Color

Computer Graphics
CMU 15-462/15-662

Why do we need to be able to talk precisely about color?


"An easy choice"
WIN
CONIA
SONY
LG
$1,199

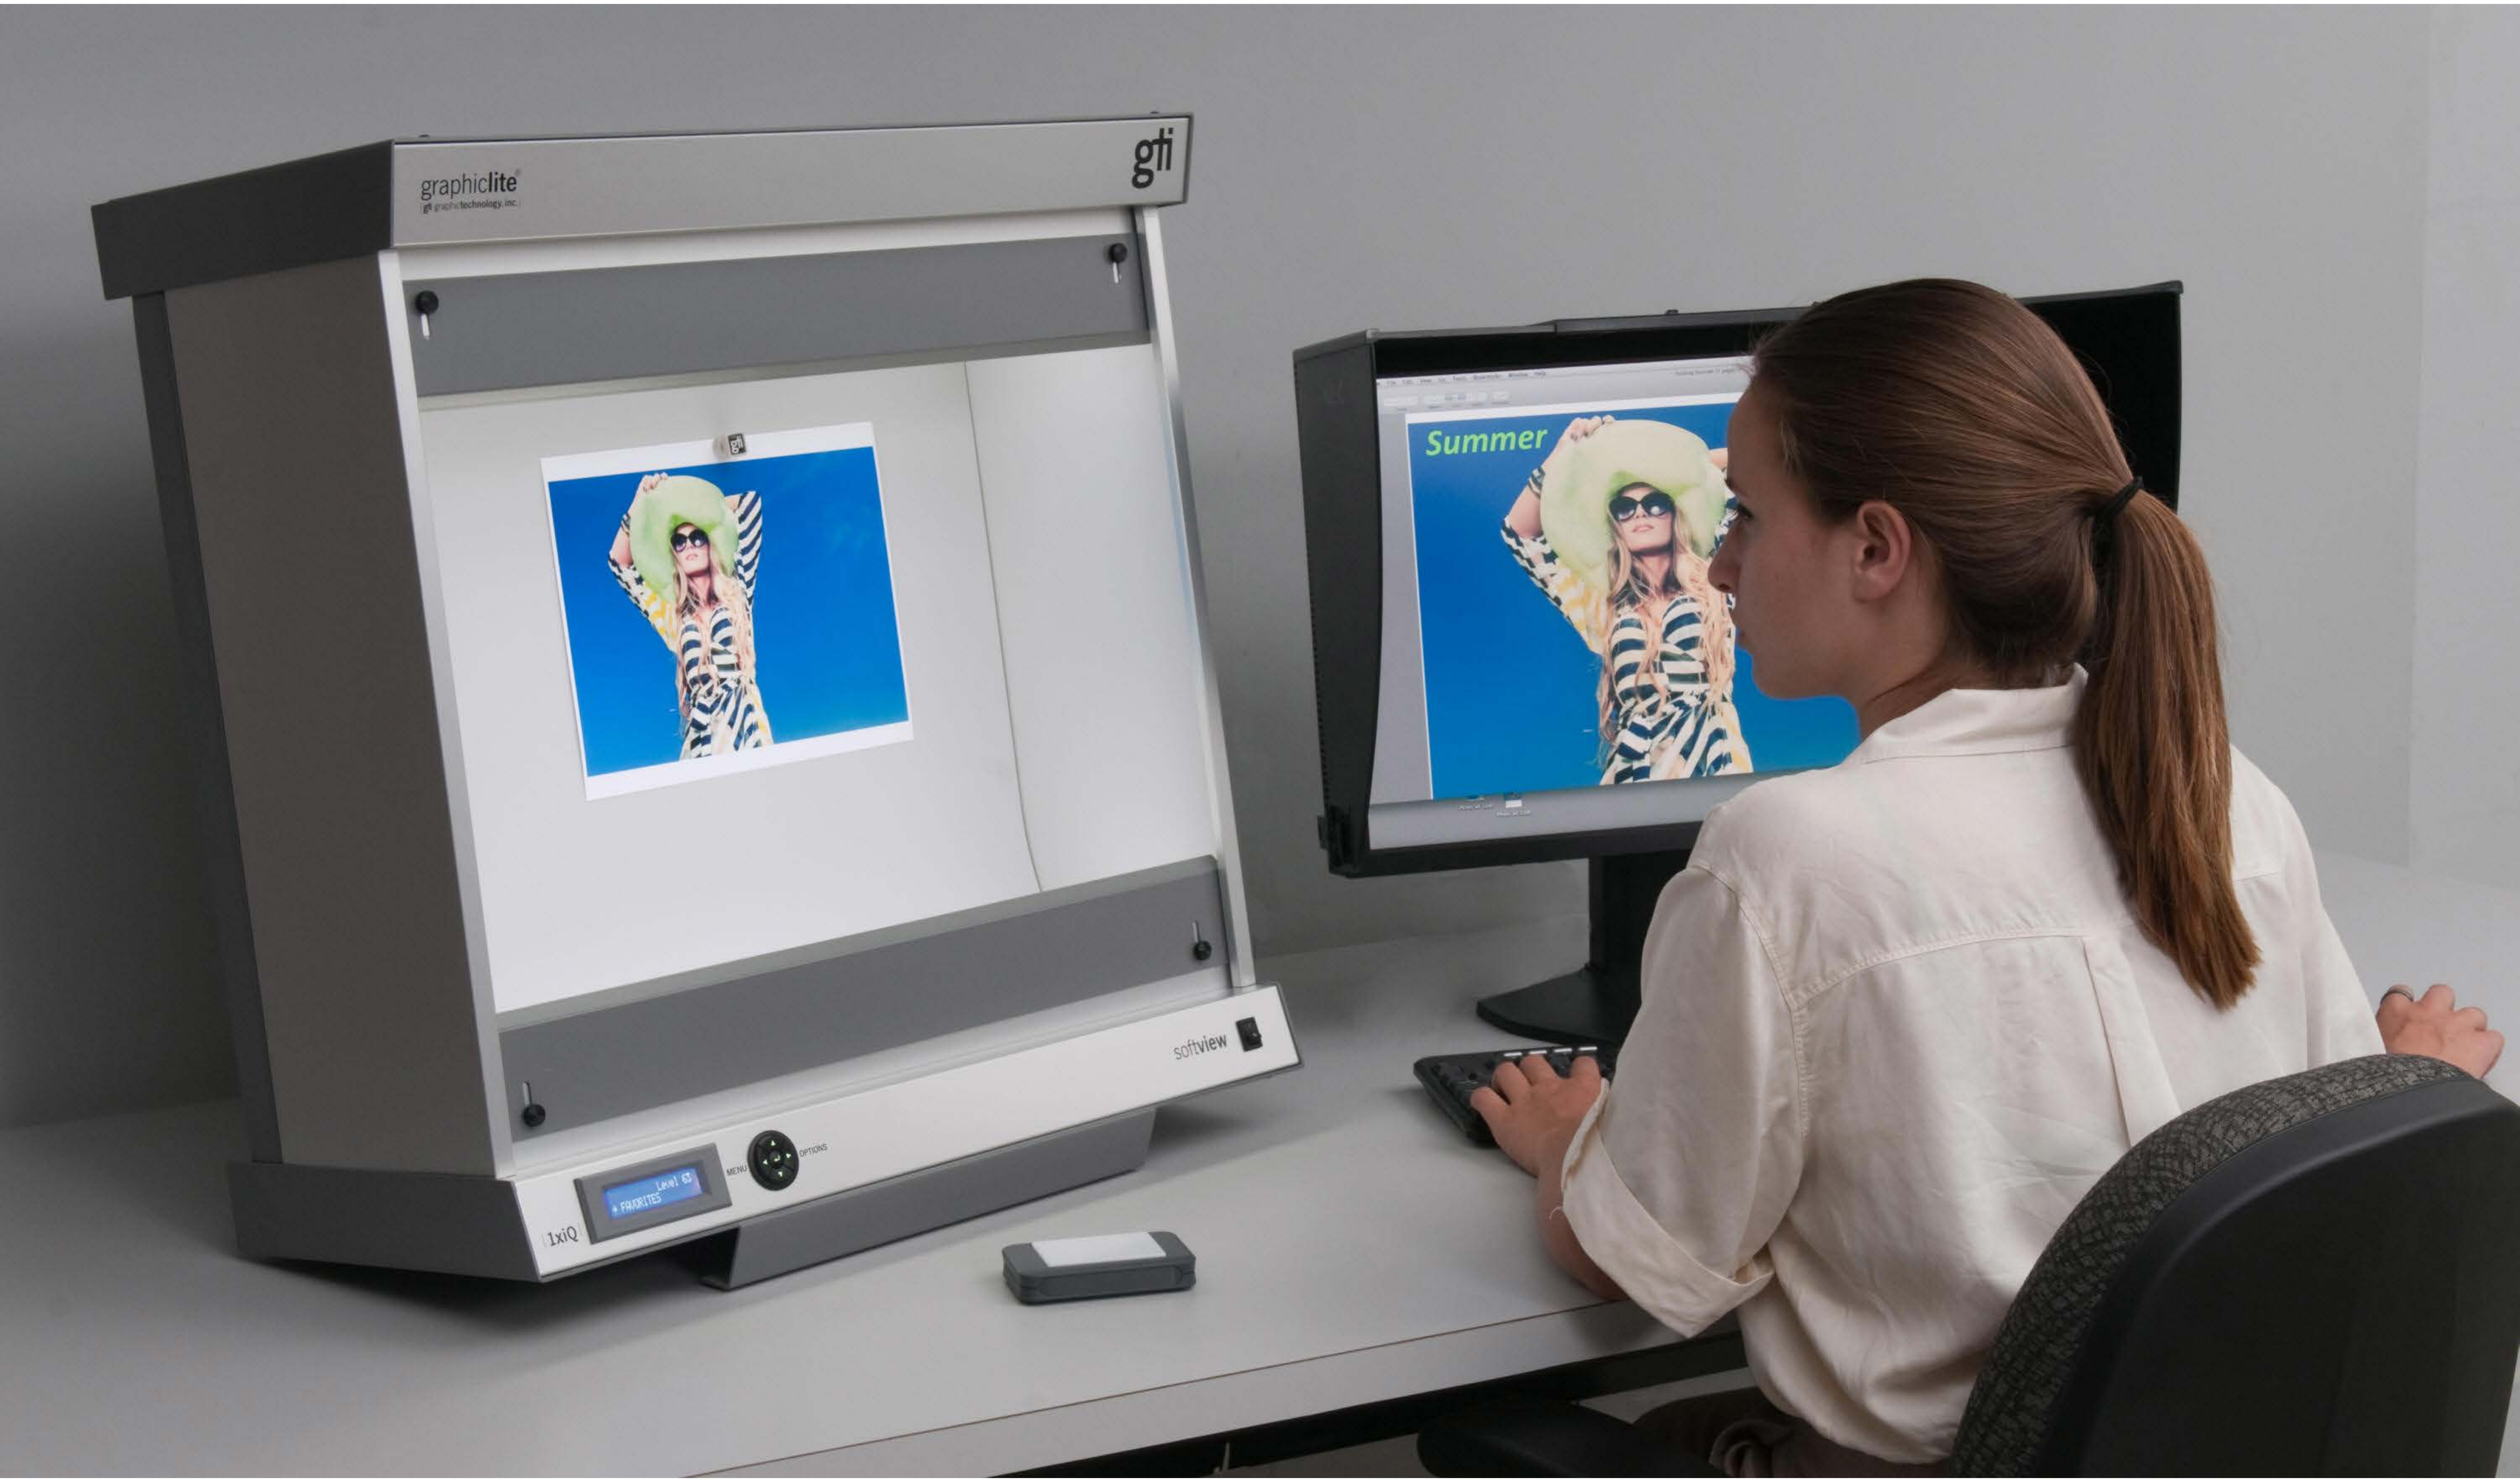
graphiclite
gti graphictechnology, inc.
gti
softview
Level 63
+ FAVORITES
MENU
OPTIONS
1xiQ
Summer

on screen

printed

Zhangye Danxia Geological Park, China

Credit: http://parade.com/63549/linzlowe/where-in-the-world-are-these-incredible-rainbow-mountains

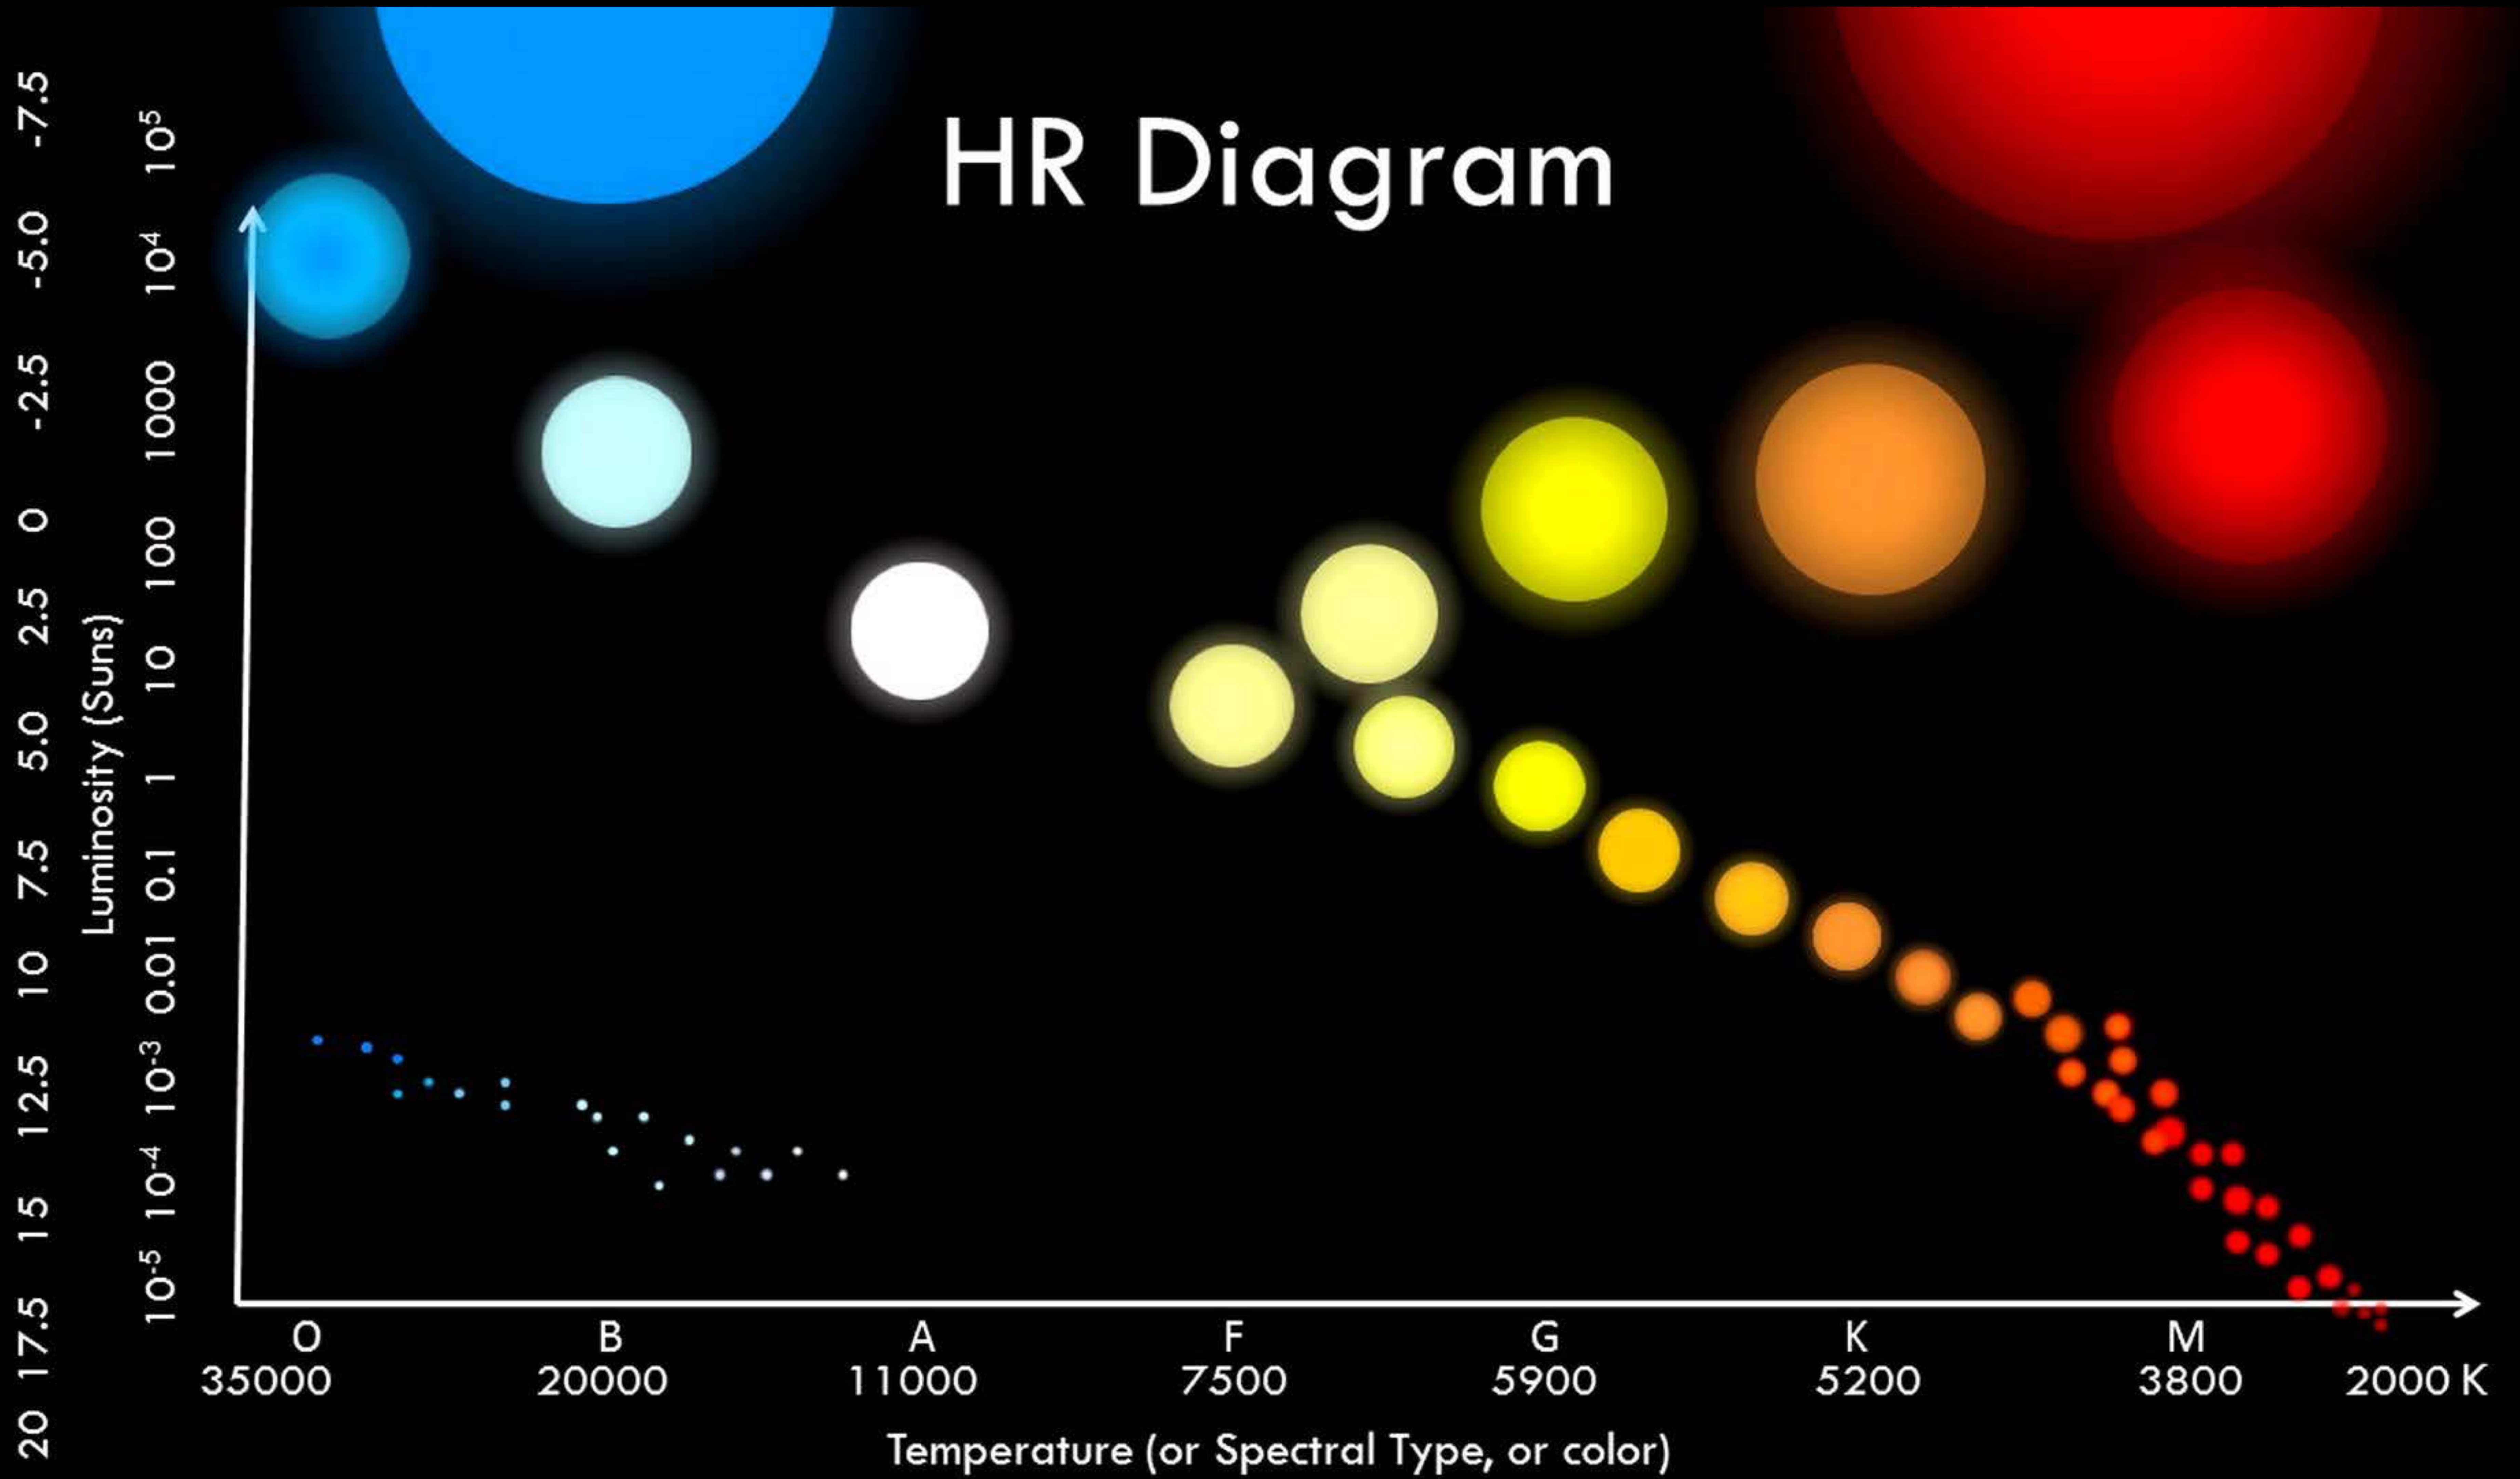

Hertzsprung-Russell diagram

Credit: Adric Riedel https://www.youtube.com/watch?v=588oz17Xj8o

Starry Night, Van Gogh

Cannon Beach, Oregon

What is color?

Light is EM Radiation; Color is Frequency

- Light is oscillating electric & magnetic field
- KEY IDEA: frequency determines color of light
- Q: What is the difference between *frequency* and *wavelength*?

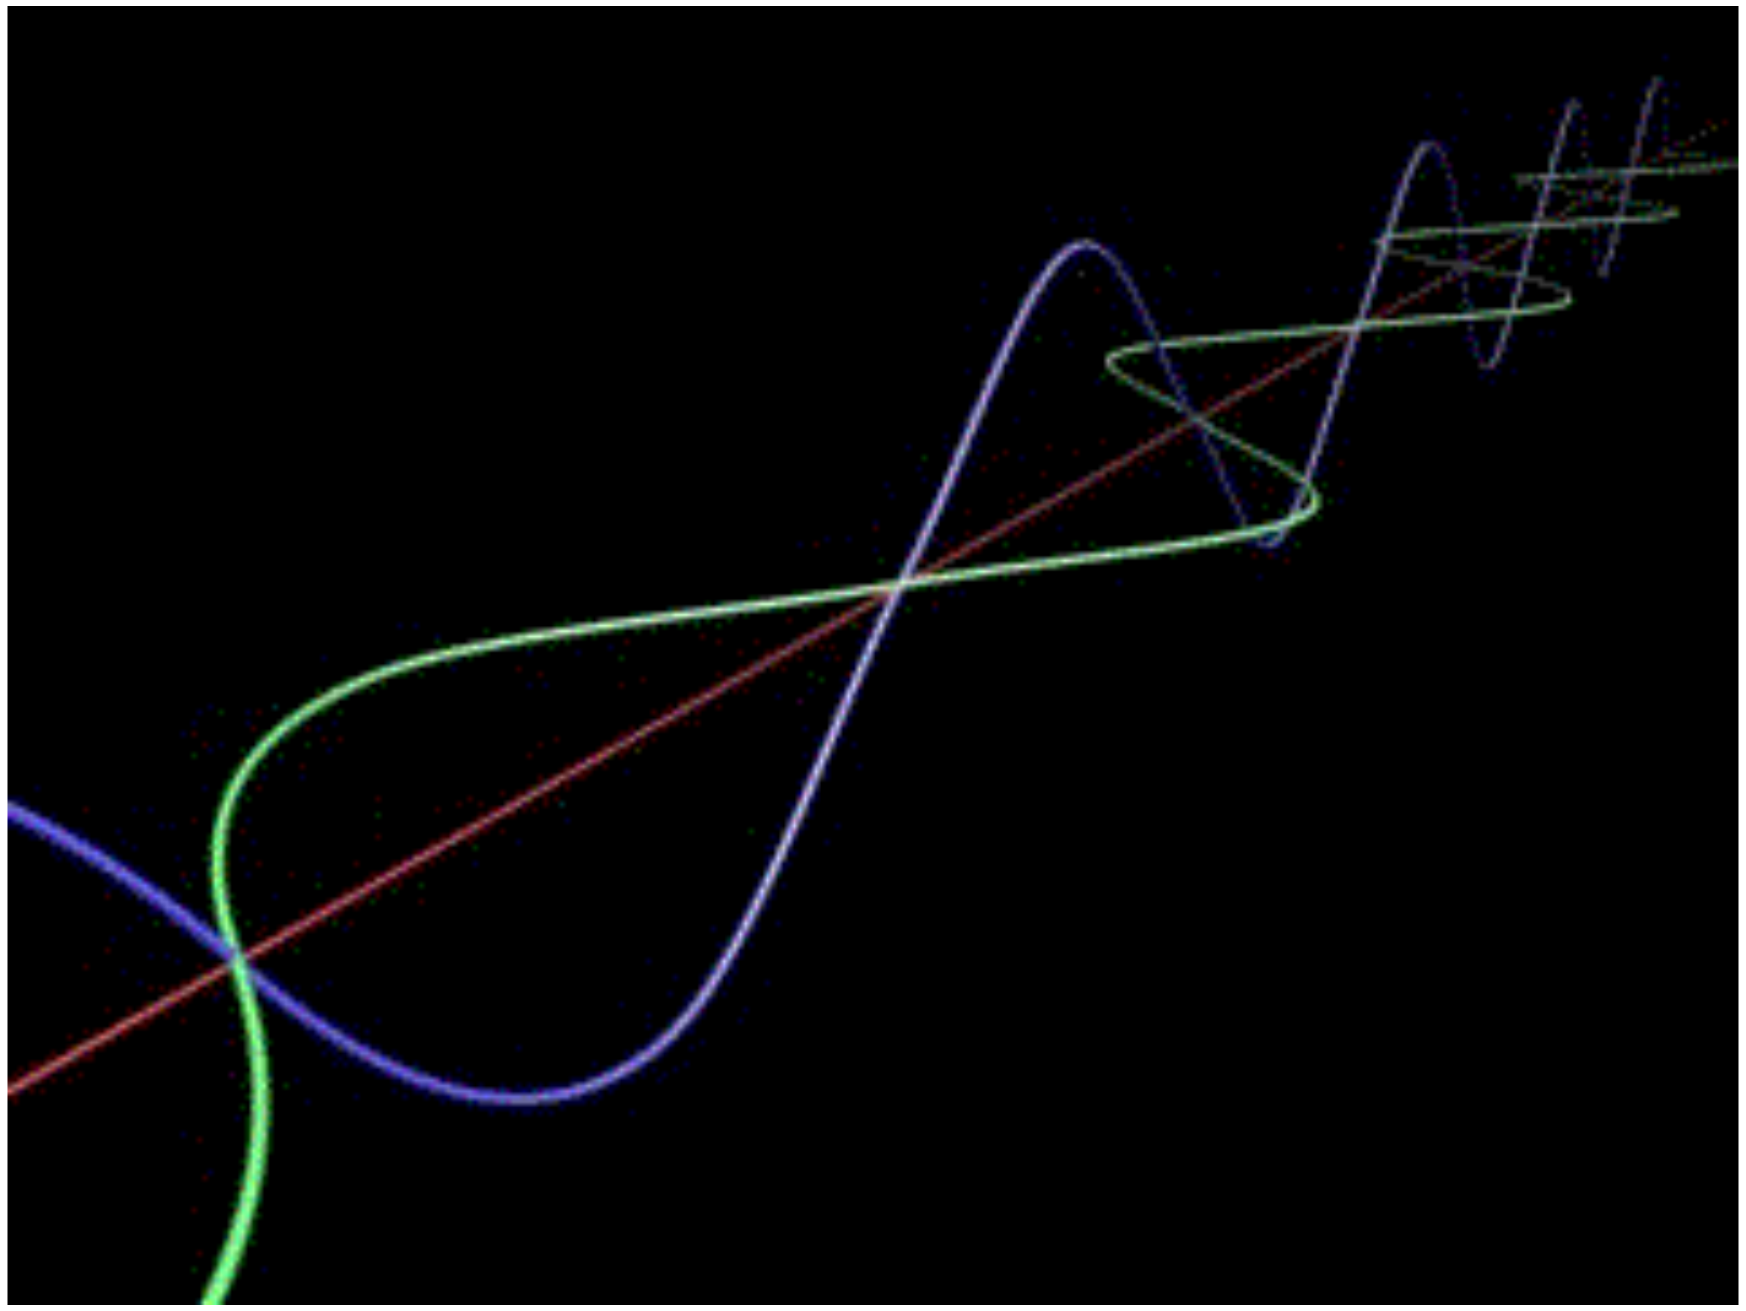

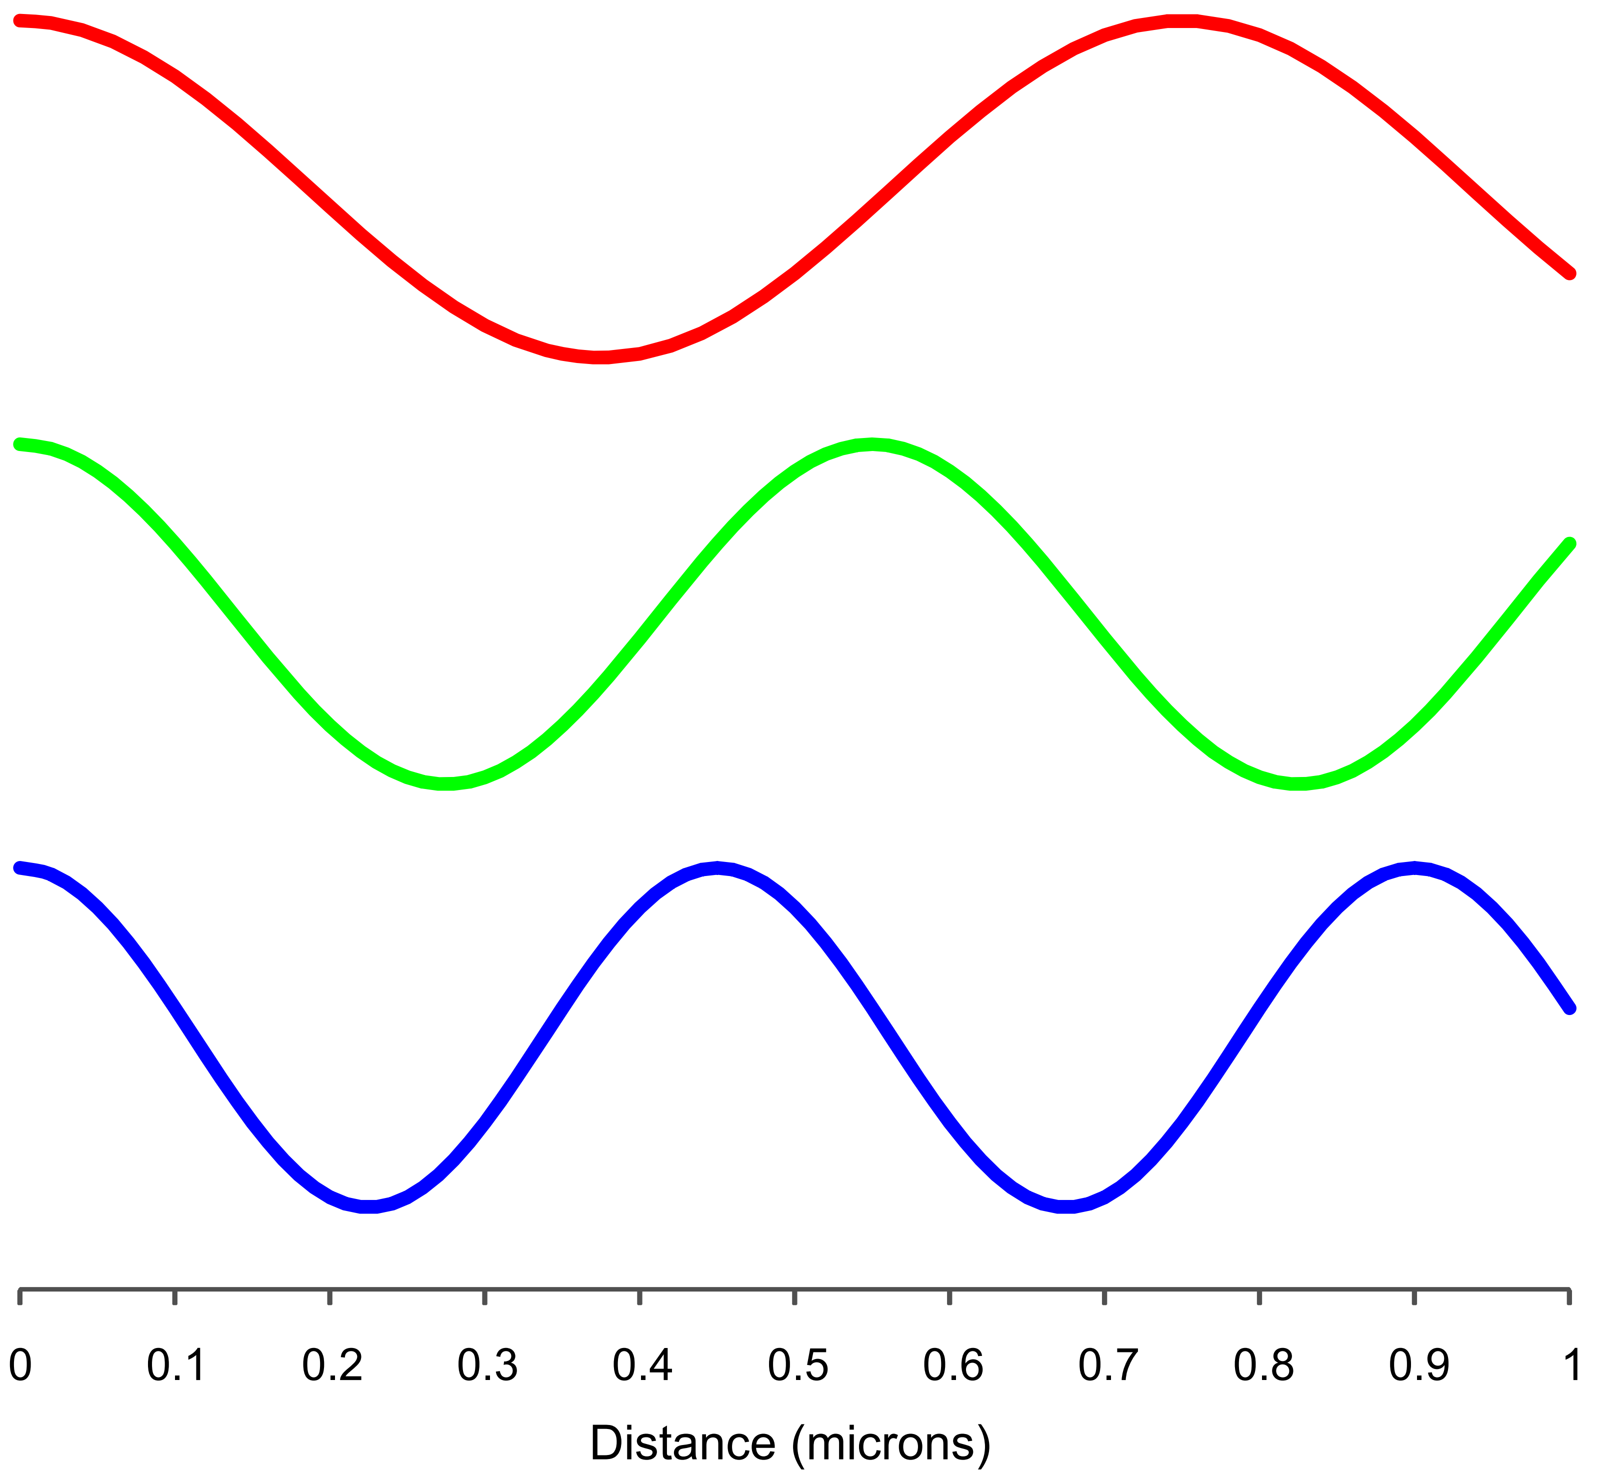

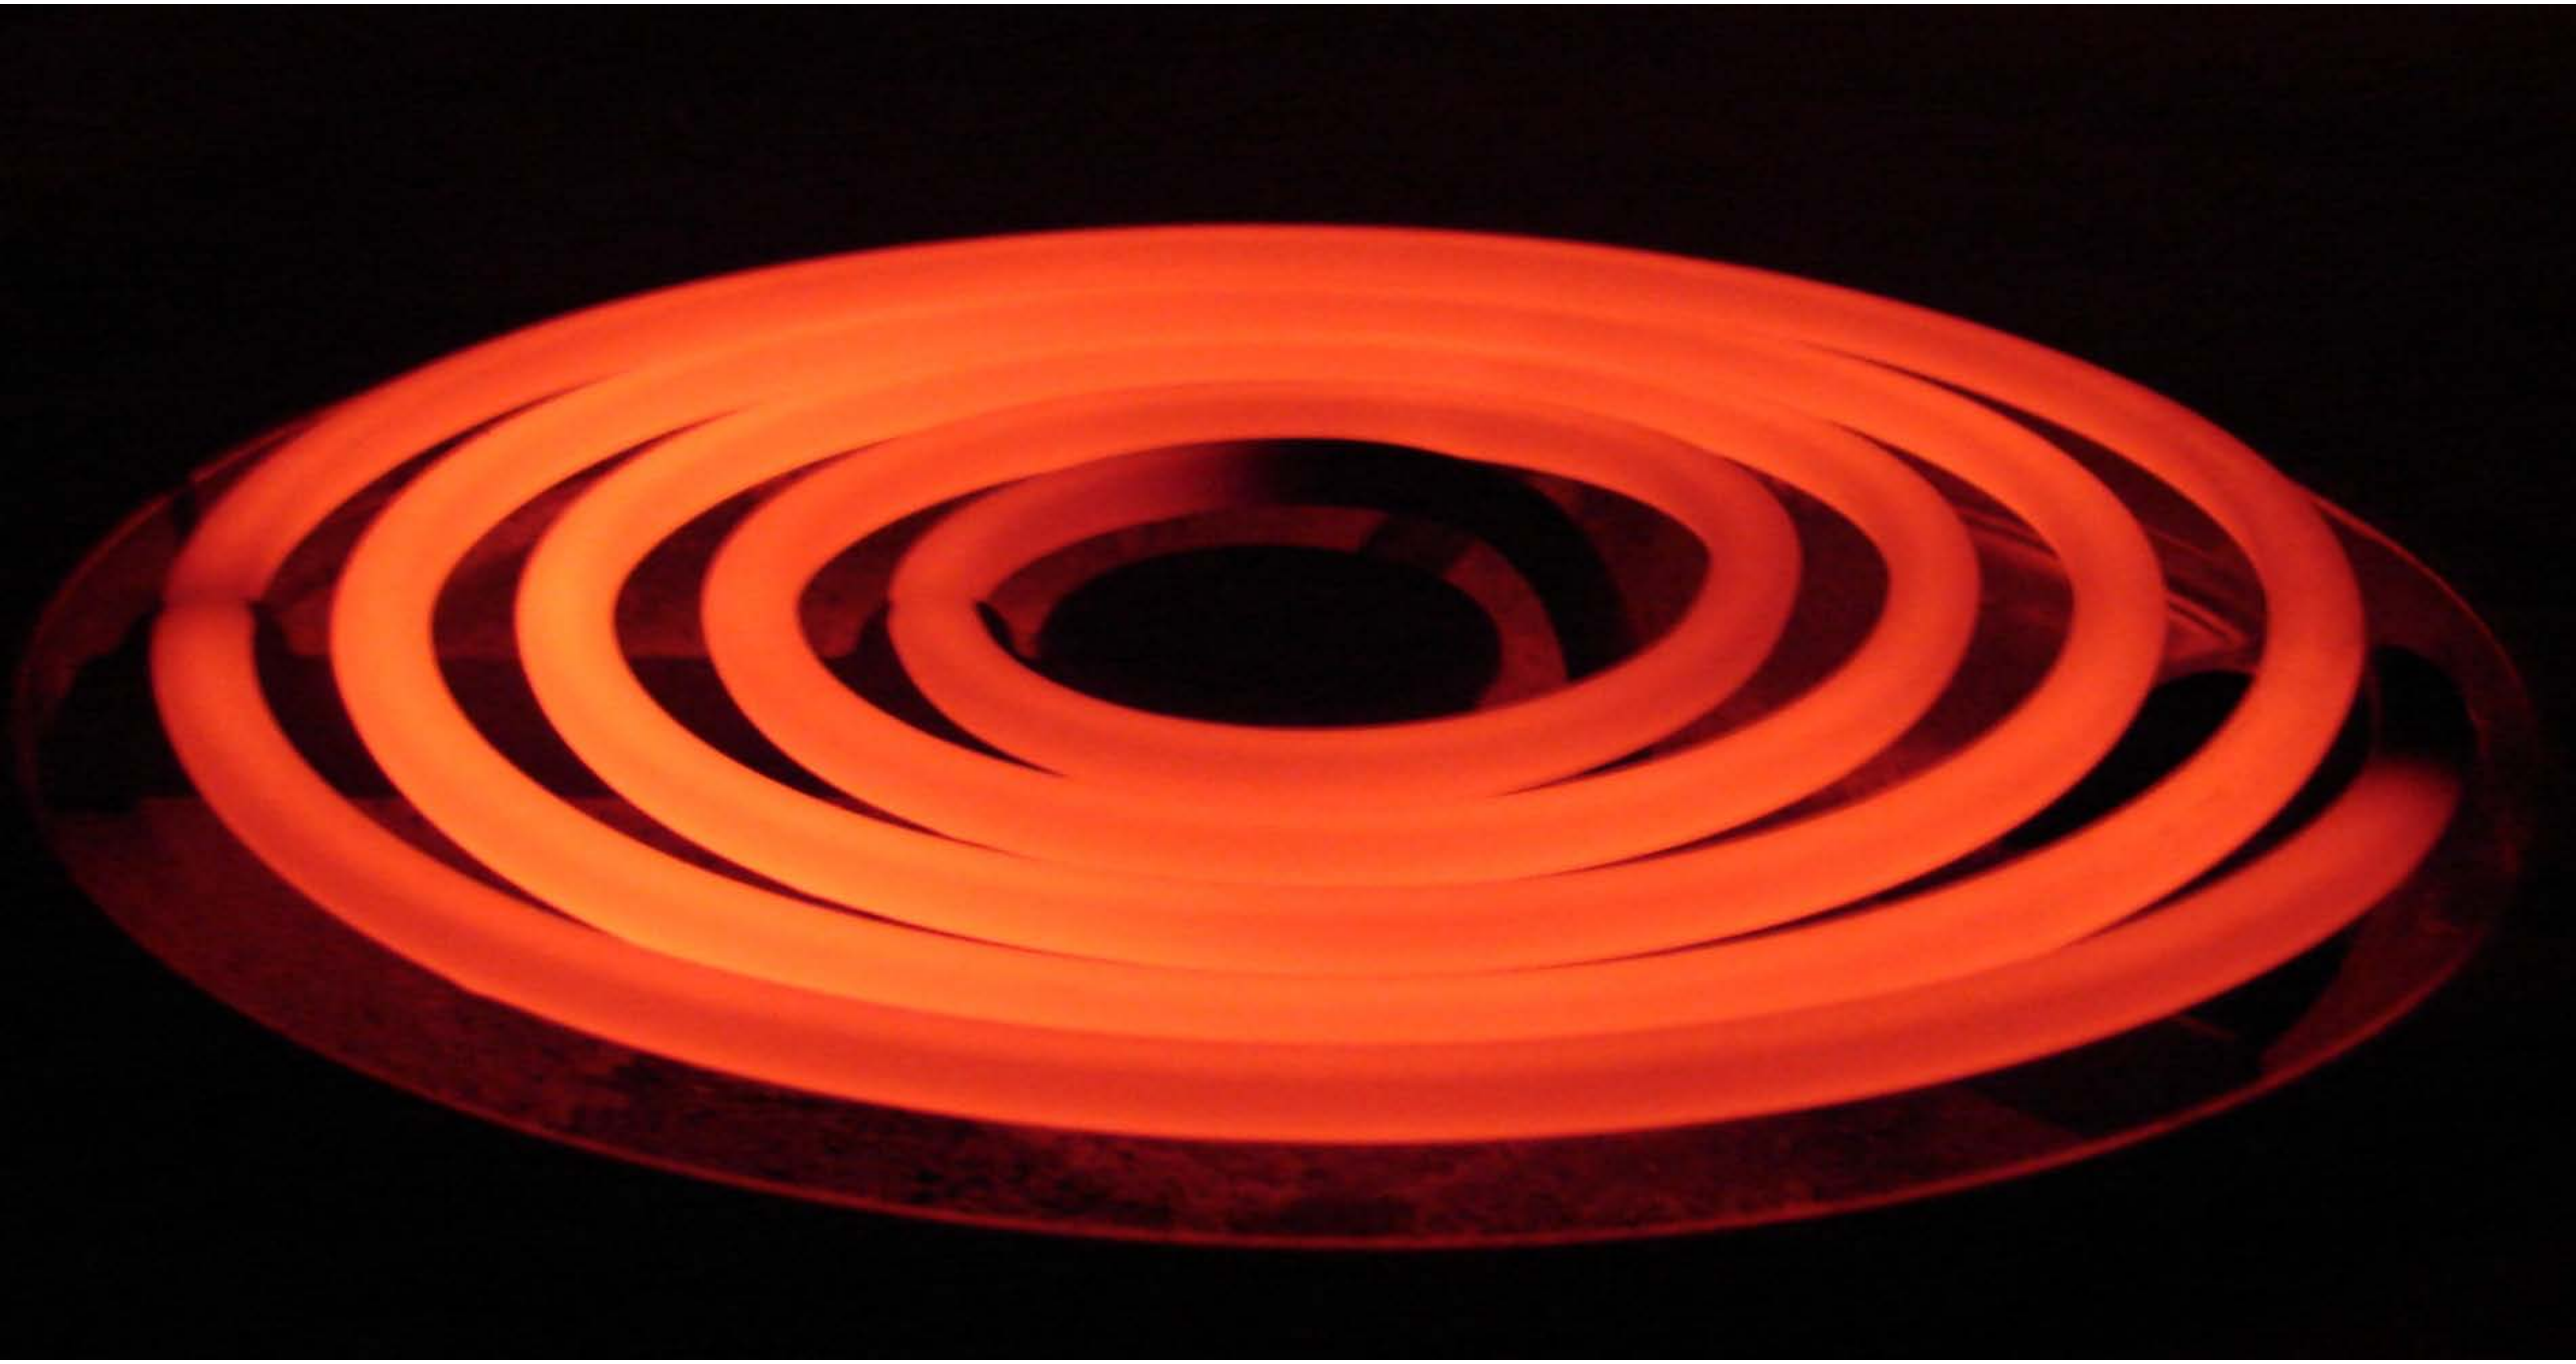

Q: Why does your stove turn red when it heats up?

Heat generates light

- One of *many* ways light is produced:
- Maxwell: motion of charged particles creates EM field
- Thermodynamics: …particles jiggle around!
- Hence, anything moving generates light
- In other words:
  - *every* object around you is producing color!
  - frequency determined by temperature

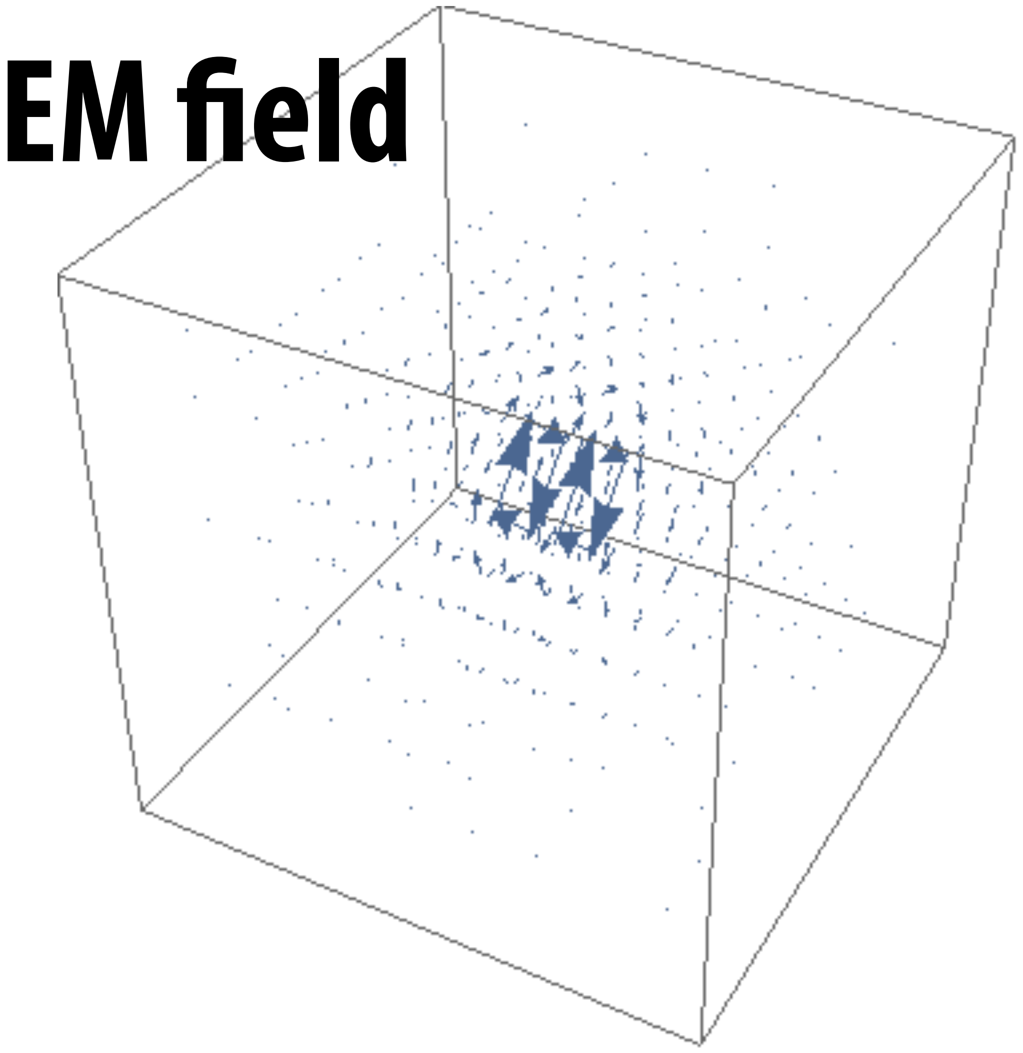

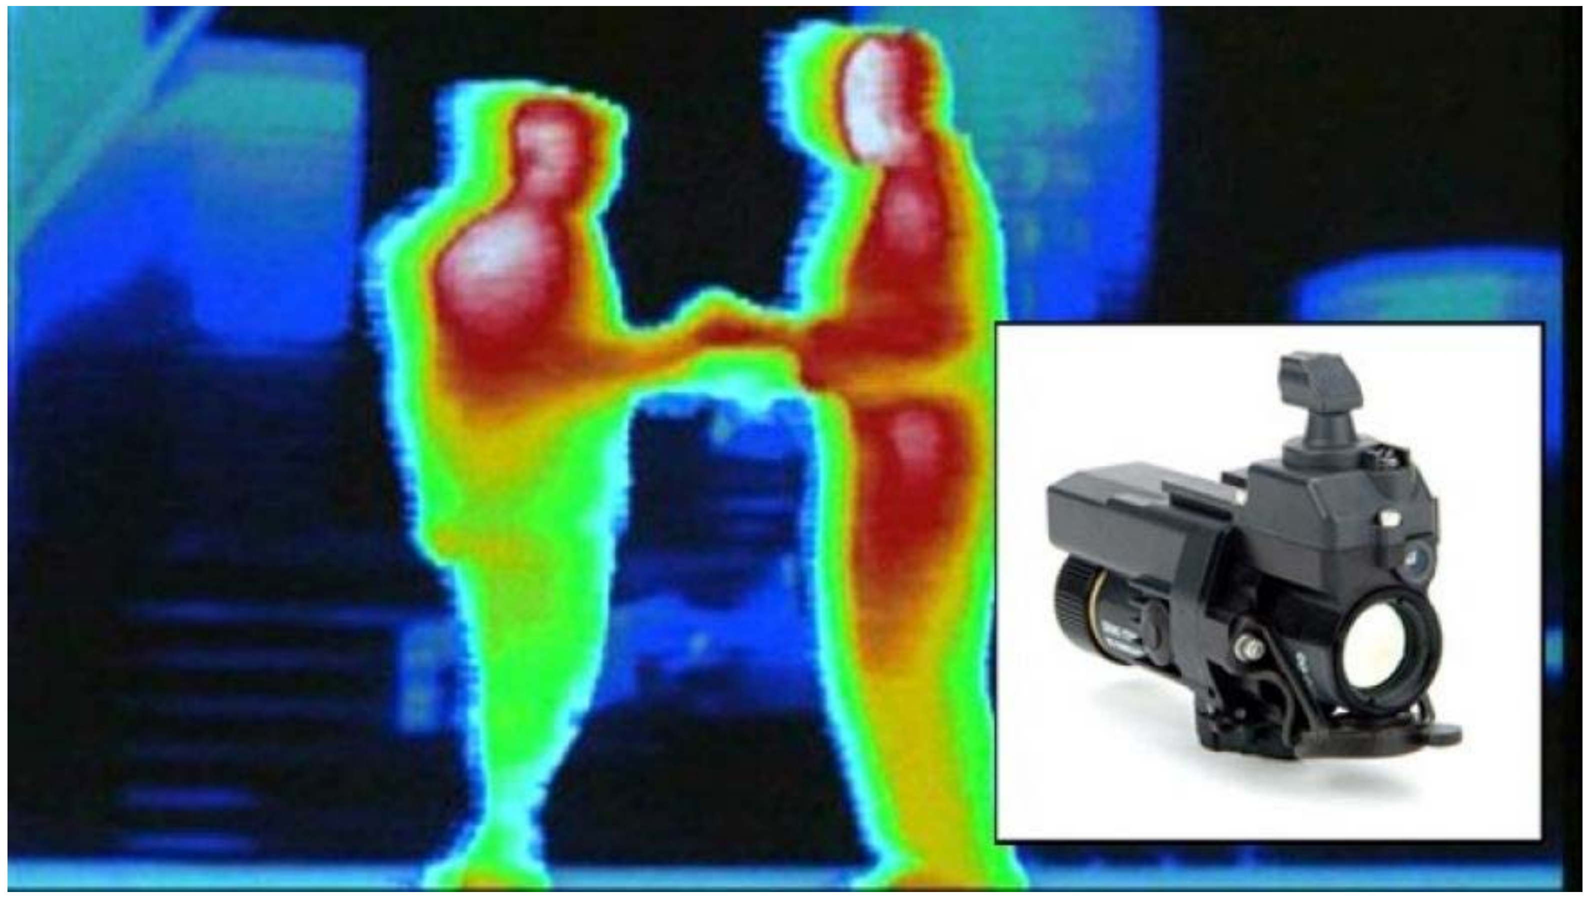

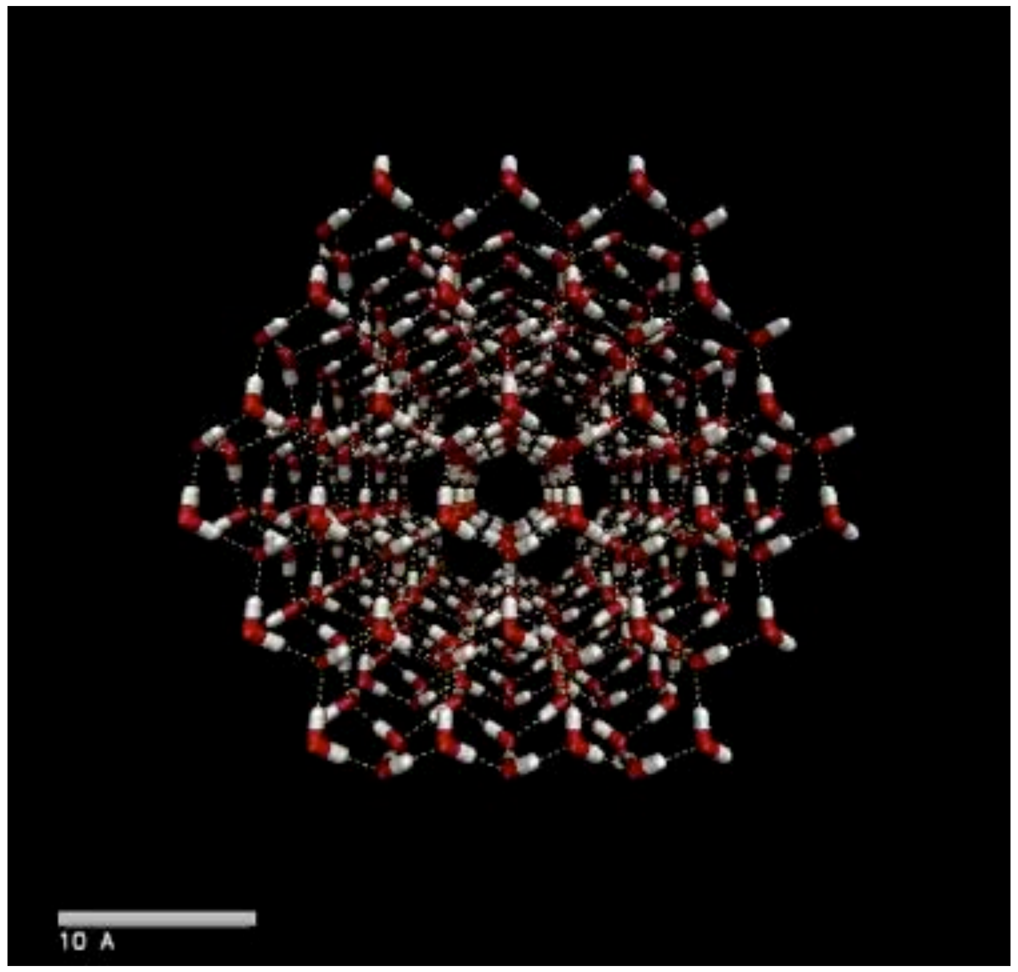

Most light is not visible!

- Frequencies visible by human eyes are called "visible spectrum"
- These frequencies what we normally think of as "color"

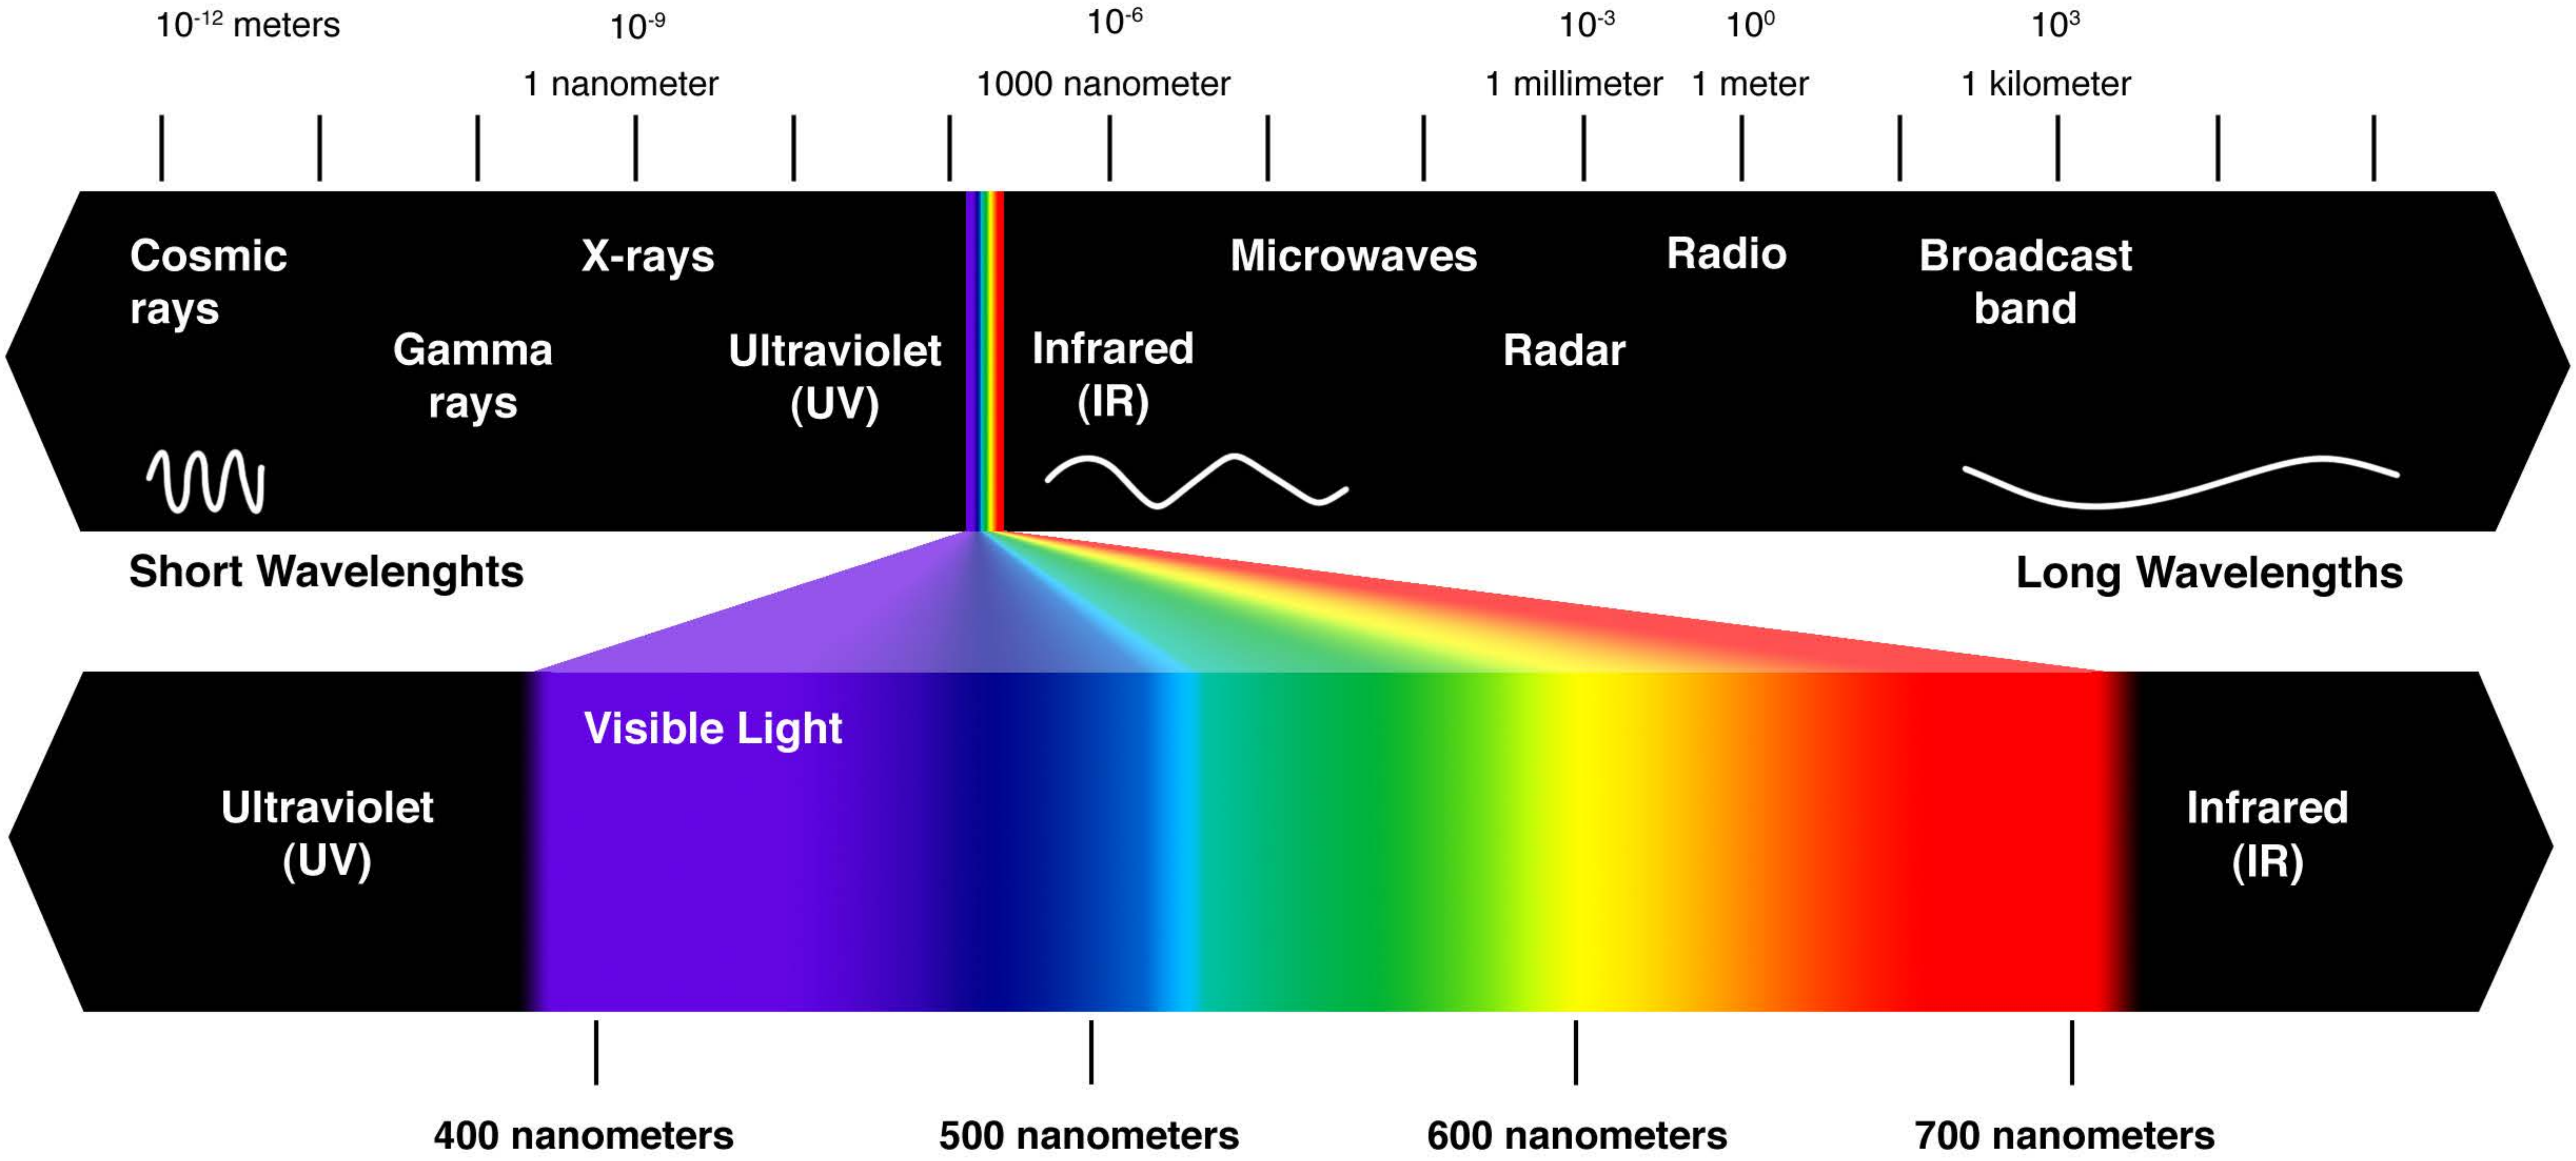

Natural light is a mixture of frequencies

- "White" light is really a mixture of all (visible) frequencies
- E.g., the light from our sun

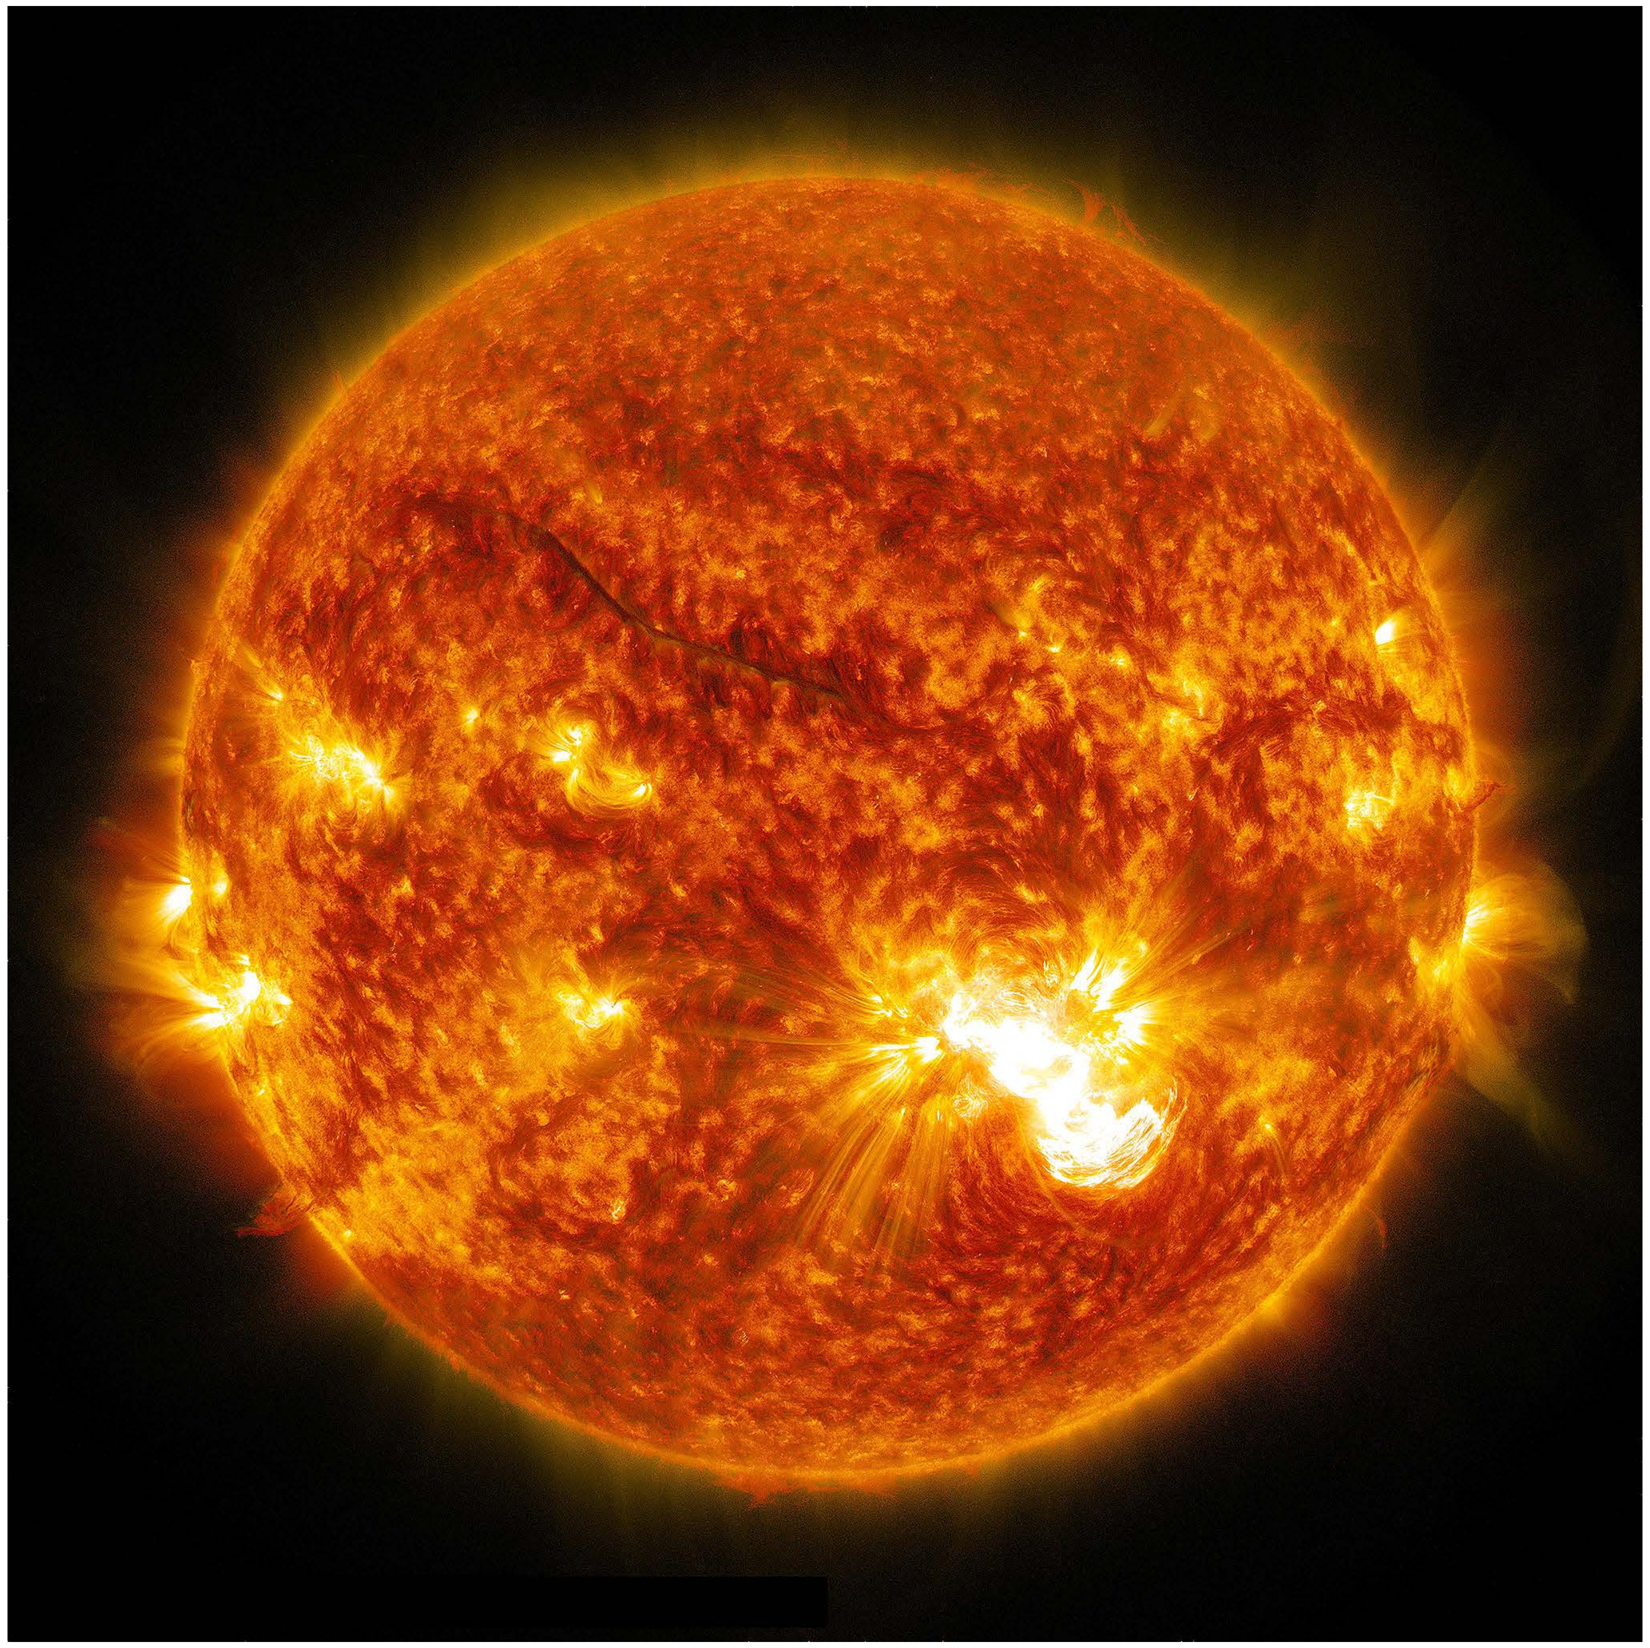

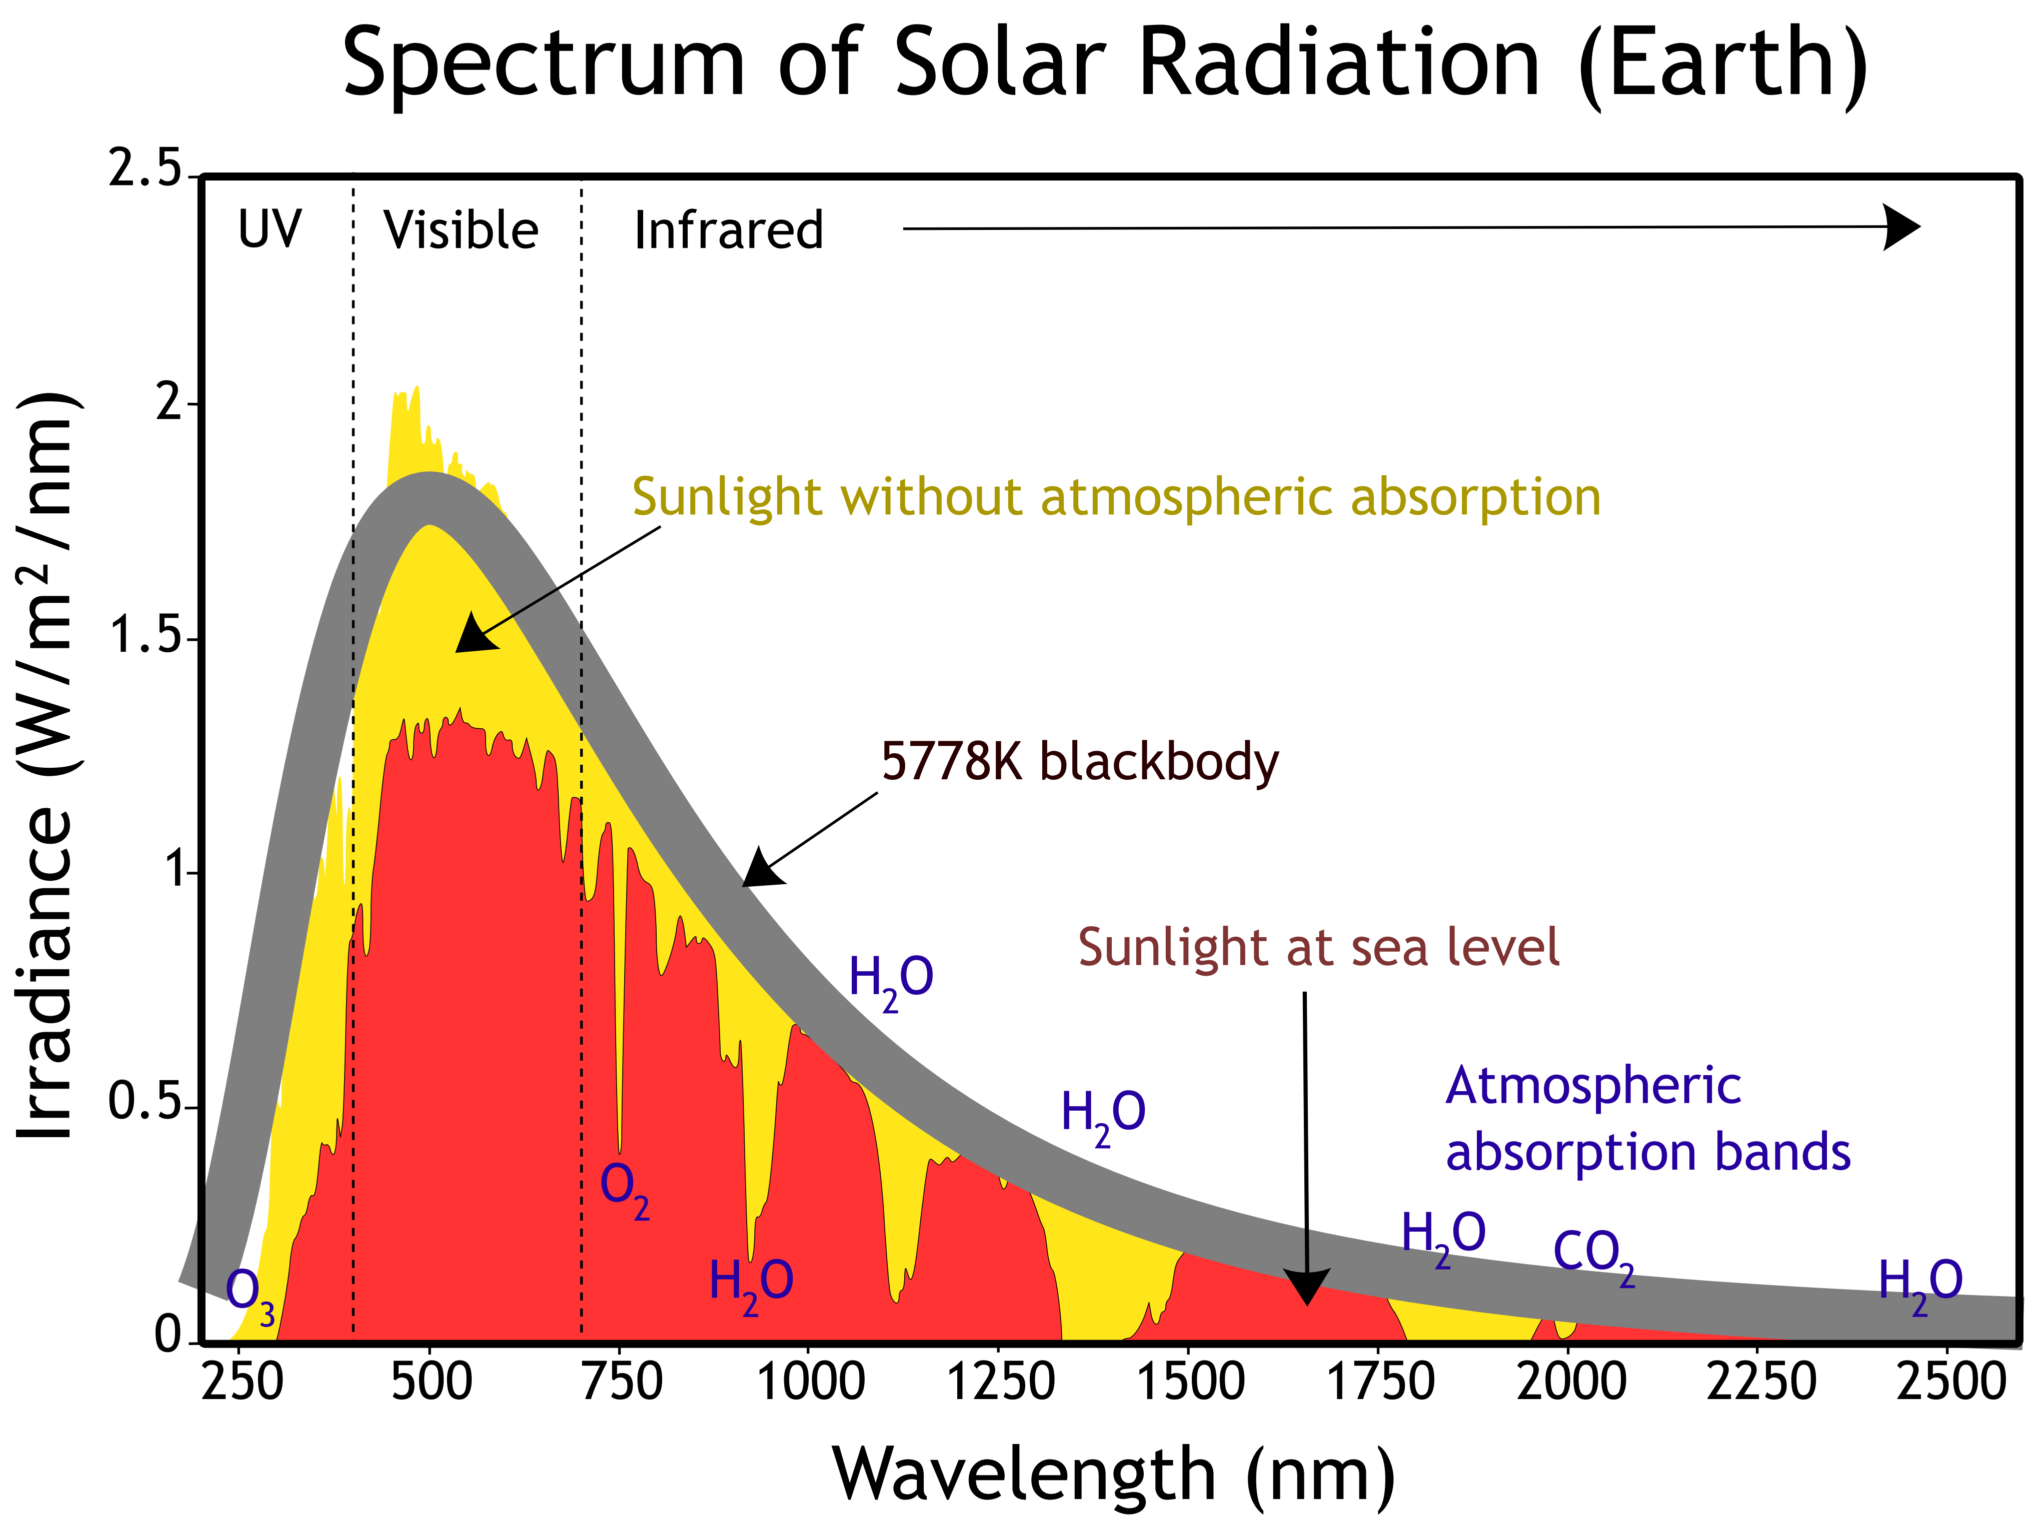

Additive vs. Subtractive Models of Light

- Spectrum we just saw for the sun *"emission spectrum"*
  - How much light is *produced* (by heat, fusion, etc.)
  - Useful for, e.g., characterizing color of a lightbulb
- Another useful description: *"absorbtion spectrum"*
  - How much light is *absorbed* (e.g., turned *into* heat)
  - Useful for, e.g., characterizing color of paint, ink, etc.

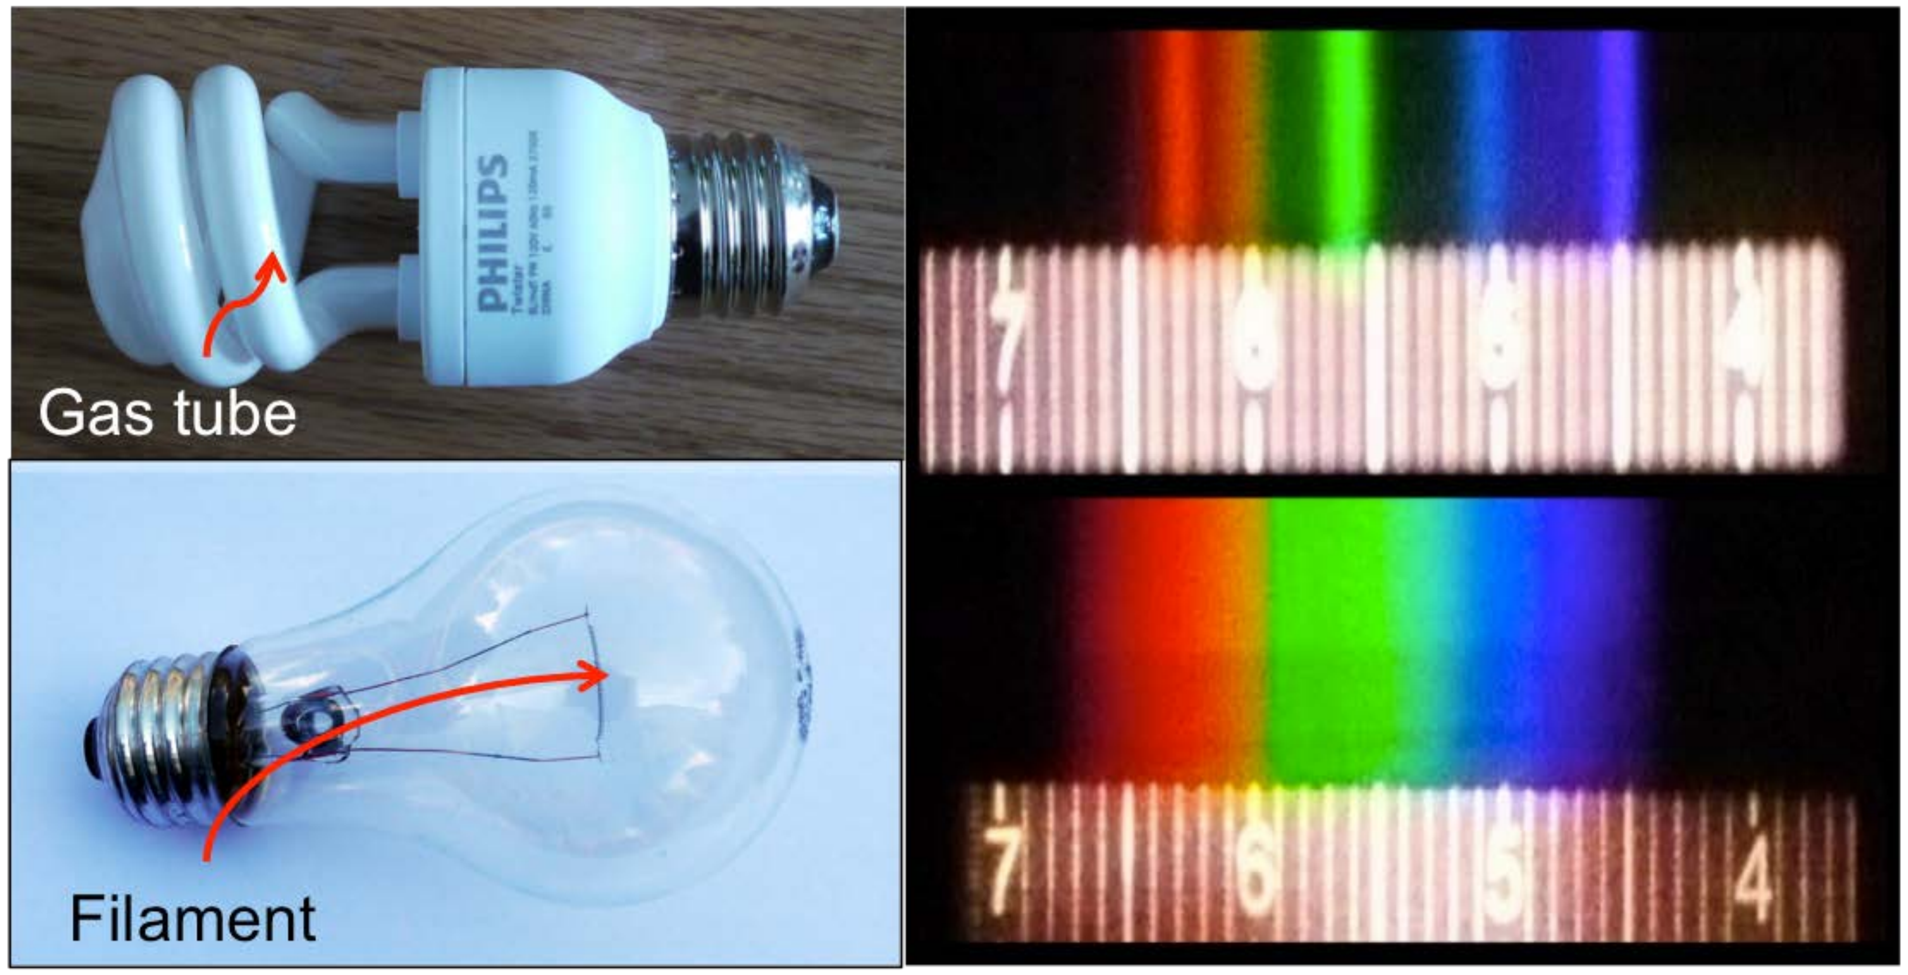

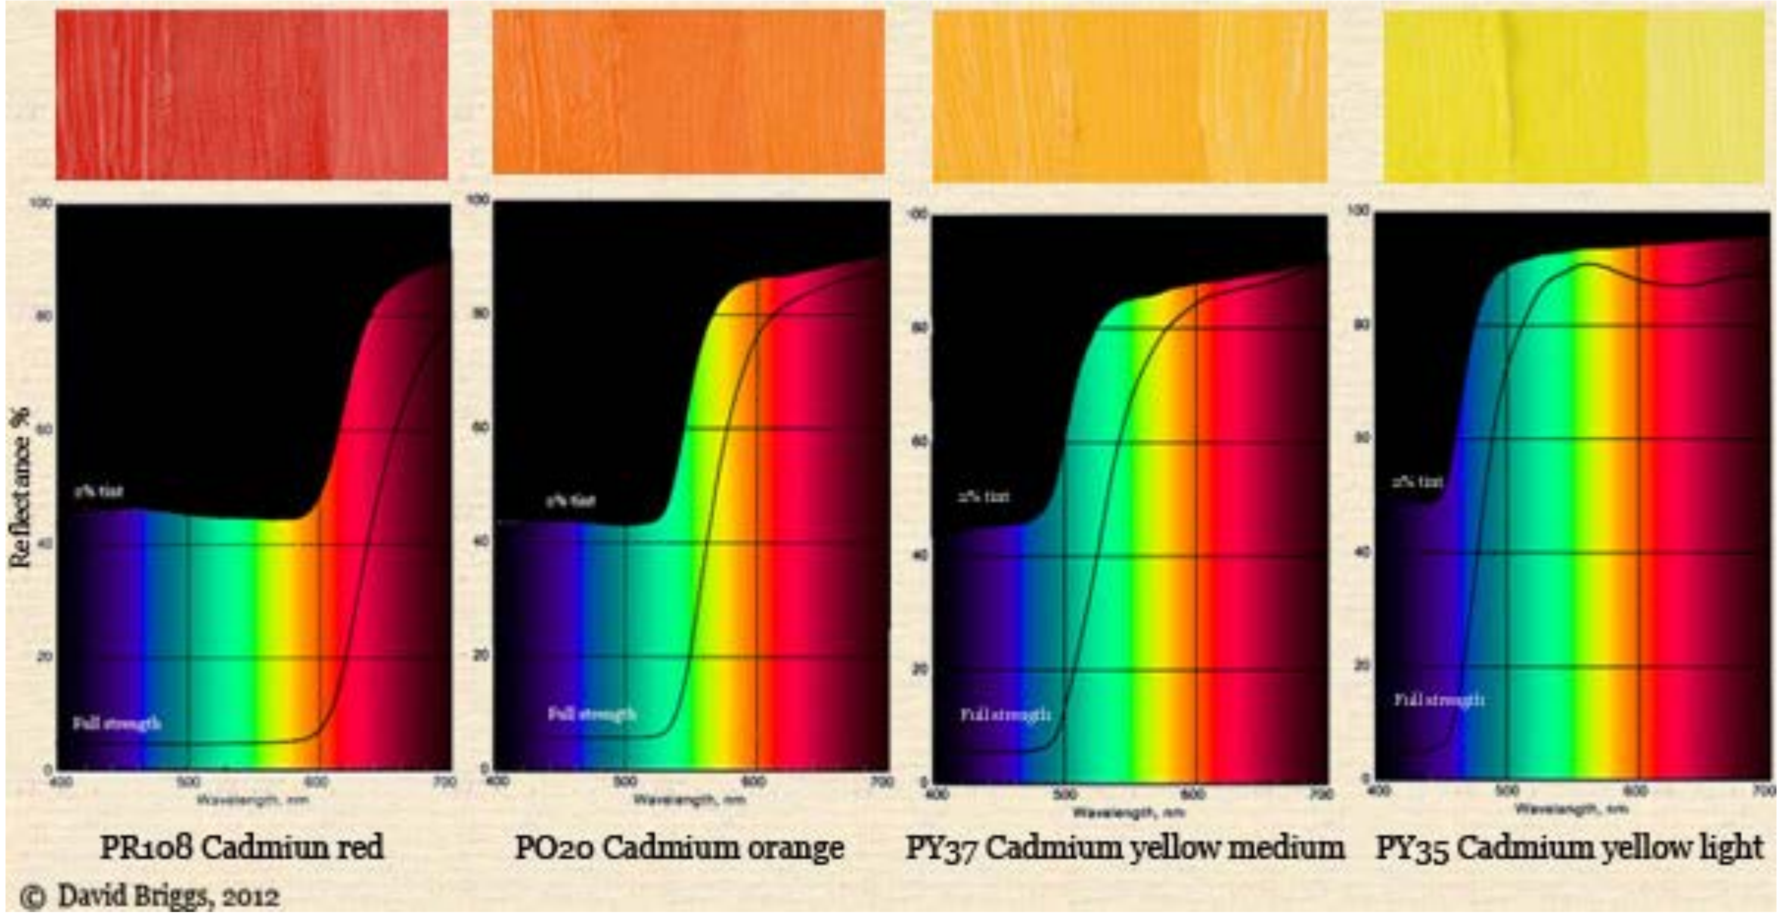

© David Briggs, 2012

Emission Spectrum

Describes light intensity as a function of frequency

Below: spectrum of various common light sources:

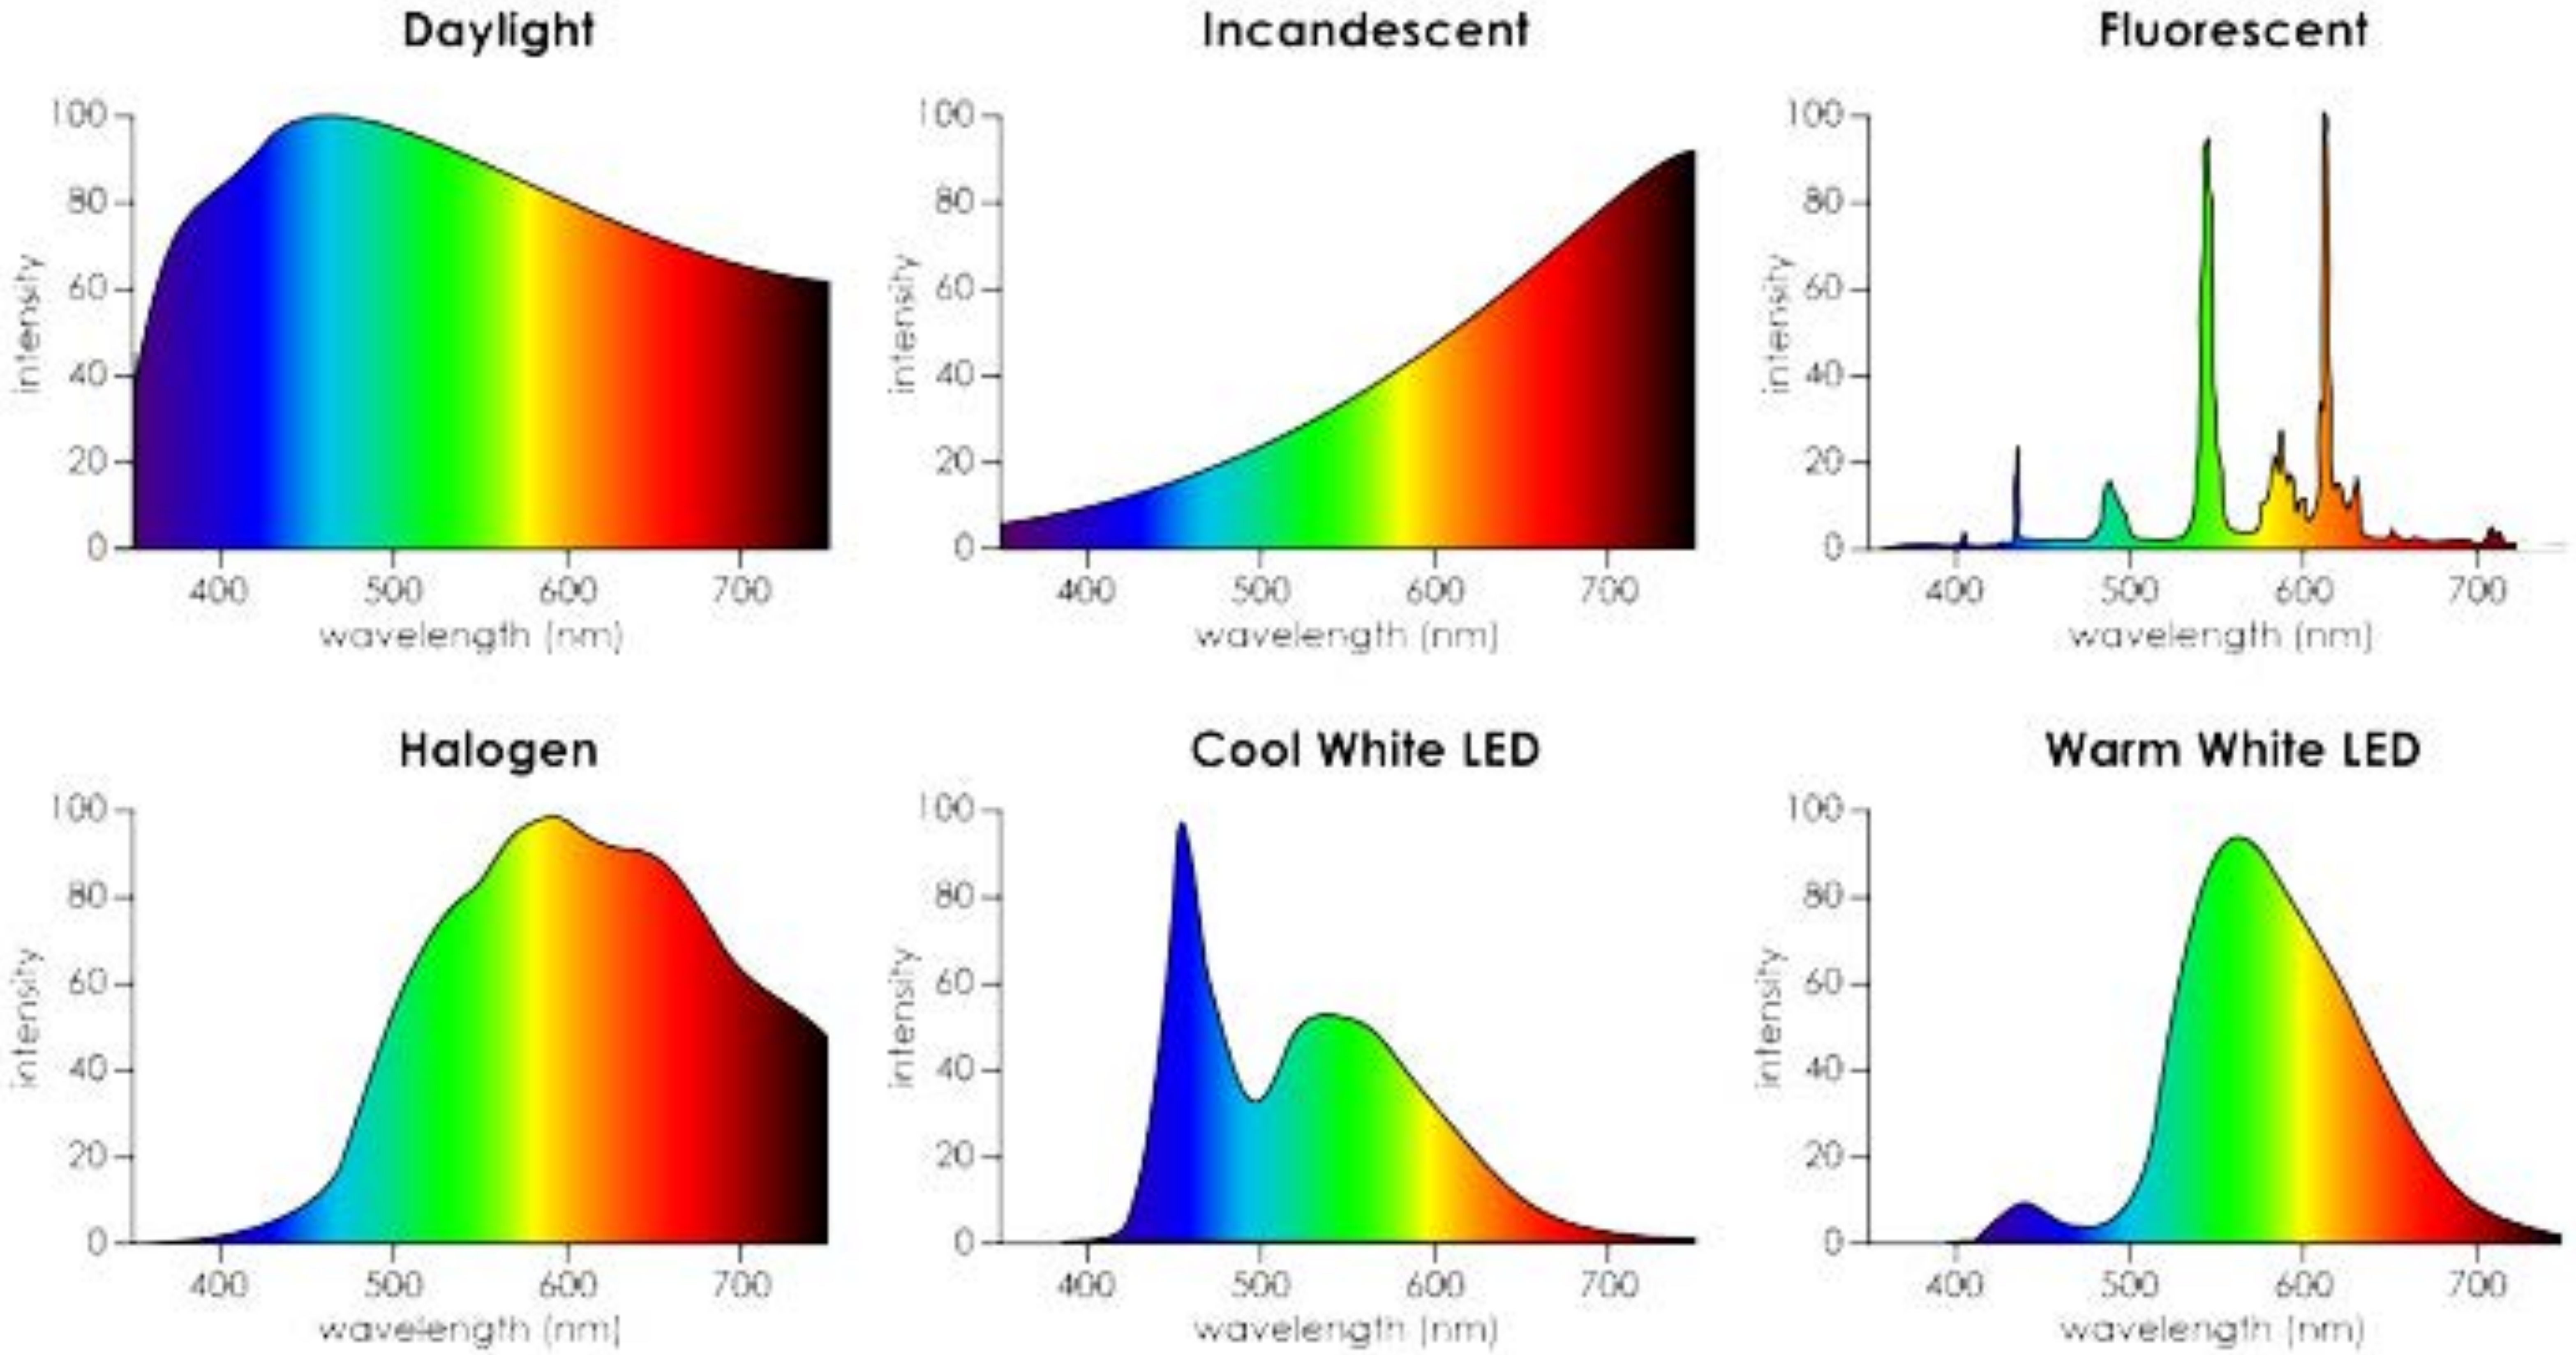

Figure credit:

admesy
ADVANCED MEASUREMENT SYSTEMS

Emission Spectrum—Example

- Why so many different kinds of lightbulbs on the market?
- "Quality" of light:

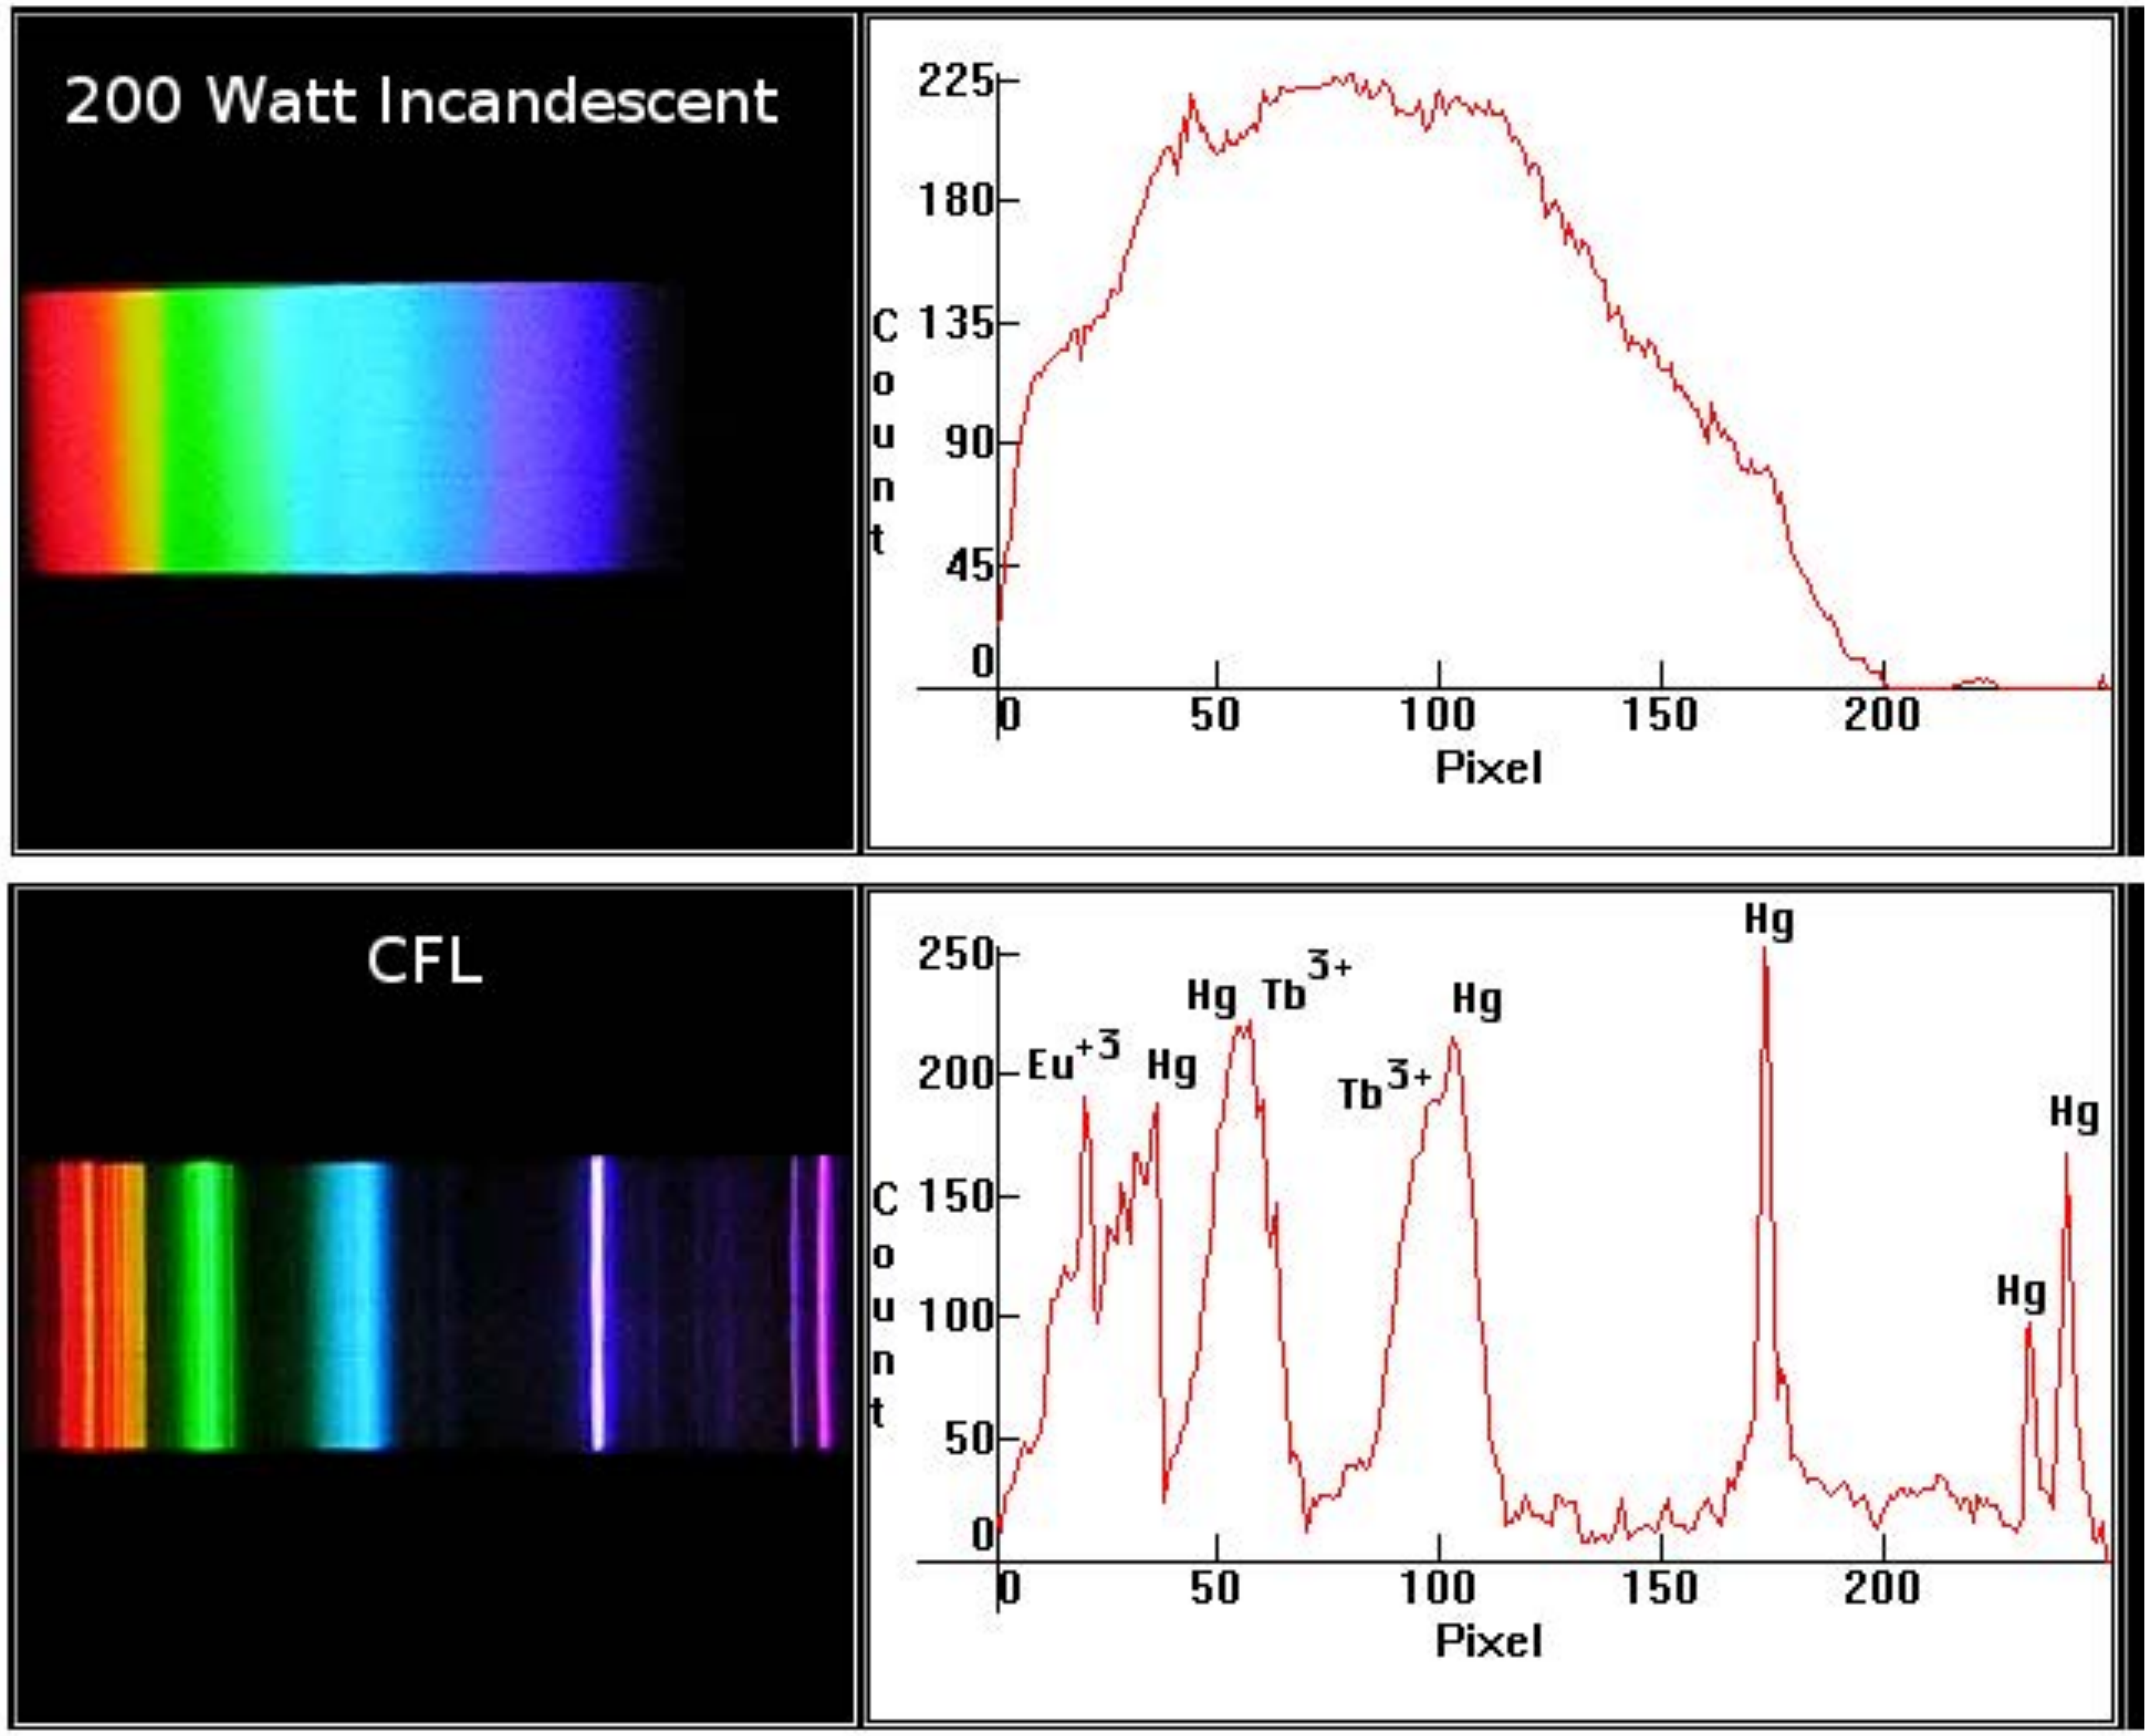

Incandescent:

+more sun-like

-power-hungry

CFL:

-"choppy" spectrum

+power efficient

Absorption Spectrum

- **Emission spectrum is *intensity* as a function of frequency**
- **Absorption spectrum is *fraction absorbed* as function of frequency**

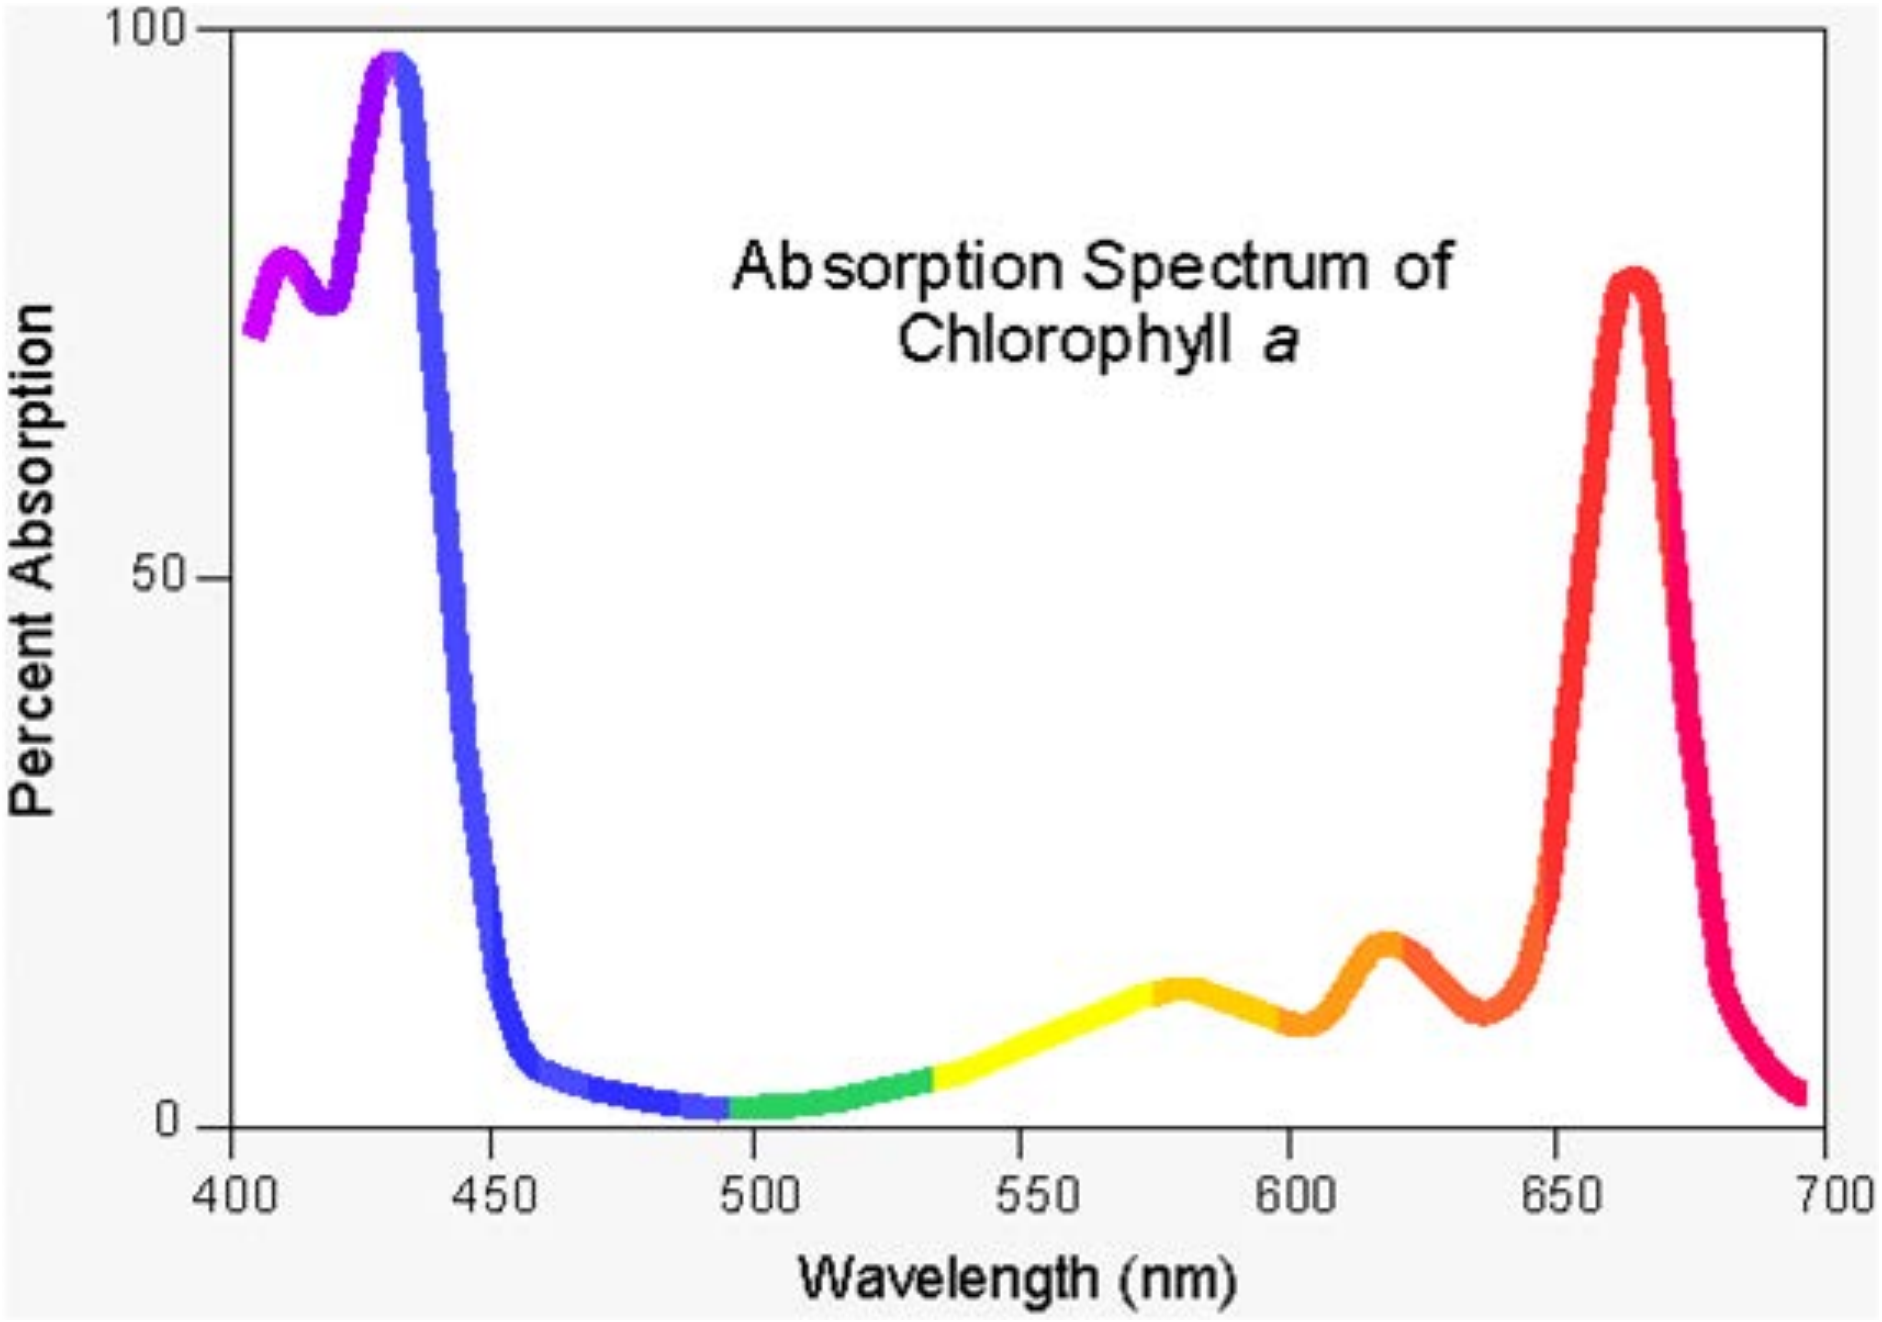

Q: What color is an object with this absorption spectrum?

This is the fundamental description of color: *intensity or absorption as a function of frequency*

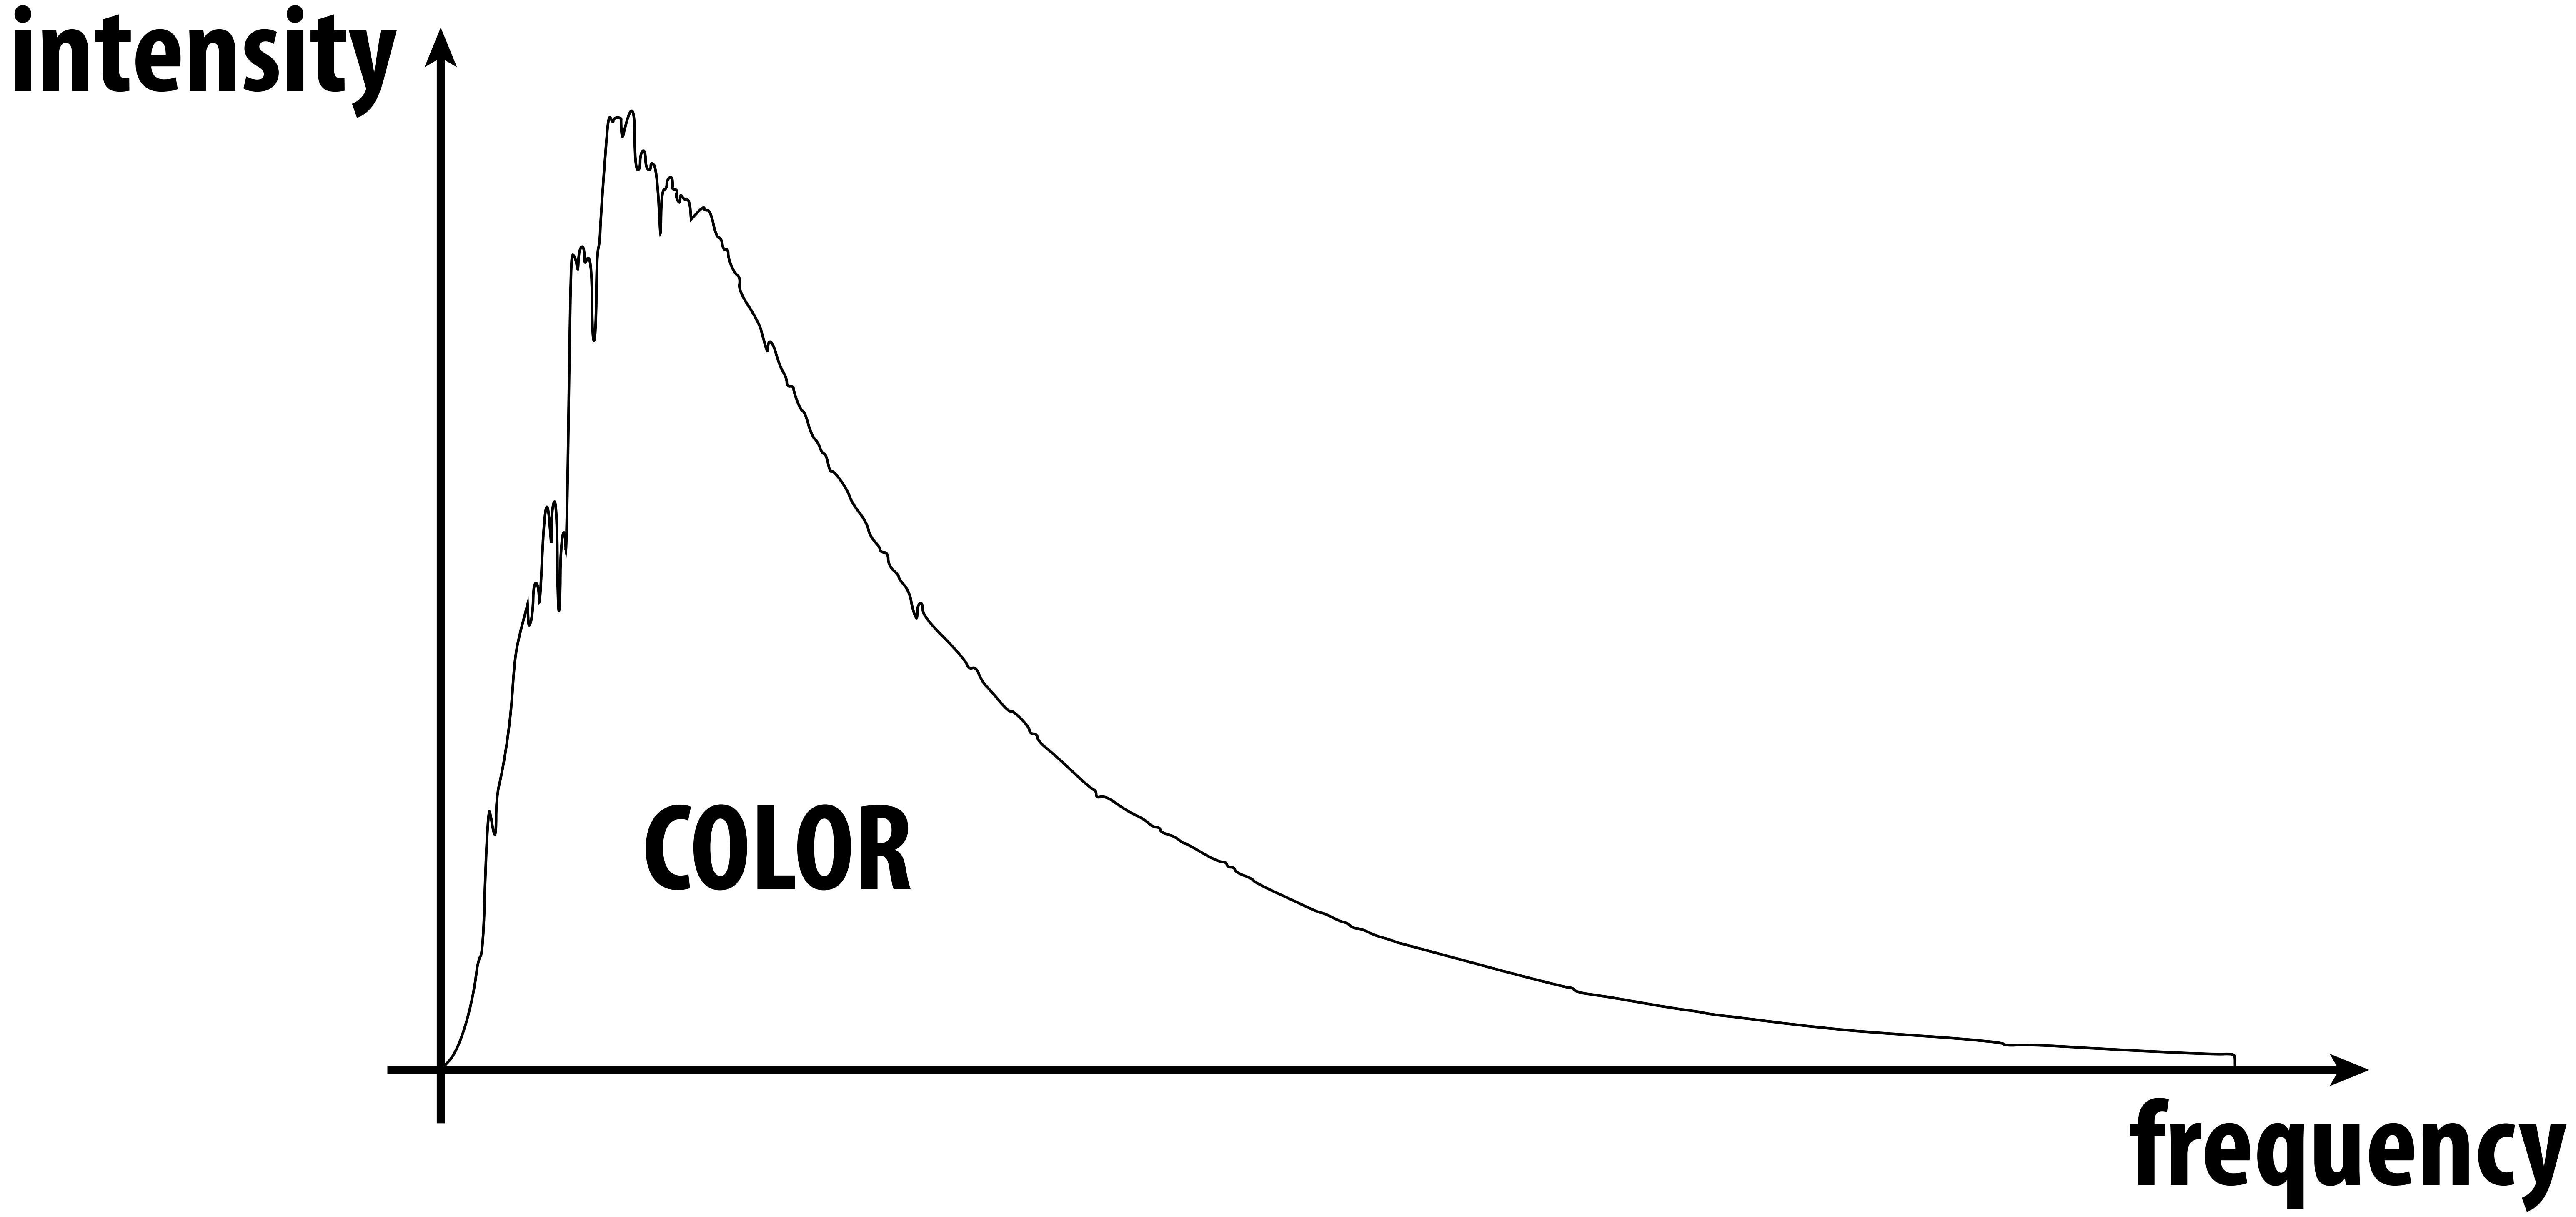

***Everything* else is merely a convenient approximation!**

If you remember to use spectral description as a starting point, the issues surrounding color theory/ practice will make a *lot* more sense!

If on the other hand you always think of color in terms of approximate digital encodings (RGB, CMYK) etc., there are certain phenomena you simply cannot explain/understand!

Interaction of emission and reflection

- Toy model for what happens when light gets reflected
  - ν—frequency (Greek “nu”)
  - Light source has emission spectrum f(ν)
  - Surface has reflection spectrum g(ν)
  - Resulting intensity is the *product* f(ν)g(ν)

f(ν)

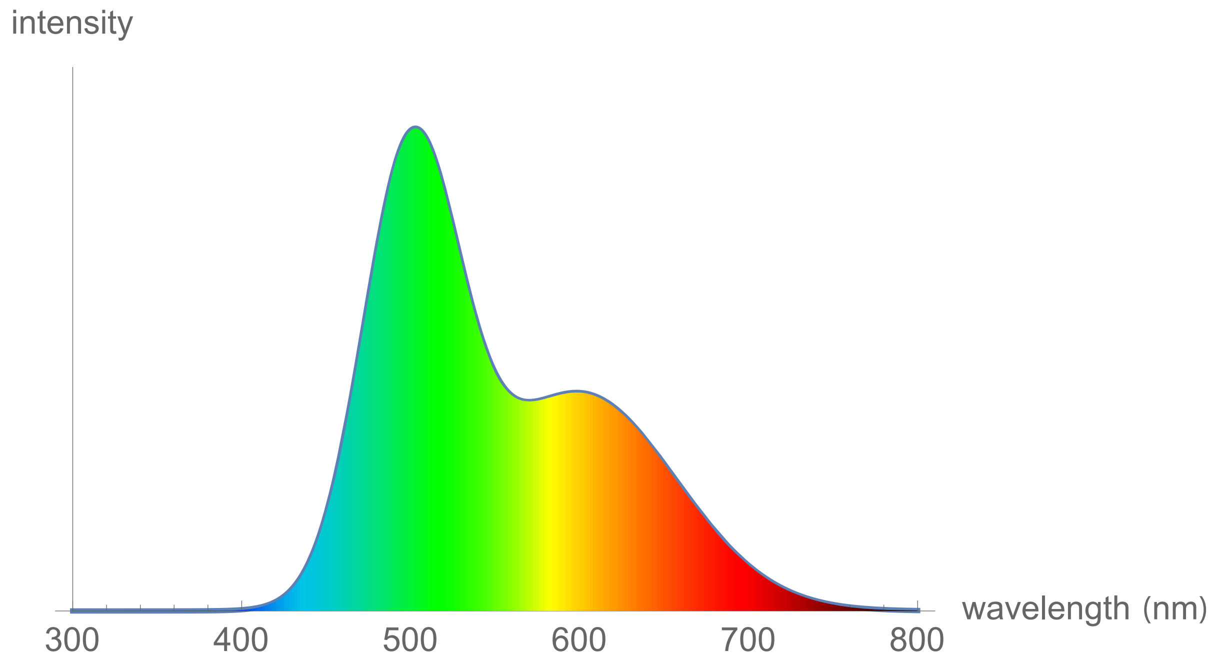

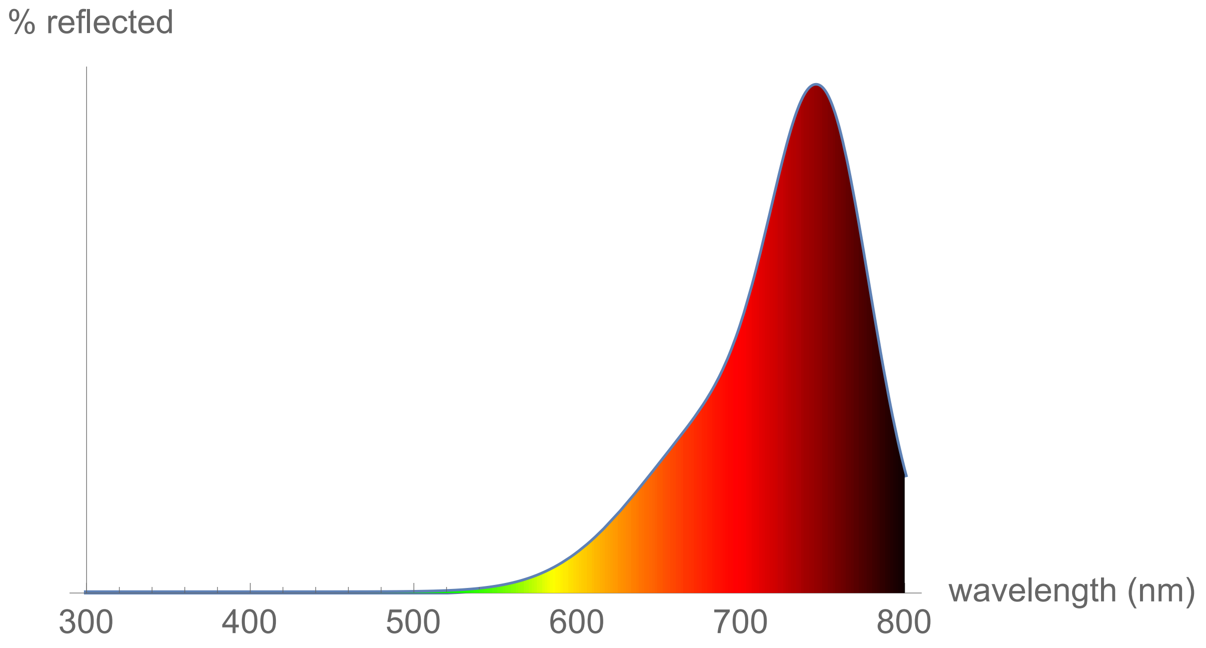

g(ν)

fg(ν)

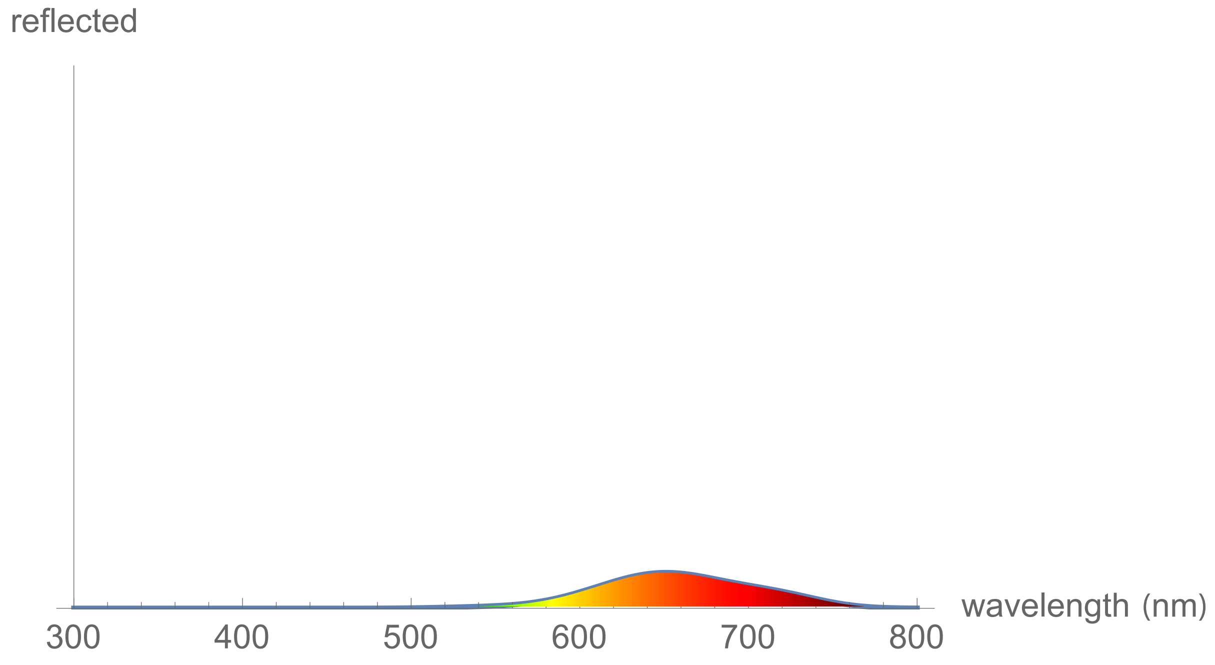

Color reproduction is hard!

- Color clearly starts to get complicated as we start combining emission and absorption/reflection (real-world challenge!)

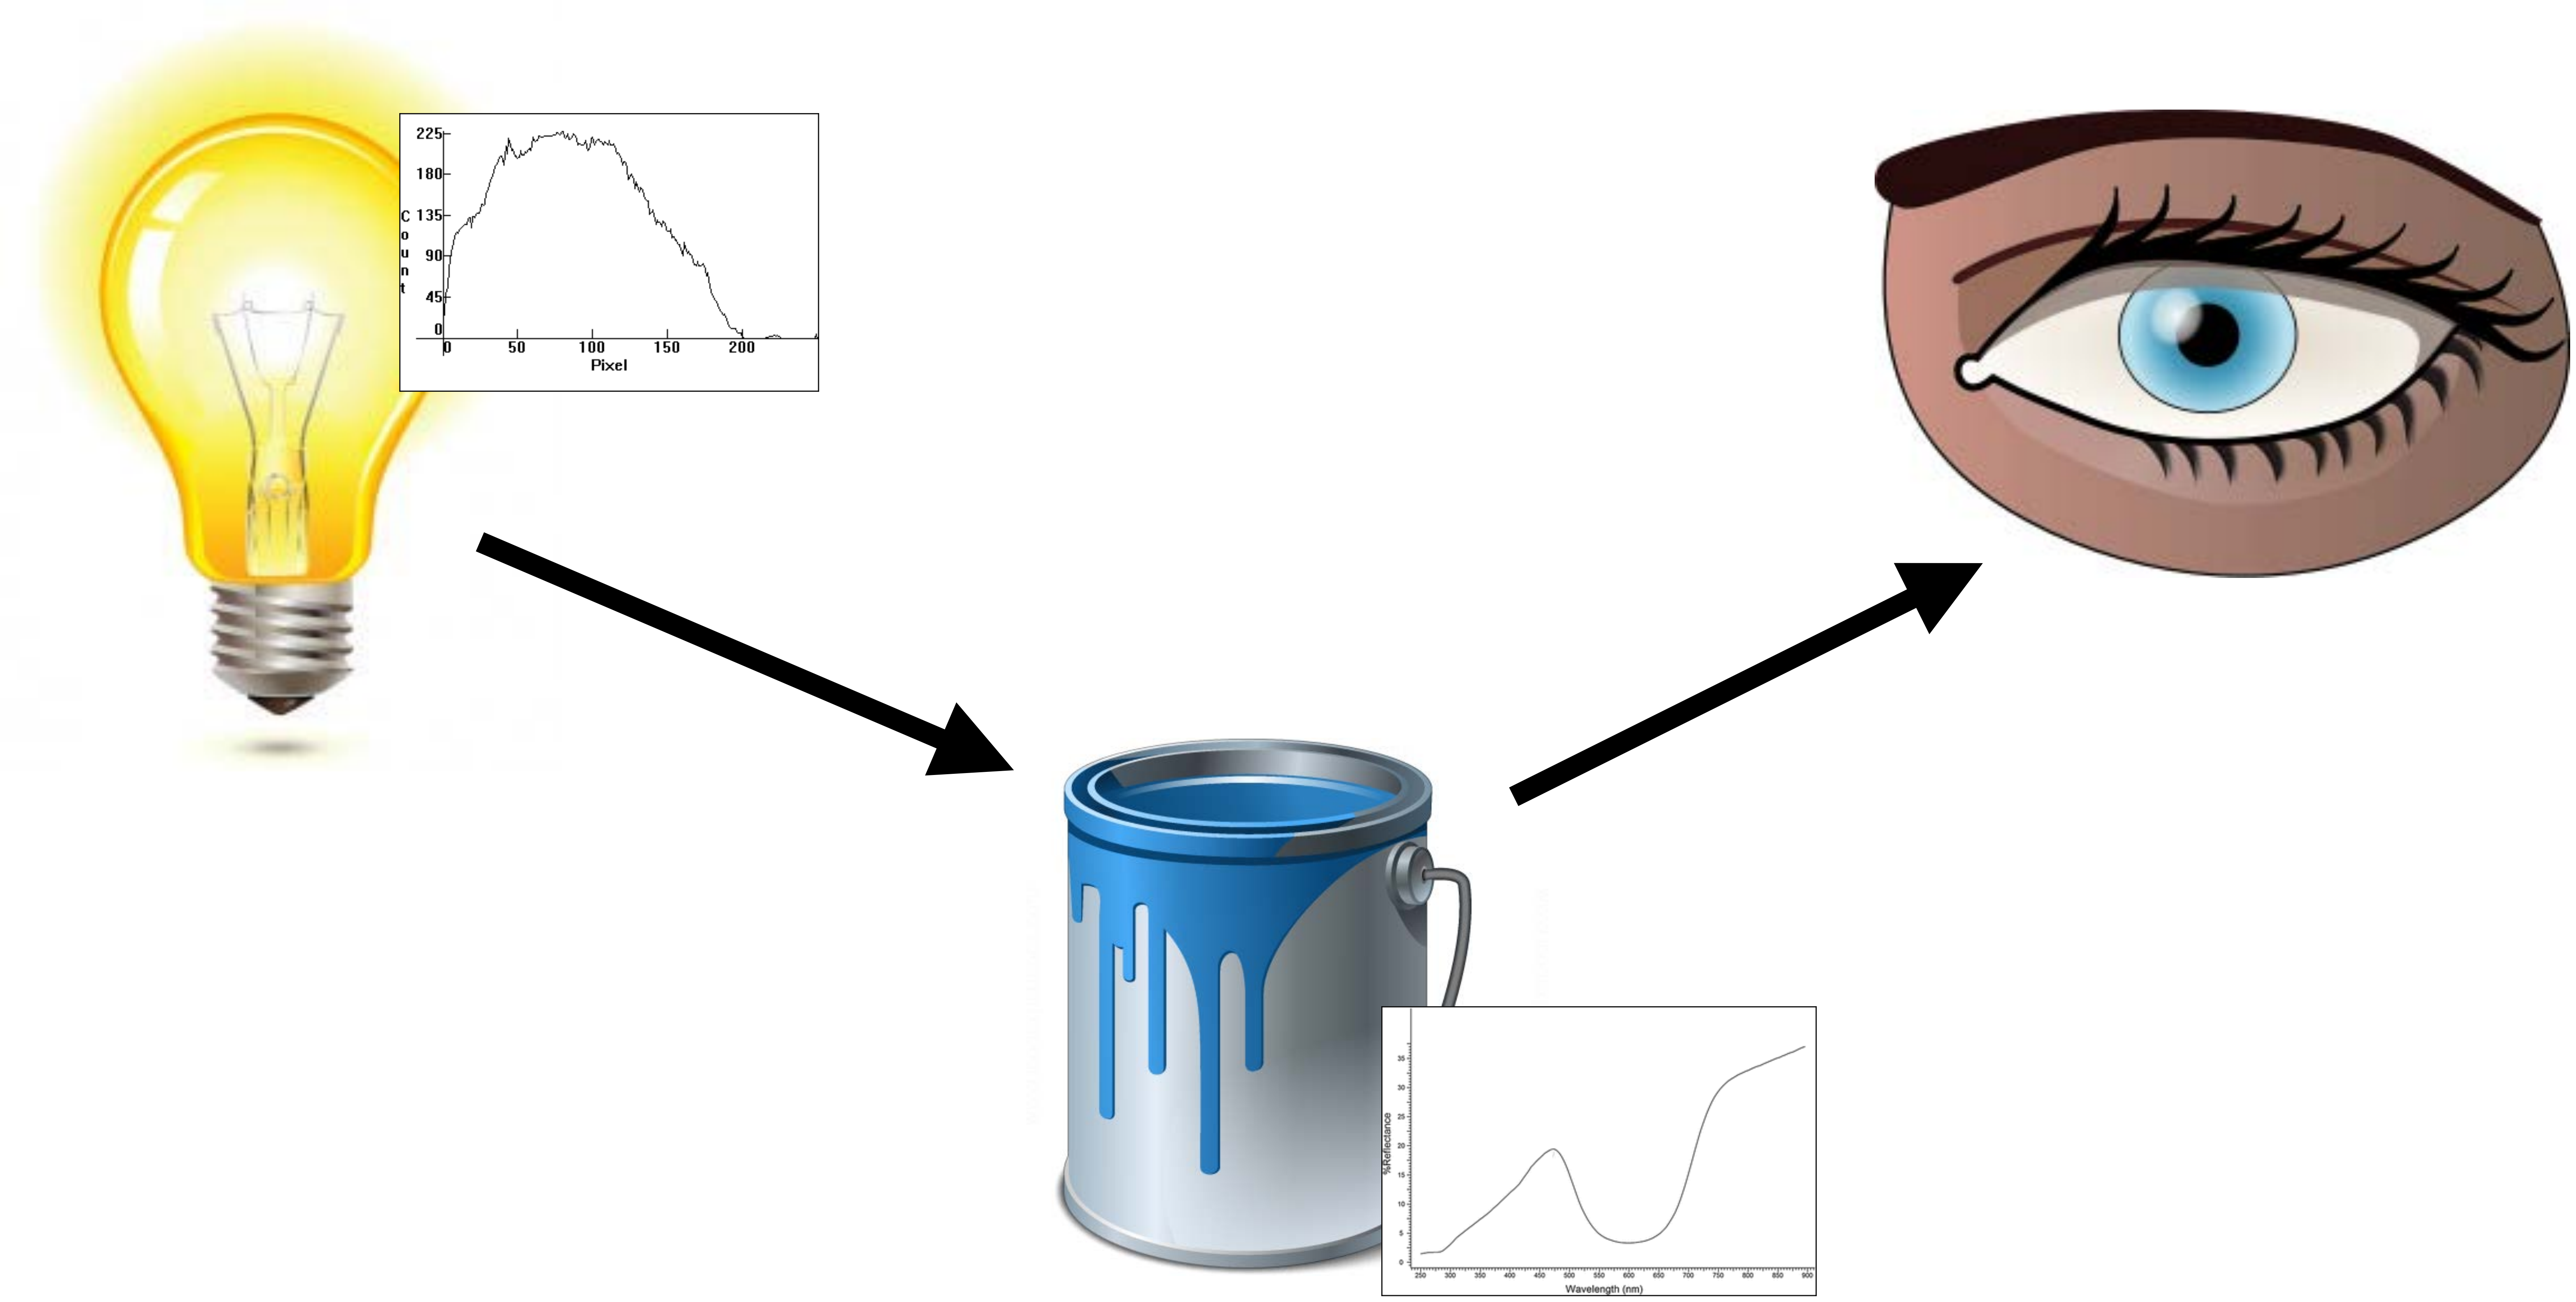

(What color ink should we use to get the desired appearance?)

…And what about perception?

Q: What color is this dress?

How does electromagnetic radiation (with a given power distribution) end up being **perceived by** a human as a certain color?

The eye

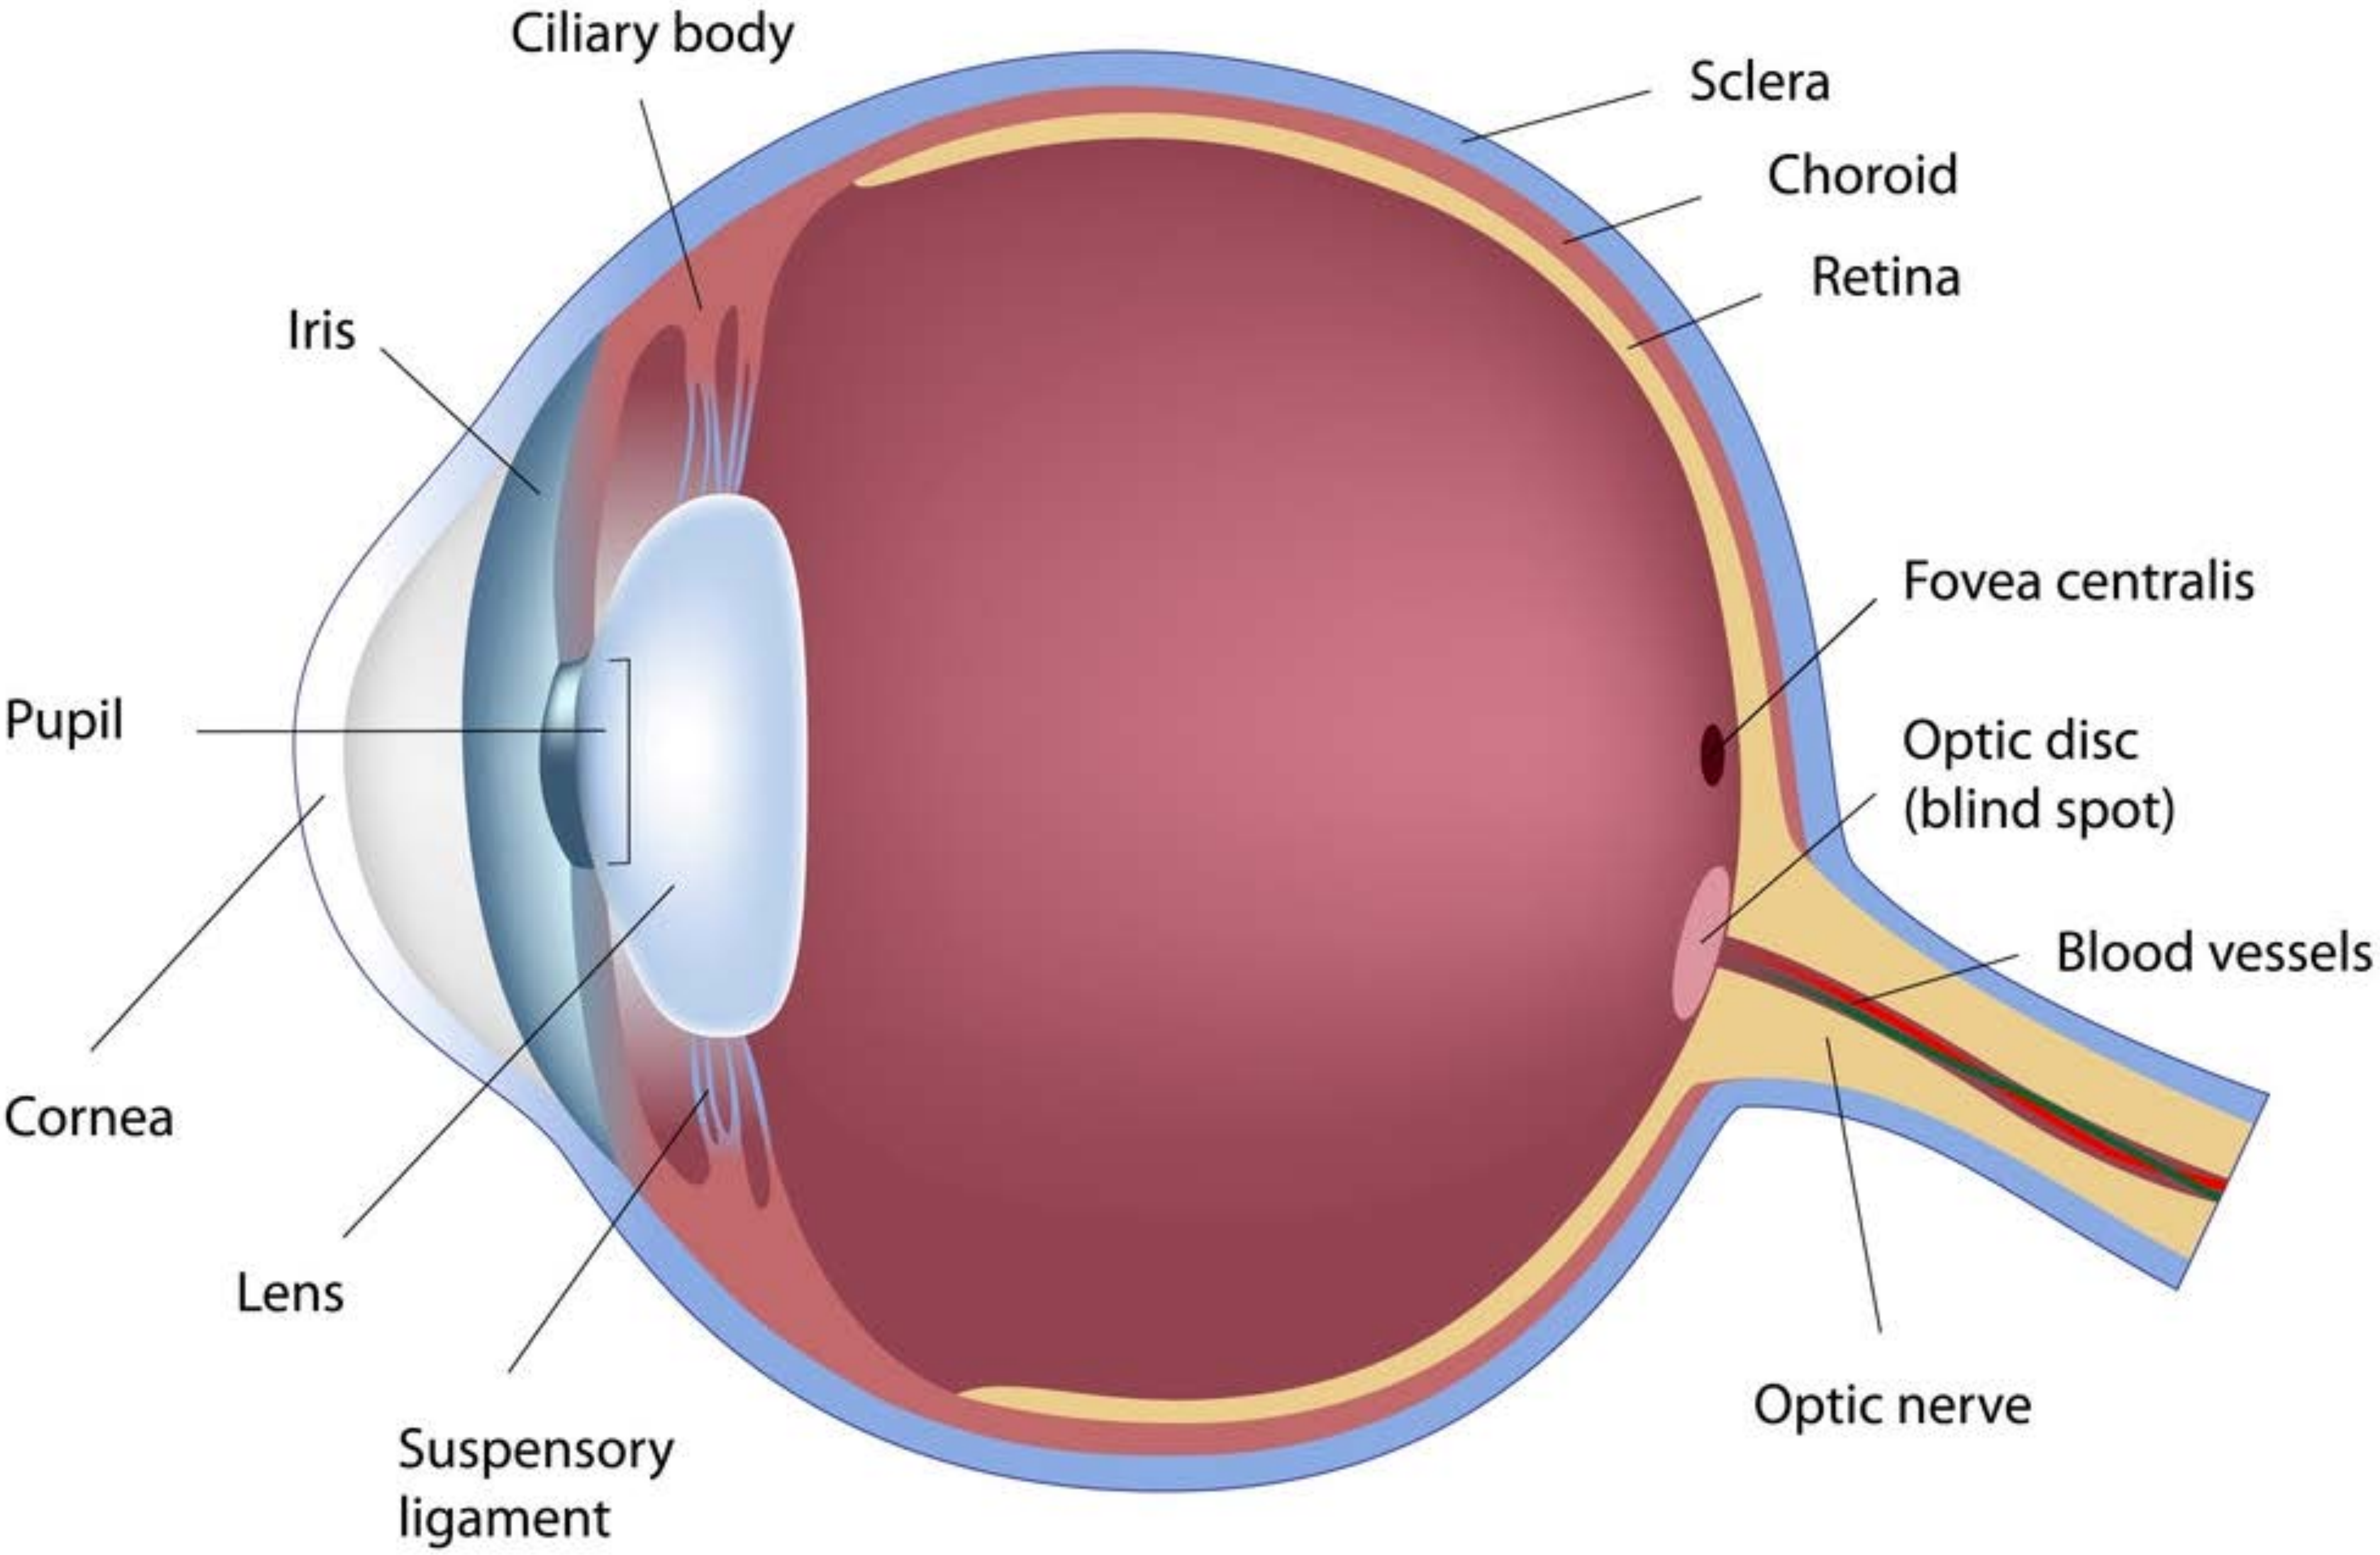

Image credit: Georgia Retina (http://www.garetina.com/about-the-eye)

The eye (optics)

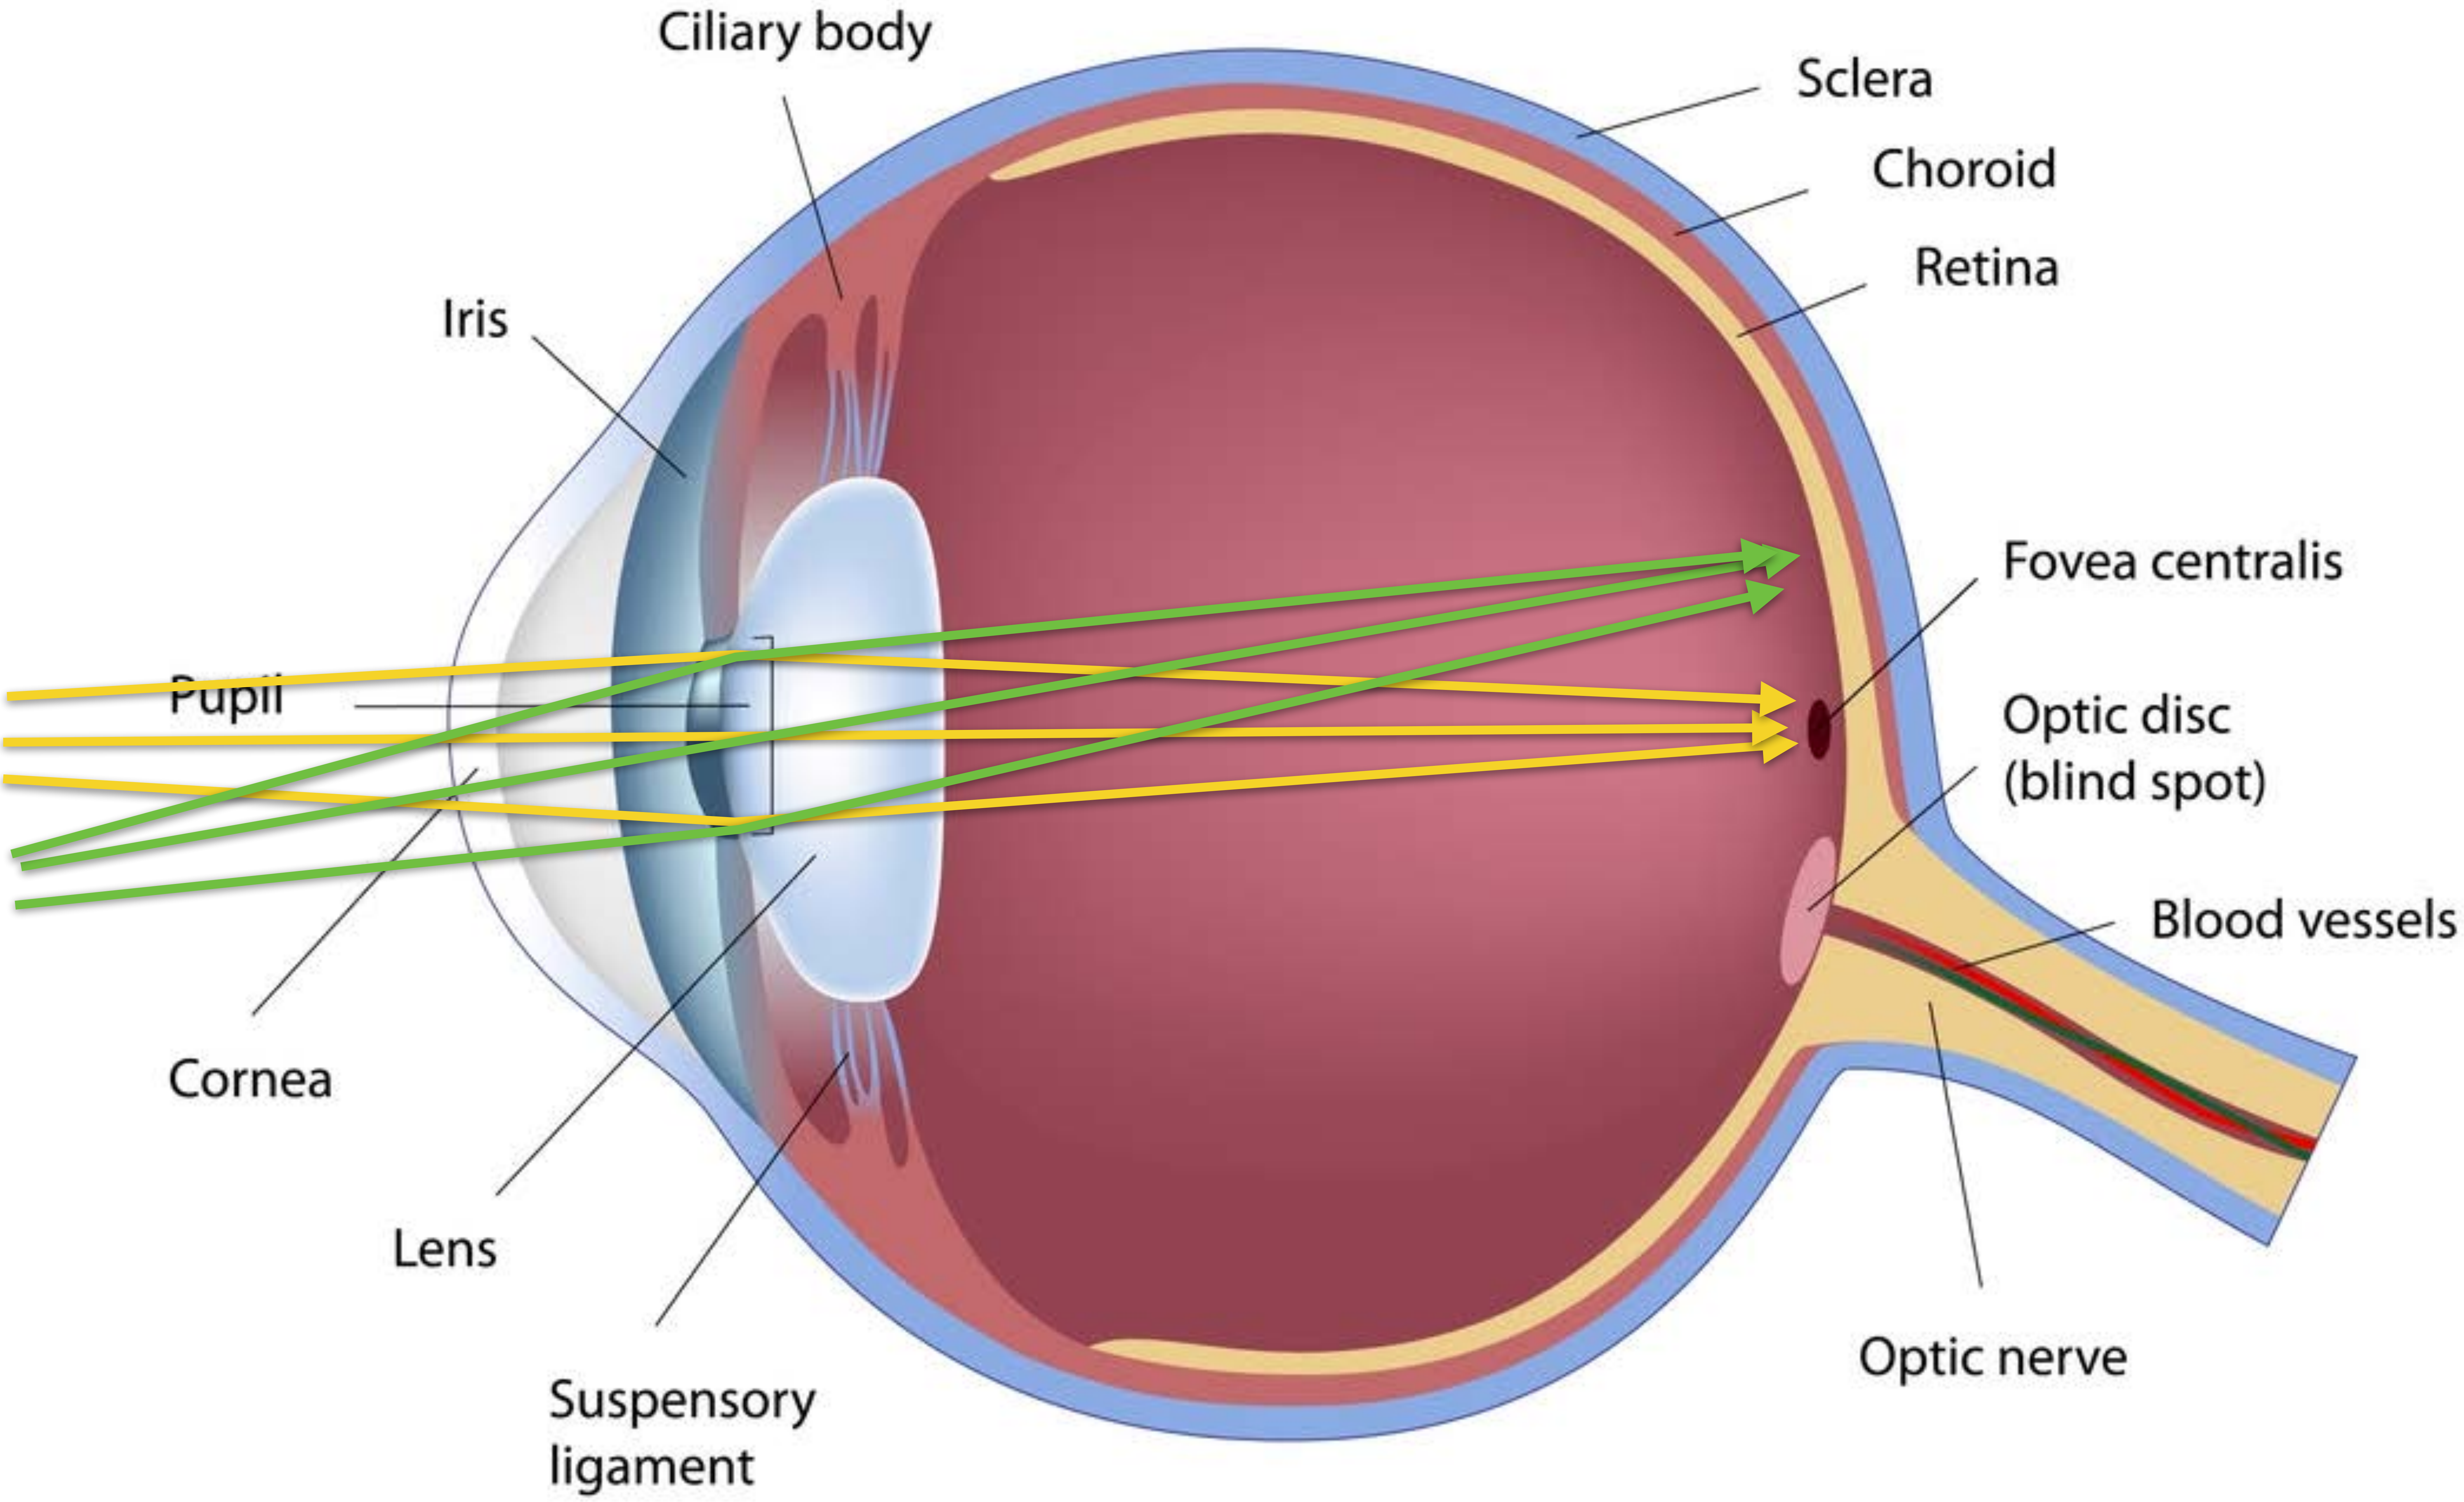

Image credit: Georgia Retina (http://www.garetina.com/about-the-eye)

Photosensor response (eye, camera, …)

- **Photosensor input: light**
  - **Electromagnetic power distribution over wavelengths:** $\Phi(\lambda)$
- **Photosensor output: a "response" … a number**
  - **e.g., encoded in electrical signal**
- **Spectral response function:** $f(\lambda)$
  - **Sensitivity of sensor to light of a given wavelength**
  - **Greater $f(\lambda)$ corresponds to more a efficient sensor (when $f(\lambda)$ is large, a small amount of light at wavelength λ will trigger a large sensor response)**
- **Total response of photosensor:**

$$R = \int_\lambda \Phi(\lambda) f(\lambda) d\lambda$$

The eye's photoreceptor cells: rods & cones

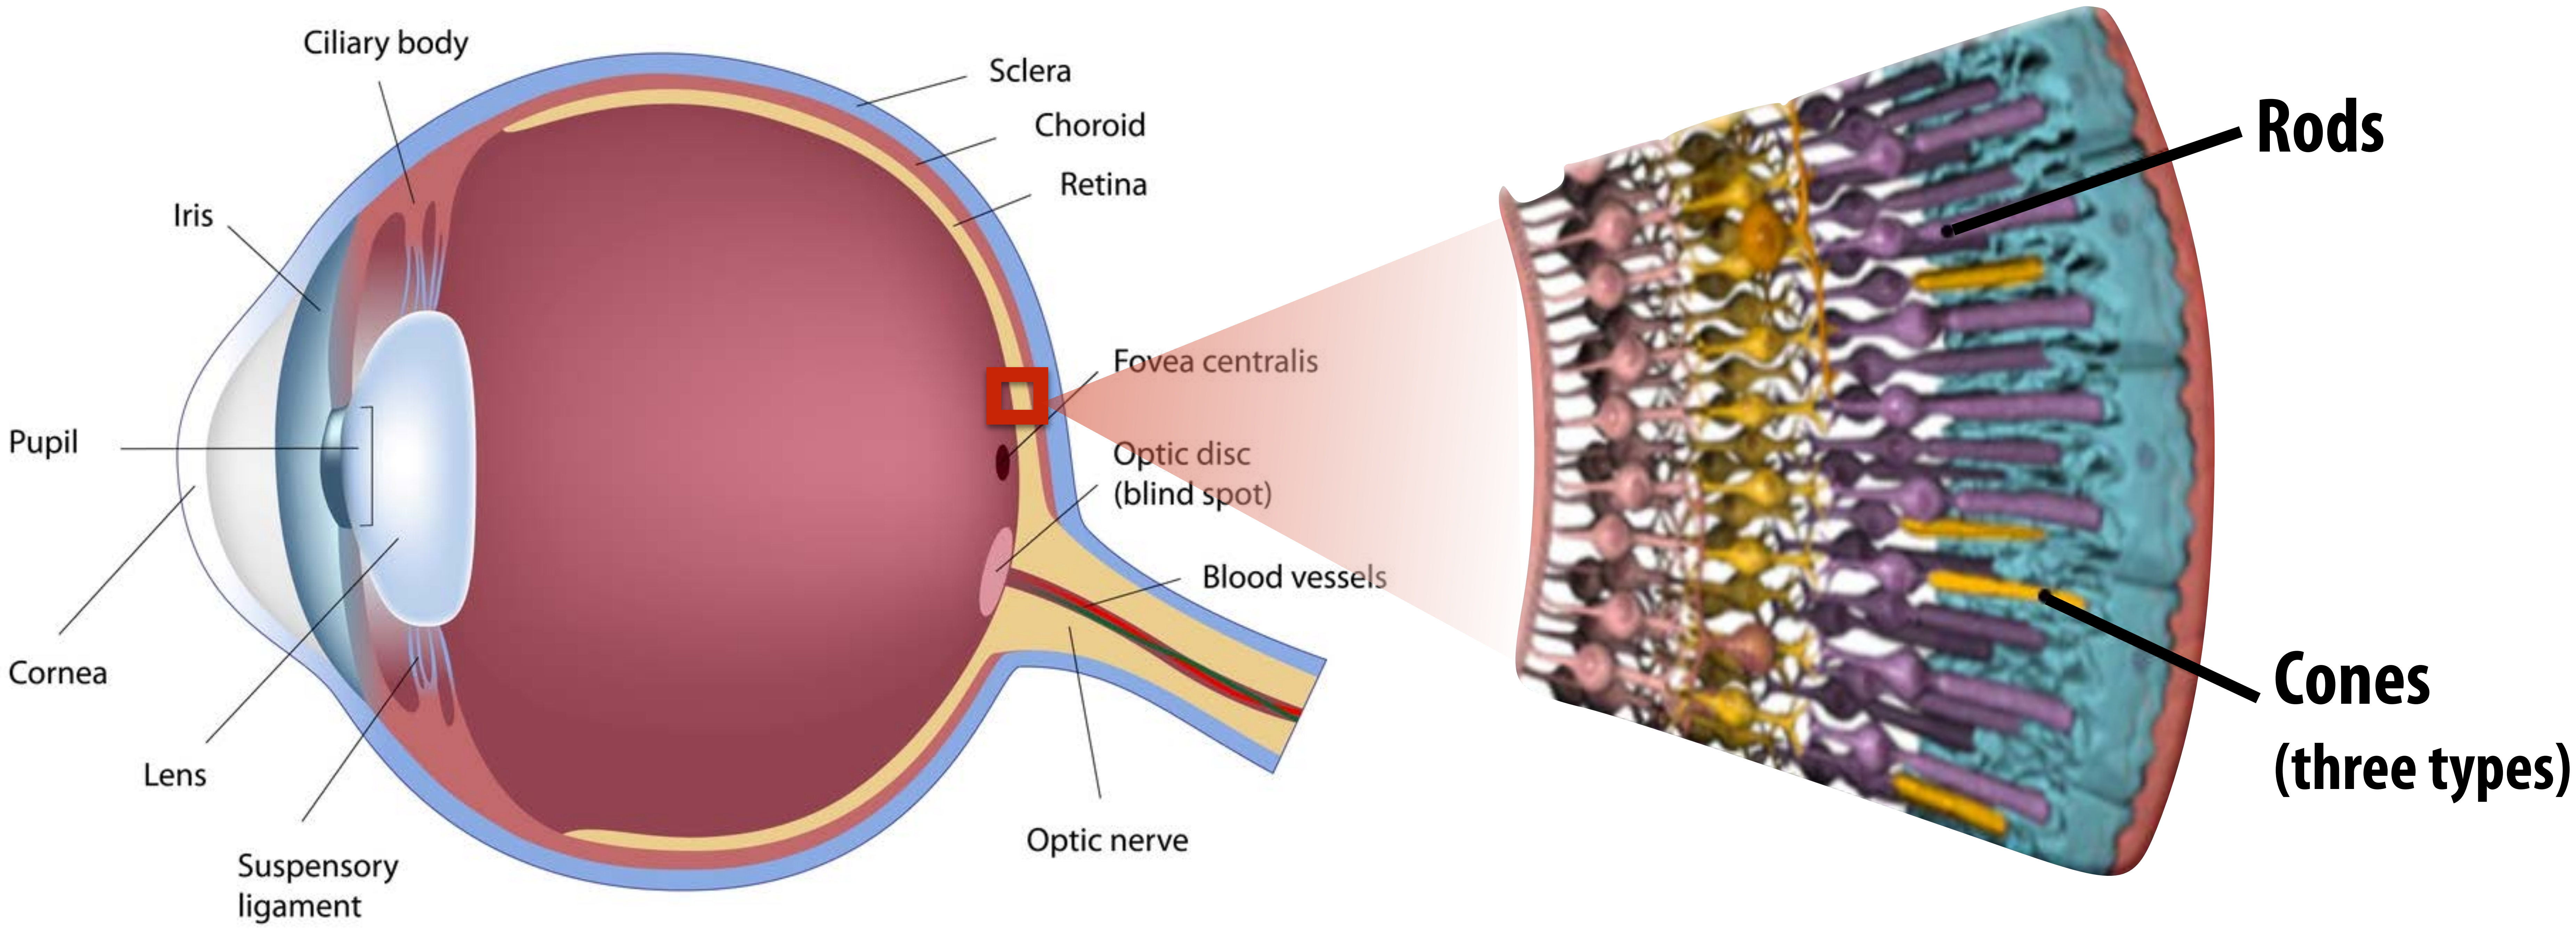

- **Rods are primary receptors under dark viewing conditions (scotopic conditions)**
  - **Approx. 120 million rods in human eye**
- **Cones are primary receptors under high-light viewing conditions (photopic conditions, e.g., daylight)**
  - **Approx. 6-7 million cones in the human eye**
  - **Each of the three types of cone feature a different spectral response. This will be critical to color vision (much more on this in the coming slides)**

Density of rods and cones in the retina

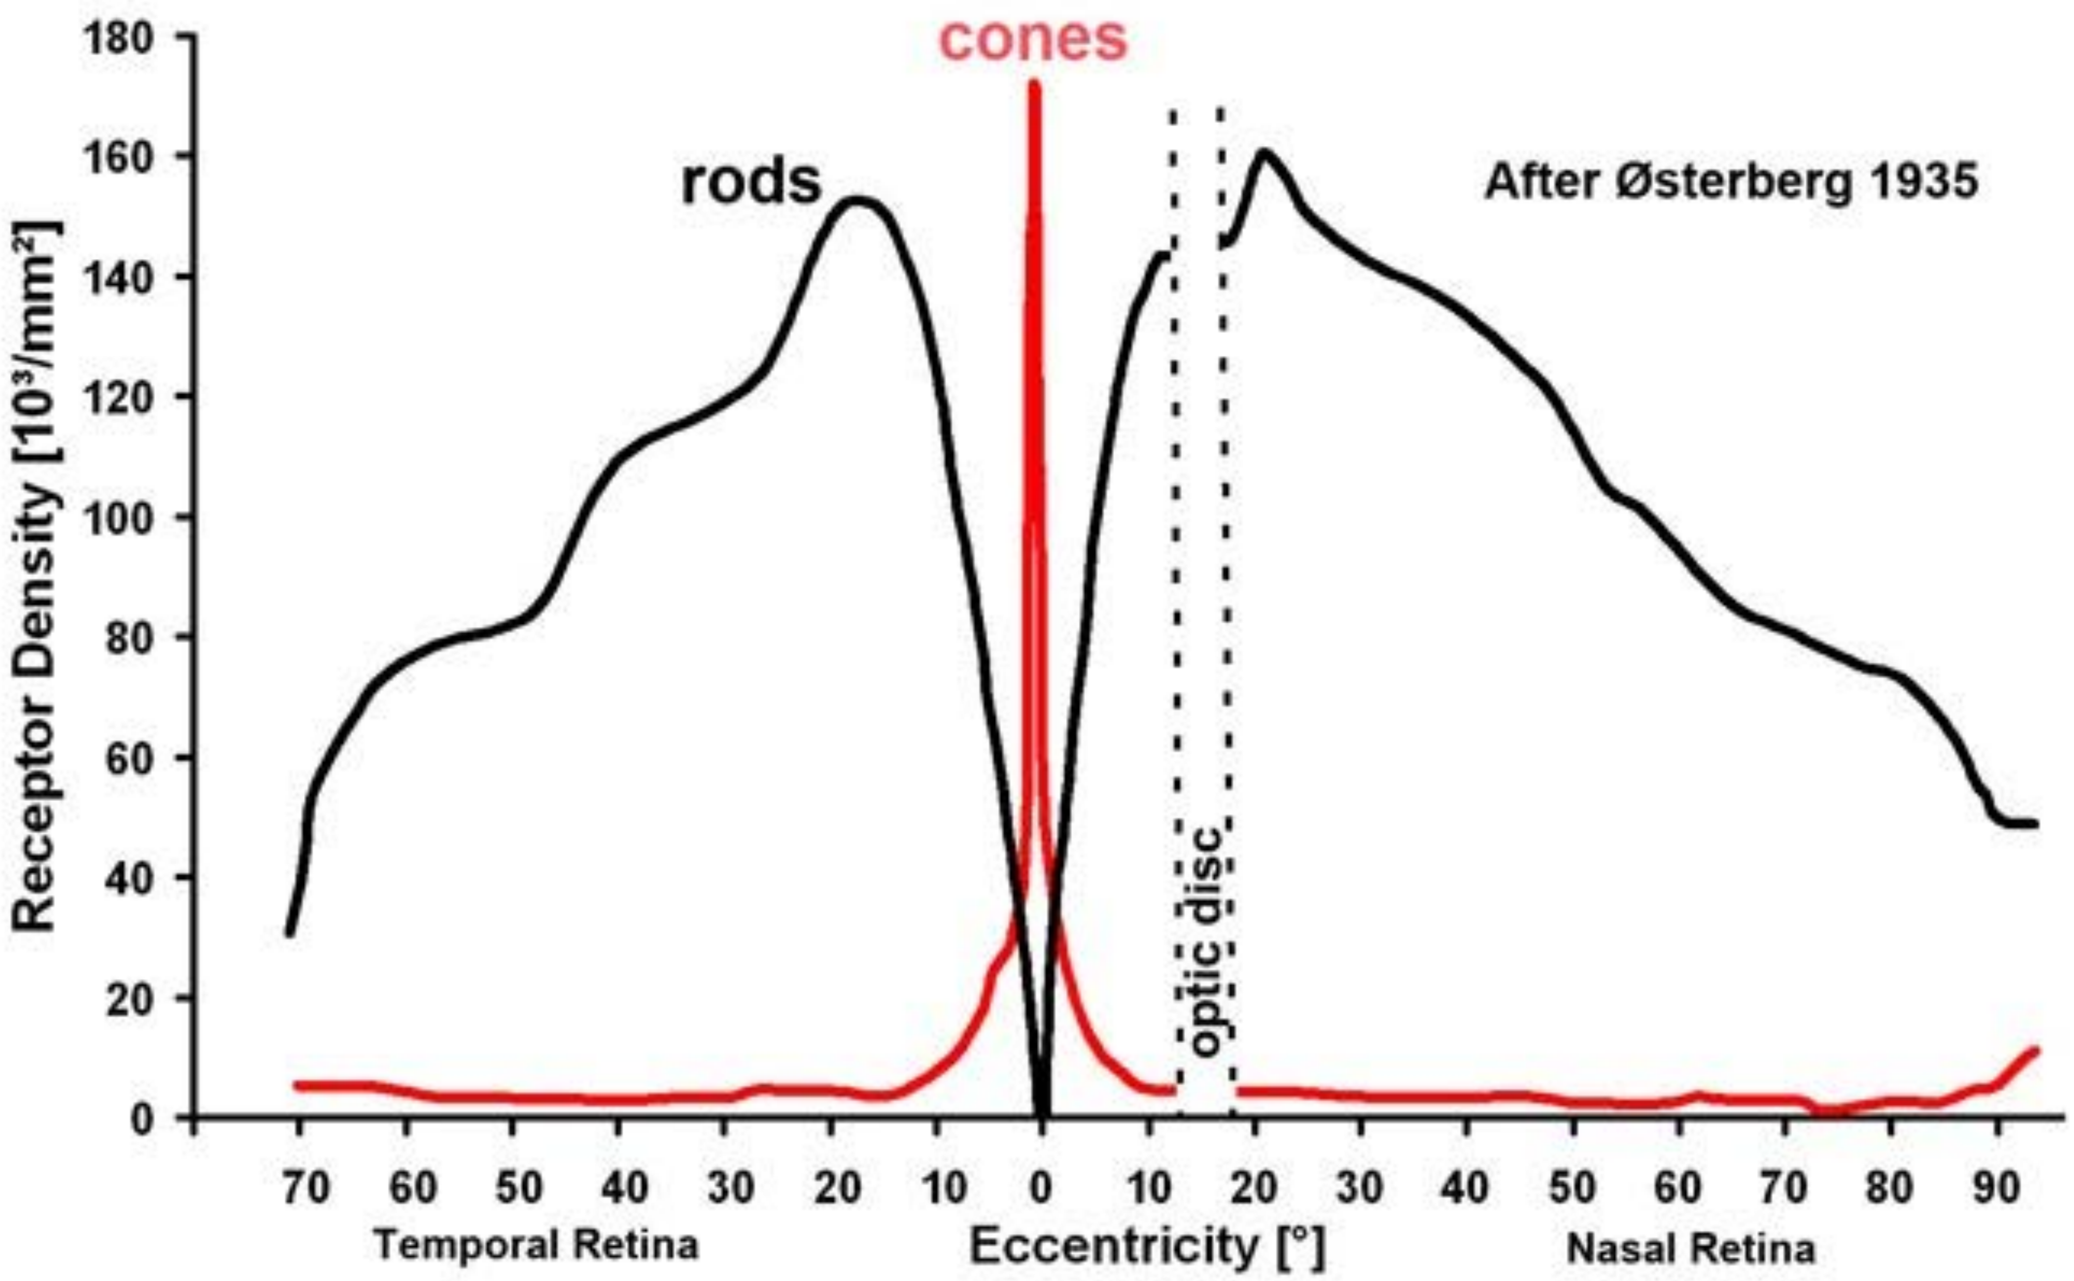

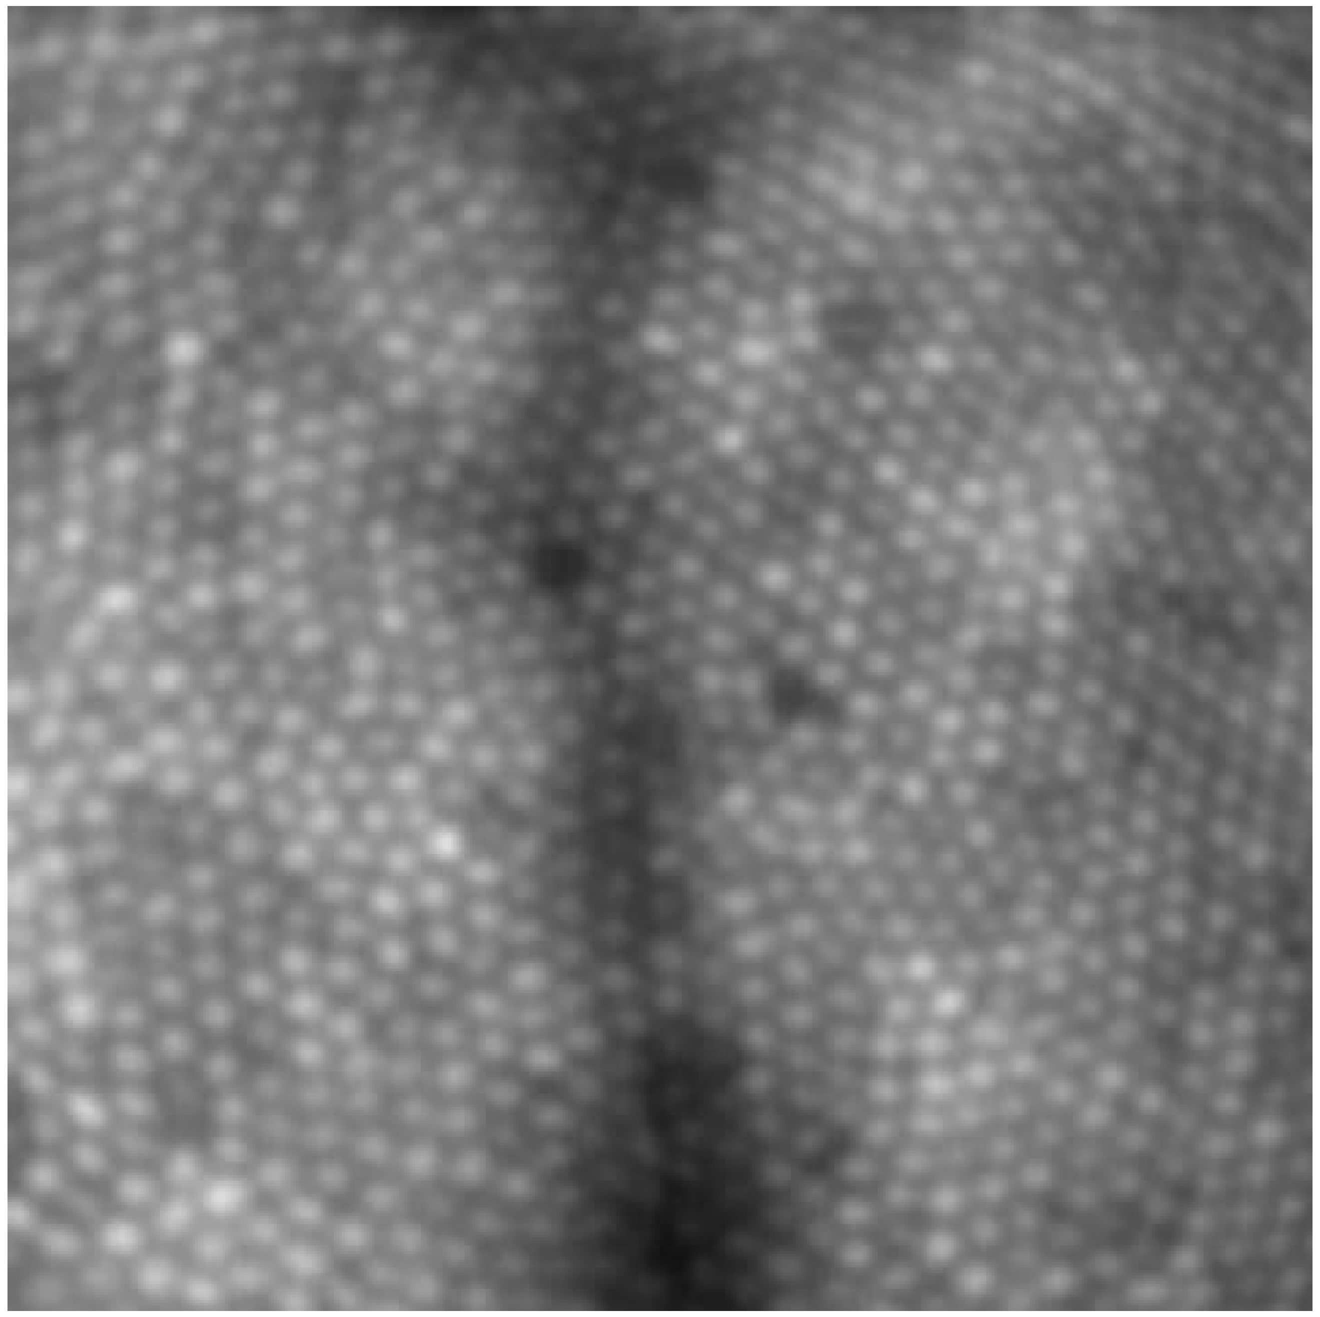

[Roorda 1999]

- Highest density of cones is in fovea
  (best color vision at center of where human is looking)
- Note "blind spot" due to optic nerve

ACTIVITY: Rods vs. Cones

- Grab someone and try it at home!
  - Have them hold up colored markers in peripheral vision
  - All you have to do is say what color it is (easy!)

Spectral response of cones

Three types of cones: S, M, and L cones (corresponding to peak response at short, medium, and long wavelengths)

$$S = \int_\lambda \Phi(\lambda) S(\lambda) d\lambda$$

$$M = \int_\lambda \Phi(\lambda) M(\lambda) d\lambda$$

$$L = \int_\lambda \Phi(\lambda) L(\lambda) d\lambda$$

Response functions for S, M, and L cones

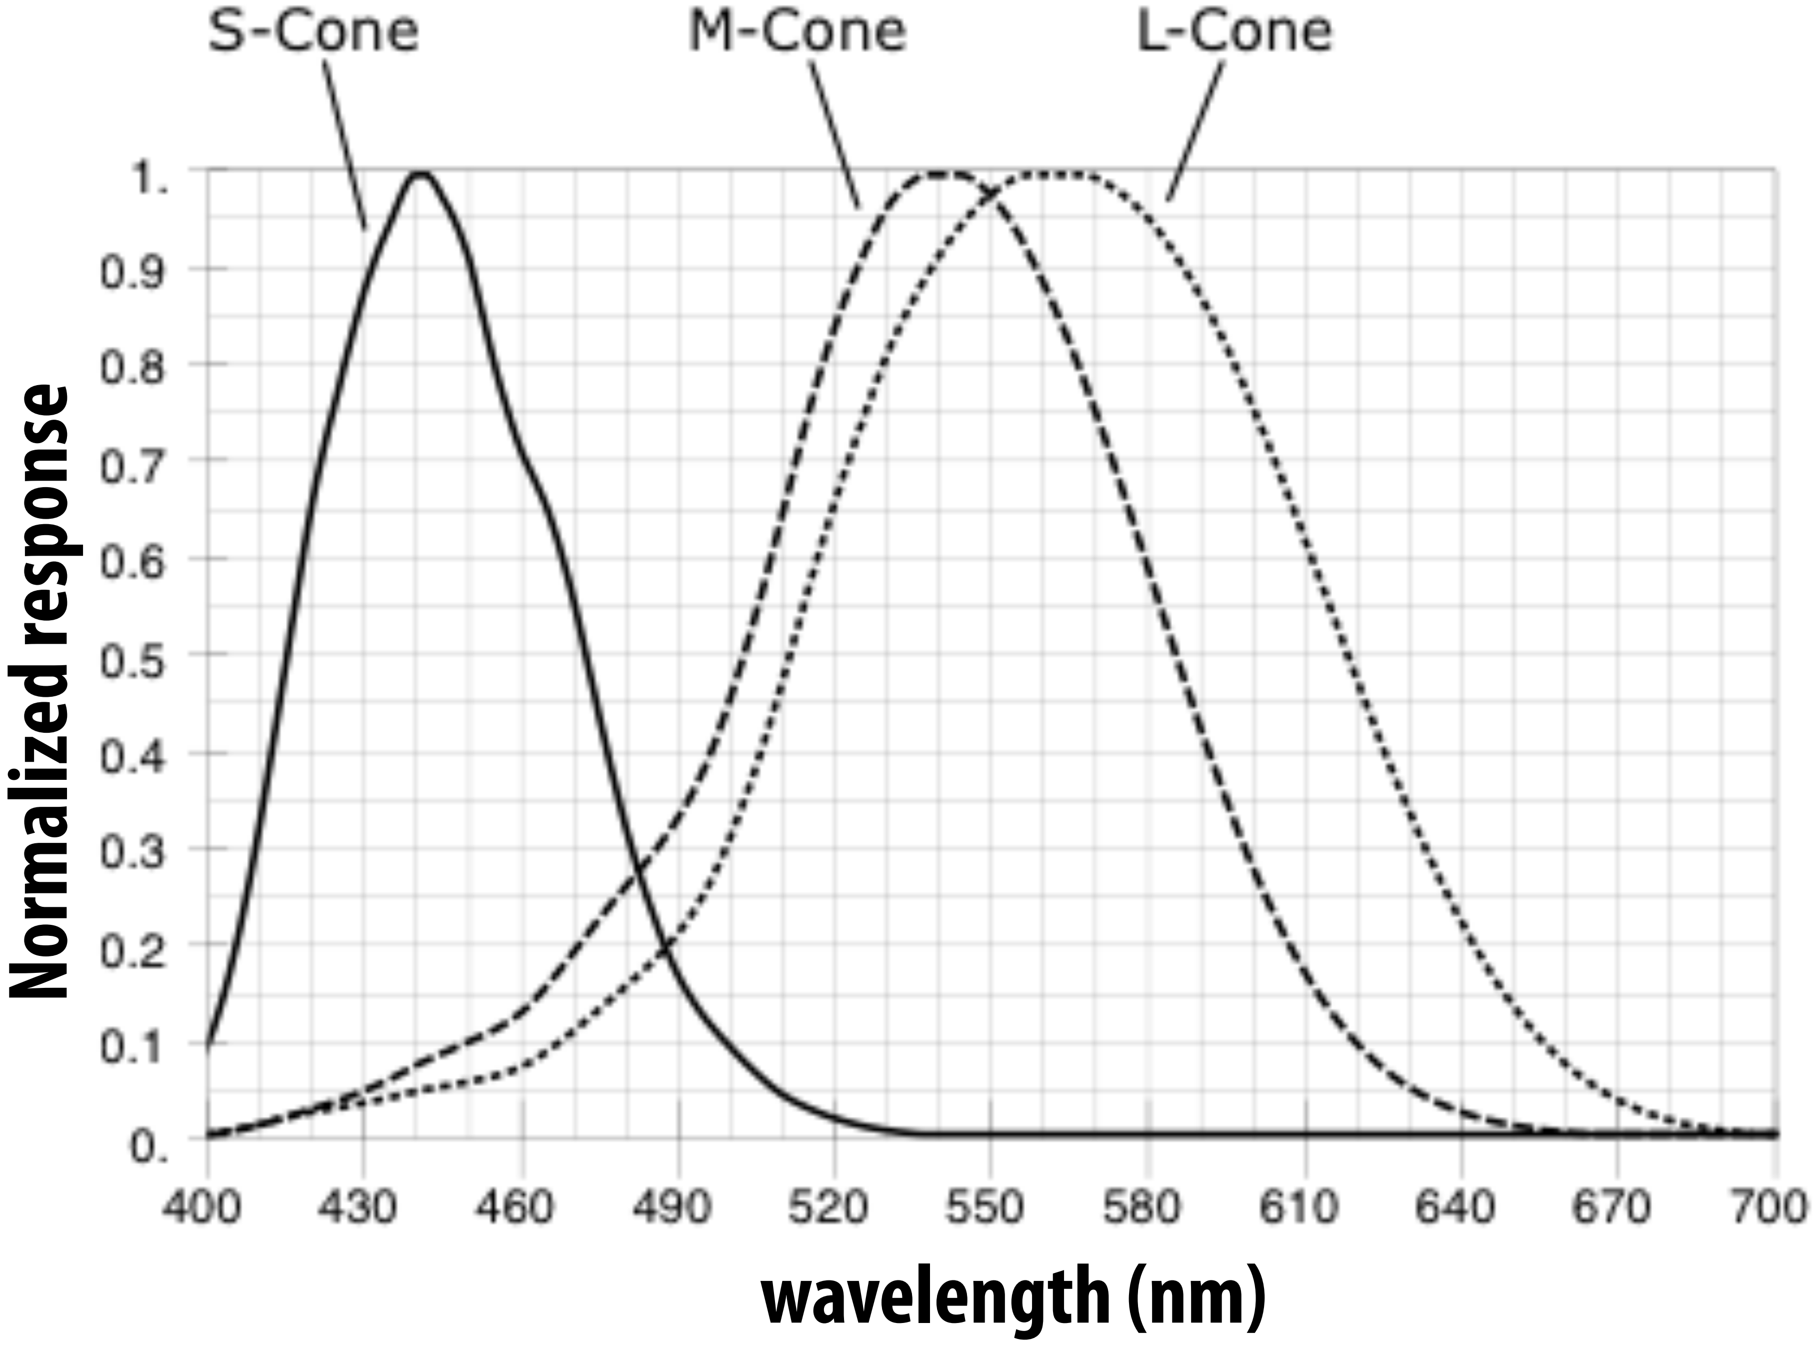

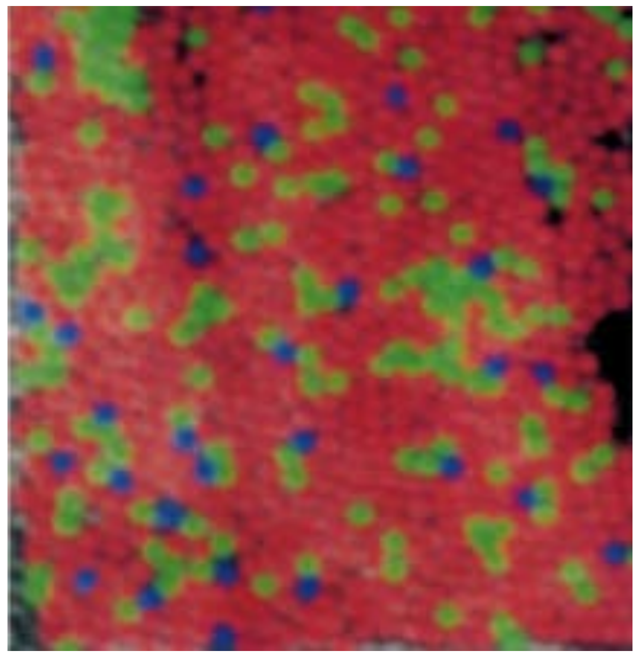

Uneven distribution of cone types in eye
~64% of cones are L cones, ~ 32% M cones

Response of S,M,L cones to monochromatic light

Figure visualizes cone's response to monochromatic light (light with energy in a single wavelength) as points in 3D space

(plots value of S, M, L response functions as a point in 3D space)

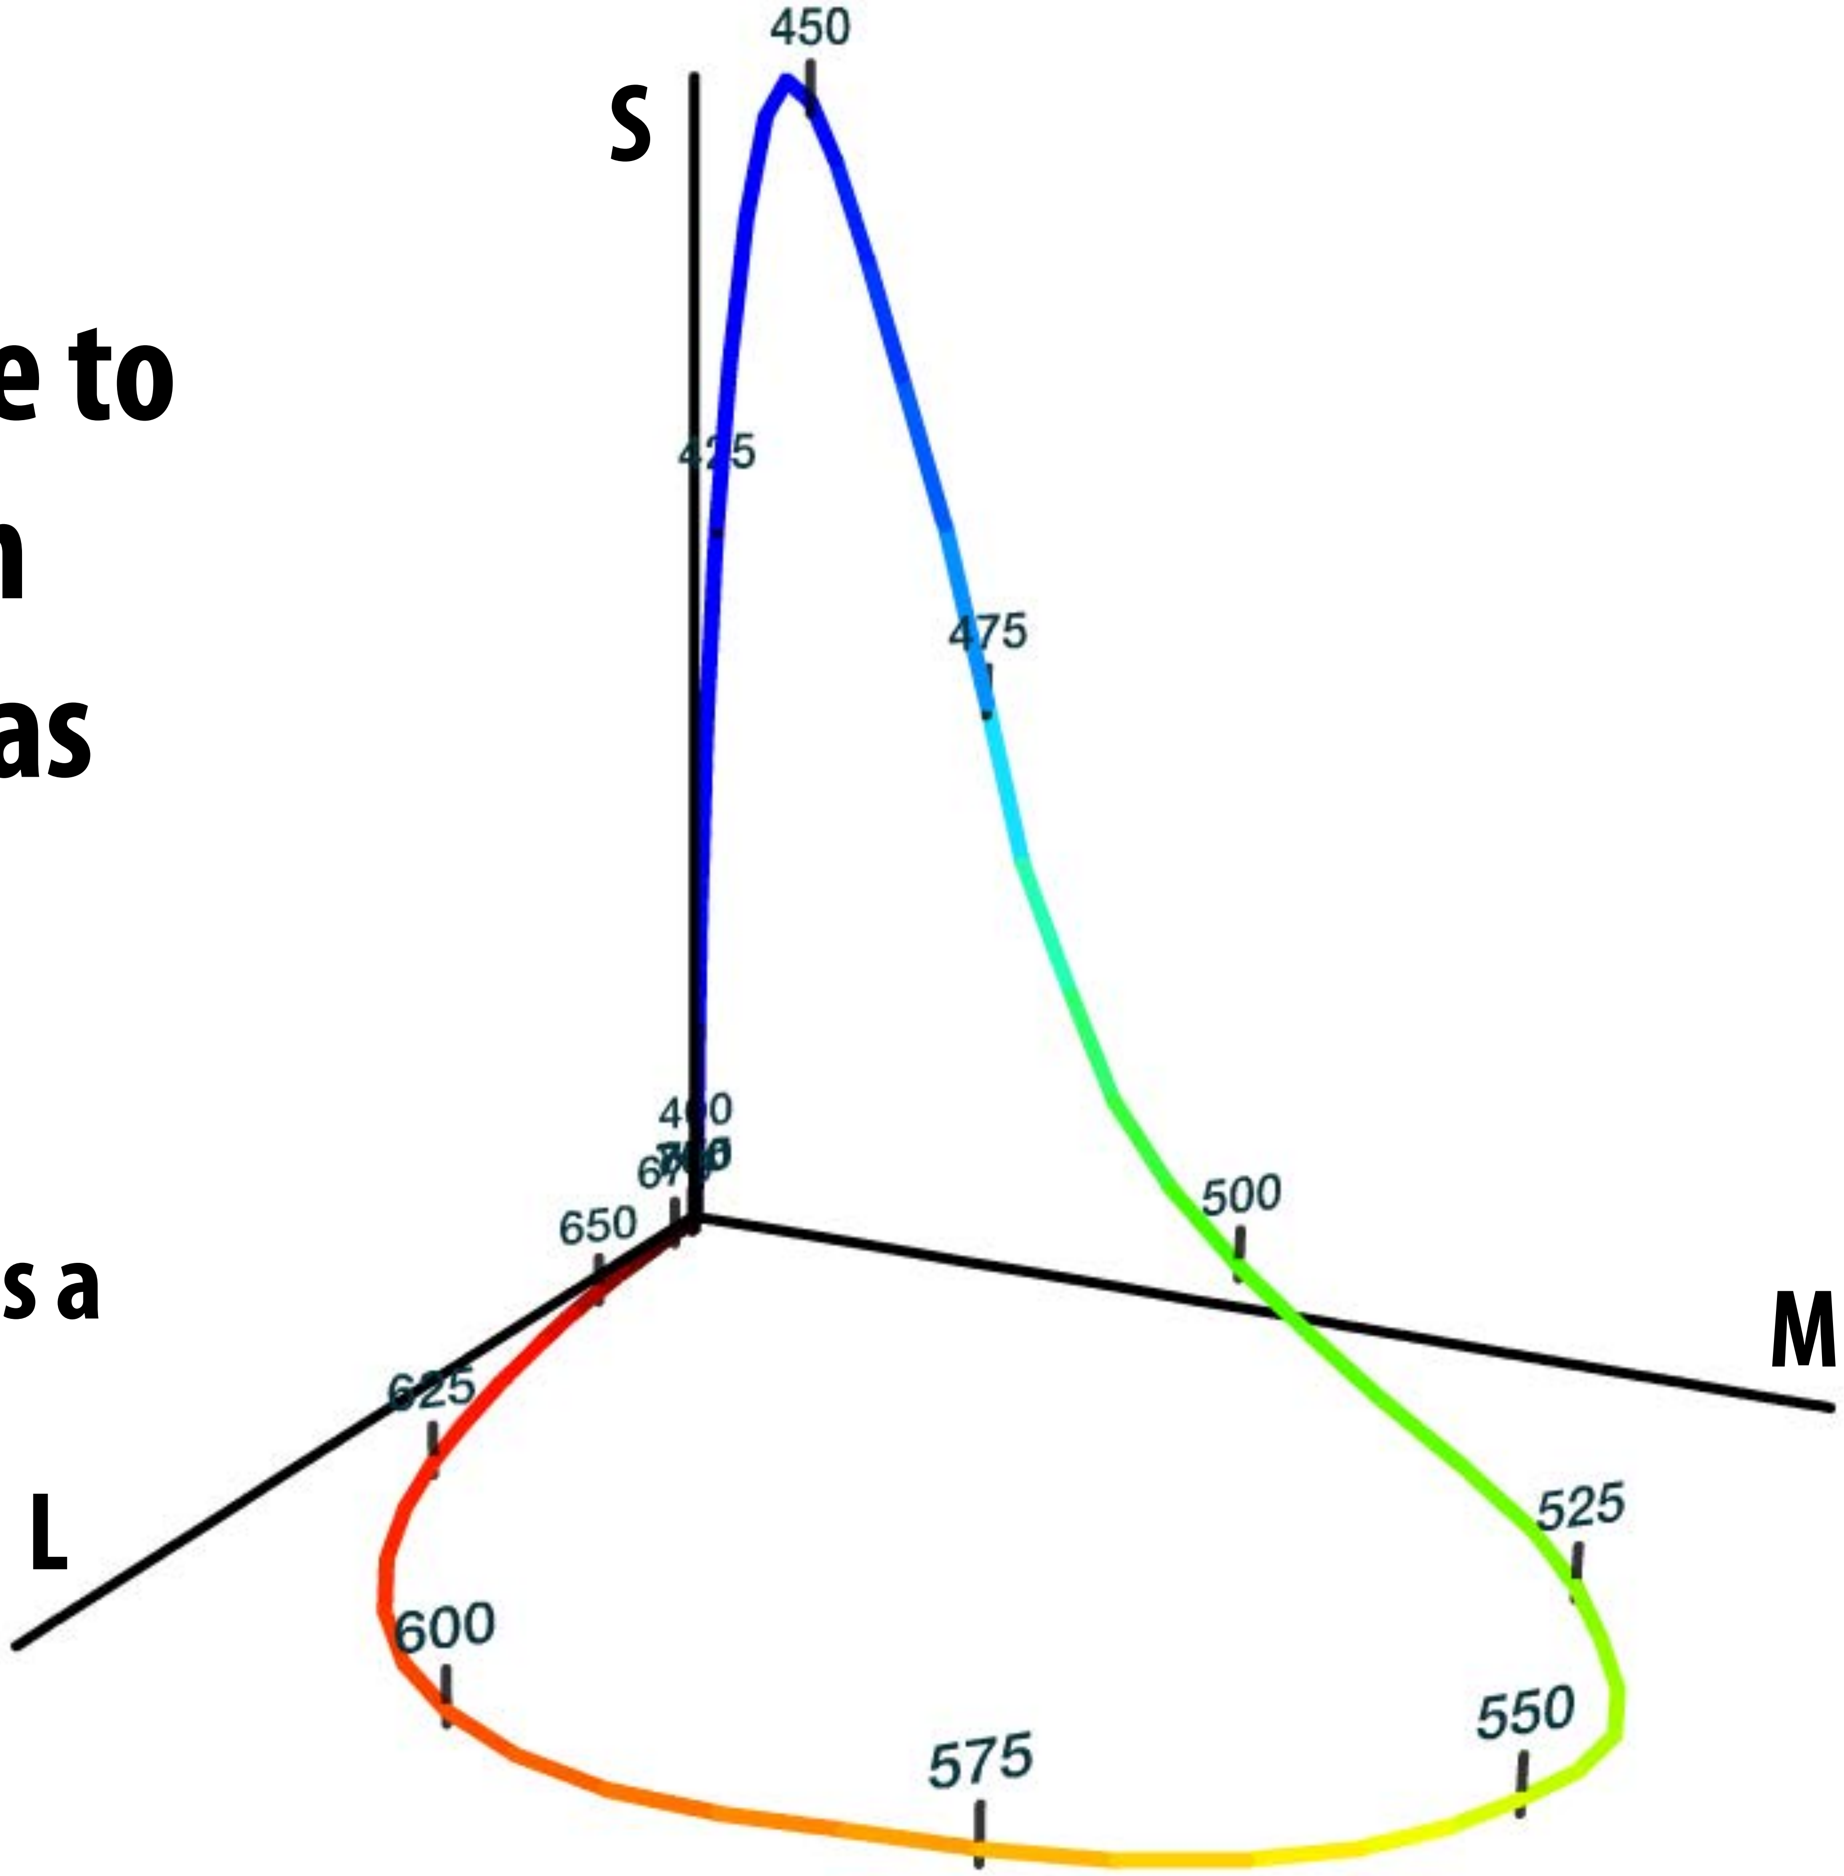

The human visual system

- **Human eye does not directly measure the spectrum of incoming light**
  - **i.e., the brain does not receive "a spectrum" from the eye**
- **The eye measures three response values = (S, M, L). The result of integrating the incoming spectrum against response functions of S, M, L-cones**

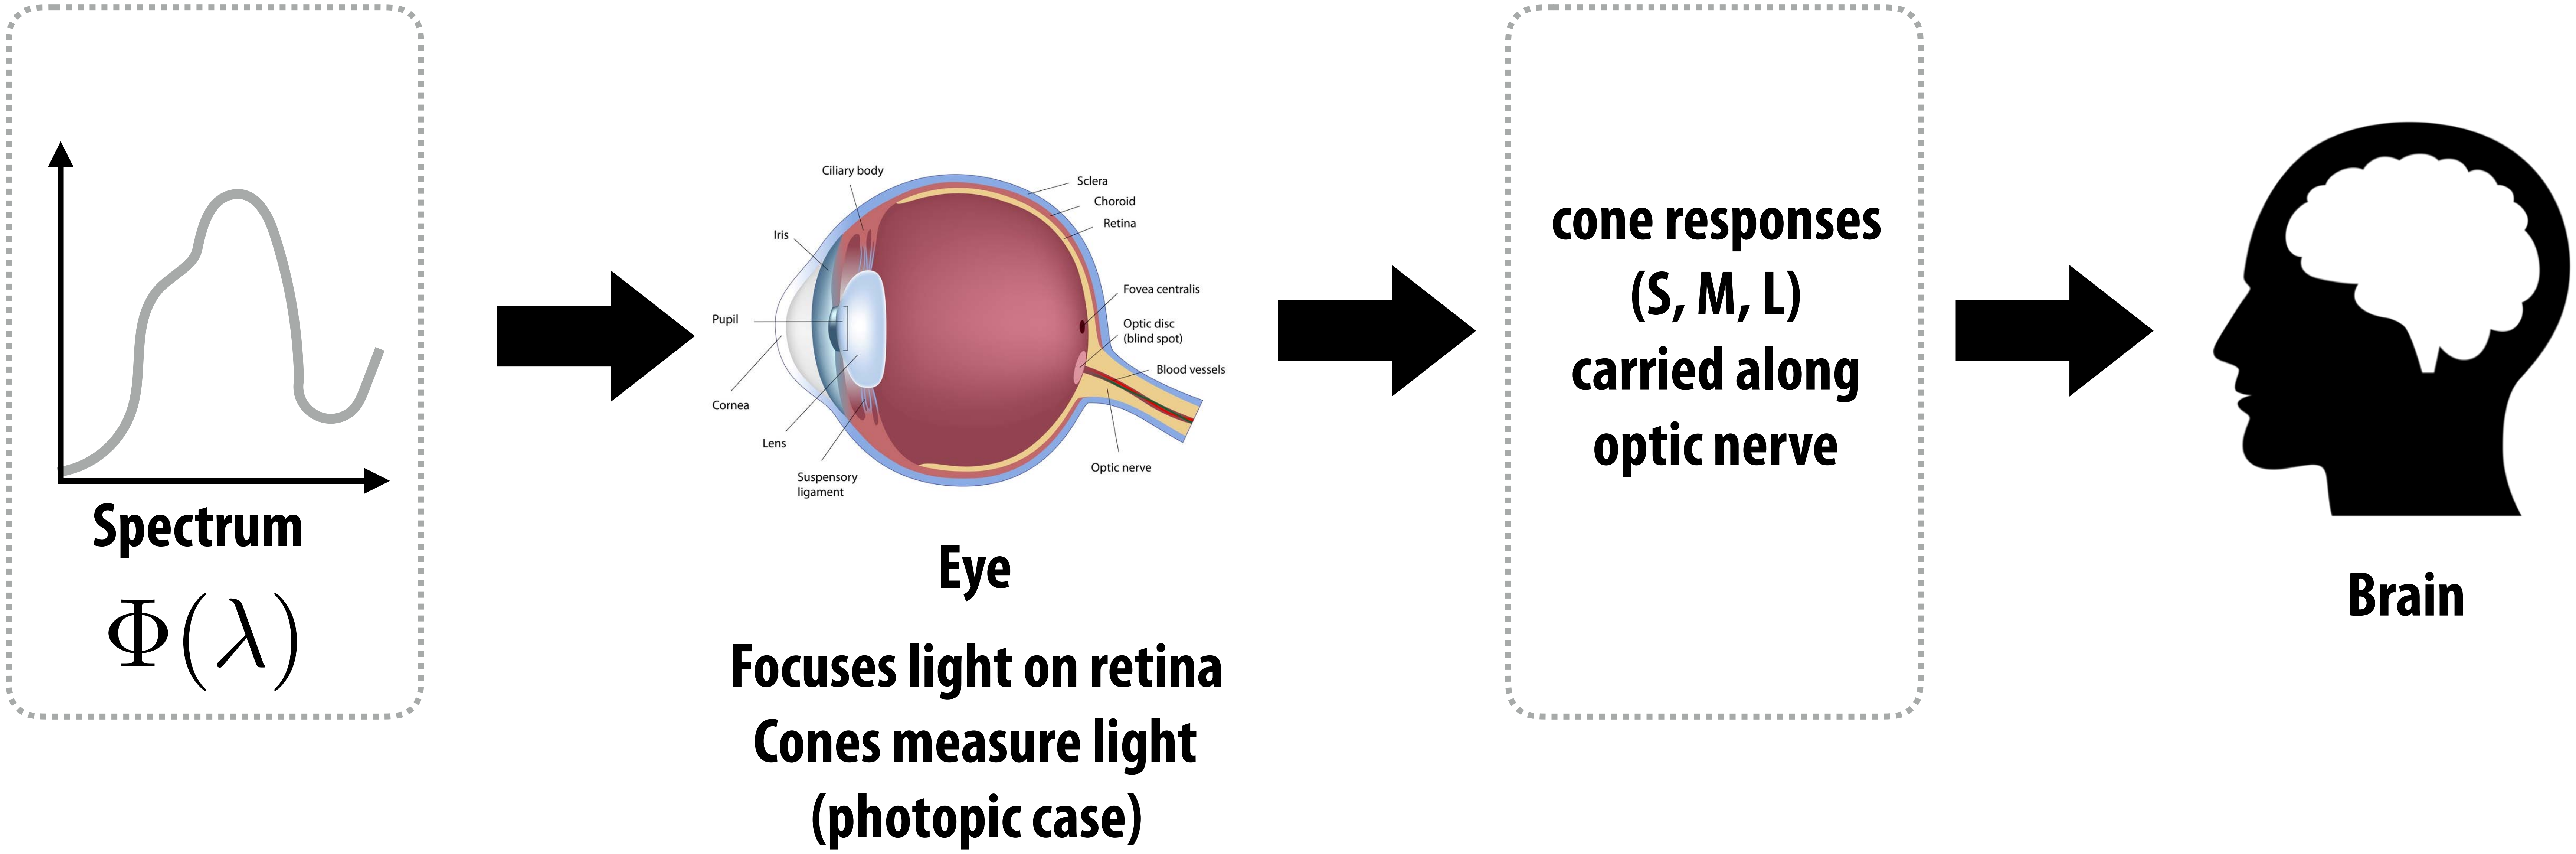

Spectrum

$\Phi(\lambda)$

Eye

Focuses light on retina
Cones measure light
(photopic case)

cone responses
(S, M, L)
carried along
optic nerve

Brain

Q: Is it possible for two functions to integrate to the same value?

Metamers

- Metamers = two different spectra that integrate to the same (S,M,L) response!
- The fact that metameters exist is critical to color reproduction: we don't have to reproduce the exact same spectrum that was present in a real world scene in order to reproduce the *perceived* color on a monitor (or piece of paper, or paint on a wall)
- …On the other hand, combination of light & paint could still cause trouble—different objects appearing "wrong" under different lighting conditions.

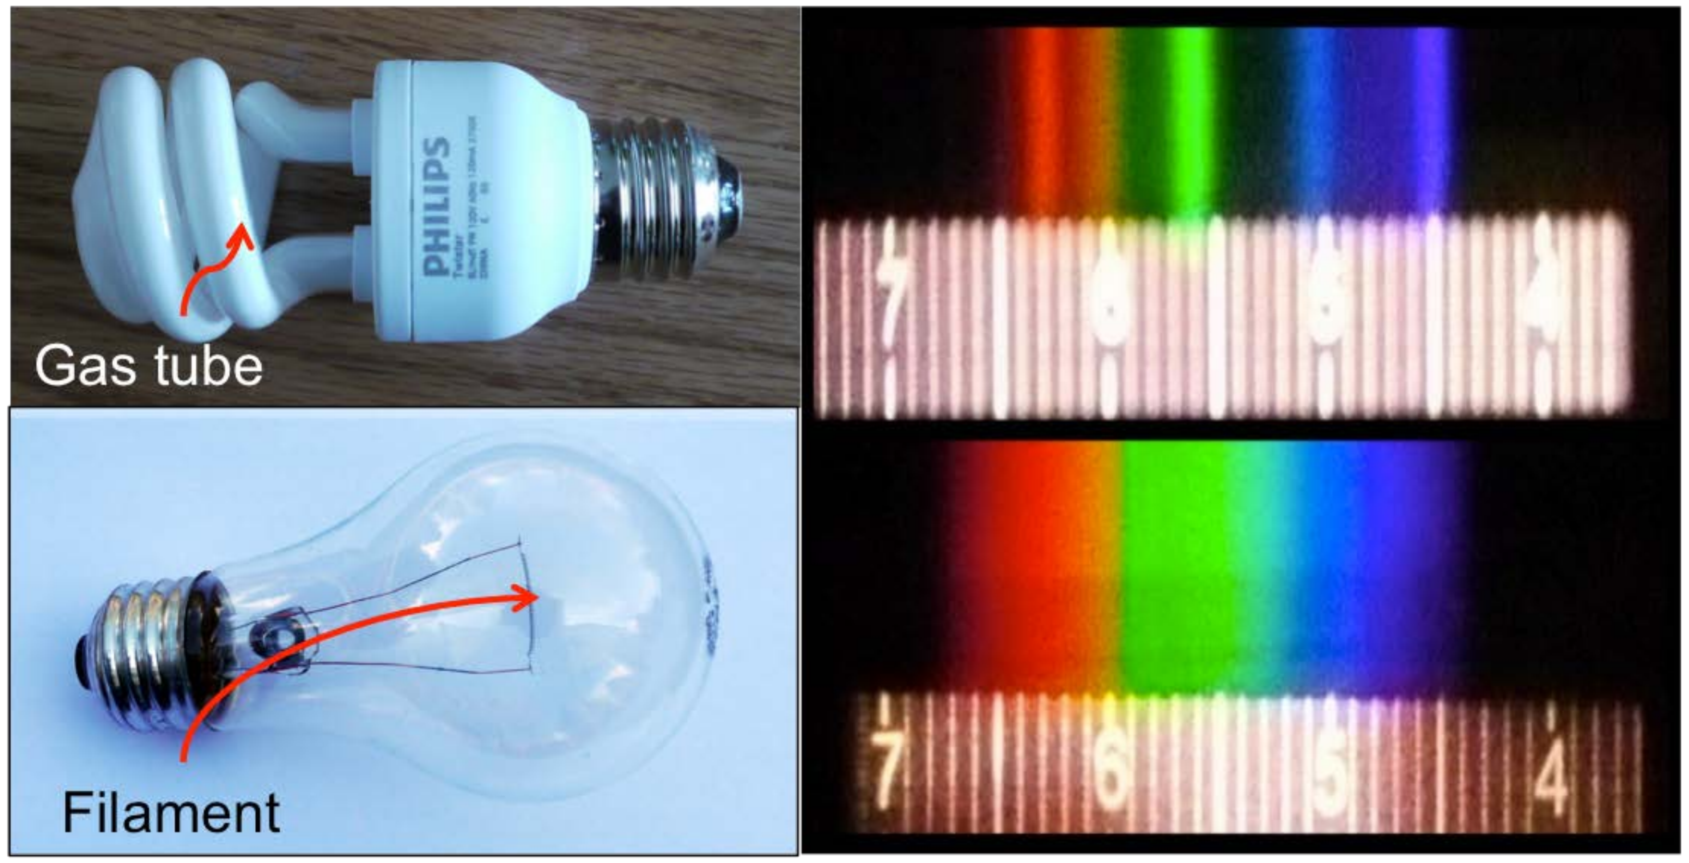

Example: Counterfeit Detection

- **Many countries print currency, passports, etc., with special inks that yield different appearance under UV light:**

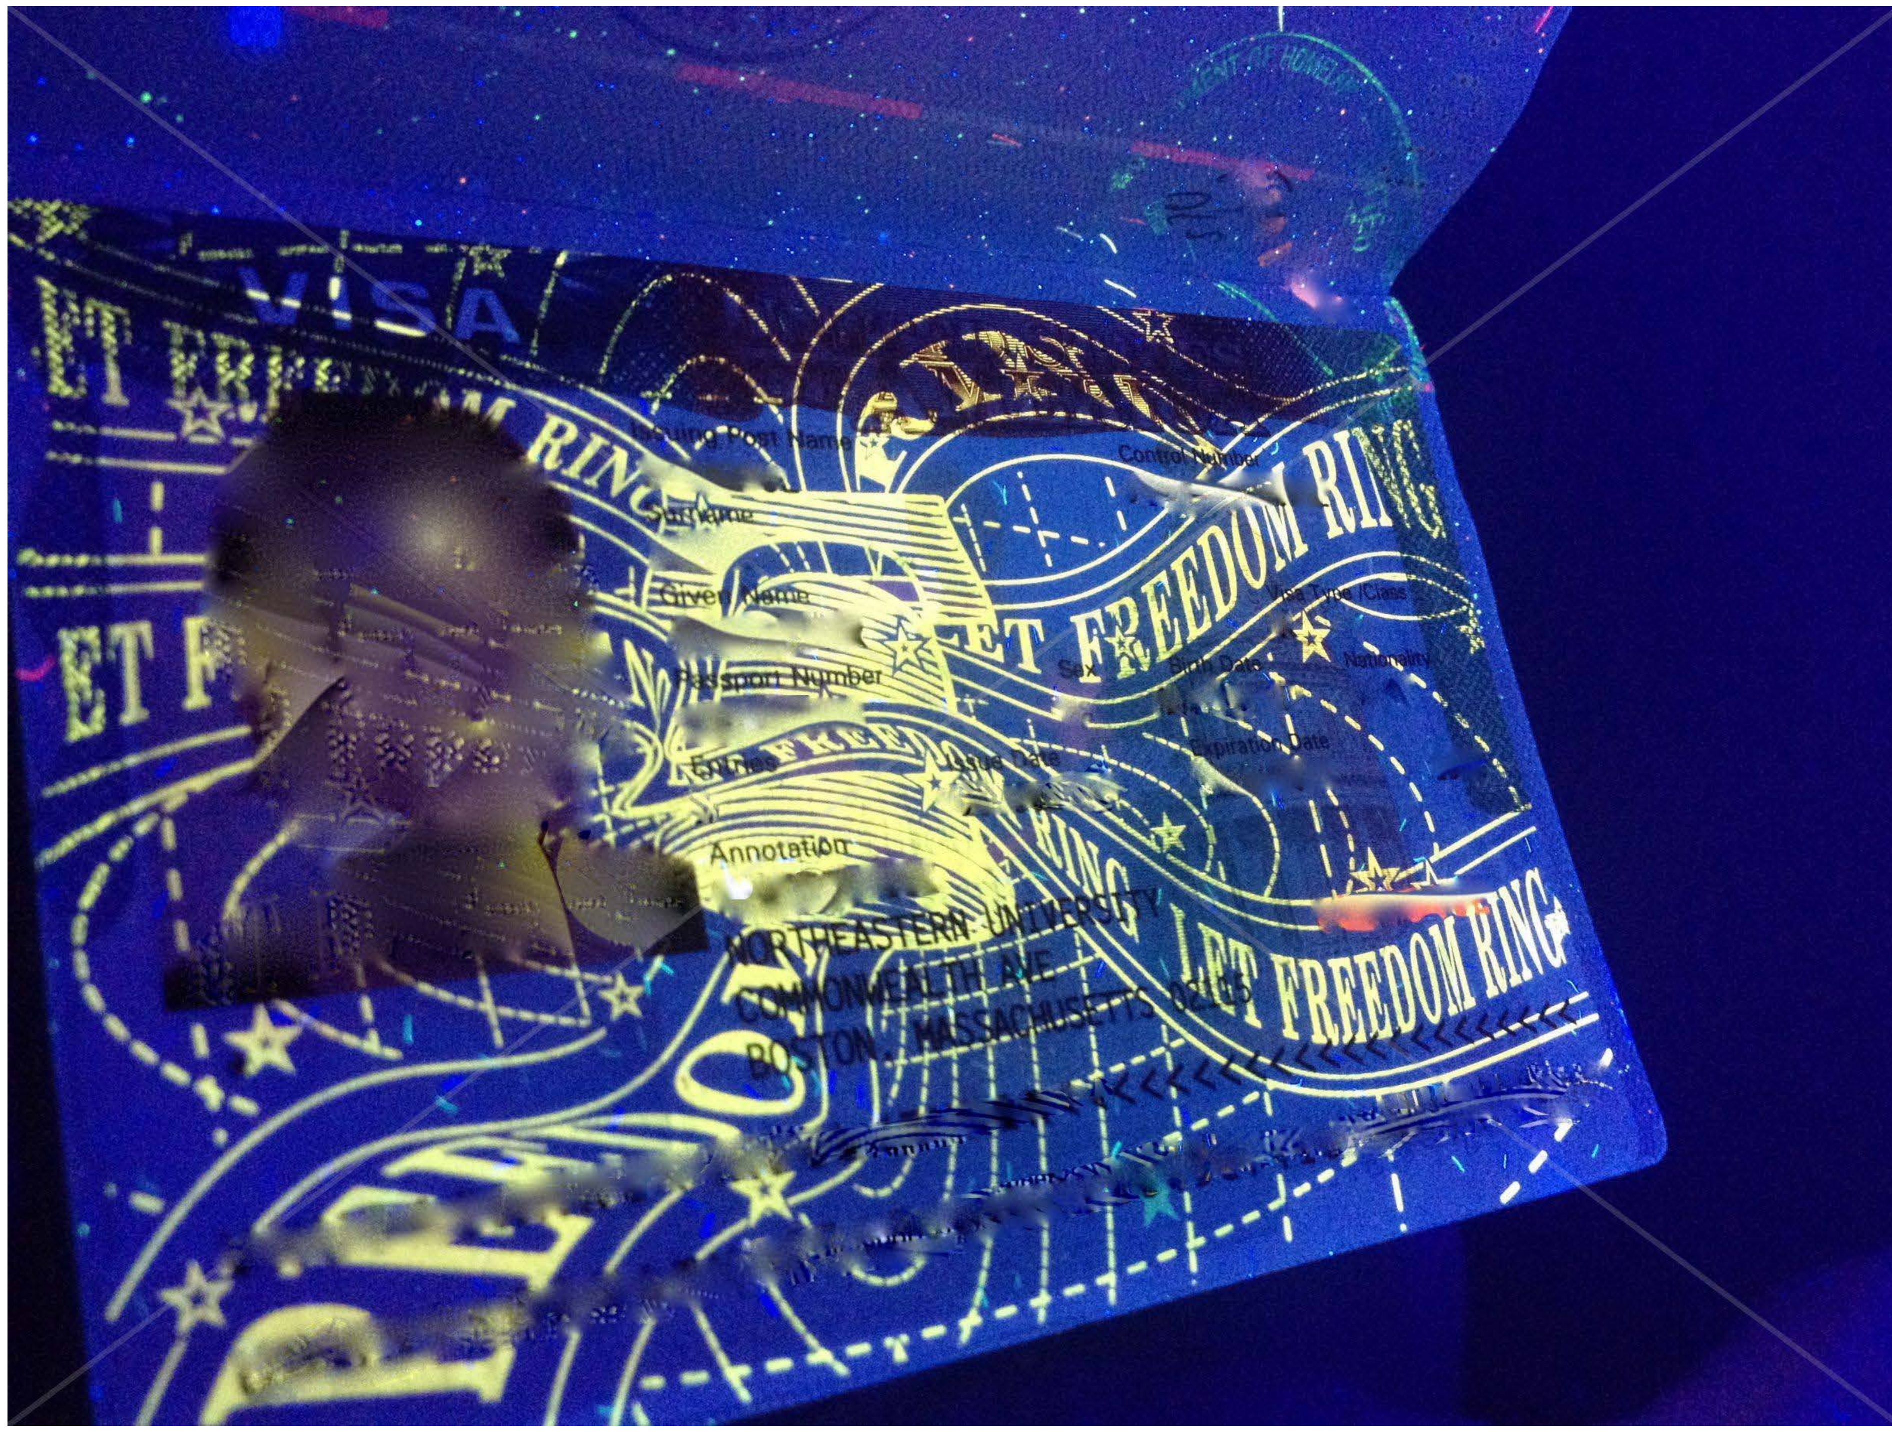

Ok, so color can get pretty complicated!

How do we encode it in a simple(r) way?

Color Spaces and Color Models

- Many ways to specify a color
    - storage
    - convenience
- In general, specify a color from some *color space* using a *color model*
- *Color space* is like artist's palette: full range of colors we can choose from
- *Color model* is the way a particular color in a color space is specified:
    - artist's palette: "yellow ochre"
    - RGB color model: 204, 119, 34

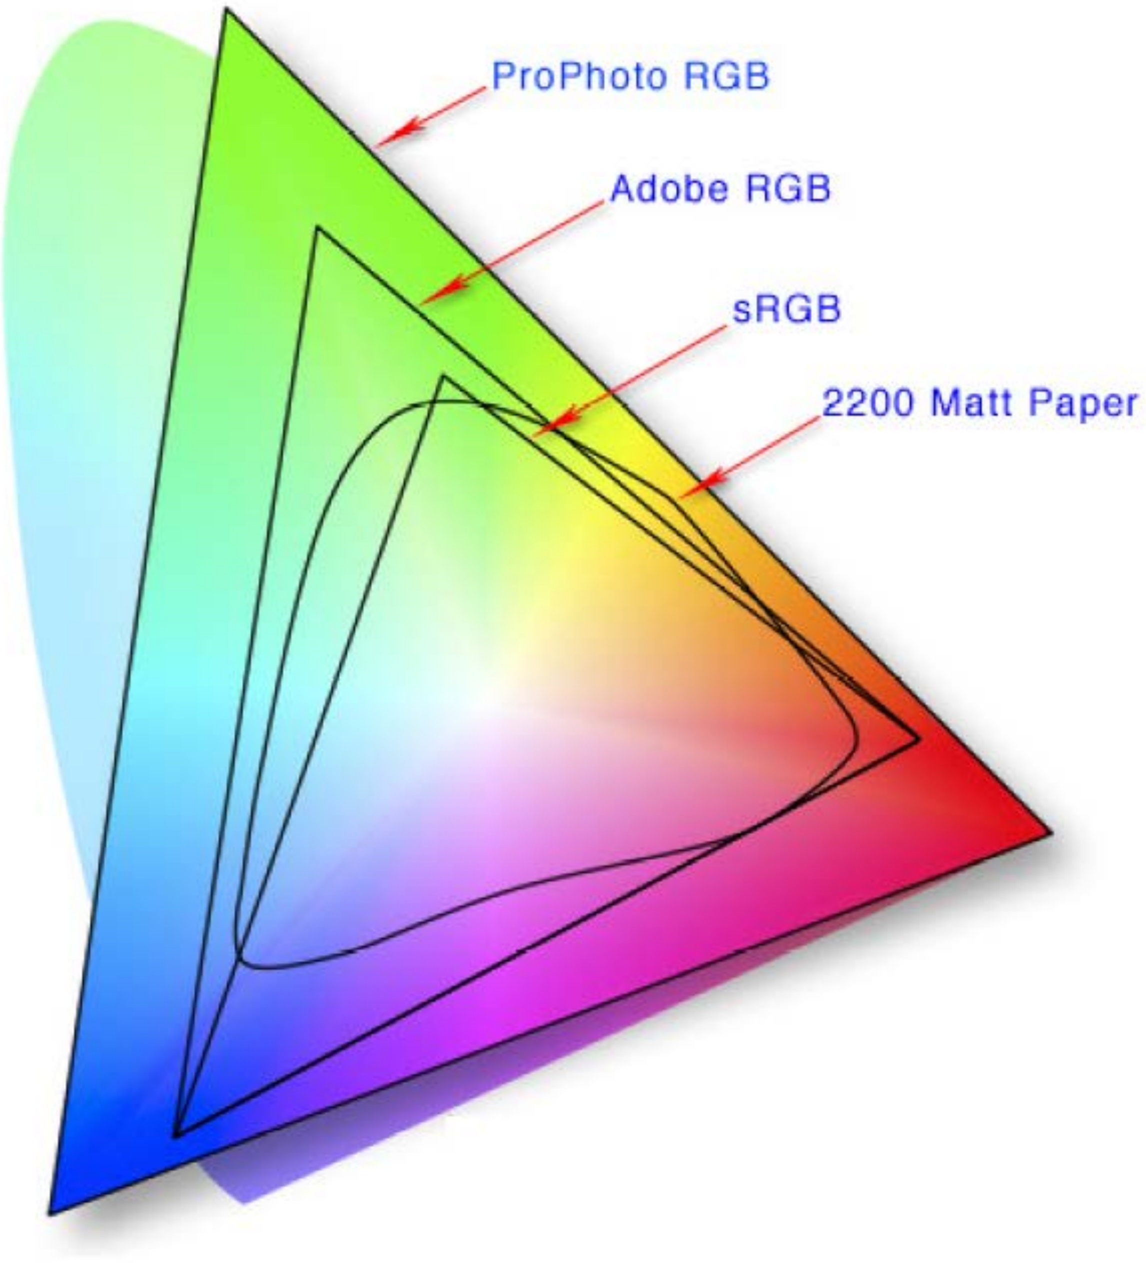

Additive vs. Subtractive Color Models

- Just like we had emission & absorption spectra, we have *additive* and *subtractive** color models
- Additive
  - Used for, e.g., combining colored lights
  - Prototypical example: RGB
- Subtractive
  - Used for, e.g., combining paint colors
  - Prototypical example: CMYK

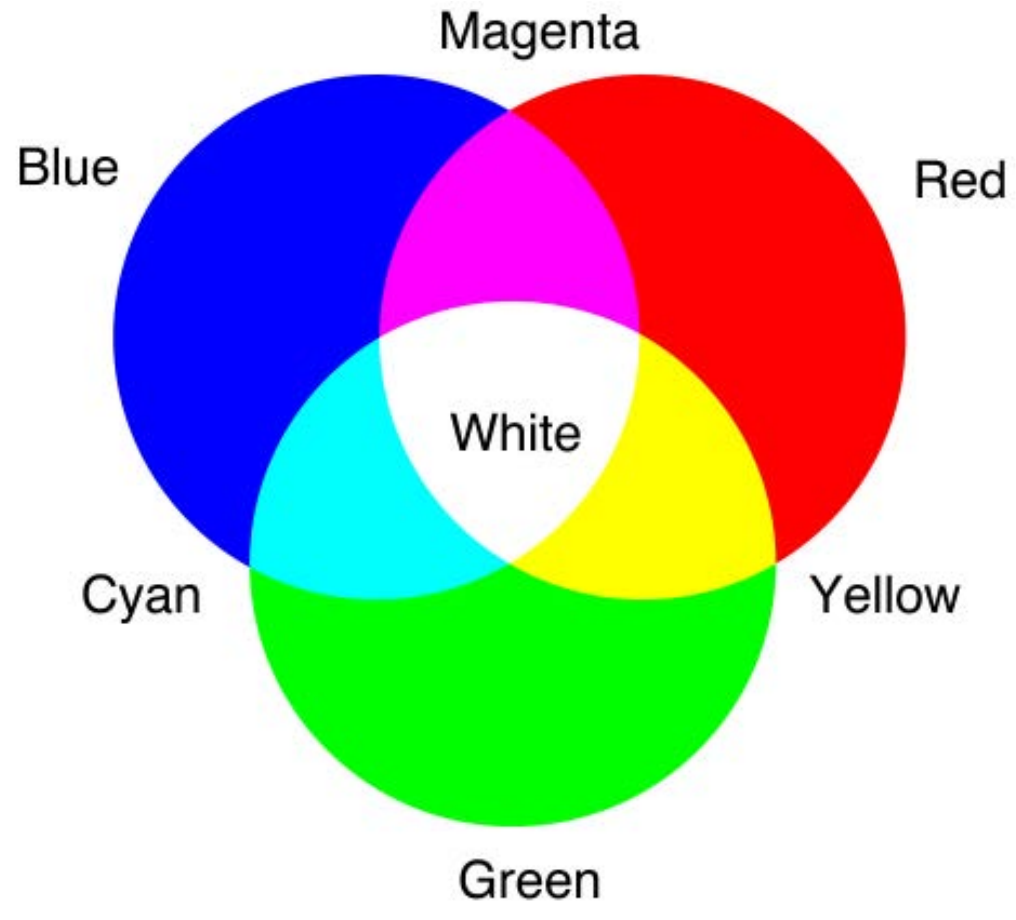

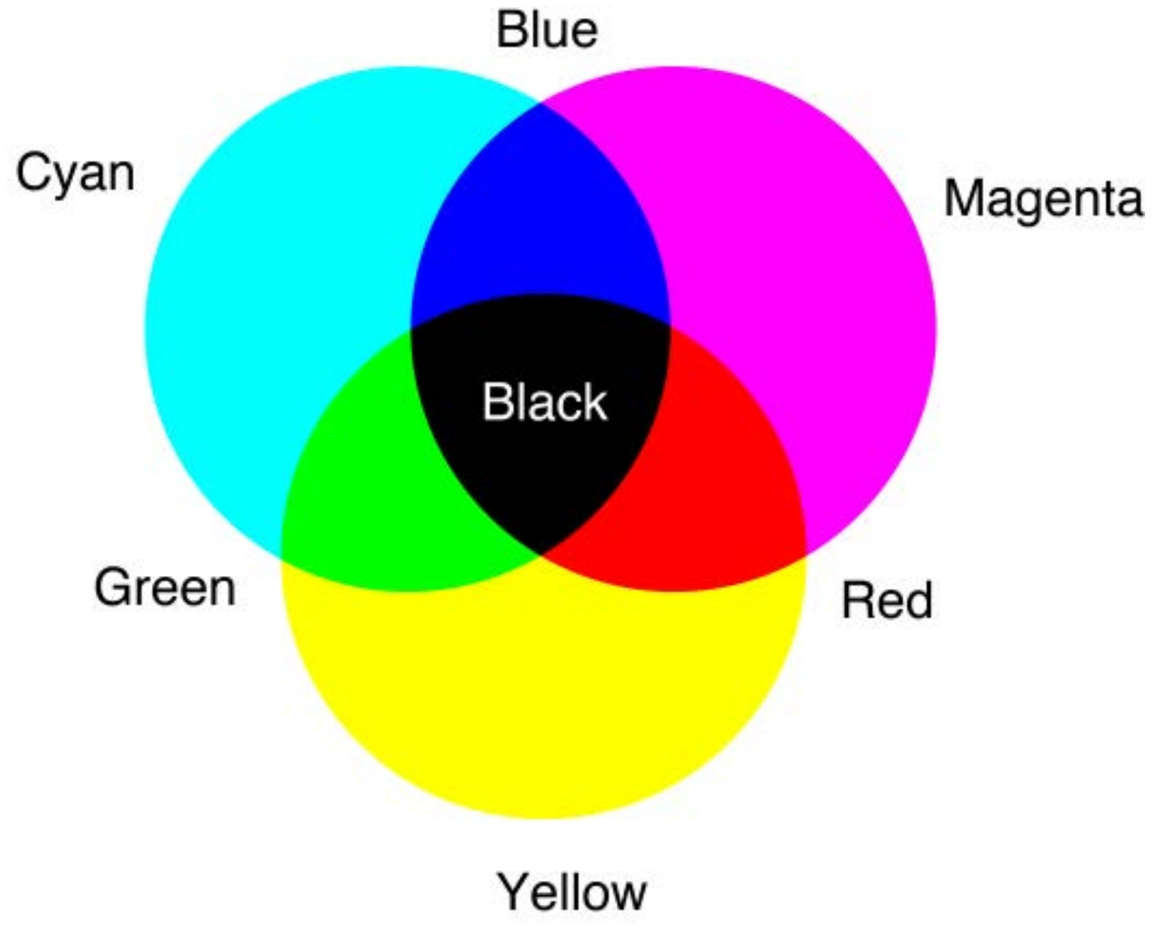

*A better name than subtractive might be *multiplicative*, since we *multiply* to get the final color!

Let’s shed some light on this picture…

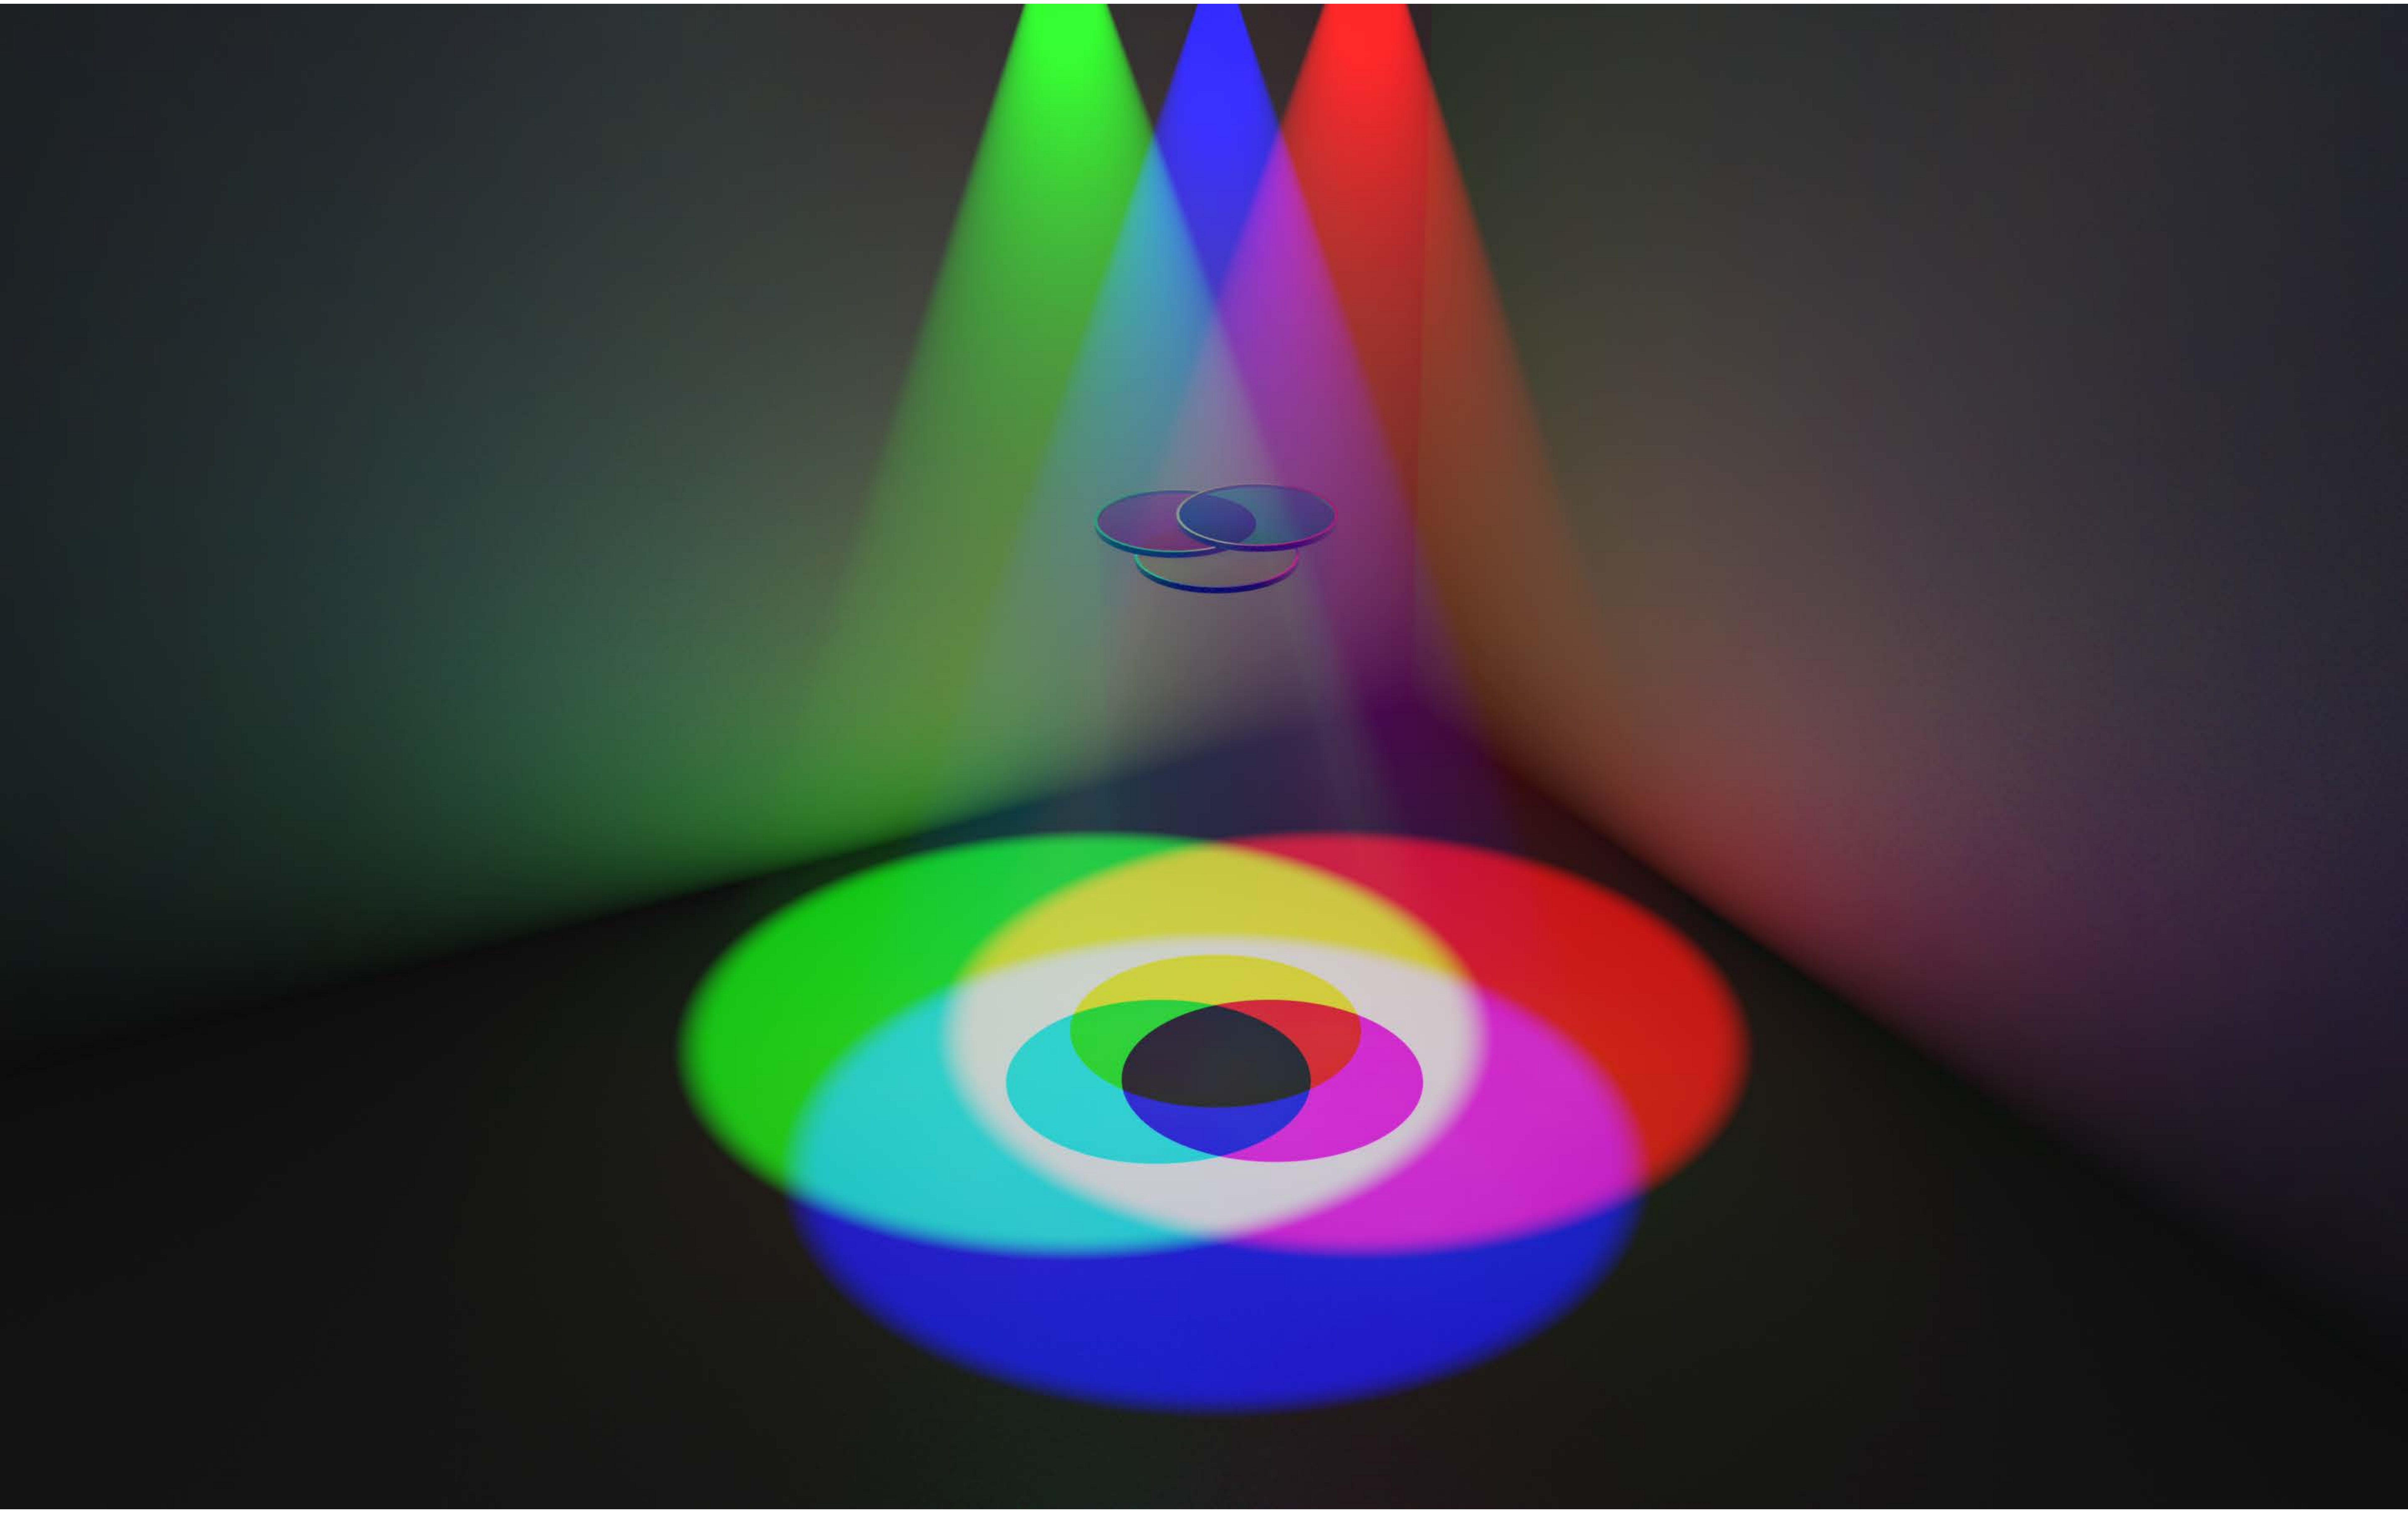

Other Common Color Models

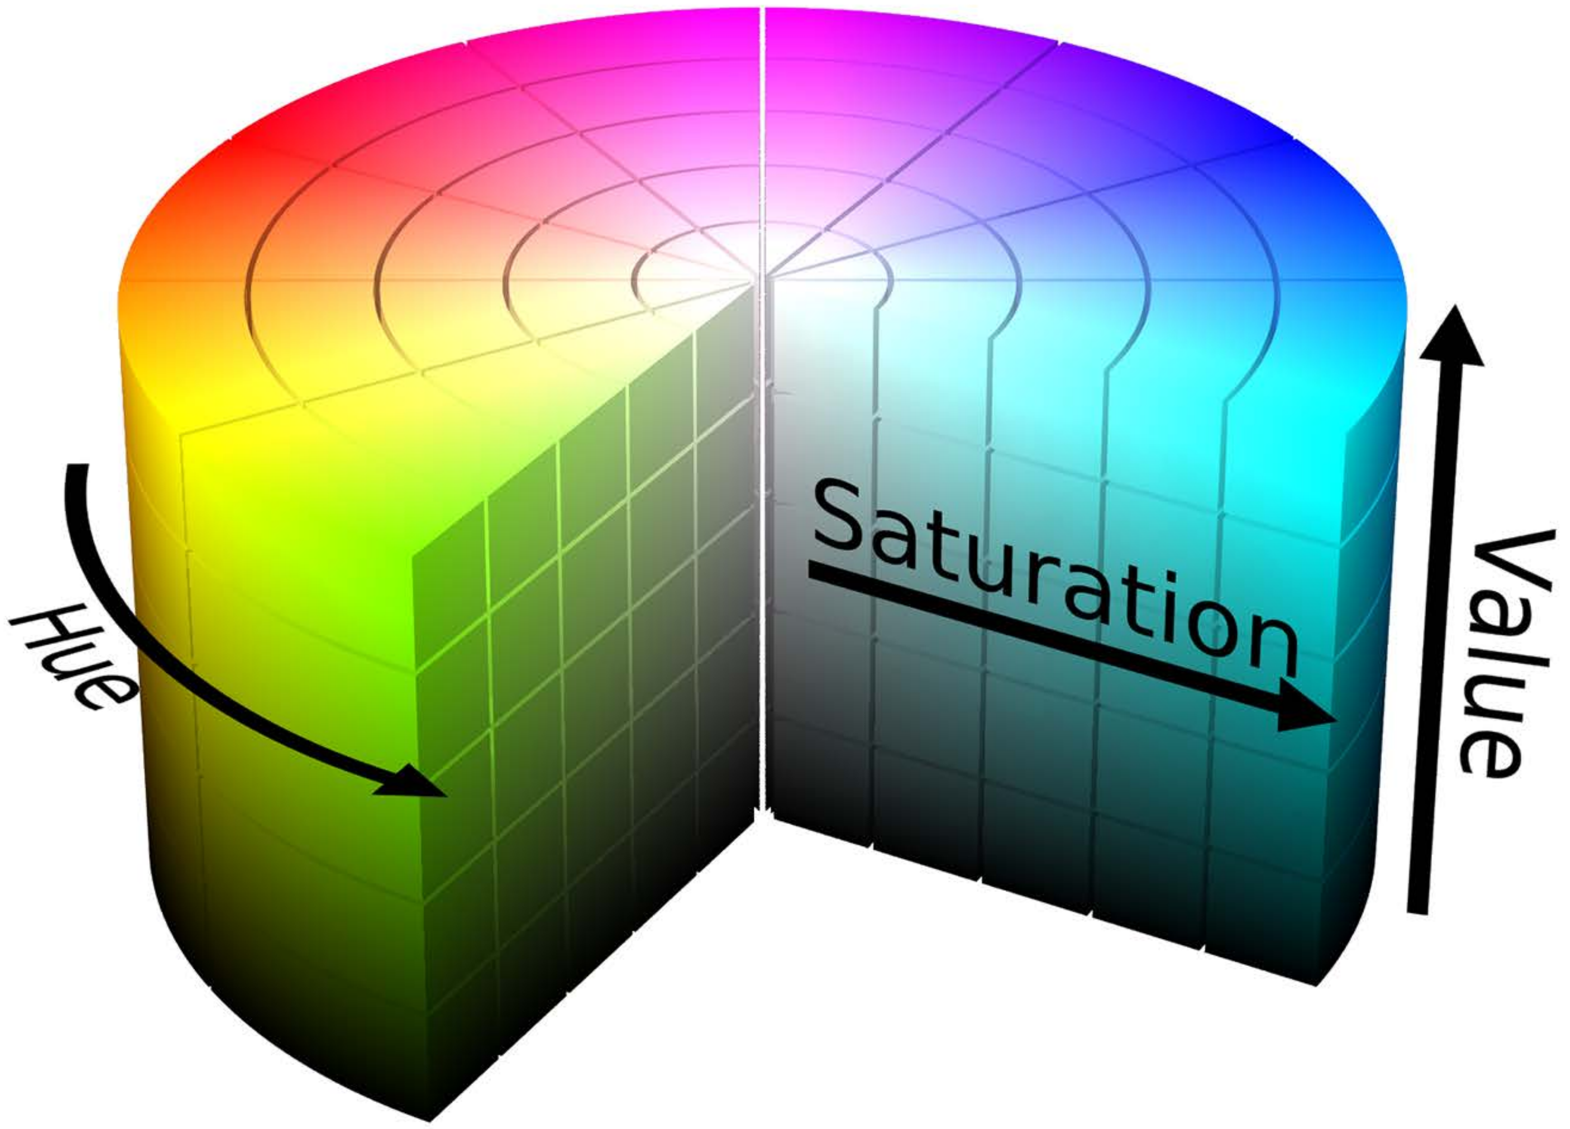

- HSV
  - hue, saturation, value
  - more intuitive than RGB/CMYK
- SML—physiological model
  - corresponds to stimulus of cones
  - not practical for most color work
- XYZ—preceptually-driven model
  - Y captures *luminance* (intensity)
  - X,Z capture *chromaticity* (color)
  - related to, but different from, SML
- Lab—“perceptually uniform” modification of XYZ

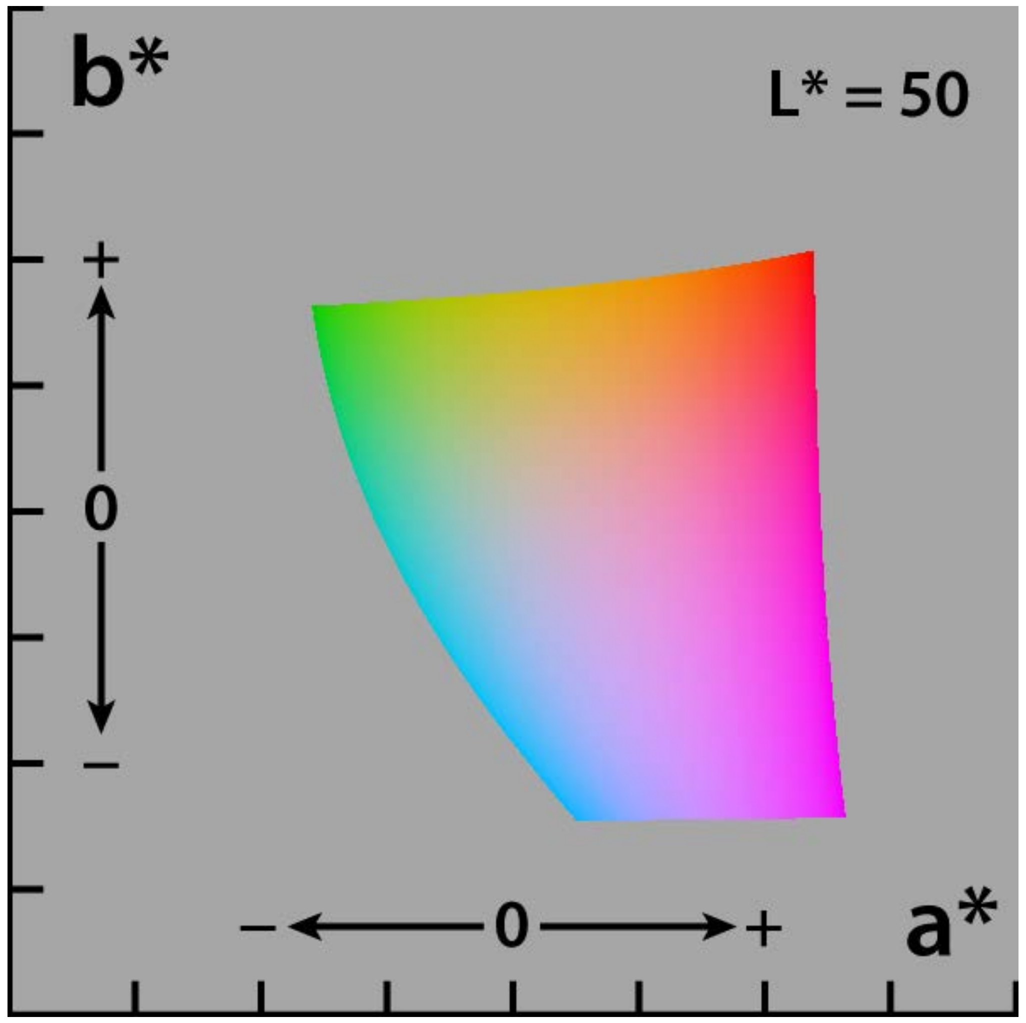

Practical Encoding of Color Values

- How do colors actually get encoded digitally?
- One common encoding (e.g., HTML): *8bpc hexadecimal values**:

`#1B1F8A`

- What does this string mean? Common encoding of RGB.
- Want to store 8-bits per channel (red, green, blue), corresponding to 256 possible values
- Rather than use digits 0-9, use 0, 1,2,3,4,5,6,7,8,9,A,B,C,D,E,F
- Single character now encodes 16 values, two characters encode 16*16 = 256 values
- Q: Roughly what color is `#ff6600`?

*Upper vs. lowercase letters? Makes absolutely *no* difference!

Other Ways of Specifying Color?

- Other color specifications not based on continuous color space
- E.g., Pantone Matching System
  - industry standard (proprietary)
  - 1,114 colors
  - Combination of 13 base pigments
- And not to forget…

Why use different color models?

- Convenience
  - Is it easy for a user to choose the color they want?
- Color compositing/processing
  - Does it matter what color space we interpolate / blend in?
- Efficiency of encoding
  - E.g., use more of numerical range for perceptually significant colors
  - Compression!

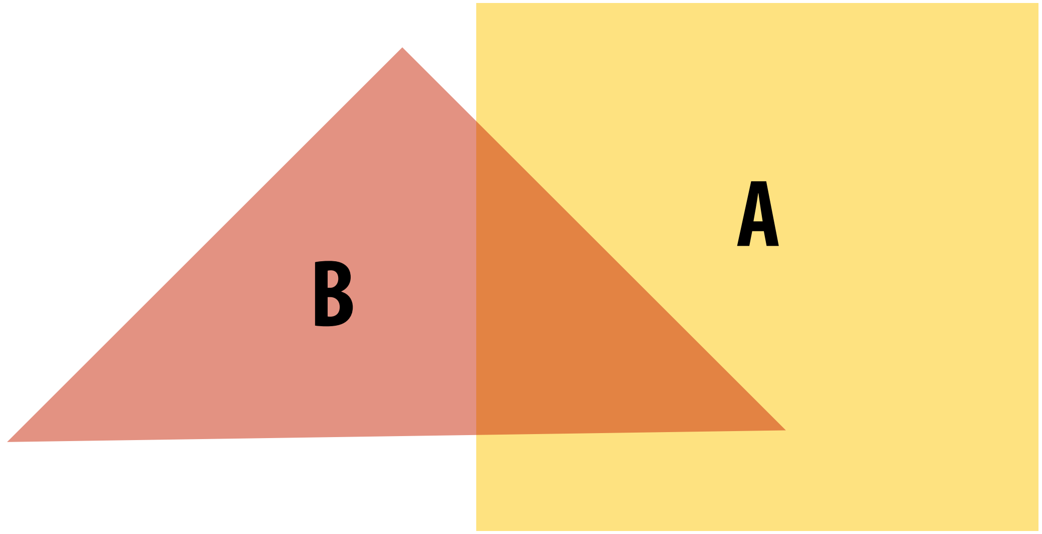

B over A

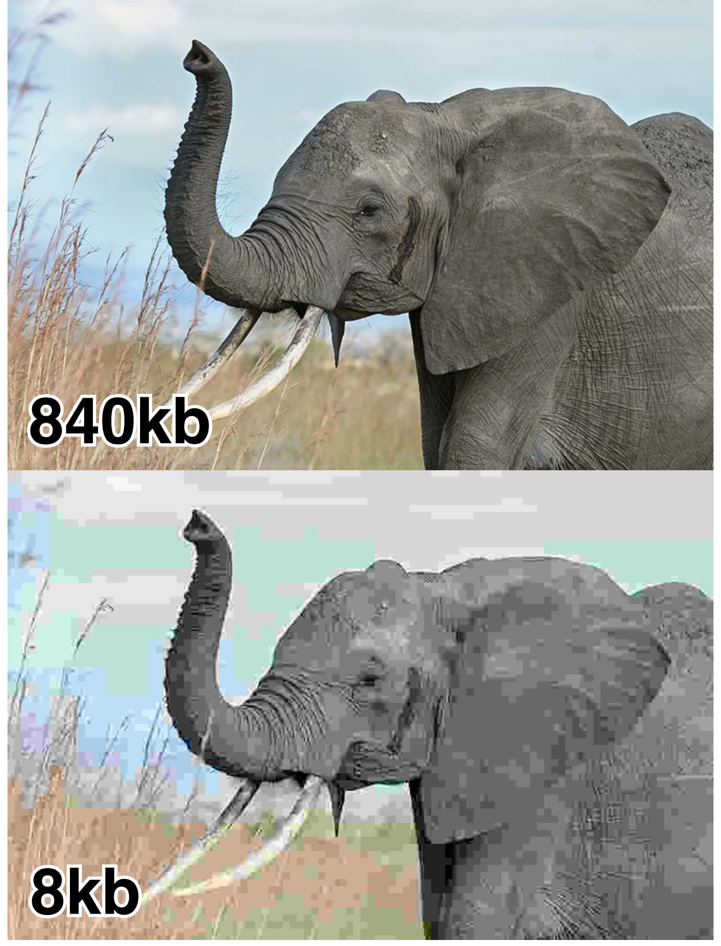

Example: Y'CbCr color model

- Common for modern digital video
- Y' = luma: perceived luminance (same as L* in CIELAB)
- Cb = blue-yellow deviation from gray
- Cr = red-cyan deviation from gray

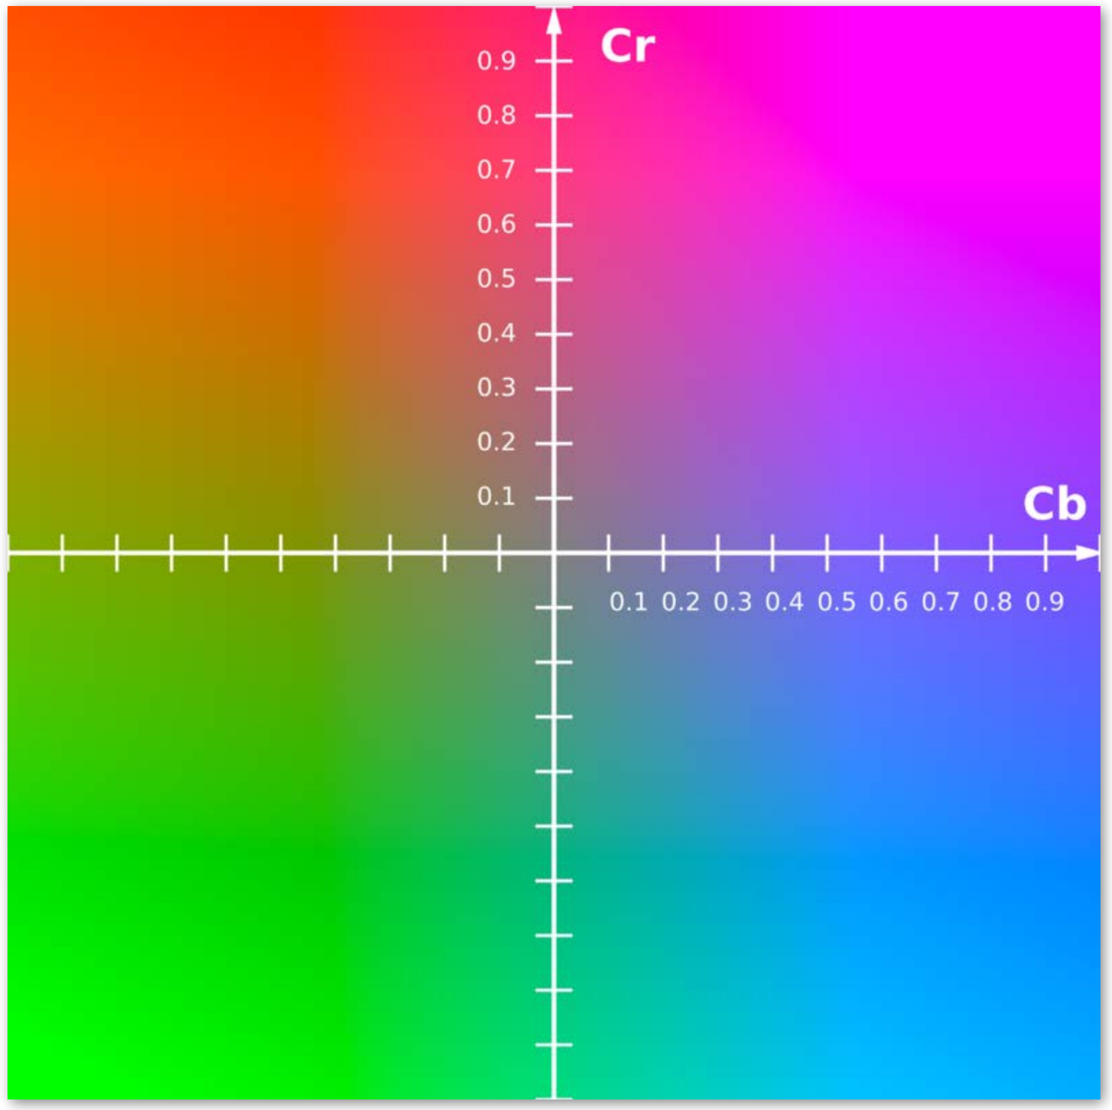


(input)


Y'


Cb


Cr

Original picture

Contents of CbCr color channels downsampled by a factor of 20 in each dimension
(400x reduction in number of samples)

Full resolution sampling of luma (Y′)

Reconstructed result
(looks pretty good)

Original picture

By the way, how might we reduce this artifact?

Reconstructed result

Why use different color models? (cont.)

- **Convenience**
  - **Is it easy for a user to choose the color they want?**
- **Efficiency of encoding**
  - **E.g., use more of numerical range for perceptually significant colors**
  - **Do color images *compress* well?**
- ***Gamut***
  - **Which colors can be expressed using a given model?**
  - **Very different for print vs. display**

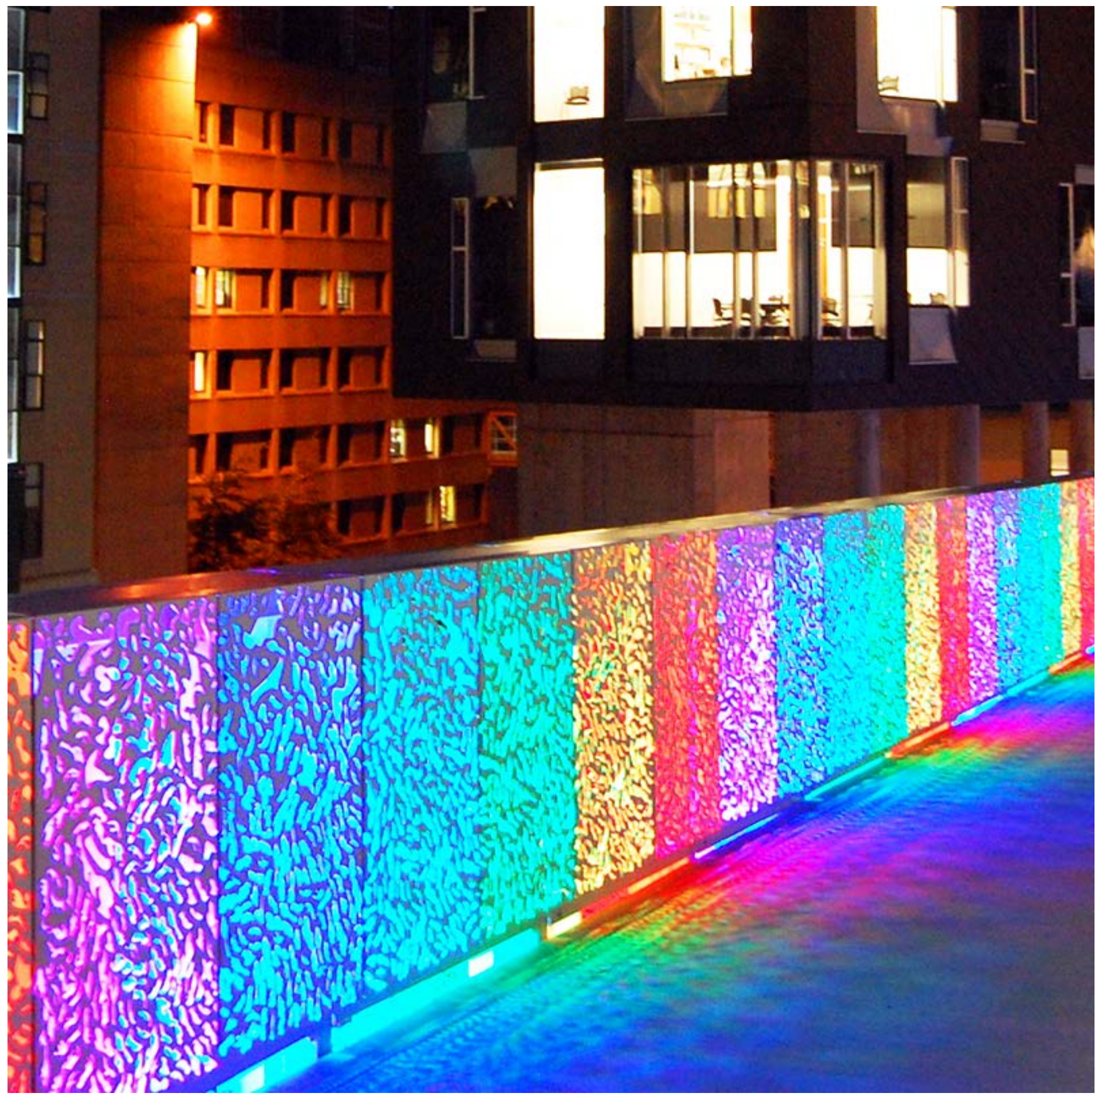

RGB

CMYK

Which raises a very important question:

Which actual colors (i.e., *spectra*) do these values get mapped to?

CIE 1931* Color Space

- Standard "reference" color space
- Encompasses all colors visible by "most" human observers
  - associated color *model* (XYZ) captures perceptual effects
  - e.g., perception of color ("chromaticity") changes w/ brightness ("luminosity")
  - different from specifying direct simulation of cones (SML)

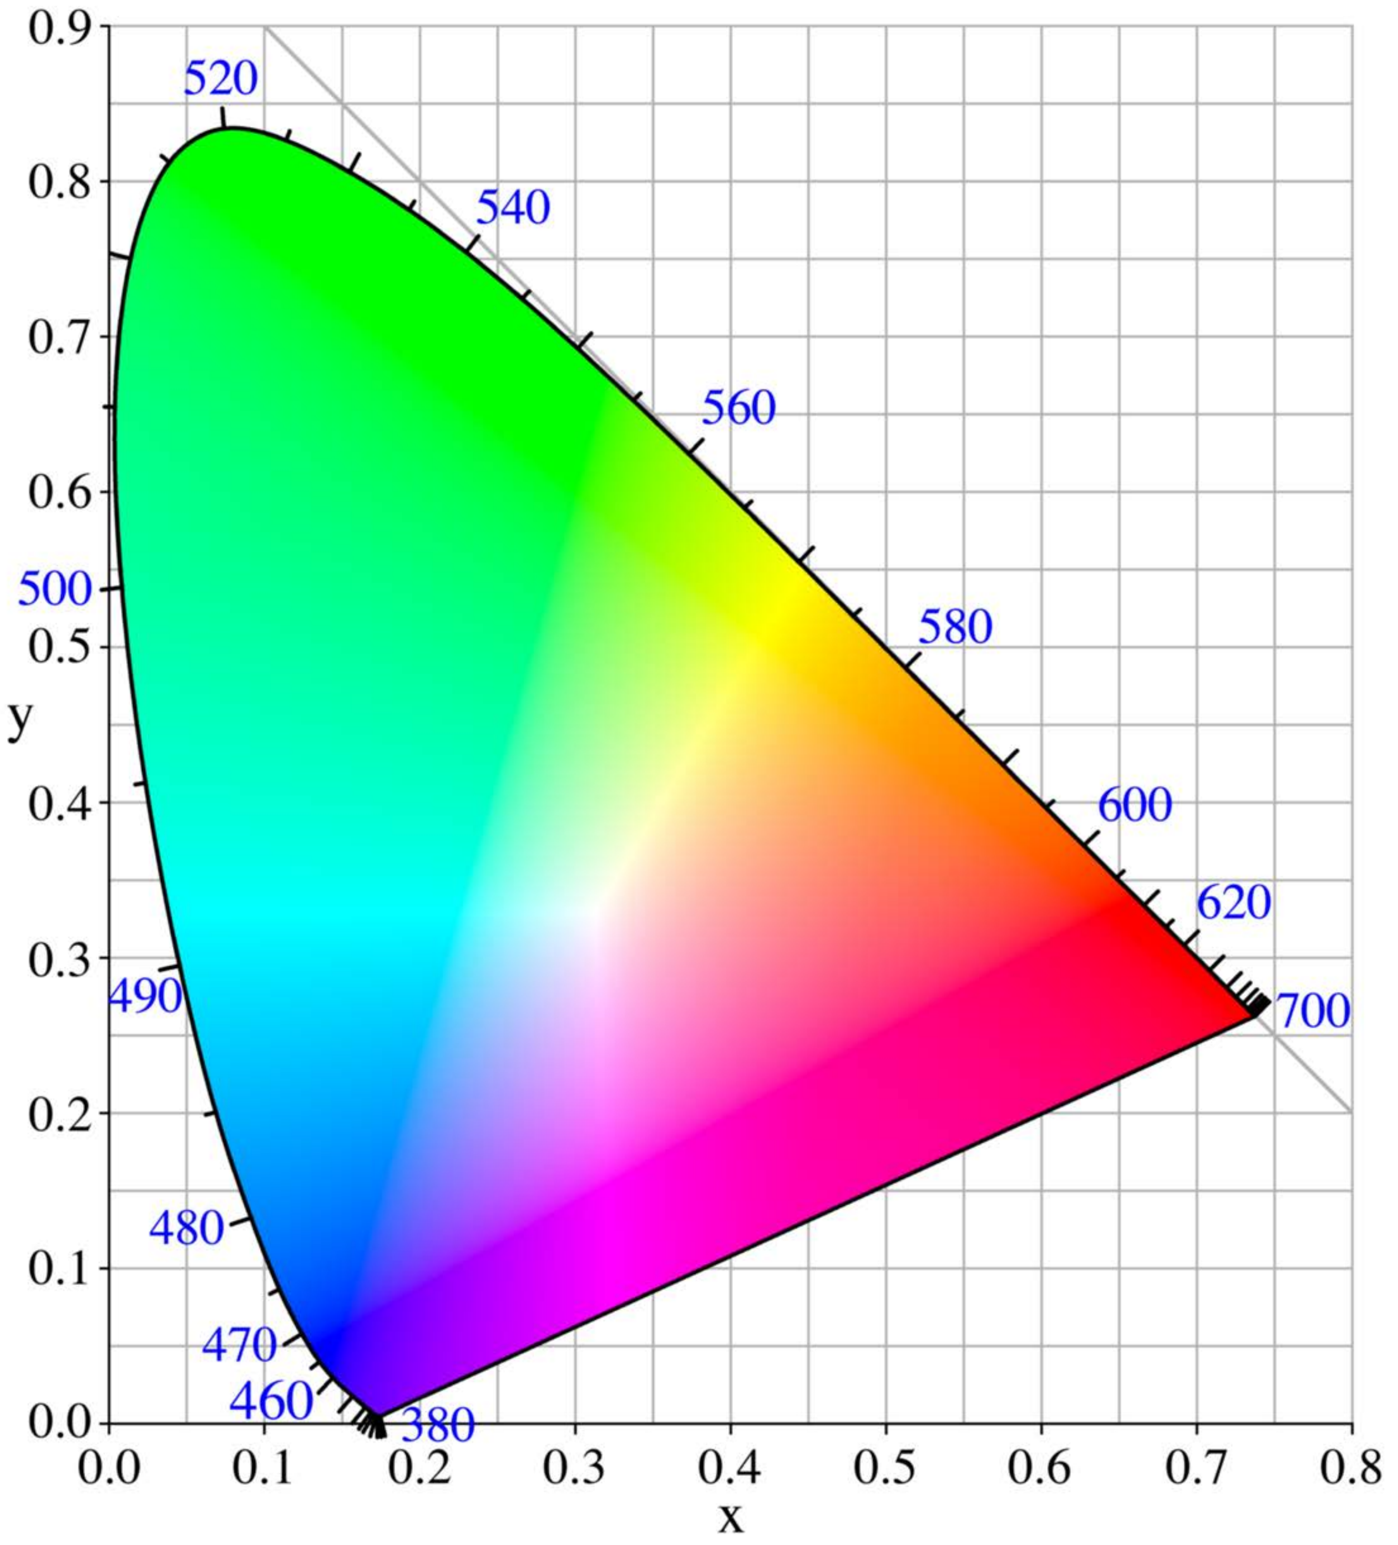

*CIE 1931 does not mean anything important: "created in 1931 by the Commission Internationale de l'Éclairage"

Chromaticity Diagrams

- Chromaticity is the intensity-independent component of a color
- Chromaticity diagram used to visualize extent of a color space

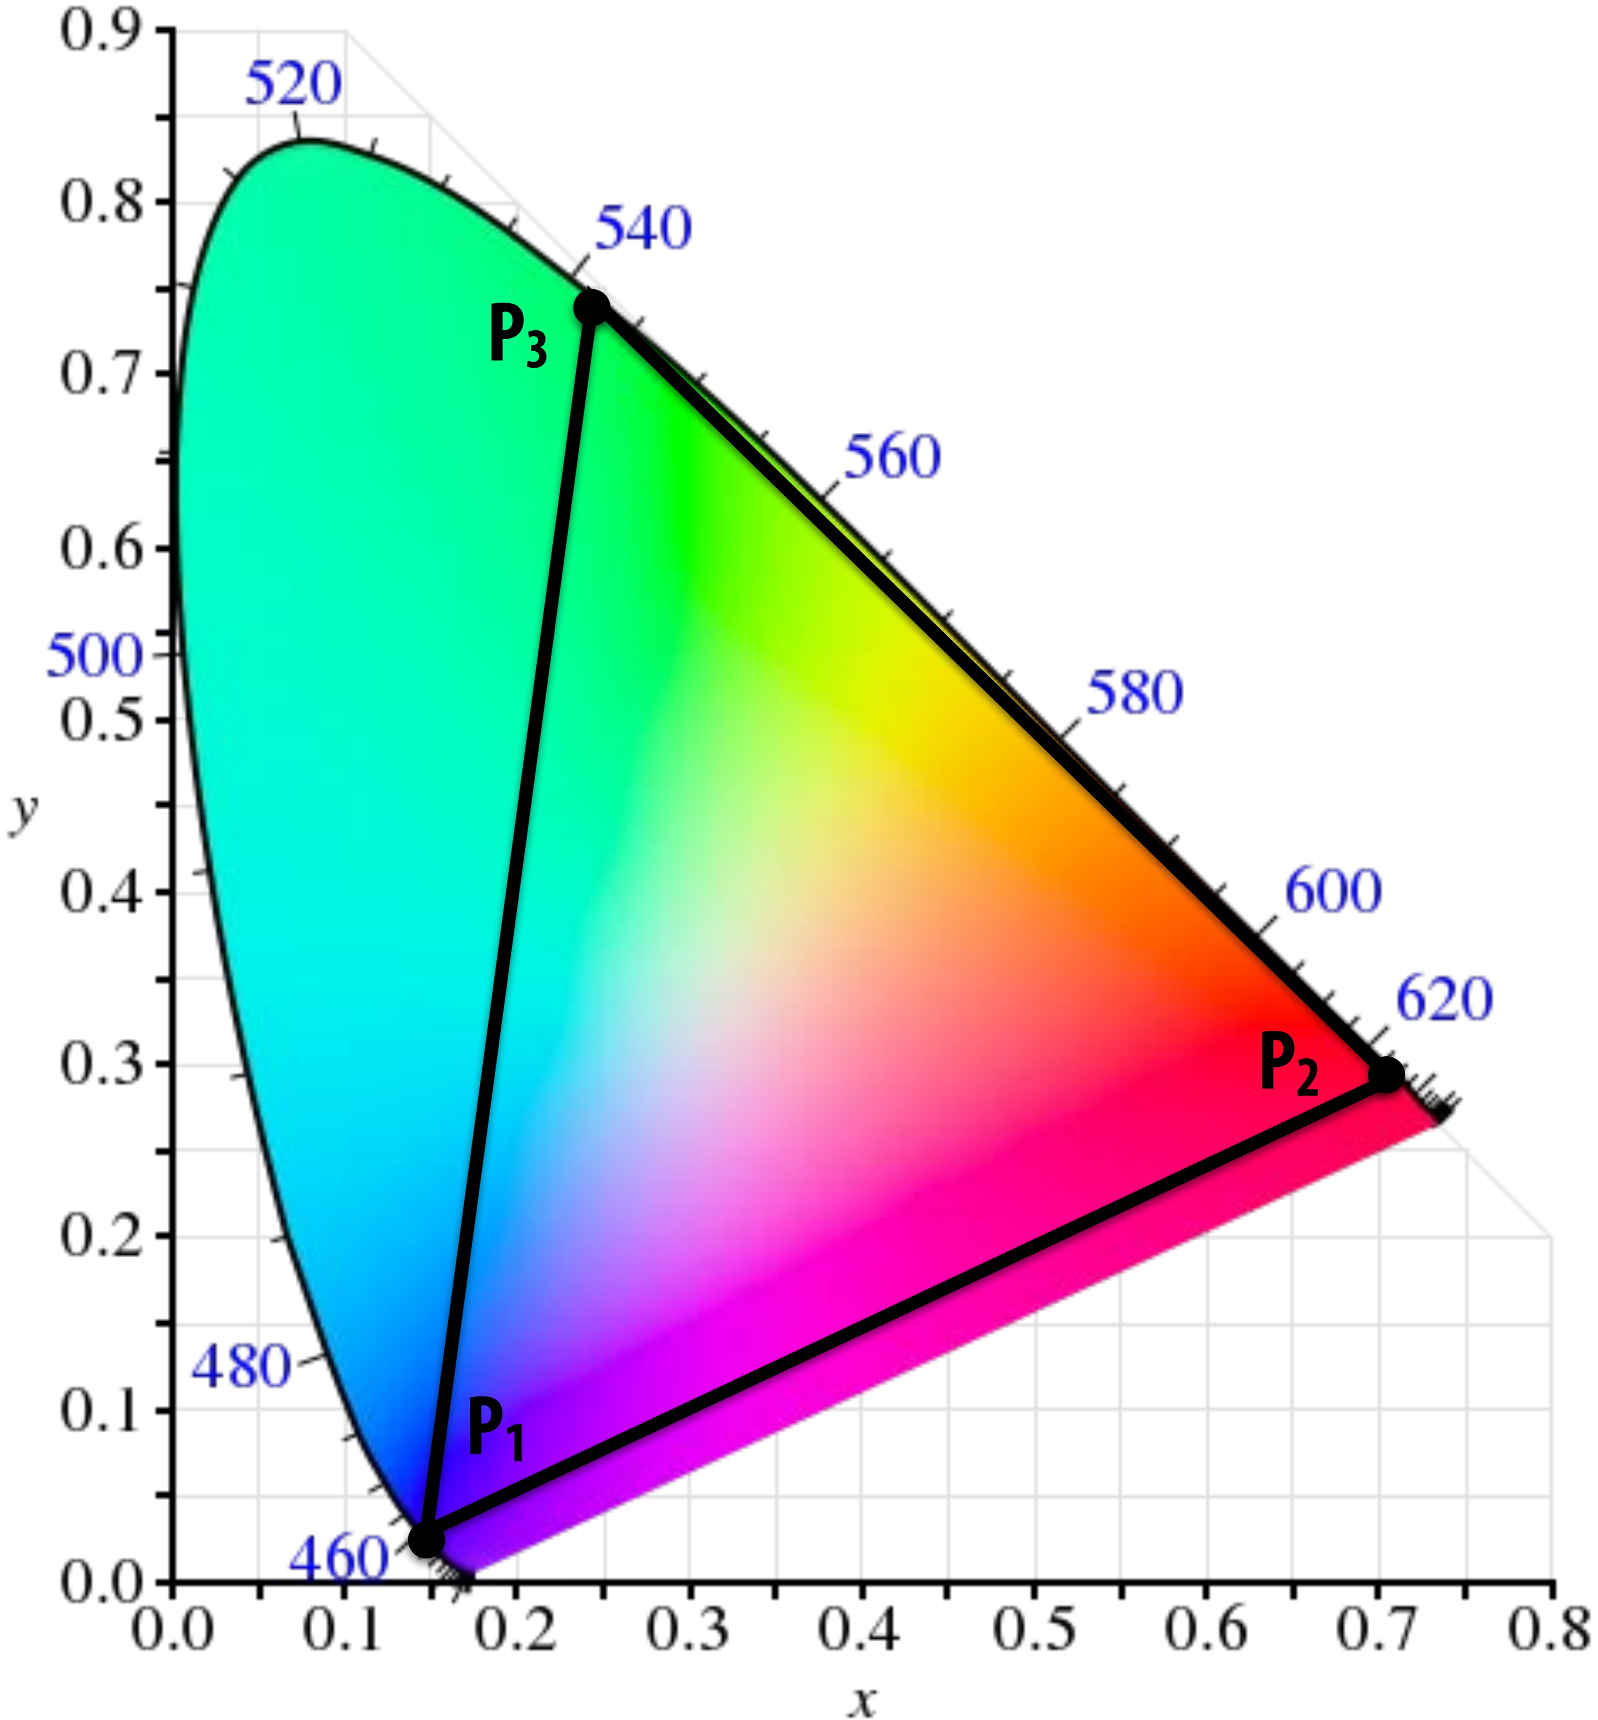

A display with primaries with chromacities P_1, P_2, P_3 can create colors that are combinations of these primaries (colors that fall within the triangle)

sRGB Color Space

- CIE 1934 captured all possible human-visible colors
- sRGB (roughly) subset of colors available on displays, printers, …
- Nonlinear relationship between stored RGB values & intensity
  - Makes better use of limited set of numerical values

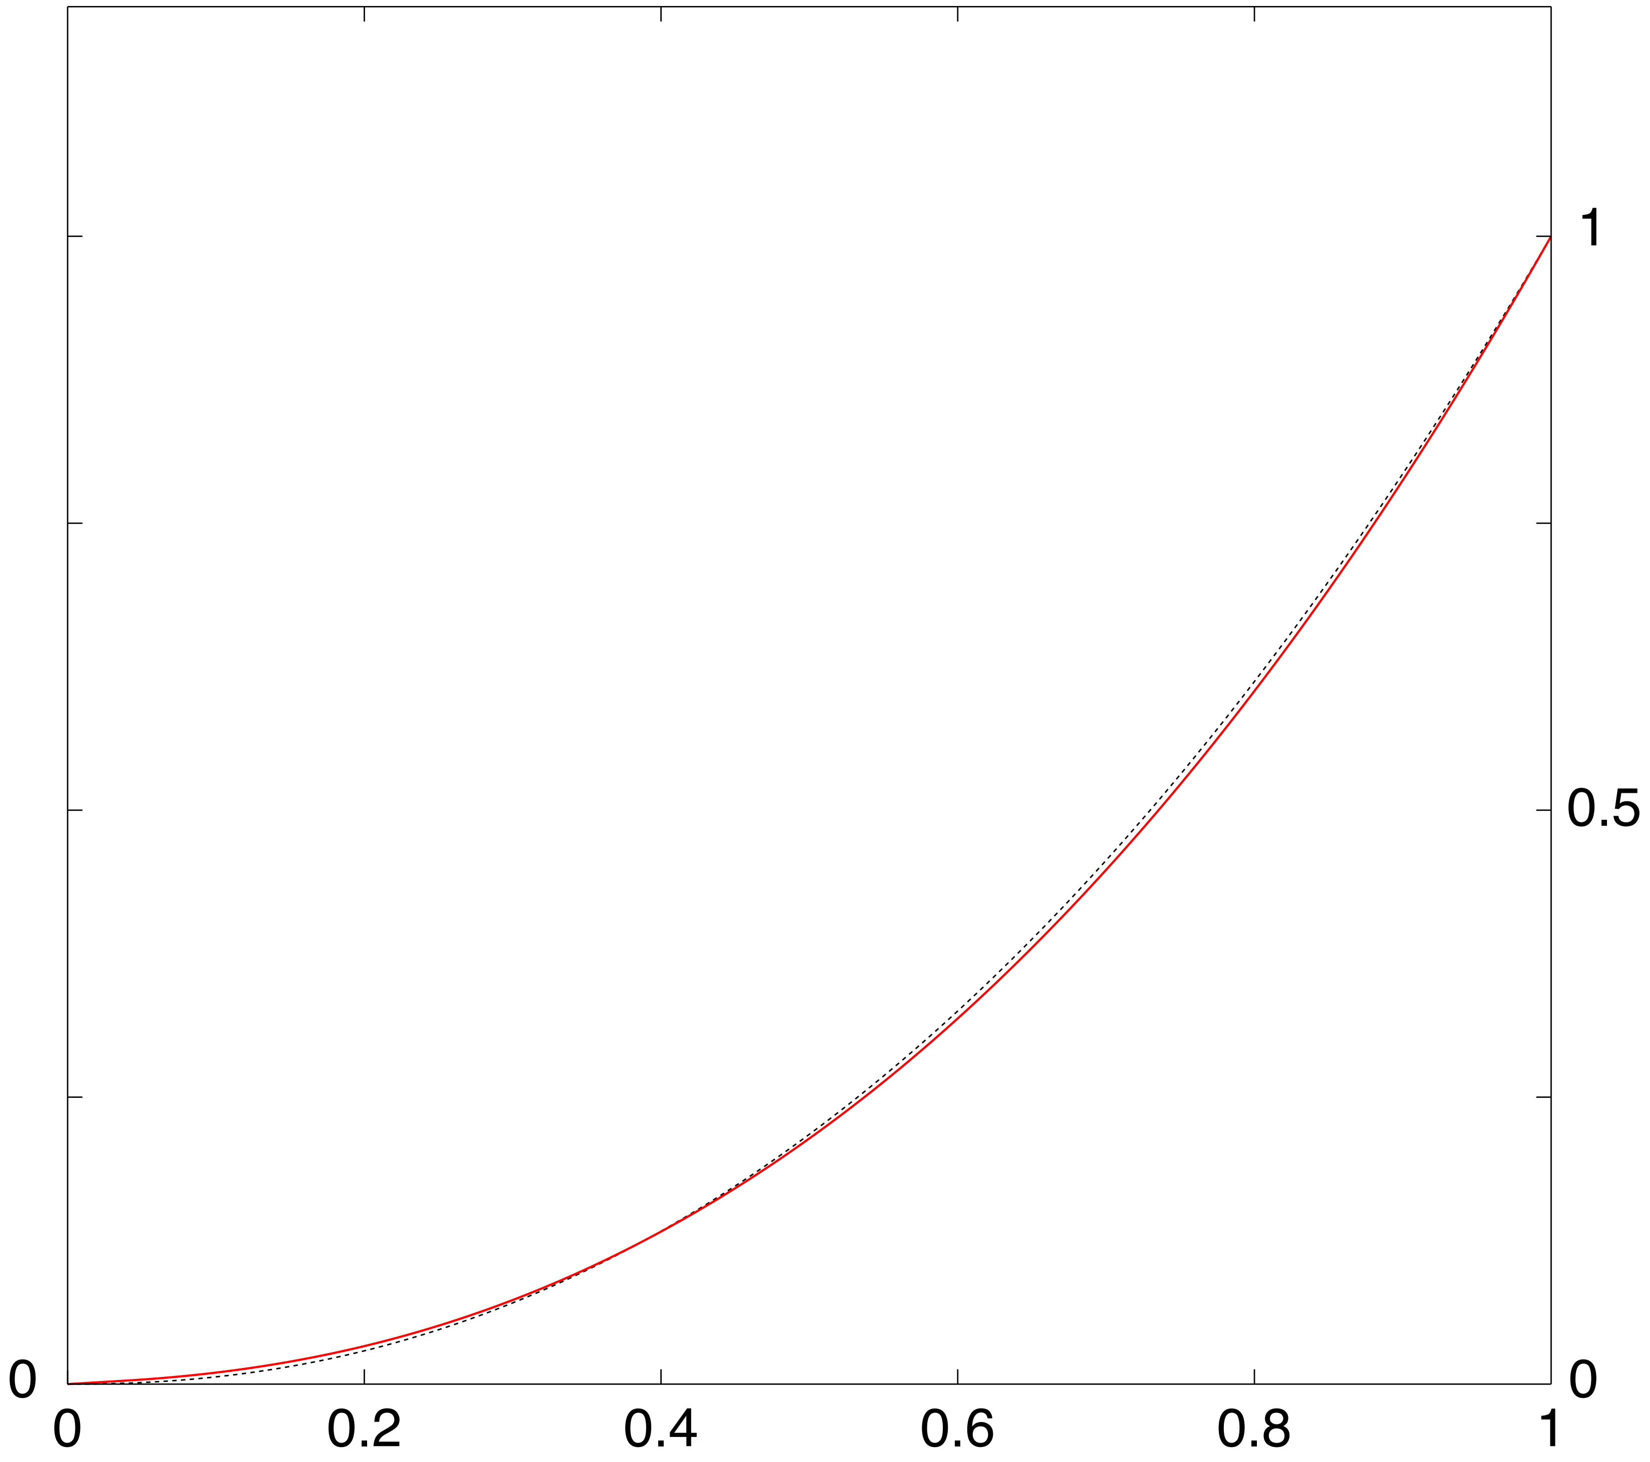

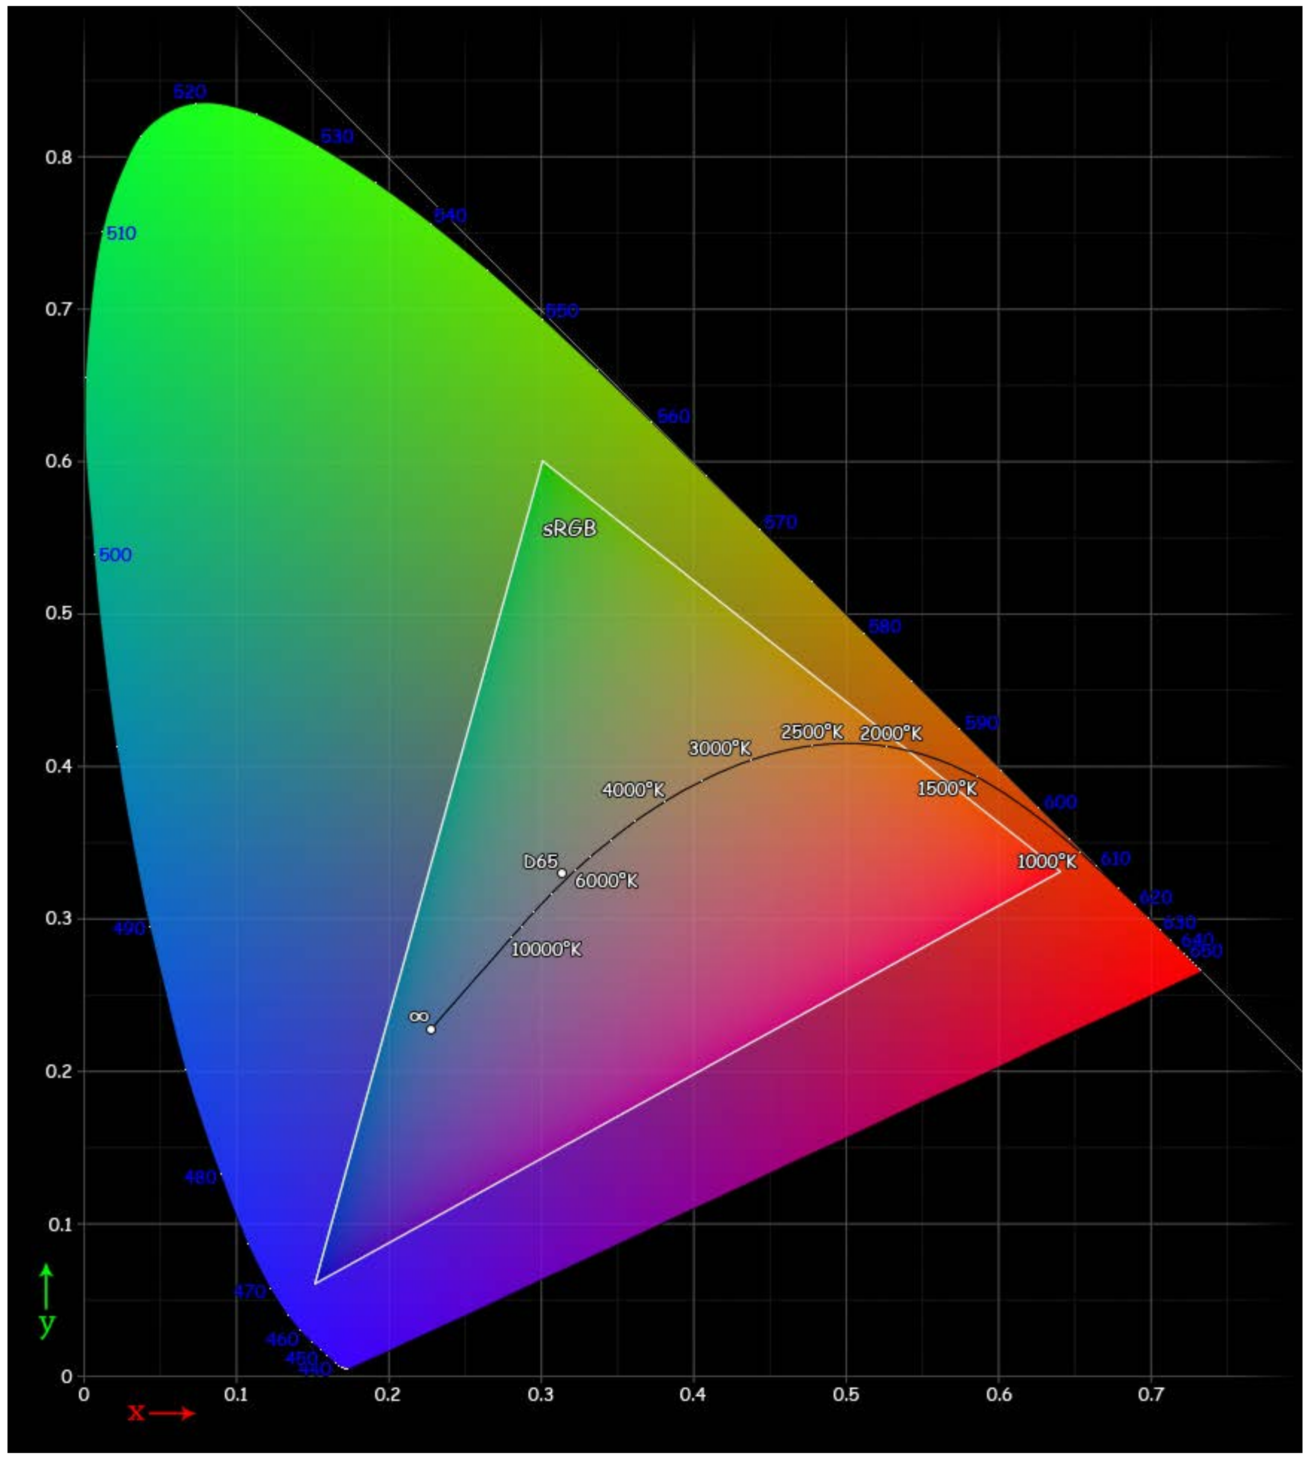

Color Acuity (MacAdam Ellipse)

- In addition to *range* of colors visible, one might be interested in how *sensitive* people are to changes in color
- Each ellipse corresponds to a region of "just noticeable differences" of color (chromaticity)
- So, if you want to make two colors distinct, at bare minimum should avoid overlapping ellipses…

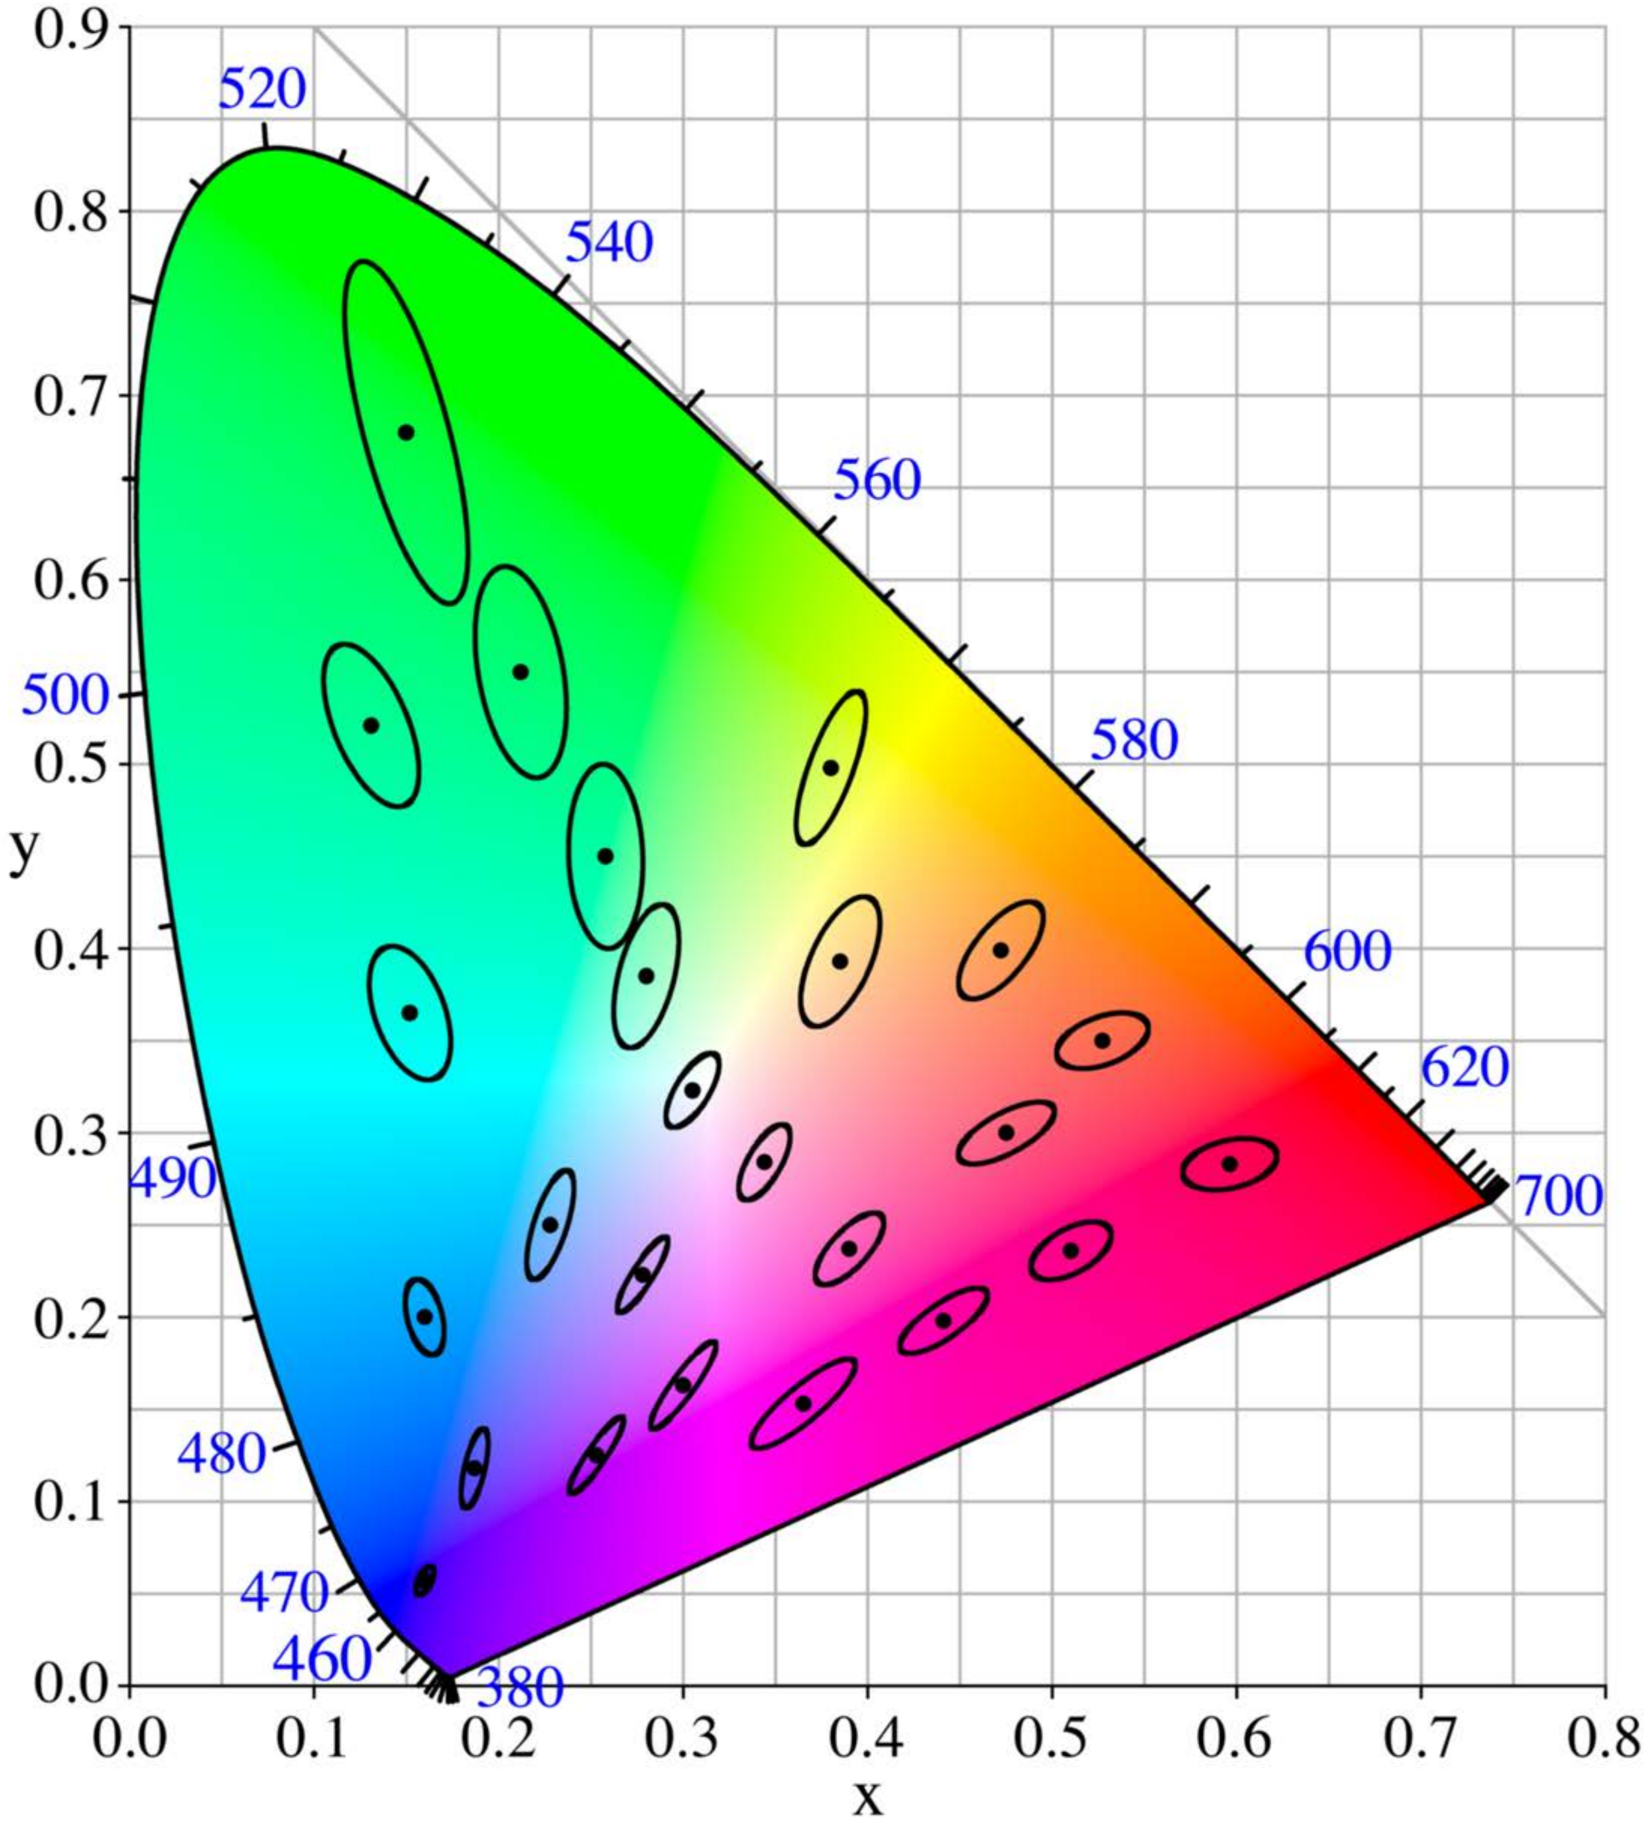

Nonstandard Color Vision

- **Morphological differences in eye can cause people (& animals) to see different ranges of color (e.g., more/fewer cone types)**
- **Alternative chromaticity diagrams help visualize color gamut, useful for designing, e.g., widely-accessible interfaces**

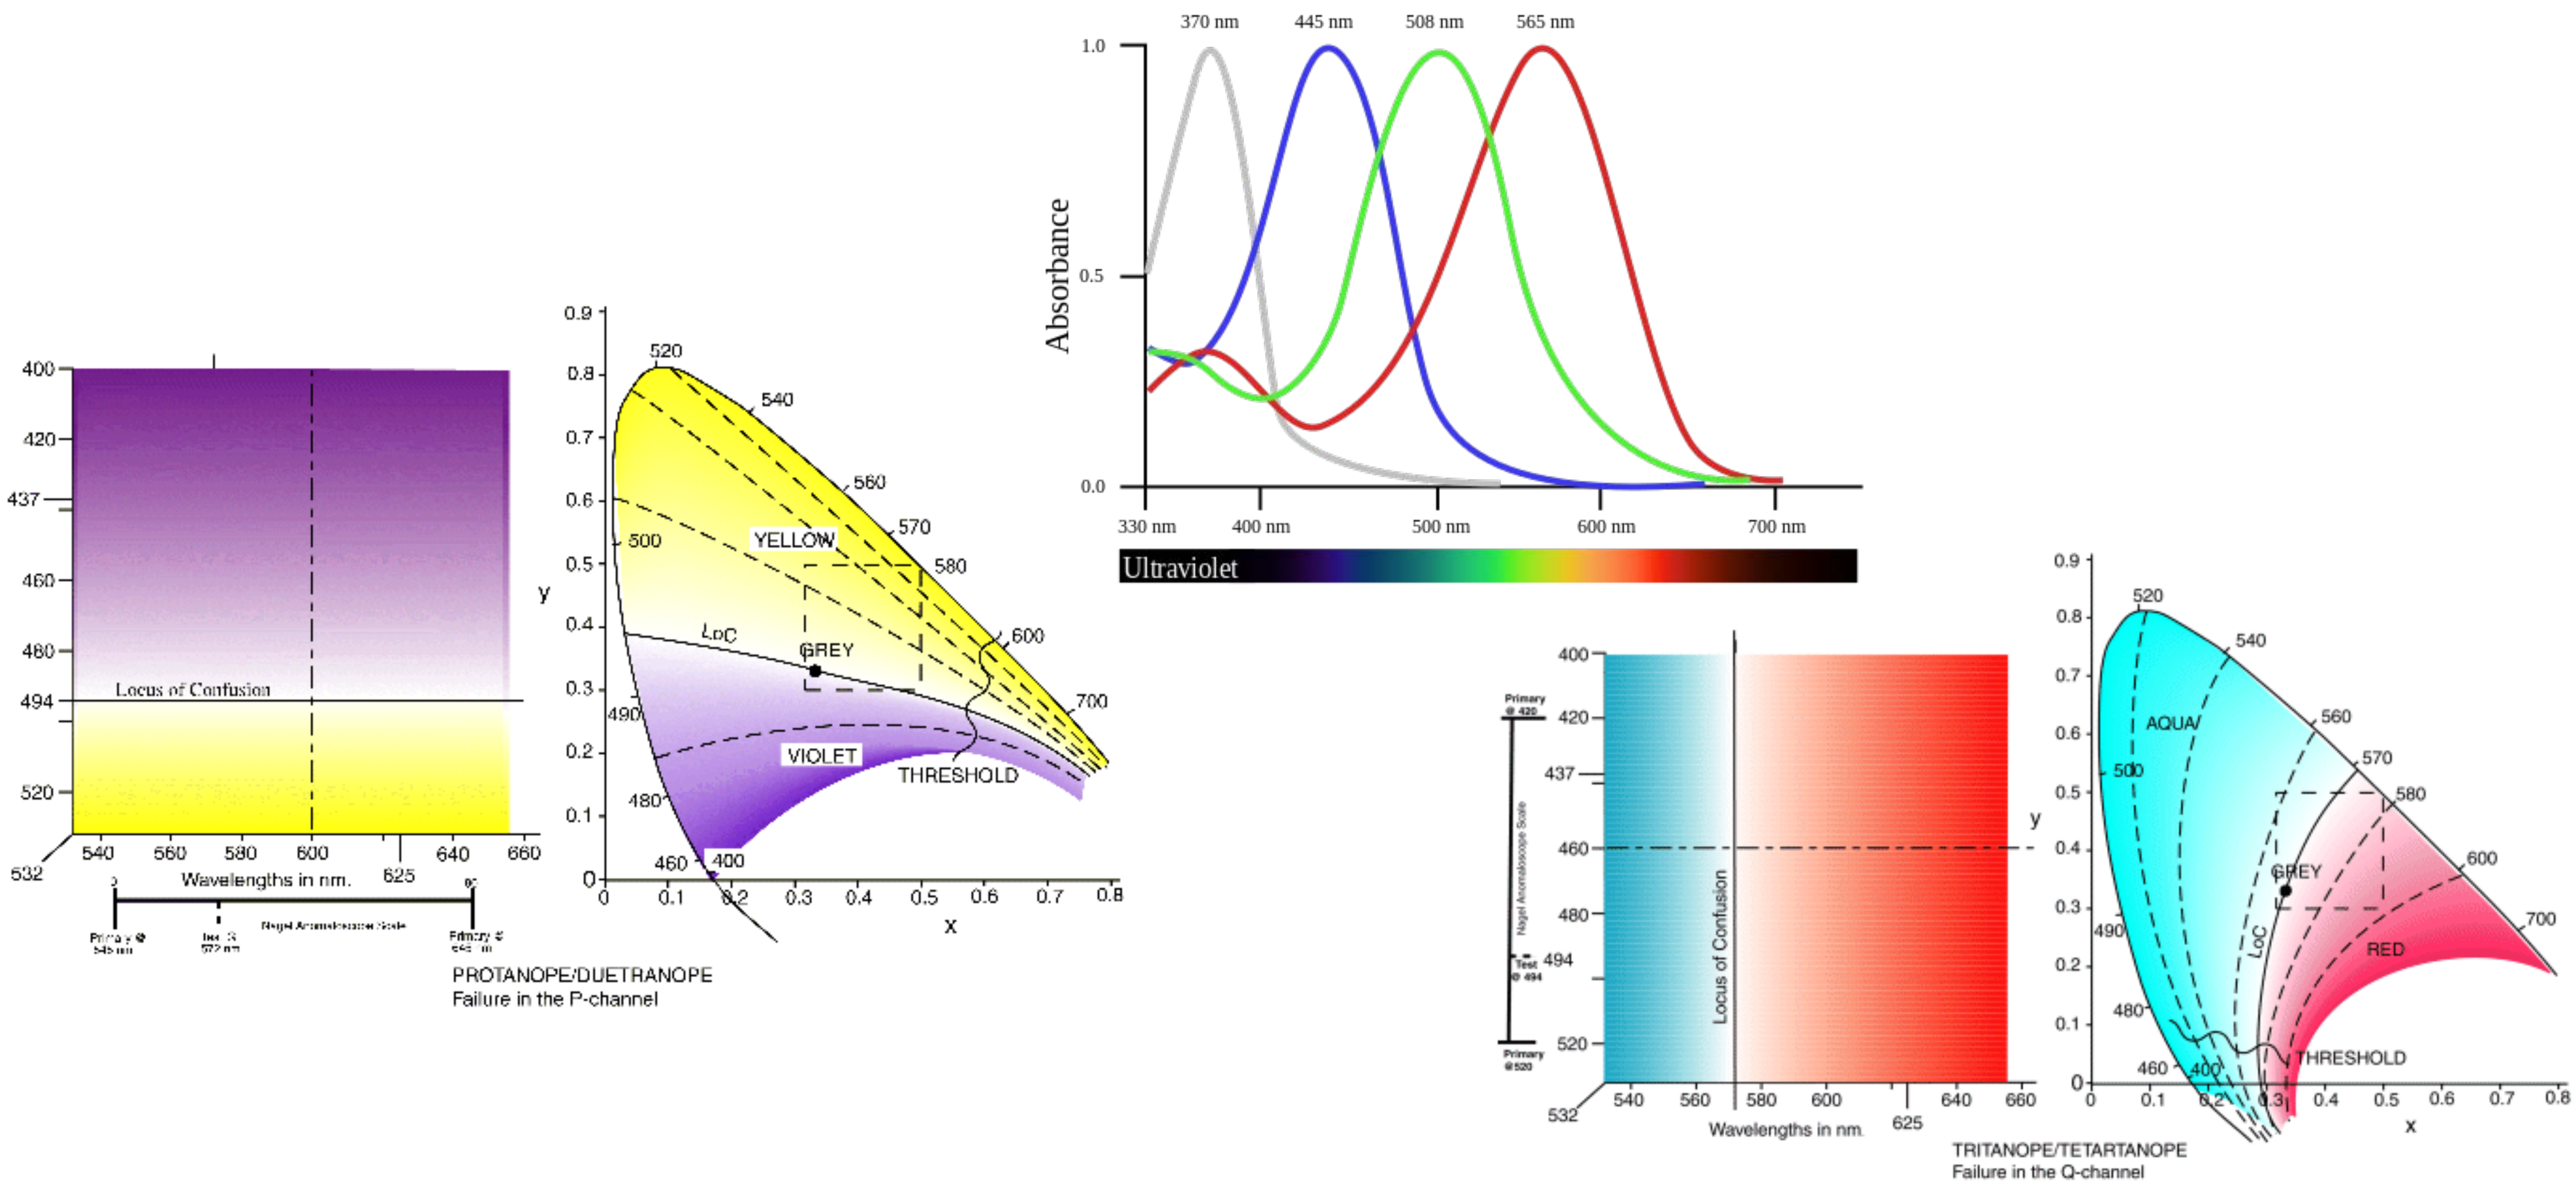

PROTANOPE/DUETRANOPE
Failure in the P-channel

TRITANOPE/TETARTANOPE
Failure in the Q-channel

Color Conversion

- Given a color specified in one model/space (e.g., sRGB), try to find corresponding color in another model (e.g., CMYK)
- In a perfect world: want to match *output spectrum*
- Even matching *perception* of color would be terrific (metamers)
- In reality: may not always be possible!
    - Depends on the gamut of the output device
    - E.g., VR headset vs. inkjet printer
- Complicated task!
- Lots of standards & software
    - ICC Profiles
    - Adobe Color Management, …

Gamma correction

(non-linear correction for CRT display)

Old CRT display: ⟶

1. Image contains value X
2. CRT display converts digital signal to an electron beam voltage V(x) (linear relationship)
3. Electron beam voltage converted to light: (non-linear relationship)

$$Y \propto V^{\gamma}$$

Where: $\gamma \approx 2.5$

(Doesn't apply to modern LCD displays, whose luminance output is linearly proportional to input; DOES still apply to other devices, like sensors, etc.)

So if pixels store Y, what will the display's output look like?

Fix: pixels sent to display must store:

$$Y^{1/2.5} = Y^{0.4}$$

Observed display output

Desired display output

Image credit: http://creativebits.org/mac_os_x/windows_vs_mac_monitor_gamma

Human Perception—Acommodation Effect

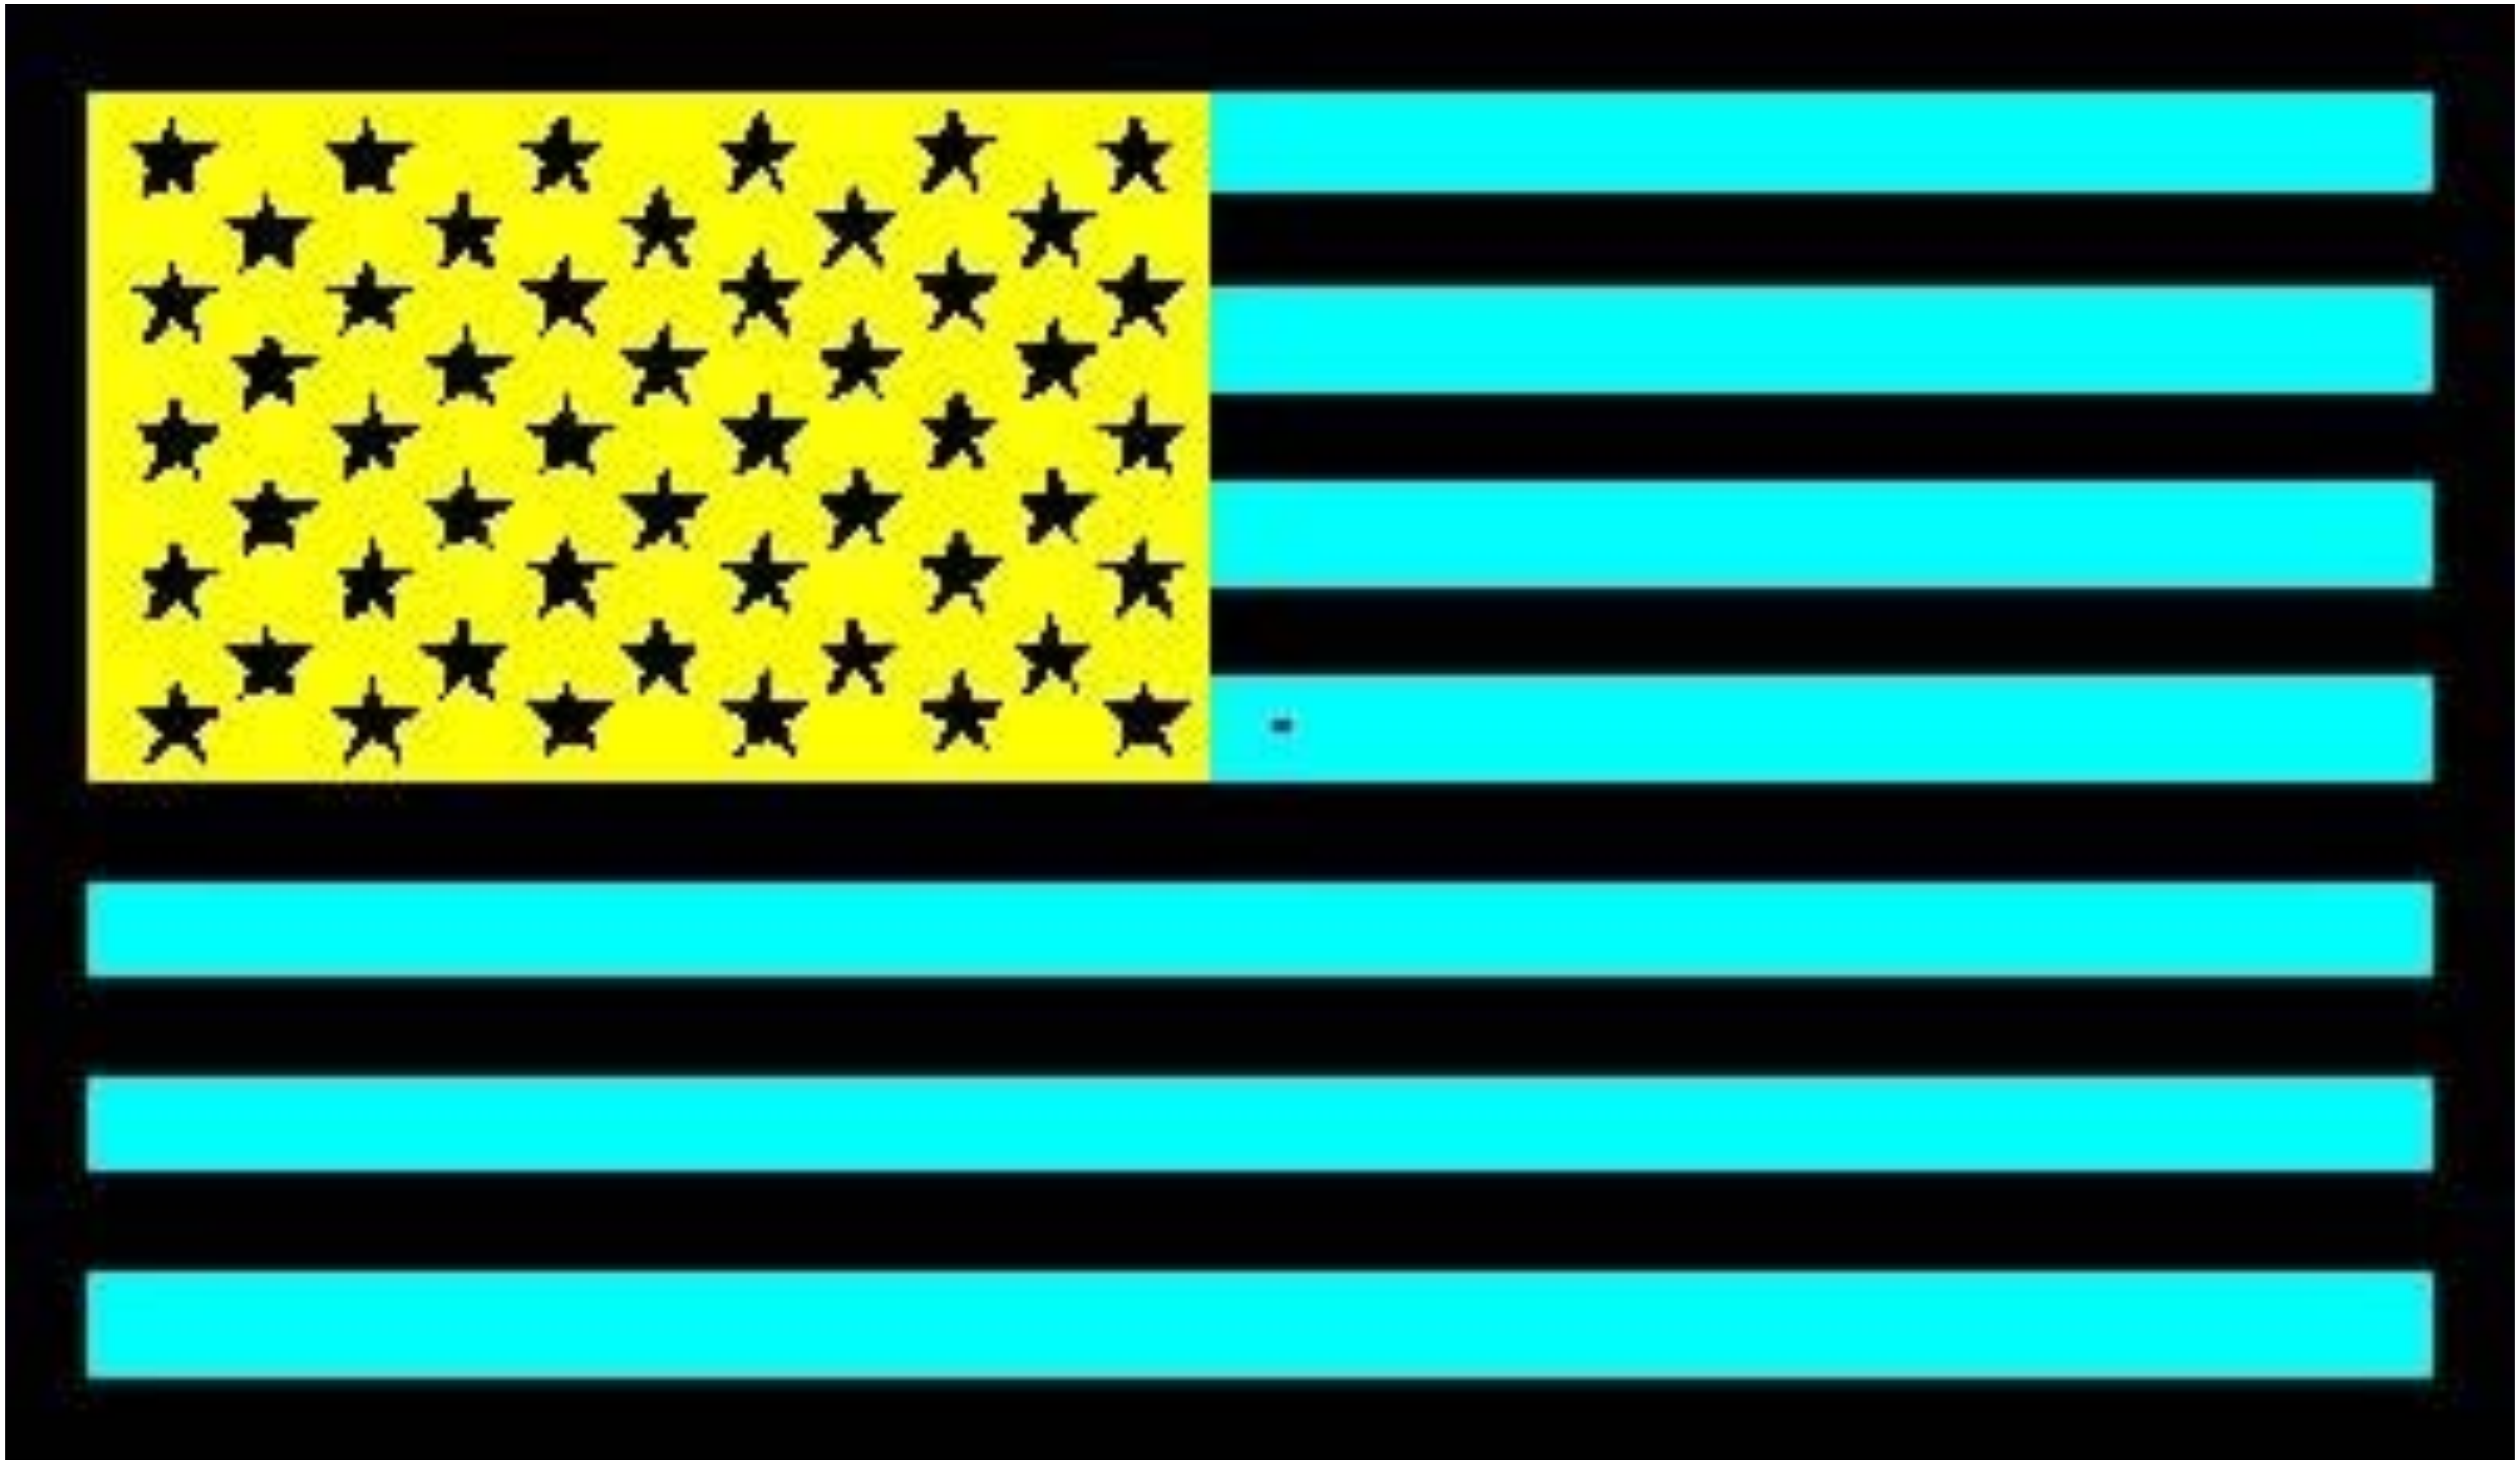

Human Perception—Acommodation Effect

Next time…

- A whole spectrum of things to know about light & color
- In the next few lectures we'll talk more about
  - radiometry
  - cameras
  - scattering
  - …