1950

Future Vision

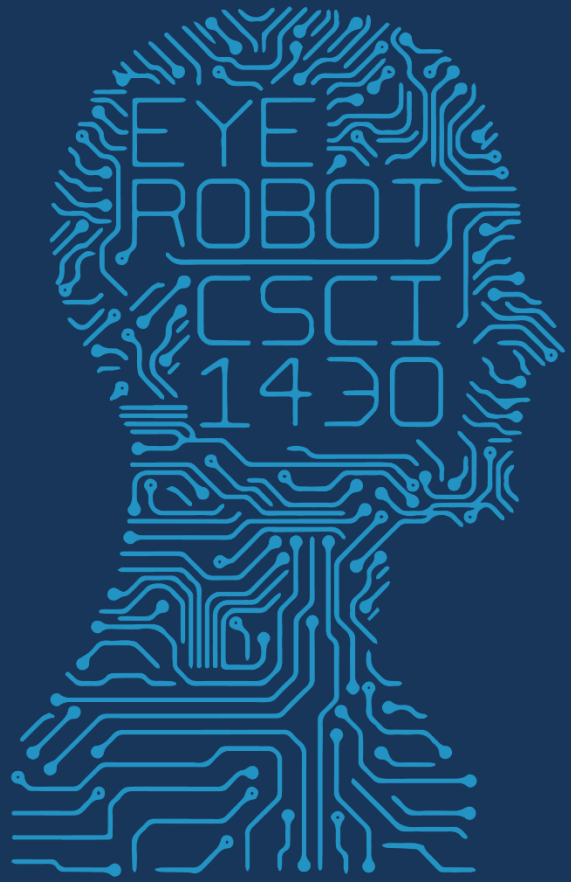

2020

Computer Vision

Bela Borsodi

Bela Borsodi

Recap so far

- What is an image?
  - Sampling and aliasing
  - Thinking in frequency
- Image filtering
  - Kernels and their responses
  - Kernel orientation
  - Kernel sizes and scale spaces
- Local features
  - Detection by finding corners via peaks
  - Description by hand-coding local ‘texture’
  - Matching robustly

Elephant in room:

COLOR

Grayscale intensity

Color

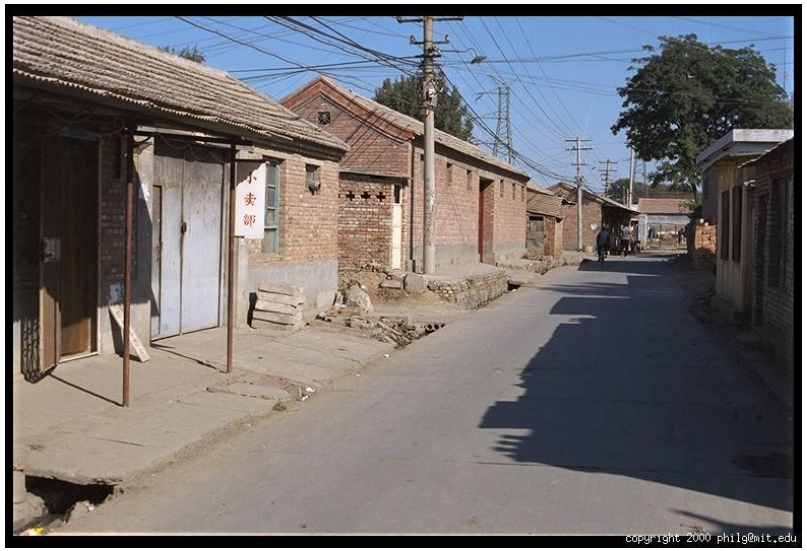

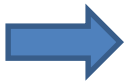

James Hays

Images in Python Numpy

N x M RGB image “im”

- im[0,0,0] = top-left pixel value in R-channel
- Im[x, y, b] = x pixels to right, y pixels down in the b^{th} channel
- Im[N-1, M-1, 3] = bottom-right pixel in B-channel

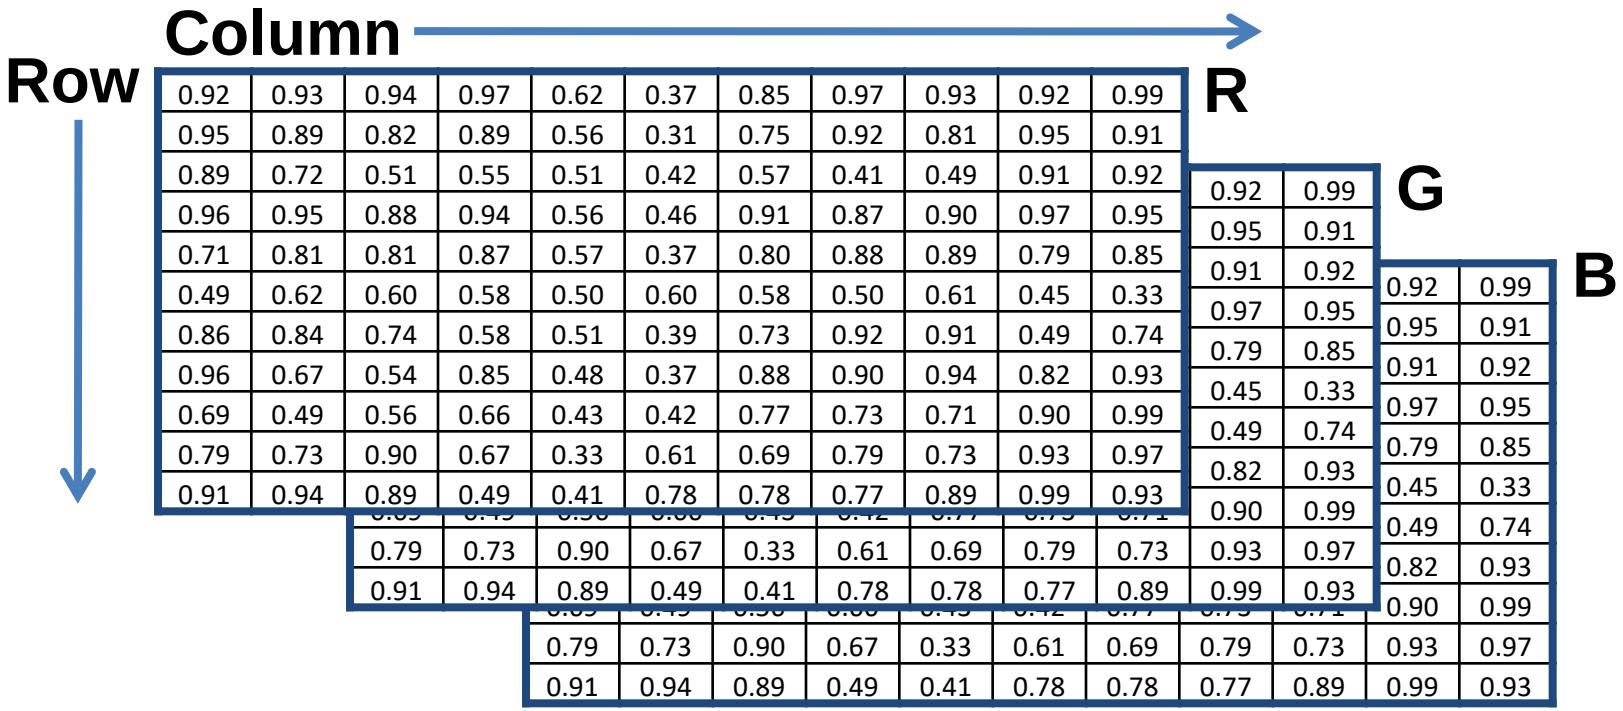

But what is color?

ANATOMY

The Eye

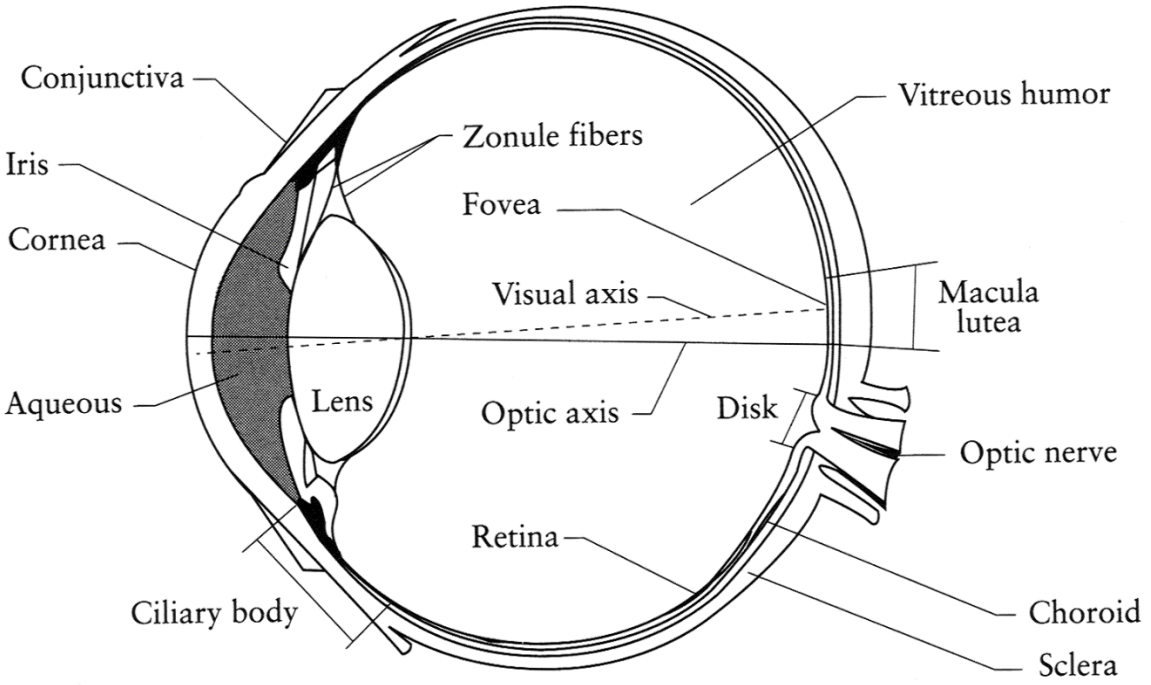

- The human eye is a camera
  - **Iris** - colored annulus with radial muscles
  - **Pupil** - the hole (aperture) whose size is controlled by the iris
  - What's the sensor?
    - photoreceptor cells (rods and cones) in the **retina**

Slide by Steve Seitz

The Retina

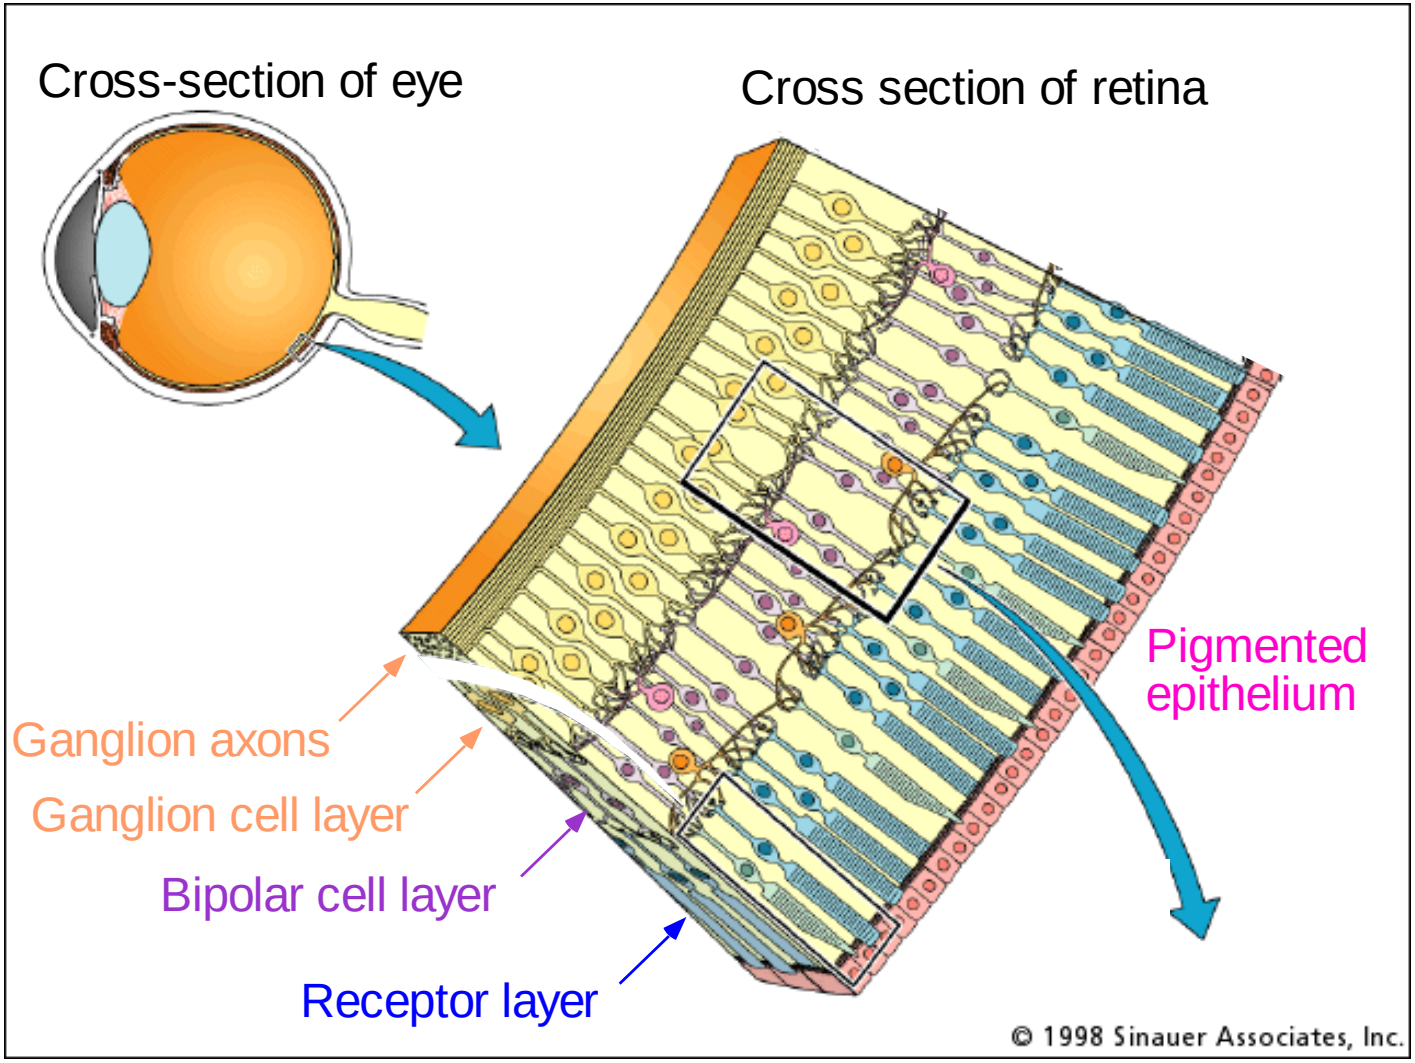

Axon = nerve fibre

Ganglion cell = a neuron (nerve cell)

Bipolar cell = pass ‘graded signal changes’

Wait, the blood vessels are in front of the photoreceptors??

https://www.youtube.com/watch?v=L_W-IXqoxHA

What humans don’t have: tapetum lucidum

Human eyes can reflect a tiny bit and blood in the retina makes this reflection red.

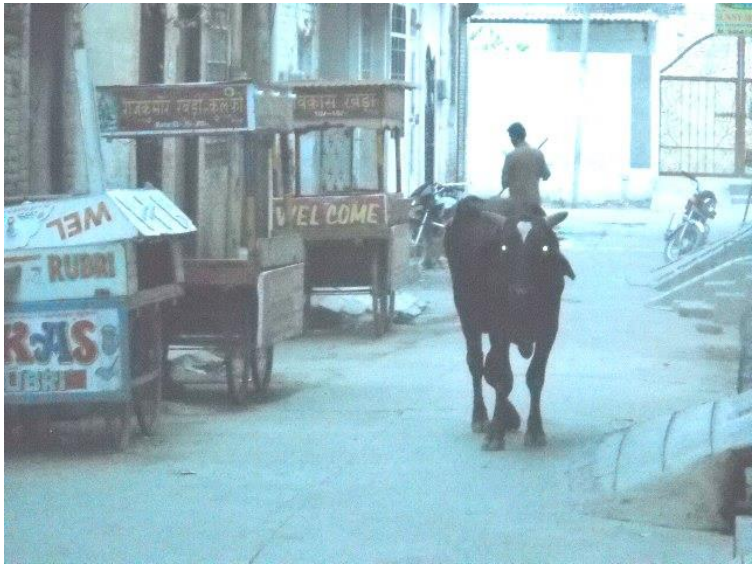

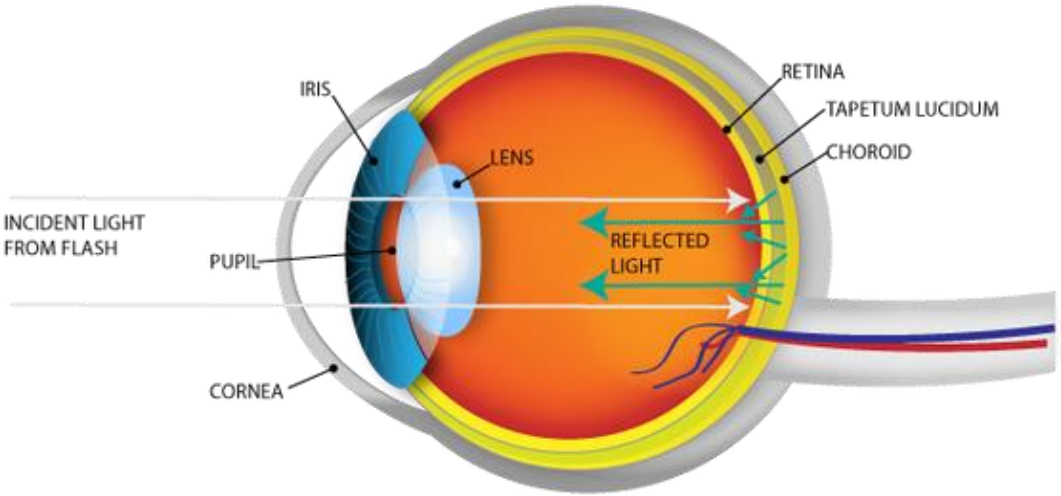

Tapetum lucidum exposed (cow eye)

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

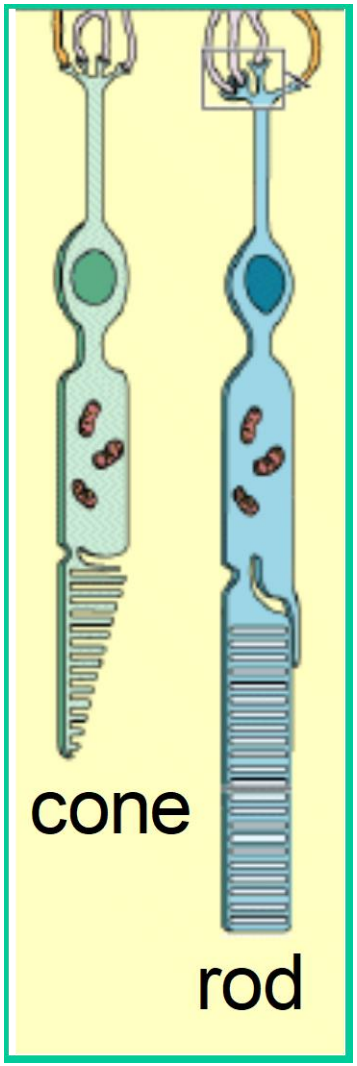

Distribution of Rods and Cones

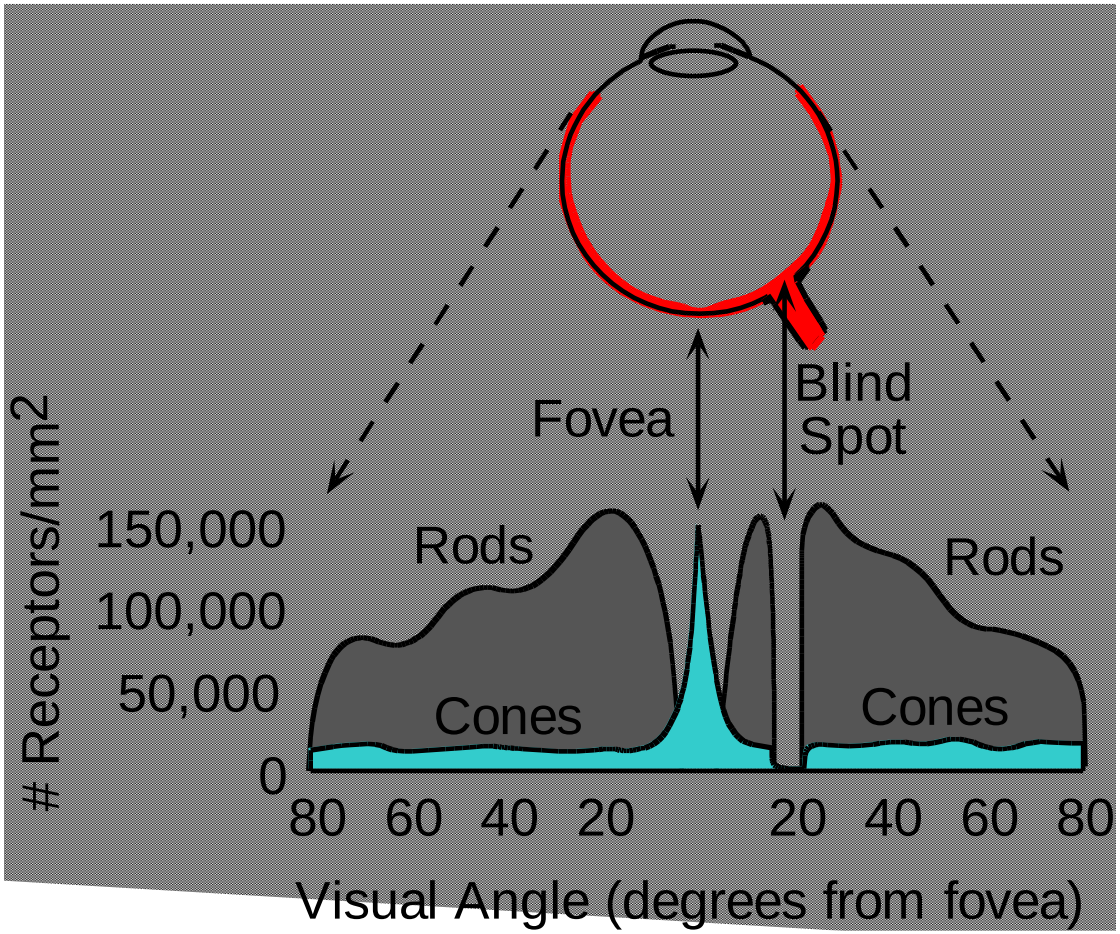

Night Sky: why are there more stars off-center?
Averted vision: http://en.wikipedia.org/wiki/Averted_vision

Rod / Cone sensitivity

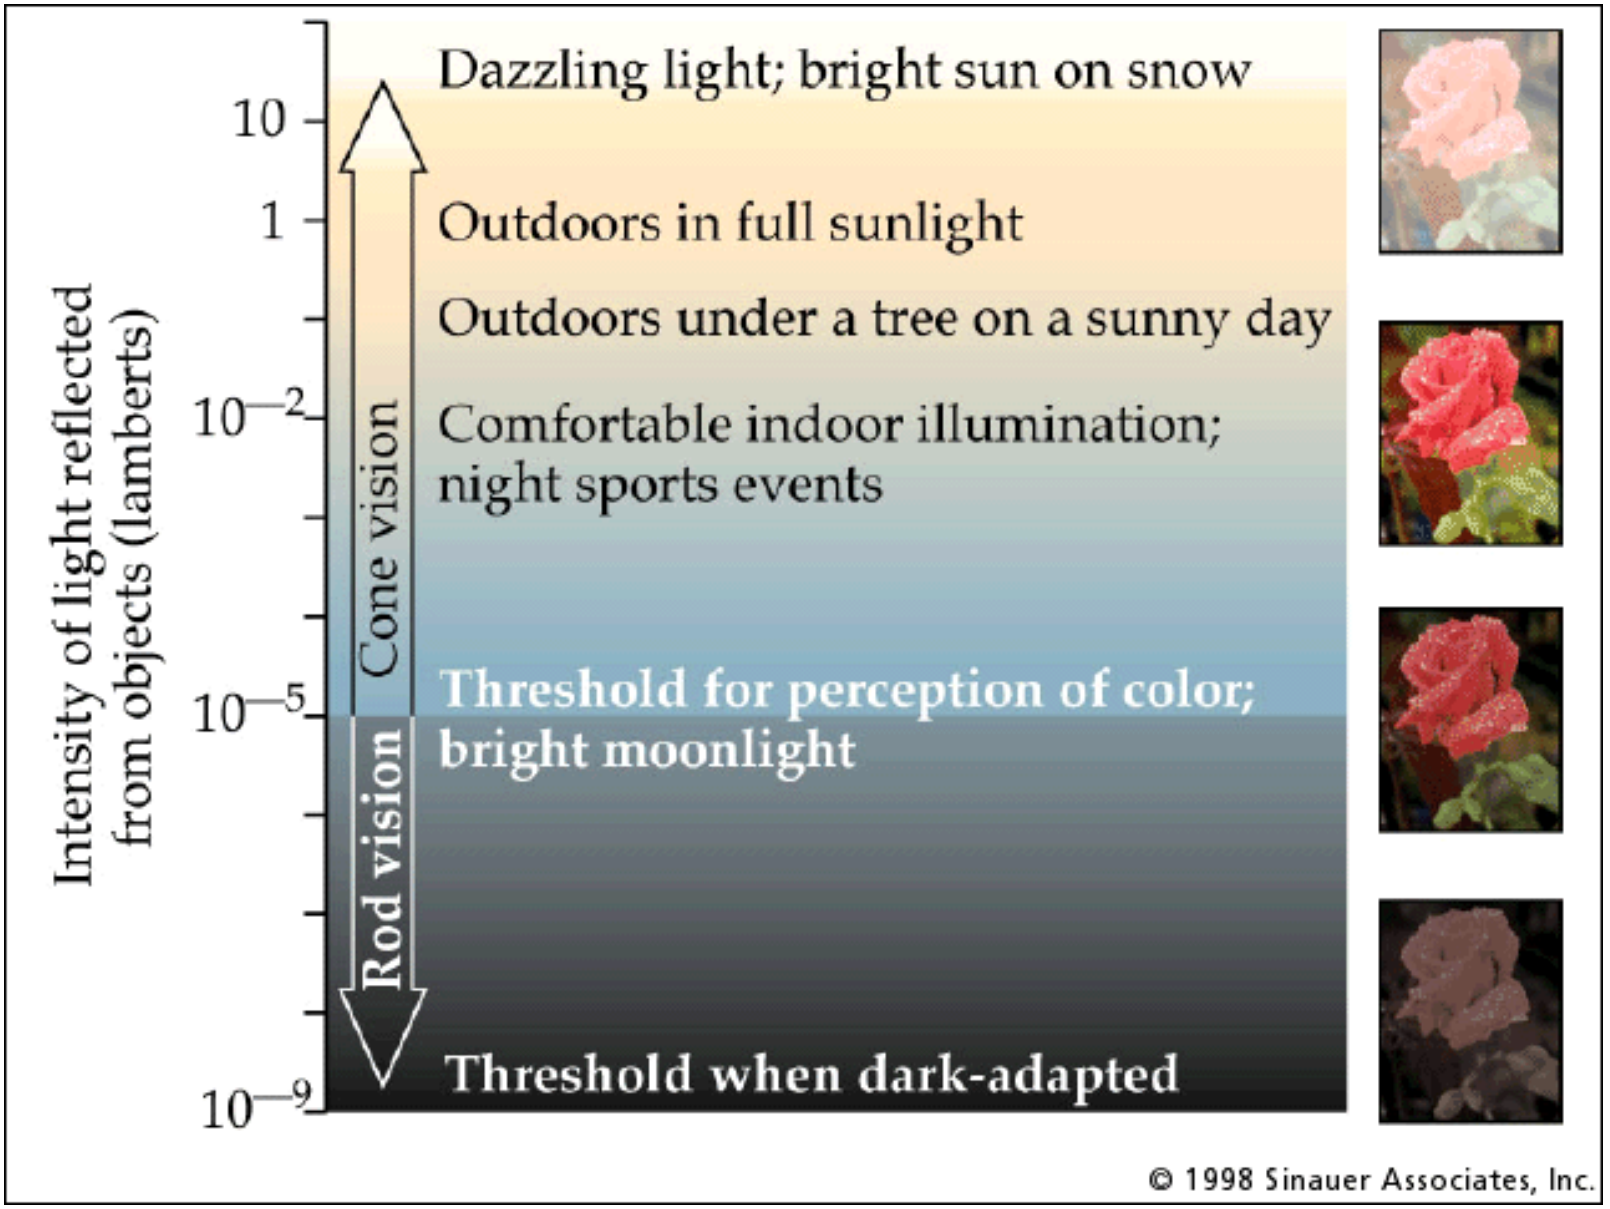

Does the eye alias?

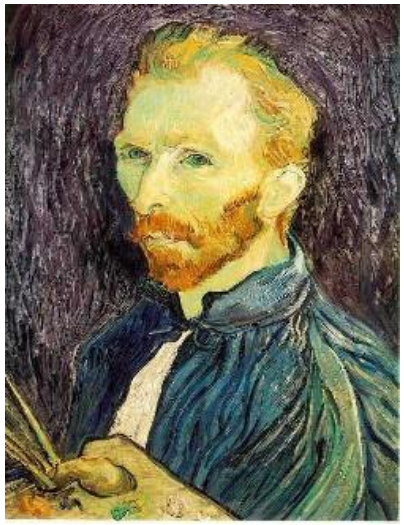

4x downsample
nearest neighbor

Spatially, apparently not.

The retina (sensor) has high resolution, but the optics (lens) of the eye cannot meet that resolution.

The image is blurred optically before being sampled (removes high-frequency content!)

[Thanks to Leslie Bresnahan]

Electromagnetic Spectrum

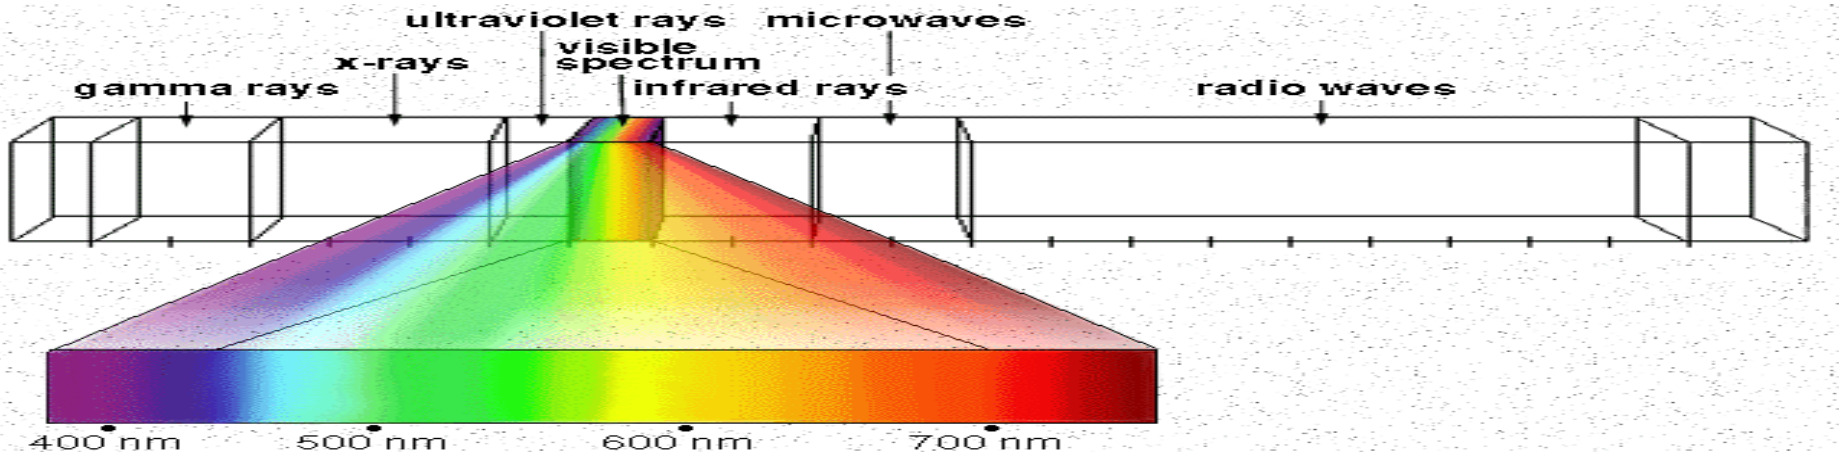

Wavelength of light and its perceived color

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

Physiology of Color Vision

Three kinds of cones:

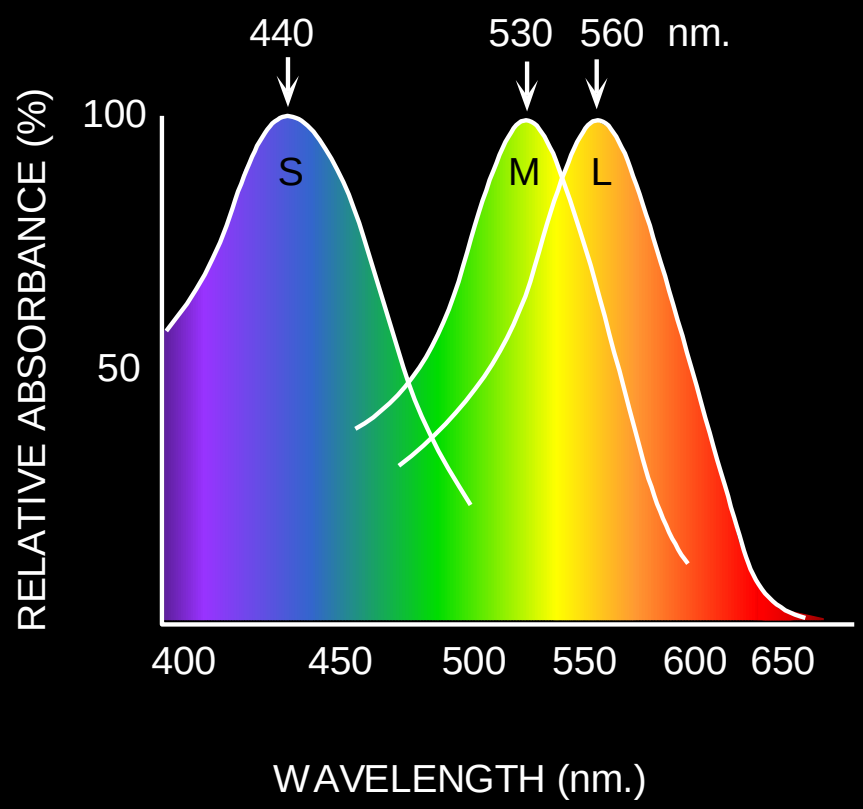

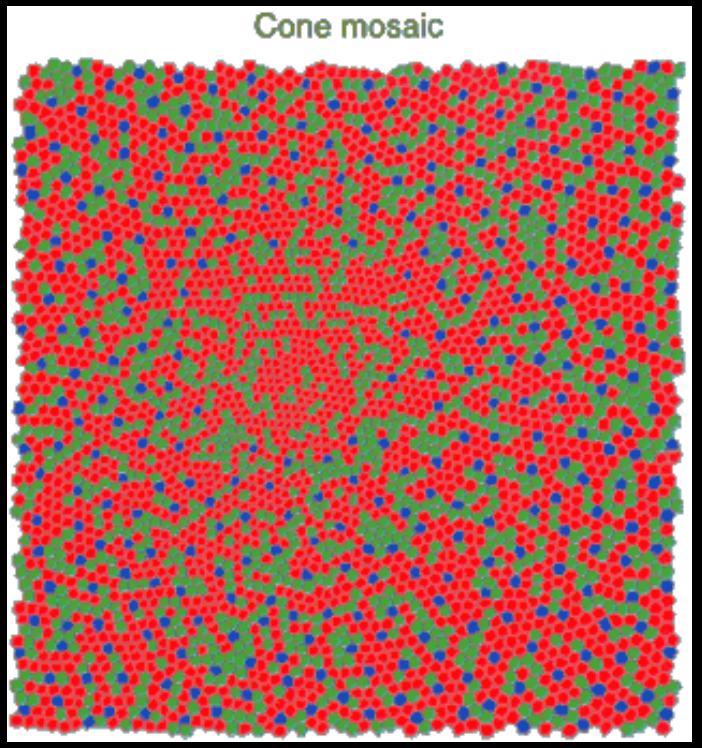

The Physics of Light

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.

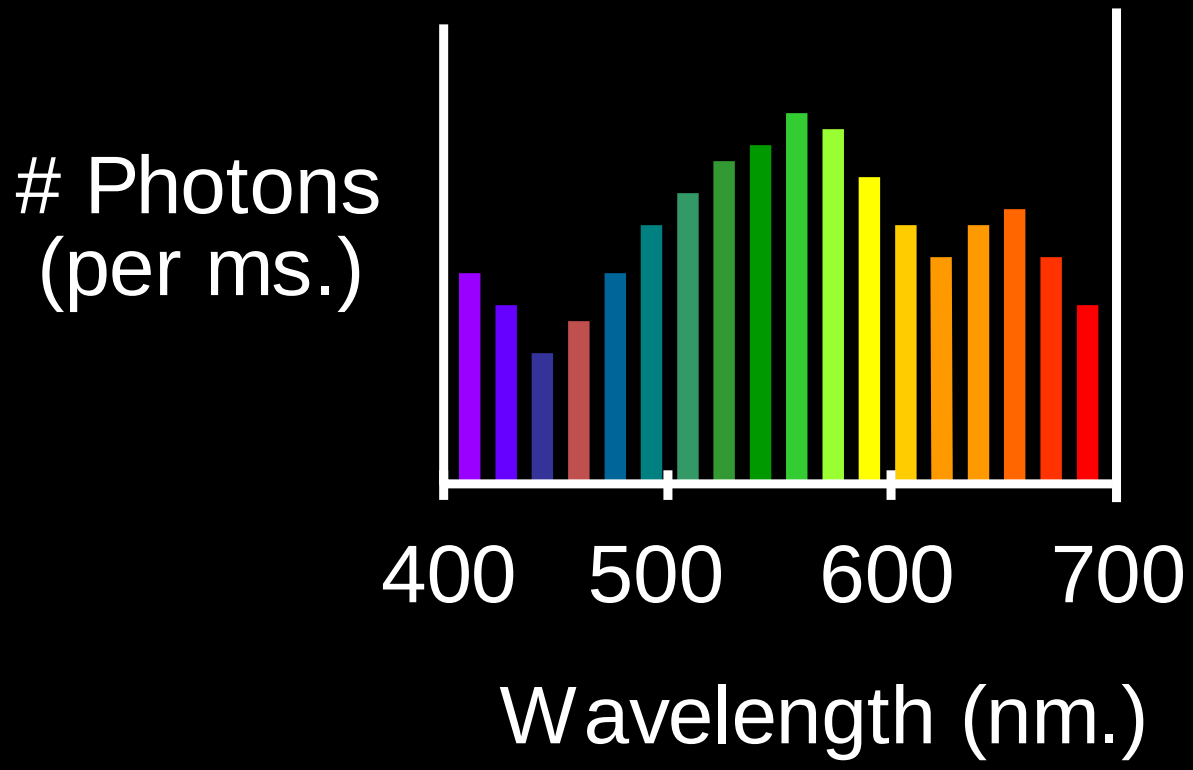

The Physics of Light

Some examples of the spectra of light sources

A. Ruby Laser

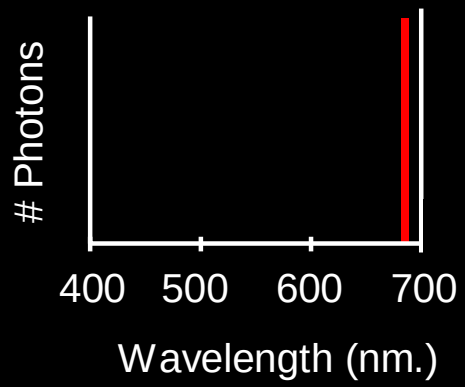

B. Gallium Phosphide Crystal

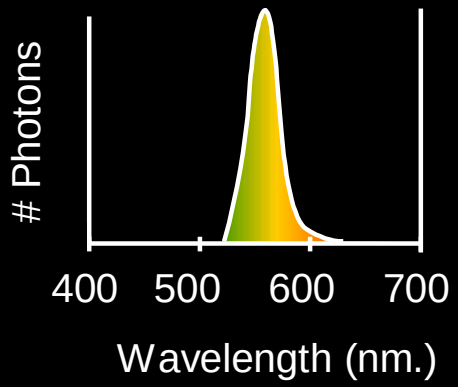

C. Tungsten Lightbulb

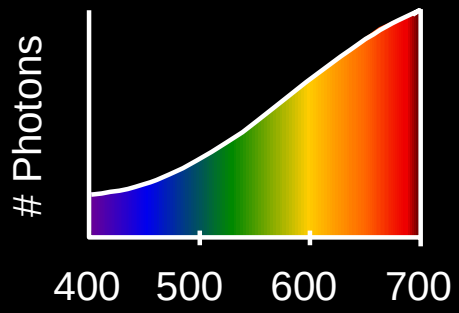

D. Normal Daylight

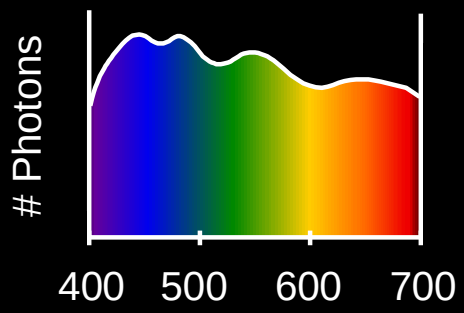

The Physics of Light

Some examples of the reflectance spectra of surfaces




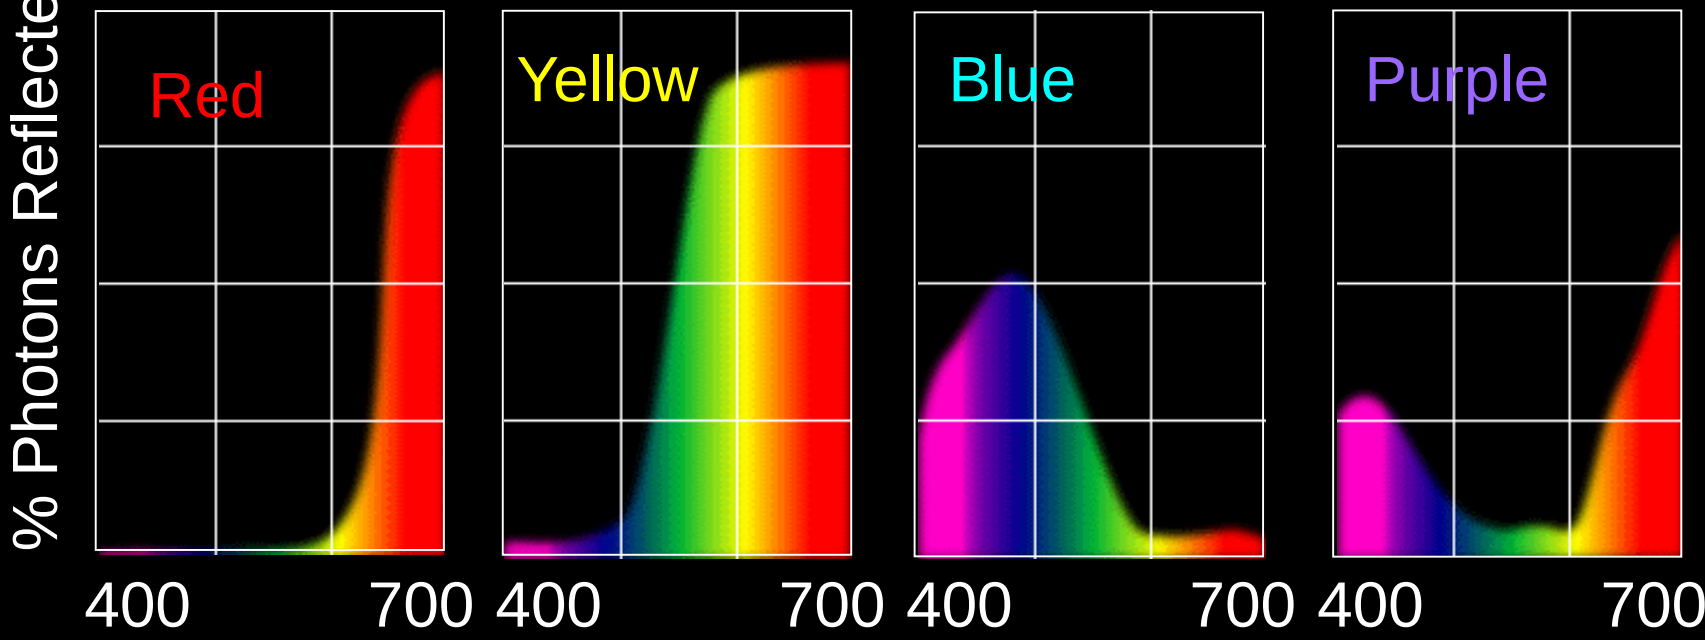

Physiology of Color Vision

Three kinds of cones:

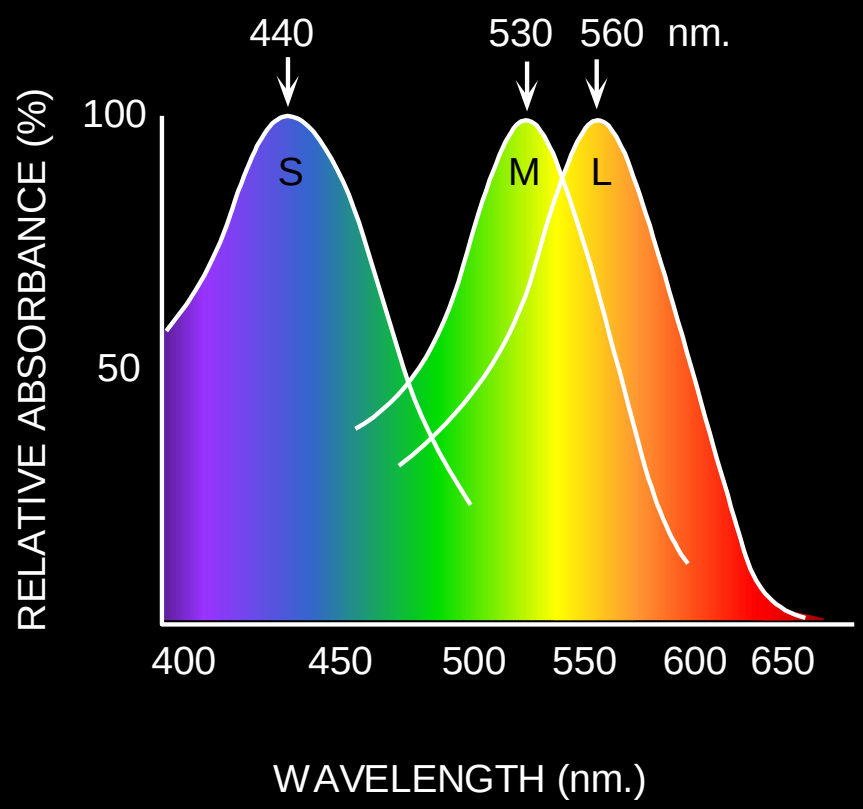

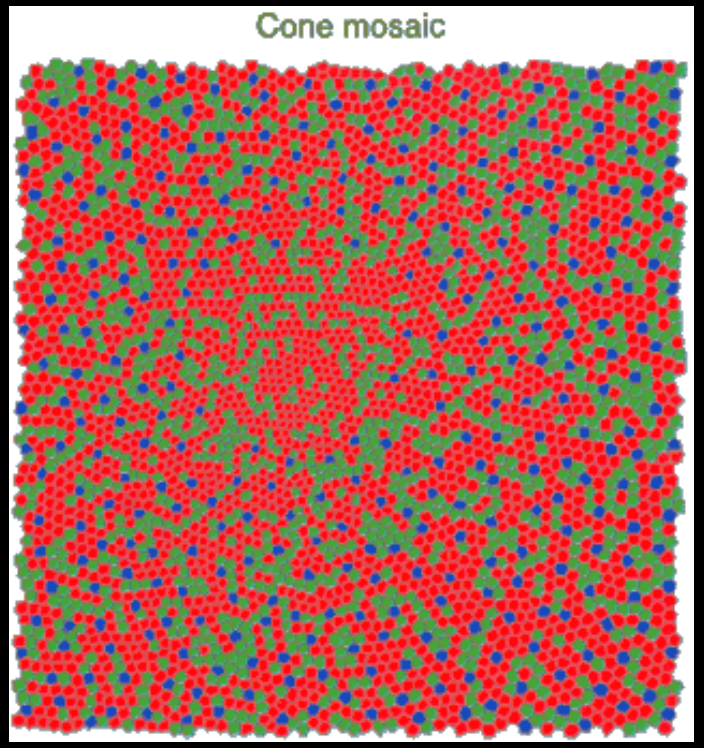

- Why are M and L cones so close?
- Why are there 3?

Electromagnetic Spectrum

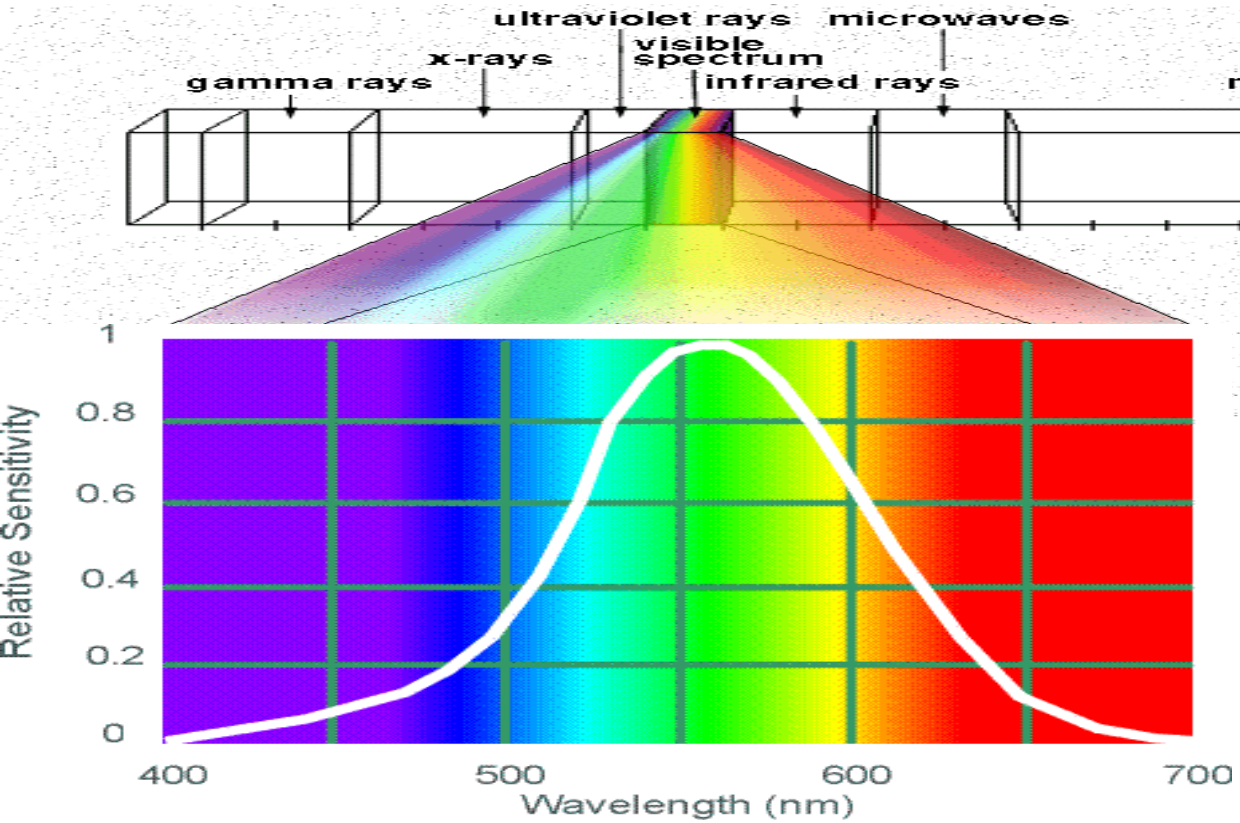

Human Luminance Sensitivity Function

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

Tetrachromatism

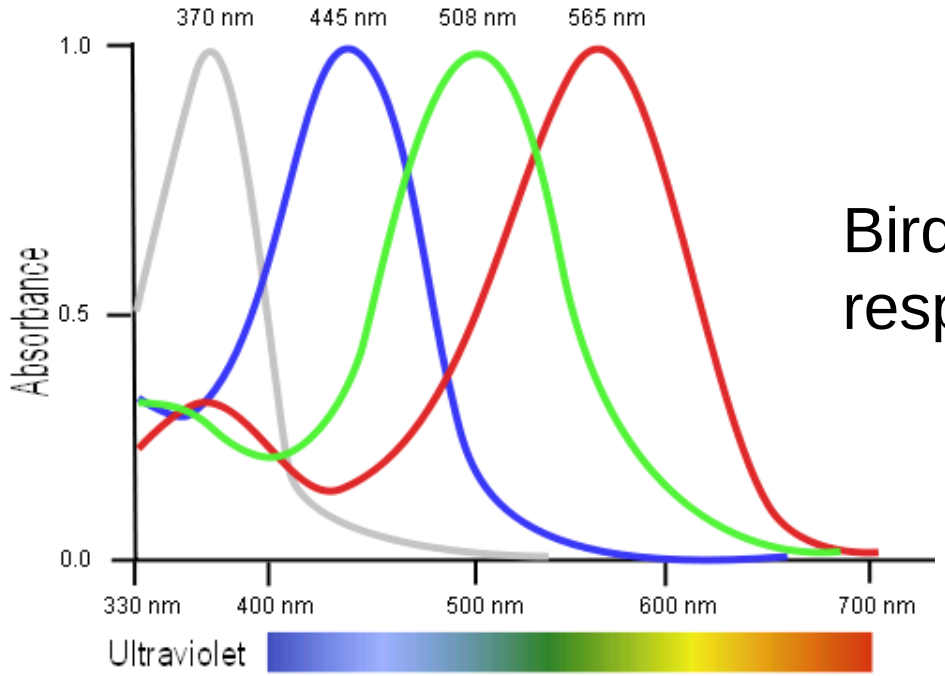

Bird cone responses

- Most birds, and many other animals, have cones for ultraviolet light.
- Some humans seem to have four cones (12% of females).
- True tetrachromatism is _rare_; requires learning.

James Hays

Bee vision

Mantis Shrimp

- 16 photoreceptor types
- ‘Mantis Shrimp’ camera: https://www.theatlantic.com/science/archive/2018/04/mantis-shrimp-eye-camera/557195/

Metamers

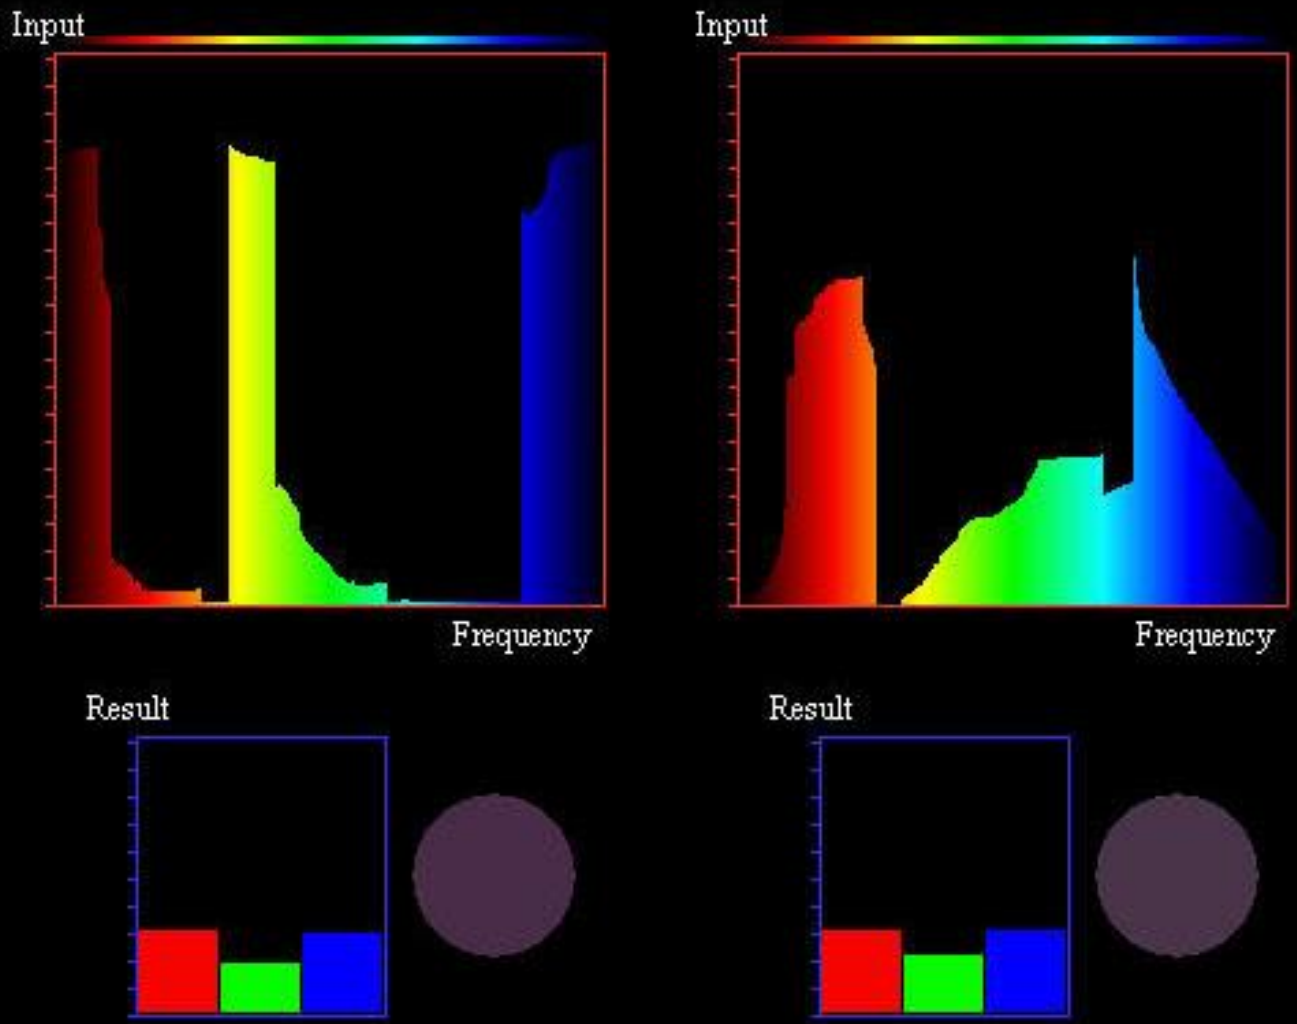

by Jeff Beall, Adam Doppelt and John F. Hughes

Impossible Colors

Can you make the cones respond in ways that typical light spectra never would?

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

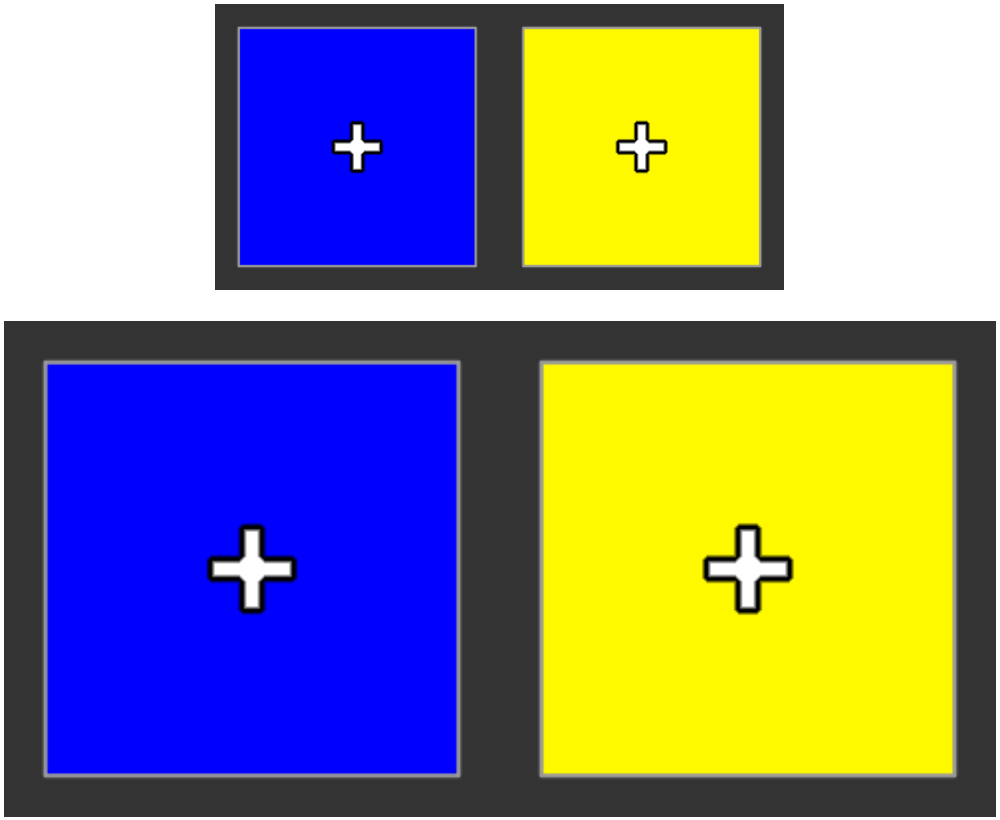

CHIMERICAL COLOR DEMO TEMPLATES

Fatigue template (stare at "x")	Target field (glance at "x")	Approximate Rendering	
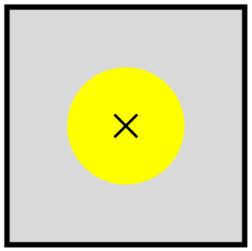	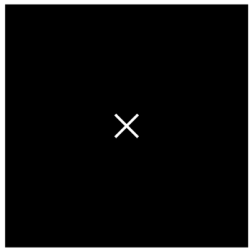	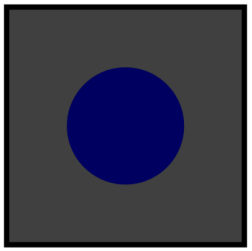	STYGIAN BLUE (simultaneously deep blue and black)
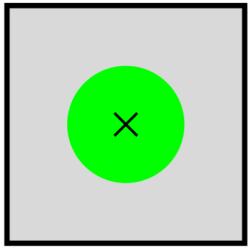	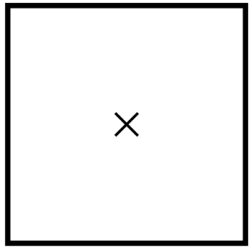	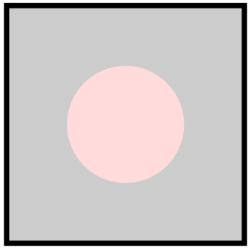	SELF-LUMINOUS RED (simultaneously red and brighter than white)
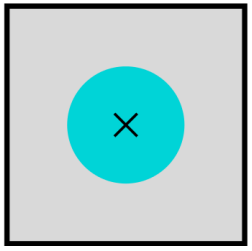	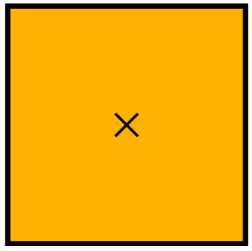	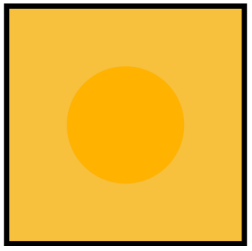	HYPERBOLIC ORANGE (more than 100% color saturation)

What is color?

Why do we even care about human vision in this class?

Why do we care about human vision?

- We don’t, necessarily.

- But biological vision shows that it is possible to make important judgements from images.

Why do we care about human vision?

- We don’t, necessarily.
- But biological vision shows that it is possible to make important judgements from images.
- It’s a human world -> cameras imitate the frequency response of the human eye to try to see as we see.

Ornithopters

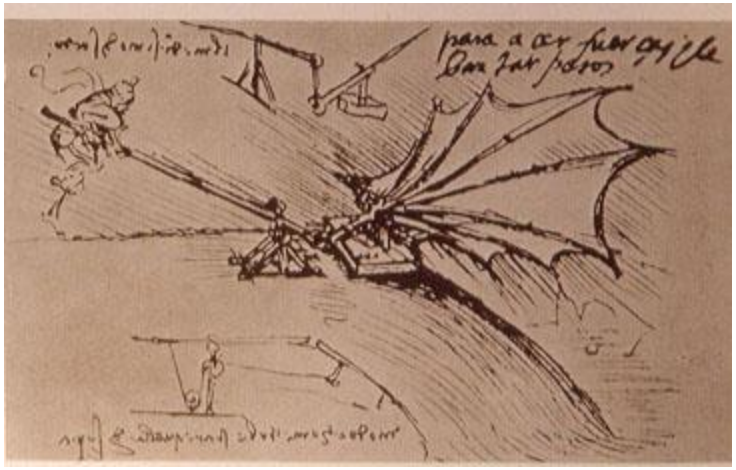

"Can machines fly like a bird?"
No, because airplanes don’t flap.

"Can machines fly?"
Yes, but airplanes use a different mechanism.

"Can machines perceive?"
Is this question like the first, or like the second?

Adapted from Peter Norvig

Cameras with Three Sensors

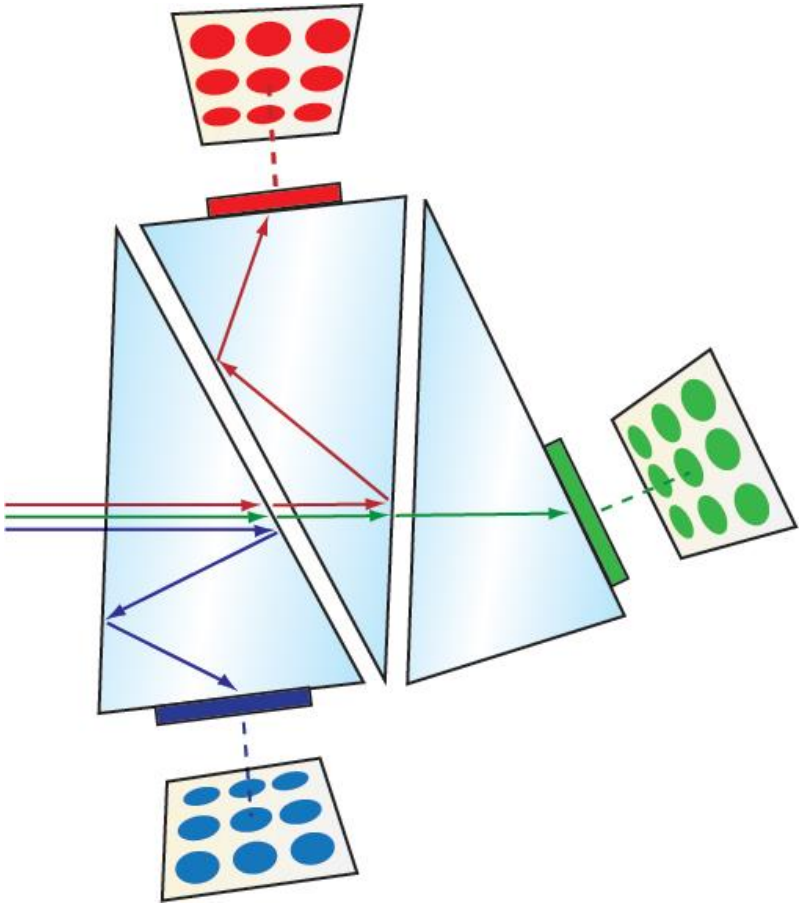

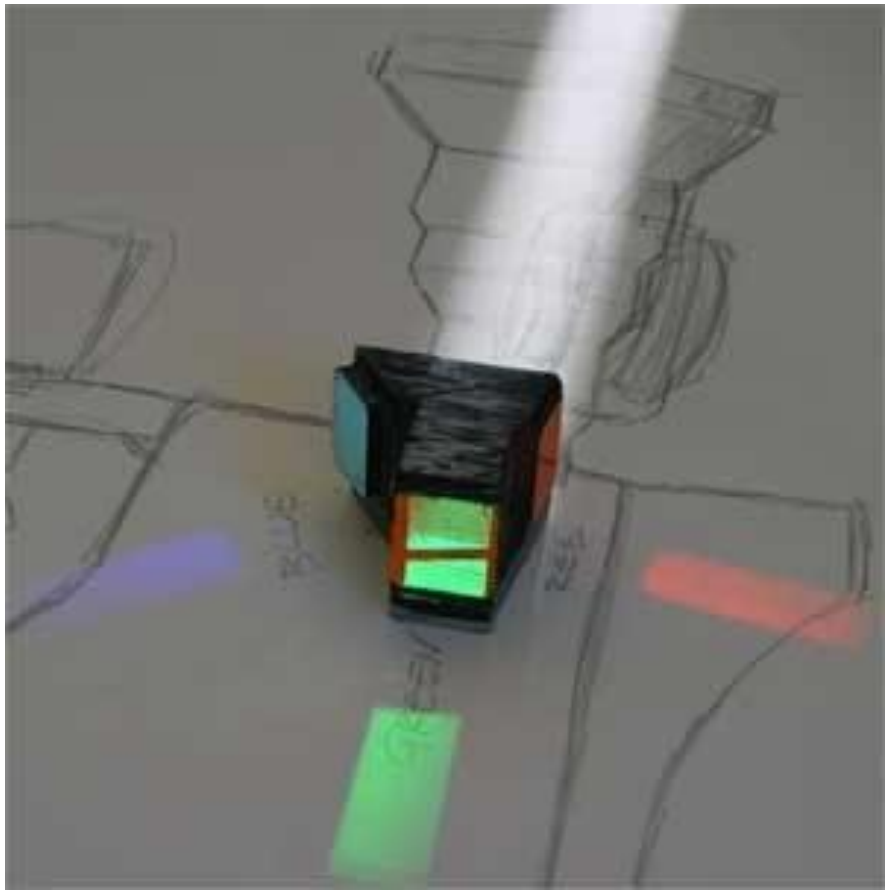

[Edmund Optics; Adam Wilt]

Color Sensing in Camera (RGB)

- 3-chip vs. 1-chip: quality vs. cost
- Why more green?

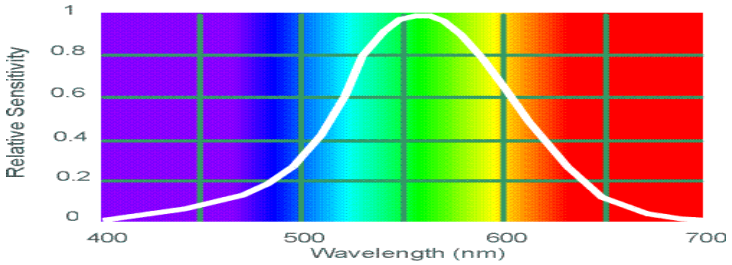

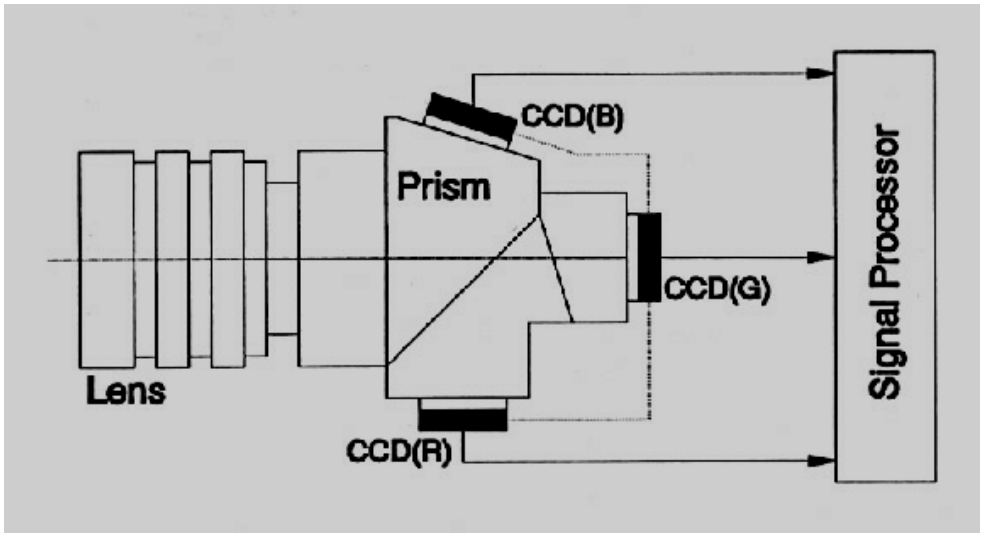

Why 3 colors?

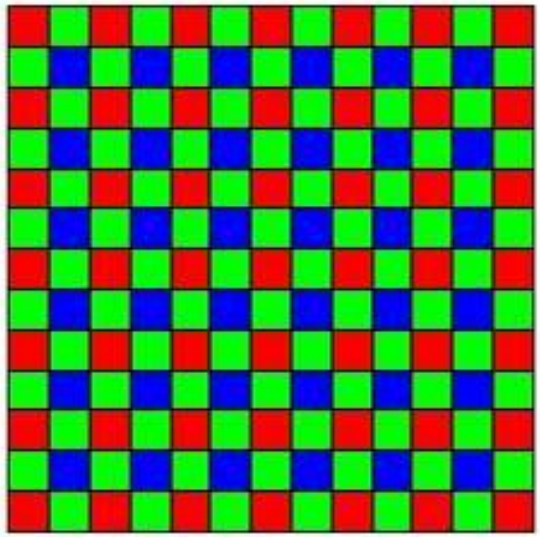

Bayer filter

Stuff Works

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

Slide by Steve Seitz

Cheaper/More Compact Color Sensing: Bayer Grid

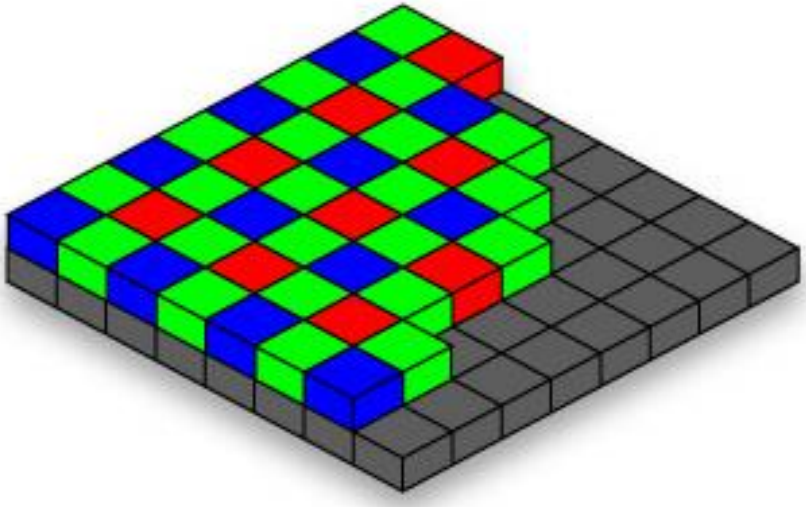

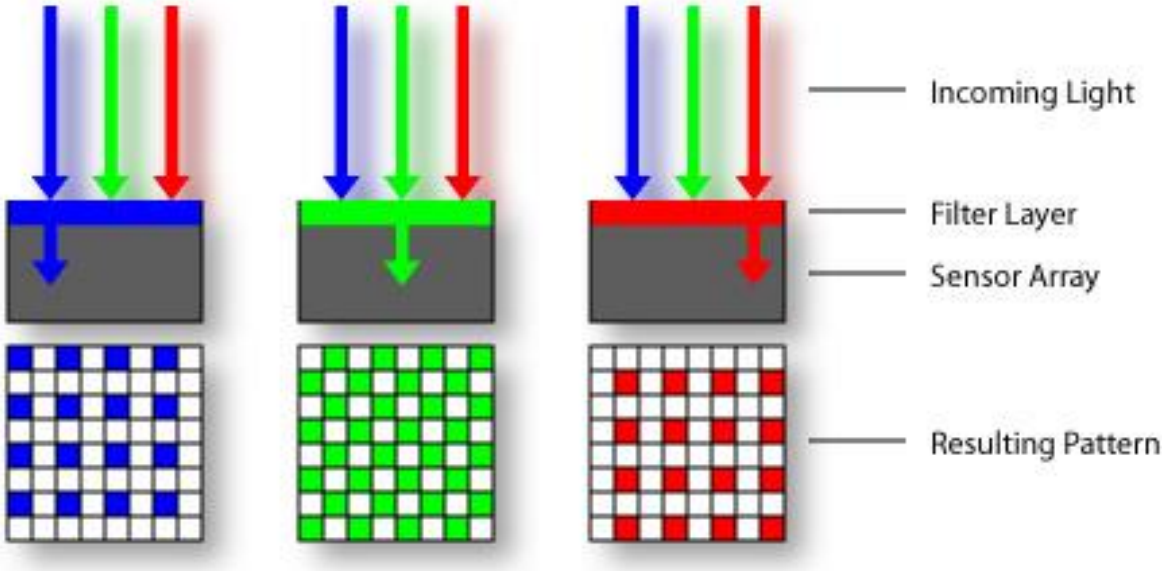

Estimate RGB at ‘G’ cells from neighboring values

Slide by Steve Seitz

Why more green?

Approximate human spectral sensitivity

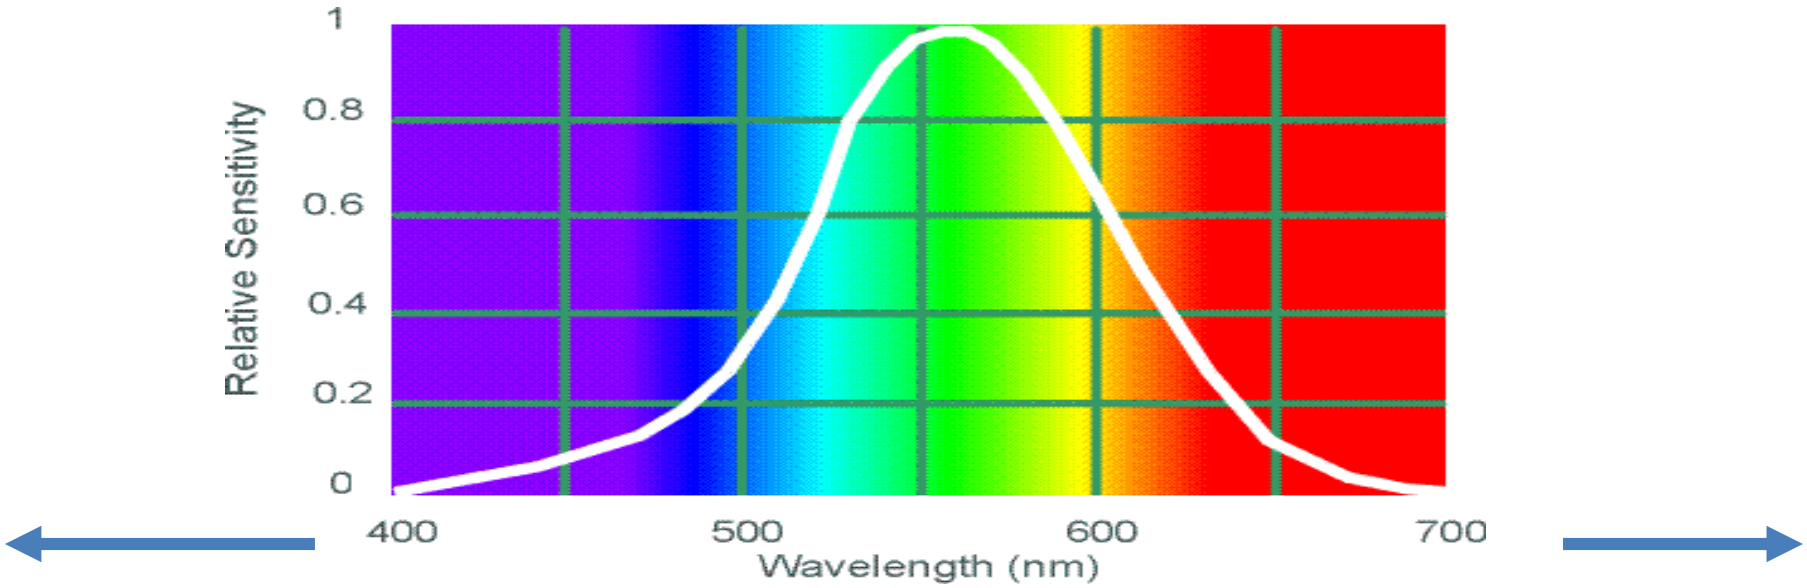

Less than
~400nm to 10nm
= ultraviolet (UV)

Human visible portion
of electromagnetic
(EM) spectrum

Greater than
~700nm to 1mm
= infrared (IR)

RGB Camera Color Response

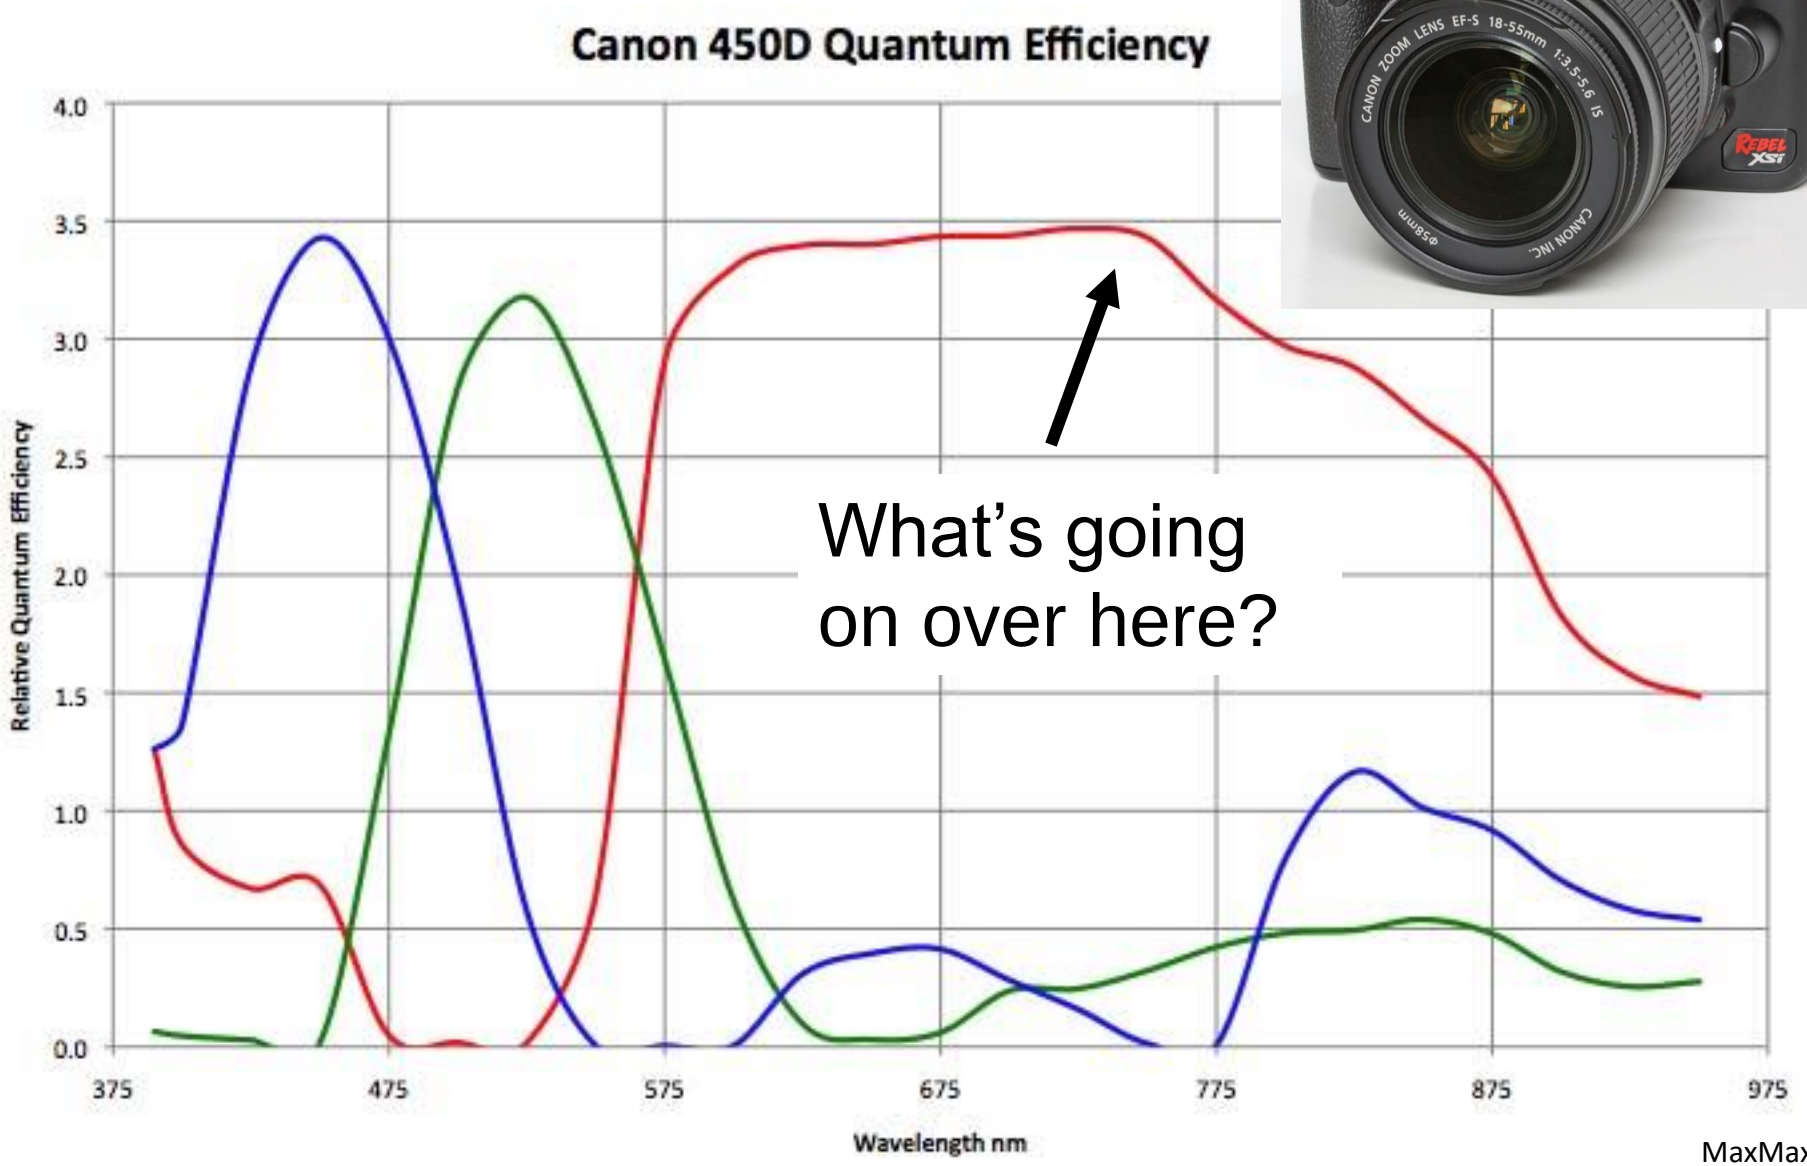

Display Color Response

ASUS

PA248Q LCD Monitor
Color Calibration Testing Report

Every ASUS PA248Q is equipped with pre-tuned sRGB and has undergone rigorous tests and calibration processes to ensure that color difference, ΔE, is less than 5, thus preventing color inaccuracy and inconsistency on screen.
ASUS advanced gray-scale tracking technology ensures smoother color gradation delivered by every ASUS PA248Q.

Serial No.	J3LMQS158406
Test Equipment	Minolta Color Analyzer CA210

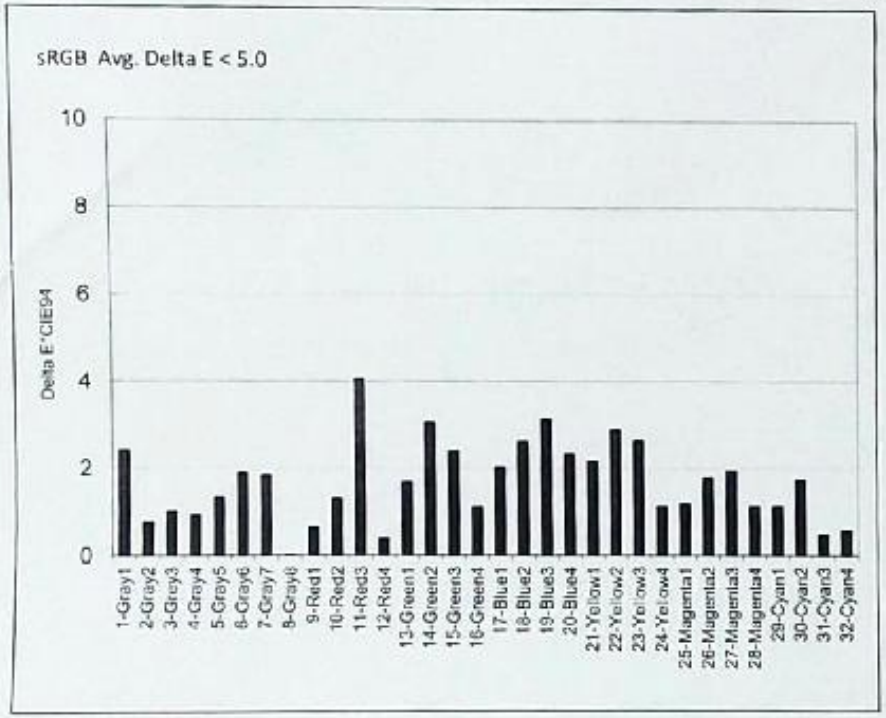

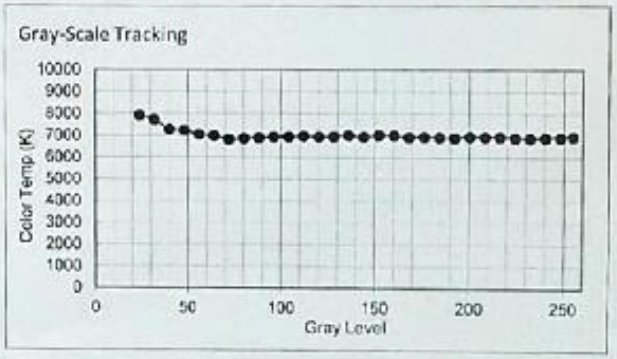

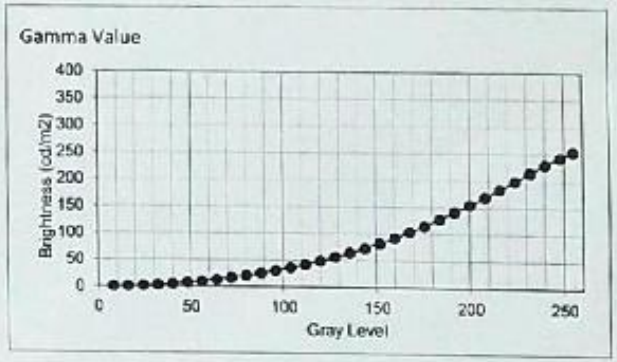

Note: The sRGB calibrations of every ASUS PA248Q are pre-tuned and tested under ASUS standardized procedures using calibrated facilities at the factory manufacturing line. This report is a certificate only for the newly manufactured ASUS PA248Q monitor unit. Test results may vary under different test procedures, equipment and patterns.

Display Color Response

ASUS

PA248Q LCD Monitor
Color Calibration Testing Report

Every ASUS PA248Q is equipped with pre-tuned sRGB and has undergone rigorous tests and calibration processes to ensure that color difference, ΔE, is less than 5, thus preventing color inaccuracy and inconsistency on screen.
ASUS advanced gray-scale tracking technology ensures smoother color gradation delivered by every ASUS PA248Q.

Serial No.	J4LMQS157085
Test Equipment	Minolta Color Analyzer CA210

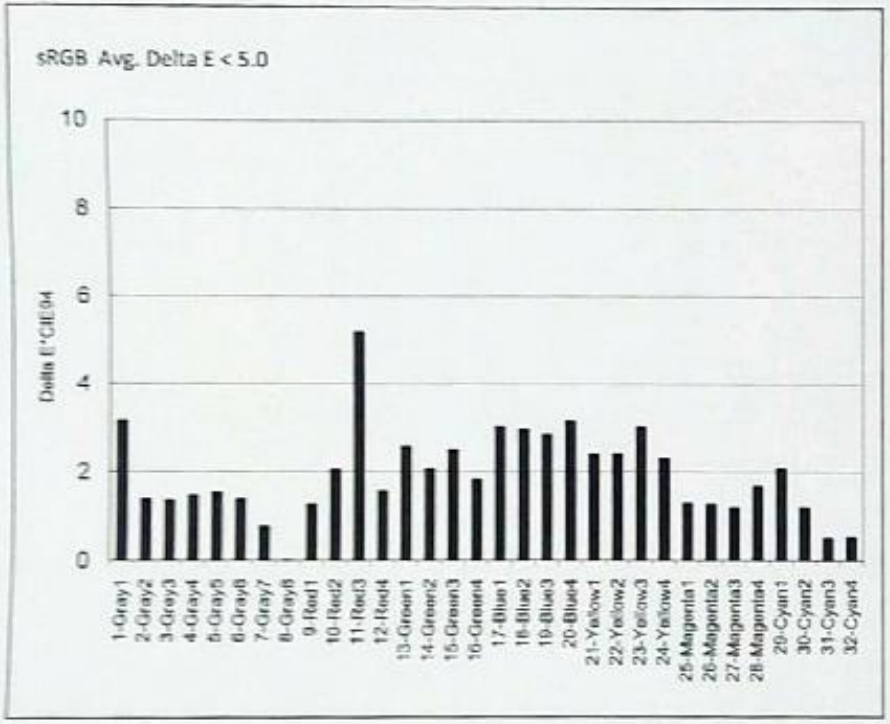

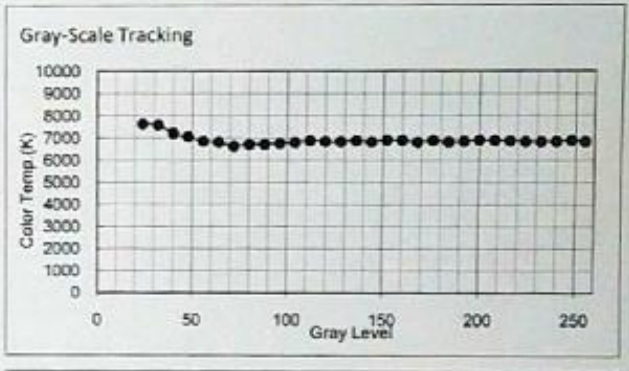

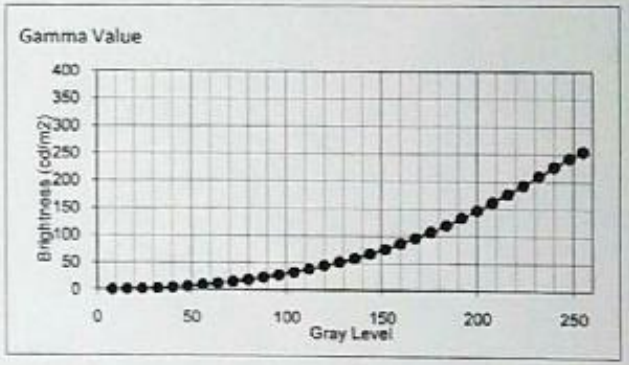

Note: The sRGB calibrations of every ASUS PA248Q are pre-tuned and tested under ASUS standardized procedures using calibrated facilities at the factory manufacturing line. This report is a certificate only for the newly manufactured ASUS PA248Q monitor unit. Test results may vary under different test procedures, equipment and patterns.

Color spaces

How can we represent color?

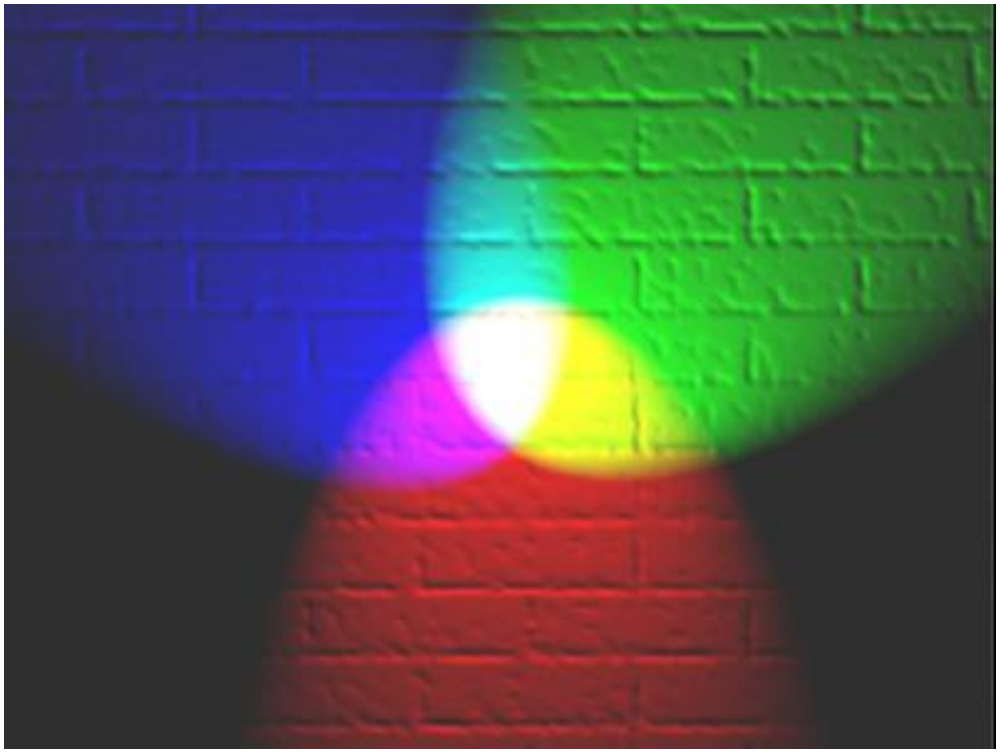

http://en.wikipedia.org/wiki/File:RGB_illumination.jpg

Color spaces: RGB

Default color space

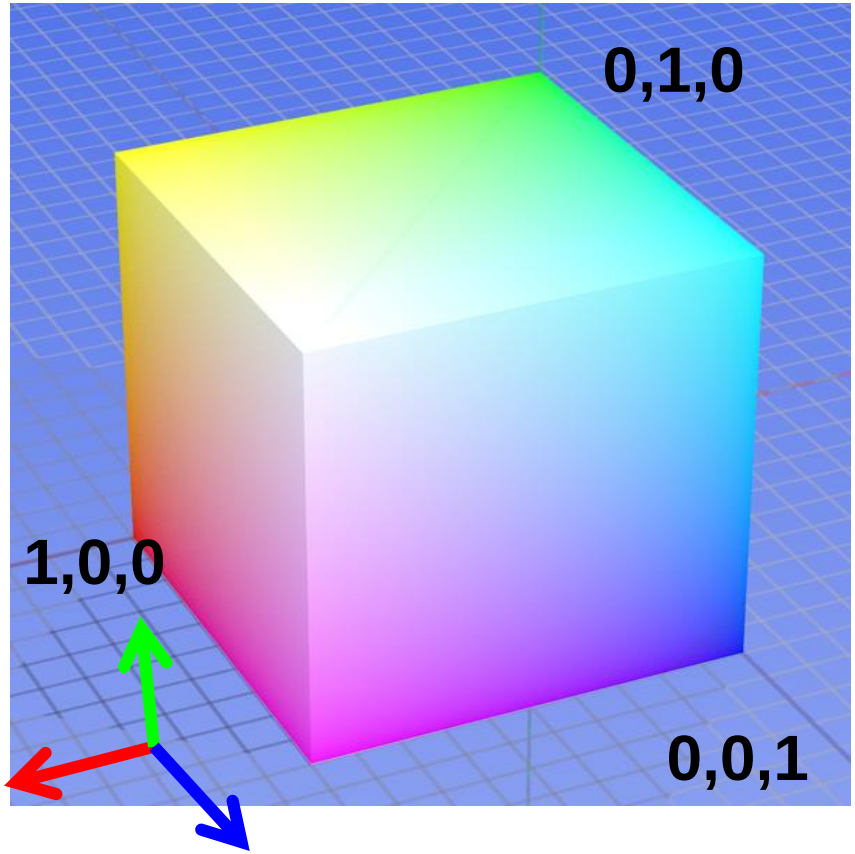

R = 1
(G=0,B=0)

G = 1
(R=0,B=0)

B = 1
(R=0,G=0)

Any color = r*R + g*G + b*B

- Strongly correlated channels
- Non-perceptual

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

Got it. $C = r*R + g*G + b*B$

IS COLOR A VECTOR SPACE? THINK-PAIR-SHARE

Color spaces: HSV

Intuitive color space

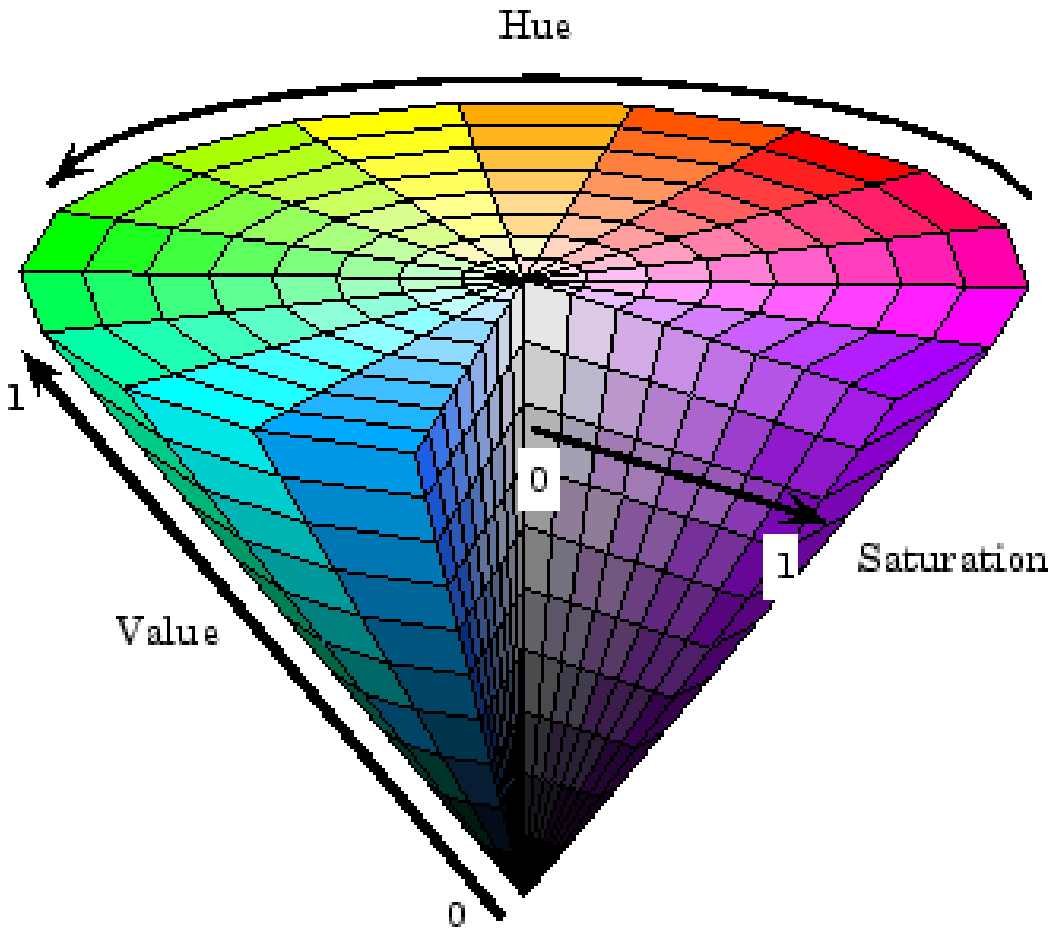

If you had to choose, would you rather go without:
- intensity (‘value’), or
- hue + saturation (‘chroma’)?

Think-Pair-Share

Most information in intensity

Only color shown – constant intensity

Most information in intensity

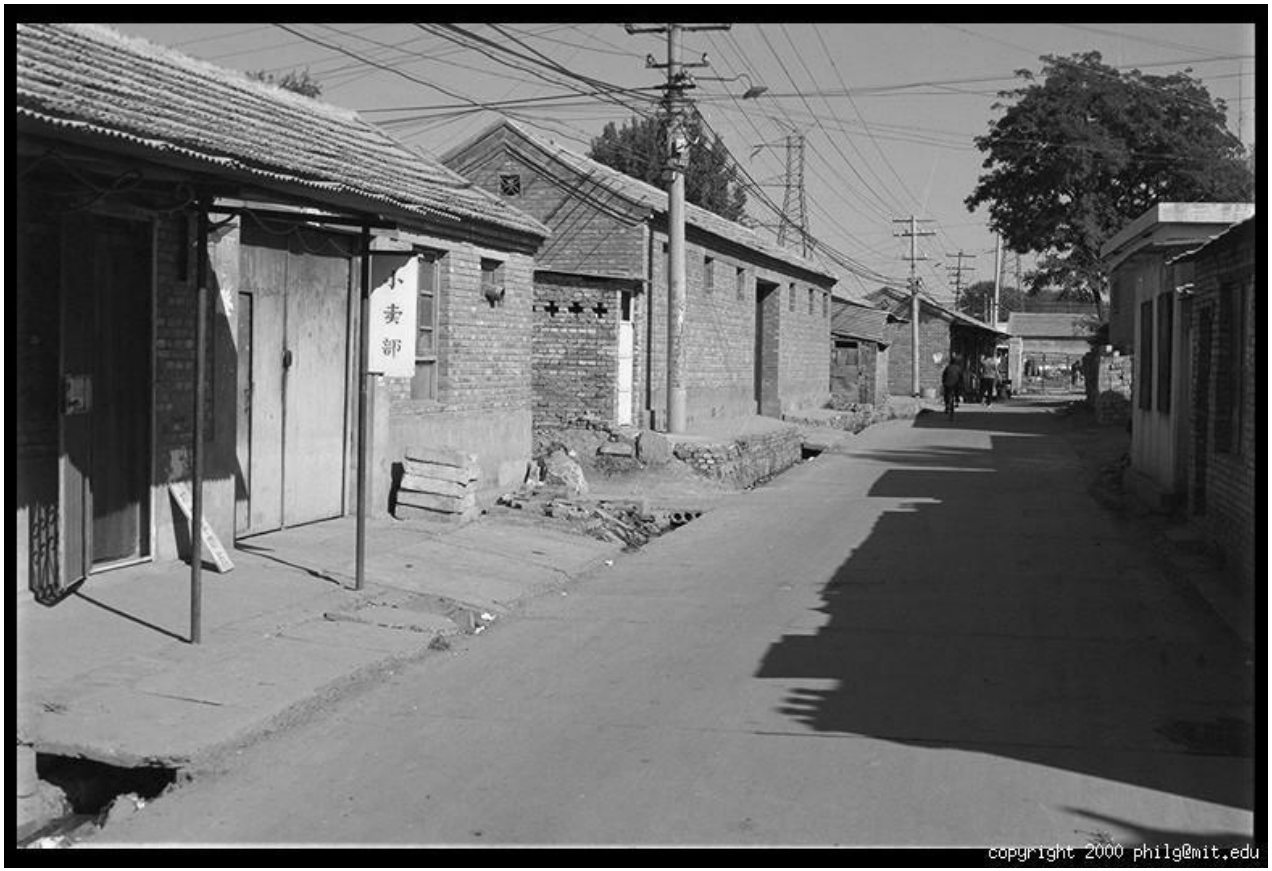

Only intensity shown – constant color

Most information in intensity

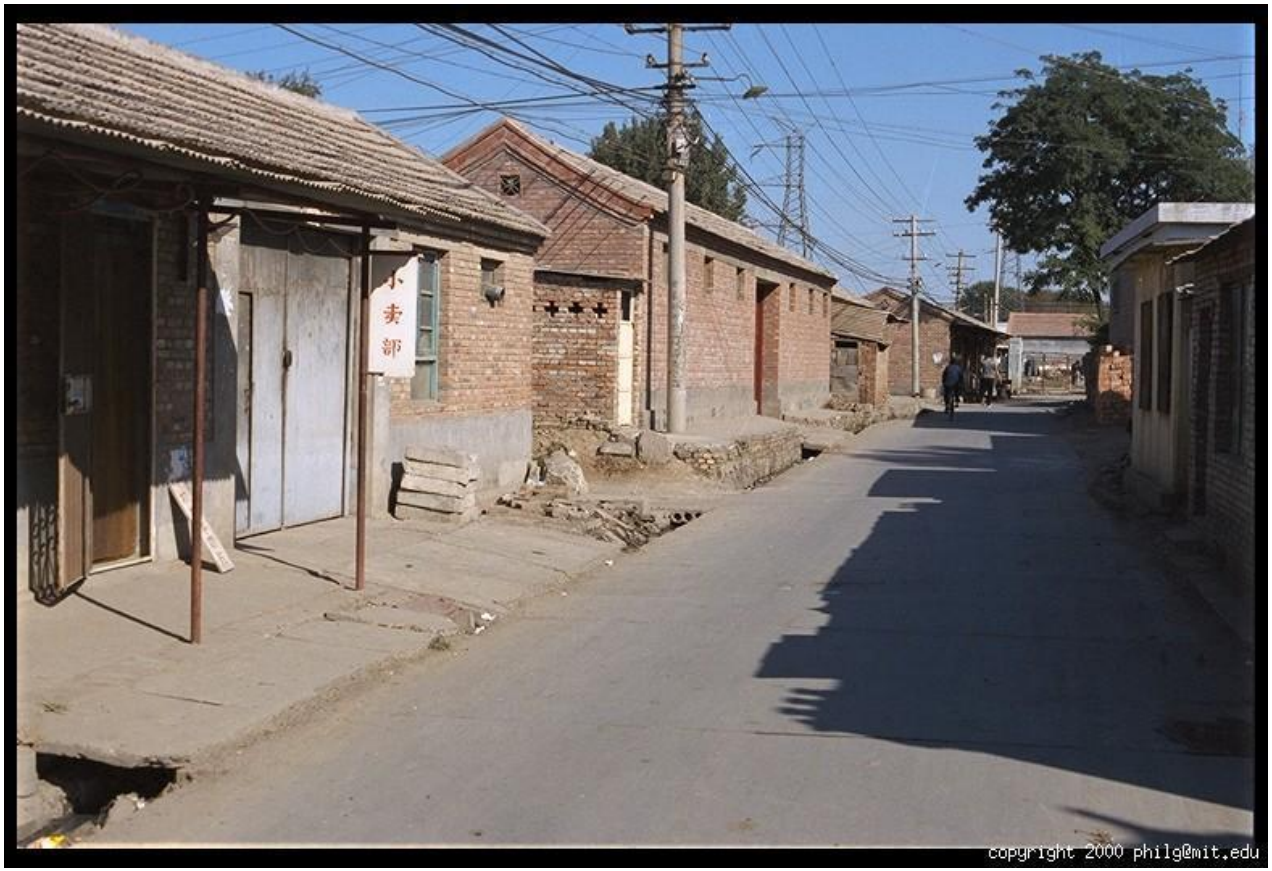

Original image

Color spaces: HSV

Intuitive color space

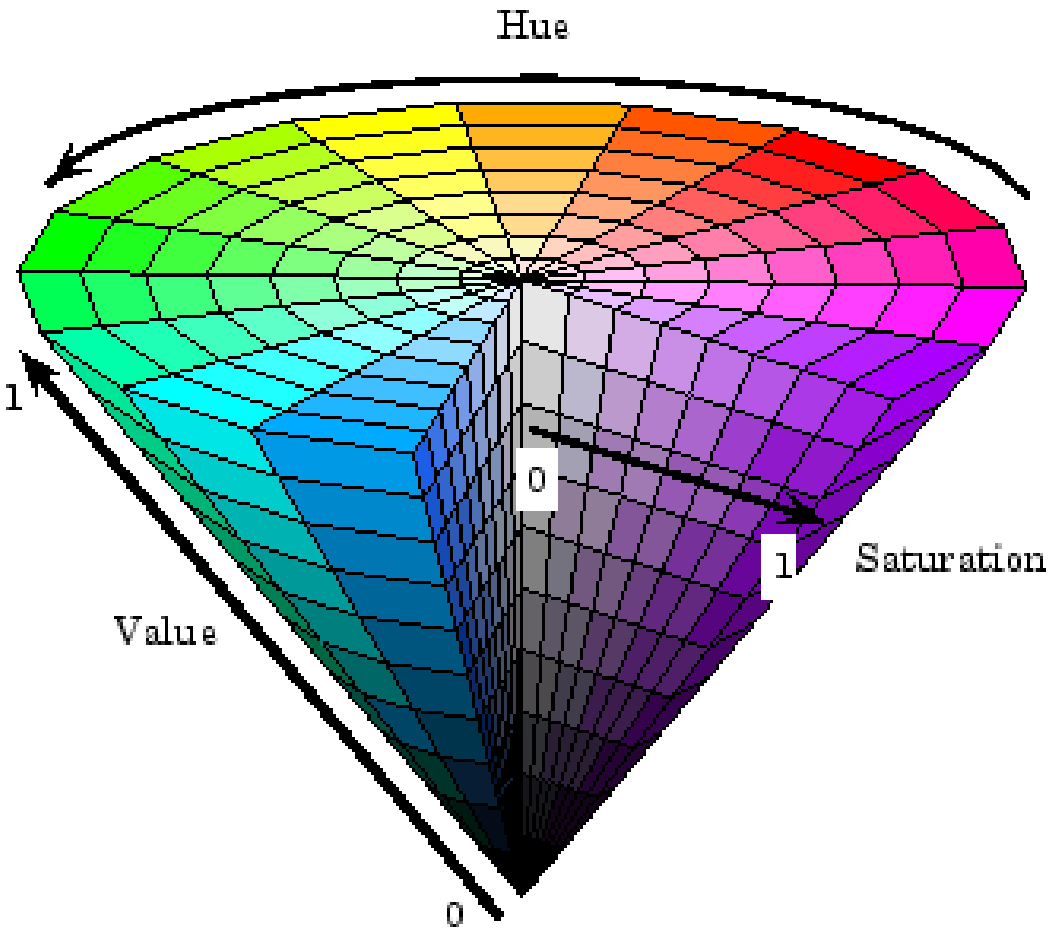

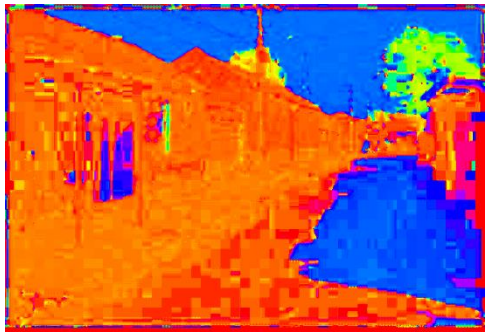

H
(S=1,V=1)

S
(H=1,V=1)

V
(H=1,S=0)

James Hays

Color spaces: YCbCr

Fast to compute, good for compression, used by TV

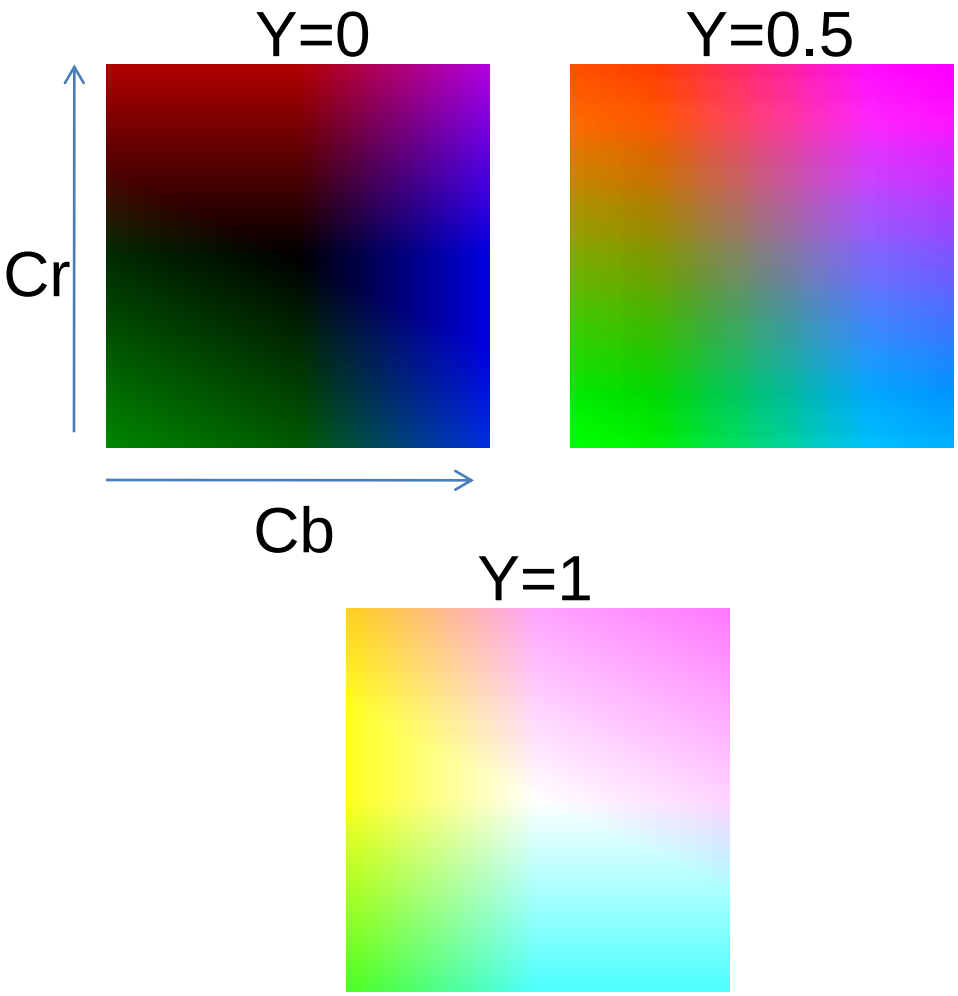

Y
(Cb=0.5,Cr=0.5)

Cb
(Y=0.5,Cr=0.5)

Cr
(Y=0.5,Cb=05)

Most JPEG images & videos subsample chroma

PSP Comp 3
2x2 Chroma subsampling
285K

Original
1,261K lossless
968K PNG

IS COLOR PERCEPTION A VECTOR SPACE?

Color spaces: L*a*b*

"Perceptually uniform"* color space

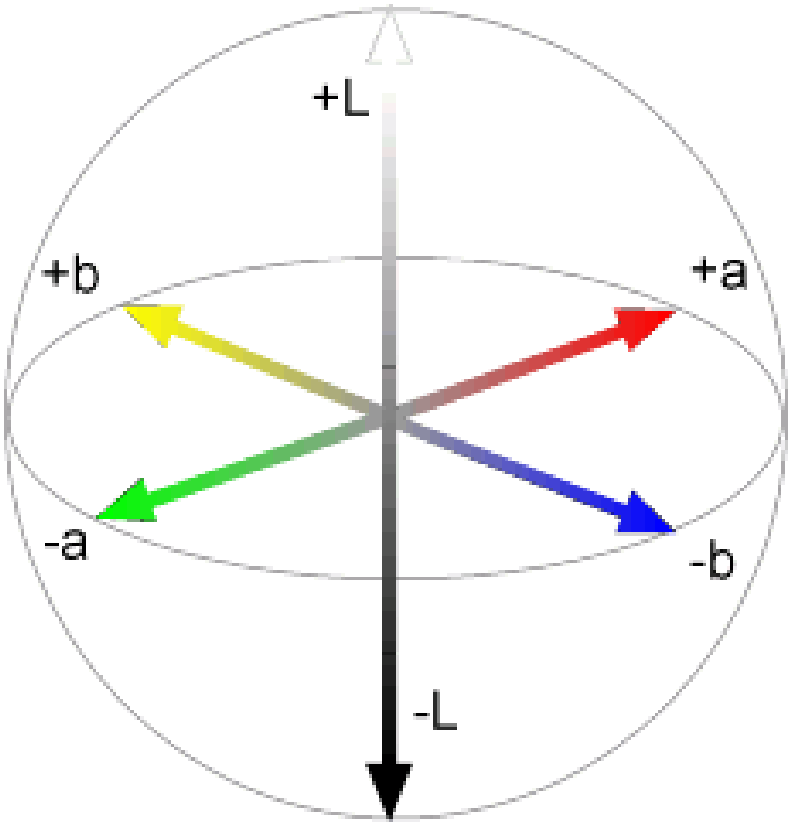

L
(a=0,b=0)

a
(L=65,b=0)

b
(L=65,a=0)

James Hays

EVOLUTION OF MY UNDERSTANDING OF COLOR OVER TIME:

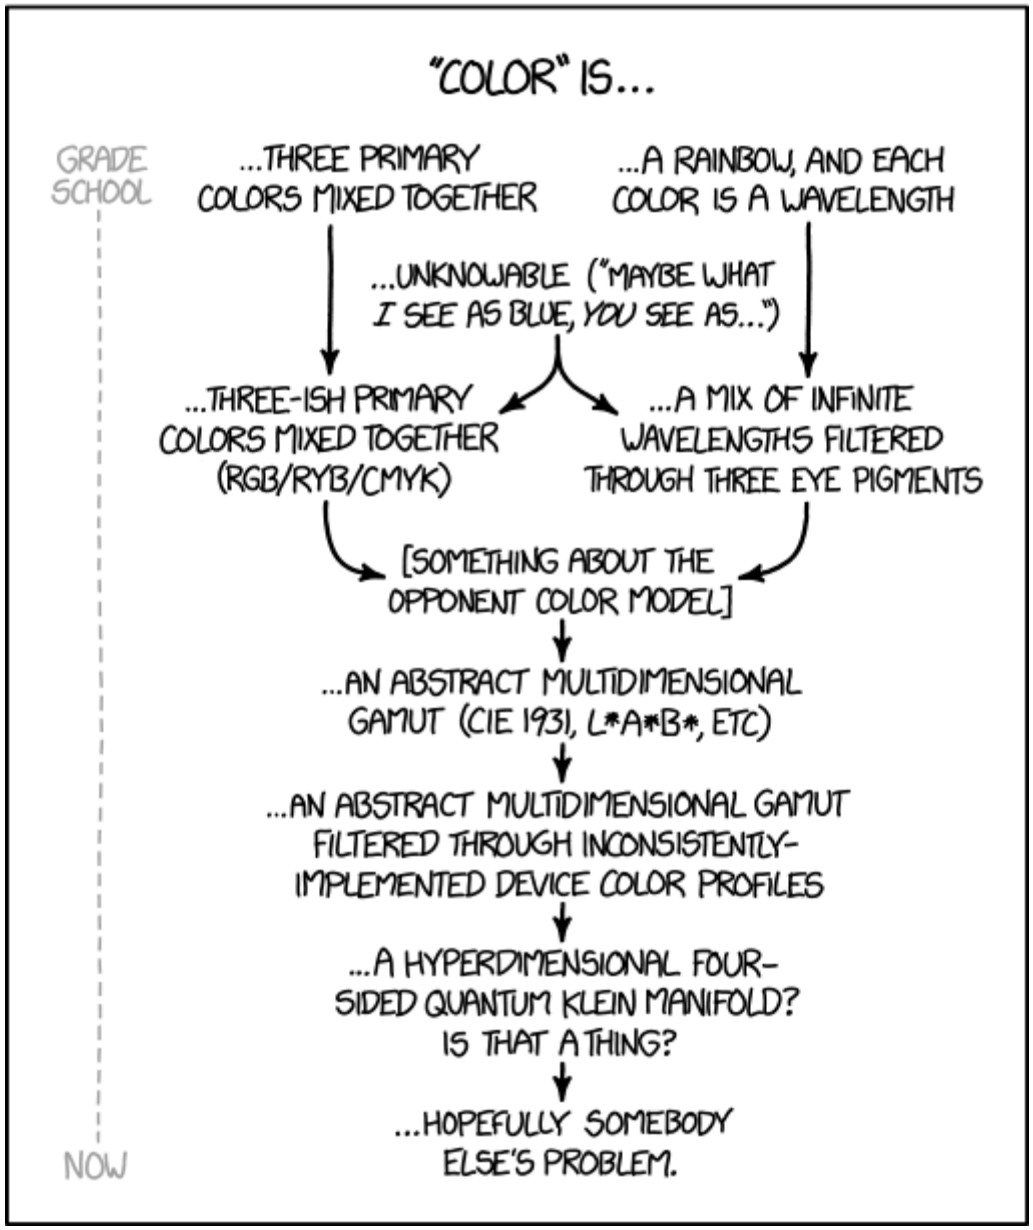

XKCD

More references

- https://www.colorsystem.com/

- A description of many different color systems developed through history.
- Navigate from the right-hand links.

- Thanks to Alex Nibley!

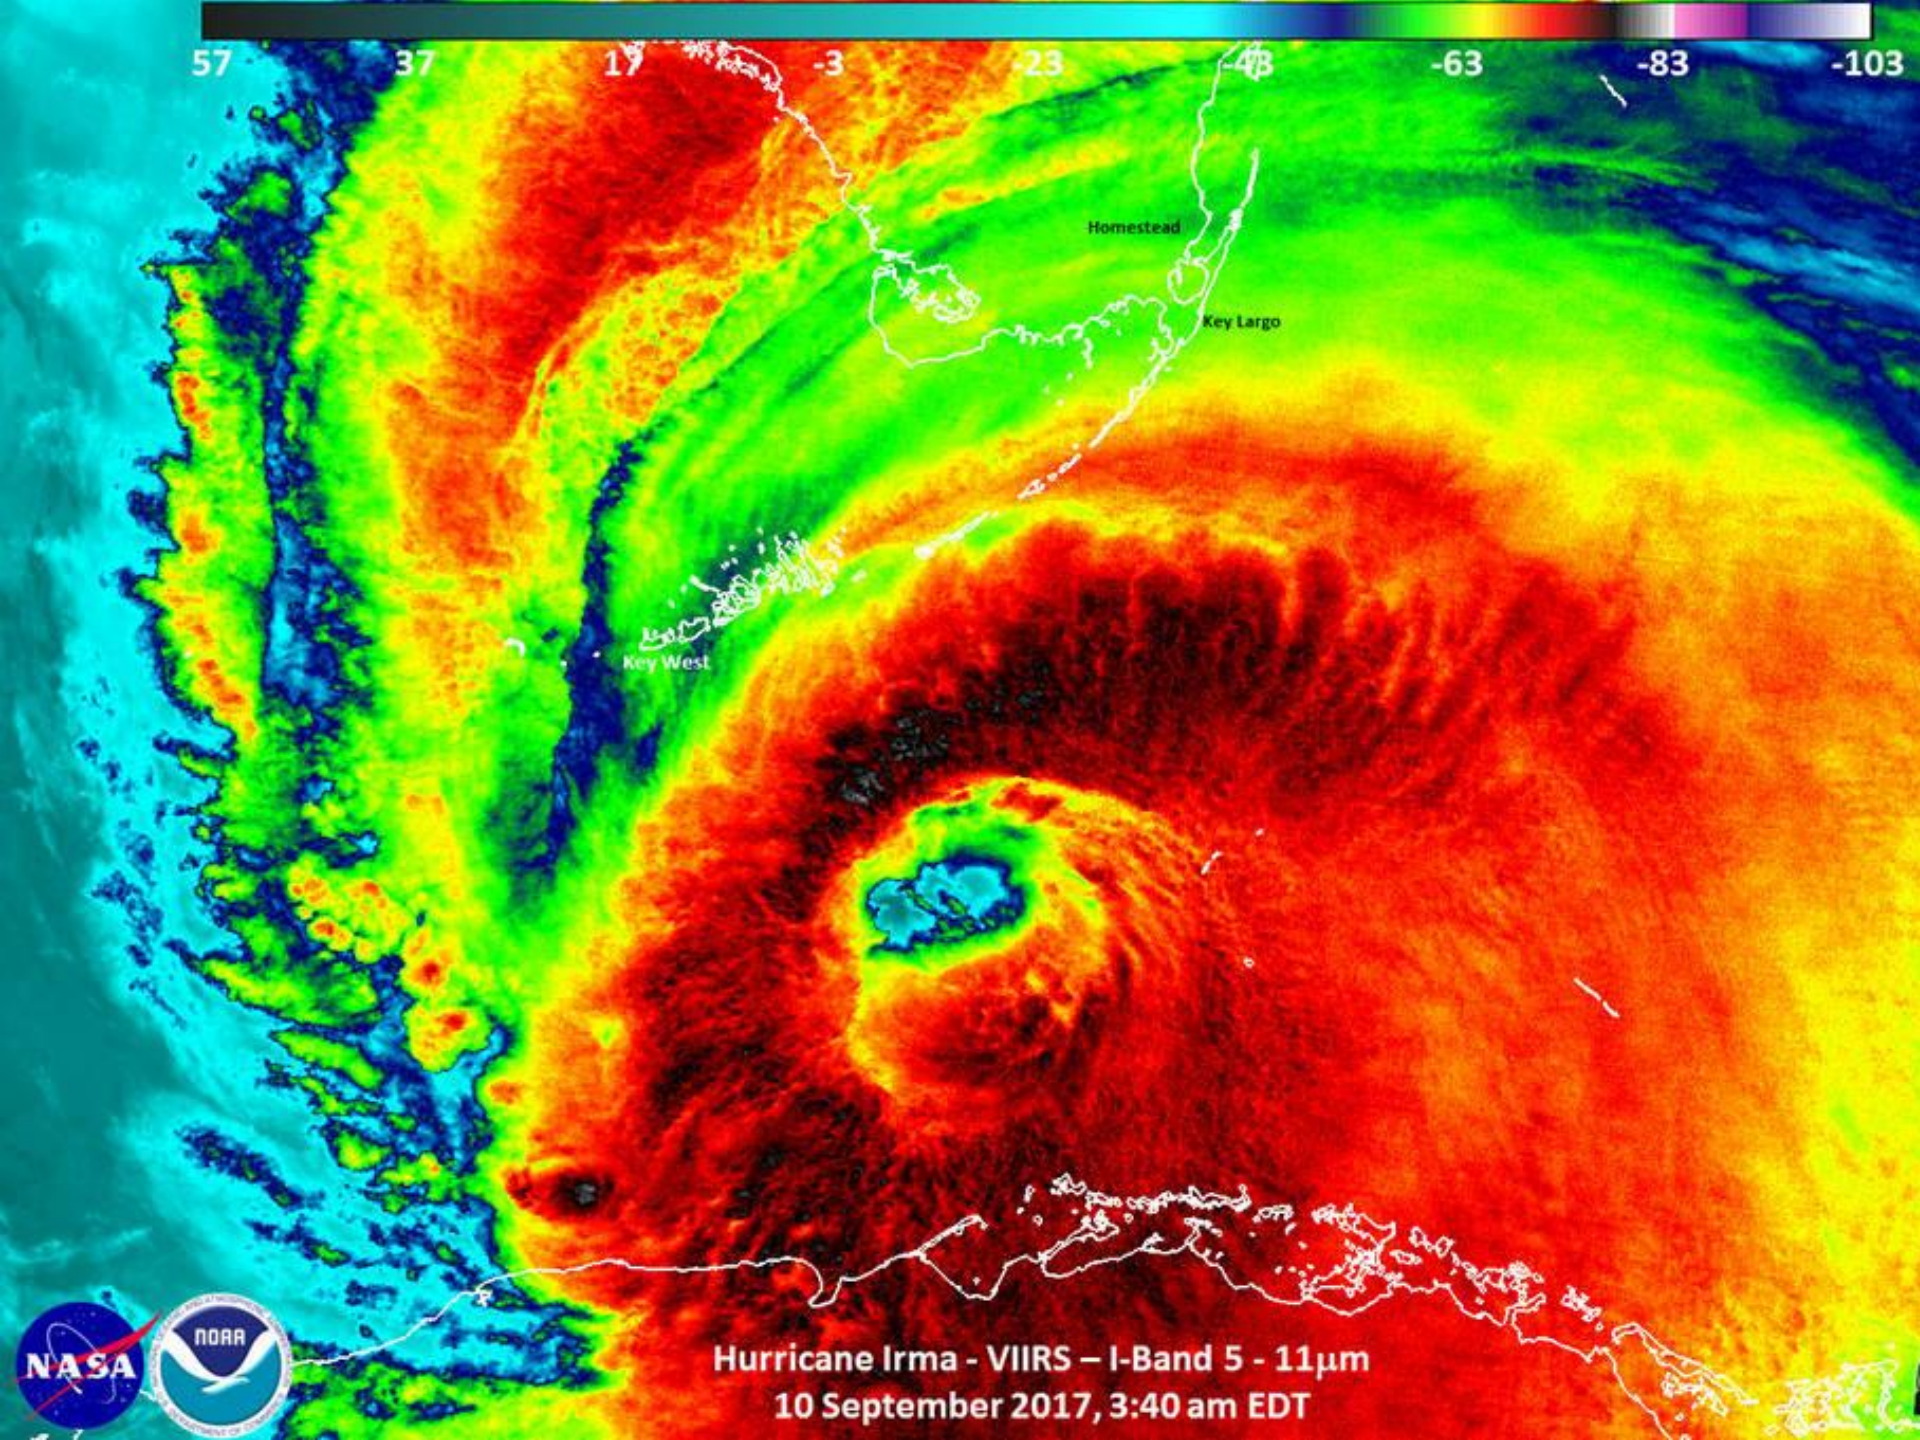
57
37
17
-3
-23
-43
-63
-83
-103
Homestead
Key Largo
Key West
NASA
NOAA
Hurricane Irma - VIIRS – I-Band 5 - 11µm
10 September 2017, 3:40 am EDT

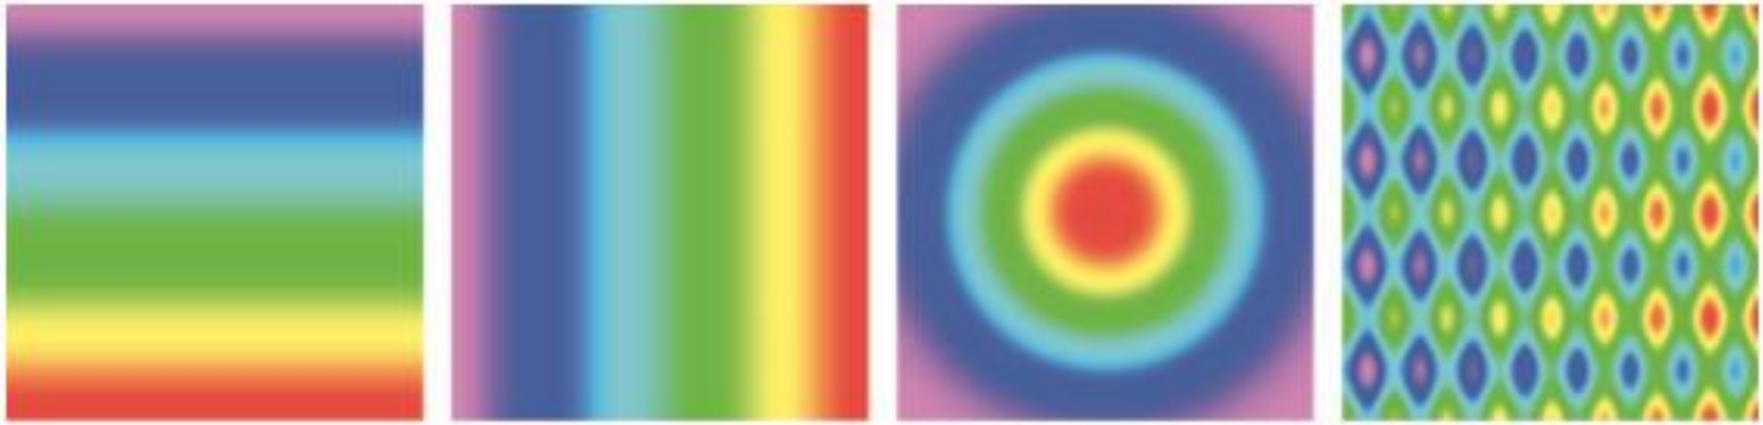

Rainbow color map considered harmful

Borland and Taylor

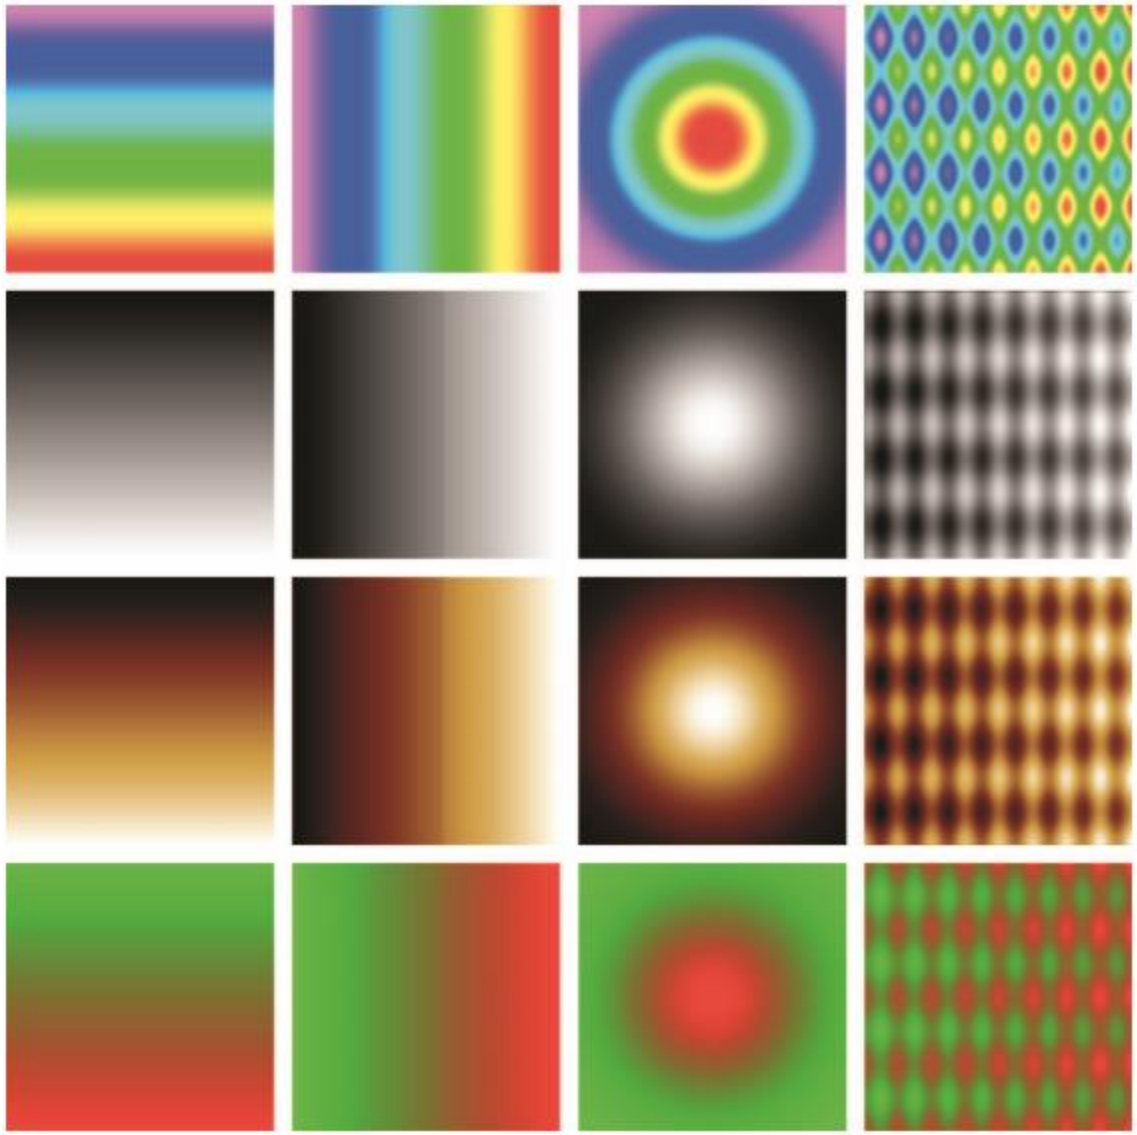

“Intuitive” color space?

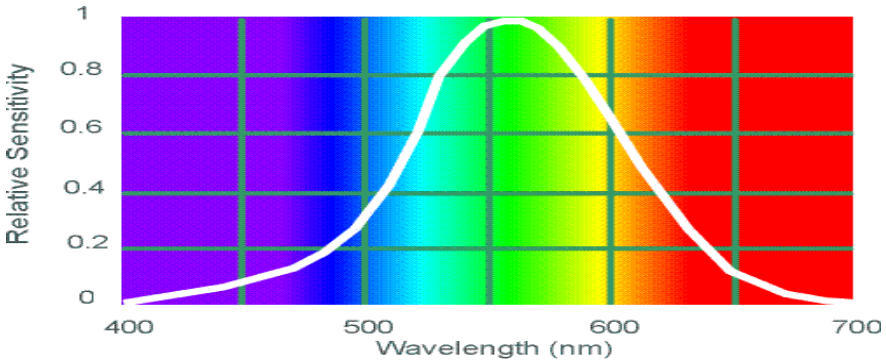

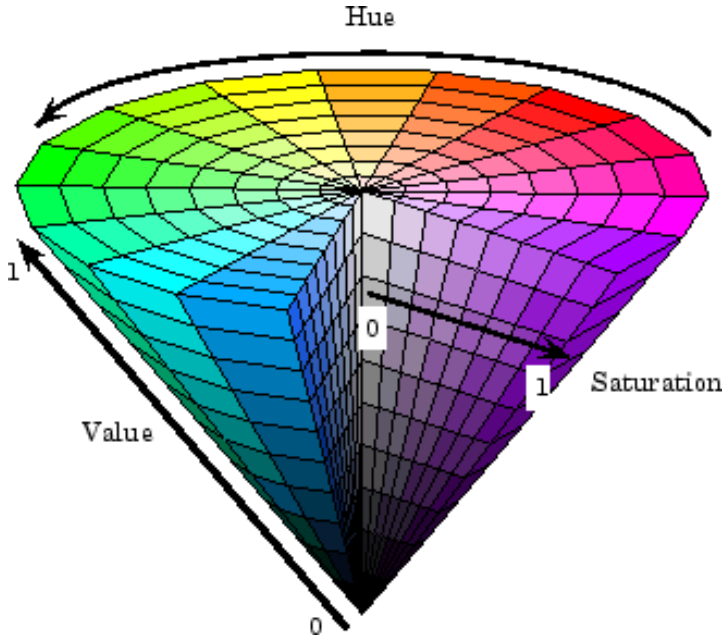

Wait a minute…

WHY DOES COLOR LOOK LIKE IT MAPS SMOOTHLY TO A CIRCLE?

‘Color’ != position on EM spectrum

Our cells induce color perception
by interpreting spectra.

Most mammals are dichromats:

- Lack ‘L’ cone; cannot distinguish green-red
- 1% of men (protanopia color blindness)

Trichromaticity *evolved.*
No implicit reason for effect of
extra cone to be linear.

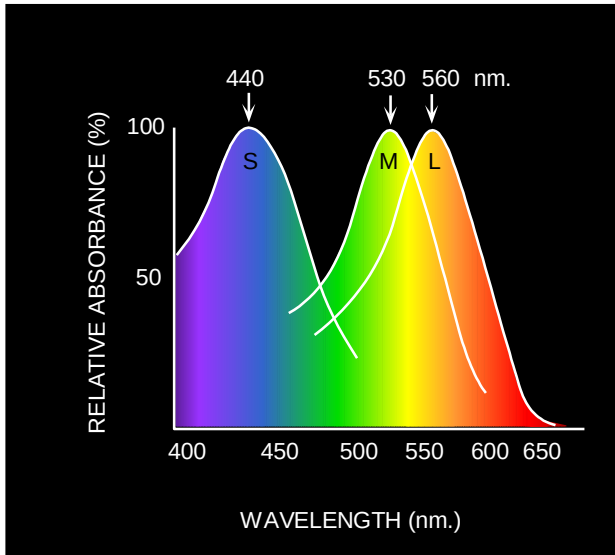

Thanks to Cam Allen-Lloyd

‘Color’ != position on EM spectrum

Many different ways to parameterize color.
Ask Prof. Thomas Serre for a qualified answer.

Or…

“When some primates started growing a *third* cone in their retinas, the old bipolar system remained, with the third cone adding a 2nd dimension of color encoding: red versus green. since color is now encoded in a 2d space, you find that you can draw a circle of colors in that space, which when you think about the fact that wavelength is 1d is really weird.”

- aggasalk, Reddit.

Thanks to Alexander Nibley

Held and Hein (1963)

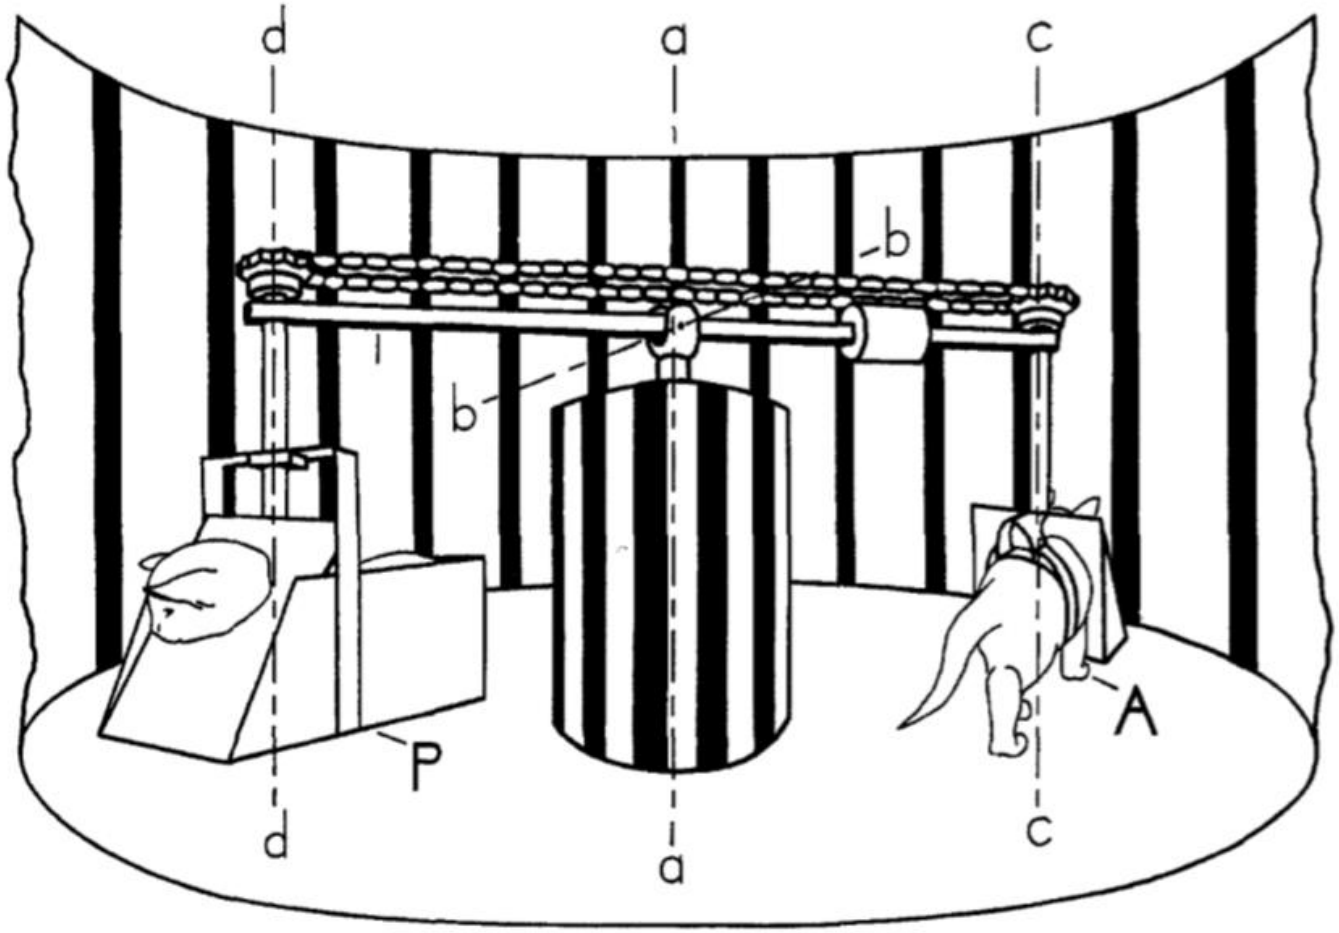

FIG. 1. Apparatus for equating motion and consequent visual feedback for an actively moving (A) and a passively moved (P) *S*.

A Bird's-Eye View of Nature's Hidden Order

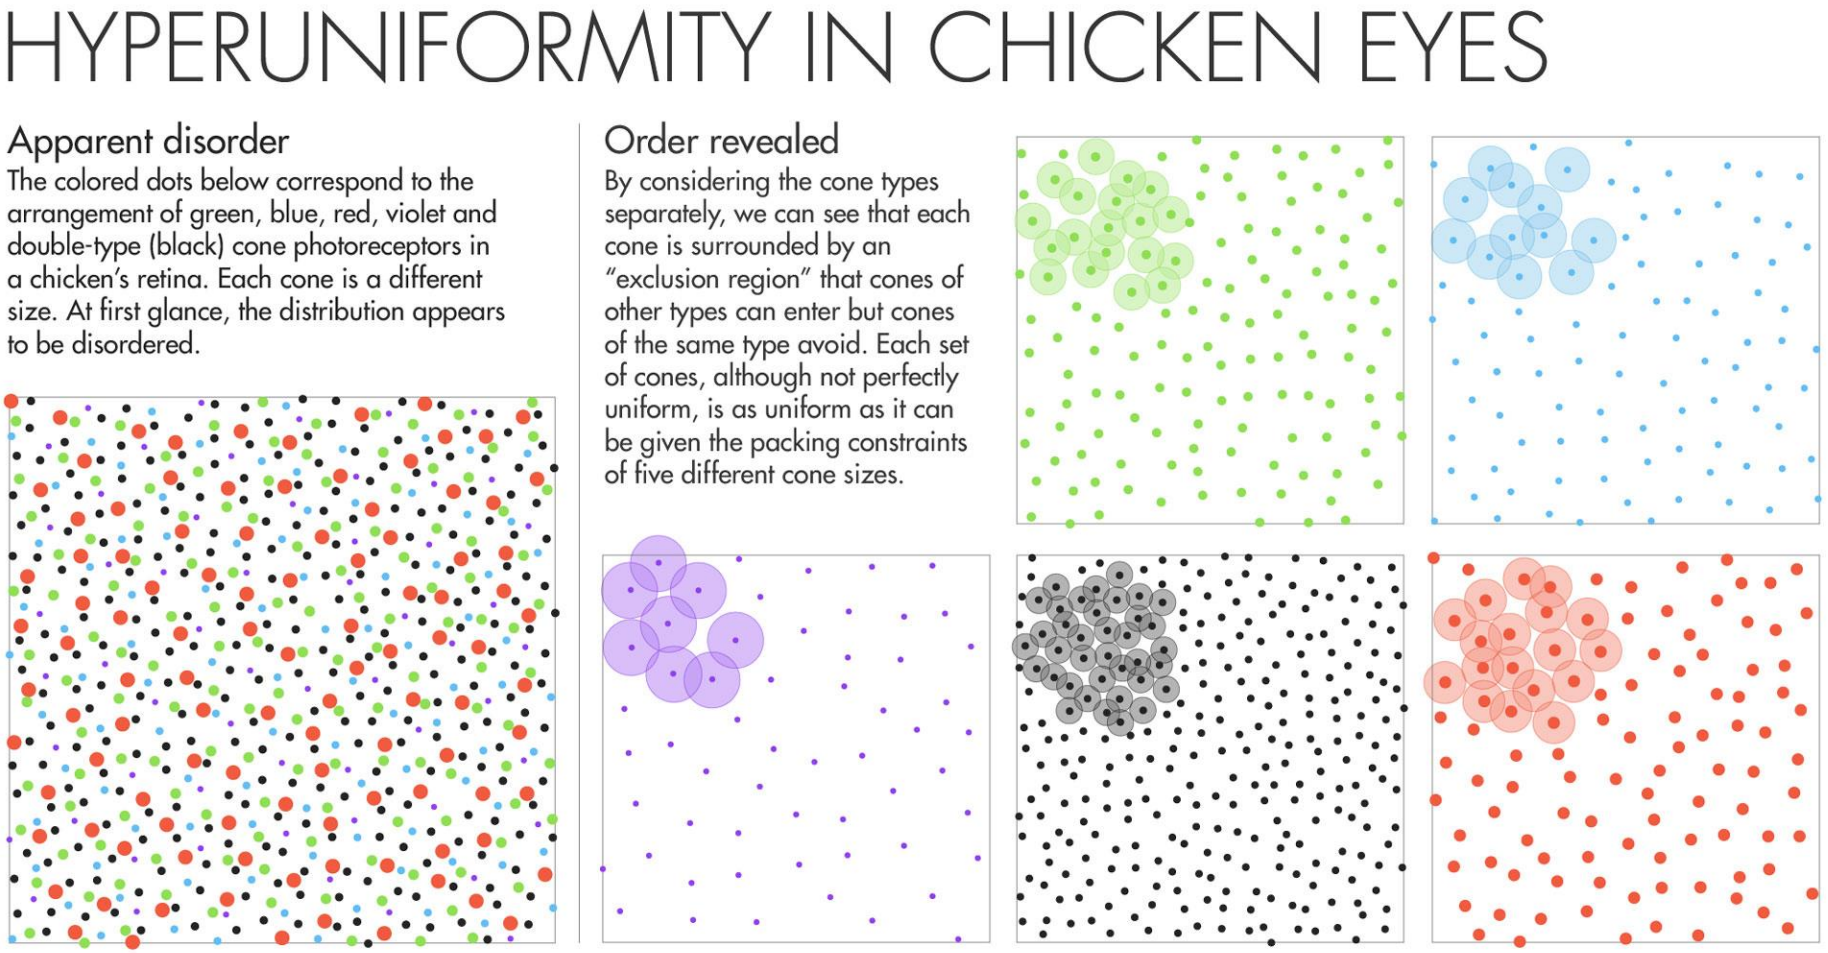

Natalie Wolchover, Quanta Magazine 2016