THE UNIVERSITY OF BRITISH COLUMBIA

CPSC 425: Computer Vision

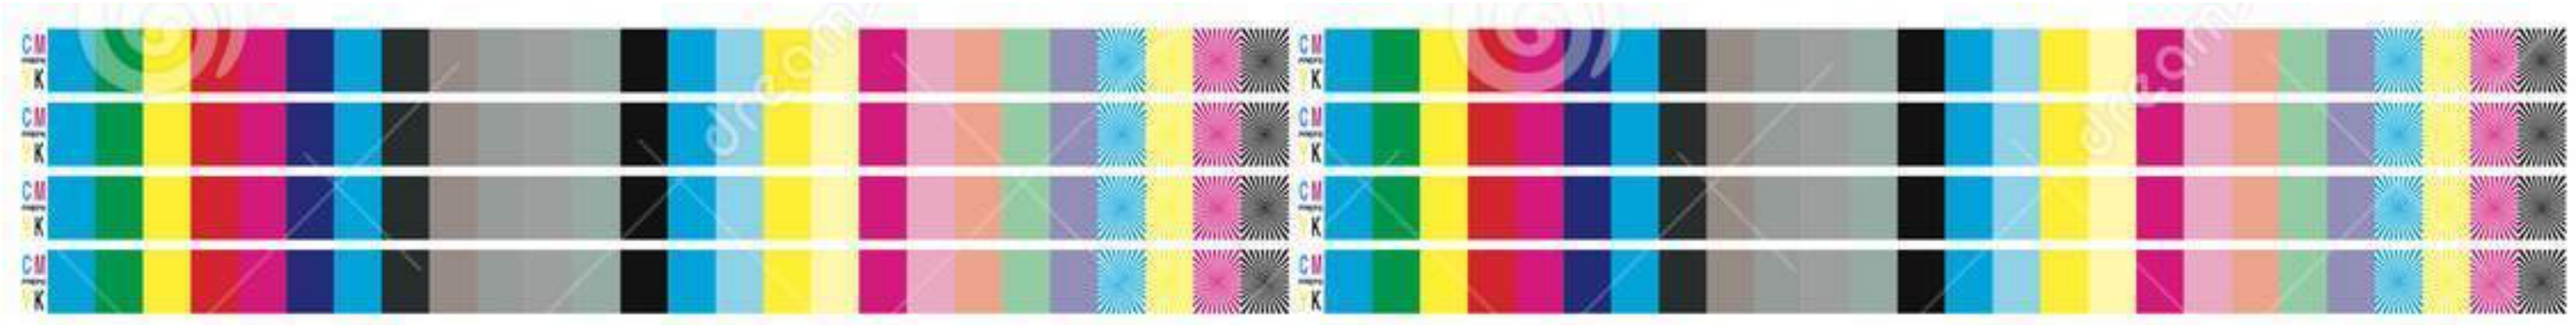

Lecture 12: Color

(unless otherwise stated slides are taken or adopted from **Bob Woodham, Jim Little** and **Fred Tung**)

Overview: Image Formation, Cameras and Lenses

The **image formation process** that produces a particular image depends on

- **Lightening** condition
- Scene **geometry**
- **Surface** properties
- Camera **optics**

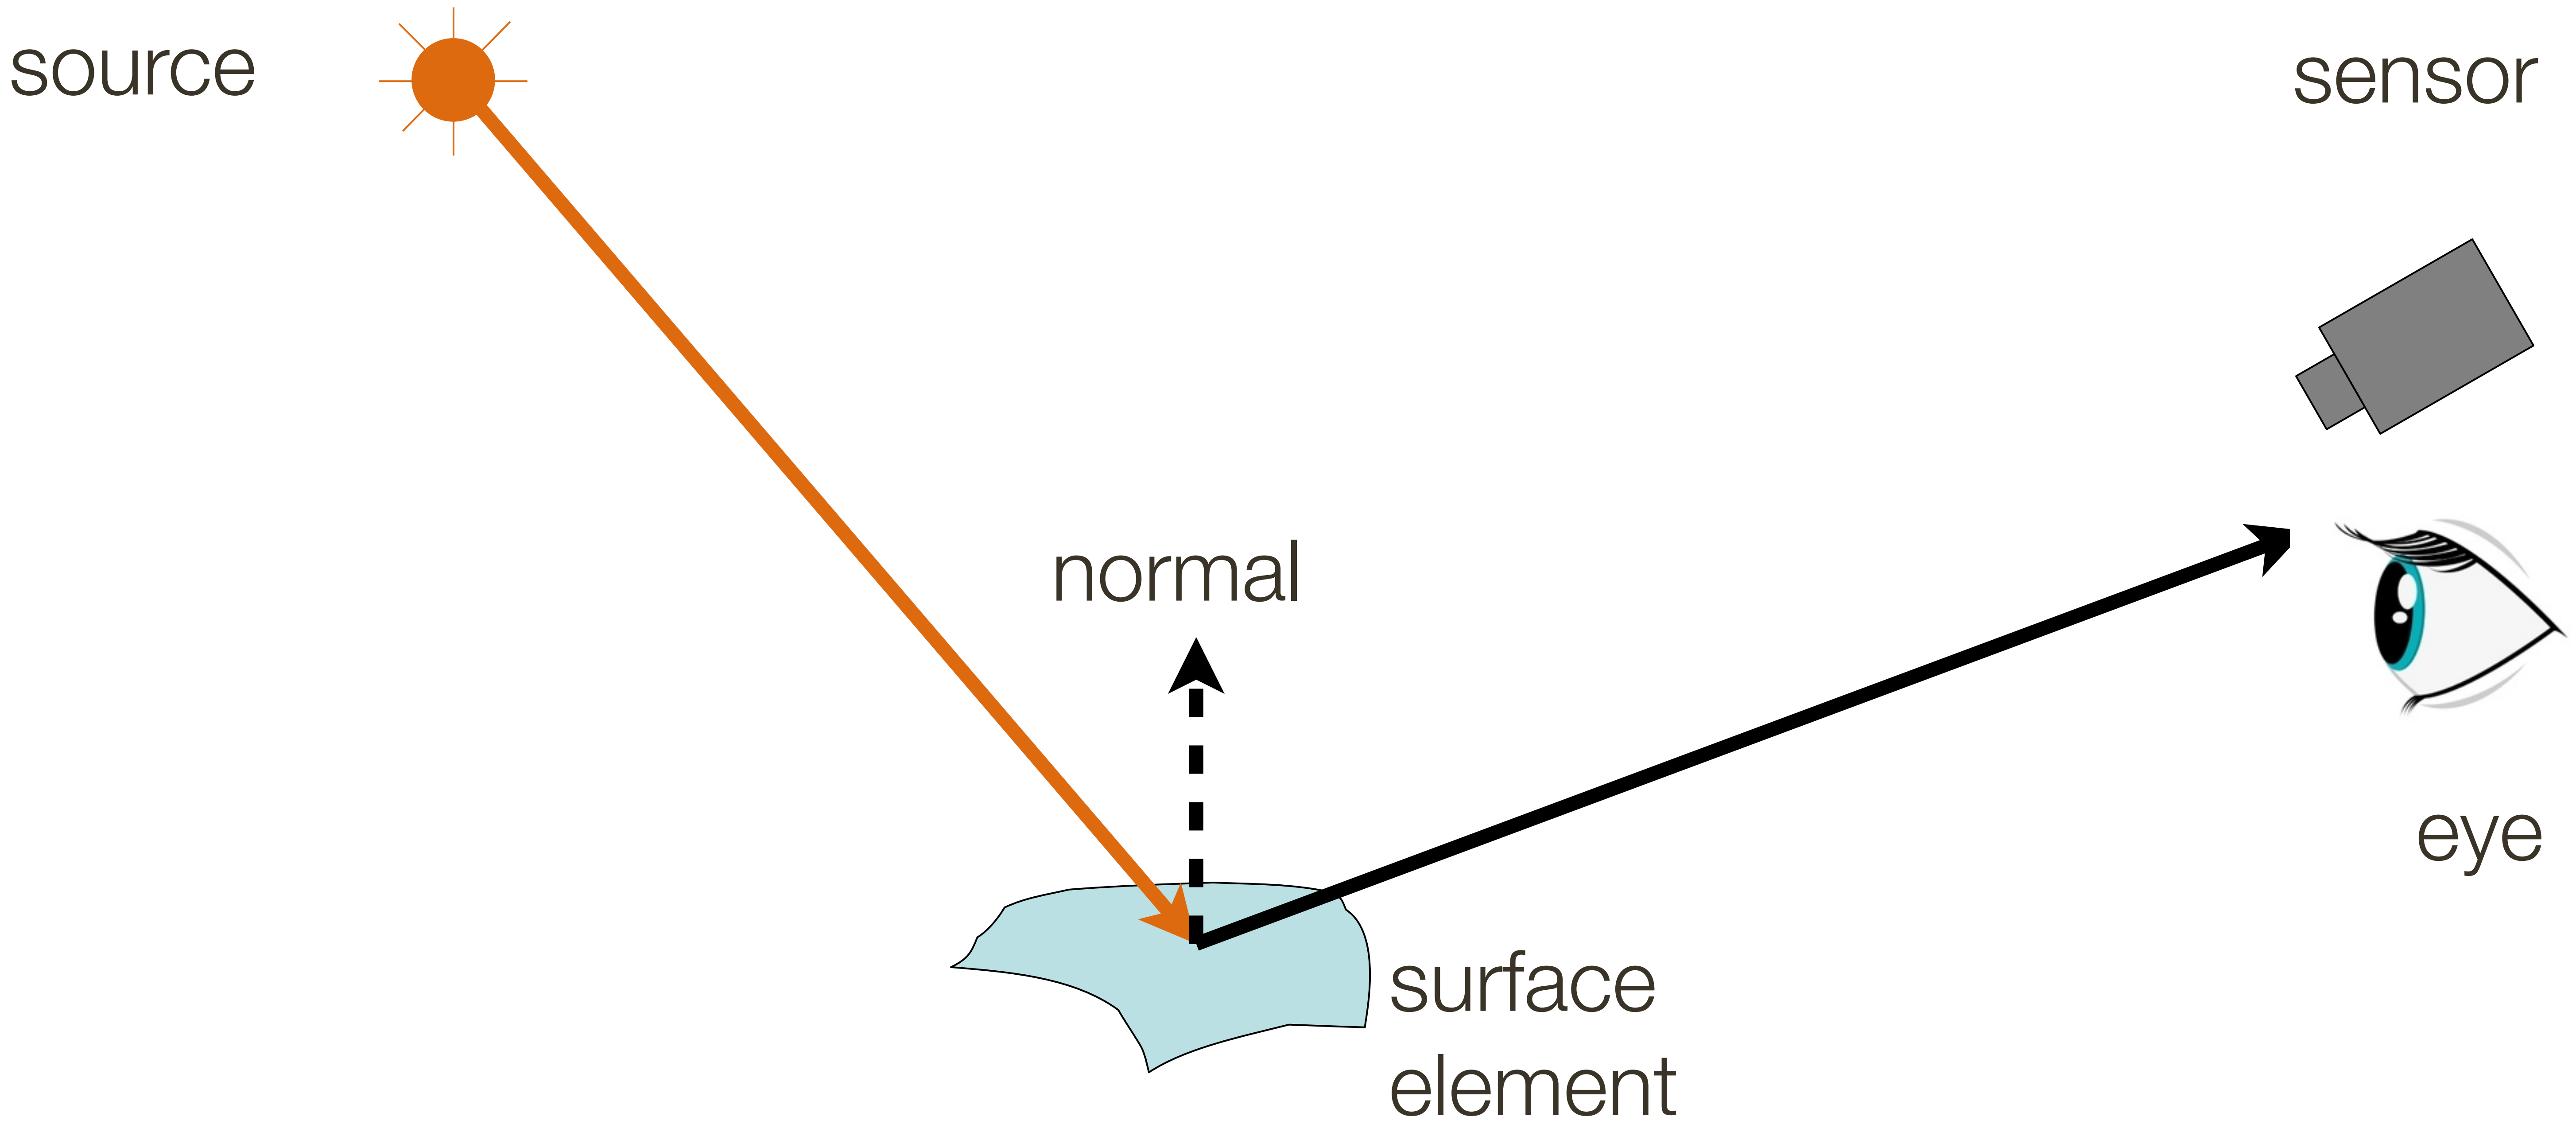

Sensor (or eye) **captures amount of light** reflected from the object

Colour

— Light is produced in different amounts at different wavelengths by each light source

— Light is differentially reflected at each wavelength, which gives objects their natural colour (**surface albedo**)

— The sensation of colour is determined by the human visual system, based on the product of light and reflectance

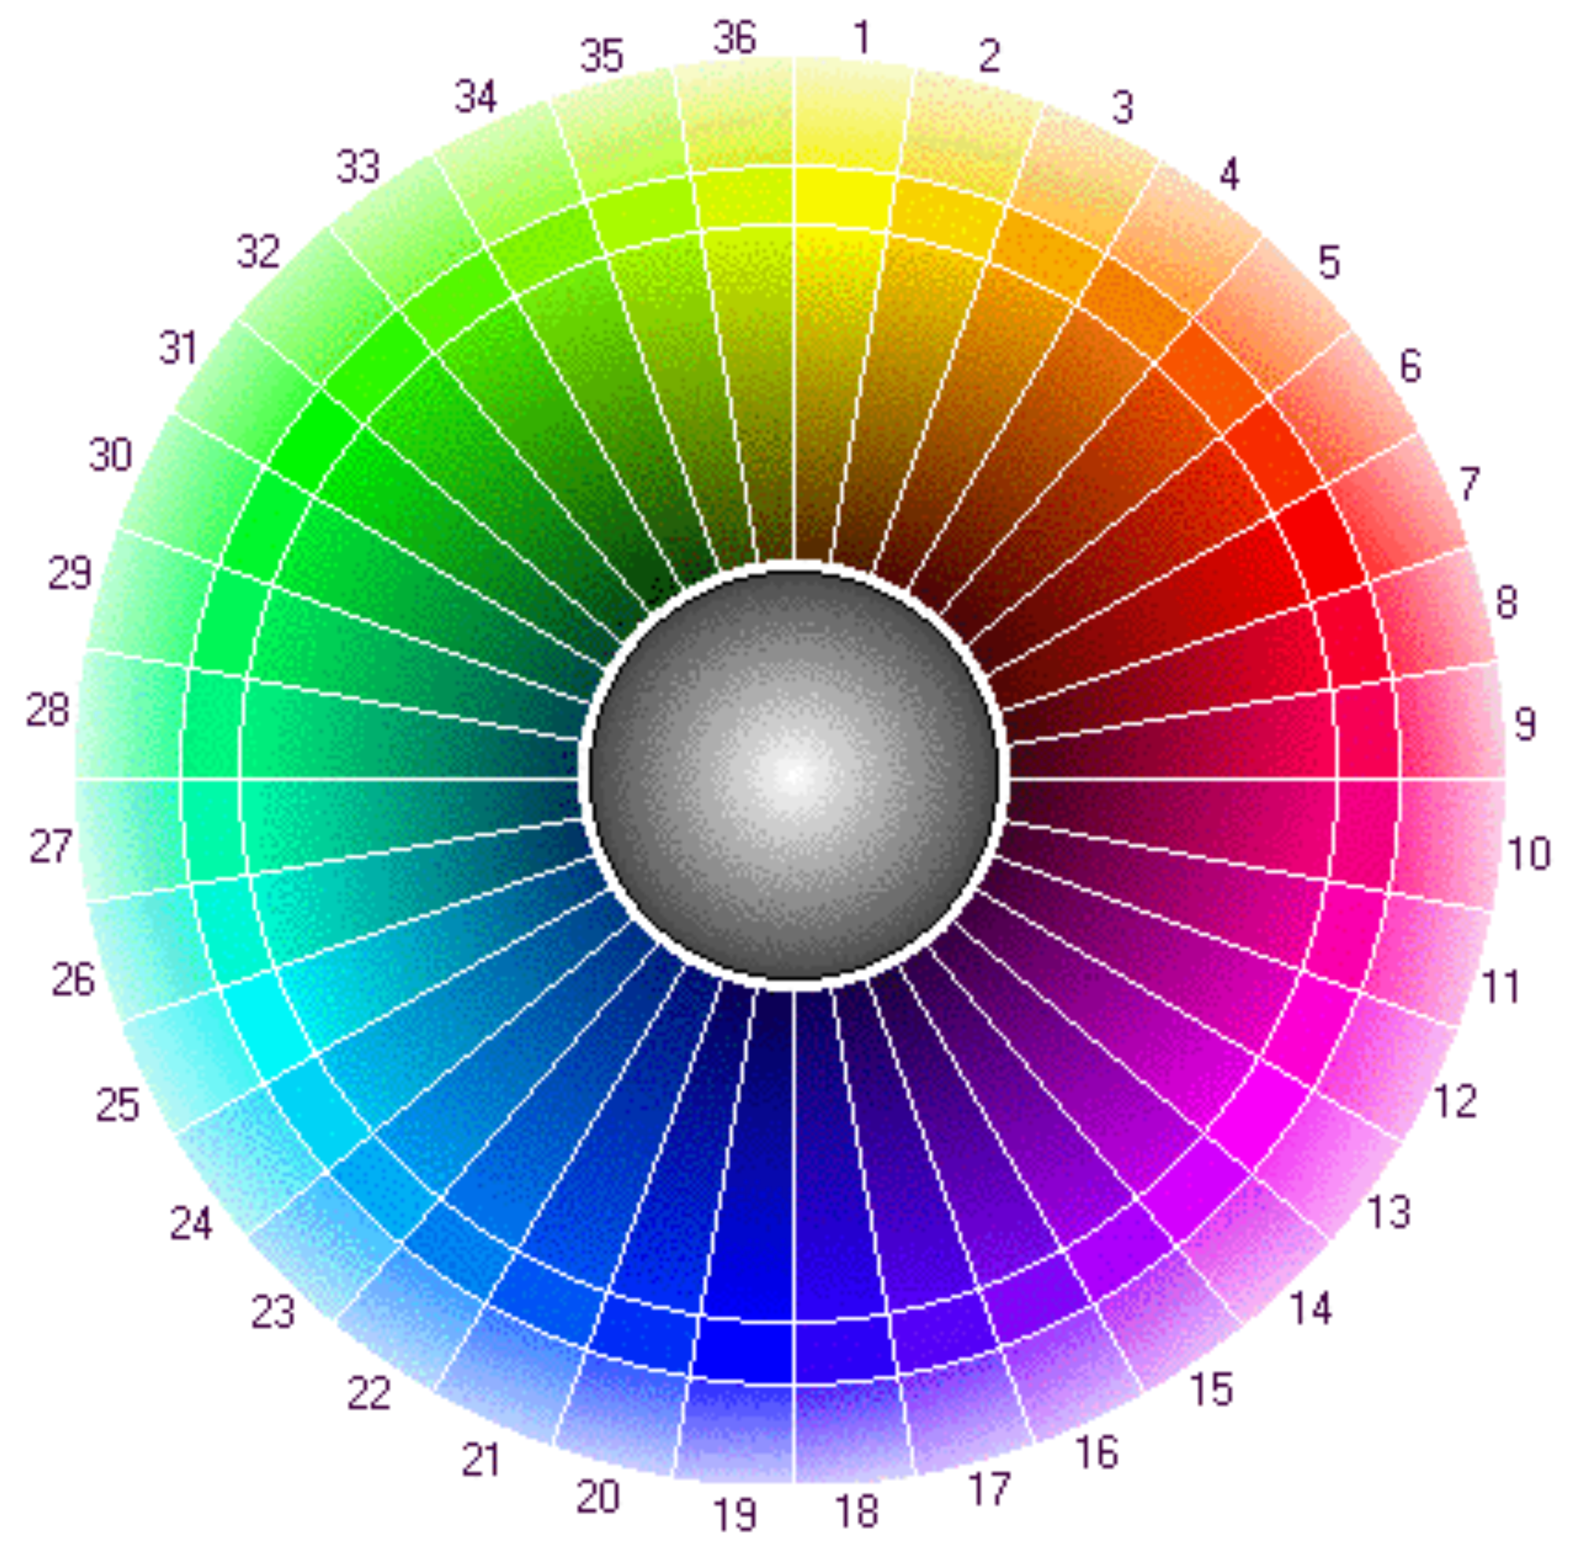

Relative **Spectral Power** of Two Illuminants

Relative spectral power plotted against wavelength in nm

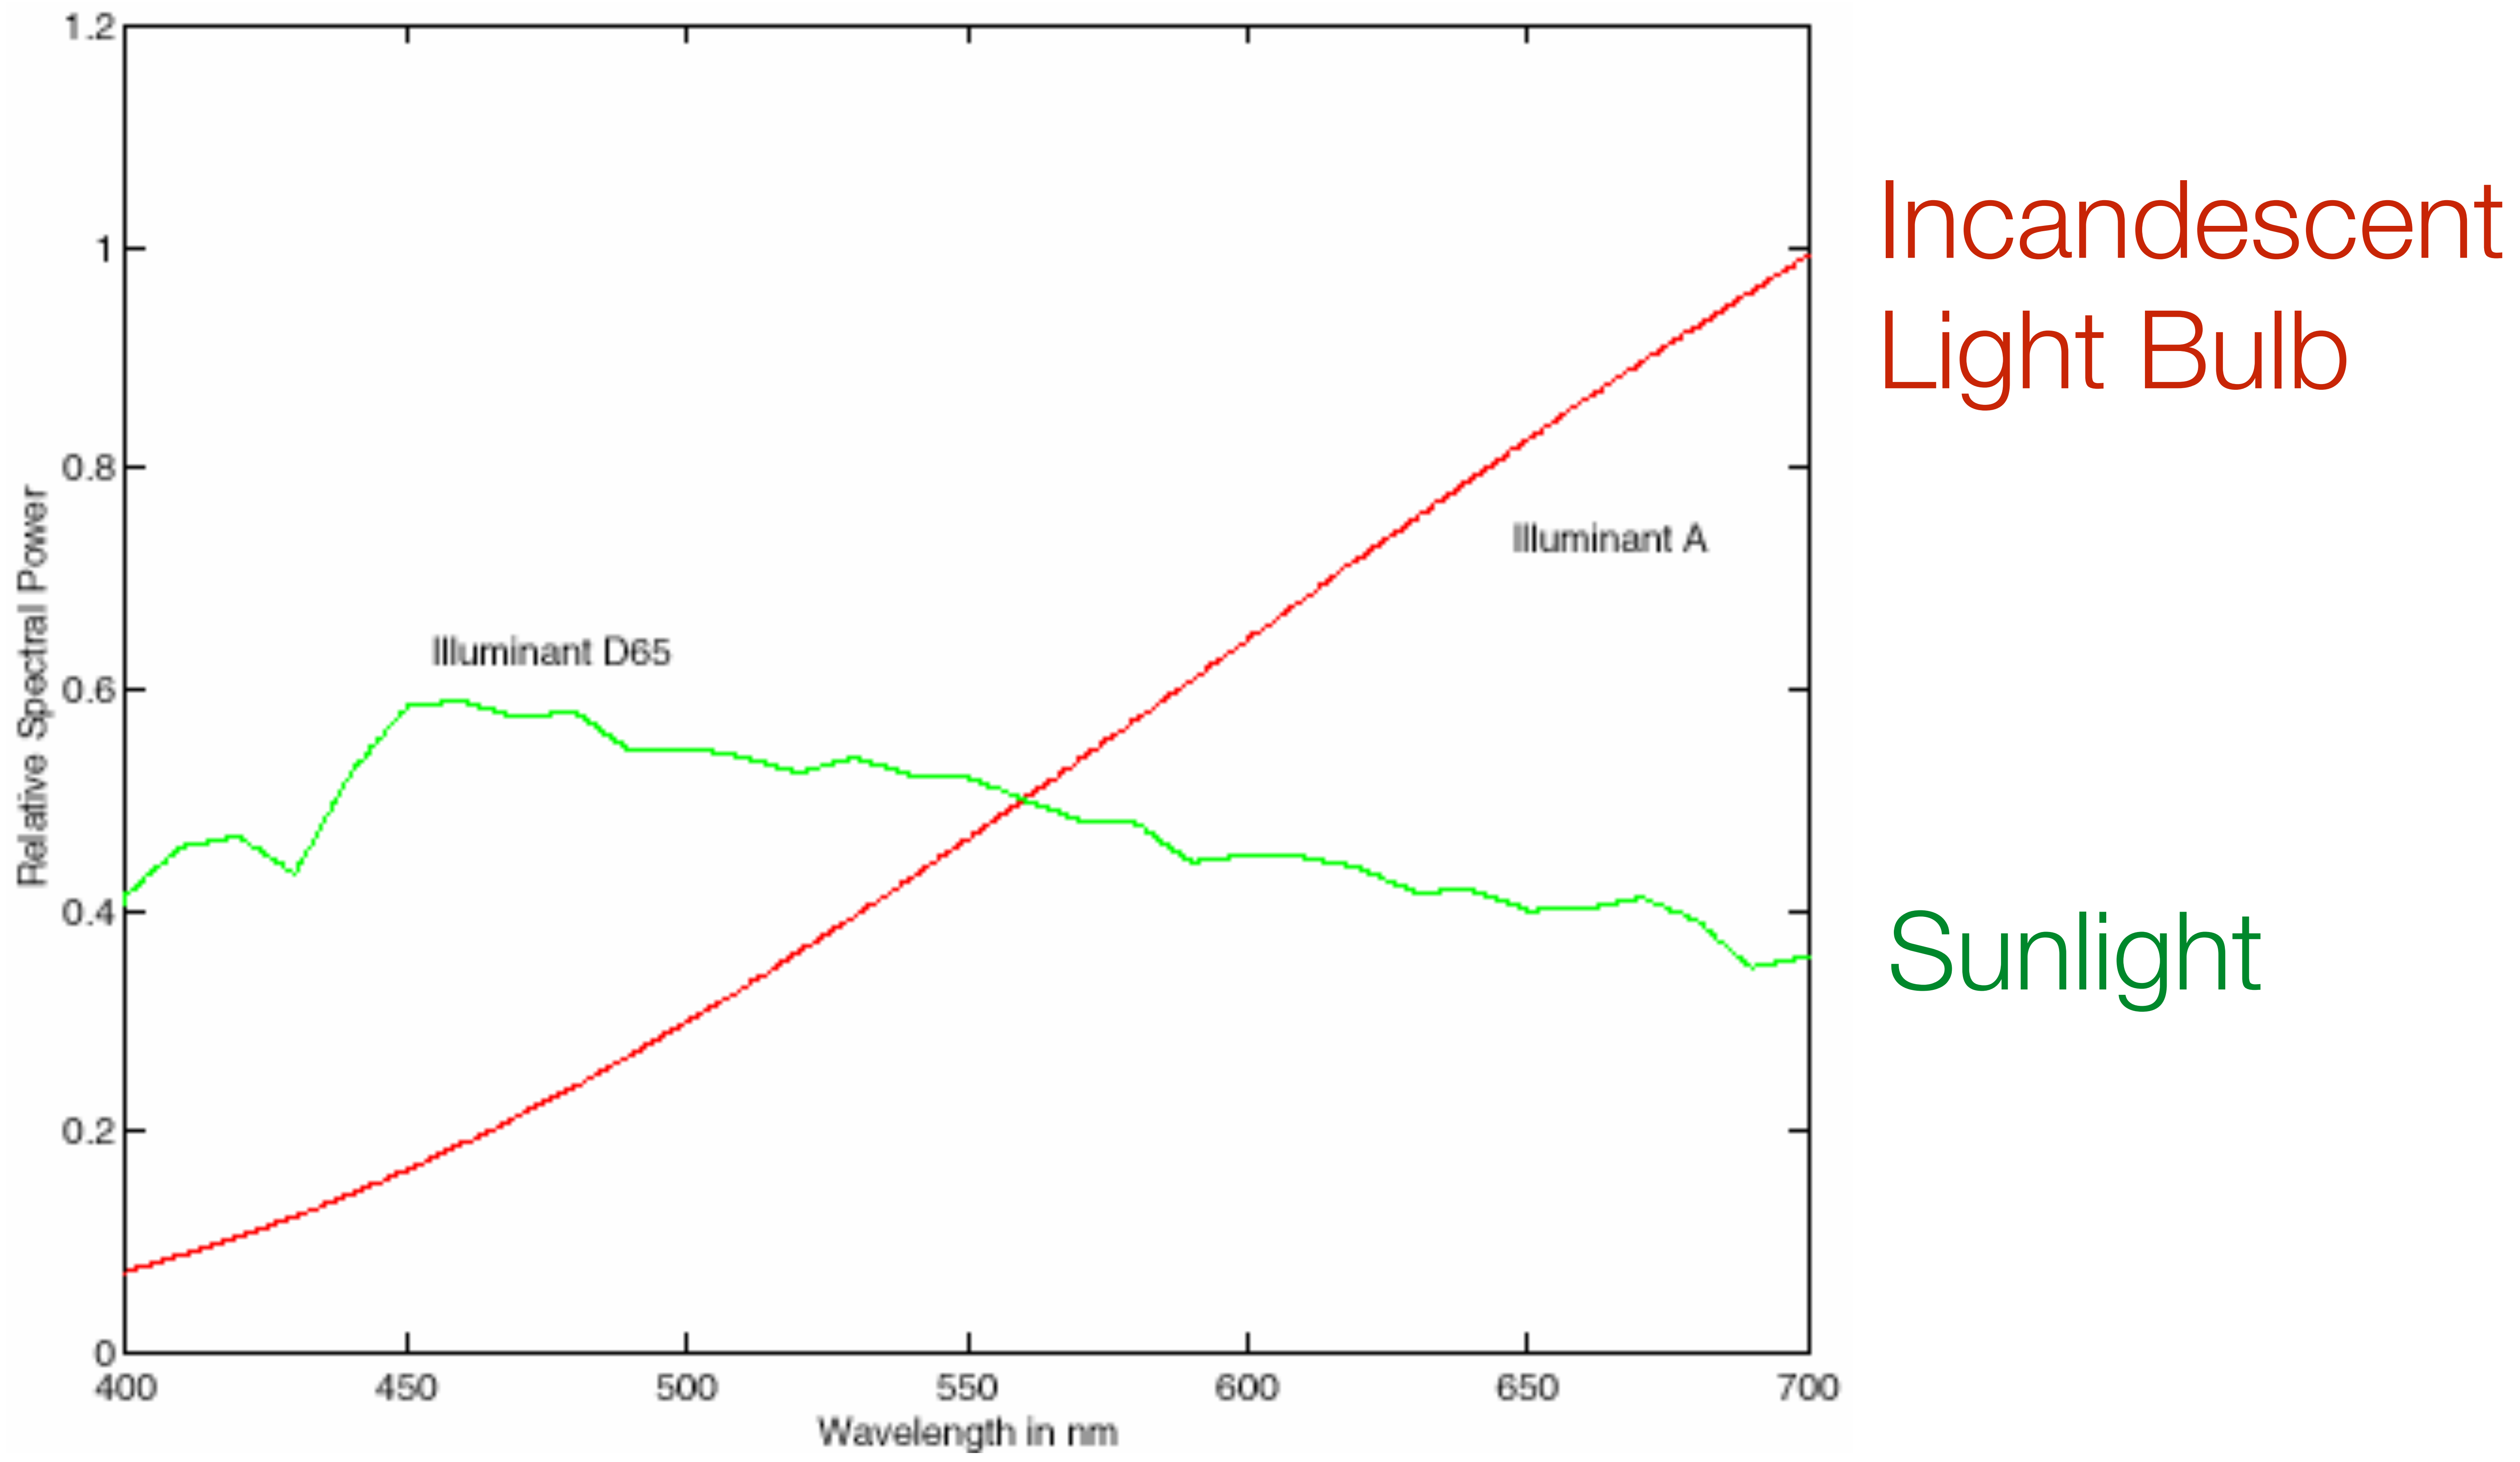

Forsyth & Ponce (2nd ed.) Figure 3.4

(small) **Graphics** Review

Surface reflection depends on both the **viewing** (θ_v, ϕ_v) and **illumination** (θ_i, ϕ_i) direction, with Bidirectional Reflection Distribution Function: $\mathbf{BRDF}(\theta_i, \phi_i, \theta_v, \phi_v)$

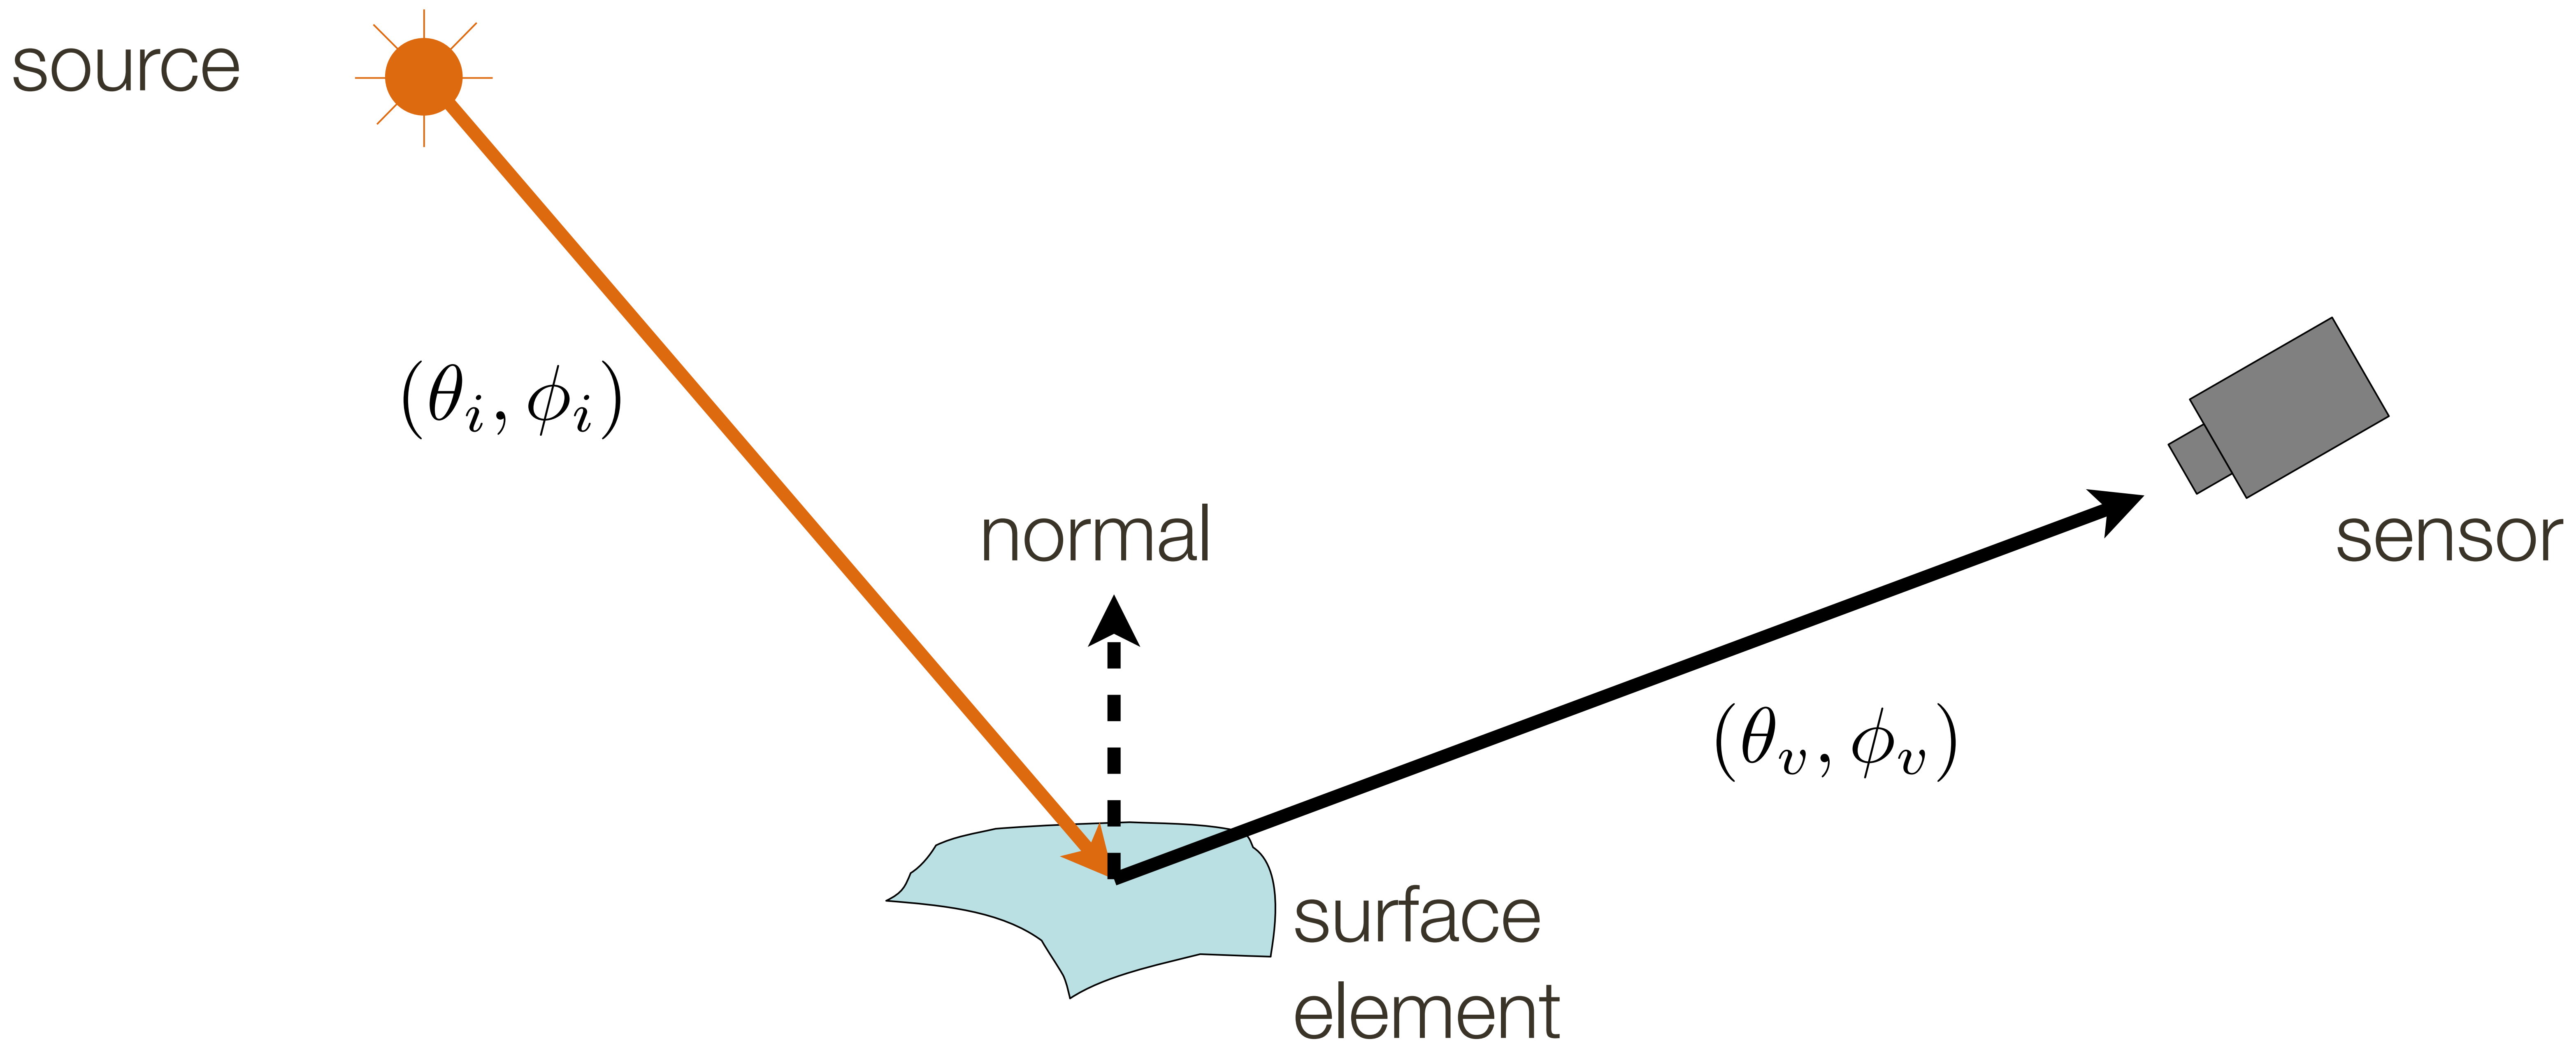

Slide adopted from: Ioannis (Yannis) Gkioulekas (CMU)

(small) **Graphics** Review

Surface reflection depends on both the **viewing** (θ_v, ϕ_v) and **illumination** (θ_i, ϕ_i) direction, with Bidirectional Reflection Distribution Function: $\mathbf{BRDF}(\theta_i, \phi_i, \theta_v, \phi_v)$

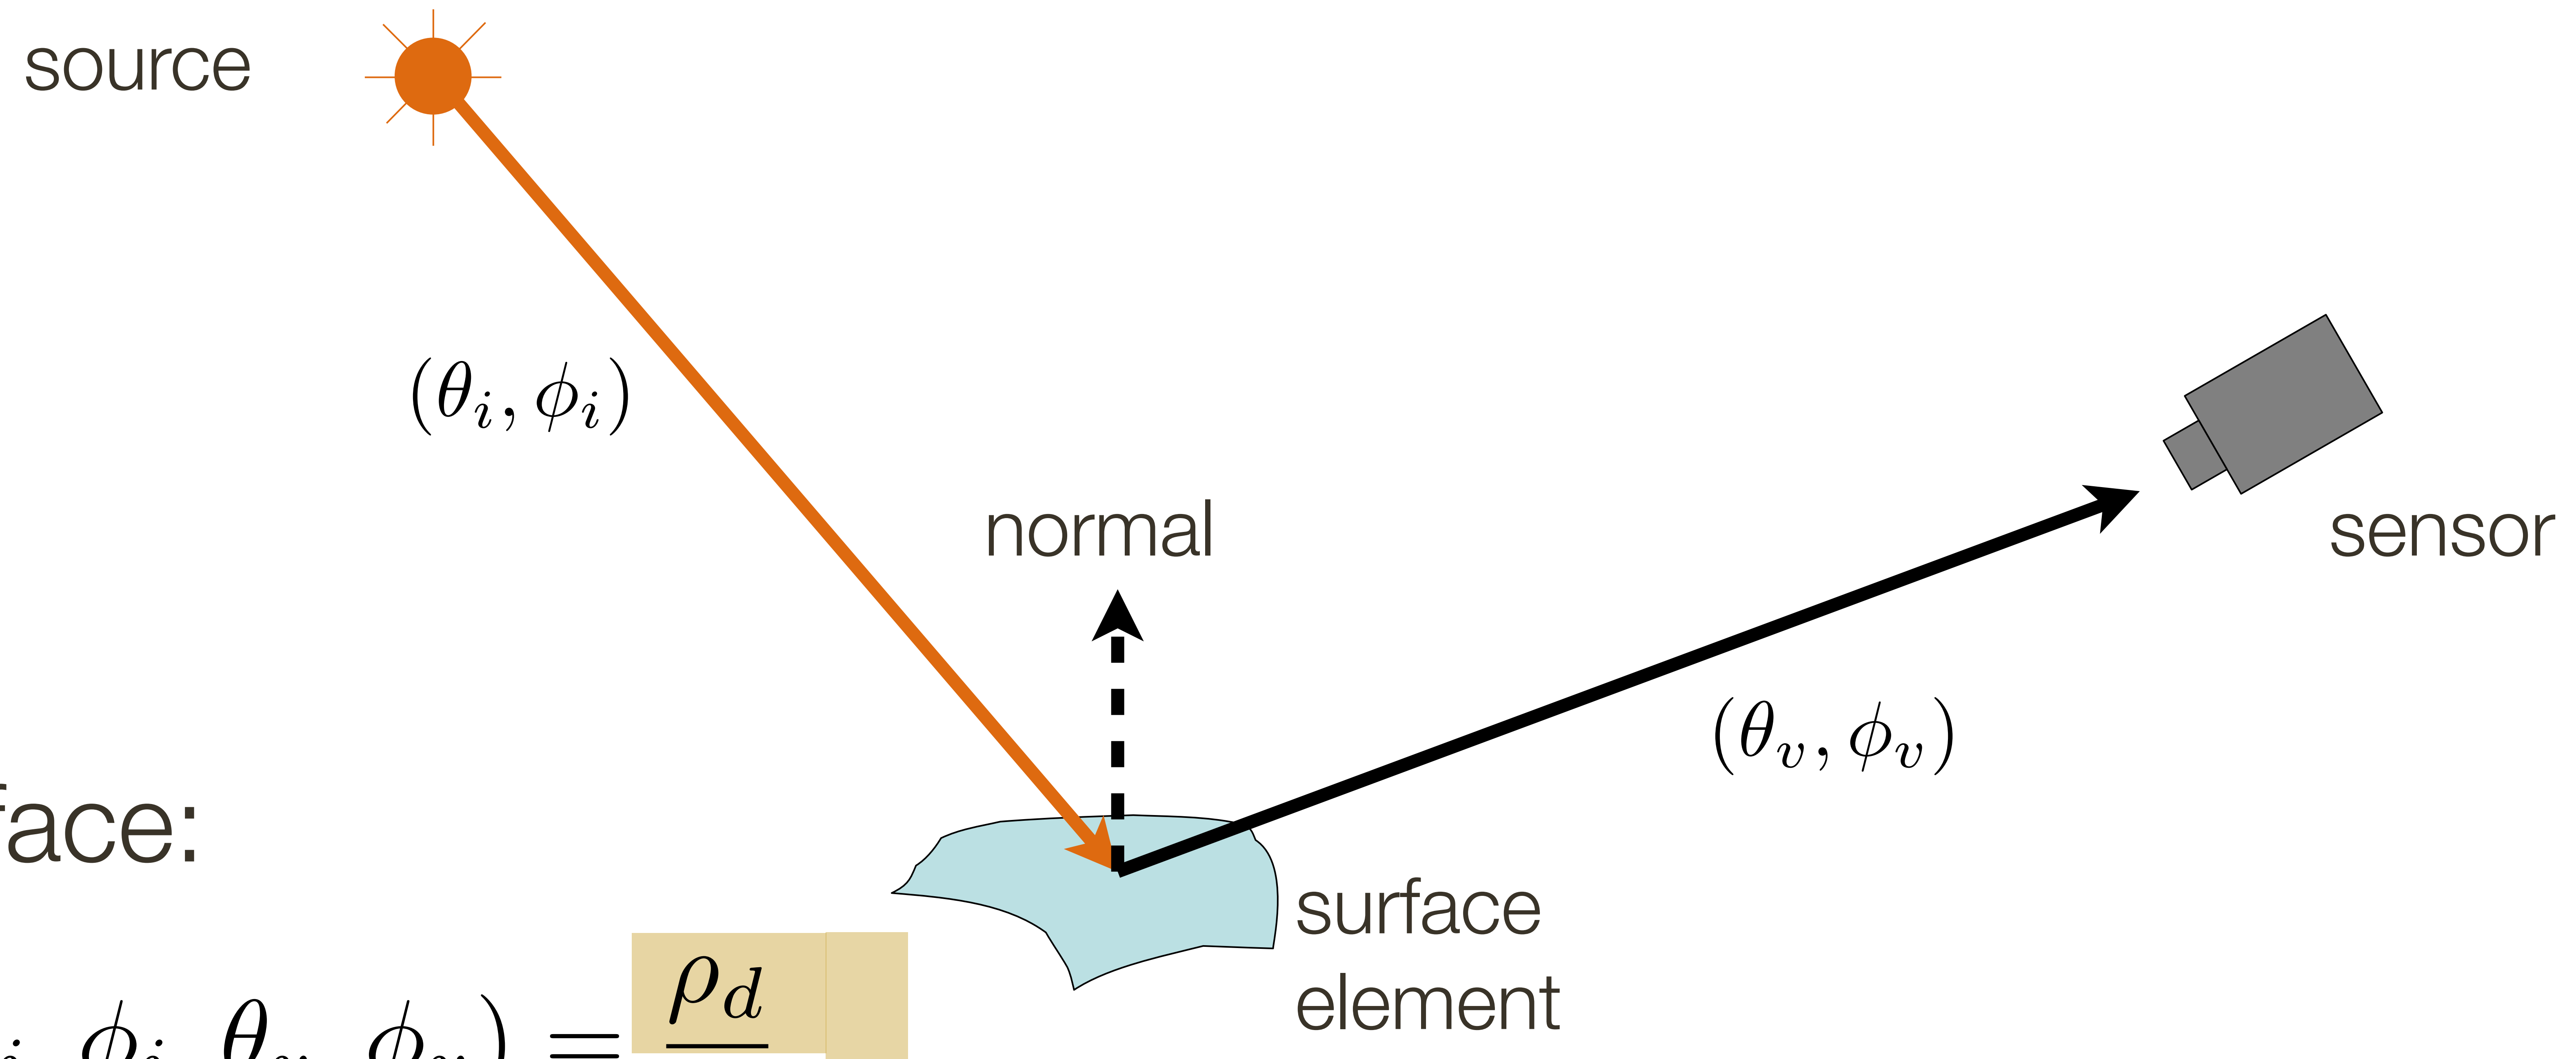

Lambertian surface:

$$\mathbf{BRDF}(\theta_i, \phi_i, \theta_v, \phi_v) = \frac{\rho_d}{\pi}$$

constant, called **albedo**

Slide adopted from: Ioannis (Yannis) Gkioulekas (CMU)

(small) **Graphics** Review

Question: What are the simplifying assumptions we are making here?

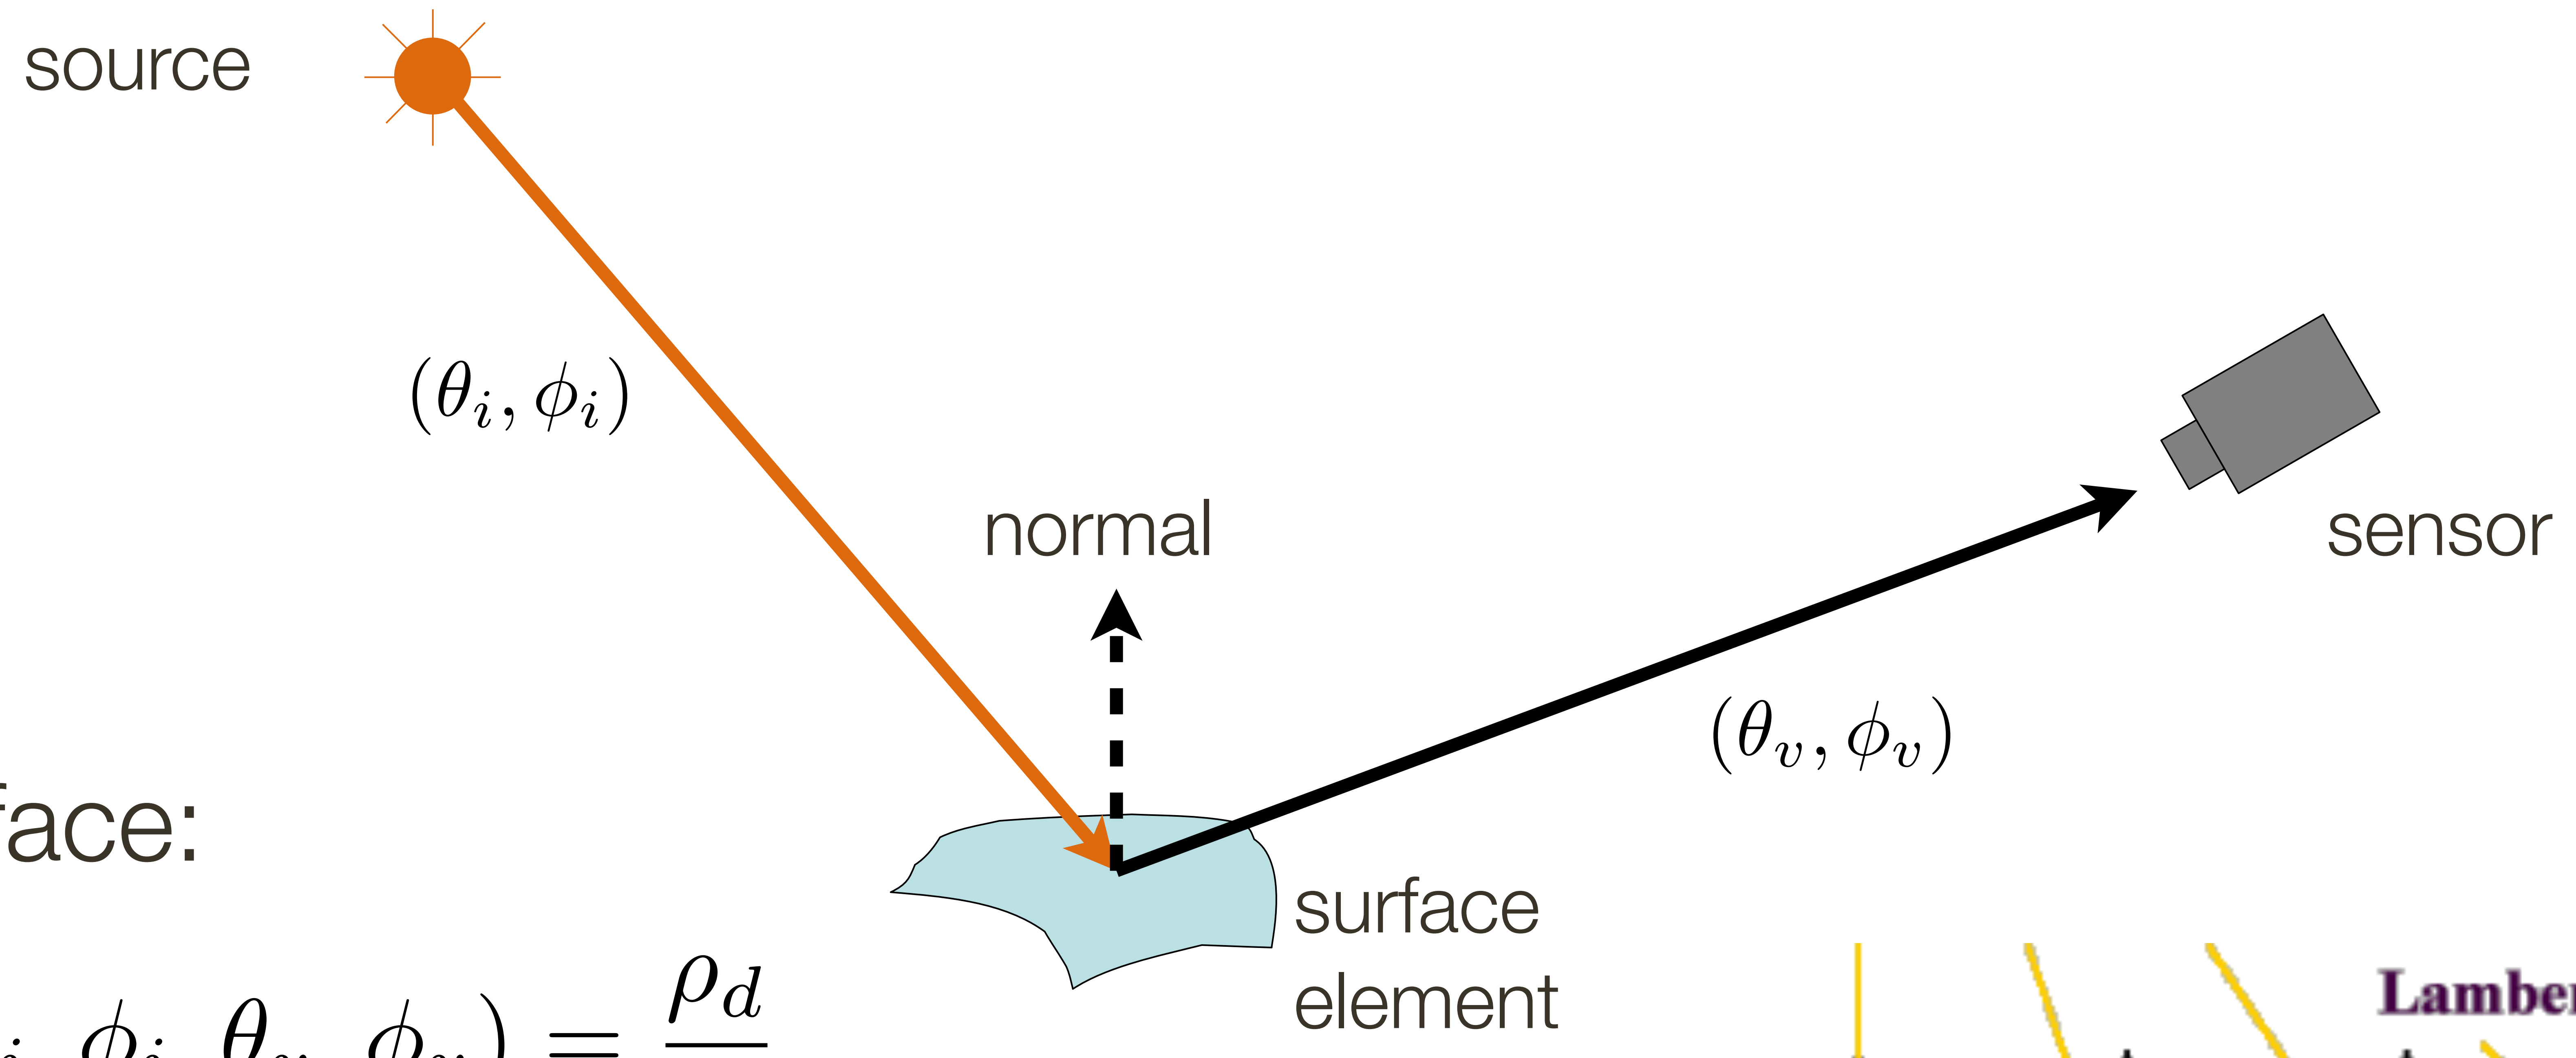

Lambertian surface:

$$\mathbf{BRDF}(\theta_i, \phi_i, \theta_v, \phi_v) = \frac{\rho_d}{\pi}$$

$$L = \frac{\rho_d}{\pi} I(\vec{i} \cdot \vec{n})$$

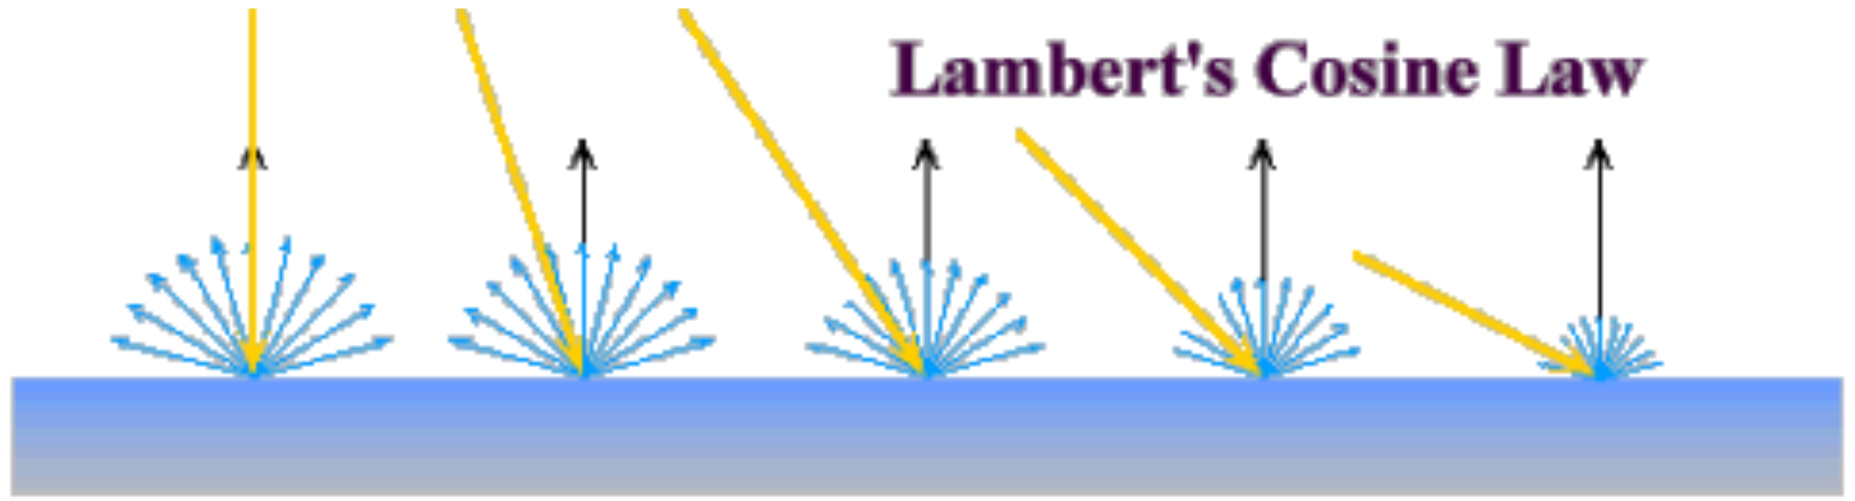

Slide adopted from: Ioannis (Yannis) Gkioulekas (CMU)

(small) **Graphics** Review

Question: What are the simplifying assumptions we are making here?

1. Light spectra is absorbed uniformly by the surface (no change in color)

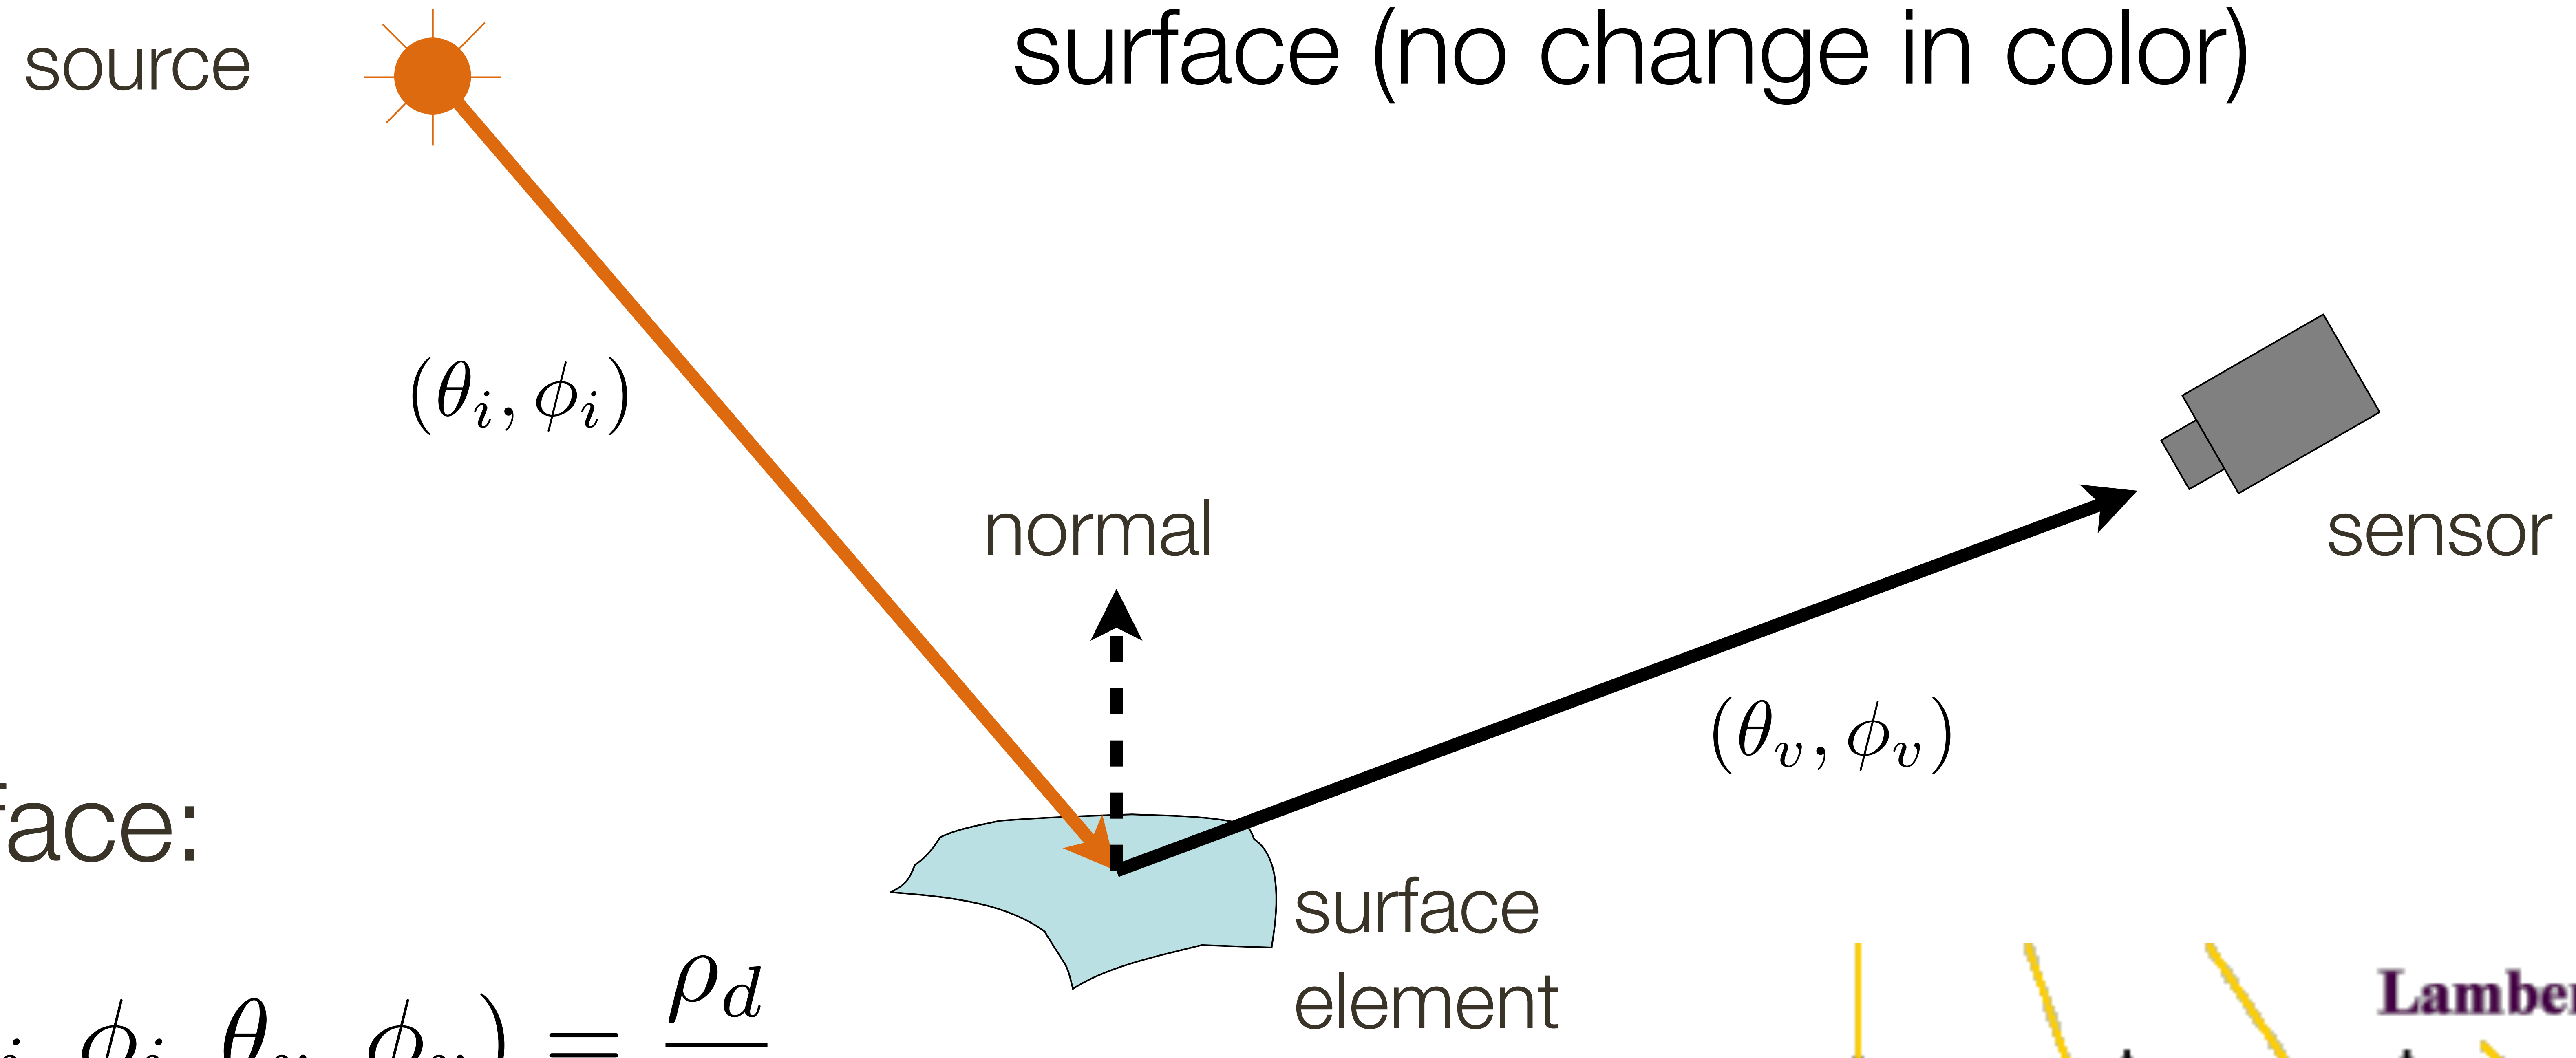

Lambertian surface:

$$\mathbf{BRDF}(\theta_i, \phi_i, \theta_v, \phi_v) = \frac{\rho_d}{\pi}$$

$$L = \frac{\rho_d}{\pi} I(\vec{i} \cdot \vec{n})$$

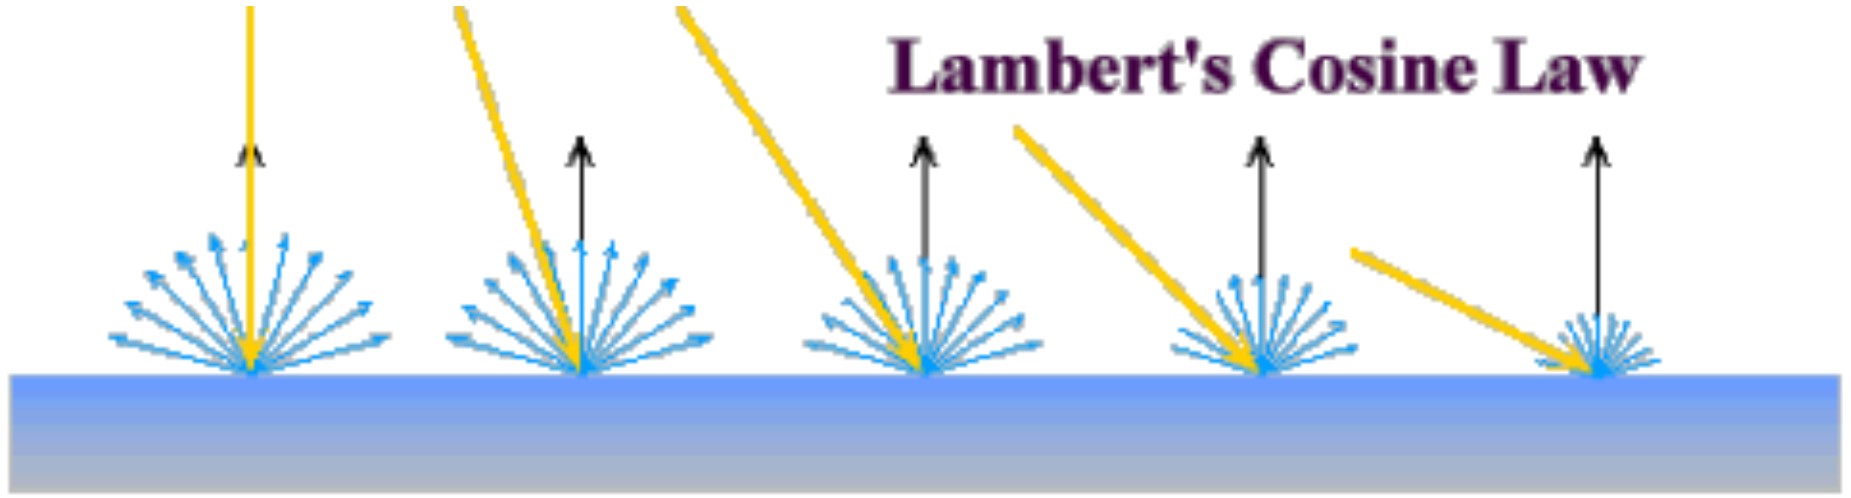

Slide adopted from: Ioannis (Yannis) Gkioulekas (CMU)

Spectral **Albedo** of Natural Surfaces

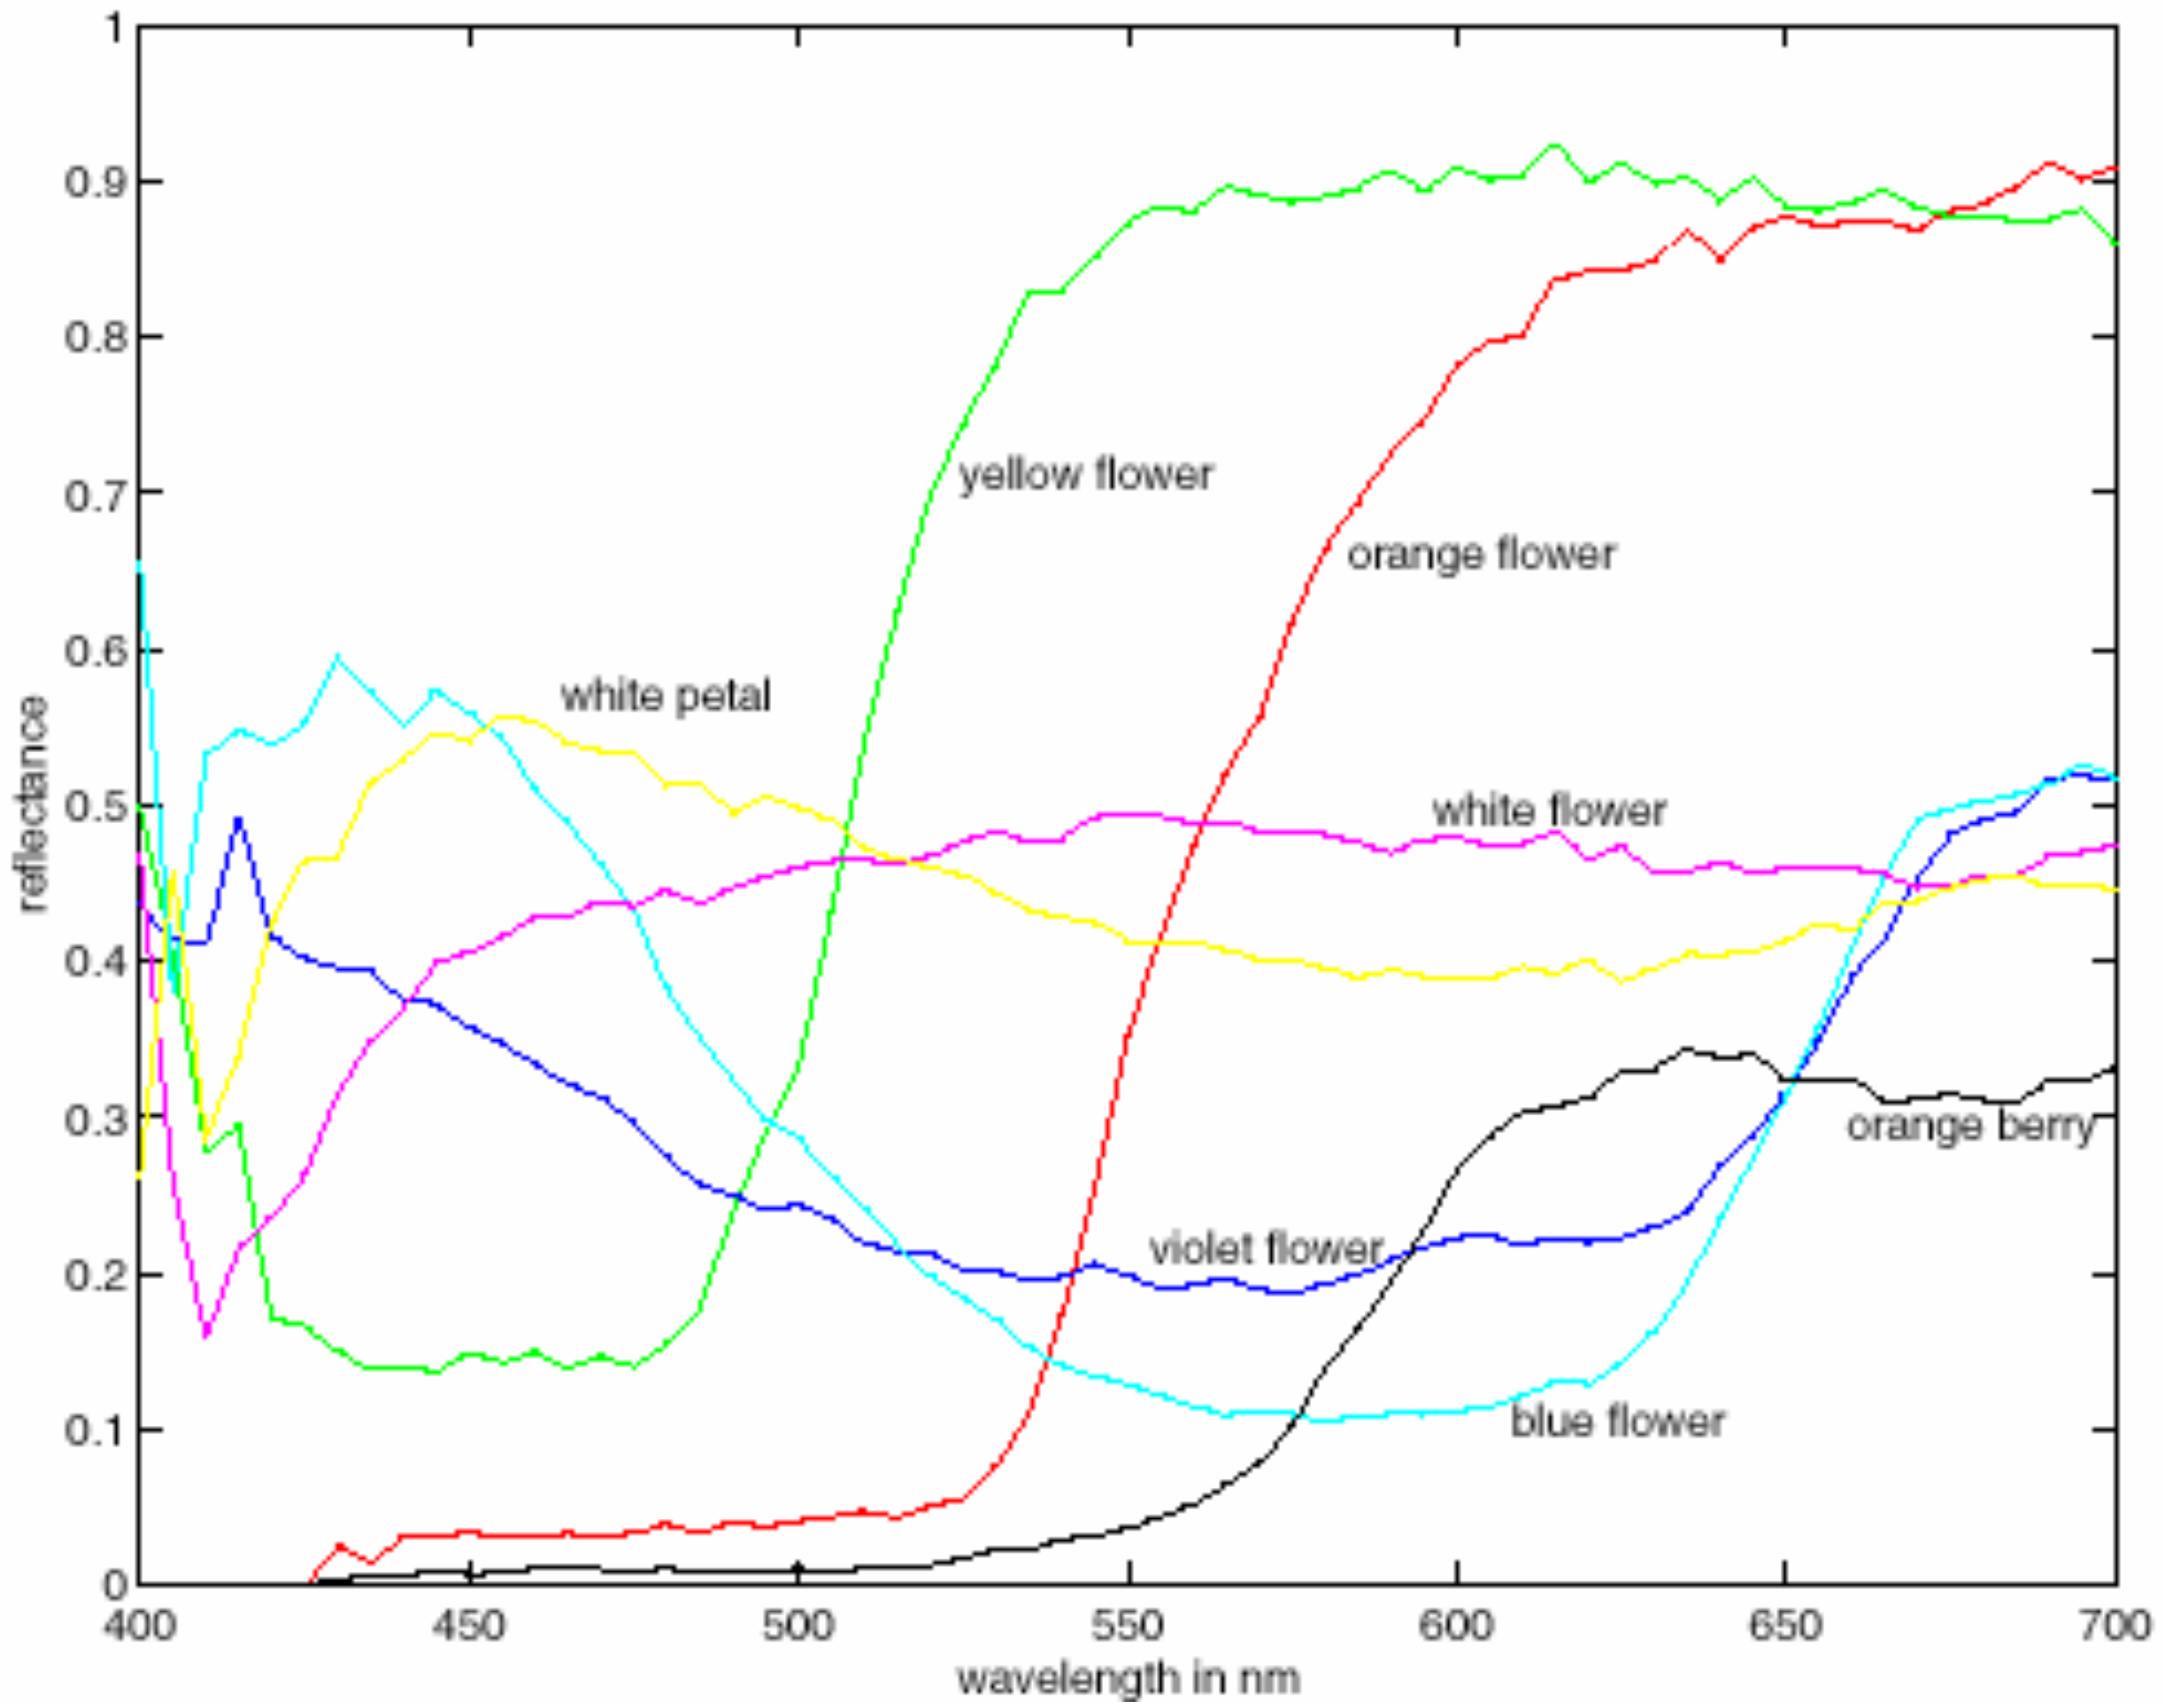

Forsyth & Ponce (2nd ed.) Figure 3.6

Colour Appearance

Reflected light **at each wavelength** is the product of illumination and surface reflectance at that wavelength

Surface reflectance often is modeled as having two components:

— **Lambertian** reflectance: equal in all directions (diffuse)

— **Specular** reflectance: mirror reflectance (shiny spots)

(small) **Graphics** Review

Surface reflection depends on both the **viewing** (θ_v, ϕ_v) and **illumination** (θ_i, ϕ_i) direction, with Bidirectional Reflection Distribution Function: $\mathbf{BRDF}(\theta_i, \phi_i, \theta_v, \phi_v)$

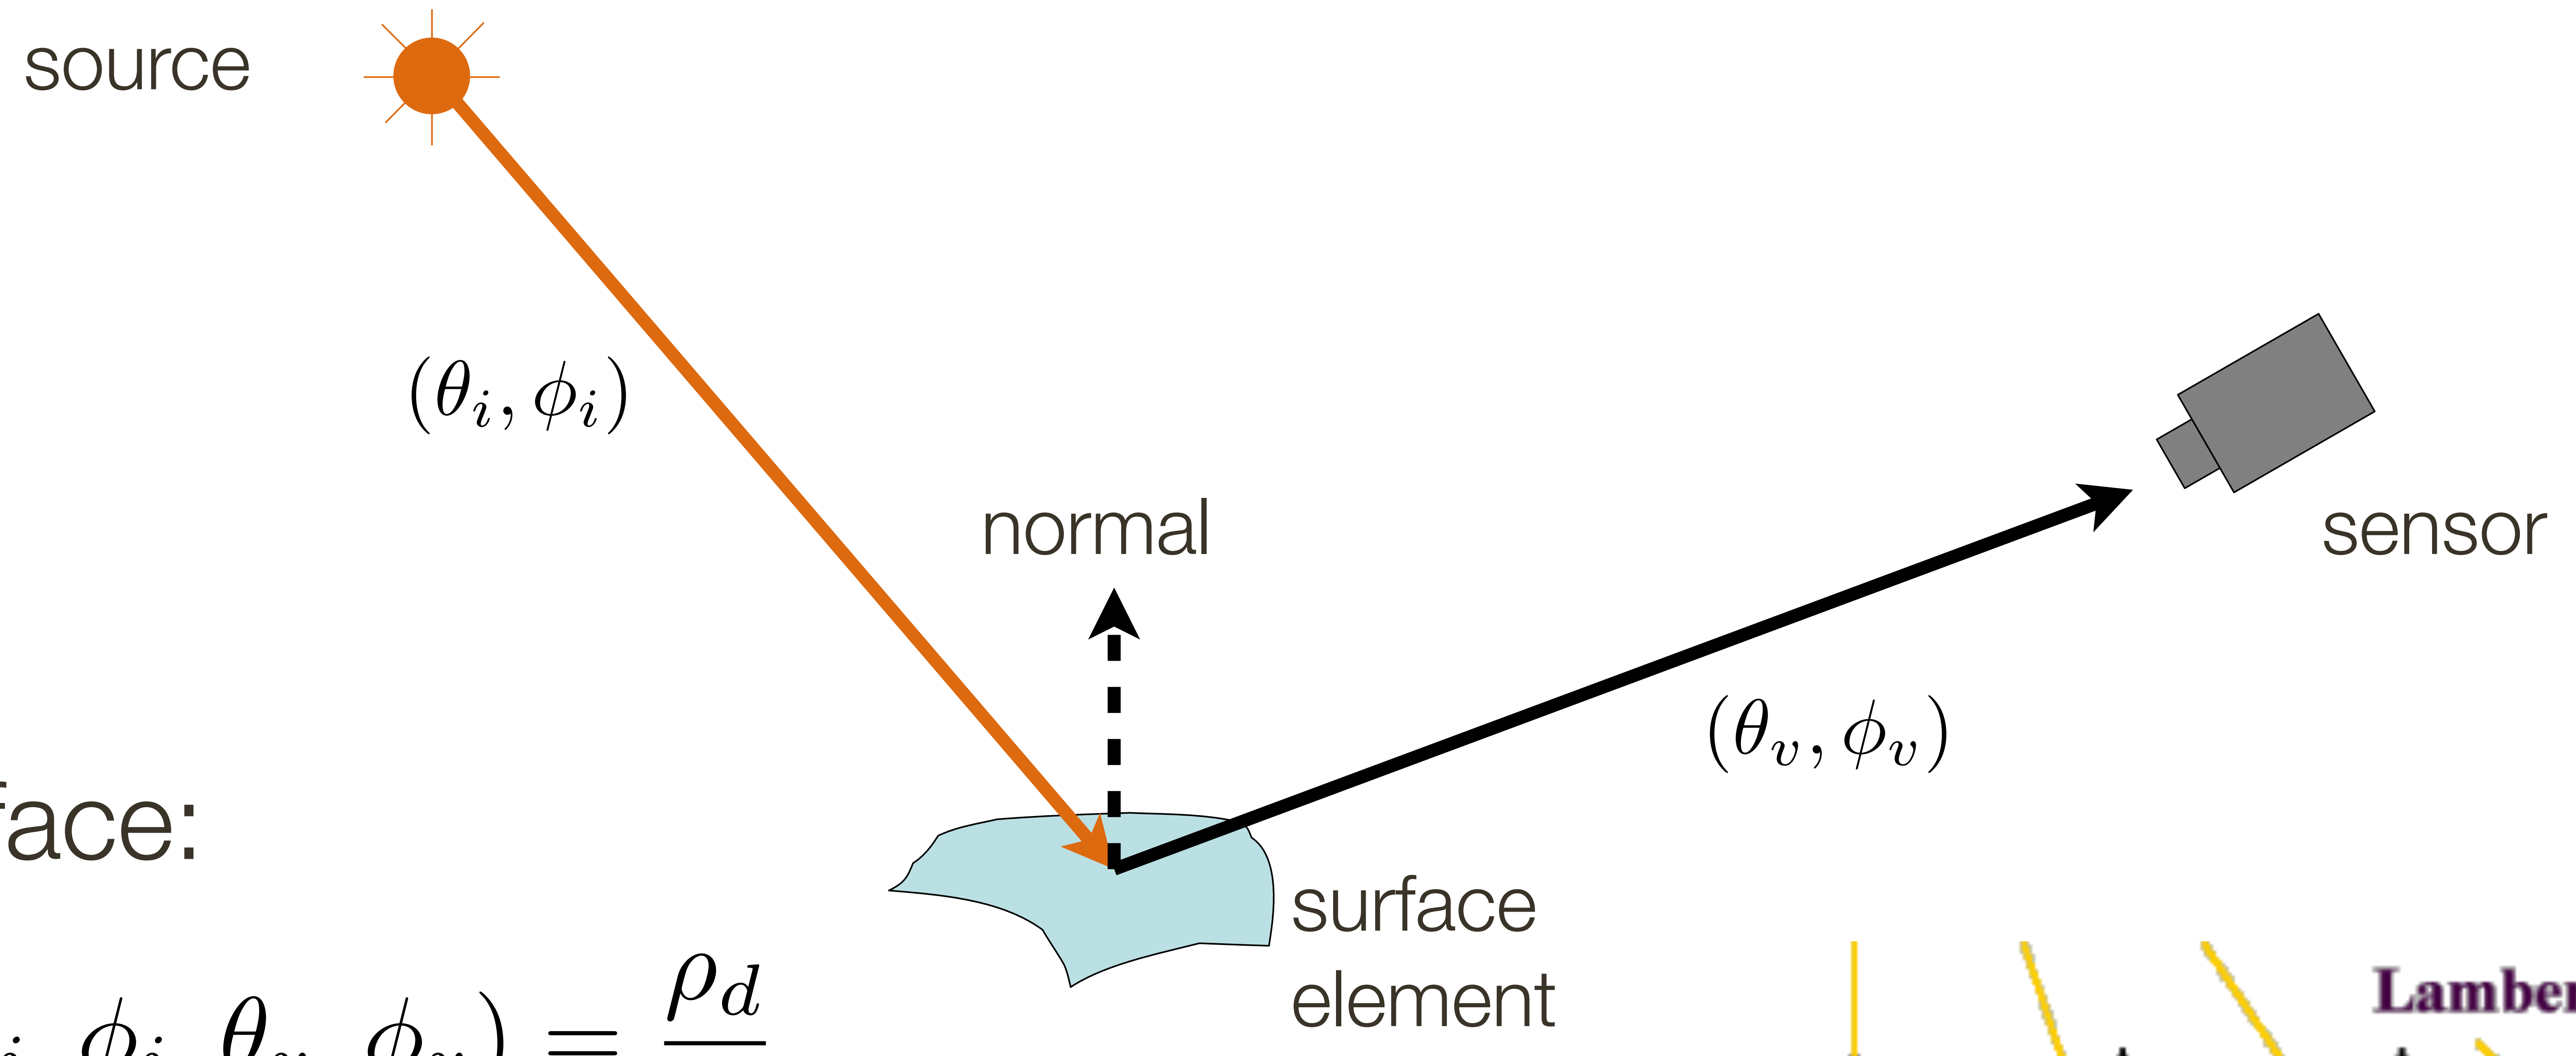

Lambertian surface:

$$\mathbf{BRDF}(\theta_i, \phi_i, \theta_v, \phi_v) = \frac{\rho_d}{\pi}$$

$$L = \frac{\rho_d}{\pi} I(\vec{i} \cdot \vec{n})$$

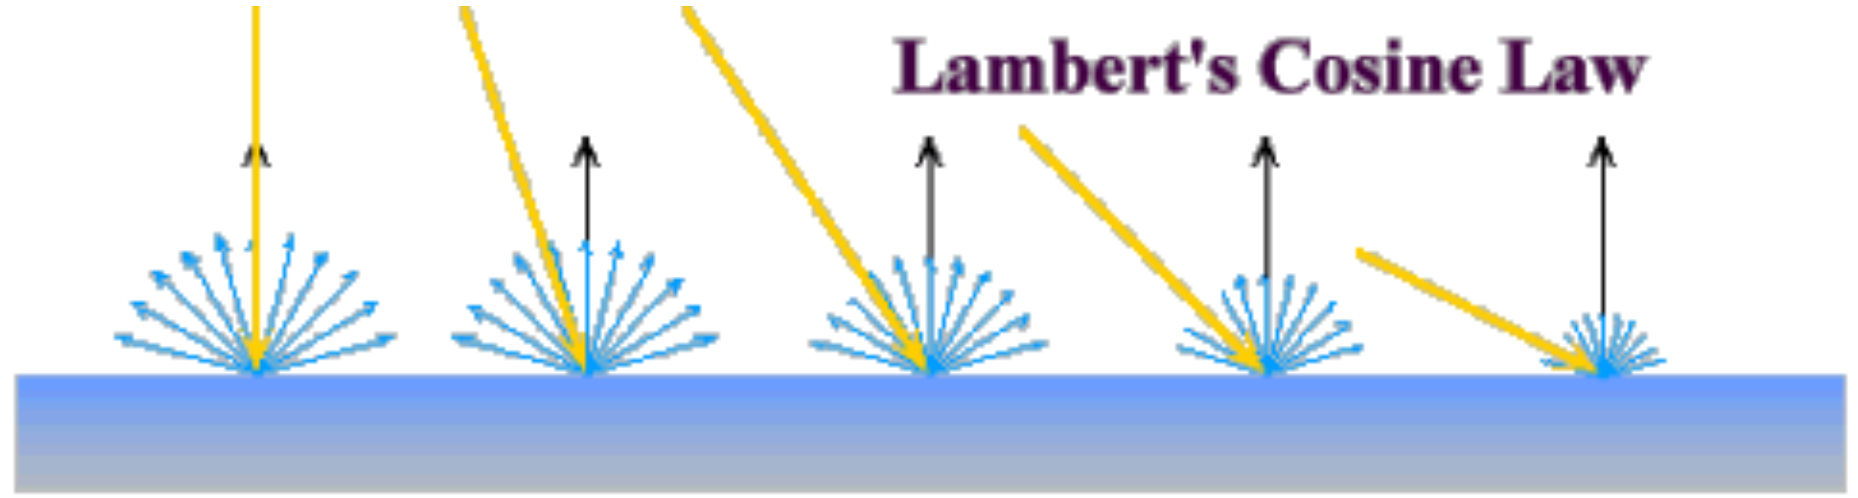

Slide adopted from: Ioannis (Yannis) Gkioulekas (CMU)

(small) **Graphics** Review

Surface reflection depends on both the **viewing** (θ_v, ϕ_v) and **illumination** (θ_i, ϕ_i) direction, with Bidirectional Reflection Distribution Function: $\mathbf{BRDF}(\theta_i, \phi_i, \theta_v, \phi_v)$

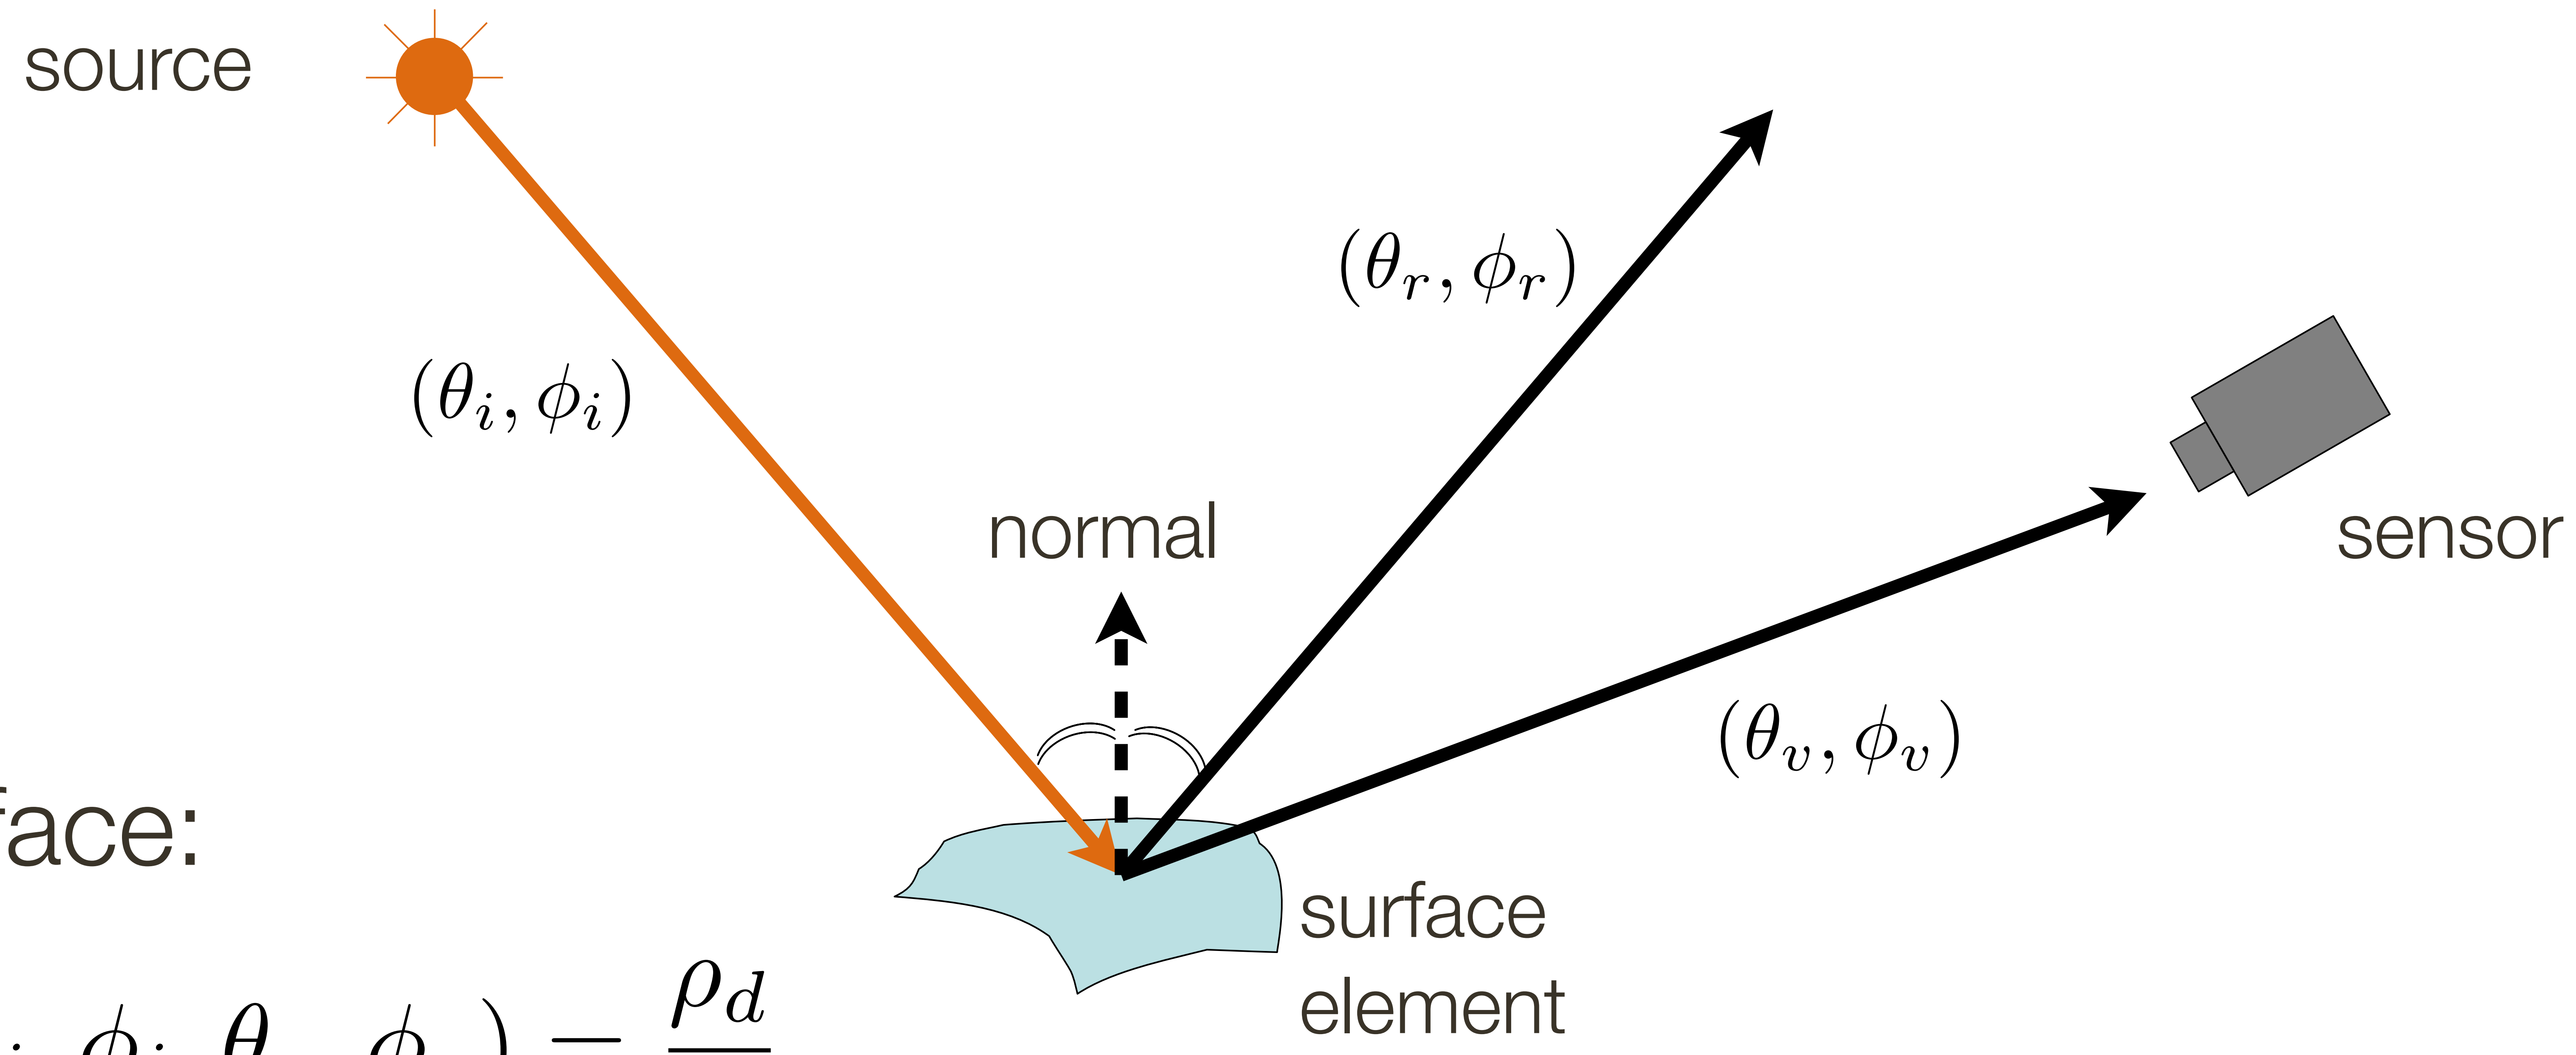

Lambertian surface:

$$\mathbf{BRDF}(\theta_i, \phi_i, \theta_v, \phi_v) = \frac{\rho_d}{\pi}$$

Mirror surface: all incident light reflected in one directions $(\theta_v, \phi_v) = (\theta_r, \phi_r)$

Slide adopted from: Ioannis (Yannis) Gkioulekas (CMU)

Color Matching Experiments

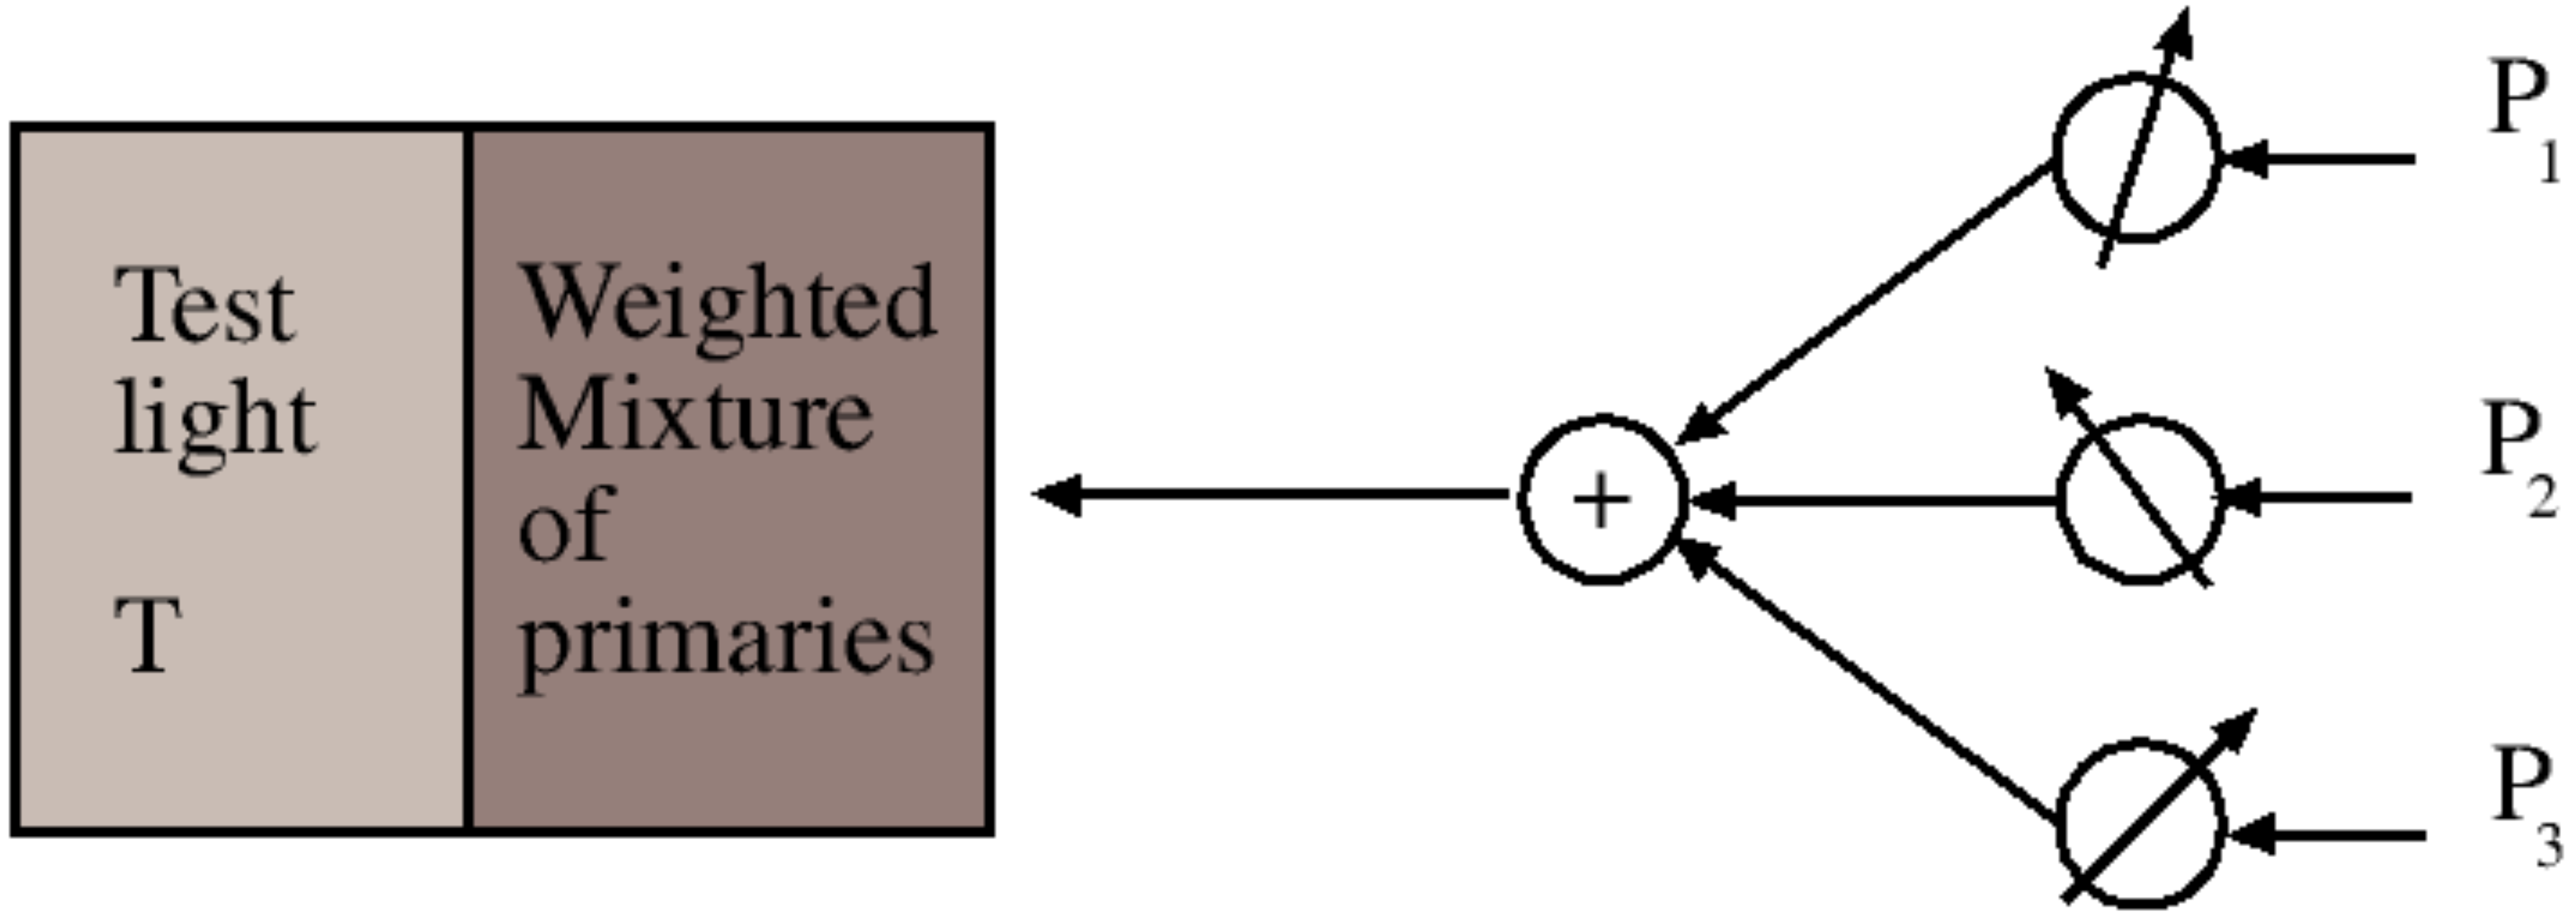

Forsyth & Ponce (2nd ed.) Figure 3.2

Show a split field to subjects. One side shows the light whose colour one wants to match. The other a weighted mixture of three primaries (fixed lights)

$$T = w_1 P_1 + w_2 P_2 + w_3 P_3$$

Recall: Color is an Artifact of Human Perception

"Color" is **not** an objective physical property of light (electromagnetic radiation).

Instead, light is characterized by its wavelength.

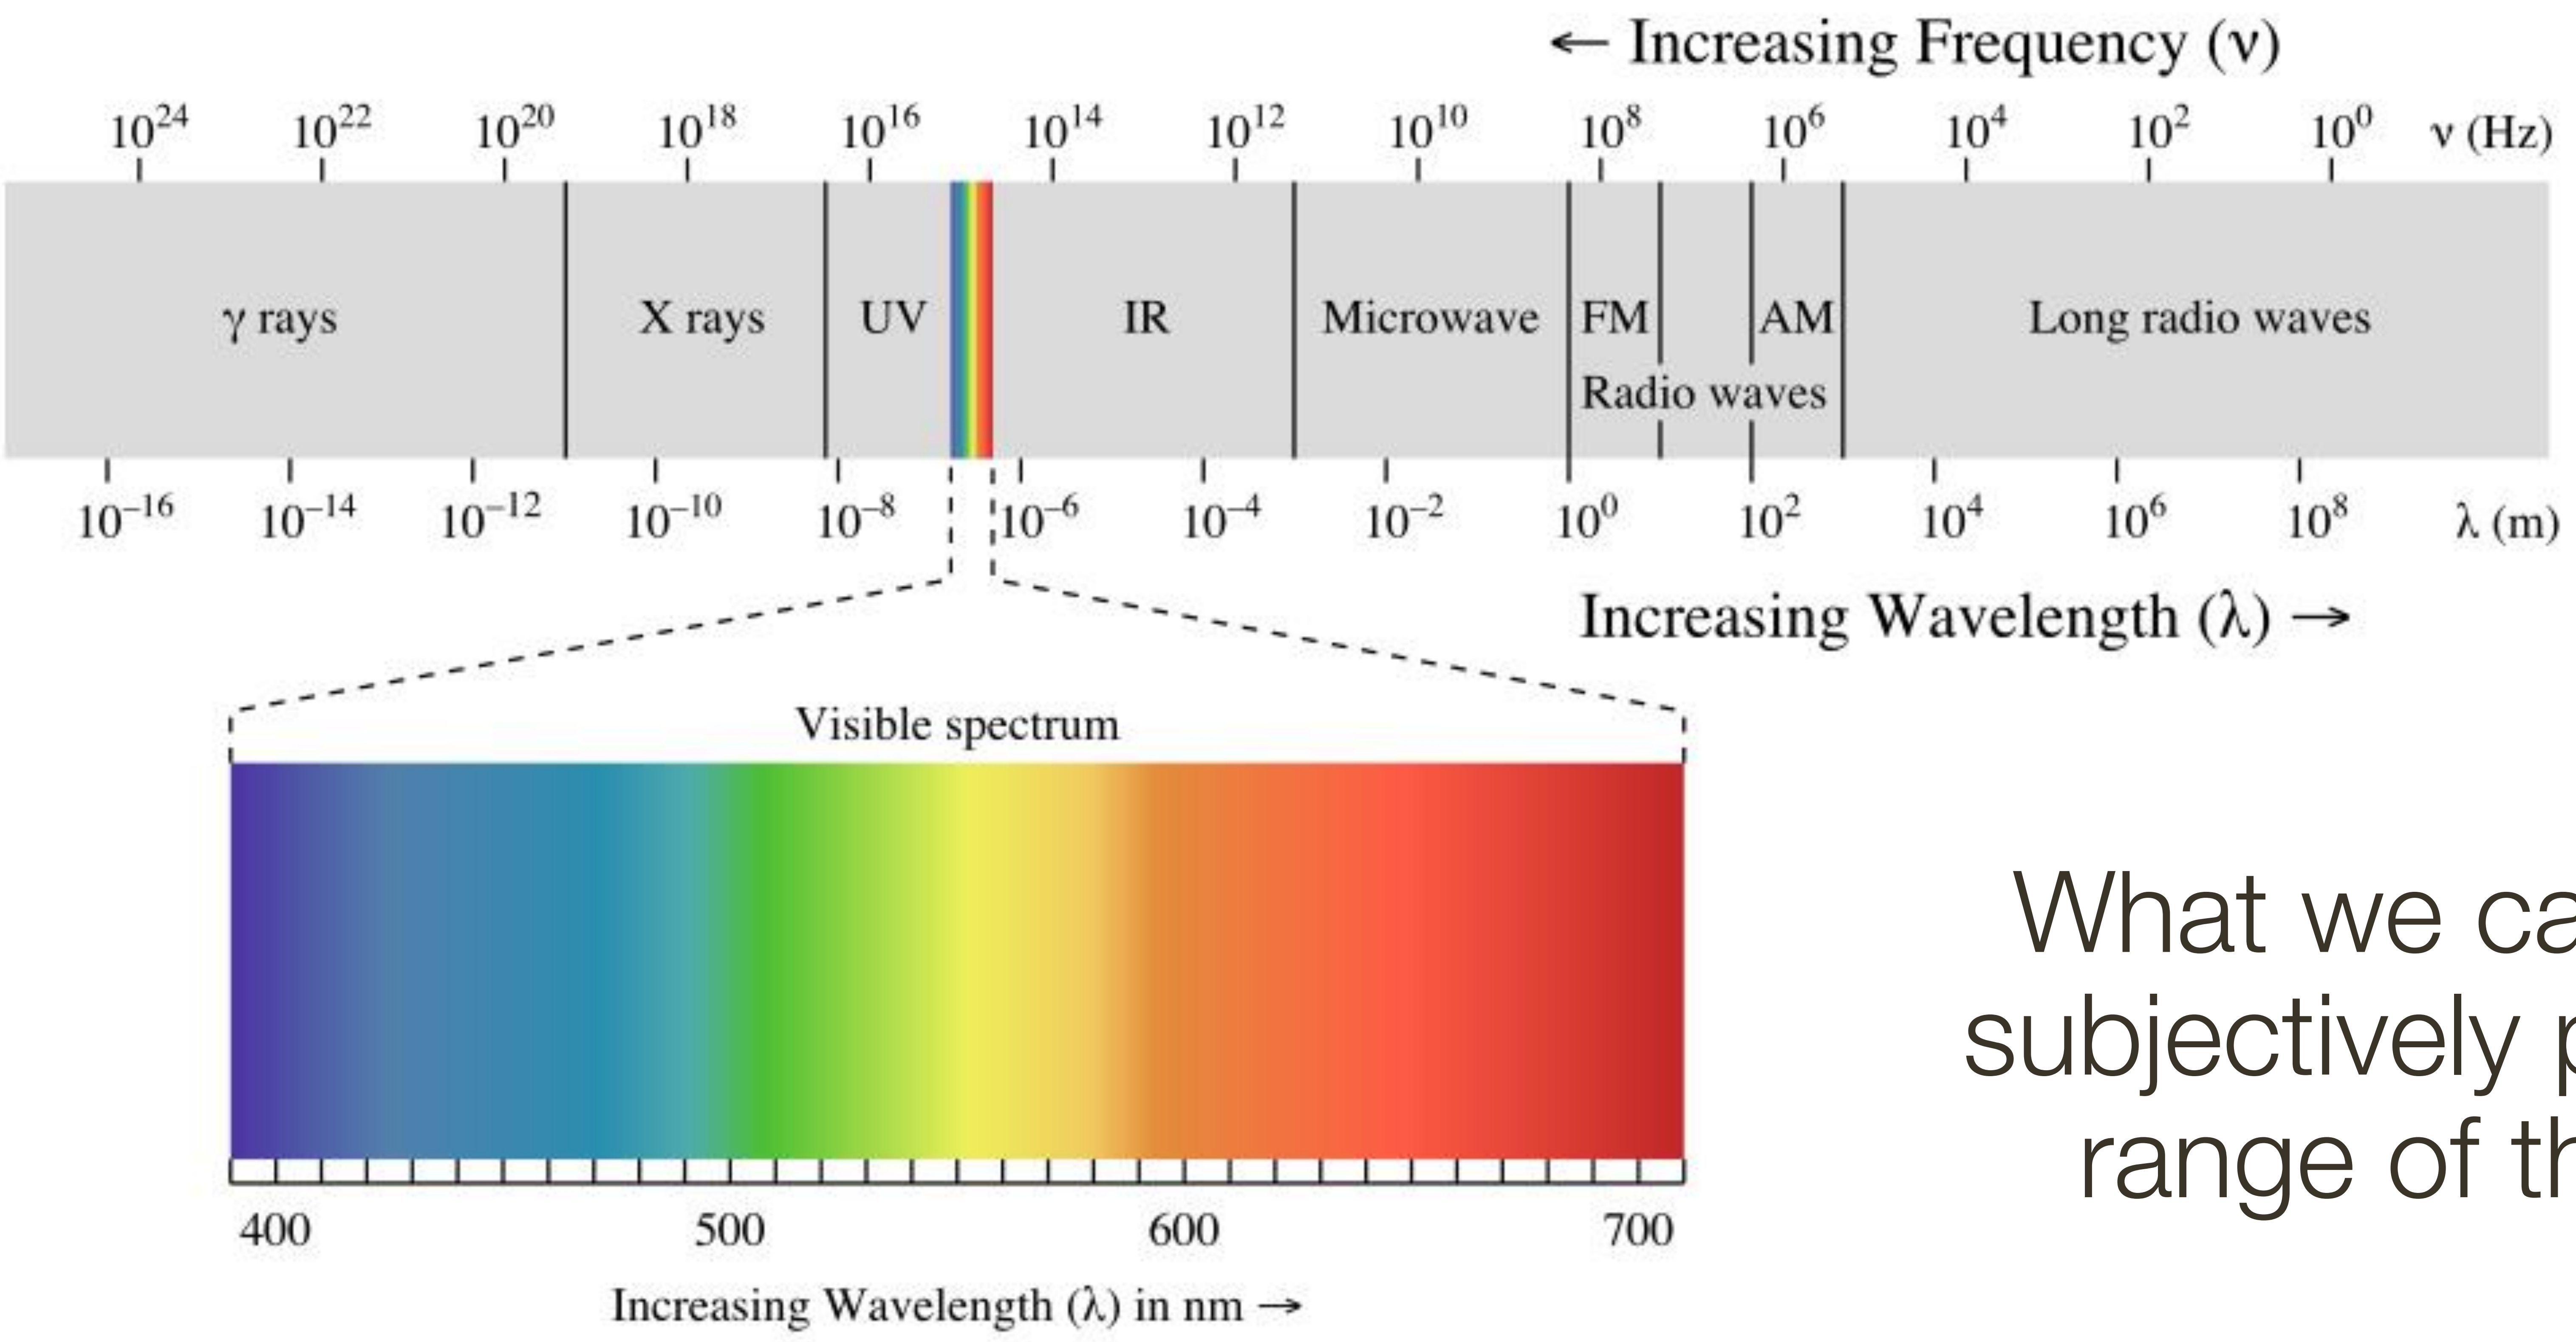

electromagnetic spectrum

What we call "color" is how we subjectively perceive a very small range of these wavelengths.

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Recall: Color is an Artifact of Human Perception

"Color" is **not** an objective physical property of light (electromagnetic radiation).
Instead, light is characterized by its wavelength.

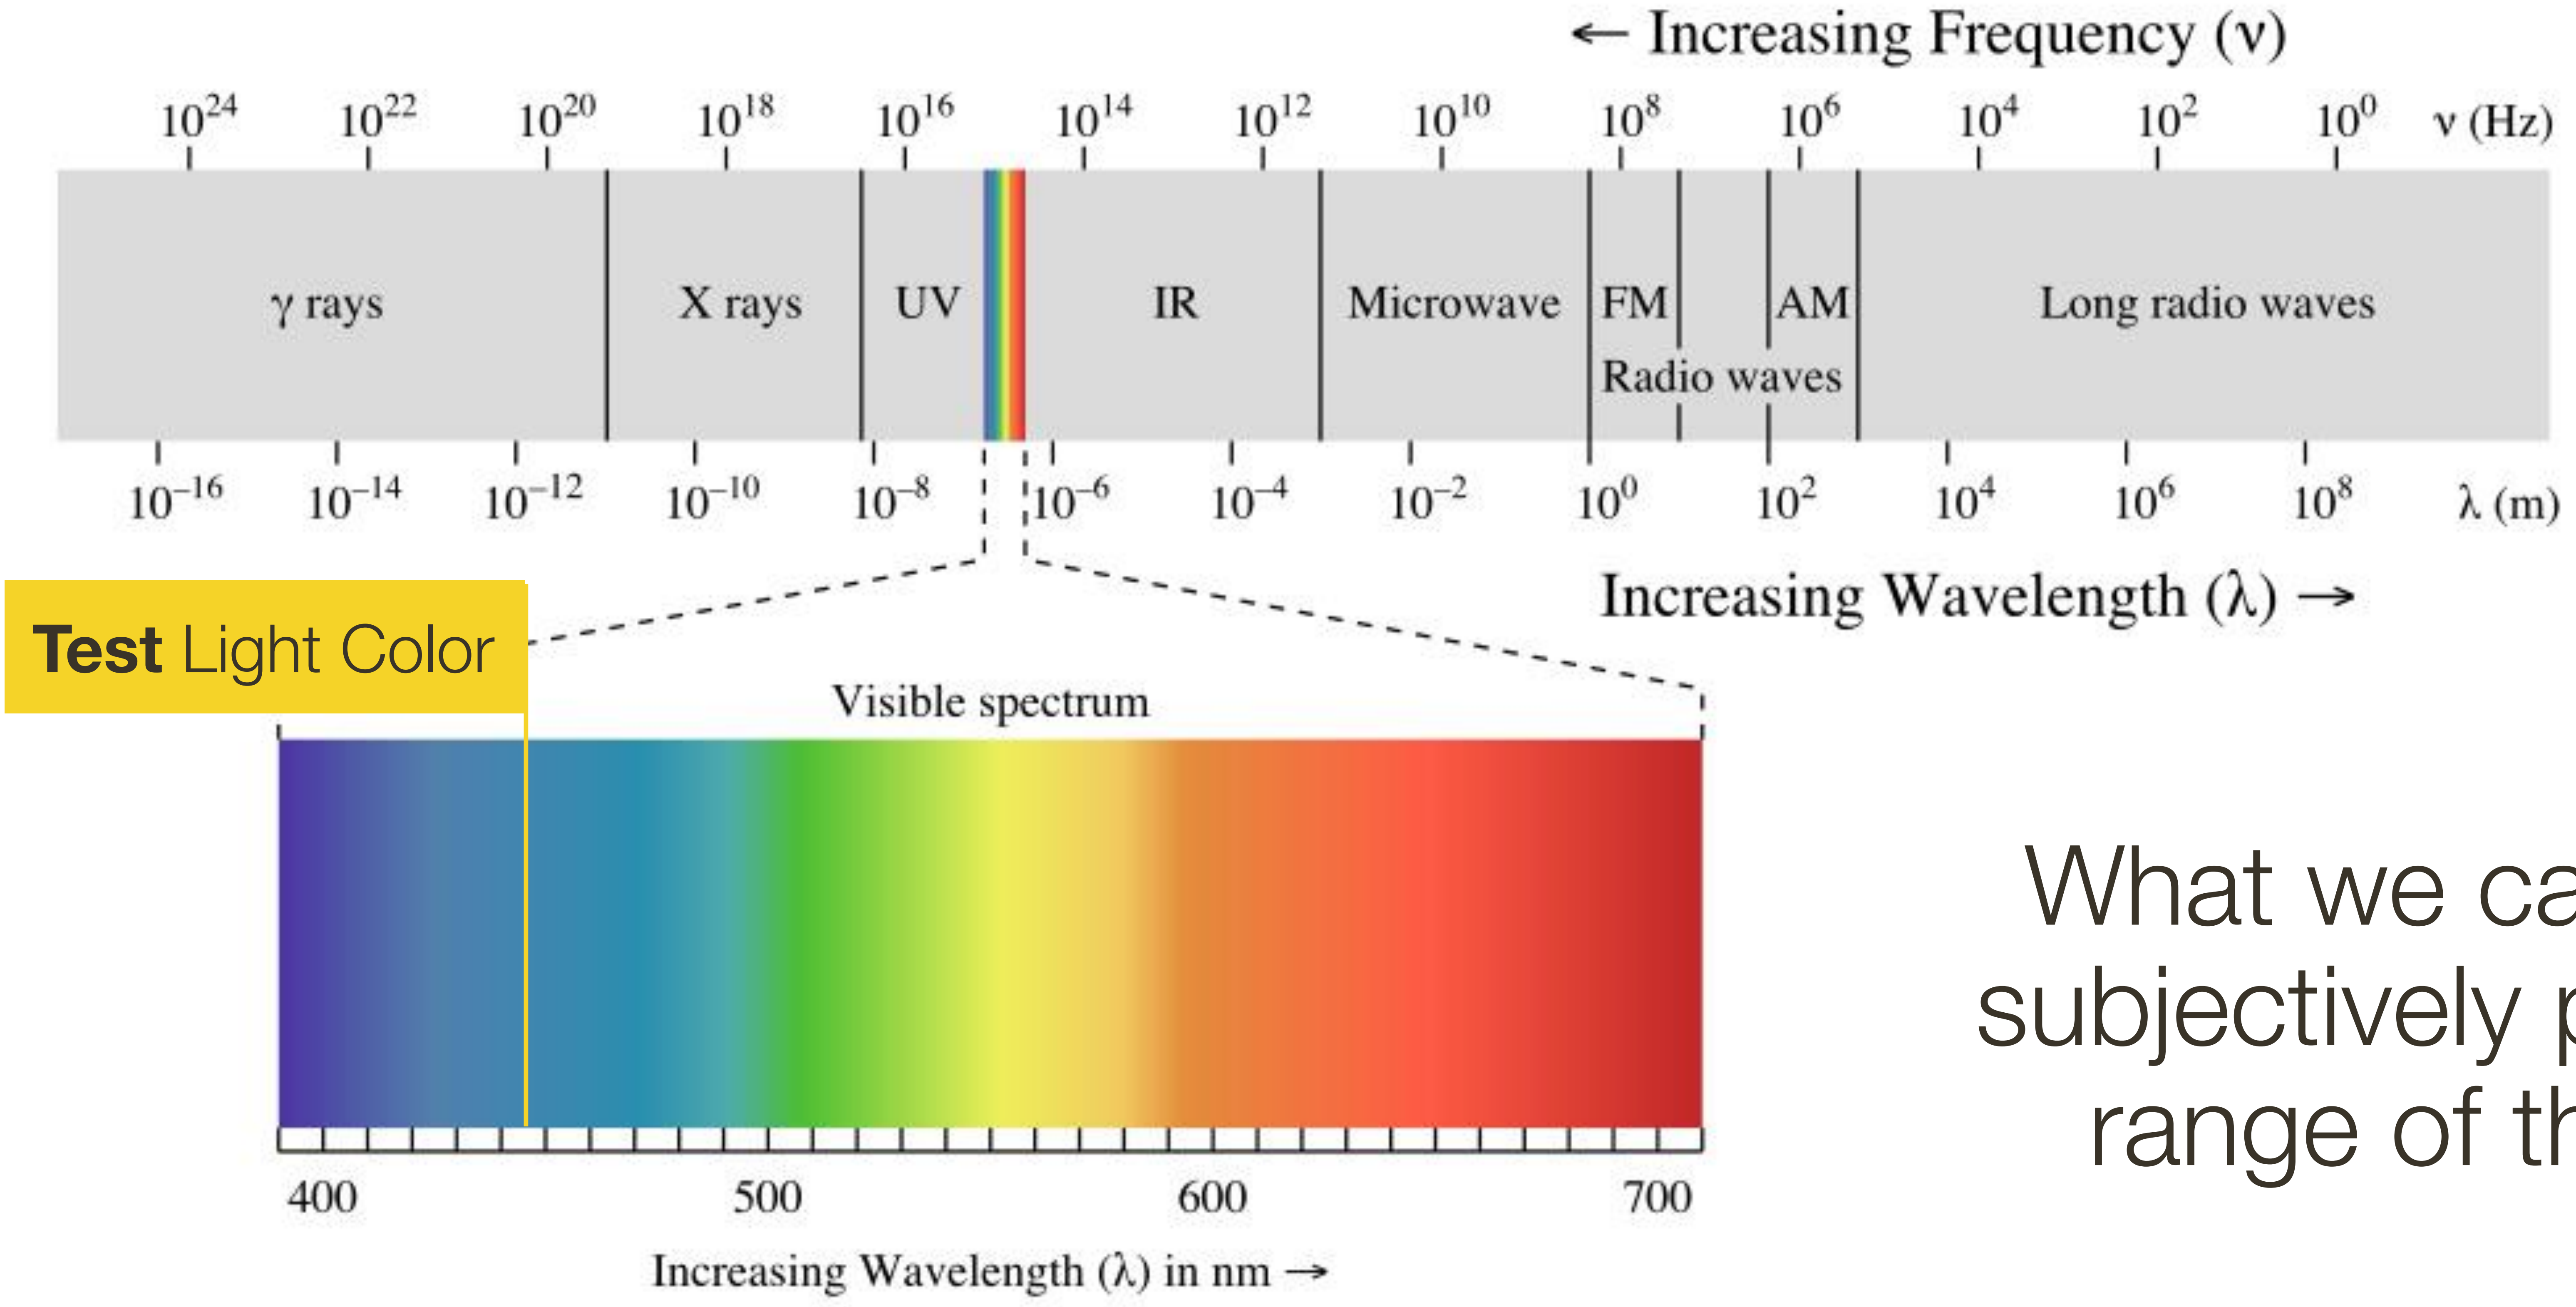

electromagnetic spectrum

What we call "color" is how we subjectively perceive a very small range of these wavelengths.

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Recall: Color is an Artifact of Human Perception

"Color" is **not** an objective physical property of light (electromagnetic radiation).
Instead, light is characterized by its wavelength.

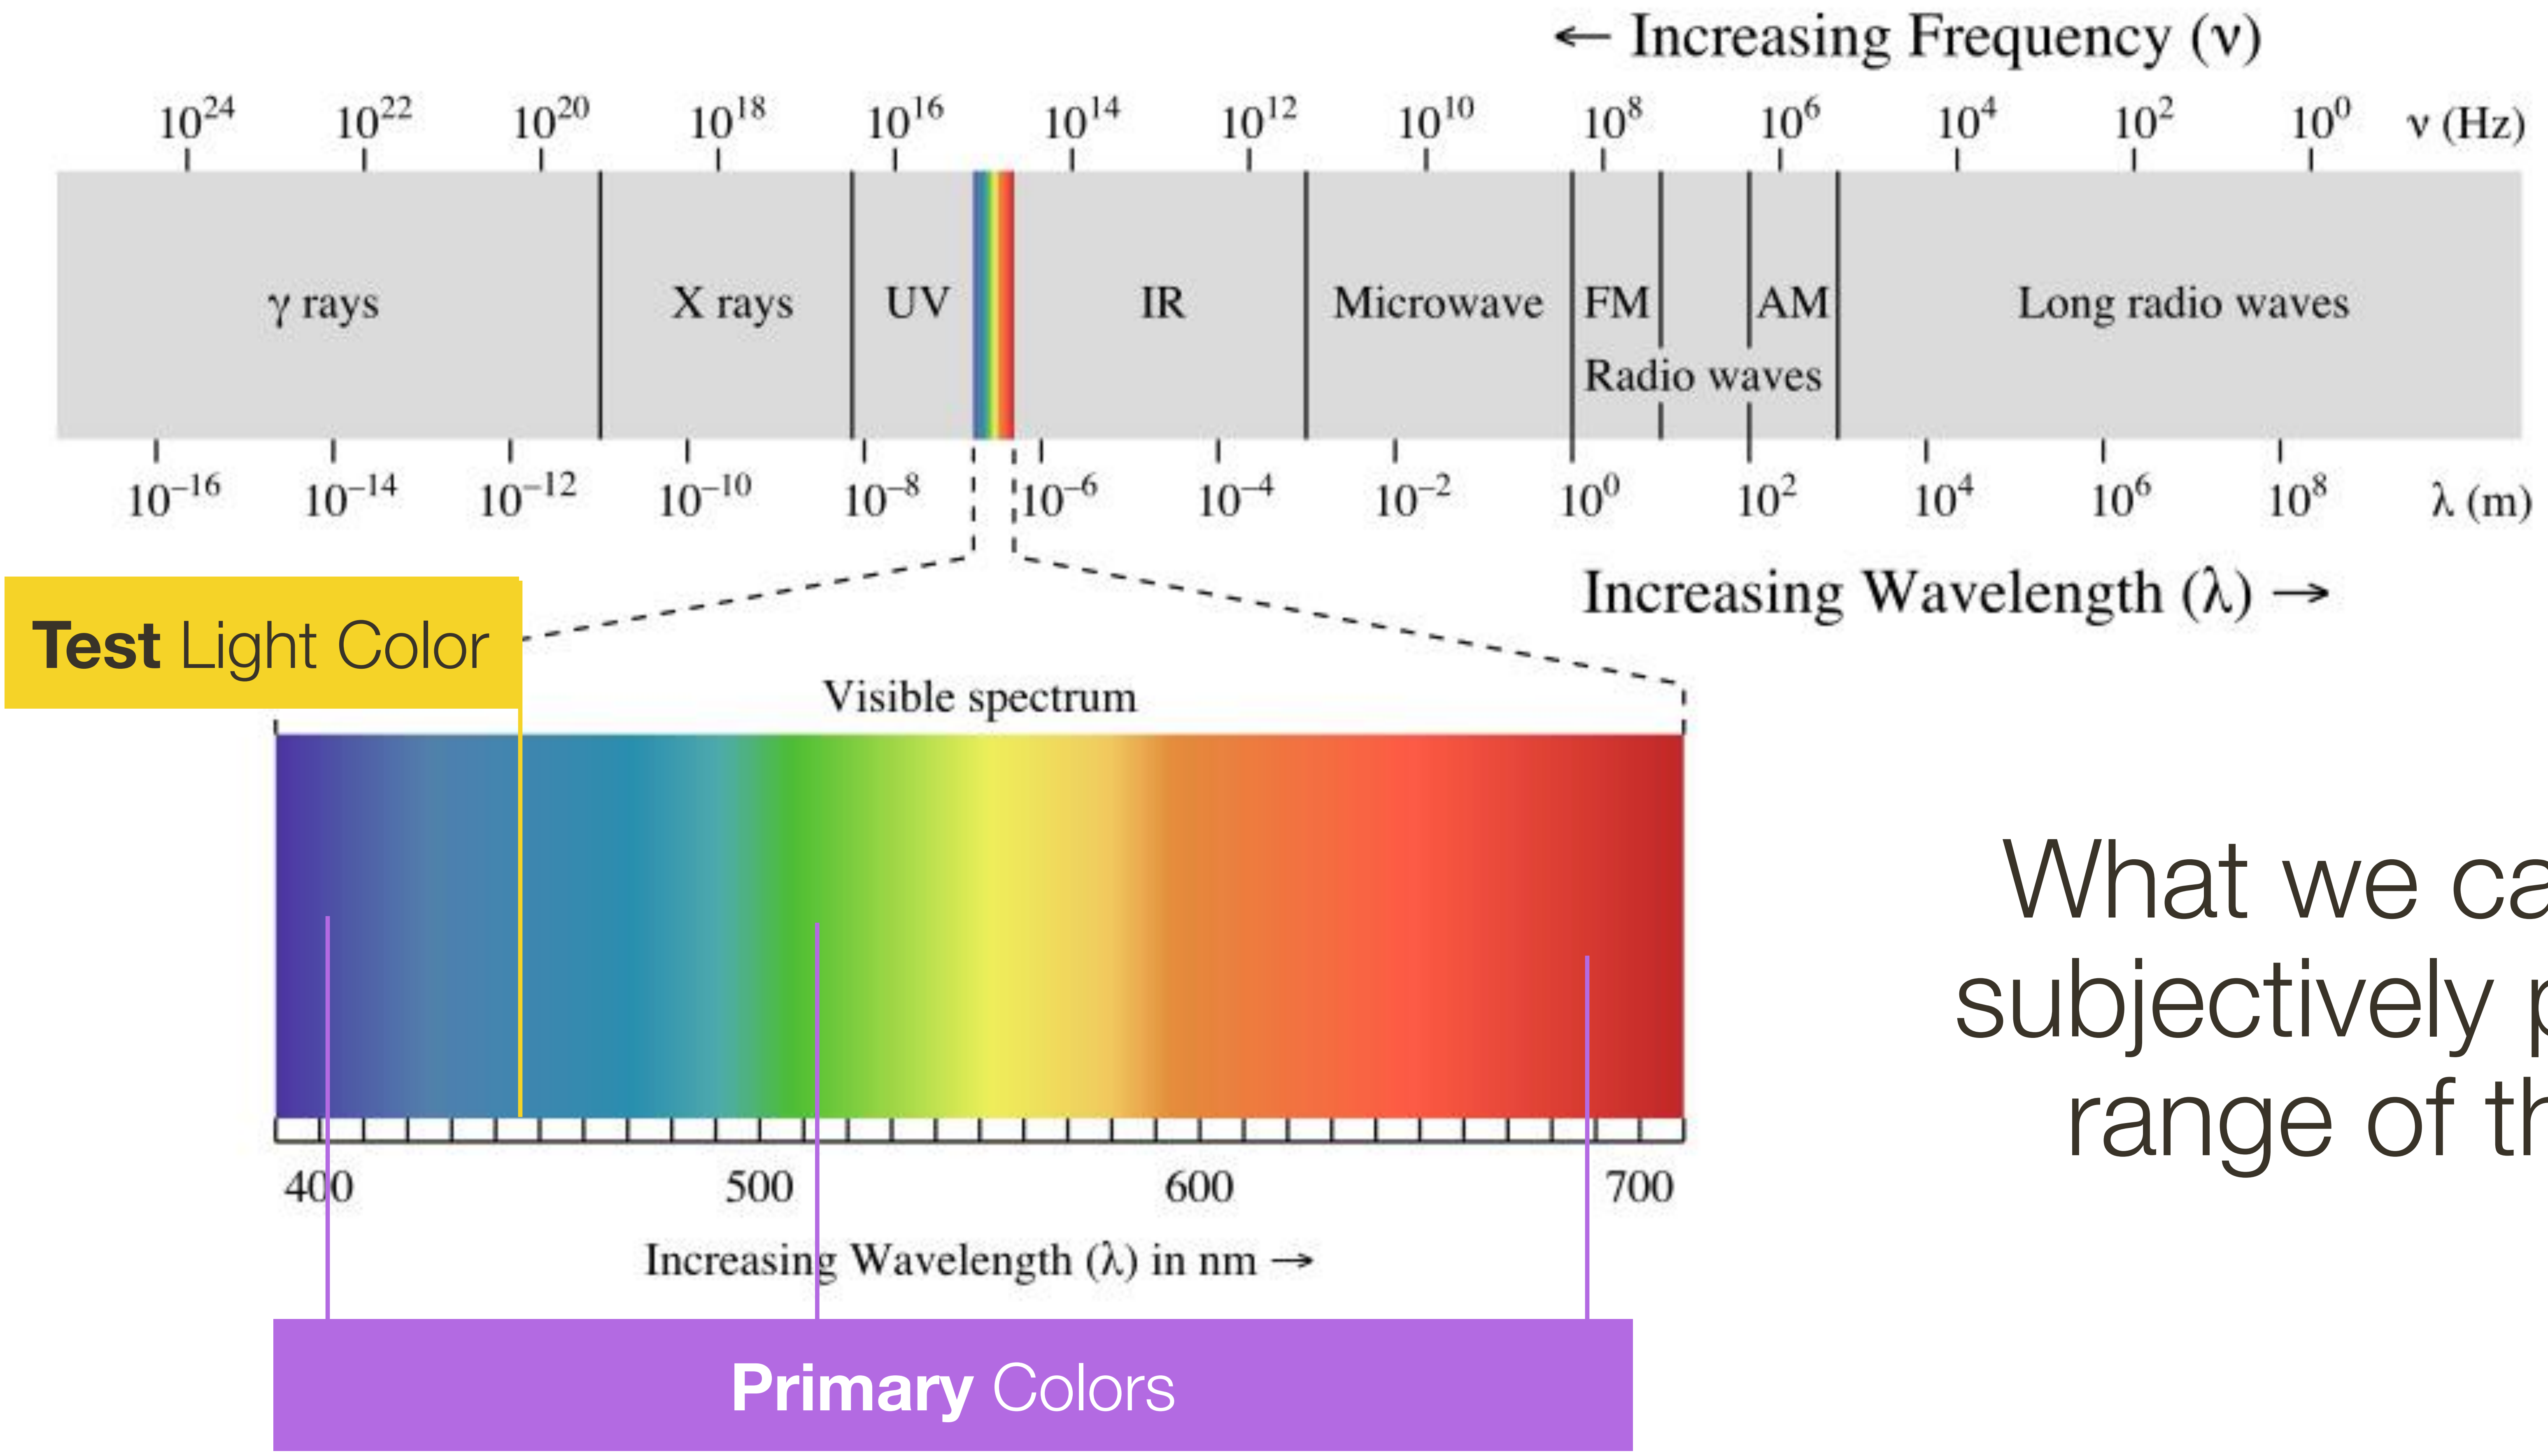

electromagnetic spectrum

What we call "color" is how we subjectively perceive a very small range of these wavelengths.

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Recall: Color is an Artifact of Human Perception

"Color" is **not** an objective physical property of light (electromagnetic radiation).

Instead, light is characterized by its wavelength.

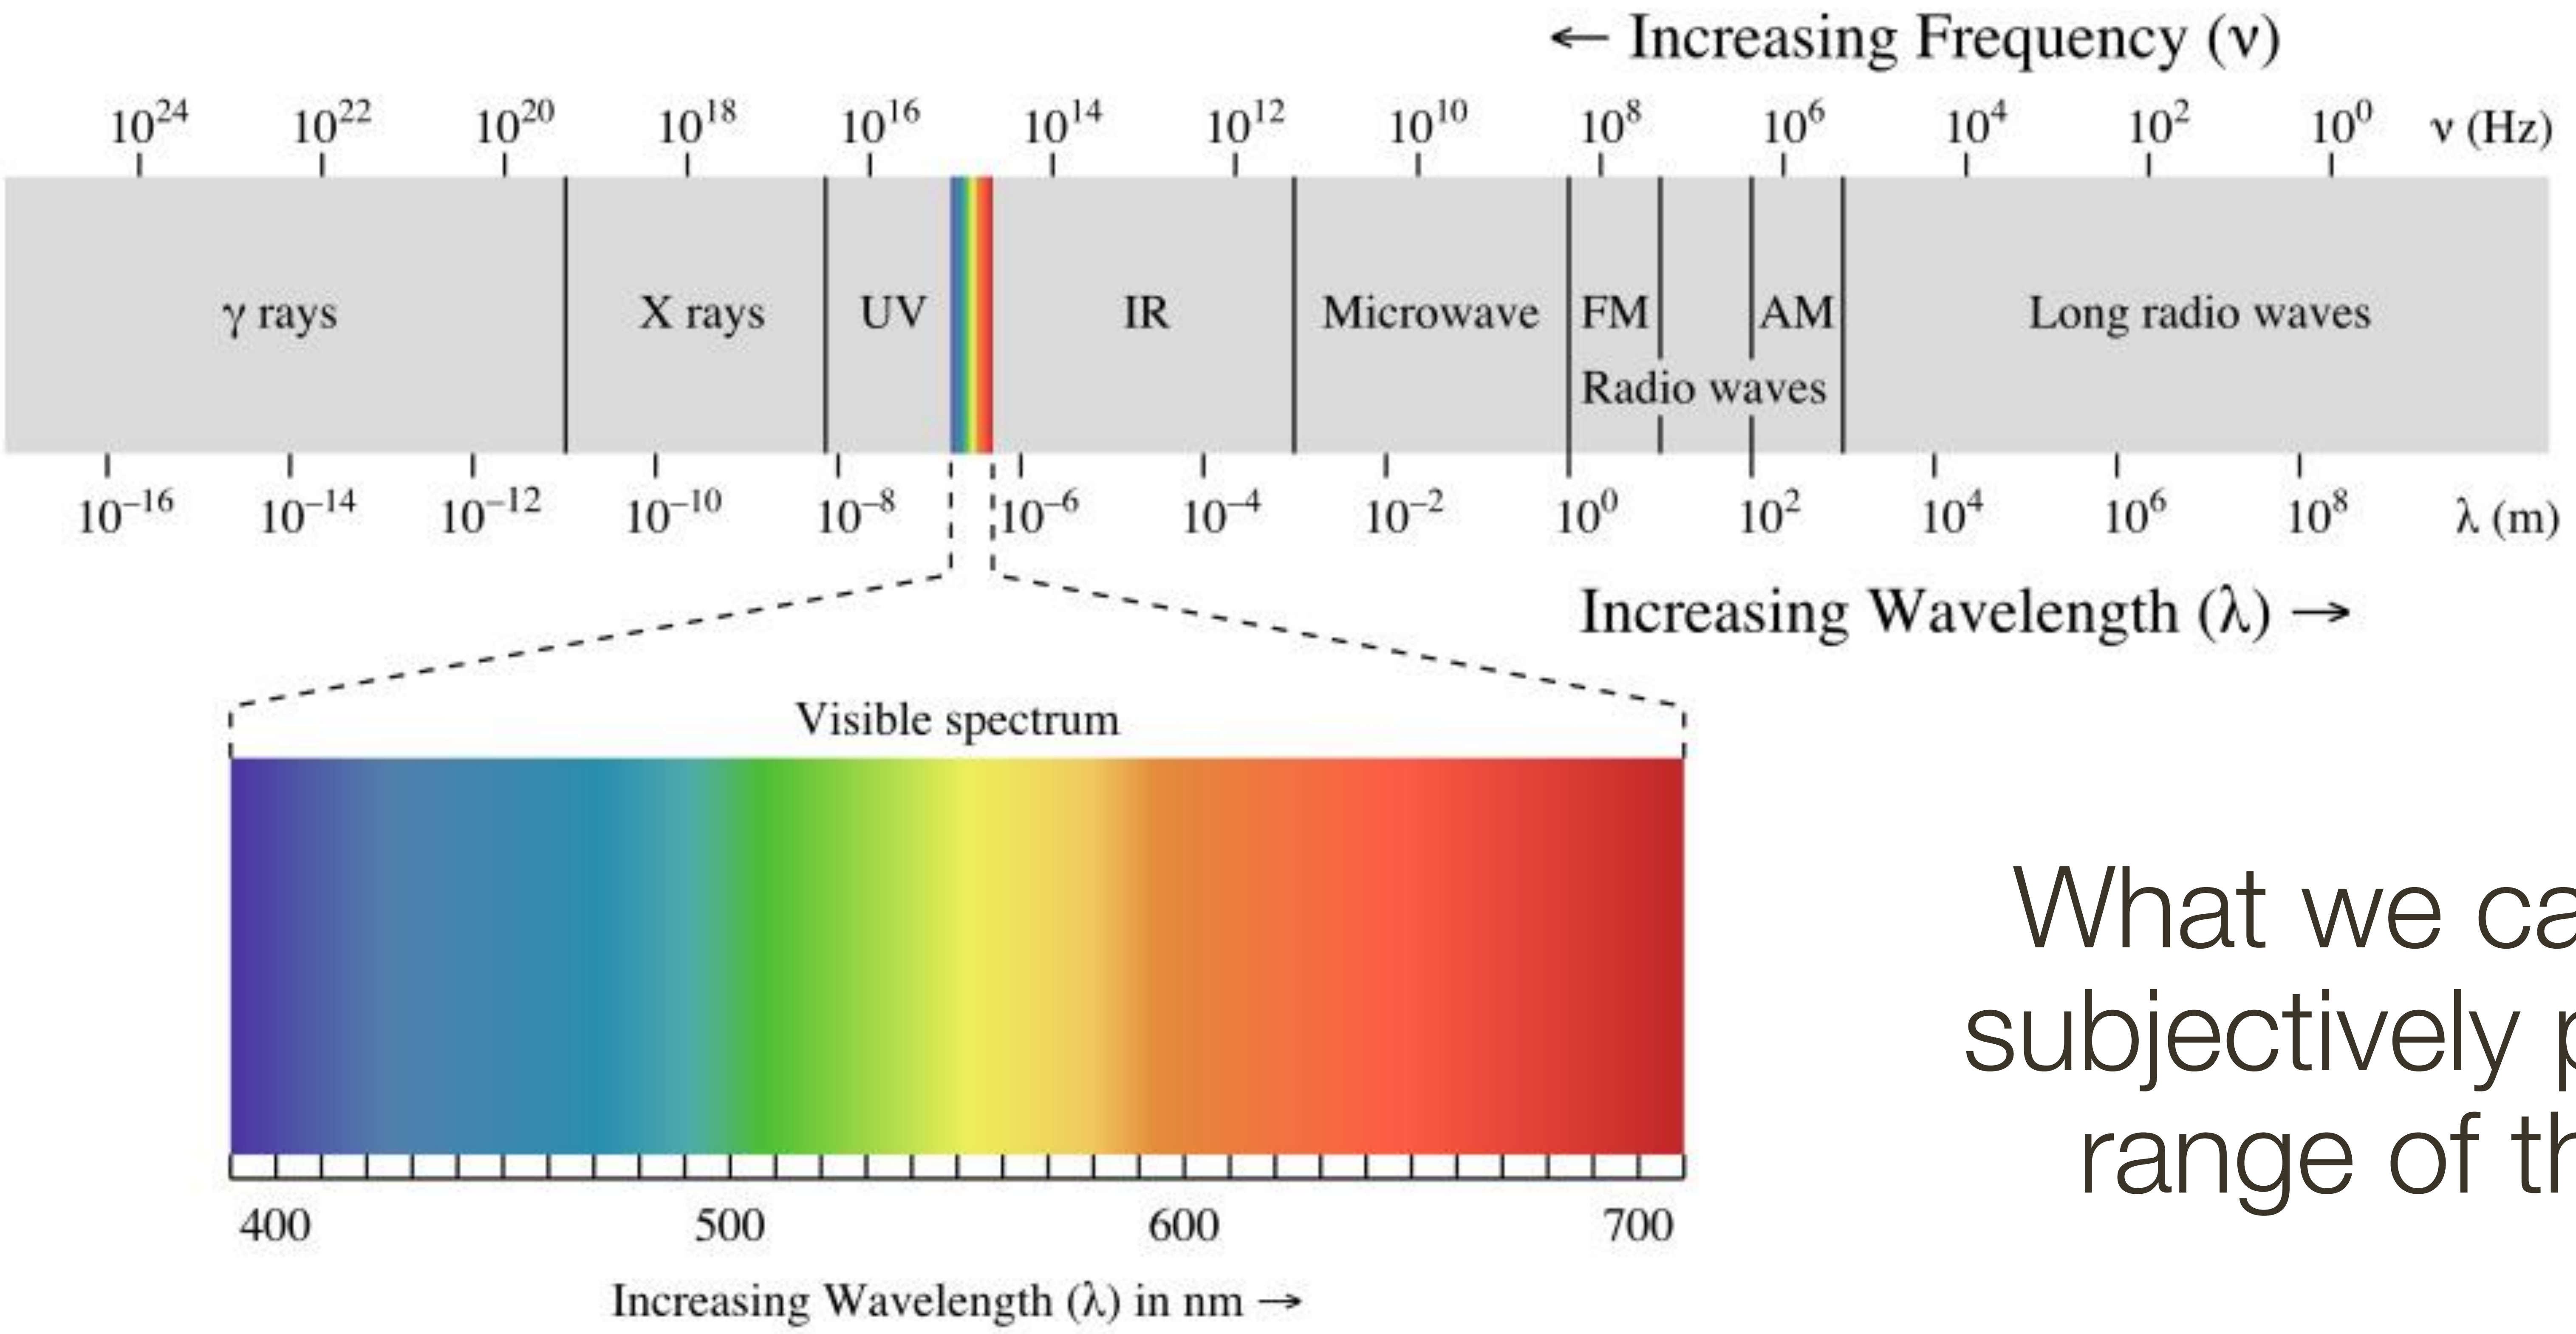

electromagnetic spectrum

What we call "color" is how we subjectively perceive a very small range of these wavelengths.

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Recall: Color is an Artifact of Human Perception

"Color" is **not** an objective physical property of light (electromagnetic radiation).
Instead, light is characterized by its wavelength.

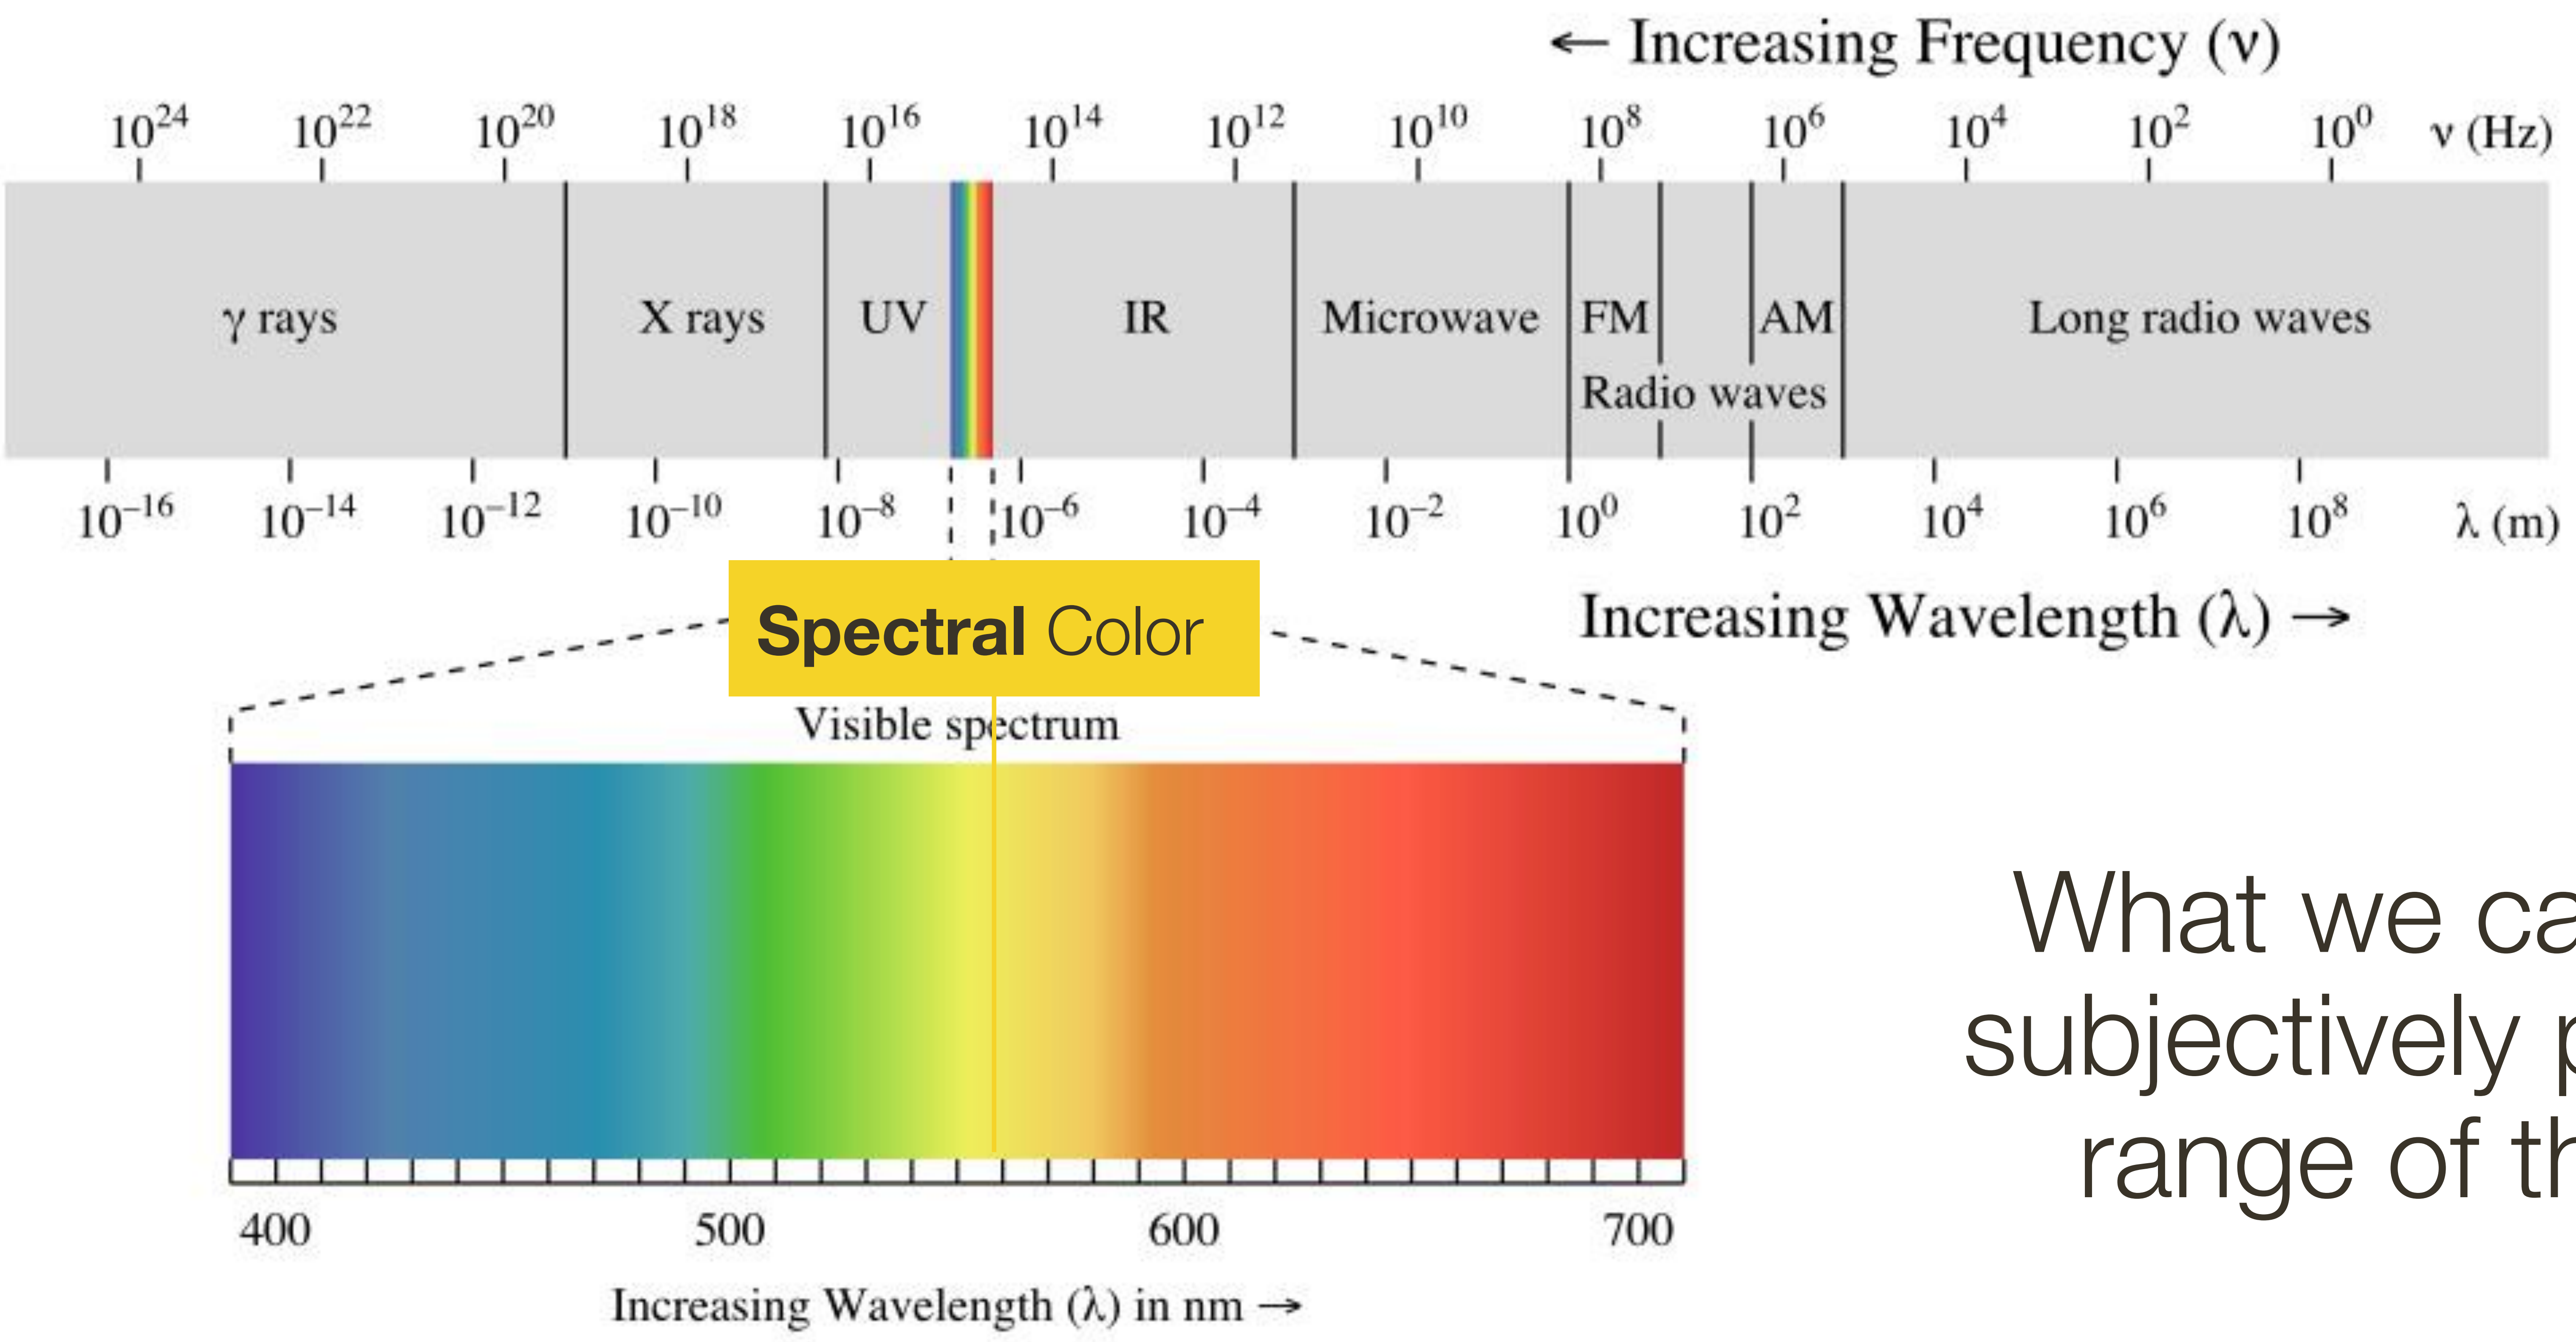

electromagnetic spectrum

What we call "color" is how we subjectively perceive a very small range of these wavelengths.

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Recall: Color is an Artifact of Human Perception

"Color" is **not** an objective physical property of light (electromagnetic radiation).
Instead, light is characterized by its wavelength.

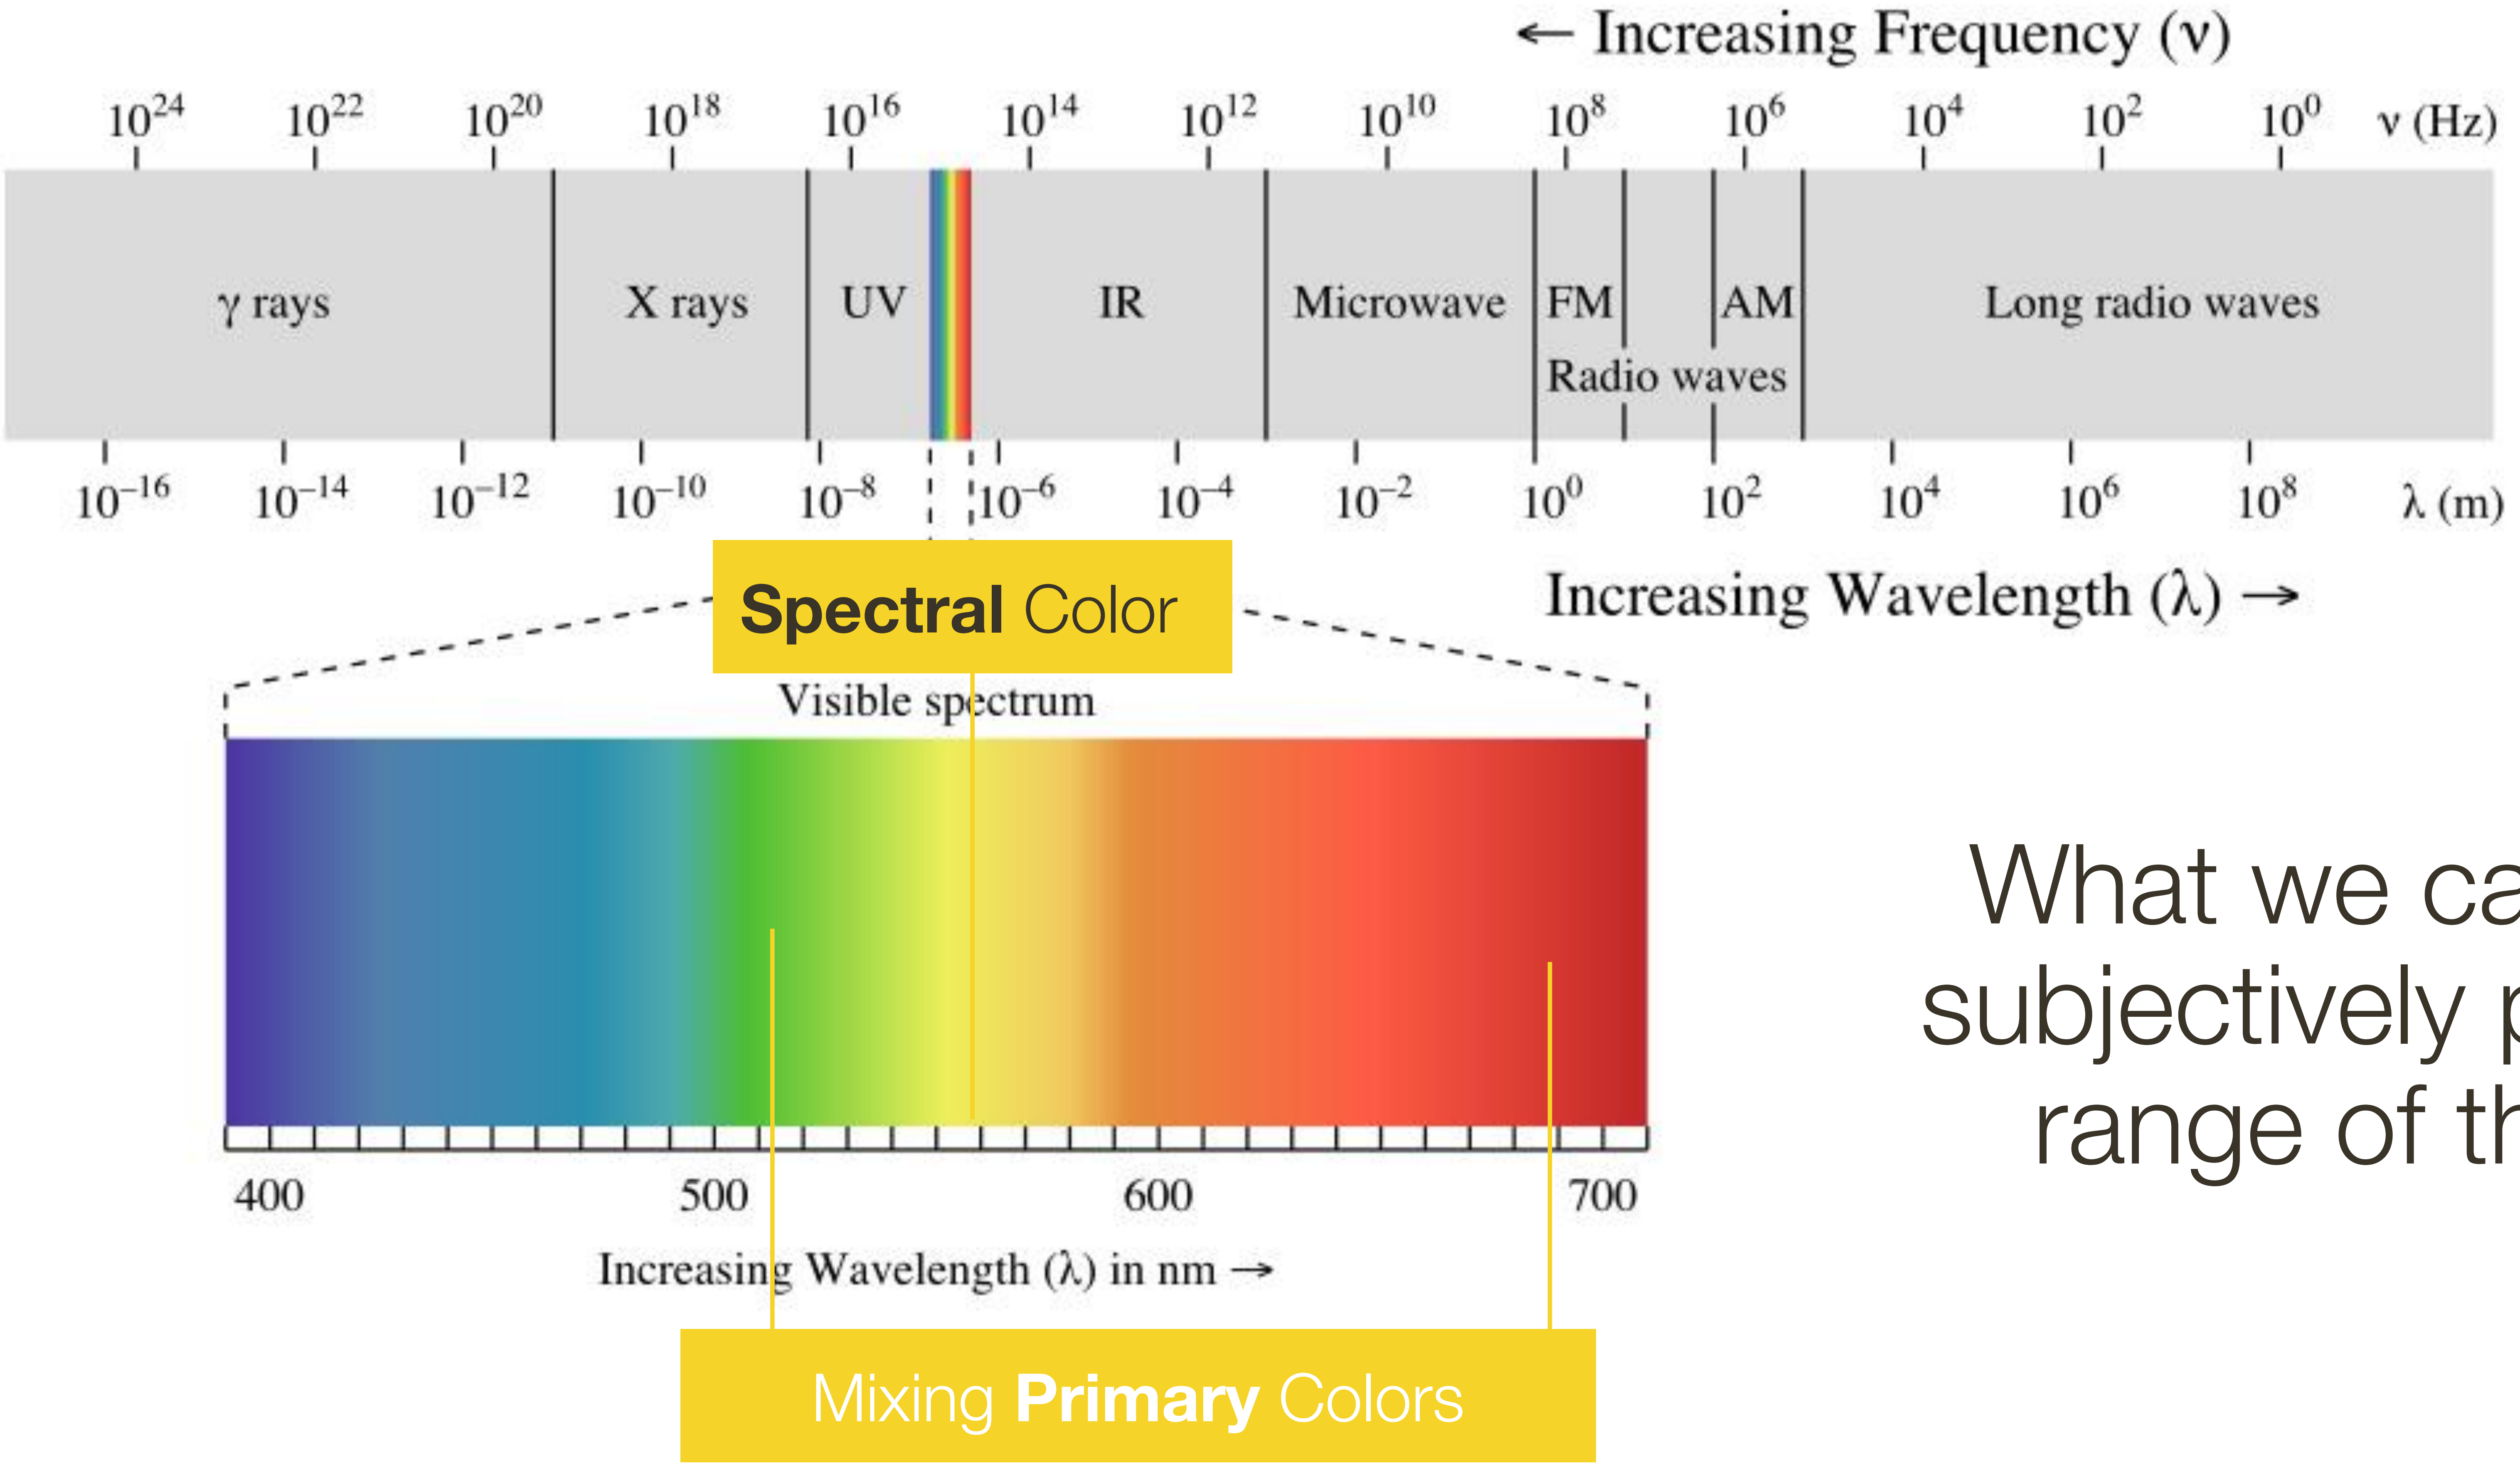

electromagnetic spectrum

What we call "color" is how we subjectively perceive a very small range of these wavelengths.

Slide Credit: Ioannis (Yannis) Gkioulekas (CMU)

Color Matching Experiments

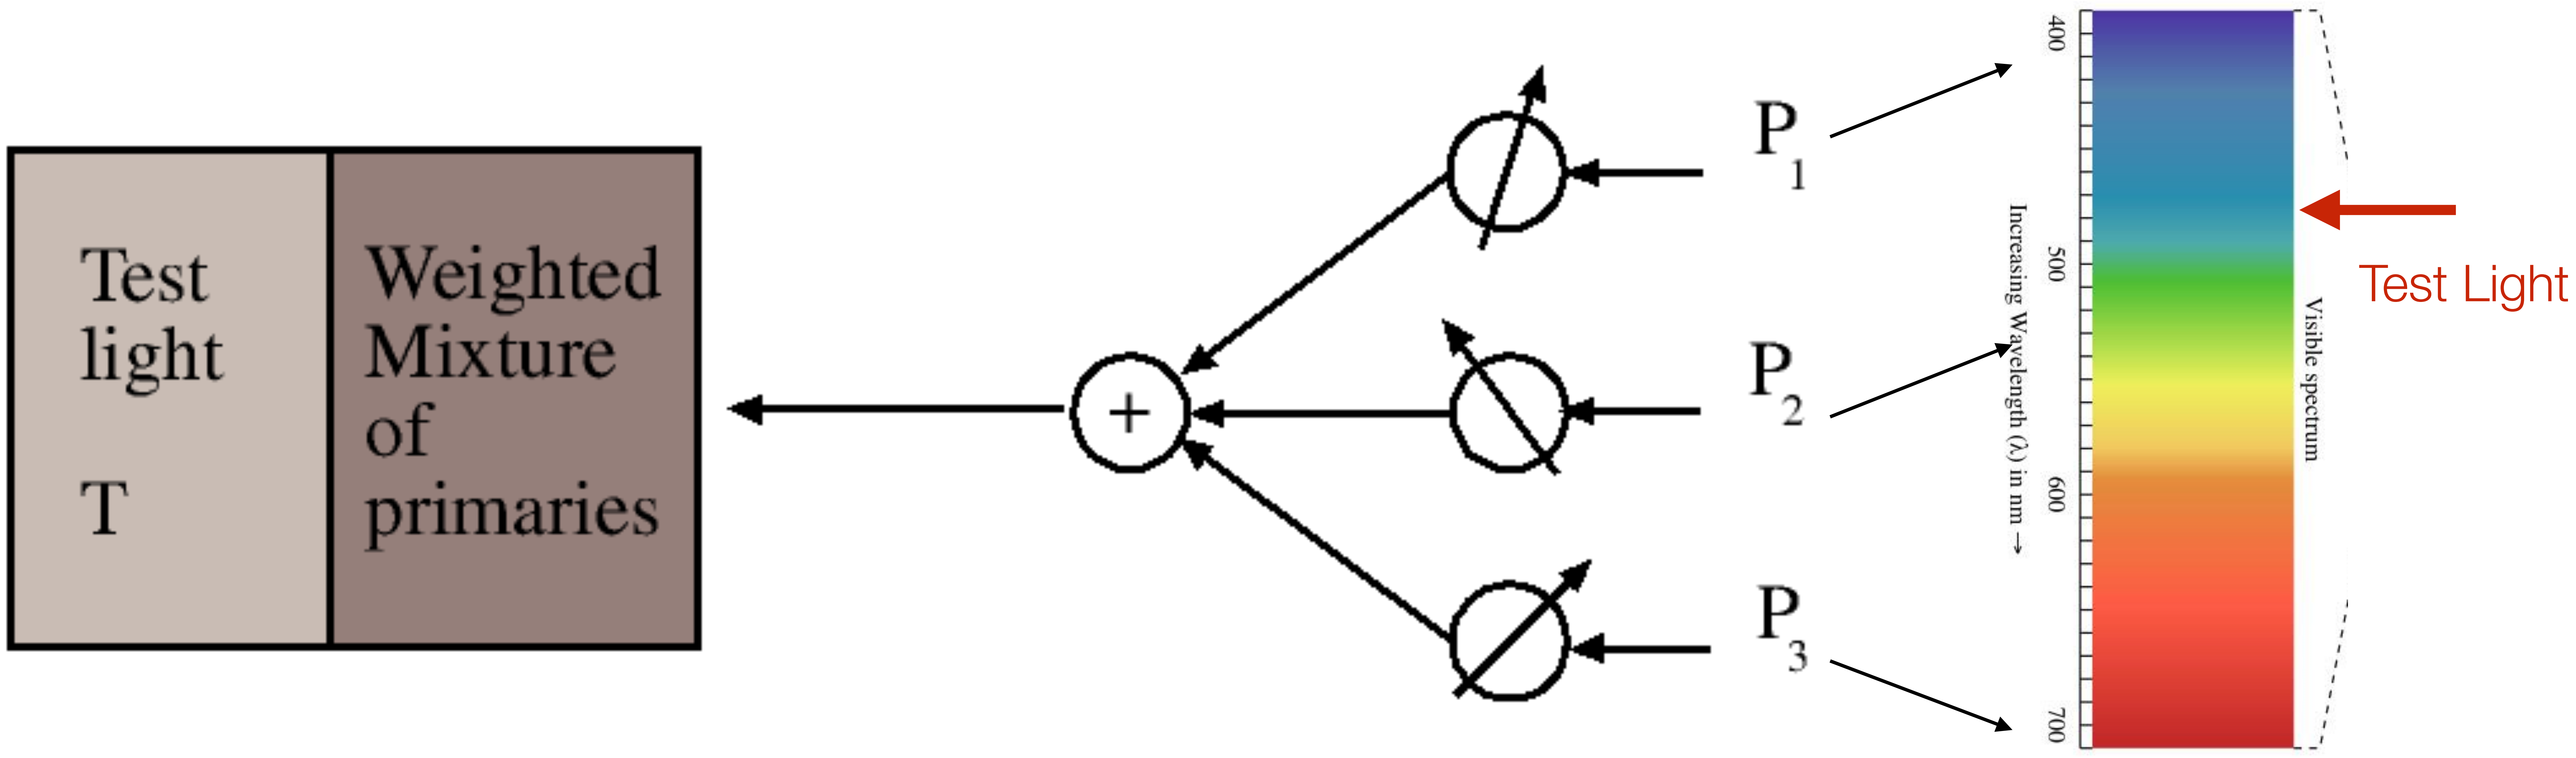

Forsyth & Ponce (2nd ed.) Figure 3.2

Show a split field to subjects. One side shows the light whose colour one wants to match. The other a weighted mixture of three primaries (fixed lights)

$$T = w_1 P_1 + w_2 P_2 + w_3 P_3$$

Maxwell Colour Matching Experiments

Maxwell mixed colours by rapidly spinning a top with different fractions of primaries, e.g., to match a central colour

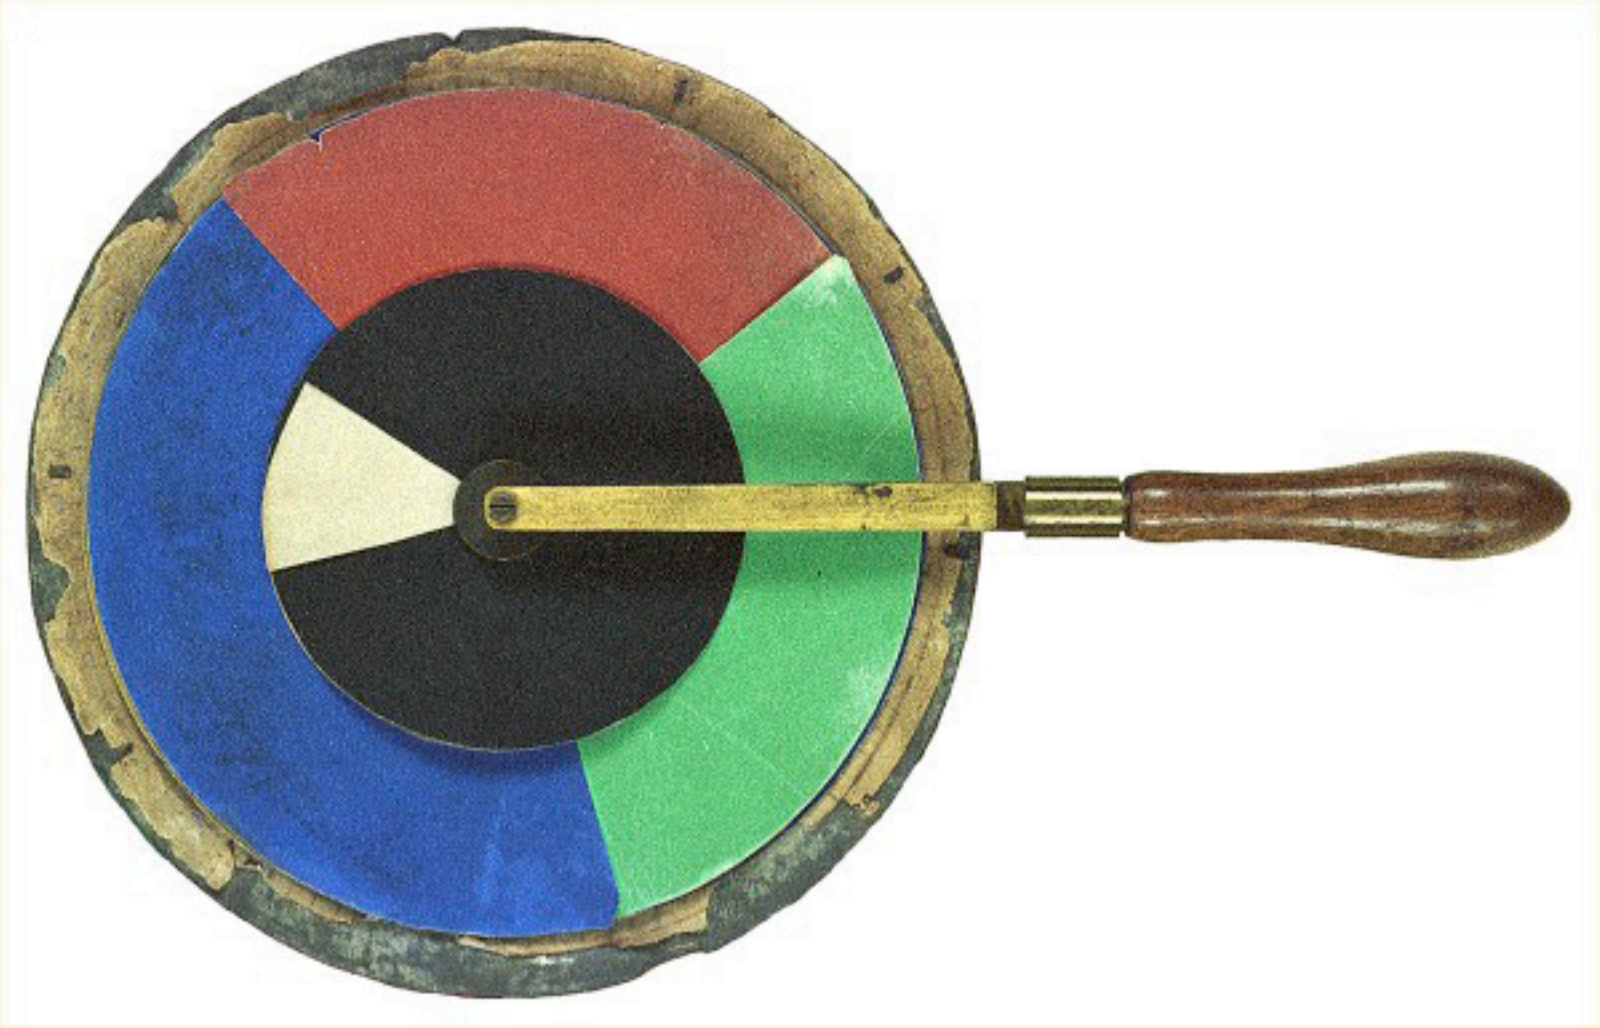

https://designblog.rietveldacademie.nl/?p=68422

Color Matching Experiments

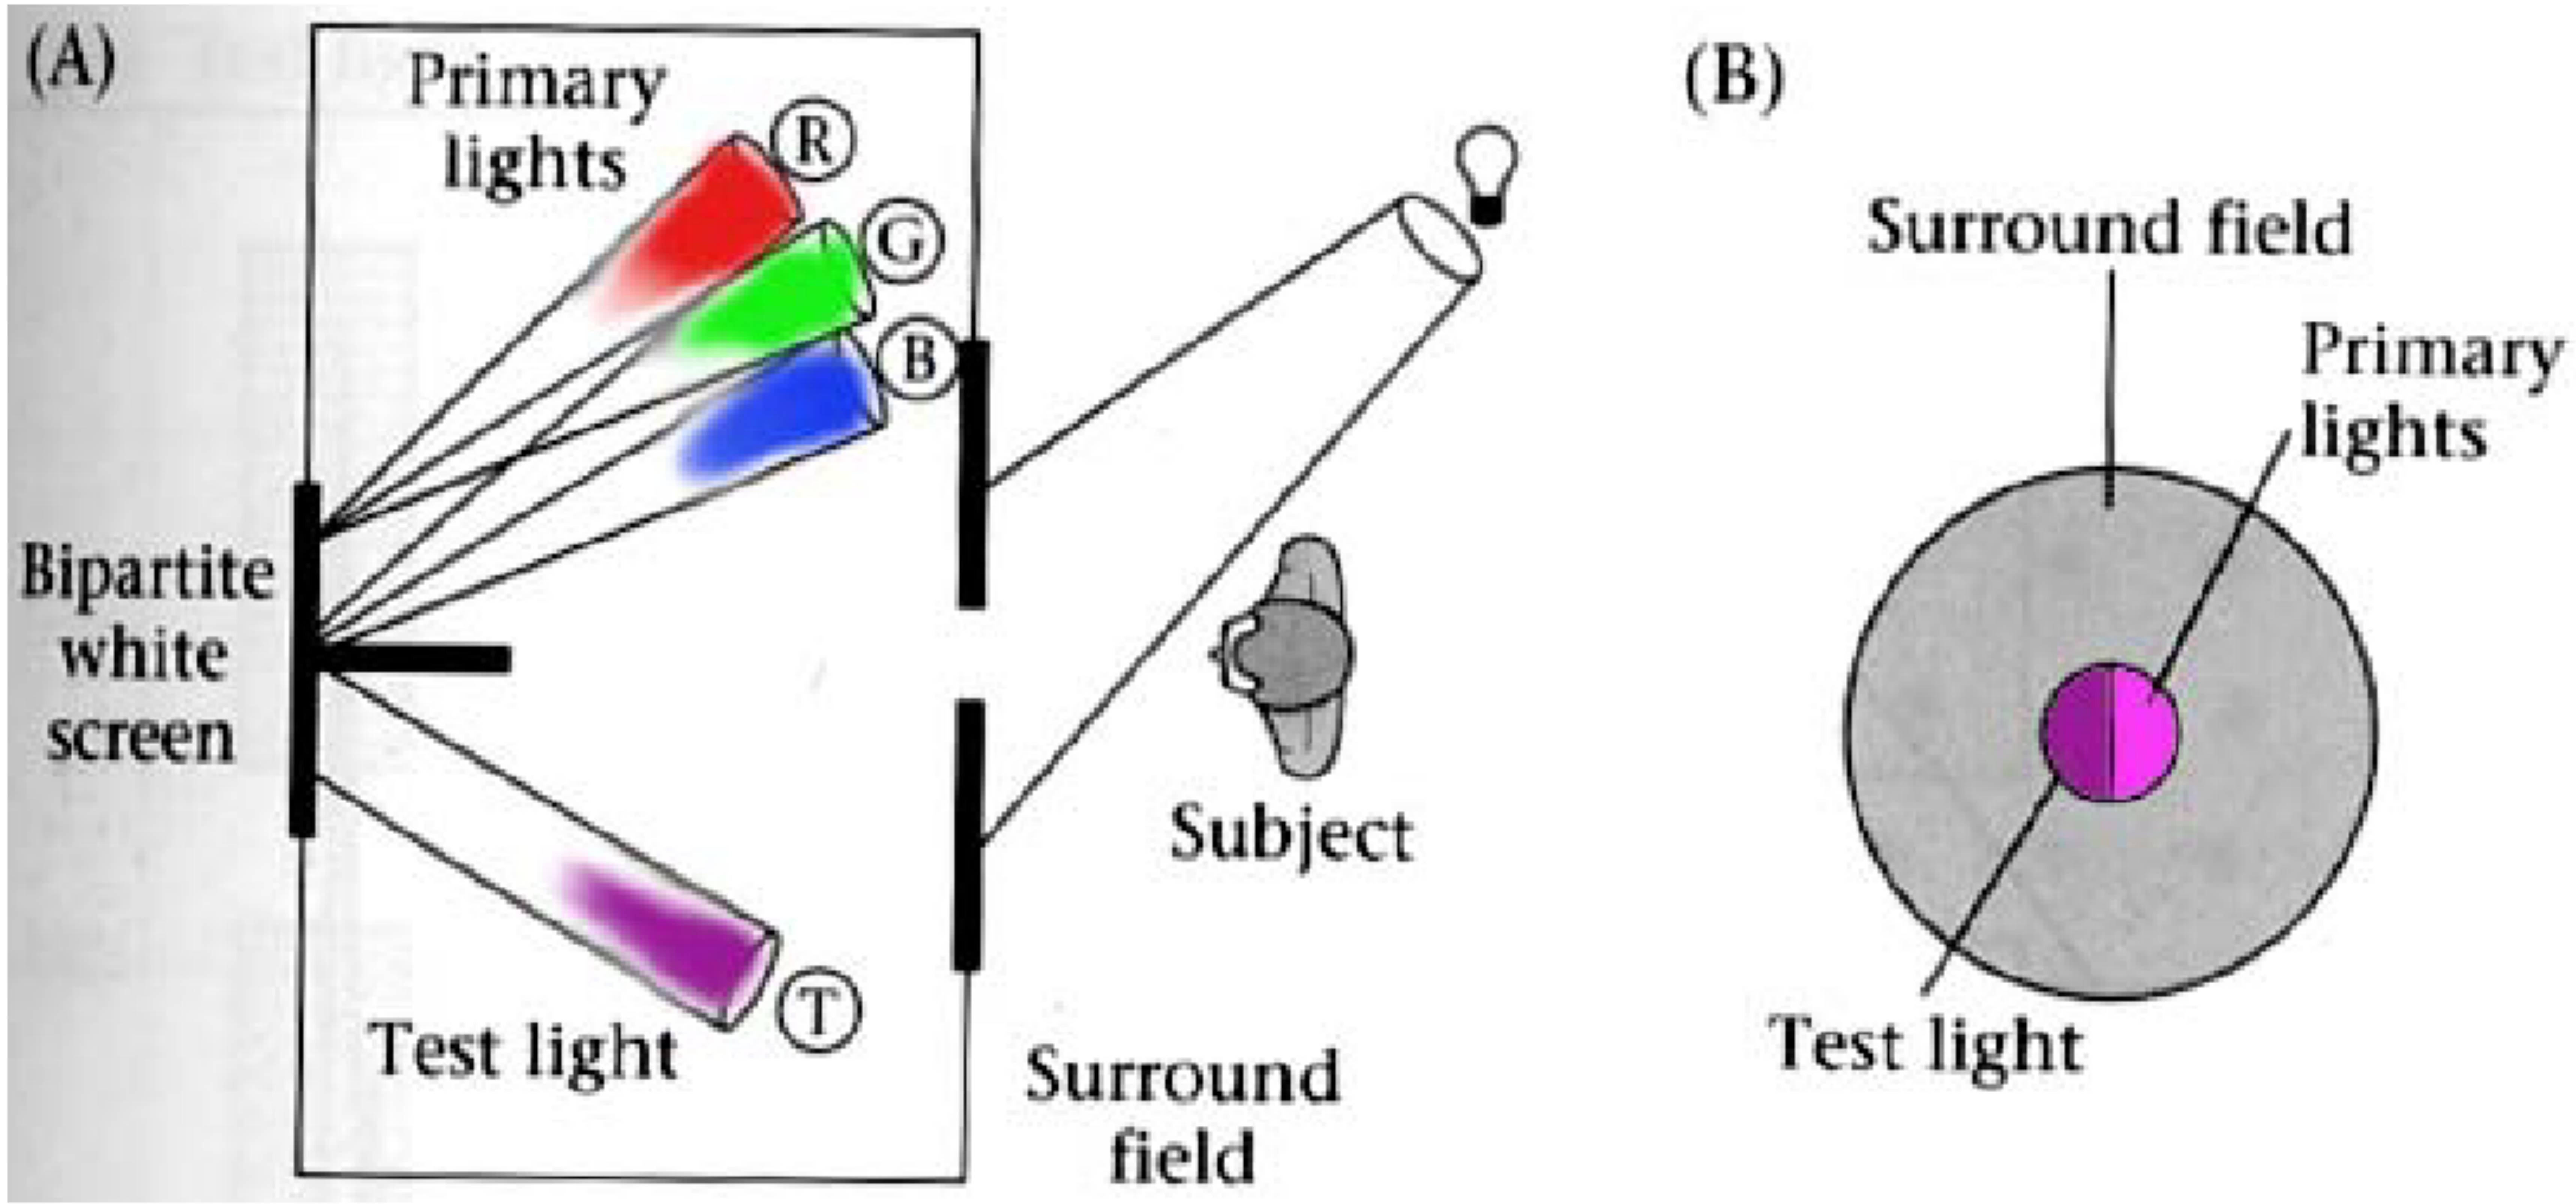

Figure Credit: Brian Wandell, Foundations of Vision, Sinauer Associates, 1995

Example 1: Color Matching Experiment

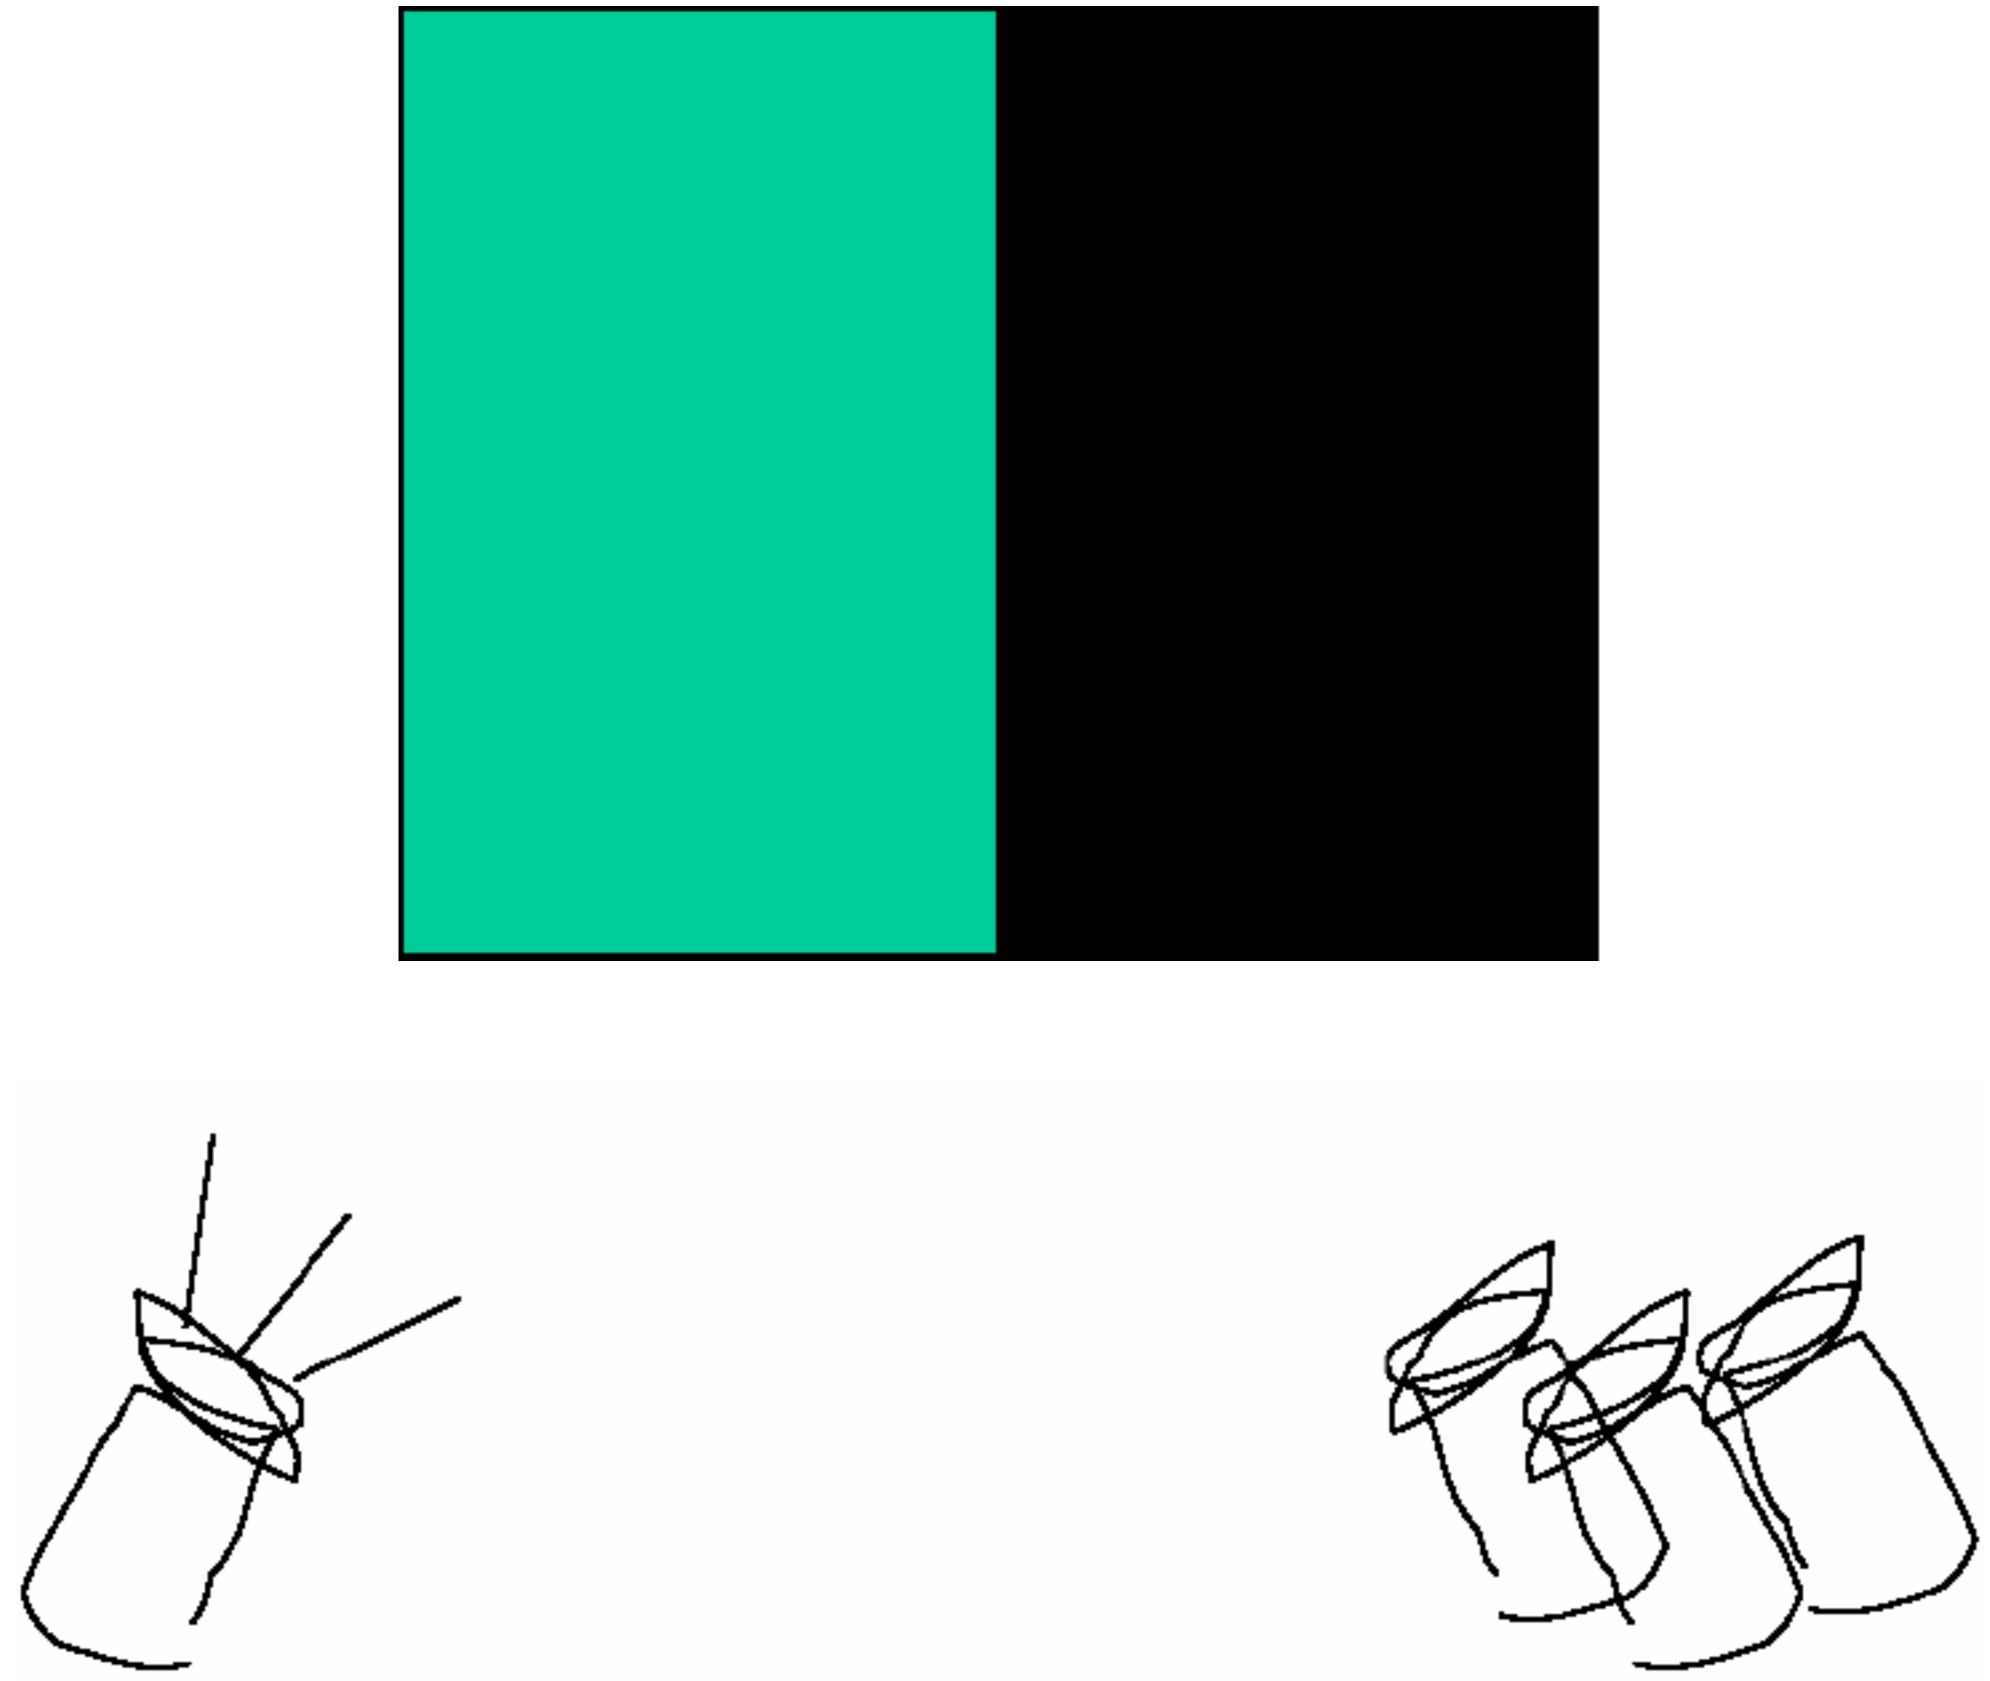

knobs here

Example Credit: Bill Freeman

Example 1: Color Matching Experiment

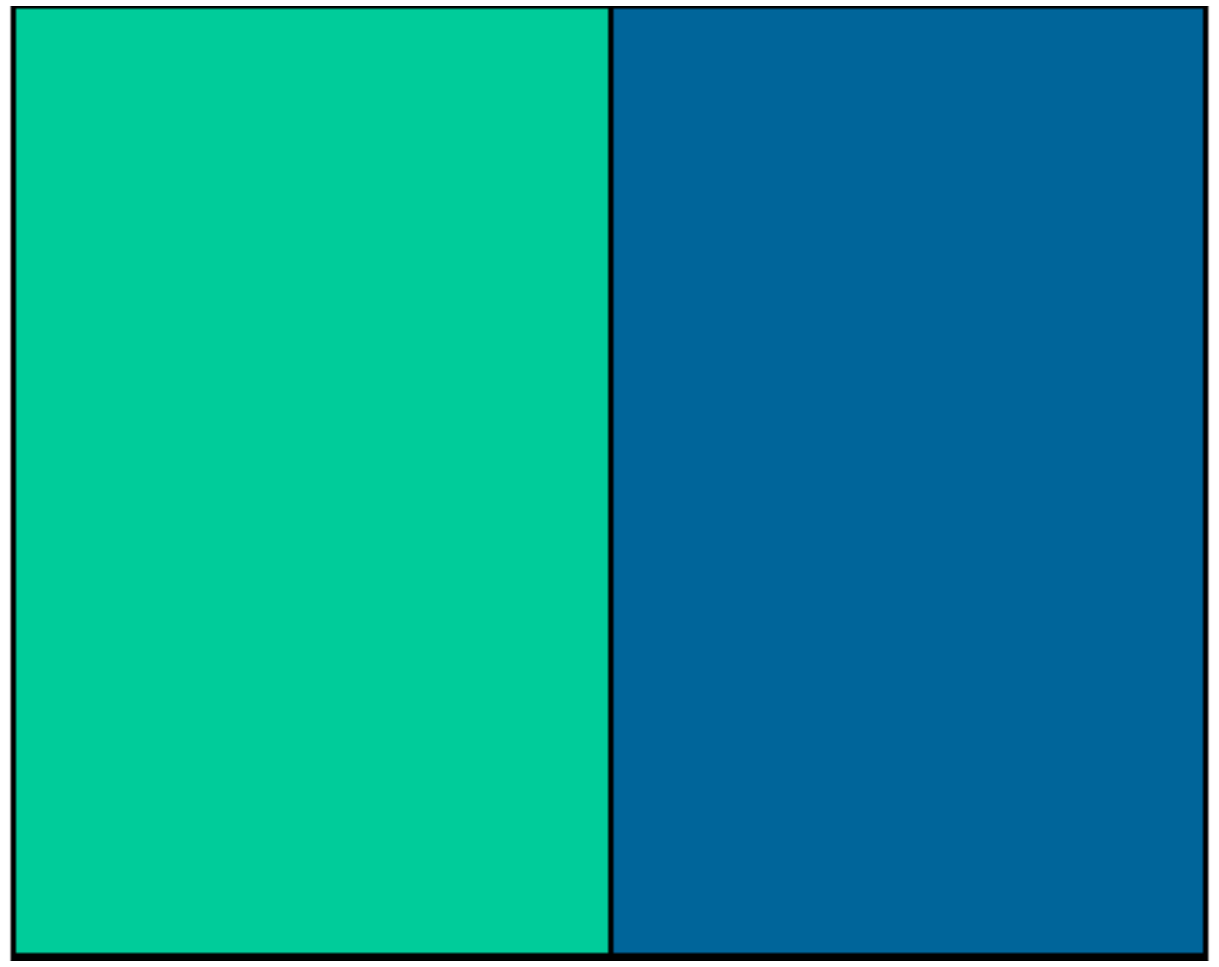

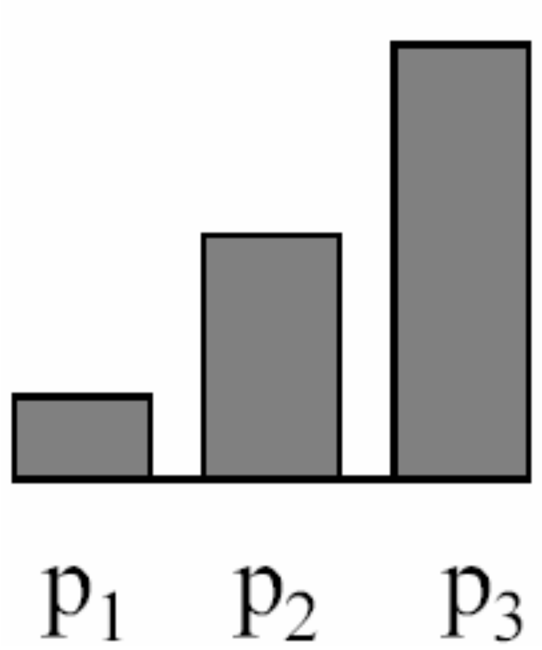

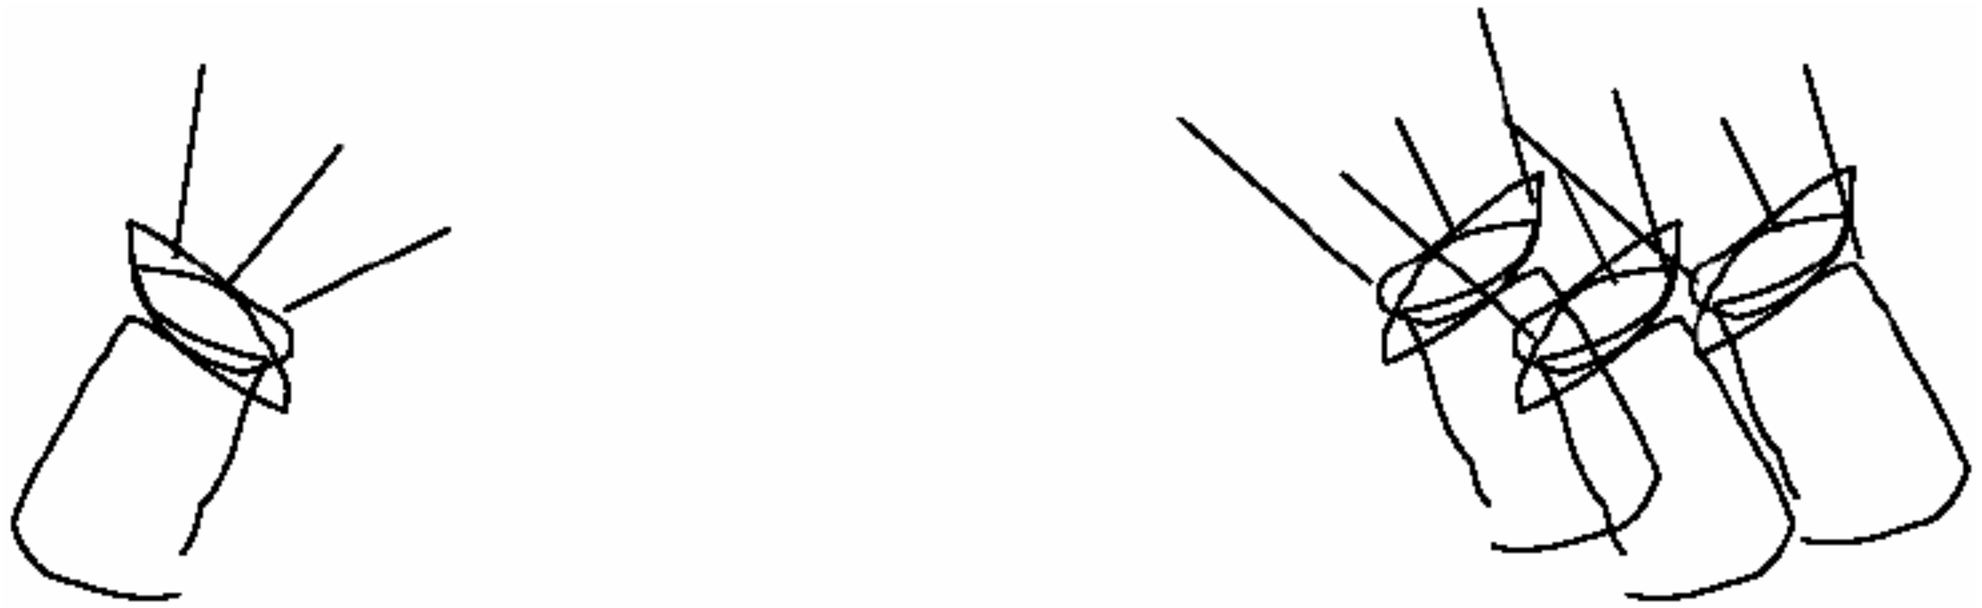

knobs here

Example Credit: Bill Freeman

Example 1: Color Matching Experiment

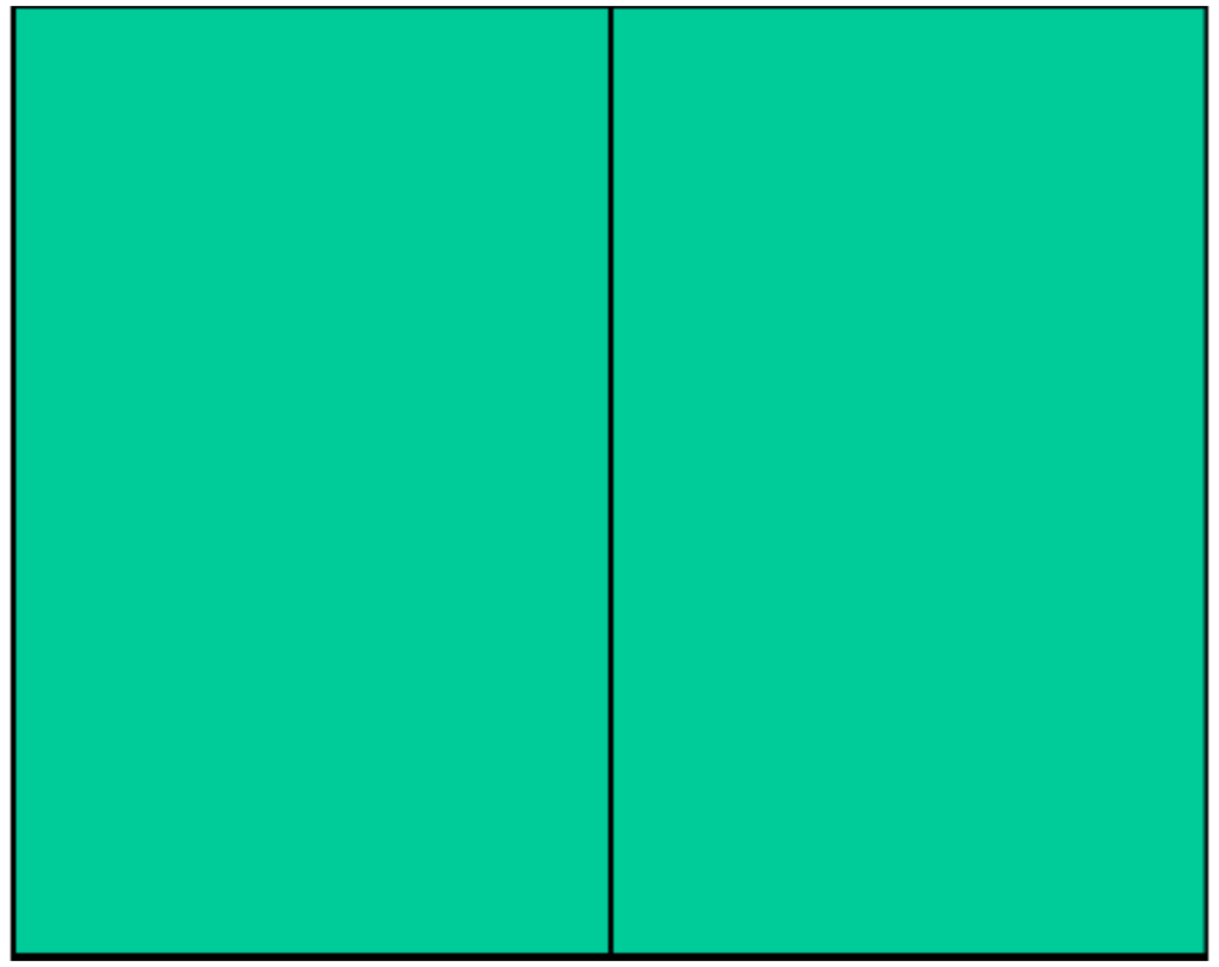

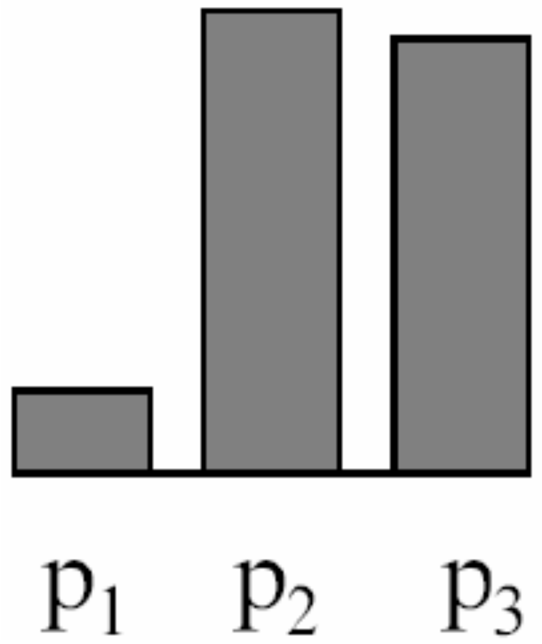

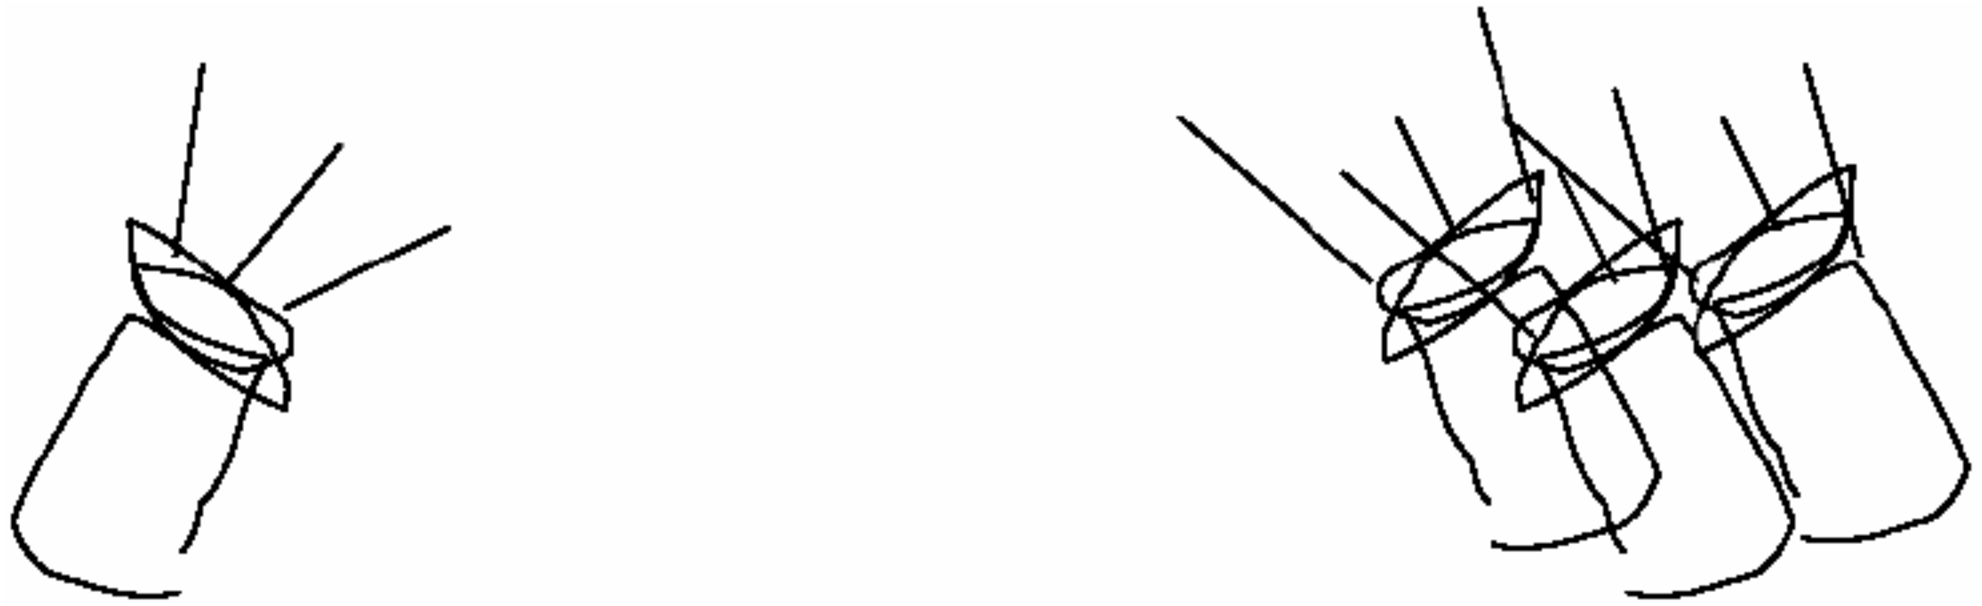

knobs here

Example Credit: Bill Freeman

Example 1: Color Matching Experiment

$$T = w_1 P_1 + w_2 P_2 + w_3 P_3$$

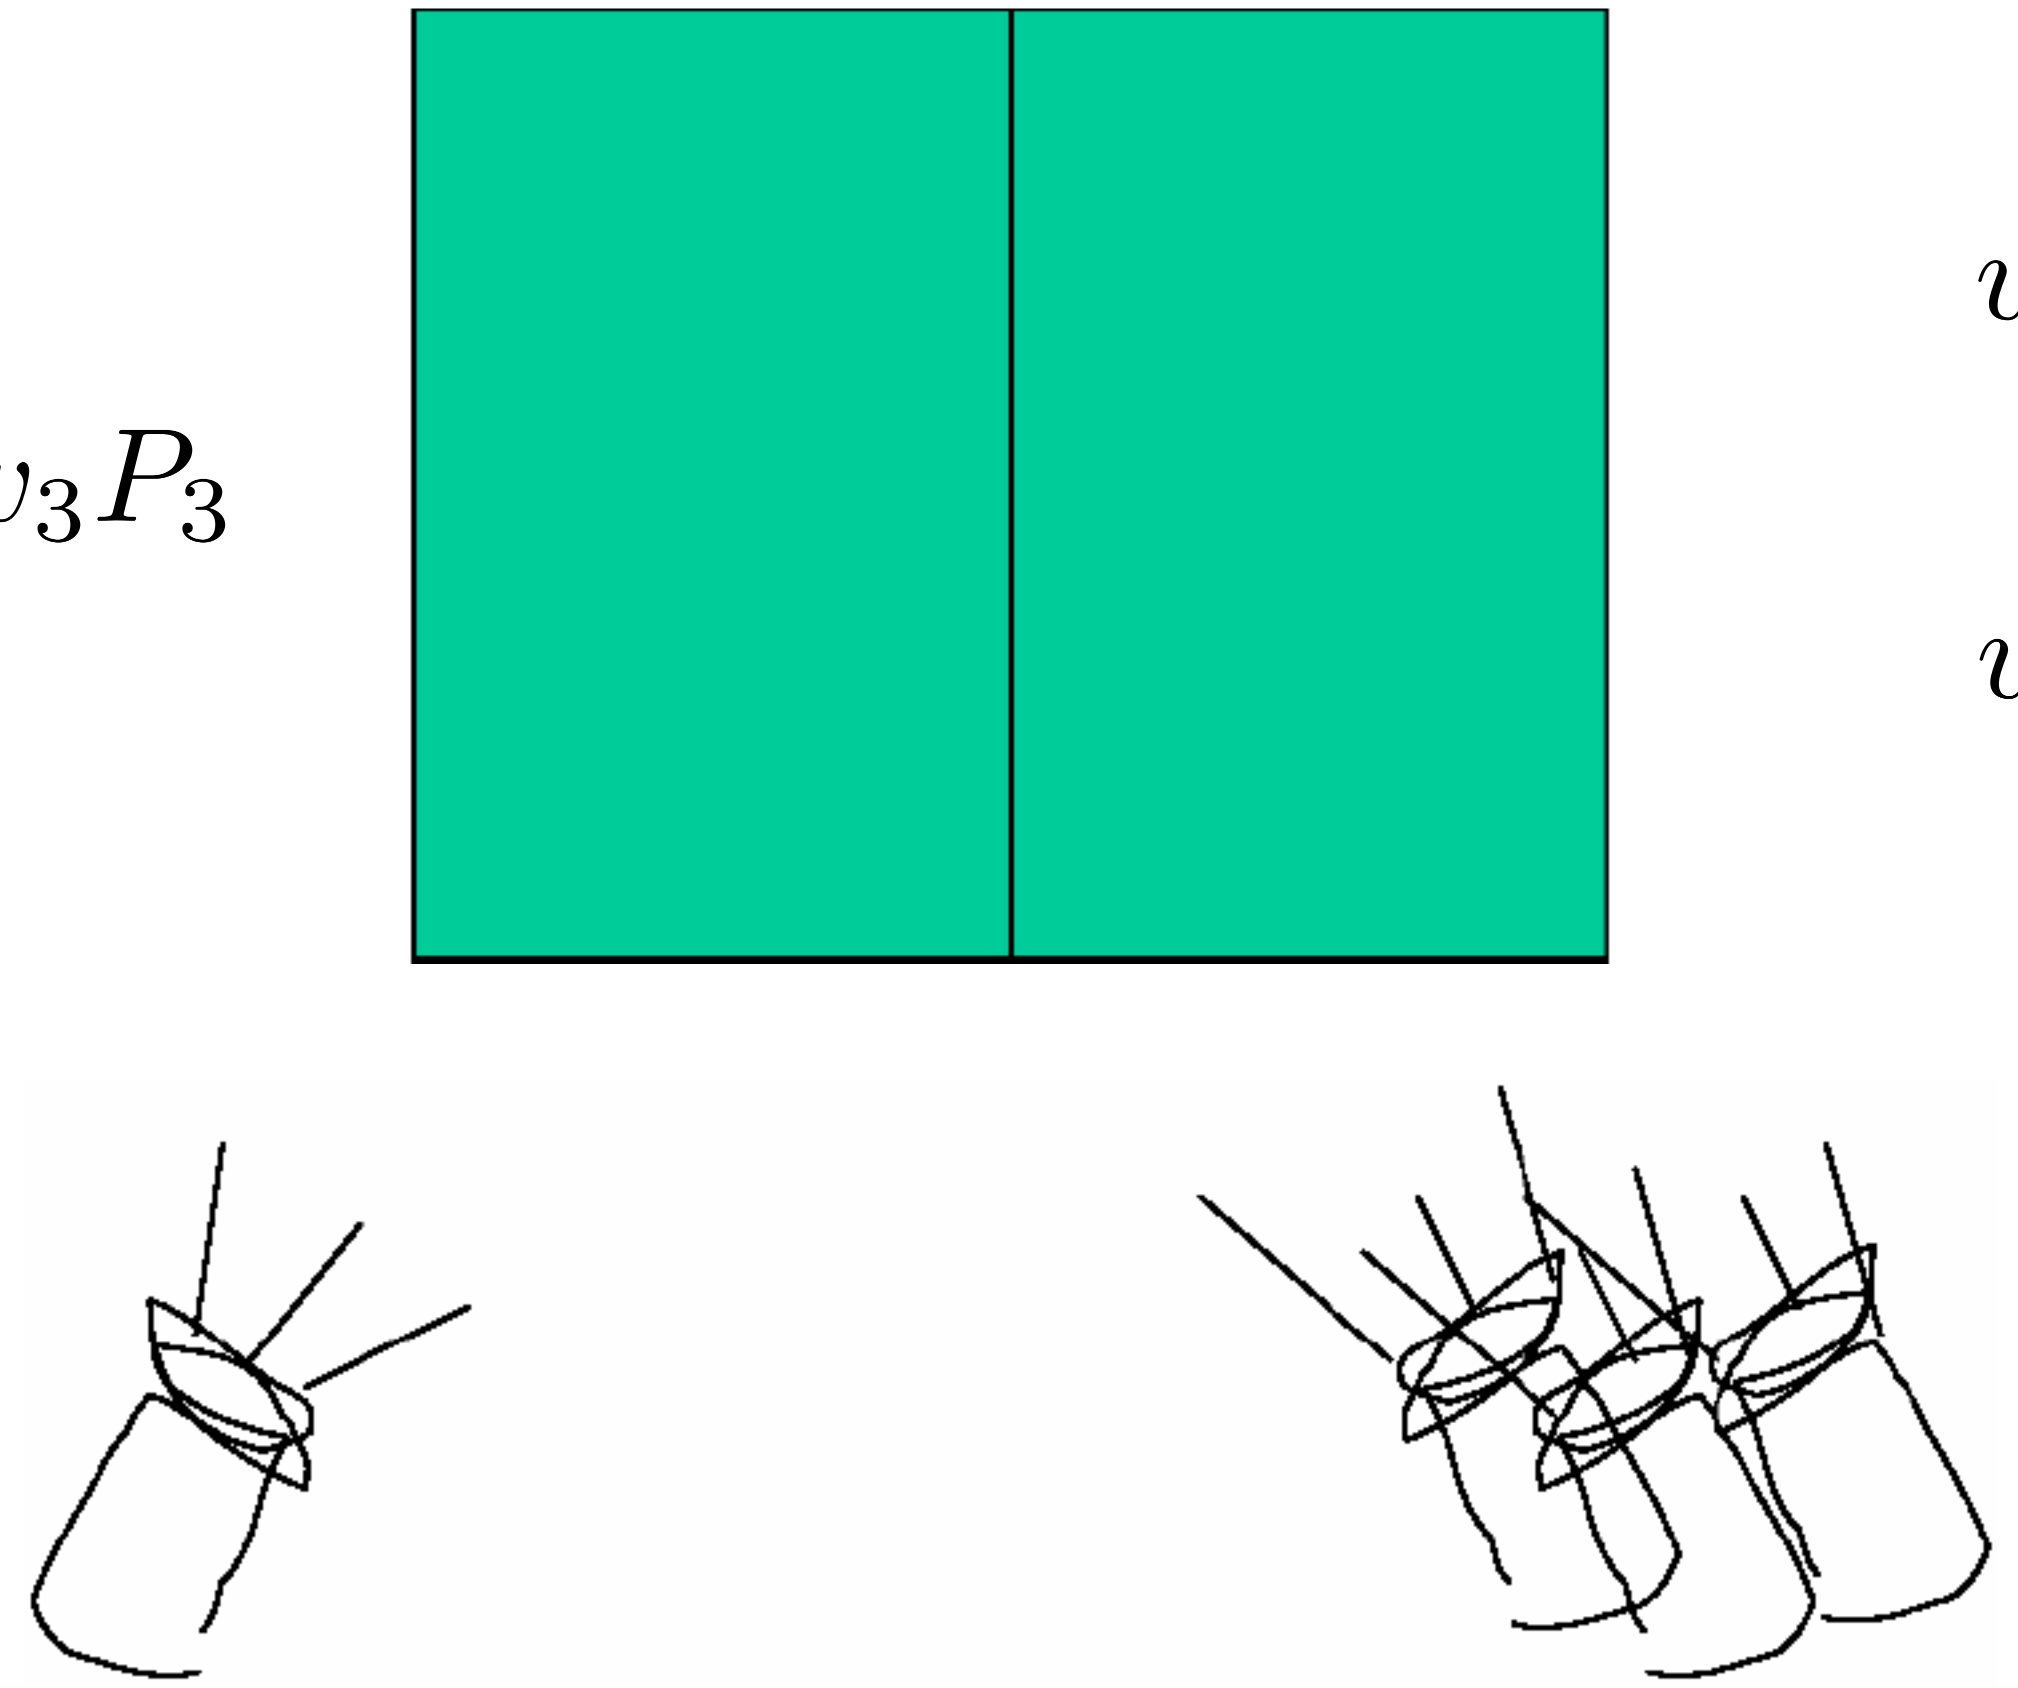

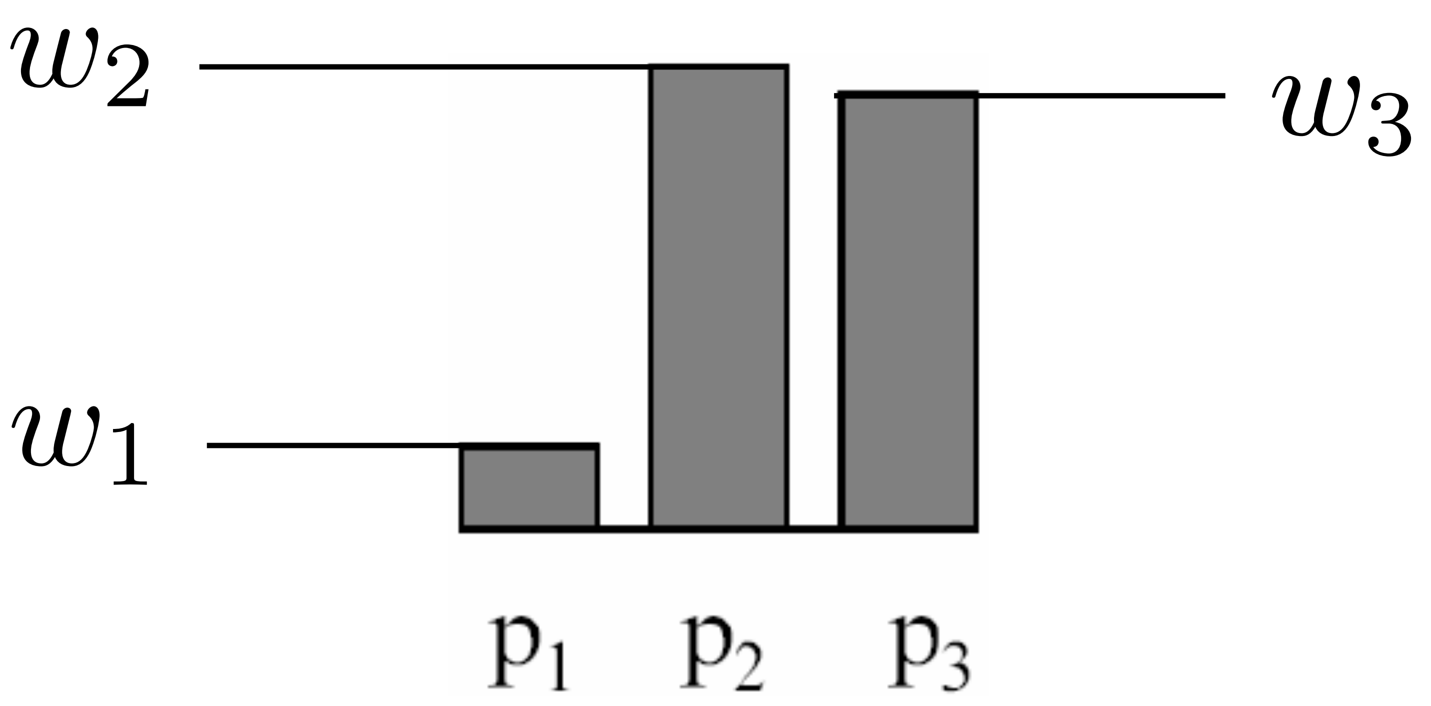

knobs here

Example Credit: Bill Freeman

Example 2: Color Matching Experiment

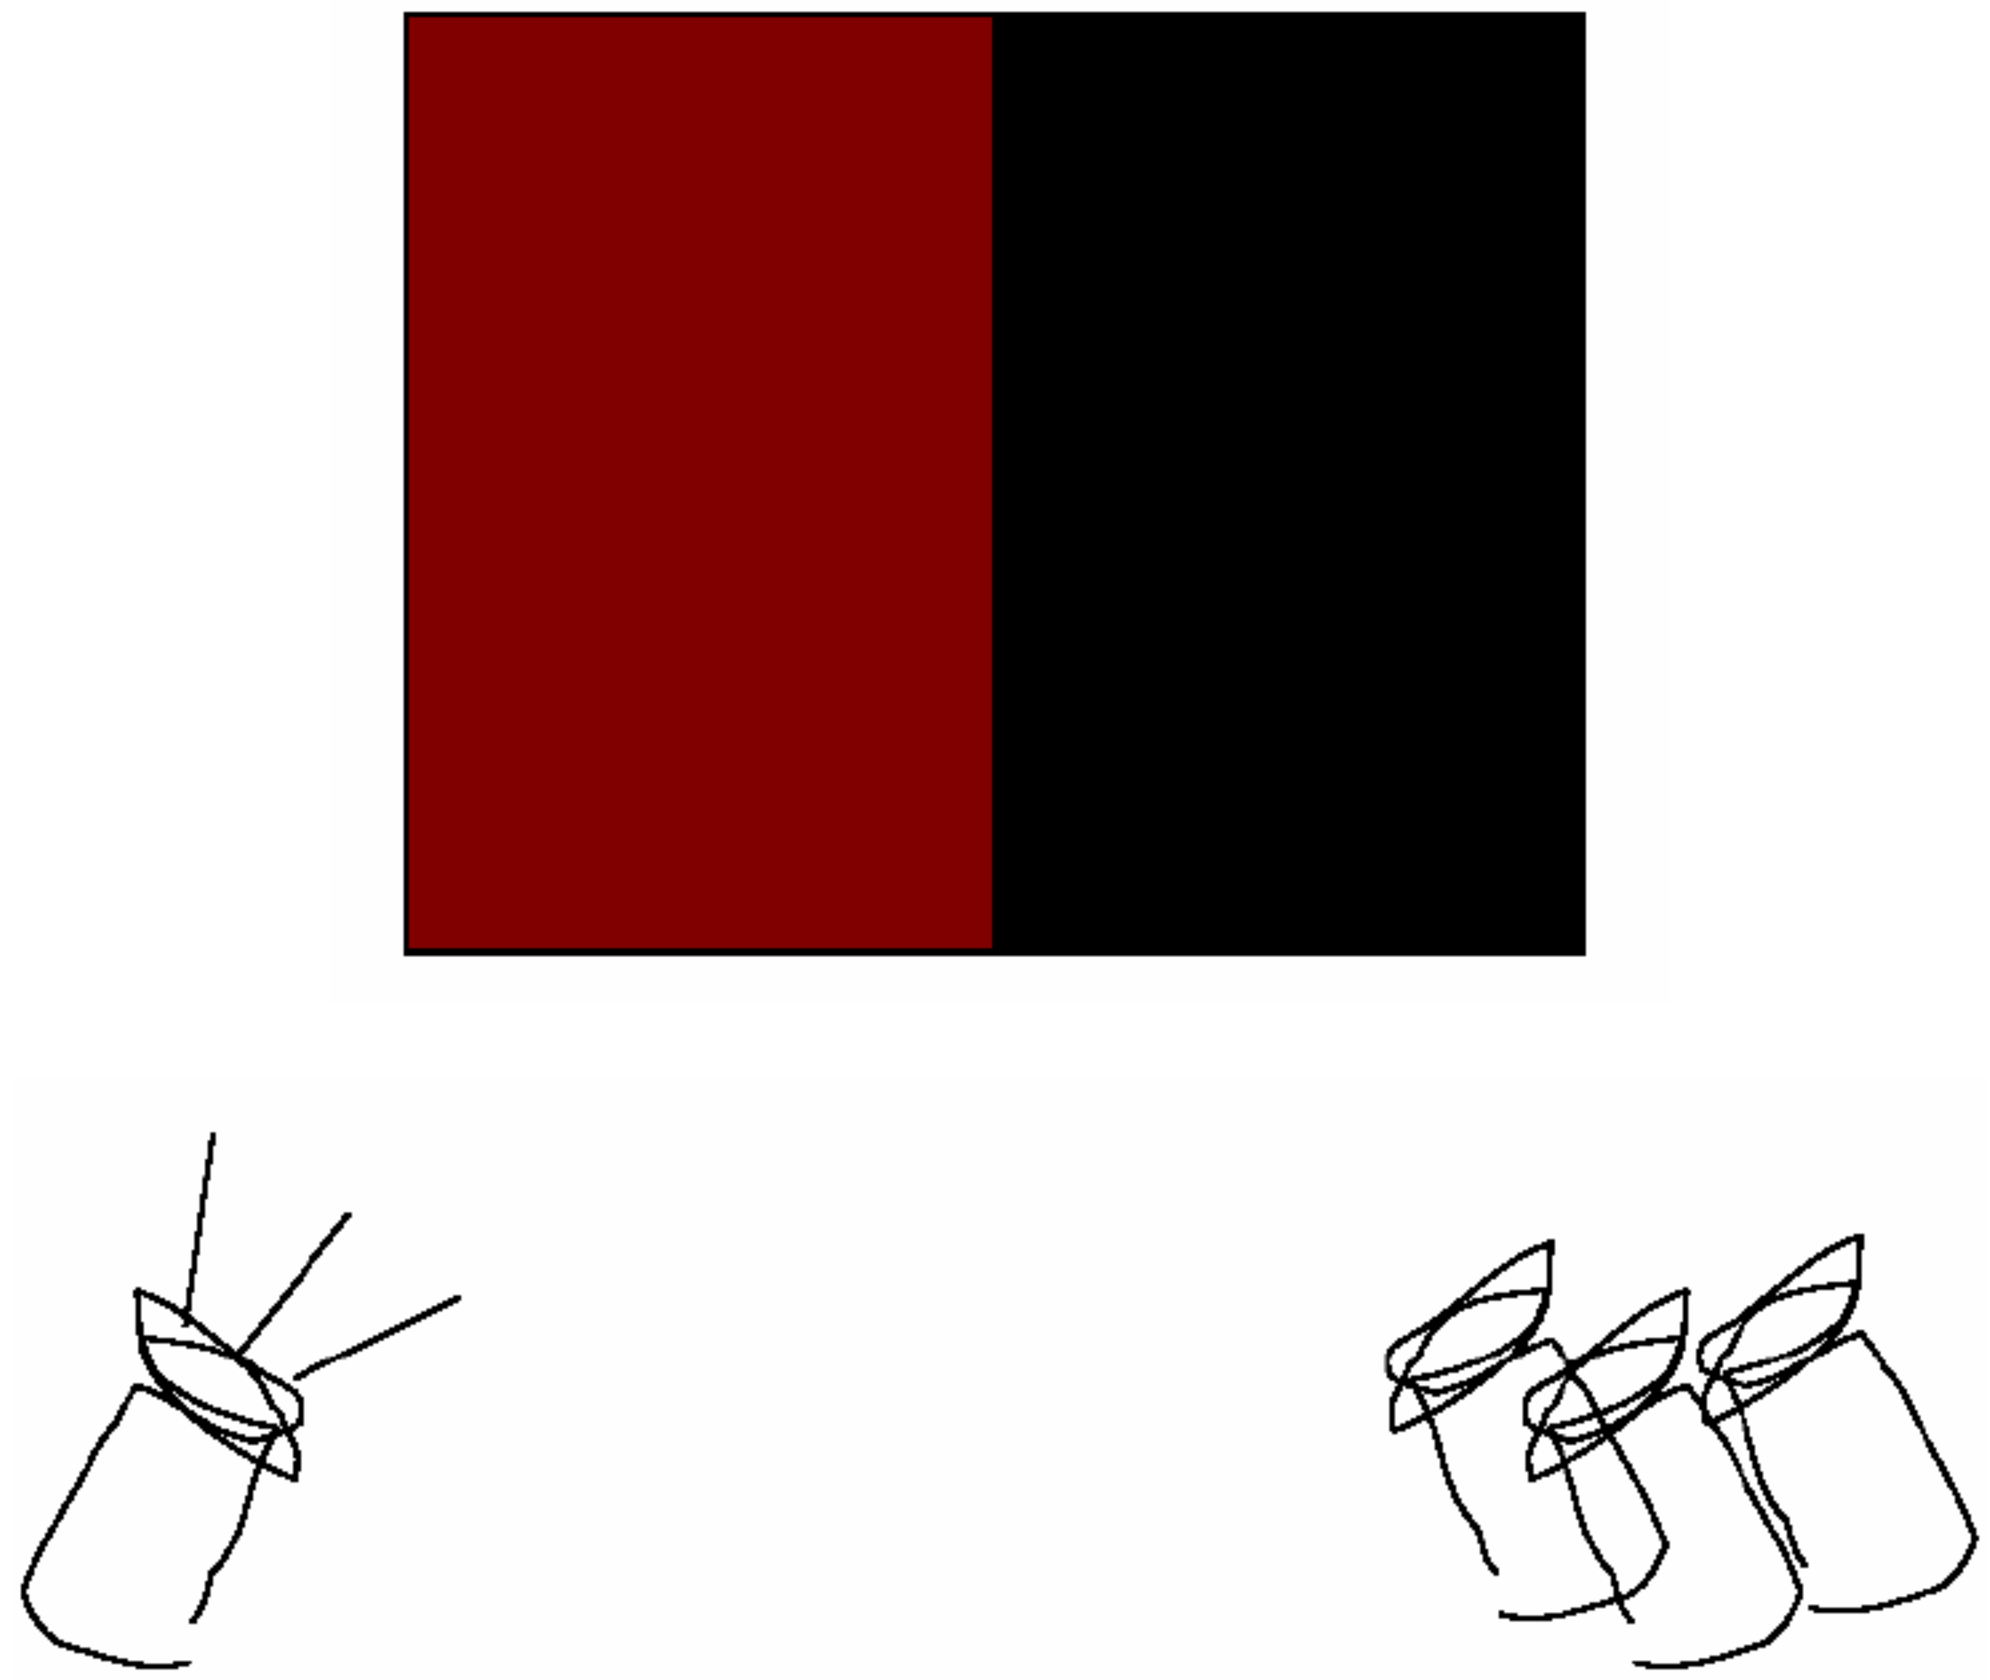

Example Credit: Bill Freeman

Example 2: Color Matching Experiment

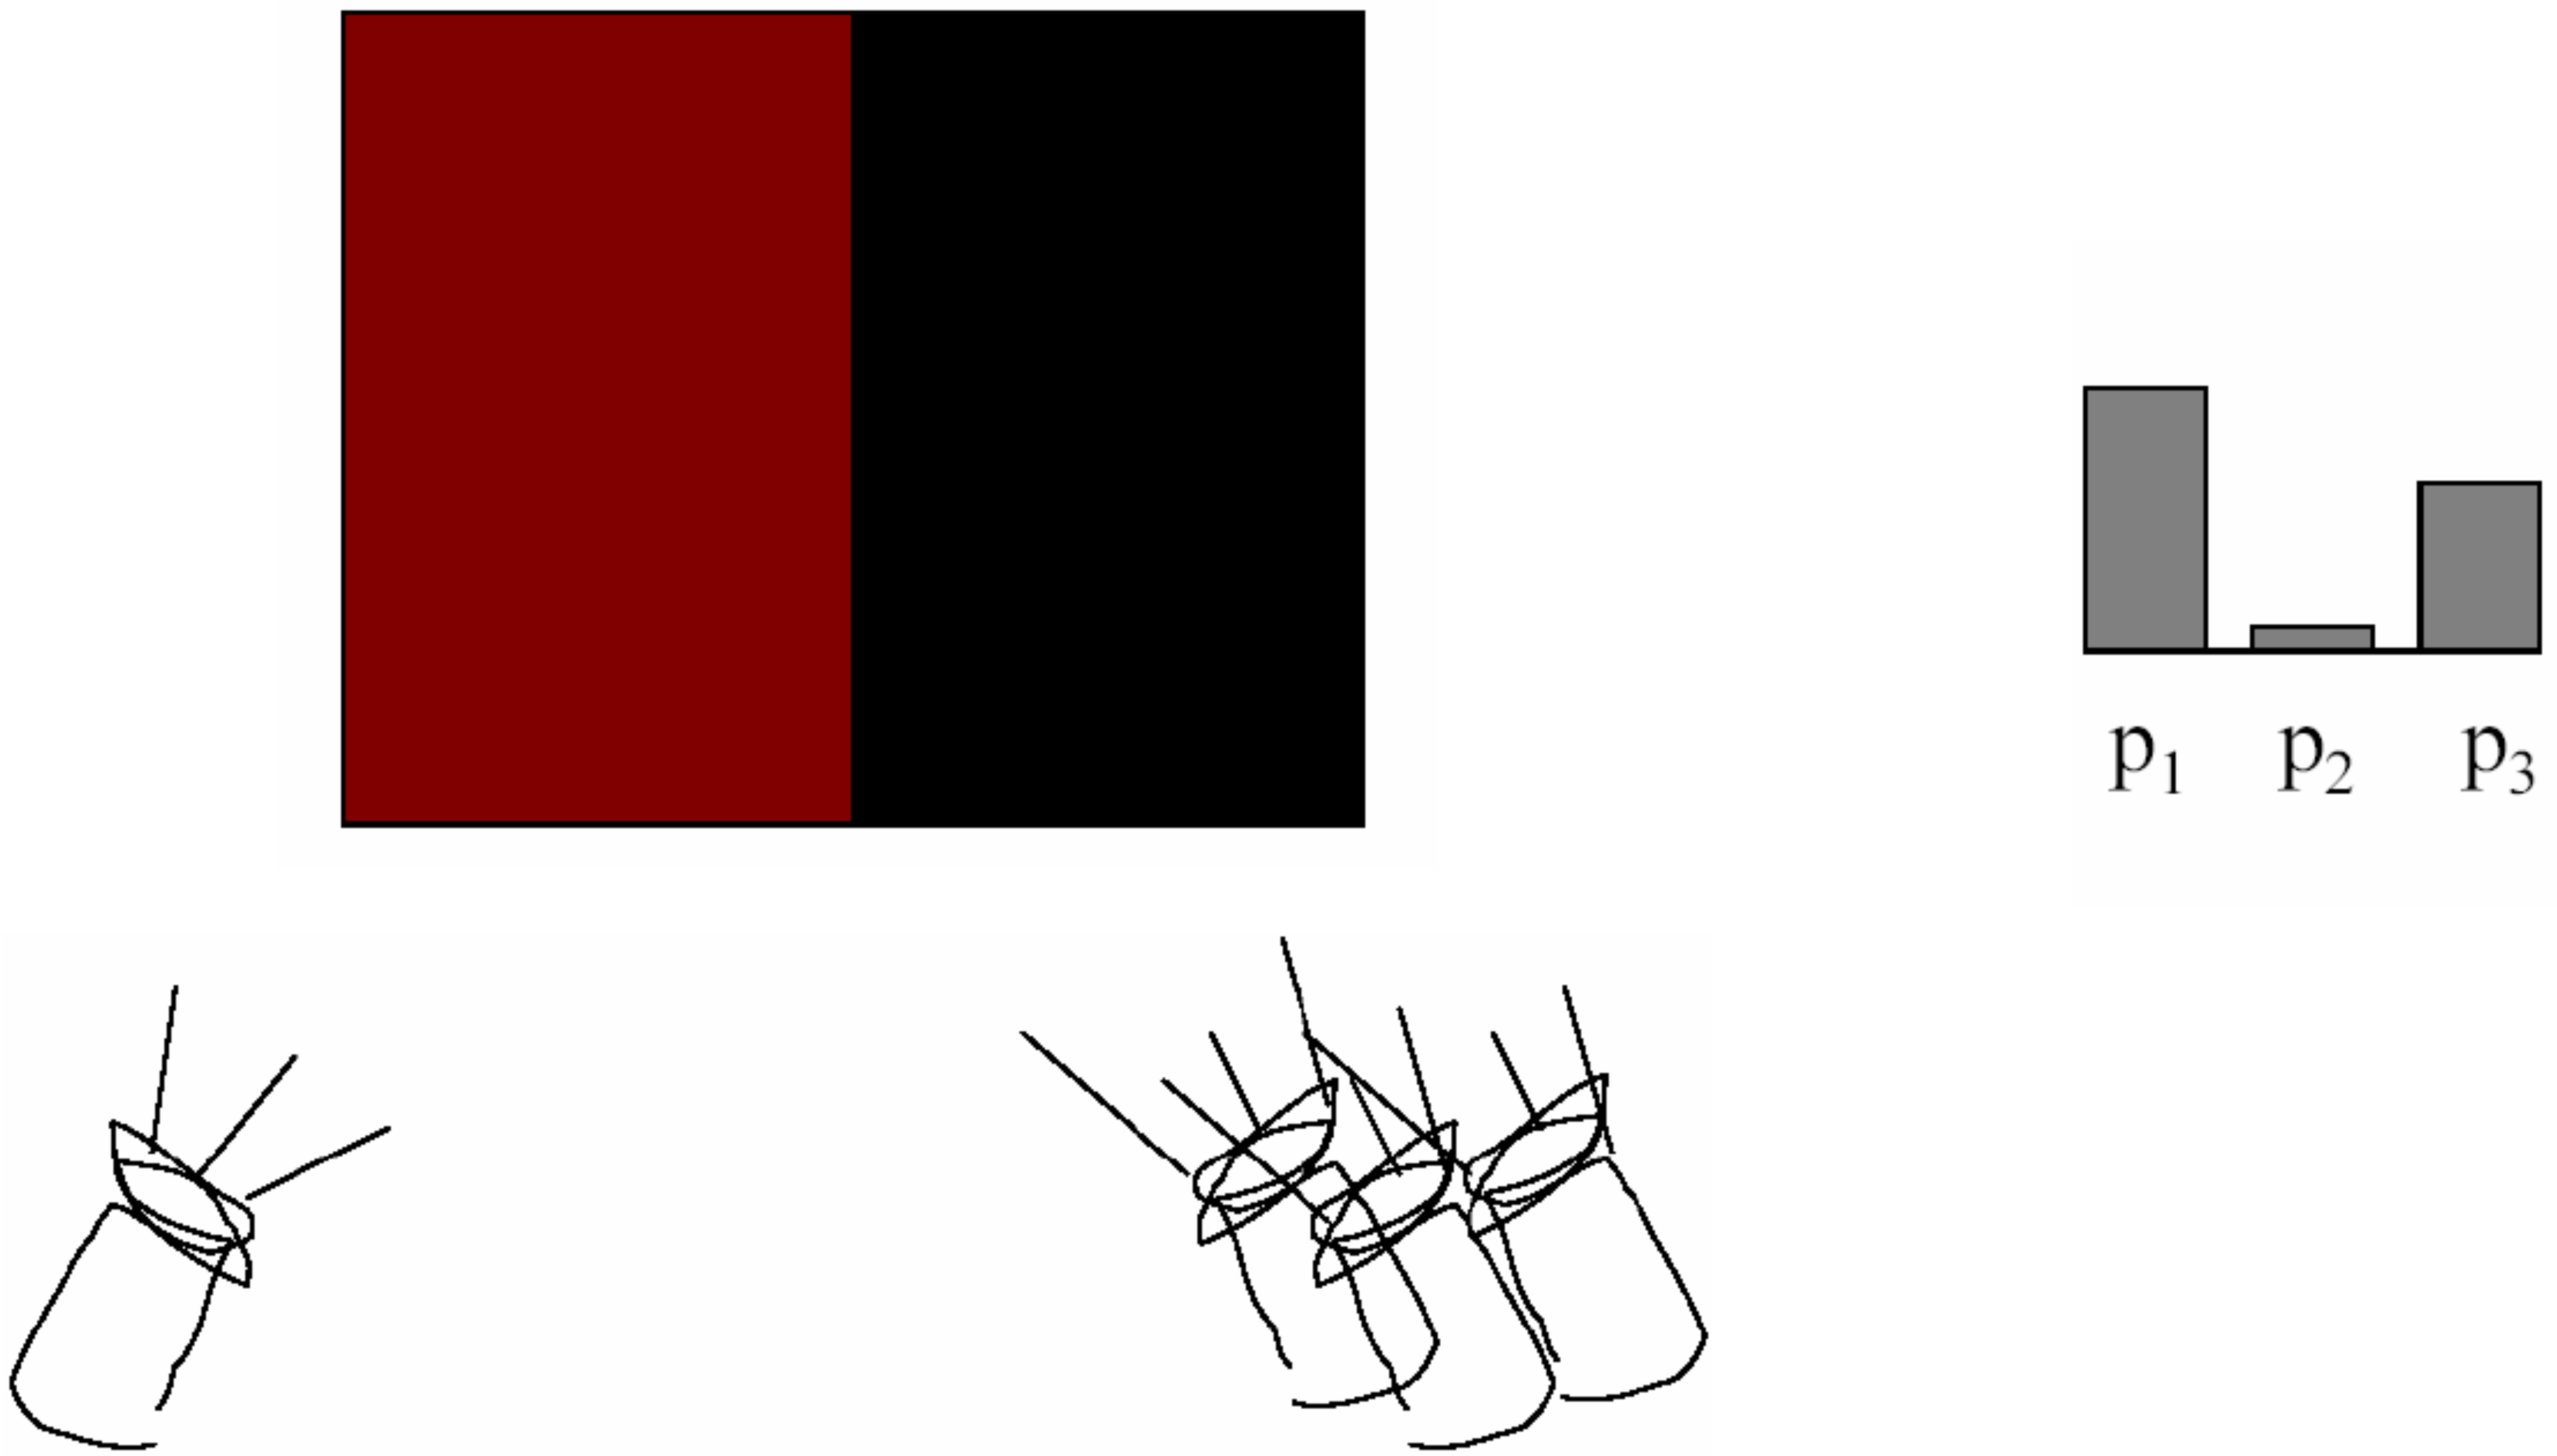

Example Credit: Bill Freeman

Example 2: Color Matching Experiment

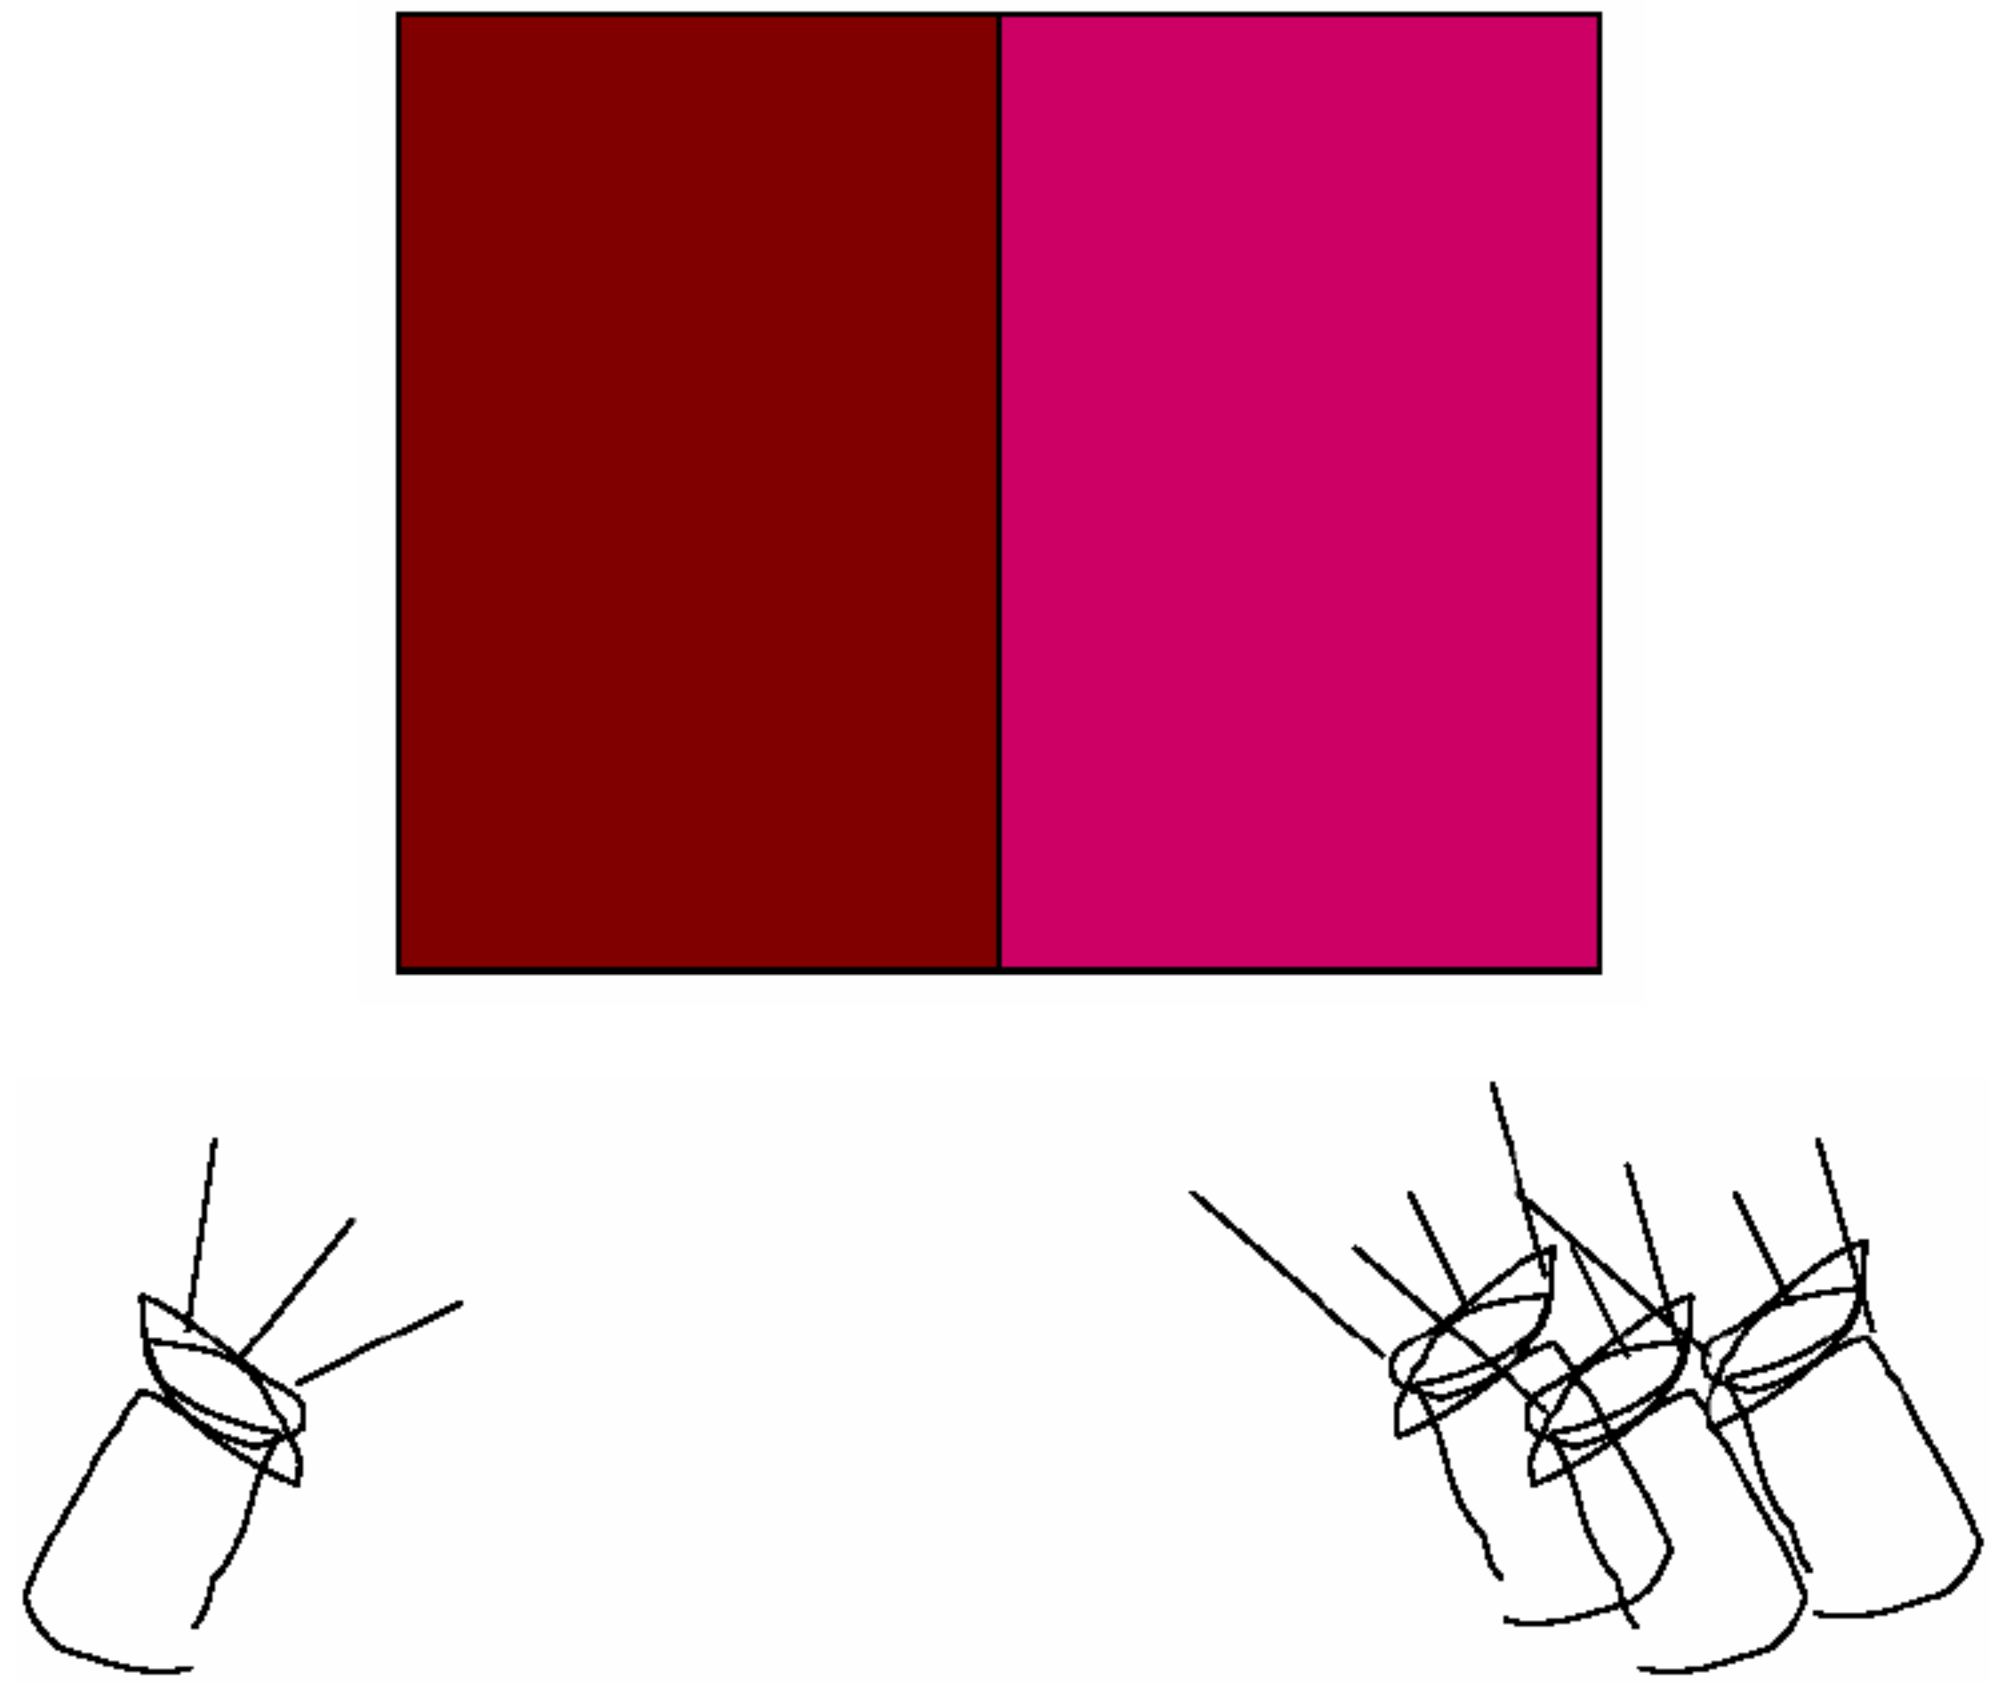

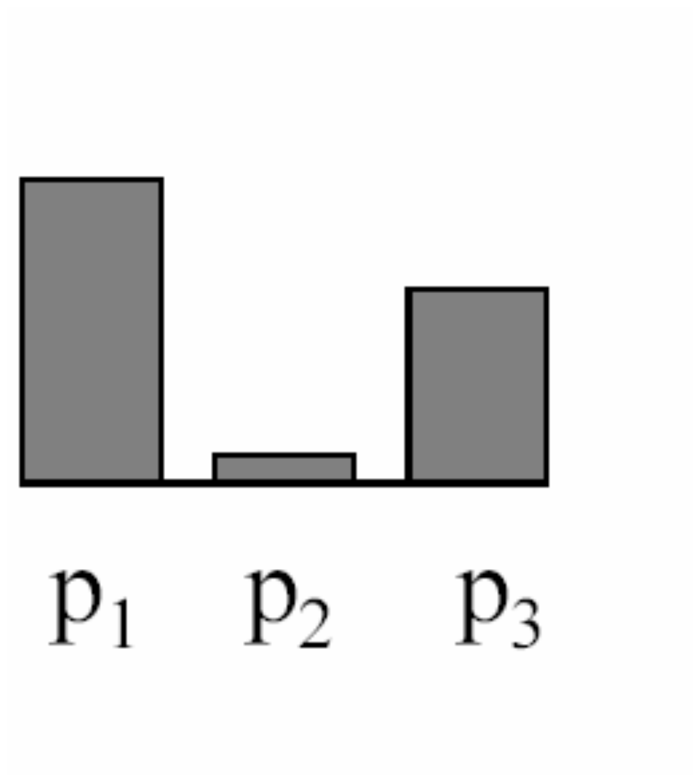

Example Credit: Bill Freeman

Example 2: Color Matching Experiment

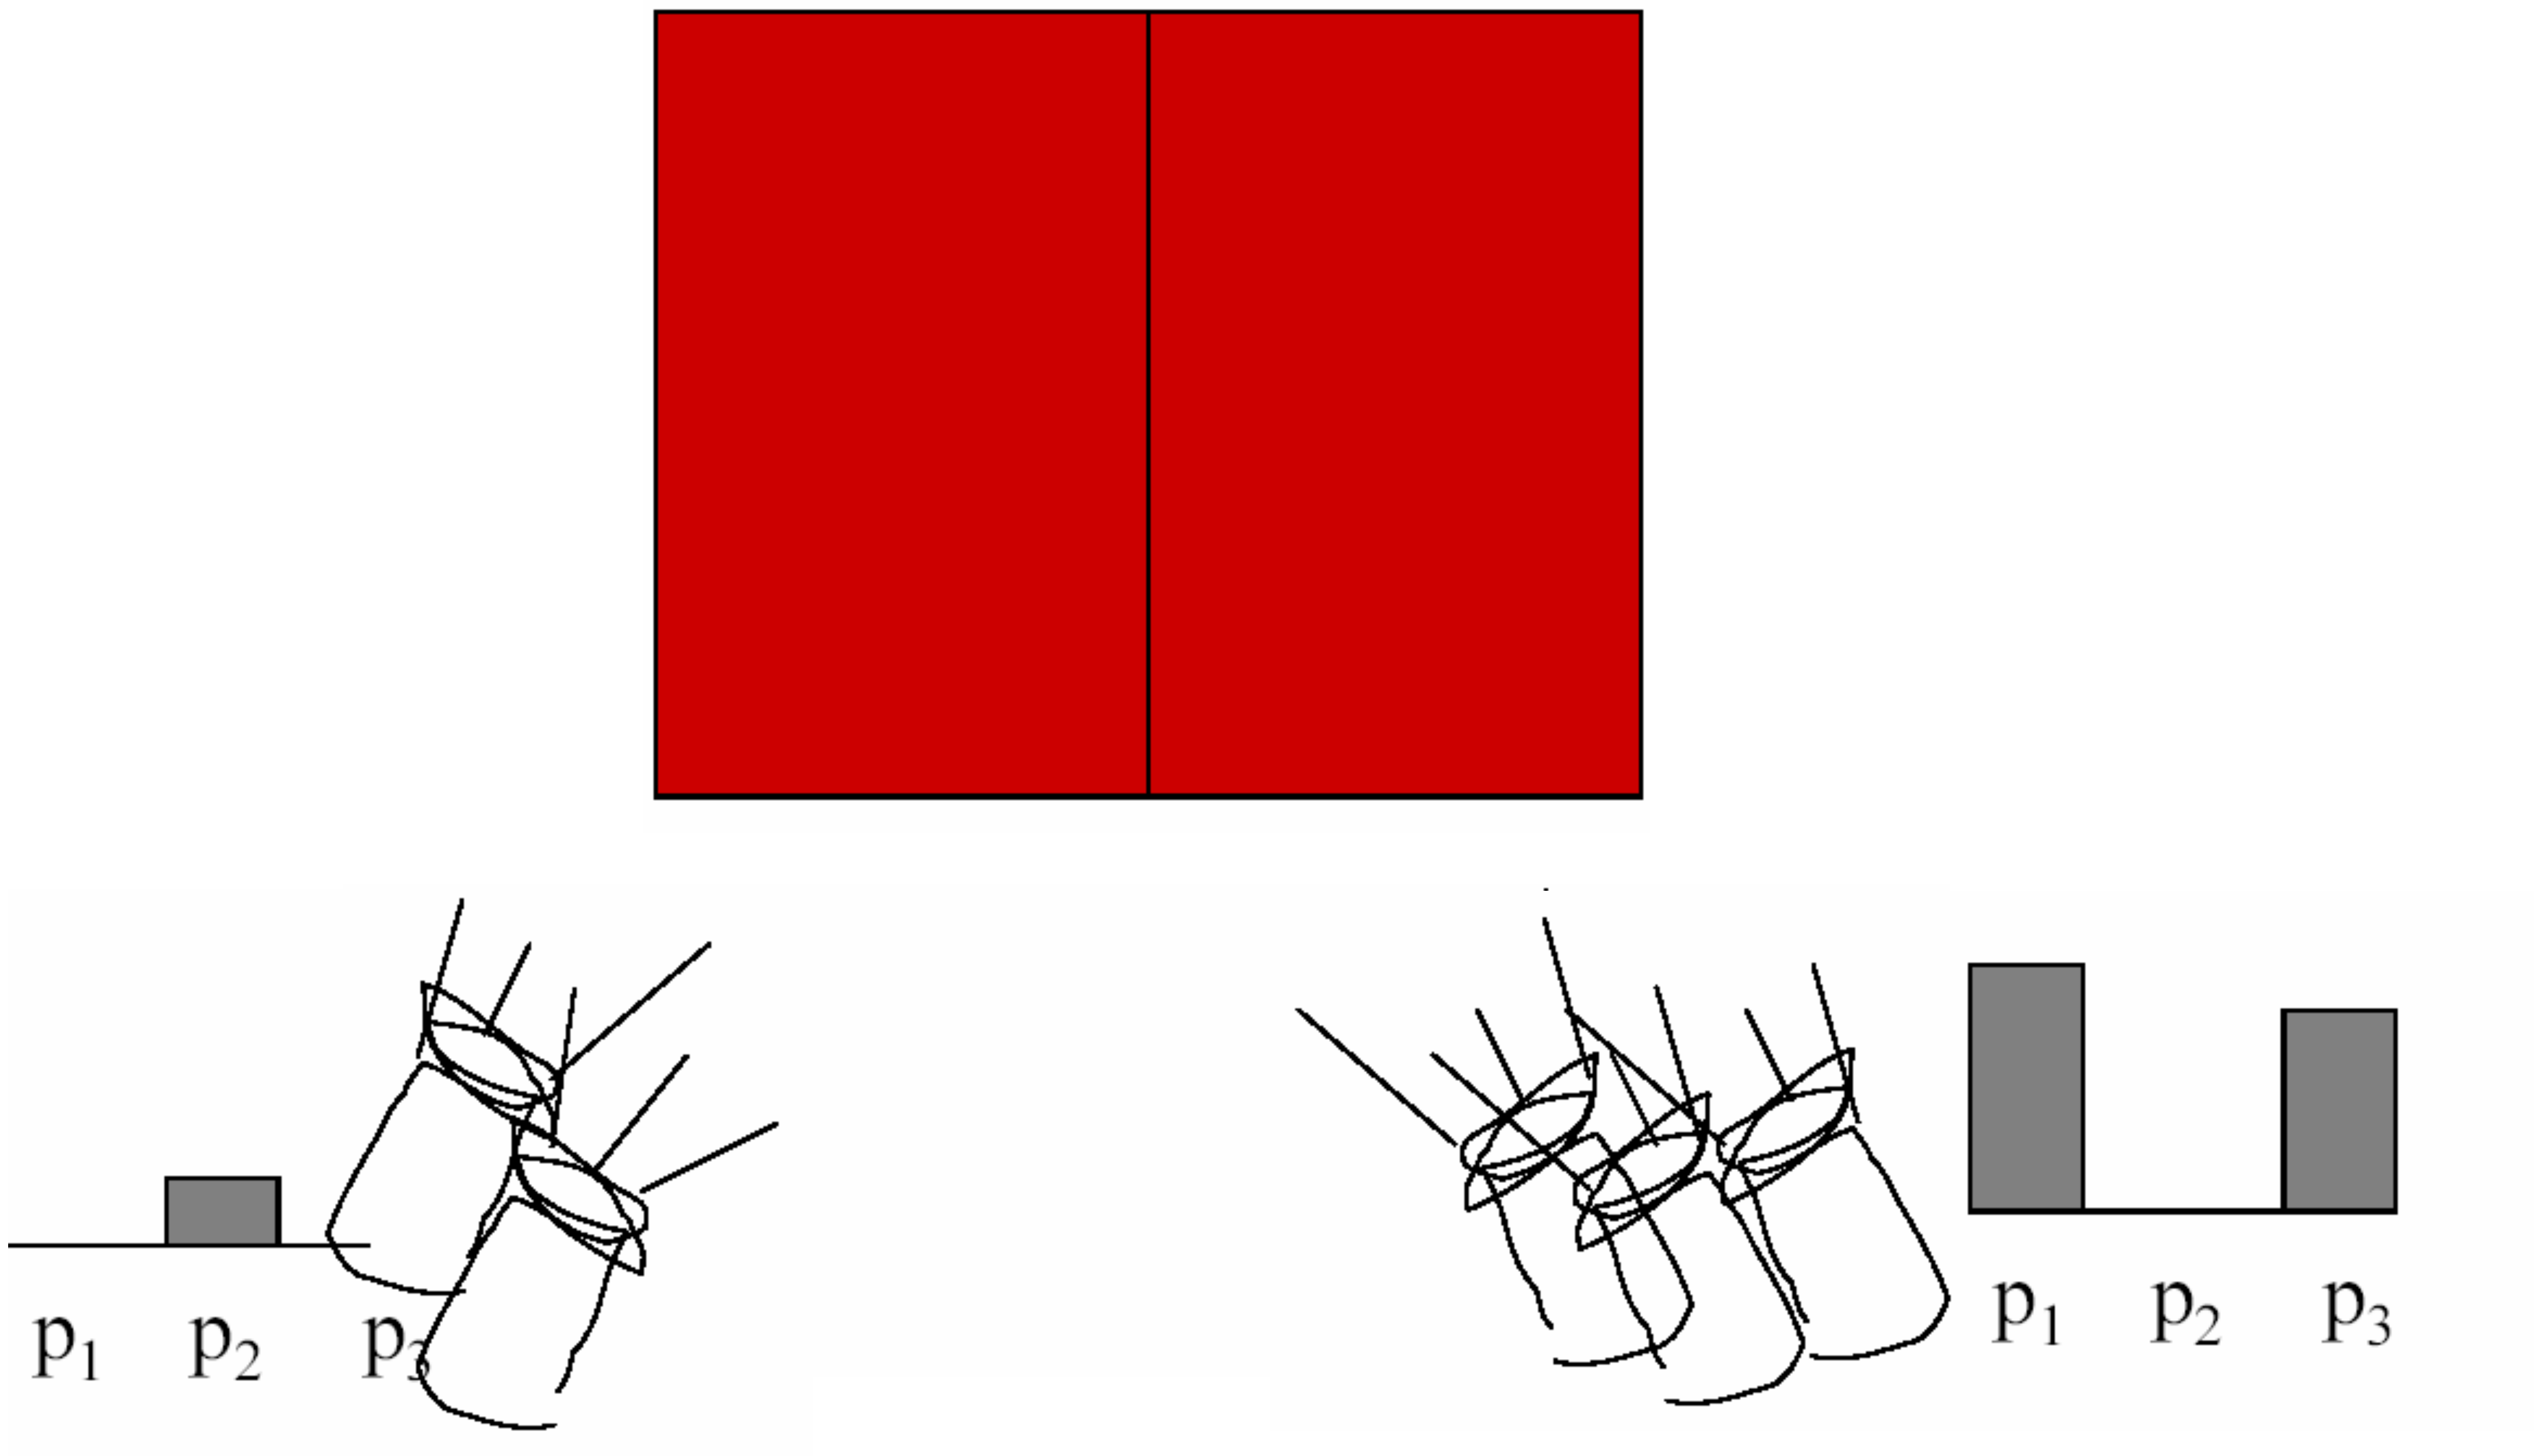

Example Credit: Bill Freeman

Example 2: Color Matching Experiment

We say a “negative” amount of P_2 was needed to make a match , because we added it to the test color side

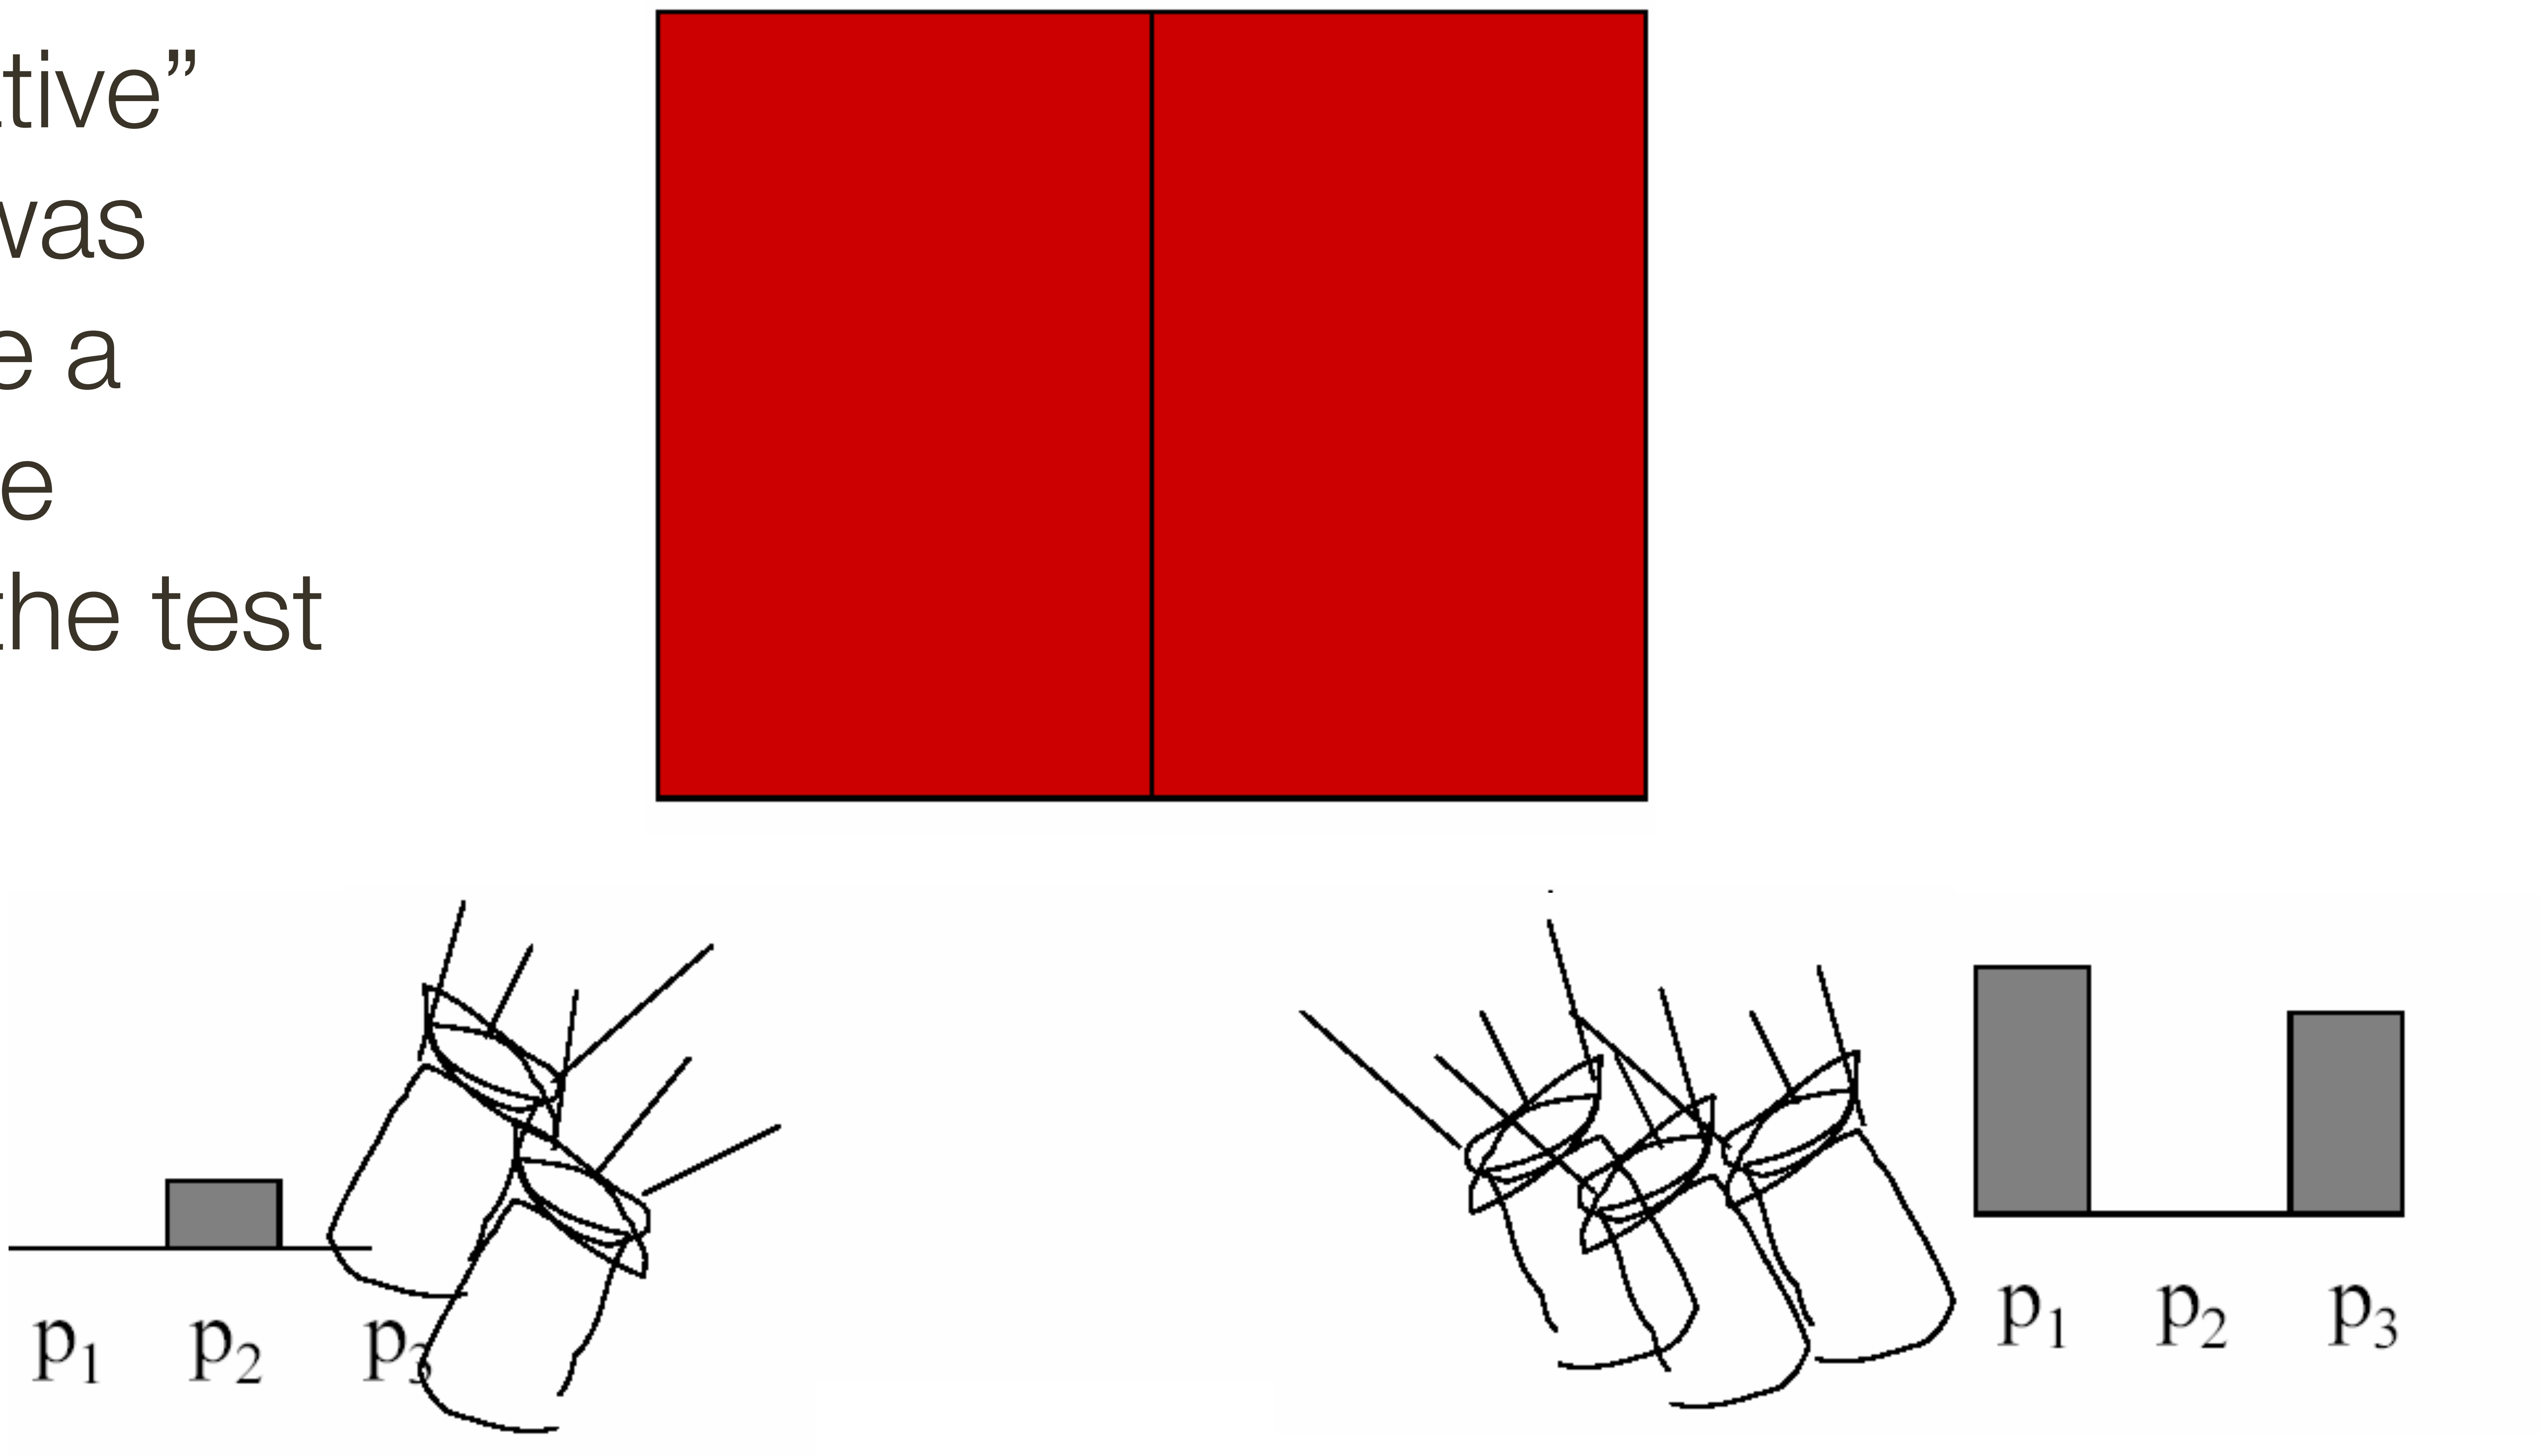

Example Credit: Bill Freeman

Example 2: Color Matching Experiment

We say a “negative” amount of P_2 was needed to make a match , because we added it to the test color side

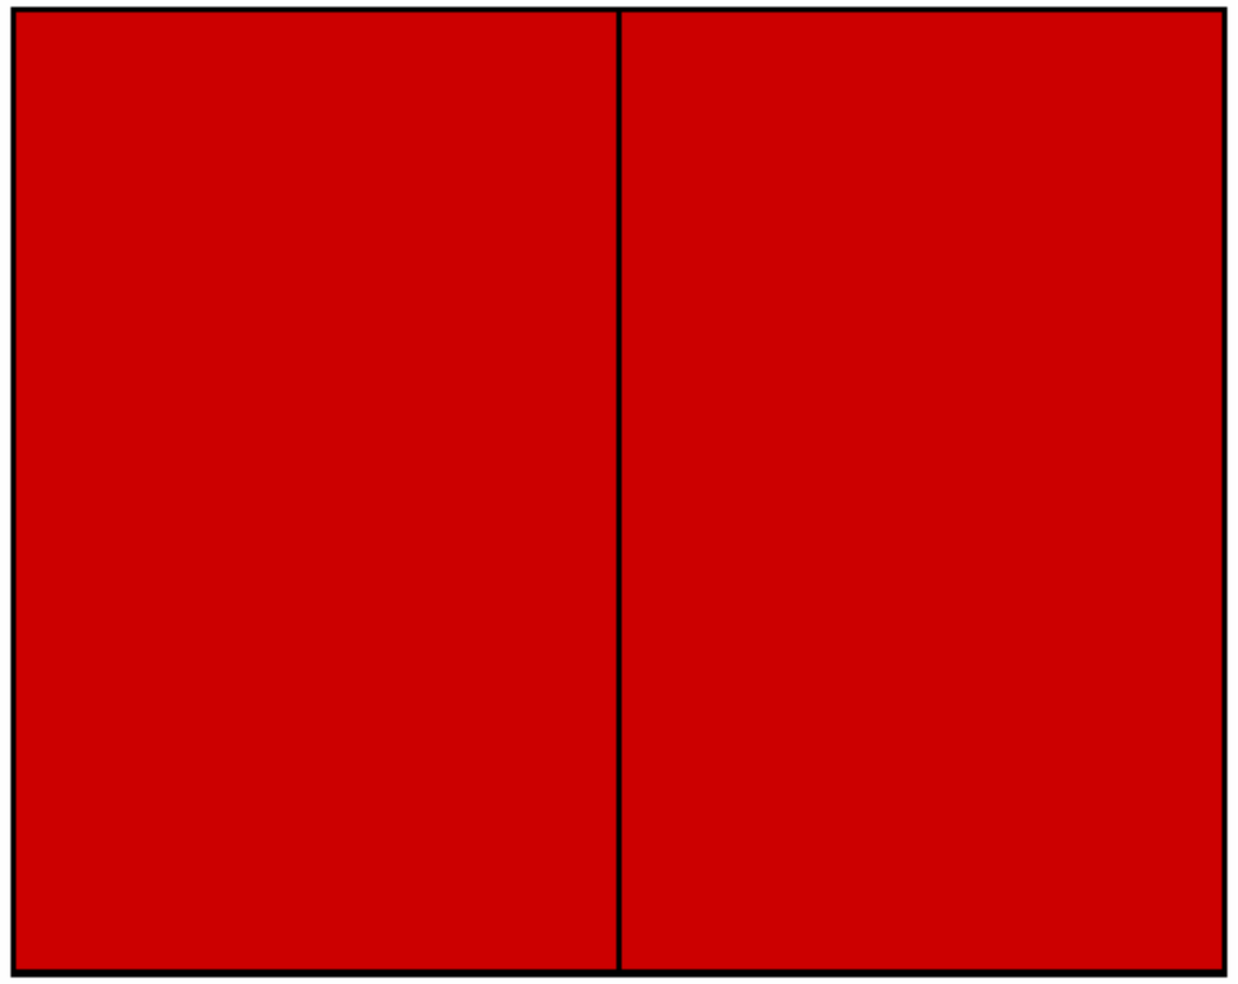

The primary color amount needed to match:

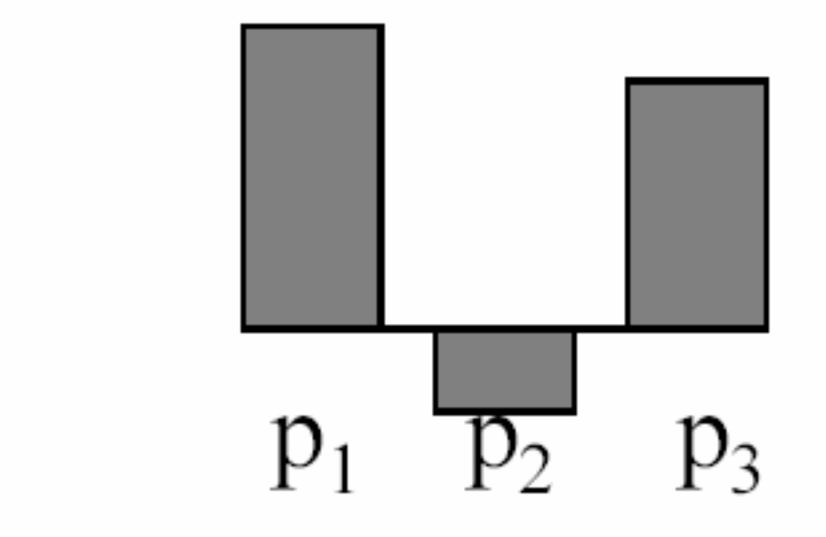

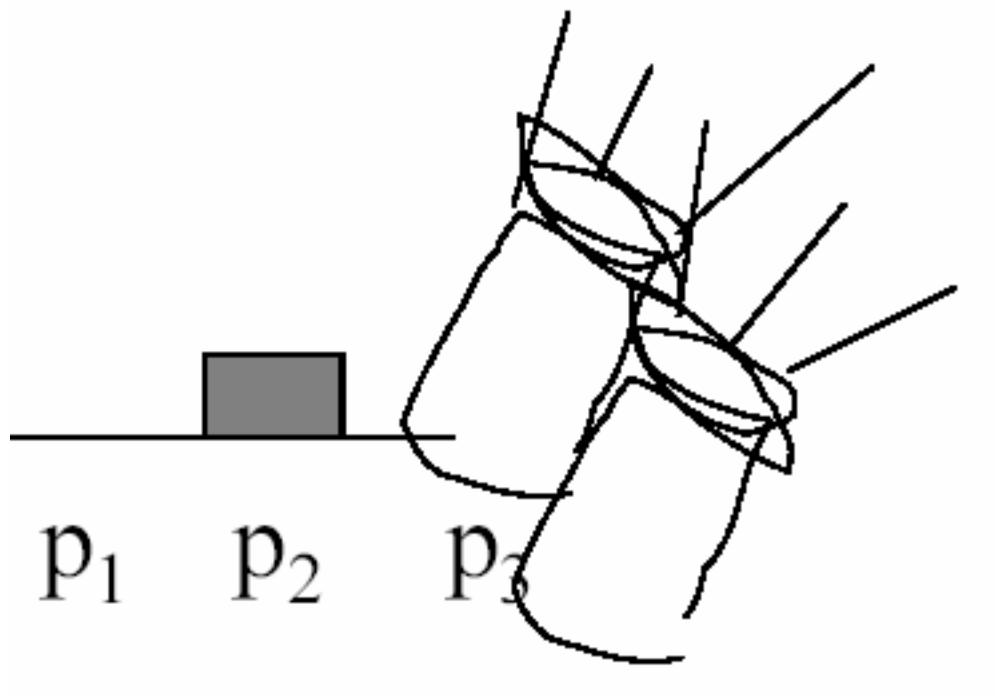

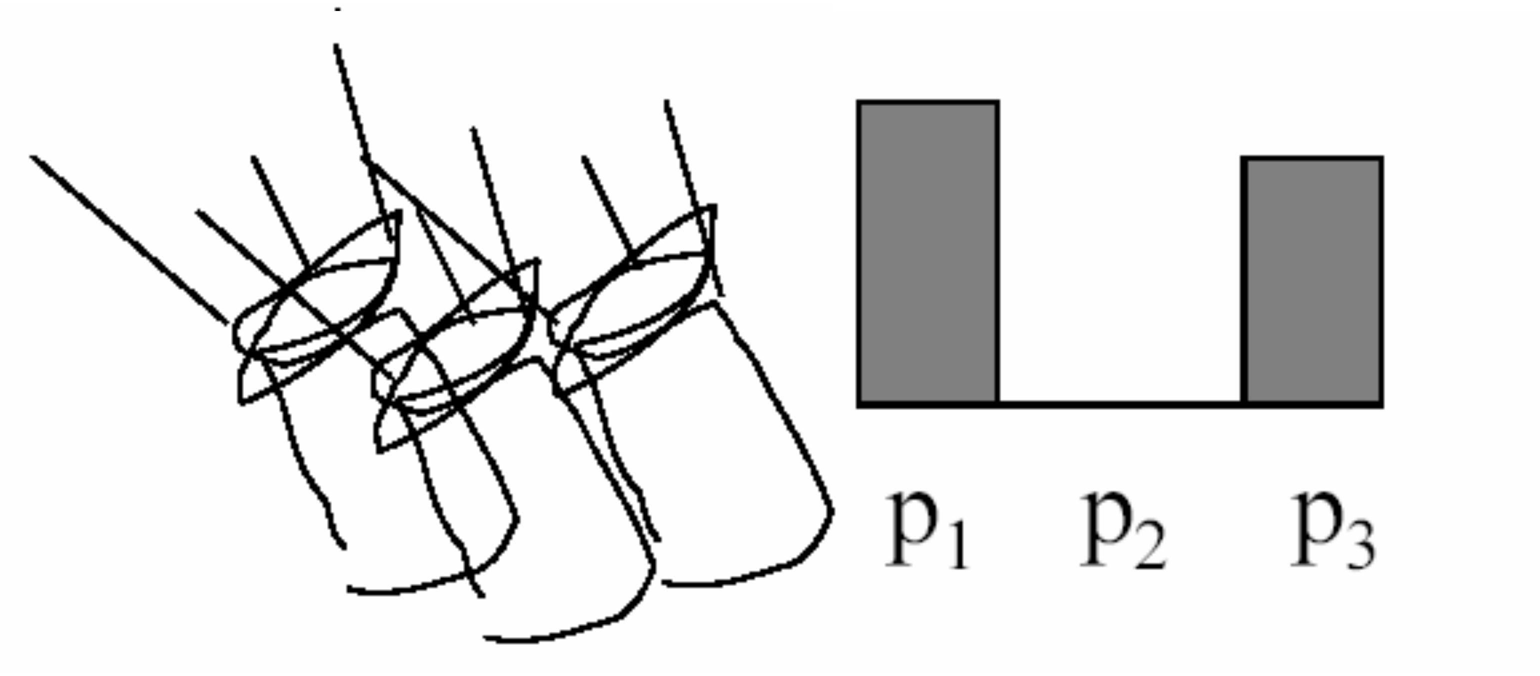

$$T + w_2 P_2 = w_1 P_1 + w_3 P_3$$

Example Credit: Bill Freeman

Important Implication

Most televisions and monitors that are tri-chromatic cannot produce the full spectrum of colors we as humans can perceive (e.g., there are natural colors in bluish-greenish range that we cannot generally produce using RGB)

Sharp aquos

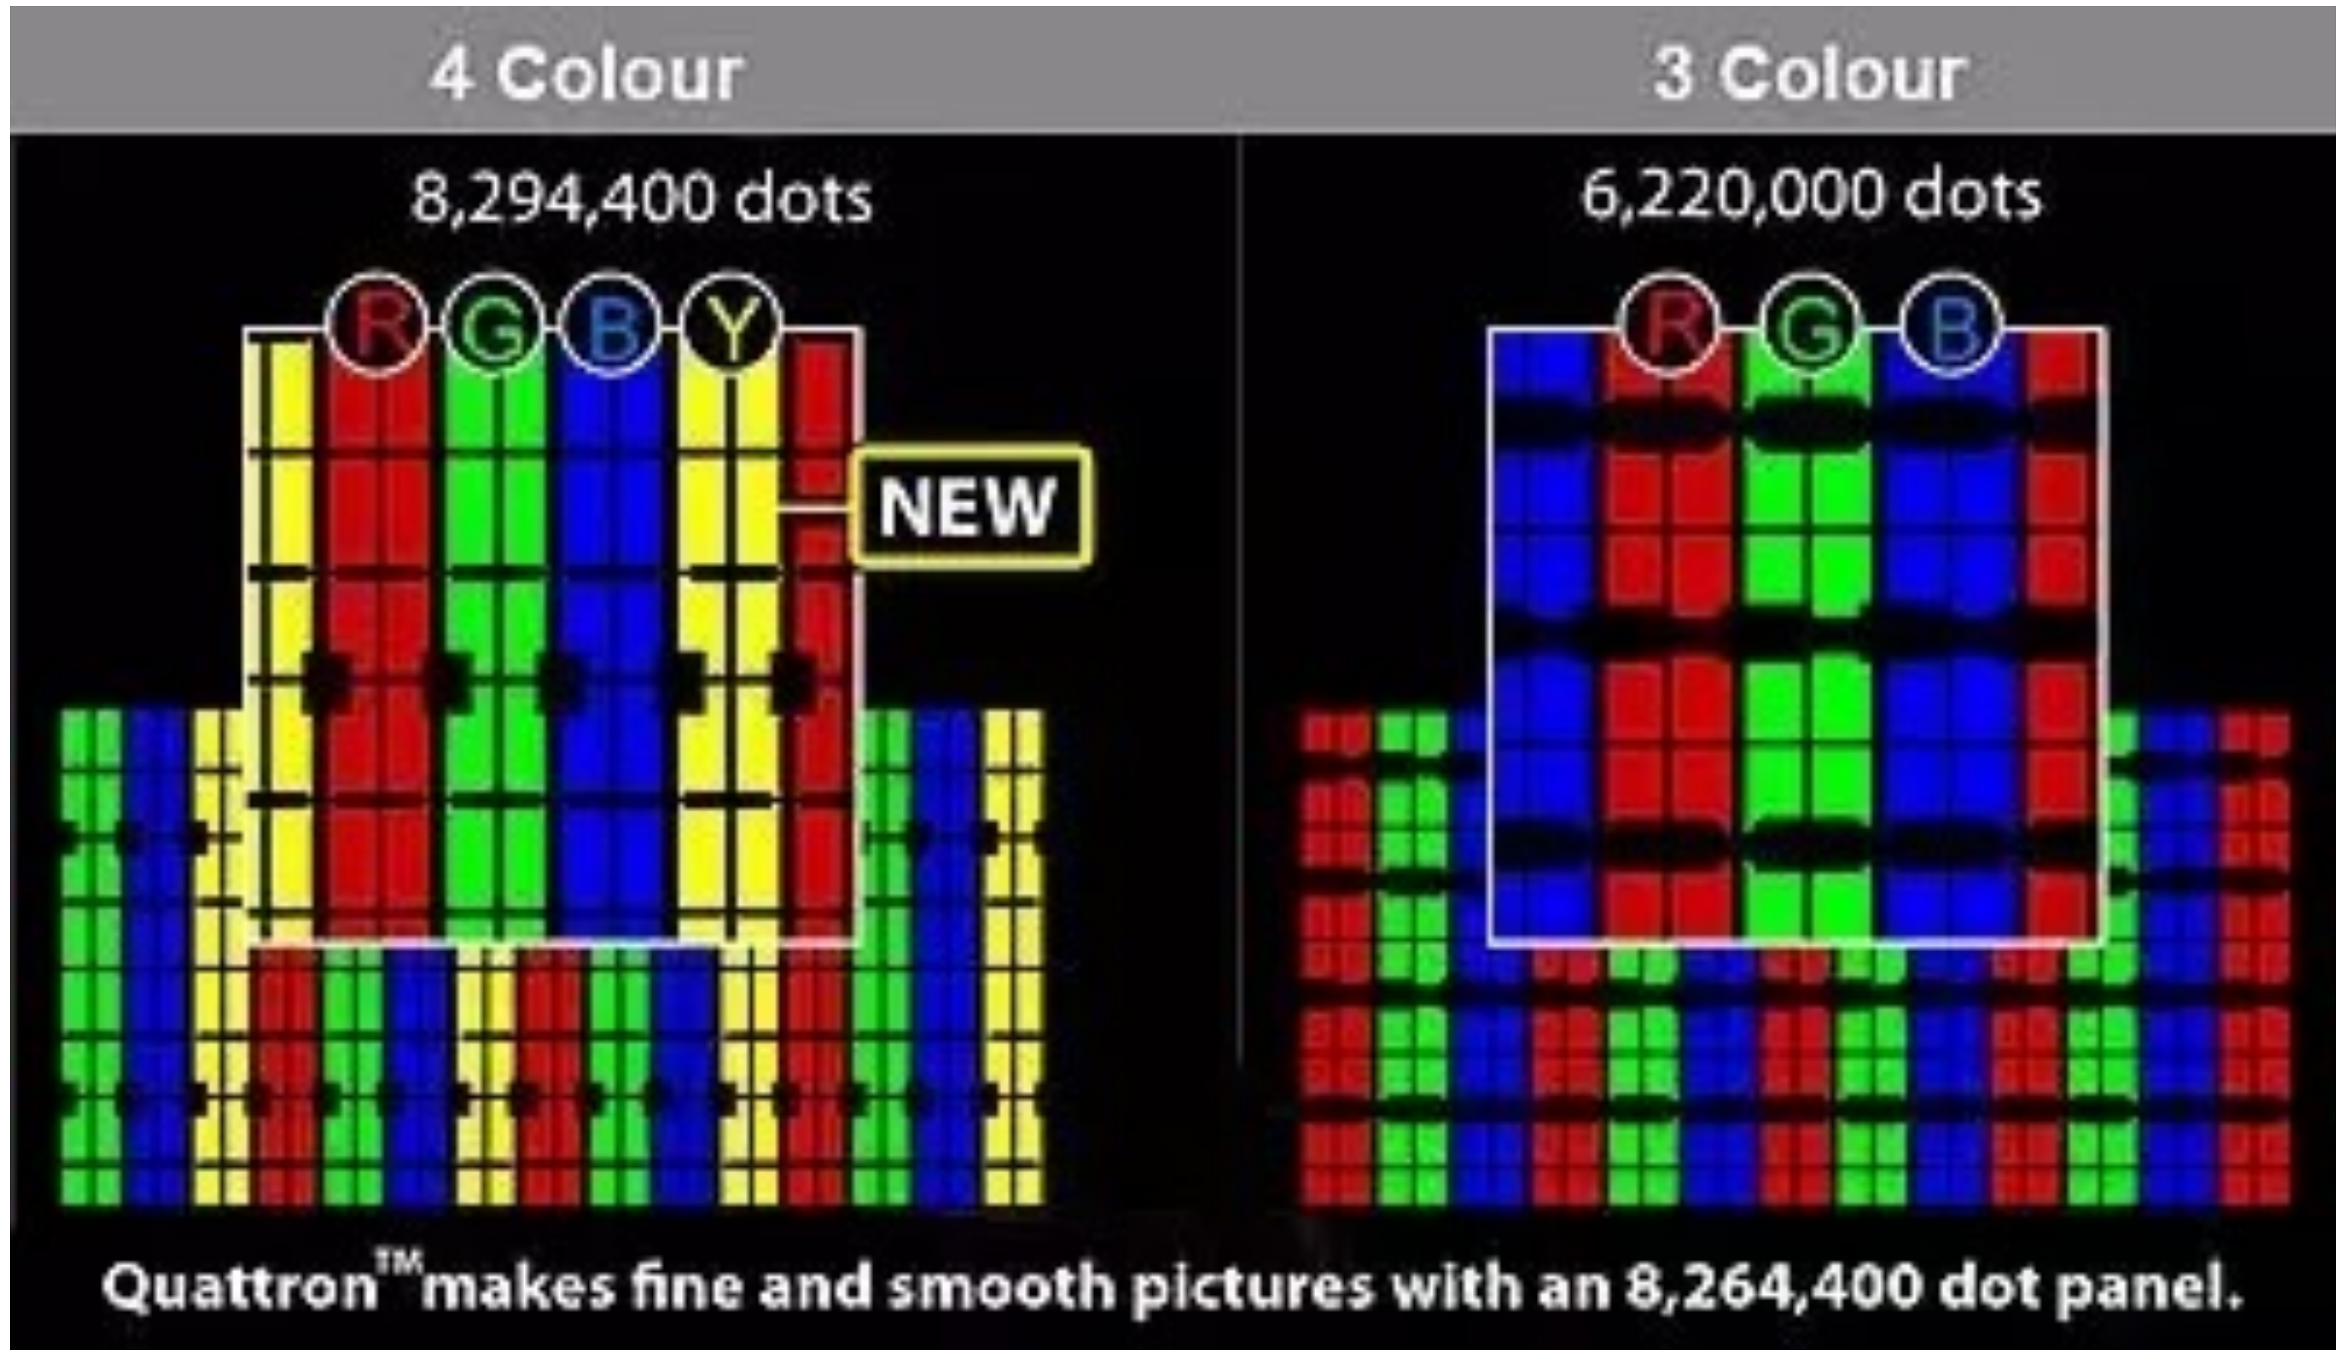

Color Matching Experiments

— Many colours can be represented as a positive weighted sum of A, B, C

— Write

$$M = aA + bB + cC$$

where the = sign should be read as “matches”

— This is **additive** matching

— Defines a colour description system
  — two people who agree on A, B, C need only supply (a, b, c)

Color Matching Experiments

— Some colours can’t be matched this way

— Instead, we must write

$$M + aA = bB + cC$$

where, again, the = sign should be read as “matches”

— This is **subtractive** matching

— Interpret this as (–a, b, c)

Color Matching Experiments

— Some colours can't be matched this way

— Instead, we must write

$$M + aA = bB + cC$$

where, again, the = sign should be read as "matches"

— This is **subtractive** matching

— Interpret this as (–a, b, c)

Problem for **designing displays**: Choose phosphors R, G, B so that **positive linear combinations** match a large set of colours

Principles of **Trichromacy**

Experimental facts:

Three primaries work for most people, provided we allow subtractive matching

— Exceptional people can match with two or only one primary

— This likely is caused by biological deficiencies

Most people make the same matches

— There are some anomalous trichromats, who use three primaries but match with different combinations

Grassman's Laws

For colour matches:

- **symmetry**: $U = V \Leftrightarrow V = U$
- **transitivity**: $U = V \text{ and } V = W \Rightarrow U = W$
- **proportionality**: $U = V \Leftrightarrow tU = tV$
- **additivity**: if any two of the statements are true, then so is the third

$$U = V,$$
$$W = X,$$
$$(U + W) = (V + X)$$

These statements mean that colour matching is, to an accurate approximation, linear.

Additive vs. **Subtractive** Color

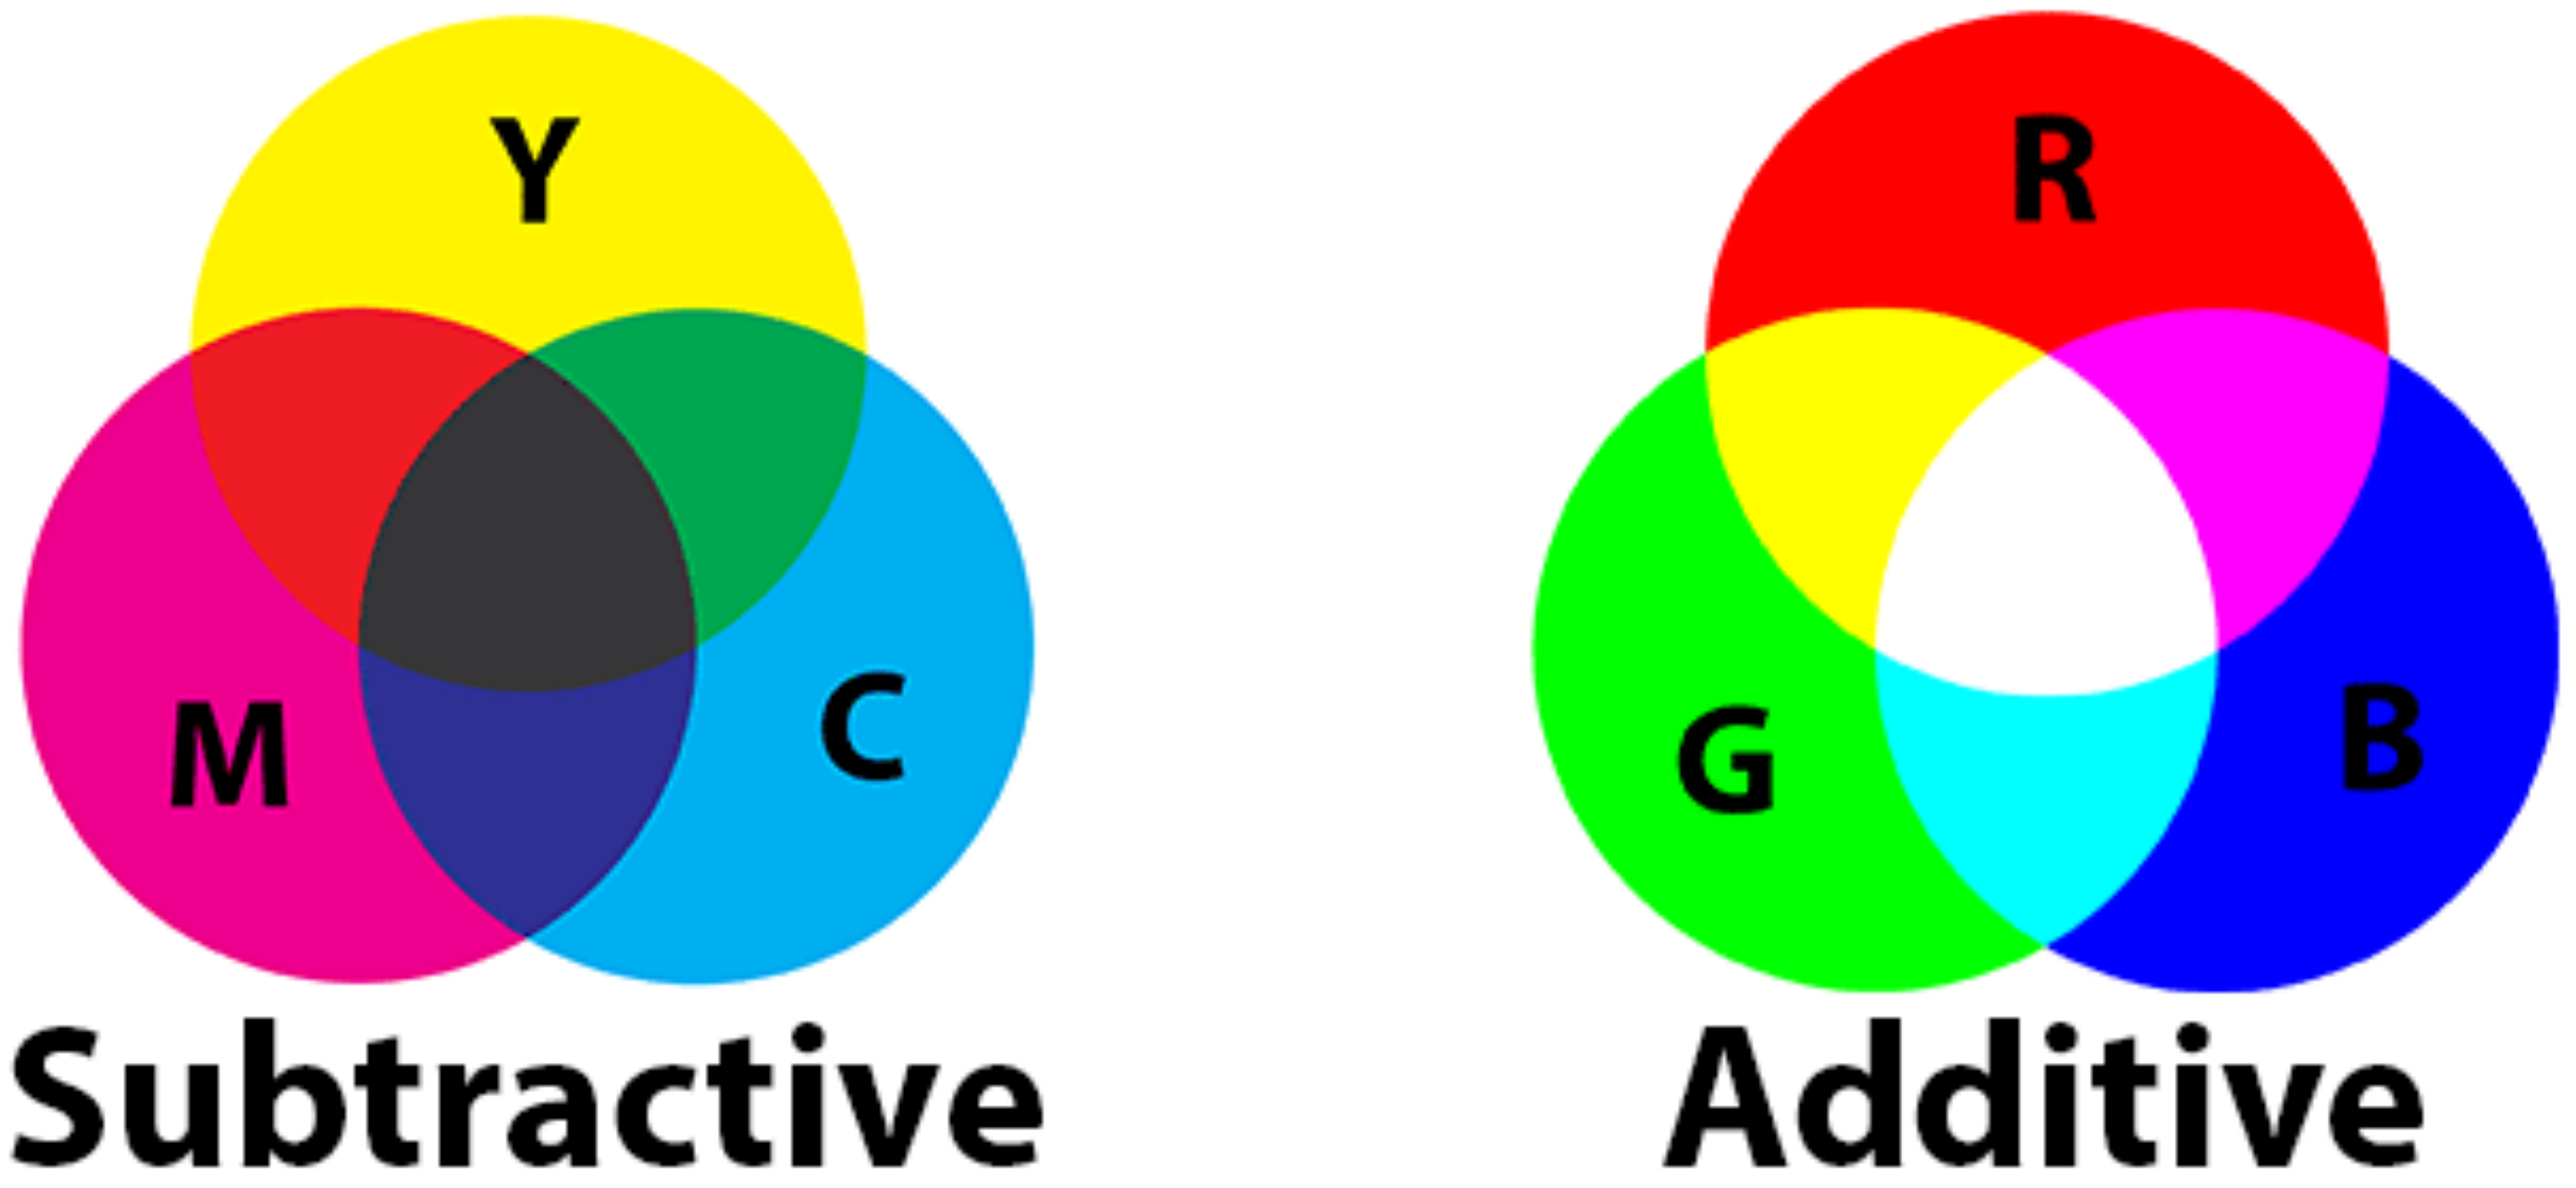

Human **Cone** Sensitivity

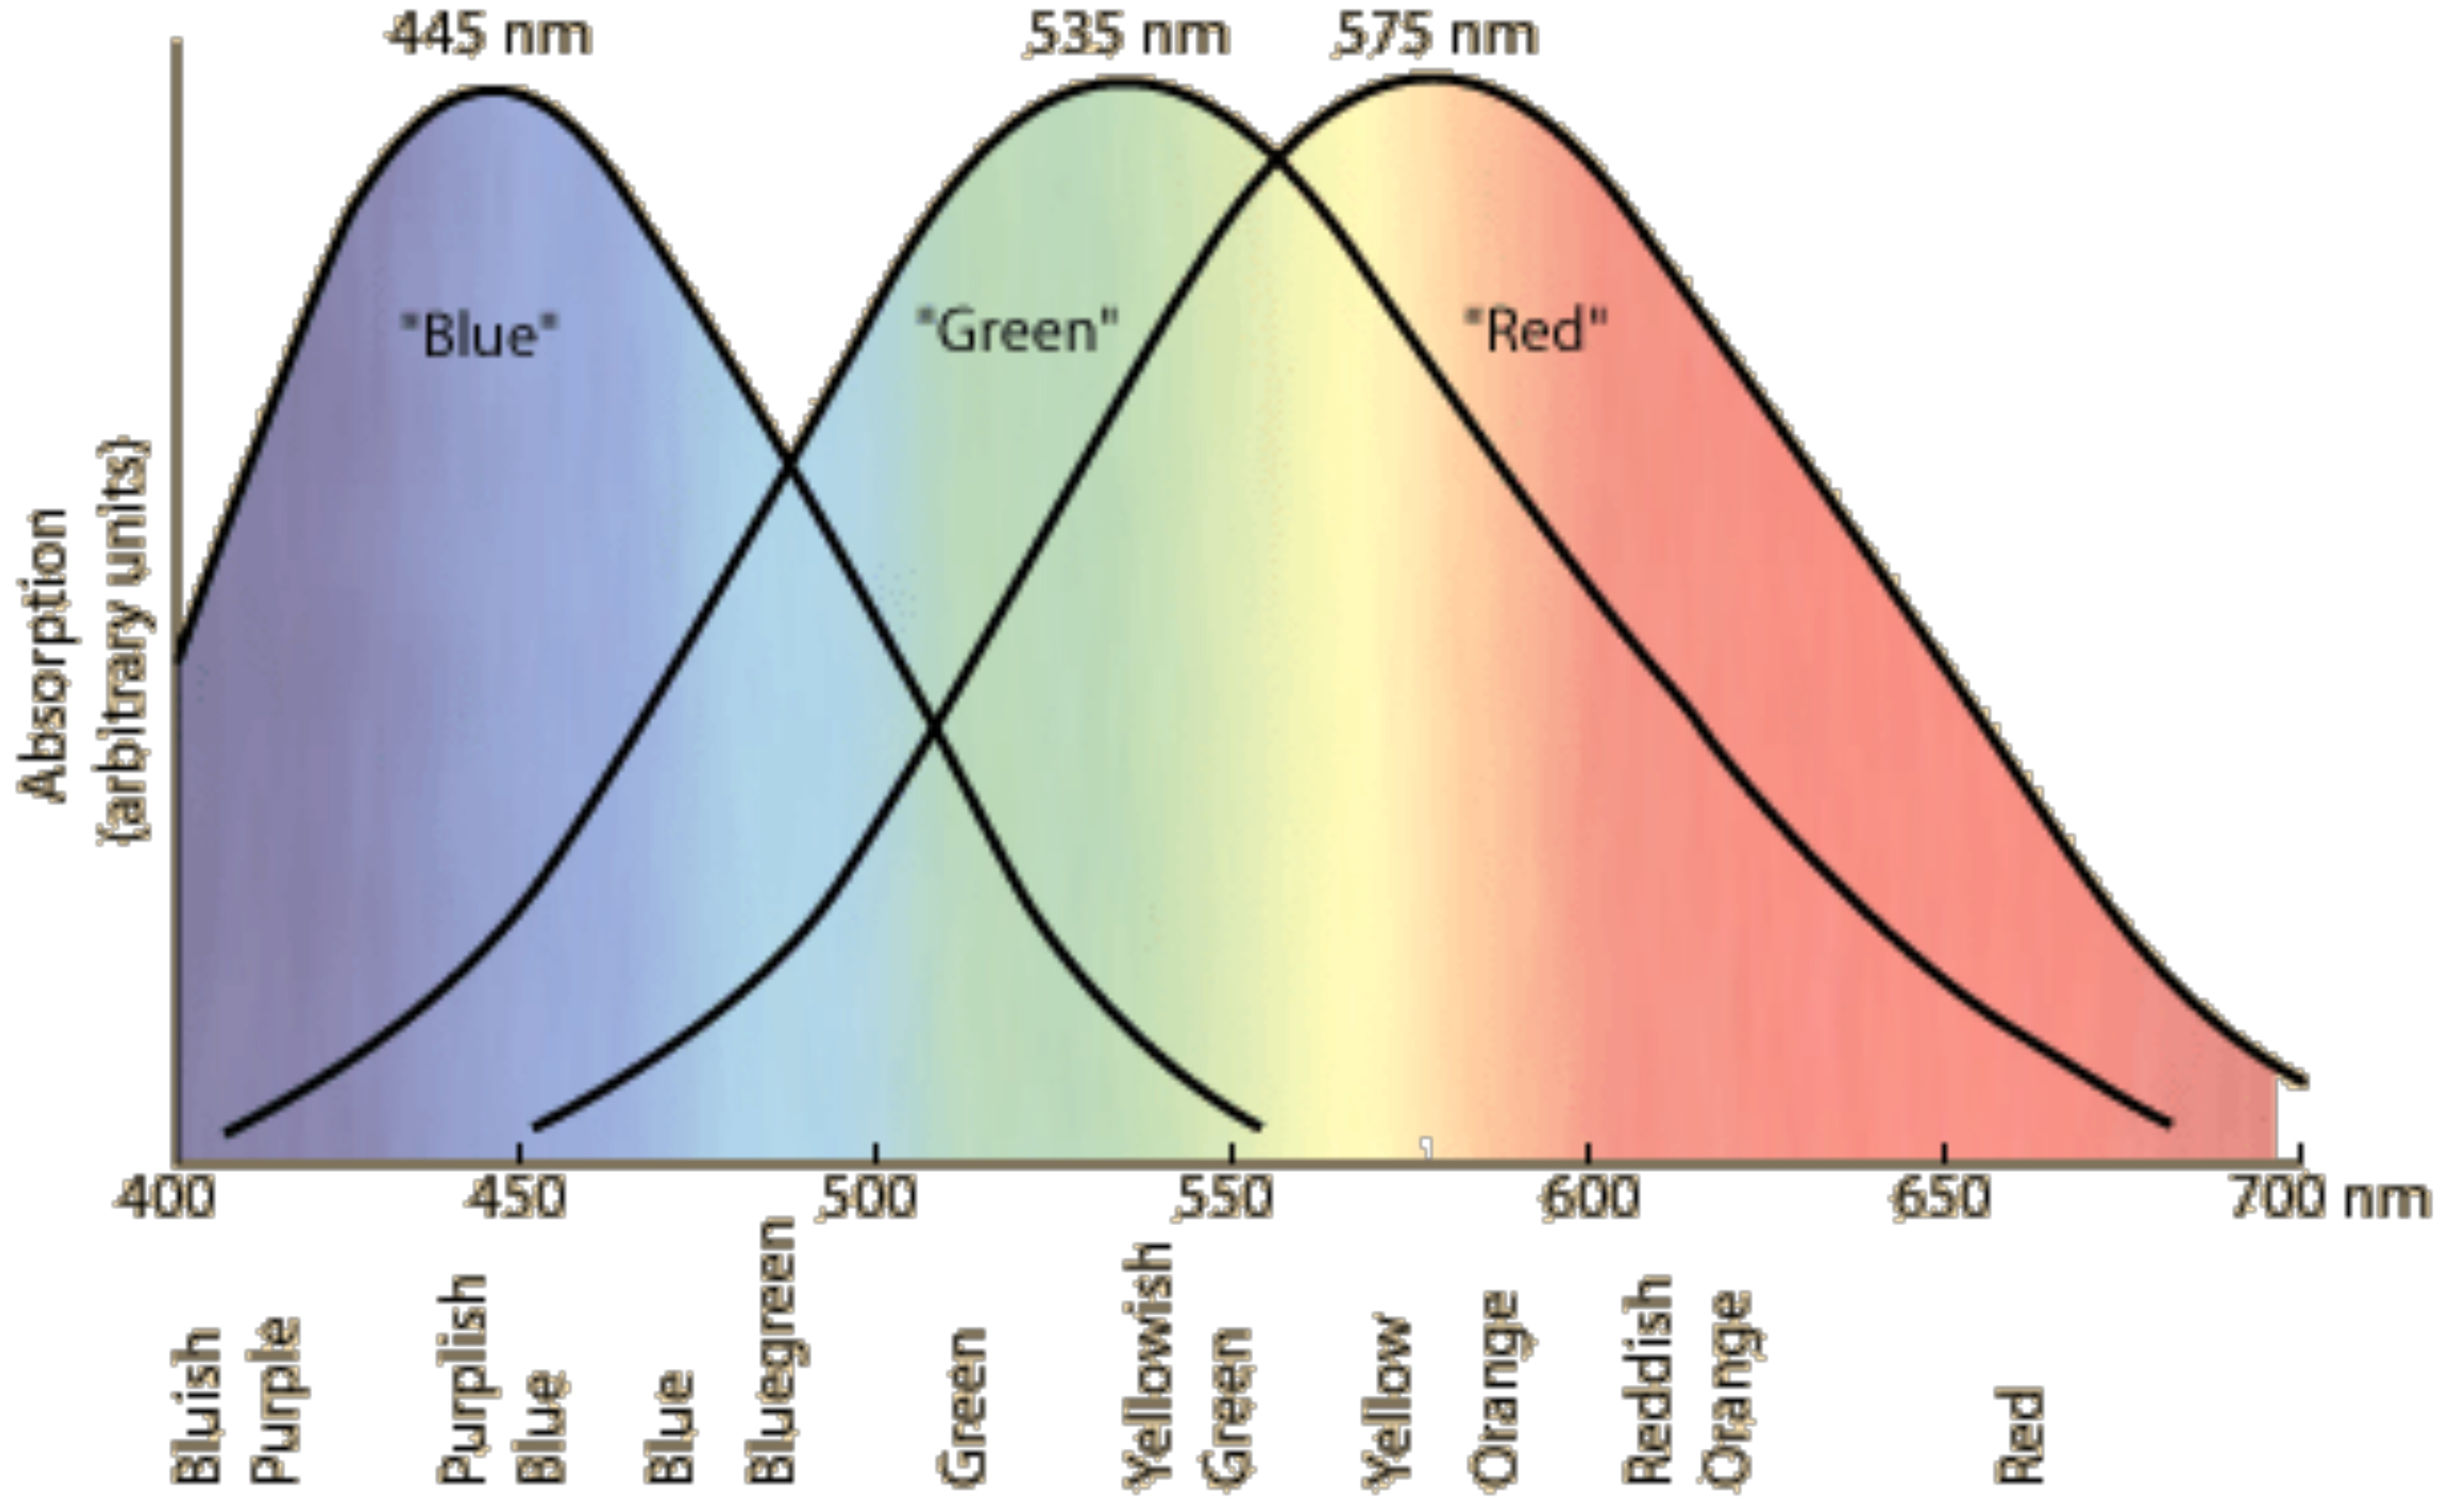

http://hyperphysics.phy-astr.gsu.edu/hbase/vision/colcon.html

Representing Colour

— Describing colours accurately is of practical importance (e.g. Manufacturers are willing to go to a great deal of trouble to ensure that different batches of their product have the same colour)

— This requires a standard system for representing colour.

Color Spaces

— **RGB**: Primaries are monochromatic energies, say 645.2 nm, 526.3 nm, 444.4 nm, standard colour space related to displays

— **CIE XYZ**: Primaries are imaginary, but have other convenient properties. Colour coordinates are (X , Y , Z), where X is the amount of the X primary, etc.

— **CIE LAB**: Equal distances in space correspond to perceptually uniform colour differences

— **HSV**: Hue, Saturation, Value a useful colour space for artists and **colour selection** applications

— **YCbCr:** Separates **luminance** (Y) and **opponent colours** (CbCr) which are

RGB Colour Matching Functions

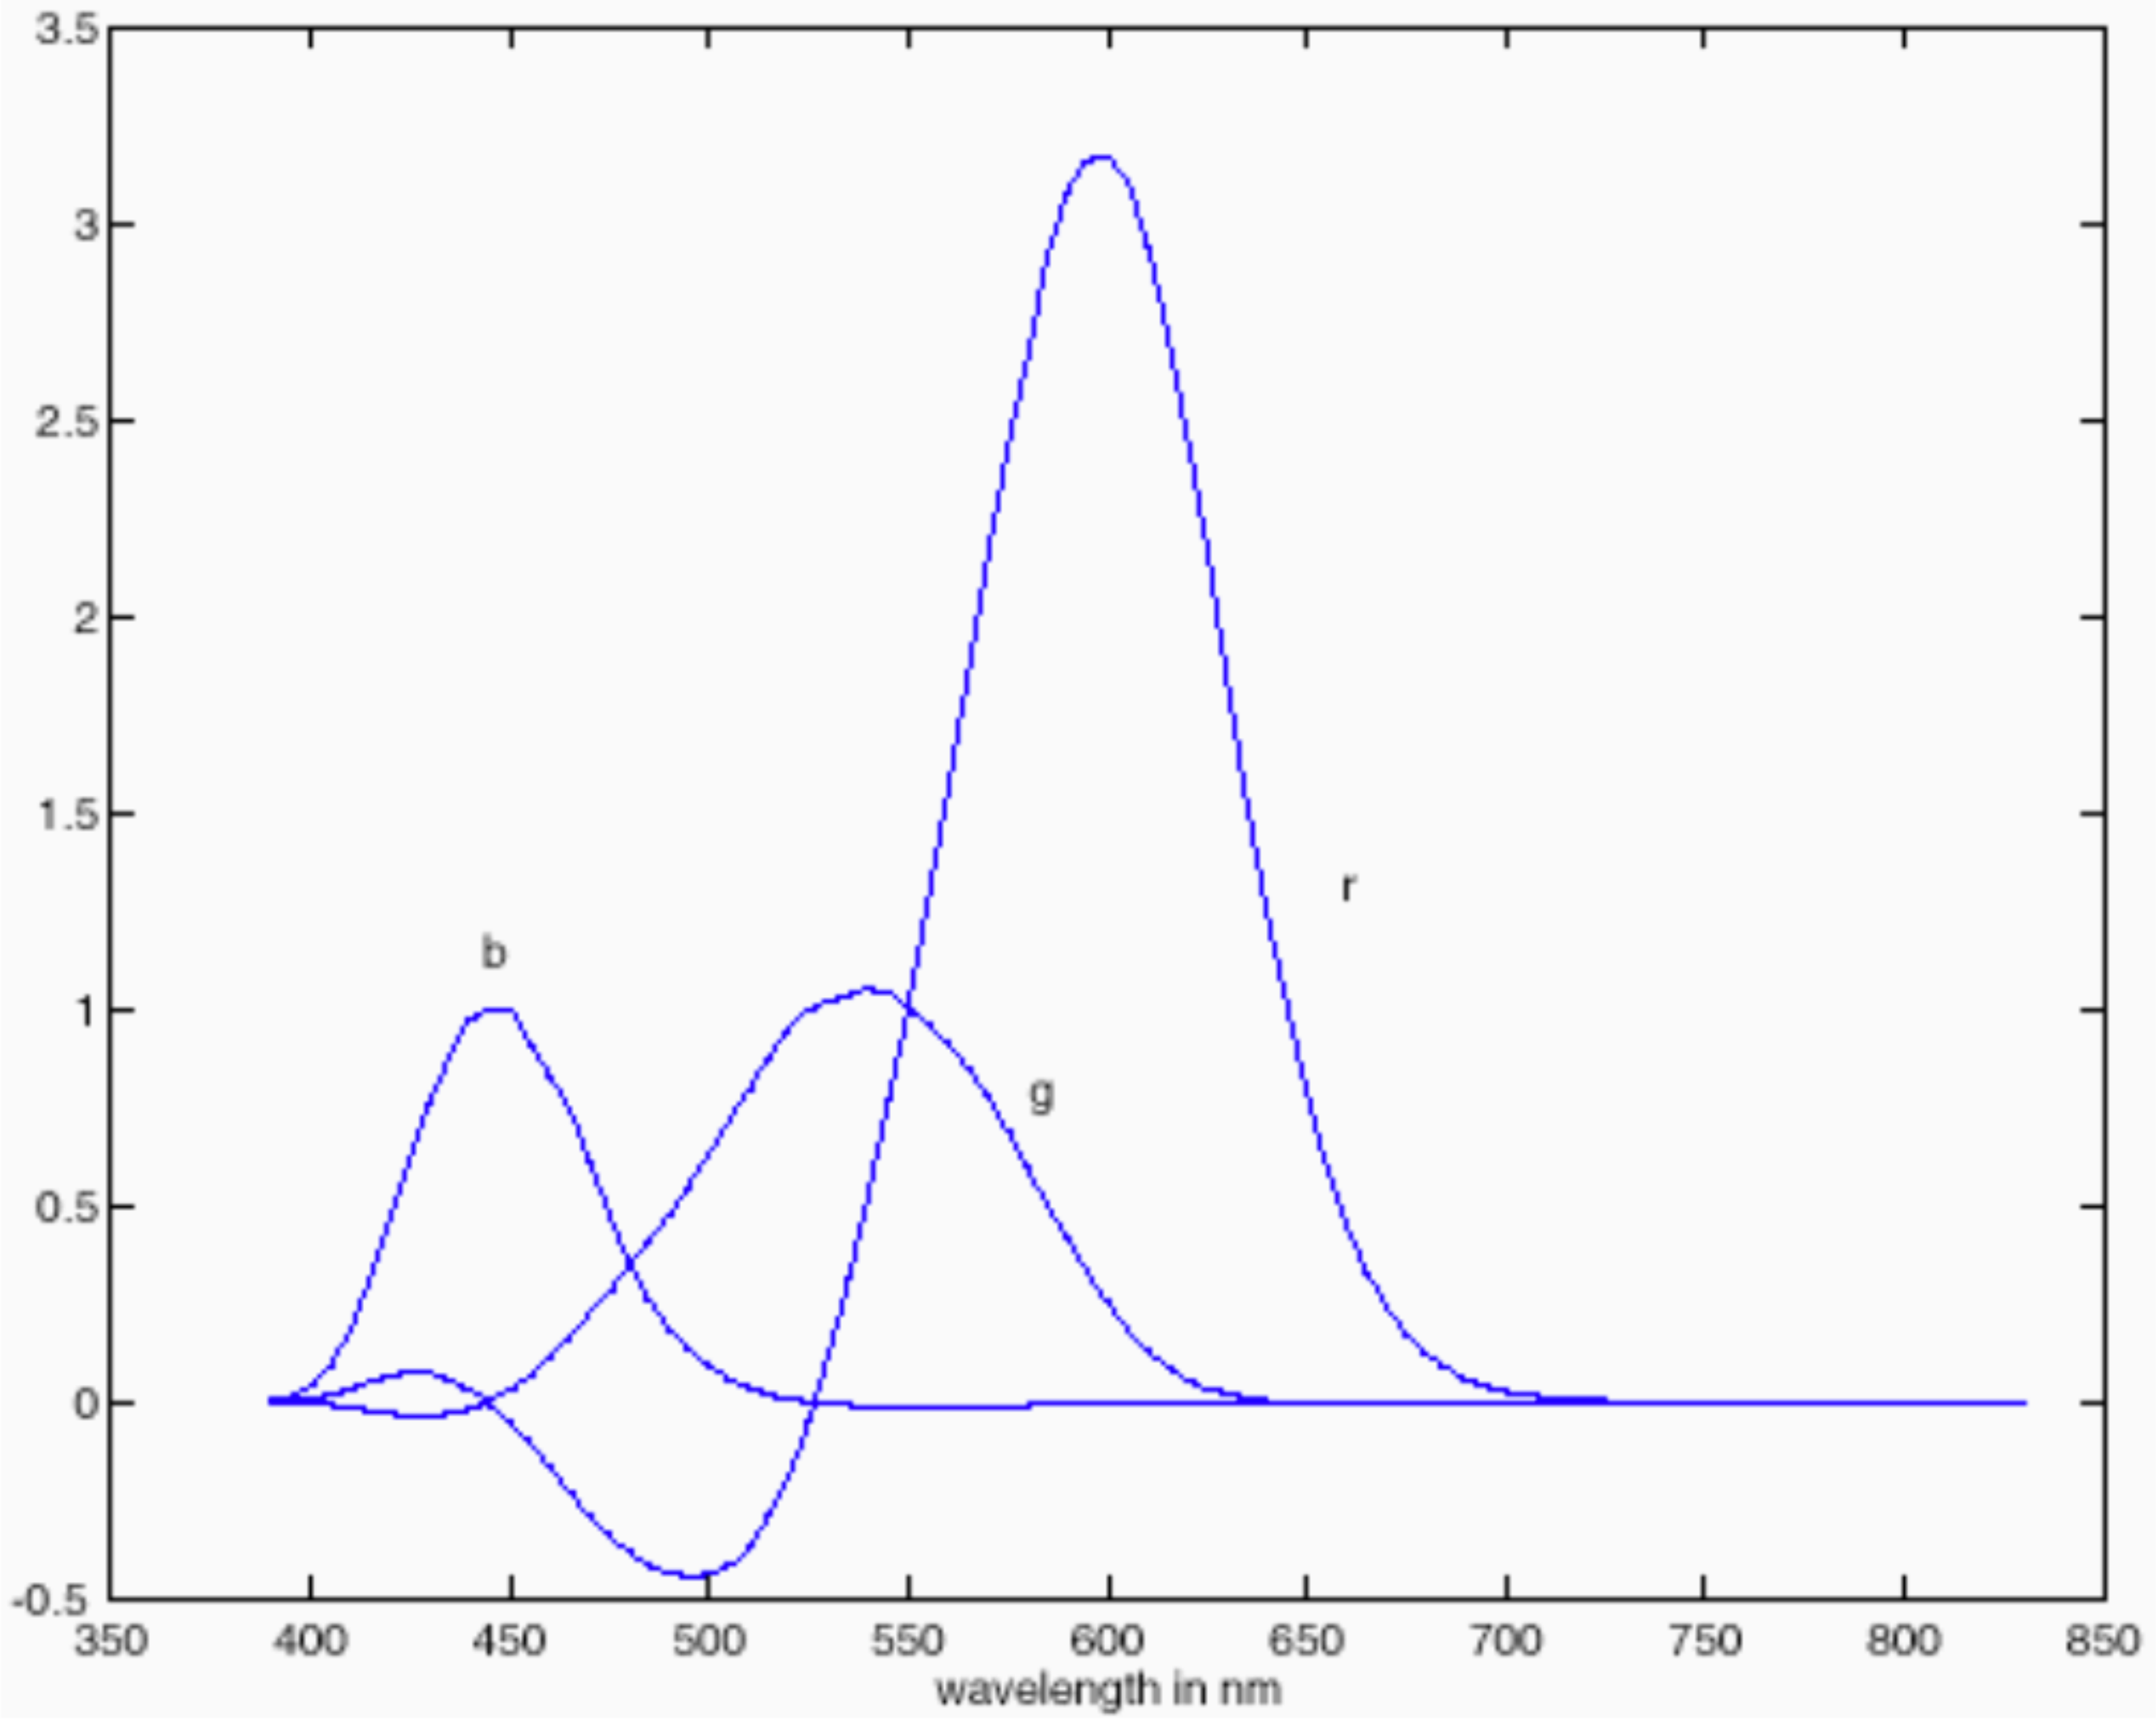

Forsyth & Ponce (2nd ed.) Figure 3.9

— Primaries monochromatic

— Wavelengths 645.2, 526.3 and 444.4 nm

— Negative parts means some colours can be matched only subtractively

RGB Color Space

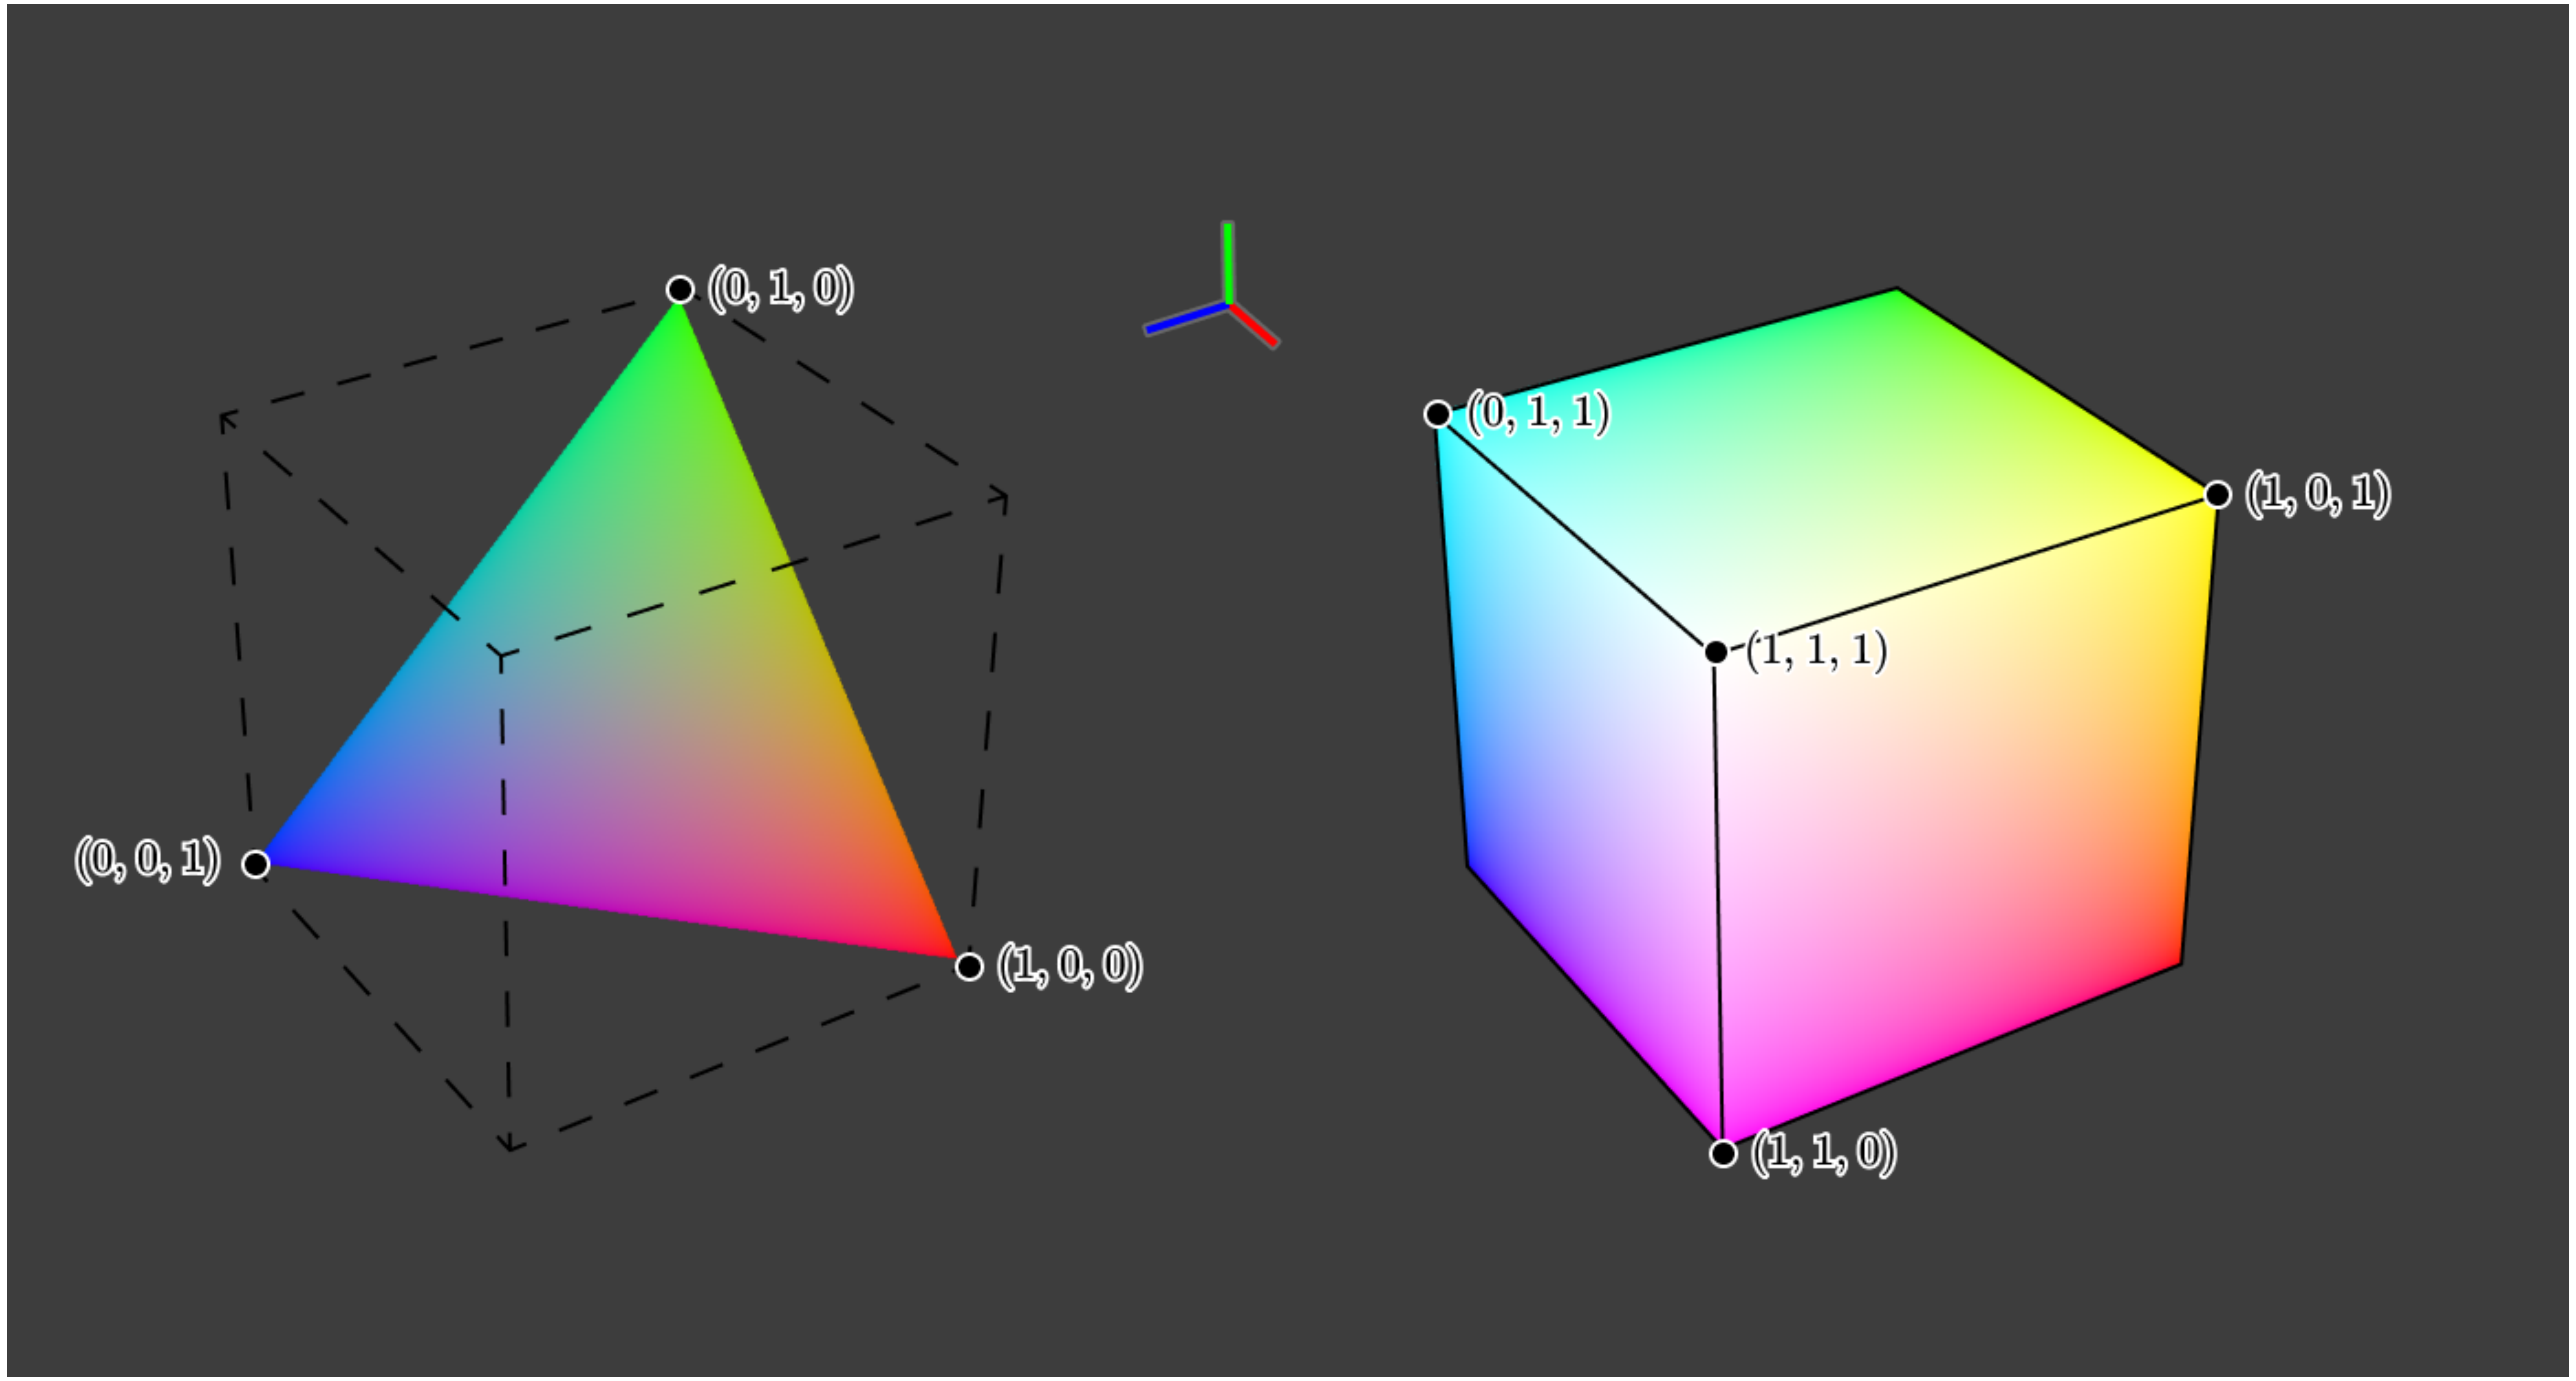

RGB Colour Matching Functions

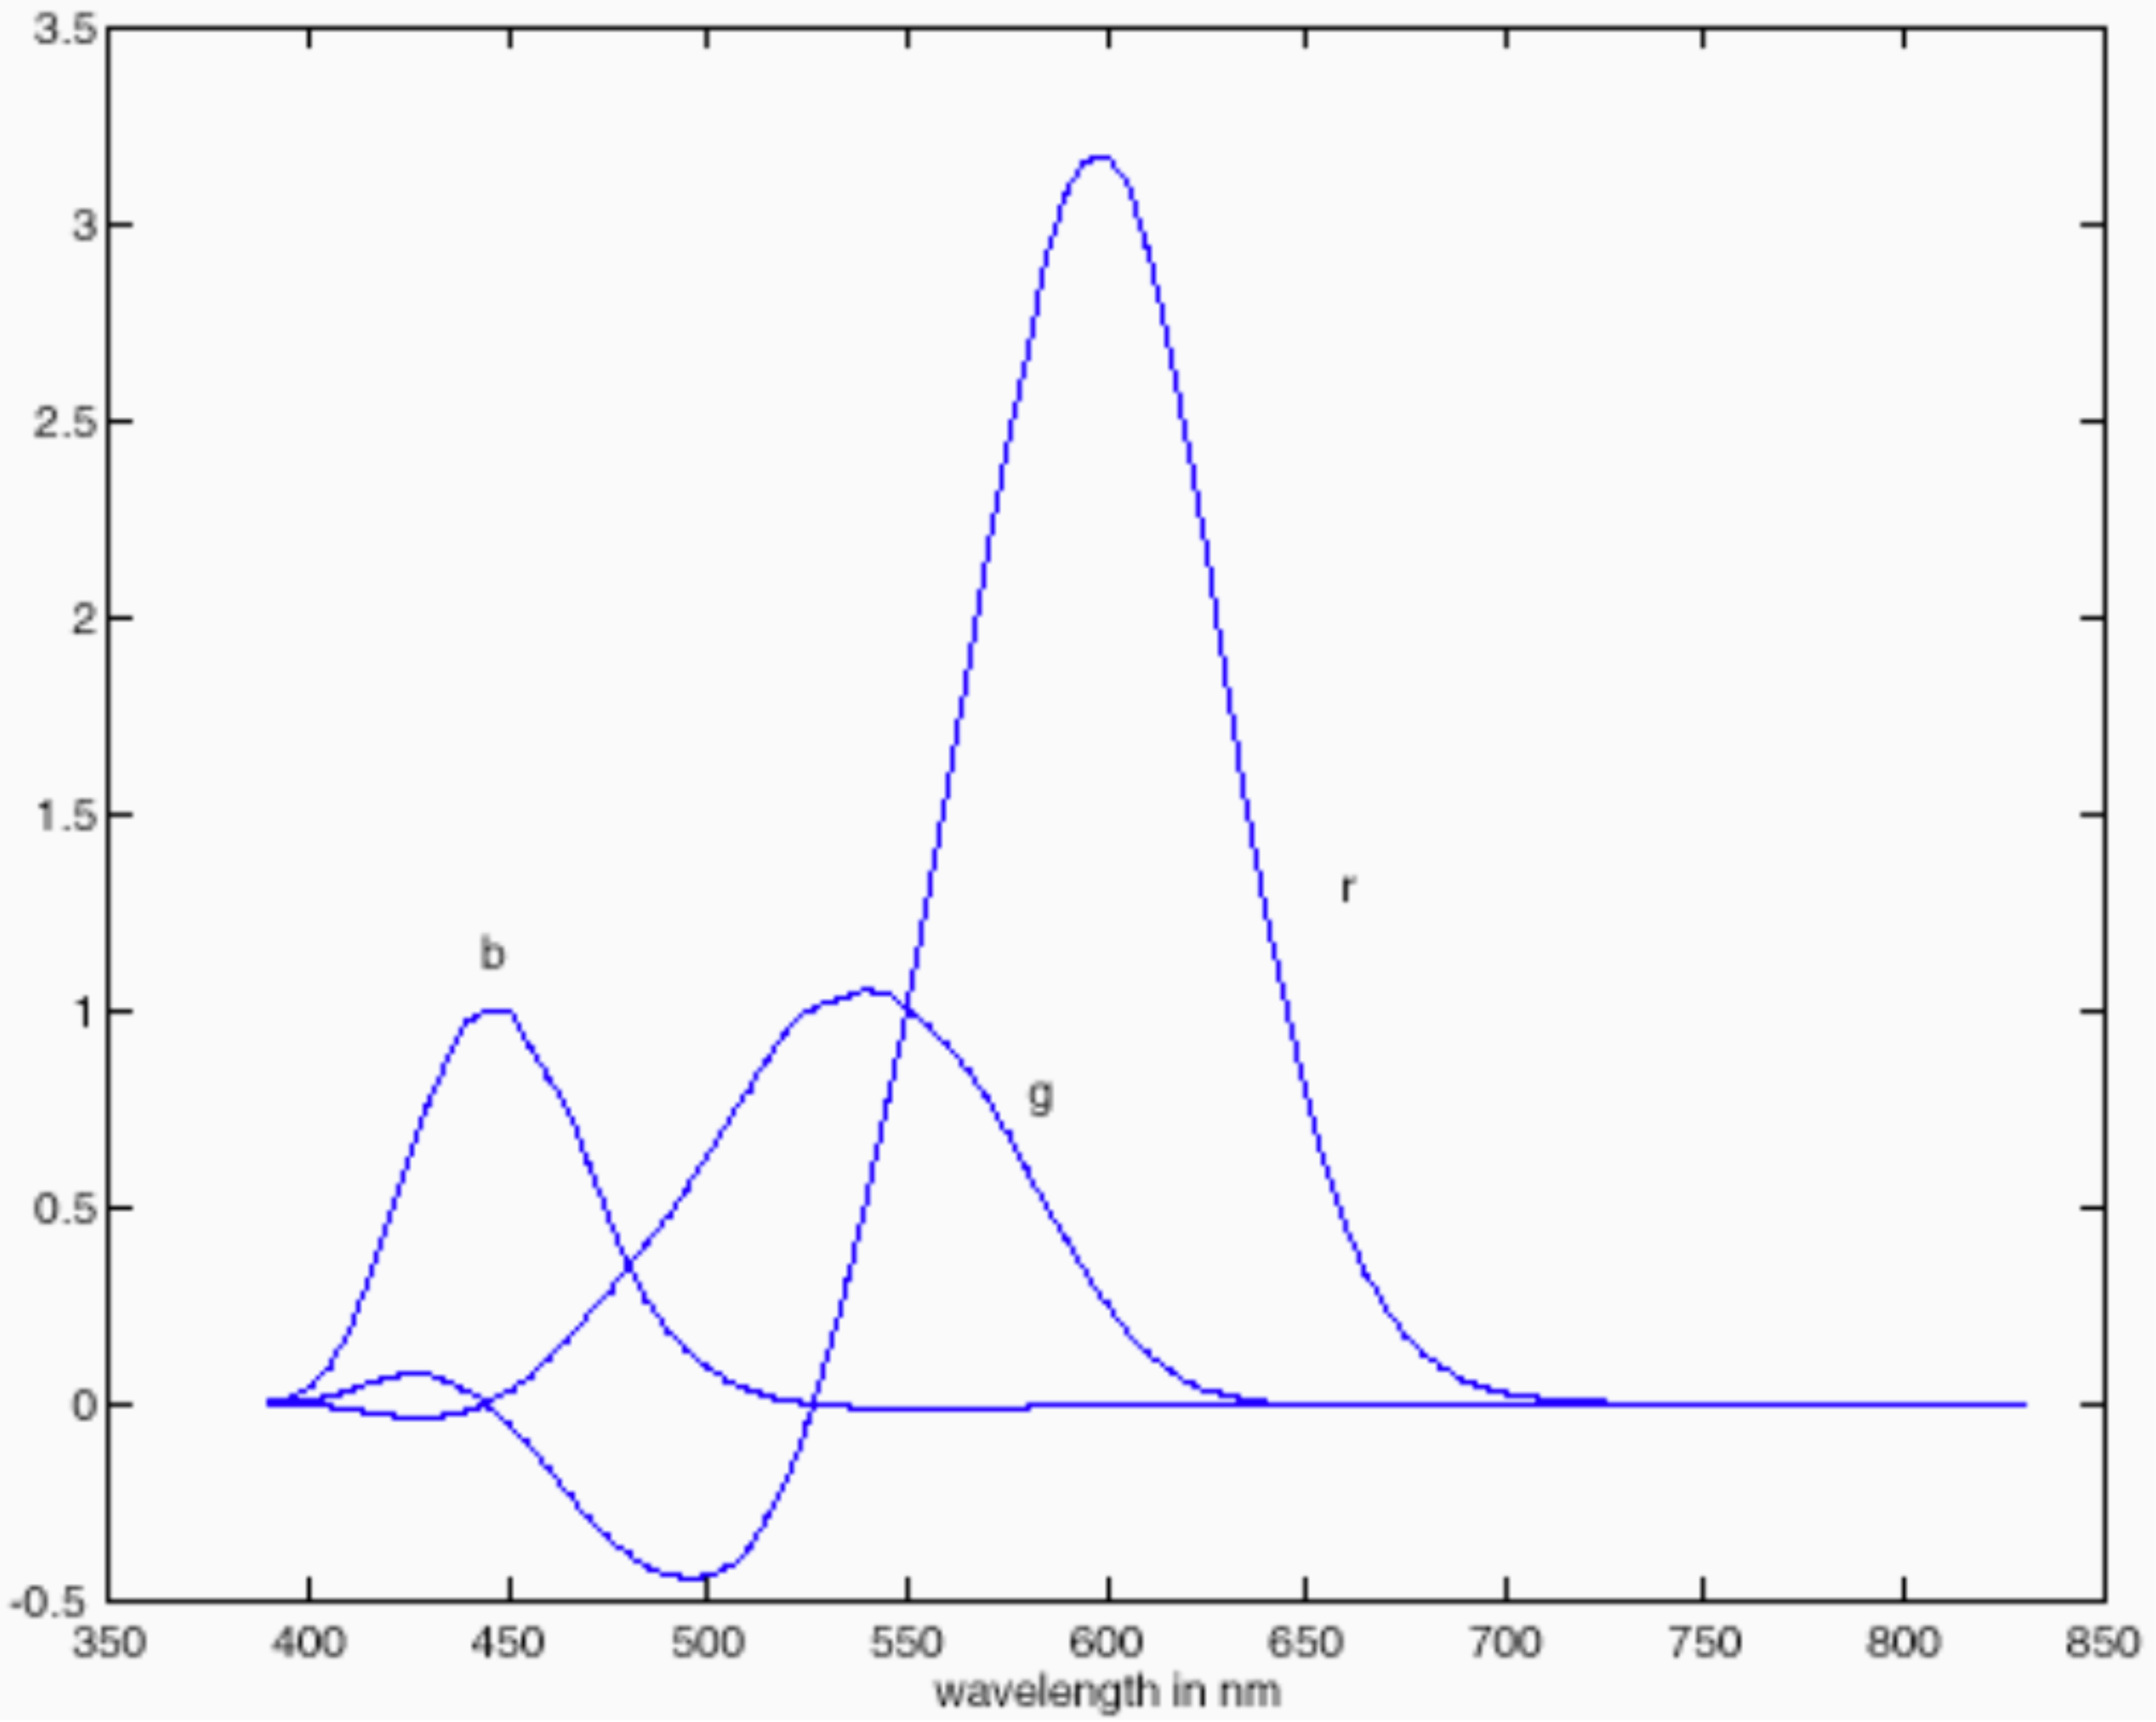

Forsyth & Ponce (2nd ed.) Figure 3.9

— Primaries monochromatic

— Wavelengths 645.2, 526.3 and 444.4 nm

— Negative parts means some colours can be matched only subtractively

RGB Colour Matching Functions

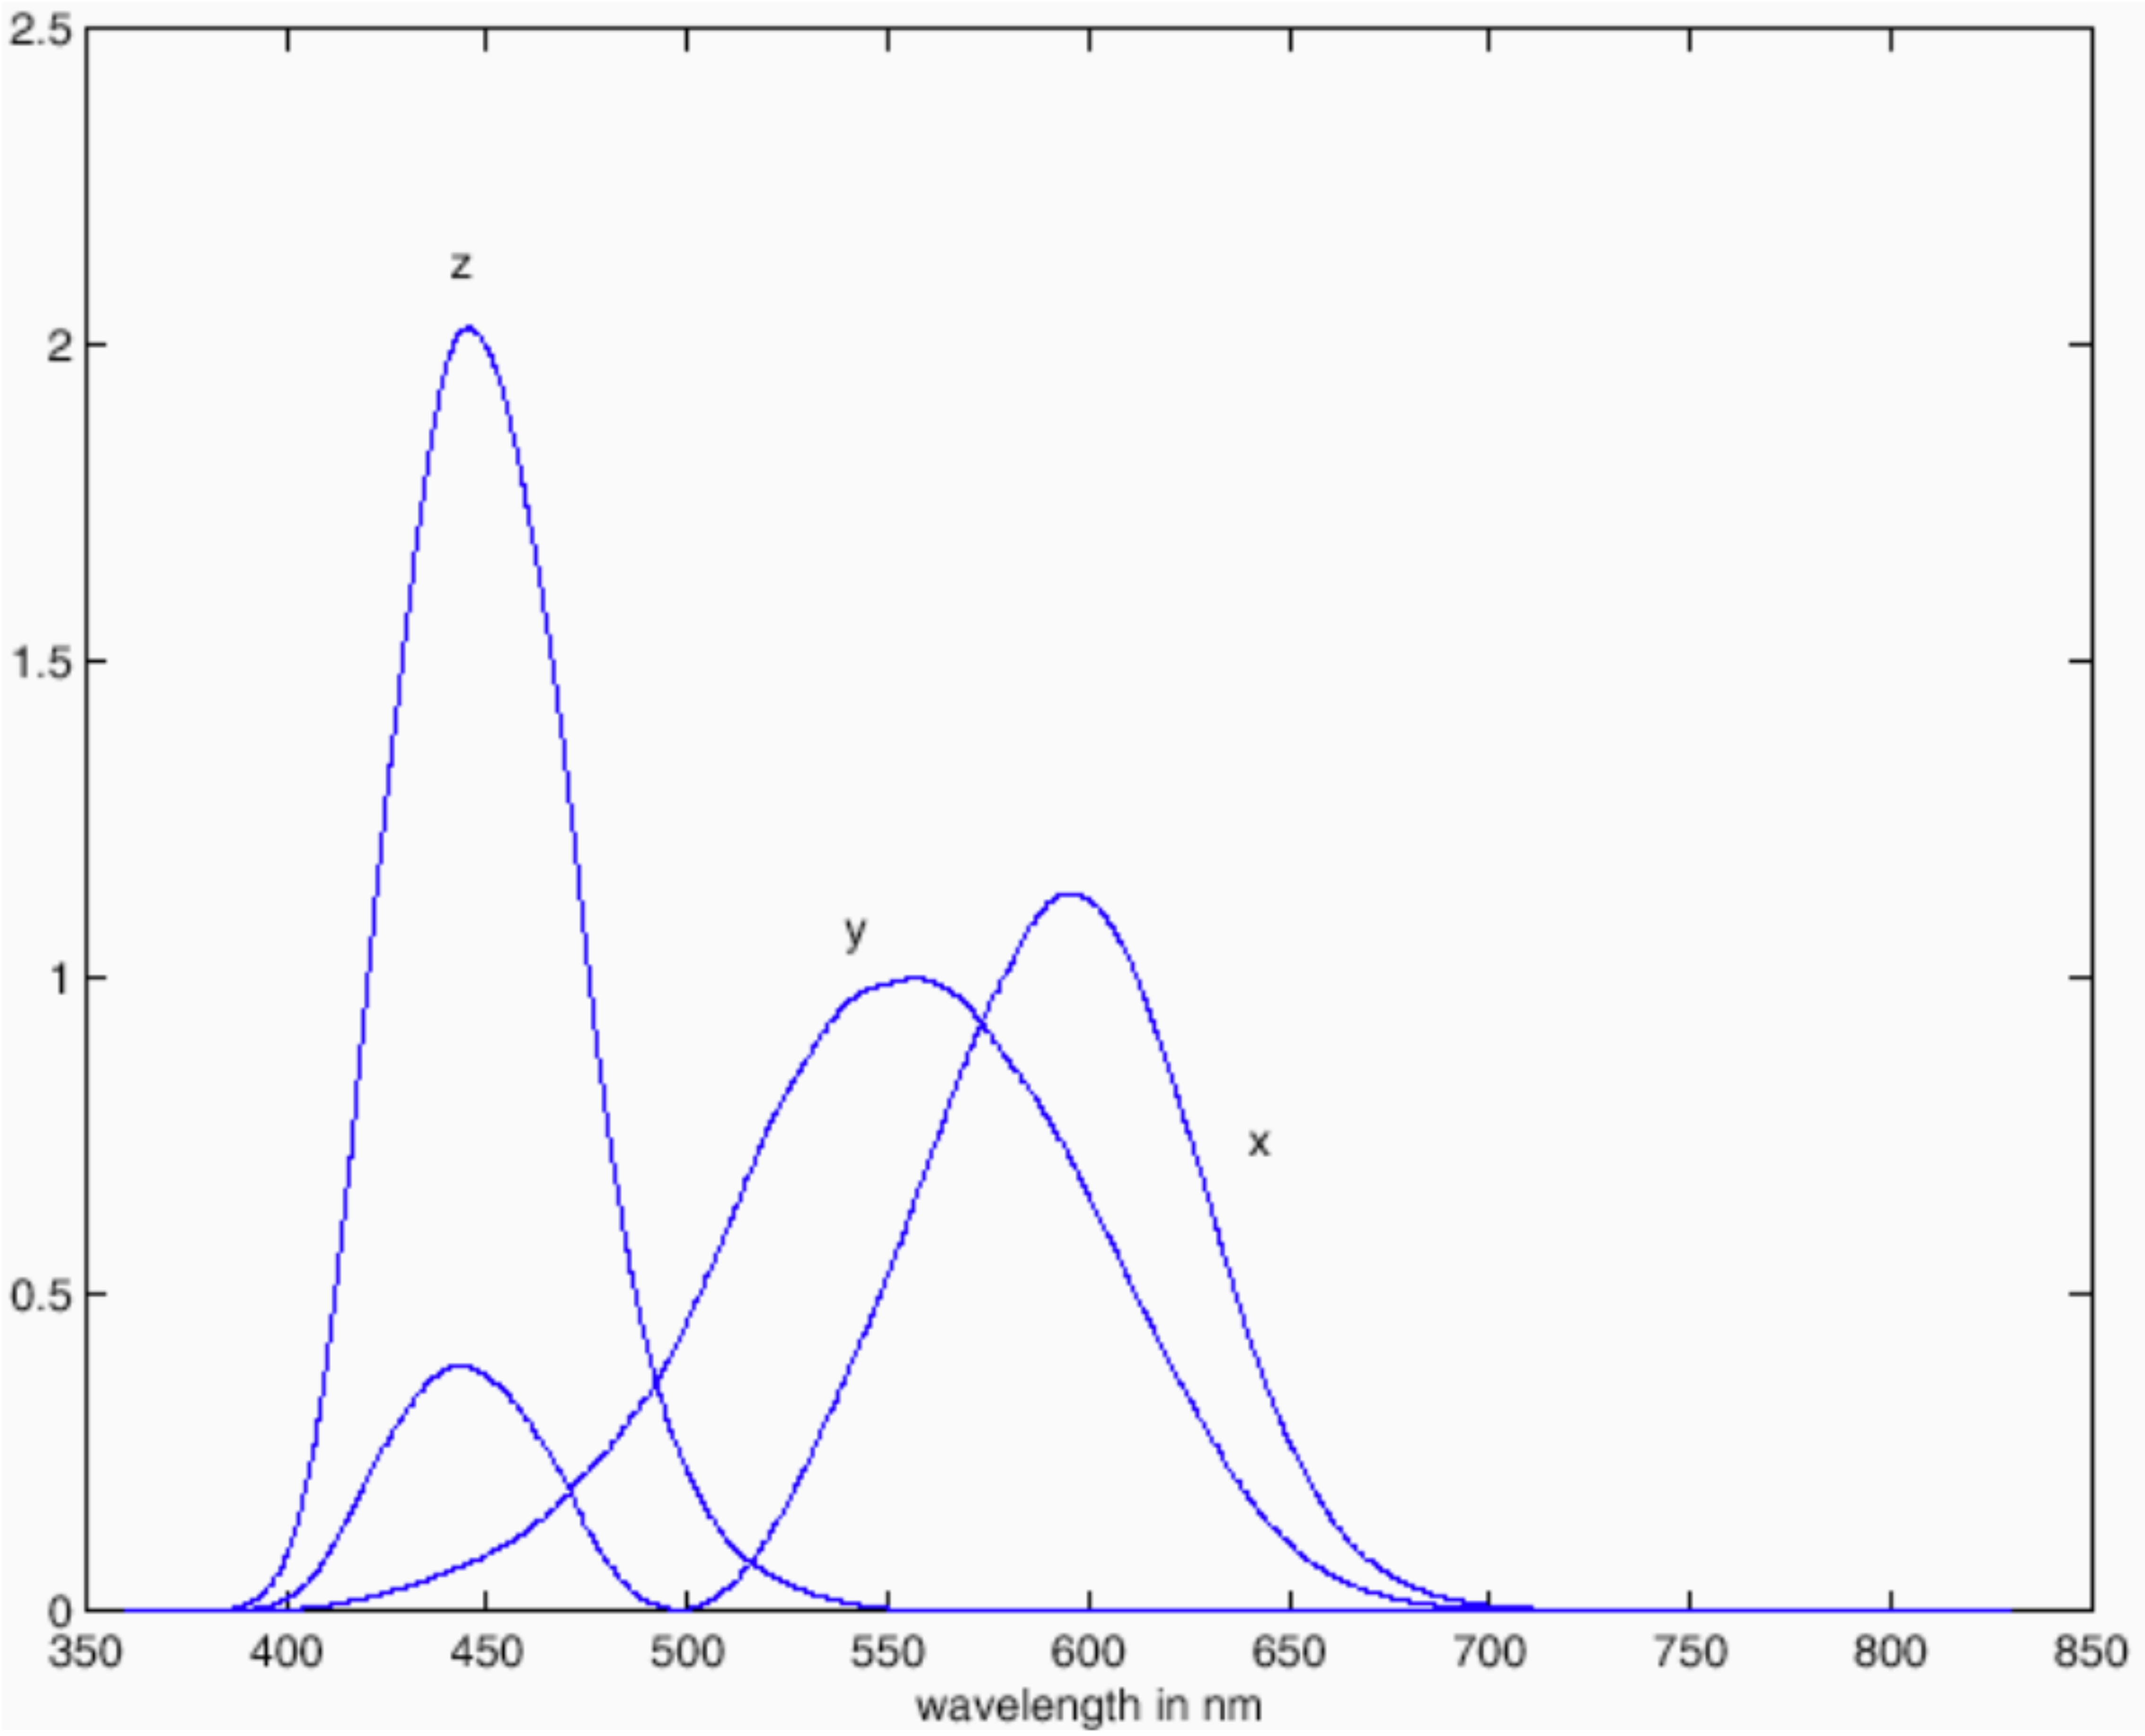

Forsyth & Ponce (2nd ed.) Figure 3.8

CIE XYZ: Colour matching functions are positive everywhere, but primaries are imaginary. Usually draw x, y, where

$$x = X/(X+Y+Z)$$
$$y = Y/(X+Y+Z)$$

Overall brightness is ignored

Geometry of Colour (CIE)

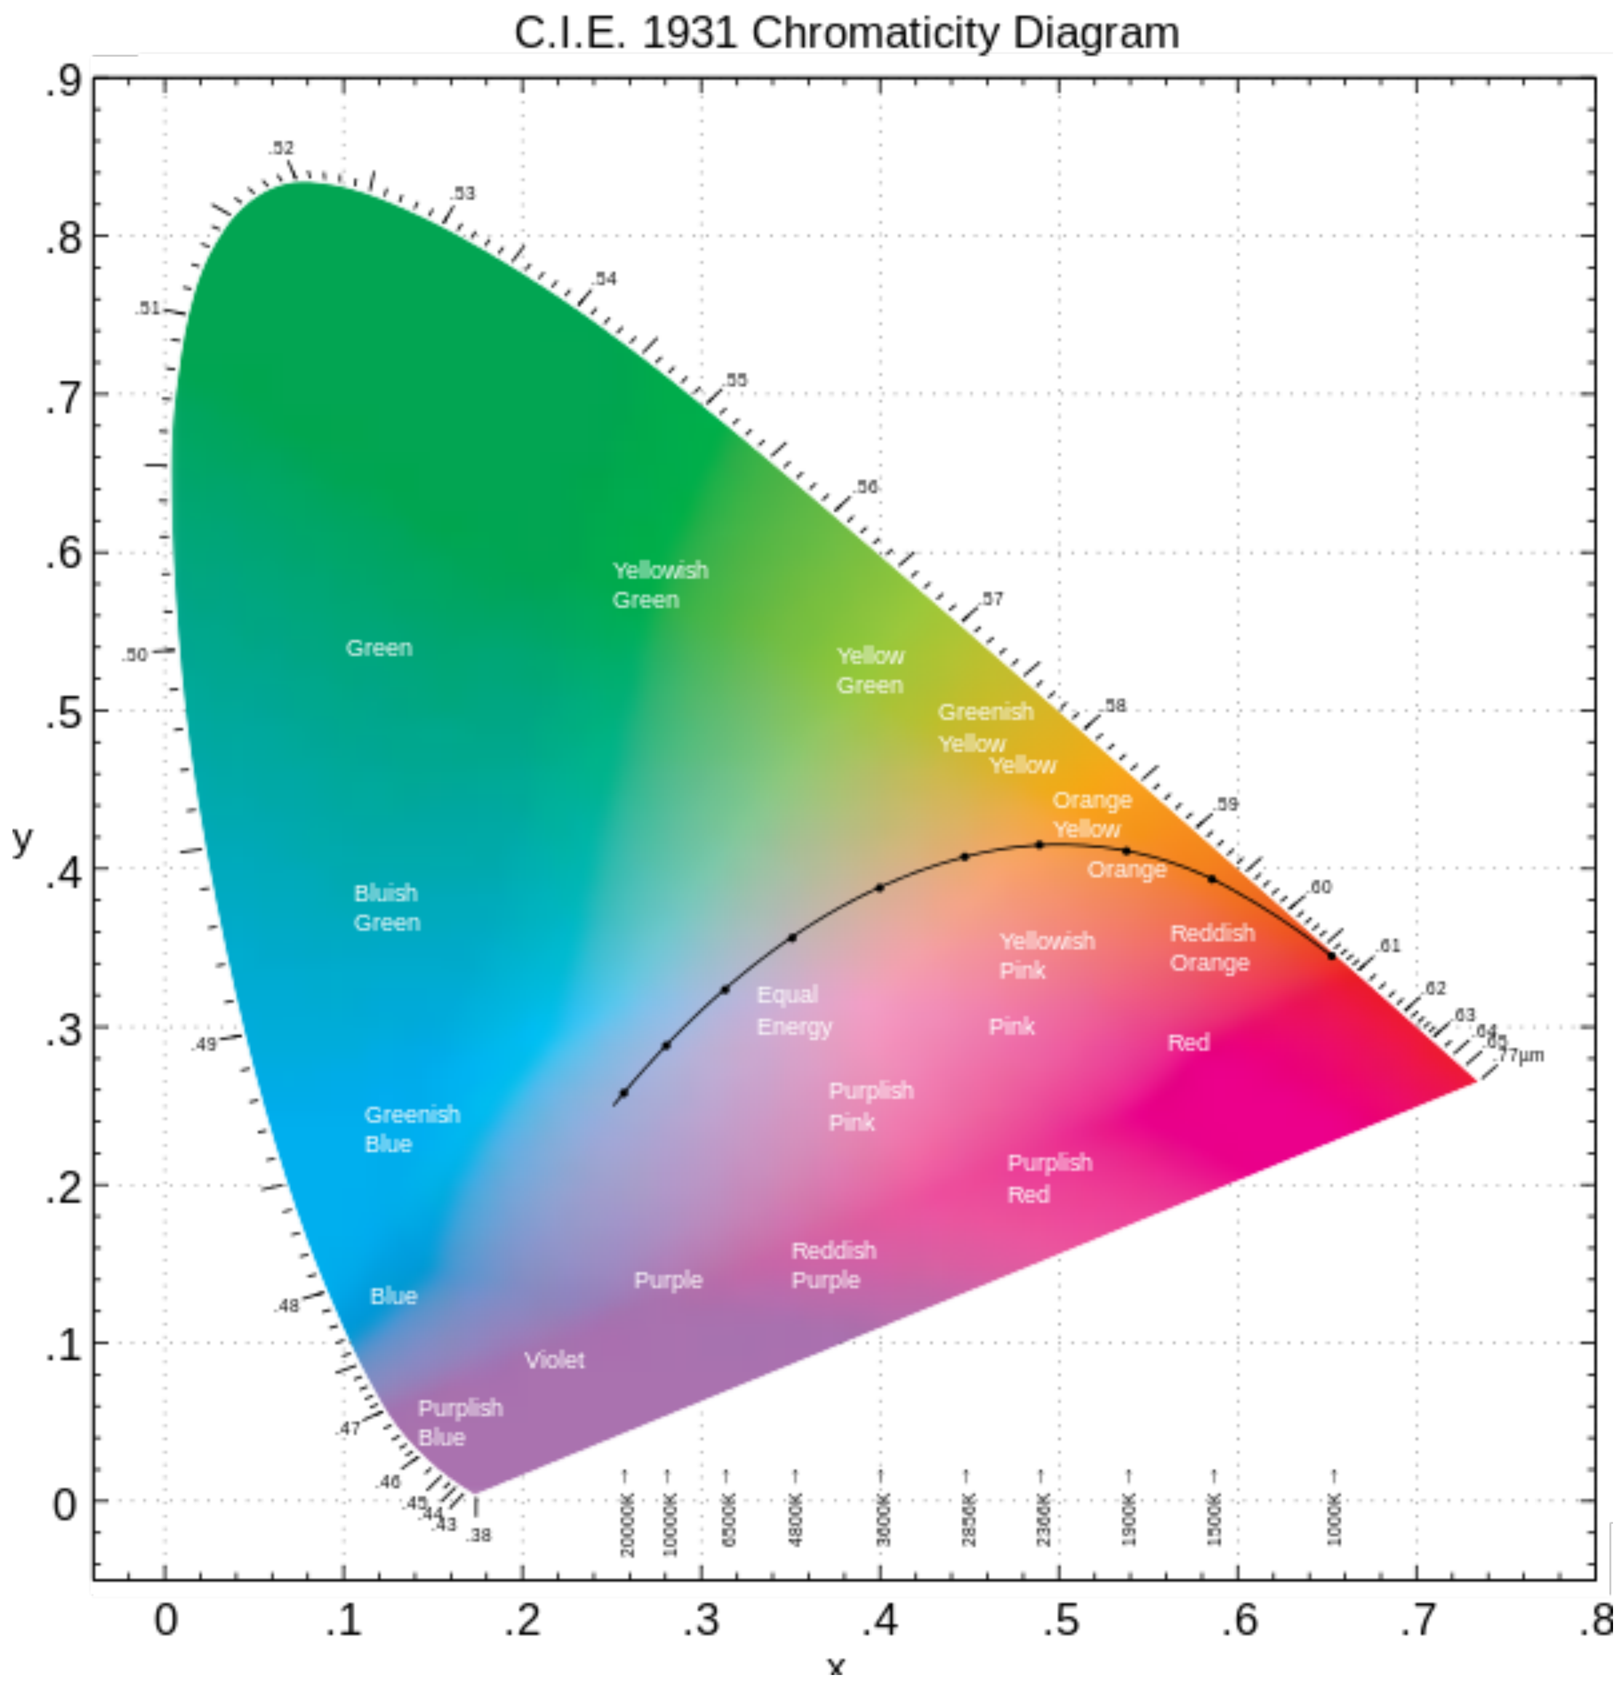

— White is in the center, with saturation increasing towards the boundary

— Mixing two coloured lights creates colours on a straight line

— Mixing 3 colours creates colours within a triangle

— Curved edge means there are no 3 actual lights that can create all colours that humans perceive!

RGB Colour Space

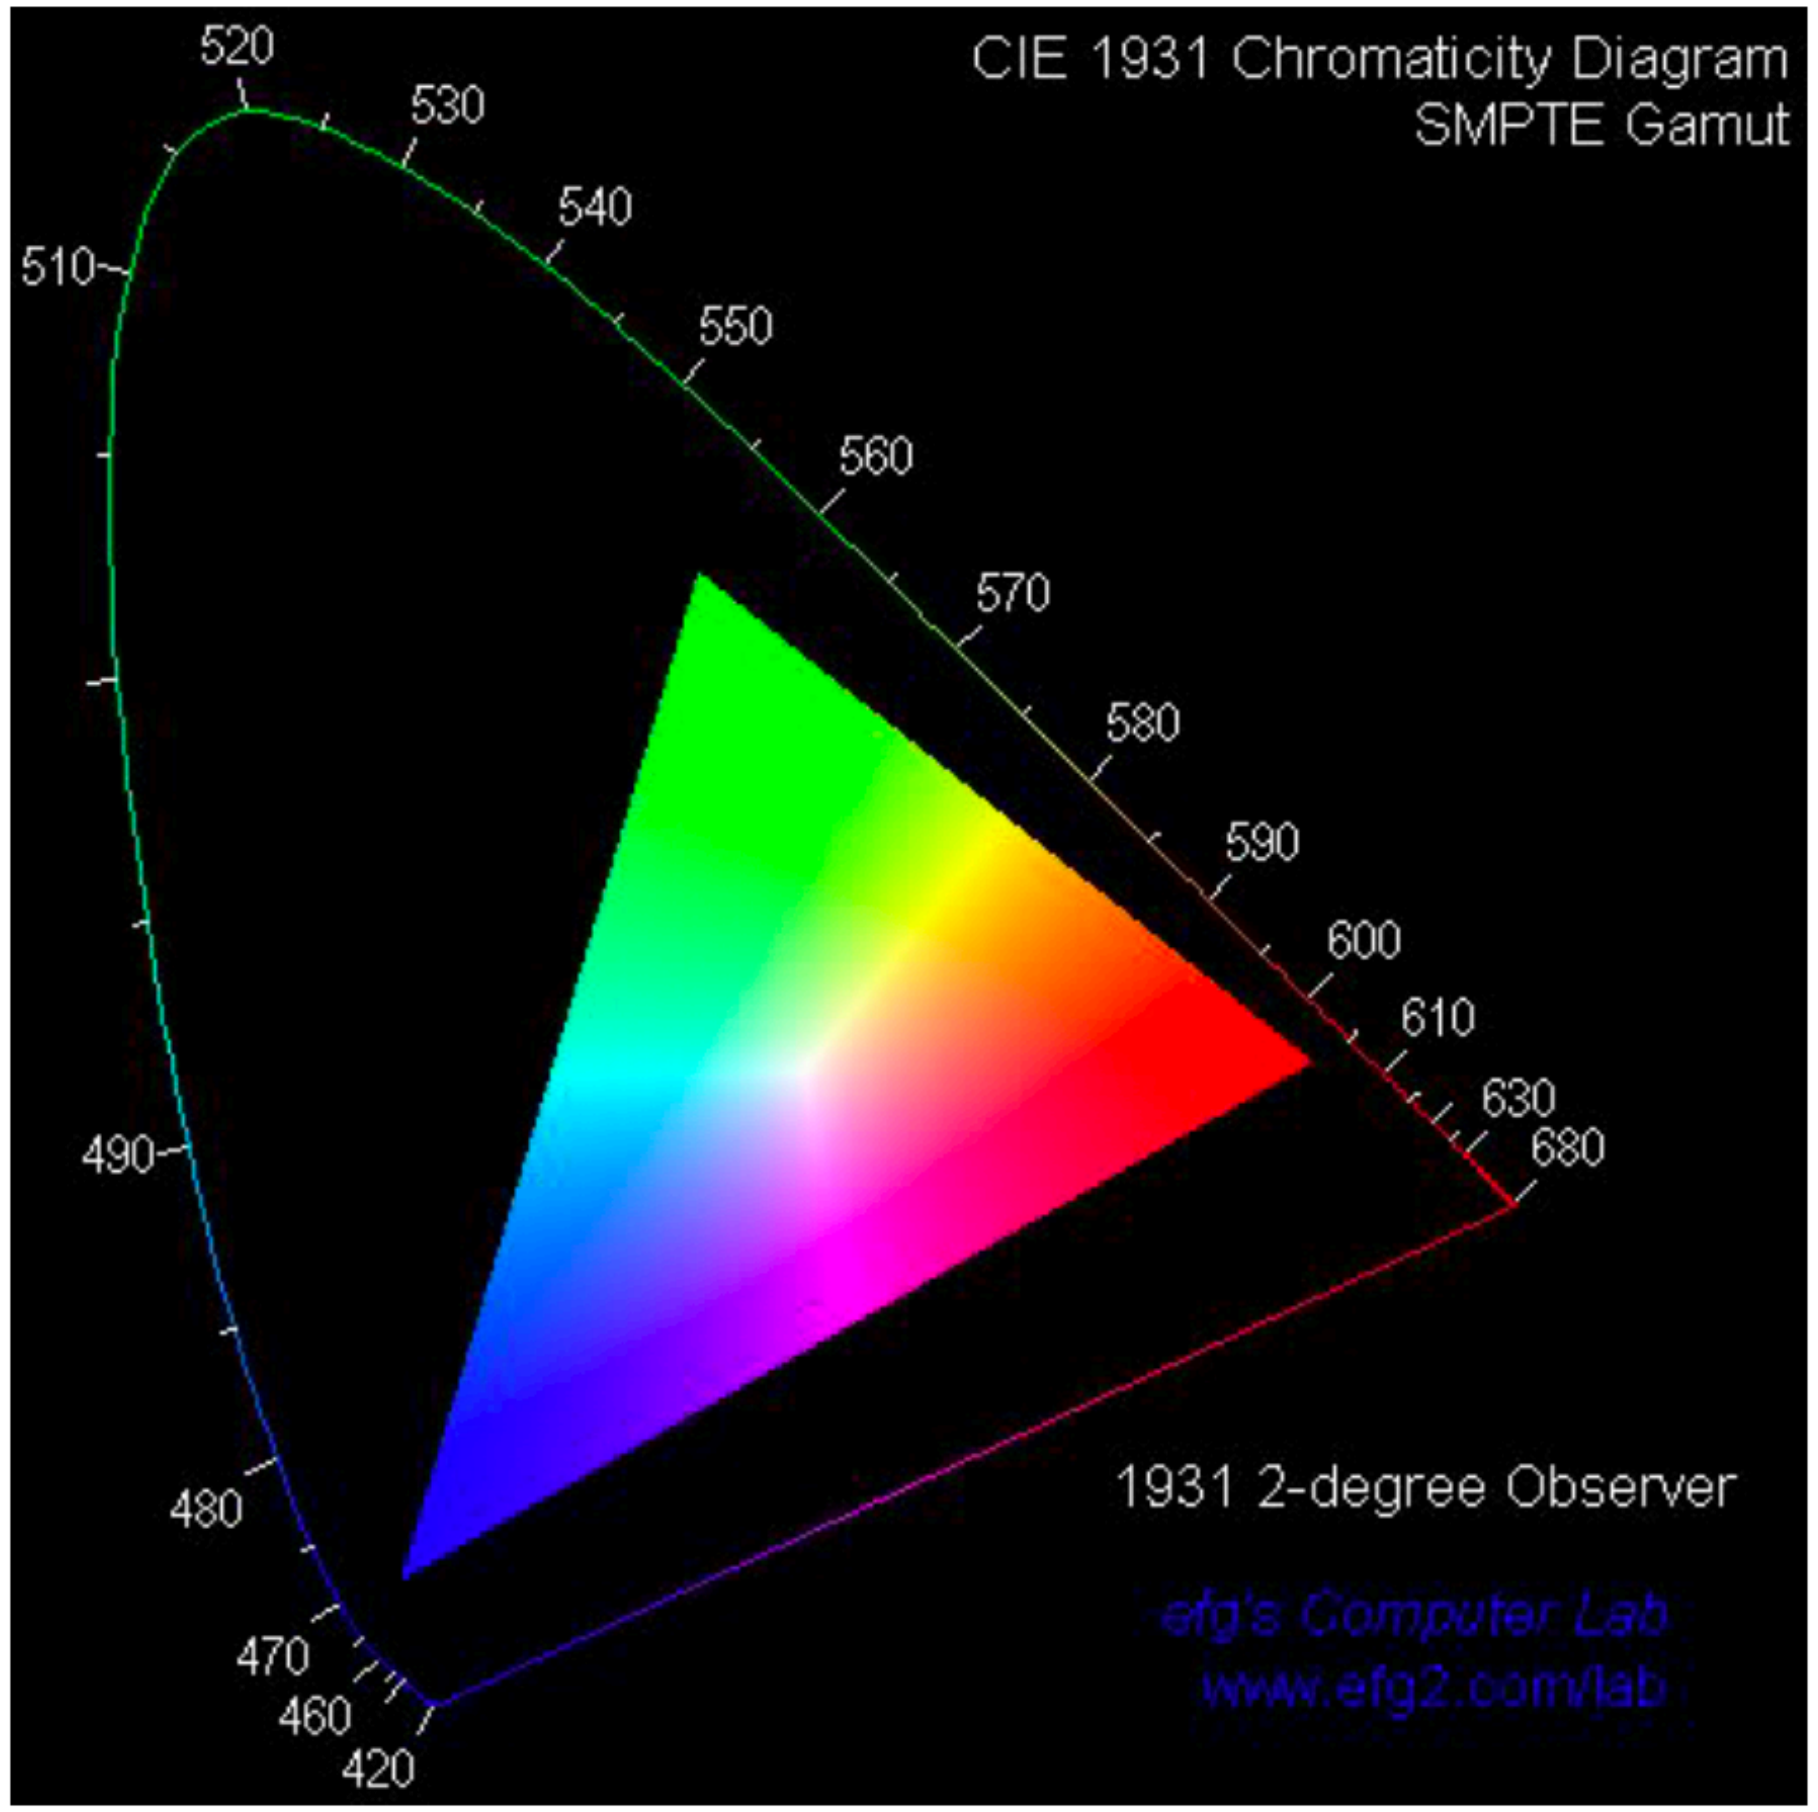

The sub-space of CIE colours that can be displayed on a typical computer monitor (phosphor limitations keep the space quite small)

RGB Colour Space

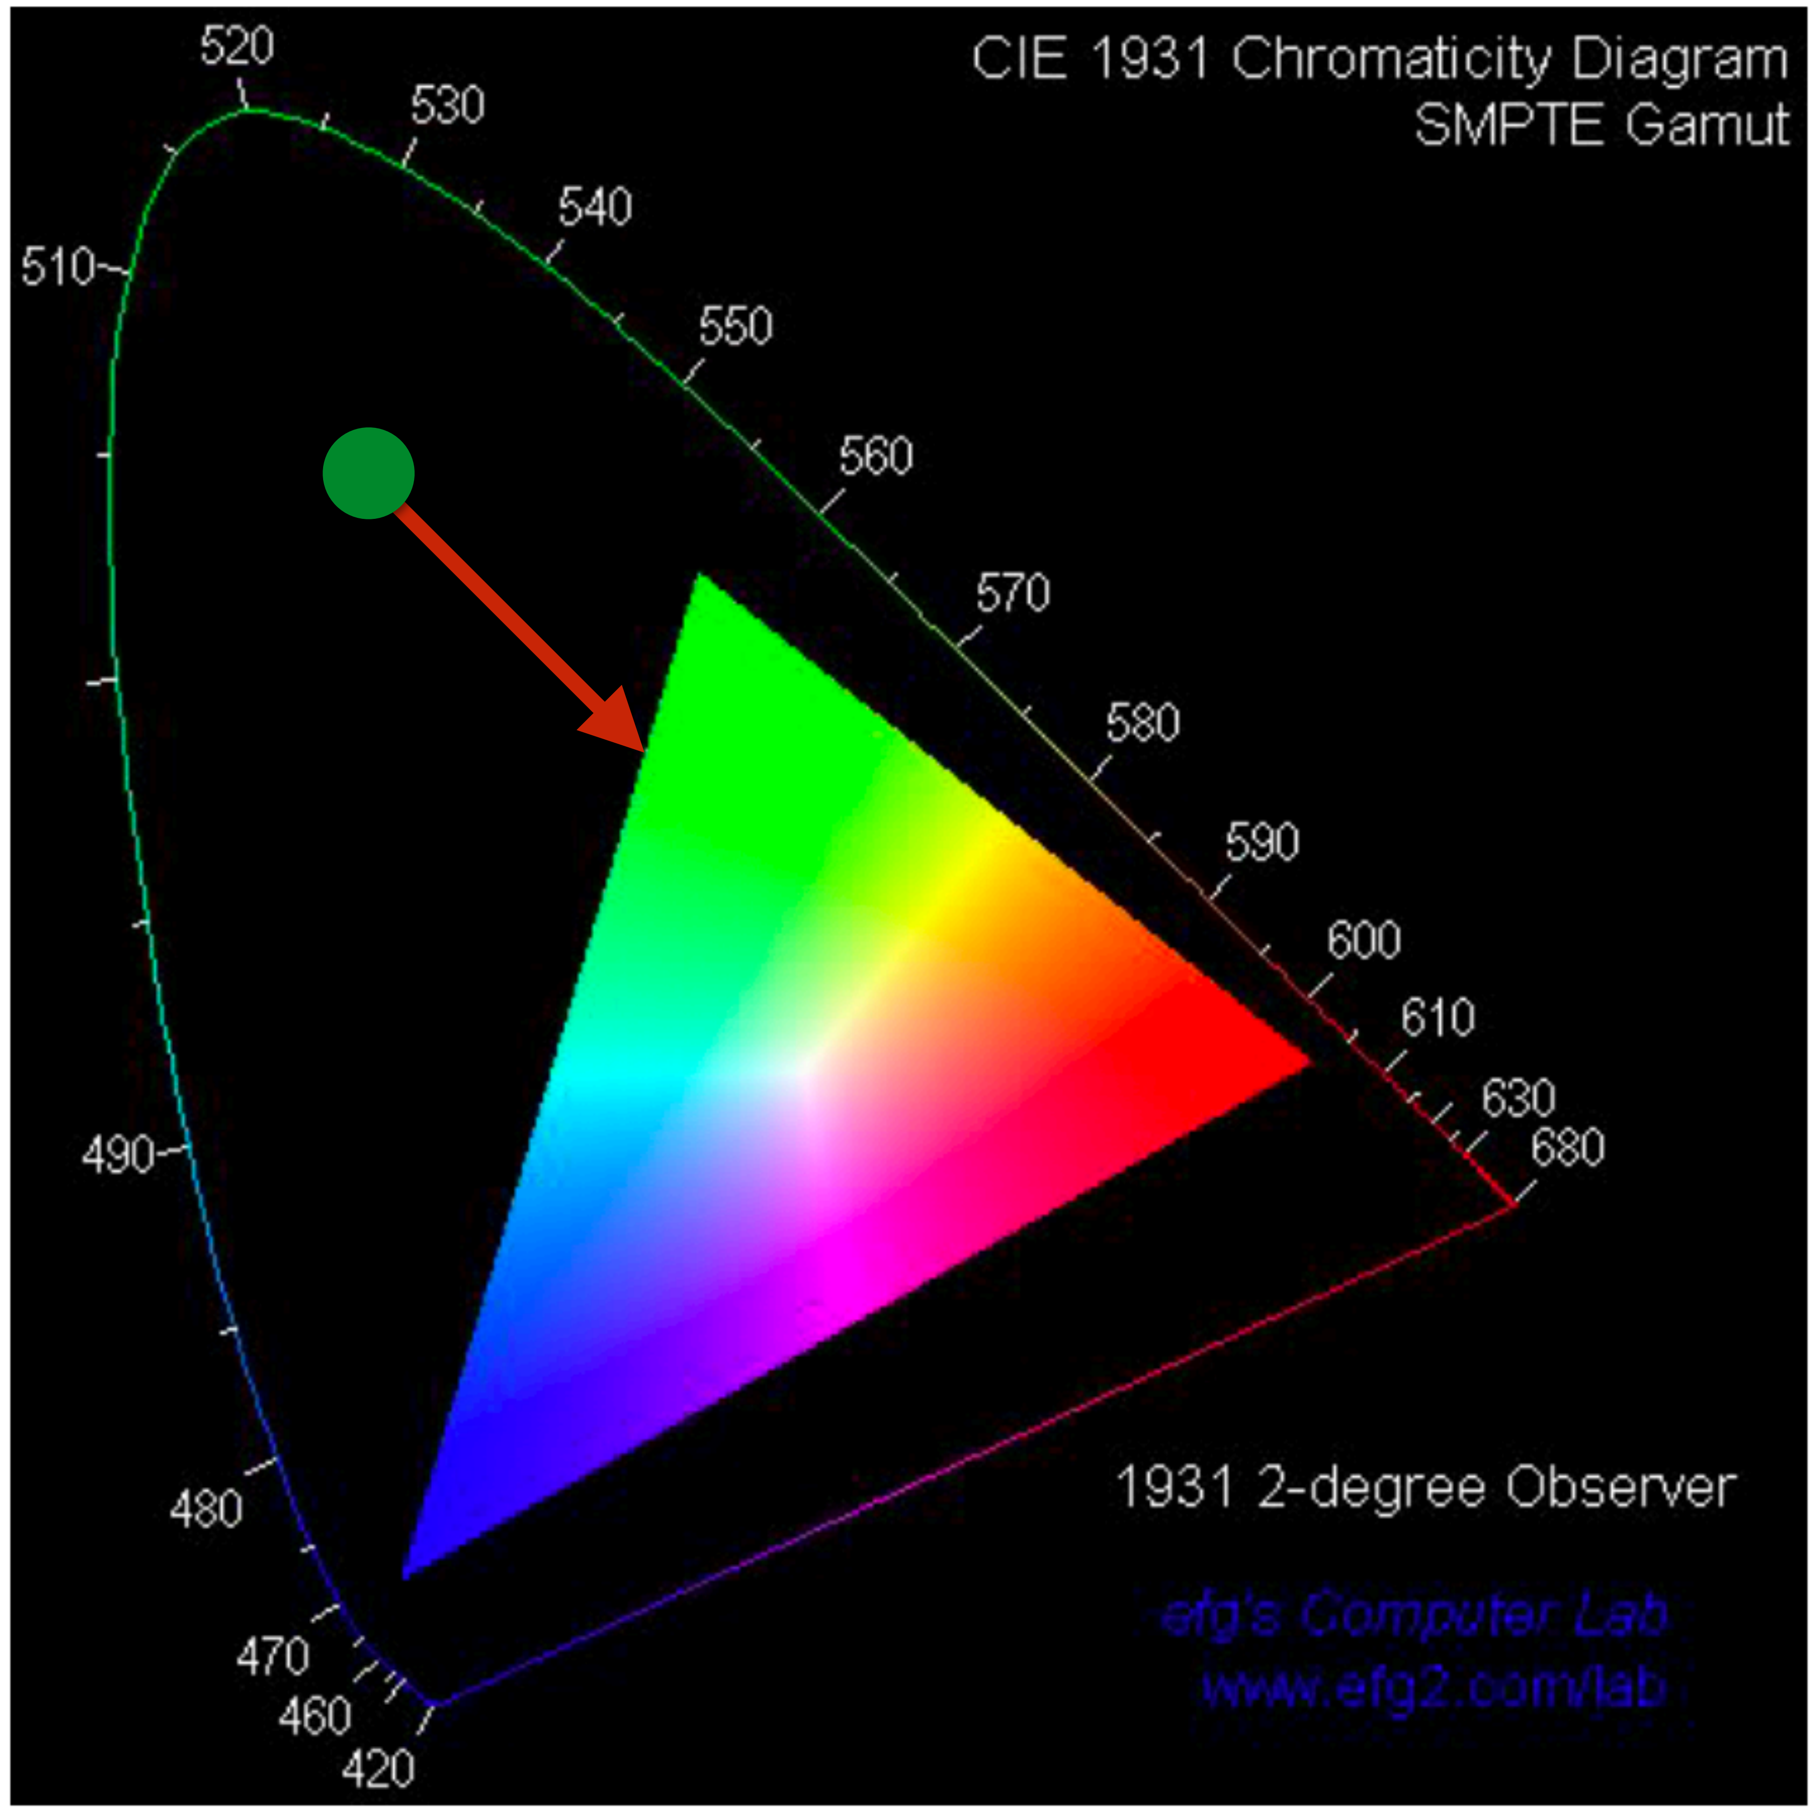

Adding **red** to the green color outside of the region brings it back to where it can be matched by **green** and **blue** RGB primaries

Uniform Colour Spaces

Usually one cannot reproduce colours exactly

This means it is important to know whether a colour difference would be noticeable to a human viewer

Uniform Colour Spaces

McAdam Ellipses: Each ellipse shows colours perceived to be the same

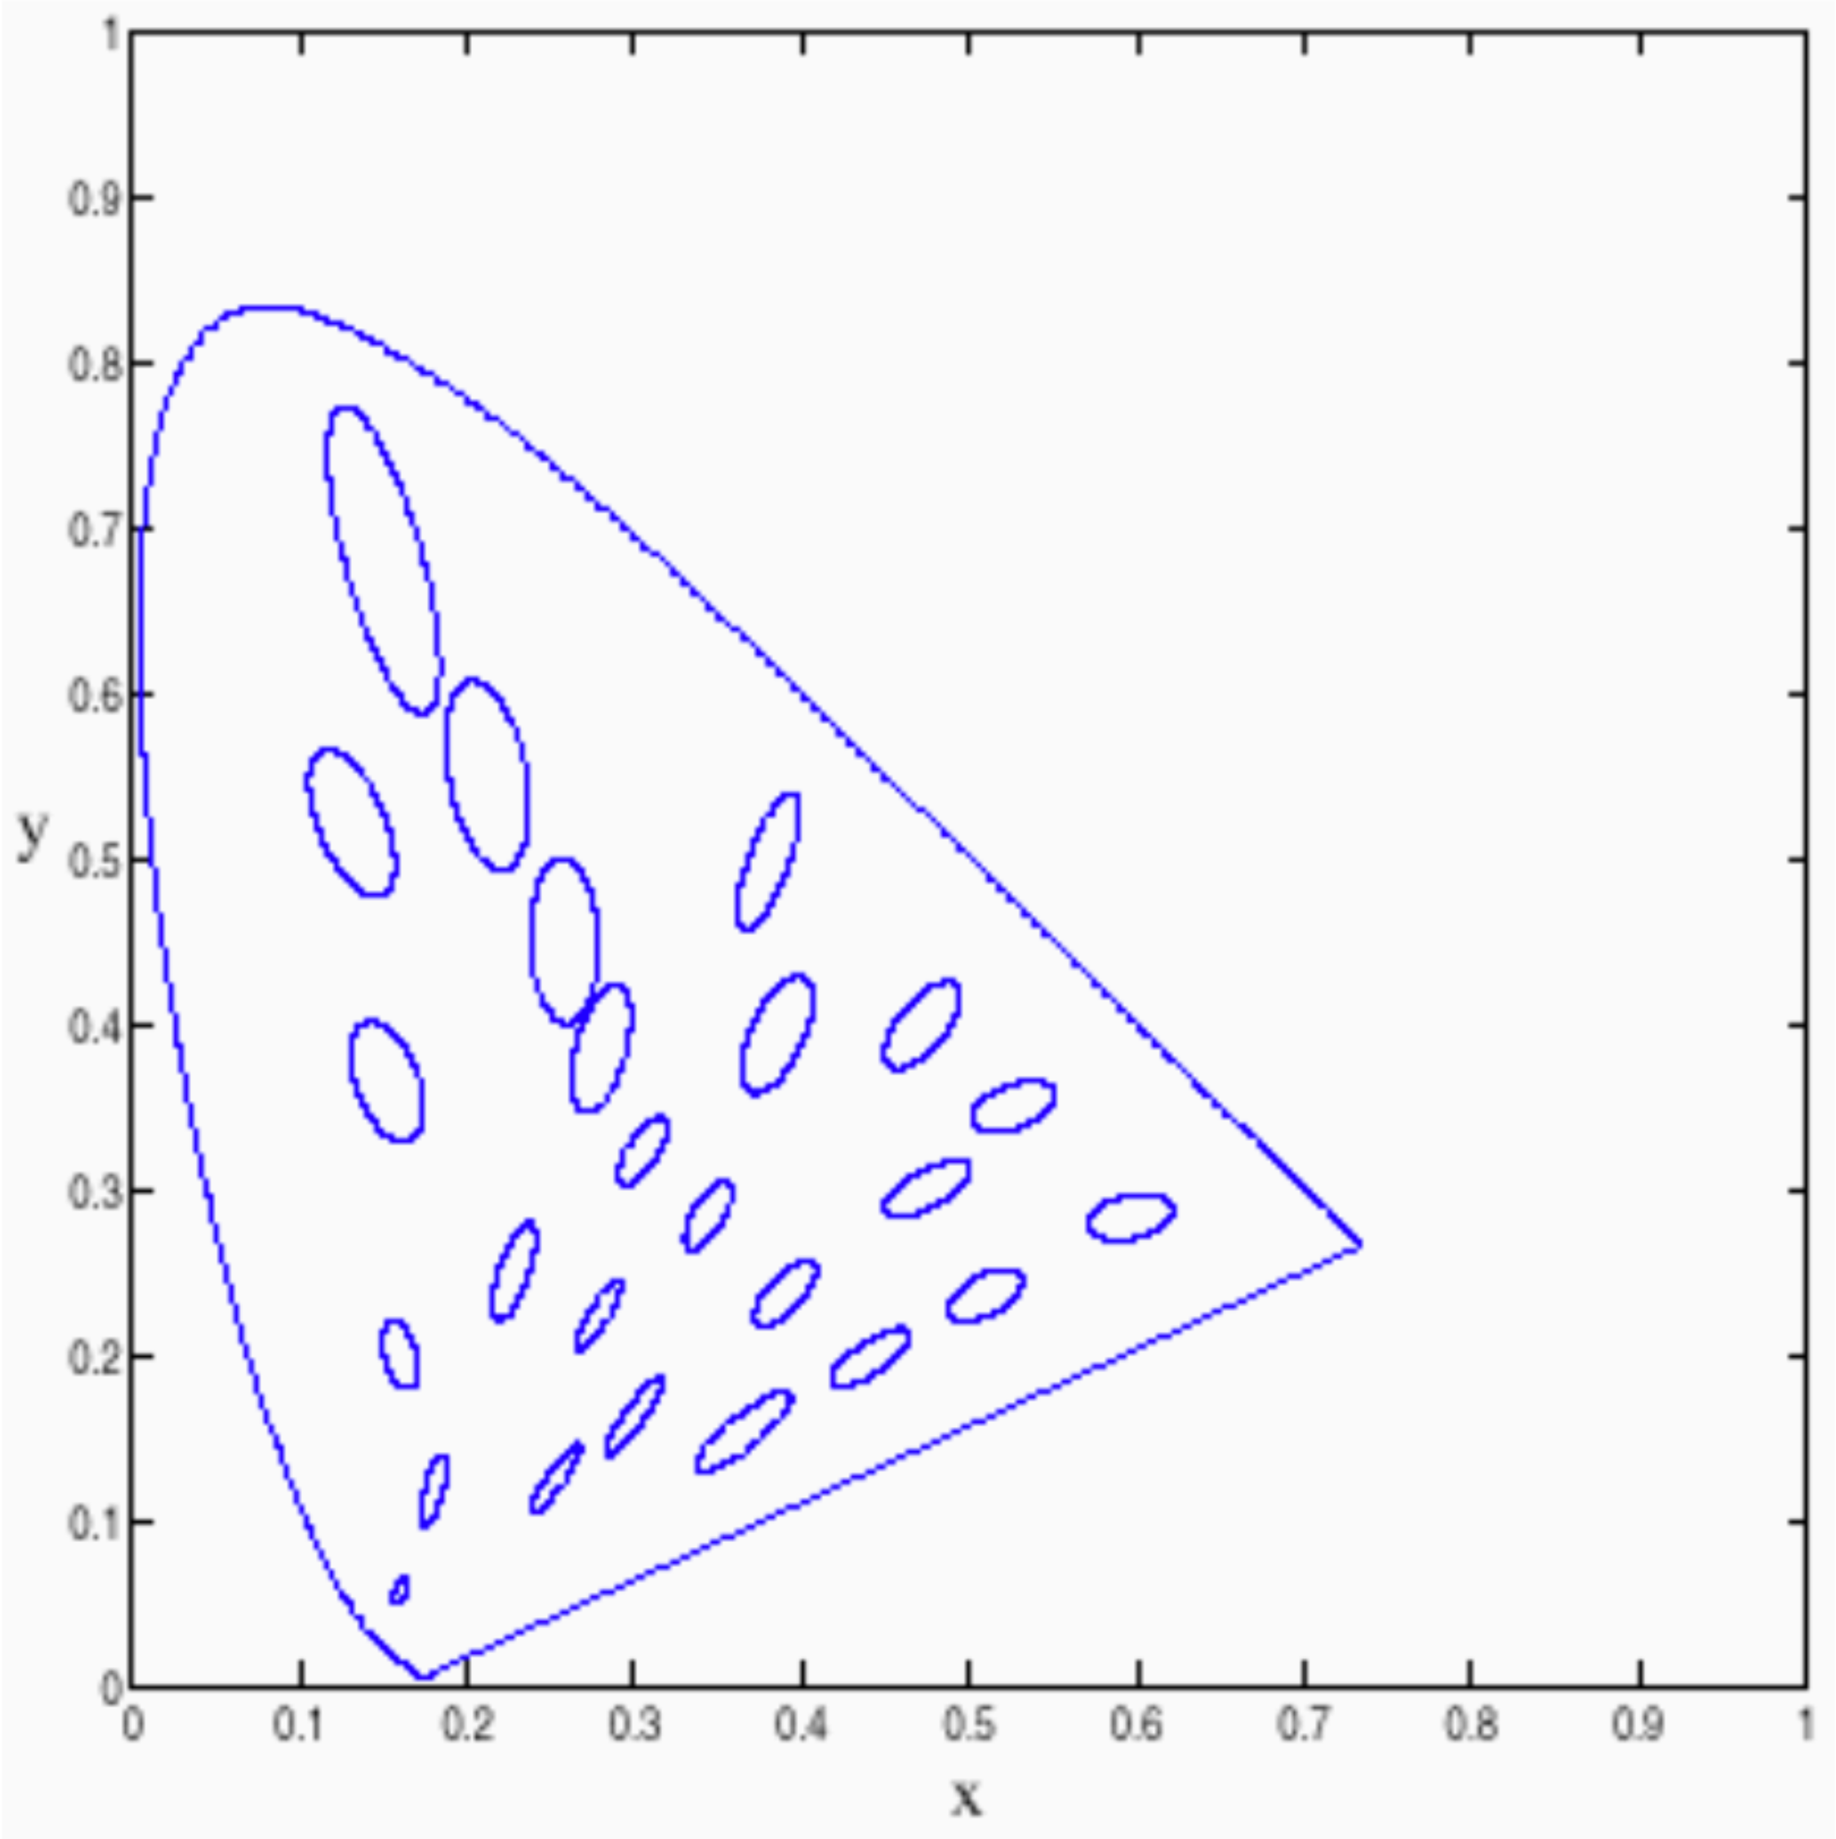

10 times actual size

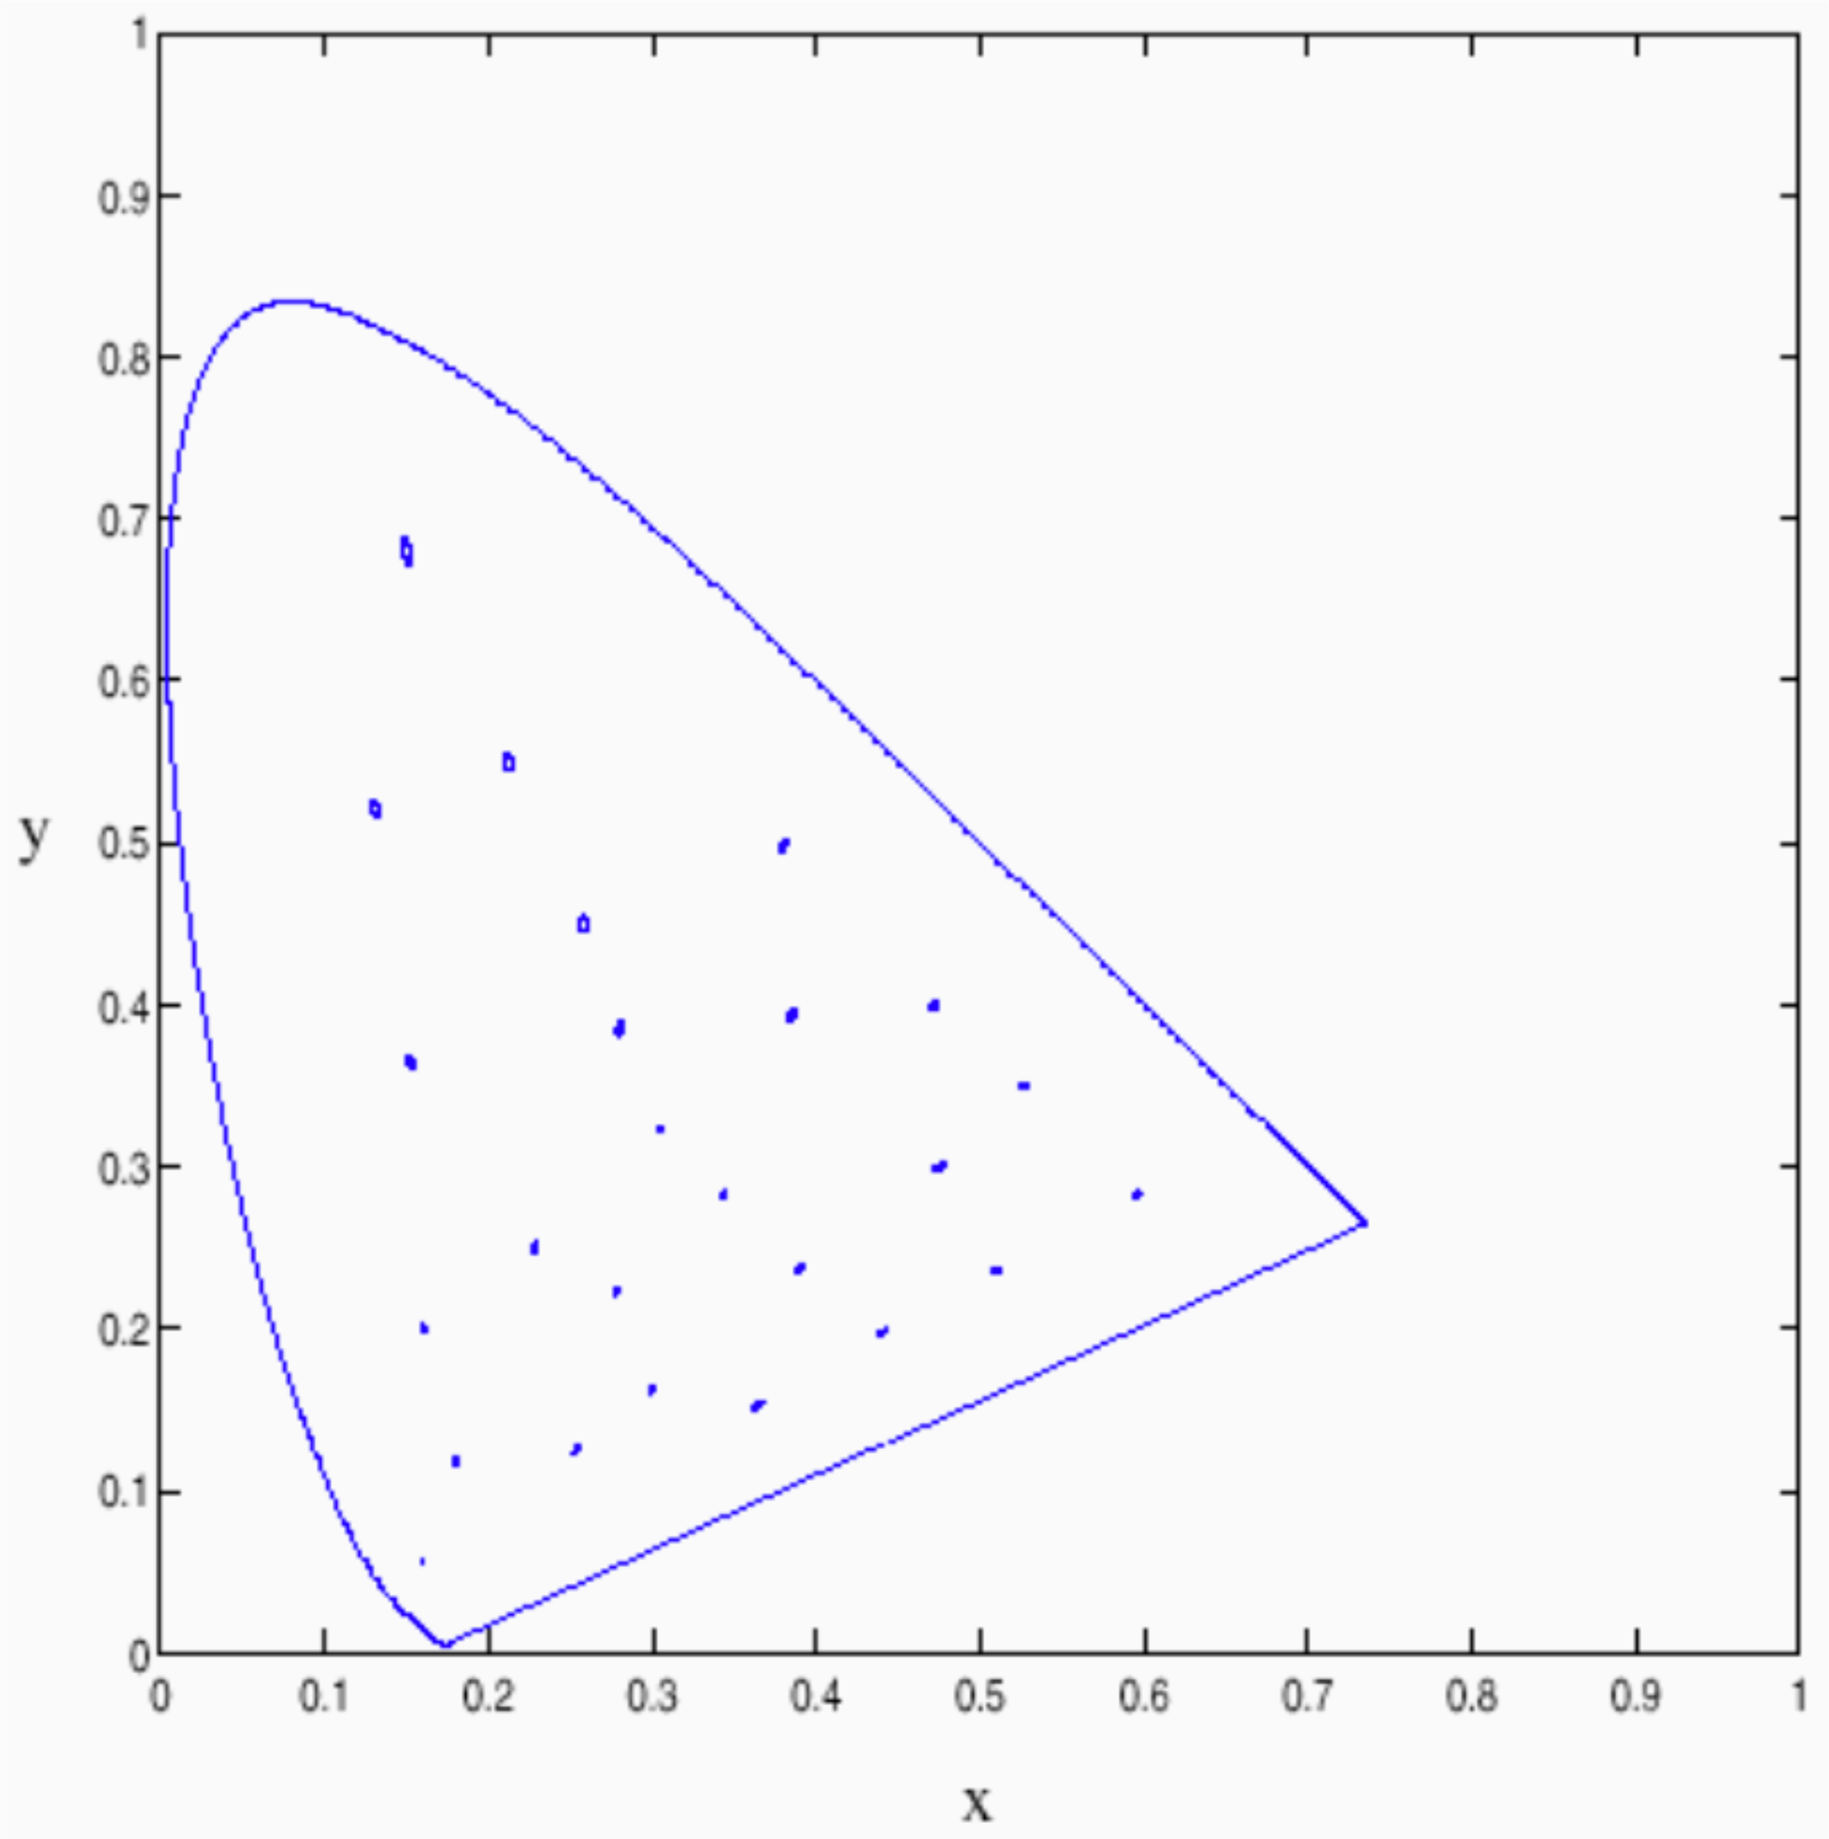

Actual Size

Forsyth & Ponce (2nd ed.) Figure 3.14

Uniform Colour Space **CIE-LAB**

McAdam Ellipses: Each ellipse shows colours perceived to be the same

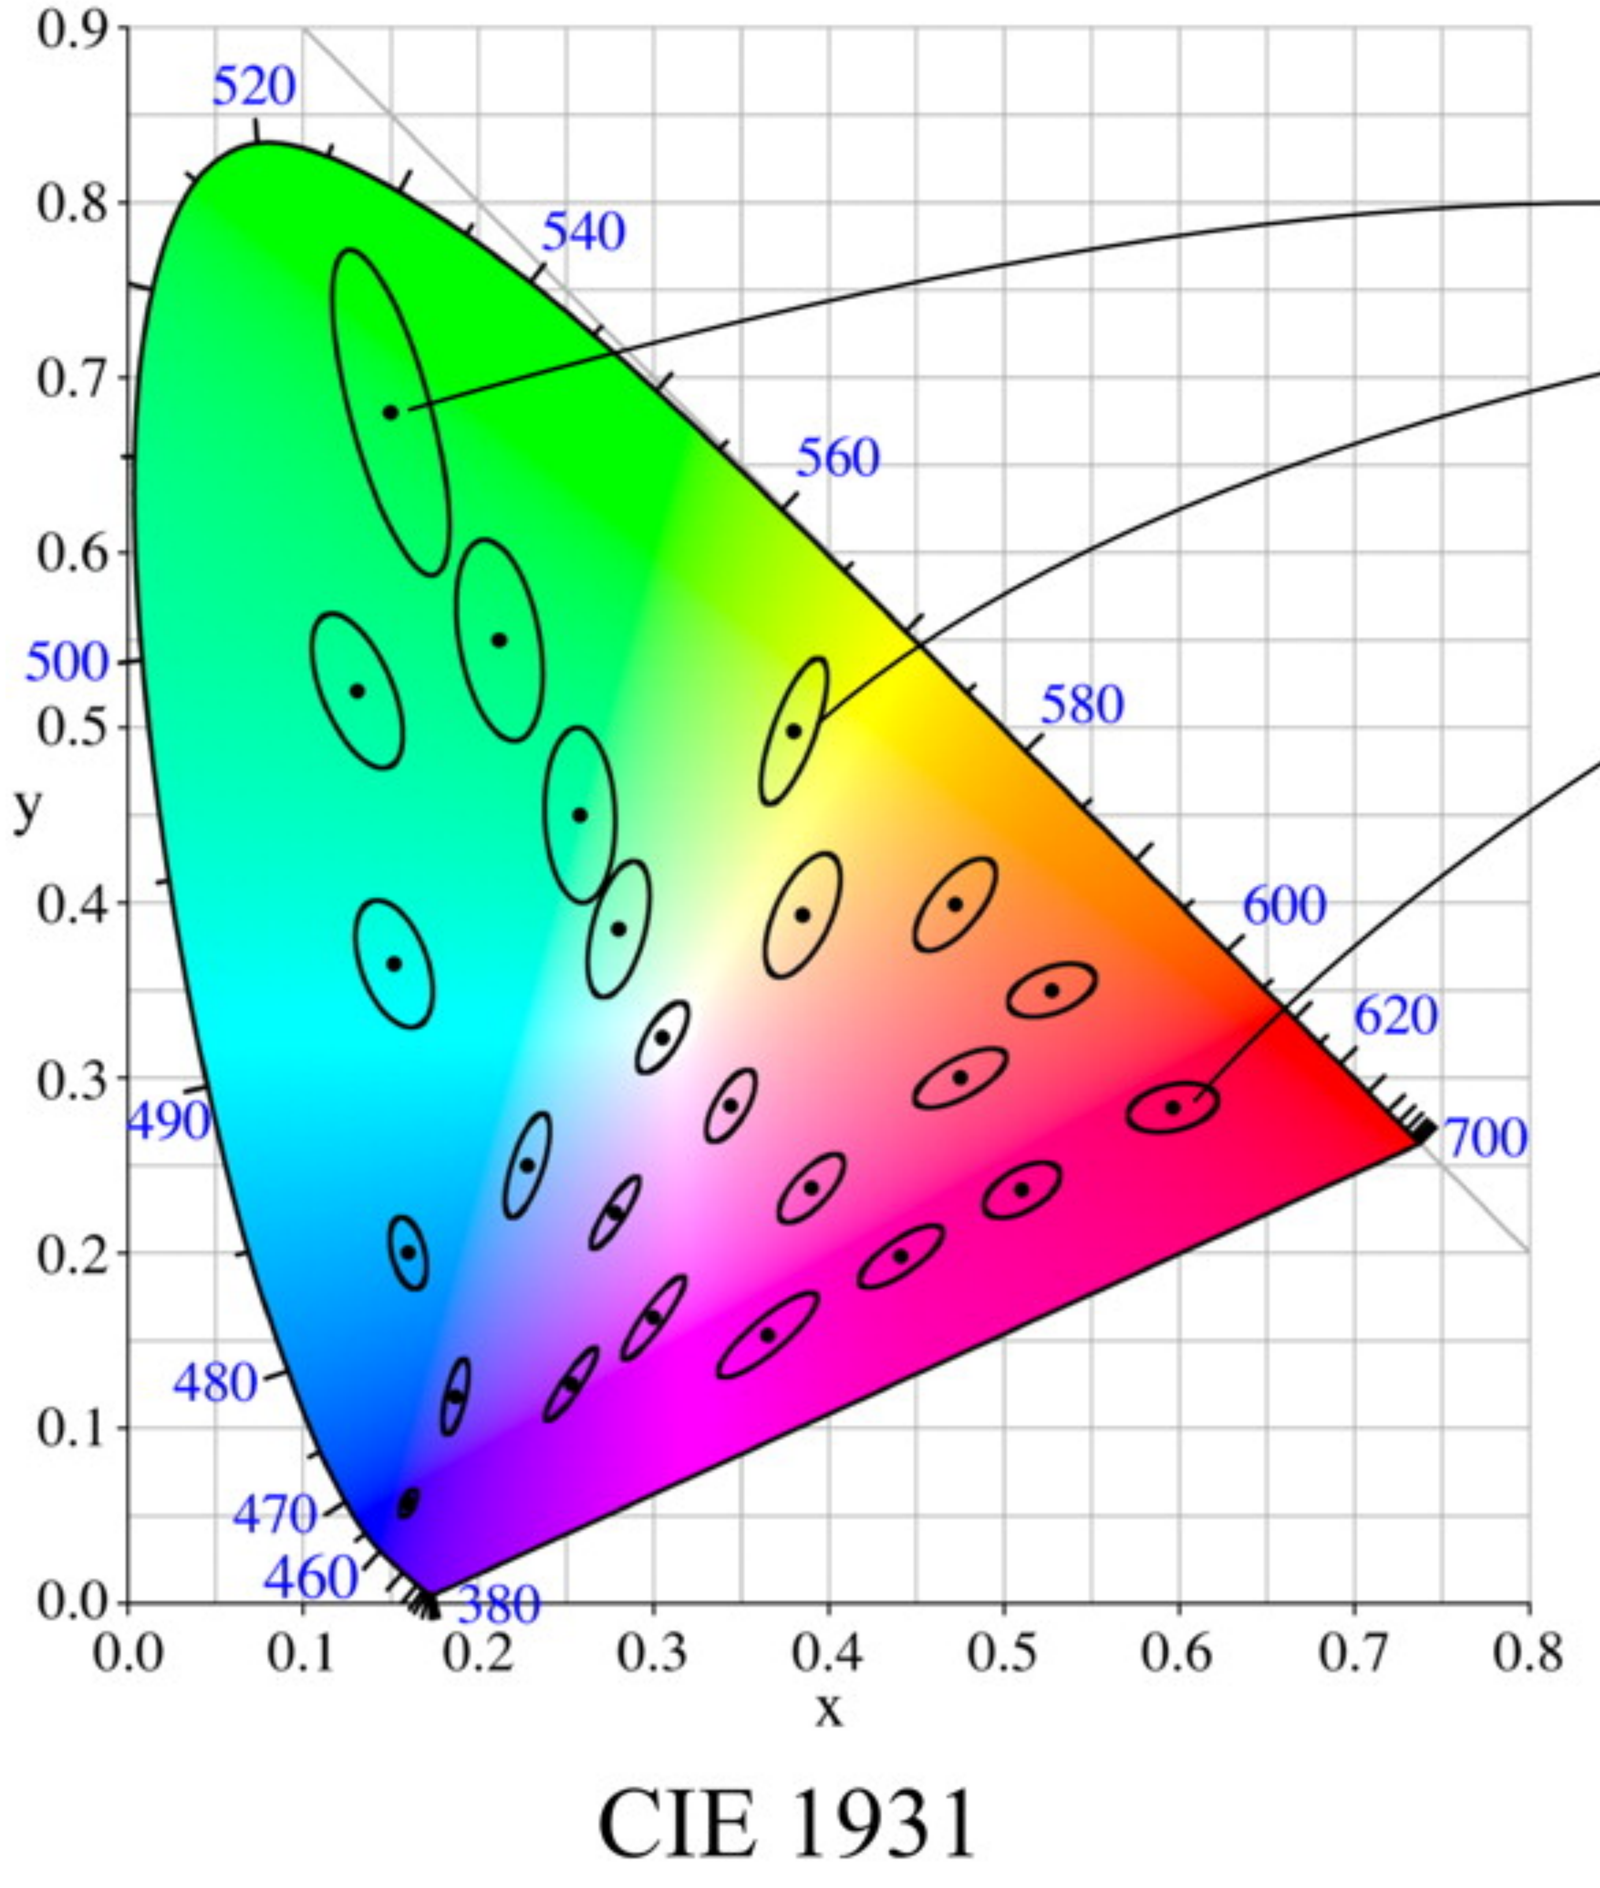

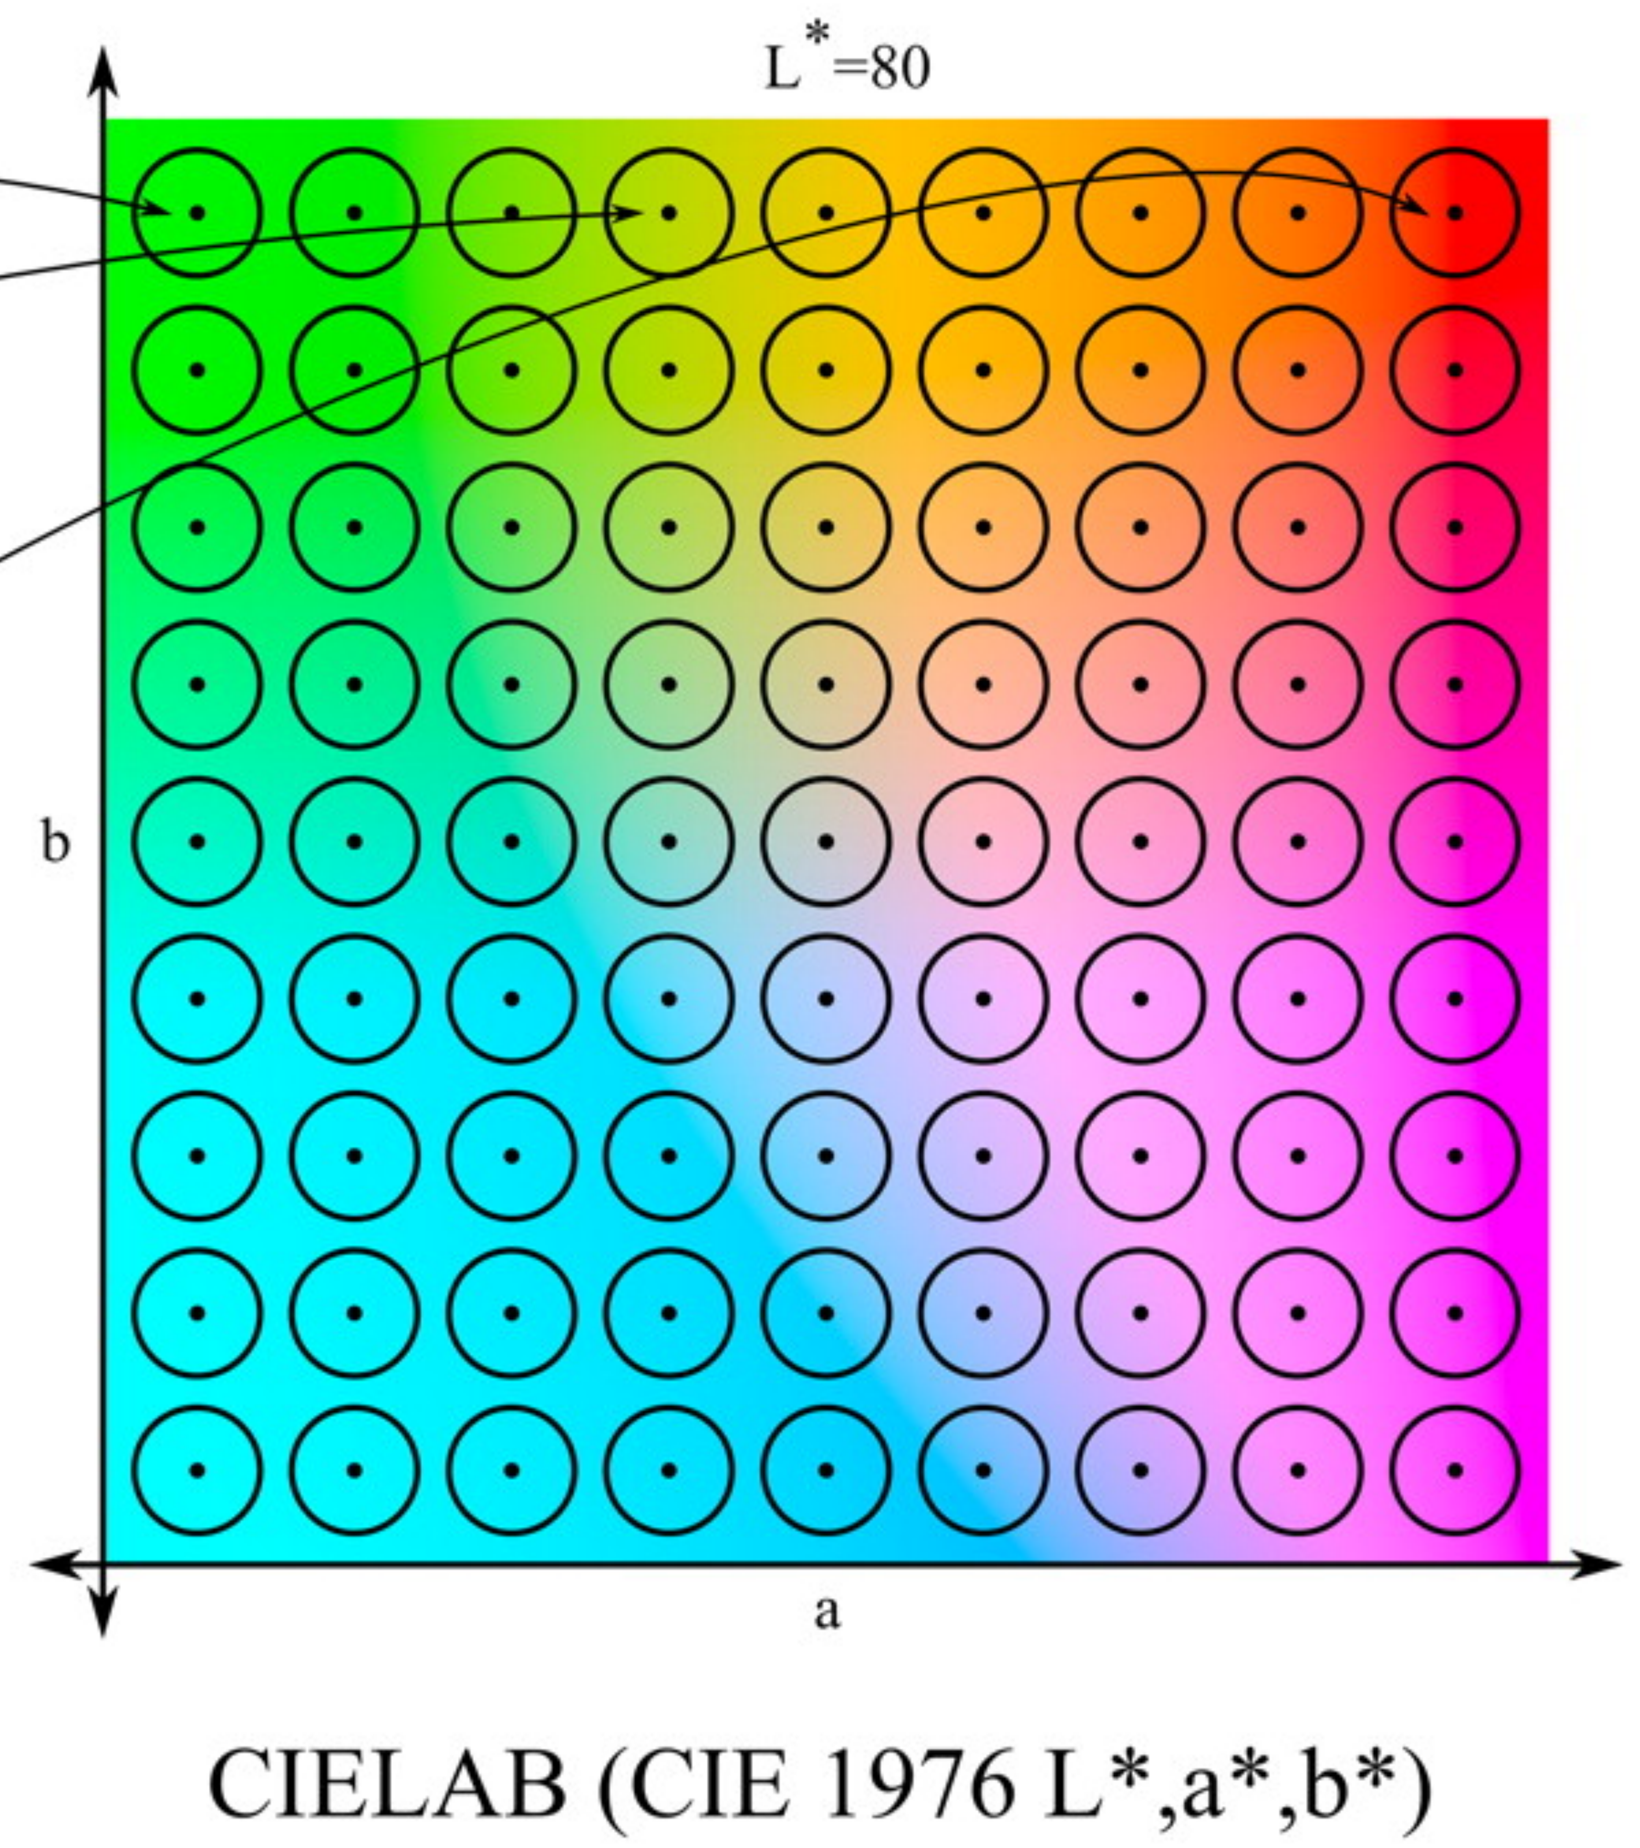

From Henrich et al. 2011

https://iovs.arvojournals.org/article.aspx?articleid=2187751

Uniform Colour Spaces

McAdam ellipses demonstrate that differences in x , y are a poor guide to differences in perceived colour

A **uniform colour space** is one in which differences in coordinates are a good guide to differences in perceived colour

— example: CIE LAB

Why **should** you care about all this?

A state of the art super-res network trained with L2 loss is good at sharpening edges, but results lack realistic texture

[Ledig et al 2017]

Why **should** you care about all this?

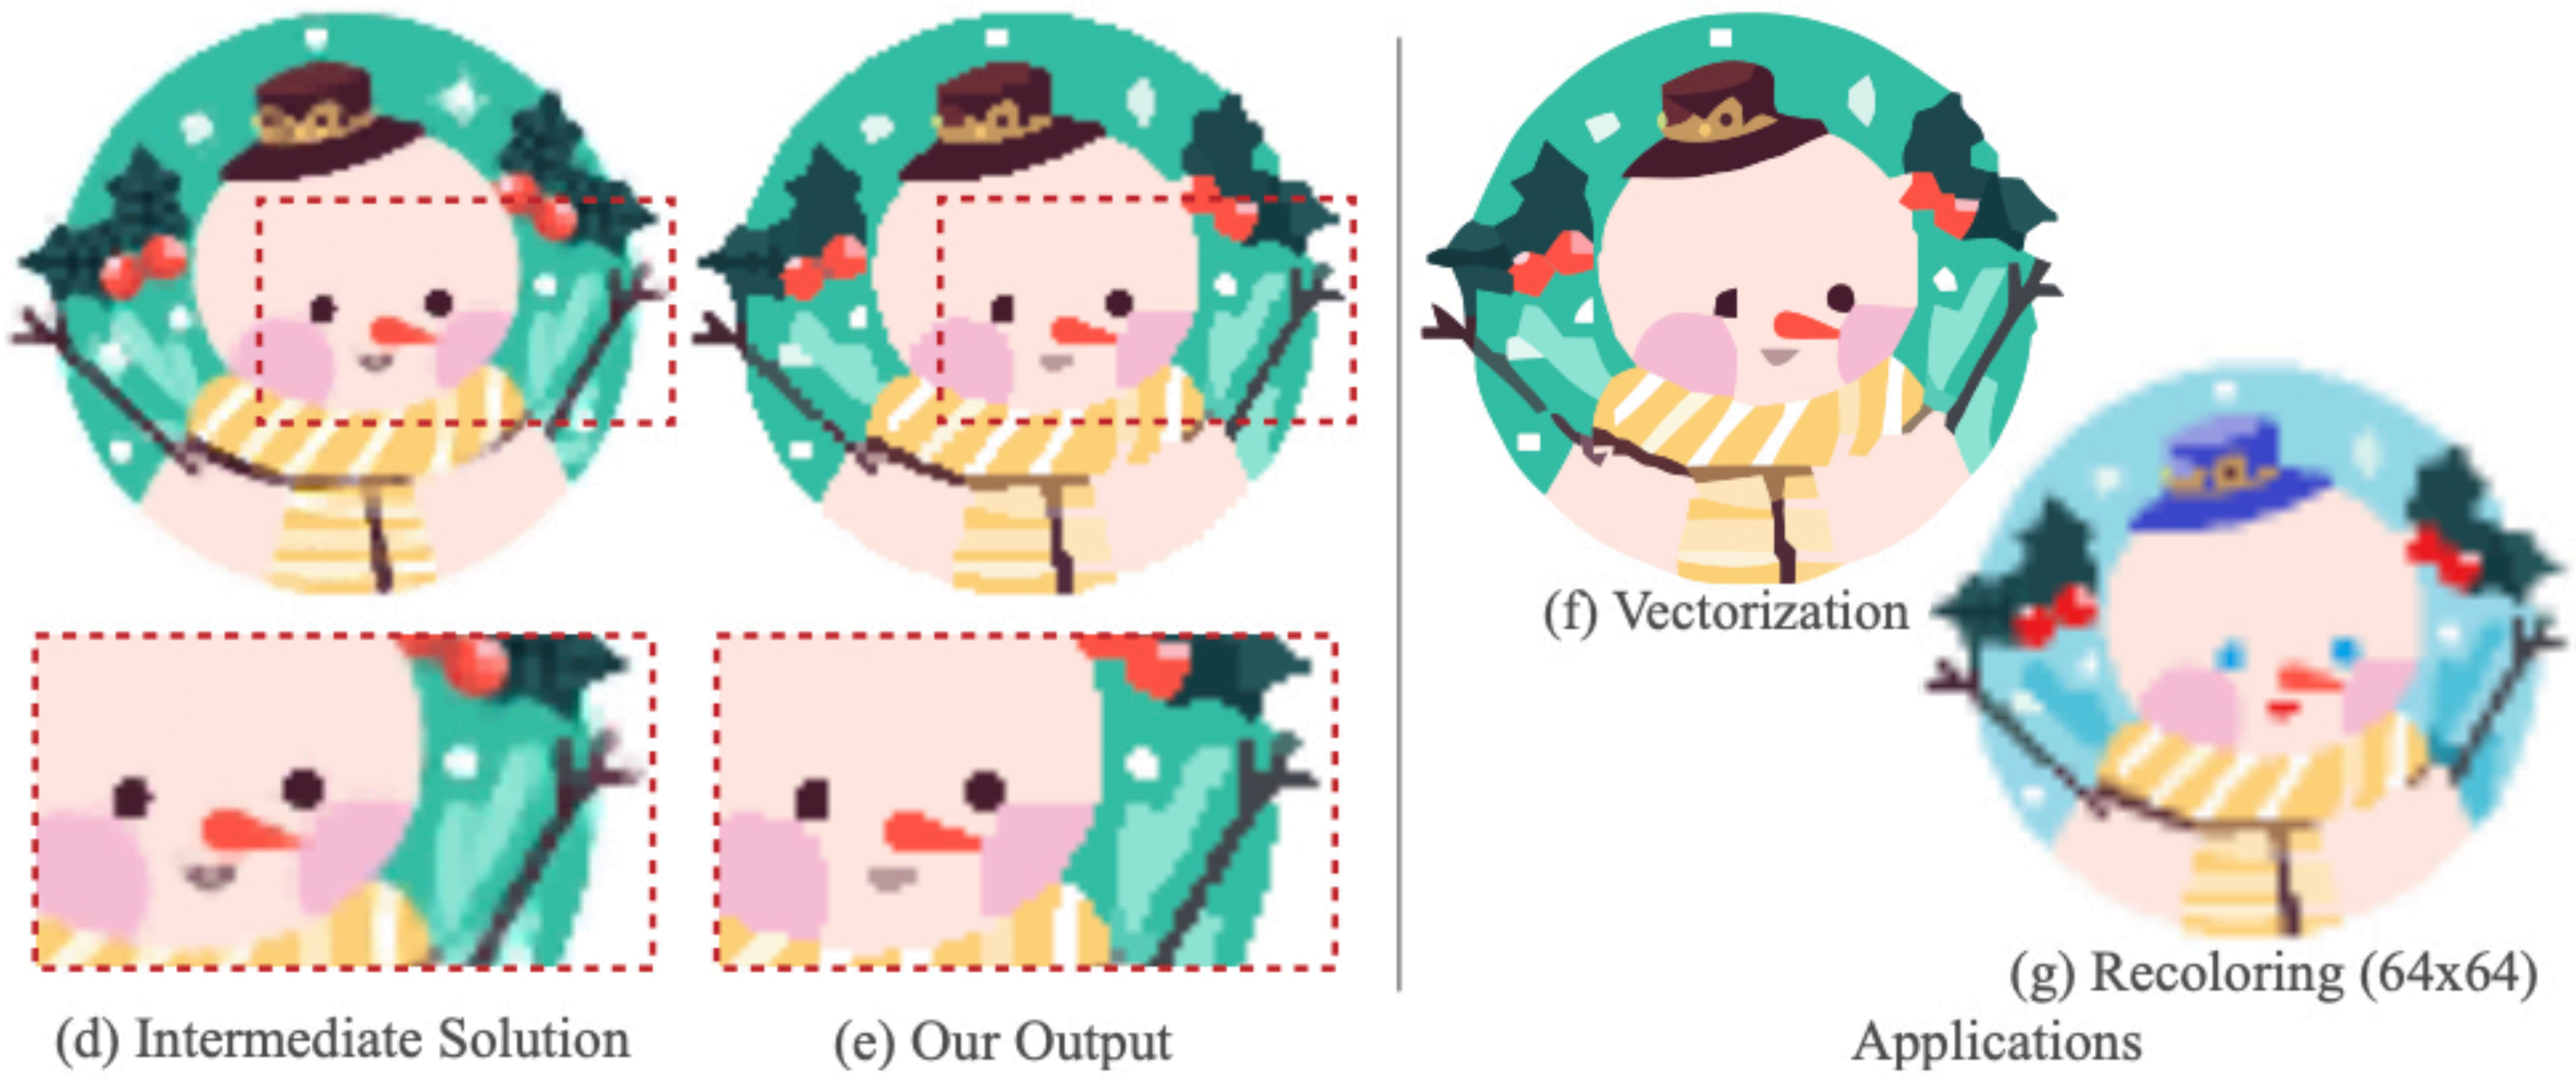

[Yang et al., Eurographics 2023]

YCbCr Color Space

- Separates luminance (**Y**) from chrominance (**Cr**, **Cb**)
- Chrominance can be compressed more (e.g. 1/2 size in JPG)

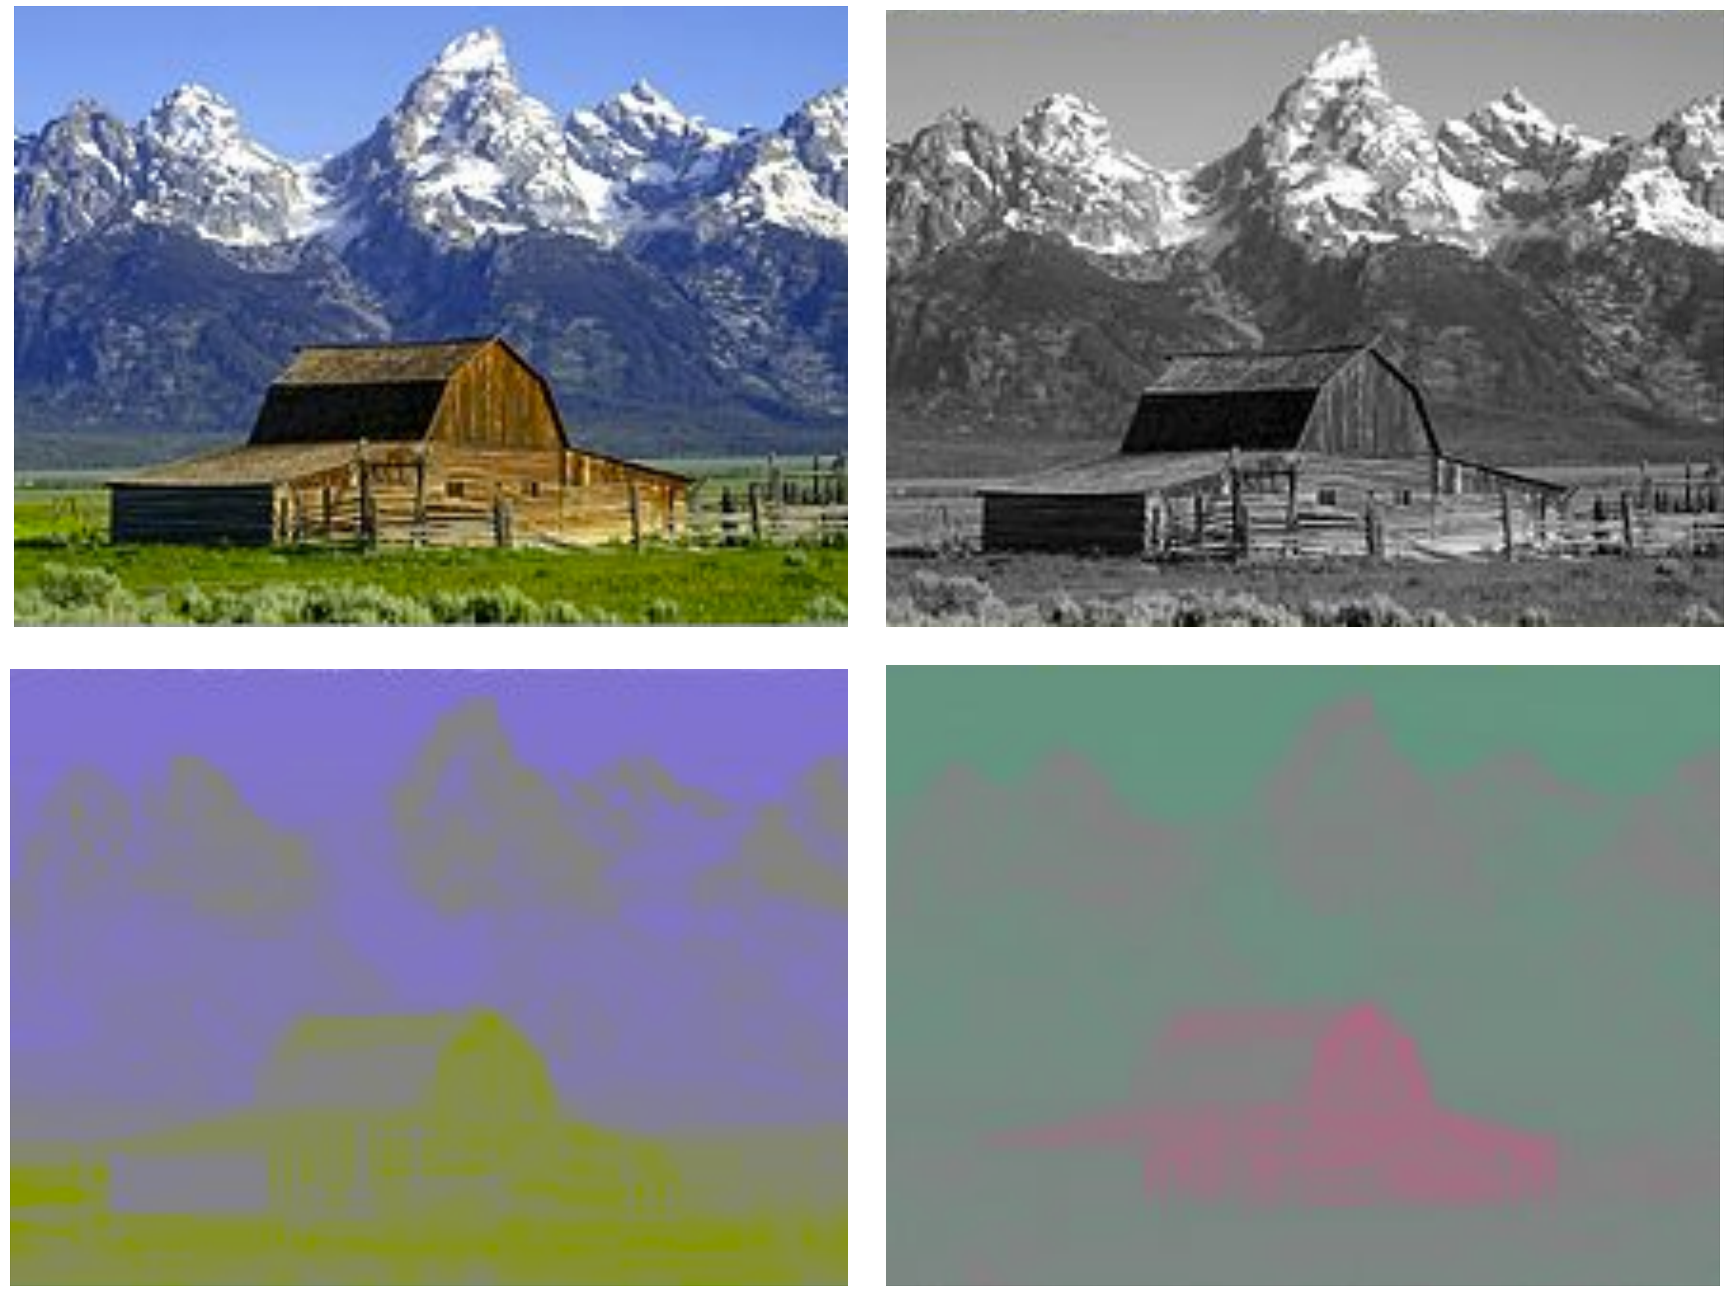

$$Y' = 16 + 65.5R' + 128.6G' + 25.0B'$$
$$Cb = 128 - 37.8R' - 74.2G' + 112B'$$
$$Cr = 128 + 112.0R' - 93.8G' - 18.2B'$$

Linear transform of RGB

YCrCb is used for image and video coding. Human vision uses a similar transform (opponent colours) and we have more rods than cones

RGB Color Space

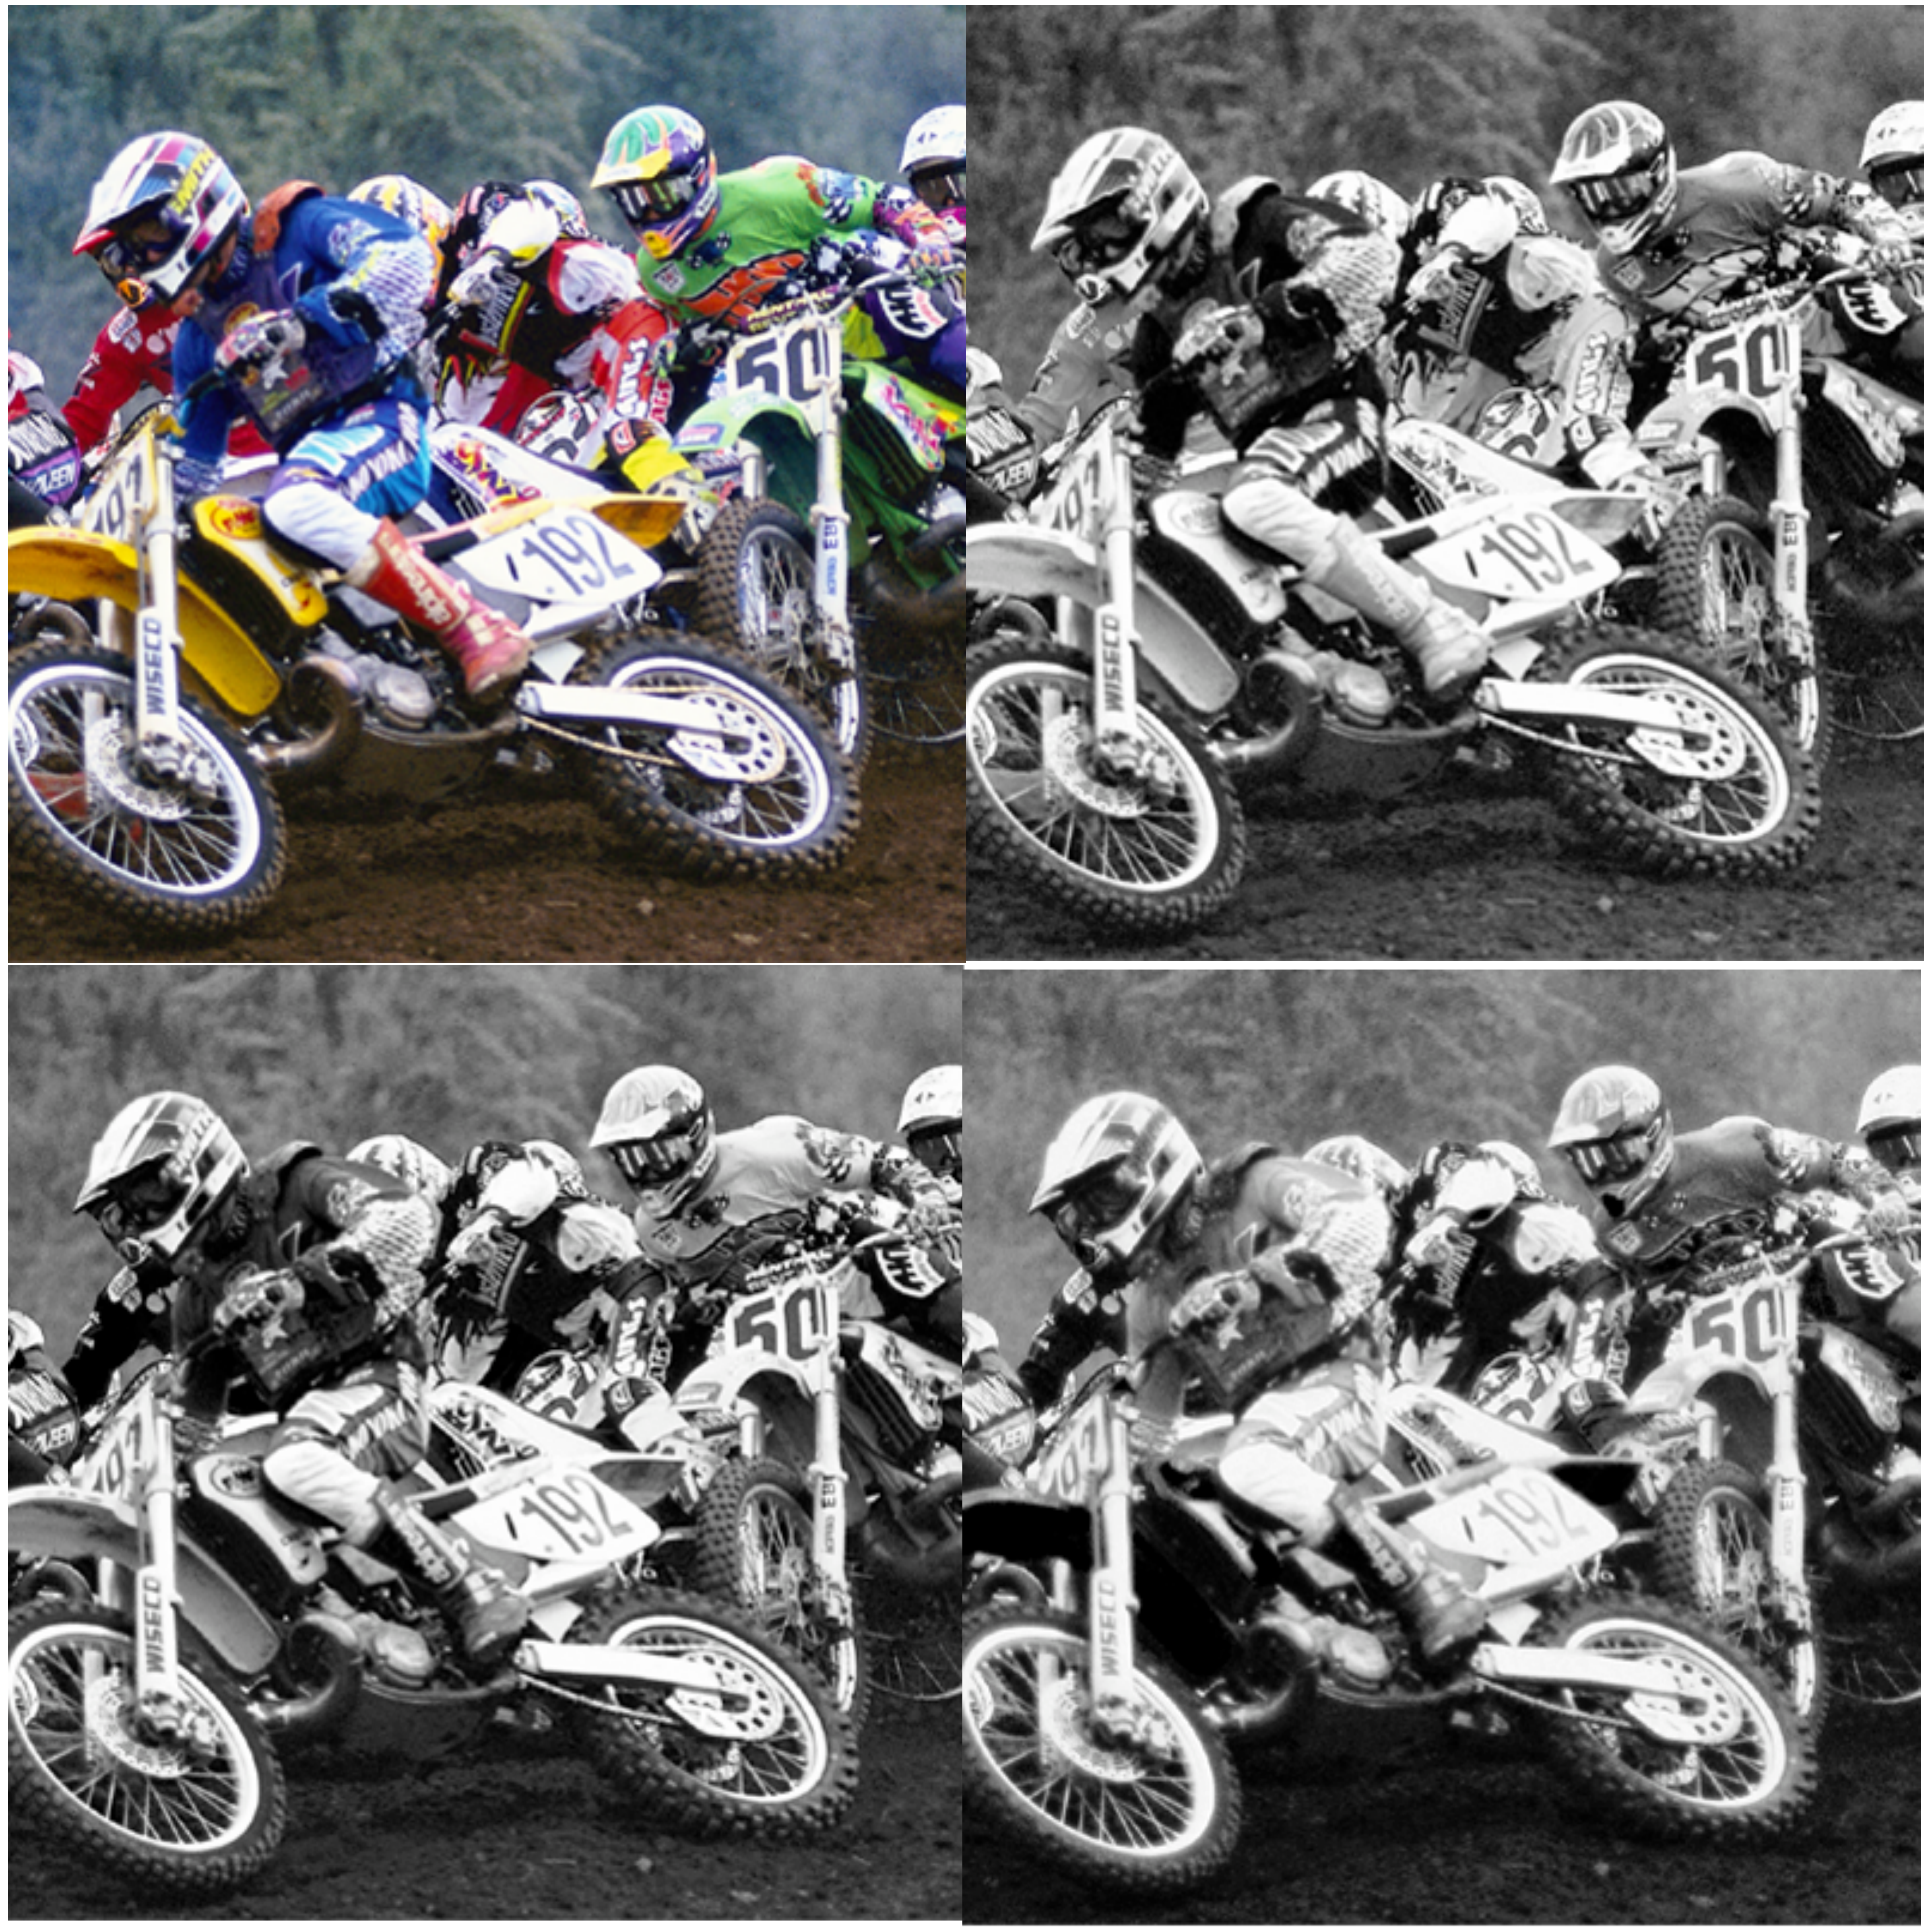

Red

Green

Blue

YCbCr Color Space

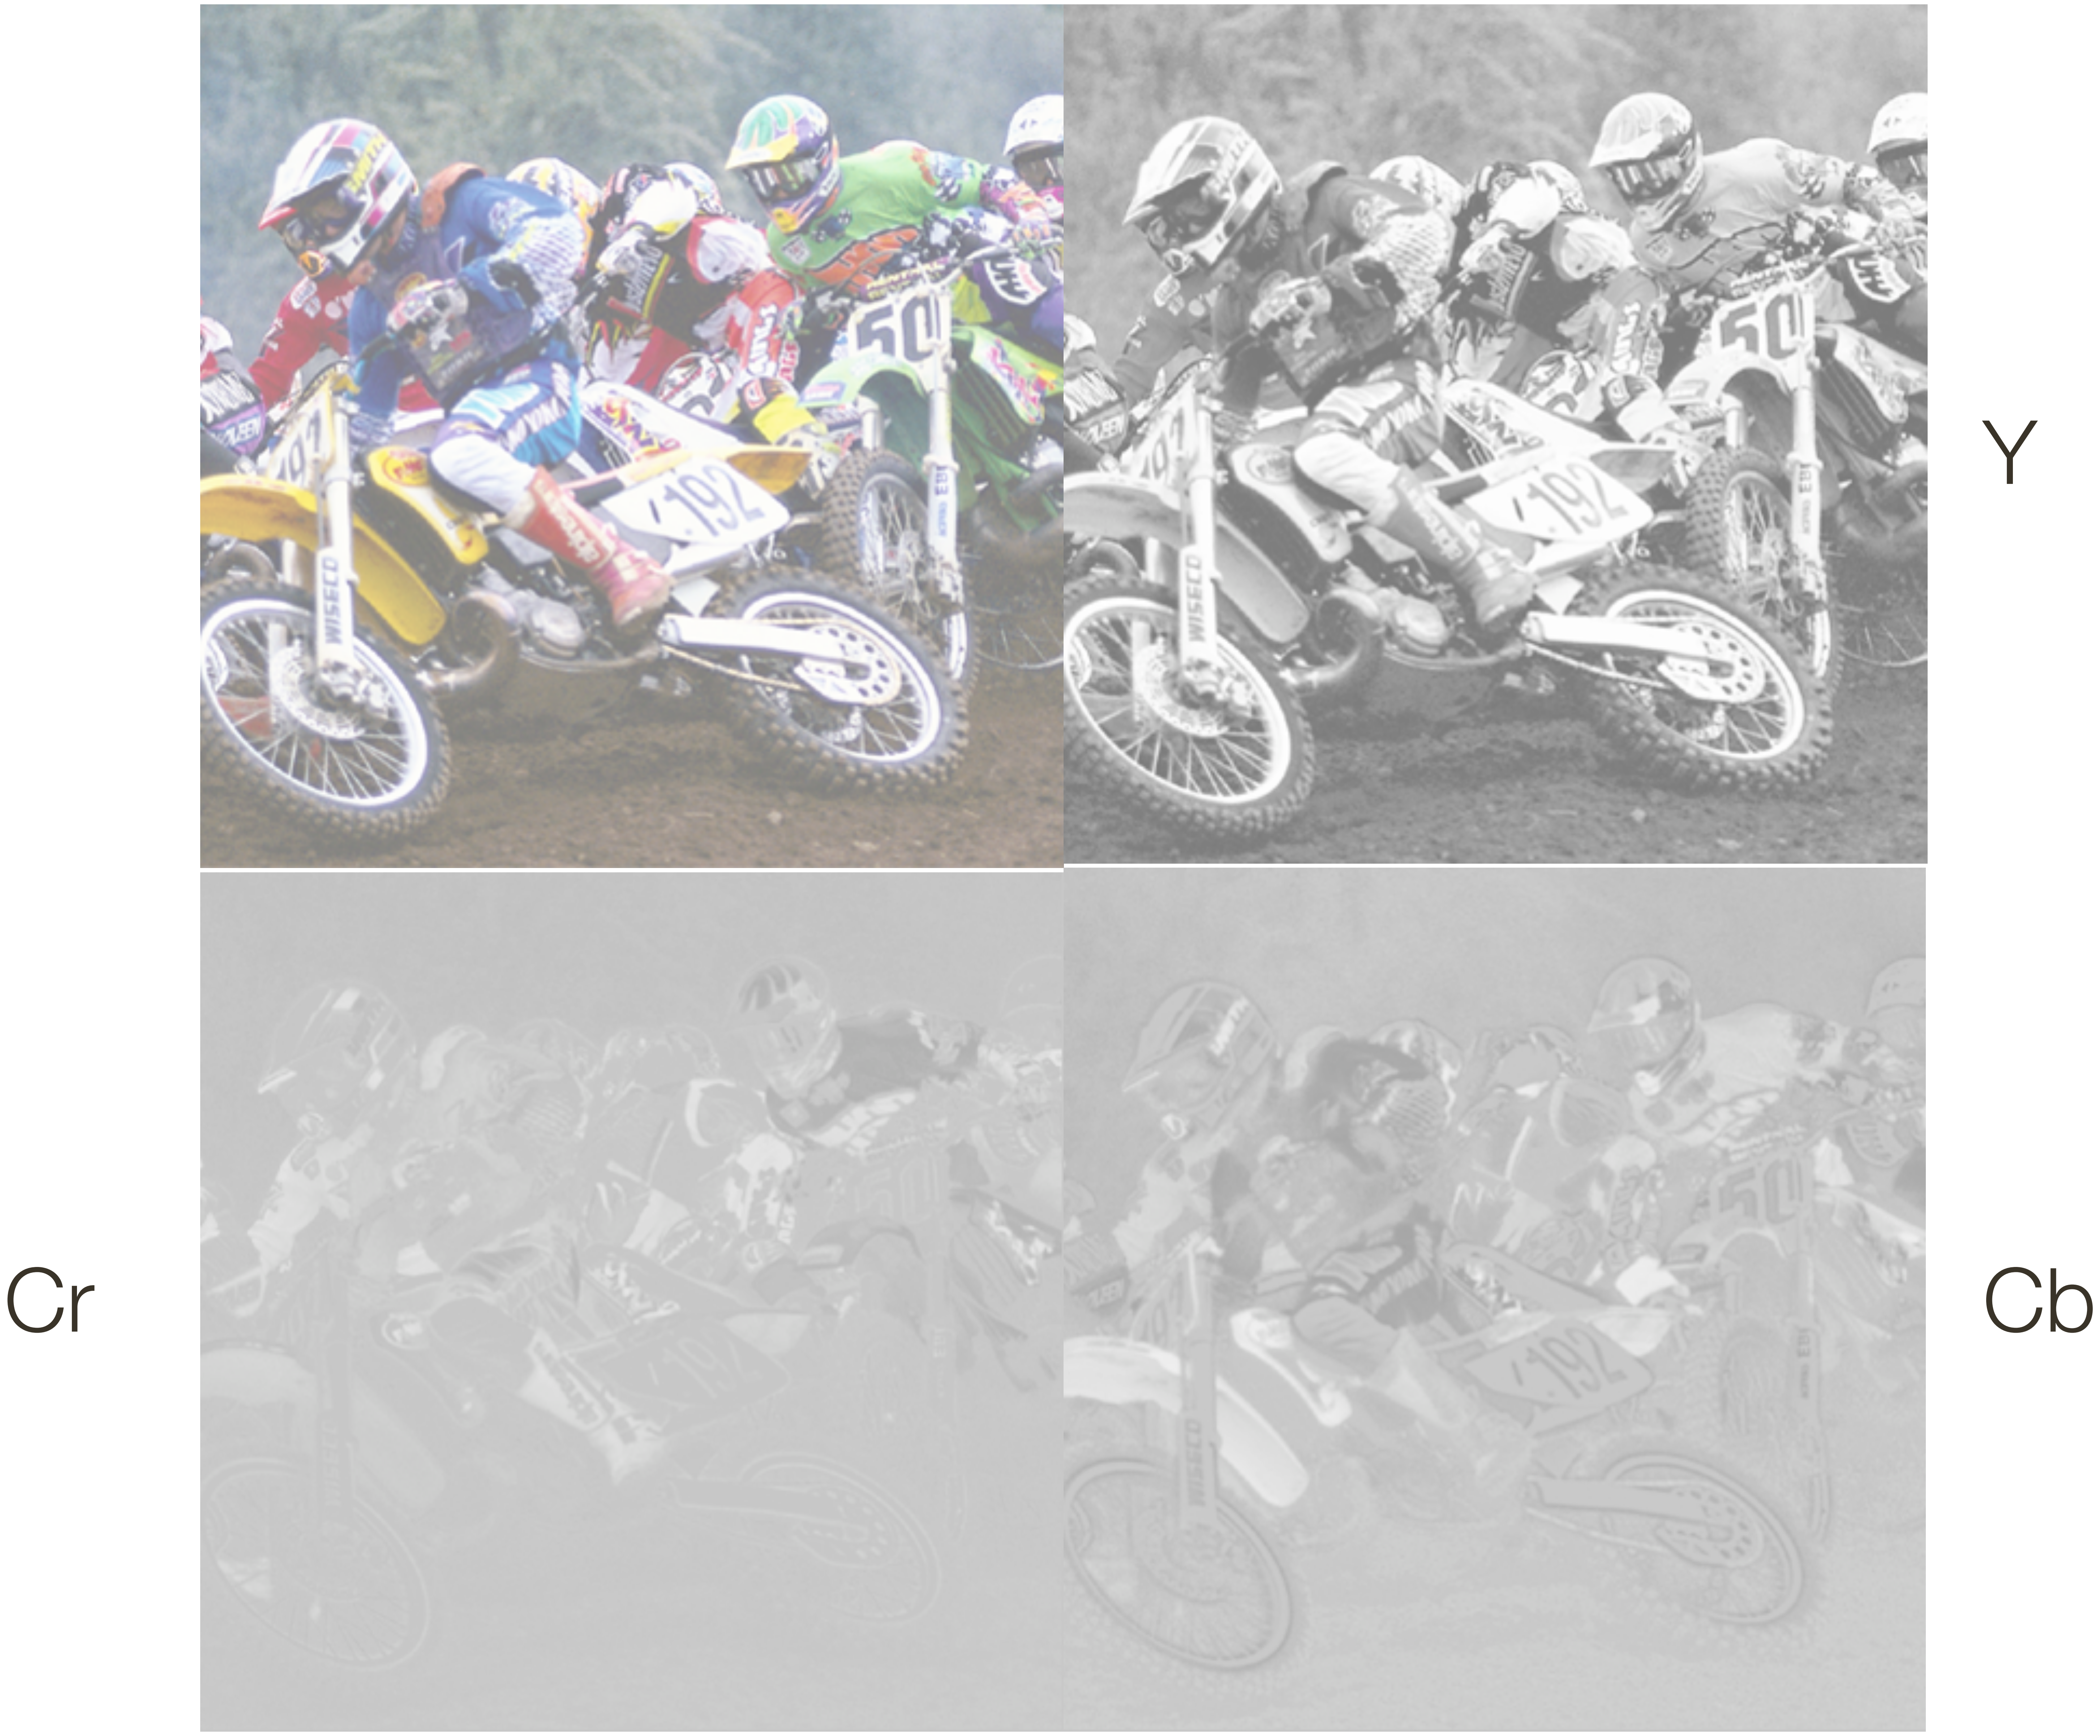

Blurring **CbCr**

sigma = 1.0

Blurring **CbCr**

sigma = 2.0

Blurring **CbCr**

sigma = 4.0

Blurring **CbCr**

sigma = 8.0

Blurring **CbCr**

sigma = 16.0

Blurring **CbCr**

sigma = 32.0

Blurring **Y**

sigma = 1.0

Blurring **Y**

sigma = 2.0

Blurring **Y**

sigma = 4.0

Blurring **Y**

sigma = 8.0

Blurring **Y**

sigma = 16.0

Subsampling **CbCr** vs **Y**

Original

Chrominance
1/8 scale

Luminance
1/8 scale

Compressibility …

Cb+Cr are transmitted at 1/2 size for JPEG

Note that **human vision** uses a similar transform to this (opponent colours), also we have fewer cones than rods

Colour **Constancy**

Image colour depends on both light colour and surface colour

Colour constancy: determine perceived colour under different colours of lighting

It is surprisingly difficult to predict what colours a human will perceive in a complex scene
— depends on context, other scene information

Humans can usually perceive
— the colour a surface would have under white light

Environmental Effects

Chromatic adaptation: If the human visual system is exposed to a certain colour light for a while, colour perception starts to skew

Contrast effects: Nearby colours affect what is perceived

Summary

- Human colour **perception**
  - principle of trichromacy
  - colour matching experiments
- Colour **reproduction**
  - linear colour spaces
  - colour matching functions
- Colour **spaces**
  - multiple objectives: art/design orientation, perceptually uniform, image coding etc.