Capturing Light… in human and machine

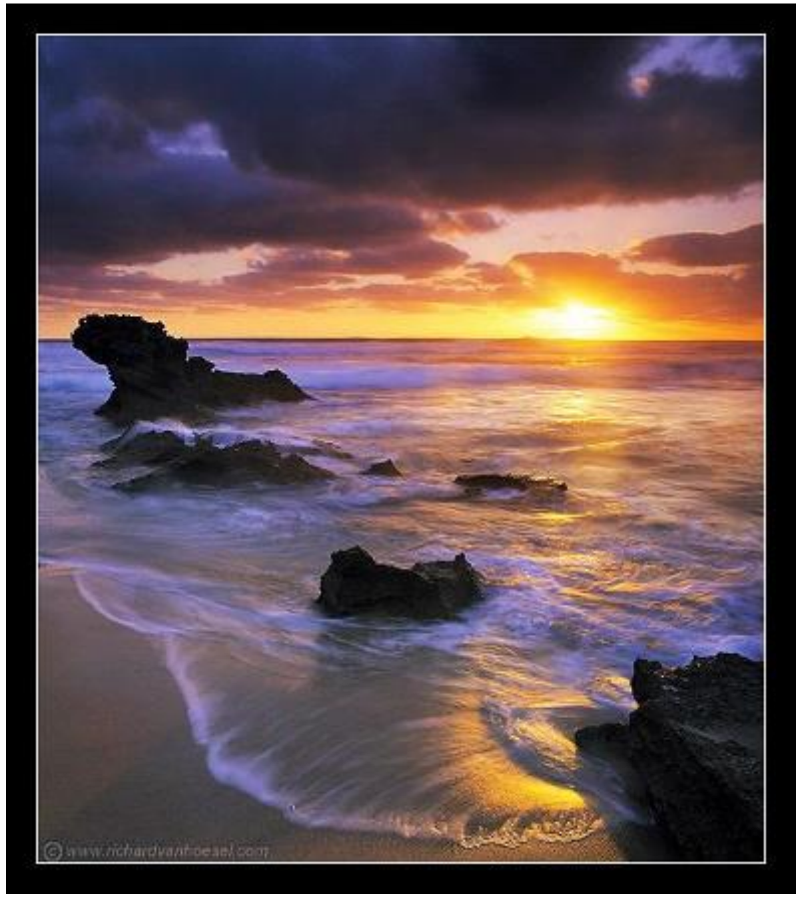

CS180: Computer Vision and Comp. Photography
Efros & Kanazawa, UC Berkeley, Fall 2025

Image Formation

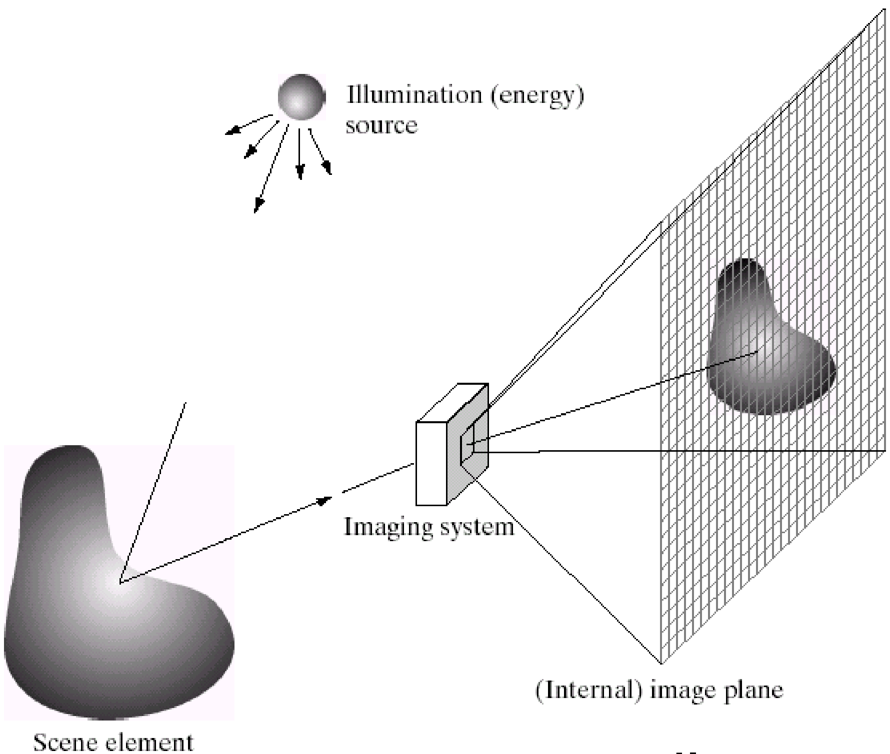

Film

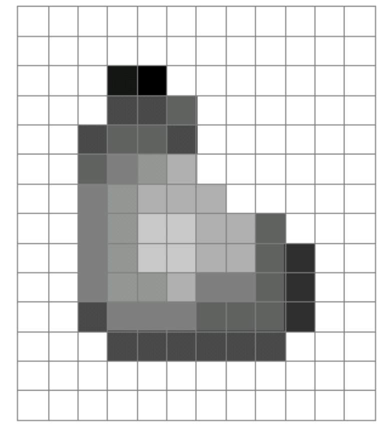

Digital Camera

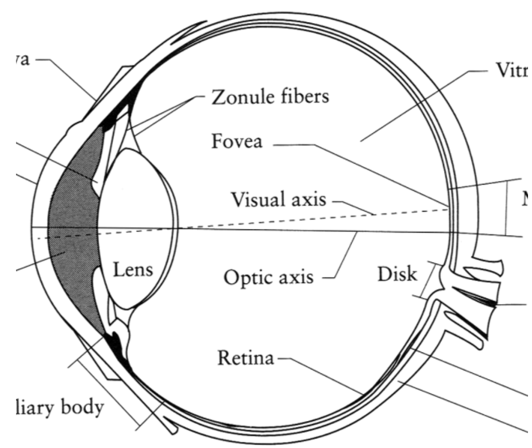

The Eye

Pixel (**Pic**ture **El**ement) Array

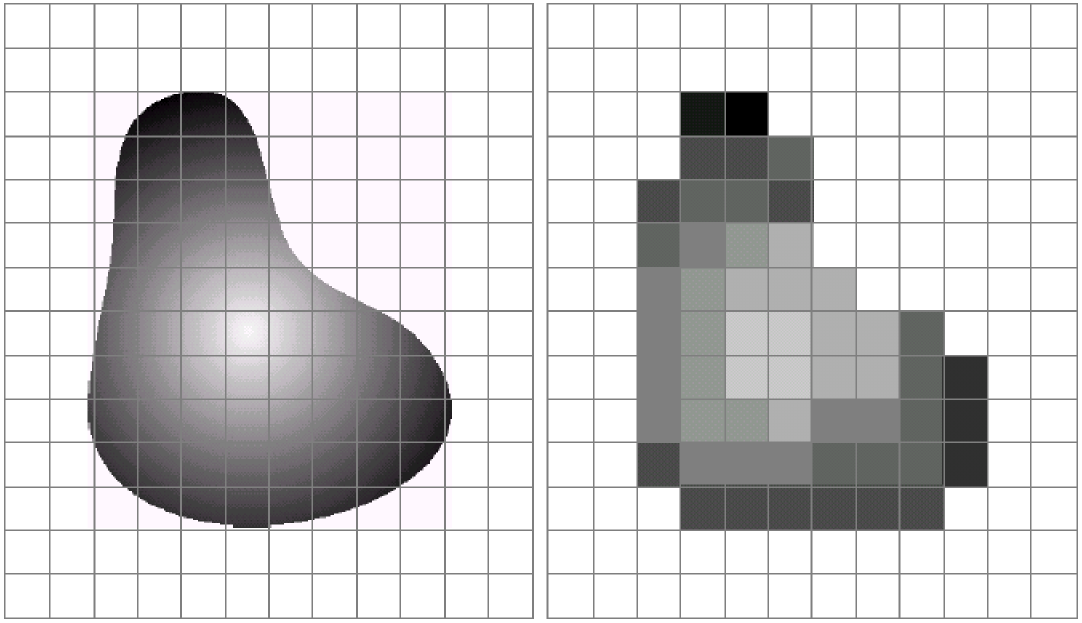

a b

FIGURE 2.17 (a) Continuos image projected onto a sensor array. (b) Result of image sampling and quantization.

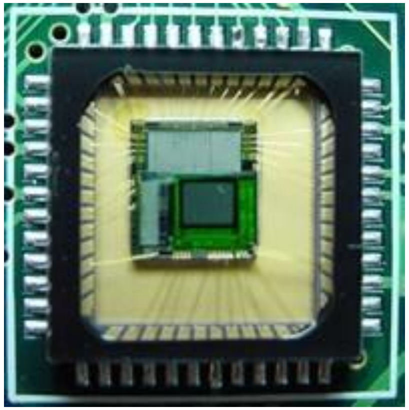

CMOS sensor

Rolling Shutter

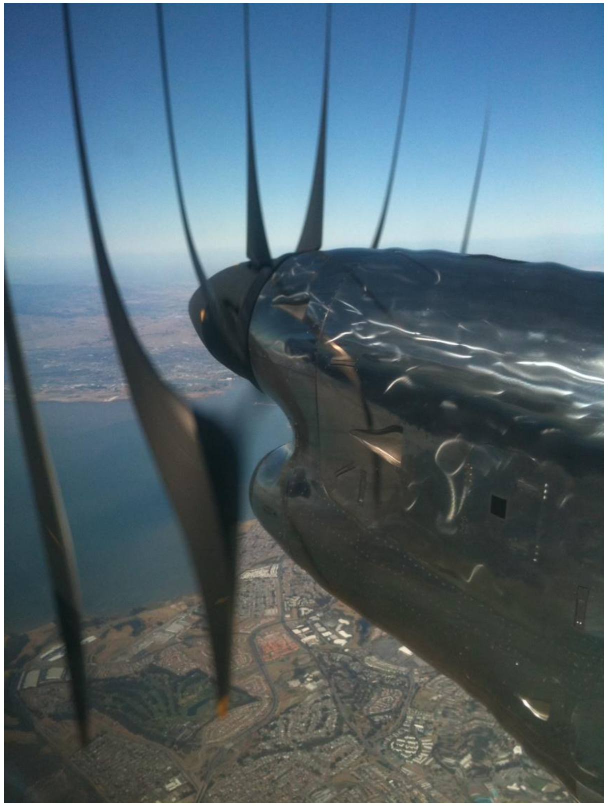

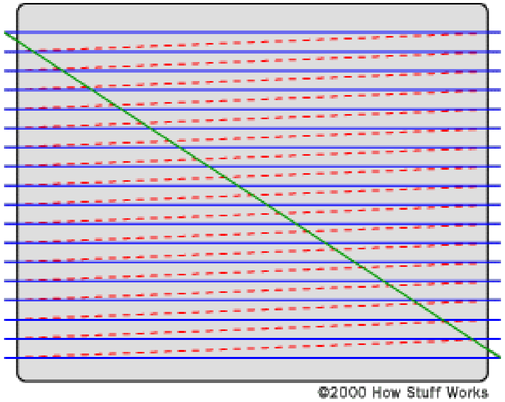

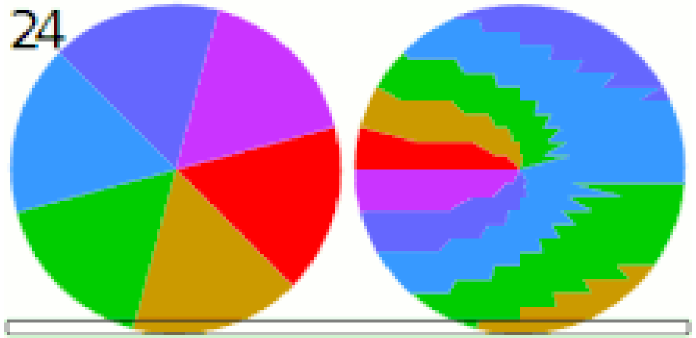

http://en.wikipedia.org/wiki/Rolling_shutter

Saccadic eye movement

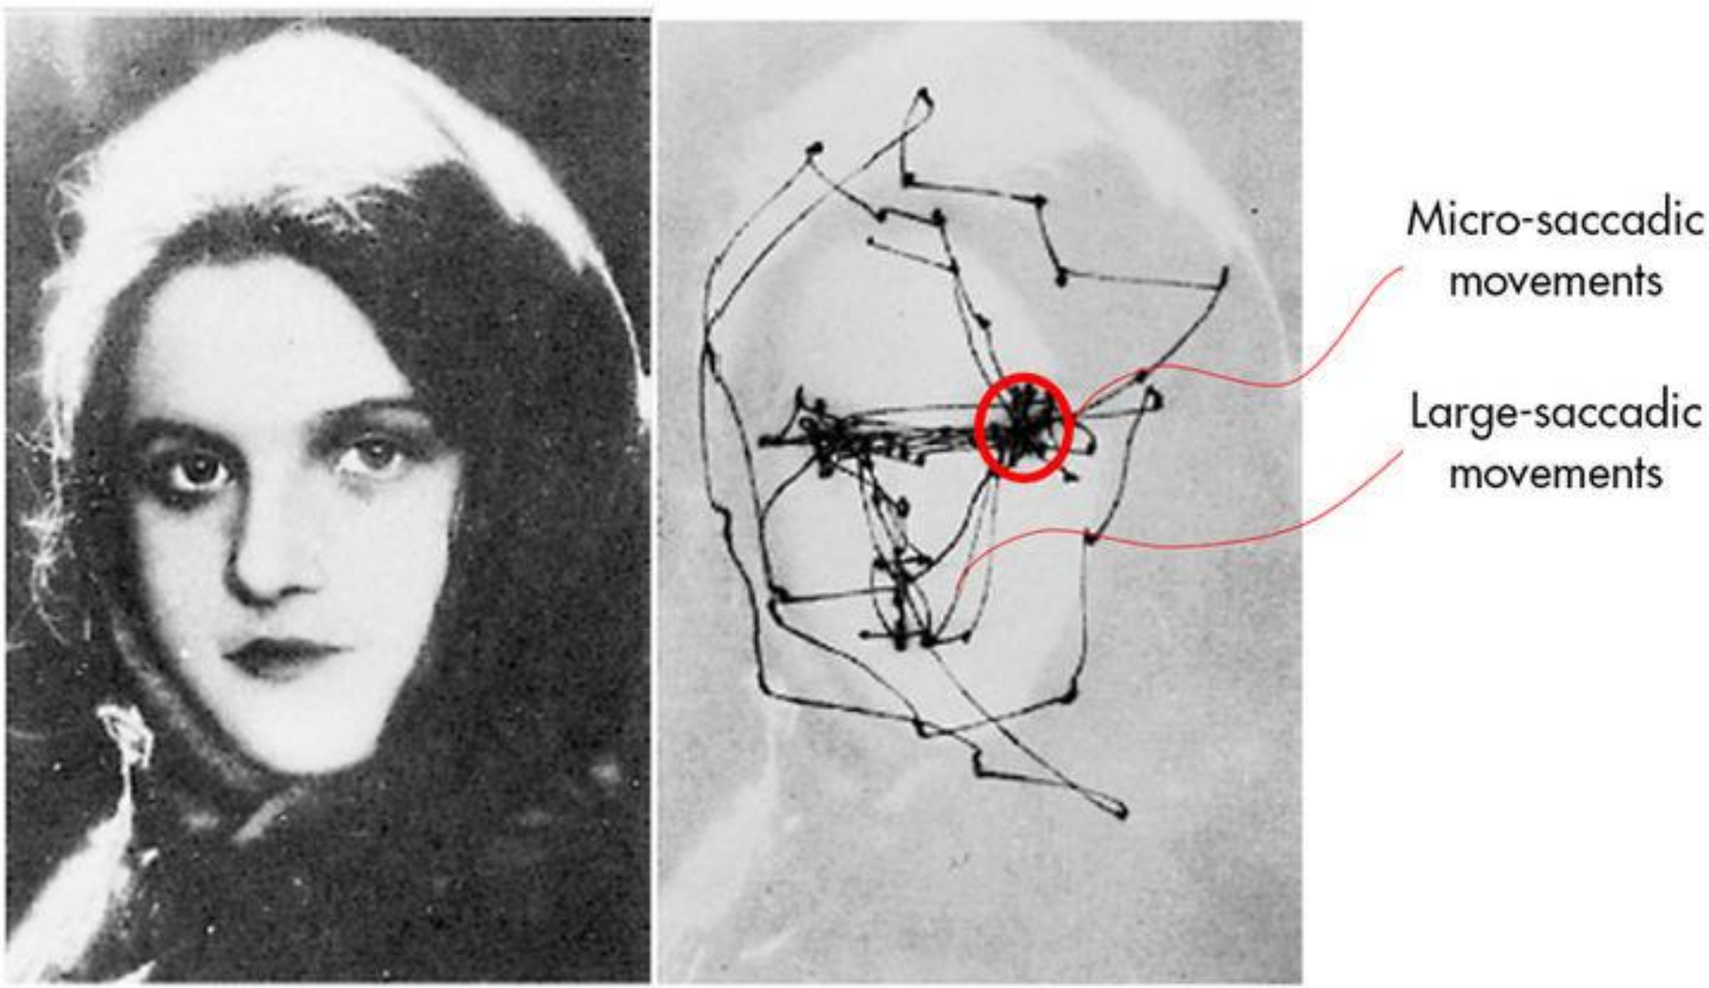

Saccadic eye movement

The Eye

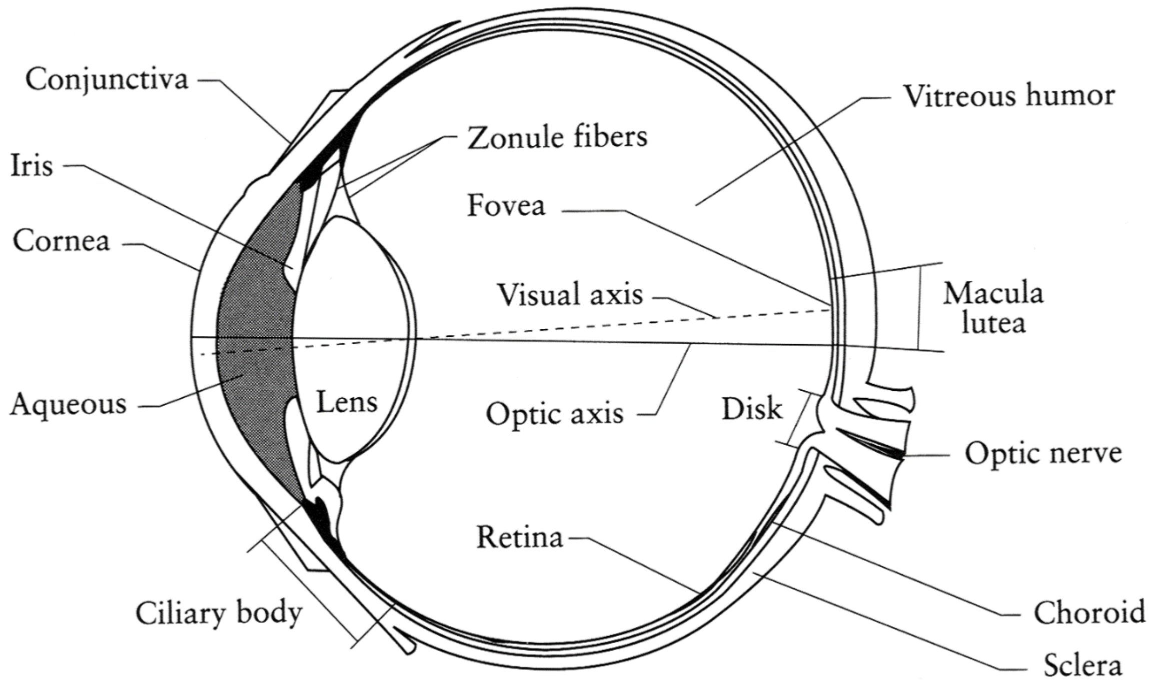

The human eye is a camera!

- **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

The Retina

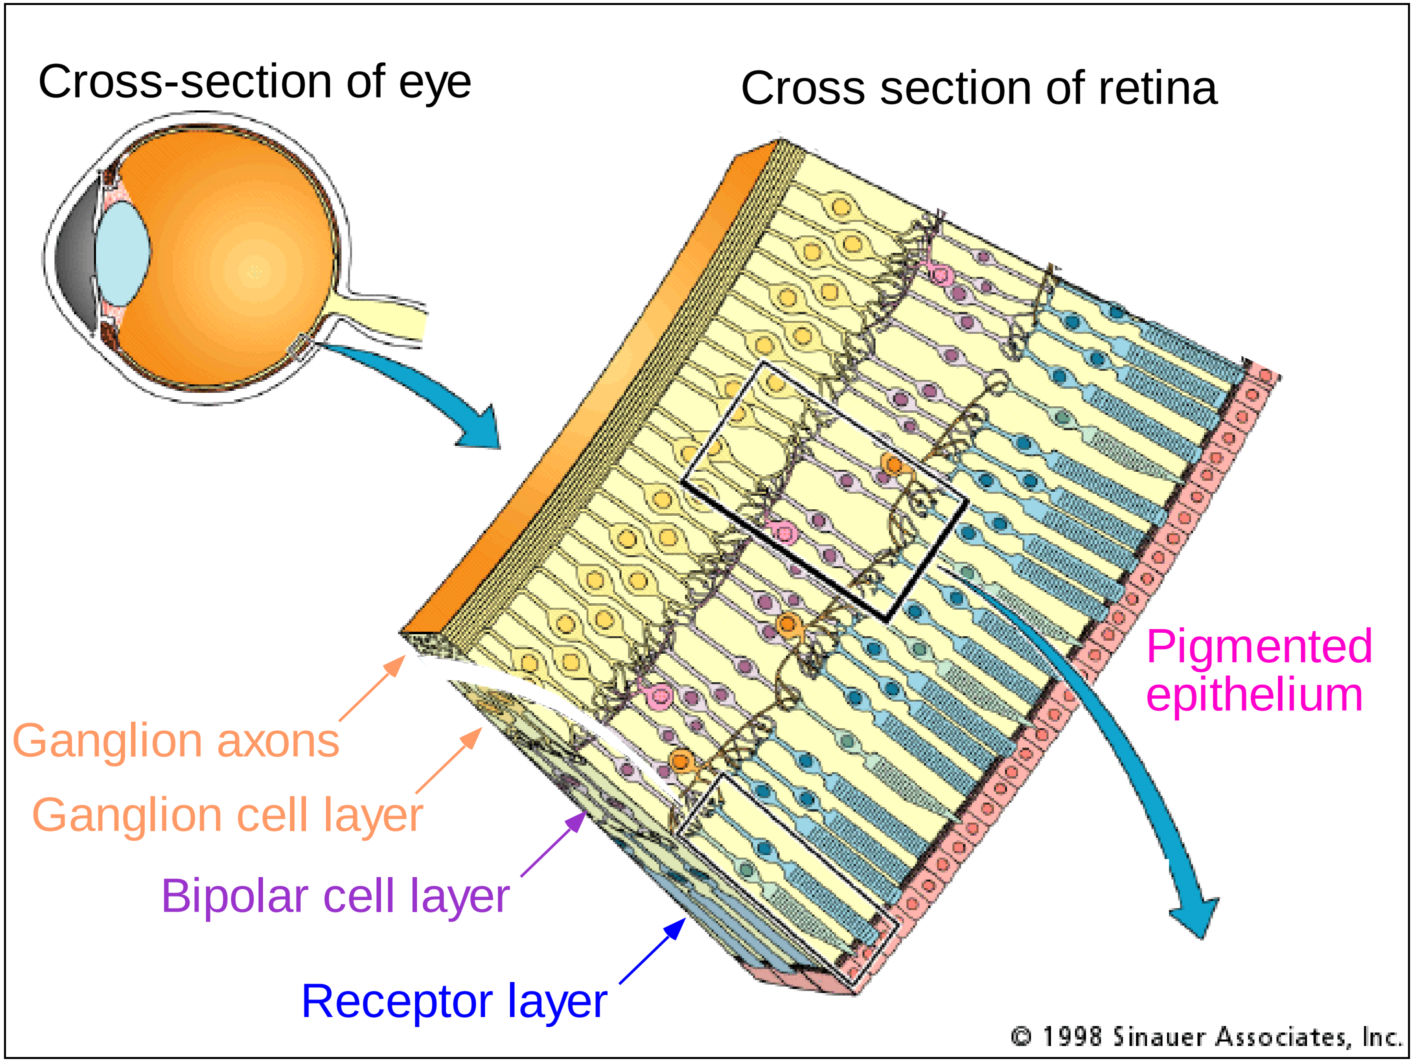

Cross-section of eye
Cross section of retina
Pigmented
epithelium
Ganglion axons
Ganglion cell layer
Bipolar cell layer
Receptor layer
© 1998 Sinauer Associates, Inc.

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

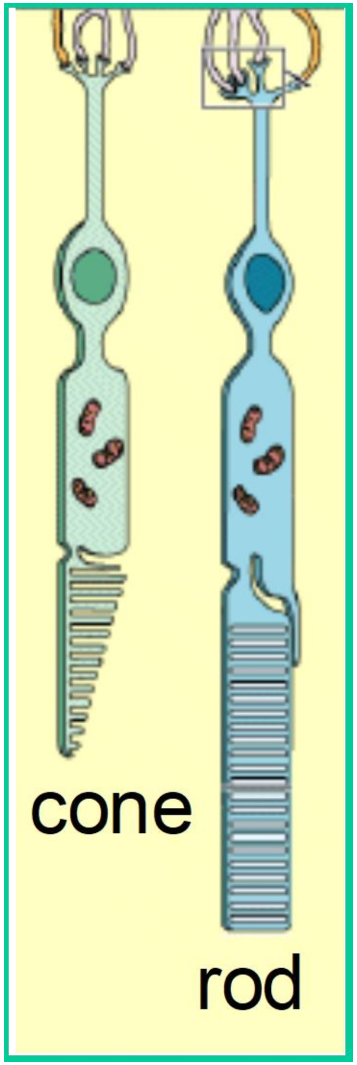

Rod / Cone sensitivity

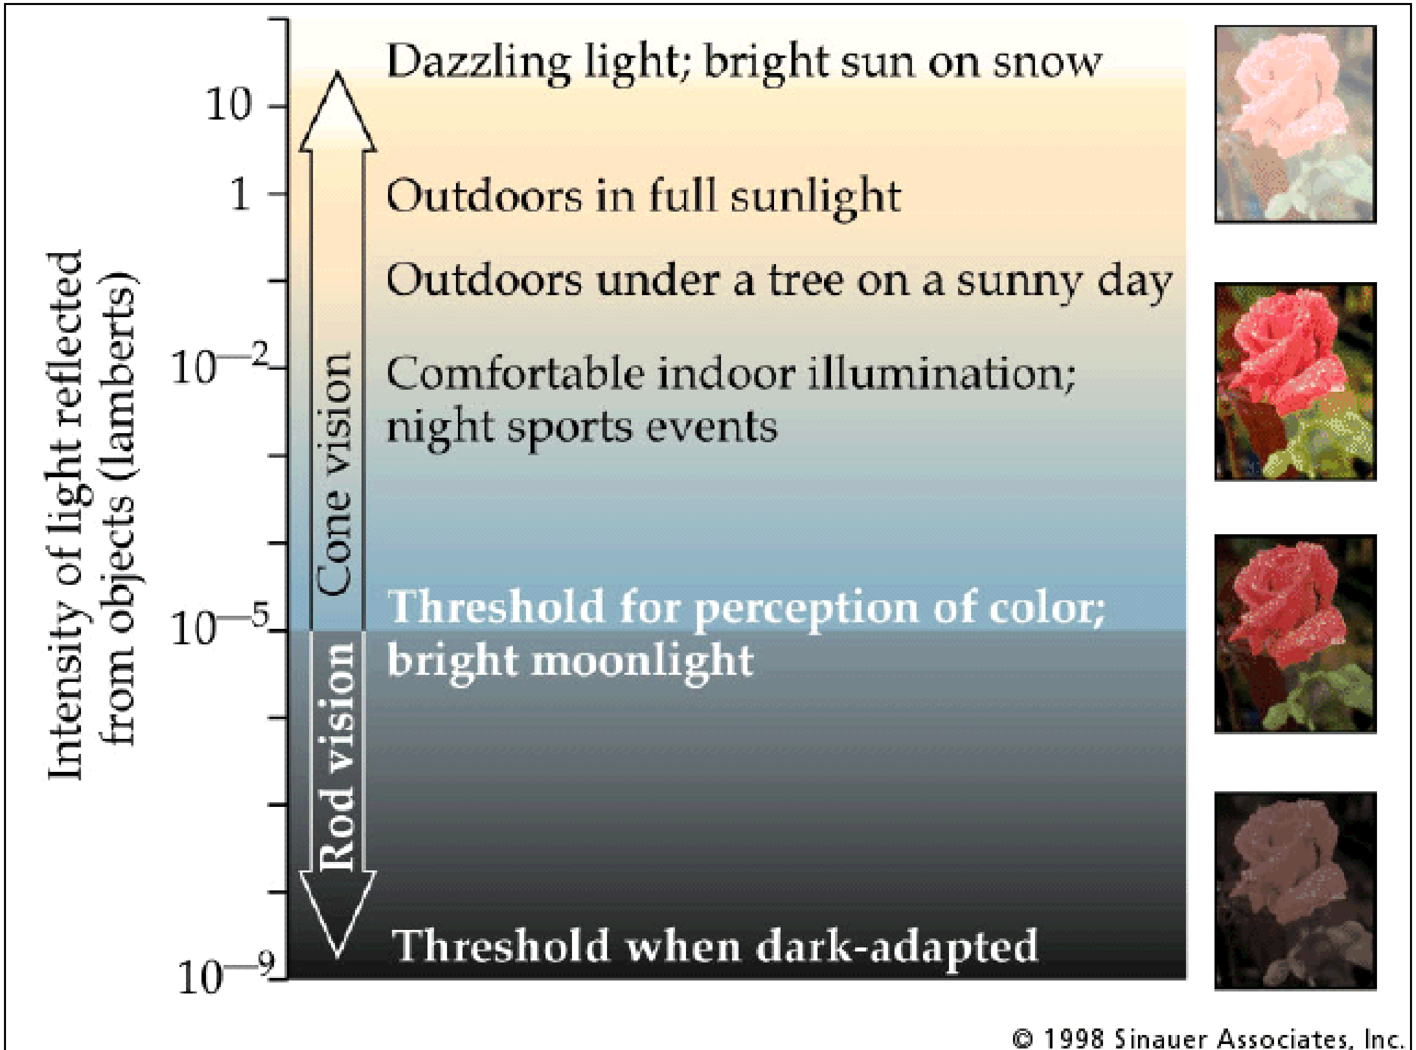

The famous sock-matching problem…

Retina up-close

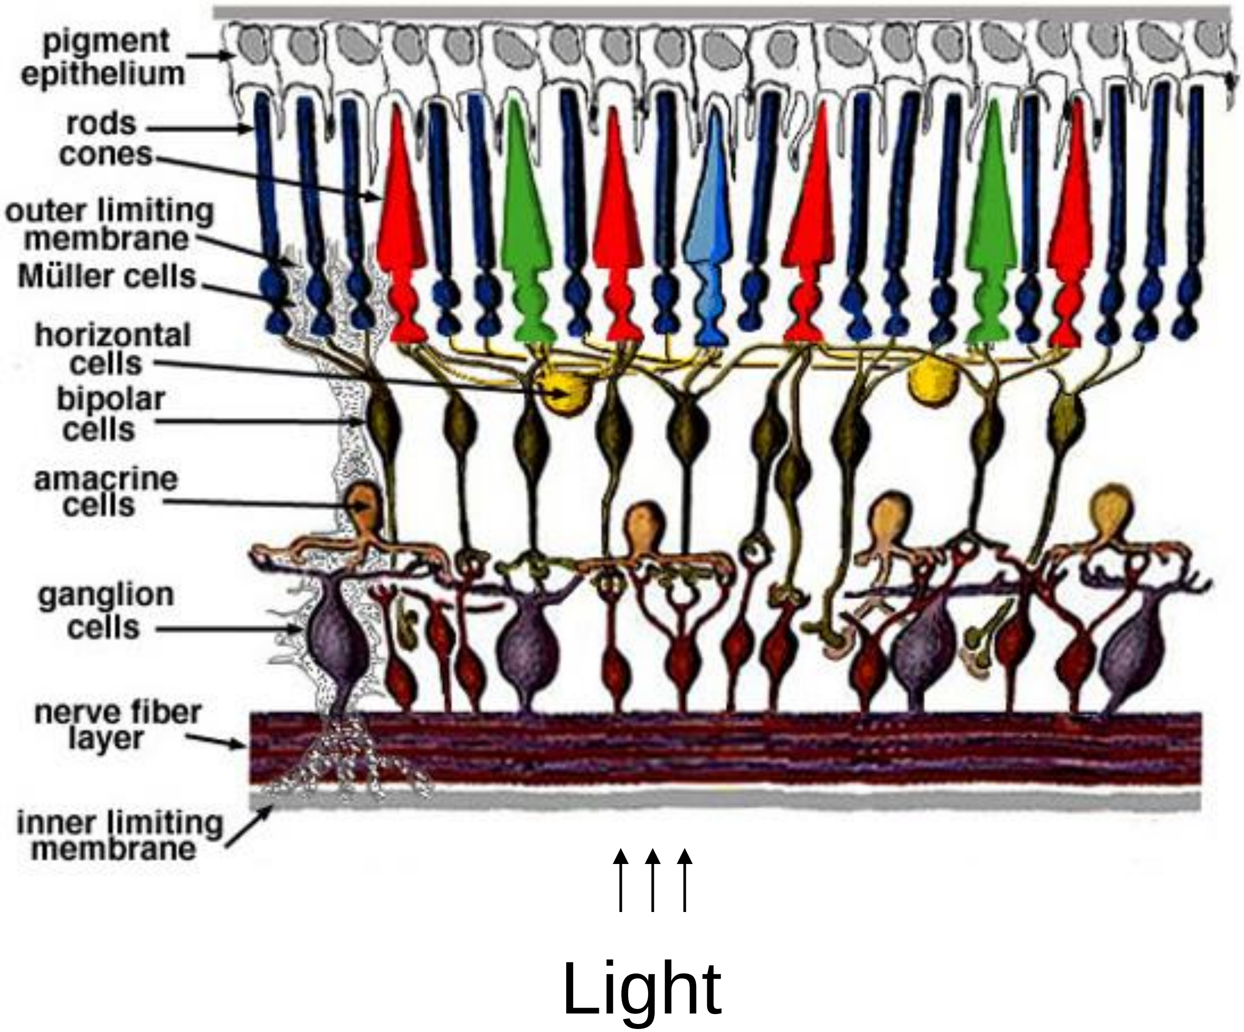

Distribution of Rods and Cones

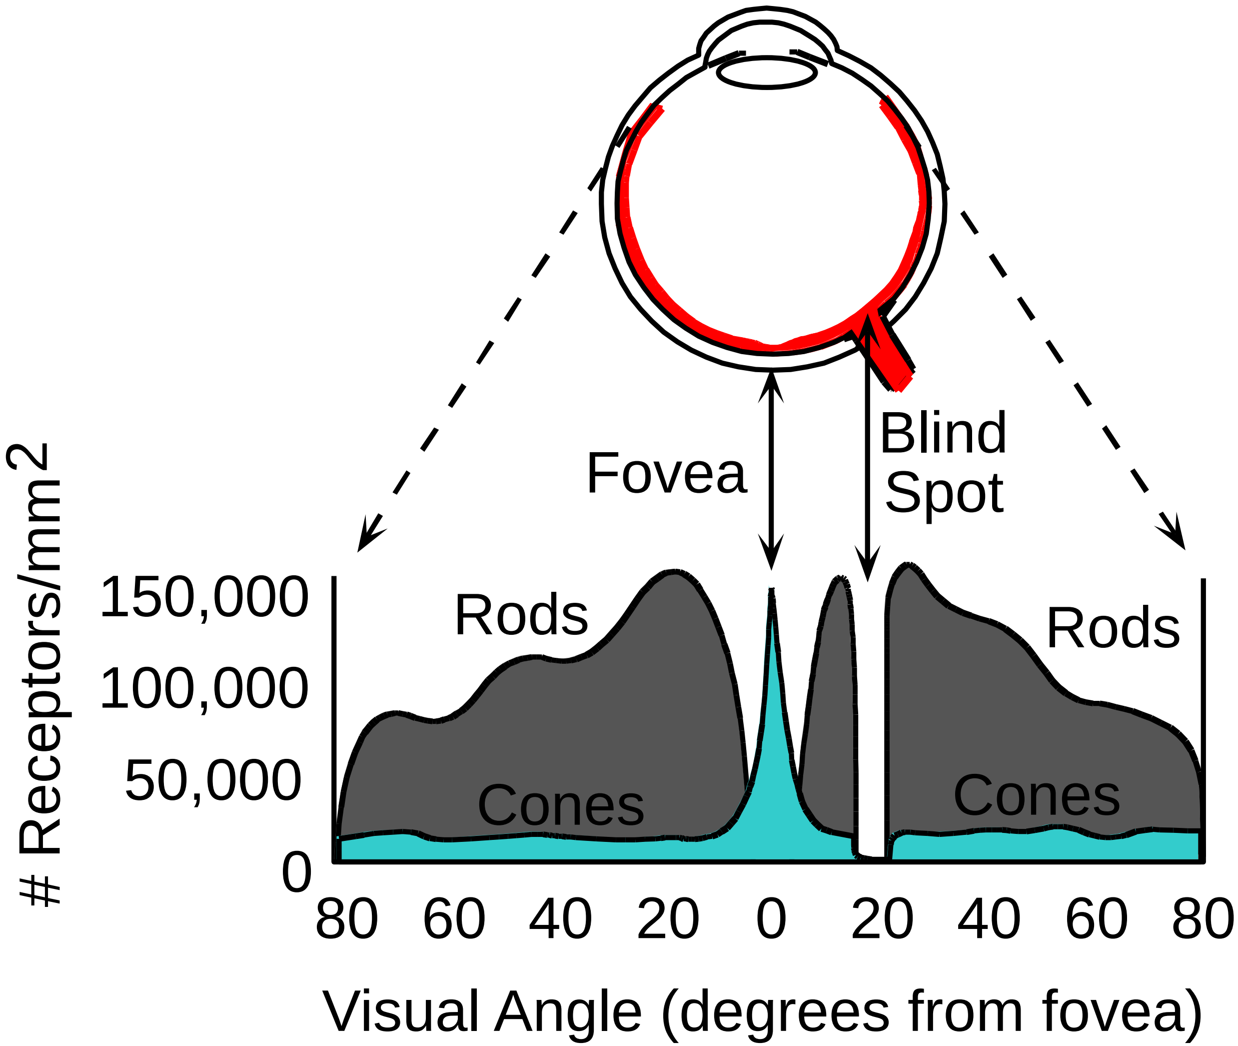

Night Sky: why are there more stars off-center?

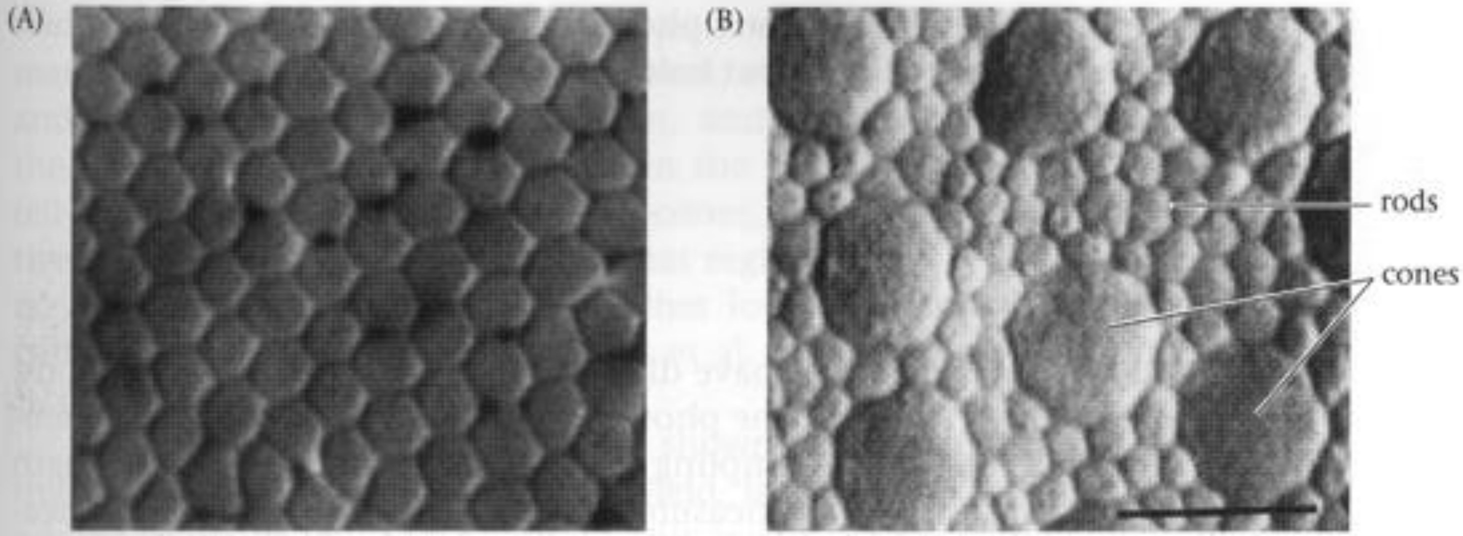

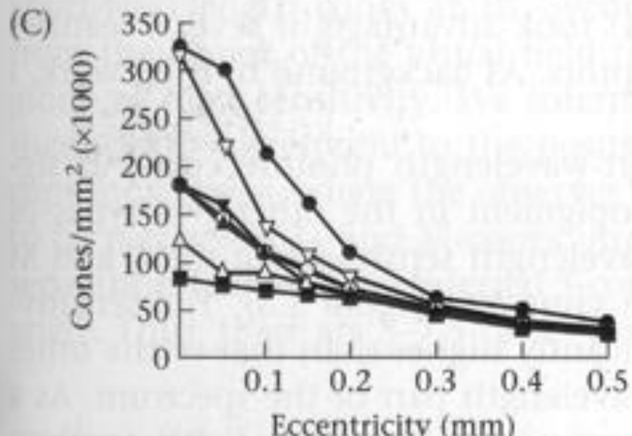

3.4 THE SPATIAL MOSAIC OF THE HUMAN CONES. Cross sections of the human retina at the level of the inner segments showing (A) cones in the fovea, and (B) cones in the periphery. Note the size difference (scale bar = 10 μm), and that, as the separation between cones grows, the rod receptors fill in the spaces. (C) Cone density plotted as a function of distance from the center of the fovea for seven human retinas; cone density decreases with distance from the fovea. Source: Curcio et al., 1990.

Foundations of Vision, by Brian Wandell, Sinauer Assoc., 1995

Electromagnetic Spectrum

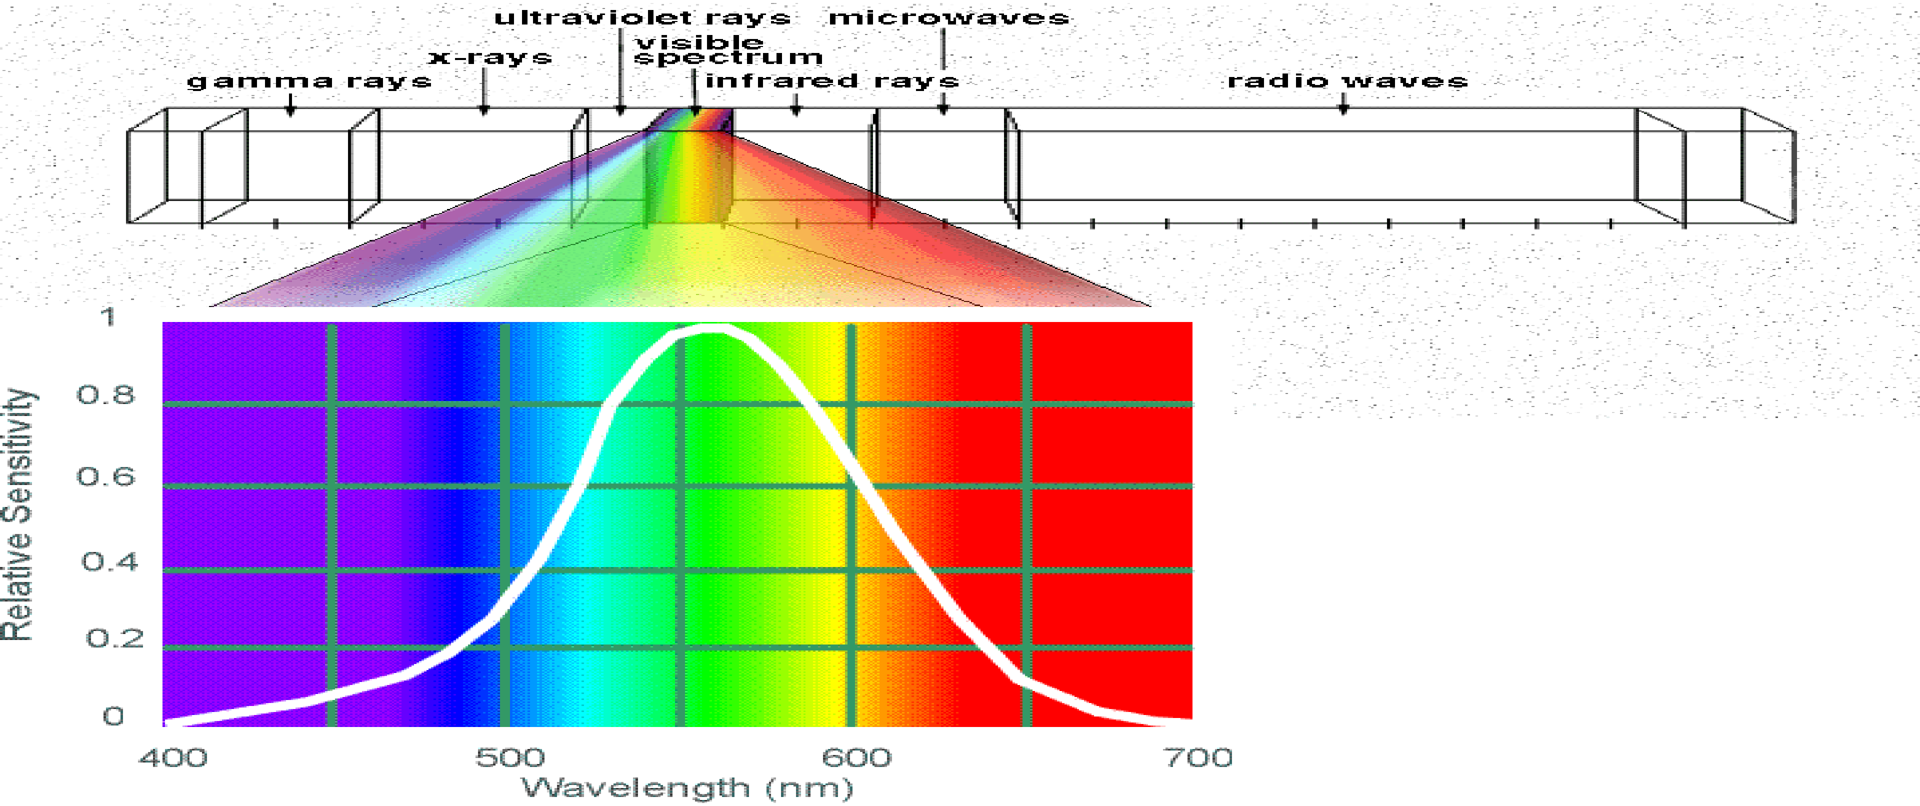

Human Luminance Sensitivity Function

http://www.yorku.ca/eye/photopik.htm

Visible Light

Why do we see light of these wavelengths?

…because that's where the Sun radiates EM energy

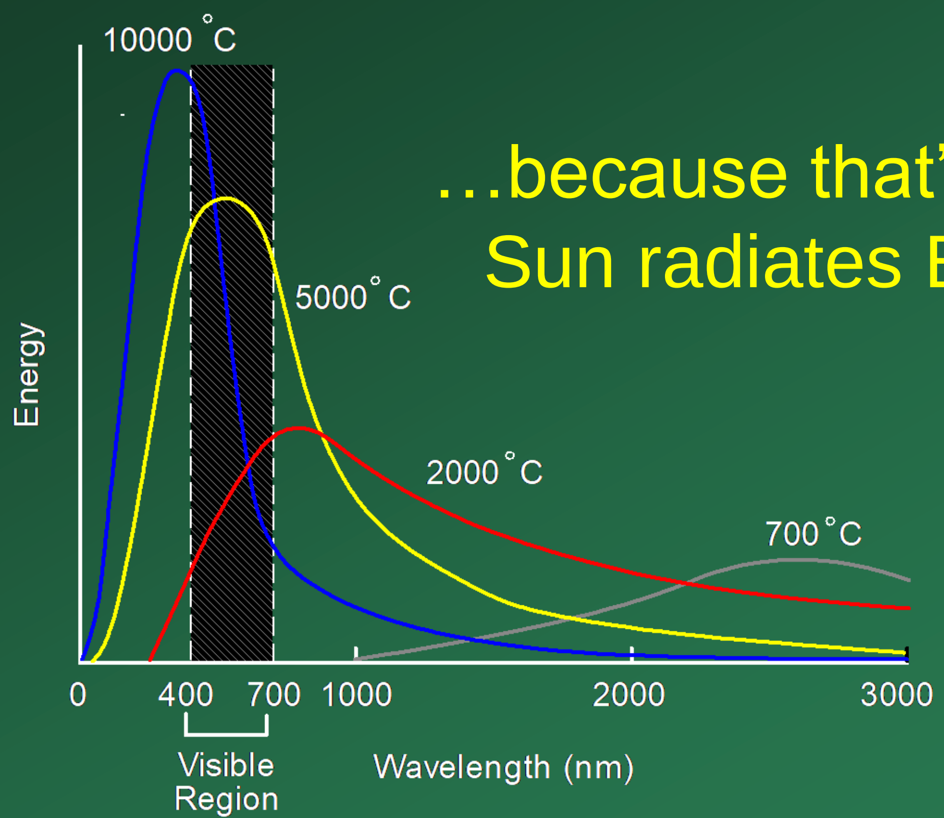

Curiouser and Curiouser

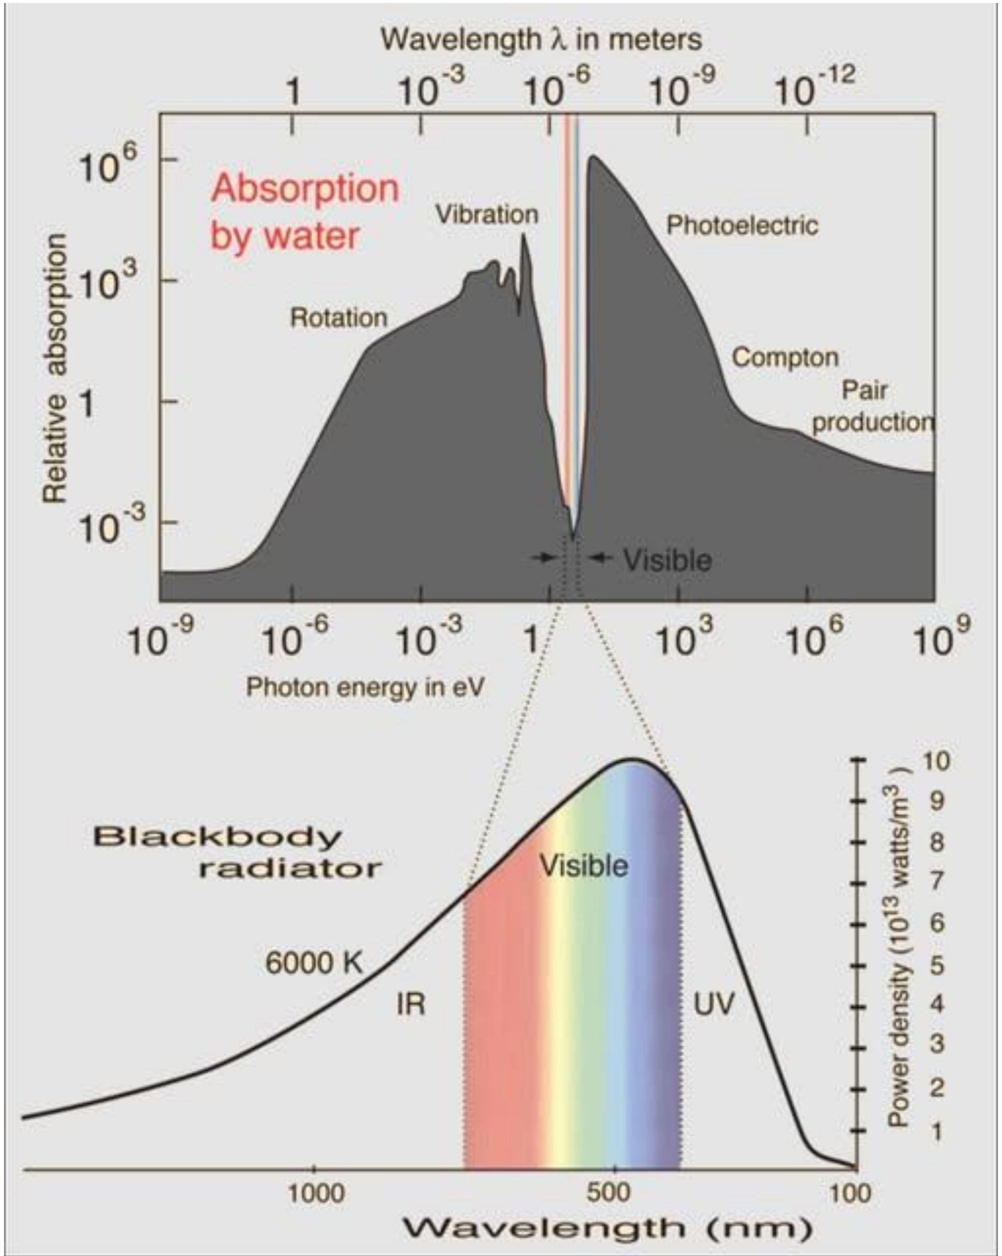

The Physics of Light

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.

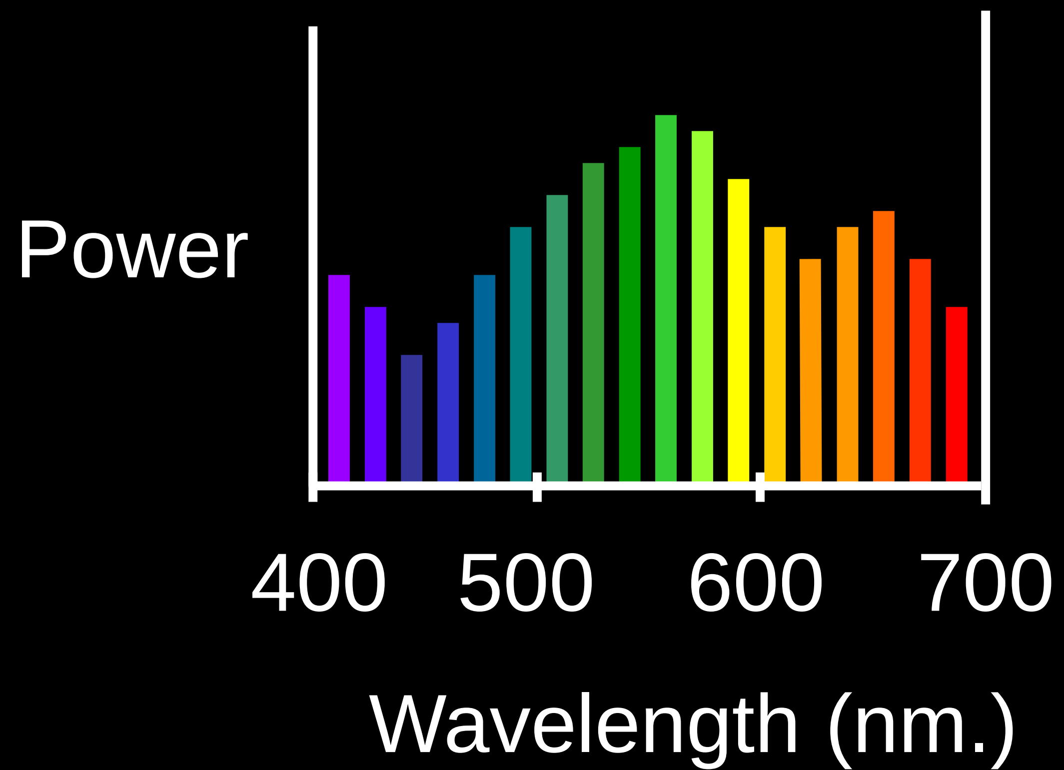

The Physics of Light

Some examples of the spectra of light sources

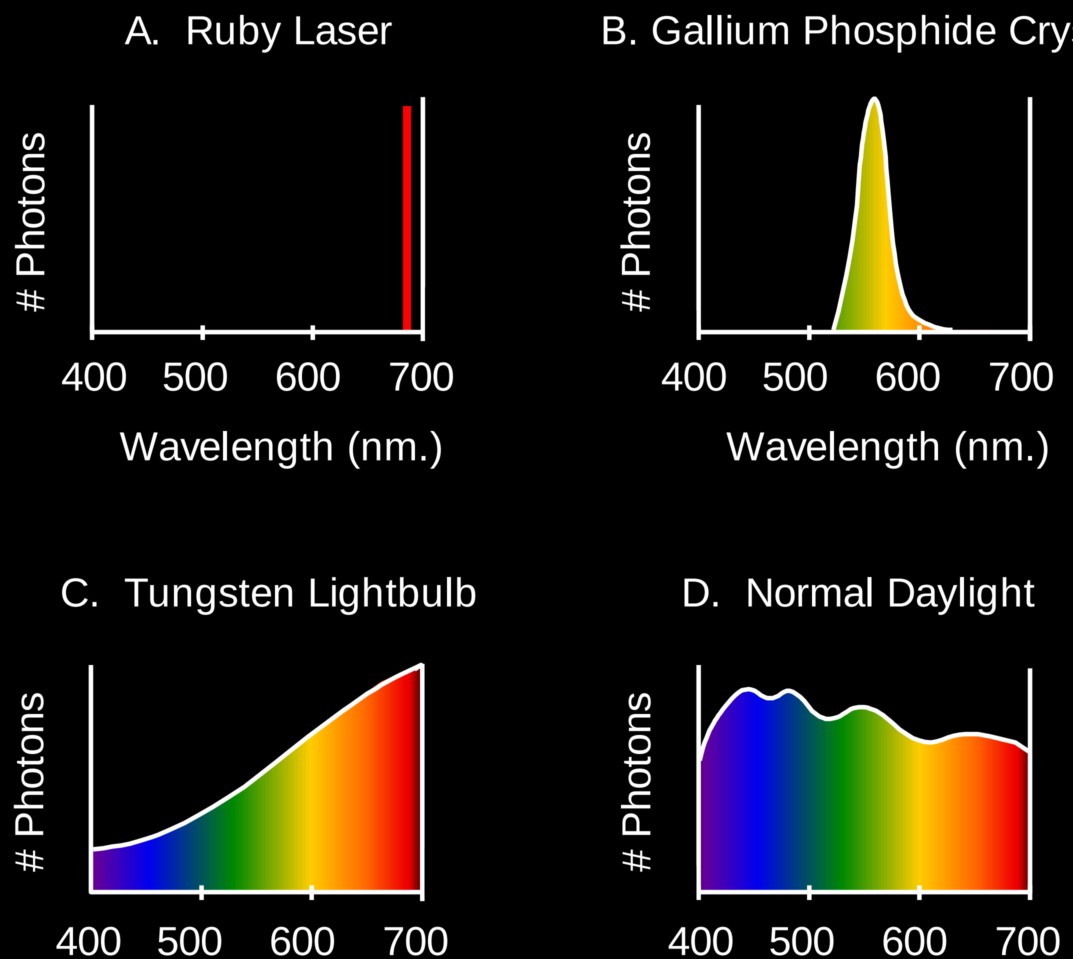

The Physics of Light

Some examples of the reflectance spectra of surfaces

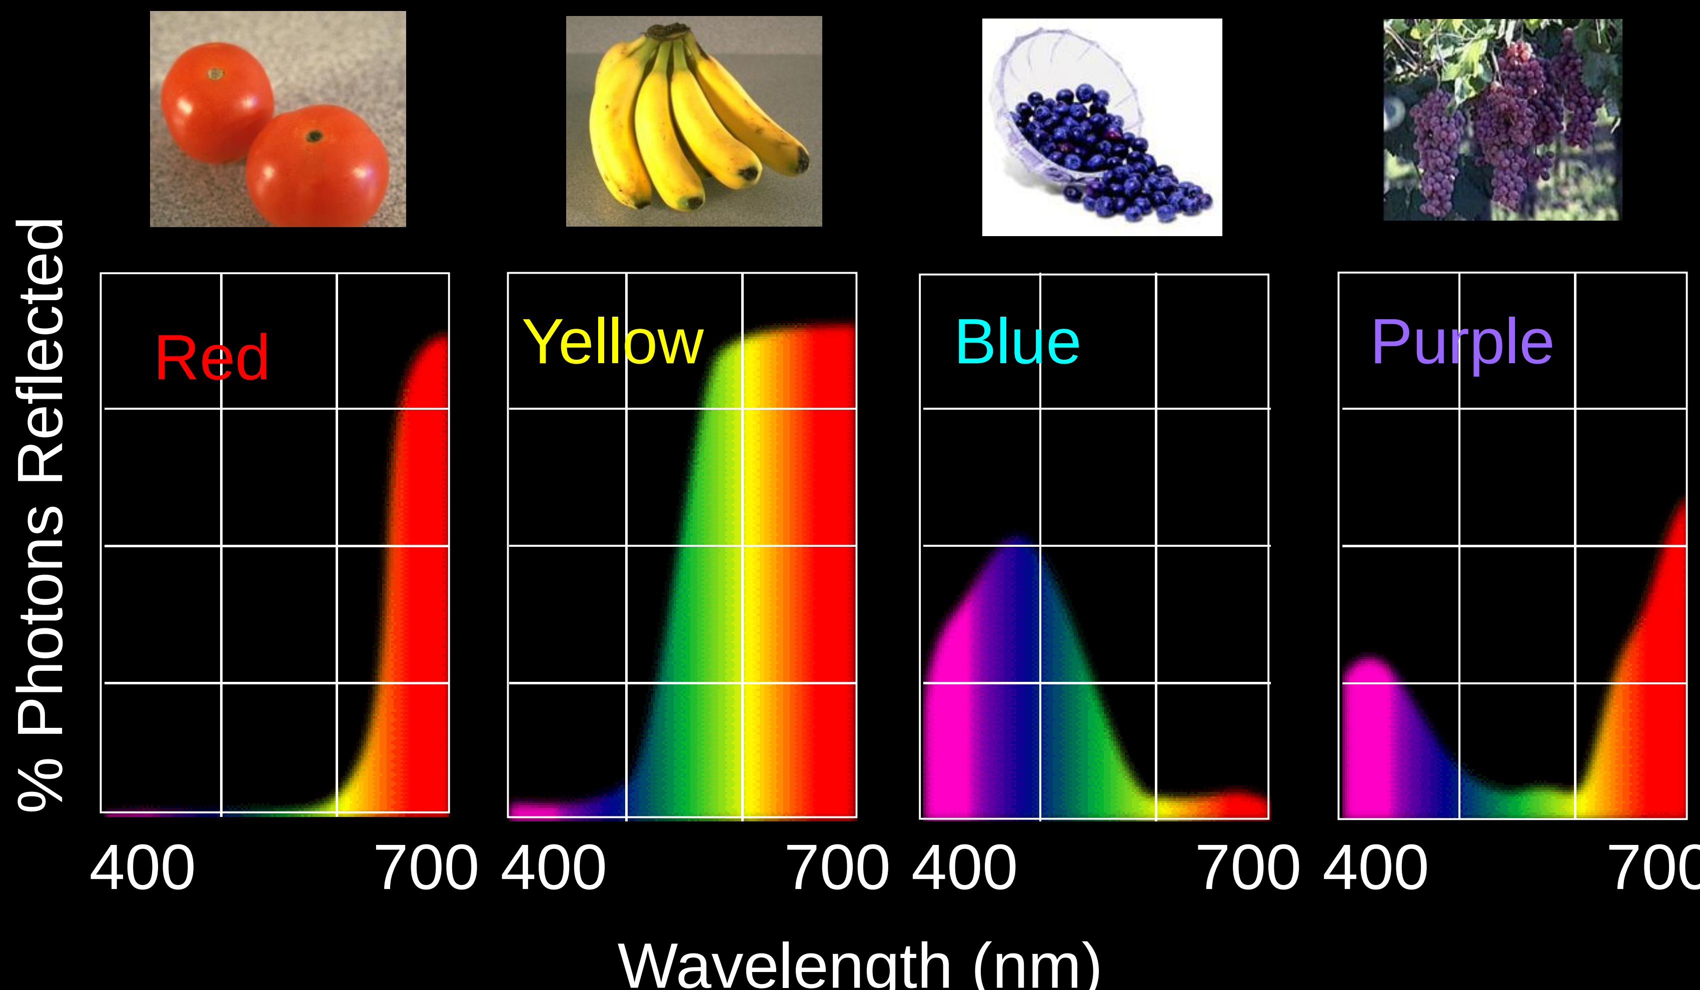

Physiology of Color Vision

Three kinds of cones:

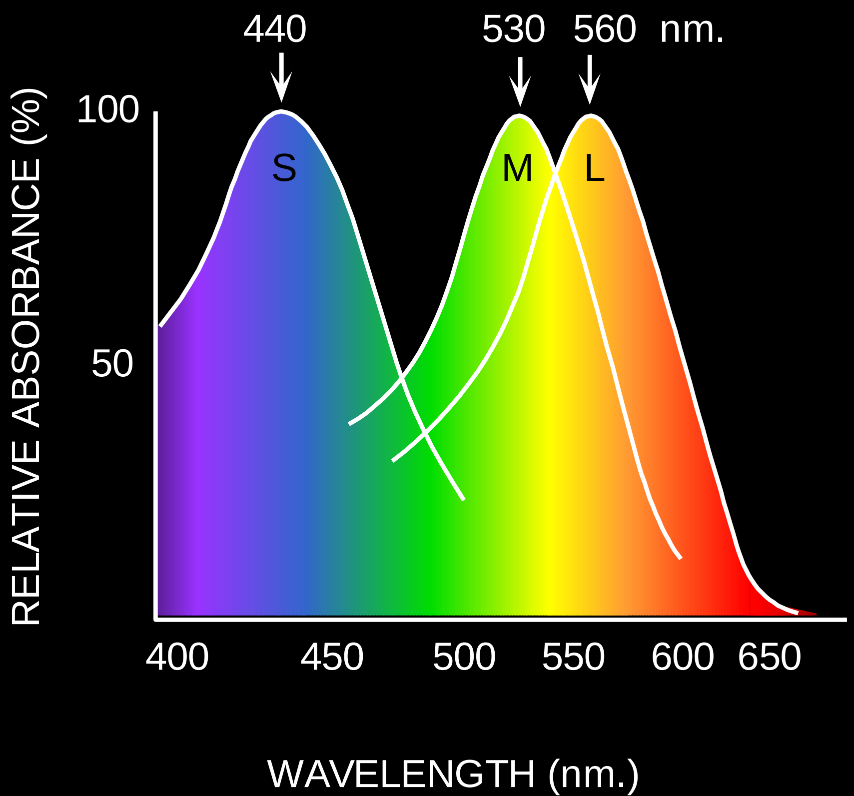

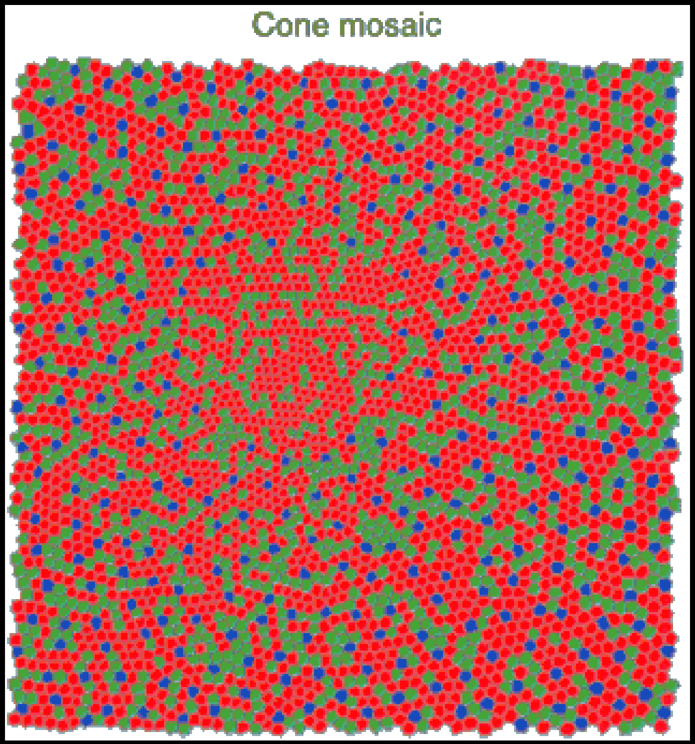

- What's up with S cones?
- Why are there 3?

Not everyone is trichromat

- Types of Dichromacy:
  - Deuteranopia: missing M cones
  - Protanopia: missing L cones
  - Tritanopia: missing S cones
- "M" and "L" on the X-chromosome
  - Why men (XY) are more likely to be color blind
  - "L" has high variation, so some women (XX) are tetrachromatic
- Some animals have
  - 1 (night animals)
  - 2 (e.g., dogs)
  - 4 (fish, birds)
  - 5 (pigeons, some reptiles/amphibians)
  - 12 (mantis shrimp)
- See Ren Ng's class for whole semester on this!

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

Color blindness is common in men

- 5-8% in males (XY)
- 0.5% to 1.0% in females (XX)

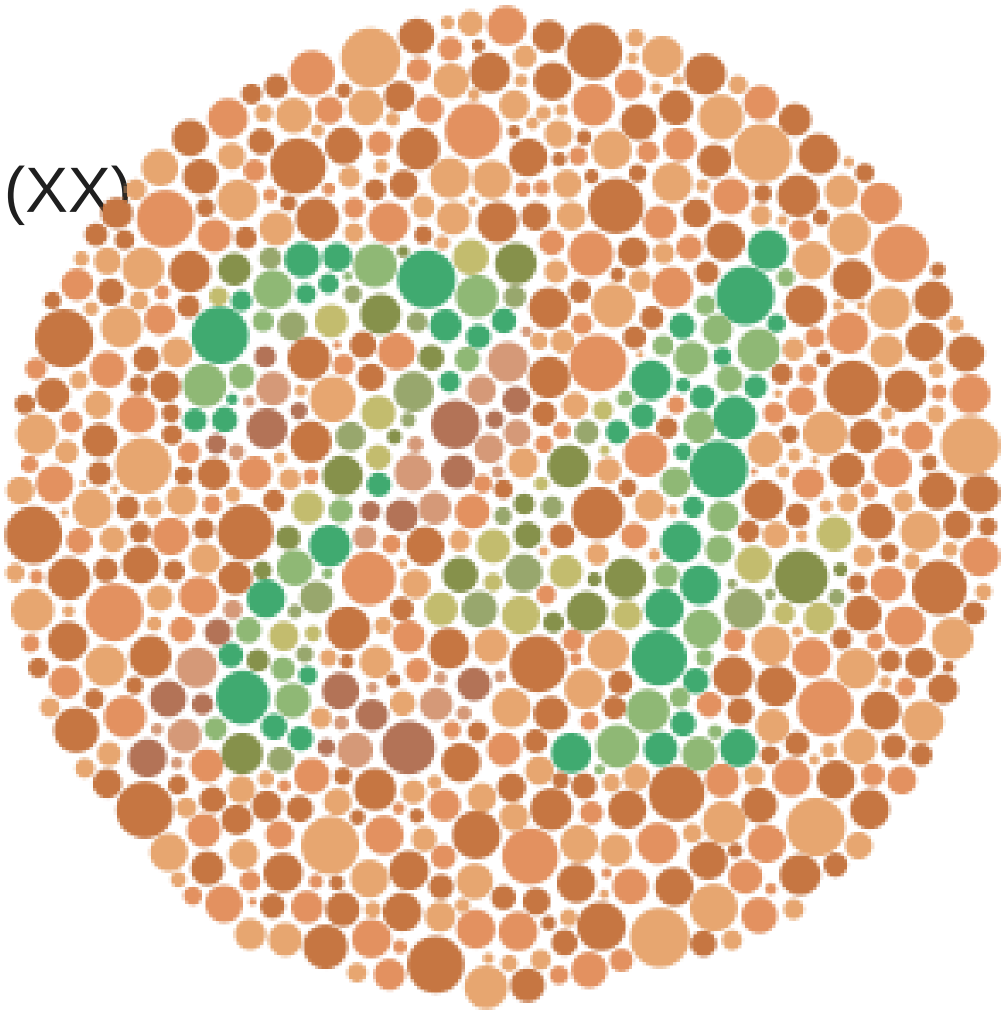

Ishihara color test plate. The number "74" should be clearly visible to viewers with normal color vision. Viewers with red-green color blindness will read it as "21", and viewers with monochromacy may see nothing.

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

Trichromacy

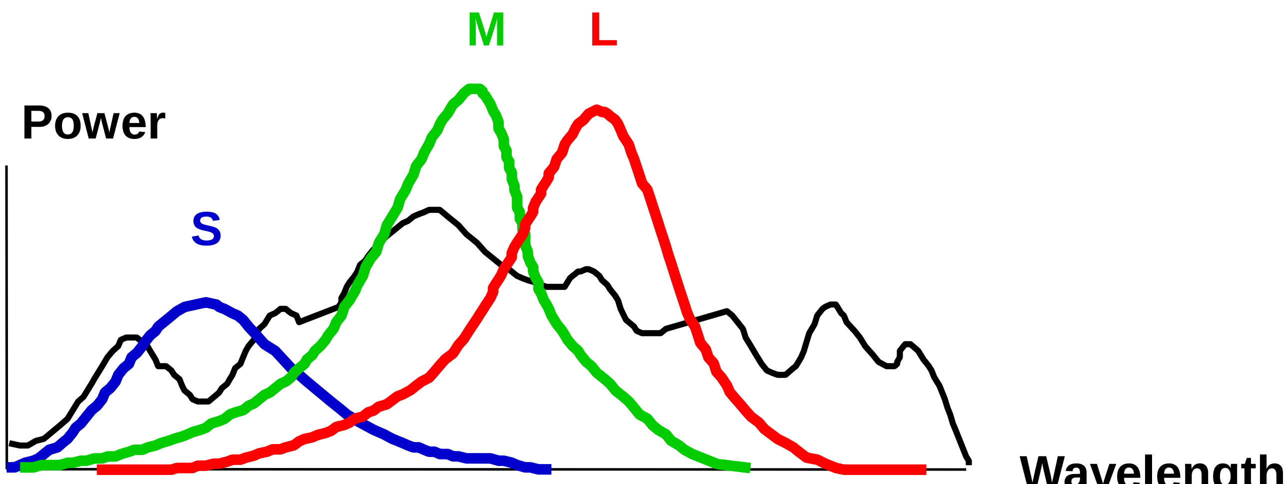

Rods and cones act as *filters* on the spectrum

- 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! A lot of the information is lost
  - As a result, two different spectra may appear indistinguishable
    - » such spectra are known as **metamers**

Slide by Steve Seitz

We are all color-blind!

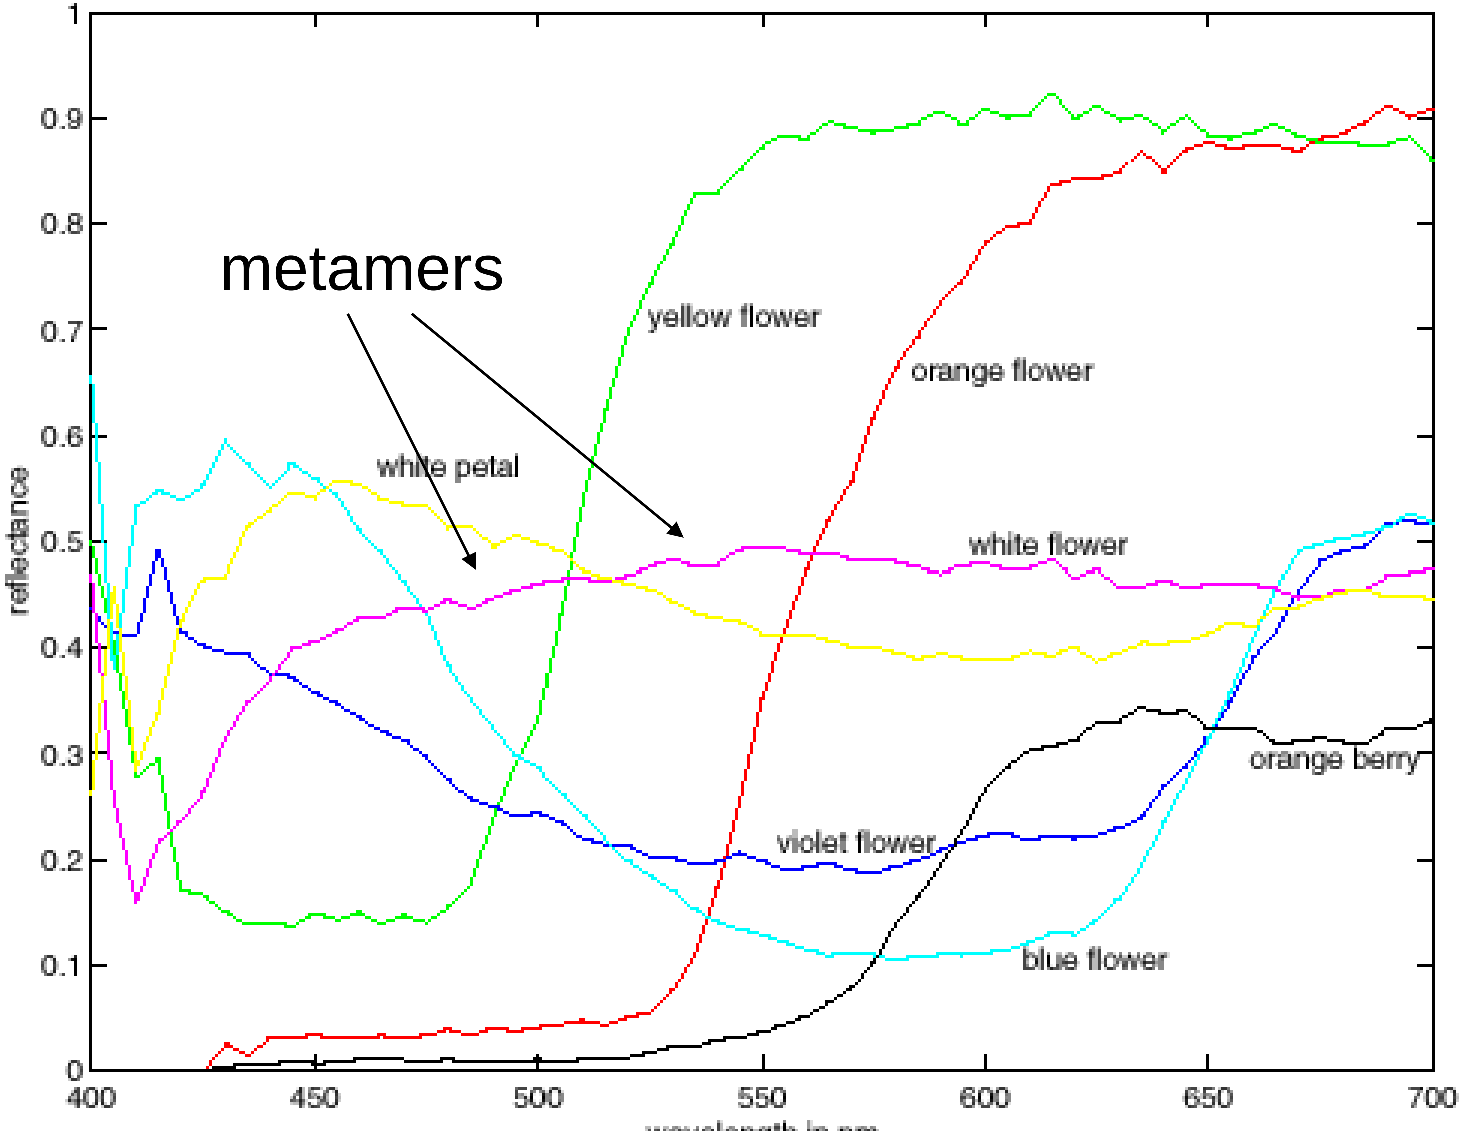

The Psychophysics of “Color”

There is no simple functional description for the perceived color of all lights under all viewing conditions, but …...

A helpful constraint:
Consider only physical spectra with normal distributions

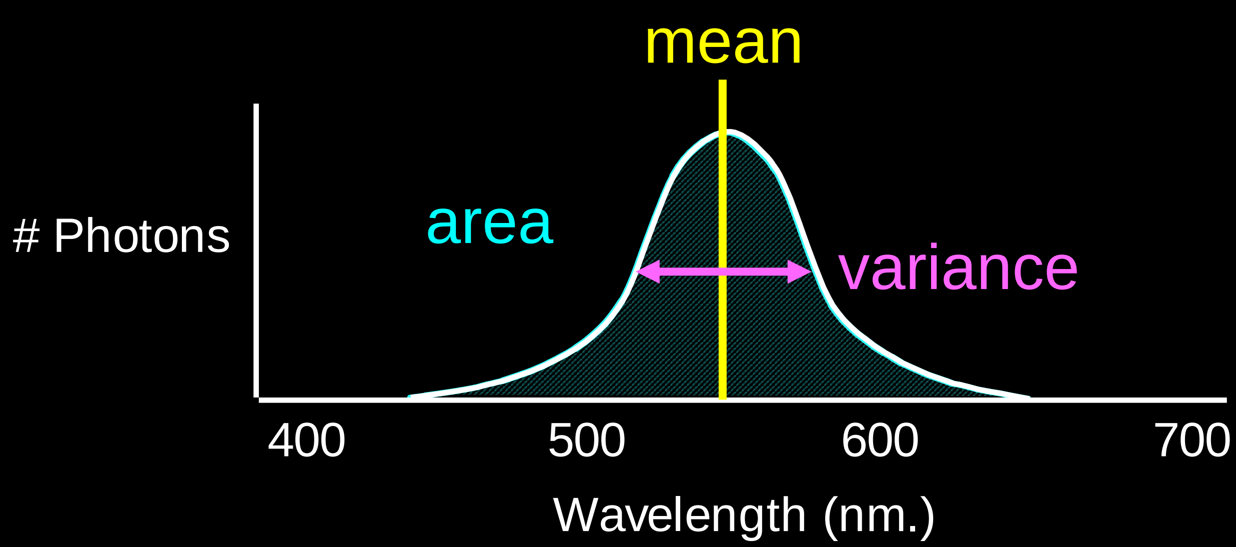

The Psychophysical Correspondence

Mean ⟷ Hue

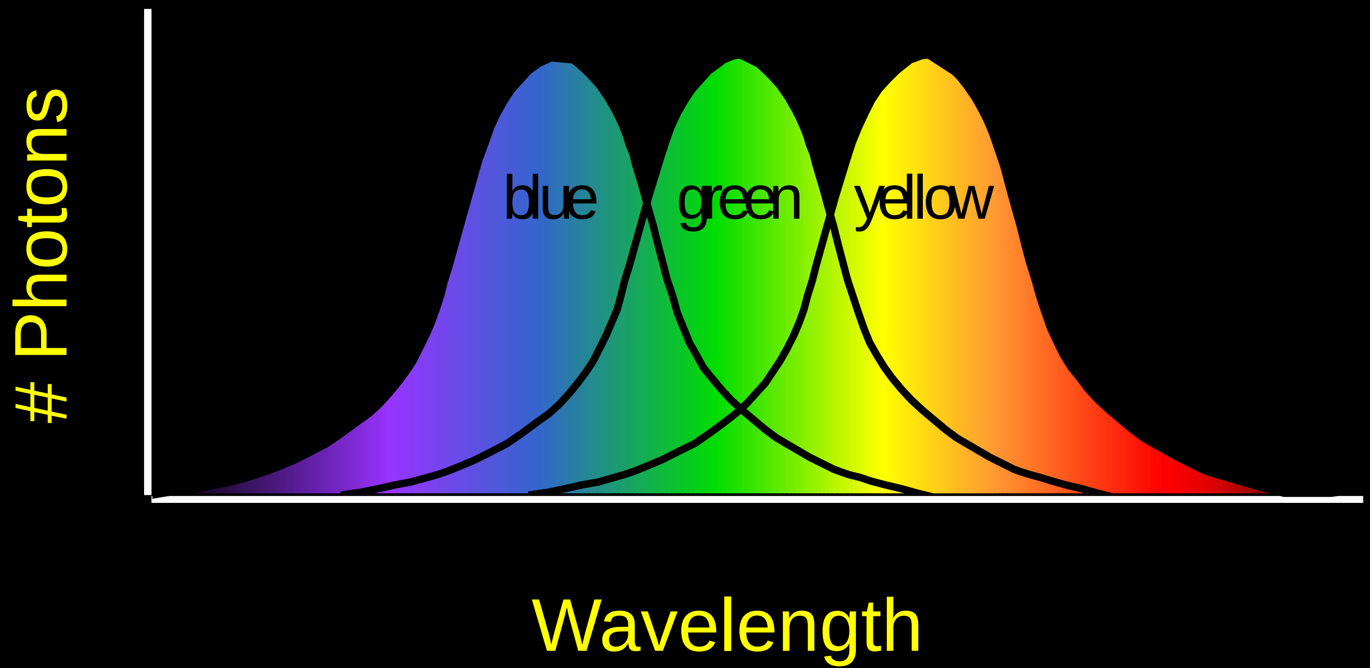

The Psychophysical Correspondence

Variance ⟷ Saturation

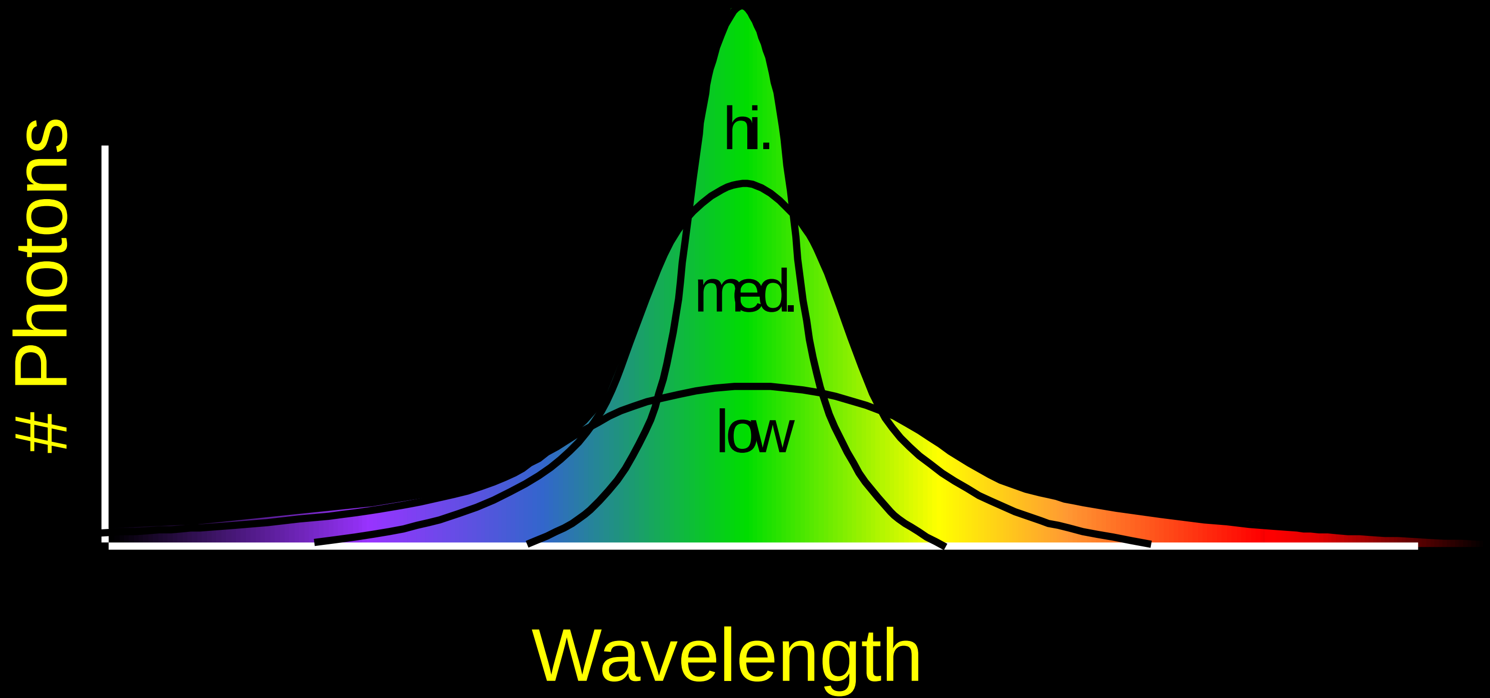

The Psychophysical Correspondence

Area ⟷ Brightness

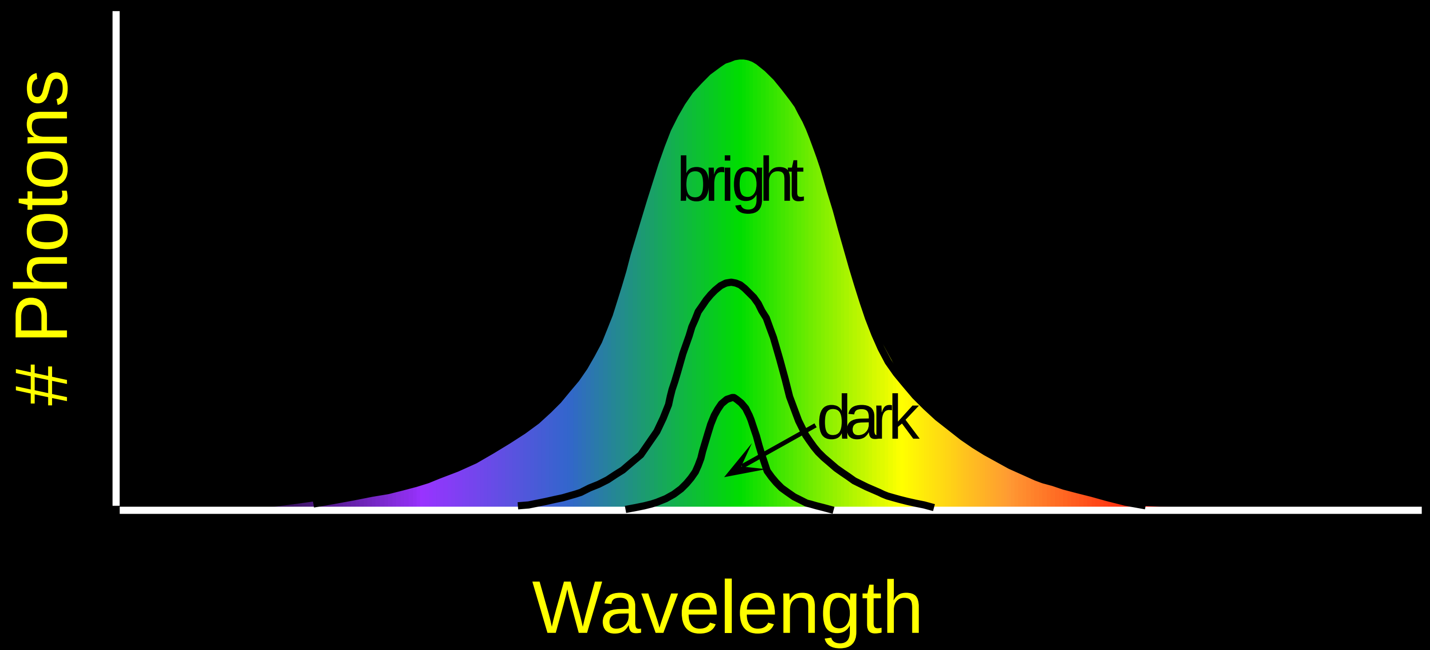

Color Constancy

<u>The "photometer metaphor" of color perception</u>:
Color perception is determined by the spectrum of light on each retinal receptor (as measured by a photometer).

Color Constancy

The "photometer metaphor" of color perception:
Color perception is determined by the spectrum of light on each retinal receptor (as measured by a photometer).

Color Constancy

<u>The "photometer metaphor" of color perception</u>:
Color perception is determined by the spectrum of light on each retinal receptor (as measured by a photometer).

What color is the "The Dress"?

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

Two Scene Interpretations of #thedress

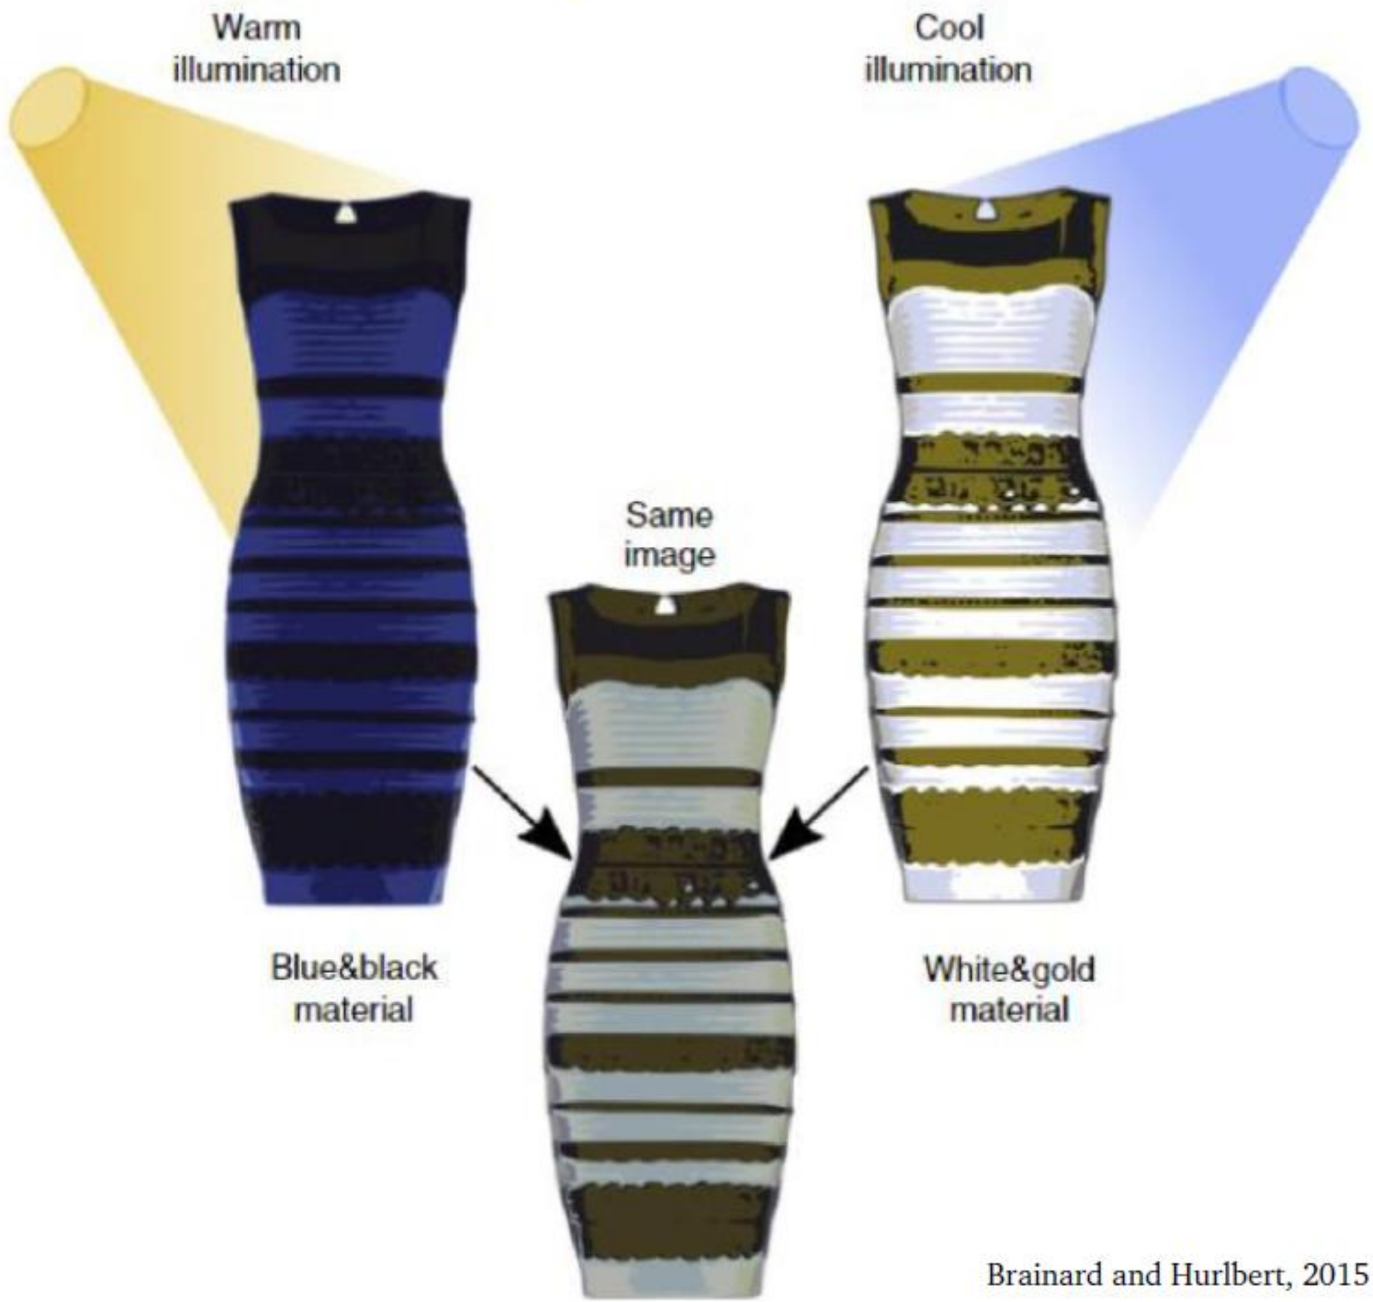

Brainard and Hurlbert, 2015

Color Constancy

~~Do we have constancy over _all_ global color transformations?~~

60% blue filter

Complete inversion

Color Constancy

Color Constancy: the ability to perceive the invariant color of a surface despite ecological Variations in the conditions of observation.

Another of these hard inverse problems:
Physics of light emission and surface reflection underdetermine perception of surface color

Camera White Balancing

- Manual
  - Choose color-neutral object in the photos and normalize
- Automatic (AWB)
  - Grey World: force average color of scene to grey
  - White World: force brightest object to white

Color Sensing in Camera (RGB)

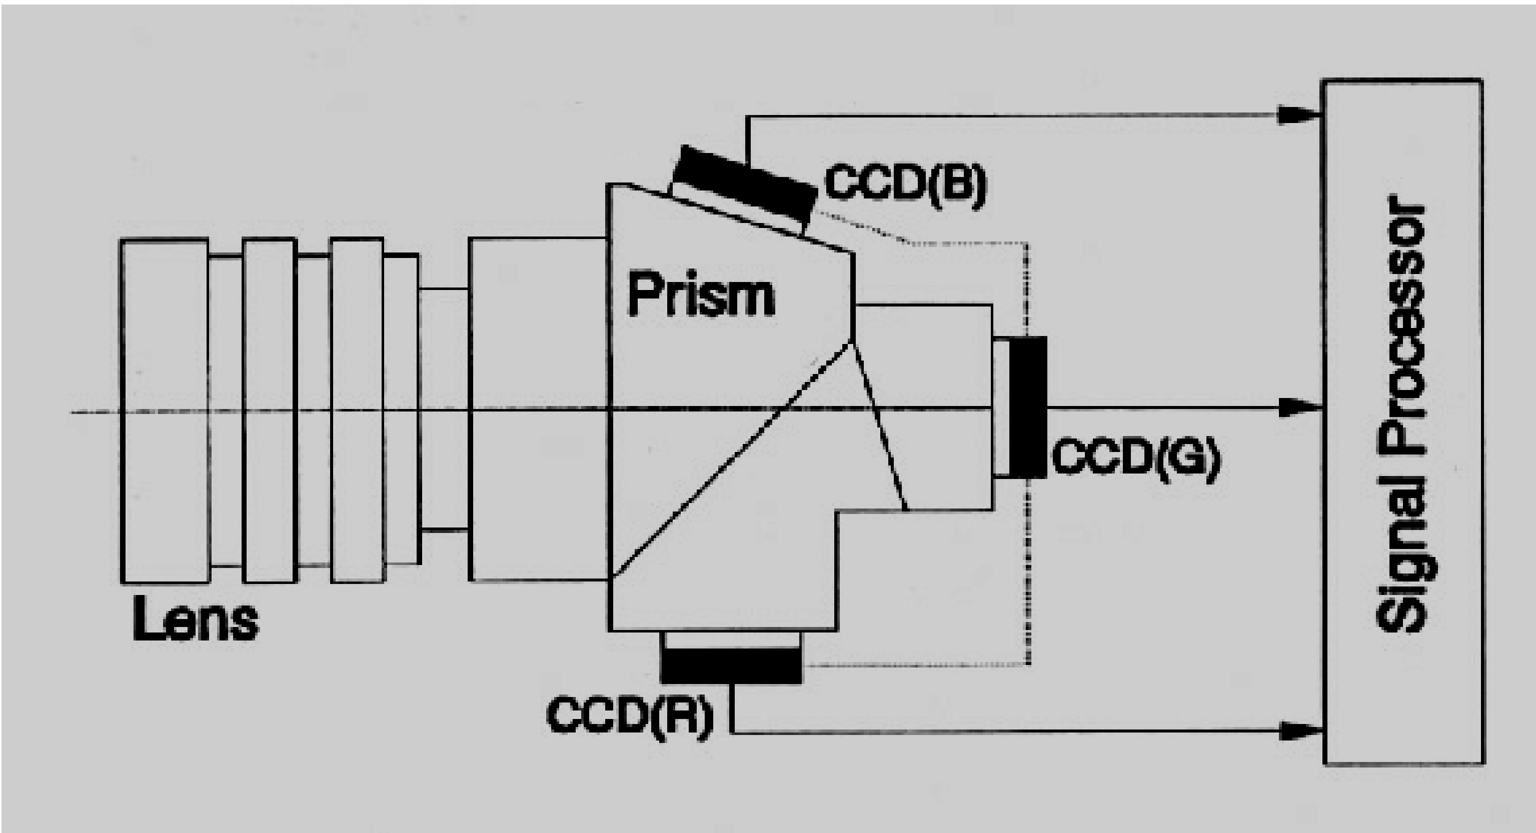

Why 3 colors?

Color Image

Image representation

- Images represented as a matrix
- Suppose we have a NxM RGB image called “im”
  - im(1,1,1) = top-left pixel value in R-channel
  - im(y, x, b) = y pixels down, x pixels to right in the b^{th} channel
  - im(N, M, 3) = bottom-right pixel in B-channel

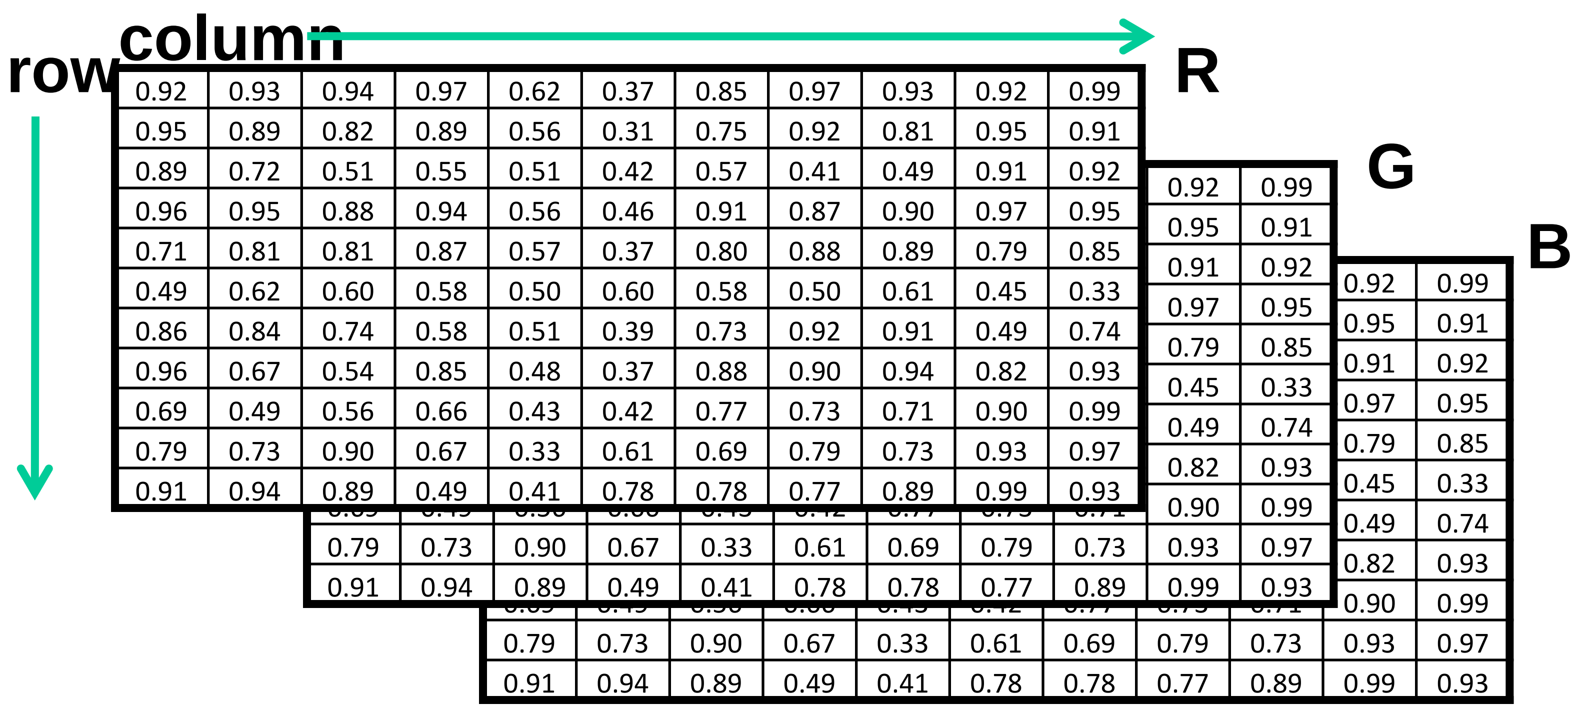

0.92	0.93	0.94	0.97	0.62	0.37	0.85	0.97	0.93	0.92	0.99
0.95	0.89	0.82	0.89	0.56	0.31	0.75	0.92	0.81	0.95	0.91
0.89	0.72	0.51	0.55	0.51	0.42	0.57	0.41	0.49	0.91	0.92
0.96	0.95	0.88	0.94	0.56	0.46	0.91	0.87	0.90	0.97	0.95
0.71	0.81	0.81	0.87	0.57	0.37	0.80	0.88	0.89	0.79	0.85
0.49	0.62	0.60	0.58	0.50	0.60	0.58	0.50	0.61	0.45	0.33
0.86	0.84	0.74	0.58	0.51	0.39	0.73	0.92	0.91	0.49	0.74
0.96	0.67	0.54	0.85	0.48	0.37	0.88	0.90	0.94	0.82	0.93
0.69	0.49	0.56	0.66	0.43	0.42	0.77	0.73	0.71	0.90	0.99
0.79	0.73	0.90	0.67	0.33	0.61	0.69	0.79	0.73	0.93	0.97
0.91	0.94	0.89	0.49	0.41	0.78	0.78	0.77	0.89	0.99	0.93

More realistic color camera

3-chip vs. 1-chip: quality vs. cost

Why more green?

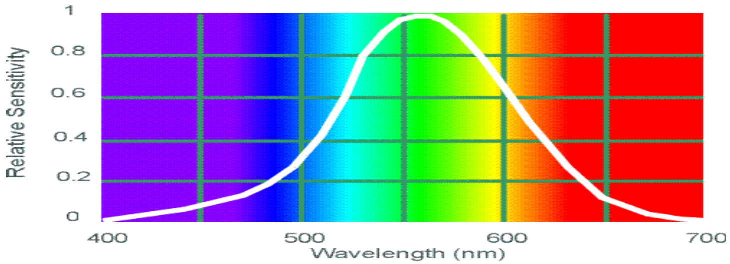

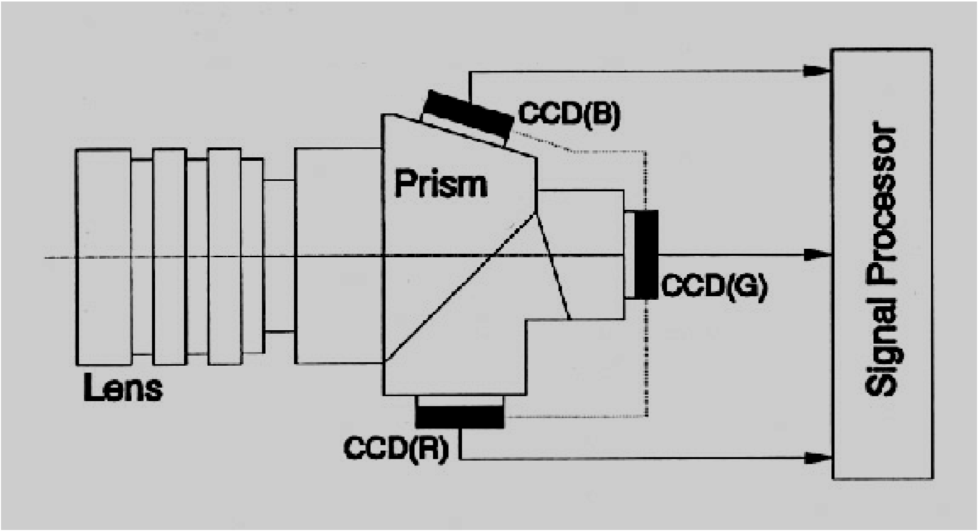

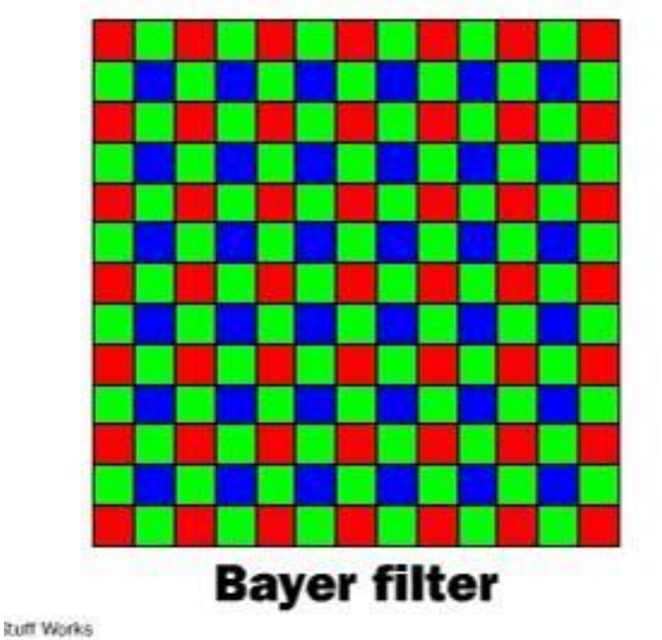

http://www.cooldictionary.com/words/Bayer-filter.wikipedia

Slide by Steve Seitz

Practical Color Sensing: Bayer Grid

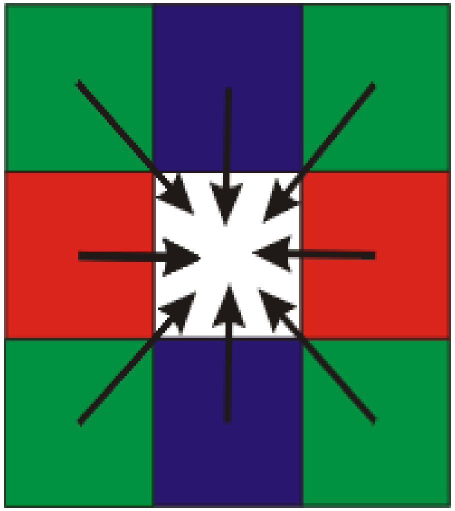

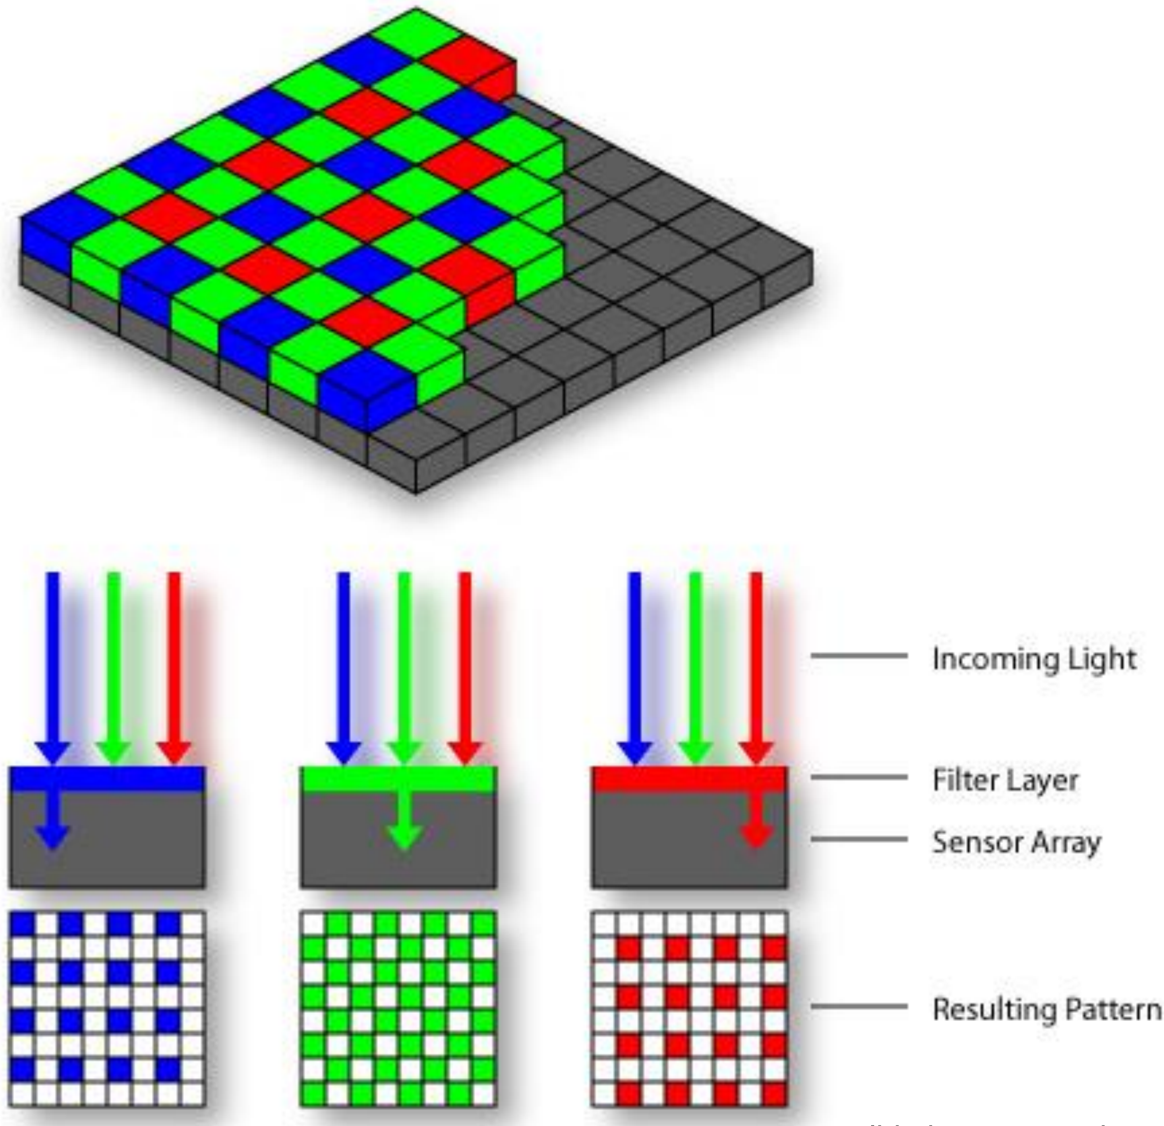

Estimate RGB
at 'G' cels from neighboring values

http://www.cooldictionary.com/words/Bayer-filter.wikipedia

Slide by Steve Seitz

Color spaces

How can we represent color?

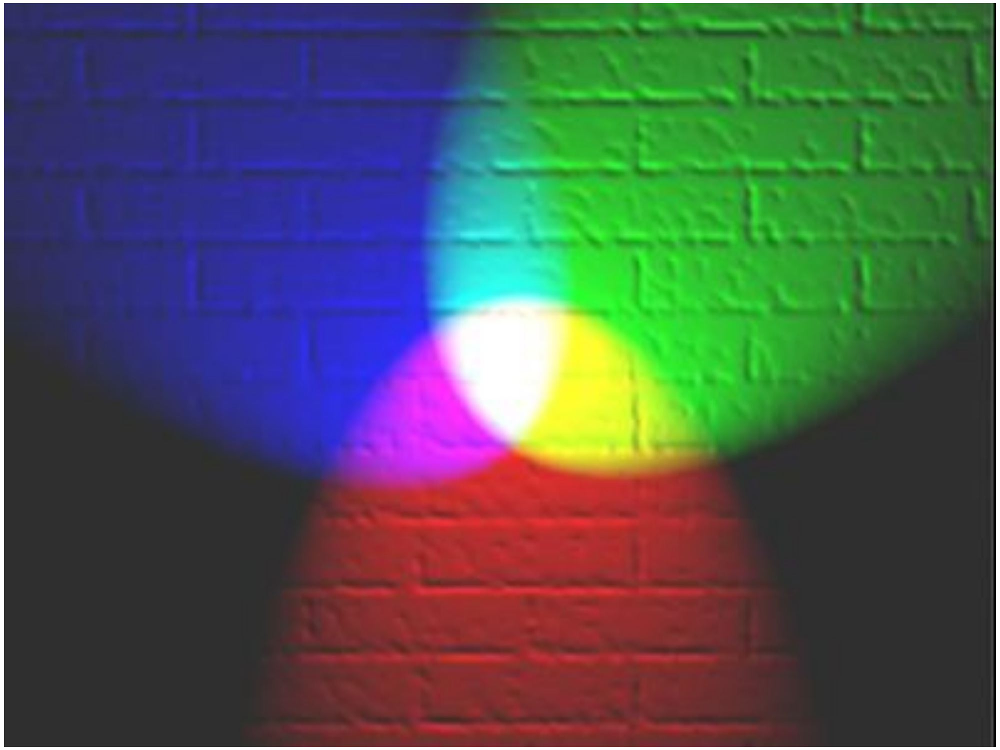

http://en.wikipedia.org/wiki/File:RGB_illumination.jpg

Color spaces: RGB vs. CMY(K)

Light projection vs paint

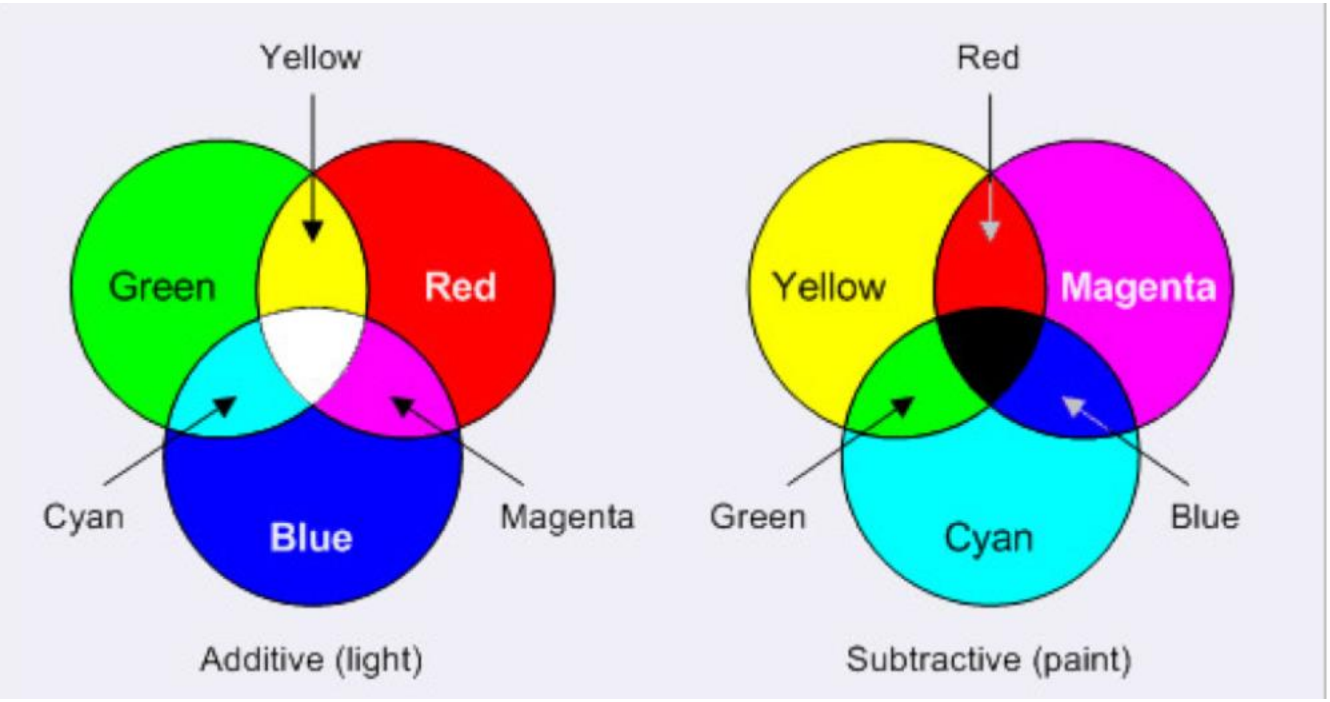

Source: https://intranet.mcad.edu/kb/cmyk-vs-rgb-what-color-space-should-i-work

Color spaces: RGB

Default color space

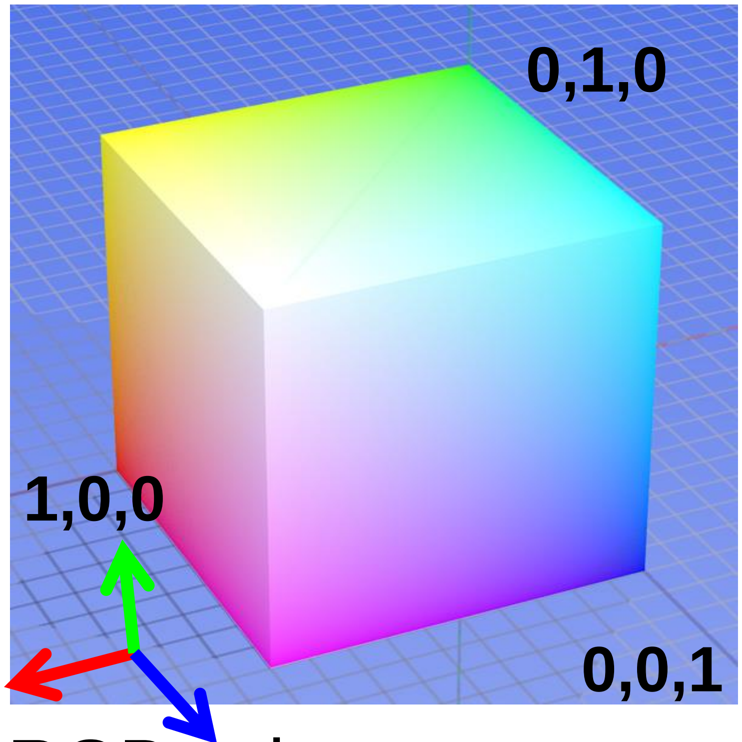

RGB cube

- Easy for devices
- But not perceptual
- Where do the grays live?
- Where is hue and saturation?

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

HSV

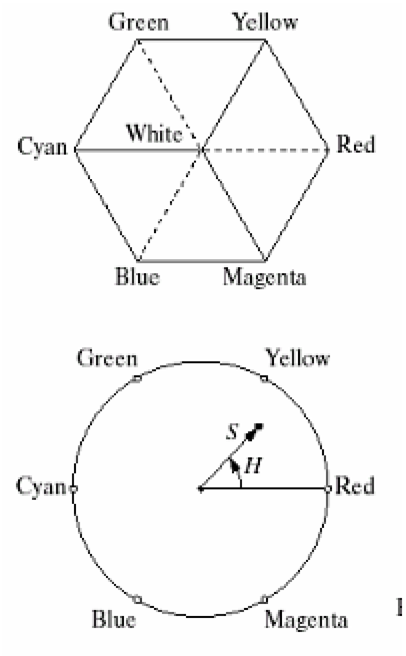

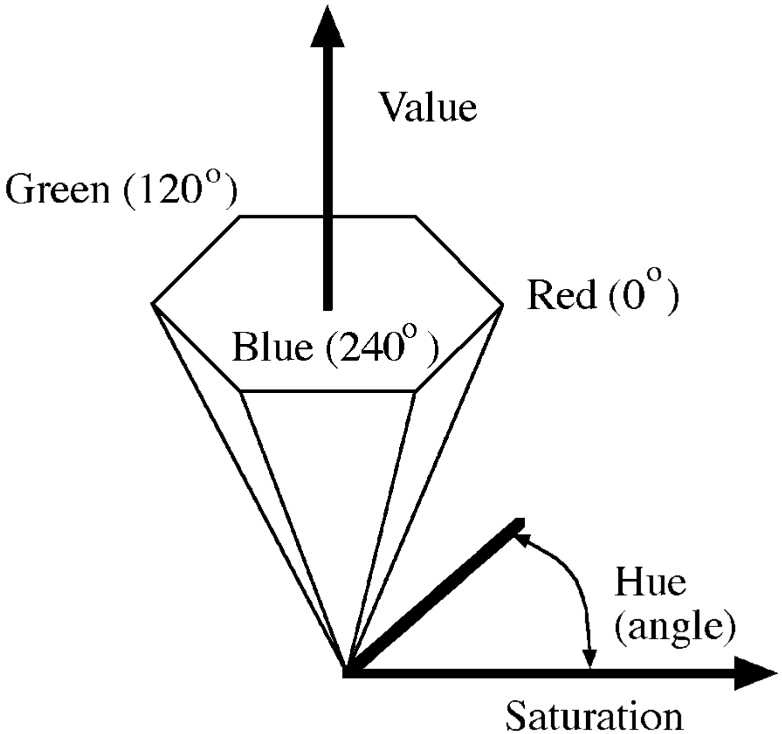

Hue, Saturation, Value (Intensity)

- RGB cube on its vertex

Decouples the three components (a bit)

Use rgb2hsv() and hsv2rgb()

Slide by Steve Seitz

Color spaces: HSV

Intuitive color space

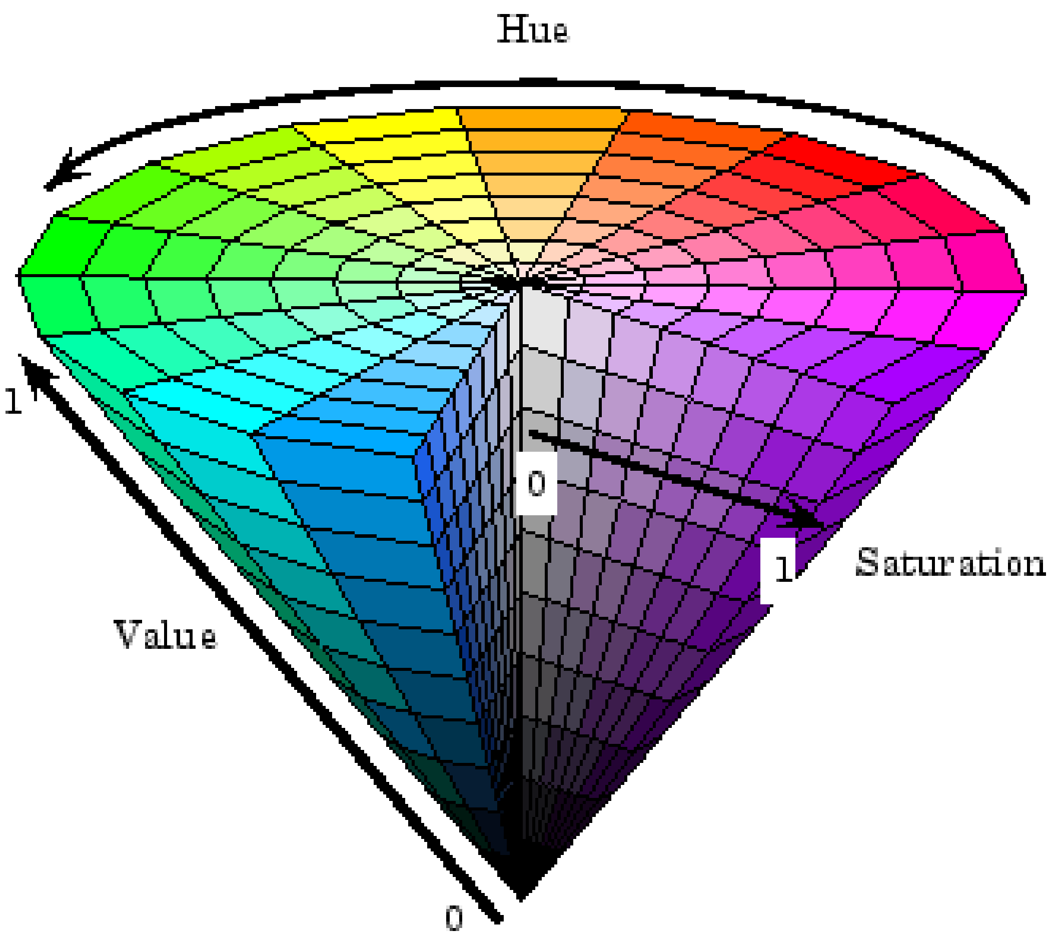

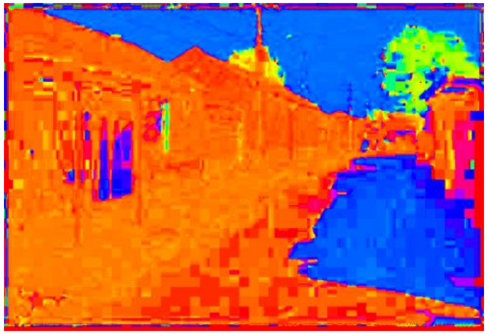

H
(S=1,V=1)

S
(H=1,V=1)

V
(H=1,S=0)

Color spaces: L*a*b*

“Perceptually uniform”* color space

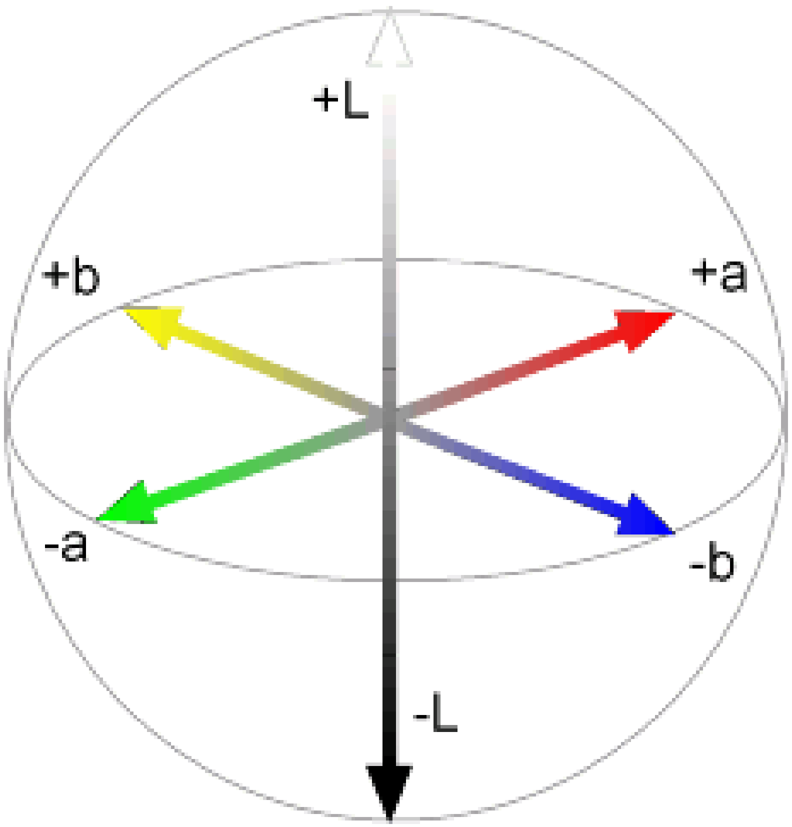

L
(a=0,b=0)

a
(L=65,b=0)

b
(L=65,a=0)

Programming Project #1

Prokudin-Gorskii's Color Photography (1907)

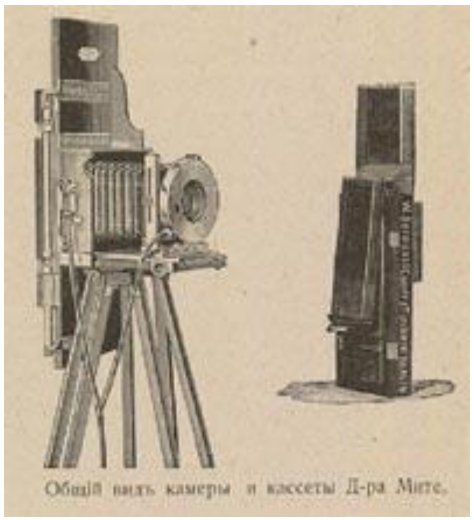

Programming Project #1

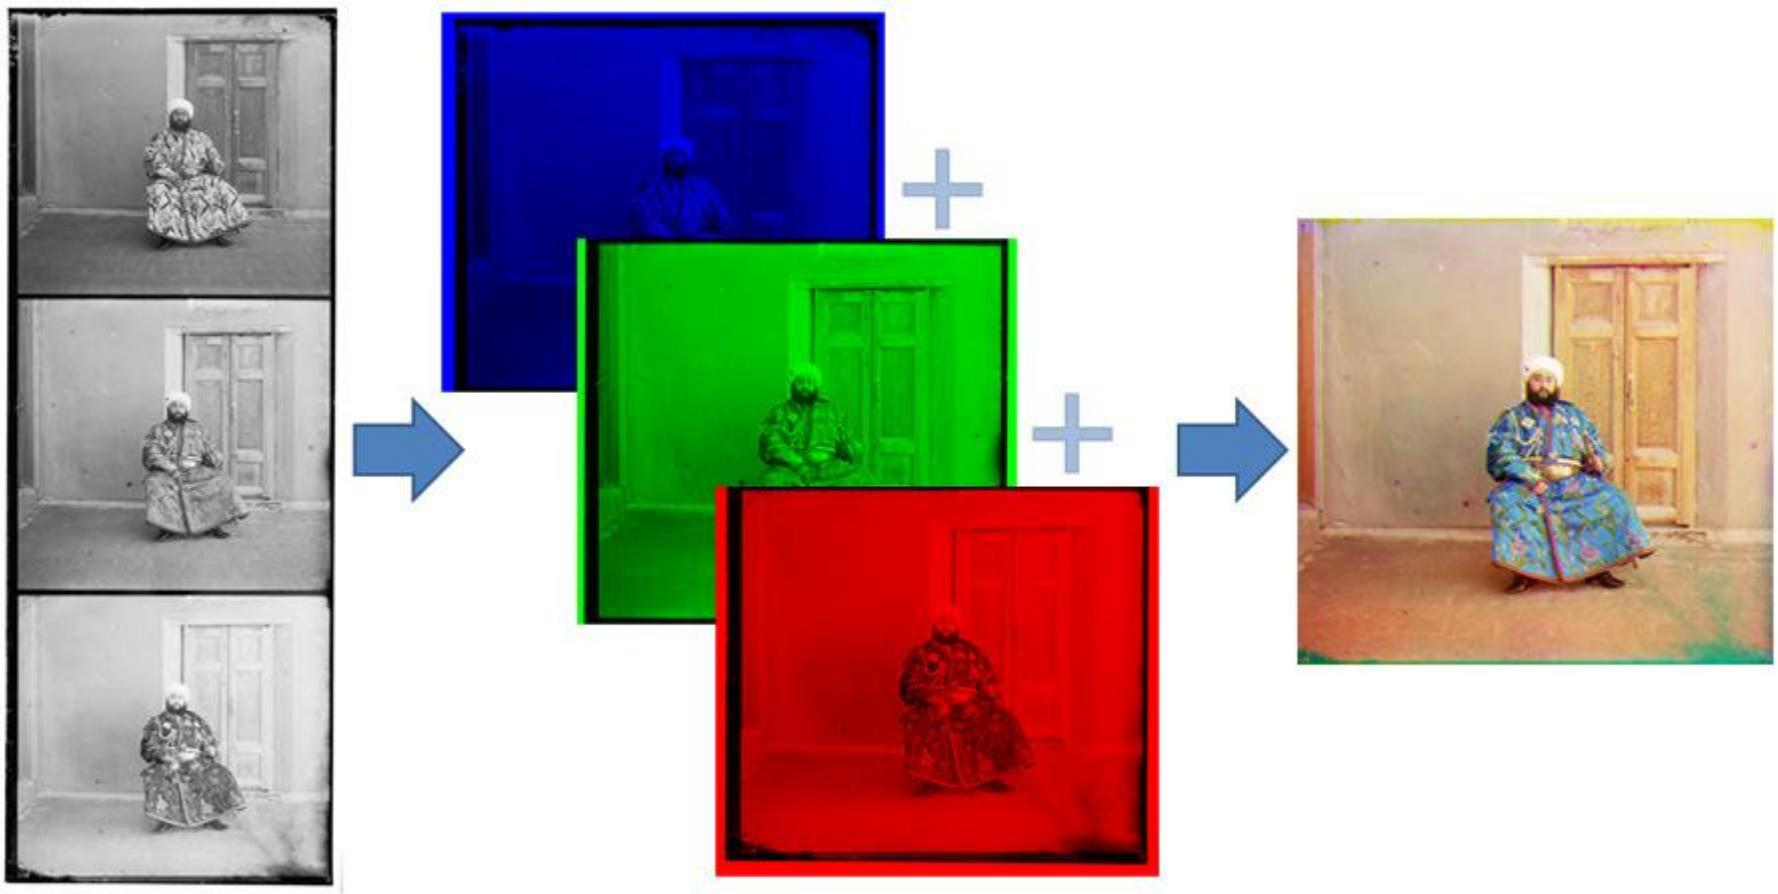

Programming Project #1

- How to compare R,G,B channels?
- No right answer
  - L2 norm / Euclidian Distance (sqrt dropped):

$$ssd(u,v) = \sum_{(x,y)\in N} \left[I(u+x, v+y) - P(x,y)\right]^2$$

  - Normalized Correlation (NCC):

$$ncc(u,v) = \frac{\sum_{(x,y)\in N} \left[I(u+x, v+y) - \bar{I}\right]\left[P(x,y) - \overline{P}\right]}{\sqrt{\sum_{(x,y)\in N} \left[I(u+x, v+y) - \bar{I}\right]^2 \sum_{(x,y)\in N} \left[P(x,y) - \overline{P}\right]^2}}$$