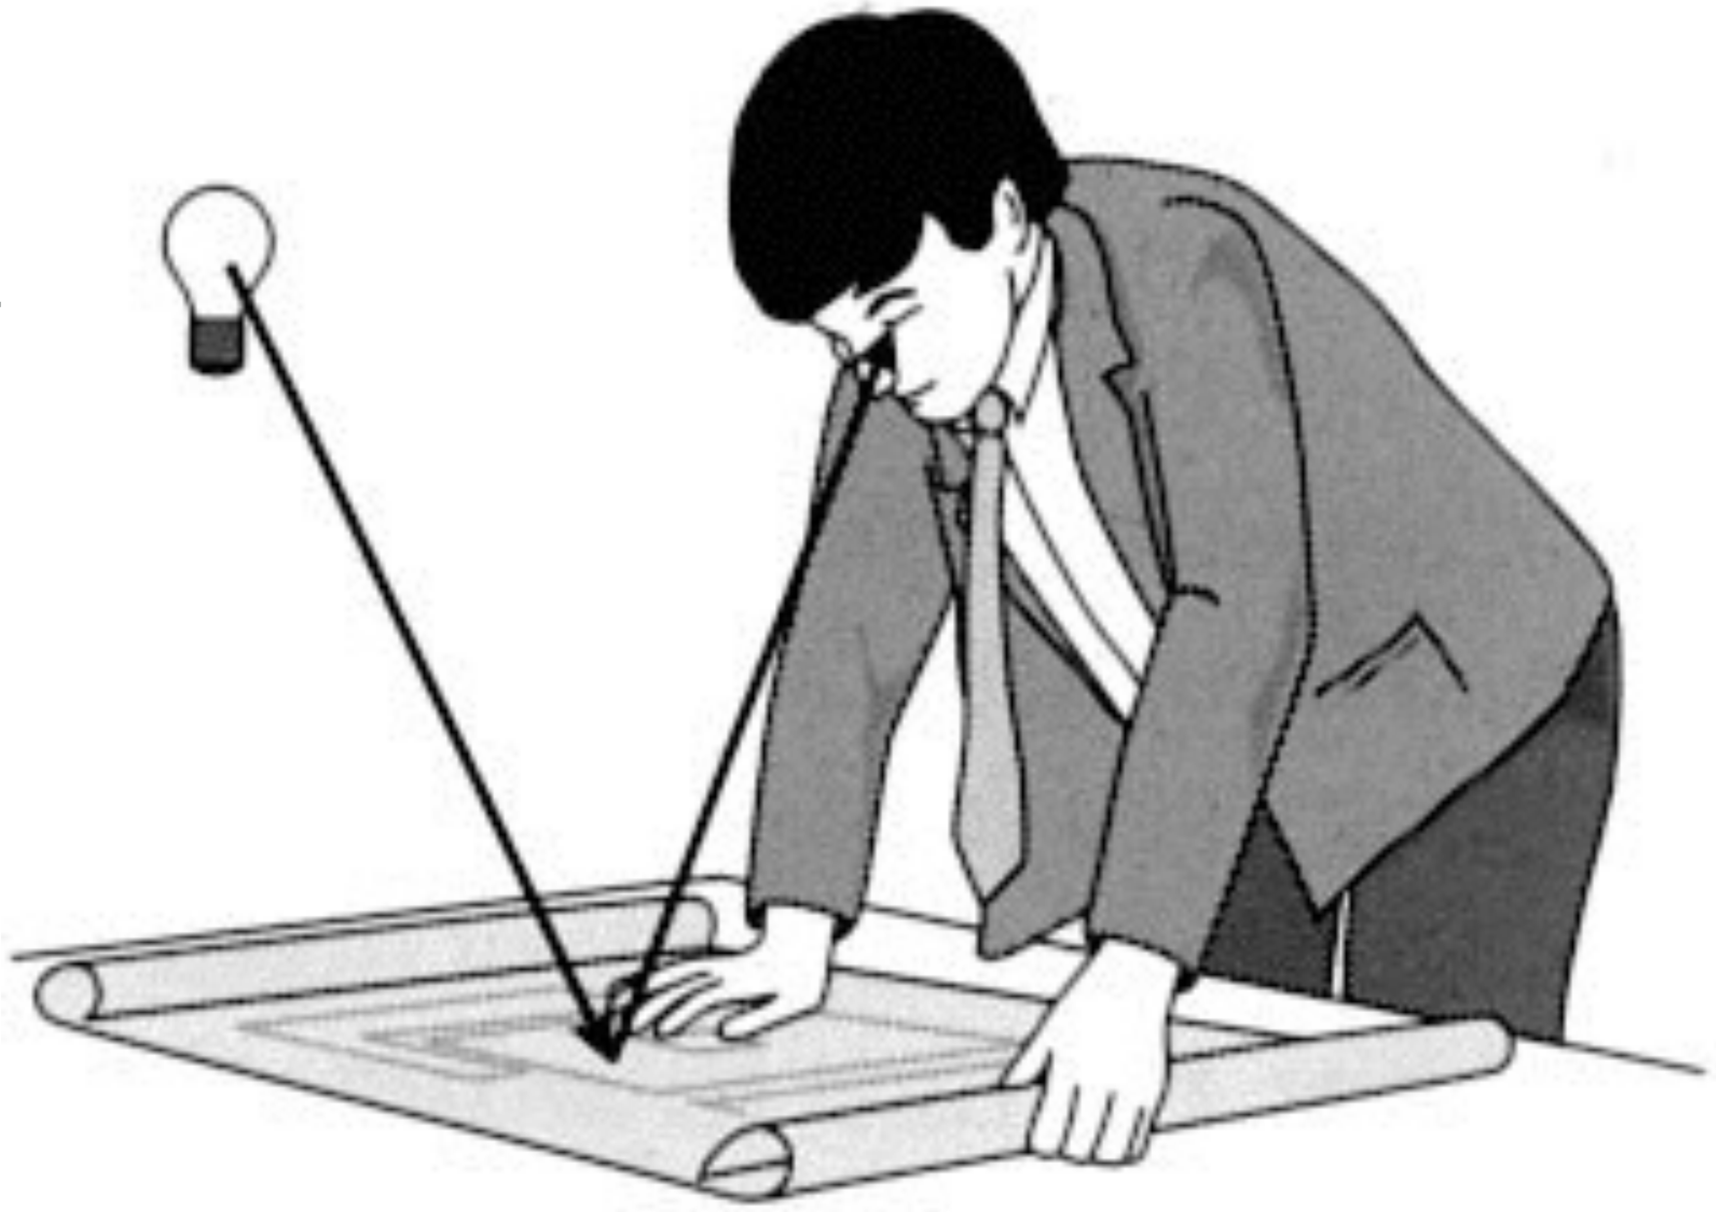

Where spectra come from:

- light source spectrum
- object reflectance (aka spectral albedo)
- multiplied wavelength by wavelength

There are different physical processes that explain this multiplication e.g. absorption, interferences

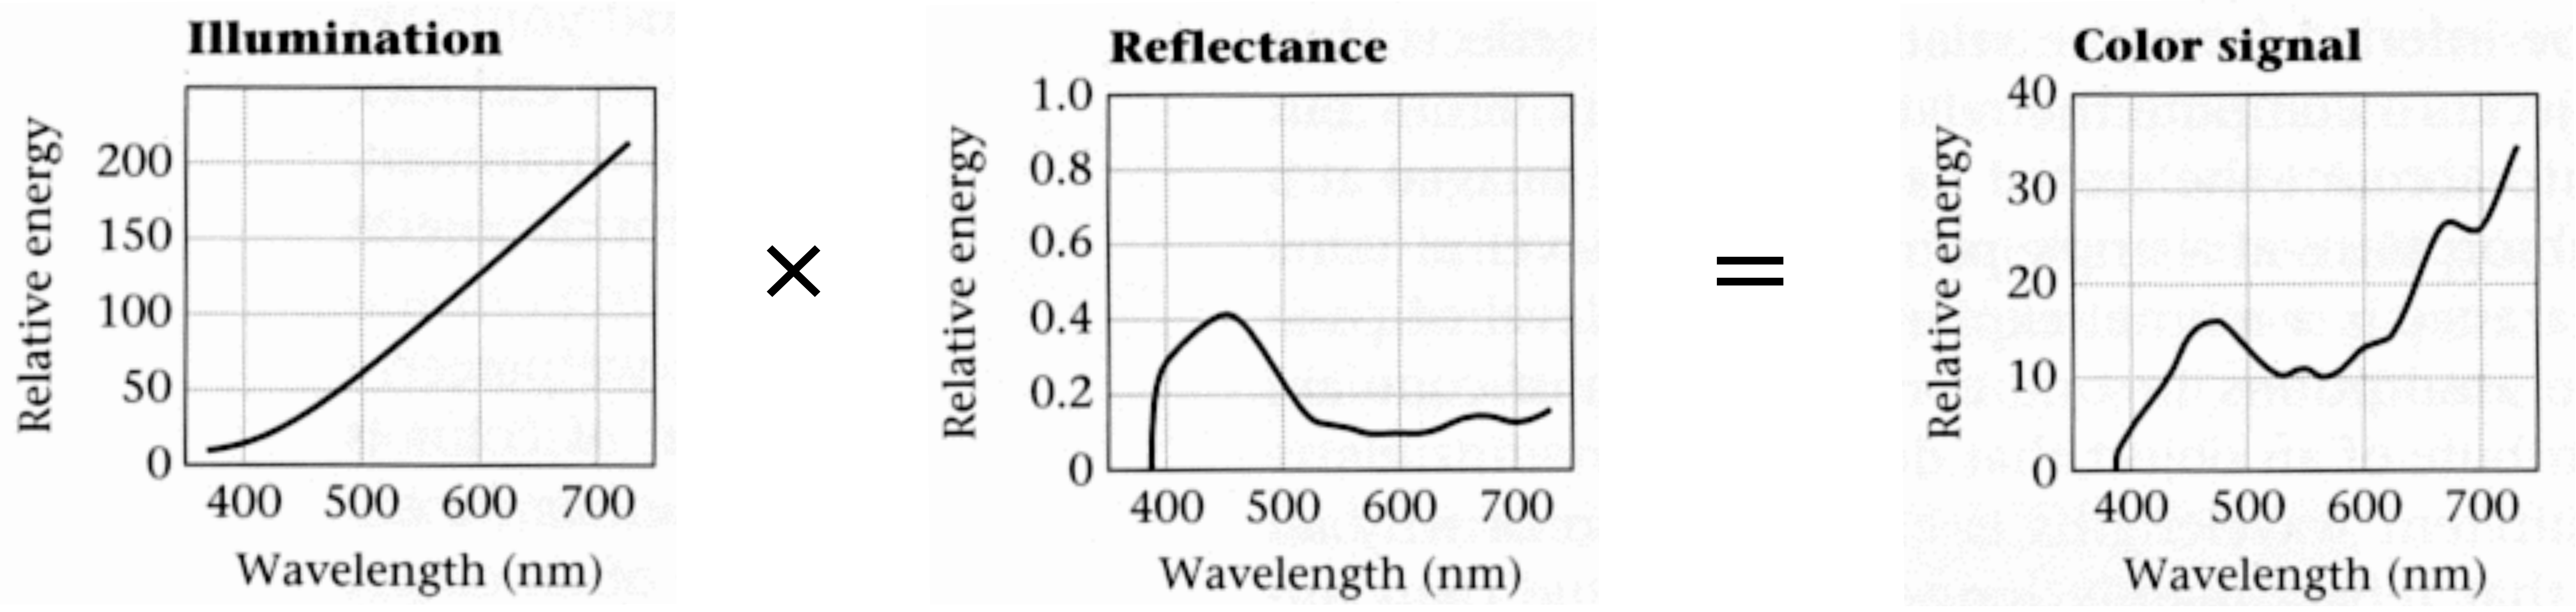

What is color?

Colors are the sensations that arise from light energy with different wavelength distributions

Color is a phenomenon of human perception; it is not a universal property of light

Roughly speaking, things appear “colored” when they depend on wavelength and “gray” when they do not

The problem of color science

Build a model for human color perception

That is, map a physical light description to a perceptual color sensation

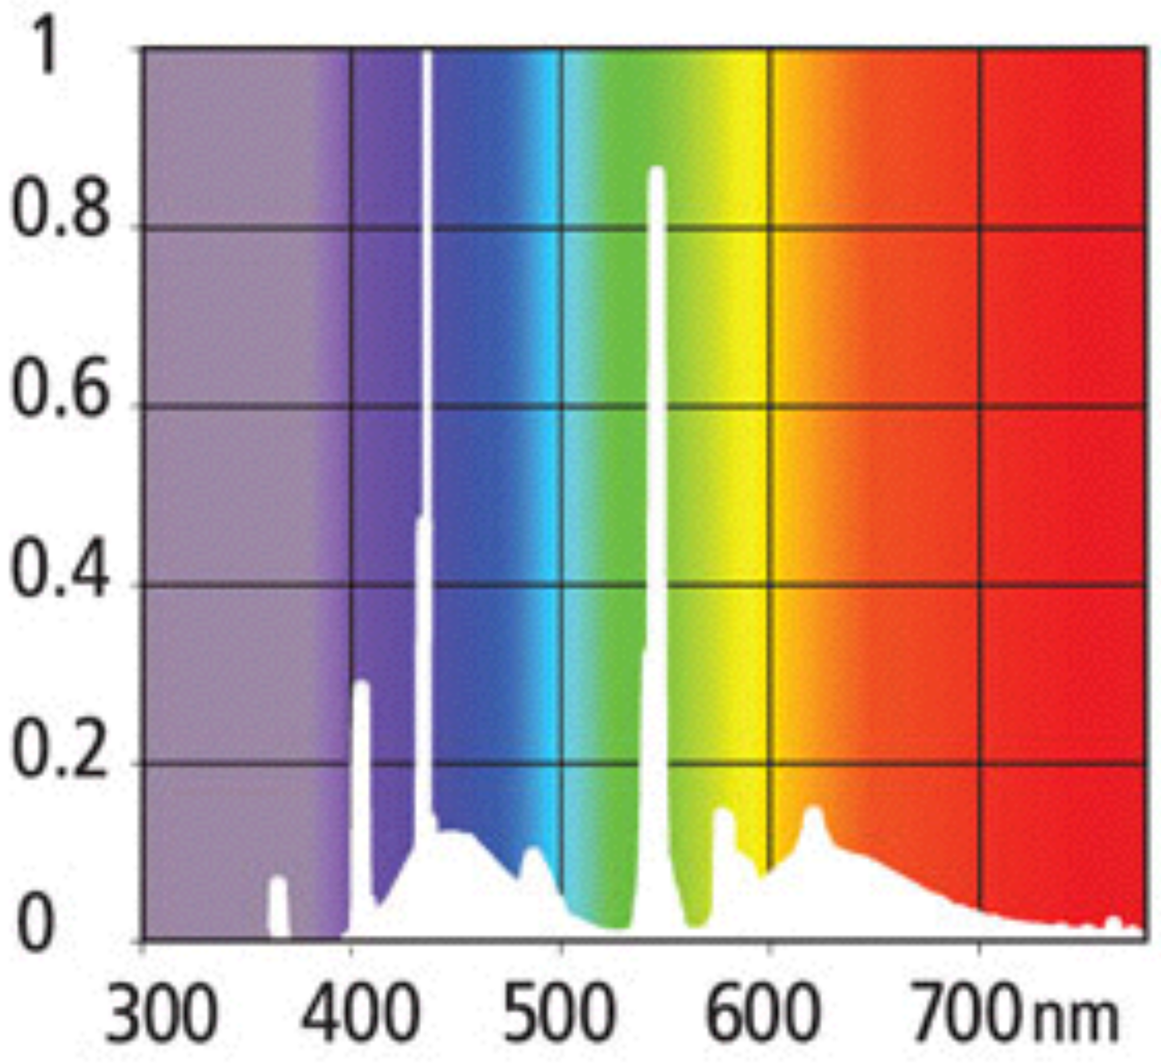

Physical

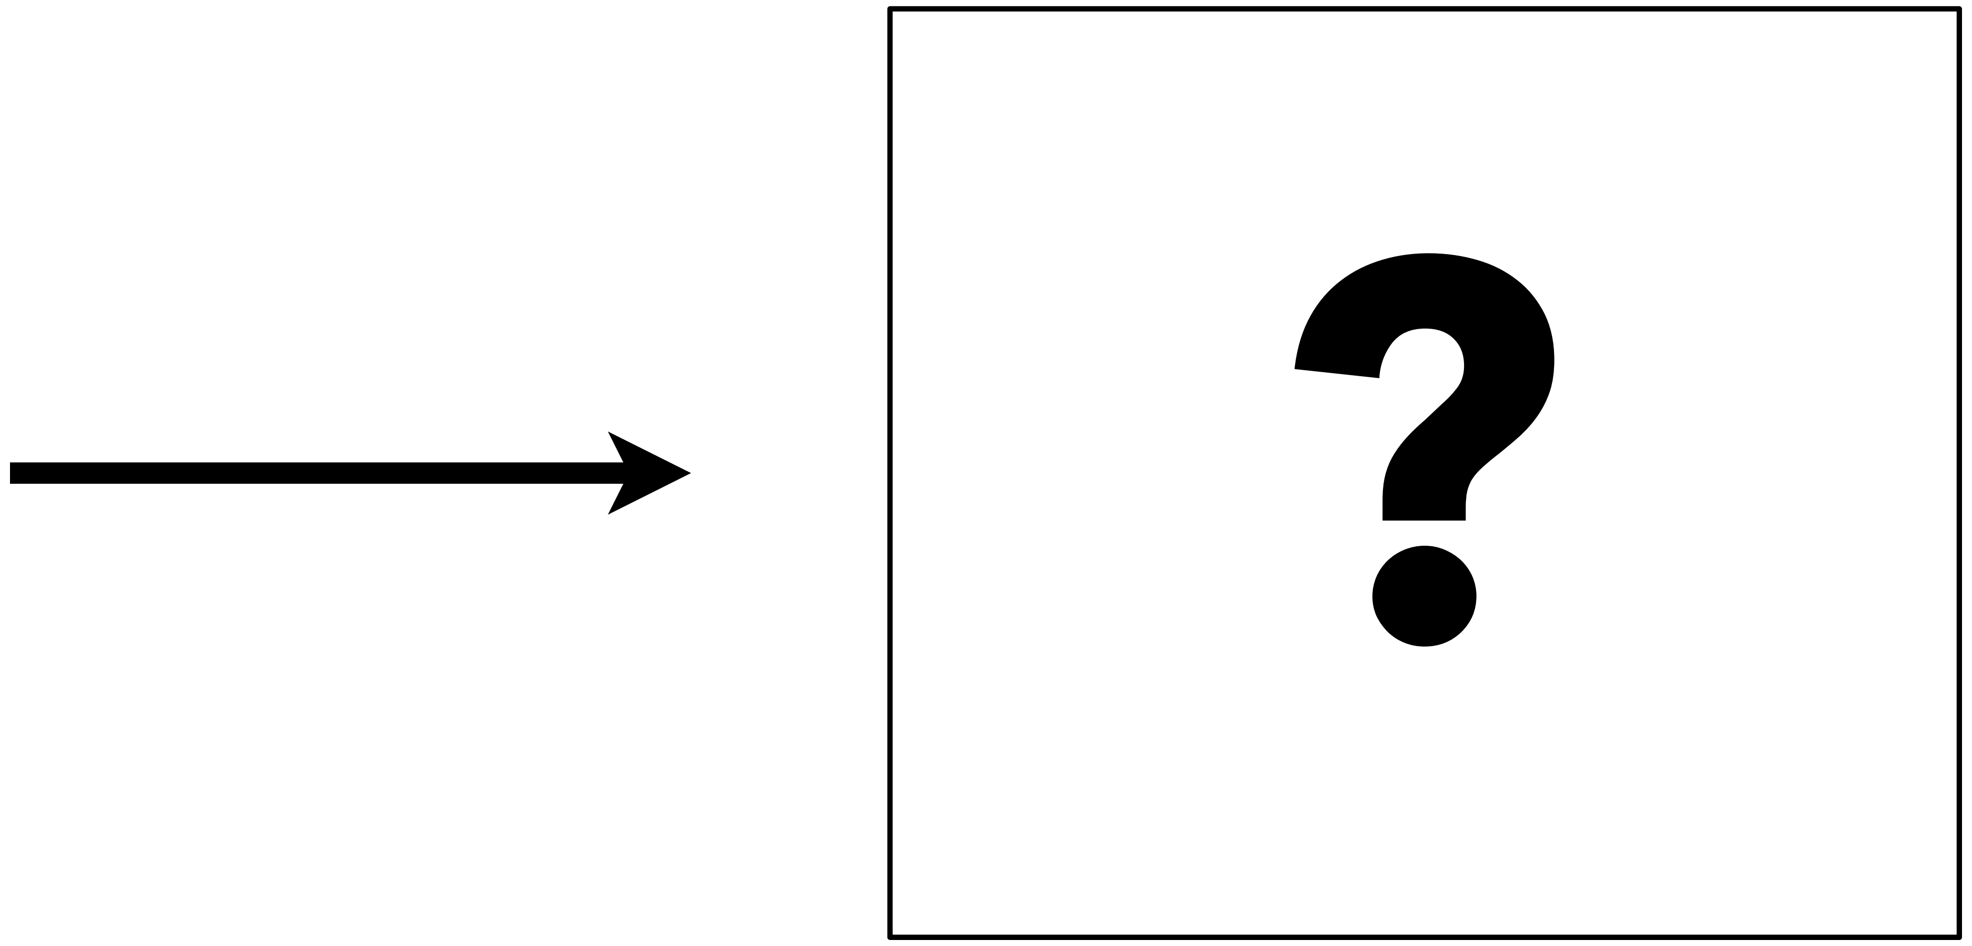

Perceptual

Today

Light & Color

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

The eye as a measurement device

We can model the low-level behavior of the eye by thinking of it as a light-measuring machine

- optics are much like a camera
- its detection mechanism is also much like a camera

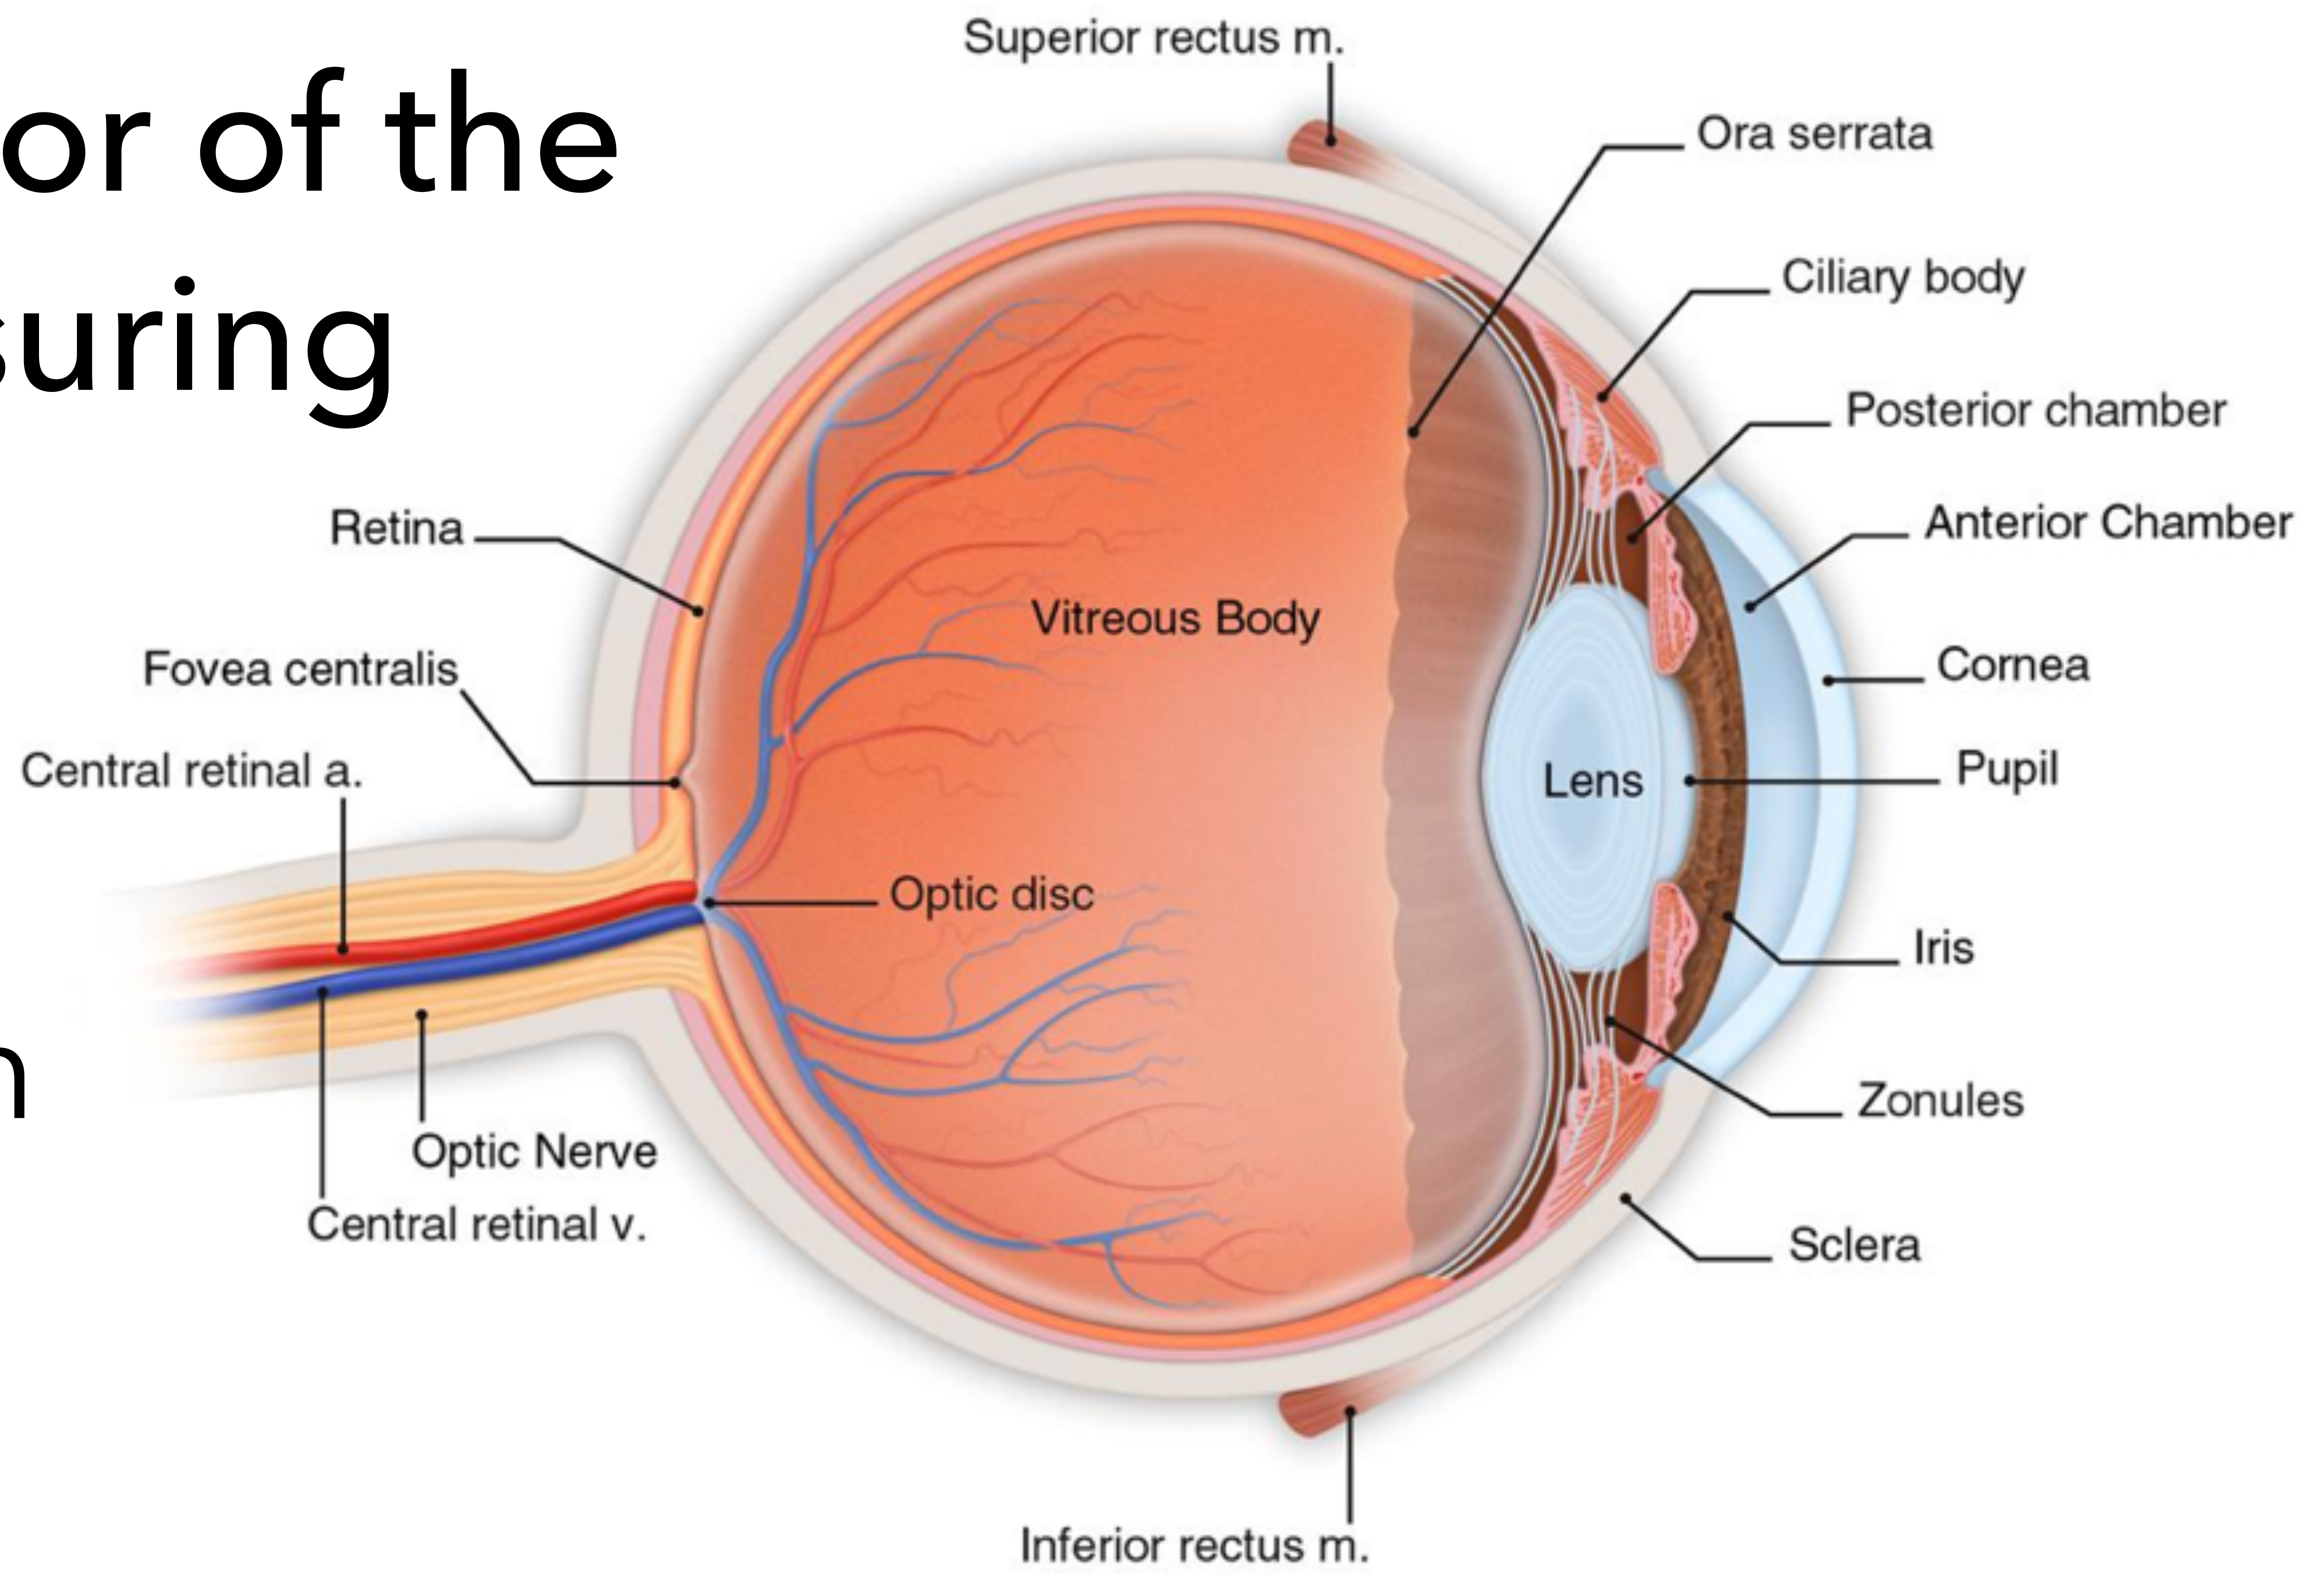

Lens focuses light on retina

- cells in retina respond to light
- different types respond to different wavelengths

Retinal composition: two kinds of cells

Cones are concentrated in fovea

- high acuity, require more light
- “respond to color”

Rods concentrated outside fovea

- lower acuity, require less light
- roughly 10x more sensitive
- “respond to intensity only”

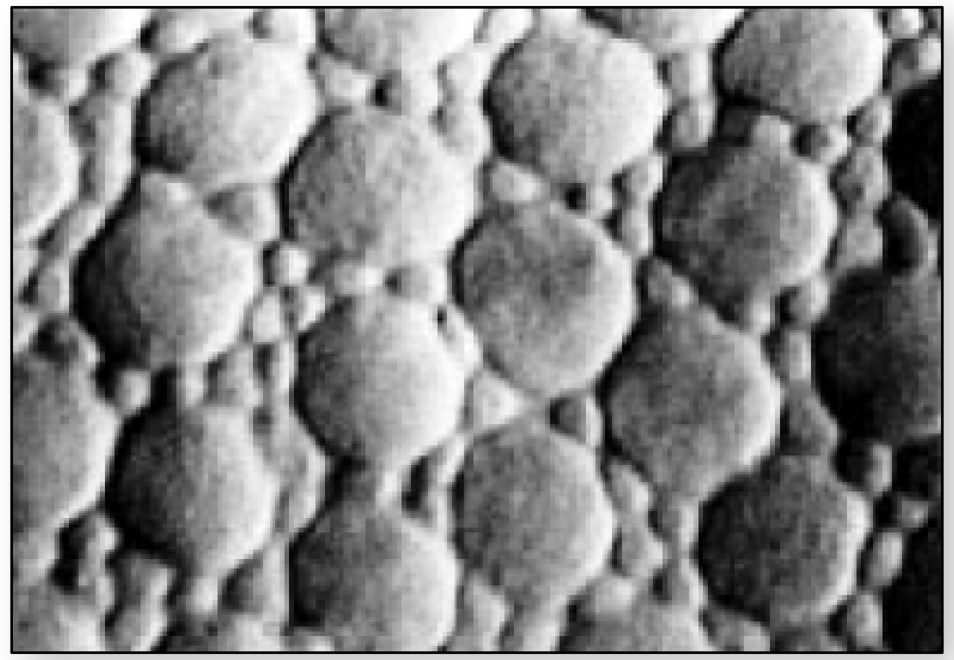

near fovea

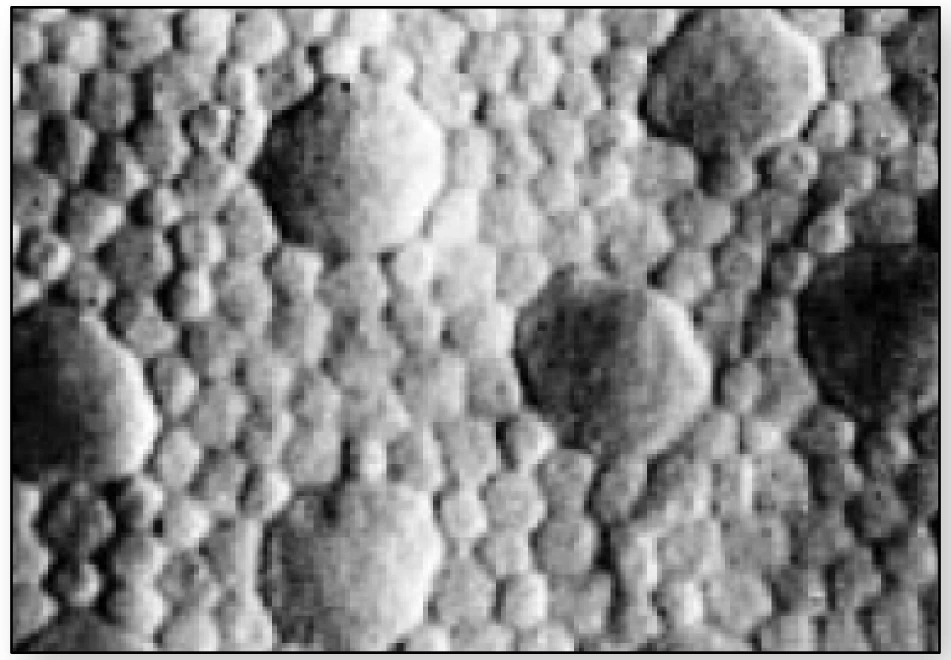

away from fovea

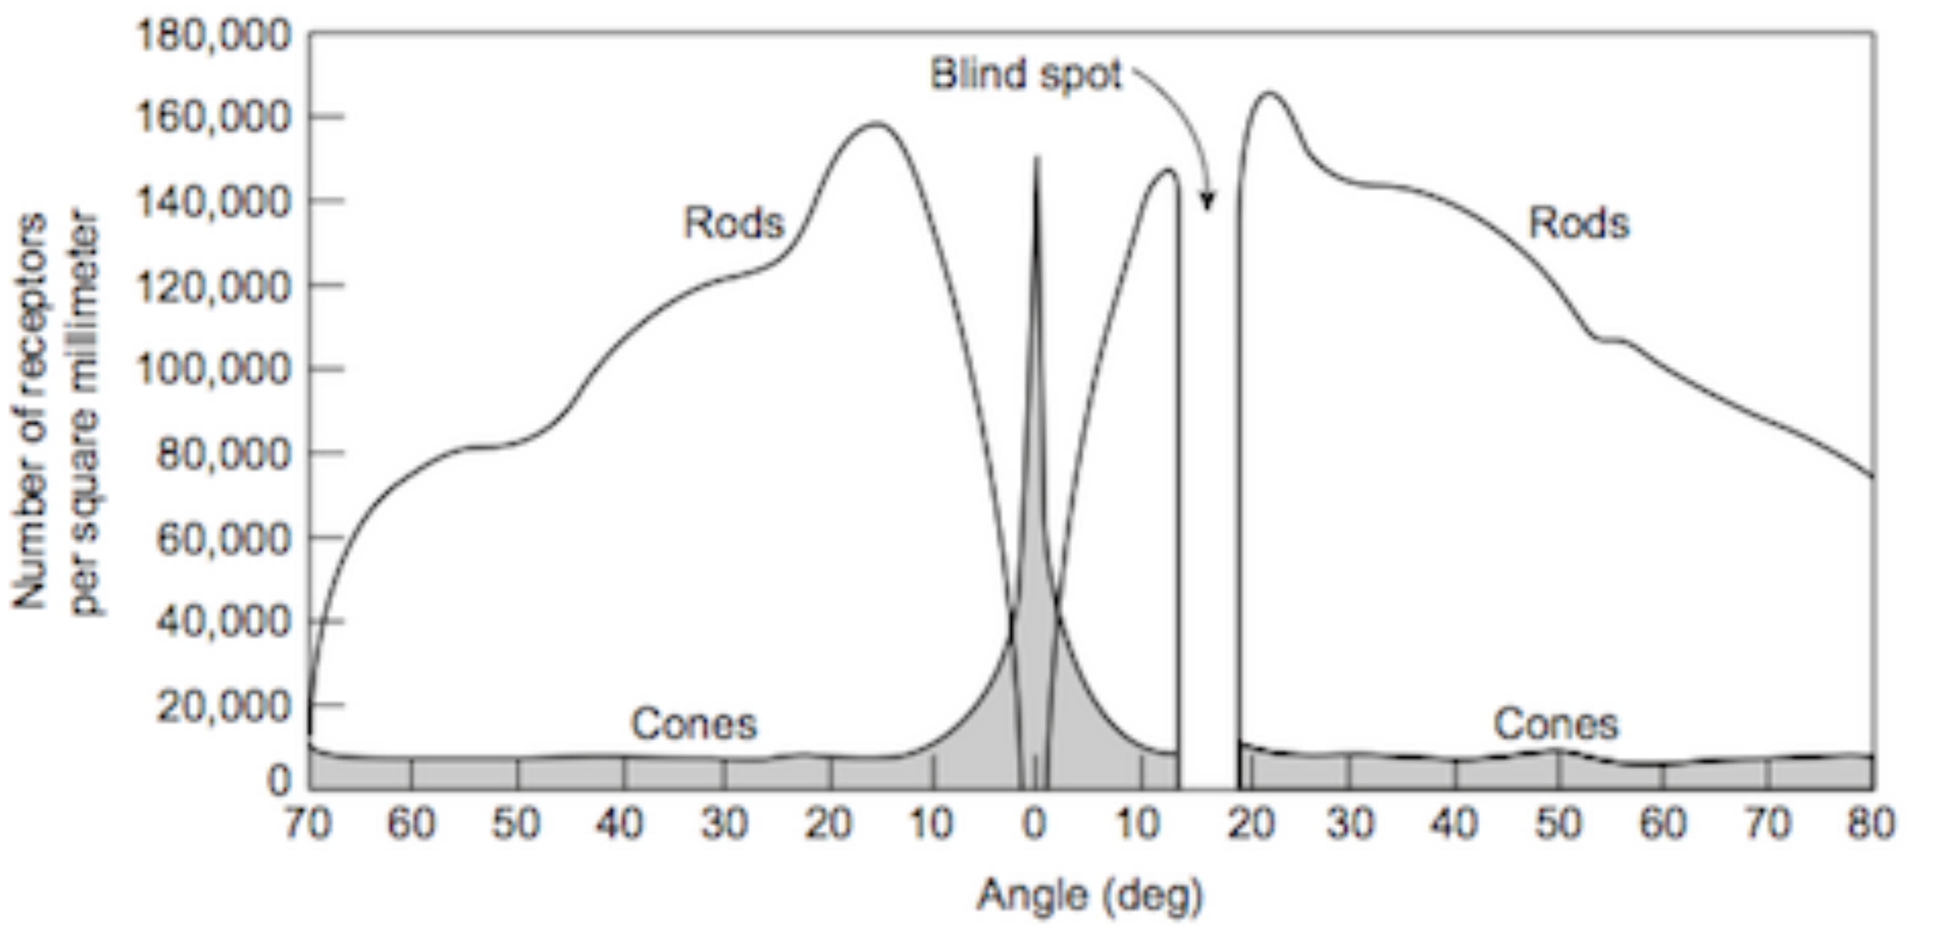

A simple light detector

Produces a scalar value (a number) when photons land on it

- this value depends strictly on the number of photons detected
- each photon has a probability of being detected that depends on the wavelength
- there is no way to tell the difference between signals caused by light of different wavelengths: there is just a number

This is a reasonable model for many detectors:

- based on semiconductors (such as in a digital camera)
- based on visual photopigments (such as in human eyes)

A simple light detector

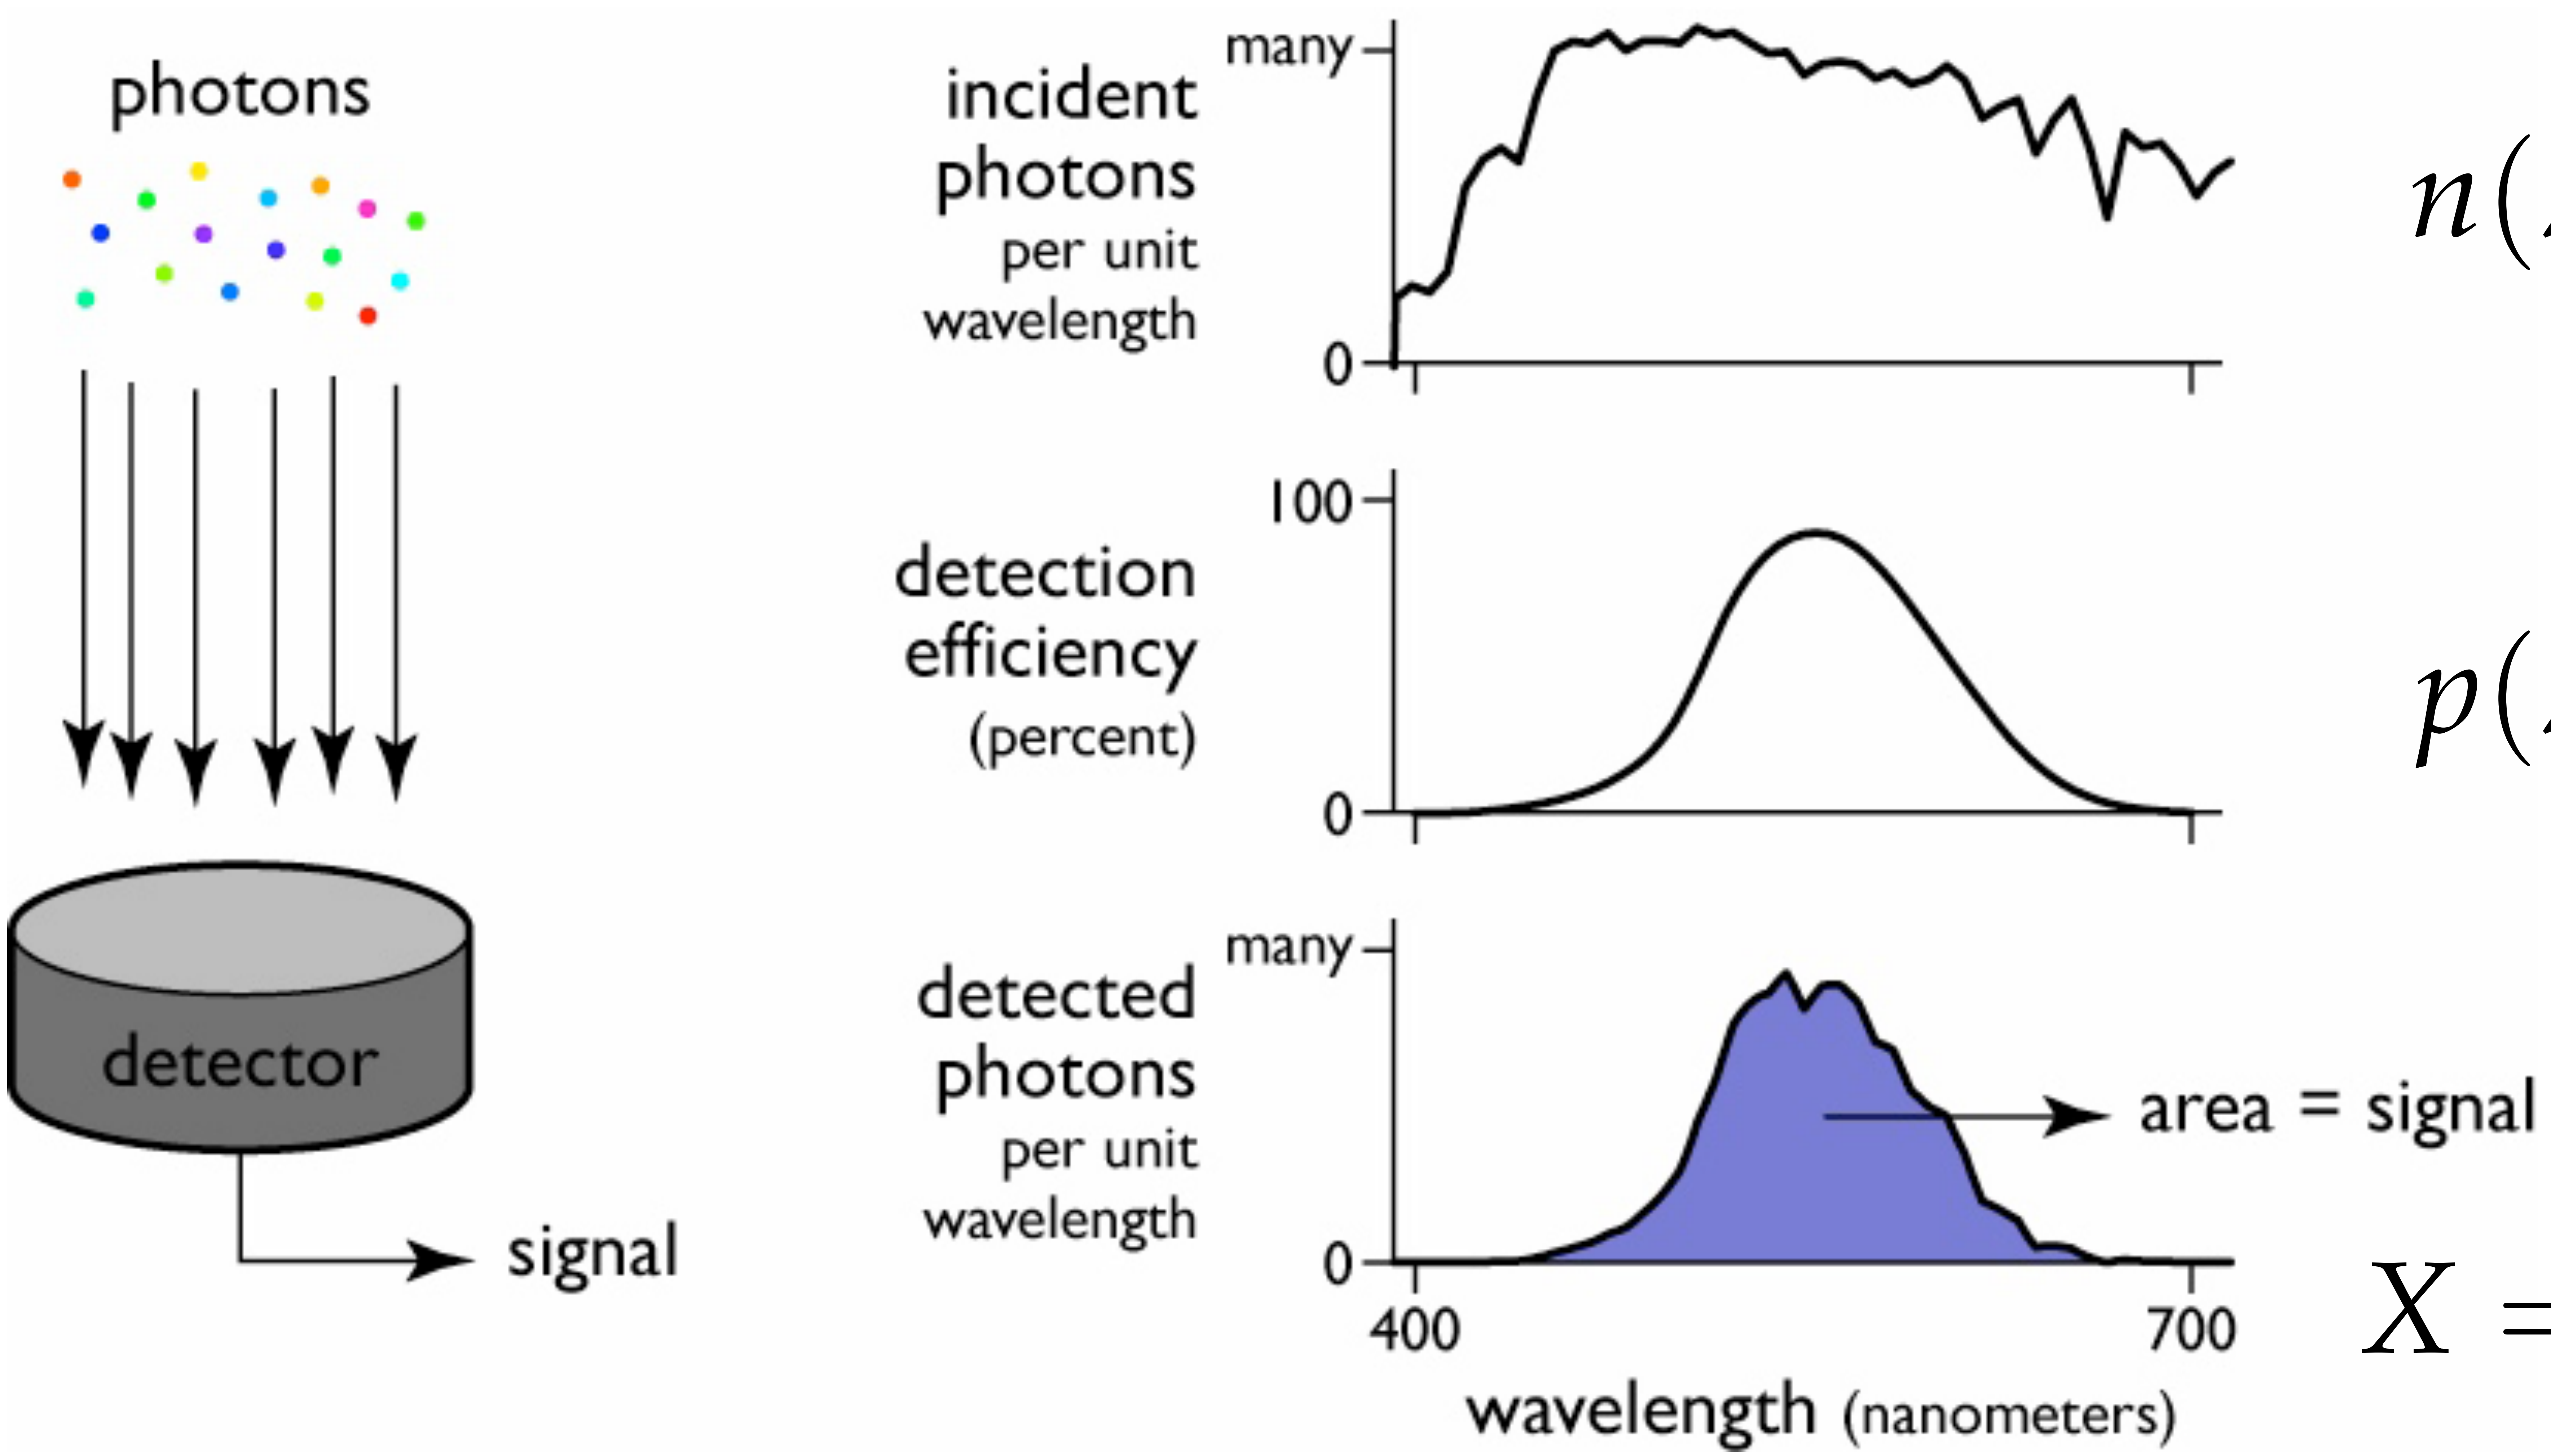

$n(\lambda)$

$p(\lambda)$

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

Light detection math

Same math carries over to spectral distributions

- spectrum entering the detector has its spectral power distribution (SPD), $s(\lambda)$
- detector has its spectral sensitivity or spectral response, $r(\lambda)$

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

- X: measured signal
- $s(\lambda)$: input spectrum
- $r(\lambda)$: detector's sensitivity

Cone responses

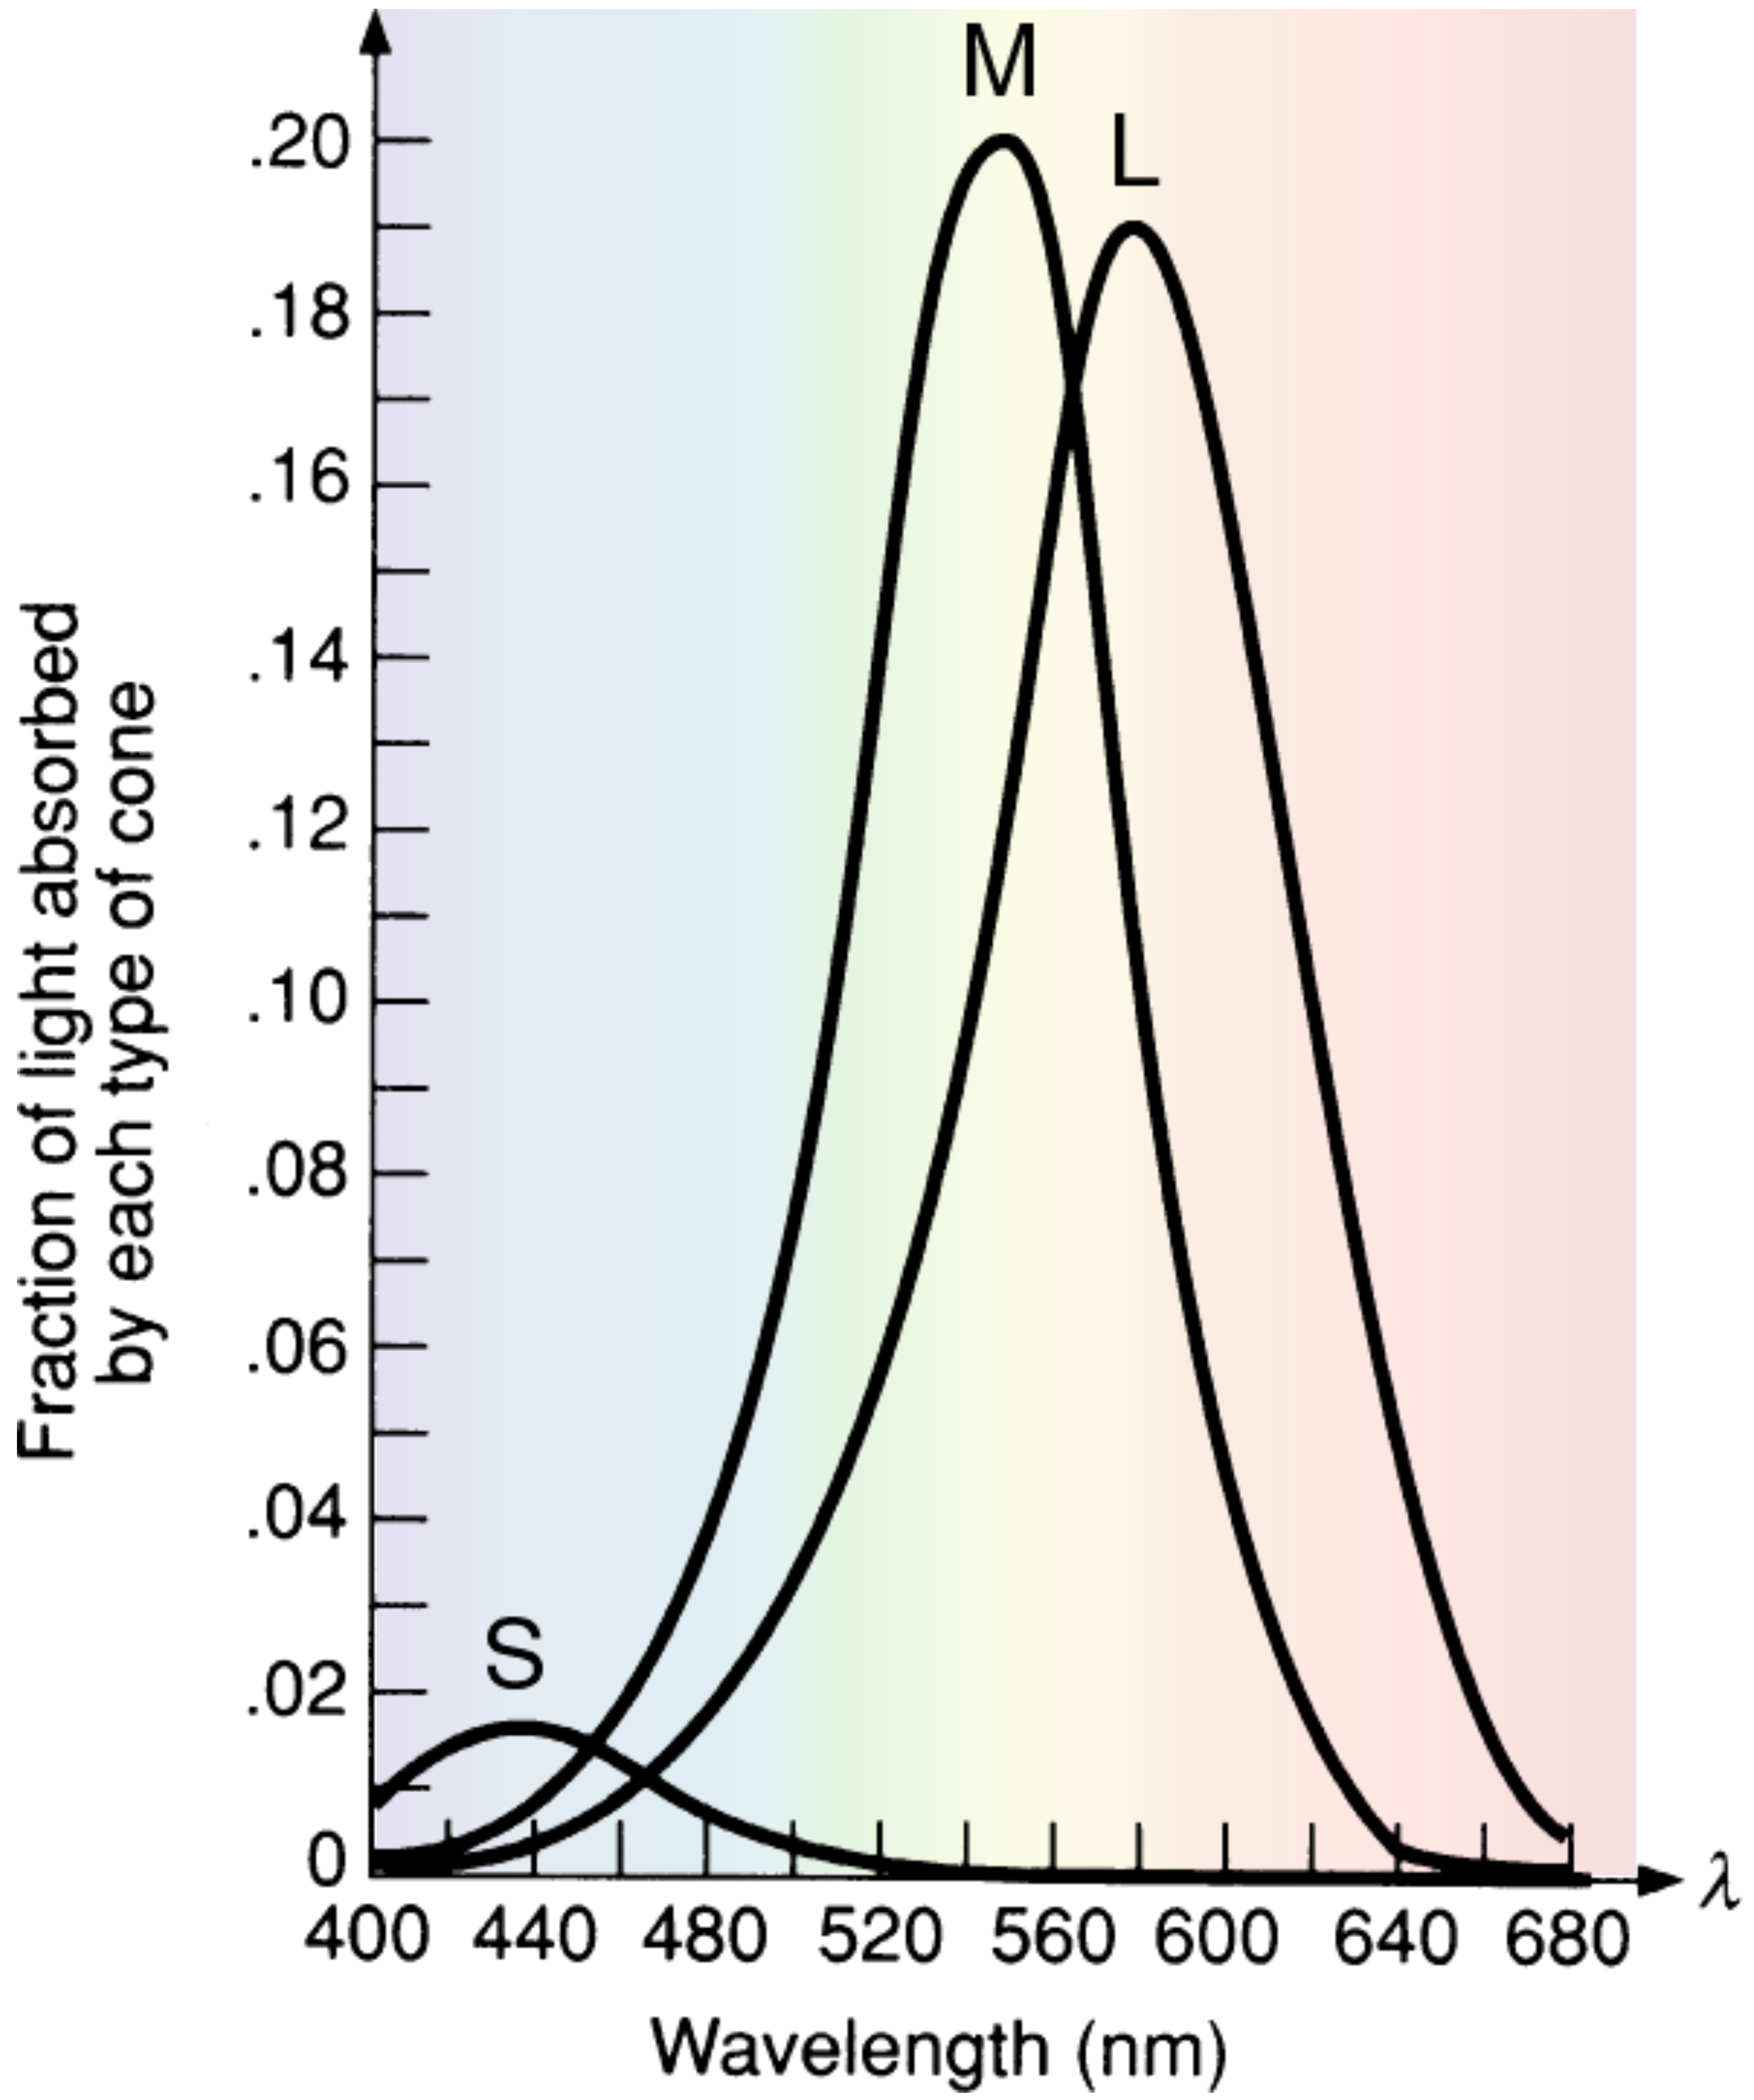

Three types of cones with broadband spectral sensitivity

- S cones respond to short-wavelengths ("blue")
- M cones respond to medium-wavelengths ("green")
- L cones respond to long-wavelengths ("red")
- Experimentally determined in the 1980s [link]

S,M,L neural response is integrated w.r.t. λ

- we'll call the response functions r_S, r_M, r_L

Results in a trichromatic visual system

S, M, and L are *tristimulus values*

Cone responses to a spectrum s (Math)

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

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

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

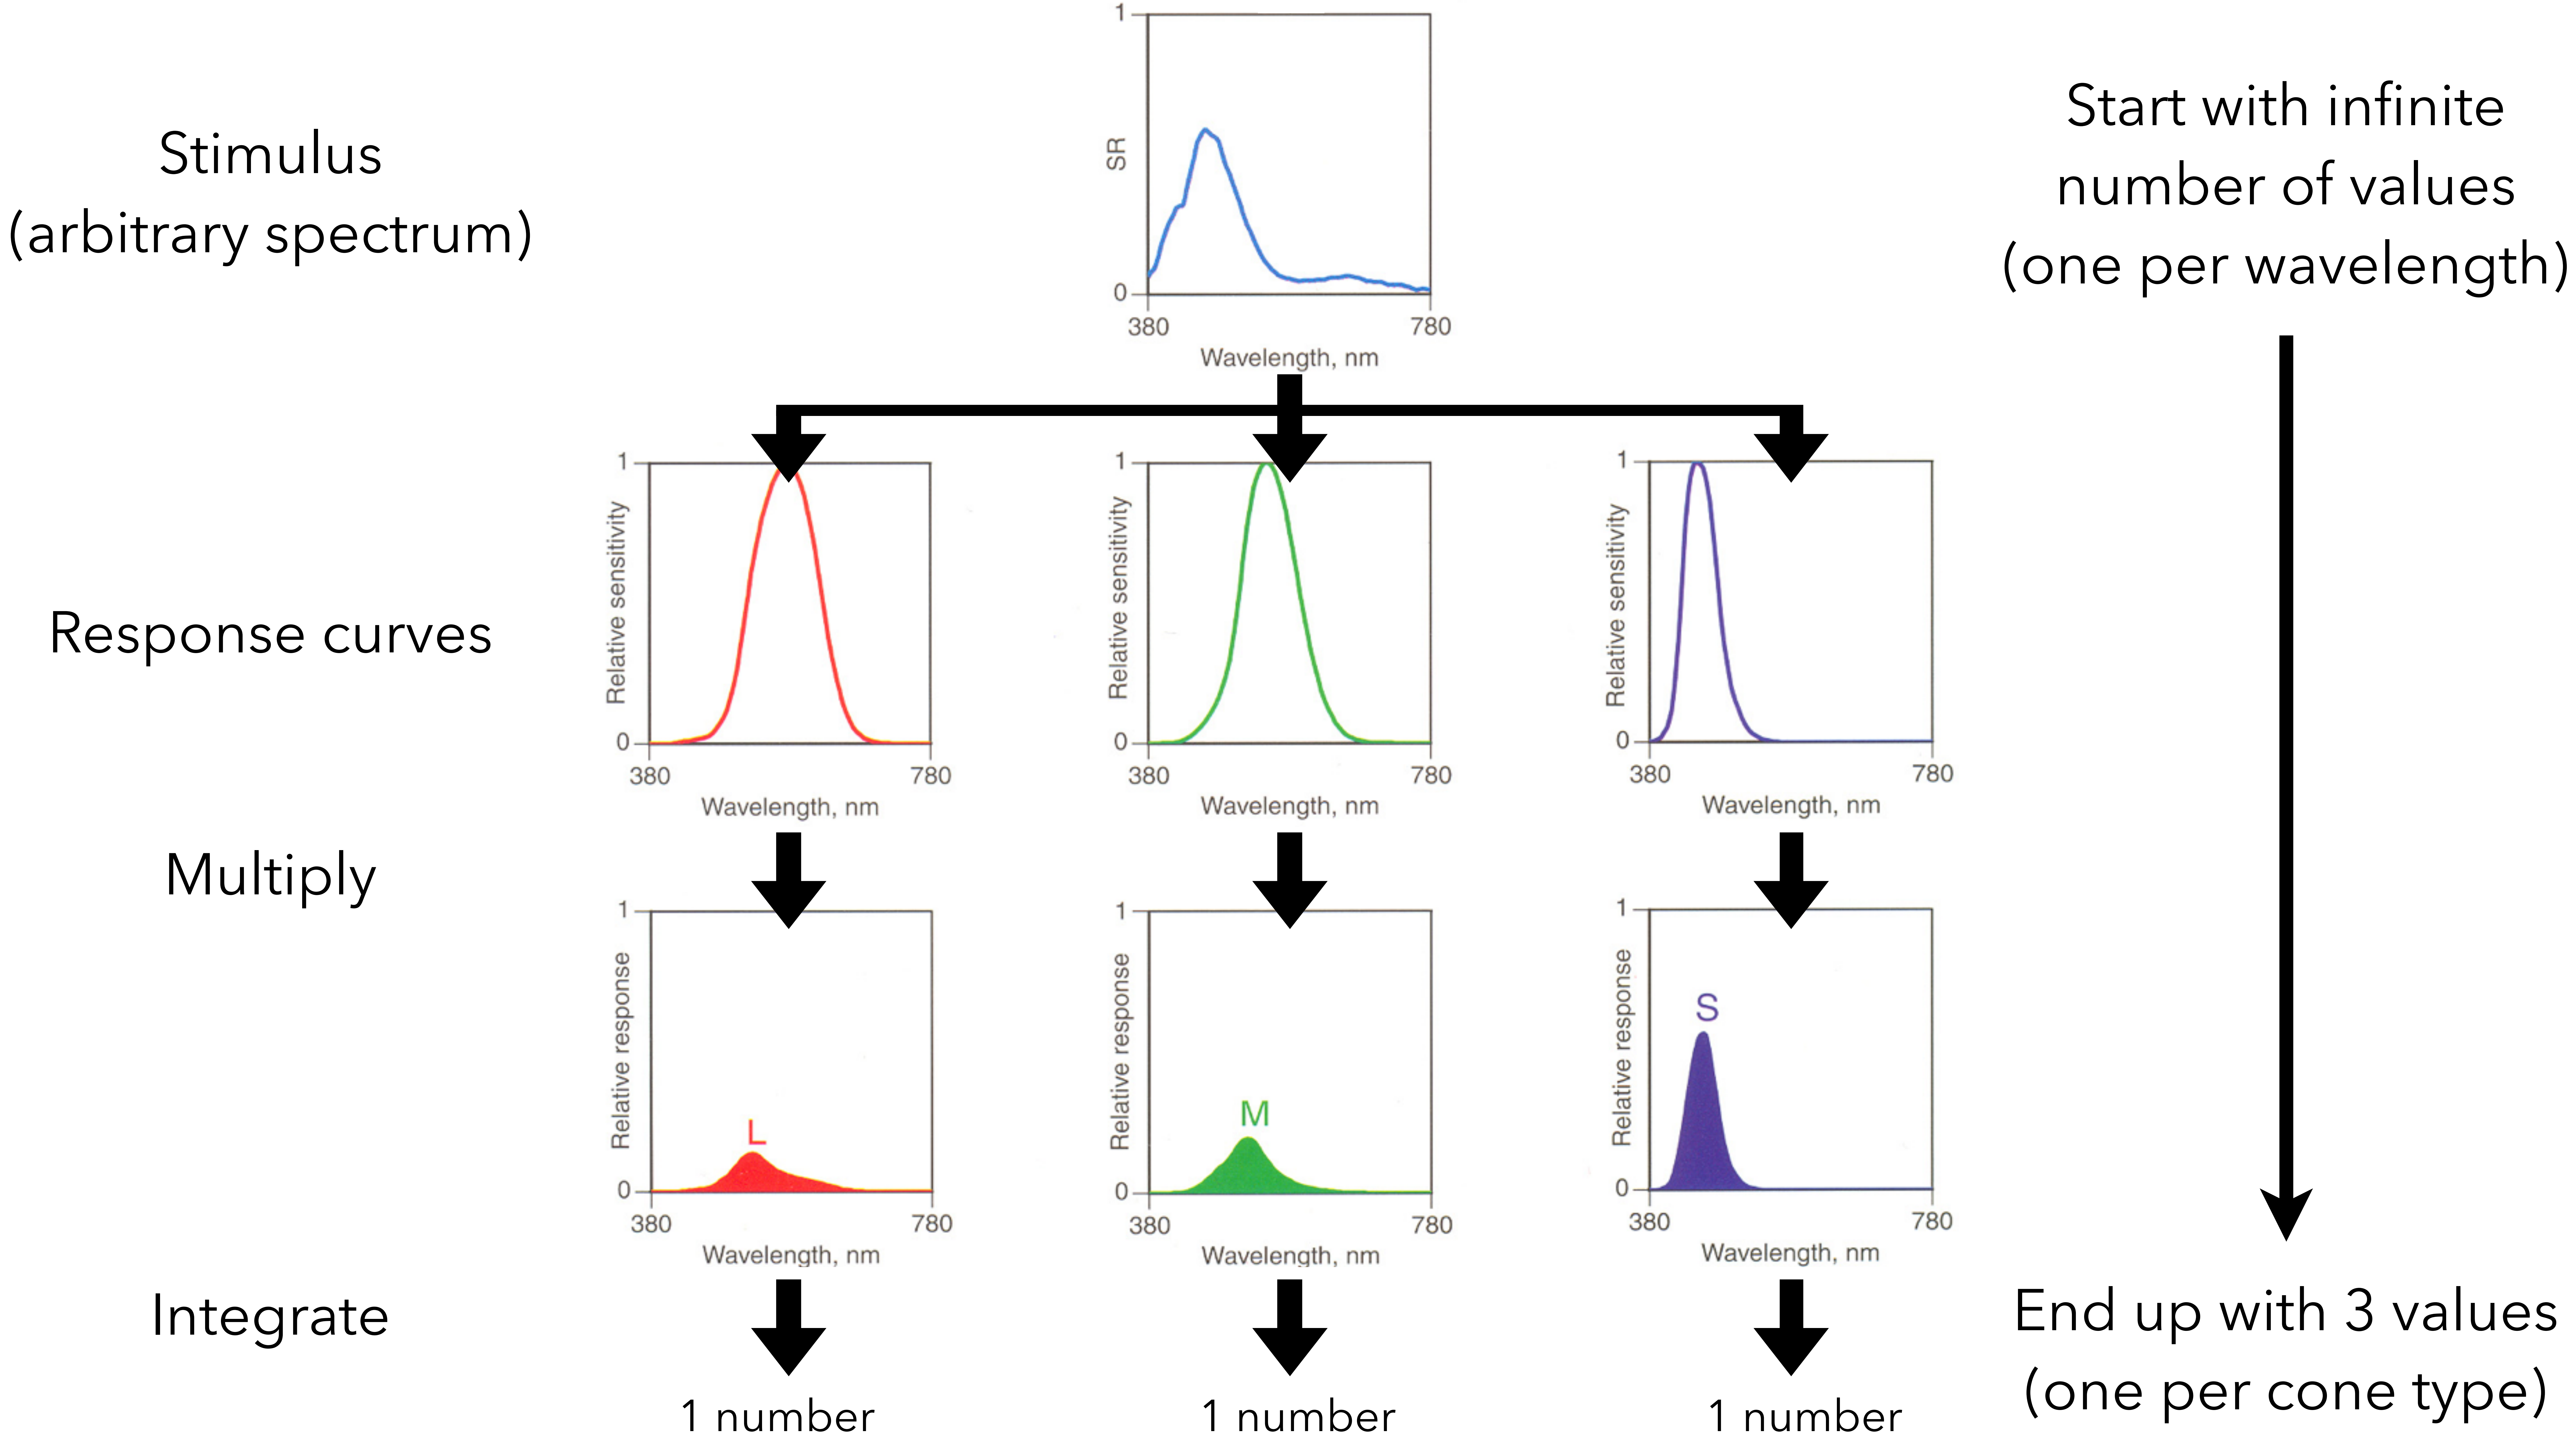
Stimulus
(arbitrary spectrum)
Response curves
Multiply
Integrate
SR
Relative sensitivity
Relative response
Wavelength, nm
380
780
0
1
L
M
S
1 number
1 number
1 number
Start with infinite number of values (one per wavelength)
End up with 3 values (one per cone type)

Linear algebra interpretation

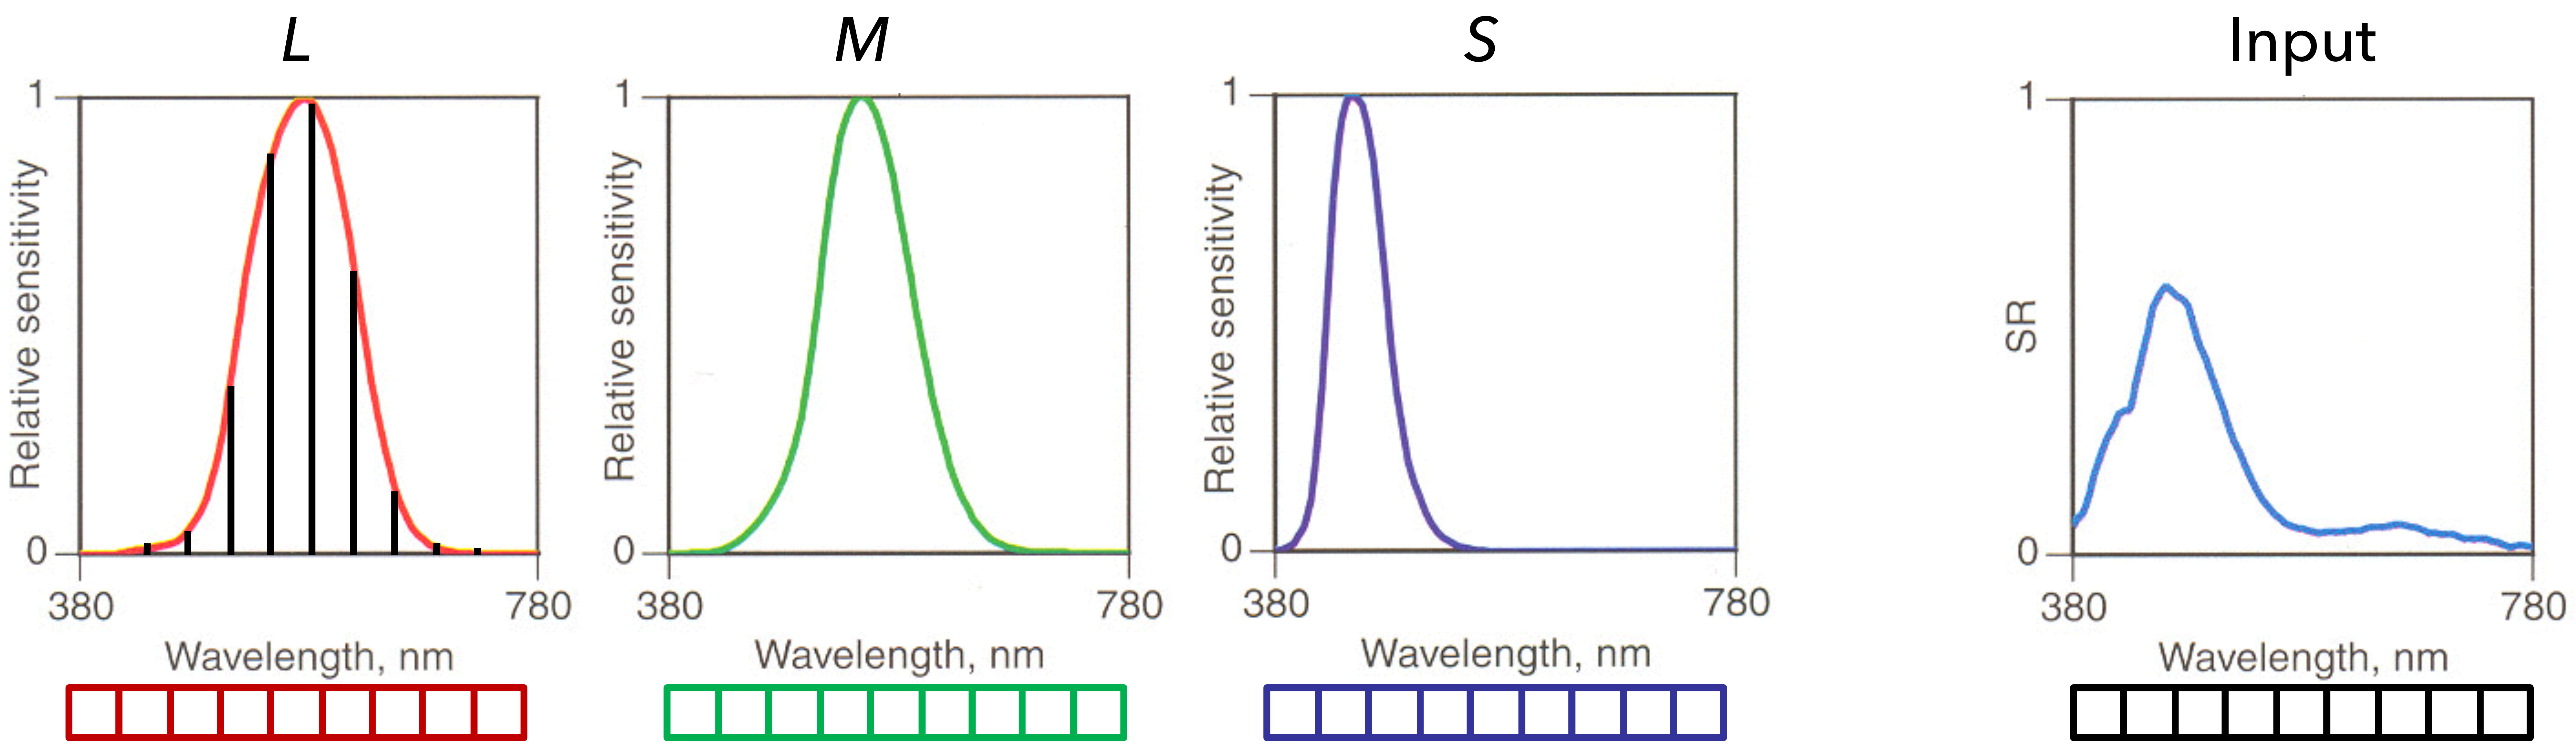

Discrete representation of cones and input spectrum as vectors

Linear algebra interpretation

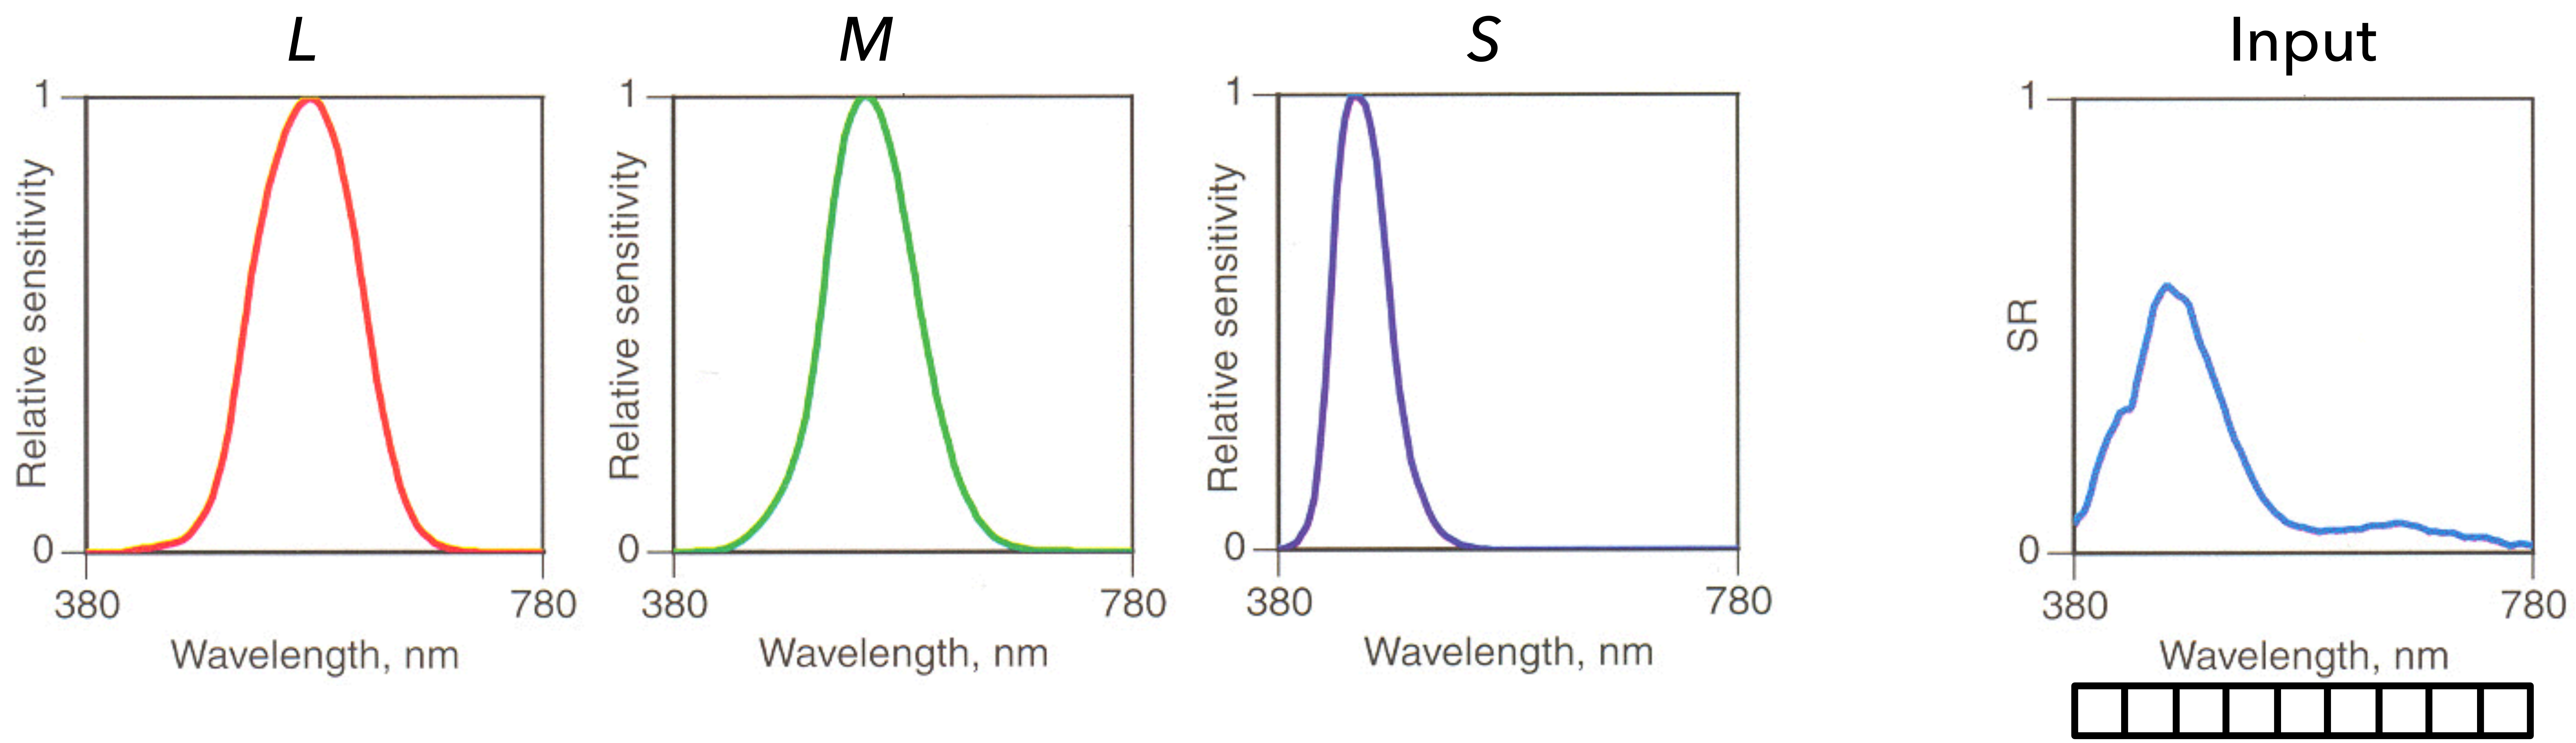

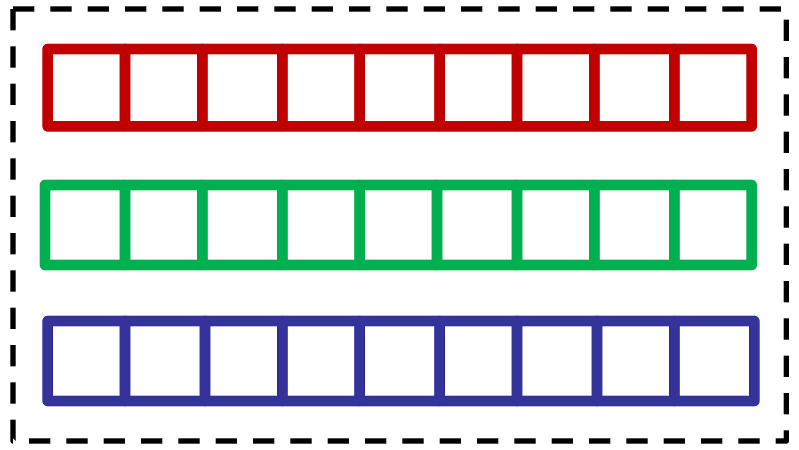

cone sensitivities

Linear algebra interpretation

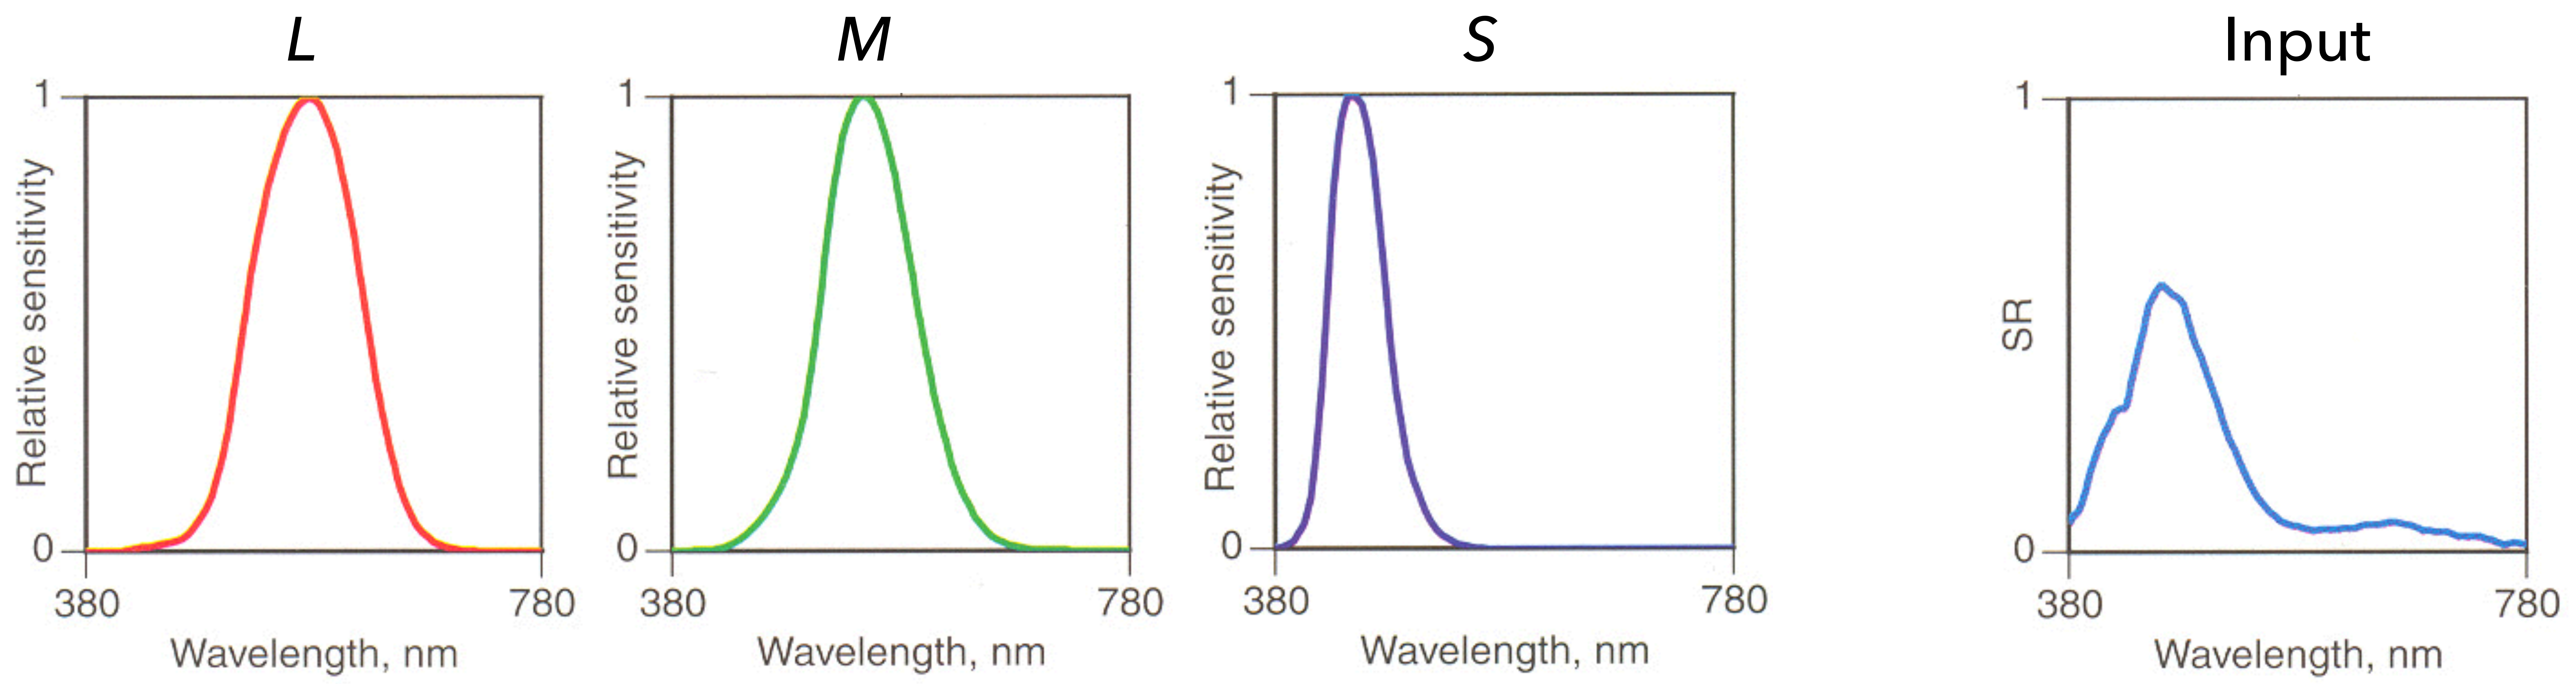

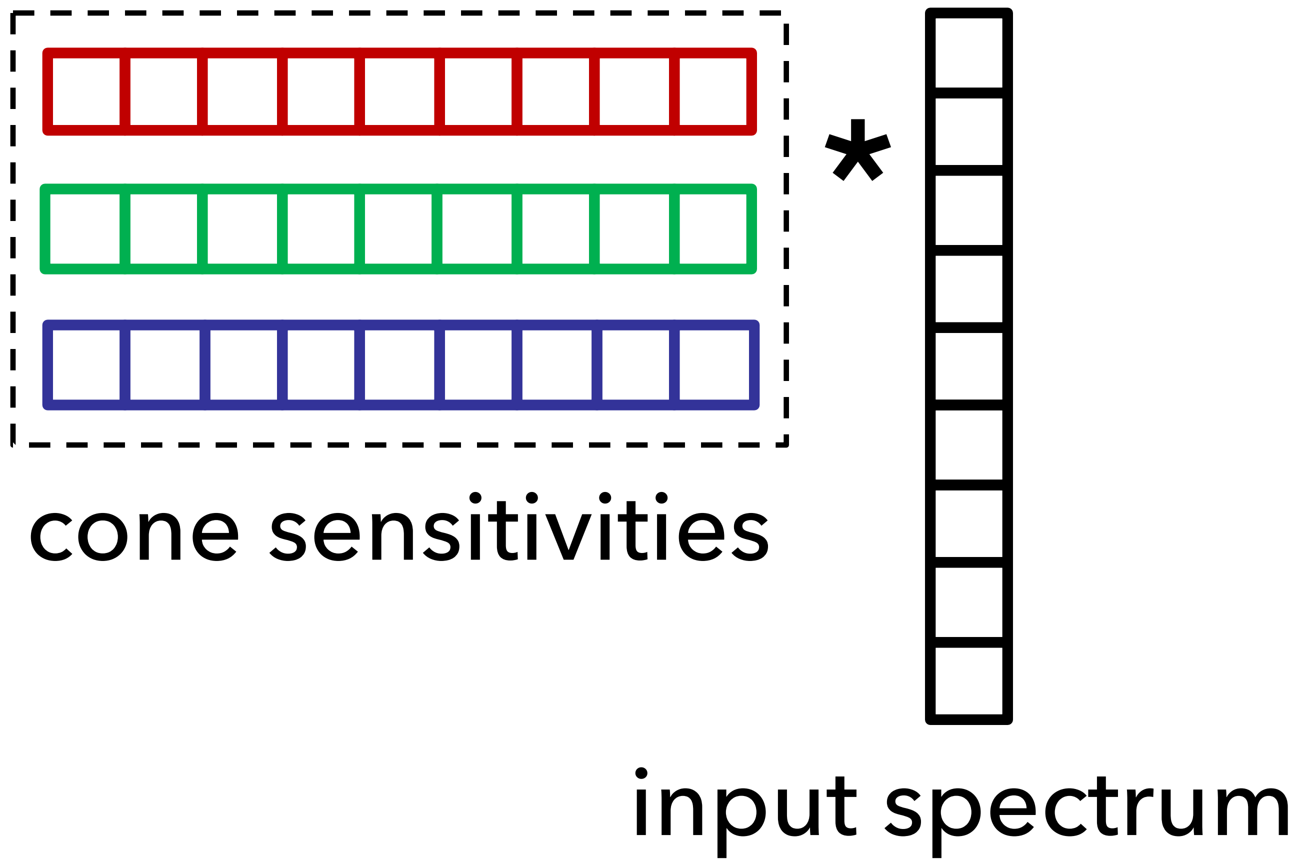

Linear algebra interpretation

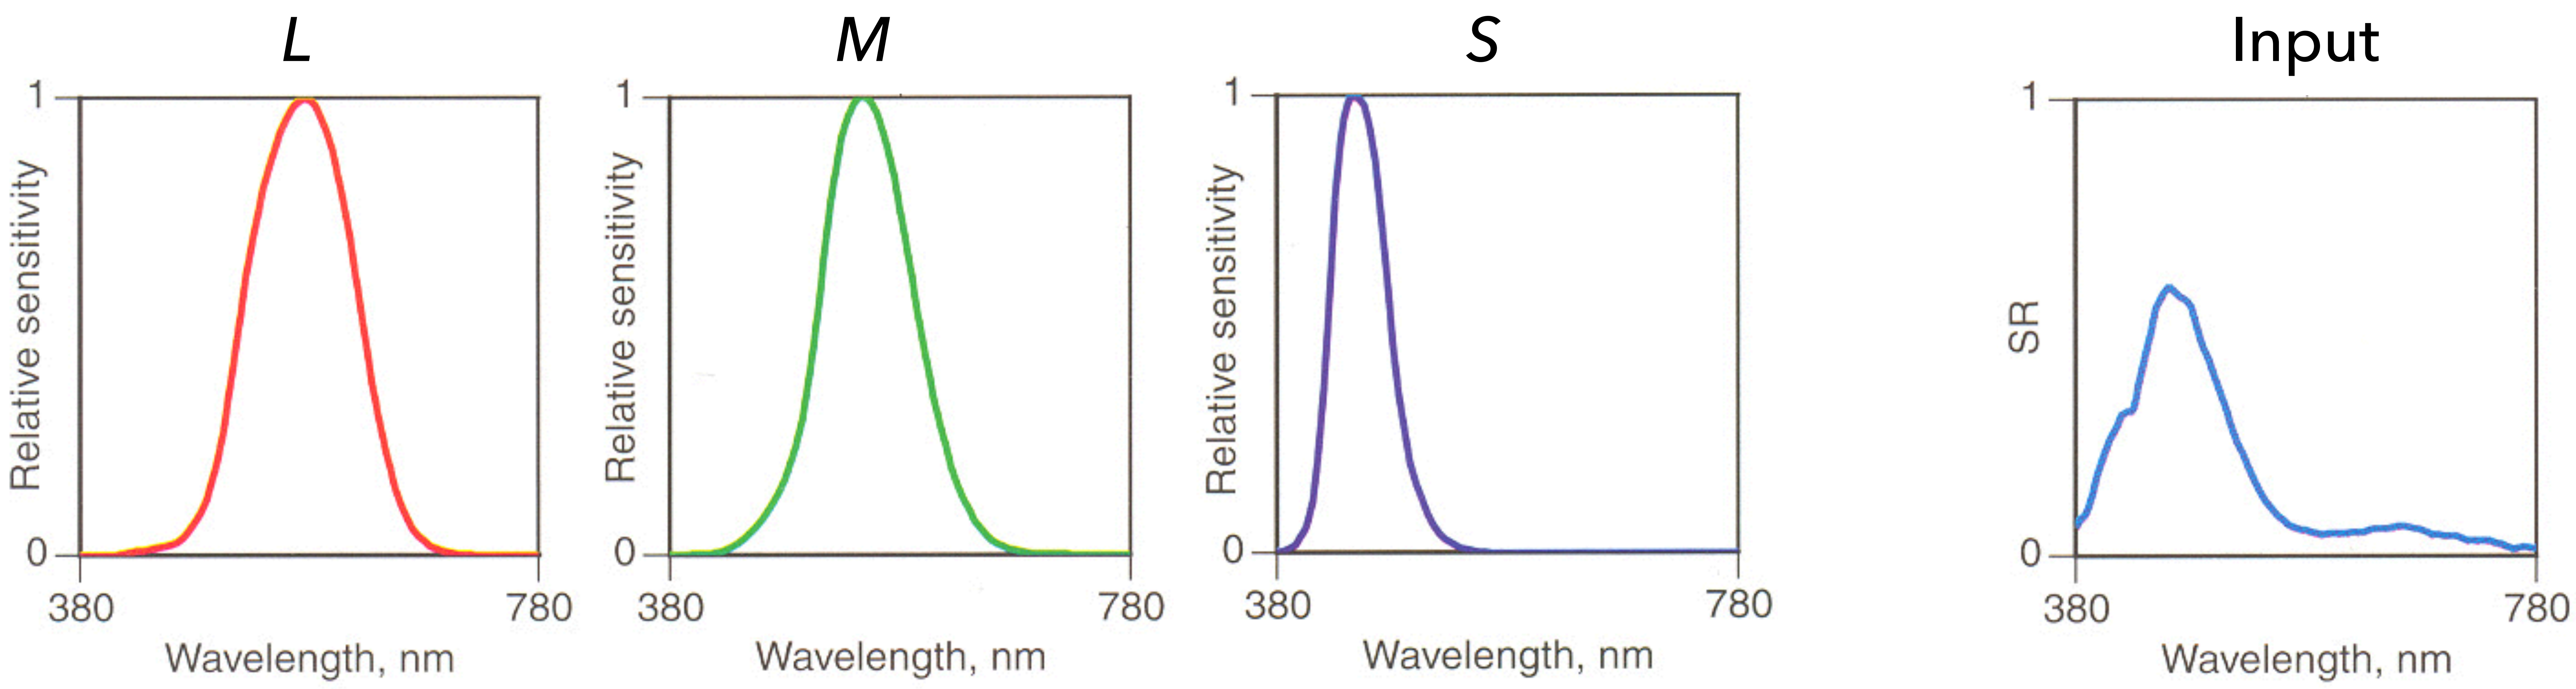

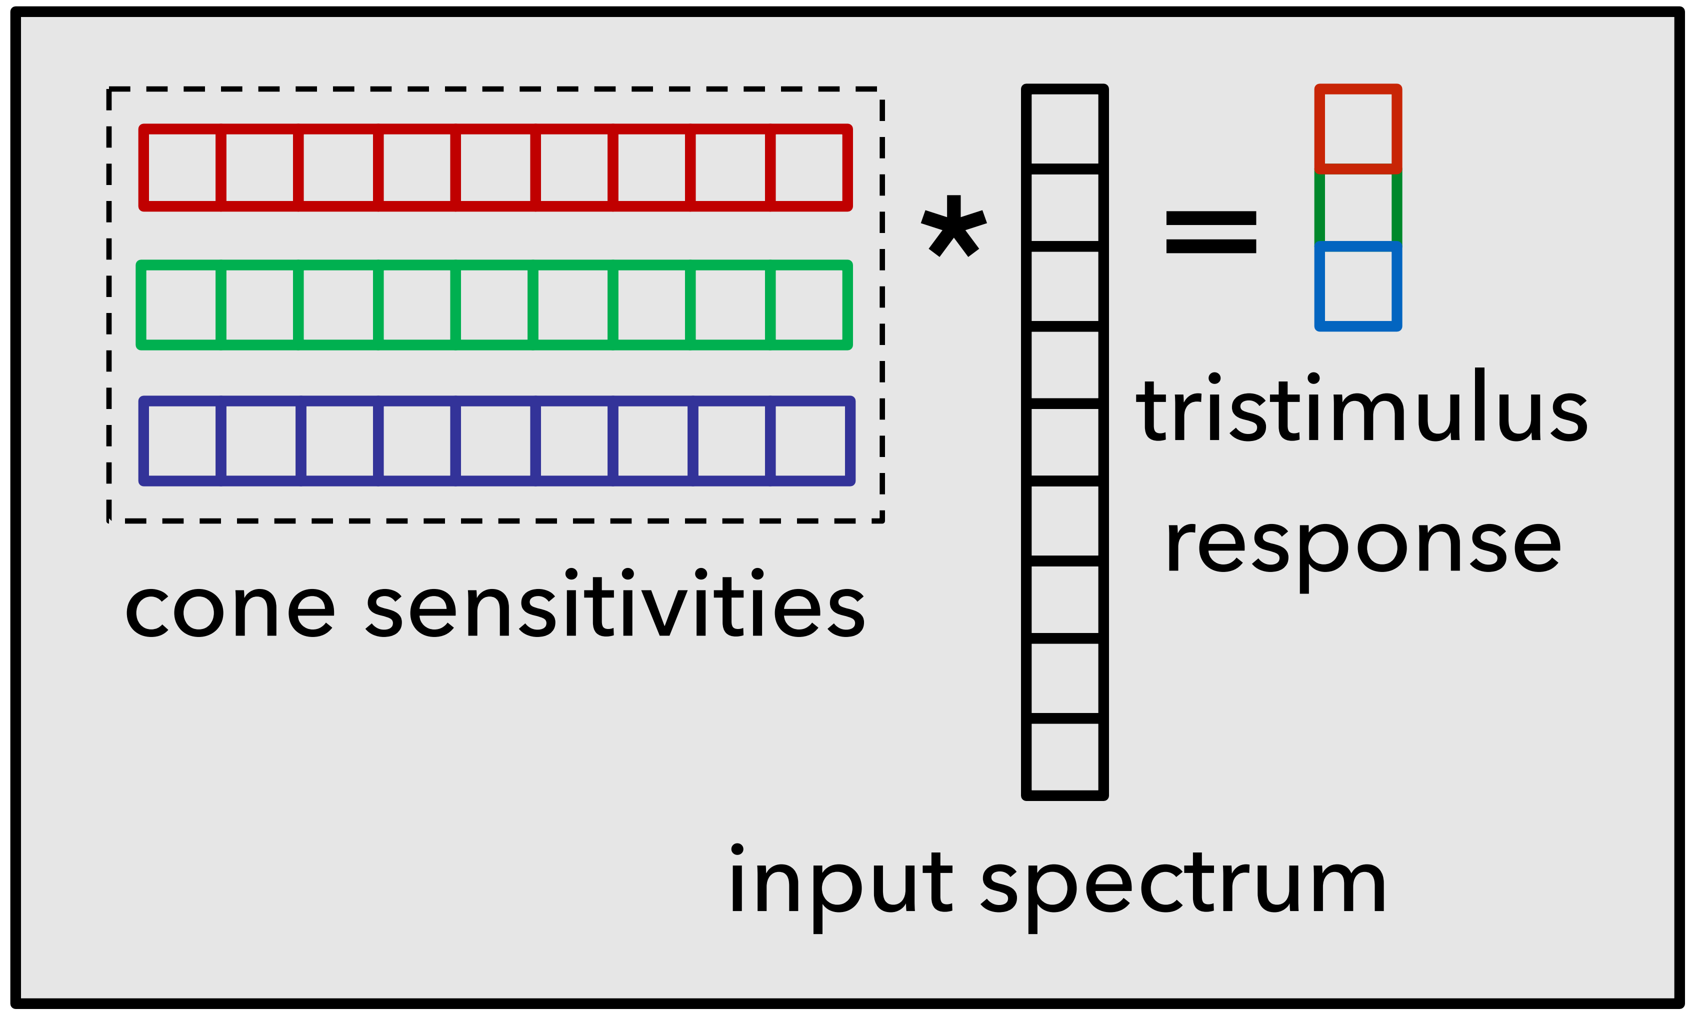

Tristimulus response is a matrix-vector multiplication

Integration is now summation

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:

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

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

Colorimetry: an answer to the problem

Wanted to map a *physical light description* to a *perceptual color sensation*

Basic solution was known and standardized by 1930

- Though not quite in this form – more on that later

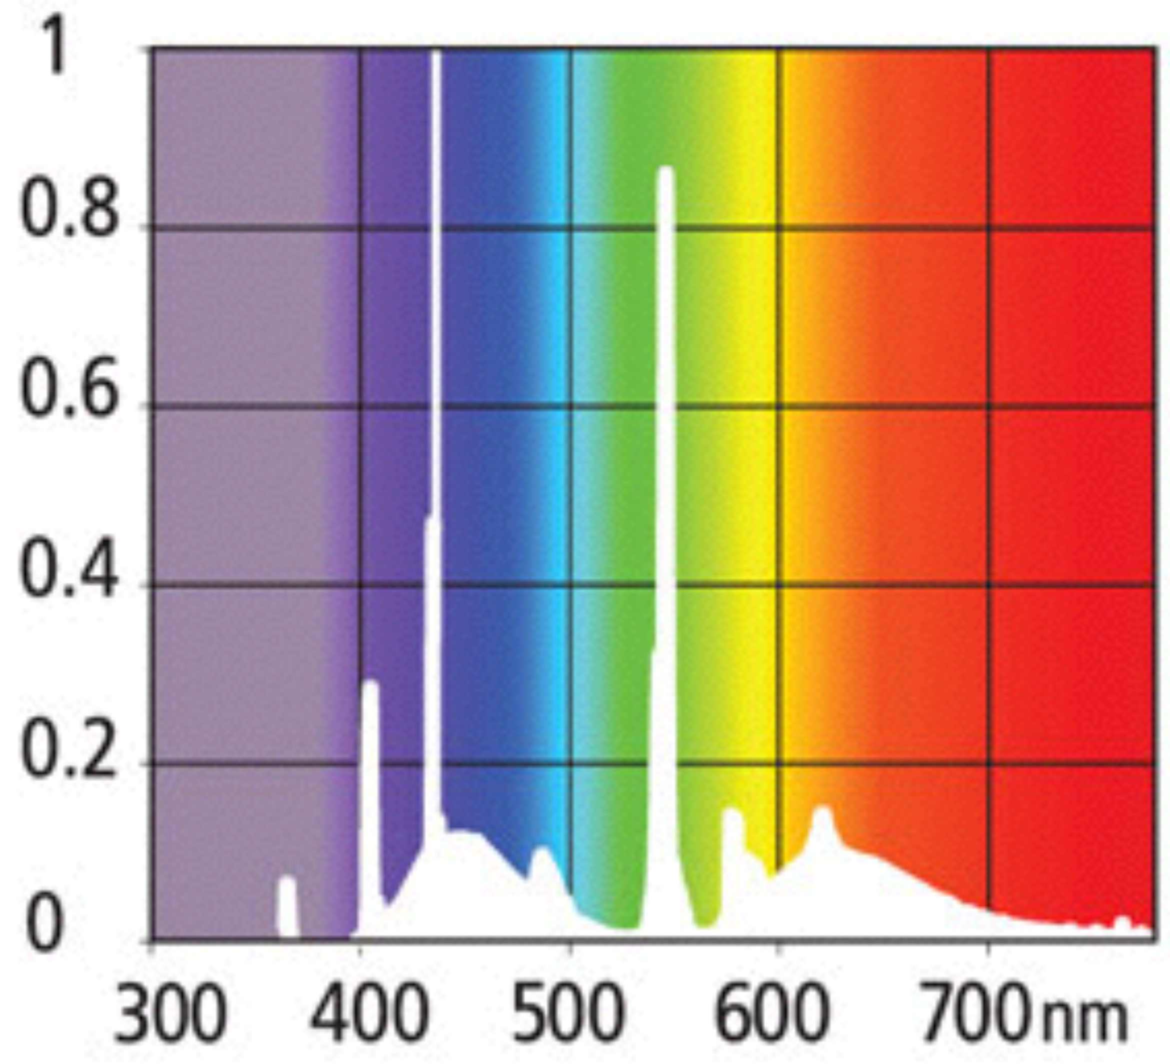

Physical

→

$$S = r_S \cdot s$$
$$M = r_M \cdot s$$
$$L = r_L \cdot s$$

Perceptual

Basic colorimetric concepts

Luminance

- the overall magnitude of the the visual response to a spectrum (independent of its color)
    - corresponds to the everyday concept “brightness”
- determined by product of SPD with the *luminous efficiency function* V_λ that describes the eye’s overall ability to detect light at each wavelength
- e.g. lamps are optimized to improve their luminous efficiency (tungsten vs. fluorescent vs. sodium vapor)

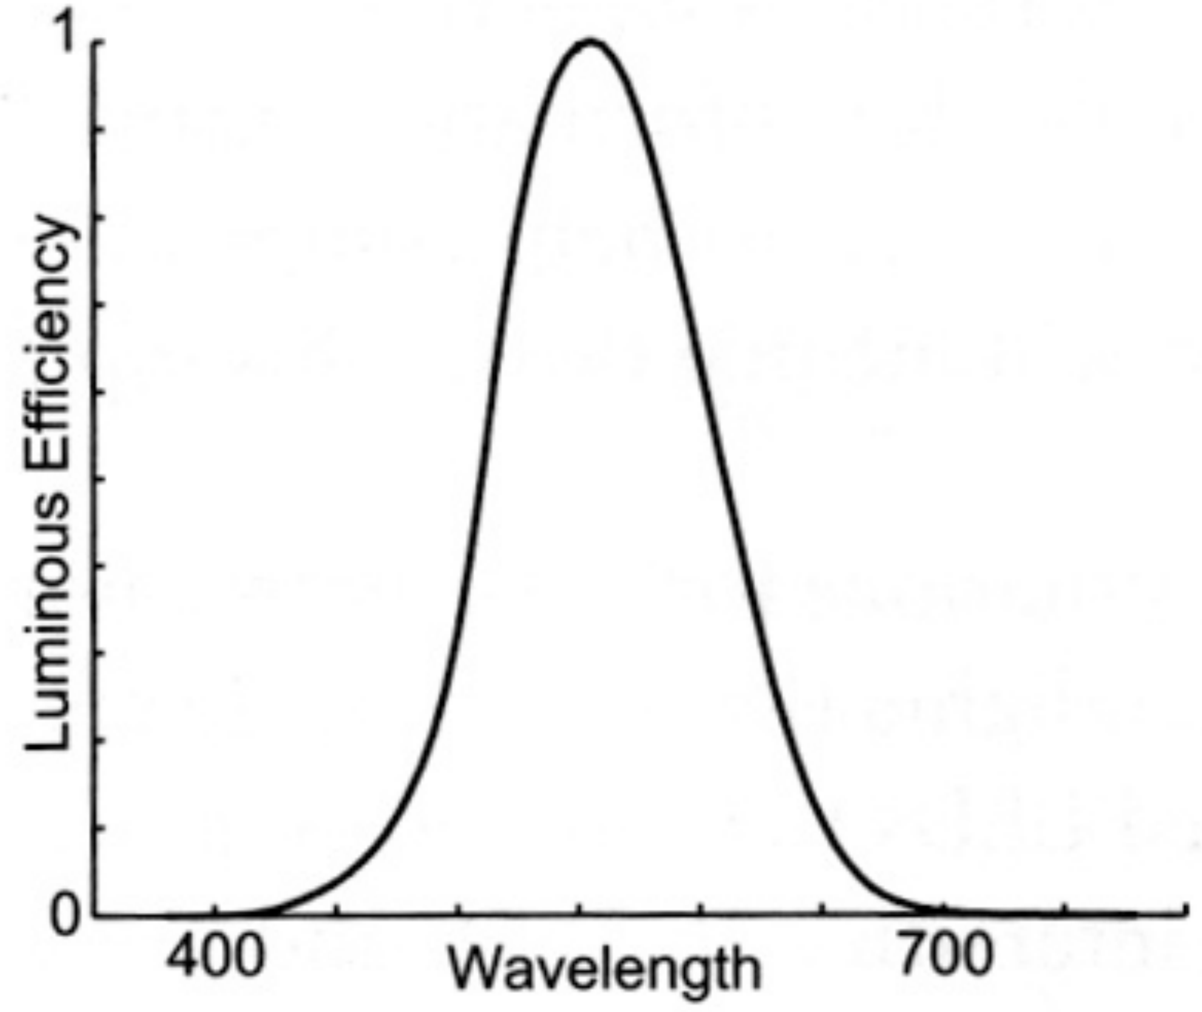

Luminance, mathematically

Y: just another response curve (like *S*, *M*, and *L*)

$$Y = r_Y \cdot s$$

- r_Y is really called V_λ

V_λ is a linear combination of *S*, *M*, and *L*

- has to be, since it's derived from cone outputs

More basic colorimetric concepts

Chromaticity

- what’s left after luminance is factored out (the color without regard for overall brightness)
- scaling a spectrum up or down leaves chromaticity alone

A cone does not “see” colors

Different wavelength, different intensity

Same response

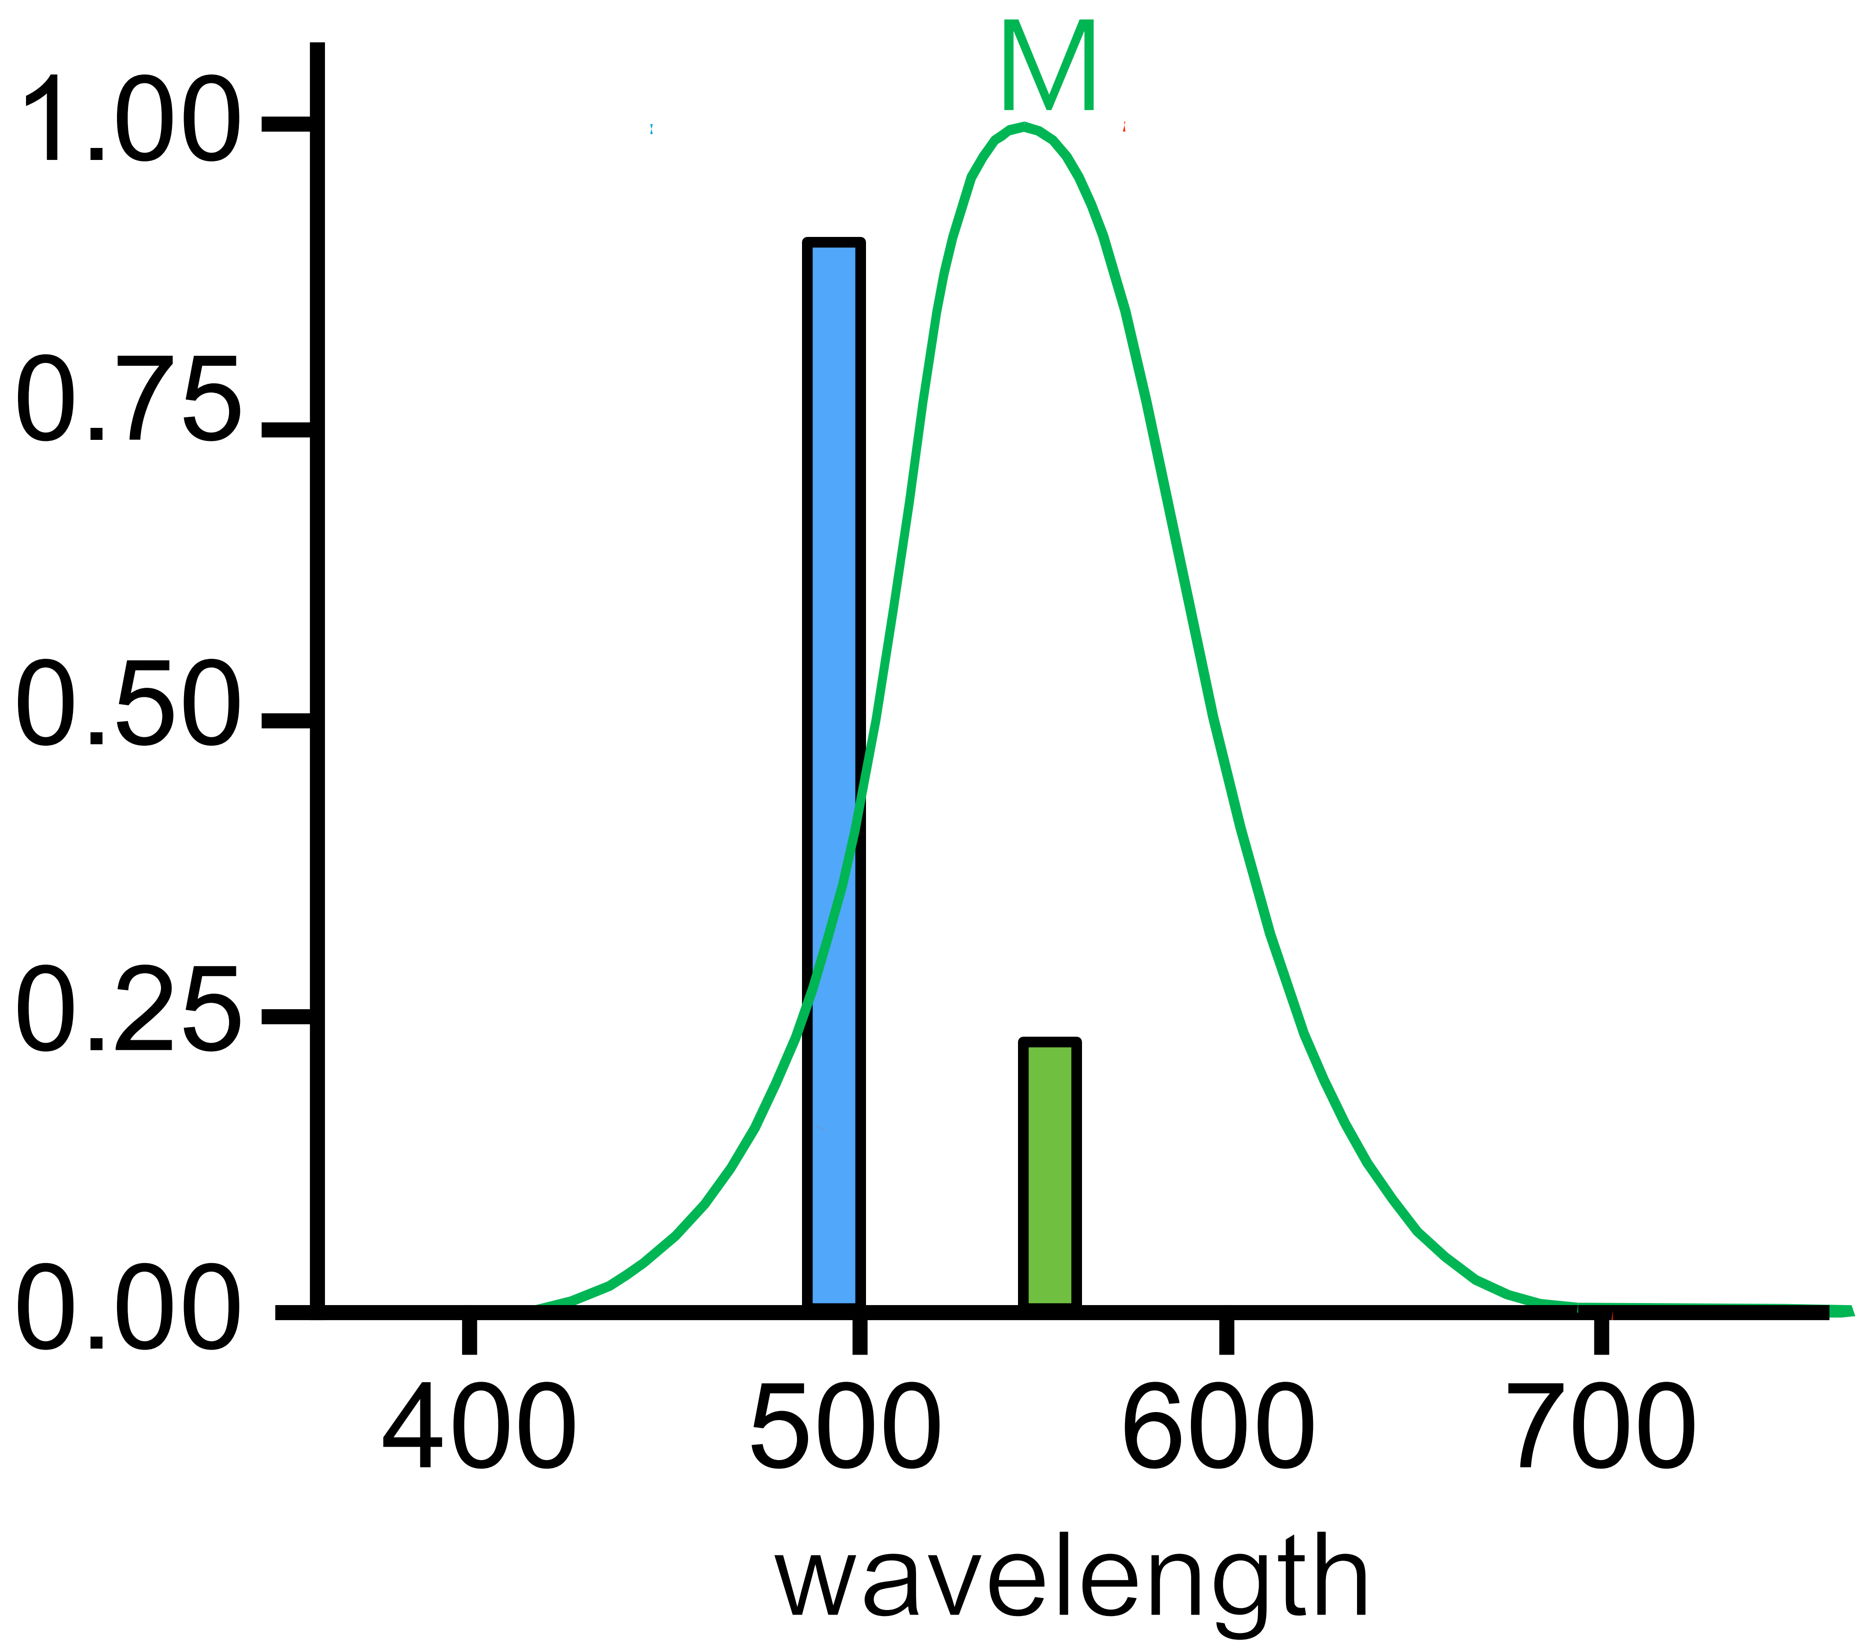

Response comparison

Different wavelength, different intensity

But different response for different cones

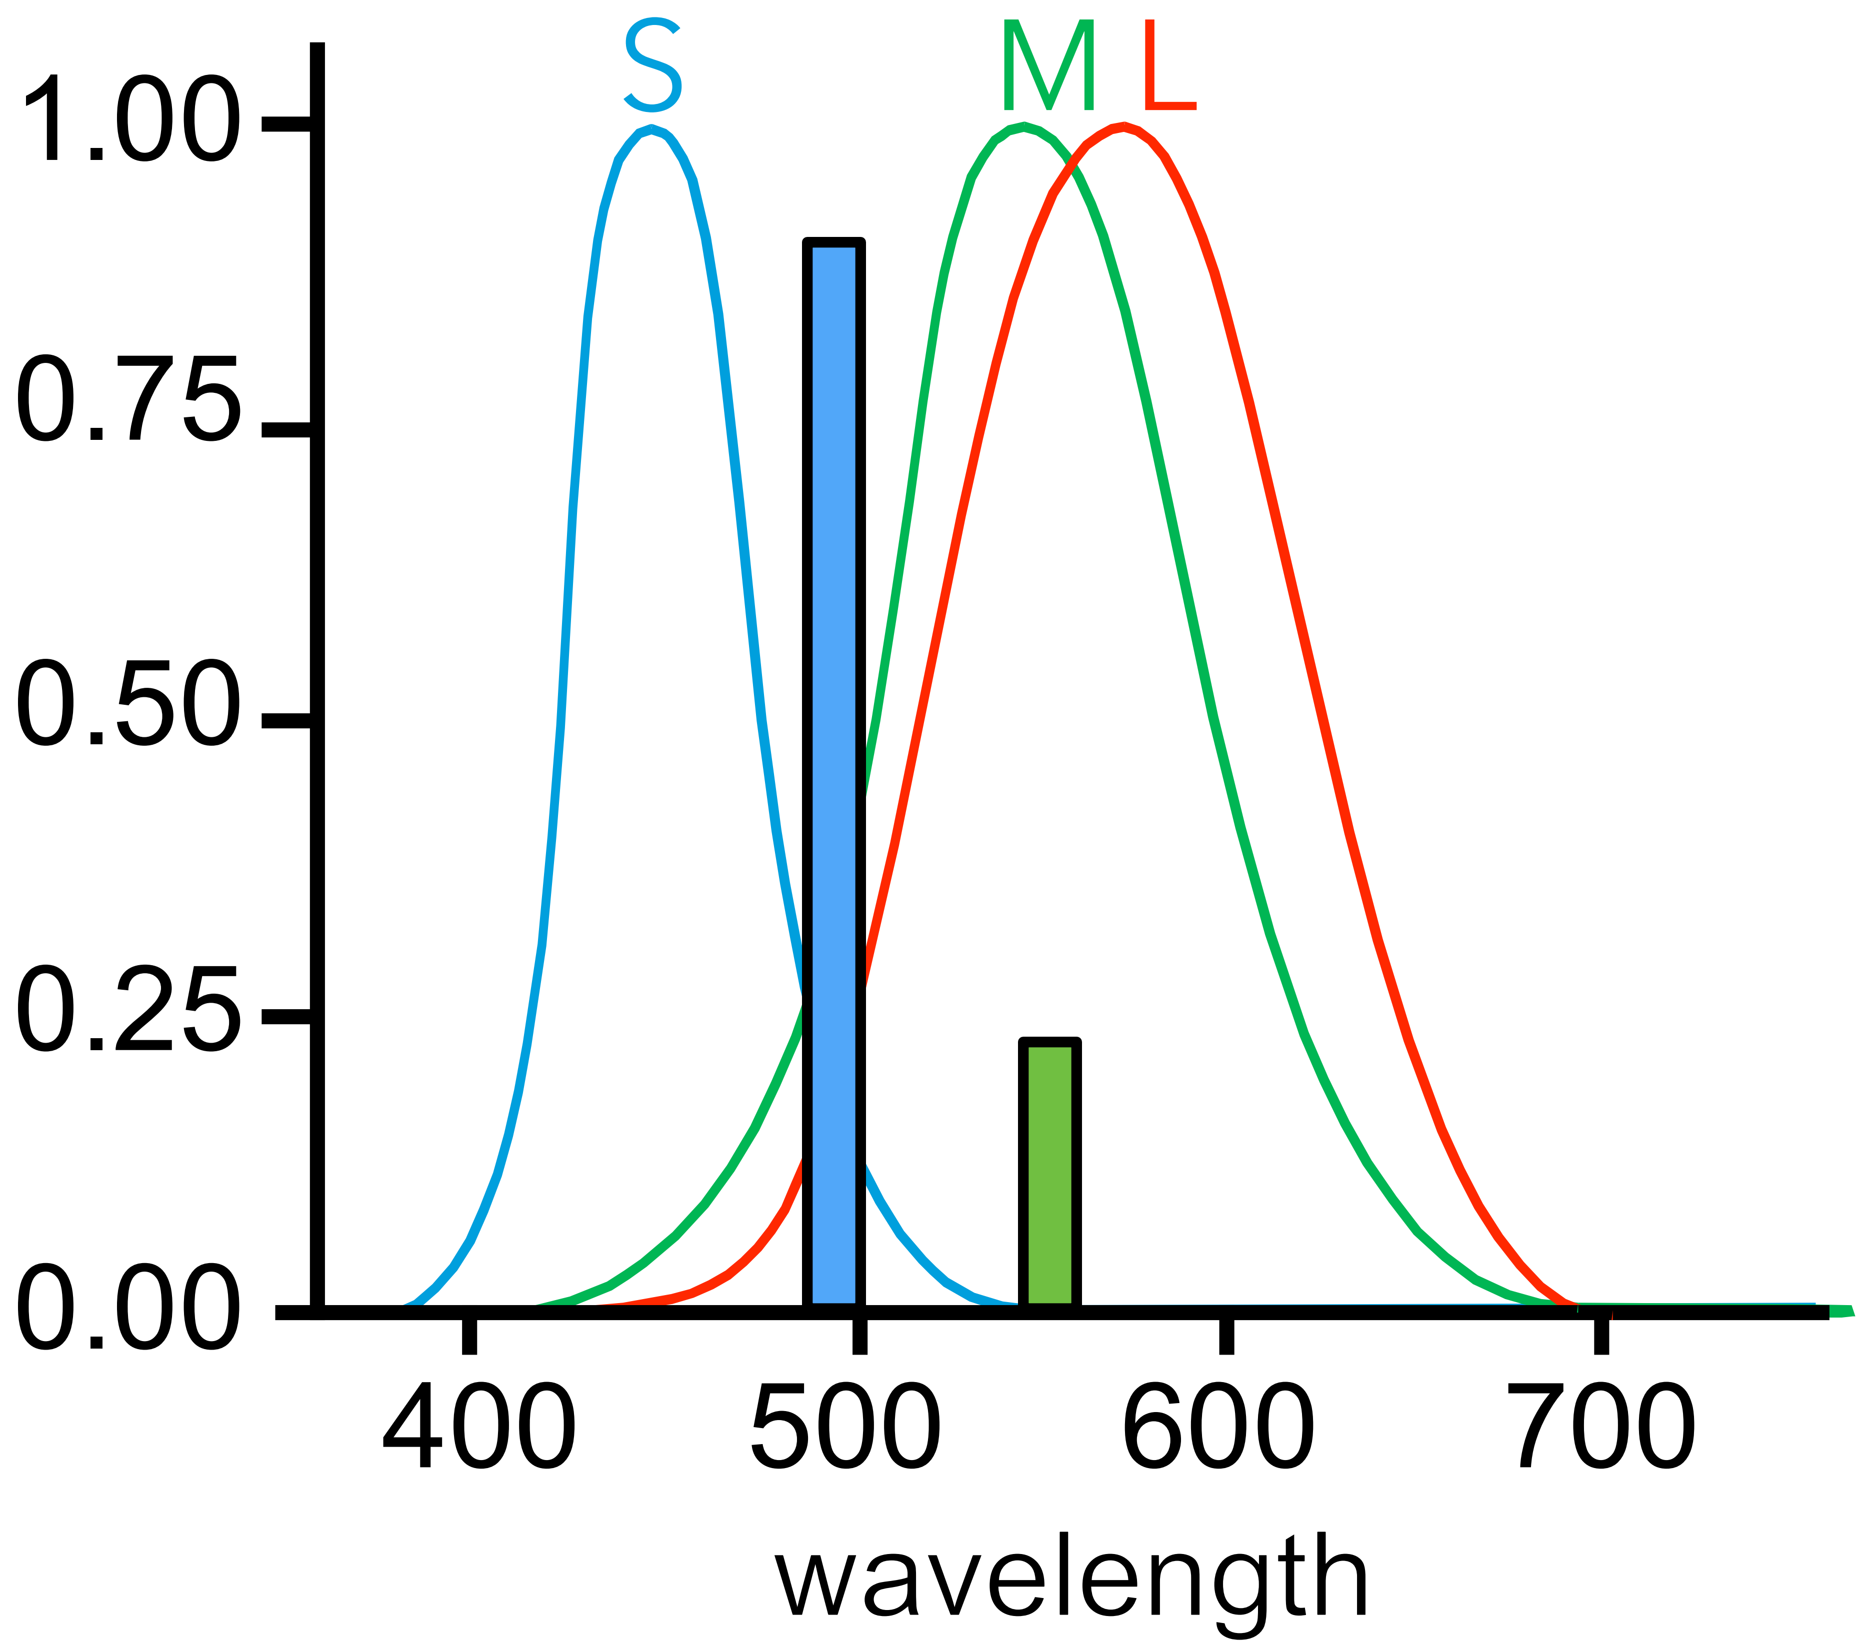

Color blindness

Color blindness

Classical case: 1 type of cone is missing (e.g. L)

Makes it impossible to distinguish some spectra

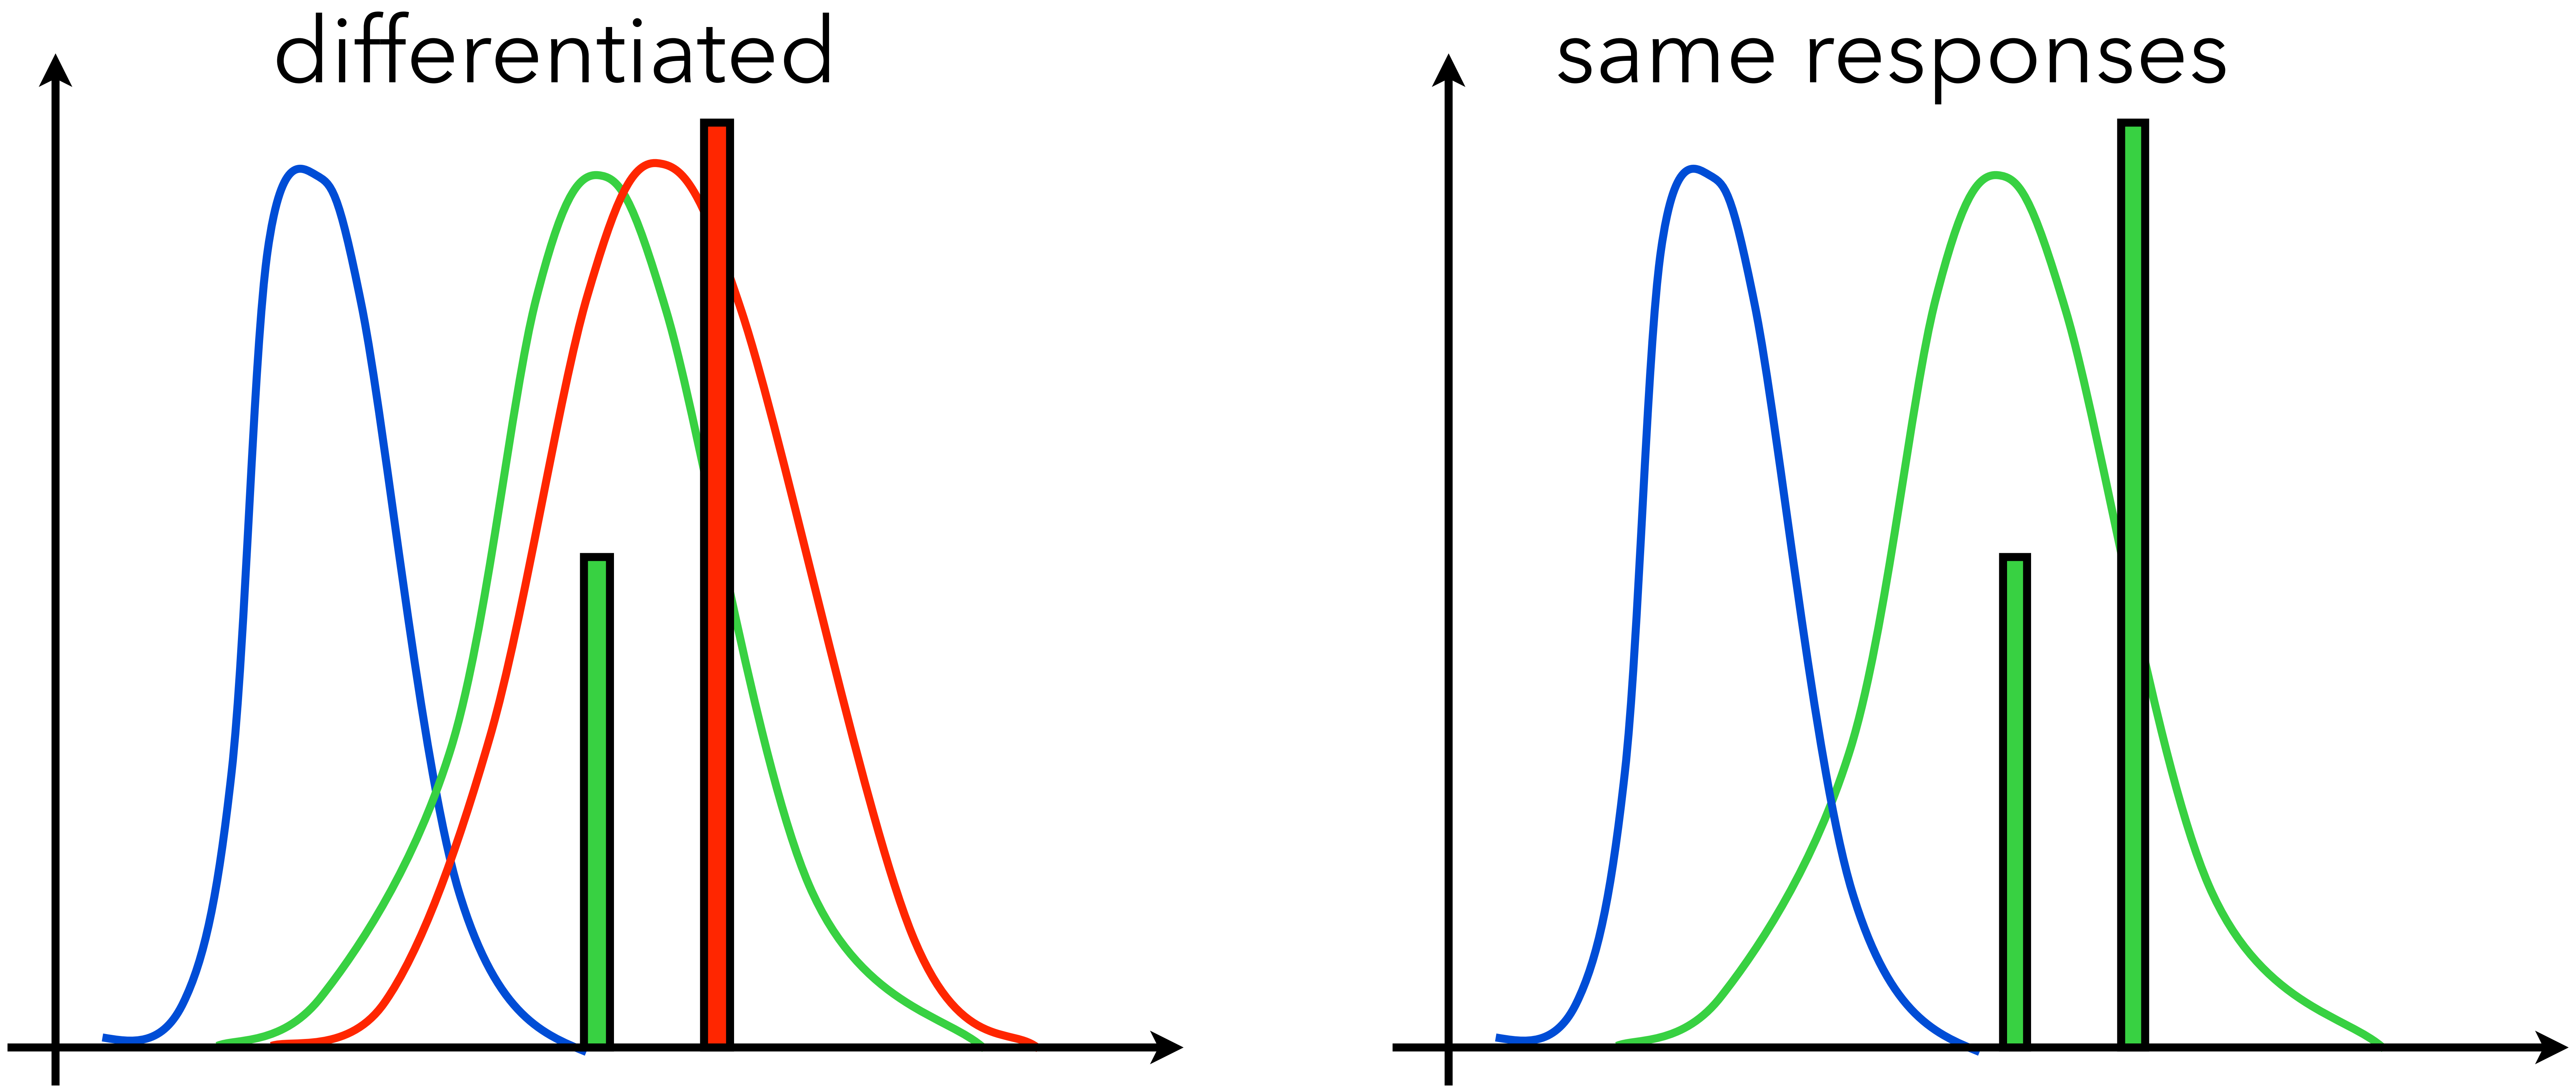

Color blindness – more general

8% male, 0.6% female

Genetic

Dichromate (strong color blind) – 2% male

- One type of cone missing
- L (protanope), M (deuteranope), S (tritanope)

Anomalous trichromat (weak color blind)

- Shifted sensitivity

More at, e.g. http://en.wikipedia.org/wiki/Color_blindness

Color blindness test

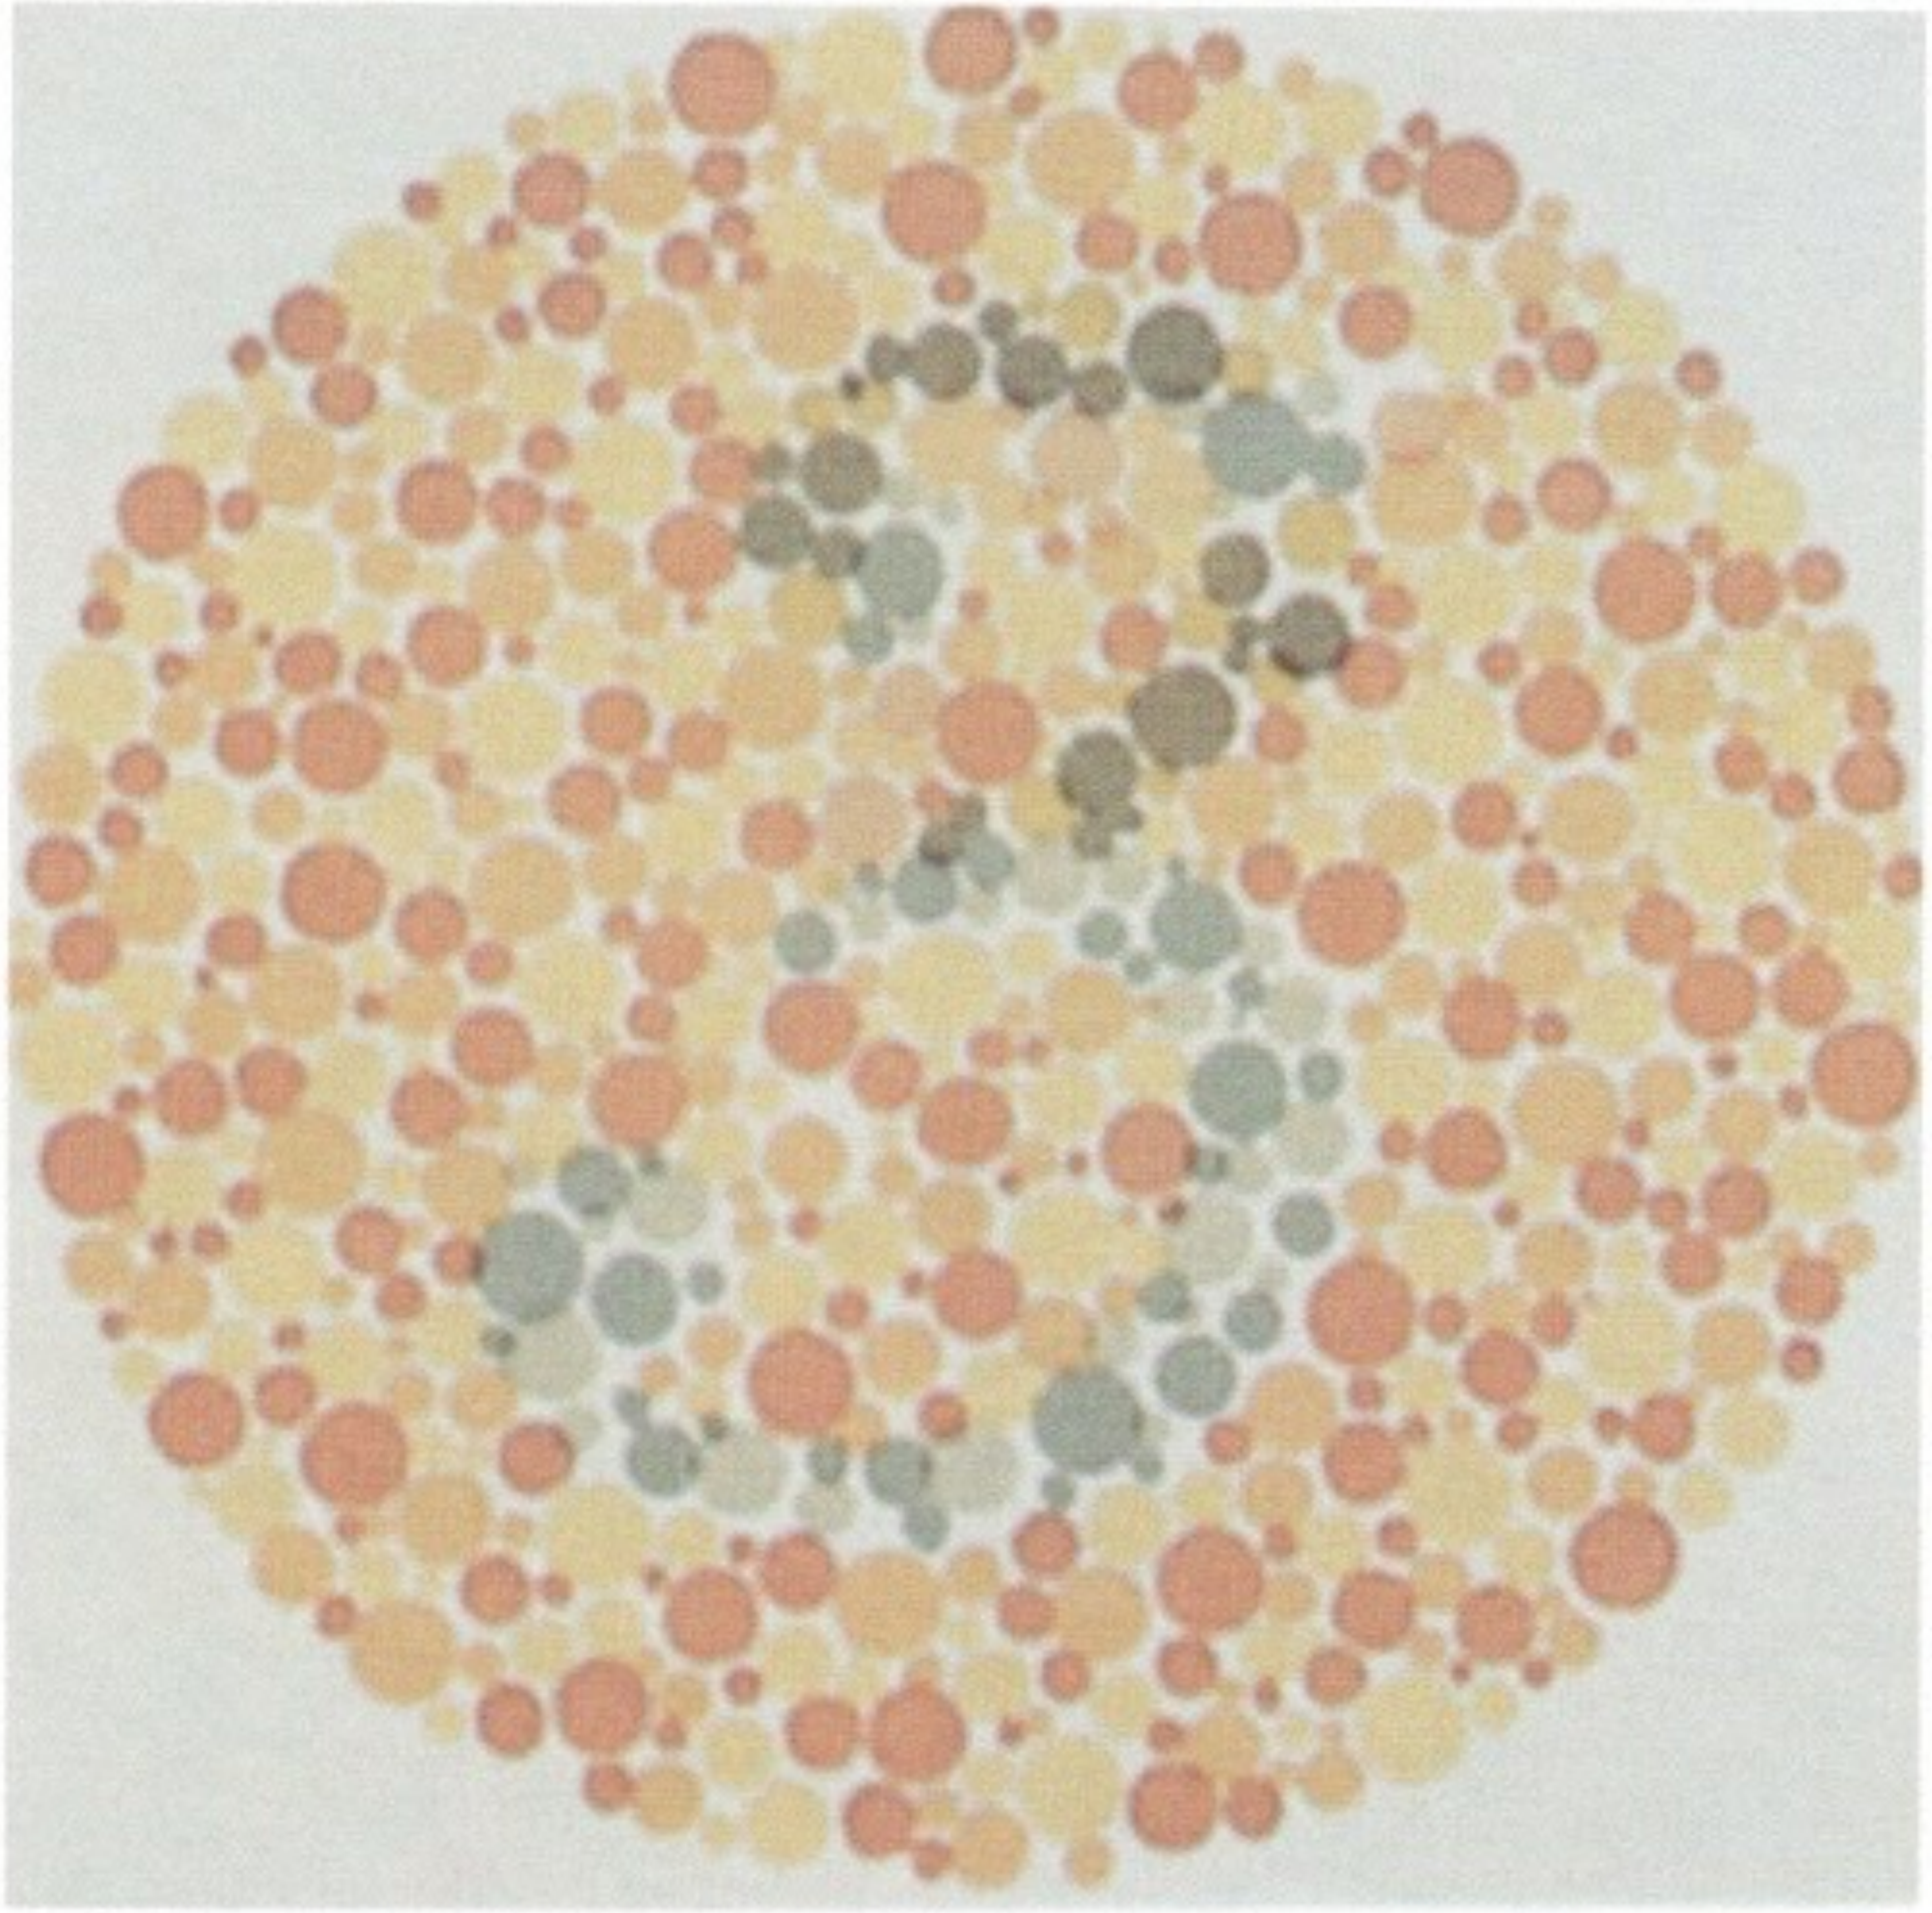

A

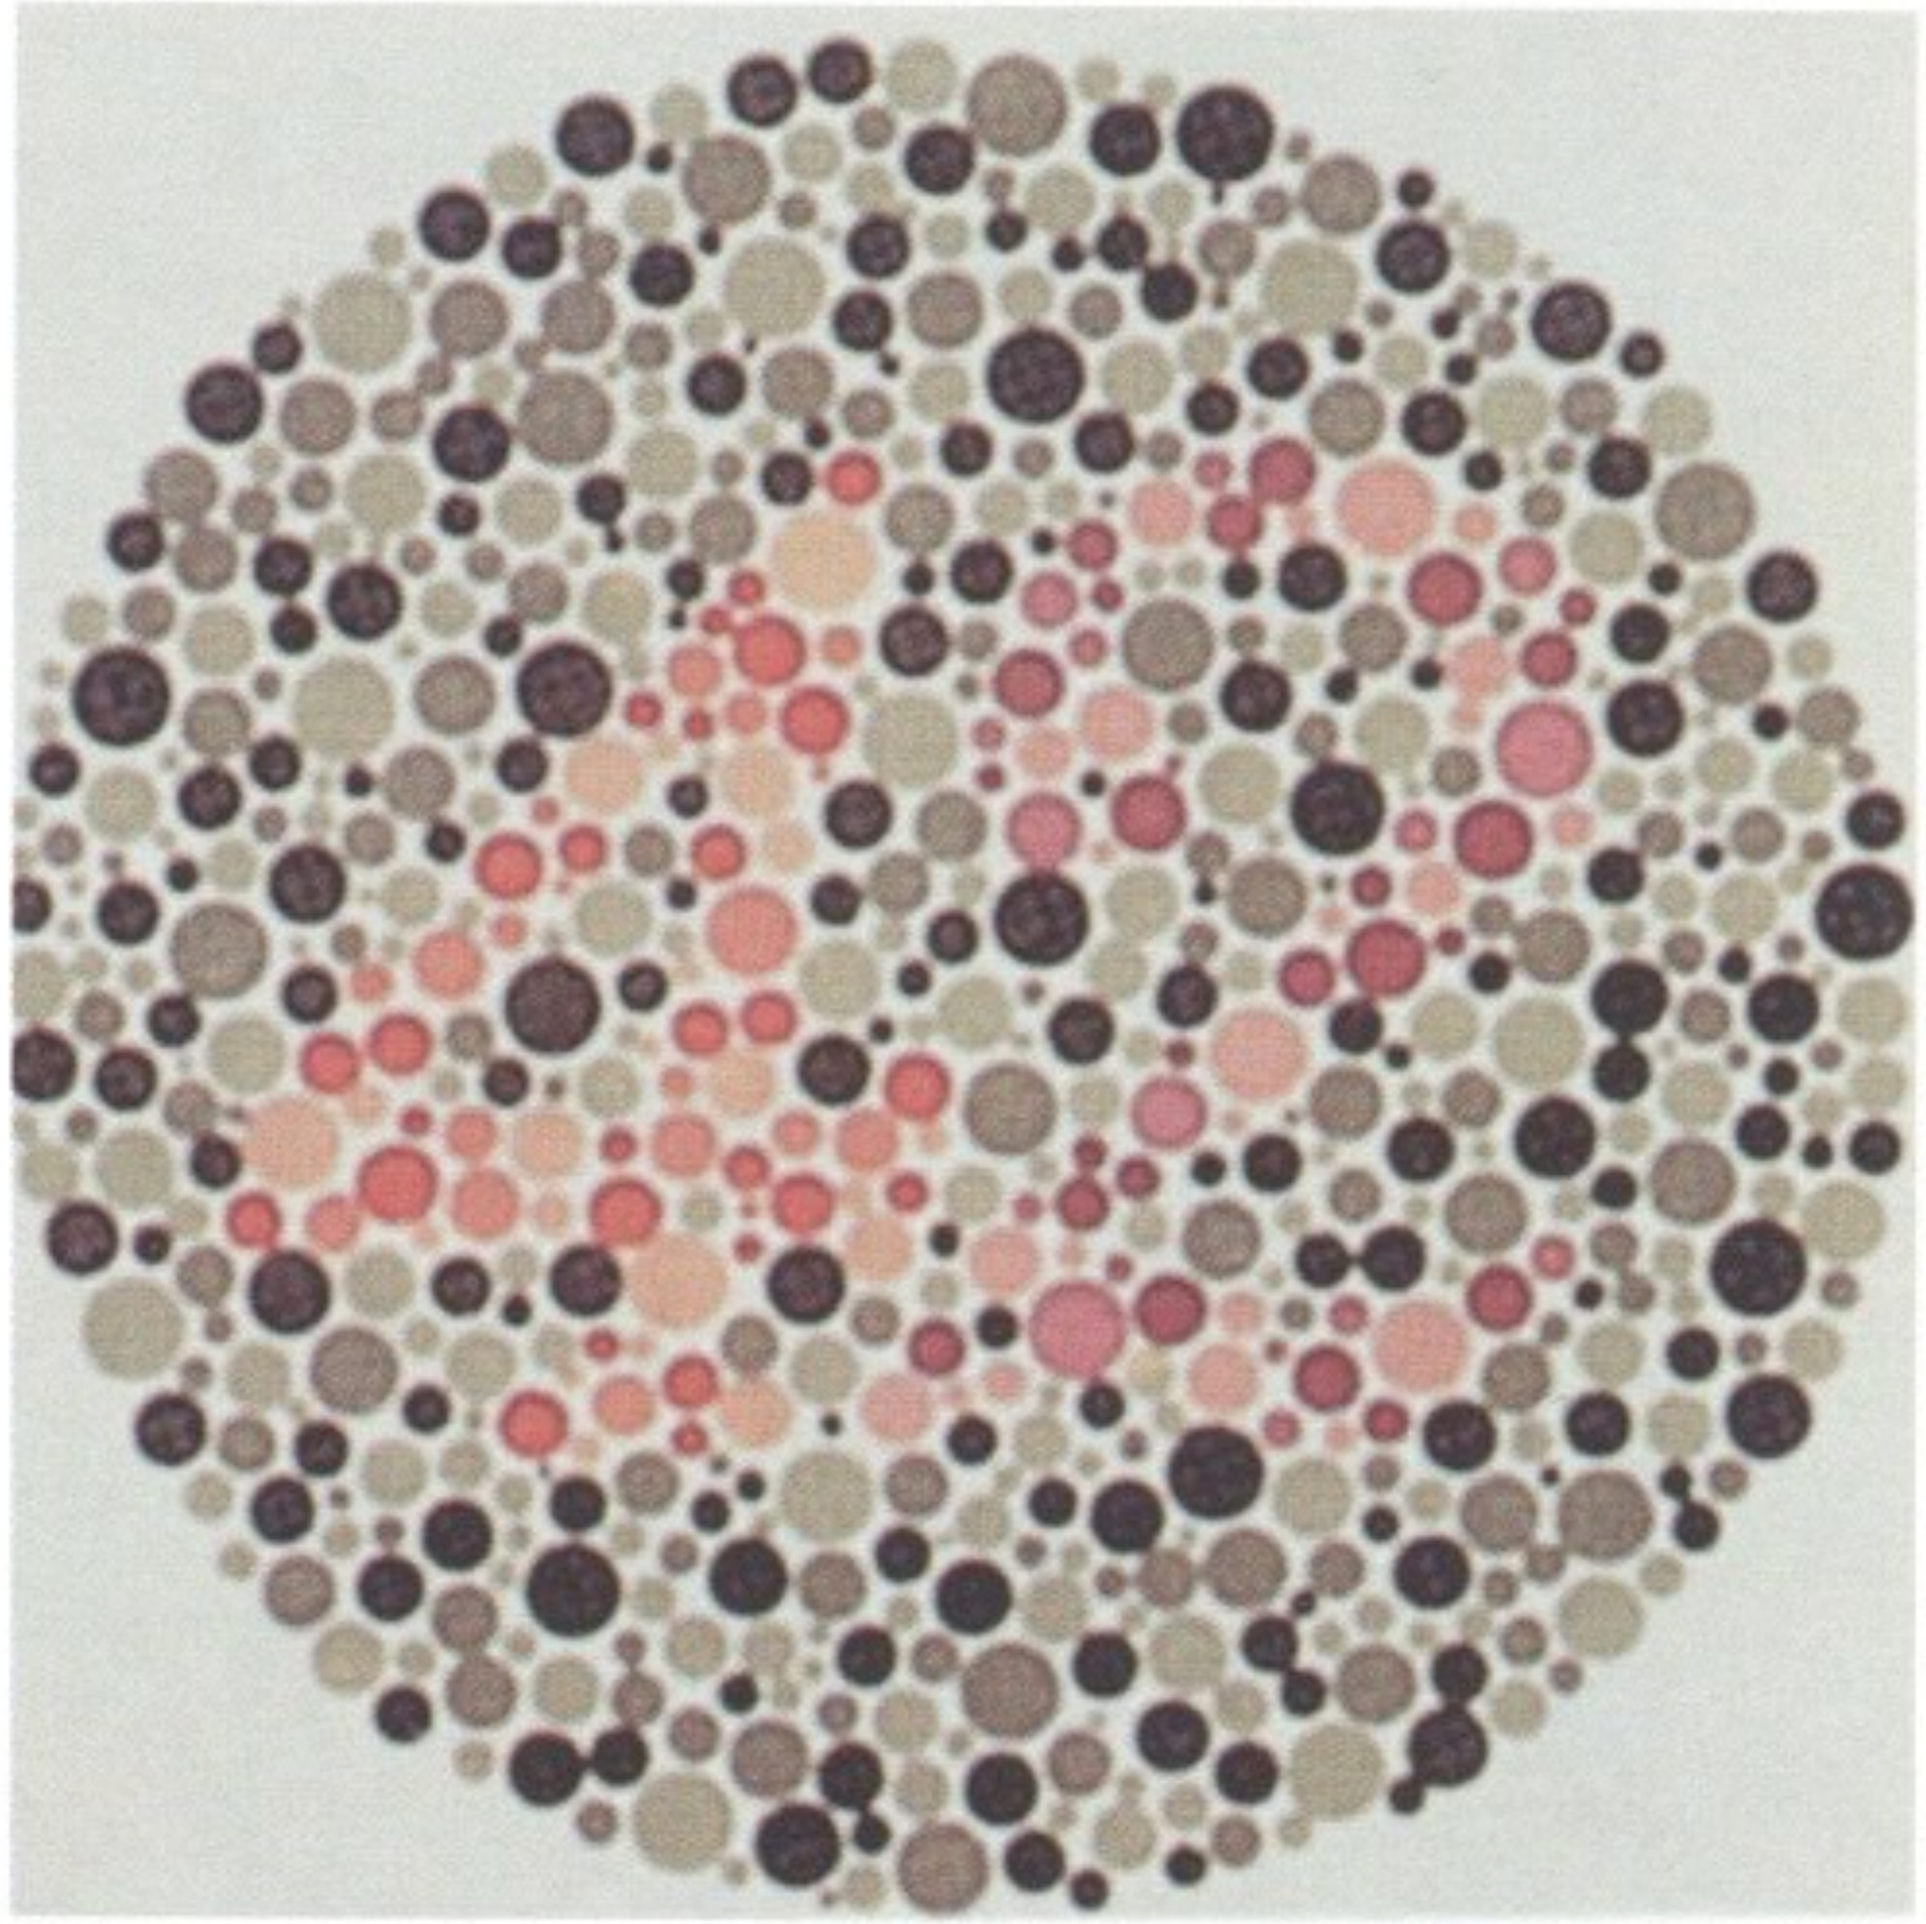

B

Color blindness test

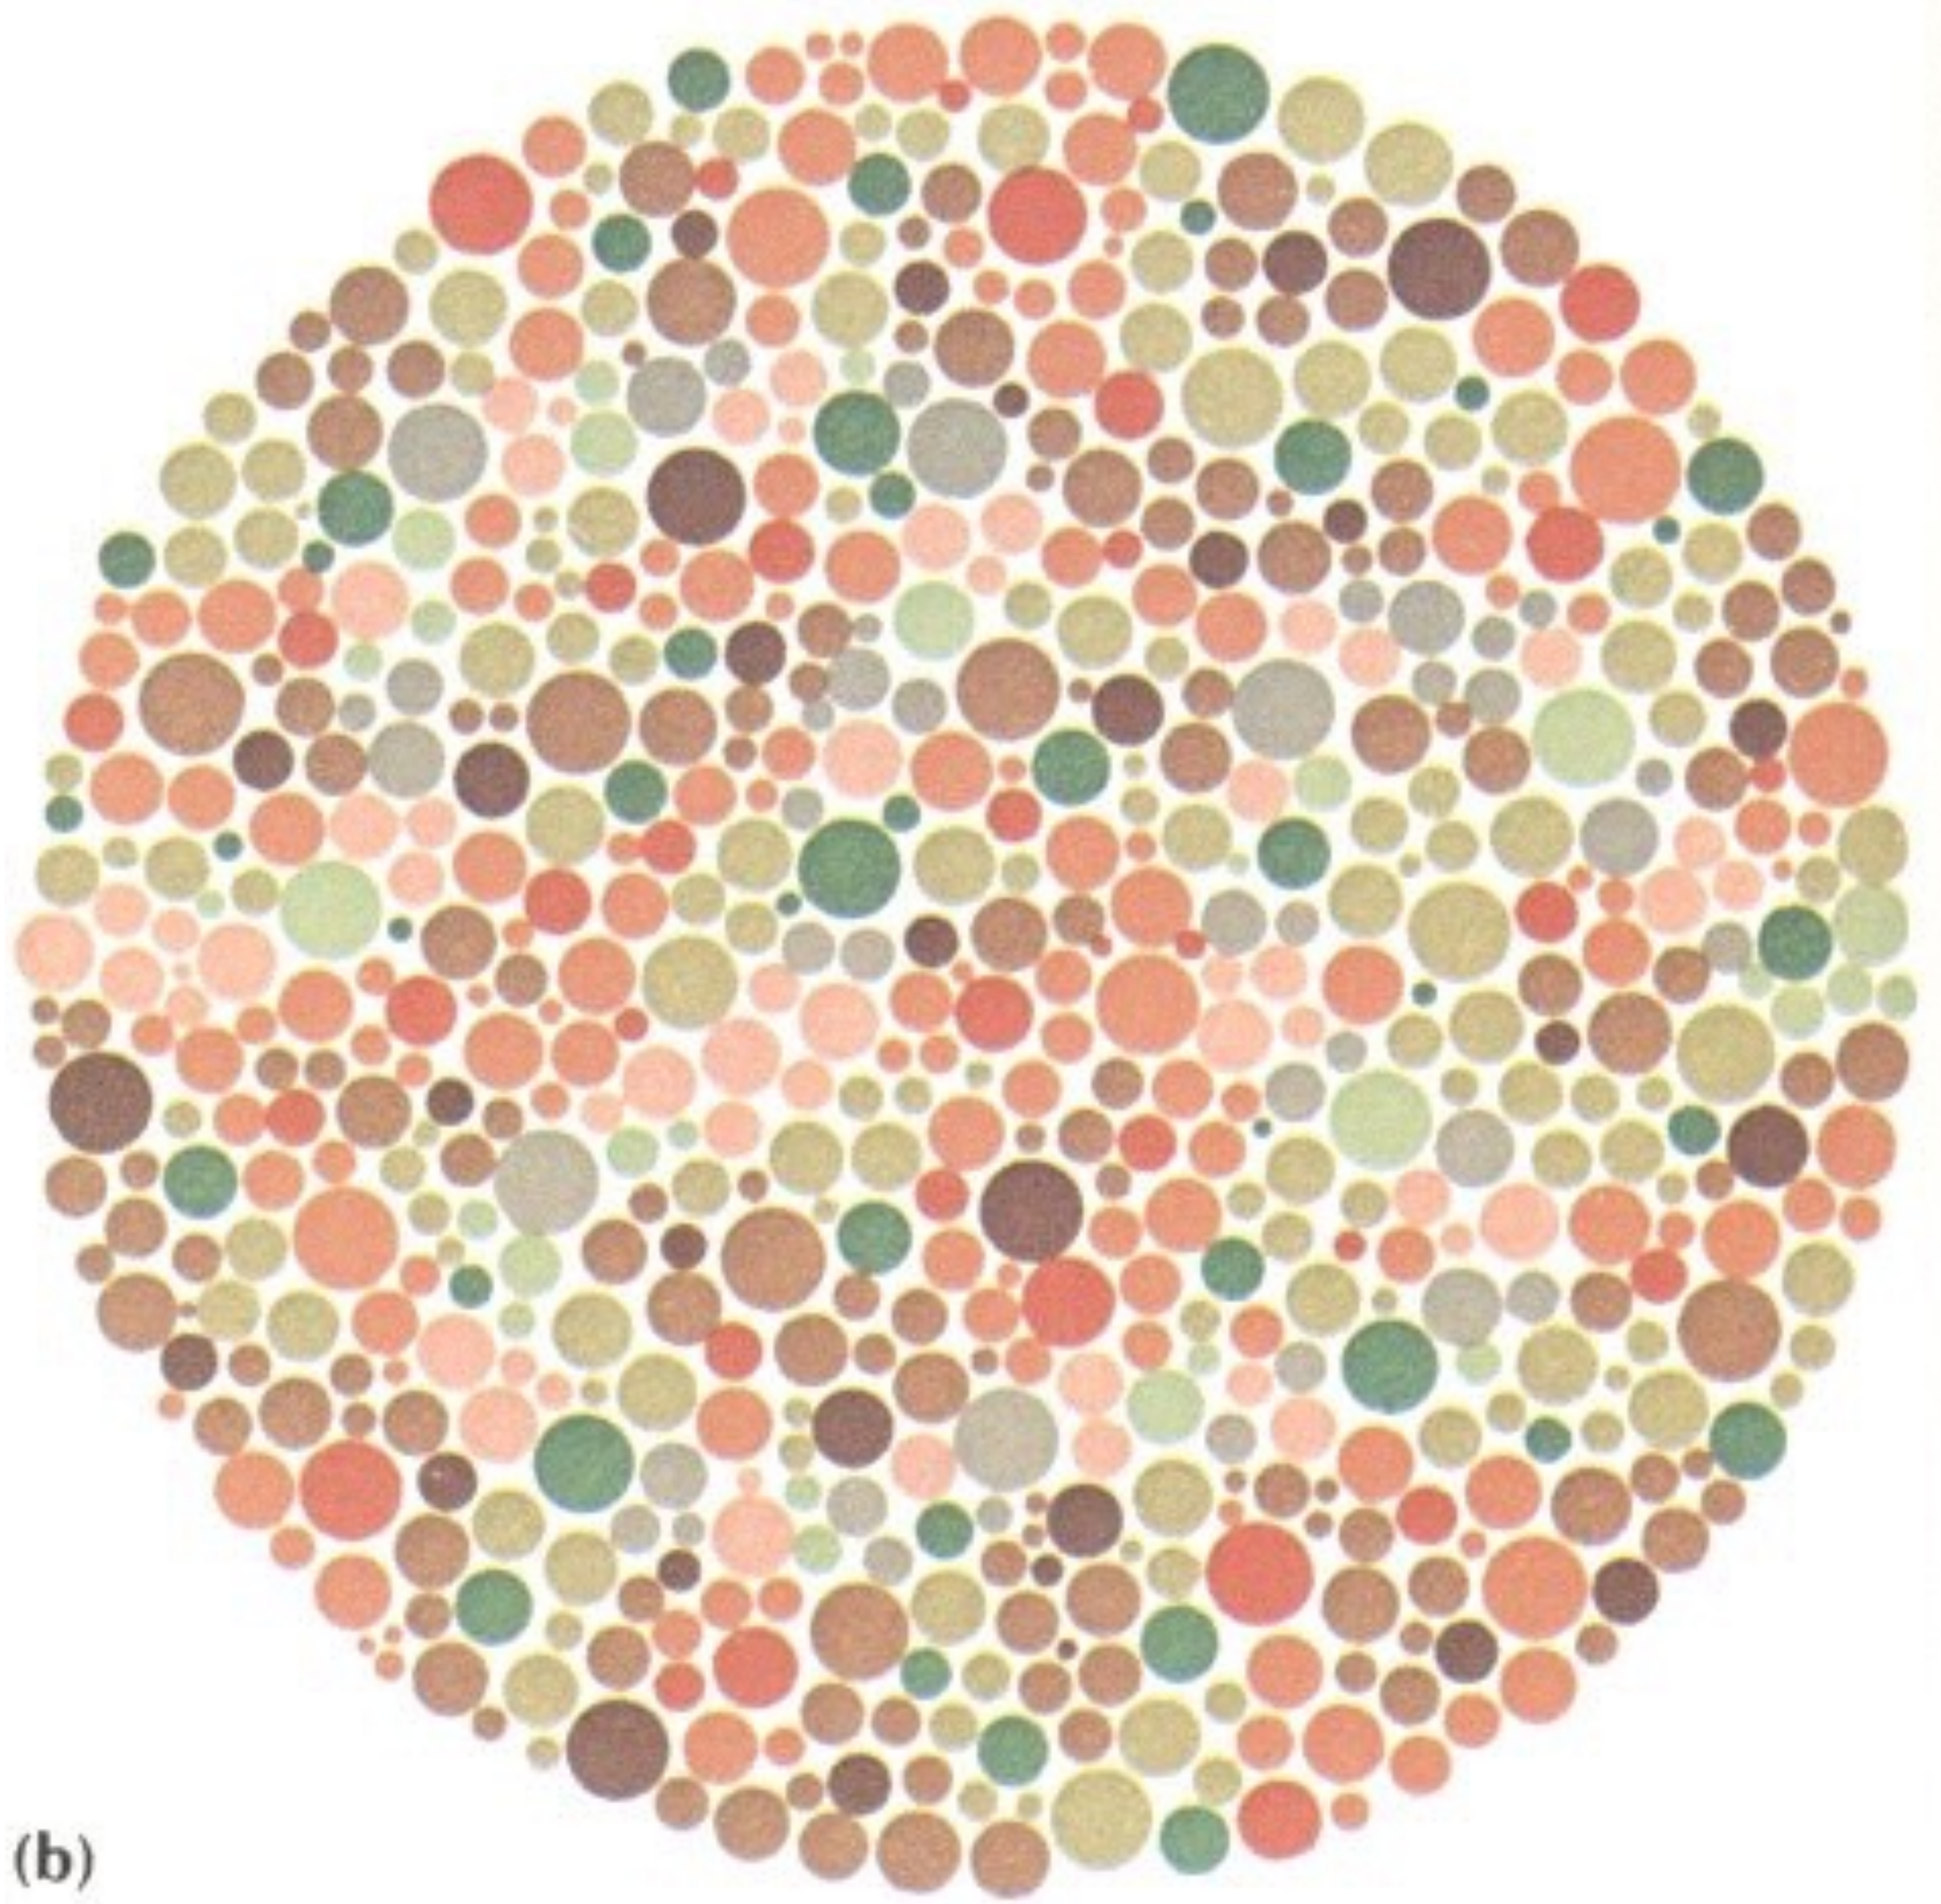

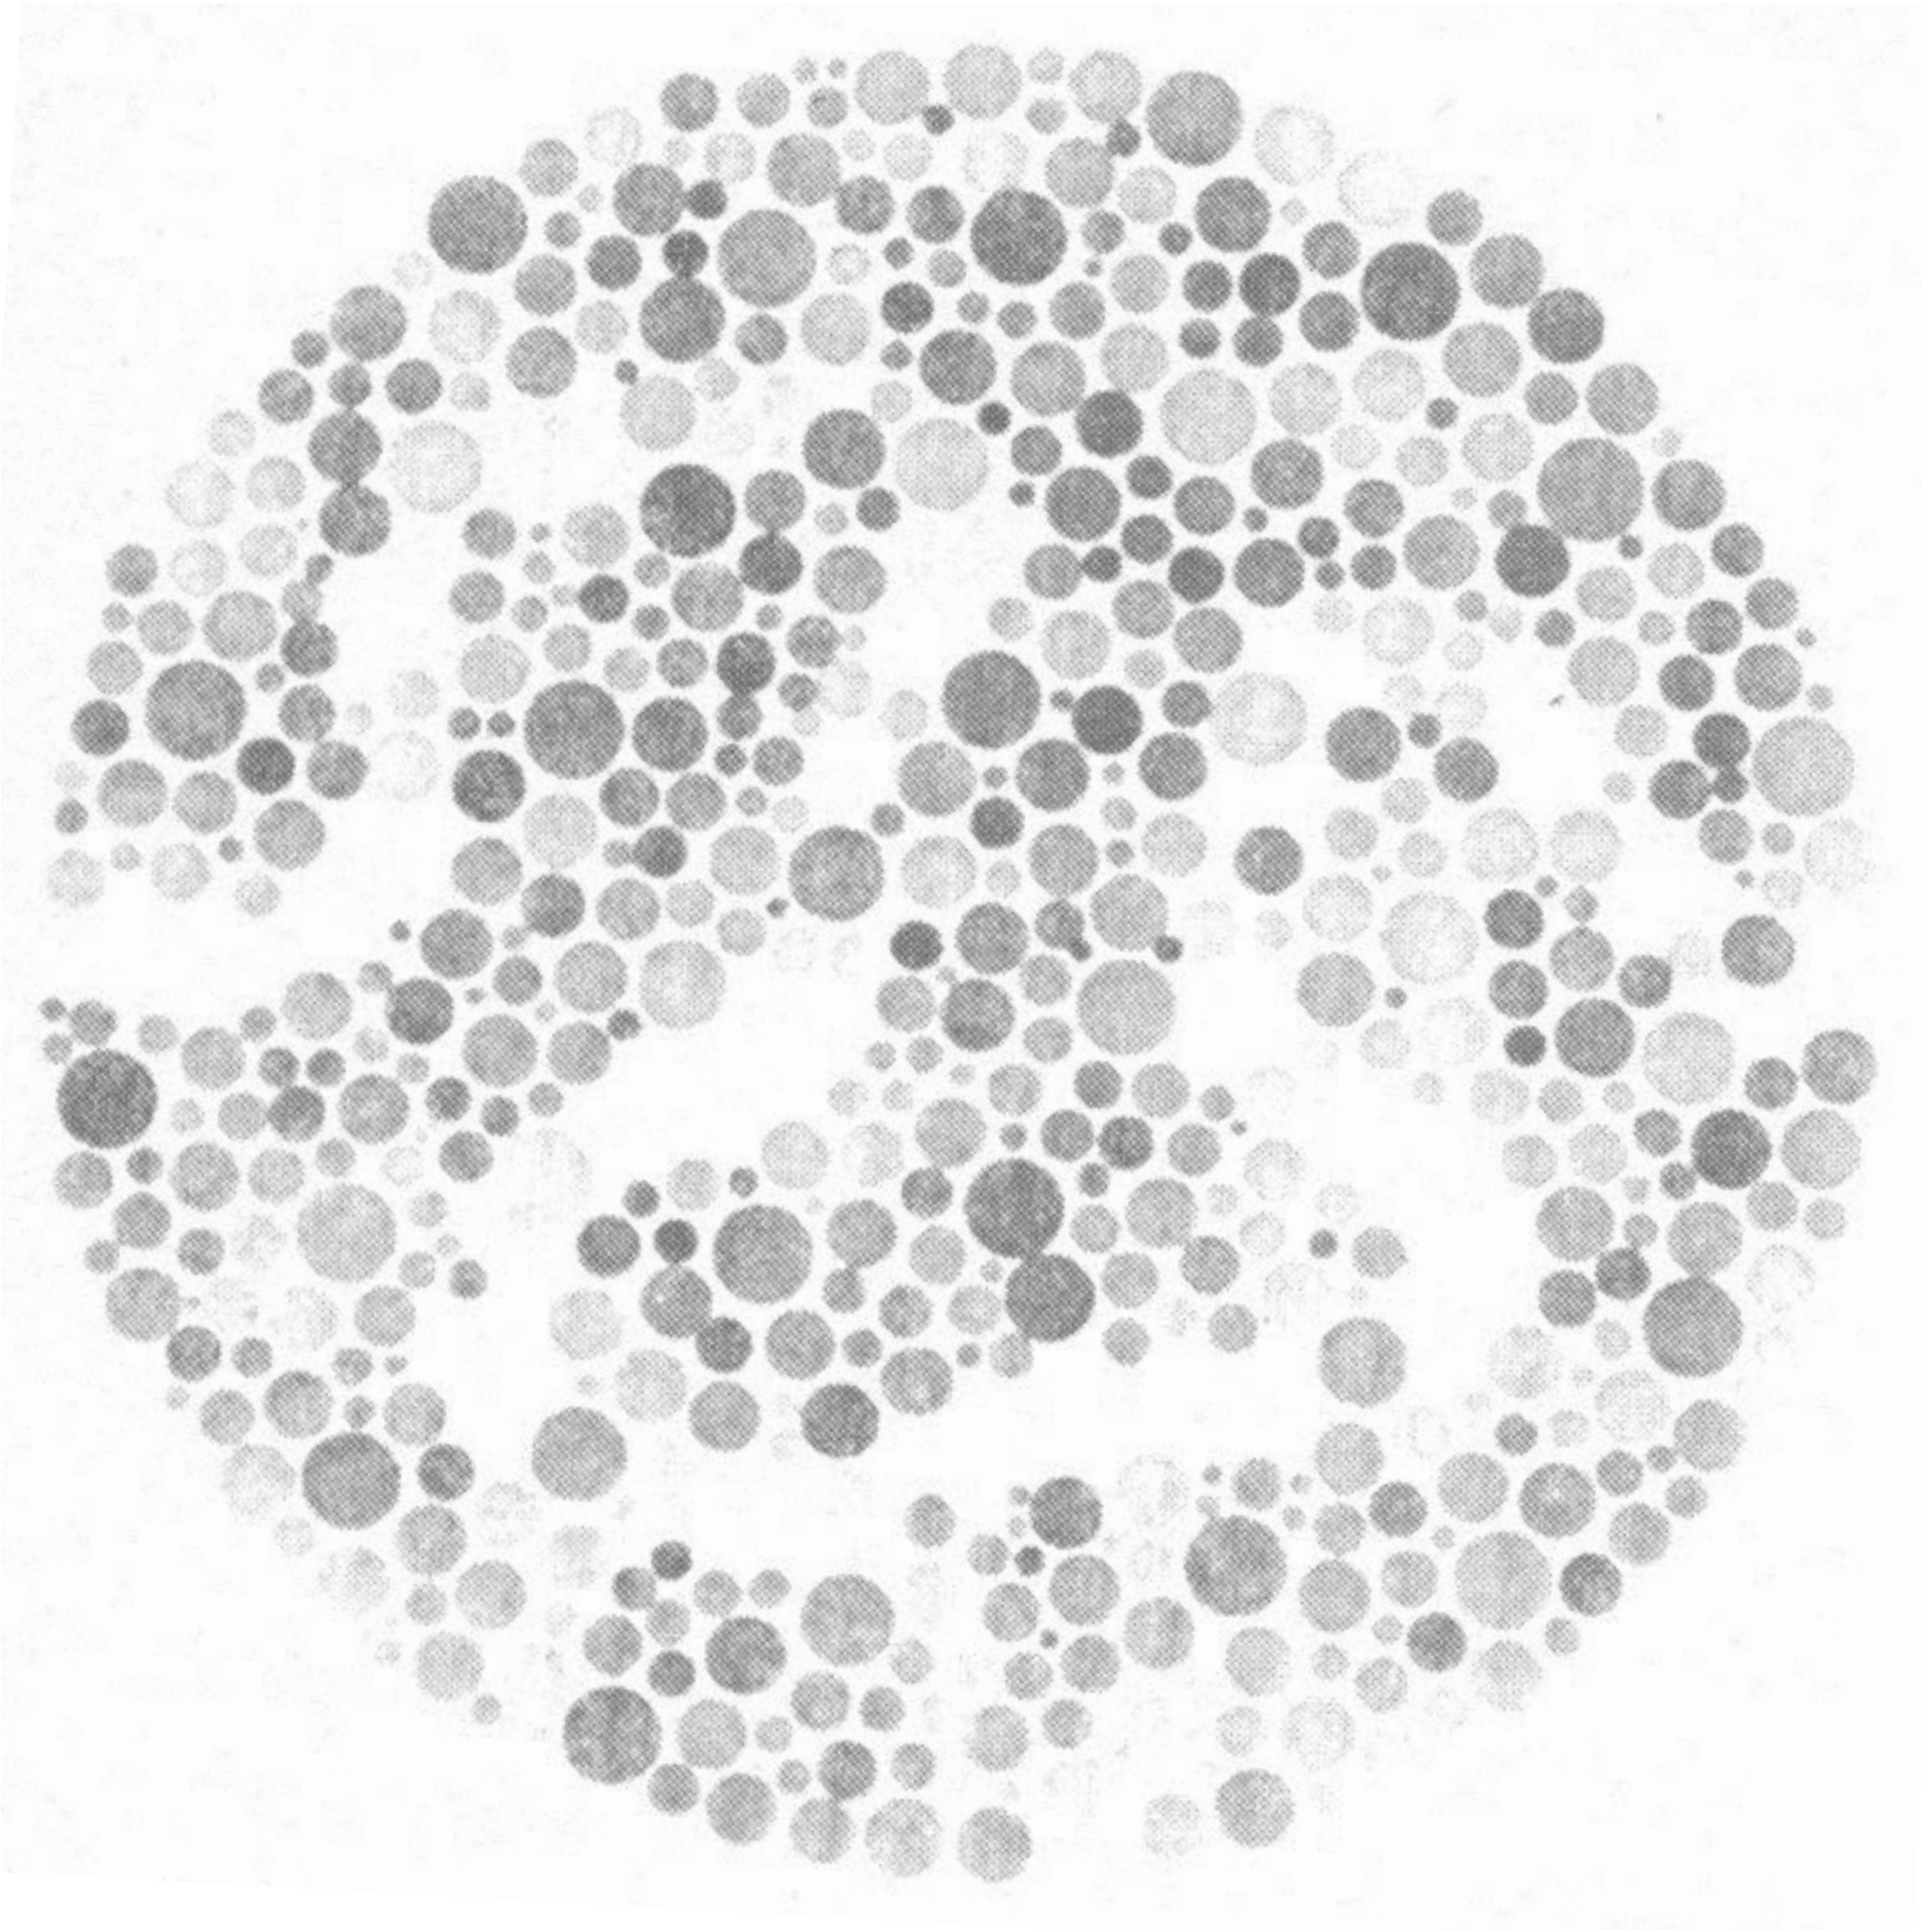

Color blindness test

Maze in subtle intensity contrast

Visible only to color blinds

Color contrast overrides intensity otherwise

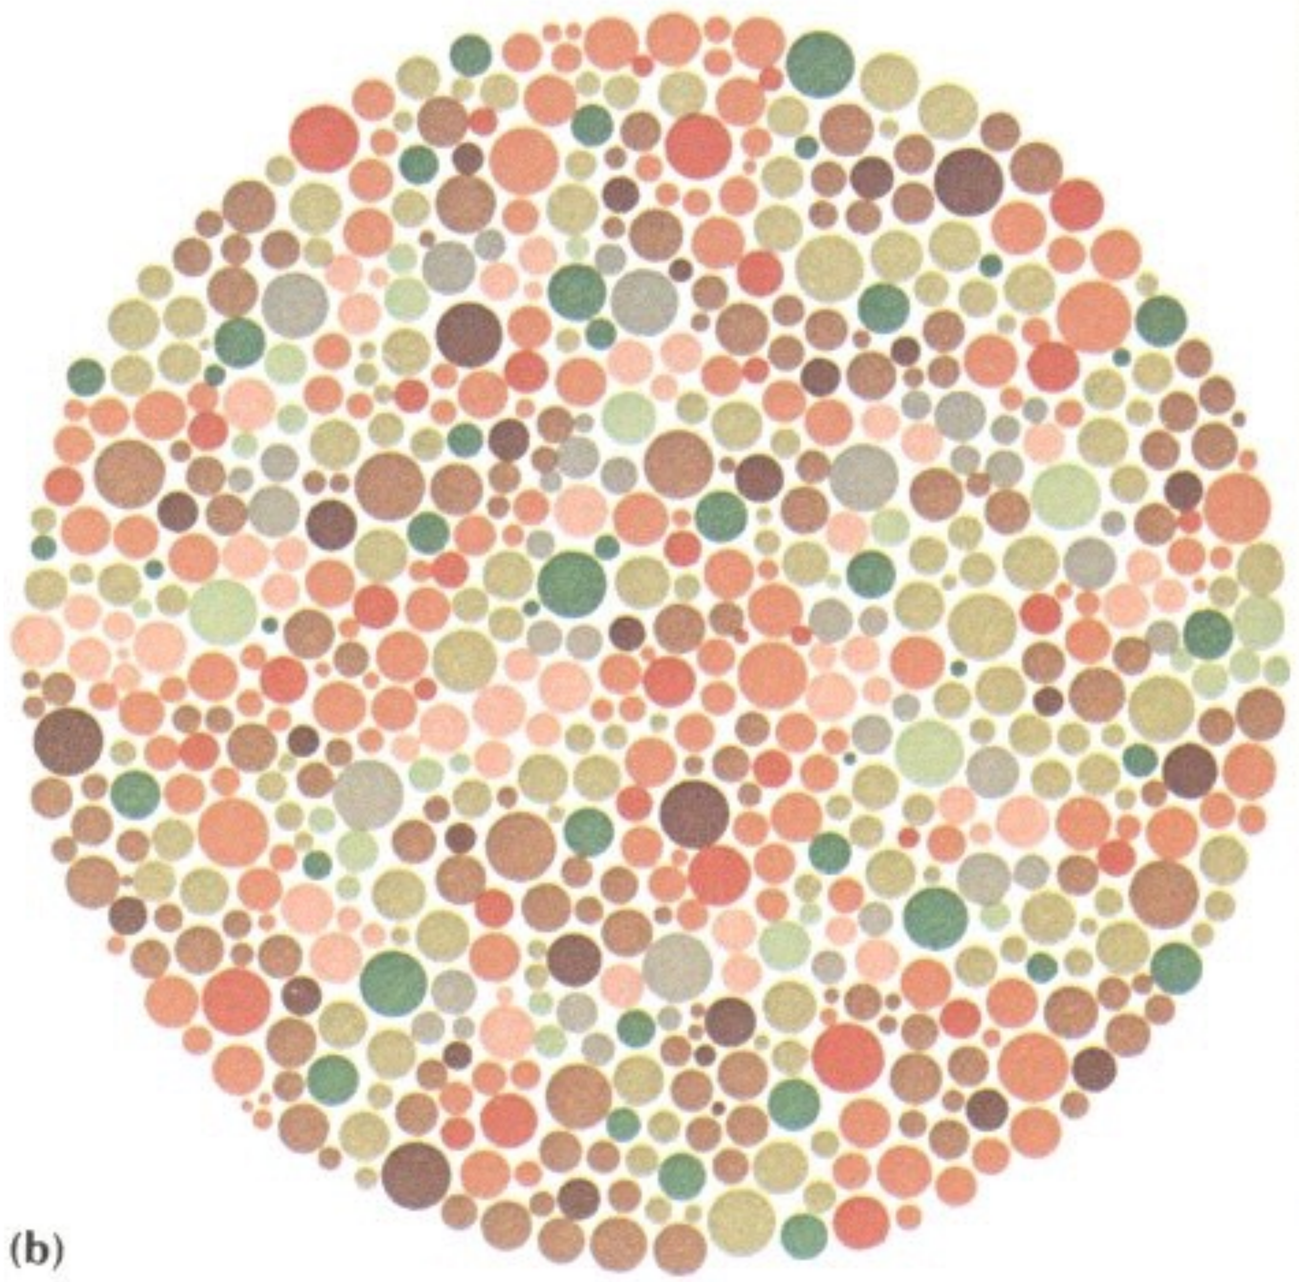

(b)

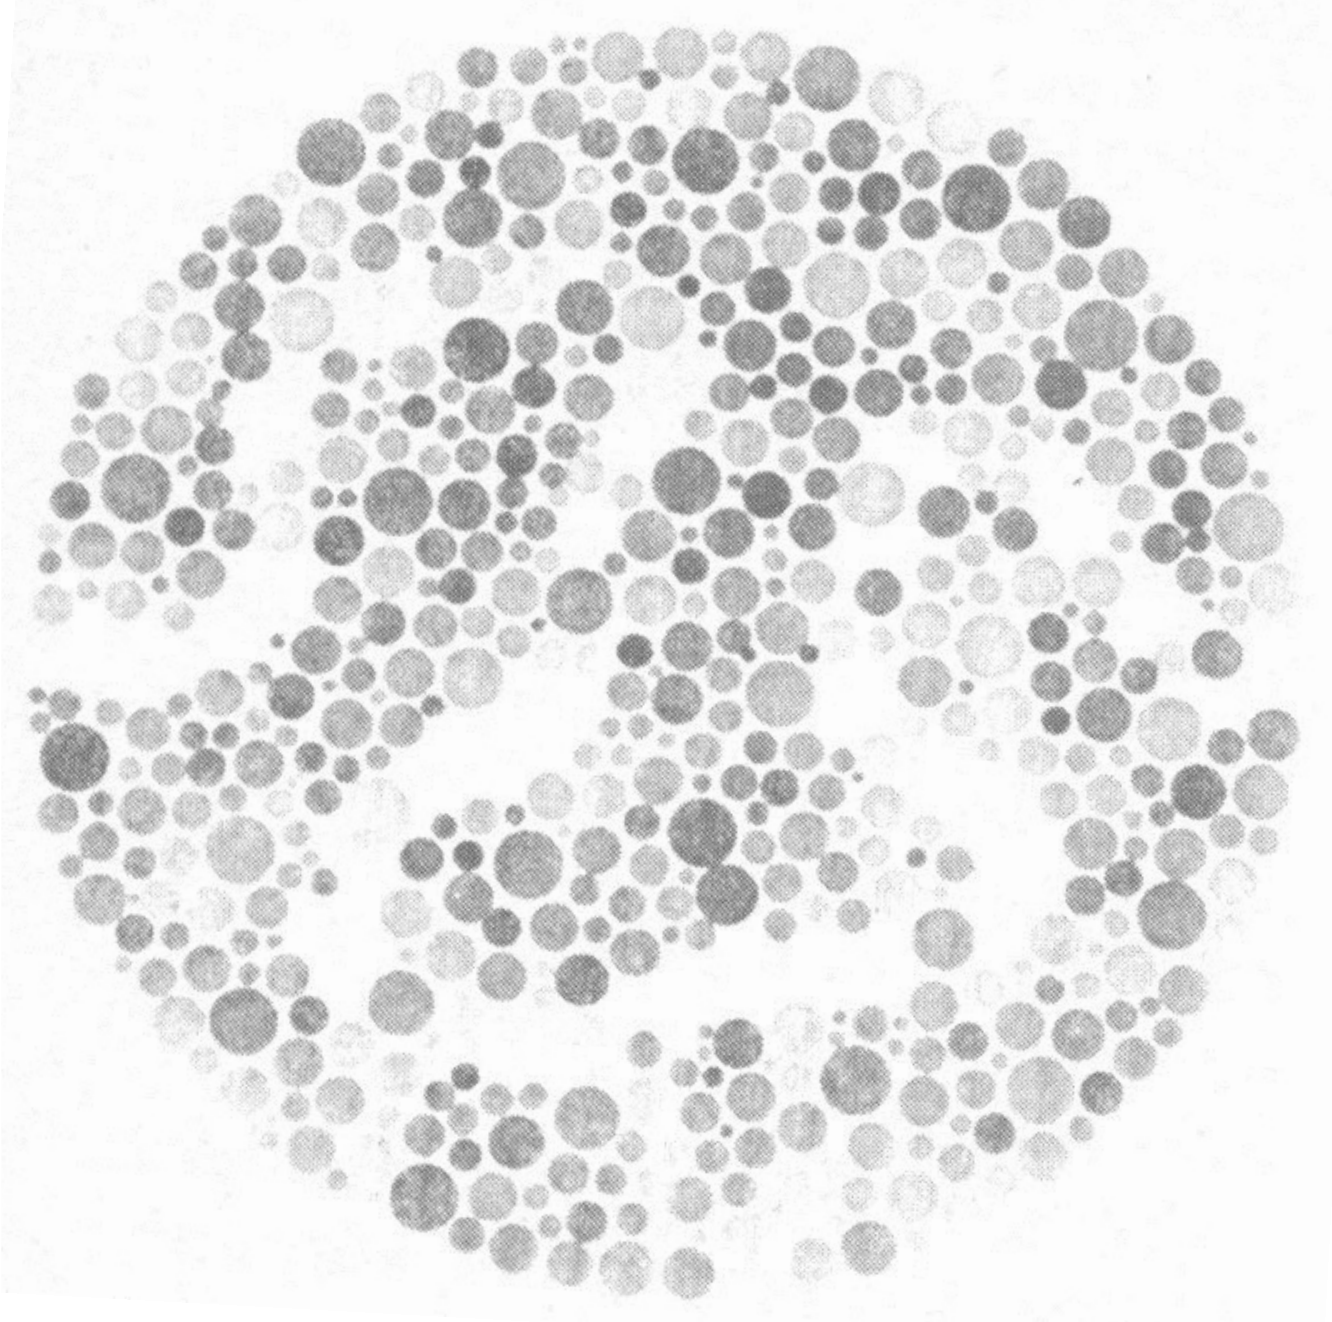

THE DIFFERENT APPEARANCES OF THE VISIBLE SPECTRUM

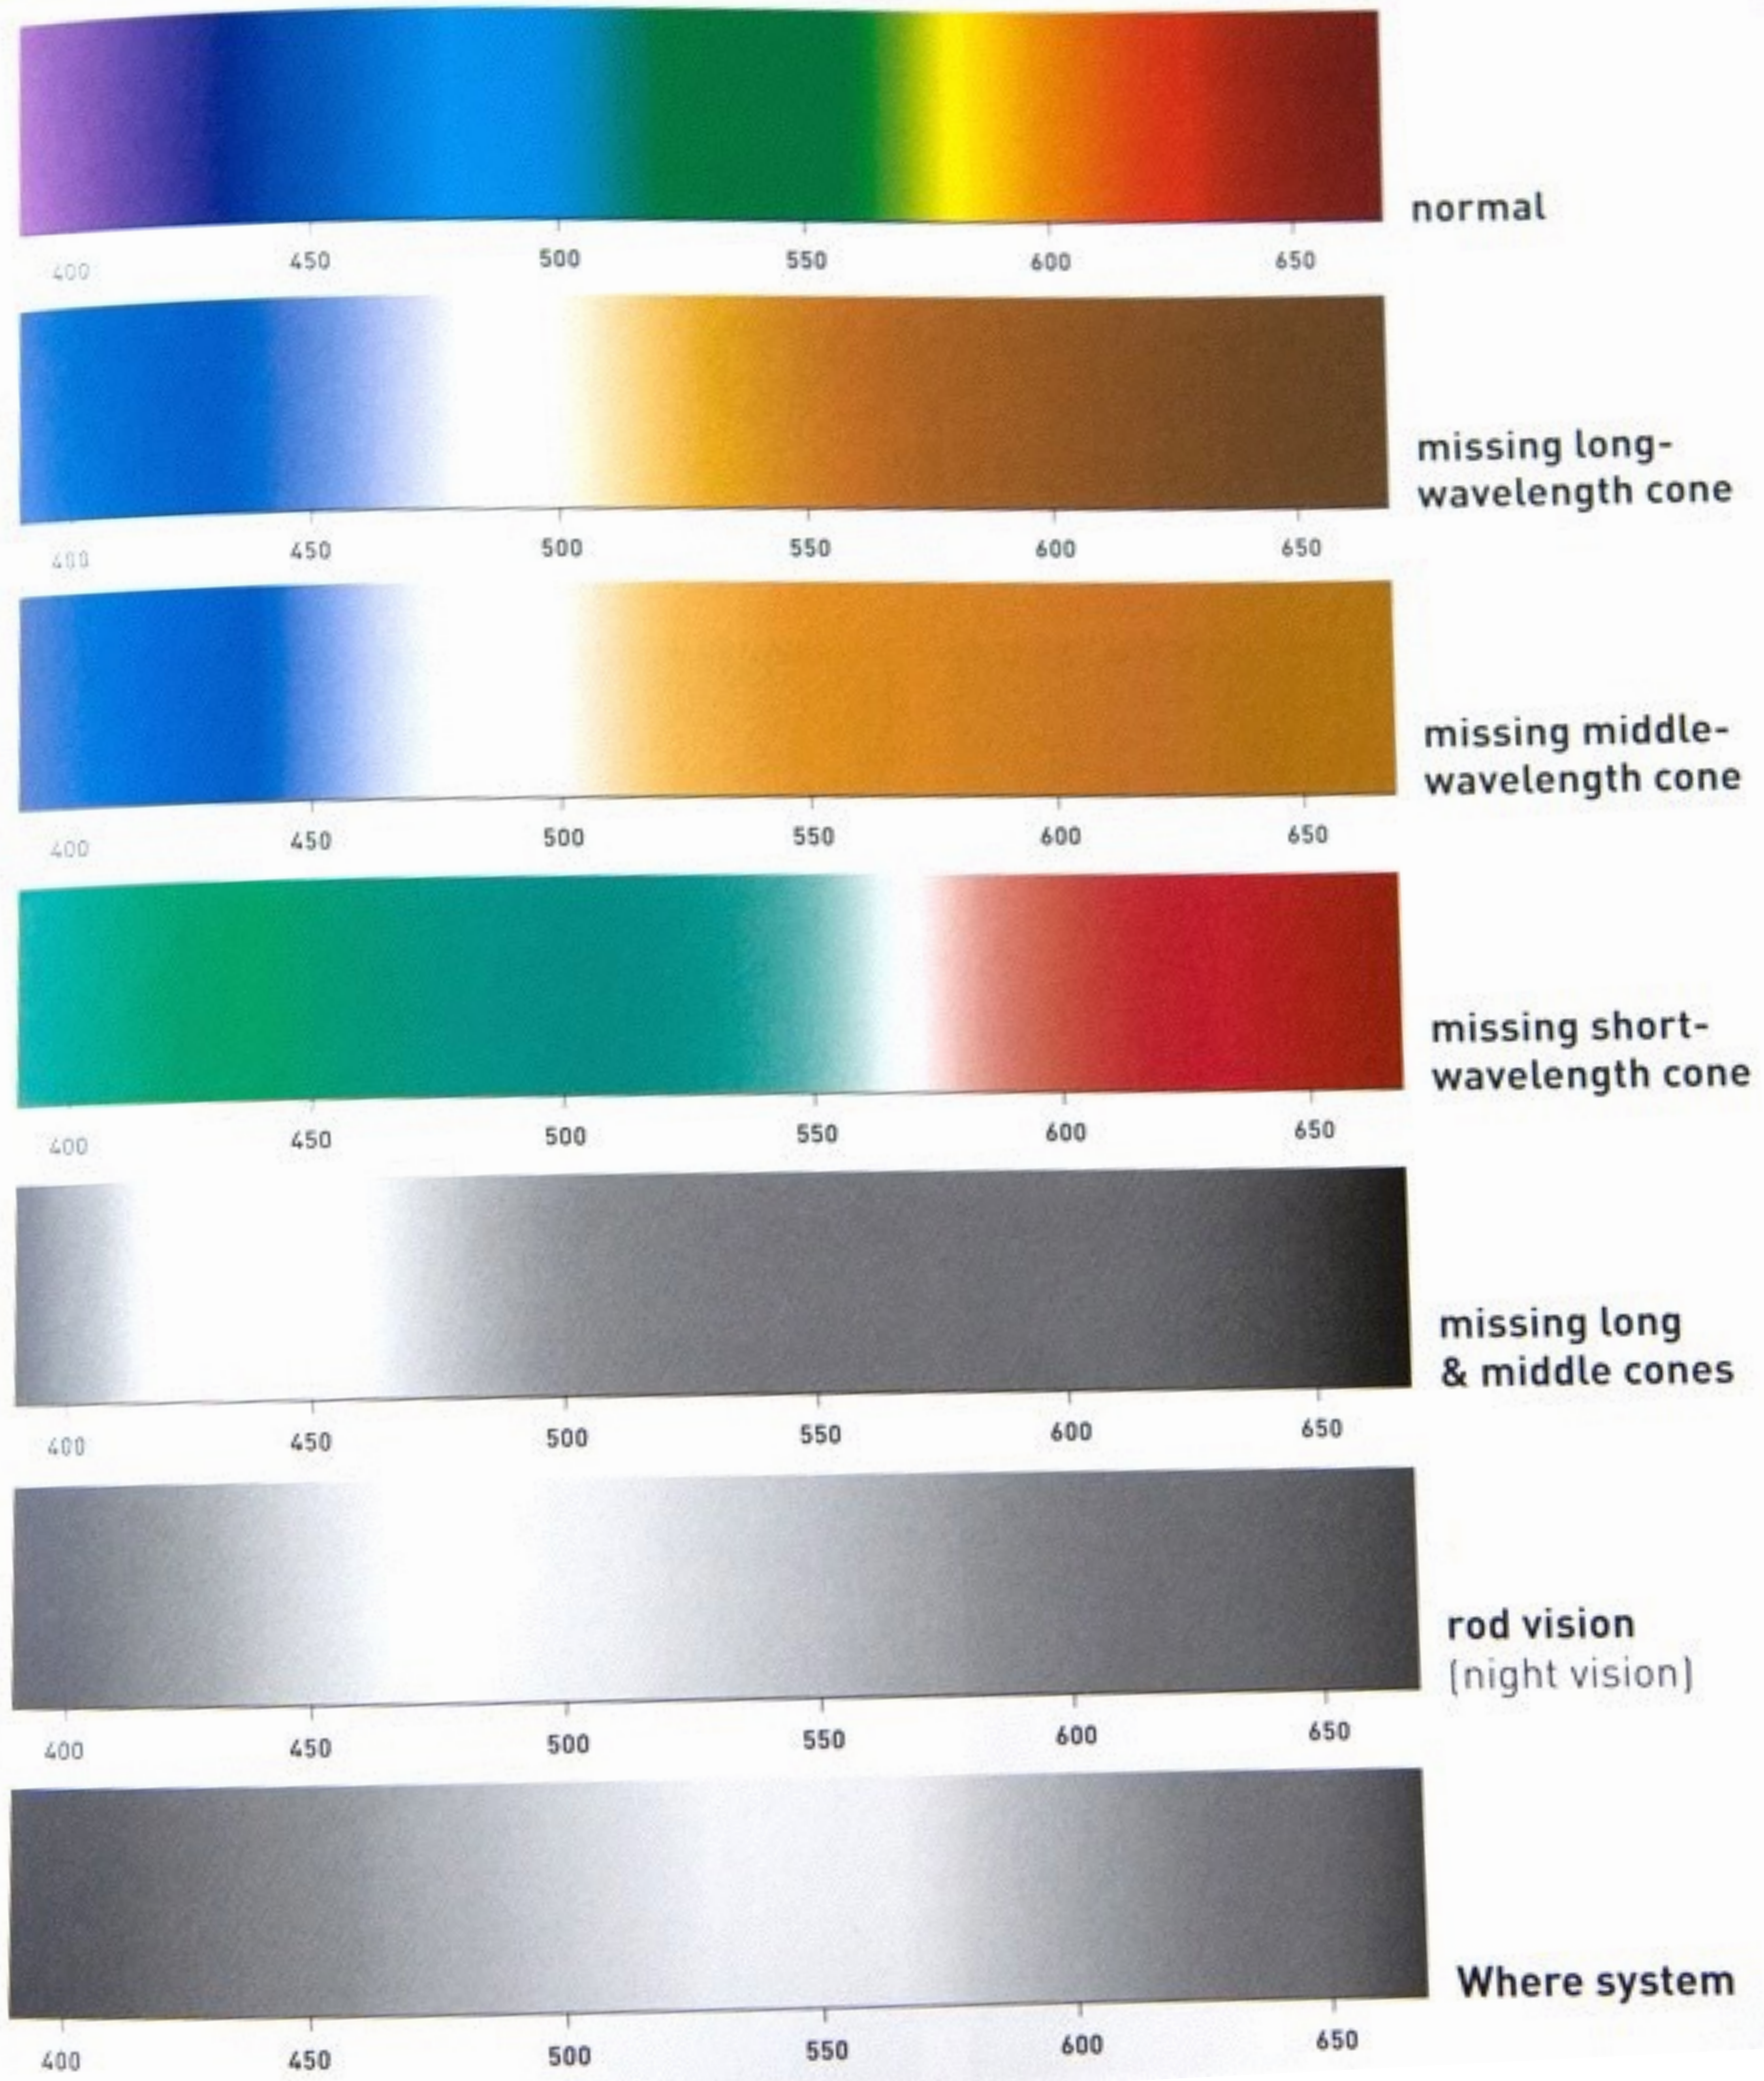

Questions?

Links:

- Vischeck shows you what an image looks like to someone who is colorblind

Metamers

Metamers

We are all color blind!

These two different spectra elicit the same cone responses

Called *metamers*

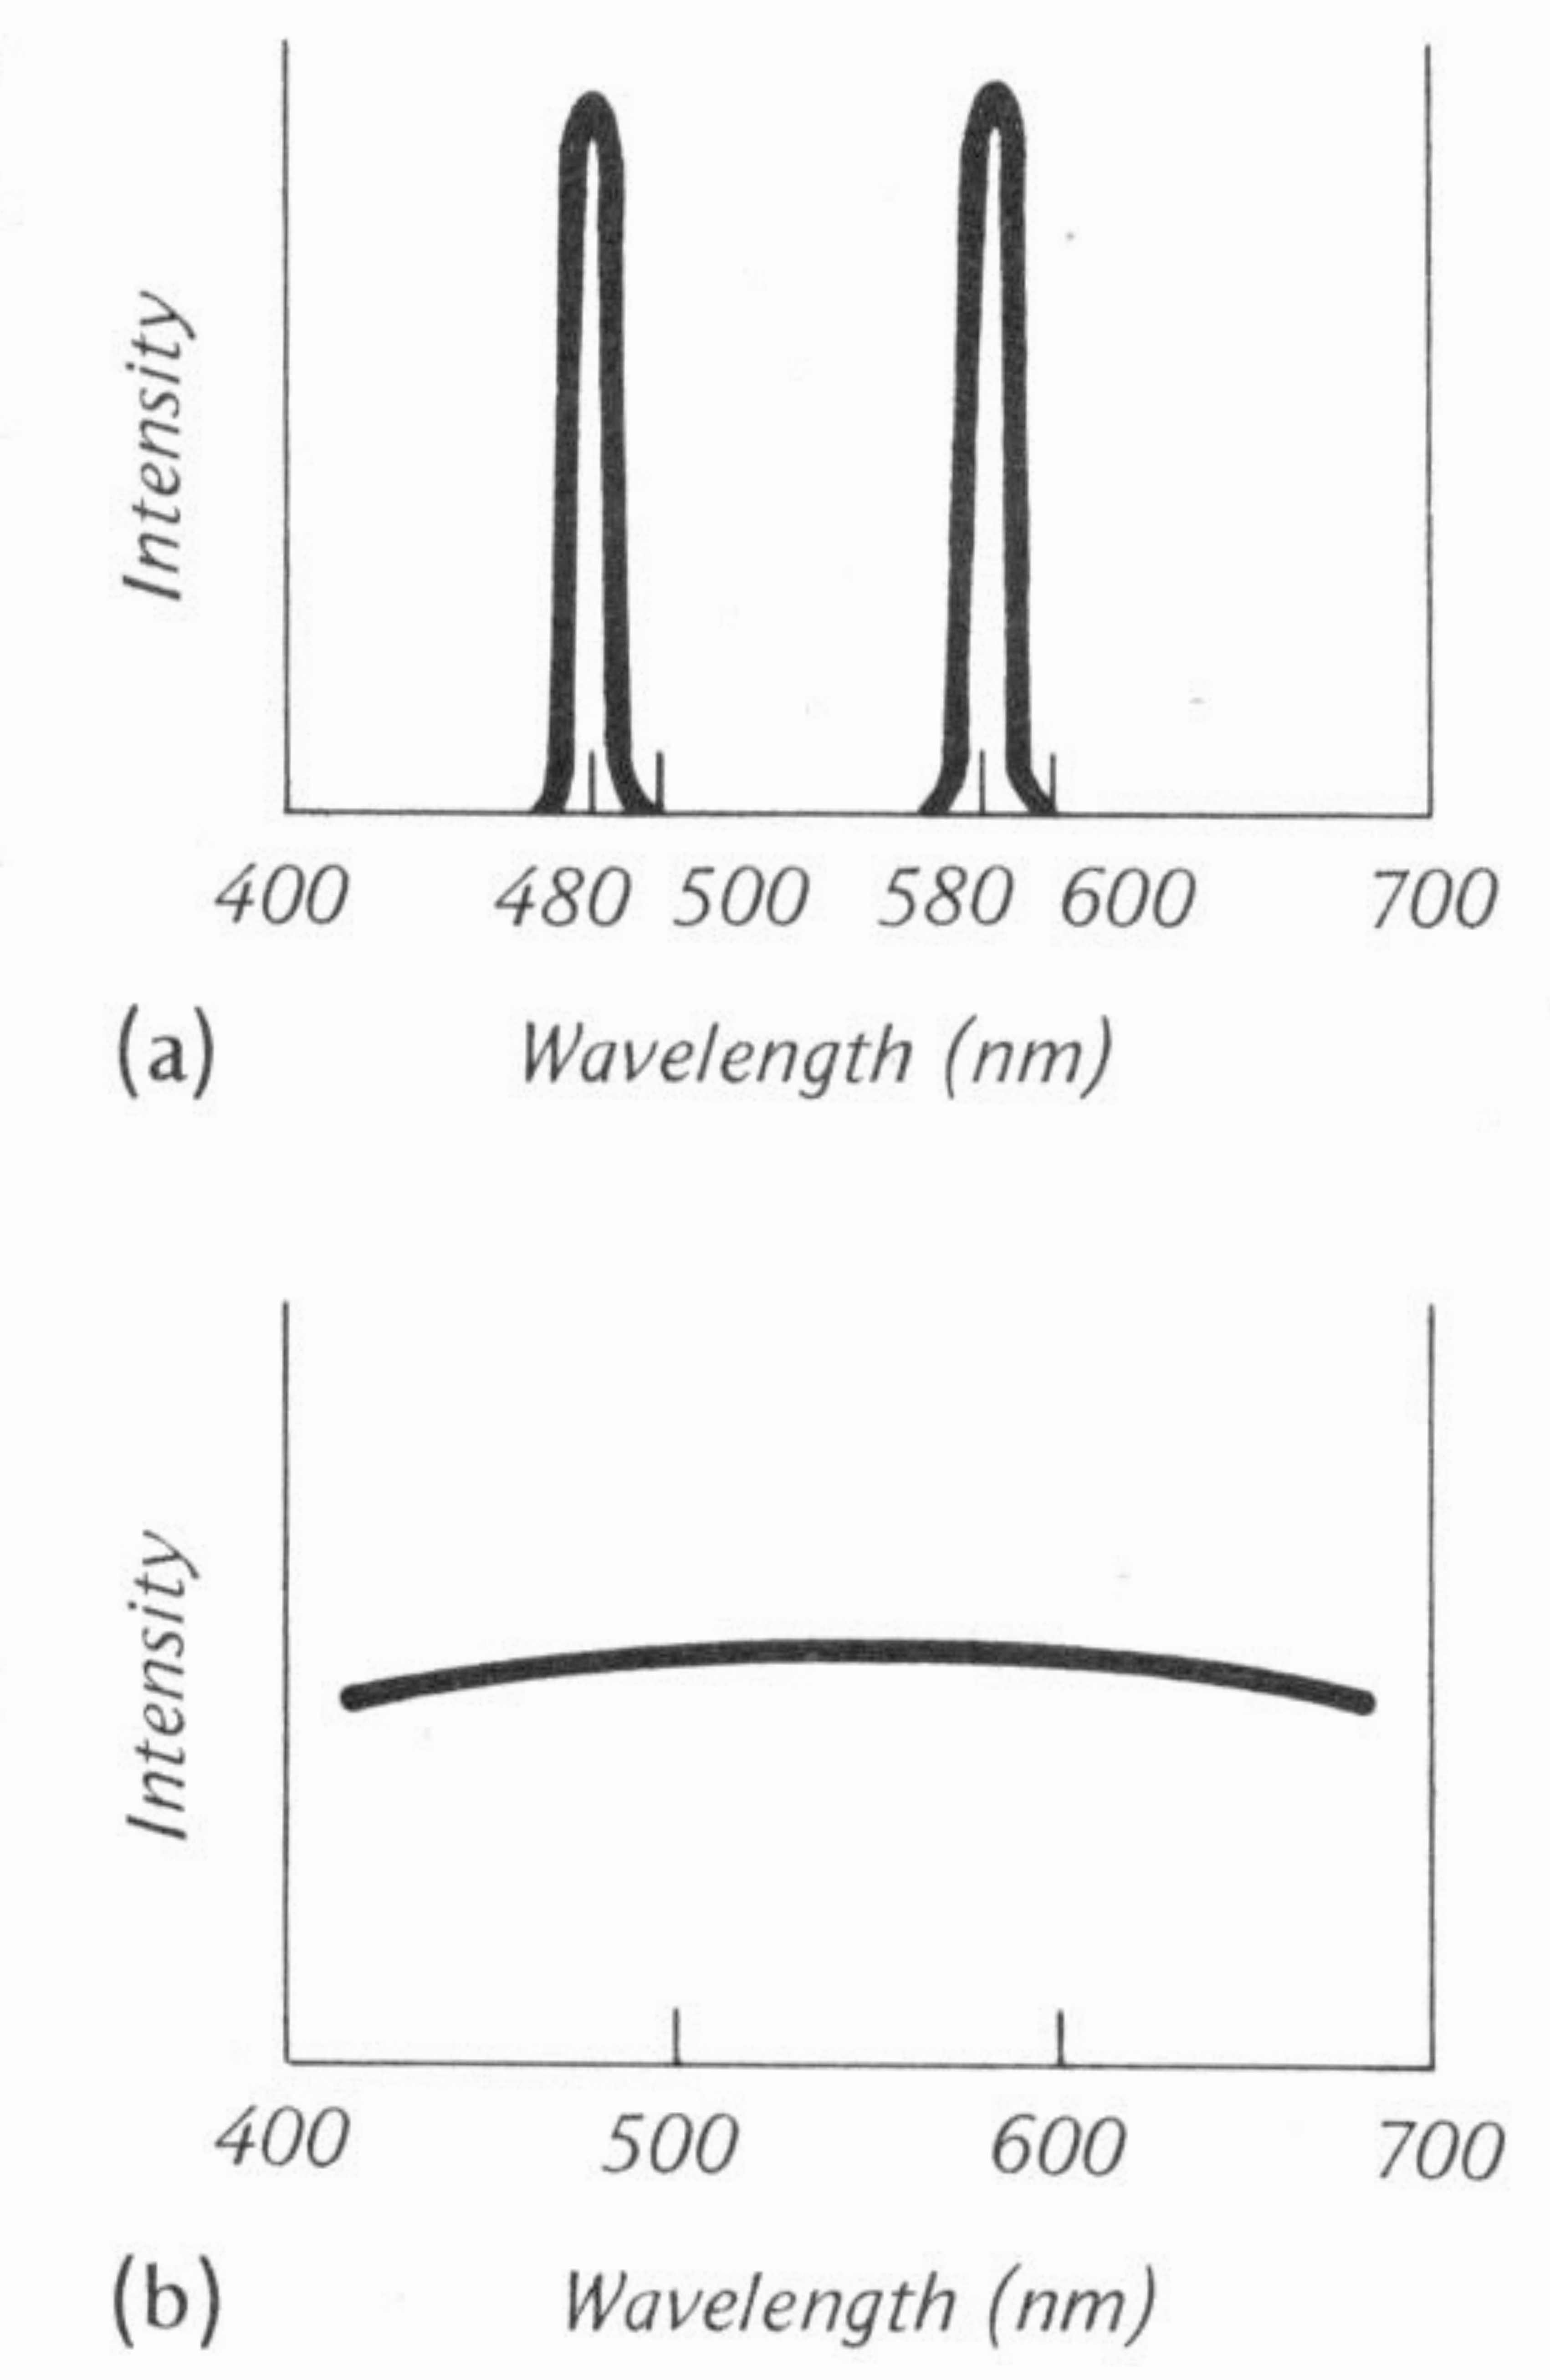

Basic fact of colorimetry

Take a spectrum (which is a function)

Eye produces three numbers

This throws away a lot of information!

- Quite possible to have two different spectra that have the same S, M, L tristimulus values
- Two such spectra are *metamers*

Pseudo-geometric interpretation

A dot product is a projection

Humans project an infinite dimensional vector (the SPD) onto a 3-D subspace

- differences that are perpendicular to all 3 vectors are not detectable

For intuition, we can imagine a 3D analog

- 3D stands in for the infinite-dimensional vectors
- 2D stands in for 3D
- Then color perception is just projection onto a plane

Pseudo-geometric interpretation

The information available to the visual system about a spectrum is just 3 numbers!

Two spectra that project to the same response are *metamers*

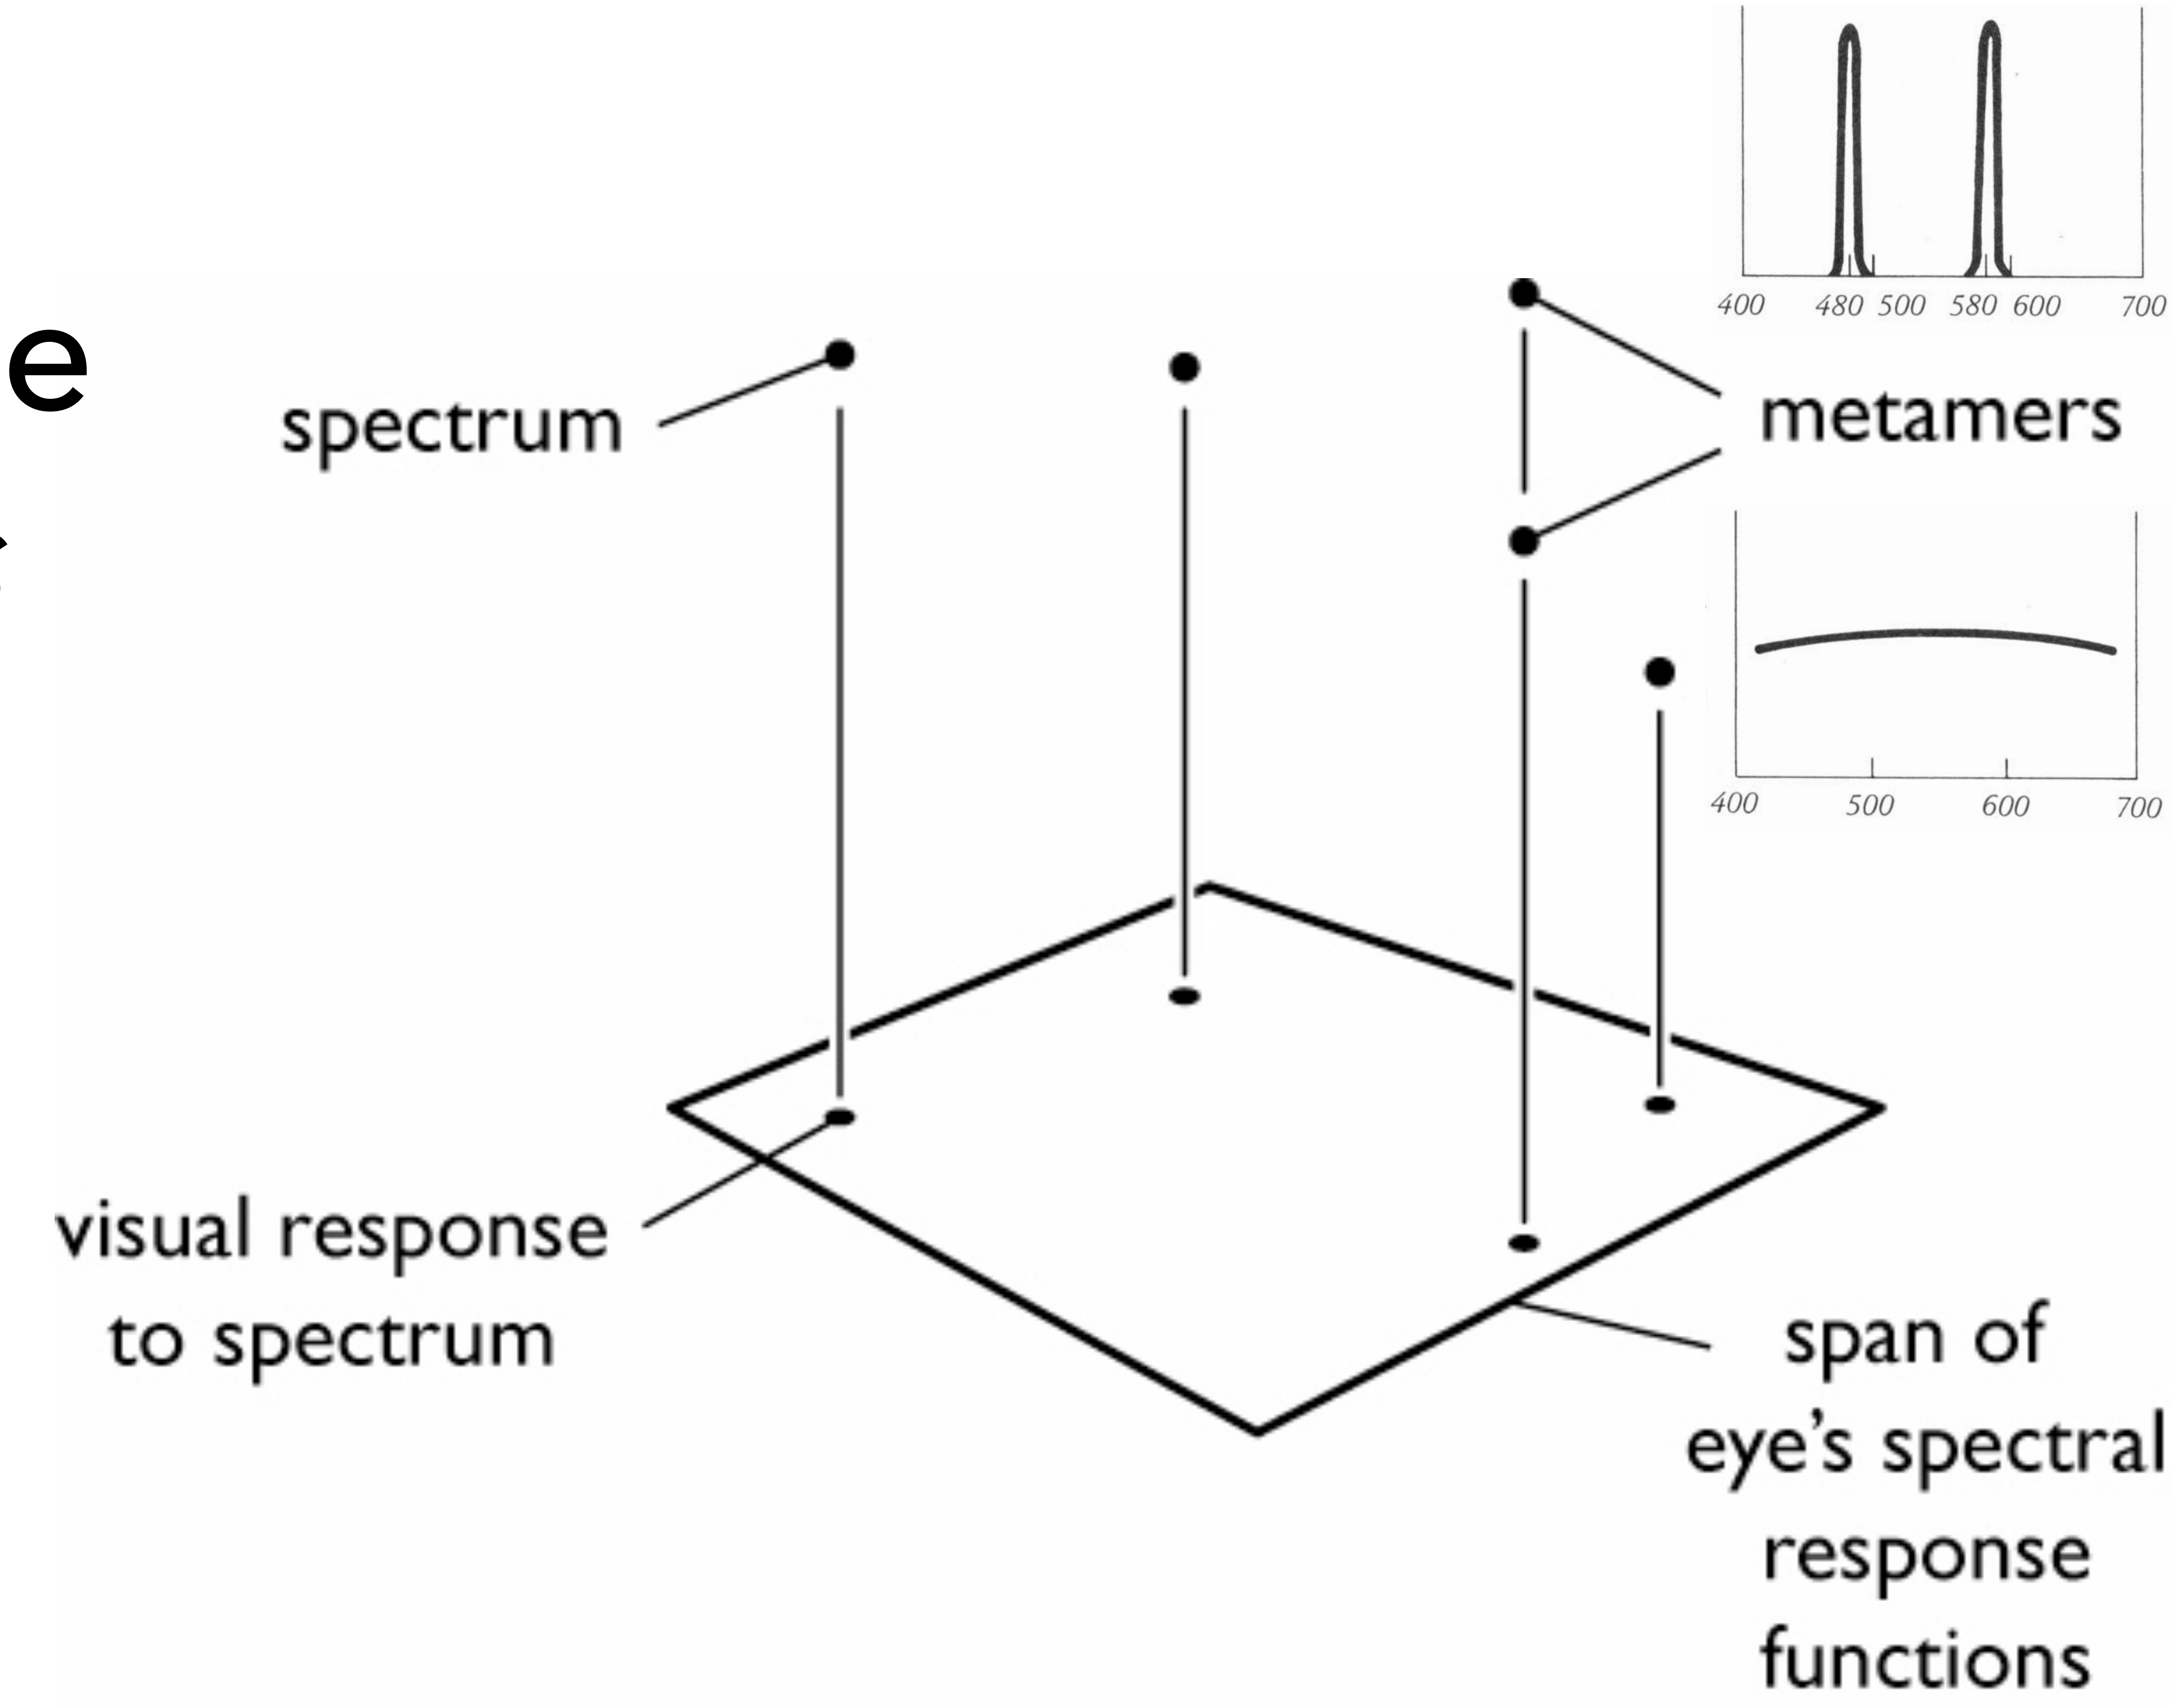

Metamers

Which stimuli are metamers?

Response curve

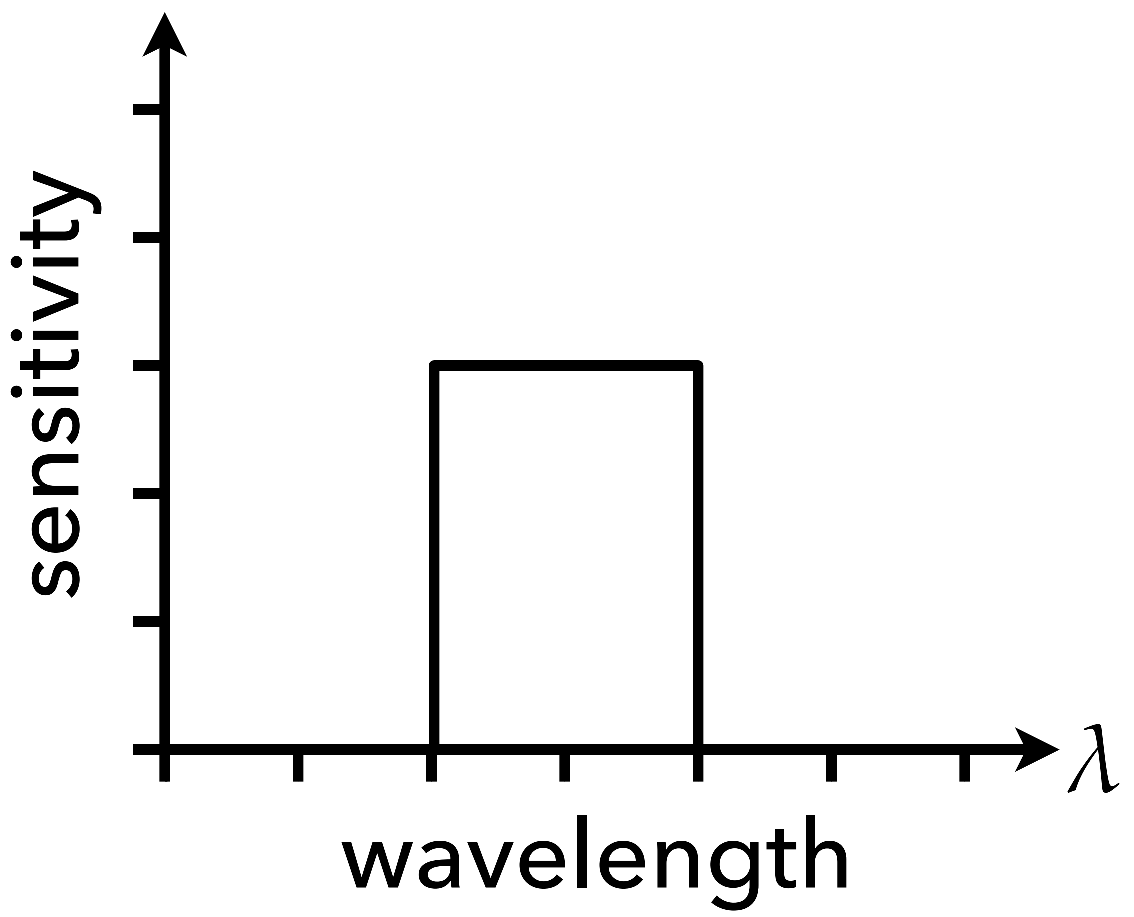

Stimuli

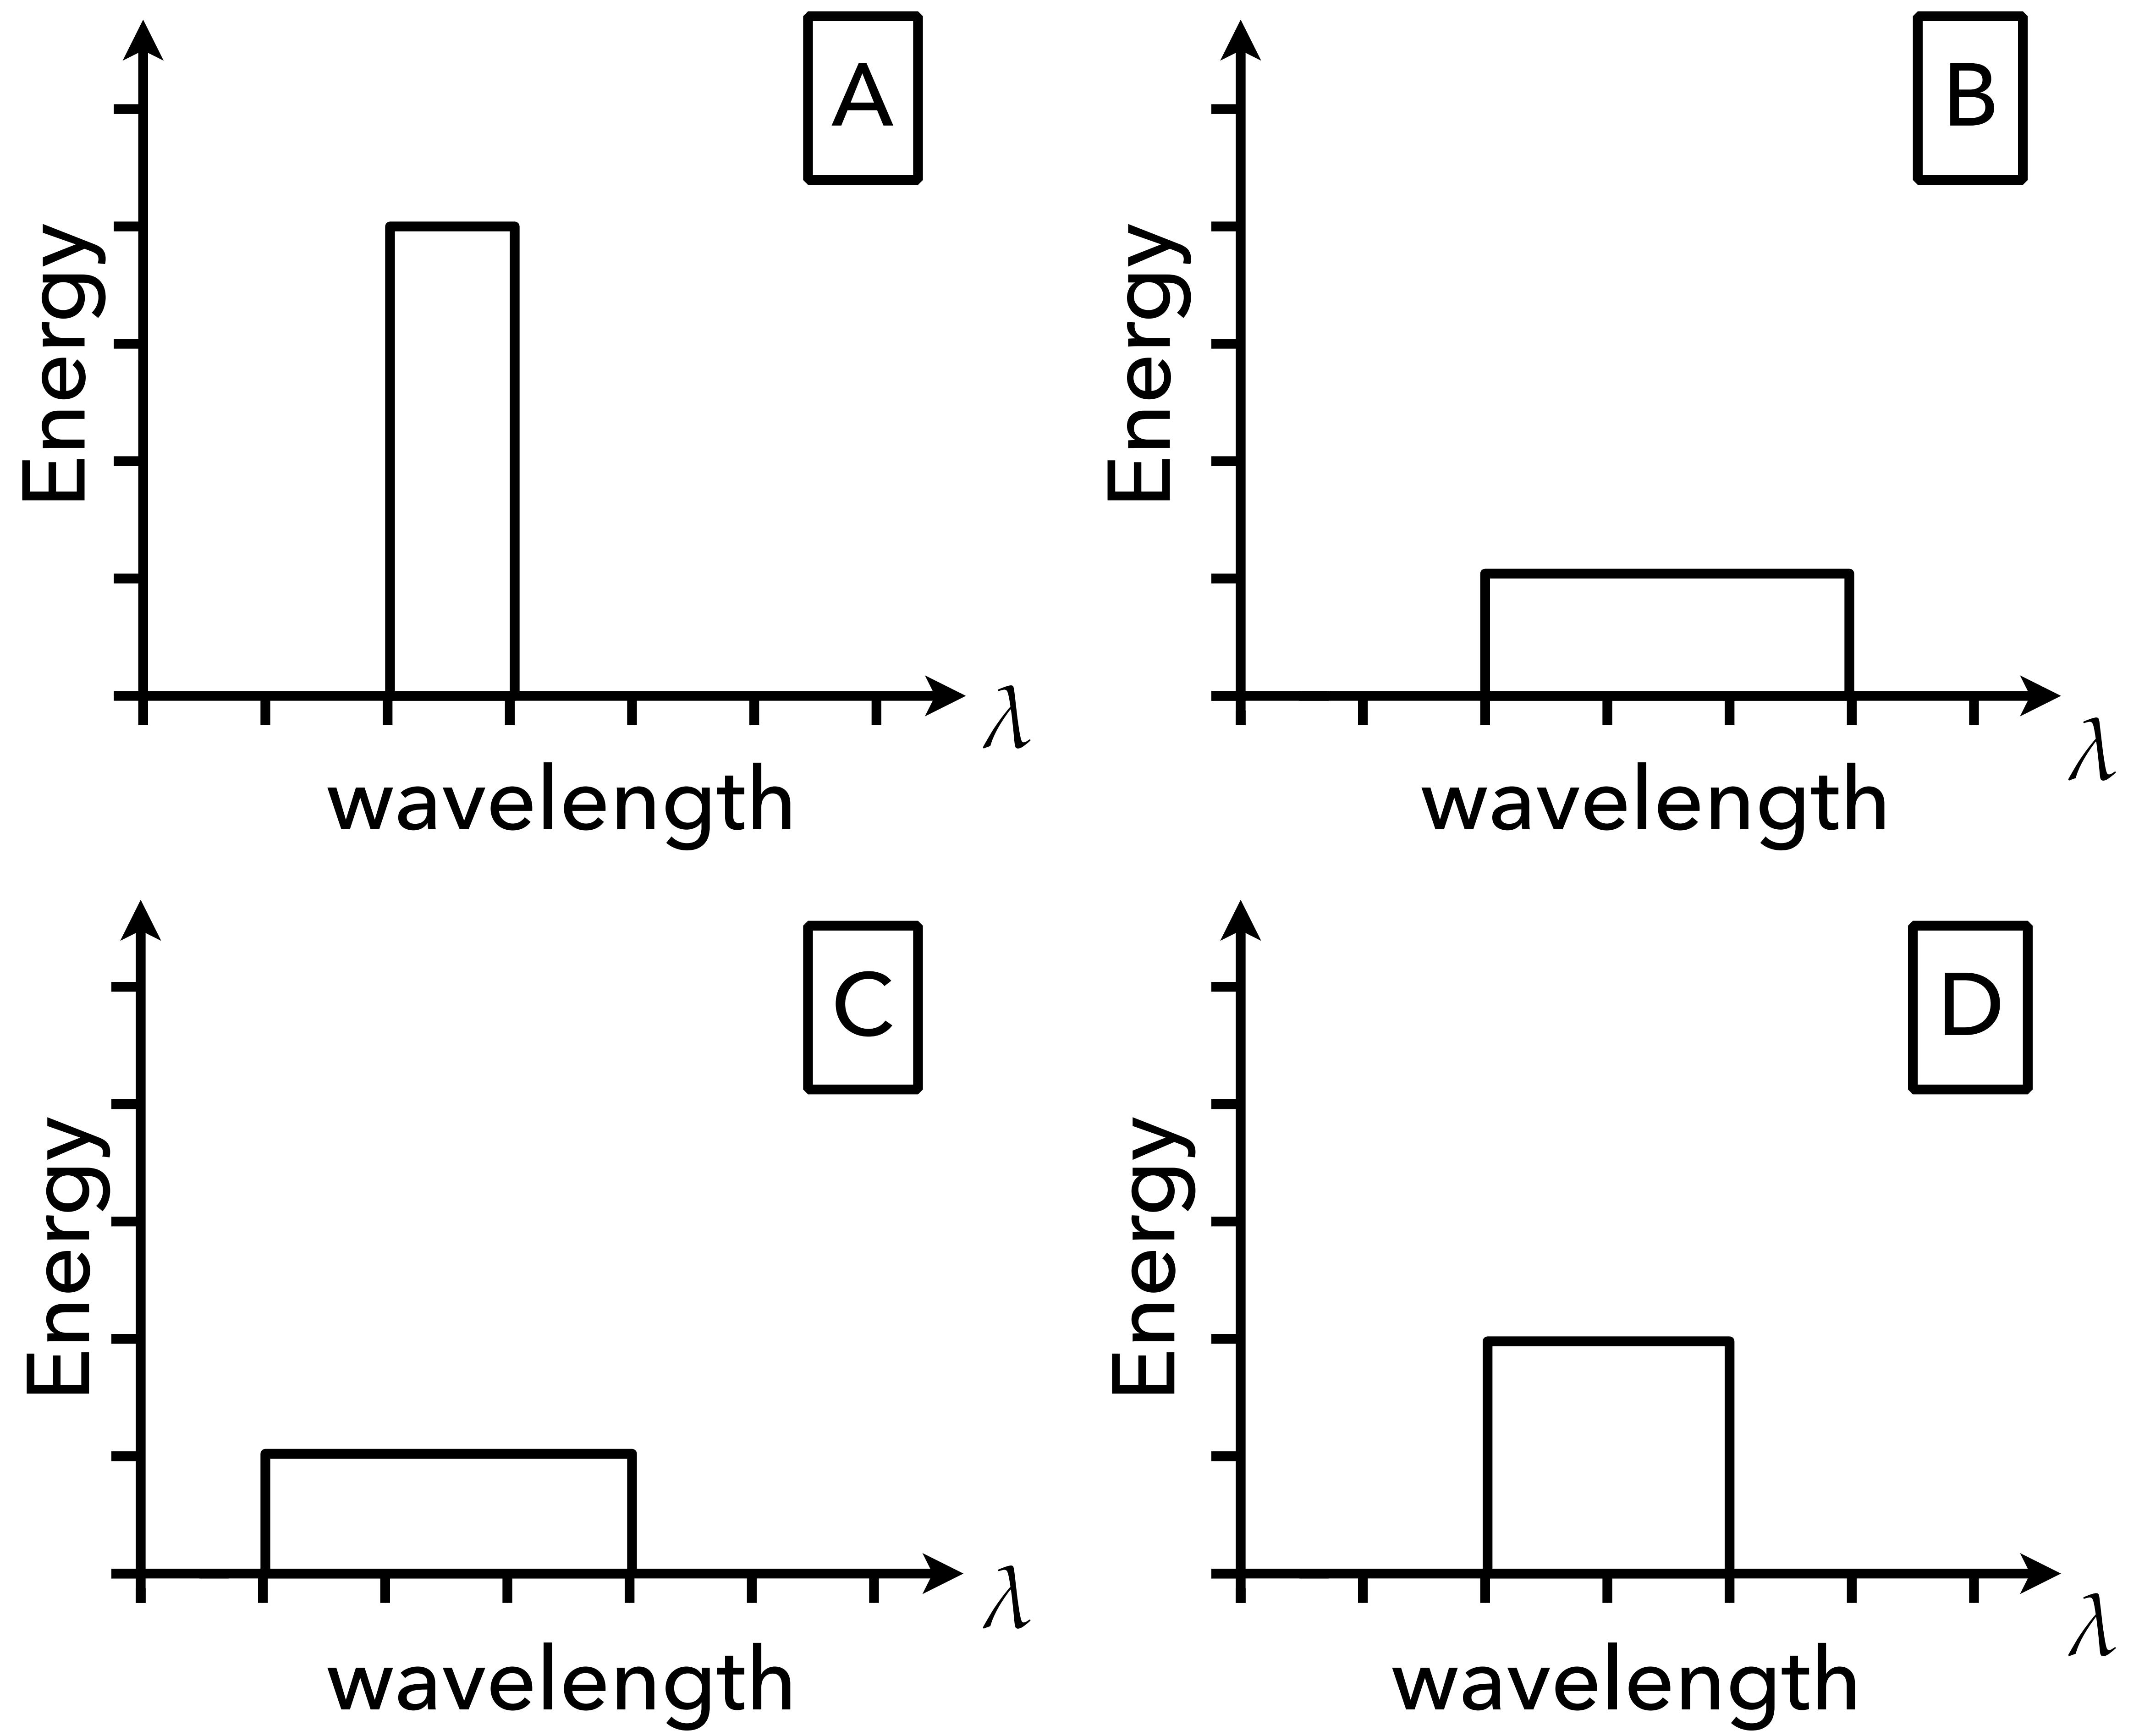

There is an infinity of metamers

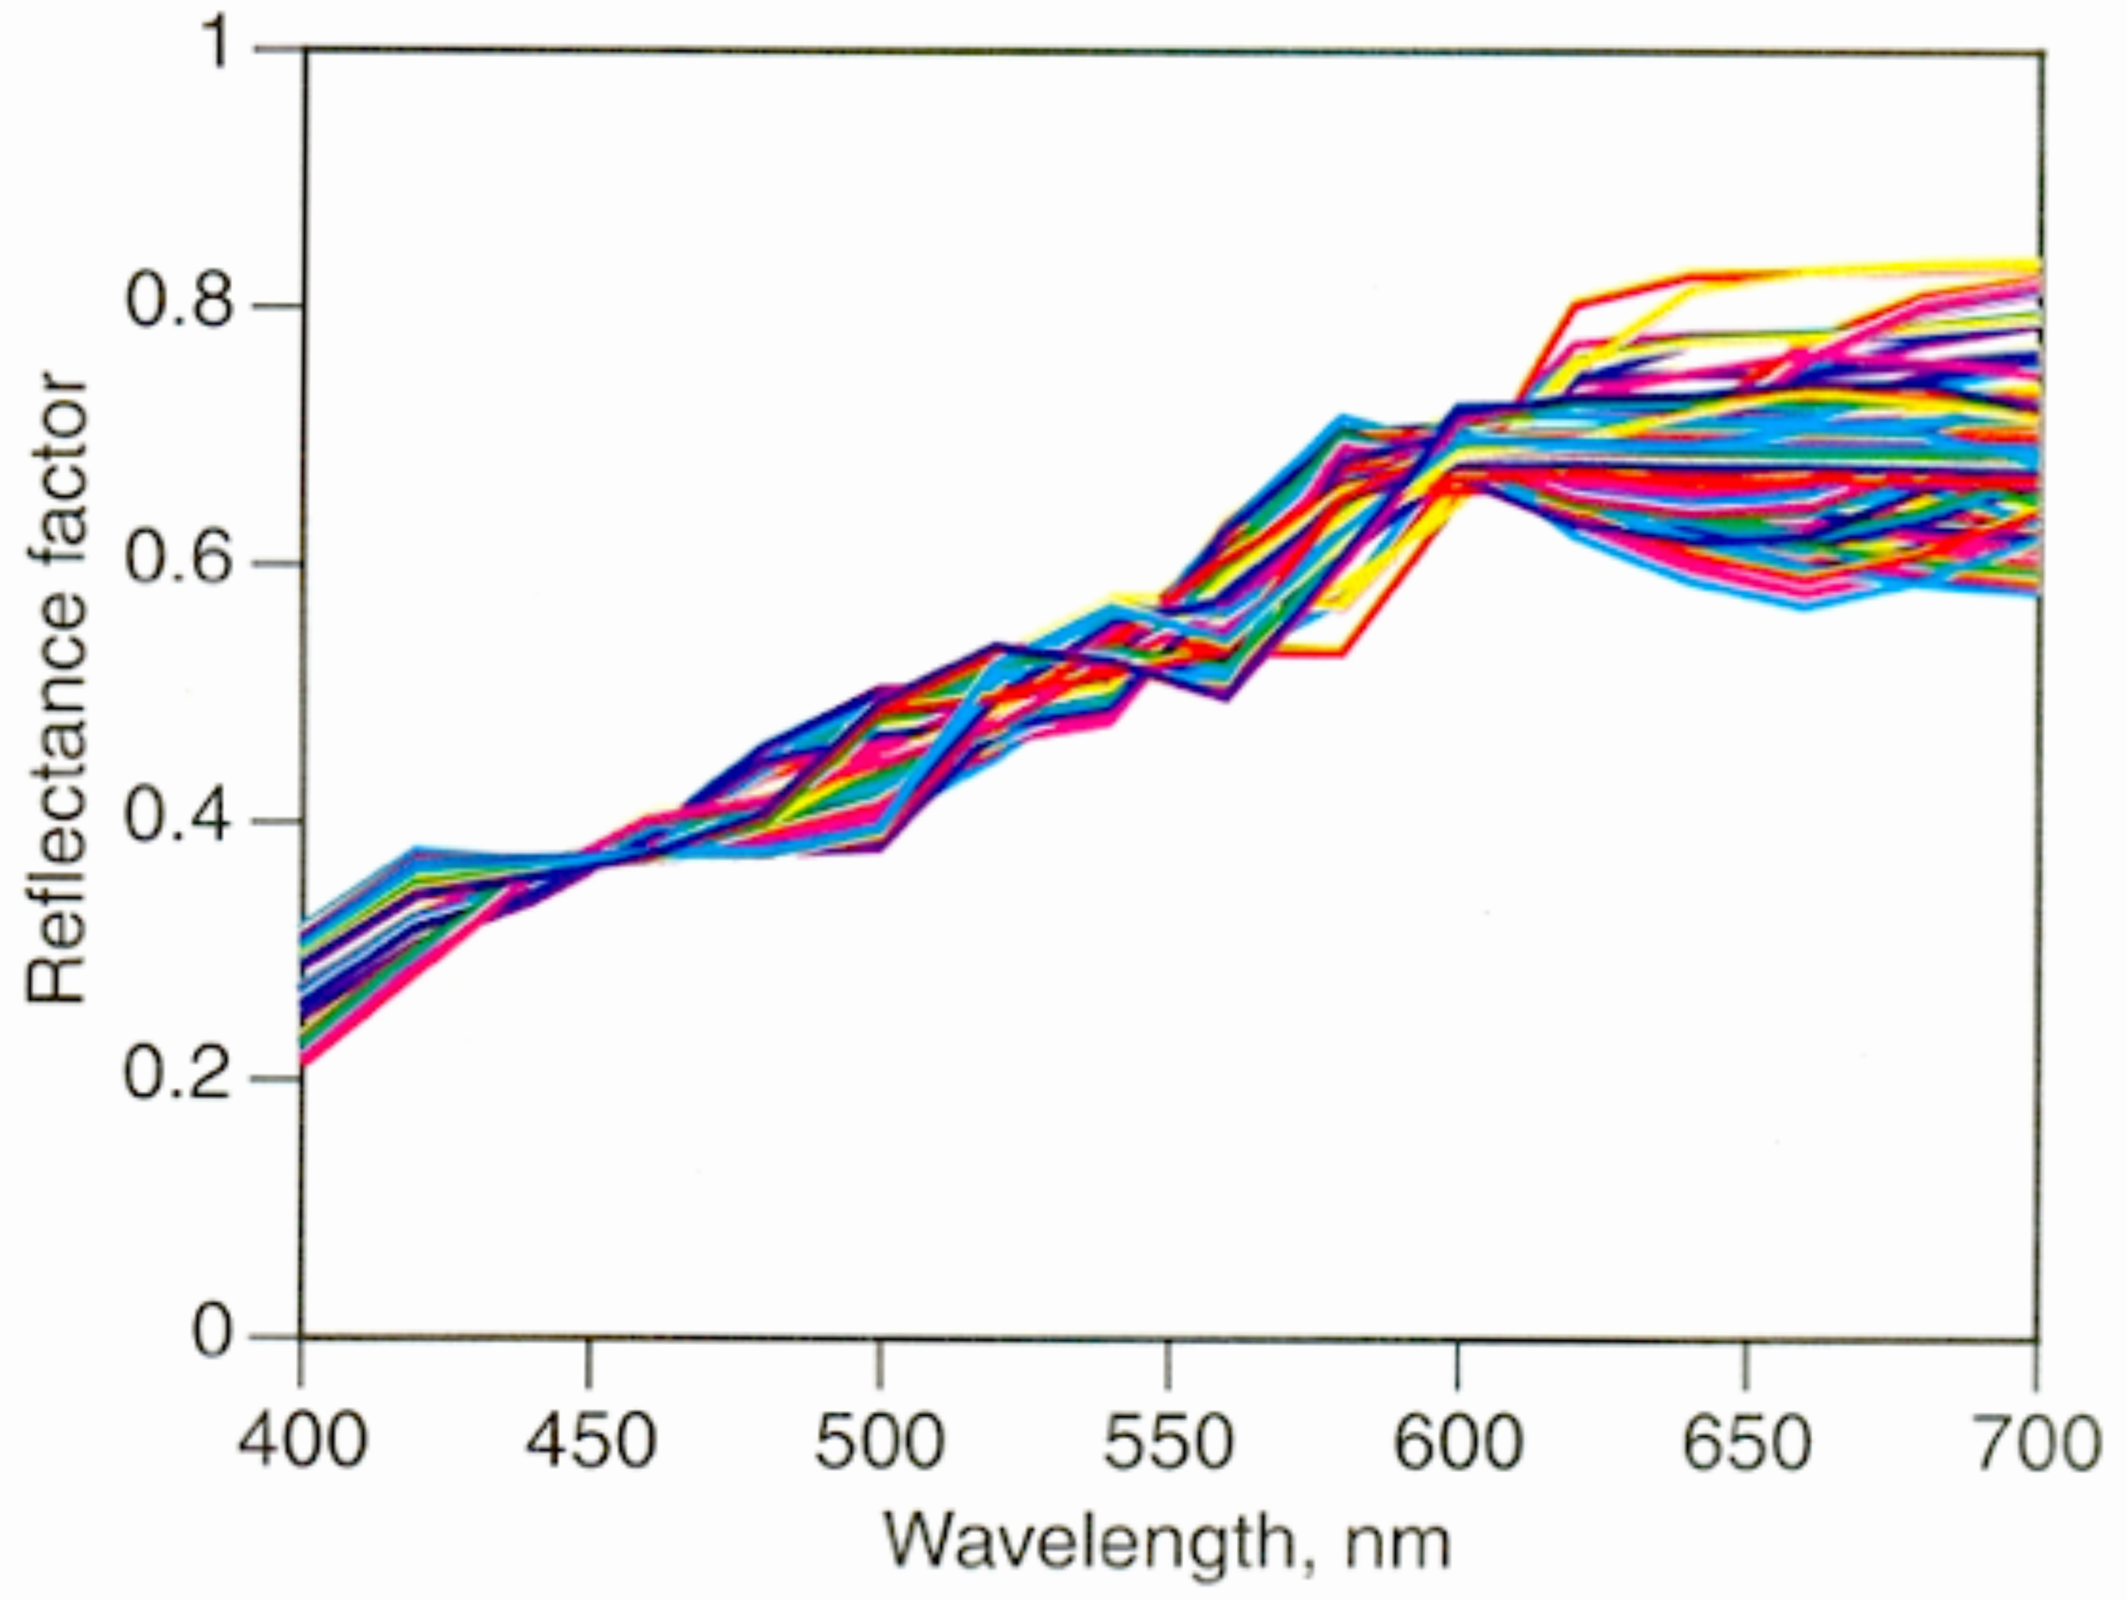

Ensemble of spectral reflectance curves corresponding to three chromatic-pigment recipes all matching a tan material when viewed by an average observer under daylight illumination. [Based on Berns (1988b).]

Good news: color reproduction

3 primaries are (to a first order) enough to reproduce all colors!

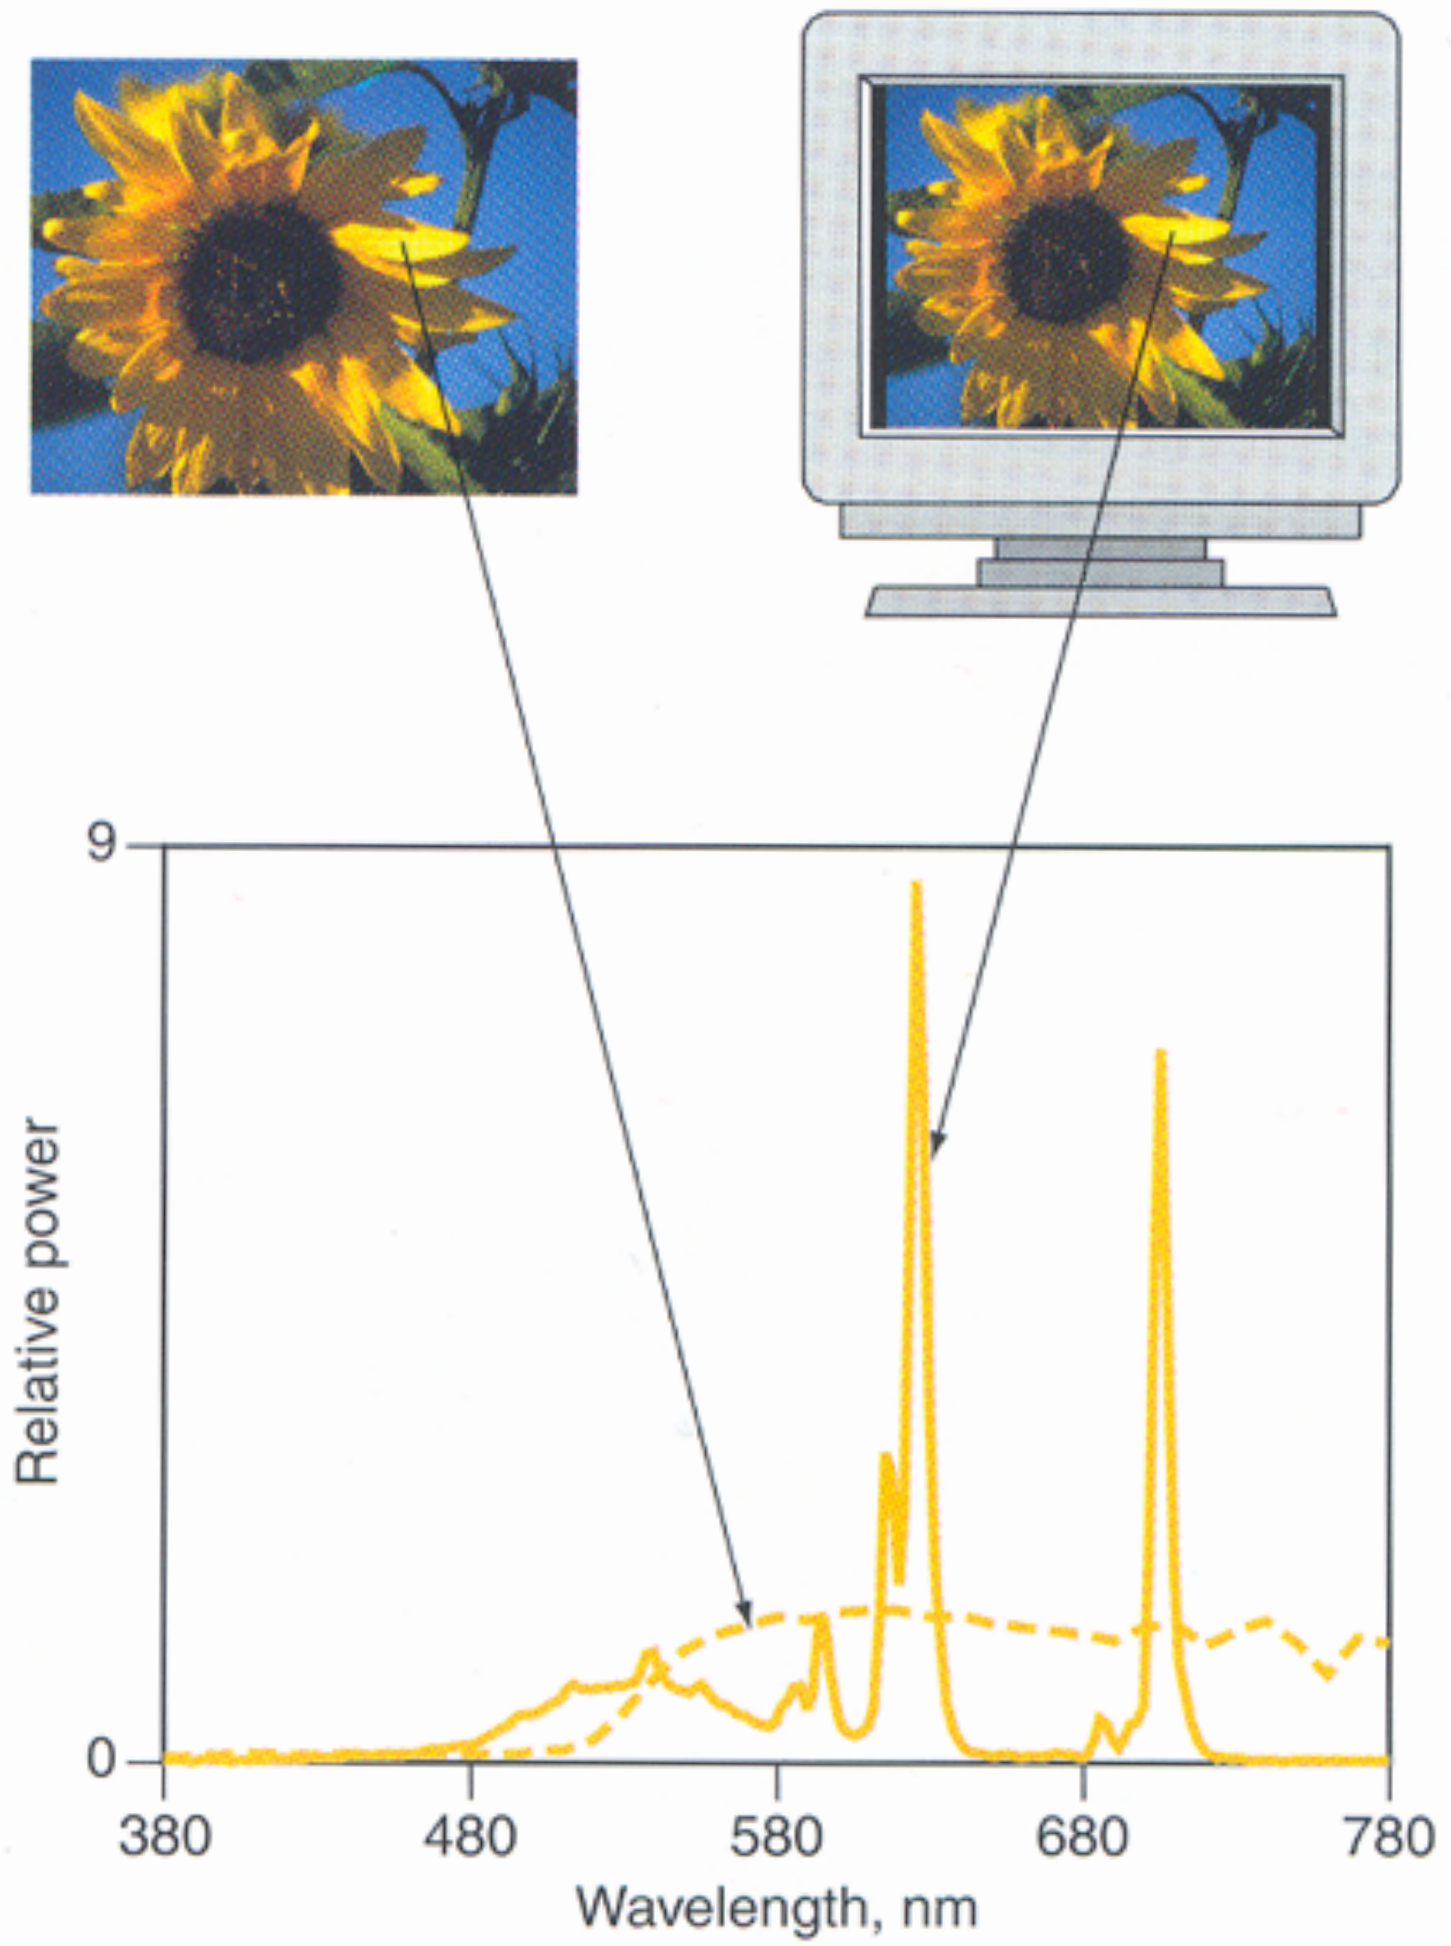

The dashed line represents daylight reflecting from sunflower petals, while the solid line represents the light emitted by a color CRT display adjusted to match the color of the sunflower.

Metamerism & light sources

Metamers under a given light source

May not be metamers under a different lamp

Illuminant metamerism example

Two grey patches in Billmeyer & Saltzman's book look the same under daylight but different under neon or halogen

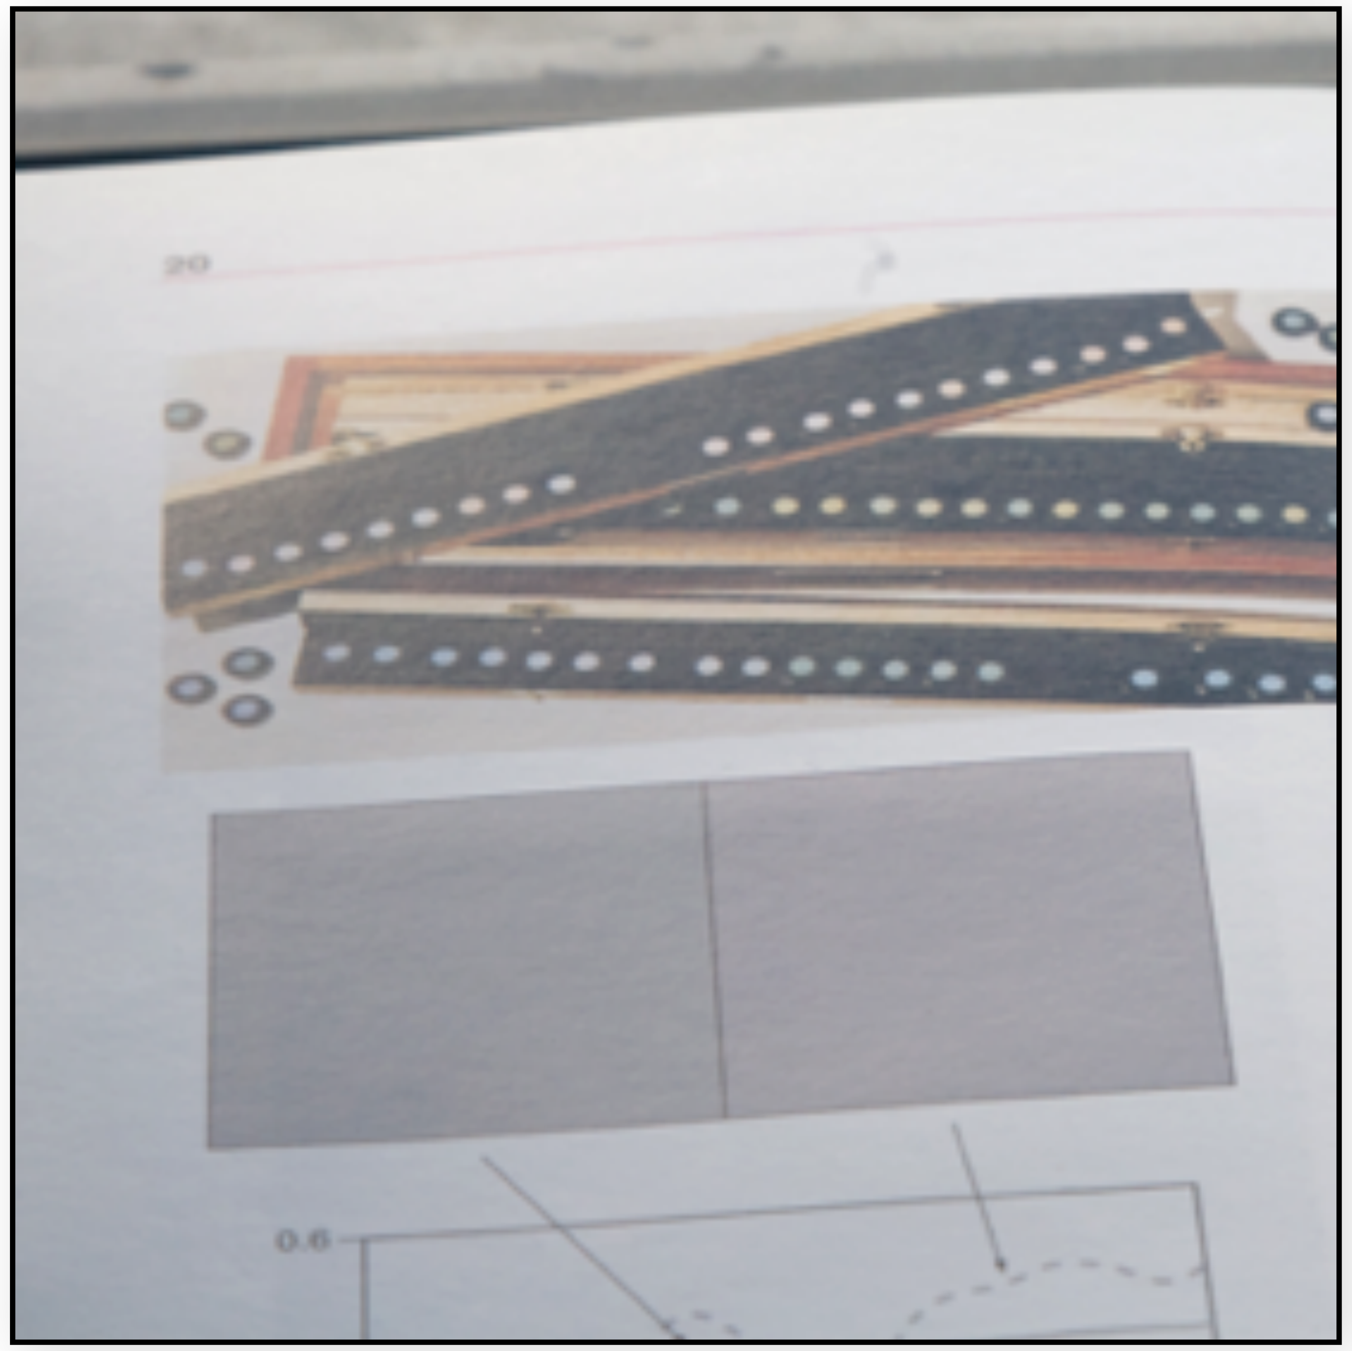

Daylight

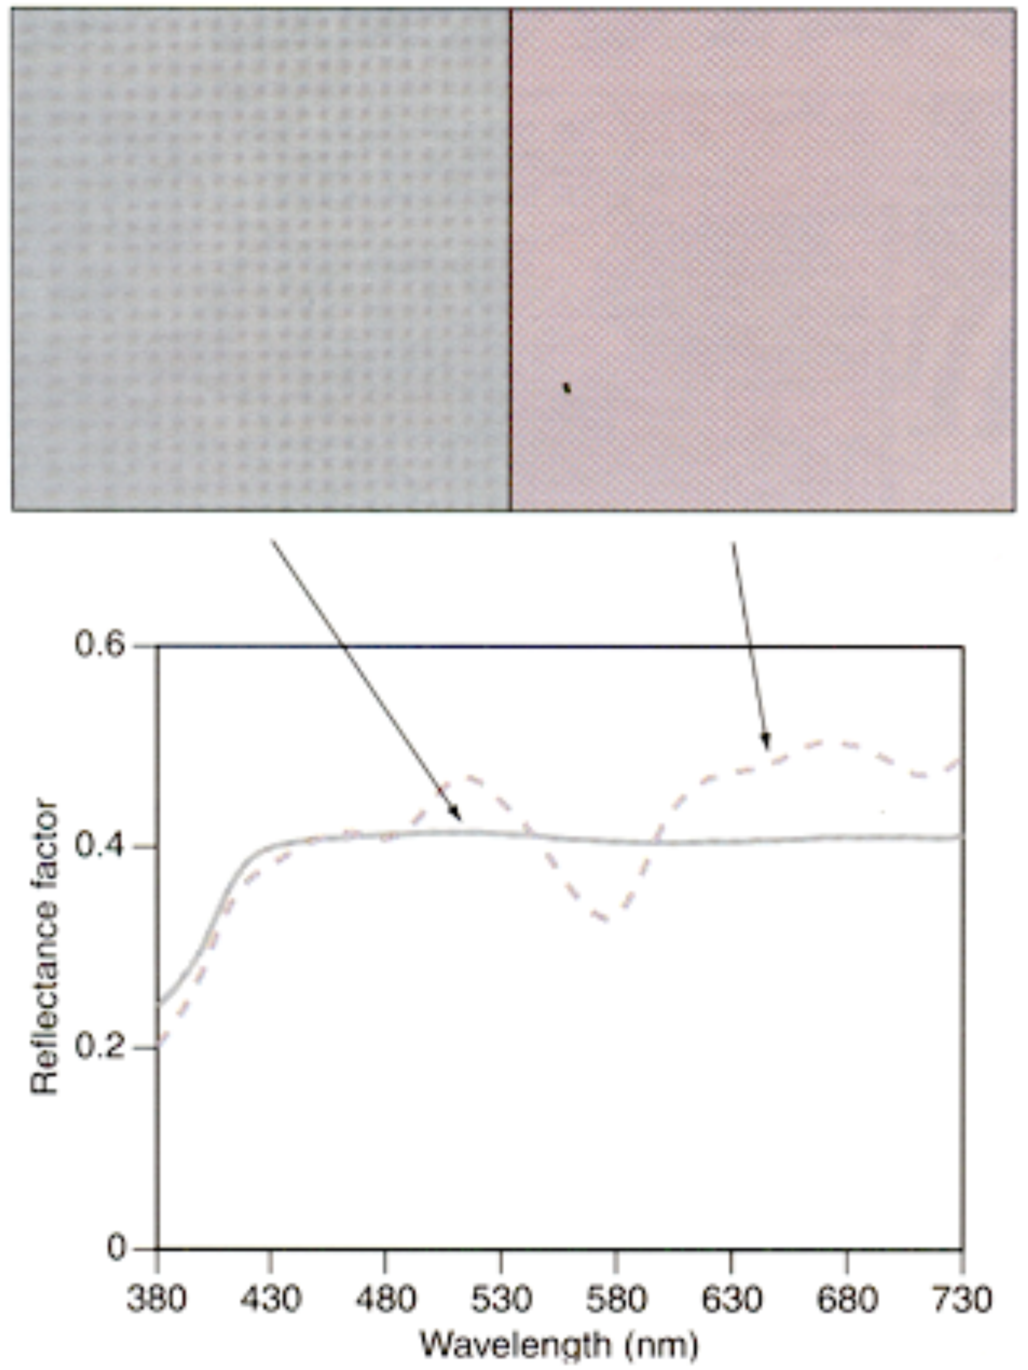

Scan (neon)

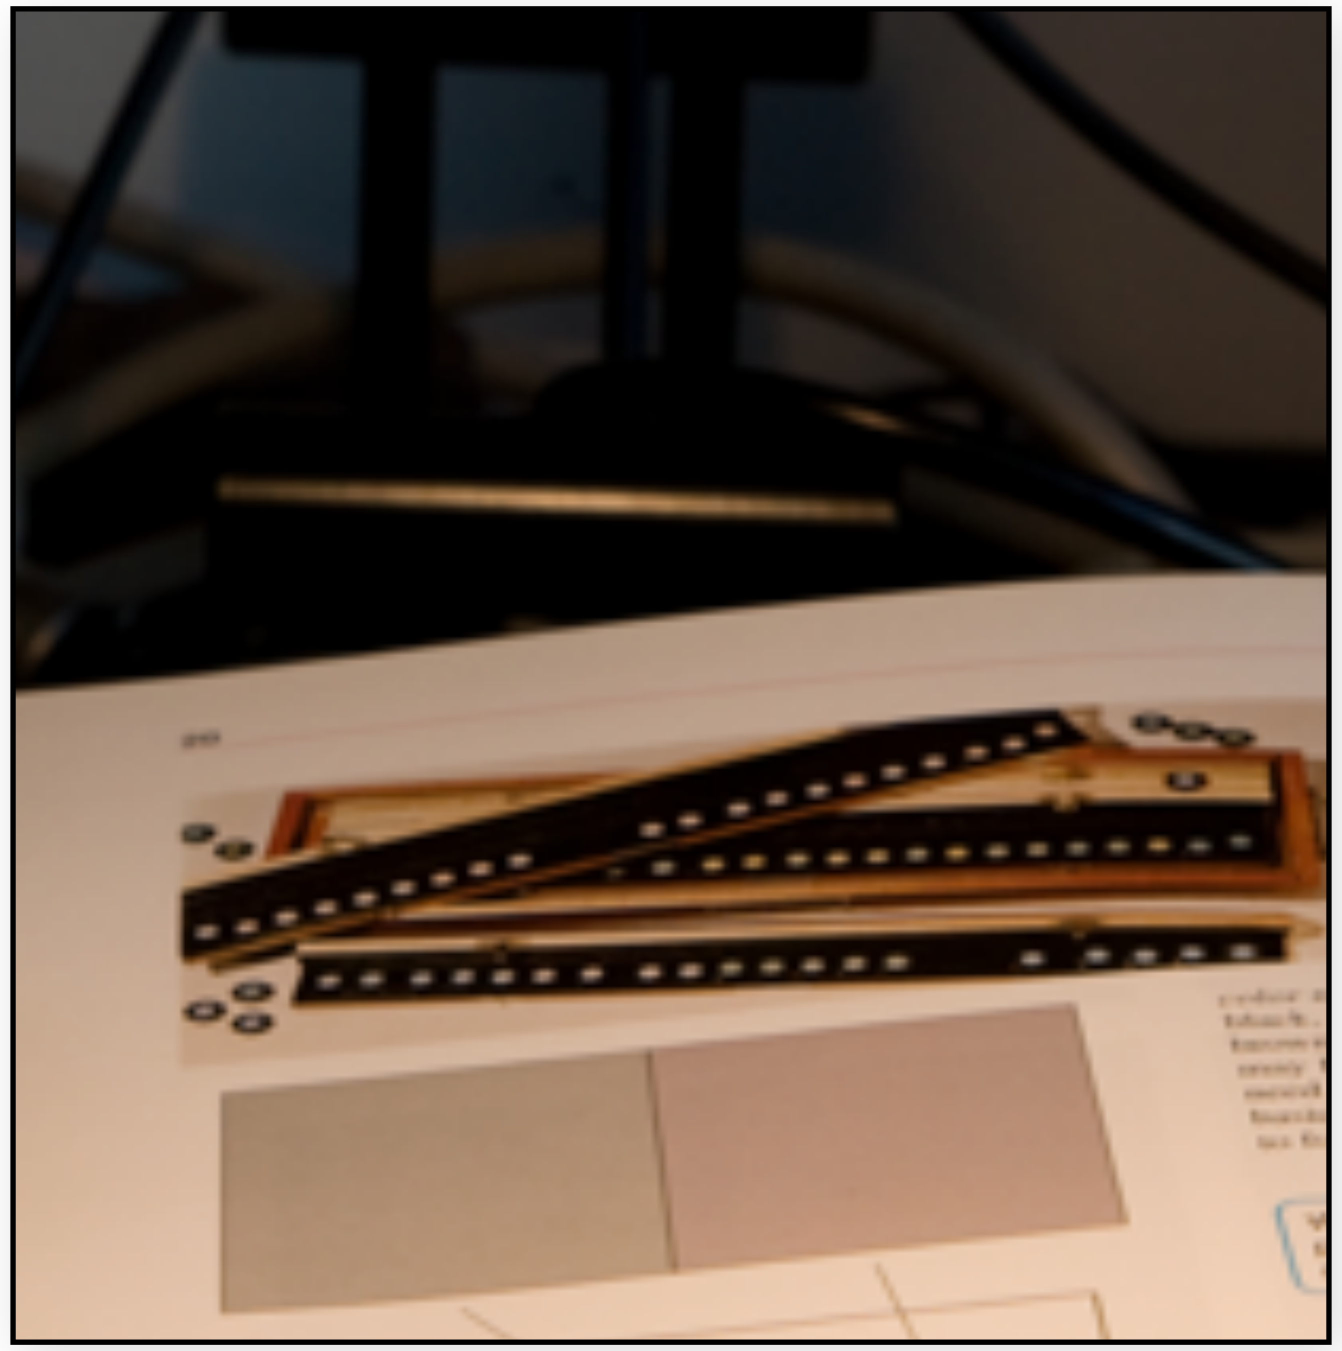

Hallogen

Bad consequence: cloth matching

Clothes appear to match in store (e.g. under fluorescent)

Don't match outdoors

The sun (a "blackbody")

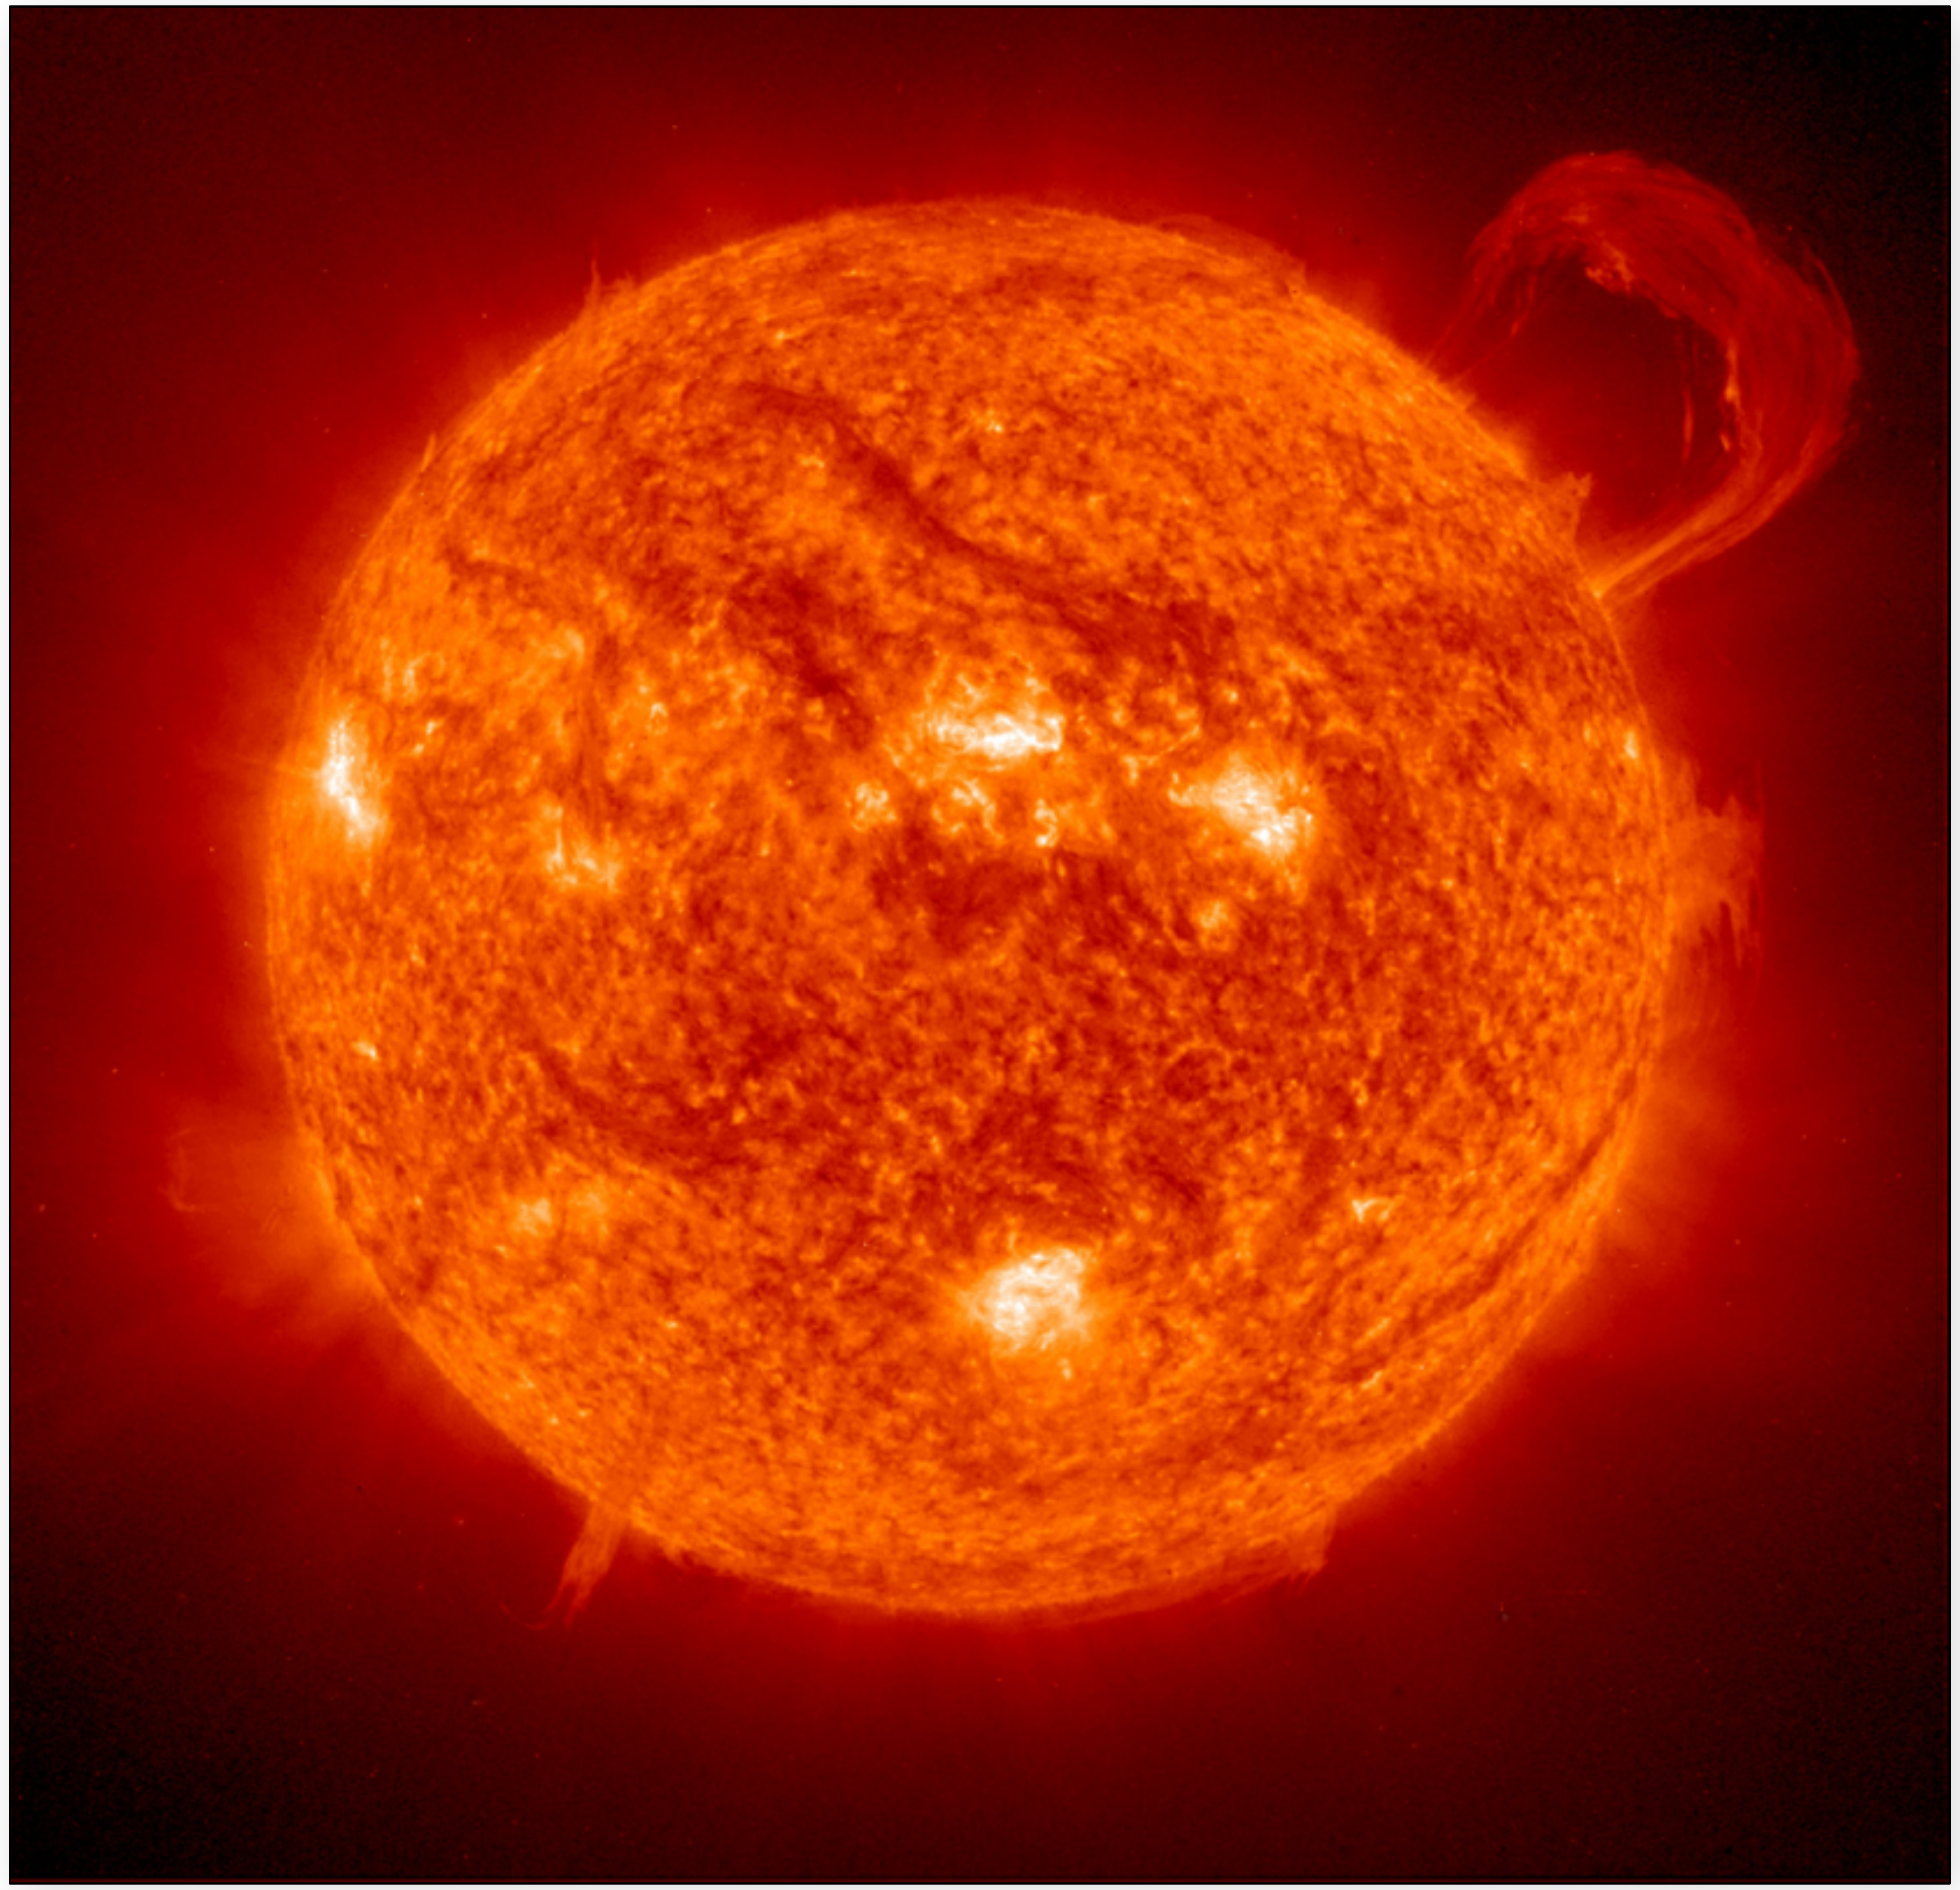

Blackbody Spectrum

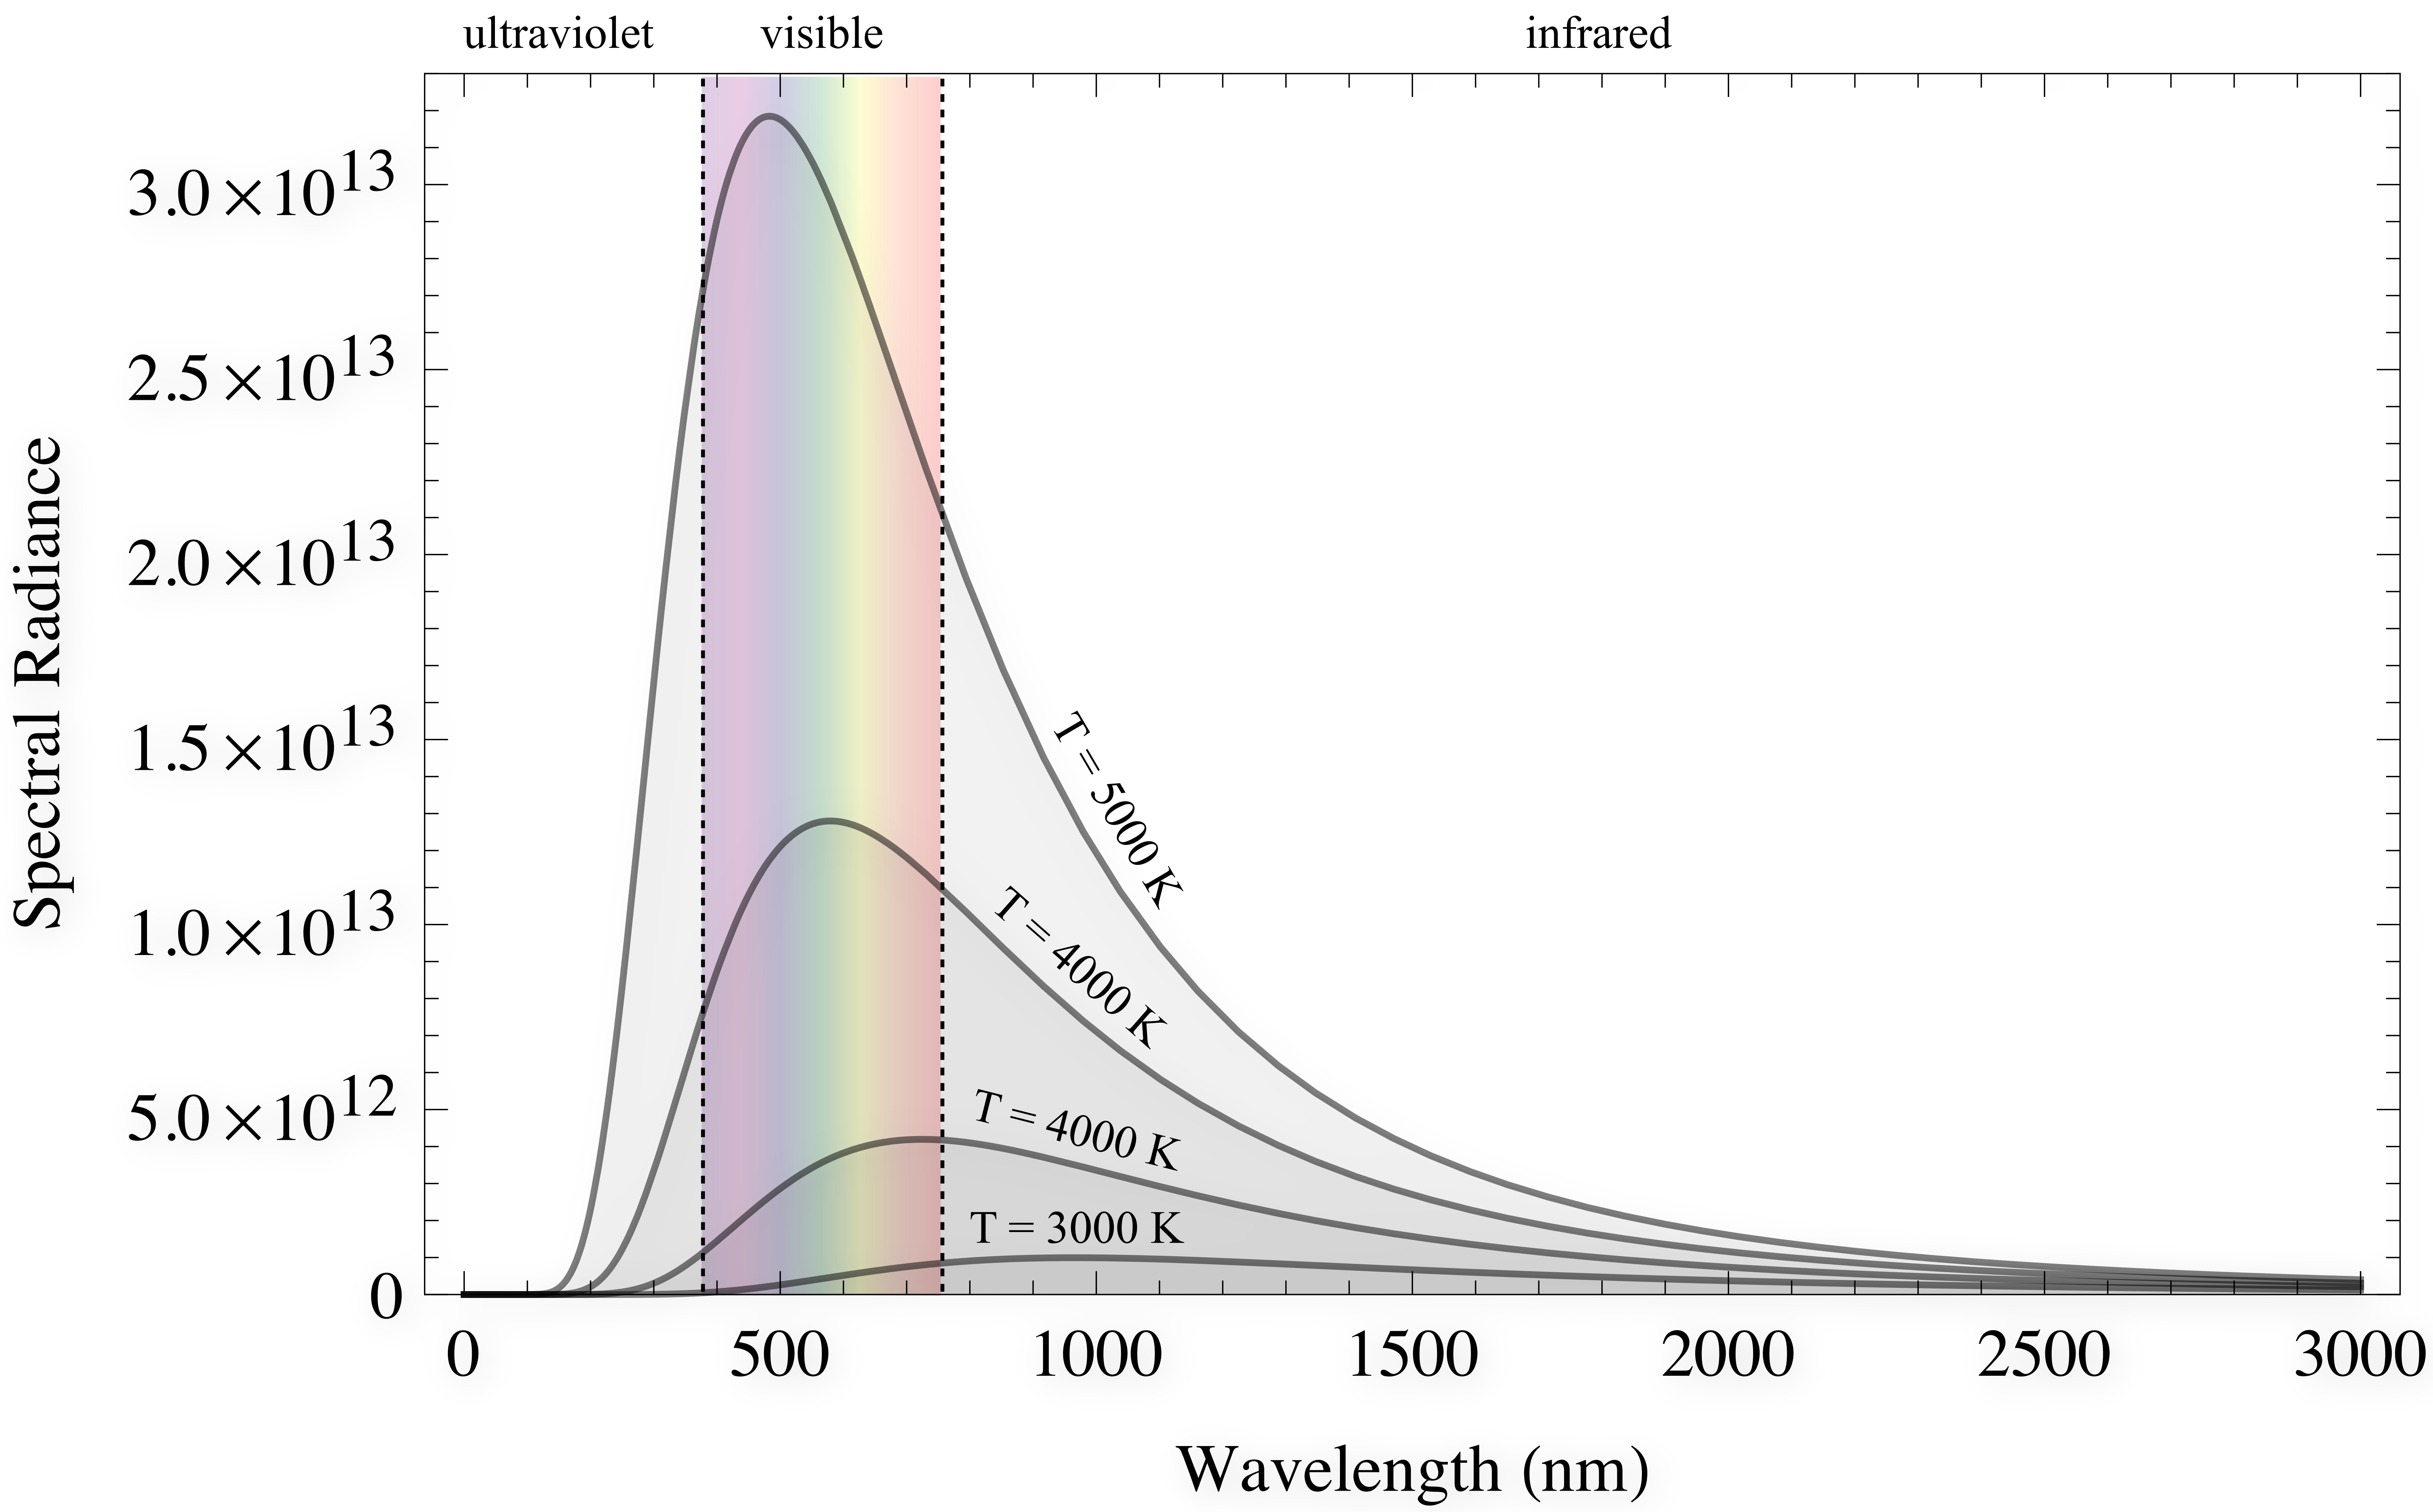

Atomic Emission

Emission spectrum of Hydrogen

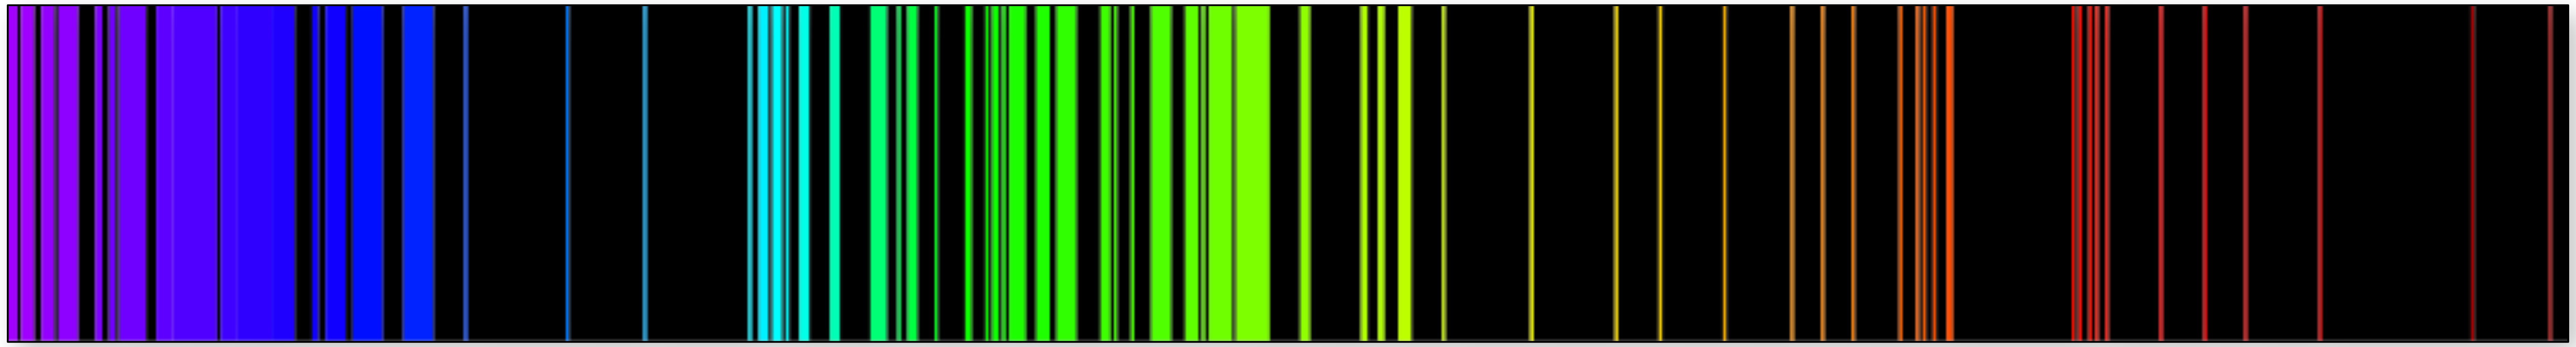

Emission spectrum of Iron

Sodium Vapor Lights

Light emitted at 589nm and 589.6nm

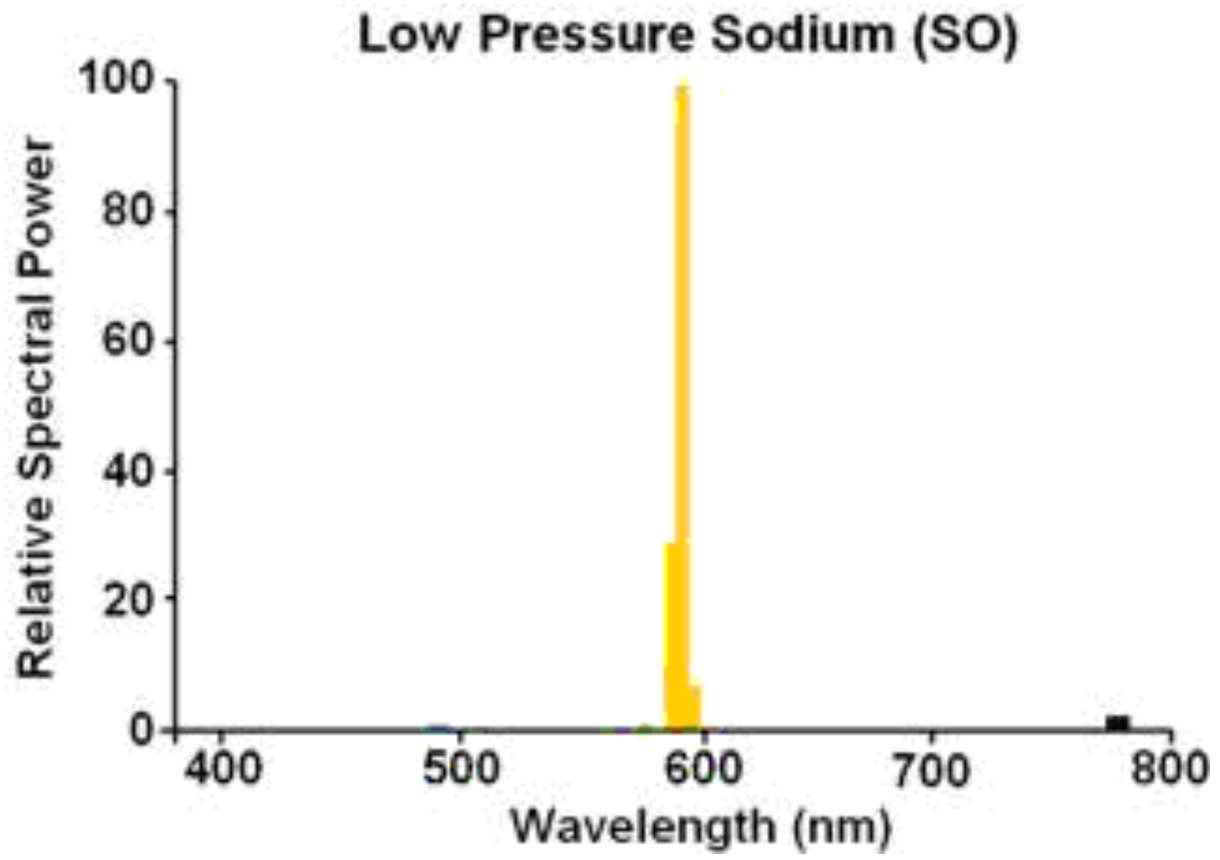

Recap

Spectrum is an infinity of numbers

Projected to 3D cone-response space

- for each cone, multiply per wavelength and integrate
- a.k.a. dot product

Metamerism: infinite-D points projected to the same 3D point (different spectrum, same perceived color)

- affected by illuminant
- enables color reproduction with only 3 primaries

Color perception in the animal kingdom

Humans project s(λ) into a 3D subspace

- Some people (only women) are tetrachromats (4 types of cones)!

Most mammals have 2 types of cones (2D subspace)

Many birds have UV receptors, some can see magnetic fields

Some animals have even more:

- Mantis Shrimp use an 8D subspace!

DOGHOUSEDIARIES

Today

Light & Color

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

Analysis & Synthesis

We want to measure & reproduce color as seen by humans

No need for full spectrum!

Only need to match up to metamerism

Additive color

We will focus on additive color

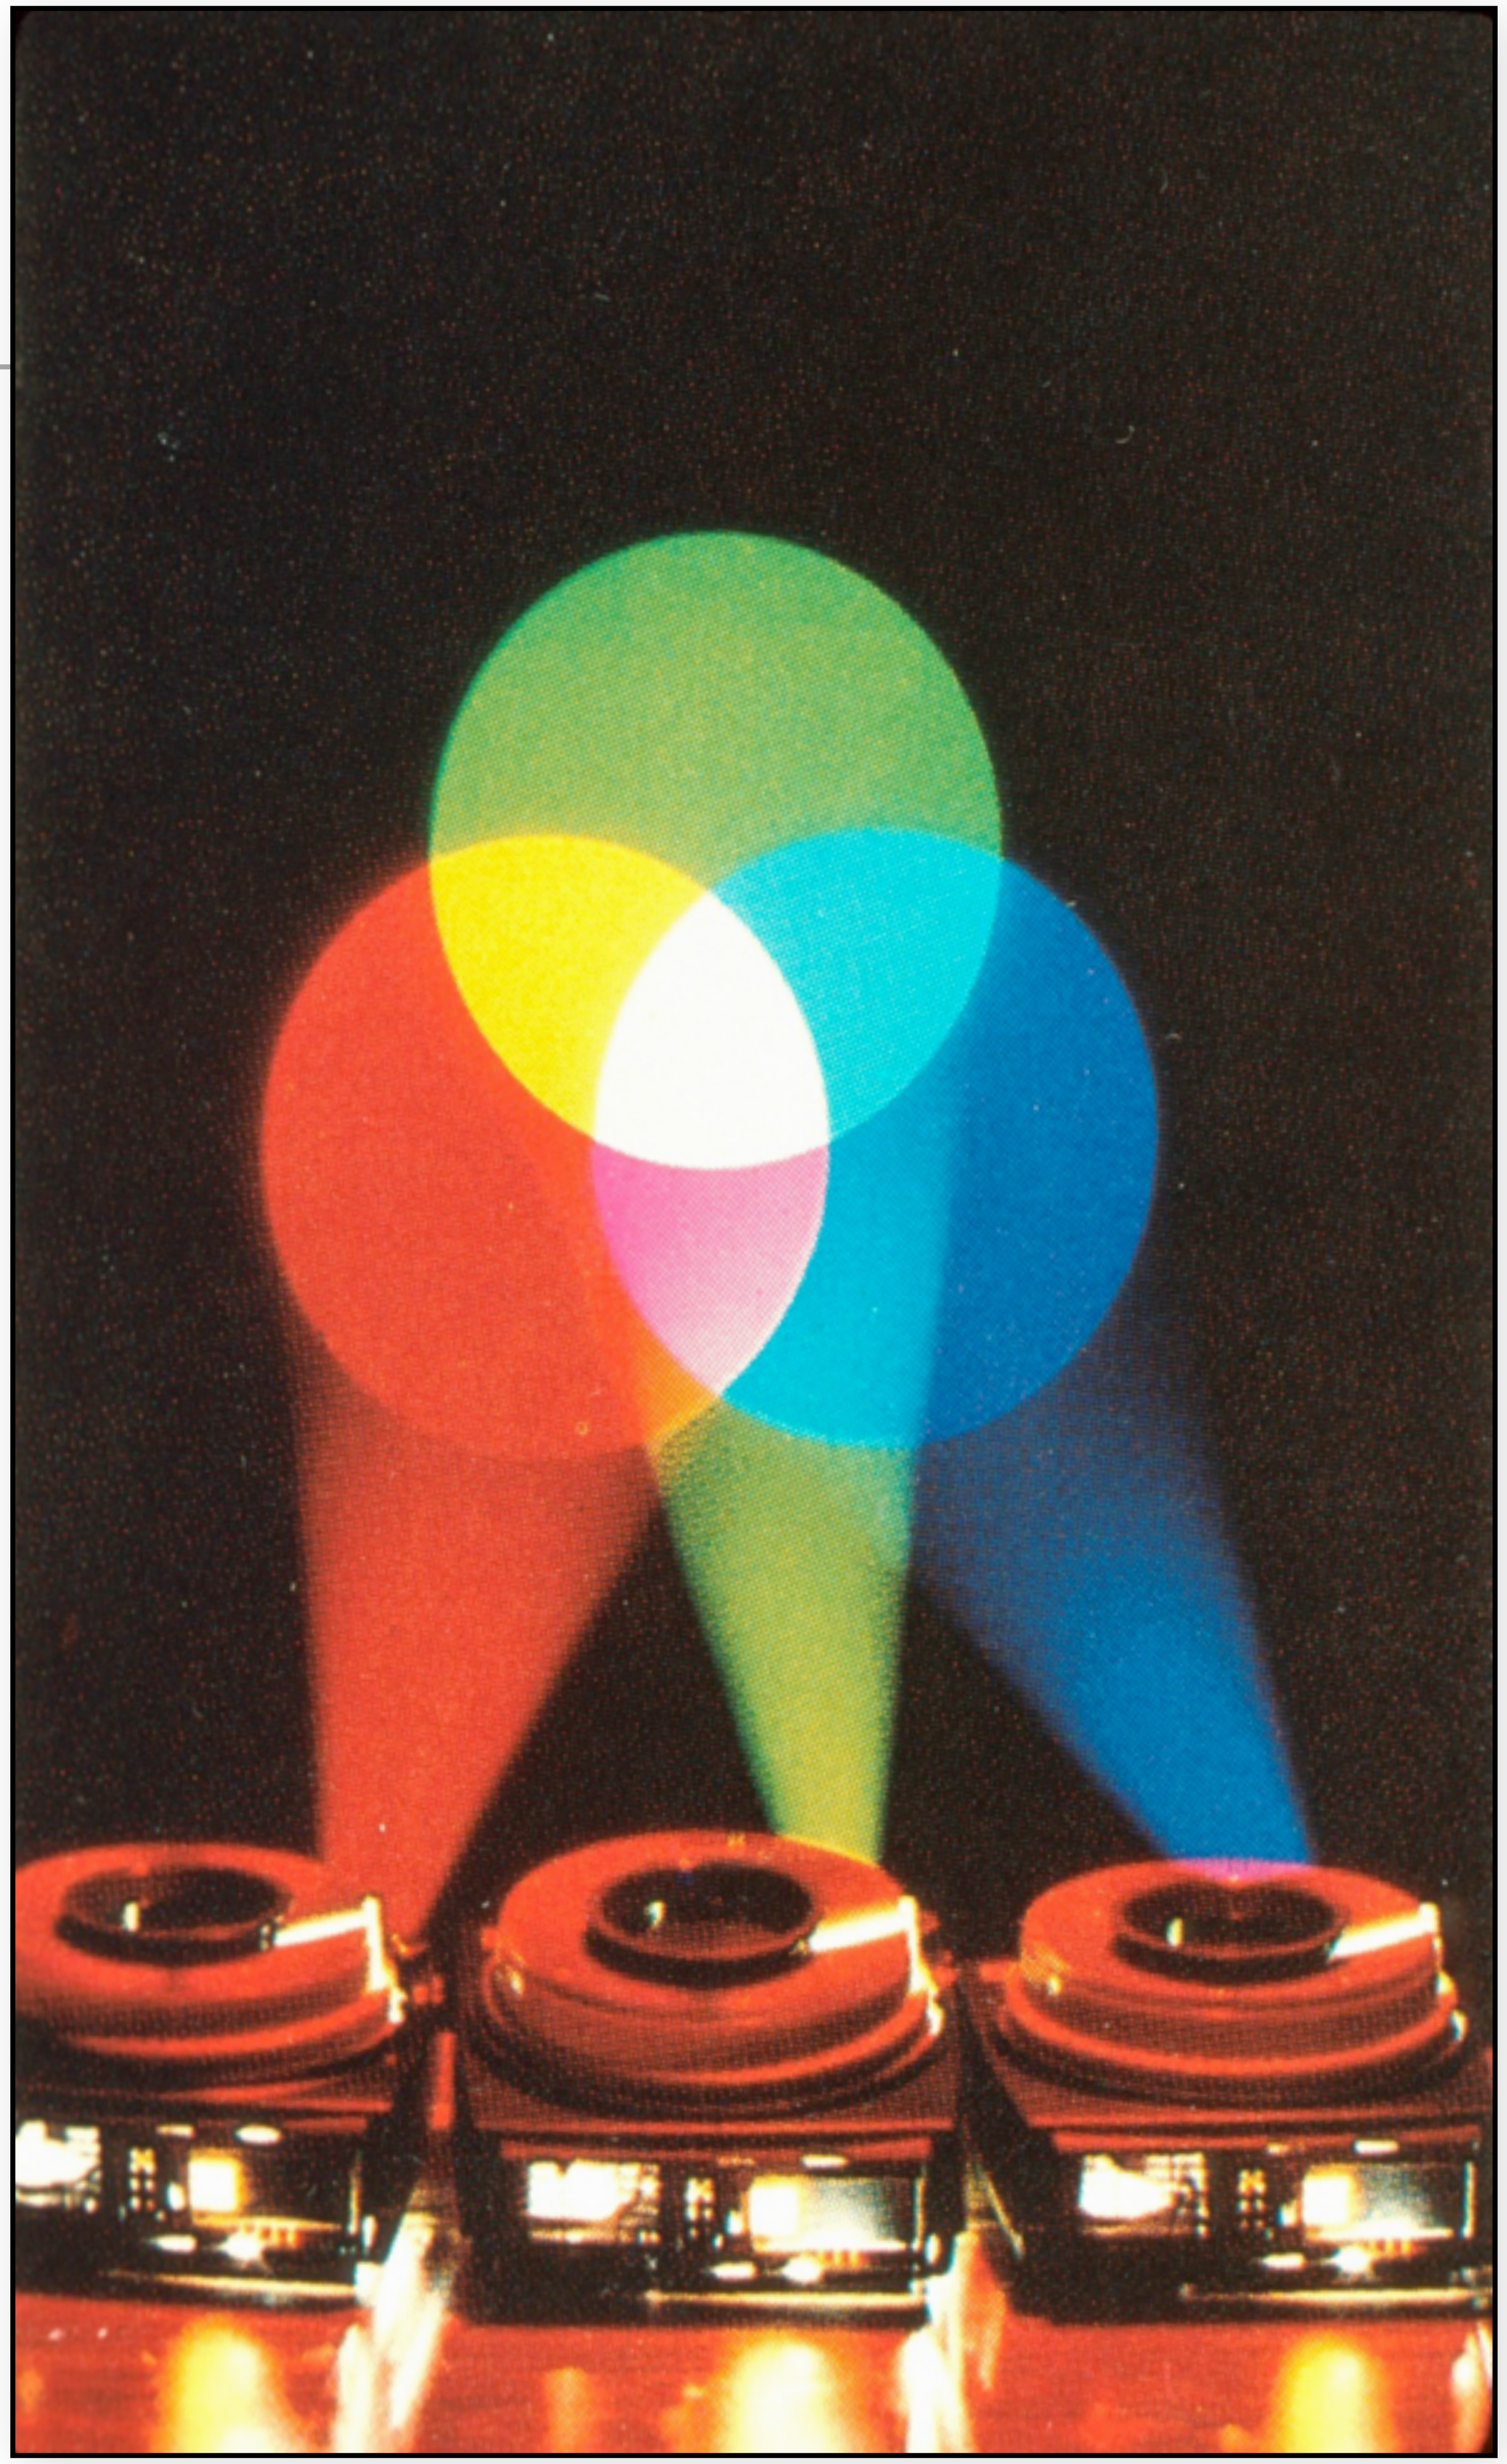

Analysis & Synthesis

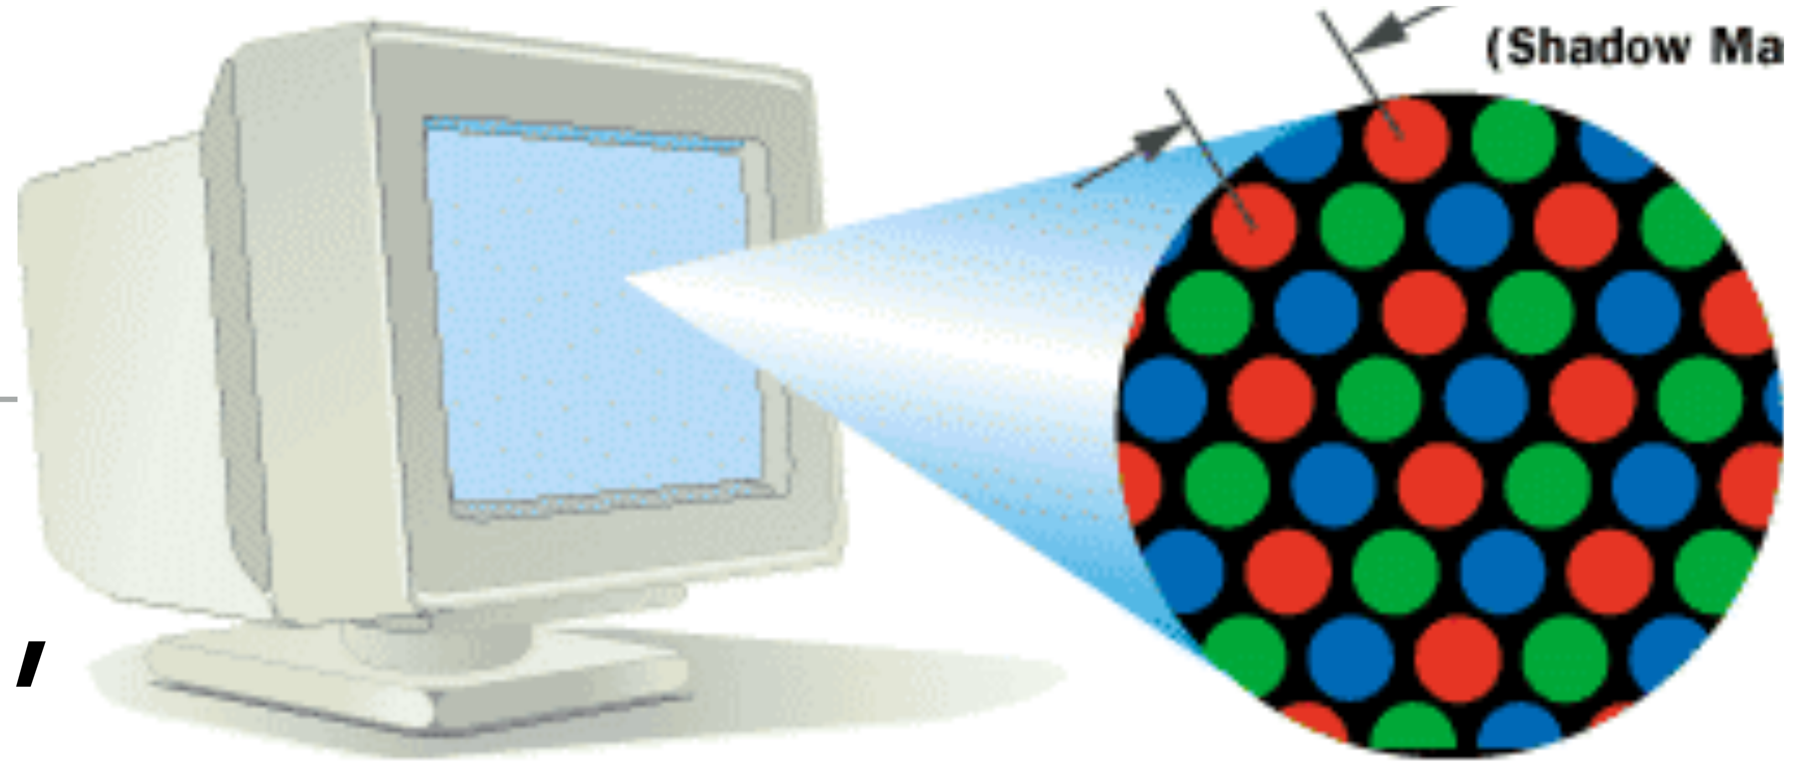

We'll use 3 primaries (e.g. red, green, and blue) to match all colors

- What should those primaries be?
- How do we tell the amount of each primary needed to reproduce a given target color?

Additive Synthesis (the wrong way!)

Take a given stimulus and the corresponding responses S, M, L (here 0.5, 0, 0)

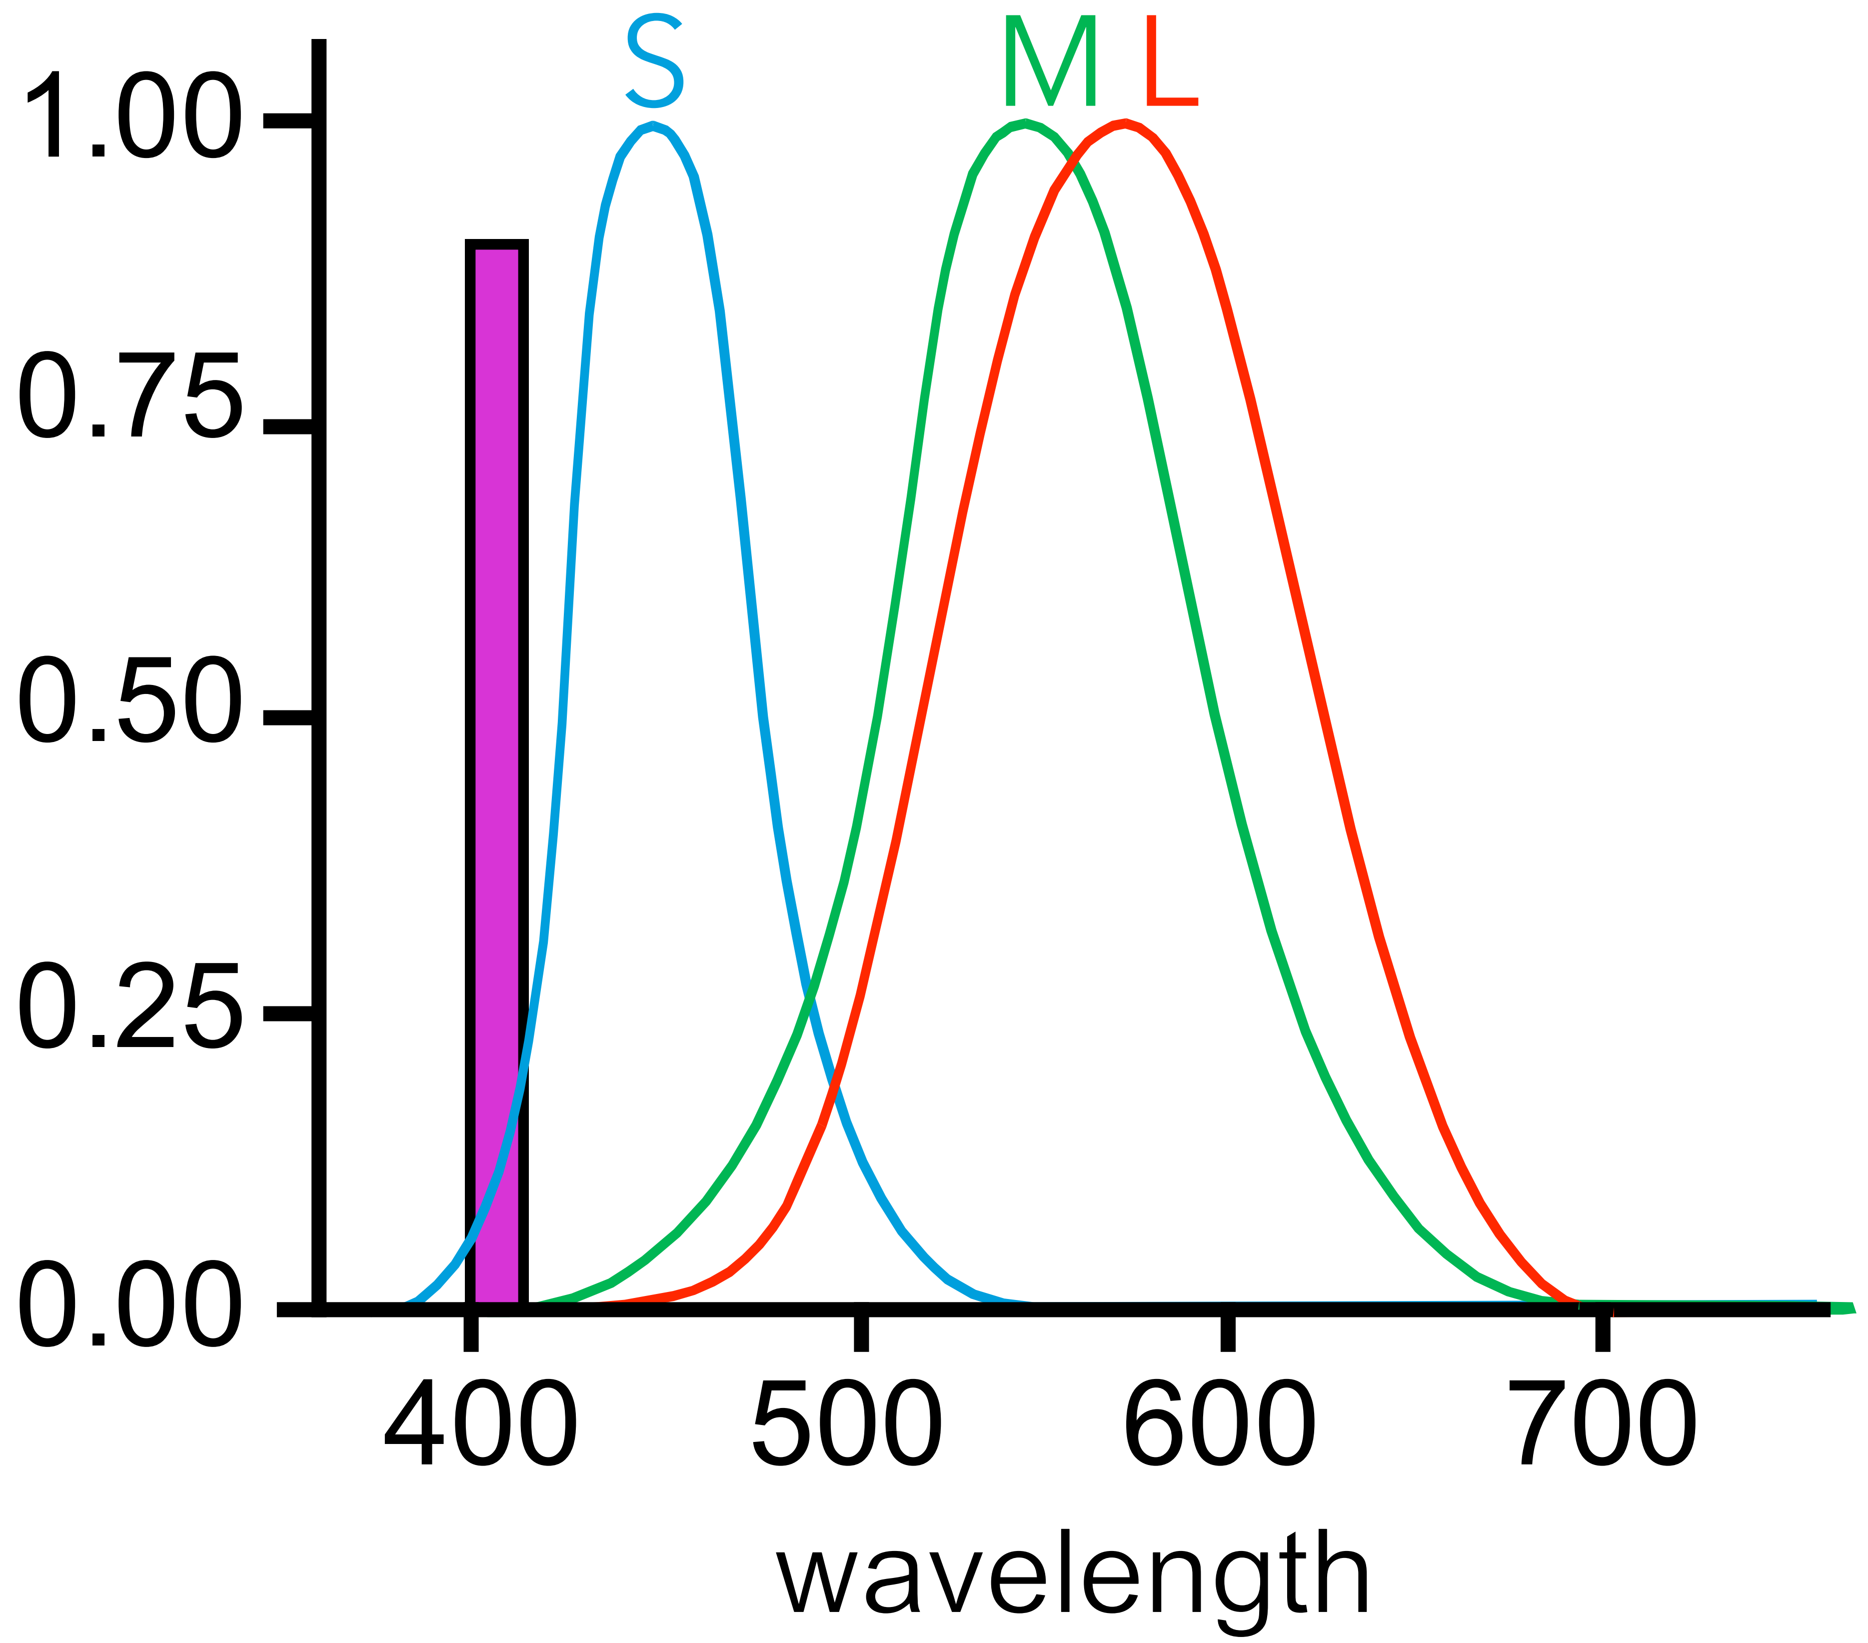

Additive Synthesis (the wrong way!)

Use it to scale the cone spectra (here 0.5 * S)

You don't get the same cone response!
(here 0.5, 0.1, 0.1)

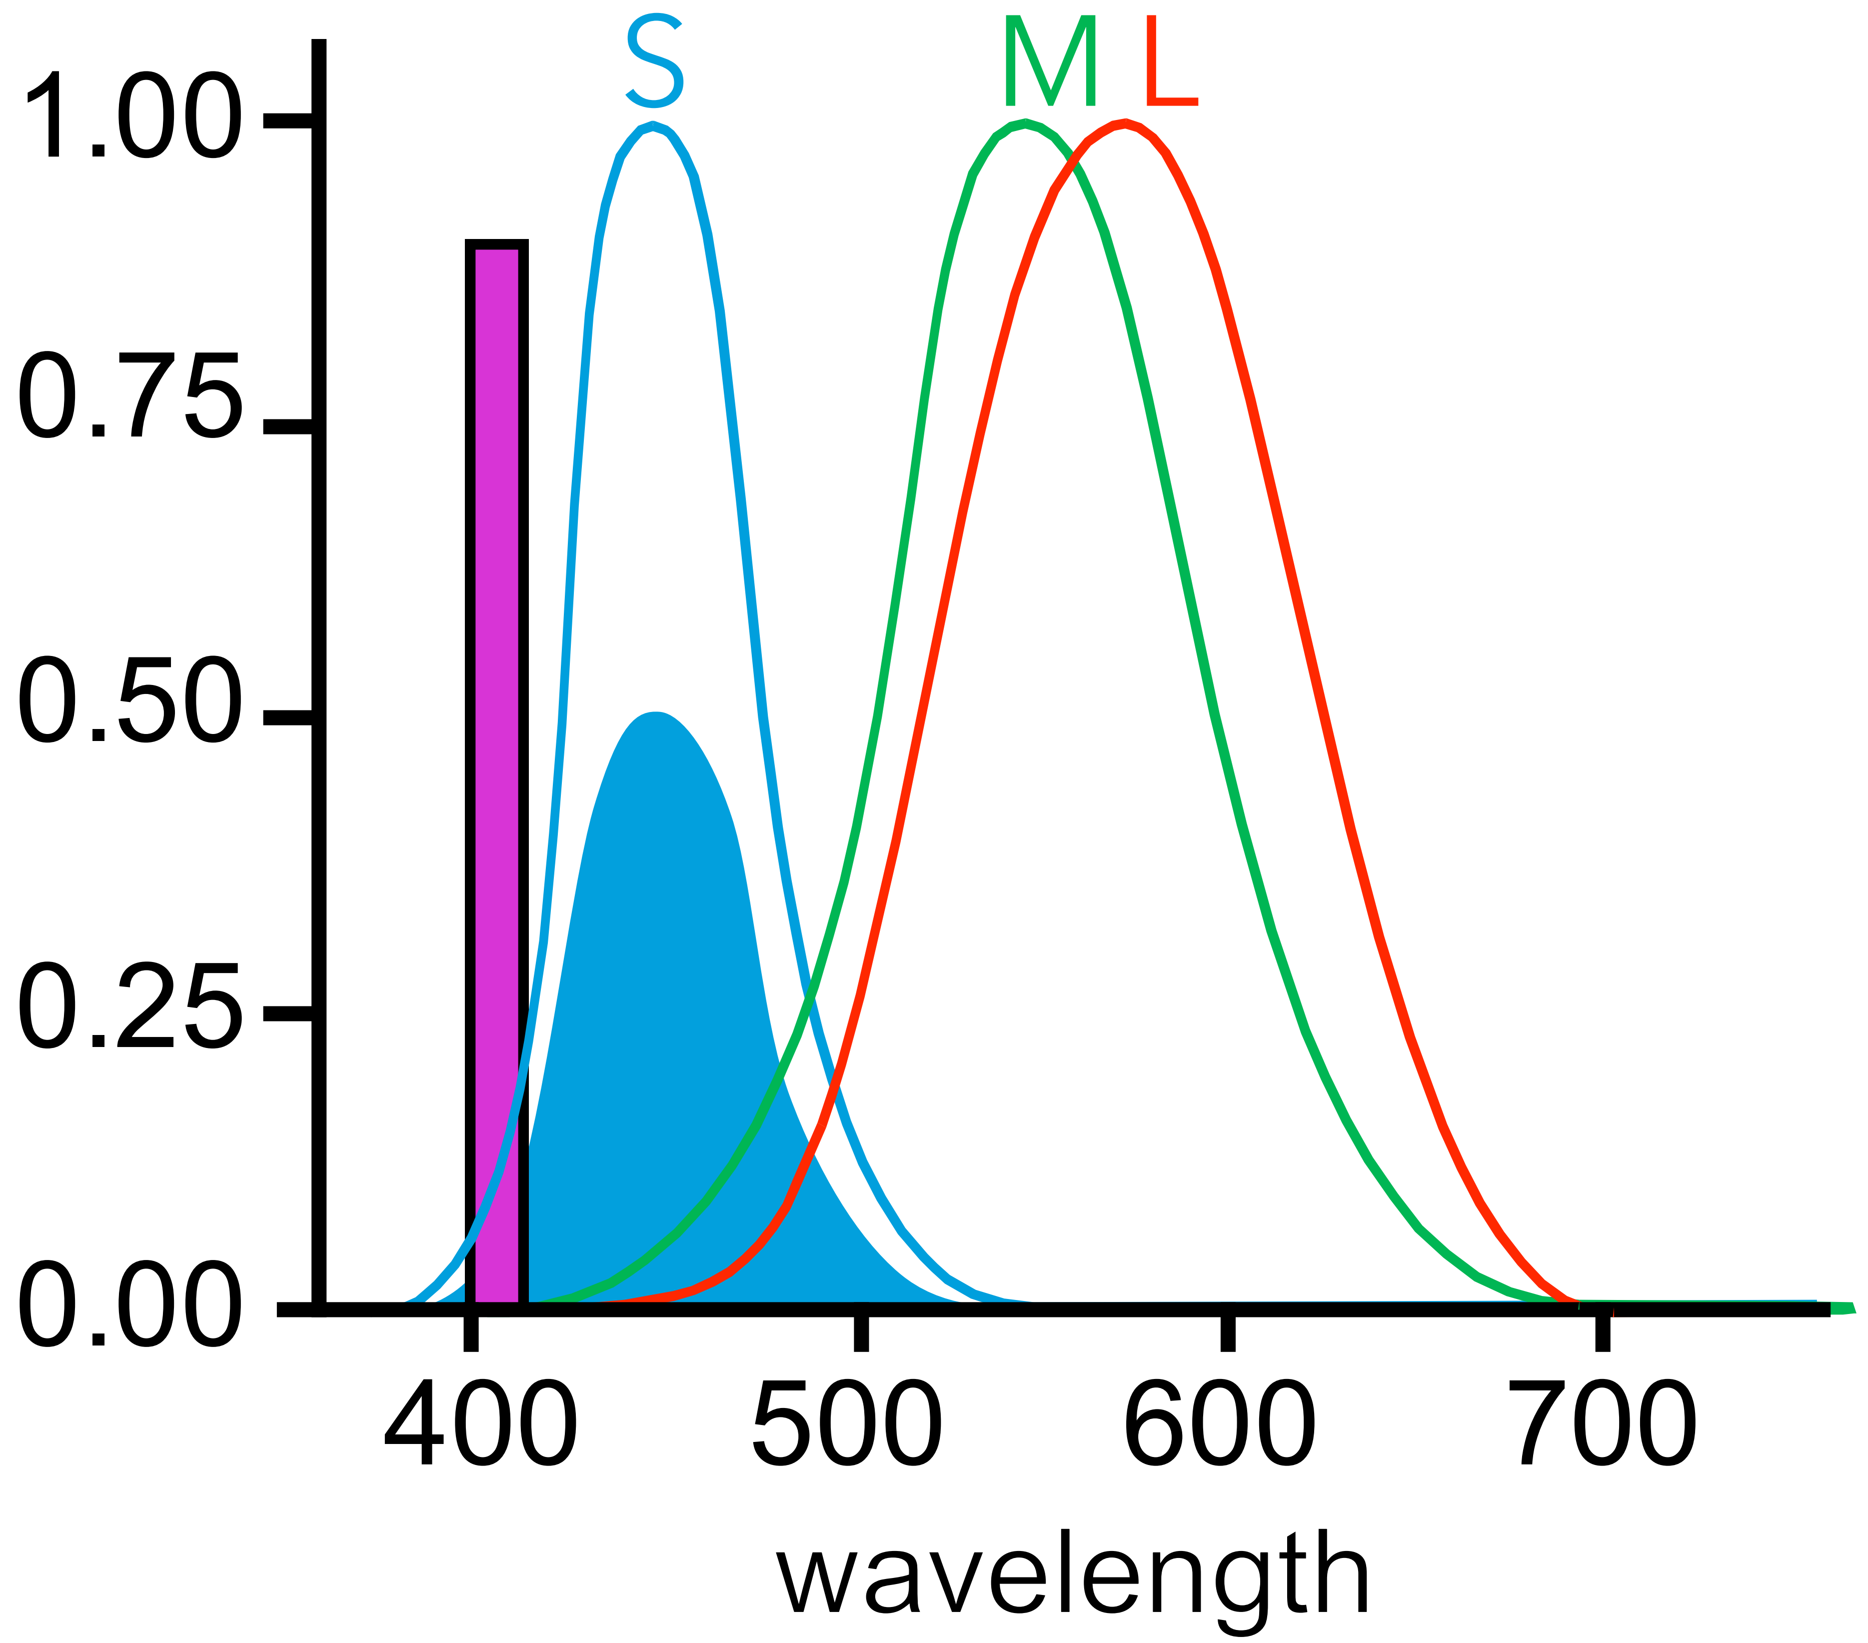

What’s going on?

The three cone responses are not orthogonal

i.e. they overlap and “pollute” each other

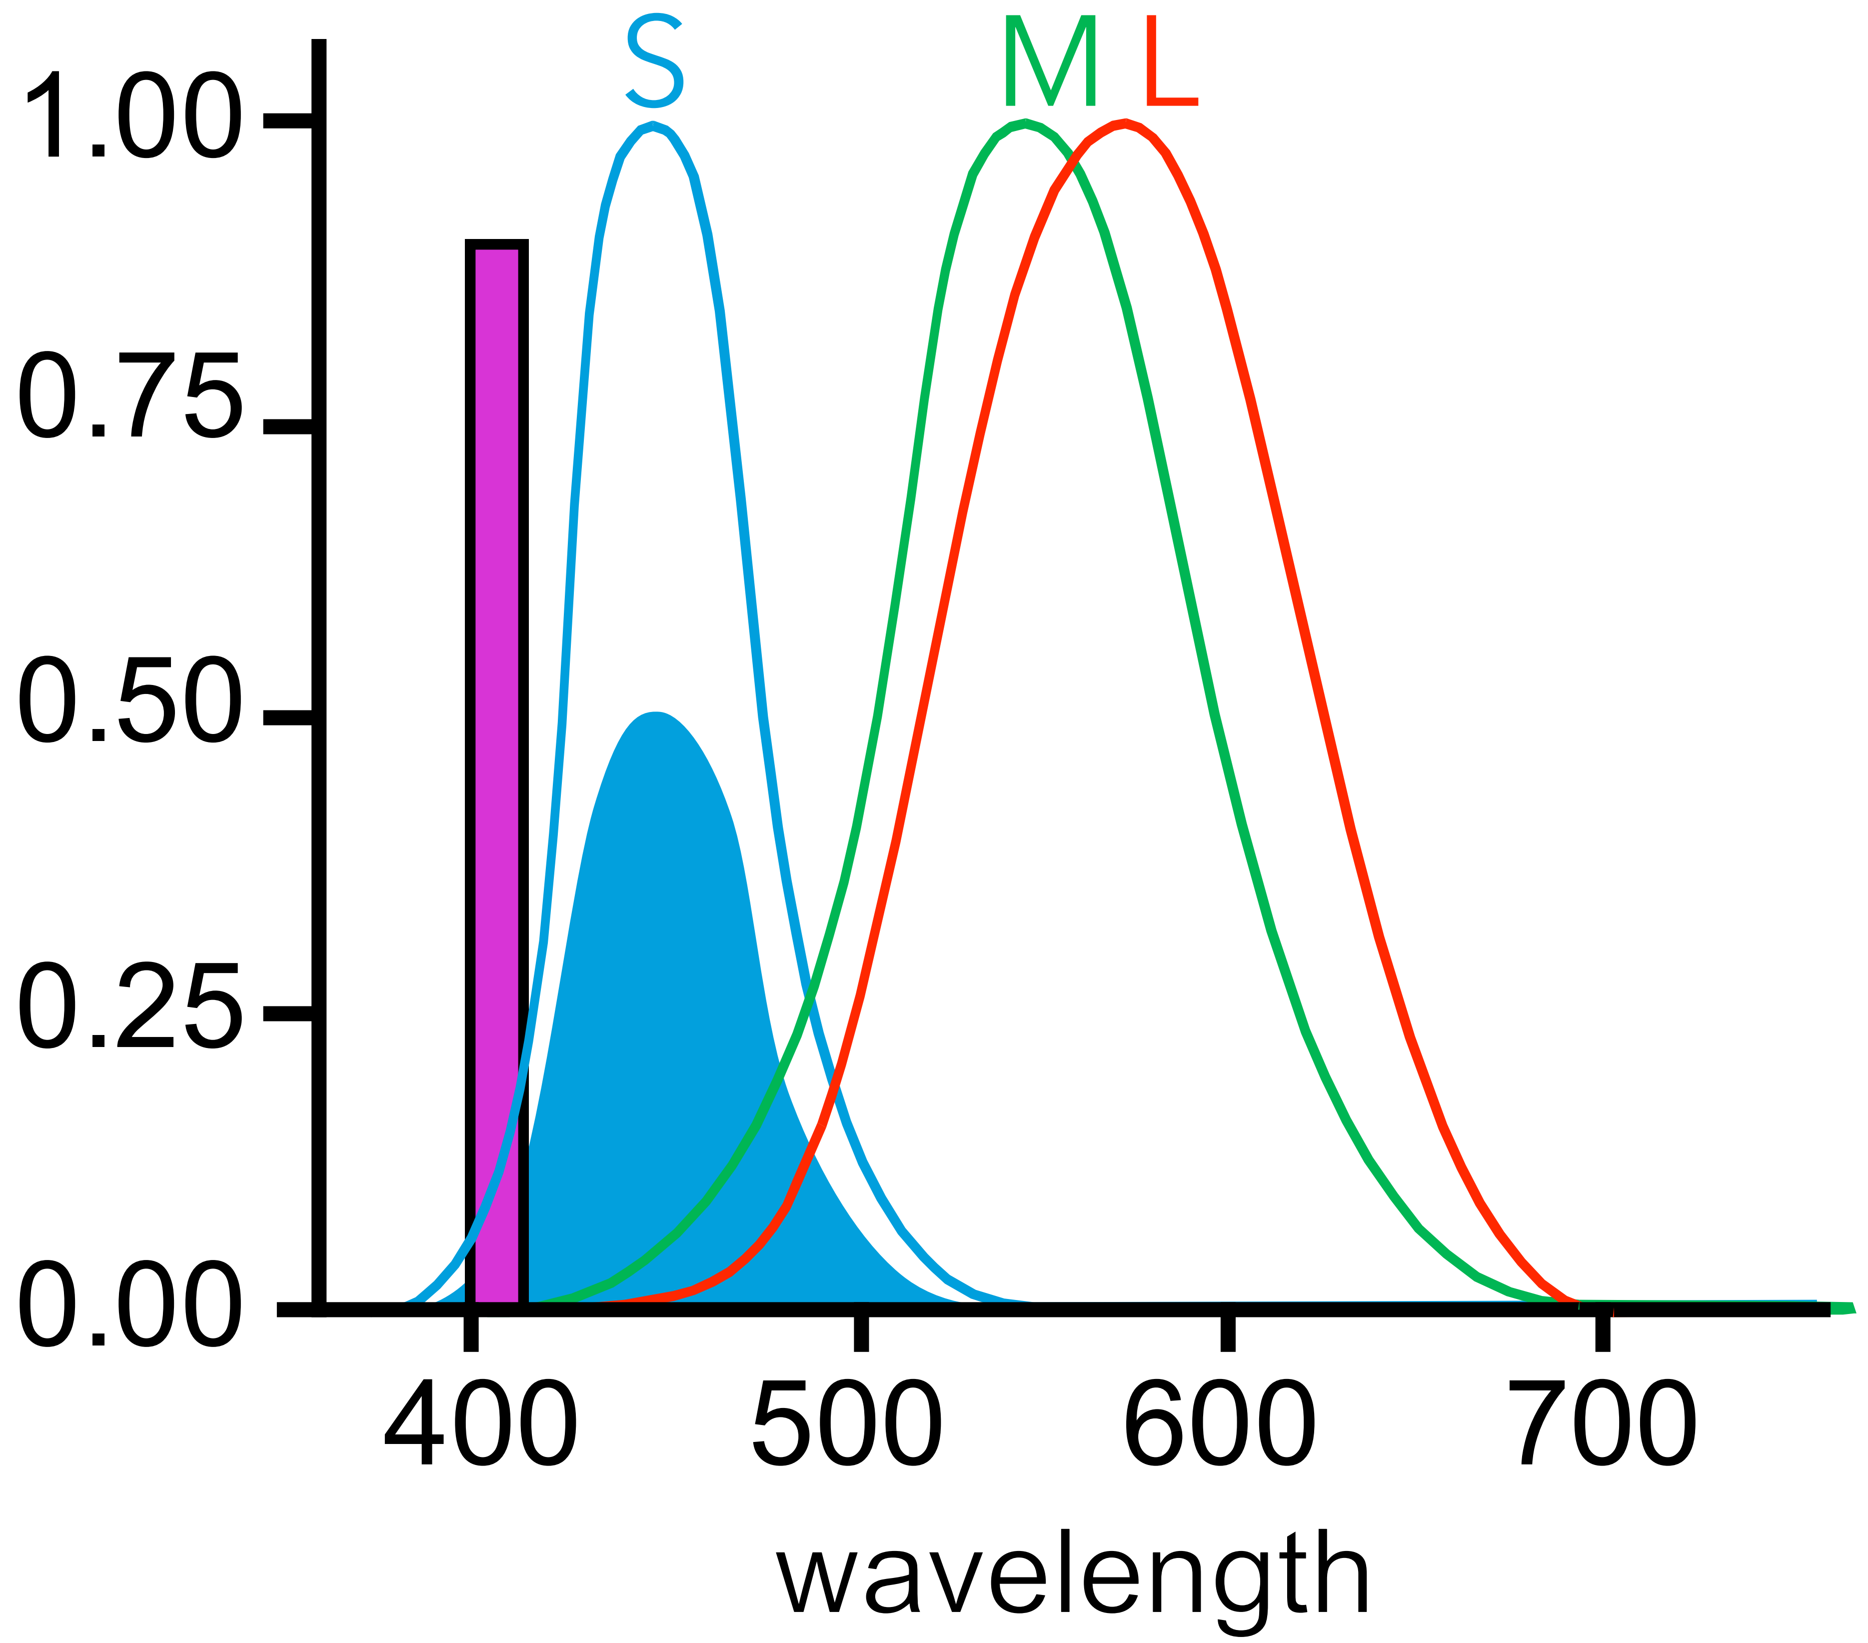

Same as non-orthogonal bases

Non-orthogonal bases are harder to handle

Can't use dot product on same vector to infer coordinates

- Same problem as with cones, the i & j components pollute each other

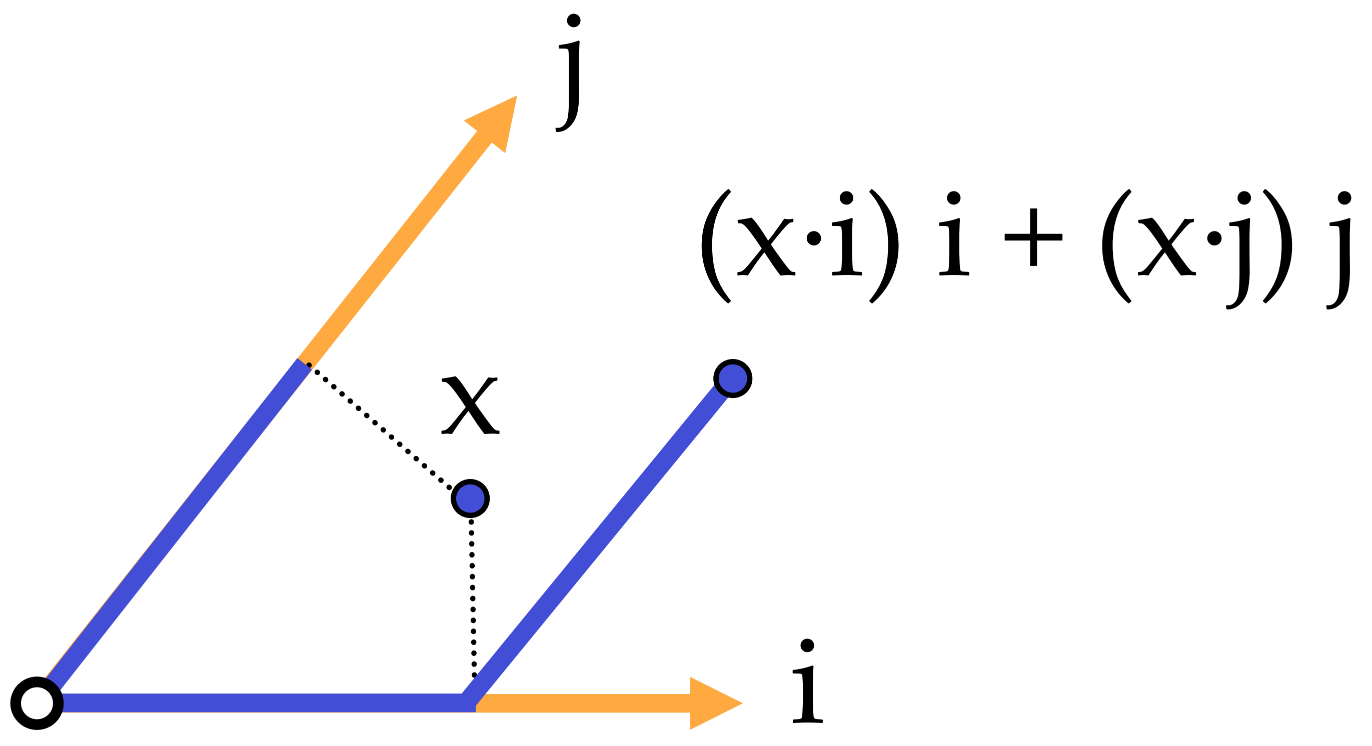

$x \neq (x \cdot i)\, i + (x \cdot j)\, j$

Same as non-orthogonal bases

Non-orthogonal bases are harder to handle

Can't use dot product on same vector to infer coordinates

Need a so-called dual basis

- Same for color: different basis for analysis/synthesis

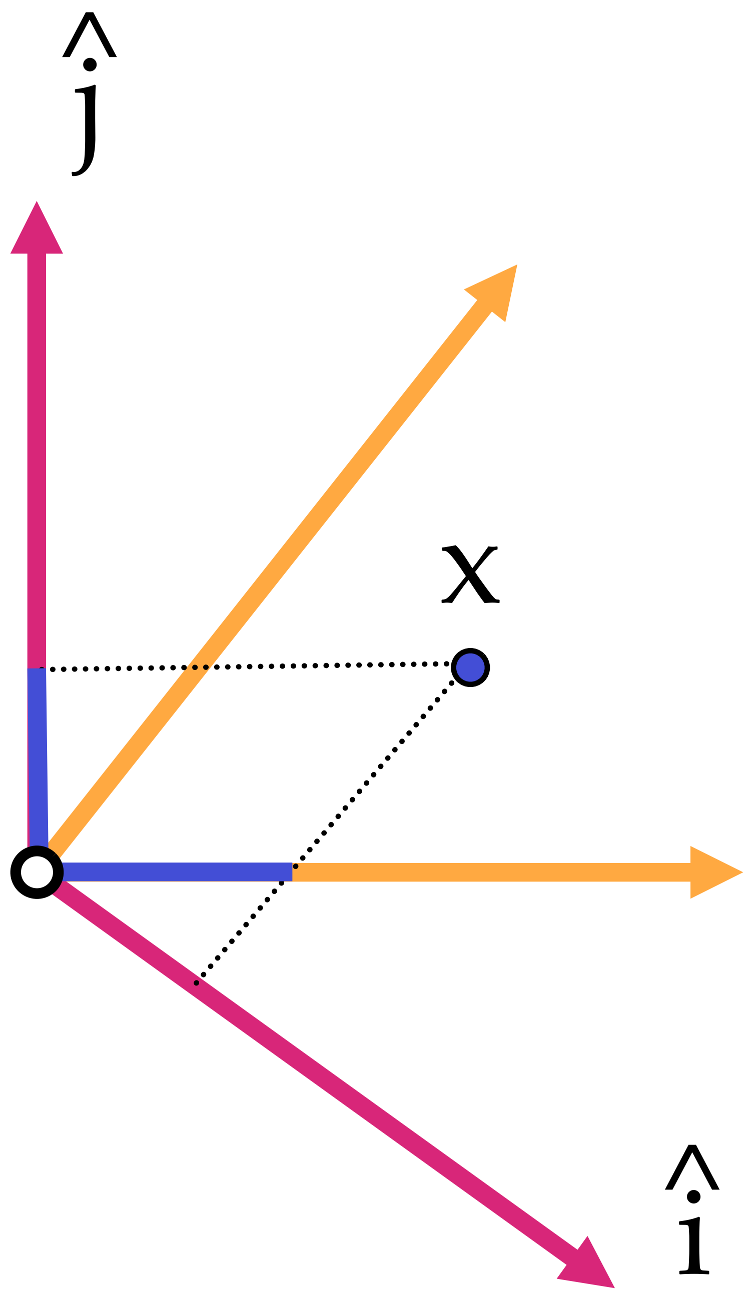

$(x \cdot \hat{i})\, i + (x \cdot \hat{j})\, j$

Note that $\hat{i}$ has negative coordinates

Warning: tricky thing with color

Spectrum for the stimulus / synthesis

- Light, monitor, reflectance

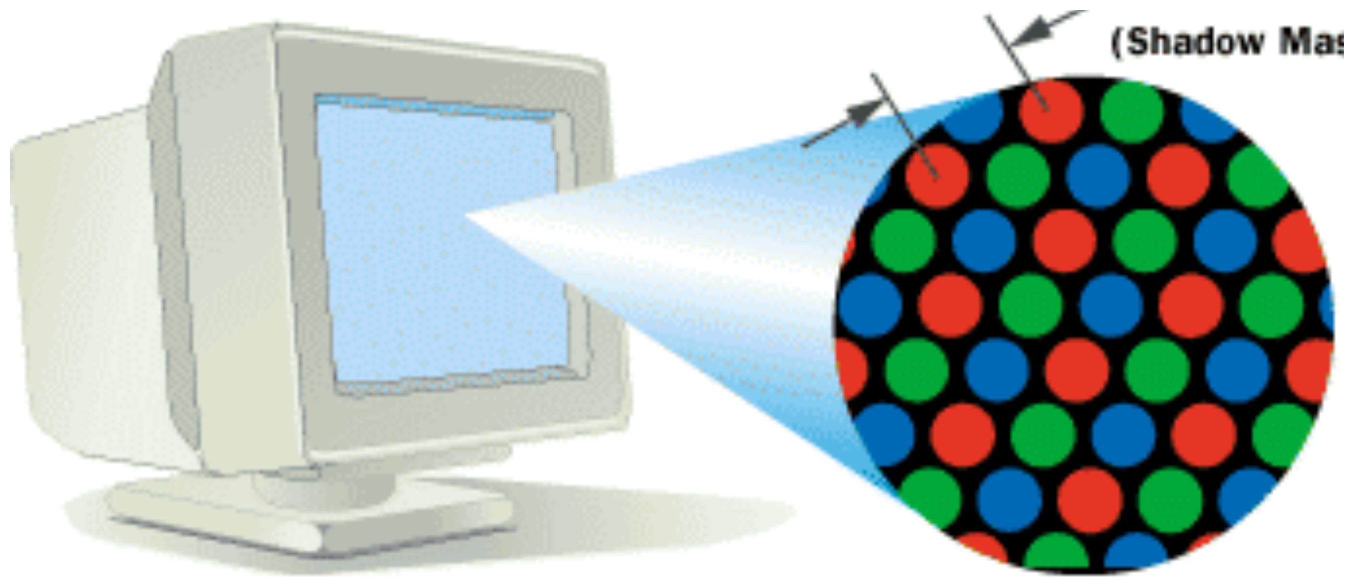

Response curve for receptor / analysis

- Cones, camera, scanner

Usually not the same

Because cone responses are not orthogon

Color reproduction (the right way)

Have a spectrum s; want to match on RGB monitor

- "match" means it looks the same
- any spectrum that projects to the same point in the visual color space is a good reproduction

So, we want to find a spectrum that the monitor can produce that matches s

- that is, we want to display a metamer of s on the screen

LCD display primaries

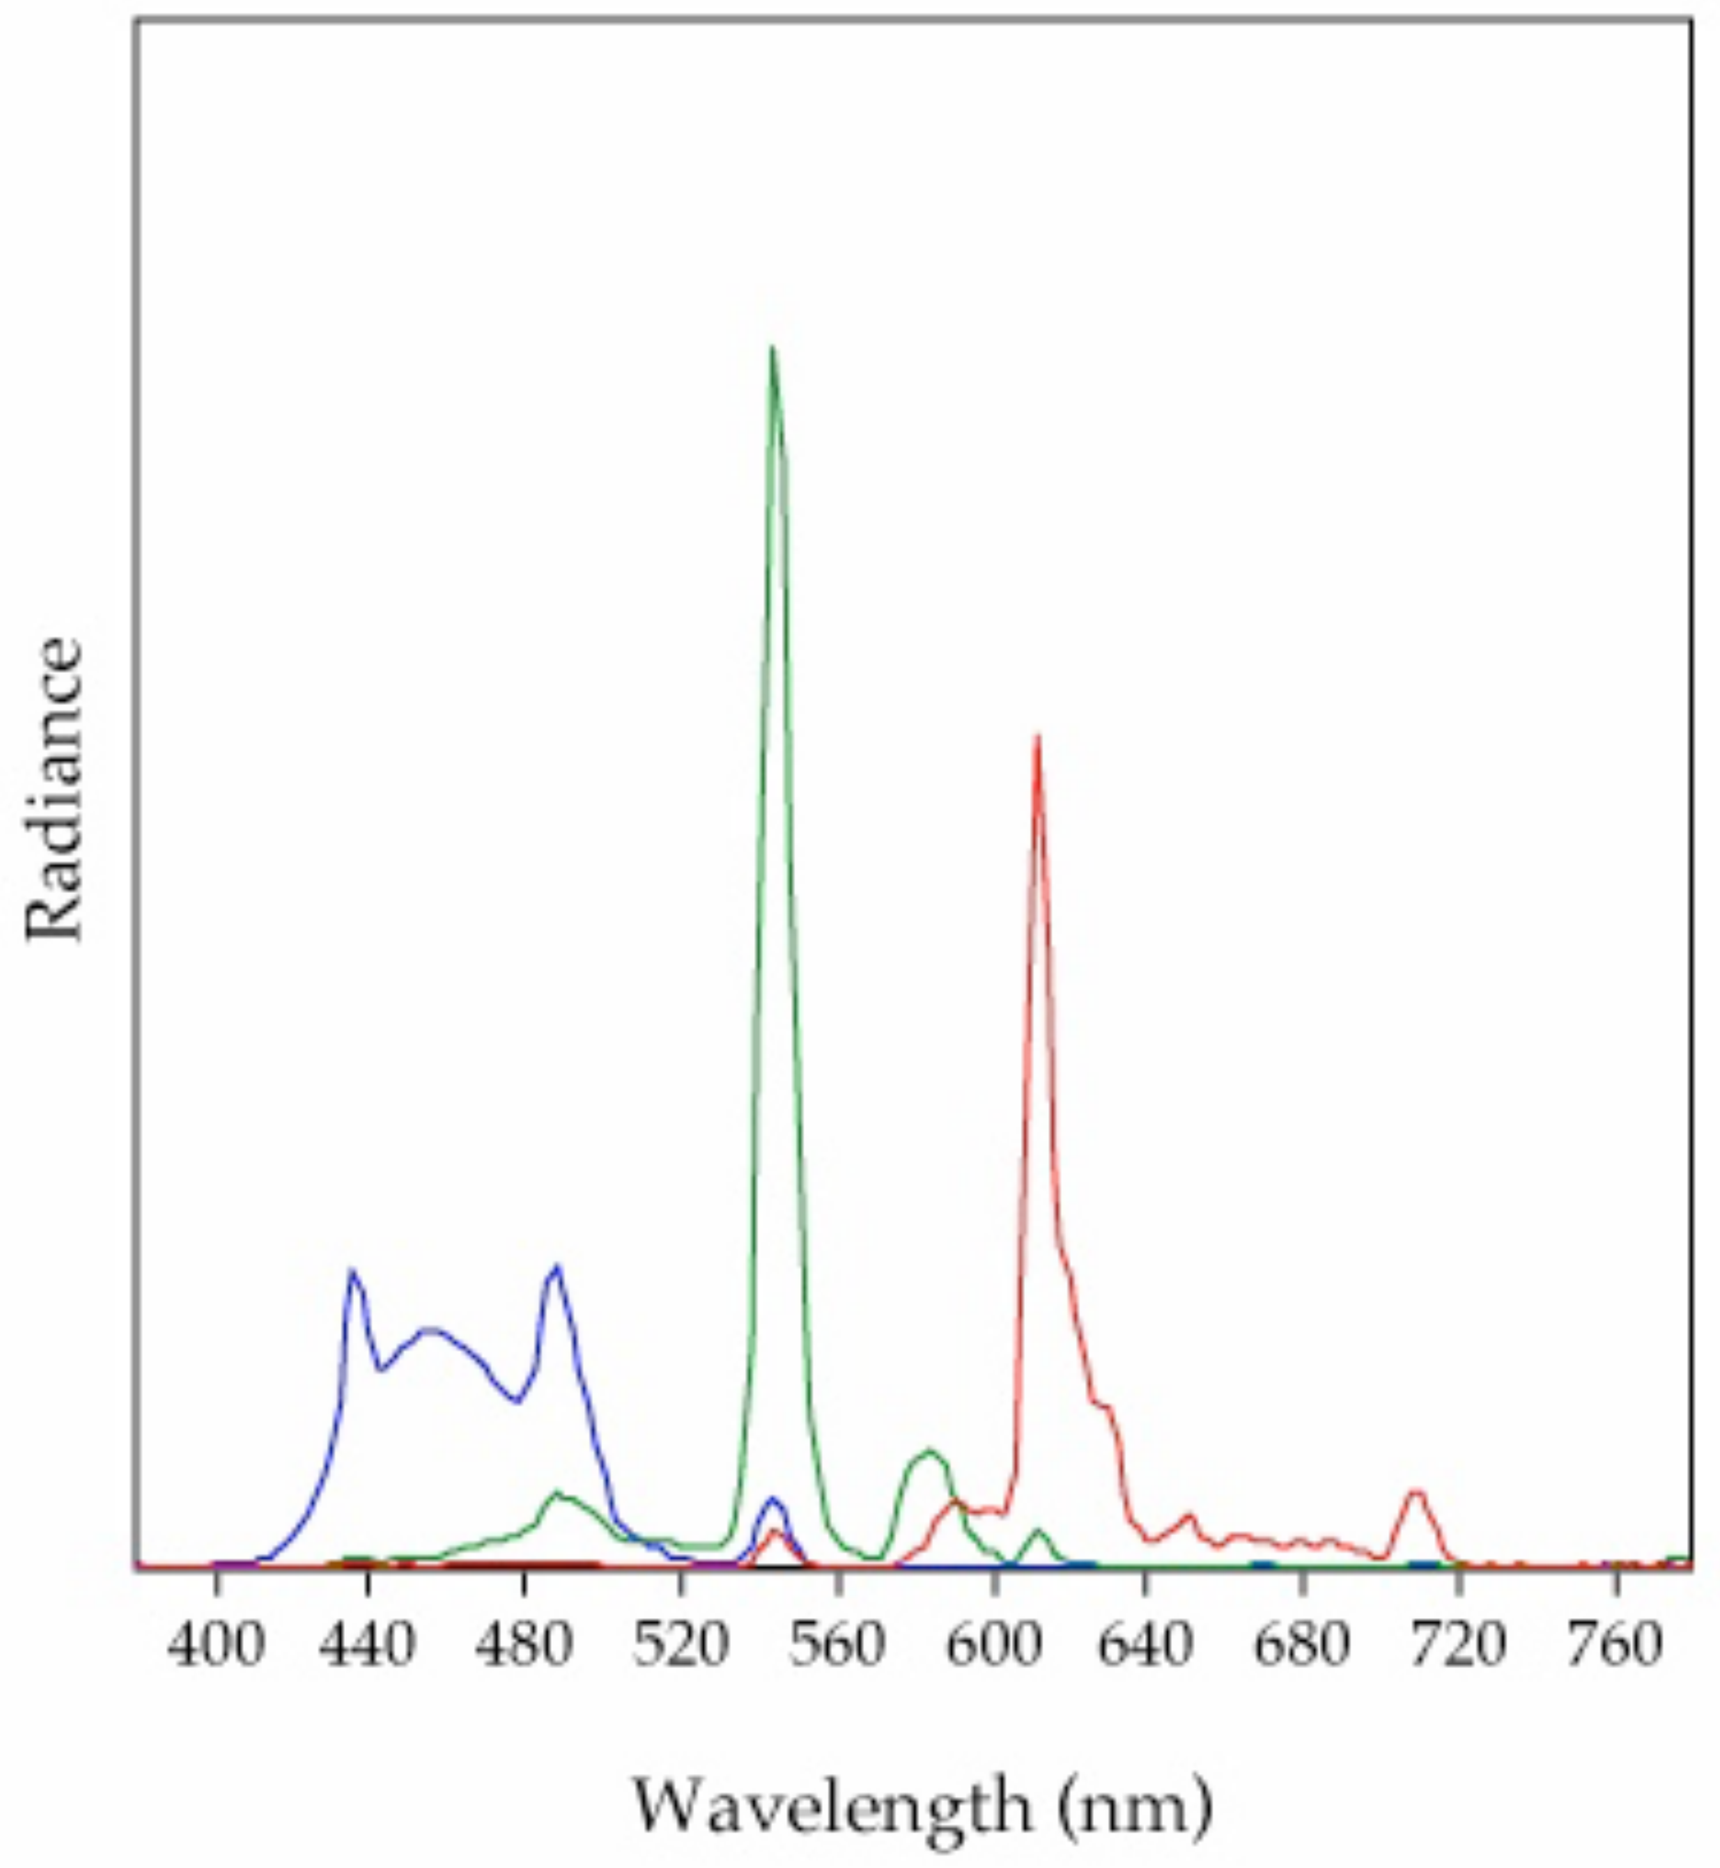

[Fairchild 97]

Curves determined by (fluorescent or LED) backlight and filters

Additive color

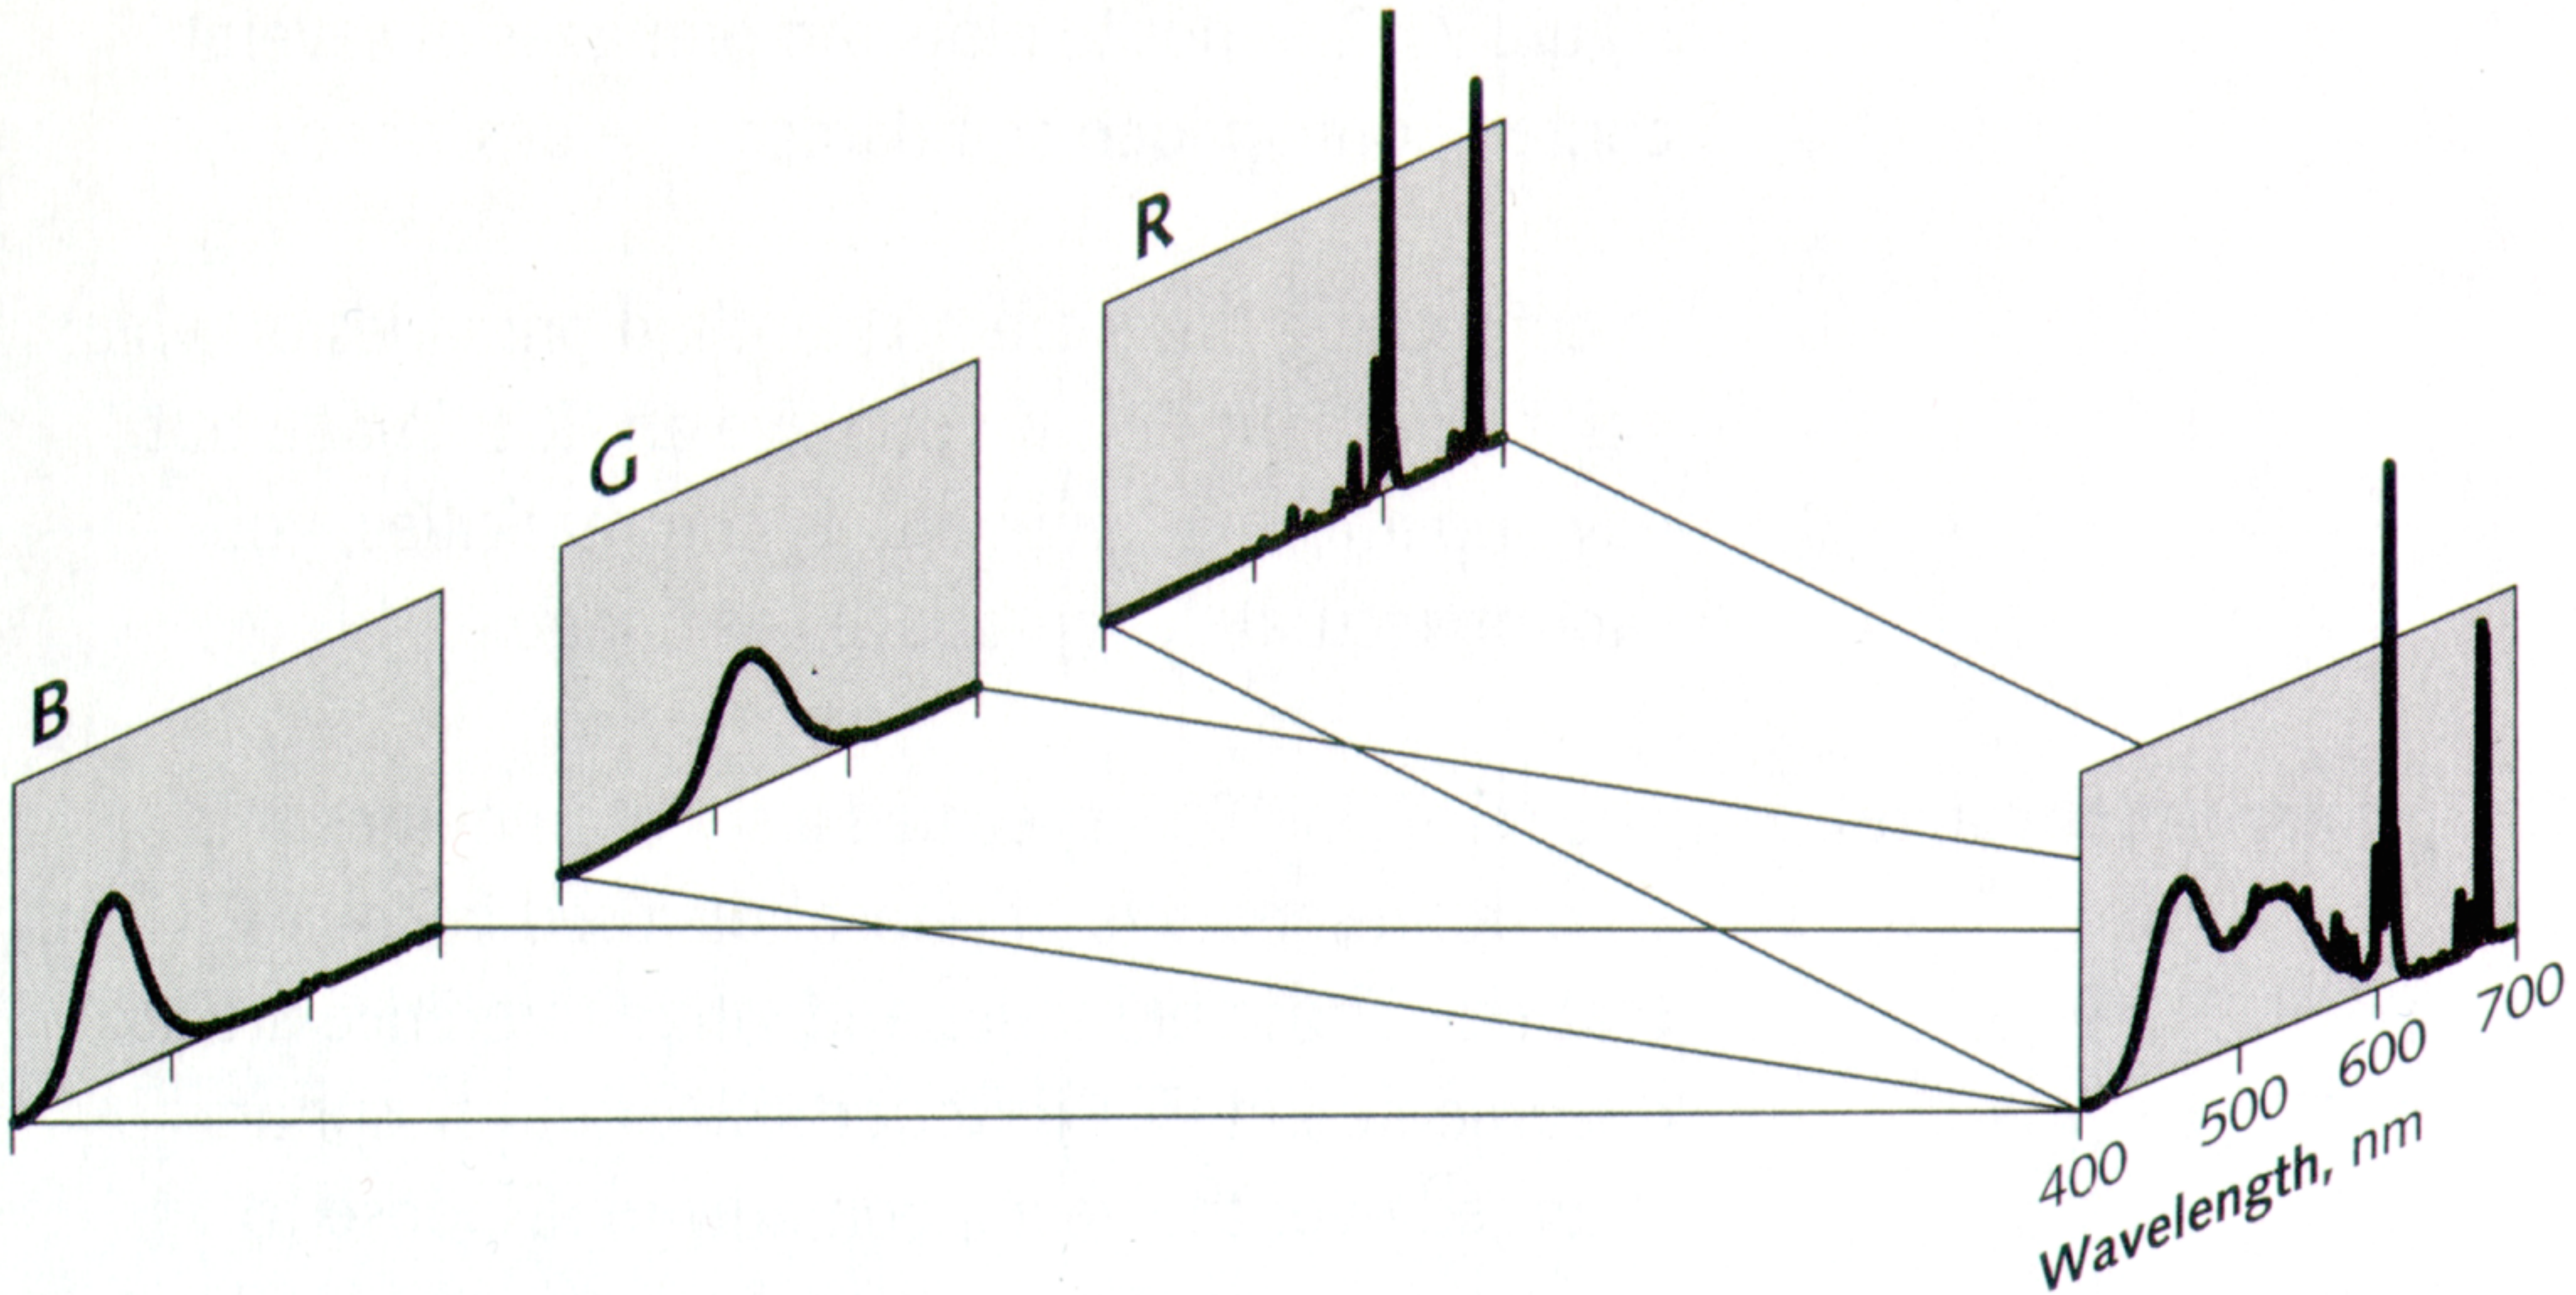

Color reproduction (the right way)

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

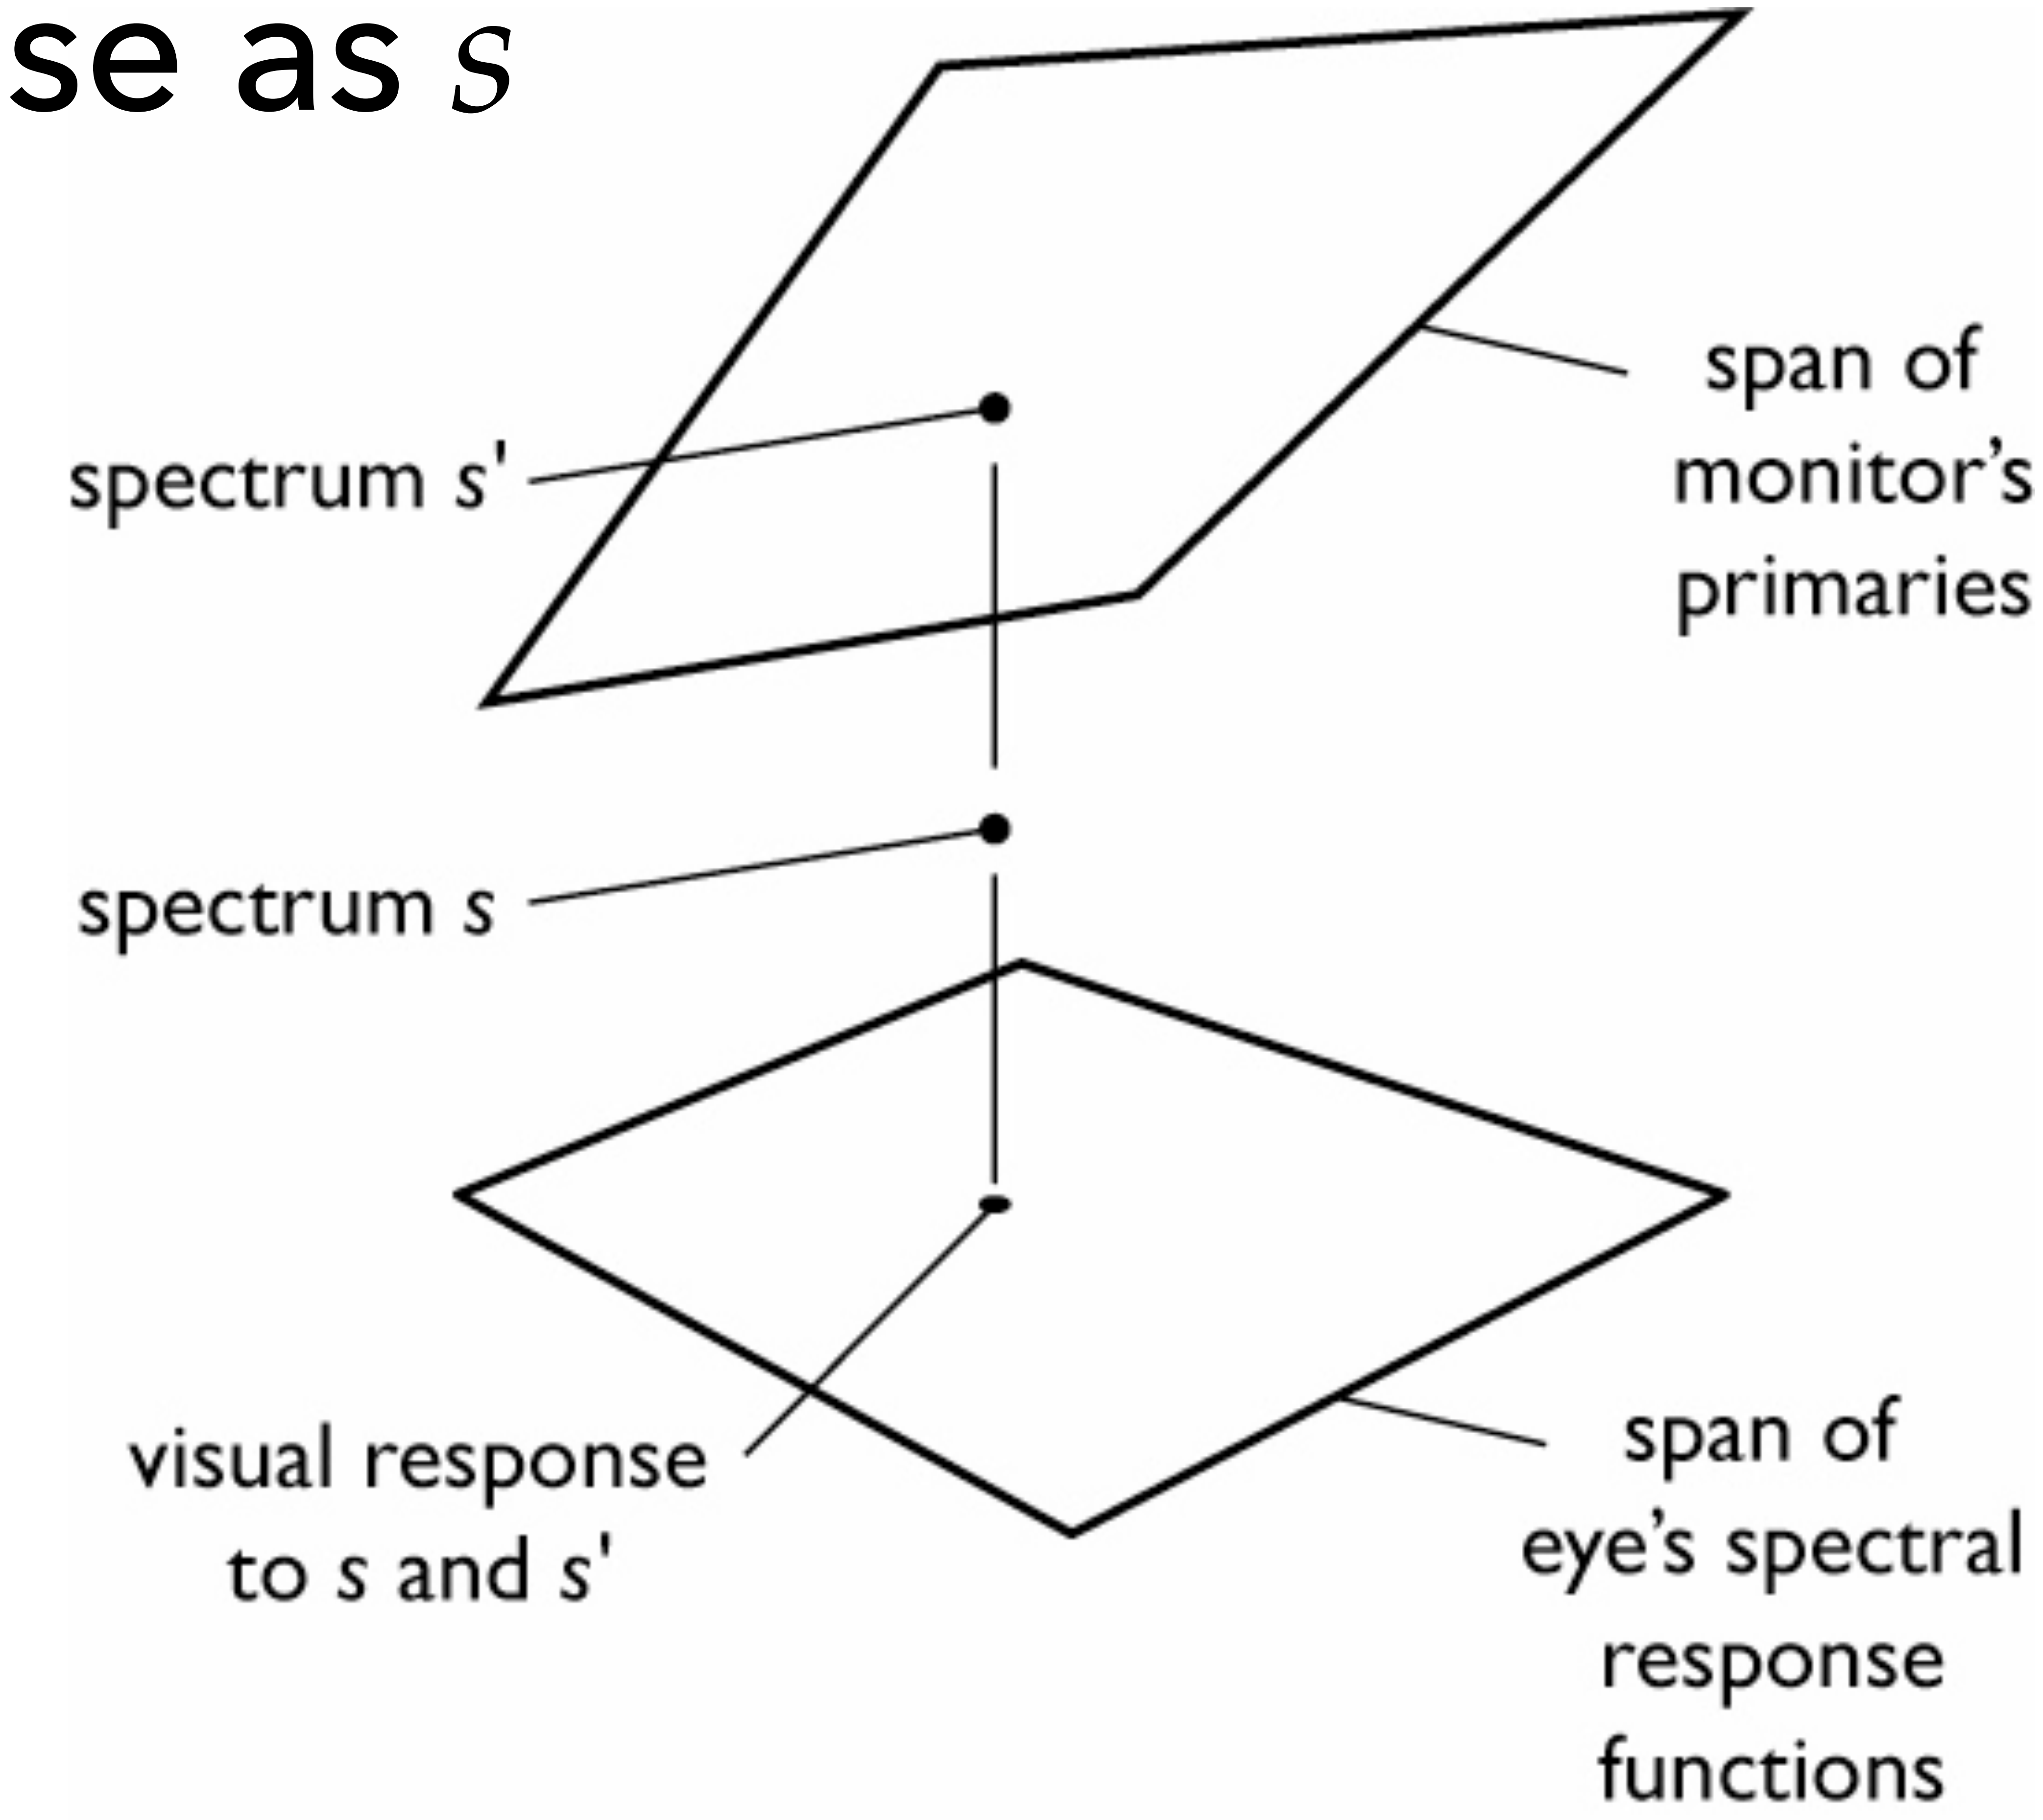

Color reproduction as linear algebra

The projection onto the three response functions can be written in matrix form:

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

or,

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

Color reproduction as linear algebra

The spectrum that is produced by the monitor for the color signals *R*, *G*, and *B* is:

$$s_a(\lambda) = R\, s_R(\lambda) + G\, s_G(\lambda) + B\, s_B(\lambda)$$

Again, the discrete form can be written as a matrix:

$$\begin{bmatrix} | \\ s_a \\ | \end{bmatrix} = \begin{bmatrix} | & | & | \\ s_R & s_G & s_B \\ | & | & | \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix}$$

or,

$$s_a = M_{RGB}\, C$$

Color reproduction as linear algebra

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

- Feed C to display

$$C$$

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,

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

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

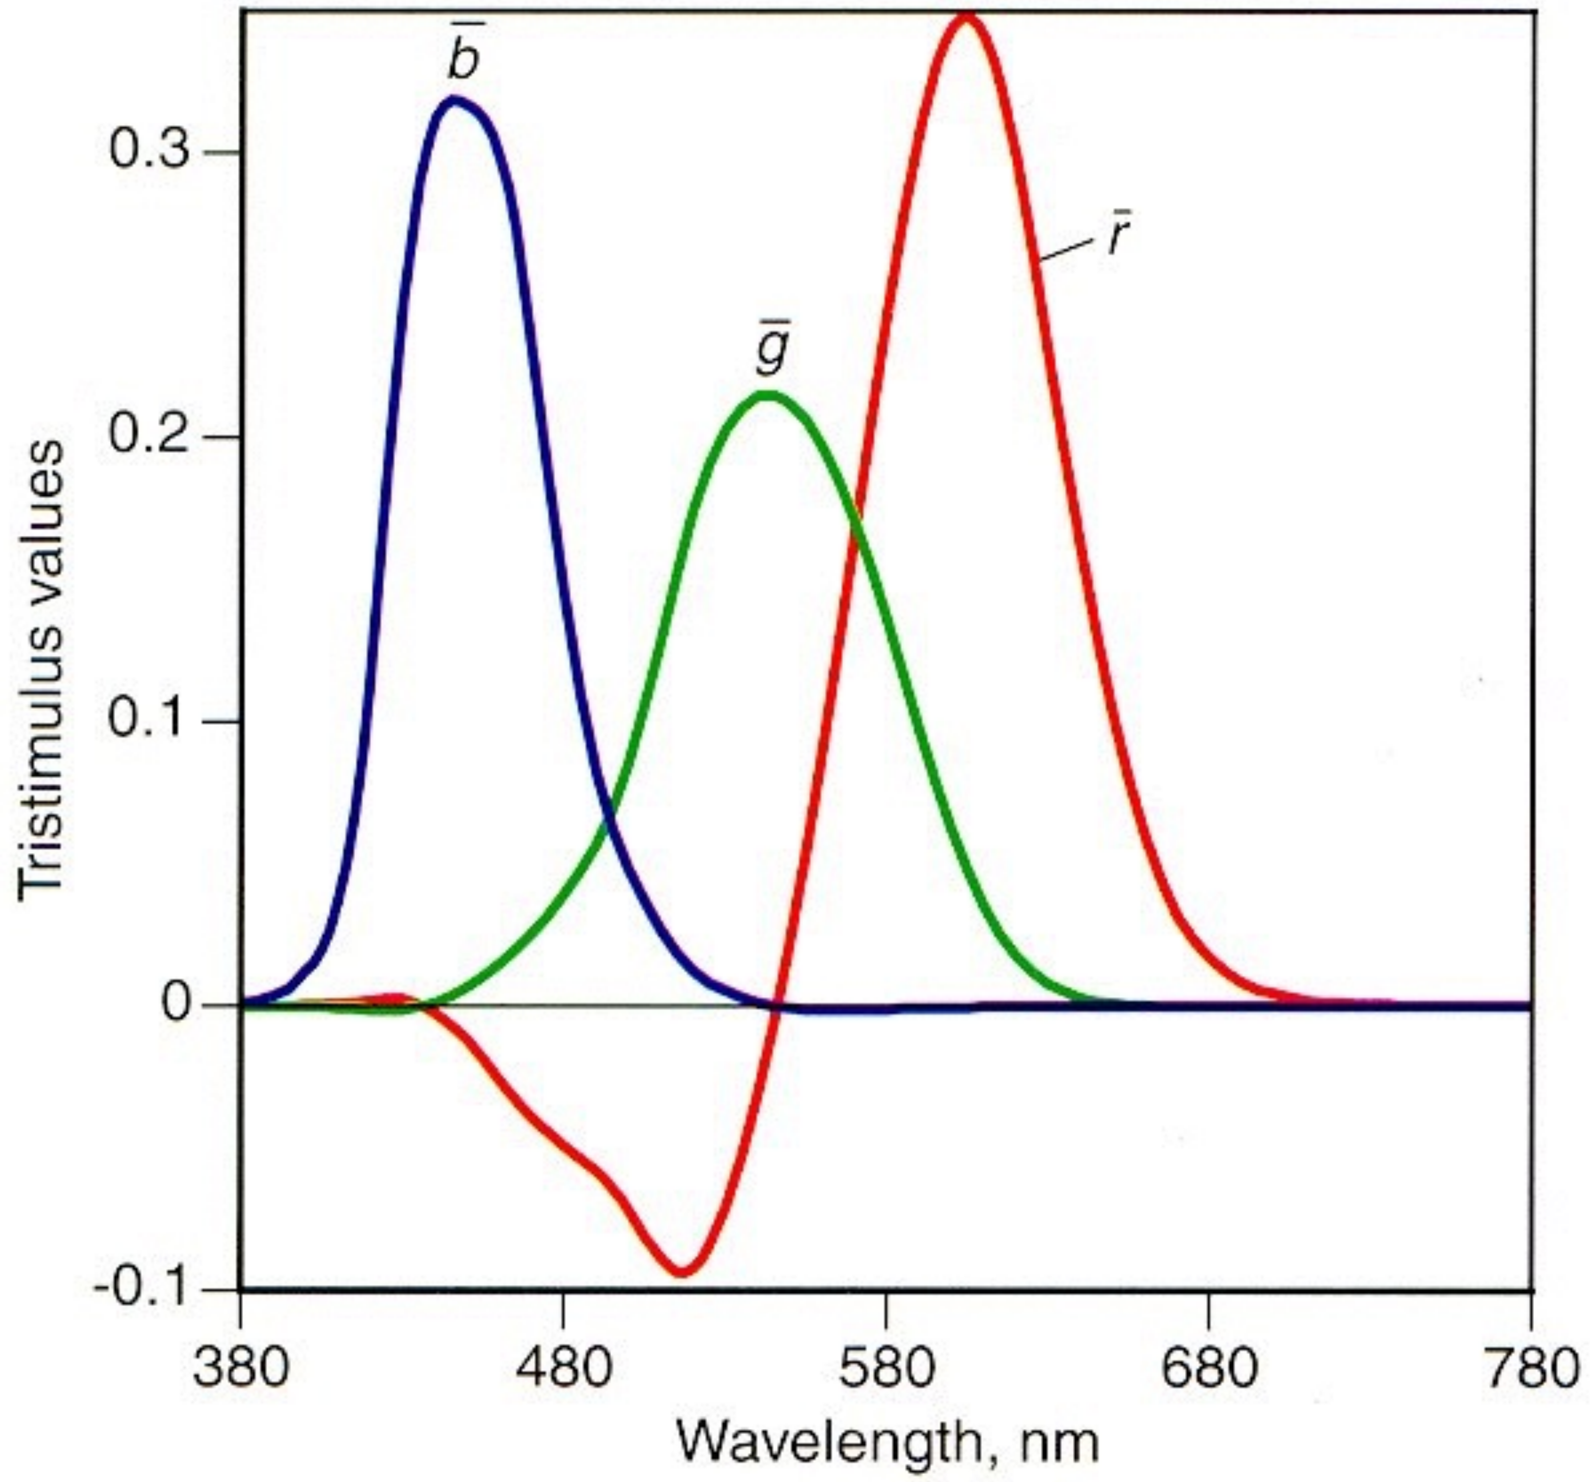

These curves are the color-matching functions for the 1931 standard observer, The average results of 17 color-normal observers having matched each wavelength of the equal-energy spectrum with primaries of 435.8 nm, 546.1 nm, and 700 nm.

Color reproduction recap

We now know how to match any color from the real world on a display

We don't need to know the whole spectrum, only the projections onto S, M, and L response functions

There is then a simple linear procedure to work out the combination of any 3 primaries to match the color

But there is a catch. More on that later.

Summary

Physical color

- Spectrum
- multiplication of light & reflectance spectrum

Perceptual color

- Cone spectral response: 3 numbers
- Metamers: different spectrum, same responses
    - Color matching, enables color reproduction with 3 primaries

Fundamental difficulty

- Spectra are infinite-dimensional (full function)
- Projected to only 3 types of cones
- Cone responses overlap / they are non-orthogonal
    - Means different primaries for analysis and synthesis
- Negative numbers are not physical

Today

Light & Color

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

Color spaces

How can we quantitatively represent, reproduce color?

Brute force: store, reproduce full spectral energy distribution

- Disadvantages?

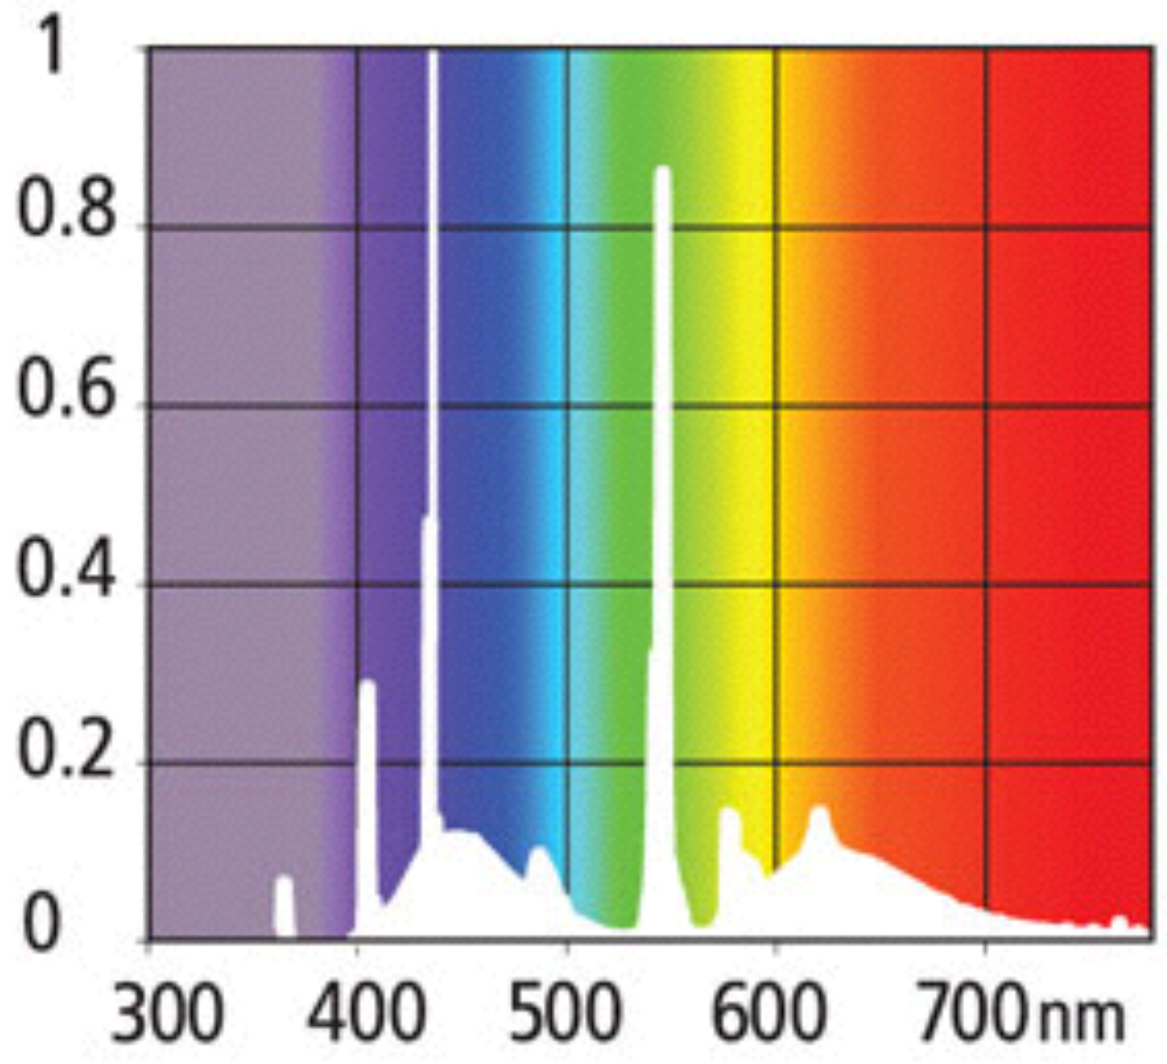

Color spaces

Representation should be complete, but as compact as possible

- Any pair of colors that can be distinguished by humans should have two different representations
- Any pair of colors that appears the same to humans should have the same representation

Standard color spaces

We need a principled color space

Three types of cones, so expect three parameters to be sufficient

Many possible definitions

- Including cone response (SML)
- Unfortunately not really used (unknown when colorimetry was invented)

Good news: color vision is linear and 3-dimensional, so any new color space based on color matching can be obtained using 3x3 matrix

- but there are also non-linear color spaces (e.g. Hue Saturation Value, Lab)

Overview

Most standard color space: CIE XYZ

SML and the various flavors of RGB are just linear transformations of the XYZ basis

- 3x3 matrices

Why not measure cone sensitivity?

Less directly measurable

- electrode in photoreceptor?
- not available when color spaces were defined

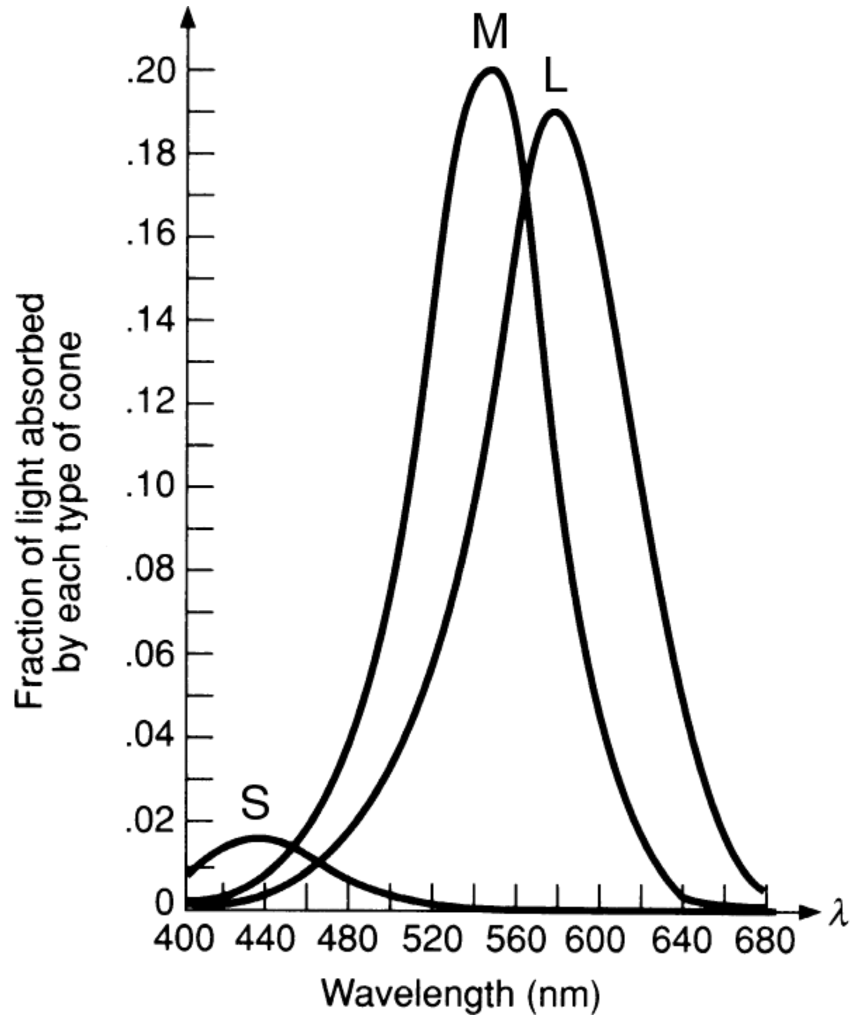

Most directly available measurement:

- notion of metamers & color matching
- directly in terms of color reproduction:
  given an input color, how to reproduce it with 3 primary colors?
- CIE: **C**ommission **I**nternationale de l'**E**clairage
  (International Lighting Commission)
- Circa 1920

CIE color matching experiment

Given an input color, how to reproduce it with 3 primary colors?

(Idea by Maxwell)

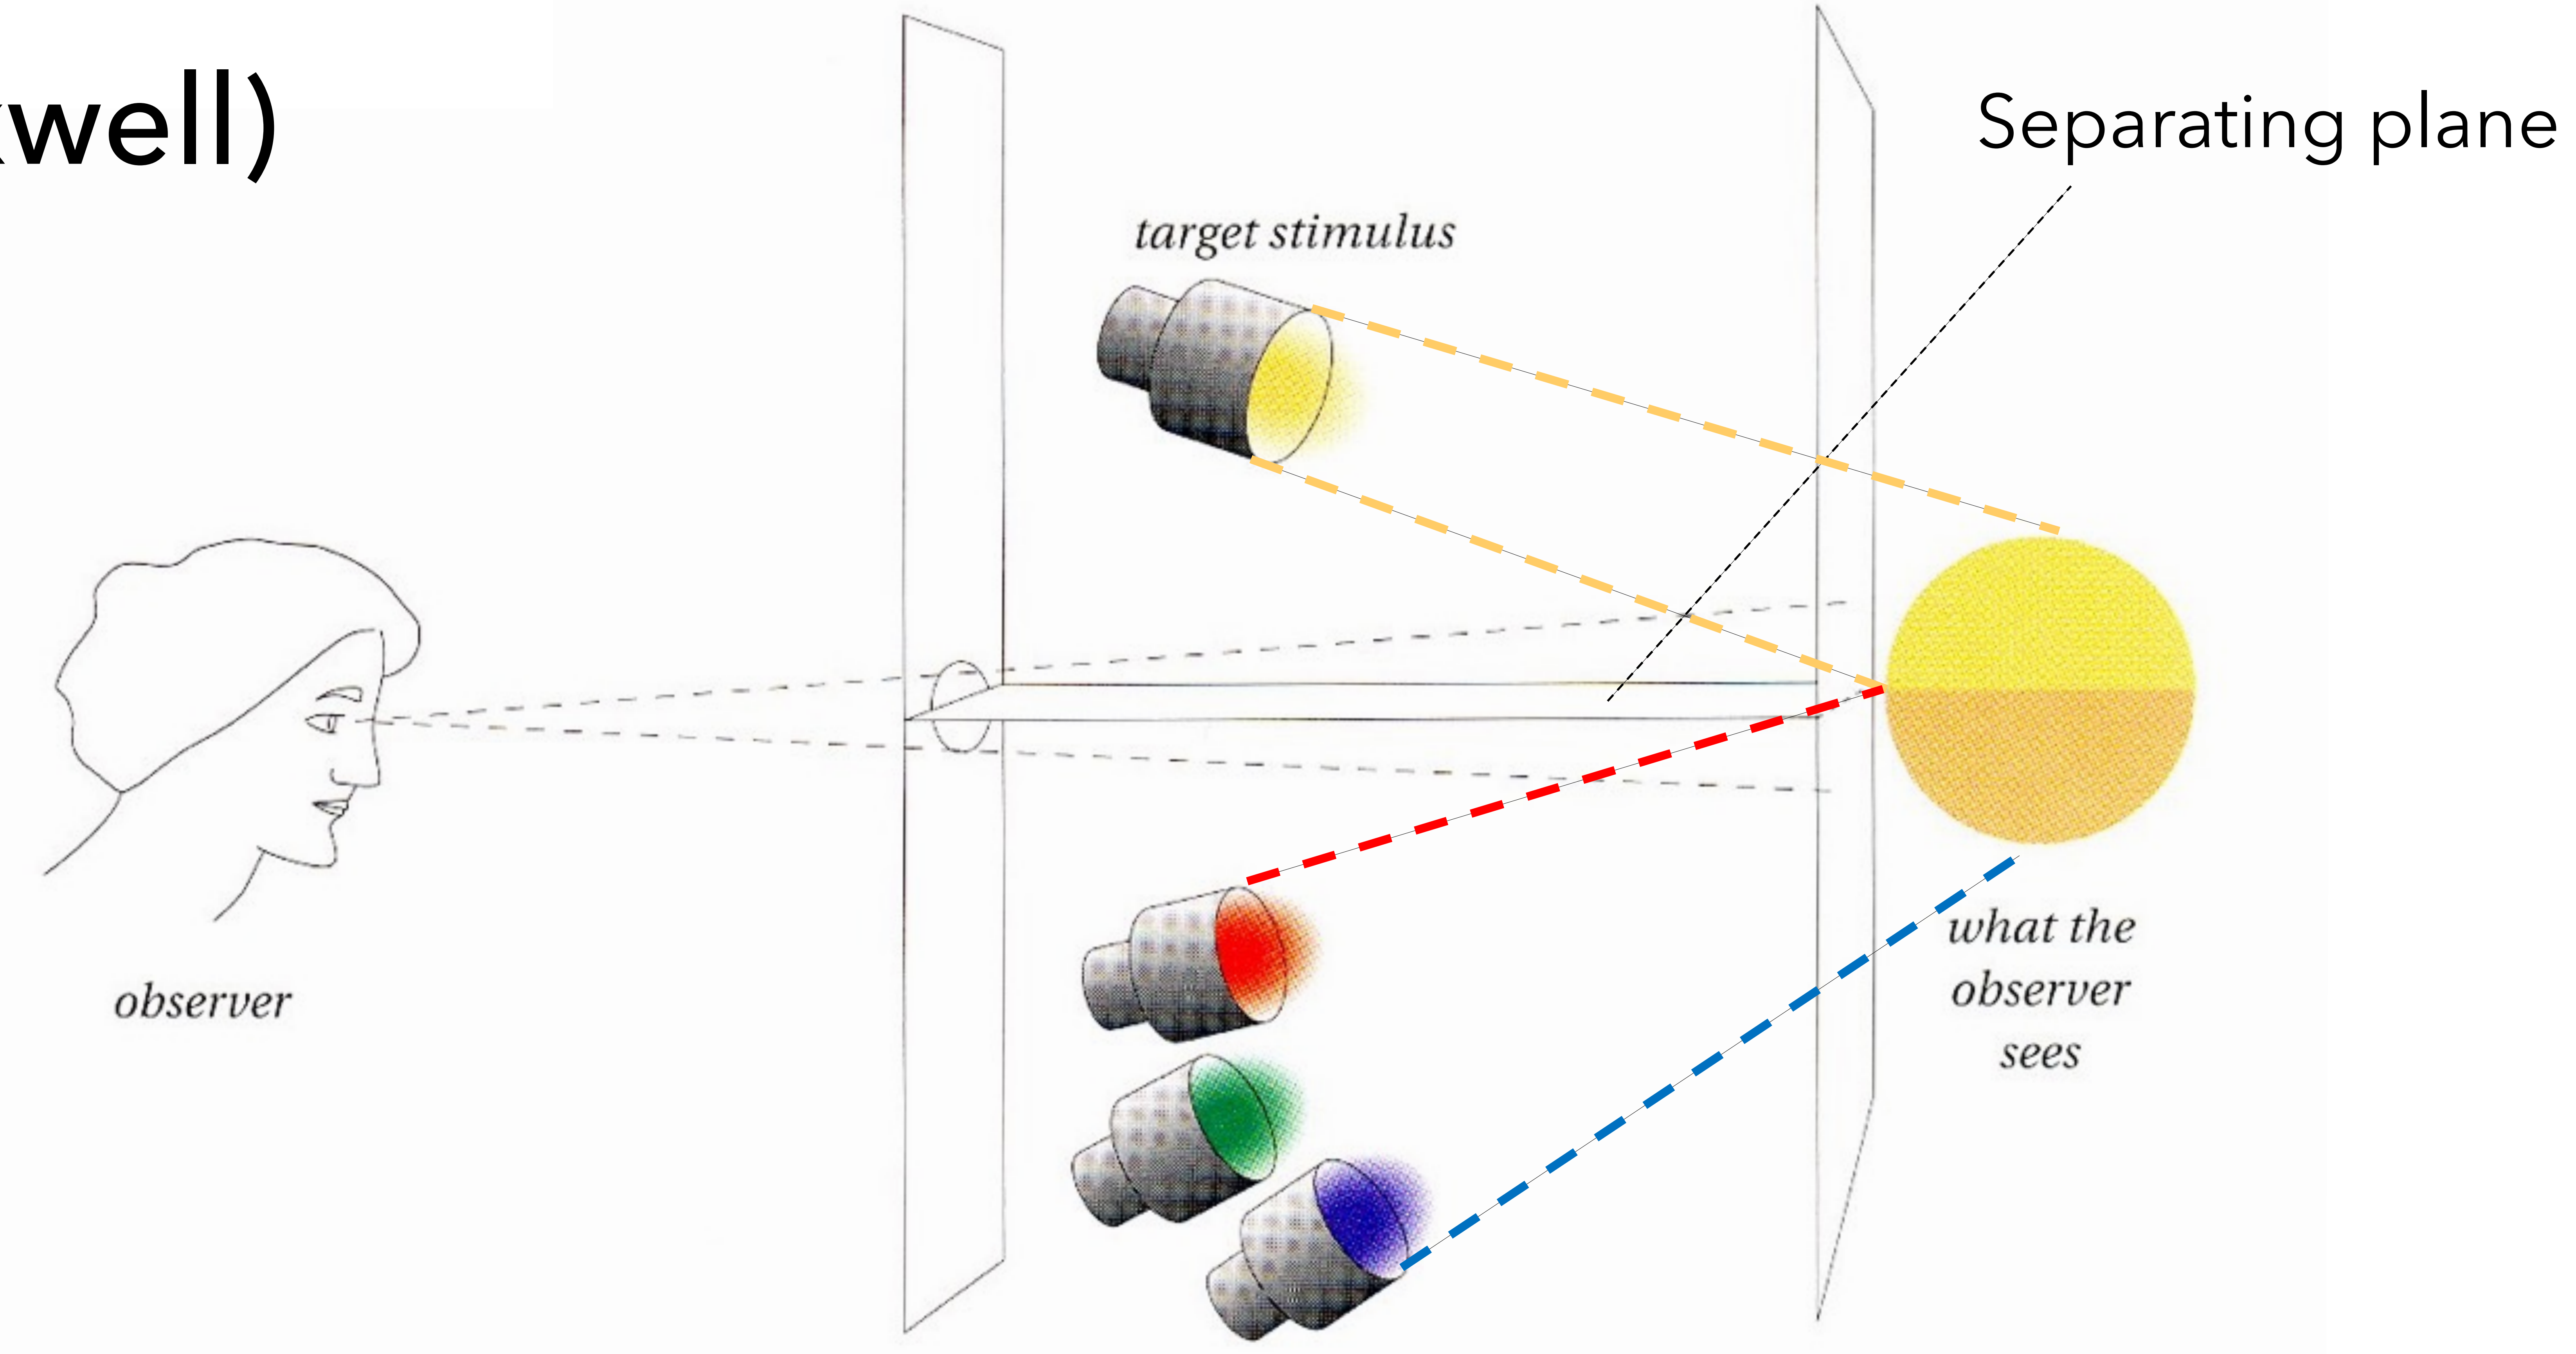

The observer adjusts the intensities of the red, green, and blue lamps until they match the target stimulus on the split screen.

CIE color matching experiment

Primaries (synthesis) at 435.8, 546.1 and 700nm

- Chosen for robust reproduction, good separation in red-green
- Don’t worry, we’ll be able to convert it to any other set of primaries (Linear algebra to the rescue!)

Resulting 3 weights for each primary are called tristimulus values

Applet

http://graphics.stanford.edu/courses/cs178-10/applets/colormatching.html

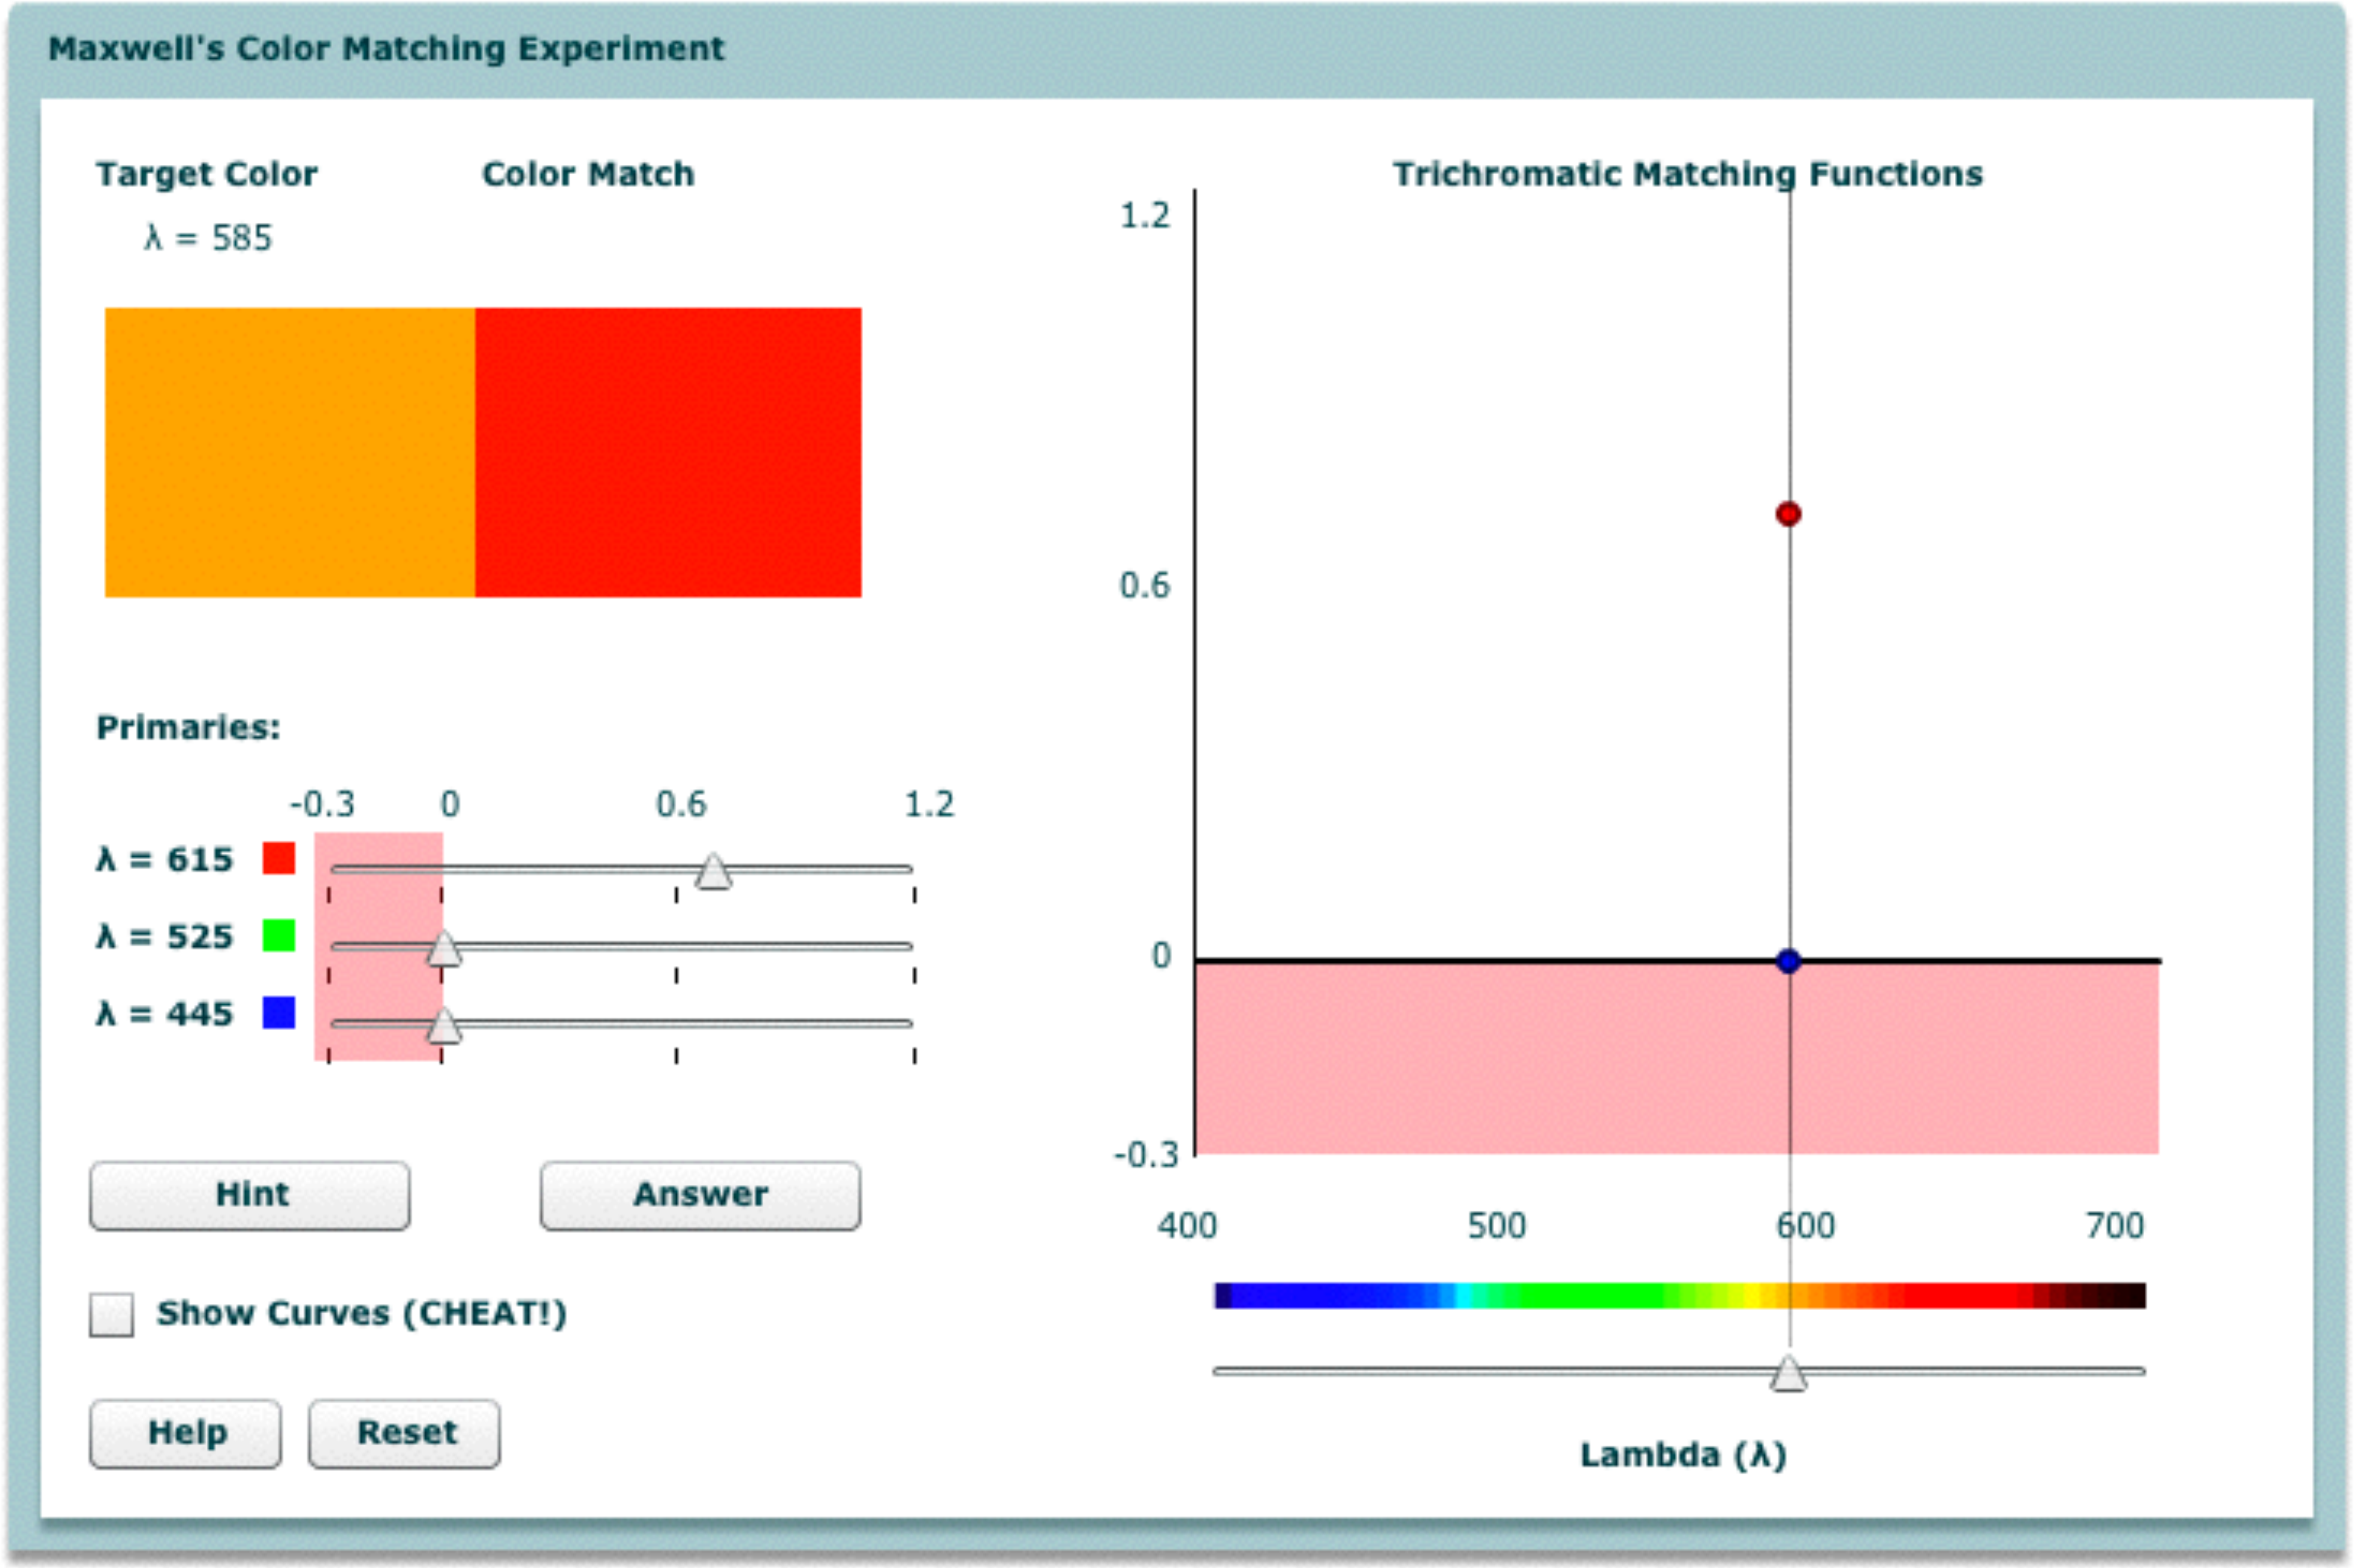

CIE color matching

Meaning of these curves: a monochromatic wavelength λ can be reproduced with:

 b(λ) amount of the 435.8nm primary,

+g(λ) amount of the 546.1 primary,

+r(λ) amount of the 700 nm primary

This fully specifies the color perceived by a human

What is this!?

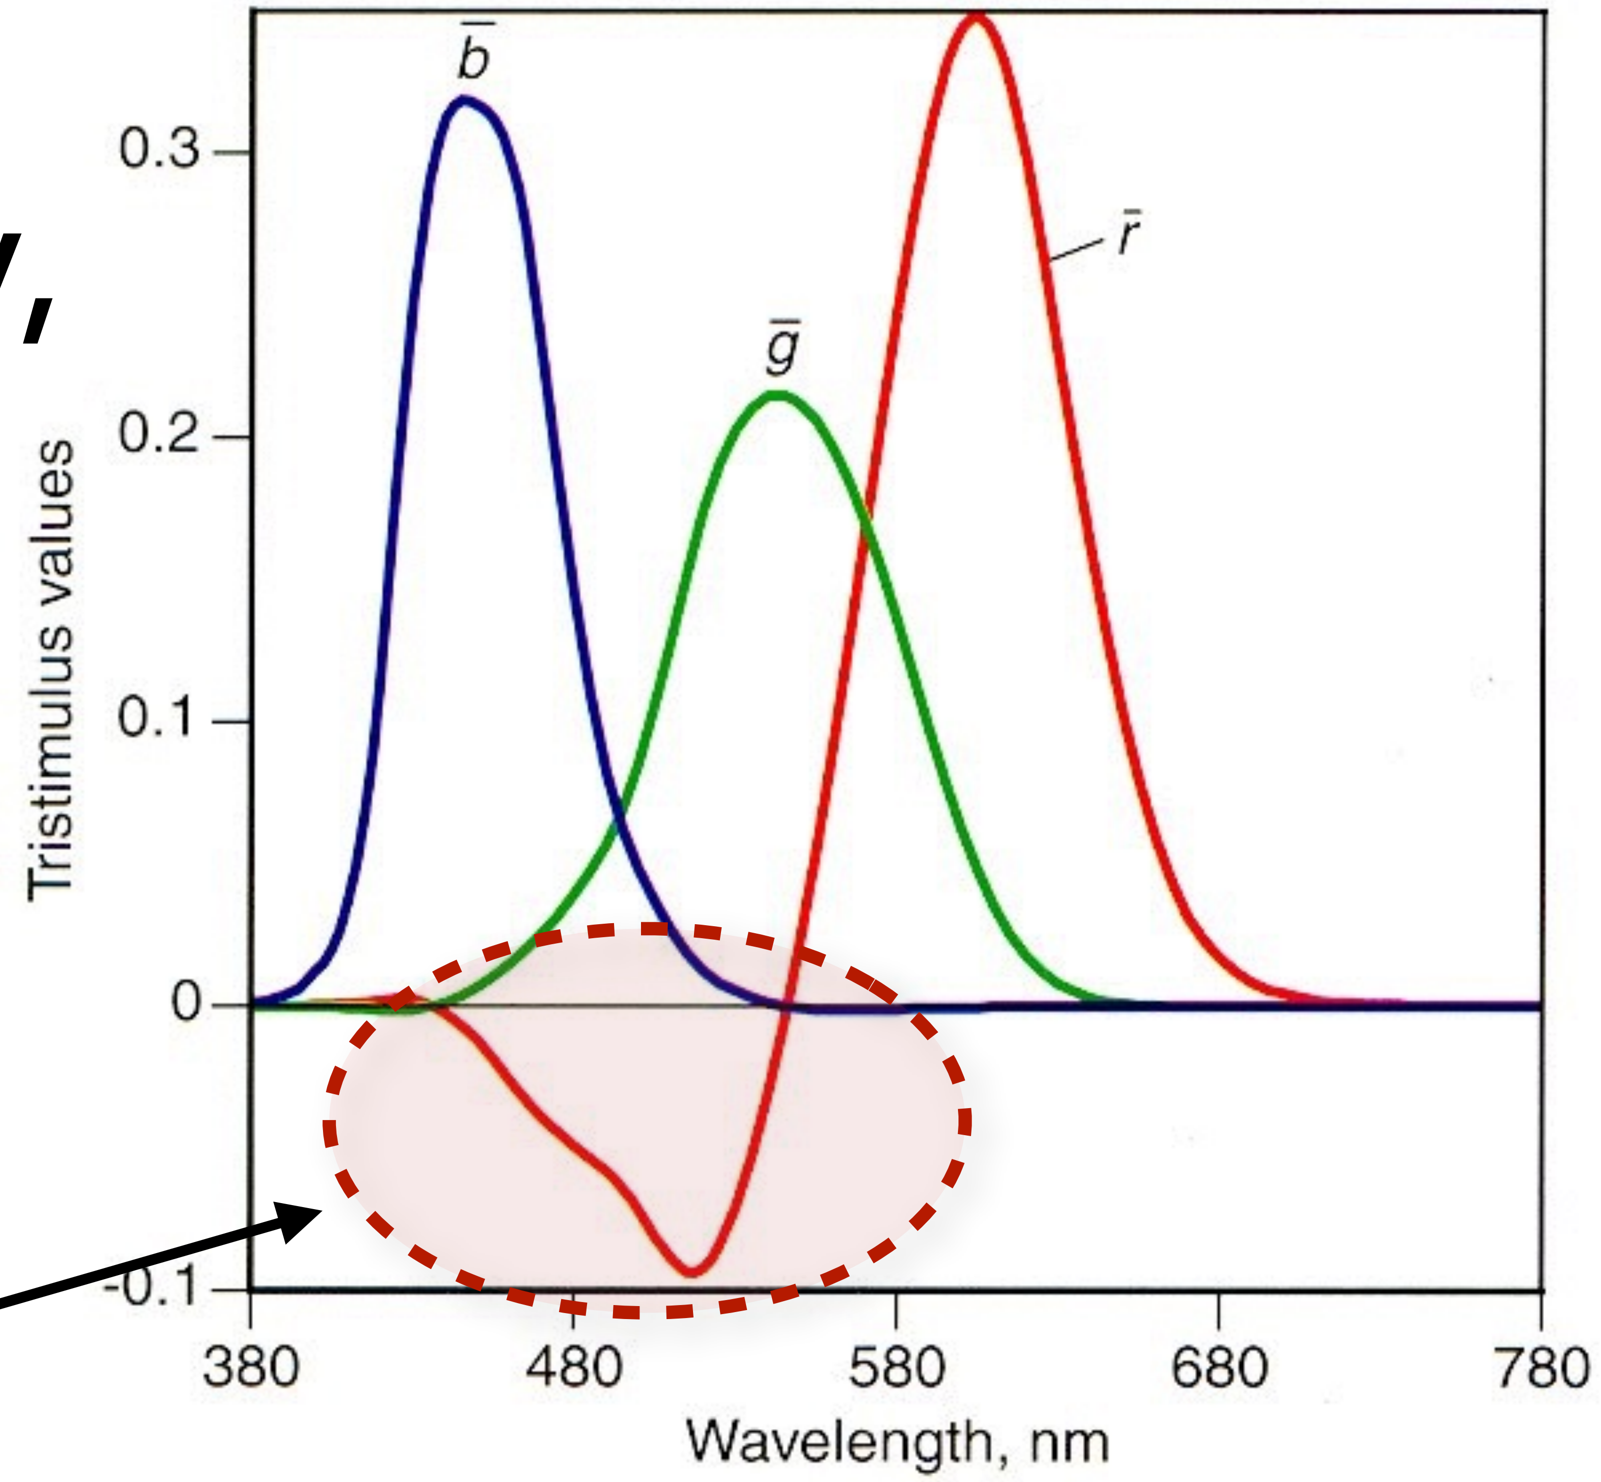

These curves are the color-matching functions for the 1931 standard observer, The average results of 17 color-normal observers having matched each wavelength of the equal-energy spectrum with primaries of 435.8 nm, 546.1 nm, and 700 nm.

Negative matching values?

Negative light doesn't exist, so what do these mean?

Some spectral colors could not be matched by primaries in the experiment

The "Trick":

- One primary could be added to the source
- Match with the remaining two
- Weight of primary added to the source is considered negative

But negative light is…inconvenient

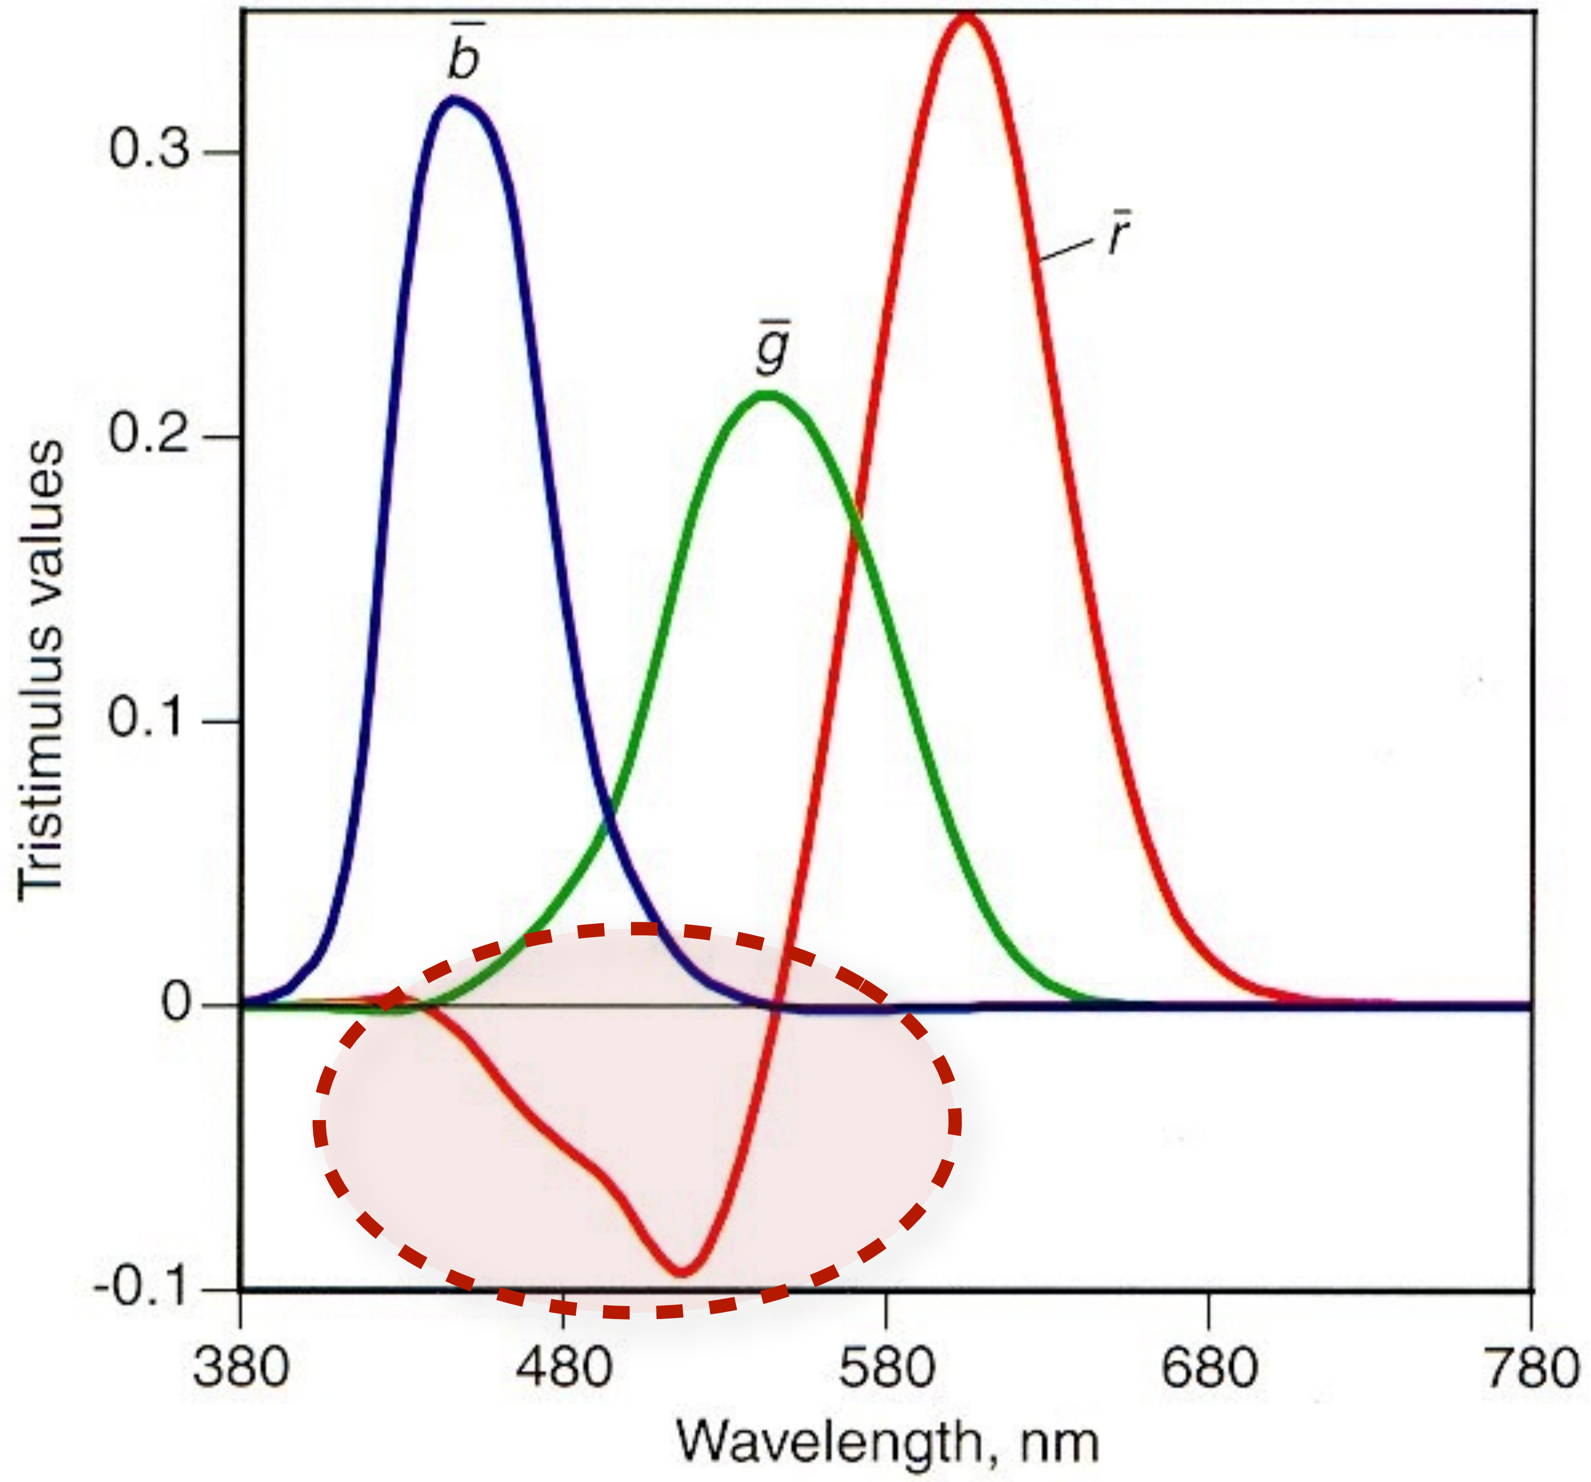

These curves are the color-matching functions for the 1931 standard observer, The average results of 17 color-normal observers having matched each wavelength of the equal-energy spectrum with primaries of 435.8 nm, 546.1 nm, and 700 nm.

CIE color spaces

CIE was not satisfied with range of RGB values for visible colors

- Negative tristimulus values

Defined CIE XYZ color space via simple mathematical transformation

- http://en.wikipedia.org/wiki/CIE_1931_color_space#Definition_of_the_CIE_XYZ_color_space

Most common color space still today

CIE XYZ color space

Infinitely many ways to obtain non-negative matching functions!

Let's call ours XYZ

- Y measures *brightness* or *luminance*
- Set white to XYZ=(1/3,1/3,1/3)
- imaginary primaries "supersaturated"

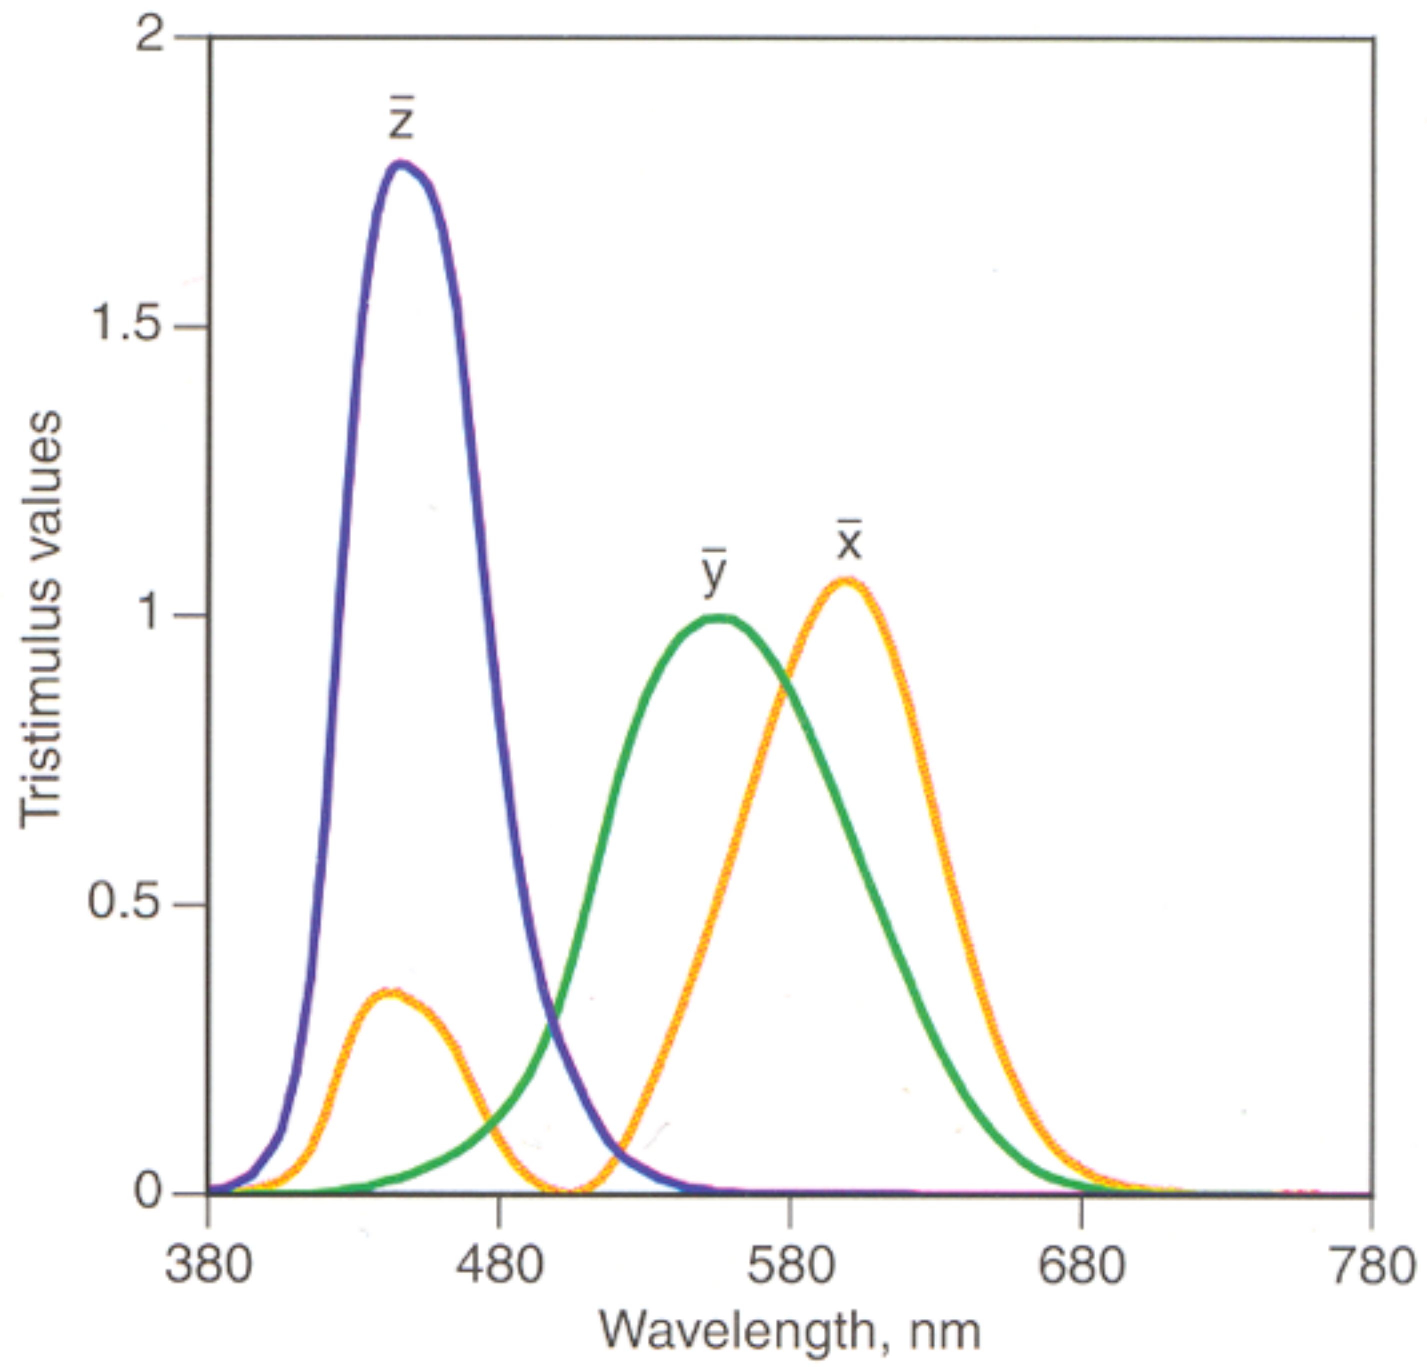

Linear transformation of CIE RGB

$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = \frac{1}{b_{21}} \begin{bmatrix} b_{11} & b_{12} & b_{13} \\ b_{21} & b_{22} & b_{23} \\ b_{31} & b_{32} & b_{33} \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix} = \frac{1}{0.17697} \begin{bmatrix} 0.49 & 0.31 & 0.20 \\ 0.17697 & 0.81240 & 0.01063 \\ 0.00 & 0.01 & 0.99 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix}$$

CIE XYZ ←- CIE RGB

XYZ to RGB & back

sRGB to XYZ

0.412424	0.212656	0.0193324
0.357579	0.715158	0.119193
0.180464	0.0721856	0.950444

XYZ to sRGB

3.24071	-0.969258	0.0556352
-1.53726	1.87599	-0.203996
-0.498571	0.0415557	1.05707

Adobe RGB to XYZ

0.576700	0.297361	0.0270328
0.185556	0.627355	0.0706879
0.188212	0.0752847	0.991248

XYZ to Adobe RGB

2.04148	-0.969258	0.0134455
-0.564977	1.87599	-0.118373
-0.344713	0.0415557	1.01527

NTSC RGB to XYZ

0.606734	0.298839	0.000000
0.173564	0.586811	0.0661196
0.200112	0.114350	1.11491

XYZ to NTSC RGB

1.91049	-0.984310	0.0583744
-0.532592	1.99845	-0.118518
-0.288284	-0.0282980	0.898611

http://www.brucelindbloom.com/index.html?Eqn_RGB_XYZ_Matrix.html

CIE color matching Recap

CIE performed color matching experiments

- chose primaries for reproduction (synthesis)
- for each wavelength, how much of each primary do we need
    - 3 analysis curves
- Then a little bit of linear algebra to make everything positive
    - 3 new analysis curves

Gives us XYZ color space

Linear transform to/from LMS, RGB

CIE XYZ Recap

The most widely recognized color space

Y corresponds to brightness (1924 CIE standard photometric observer)

No negative values in matching curves

But no physically-realizable primary (negative values in primary rather than in matching curve)

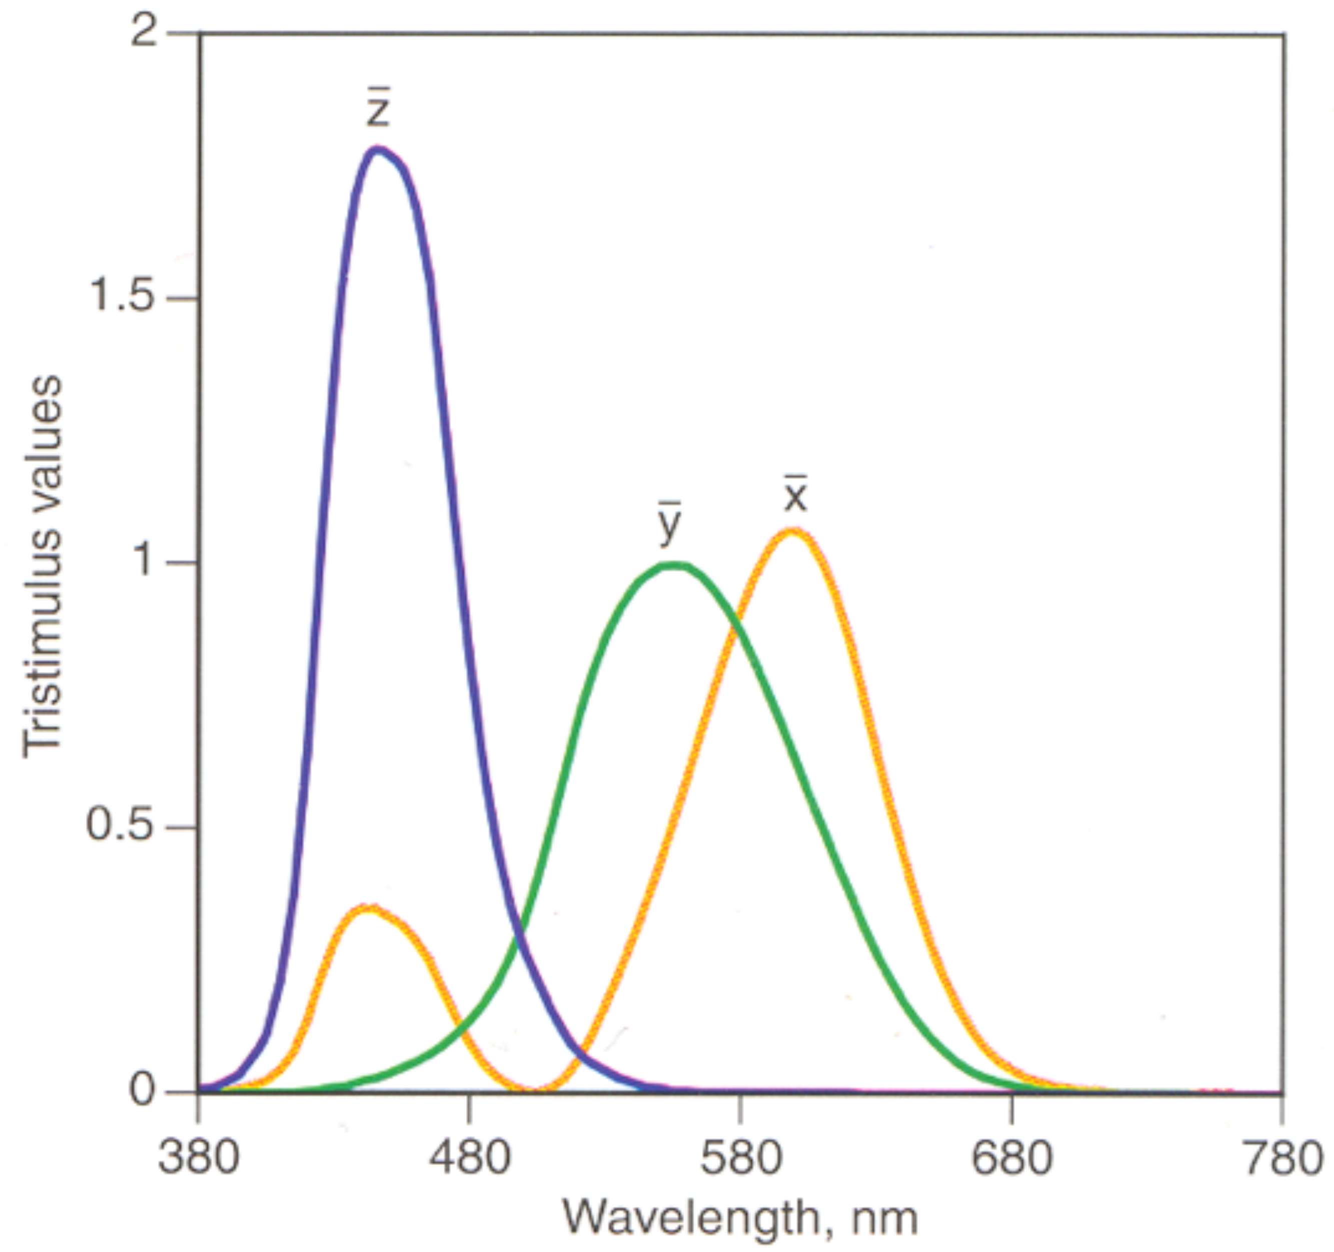

The 1931 standard observer, as it is usually shown.

Chromaticity Diagram

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

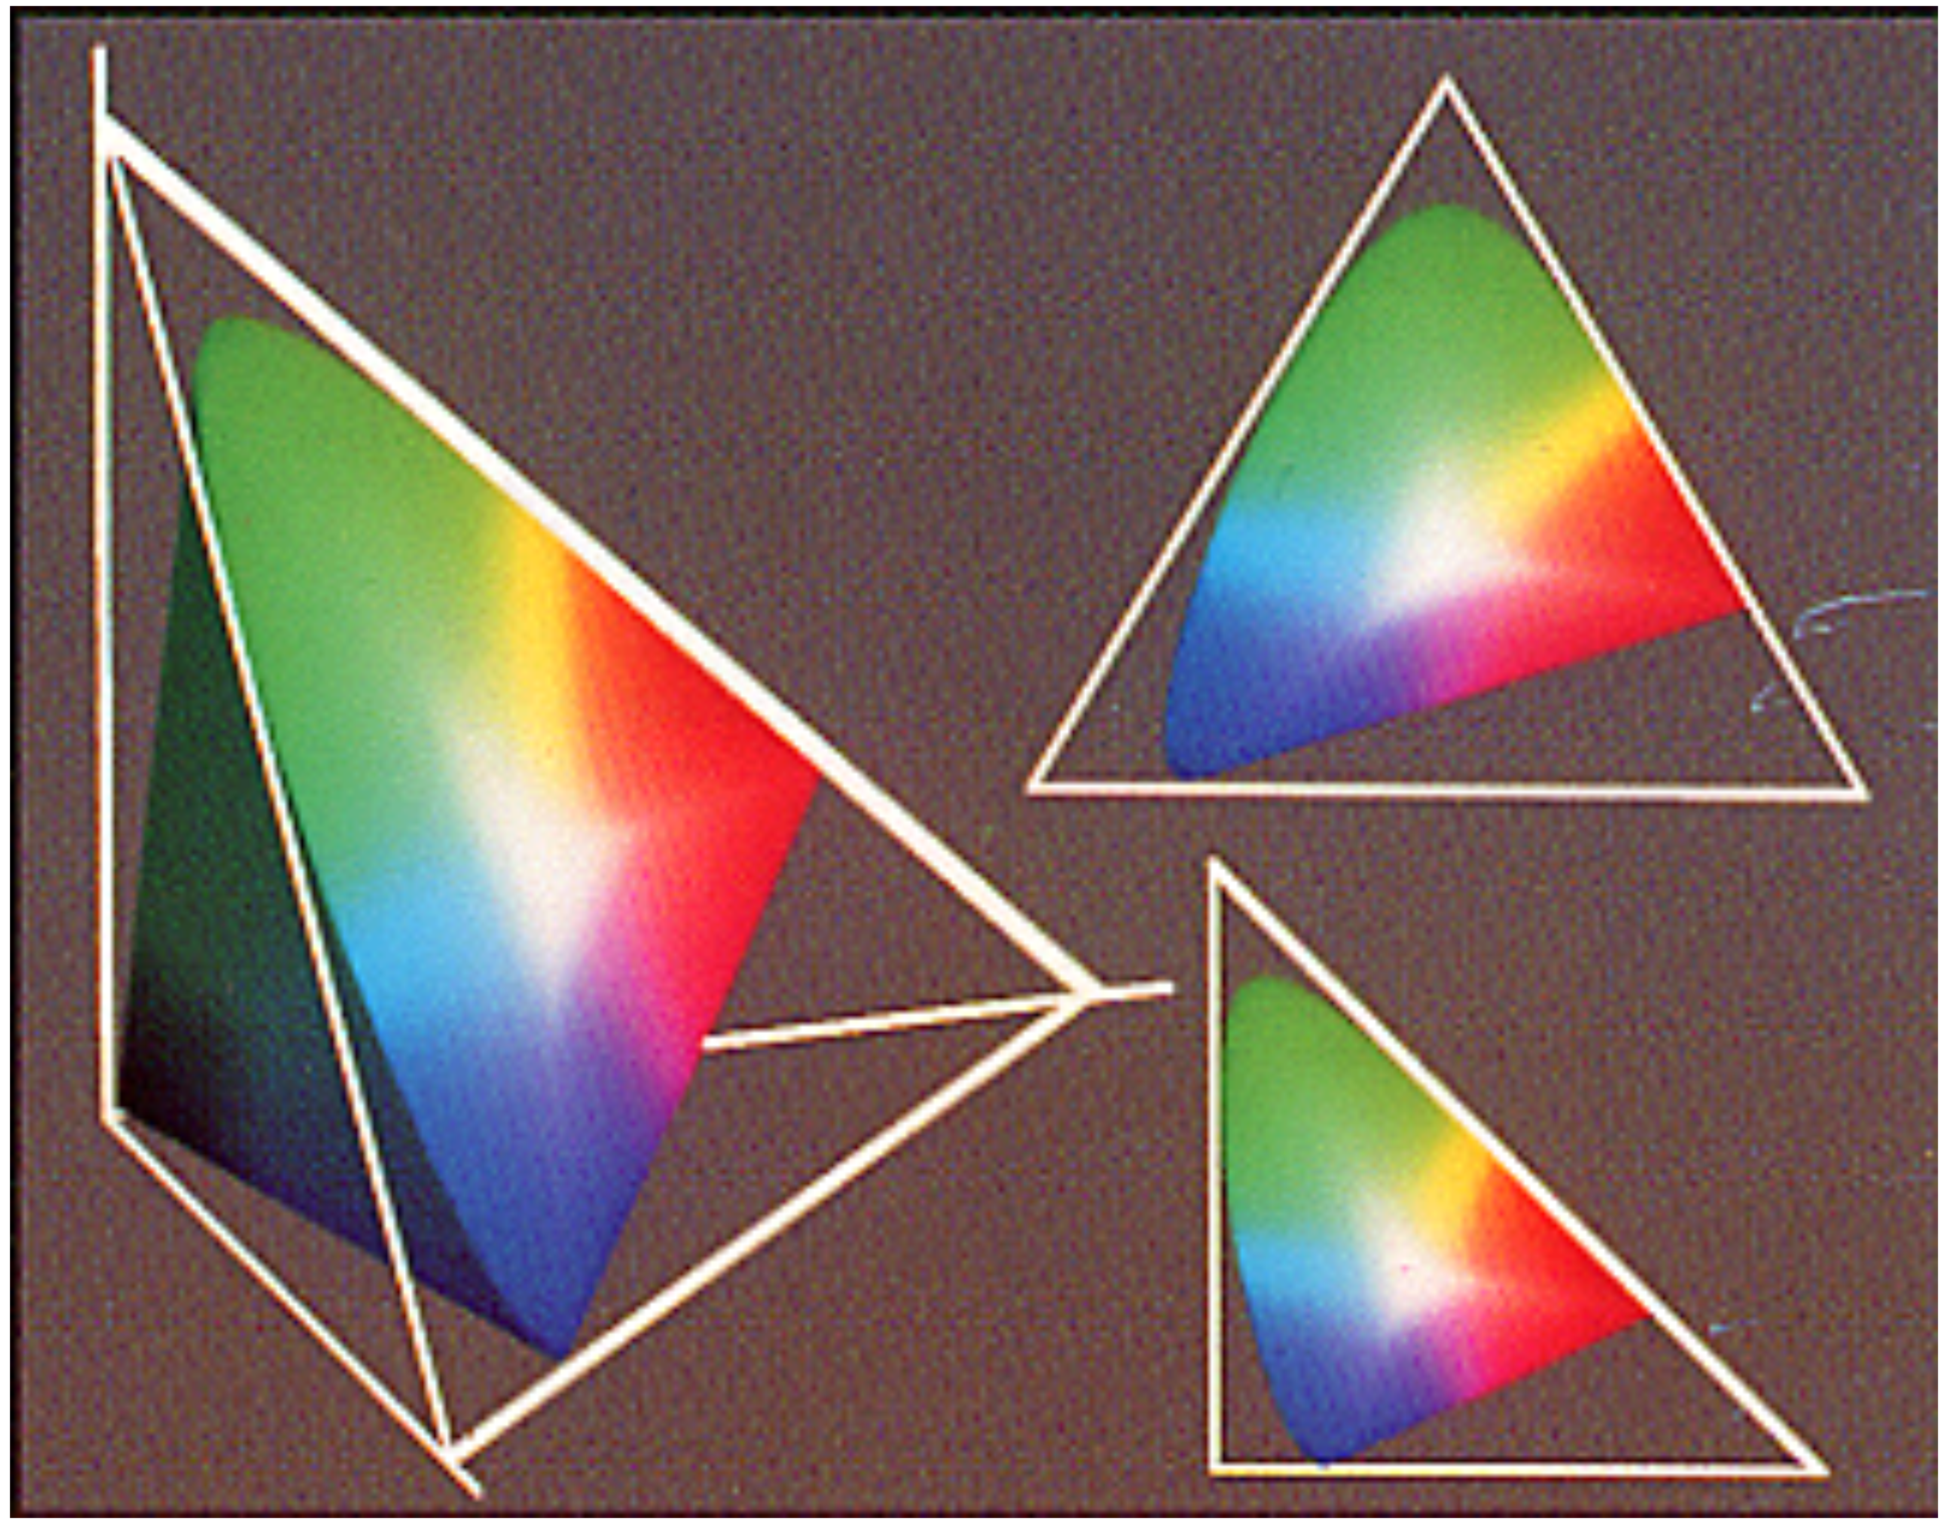

The CIE xyY Color Space

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)

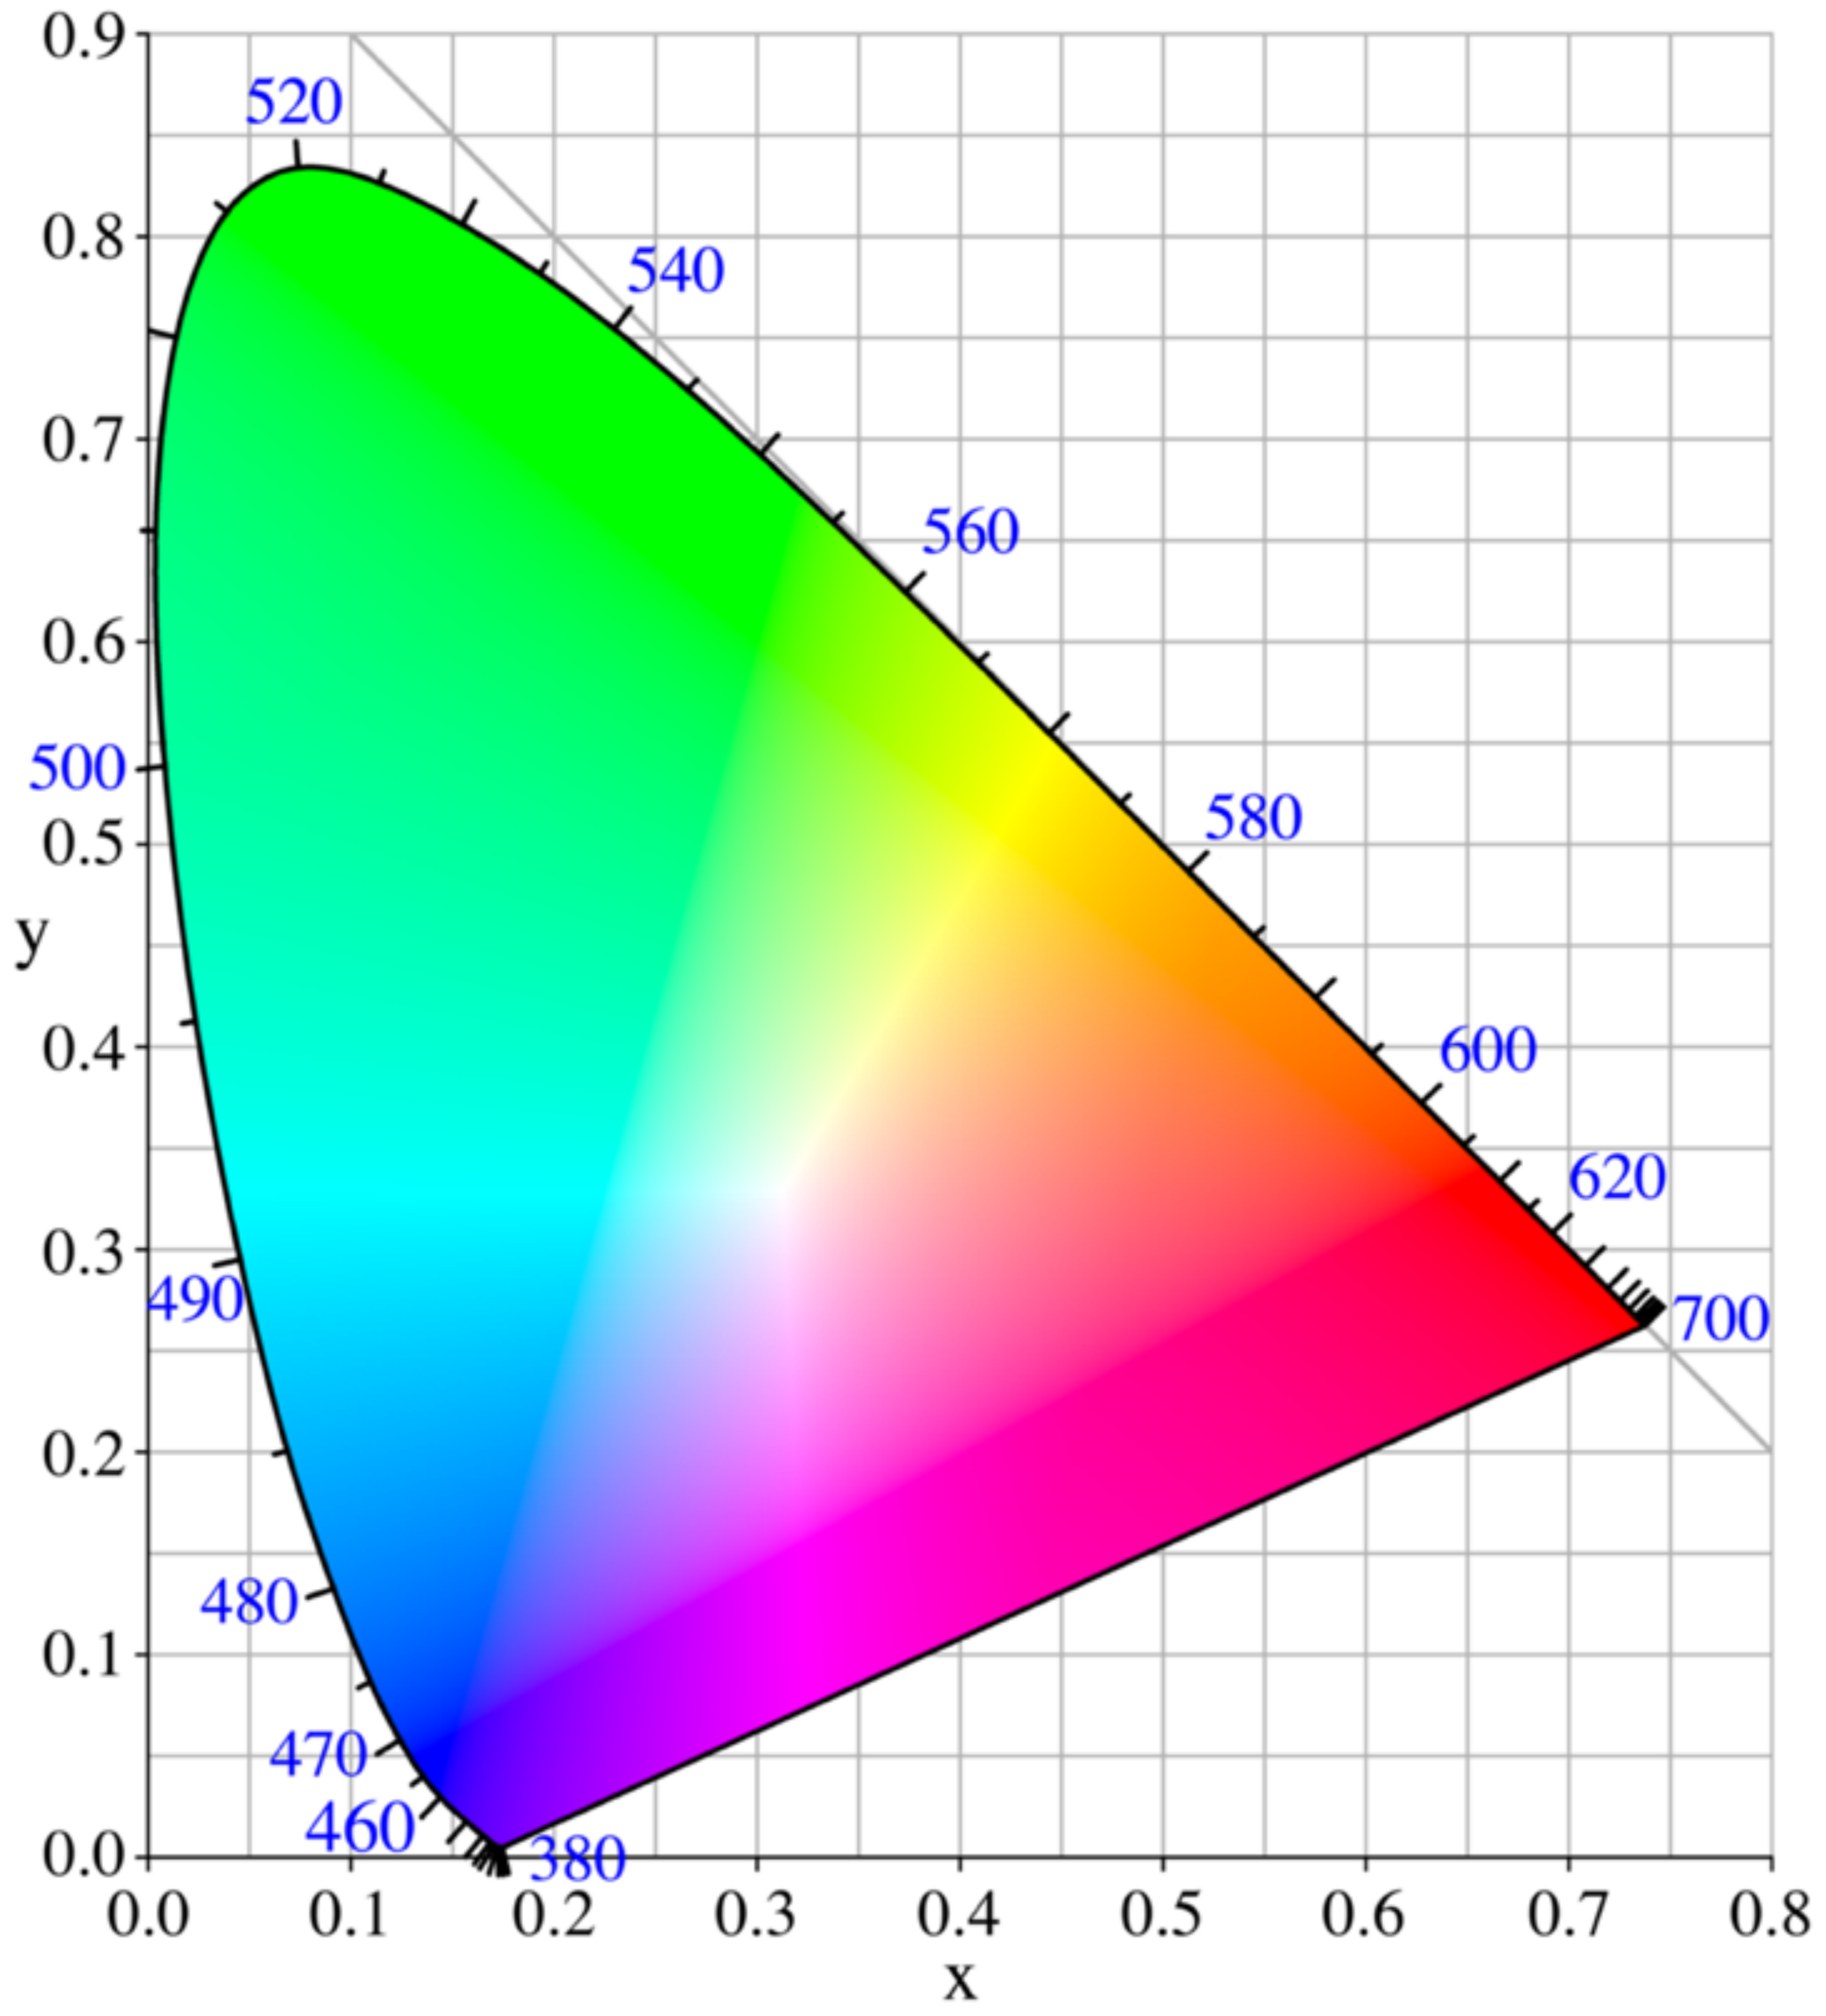

CIE RGB Color Space

Color primaries at: 435.8, 546.1, 700.0 nm

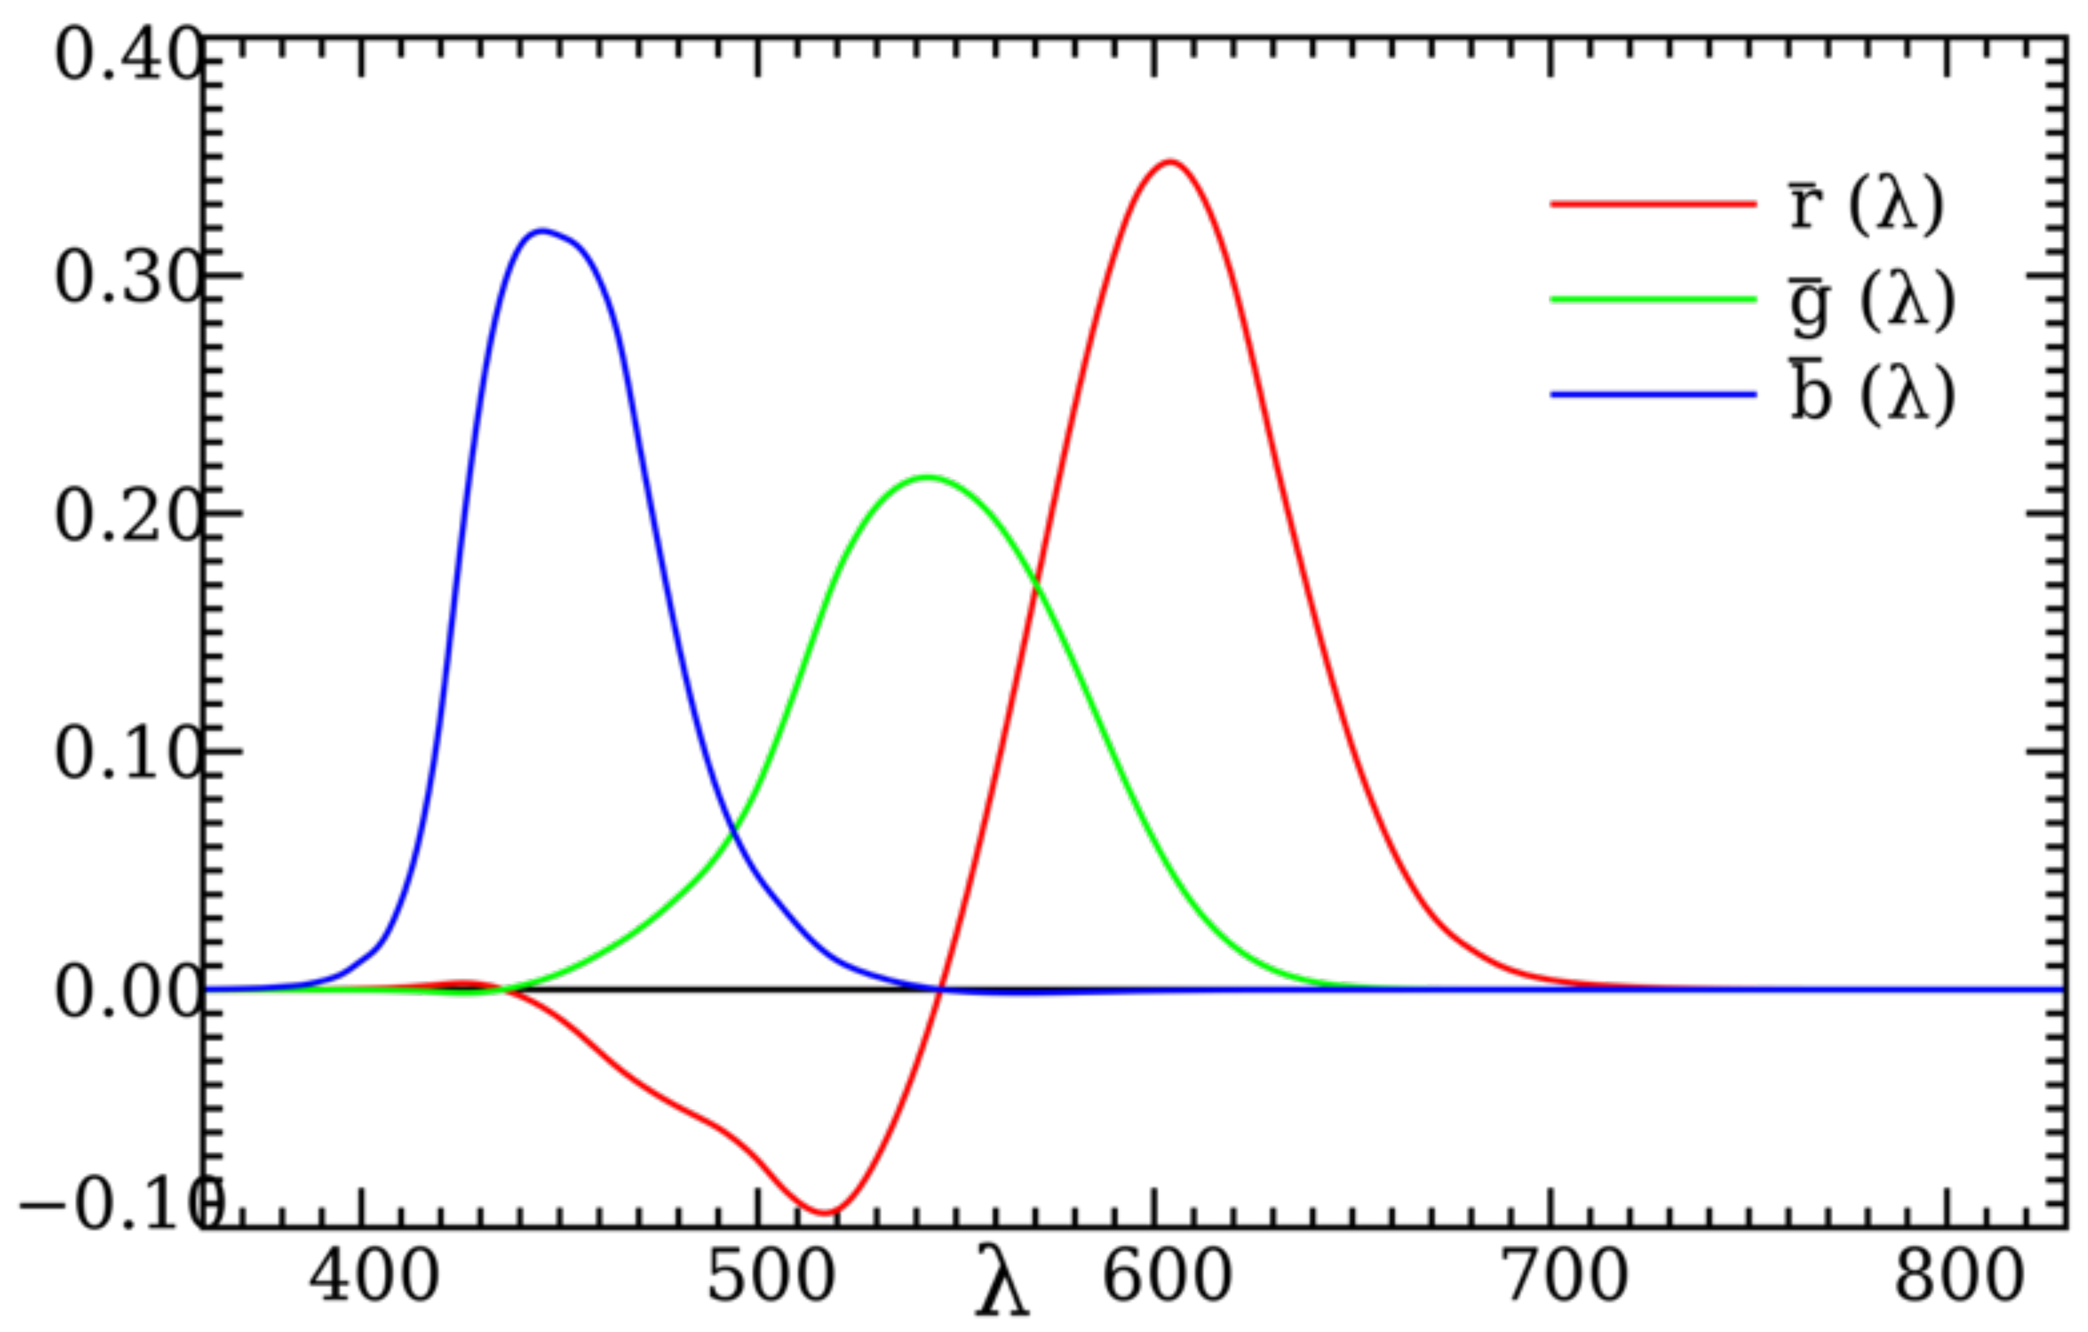

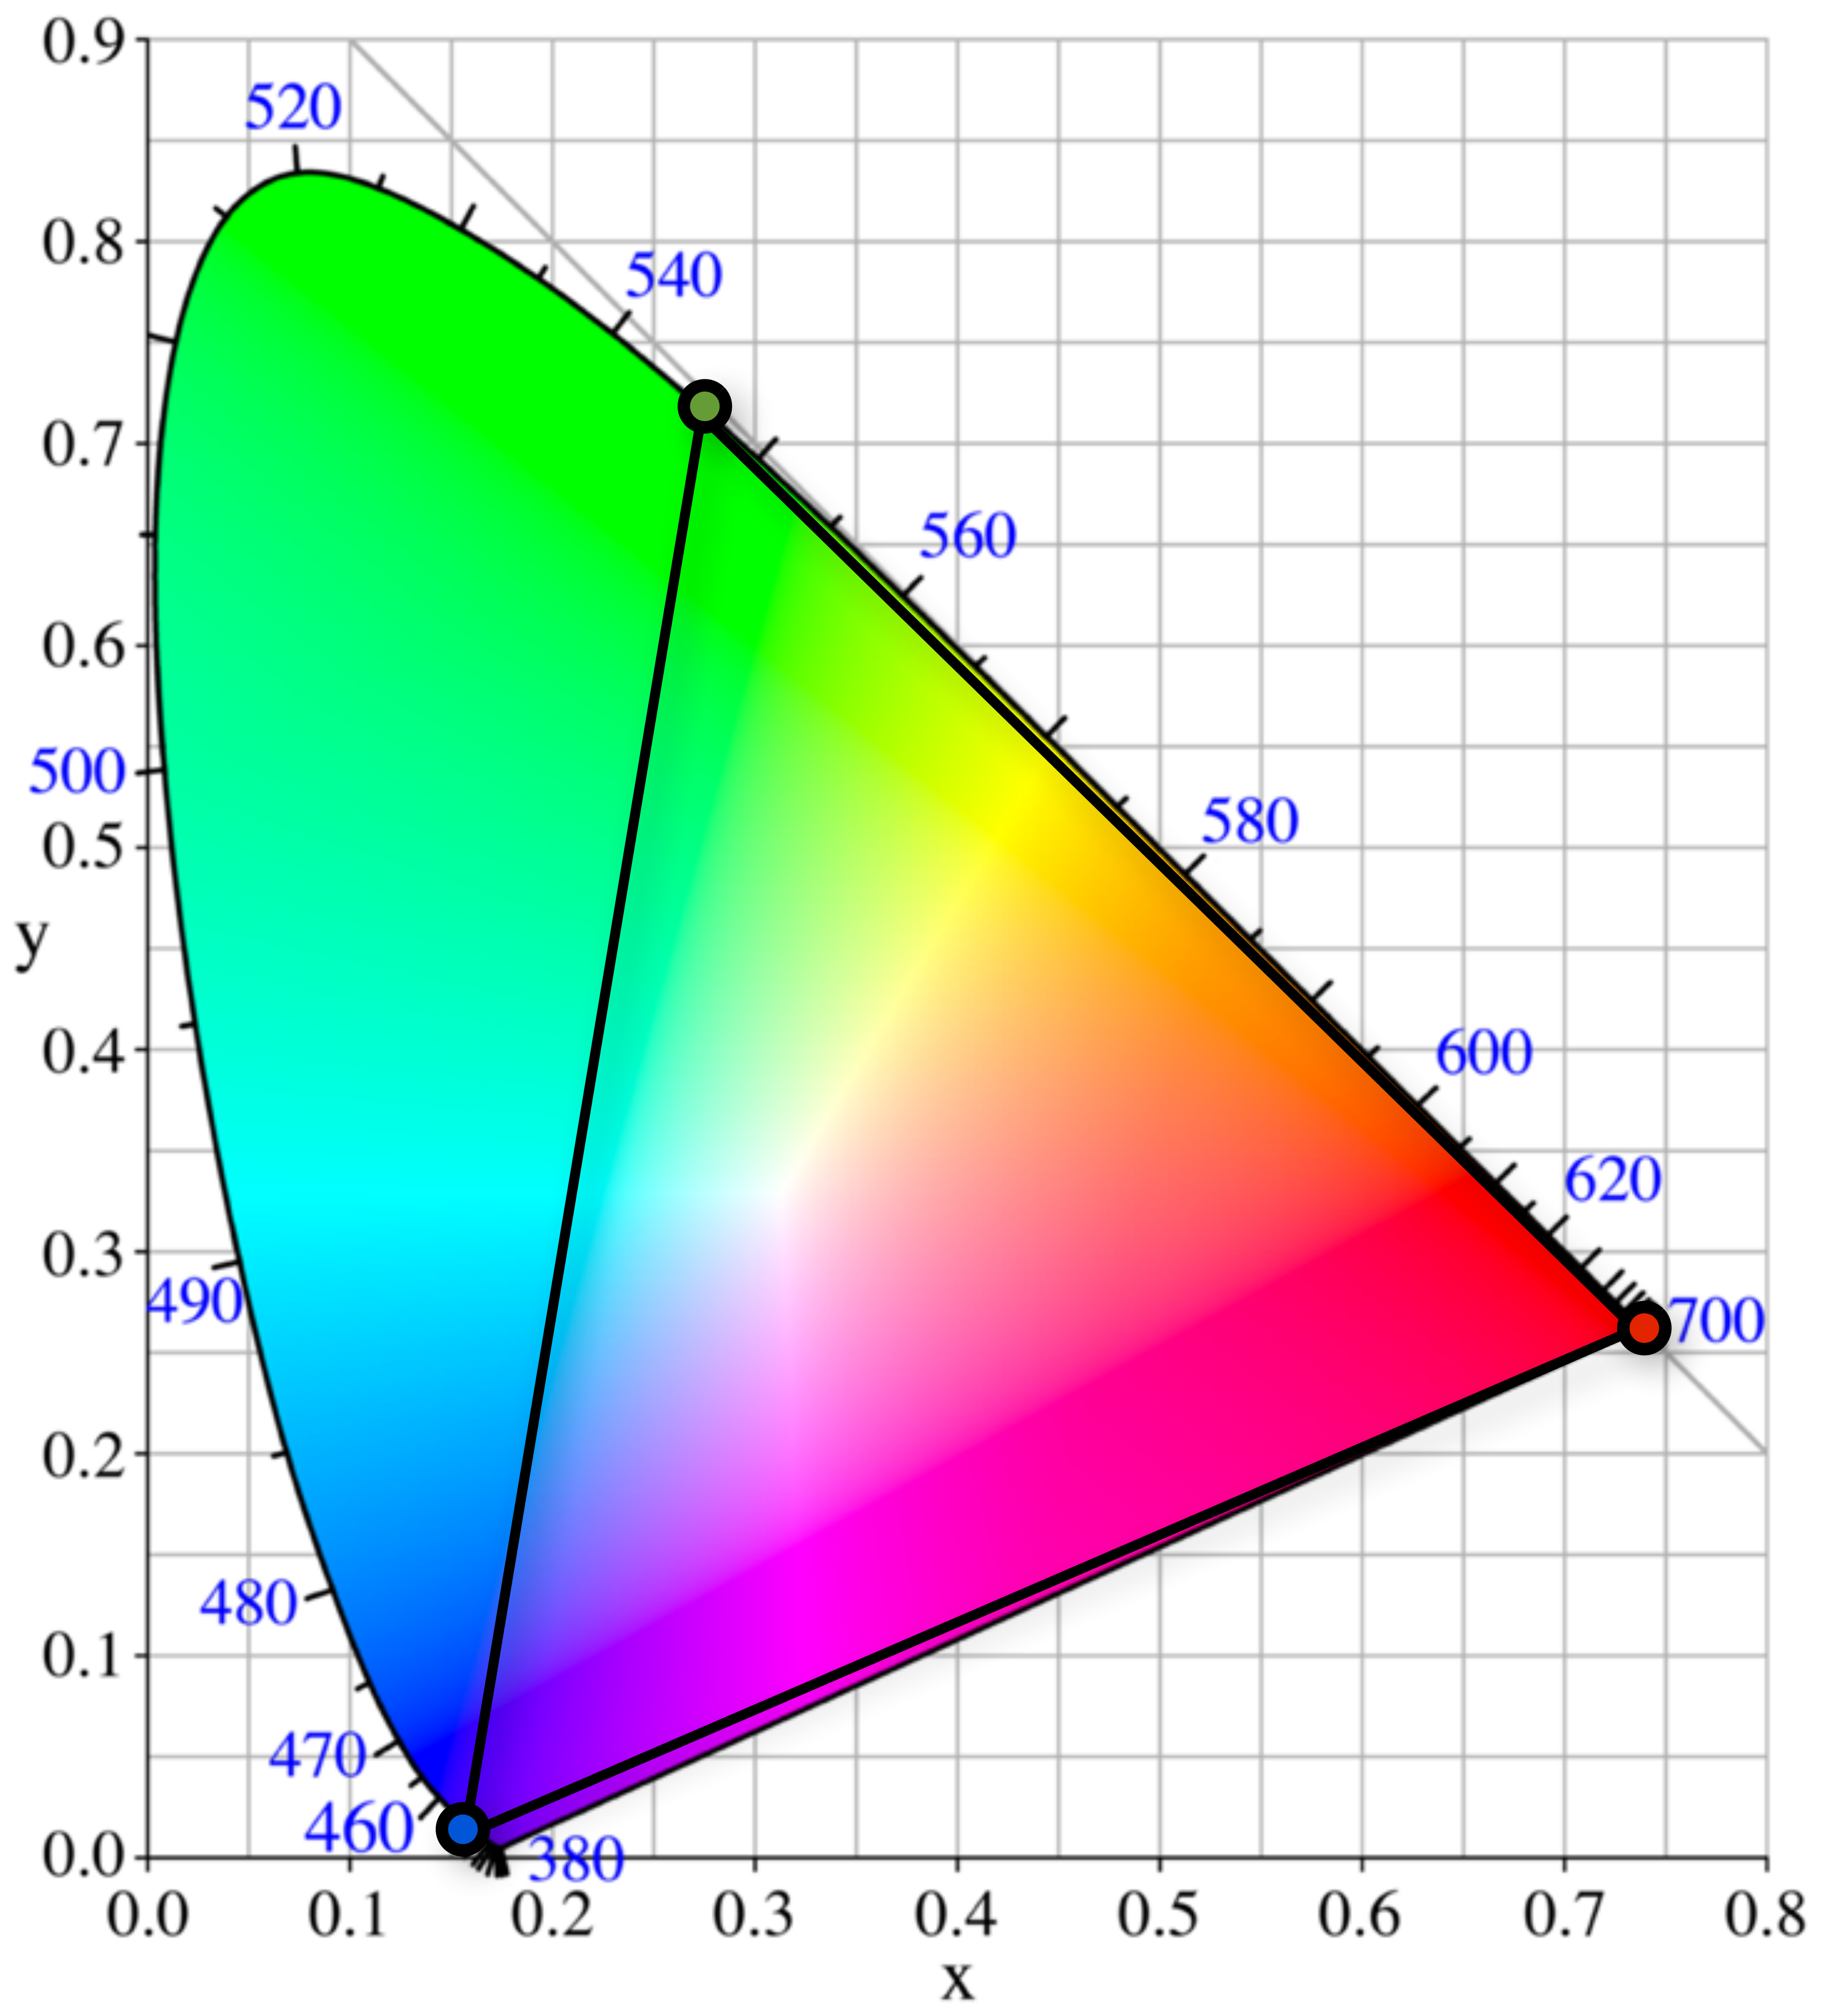

Color Gamuts

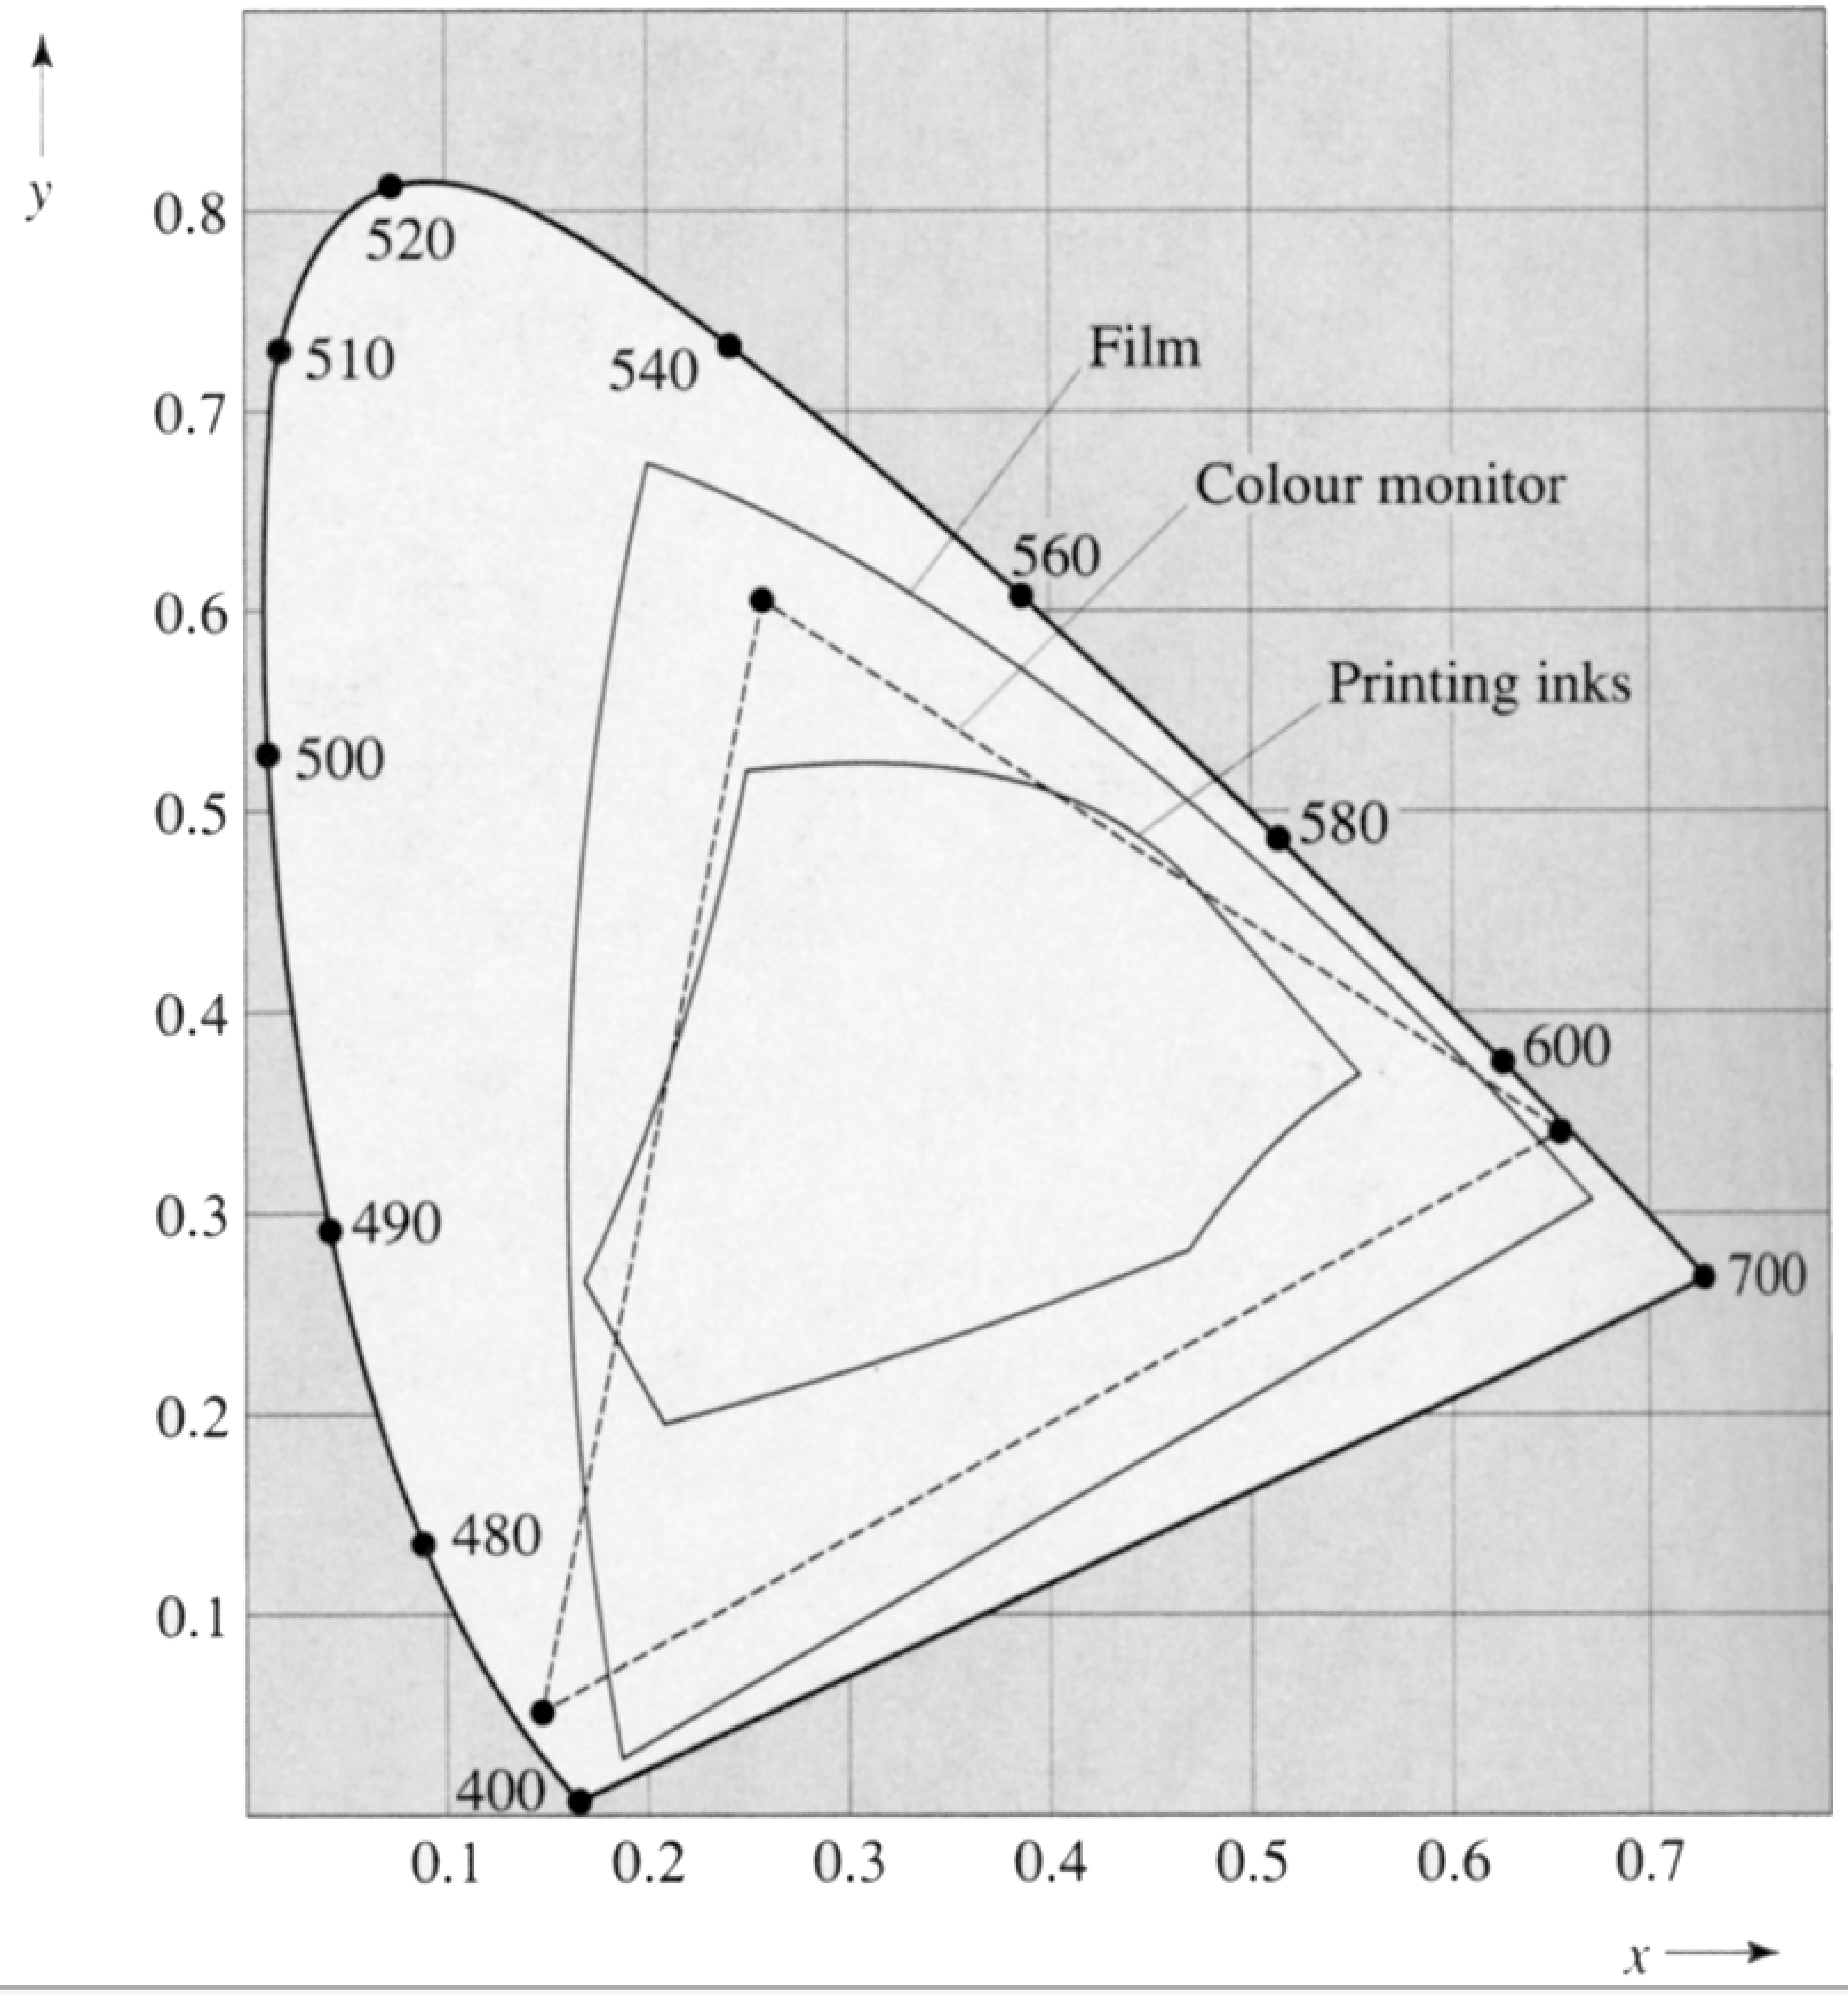

CIE Chromaticity Chart Features

White Point

Dominant wavelength

Inverse color

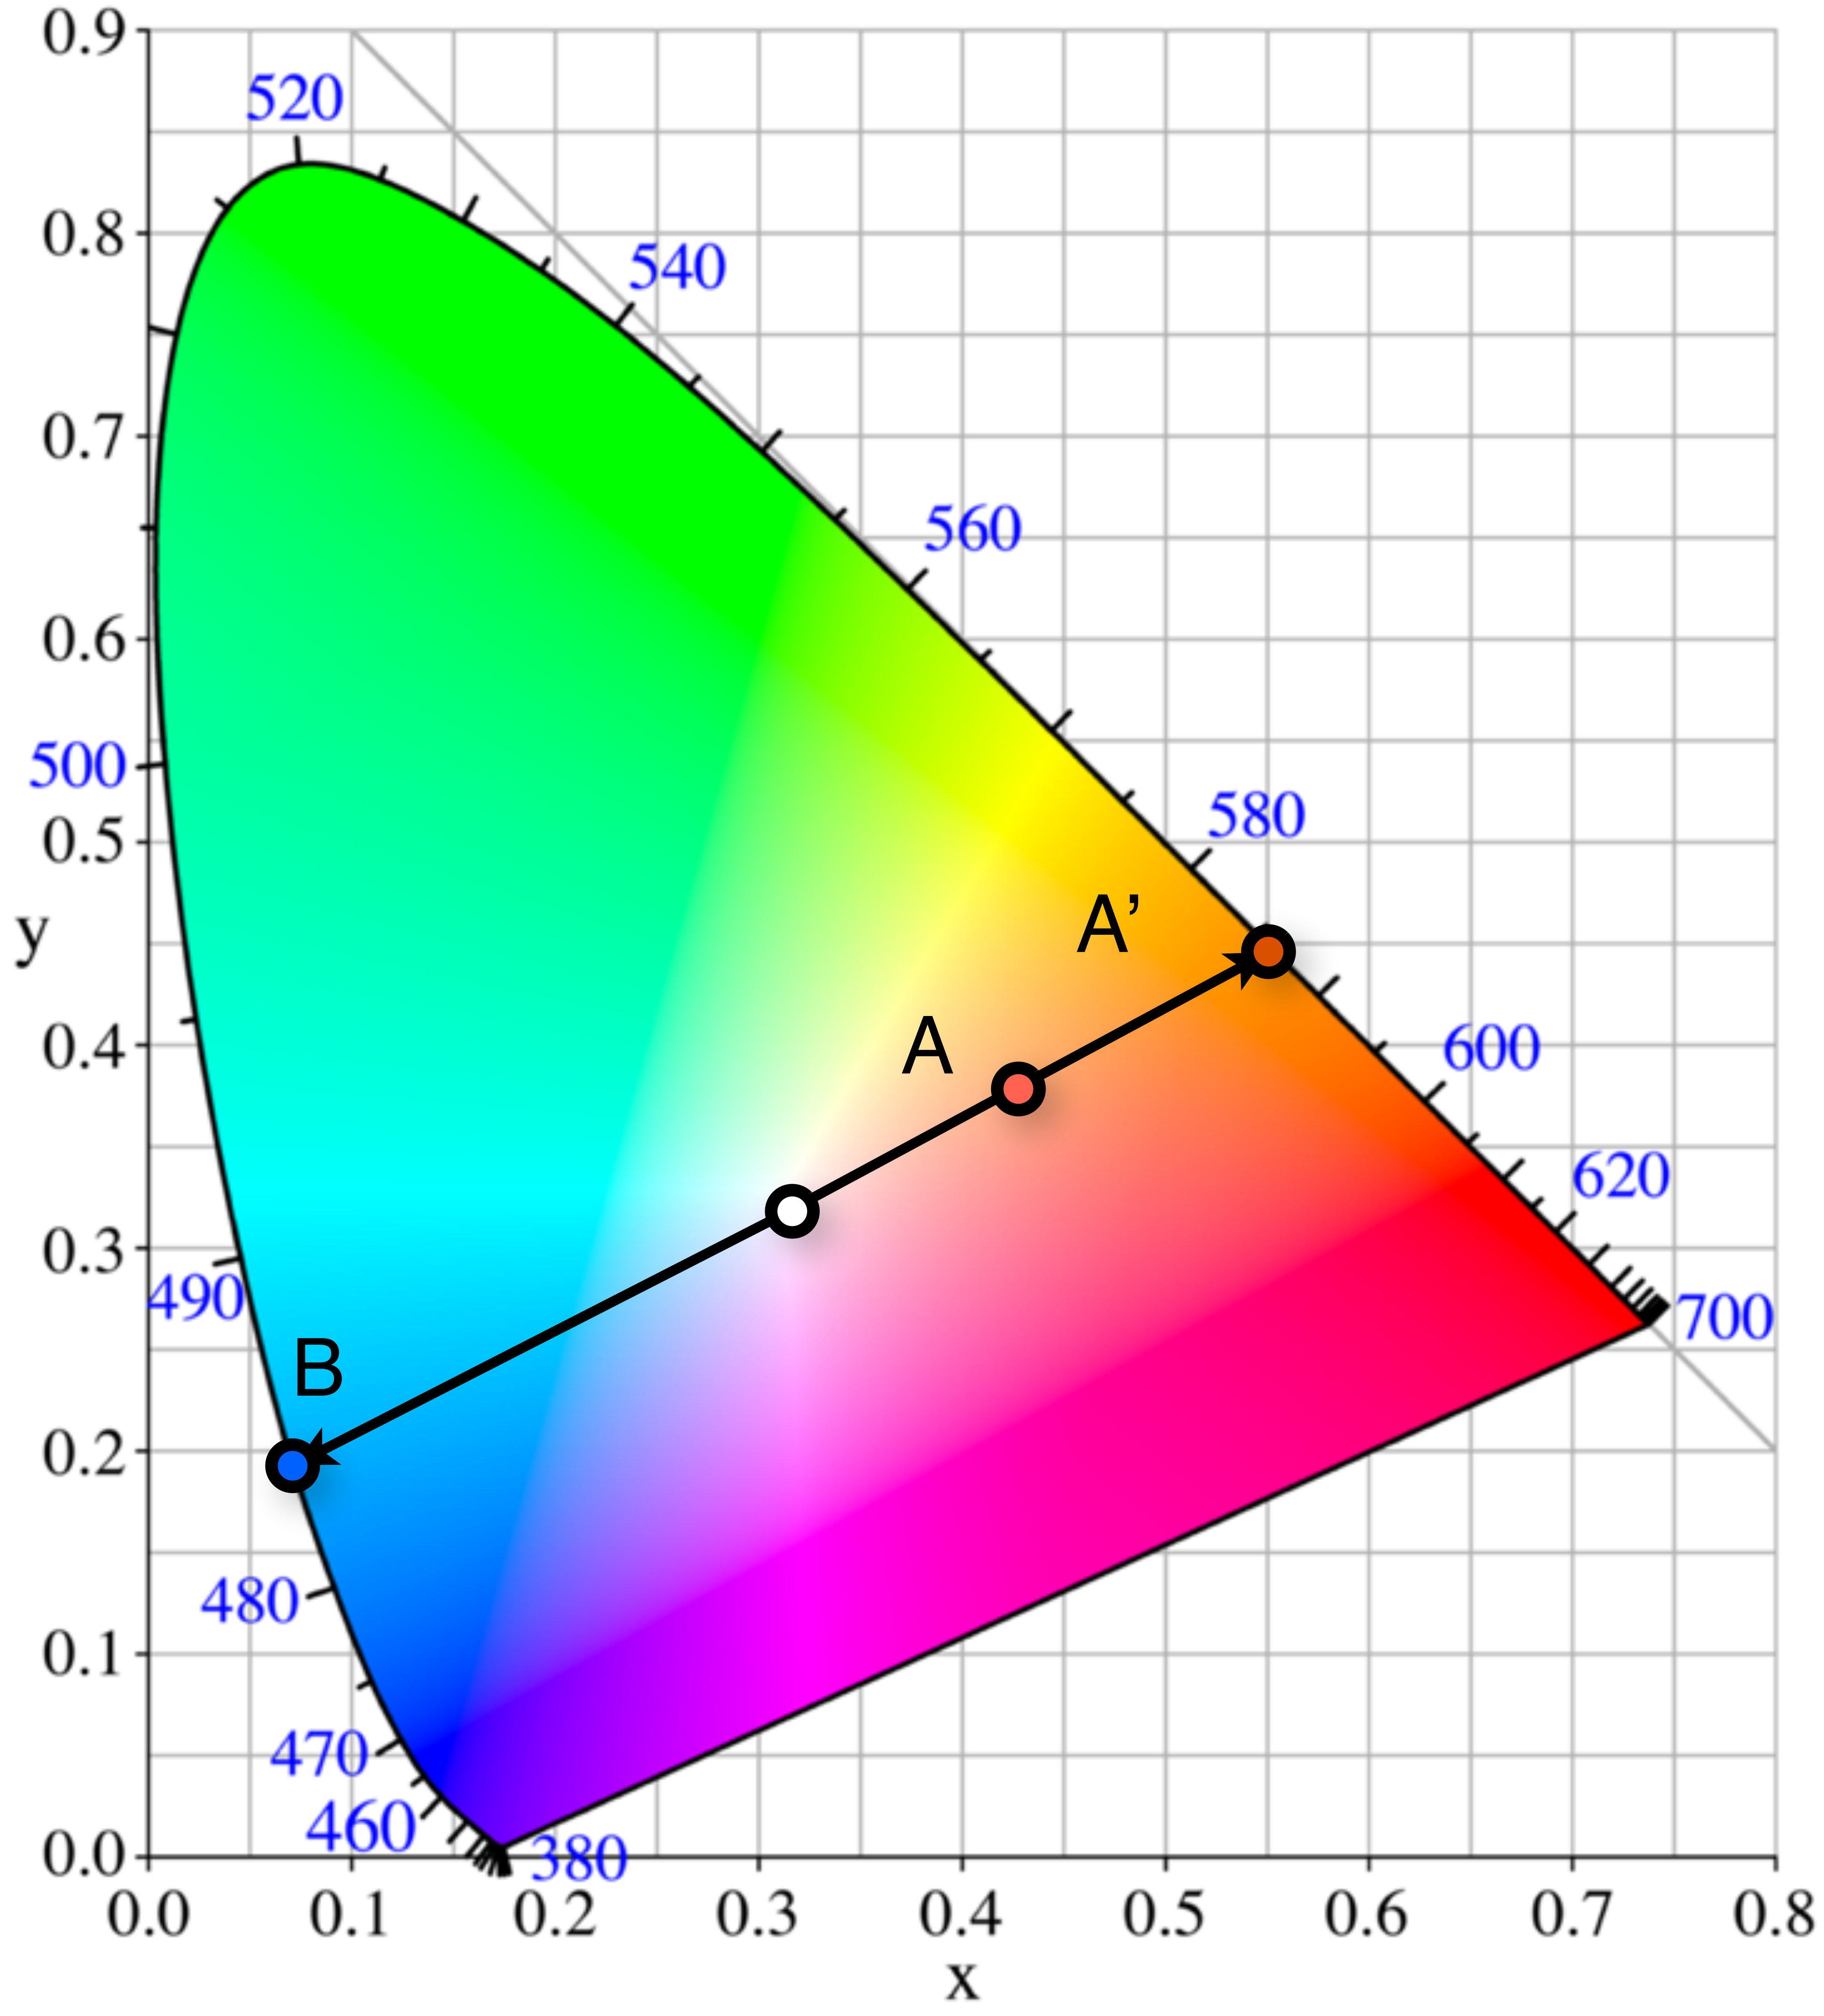

Perceptually-Uniform Color Spaces

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 and not necessarily very different!

Measuring “perceptual distance” in color spaces important for many industries

Experiments by MacAdams

MacAdams Color Ellipses

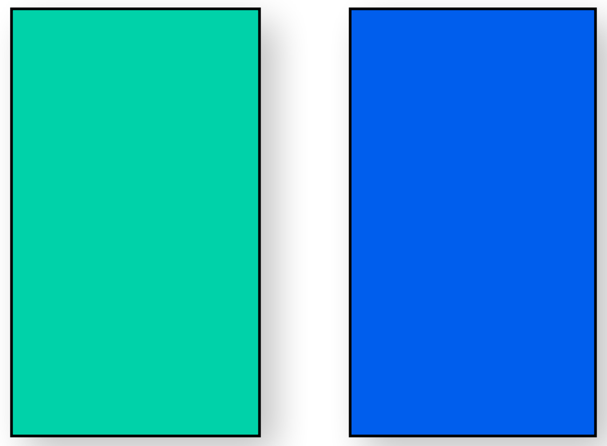

Test patches

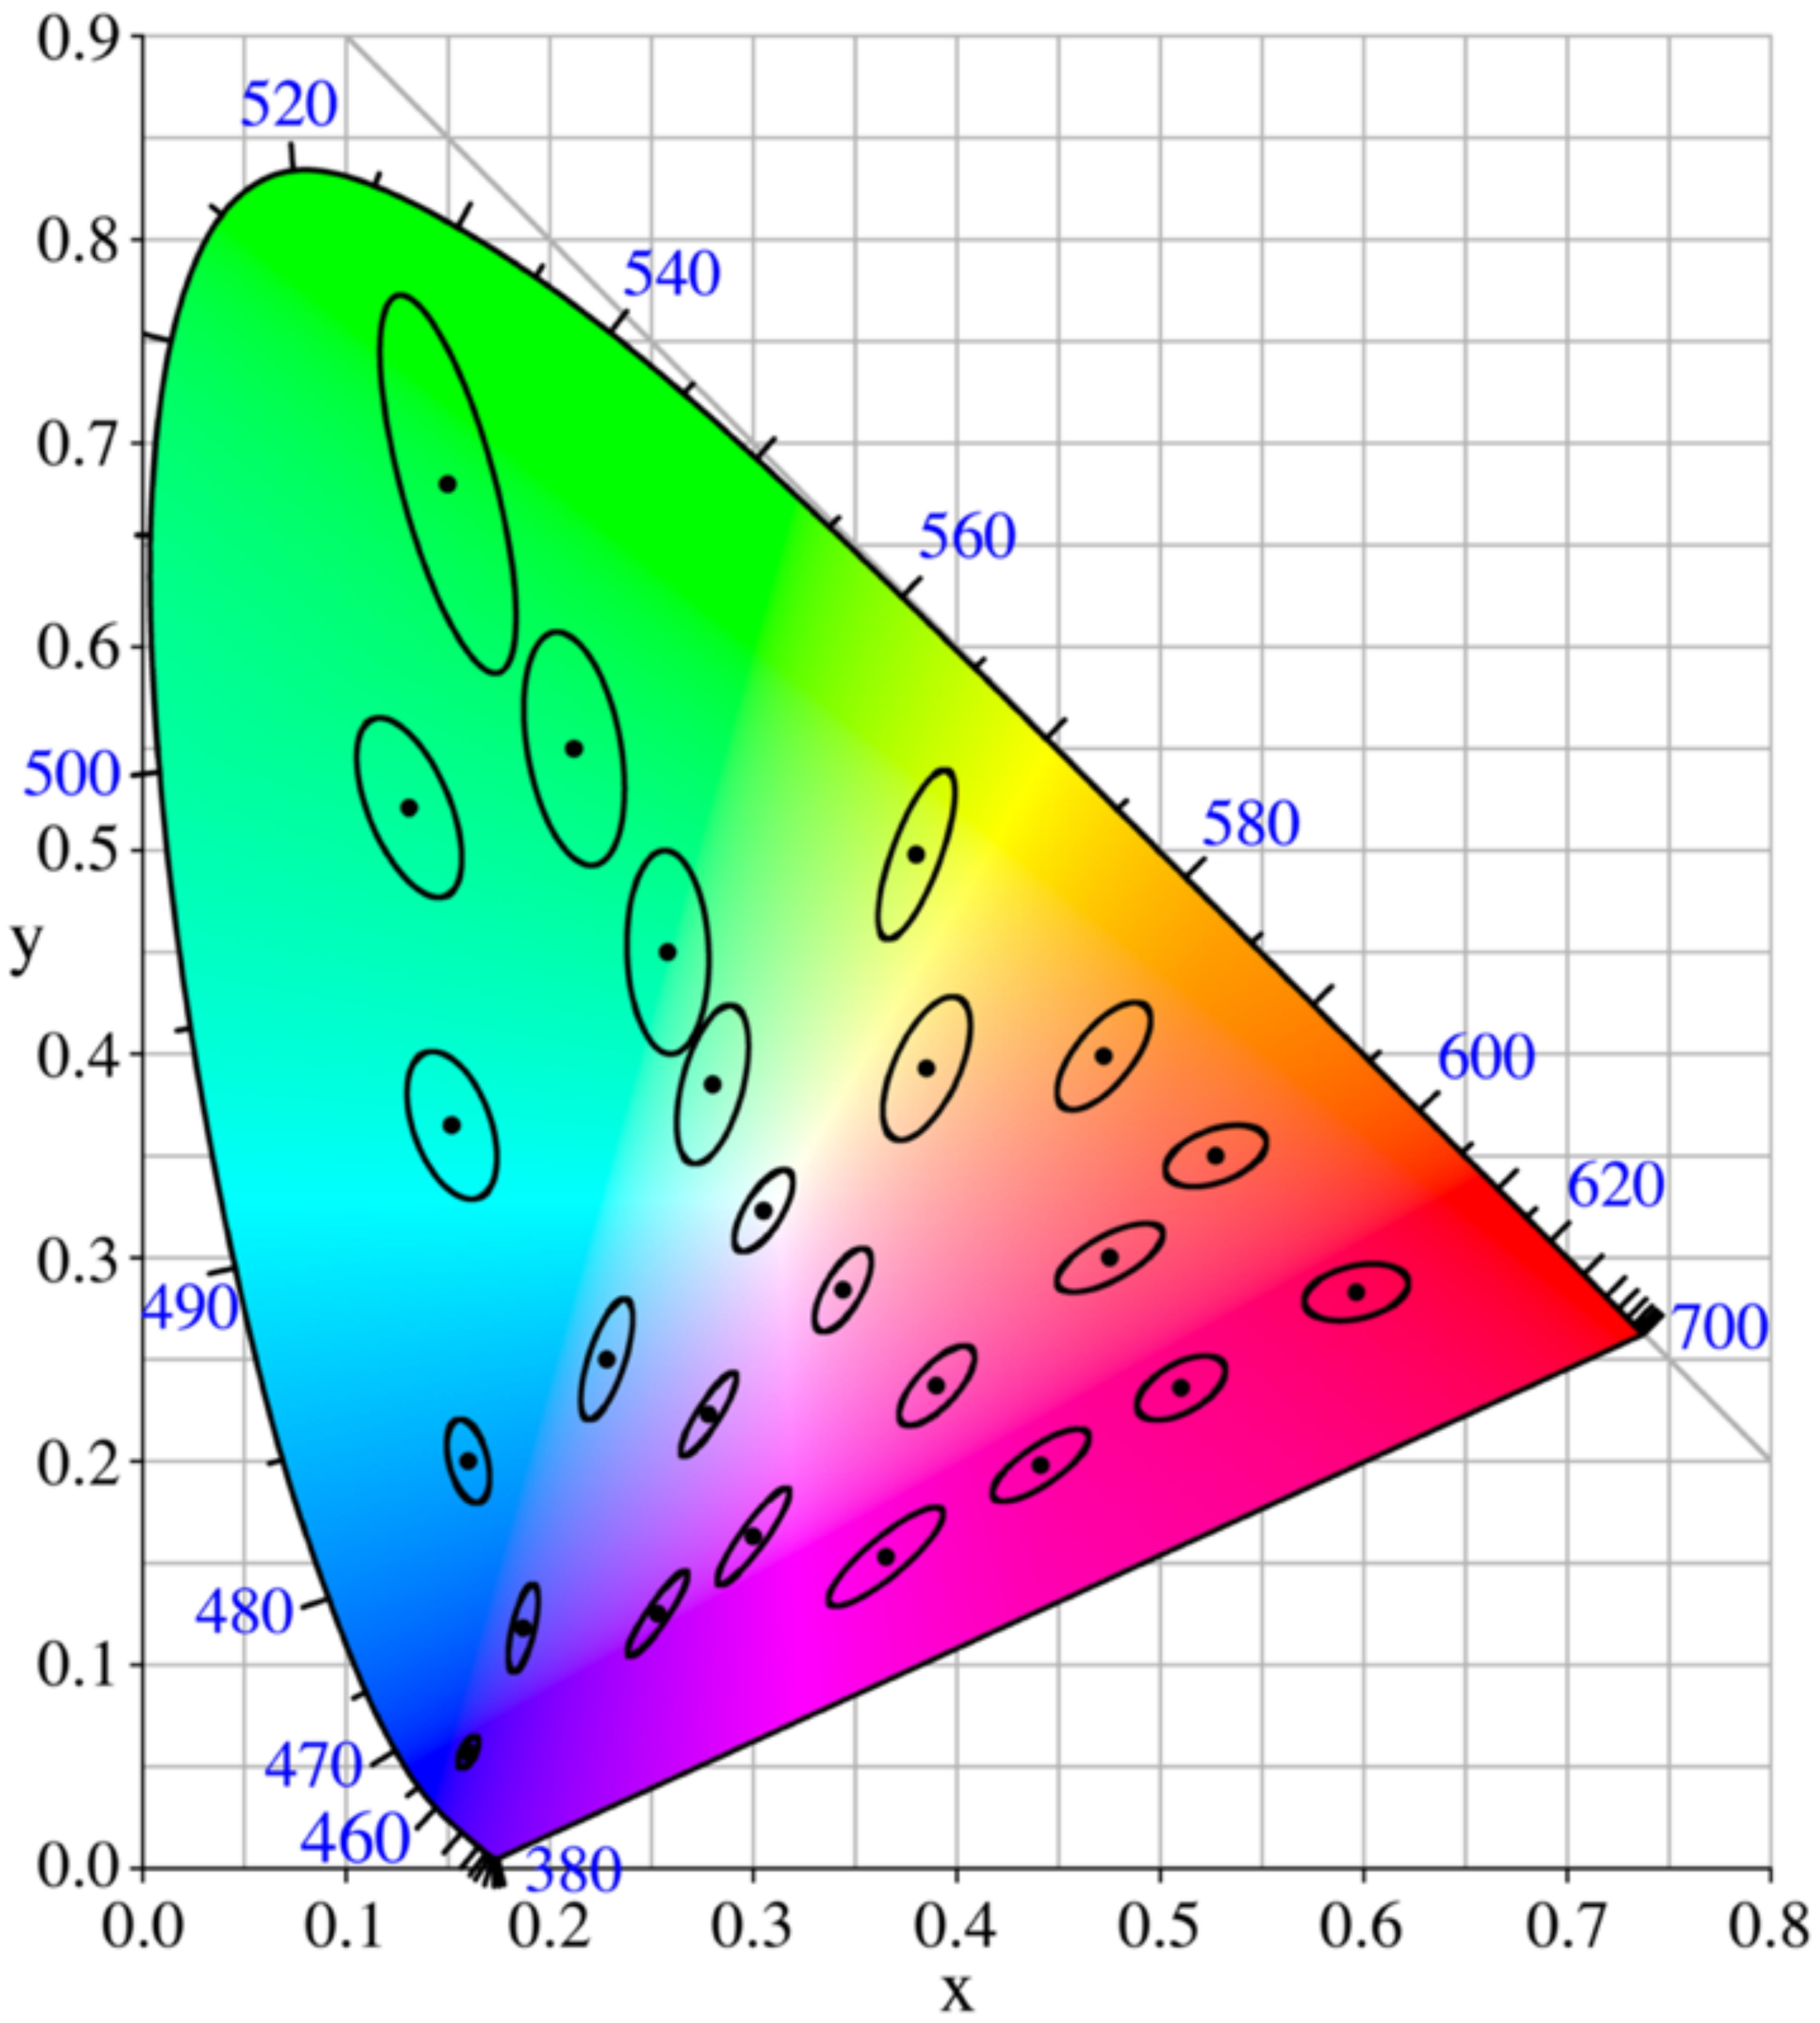

CIELab and CIELuv Color Spaces

Two attempts to make a perceptually-uniform color space

MacAdams ellipses become nearly (but not perfectly) circular

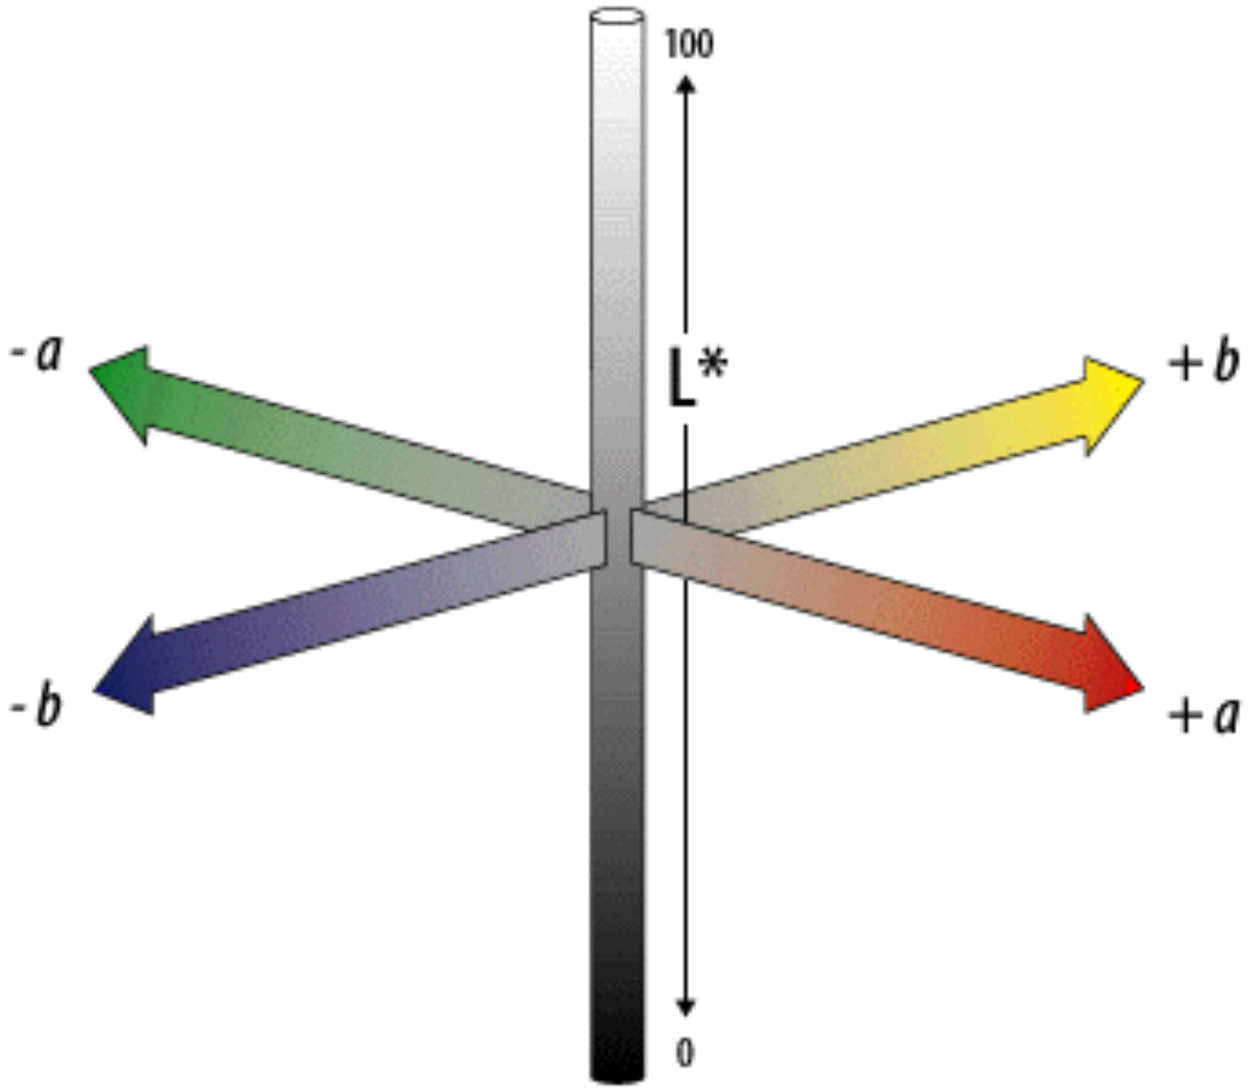

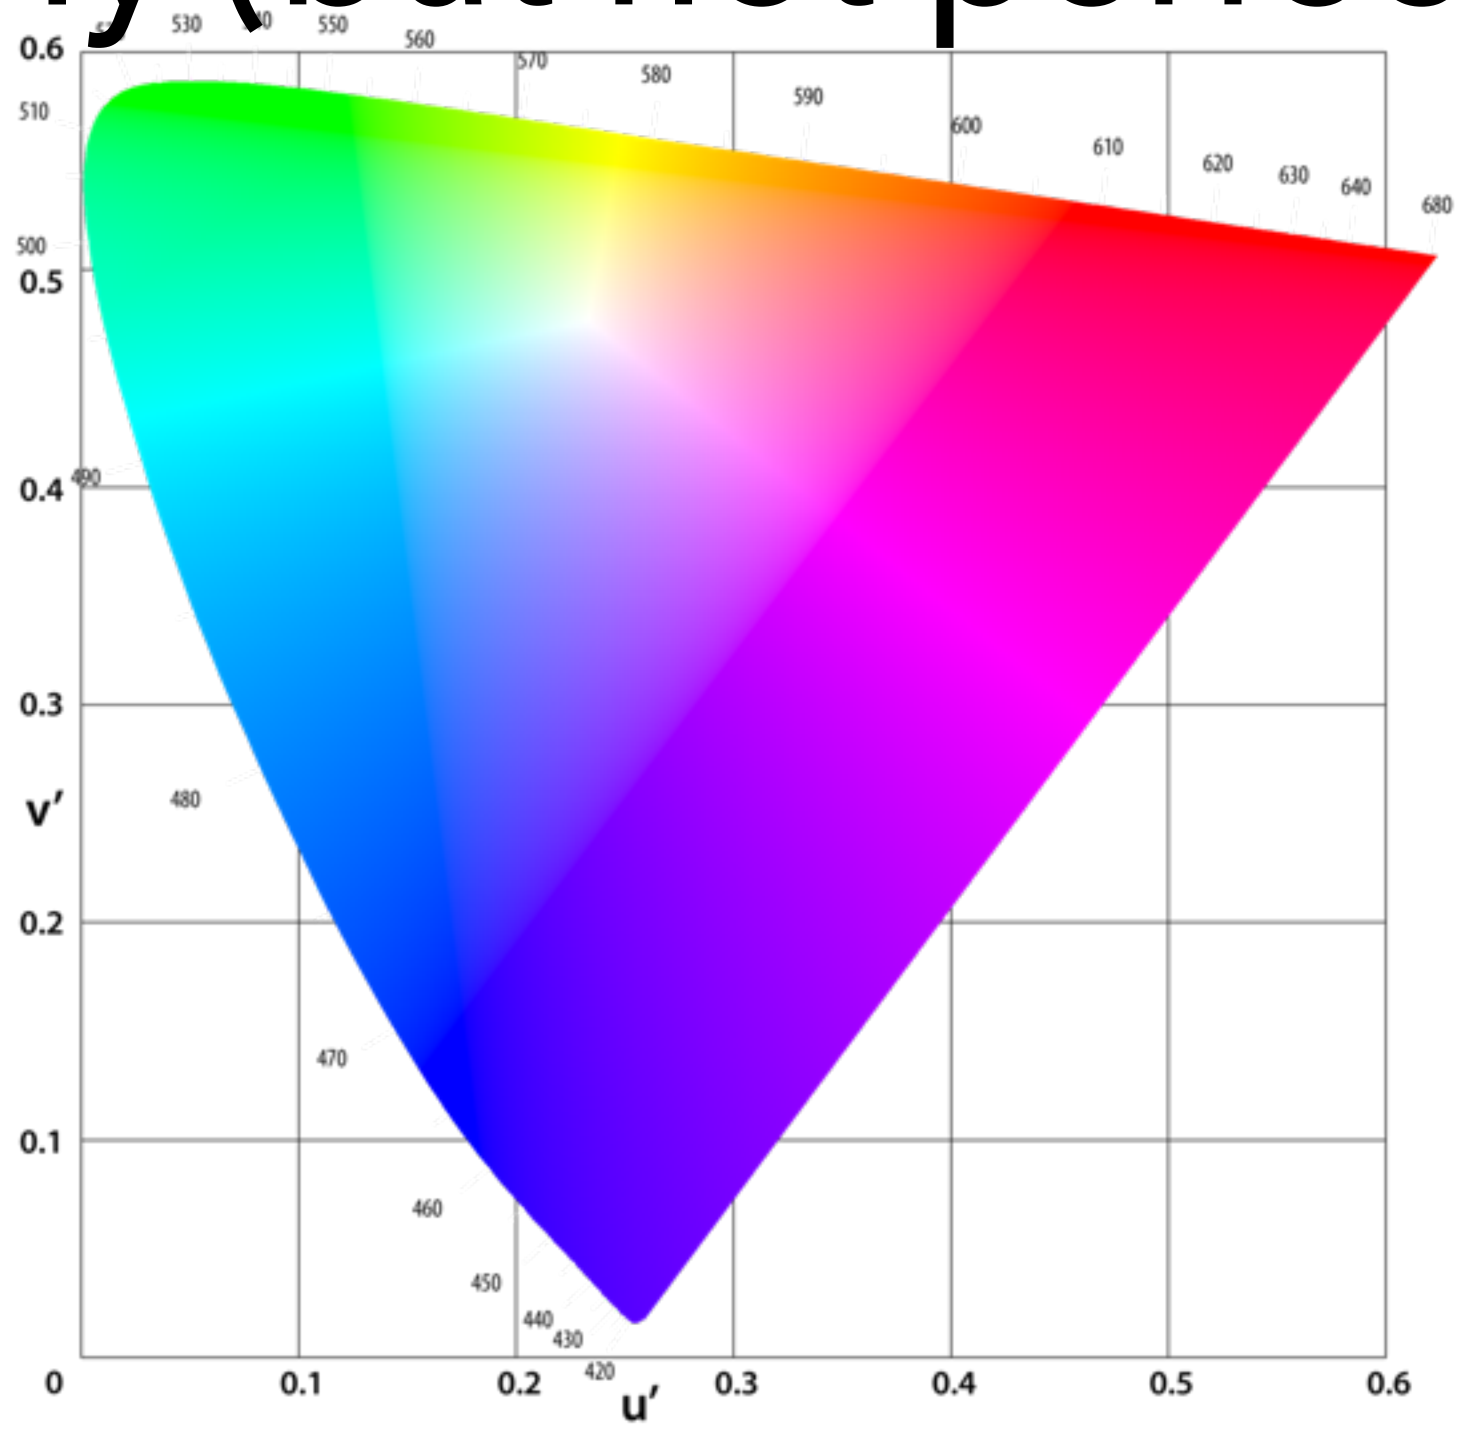

Higher-level color perception

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

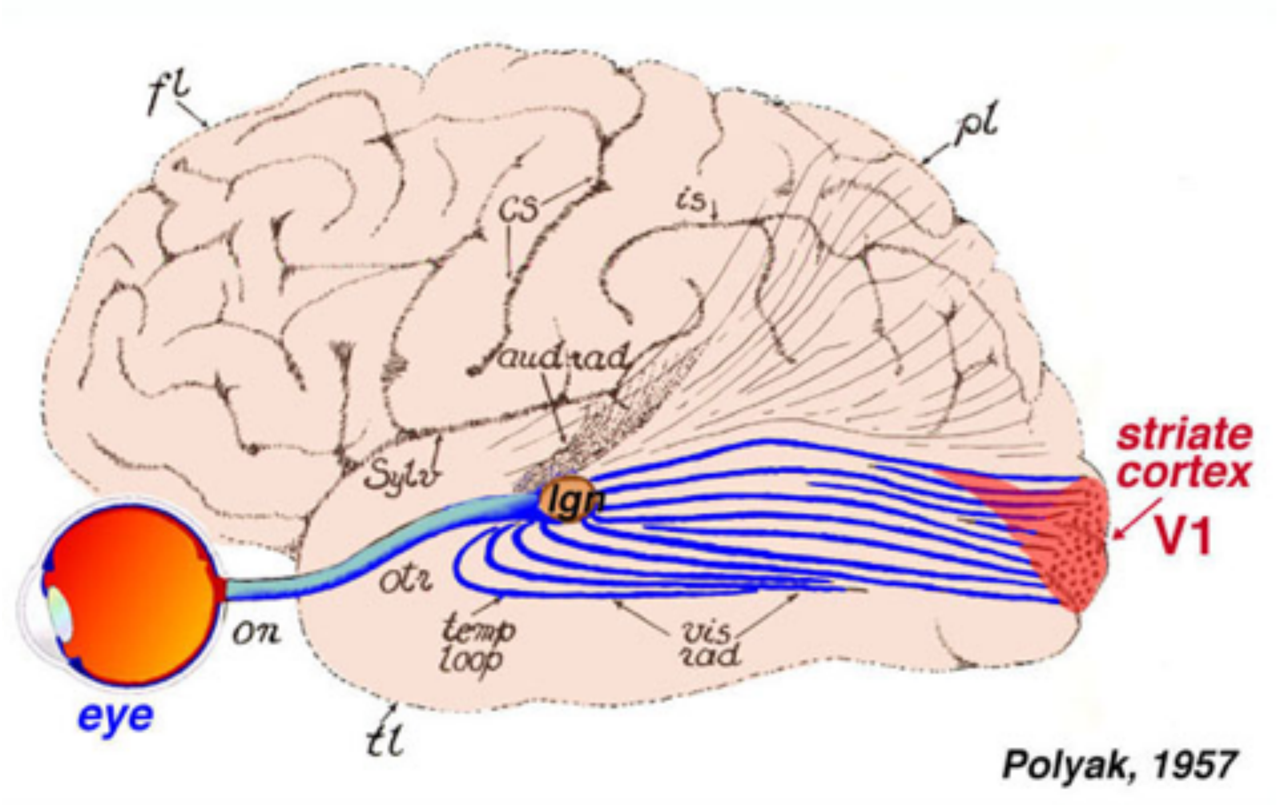

Color constancy

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

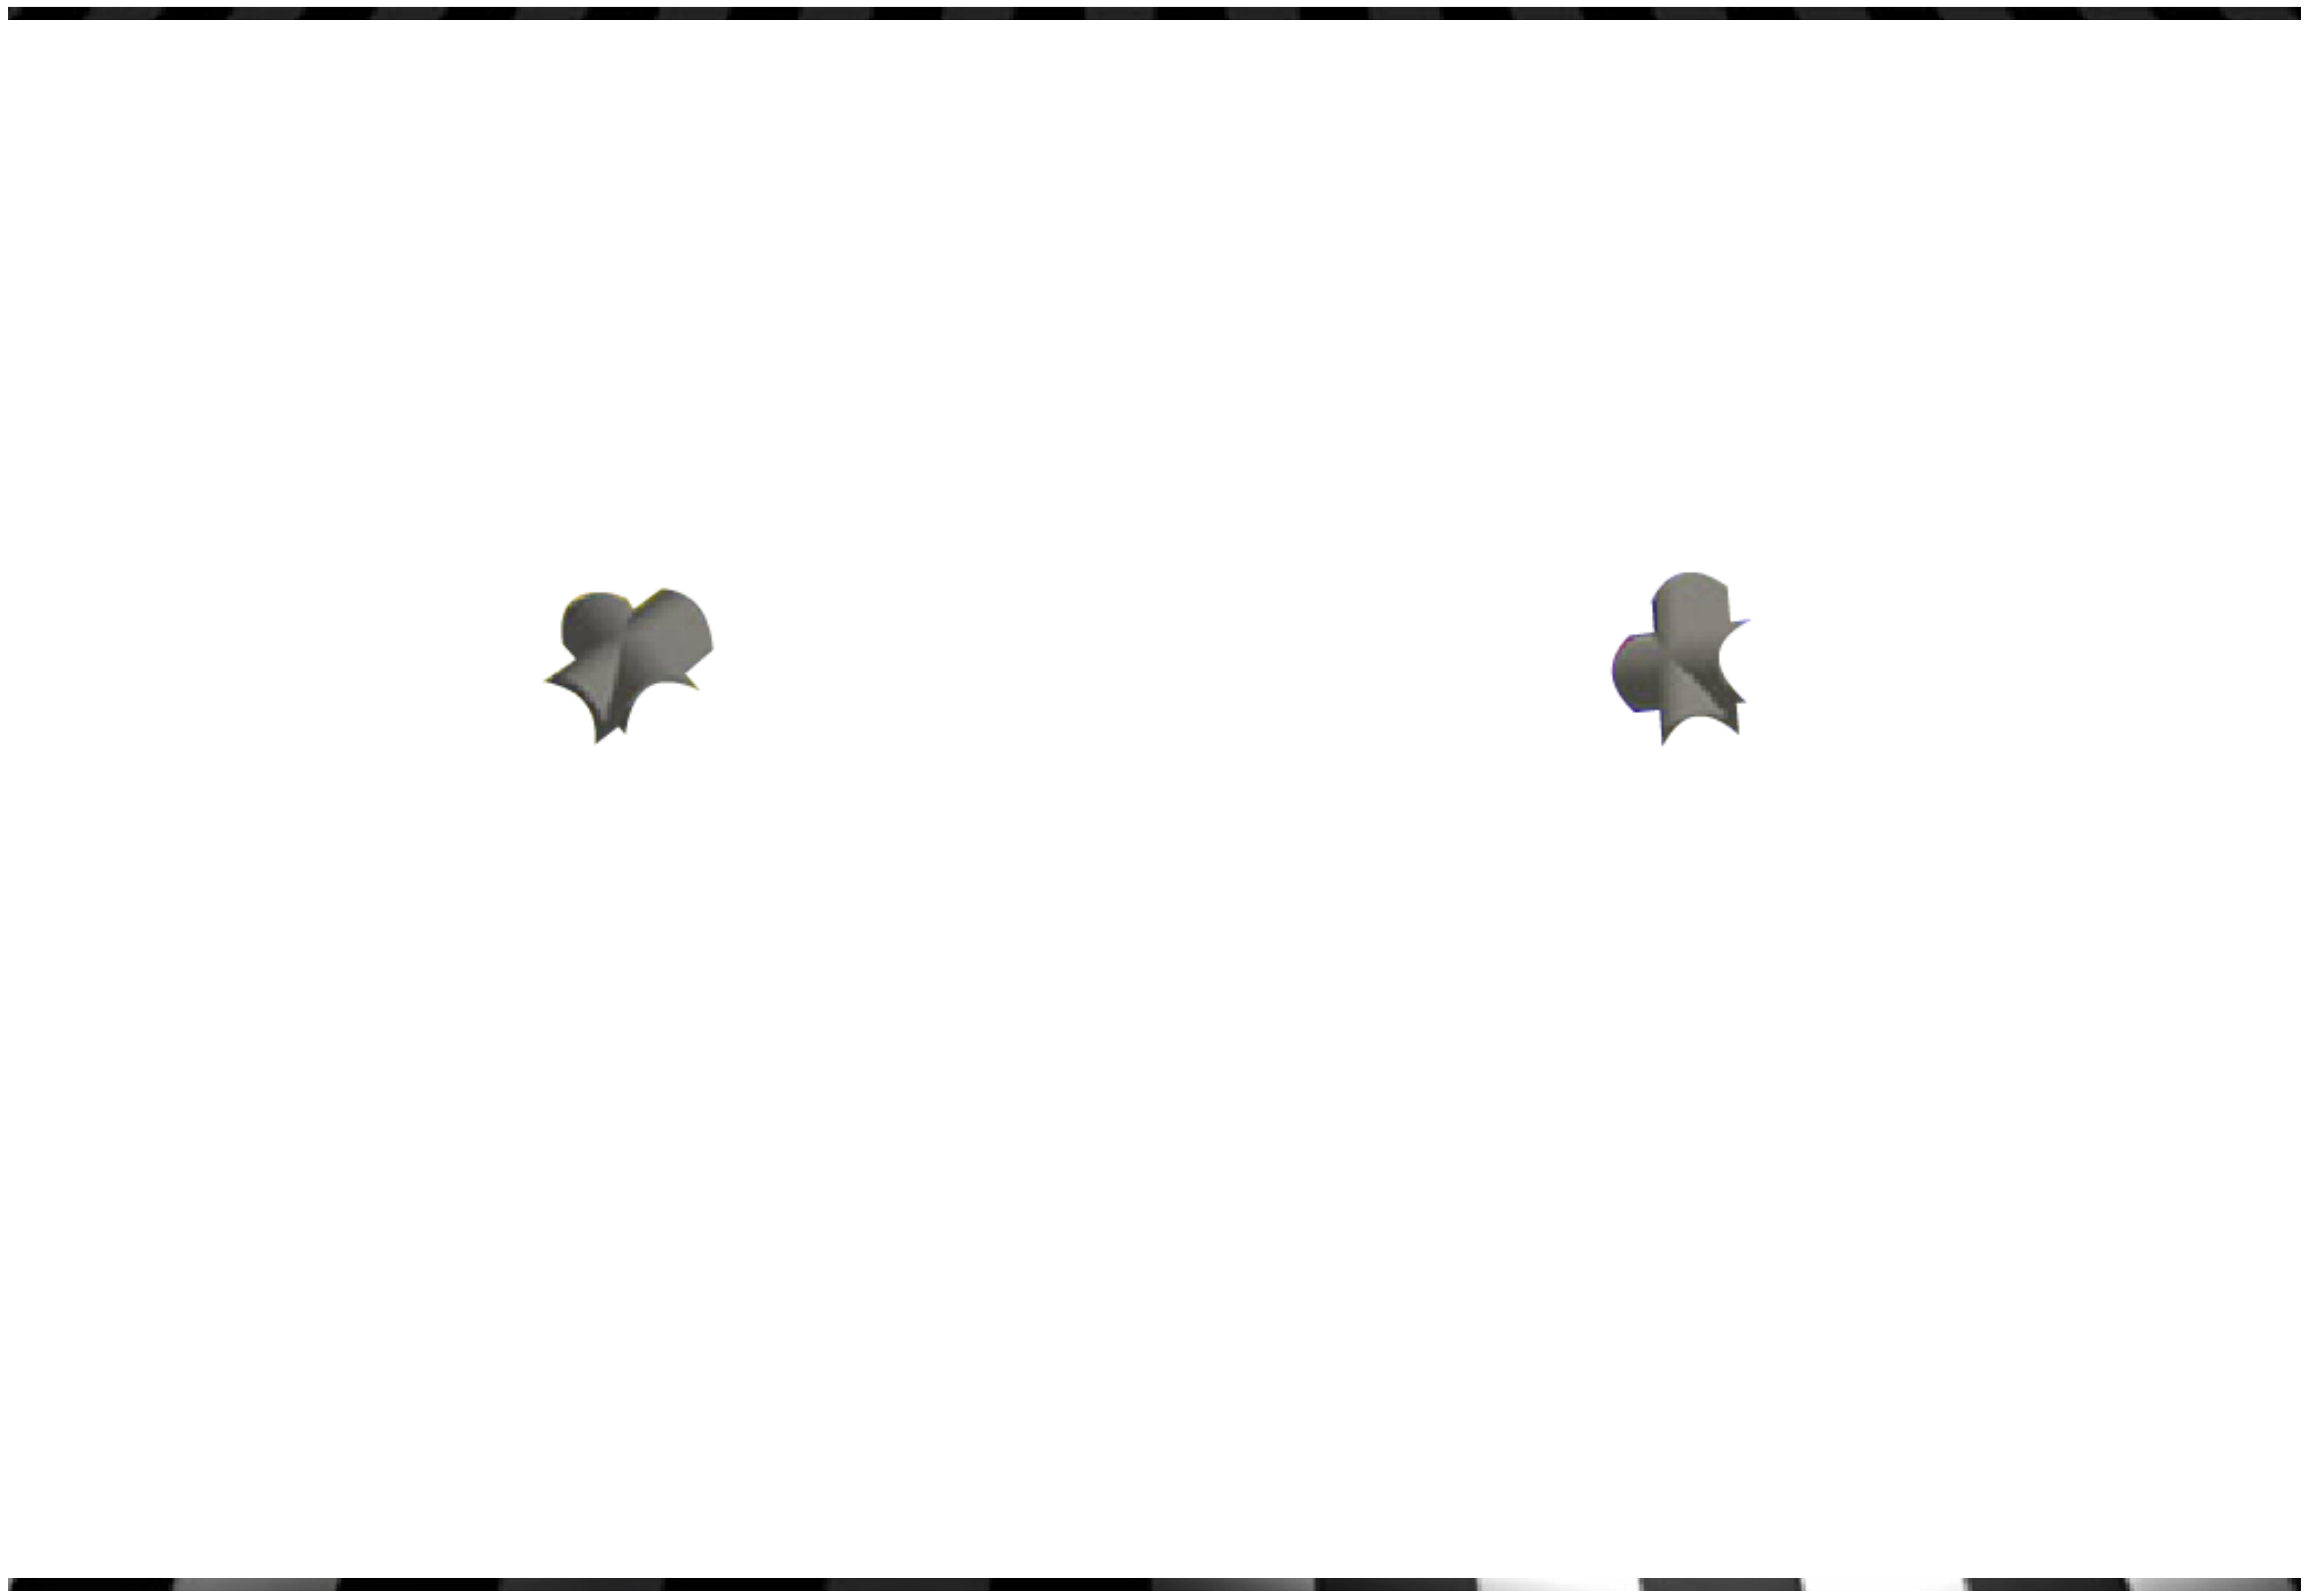

blue and black?

or

white and gold?

Color constancy failure

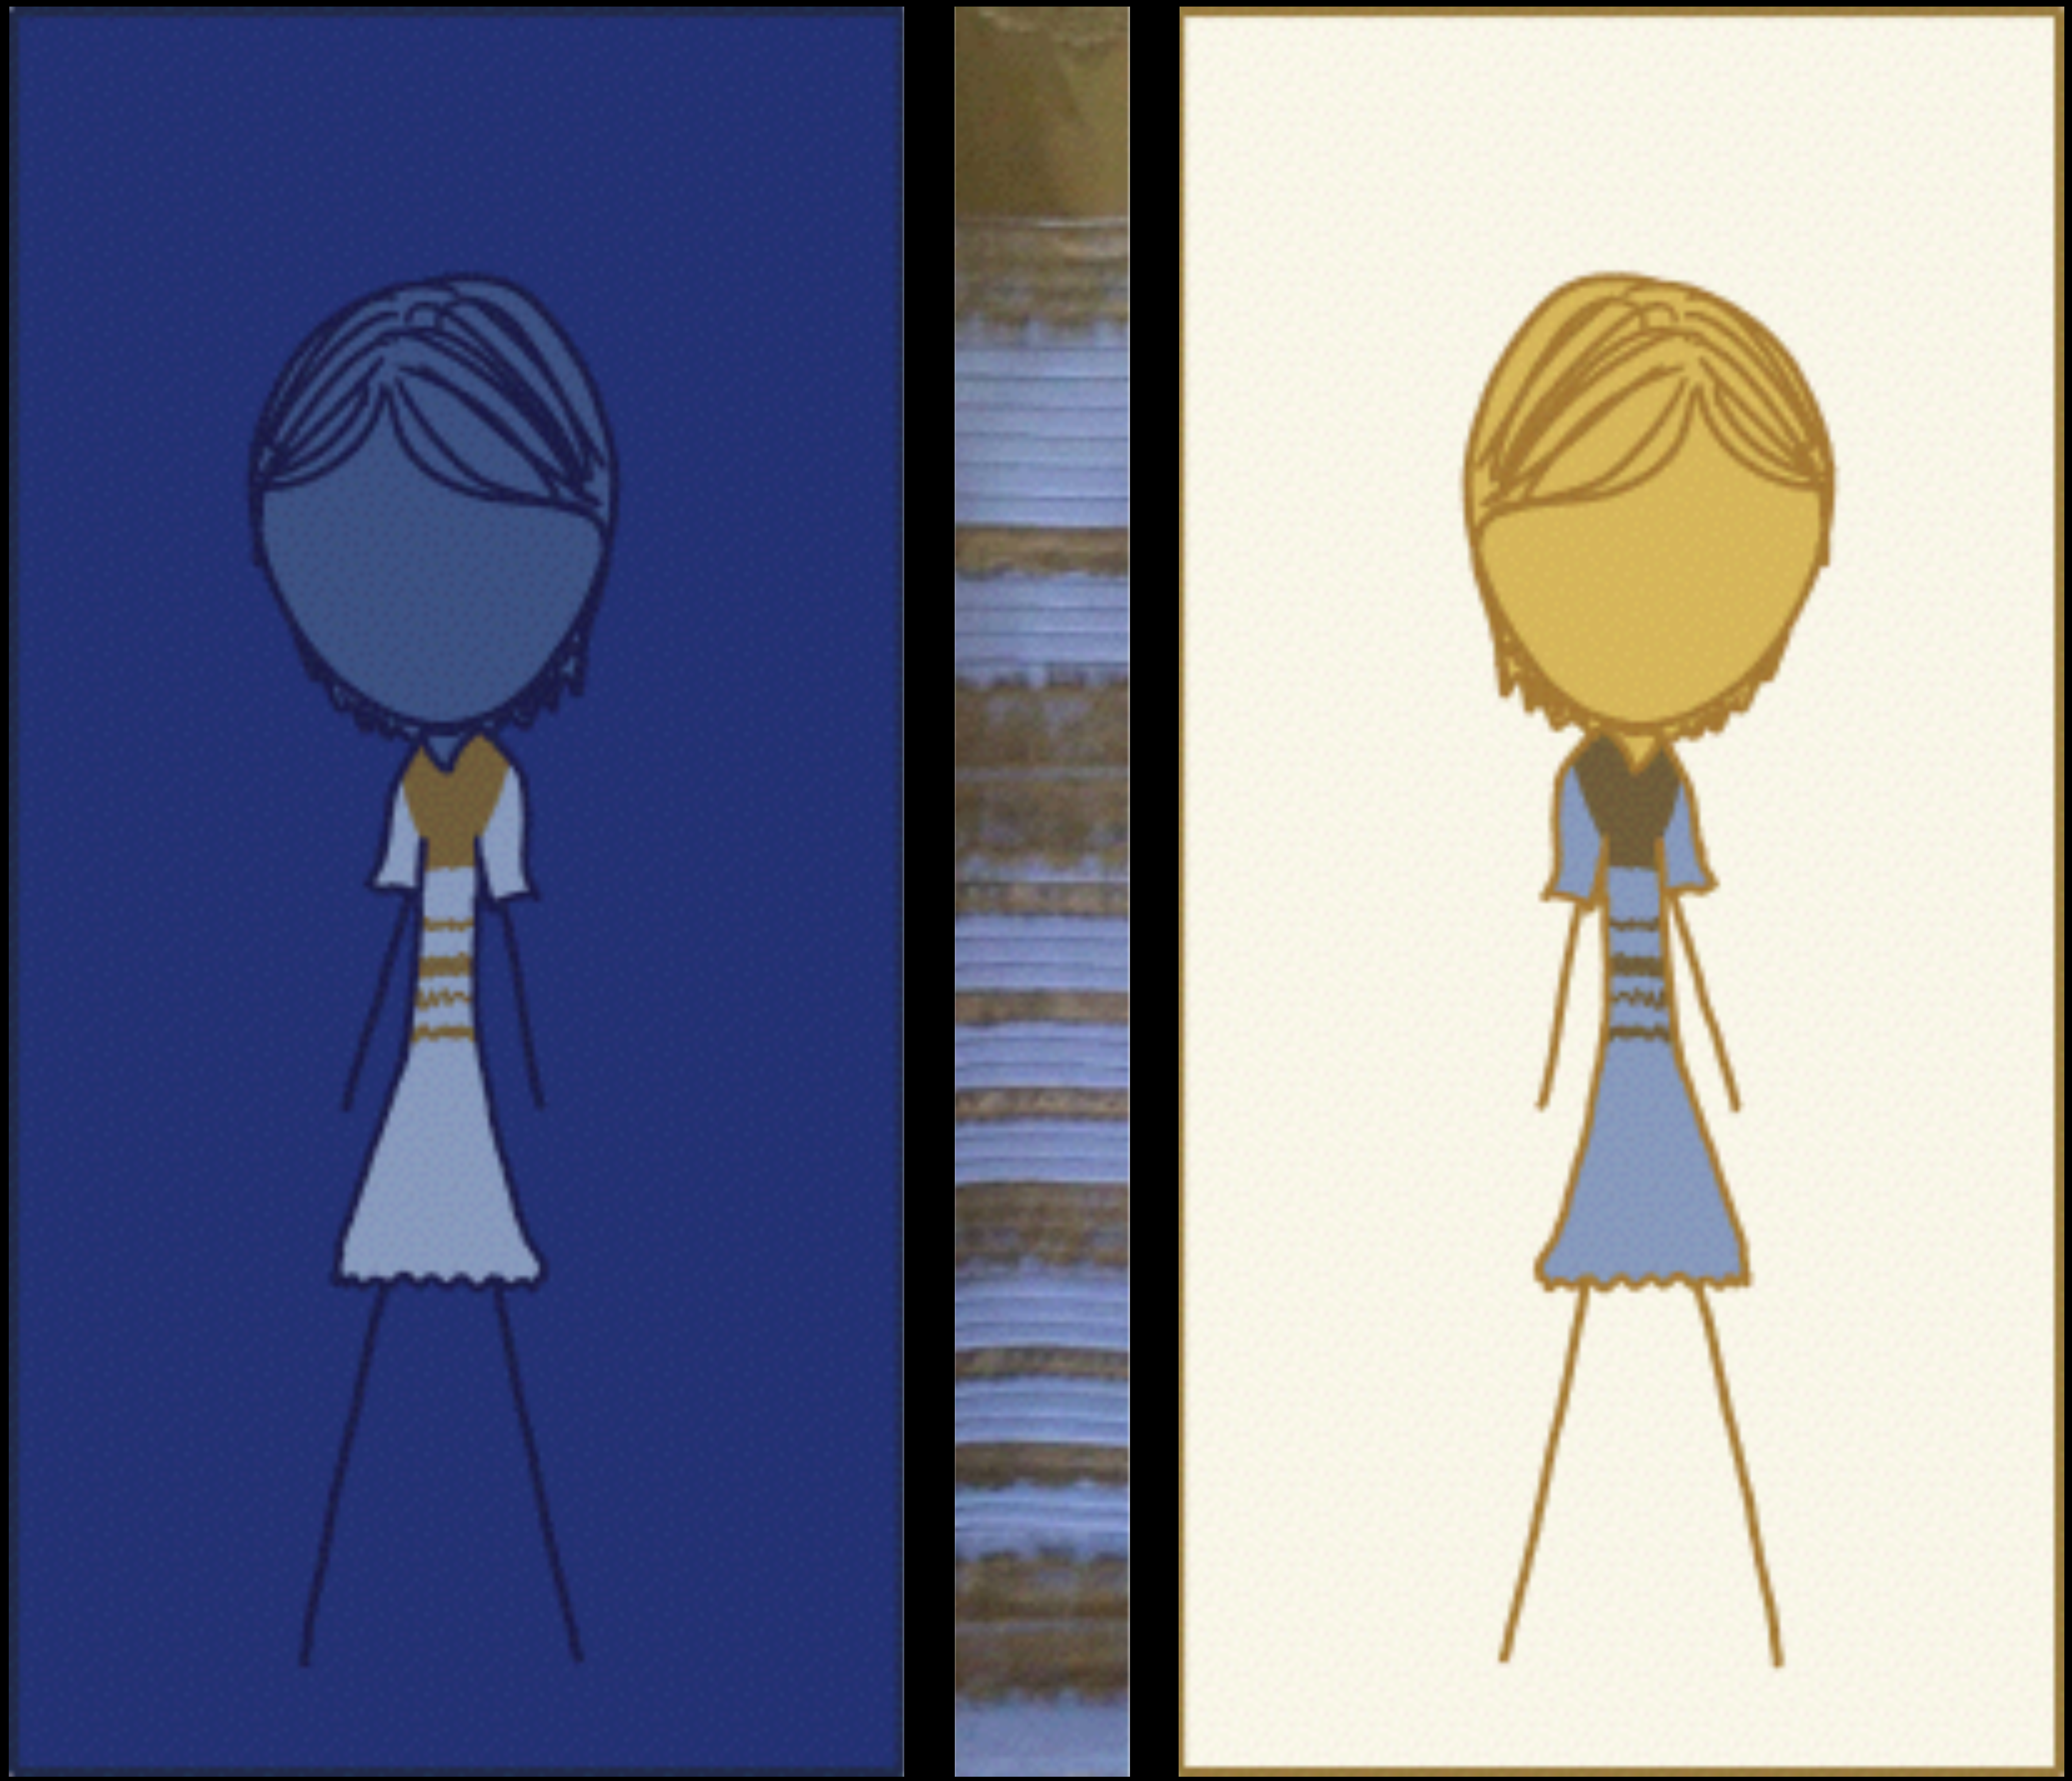

http://xkcd.com/1492/

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

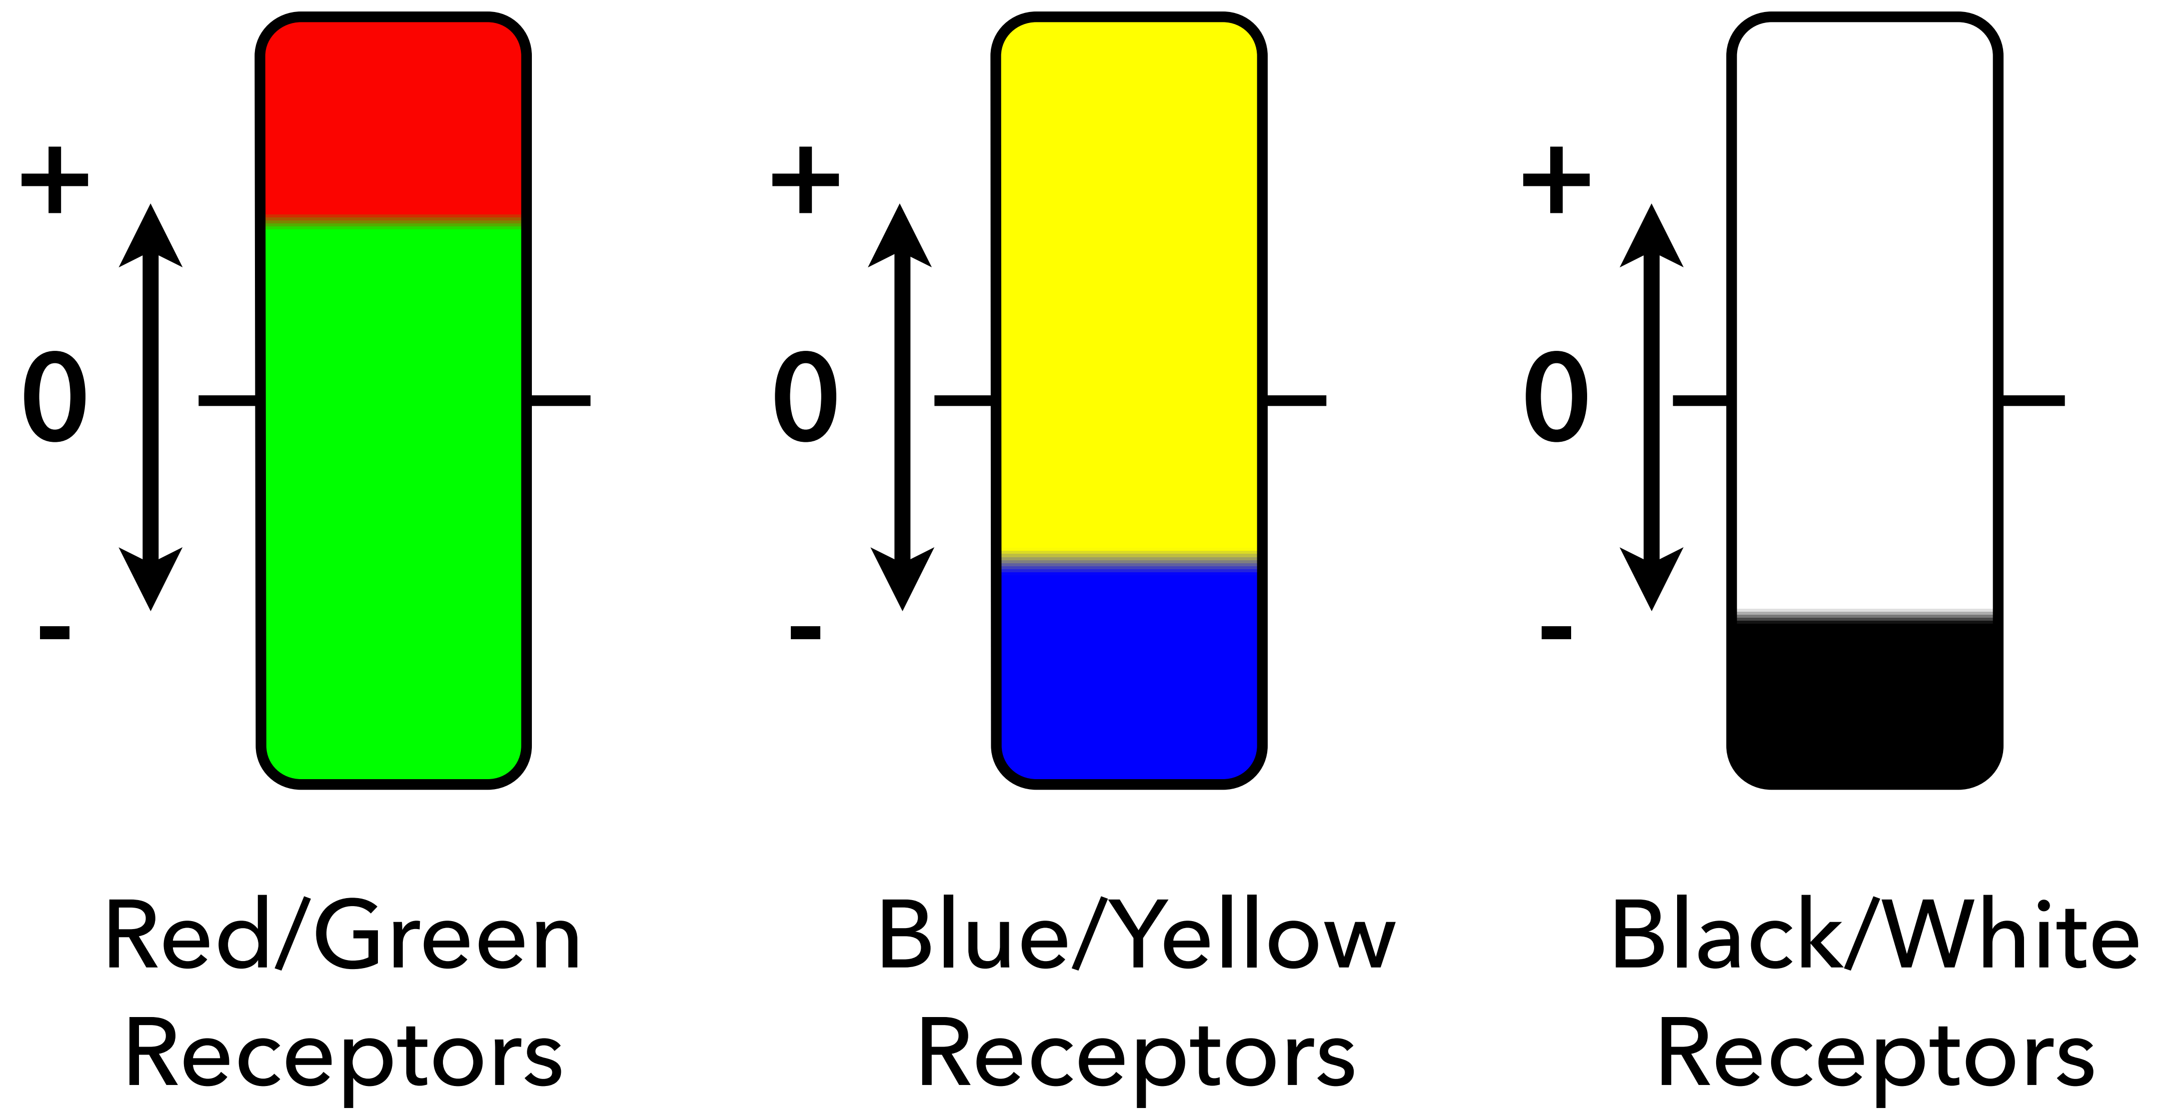

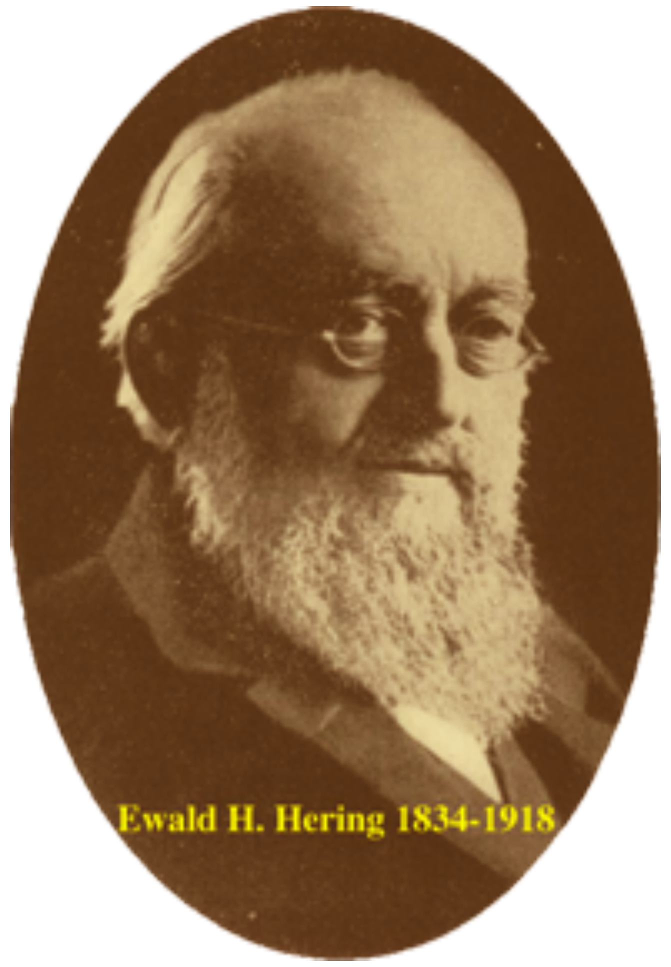

Dual process theory

Input is LMS

Output has a different parameterization:

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

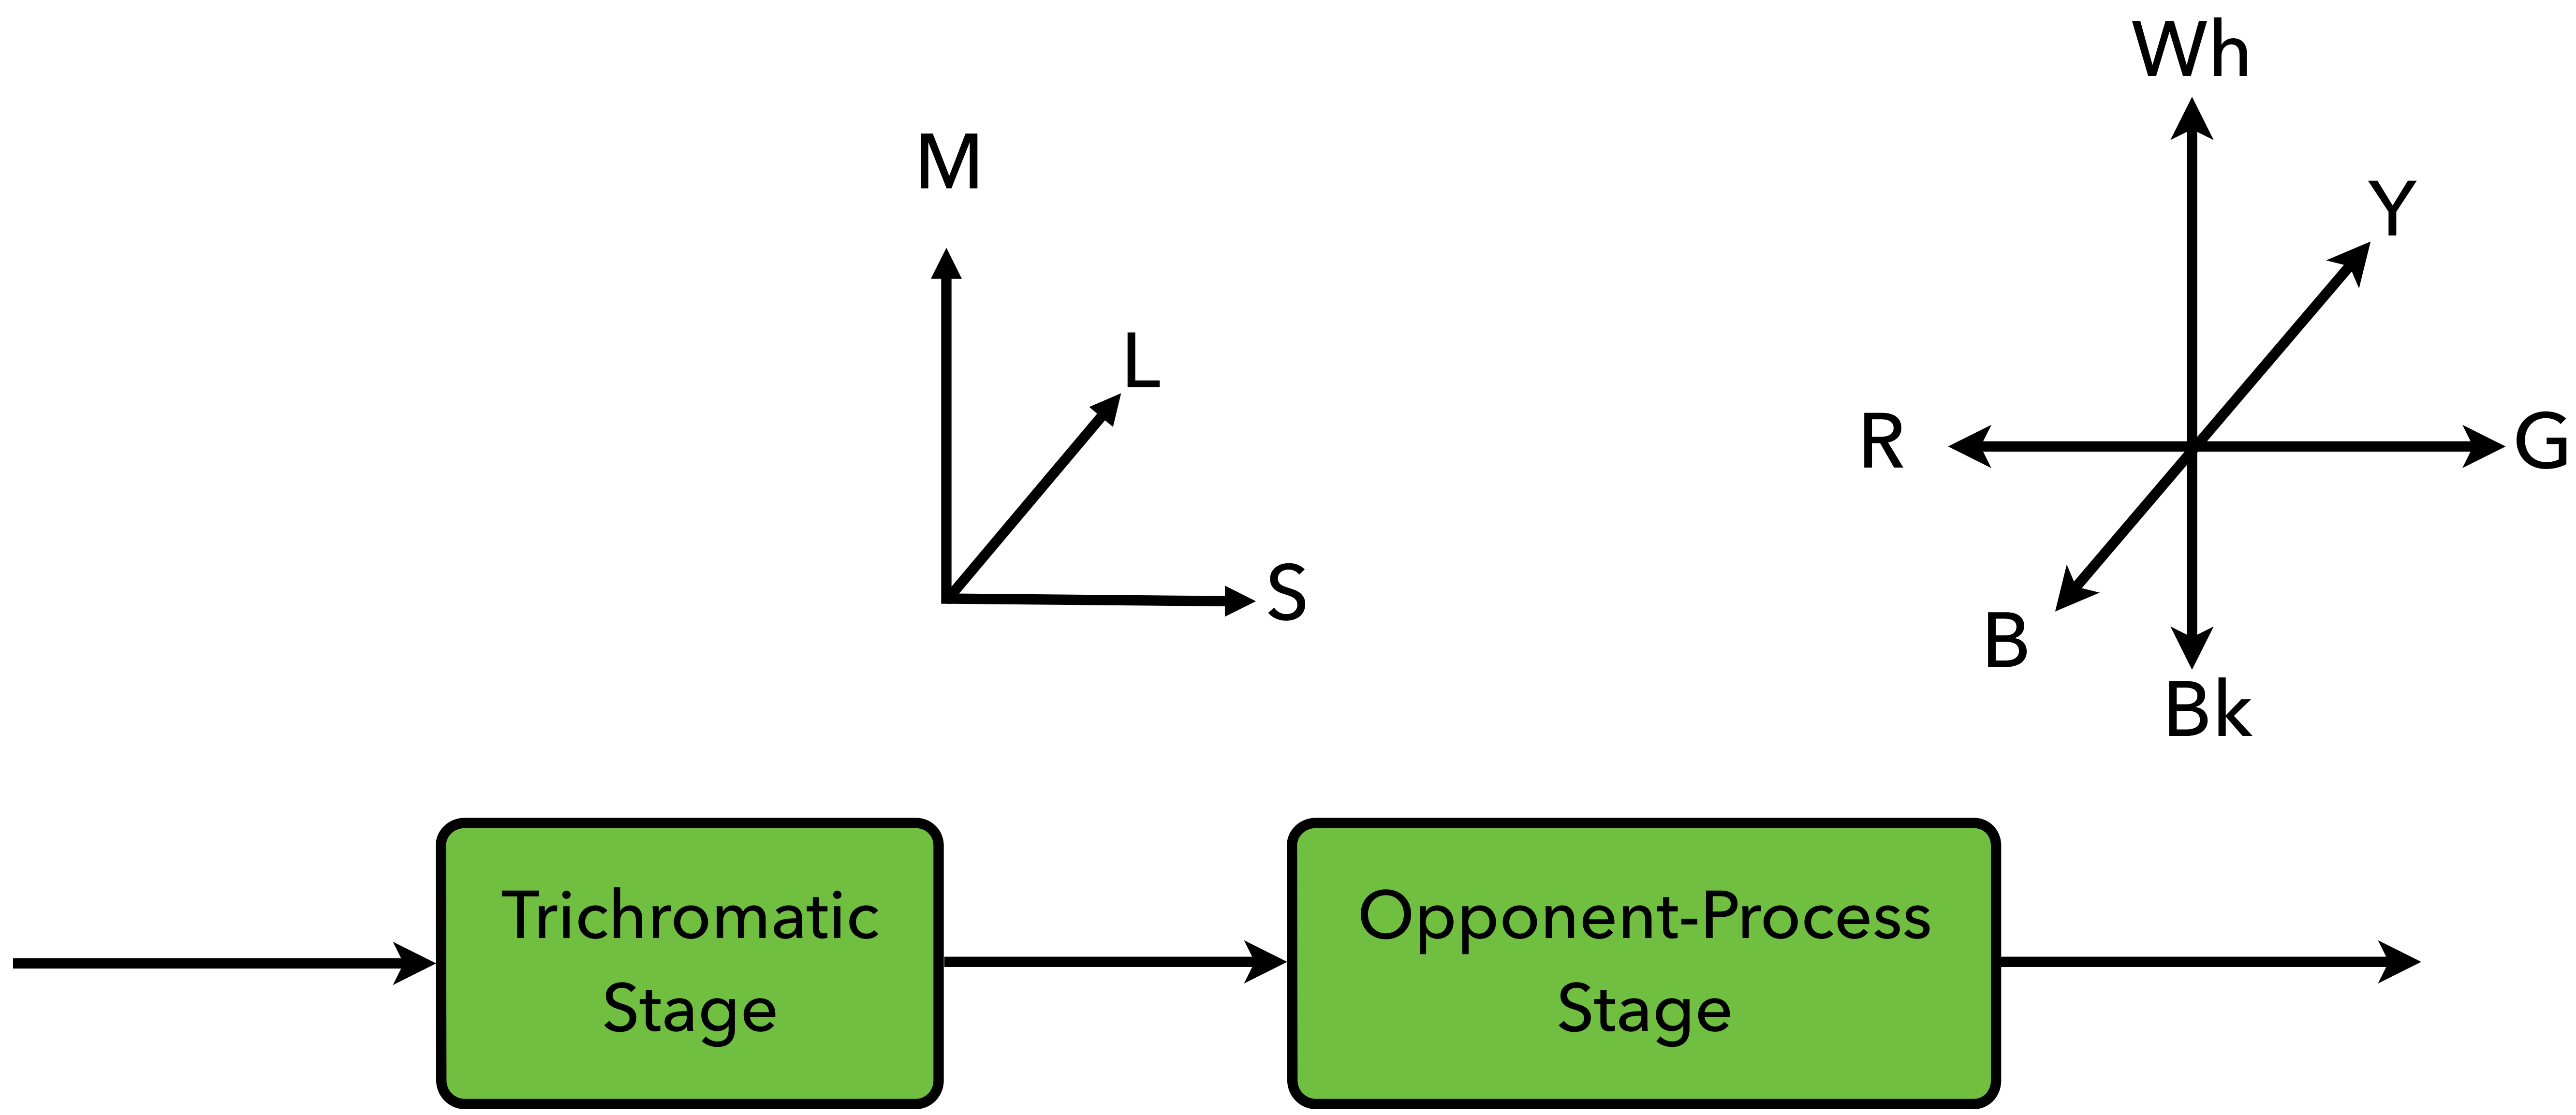

Color opponents wiring

Sums for brightness

Differences for color opponents

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

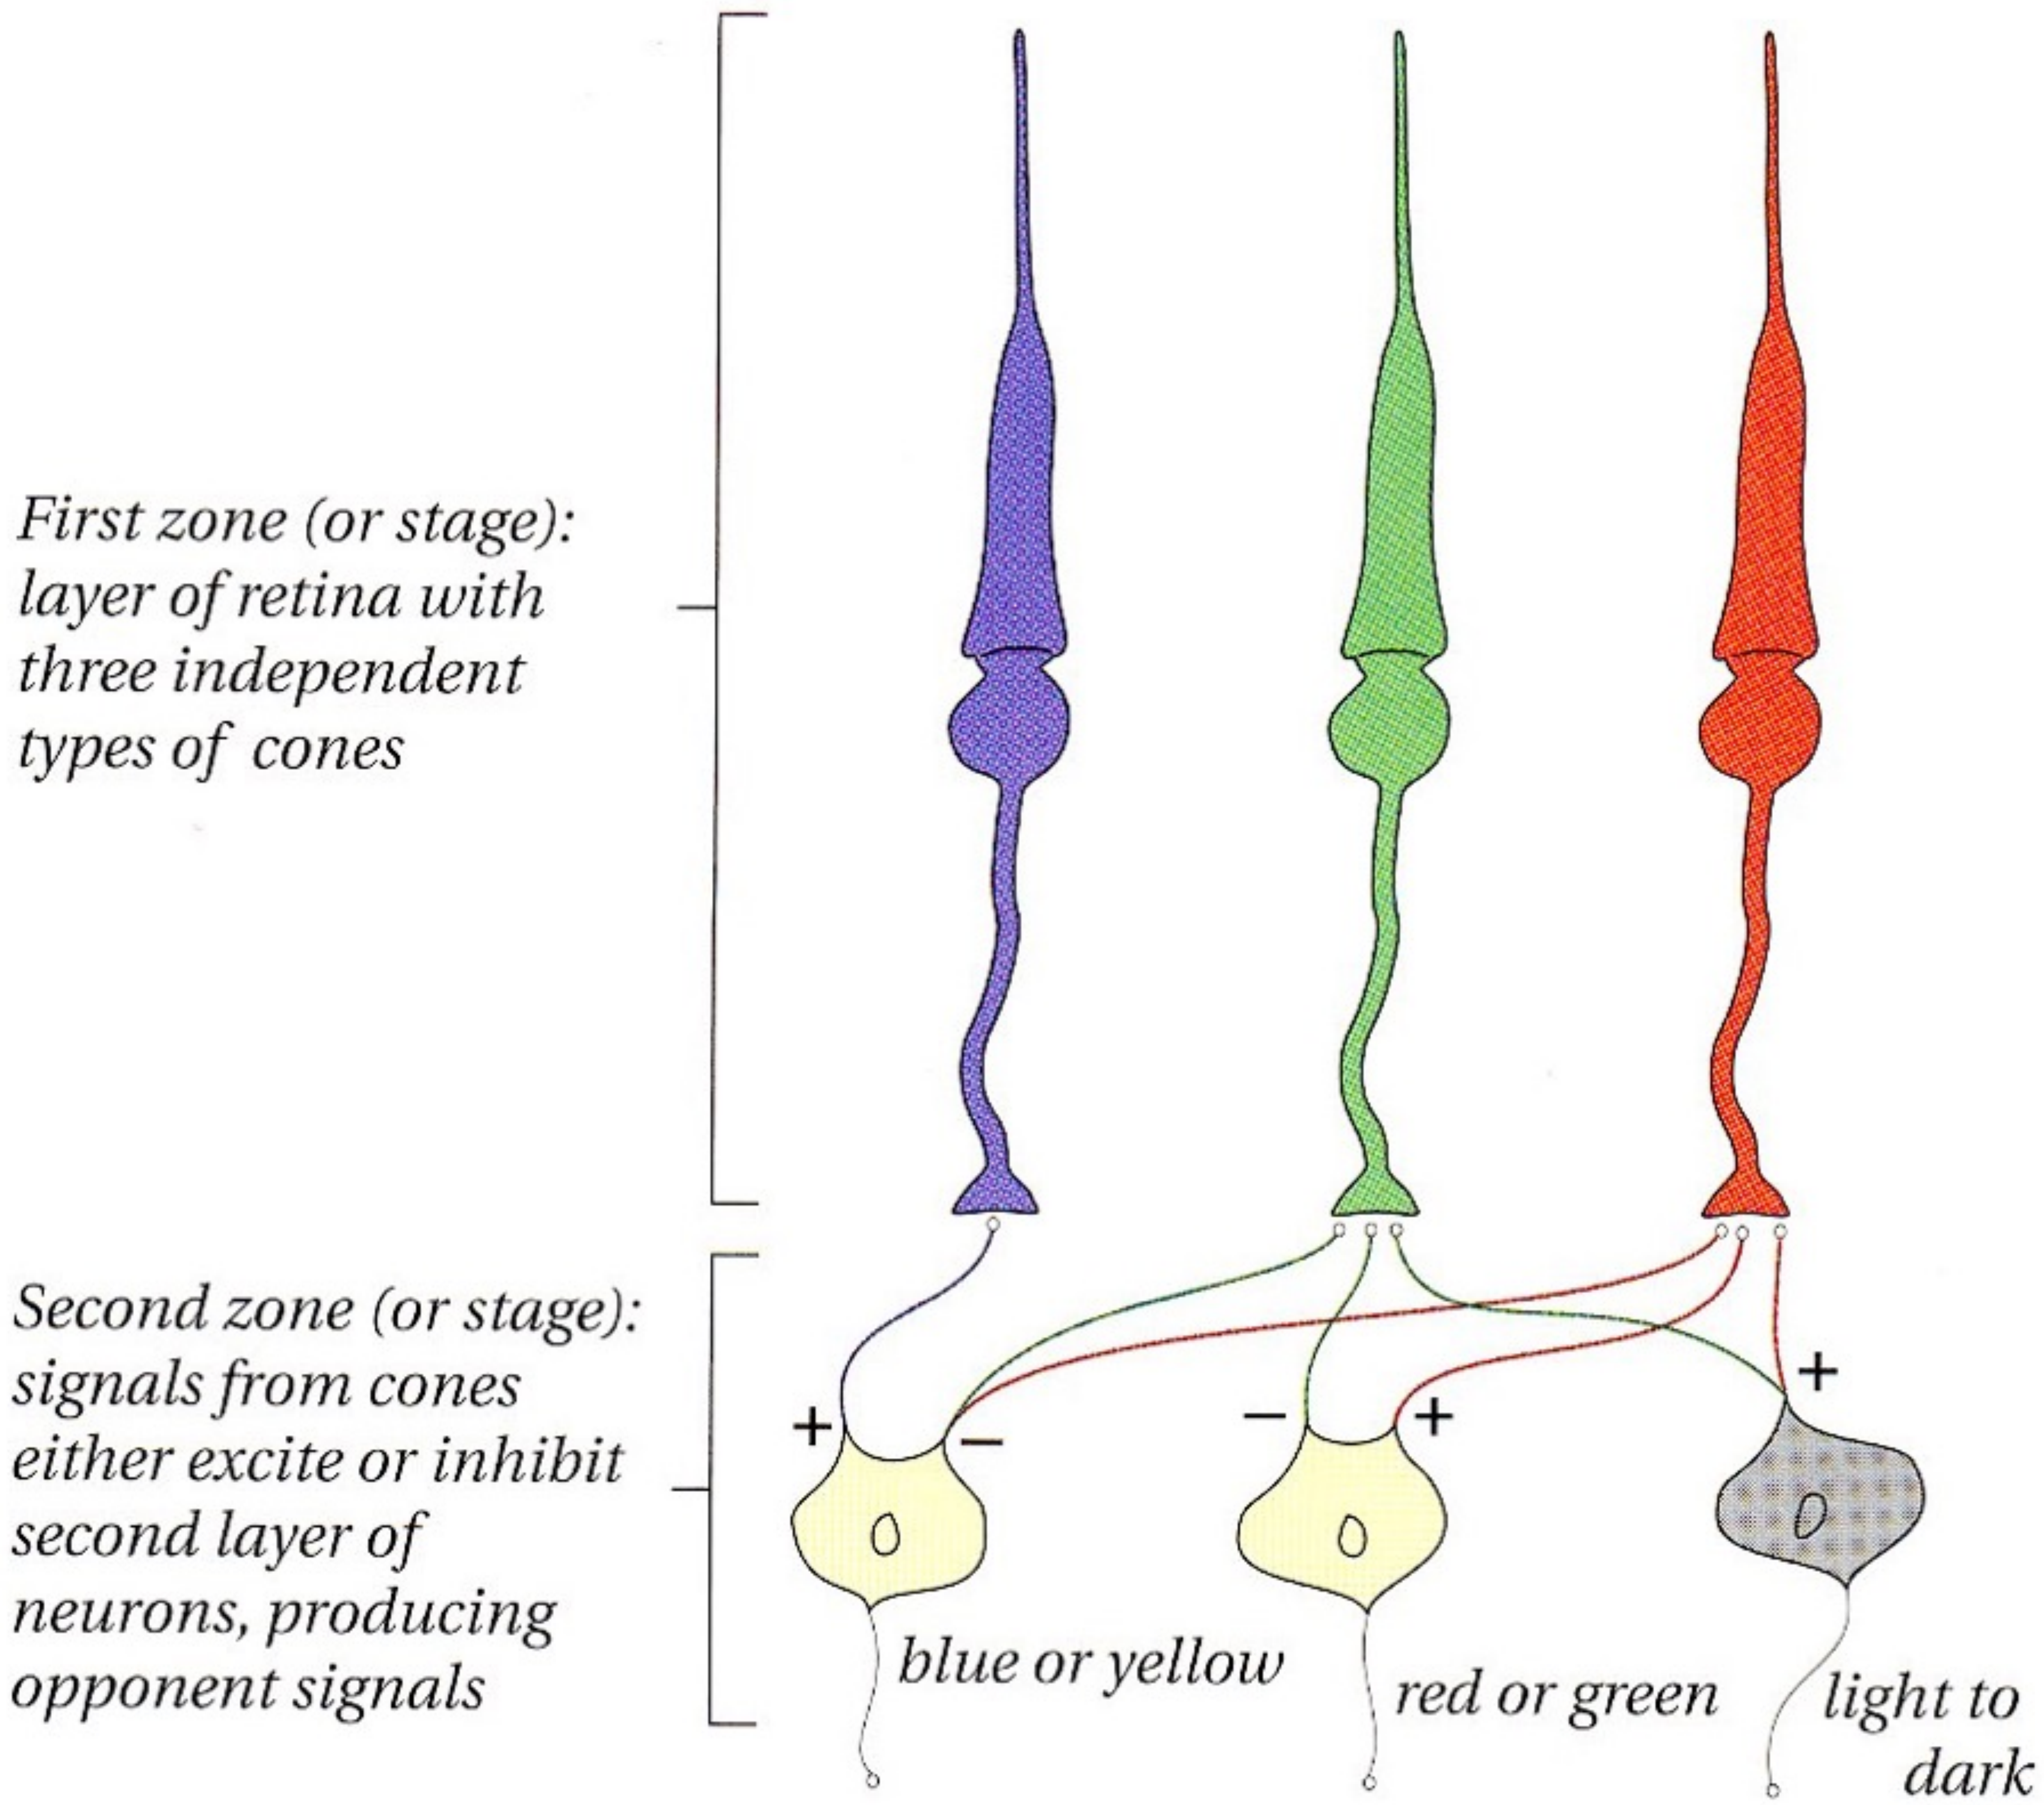

Opponent Colors

Image | Afterimage

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

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

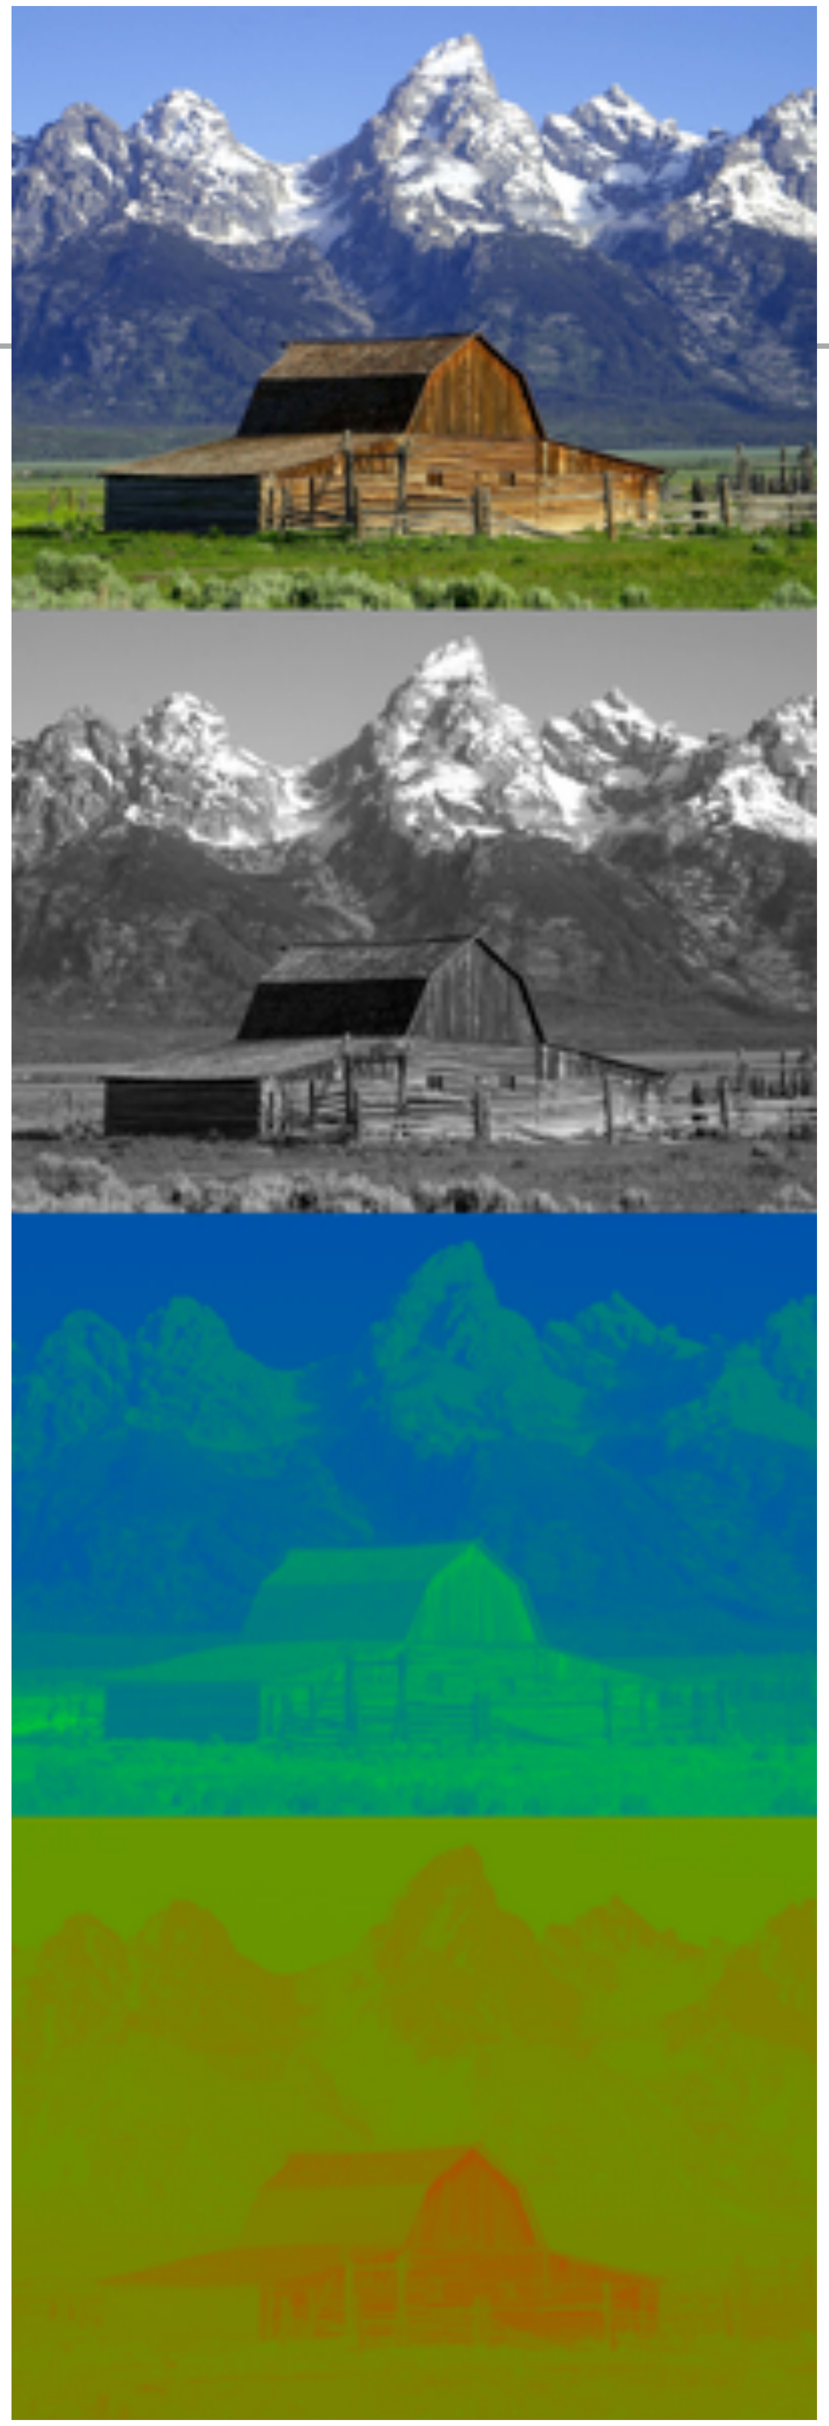

YUV

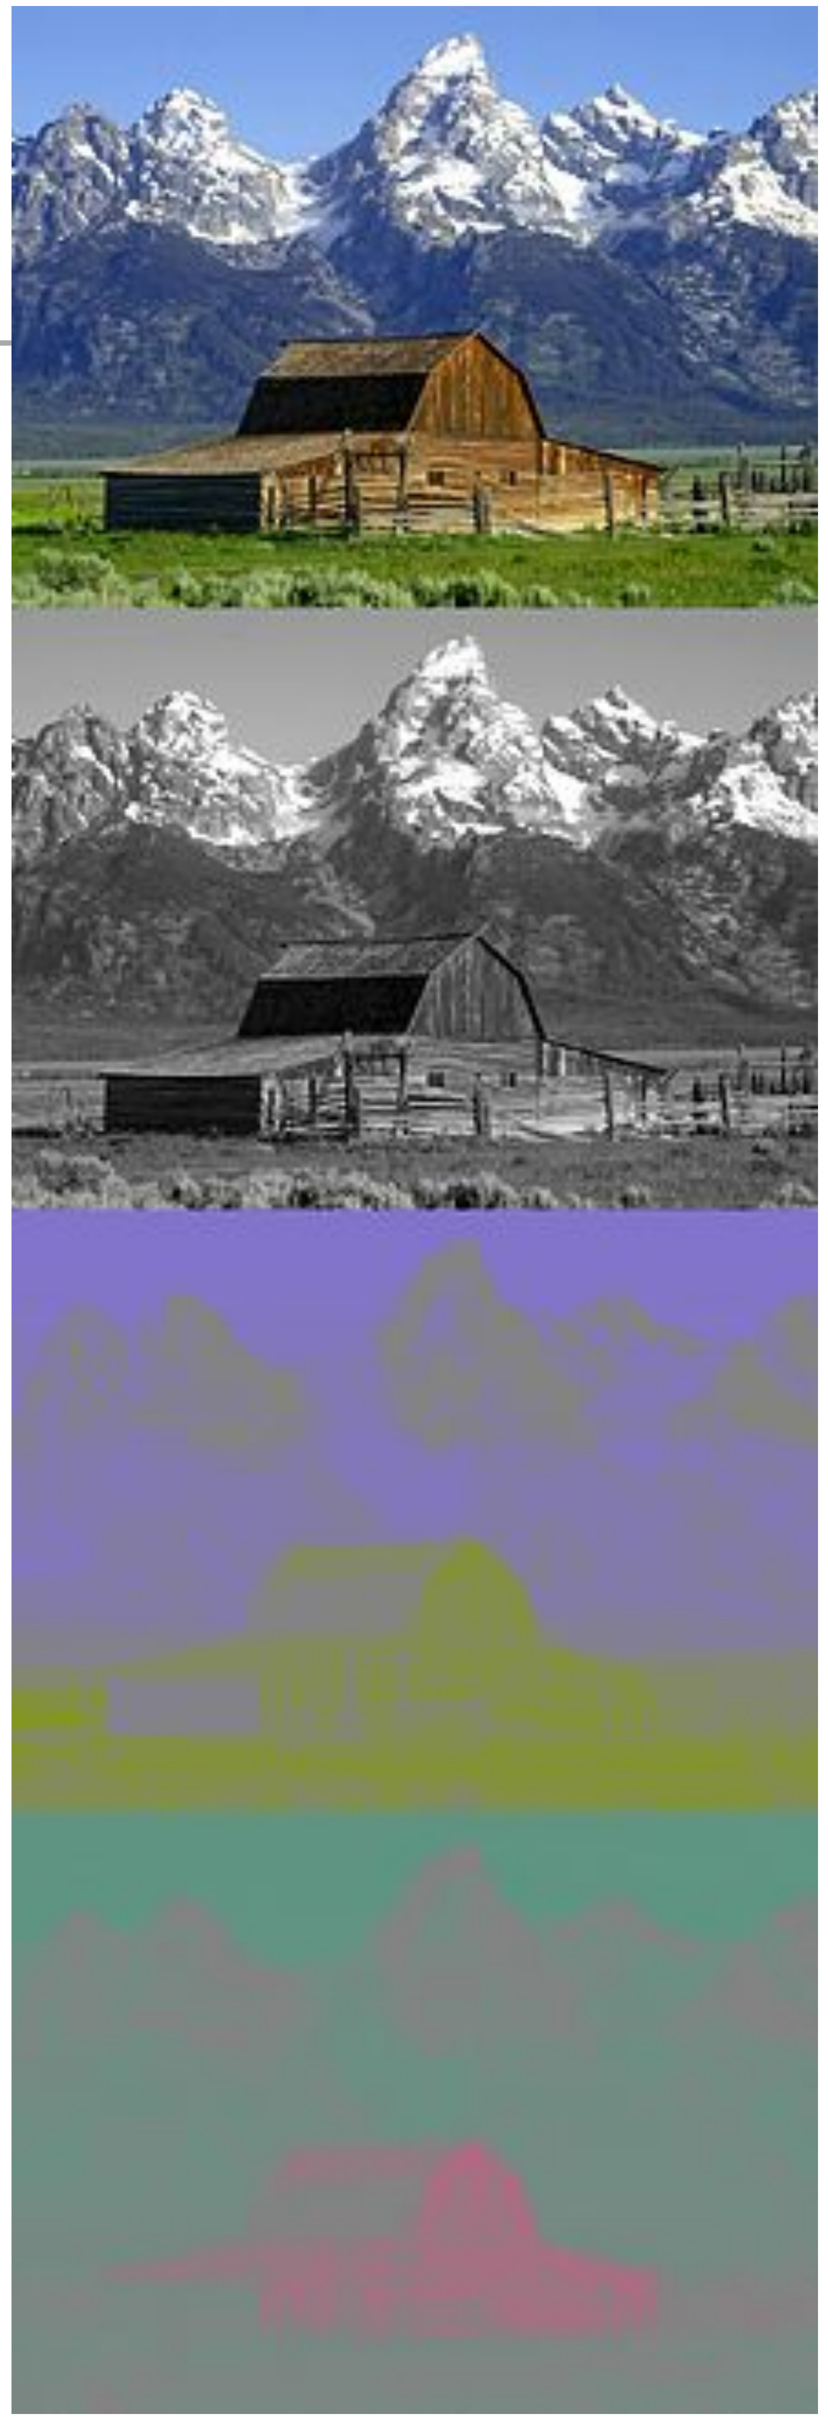

YCrCb

Slide credits

Frédo Durand

Steve Marschner

Matthias Zwicker