Cameras and Imaging
http://www.ugrad.cs.jhu.edu/~cs461

Professor Hager

http://www.cs.jhu.edu/~hager

Unit Outline

- Camera imaging sensor basics
- Storing images
- Light transport
- Color

- Reading:
  - F&P Chap 1.
  - F&P Chap 4.
  - F&P Chaper 6 (mostly 6.1-6.4)
  - Background section of the WIKI

A Word On Computer-Imaging

- Video imaging has gone from an exotic technology to everyday commodity.

- Originally (since ~1930) NTSC standard
  - 480 x 640 YUV
  - Interlaced

- Now, a wide variety of resolutions and quality
  - VGA (= NTSC)
  - SVGA (= 600x800)
  - XVGA (= 768x1024)
  - SXGA (=1024x1280)
  - UGA (= 1200x1600)

How Cameras Produce Images

- Basic process:
  - photons hit a detector
  - the detector becomes charged
  - the charge is read out as brightness

- Sensor types:
  - CCD (charge-coupled device)
    - most common
    - high sensitivity
    - high power
    - cannot be individually addressed
    - blooming
  - CMOS
    - simple to fabricate (cheap)
    - lower sensitivity, lower power
    - can be individually addressed

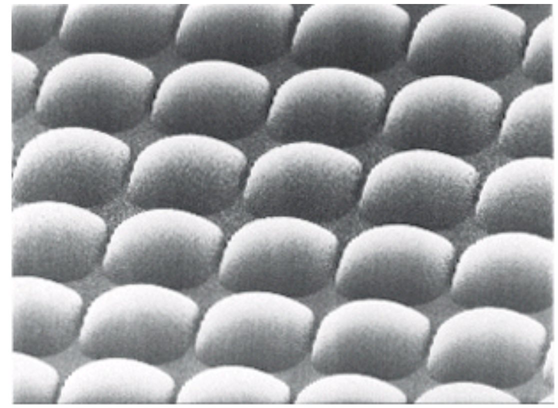

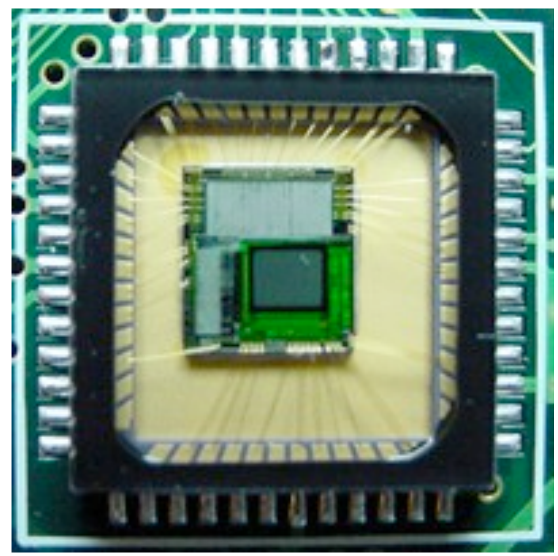

How Color Cameras Work

- 1 CCD cameras
  - A **Bayer** pattern is placed in front of the CCD
  - A **Demosaicing** process reads the pixels in a region and computes color and intensity

- 3 CCD camera use a beam splitter and 3 separate CCDs
  - higher color fidelity
  - requires careful alignment of ccds

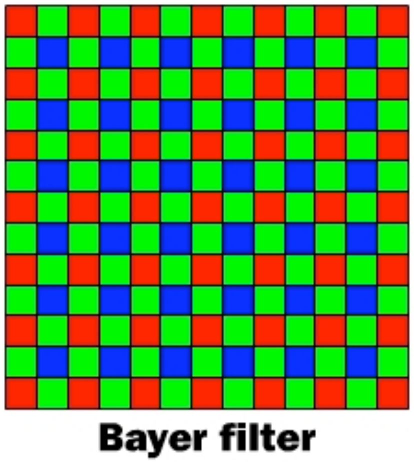

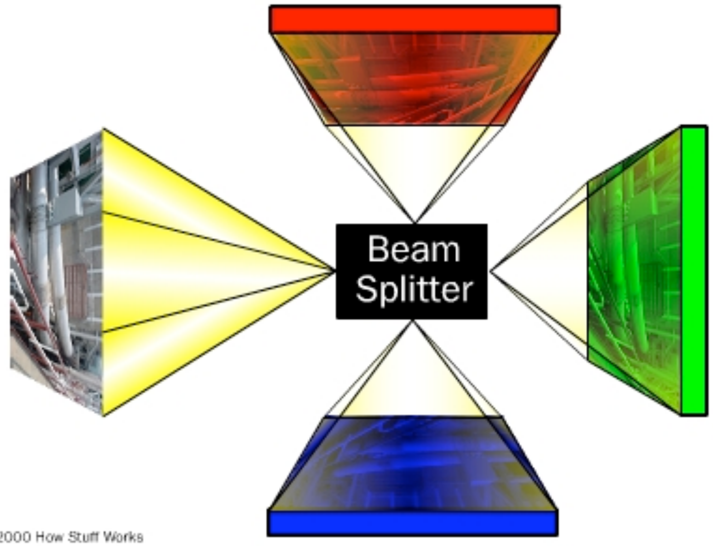

A “Traditional” Camera

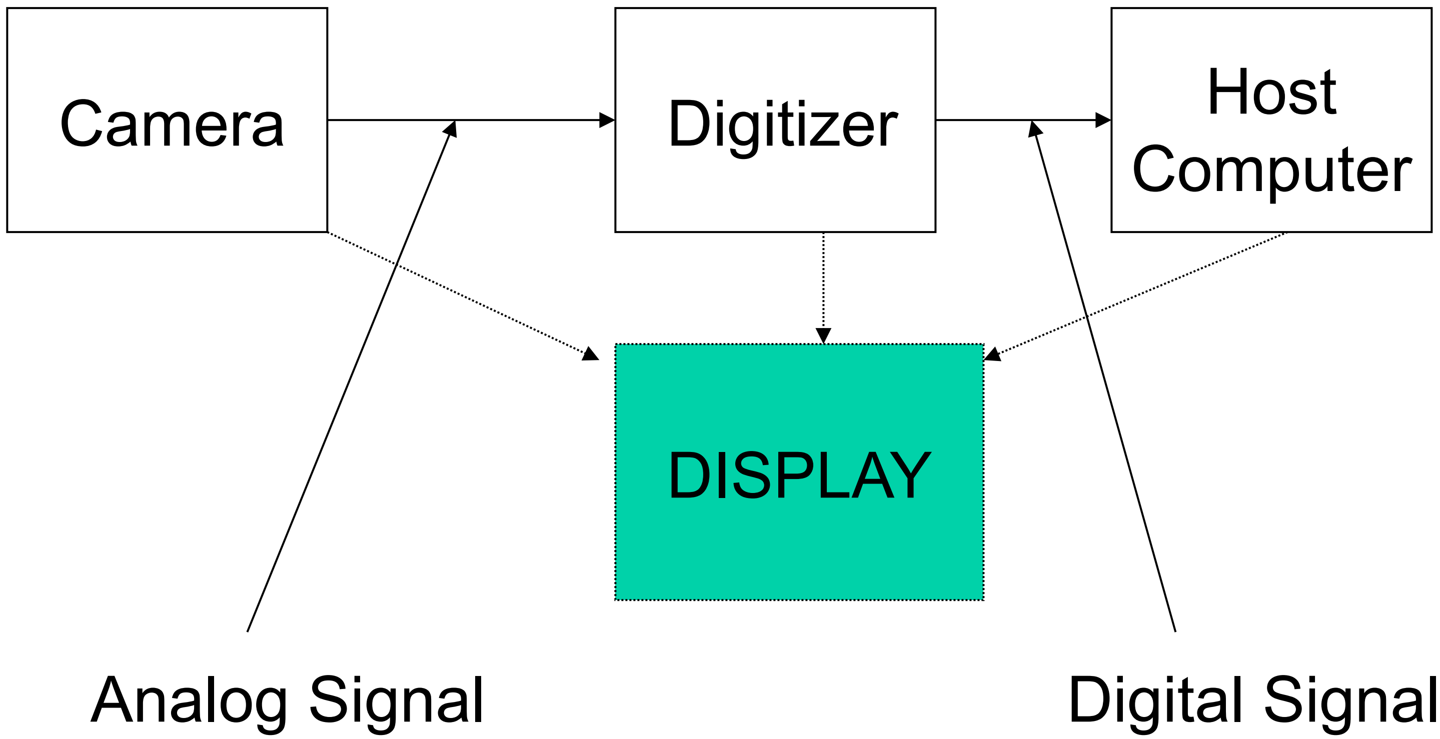

What's under the Hood

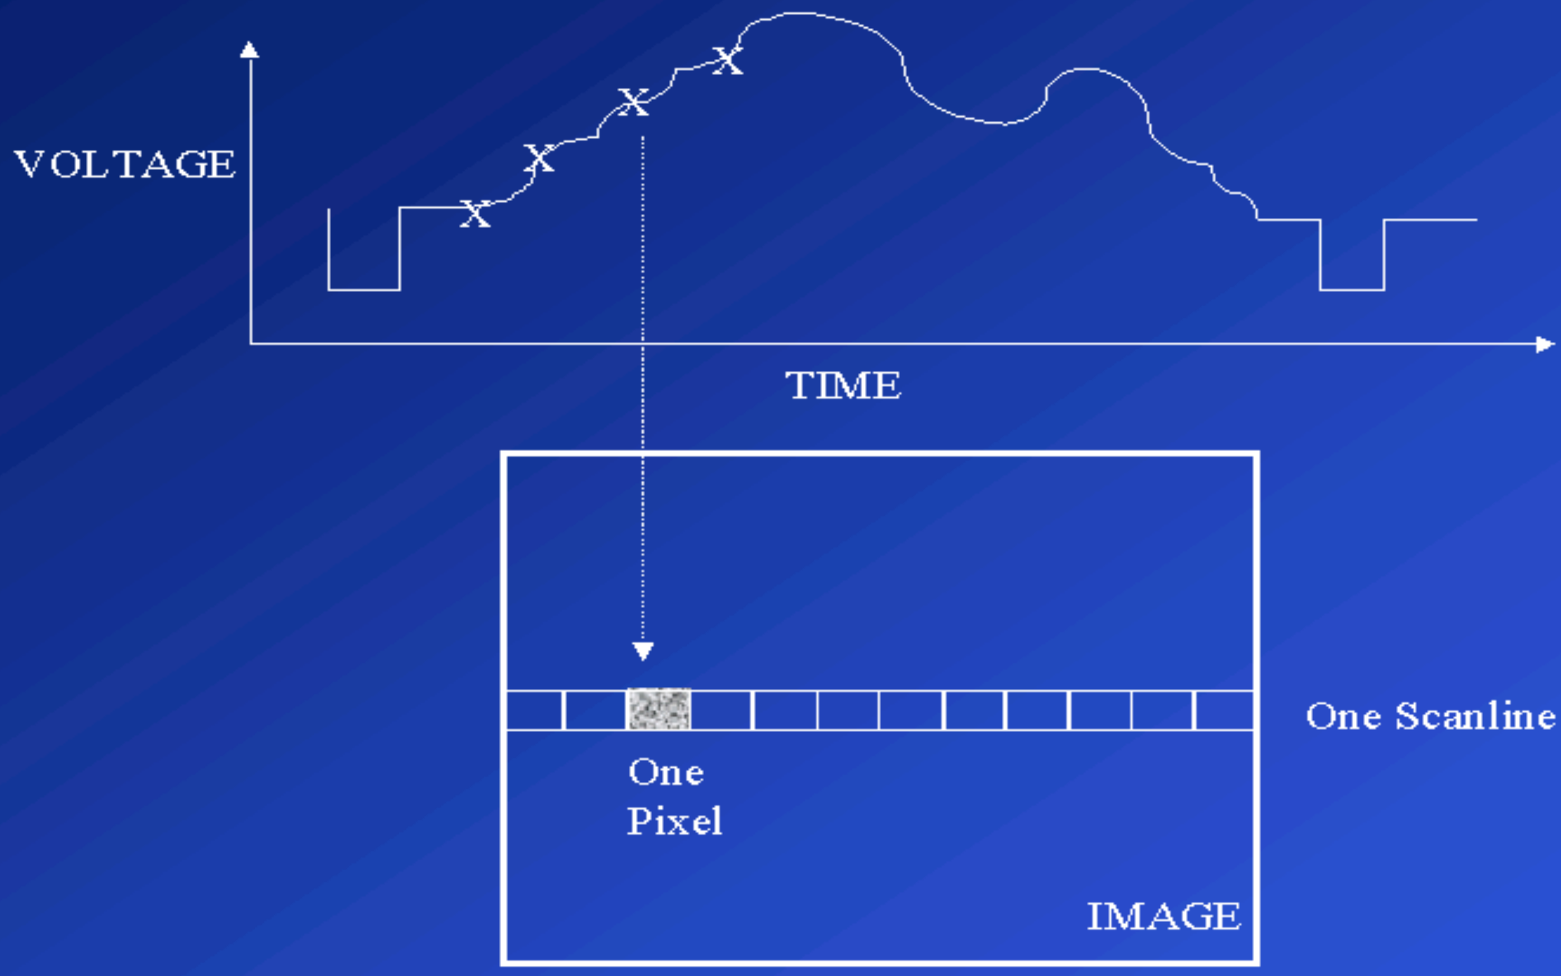

A Modern Digital Camera

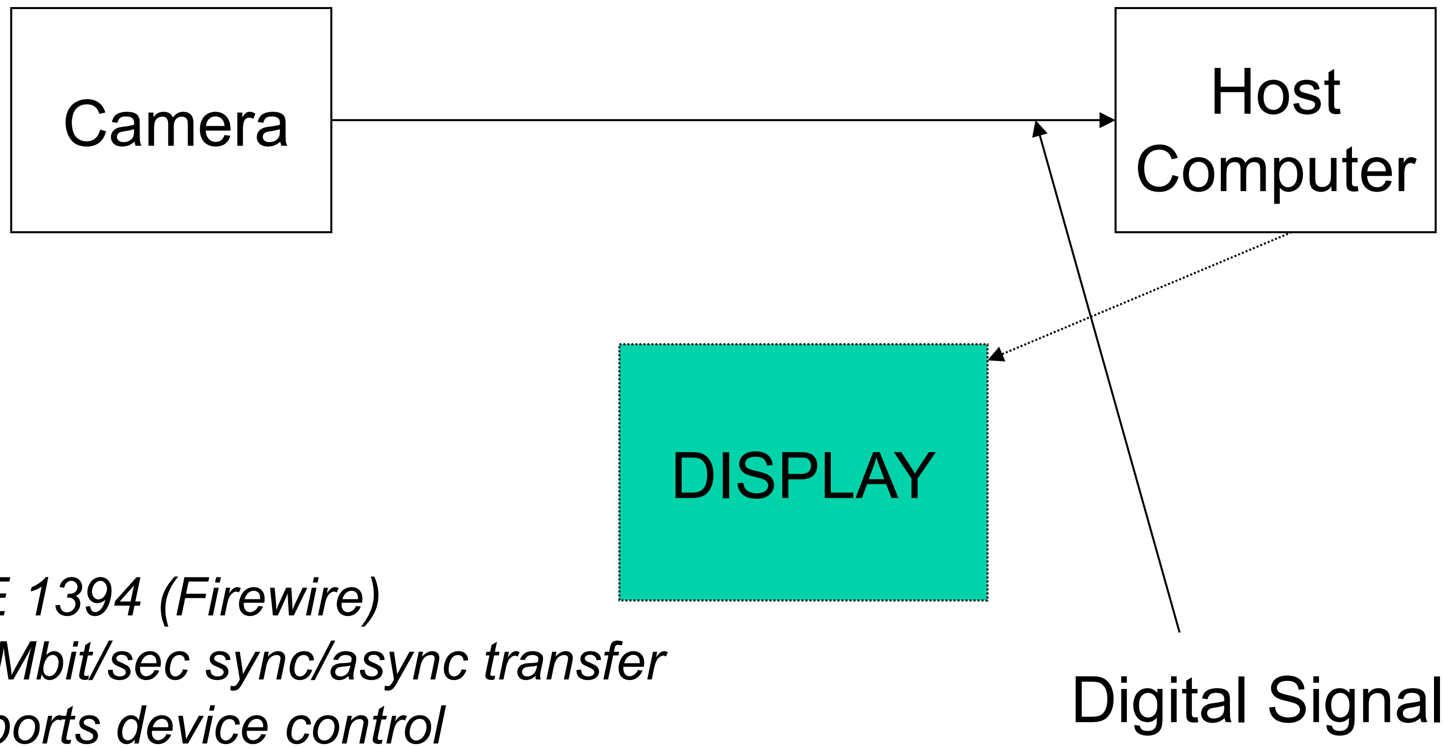

IEEE 1394 (Firewire)
400 Mbit/sec sync/async transfer
Supports device control

USB 2.0
480 Mbit/sec (~280Mbit/sec in practice)
Less flexible, but simpler to implement

Other Issues

- Automatic Gain Control (AGC): adjusting amplification and black level to get a “good fit” of the incident light power to the range of the image
- Shuttering: Electronic “switch” that controls how long the CCD is “exposed.”
- White balance: Adjustment of the mapping from measured spectral quantities to image RGB quantities (we’ll talk about this more when we get to color).

THE ORGANIZATION OF A 2D IMAGE

Pixel

<u>Binary</u>
1 bit

<u>Grey</u>
1 byte

<u>Color</u>
3 bytes

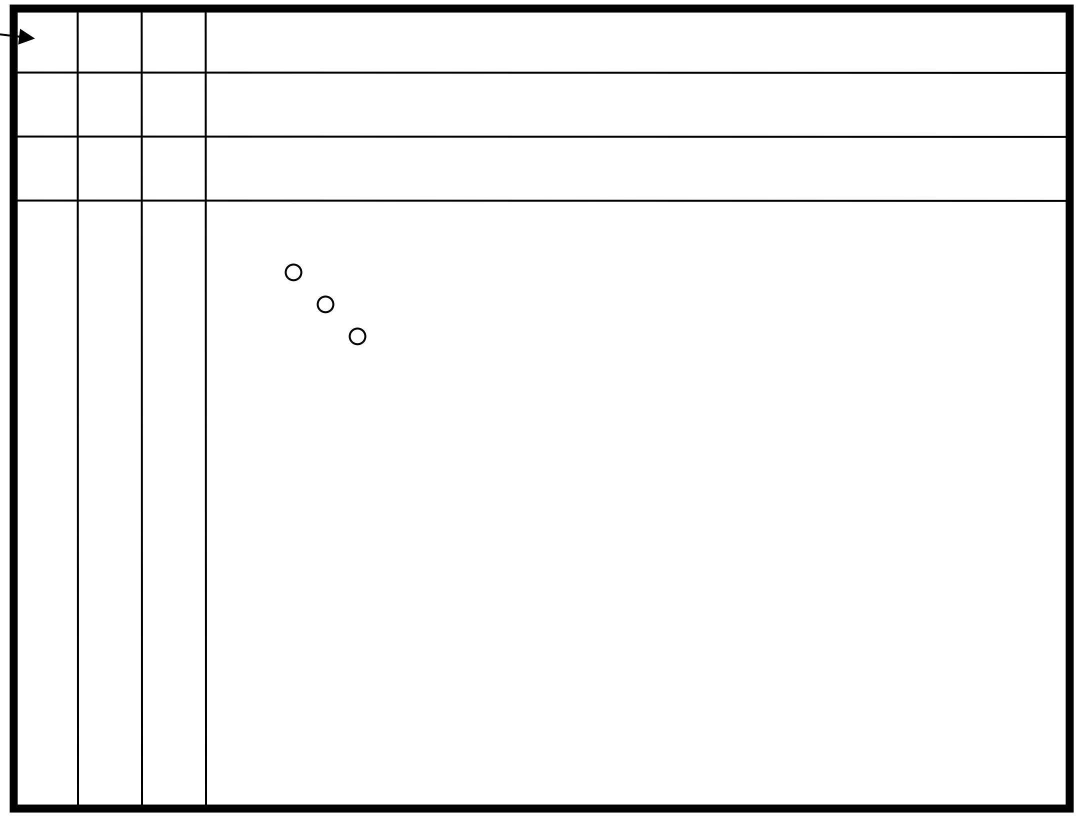

Storing Images

- Non-lossy schemes
  - pbm/pgm/ppm/pnm
    - code for file type, size, number of bands, and maximum brightness
  - tif (lossless and lossy versions)
  - bmp
  - gif (grayscale)

- Lossy schemes
  - gif (color)
  - jpg
    - uses Y Cb Cr color representation; subsamples the color
    - Uses DCT on result
    - Uses the fact the human system is less sensitive to color than spatial detail

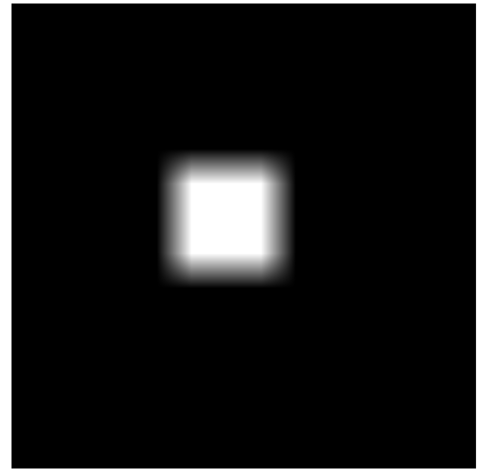

GIF IMAGE FORMAT

- GIF (Graphics Interchange Format)
  - Limited to 8 bits/pixel for both color and gray-scale.

8-bit index	RED	GREEN	BLUE
0	R0	G0	B0
1	R1	G1	B1
2	R2	G2	B2
	⋮	⋮	⋮
254	R254	G254	B254
255	R255	G255	B255

TIFF IMAGE FORMAT

- TIFF (Tagged Image File Format)
  - More general than GIF
  - Allows 24 bits/pixel
  - Supports 5 types of image compression including:
    - RLE (Run length encoding)
    - LZW (Lempel-Ziv-Welch)
    - JPEG (Joint Photographic Experts Group)

The Imaging process

- Light is generated by some source
  - point source
  - extended source
  - white/colored ....
- Light is reflected from some surface
  - matte,
  - mirrorlike
  - colored/light/dark ....
- Light is sensed by some instrument
  - sensitivity
  - field of view
  - gray scale/color/....

Basic Radiometry

- Core Questions:
  - how “bright” will surfaces be?
  - what is “brightness”?
    - measuring light
    - interactions between light and surfaces
- Core idea - think about light arriving at a surface
- around any point is a hemisphere of directions
- Simplest problems can be dealt with by reasoning about this hemisphere

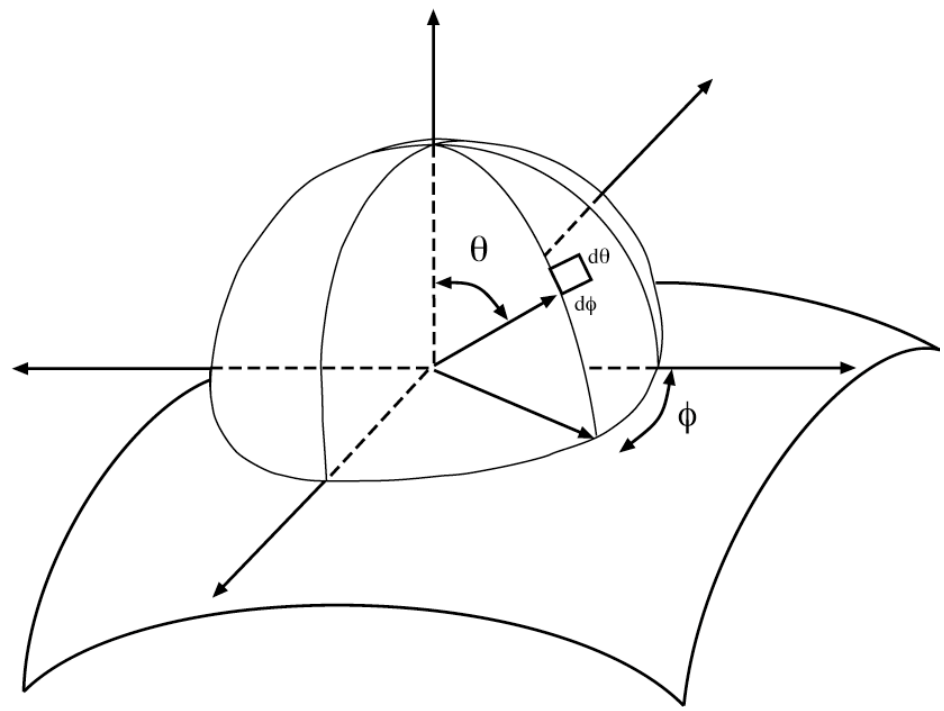

Solid Angle

- Defined by analogy with angle (in radians)
  - solid angle subtended by a patch is the area covered by the patch on a unit sphere
- The solid angle subtended by an infinitesimal patch area dA is given by

$$d\omega = \frac{dA\cos\theta}{r^2}$$

- Another useful expression:

$$d\omega = \sin\theta\left(d\theta\right)\left(d\phi\right)$$

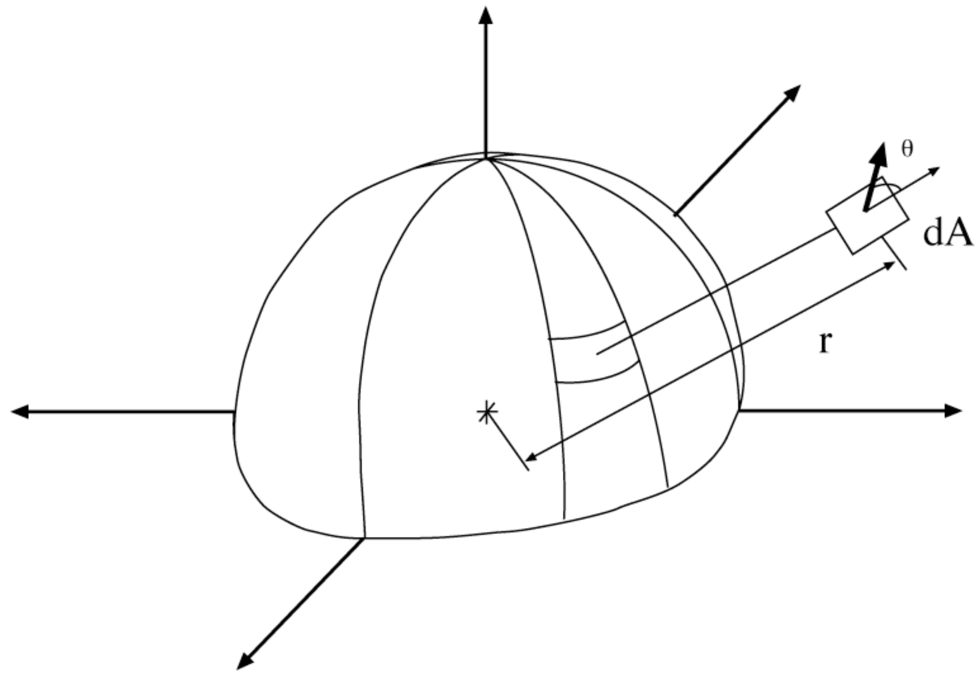

Radiance

- Measure the "amount of light" at a point, in a direction
- Definition: *Radiant power per unit* ***foreshortened*** *area (area perpendicular to the direction of travel) per unit solid angle*
- Units: watts per square meter per steradian ($wm^{-2}sr^{-1}$)
- Usually written as:
- Crucial property: In a vacuum, radiance leaving p in the direction of q is the same as radiance arriving at q from p

point in space

direction of travel

$$L\left(\underline{x}, \theta, \varphi\right)$$

Alternatively L(x, n) for direction vector n

What is A Source?

- Anything that generates light
- Characterized by
  - Spectrum
  - Power distribution over sphere (i.e. radiance distribution it creates)
- Common to abstract as a “point source at infinity” or as a completely uniform illuminant although in the real world, this is never the case

REFLECTANCE MODELS

- Description of how light energy incident on an object is transferred from the source to the object to the camera sensor

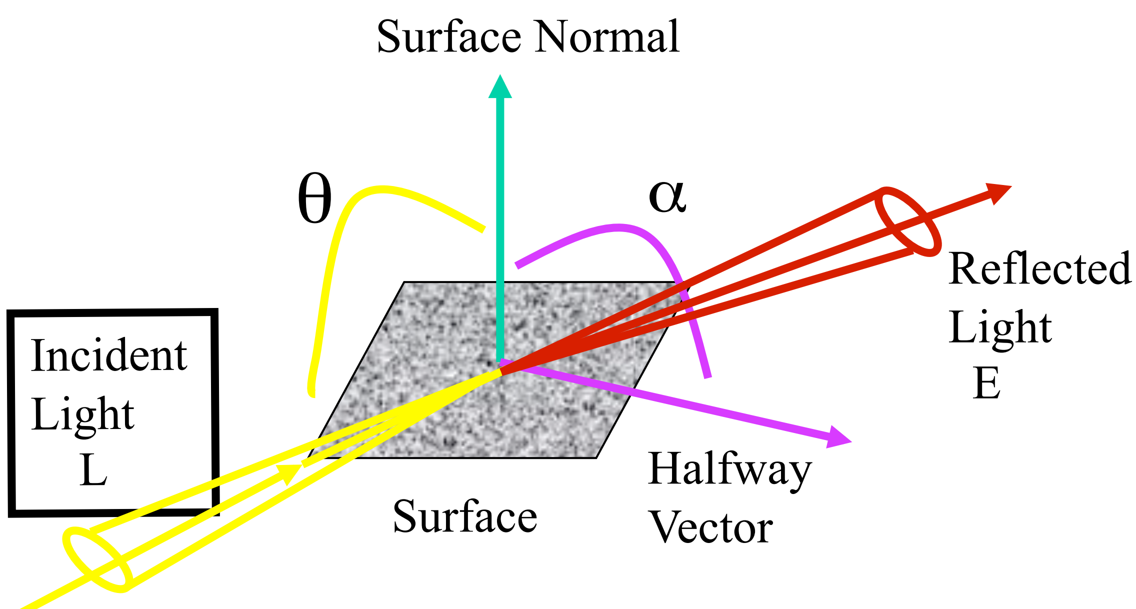

Irradiance

- How much light is arriving at a surface?
- Unit is *Irradiance* = incident power per unit area *not foreshortened*
- This is a function of incoming angle.
- A surface experiencing radiance $L(x,\theta,\phi)$ coming in from $d\omega$ experiences irradiance

this term deals with the fact that radiance is in terms of *foreshortened* area

$$L\left(\underline{x},\theta,\varphi\right)\cos\theta\, d\omega$$

Irradiance

- Total power arriving at the surface is given by adding irradiance over all incoming angles

$$\int_{\Omega} L\left(\underline{x},\theta,\varphi\right)\cos\theta\sin\theta d\theta d\varphi$$

radiance

x foreshortening = irradiance

summed over sphere = power

Example: How Light Goes to the Camera

- Question: How much light reaches a camera sensor?
  - In particular, note that camera pixels are *foreshortened* as we move away from the optical center.

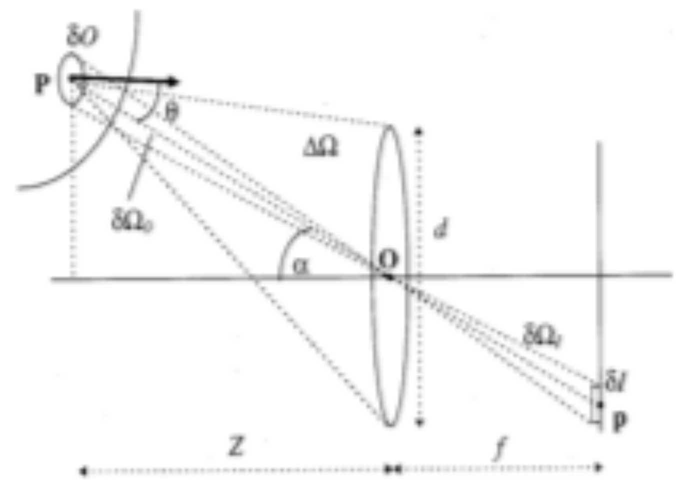

$$E(p) = L(P)\pi / 4(d / f)^2 \cos^4 \alpha$$

  - L(P) --- radiance at P directed toward lens; p point on image sensor
  - f -- focal length
  - d -- aperture
  - α --- angle between optical axis and p-o (optical center)

This implies that we expect objects away from the axis to be darker than those near the axis. E.g. for 25 degrees off center, a drop-off of approximately 33%.

Typical 1/4” sensor (= 6mm) w 12mm focal length = 26.5 degrees

Typically, this effective is dominated by other lens imperfections (vignetting)

Surfaces and the BRDF

- Many effects when light strikes a surface -- could be:
  - absorbed
  - transmitted
  - reflected
  - scattered
- Assume that
  - surfaces don't fluoresce
  - surfaces don't emit light (i.e. are cool)
  - all the light leaving a point is due to that arriving at that point

The BRDF

- Given assumptions, we can model effects at a surface with a record of outgoing vs incoming illumination
  - the Bidirectional Reflectance Distribution Function (BRDF)
- Definition:
  - the ratio of the radiance in the outgoing direction to the incident irradiance
- Units: Inverse steradians

$$\rho_{bd}\left(\underline{x},\vartheta_o,\varphi_o,\vartheta_i,\varphi_i,\right)=\frac{L_o\left(\underline{x},\vartheta_o,\varphi_o\right)}{L_i\left(\underline{x},\vartheta_i,\varphi_i\right)\cos\vartheta_i d\omega}$$

or equivalently

$$L_o\left(\underline{x},\vartheta_o,\varphi_o\right)=\rho_{bd}\left(\underline{x},\vartheta_o,\varphi_o,\vartheta_i,\varphi_i,\right)L_i\left(\underline{x},\vartheta_i,\varphi_i\right)\cos\vartheta_i d\omega$$

$$\boxed{L(x,n_o)=p_{bd}(x,n_o,n_i)L(x,n_i)n\cdot n_i}$$

Some Observations about BRDF

- BRDF is symmetric in incoming and outgoing directions
  - this is referred to as Helmholtz reciprocity; recent work has shown that this property along can be used to compute 3D geometry for surfaces with arbitrary BRDF

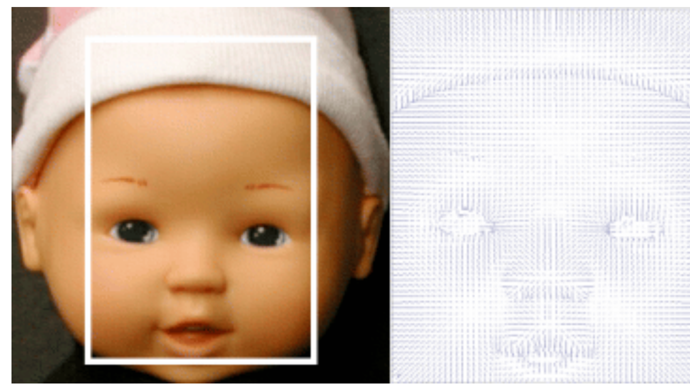

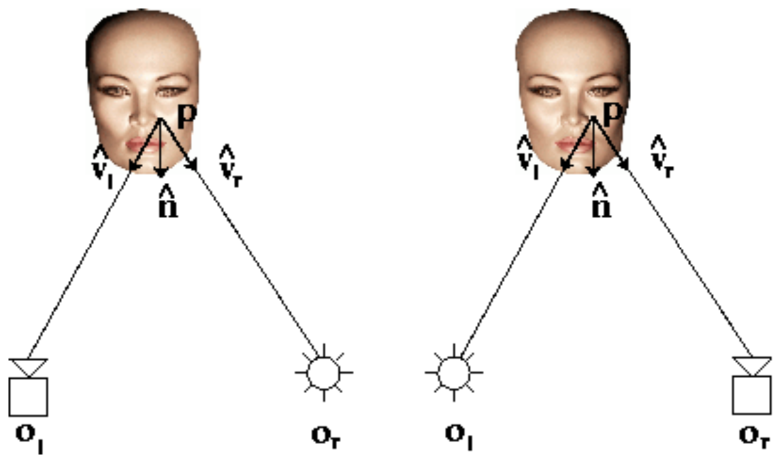

Courtesy Todd Zickler

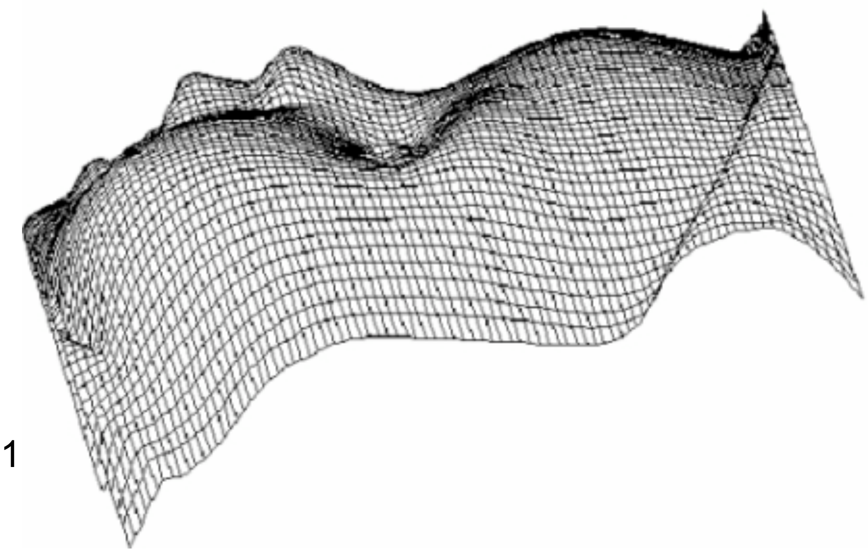

Some Observations about BRDF

- BRDF is symmetric in incoming and outgoing directions
  - this is referred to as Helmholtz reciprocity; recent work has shown that this property along can be used to compute 3D geometry for surfaces with arbitrary BRDF

- In general, BRDFs are difficult to measure, so are usually not known (hence the interest in methods that are independent of BRDF).

- As a result, we often consider special cases of BRDF when analyzing algorithms

Lambertian Surfaces and Albedo

For a Lambertian surface, BRDF is independent of angle, thus:

$$L_o\left(\underline{x}\right) = \rho_{brdf}(x) L_i\left(\underline{x}, \vartheta_i, \varphi_i\right) \cos\vartheta_i d\omega$$

A natural value to use in describing the surface is the *albedo* defined as the fraction of incident irradience that is reflected, whatever the direction. If ρ_d is the albedo, then

$$\rho_{brdf} = \rho_d / \pi$$

Suppose now that we have an ideal distant pinhole source in a given direction. Then we if look at the entire power (i.e. integrate over the sphere we have

$$L_o\left(\underline{x}\right) = \rho(x)_{brdf} L_{source} \cos\vartheta_{source}$$

The Geometry of Lambertian Surfaces

Recalling the Lambertian, point source case....

$$L_o\left(\underline{x}\right) = \rho(x)_{brdf} L_{source} \cos\vartheta_{source}$$

Equivalently, we can say

$L_o(x) = \rho\ (x)\ L_s\ (s \cdot n(x)\ )$ where s is the direction of the source and n is the surface normal at x

Suppose now we pick three different source directions, each of the same brightness to the surface (L_s is constant)

This is a simple version of *photometric stereo*

Specular surfaces

- Another important class of surfaces is specular, or mirror-like.
  - radiation arriving along a direction leaves along the specular direction
  - reflect about normal
  - some fraction is absorbed, some reflected
  - on real surfaces, energy usually goes into a lobe of directions
  - can write a BRDF, but requires the use of funny functions

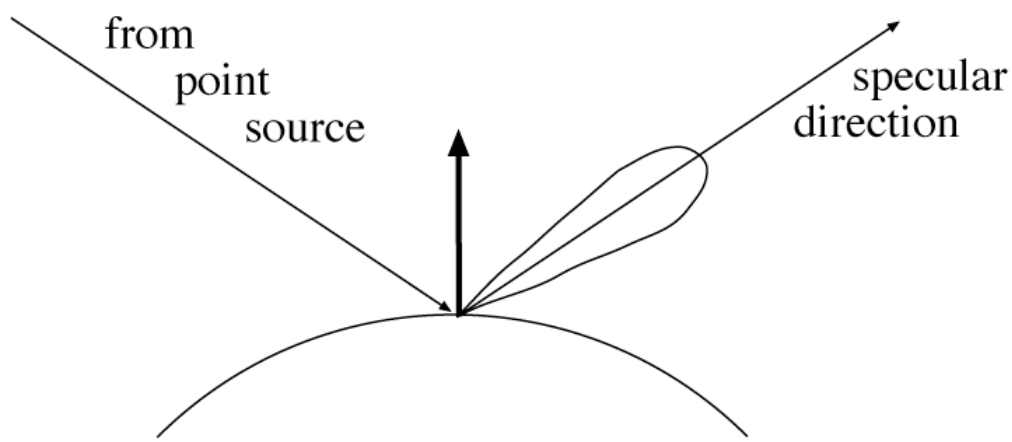

Phong's model

- There are very few cases where the exact shape of the specular lobe matters.
- Typically:
  - very, very small --- mirror
  - small -- blurry mirror
  - bigger -- see only light sources as "specularities"
  - very big -- faint specularities
- Phong's model
  - reflected energy falls off with

$$\cos^n\left(\delta\theta\right)$$

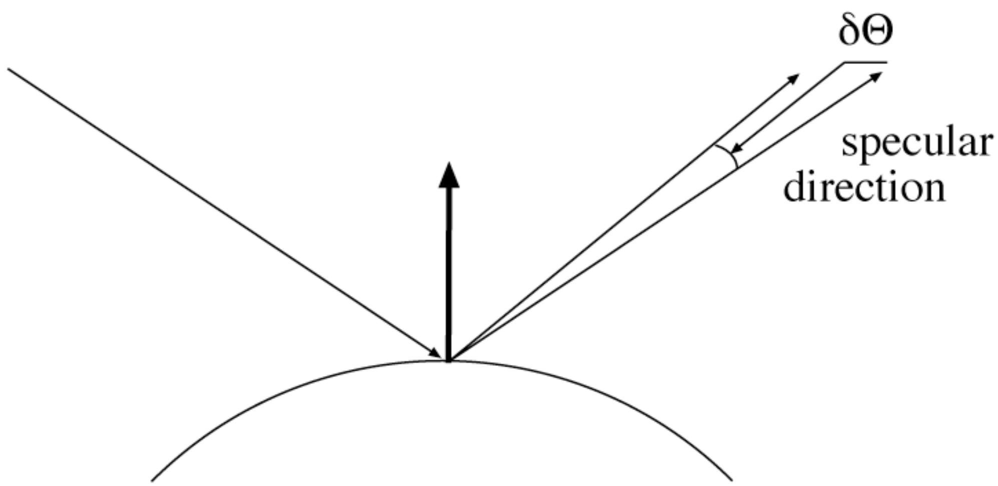

Lambertian + specular model

- Widespread model
  - many surfaces are Lambertian plus specular component
- Advantages
  - easy to manipulate
  - very often quite close true
- Disadvantages
  - some surfaces are not
    - e.g. underside of CD's, feathers of many birds, blue spots on many marine crustaceans and fish, most rough surfaces, oil films (skin!), wet surfaces
  - Generally, very little advantage in modeling behavior of light at a surface in more detail -- it is quite difficult to understand behaviour of L+S surfaces

$$L(x,\theta_0,\phi_0) = \rho_l(x)\cos\theta_i + \rho_s(x)L(x,\theta_s,\phi_s)\cos^n(\theta_s - \theta_o)$$

$$L(x, n_o) = \rho_l(x)\, n \bullet n_i + \rho_s(x)\, L(x, n_s)\, (n_s \bullet n_o)^k$$

Color

What is Color?

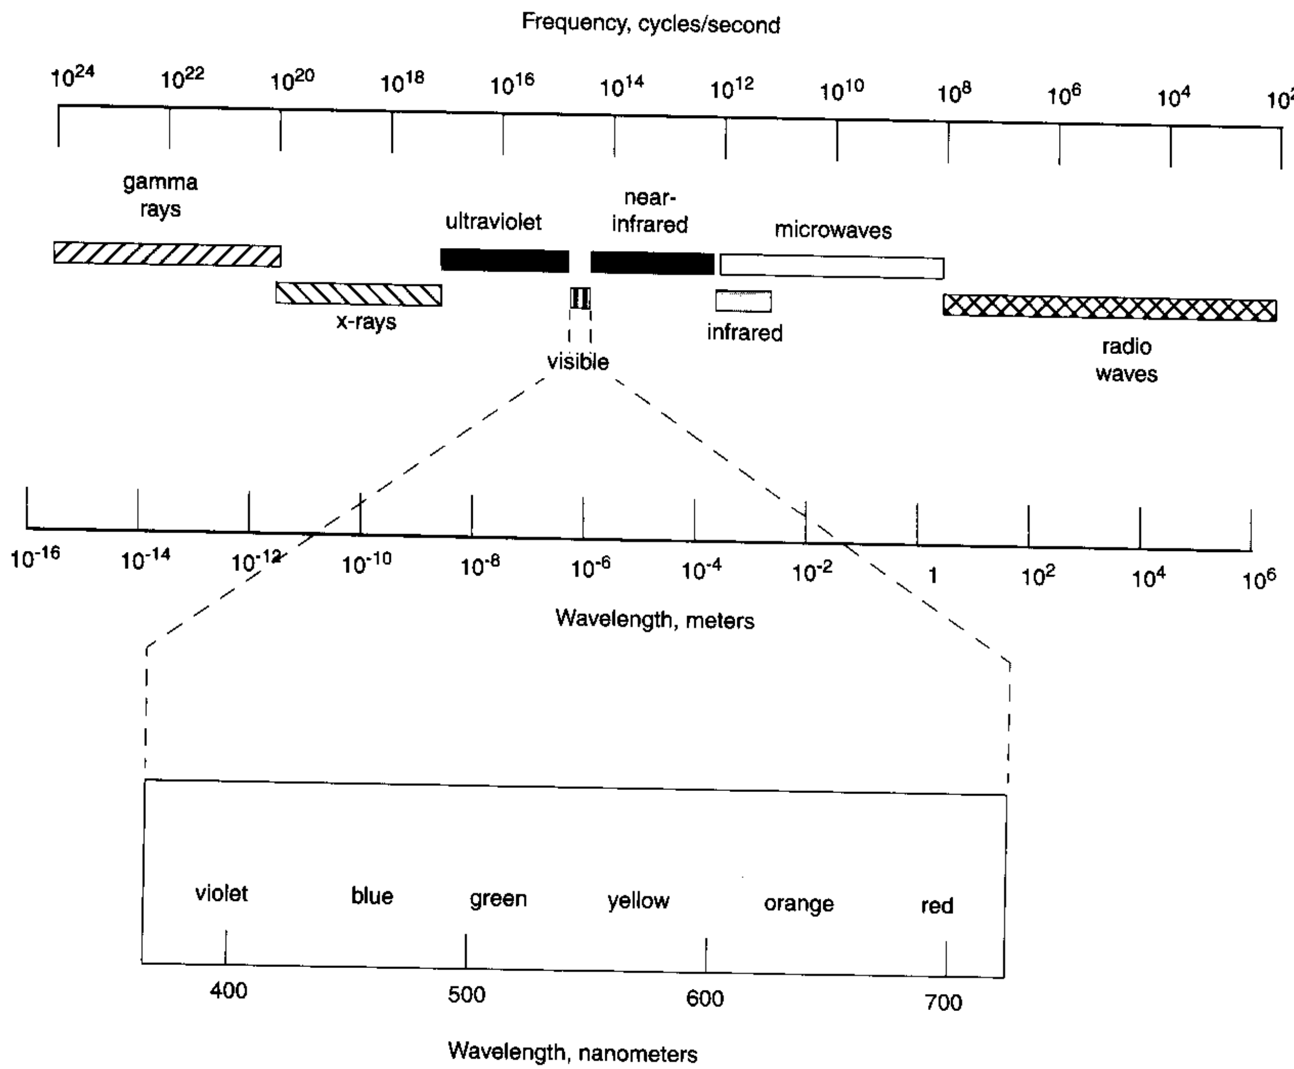

What is Color?

- We almost never see a “pure” wavelength of light; rather a mixture of wavelengths, each with a different “power”
- Only some colors occur as pure wavelengths; many are mixtures of pure colors (e.g. white)

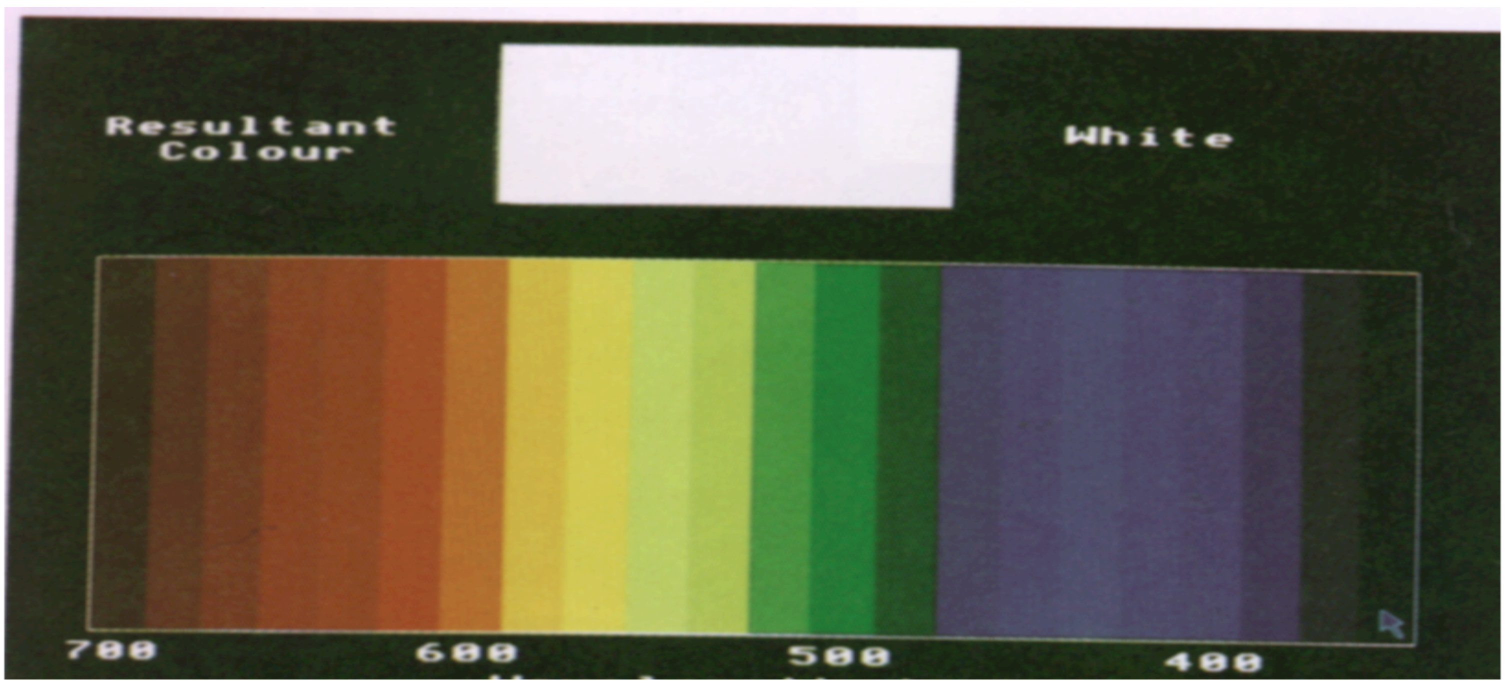

Sunlight

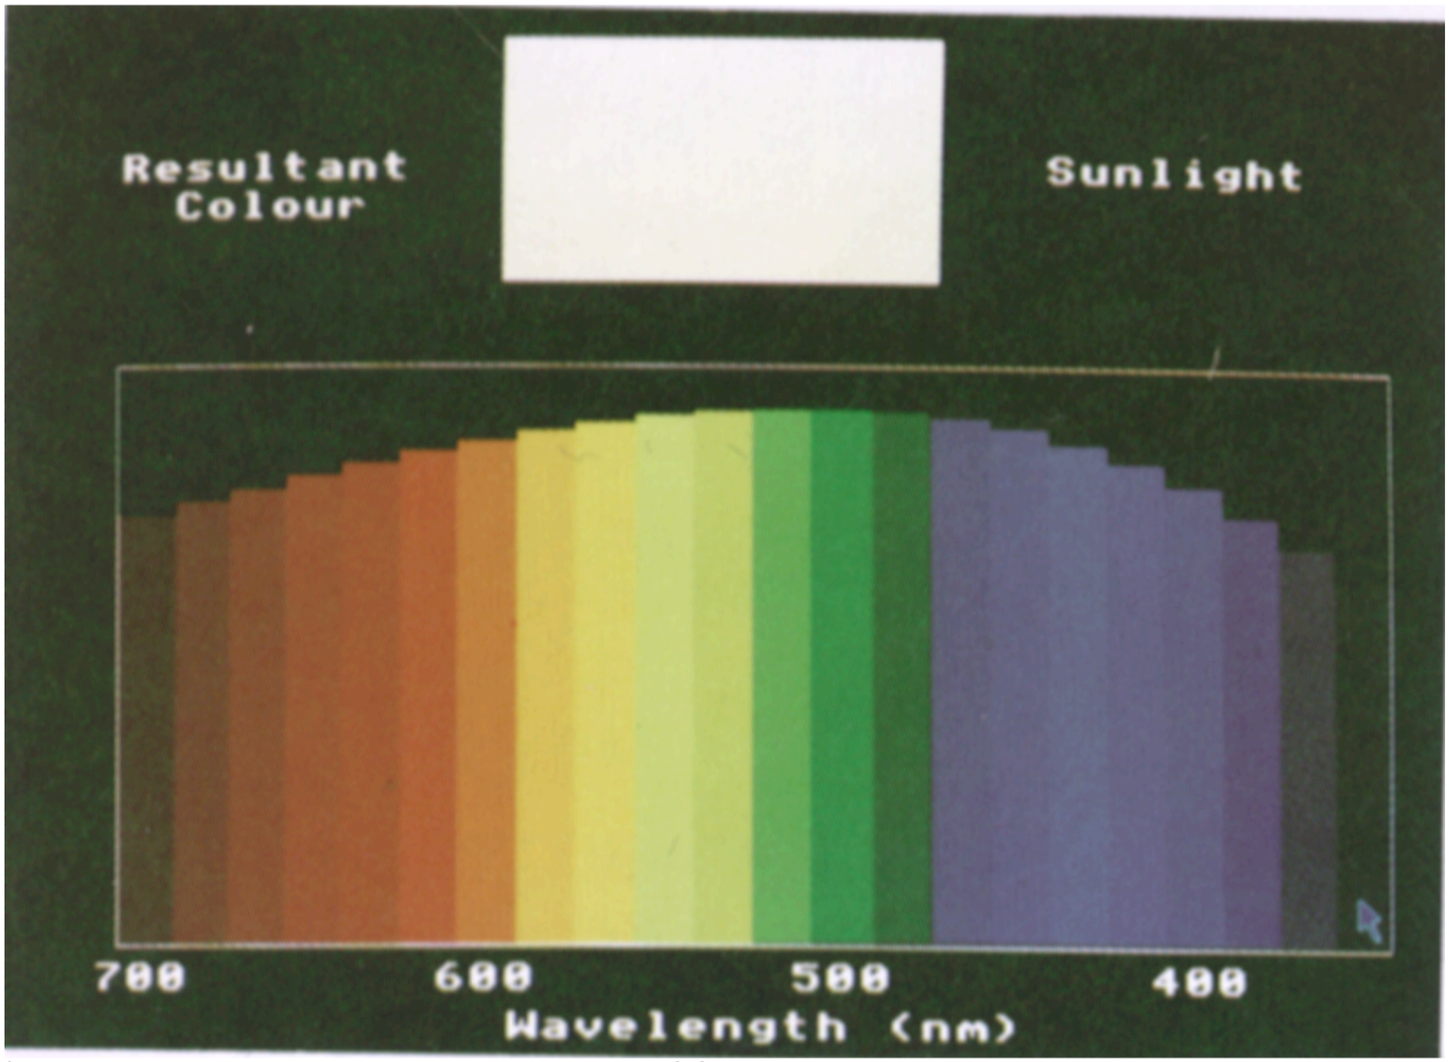

Thermal Imaging for Tracking Humans

Color Measurement

- Let λ denote wavelength
- Let $E(\lambda)$ denote the spectral power at a given wavelength
- Let $\rho_k(\lambda)$ denote the responsiveness of a sensor k to a given wavelength of light
- Then we can compute the "response" r_k of k as

$$r_k = \int \rho(\lambda) E(\lambda) d\lambda$$

EXAMPLE: THE HUMAN EYE

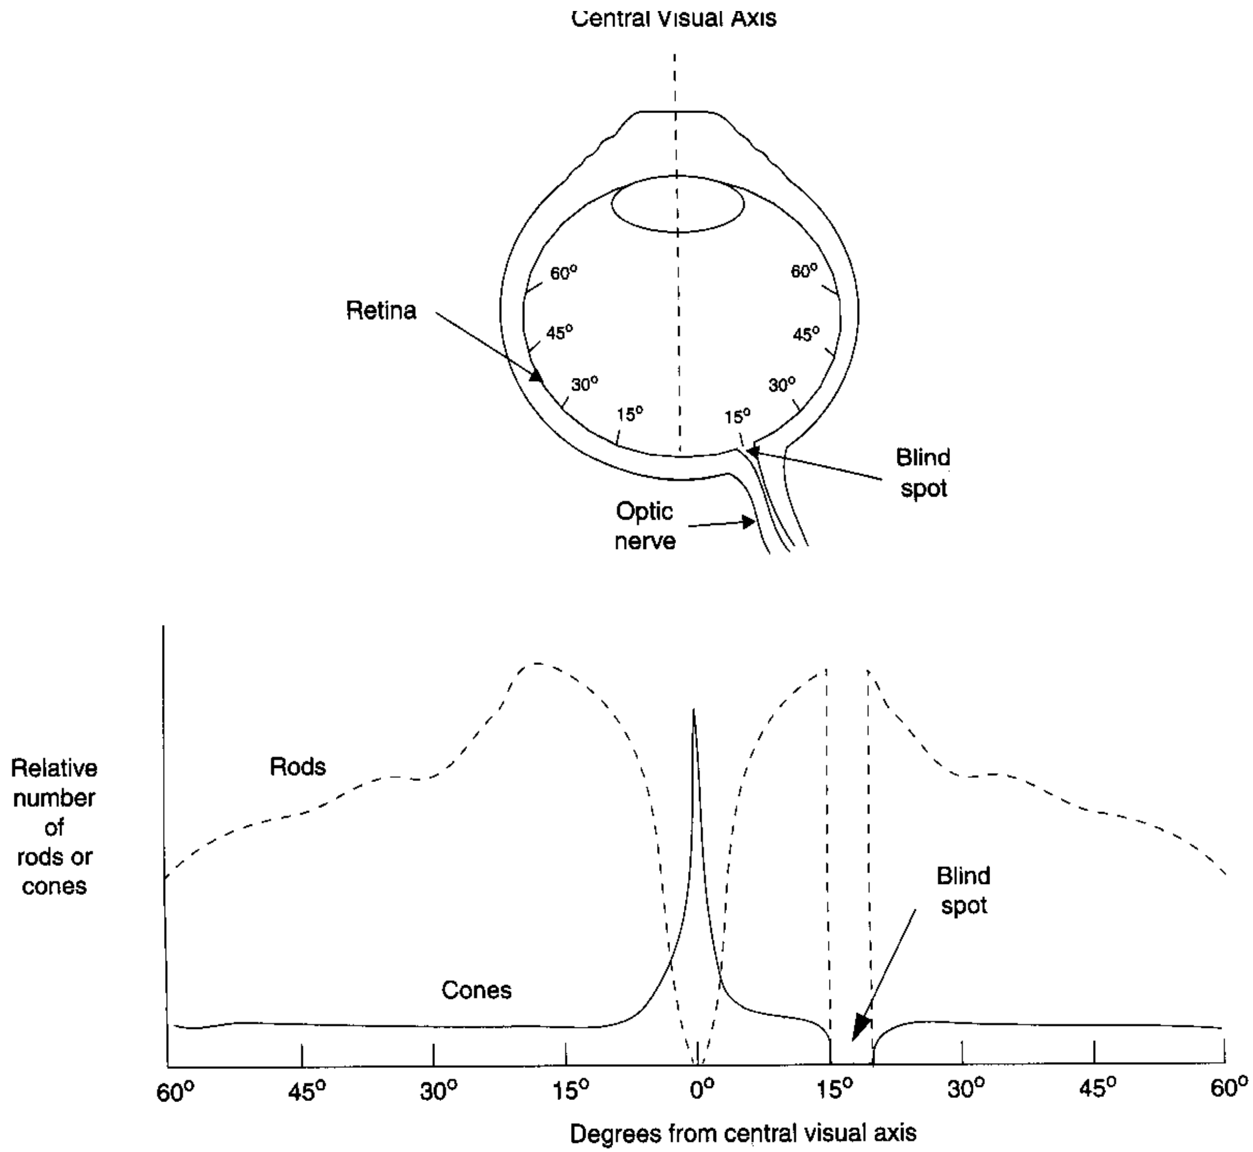

THE HUMAN EYE: RESPONSE

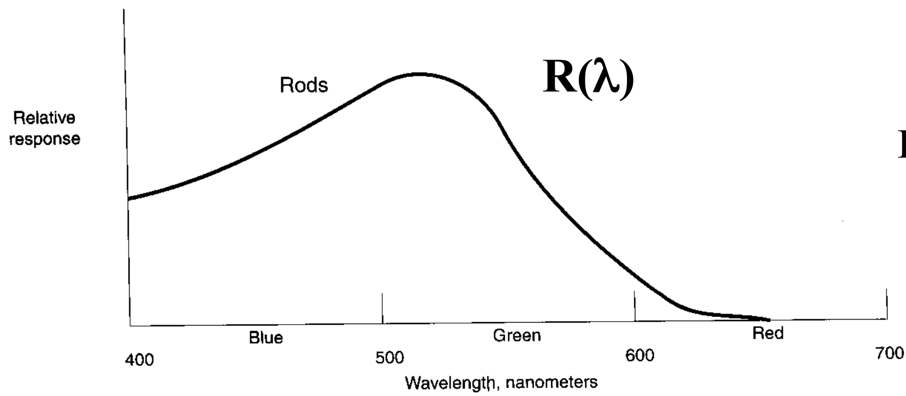

$$\text{BRIGHTNESS} = \int_{\lambda=400nm}^{\lambda=700nm} R(\lambda) I(\lambda) d\lambda$$

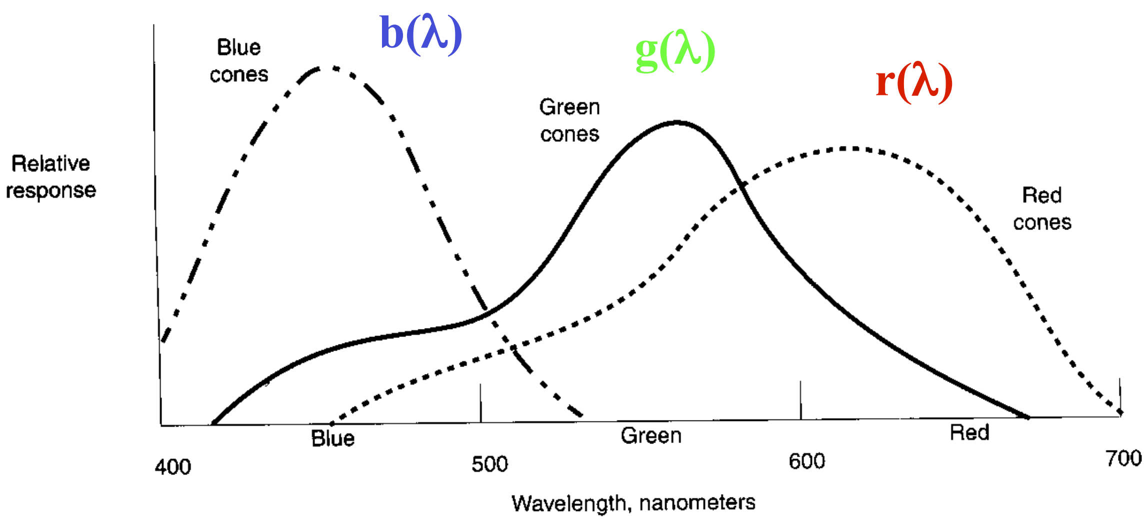

$$\text{RED} = \int_{\lambda=400nm}^{\lambda=700nm} r(\lambda) I(\lambda) d\lambda$$

$$\text{GREEN} = \int_{\lambda=400nm}^{\lambda=700nm} g(\lambda) I(\lambda) d\lambda$$

$$\text{BLUE} = \int_{\lambda=400nm}^{\lambda=700nm} b(\lambda) I(\lambda) d\lambda$$

Color receptors

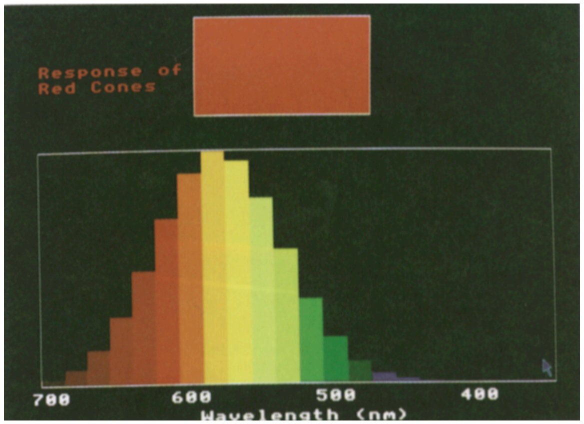

“Red” cone

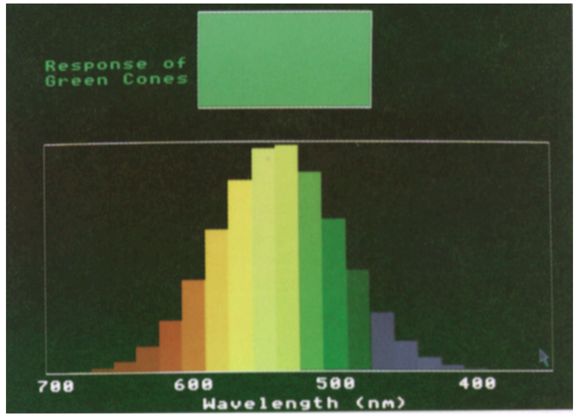

“Green” cone

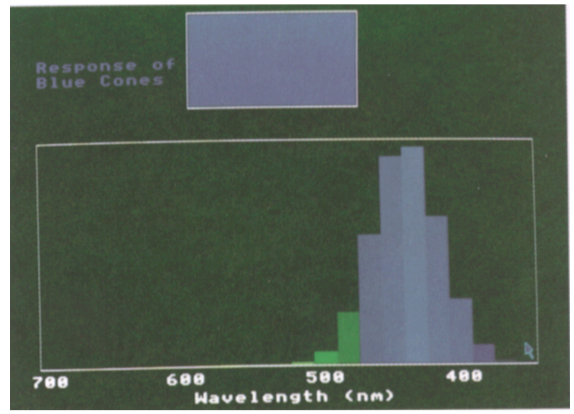

“Blue” cone

Principle of univariance: cones give the same amount of response to different wavelengths -- a single cone cannot distinguish color. Output of cone is obtained by summing probability of absorption over wavelengths.

METAMERISM

- Two different Spectral Energy Distributions with the same RED, GREEN, BLUE response are termed ***metamers***.

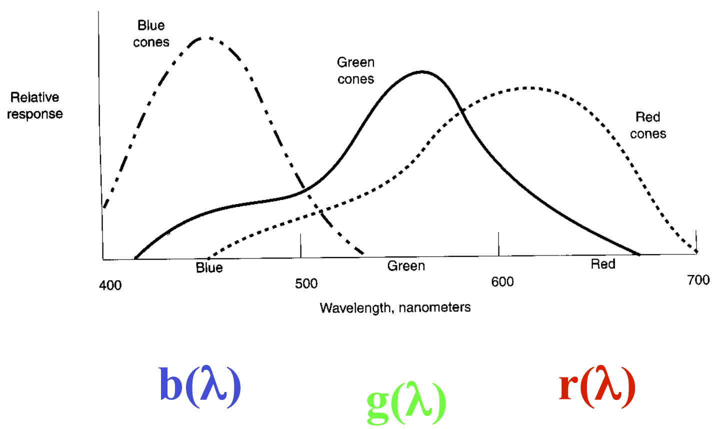

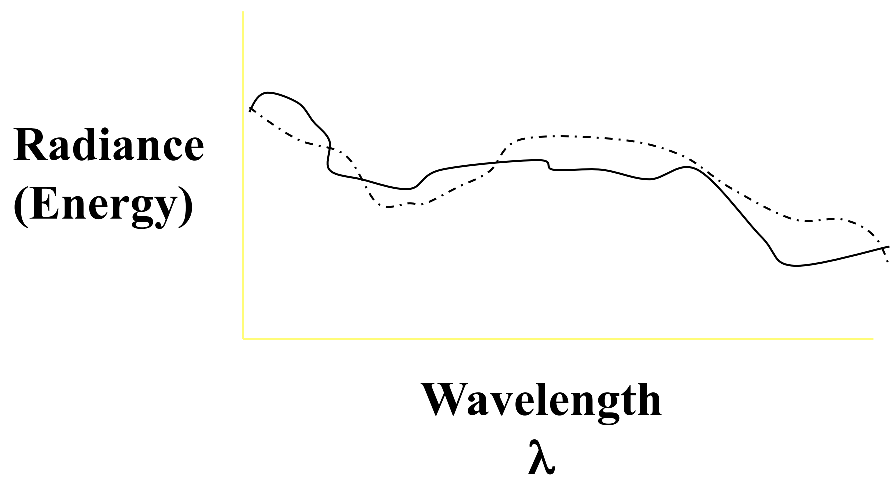

*Metamerism is important to many industries, but is fickle; it can vary person to person and lighting situation to lighting situation. Metameric **failure** can be a large problem.*

How Color Cameras Work

- 1 CCD cameras
  - A **Bayer** pattern is placed in front of the CCD
  - A **Demosaicing** process reads the pixels in a region and computes color and intensity

- 3 CCD camera use a beam splitter and 3 separate CCDs
  - higher color fidelity
  - needs lots of light
  - requires careful alignment of ccds

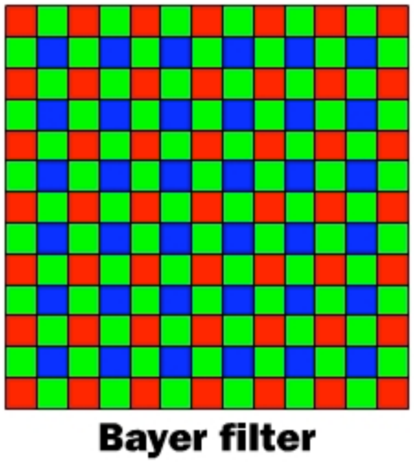

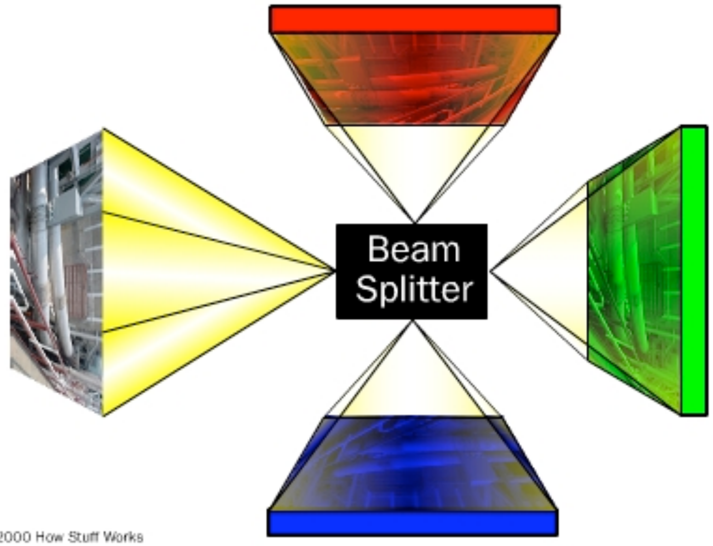

Unfiltered CCD Response

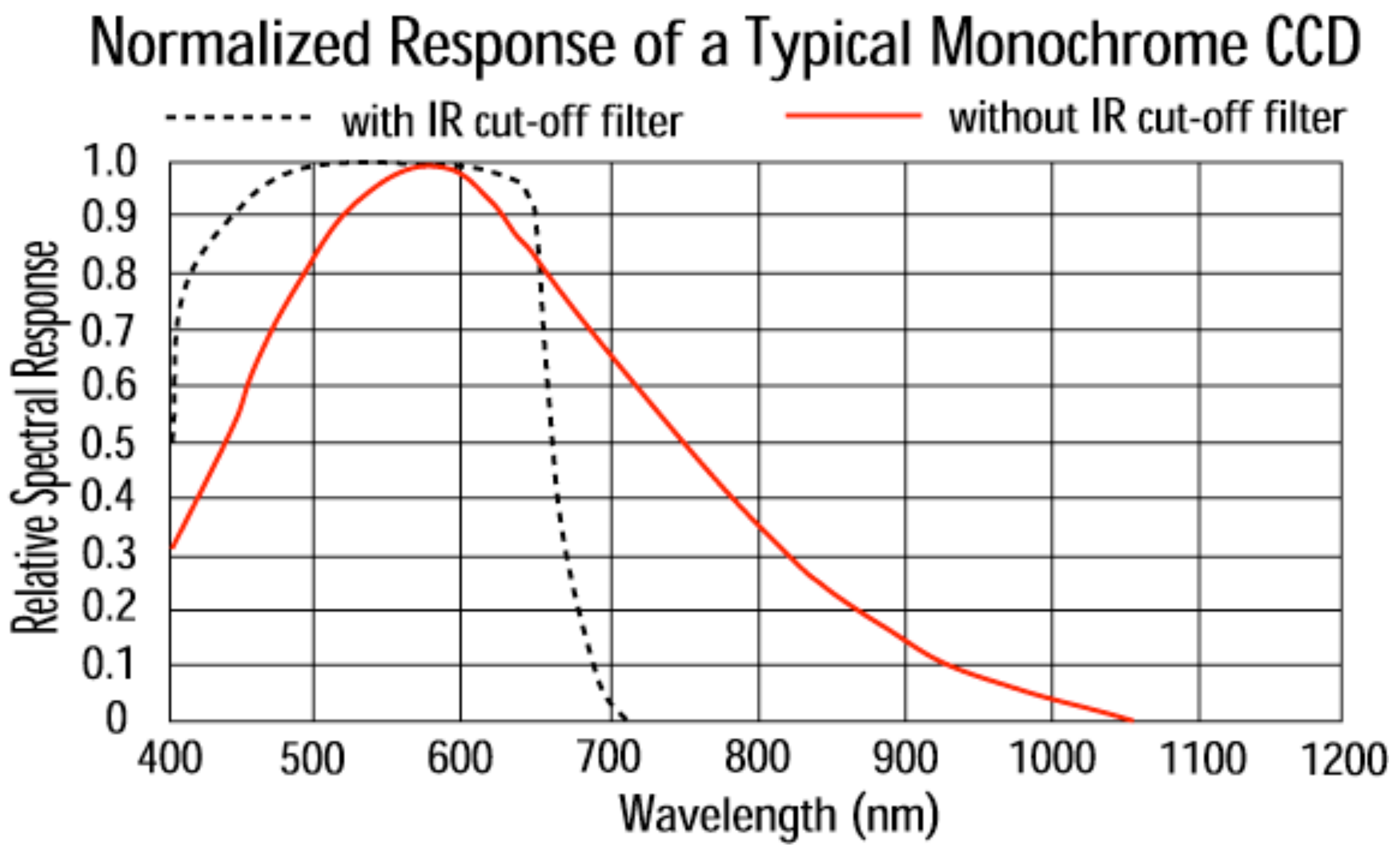

One Chip CCD Response (Sony DFW V500)

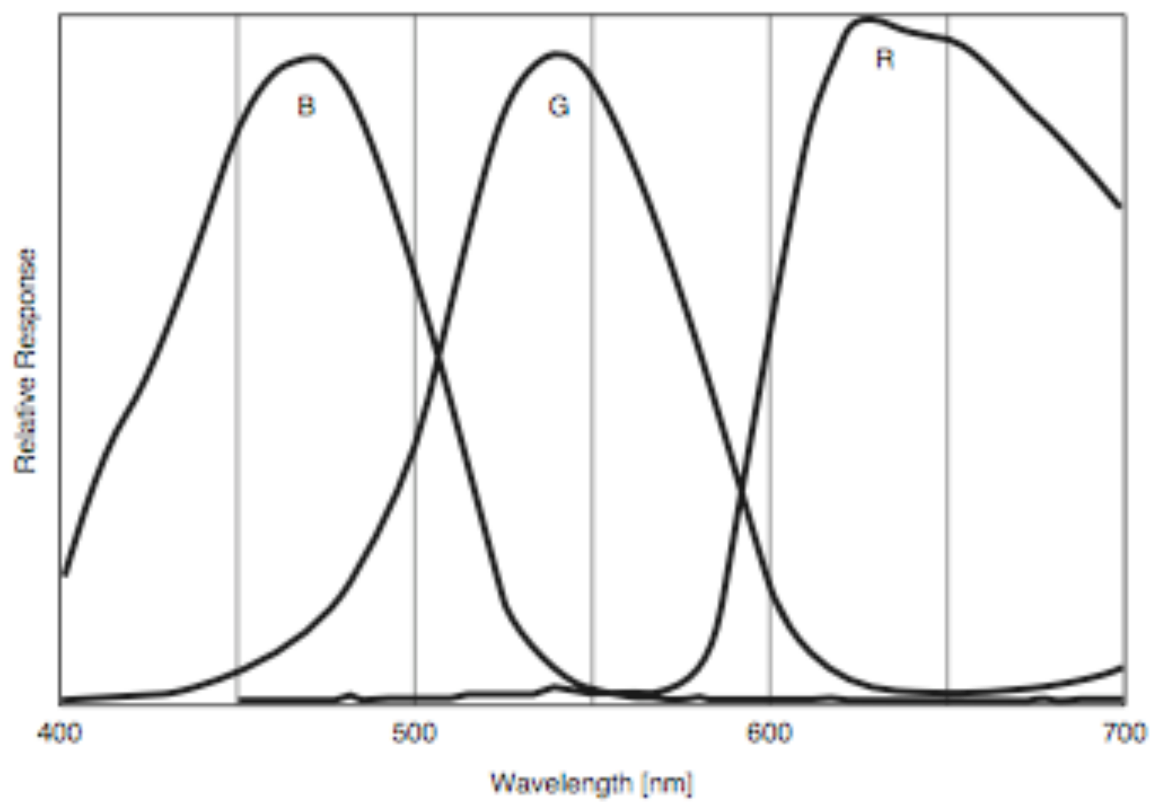

Filtering Colors

Green Red

Standard Linear Color Systems

- Several standards are used to define “color” based on specific spectral response functions
  - CIE (Commission International d’Eclairage) establishes standards
  - CIE XYZ is a popular standard with everywhere positive response
  - RGB requires a negative (subtractive) component in R response to render the complete color gamut of CIE XYZ

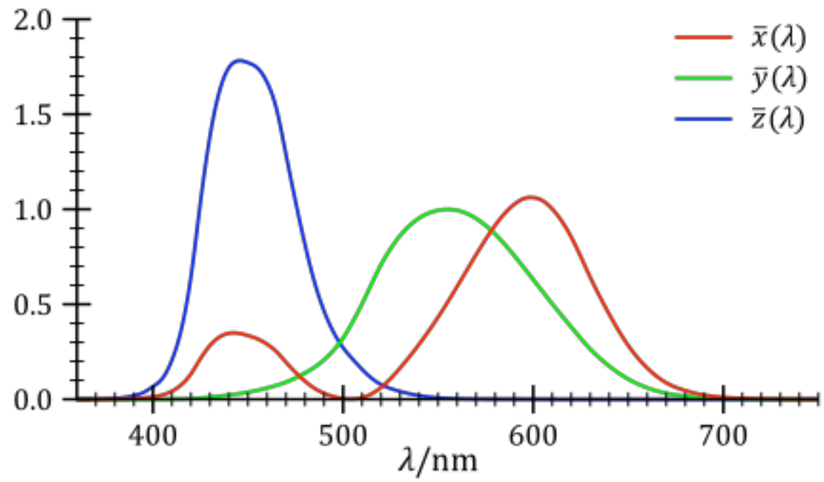

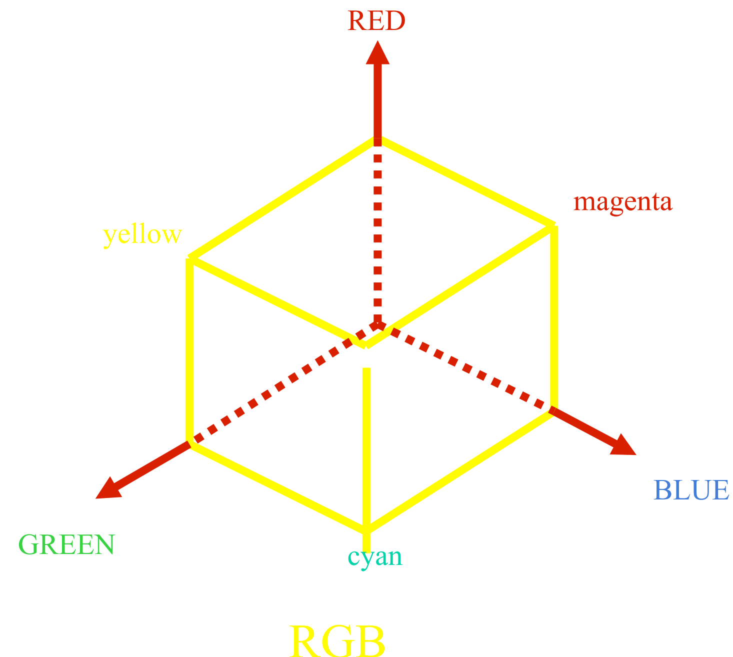

ANOTHER LINEAR SCHEME FOR REPRESENTING COLOR

- Invented for color television (NTSC)
- Backward compatible with B/W TV
- Y given higher bandwidth than I/Q

$$\begin{bmatrix} Y \\ I \\ Q \end{bmatrix} = \begin{bmatrix} .3 & .59 & .11 \\ .6 & -.28 & -.32 \\ .21 & -.52 & .31 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix}$$

YUV is similar to YIQ; PAL vs. NTSC
YCbCr is YUV but with a different reference level for Chromiance

Dividing Up Color Space

HSI is a nonlinear representation of color space. Note the non-uniform treatment of color

$I = (R+G+B)/3$ or
$L = .3R + .6G + .1B$

$S = 1 - 3\min(R,G,B)/I$

$$H = \begin{cases} \cos^{-1}(x) & \text{if } G > B \\ \pi - \cos^{-1}(x) & \text{if } G < B \end{cases}$$

$x = (R-G) + (R-B)/((R-G)^2 + (R-B)(G-B))$

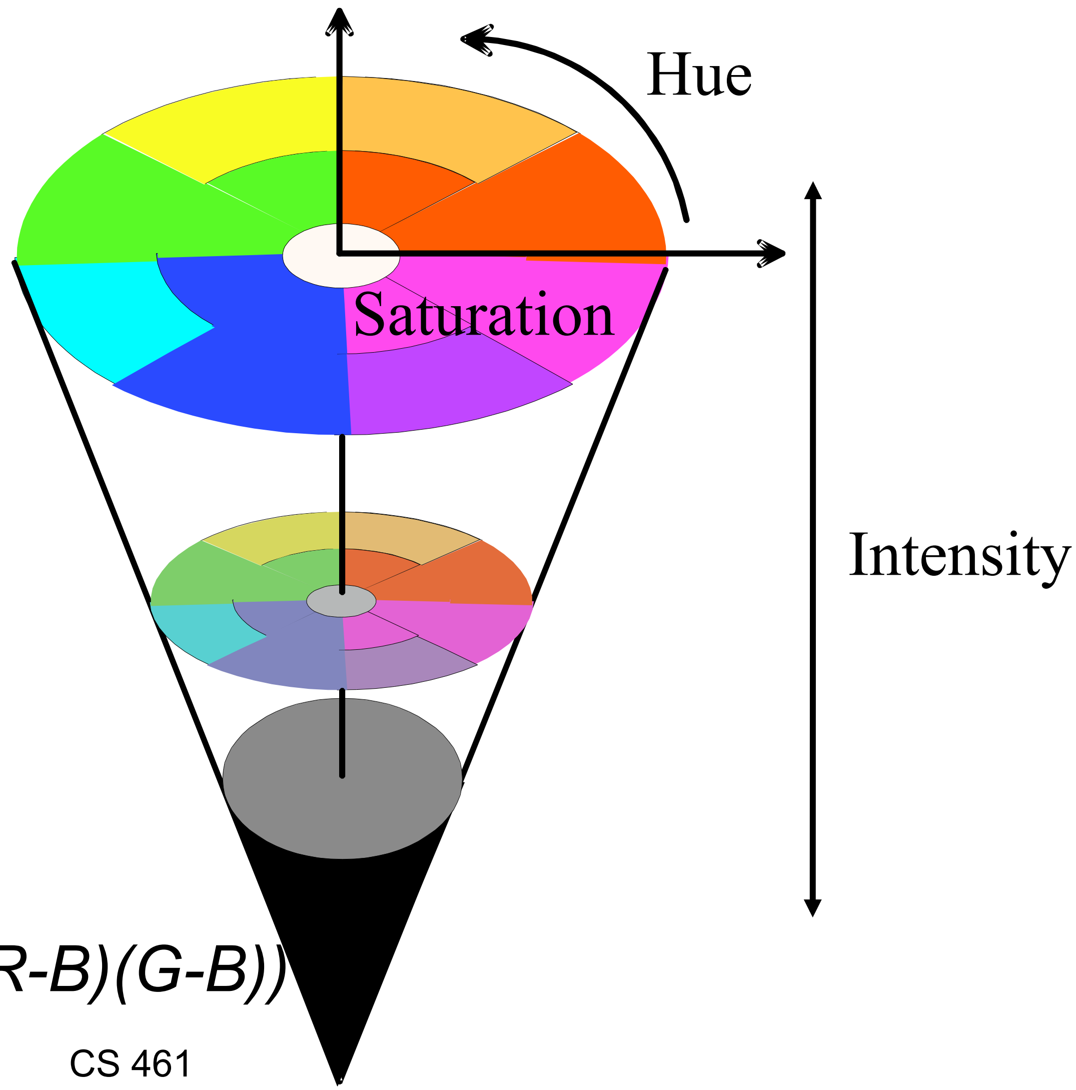

RG Chromaticity

- $r = R/(R+B+G)$ $g = G/(R+B+G)$

- The implied third variable, $b = B / (R + G + B)$, can be omitted from the representation since $r + g + b = 1$, so the blue portion of the color can also be recovered from just r and g.

- The lighting model under which invariance is achieved assumes that changes in the lighting of an object will result in multiplication of its RGB values by a constant.

Tieing Things Together: Color Imaging

- Color imaging is a complex process involving the intrinsic color of the object, in color of the illumination, inter-reflection of surfaces ....

- As a result, what we might perceive as "constant" color is in fact widely varying depending on the vantage point

- There are several extant theories of color constancy, but this is still an open problem: there is no algorithm that comes close to human performance

Example Use: Finding Specularities

- Assume we are dealing with dielectrics
  - specularly reflected light is the same colour as the source
- Reflected light has two components
  - diffuse
  - specular
  - and we see a weighted sum of these two

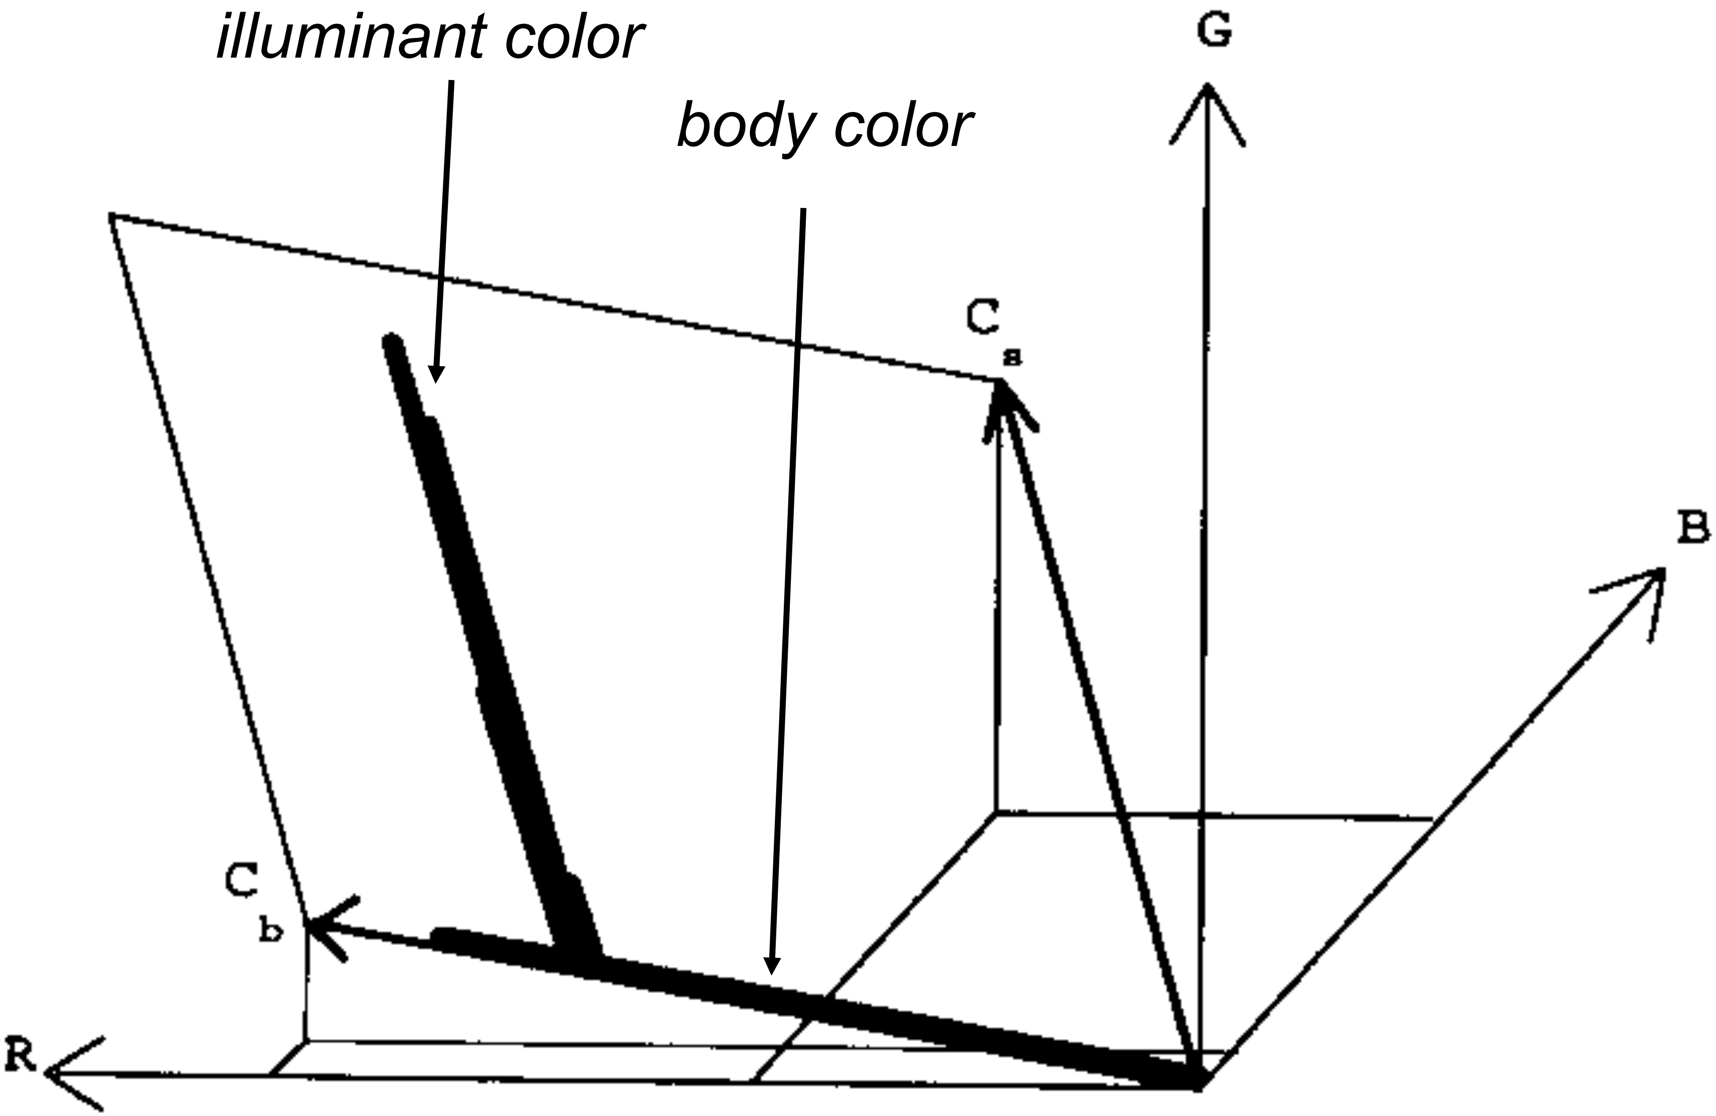
illuminant color
body color
G
B
R
C_s
C_b

Initial Image

Desk Alone

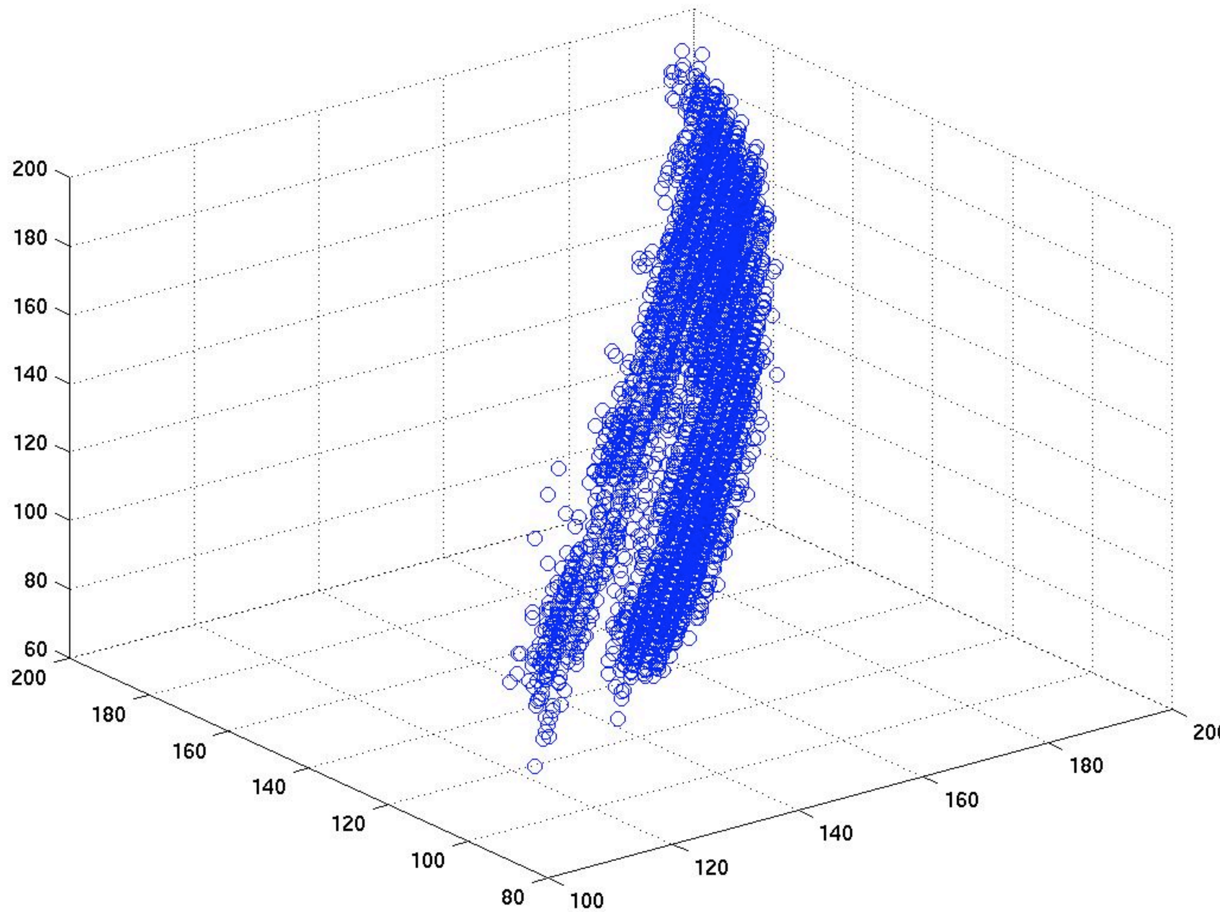
200
180
160
140
120
100
80
60
200
180
160
140
120
100
80
100
120
140
160
180
200

Desk with Specularity

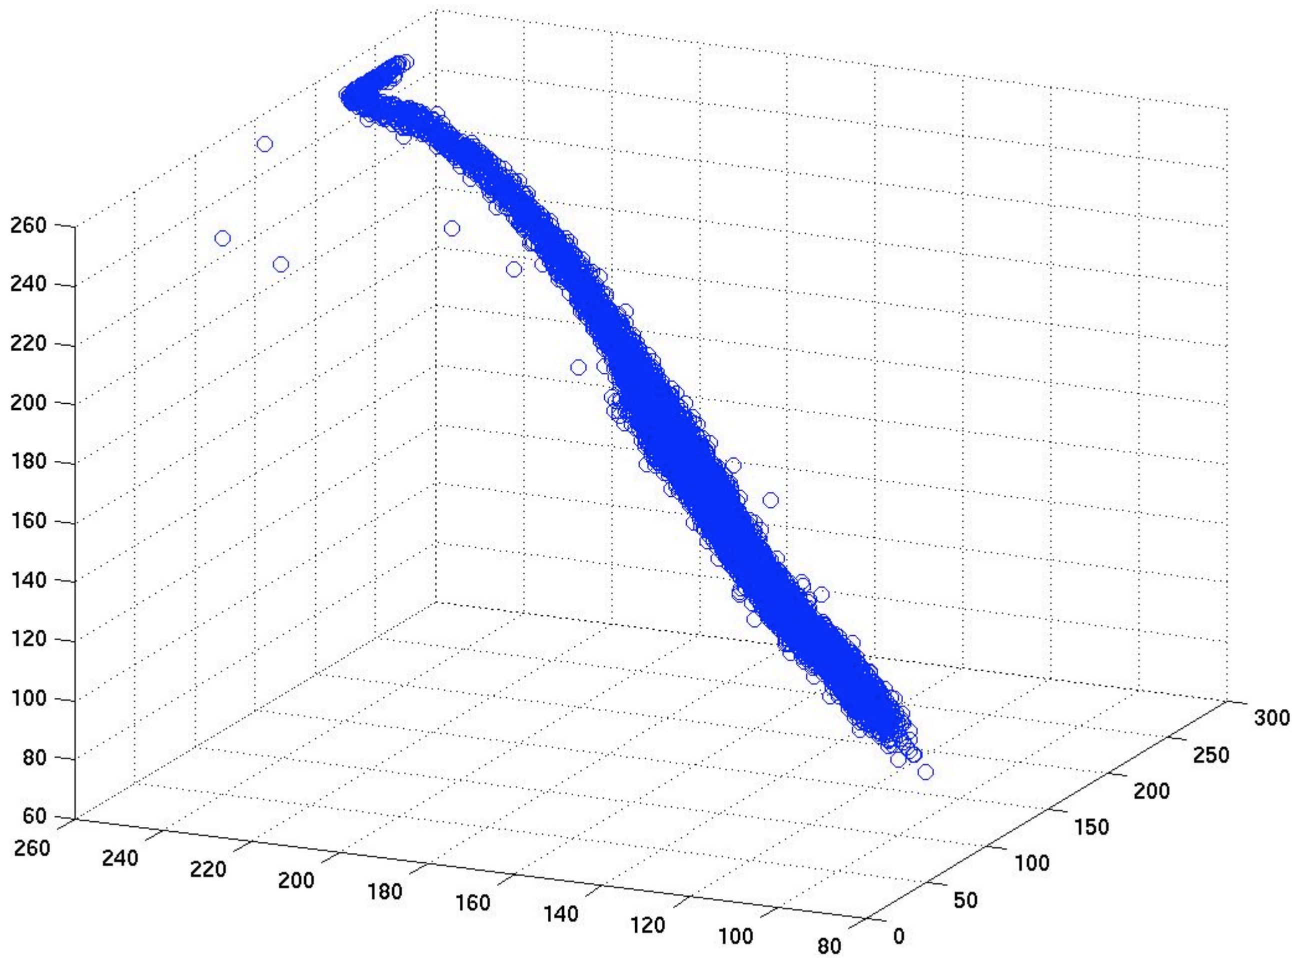

Complete Image

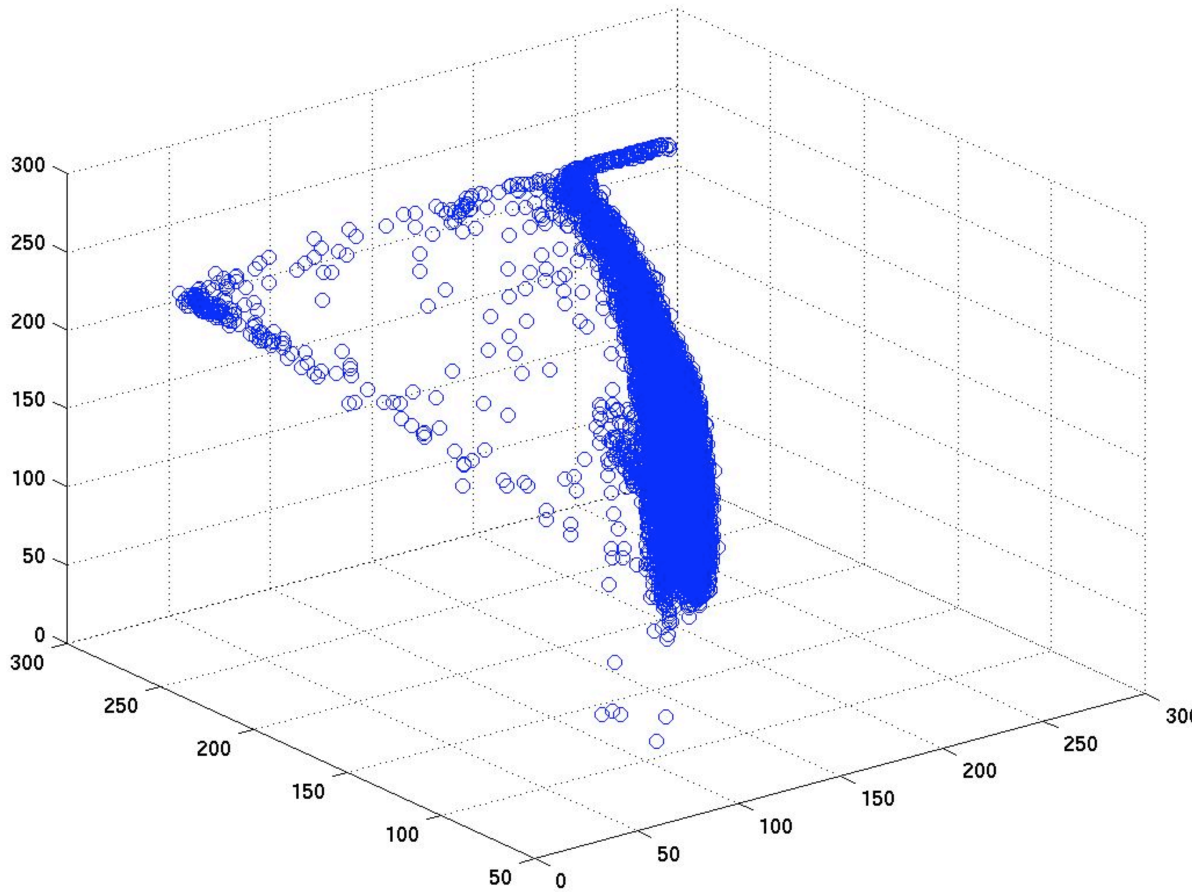

A "Color Space" For Specularities

(Mallick et al. CVPR, 2006)

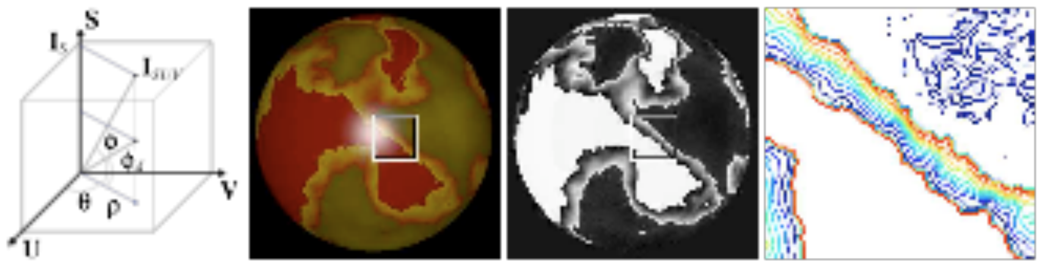

(a) (b) (c) (d)

- Create an YUV-like "color space" where the central axis points in the direction of the source: SUV
- The S axis measures the illumination brightness; off-axis are the remaining color components
- Now, run a diffusion process to "smooth" the S component throughout the image (this is tricky!)

Specularity Removal Result

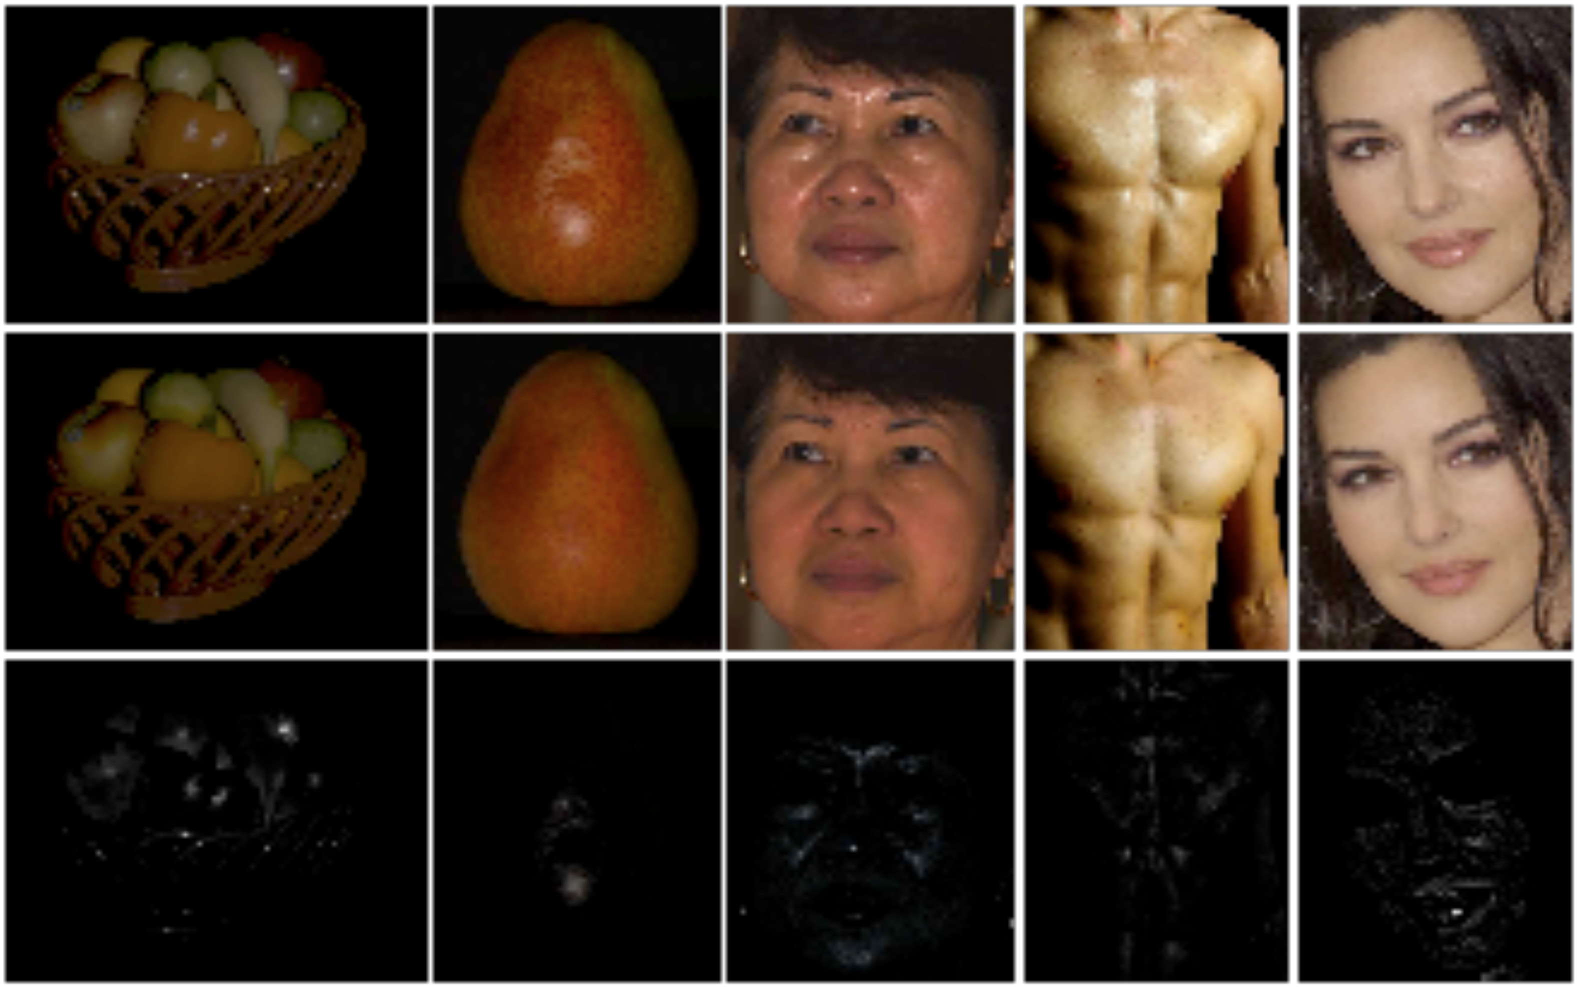

Check out taaz.com

Summary

- Cameras and how they work --> plusses and pitfalls
- Radiance/irradiance and BRDF --> photometric stereo
- Color imaging --> color systems and color constancy
- Images and storage ---> how matlab handles things