COMP 590/776: Computer Vision

Lecture 3: Colors: Human Vision & Computer Vision

Instructor: Soumyadip (Roni) Sengupta

ULA: Andrea Dunn, William Li, Liujie Zheng

Course Website:
Scan Me!

Few Announcements

- Assignment due Friday midnight, free extension until next Tuesday midnight.
- TA Office hour: Thursday 4.30-5.30 (SN 252), today via Zoom.
- Few changes will be made to the schedule, to be announced next week.

Etymology

PHOTOGRAPHY

light

drawing / writing

Camera = light-measuring device

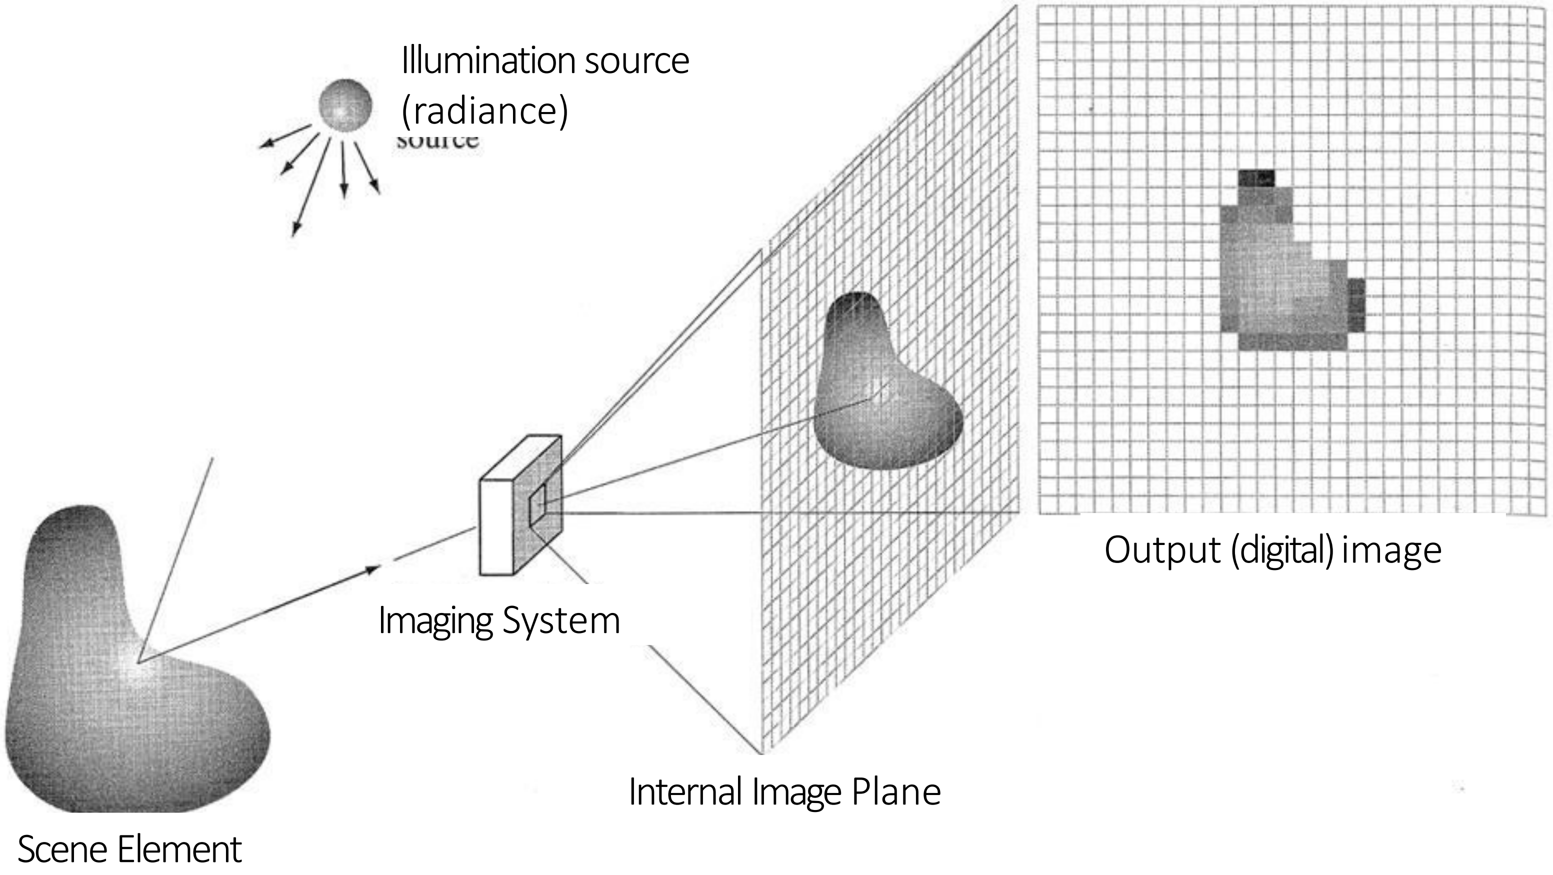

Simple models of a camera assumes an image is a “quantitative measurement” of scene radiance.

Figure from Digital Image Processing, Gonzales/Woods

This assumption is made often in computer vision

- Shape from shading
- HDR imaging
- Image matching
- Color constancy
- Applications relying on color
- Image delubrring
- Etc …

Shape-from-shading

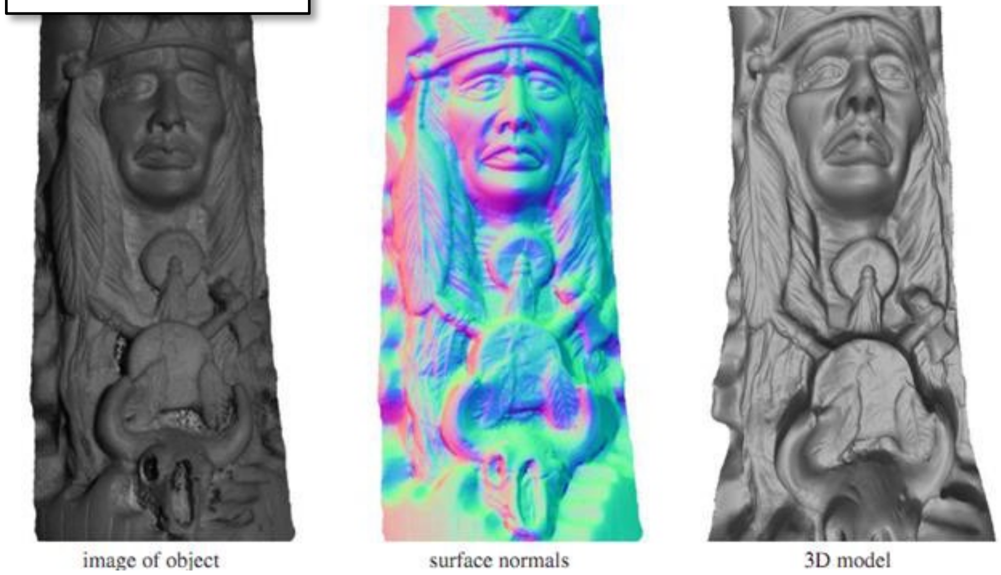

From Lu et al, CVPR'10

Image matching

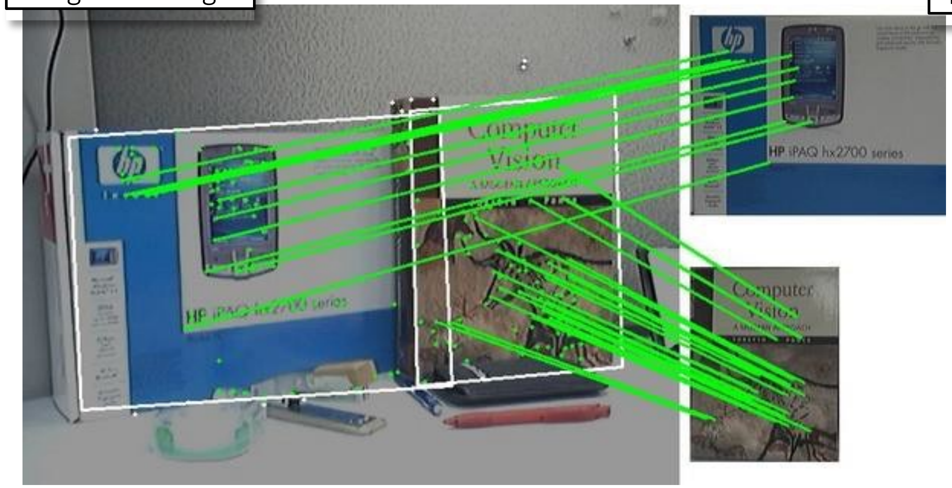

From Jon Mooser, CGIT Lab, USC

HDR imaging

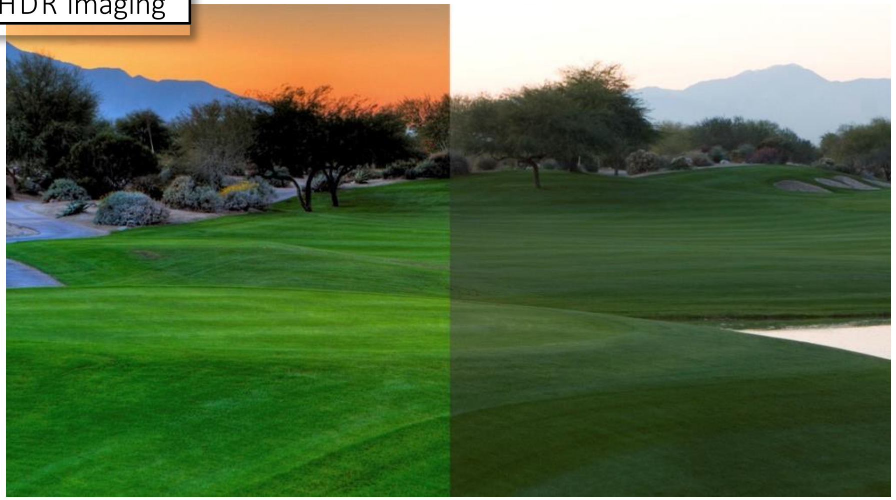

From O'Reilly's digital media forum

Camera = light-measuring device?

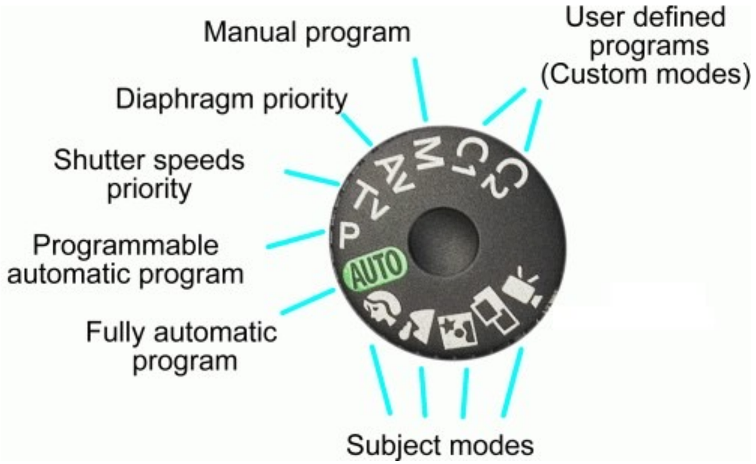

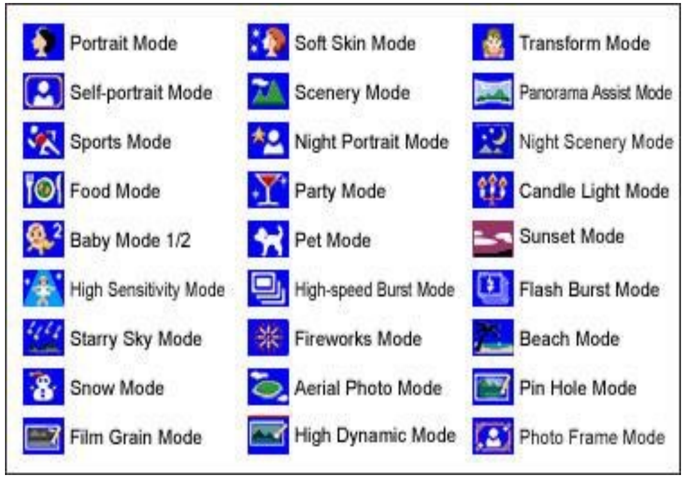

Camera pipeline photo-finishing routines

“Secret recipe” of a camera

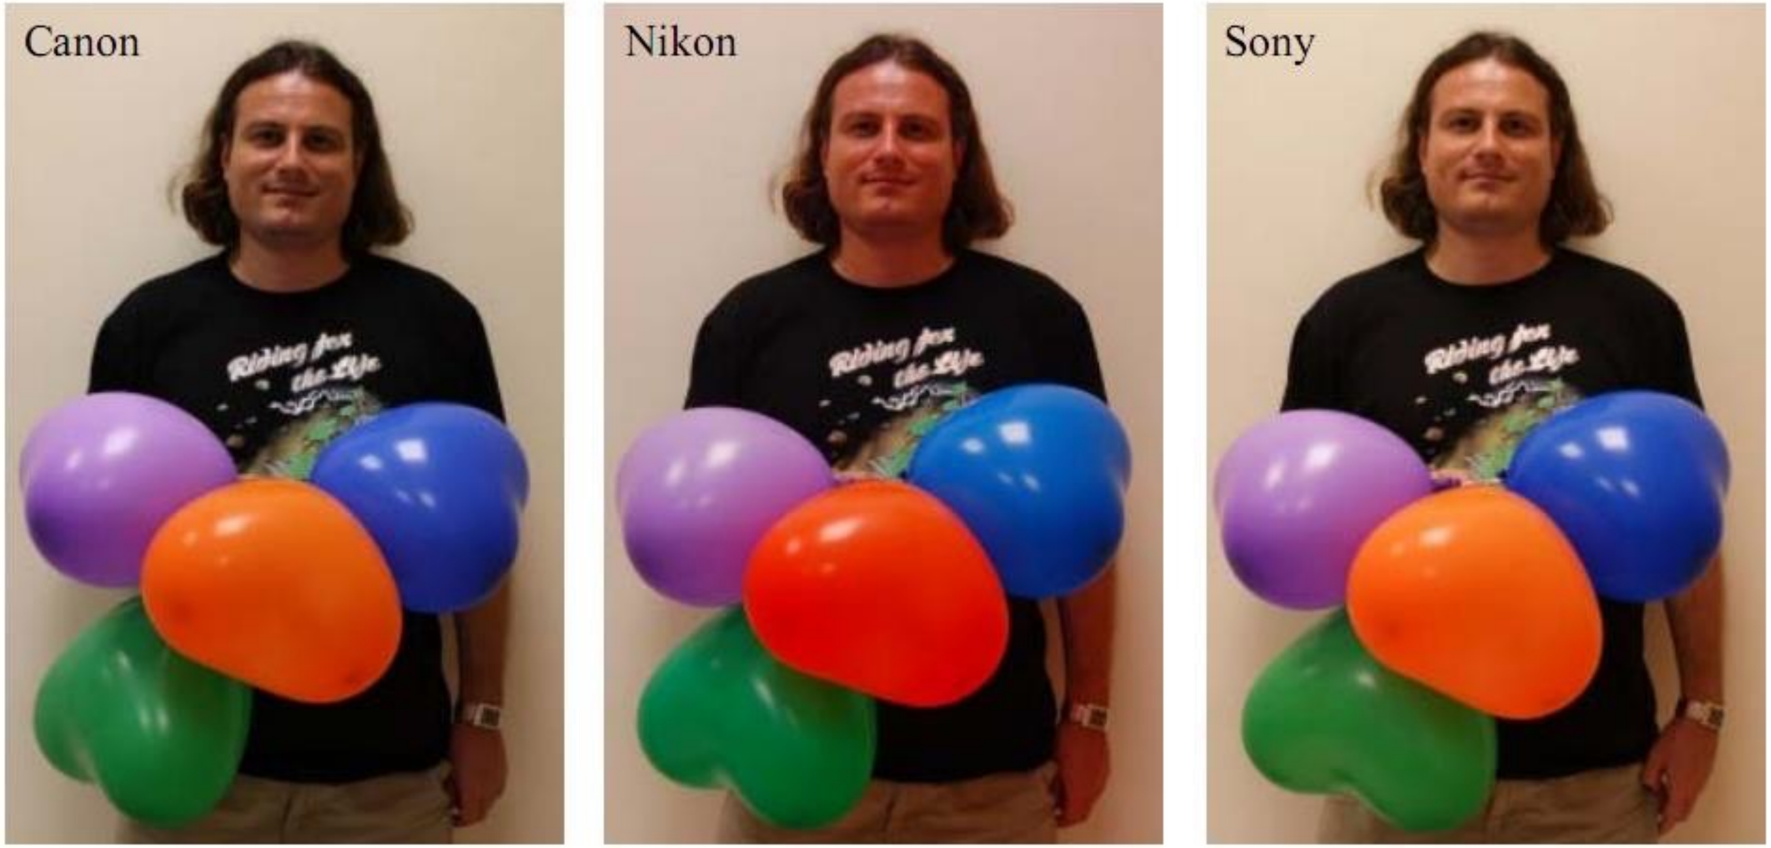

Photographs taken from three different cameras with the same aperture, shutter speed, white-balance ,ISO, and picture style.

Modern photography pipeline

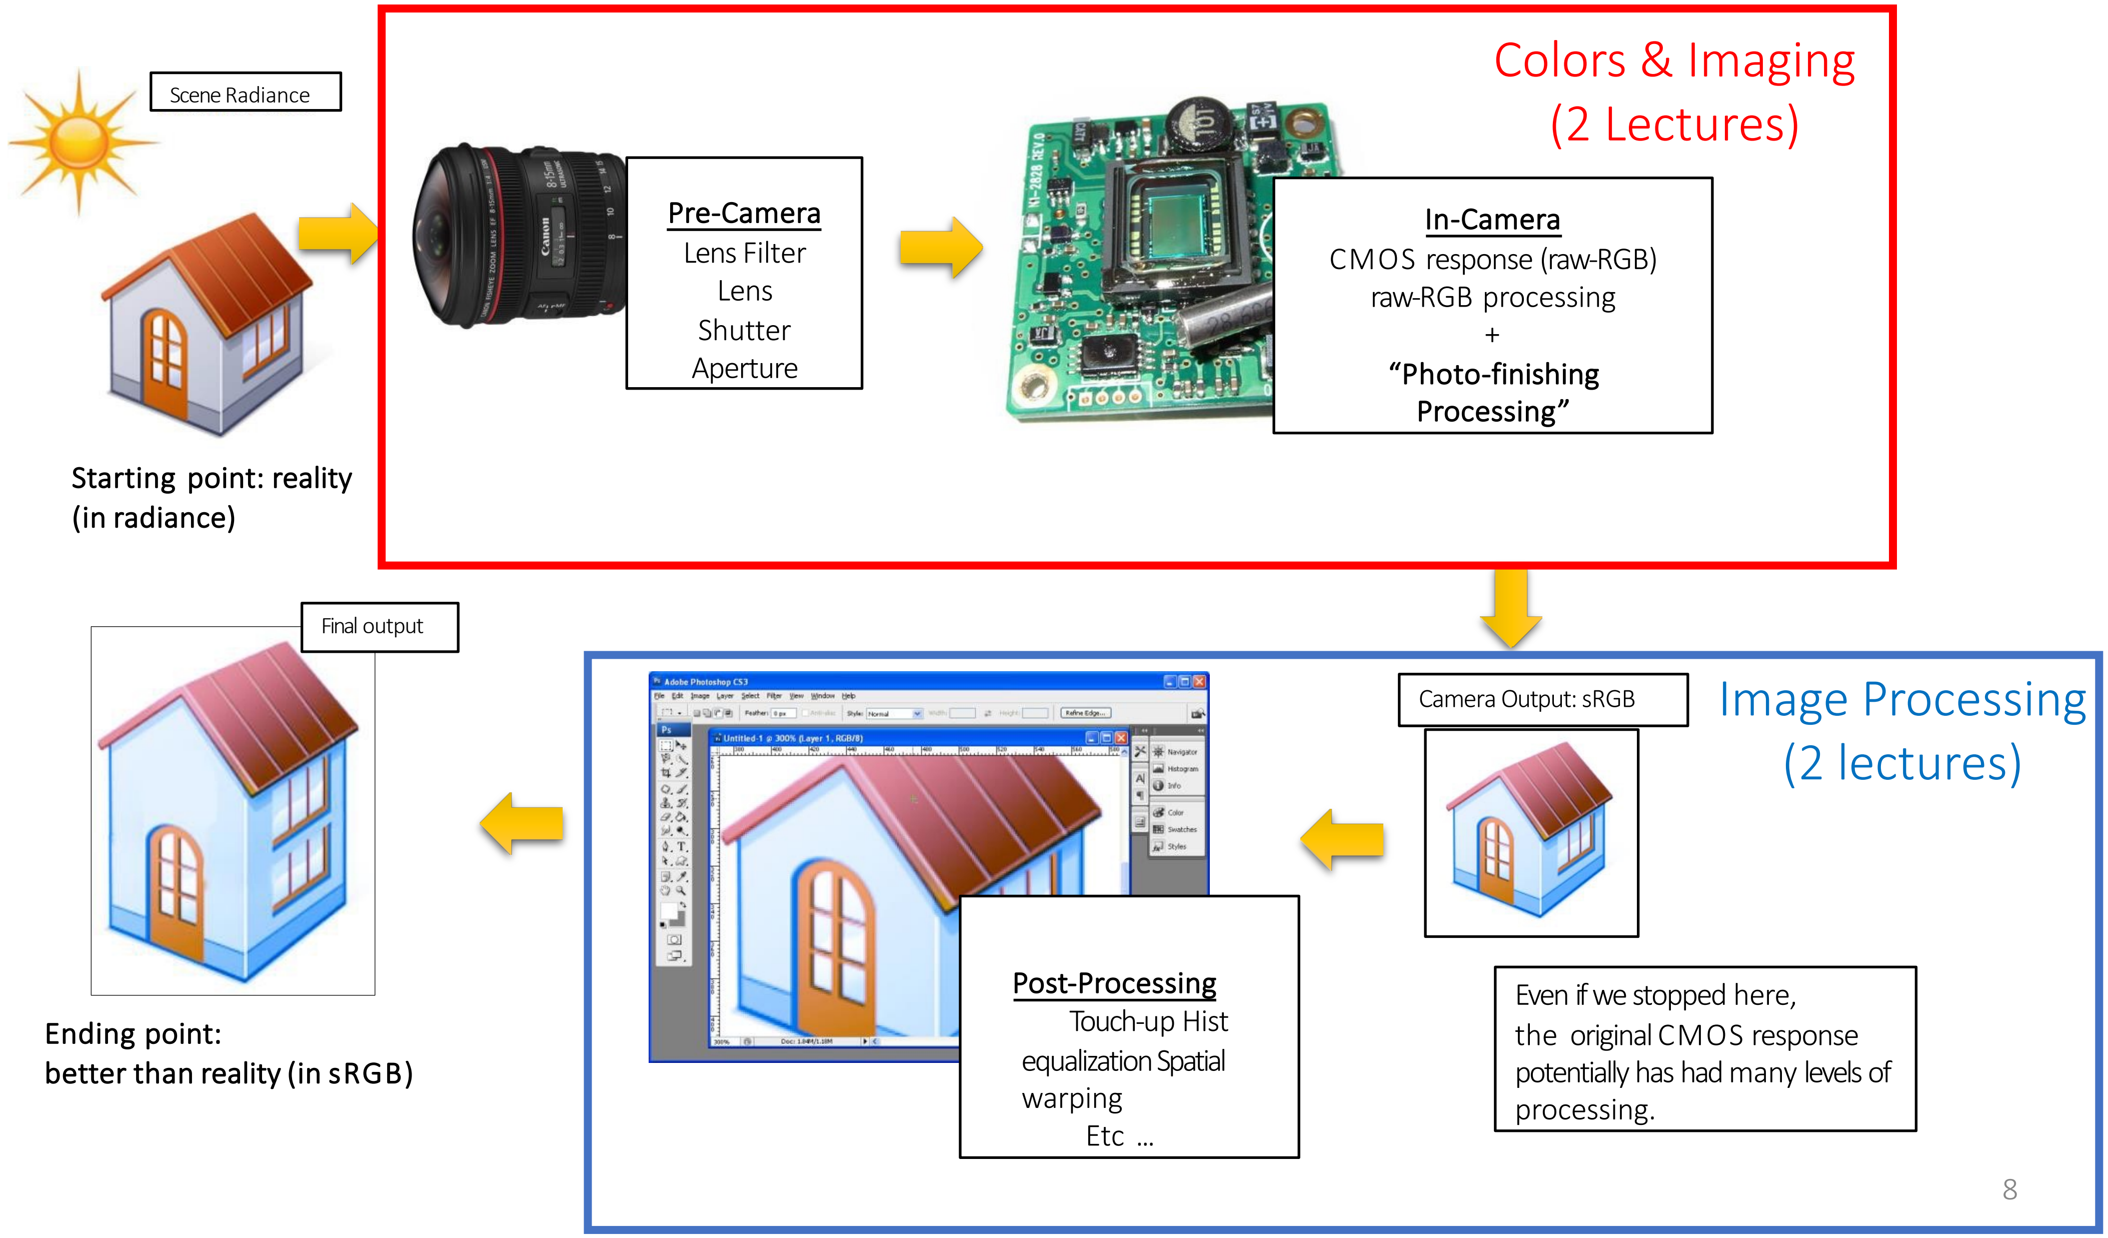

Digital cameras

- Digital cameras ***are not designed to be*** light-measuring devices
- **They are designed to produce visually pleasing photographs**
- There is a great deal of processing (photo-finishing) applied in the camera hardware

But what about other forms of cameras?

Do all ‘cameras’ have same pipeline?

Other Cameras

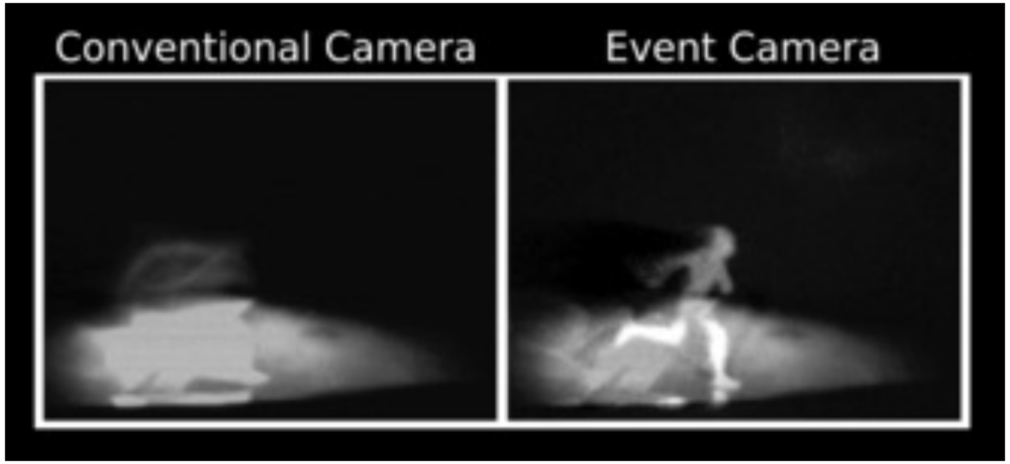

Event Cameras

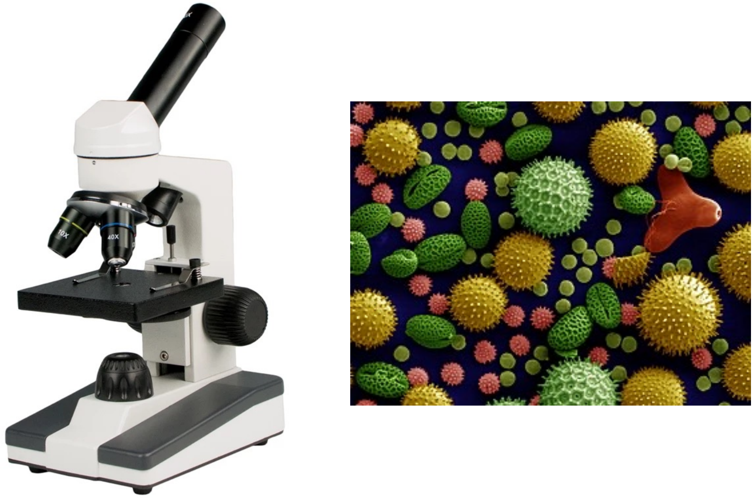

Microscope

"Computational Cameras and Displays"
Spring 2024, by Praneeth Chakravarthula

Perkin's Observatory Telescope

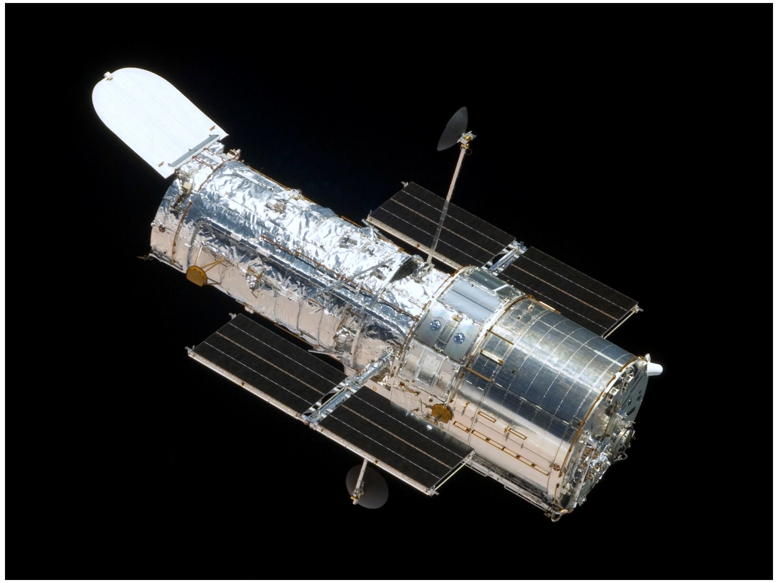

Hubble Space Telescope

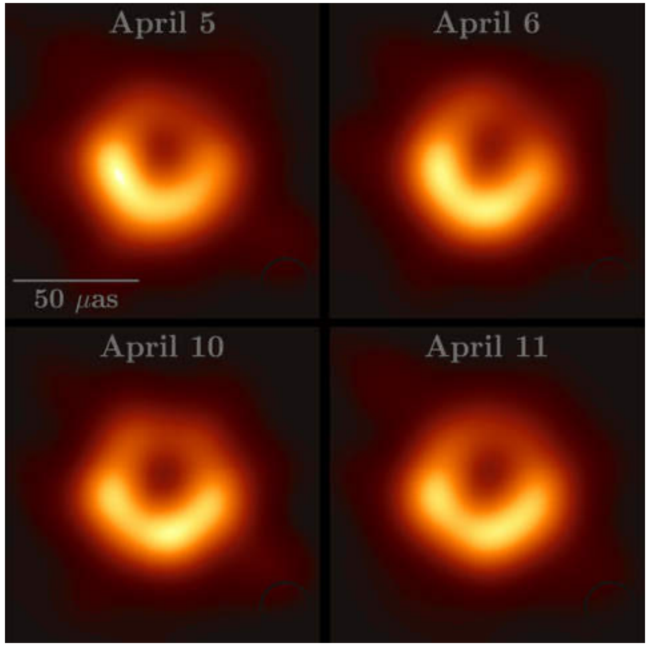

Imaging the black hole

Today’s class

- Eye & Human Vision
- Pinhole Camera
- Perception of Color
- Quantifying Color
- Commonly Used Color Spaces
- Color Constancy

Today's class

- Eye & Human Vision
- Pinhole Camera
- Perception of Color
- Quantifying Color
- Commonly Used Color Spaces
- Color Constancy

The Eye

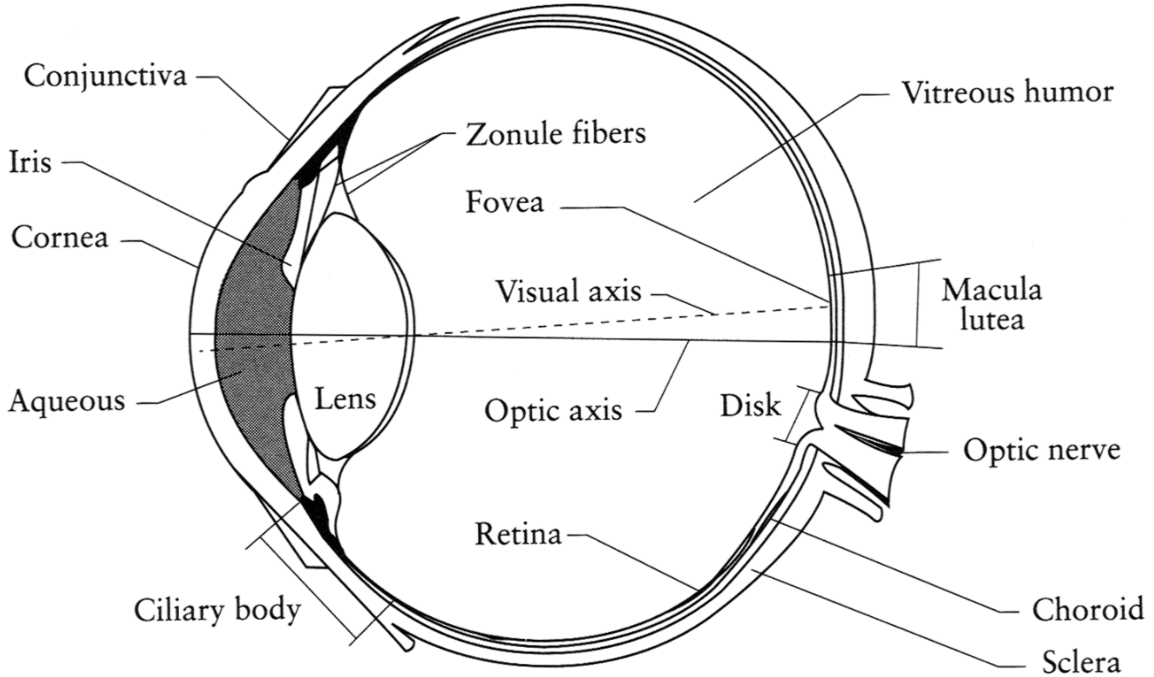

- The human eye is a camera!
  - **Iris** - colored annulus with radial muscles
  - **Pupil** - the hole (aperture) whose size is controlled by the iris
  - What's the "film"?
    - – photoreceptor cells (rods and cones) in the **retina**

Slide by Steve Seitz

Biology of color sensations

- Our eye has three receptors (cone cells) that respond to visible light and give the sensation of color

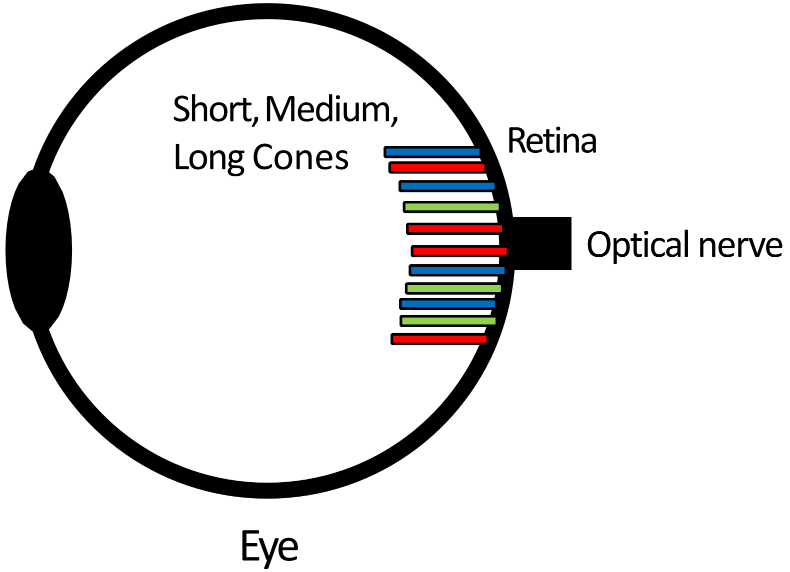

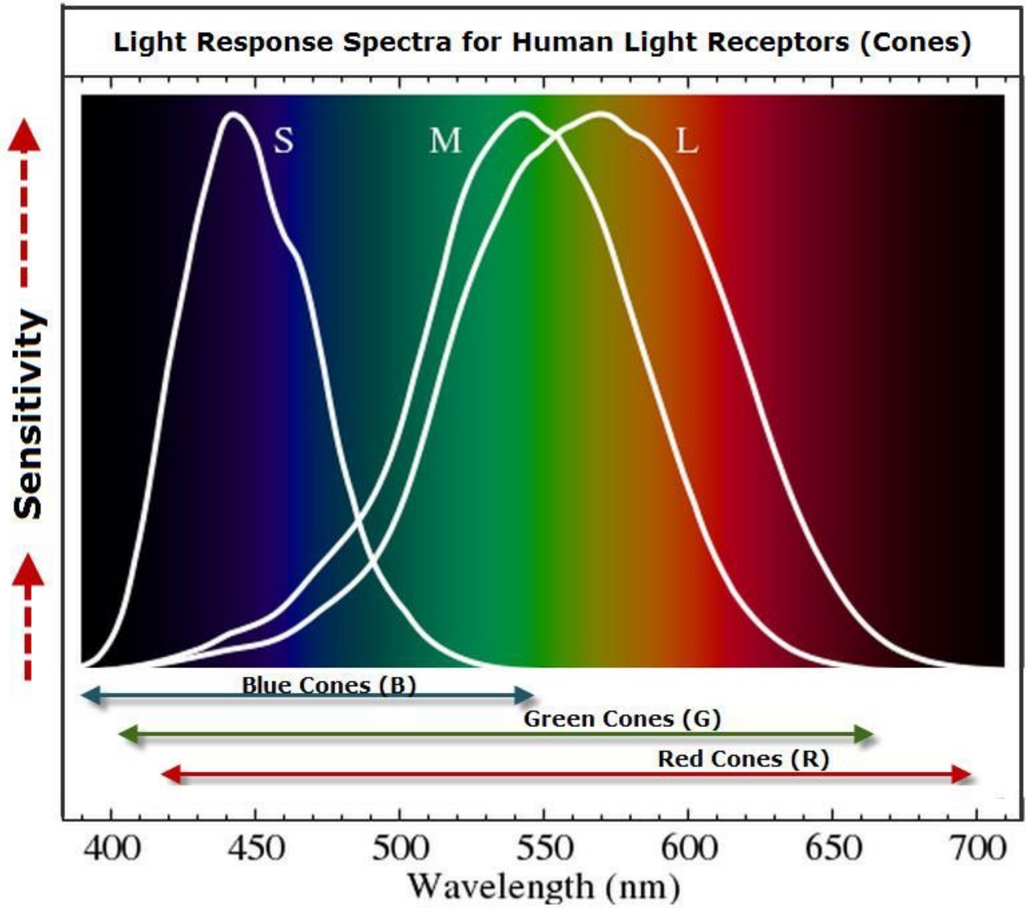

Retina up-close

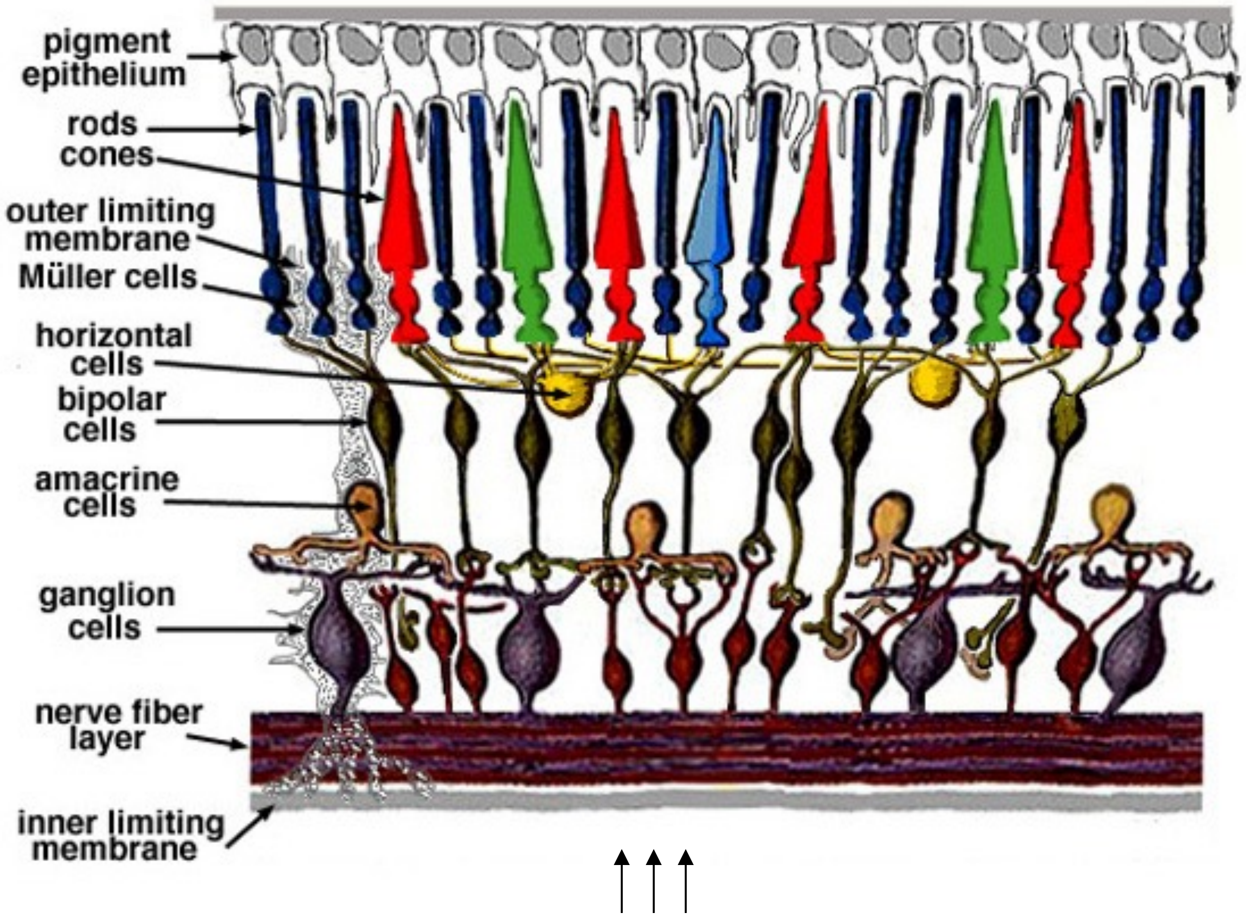

Light

Cones and rods

- We have additional light sensitive cells called *rods* that are not responsible for color.
- Rods are used in low-light vision.
- Cone cells are most concentrated around the fovea of the eye

Fovea = region of retina where visual activity is the highest, responsible for central vision.

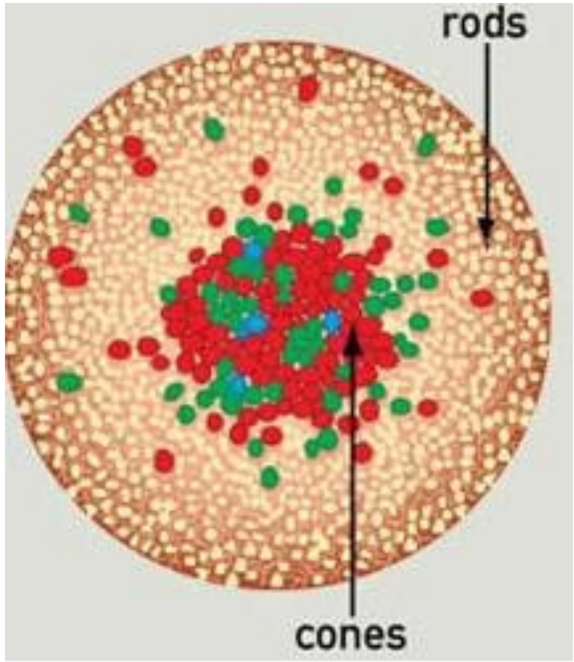

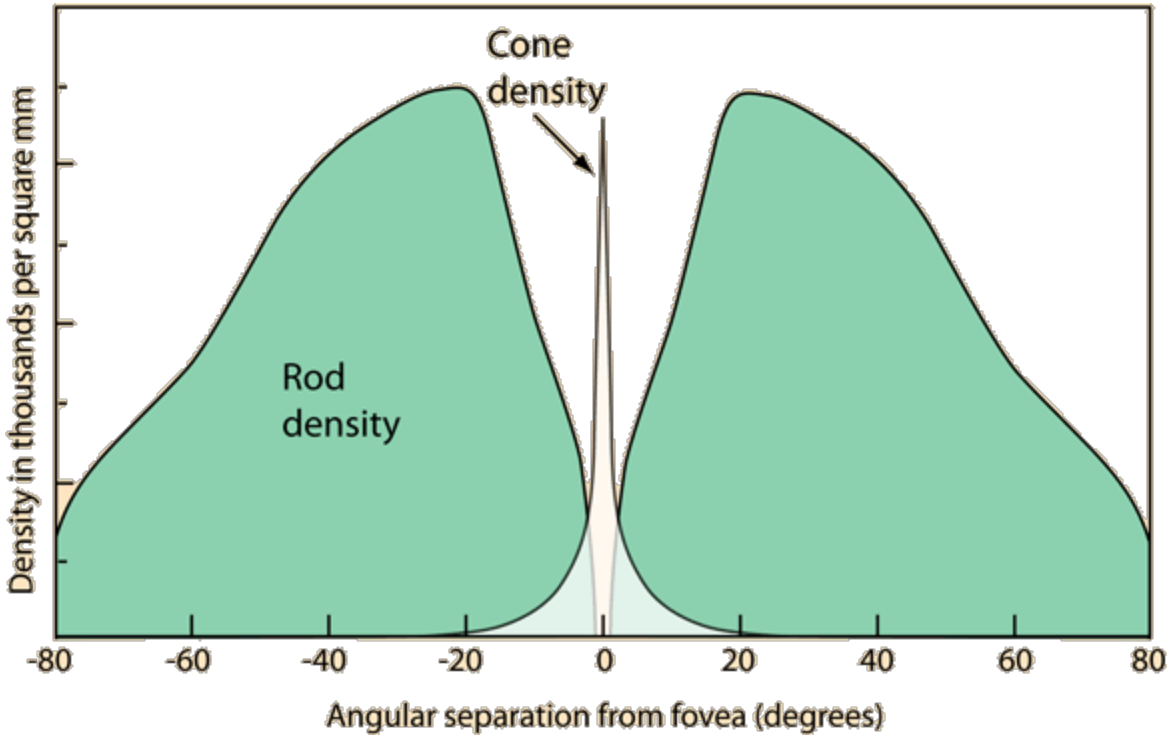

Distribution of Rods and Cones

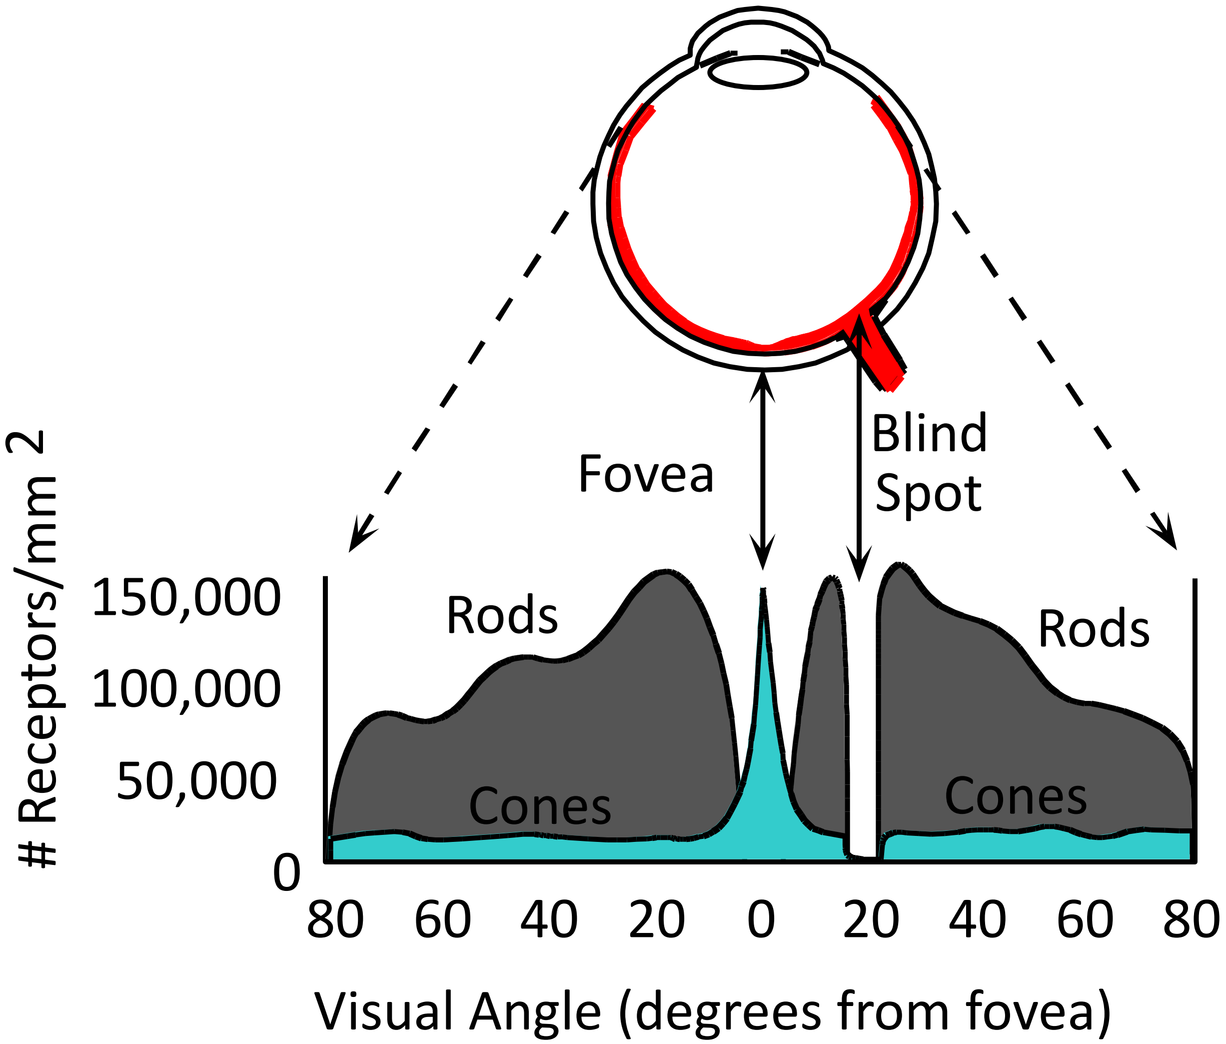

Night Sky: why are there more stars off-center?

Rod / Cone sensitivity

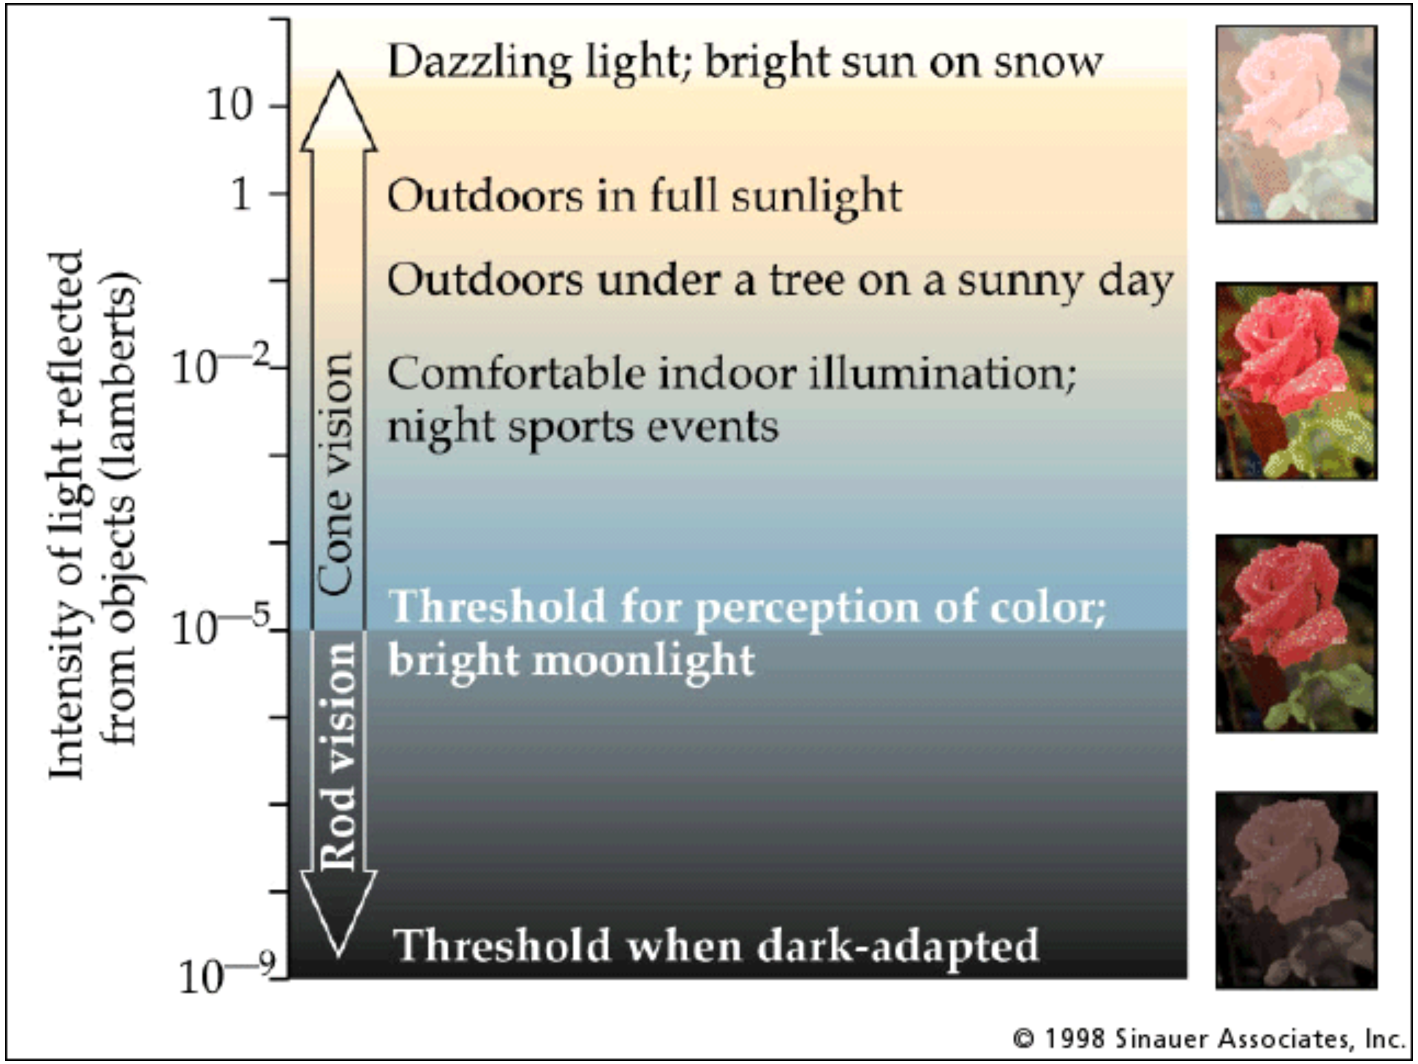

Today’s class

- Eye & Human Vision
- Pinhole Camera
- Perception of Color
- Quantifying Color
- Commonly Used Color Spaces
- Color Constancy

We will learn more about mathematical models of Pinhole camera and projection in 3D section

Eye vs Camera

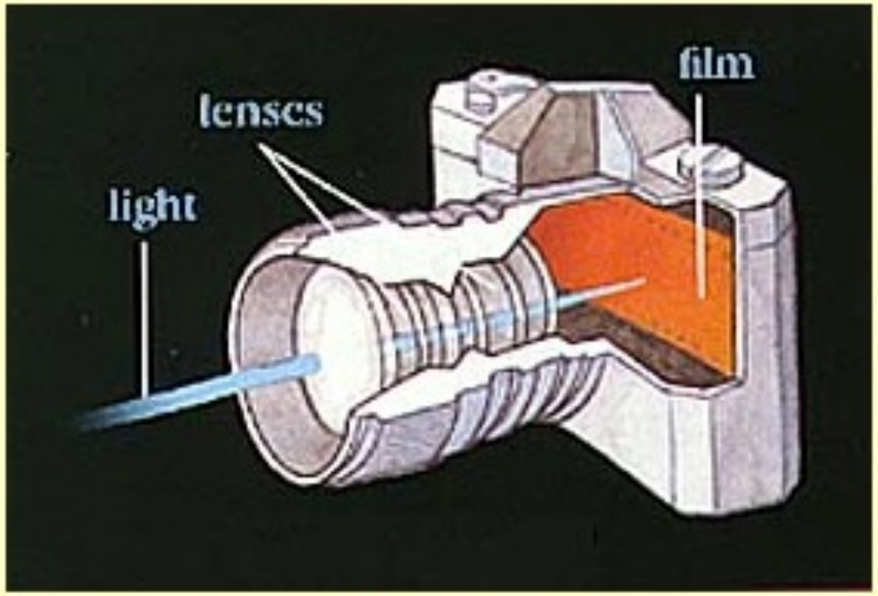

The Camera

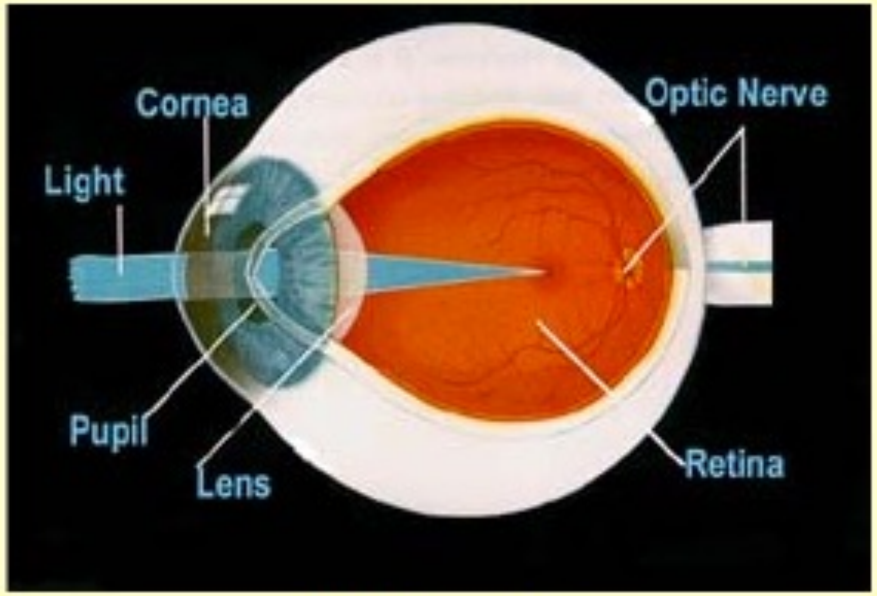

The Human Eye

	Human Eye	Camera
How does light enter?	Pupil	Aperture
What controls the amount of light?	Iris	Diaphragm
What interprets the image?	Retina	Film
How is the light focused?	Lens	Lens

Credits: https://slideplayer.com/slide/6029035/ & https://www.knowswhy.com/similarities-between-camera-and-human-eye/

Today's class

- Eye & Human Vision
- Pinhole Camera
- Perception of Color
- Quantifying Color
- Commonly Used Color Spaces
- Color Constancy

Color

Def *Color* (noun):The property possessed by an object of producing different sensations on the eye as a result of the way it reflects or emits light.

Oxford Dictionary

Color is perceptual

- **Color is not** a primary *physical* property on an object
- Red, Green, Blue, Pink, Orange, Atomic Tangerine, Baby Pink, etc., are just words we assign to human color sensations

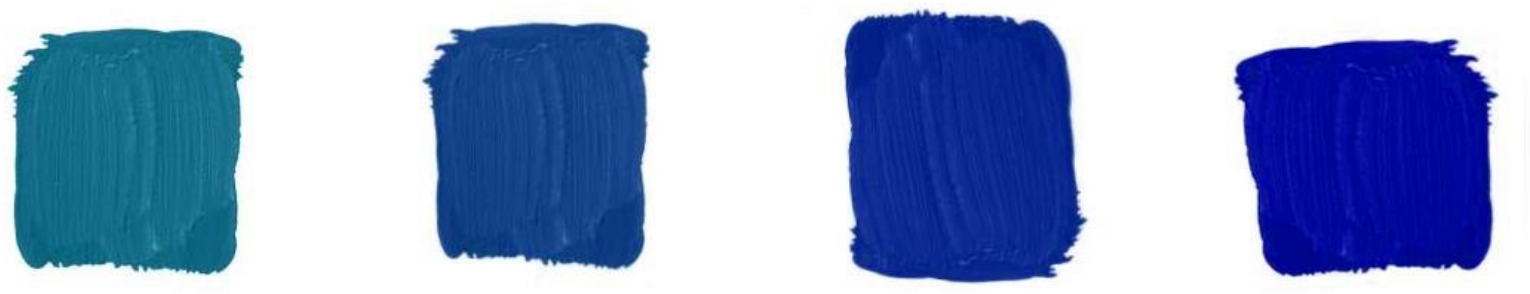

Which is the "true blue"?

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

Two types of light-sensitive receptors

Cones

- cone-shaped
- less sensitive
- operate in high light
- color vision

Rods

- rod-shaped
- highly sensitive
- operate at night
- gray-scale vision

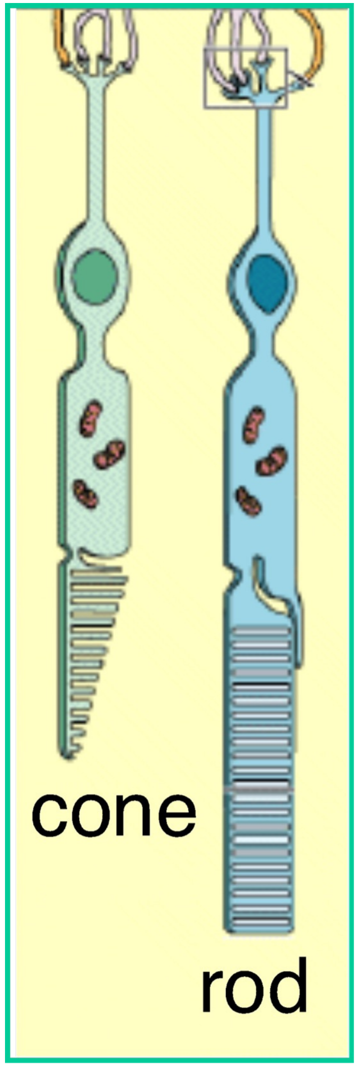

Where do "color sensations" come from?

A *very* small range of electromagnetic radiation

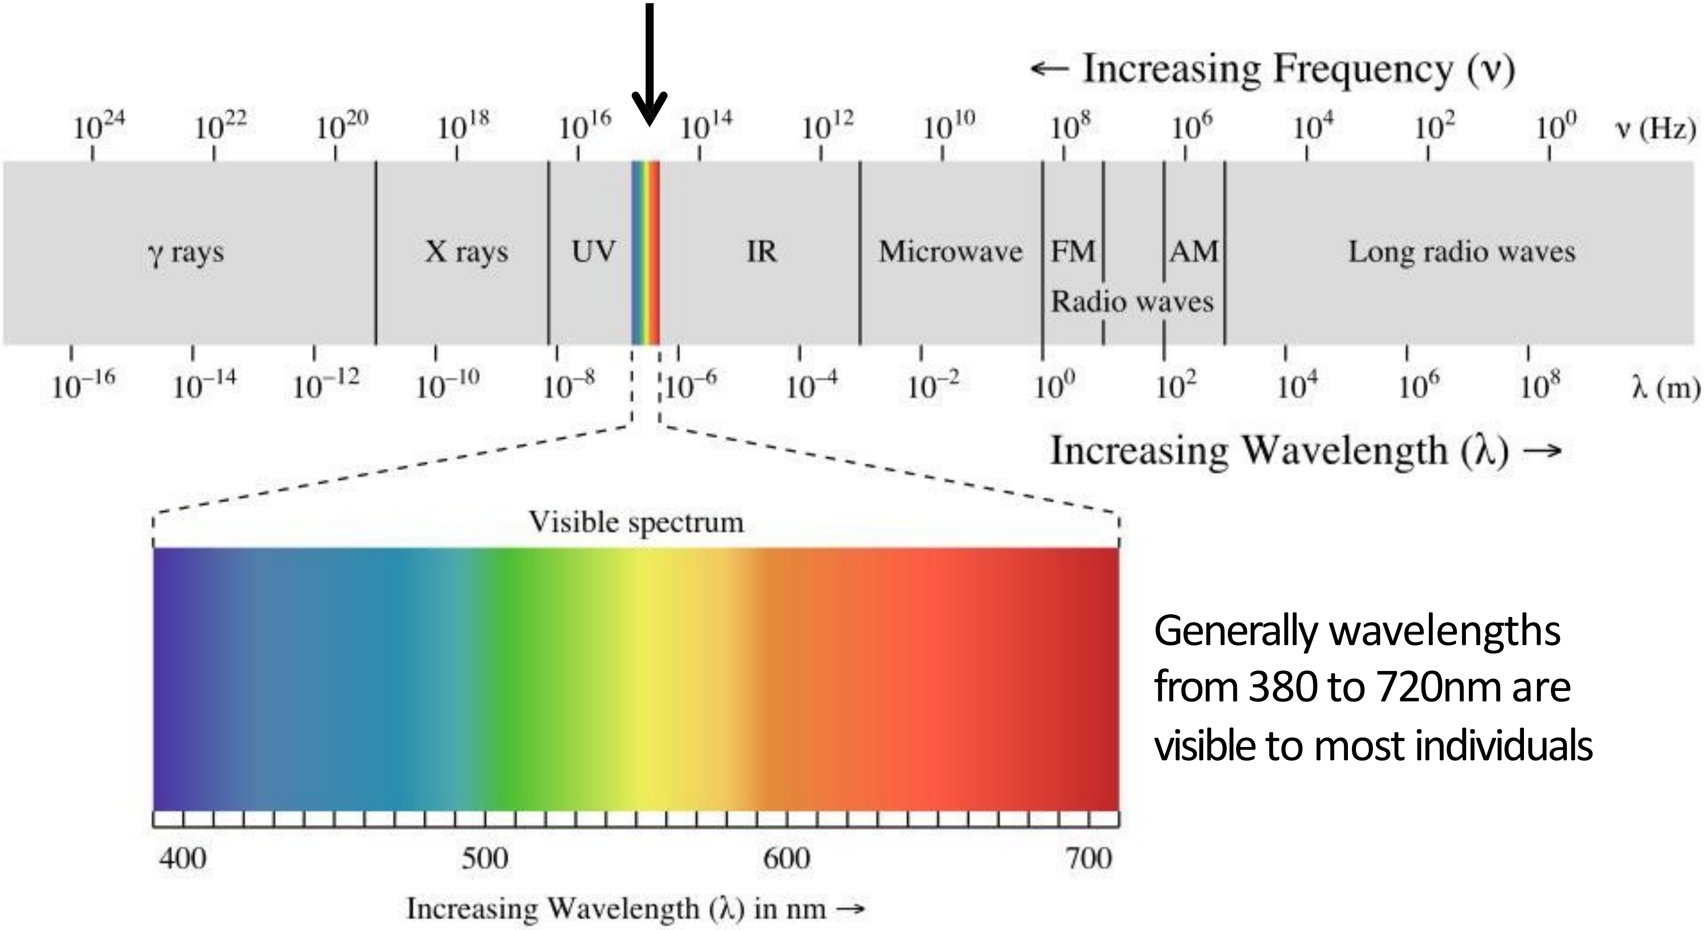

Generally wavelengths from 380 to 720nm are visible to most individuals

The Physics of Light

Any patch of light can be completely described physically by its spectrum: the number of photons (per time unit) at each wavelength 400 - 700 nm.

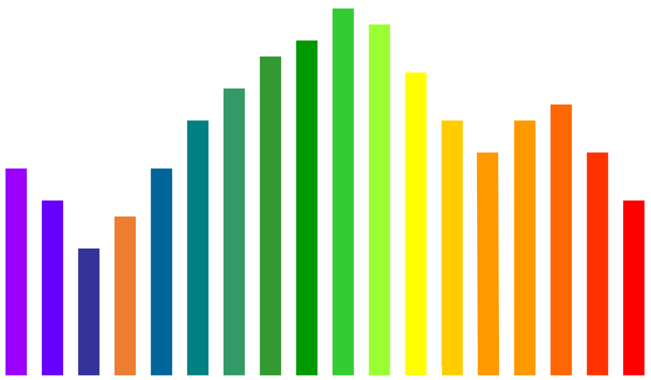

White light through a prism

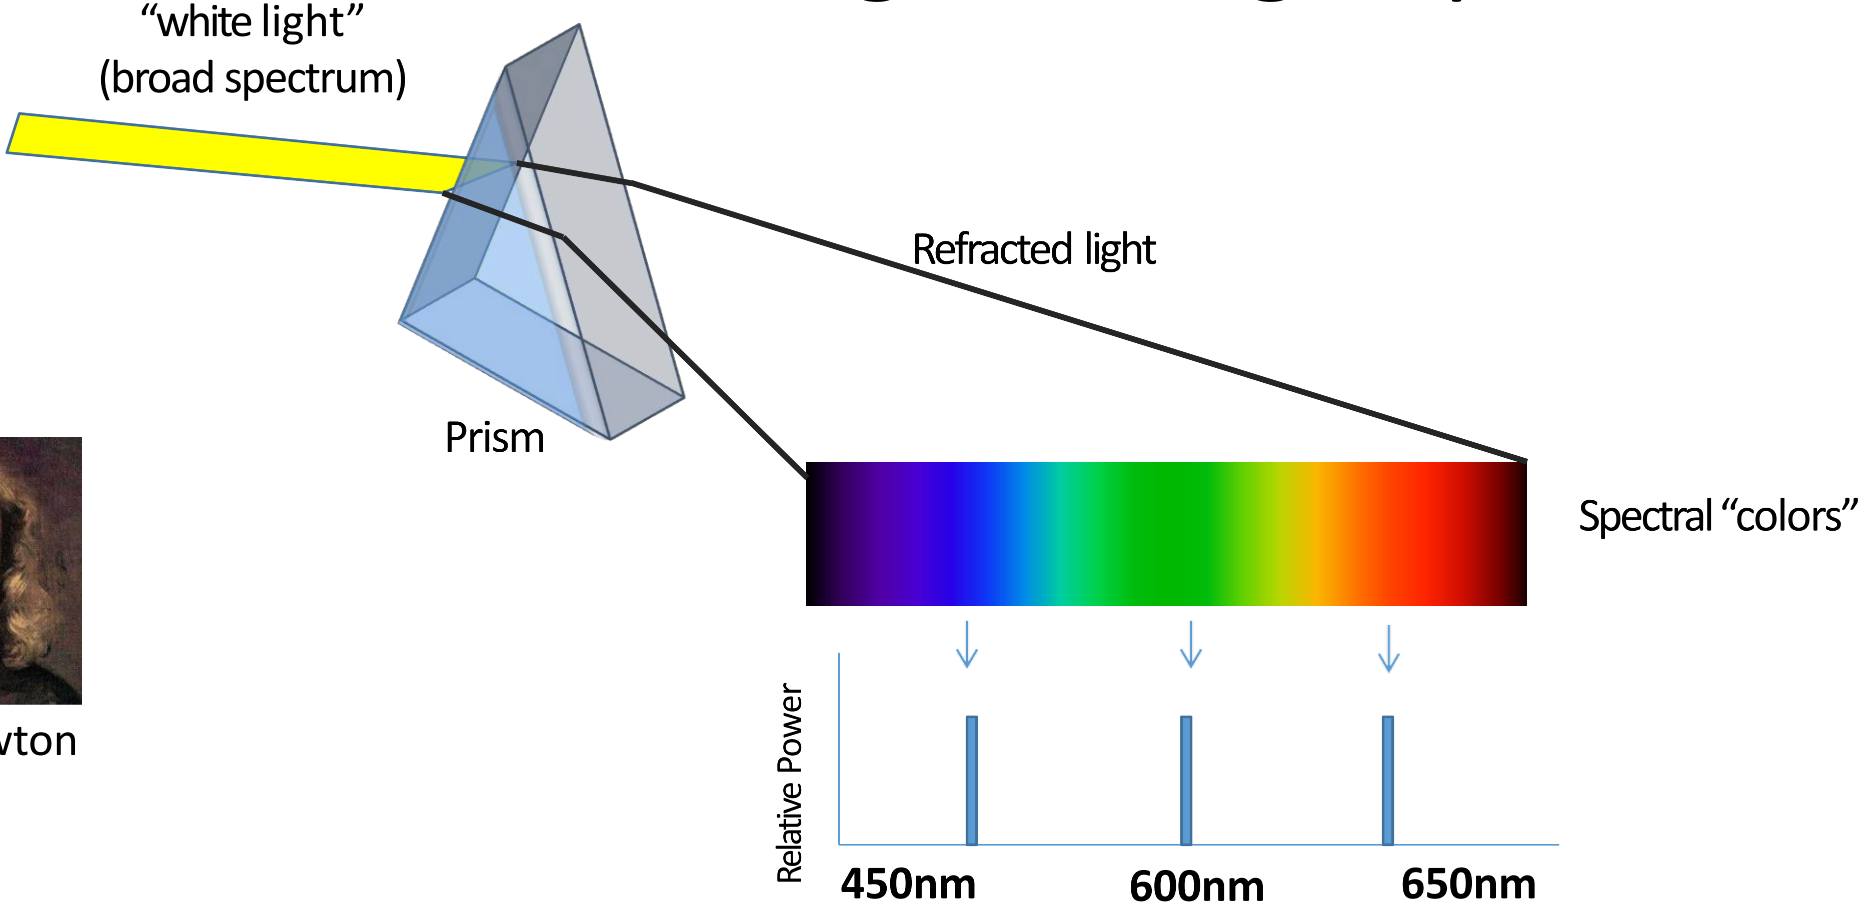

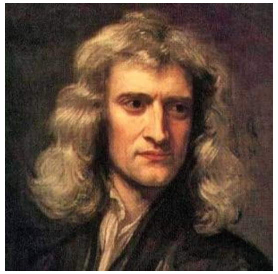
Isaac Newton

Light is separated into “monochromatic” light at different wave lengths.

Light-material interaction

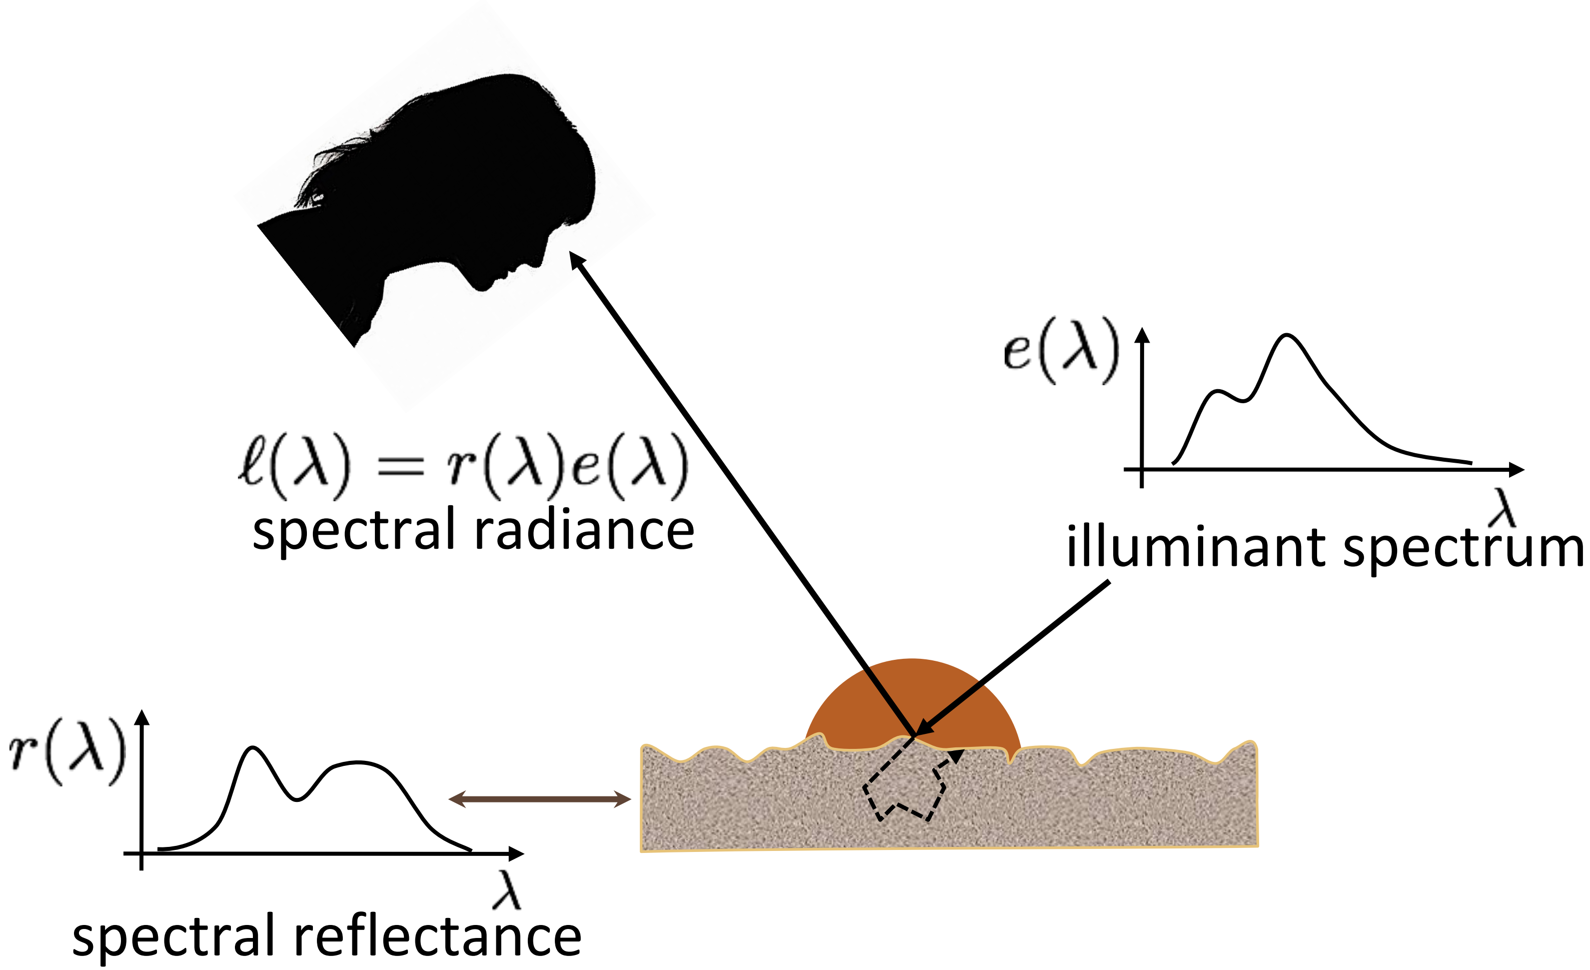

Light-material interaction

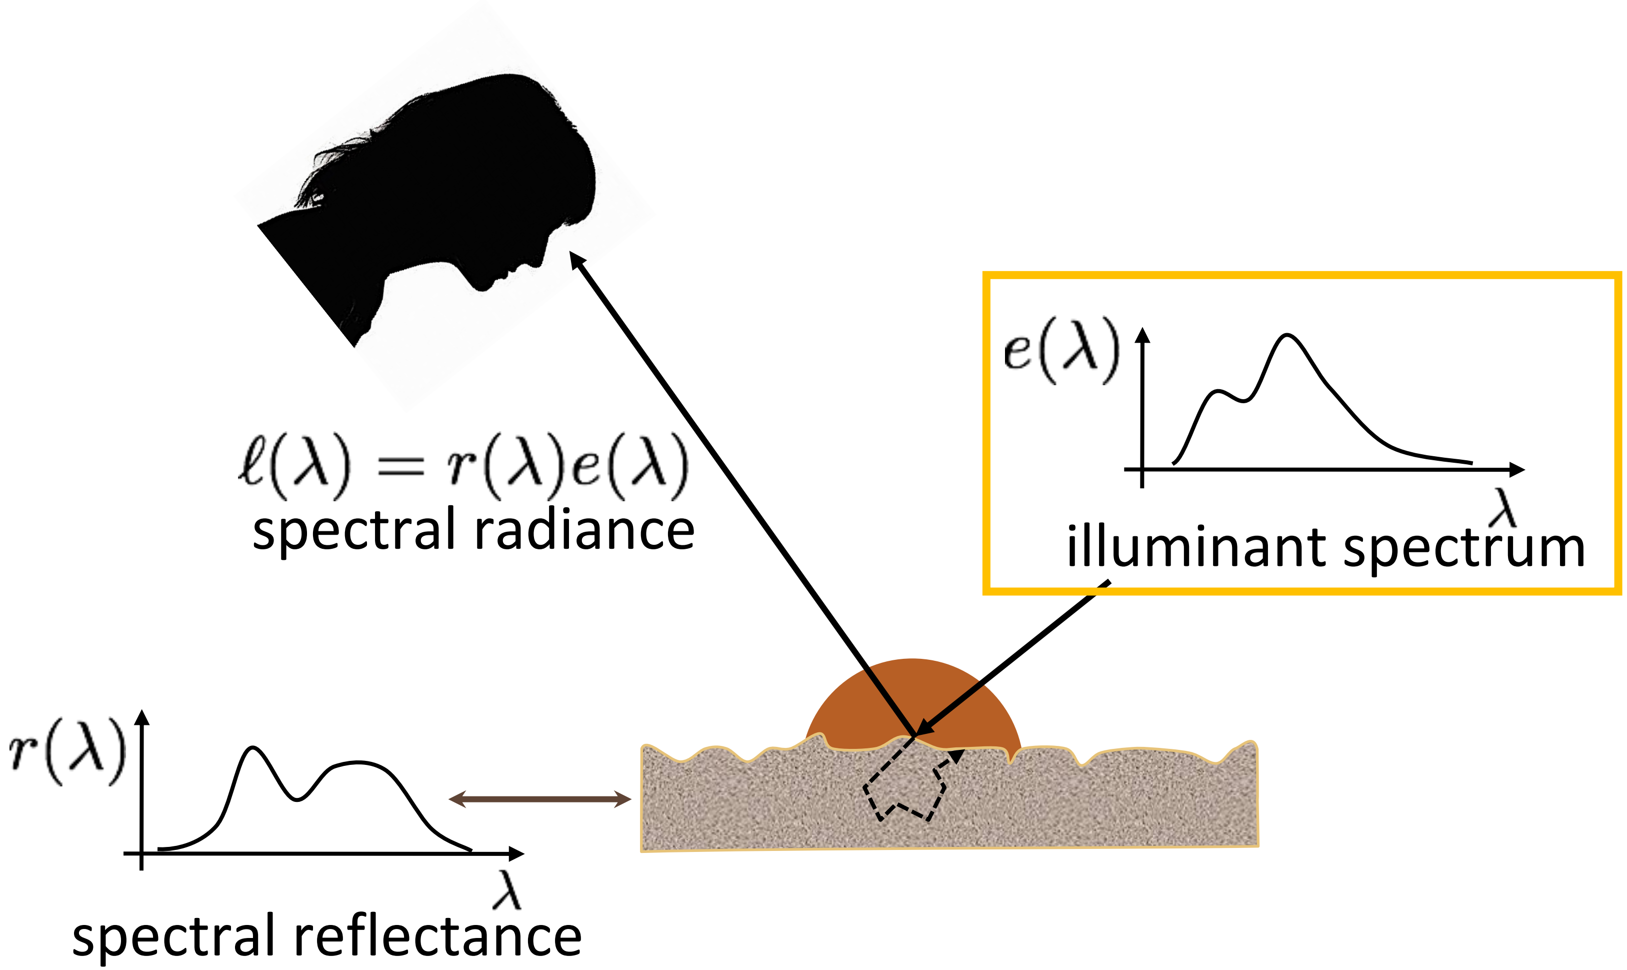

Illuminant Spectral Power Distribution (SPD)

- Most types of light “contain” more than one wavelengths.
- We can describe light based on the distribution of power over different wavelengths.

We call our sensation of all of these distributions “white”.

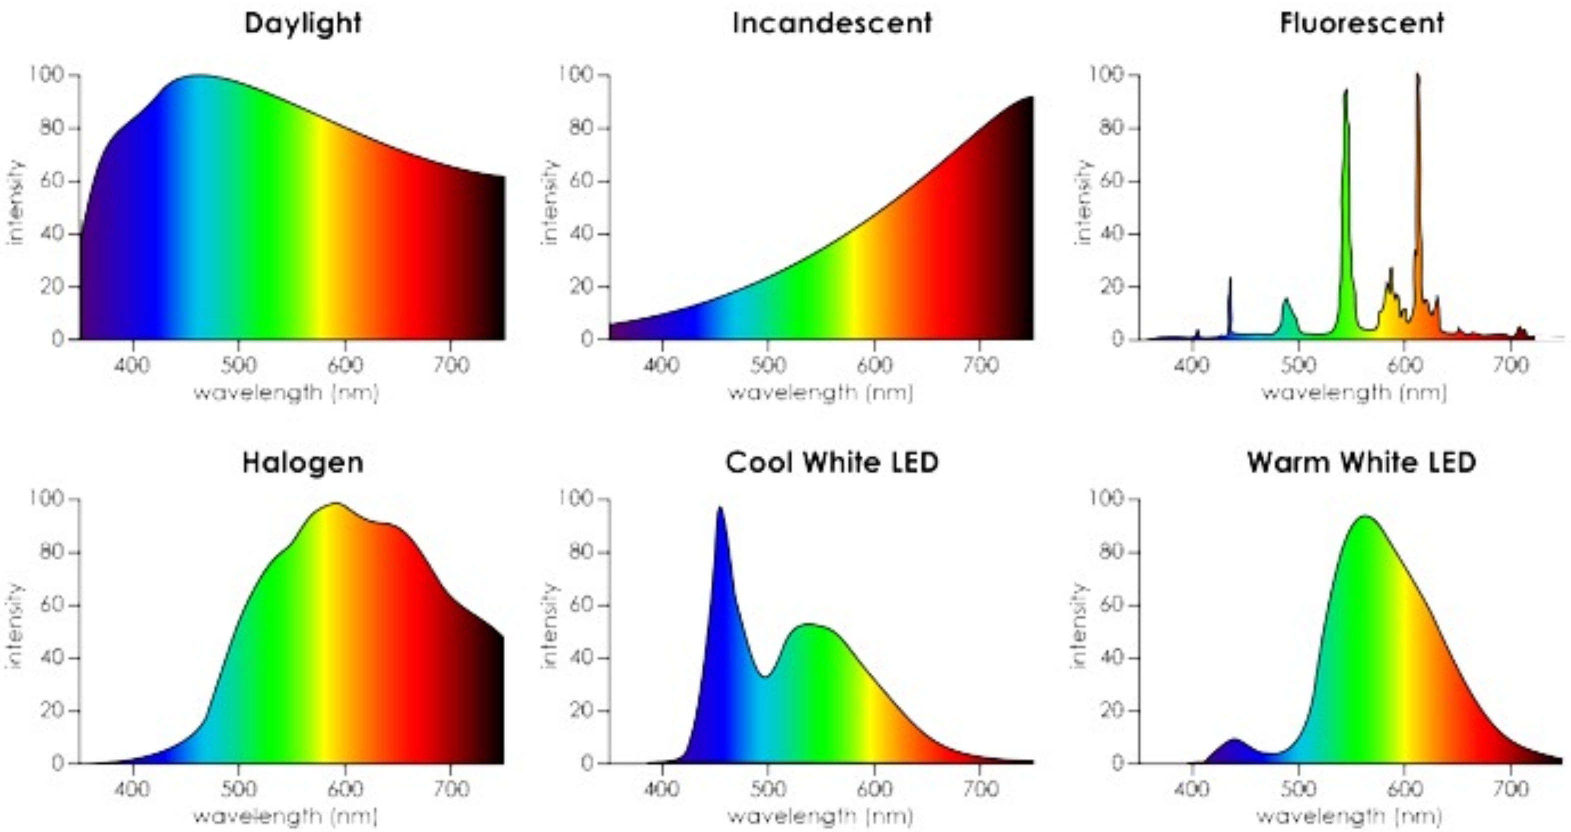

Color temperature

- Illuminants are often described by their "color temperature"
- This mapping is based on theoretical "blackbody radiators" that produce SPDs for a given temperature -- expressed in Kelvin (K)
- We map light sources (both real and synthetic) to their closest color temperature (esp in Photography/Video production)

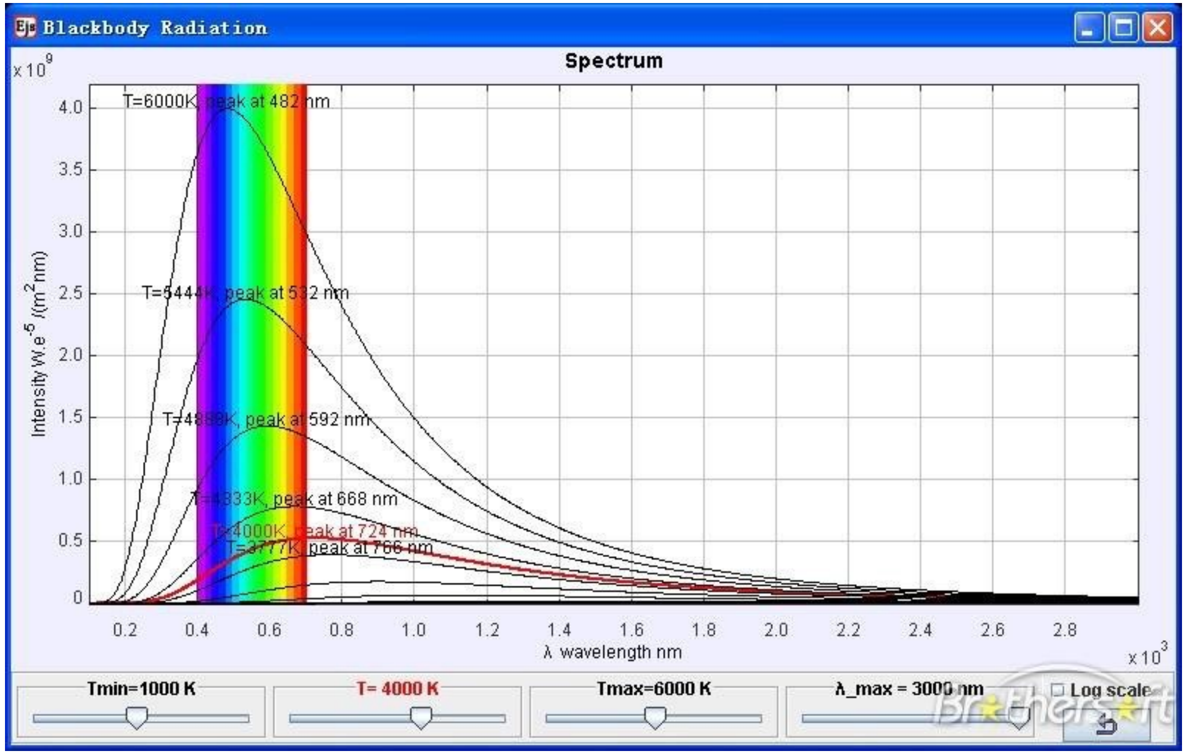

$$B_\lambda(\lambda, T) = \frac{2hc^2}{\lambda^5} \frac{1}{e^{\frac{hc}{\lambda k_B T}} - 1}$$

Plank's law
Spectral density of electromagnetic radiation emitted by a blackbody radiator at a given temperature T.

Color temperature

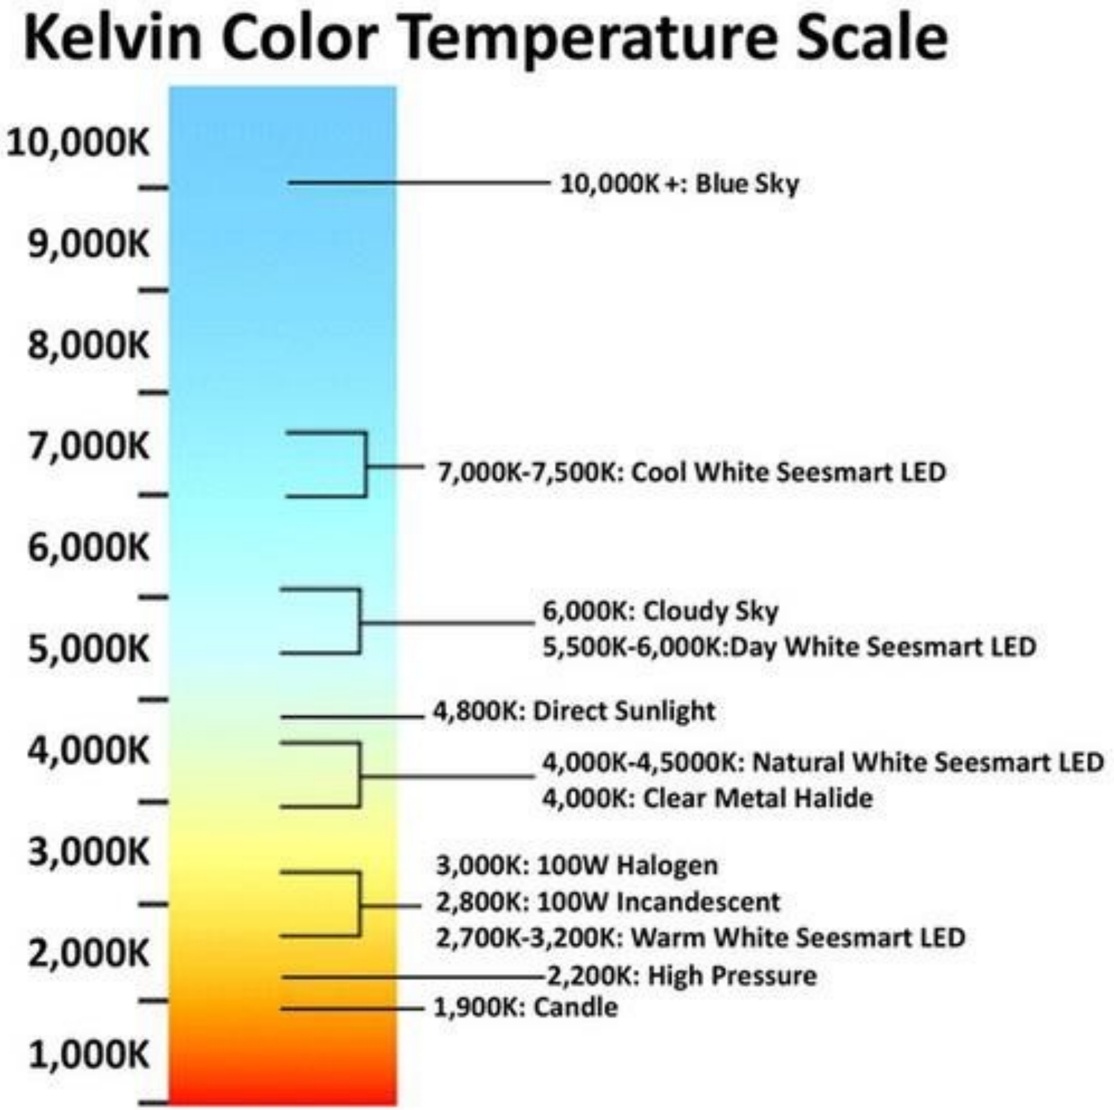

Typical description of color temperature used in photography & lighting sources.

From B&H Photo

Lighting industry uses color temperature

Usage of color temperature in these ads relate to the perceived color of the bulb's light.The heat output of a typical LED bulb is between 60C-100C (~333-373K).

Light-material interaction

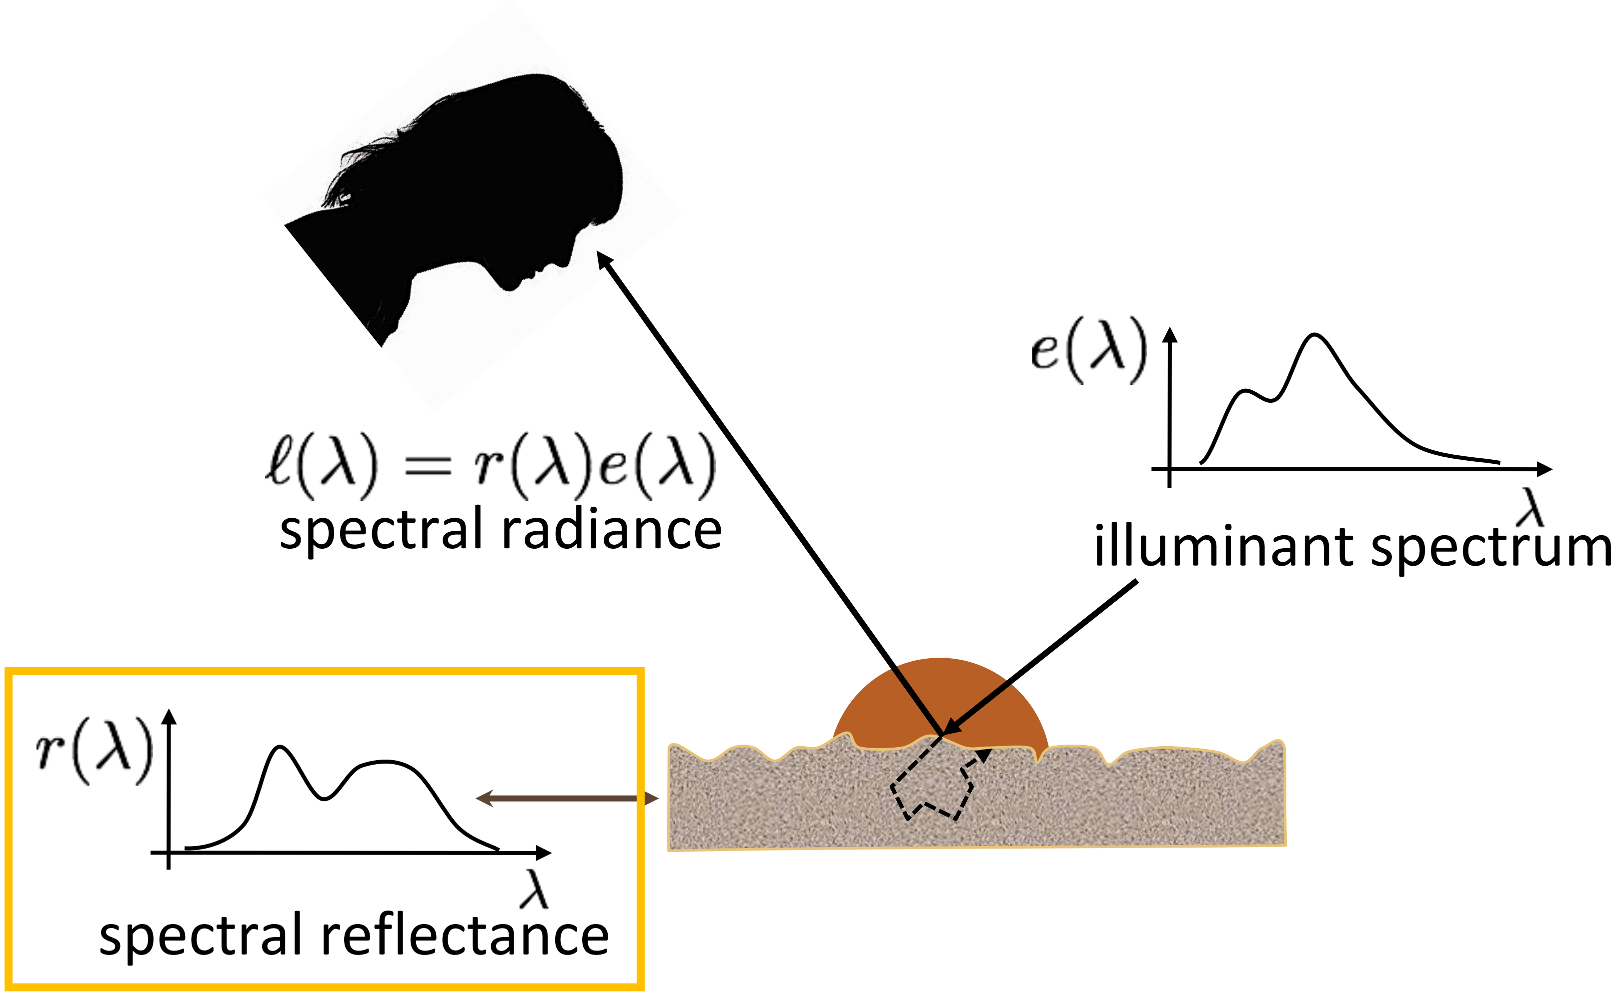

Spectral reflectance

- Most materials absorb and reflect light differently at different wavelengths.
- We can describe this as a ratio of reflected vs incident light over different wavelengths.

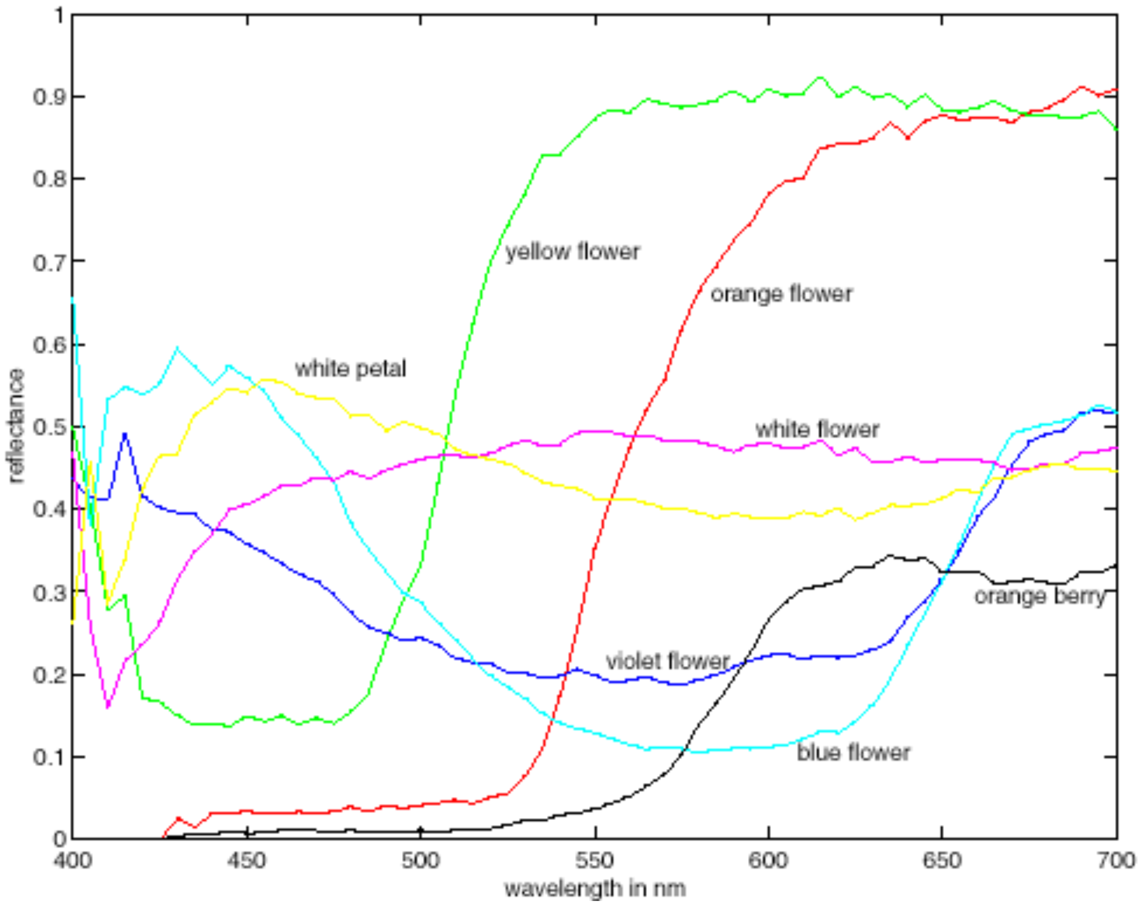

Light-material interaction

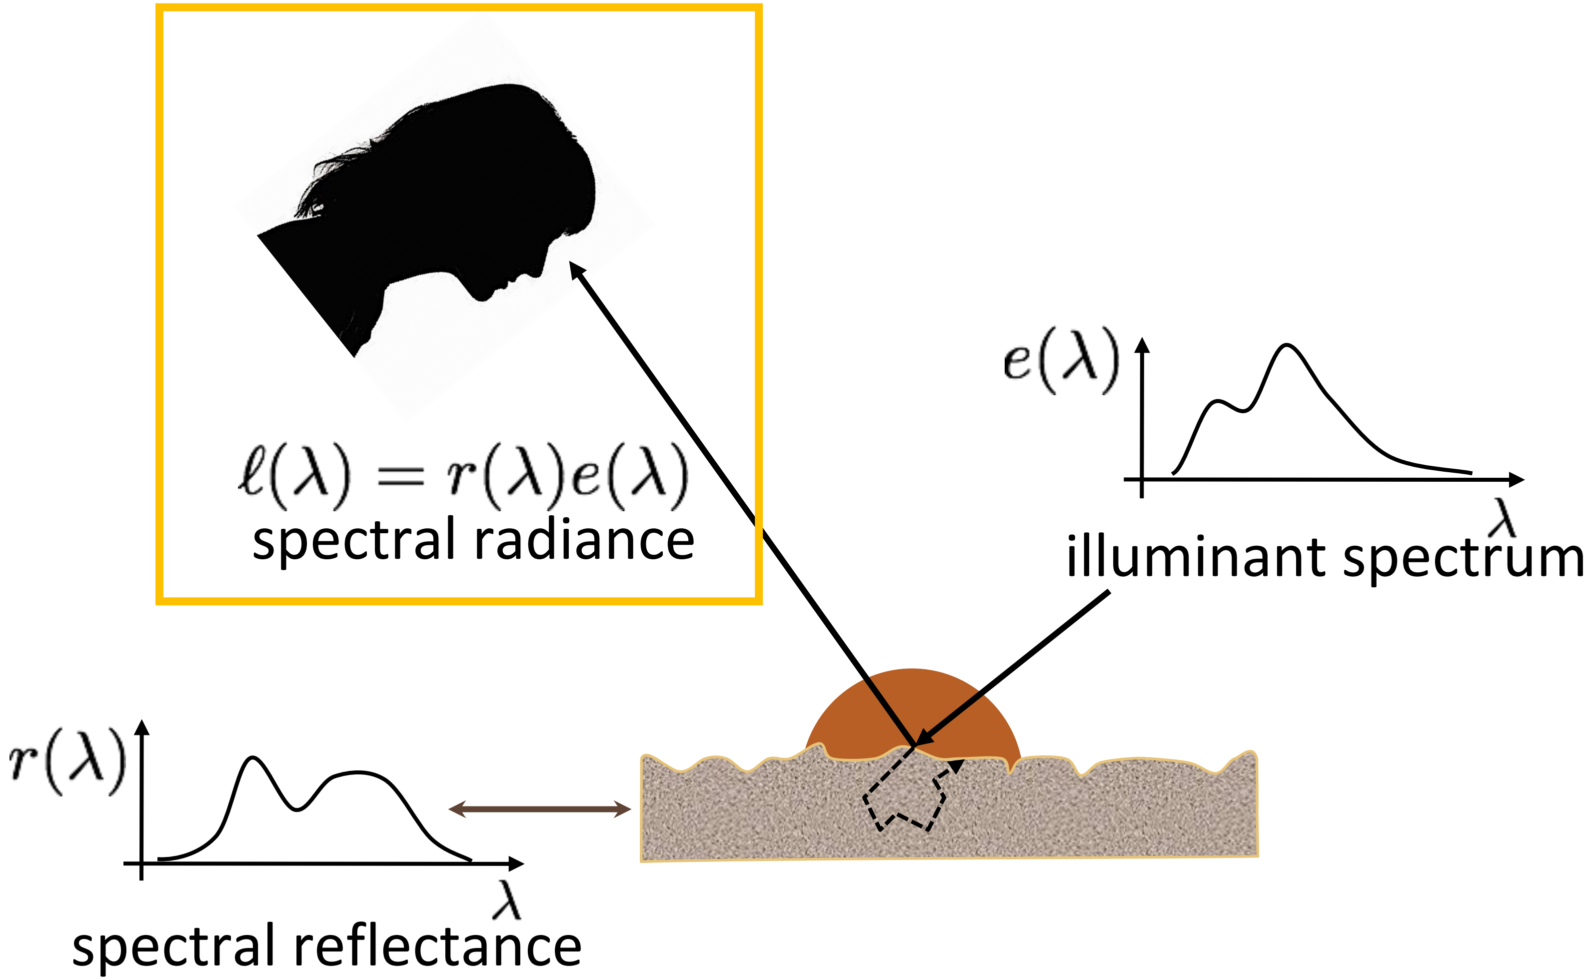

Physiology of Color Vision

Three kinds of cones:

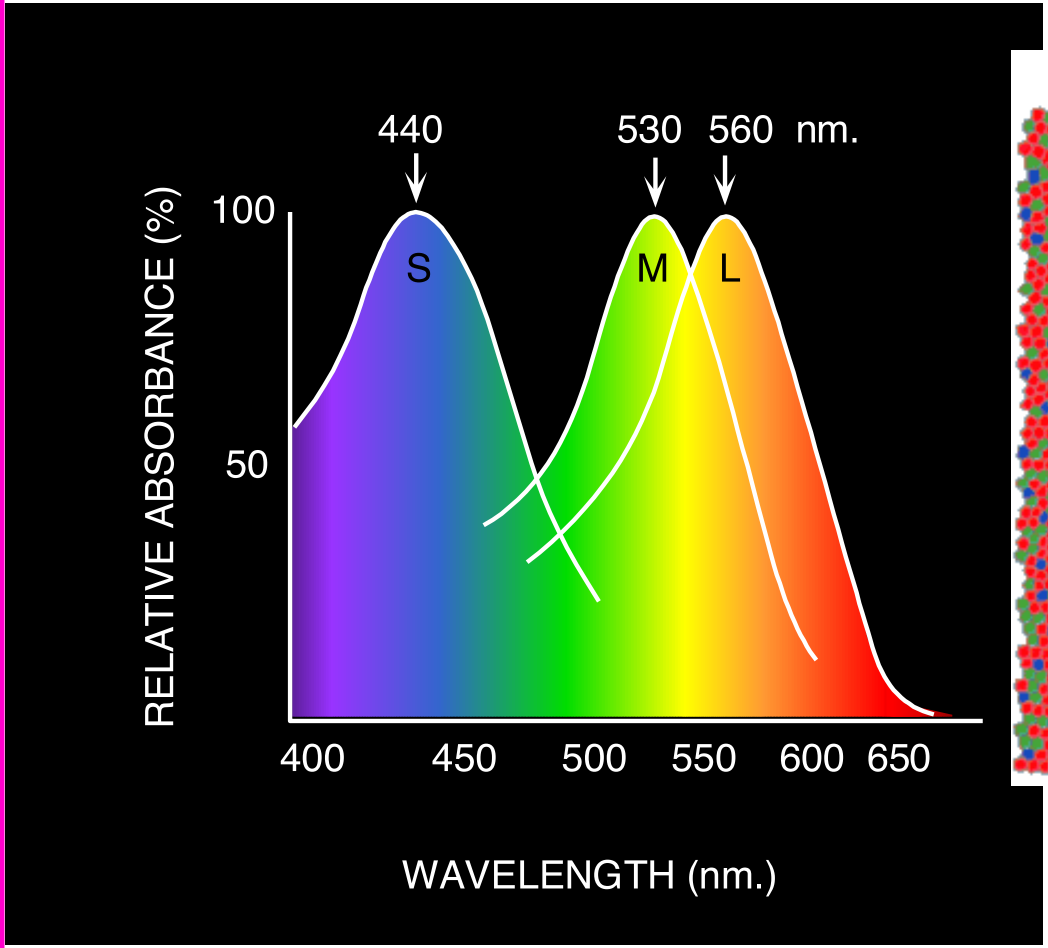

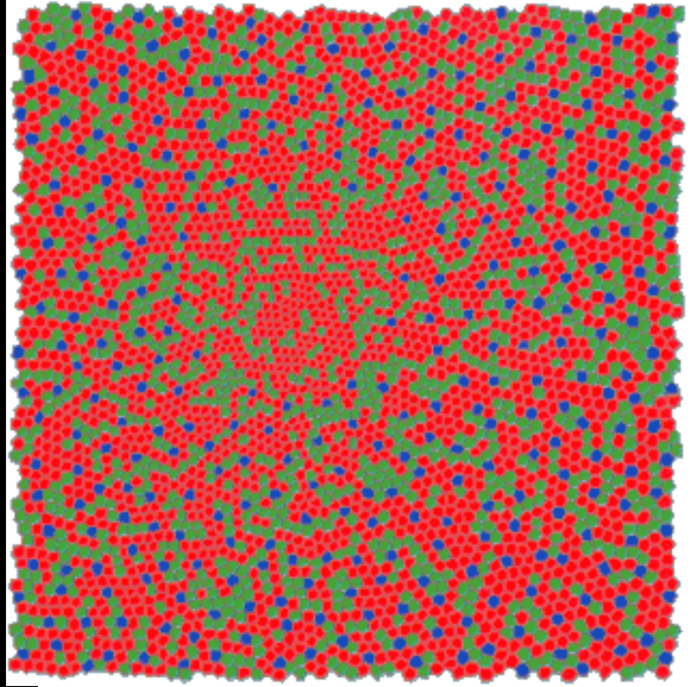

Human color vision

retinal color

$$\mathbf{c}(\ell(\lambda)) = (c_s, c_m, c_l)$$

$$c_s = \int k_s(\lambda)\ell(\lambda)d\lambda$$

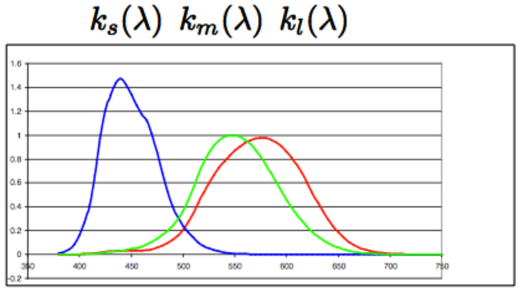

LMS senstivity functions

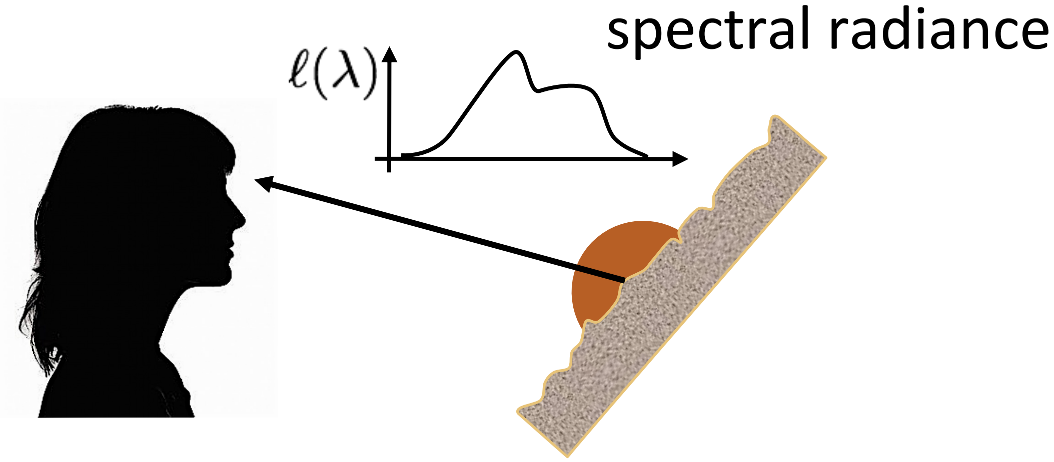

perceived color

object color

color names

Retinal vs perceived color

Retinal vs perceived color.

This is known as Color Constancy: more later in this lecture

Perceived vs measured brightness by human eye

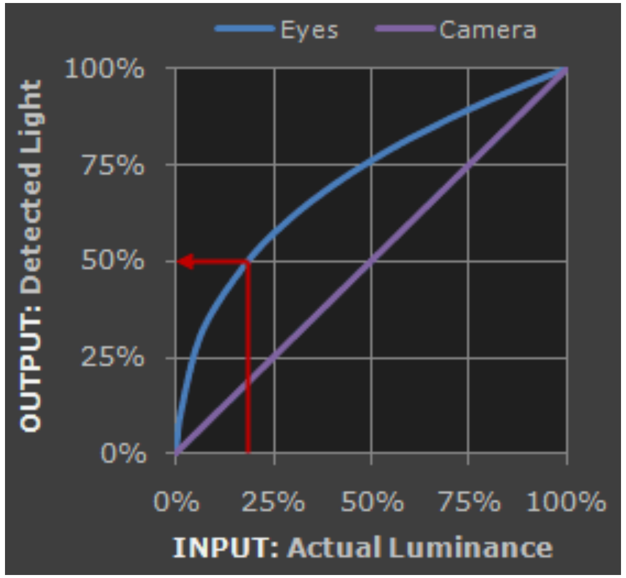

Human-eye *response* (measured brightness) is linear.

However, human-eye *perception* (perceived brightness) is *non-linear*:

- More sensitive to dark tones.
- Approximately a Gamma function.

Not everyone is trichromat

- Types of color blindness:
  - Deuteranopia: missing M cones
  - Protanopia: missing L cones
  - Tritanopia: missing S cones
- "M" and "L" on the X-chromosome
  - Why men are more likely to be color blind
  - "L" has high variation, so some women are tetrachromatic
- Some animals have
  - 1 (night animals)
  - 2 (e.g., dogs)
  - 4 (fish, birds)
  - 5 (pigeons, some reptiles/amphibians)
  - 12 (mantis shrimp)

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

Trichromacy

Cone primaries = L, M, S
Image primaries = R, G, B

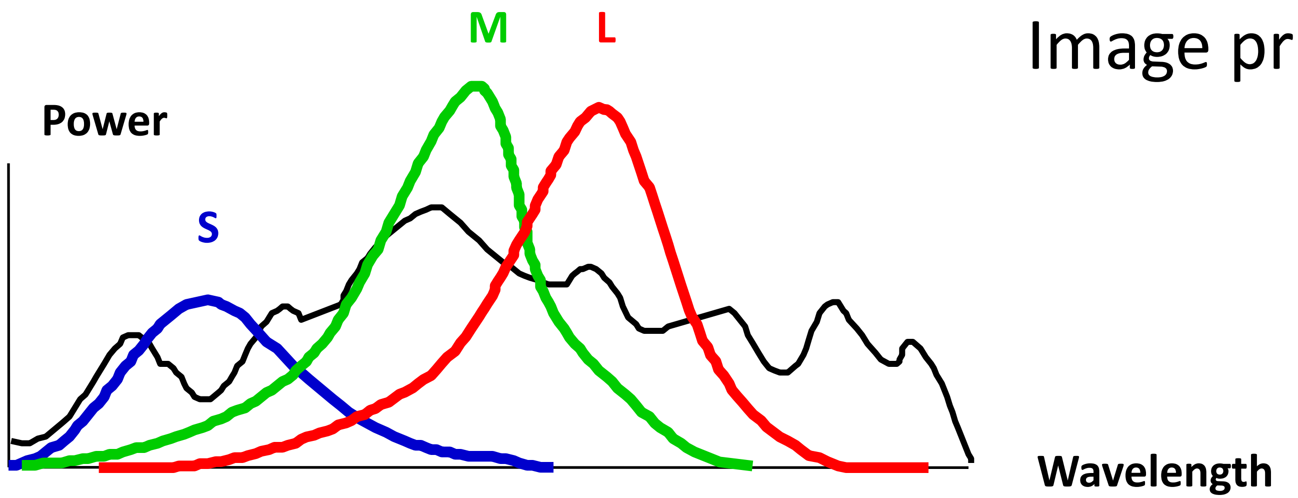

Rods and cones act as *filters* on the spectrum

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

- How can we represent an entire spectrum with 3 numbers?
- We can't! Most of the information is lost
  - As a result, two different spectra may appear indistinguishable
    - » such spectra are known as **metamers**

Slide by Steve Seitz

Today's class

- Eye & Human Vision
- Pinhole Camera
- Perception of Color
- Quantifying Color
- Commonly Used Color Spaces
- Color Constancy

Radiometry vs. photometry

- **Radiometry**
  - Quantitative measurements of radiant energy
  - Often shown as spectral power distributions (SPD)
  - Measures either light coming from a source (radiance) or light falling on a surface (irradiance)
- **Photometry/ colorimetry**
  - Quantitative measurement of **perceived** radiant energy based on human's sensitivity to light
  - Perceived in terms of "brightness" (photometry) and color (colorimetry)

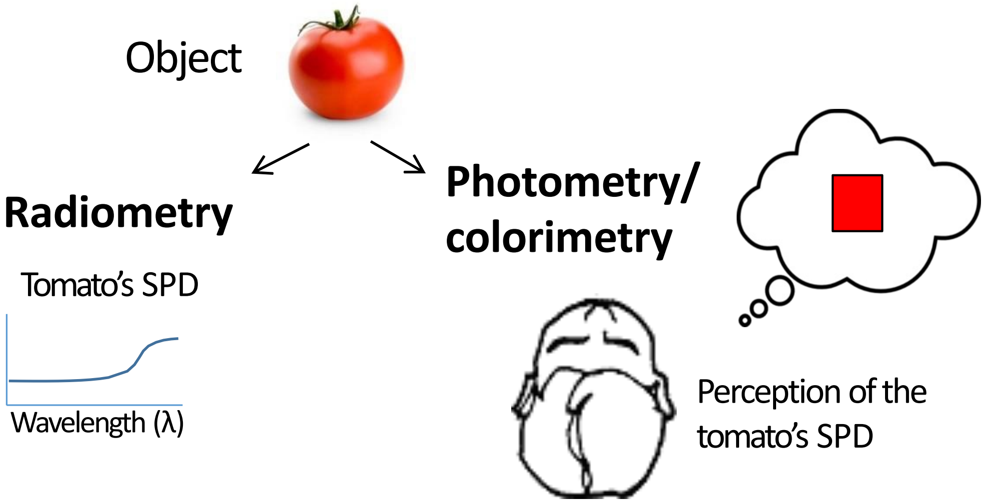

Quantifying color

- We still need a way to quantify color & brightness
- SPDs go through a "black box" (human visual system) and are perceived as color
- The only way to quantify the "black box" is to perform a human study

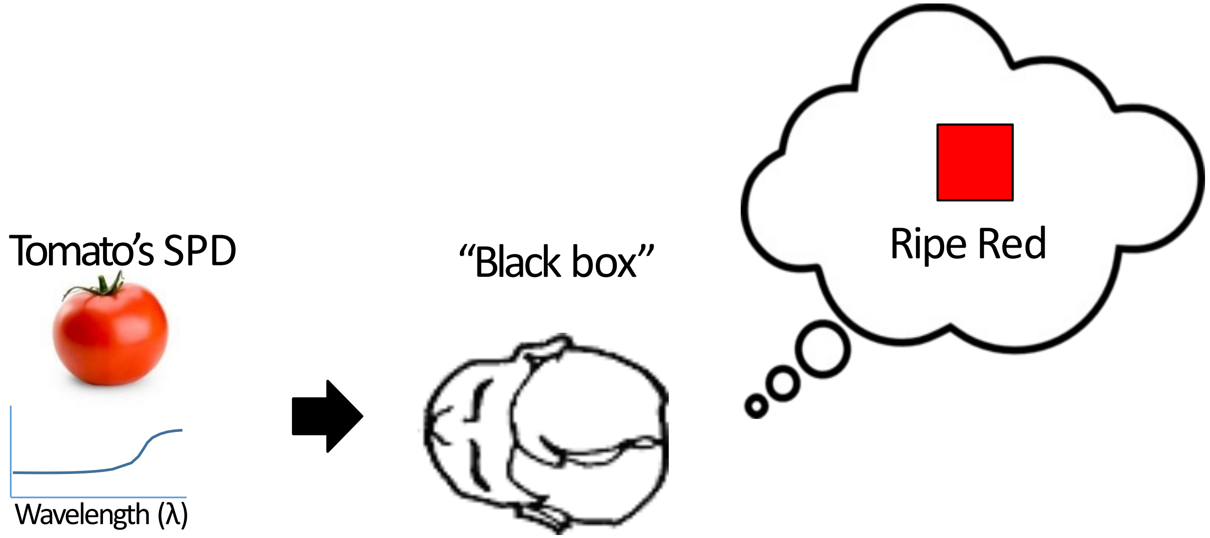

CIE RGB color matching

Commission Internationale de l'Elcairage (CIE)

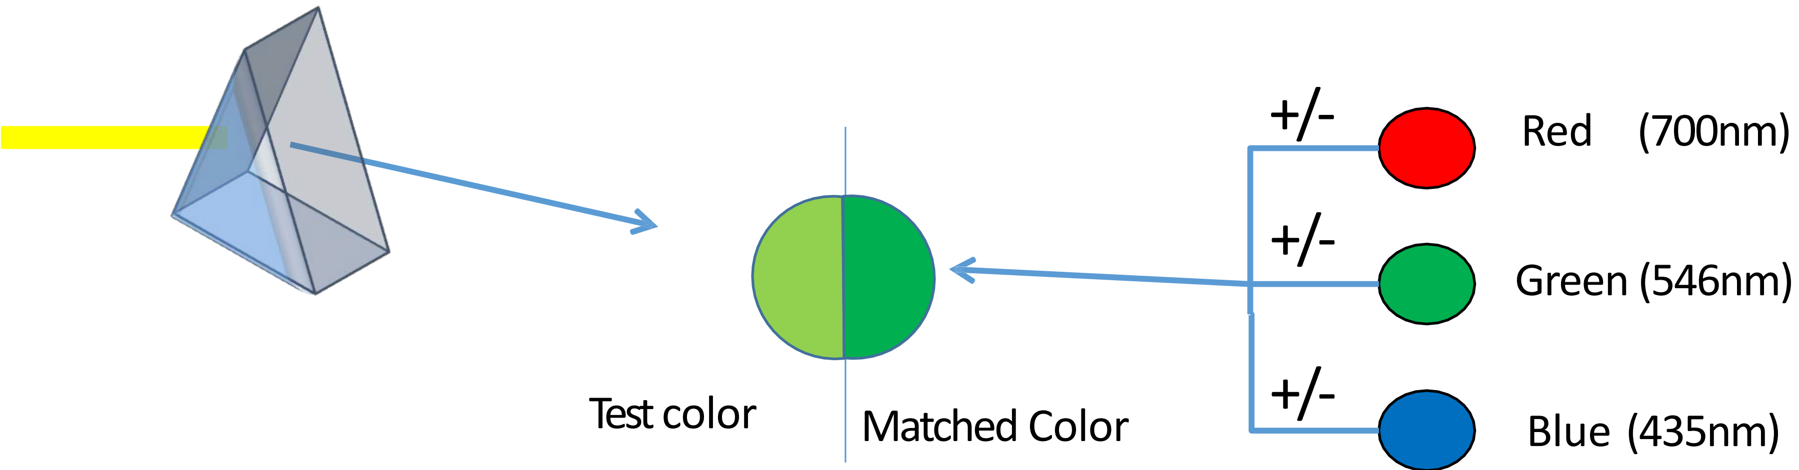

Human subjects matched test colors by add or subtracting three primaries.

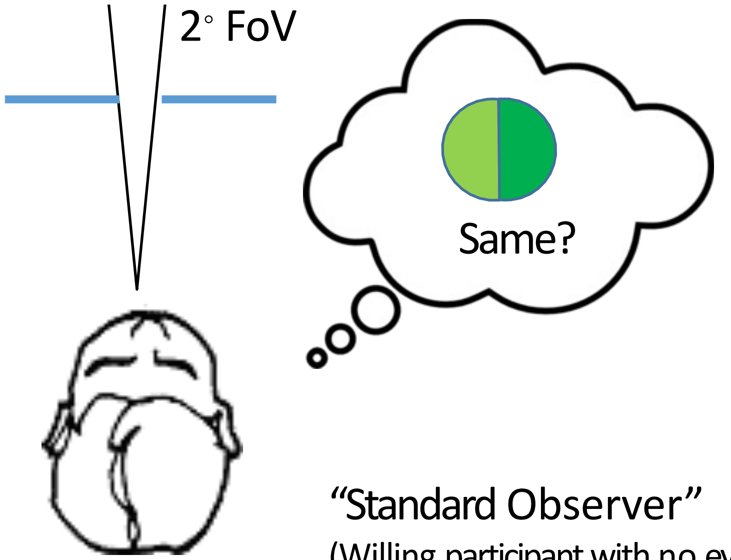

Field of view was 2-degrees (where color cones are most concentrated)

“Standard Observer”
(Willing participant with no eye disease)

Experiments carried out by
W. David Wright (Imperial College) and John Guild (National Physical Laboratory, London) – Late 1920s

CIE RGB color matching

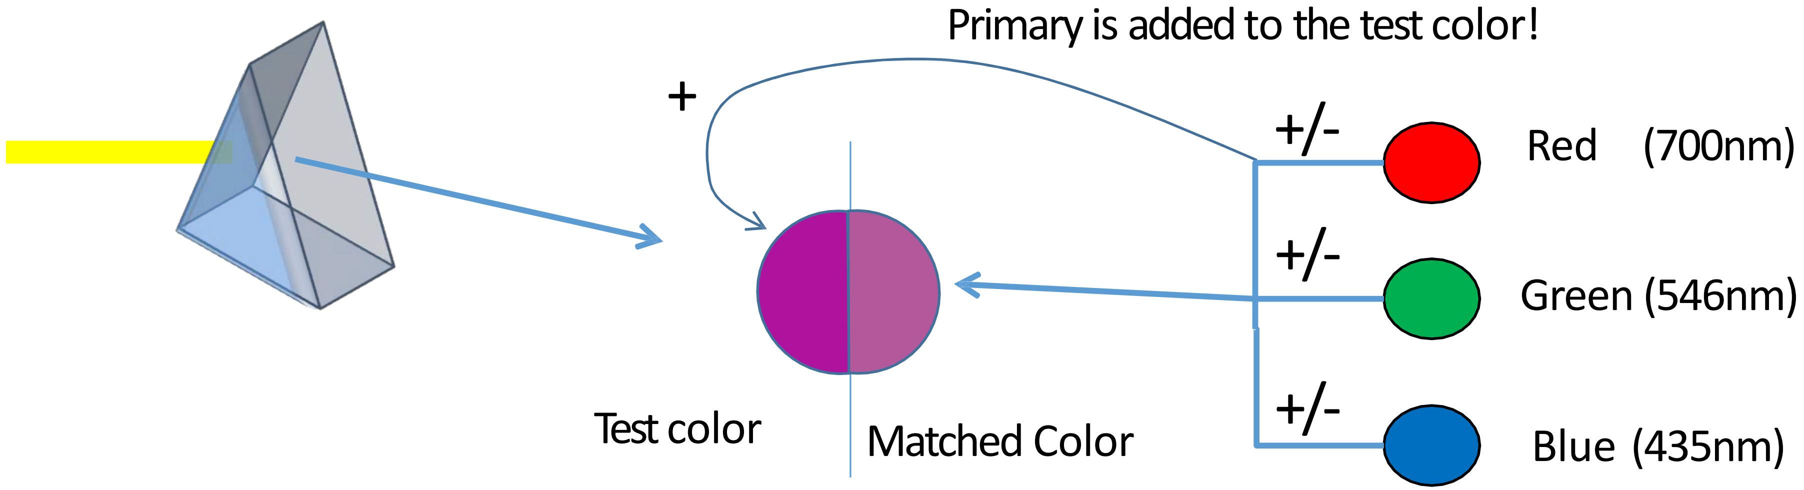

For some test colors, no mix of the primaries could give a match! For these cases, the subjects were ask to add primaries to the test color to make the match.

This was treated as a negative value of the primary added to the test color.

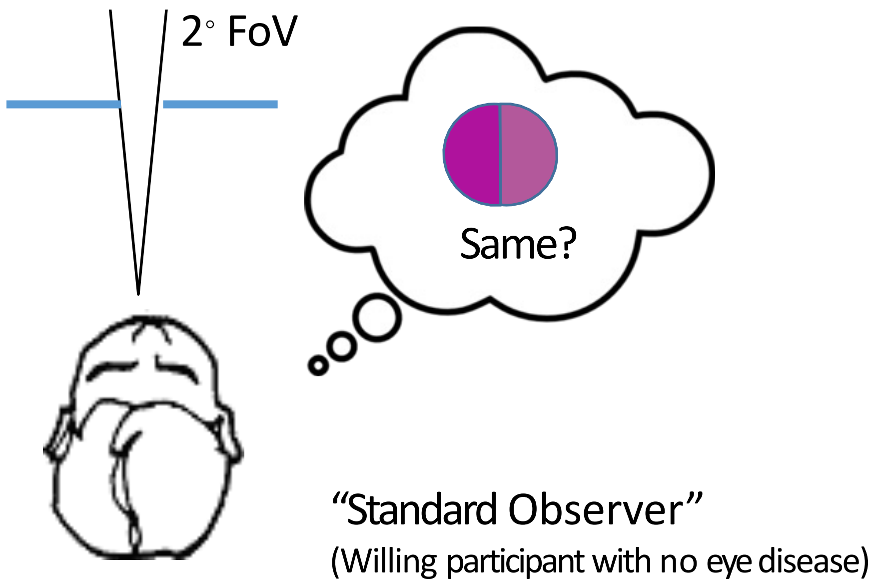

CIE RGB results

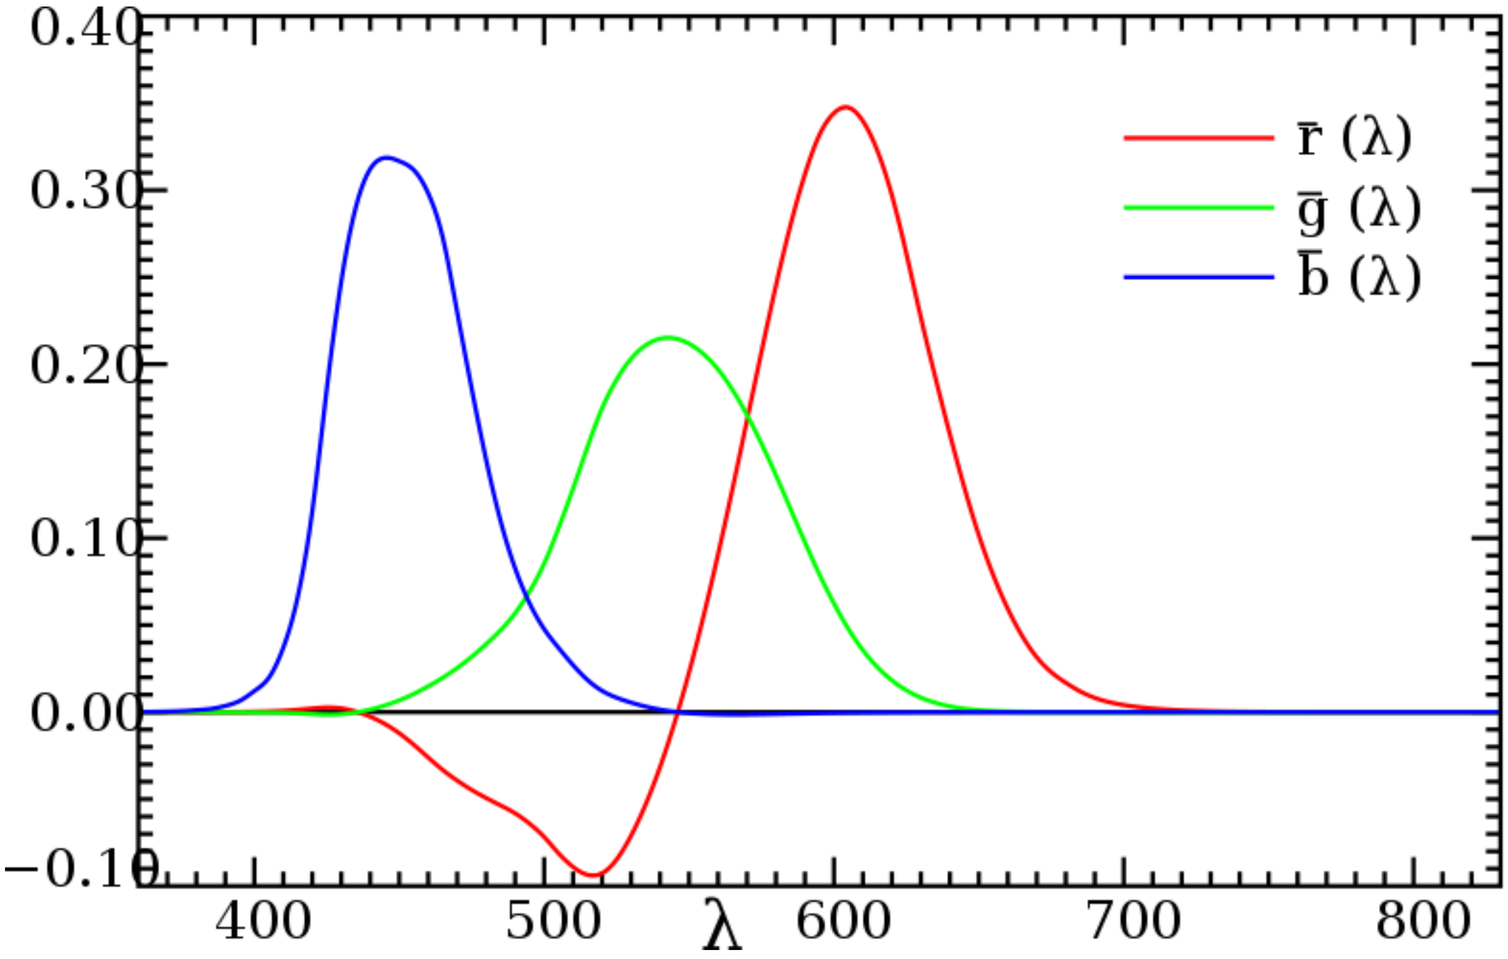

CIE RGB 2-degree Standard Observer
(based on Wright/Guild's data)

Plots are of the mixing coefficients of each primary needed to produce the corresponding monochromatic light at that wavelength.

Note that these functions have been scaled such that area of each curve is equal.

CIE RGB results

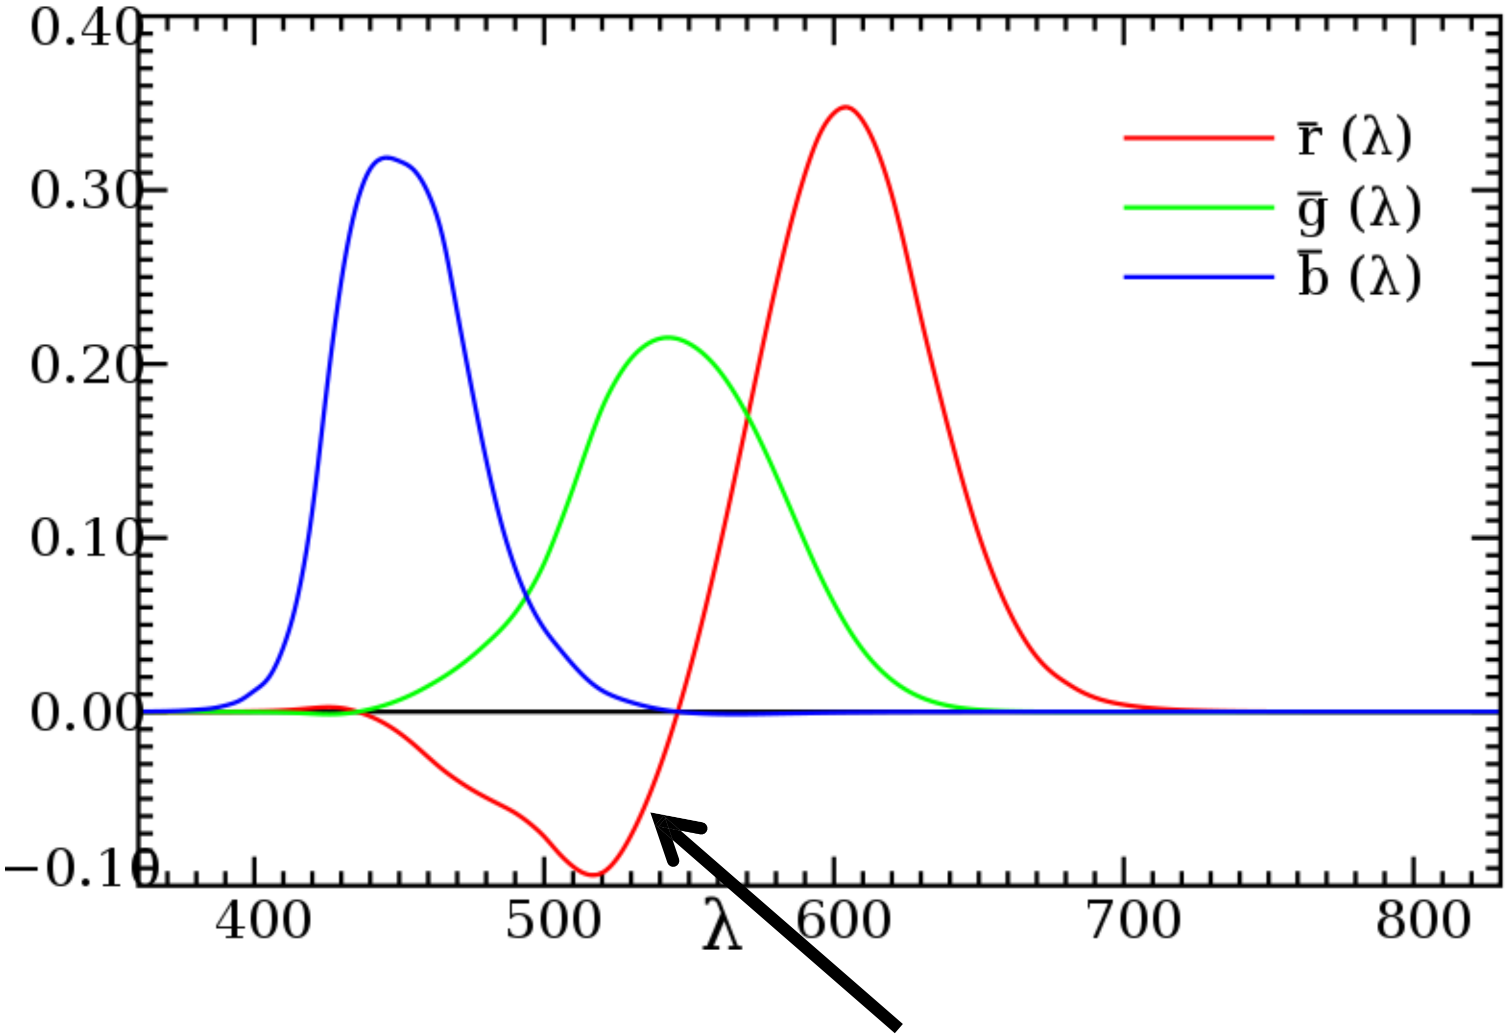

RGB filter response is device specific. Different cameras will have slightly different RGB filter.

Then CIE RGB is not a standard across different cameras!

Negative values -- the three primaries used did not span the full range of perceptual colors.

CIE 1931 XYZ

- In 1931, the CIE met and approved defining a new canonical basis, termed XYZ that would be derived from Wright-Guild's CIE RGB data
- Properties desired in this conversion:
  - White point defined at X=1/3,Y=1/3, Z=1/3
  - Y would be the luminosity function ($V(\lambda)$)
  - Quite a bit of freedom in selecting these XYZ basis
  - In the end, the adopted transform was:

$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = \begin{bmatrix} 0.4887180 & 0.3106803 & 0.2006017 \\ 0.1762044 & 0.8129847 & 0.0108109 \\ 0.0000000 & 0.0102048 & 0.9897952 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix}$$

← CIE 1931 RGB

Nice article see: Fairman et al "How the CIE 1931 Color-Matching Functions Were Derived from Wright–Guild Data", Color Research & Application, 1997

CIE xy (chromaticity)

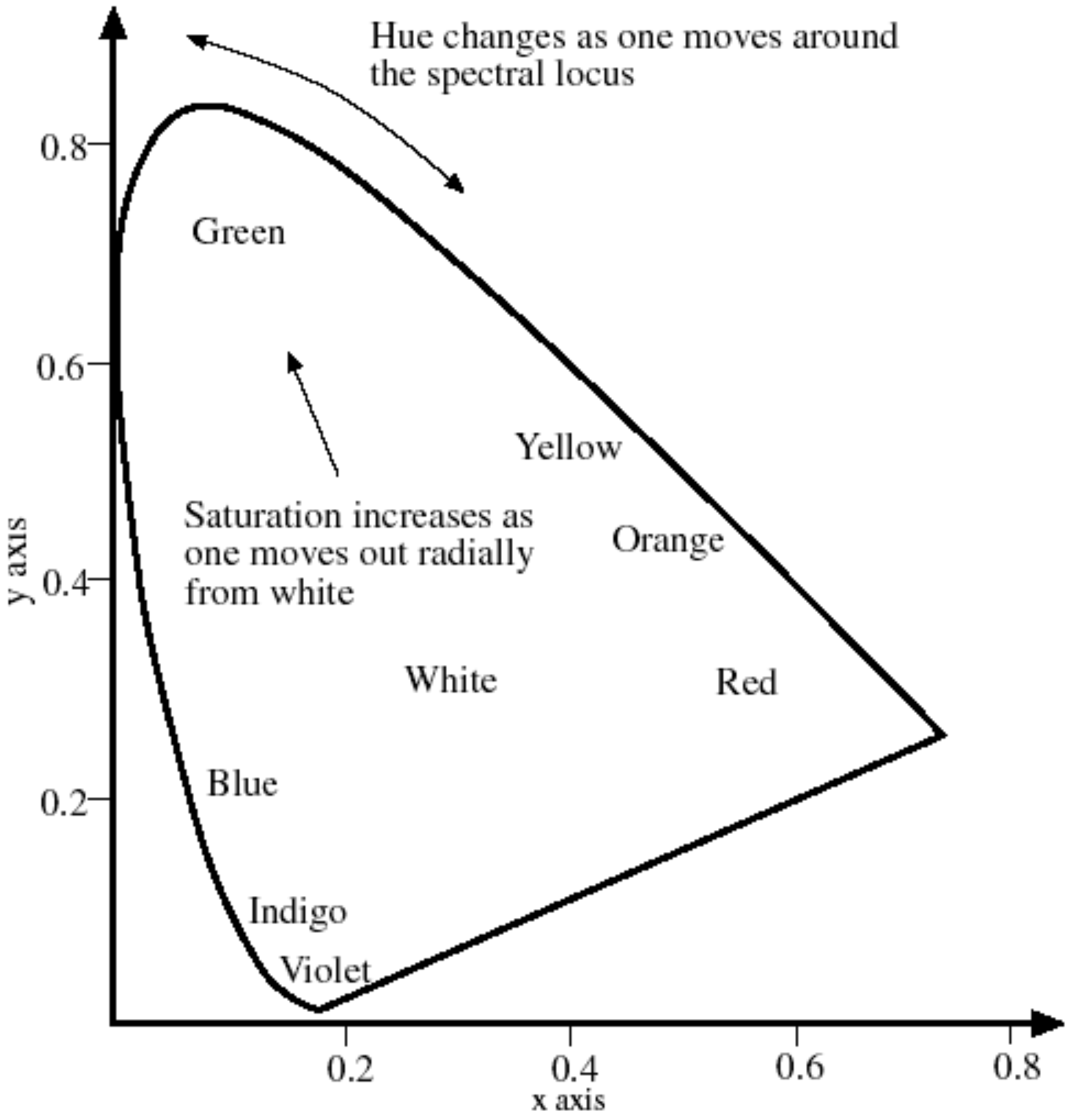

$$x = \frac{X}{X+Y+Z}$$

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

$$(X, Y, Z) \longleftrightarrow (\underline{x, y}, Y)$$

chromaticity

luminance/brightness

Perspective projection of 3D retinal color space to two dimensions.

CIE xy chromaticity diagram

This gives us the familiar horseshoe shape of visible colors as a 2D plot. Note the axis are x & y.

Point "E" represents where X=Y=Z have equal energy (X=0.33,Y=0.33, Z=0.33)

CIE XYZ "white point"

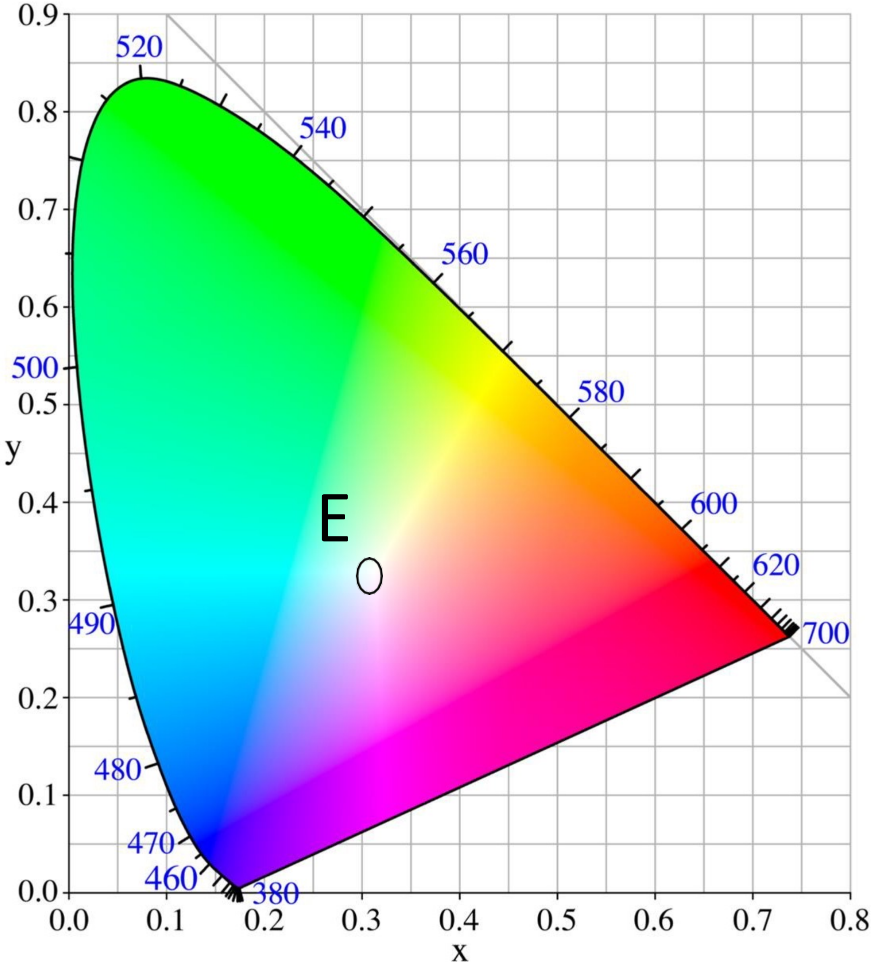

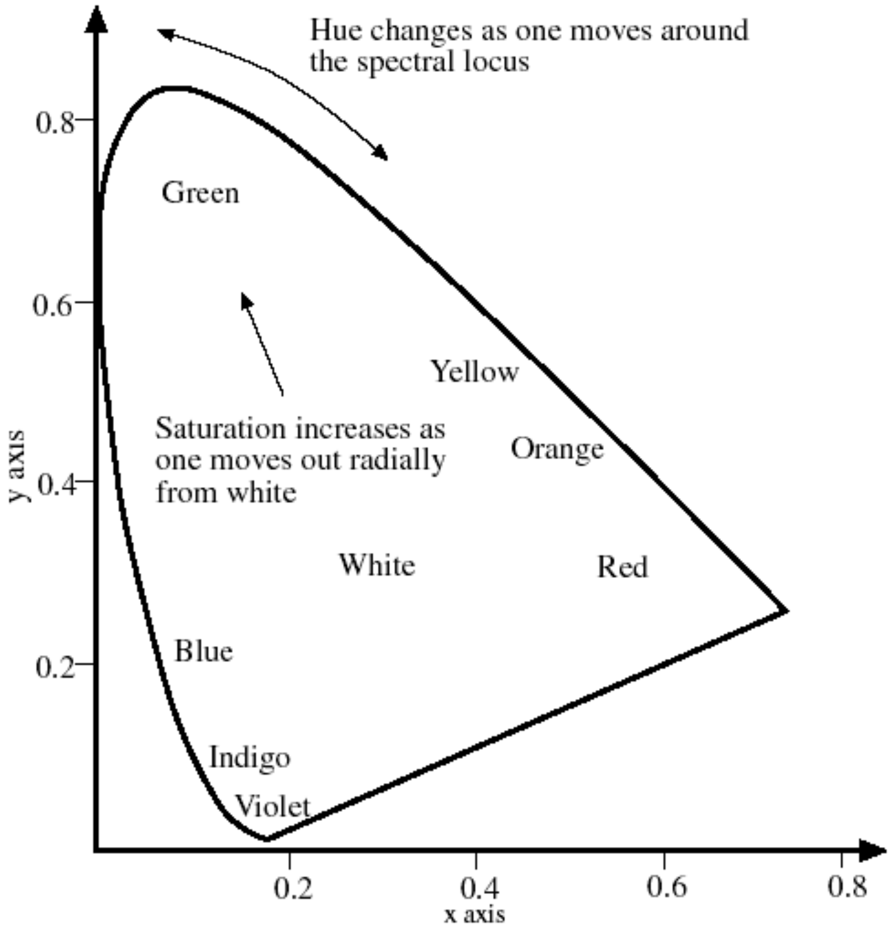

In the 1930s, CIE had a bad habit of over using the variables X,Y. Note that x, y are chromaticity coordinates, x, y (with the bar above) are the matching functions, and X,Y are the imaginary SPDs of CIE XYZ.

Fast forward 80+ years

- CIE 1931 XYZ, CIE 1931 xyY (2-degree standard observer) color spaces have stood the test of time
- Many other studies have followed (most notably - CIE 1965 XYZ 10- degree standard observer), …
- But in the literature (and in this tutorial) you'll find CIE 1931 XYZ color space remains the preferred standard

What is perhaps most amazing?

- 80+ years of CIE XYZ and it is all based on the experiments by the "standard observers"

- How many standard observers were used? 100, 500, 1000?

A Standard Observer

CIE XYZ is based on 17 standard observers

10 by Wright, 7 by Guild

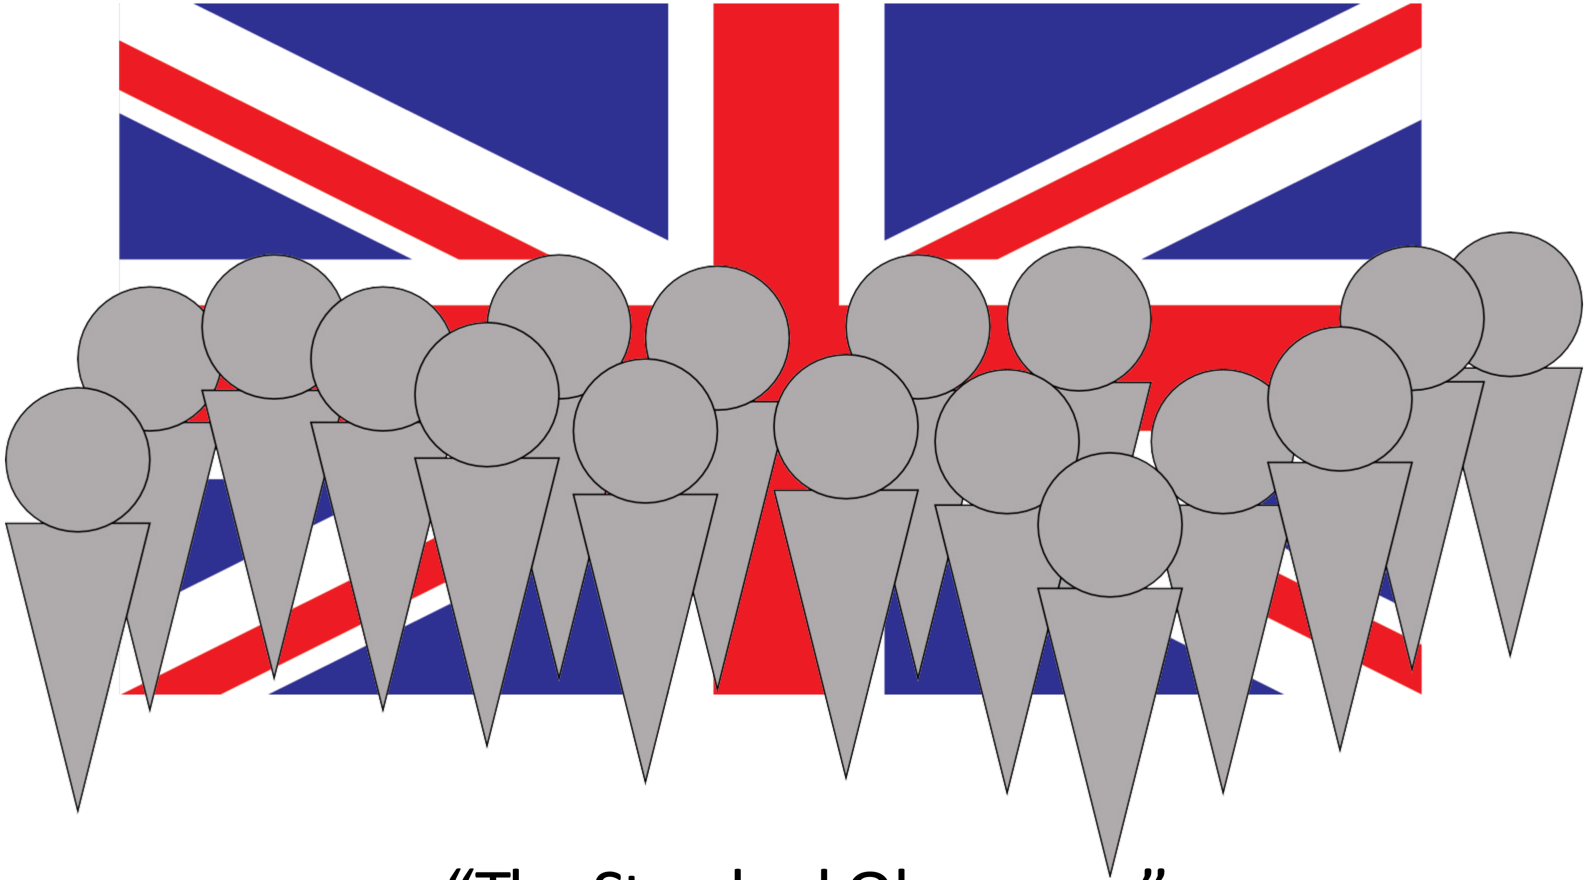

"The Standard Observers"

Today’s class

- Eye & Human Vision
- Pinhole Camera
- Perception of Color
- Quantifying Color
- Commonly Used Color Spaces
- Color Constancy

CIE XYZ and RGB

- While CIE XYZ is a canonical color space, images/devices rarely work directly with XYZ
- XYZ are not real primaries
- RGB primaries dominate the industry
- We are all familiar with the RGB color cube

But by now, you should realize that Red, Green, Blue have no quantitative meaning. We need to know their corresponding SPDs or CIE XYZ values.

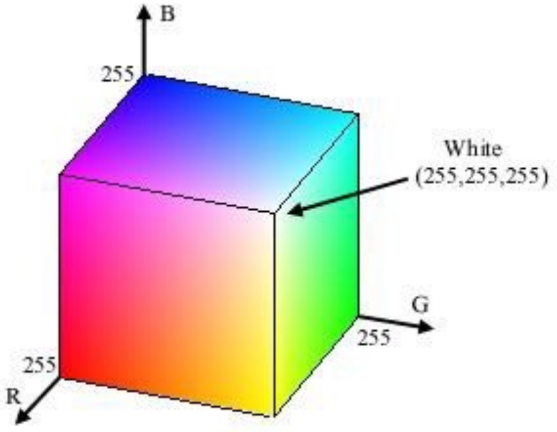

RGB color space is device specific, XYZ is device independent.

[X,Y,Z] = M * [R,G,B]

M is 3x3 transformation matrix provided by the manufacturor of the camera or can be calibrated.

Color spaces: RGB

Default color space

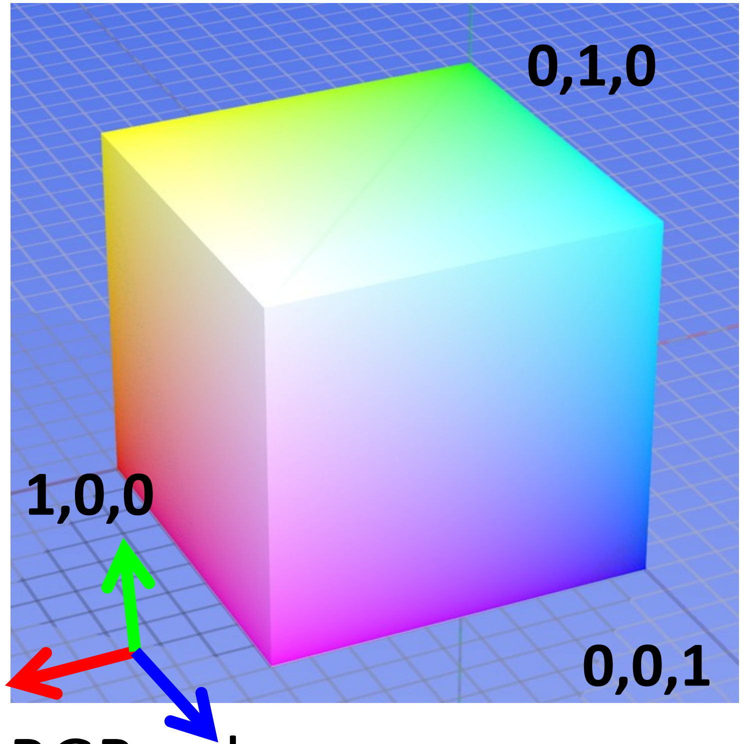

RGB cube

- Easy for devices
- But not perceptual

R
(G=0,B=0)

G
(R=0,B=0)

B
(R=0,G=0)

Image from: http://en.wikipedia.org/wiki/File:RGB_color_solid_cube.png

Color gamuts

sRGB color gamut

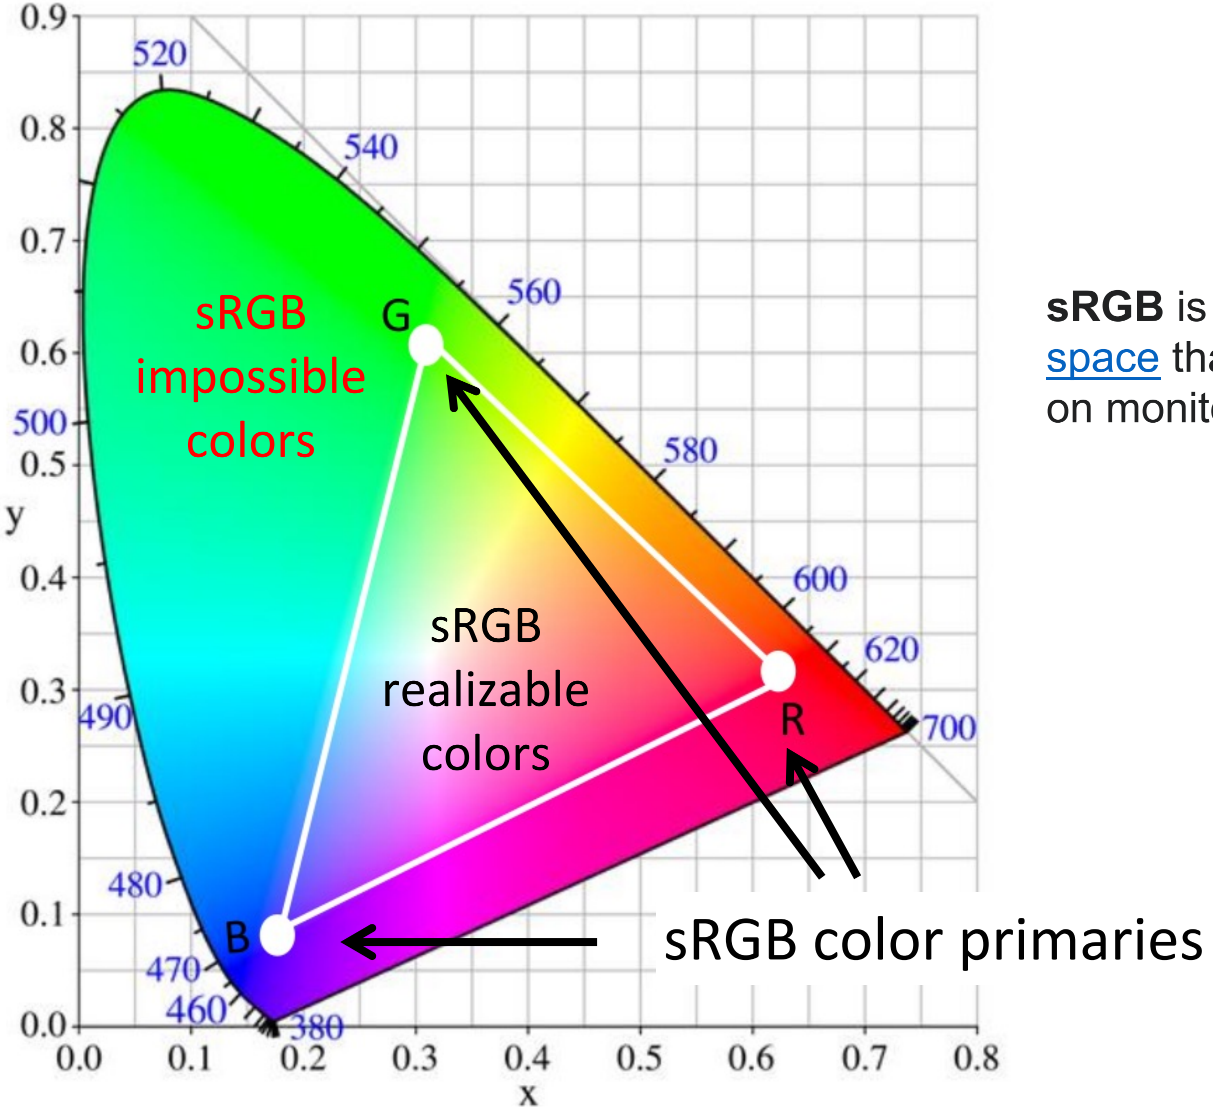

The RGB values span a subspace of CIE-XYZ to define the devices gamut.

sRGB is a standard RGB (red, green, blue) color space that HP and Microsoft created cooperatively in 1996 to use on monitors, printers, and the World Wide Web.

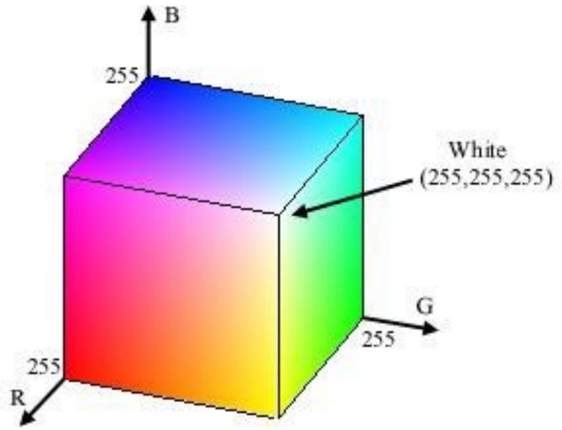

Color gamuts

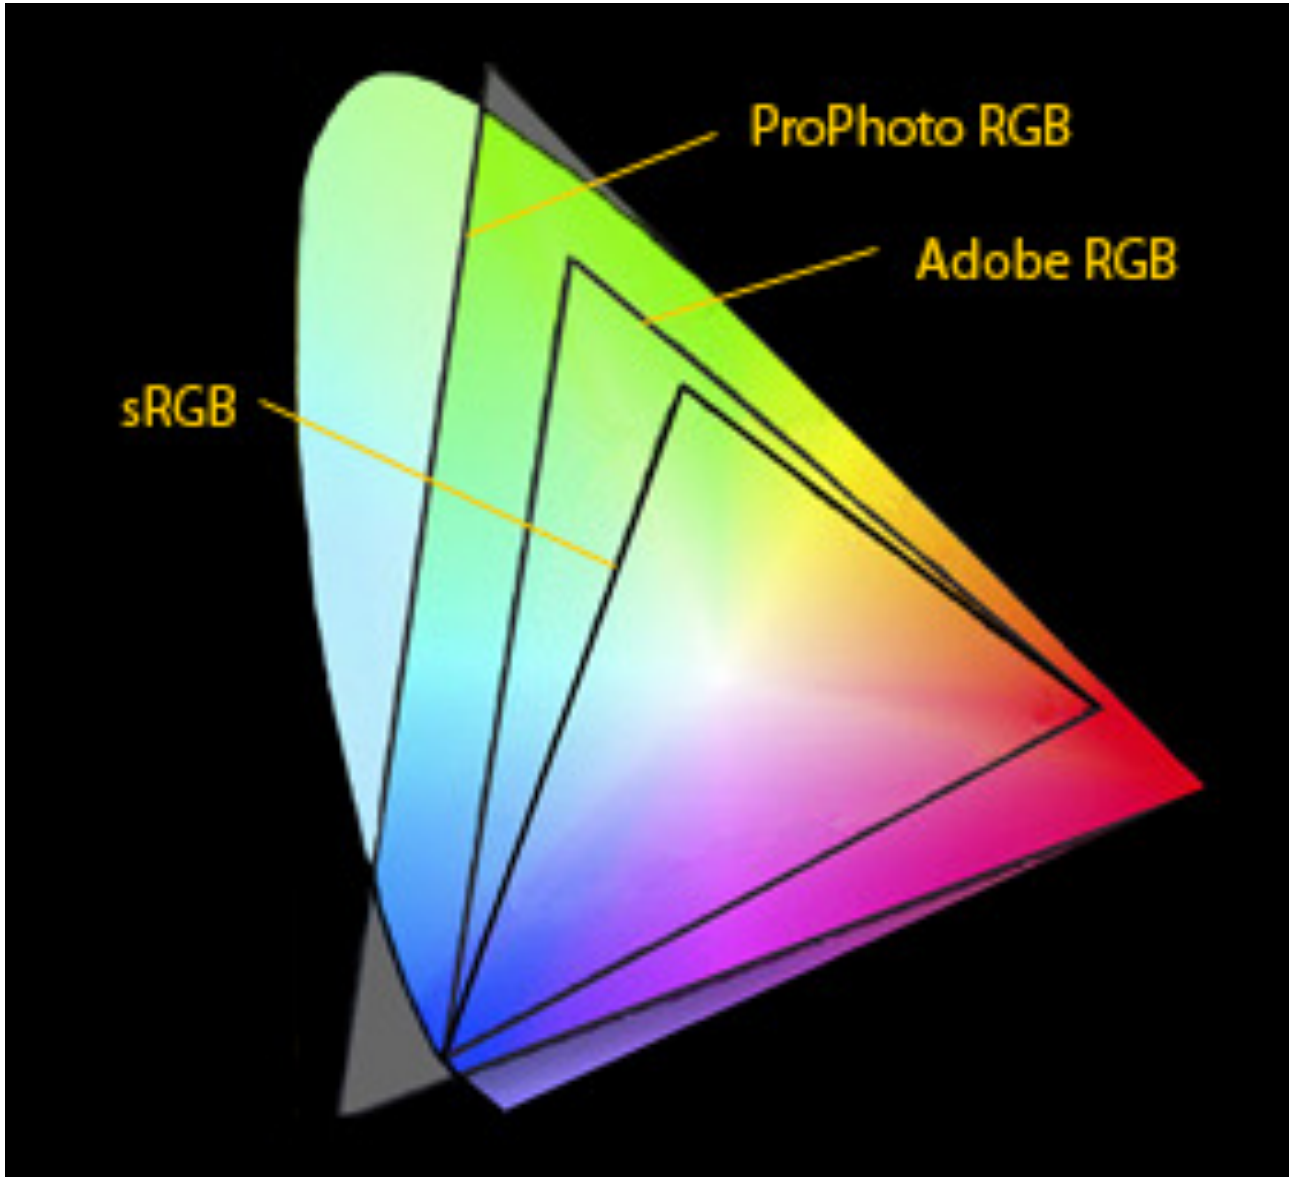

Gamuts of various common industrial RGB spaces

The problem with RGBs visualized in chromaticity space

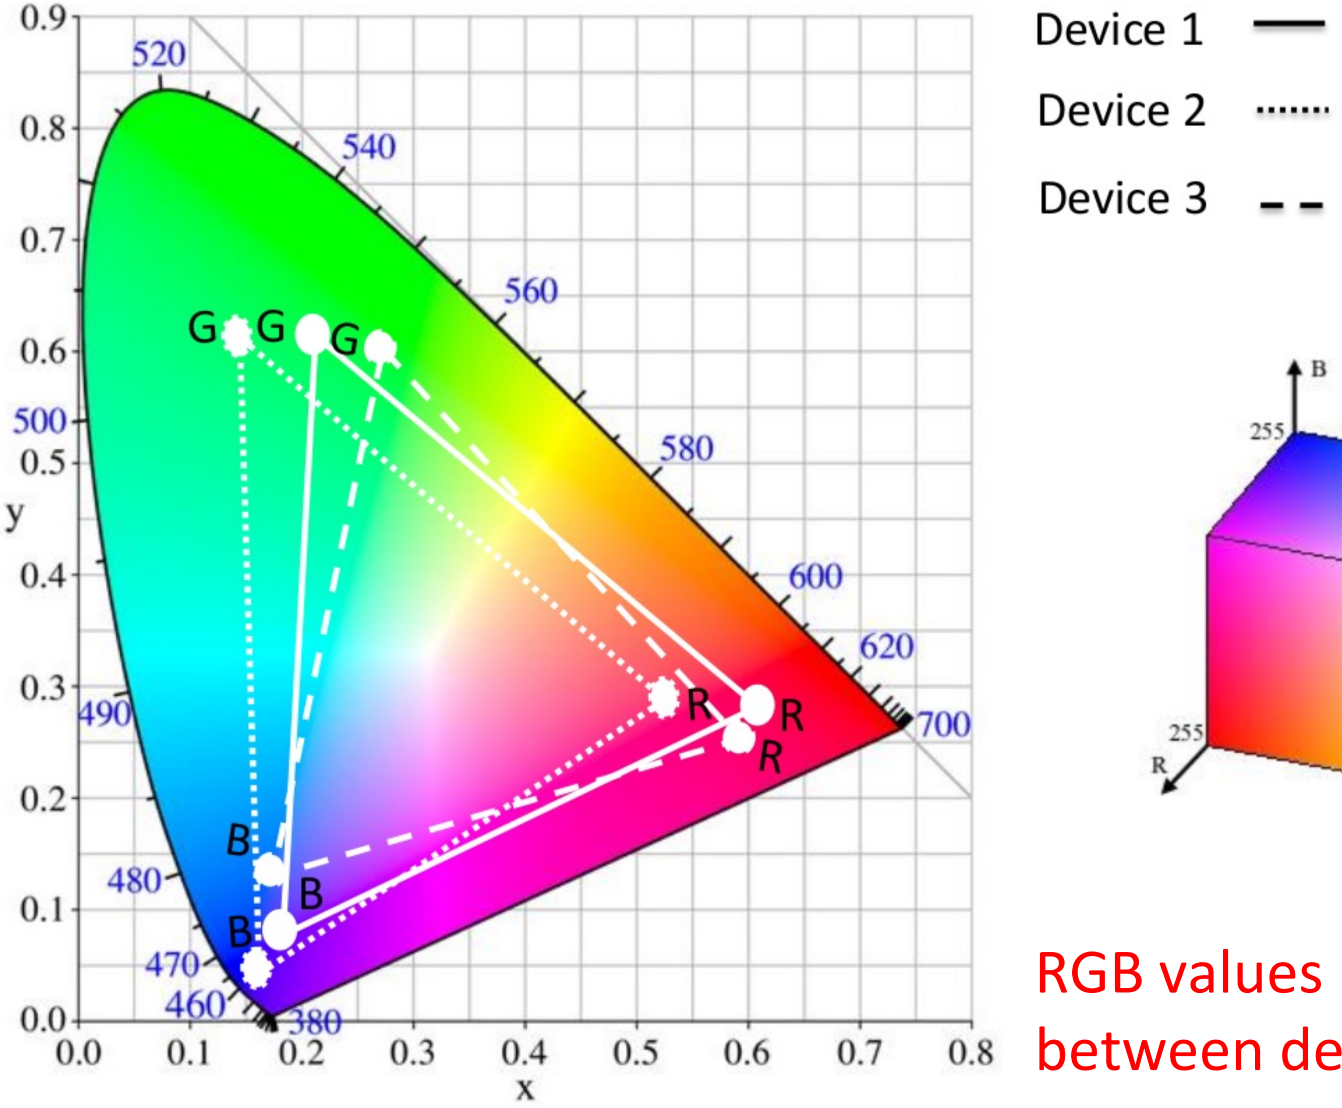

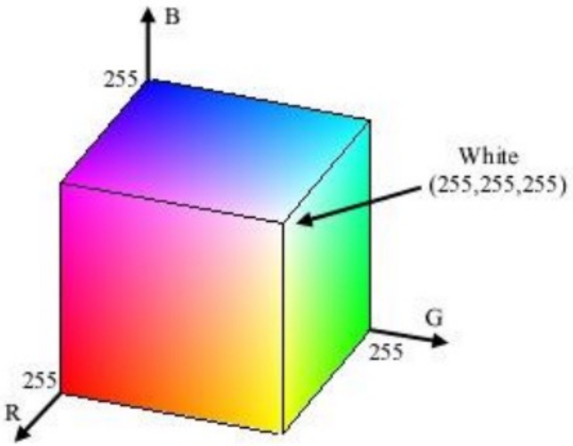

RGB values have no meaning if the primaries between devices are not the same!

HSV: Perceptual Color Space

Hue

Name of the color
(yellow, red, blue, green, ...)
the "kind" of color, regardless of its attributes

Value/Lightness/Brightness

How light or dark a color is.
total amount of light

Saturation/Chroma/Color Purity

How "strong" or "pure" a color is.
Purity, "colorfulness"

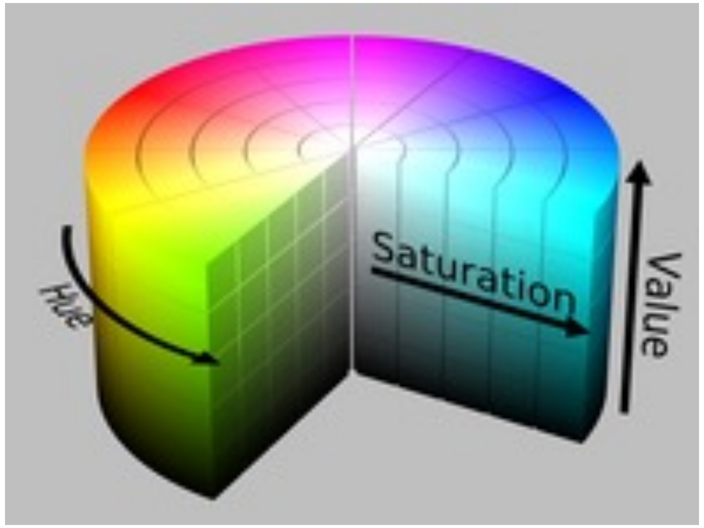

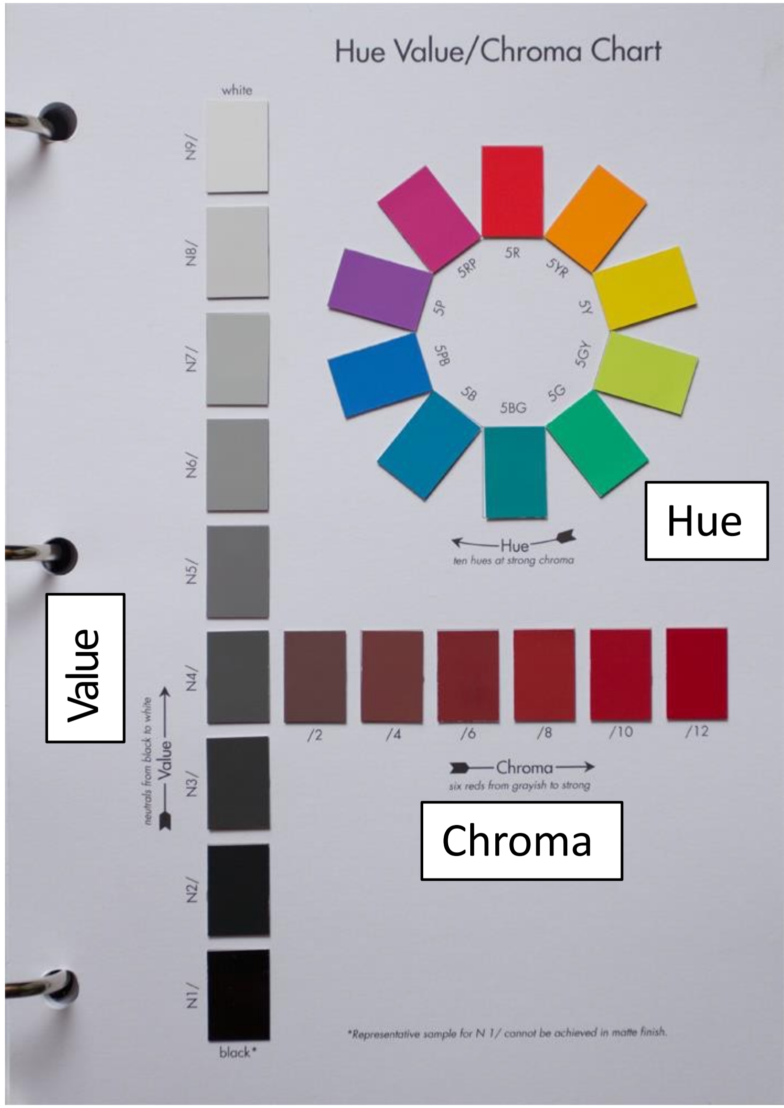

Image from Benjamin Salley
A page from a Munsell Student Color Set

Color spaces: HSV

Intuitive color space

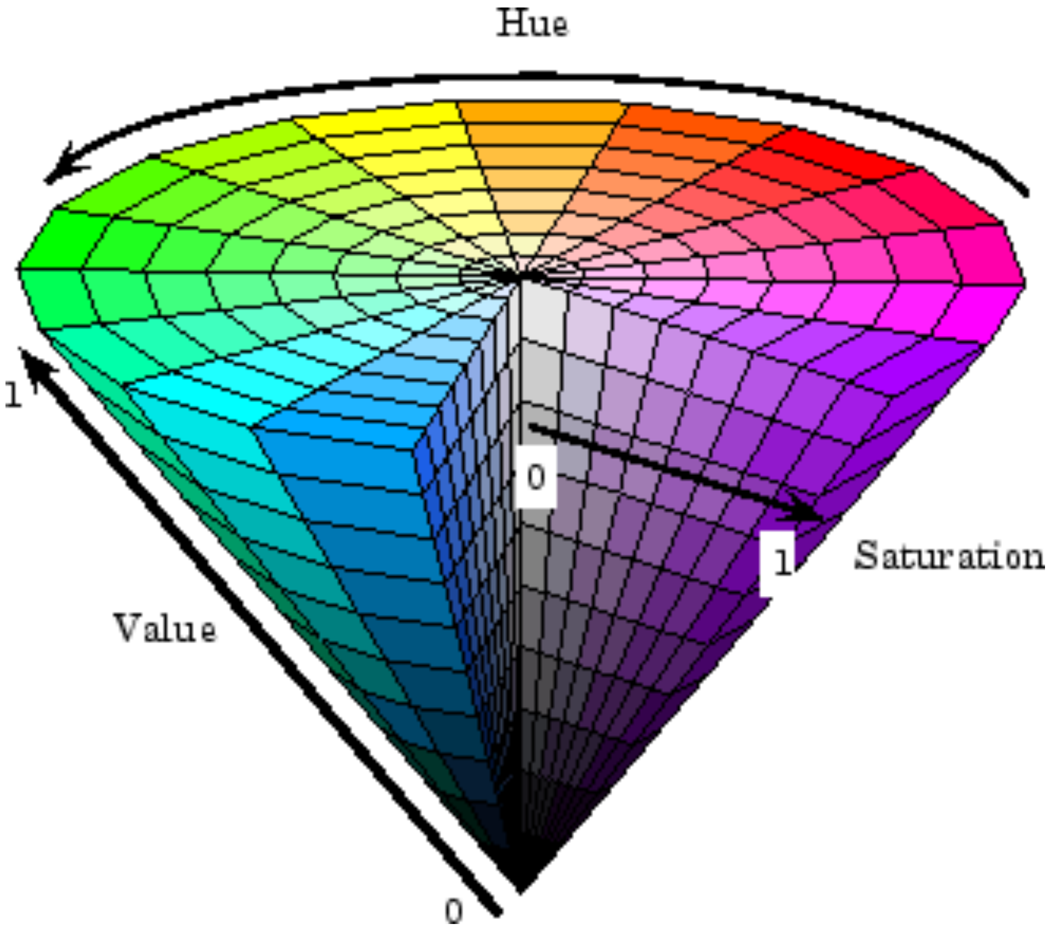

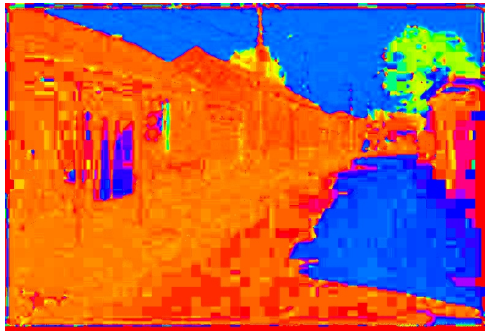

H
(S=1,V=1)

S
(H=1,V=1)

V
(H=1,S=0)

Color spaces: L*a*b*

"Perceptually uniform"* color space

L = Luma (Lightness)
a = Red to Green
b = Blue to yellow

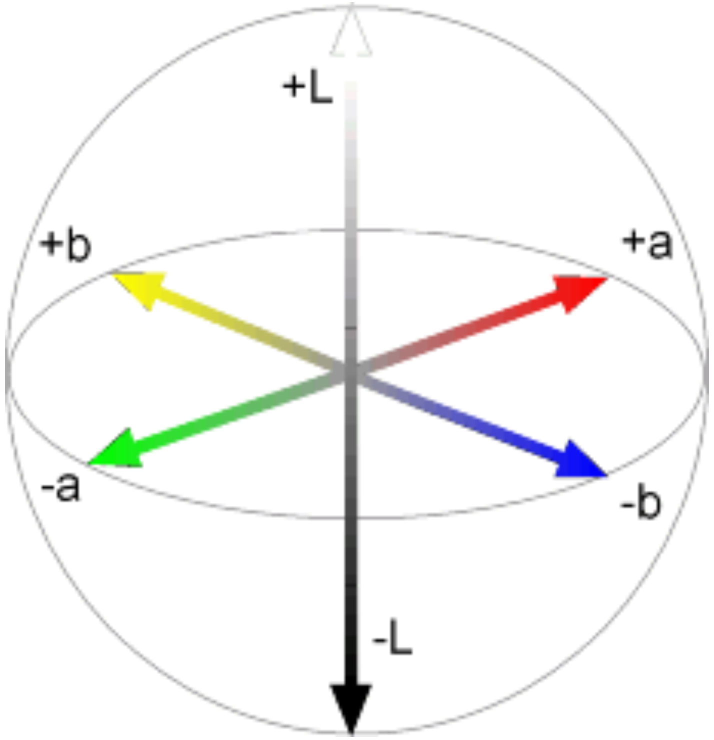

L
(a=0,b=0)

a
(L=65,b=0)

b
(L=65,a=0)

CIE LAB space

- CIE LAB space (also written as CIE L*a*b*) was introduced as a perceptually uniform color space
- CIE Lab transforms CIE to a new space where color (and brightness) differences are more uniform.

The ellipses shows the range of colors (around the center of the ellipse) that would be perceived as the same. We can see that CIE XYZ this is not uniform.

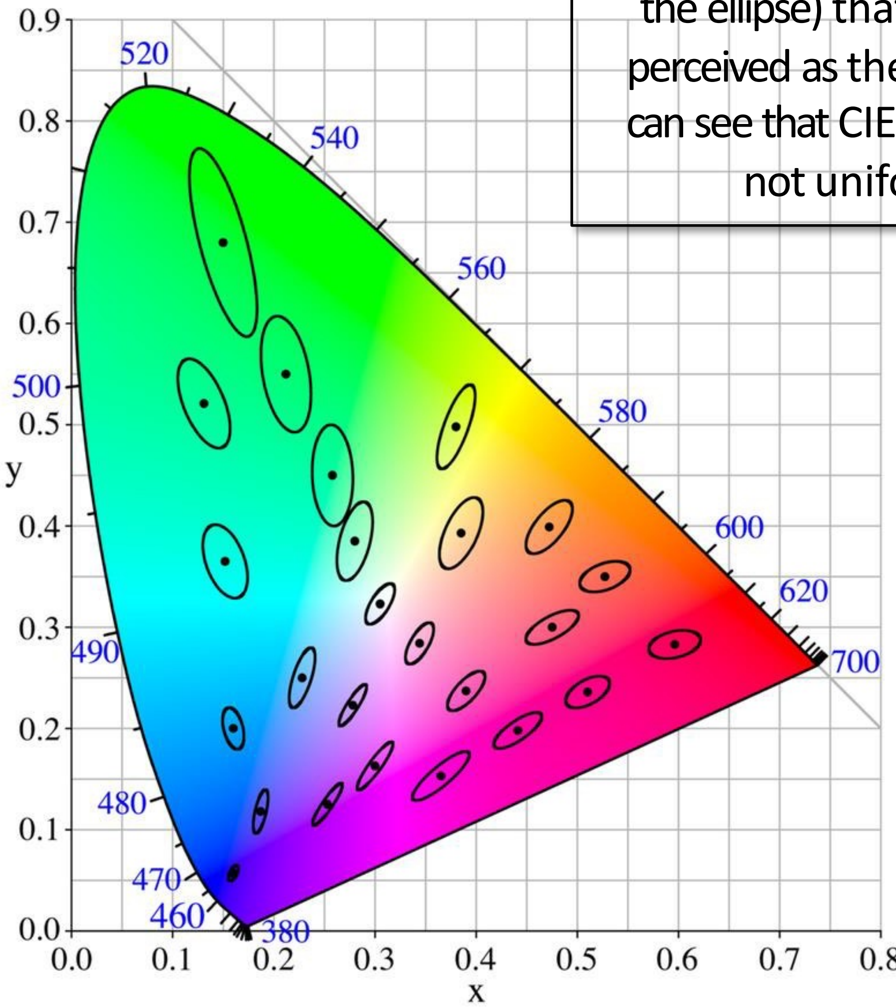

David MacAdam performed experiments into color perception. This plot is known as the MacAdam ellipses.

CIE LAB

CIE-L*ab space

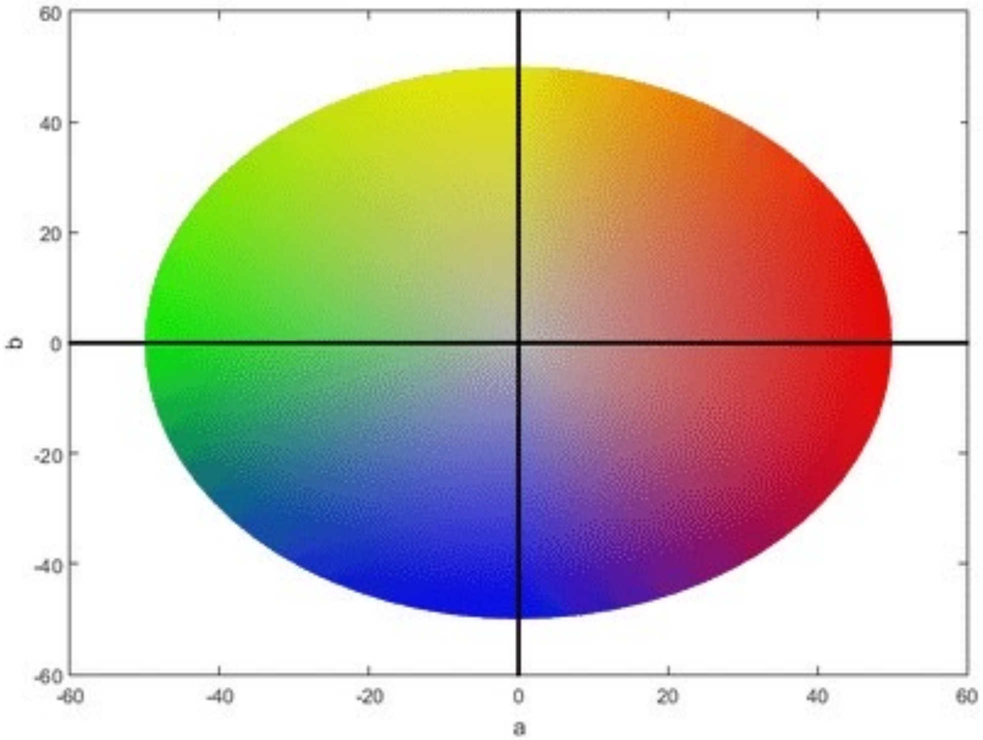

CIE-xyY space

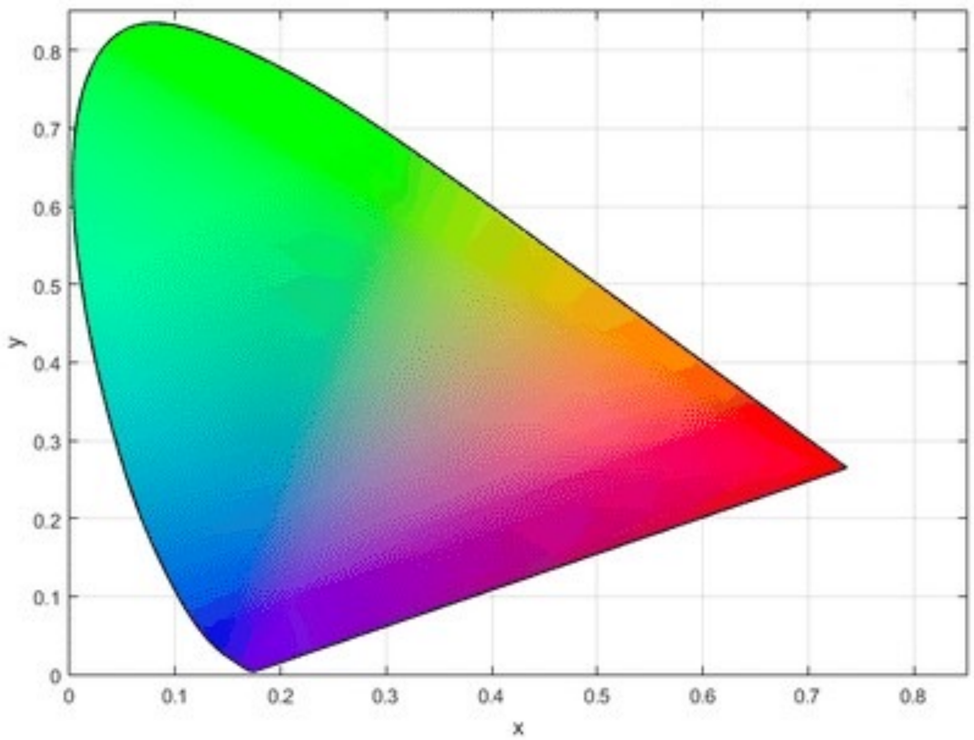

Chromaticity comparison's between CIE LAB and CIE XYX

Image from Bagdasar et al ICSTCC'17

I want to train Deep Learning algorithms, why should I care beyond RGB color spaces?

- If you are training ML systems to do object recognition, detection, etc, you shouldn’t care.

- But if you are solving different Image Processing & Computational Photography tasks, you should!

- CIE Lab color space is commonly used for tasks like:
  - Image Colorization
  - Intrinsic Image Decomposition
  - Reflectance Estimation

- Many papers often use other color spaces instead of RGB. It can help in easier learning.

Color error metric – CIE 2000 Delta E (ΔE)

- The CIE defined a color error metric in 2000 based on the CIE LAB space. This returns a color error between 0-100.
- You will see this referred to as CIEDE2000, CIEDE, ΔE, Delta E, DE, ..
- Delta E 2000 interpretation:

Delta E	Perception
<= 1.0	Not perceptible by human eyes.
1 - 2	Perceptible through close observation.
2 - 10	Perceptible at a glance.
11 - 49	Colors are more similar than opposite
100	Colors are exact opposite

In general, DE of 2 or less is considered to be very good. It means a standard observer could not tell that two colors are different unless they observed them very closely.

Table from
https://zschuessler.github.io/DeltaE/learn/

Interested in Color & Perception of Color in Data Visualization?

UNC VisuaLab, PI: Danielle Szafir

danielle.szafir@cs.unc.edu | https://cu-visualab.org/

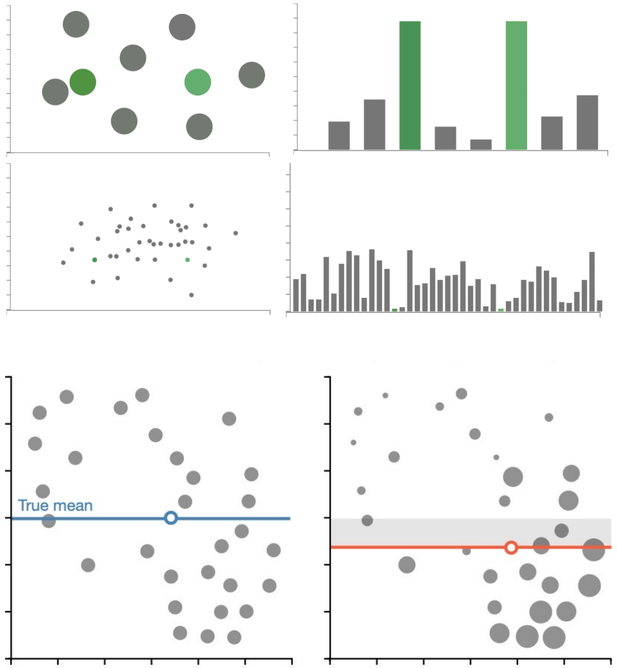

Measure color & other aspects of perception in data visualization

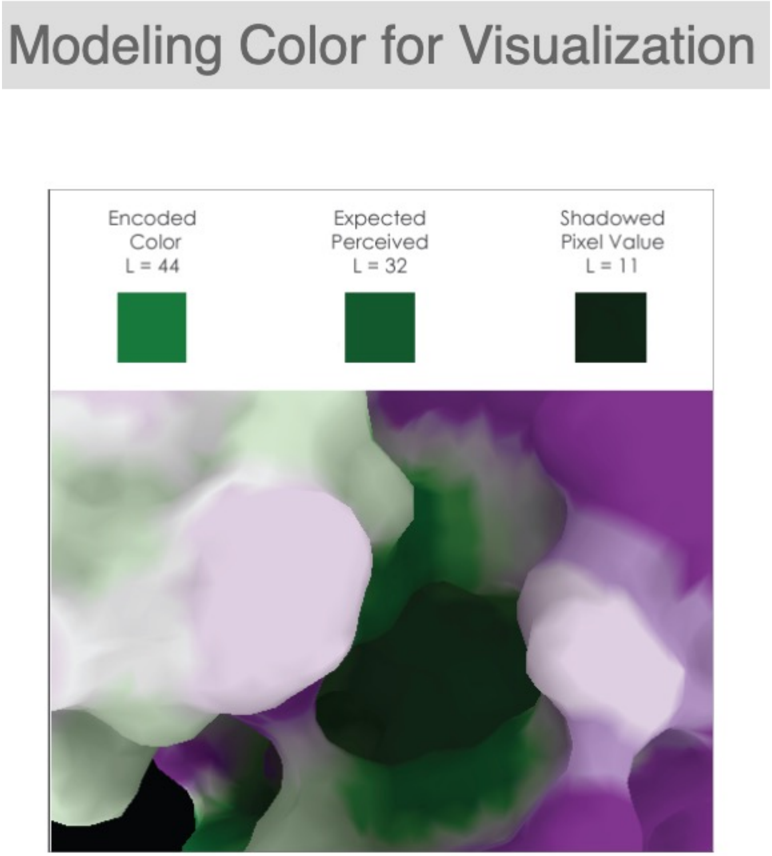

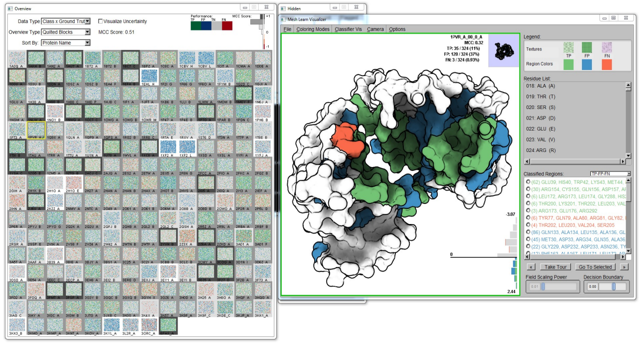

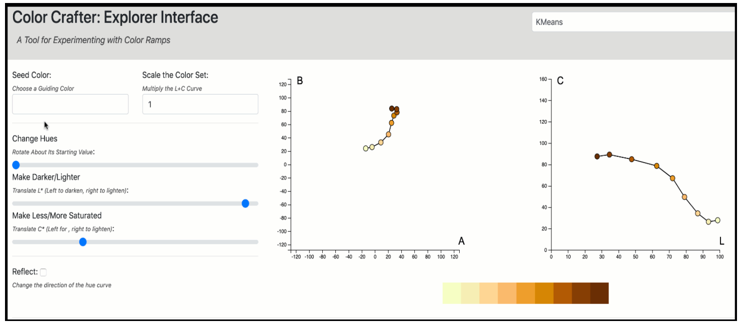

Use models of perception to create tools for data science

Today’s class

- Eye & Human Vision
- Pinhole Camera
- Perception of Color
- Quantifying Color
- Commonly Used Color Spaces
- Color Constancy

An object's SPD

- In a real scene, an object's SPD is a combination of the its reflectance properties **and** scene illumination

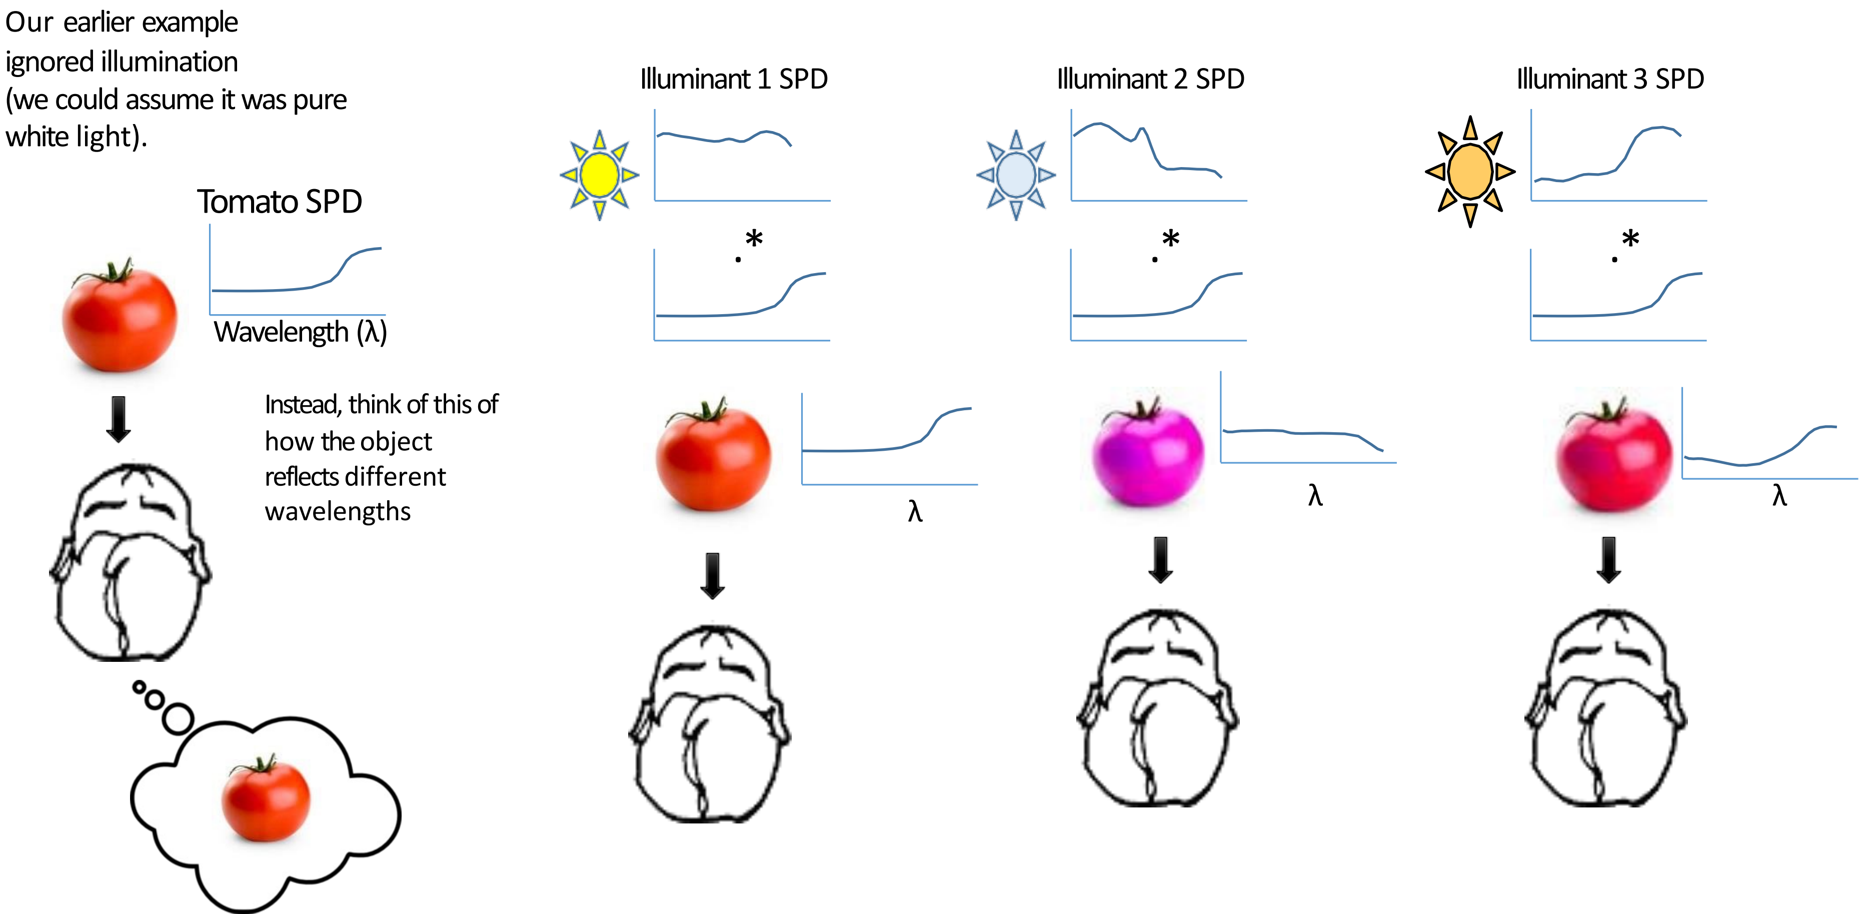

Color constancy

- Our visual system is able to compensate for the illumination

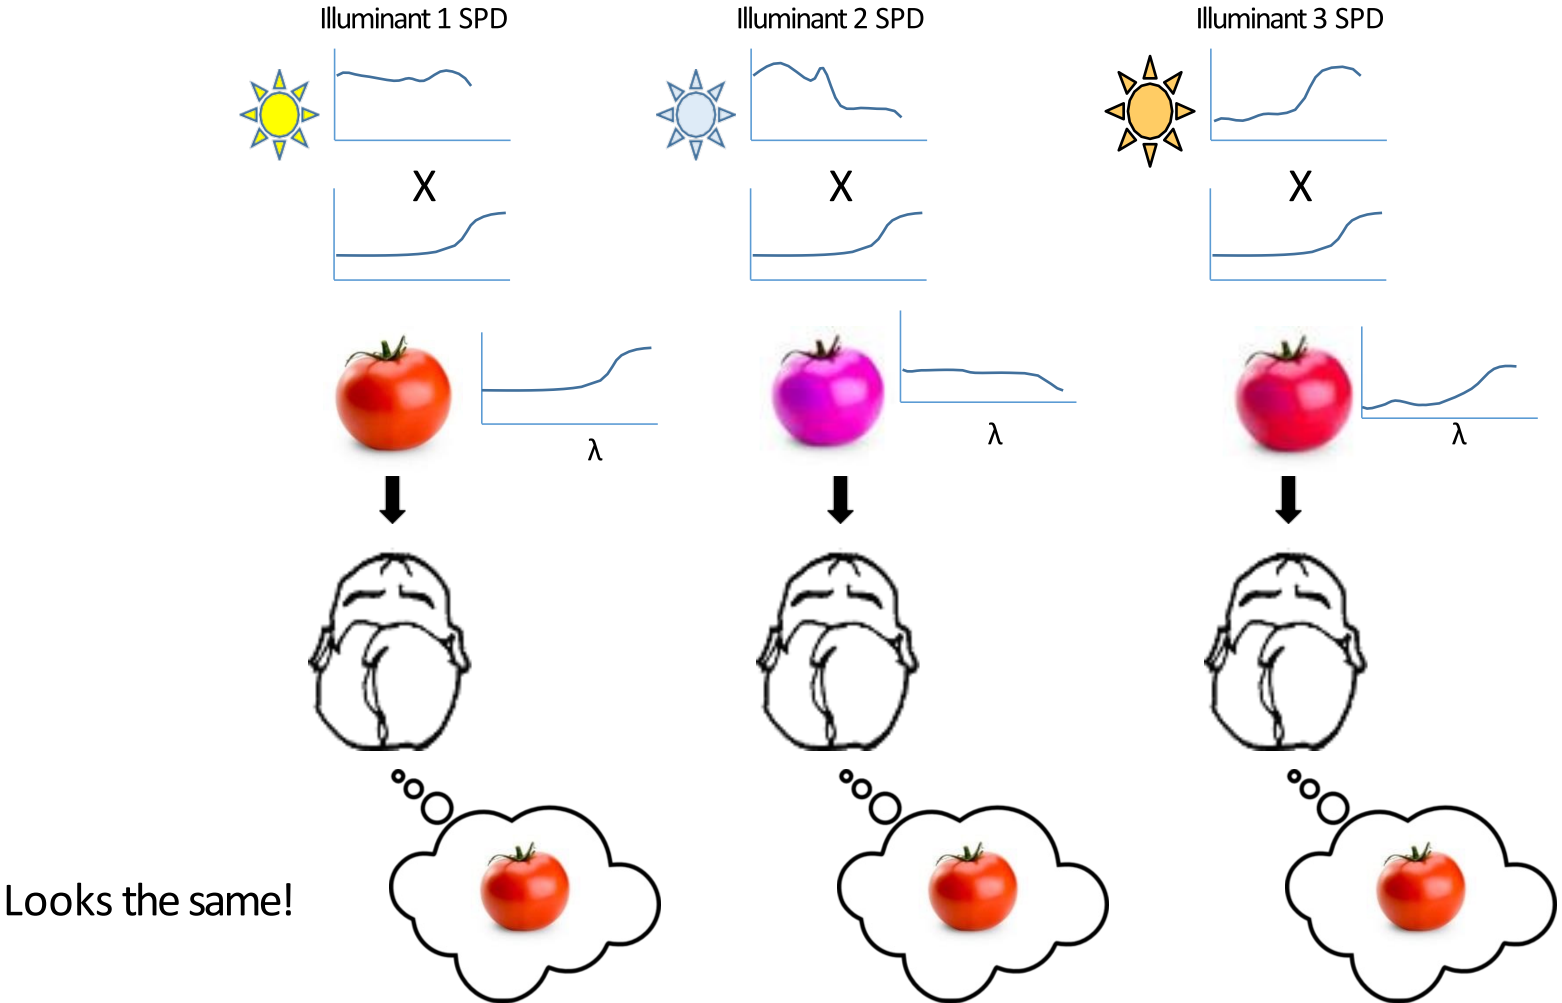

Chromatic adaptation example

Chromatic adaptation example

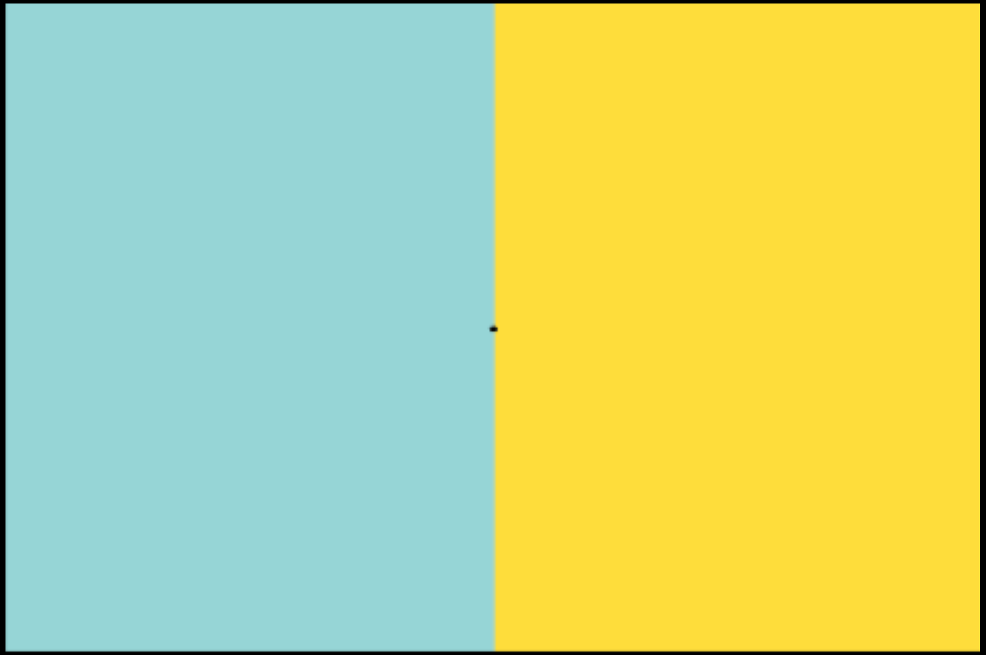

What color is the "The Dress"?

https://en.wikipedia.org/wiki/The_dress

Two Scene Interpretations of #thedress

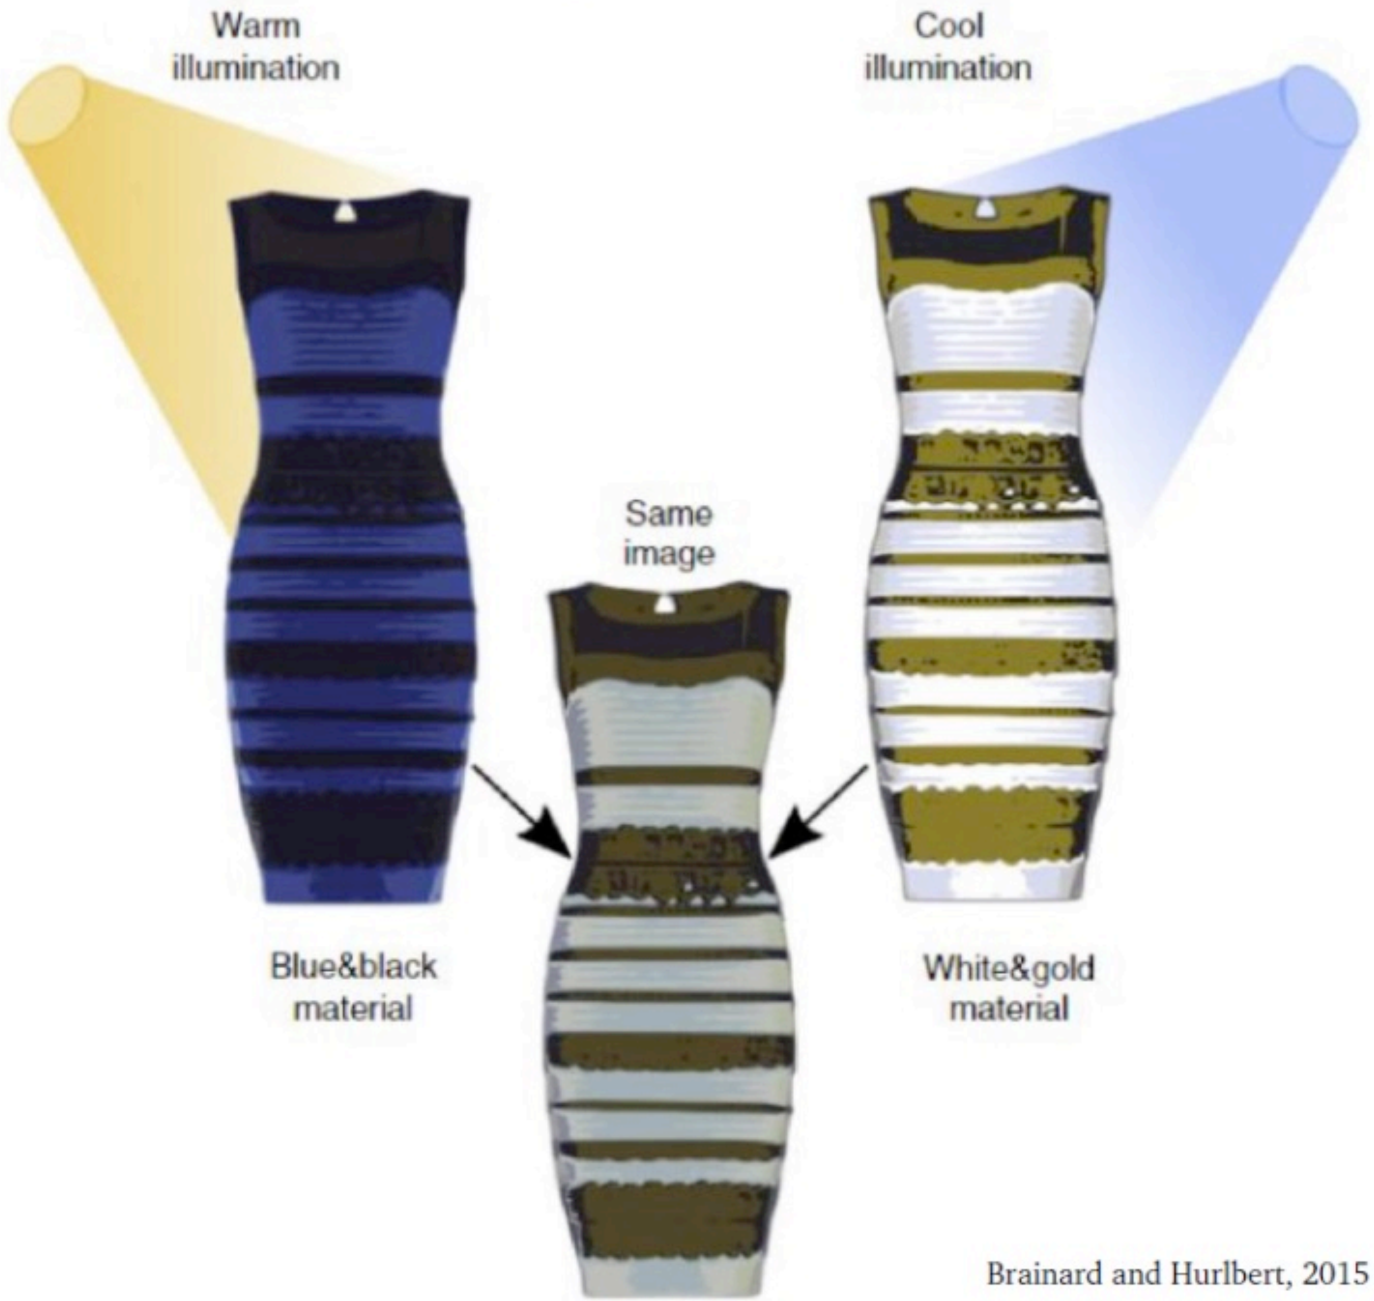

Brainard and Hurlbert, 2015

Color constancy/chromatic adaptation

- Color constancy (also called *chromatic adaptation*) is the ability of the human visual system to adapt to scene illumination
- This ability is not perfect, but it works fairly well
- Image sensors do not have this ability (it must be performed as a processing step, i.e. "white balance")

Note: Our eyes do not adjust to the illumination in the photograph -- we adjust to the viewing conditions of the scene we are viewing the photograph!

White point

- A white point is a CIE XYZ or CIE xyY value of an ideal “white target” or “white reference”
- The idea of chromatic adaptation is to make white points the same between scenes.

Color constancy (at its simplest)


Johannes von Kries

- (Johannes) *Von Kries* transform
- Compensate for each channel corresponding to the L, M, S cone response

$$\begin{bmatrix} L_2 \\ M_2 \\ S_2 \end{bmatrix} = \begin{bmatrix} 1/L_{1w} & 0 & 0 \\ 0 & 1/M_{1w} & 0 \\ 0 & 0 & 1/S_{1w} \end{bmatrix} \begin{bmatrix} L_1 \\ M_1 \\ S_1 \end{bmatrix}$$

L_2, M_2, S_2 is the new LMS response with the illuminant divided "out". In this case white is equal to [1,1,1]

L_{1w}, M_{1w}, S_{1w} is the LMS response to "white" under this illuminant

L_1, M_1, S_1 are the input LMS space under an illuminant.

Example

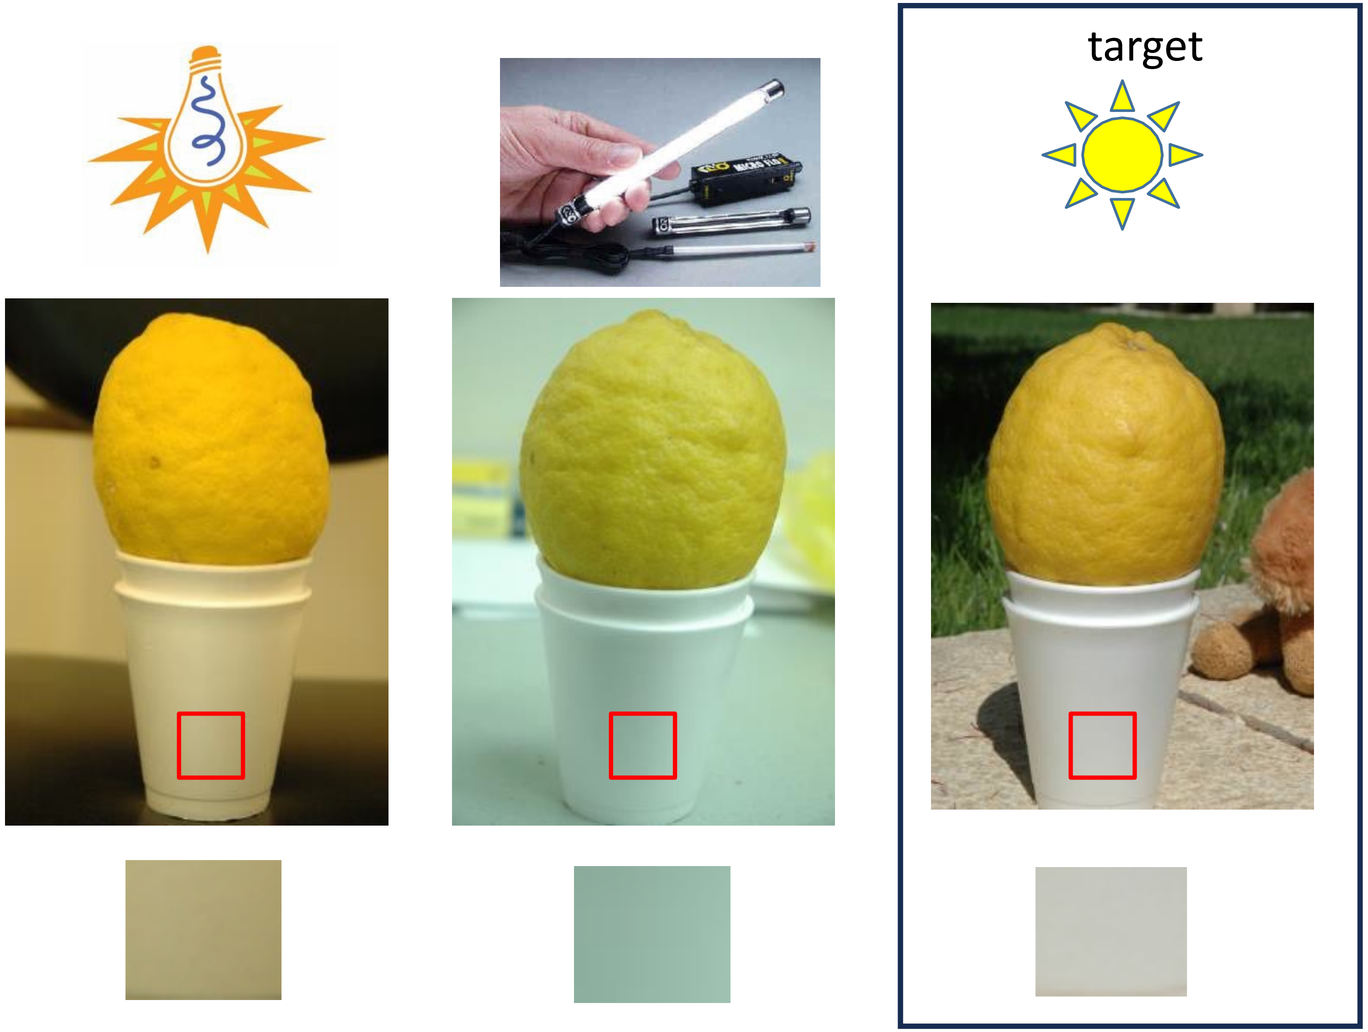

Simulation of different "white points" by photographing a "white" object under different illumination.

Images courtesy of Sharon Albert (Weizmann Institute)

Example

Input

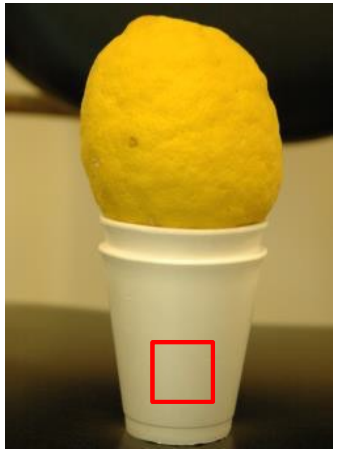

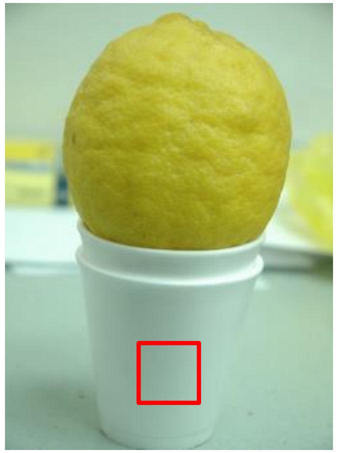

Here, we have mapped the two input images to one below to mimic chromatic adaptation. The "white" part of the cup is shown before and after to help show that the illumination falling on white appears similar after the "chromatic adaptation".

Adapted to "target" illuminant

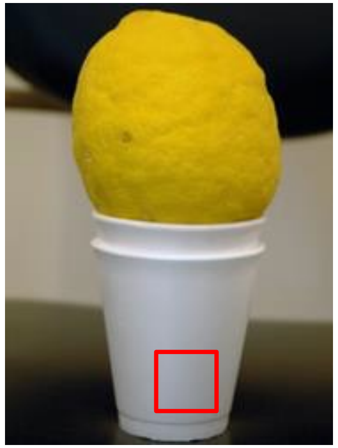

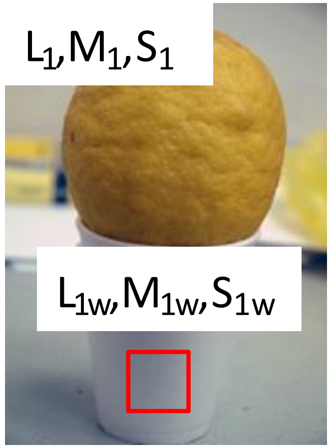

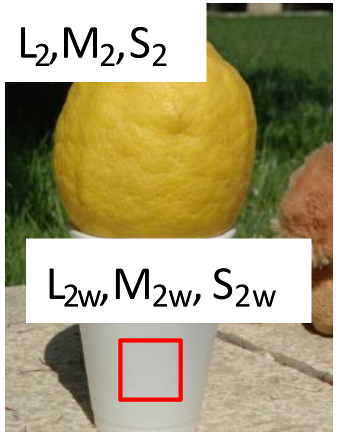

Before

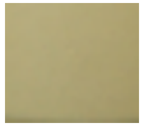

After

Before

After

Target Illumination

Illuminant to illuminant mapping

- More appropriate would be to map to another illuminant's LMS response (e.g. in the desired viewing condition)
- $(LMS)_1$ under an illuminant with white-response (L_{1w}, M_{1w}, S_{1w})
- $(LMS)_2$ under an illuminant with white-response (L_{2w}, M_{2w}, S_{2w})

$$\begin{bmatrix} L_2 \\ M_2 \\ S_2 \end{bmatrix} = \begin{bmatrix} L_{2w}/L_{1w} & 0 & 0 \\ 0 & M_{2w}/M_{1w} & 0 \\ 0 & 0 & S_{2w}/S_{1w} \end{bmatrix} \begin{bmatrix} L_1 \\ M_1 \\ S_1 \end{bmatrix}$$

L_2, M_2, S_2 is the new LMS response with the illuminant divided "out" and scaled to LMS_2 illuminant

L_{1w}, M_{1w}, S_{1w} is the LMS response to "white" the input illuminant, L_{2w}, M_{2w}, S_{2w} response to "white" of output illuminant

L_1, M_1, S_1 are the input LMS space under an illuminant.

White Balancing

- The above method is a very rudimentary way of performing white balancing. It requires capturing the ‘white point’ under target illumination!
- Modern methods do not require a target ‘white point’. They are called Auto White Balancing (AWB)
- Most modern methods use Deep Learning to perform white-balancing from a single image only.
- More about this in next class … as a part of camera pipeline.

Today's class

- Eye & Human Vision
- Pinhole Camera
- Perception of Color
- Quantifying Color
- Commonly Used Color Spaces
- Color Constancy

Next Class: A typical color imaging pipeline

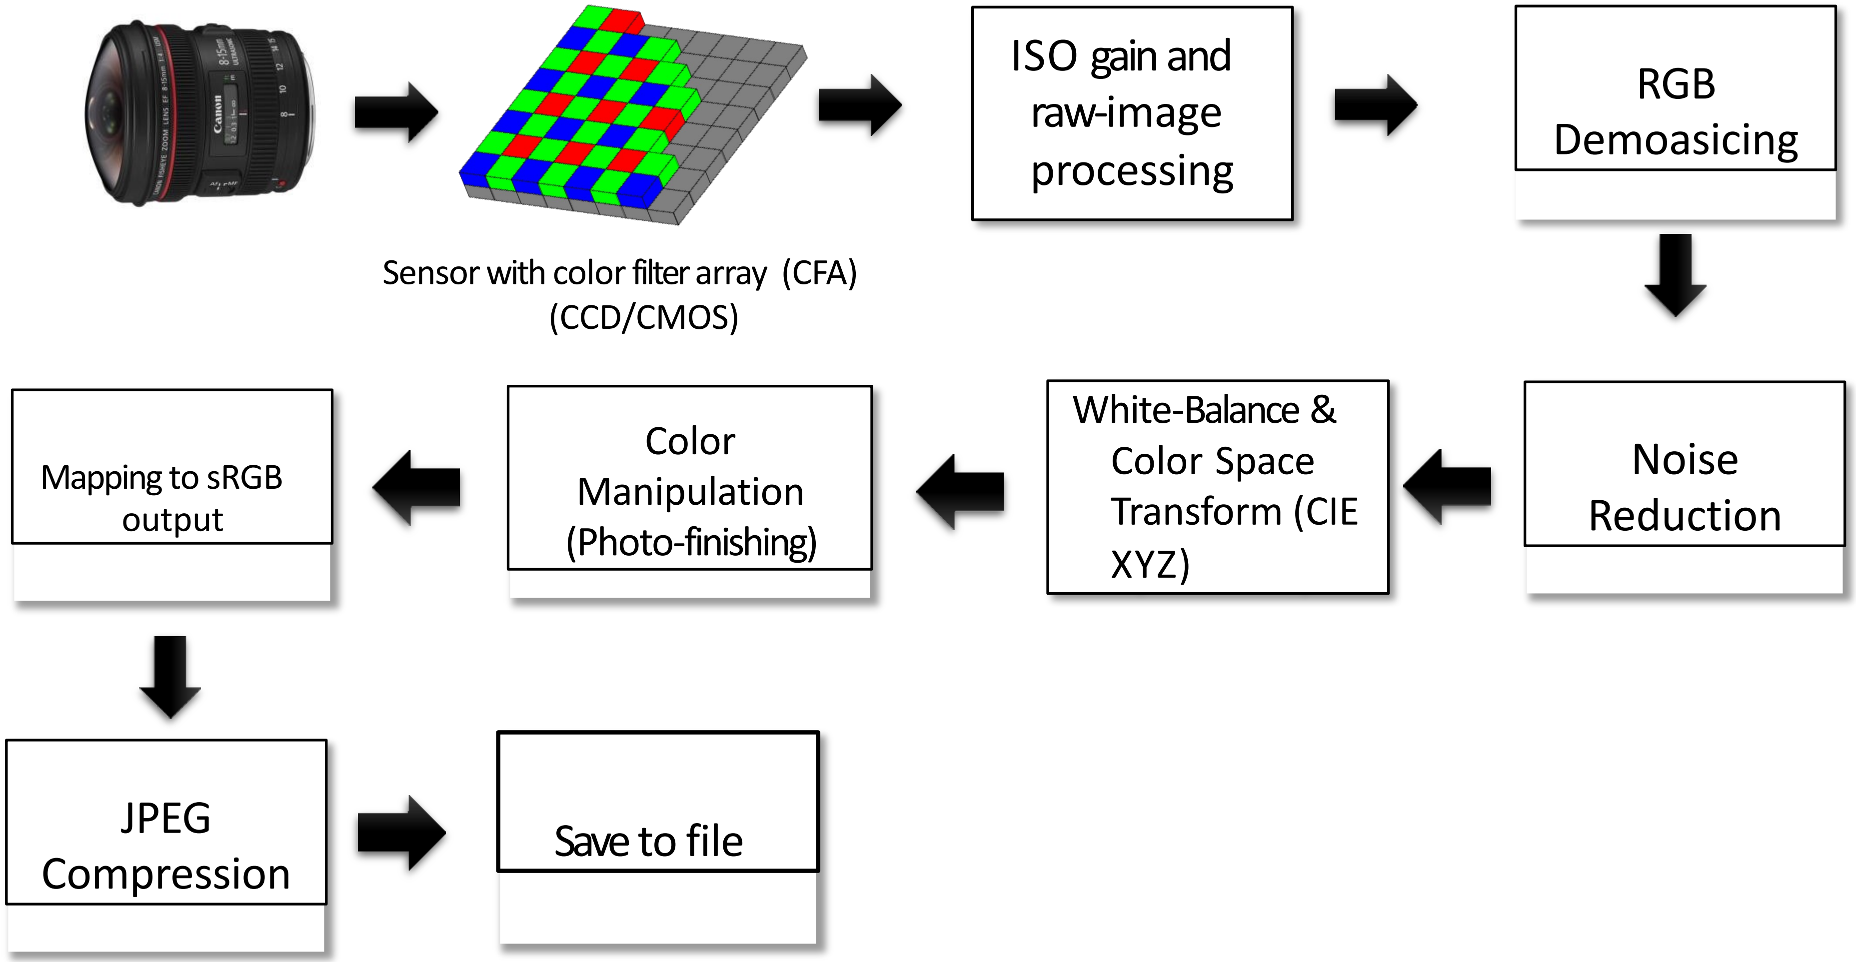

NOTE:This diagram represents the steps applied on a typical consumer camera pipeline. ISPs may apply these steps in a different order or combine them in various ways. A modern camera ISP will undoubtedly be more complex, but will almost certainly implement these steps in some manner.

Watch these 5 min videos!

- Color in 5 min: https://youtu.be/6tTNgvAl1y4
- Displays in 5 min: https://youtu.be/1albYPL9Cfg
- Pinhole Camera in 5 min: https://youtu.be/F5WA26W4JaM

Additional Reading

Sec 3.1, 3.2, 3.3 from Forsyth & Ponce

Slide Credits

- **"Understanding Color and the In-Camera Image Processing Pipeline for Computer Vision", by Michael S. Brown, ICCV 2019 tutorial.**
- CS 194-26/294-26: Intro to Computer Vision and Computational Photography**, UC Berkeley, by Alyosha Efros.**
- **CS 15-463, 663, 862, CMU, by Computational Photography, Ioannis Gkioulekas.**