Lecture 2

Image formation and Color

Administrative

A0 is out.

- It is ungraded
- Meant to help you with python and numpy basics
- Learn how to do homeworks and submit them on gradescope.

Administrative

Recitation sections on fridays

- (optional)
- Fridays 12:30pm-1:20pm
- JHN 102
- It will be recorded

This week:
Mahtab will go over Linear algebra recap

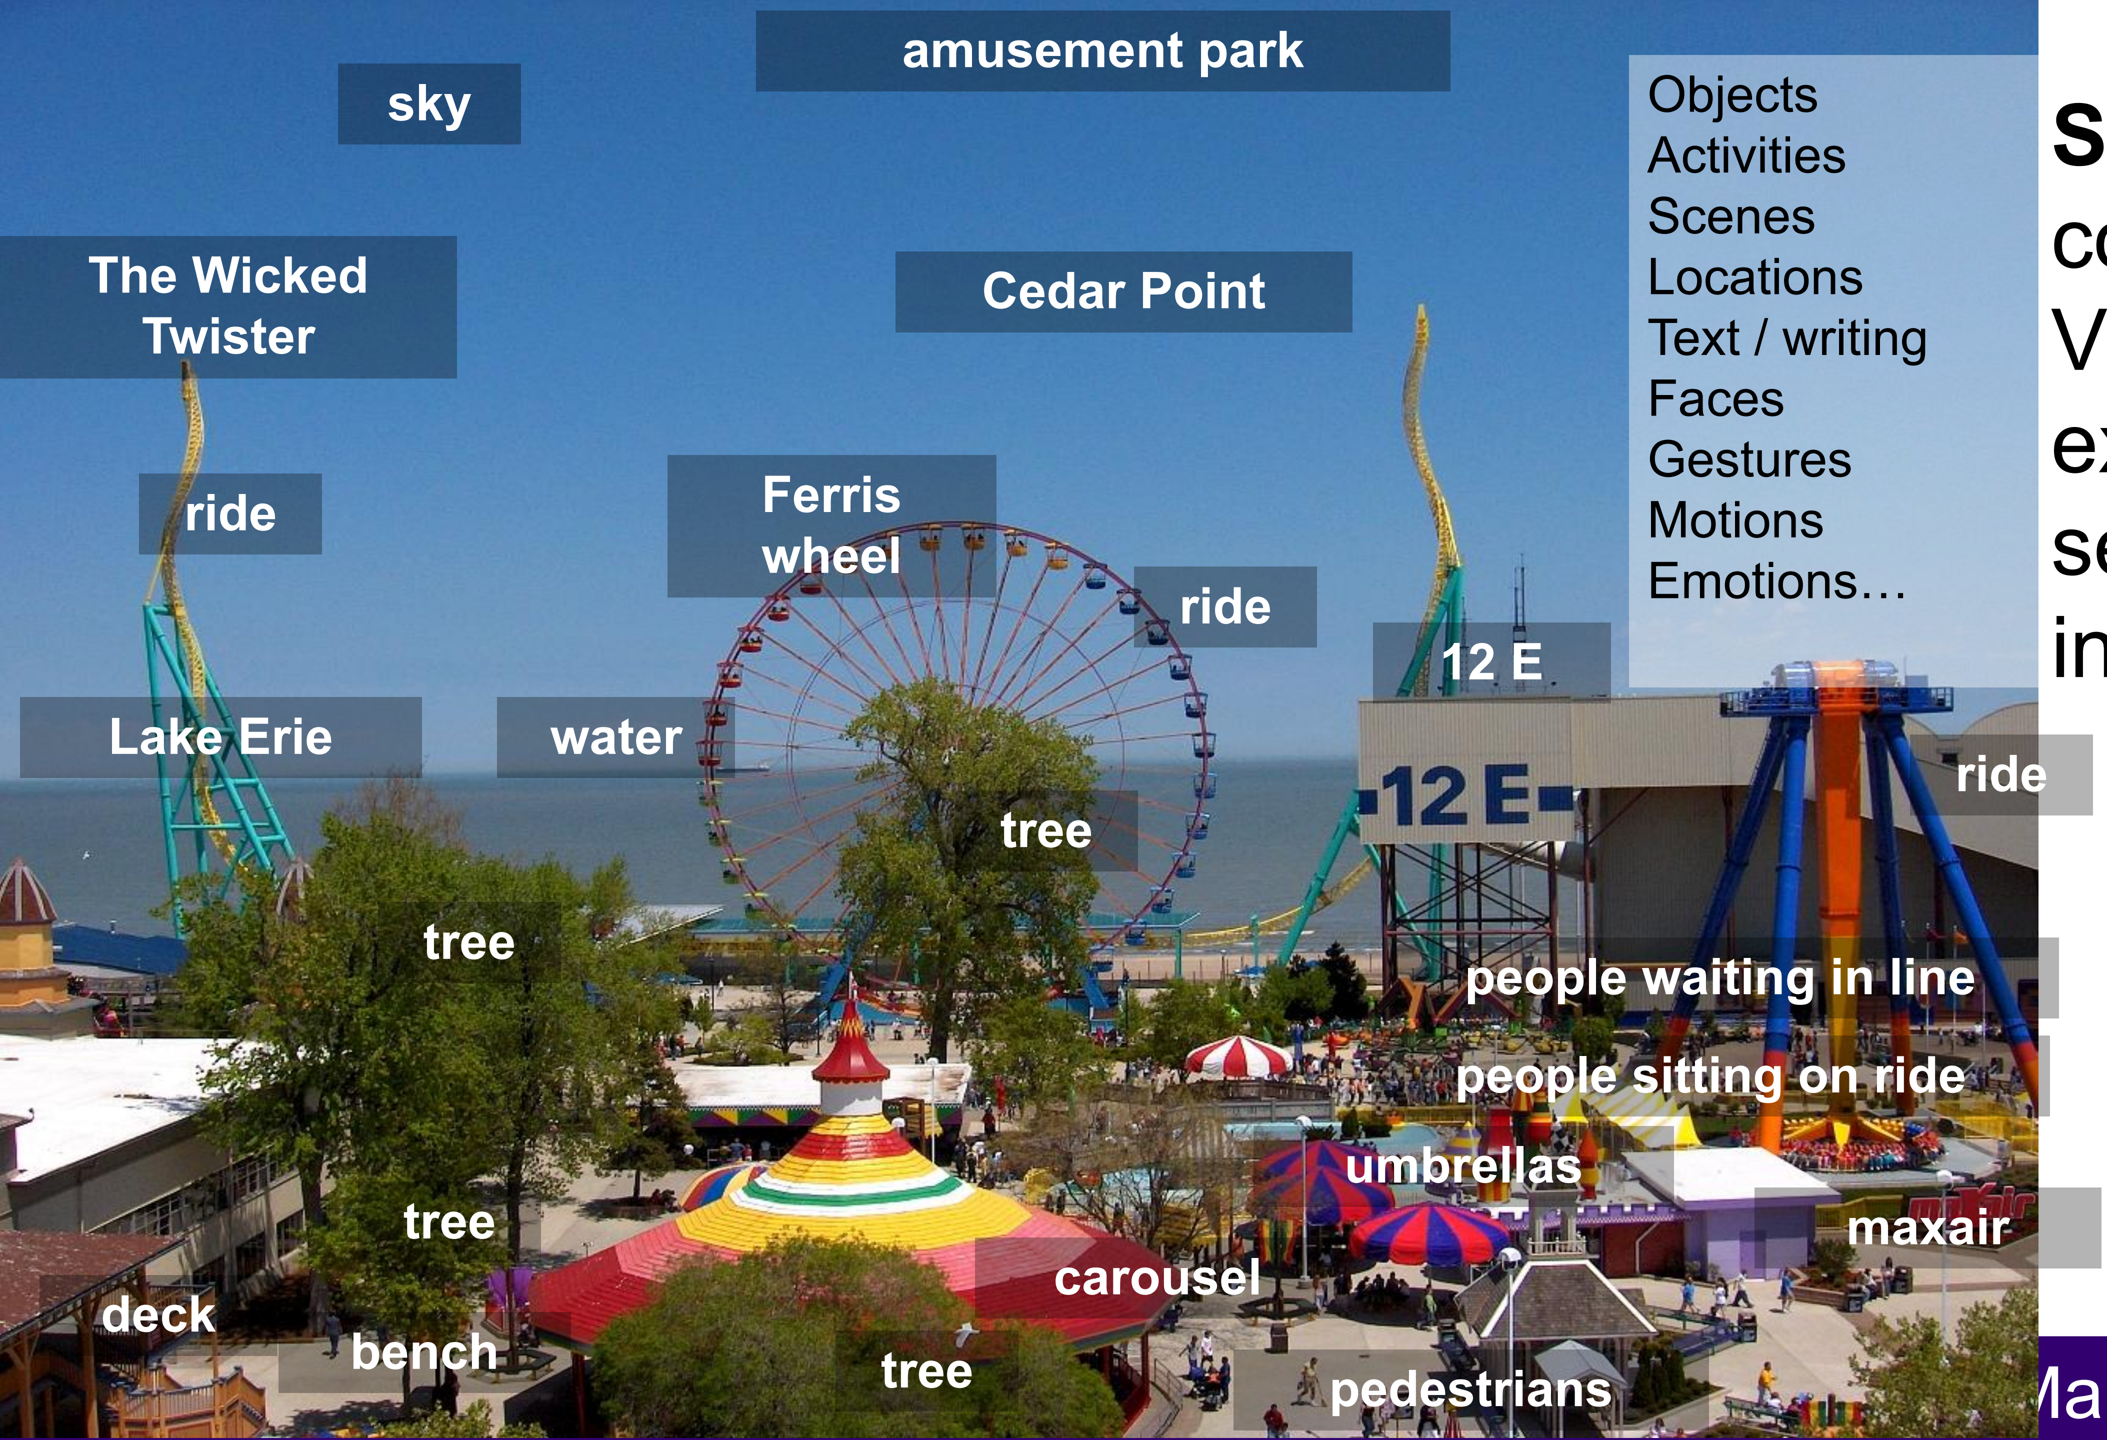

So far: computer Vision extracts semantic information

March 28, 2024

So far: Computer vision extracts geometric 3D information from 2D images

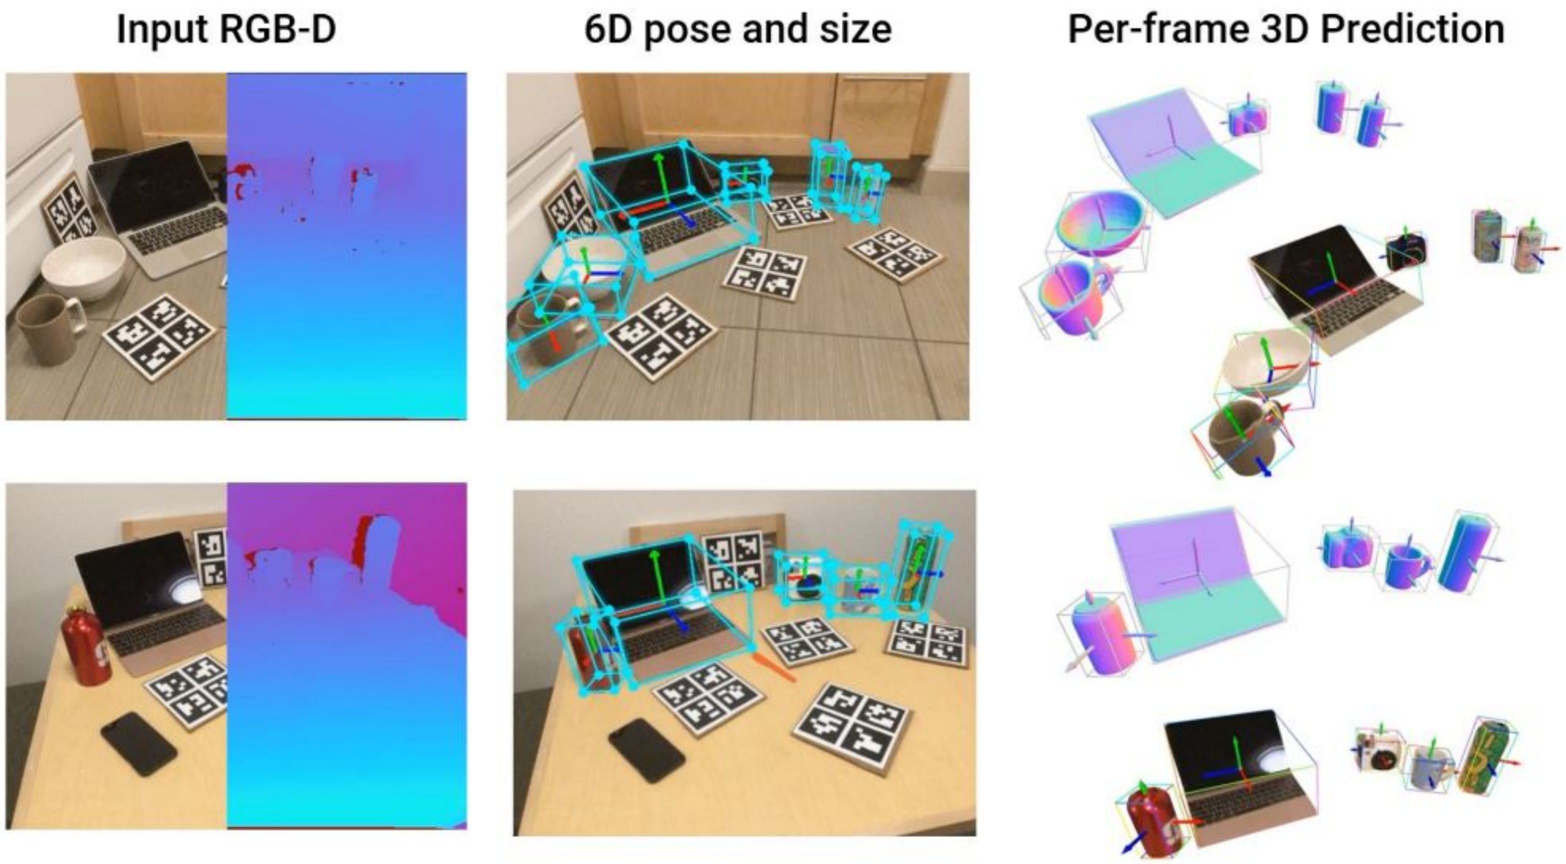

TRI & GATech's ShaPO (ECCV'22): https://zubair-irshad.github.io/projects/ShAPO.html

So far: why is computer vision hard?

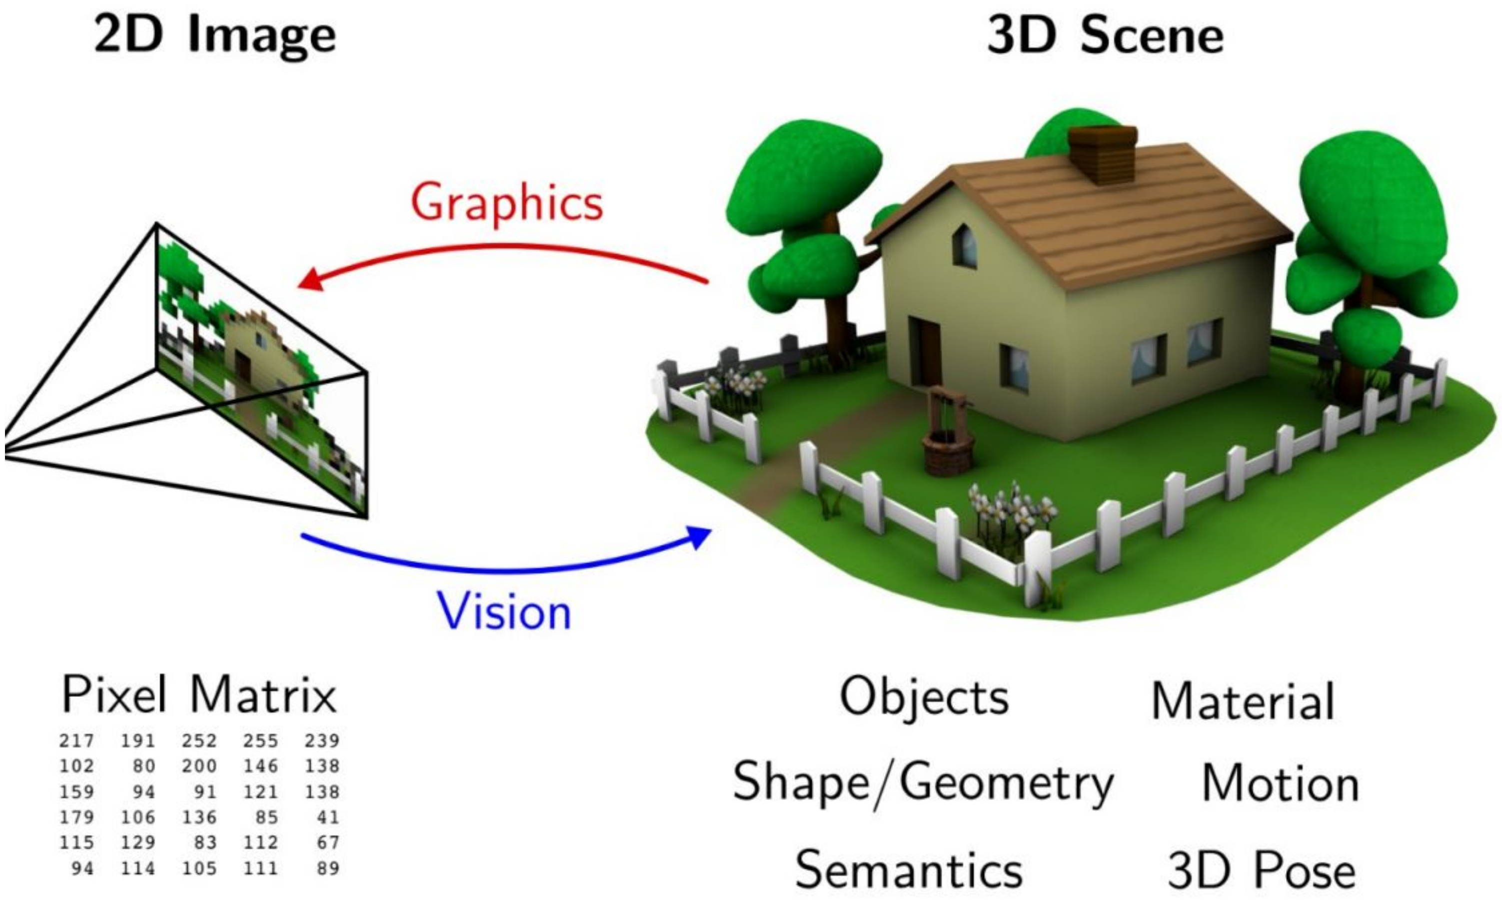

It is an ill posed problem

Today: let's chat about image formation and color?

Q. Why is color important? When does color play a role in helping you make decisions in your day to day?

http://www.hobbylinc.com/gr/pll/pll5019.jpg

Why does a visual system need color? (an incomplete list…)

- To identify edible food

- To distinguish material changes from shading changes.

- To group parts of one object together in a scene.

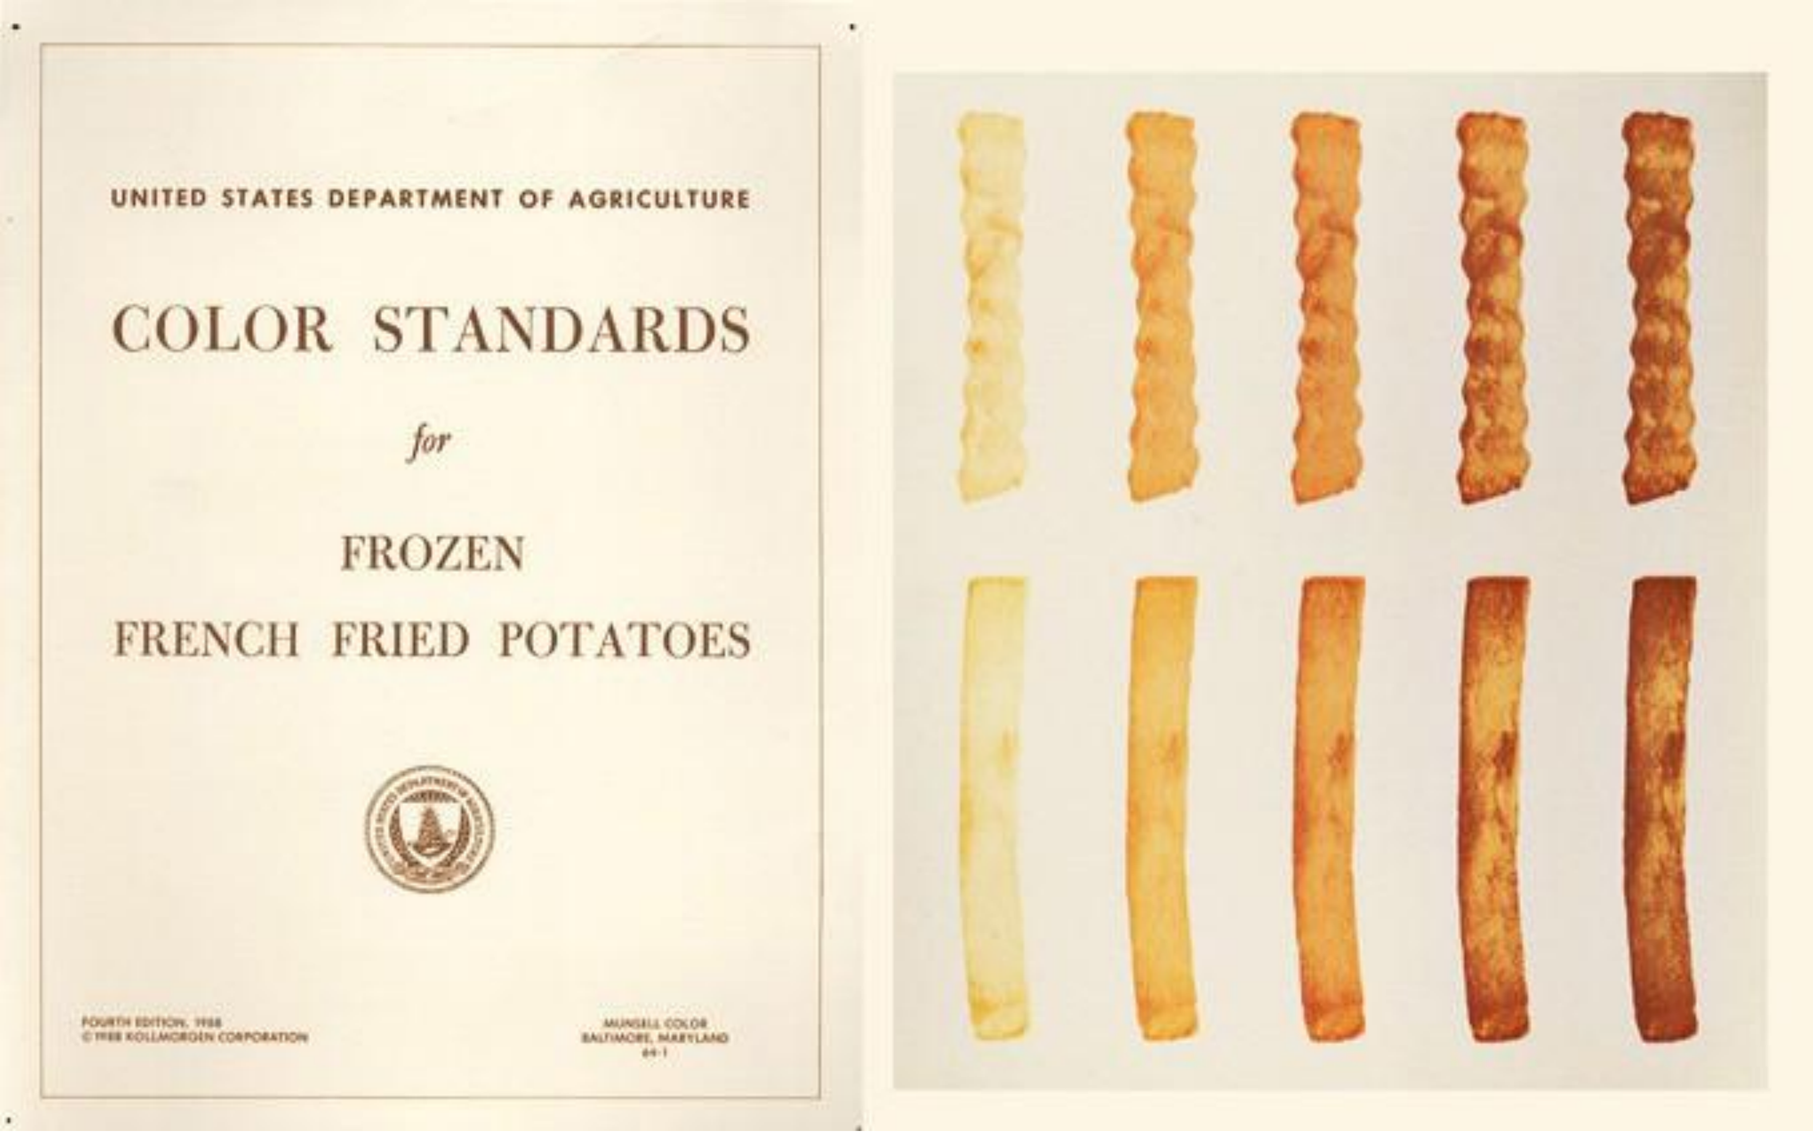

Fun link to the government's recommendations for visual characteristics of agricultural produce

Color standards are important in industry

APL-CP-3 (February 1990)
Excessively rough or bark-like russeting in calyx basin or stem cavity of apples - U.S. Extra Fancy, U.S. Fancy, and U.S. No. 1.

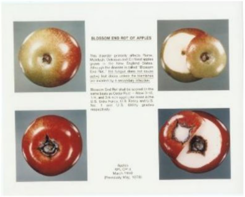

APL-CP-4 (March 1990)
Blossom End Rot of apples.

APL-CP-5 (July 2015)
Internal Stem Bowl Cracking
Scoring guide for injury and damage.

APL-CP-6 (July 2015)
Internal Stem Bowl Cracking
Scoring guide for serious damage.

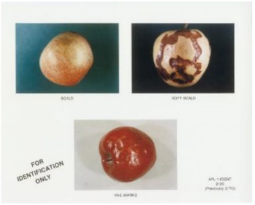

APL-1-IDENT (February 1990)
ID only: Scald; Soft Scald; Hail Marks.

Index of Official Visual Aids (January 2017)

IP Color Trademarks

http://blog.patents-tms.com/?p=52

CURRENTLY REGISTERED COLOR TRADEMARKS

A color trademark is a non-conventional trademark where at least one color is used to identify the commercial origin of a product or service. A color trademark must meet the same requirements of a conventional trademark. Thus, the color trademark must either be inherently distinctive or have acquired secondary meaning. To be inherently distinctive, the color must be arbitrarily or suggestively applied to a product or service. In contrast, to acquire secondary meaning, consumers must associate the color used on goods or services as originating from a single source. Below is a selection of some currently registered color trademarks in the U.S. Trademark Office:

MARK/COLOR(S)/OWNER:

BANK OF AMERICA 500
blue, red & grey
Bank of America Corporation

NATIONAL CAR RENTAL
green
NCR Affiliate Servicer, Inc.

FORD
blue
Ford Motor Company

VISTEON
orange
Ford Motor Company

76
red & blue
ConocoPhillips Company

THE HOME DEPOT
orange
Homer TLC, Inc.

HONDA
red
Honda Motor Co., Ltd.

M MARATHON
brown, orange, yellow
Marathon Oil Company

M MARATHON
gray, black & white
Marathon Oil Company

COSTCO
red
Costco Wholesale Membership, Inc.

TARGET
red
Target Brands, Inc.

AT&T
light blue, dark blue and gray
AT&T Corp.

Filed under: Trademark by admin

TEENAGE MUTANT NINJA TURTLES MUTANTS & MONSTERS
red, green, yellow, black, grey and
white Mirage Studios, Inc.

VW
silver, metallic blue, black and white
Volkswagen Aktiengesellschaft Corp.

Uses of color in computer vision

Color histograms for indexing and

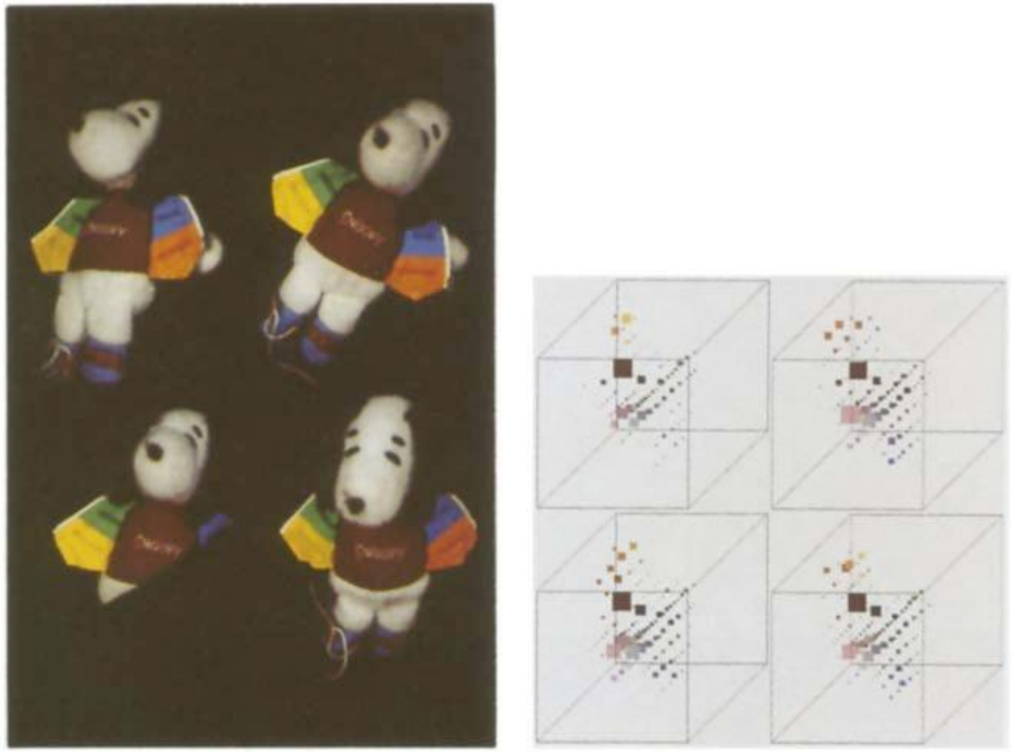

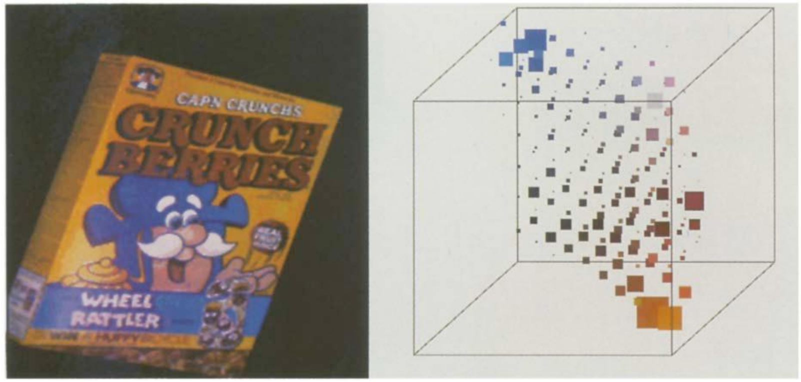

Swain and Ballard, Color Indexing, IJCV 1991.

Uses of color in computer vision

Skin detection for
content moderation

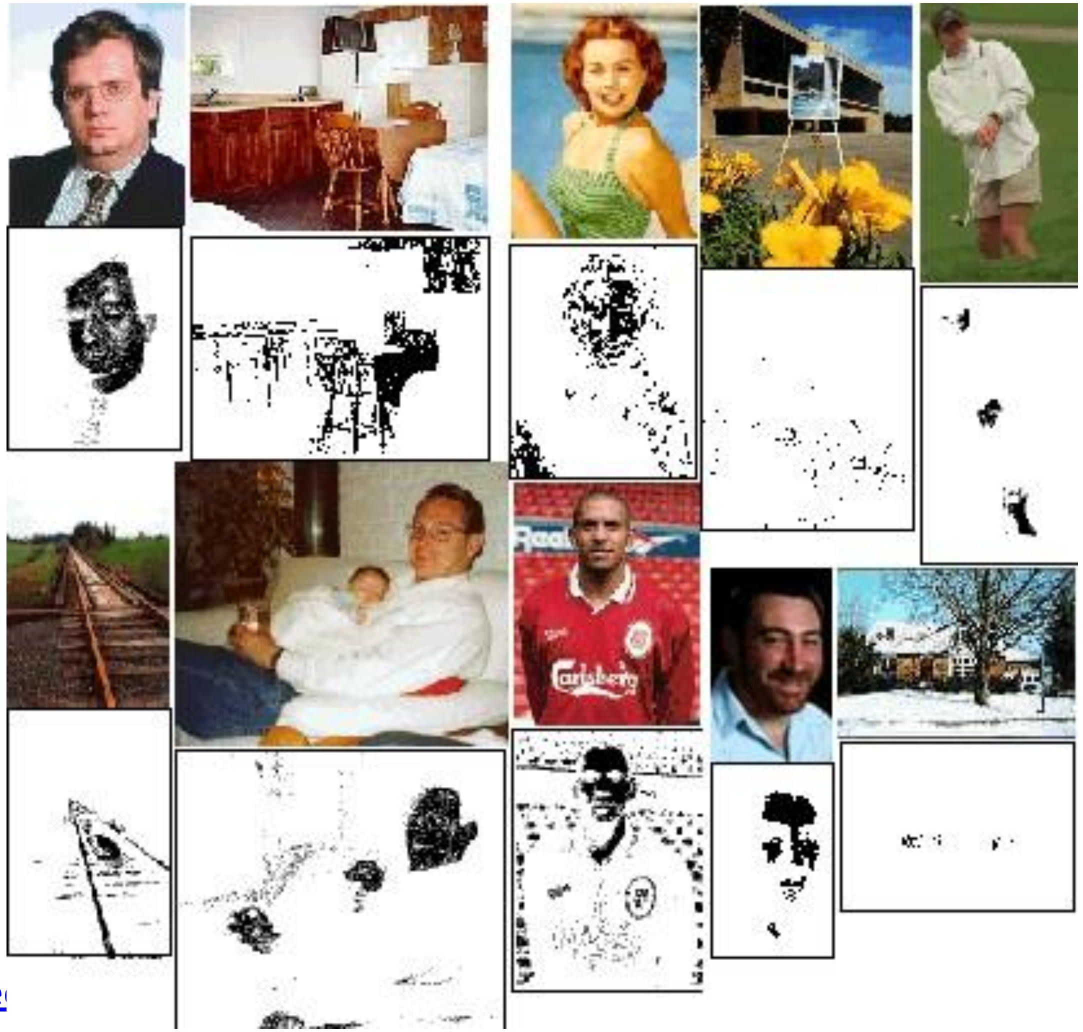

M. Jones and J. Rehg, Statistical Color Models with Application to Skin Dete

8

Uses of color in computer vision

Image segmentation and retrieval

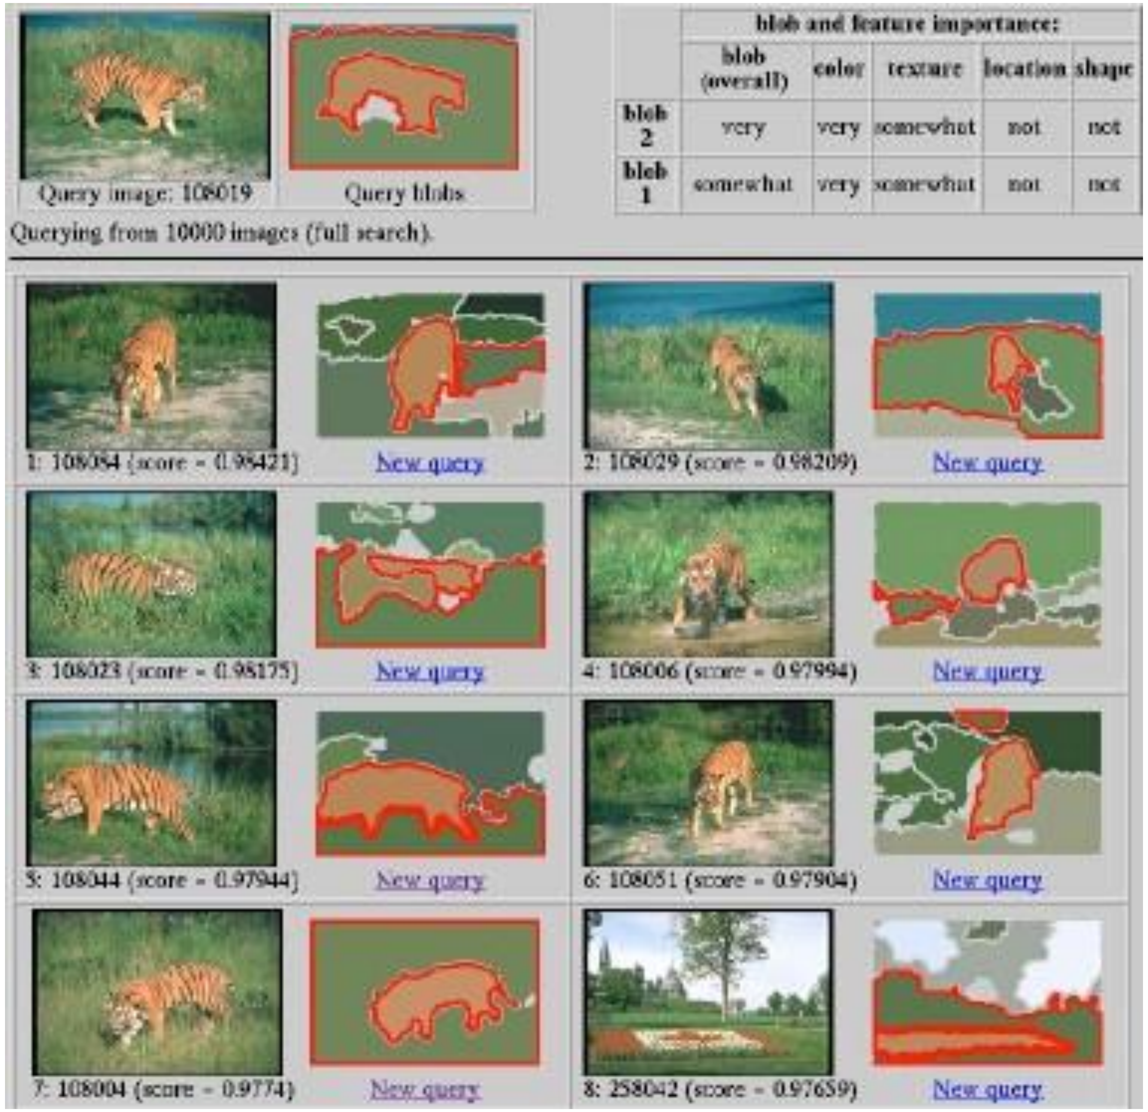

C. Carson, S. Belongie, H. Greenspan, and Ji. Malik, Blobworld: Image segmentation using Expectation-Maximization and its application to image querying, ICVIS 1999.

Uses of color in computer vision

Robot soccer

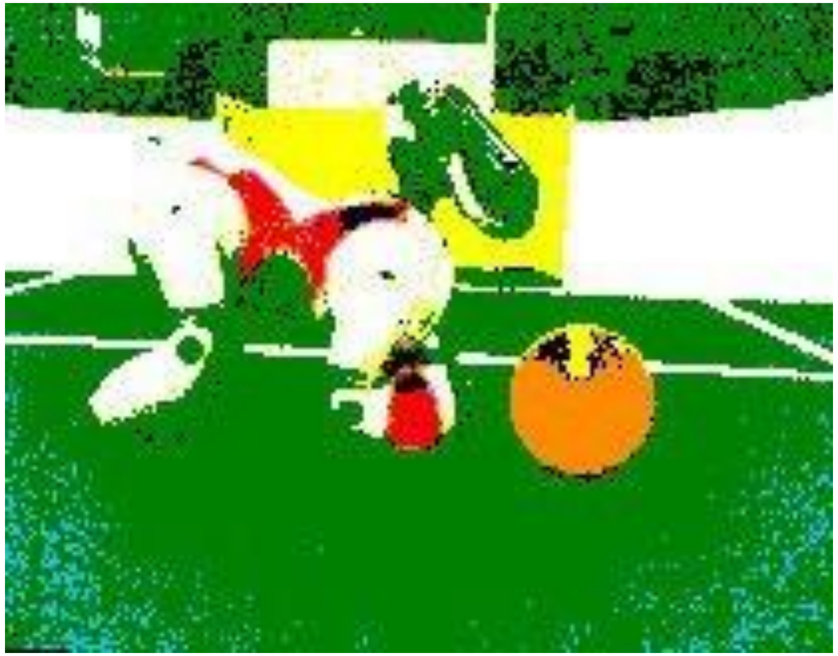

M. Sridharan and P. Stone, Towards Eliminating Manual Color Calibration at RoboCup. RoboCup-2005: Robot Soccer World Cup IX, Springer Verlag, 2006

Uses of color in computer vision

Building appearance models for tracking

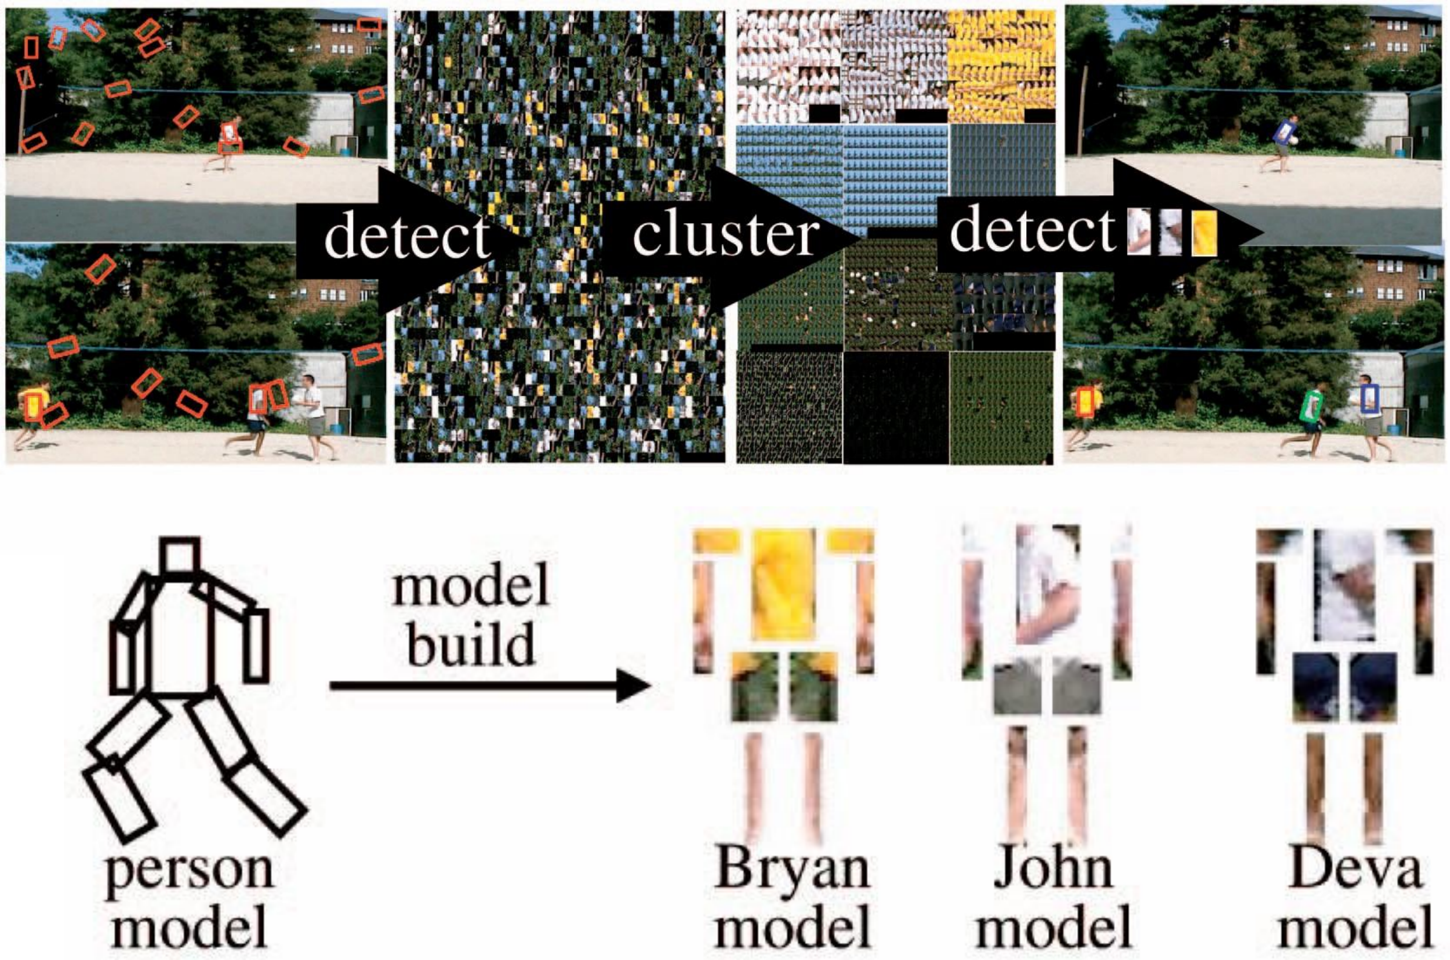

D. Ramanan, D. Forsyth, and A. Zisserman. Tracking People by Learning their Appearance. PAMI 2007.

Today's agenda

- **Image formation**
- Physics of Color
- Color matching
- Color spaces
- Image sampling and quantization

Image Formation

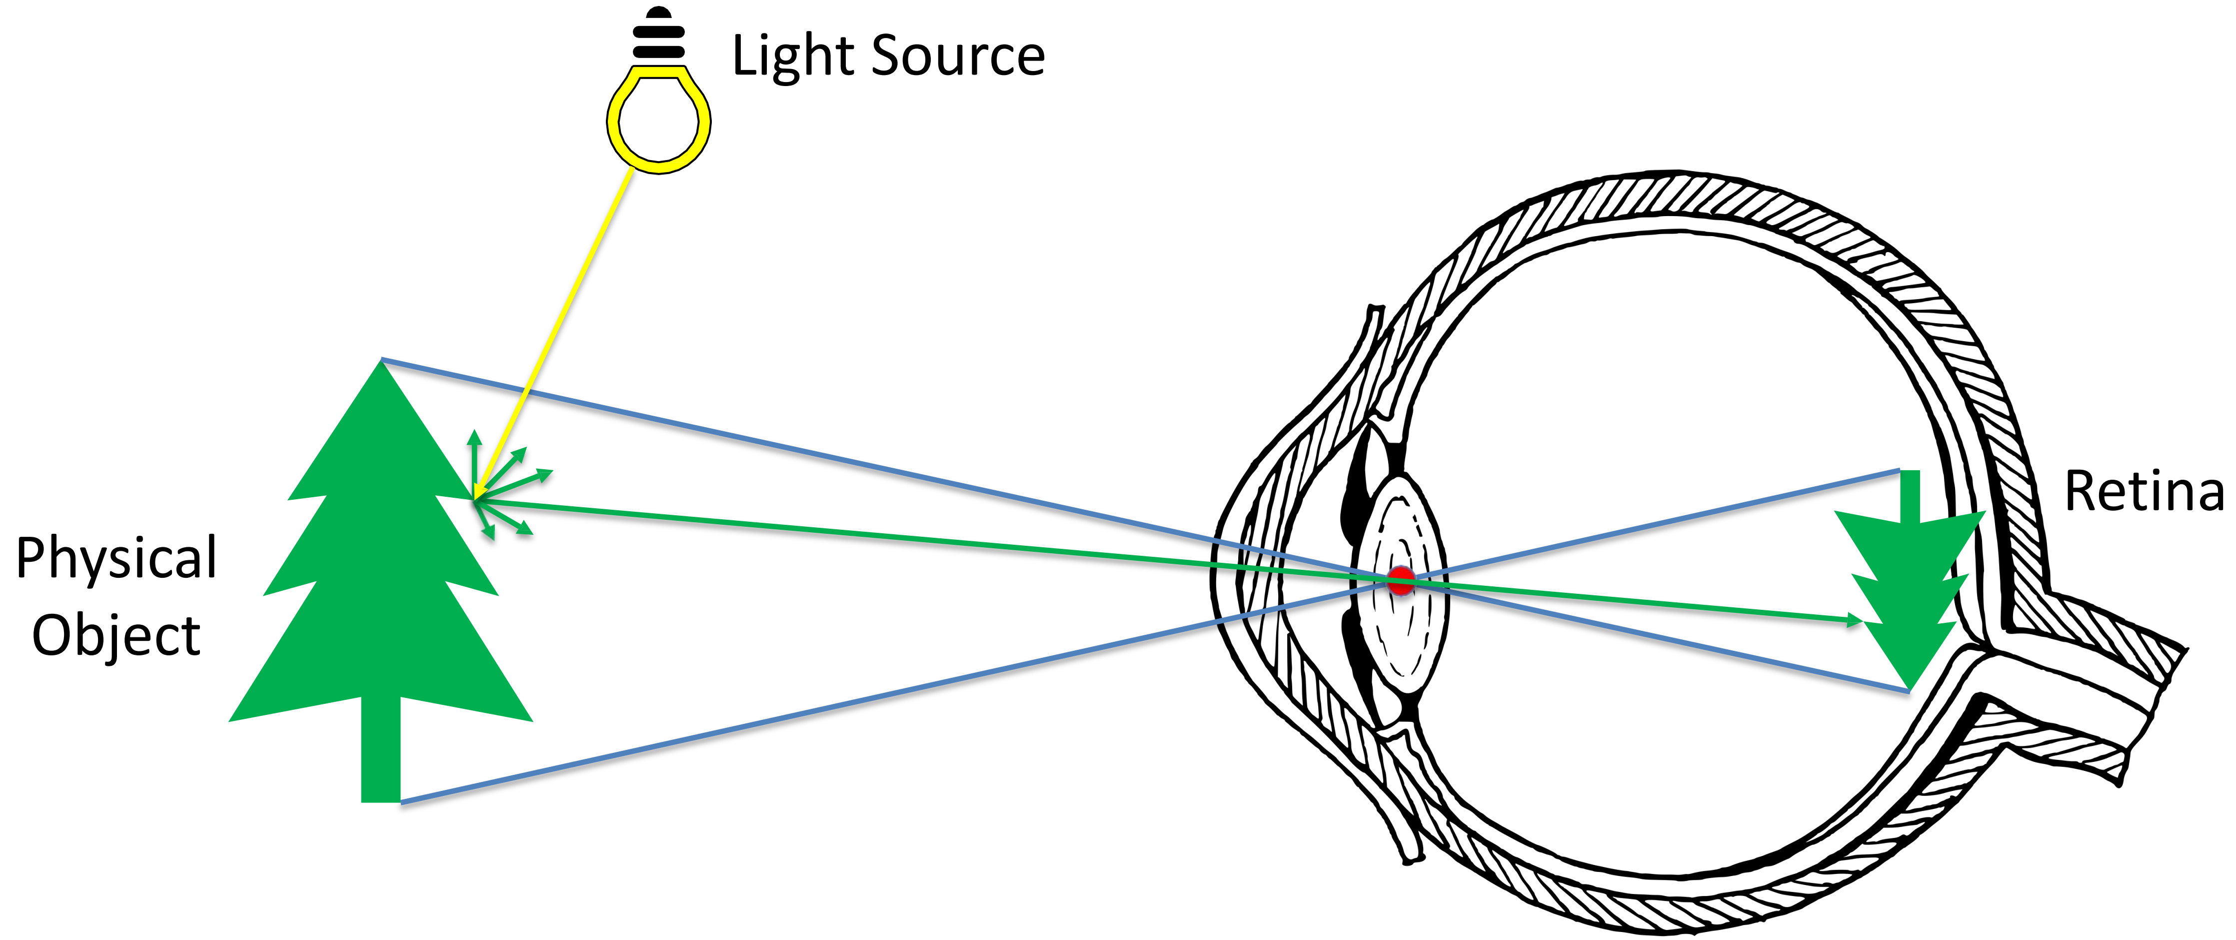

Image Formation

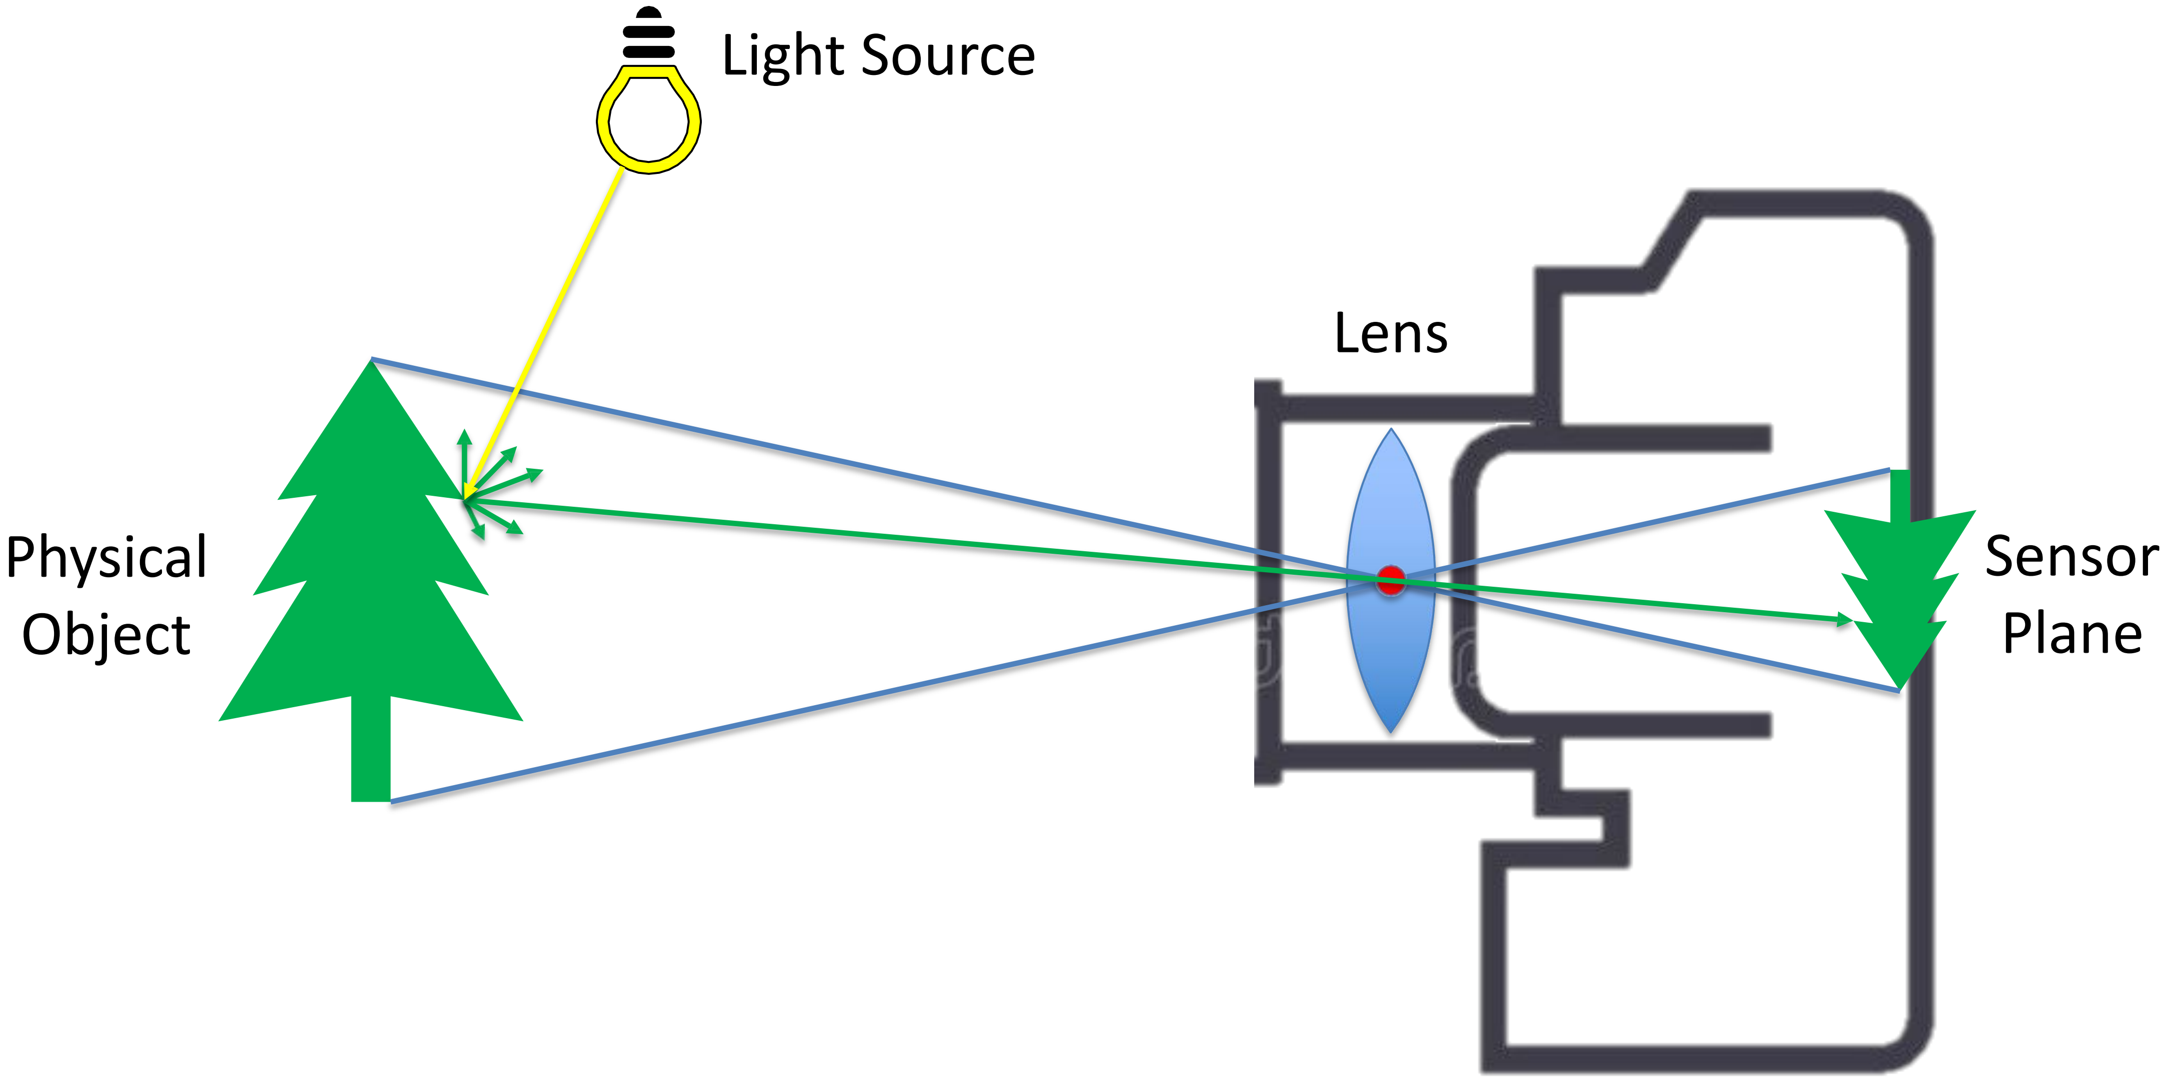

Today's agenda

- Image formation
- **Physics of Color**
- Color matching
- Color spaces
- Image sampling and quantization

Image Formation

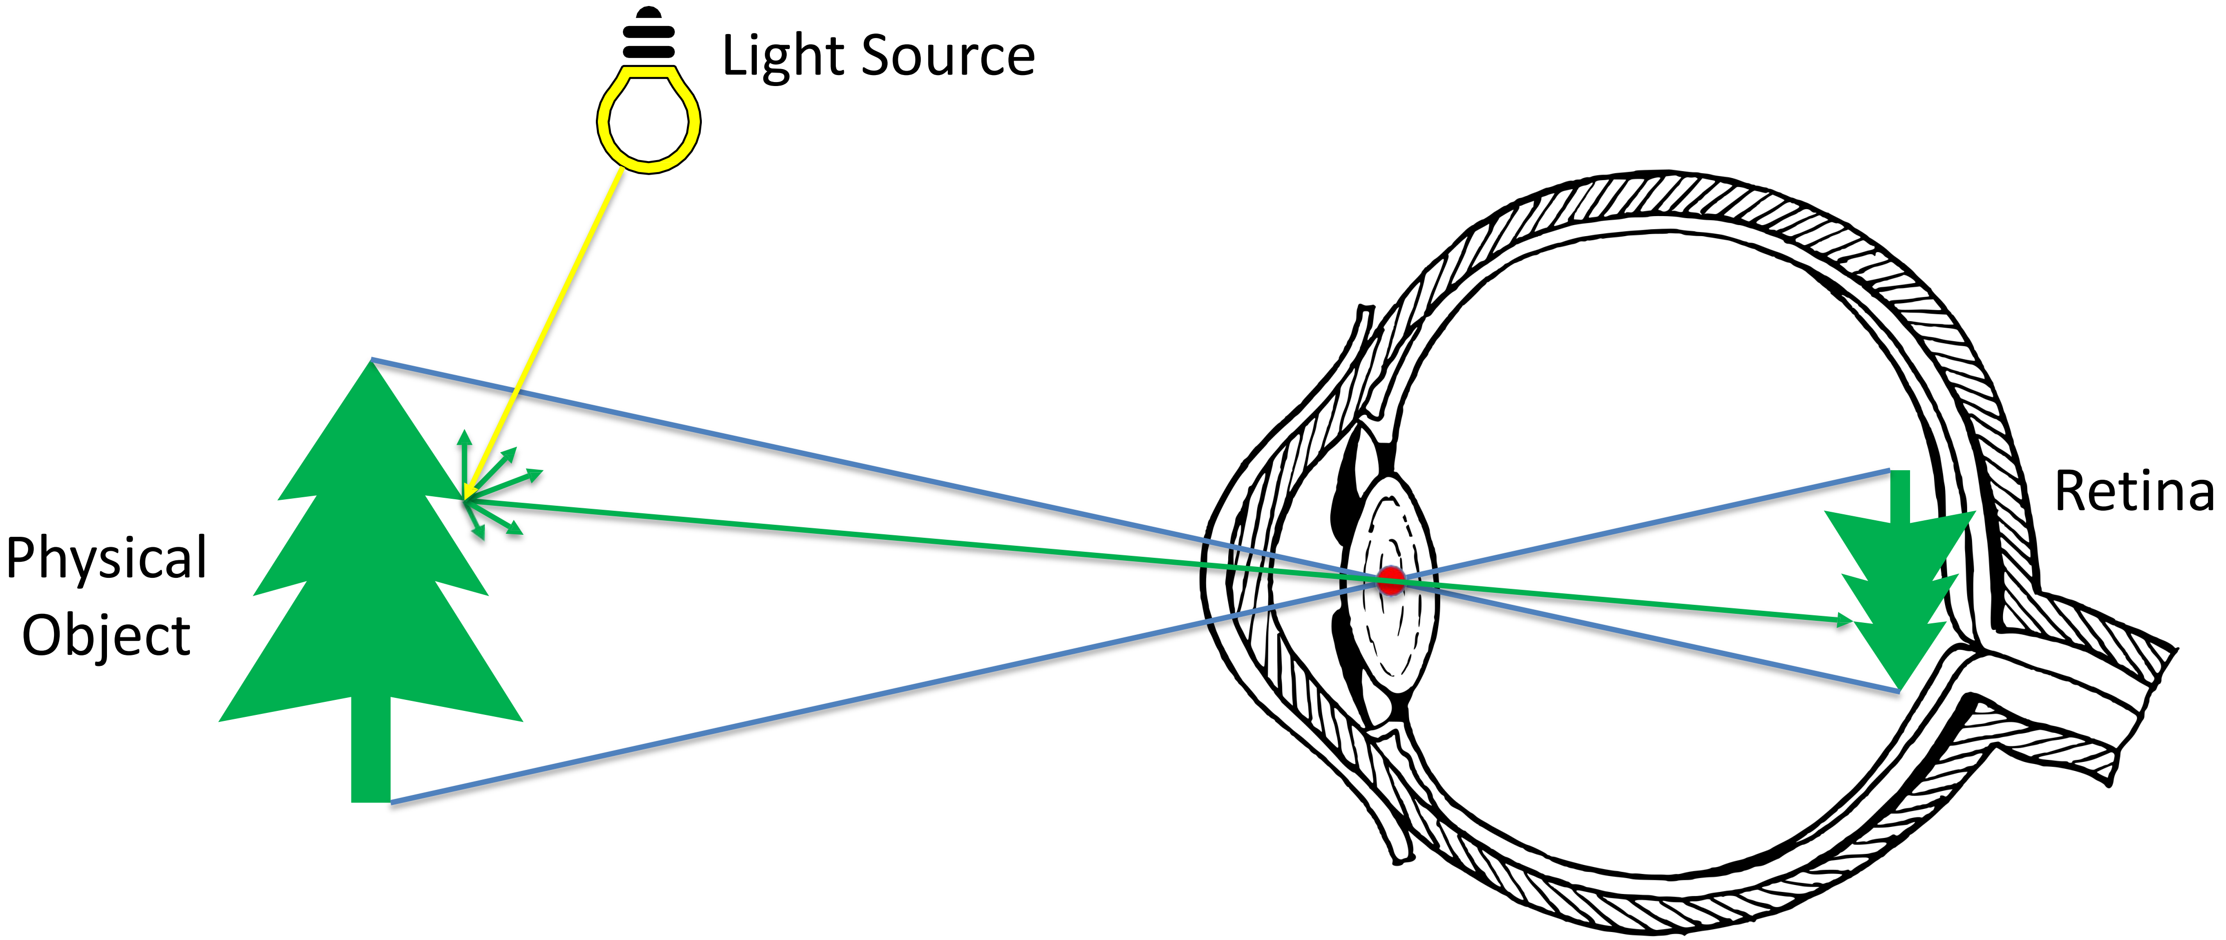

Image Formation

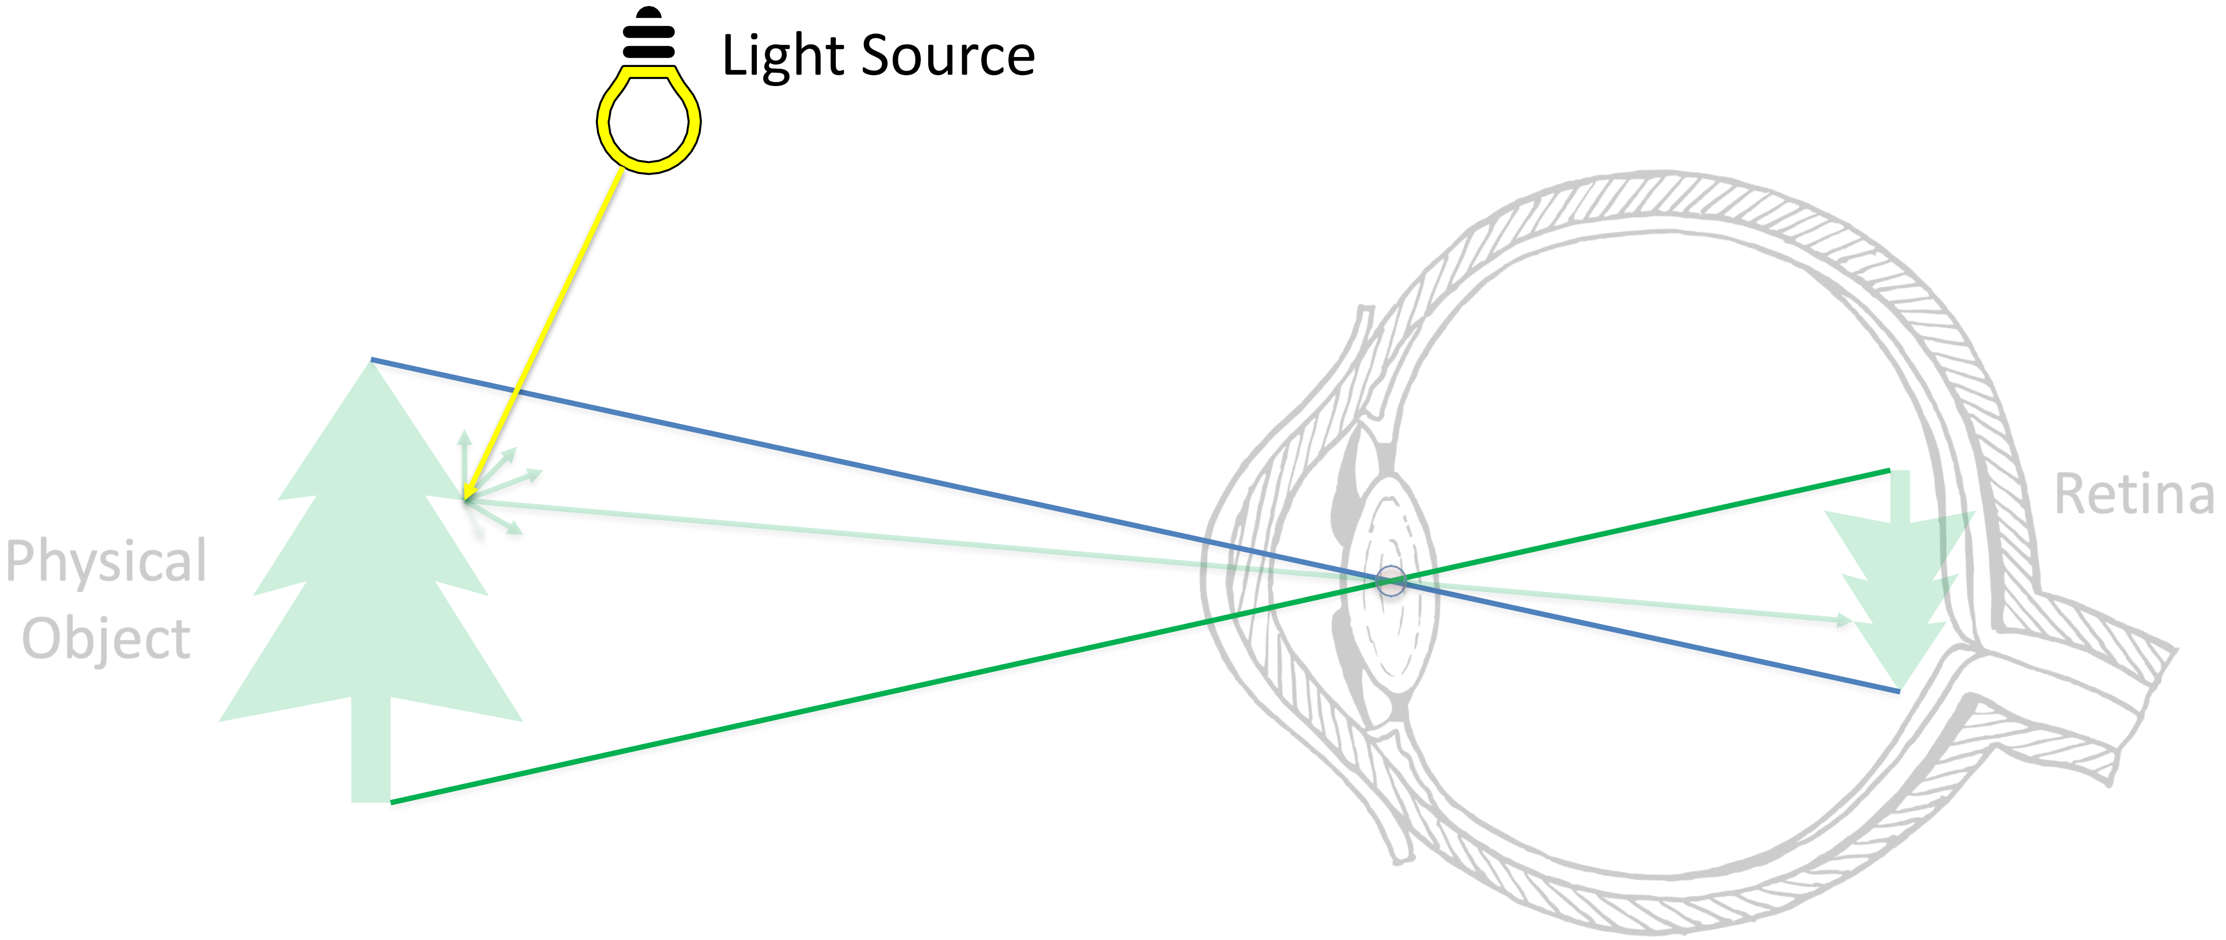

Color and light

White light: composed of almost equal energy in all wavelengths of the visible spectrum

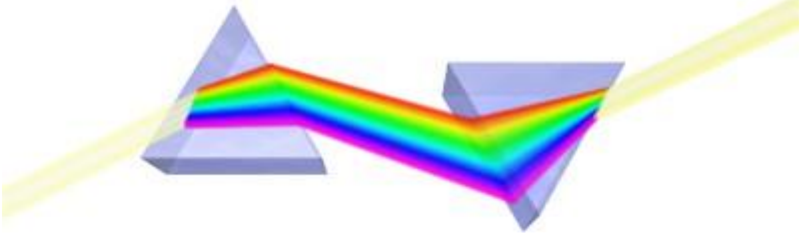

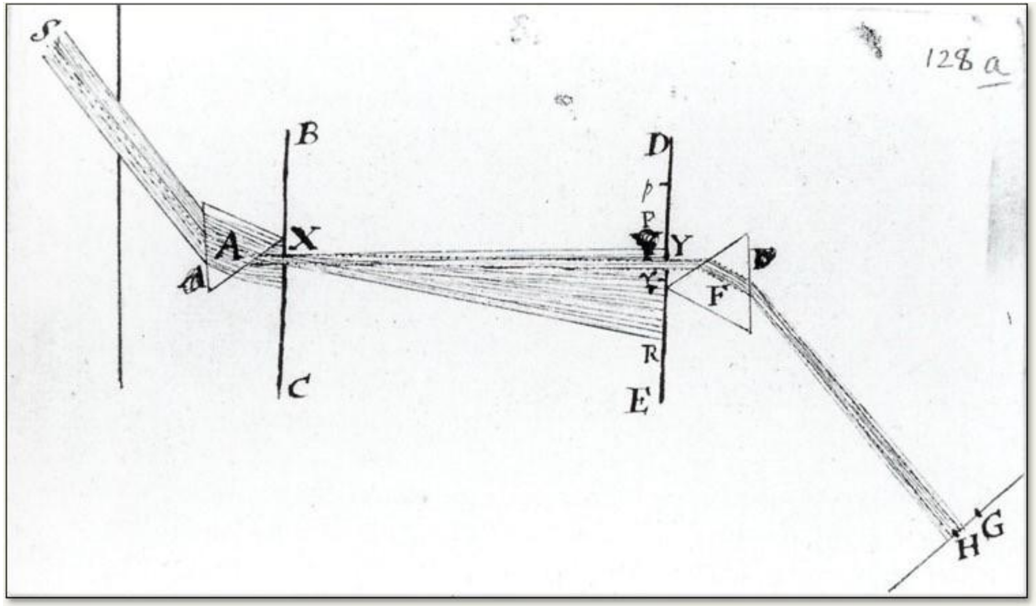

Newton 1665

Image from http://micro.magnet.fsu.edu/

Electromagnetic Spectrum & Colors

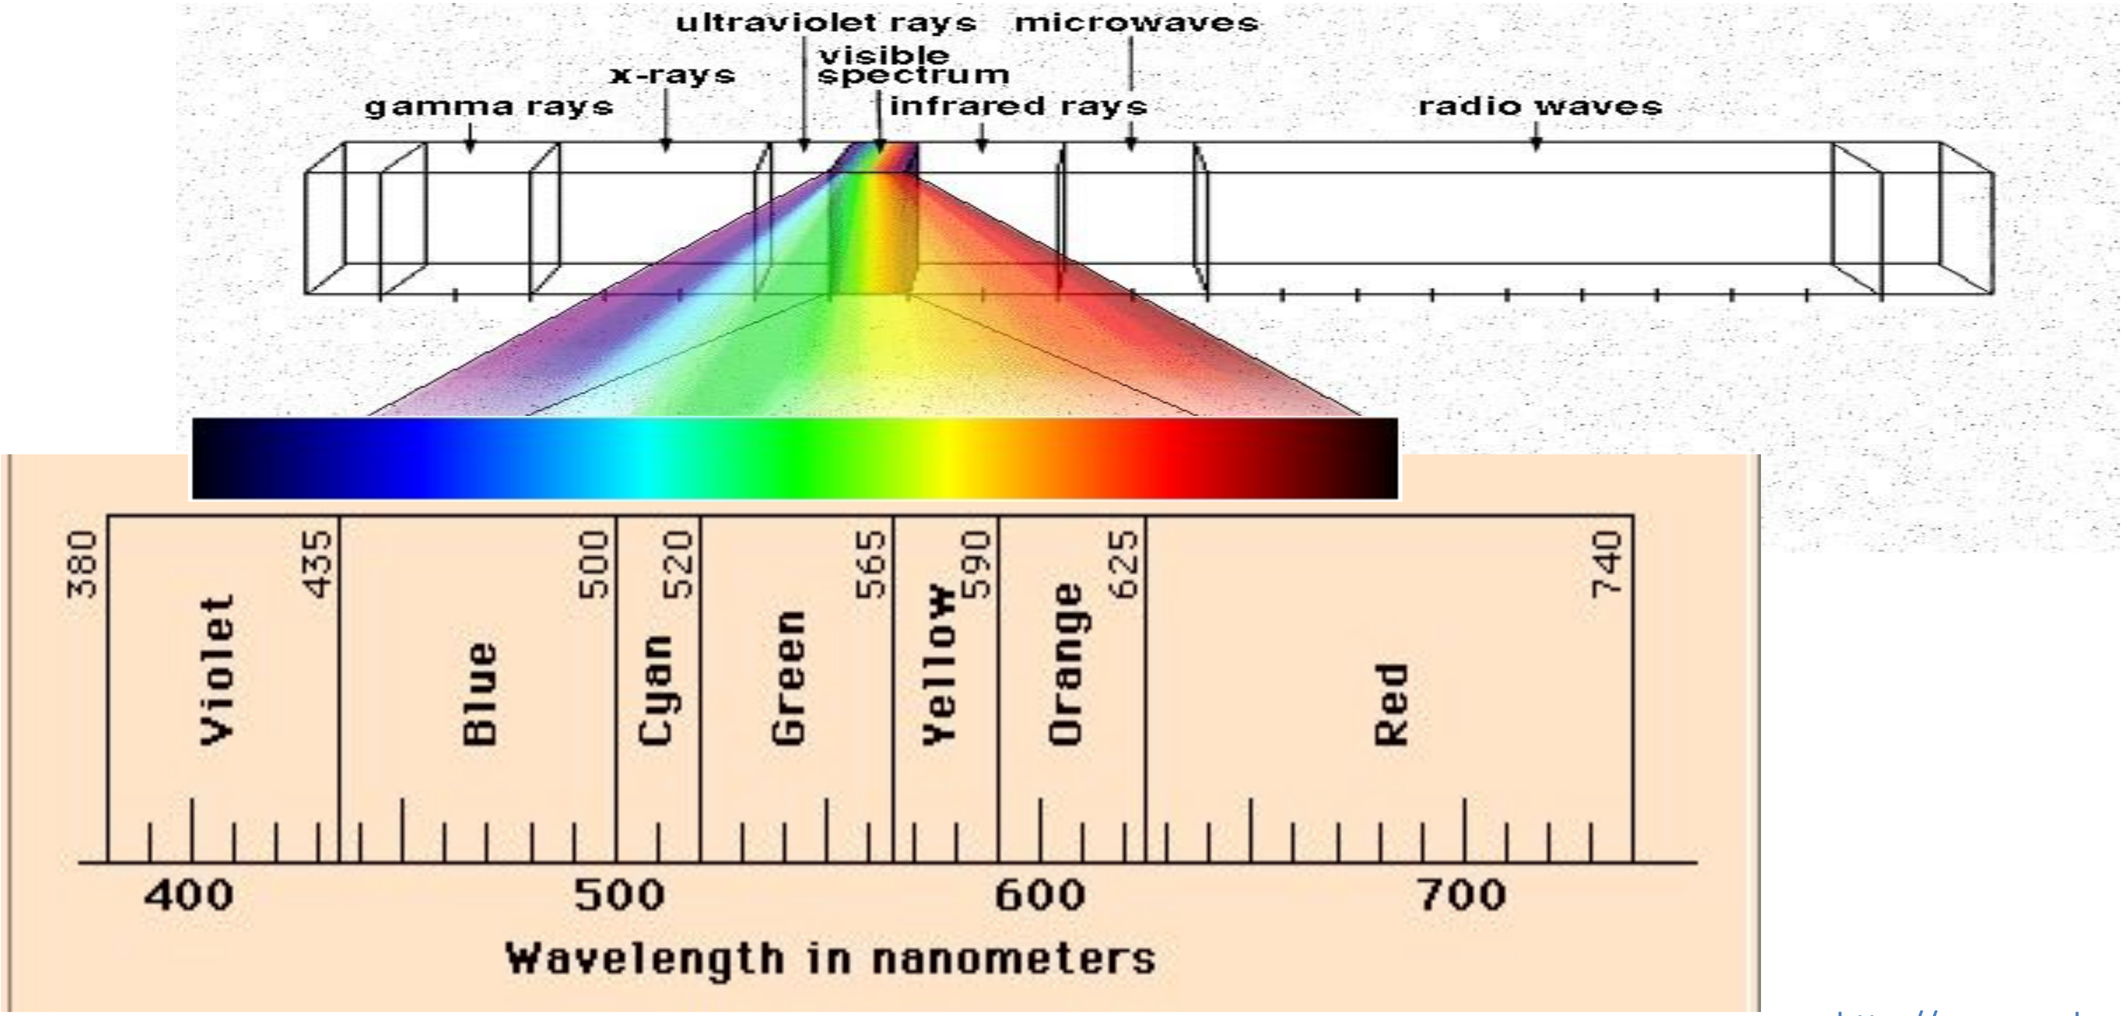

http://www.yorku.ca/eye/photopik.htm

http://hyperphysics.phy-astr.gsu.edu/hbase/vision/specol.html#c2

The Physics of Light

Any patch of light can be completely described physically by its **spectrum**.

A spectrum is the number of photons (per time unit) at each wavelength 400 - 700 nm.

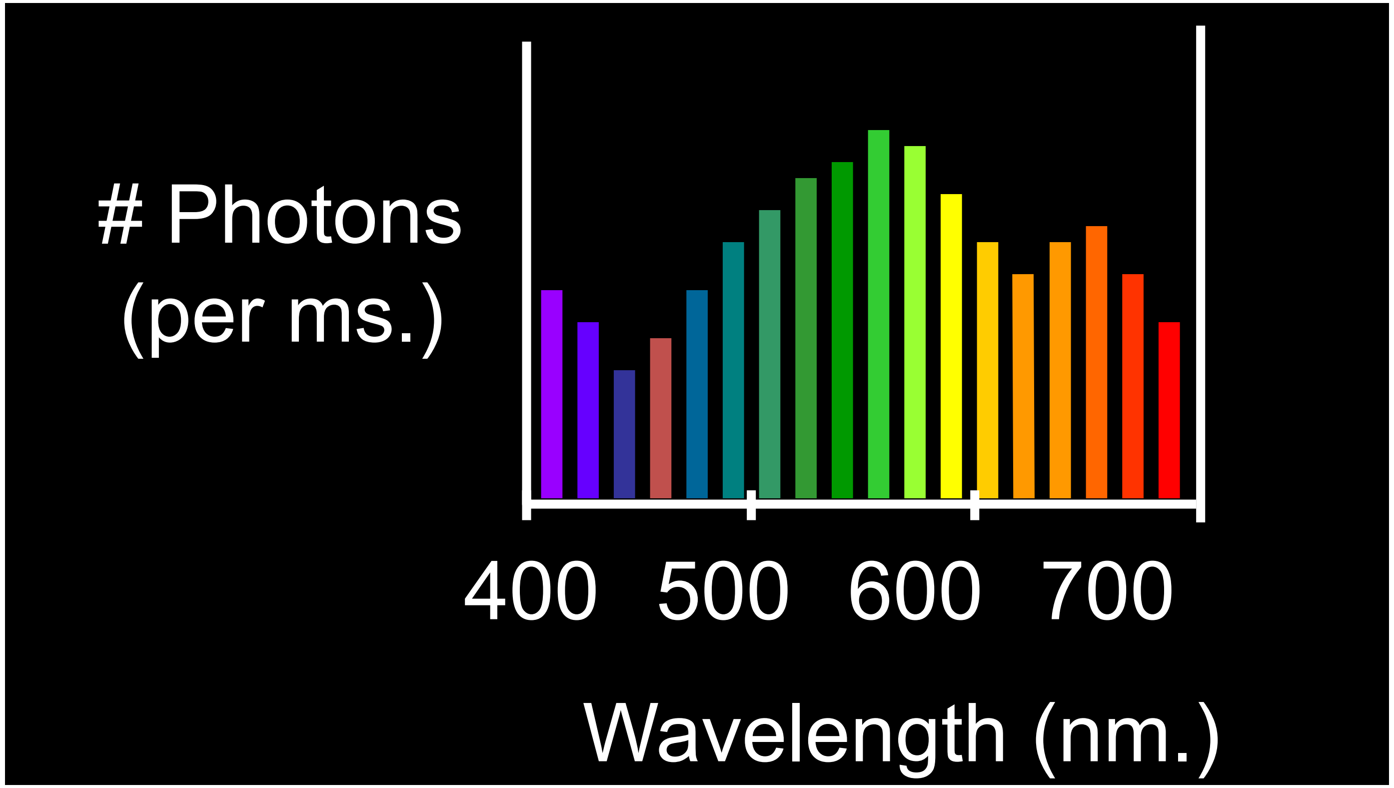

The Physics of Light

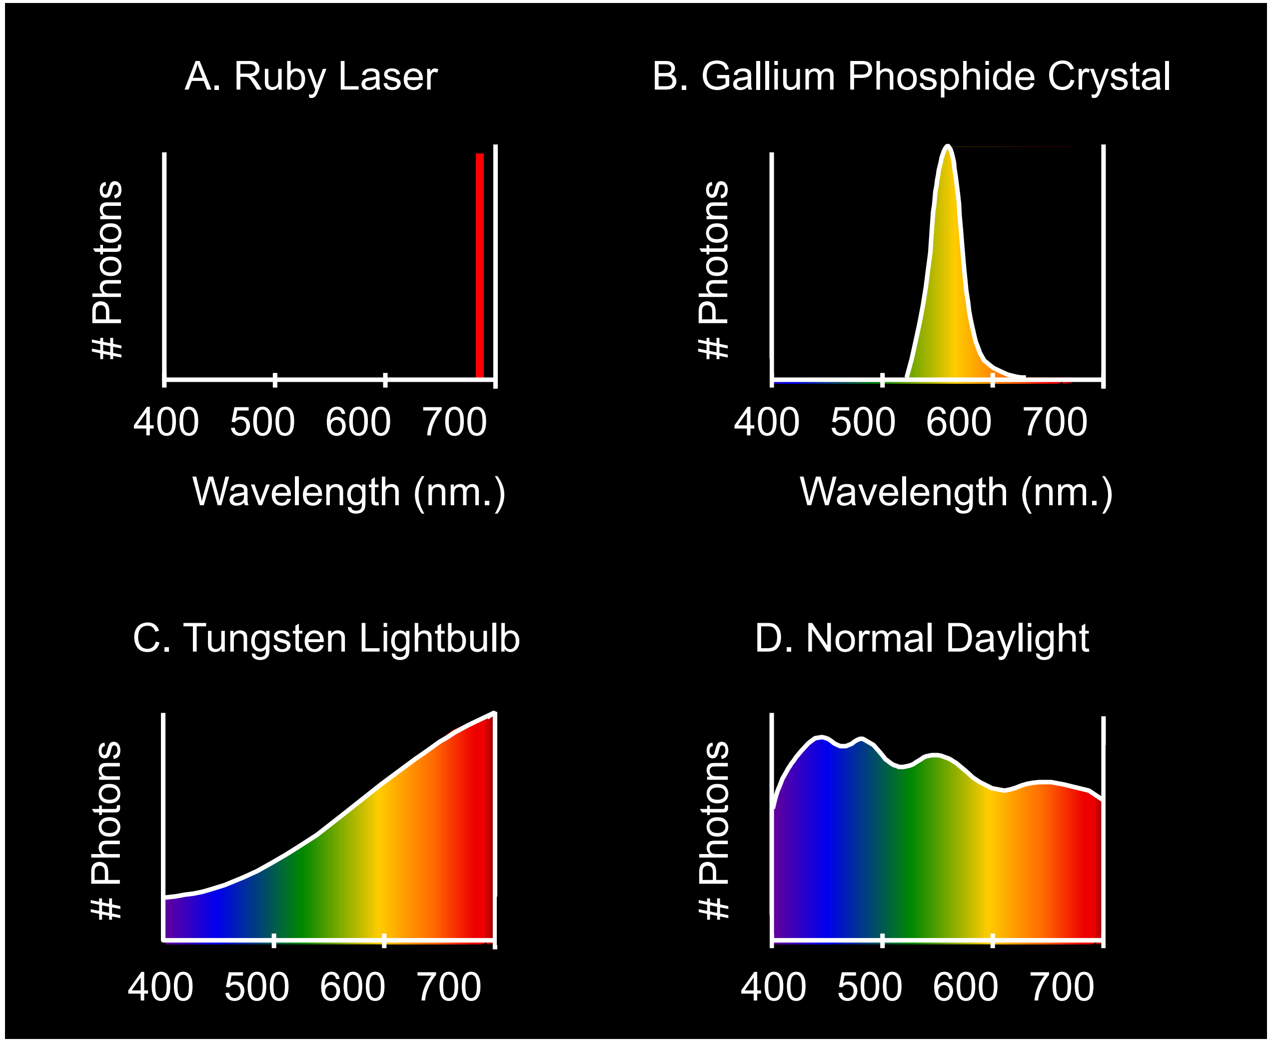

Some examples of the spectra of light sources

Color names for cartoon spectra

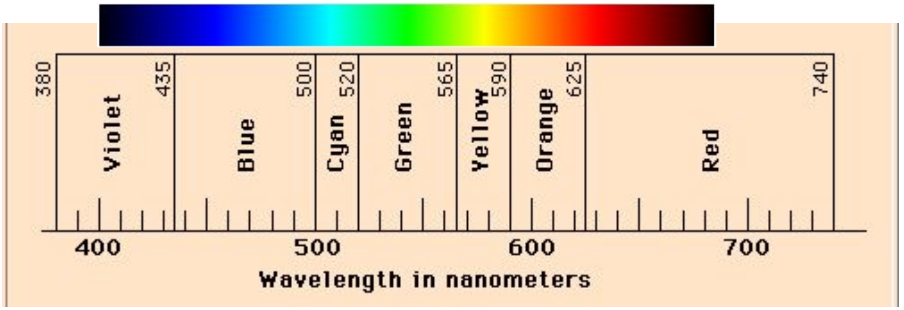

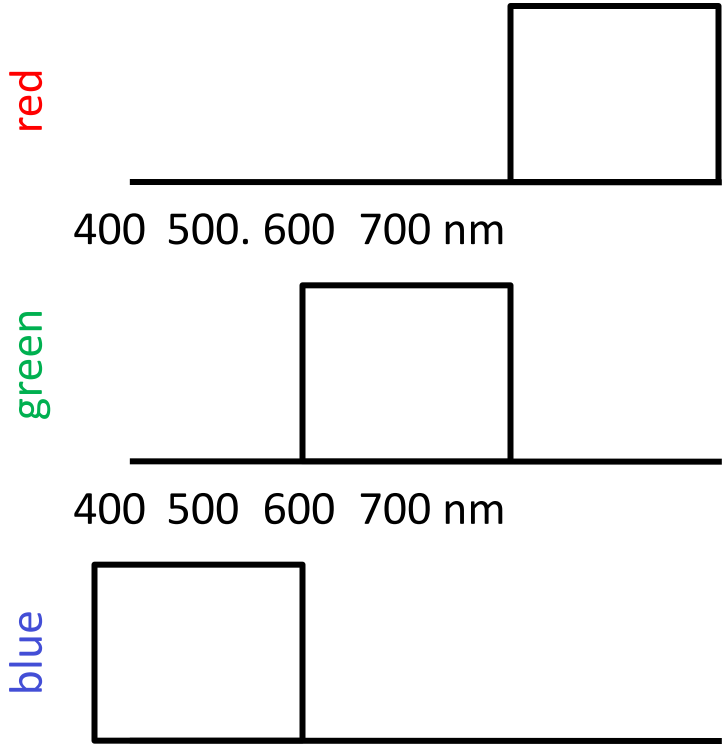

Color names for cartoon spectra

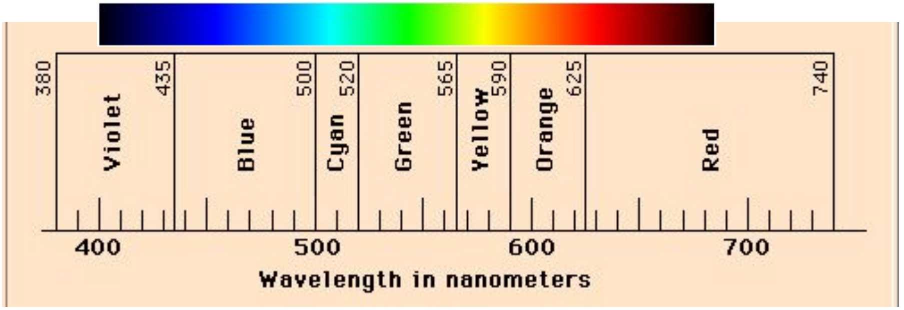

red

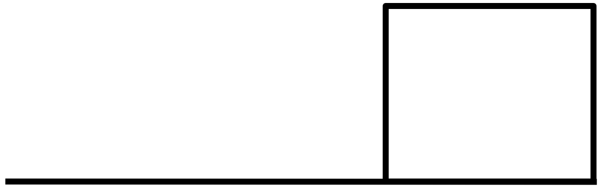

400 500. 600 700 nm

green

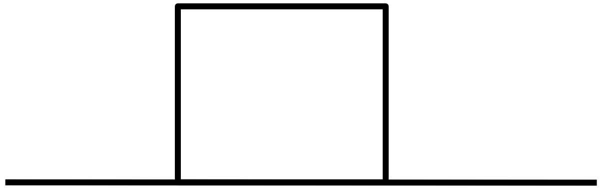

400 500 600 700 nm

blue

minus red

400 500 600 700 nm

Q. What color do we get when we subtract red?

Color names for cartoon spectra

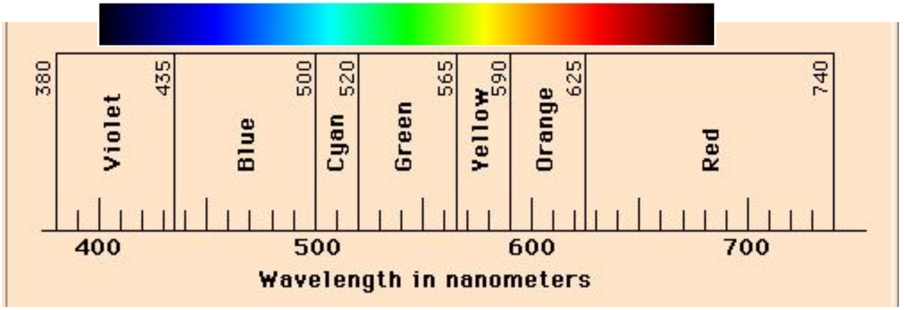

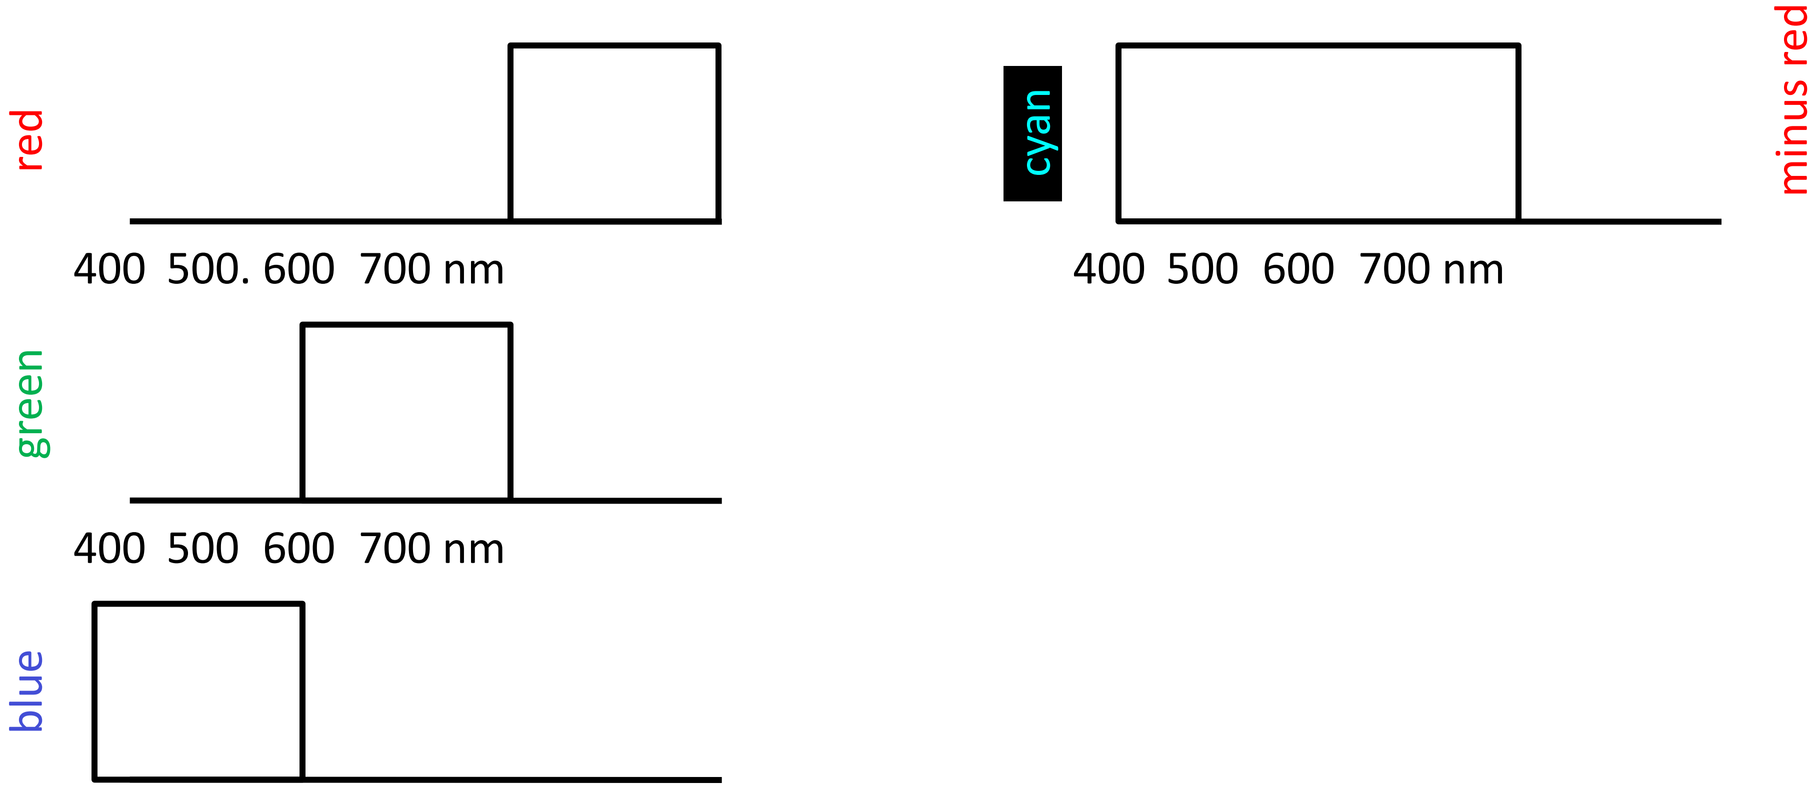

Color names for cartoon spectra

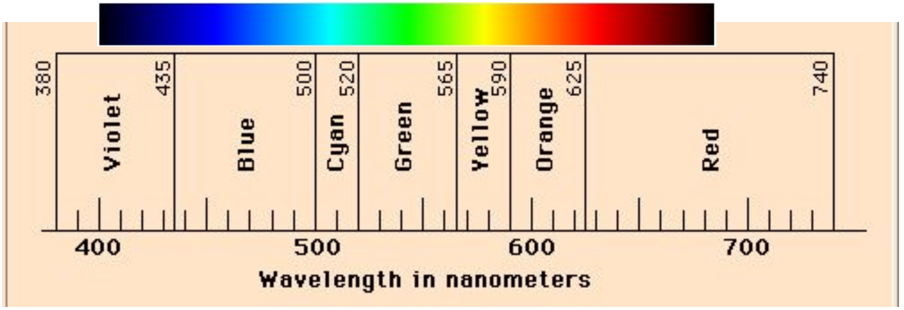

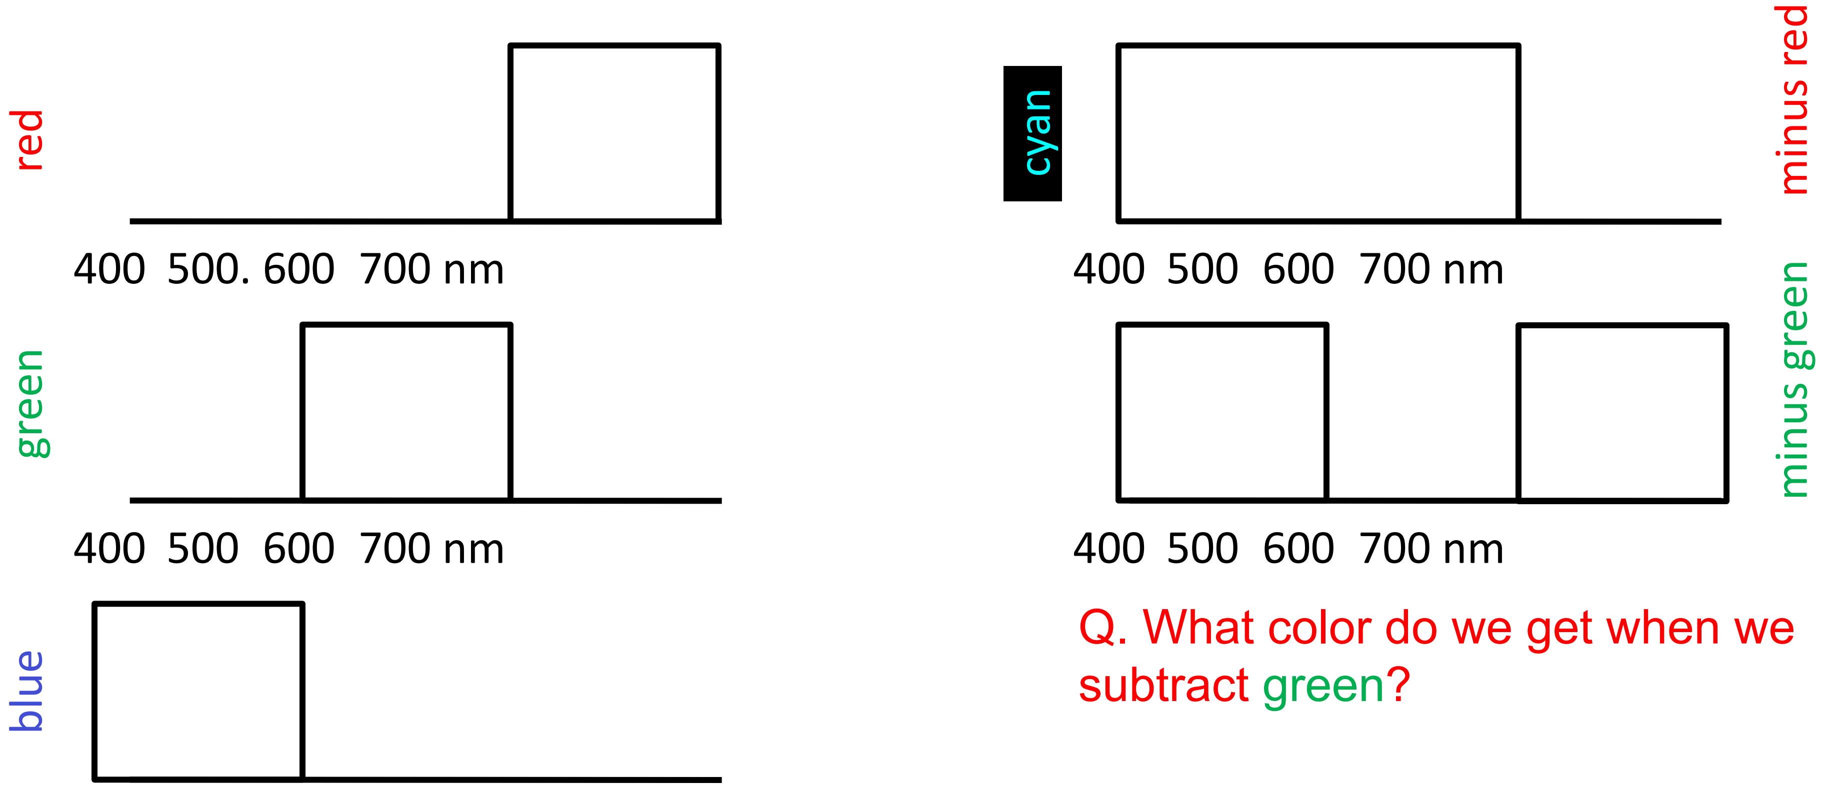

Q. What color do we get when we subtract green?

Color names for cartoon spectra

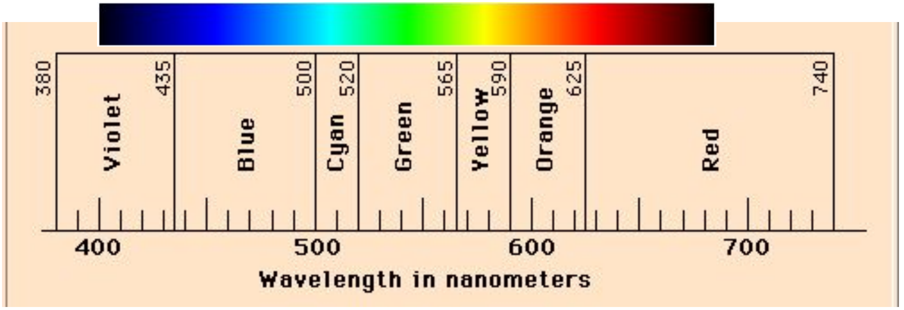

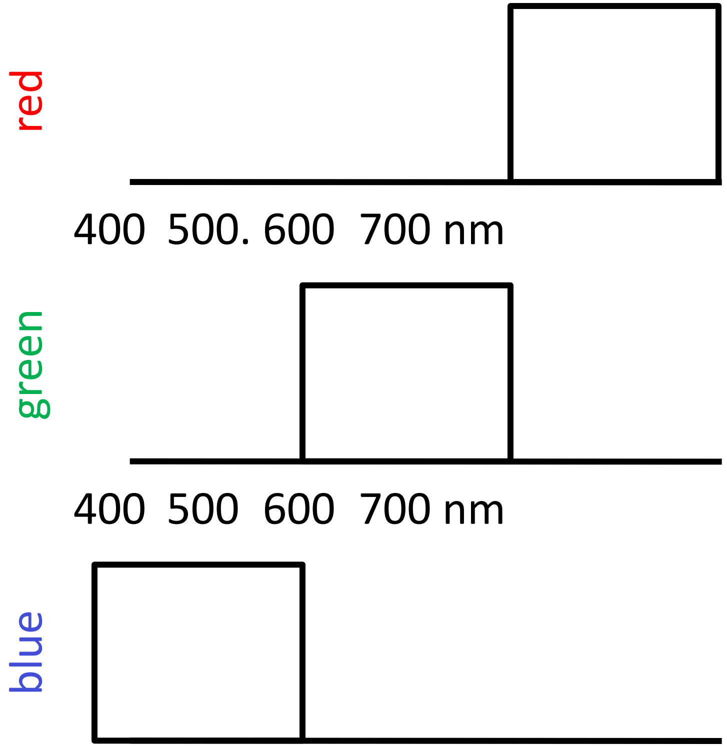

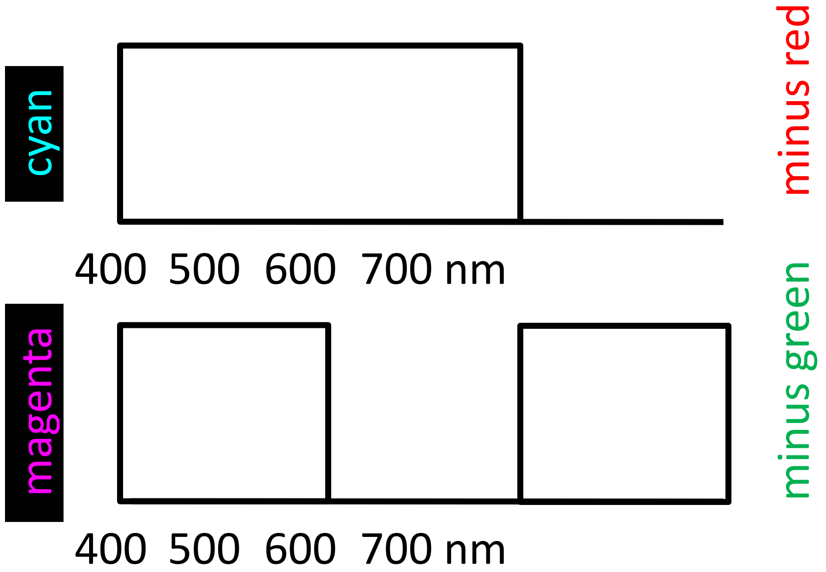

Q. What color do we get when we subtract green?

Color names for cartoon spectra

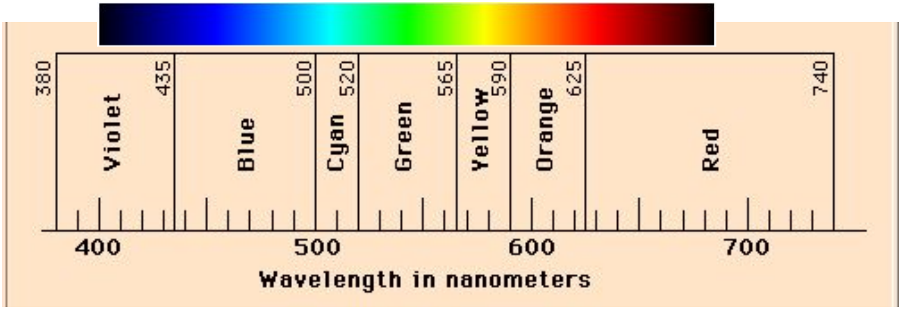

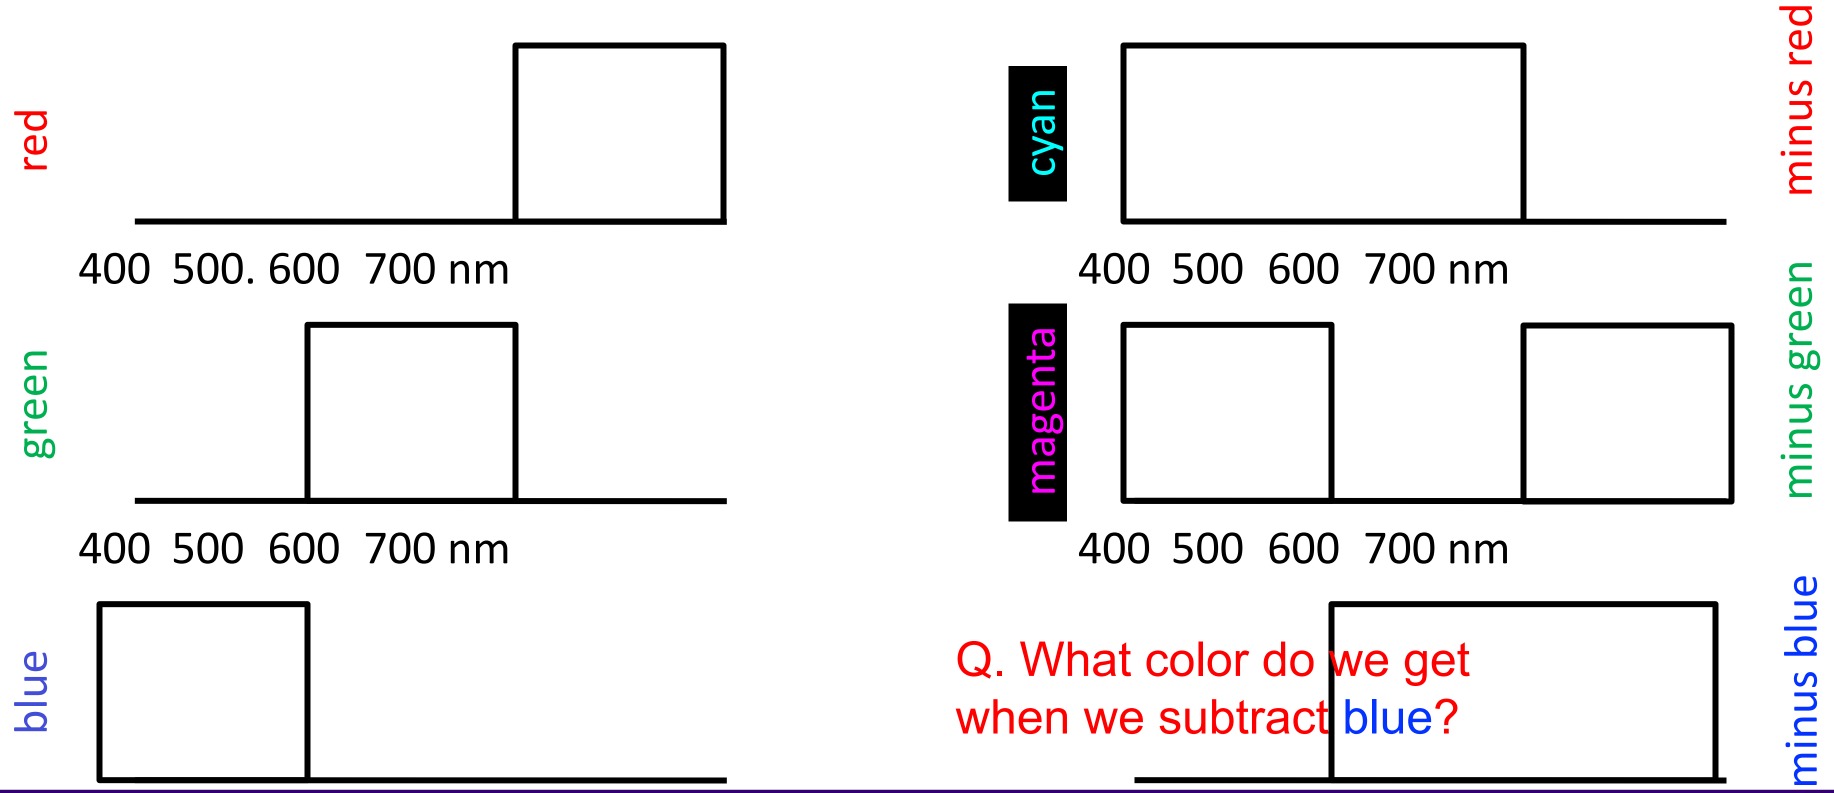

Color names for cartoon spectra

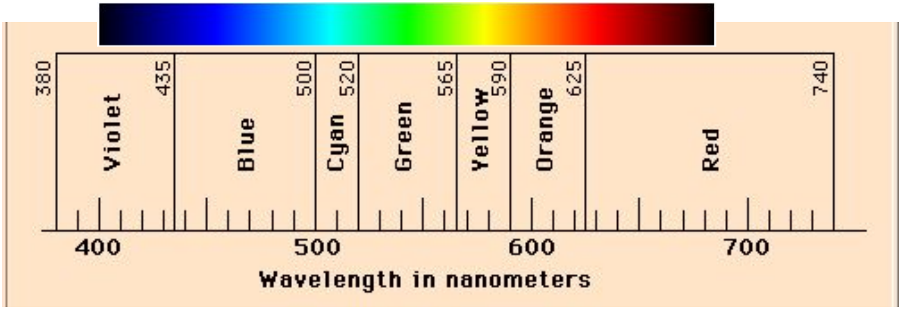

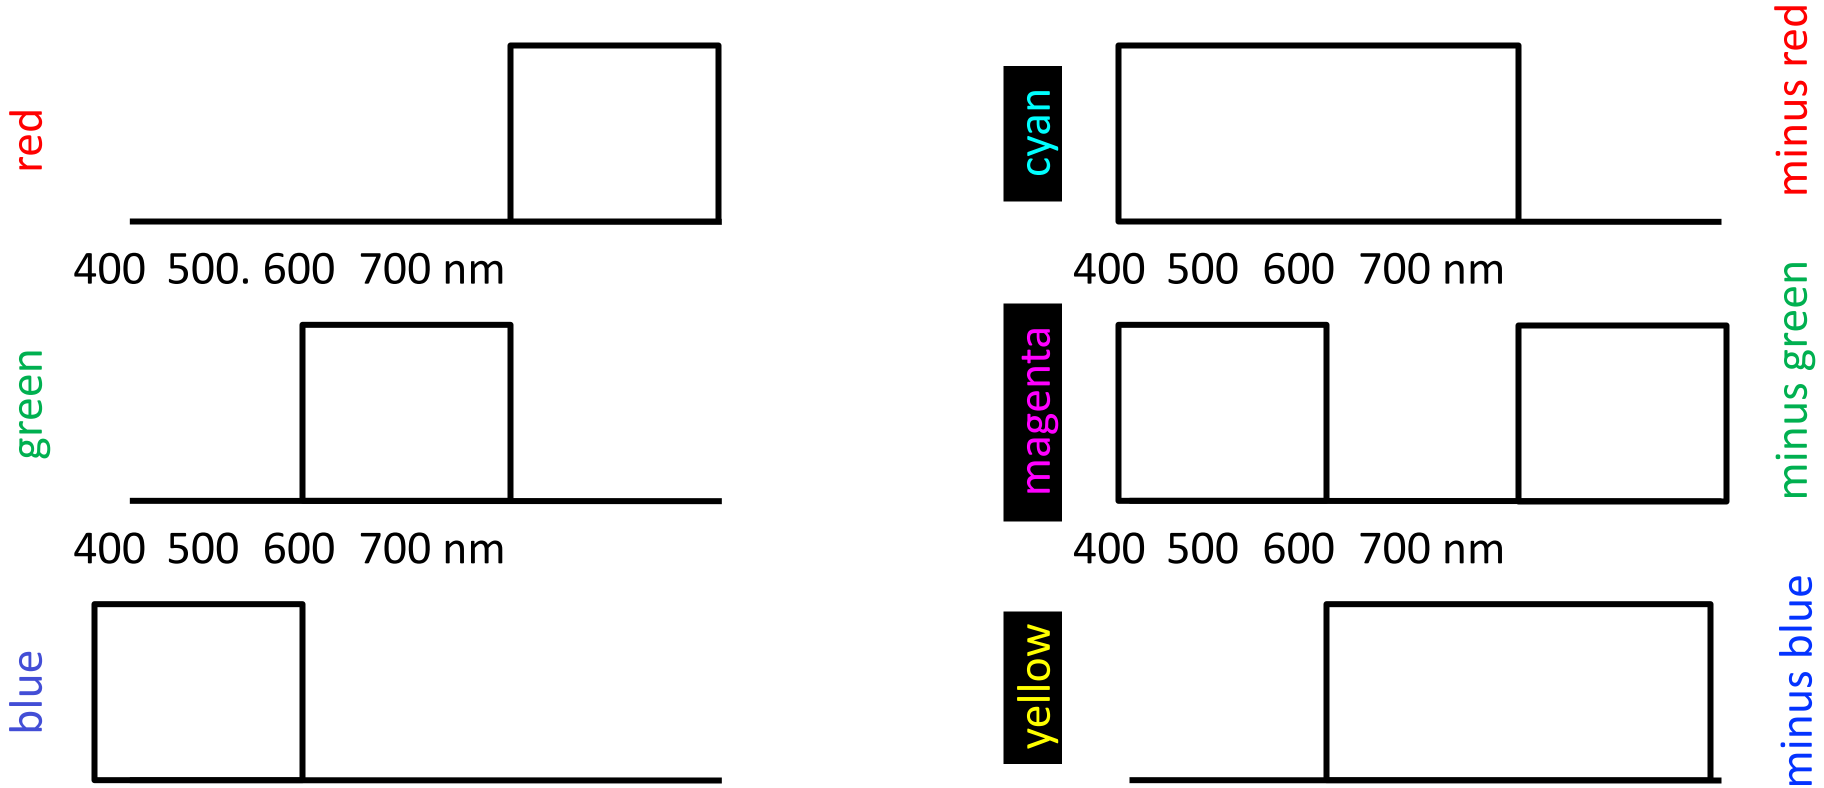

Additive color mixing

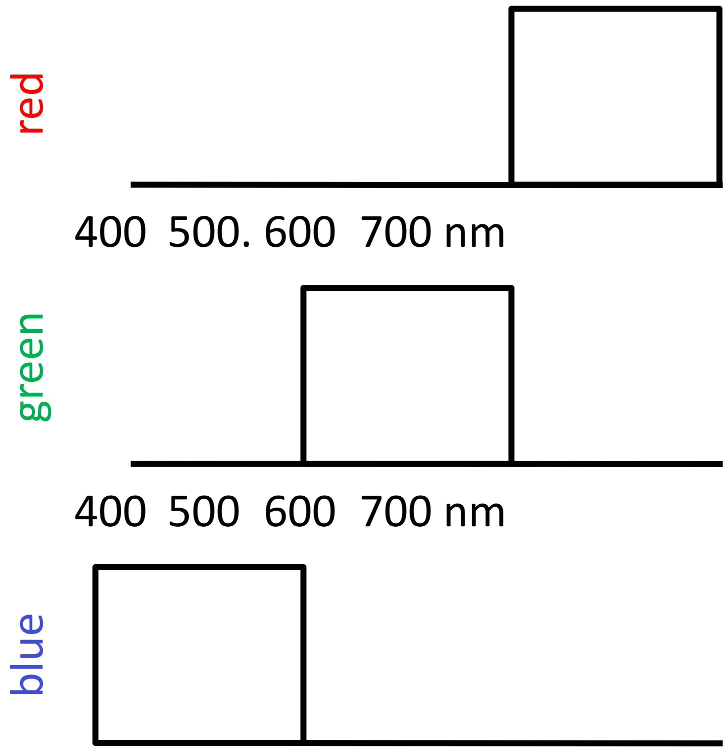

Colors combine by *adding* color spectra

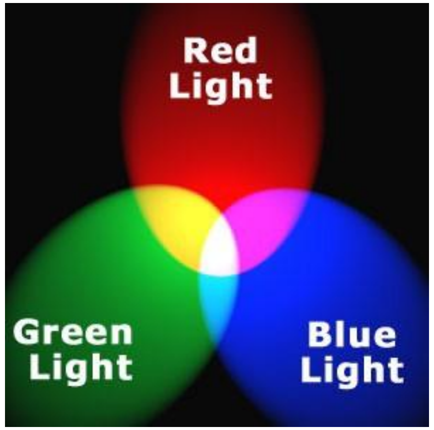

Light *adds* to white.

Examples of additive color systems

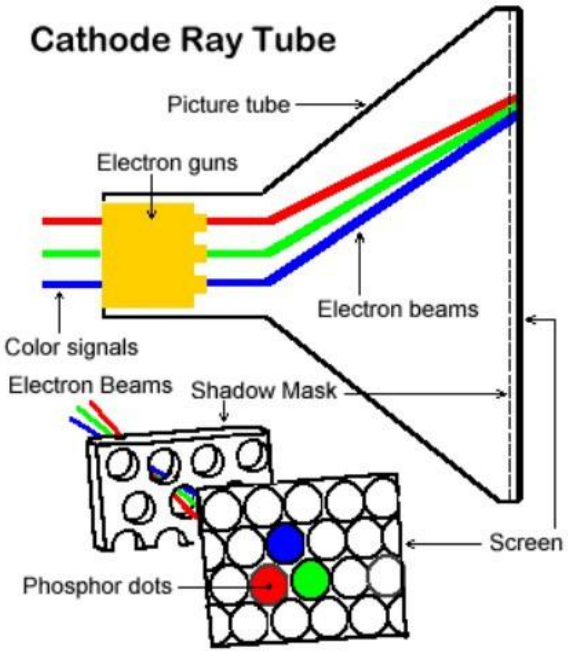

CRT phosphors

multiple projectors

http://www.jegsworks.com

http://www.crtprojectors.co.uk/

Subtractive color mixing

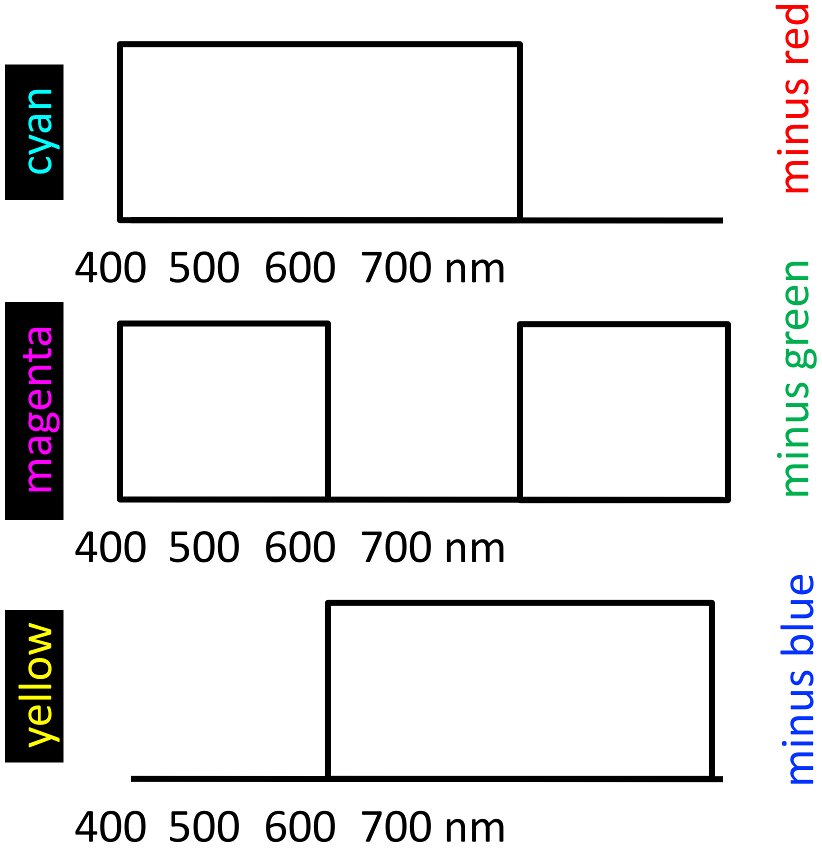

Colors combine by *multiplying* color spectra.

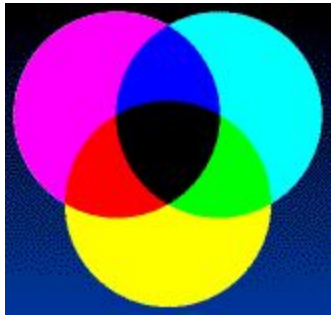

Pigments *remove* color from incident (white) light.

Examples of subtractive color systems

- Printing on paper
- Crayons
- Photographic film

Image Formation

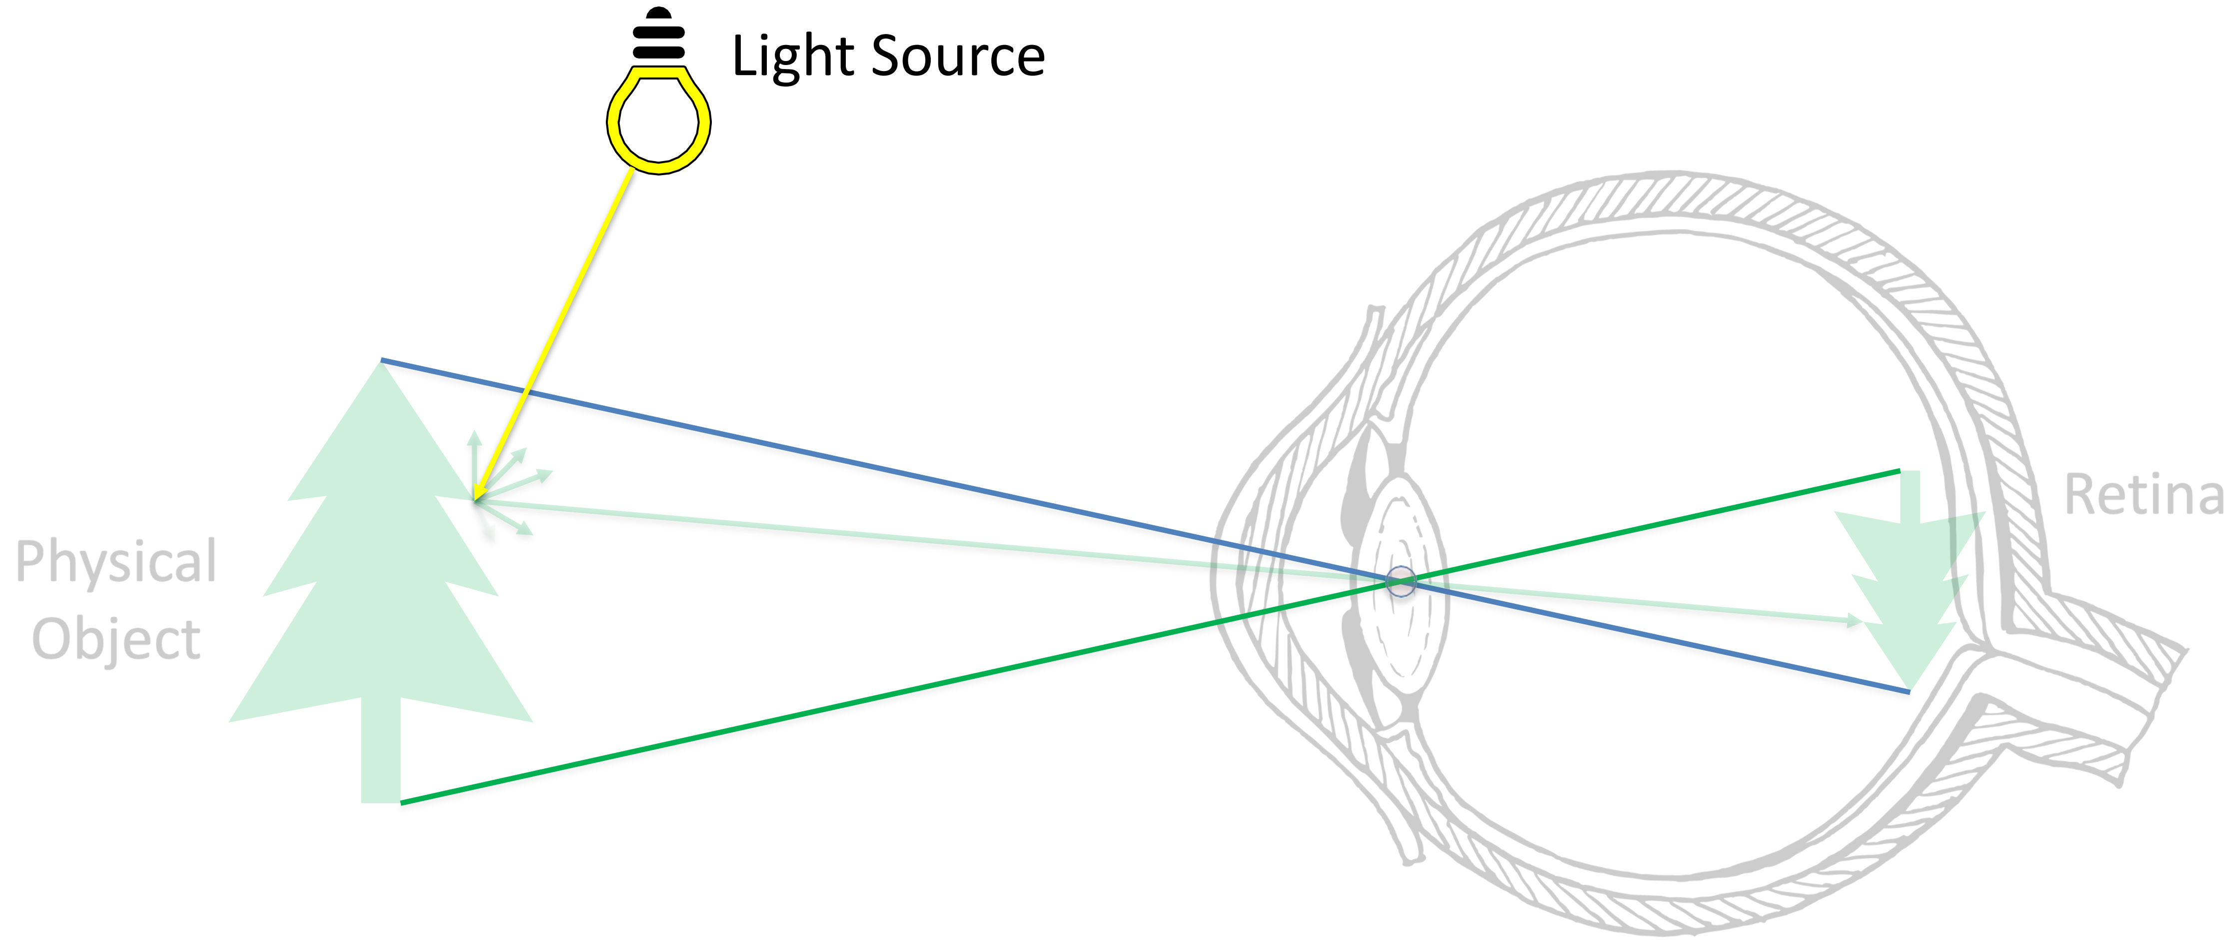

Image Formation

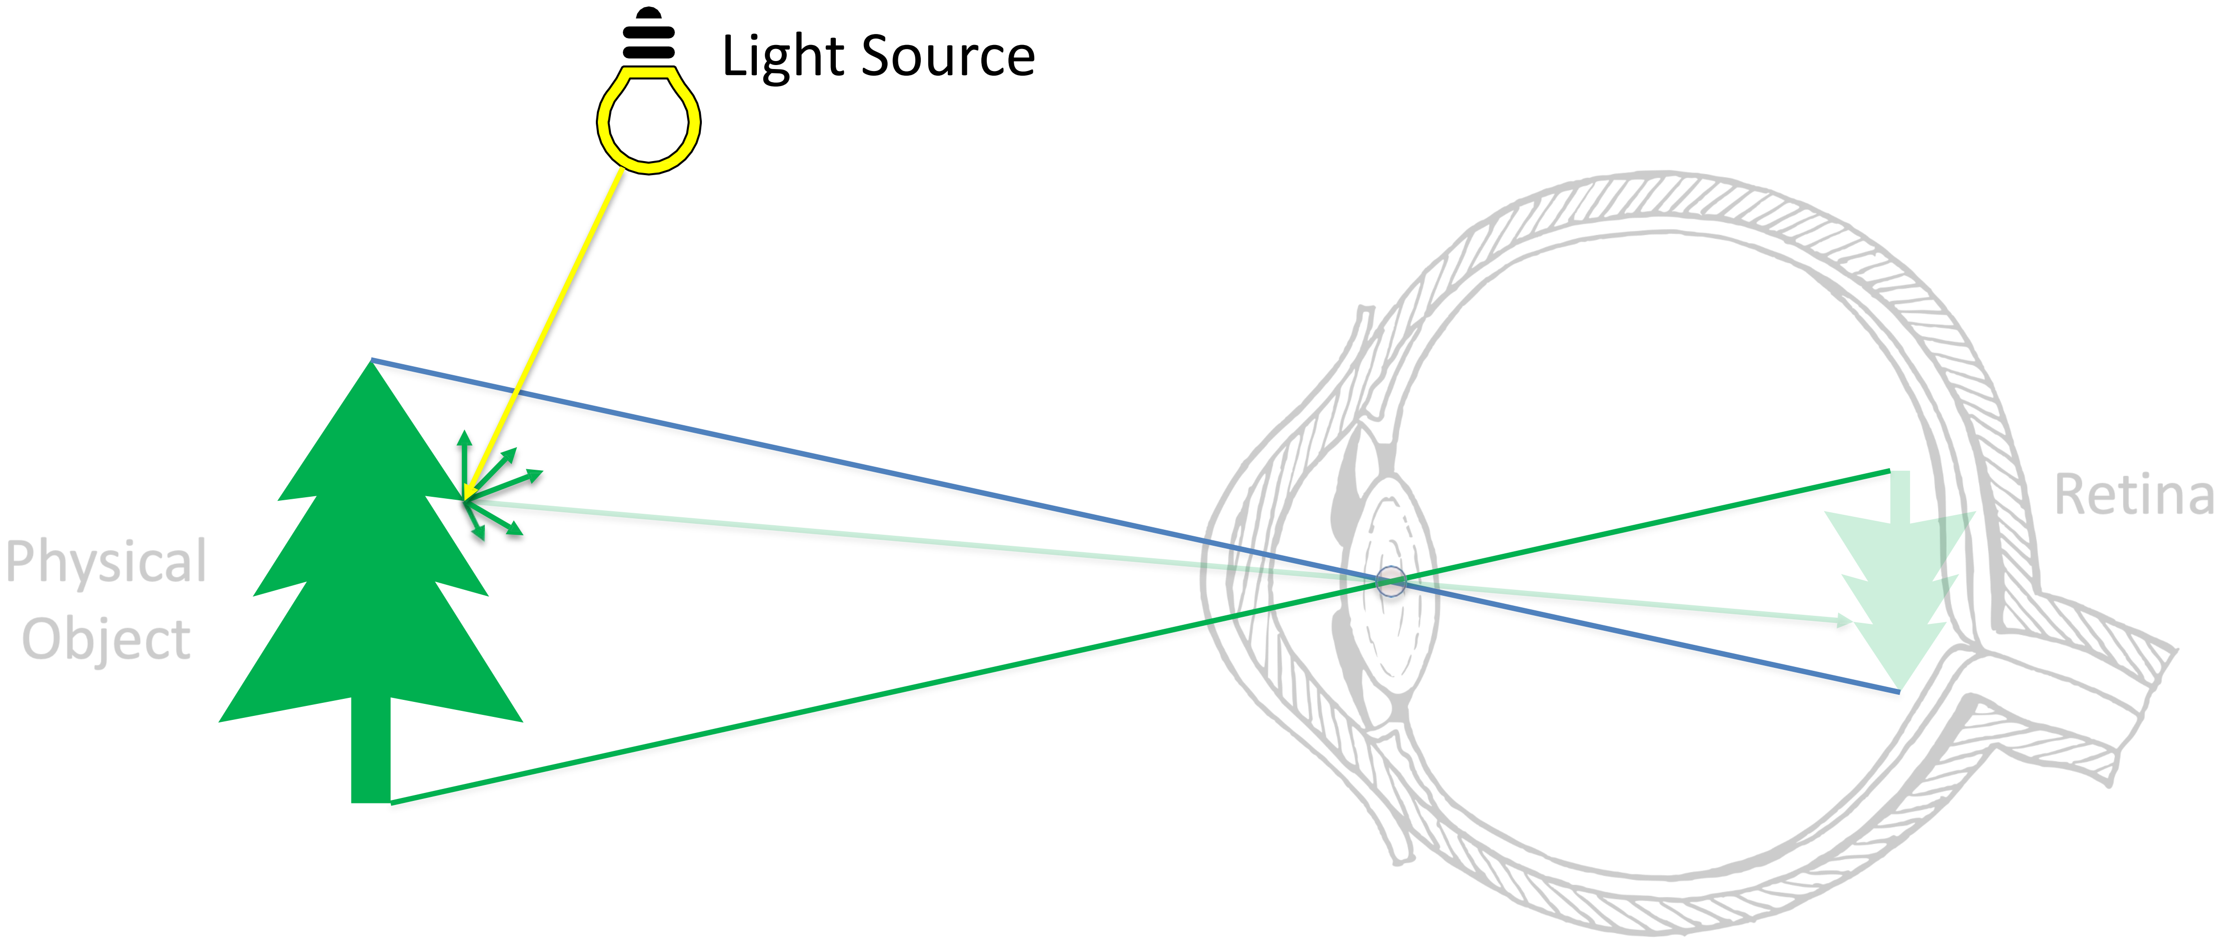

Photon

- Absorption
- Diffusion
- Reflection
- Transparency
- Refraction
- Fluorescence
- Subsurface scattering
- Phosphorescence
- Interreflection

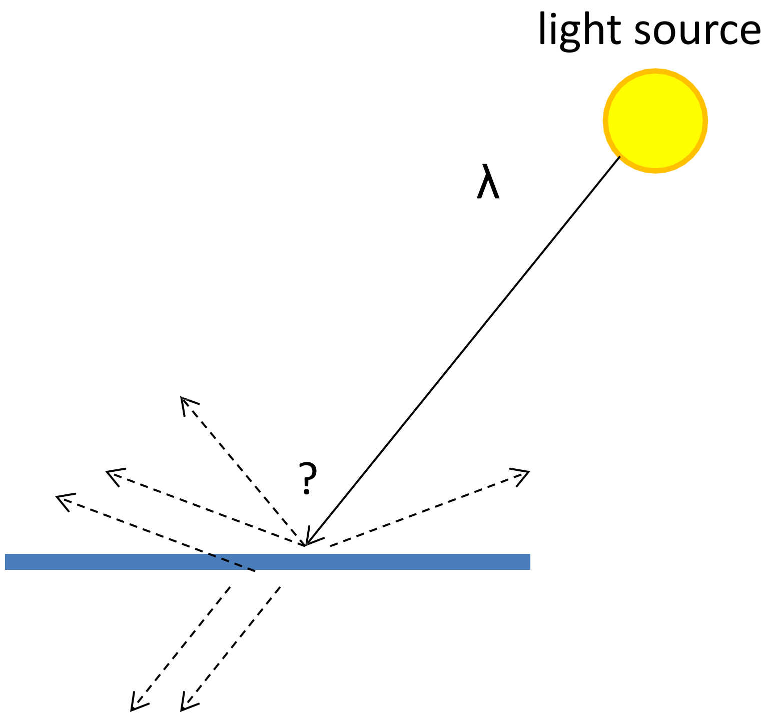

Photon

- **Absorption**
- Diffusion
- Reflection
- Transparency
- Refraction
- Fluorescence
- Subsurface scattering
- Phosphorescence
- Interreflection

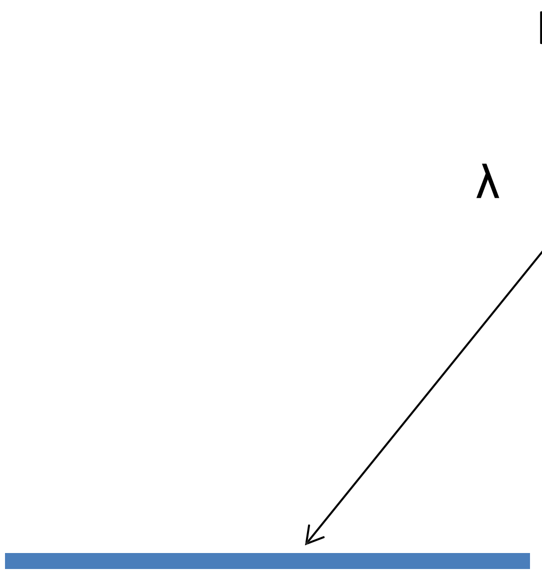

Photon

- Absorption
- **Diffuse Reflection**
- Reflection
- Transparency
- Refraction
- Fluorescence
- Subsurface scattering
- Phosphorescence
- Interreflection

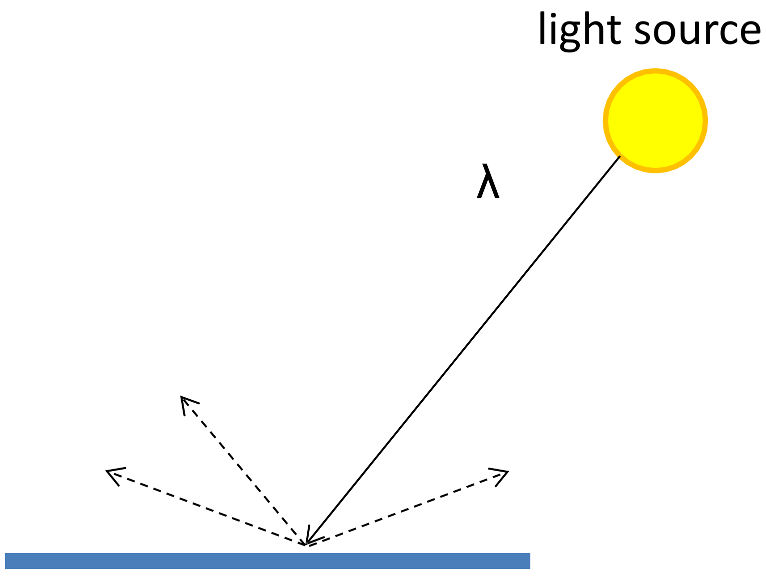

Photon

- Absorption
- Diffusion
- **Specular Reflection**
- Transparency
- Refraction
- Fluorescence
- Subsurface scattering
- Phosphorescence
- Interreflection

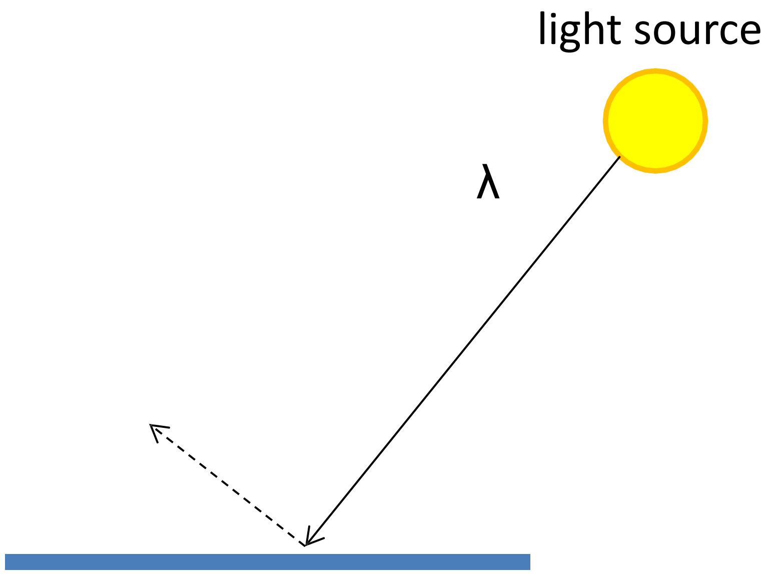

Photon

- Absorption
- Diffusion
- Reflection
- **Transparency**
- Refraction
- Fluorescence
- Subsurface scattering
- Phosphorescence
- Interreflection

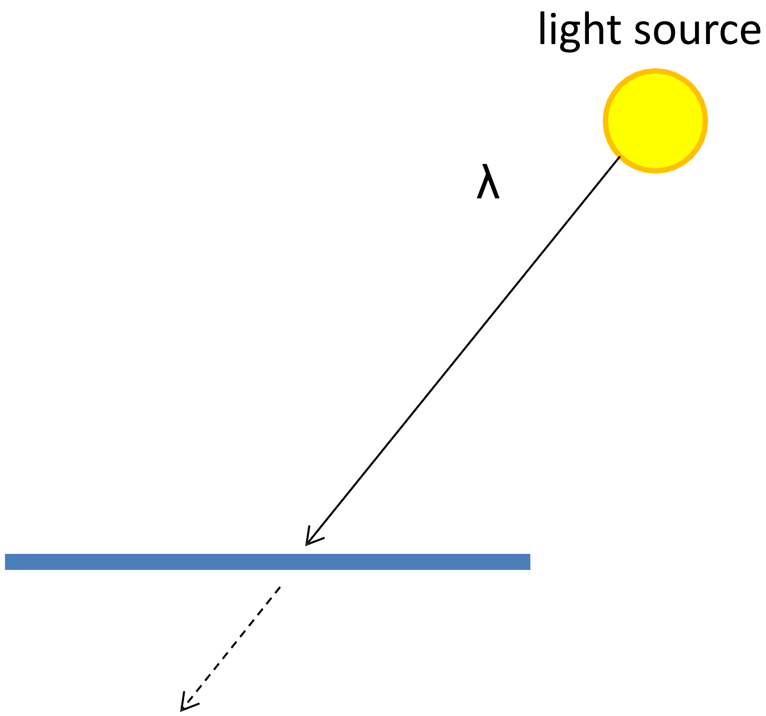

Photon

- Absorption
- Diffusion
- Reflection
- Transparency
- **Refraction**
- Fluorescence
- Subsurface scattering
- Phosphorescence
- Interreflection

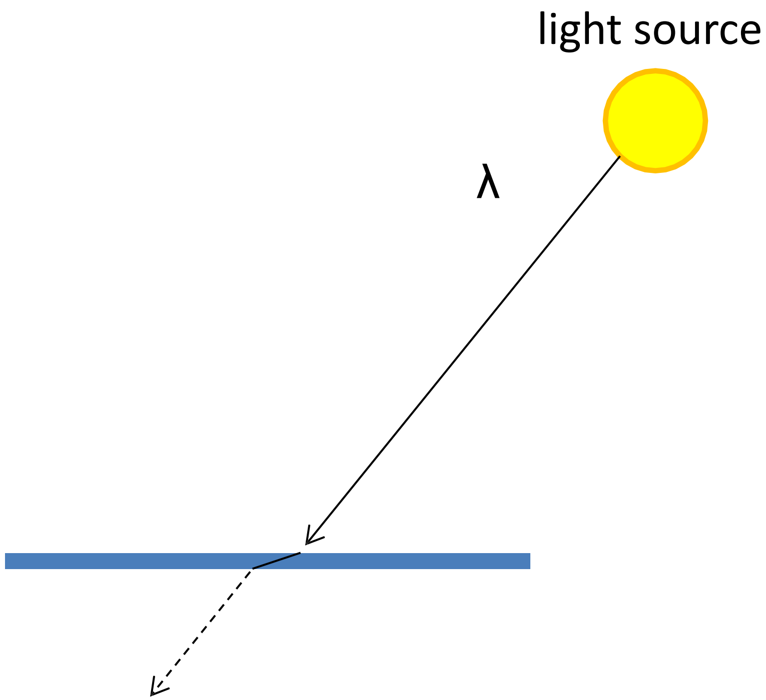

Photon

- Absorption
- Diffusion
- Reflection
- Transparency
- Refraction
- **Fluorescence**
- Subsurface scattering
- Phosphorescence
- Interreflection

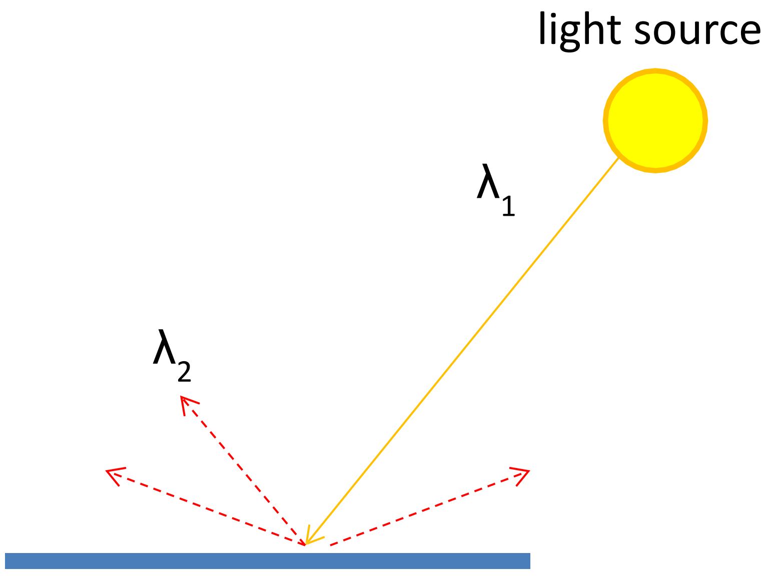

Photon

- Absorption
- Diffusion
- Reflection
- Transparency
- Refraction
- Fluorescence
- **Subsurface scattering**
- Phosphorescence
- Interreflection

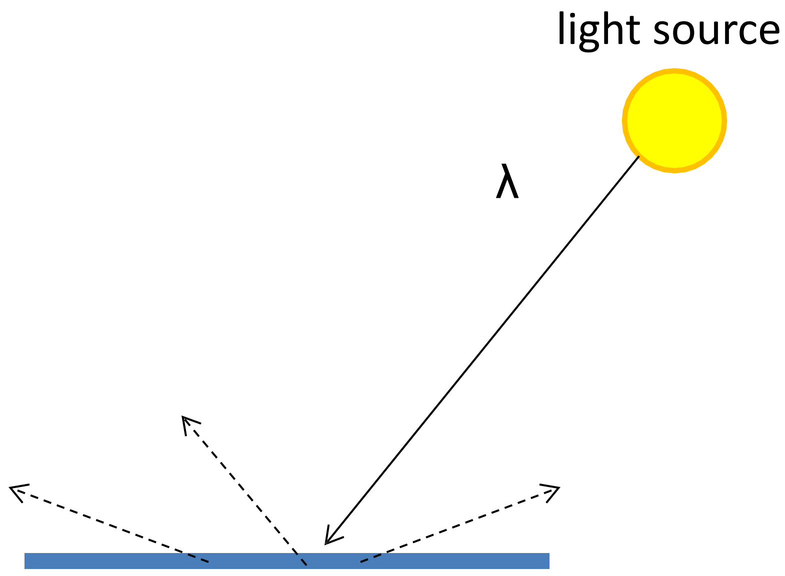

Photon

- Absorption
- Diffusion
- Reflection
- Transparency
- Refraction
- Fluorescence
- Subsurface scattering
- **Phosphorescence**
- Interreflection

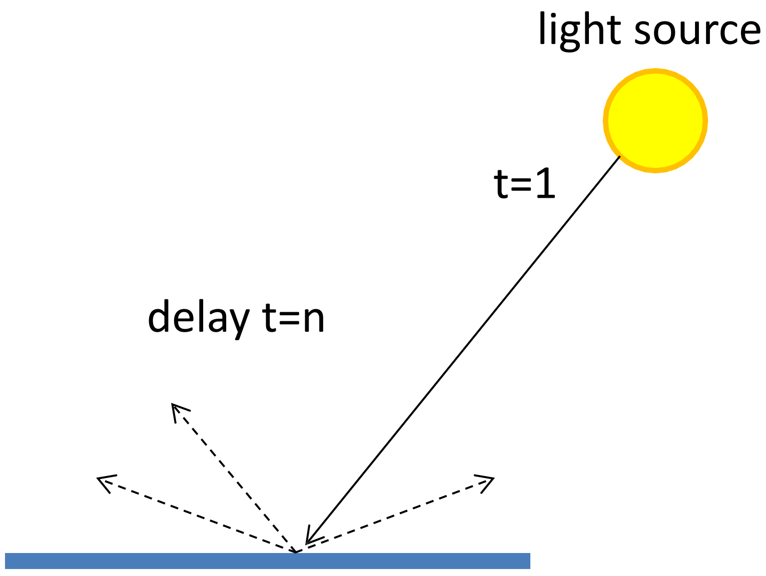

Photon

- Absorption
- Diffusion
- Reflection
- Transparency
- Refraction
- Fluorescence
- Subsurface scattering
- Phosphorescence
- **Interreflection**

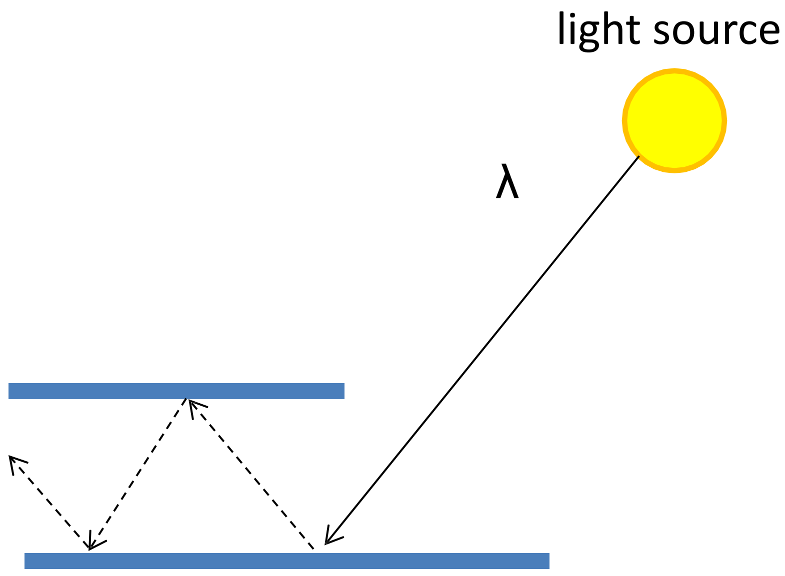

(Specular Interreflection)

Image Formation

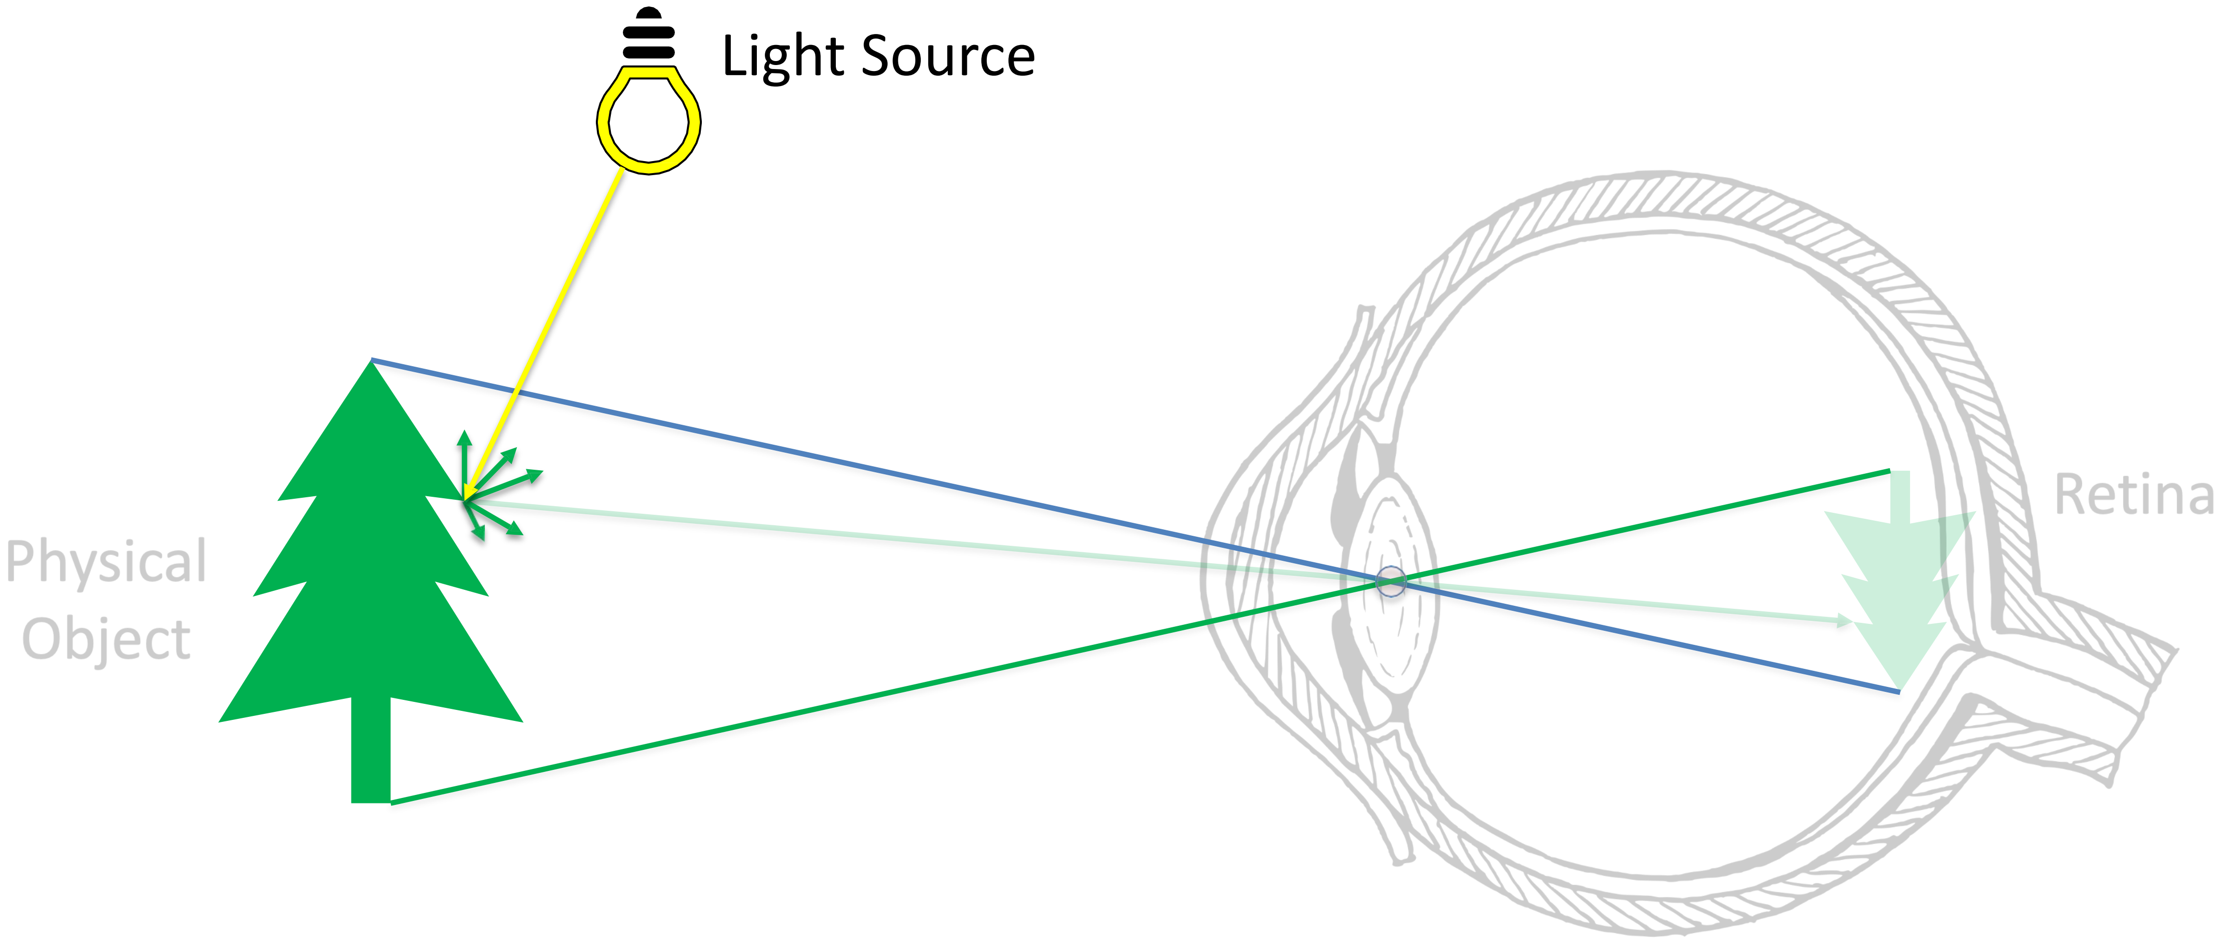

Image Formation

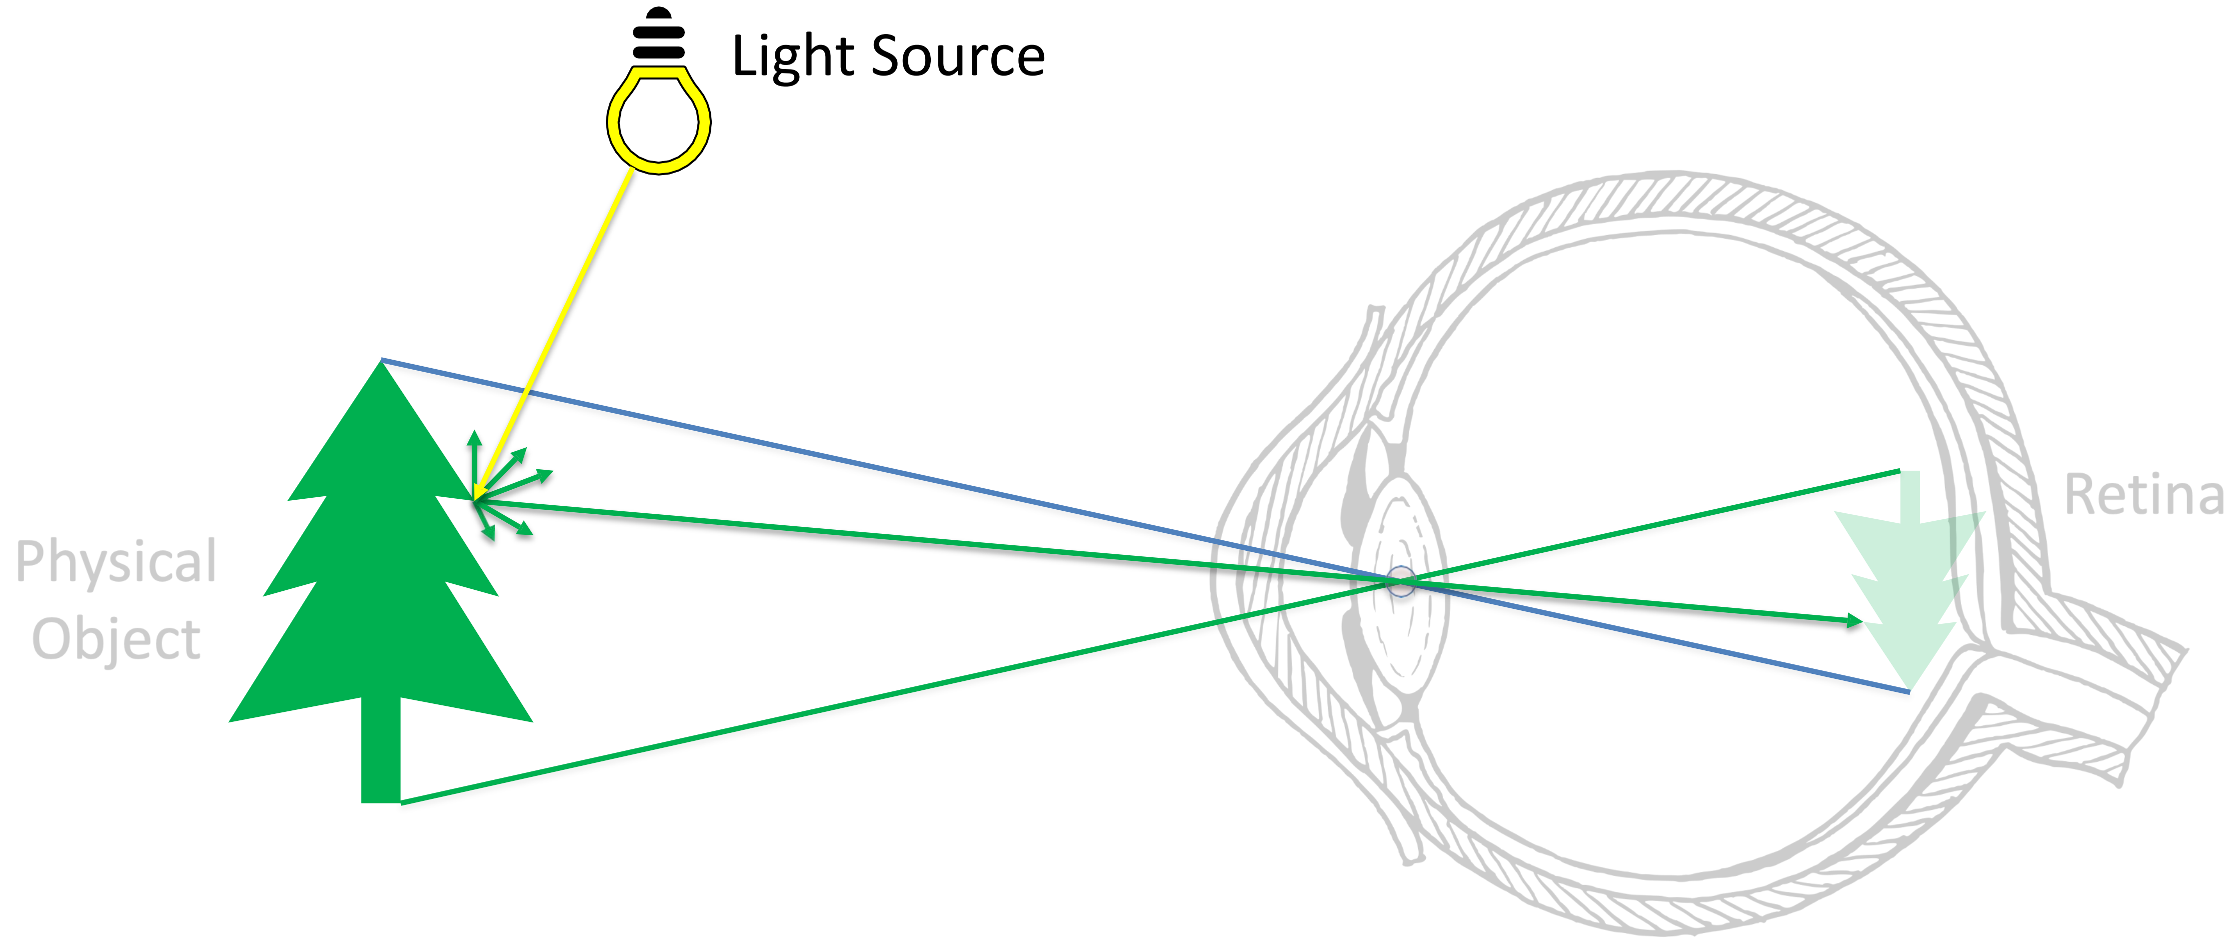

Radiometry: some definitions

- **Radiance**: power emitted per unit area in a direction
- **Irradiance**: total incident power falling on a surface

Directions specified by (polar angle, azimuth)

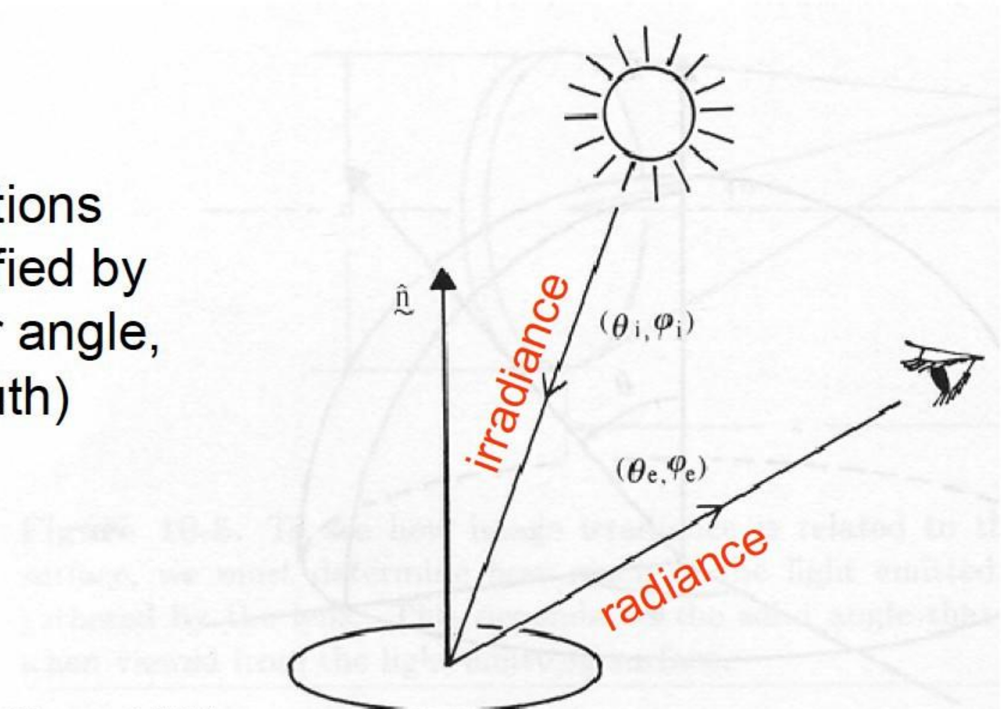

Horn, 1986

Radiometry: BRDF

- **Bidirectional reflectance distribution function**: Model of local reflection that tells how **bright** a surface appears when viewed from one direction when light falls on it from another.

Directions specified by (polar angle, azimuth)

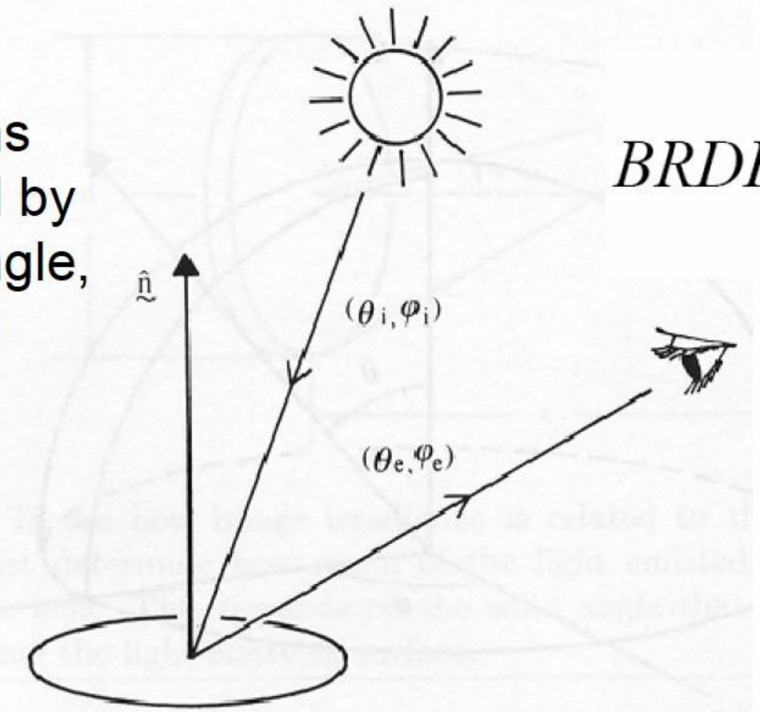

Horn, 1986

$$BRDF = f(\theta_i, \phi_i, \theta_e, \phi_e) = \frac{L(\theta_e, \phi_e)}{E(\theta_i, \phi_i)}$$

radiance / irradiance

$$BRDF = f(\theta_i, \phi_i, \theta_e, \phi_e) = \frac{L(\theta_e, \phi_e)}{E(\theta_i, \phi_i)}$$

$$BRDF = f(\theta_i, \phi_i, \theta_e, \phi_e, \lambda) = \frac{L(\theta_e, \phi_e, \lambda)}{E(\theta_i, \phi_i, \lambda)}$$

Spectral radiance / spectral irradiance

…extend radiometry terms to incorporate *spectral* units (per unit wavelength)

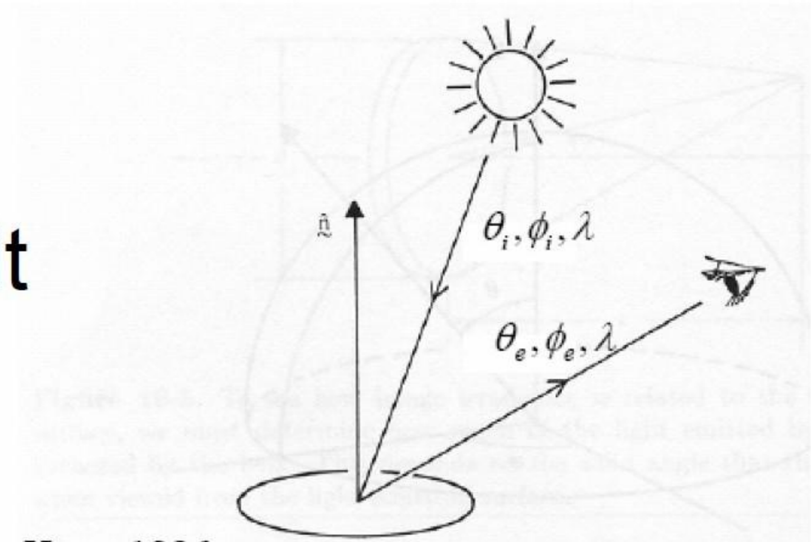

Horn, 1986

Radiometry: BRDF

- BRDF is a very general notion
  - some surfaces need it (underside of a CD; tiger eye; etc)
  - very hard to measure
    - illuminate from one direction, view from another, repeat
  - very unstable
    - minor surface damage can change the BRDF
    - e.g. ridges of oil left by contact with the skin can act as lenses
- For many surfaces, light leaving the surface is largely independent of exit angle

Lambertian Reflectance

- Lambertian/diffuse surfaces
  - Appear equally bright from all viewing directions

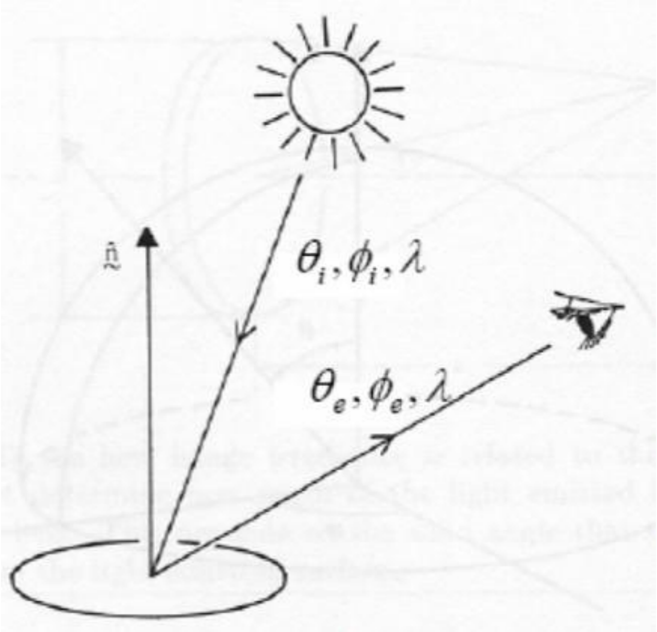

http://en.wikipedia.org/wiki/Lambertian_reflectance

Simplified rendering models: BRDF reflectance

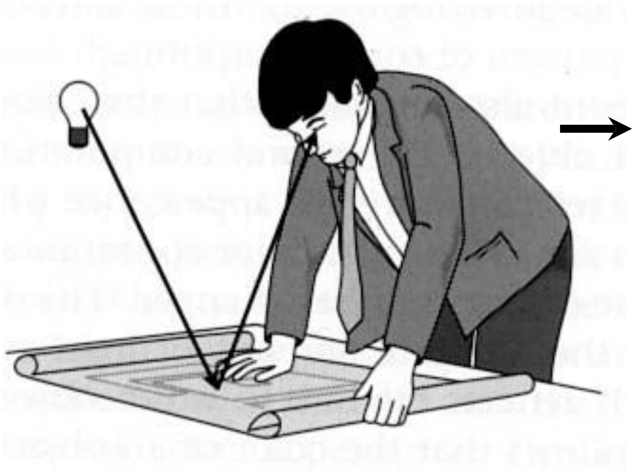

For diffuse reflections, we replace the BRDF calculation with a wavelength-by-wavelength scalar multiplication

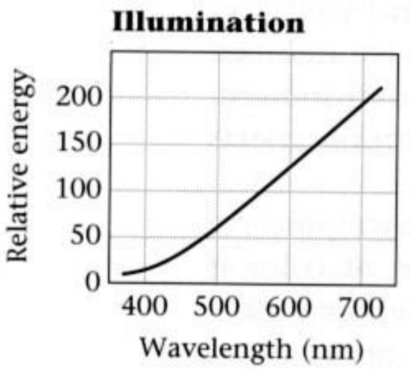

.*

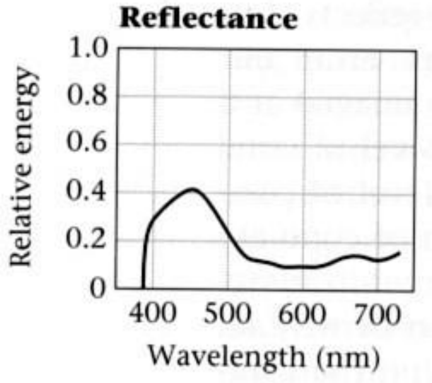

=

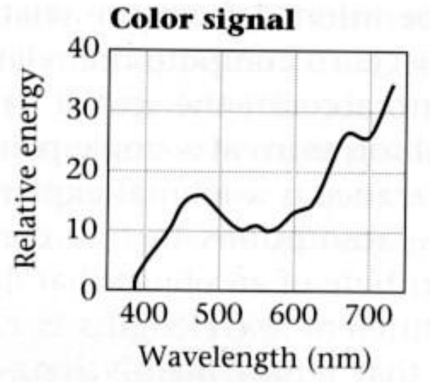

Foundations of Vision, by Brian Wandell, Sinauer Assoc., 1995

The Physics of Light

Some examples of the spectra of light sources

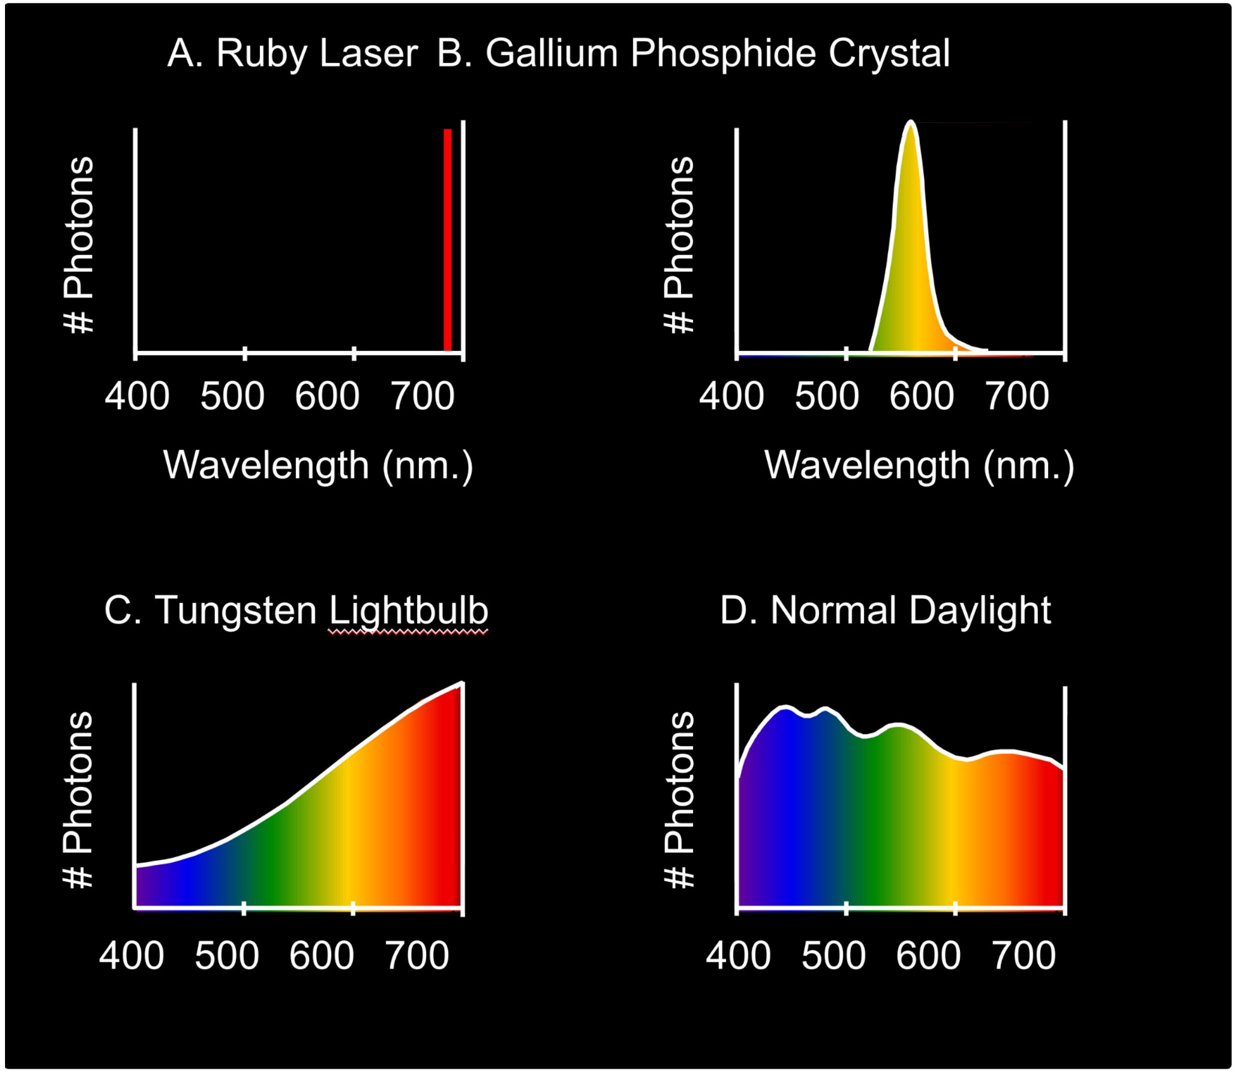

The Physics of Light

Some examples of the reflectance spectra of surfaces

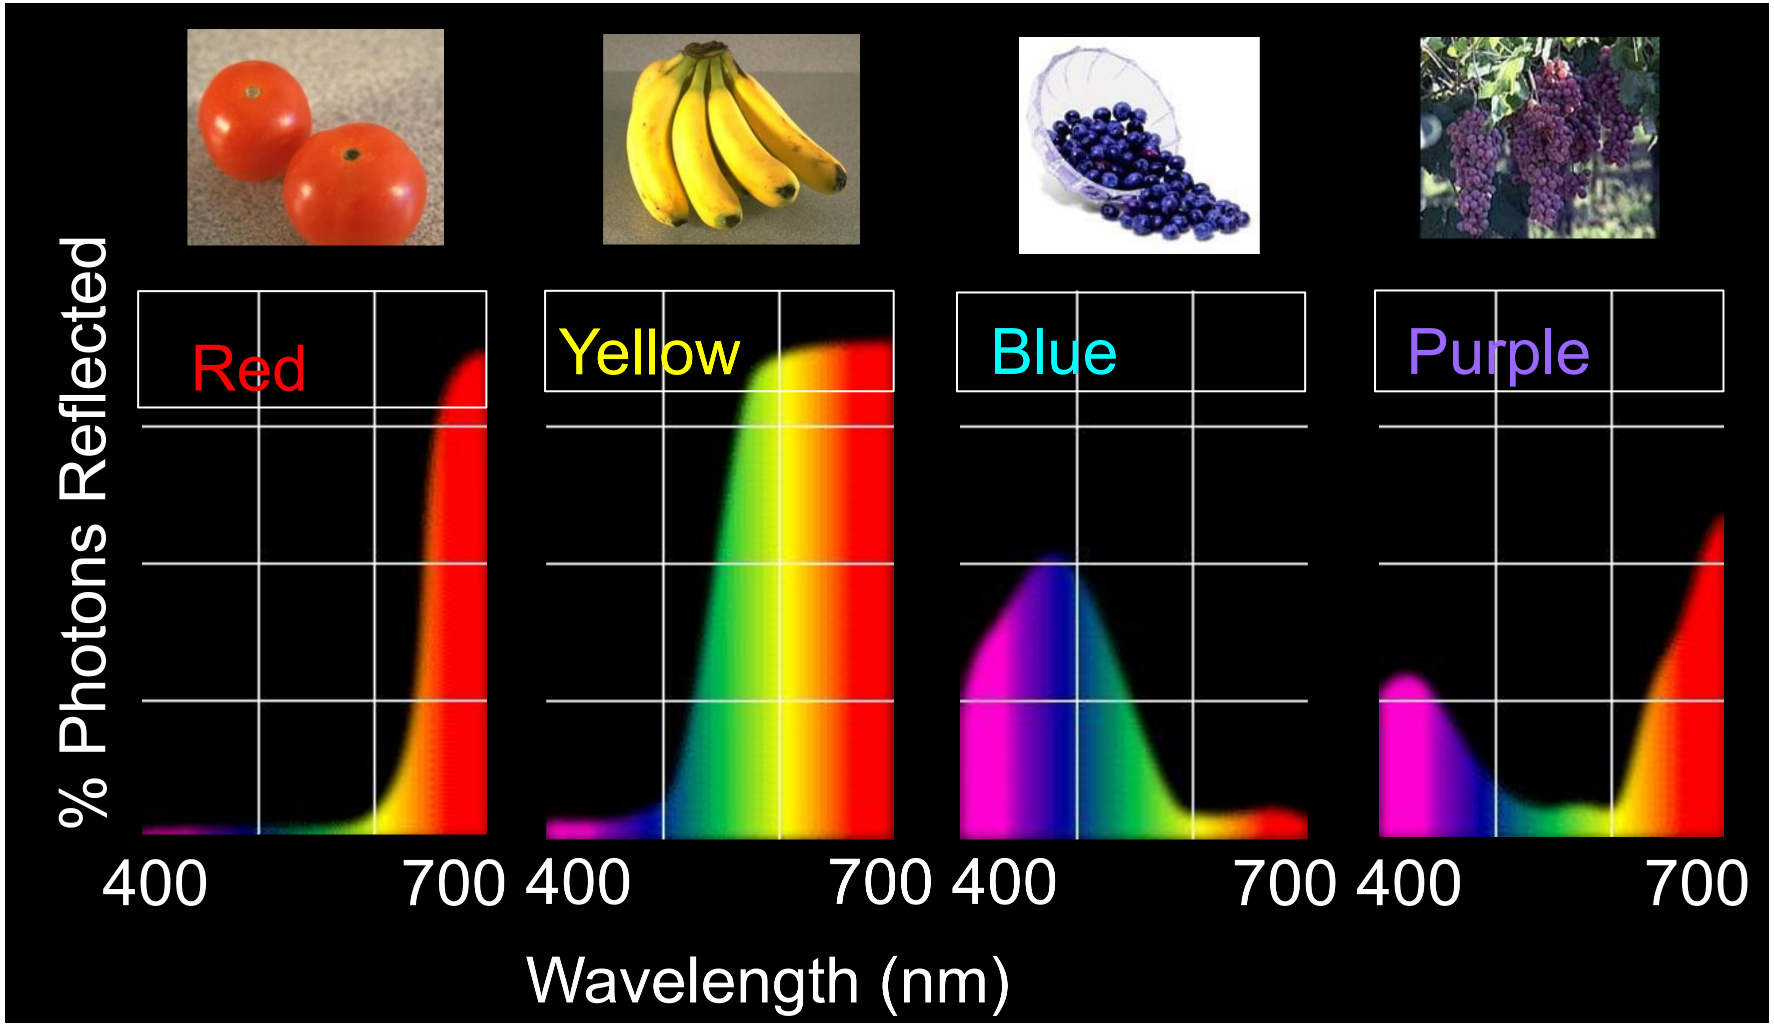

Some reflectance spectra for different leaves

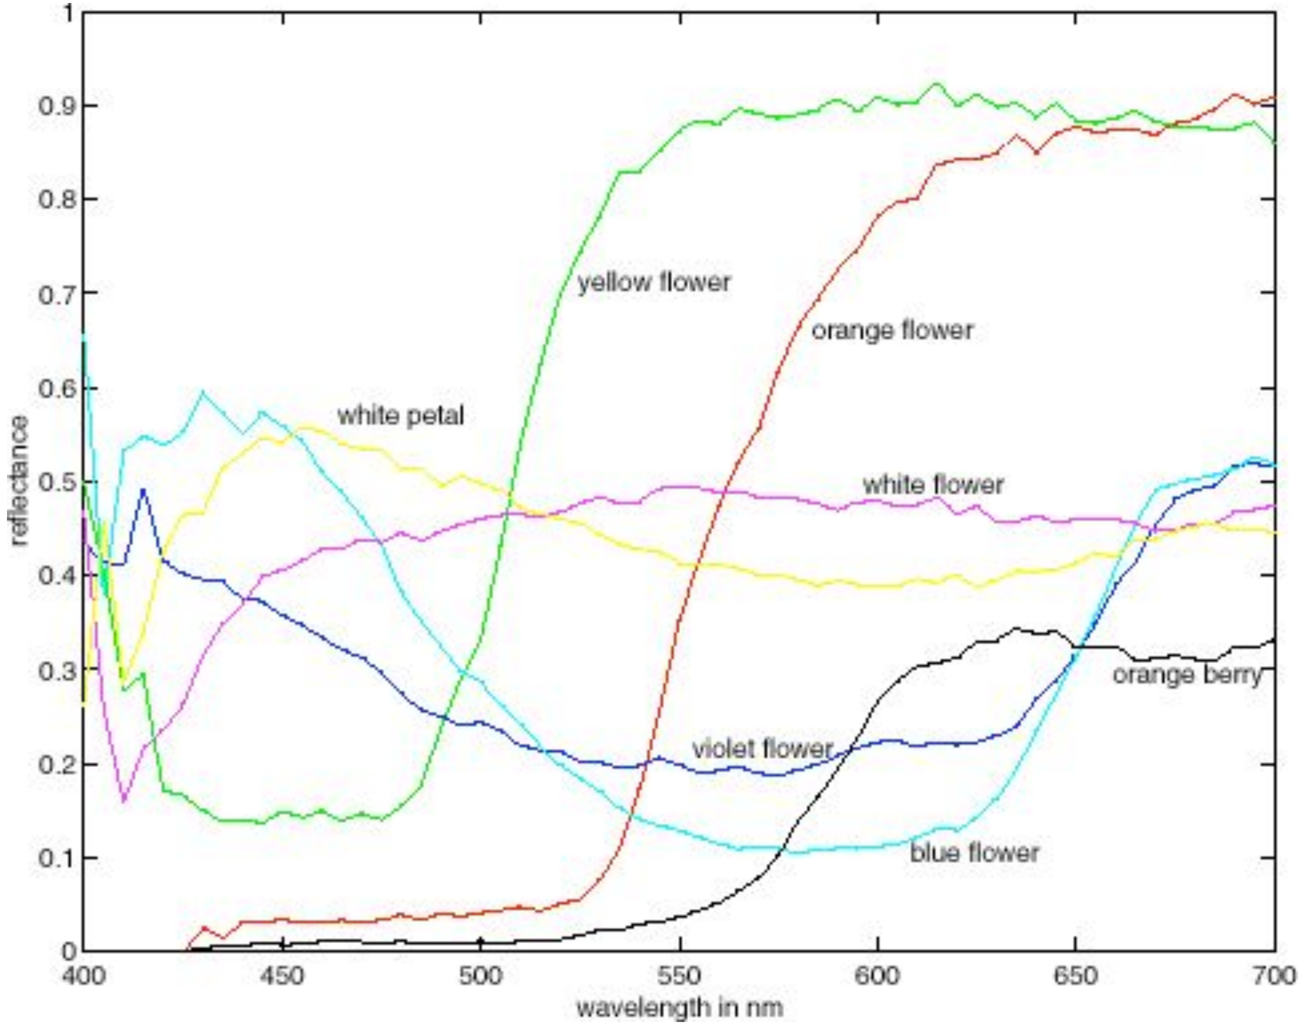

Albedo is the fraction of light that a surface reflects

Here are the spectral albedoes for several different leaves, with their perceived colors.

Naturally, different colours have different spectral albedo.’

But different spectral albedoes sometimes result in the same perceived color (compare the two whites).

Spectral albedoes are typically quite smooth functions.

Source: Measurements by E.Koivisto.

Today’s agenda

- Image formation
- Physics of Color
- **Color matching**
- Color spaces
- Image sampling and quantization

Image Formation

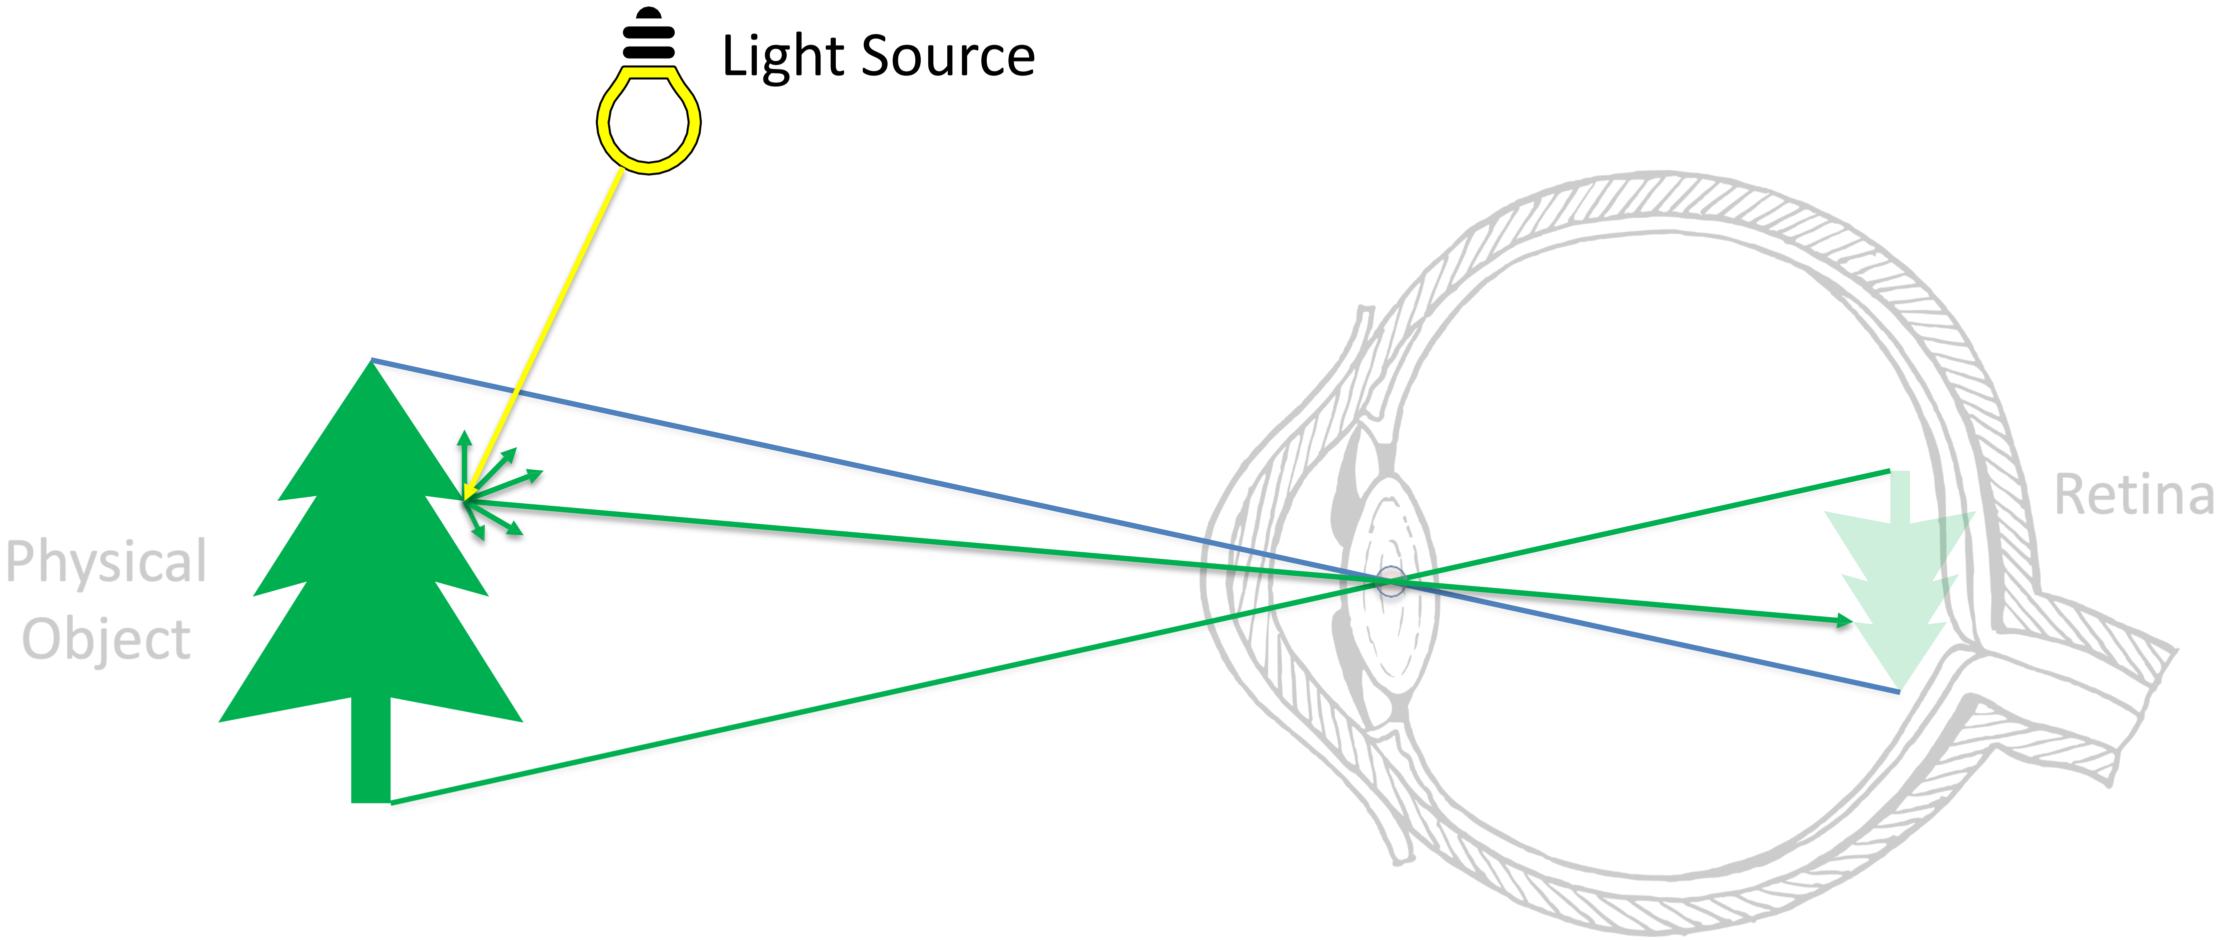

Image Formation

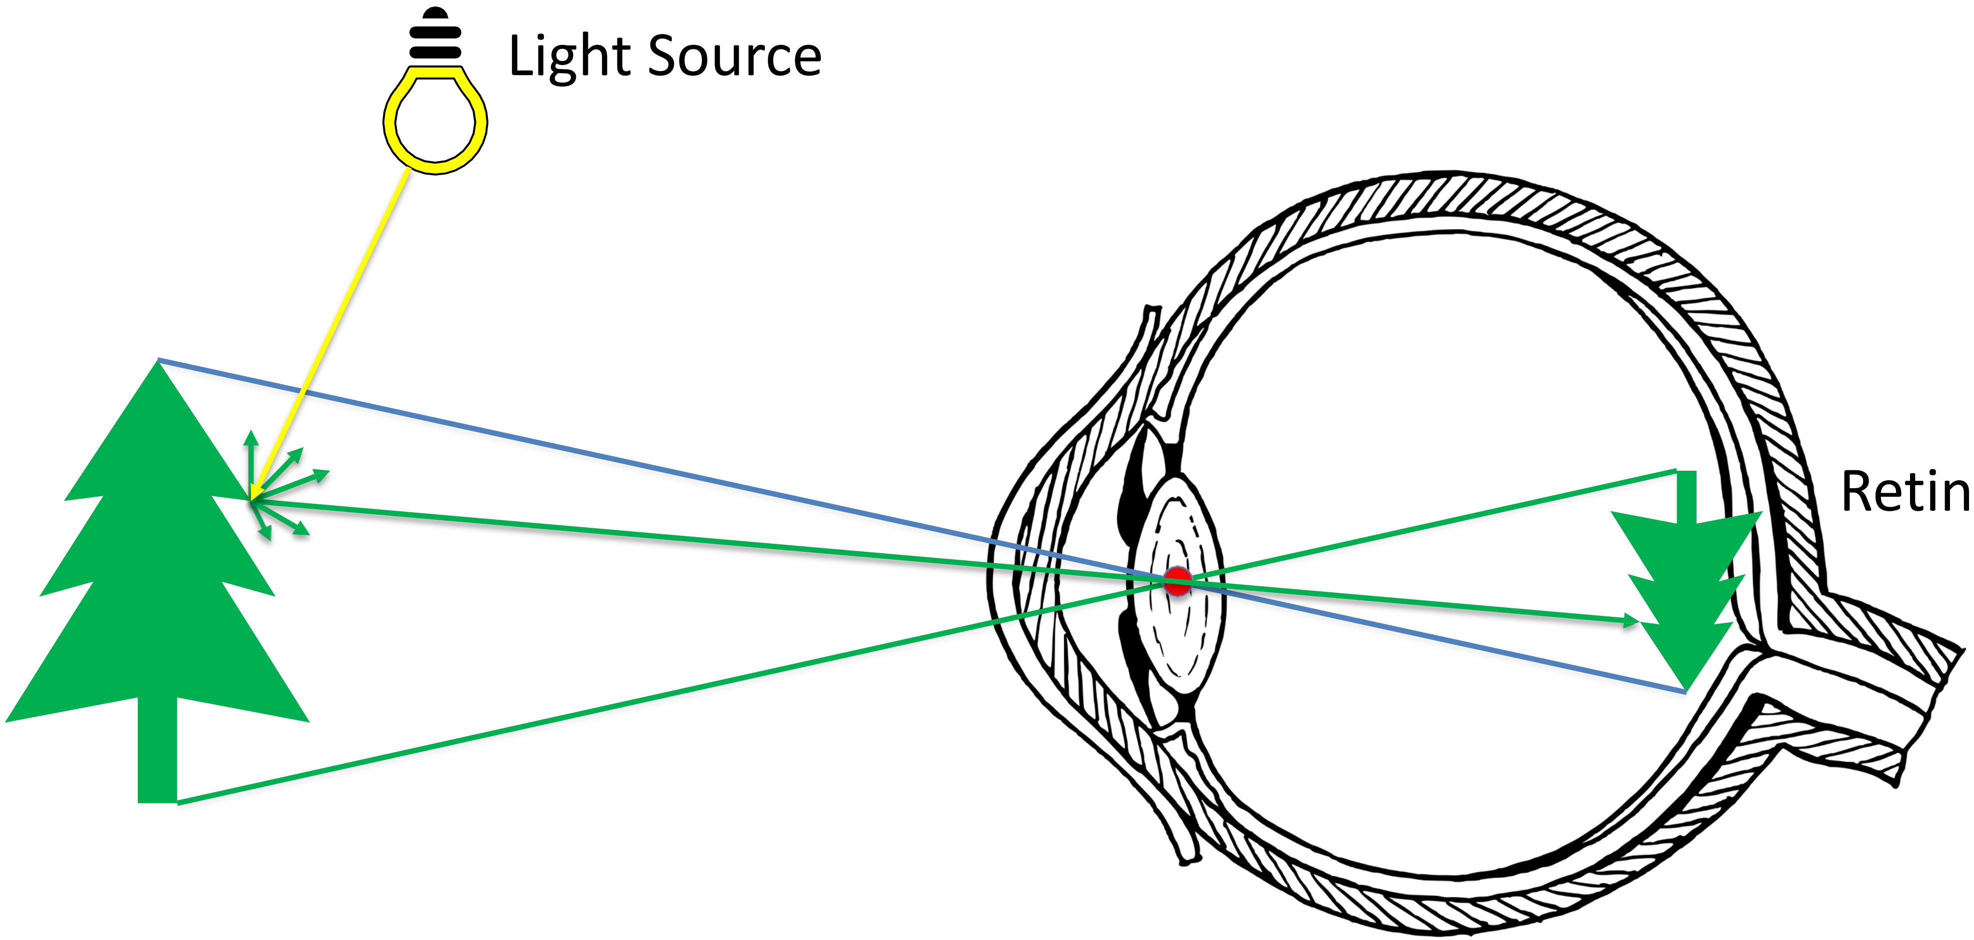

Why specify color numerically?

- People will describe the same color using different words:
  - Be able to tell if we are talking about the same color
- Accurate color reproduction is commercially valuable
  - Many products are identified by color ("golden" arches) 
- Few color names are widely recognized by English speakers
  - About 10; other languages have fewer/more, but not many more.
  - Common to disagree on appropriate color names.
- Color reproduction problems increased by prevalence of digital imaging – e.g. digital libraries of art.
  - How to ensure that everyone perceives the same color?
  - What spectral radiances produce the same response from people under simple viewing conditions?

The assumption for color perception

- We know color appearance really depends on:
  - The illumination
  - Your eye's adaptation level
  - The colors and scene interpretation surrounding the observed color.

- But for now we will assume that the spectrum of the light arriving at your eye completely determines the perceived color.

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

What's the machinery in the eye? Human Photoreceptors

Cones and rods near the

near the fovea periphery

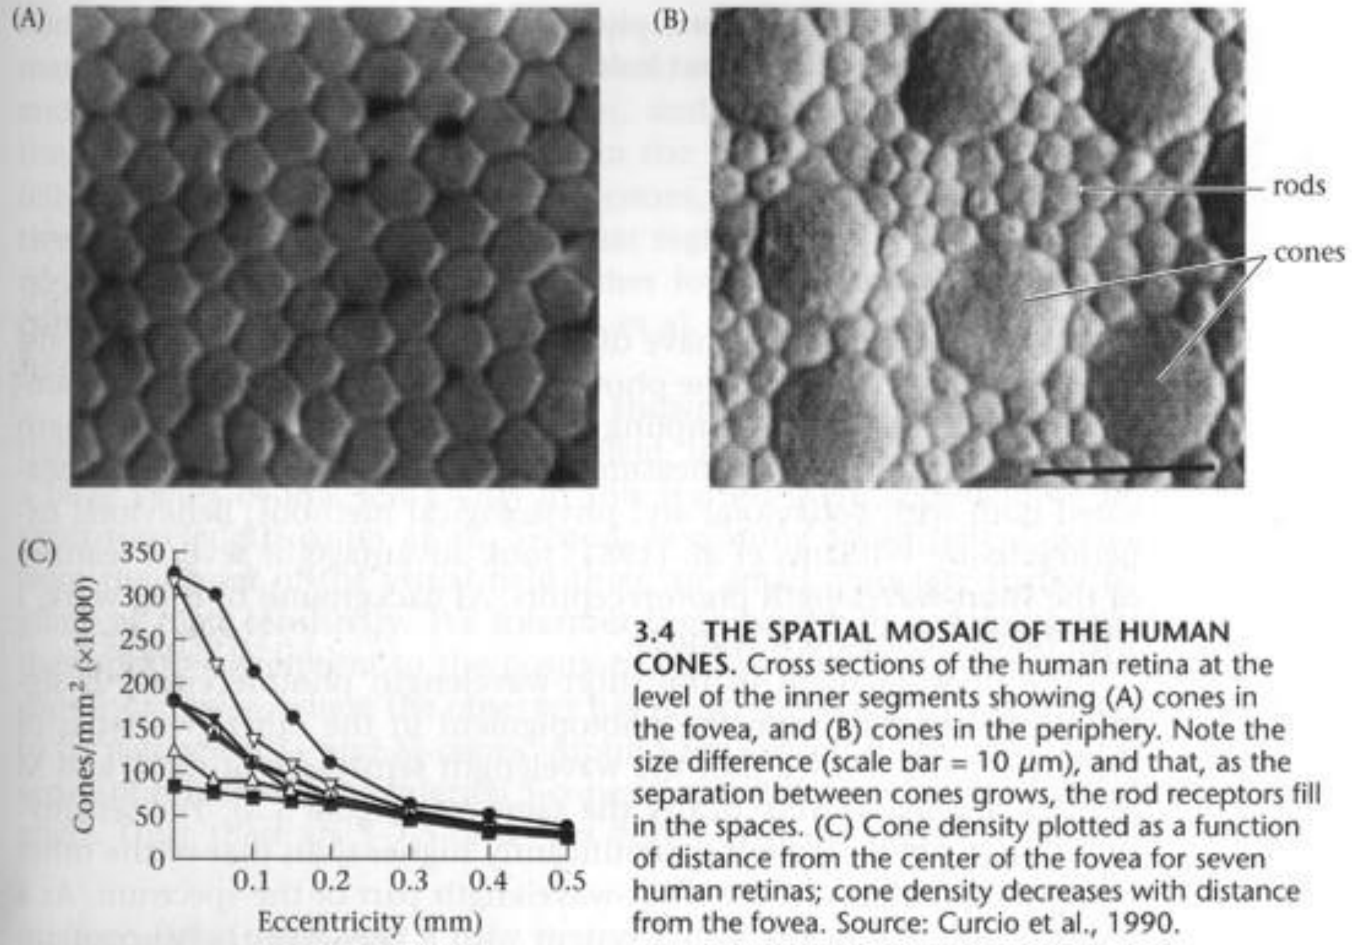

3.4 THE SPATIAL MOSAIC OF THE HUMAN CONES. Cross sections of the human retina at the level of the inner segments showing (A) cones in the fovea, and (B) cones in the periphery. Note the size difference (scale bar = 10 μm), and that, as the separation between cones grows, the rod receptors fill in the 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

Physiology of Color Vision

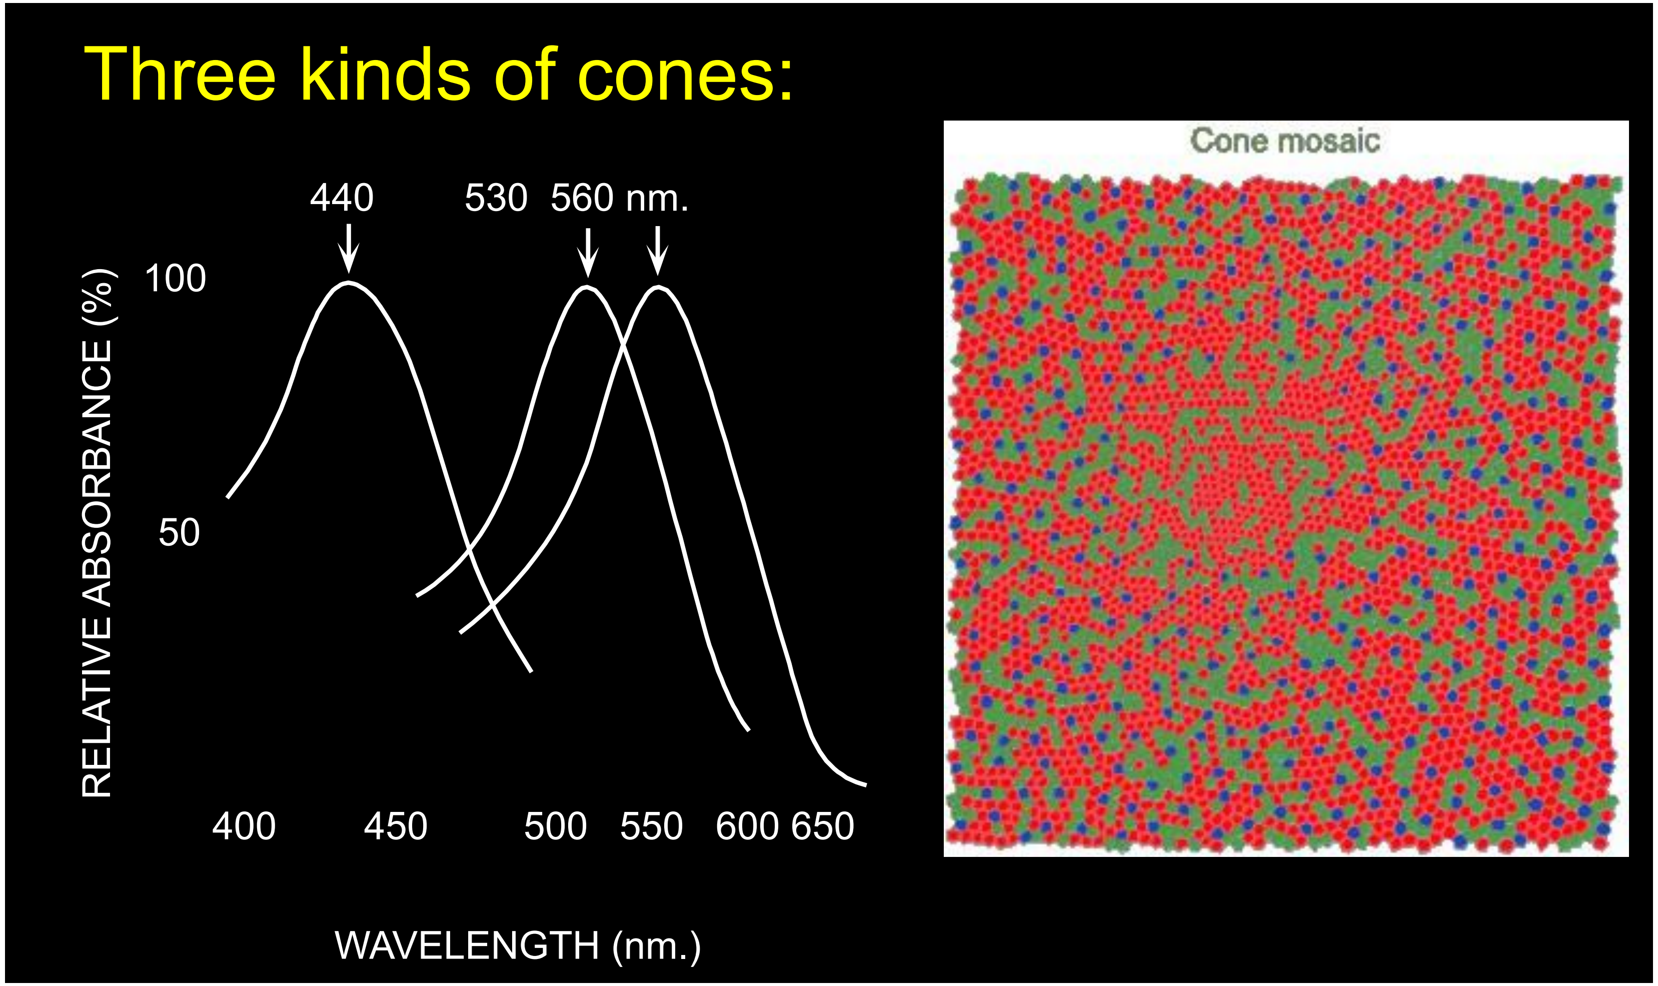

Human eye photoreceptor spectral sensitivities

3.3 SPECTRAL SENSITIVITIES OF THE L-, M-, AND S-CONES in the human eye. The measurements are based on a light source at the cornea, so that the wavelength loss due to the cornea, lens, and other inert pigments of the eye plays a role in determining the sensitivity. Source: Stockman and MacLeod, 1993.

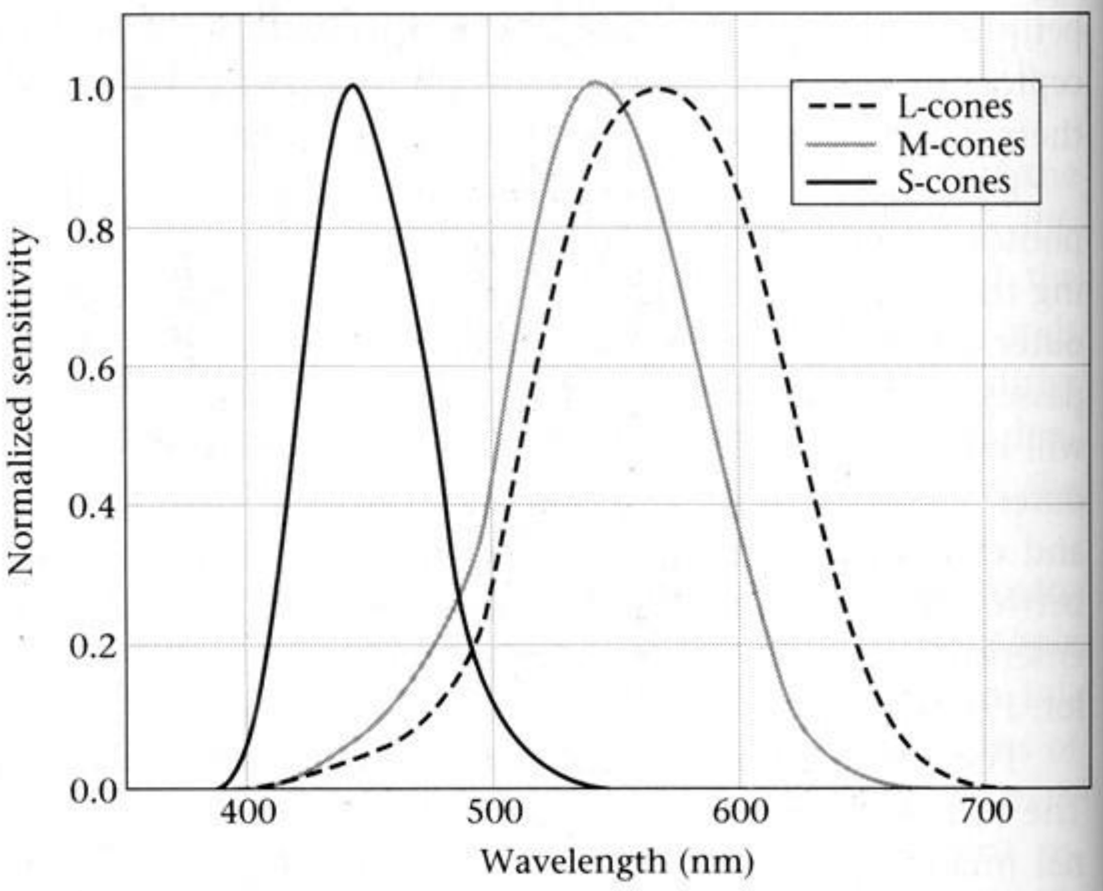

Foundations of Vision, by Brian Wandell, Sinauer Assoc., 1995

How we sense light spectra

test light

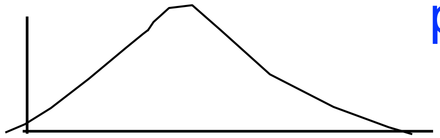

project →

Cone sensitivities

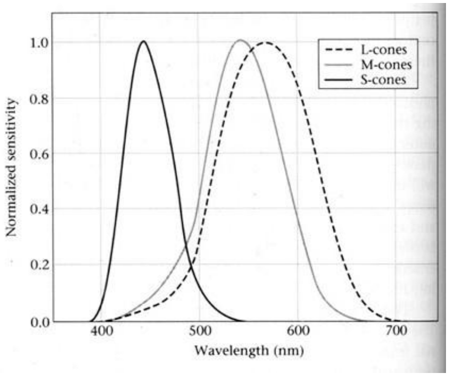

→ L, M, S responses

biophysics: integrate the response over all wavelengths, weighted by the photosensor's sensitivity at each wavelength.

mathematically: take dot product of input spectrum with the cone sensitivity basis vectors. Project the high-dimensional test light into a 3-d space.

Cone response curves as basis vectors in a 3- d subspace of light power spectra

3-d depiction of the high-dimensional space of all possible power spectra

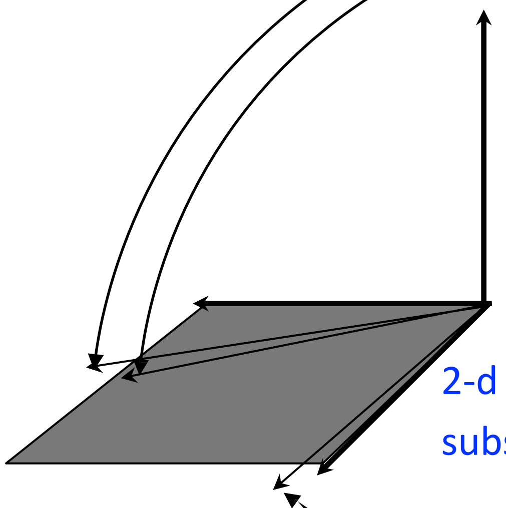

2-d depiction of the 3-d subspace of sensor responses

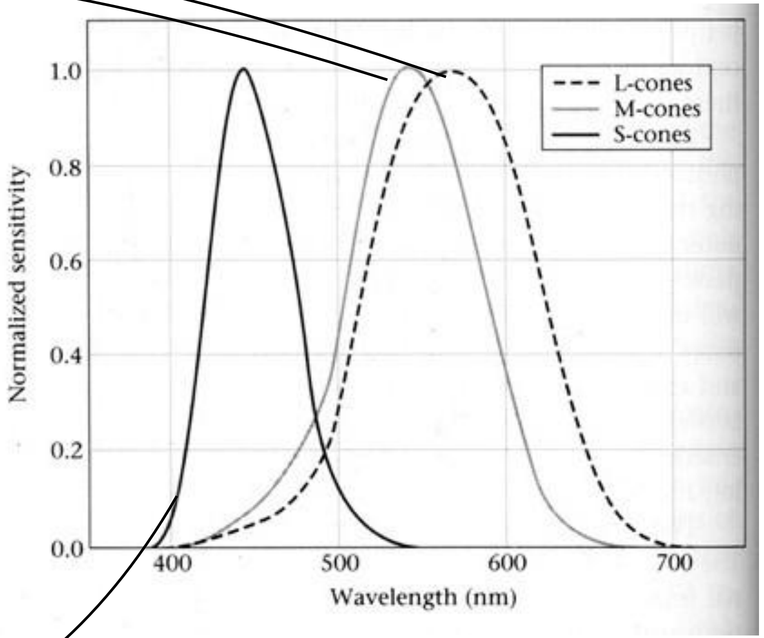

Spectral sensitivities of L, M, and S cones

Color matching experiment

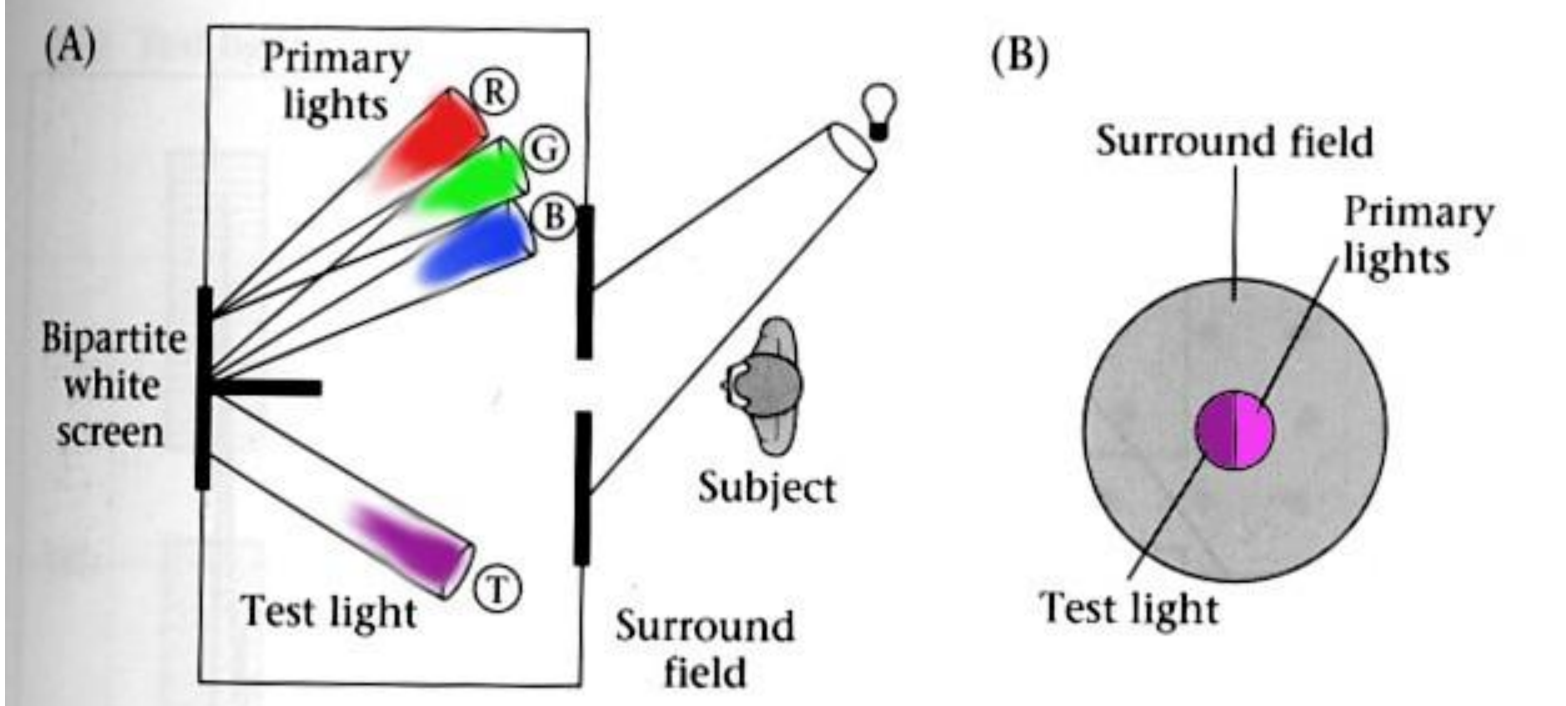

4.10 THE COLOR-MATCHING EXPERIMENT. The observer views a bipartite field and adjusts the intensities of the three primary lights to match the appearance of the test light. (A) A top view of the experimental apparatus. (B) The appearance of the stimuli to the observer. After Judd and Wyszecki, 1975.

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

Color matching experiment 1

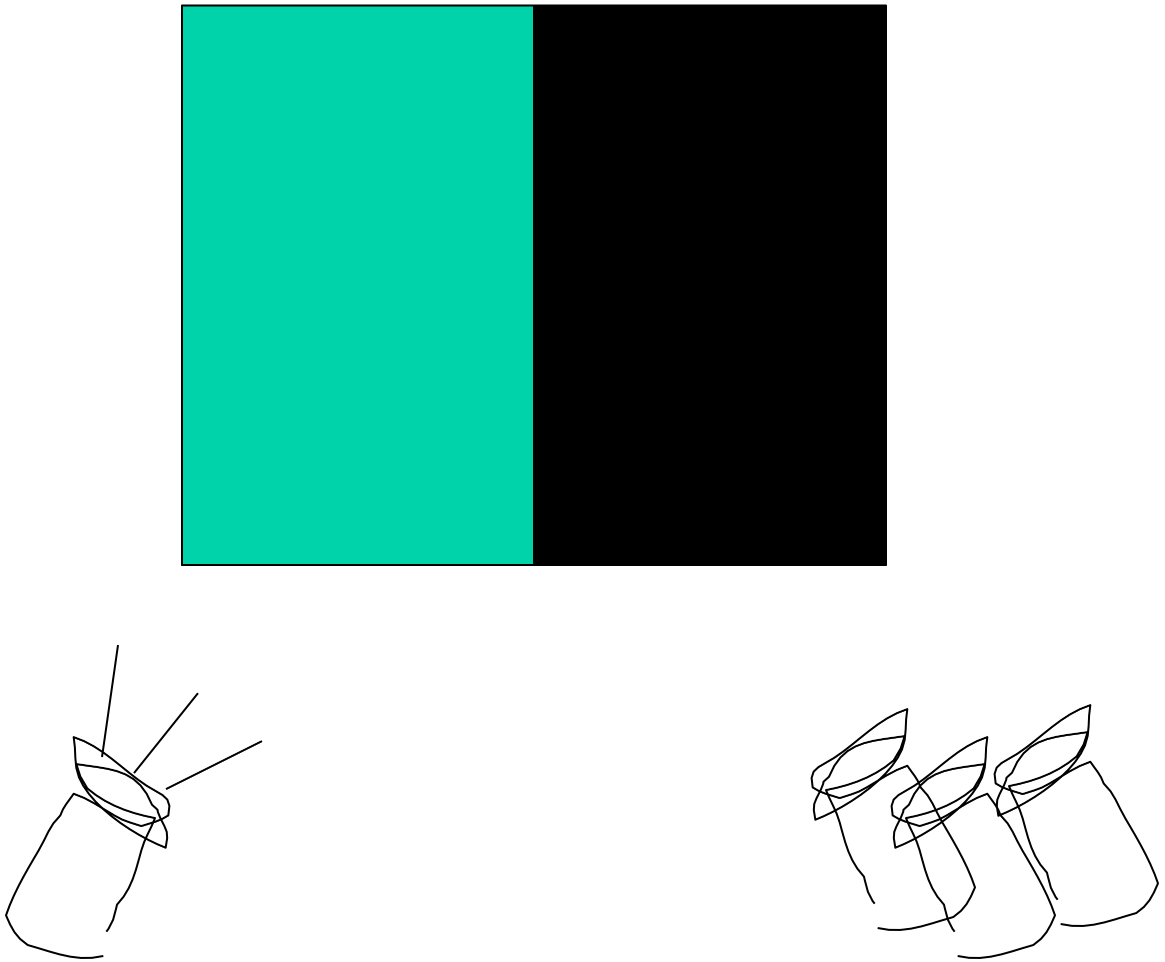

Color matching experiment 1

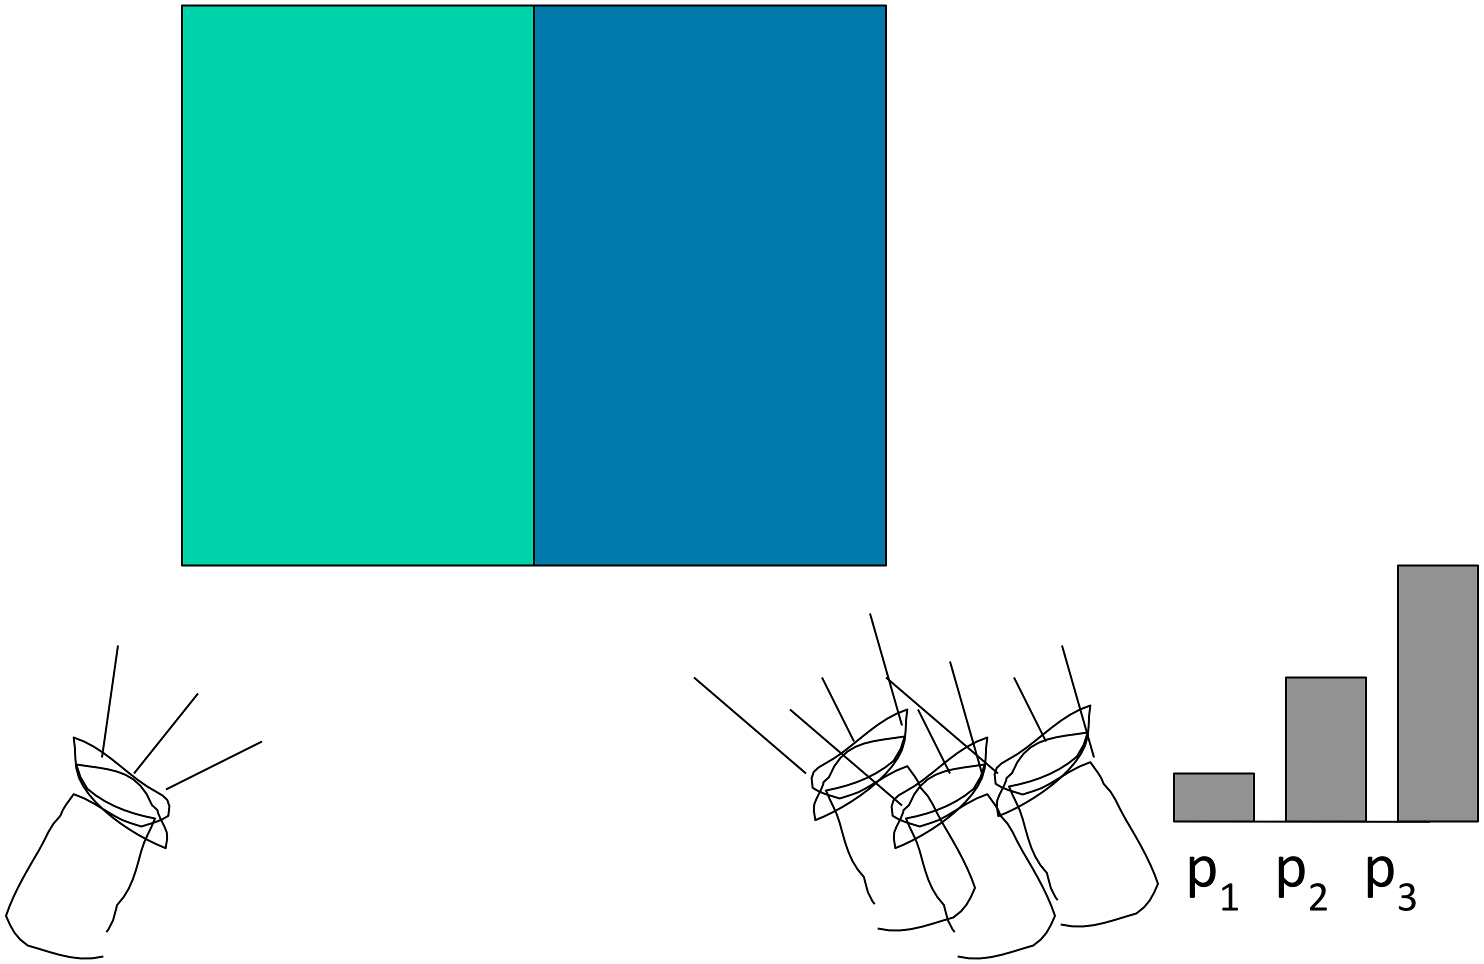

Color matching experiment 1

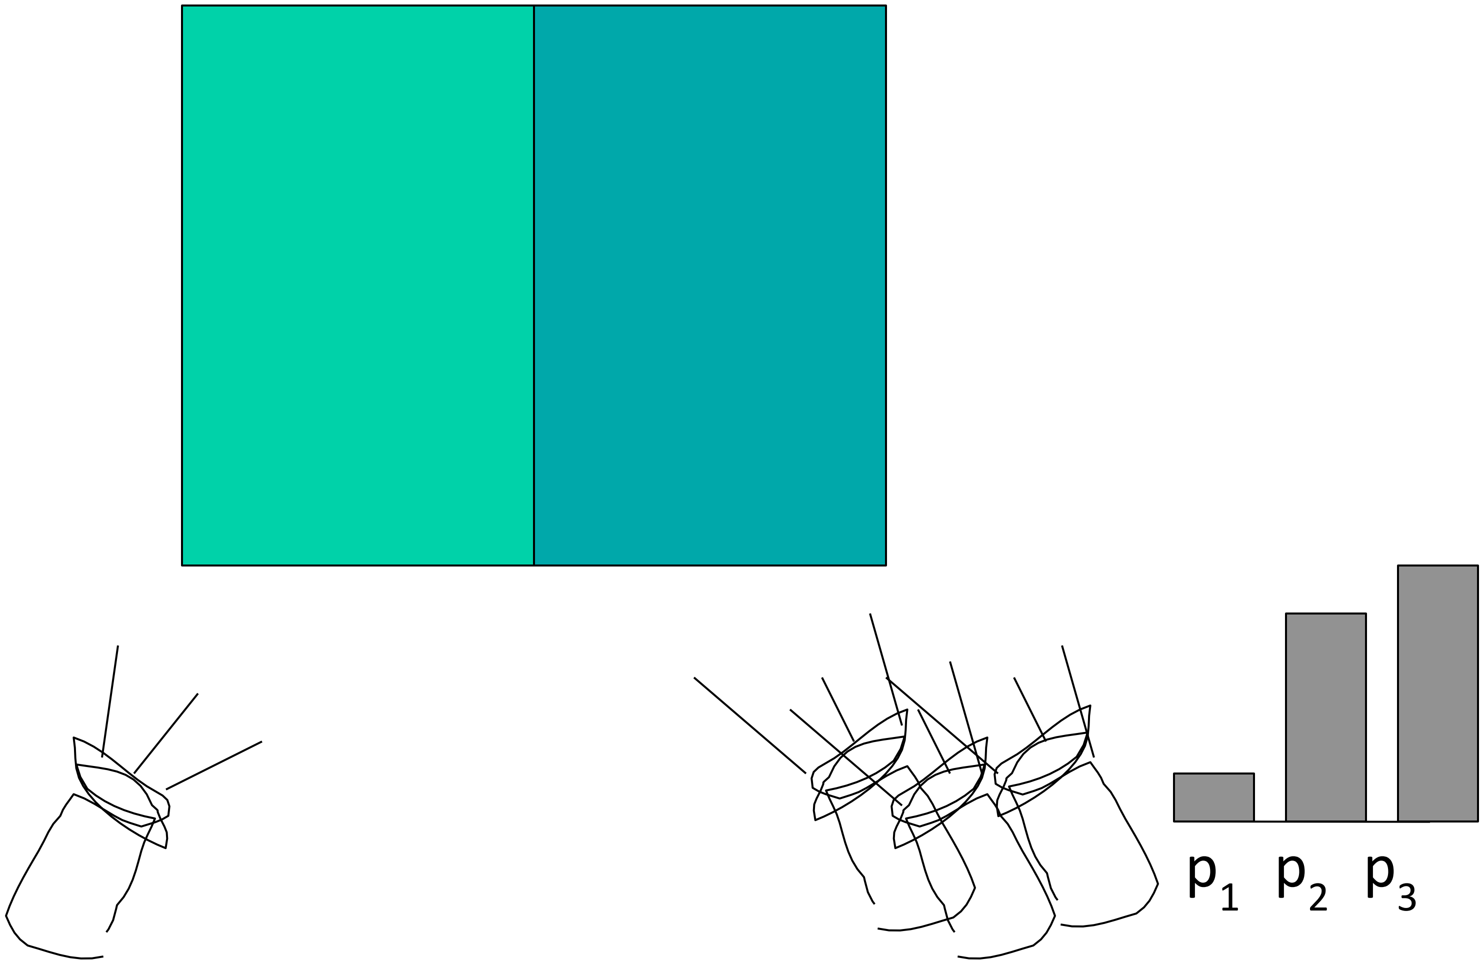

Color matching experiment 1

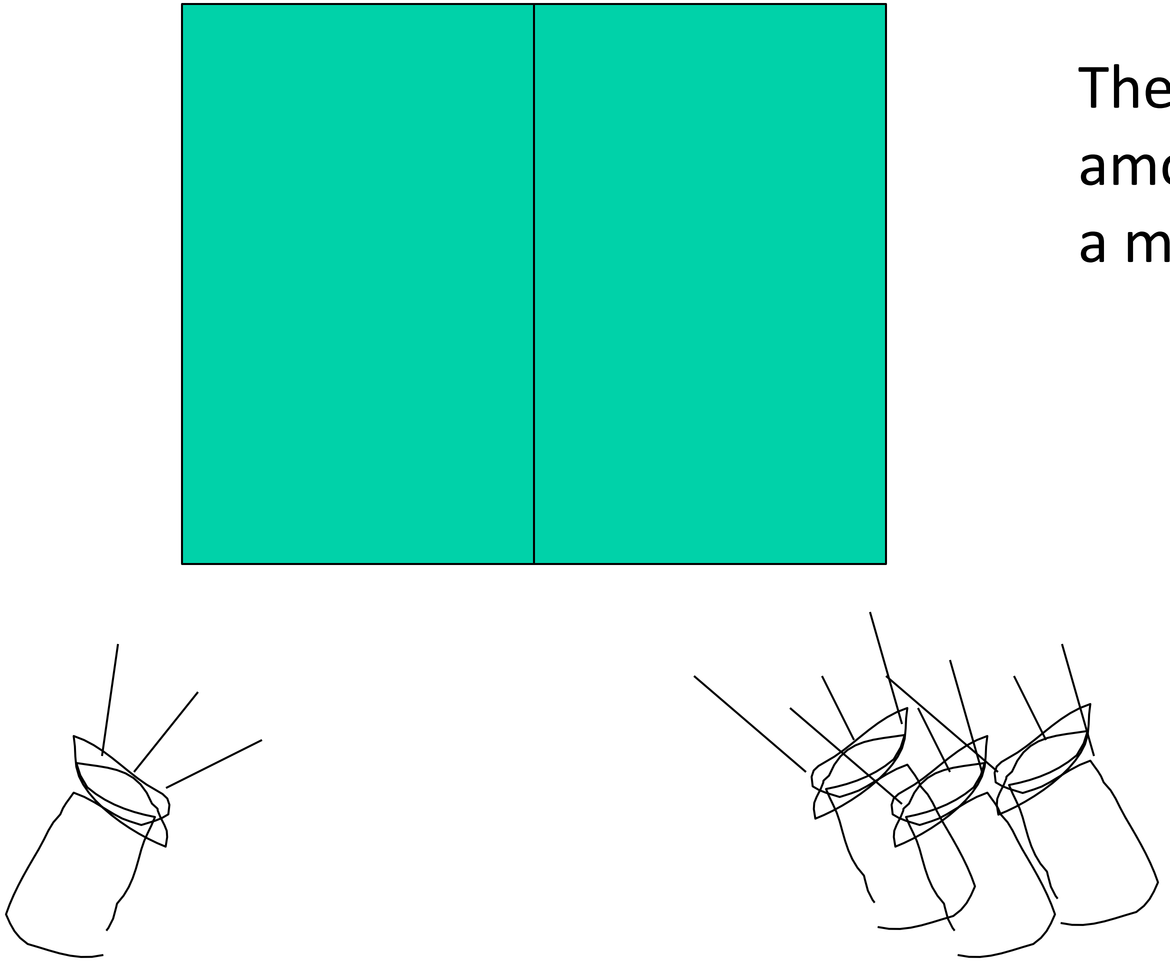

The primary color amounts needed for a match

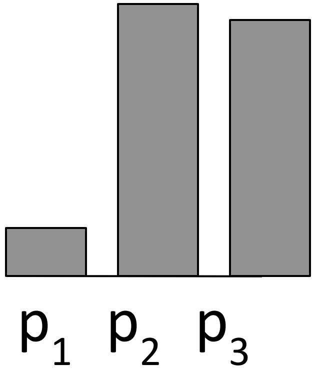

Color matching experiment 2

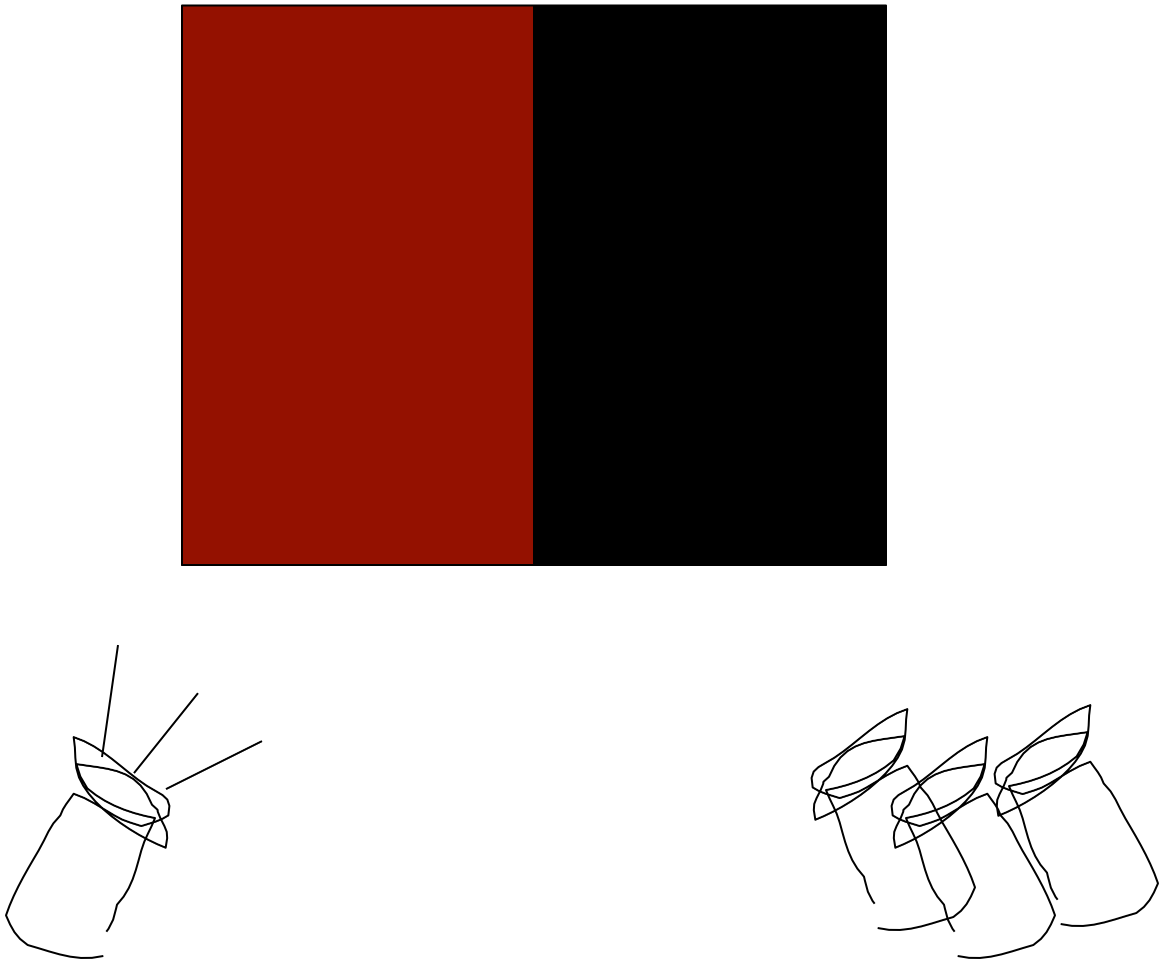

Color matching experiment 2

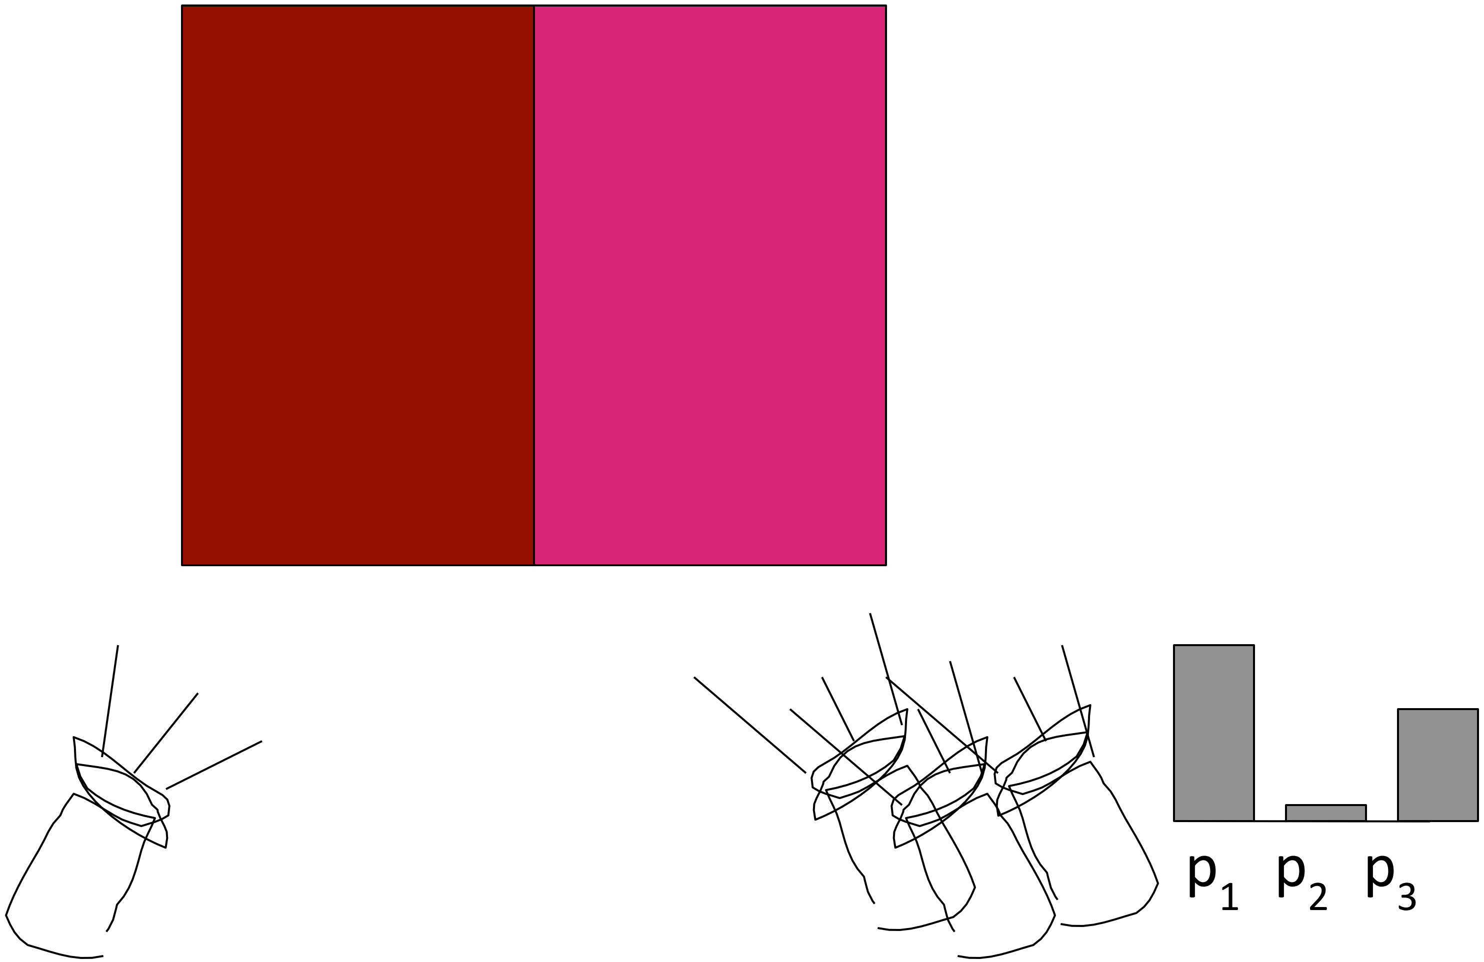

Color matching experiment 2

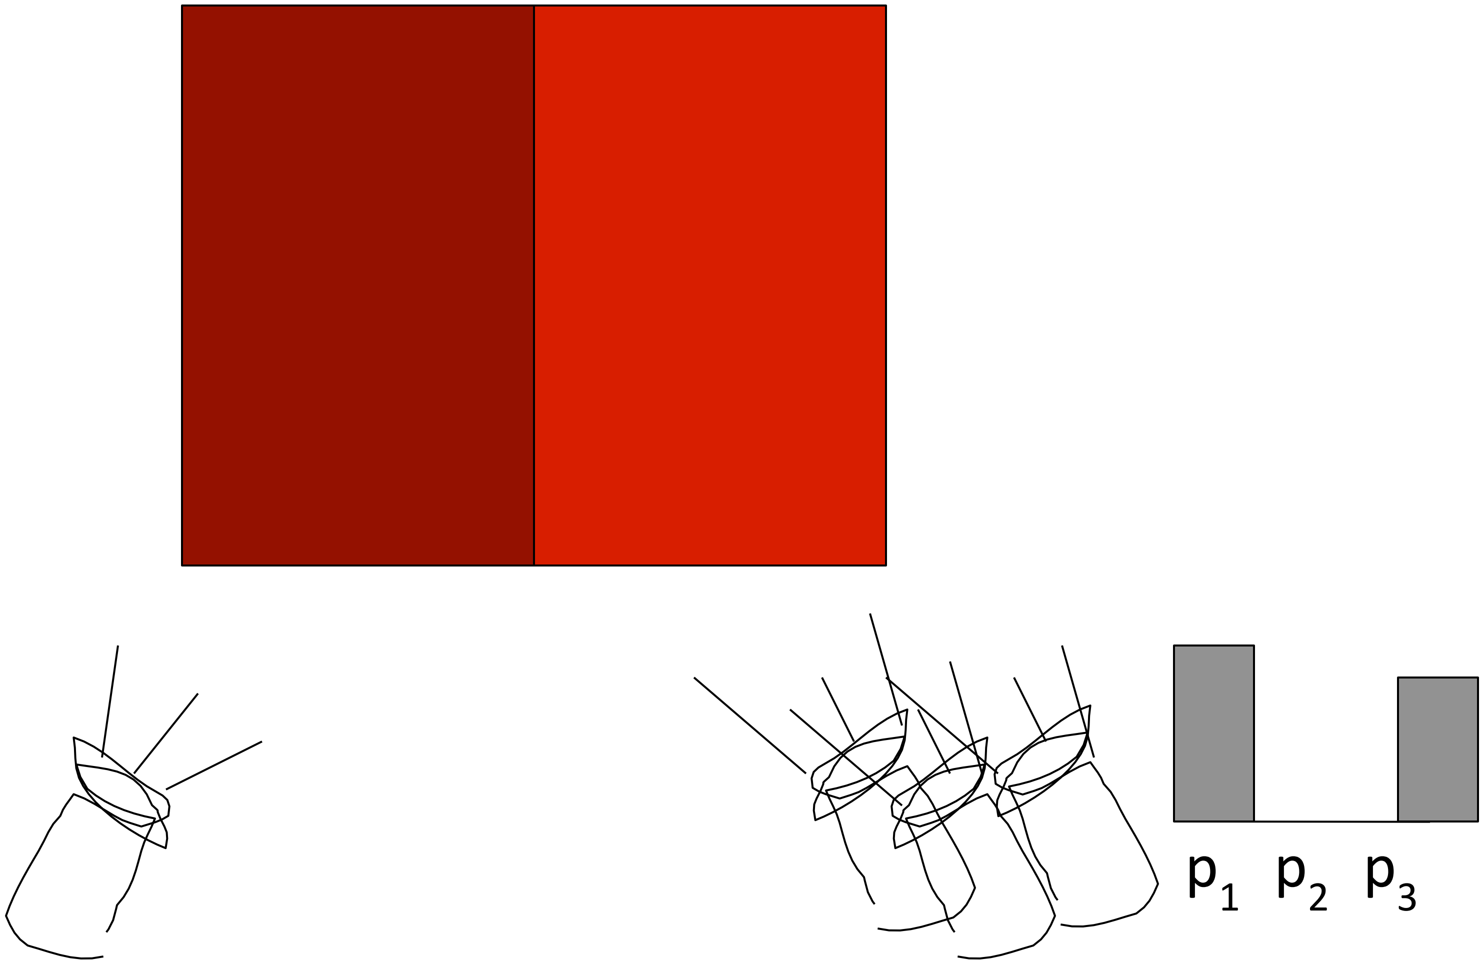

Color matching with positive amounts of the primaries

Match the sensors' response to the target light to the sum of responses to the primary lights

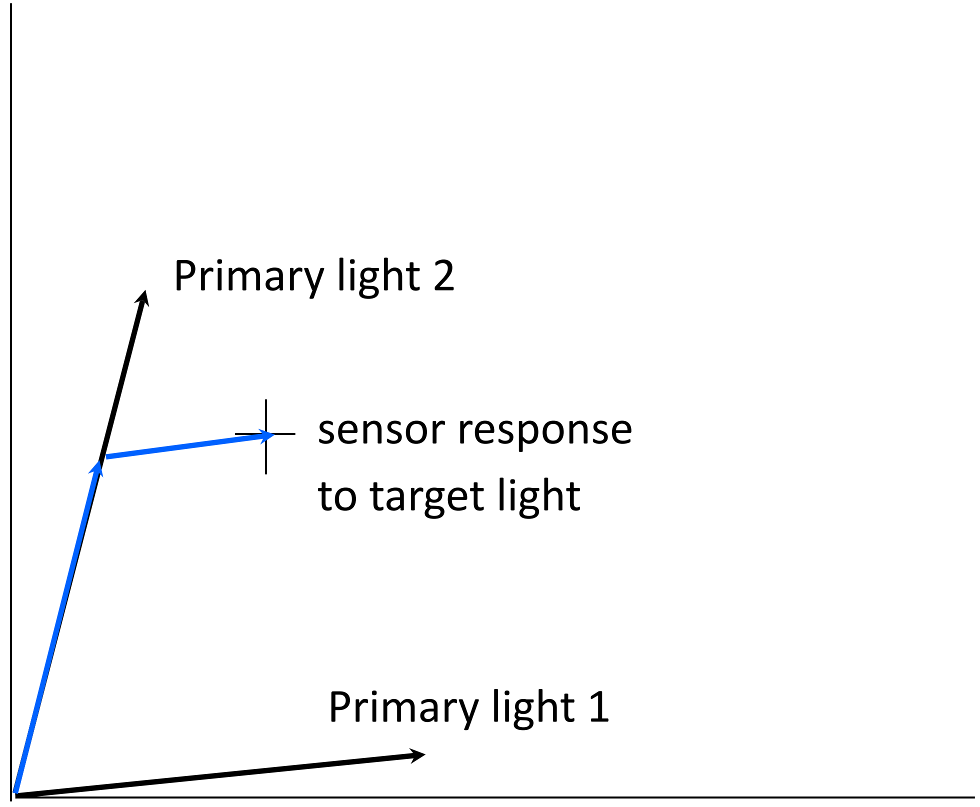

Color matching with a negative amount of primary 1

Match sensors' response to the target light plus some amount of primary light 1 to the response to some of primary light 2

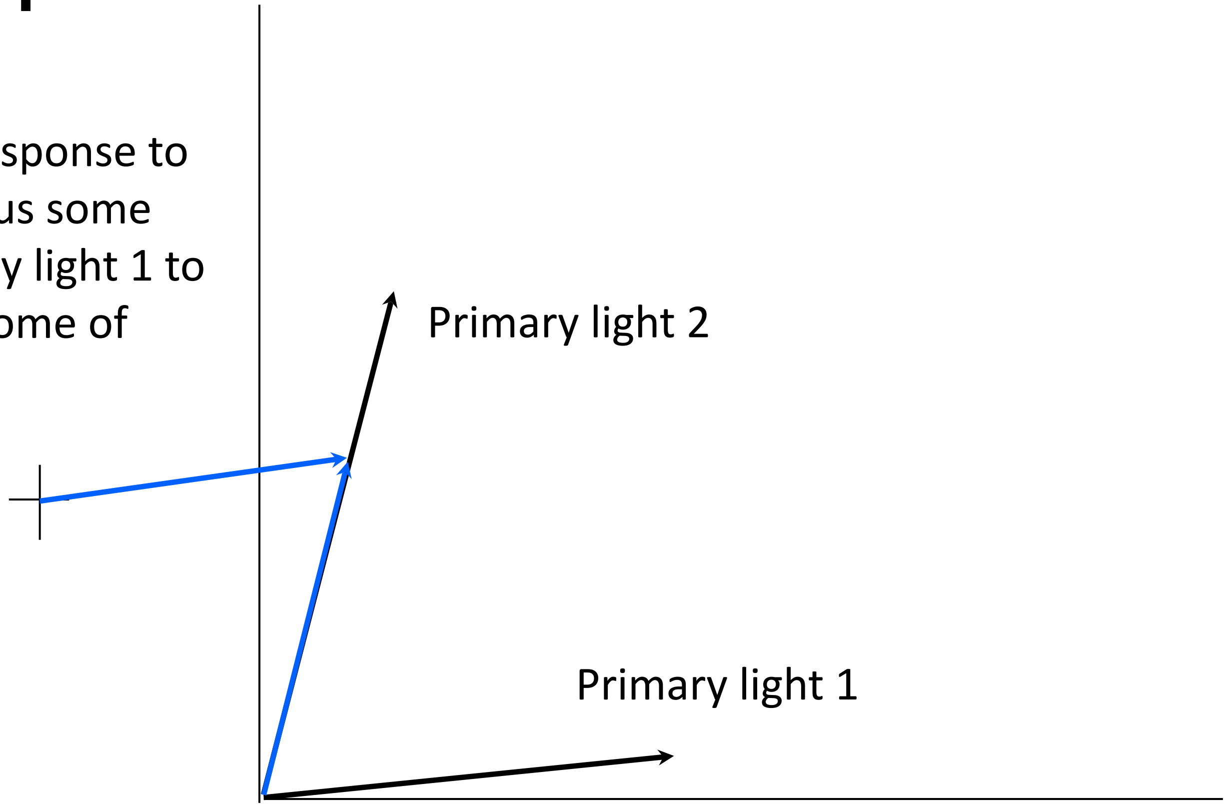

Color matching experiment--
handle negative light by adding light to the test.

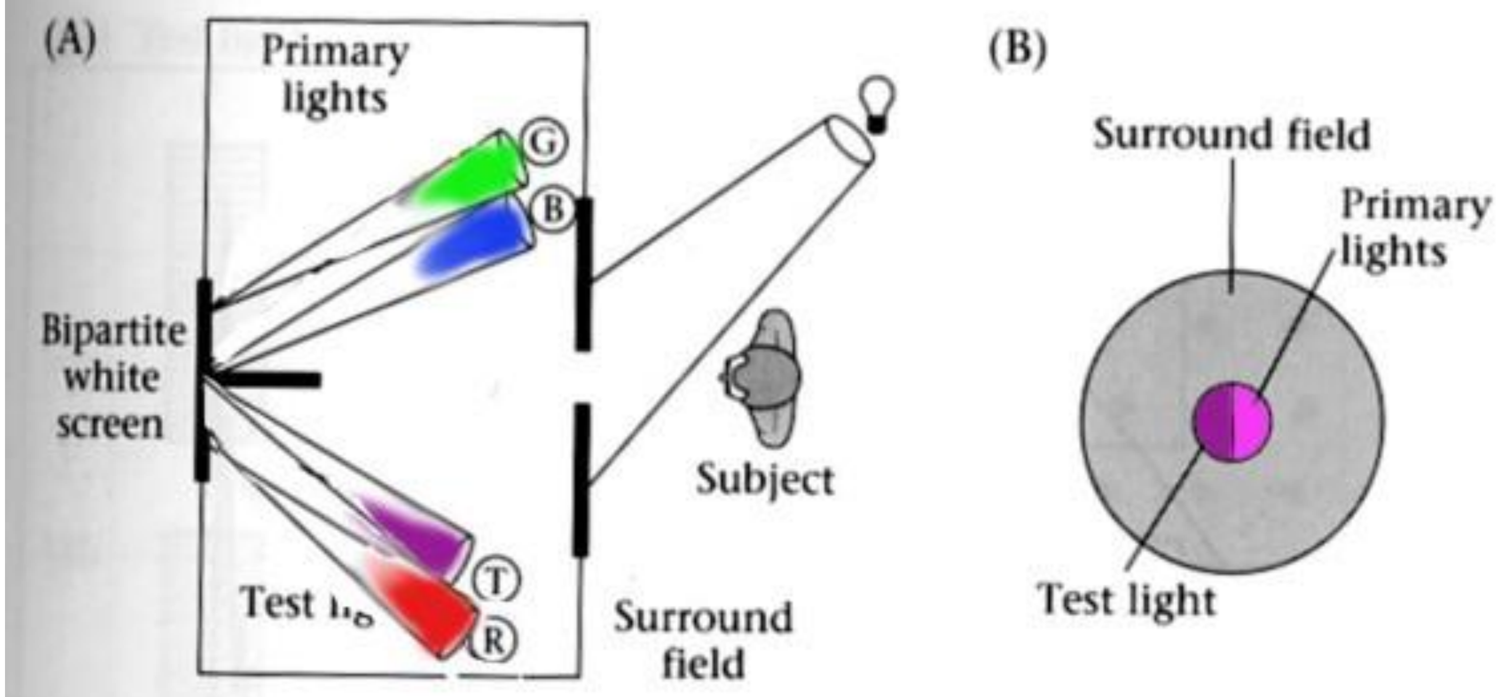

4.10 **THE COLOR-MATCHING EXPERIMENT.** The observer views a bipartite field and adjusts the intensities of the three primary lights to match the appearance of the test light. (A) A top view of the experimental apparatus. (B) The appearance of the stimuli to the observer. After Judd and Wyszecki, 1975.

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

Let's add some green light to the test color

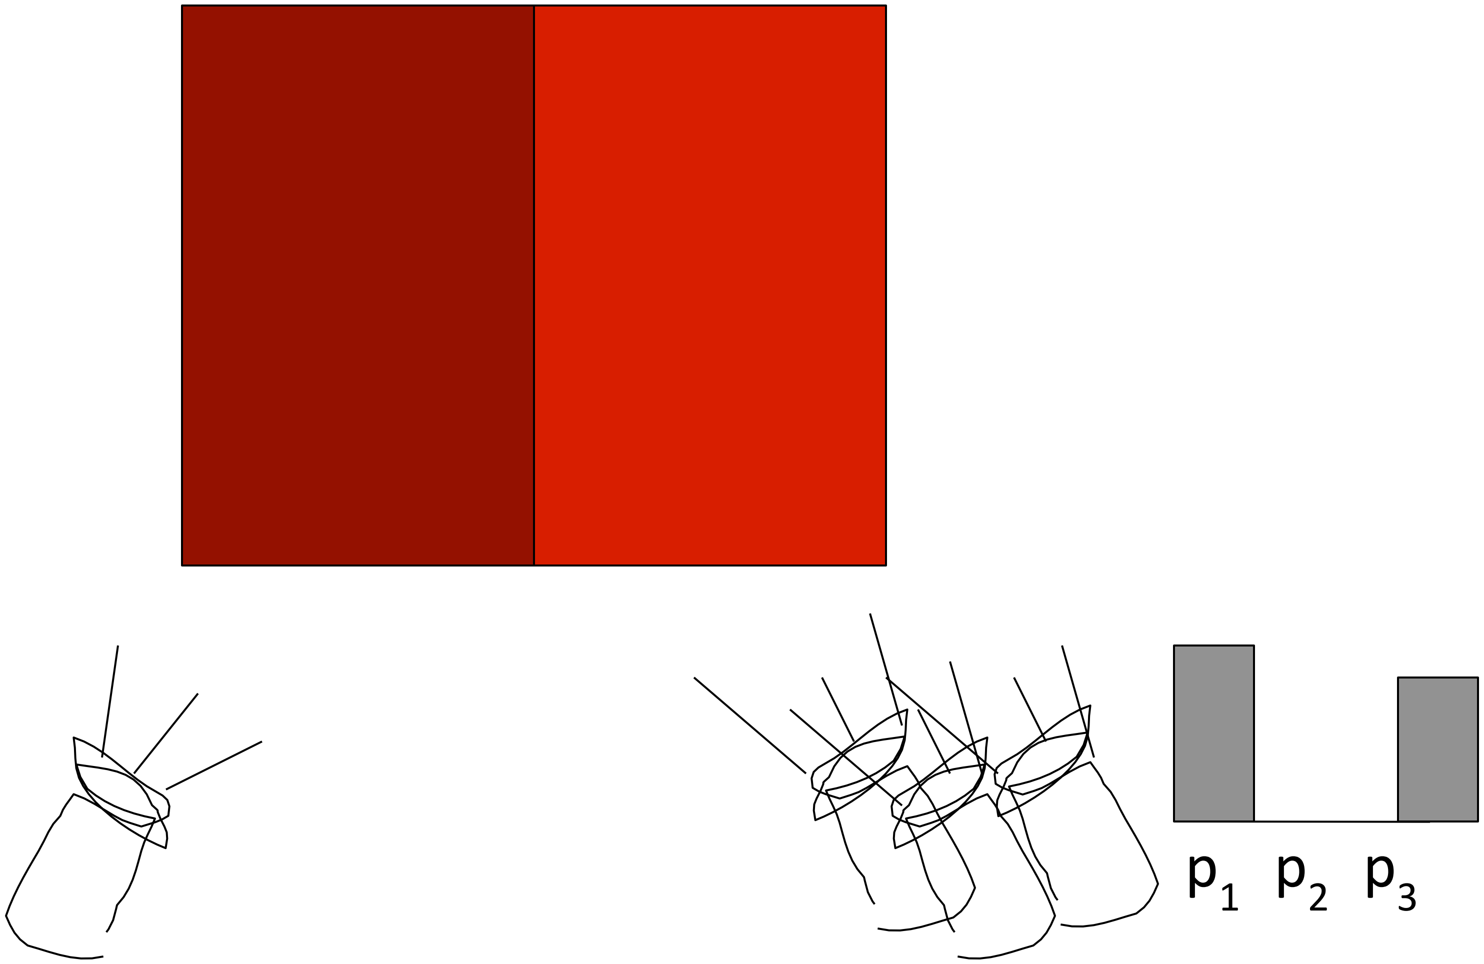

Color matching experiment 2

We say a “negative” amount of p_2 was needed to make the match, because we added it to the test color’s side.

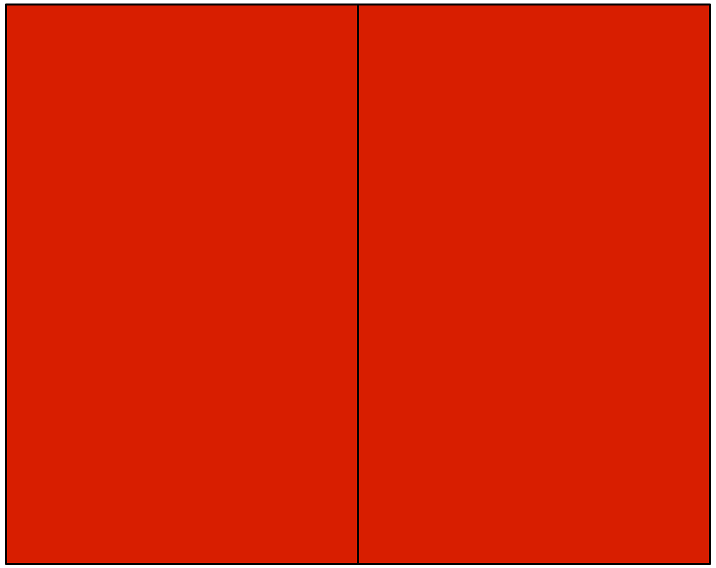

The primary color amounts needed for a match:

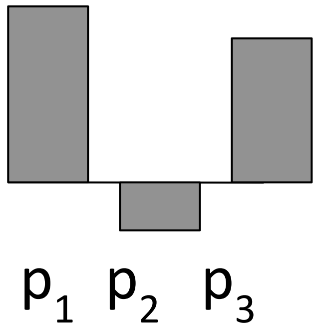

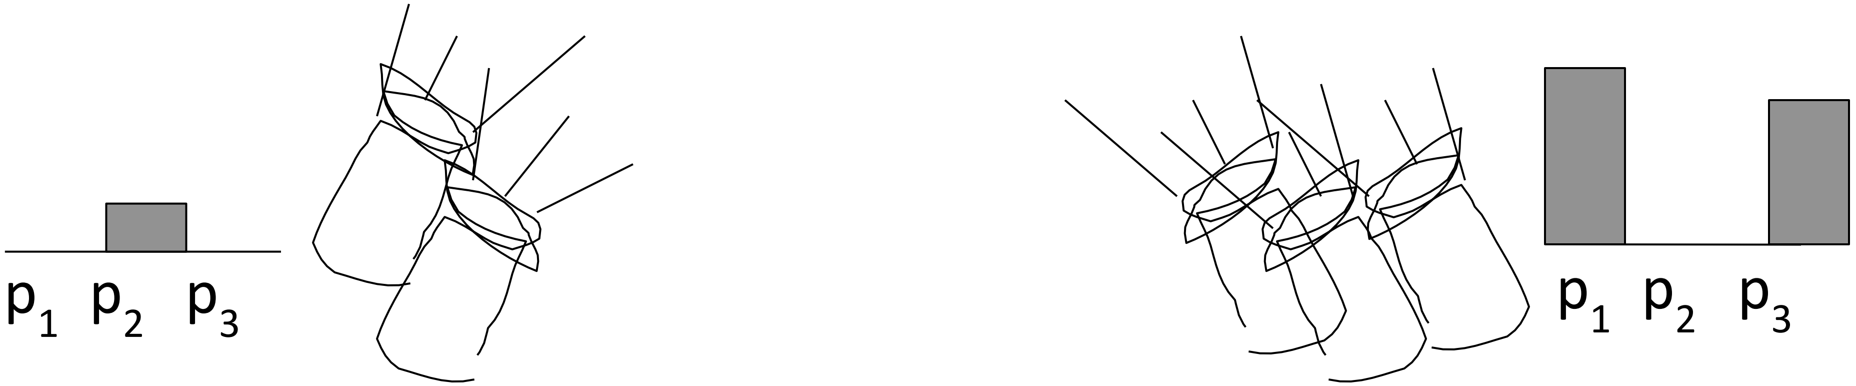

Color matching superposition (Grassman's laws)

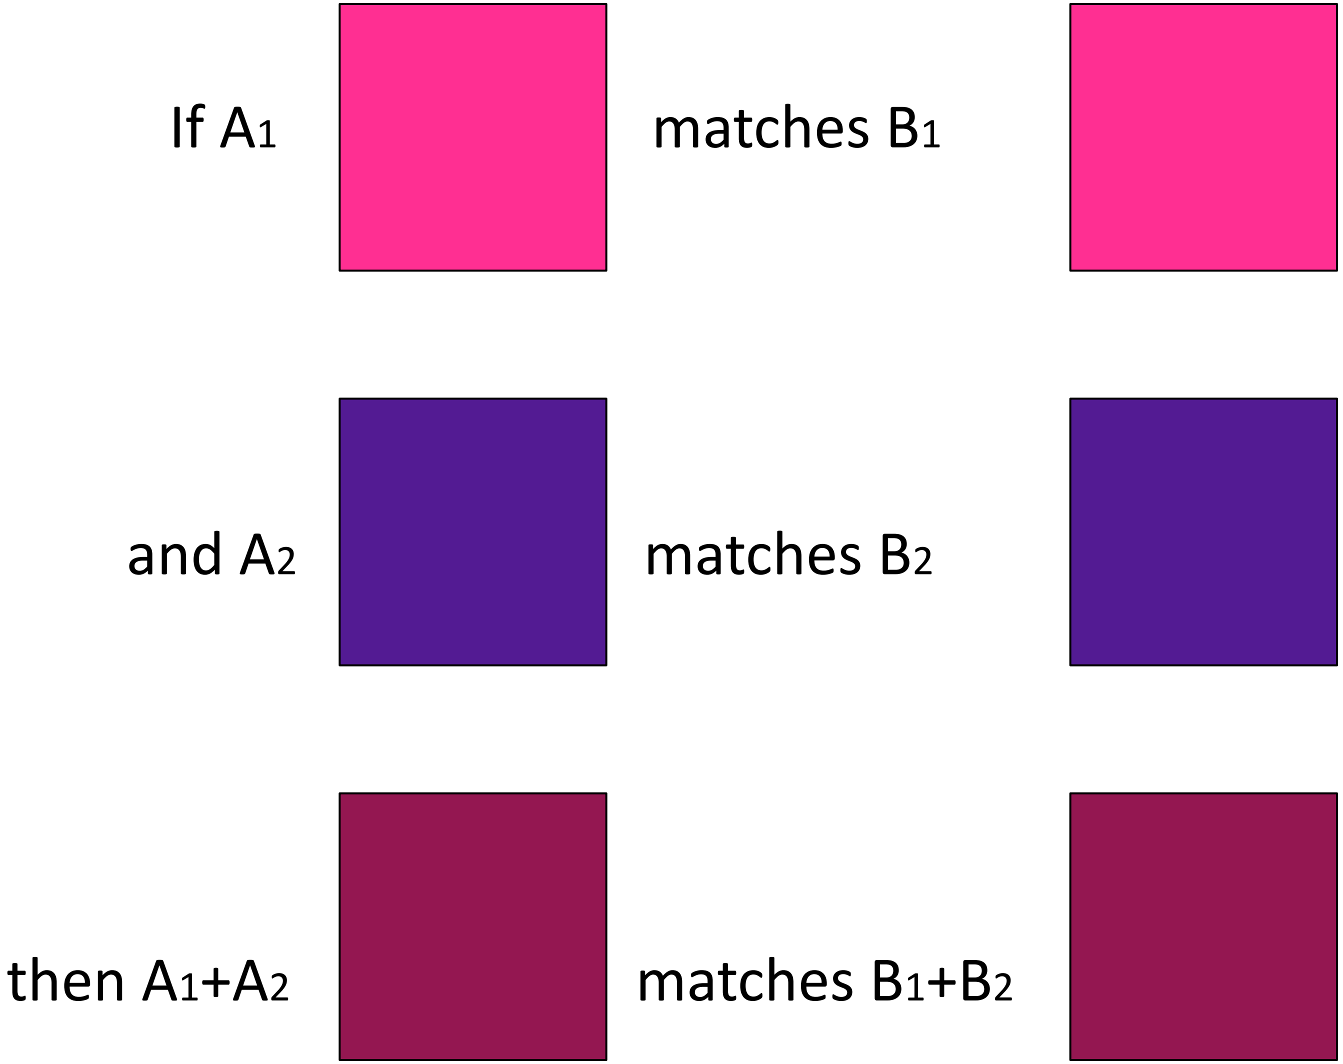

To reproduce a color

1. Choose a set of **3 primary colors** (three power spectra).
2. Determine how much of each **primary needs to be added** to a probe signal to match the test light.

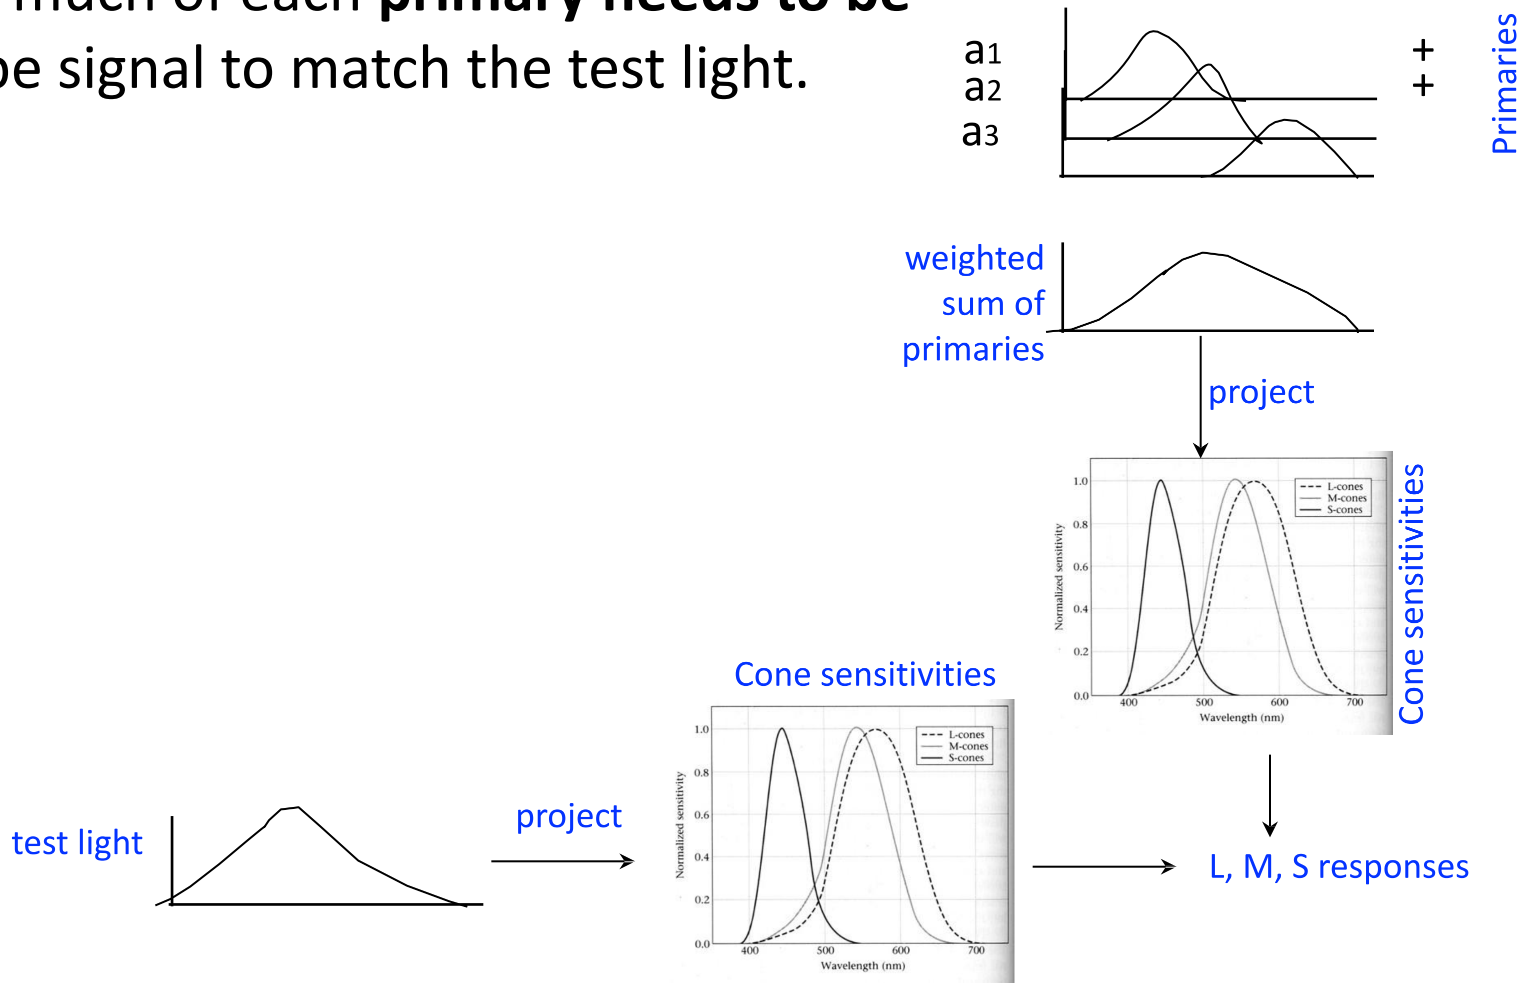

The assumption for color perception

- We know color appearance really depends on:
  - The illumination
  - Your eye’s adaptation level
  - The colors and scene interpretation surrounding the observed color.

- But for now we will assume that the spectrum of the light arriving at your eye completely determines the perceived color.

What is color?

- The result of interaction between physical light in the environment and our visual system.

- A *psychological property* of our visual experiences when we look at objects and lights, *not a physical property* of those objects or lights.

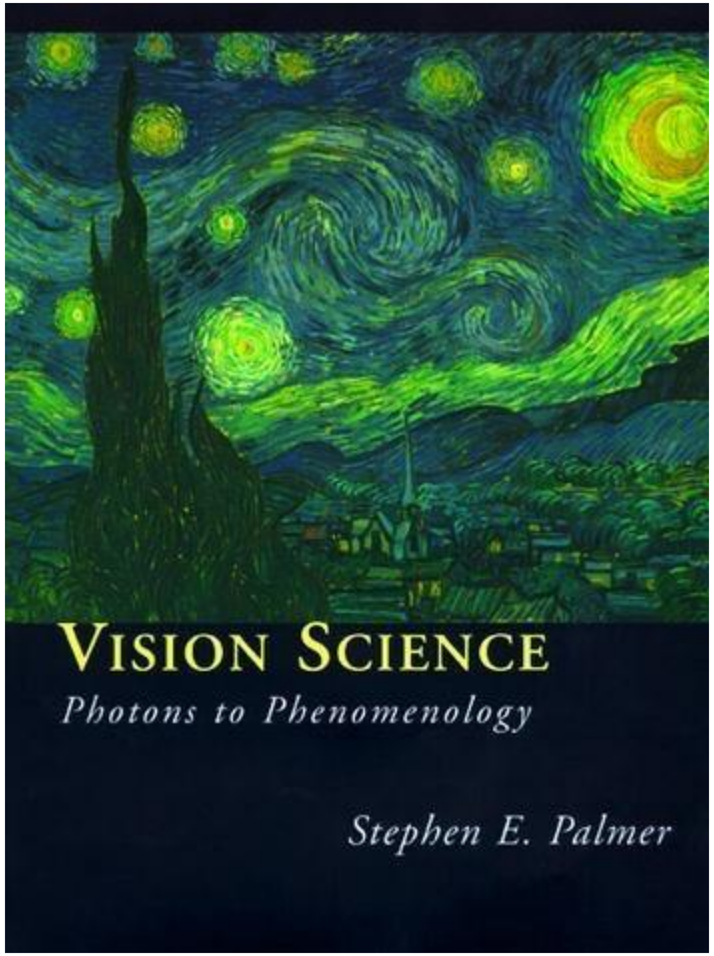

Slide credit: Lana Lazebnik

Image Formation

Perceived colors depend on: Light source, Surface properties (BRDF, reflectance), Light reaching the retina, spectral response of the retina (L-,M-, S-cones), Context/scene.

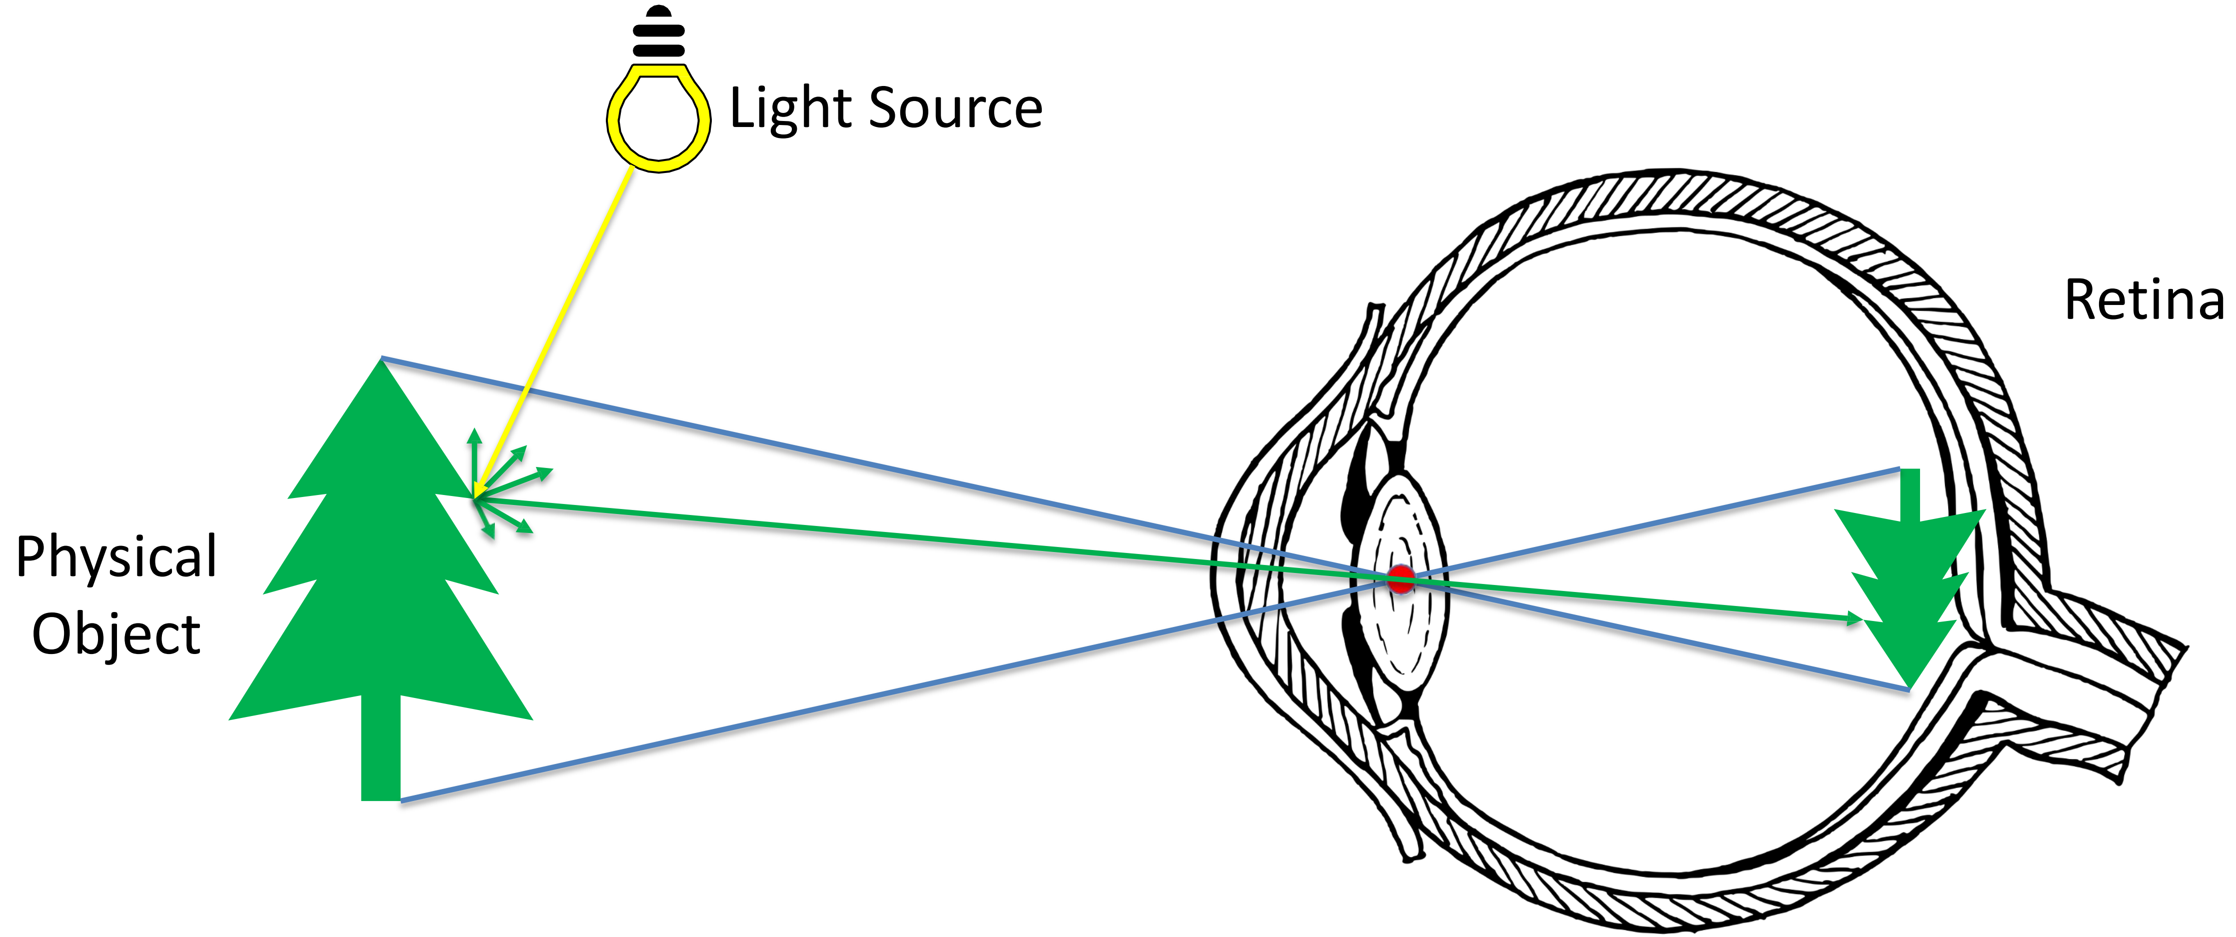

Image Formation

Acquired pixel colors will depend on: Light source, Surface properties (BRDF, reflectance), Light reaching the sensor, Spectral response of sensor.

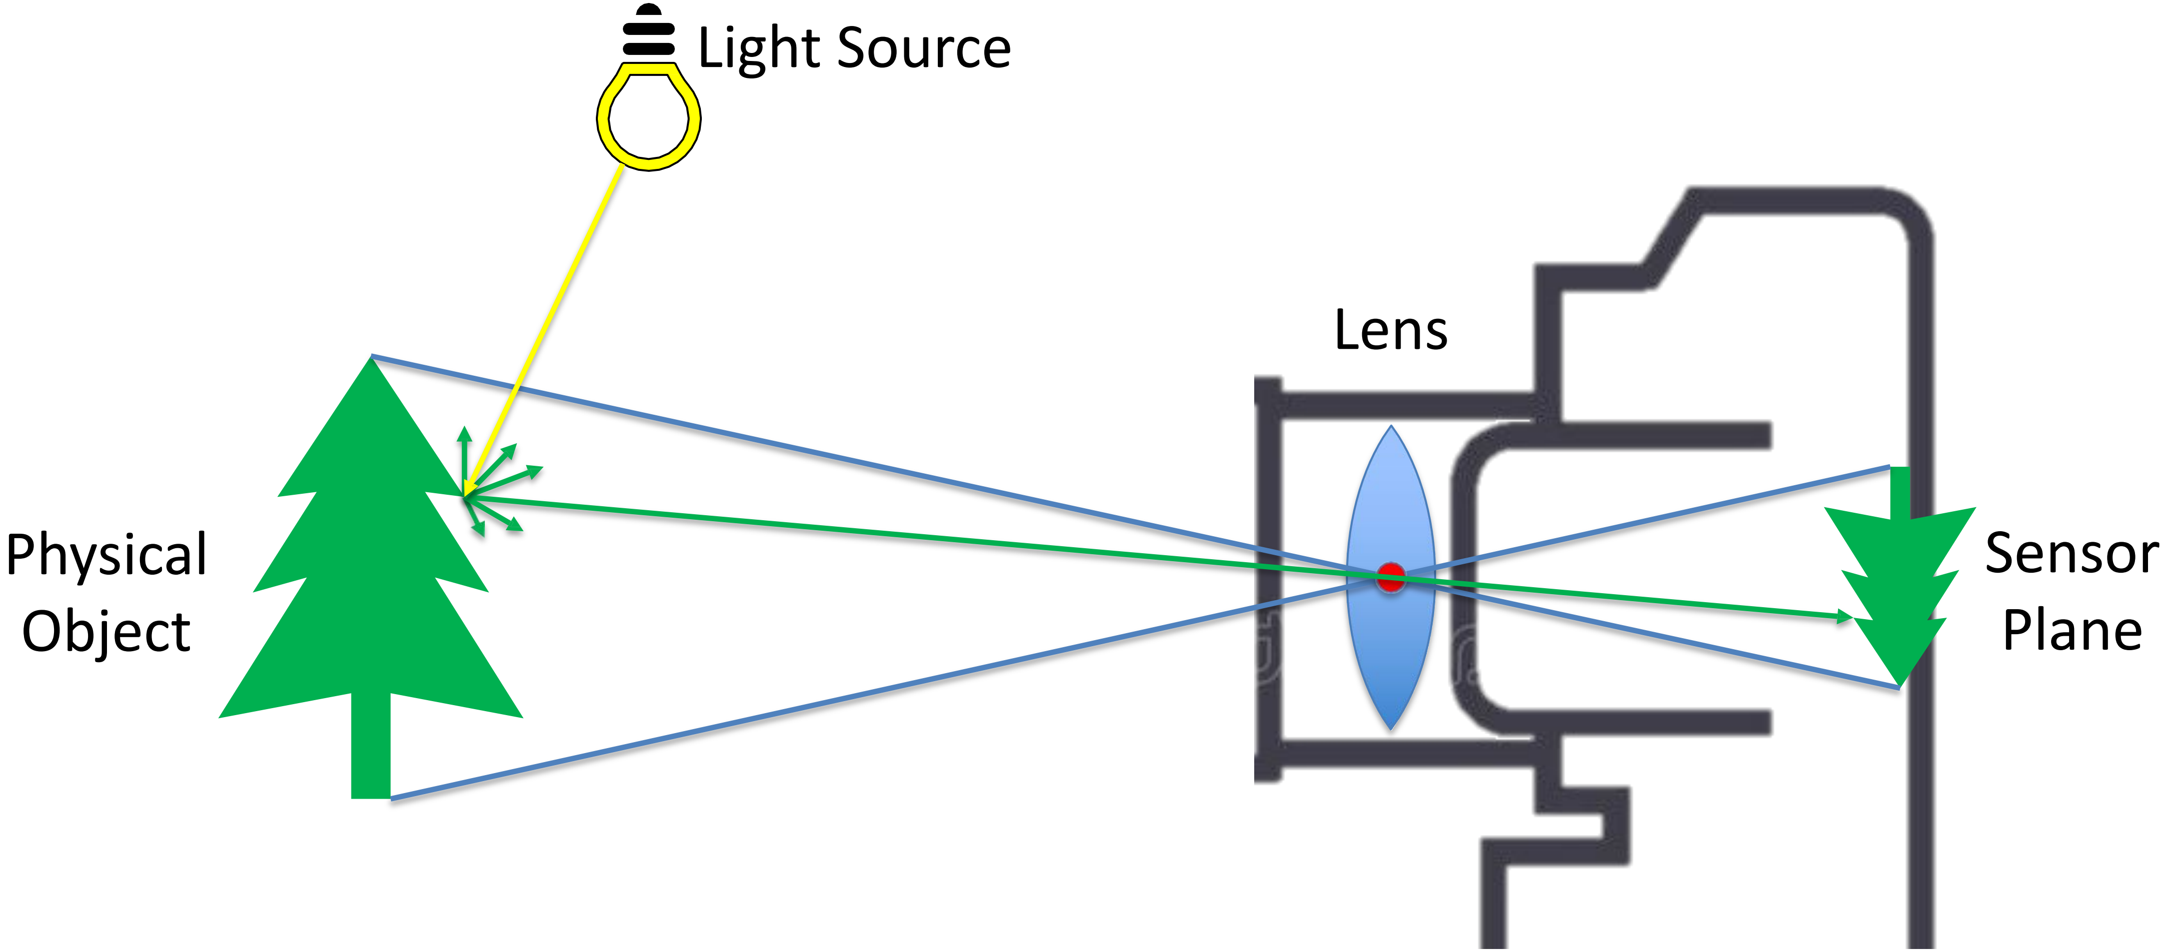

Today's agenda

- Image formation
- Physics of Color
- Color matching
- **Color spaces**
- Image sampling and quantization

Linear color spaces

- Defined by a choice of three *primaries*
- The coordinates of a color are given by the weights of the primaries used to match it

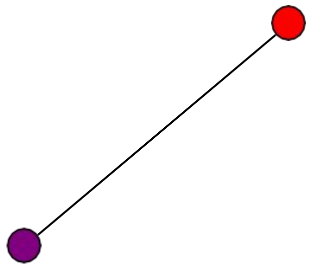

mixing two lights produces colors that lie along a straight line in color space

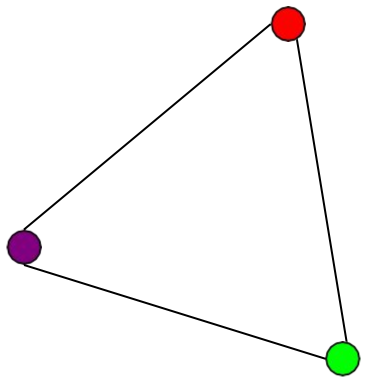

mixing three lights produces colors that lie within the triangle they define in color space

How to compute the weights of the primaries to match any spectral signal

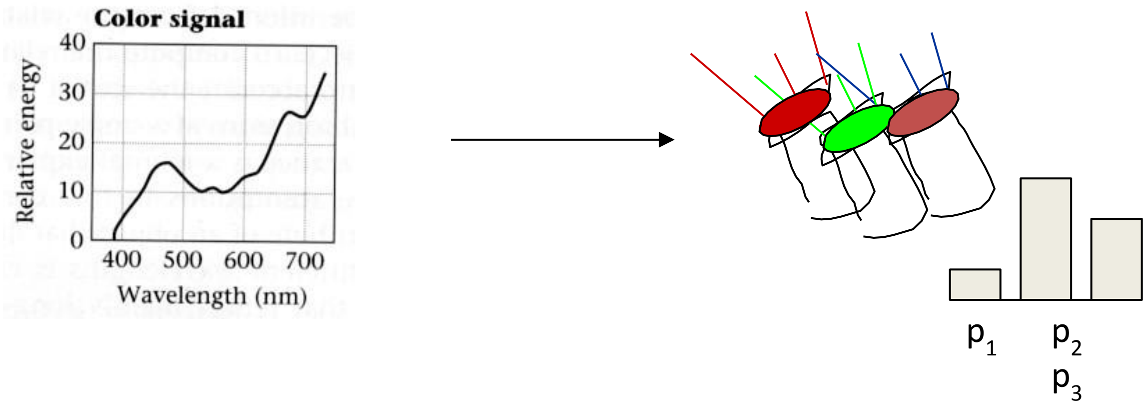

Matching functions: the amount of each primary needed to match a monochromatic light source at each wavelength

RGB space

Primaries are monochromatic lights (for monitors, they correspond to the three types of phosphors)

Subtractive matching required for some wavelengths

RGB primaries

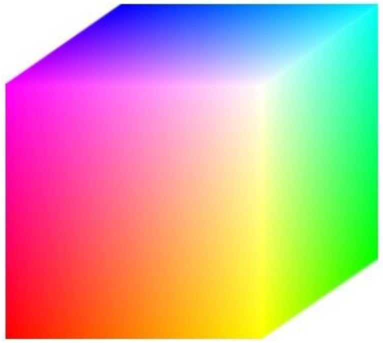

$p_1 = 645.2$ nm
$p_2 = 525.3$ nm
$p_3 = 444.4$ nm

RGB matching functions

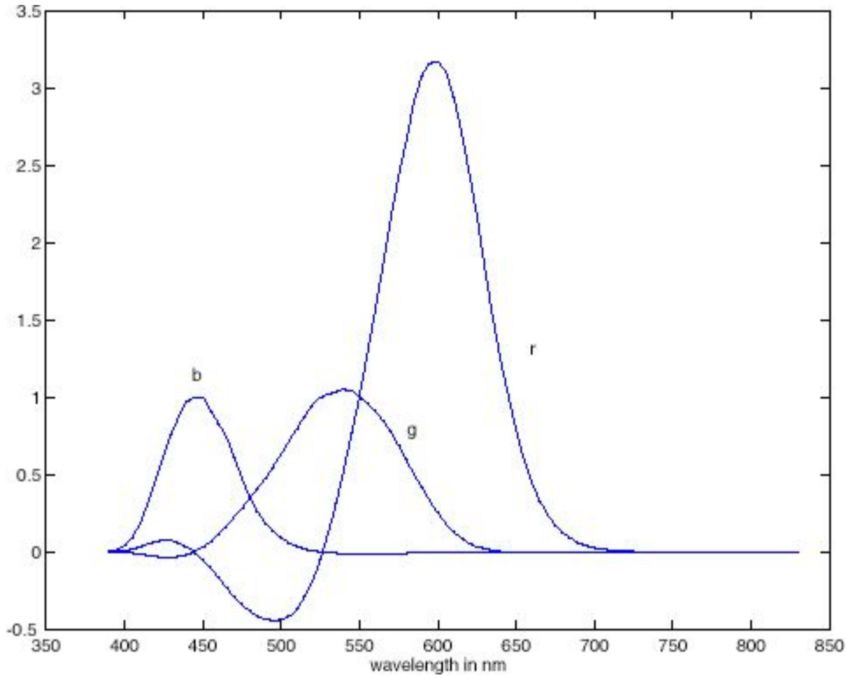

http://en.wikipedia.org/wiki/CIE_1931_color_space

Linear color spaces: CIE XYZ

- Primaries are imaginary, but matching functions are everywhere positive
- The Y parameter corresponds to brightness or *luminance* of a color
- 2D visualization: draw (x,y), where $x = X/(X+Y+Z)$, $y = Y/(X+Y+Z)$

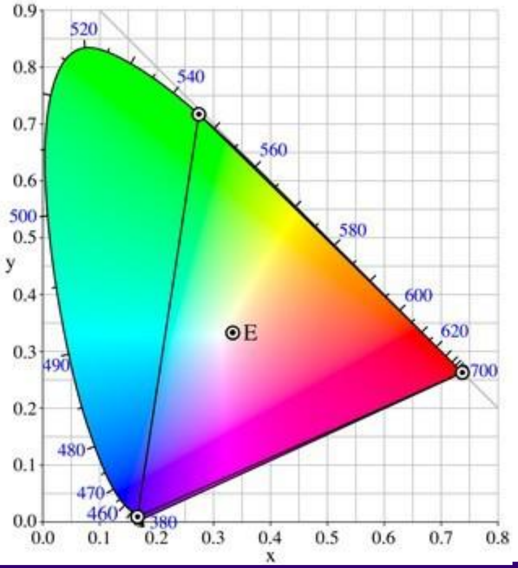

Matching functions

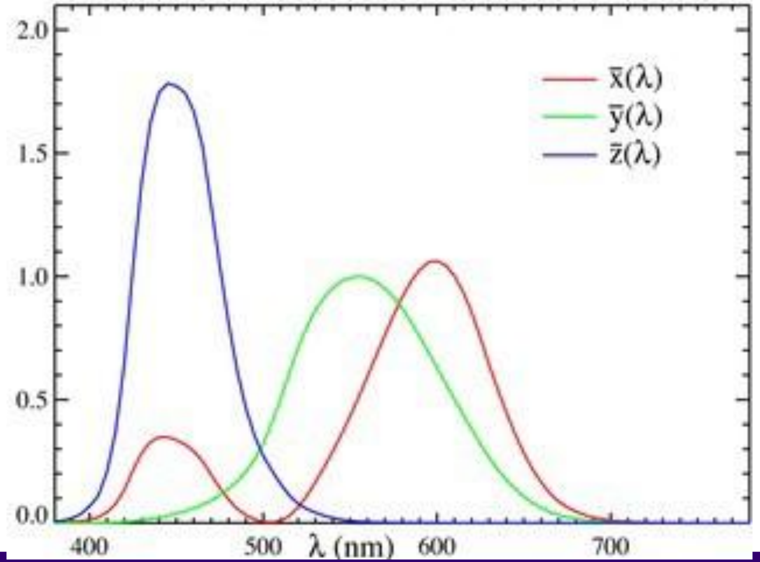

Nonlinear color spaces: HSV

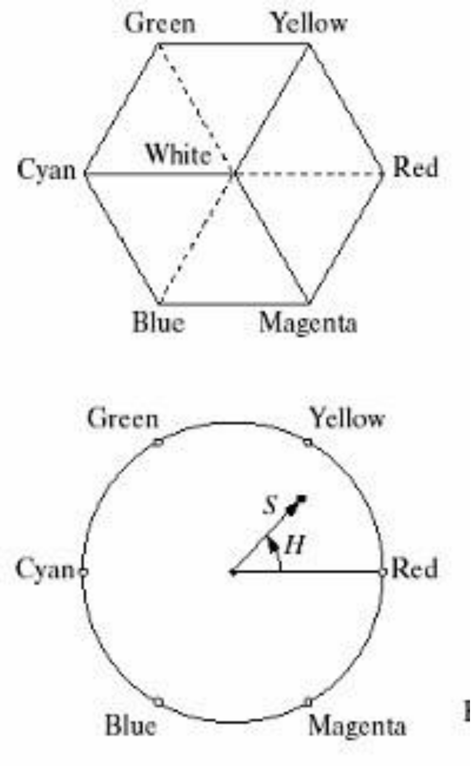

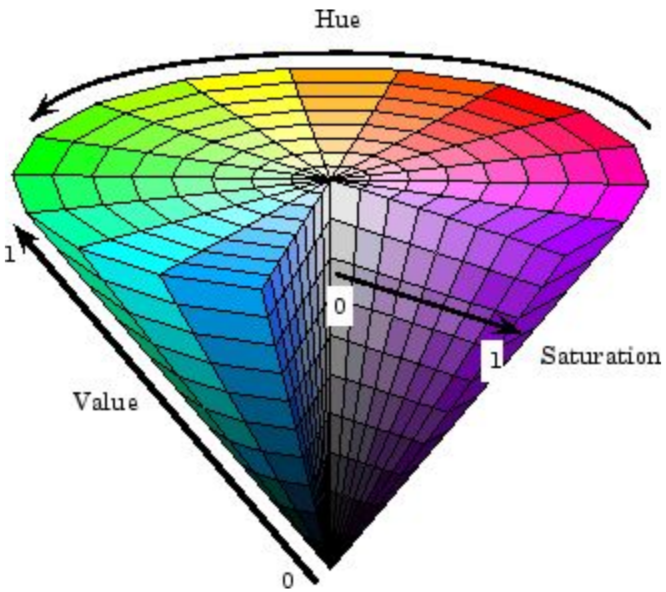

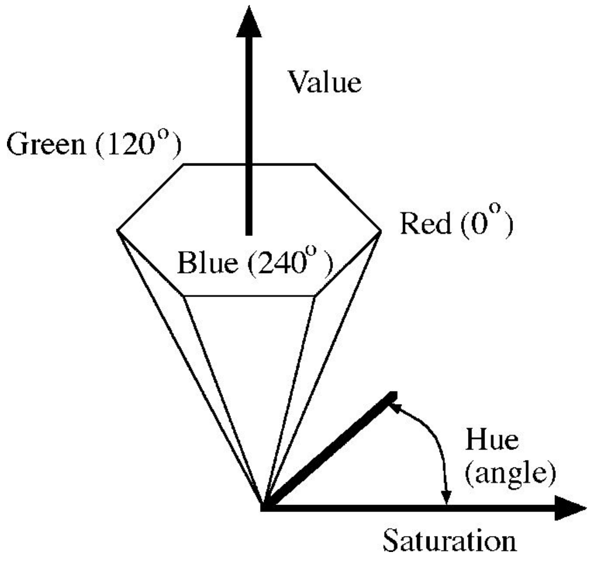

- Perceptually meaningful dimensions: Hue, Saturation, Value (Intensity)

Nonlinear color spaces: HSV

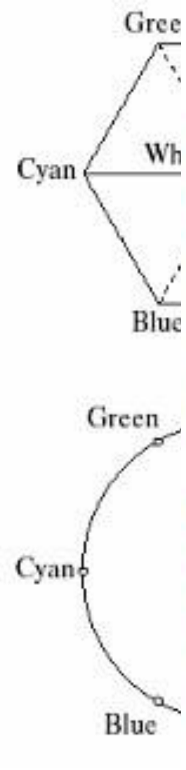

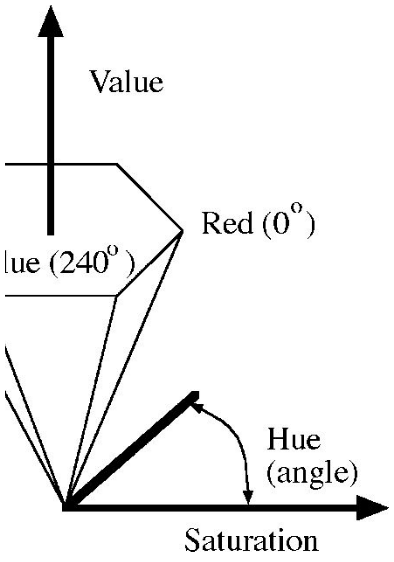

- Perceptually me … tensity)

Today's agenda

- Image formation
- Physics of Color
- Color matching
- Color spaces
- **Image sampling and quantization**

Image Formation

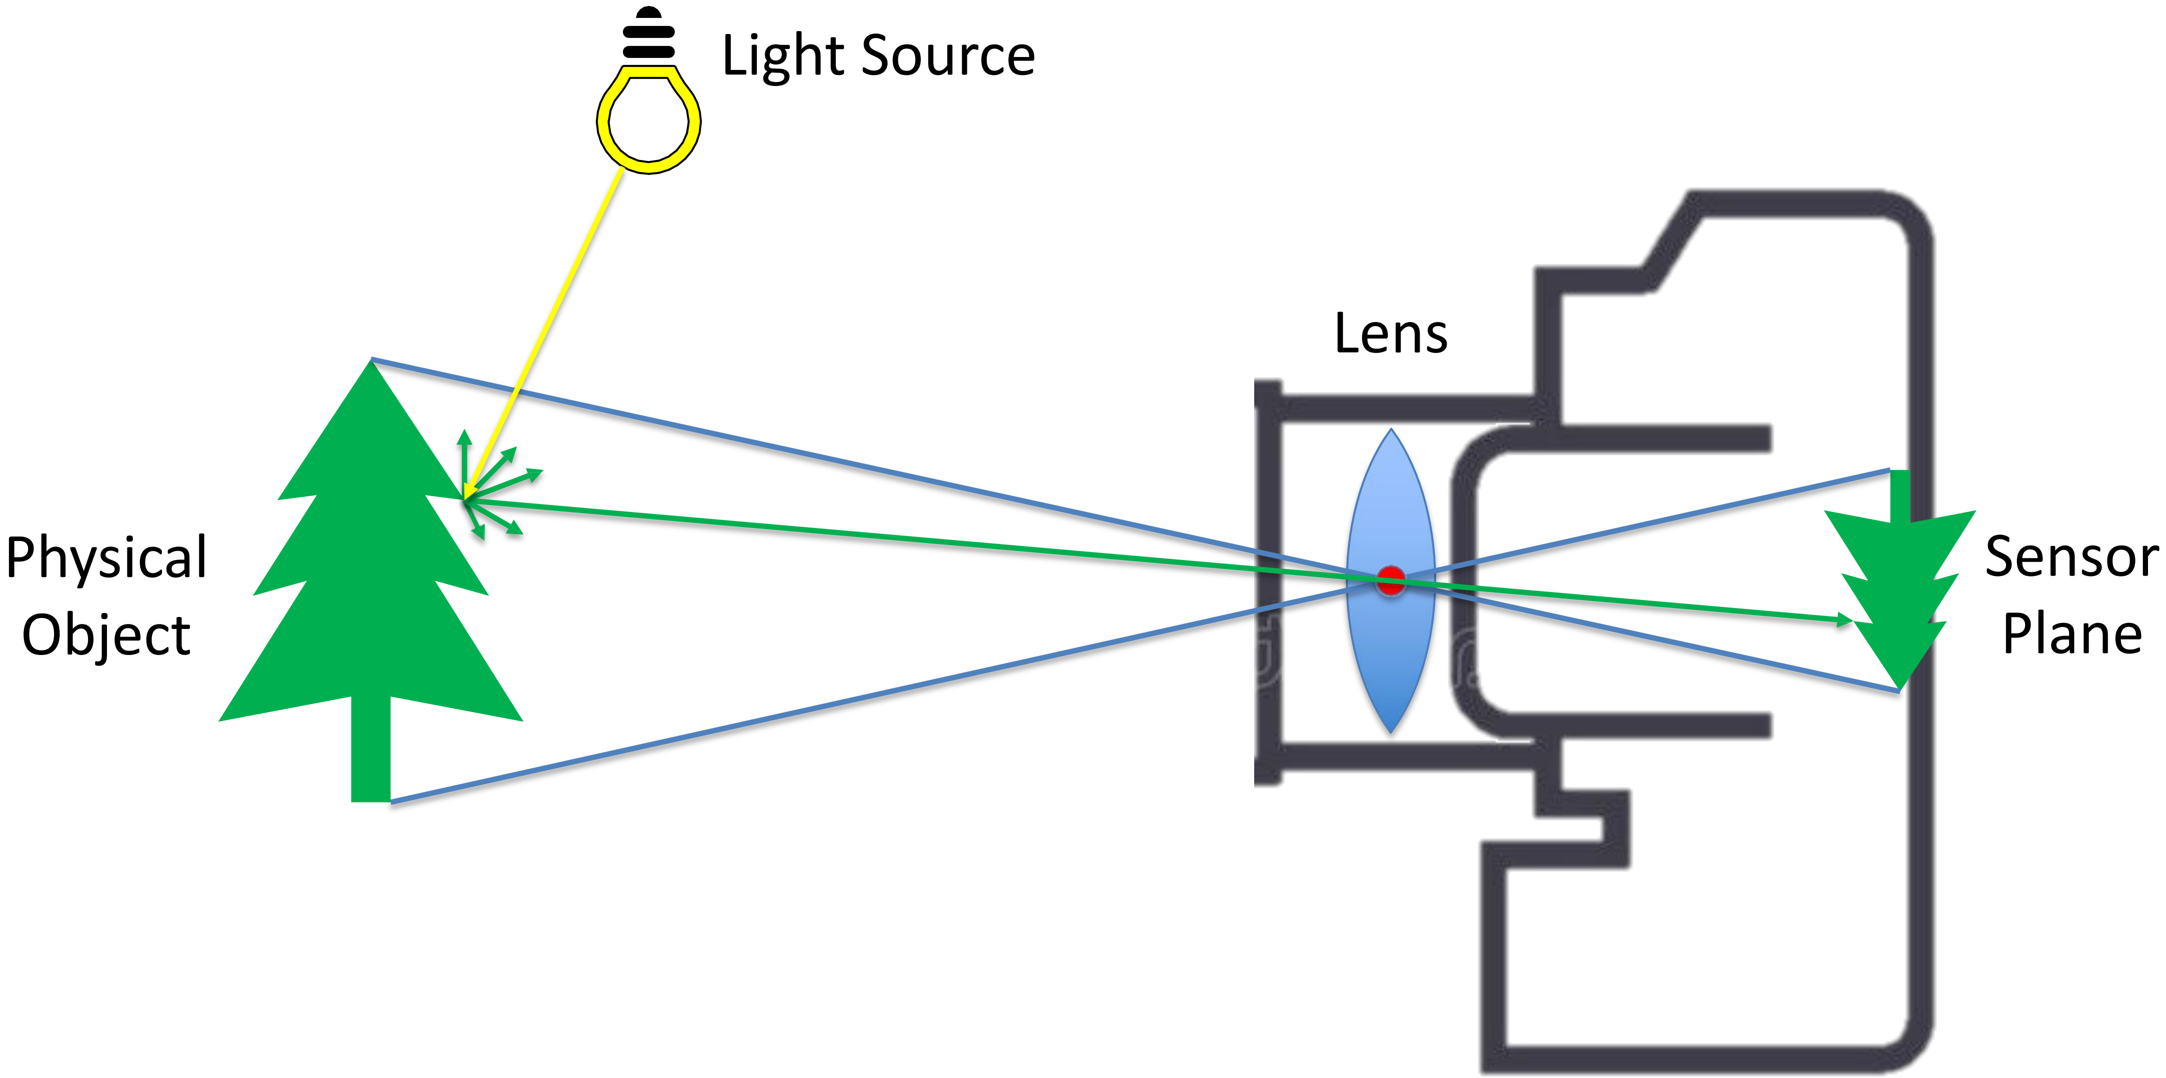

Camera Sensor produce discrete outputs

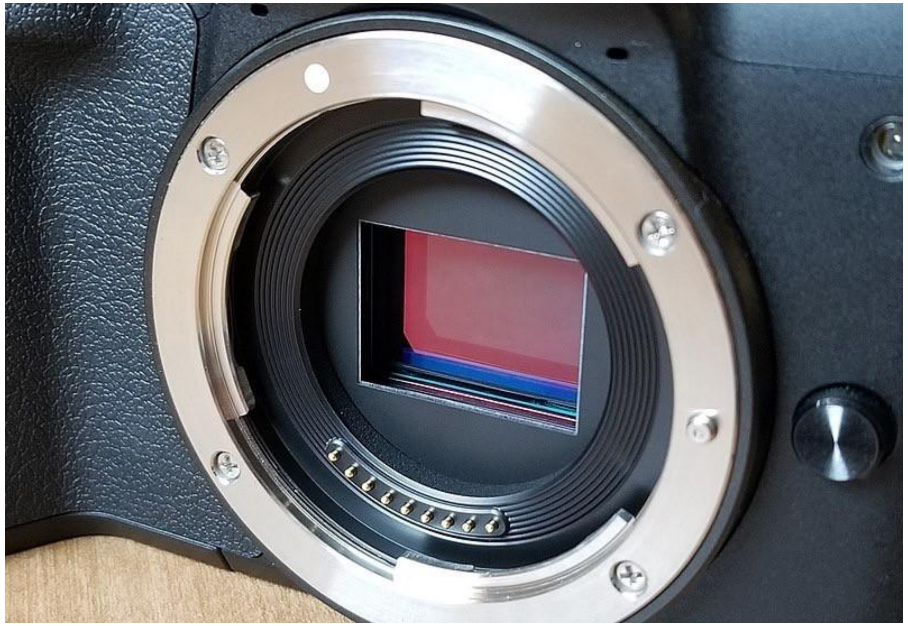

https://commons.wikimedia.org/wiki/File:Mirrorless_Camera_Sensor.jpg

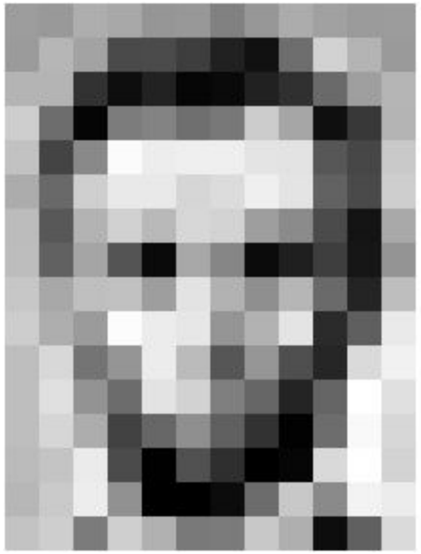

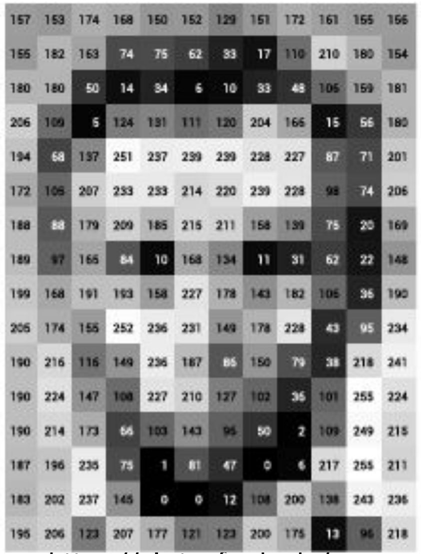

157	153	174	168	150	152	129	151	172	161	155	156
155	182	163	74	75	62	33	17	110	210	180	154
180	180	50	14	34	6	10	33	48	106	159	181
206	109	5	124	131	111	120	204	166	15	56	180
194	68	137	251	237	239	239	228	227	87	71	201
172	105	207	233	233	214	220	239	228	98	74	206
188	88	179	209	185	215	211	158	139	75	20	169
189	97	165	84	10	168	134	11	31	62	22	148
199	168	191	193	158	227	178	143	182	106	36	190
205	174	155	252	236	231	149	178	228	43	95	234
190	216	116	149	236	187	86	150	79	38	218	241
190	224	147	108	227	210	127	102	36	101	255	224
190	214	173	66	103	143	96	50	2	109	249	215
187	196	235	75	1	81	47	0	6	217	255	211
183	202	237	145	0	0	12	108	200	138	243	236
195	206	123	207	177	121	123	200	175	13	96	218

https://ai.stanford.edu/~syyeung/cvweb/Pictures1/imagematrix.png

Types of Images

Binary

Grayscale

Color

Binary image representation

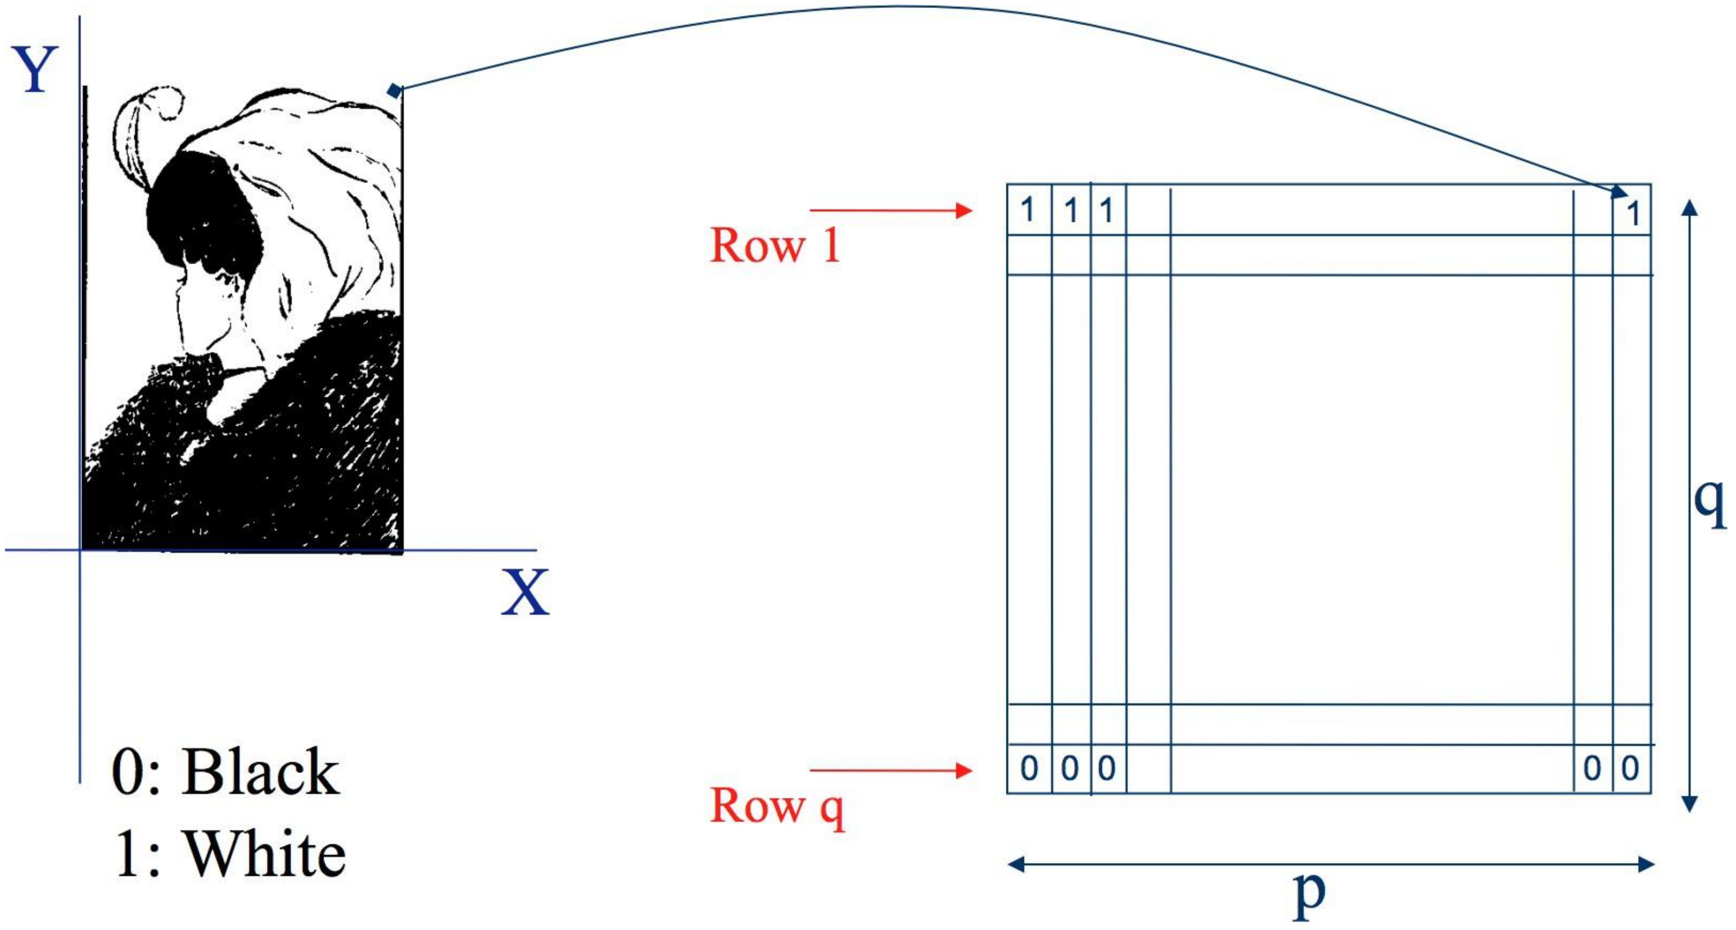

Grayscale image representation

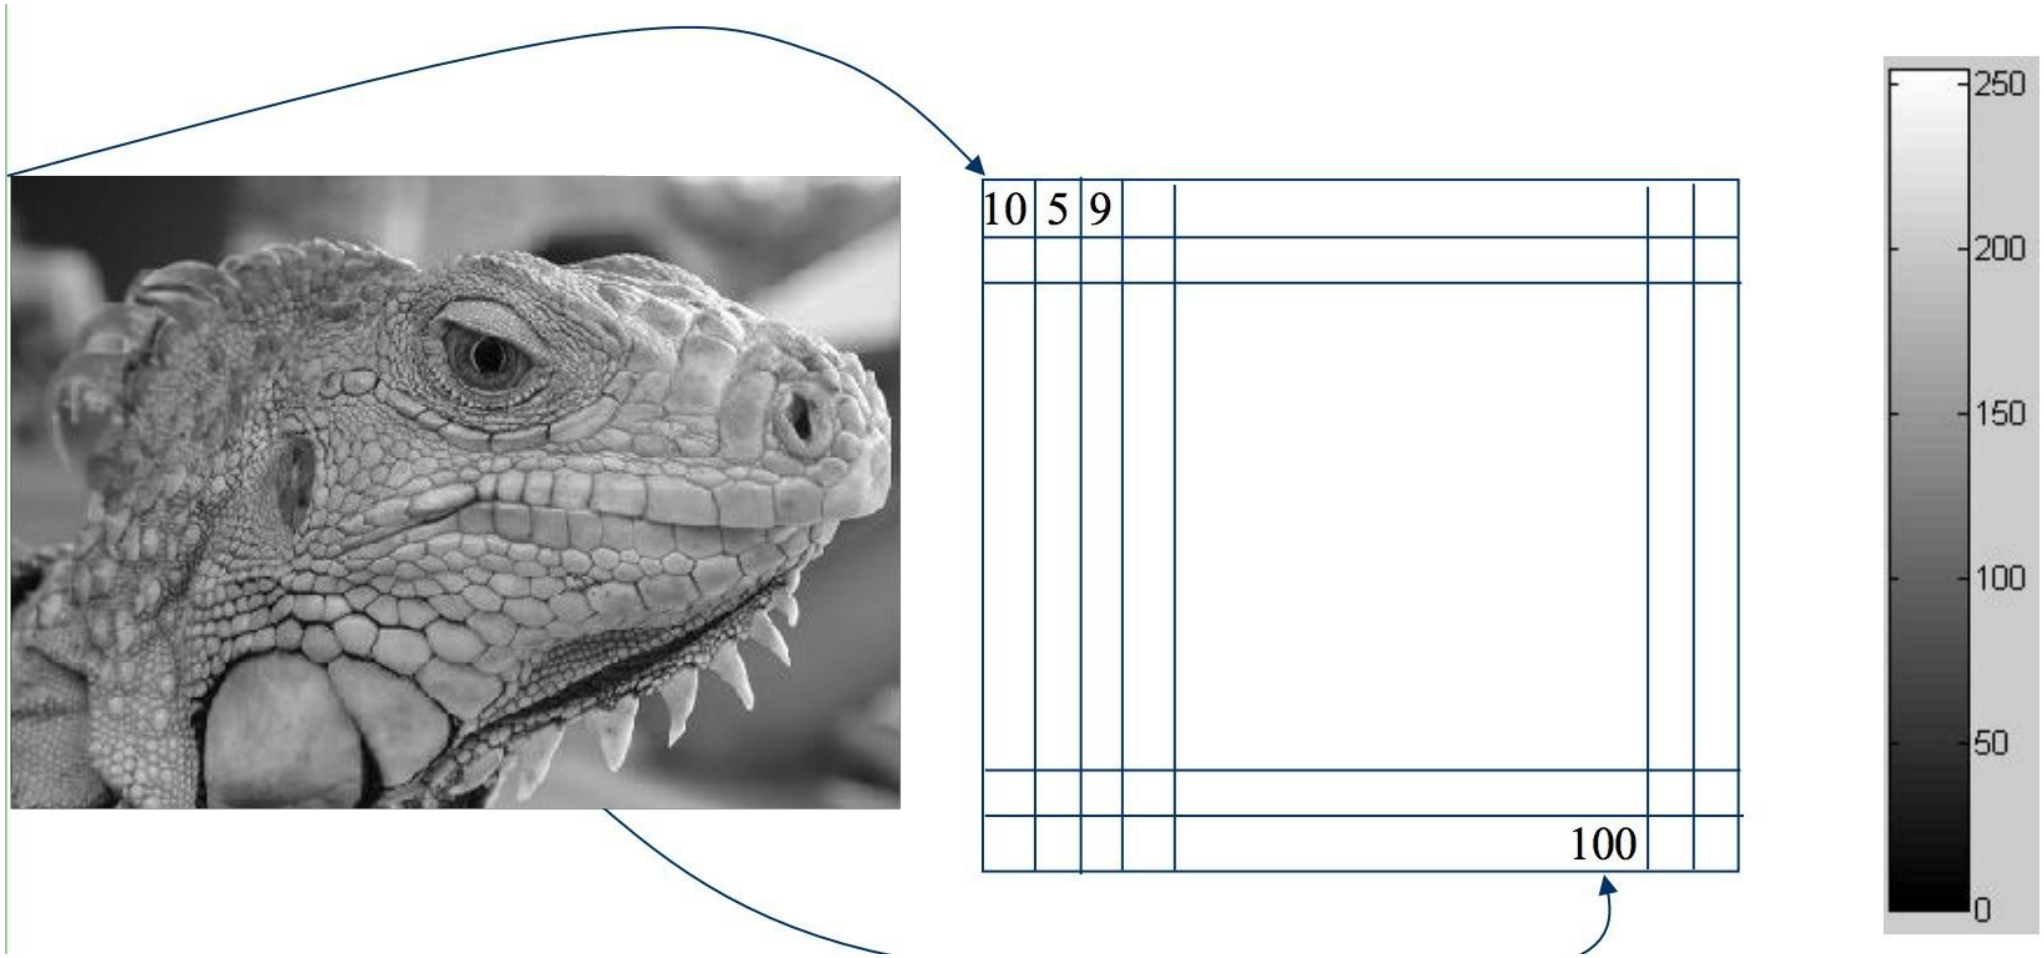

Grayscale image representation

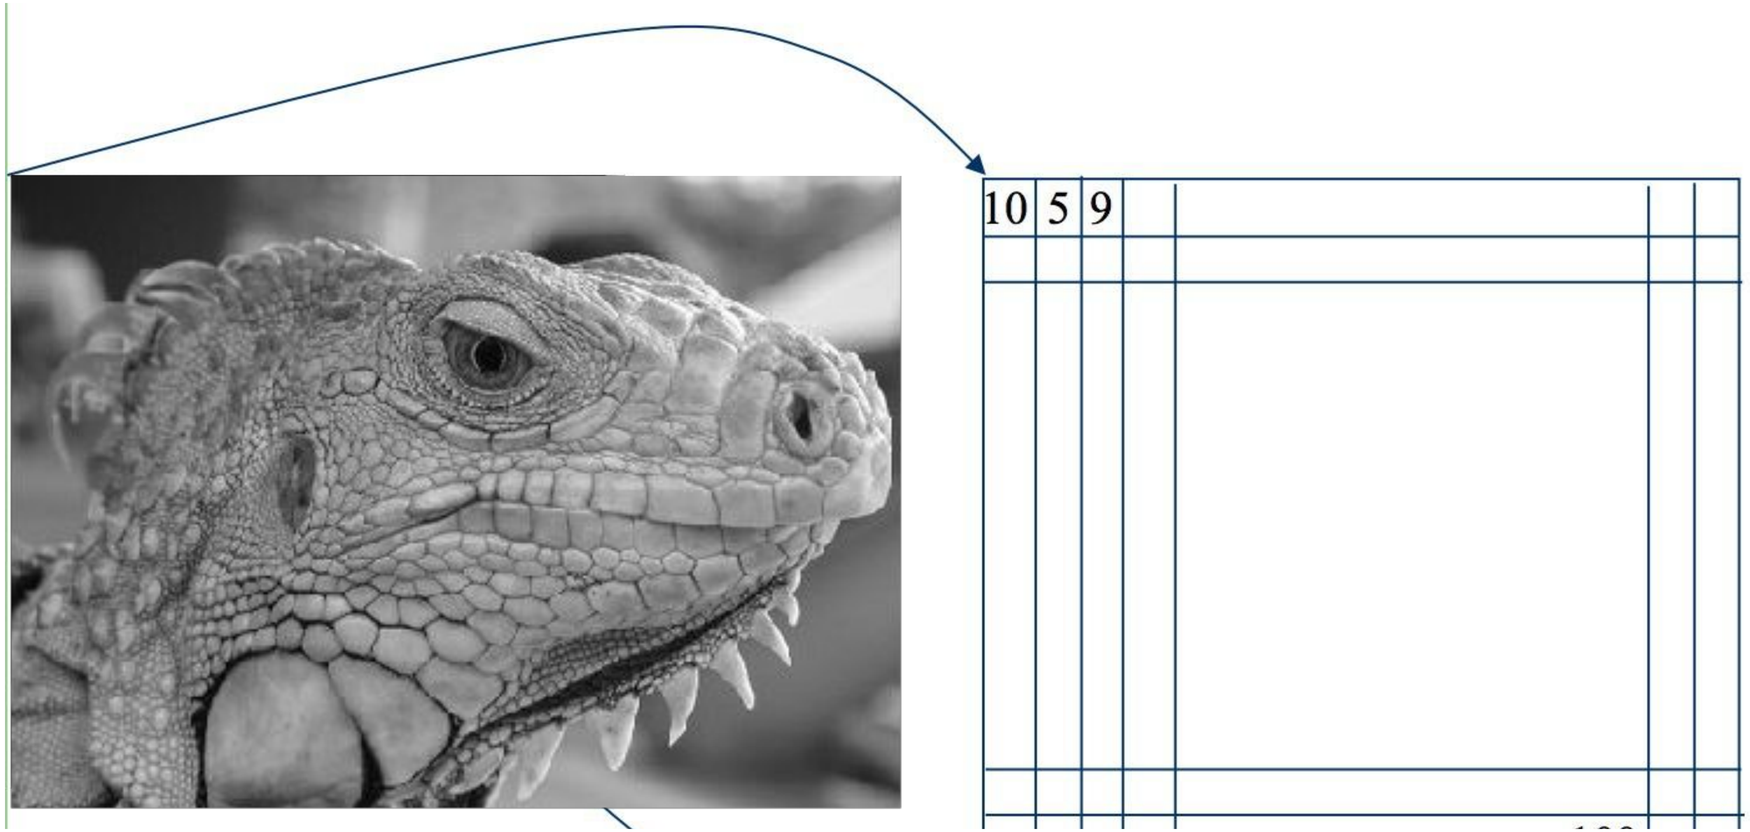

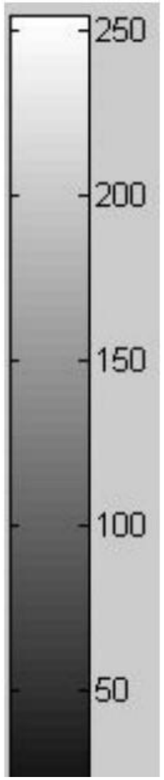

Q. If you used HSV to represent grayscale images, is the slider representing hue? Or saturation? Or value?

Color image representation

B channel

G channel

R channel

Color image - one channel

R channel

Types of Images

Binary

Grayscale

Color

Phil Noble / AP

Digital Images are sampled

What happens when we zoom into the images we capture?

Errors due to Sampling

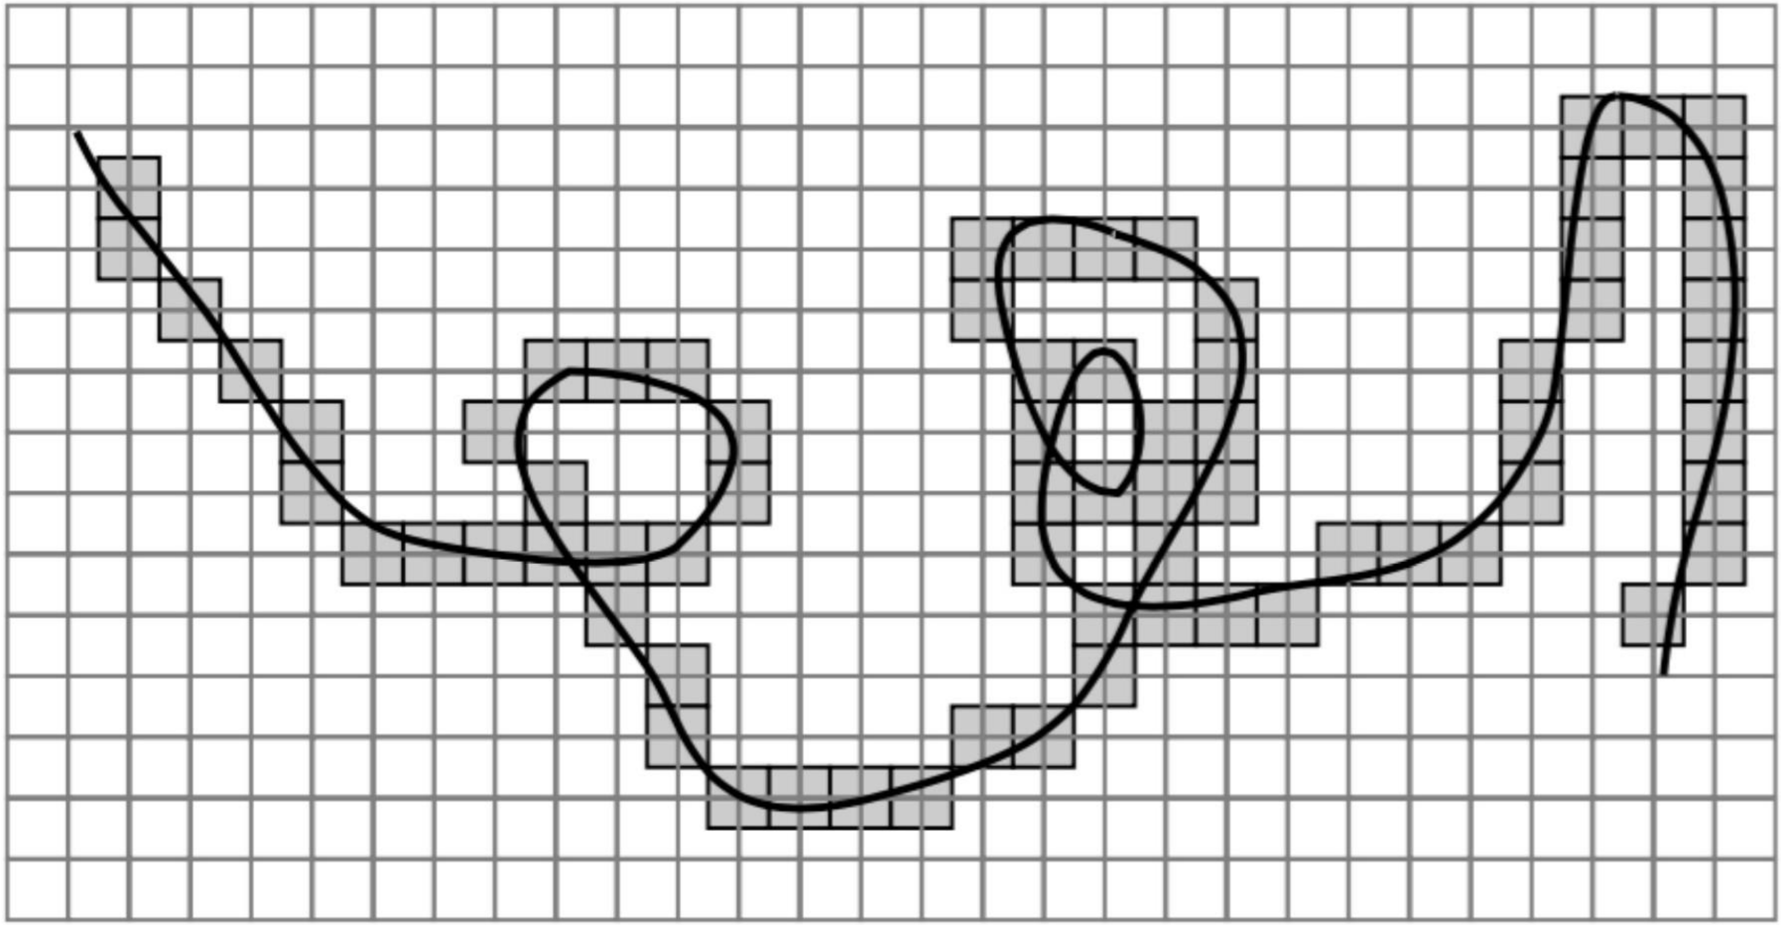

Resolution

is a **sampling** parameter, defined in dots per inch (DPI) or equivalent measures of spatial pixel density

Images are Sampled and Quantized

- An image contains discrete number of pixels
  - Pixel value:
    - "grayscale"
      (or "intensity"): [0,255]

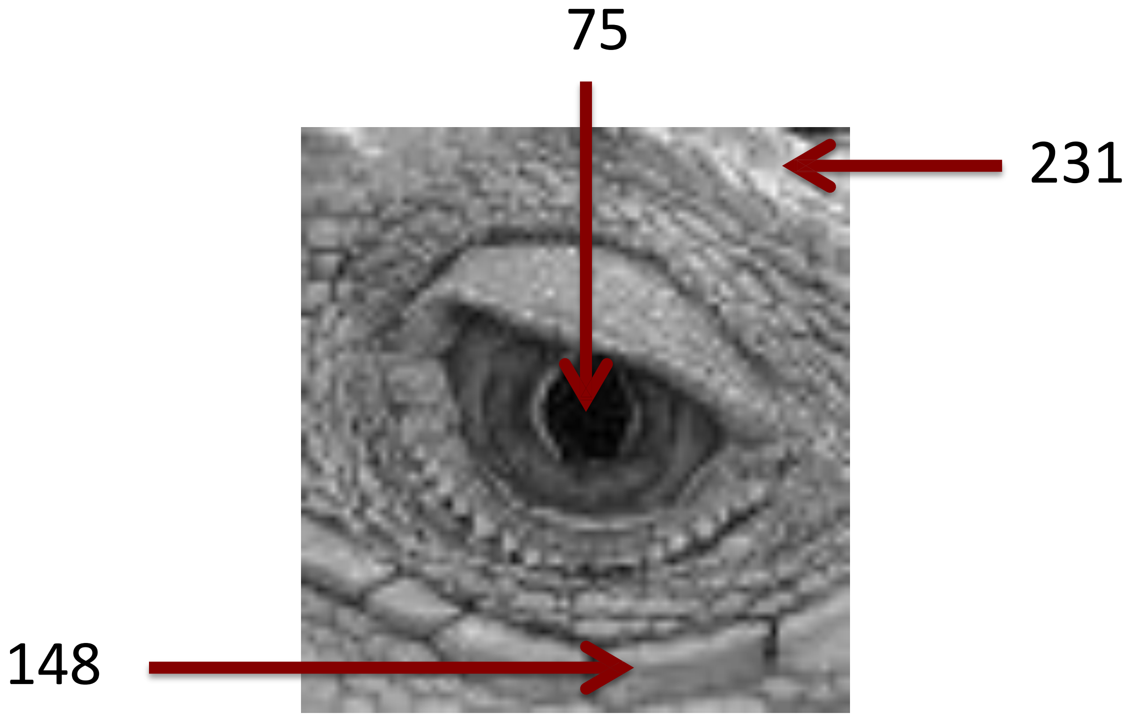

Images are Sampled and Quantized

- An image contains discrete number of pixels
  - –Pixel value:
    - •“grayscale”
      (or “intensity”): [0,255]
    - •“color”
      - –RGB: [R, G, B]

With this loss of information (from sampling and quantization),

Can we still use images for useful tasks?

Summary

- Image formation
- Physics of Color
- Color matching
- Color spaces
- Image sampling and quantization