COLLEGE OF ENGINEERING
ELECTRICAL ENGINEERING & COMPUTER SCIENCE
UNIVERSITY OF MICHIGAN

Light, Shadows & Color: Photometric Image Formation

EECS 598-08 Fall 2014
Foundations of Computer Vision

Instructor: Jason Corso (jjcorso)
web.eecs.umich.edu/~jjcorso/t/598F14

Readings: FP 2 and 3; SZ 2.2 and 2.3

Date: 9/22/14

Materials on these slides have come from many sources in addition to myself; I am infinitely grateful to these, especially Greg Hager, Silvio Savarese, and Steve Seitz.

Plan

- Radiometry and the BRDF
- Light and shading models
- The digital camera

Quantum “Catception”

Source: https://www.youtube.com/watch?v=2dRr-fnPCwM

Photometric Image Formation

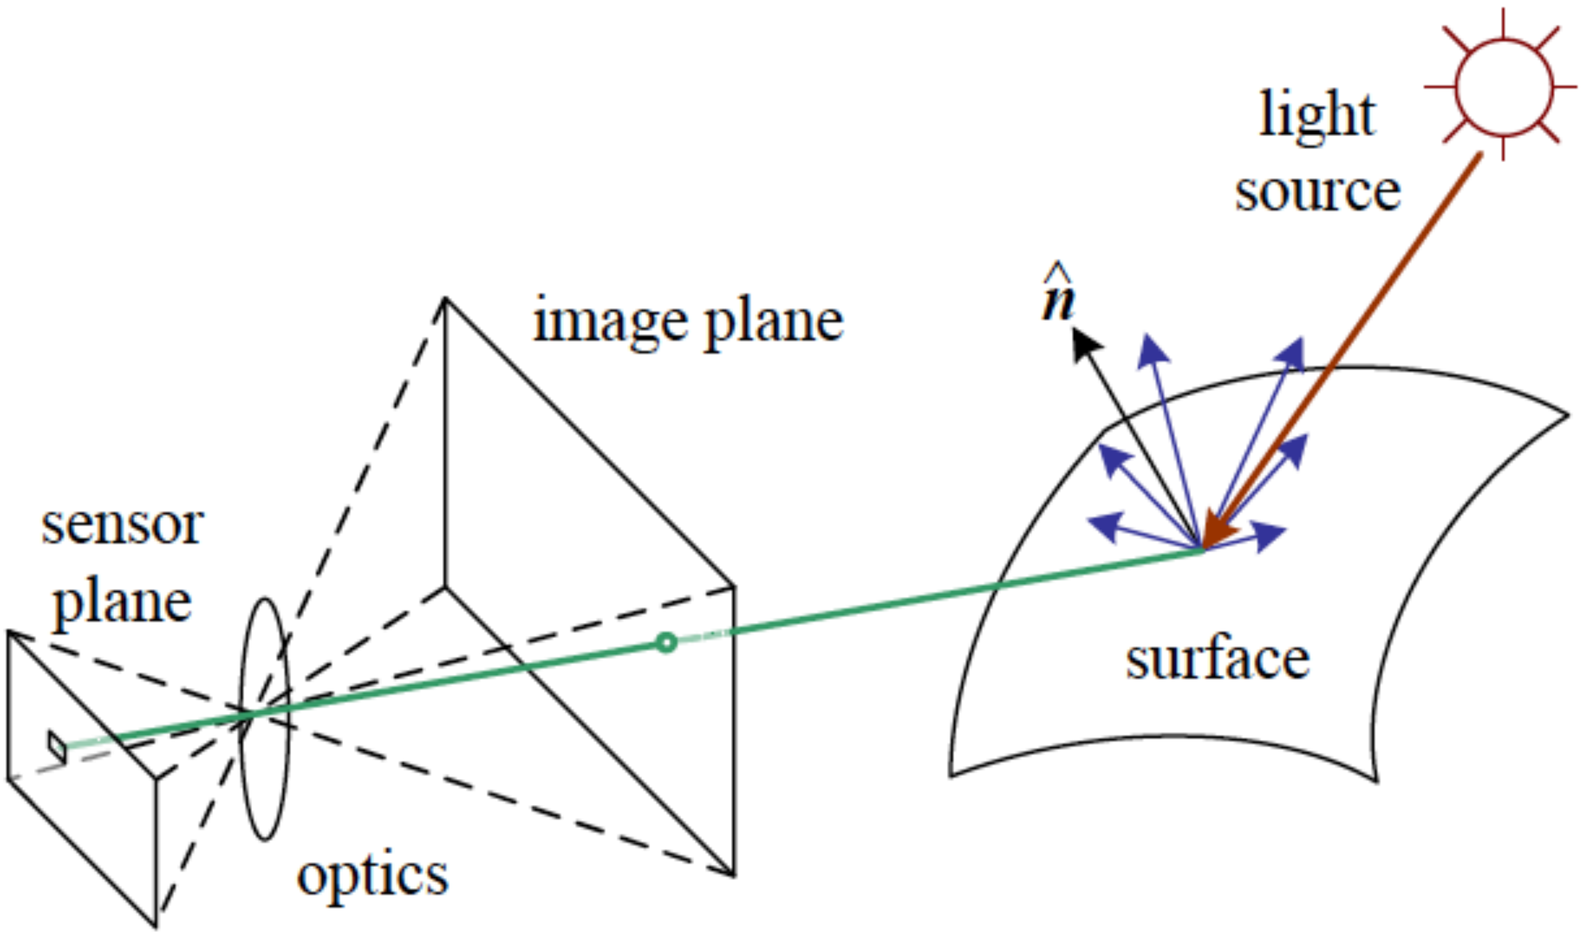

Source: Szeliski book (Fig. 2.14).

Photometric Image Formation

- Three components to pixel brightness
  - Illumination and light sources

  - Surfaces and reflection

  - Camera response

Lights in the real-world

- Real lights are complicated
  - Sun-light, incandescent bulbs, fluorescent bulbs

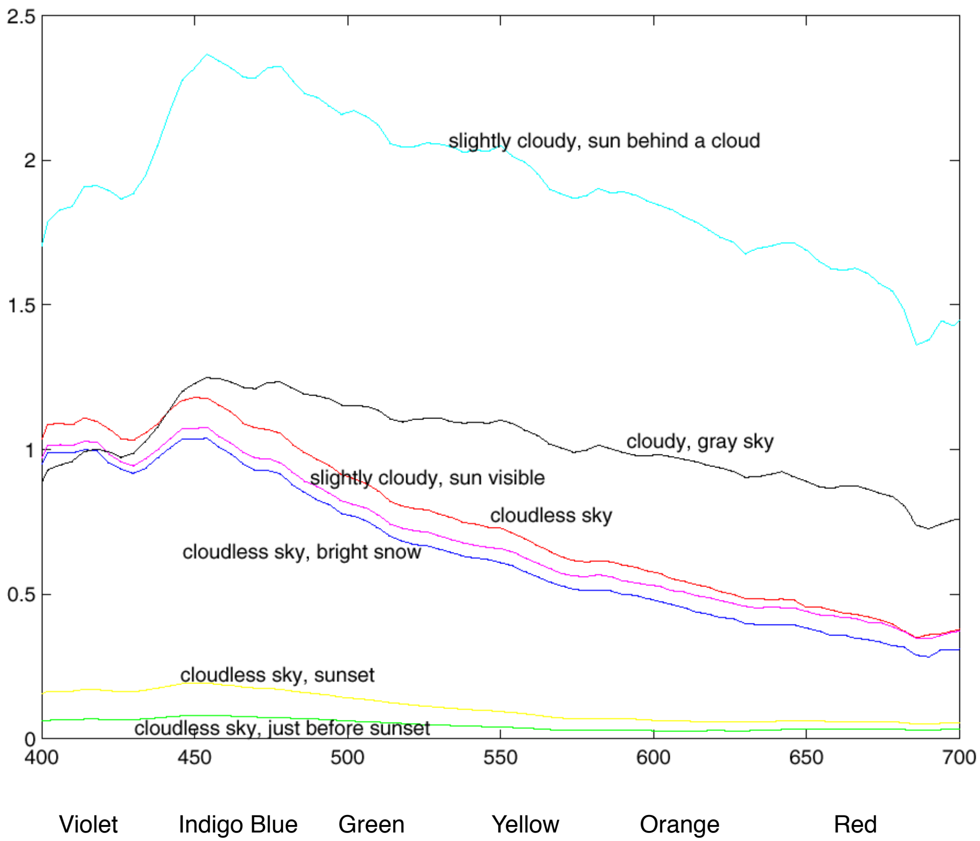

Measurements of relative spectral power of sunlight, made by J. Parkkinen and P. Silfsten. Relative spectral power is plotted against wavelength in nm. The visible range is about 400nm to 700nm. The color names on the horizontal axis give the color names used for monochromatic light of the corresponding wavelength --- the “colors of the rainbow”. Mnemonic is “Richard of York got blisters in Venice”.

Source: Forsyth and Ponce slides

Lights in the real-world

- Real lights are complicated
  - Sun-light, incandescent bulbs, fluorescent bulbs
  - Different spectra
  - Different directions
  - Time-varying
- Fluorscence and biochemistry as well.

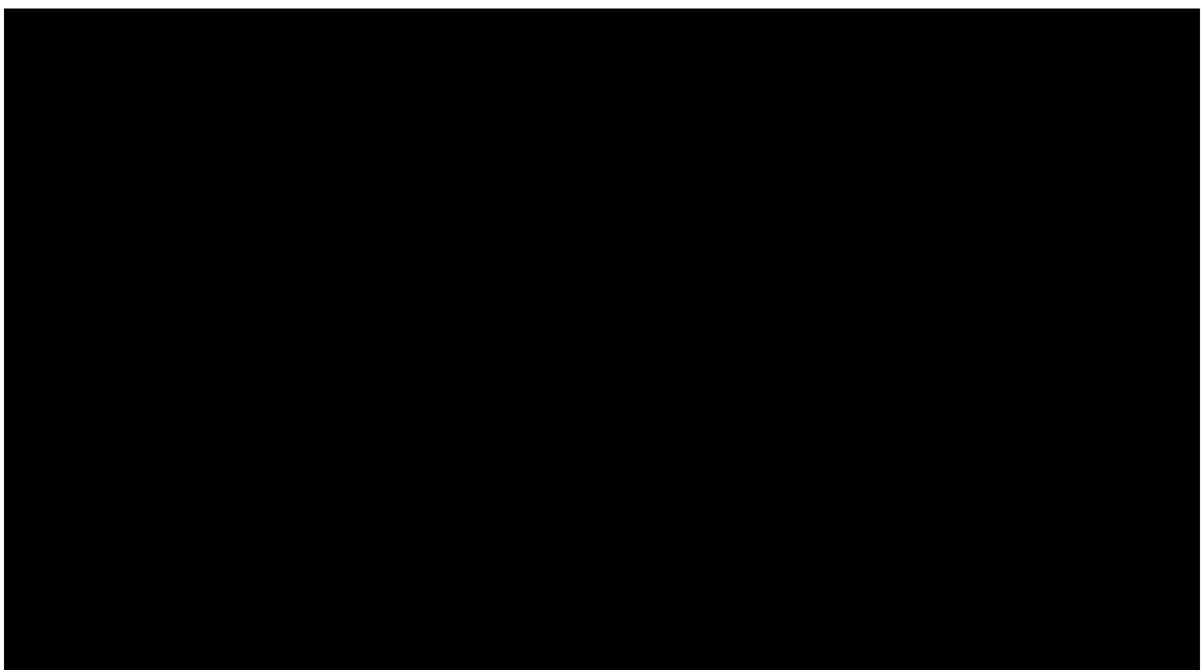

Source: https://www.youtube.com/watch?v=WnWIt0iz00A

Light Models

- Coarse approximations to real light
  - Point light
    - Directional
    - Spot
    - Has a location in space and a distribution over wavelengths $L(\lambda)$
  - Area lights
    - Light-fields
- Environment Map $L(\hat{\boldsymbol{v}}; \lambda)$
  - Maps incident light directions to color values

Source: Image from Pat Hanrahan slides on environment maps

Point Light

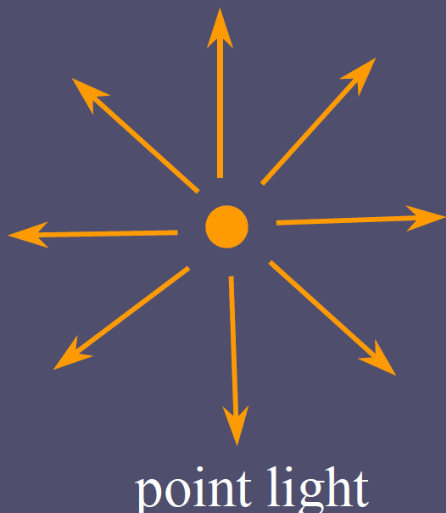

Specified by:

- **position (x,y,z)**
- **intensity (r,g,b)**

Radiates equal intensity in all directions

$$L = P_{light} - P_{surface}$$

Directional Light

- point light at infinity

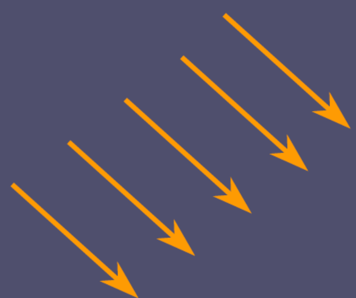

Point light at infinity

Specified by:

- direction (x,y,z)
- intensity (r,g,b)

All light rays are parallel

L = -direction

Spot (Warn) Light

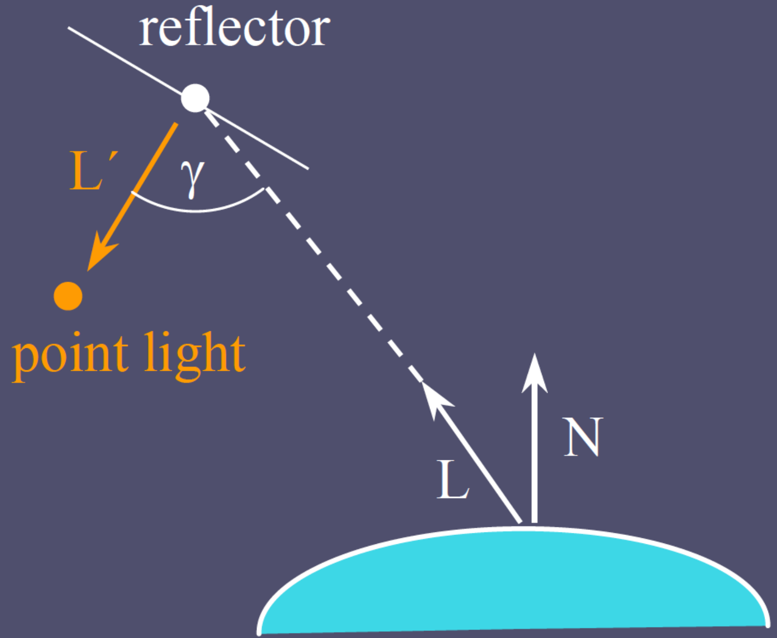

Specular reflection of point light source

Specified by:

- **position of reflector**
- **position of point light (or direction to point light)**
- **intensity of point light**
- **falloff exponent**

$$I_{warn} = I_{point} \cos^p \gamma = I_{point} (V'.R')^p = I_{point} (-L.L')^p$$

Spot (Warn) Light

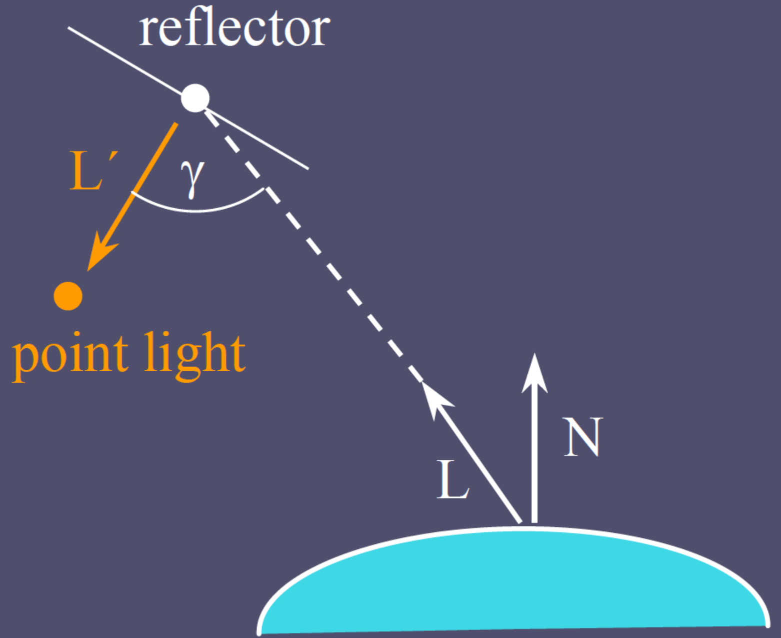

Specular reflection of point light source

Specified by:

- **position of reflector**
- **position of point light (or direction to point light)**
- **intensity of point light**
- **falloff exponent**

$$I_{warn} = I_{point}\cos^p\gamma = I_{point}(V'.R')^p = I_{point}(-L.L')^p$$

Warn Light Profile and Examples

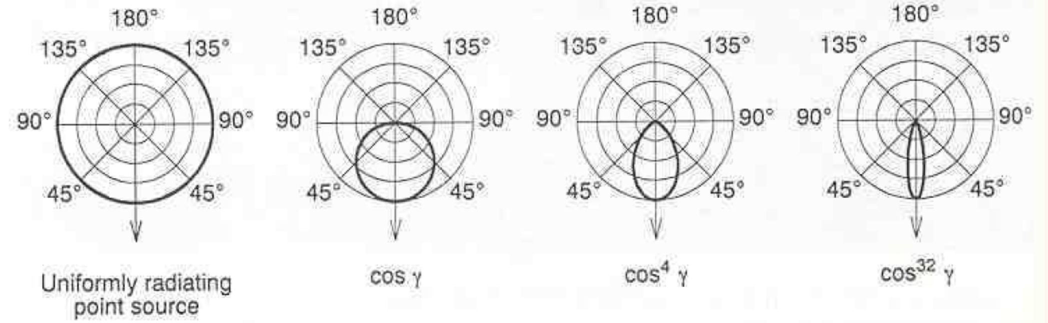

Fig. 16.14 Intensity distributions for uniformly radiating point source and Warn light source with different values of *p*.

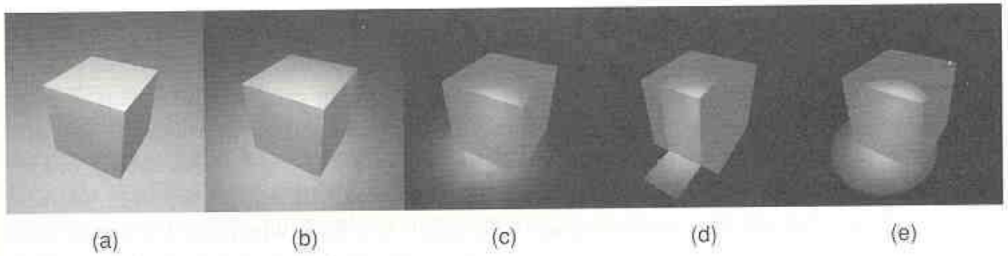

Fig. 16.15 Cube and plane illuminated using Warn lighting controls. (a) Uniformly radiating point source (or $p = 0$). (b) $p = 4$. (c) $p = 32$. (d) Flaps. (e) Cone with $\delta = 18°$. (By David Kurlander, Columbia University.)

From Foley, vanDam, Feiner, and Hughes, **Computer Graphics: Principles and Practice, 2nd edition, page 732, 733**

Surfaces and Reflectance

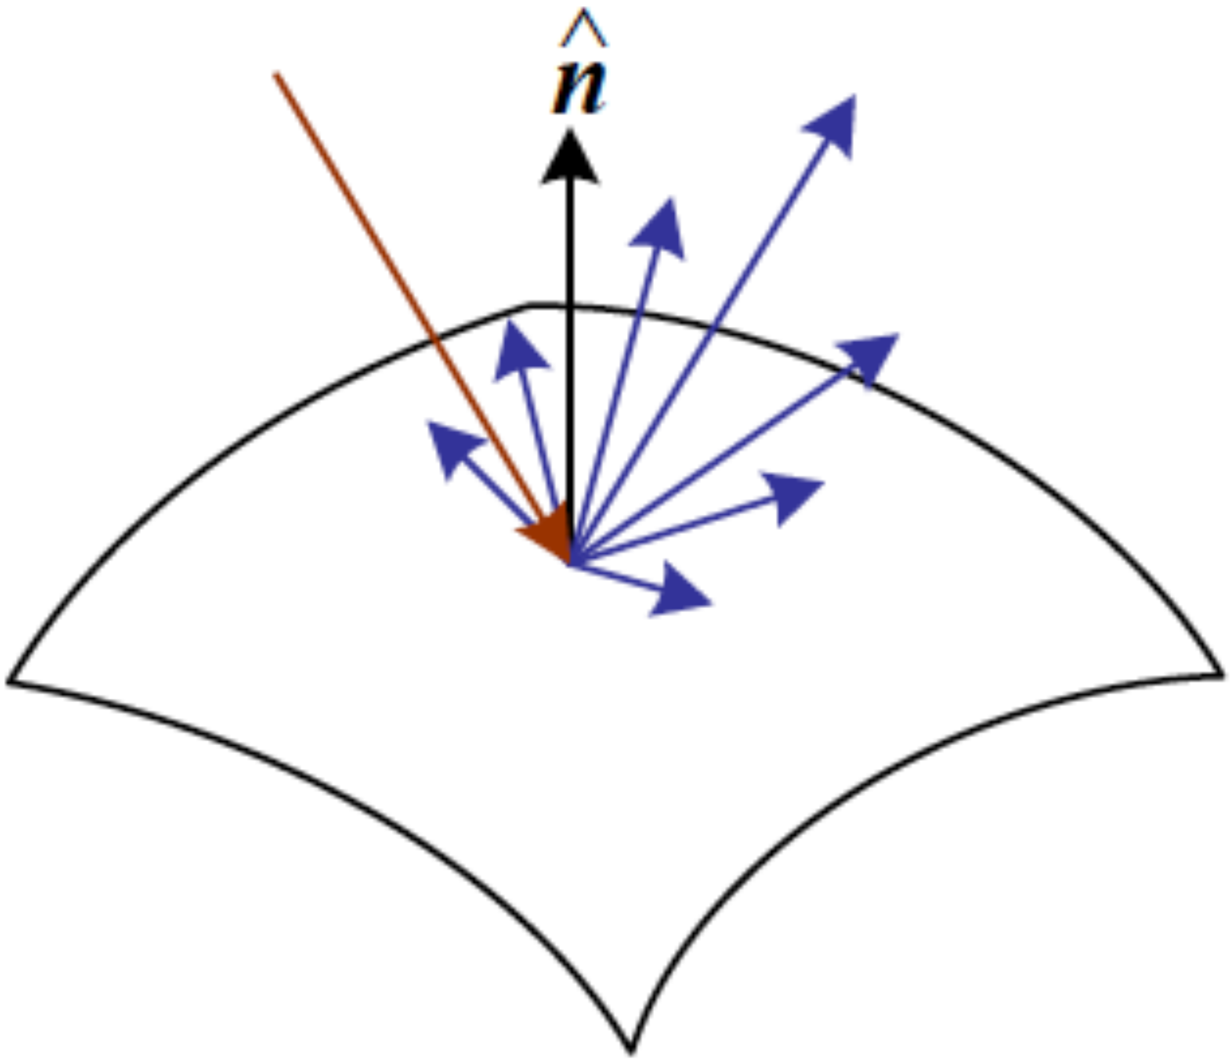

Source: Szeliski book.

Surfaces and Reflectance: The BRDF

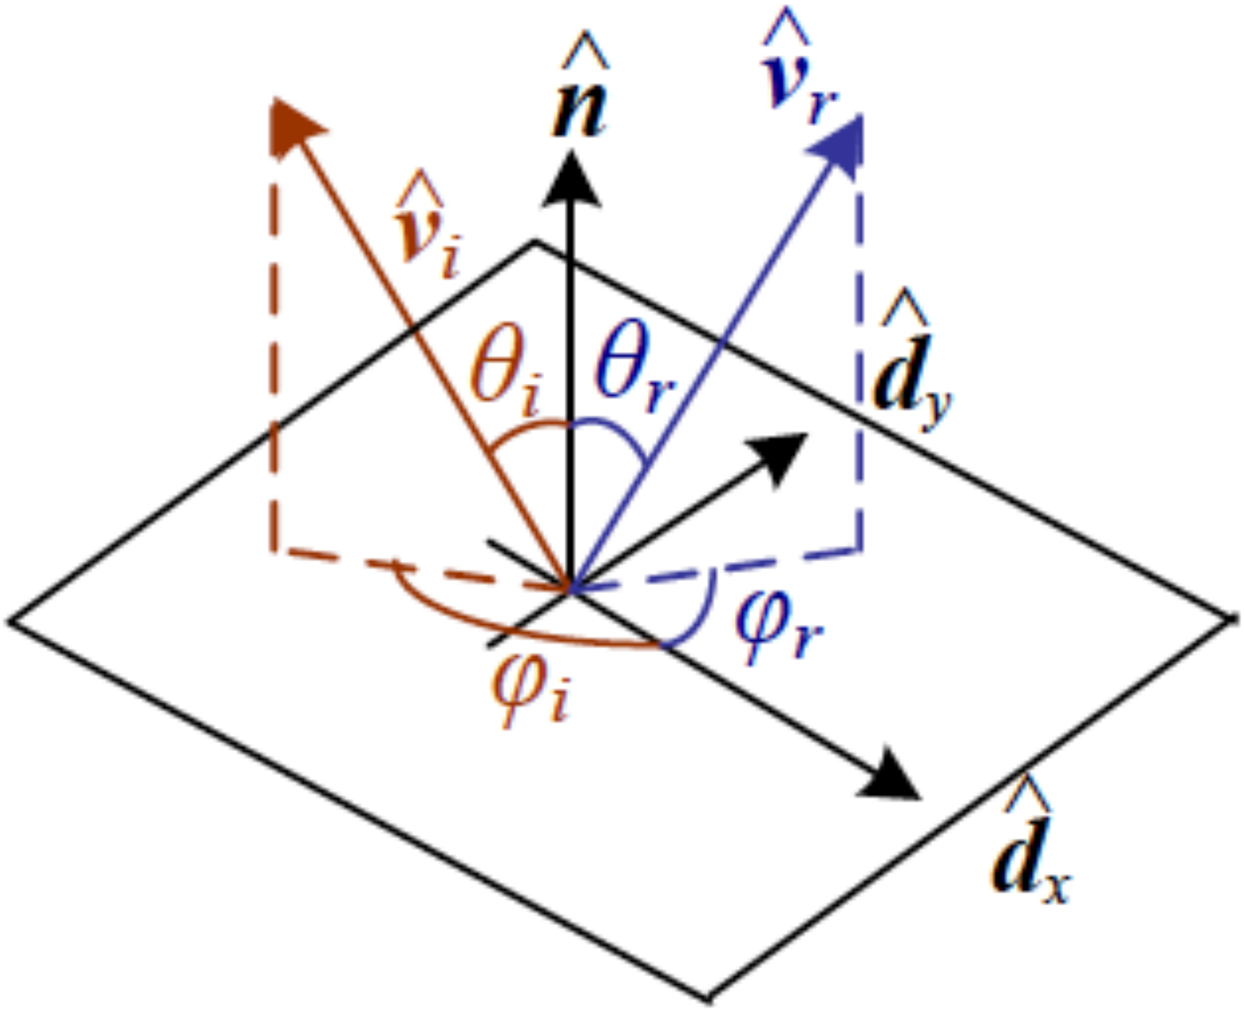

Source: Szeliski book.

The Bidirectional Reflectance Distribution Function

- A general model of light scattering

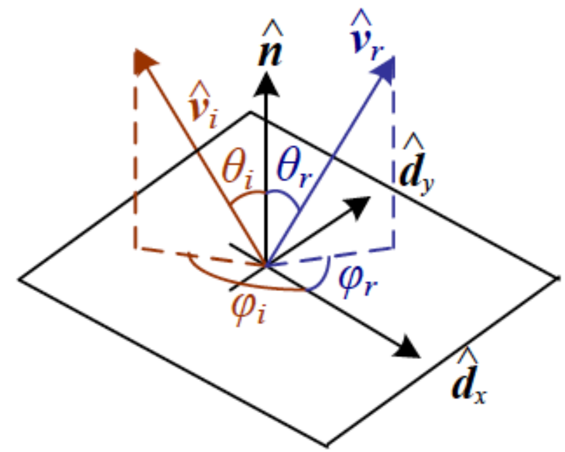

$$f_r(\boxed{\theta_i, \varphi_i}, \boxed{\theta_r, \varphi_r}; \boxed{\lambda})$$

Incident light parameters

Light parameter

Reflected light parameters

- Helmholtz reciprocity
- Simplified form for isotropic materials

$$f_r(\hat{\boldsymbol{v}}_i, \hat{\boldsymbol{v}}_r, \hat{\boldsymbol{n}}; \lambda)$$

BRDF

- Isotropic vs. Anisotropic

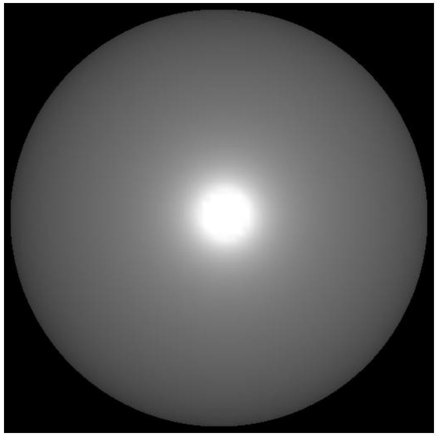

Isotropic

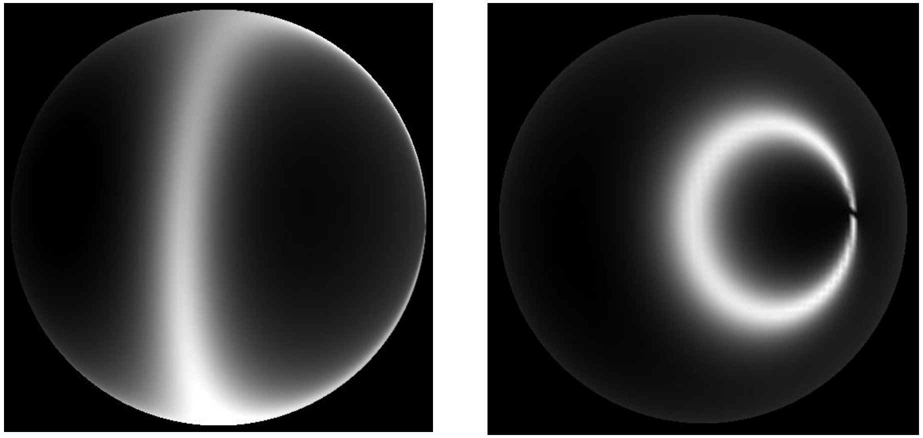

Anisotropic

Source: Hanrahan slides.

BRDF

- Light exiting a surface point in a direction under a given lighting condition:

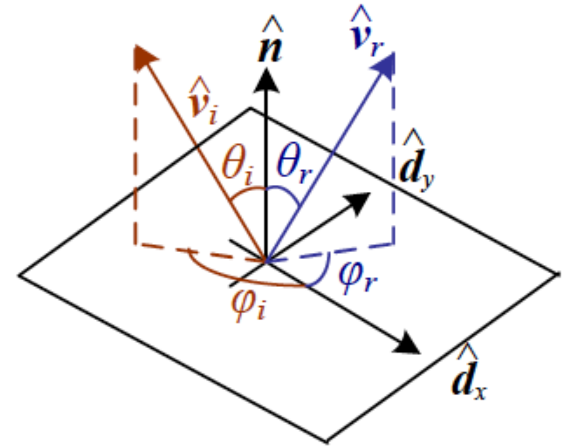

$$L(\hat{\boldsymbol{v}}_r;\lambda) = \int L_i(\hat{\boldsymbol{v}}_i;\lambda) f_r(\hat{\boldsymbol{v}}_i, \hat{\boldsymbol{v}}_r, \hat{\boldsymbol{n}};\lambda) \boxed{\cos^+\theta_i} d\hat{\boldsymbol{v}}_i$$

$$\cos^+\theta_i = \max(0, \cos\theta_i)$$

Foreshortening effect due to surface orientation

BRDF

- Understanding and modeling the BRDF is critical to realism in graphics as well as various computer vision applications.

- http://www.disneyanimation.com/technology/brdf.html

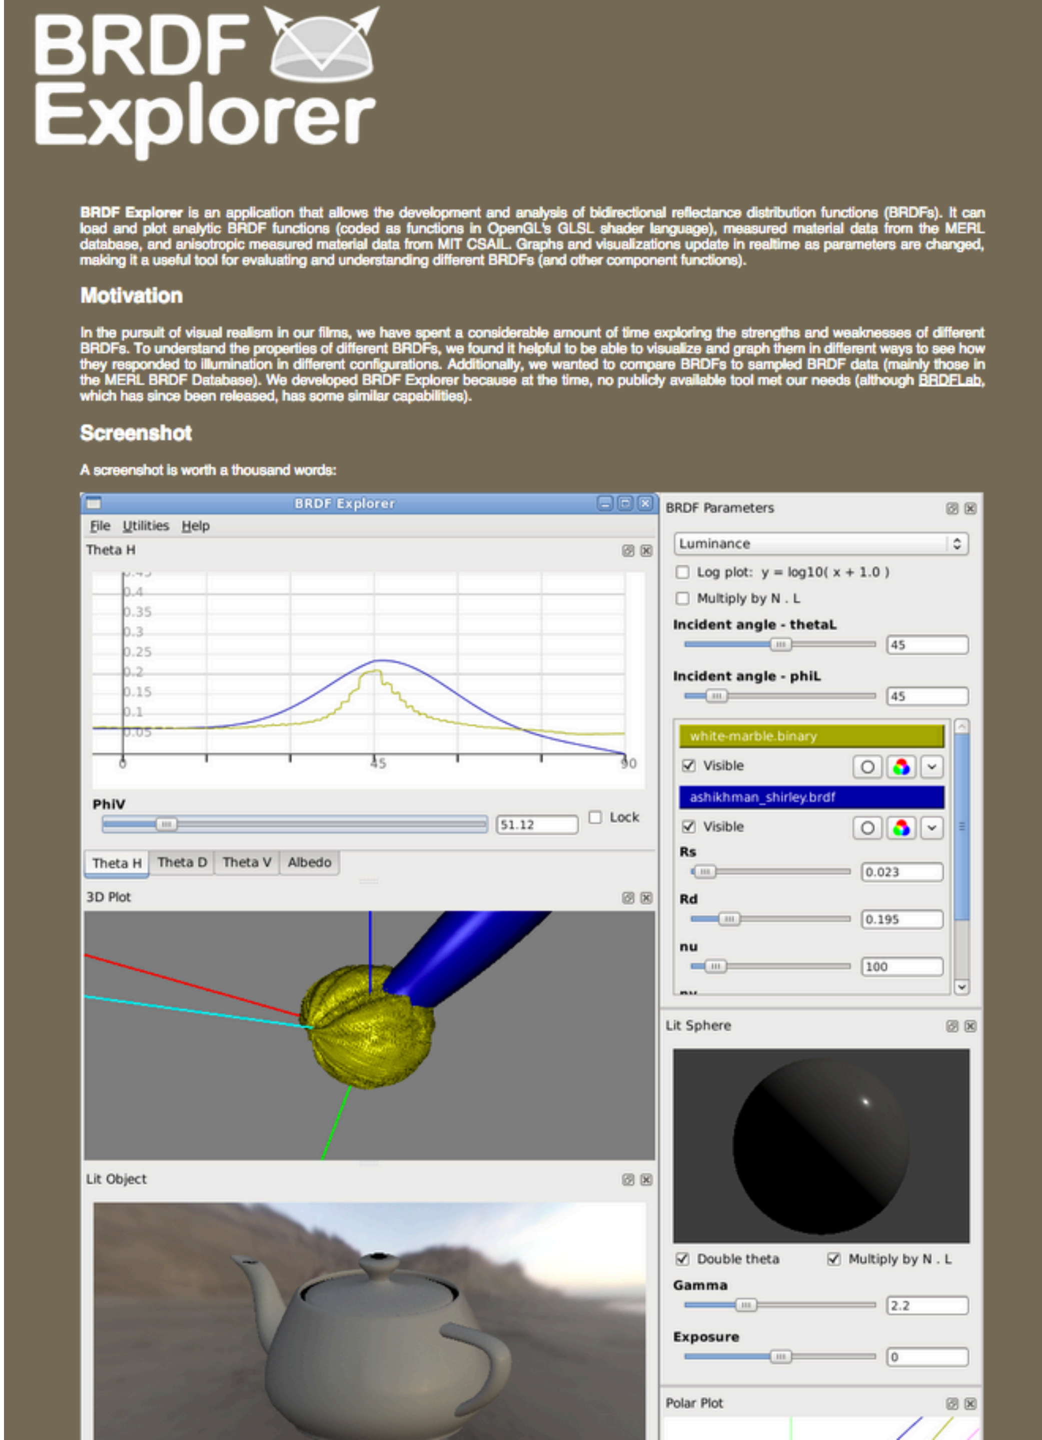

Diffuse or Lambertian Reflection

- Light is scattered uniformly in all directions.

$$f_d(\hat{\boldsymbol{v}}_i, \hat{\boldsymbol{r}}_i, \hat{\boldsymbol{n}}; \lambda) = f_d(\lambda)$$

- The amount of light depends on the angle between the incident light direction and the surface normal.
  - Lambert's cosine law

$$\begin{aligned} L_d(\hat{\boldsymbol{v}}_r; \lambda) &= \sum_i L_i(\lambda) f_d(\lambda) \cos^+ \theta_i \\ &= \sum_i L_i(\lambda) f_d(\lambda) \left[\hat{\boldsymbol{v}}_i^\mathsf{T} \hat{\boldsymbol{n}}\right]^+ \end{aligned}$$

Diffuse Reflection Examples

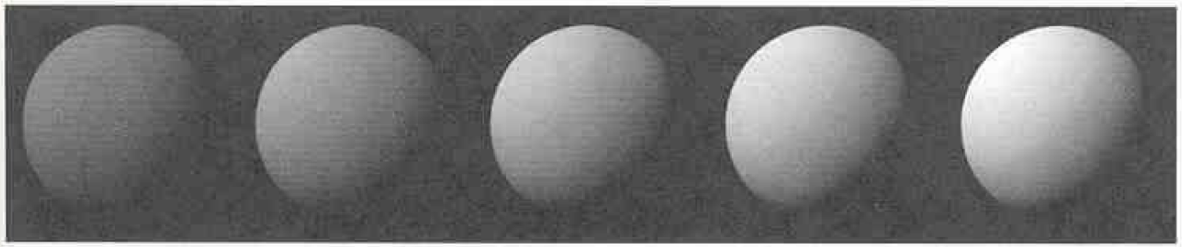

Fig. 16.3 Spheres shaded using a diffuse-reflection model (Eq. 16.4). For all spheres, $I_p = 1.0$. From left to right, $k_d = 0.4, 0.55, 0.7, 0.85, 1.0$. (By David Kurlander, Columbia University.)

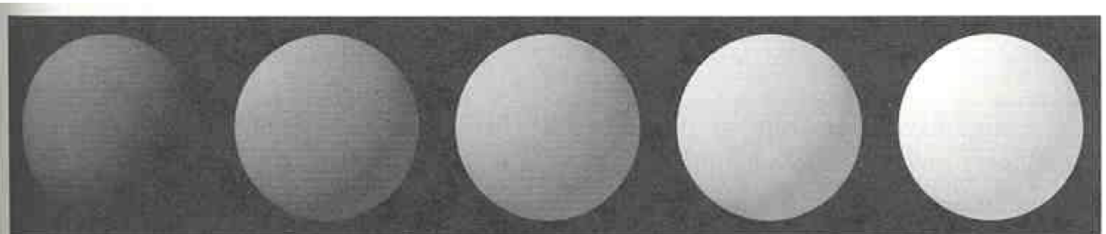

Fig. 16.4 Spheres shaded using ambient and diffuse reflection (Eq. 16.5). For all spheres, $I_a = I_p = 1.0$, $k_d = 0.4$. From left to right, $k_a = 0.0, 0.15, 0.30, 0.45, 0.60$. (By David Kurlander, Columbia University.)

***From Foley, vanDam, Feiner, and Hughes,* Computer Graphics: Principles and Practice, 2nd edition, page 725**

Diffuse Vs. Specularity

Source: Paolina Image from Caltech Vision Lab.

Specular (or Mirror) Reflections

- Specularity depends strongly on the direction of the outgoing light.
- Mirror-like reflection: incoming light is reflected off the surface in a single direction (which is the rotation of 180 degree around the surface normal).

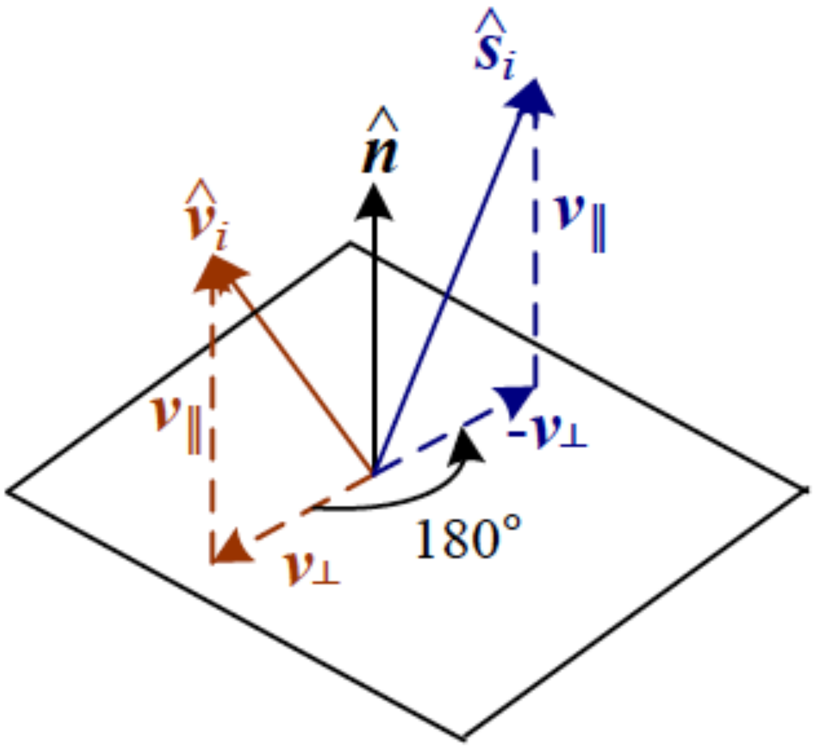

Phong Shading Model

Ambient Light

Diffuse Reflectance

Specular Reflectance

$$L_{\text{Phong}}(\hat{\boldsymbol{v}}_r;\lambda) = \boxed{k_a(\lambda)L_a(\lambda)} + \boxed{k_d(\lambda)\sum_i L_i(\lambda) f_d(\lambda)\left[\hat{\boldsymbol{v}}_i^{\mathsf{T}}\hat{\boldsymbol{n}}\right]^+} + \boxed{k_s(\lambda)\sum_i L_i(\lambda)\left[\hat{\boldsymbol{v}}_r^{\mathsf{T}}\hat{\boldsymbol{s}}_i\right]^{k_e}}$$

Phong Illumination Example

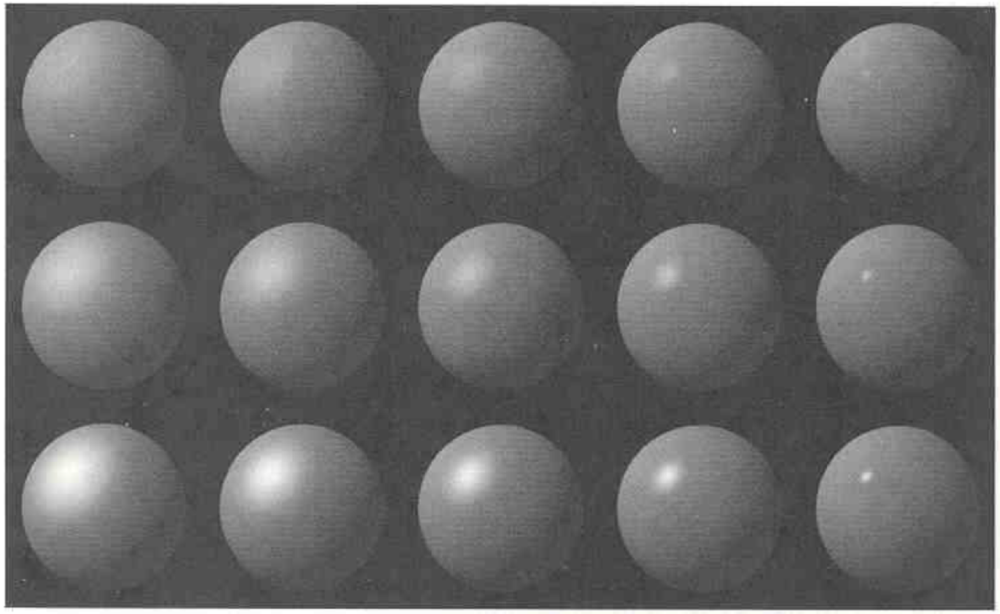

Fig. 16.10 Spheres shaded using Phong's illumination model (Eq. 16.14) and different values of k_s and n. For all spheres, $I_a = I_p = 1.0$, $k_a = 0.1$, $k_d = 0.45$. From left to right, n = 3.0, 5.0, 10.0, 27.0, 200.0. From top to bottom, k_s = 0.1, 0.25, 0.5. (By David Kurlander, Columbia University.)

From Foley, vanDam, Feiner, and Hughes, **Computer Graphics: Principles and Practice, 2nd edition, page 730**

Torrance and Sparrow Shading

- Phong shading used a power of the cosine of the angle law

$$f_s(\theta_s; \lambda) = k_s(\lambda) \cos^{k_e} \theta_s$$

- The Torrance and Sparrow model uses a Gaussian

$$f_s(\theta_s; \lambda) = k_s(\lambda) \exp\left(-c_s^2 \theta_s^2\right)$$

Phong vs. Cook/Torrance Example

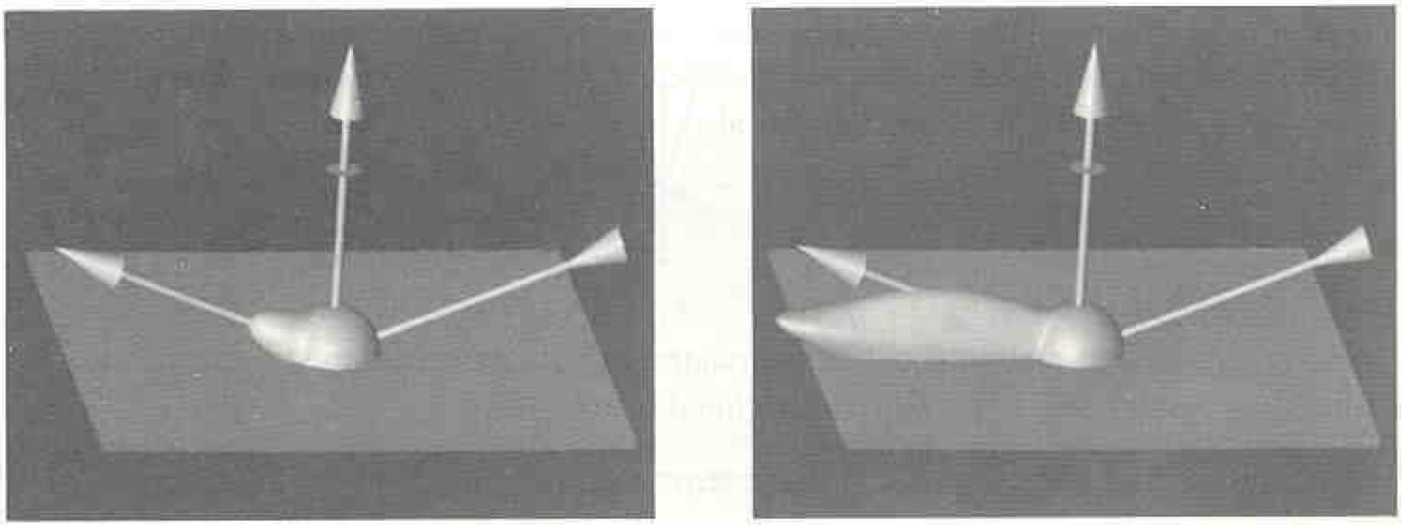

(a) Phong model (b) Torrance-Sparrow model

Fig. 16.44 Comparison of Phong and Torrance–Sparrow illumination models for light at a 70° angle of incidence. (By J. Blinn [BLIN77a], courtesy of the University of Utah.)

From Foley, vanDam, Feiner, and Hughes, **Computer Graphics: Principles and Practice, 2nd edition, page 768**

Wait a minute; light bounces…a lot

No Recursion

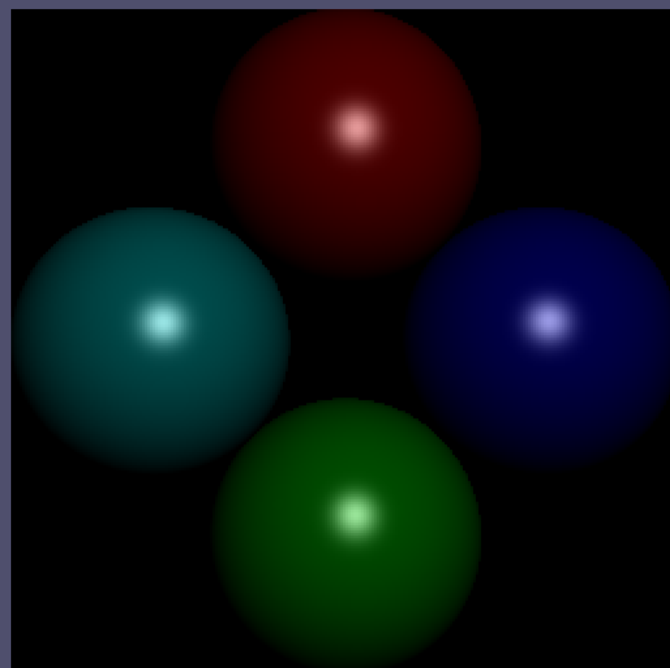

Johns Hopkins Department of Computer Science
Course 600.456: Rendering Techniques, Professor: Jonathan Cohen

Wait a minute; light bounces…a lot

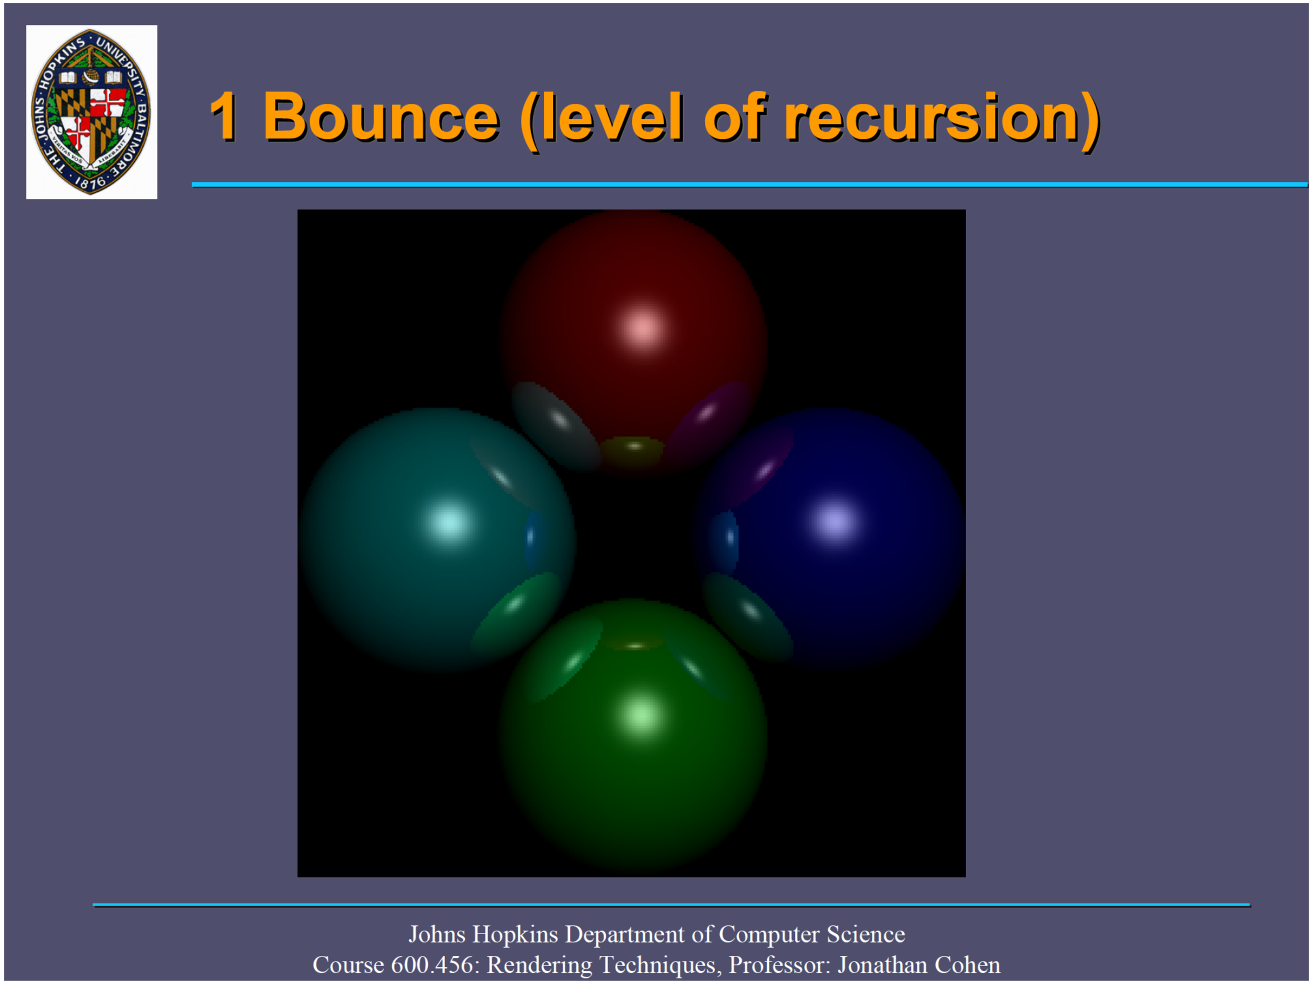

Wait a minute; light bounces…a lot

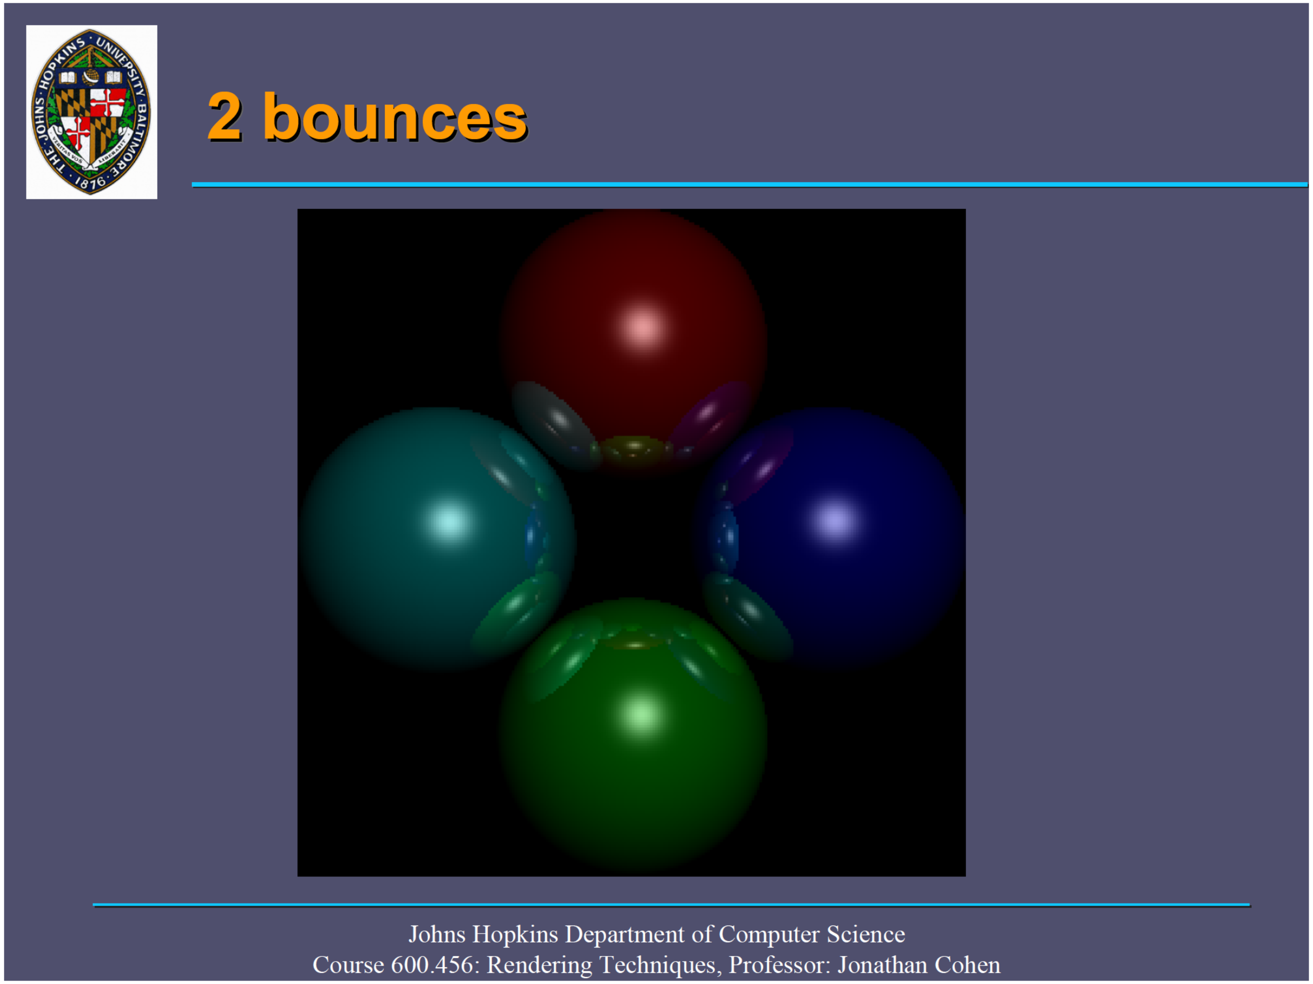

Wait a minute; light bounces…a lot

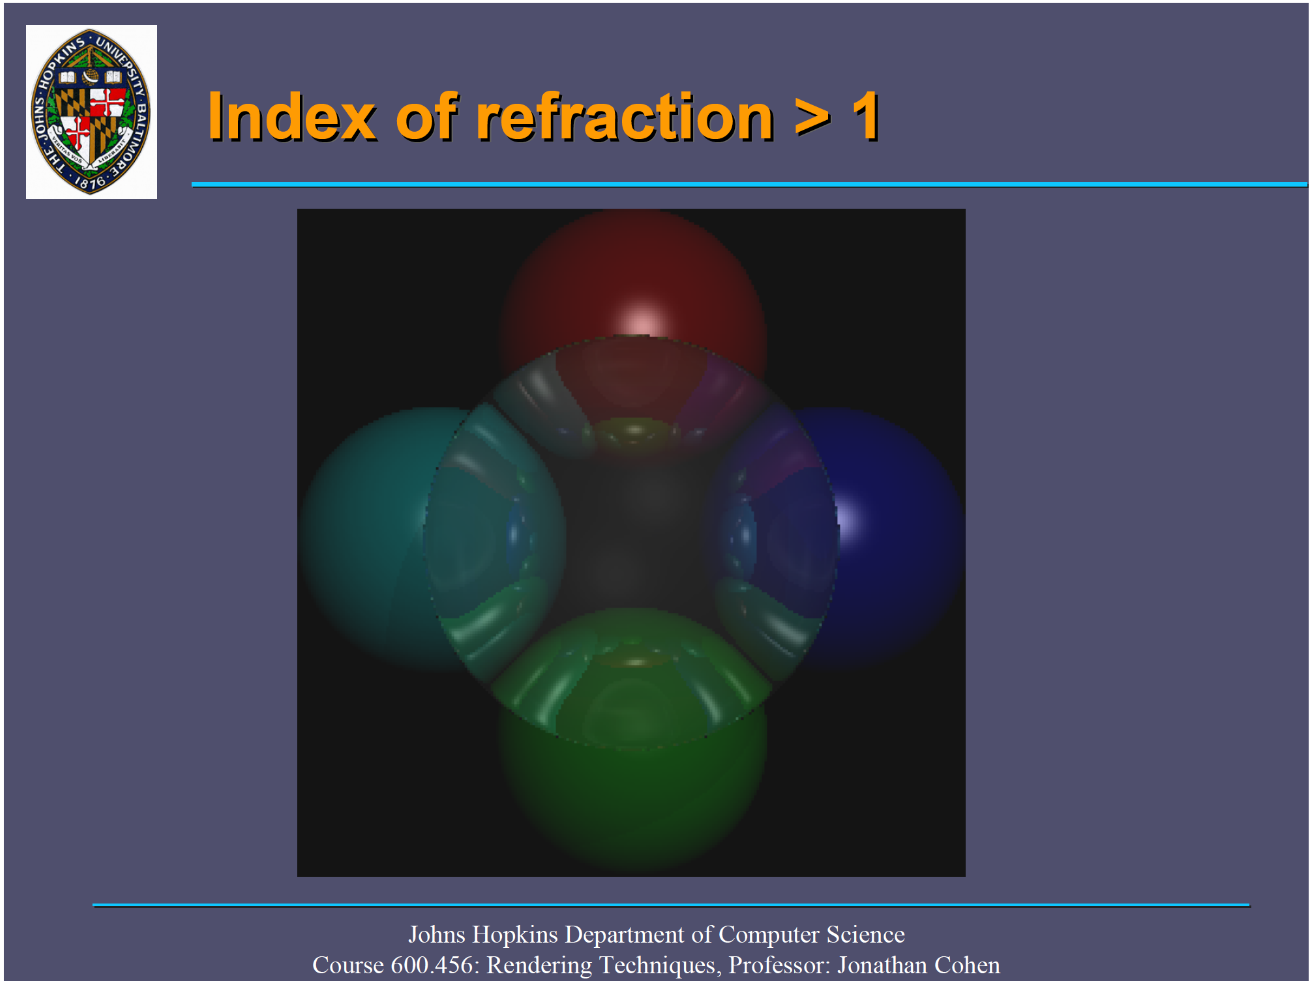

A Word On Computer-Imaging

- 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)
  - HD (= 1080x1960)
  - SHD (=1080x1960x2)

Source: G Hager Slides

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

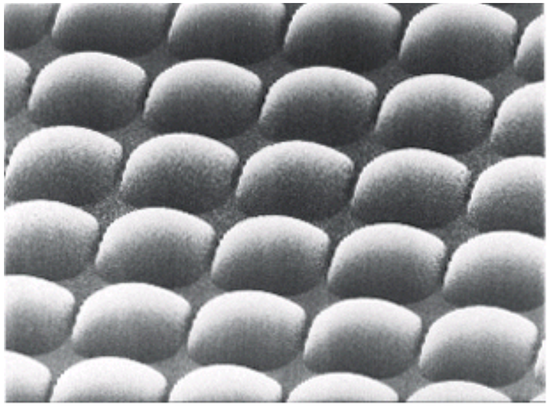

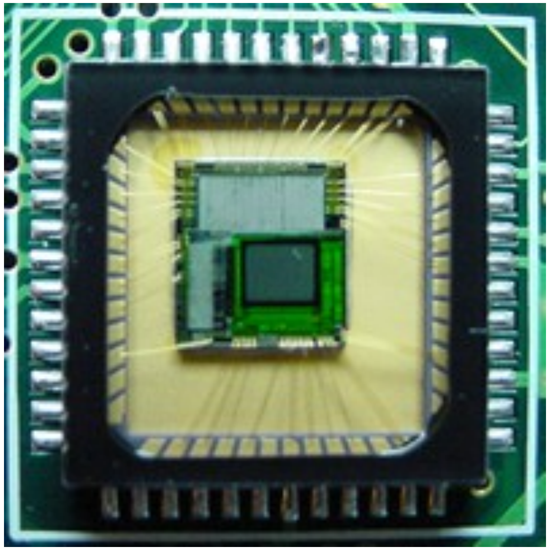

Source: G Hager Slides

A Modern Digital Camera

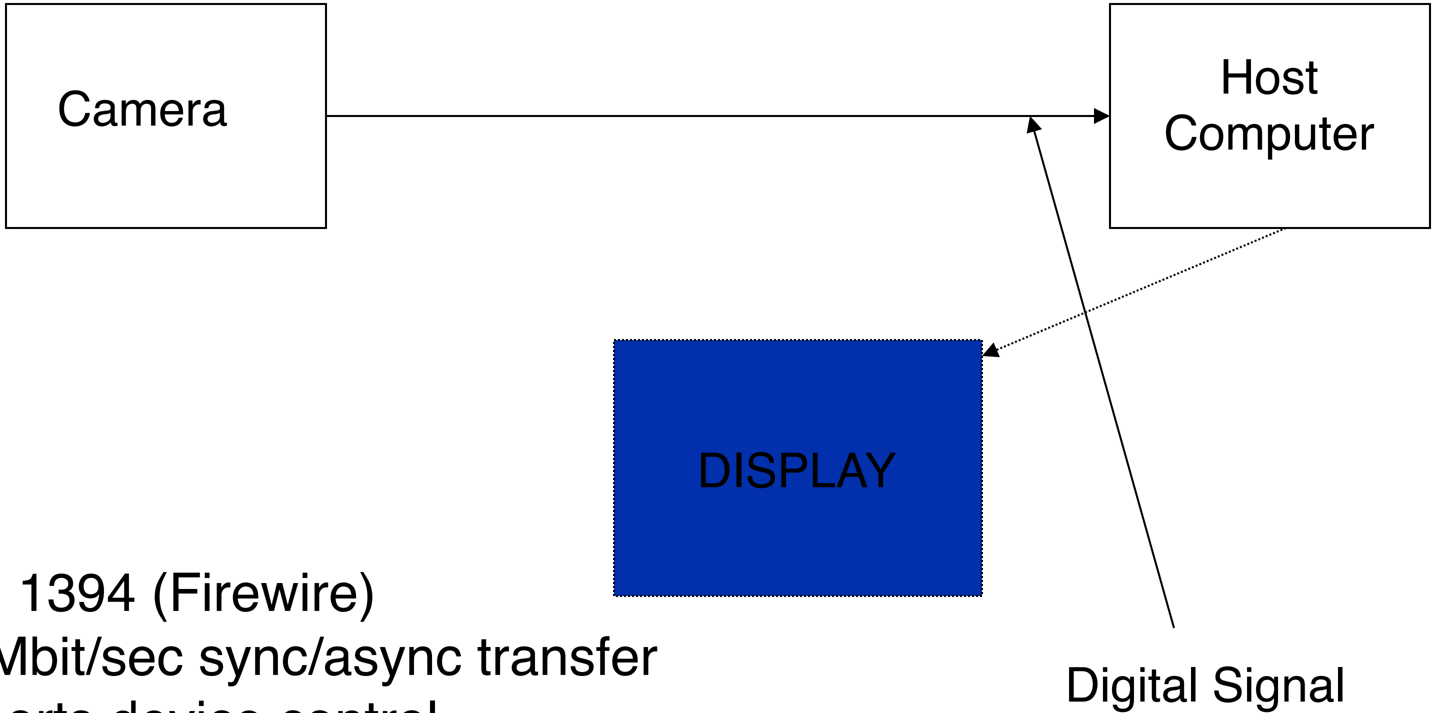

IEEE 1394 (Firewire)
400 Mbit/sec sync/async transfer
Supports device control

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

Source: G Hager Slides

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).

What’s going on here?

Source: G Hager Slides

THE ORGANIZATION OF A 2D IMAGE

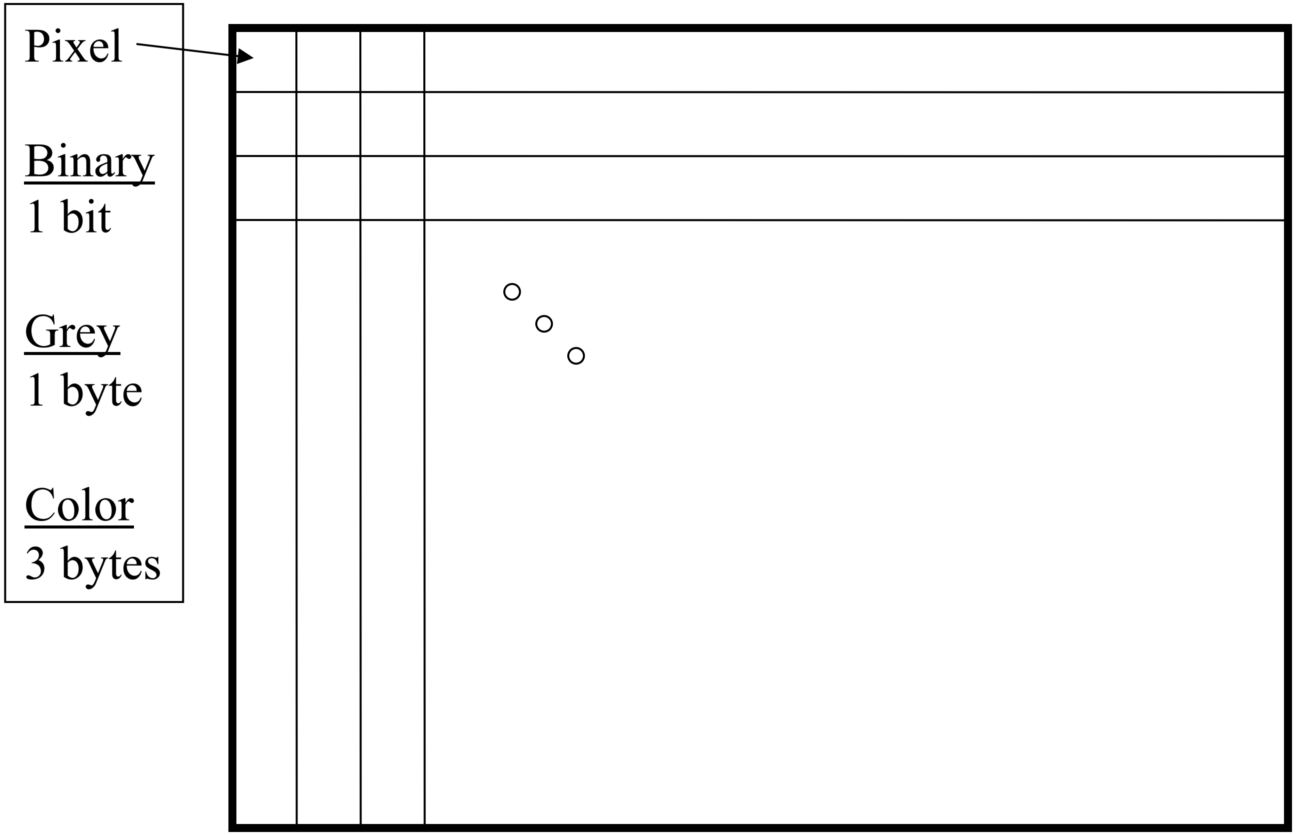

Source: G Hager Slides

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

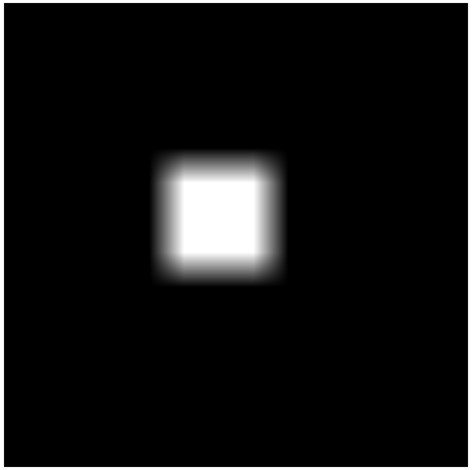

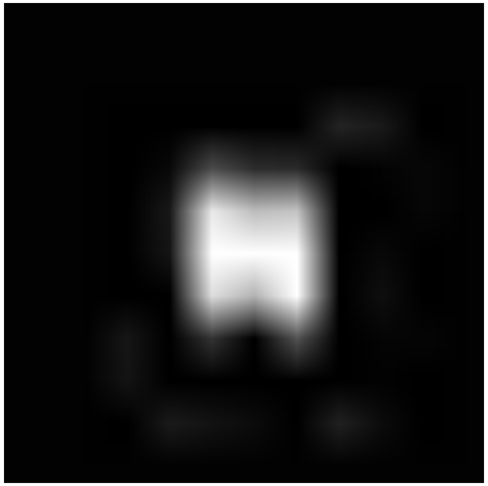

Source: G Hager Slides

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 sche
  - gif (colo
  - jpg
    - use on;
      sub
    - Use
    - Use n is less
      sen etail

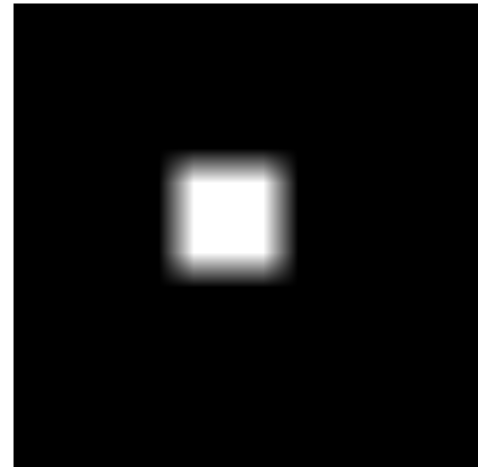

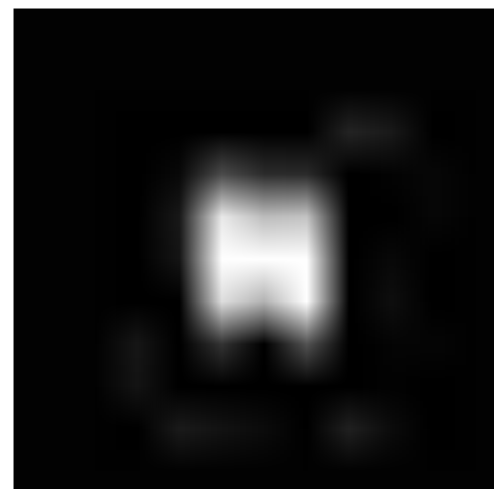

Source: G Hager Slides

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

Source: G Hager Slides

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)

Color

Source: G Hager Slides

What is Color?

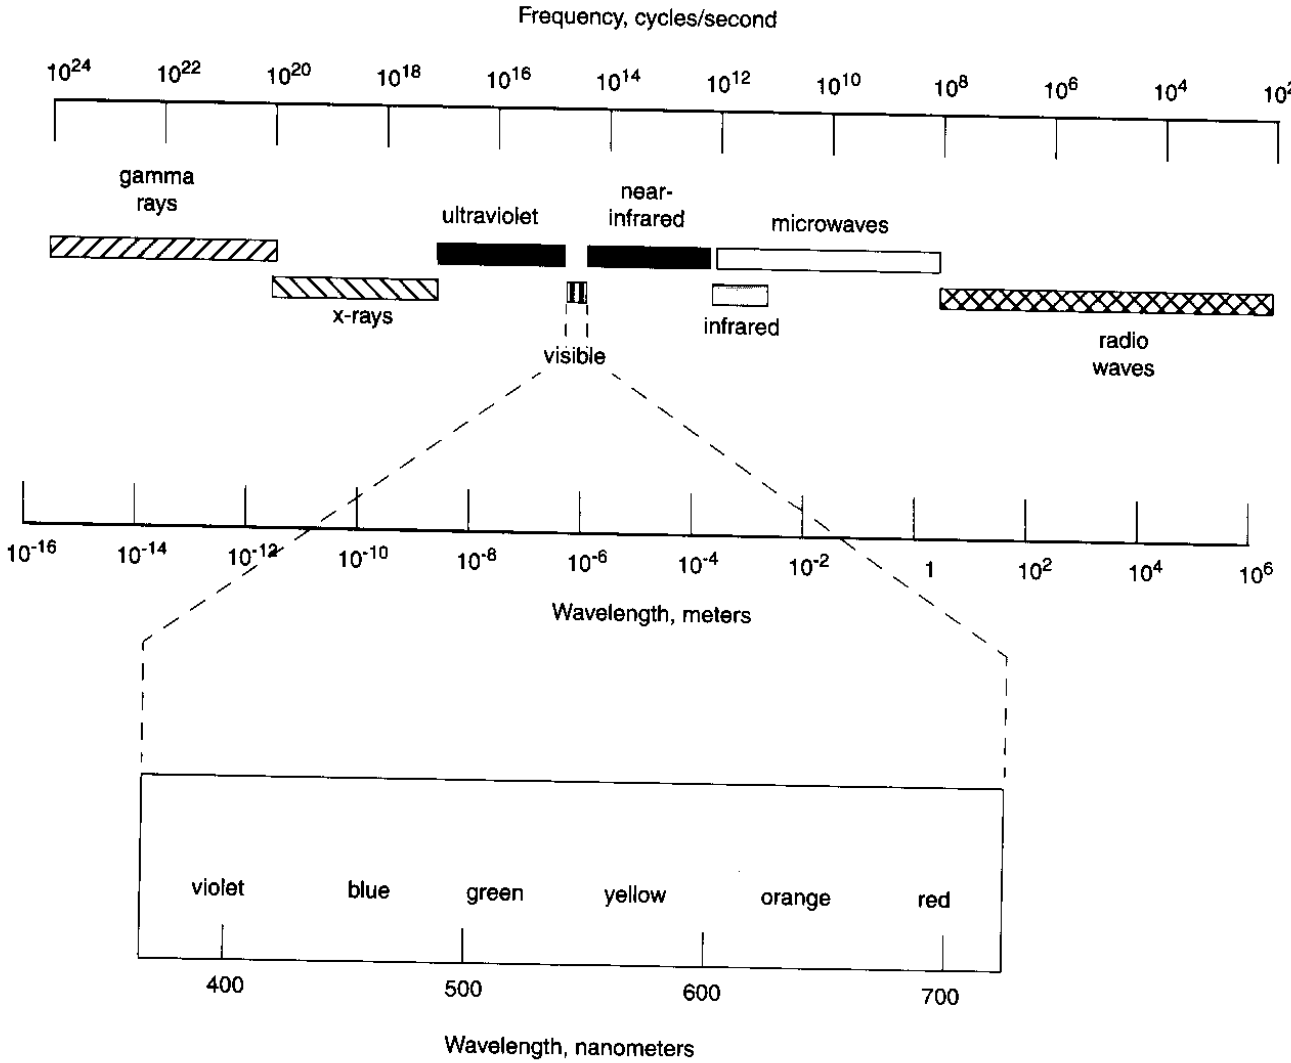

Source: G Hager Slides

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)

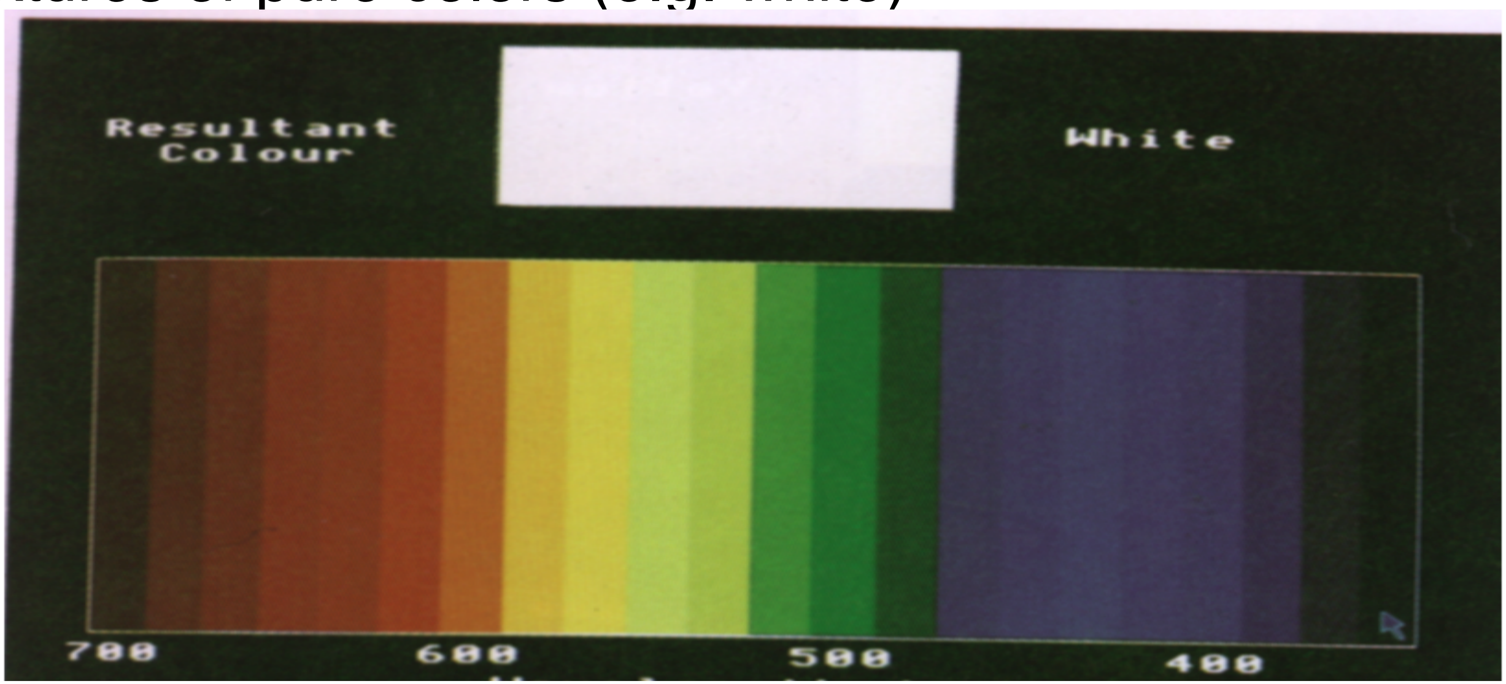

Source: G Hager Slides

Sunlight

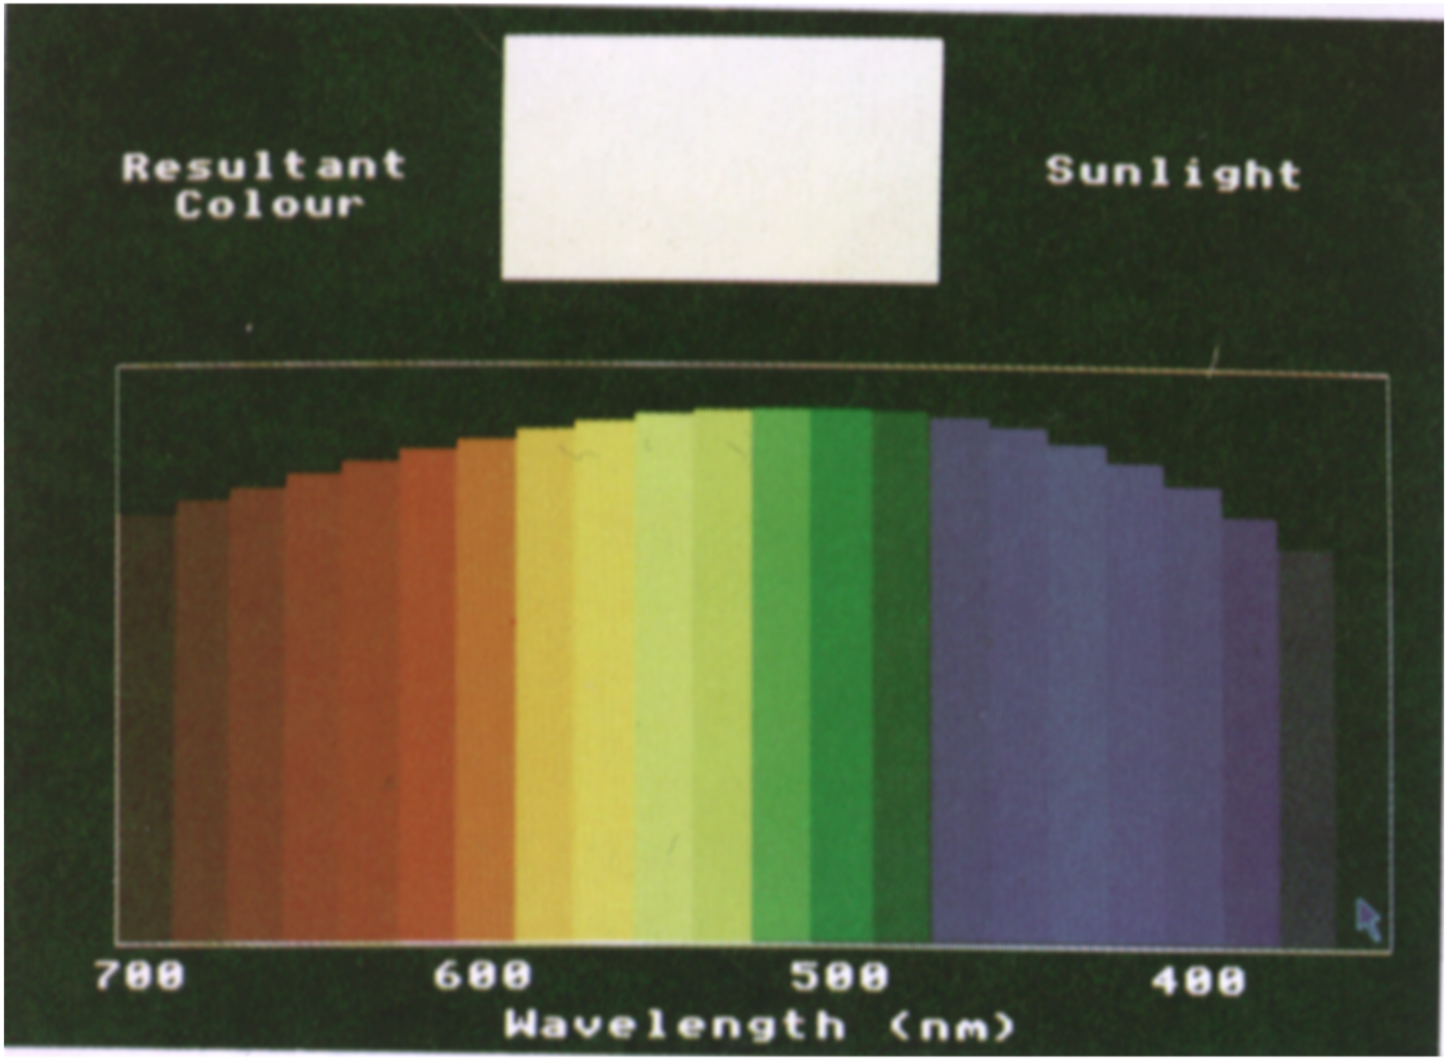

Source: G Hager Slides

Example: The Human Eye

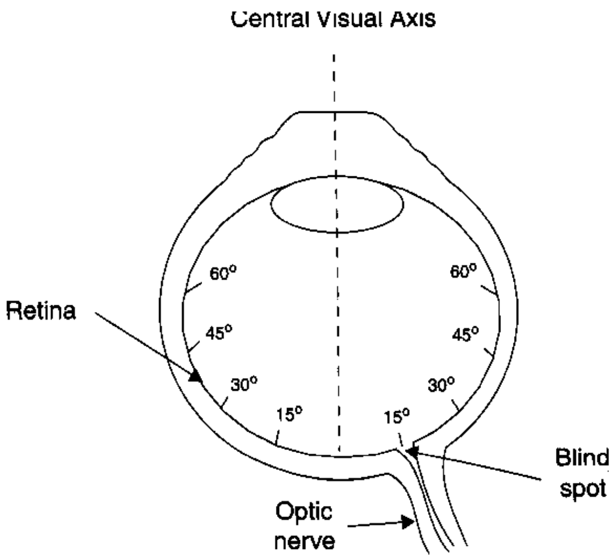

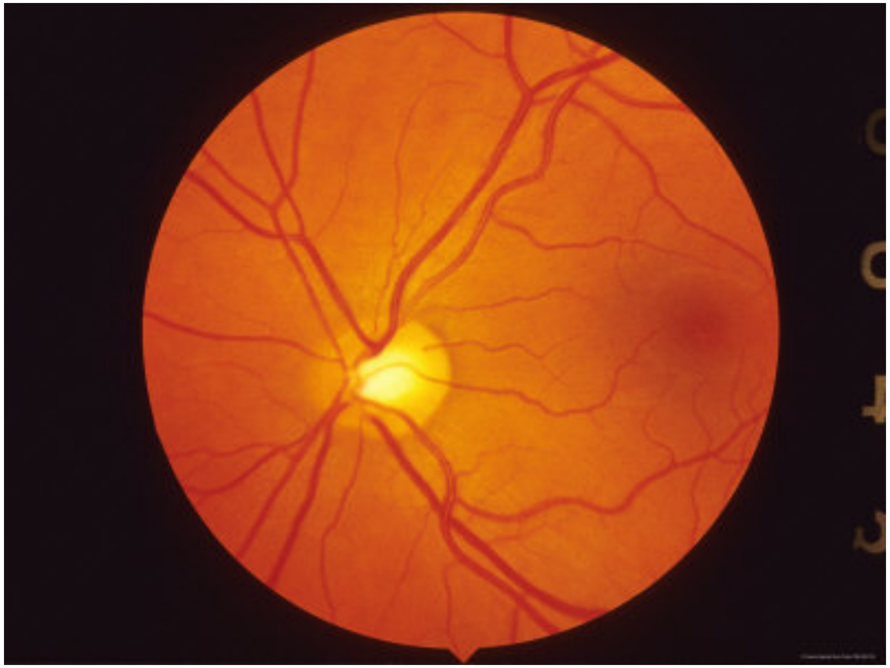

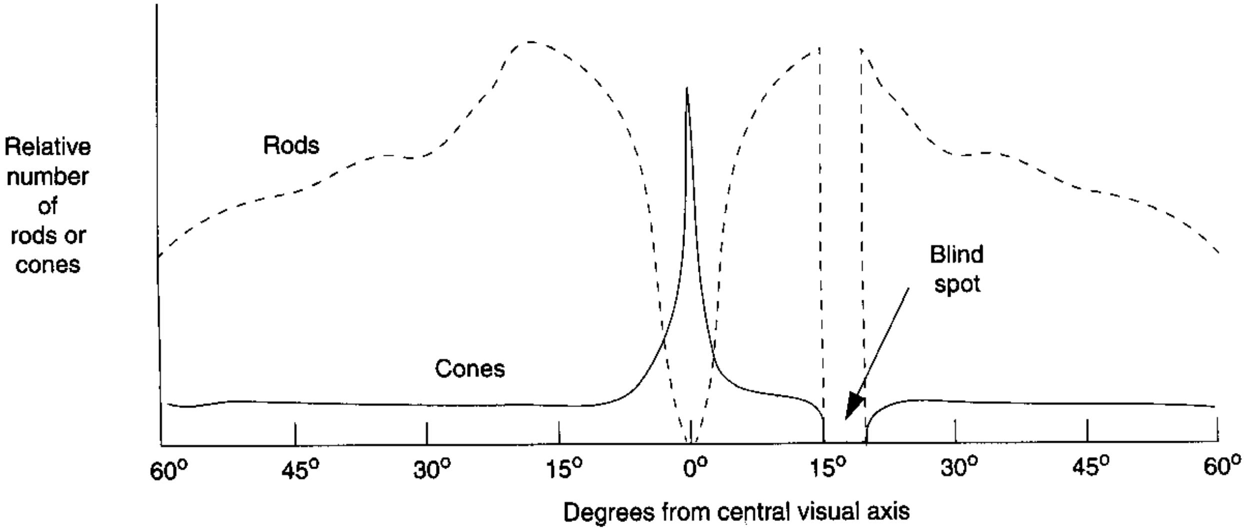

Source: G Hager Slides

The Human Eye Response

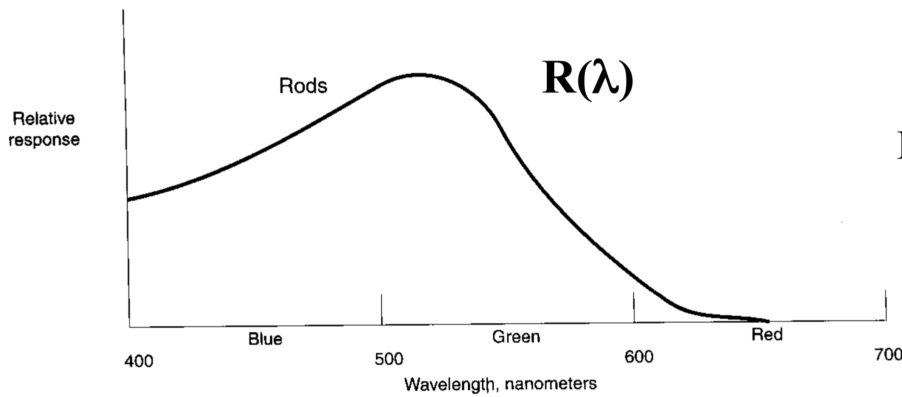

$$\text{BRIGHTNESS} = \int_{\lambda=400nm}^{\lambda=700nm} R(\lambda) I(\lambda) d\lambda$$

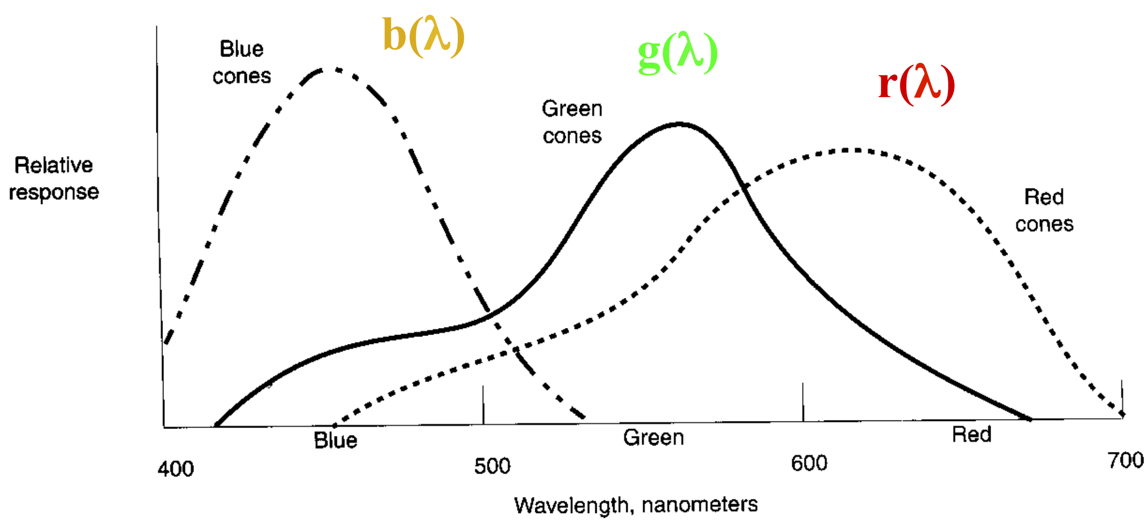

$$\text{RED} = \int_{\lambda=400nm}^{\lambda=700nm} r(\lambda) I(\lambda) d\lambda$$

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

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

Source: G Hager Slides

Color receptors

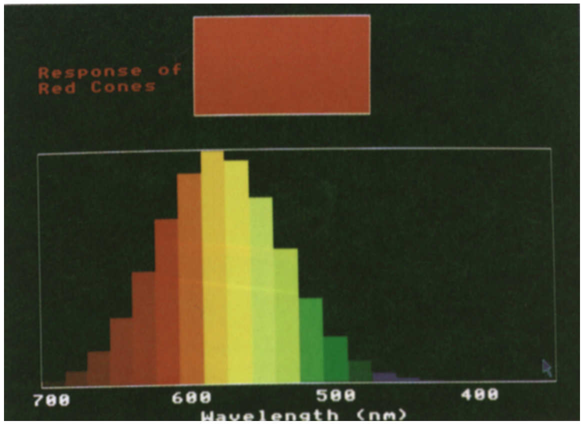

"Red" cone

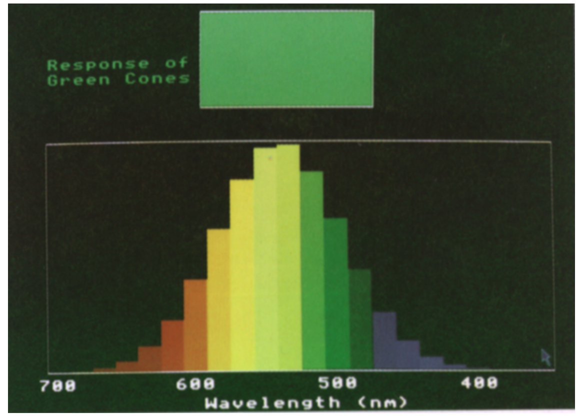

"Green" cone

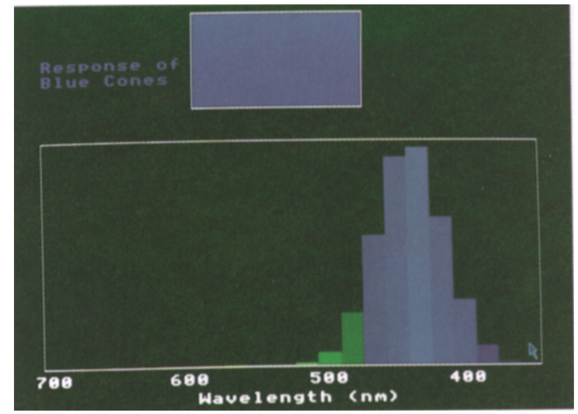

"Blue" cone

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.

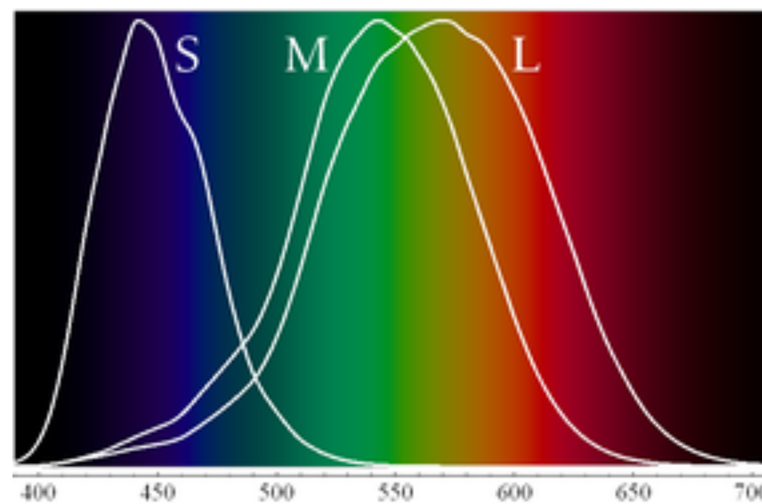

Source: G Hager Slides

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

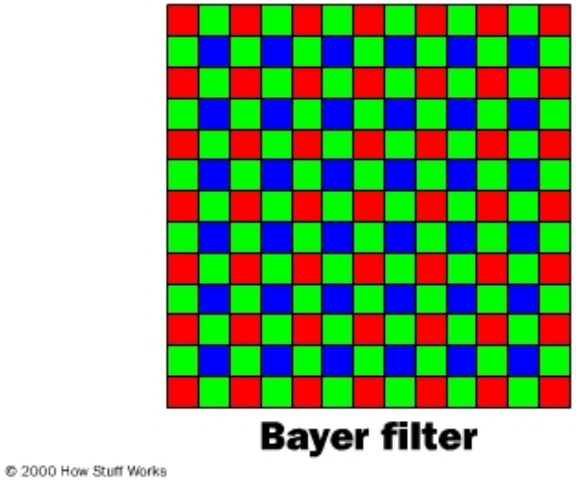

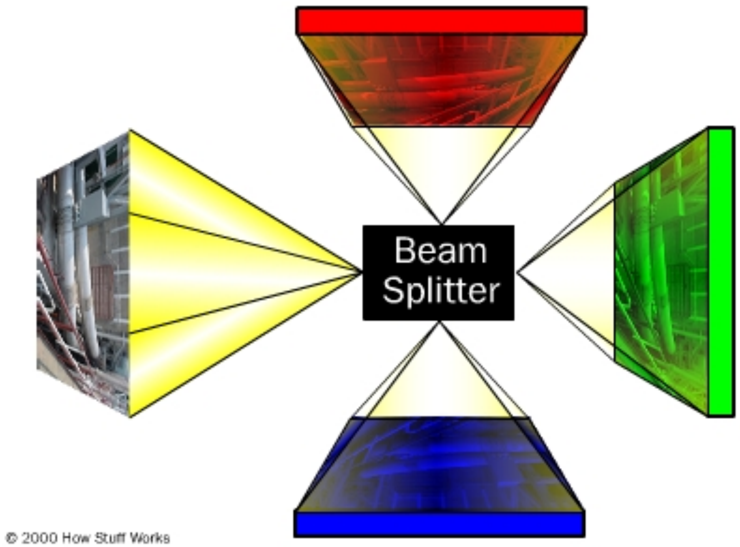

Source: G Hager Slides

Unfiltered CCD Response

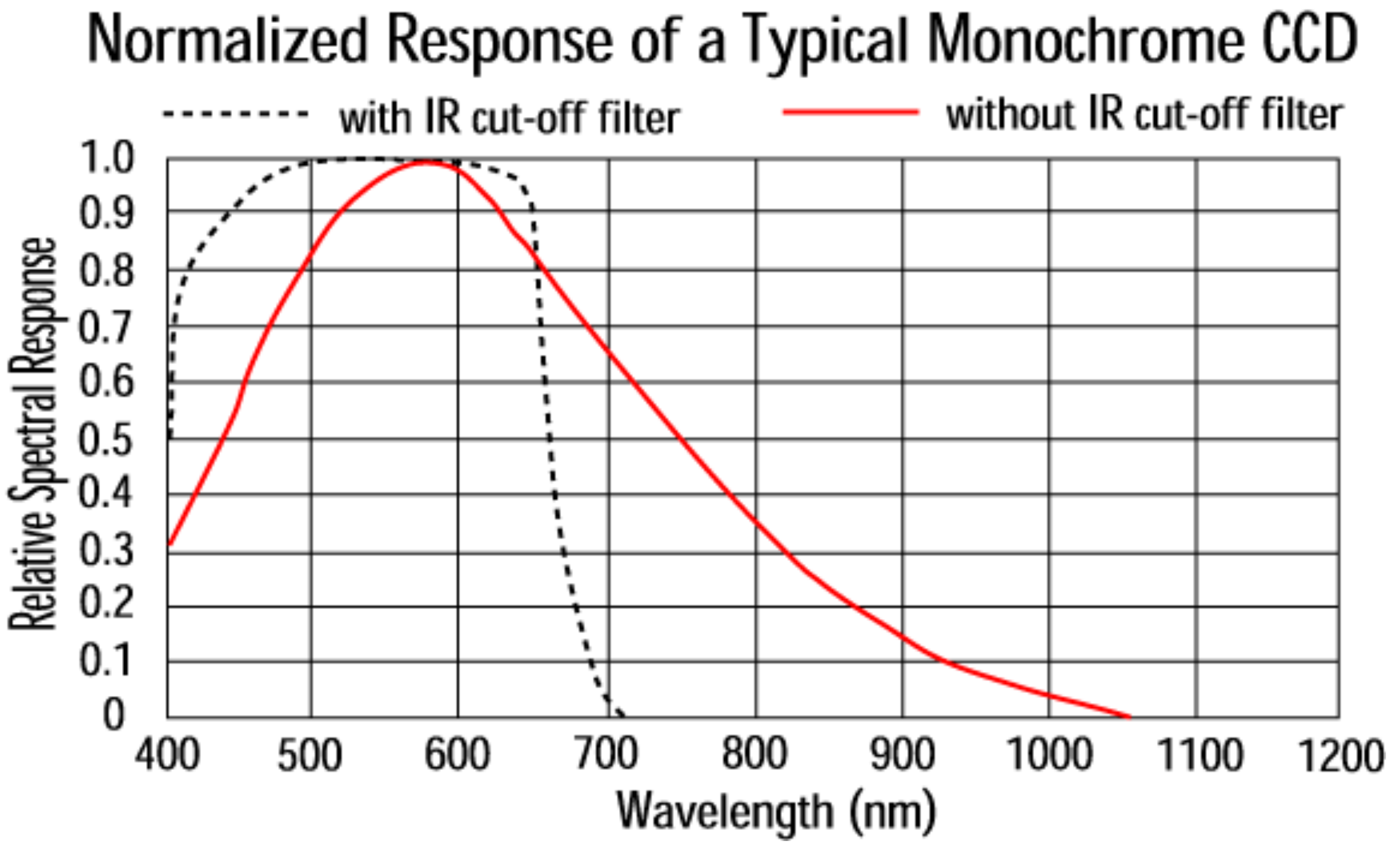

Source: G Hager Slides

One Chip CCD Response (Sony DFW V500)

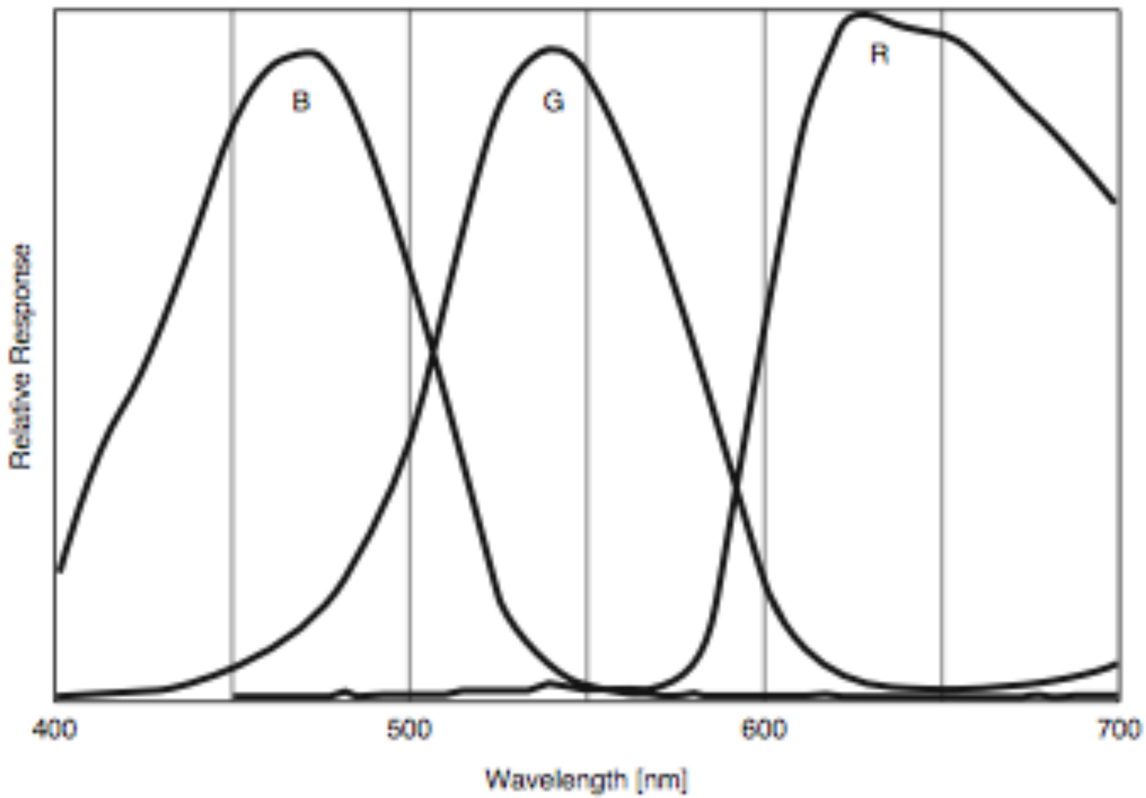

Source: G Hager Slides

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

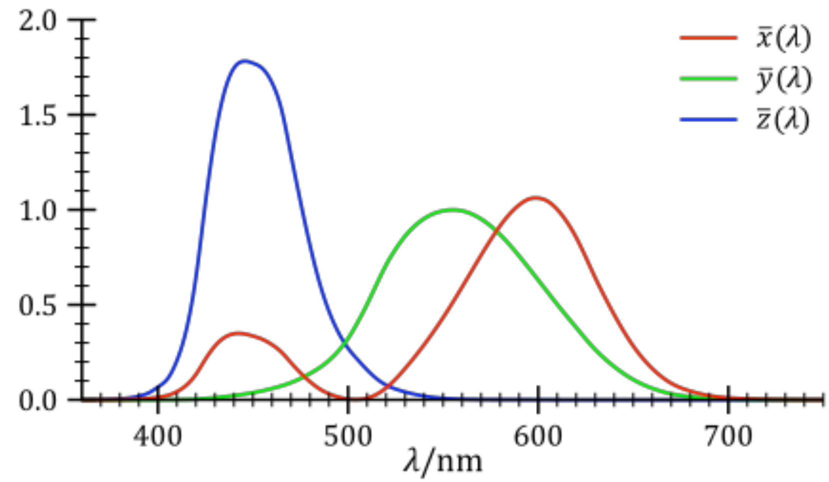

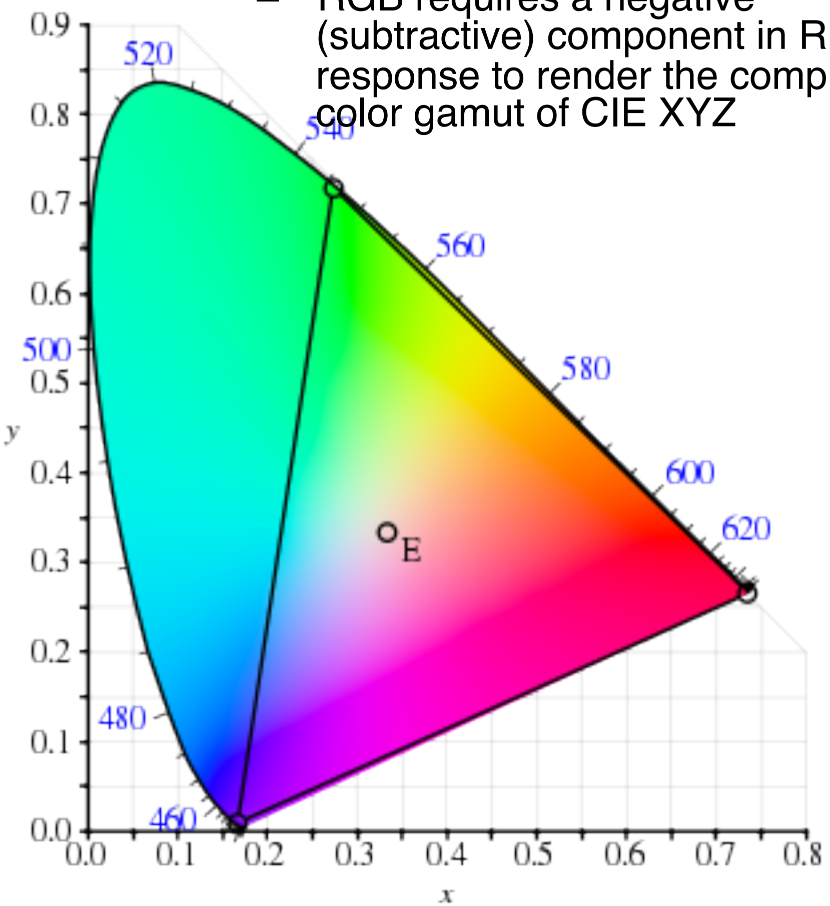

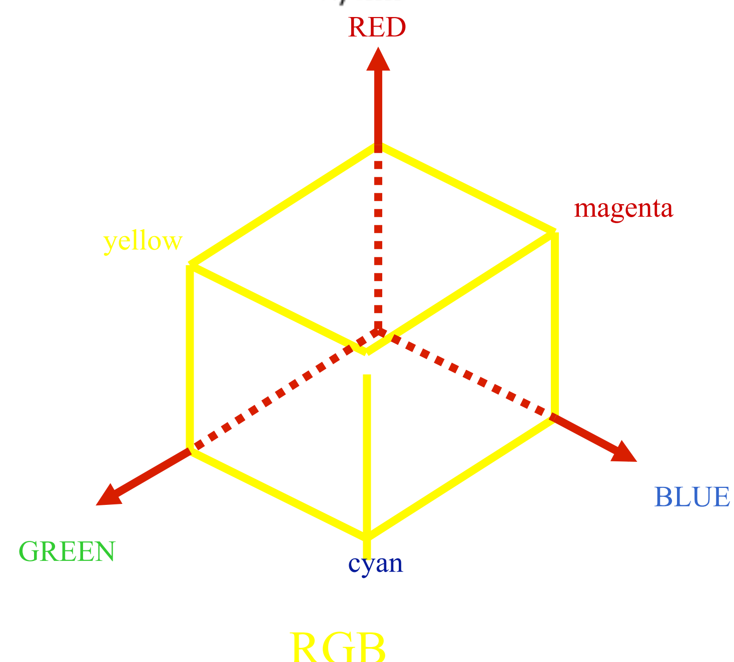

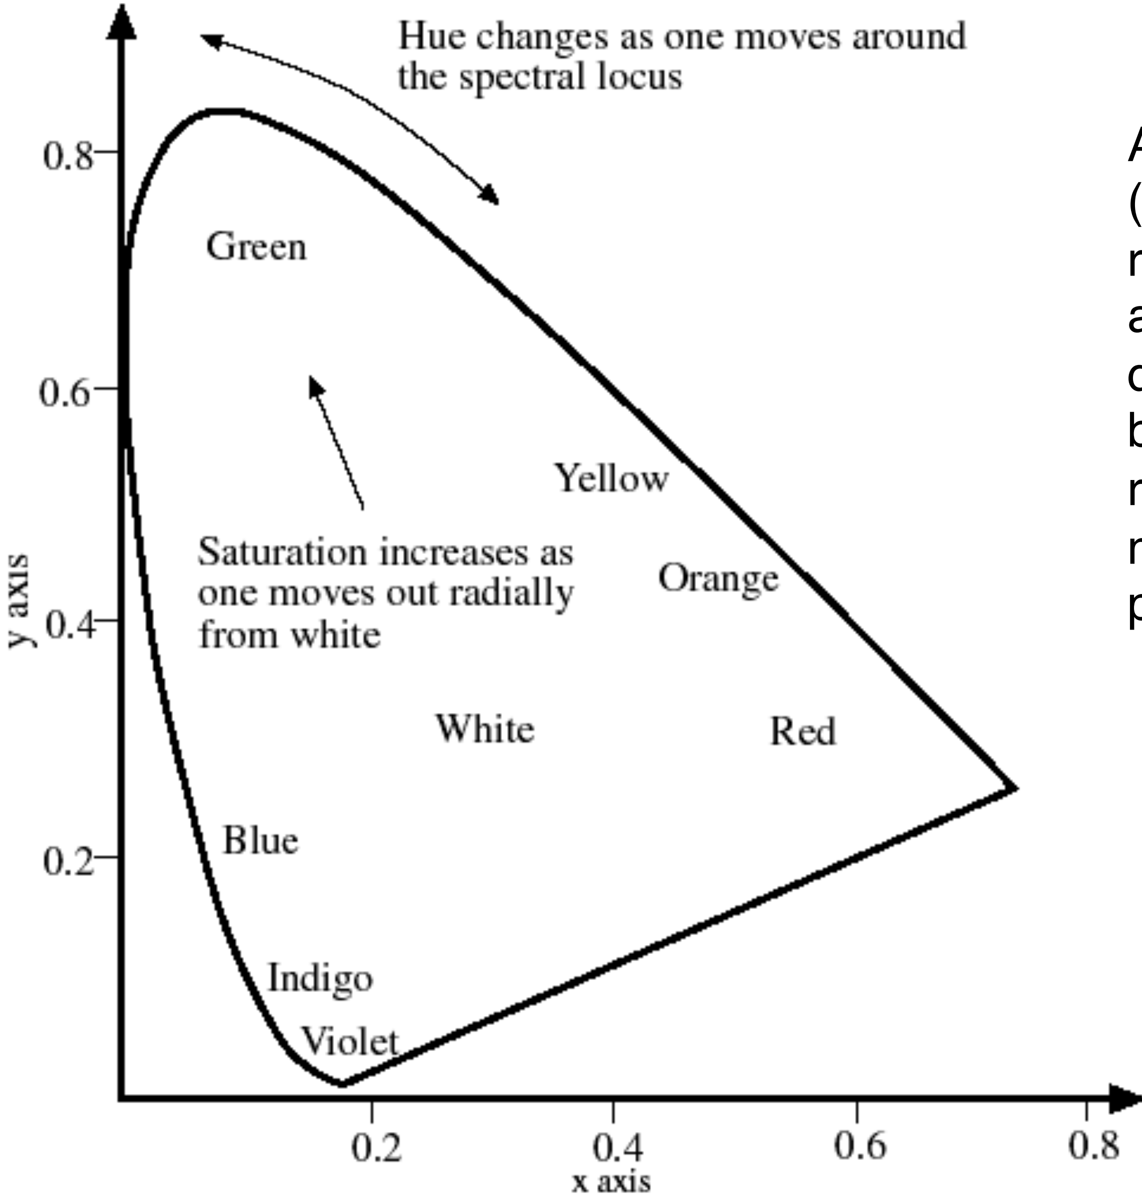

A qualitative rendering of the CIE (x,y) space. The blobby region represents visible colors. There are sets of (x, y) coordinates that don't represent real colors, because the primaries are not real lights (so that the color matching functions could be positive everywhere).

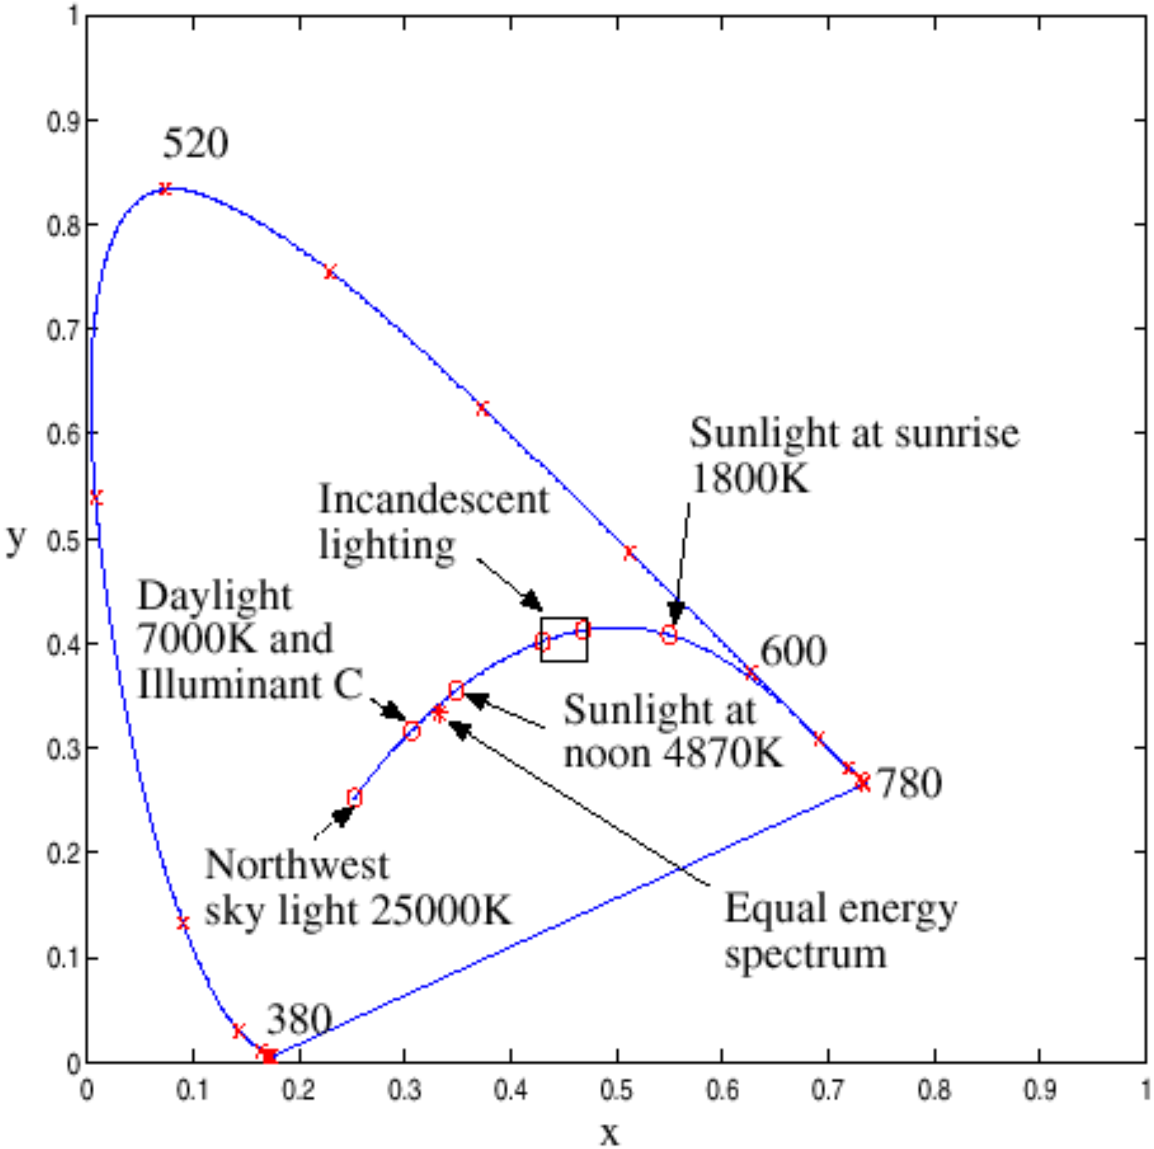

A plot of the CIE (x,y) space. We show the spectral locus (the colors of monochromatic lights) and the black-body locus (the colors of heated black-bodies). I have also plotted the range of typical incandescent lighting.

Why specify color numerically?

- Accurate color reproduction is commercially valuable
  - Many products are identified by color ("golden" arches;
- Few color names are widely recognized by English speakers -
  - About 10; other languages have fewer/more, but not many more.
  - It's common to disagree on appropriate color names.
- Color reproduction problems increased by prevalence of digital imaging - eg. digital libraries of art.
  - How do we ensure that everyone sees the same color?

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 Chrominance

Source: G Hager Slides

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
$L = .3\,R + .6G + .1B$

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

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

$x = (R-G) + (R-B)/((R-G)^2 + (R-B)(G-B))^{1/2}$

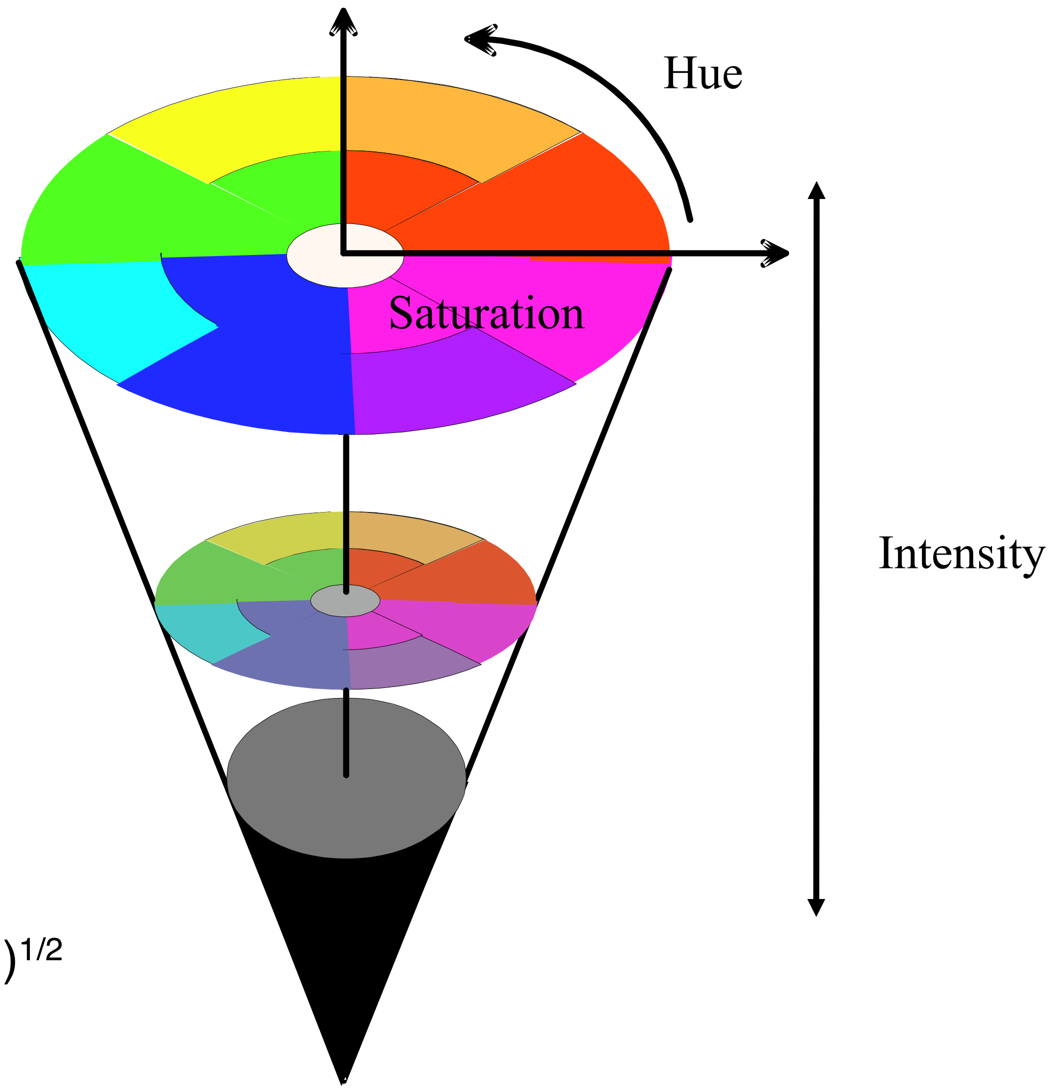

Source: G Hager Slides

HSV hexcone

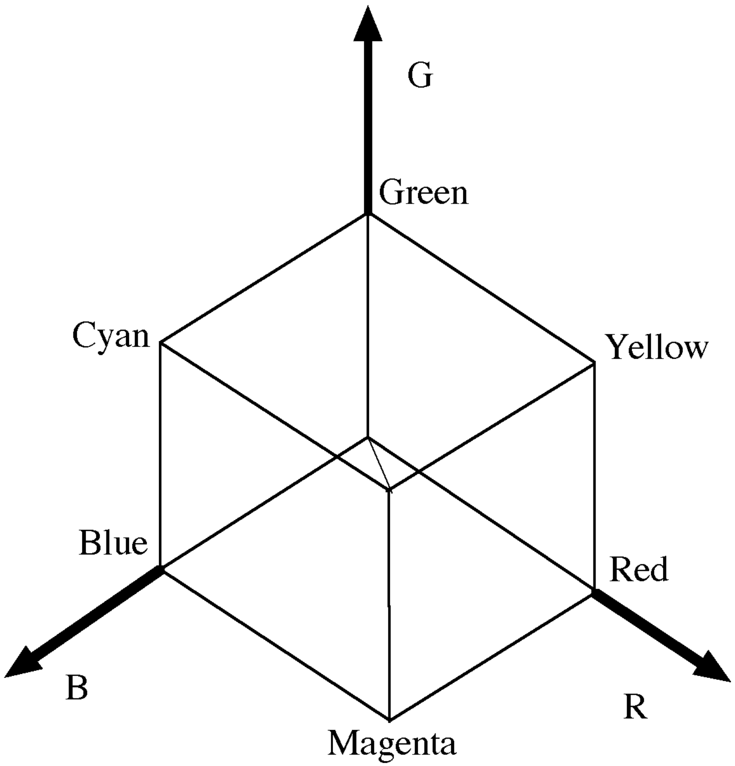
G
Green
Cyan
Yellow
Blue
Red
B
R
Magenta

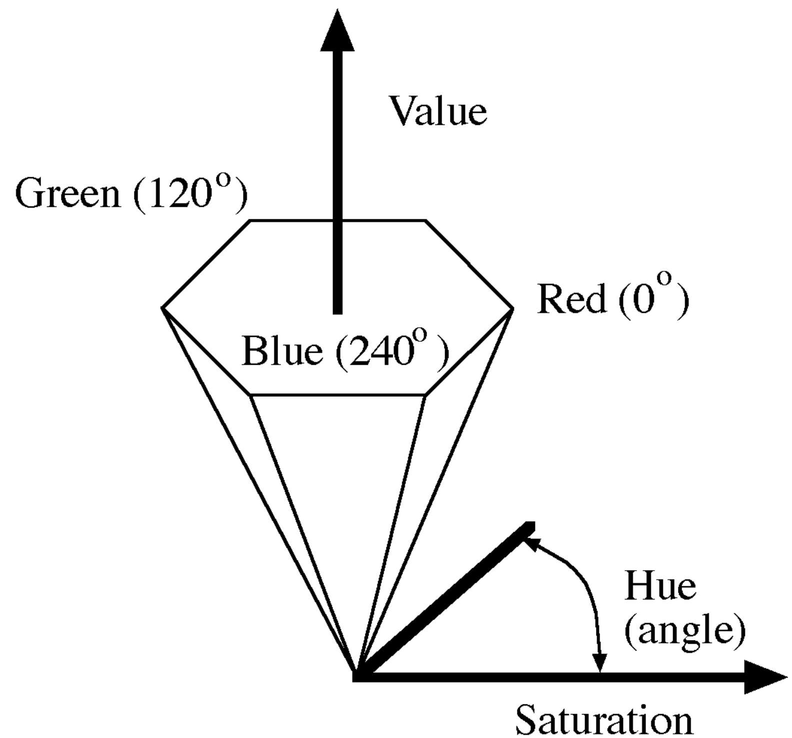
Value
Green (120°)
Red (0°)
Blue (240°)
Hue
(angle)
Saturation

Color perception...

- It’s not all physics: as the following samples show.

Brightness contrast and constancy

- The apparent brightness depends on the surrounding region
  - **brightness contrast**: a constant colored region seem lighter or darker depending on the surround:

- http://www.sandlotscience.com/Contrast/Checker_Board_2.htm

  - **brightness constancy**: a surface looks the same under widely varying lighting conditions.

Source: Szeliski Slides

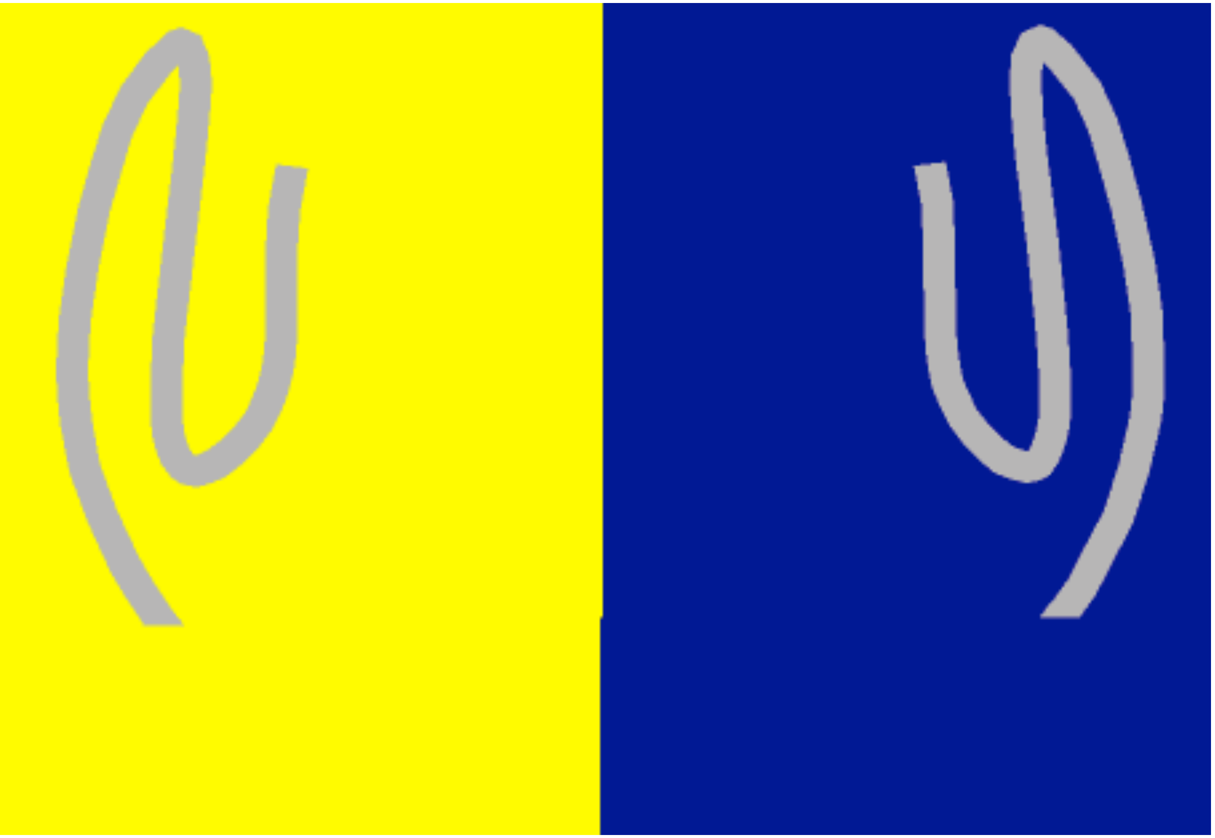

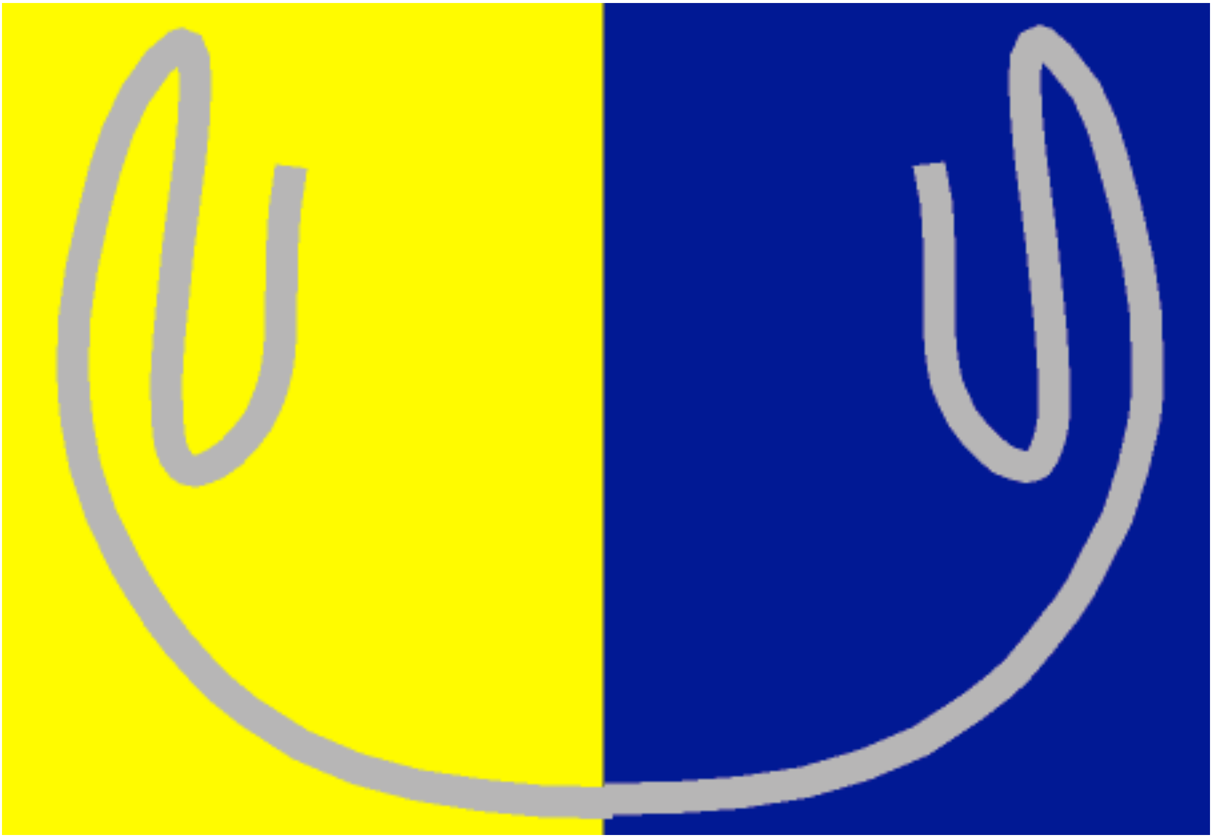

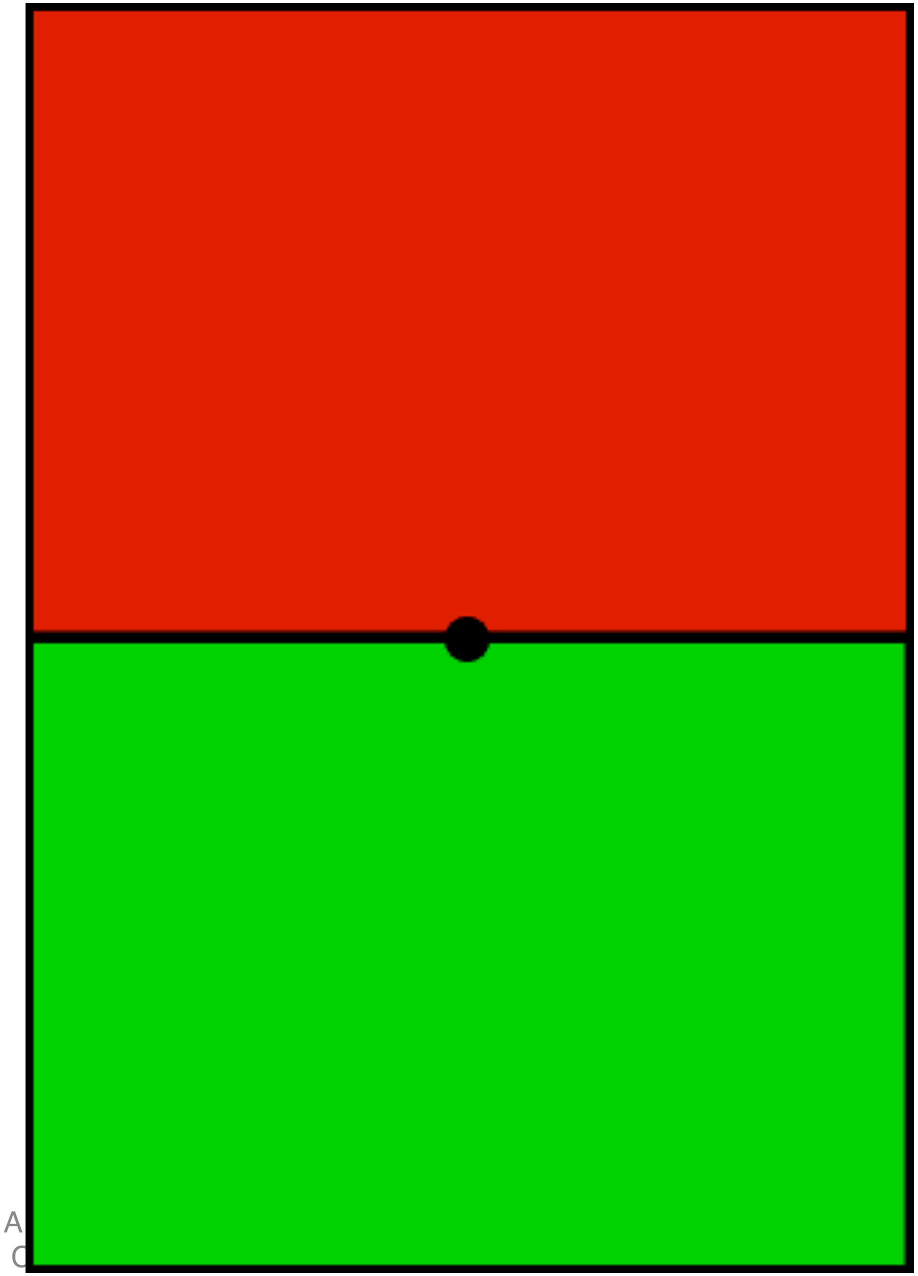

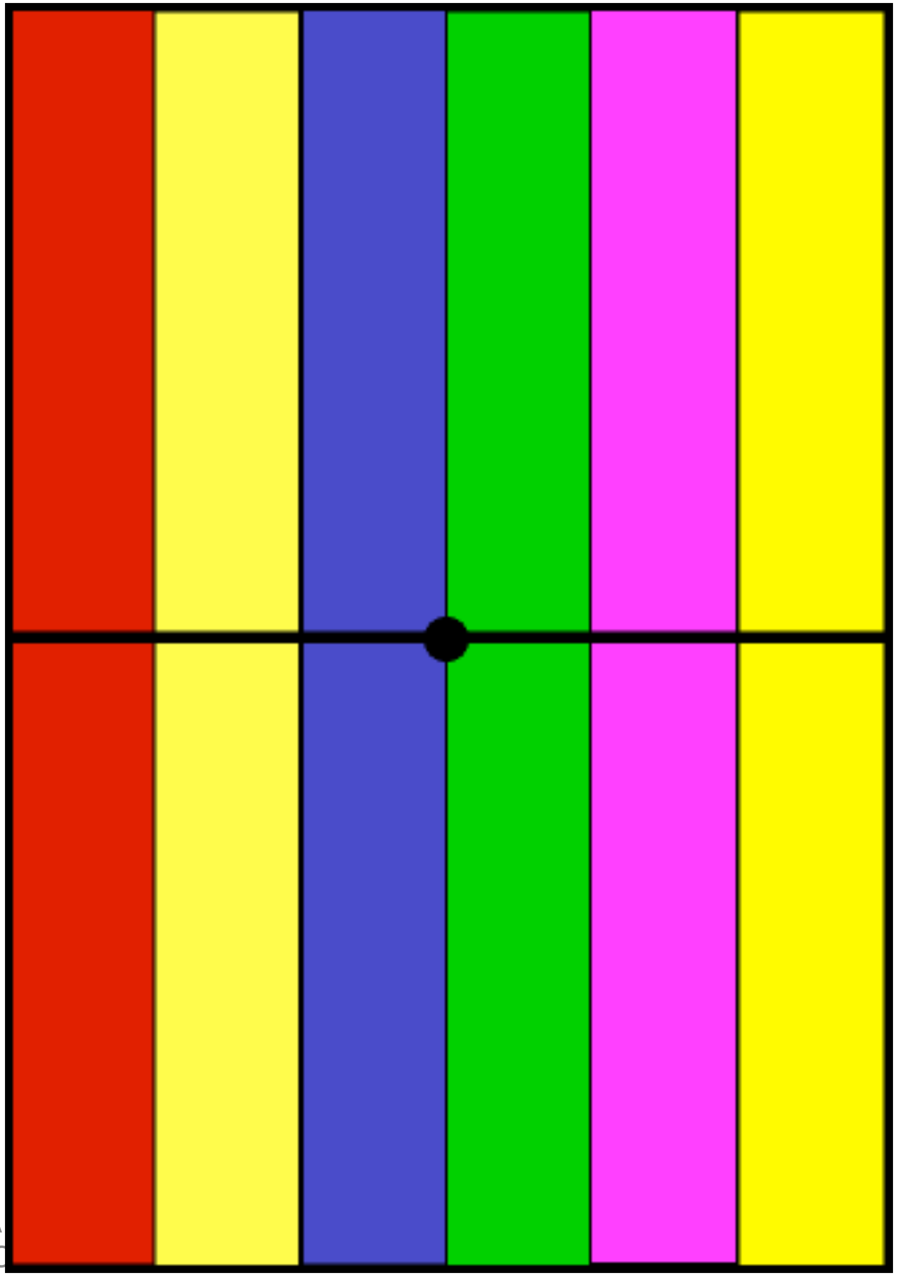

XXXXXX	GREEN	GREEN
XXXXXX	BLUE	BLUE
XXXXXX	YELLOW	YELLOW
XXXXXX	PURPLE	PURPLE
XXXXXX	ORANGE	ORANGE
XXXXXX	RED	RED
XXXXXX	WHITE	WHITE
XXXXXX	PURPLE	PURPLE
XXXXXX	ORANGE	ORANGE
XXXXXX	BLUE	BLUE
XXXXXX	RED	RED
XXXXXX	GREEN	GREEN
XXXXXX	WHITE	WHITE
XXXXXX	YELLOW	YELLOW
XXXXXX	PURPLE	PURPLE
XXXXXX	RED	RED
XXXXXX	GREEN	GREEN
XXXXXX	BLUE	BLUE

One Application: Photometric Stereo

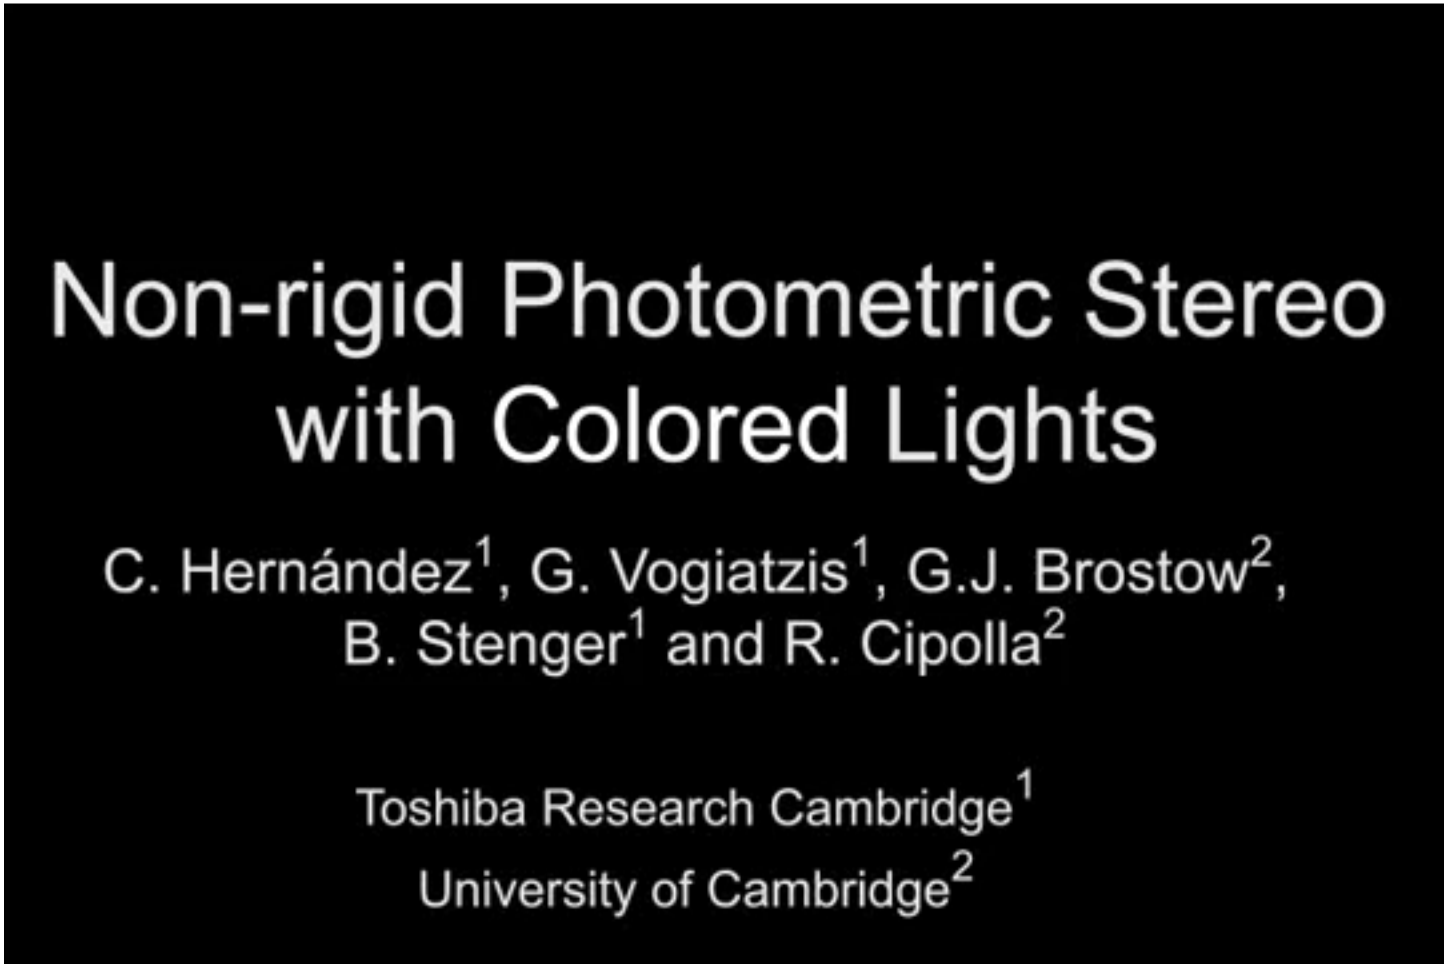

Next Lecture: Linear Filters and Image Processing

- Reading: FP 4; SZ 3