1950

Future Vision

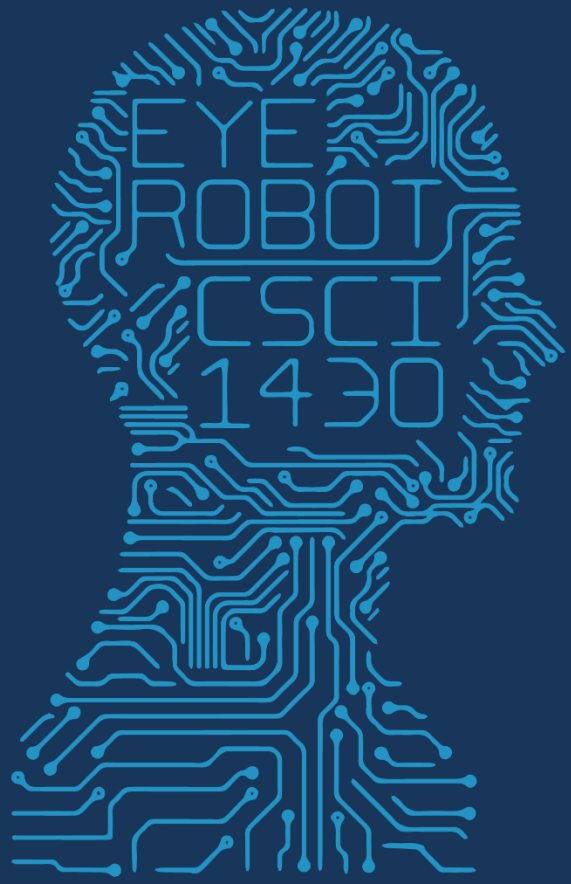

2017 MWF 1pm

Computer Vision

Bela Borsodi

Bela Borsodi

Waitlist

- We’ll let you know as soon as we can.
- Biggest issue is TAs

CS 143 – James Hays

- Many materials, courseworks, based from him + previous TA staff – serious thanks!

Textbook

Computer Vision: Algorithms and Applications

© 2010 Richard Szeliski, Microsoft Research

http://szeliski.org/Book/

Textbook

- Can I get a PDF of this book?

  No, our contract with MIT Press forbids distribution of too easily copied electronic formats of the book.

- Why are you using HTML format for the web version of the book?

  This format is a sort of weak DRM required by our contract with MIT Press. It's intended to discourage unauthorized copying/editing of the book.

- What is the best way to print the HTML format?

  Printing seems to work best printing directly from the browser, using Chrome. Other browsers do not work as well.

Class experience

- Linear algebra
- Probability

- Graphics course?
- Vision/image processing course before?
- Machine learning?

WHAT IS AN IMAGE?

```
>> I = rand(256,256);
```

Think-Pair-Share:

- What is this? What does it look like?
- Which values does it take?
- How many values can it take?

- Is it an image?

```
>> I = rand(256,256);
>> imshow(I);
```

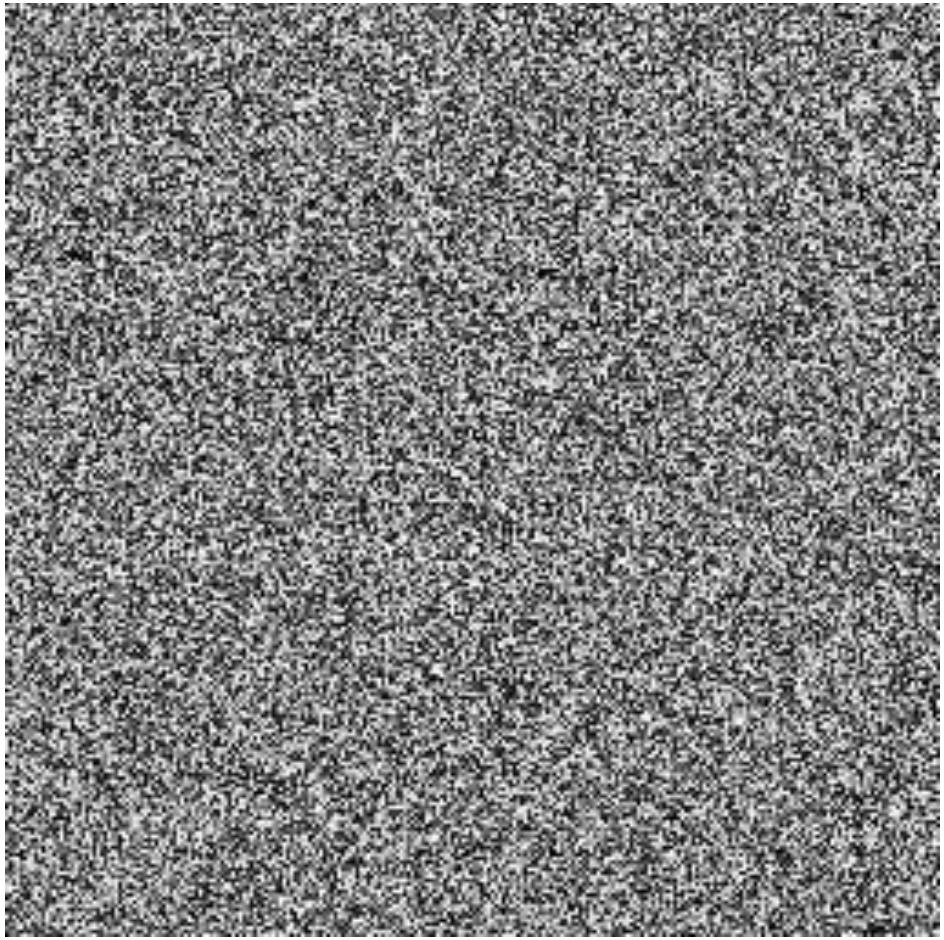

Dimensionality of an Image

- @ 8bit = 256 values ^ 65,536
  - Computer says ‘Inf’ combinations.

- Some depiction of all possible scenes would fit into this memory.

Dimensionality of an Image

- @ 8bit = 256 values ^ 65,536
  - Computer says ‘Inf’ combinations.

- Some depiction of all possible scenes would fit into this memory.
- Computer vision as making sense of an extremely high-dimensional space.
  - Subspace of ‘natural’ images.
  - Deriving low-dimensional, explainable models.

What is each part of an image?

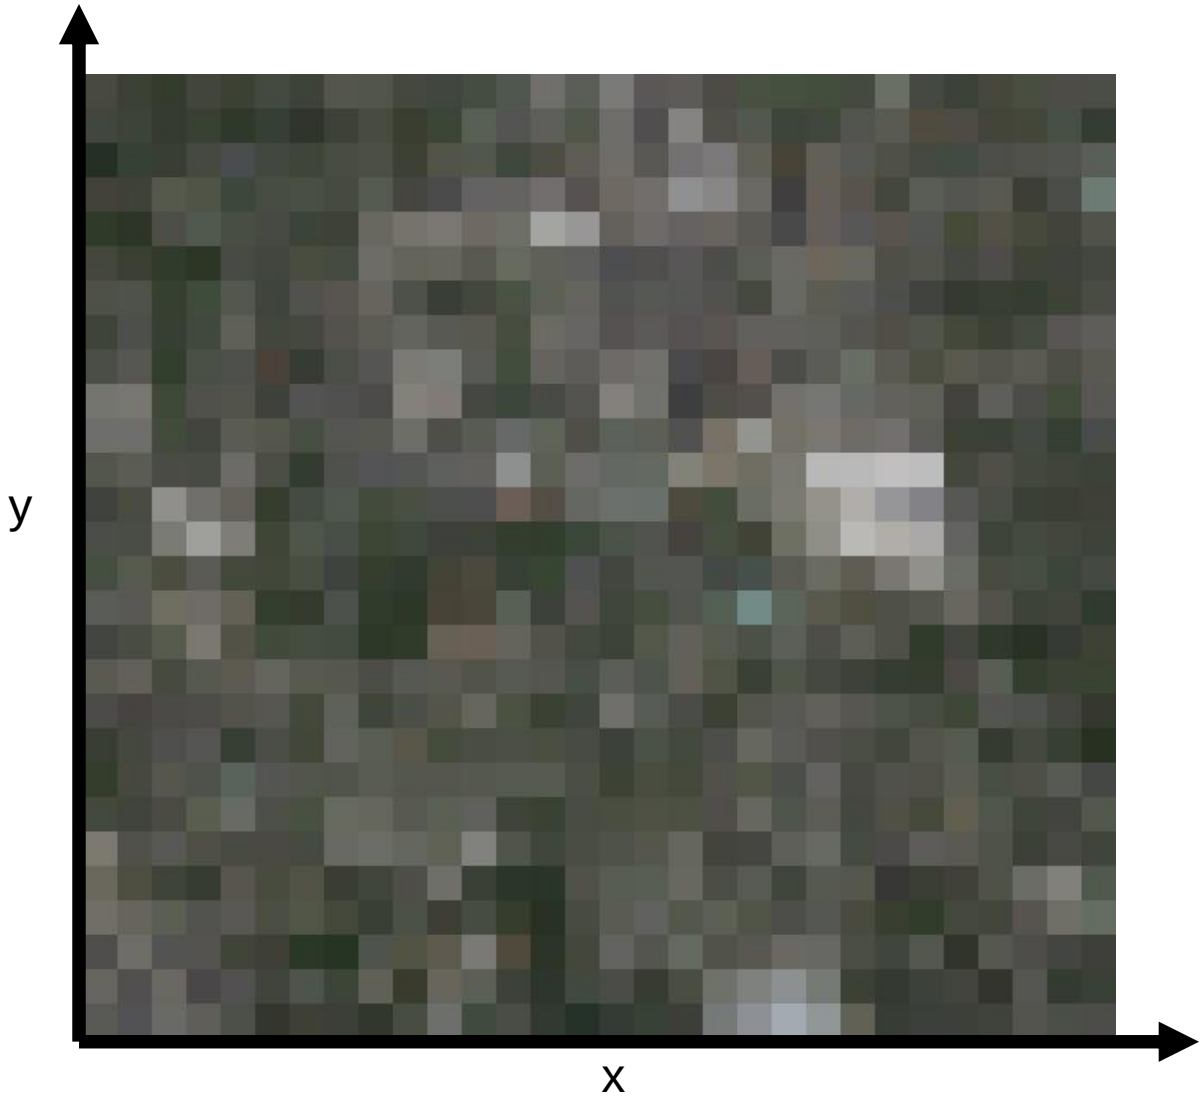

What is each part of an image?

- Pixel -> picture element

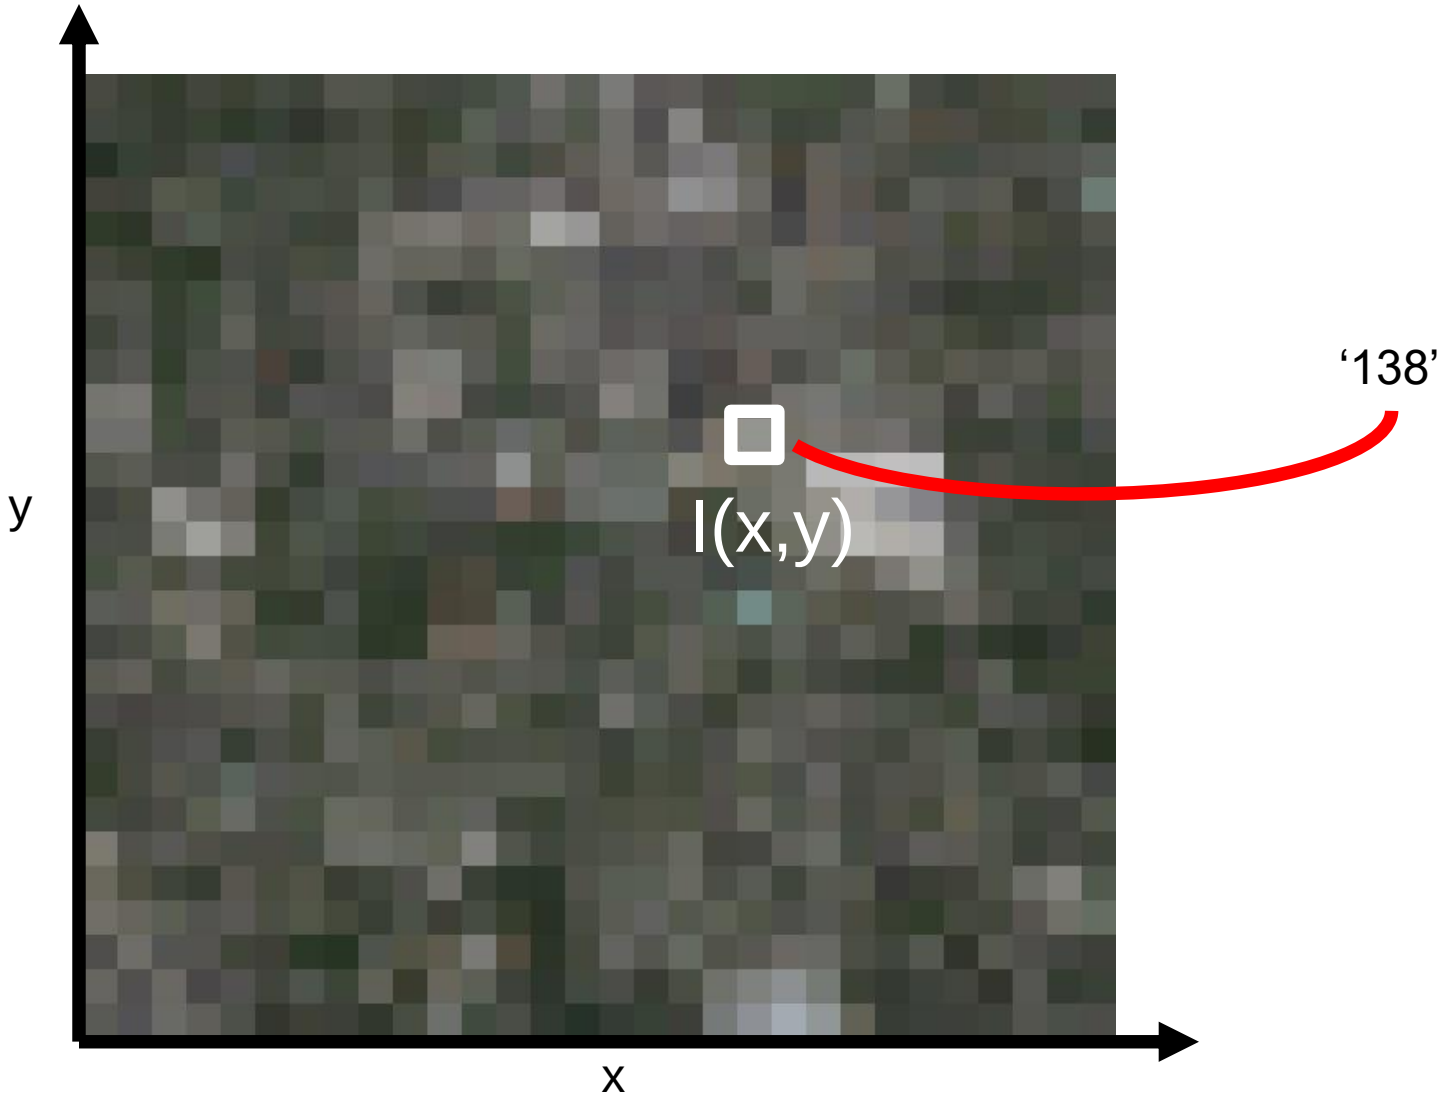

Image as a 2D sampling of signal

- Signal: function depending on some variable with physical meaning.

- Image: sampling of that function.
  - 2 variables: xy coordinates
  - 3 variables: xy + time (video)
  - ‘Brightness’ is the value of the function for visible light

- Can be other physical values too: temperature, pressure, depth …

Example 2D Images

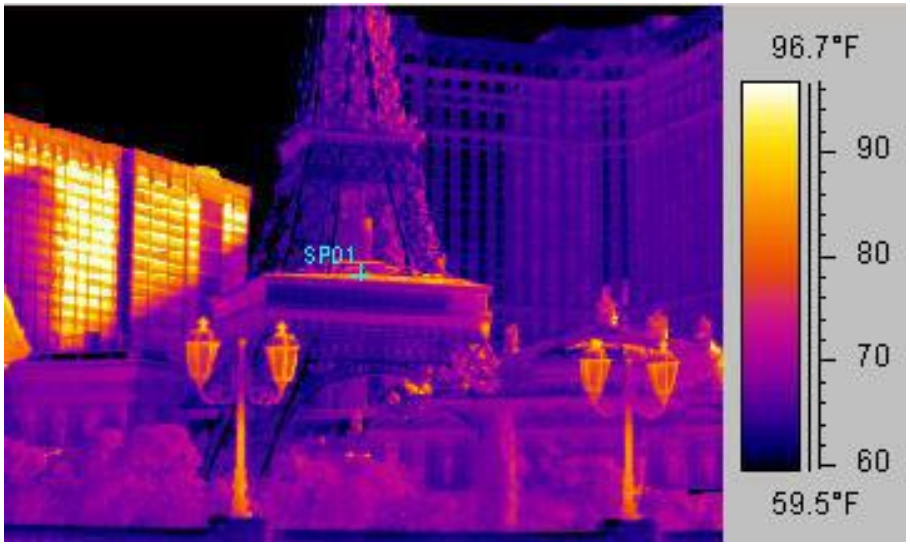

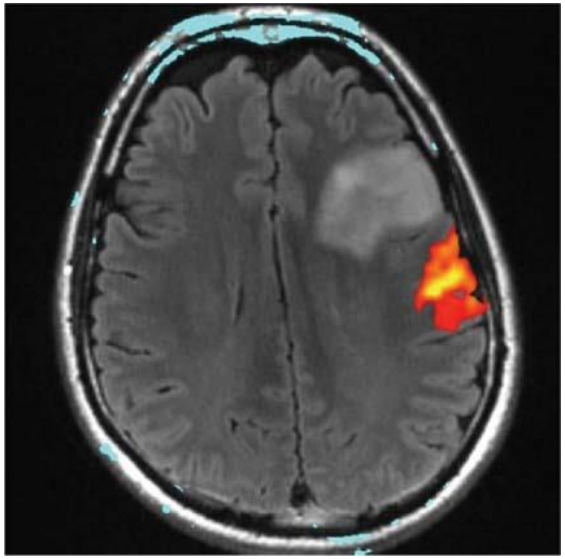

Sampling in 1D

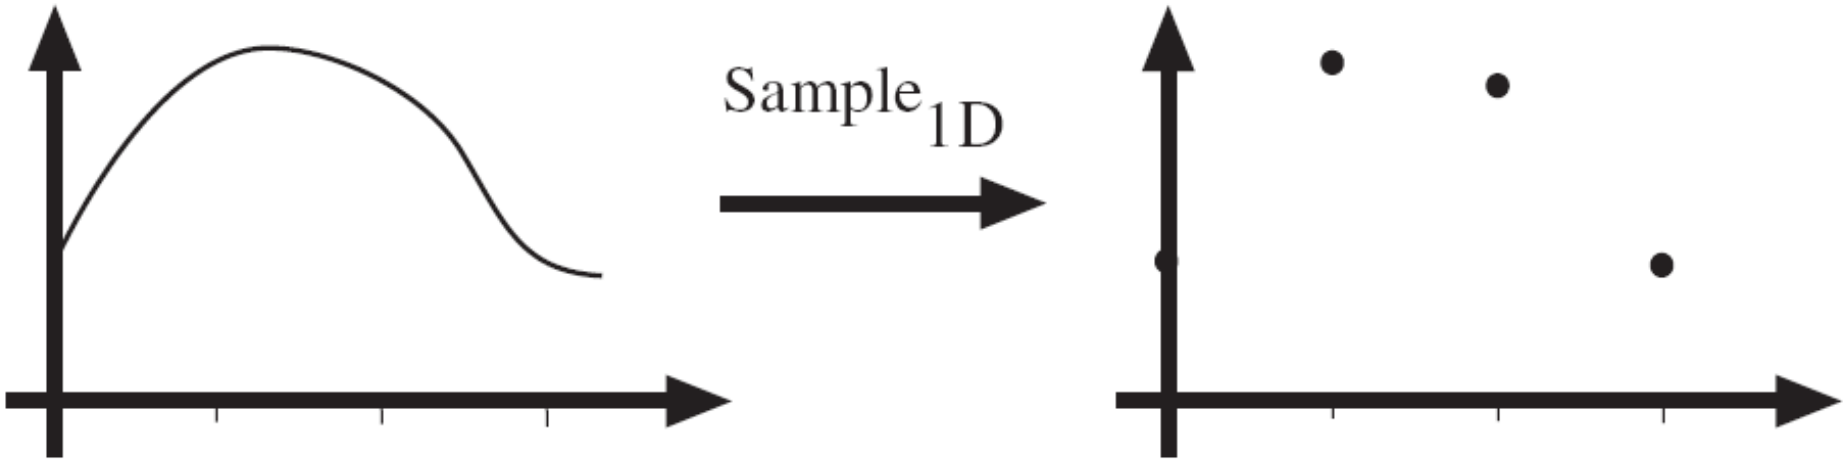

- Sampling in 1D takes a function, and returns a vector whose elements are values of that function at the sample points.

Sampling in 2D

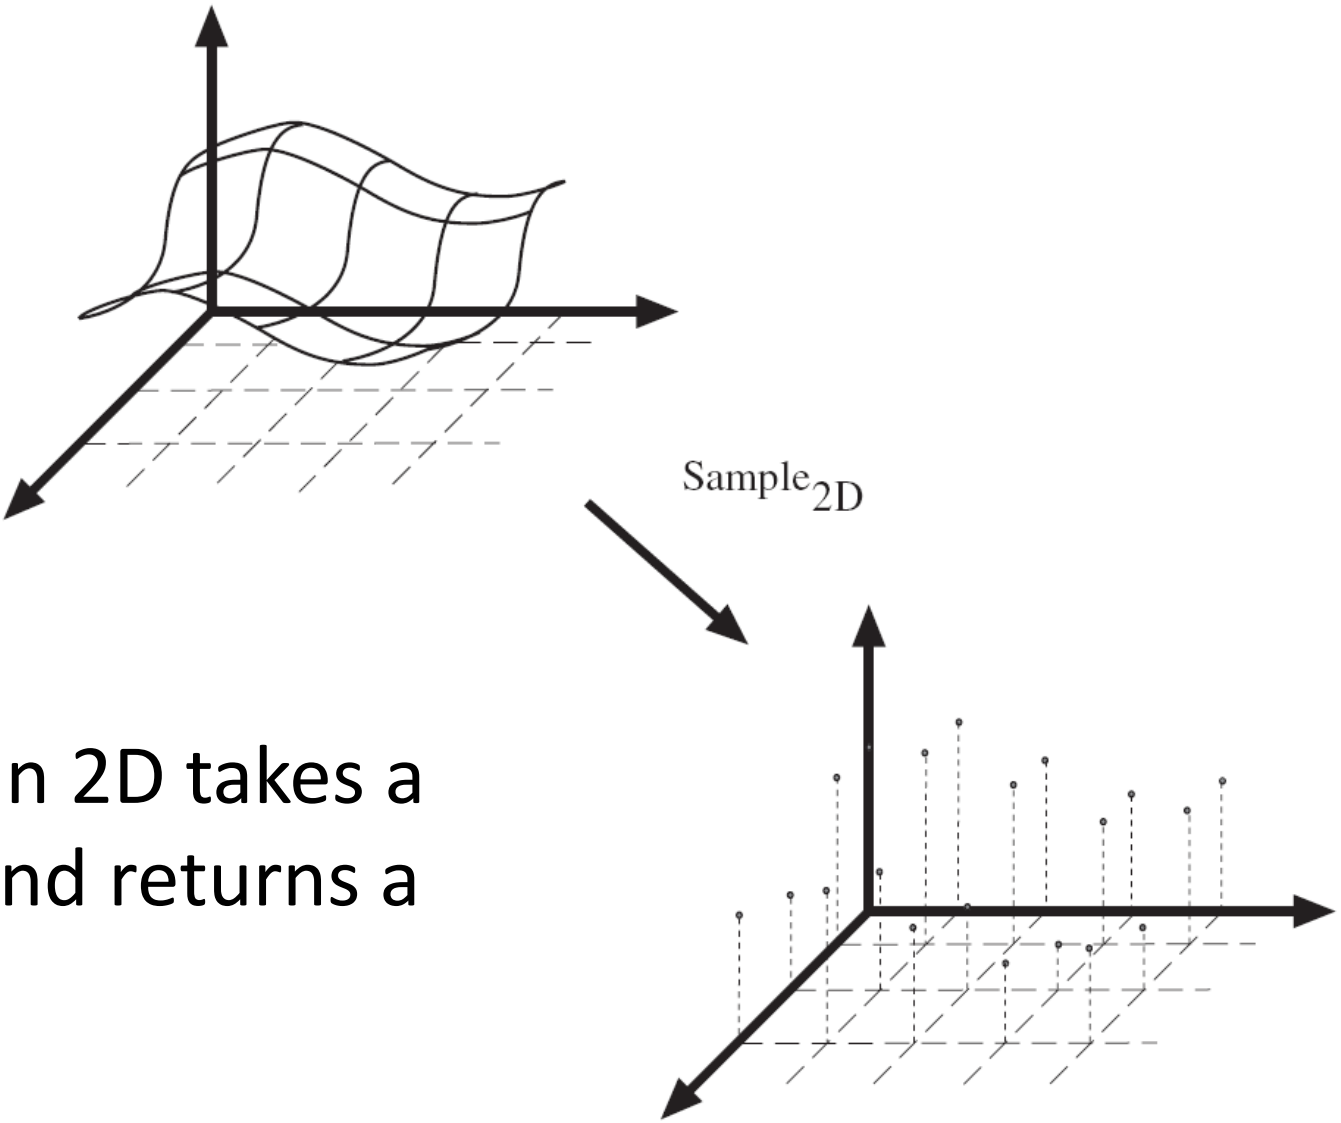

- Sampling in 2D takes a function and returns a matrix.

Grayscale Digital Image

Brightness or intensity

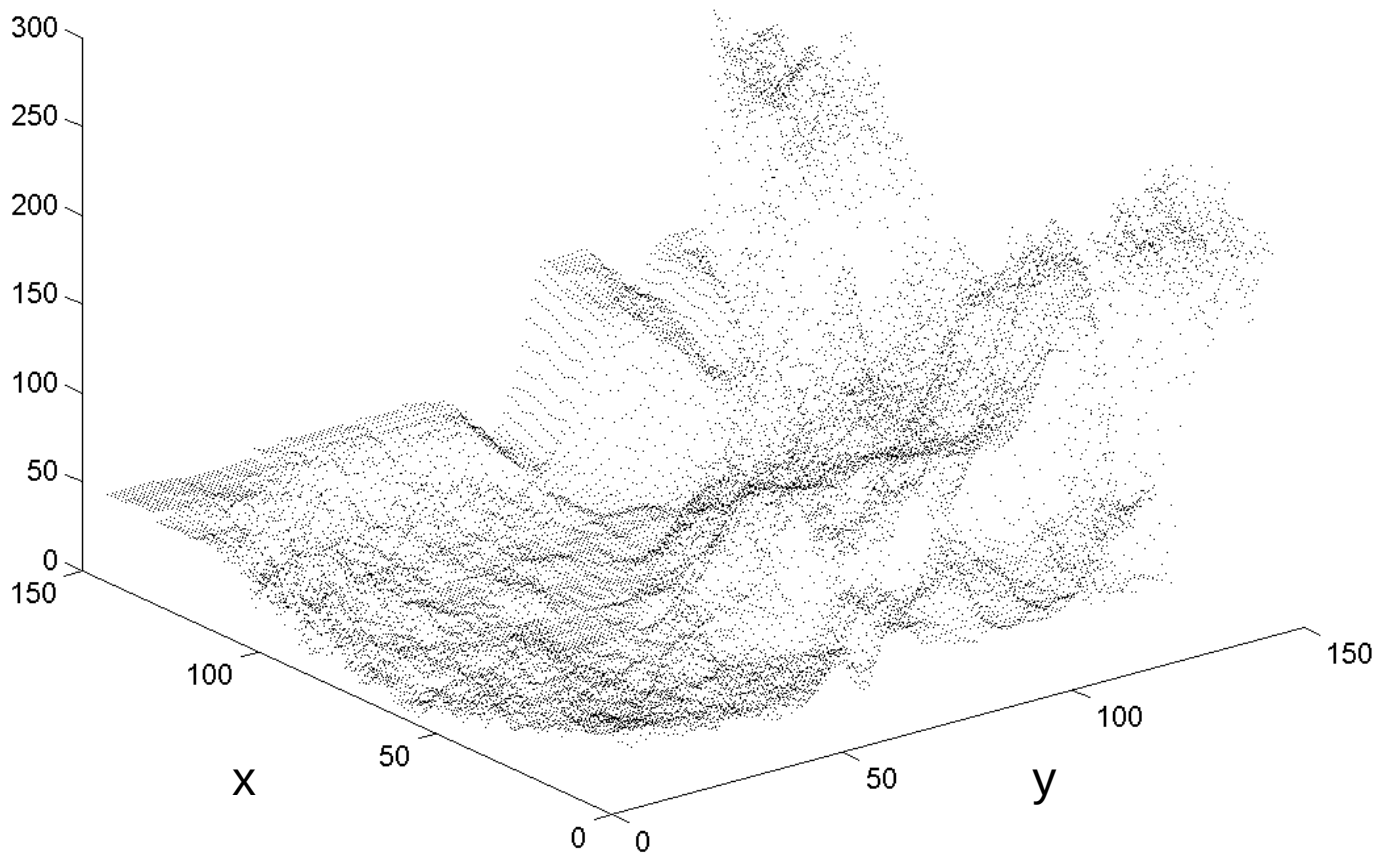

What is each part of a photograph?

- Pixel -> picture element

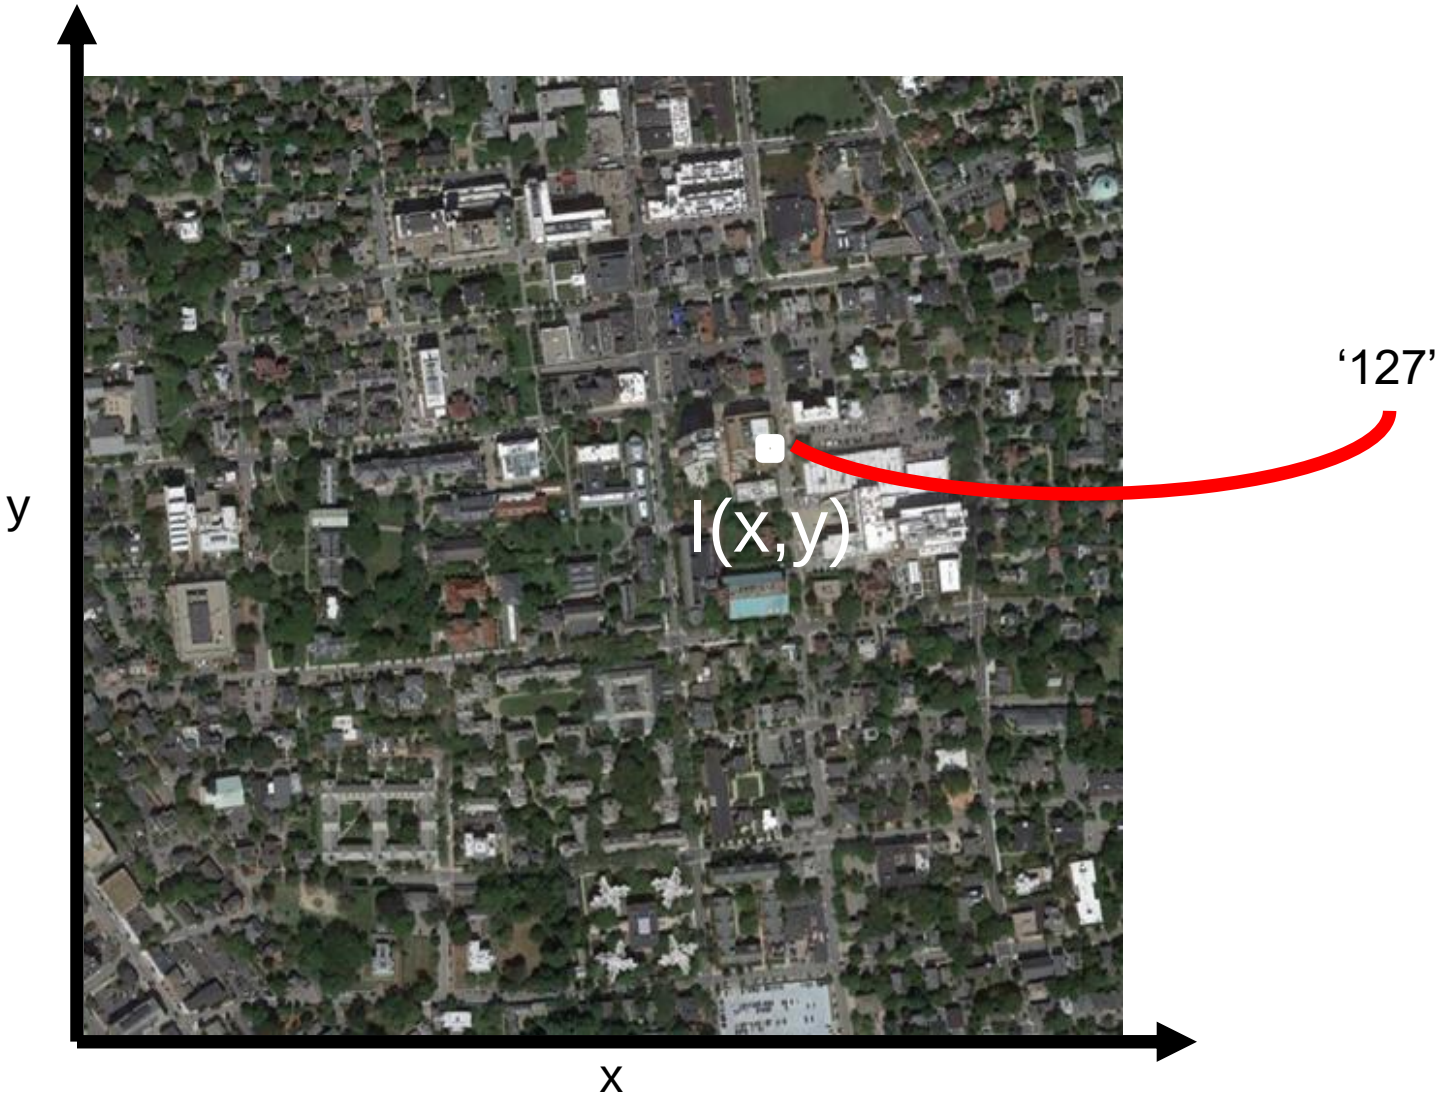

Integrating light over a range of angles.

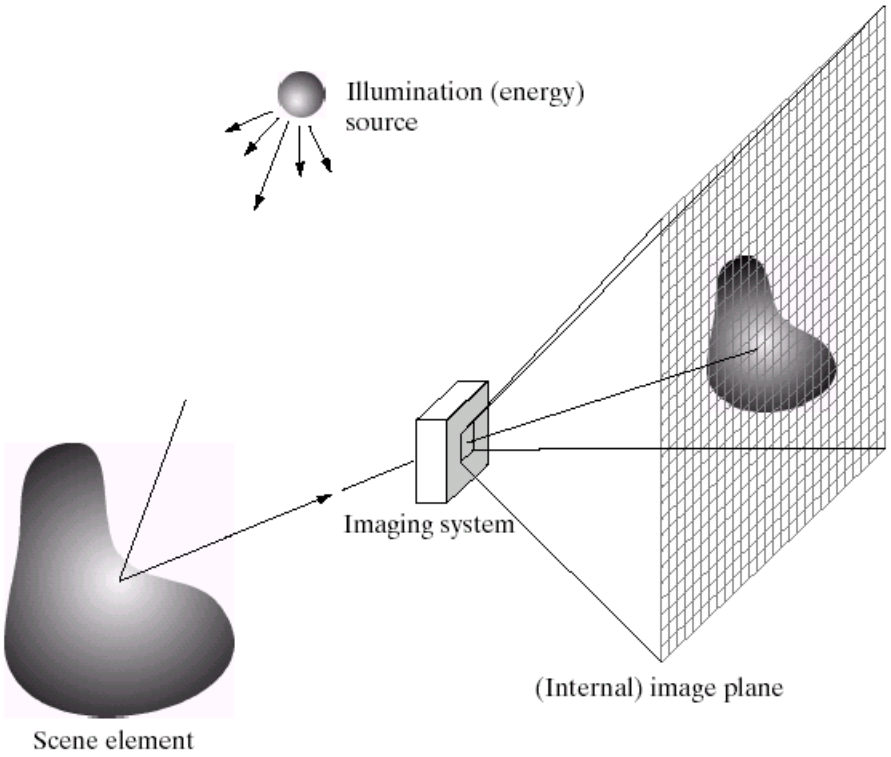

Camera Sensor

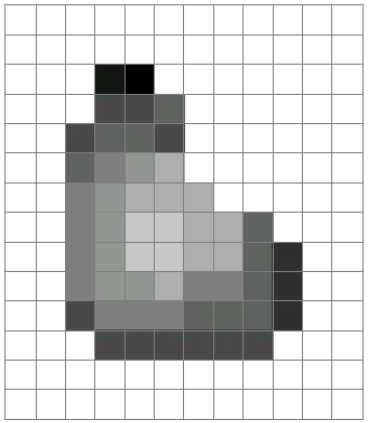

Output Image

Resolution – geometric vs. spatial resolution

Both images are ~500x500 pixels

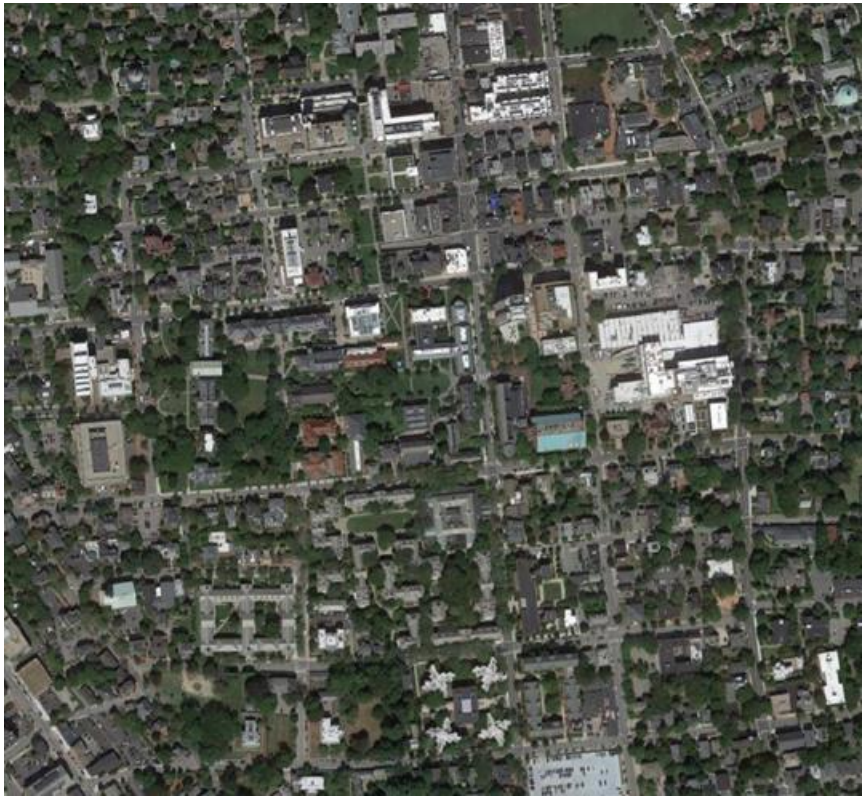

Quantization

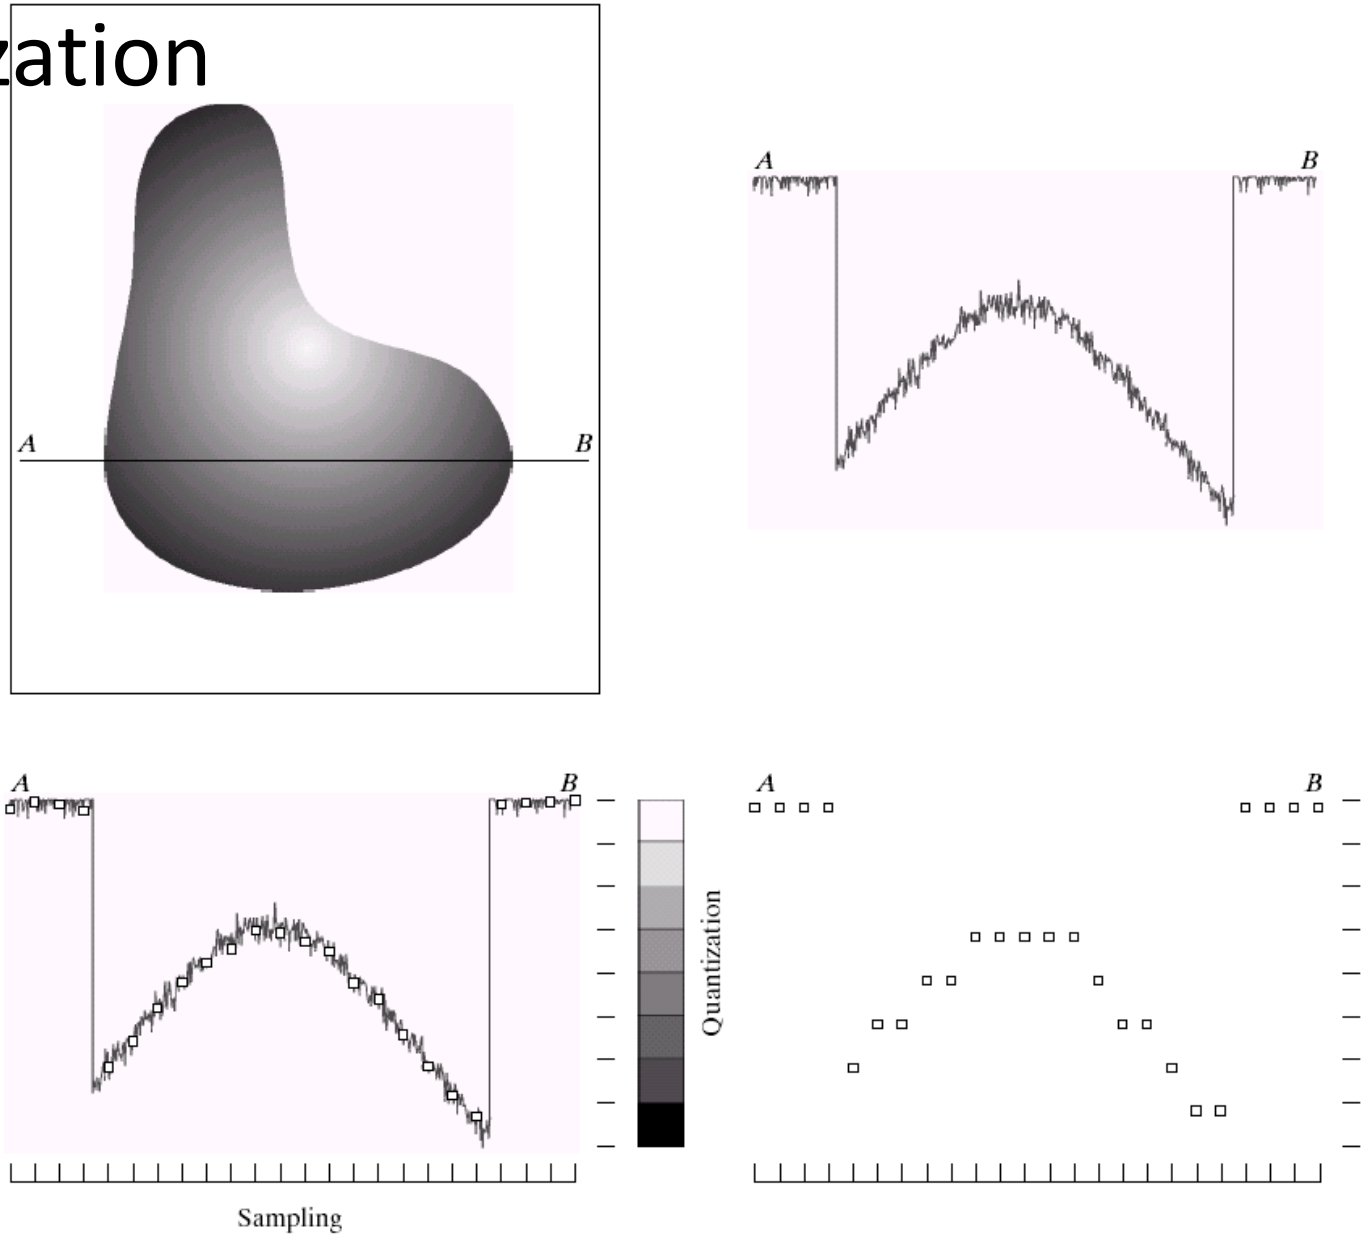

a b
c d

FIGURE 2.16 Generating a digital image. (a) Continuous image. (b) A scan line from *A* to *B* in the continuous image, used to illustrate the concepts of sampling and quantization. (c) Sampling and quantization. (d) Digital scan line.

Quantization Effects – Radiometric Resolution

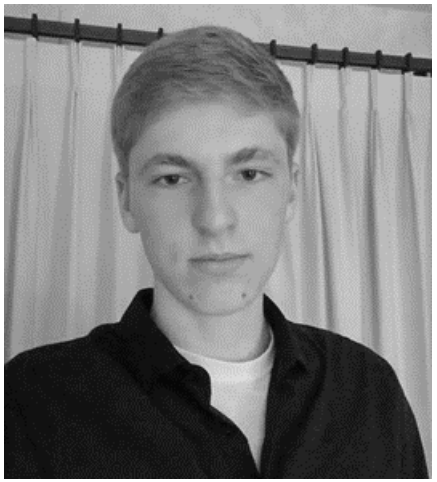

8 bit – 256 levels

4 bit – 16 levels

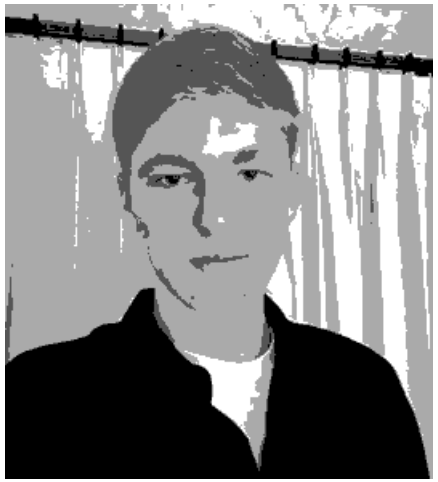

2 bit – 4 levels

1 bit – 2 levels

Color

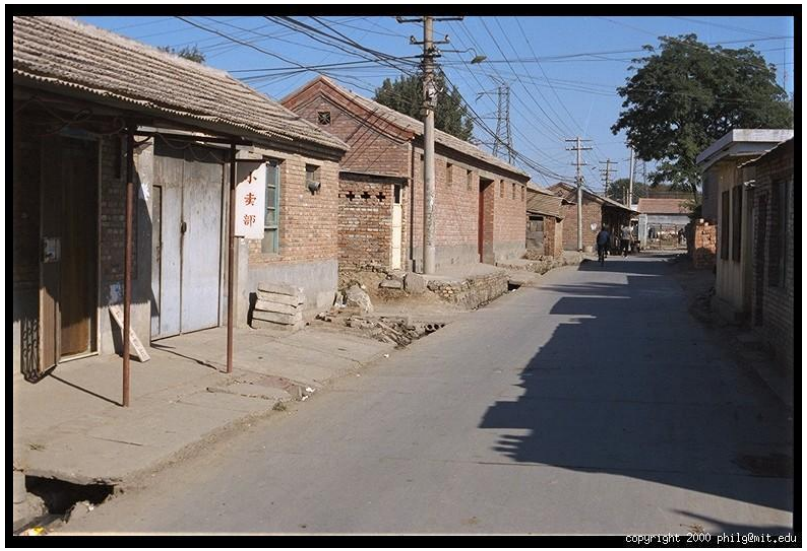

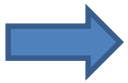

James Hays

Images in Matlab

- NxM RGB "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

- imread(filename) returns a uint8 image (values 0 to 255)
  - Convert to double format (values 0 to 1) with im2double

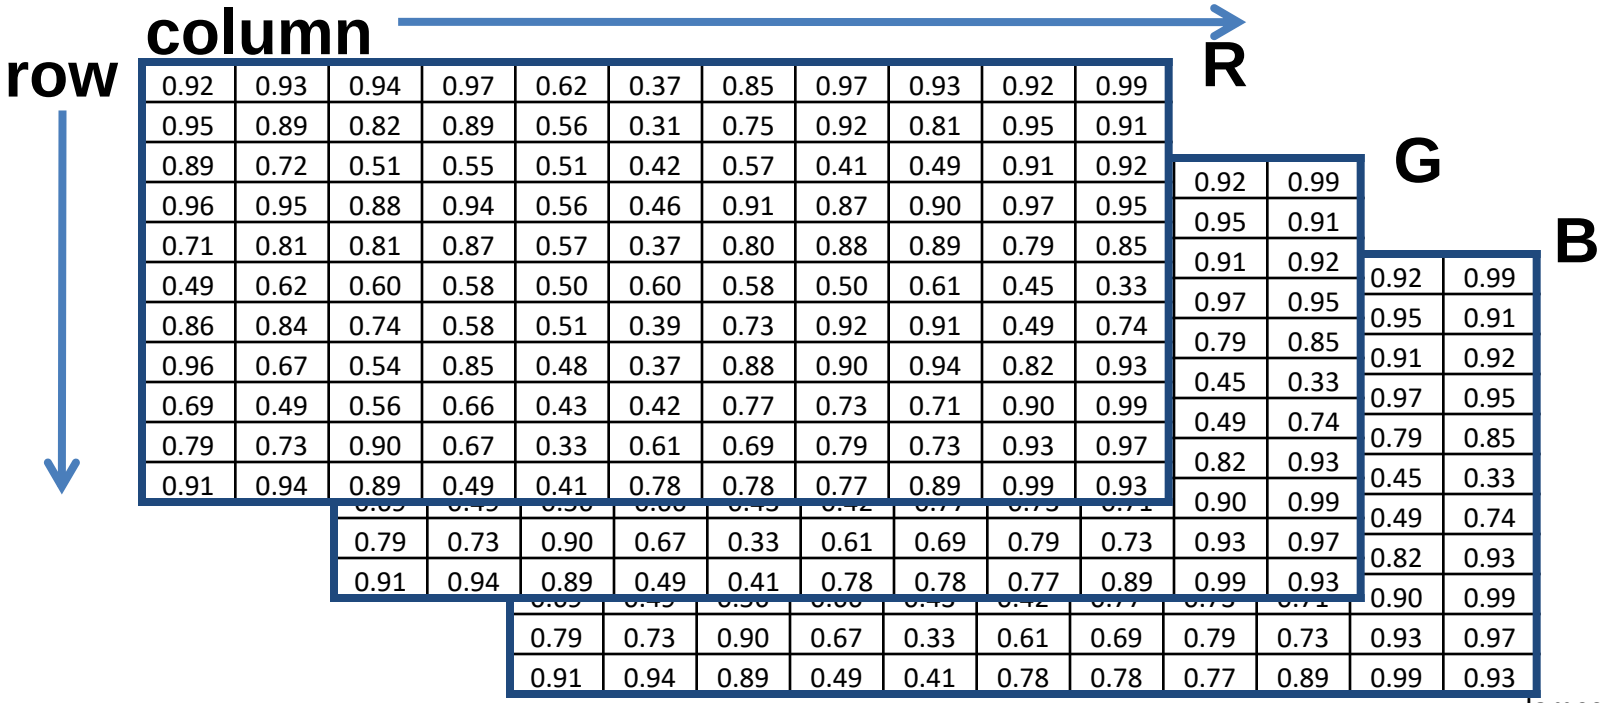

But what is color?

ANATOMY

The Eye

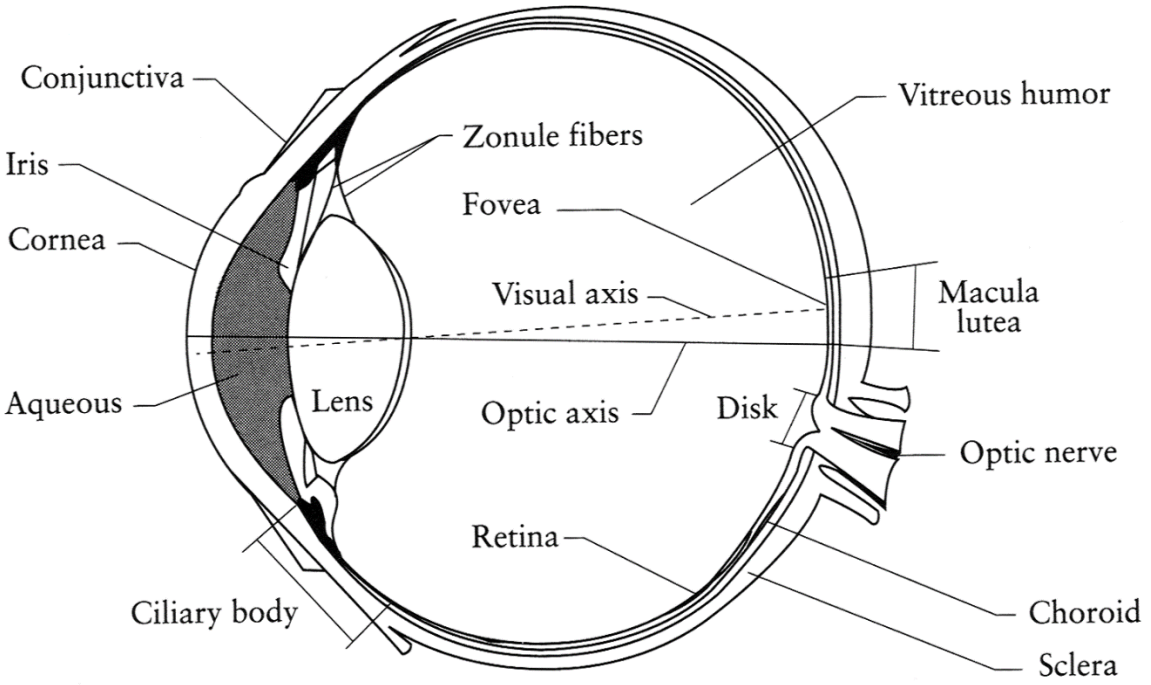

- The human eye is a camera
  - **Iris** - colored annulus with radial muscles
  - **Pupil** - the hole (aperture) whose size is controlled by the iris
  - What's the sensor?
    - photoreceptor cells (rods and cones) in the **retina**

Slide by Steve Seitz

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

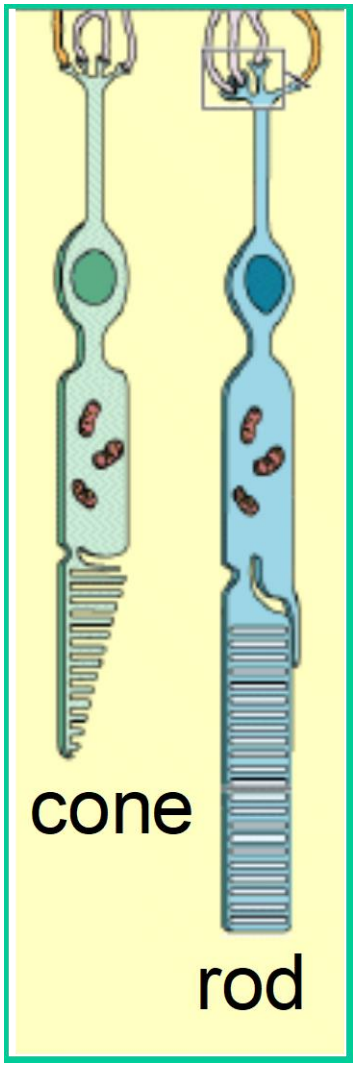

Distribution of Rods and Cones

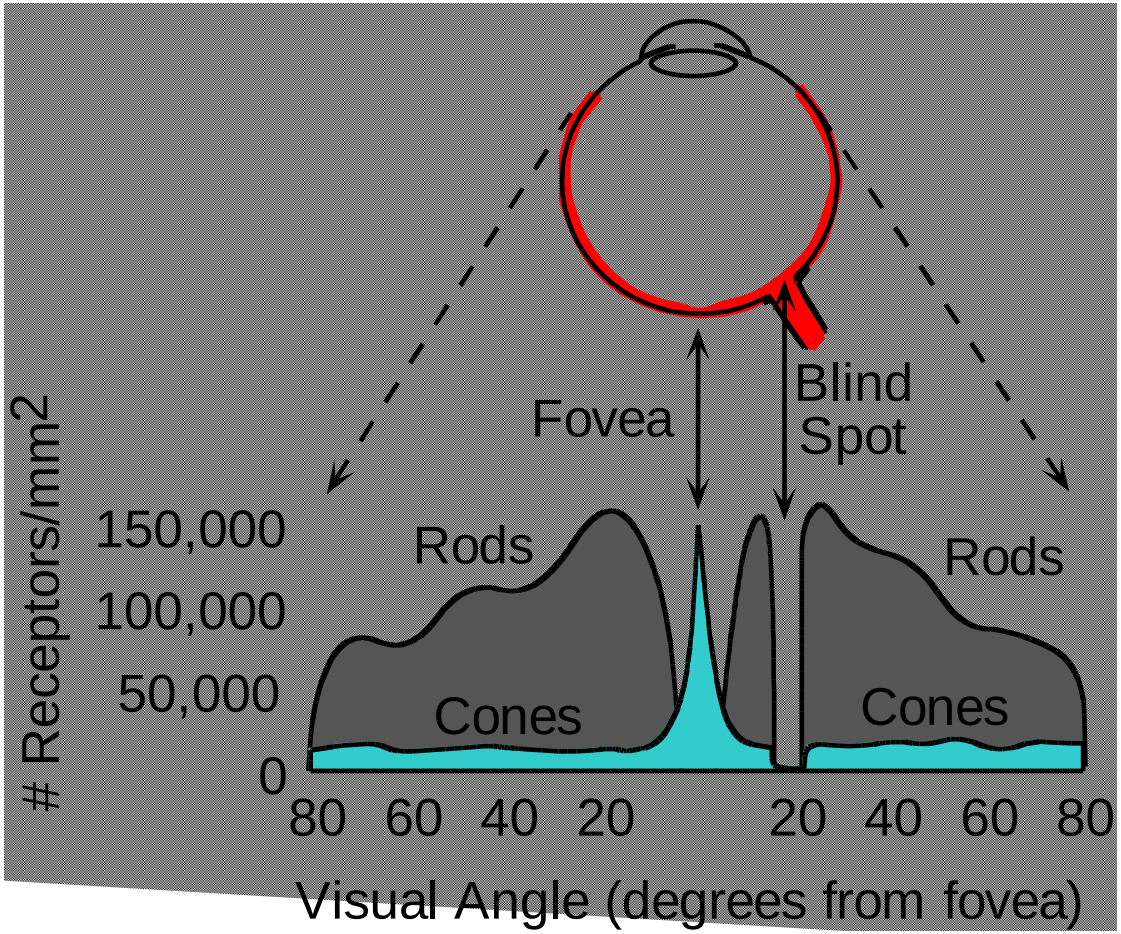

Night Sky: why are there more stars off-center?
Averted vision: http://en.wikipedia.org/wiki/Averted_vision

Rod / Cone sensitivity

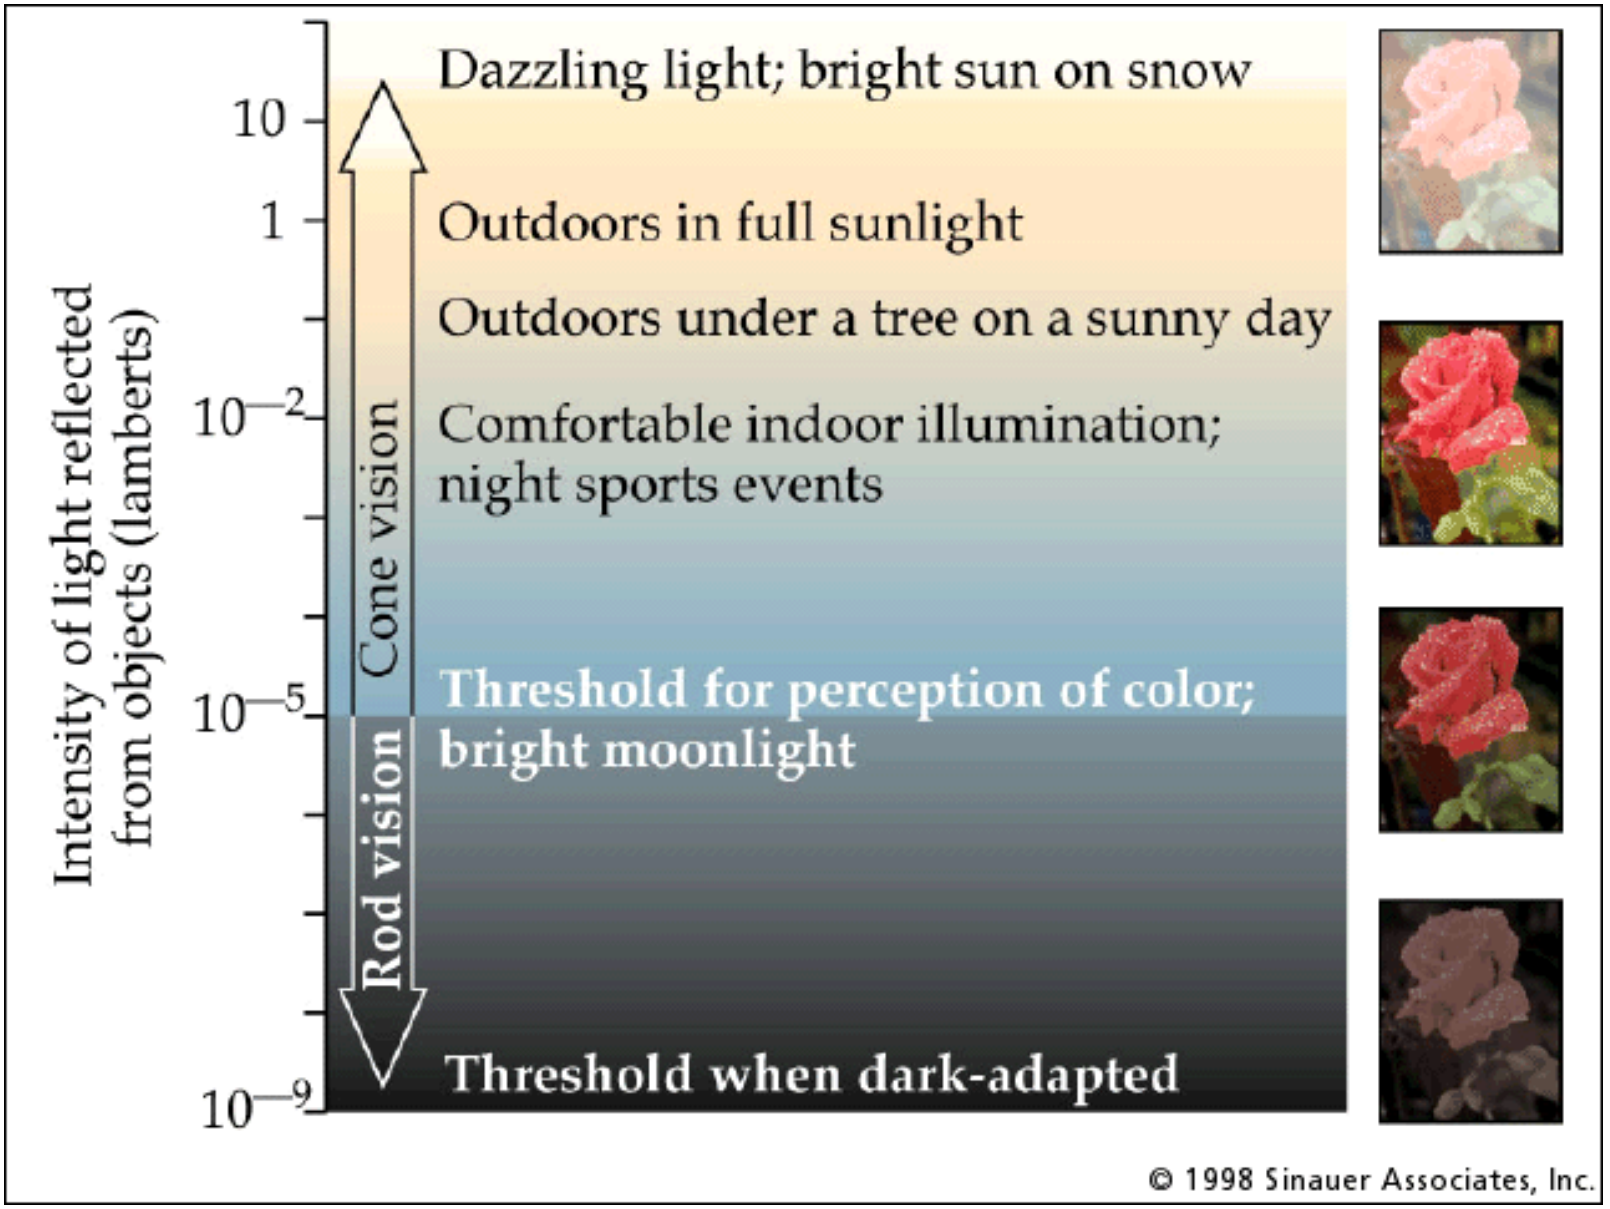

Electromagnetic Spectrum

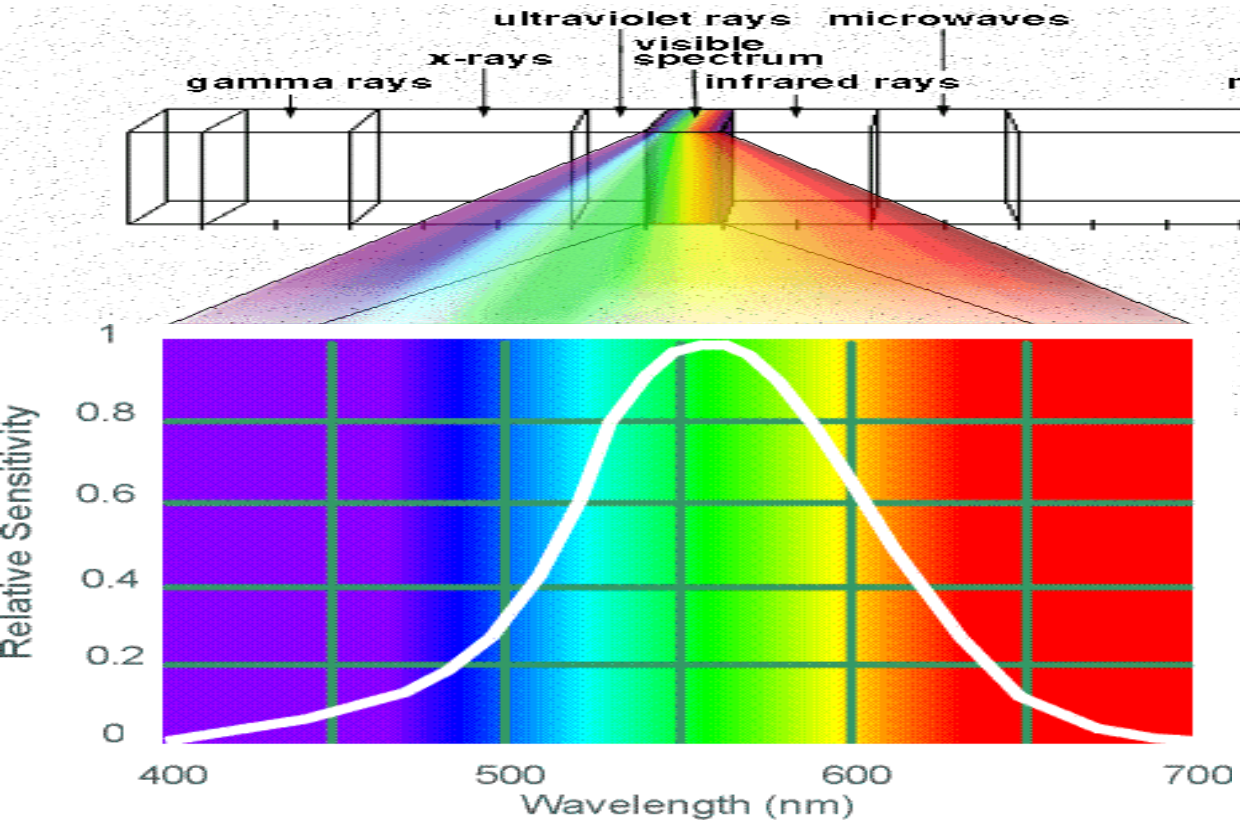

Human Luminance Sensitivity Function

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

Physiology of Color Vision

Three kinds of cones:

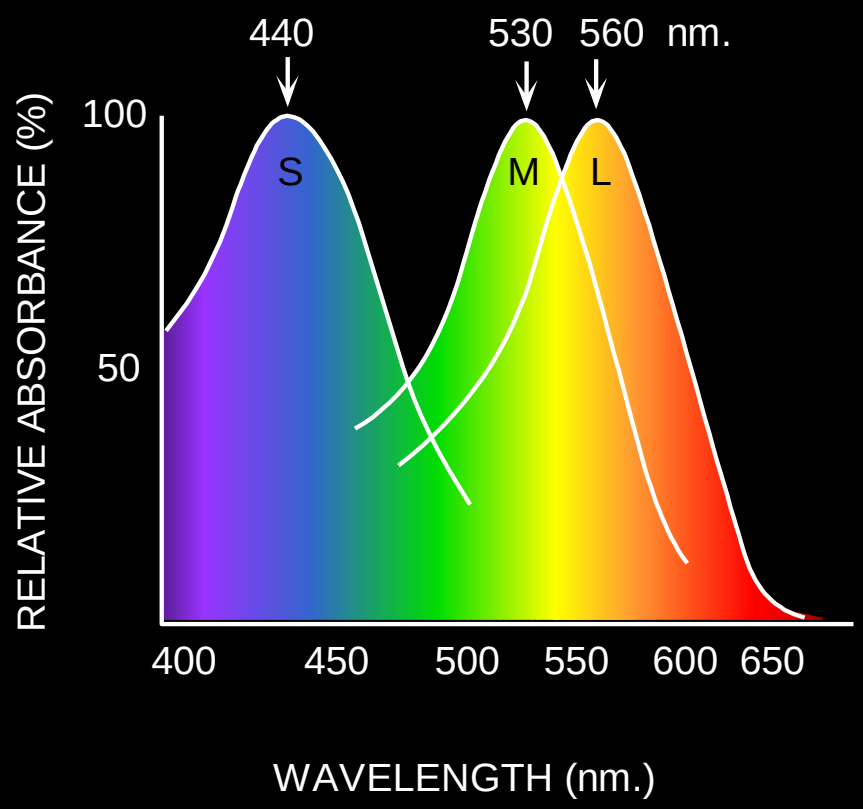

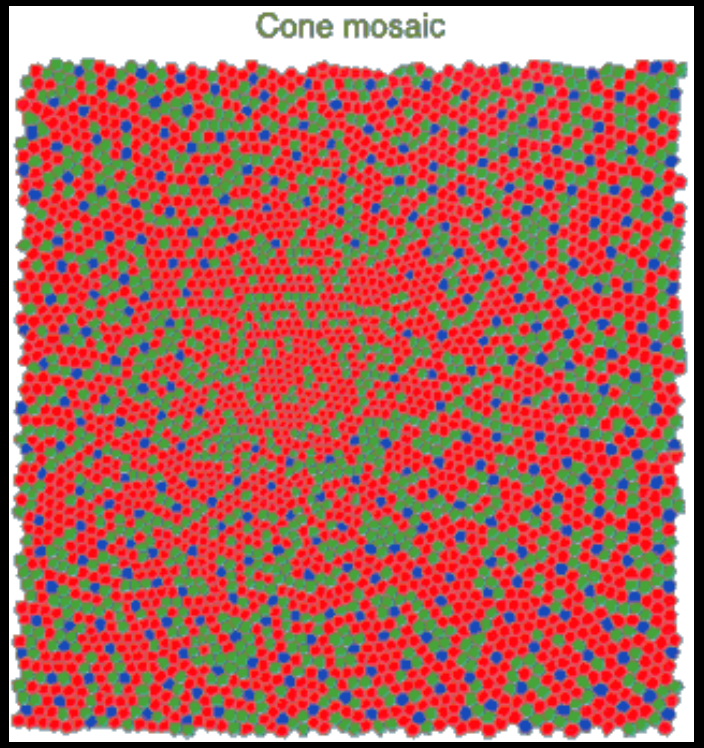

The Physics of Light

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.

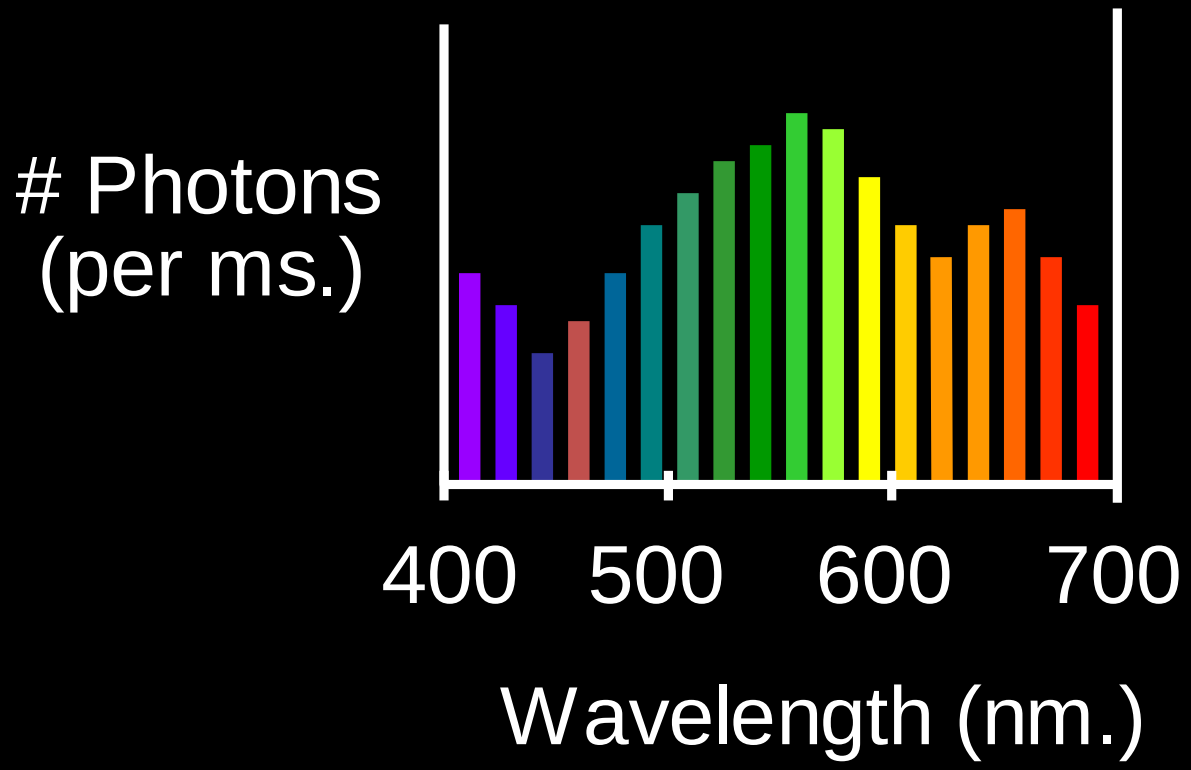

The Physics of Light

Some examples of the spectra of light sources

A. Ruby Laser

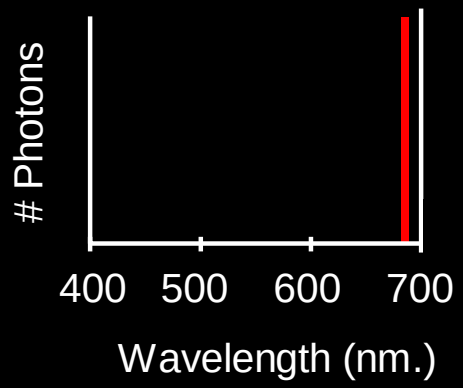

B. Gallium Phosphide Crystal

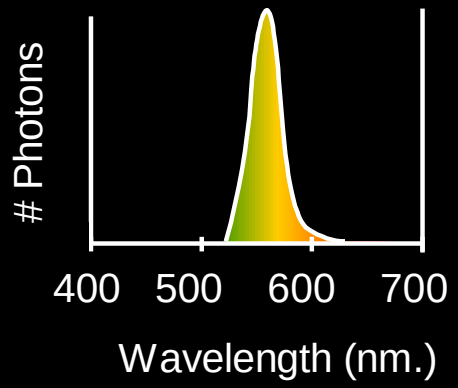

C. Tungsten Lightbulb

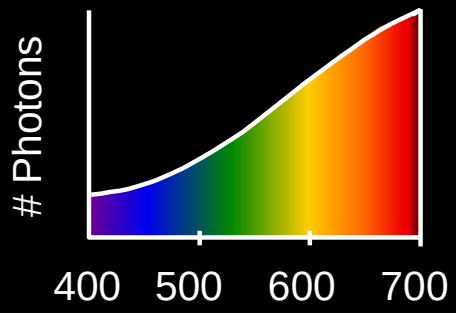

D. Normal Daylight

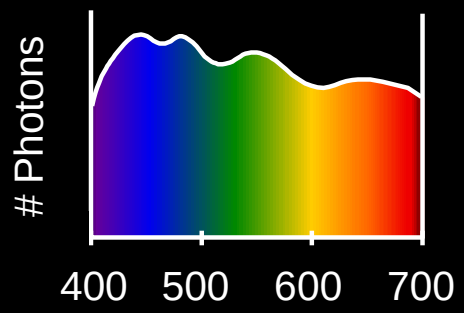

The Physics of Light

Some examples of the reflectance spectra of surfaces




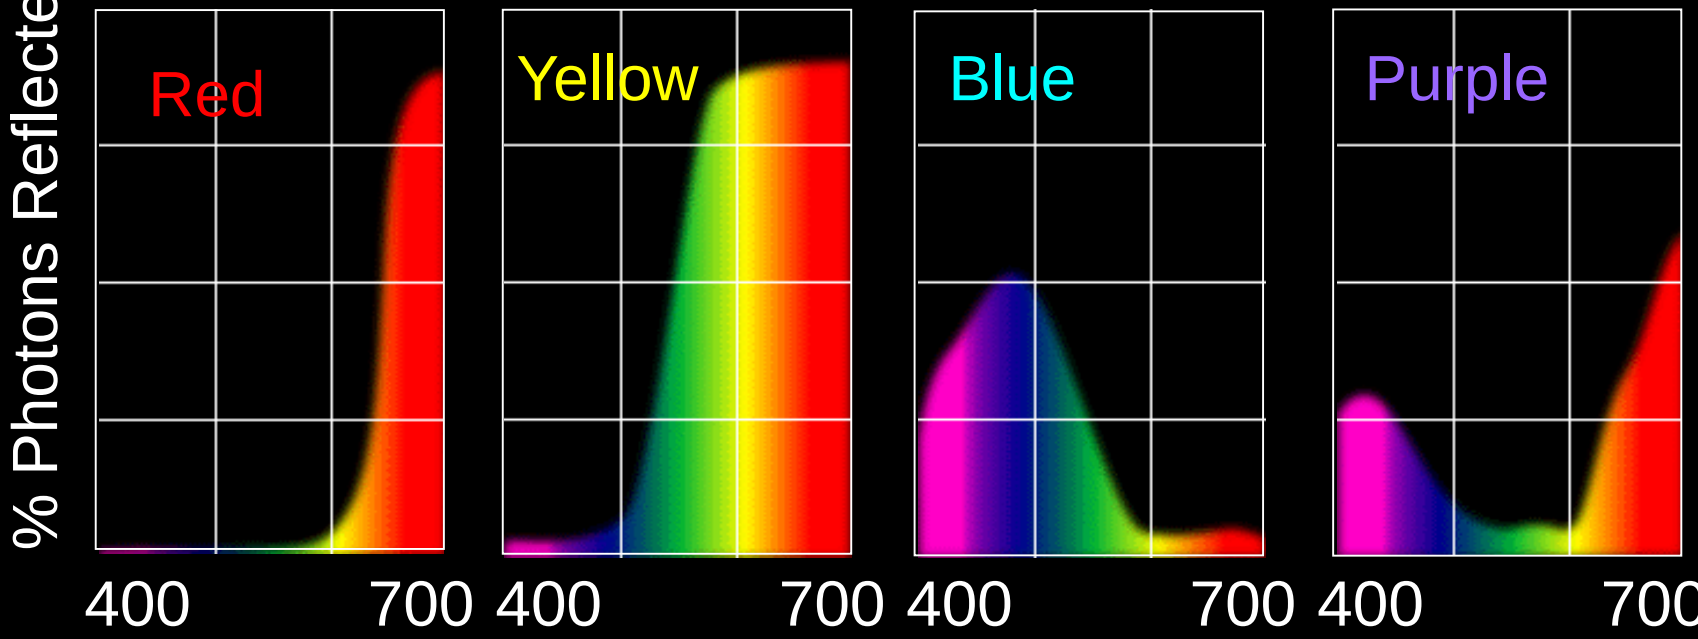

Physiology of Color Vision

Three kinds of cones:

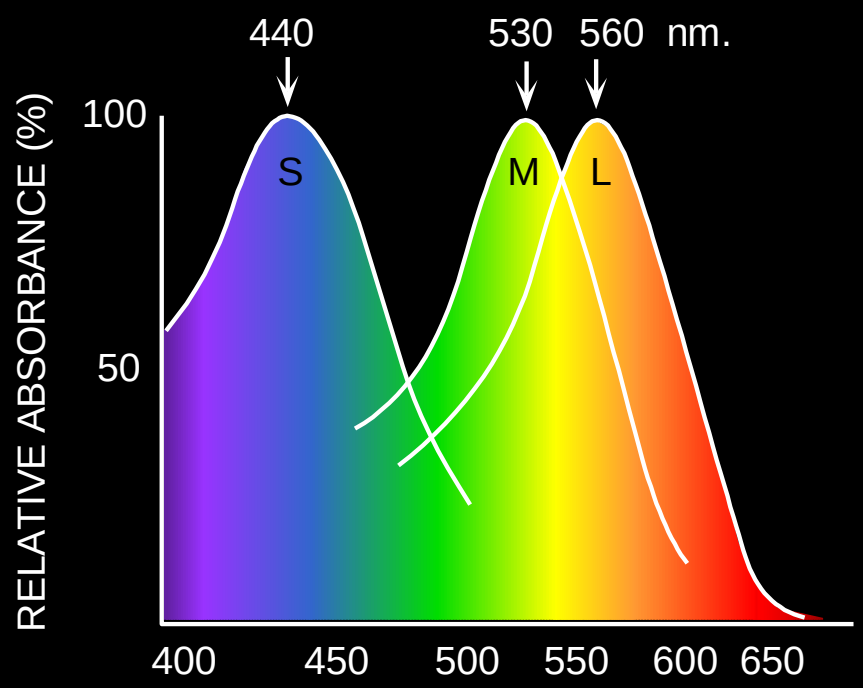

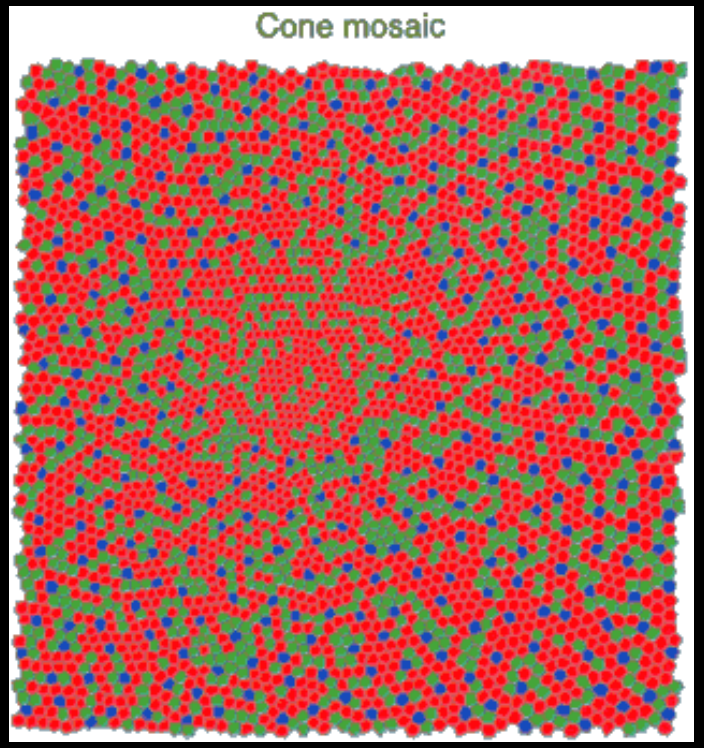

- Why are M and L cones so close?
- Why are there 3?

Tetrachromatism

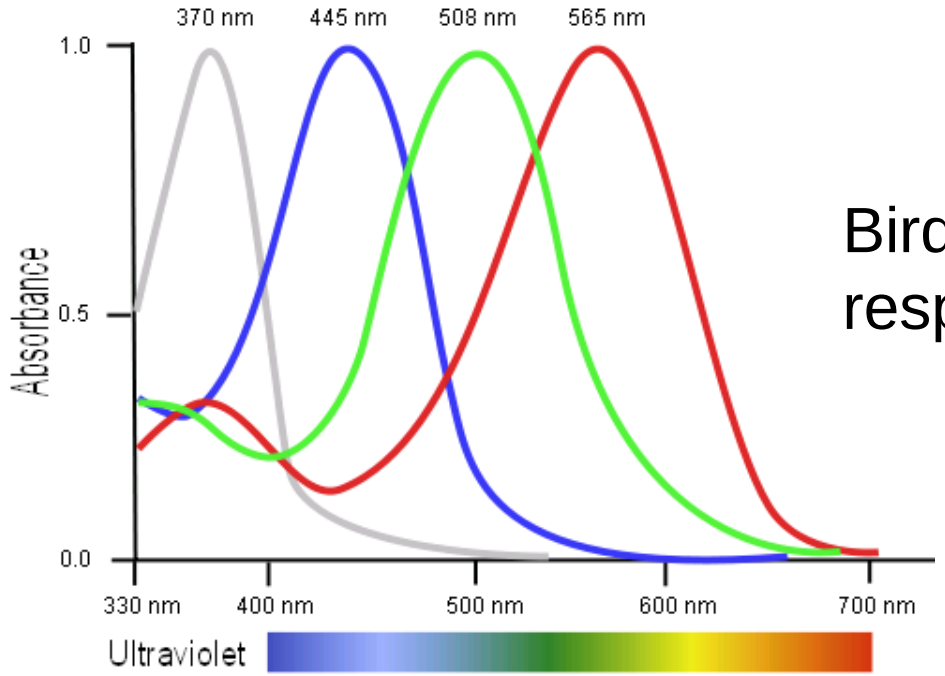

Bird cone responses

- 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

Metamers

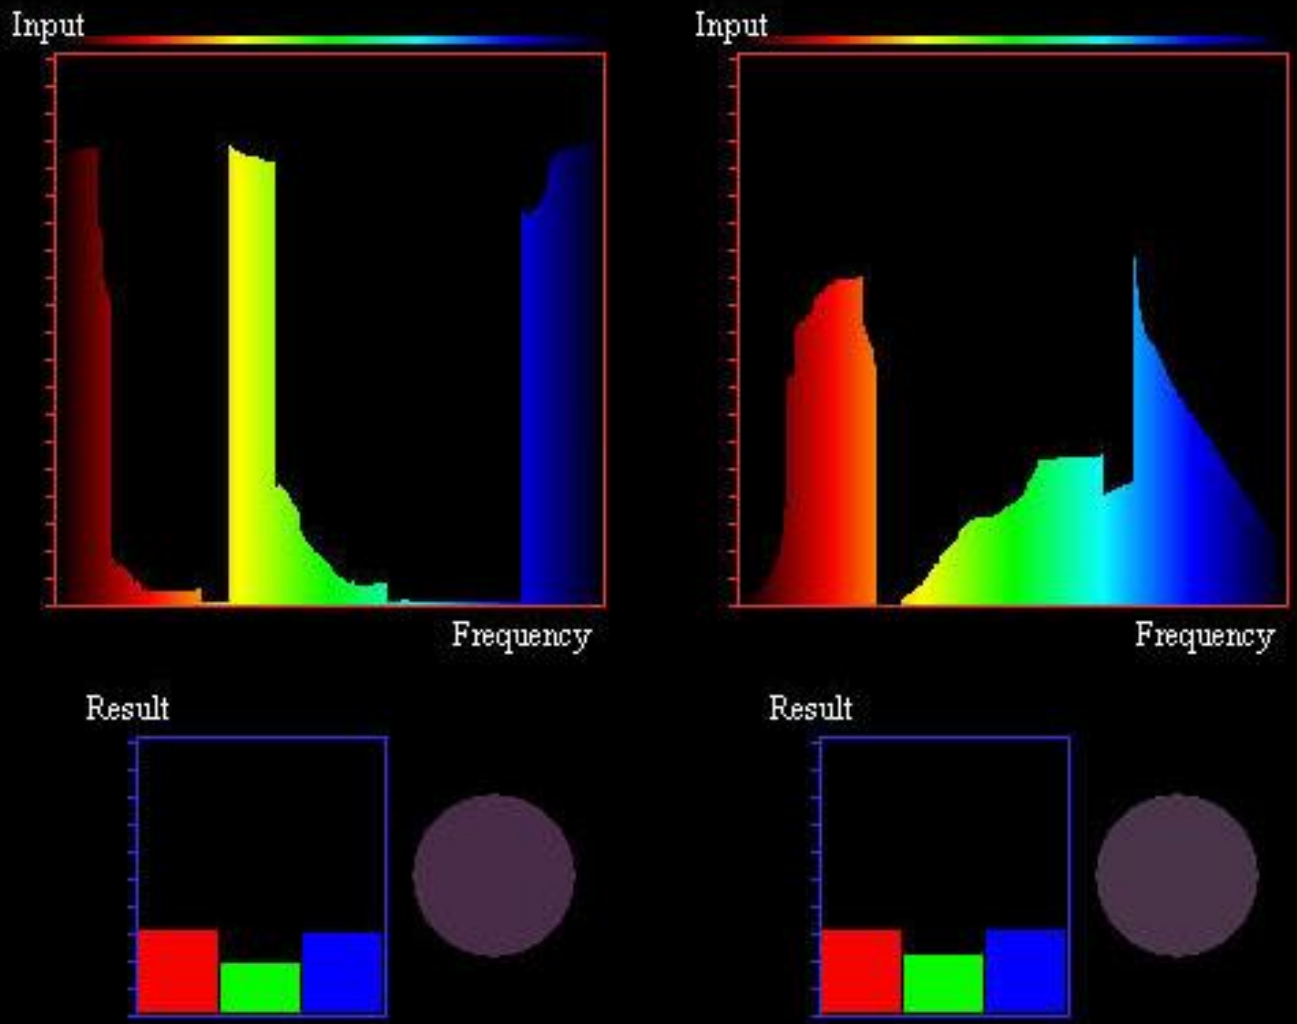

by Jeff Beall, Adam Doppelt and John F. Hughes

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

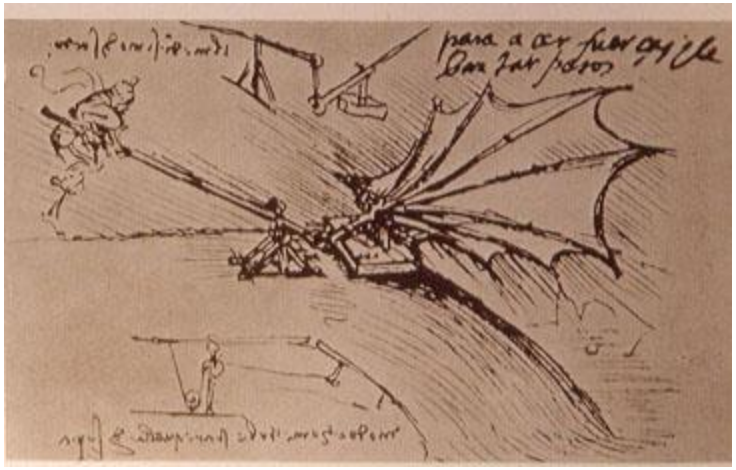

"Can machines fly like a bird?"
No, because airplanes don’t flap.

"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

Color Sensing in Camera (RGB)

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

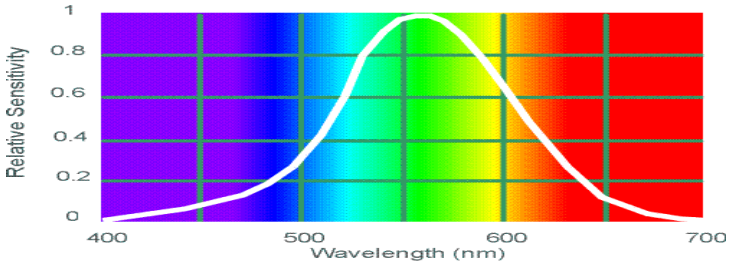

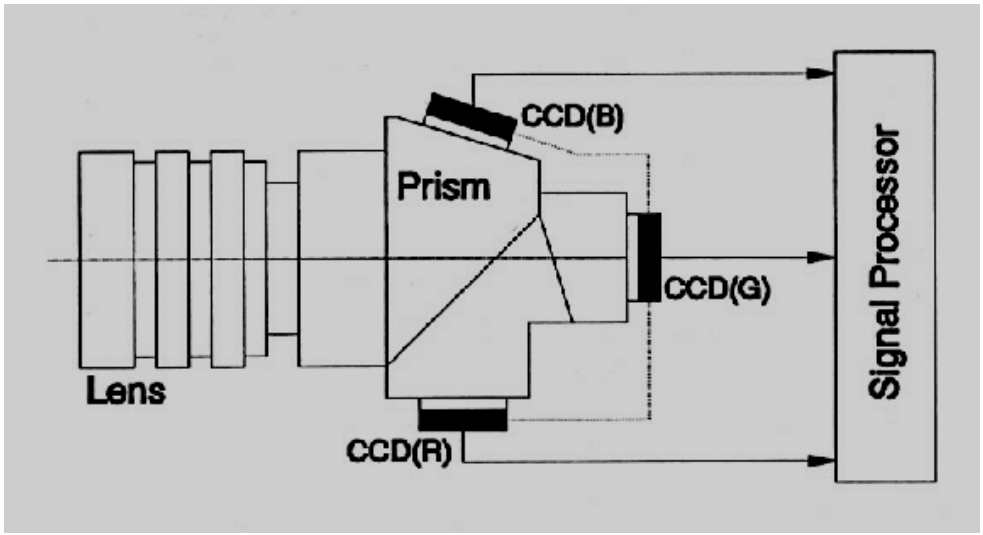

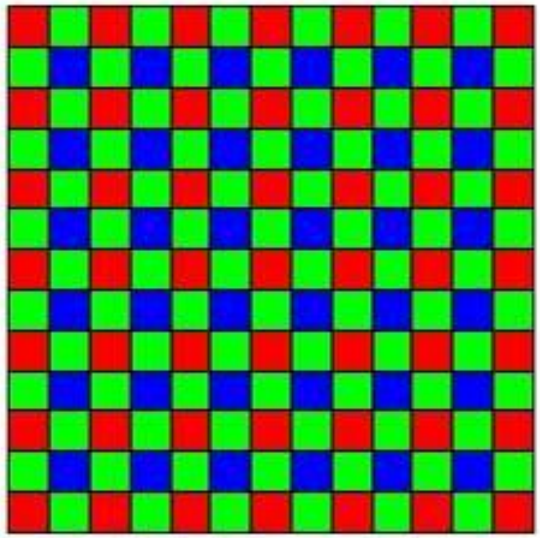

Bayer filter

Why 3 colors?

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

Slide by Steve Seitz

Practical Color Sensing: Bayer Grid

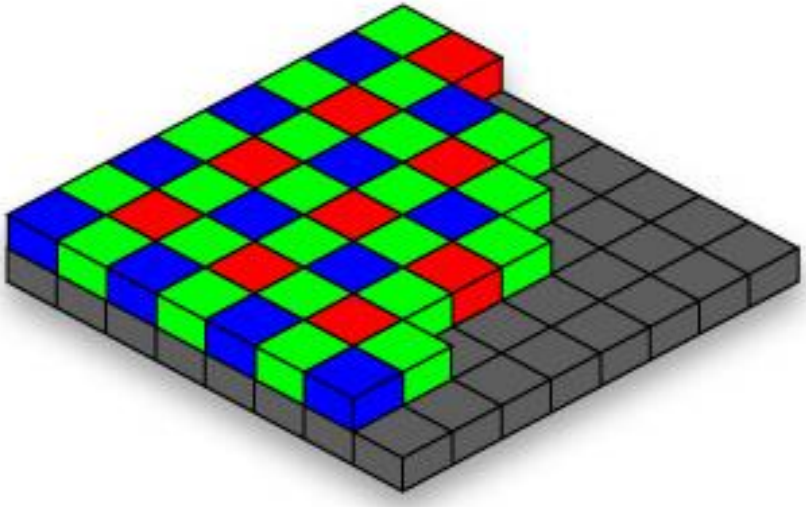

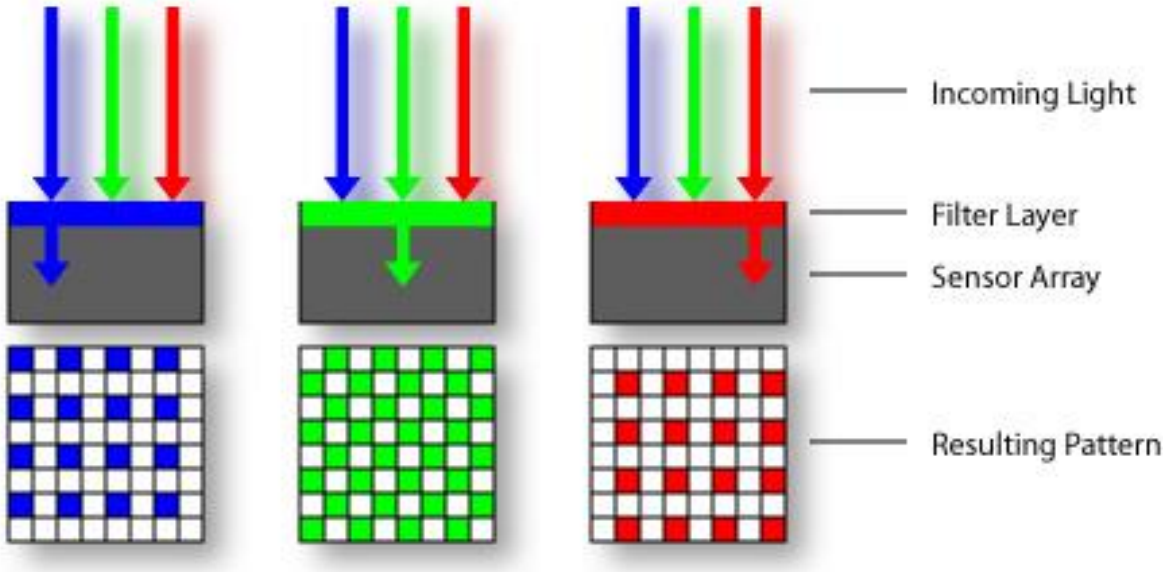

- Estimate RGB at ‘G’ cells from neighboring values

Slide by Steve Seitz

Camera Color Response

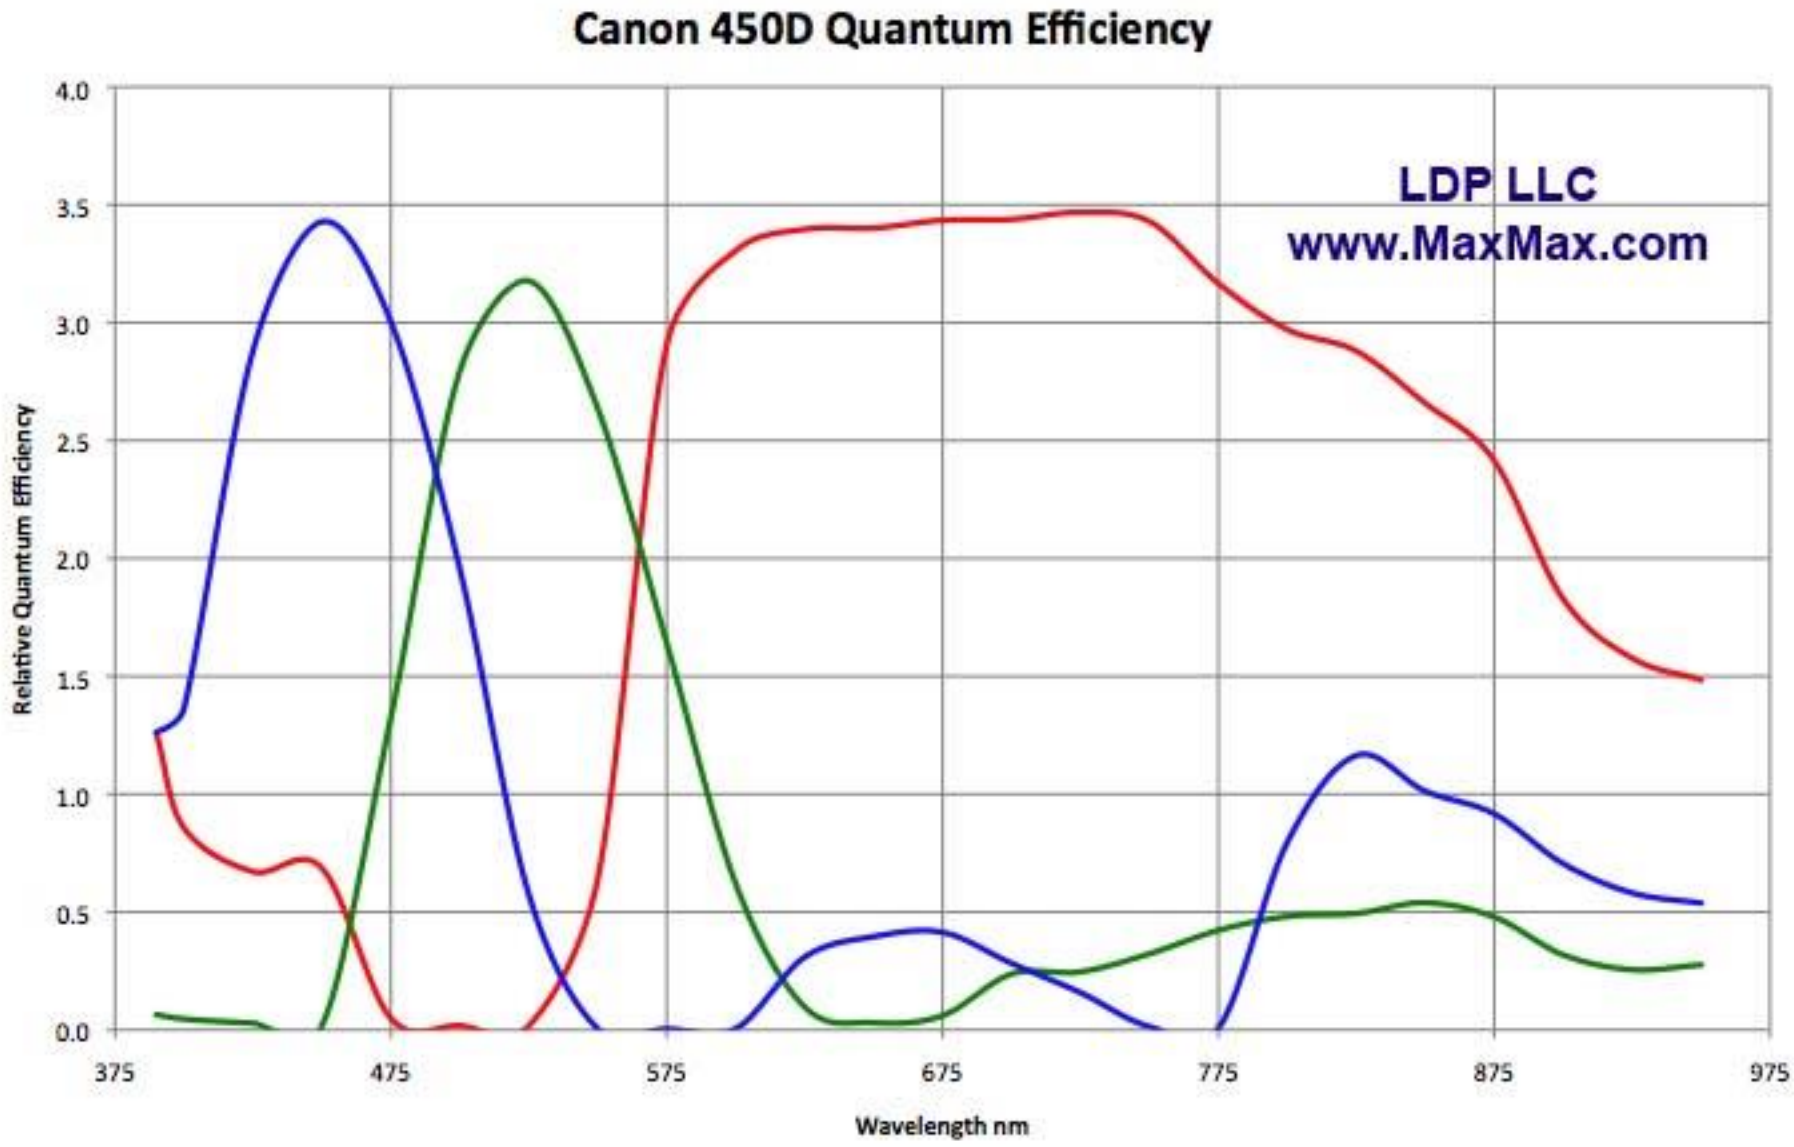

Color spaces

- How can we represent color?

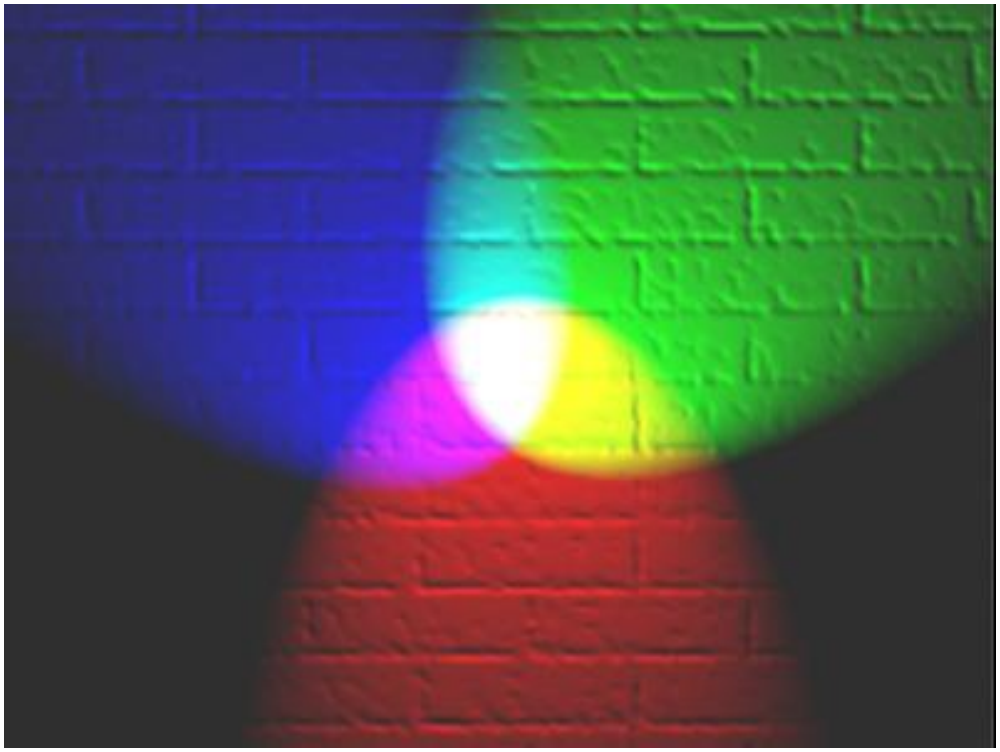

http://en.wikipedia.org/wiki/File:RGB_illumination.jpg

Color spaces: RGB

Default color space

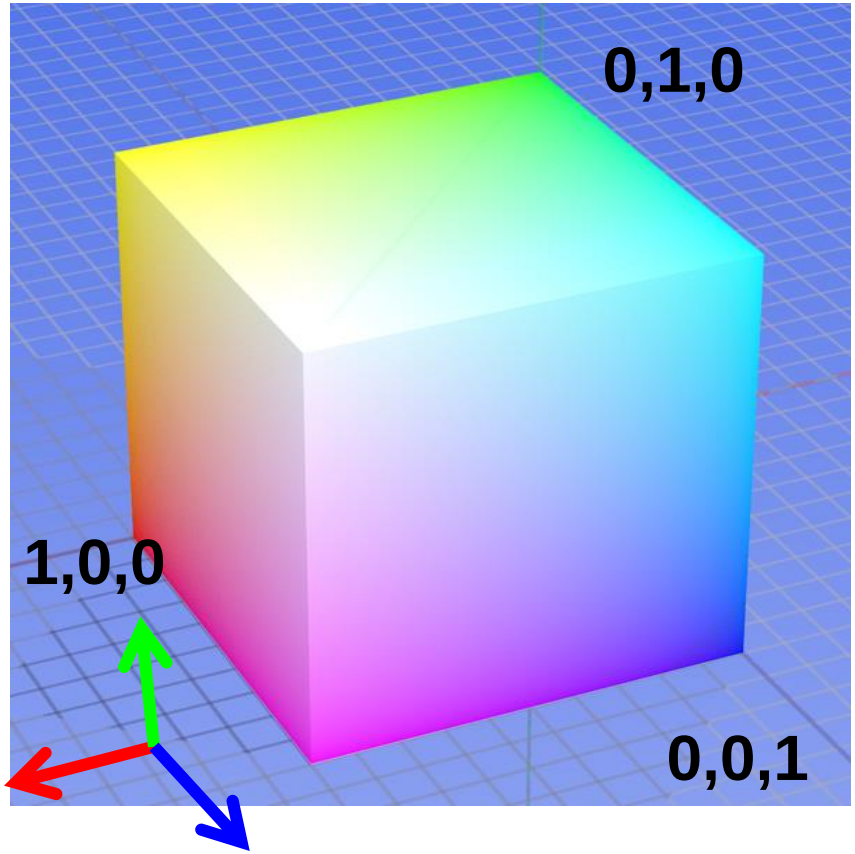

R = 1
(G=0,B=0)

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

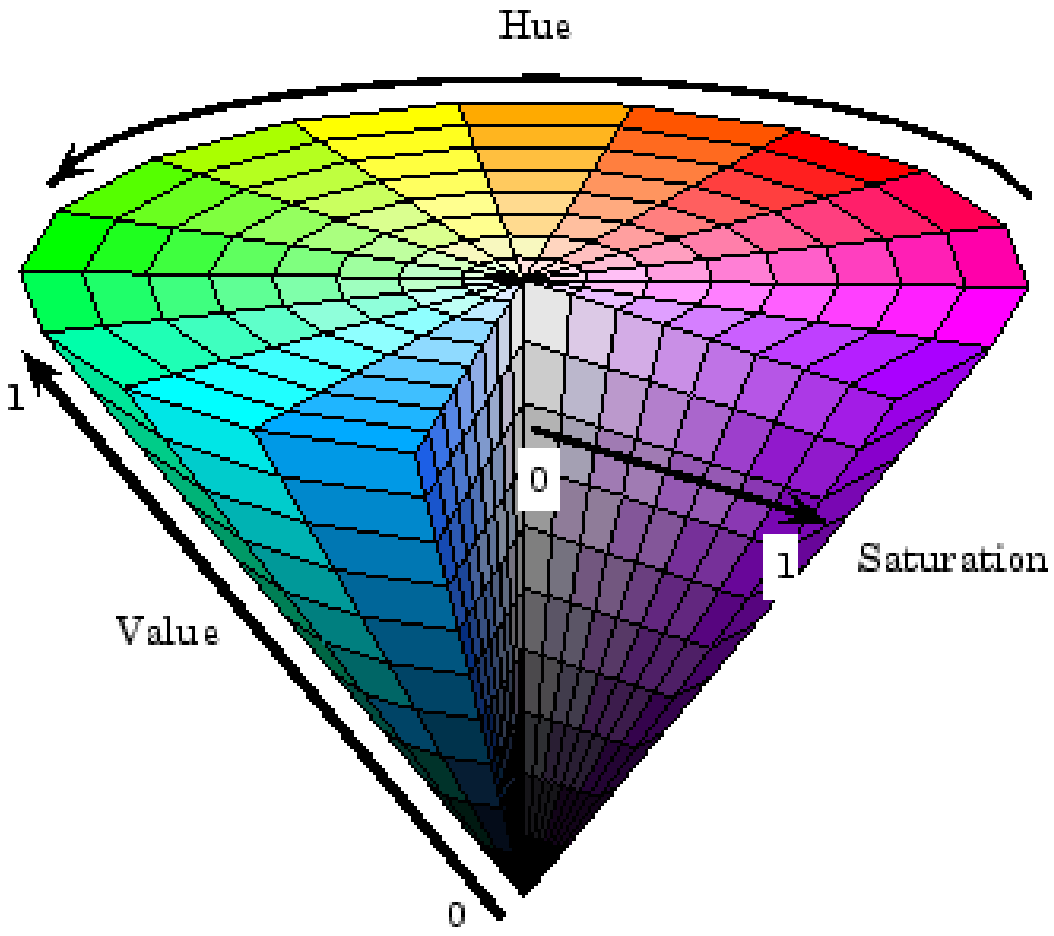

If you had to choose, would you rather go without:
- intensity (‘value’), or
- hue + saturation (‘chroma’)?

Think-Pair-Share

Most information in intensity

Only color shown – constant intensity

Most information in intensity

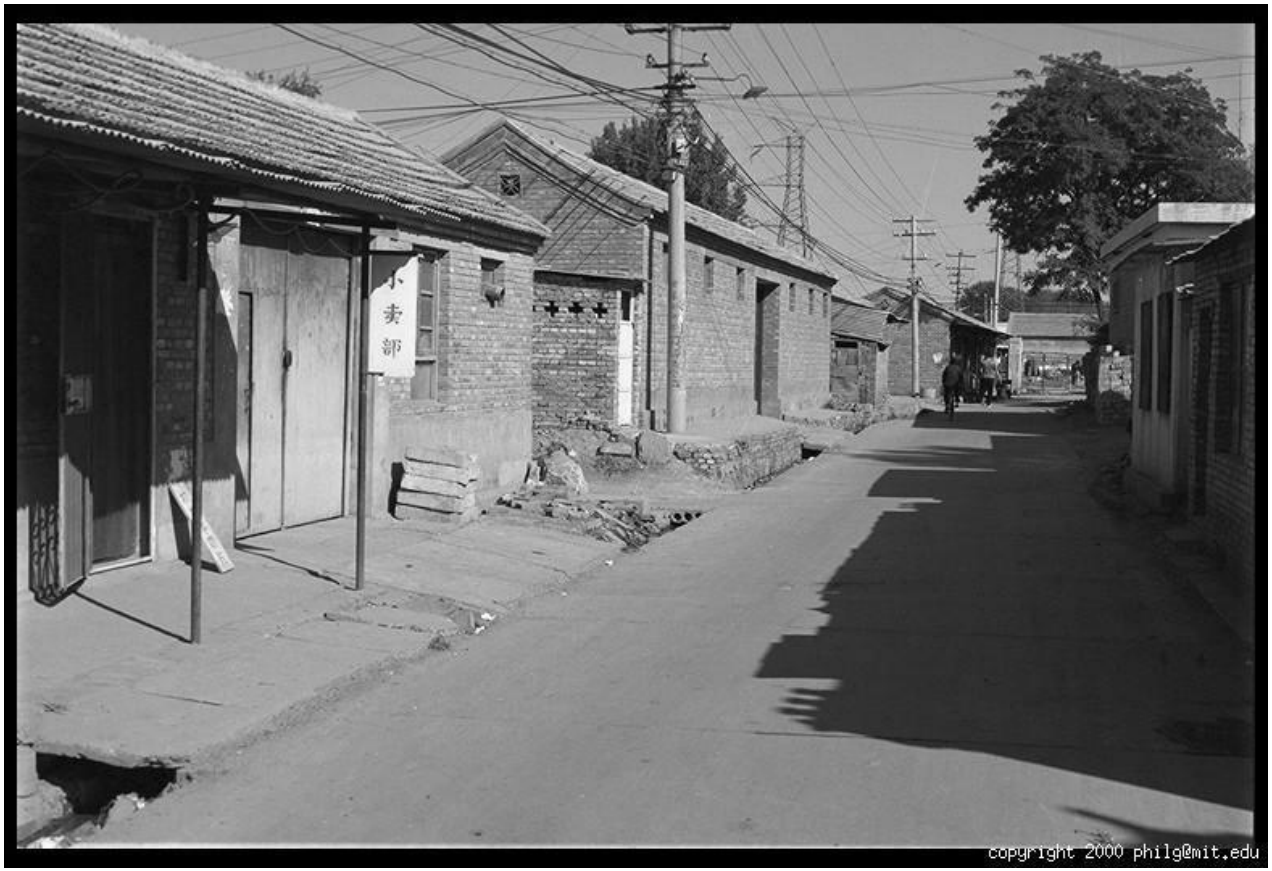

Only intensity shown – constant color

James Hays

Most information in intensity

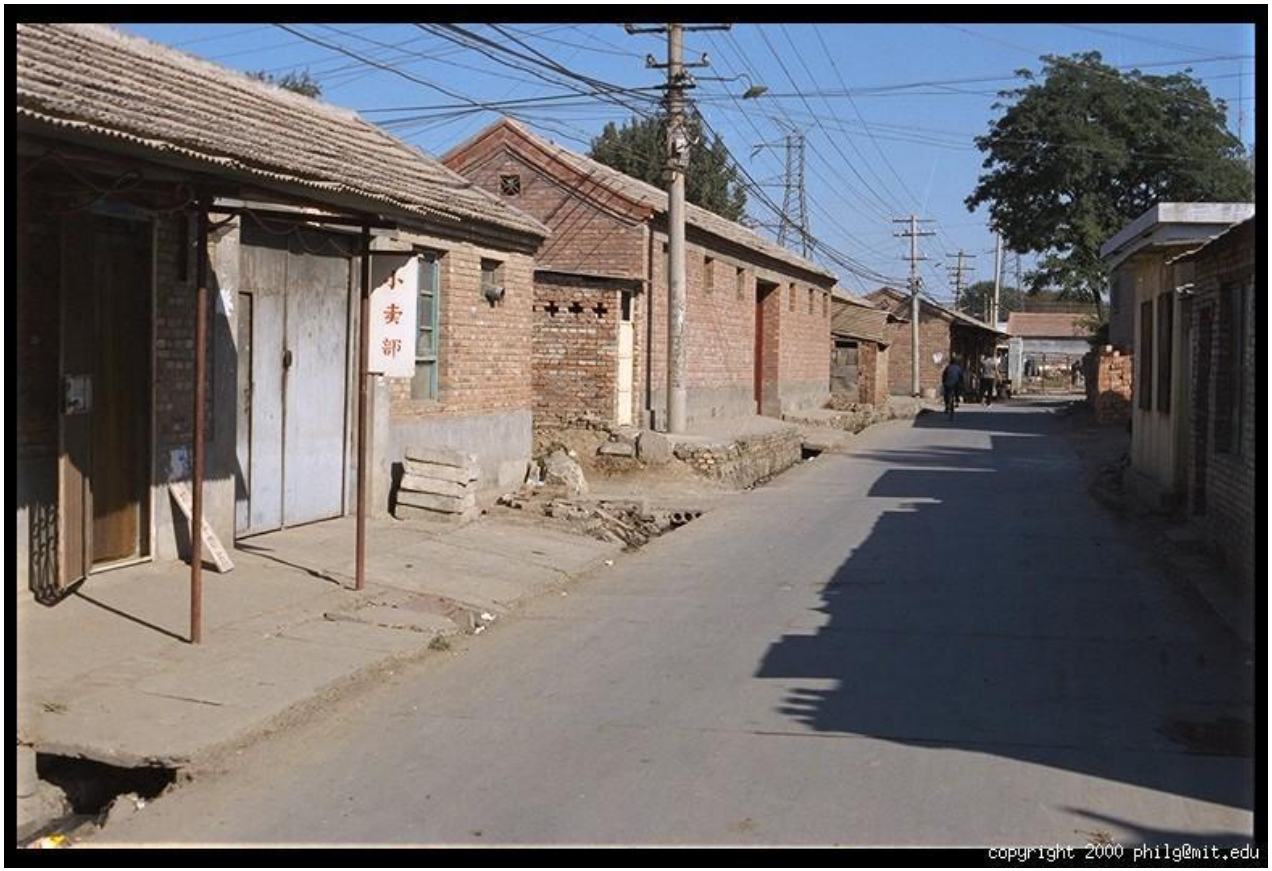

Original image

Color spaces: HSV

Intuitive color space

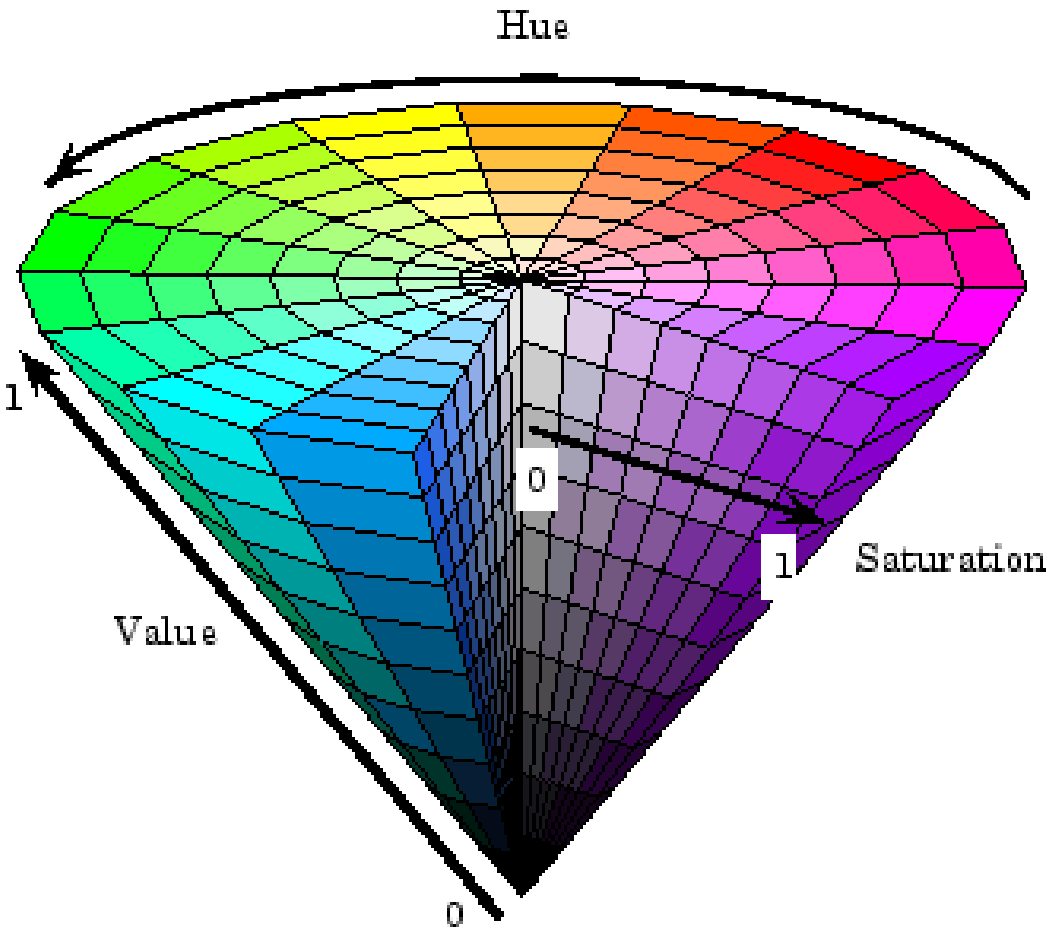

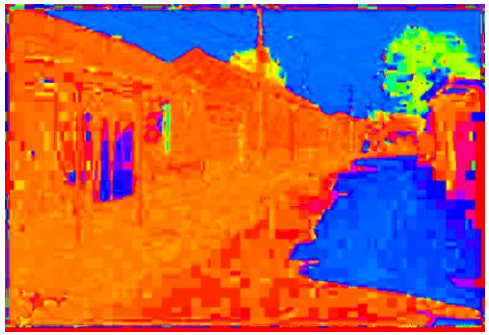

H
(S=1,V=1)

S
(H=1,V=1)

V
(H=1,S=0)

James Hays

Color spaces: YCbCr

Fast to compute, good for compression, used by TV

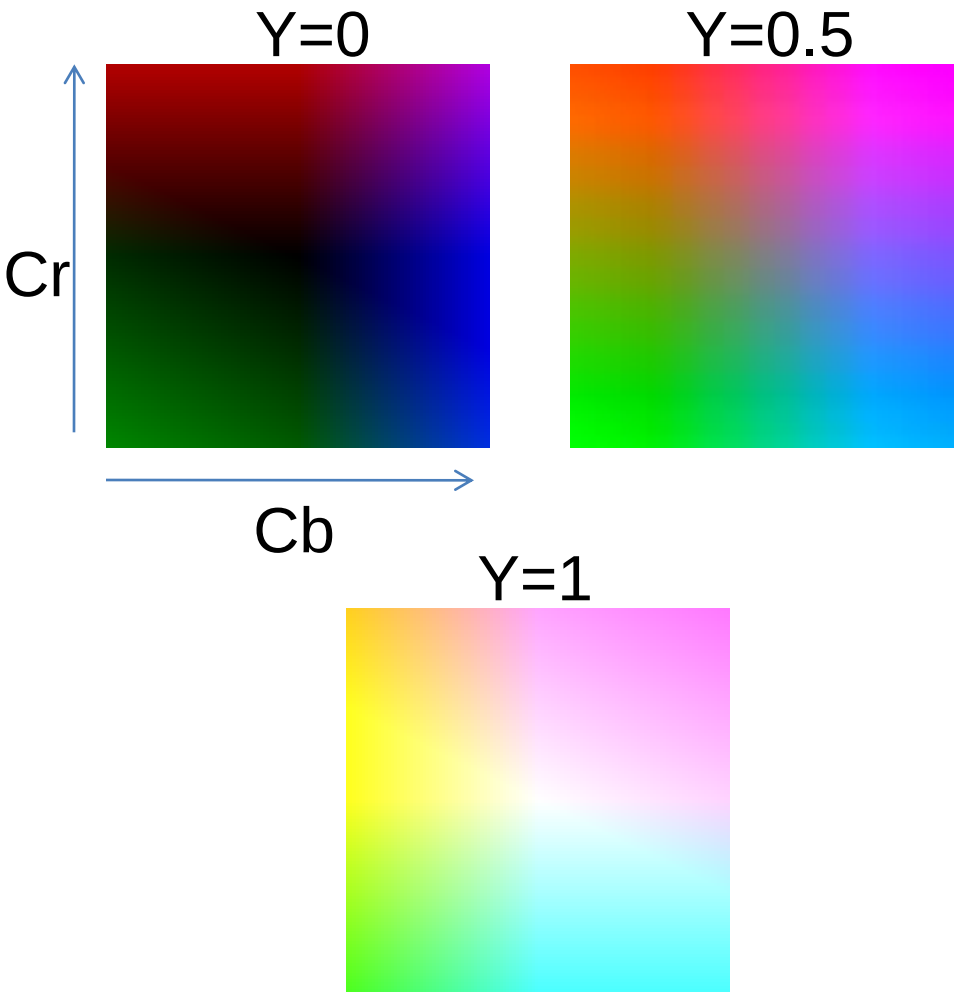

Y
(Cb=0.5,Cr=0.5)

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

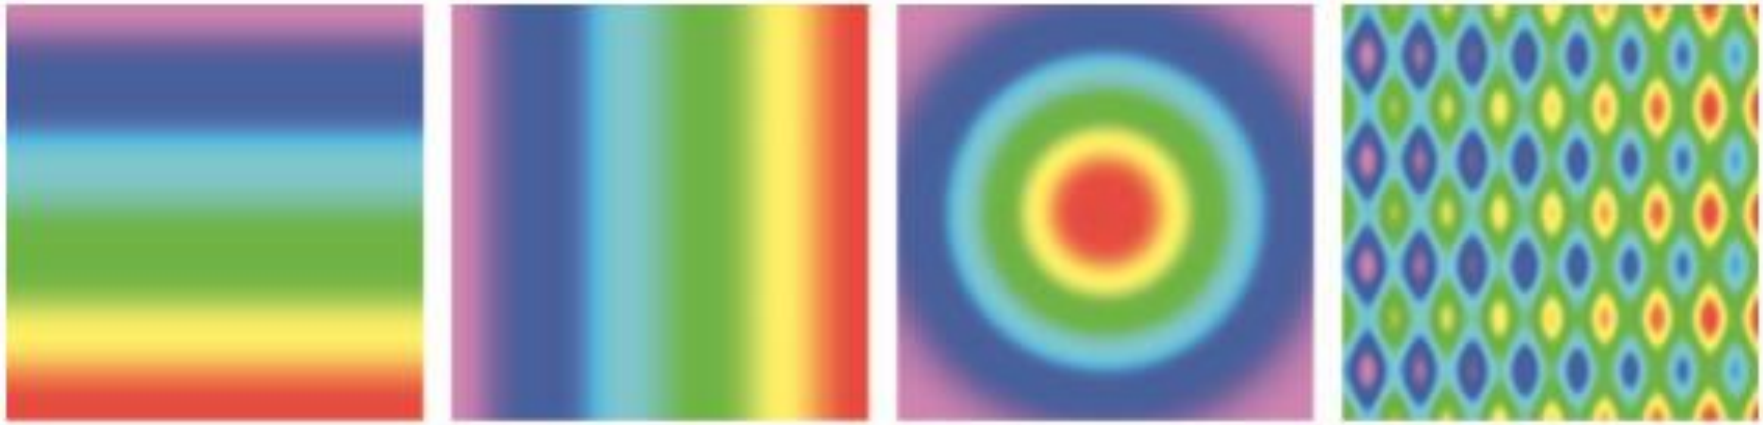

Rainbow color map considered harmful

Borland and Taylor

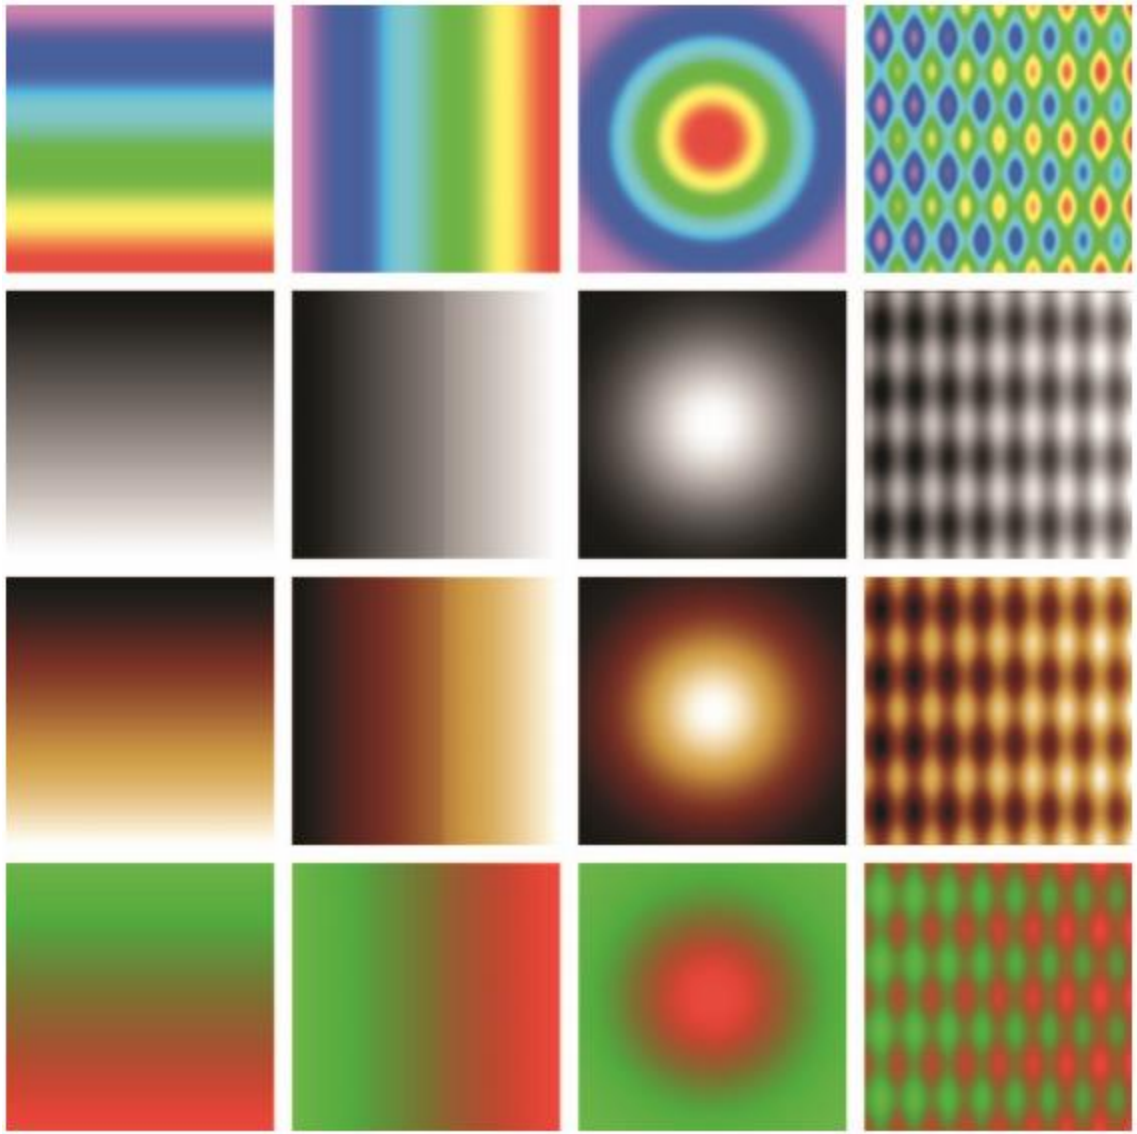

IS COLOR PERCEPTION A VECTOR SPACE?

Color spaces: L*a*b*

"Perceptually uniform"* color space

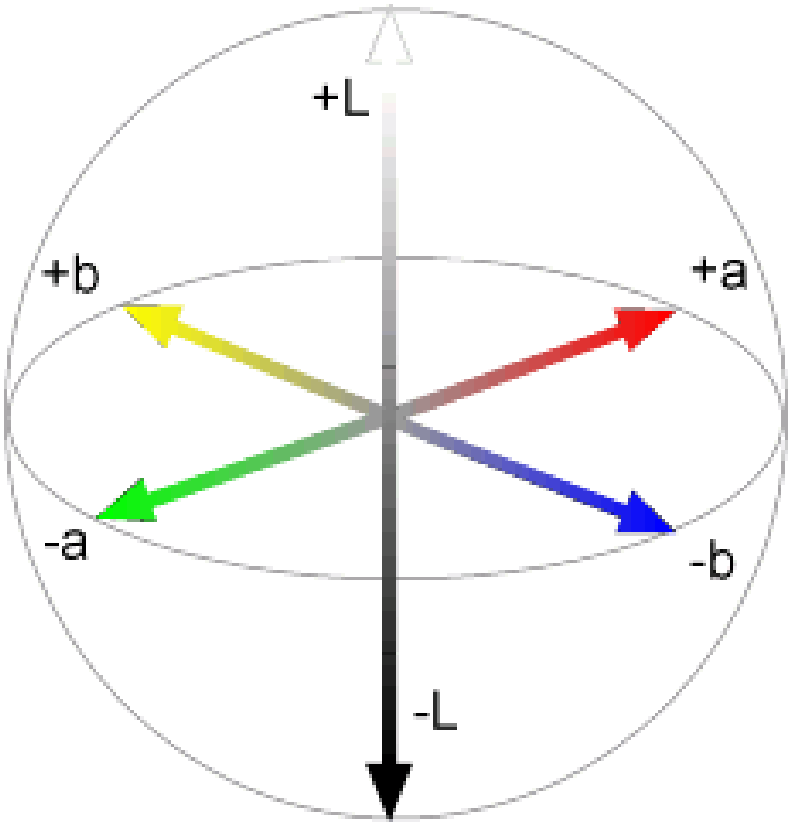

L
(a=0,b=0)

a
(L=65,b=0)

b
(L=65,a=0)

James Hays

“Intuitive” color space?

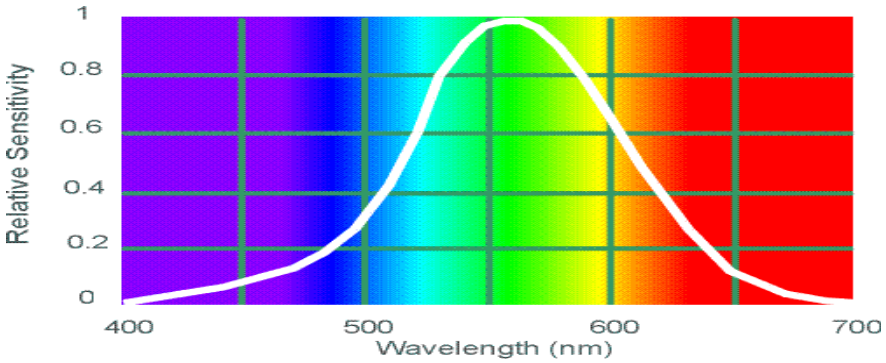

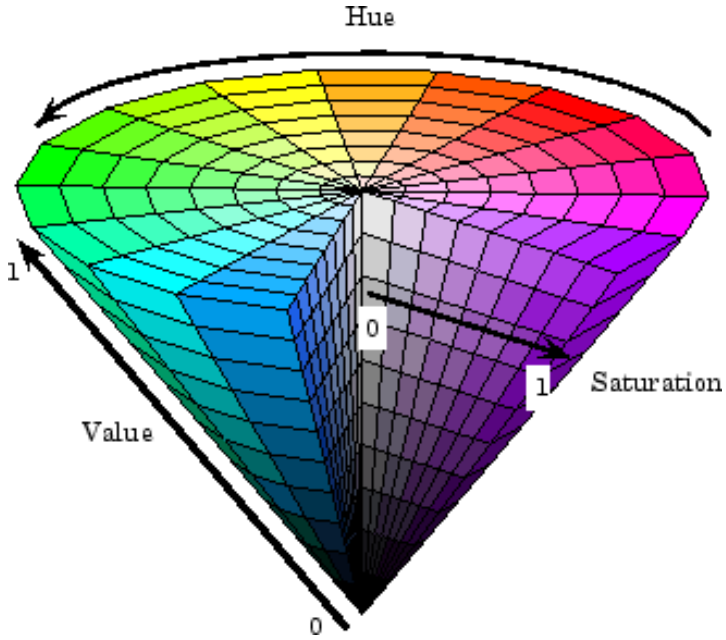

Wait a minute…

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

Project 0:

Tonight Sunlab 6pm-9pm

Next week: Project 1

- Convolution
- Filtering
- Image Pyramids
- Frequencies

EVOLUTION OF MY UNDERSTANDING OF COLOR OVER TIME:

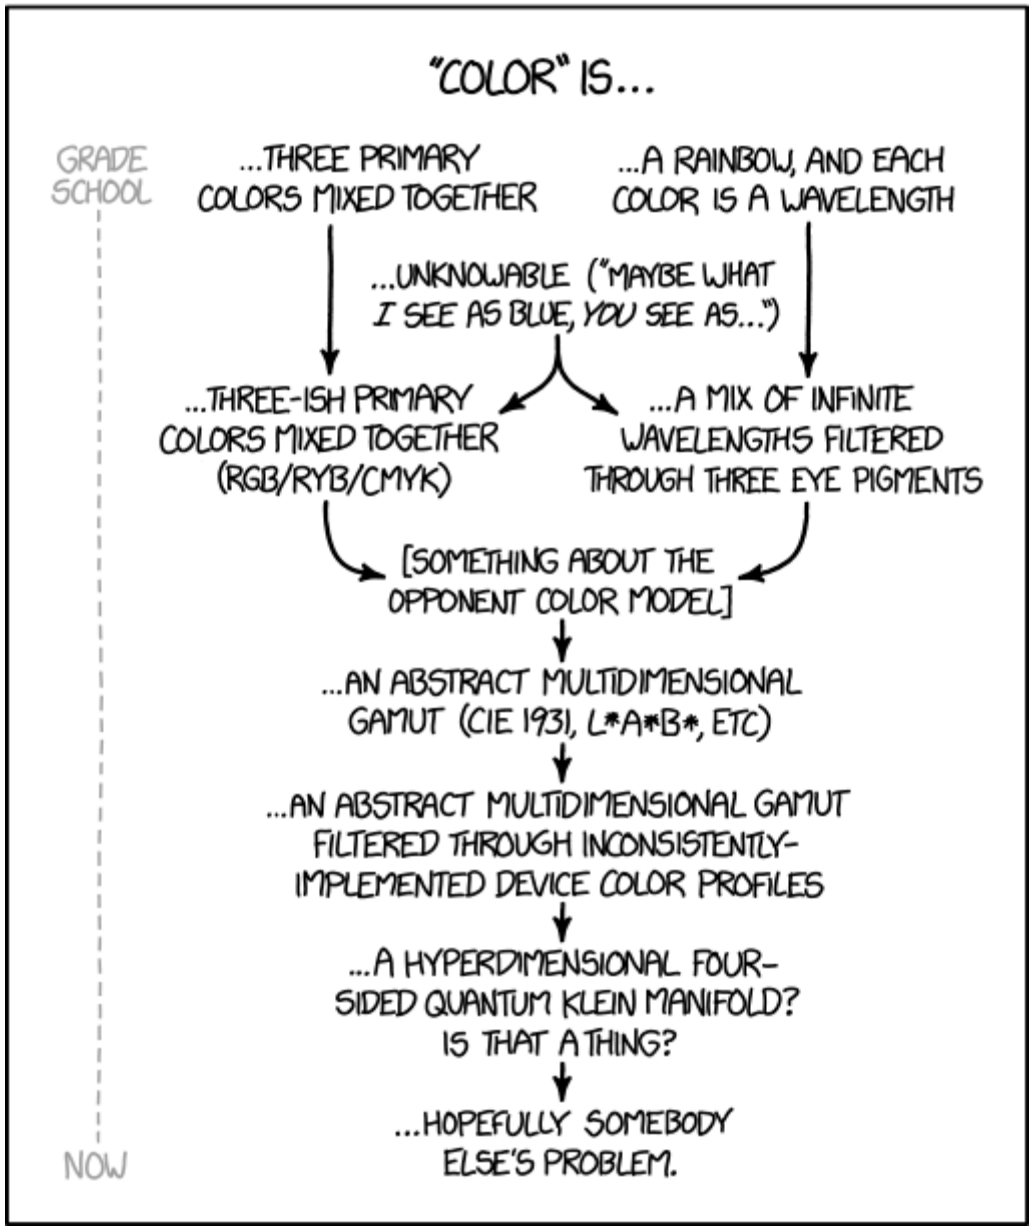

XKCD

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!

Proj 1: Image Filtering and Hybrid Images

- Implement image filtering to separate high and low frequencies.
- Combine high frequencies and low frequencies from different images to create a scale-dependent image.

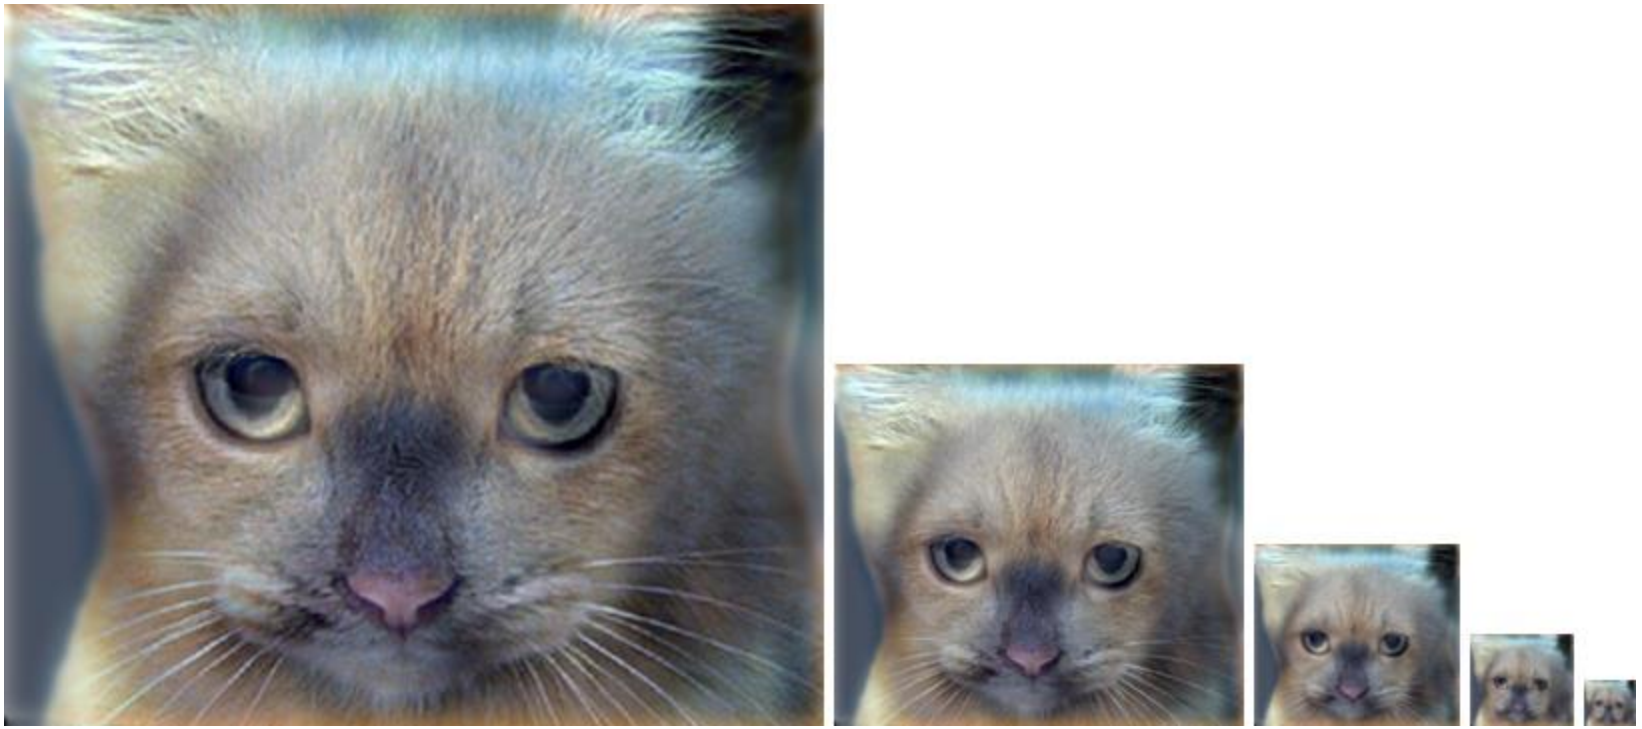

James Hays