Cameras, Radiometry, and Color
http://www.ugrad.cs.jhu.edu/~cs461

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

Unit Outline

- Camera imaging sensor basics
- Radiometry
- Photometric Stereo
- Storing images
- Matlab (if we get to it)
- Color

- Reading:
  – F&P Chap 1.
  – F&P Chap 4 (focus on understanding terms in table 4.1)
  – F&P Chap 5.4
  – F&P Chaper 6 (mostly 6.1-6.4)

A "Traditional" Camera

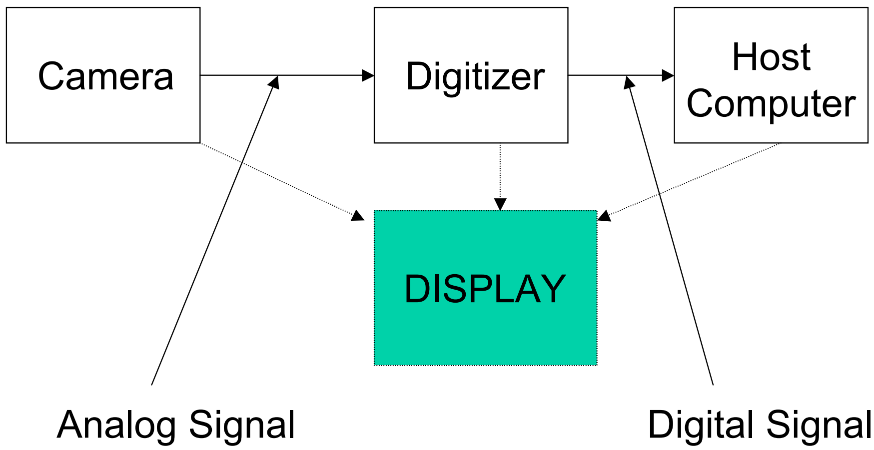

A Modern Digital Camera (Firewire)

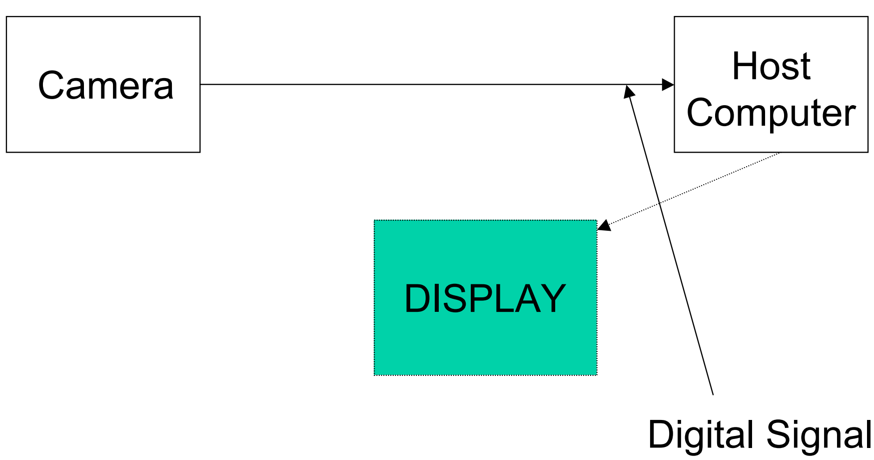

How Cameras Produce Images

- 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

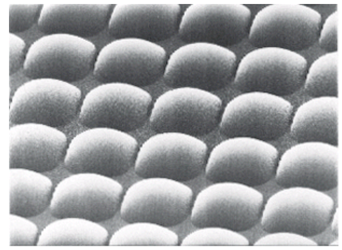

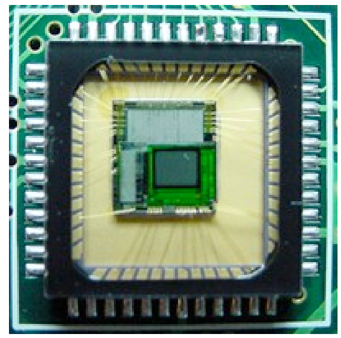

What's under the Hood

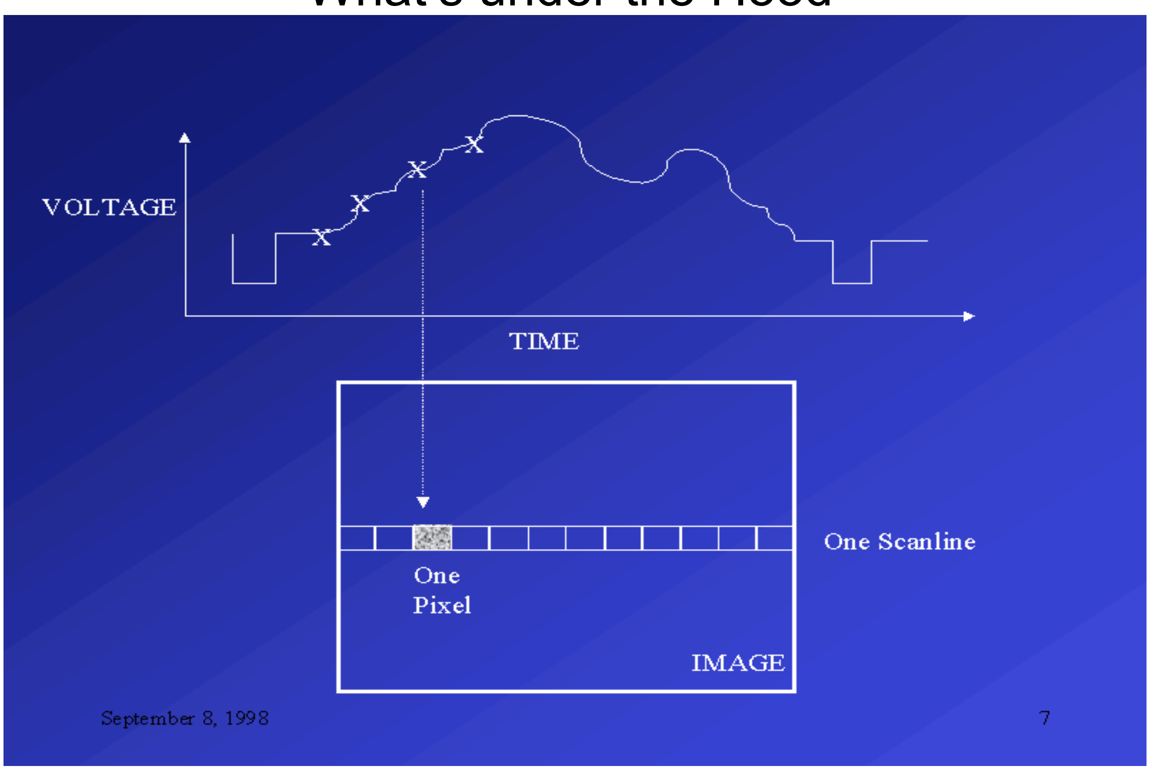

Other Issues

- 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).

REFLECTANCE MODELS

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

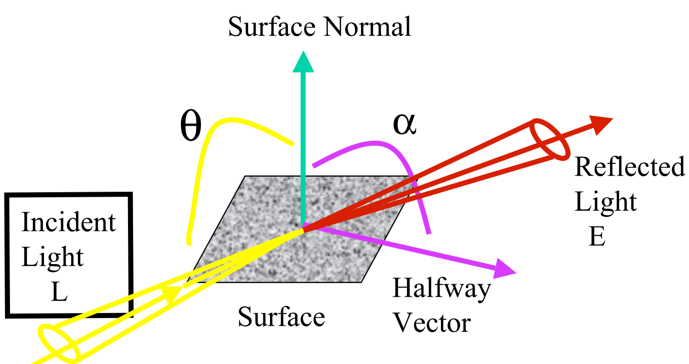

Radiometry

- 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

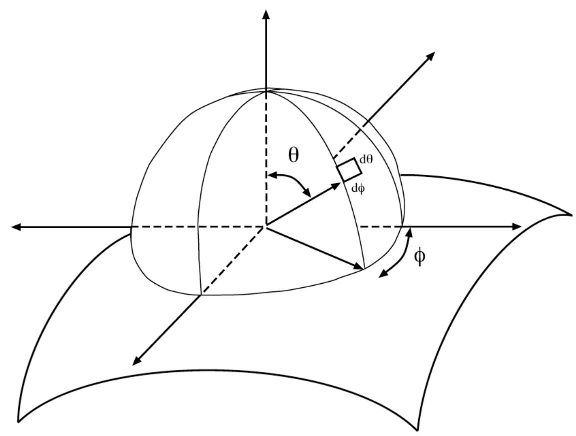

Solid Angle

- 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:

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

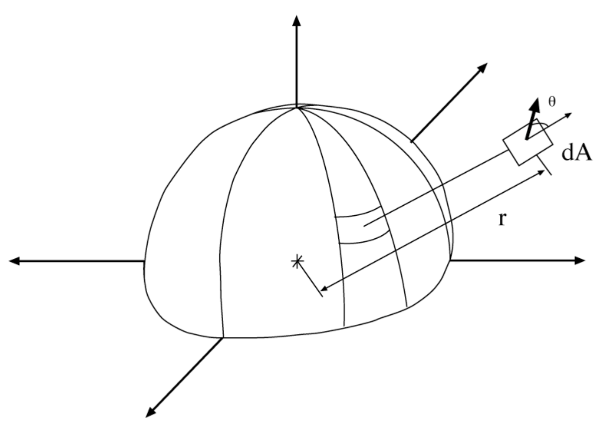

Radiance

- 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

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

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 *<u>foreshortened</u>* <u>area</u>

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

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.

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

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%.

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

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

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

$$P_o\left(\underline{x}\right) = \rho_{brdf}(x) \int_{\Omega} L_i\left(\underline{x}, \vartheta_i, \varphi_i\right) \cos\vartheta_i d\omega = \rho(x)_{brdf} L_{source} \cos\vartheta_{source}$$

Visualizing a Lambertian Refectance Map

$$\mathbf{L_o = L\,\rho\,COS\,\theta}$$

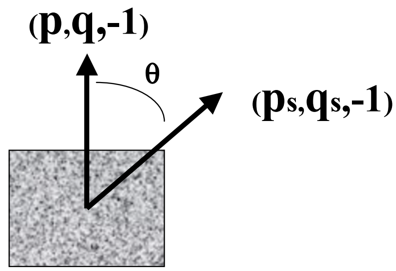

$$COS\theta = \frac{1 + pp_s + qq_s}{\sqrt{1+p^2+q^2}\sqrt{1+p_s{}^2+q_s{}^2}} \qquad L_o = L\rho\frac{1 + pp_s + qq_s}{\sqrt{1+p^2+q^2}\sqrt{1+p_s{}^2+q_s{}^2}}$$

LAMBERTIAN REFLECTANCE MAP

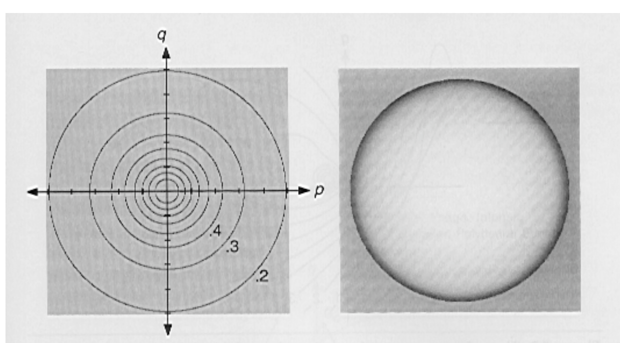

$\mathbf{p_{s=0}\quad q_{s=0}}$

LAMBERTIAN REFLECTANCE MAP

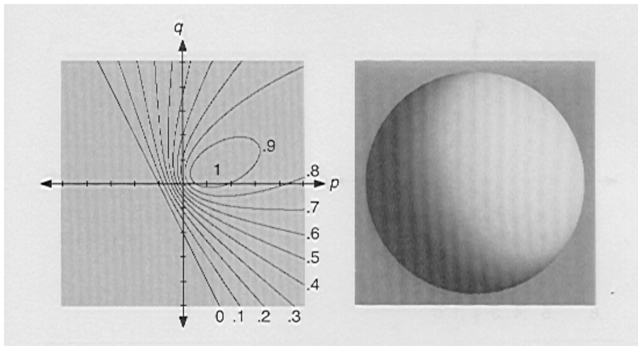

$p_s=0.7$ $q_s=0.3$

LAMBERTIAN REFLECTANCE MAP

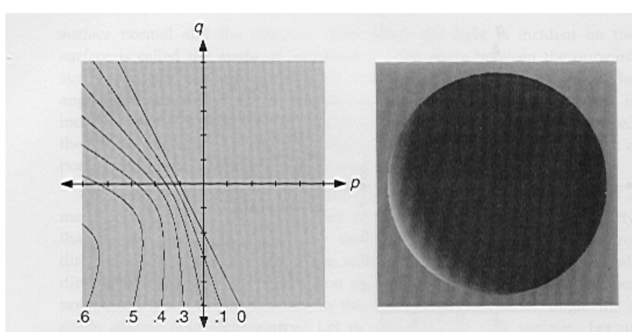

$p_s= -2$ $q_s= -1$

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

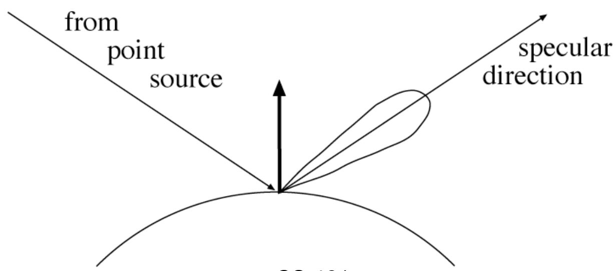

Phong's model

- 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 $\cos^n(\delta\theta)$

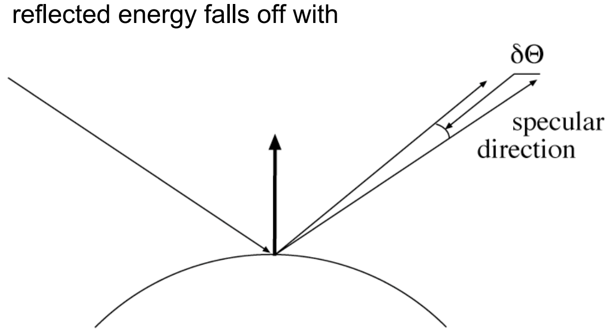

Lambertian + specular model

- 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

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

A Slightly More Elaborate Approach

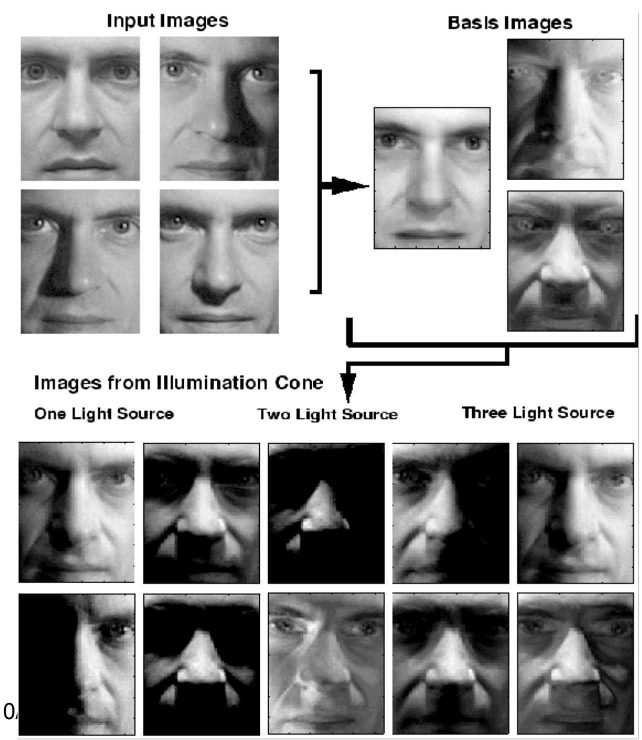

One issue is dealing with occlusion; there are methods that work around that.

Recent work on *Lumigraphs* is an example of taking this general approach of rendering from image to the extreme.

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*

Another View

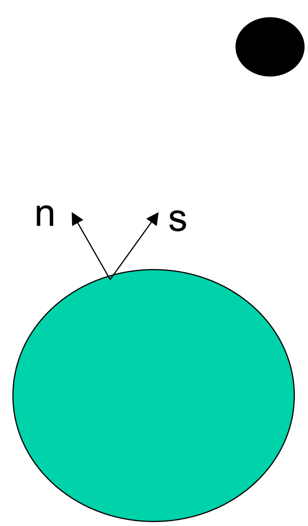

- Recall that $\cos(\theta) = n \cdot s$
  - n unit normal to surface
  - s unit vector directed toward source
- So, we can also write $L_0 = \rho_{brdf}\, n \cdot s$
- An interesting exercise:
  - show that three pinhole light sources suffice to fully determine the photometric structure of a Lambertian surface
  - hint: combine ρ and n into a single (now not unit vector) quanitity and consider varying the direction s
- A second interesting execise: show that given this you can synthesize a view of the scene from a new direction

Photometric Stereo

In the general case, assume we have N points and n illumination directions. For each observed point we have

I(x) = c(x) n(x)' ***S*** *where* ***S*** *is a matrix* $[s_1, s_2, \ldots s_n]$

For all N points, we then have

I = ***CN S*** *where* ***CN*** *is the (unknown)*
N by 3 matrix of c(x) n(x) for each point

A simple least squares solution (assuming at least 3 points) gives

*CN = I S' (S*S')*$^{-1}$ *(here ' denotes transpose)*

Photometric Stereo

*CN = I S' (S*S')*$^{-1}$ *(here ' denotes transpose)*

Now, each row of CN is a 3 vector where

$n(x) = CN_x / ||CN_x||$

Section 5.4.2 of F&P then describes how to get surface geometry from this information

Monge Patches: Normals to Surfaces

- Let's assume (quite reasonably) that the surface is of the form (x,y,f(x,y))
- It follows that the surface normal direction is (-df/dx, -df/dy,1) and the normal is the unitized version of this
- Given a normal vector (a,b,c) at location x,y, it follows that
  - df/dx = a/c and df/dy = b/c
  - note also that it should be the case that d^2 f / dx dy = d^2 f /dy dx; we can use this to either check the data, or regularize our reconstruction
- To find the height at a location u, v just do
  - f(u,v) = s_0^v df/dy(0,y) dy + s_0^u df/dx(x,v) dx

Color

What is Color?

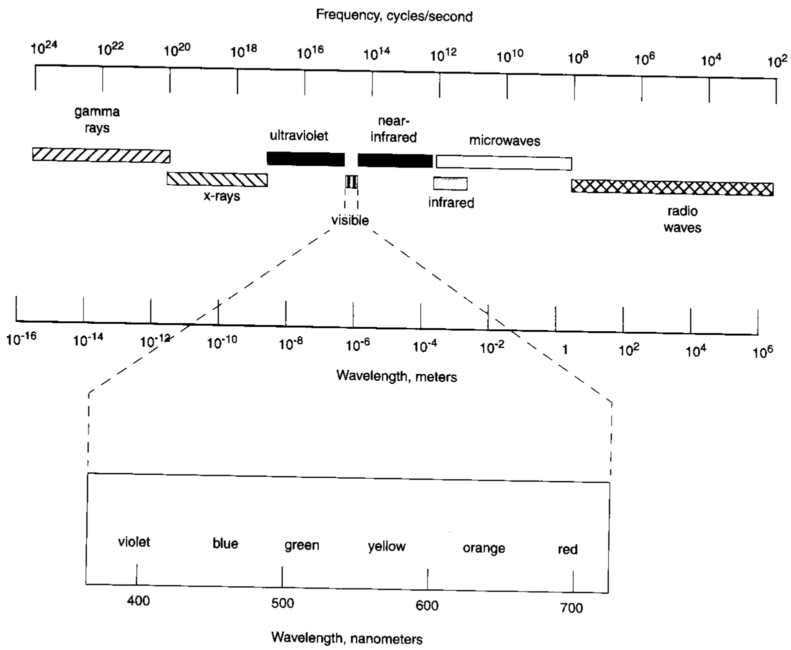

Region Tracking Video

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)

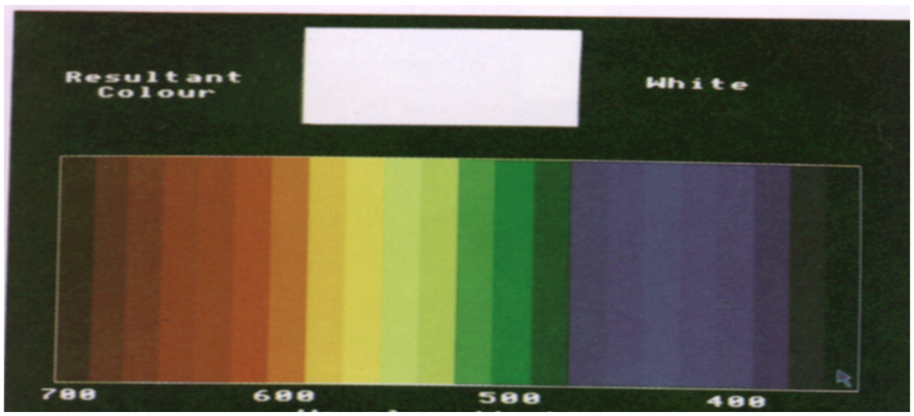

Sunlight

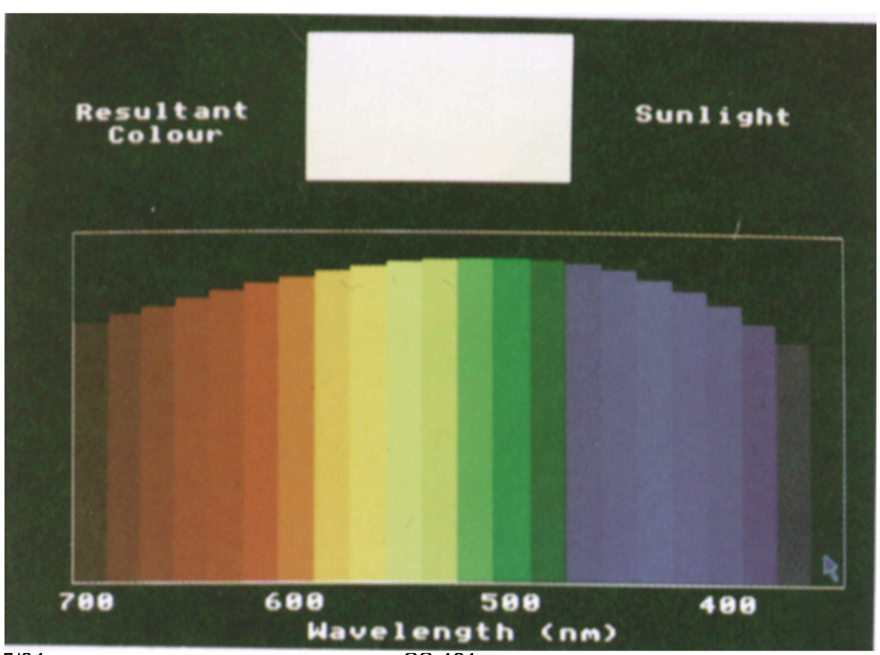

Color Measurement

- Let λ denote wavelength
- Let E(λ) 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

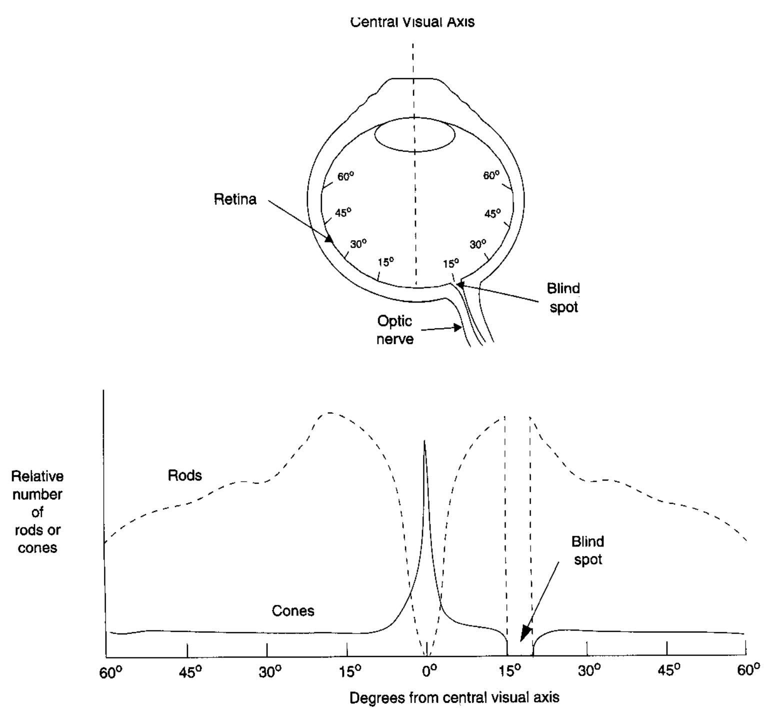

THE HUMAN EYE: RESPONSE

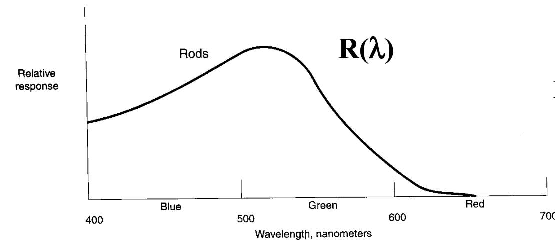

BRIGHTNESS $= \int_{\lambda=400nm}^{\lambda=700nm} R(\lambda)I(\lambda)d\lambda$

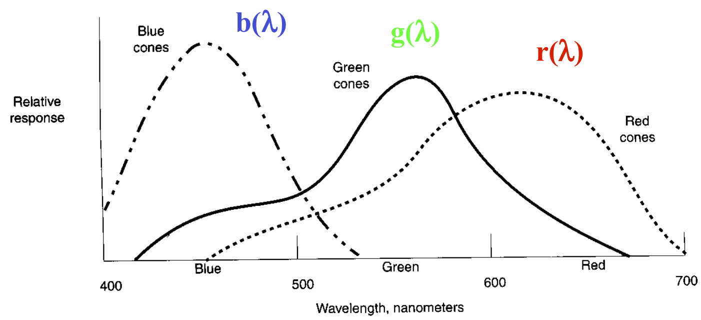

RED $= \int_{\lambda=400nm}^{\lambda=700nm} r(\lambda)I(\lambda)d\lambda$

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

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

Color receptors

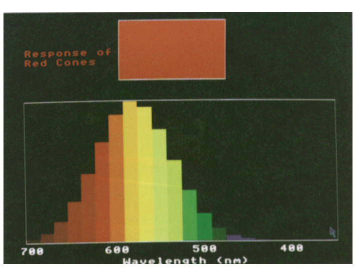

"Red" cone

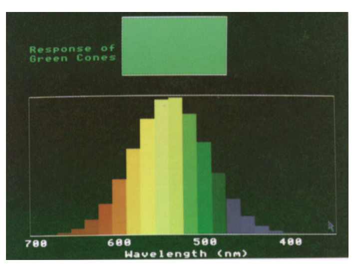

"Green" cone

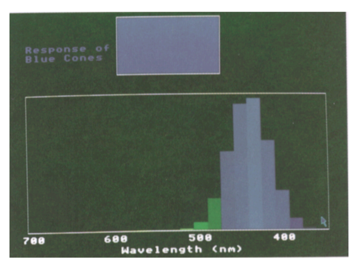

"Blue" cone

Principle of univariance: cones give the same amount
of response, to different wavelengths. 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***.

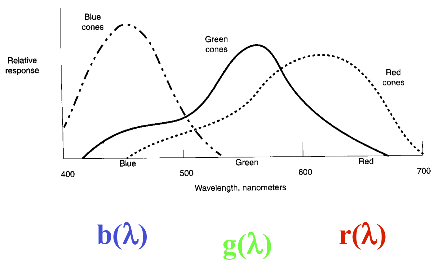

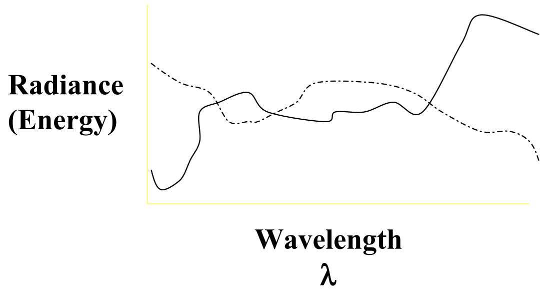

How Color Cameras Work

- 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

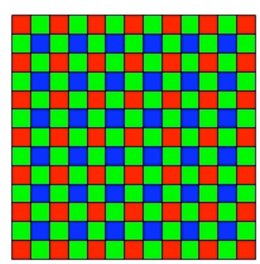
Bayer filter

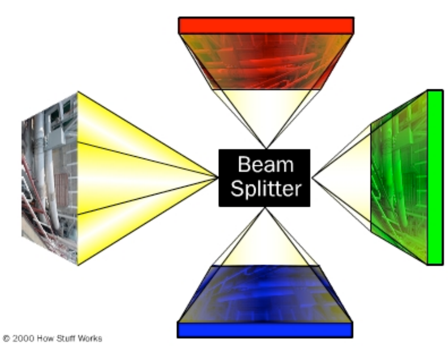

Unfiltered CCD Response

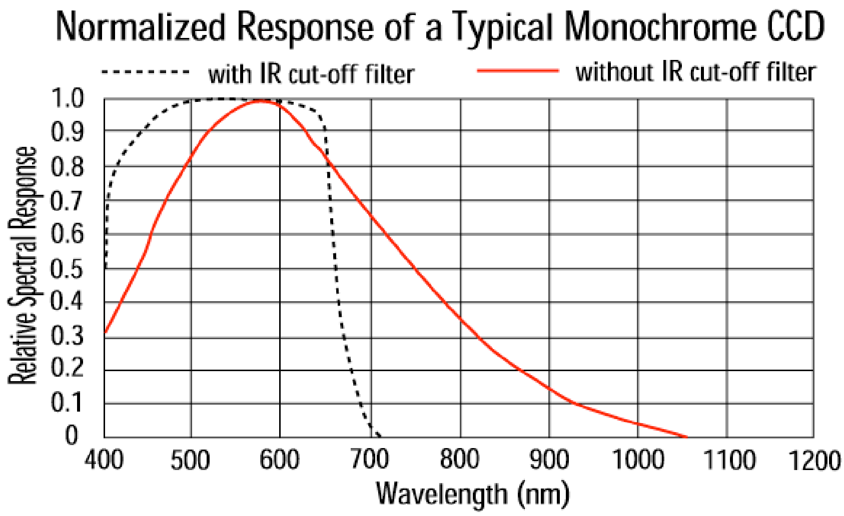

One Chip CCD Response (Sony DFW V500)

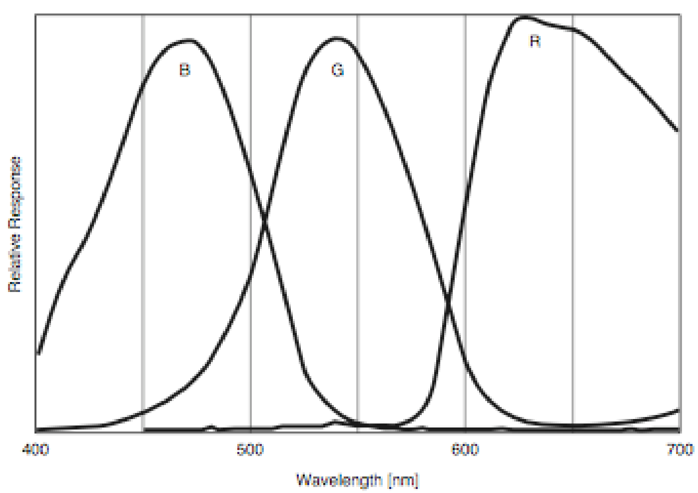

Filtering Colors

Green Red

Standard Linear Color Systems

- Several standards are used to define “color” based on 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

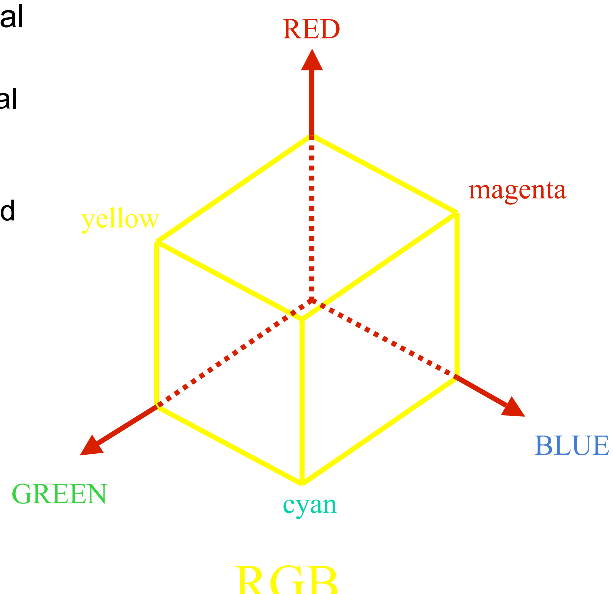

ANOTHER LINEAR SCHEME FOR REPRESENTING COLOR

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

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

YIQ (also YUV)

NONLINEAR SCHEMES FOR REPRESENTING COLOR

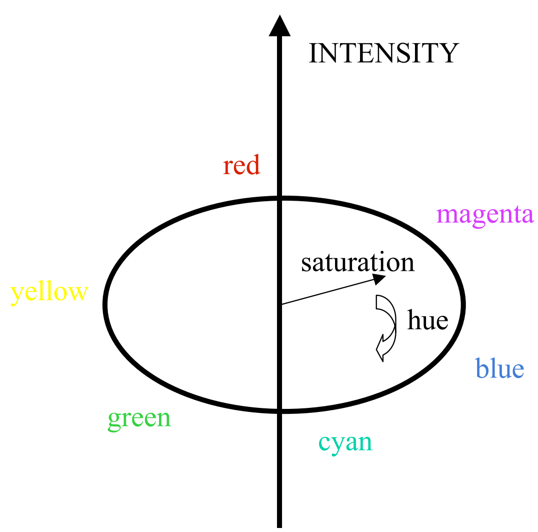

HSI

Dividing Up Color Space

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

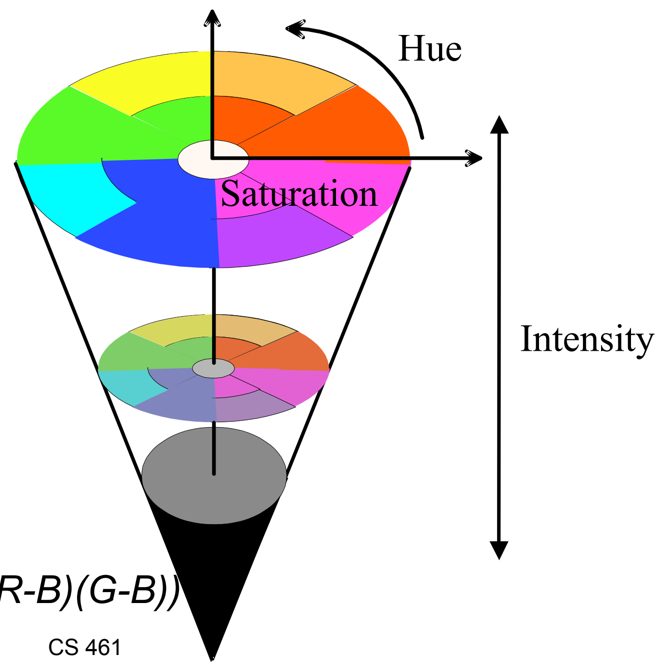

$I = (R+G+B)/3$ or
$L = .3R + .6G + .1B$

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

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

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

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

Homogeneous Region In RGB Space

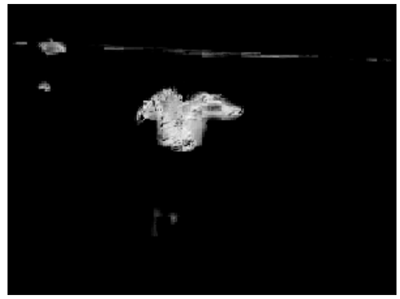

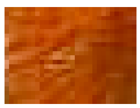

Sample

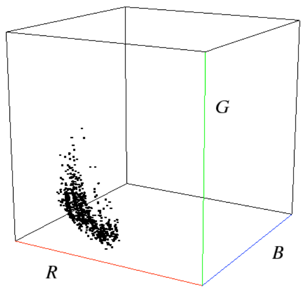

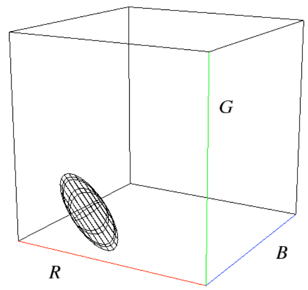

PCA-fitted ellipsoid

Homogeneous Region In RGB Space

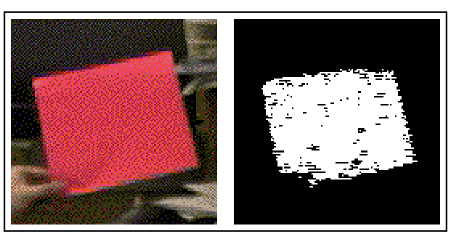

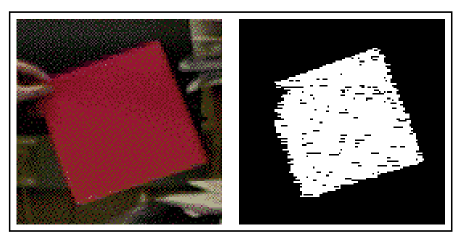

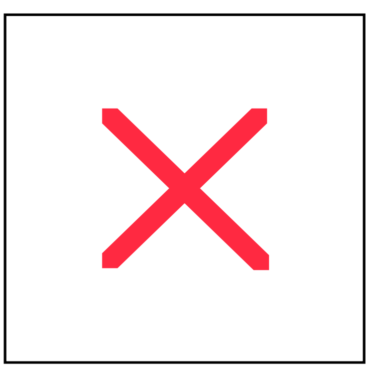

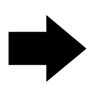

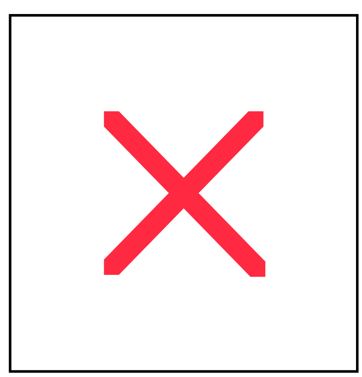

Example Use: Finding Specularities

- 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

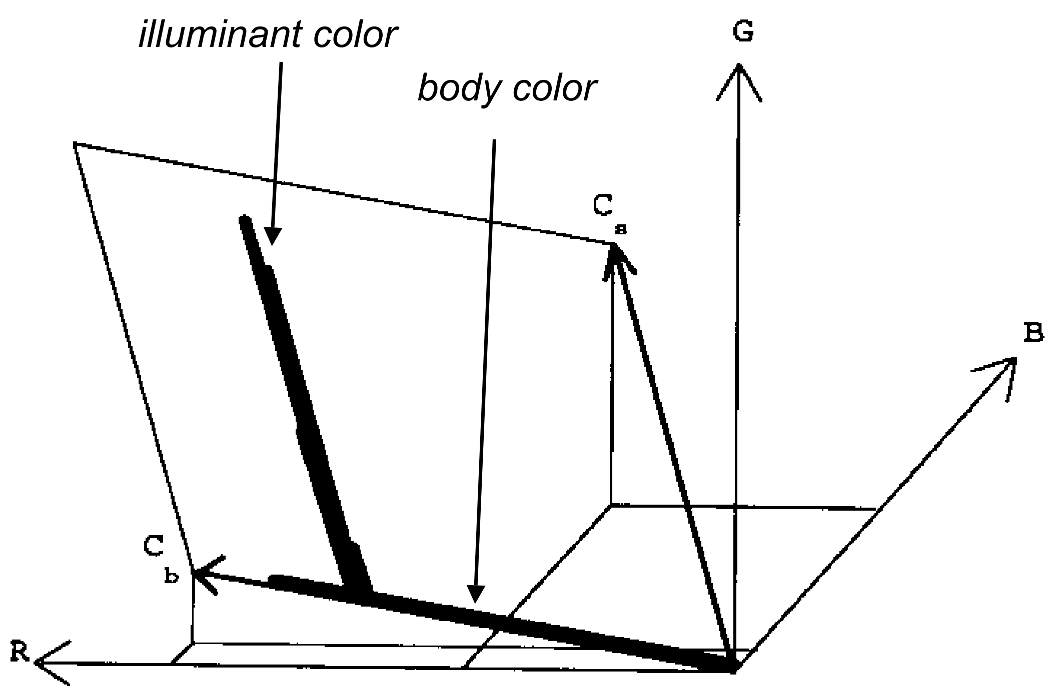

Initial Image

Desk Alone

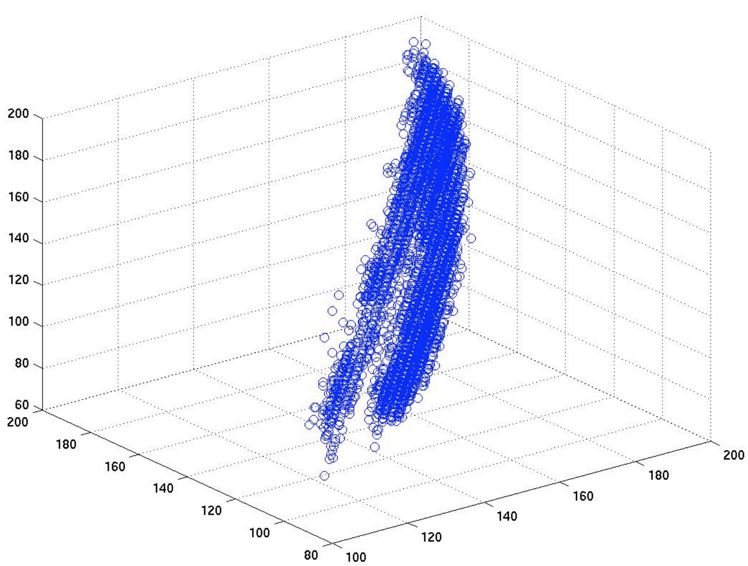

Desk with Specularity

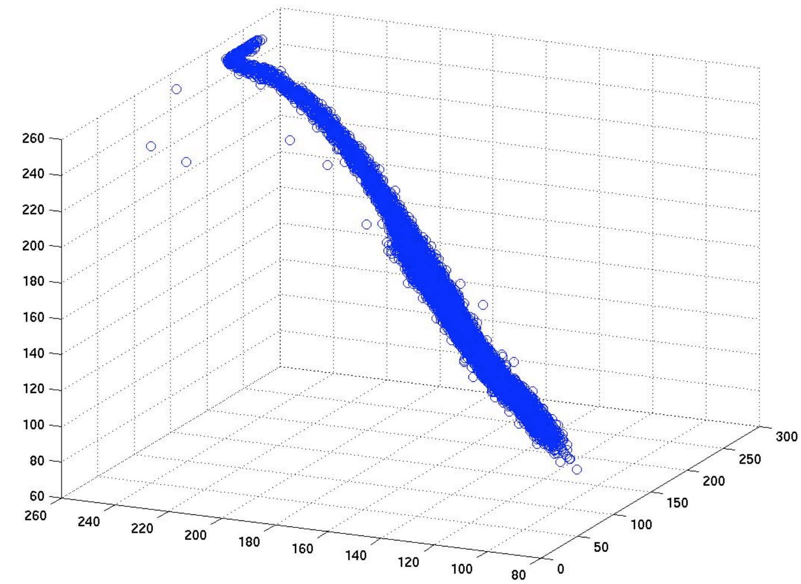

Complete Image

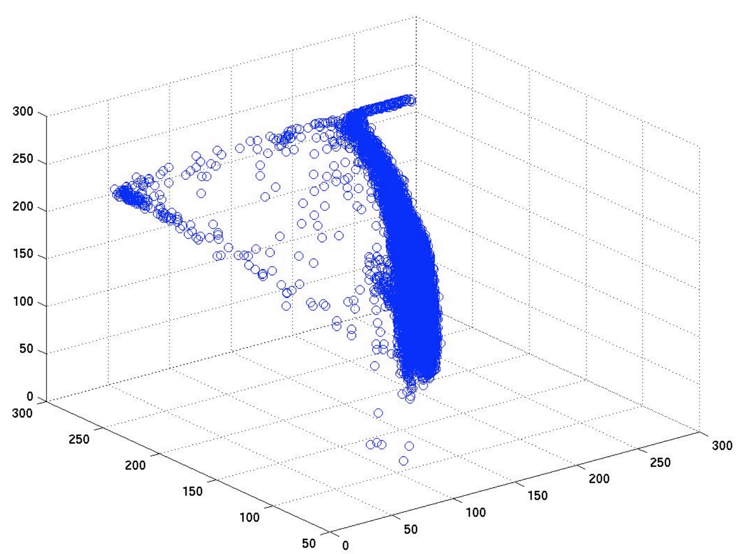

10/15/04 CS 461

THE ORGANIZATION OF A 2D IMAGE

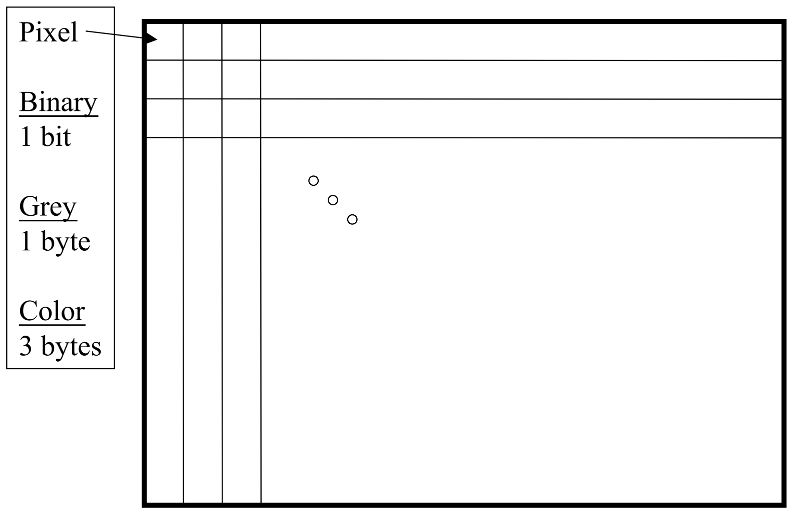

Storing Images

- 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

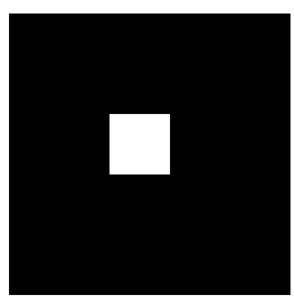

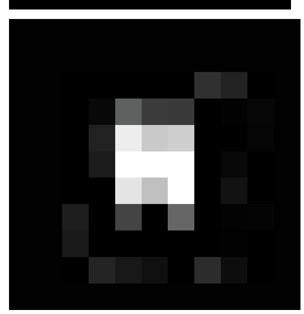

GIF IMAGE FORMAT

- 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)

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