CS4620/5620: Lecture 37

Ray Tracing, Color, Compositing

© 2012 Kavita Bala • 1
(with previous instructors James/Marschner)

Announcements

- Prelim on Thu in B17 at 7:30pm

© 2012 Kavita Bala • 2
(with previous instructors James/Marschner)

Implementing a bvol hierarchy

- A BoundedSurface can contain a list of Surfaces
- Some of those Surfaces might be more BoundedSurfaces
- Voilà! A bounding volume hierarchy
  - And it's all still transparent to the renderer

© 2012 Kavita Bala •
(with previous instructors James/Marschner)

BVH construction example

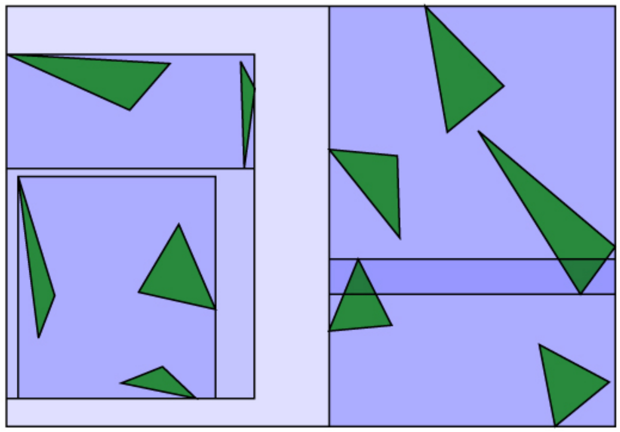

© 2012 Kavita Bala •
(with previous instructors James/Marschner)

Building a hierarchy

- Can do it top down or bottom up

- Top down
  – Make bbox for whole scene, then split into parts
    - Recurse on parts
    - Stop when there are just a few objects in your box
      – Or if you are too deep (say max depth = 24)

© 2012 Kavita Bala •
(with previous instructors James/Marschner)

Construction Possibility 1

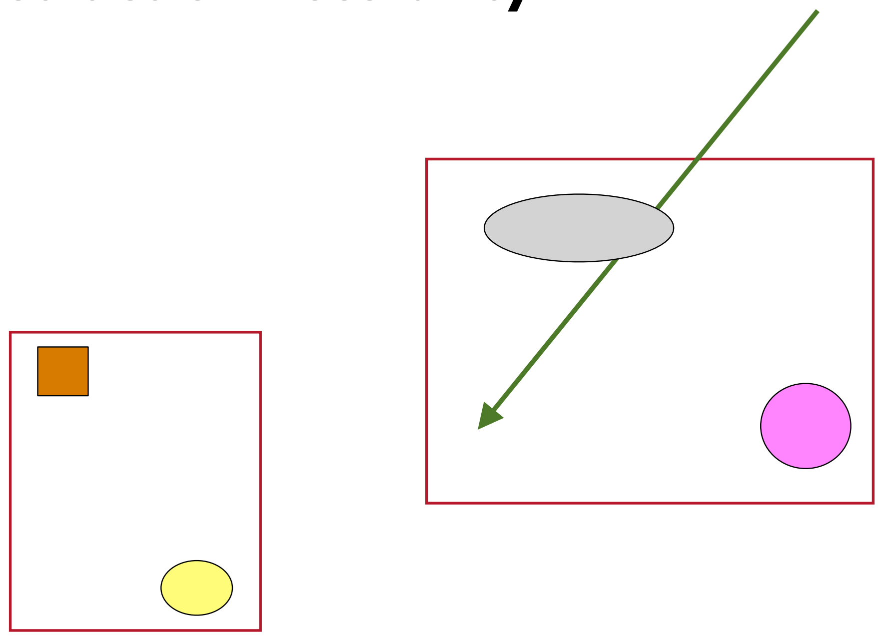

© 2012 Kavita Bala • 7
(with previous instructors James/Marschner)

Construction Possibility 2

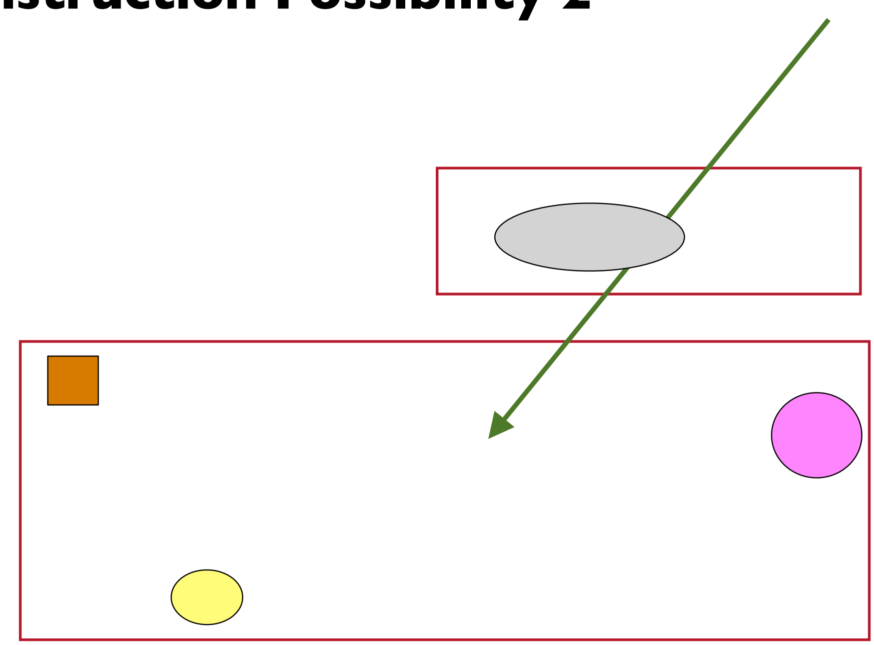

© 2012 Kavita Bala • 8
(with previous instructors James/Marschner)

Building a hierarchy

- How to partition?
  - Practical: partition along axis
    - Center partition
      - Simple
      - Unbalanced tree
    - Median partition
      - More expensive
      - More balanced tree
- Objects that cross the median partition
  - Pick one of the sides to put the object on
  - Expand the bbox to cover that object

© 2012 Kavita Bala •
(with previous instructors James/Marschner)

Hierarchical Data Structures

- From O(N) to O(log N)
  - Cluster objects hierarchically
  - Single intersection might eliminate cluster
- Bounding volume hierarchy
- Space subdivision
  - Octree
  - Kd-tree
  - Uniform

© 2012 Kavita Bala •
(with previous instructors James/Marschner)

Regular space subdivision

- An entirely different approach: uniform grid of cells

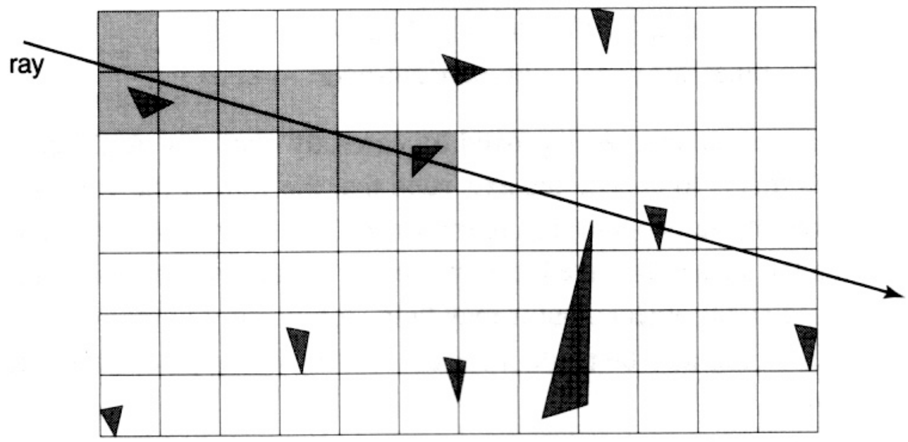

© 2012 Kavita Bala • 11
(with previous instructors James/Marschner)

Non-regular space subdivision

- *k*-d Tree
  - subdivides space, like grid
  - adaptive, like BVH

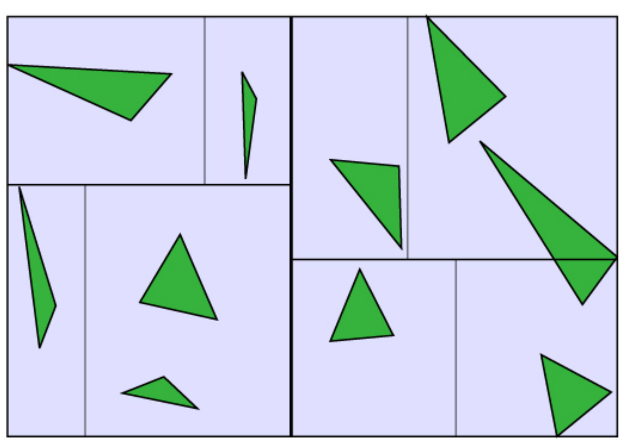

© 2012 Kavita Bala • 12
(with previous instructors James/Marschner)

Implementing acceleration structures

- Conceptually simple to build acceleration structure into scene structure
- Better engineering decision to separate them

© 2012 Kavita Bala • 13
(with previous instructors James/Marschner)

Topics we have not covered

- Compositing
- Color
- Displays
- 5625 topics
  - Focus on rendering and modeling

© 2012 Kavita Bala • 14
(with previous instructors James/Marschner)

Compositing



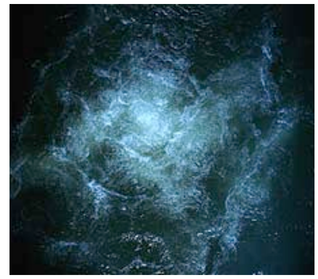
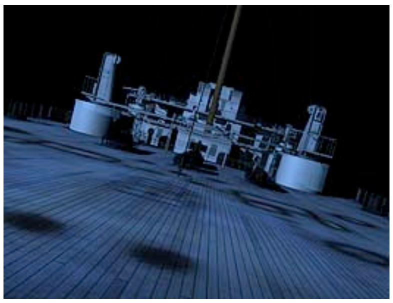


[*Titanic* ; DigitalDomain; vfxhq.com]

Combining images

- Often useful combine elements of several images
- Trivial example: video crossfade
  - smooth transition from one scene to another

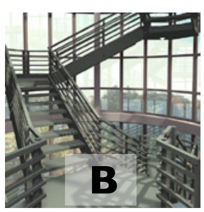

$$r_C = t r_A + (1-t) r_B$$
$$g_C = t g_A + (1-t) g_B$$
$$b_C = t b_A + (1-t) b_B$$

  - note: weights sum to 1.0
    - no unexpected brightening or darkening
    - no out-of-range results
  - this is *linear interpolation*

Foreground and background

- In many cases just adding is not enough
- Example: compositing in film production
  - shoot foreground and background separately
  - also include CG elements
  - how should we do it digitally?

© 2012 Kavita Bala • 17
(with previous instructors James/Marschner)

Foreground and background

- How we compute new image varies with position

[Chuang et al. / Corel]

- Therefore, need to store some kind of tag to say what parts of the image are of interest

© 2012 Kavita Bala • 18
(with previous instructors James/Marschner)

Binary image mask

- First idea: store one bit per pixel
  - answers question "is this pixel part of the foreground?"

[Chuang et al. / Corel]

  - causes jaggies similar to point-sampled rasterization
  - same problem, same solution: intermediate values

© 2012 Kavita Bala • 19
(with previous instructors James/Marschner)

Partial pixel coverage

- The problem: pixels near boundary are not strictly foreground or background

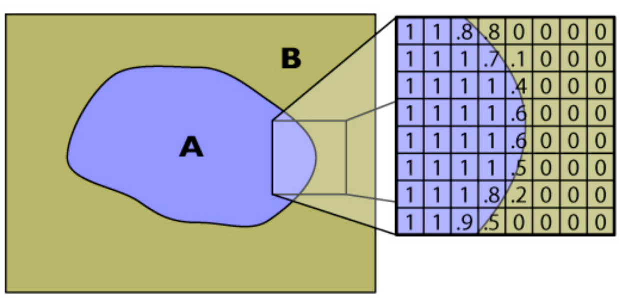

  - how to represent this simply?
  - interpolate boundary pixels between the fg. and bg. colors

© 2012 Kavita Bala • 20
(with previous instructors James/Marschner)

Datatypes for raster images

- For color or grayscale, add *alpha* channel
  - describes transparency of images

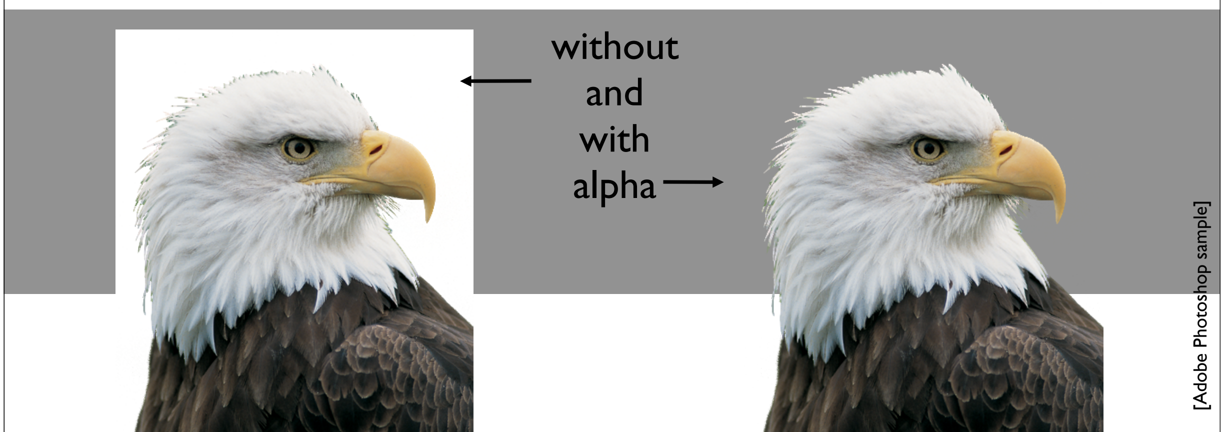

© 2012 Kavita Bala • 21
(with previous instructors James/Marschner)

Alpha compositing

- Formalized in 1984 by Porter & Duff
- Store fraction of pixel covered, called α

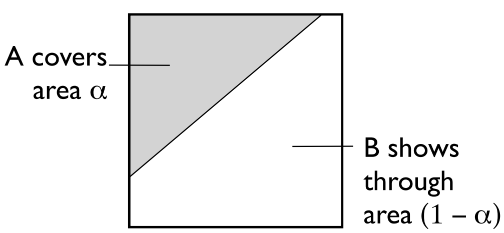

$$C = A \textbf{ over } B$$

$$\begin{aligned} r_C &= \alpha_A r_A + (1 - \alpha_A) r_B \\ g_C &= \alpha_A g_A + (1 - \alpha_A) g_B \\ b_C &= \alpha_A b_A + (1 - \alpha_A) b_B \end{aligned}$$

  - this exactly like a spatially varying crossfade
- Convenient implementation
  - 8 more bits makes 32
  - 2 multiplies + 1 add per pixel for compositing

© 2012 Kavita Bala • 22
(with previous instructors James/Marschner)

Alpha compositing—example

[Chuang et al. / Corel]

Color Science

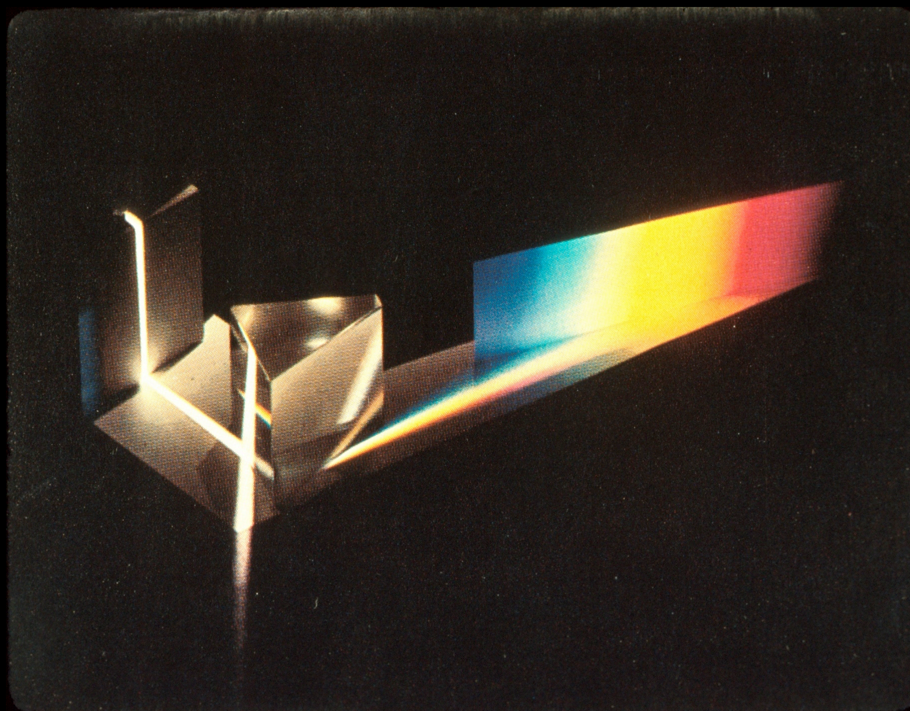

[source unknown]

What light is

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

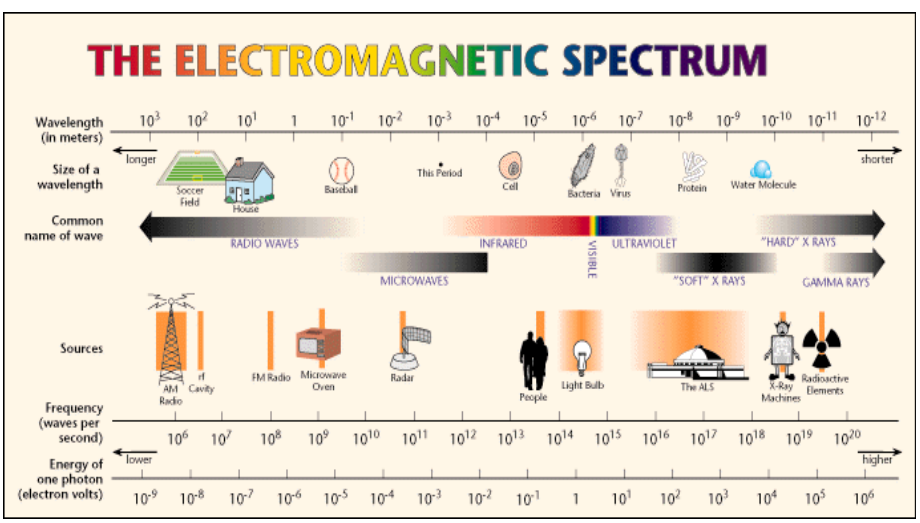

[Lawrence Berkeley Lab / MicroWorlds]

© 2012 Kavita Bala • 26
(with previous instructors James/Marschner)

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)

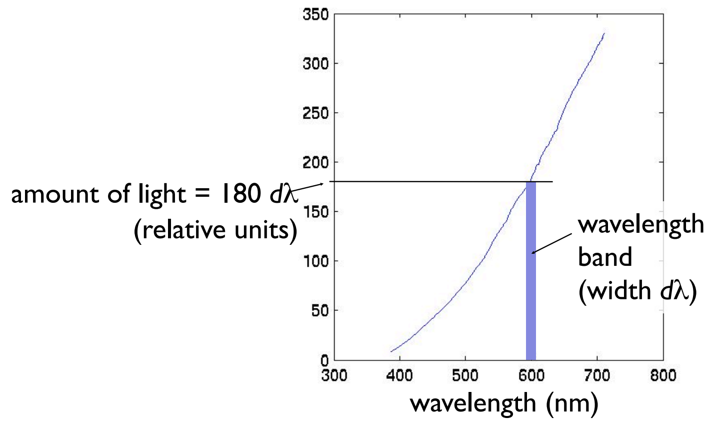

© 2012 Kavita Bala • 27
(with previous instructors James/Marschner)

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"

- Roughly speaking, things appear "colored" when they depend on wavelength and "gray" when they do not.

© 2012 Kavita Bala • 28
(with previous instructors James/Marschner)

The problem of color science

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

[Stone 2003]

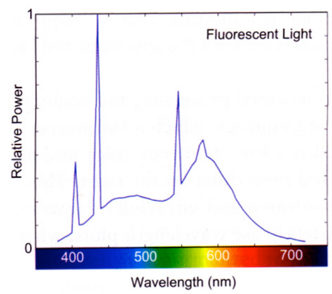

Physical

?

Perceptual

© 2012 Kavita Bala • 29
(with previous instructors James/Marschner)

The eye as a measurement device

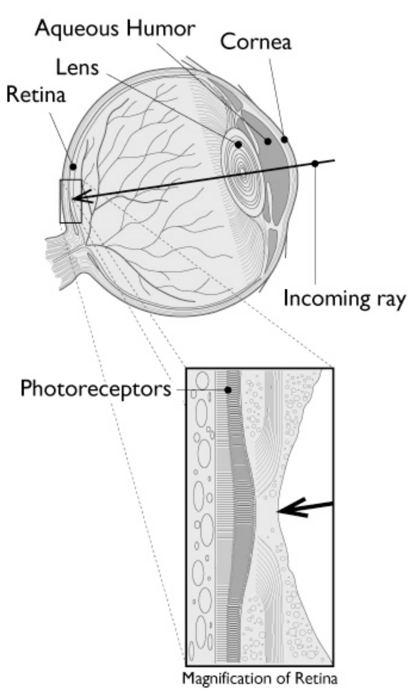

[Greger et al. 1995]

- 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

© 2012 Kavita Bala • 30
(with previous instructors James/Marschner)

Cone Responses

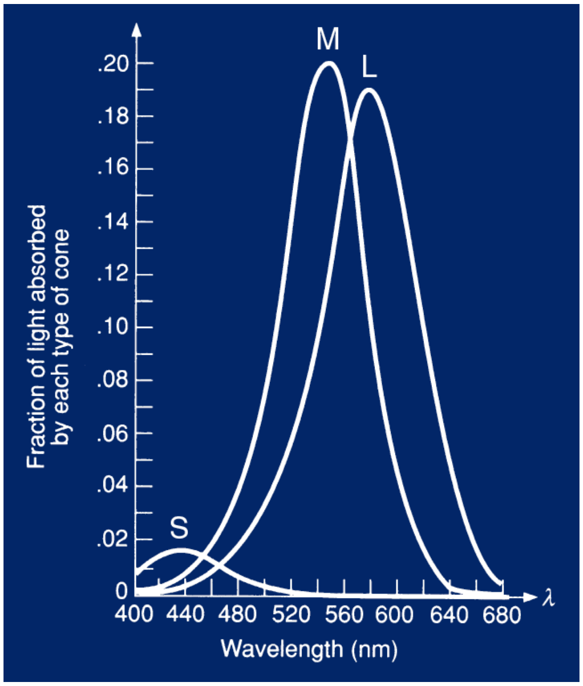

- S,M,L cones have broadband spectral sensitivity
- Results in a trichromatic visual system
- S, M, and L are *tristimulus values*

[source unknown]

© 2012 Kavita Bala

(with previous instructors James/Marschner)

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

[Stone 2003]

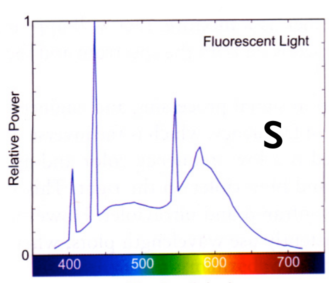

Physical

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

Perceptual

© 2012 Kavita Bala

(with previous instructors James/Marschner)

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*

© 2012 Kavita Bala • 33
(with previous instructors James/Marschner)

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)

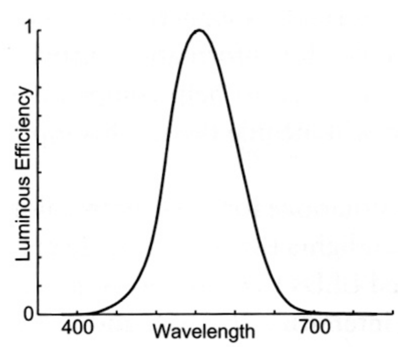

[Stone 2003]

© 2012 Kavita Bala • 34
(with previous instructors James/Marschner)

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"

© 2012 Kavita Bala • 35
(with previous instructors James/Marschner)

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*

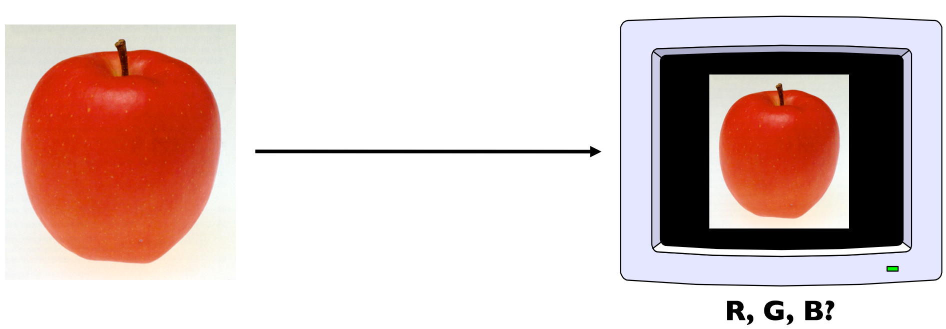

[cs417—Greenberg]

© 2012 Kavita Bala • 36
(with previous instructors James/Marschner)

Color displays

- Operating principle: humans are trichromatic
  - match any color with blend of three
  - therefore, problem reduces to producing 3 images and blending
- Additive color
  - blend images by sum
  - R, G, B make good primaries

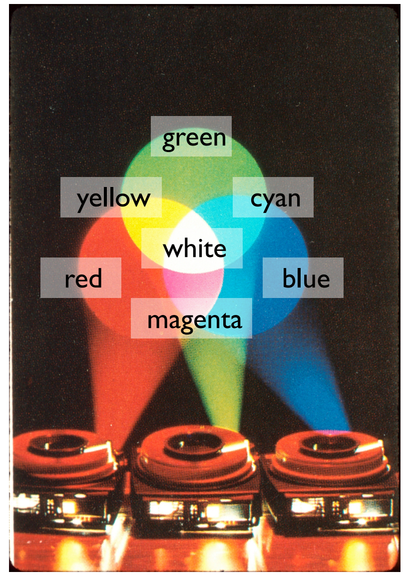

[cs417 S02 slides]

Color displays

- CRT: phosphor dot pattern to produce finely interleaved color images

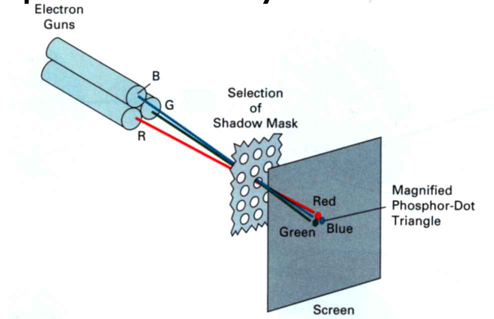

[H&B fig. 2-10]

- LCD: interleaved R,G,B pixels

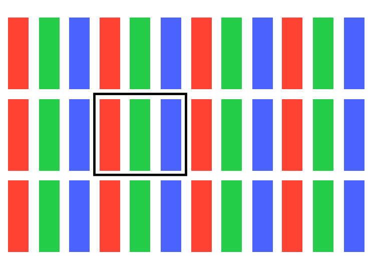

Subtractive Color

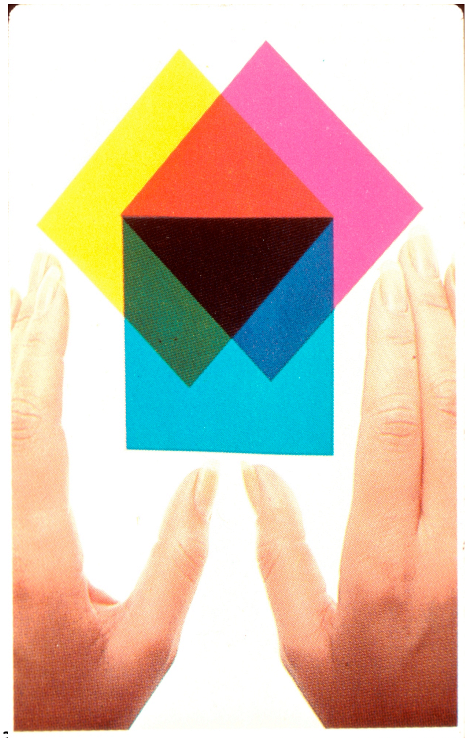

[source unknown]

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

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 = $R\mathbf{R} + G\mathbf{G} + B\mathbf{B}$ 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

© 2012 Kavita Bala • 41
(with previous instructors James/Marschner)

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

© 2012 Kavita Bala • 42
(with previous instructors James/Marschner)

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 = $X\mathbf{X} + Y\mathbf{Y} + Z\mathbf{Z}$)

  – 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

© 2012 Kavita Bala • 43
(with previous instructors James/Marschner)

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

© 2012 Kavita Bala • 44
(with previous instructors James/Marschner)

Chromaticity Diagram

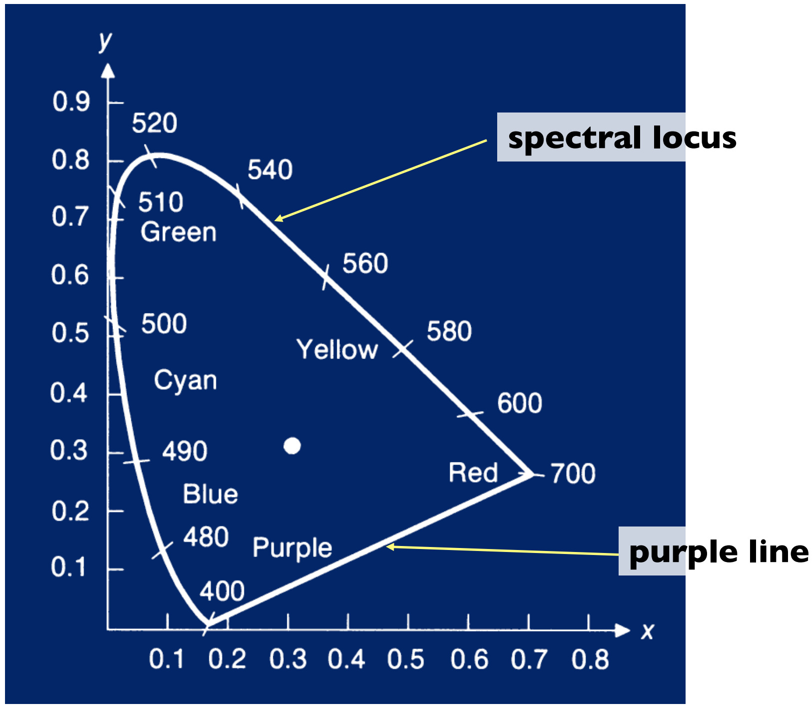

[source unknown]

© 2012 Kavita Bala

(with previous instructors James/Marschner)

Chromaticity Diagram

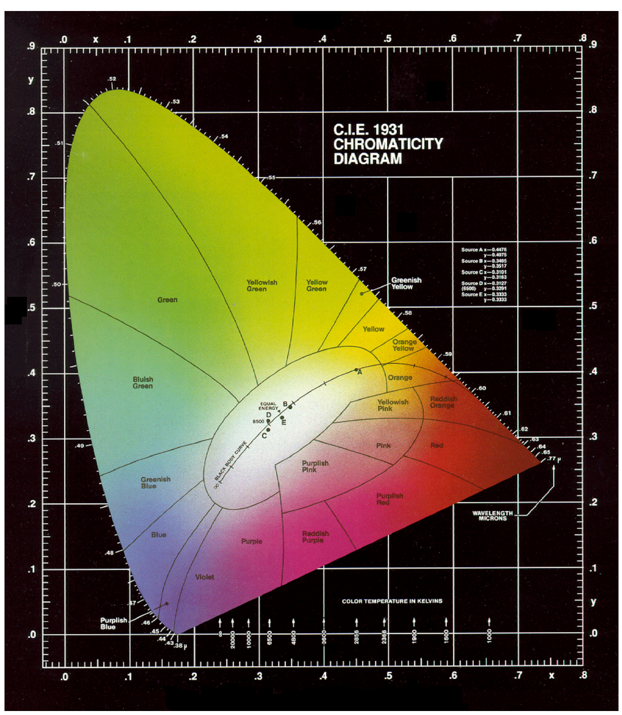

[source unknown]

© 2012 Kavita Bala

(with previous instructors James/Marschner)

Color Gamuts

[source unknown]

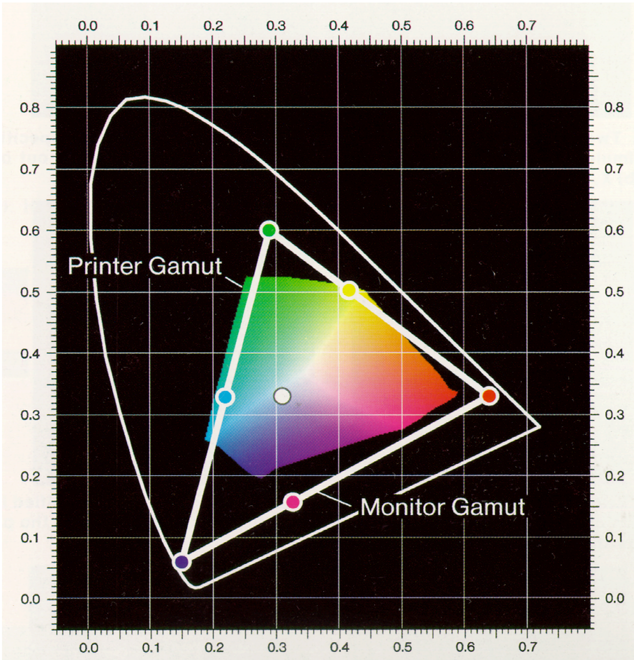

Monitors/printers can't produce all visible colors

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.

© 2012 Kavita Bala • 47
(with previous instructors James/Marschner)

Perceptually organized color spaces

- Artists often refer to colors as *tints*, *shades*, and *tones* of pure pigments
- Hue, Saturation, Lightness (HSV)
- Luv, Lab (separate luminance)

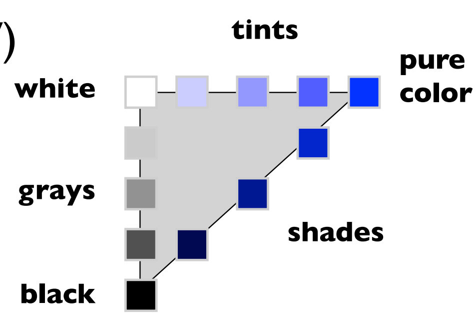

[after FvDFH]

© 2012 Kavita Bala • 48
(with previous instructors James/Marschner)