


Color Science


CS 465 Lecture 20



[source unknown]

What light is


- Light is electromagnetic radiation
 - exists as oscillations of different frequency (or, wavelength)



[Lawrence Berkeley Lab / MicroWorlds]

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




What color is


- Colors are the sensations that arise from light energy of different wavelengths
 - we are sensitive from about 380 to 760 nm—one “octave”
- 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.

The problem of color science

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



The eye as a measurement device

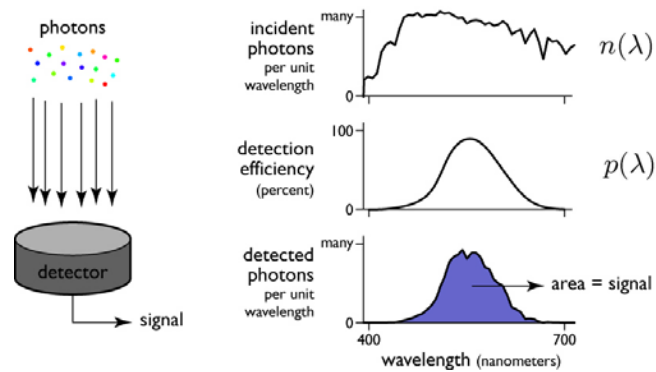


- We can model the low-level behavior of the eye by thinking of it as a light-measuring machine
 - its optics are much like a camera
 - its detection mechanism is also much like a camera
- Light is measured by the *photoreceptors* in the retina
 - they respond to visible light
 - different types respond 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



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

Cornell CS465 Fall 2007 • Lecture 20

© 2007 Doug James & Steve Marschner • 9

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

Cornell CS465 Fall 2007 • Lecture 20

© 2007 Doug James & Steve Marschner • 10

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

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


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

- this sum is very clearly a dot product

Cornell CS465 Fall 2007 • Lecture 20

© 2007 Doug James & Steve Marschner • 11

Cone Responses



- S,M,L cones have broadband spectral sensitivity
- S,M,L neural response is integrated w.r.t. λ
 - 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]

Cornell CS465 Fall 2007 • Lecture 20

© 2007 Doug James & Steve Marschner • 12

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



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




Cornell CS465 Fall 2007 • Lecture 20

© 2007 Doug James & Steve Marschner • 17

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_λ 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]

Cornell CS465 Fall 2007 • Lecture 20

© 2007 Doug James & Steve Marschner • 18

Luminance, mathematically

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

$$Y = r_Y \cdot s$$

- r_Y is really called “ V_λ ”
- V_λ is a linear combination of S, M, and L
 - Has to be, since it’s derived from cone outputs

Cornell CS465 Fall 2007 • Lecture 20

© 2007 Doug James & Steve Marschner • 19

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”

Cornell CS465 Fall 2007 • Lecture 20

© 2007 Doug James & Steve Marschner • 20

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




[cs417—Greenberg]

Cornell CS465 Fall 2007 • Lecture 20

© 2007 Doug James & Steve Marschner • 21

Additive Color




[source unknown]

Cornell CS465 Fall 2007 • Lecture 20

© 2007 Doug James & Steve Marschner • 22

CRT display primaries




- Curves determined by phosphor emission properties

Cornell CS465 Fall 2007 • Lecture 20

© 2007 Doug James & Steve Marschner • 23

LCD display primaries




[Fairchild 97]

- Curves determined by (fluorescent) backlight and filters

Cornell CS465 Fall 2007 • Lecture 20

© 2007 Doug James & Steve Marschner • 24

Combining Monitor Phosphors with Spatial Integration



Cornell CS465 Fall 2007 • Lecture 20

© 2007 Doug James & Steve Marschner • 25

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

Cornell CS465 Fall 2007 • Lecture 20

© 2007 Doug James & Steve Marschner • 26

Color reproduction

- We want to compute the combination of r, g, b that will project to the same visual response as s .



Cornell CS465 Fall 2007 • Lecture 20

© 2007 Doug James & Steve Marschner • 27

Color reproduction as linear algebra

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

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

or,

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

Cornell CS465 Fall 2007 • Lecture 20

© 2007 Doug James & Steve Marschner • 28

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:

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

- Substituting in the expression for s_a ,

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


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

Subtractive Color



[source unknown]

Reflection from colored surface



Cornell CS465 Fall 2007 • Lecture 20

© 2007 Doug James & Steve Marschner • 33

Subtractive color

- Produce desired spectrum by *subtracting* from white light (usually via absorption by pigments)
- Photographic media (slides, prints) work this way
- Leads to C, M, Y as primaries
- Approximately, 1-R, 1-G, 1-B

Cornell CS465 Fall 2007 • Lecture 20

© 2007 Doug James & Steve Marschner • 34

Color spaces

- Need three numbers to specify a color
 - but what three numbers?
 - a *color space* is an answer to this question
- 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 R, G, B)
 - device dependent (depends on gamma, phosphors, gains, ...)
 - therefore 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

Cornell CS465 Fall 2007 • Lecture 20

© 2007 Doug James & Steve Marschner • 35

Standard color spaces

- Standardized RGB (sRGB)
 - makes a particular monitor RGB standard
 - other color devices simulate that monitor by calibration
 - sRGB is usable as an interchange space; widely adopted today
 - gamut is still limited

Cornell CS465 Fall 2007 • Lecture 20

© 2007 Doug James & Steve Marschner • 36

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
 - key properties
 - 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

Separating luminance, chromaticity

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


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

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

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

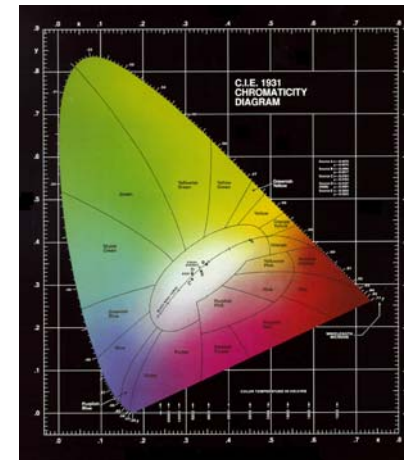
- since $x + y + z = 1$, we only need to record two of the three
 - usually choose x and y, leading to (x, y, Y) coords

Chromaticity Diagram



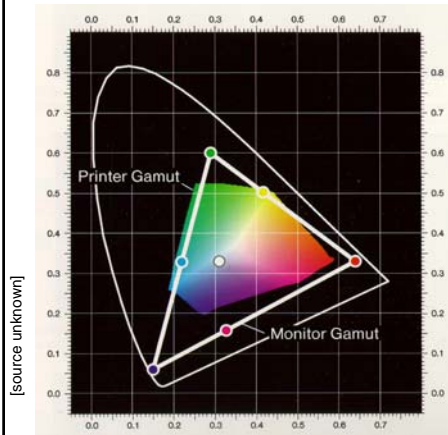
[source unknown]

Chromaticity Diagram



[source unknown]

Color Gamuts



Monitors/printers can't produce all visible colors

Reproduction is limited to a particular domain

For additive color (e.g. monitor) gamut is the triangle defined by the chromaticities of the three primaries.


Cornell CS465 Fall 2007 • Lecture 20

© 2007 Doug James & Steve Marschner • 41

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)



- This seems intuitive

- tints and shades are inherently related to the pure color
 - “same” color but lighter, darker, paler, etc.

Cornell CS465 Fall 2007 • Lecture 20

© 2007 Doug James & Steve Marschner • 42

Perceptual dimensions of color

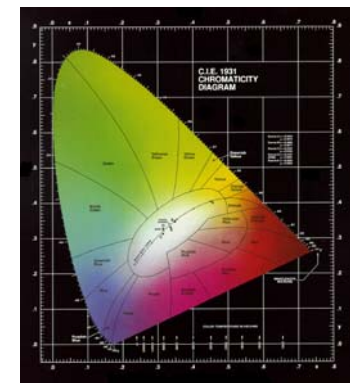
- Hue
 - the “kind” of color, regardless of attributes
 - colorimetric correlate: dominant wavelength
 - artist's correlate: the chosen pigment color
- Saturation
 - the “colorfulness”
 - colorimetric correlate: purity
 - artist's correlate: fraction of paint from the colored tube
- Lightness (or value)
 - the overall amount of light
 - colorimetric correlate: luminance
 - artist's correlate: tints are lighter, shades are darker

Cornell CS465 Fall 2007 • Lecture 20

© 2007 Doug James & Steve Marschner • 43

Perceptual dimensions: chromaticity

- In x, y, Y (or another luminance/chromaticity space), Y corresponds to lightness
- hue and saturation are then like polar coordinates for chromaticity (starting at white, which way did you go and how far?)



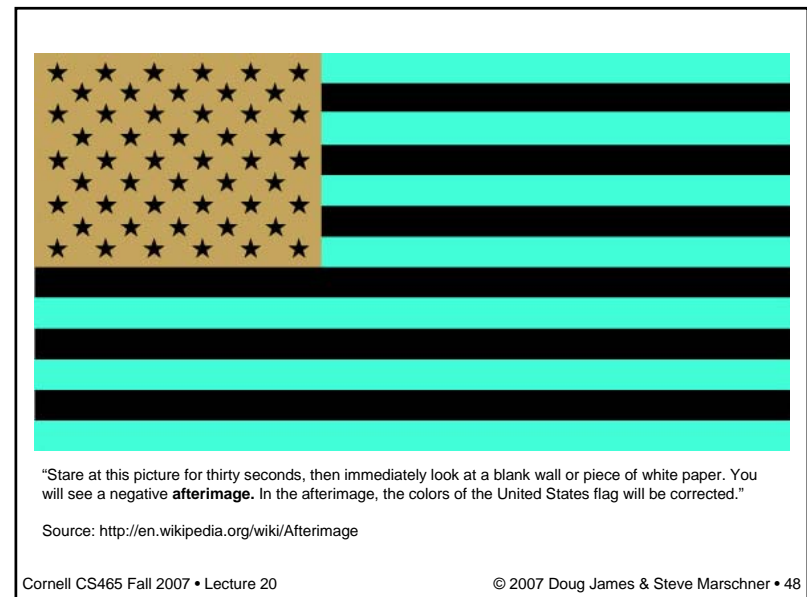

Cornell CS465 Fall 2007 • Lecture 20

© 2007 Doug James & Steve Marschner • 44

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)


March 4, 2002





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