


# **CS5630** Physically Based Rendering

16. Color science for photography  
and rendering

# What visible light is

- **One octave of the electromagnetic spectrum (380-760nm)**



# What color is

---

- **Colors are the sensations that arise from light energy with different wavelength distributions**
- **Color is a phenomenon of human perception; it is **not** a universal property of light**
- **Roughly speaking, things appear “colored” when they depend on wavelength and “gray” when they do not.**

# Measuring light

- **Salient property is the *spectral power distribution (SPD)***

the amount of light present at each wavelength

units: Watts per nanometer (tells you how much power you'll find in a narrow range of wavelengths)


for color, often use “relative units”  
when overall intensity is not important

in a physics-based renderer: all  
radiometric quantities are actually  
wavelength dependent

Spectral Power: W/nm


Spectral Radiance:  
W/(nm m<sup>2</sup> sr)

amount of light =  $180 d\lambda$   
(relative units)



# The problem of color science

- **Build a model for human color perception**
- **That is, map a *physical light description* to a *perceptual color sensation***




*Physical*

*Perceptual*

[Stone 2003]

# The eye as a measurement device



- **We can model the low-level behavior of the eye by thinking of it as a light-measuring machine**

its optics are much like a camera

its detection mechanism is also much like a camera

- **Light is measured by the photoreceptors in the retina**

they respond to visible light

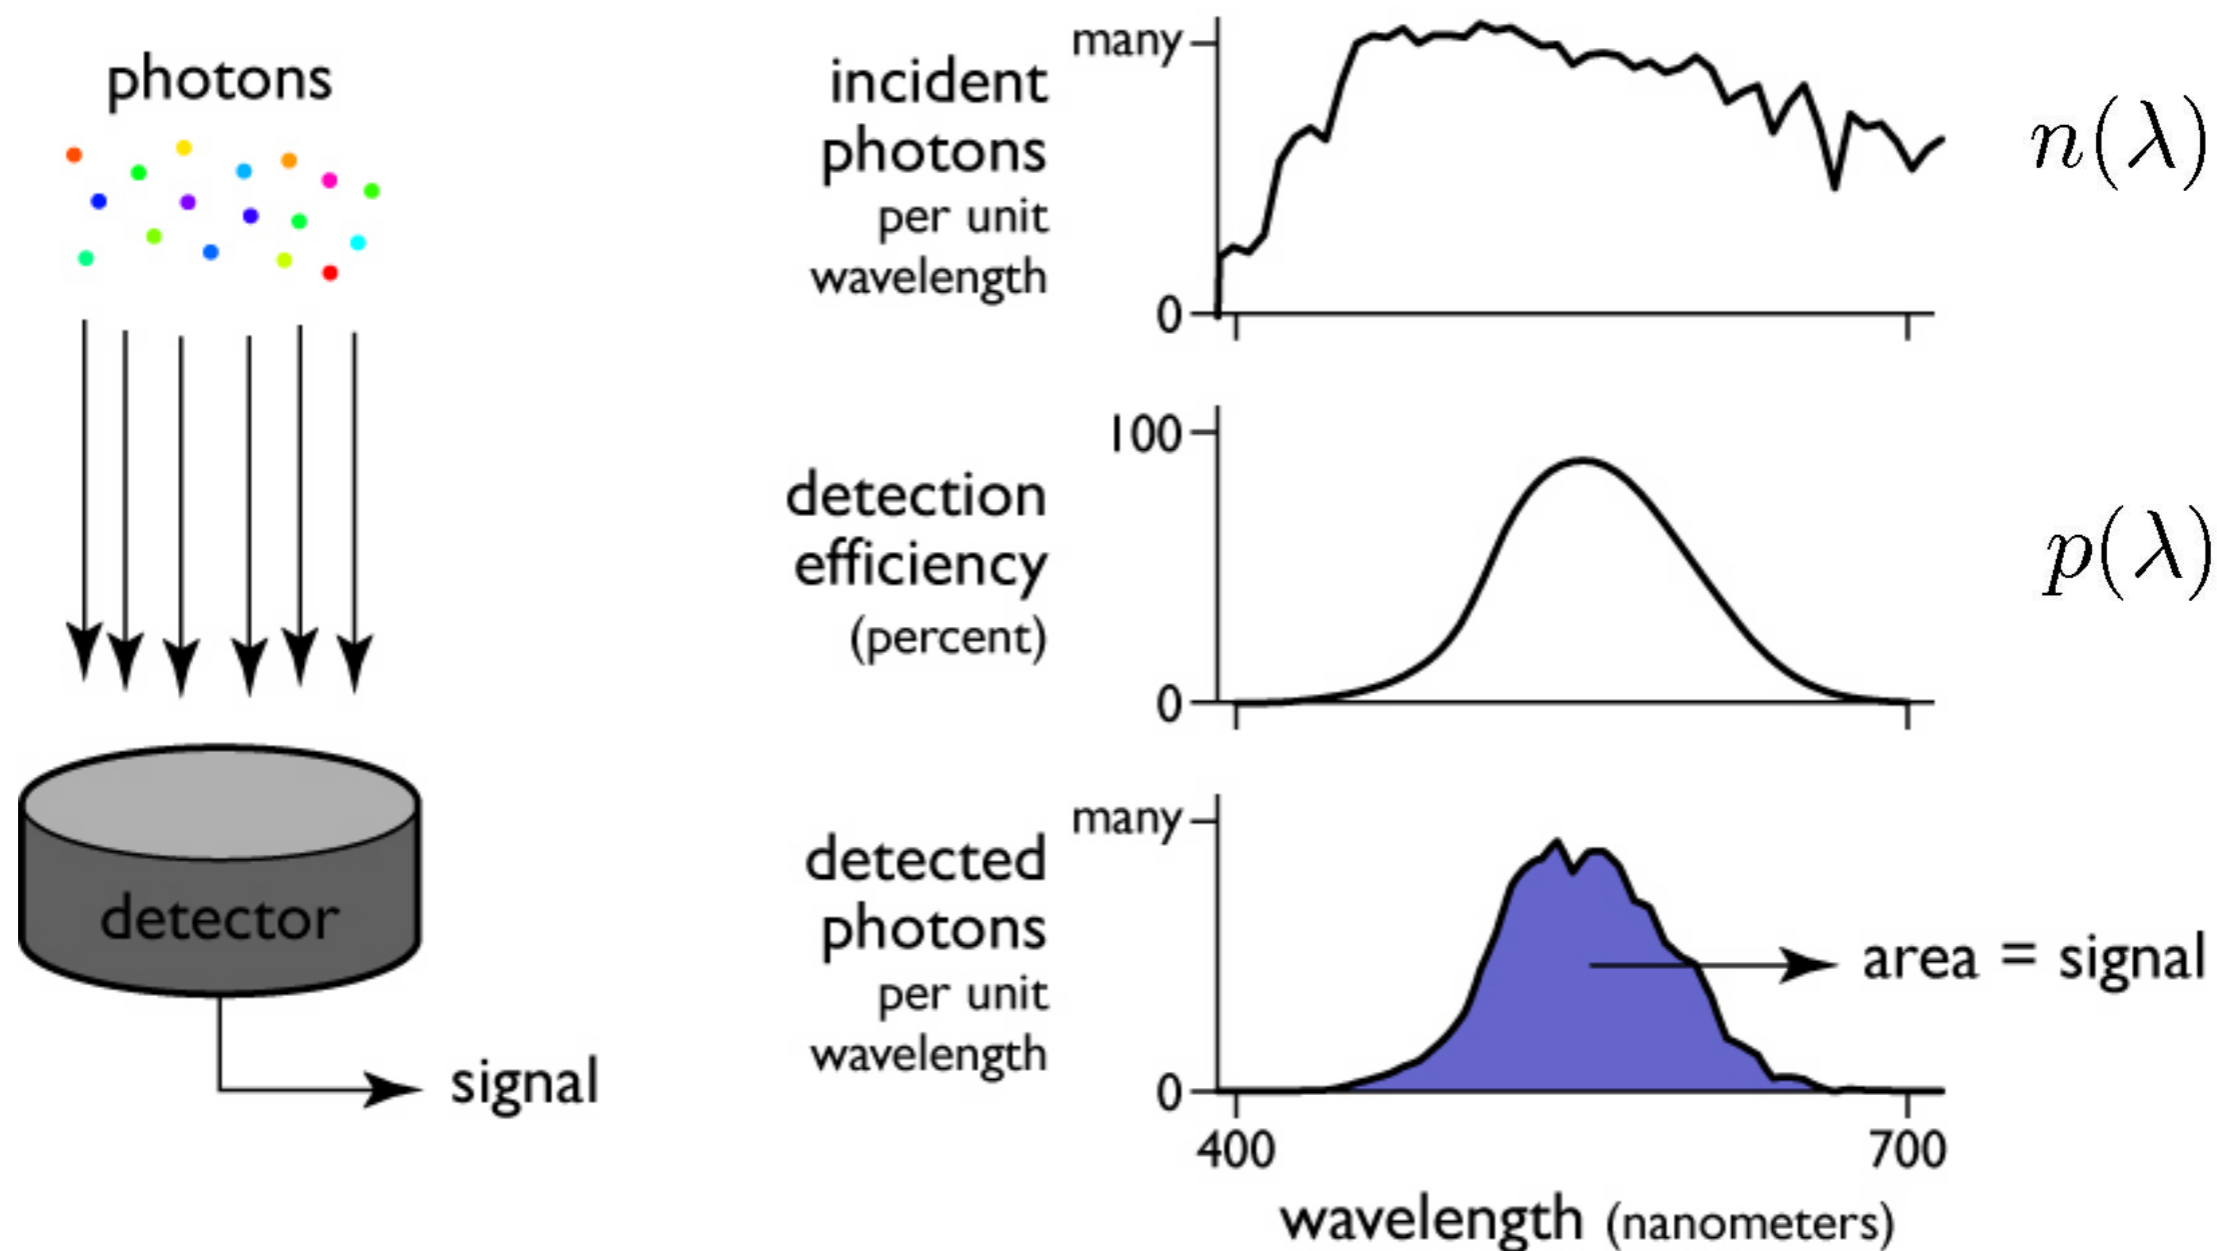
different types respond to different wavelengths

# A simple light detector

---

- **Produces a scalar value (a number) when photons land on it**
  - this value depends strictly on the number of photons detected
  - each photon has a probability of being detected that depends on the wavelength
  - there is no way to tell the difference between signals caused by light of different wavelengths: there is just a number
- **This model works for many detectors:**
  - based on semiconductors (such as in a digital camera)
  - based on visual photopigments (such as in human eyes)

# A simple light detector



$$X = \int n(\lambda)p(\lambda) d\lambda$$

# Light detection math

- **Same math carries over to power distributions**

spectrum entering the detector has its spectral power distribution (SPD),  $s(\lambda)$

detector has its *spectral sensitivity* or *spectral response*,  $r(\lambda)$

$$X = \int s(\lambda) r(\lambda) d\lambda$$

measured signal

input spectrum

detector's sensitivity

# Light detection math

---

$$X = \int s(\lambda)r(\lambda) d\lambda \quad \text{or} \quad X = s \cdot r$$

- **If we think of  $s$  and  $r$  as vectors, this operation is a dot product (aka inner product)**

in fact, the computation is done exactly this way, using sampled representations of the spectra.

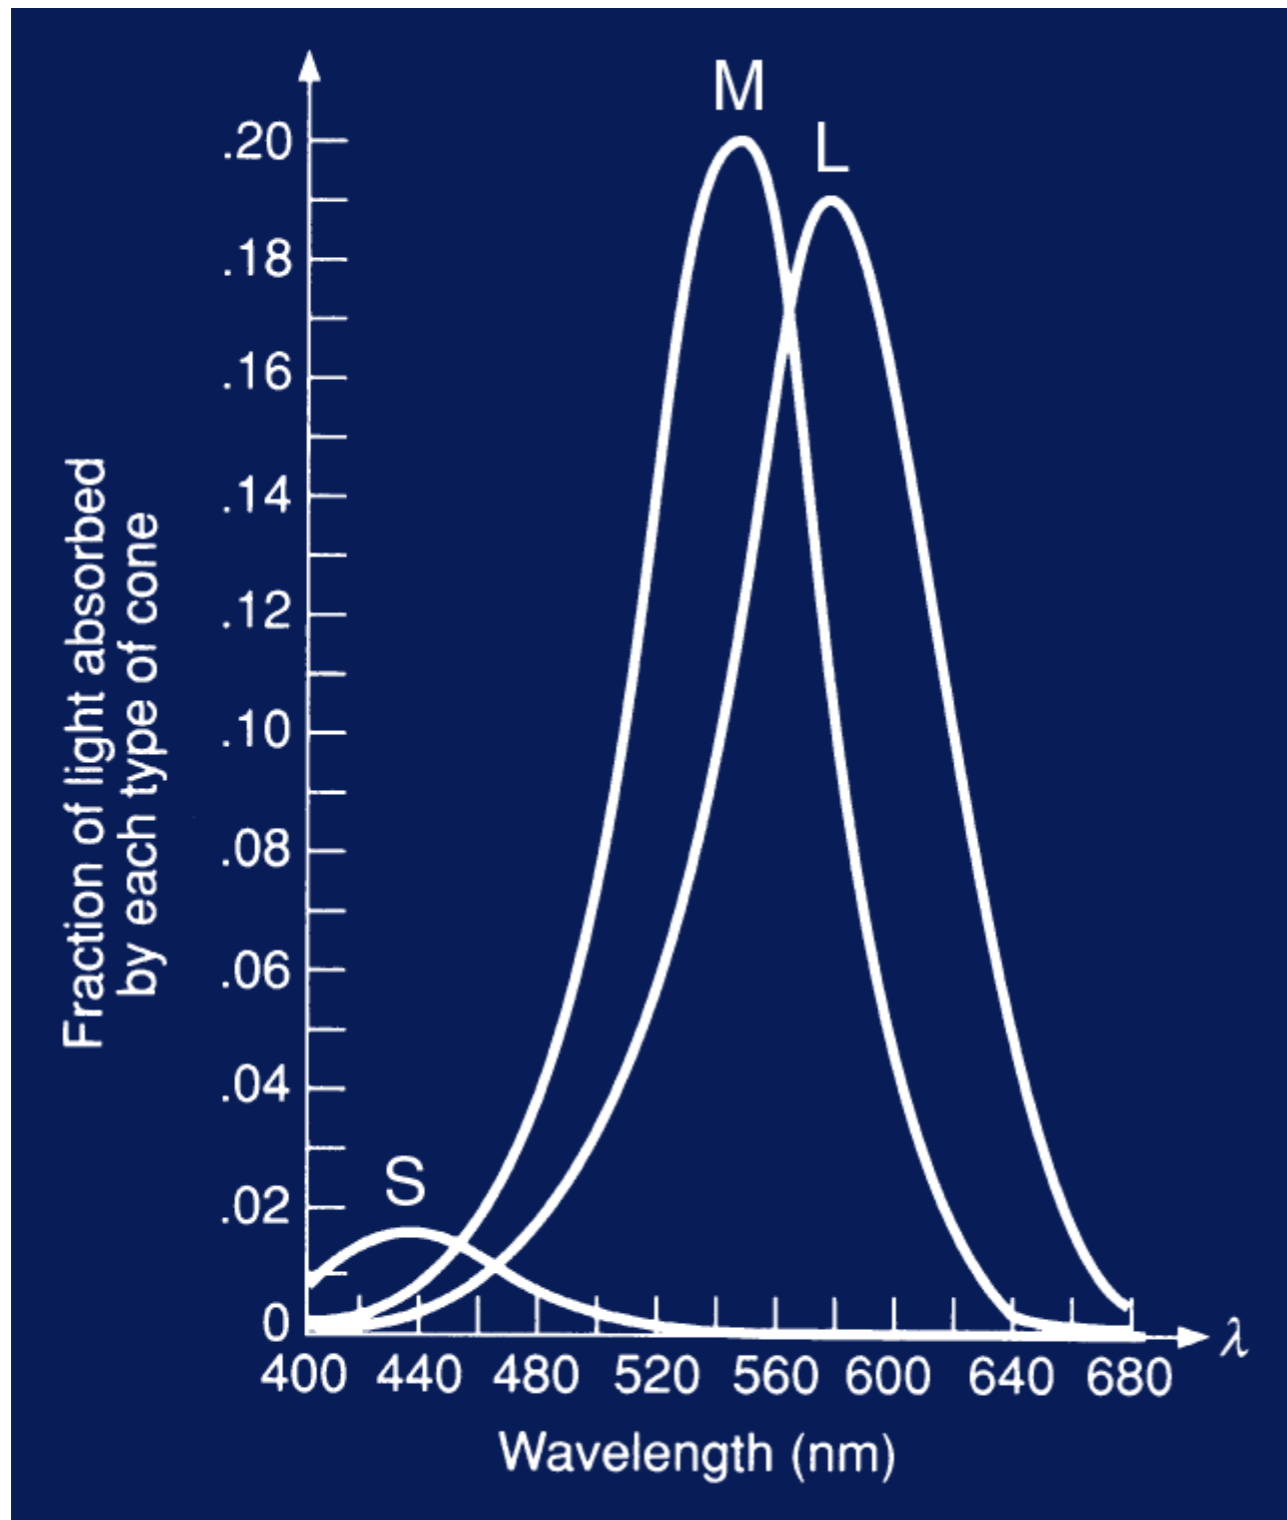
let  $\lambda_i$  be regularly spaced sample points  $\Delta\lambda$  apart; then:

$$\tilde{s}[i] = s(\lambda_i); \tilde{r}[i] = r(\lambda_i)$$

this sum is very clearly a dot product

$$\int s(\lambda)r(\lambda) d\lambda \approx \sum_i \tilde{s}[i]\tilde{r}[i] \Delta\lambda$$

# Cone Responses



- **S,M,L cones have broadband spectral sensitivity**
- **S,M,L neural response is integrated w.r.t.  $\lambda$** 
  - we'll call the response functions  $r_S, r_M, r_L$
- **Results in a trichromatic visual system**
- **S, M, and L are *tristimulus values***

[source unknown]

# Cone responses to a spectrum $s$

---

$$S = \int r_S(\lambda) s(\lambda) d\lambda = r_S \cdot s$$


$$M = \int r_M(\lambda) s(\lambda) d\lambda = r_M \cdot s$$

$$L = \int r_L(\lambda) s(\lambda) d\lambda = r_L \cdot s$$

# Colorimetry: an answer to the problem

- **Wanted to map a *physical light description* to a *perceptual color sensation***
- **Basic solution was known and standardized by 1930**

Though not quite in this form—more on that in a bit



*Physical*



$$S = r_S \cdot s$$
$$M = r_M \cdot s$$
$$L = r_L \cdot s$$

*Perceptual*

[Stone 2003]

# Basic fact of colorimetry

---

- **Take a spectrum (which is a function)**
- **Eye produces three numbers**
- **This throws away a lot of information!**

Quite possible to have two different spectra that have the same S, M, L tristimulus values

Two such spectra are *metamers*

# Pseudo-geometric interpretation

---

- **A dot product is a projection**
- **We are projecting a high dimensional vector (a spectrum) onto three vectors**
  - differences that are perpendicular to all 3 vectors are not detectable
- **For intuition, we can imagine a 3D analog**
  - 3D stands in for high-D vectors
  - 2D stands in for 3D
  - Then vision is just projection onto a plane


# Pseudo-geometric interpretation

- **The information available to the visual system about a spectrum is three values**

this amounts to a loss of information analogous to projection on a plane

- **Two spectra that produce the same response are metamers**

## Questions?



# Basic colorimetric concepts


- **Luminance**

the overall magnitude of the the visual response to a spectrum  
(independent of its color)

corresponds to the everyday concept “brightness”

determined by product of SPD with the *luminous efficiency function*  $V_\lambda$  that describes the eye’s overall ability to detect light at each wavelength

e.g. lamps are optimized to improve their luminous efficiency (tungsten vs. fluorescent vs. sodium vapor)



[Stone 2003]

# Luminance, mathematically

---

- **Y just has another response curve (like S, M, and L)**

$$Y = r_Y \cdot s$$

–  $r_Y$  is really called “ $V_\lambda$ ”

- **$V_\lambda$  is a linear combination of S, M, and L**

Has to be, since it's derived from cone outputs

# More basic colorimetric concepts

---

- **Chromaticity**

what's left after luminance is factored out (the color without regard for overall brightness)

scaling a spectrum up or down leaves chromaticity alone

- **Dominant wavelength**

many colors can be matched by white plus a spectral color

correlates to everyday concept “hue”



- **Purity**

ratio of pure color to white in matching mixture

correlates to everyday concept “colorfulness” or “saturation”


# Color reproduction

- **Have a spectrum  $s$ ; want to match on RGB monitor**
  - “match” means it looks the same
  - any spectrum that projects to the same point in the visual color space is a good reproduction
- **Must find a spectrum that the monitor *can* produce that is a metamer of  $s$**




**R, G, B?**

# Additive Color




[source unknown]

# LCD display primaries



Curves determined by (fluorescent or LED) backlight and filters

# Spatial integration



[source unknown]

# Color reproduction

---

- **Say we have a spectrum  $s$  we want to match on an RGB monitor**

“match” means it looks the same

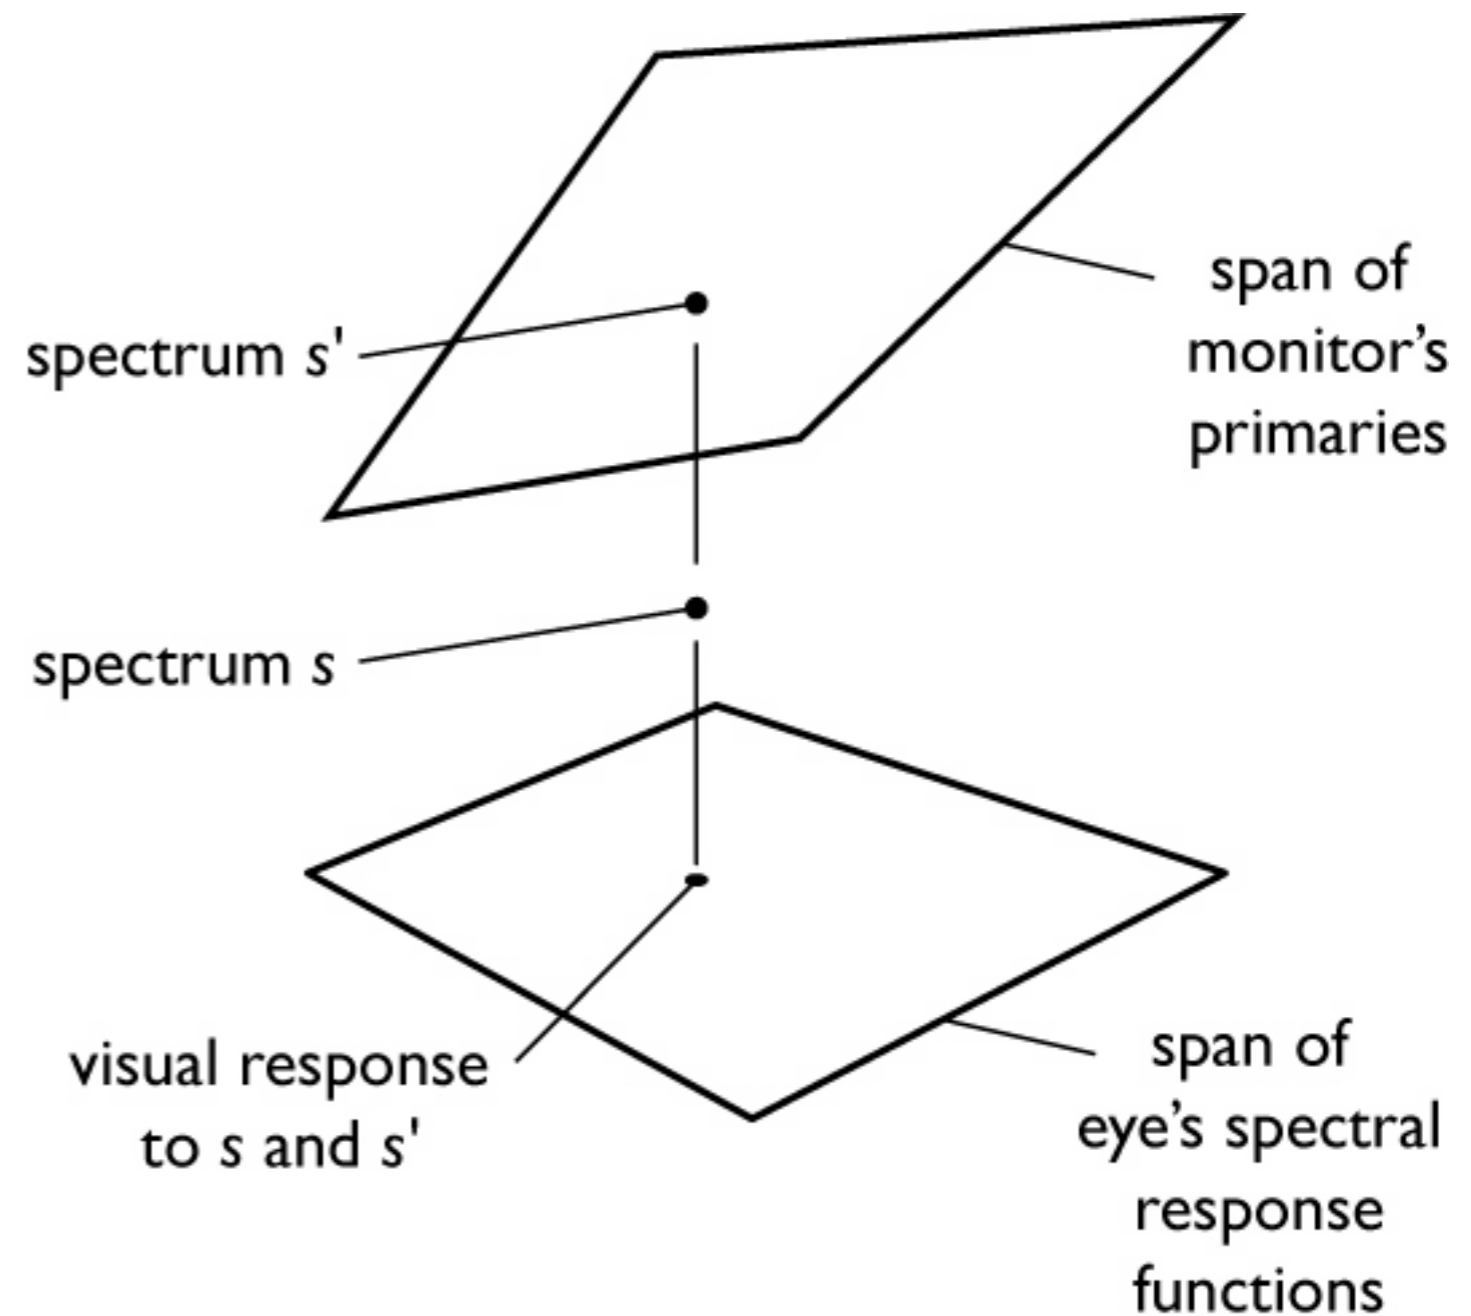
any spectrum that projects to the same point in the visual color space is a good reproduction

- **So, we want to find a spectrum that the monitor can produce that matches  $s$**

that is, we want to display a metamer of  $s$  on the screen

# Color reproduction

- **We want to compute the combination of R, G, B that will project to the same visual response as  $s$ .**



# Color reproduction as linear algebra

---

- **The projection onto the three response functions can be written in matrix form:**

$$\begin{bmatrix} S \\ M \\ L \end{bmatrix} = \begin{bmatrix} \text{---} r_S \text{---} \\ \text{---} r_M \text{---} \\ \text{---} r_L \text{---} \end{bmatrix} \begin{bmatrix} | \\ s \\ | \end{bmatrix}$$

or,

$$V = M_{SML} s.$$

# Color reproduction as linear algebra

- **The spectrum that is produced by the monitor for the color signals  $R$ ,  $G$ , and  $B$  is:**

$$s_a(\lambda) = R s_r(\lambda) + G s_g(\lambda) + B s_b(\lambda).$$

- **Again the discrete form can be written as a matrix:**

$$\begin{bmatrix} | \\ s_a \\ | \end{bmatrix} = \begin{bmatrix} | & | & | \\ s_R & s_G & s_B \\ | & | & | \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix} =$$

or,

$$s_a = M_{RGB} C.$$

# Color reproduction as linear algebra

- **What color do we see when we look at the display?**

Feed  $C$  to display

Display produces  $s_a$

Eye looks at  $s_a$  and produces  $V$

$$V = M_{SML} M_{RGB} C$$

$$\begin{bmatrix} S \\ M \\ L \end{bmatrix} = \begin{bmatrix} r_S \cdot s_R & r_S \cdot s_G & r_S \cdot s_B \\ r_M \cdot s_R & r_M \cdot s_G & r_M \cdot s_B \\ r_L \cdot s_R & r_L \cdot s_G & r_L \cdot s_B \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix}$$

# Color reproduction as linear algebra

- **Goal of reproduction: visual response to  $s$  and  $s_a$  is the same:**


$$M_{SML} \tilde{s} = M_{SML} \tilde{s}_a.$$

- **Substituting in the expression for  $s_a$ ,**

$$M_{SML} \tilde{s} = M_{SML} M_{RGB} C$$

$$C = \underbrace{(M_{SML} M_{RGB})^{-1} M_{SML}}_{\text{color matching matrix for RGB}} \tilde{s}$$

*color matching matrix for RGB*



These curves are the color-matching functions for the 1931 standard observer, The average results of 17 color-normal observers having matched each wavelength of the equal-energy spectrum with primaries of 435.8 nm, 546.1 nm, and 700 nm.

# Recap

---

- **We now know how to match any color from the real world on a display**
- **We don't need to know the whole spectrum, only the projections onto S, M, and L response functions**
- **There is then a simple linear procedure to work out the combination of any 3 primaries to match the color**

Questions?

# Color spaces

---

- **Need three numbers to specify a color**

but what three numbers?

*a color space* is an answer to this question

- **Stored numbers often map nonlinearly to intensity of primary**

enables nonuniform quantization (smaller quantization steps in dark)

common scheme is  $R = (n_R/255)^\gamma$

- **Common example: monitor RGB**

define colors by what  $R$ ,  $G$ ,  $B$  signals will produce them on your monitor

(in math,  $s = RR + GG + BB$  for some spectra  $\mathbf{R}$ ,  $\mathbf{G}$ ,  $\mathbf{B}$ )

device dependent (depends on gamma, phosphors, gains, ...)

if I choose RGB by looking at my monitor and send it to you, you may not see the same color

also leaves out some colors (limited *gamut*), e.g. vivid yellow

# Standard color spaces

---

- **Standardized RGB (sRGB)**

  - makes a particular monitor RGB standard

  - standard quantization curve is almost  $\gamma = 2.2$

  - other color devices simulate that monitor by calibration

  - sRGB is usable as an interchange space; widely adopted today

  - gamut is still limited

- **Other RGB spaces**

  - ITU Rec. 709 (older standard primaries used in sRGB)

  - ITU Rec. 2020 (wider gamut color space for HDR and 4/8K TVs)

  - DCI-P3 (standard for digital cinema, also used in computer displays)

  - ACES (newer film industry standard, imaginary primaries)


# A universal color space: XYZ

- **Standardized by CIE (Commission Internationale de l'Eclairage, the standards organization for color science)**
- **Based on three “imaginary” primaries X, Y, and Z**  
(in math,  $s = XX + YY + ZZ$ )

imaginary = only realizable by spectra that are negative at some wavelengths

*any* stimulus can be matched with positive X, Y, and Z

separates out luminance: **X**, **Z** have zero luminance, so Y tells you the luminance by itself



[source unknown]

The 1931 standard observer, as it is usually shown.

# Separating luminance, chromaticity

---

- **Luminance:**  $Y$
- **Chromaticity:**  $x, y, z$ , defined as

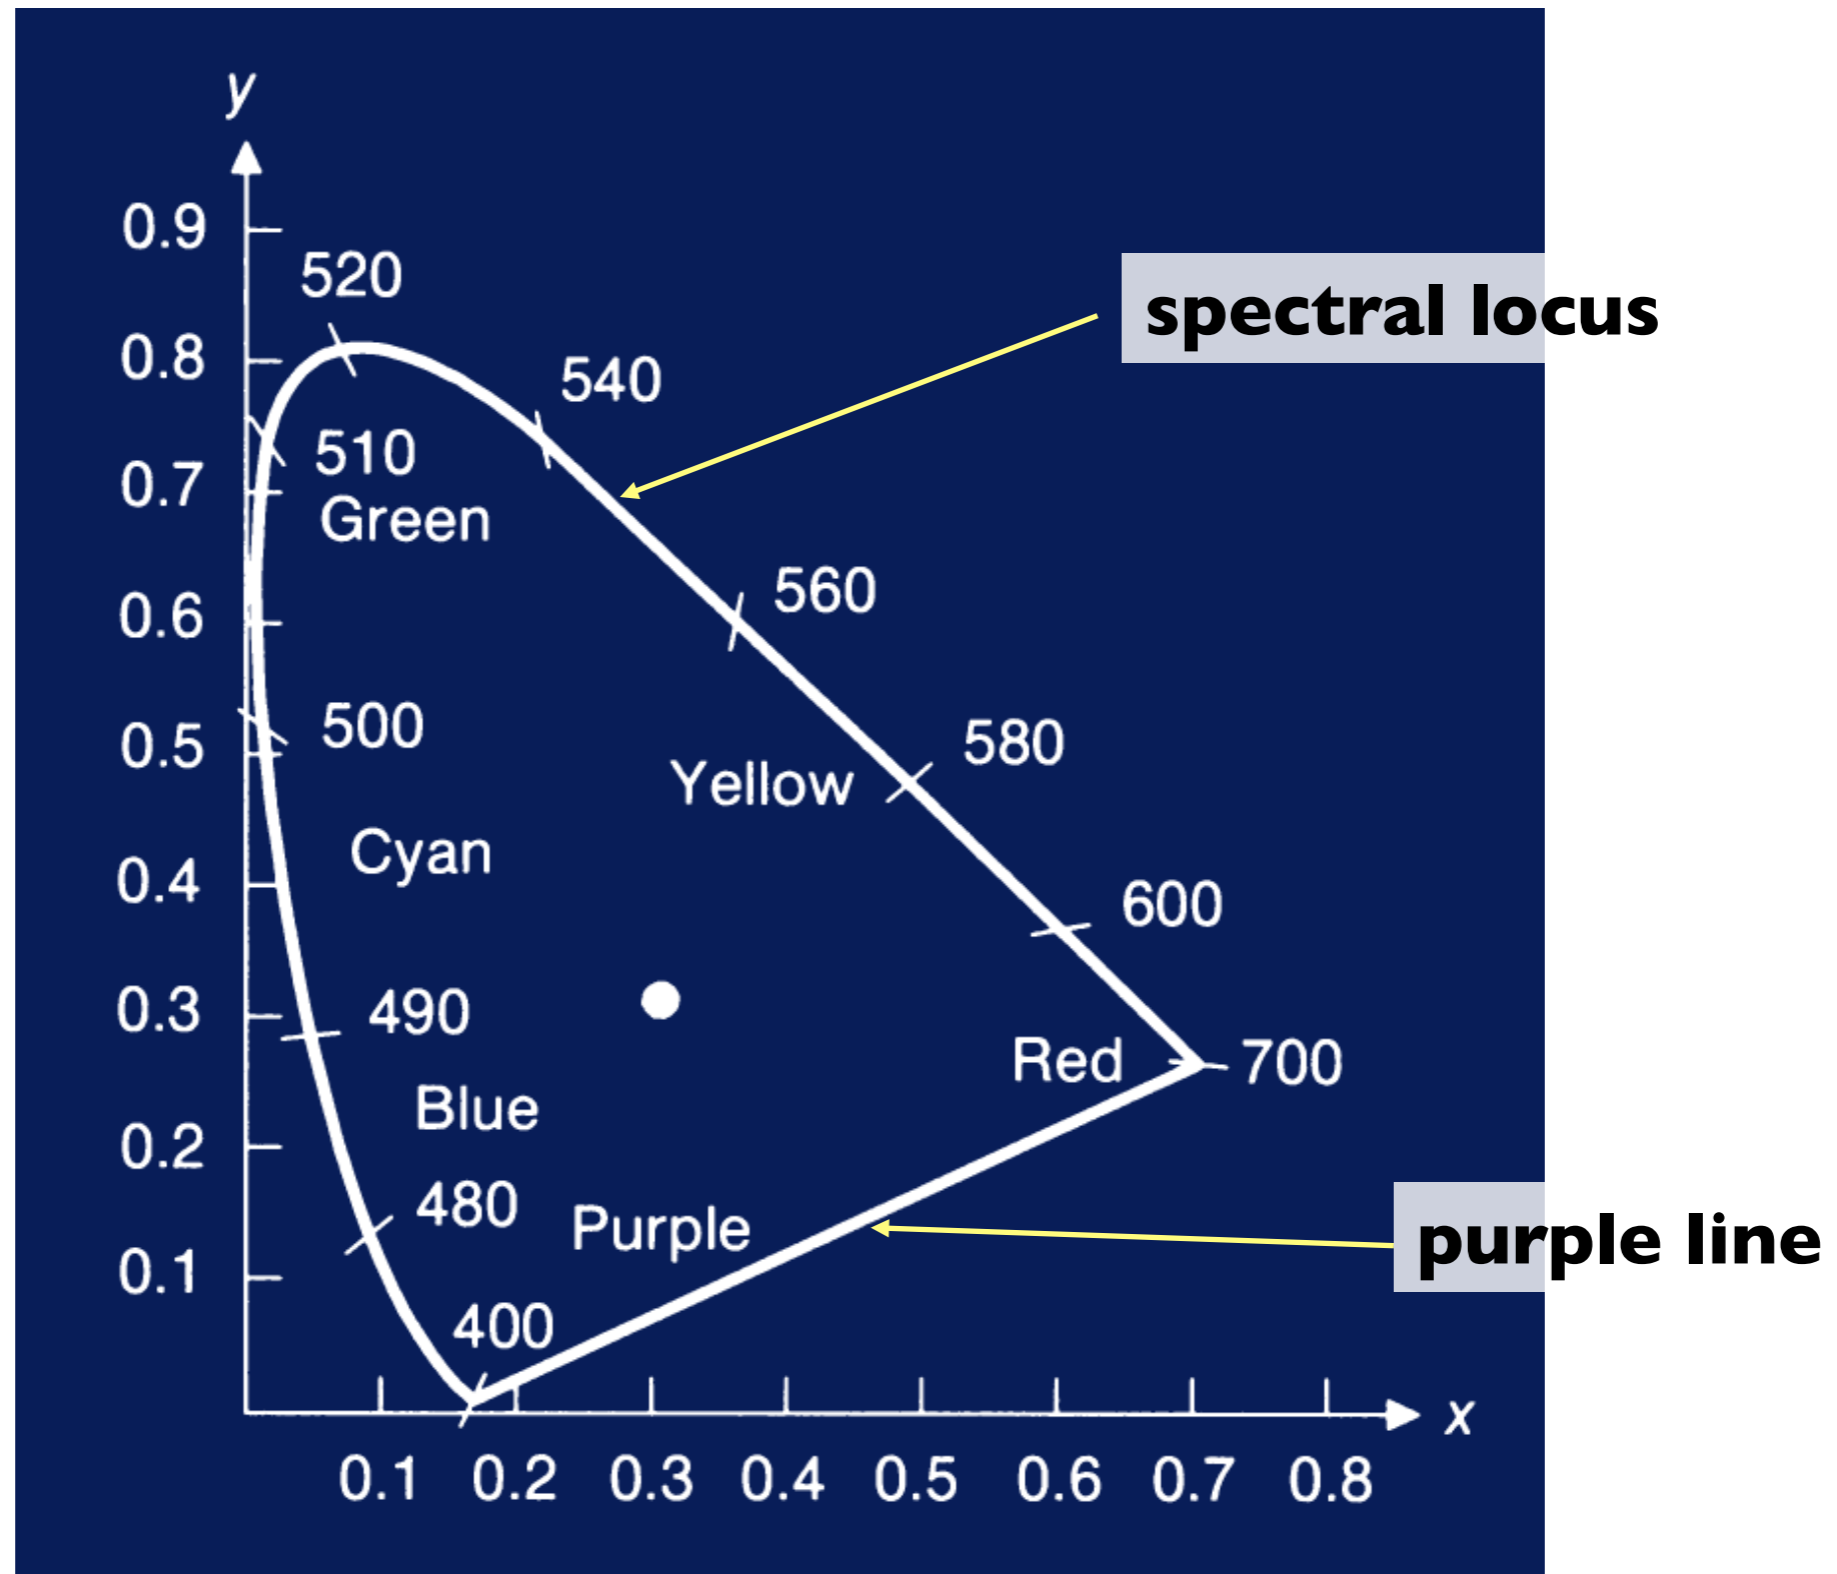
$$x = \frac{X}{X + Y + Z}$$

$$y = \frac{Y}{X + Y + Z}$$

$$z = \frac{Z}{X + Y + Z}$$

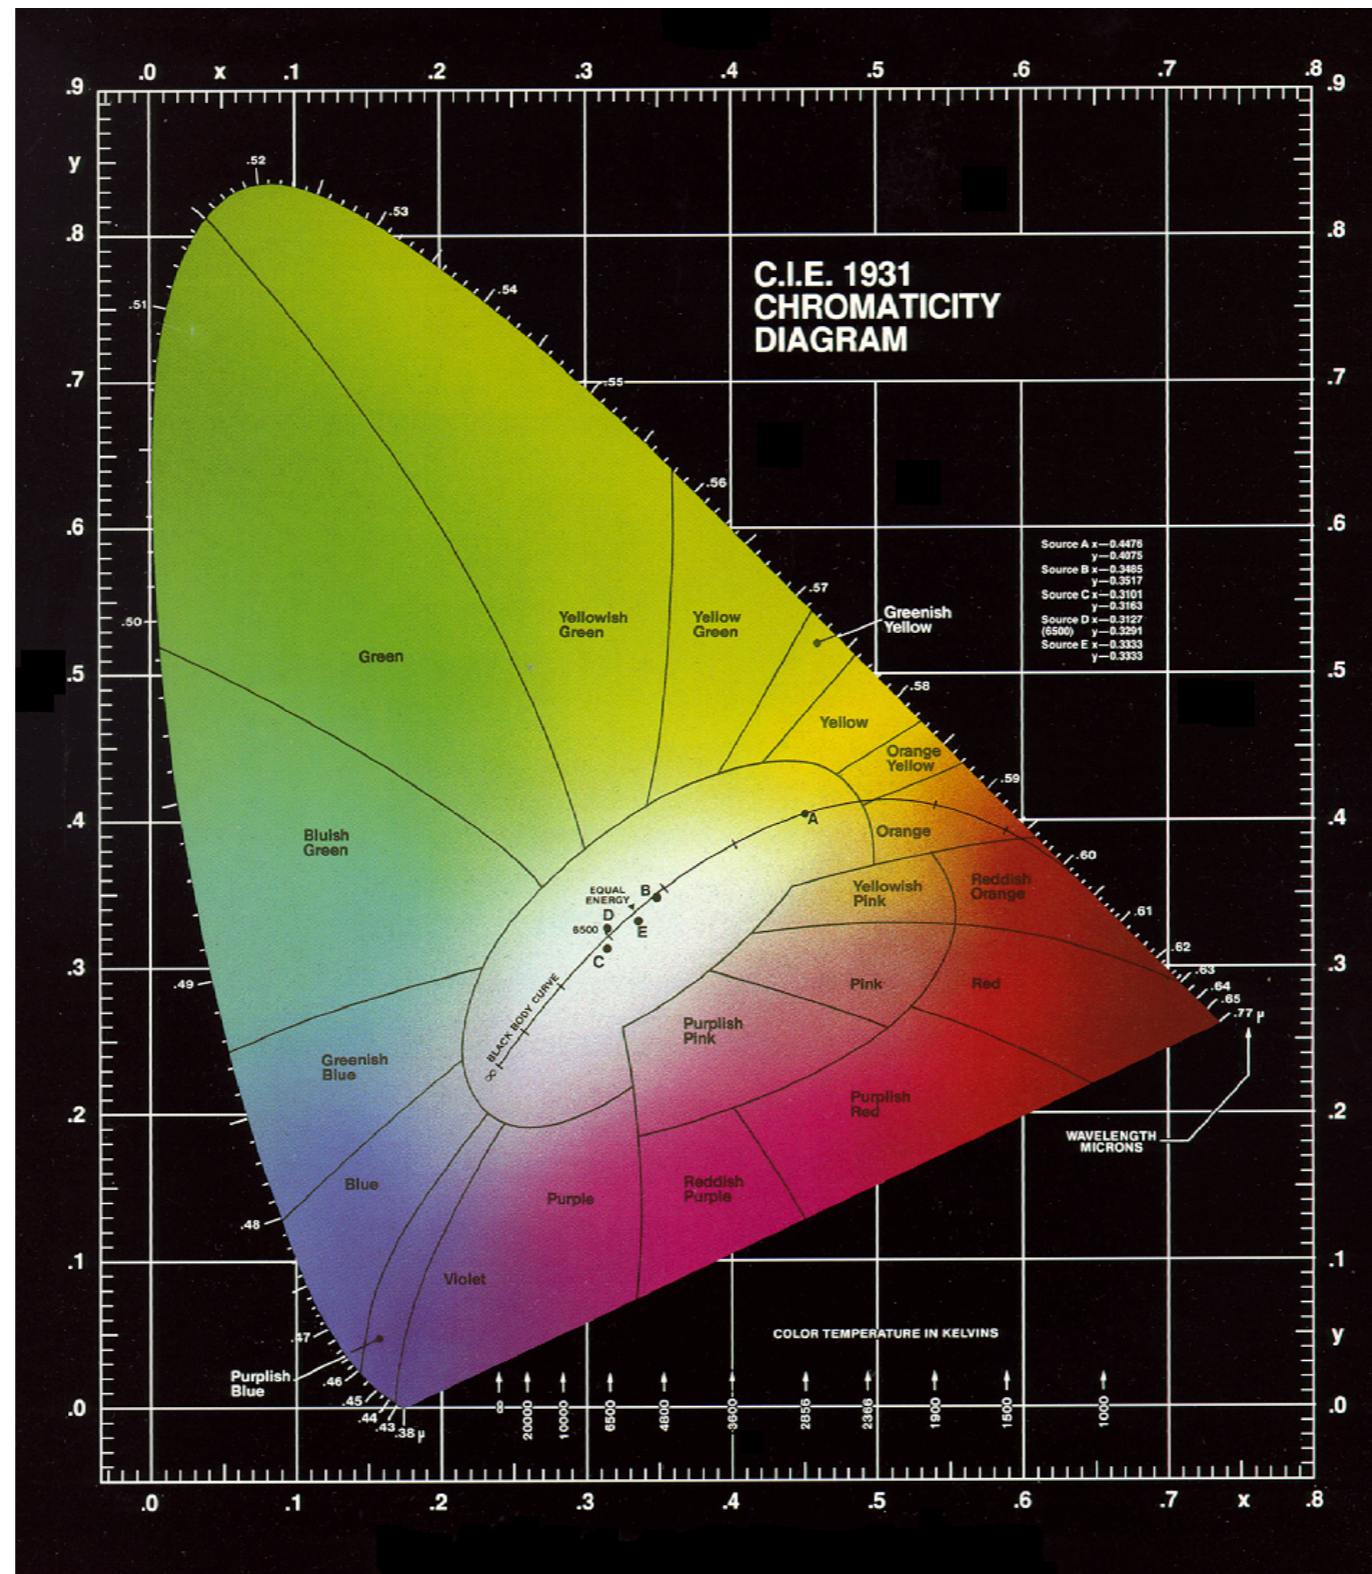
since  $x + y + z = 1$ , we only need to record two of the three  
usually choose  $x$  and  $y$ , leading to  $(x, y, Y)$  coords

# Chromaticity Diagram



[source unknown]


# Chromaticity Diagram



[source unknown]

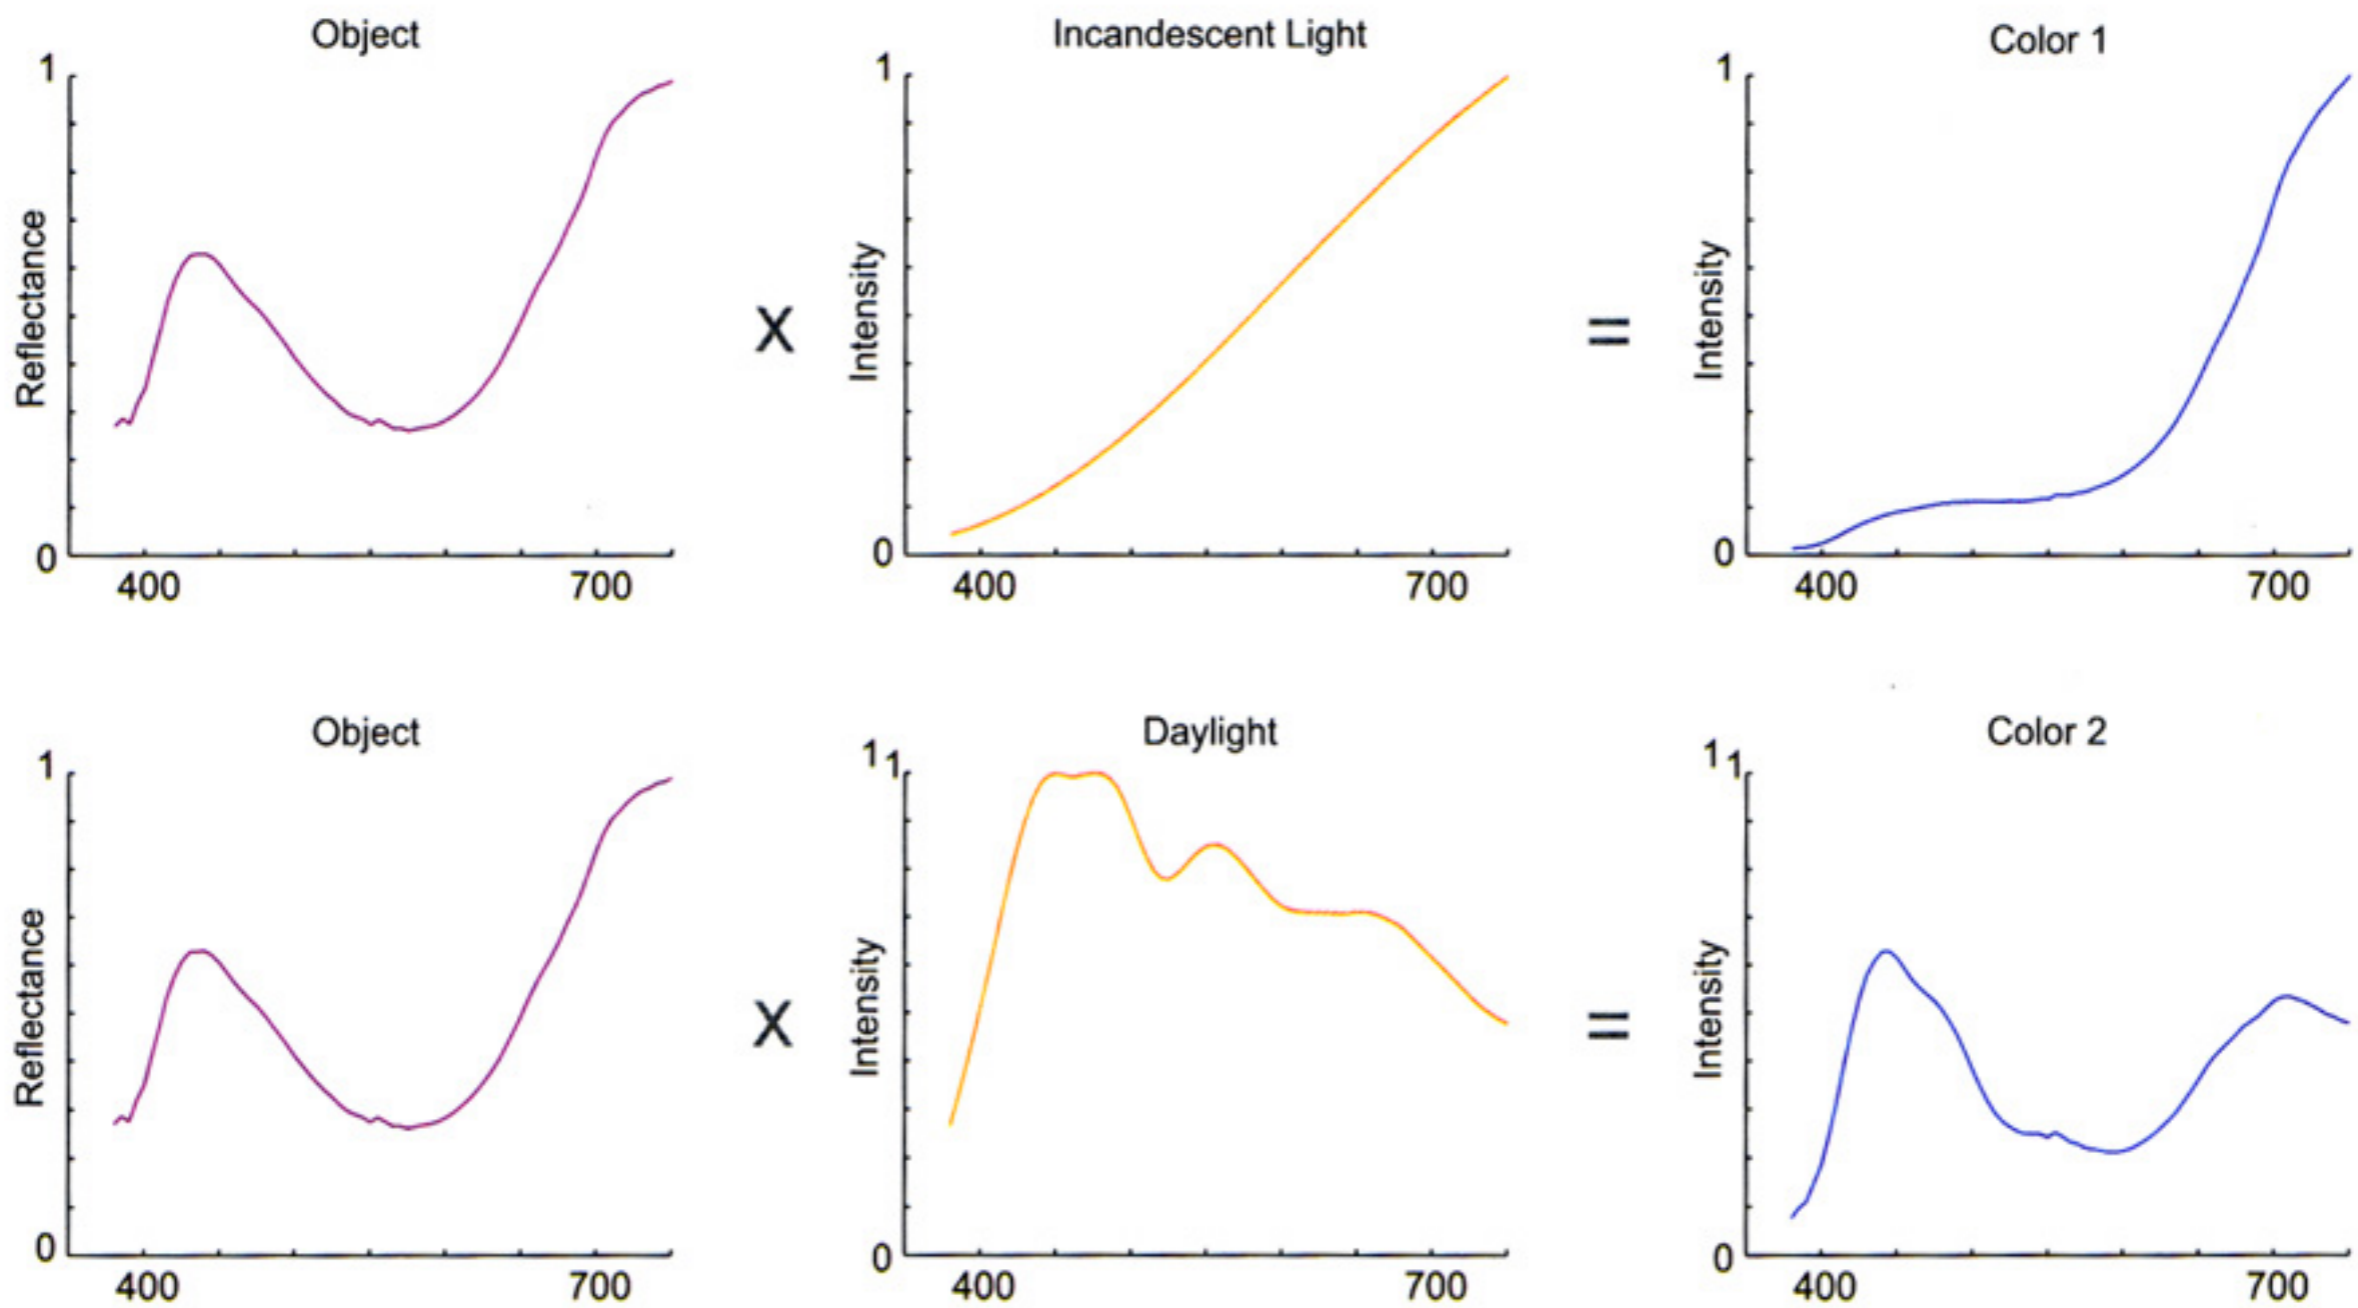
# Color Gamuts

- **Displays can't produce all visible colors**
- **For additive color, gamut is the triangle defined by the chromaticities of the three primaries.**



[PBR Fig. 4.23]

# Reflection from colored surface



[Stone 2003]

# Surface reflection in rendering

---

- **Reflection happens separately per wavelength**
  - unless there is fluorescence, which we usually don't model
- **Basic renderers use RGB for everything**
  - rays carry RGB color
  - surface reflection models return RGB colors
  - scale corresponding channels separately
  - in spectrum can think of three point samples or three piecewise-constant bands
- **This is not very accurate!**
  - more on the limitations in a bit

# Color constancy

---



# Chromatic adaptation

---

- **Objects have different spectra under different illuminants**
  - ...but your brain has no problem recognizing them anyway
- **The human visual system automatically detects the illuminant color and adjusts for it**
  - so the same object (usually) looks (roughly) the same color under a wide range of illumination conditions
  - this happens at a low level so you don't even notice
- **But color constancy is not perfect**
  - ...and indeed can't be, with just 3 color receptors
  - examples: sweater looks nice with pants in your closet, then looks different once you get out in the daylight
- **To simulate these effects we will need spectral rendering**

# Color sensing

---

- **Sensor is like eye**

- gives you projection onto a 3D (or >3D) space  
but it is the wrong space!

- **Problems with measured data**

- we have RGB, but not the right RGB

- projection onto sensitivities, not coefficients for primaries (always)

- projection onto wrong space (always in practice)

- results depend strongly on illuminant (help!)


# Sensor color properties

- Like eye, key property is the spectral sensitivity curves



KAI-2093 Image Sensor

## COLOR WITH MICROLENS QUANTUM EFFICIENCY




In a real camera, there will be a filter to block infrared

Figure 5: Quantum Efficiency Spectrum for Color Filter Array Sensors

# Camera color problem

- Given camera response, ~~determine~~ corresponding visual response  
guess
- This guess has to involve assumptions about which reflectance spectra are more likely
- Mathematical approach: assume spectra in fixed subspace
- Or, more often, just derive a transformation empirically



# White balancing

---

- **Problem with previous slide**

the camera-RGB colors depend on the illuminant

the matrix  $M$  only works for the illuminant that was used to calibrate

- **Solutions?**

calibrate separately for every illuminant?

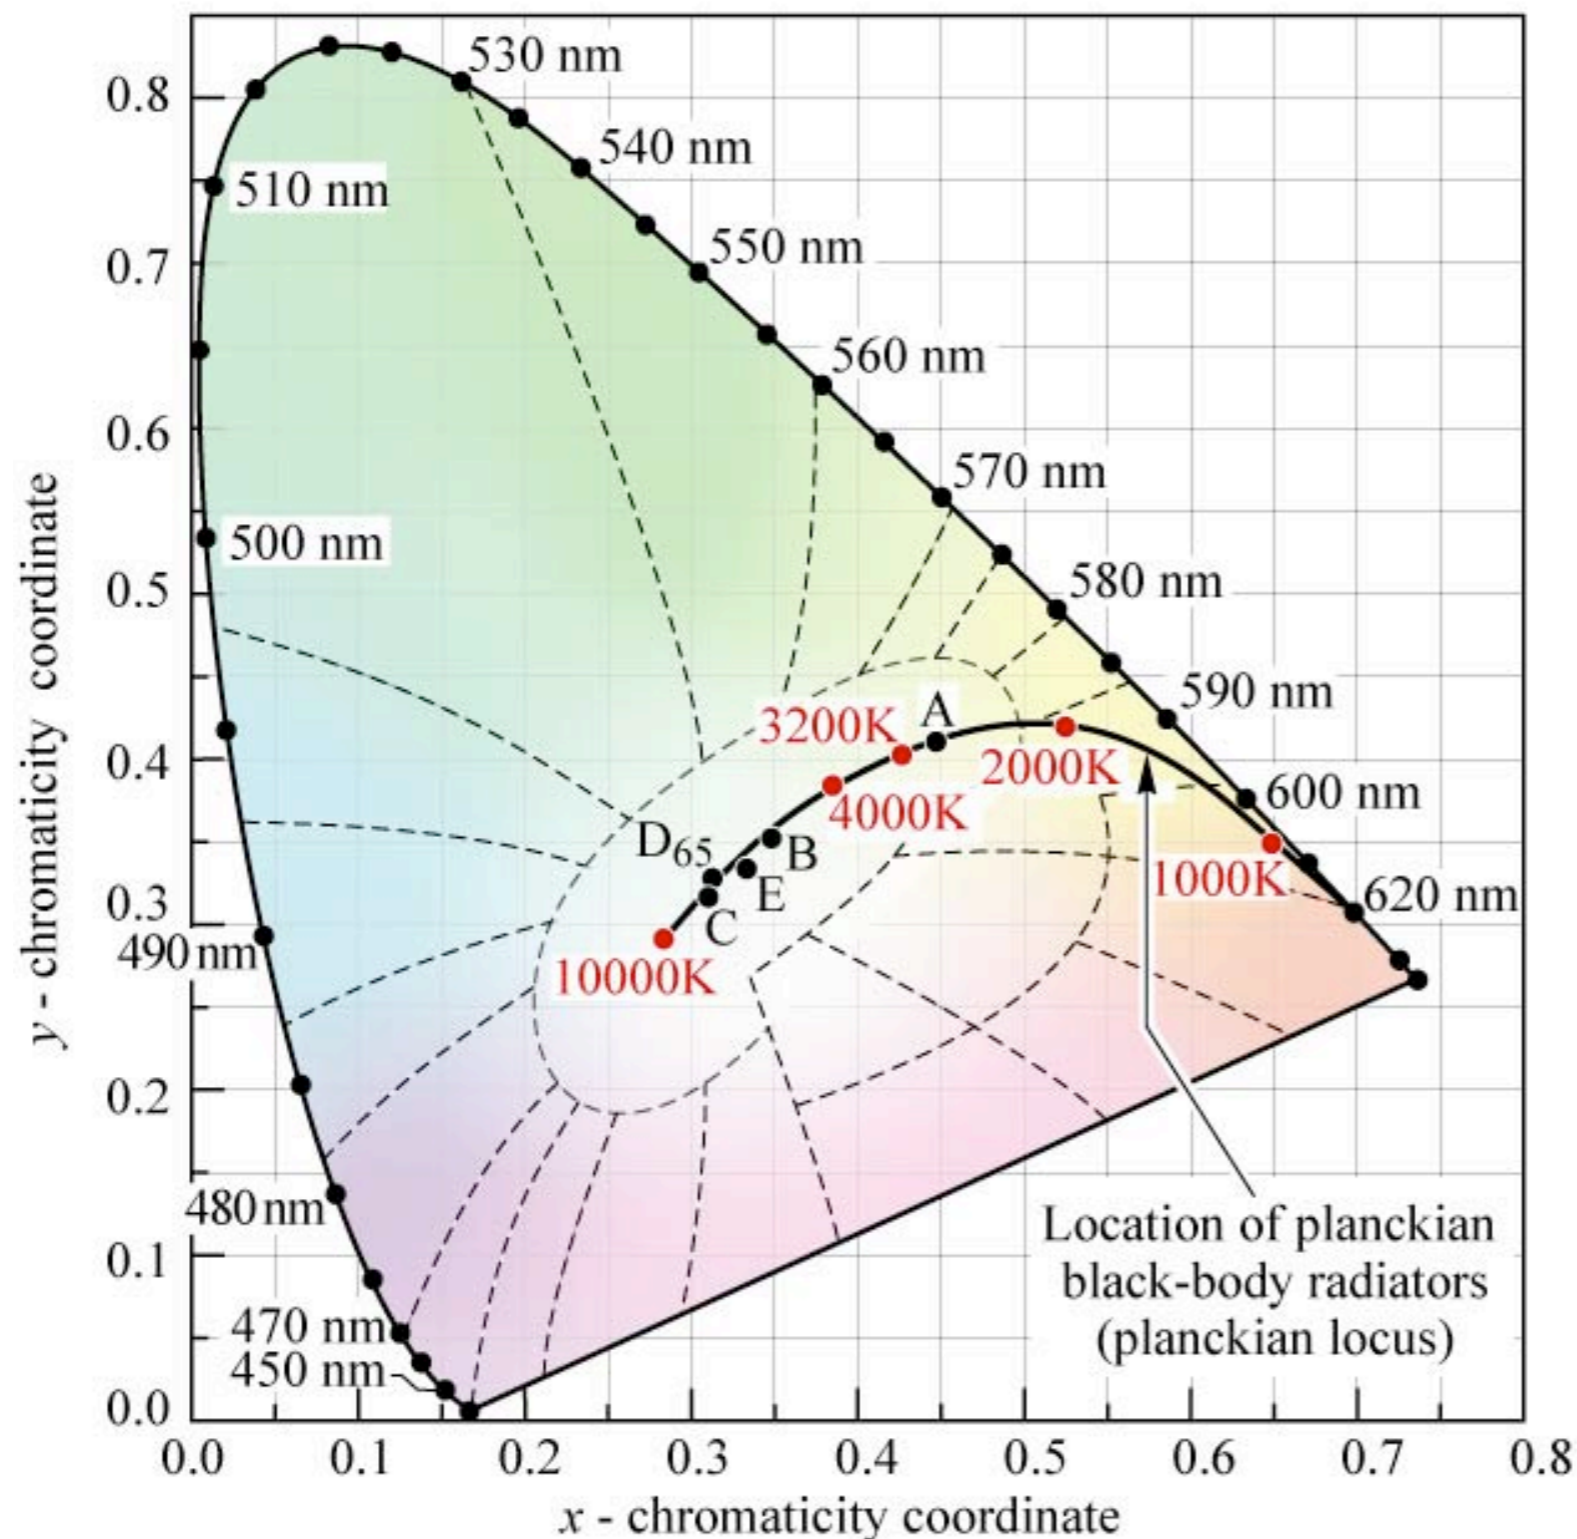
correct for illuminant first, then apply matrix!

- **Hypothesis of von Kries: eye accounts for illuminant by simply scaling the three cone signals separately**

some evidence this is a reasonable model for the eye

leads to “von Kries transform”: multiply by a diagonal matrix

# Range of illuminants



Illuminant A  
( $x, y$ ) = (0.4476, 0.4074)  
(Incandescent source,  $T = 2856$  K)

Illuminant B  
( $x, y$ ) = (0.3484, 0.3516)  
(Direct sunlight,  $T = 4870$  K)

Illuminant C  
( $x, y$ ) = (0.3101, 0.3162)  
(Overcast source,  $T = 6770$  K)

Illuminant D<sub>65</sub>  
( $x, y$ ) = (0.3128, 0.3292)  
(Daylight,  $T = 6500$  K)

Illuminant E (equal-energy point)  
( $x, y$ ) = (0.3333, 0.3333)

Fig. 18.3. Chromaticity diagram showing planckian locus, the standardized white illuminants A, B, C, D<sub>65</sub>, and E, and their color temperature (after CIE, 1978).

# White balancing steps

---

## **1. Determine the camera RGB of the illuminant (up to scale)**

professional/studio setting: photograph a gray card

poor man's version: find something gray in the image

practical solution: Auto White Balance software guesses

in a renderer: specify the illuminant in the input

## **2. Divide all the pixel values by the illuminant RGB**

undetermined scale factor

maybe fix luminance to 1

maybe scale lowest channel of illuminant to 1

### **• Now neutral colors are neutral!**

this is unbelievably important for getting nice color

# Putting it together: color processing

---

- **Calibrate your color matrix using a carefully white-balanced image**

when solving for  $M$ , constrain to ensure rows sum to 1  
(then  $M$  will leave neutral colors exactly alone)

- **For each photograph:**

1. determine illuminant
2. apply von Kries
3. apply color matrix
4. apply any desired nonlinearity
5. display the image!







# Color in spectral renderers

---

- **The right way to do it: model light as a function of  $\lambda$** 
  - ray radiance is a spectrum
  - light source radiance is a spectrum
  - BRDFs are spectral
  - pixel value is a spectrum
- **Computing a color image**
  - just like simulating the camera's spectral response
  - but: we pick the spectral response!
  - choose it to be colorimetric: use X, Y, Z color matching functions
  - result: image in linear XYZ space
- **Treat wavelength as another dimension for MC integration**

# Computing RGB colors

---

- **Inputs to the process**

  - spectrum of pixel color

  - chromaticity of light source

  - luminance that should map to output pixel value 1

  - together last two are a color that is the “white point”

- **Simple linear algebra problem**

  - earlier math gives raw RGB-to-XYZ matrix

  - solve for scale factors on X, Y, Z that map RGB [1,1,1]  
to the XYZ white point

  - invert to get XYZ-to-RGB matrix

# Perceptually organized color spaces

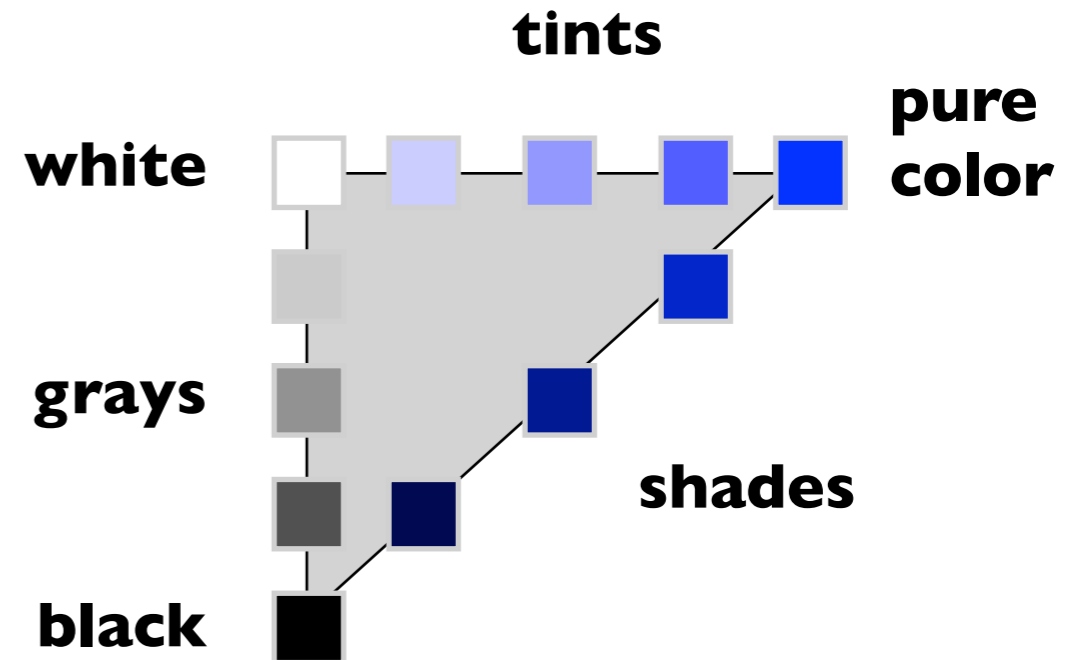
- **Artists often refer to colors as *tints*, *shades*, and *tones* of pure pigments**

tint: mixture with white

shade: mixture with black

tones: mixture with black and white

gray: no color at all  
(aka. neutral)



[after FvDFH]

- **This seems intuitive**

tints and shades are inherently related to the pure color

“same” color but lighter, darker, paler, etc.

# Perceptual dimensions of color

---

- **Hue**

- the “kind” of color, regardless of attributes
  - colorimetric correlate: dominant wavelength
  - artist’s correlate: the chosen pigment color

- **Saturation**

- the “colorfulness”
  - colorimetric correlate: purity
  - artist’s correlate: fraction of paint from the colored tube

- **Lightness (or value)**

- the overall amount of light
  - colorimetric correlate: luminance
  - artist’s correlate: tints are lighter. shades are darker



# Perceptual dimensions of color

---

- **There's good evidence (“opponent color theory”) for a neurological basis for these dimensions**

the brain seems to encode color early on using three axes:





white — black, red — green, yellow — blue

the white—black axis is lightness; the others determine hue and saturation

one piece of evidence: you can have a light green, a dark green, a yellow-green, or a blue-green, but you can't have a reddish green (just doesn't make sense)

thus red is the *opponent* to green

another piece of evidence: afterimages (next slide)






# RGB as a 3D space

---

- **A cube:**




**(demo of RGB cube)**

# Perceptual organization for RGB: HSV

- **Uses hue (an angle, 0 to 360), saturation (0 to 1), and value (0 to 1) as the three coordinates for a color**

the brightest available RGB colors are those with one of R,G,B equal to 1 (top surface)

each horizontal slice is the surface of a sub-cube of the RGB cube



**(demo of HSV color pickers)**

# Perceptually uniform spaces

---

- **Two major spaces standardized by CIE**

designed so that equal differences in coordinates produce equally visible differences in color

LUV: earlier, simpler space;  $L^*$ ,  $u^*$ ,  $v^*$

LAB: more complex but more uniform:  $L^*$ ,  $a^*$ ,  $b^*$

both separate luminance from chromaticity

including a gamma-like nonlinear component is important