Computer Graphics

Prof. Feng Liu

Fall 2021

http://www.cs.pdx.edu/~fliu/courses/cs447/

09/29/2021

Announcement

- Read *Programming Tutorial 1* before trying compiling Project 1 code
- Lab session
  - 10/06/2021 in class
  - Have a computer ready with Visual Studio 2019 installed
  - Our TA will show how to set up Project 1
- Homework 1 – due Oct. 06 before class
  - Email your homework to abhijay@pdx.edu

Last Time

- Course introduction
- Digital images
  - The difference between an image and a display
  - Ways to get them
  - Raster vs. Vector

Today

- Digital images
  - Raster vs. Vector
  - Digital images as discrete representations of reality
  - Human perception in deciding resolution and image depth
- Color
- Tri-Chromacy
- Digital Color

Raster Images

- A raster is a regular grid of *pixels* (picture elements)
  - The smallest element of an image is called a pixel
- Raster image formats store the color at each pixel, and maybe some other information
  - Easiest is to use a simple array of pixel values
  - Some formats store the pixel information in *very* different ways
  - e.g. a 5x3, floating point, grayscale image

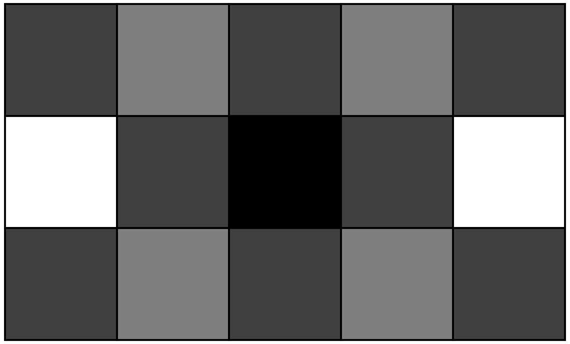

0.25	0.5	0.25	0.5	0.25
1	0.25	0	0.25	1
0.25	0.5	0.25	0.5	0.25

Vector Images

- Vector formats offer an alternative way to store images
- The most common use of vector formats are in fonts – images of characters (Postscript, TrueType)
- Store images as collections of geometric primitives
  - E.g. Lines, polygons, circles, curves, …

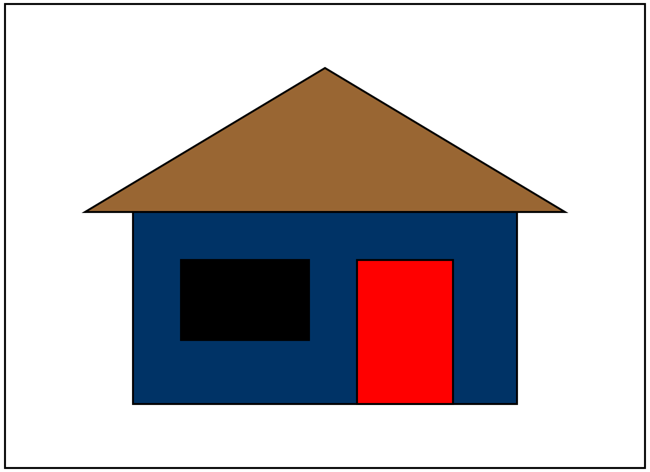

- It is possible to go from a vector image to a raster image
  - We'll learn how
- It is very hard to go the other way
  - A popular yet challenging computer vision problem

Trade-Offs

- Which format, raster or vector, is easier to:
  - Display?
  - Resize (scale bigger or smaller)?
  - Rotate?
  - Crop (cut bits off at the edges)?

Obtaining Digital Images

- What are some methods for obtaining a digital image?

Obtaining Digital Images

- What are some methods for obtaining a digital image?
  - Digital camera
  - Scanning another image
  - Other forms of scanning (e.g. medical)
  - Editing existing digital images
  - Paint or drawing programs
  - Created from abstract data (e.g. math function plot)
  - Rendered from a scene description
  - …

Ideal Images

- The information stored in images is often continuous in nature
- For example, consider the ideal photograph:
  - It captures the intensity of light at a particular set of points coming from a particular set of directions (it's called *irradiance*)
  - The intensity of light arriving at the camera can be any positive real number, and it *mostly* varies smoothly over space
  - The world we see is not pixelated
    - Where do you see spatial *discontinuities* in a photograph?

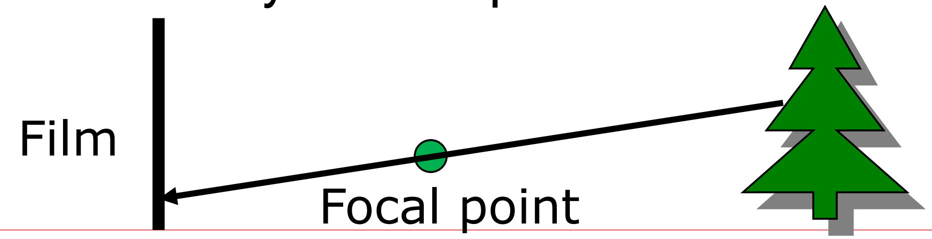

Digital Images

- Computers work with discrete pieces of information
- How do we digitize a continuous image?
  - Break the continuous space into small areas, *pixels*
  - Use a single value for each pixel - the *pixel value* (no color, yet)
  - No longer continuous in space or intensity
- This process is fraught with danger, as we shall see

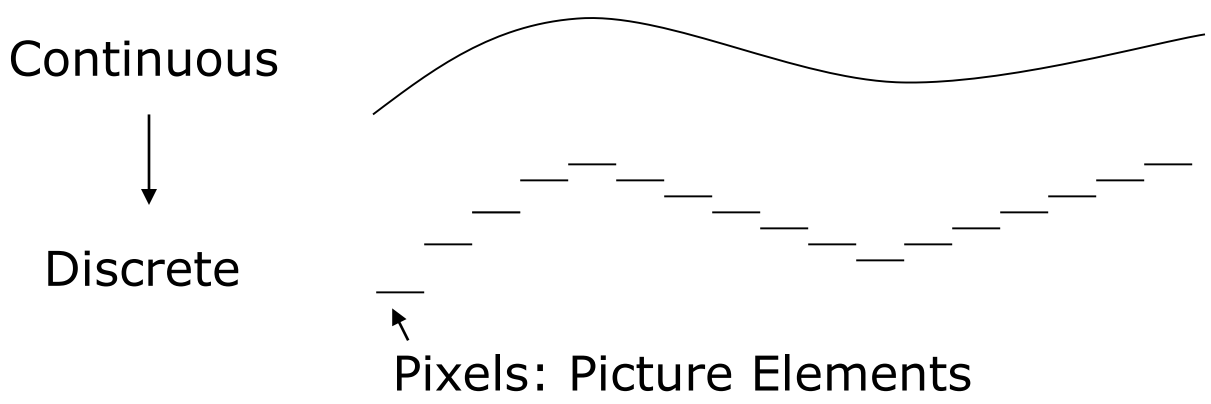

Discretization Issues

- Can only store a finite number of pixels
  - Choose your target physical image size, choose your resolution (pixels per inch, or dots per inch, dpi), determine width/height in pixels necessary
  - Storage space goes up with square of resolution
    - 600dpi has 4× more pixels than 300dpi
- Can only store a finite range of intensity values
  - Typically referred to as *depth* - number of bits per pixel
    - Directly related to the number of colors available and typically little choice
    - Most common depth is 8, but also sometimes see 16 for grey
  - Also concerned with the minimum and maximum intensity – dynamic range
- What is enough resolution and enough depth?

Perceptual Issues

- Spatially, humans can discriminate about ½ a minute of arc
  - At fovea, so only in center of view
  - At 0.5m, about 0.1mm ("Dot pitch" of monitors)
  - Sometimes limits the required number of pixels

129 128 125

- Humans can discriminate about 8 bits of intensity
  - "Just Noticeable Difference" experiments
  - Limits the required depth for typical dynamic ranges
  - Actually, it's 9-10 bits, but 8 is far more convenient
- BUT, when manipulating images much higher resolution may be required

DeepFovea: Neural Reconstruction for Foveated Rendering (Facebook Reality Lab)

https://www.youtube.com/watch?v=eTUmmW4ispA

Intensity Perception

- Humans are actually tuned to the *ratio* of intensities, not their absolute difference
  - So going from a 50 to 100 Watt light bulb looks the same as going from 100 to 200
- Most computer graphics ignores this, giving poorer perceptible intensity resolution at low light levels, and better resolution at high light levels

Dynamic Range

- Image depth refers to the number of bits available, but not how those bits map onto intensities
- We can use those bits to represent a large range at low resolution, or a small range at high resolution
- Common display devices can only show a limited dynamic range, so typically we fix the range at that of the display device and choose high resolution

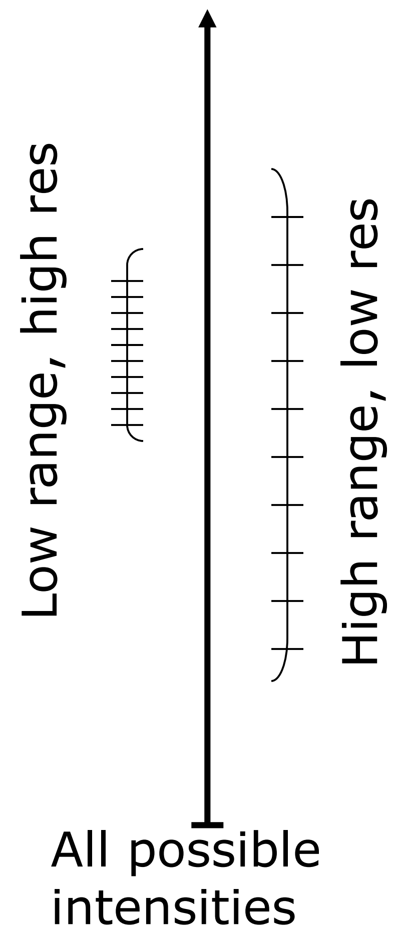

More Dynamic Range

- Real scenes have very high and very low intensities
- Humans can see contrast at very low and very high light levels
  - Can't see all levels all the time – use adaptation to adjust
  - Still, high range even at one adaptation level
- Film has low dynamic range around 100:1
- Monitors are even worse
- Many ways to deal with the problem
  - Way beyond the scope of this course

Display on a Monitor

- When images are created, a *linear* mapping between pixels and intensity is assumed
  - For example, if you double the pixel value, the displayed intensity should double
- Monitors, however, do not work that way
  - For analog monitors, the pixel value is converted to a voltage
  - The voltage is used to control the intensity of the monitor pixels
  - But the voltage to display intensity is *not linear*
  - Similar problem with other monitors, different causes
- The outcome: A linear intensity scale in memory does not look linear on a monitor
- Even worse, **different monitors do different things**

Display on a Monitor

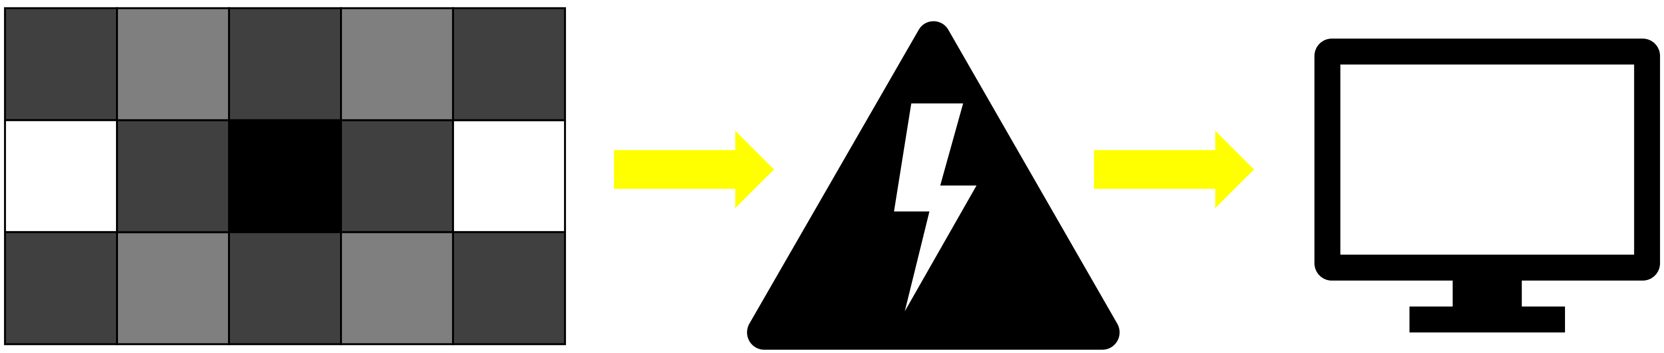

Image intensity

Voltage to monitor

Display intensity

Display on a Monitor

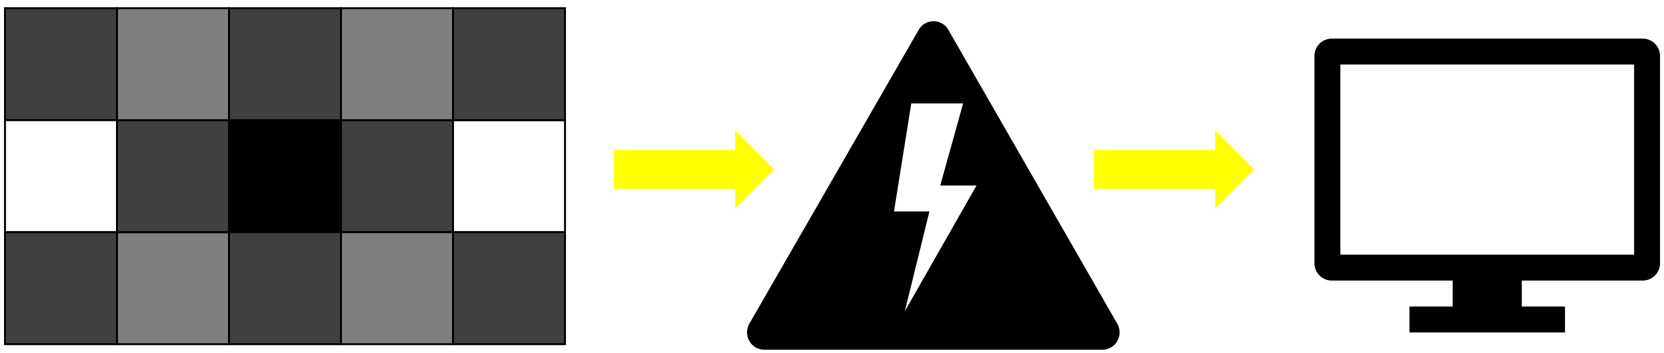

Image intensity

Voltage to monitor

Display intensity

$$I_{display} \propto I_{to-monitor}^{\gamma}$$

Display on a Monitor

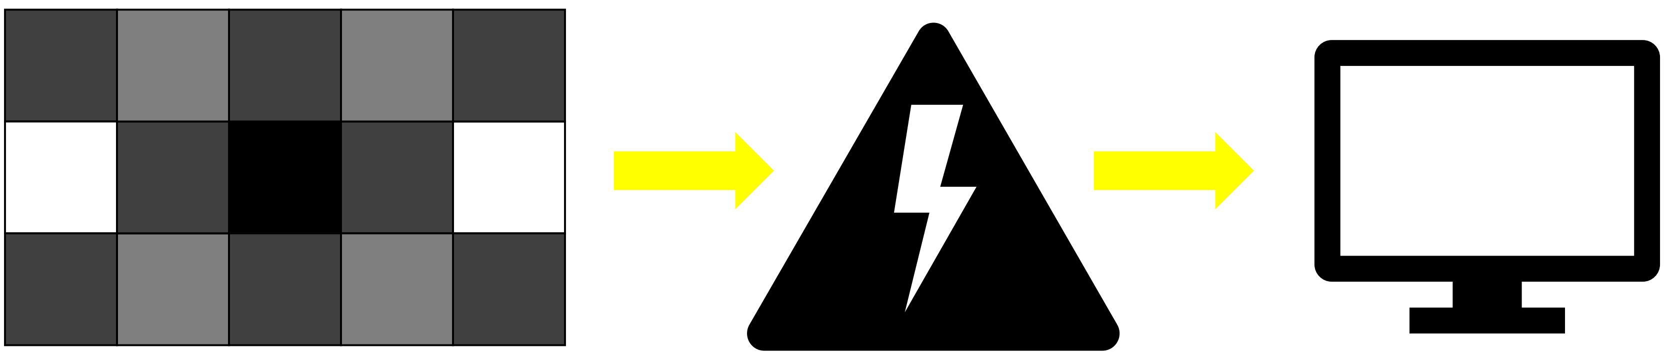

Image intensity | Voltage to monitor | Display intensity

$$I_{to-monitor} \propto I_{image}^{1/\gamma} \qquad I_{display} \propto I_{to-monitor}^{\gamma}$$

$$I_{display} \propto I_{image}$$

Gamma Control

- The mapping from voltage to display is usually an exponential function: $I_{display} \propto I_{to-monitor}^{\gamma}$
- To correct the problem, we pass the pixel values through a *gamma function* before converting them to the monitor

$$I_{to-monitor} \propto I_{image}^{1/\gamma}$$

- This process is called *gamma correction*
- The parameter, γ, is controlled by the user
  - It should be matched to a particular monitor
  - Typical values are between 2.2 and 2.5
- The mapping can be done in hardware or software

Today

- Digital images
  - Raster vs. Vector
  - Digital images as discrete representations of reality
  - Human perception in deciding resolution and image depth
- Color
- Tri-Chromacy
- Digital Color

About Color

- So far, we have only discussed intensities, so called *achromatic light* (shades of gray)
- On the order of 10 color names are widely recognized by English speakers - other languages have fewer/more, but not much more

About Color

- So far, we have only discussed intensities, so called *achromatic light* (shades of gray)
- On the order of 10 color names are widely recognized by English speakers - other languages have fewer/more, but not much more
- Accurate color reproduction is commercially valuable - e.g. painting a house, producing artwork
- E-commerce has accentuated color reproduction issues, as has the creation of digital libraries
- Color consistency is also important in user interfaces, e.g.: what you see on the monitor should match the printed version

Light and Color

- The frequency, ω, of light determines its "color"
  - Wavelength, λ, is related
  - Energy also related
- Describe incoming light by a *spectrum*
  - Intensity of light at each frequency
  - A graph of intensity vs. frequency
- We care about wavelengths in the visible spectrum: between the infra-red (700nm) and the ultra-violet (400nm)

Normal Daylight

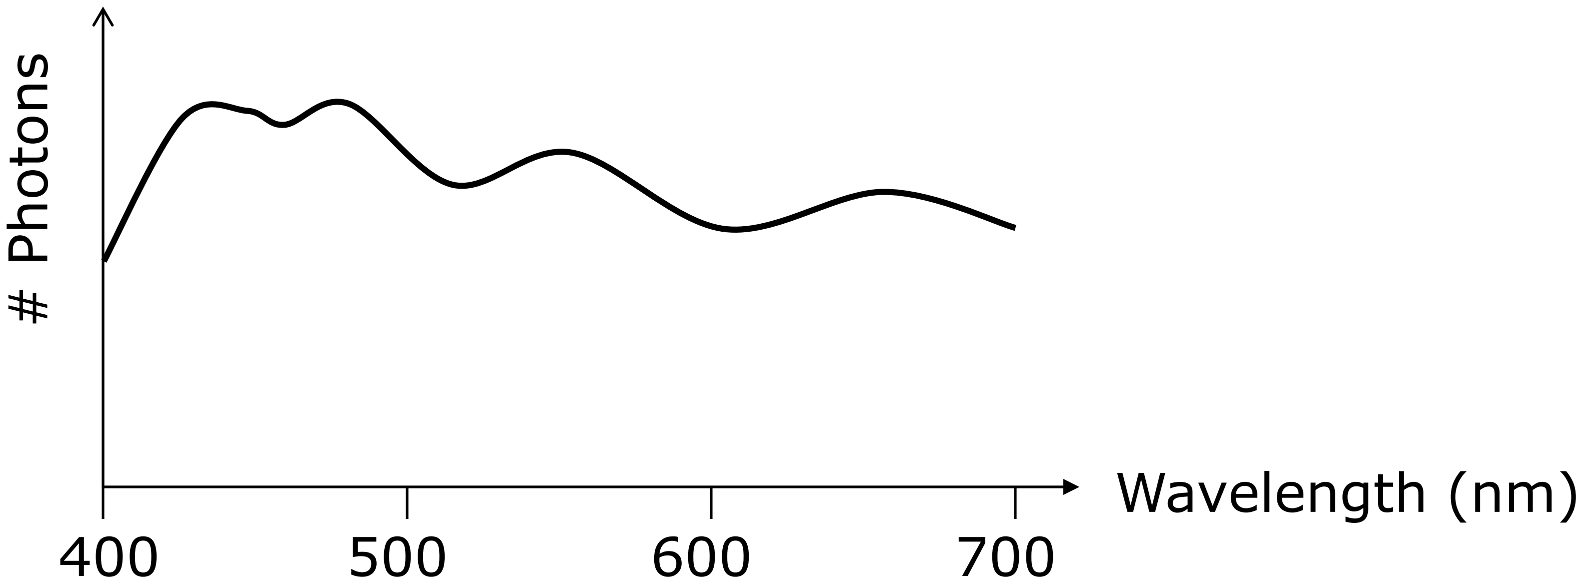

- Note the hump at short wavelengths - the sky is blue

Color and Wavelength

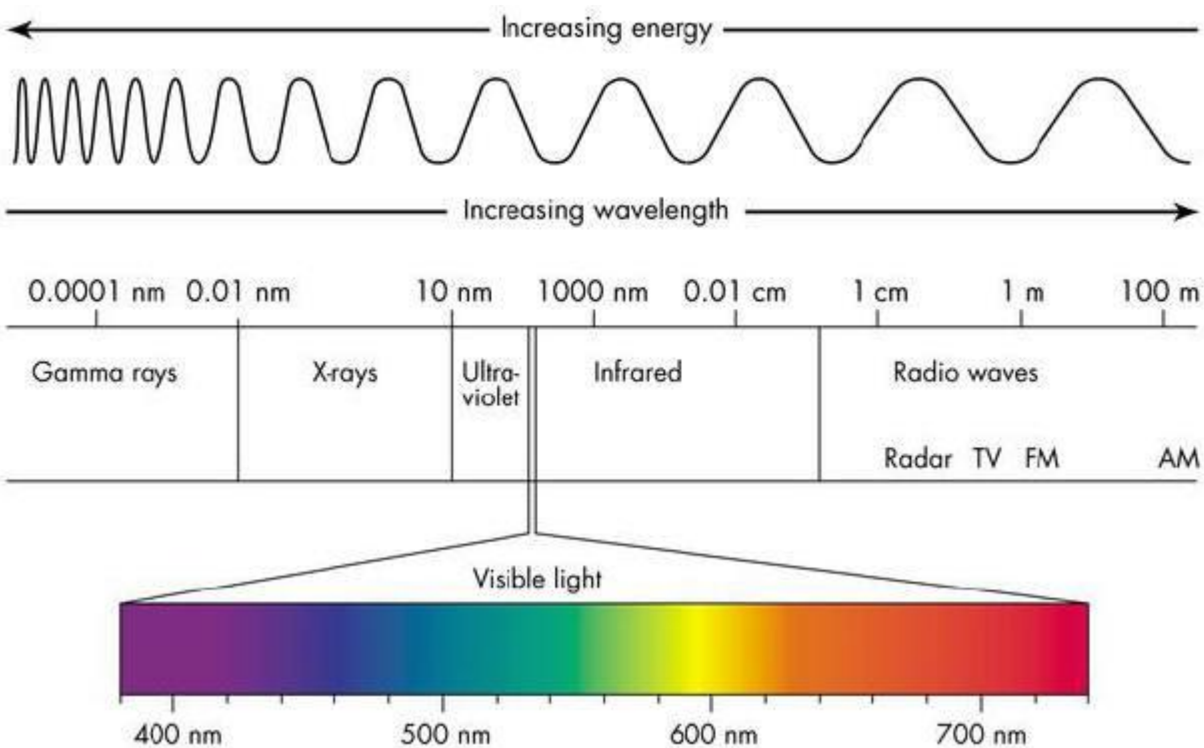

Normal Daylight

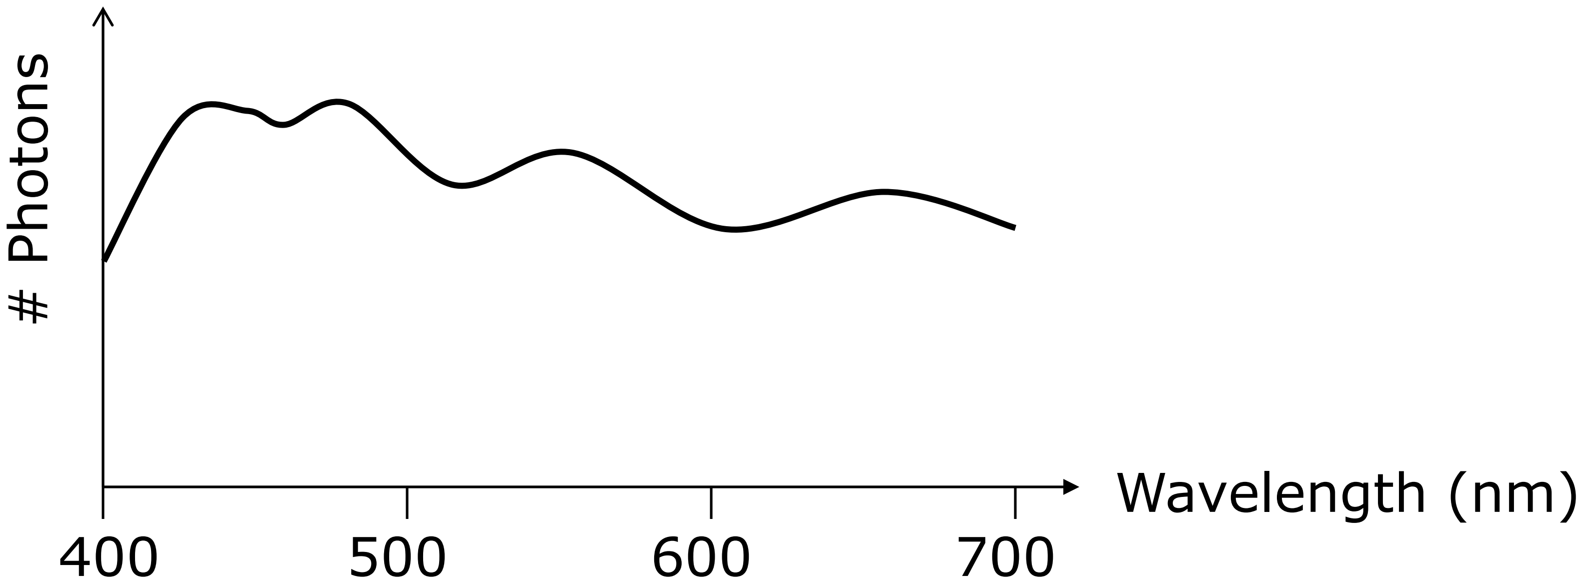

- Note the hump at short wavelengths - the sky is blue

White

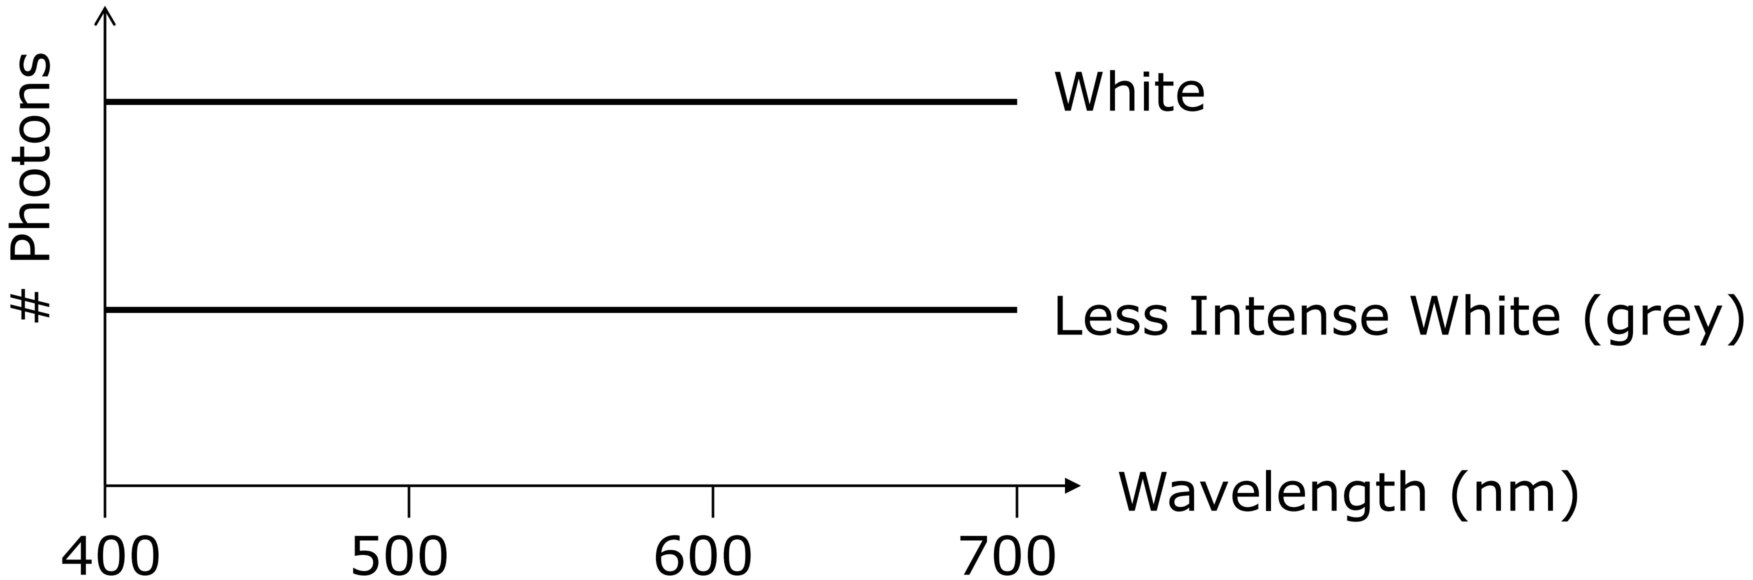

- Note that **color** and **intensity** are technically two different things
- However, in common usage we use color to refer to both
  - White = grey = black in terms of color
- You will have to use context to extract the meaning

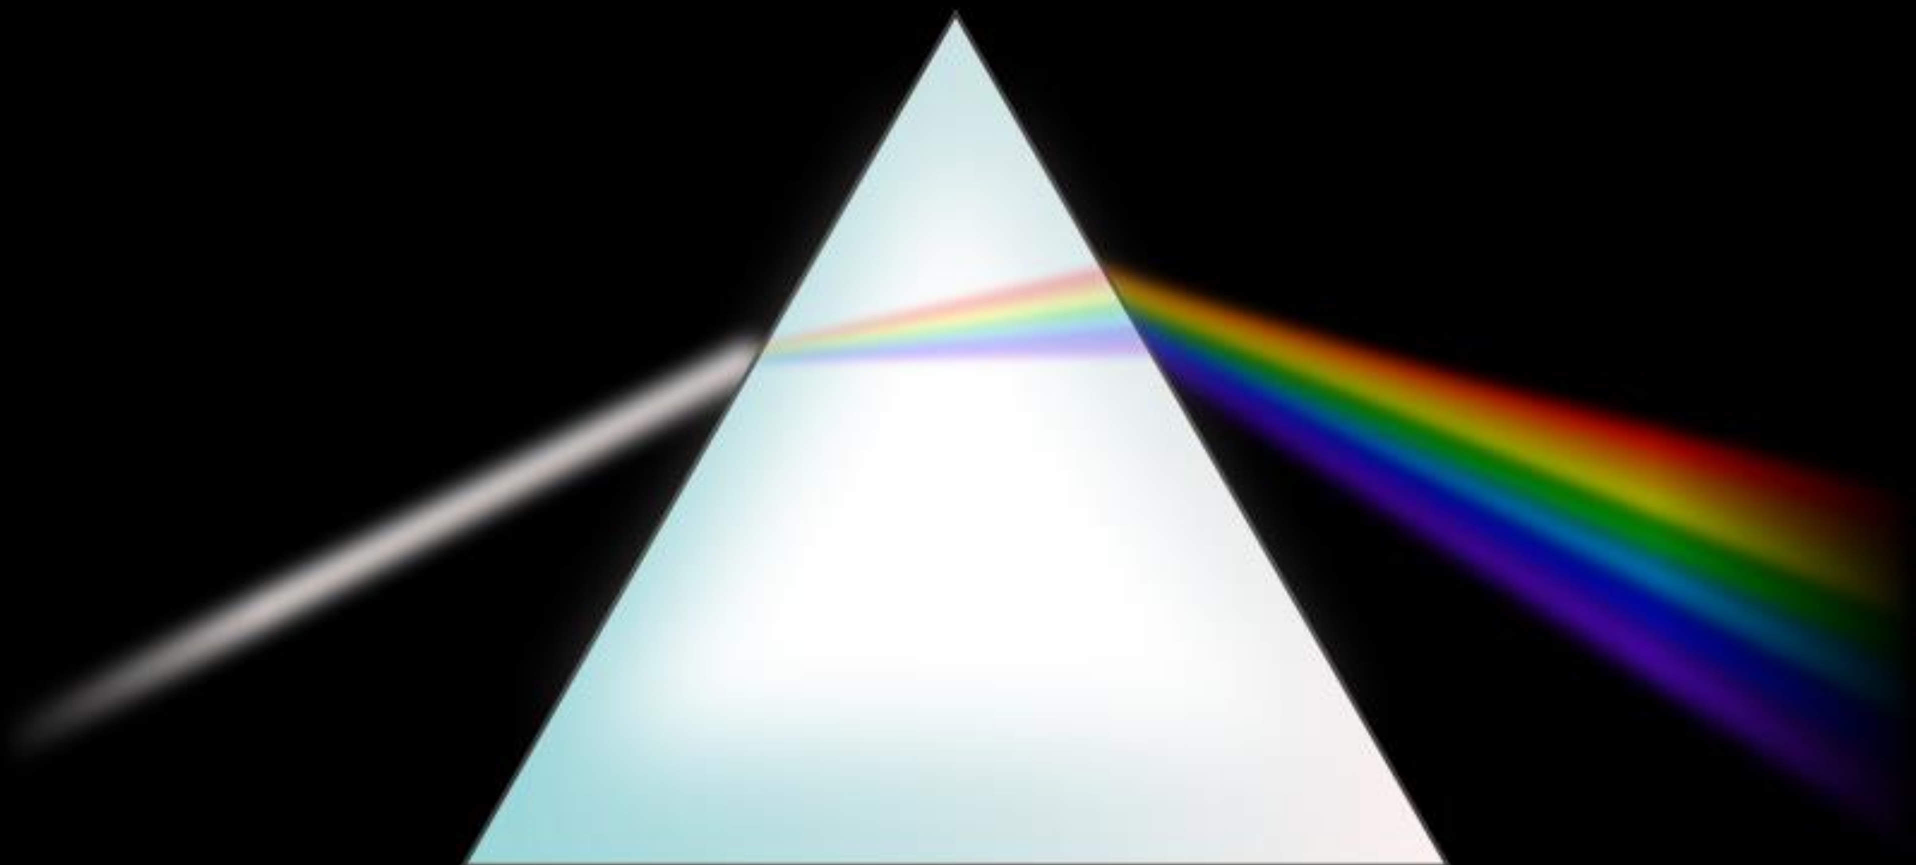

https://commons.wikimedia.org/wiki/File:Prism-rainbow-black.svg

Helium Neon Laser

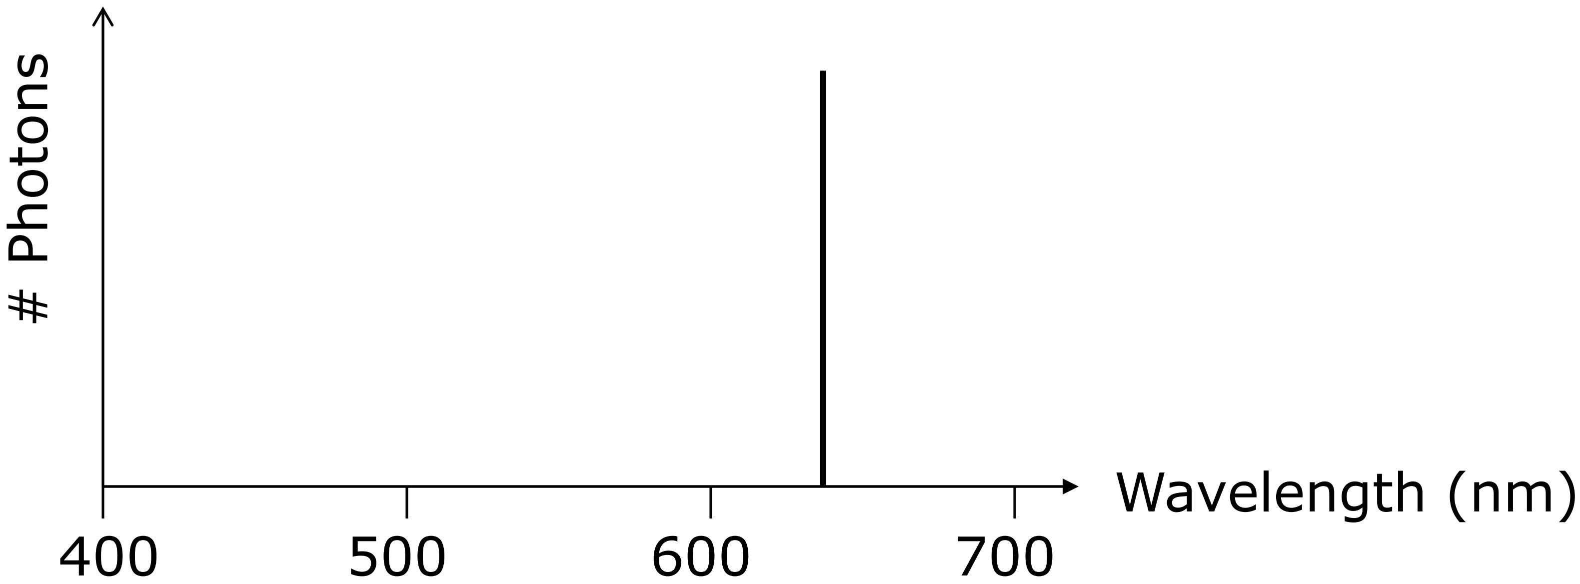

- Lasers emit light at a single wavelength, hence they appear colored in a very “pure” way

Tungsten Lightbulb

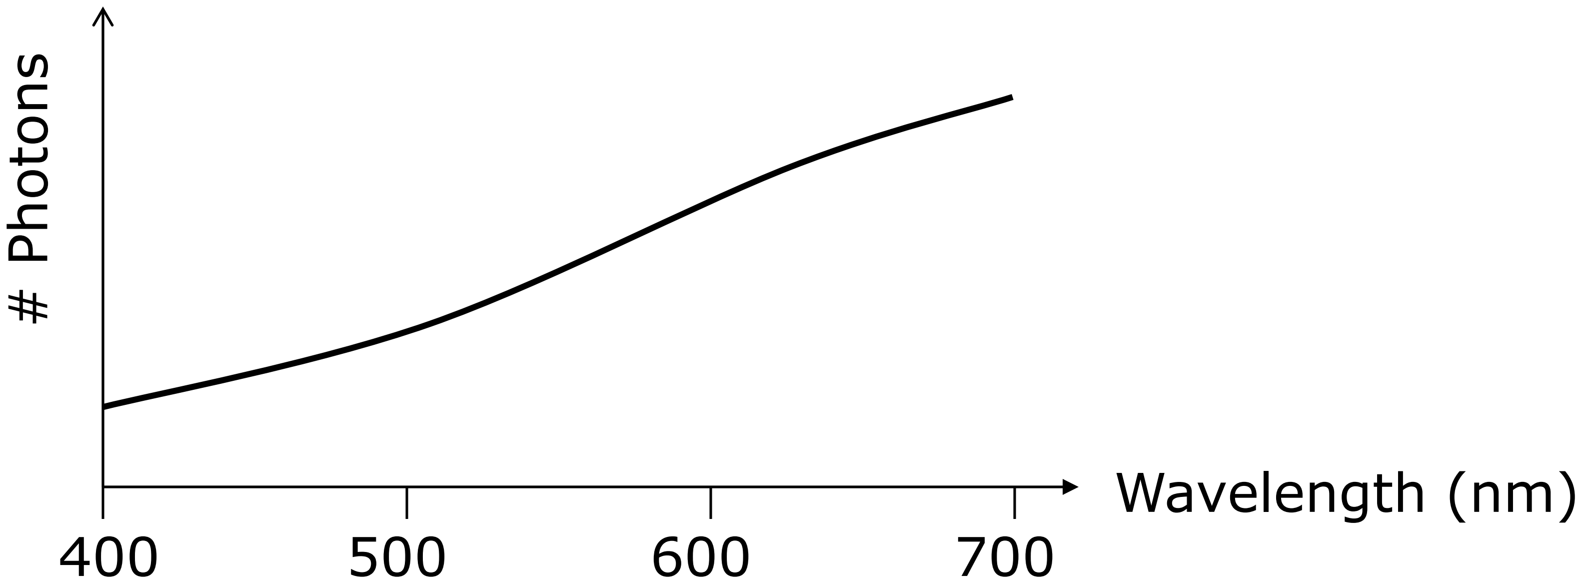

- Most light sources are not anywhere near white
- It is a major research effort to develop light sources with particular properties

Emission vs. Adsorption

- Emission is what light sources do
- Adsorption is what paints, inks, dyes etc. do
- Emission produces light, adsorption removes light
- We still talk about spectra, but now is it the *proportion* of light that is removed at each frequency
  - Note that adsorption depends on such things as the surface finish (glossy, matte) and the substrate (e.g. paper quality)
  - The following examples are qualitative at best

Adsorption Spectra

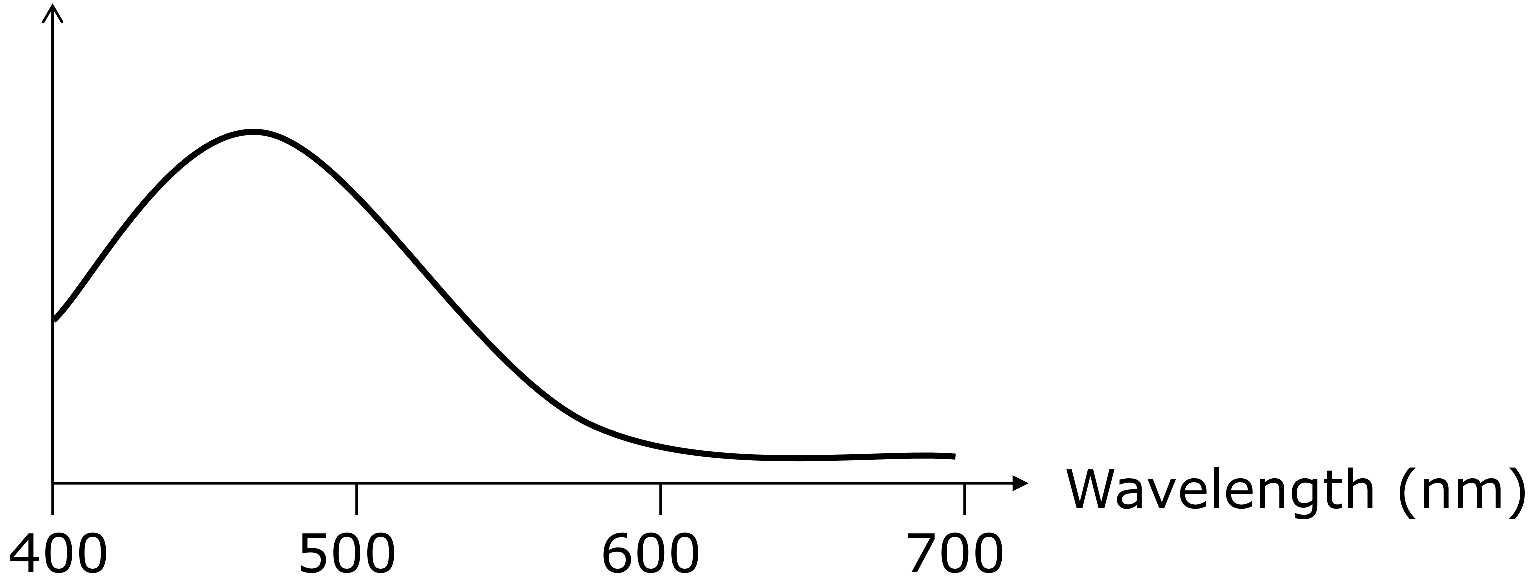

Adsorption Spectra: Red Paint

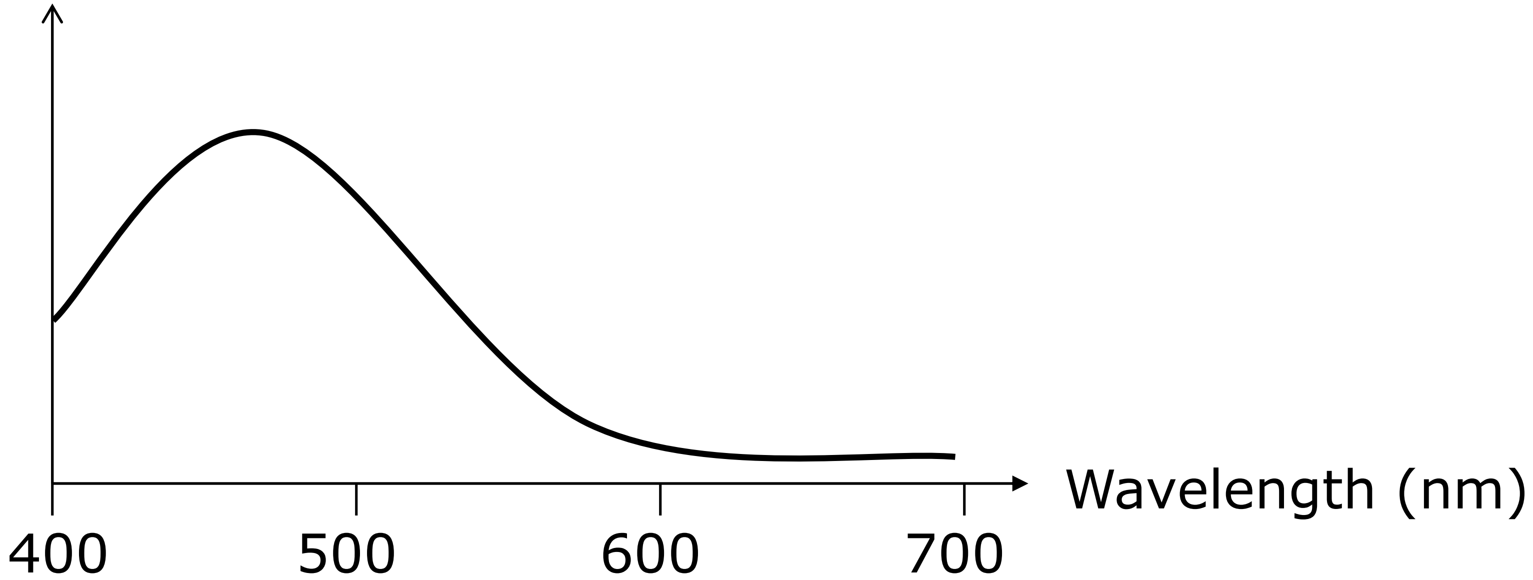

- Red paint absorbs green and blue wavelengths, and reflects red wavelengths, resulting in you seeing a red appearance

Representing Color

- Our task with digital images is to represent color
- You probably know that we use three channels: R, G and B
- We will see why this is perceptually sufficient for display and why it is computationally an approximation
- First, how we measure color

Sensors

- Any *sensor* is defined by its *response* to a frequency distribution
- Expressed as a graph of sensitivity vs. wavelength, $\rho(\lambda)$
  - For each unit of energy at the given wavelength, how much voltage/impulses/whatever the sensor provides
- To compute the response, take the integral $\int \rho(\lambda)E(\lambda)d\lambda$
  - E(λ) is the incoming energy at the particular wavelength
  - The integral multiplies the amount of energy at each wavelength by the sensitivity at that wavelength, and sums them all up

A "Red" Sensor

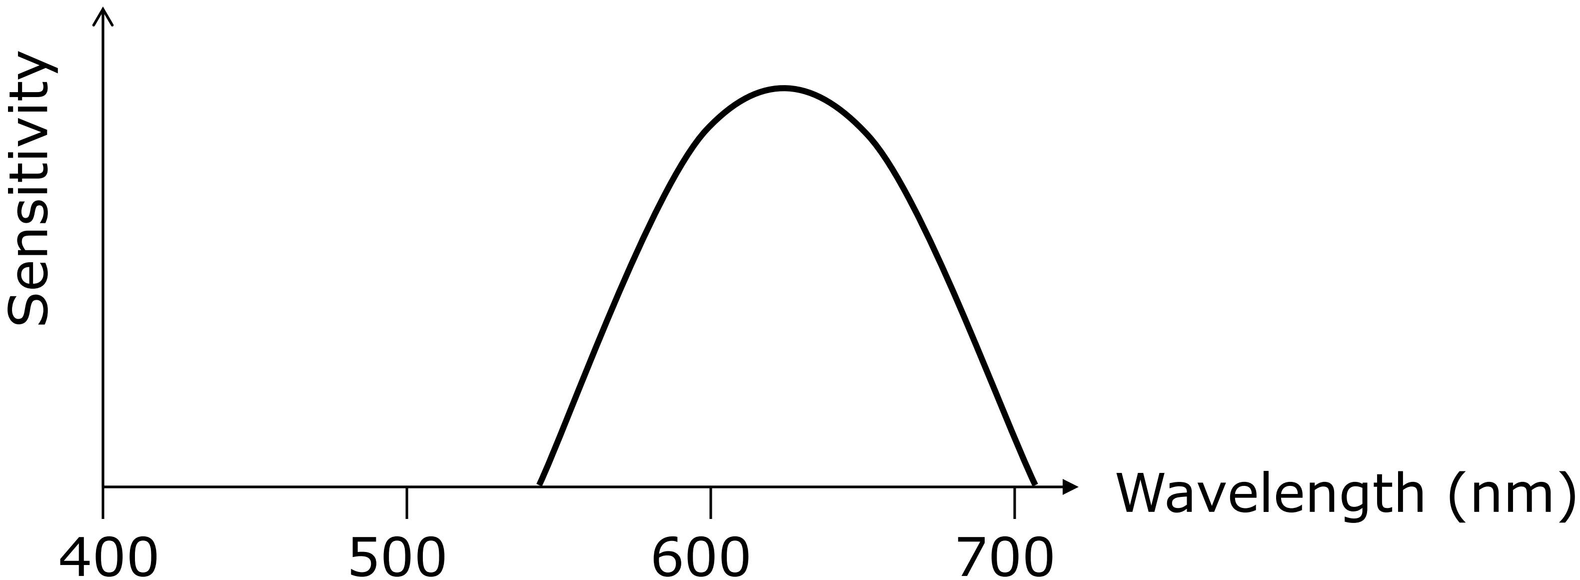

- This sensor will respond to red light, but not to blue light, and a little to green light

The "Red" Sensor Response

Sensor

Color

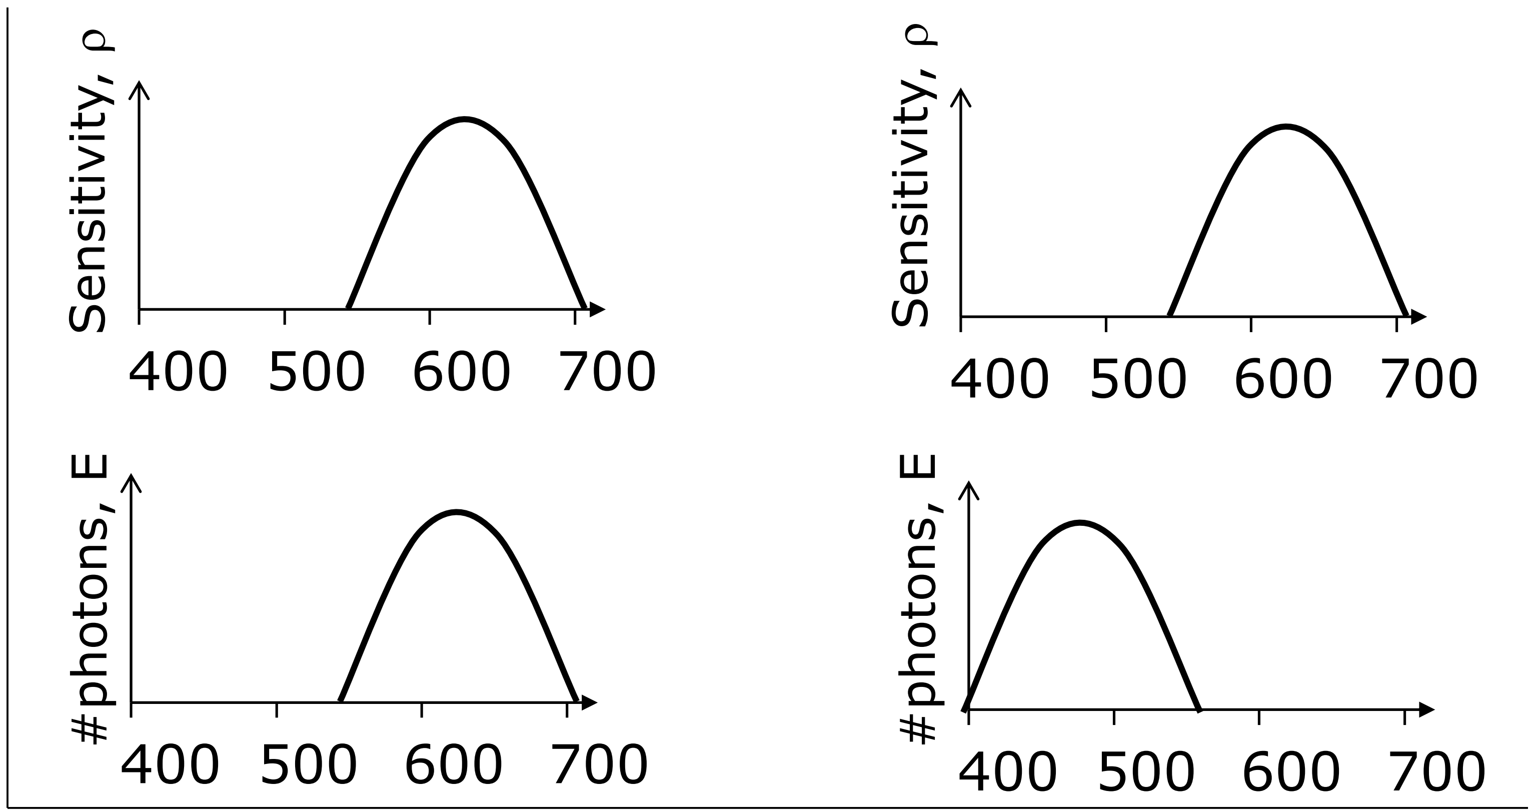

The “Red” Sensor Response

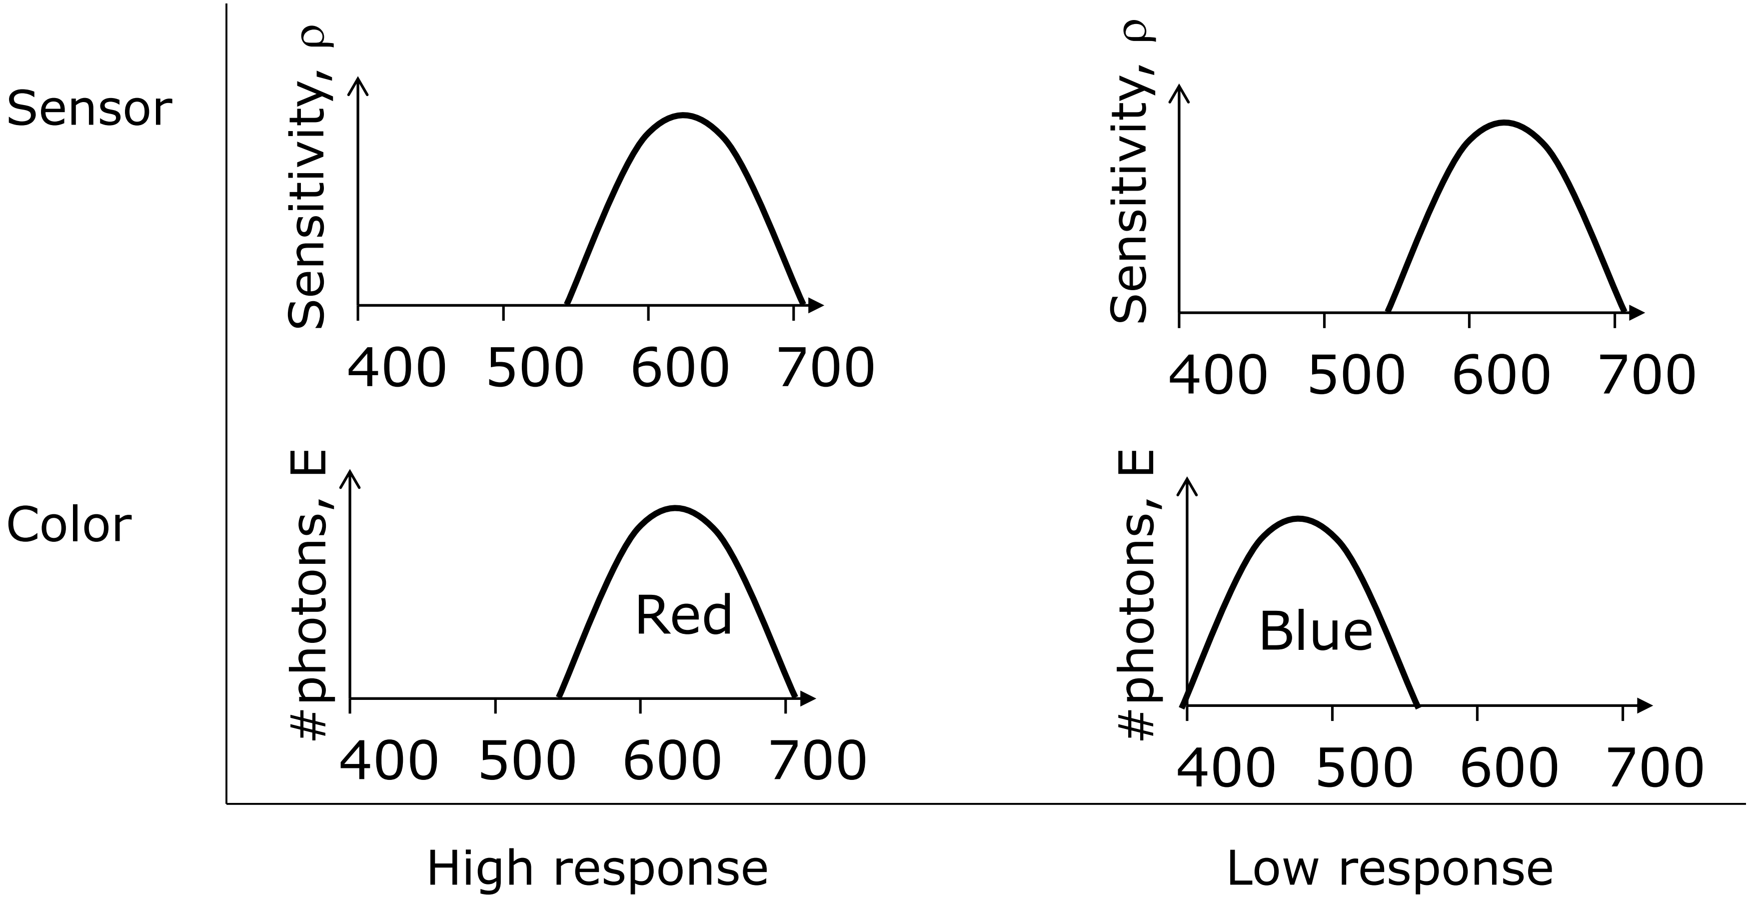

Seeing in Color

- The eye contains *rods* and *cones*
  - Rods work at low light levels and do not see color
    - That is, their response depends only on how many photons, not their wavelength
  - Cones come in three types (experimentally and genetically proven), each responds in a different way to frequency distributions

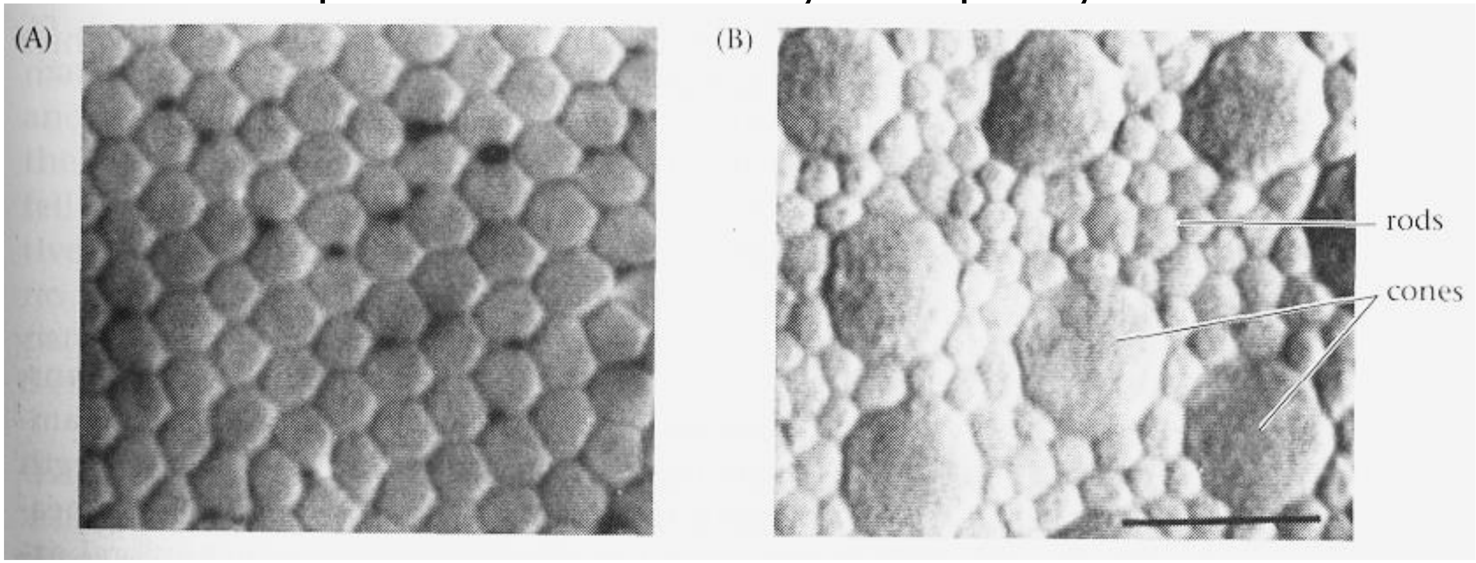

Color receptors

- Each cone type has a different sensitivity curve
  - Experimentally determined in a variety of ways
- For instance, the L-cone responds most strongly to red light
- "Response" in your eye means nerve cell firings
- How you interpret those firings is not so simple …

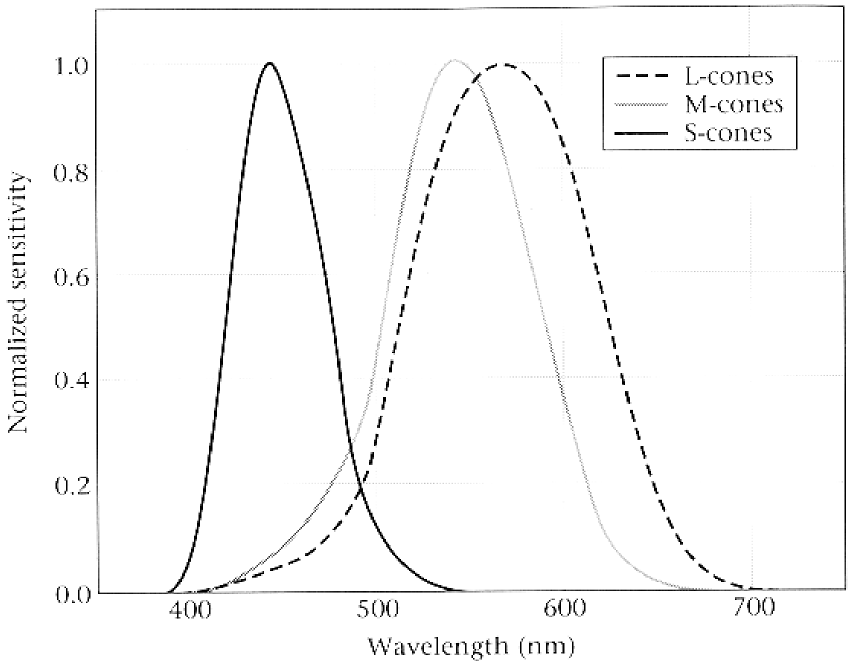

Color Perception

- How your brain interprets nerve impulses from your cones is an open area of study, and deeply mysterious
- Colors may be perceived differently:
  - Affected by other nearby colors
  - Affected by adaptation to previous views
  - Affected by "state of mind"
- Experiment:
  - Subject views a colored surface through a hole in a sheet, so that the color looks like a film in space
  - Investigator controls for nearby colors, and state of mind

The Same Color?

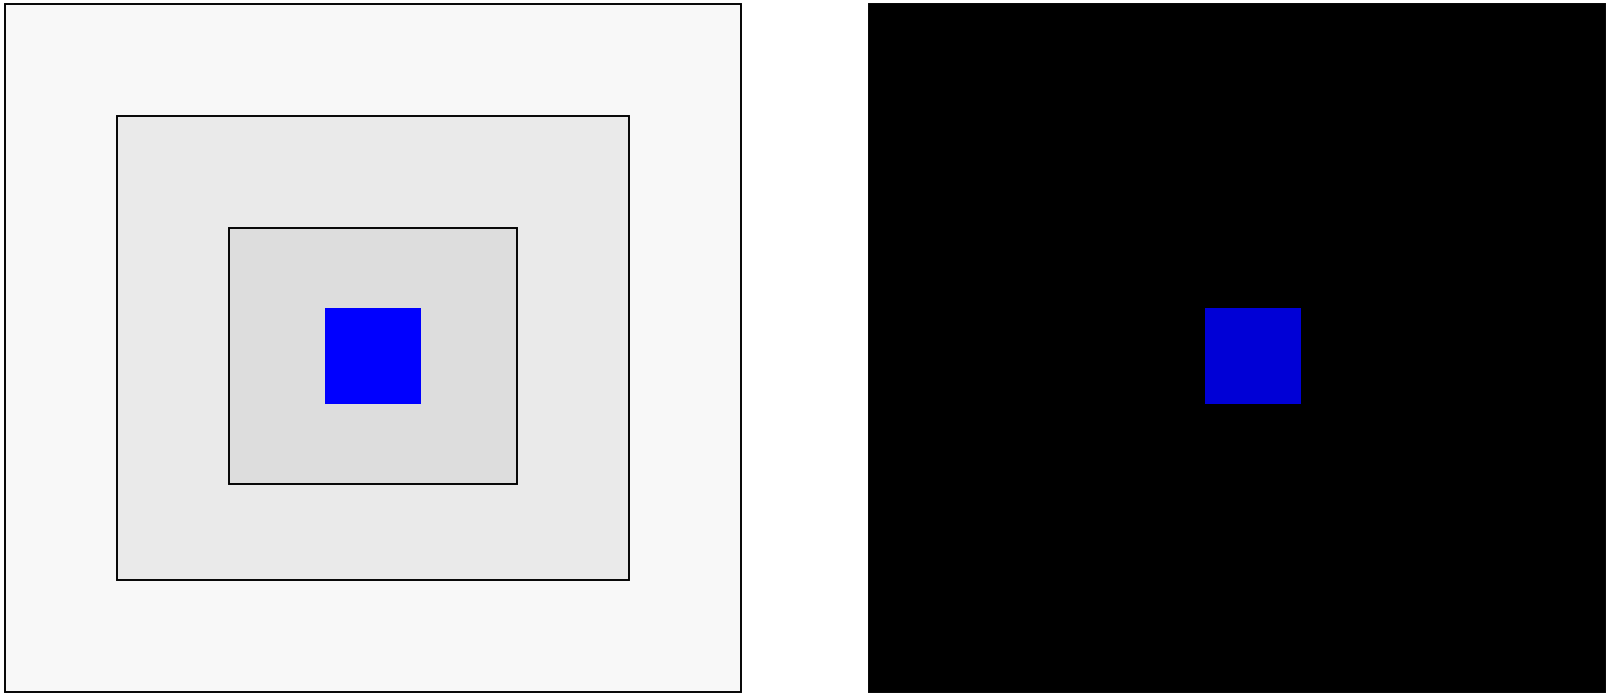

The Same Color?

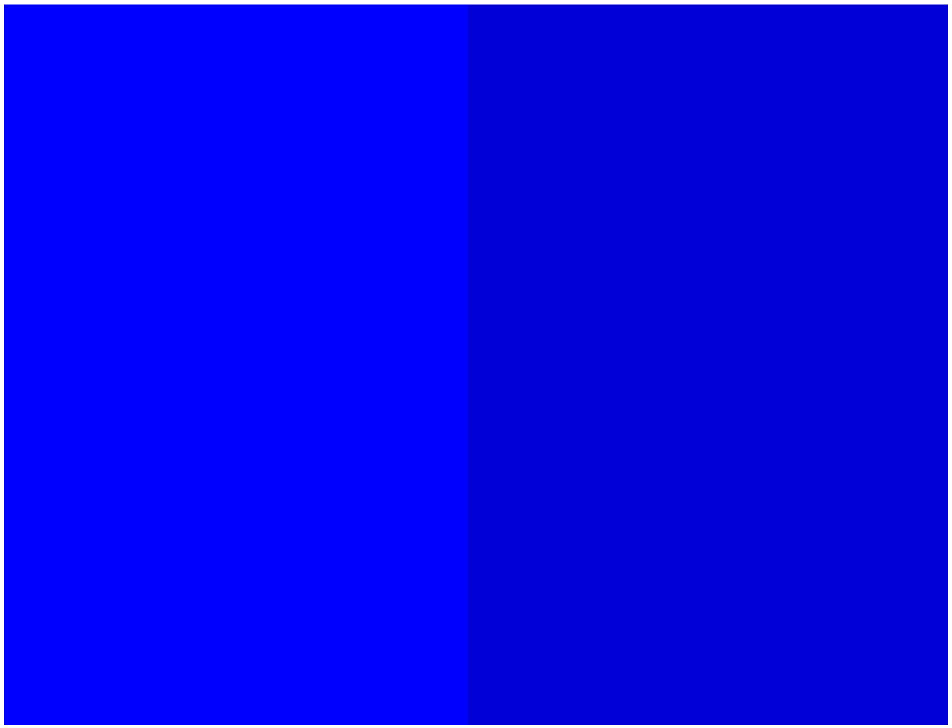

Color Deficiency

- Some people are missing one type of receptor
  - Most common is red-green color blindness in men
  - Red and green receptor genes are carried on the X chromosome - most red-green color blind men have two red genes or two green genes
- Other color deficiencies
  - Anomalous trichromacy, Achromatopsia, Macular degeneration
  - Deficiency can be caused by the central nervous system, by optical problems in the eye, injury, or by absent receptors

Color Deficiency

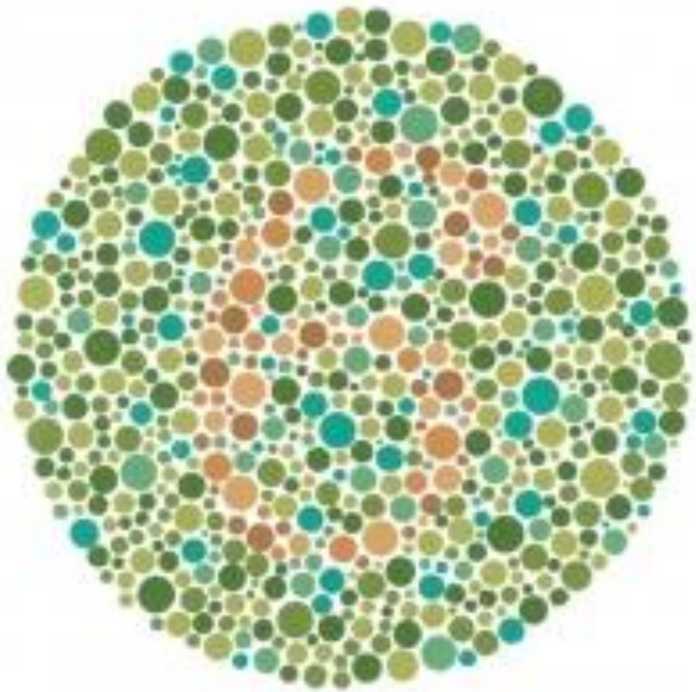

Today

- Digital images
  - Raster vs. Vector
  - Digital images as discrete representations of reality
  - Human perception in deciding resolution and image depth
- Color
- Tri-Chromacy
- Digital Color

Recall

- We're working toward a representation for digital color
- We have seen that humans have three sensors for color vision
- Now, the implications …

Trichromacy

- Experiment:
  - Show a target color *spectrum* beside a user-controlled color
  - User has knobs that adjust *primary sources* to set their color
    - Primary sources are just lights with a fixed spectrum and variable intensity
  - Ask the user to *match* the colors – make their light look the same as the target
- Experiments show that it is possible to match almost all colors using only three primary sources - *the principle of trichromacy*
- Sometimes, have to add light to the *target*
- In practical terms, this means that if you show someone the right amount of each primary, they will perceive the right color
- This was how experimentalists knew there were 3 types of cones

James Maxwell's color-matching experiment

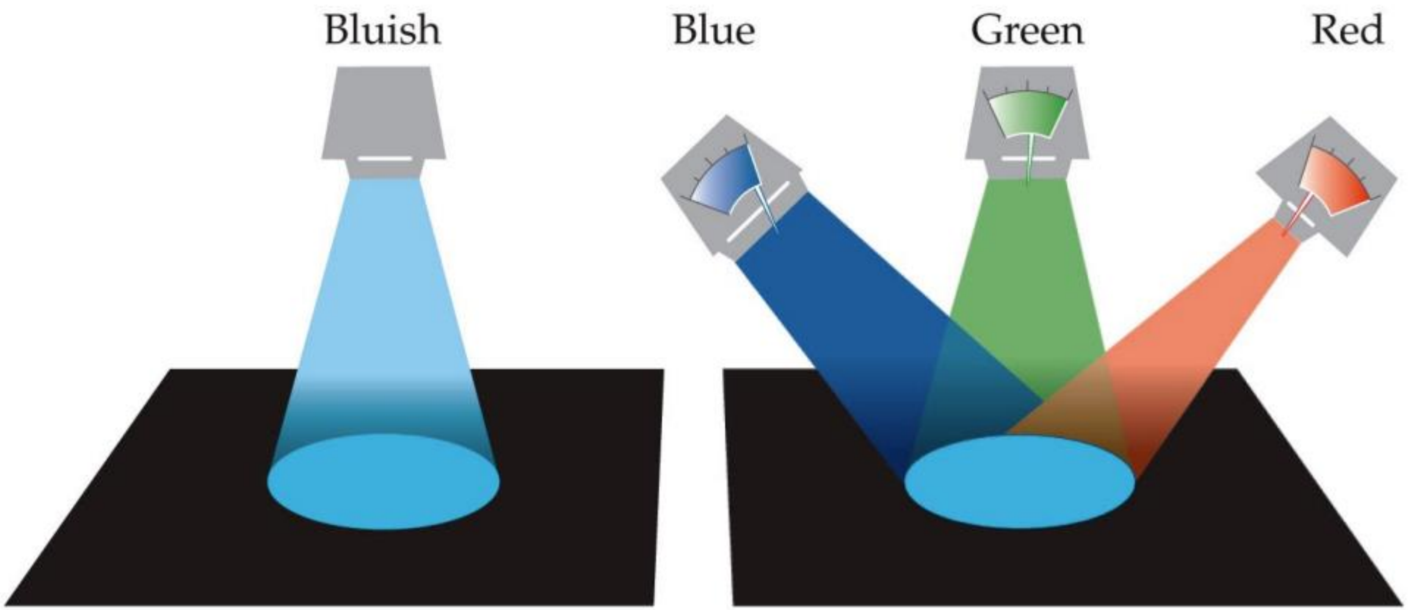

Given any "test" light, you can match it by adjusting the intensities of any three other lights

Slide Credit: Jonathan Pillow

Trichromacy Means…

Color Matching:
People think these two spectra **look the same**
(*monomers*)

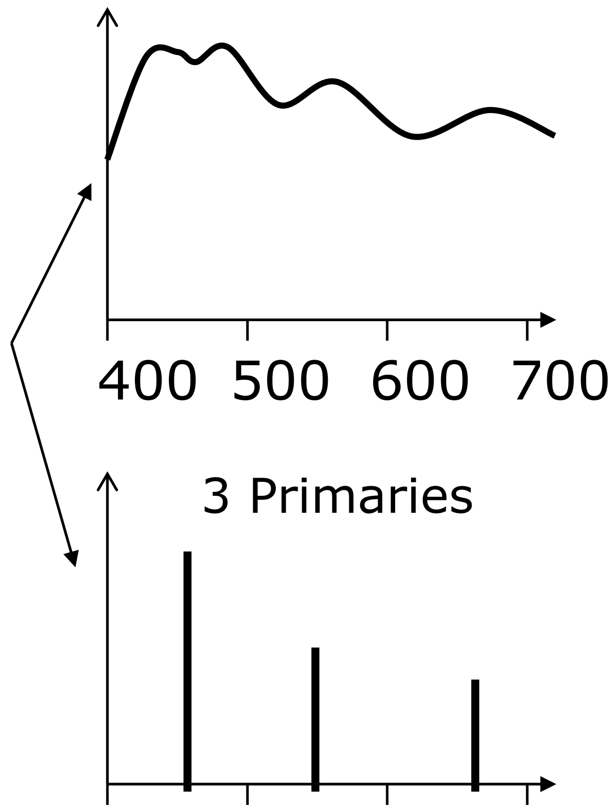

Representing color:
If you want people to "see" the continuous spectrum, you can just show the three primaries
(with varying intensities)

The Math of Trichromacy

- Write primaries as R, G and B
  - We won't precisely define them yet
- Many colors can be represented as a mixture of R, G, B: $M = rR + gG + bB$ (Additive matching)
- Gives a color description system - two people who agree on R, G, B need only supply (r, g, b) to describe a color
- Some colors can't be matched like this, instead, write: $M + rR = gG + bB$ (Subtractive matching)
  - Interpret this as (-r, g, b)
  - Problem for reproducing colors – you can't subtract light using a monitor, or add it using ink

Primaries are Spectra Too

- A primary can be a spectrum
  - Single wavelengths are just a special case

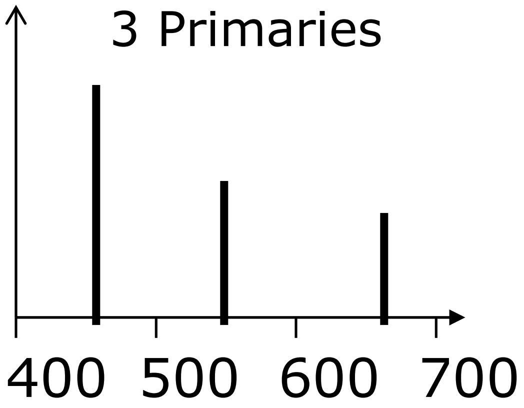

or

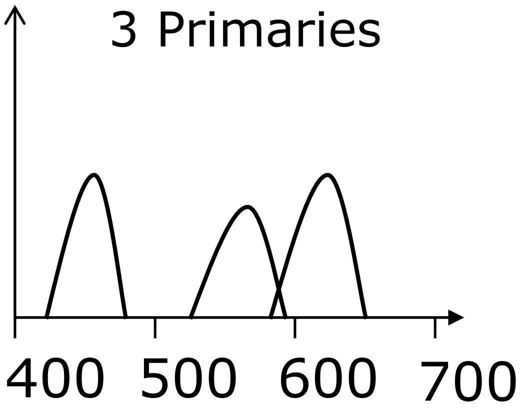

Color Matching

- Given a spectrum, how do we determine how much each of R, G and B to use to match it?
- First step:
  - For a light of unit intensity *at each wavelength*, ask people to match it using some combination of R, G and B primaries
  - Gives you, $r(\lambda)$, $g(\lambda)$ and $b(\lambda)$, the amount of each primary used for wavelength λ
  - Defined for all visible wavelengths, $r(\lambda)$, $g(\lambda)$ and $b(\lambda)$ are the RGB *color matching functions*

The RGB Color Matching Functions

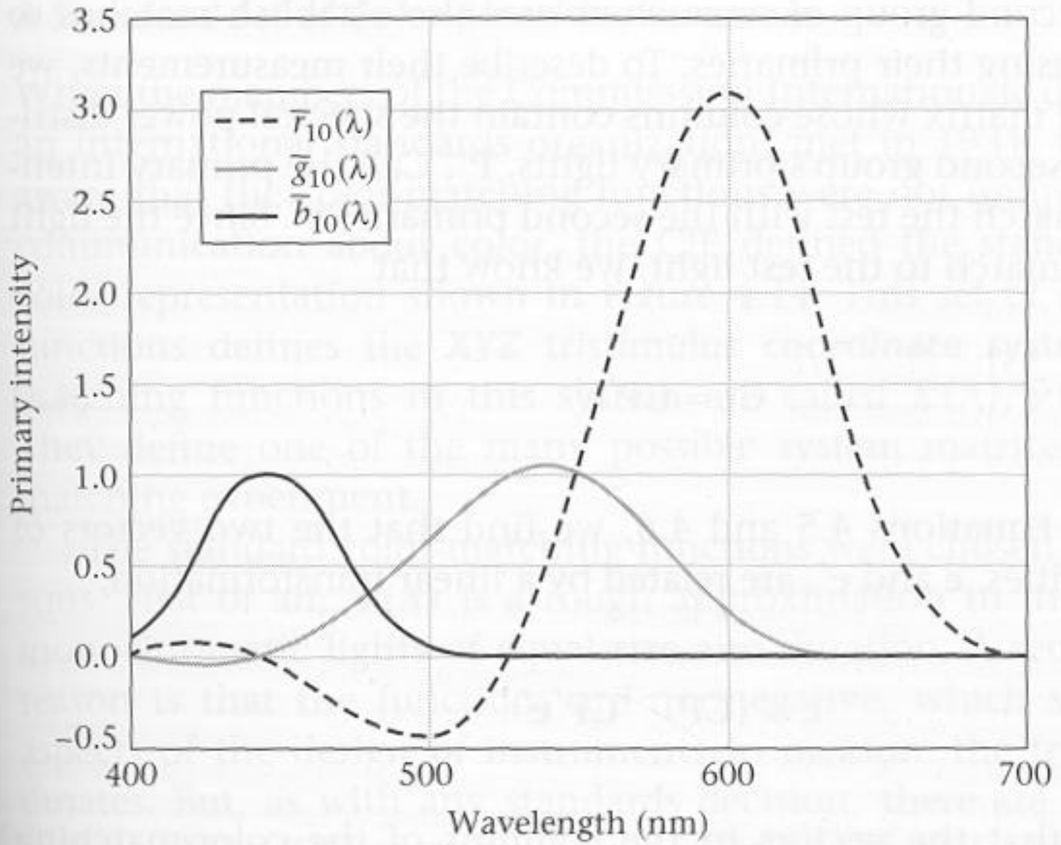

4.13 THE COLOR-MATCHING FUNCTIONS ARE THE ROWS OF THE COLOR-MATCHING SYSTEM MATRIX. The functions measured by Stiles and Burch (1959) using a 10-degree bipartite field and primary lights at the wavelengths 645.2 nm, 525.3 nm, and 444.4 nm with unit radiant power are shown. The three functions in this figure are called $\bar{r}_{10}(\lambda)$, $\bar{g}_{10}(\lambda)$, and $\bar{b}_{10}(\lambda)$.

Computing the Matching

- Given a spectrum, how do we determine how much each of R, G and B to use to match it?
- The spectrum function that we are trying to match, E(λ), gives the amount of energy at each wavelength
- The RGB matching functions describe how much of each primary is needed to give one energy unit's worth of response at each wavelength

$$E = rR + gG + bB$$

$$r = \int r(\lambda)E(\lambda)d\lambda$$

$$g = \int g(\lambda)E(\lambda)d\lambda$$

$$b = \int b(\lambda)E(\lambda)d\lambda$$

Next Time

- Color Spaces
- Color Quantization
- Dithering