Outline

1. Final Projects
2. Radiometry
3. Color

Poster presentations

http://16720.courses.cs.cmu.edu/project.html

==== Availability form ====
One person from each team (at least) is required to fill out the following form about poster printing and availability by this Wednesday (4/20) at midnight.
(this is different from the google spreadsheet set a couple weeks ago):
https://docs.google.com/forms/d/19Io5iS80RScp1uYgWWmYnnLrDPL8nYfrHXKajmQYhtc/viewform

==== Posters ====
- Posters should fit on a 30in x 40in board. Landscape format is preferred.
- A Blackboard link is now up for submitting posters. All teams must submit a version of their poster (in pdf form) by 11:59pm on Sunday evening. Teams printing their own posters should update this with the final printed version if it changes after Sunday evening.
- It is preferred that each team print their own poster. If you have an SCS account, or have a friend with one, submit your poster using these instructions: https://www.cs.cmu.edu/~help/printing/poster_printing.html
You should submit your poster 48 hours ahead of the presentation to ensure it is printed on time. We recommend saving it to a pdf when you submit it (to avoid formatting changes).
- If you cannot submit your own poster, the TAs will get the poster printed for you (using the version submitted to Blackboard by Sunday evening). You will have to pick up the poster yourself once it has been printed (you will be notified). This option should be selected in the form you're required to fill out, linked above.
- Tips: check the website for poster formatting tips to avoid common mistakes:

Poster presentations

http://16720.courses.cs.cmu.edu/project.html

==== Presentations ====

Each team will present on only one of the two days (schedule to be released). Students should arrive between 2:30 and 2:45pm to set up their posters on the day of their presentation. Presentations will be in the NSH Atrium.

==== Writeups ====

A Blackboard link is now up for the writeups. Check the website for writeup suggestions (requirements: 8 page maximum, CVPR format):

Big picture (1)

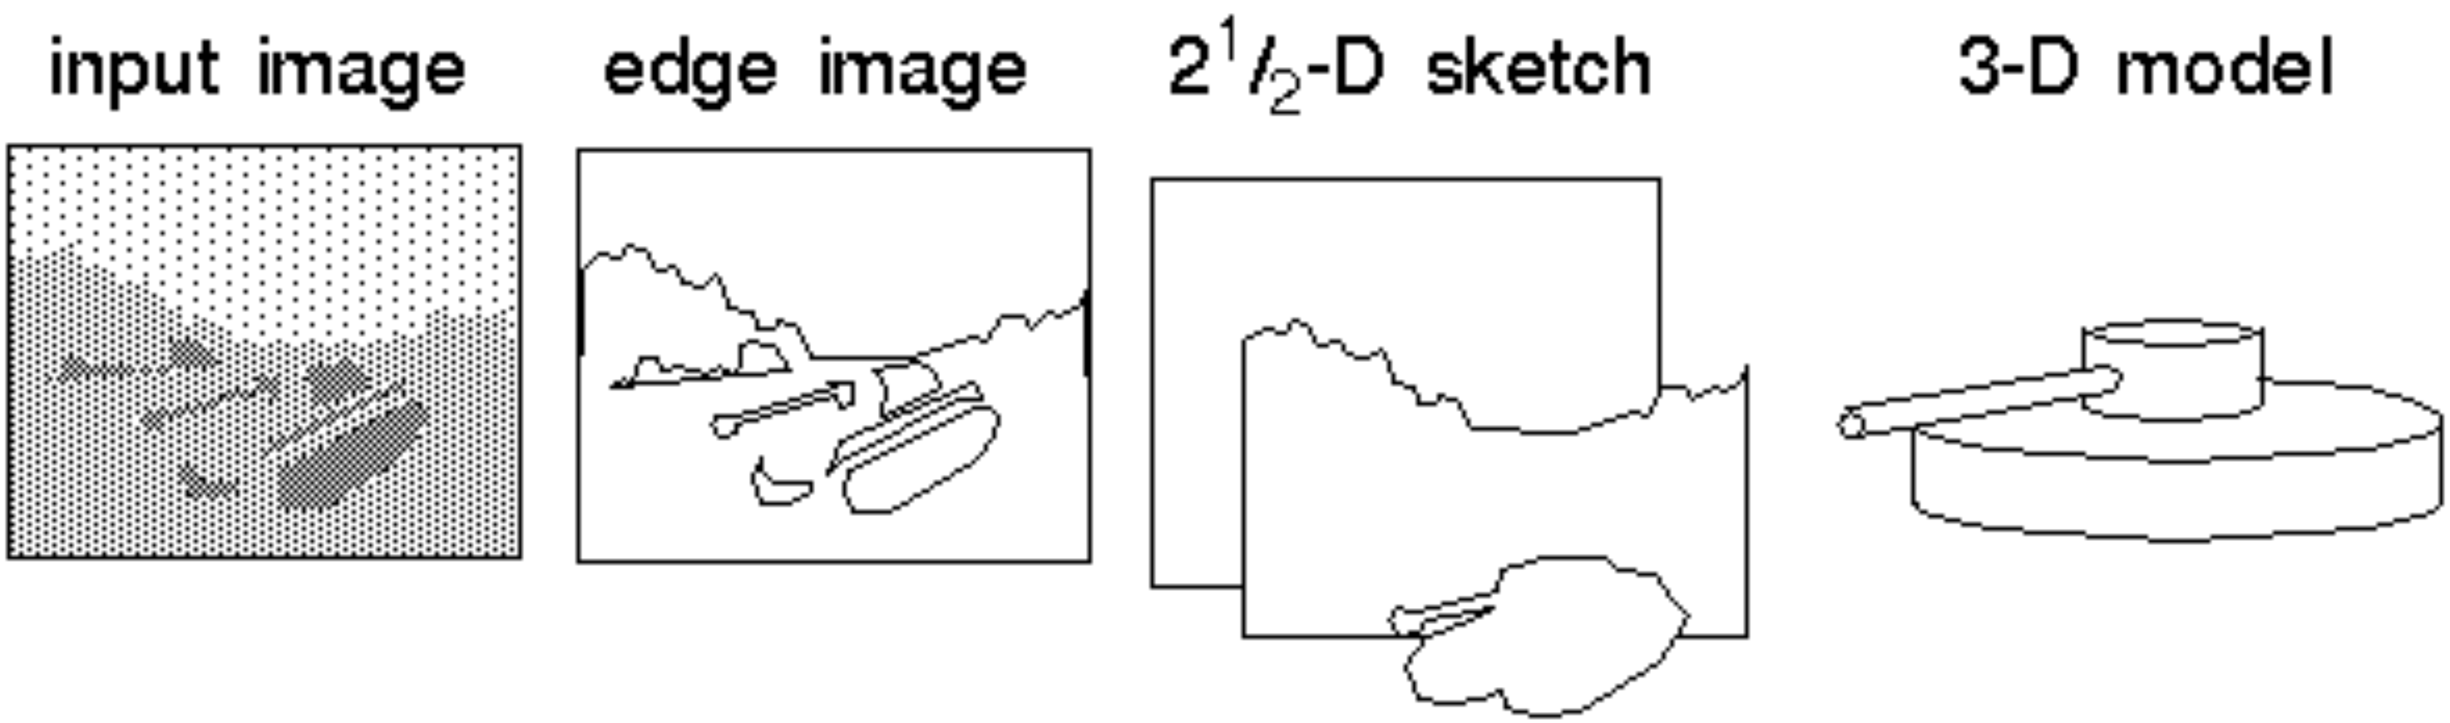

Low,Mid,High-Level Vision

Big picture (2)

Geometry

Physics of appearance

Radiometry

Study of light transport

Classically a computer graphics topic (rendering)
What is its relevance to vision?

Motivation: understanding shadows and shading

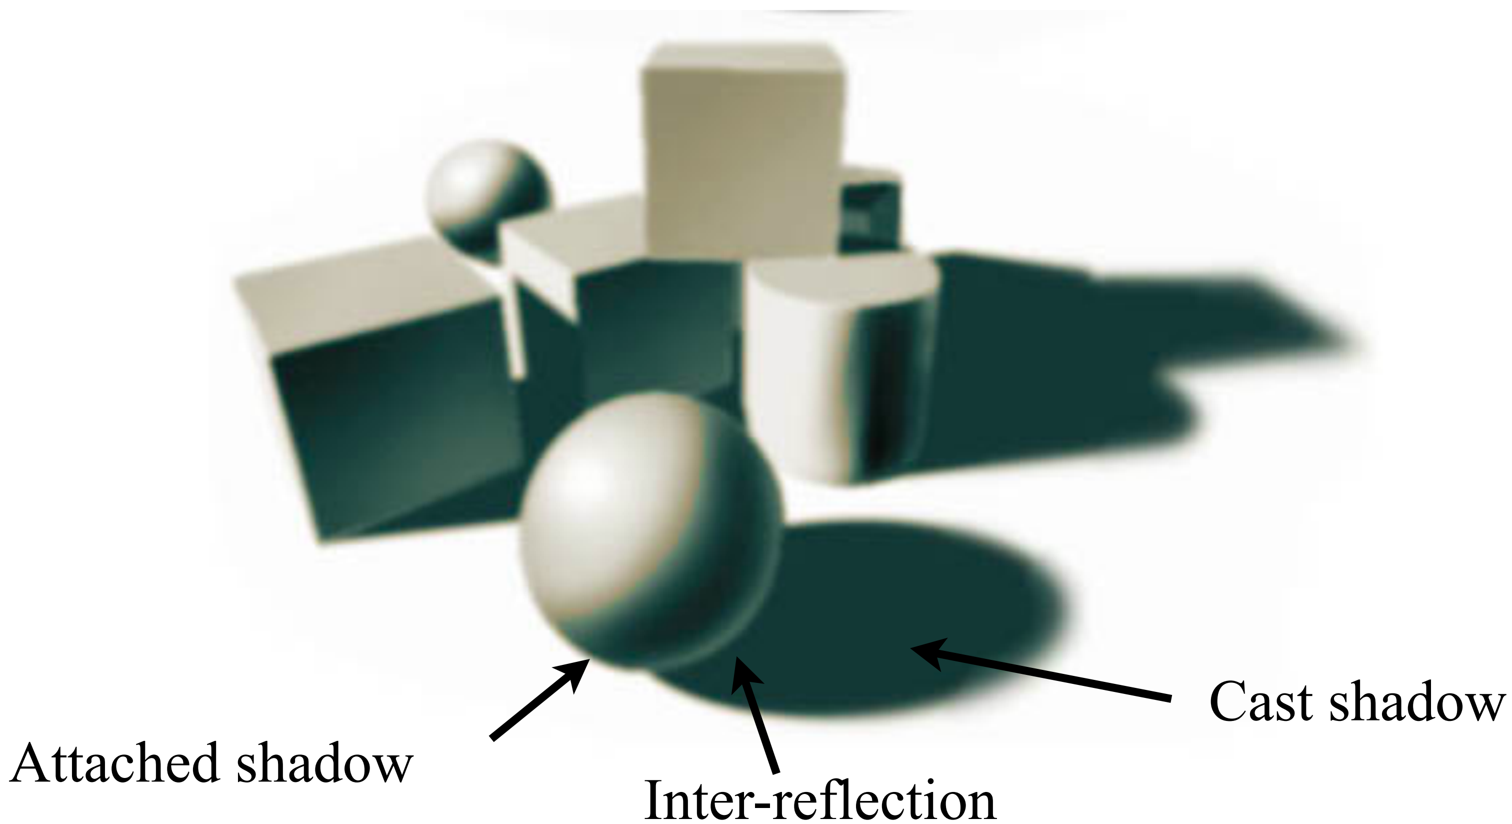

Complicating factors

1. Inter-reflections

   Other surfaces can reflect light (e.g., the wall of a room), effectively behaving as “light sources” themselves

2. Cast shadows

   Cast vs attached shadows

Motivation: dehazing

Dehazing

Motivation: photometric stereo

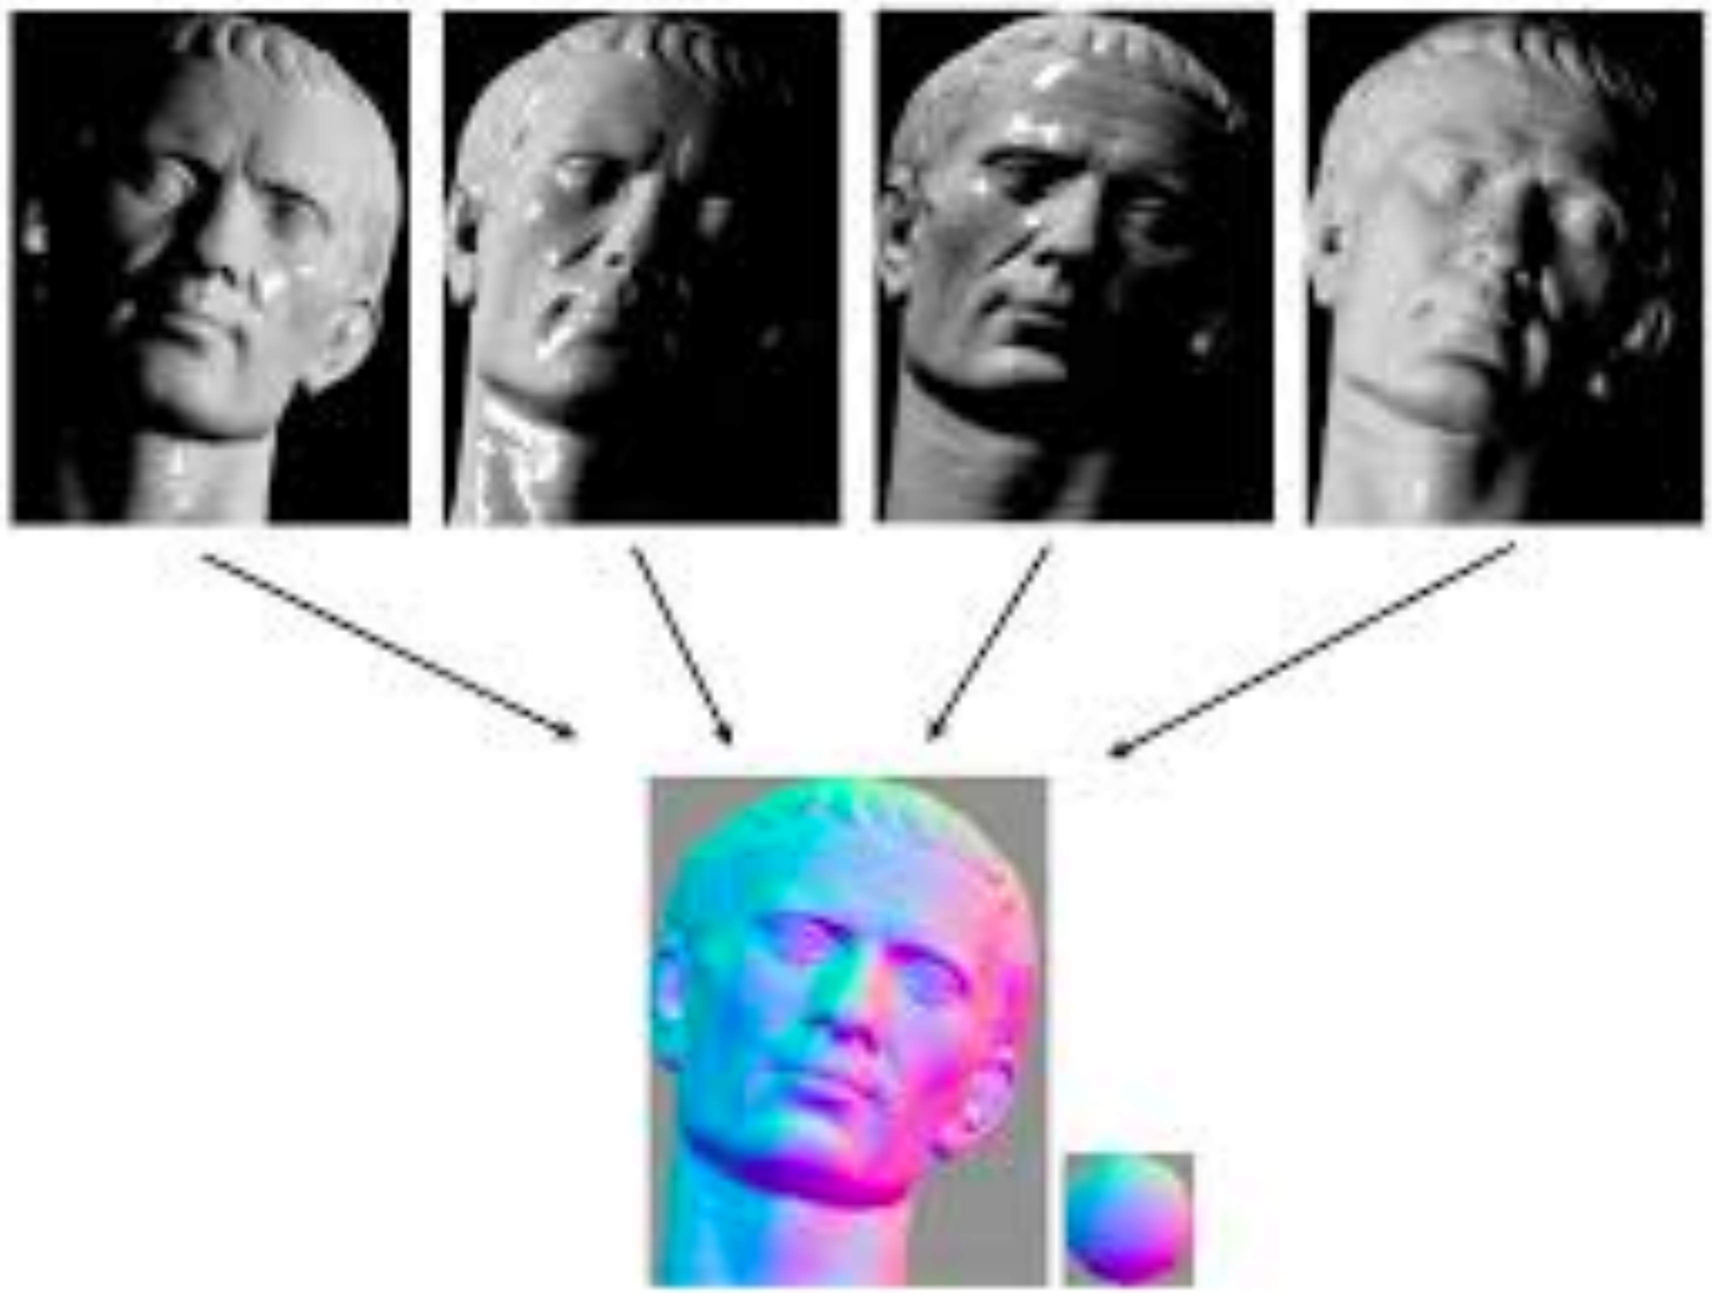

Reconstruct from a single viewpoint, but with multiple images with differing light sources

Motivation: color processing

Why can color be represented with 3 (R,G,B) values?

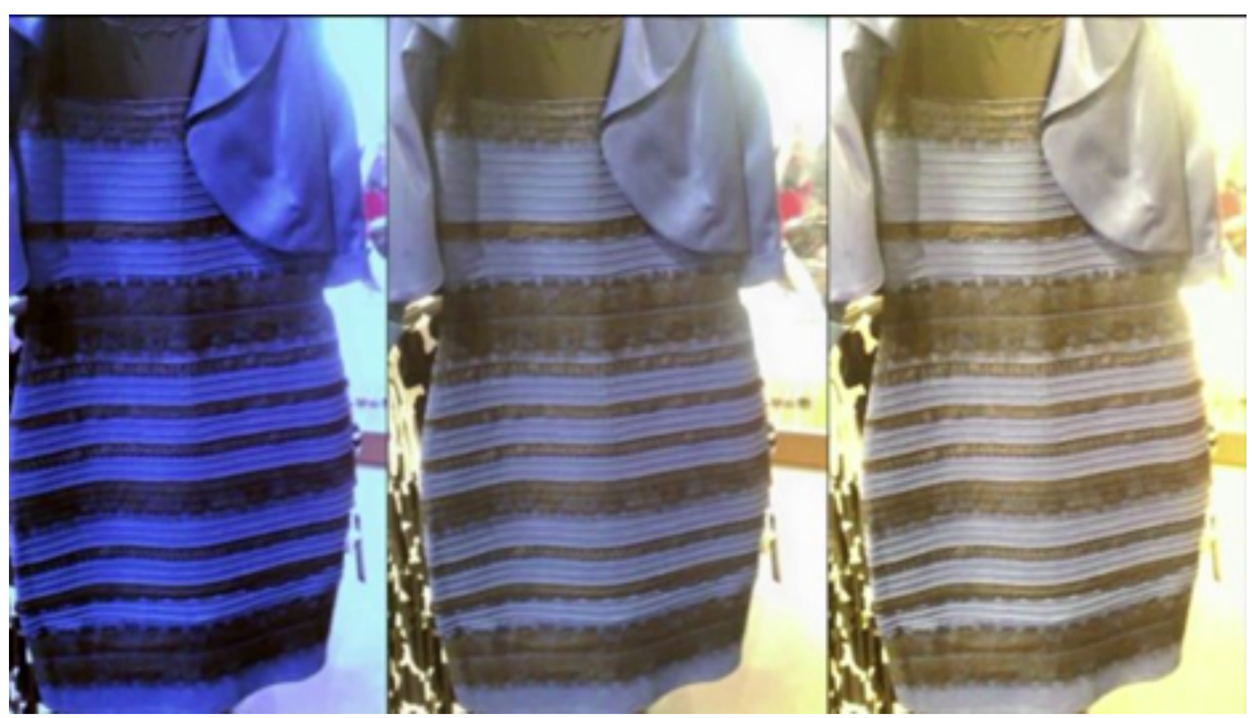

Why do objects look different under different illuminations?

Physics of illumination

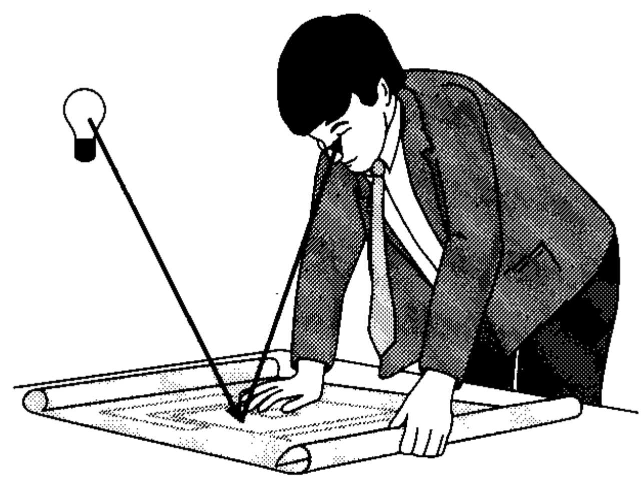

Light source
Surface
Camera

Physics of illumination

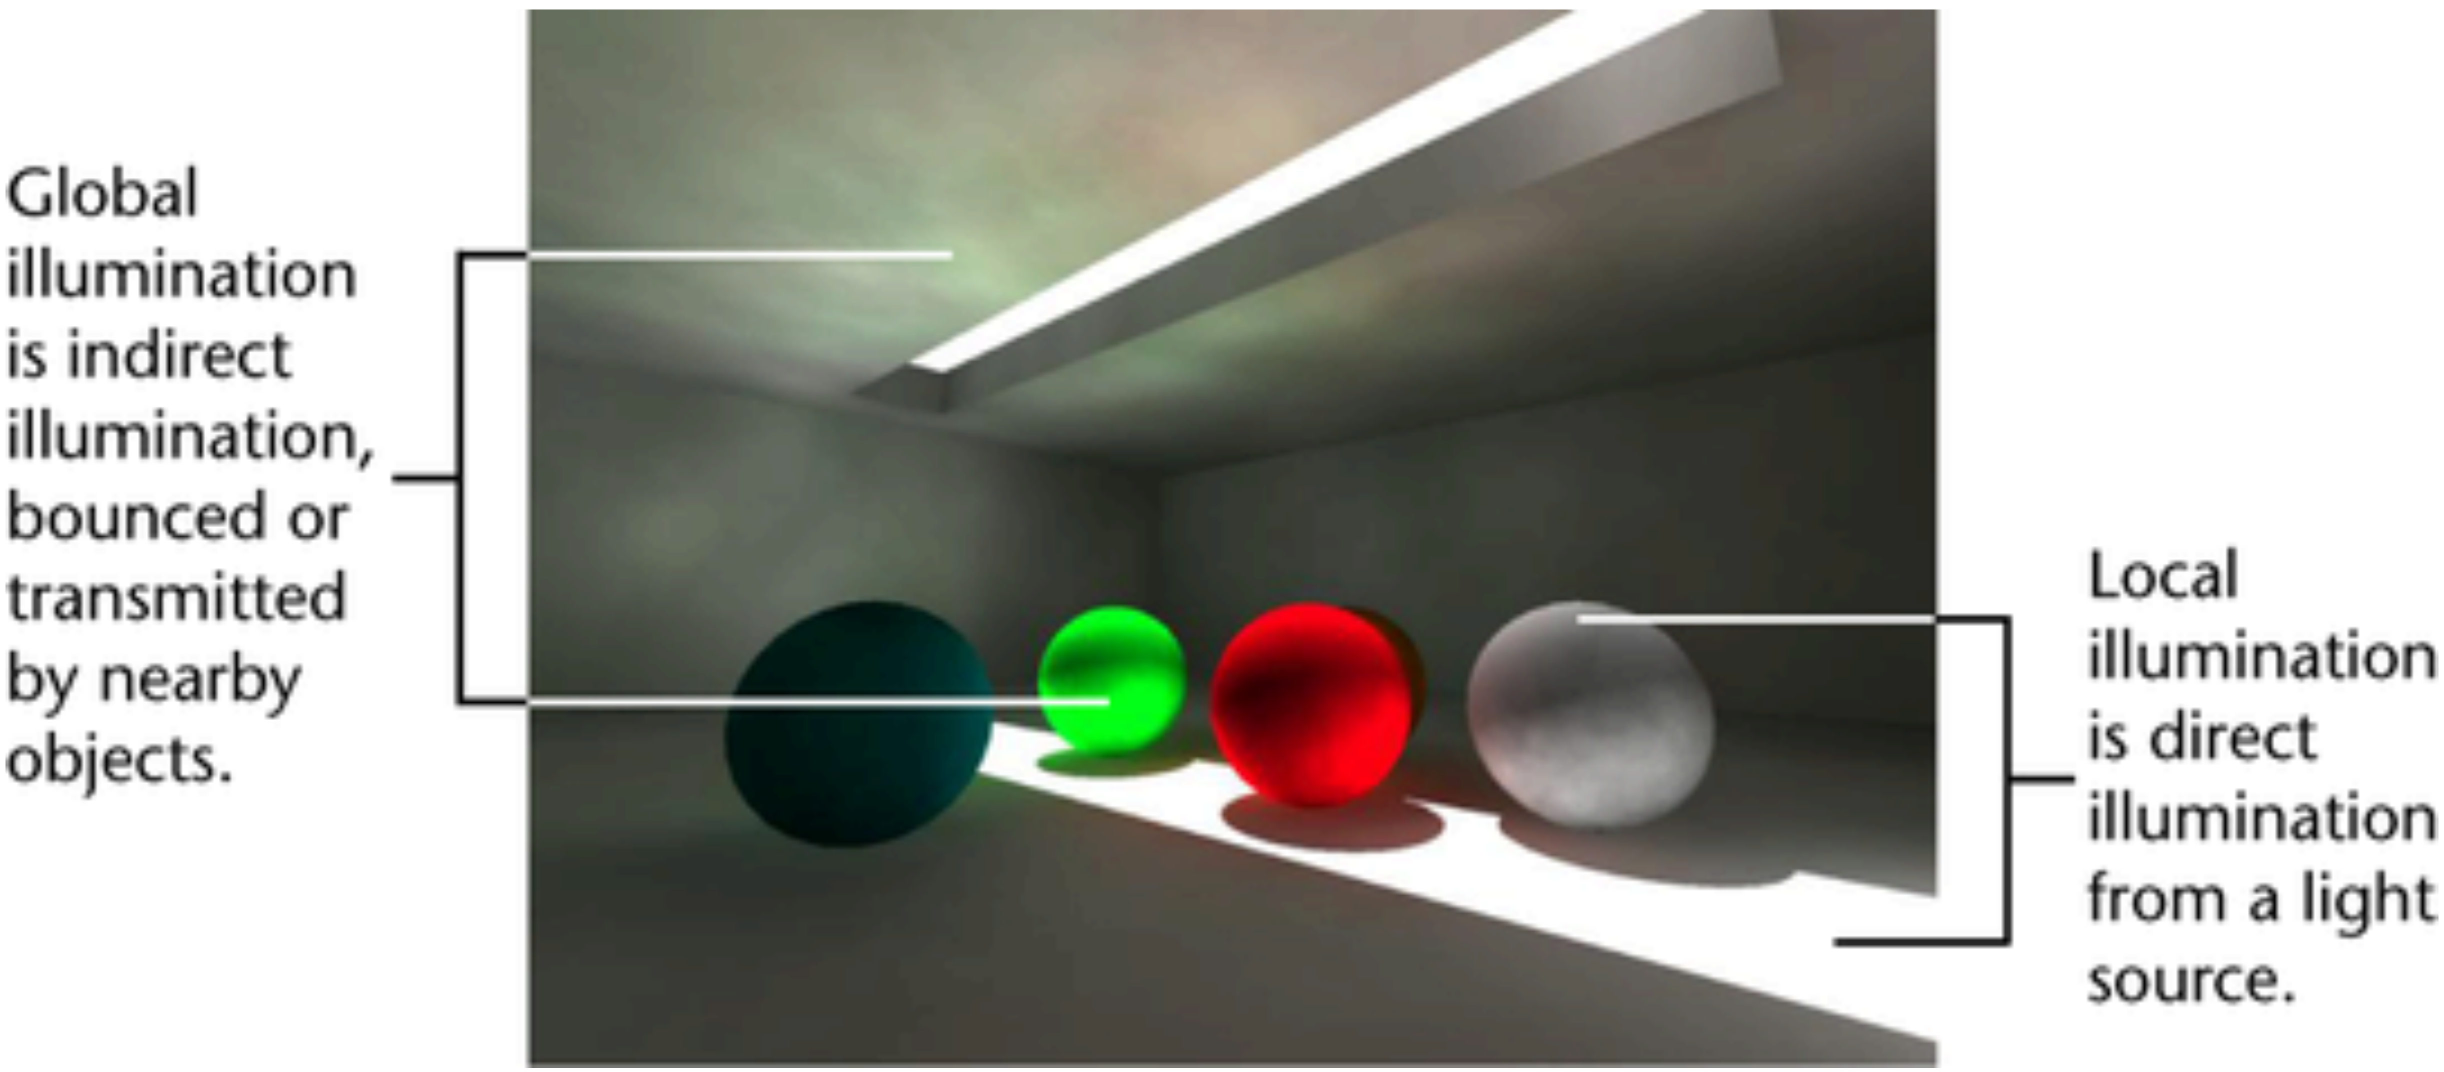

Recall:
light as a wave + particle

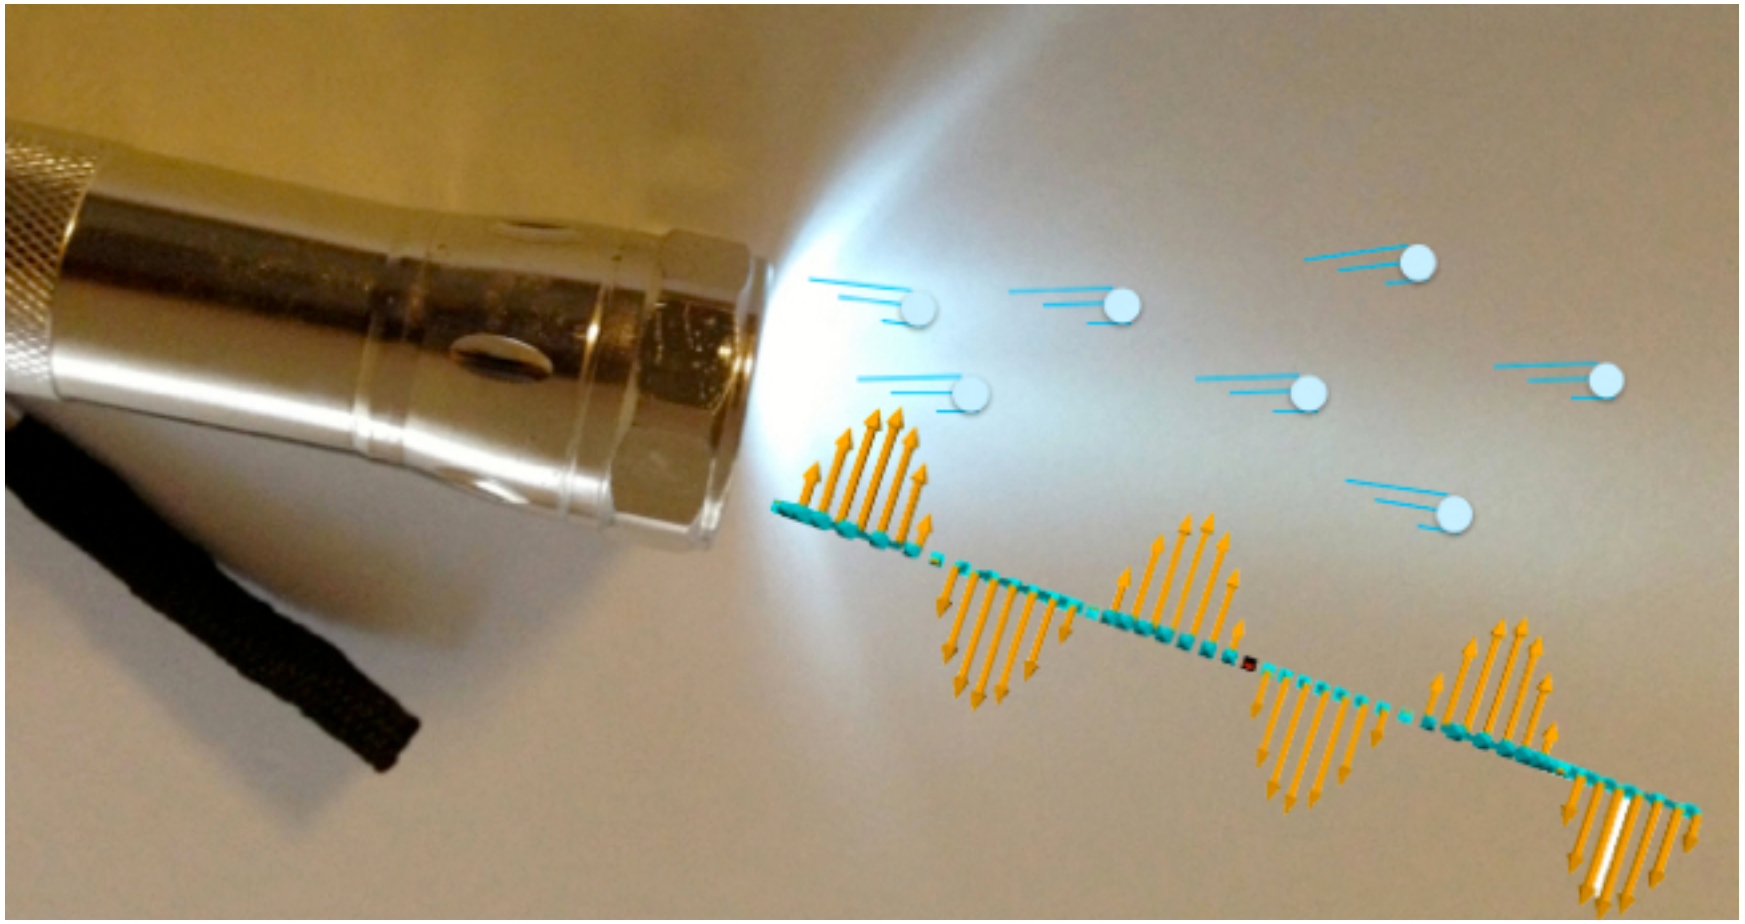

Fundamental quantity: irradiance

Measures amount of light hitting a surface

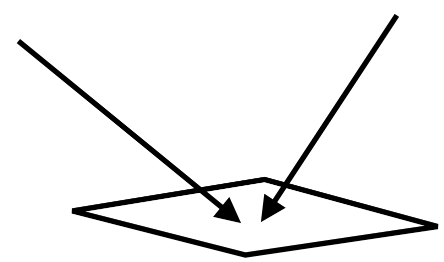

Irradiance = power of electromagnetic radiation / area

= watts / meter^2

Fundamental quantity: irradiance

Measures amount of light hitting a surface

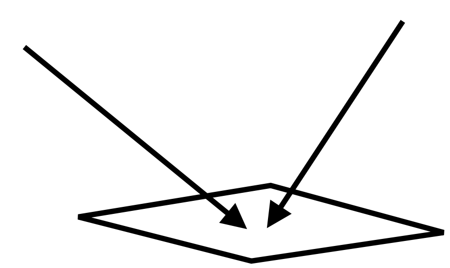

Irradiance = power of electromagnetic radiation / area

= watts / meter^2

Power = energy / sec

Fundamental quantity: irradiance

Measures amount of light hitting a surface

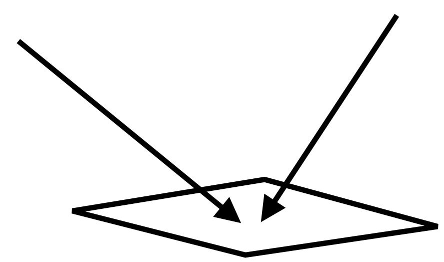

Irradiance = power of electromagnetic radiation / area

= watts / meter^2

= joules / (sec meter^2)

My intuition: # of photons hitting (infinitesimly) small area in (infinitesimly) small time step

Fundamental quantity (2): radiance

Measures amount of light traveling along rays in space

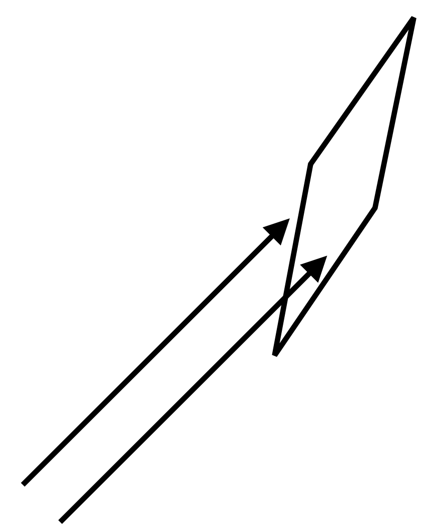

Fundamental quantity (2): radiance

Measures amount of light traveling along particular direction in space

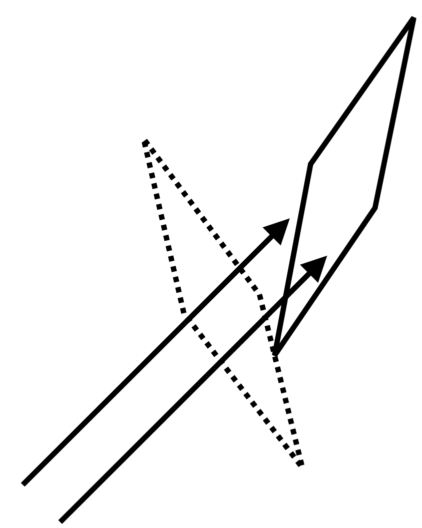

Subtety 1: Should count photons travelling along direction per foreshortened area patch

Subtety 2: Measuring infinitesimly small set of directions in 3D

Angles and Solid Angles

- **Angle** $\theta = \frac{l}{r}$

  $\Rightarrow$ **circle has 2π radians**

- **Solid angle** $\Omega = \frac{A}{R^2}$

  $\Rightarrow$ **sphere has 4π steradians**

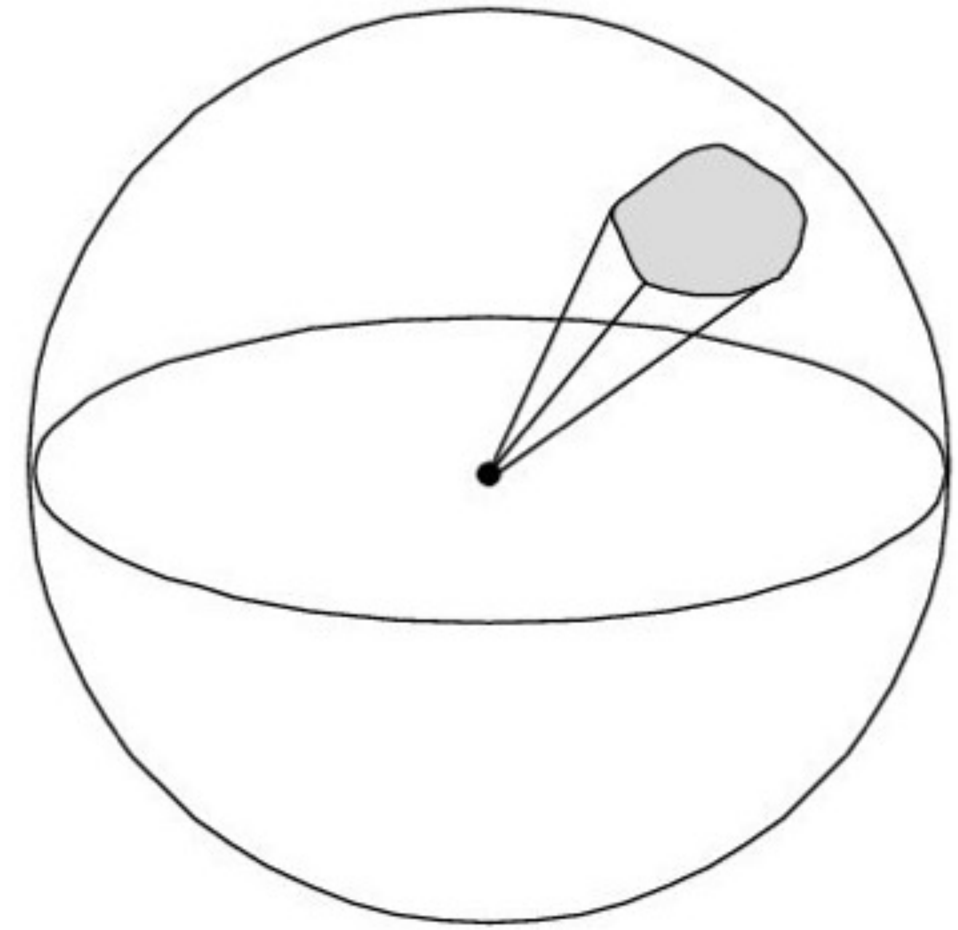

Differential solid angles

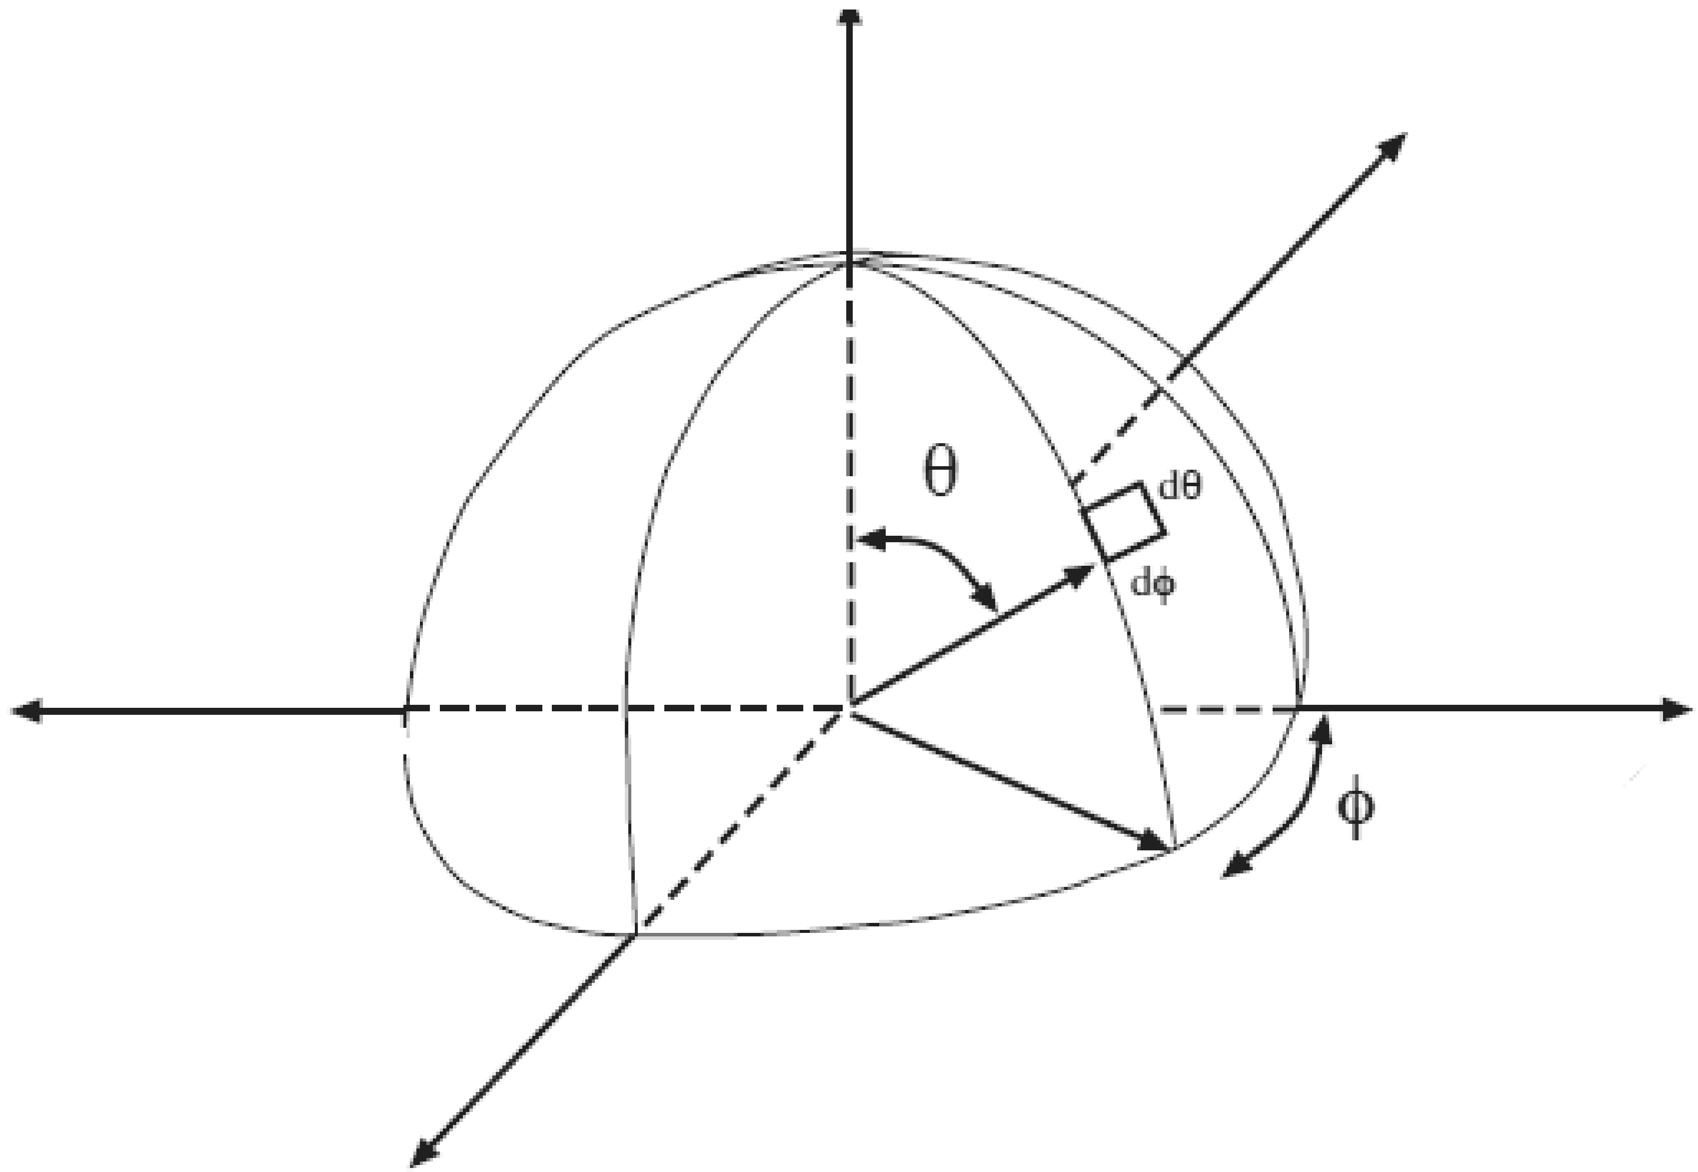

The messy math

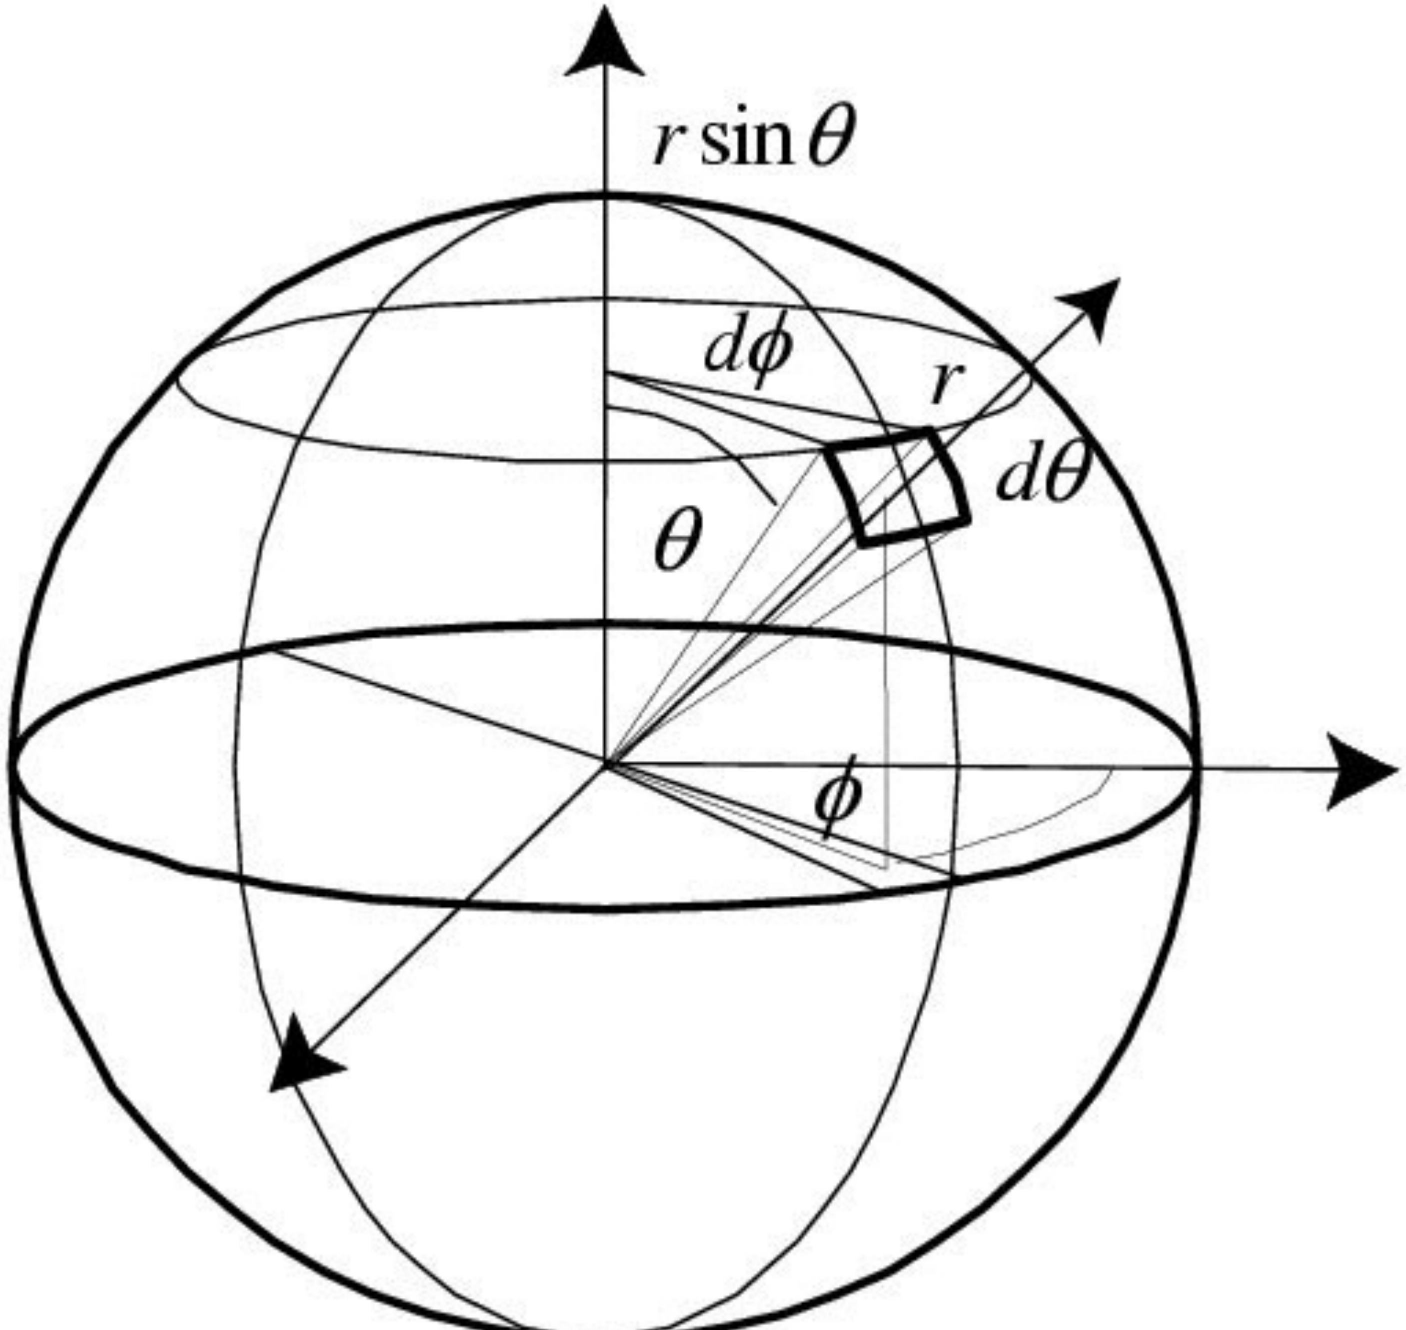

$$dA = (r\,d\theta)(r\sin\theta\,d\phi)$$
$$= r^2 \sin\theta\,d\theta\,d\phi$$

$$d\omega = \frac{dA}{r^2} = \sin\theta\,d\theta\,d\phi$$

$$S = \int_0^{\pi}\int_0^{2\pi} \sin\theta\,d\theta\,d\phi = 4\pi$$

Fundamental quantity (2): radiance

Measures amount of light traveling along particular direction in space

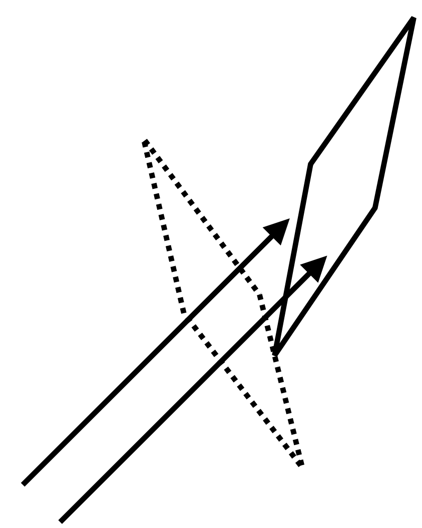

Radiance = power / (foreshortened area * solid angle)

= watt / (meter^2 * steradians)

BRDF: Bidirectional Reflectance Distribution Function

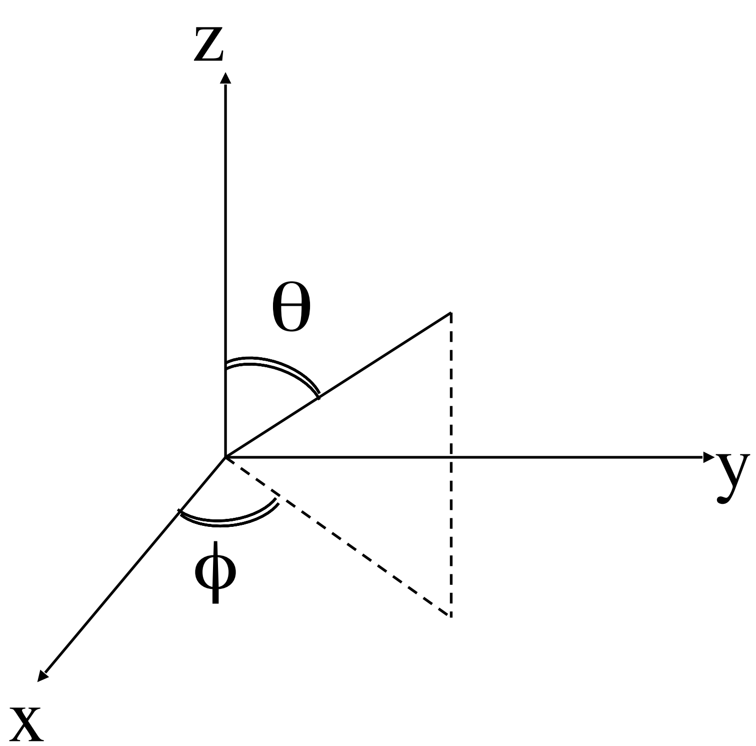

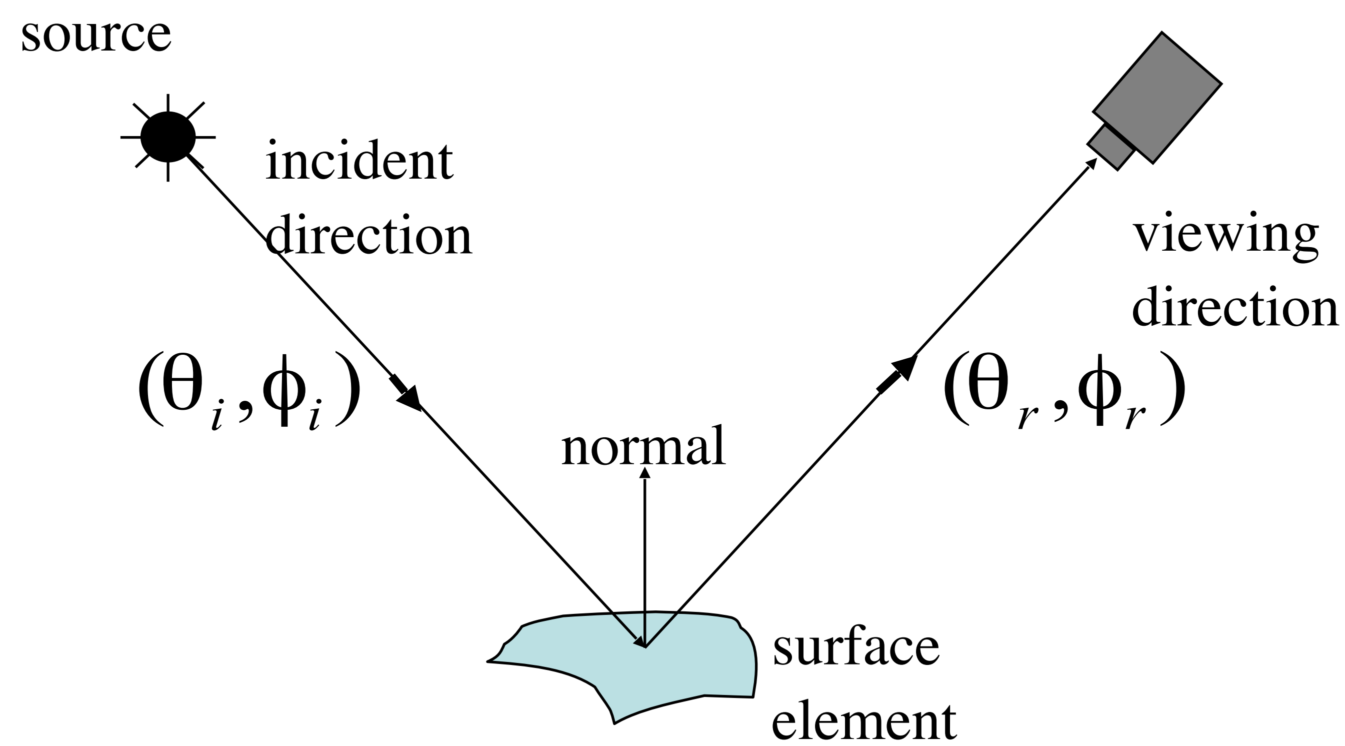

$$f(\theta_i, \phi_i; \theta_r, \phi_r) = \frac{L_r(\theta_r, \phi_r)}{L_i(\theta_i, \phi_i) \cos \theta_i}$$

BRDF: Bidirectional Reflectance Distribution Function

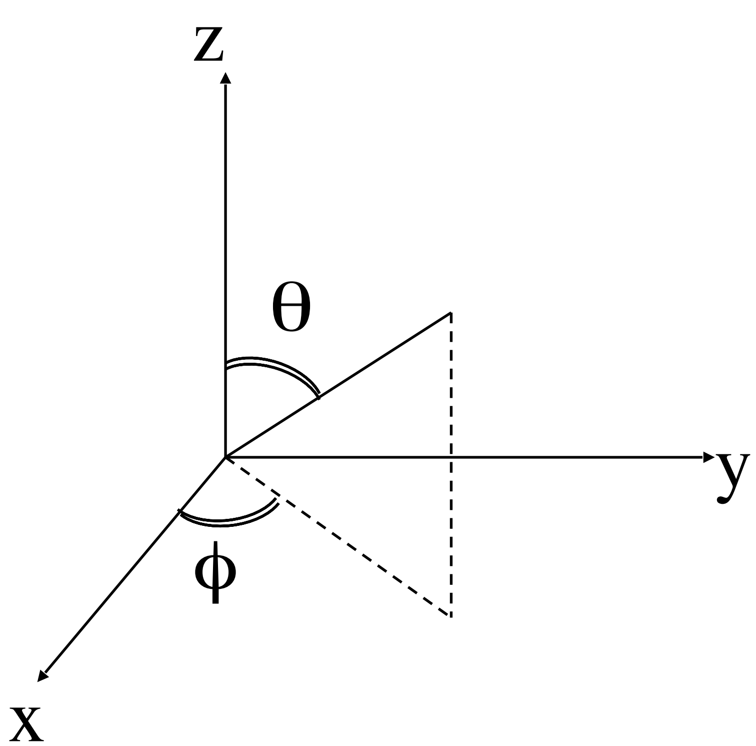

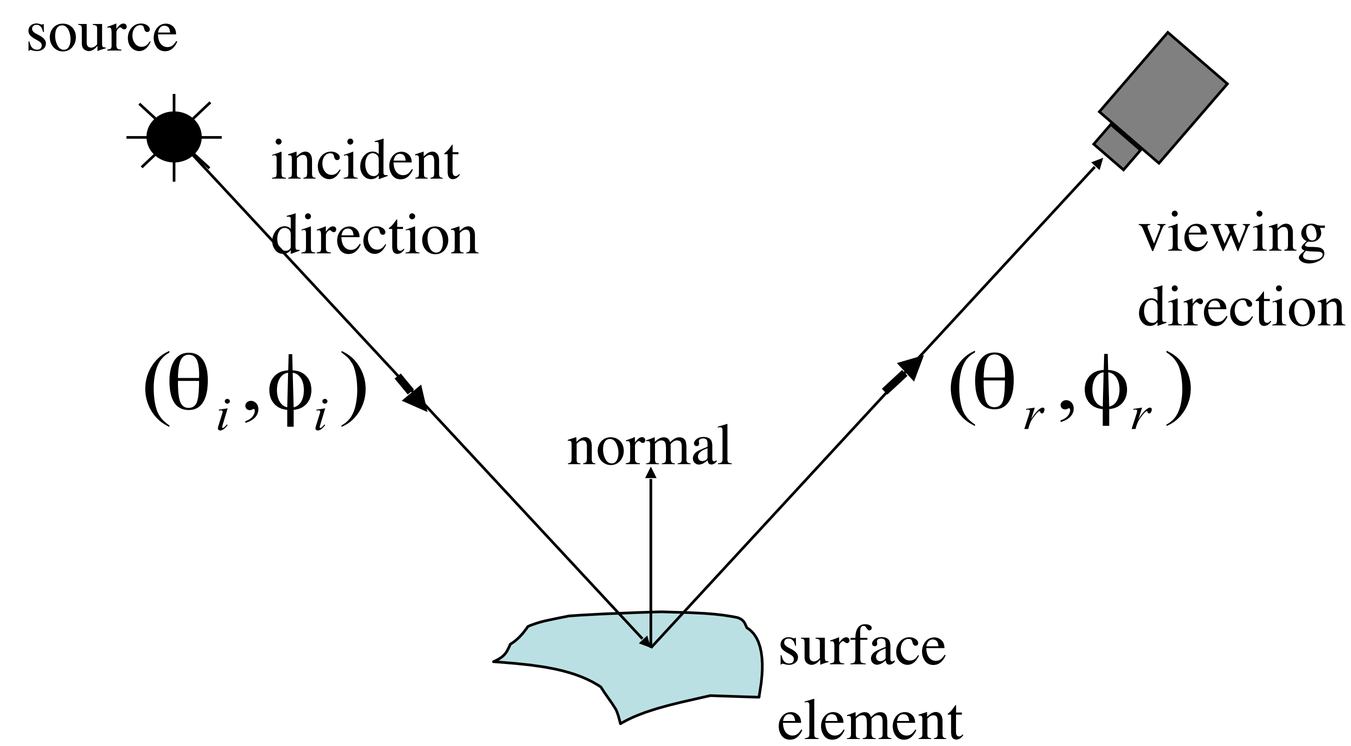

$$f(\theta_i, \phi_i; \theta_r, \phi_r) = \frac{L_r(\theta_r, \phi_r)}{L_i(\theta_i, \phi_i) \cos \theta_i}$$

$$f(\omega_i, \omega_r) = \frac{L_r(\omega_r)}{L_i(\omega_i) \cos \theta_i}$$ (solid angle notation)

BRDF: Bidirectional Reflectance Distribution Function

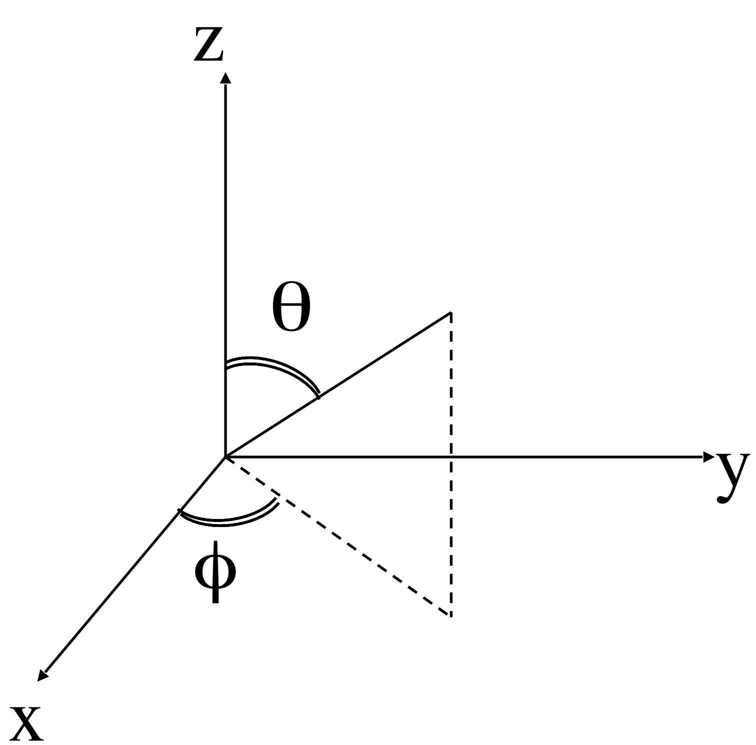

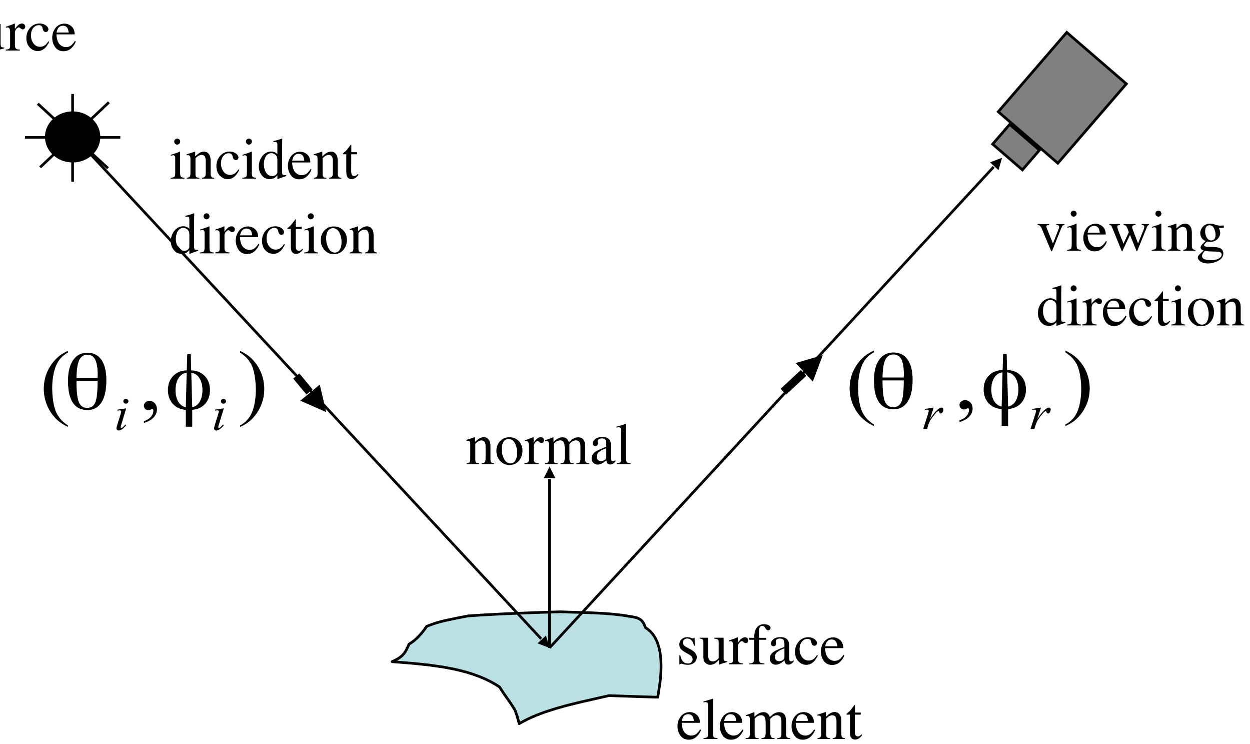

$$f(\theta_i, \phi_i, \theta_r, \phi_r) = \frac{L_r(\theta_r, \phi_r)}{L_i(\theta_i, \phi_i)\cos\theta_i}$$

$$f(\omega_i, \omega_r) = \frac{L_r(\omega_r)}{L_i(\omega_i)\cos\theta_i}$$

$E_i(\omega_i)$ (directional irradiance falling on patch)

The “rendering equation”

Add ‘x’ to model a spatially-varying BRDF

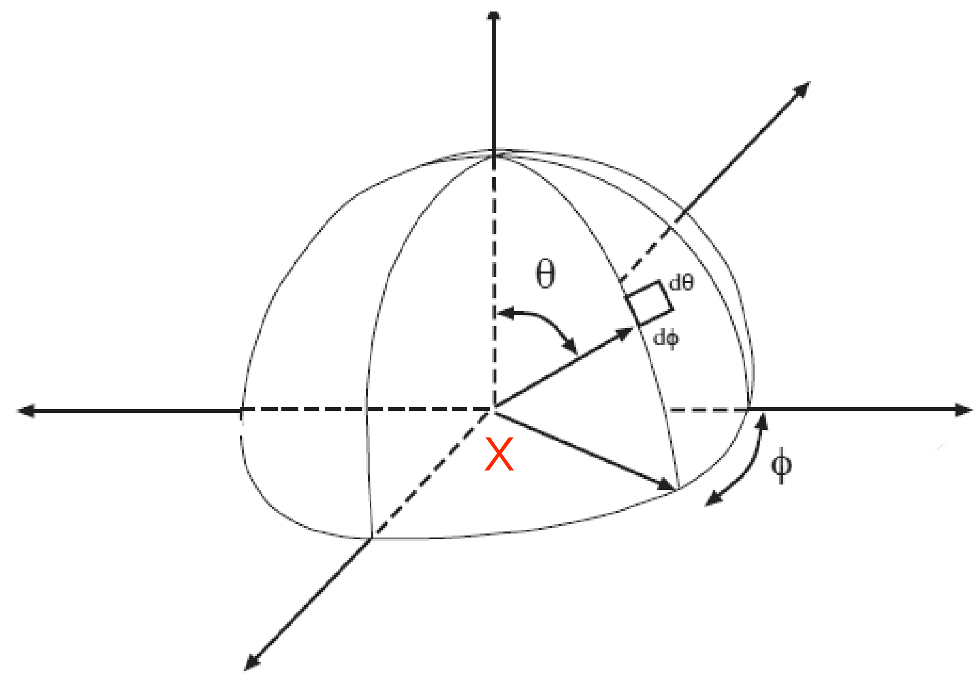

$$
\begin{aligned}
L_r(x,\omega_r) &= L_e(x,\omega_r) + \int_\Omega L_i(x,\omega_i) f(x,\omega_i,\omega_r)\cos\theta_i d\omega_i \\
&= L_e(x,\omega_r) + \int_\Omega L_r(x',-\omega_i) f(x,\omega_i,\omega_r)\cos\theta_i d\omega_i
\end{aligned}
$$

Replace incident irradiance with reflected radiance from some other surface in the scene

Global illumination: given lights (Le), visibility/geometry (x,x’), and materials (BRDFs) in a scene, solve for reflected radiance

Measuring BRDFs

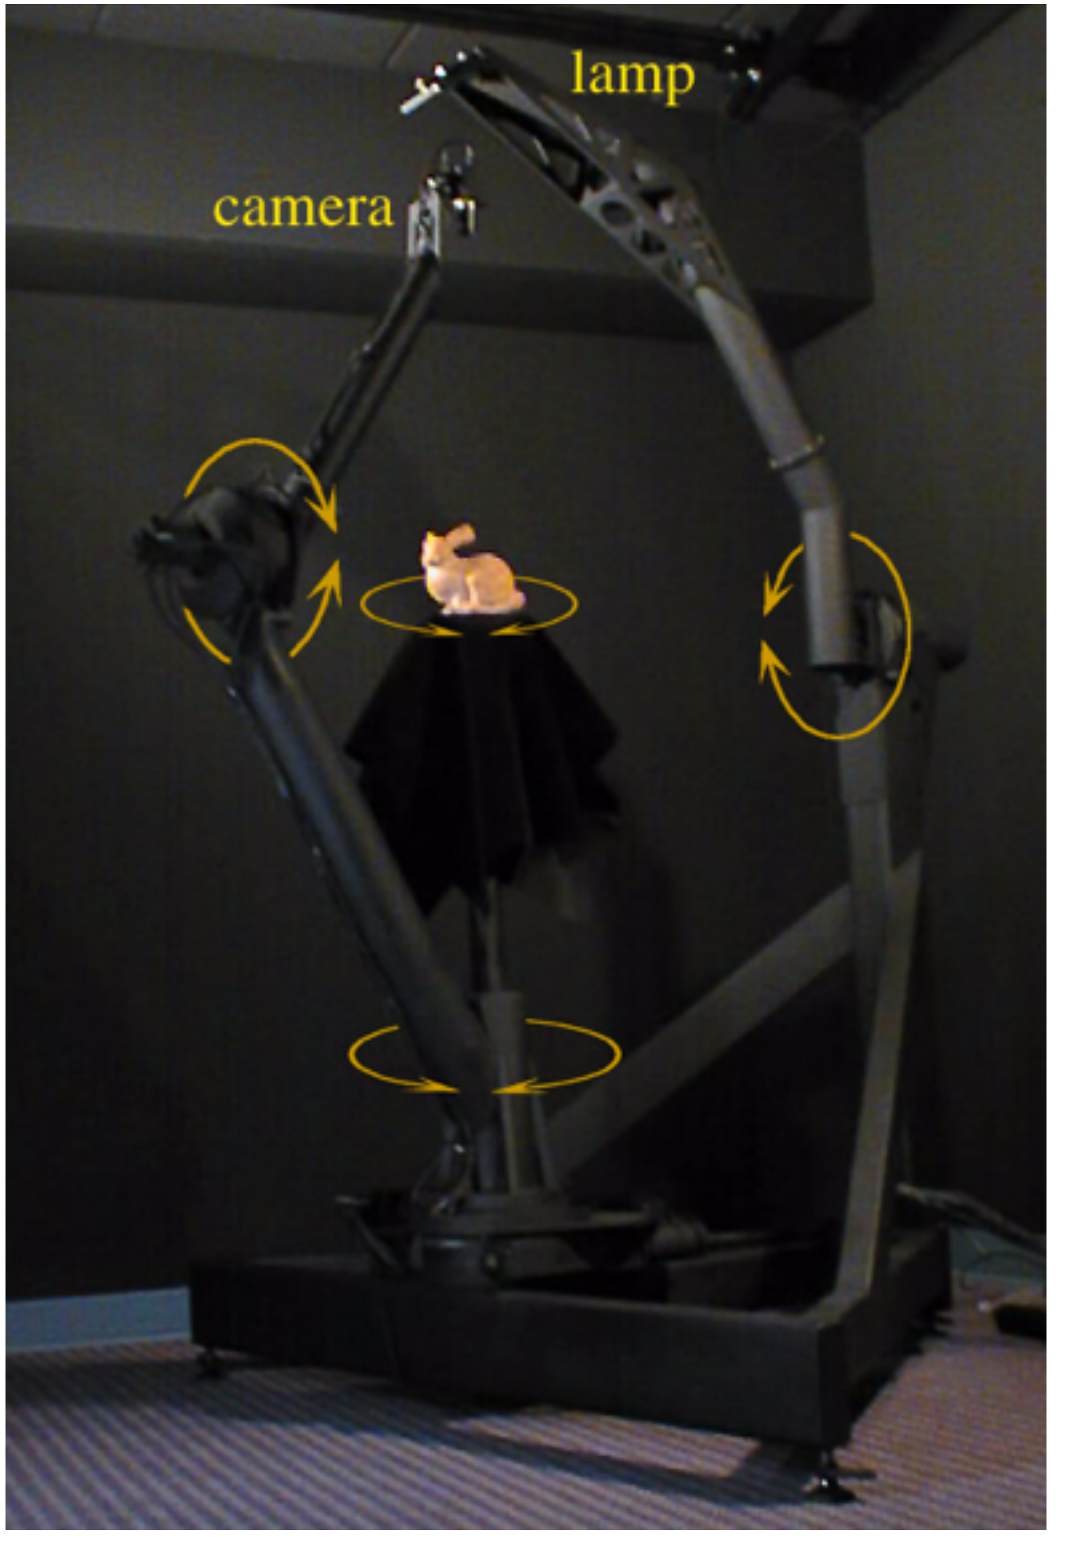

Stanford Spherical Gantry

Visualizing BRDFs

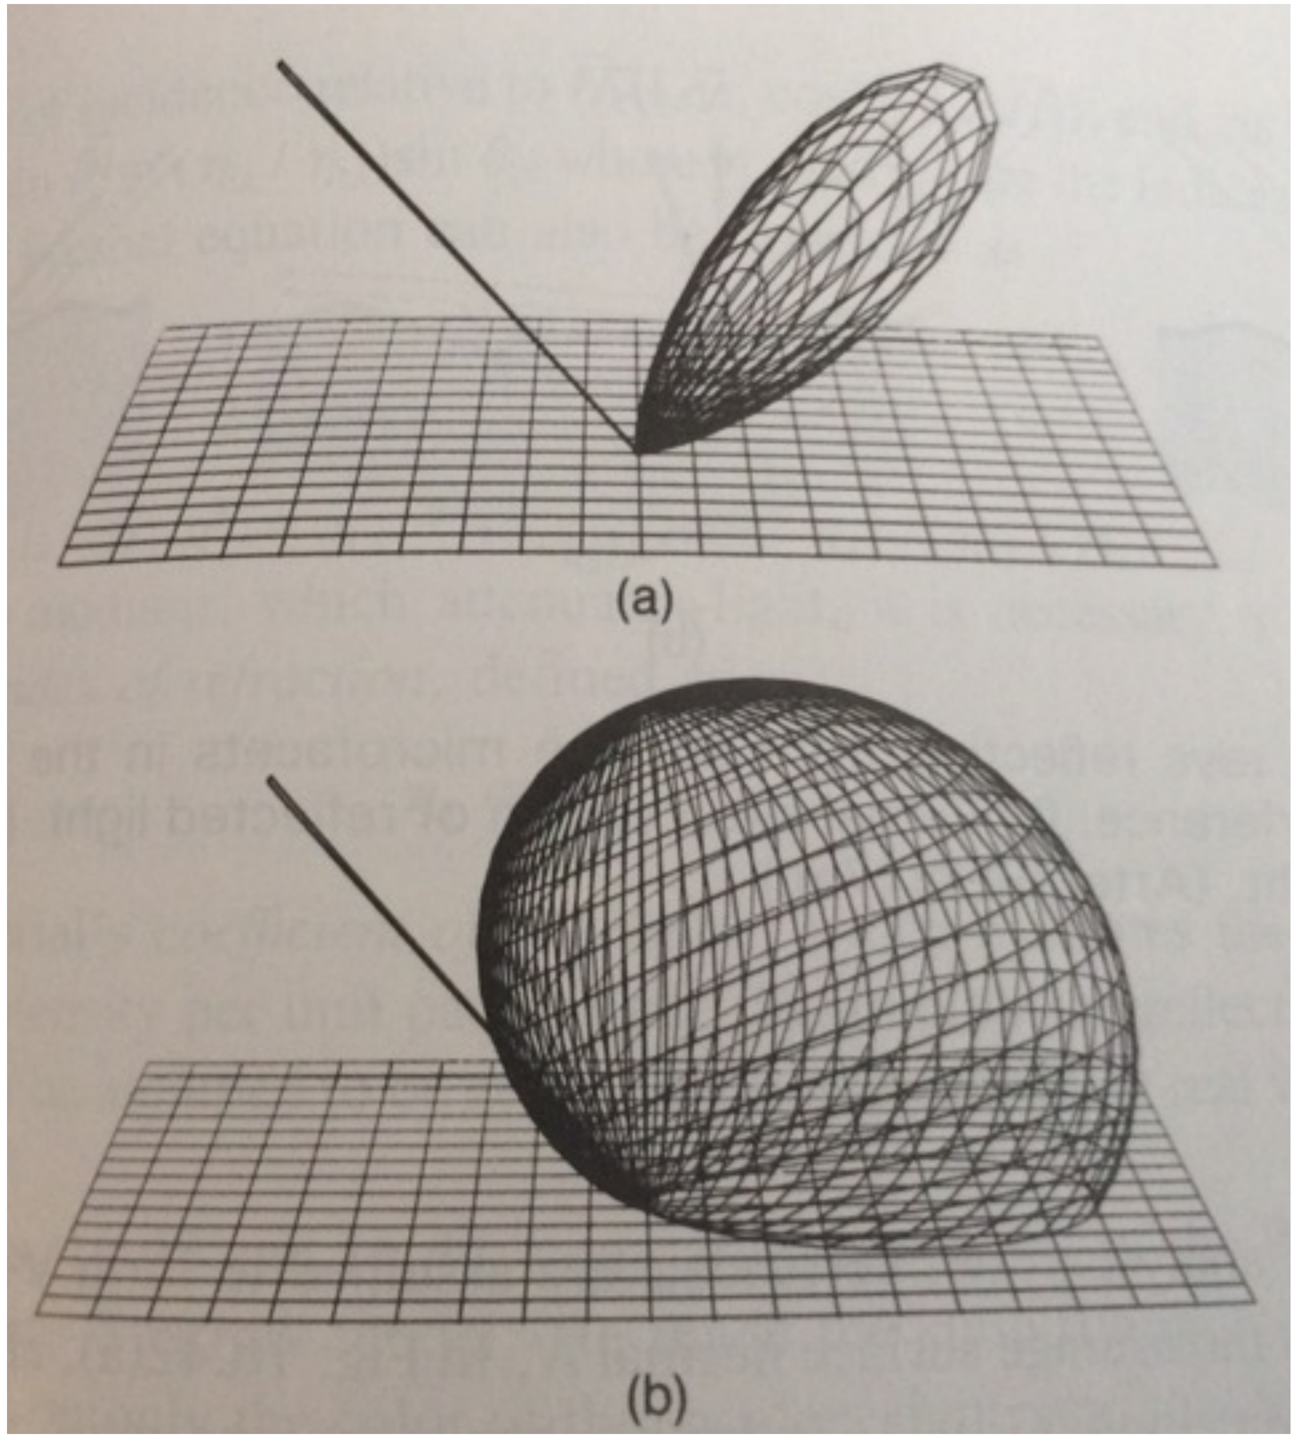

Foley & Van Dam

In order to make computations tractable, we'll make lots of simplifying assumptions: constant (albedo), mirror-like, etc…

Special cases

Isotropic

$$f(\theta_i, \phi_i, \theta_r, \phi_r) = f(\theta_i, \theta_i, \phi_r - \phi_i)$$

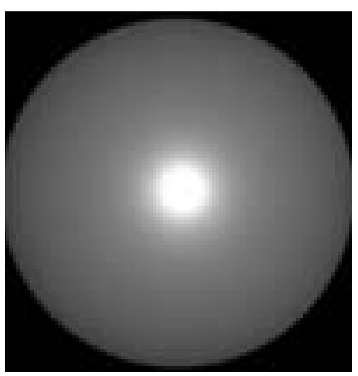

Rotate about normal without changing reflections

(most materials)

Anisotropic

$$f(\theta_i, \phi_i, \theta_r, \phi_r)$$

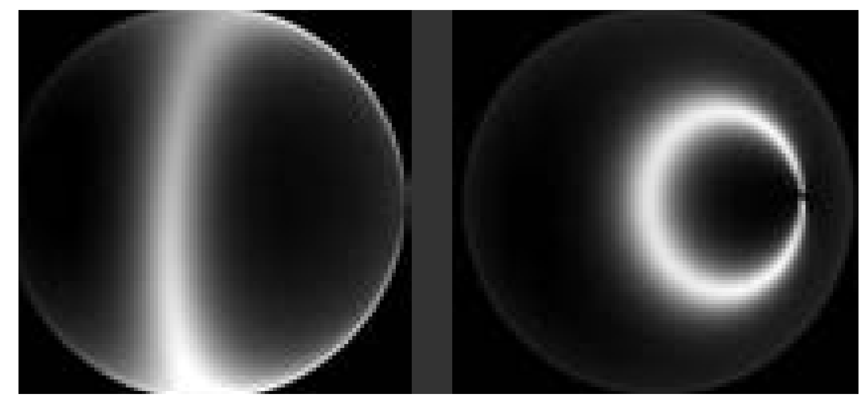

Special case (2): constant (Lambertian/diffuse) BRDF

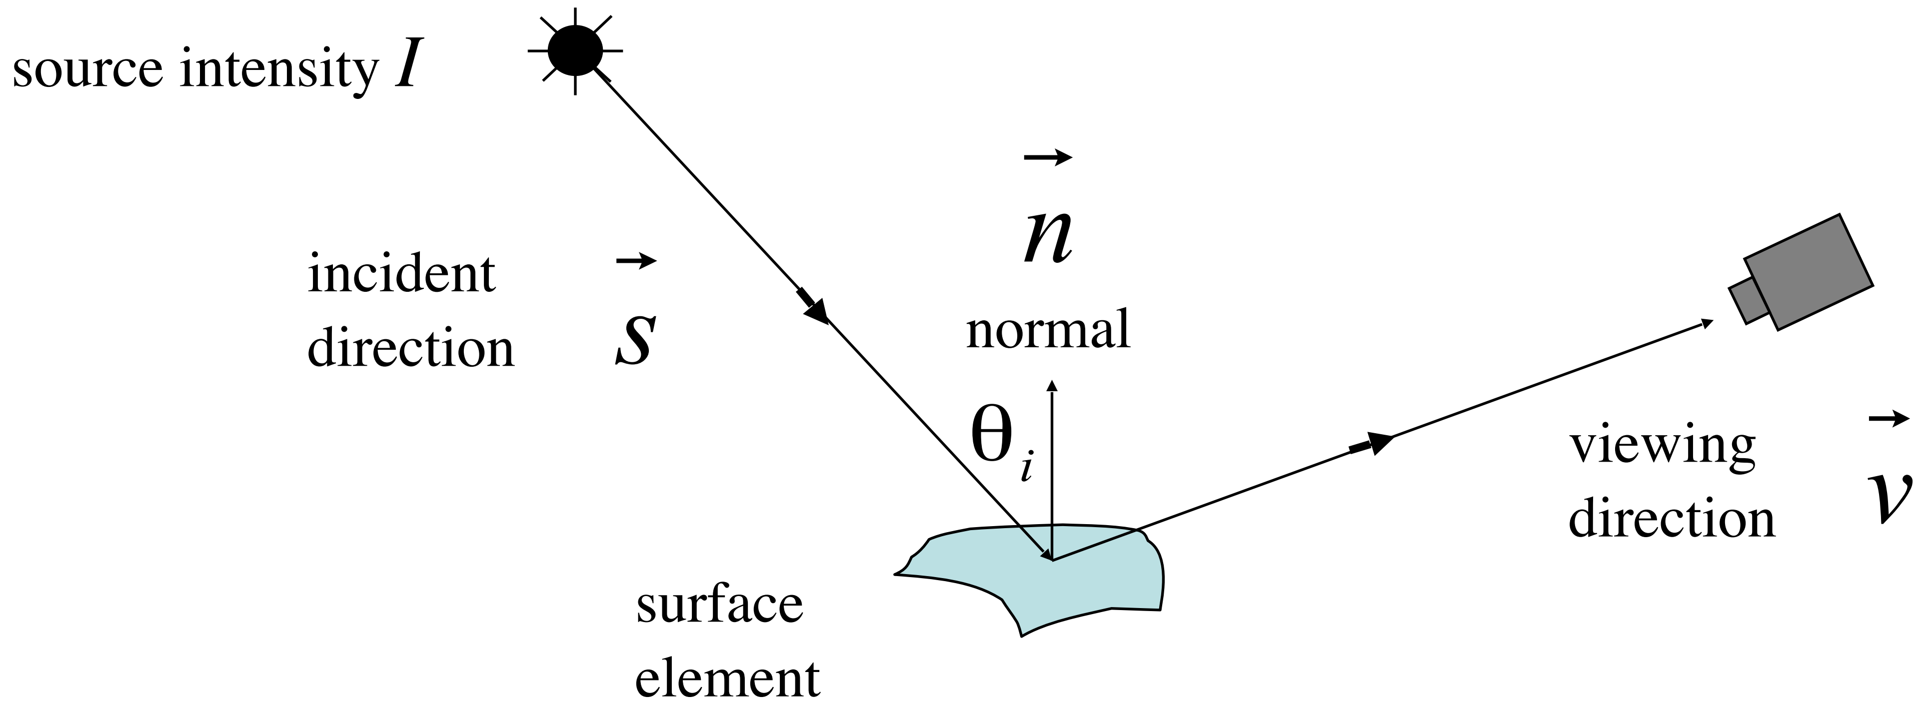

- Surface appears equally bright from ALL directions! (independent of $\vec{v}$)

- Lambertian BRDF is simply a constant : $f(\omega_i, \omega_r) = \rho_d$ (albedo)

- Surface Radiance : $L \propto I \cos\theta_i = \vec{n} \cdot \vec{s}$

- Commonly used in Vision and Graphics!

Application: photometric stereo

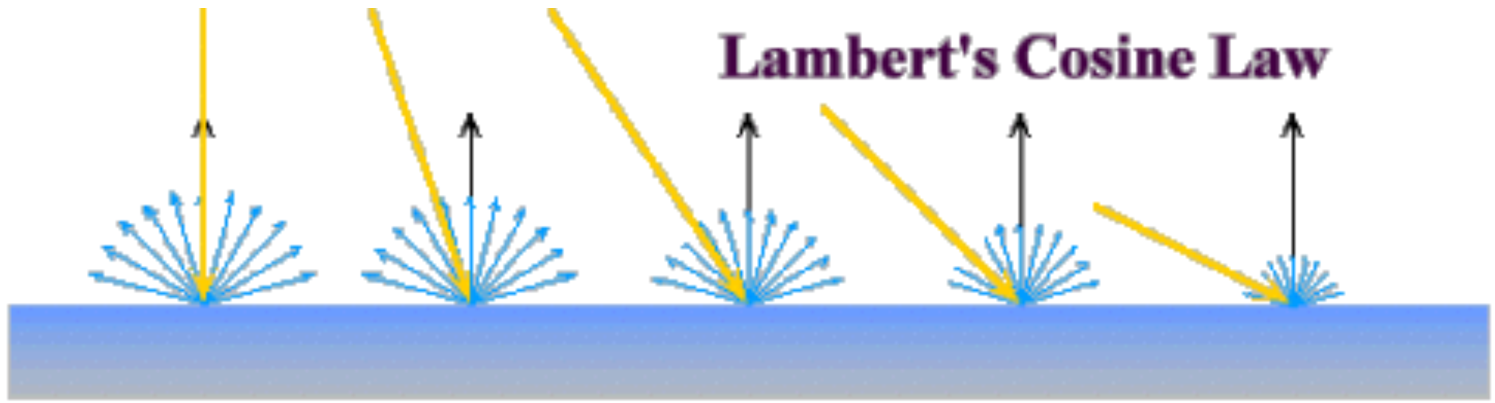

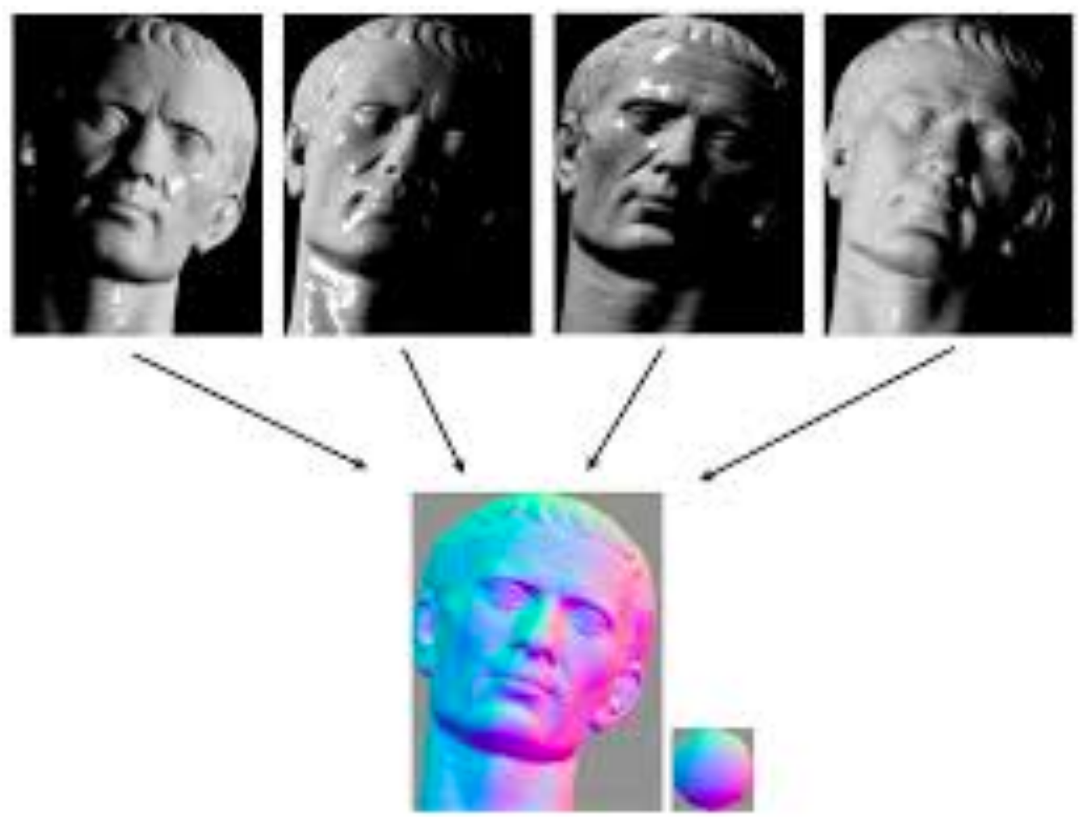

$$\text{Image Intensity} \approx \hat{n} \cdot \hat{s}$$

Given multiple measurements of images with known lighting directions (s), solve for normals with least squares

White-out Conditions from an Overcast Sky

CAN'T perceive the shape of the snow covered terrain!

CAN perceive shape in regions
lit by the street lamp!!

WHY?

Special case (2): ideal mirror

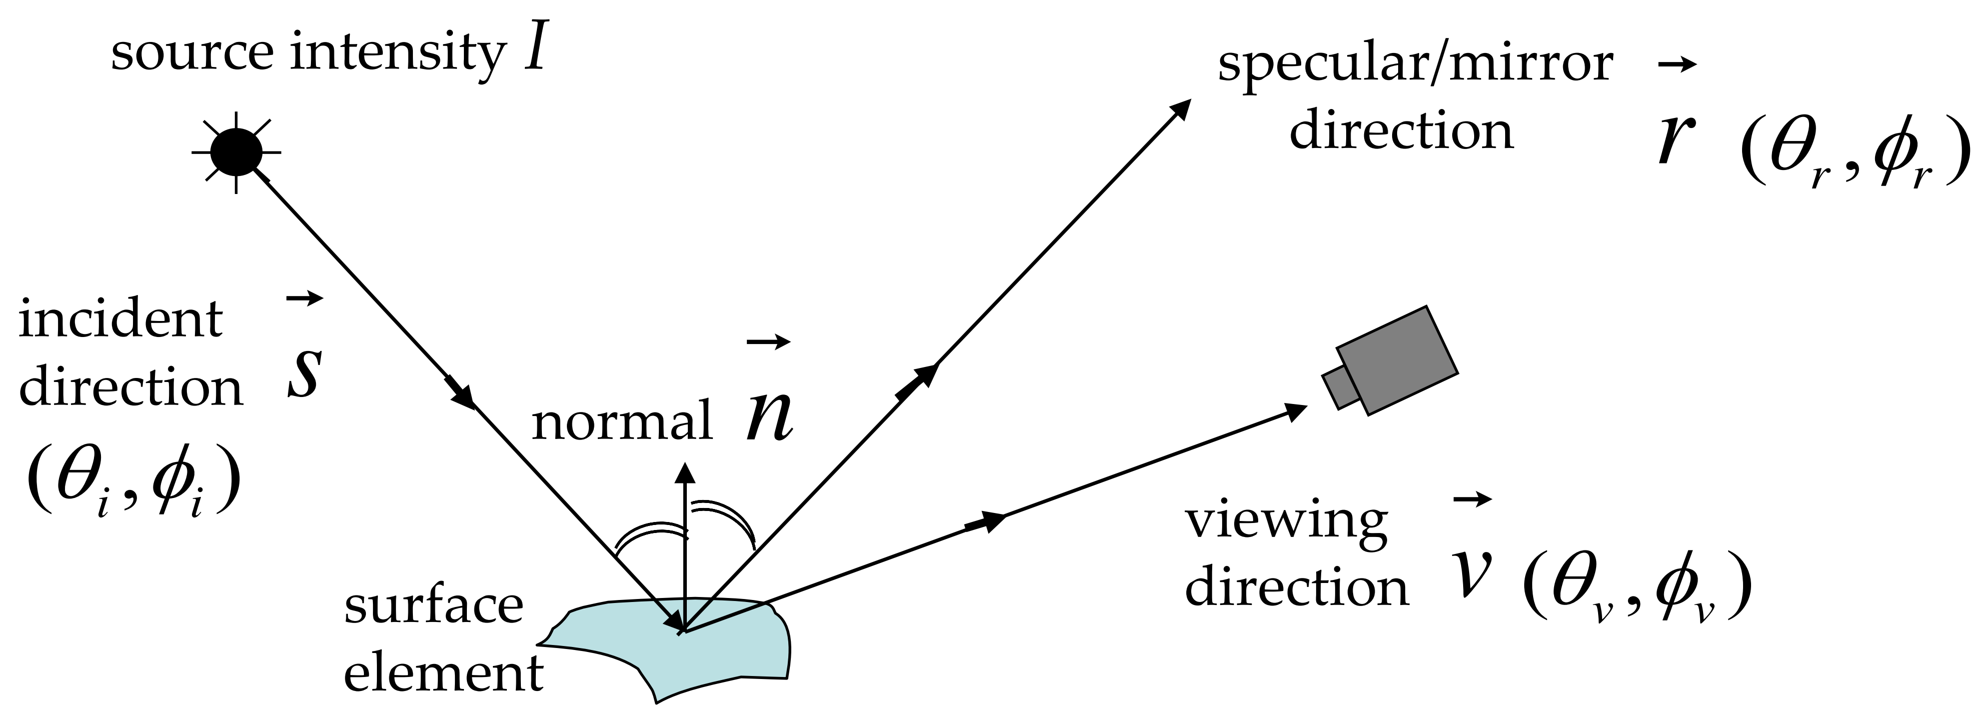

- Very smooth surface.
- All incident light energy reflected in a SINGLE direction. (only when $\vec{v} = \vec{r}$)
- Mirror BRDF is simply a double-delta function :

specular albedo

$$f(\theta_i, \phi_i; \theta_v, \phi_v) = \rho_s \ \delta(\theta_i - \theta_v)\ \delta(\phi_i + \pi - \phi_v)$$

- Surface Radiance : $L = I \ \rho_s \ \delta(\theta_i - \theta_v)\ \delta(\phi_i + \pi - \phi_v)$

Specular Reflections in Nature

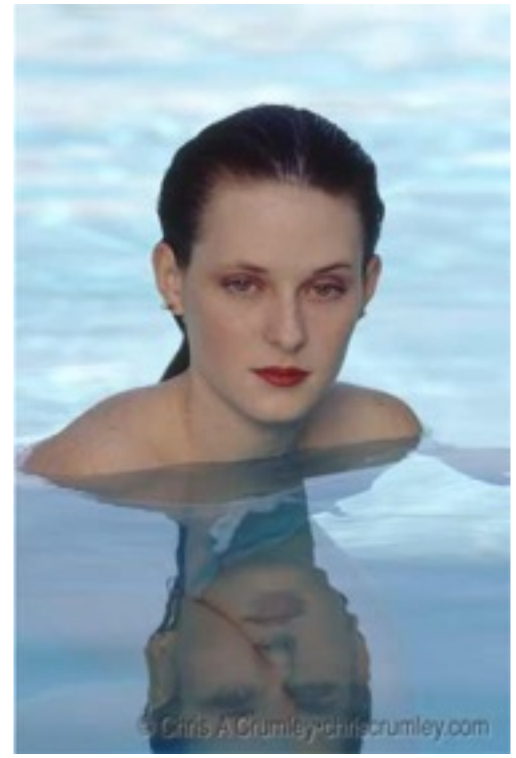

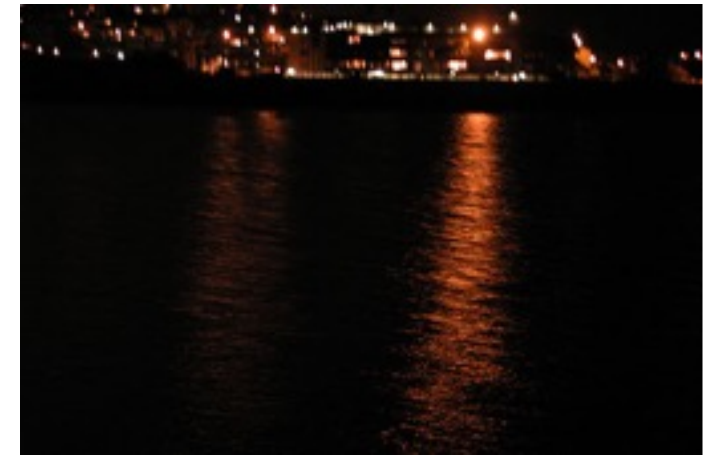

It's surprising how long the reflections are when viewed sitting on the river bank.

Compare sizes of objects and their reflections!

The reflections when seen from a lower view point are always longer than when viewed from a higher view point.

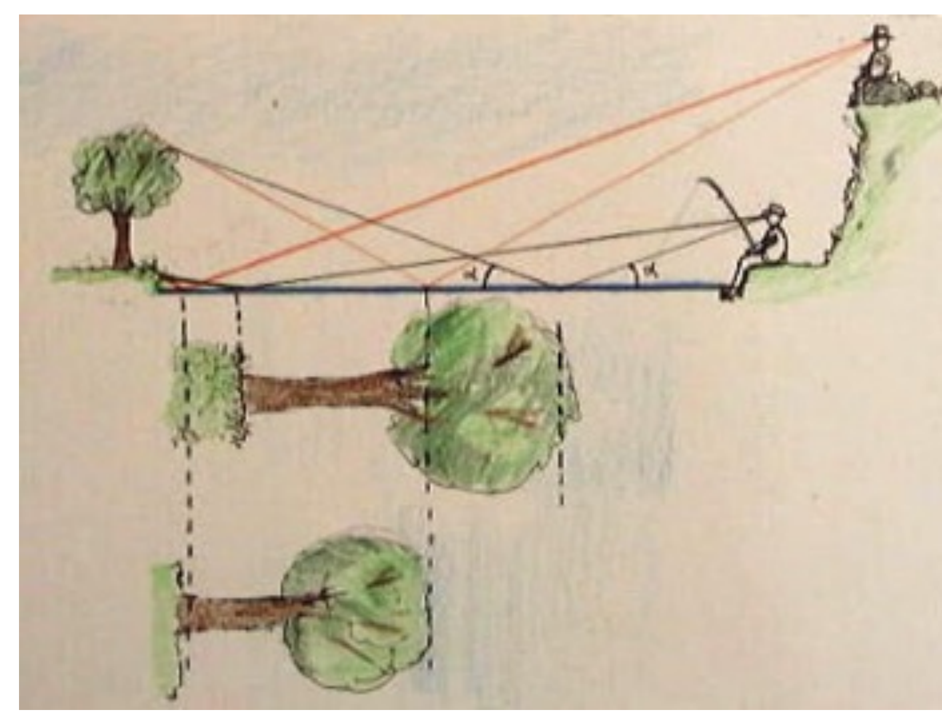

Intermeddiate regimes

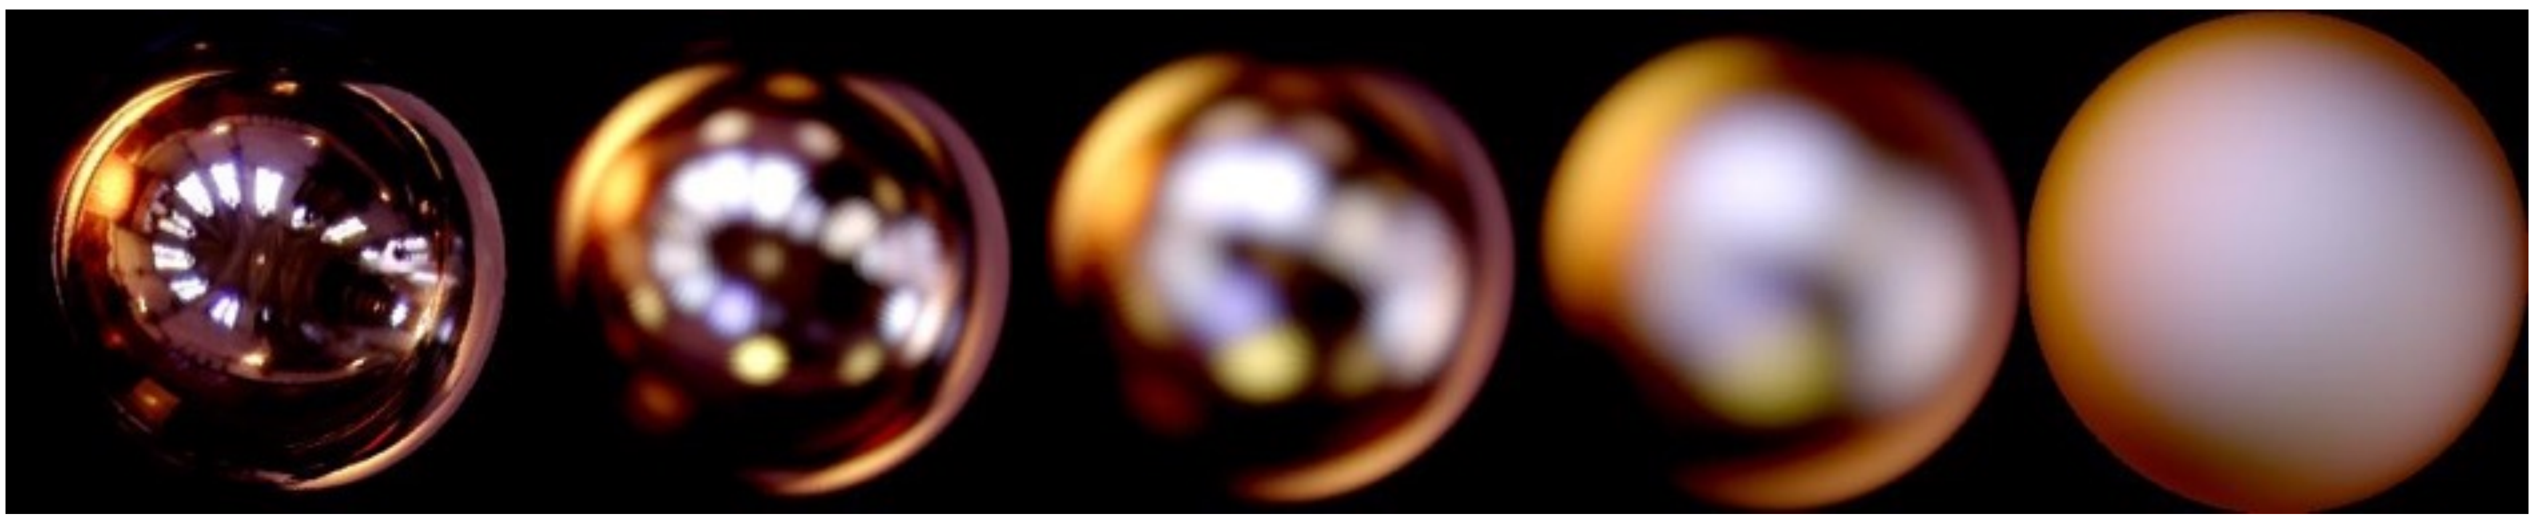

Specular/mirror

Diffuse / lambertian

Spectral BRDFs

Photon particles move with particular wavelengths, so incorporate these terms into rendering model

$$f(\omega_i, \omega_r) = \frac{L_r(\omega_r)}{L_i(\omega_i)\cos\theta_i}$$

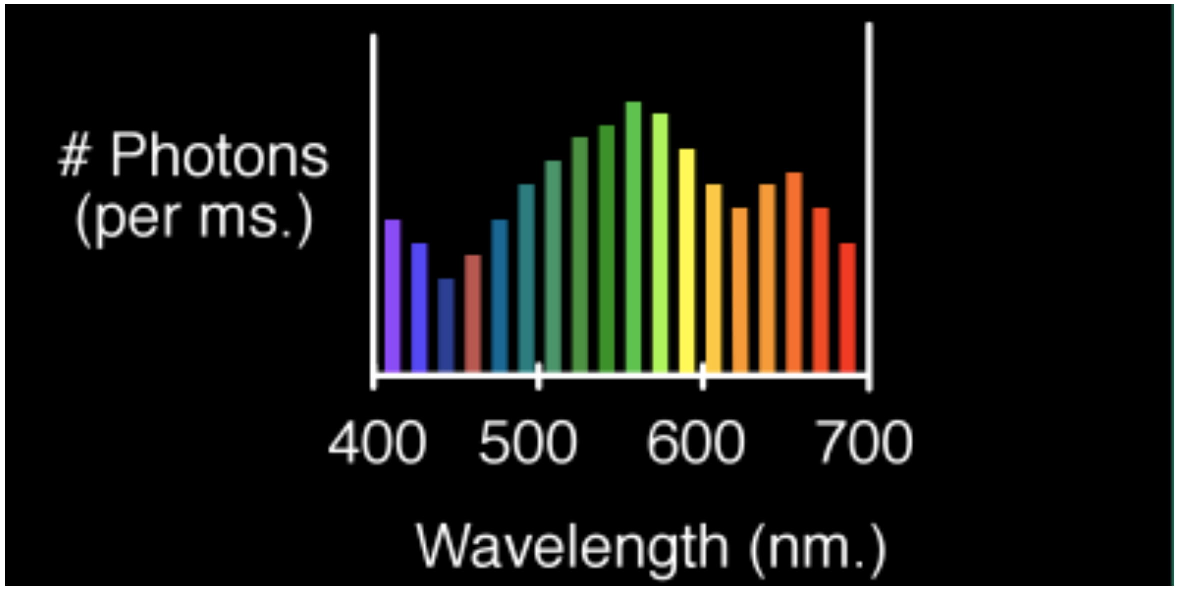

$$f(\omega_i, \omega_r, \lambda) = \frac{L_r(\omega_r, \lambda)}{L_i(\omega_i, \lambda)\cos\theta_i}$$

Outline

1. Final Projects
2. Radiometry
3. Color

Color

Two ways to talk about; color of *light* (additive) or color of *surfaces* (subtractive)

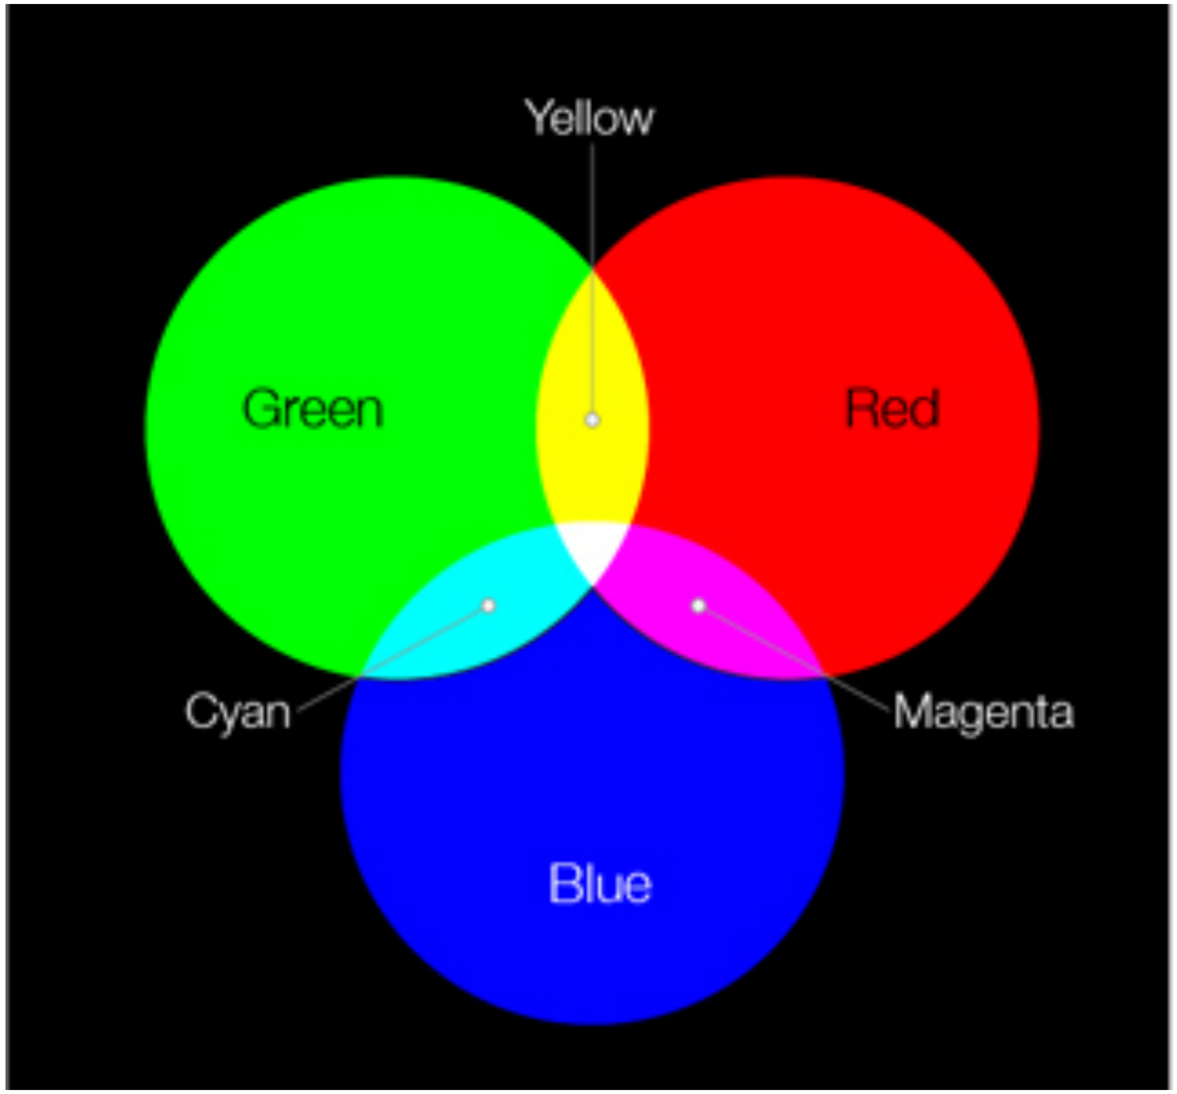

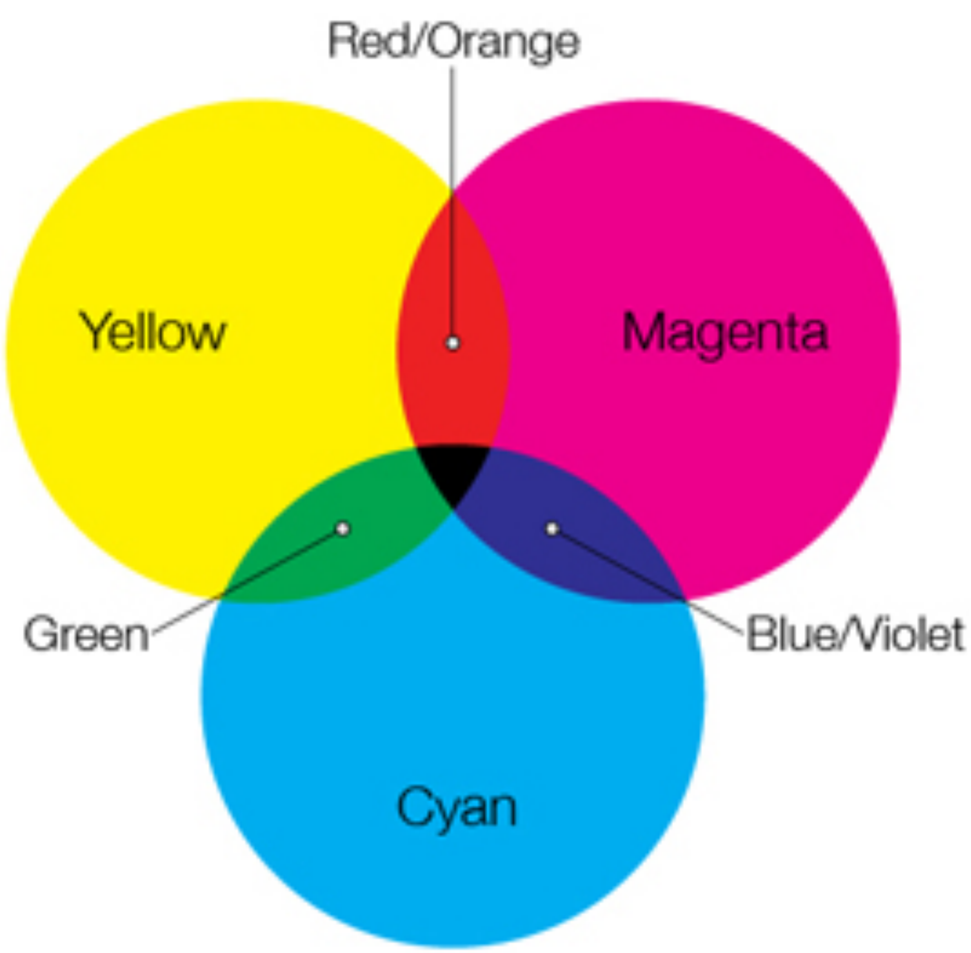

Spectral radiance

$$L(\omega, \lambda)$$

Spectral BRDFs

$$f(\omega_i, \omega_r, \lambda) \approx \rho(\lambda)$$

(spectral) abledo is a constant approximation of BRDF

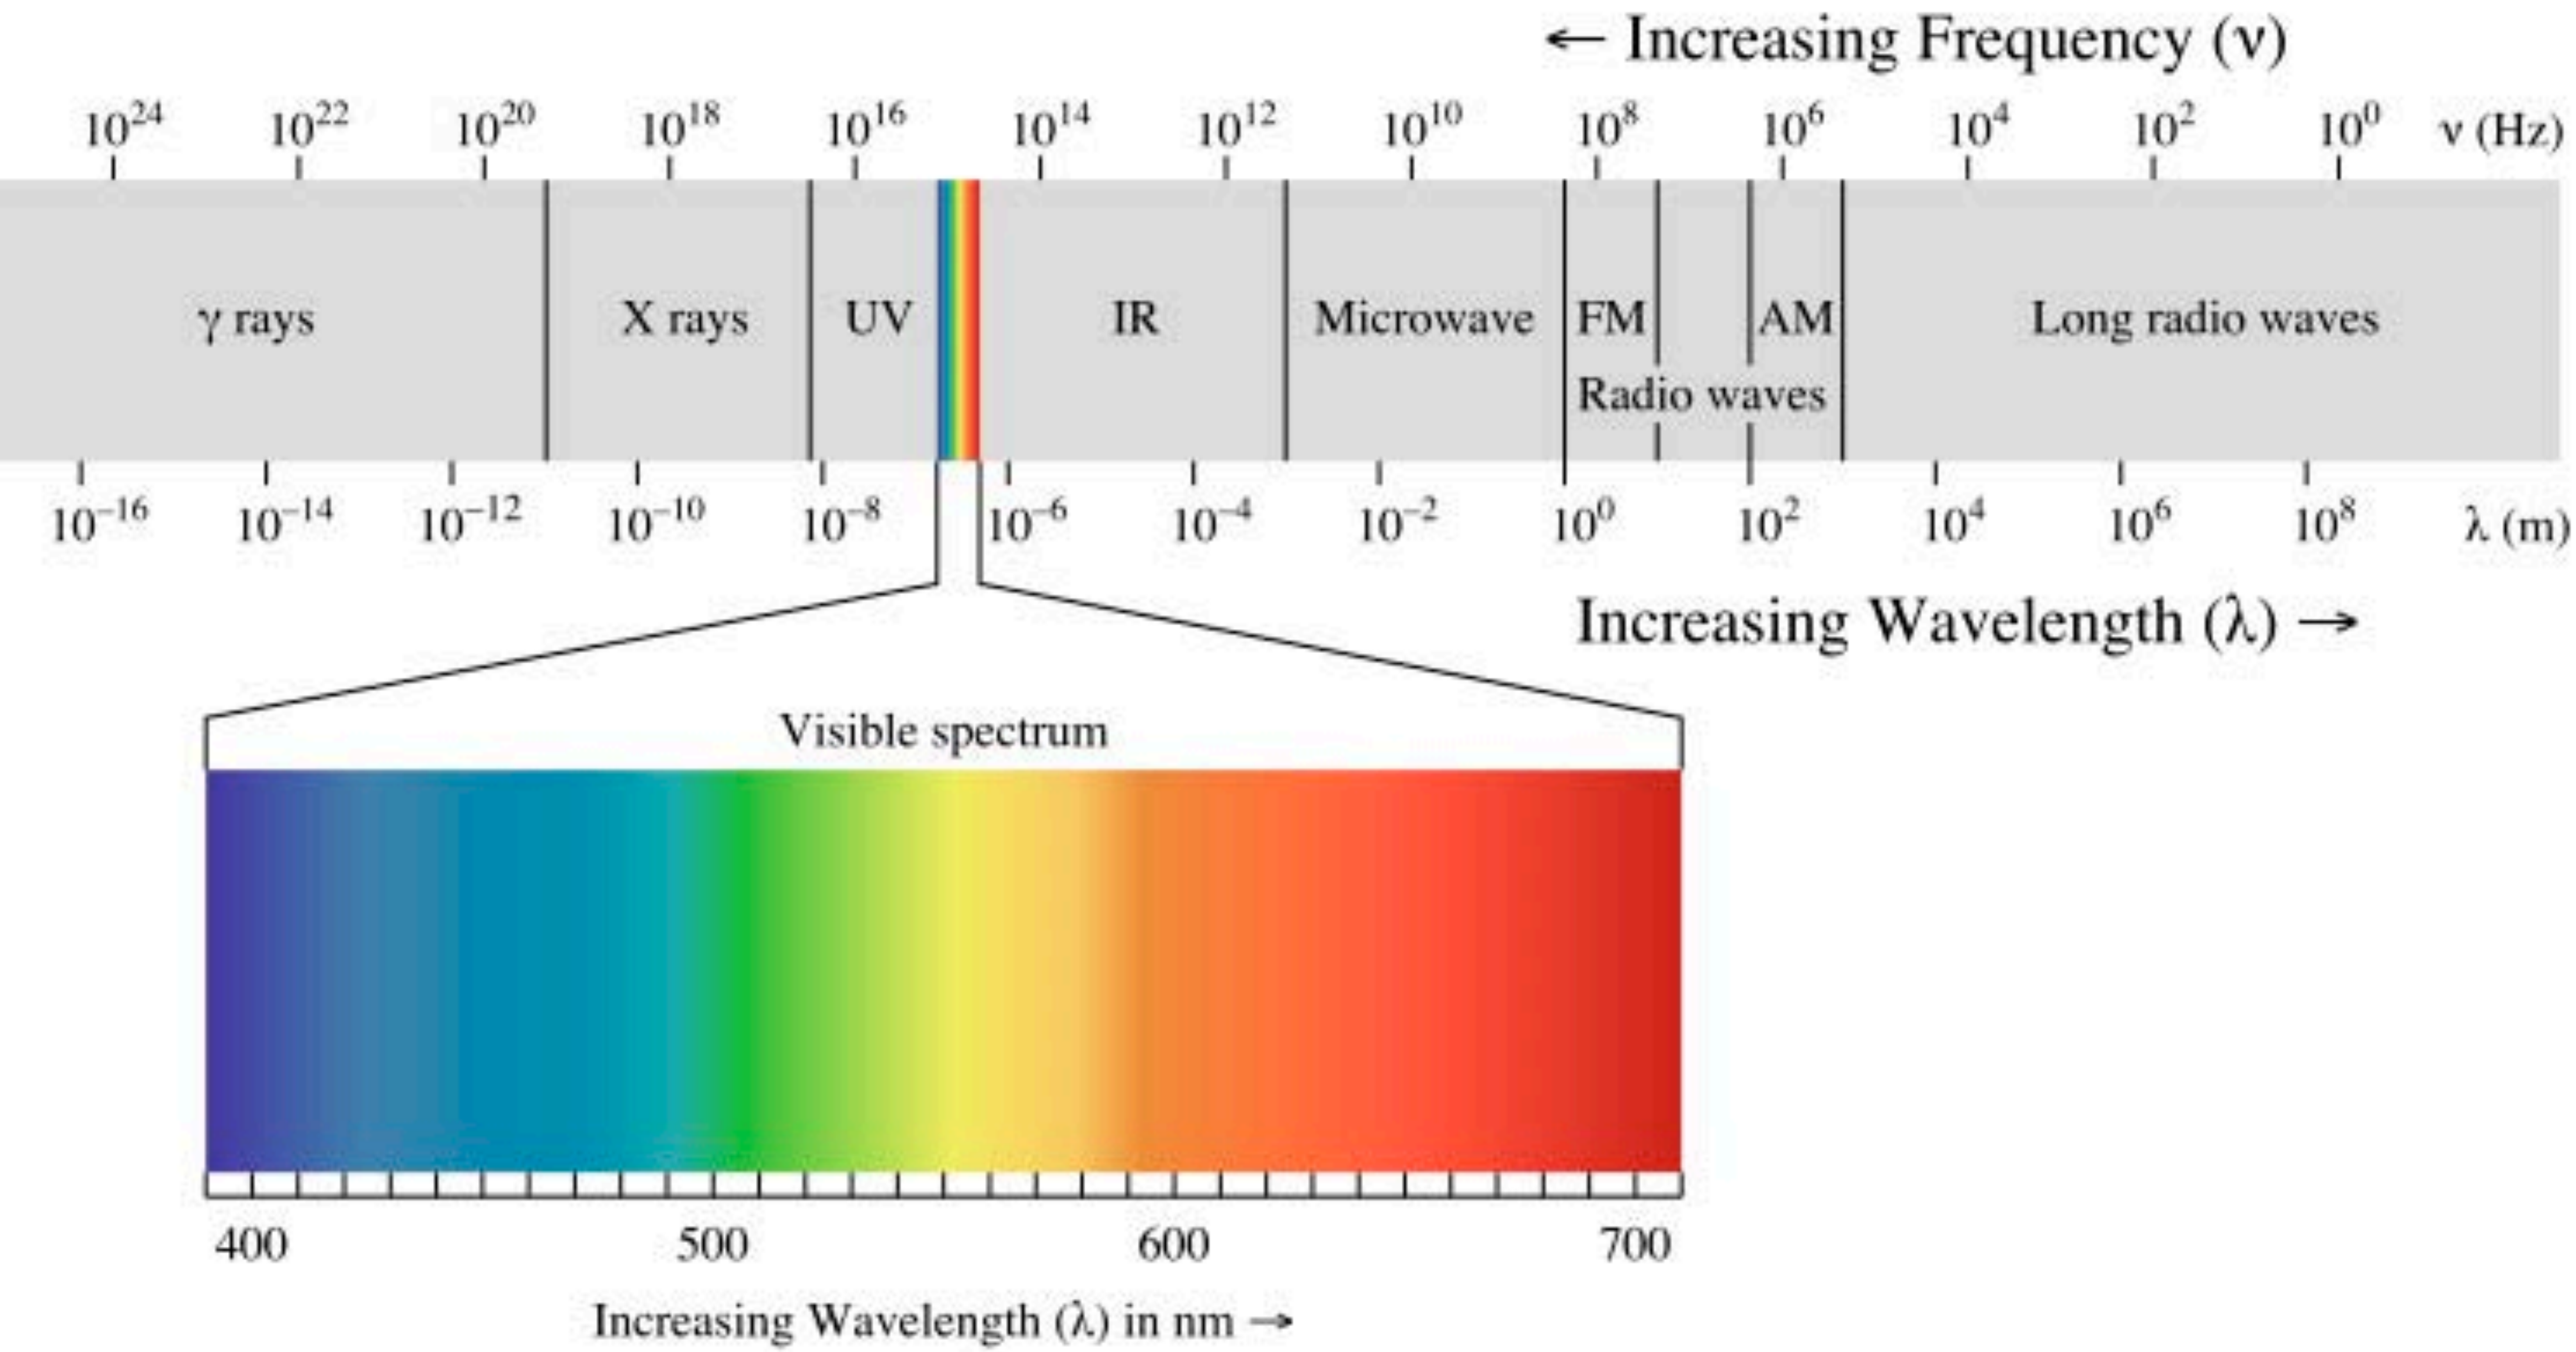
← Increasing Frequency (ν)
10^24 10^22 10^20 10^18 10^16 10^14 10^12 10^10 10^8 10^6 10^4 10^2 10^0 ν (Hz)
γ rays
X rays
UV
IR
Microwave
FM
AM
Long radio waves
Radio waves
10^−16 10^−14 10^−12 10^−10 10^−8 10^−6 10^−4 10^−2 10^0 10^2 10^4 10^6 10^8 λ (m)
Increasing Wavelength (λ) →
Visible spectrum
400
500
600
700
Increasing Wavelength (λ) in nm →

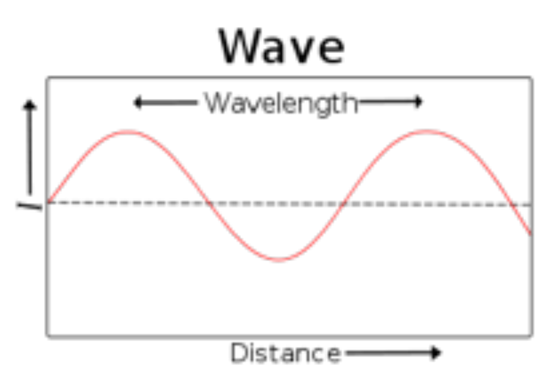
Wave
Wavelength
Distance

Any patch of light can be completely described physically by its spectrum: the number of photons (per time unit) at each wavelength 400 - 700 nm.

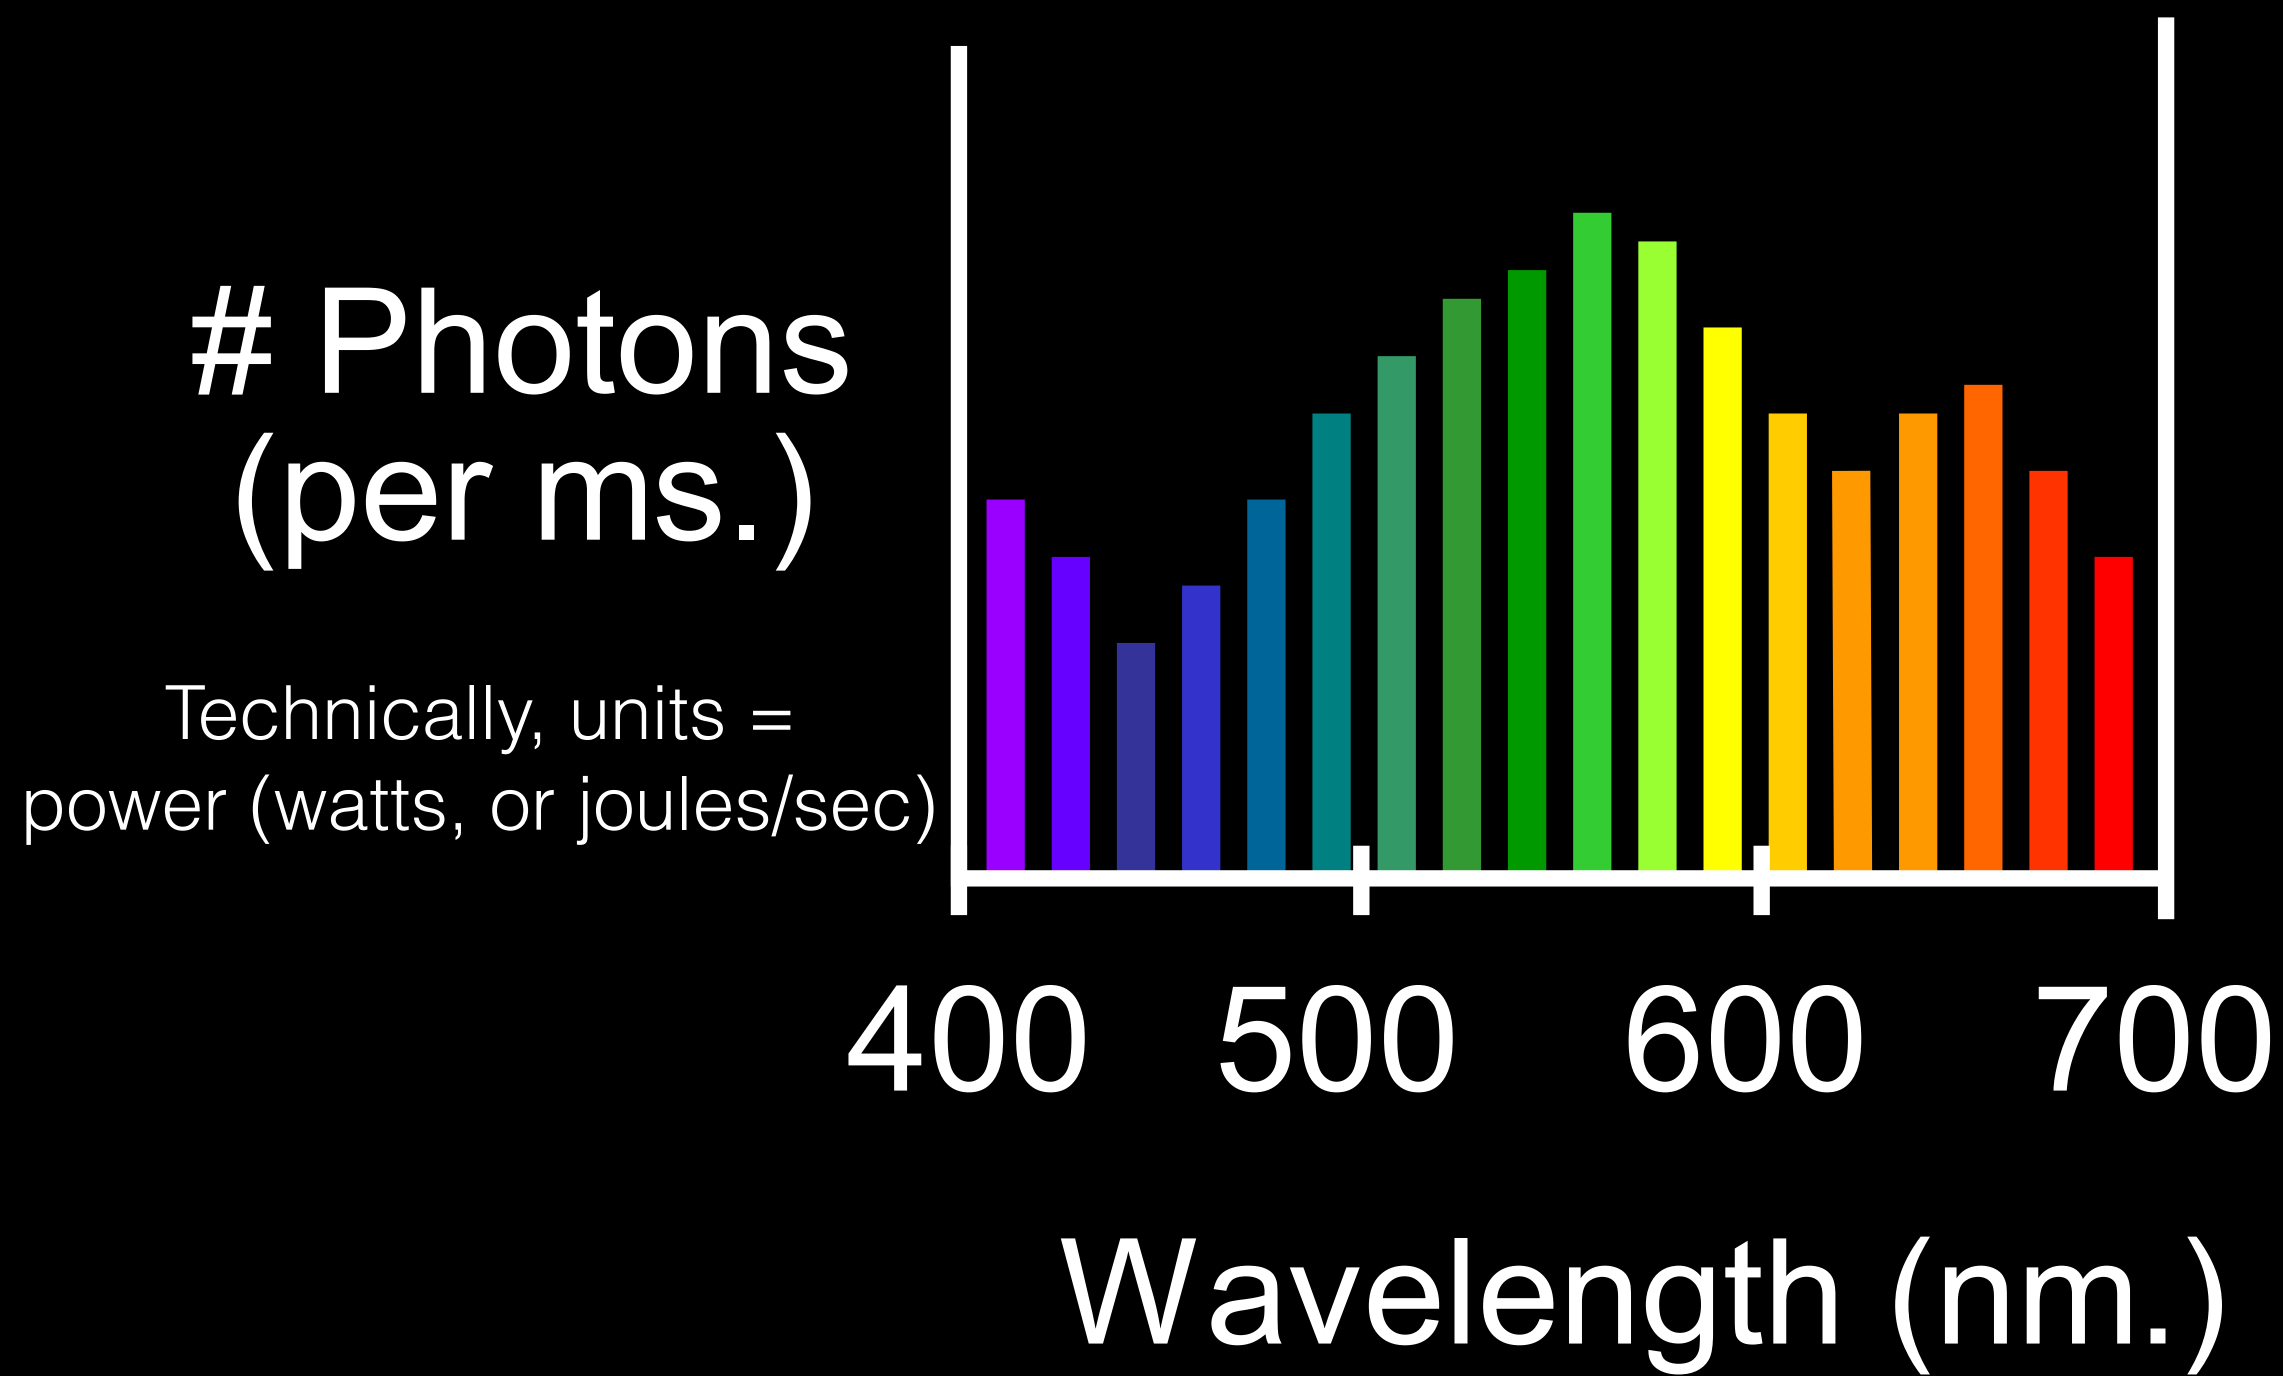

Some examples of the spectra of light sources

A. Ruby Laser

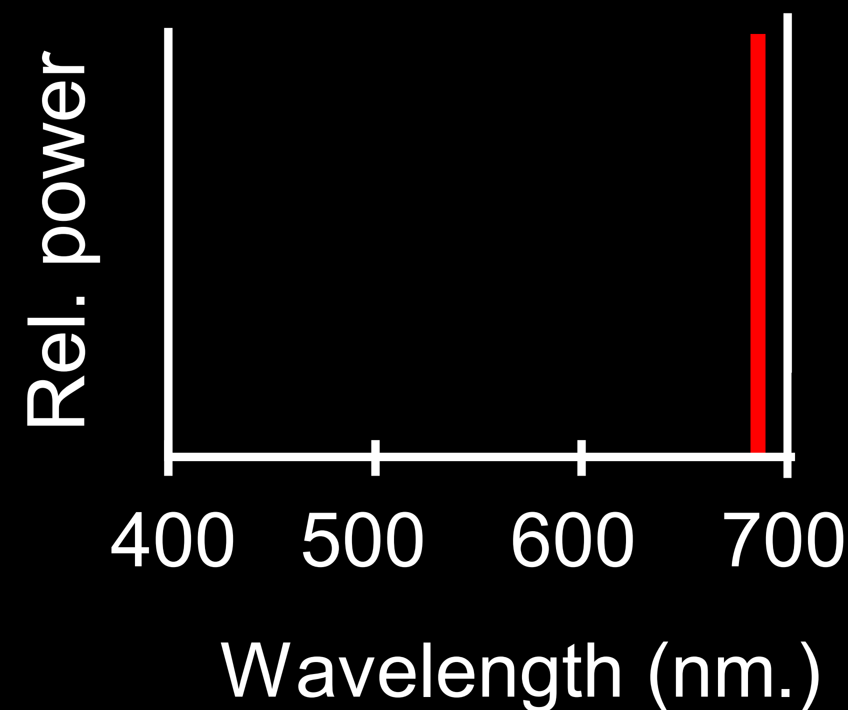

B. Gallium Phosphide Crystal

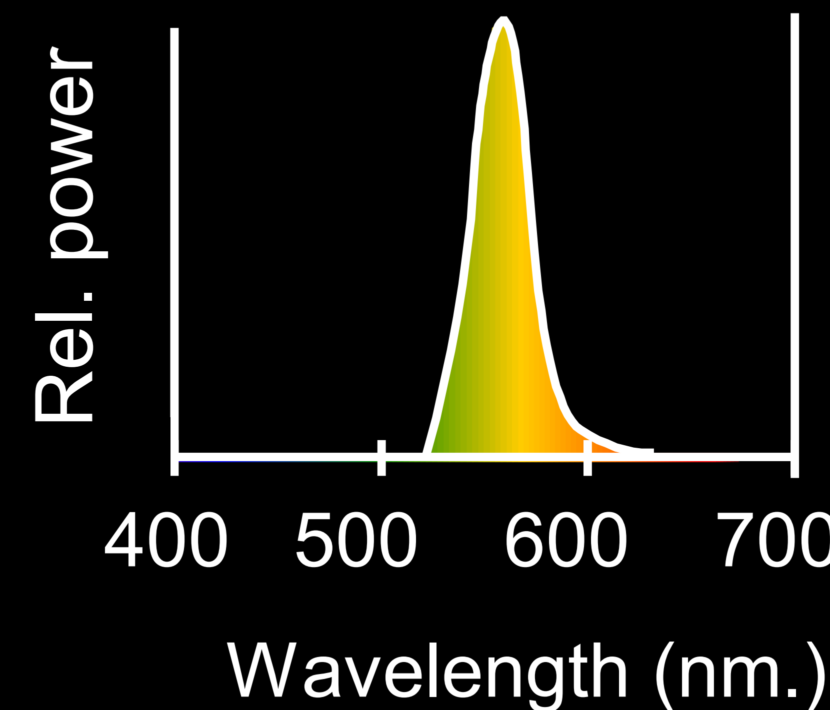

C. Tungsten Lightbulb

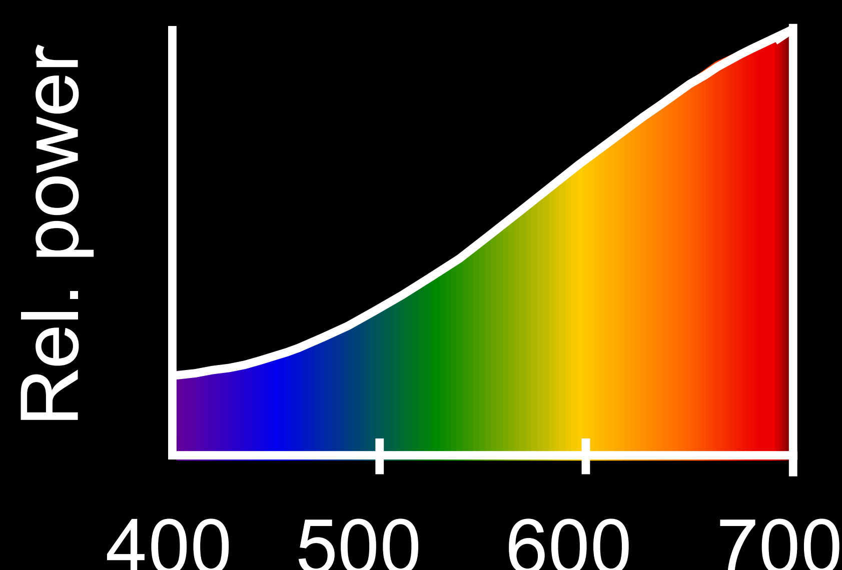

D. Normal Daylight

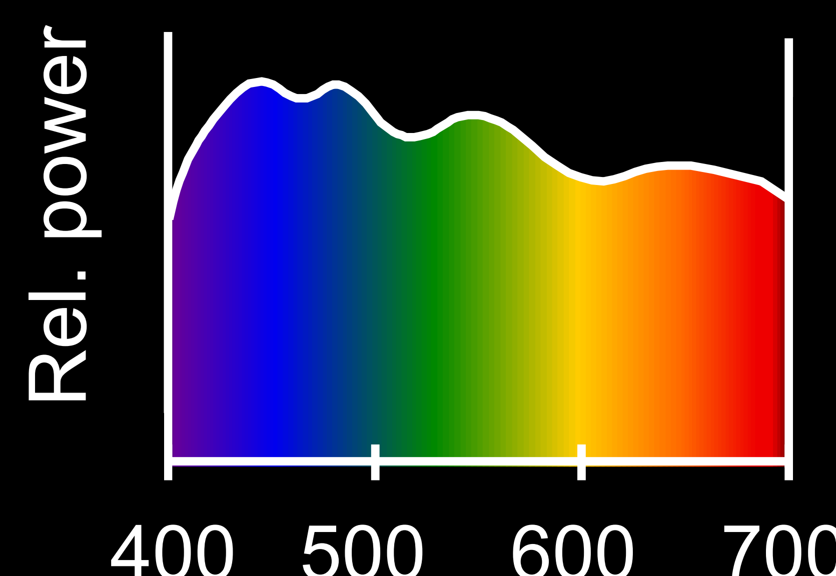

Some examples of the reflectance spectra of surfaces

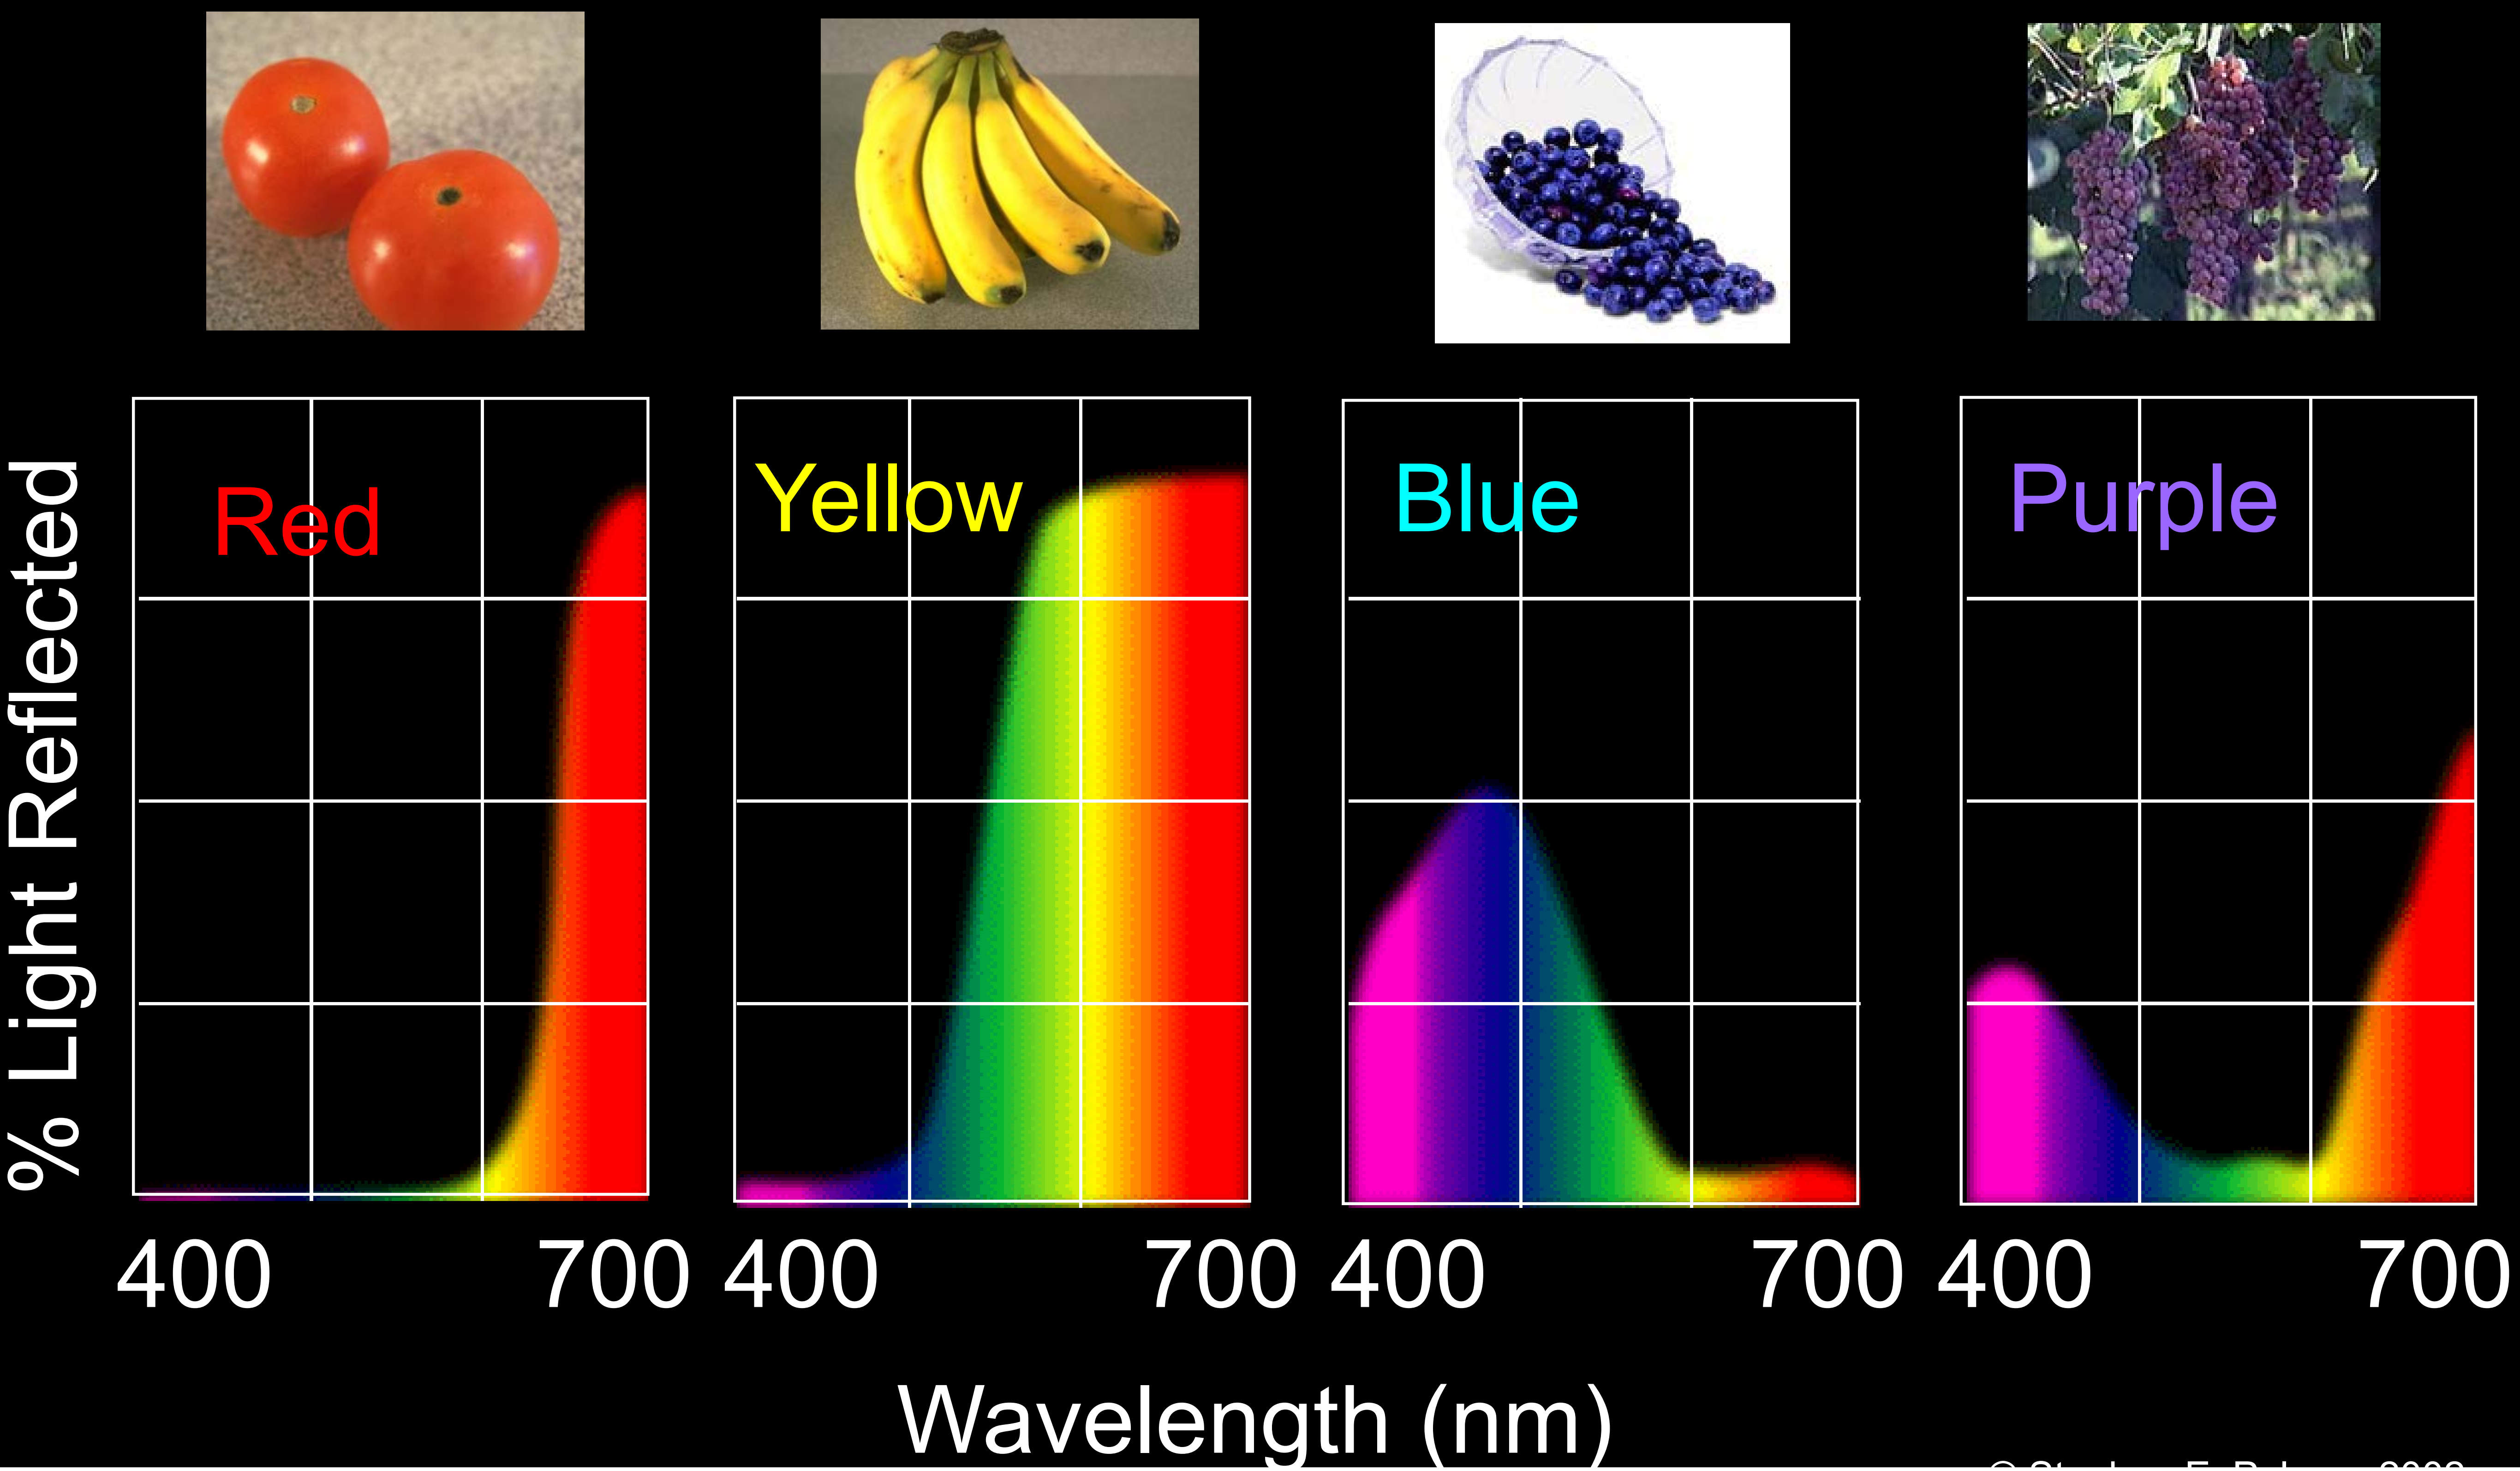

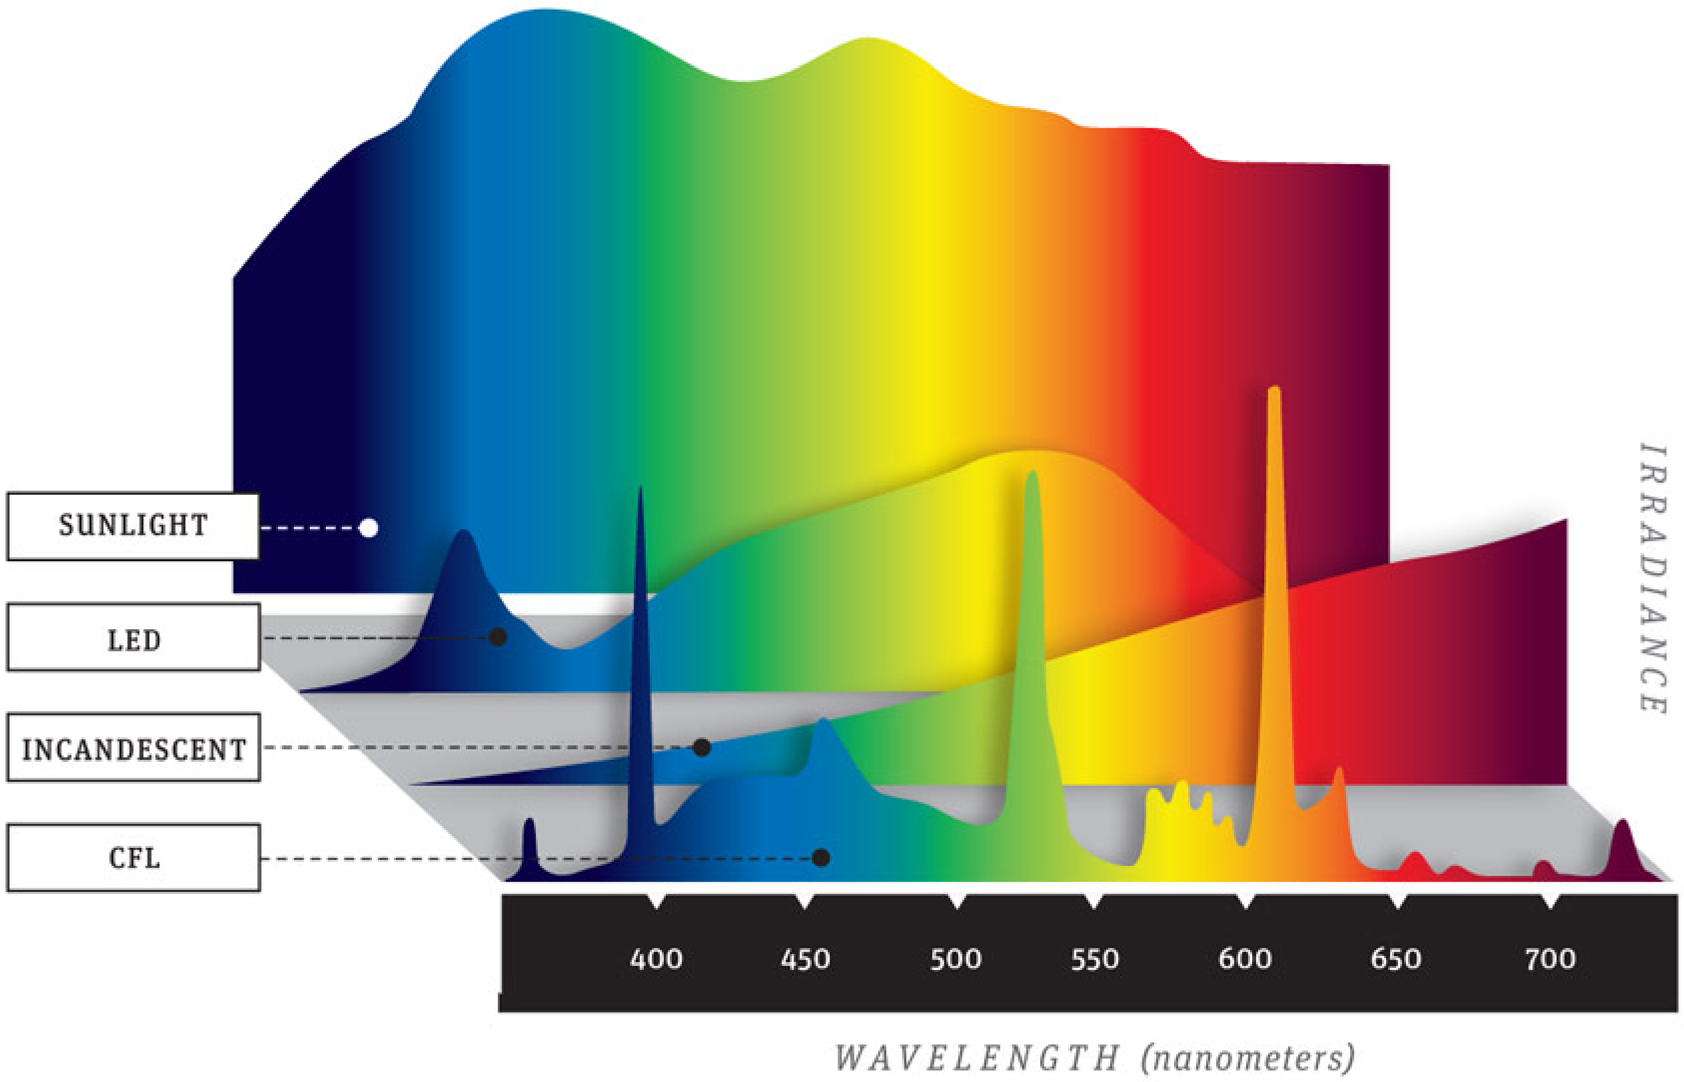

These plots suggest that 3 numbers are not sufficient to represent a color.
So what gives? (answer in a bit)

There is no simple functional description for the perceived color of all lights under all viewing conditions, but ……

A helpful constraint:

Consider only physical spectra with normal distributions

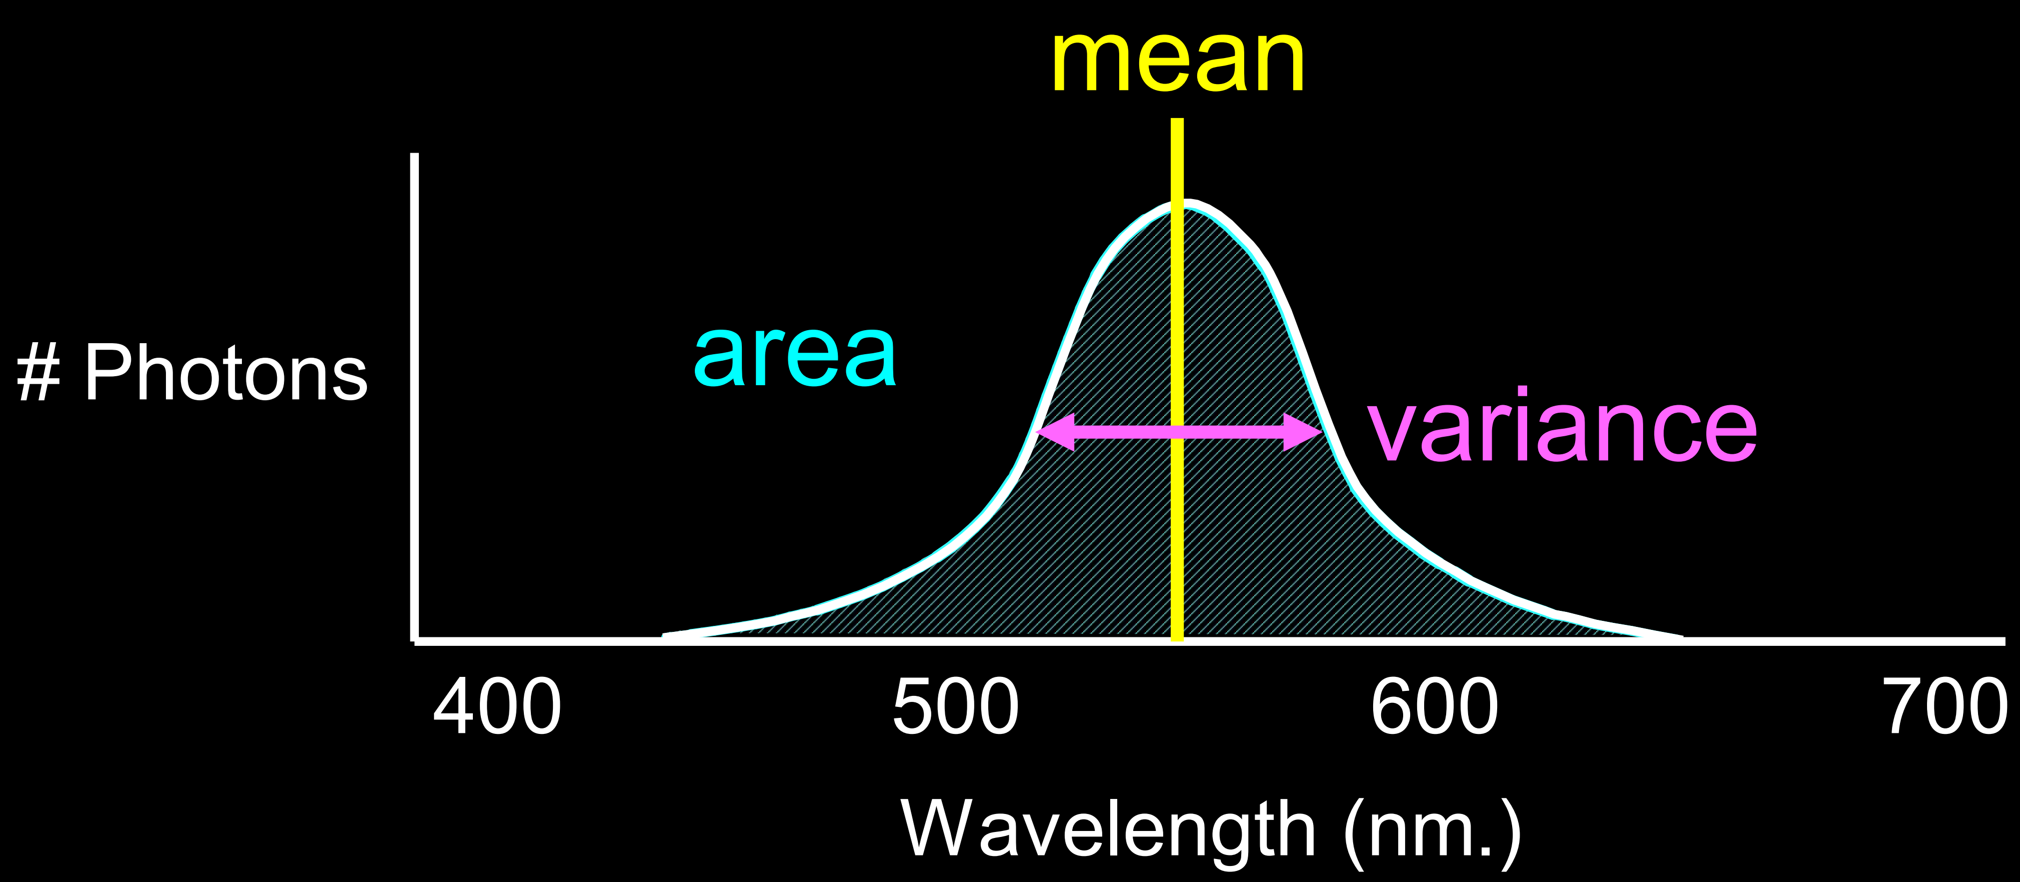

Mean ↔ Hue

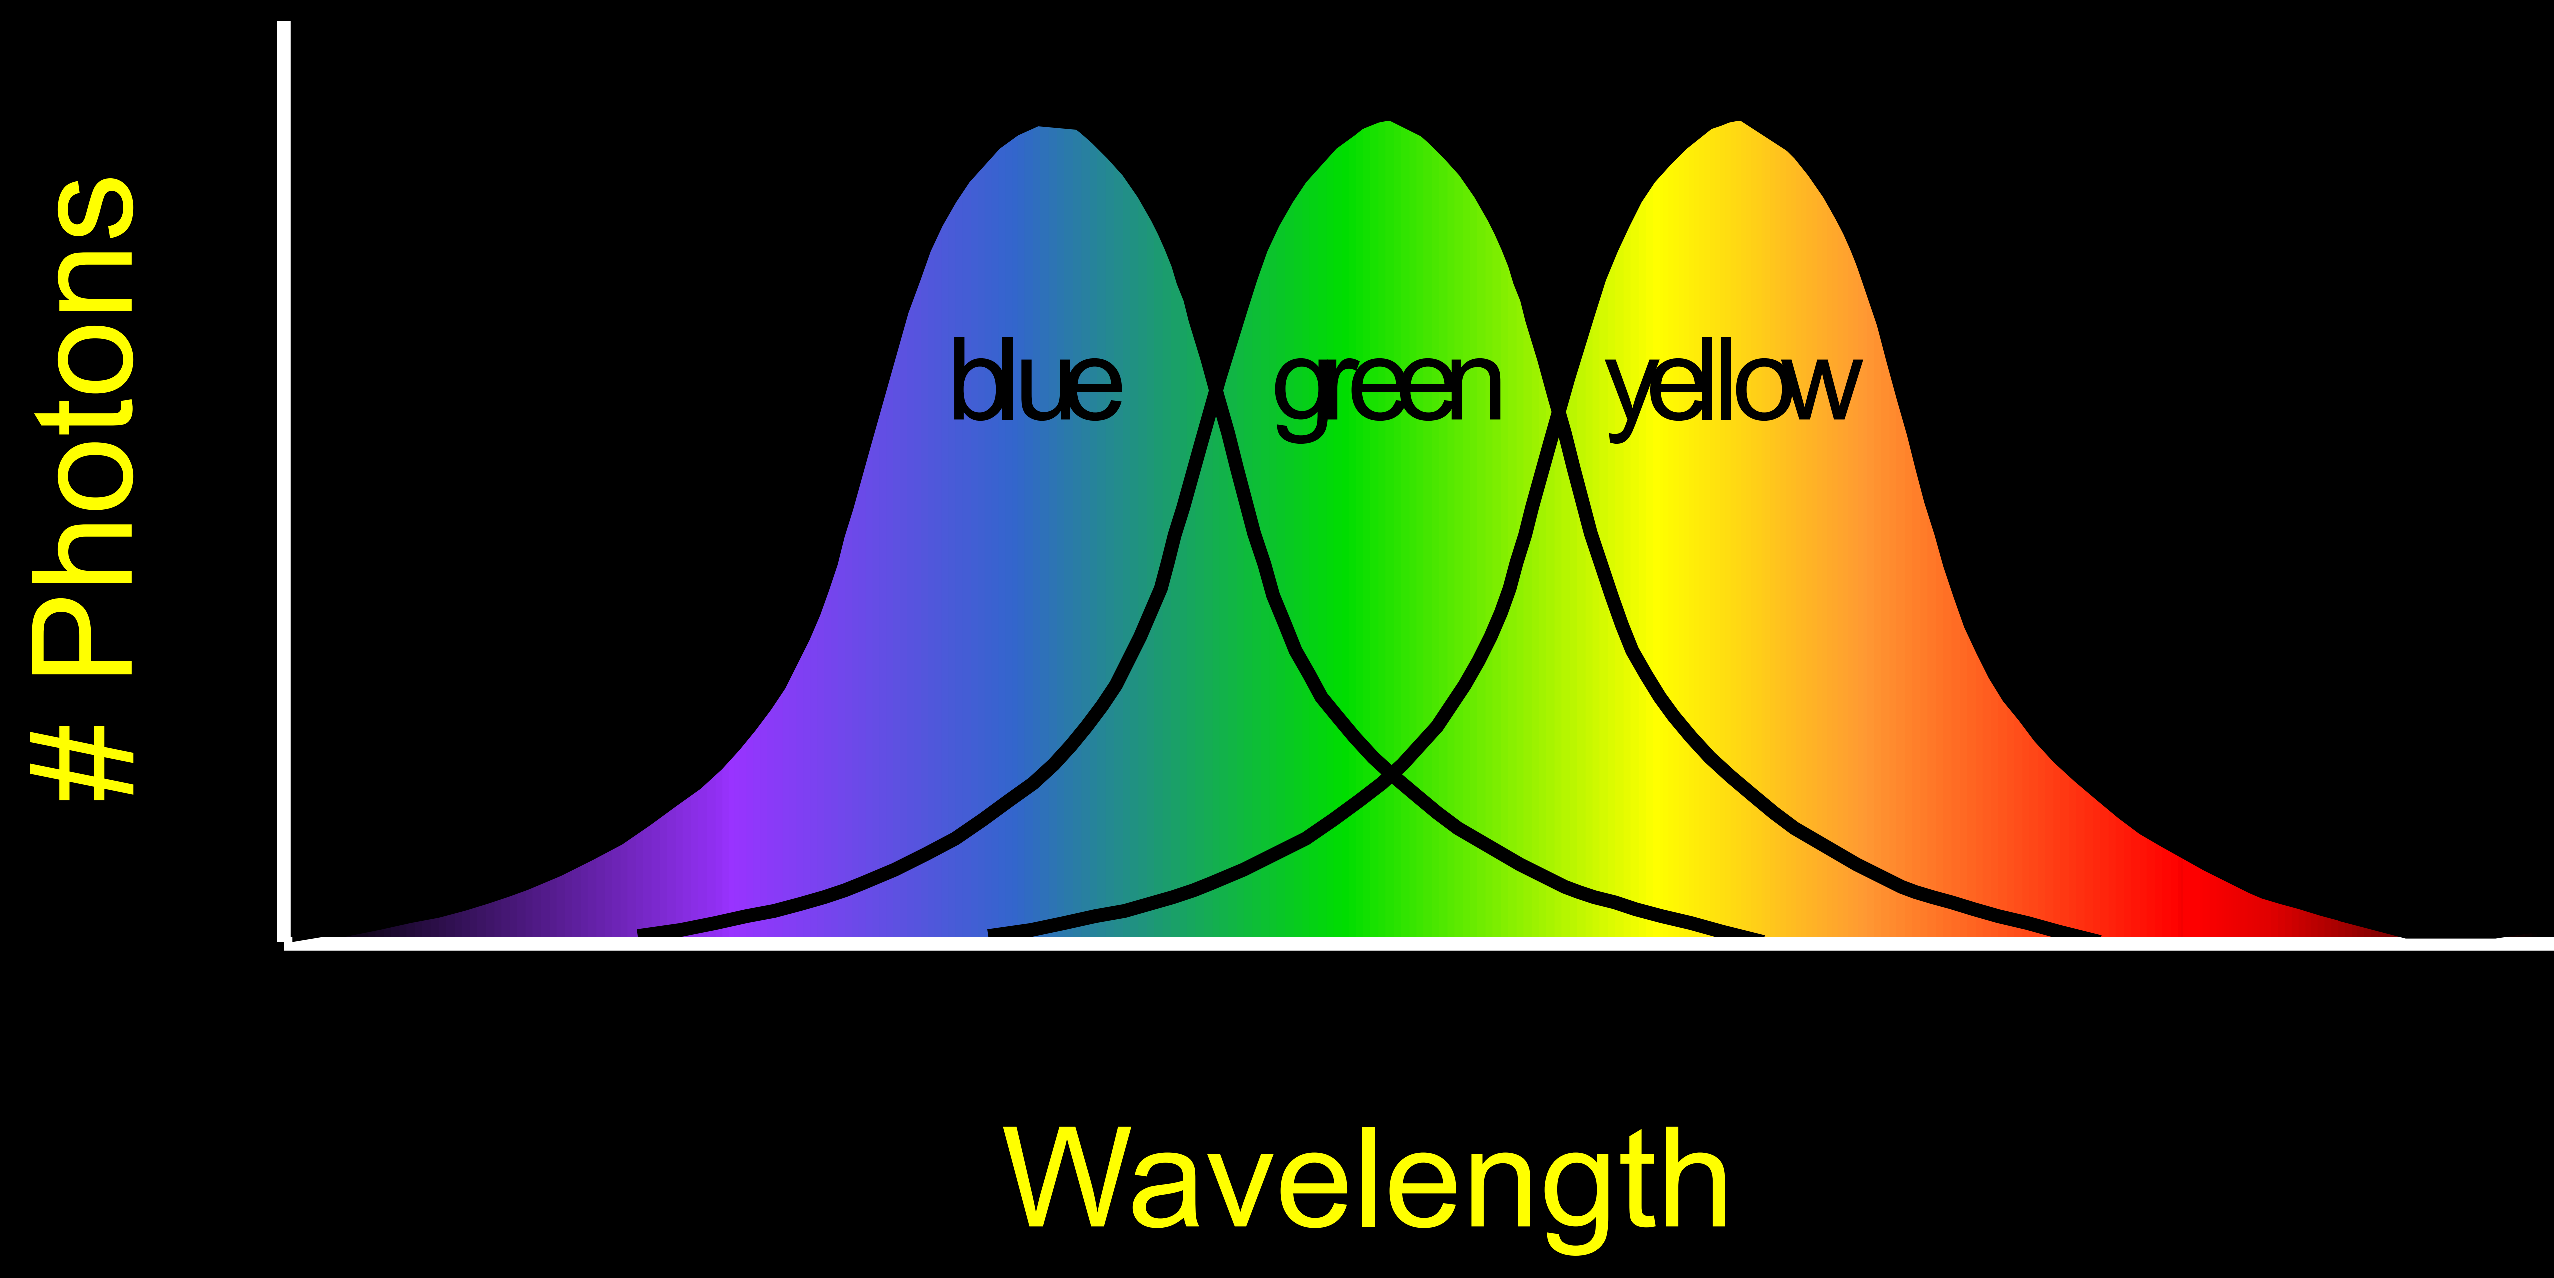

Variance ⟷ Saturation

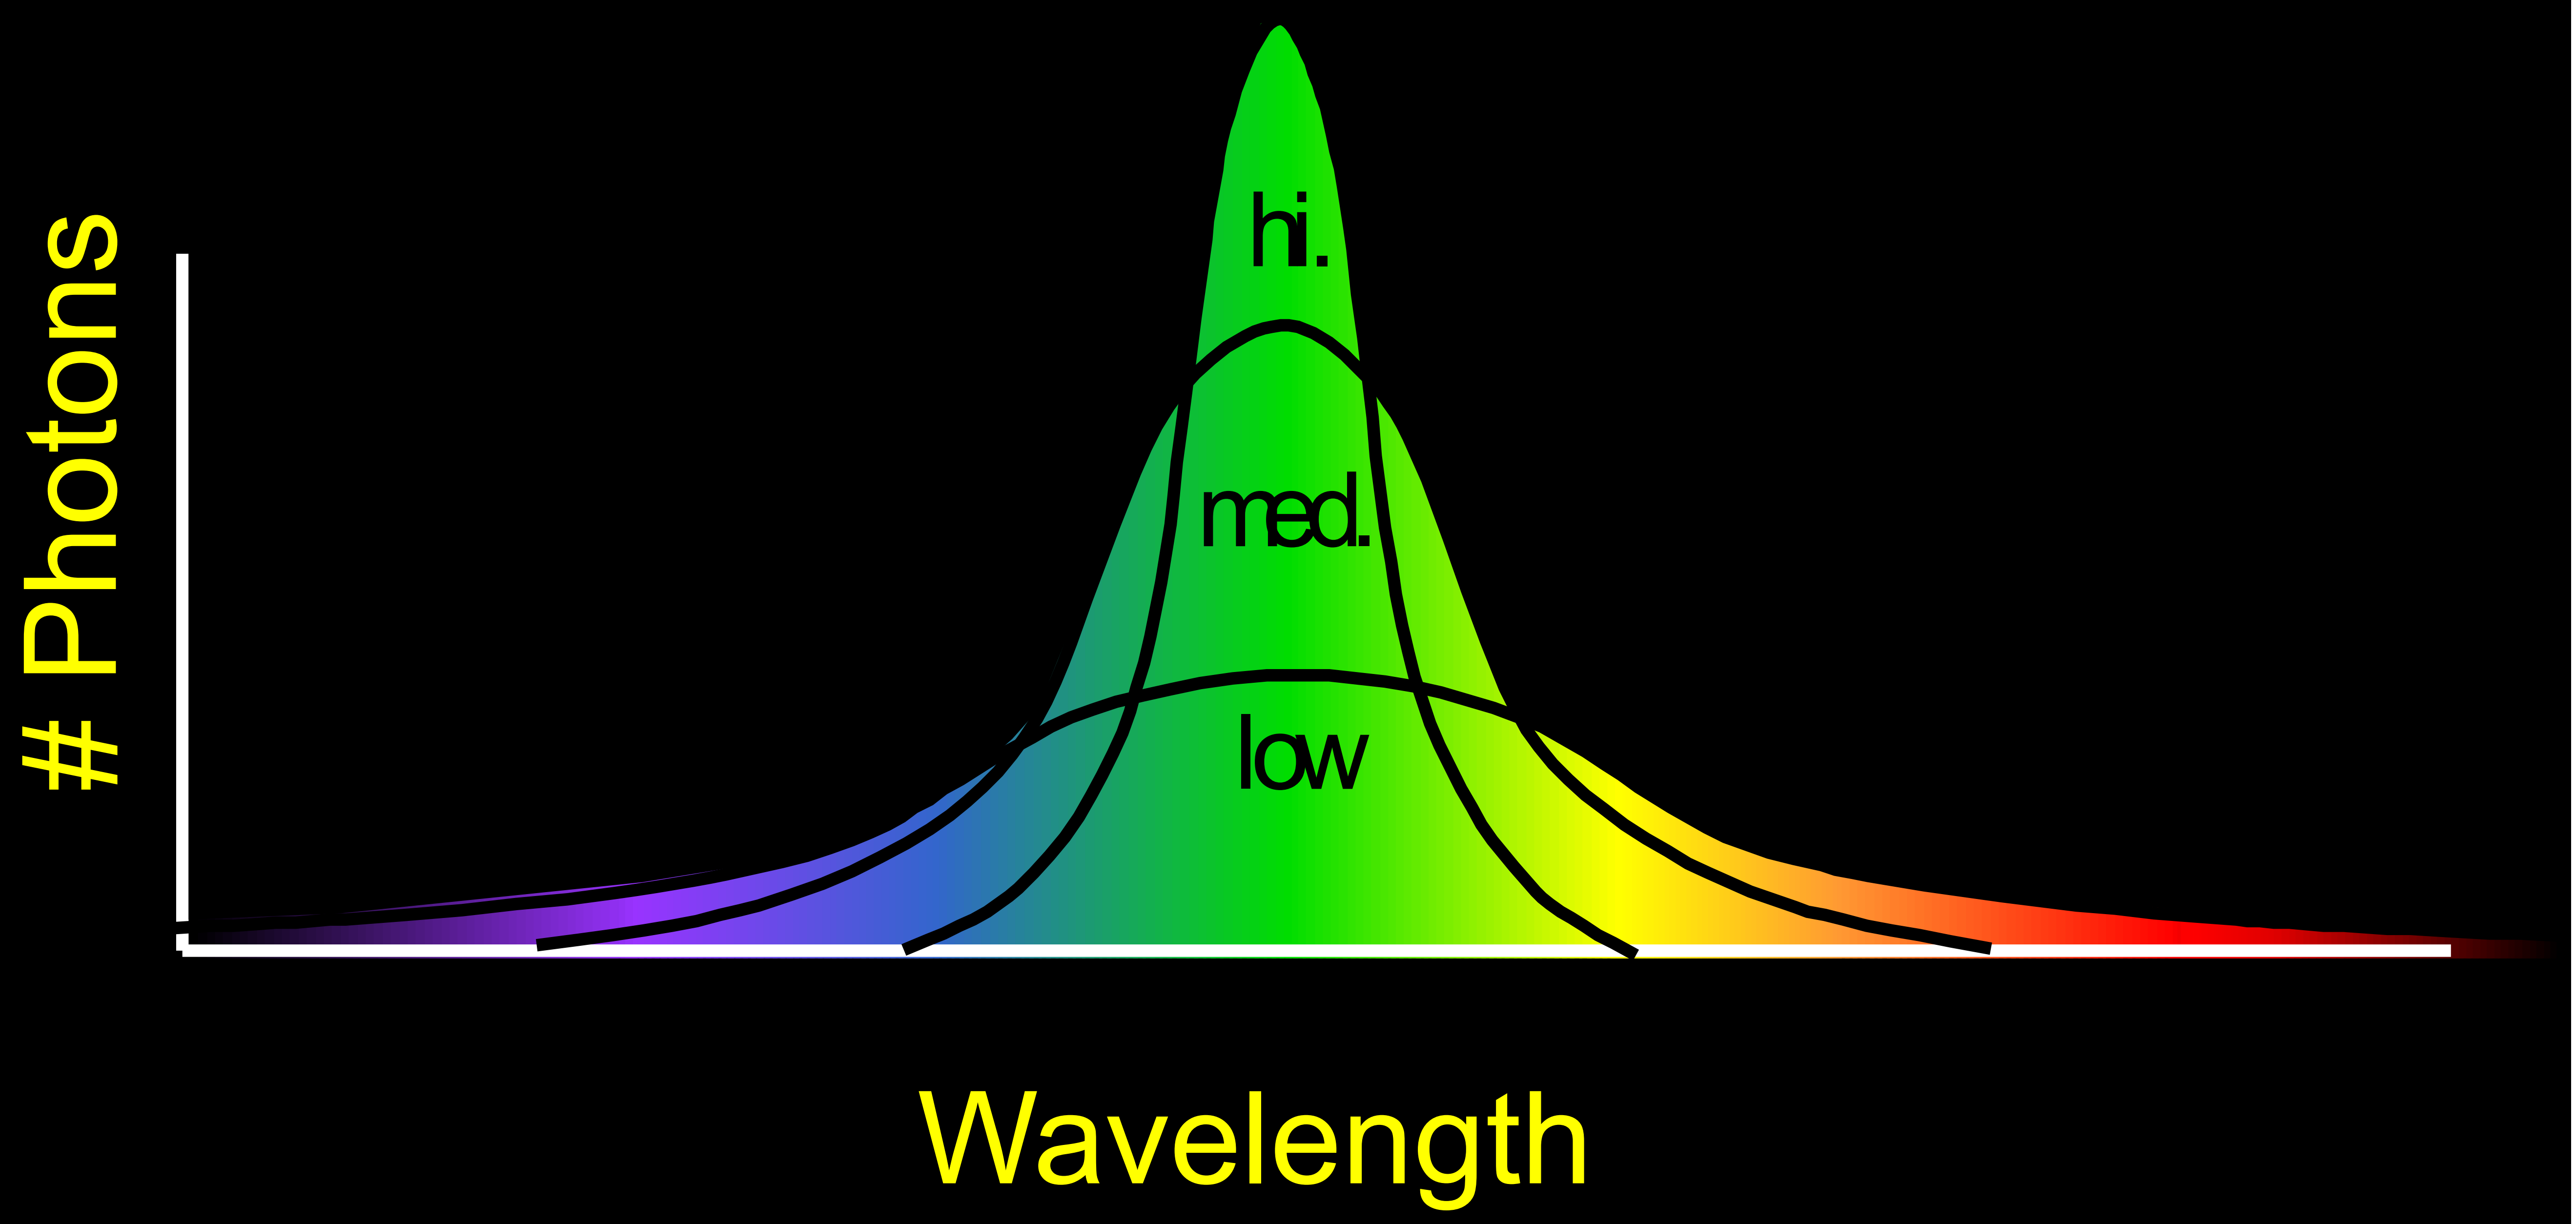

Area ⟷ Brightness

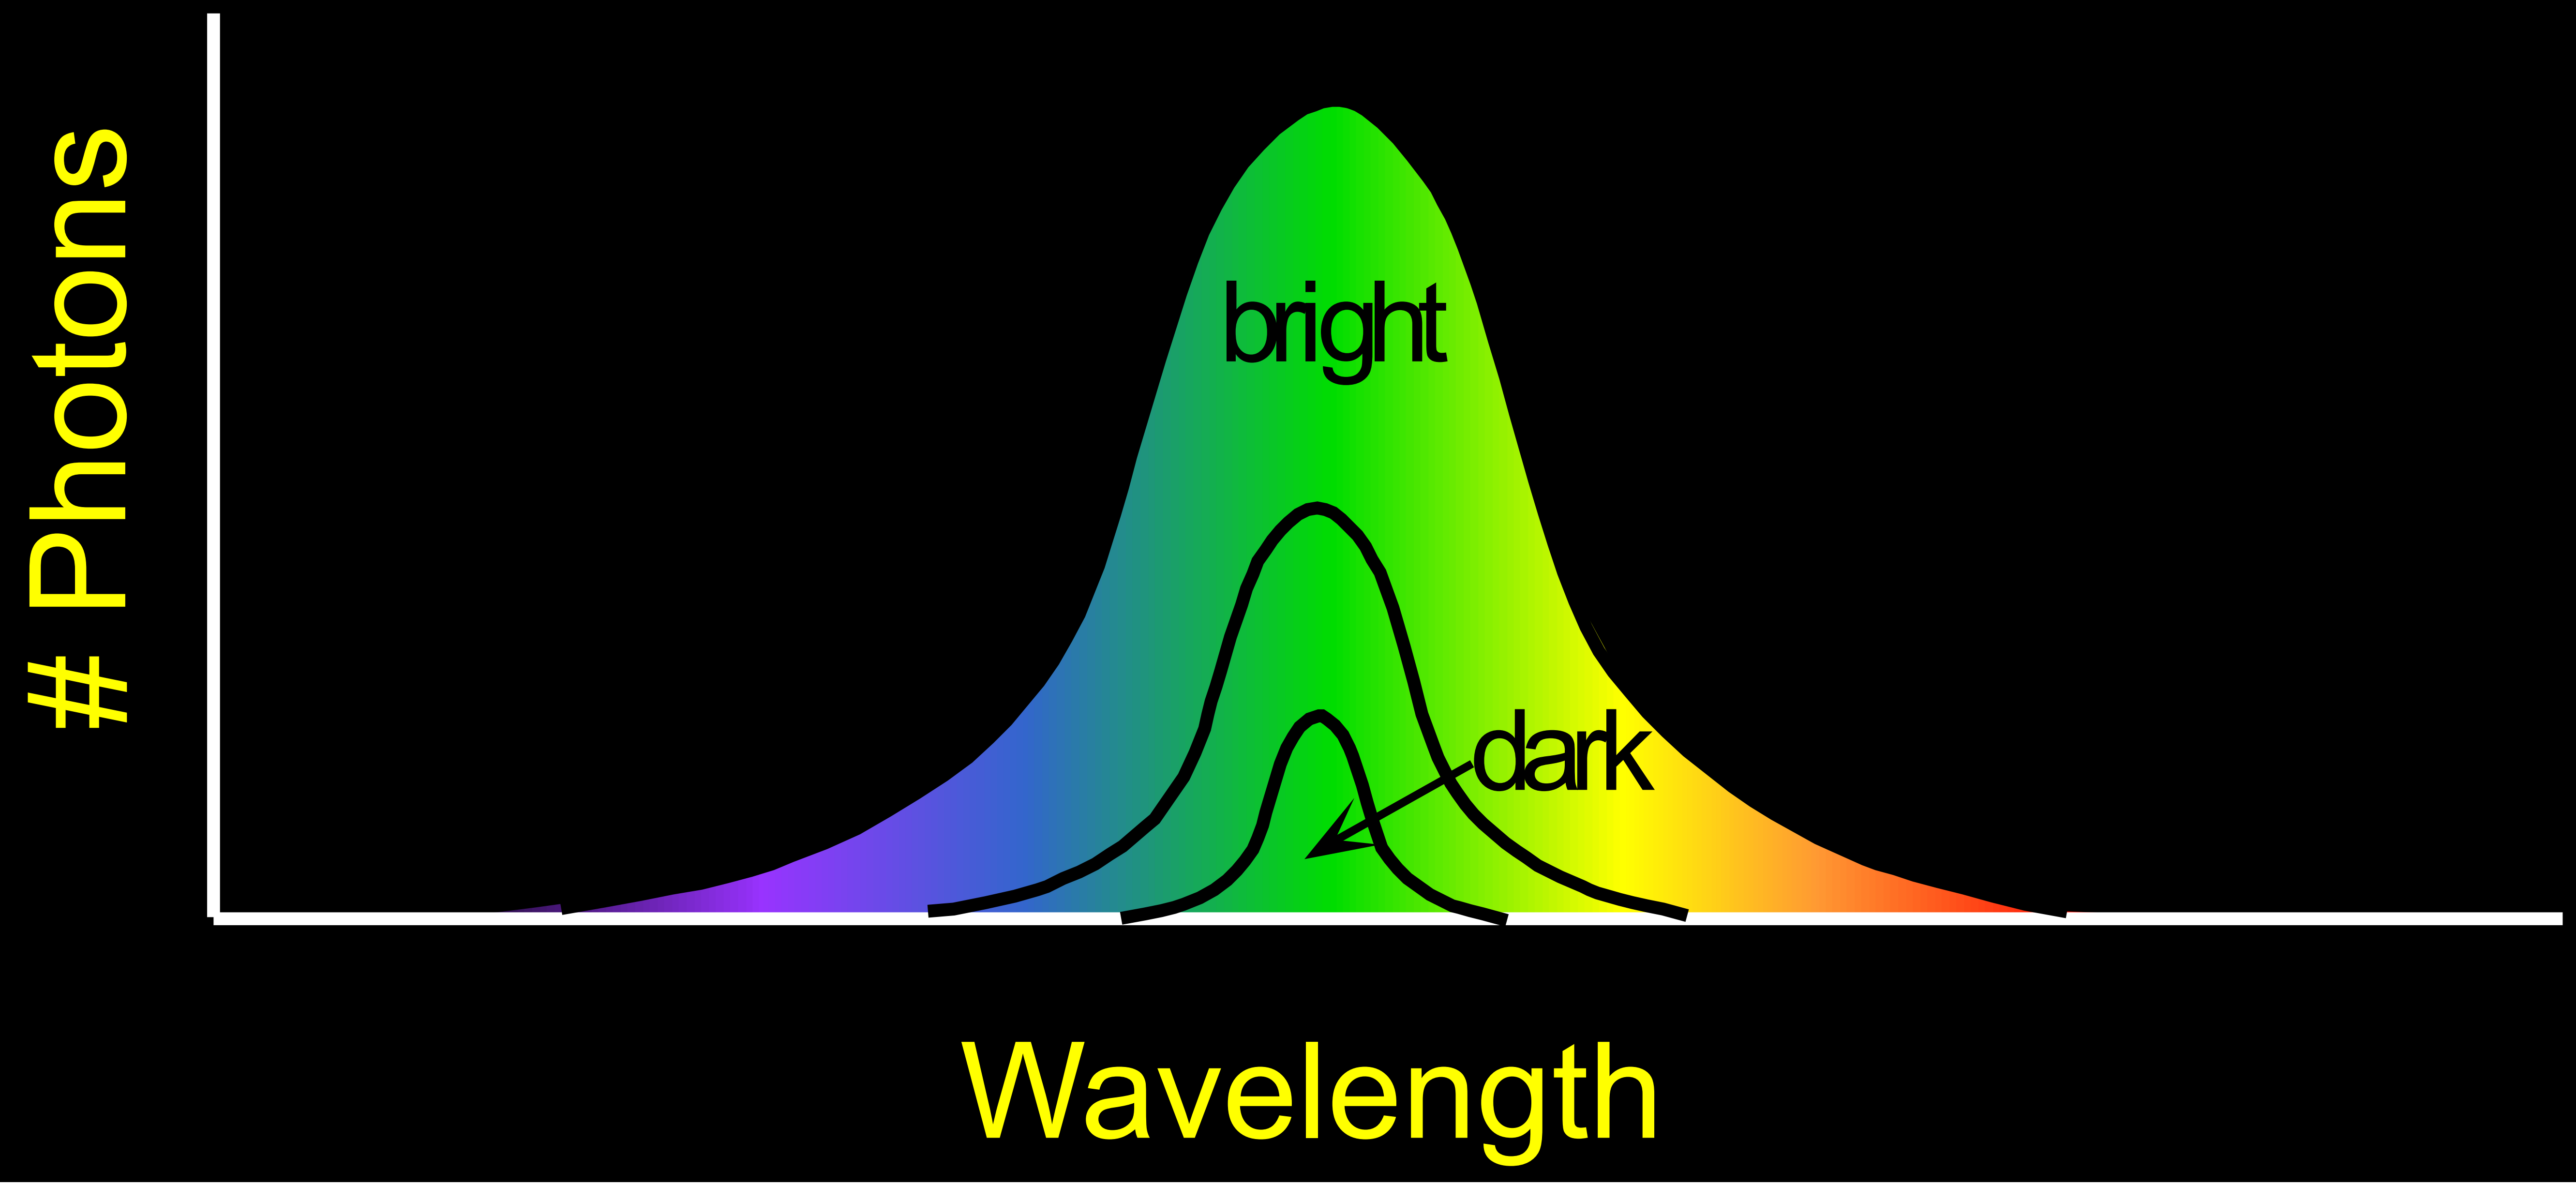

A strange observation…

Some different wavelength distributions “look” the same

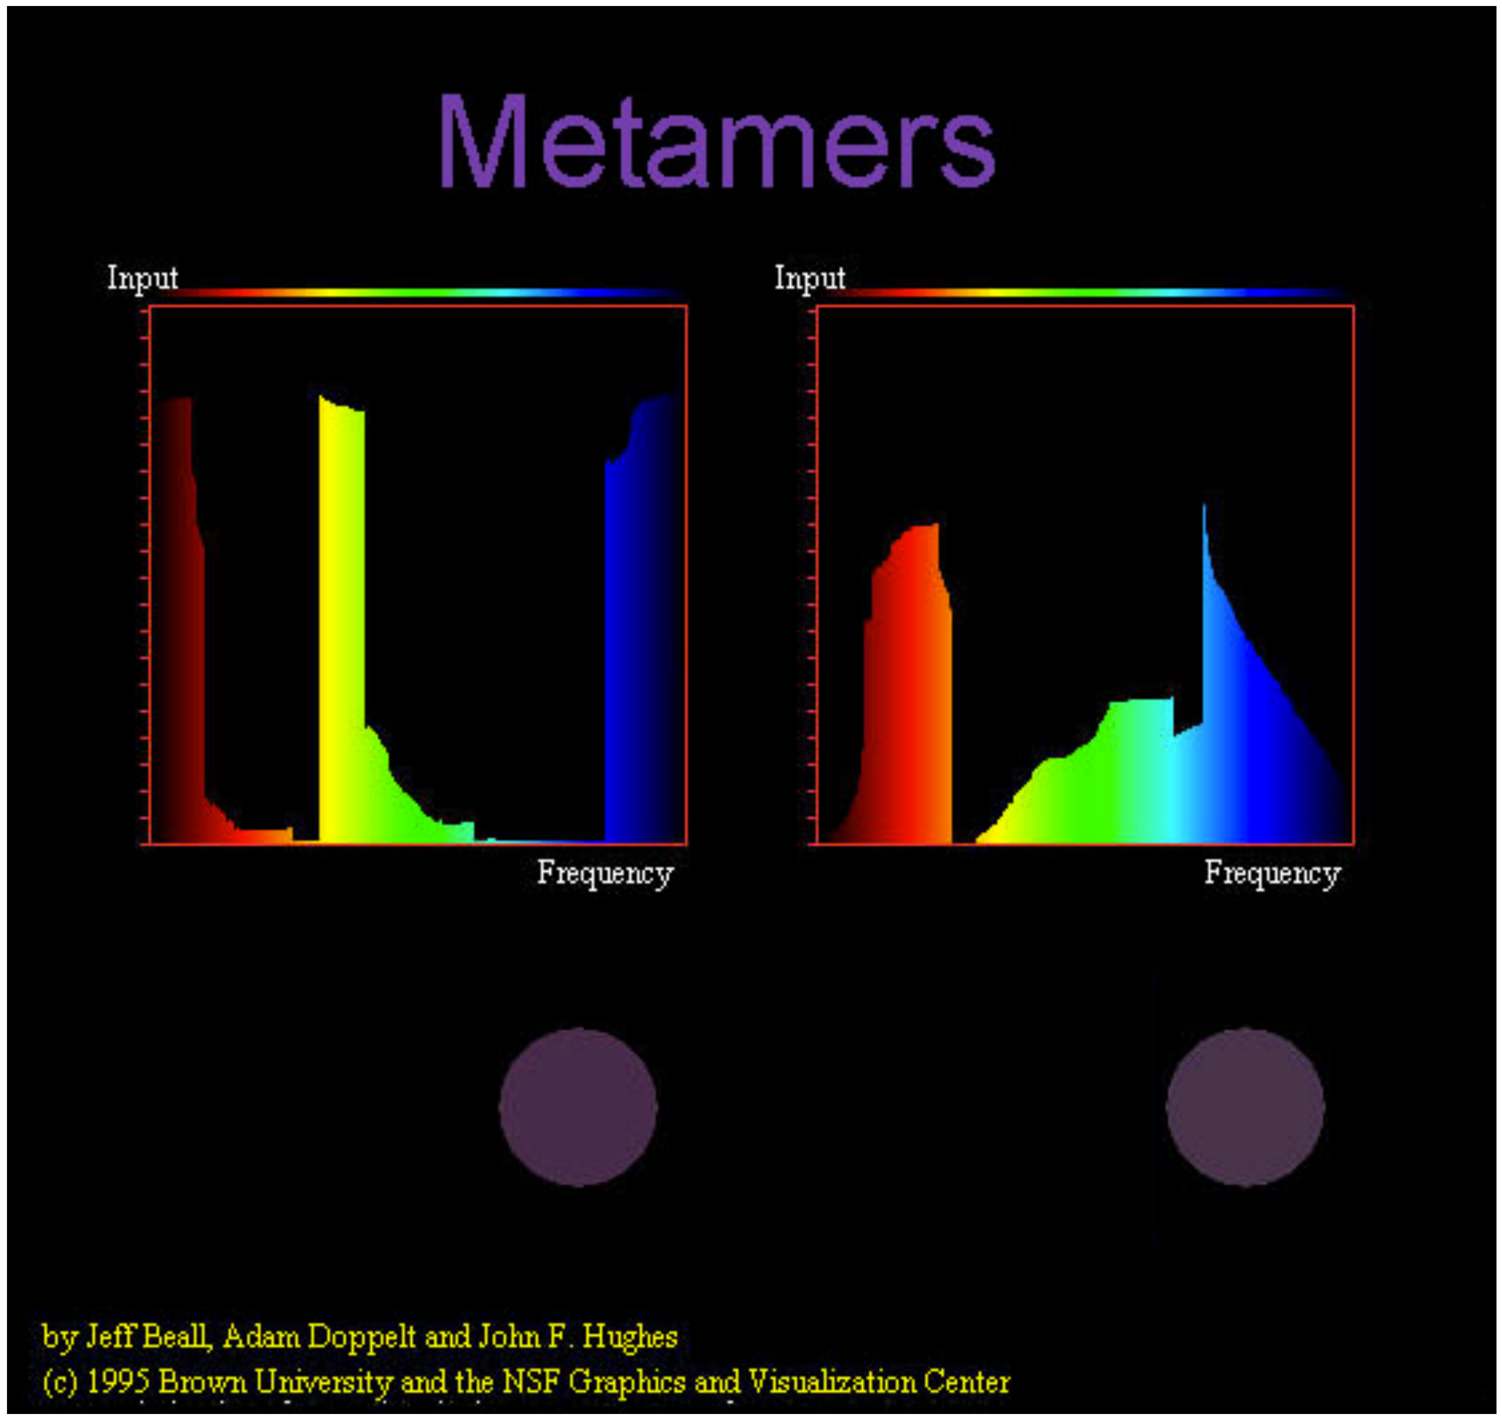

A closer look: human sensing

The Eye

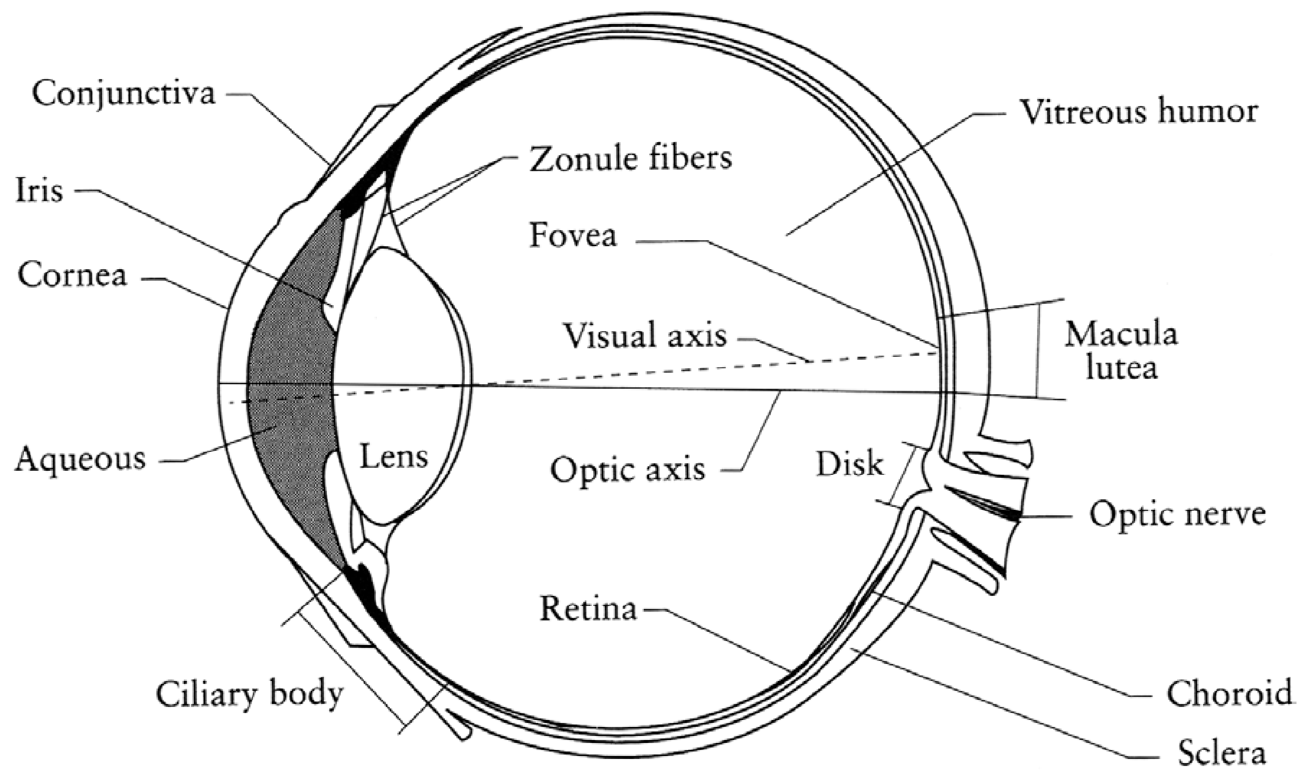

The human eye is a camera!

- **Iris** - colored annulus with radial muscles
- **Pupil** - the hole (aperture) whose size is controlled by the iris
- What's the "film"?
  - photoreceptor cells (rods and cones) in the **retina**

Slide by Steve Seitz

Human photoreceptors

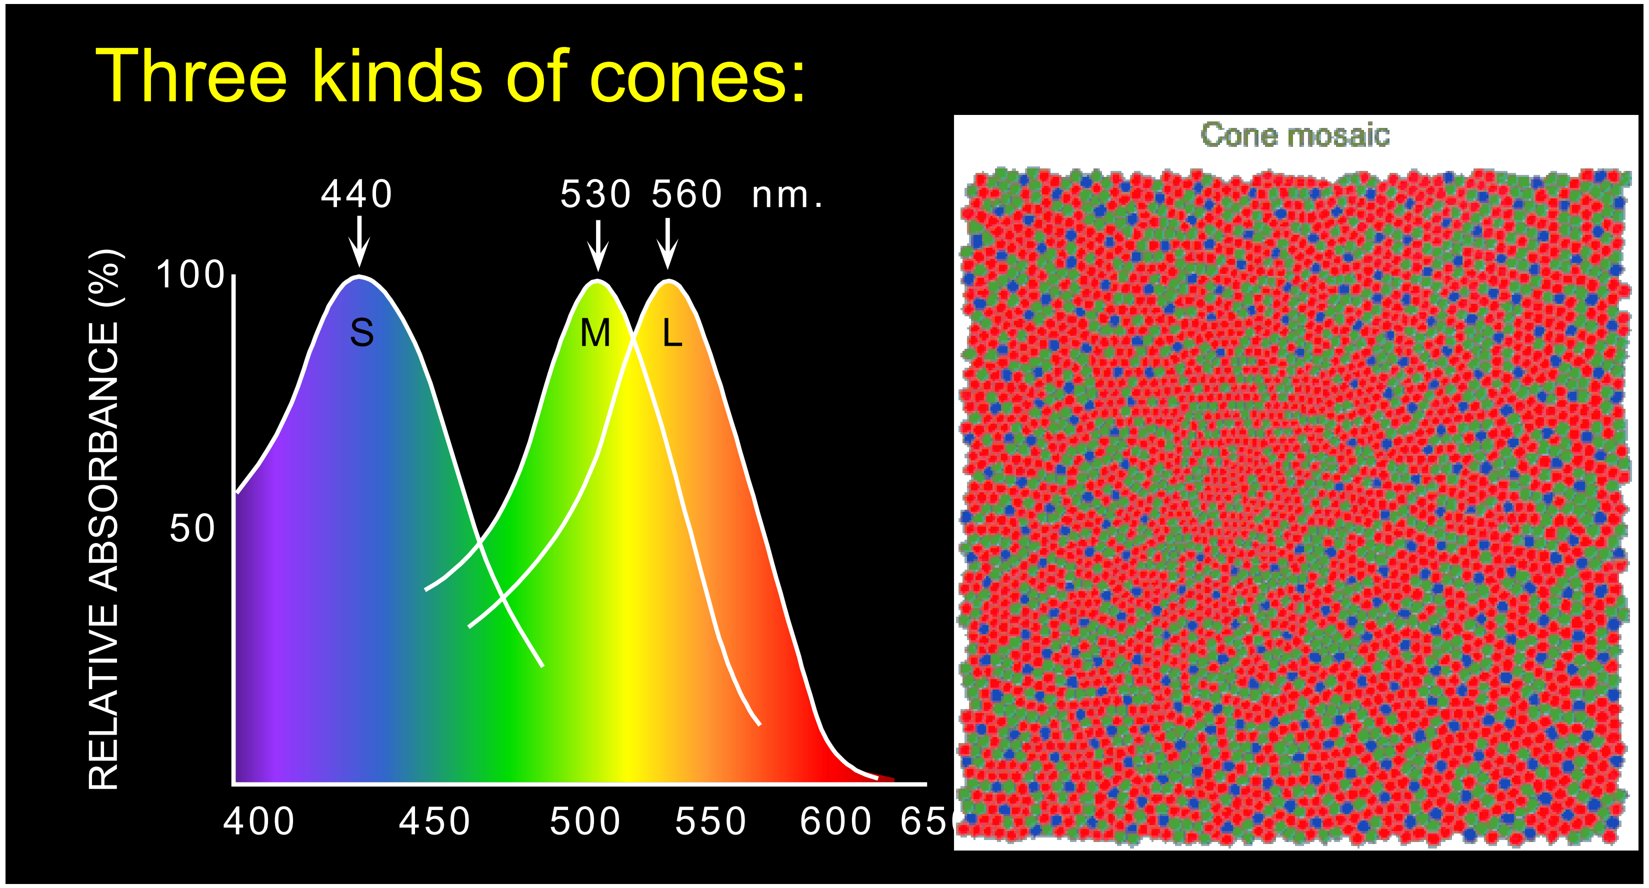

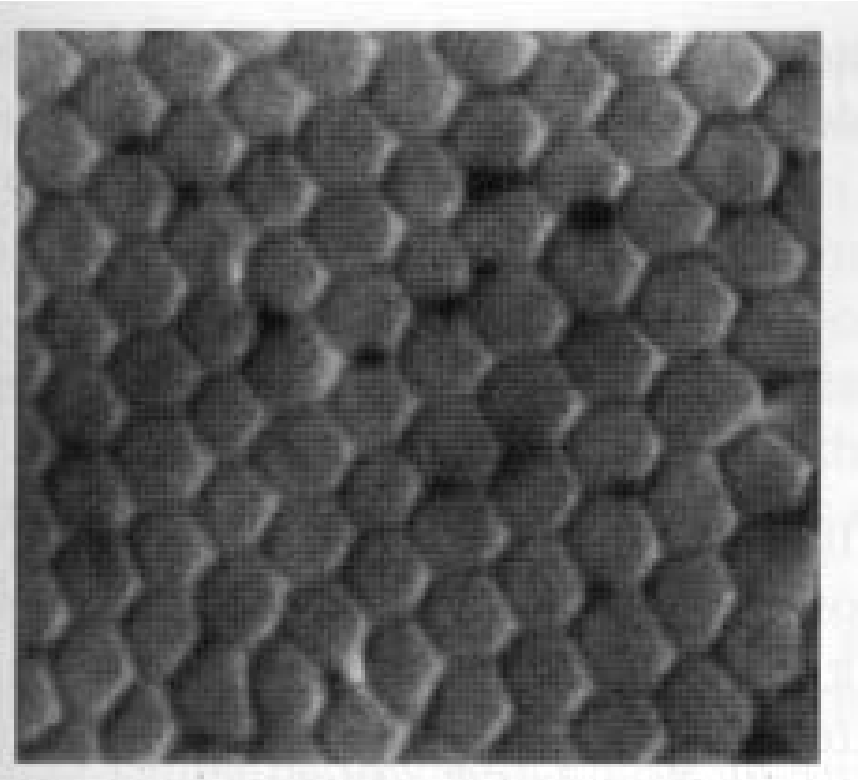

"blue" "green" "red"

Metamers revisted

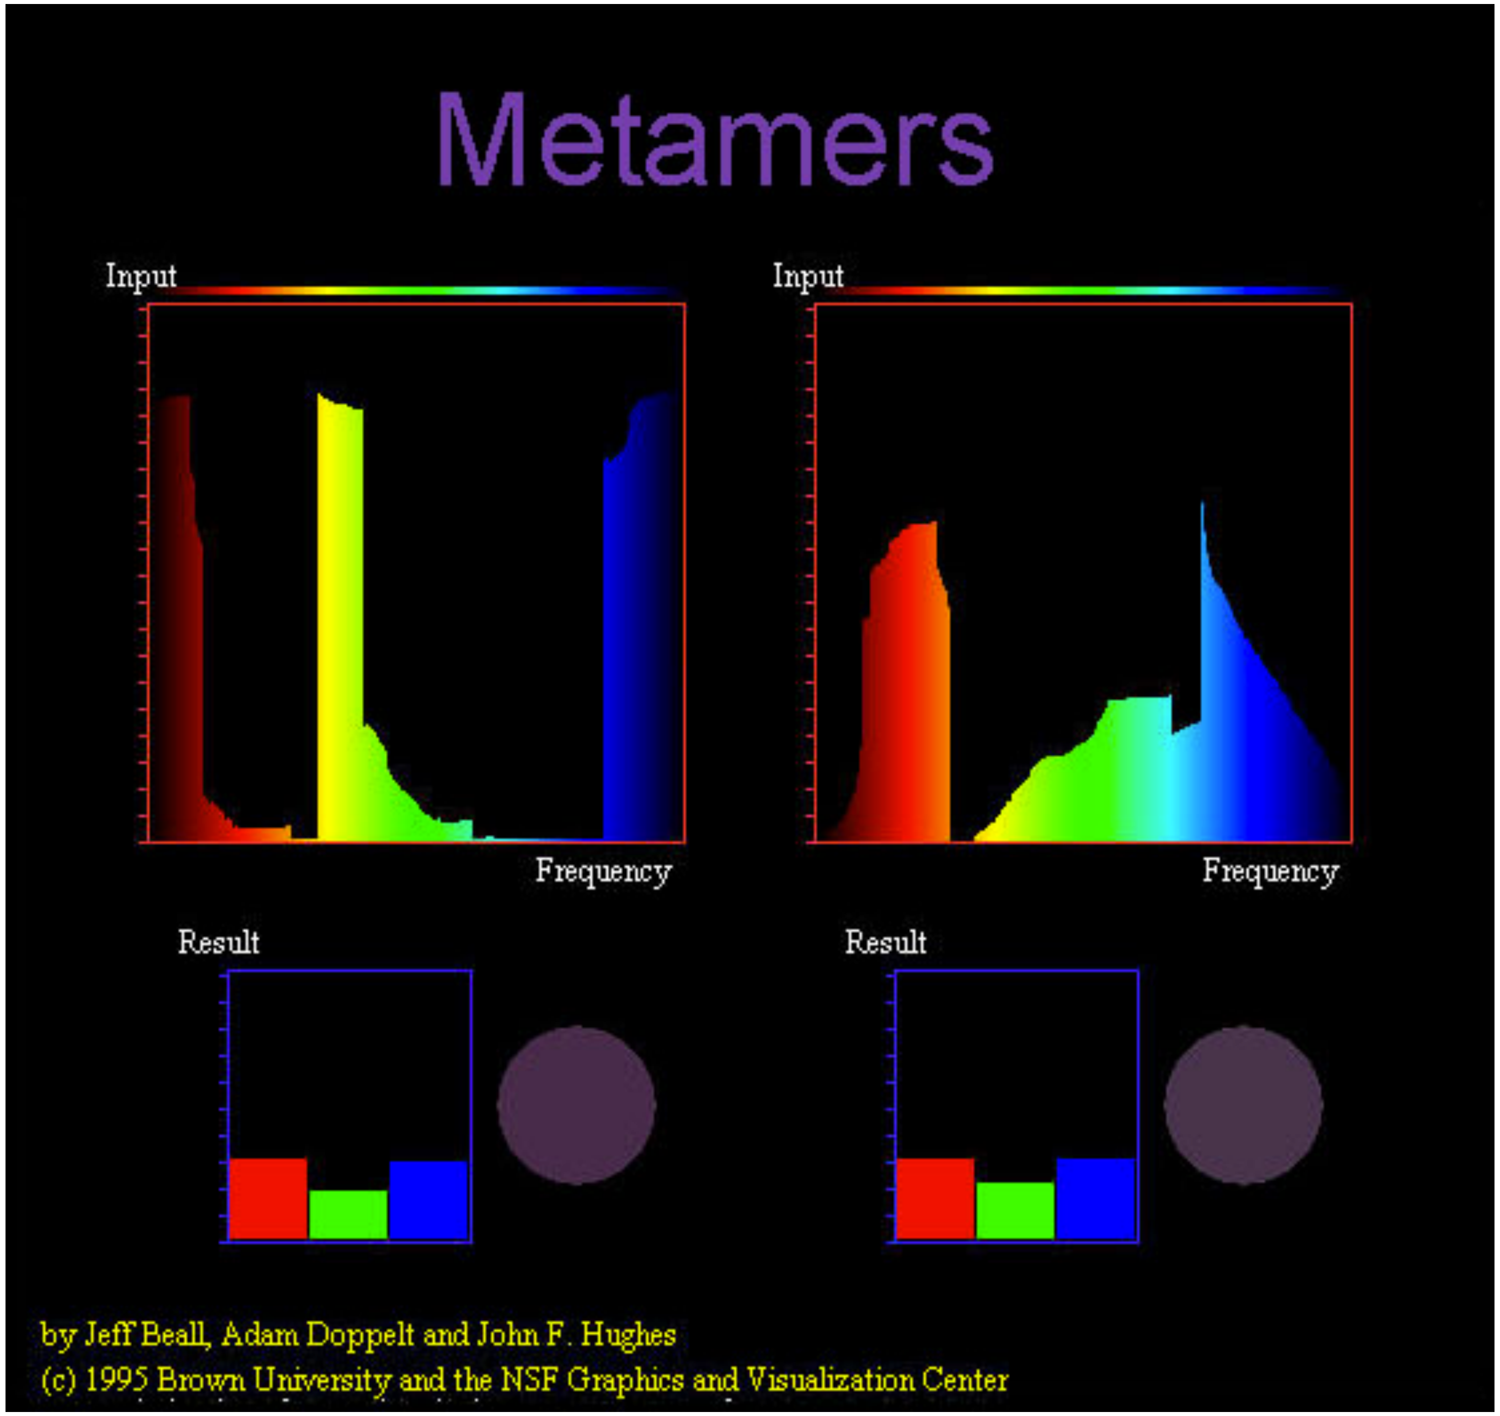

The mathematics of color

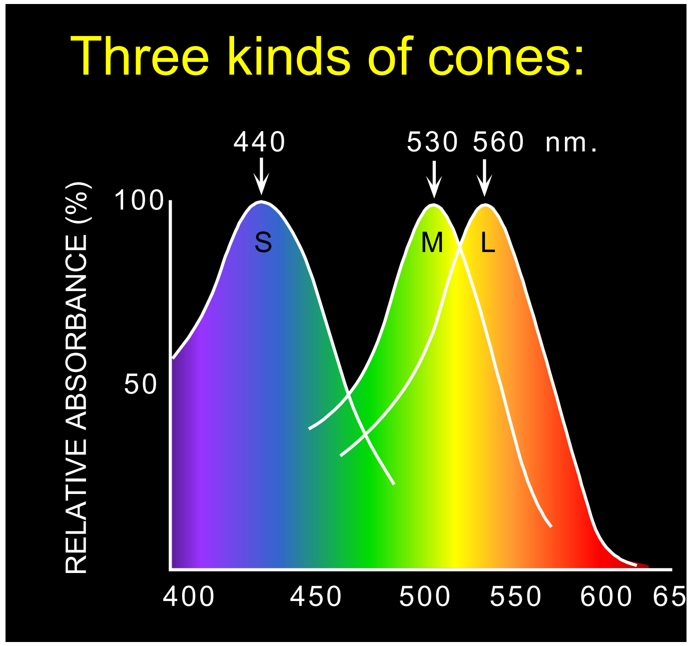

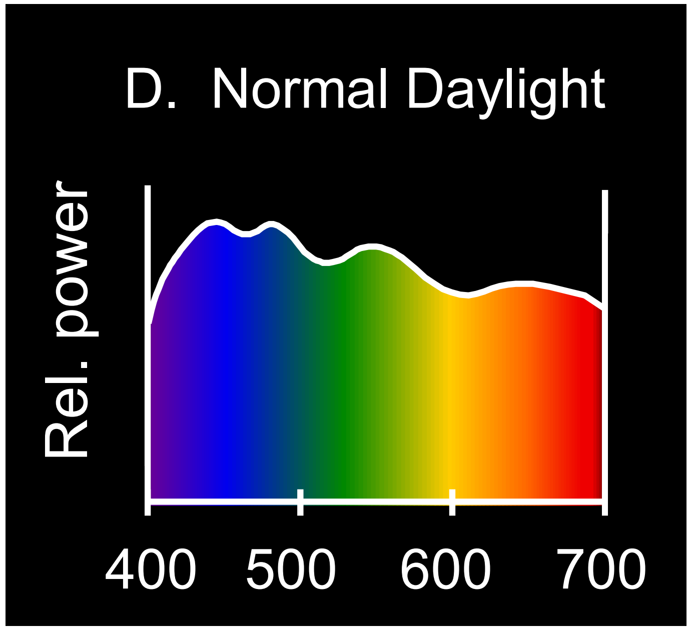

$$R_i(\lambda), \quad i = \{1, 2, 3\}$$

Photo-receptors

$$I(\lambda)$$

Incoming light

$$v_i = \int_0^\infty I(\lambda) R_i(\lambda) d\lambda \qquad i = \{1, 2, 3\}$$

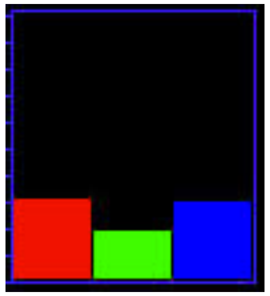

The mathematics of color

$$v_i = \int_0^\infty I(\lambda) R_i(\lambda) d\lambda$$

Means the perception of color is *linear*

$$\begin{aligned} I_1(\lambda) &\rightarrow v_1, v_2, v_3 \\ I_2(\lambda) &\rightarrow w_1, w_2, w_3 \\ \alpha I_1(\lambda) + \beta I_2(\lambda) &\rightarrow \end{aligned}$$

What about non-humans?

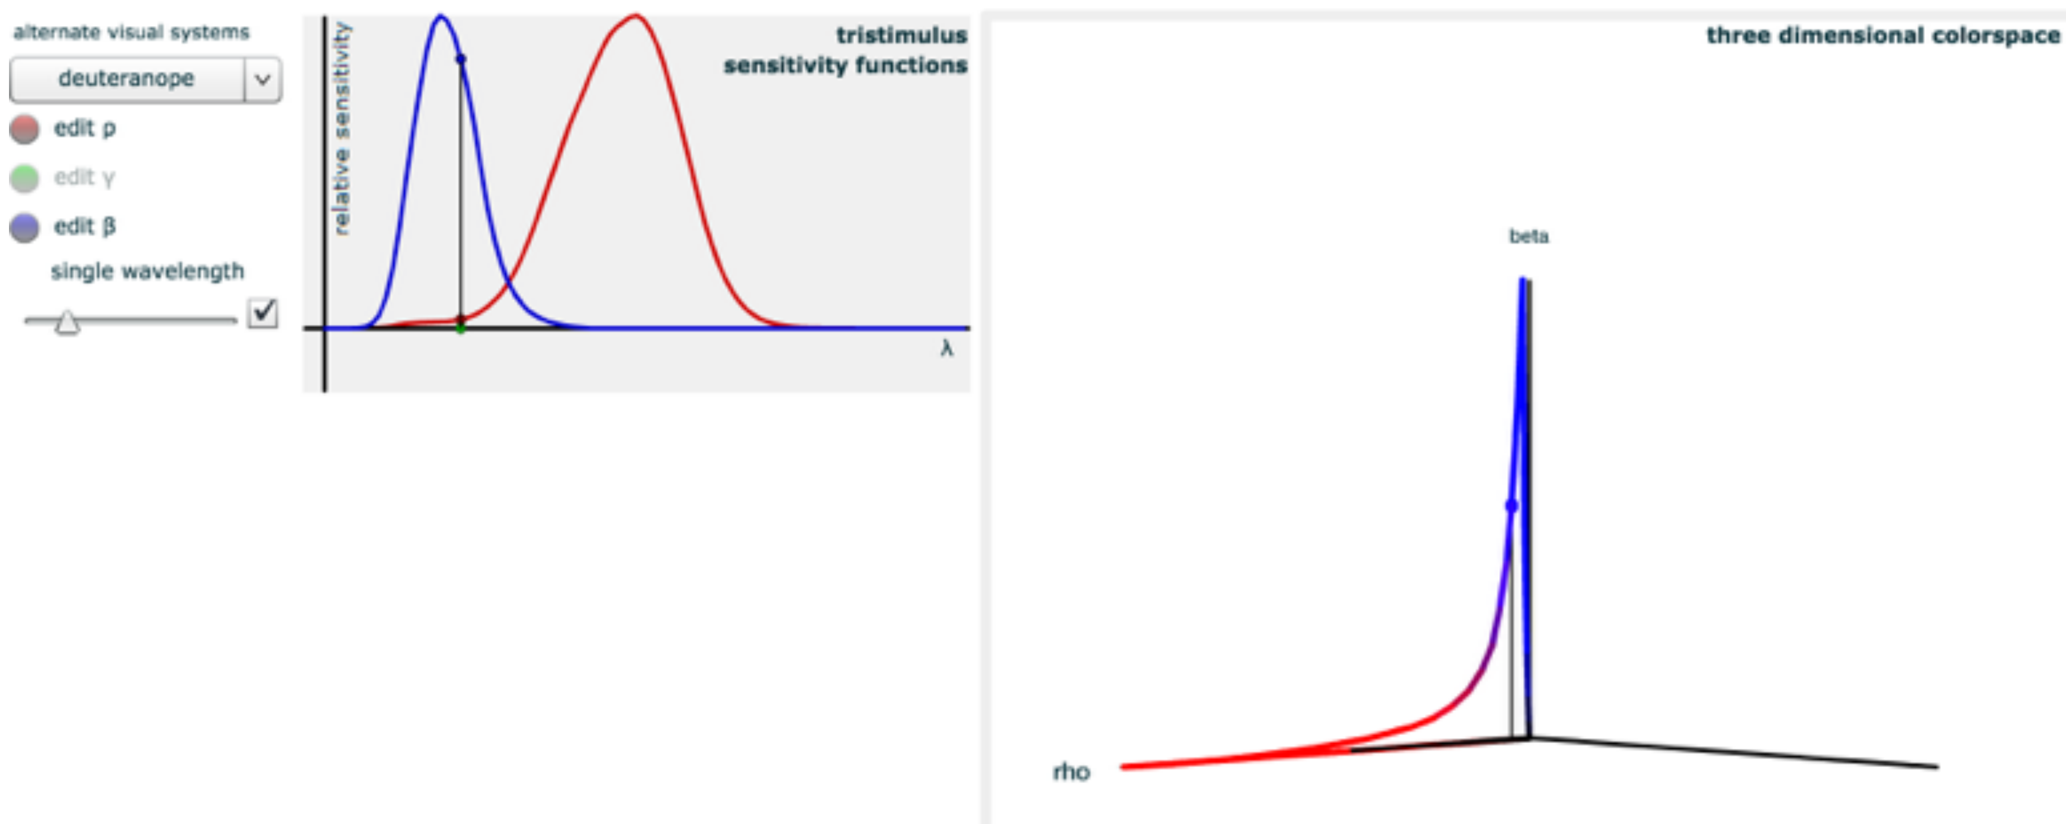

http://graphics.stanford.edu/courses/cs178/applets/locus.html

How do we numerically represent a color?

Given a candidate color, find scale factors for tuning 3 reference primaries so that they visually match

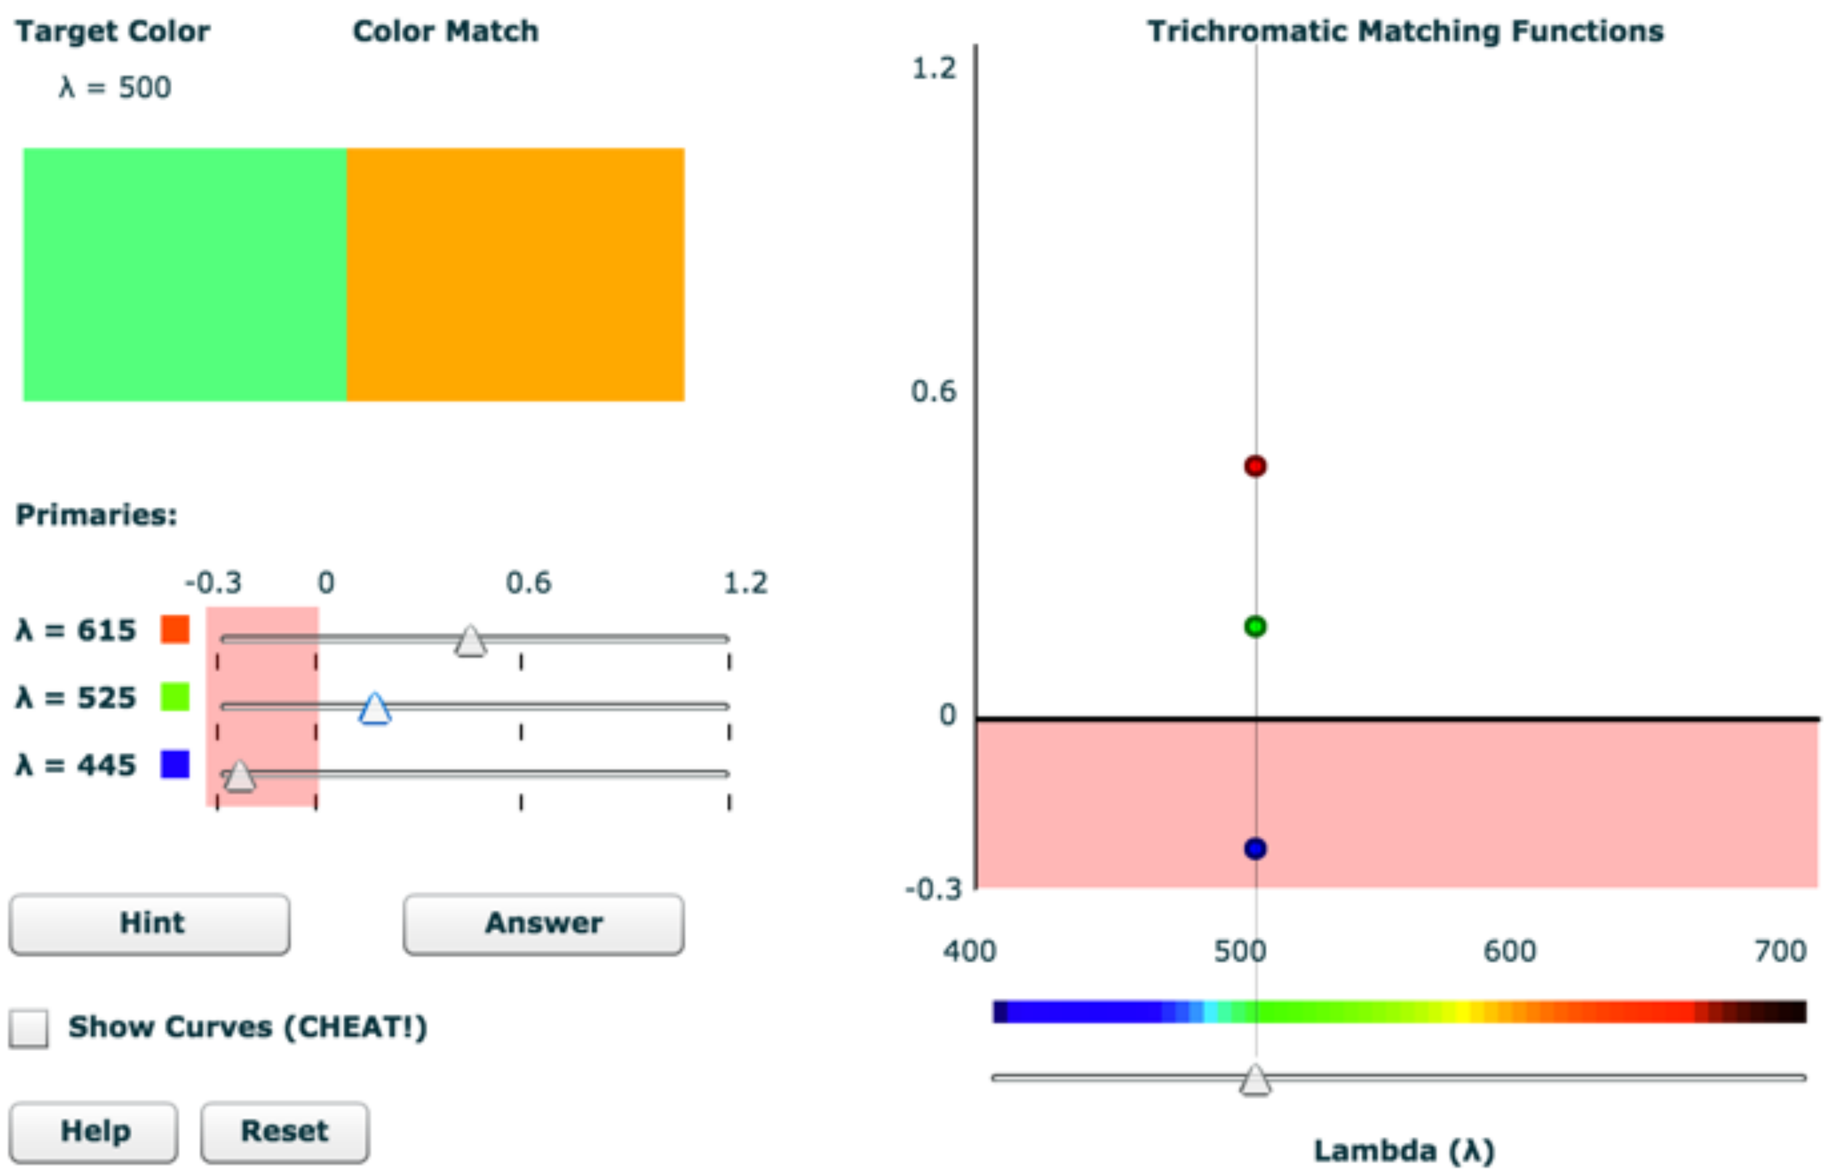

http://graphics.stanford.edu/courses/cs178/applets/colormatching.html

Summary of color matching

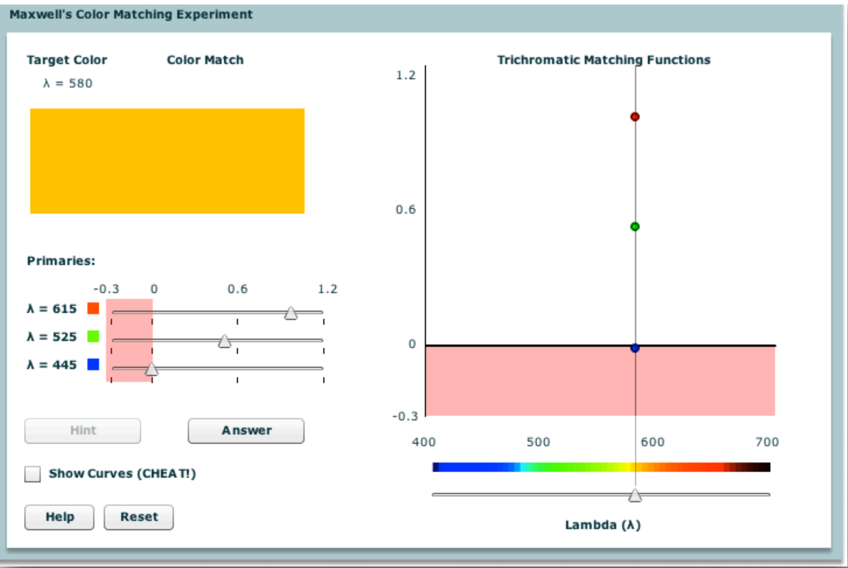

1

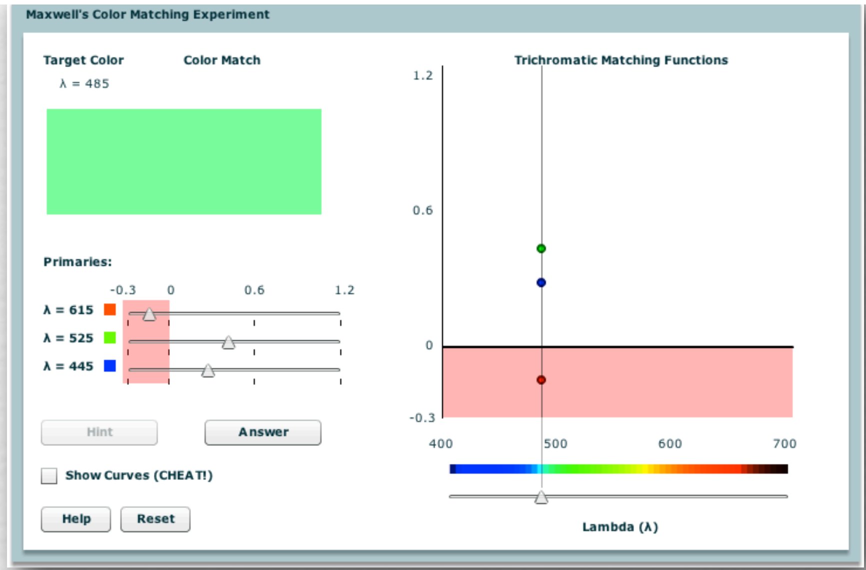

2

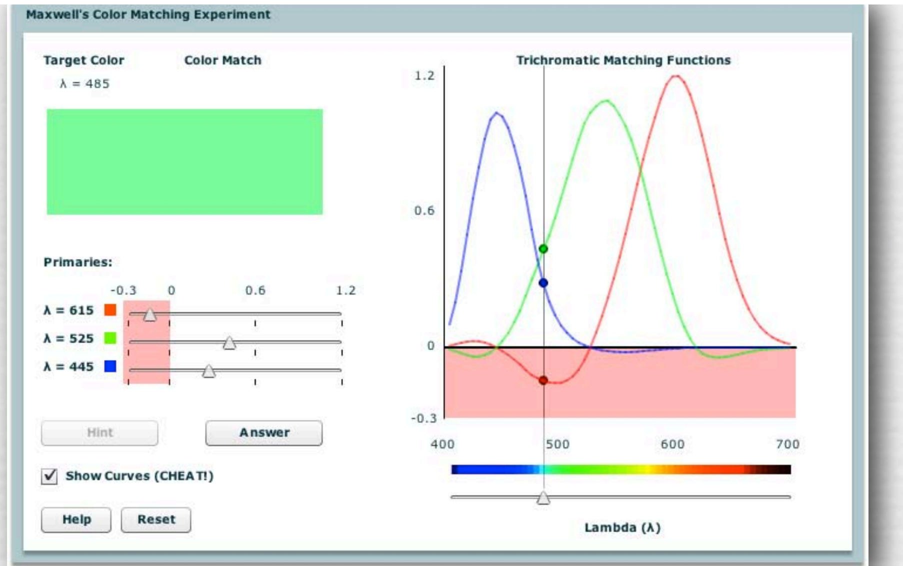

3

1. given a stimulus wavelength, the amount of each primary required to match it is given by three numbers $(\bar{r}, \bar{g}, \bar{b})$
2. some stimuli cannot be matched unless first desaturated by adding some of one primary to it before matching; the amount added is denoted by negative values of $\bar{r}$, $\bar{g}$, or $\bar{b}$
3. the sequence of $(\bar{r}, \bar{g}, \bar{b})$ values, some negative, required to match the locus of spectral colors across all λ, form the *trichromatic matching functions* $\bar{r}(\lambda)$, $\bar{g}(\lambda)$, and $\bar{b}(\lambda)$ for a particular set of 3 primaries

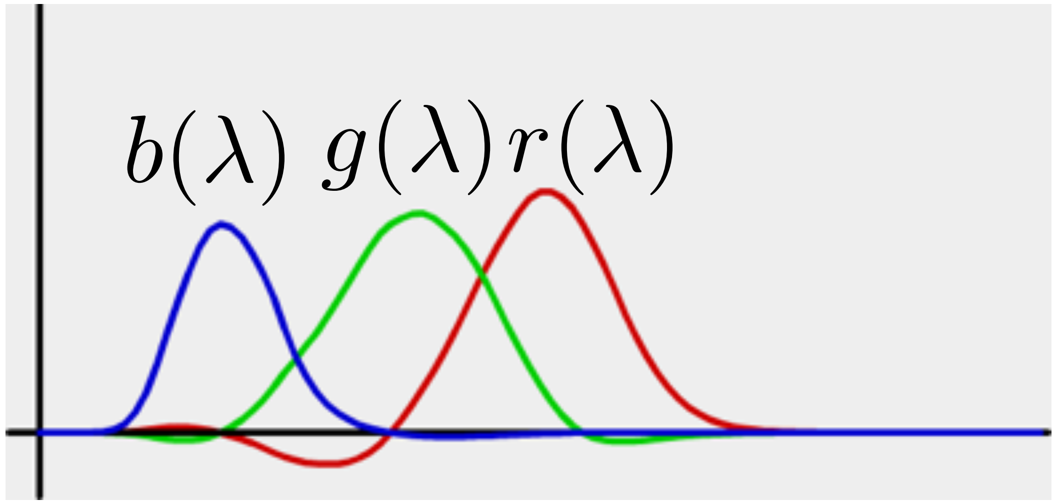

RGB color matching function

Given color-matching functions, how do I compute $\overline{r}$,$\overline{g}$,$\overline{b}$ values for a new light spectra $I(\lambda)$?

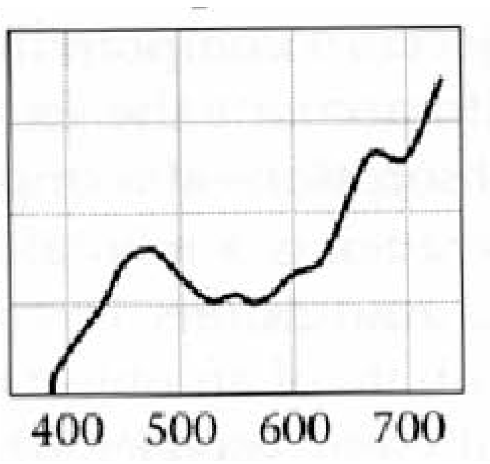

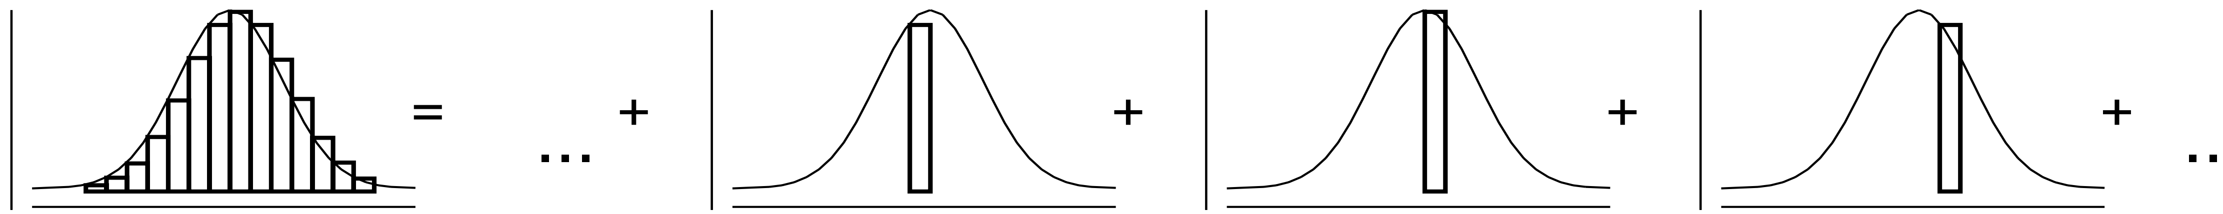

Write any spectra as a linear combination of single-wavelength signals

$$I(\lambda) \approx I(0)\delta(\lambda) + I(1)\delta(\lambda - 1) + I(2)\delta(\lambda - 2) + \ldots$$
$$= \int_0^\infty I(u)\delta(\lambda - u)du$$

To predict response to a particular (e.g., red) photoreceptor,
we need only know its sensitivty to single-wavelength signals

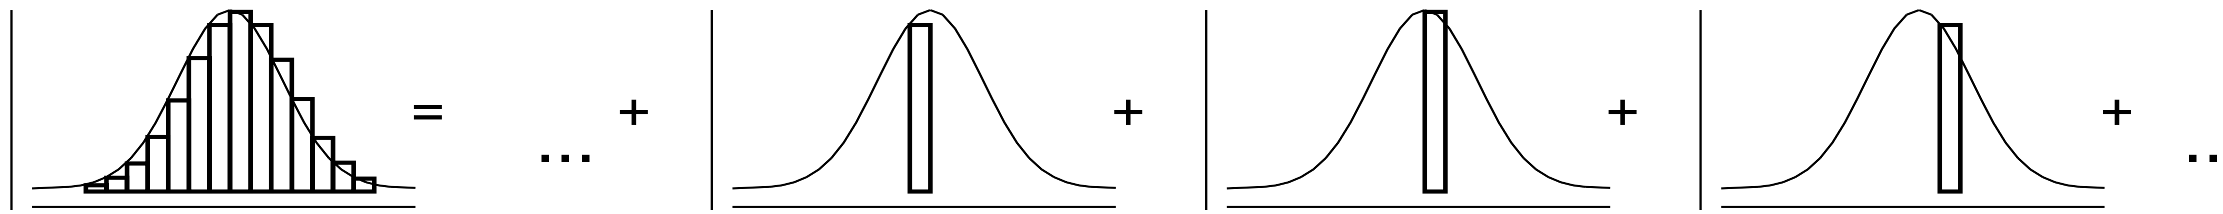

Write any spectra as a linear combination of single-wavelength signals

$$I(\lambda) \approx I(0)\delta(\lambda) + I(1)\delta(\lambda - 1) + I(2)\delta(\lambda - 2) + \ldots$$
$$= \int_0^\infty I(u)\delta(\lambda - u)du$$

$$v_i[I(\lambda)] \approx I(0)R_i(0) + I(1)R_i(1) + \ldots$$
$$= \int_0^\infty I(\lambda)R_i(\lambda)d\lambda$$

photoreceptor sensitivity function

$$i = \{1, 2, 3\}$$

Spectral locus

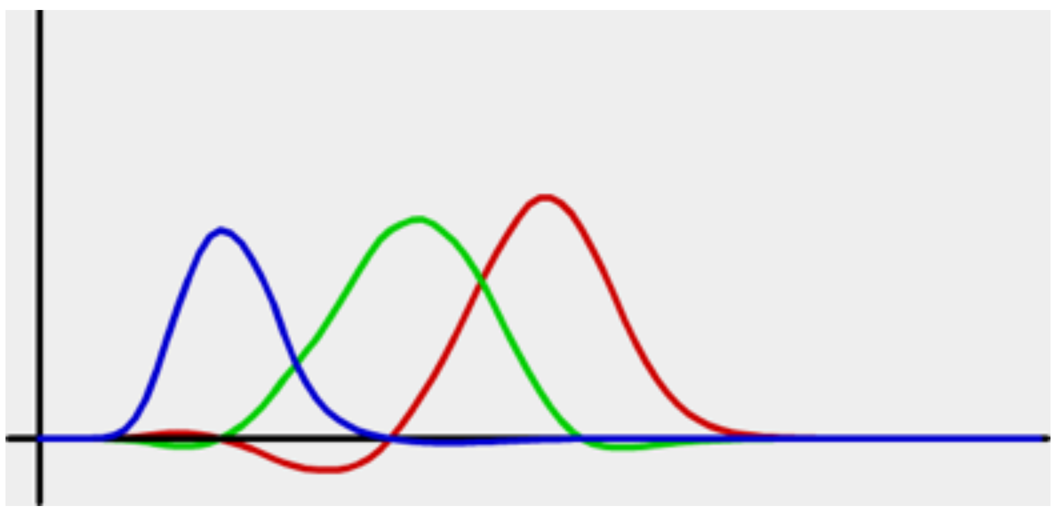

RGB color matching function

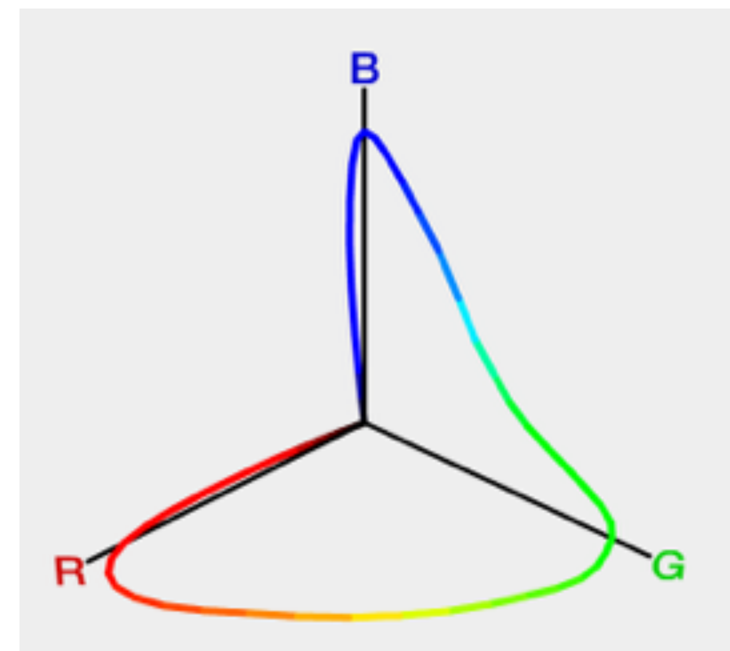

RGB values

Range of RGB values encountered for single-wavelength stimulus

RGB cube

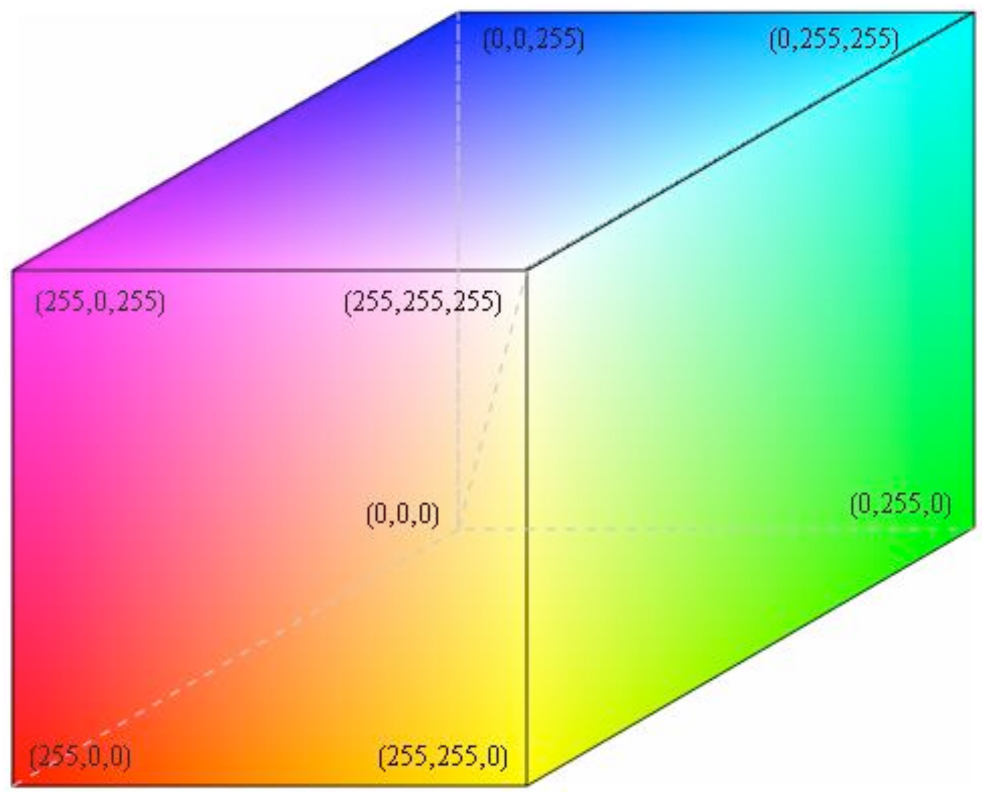

Use red, green, and blue primaries
Define RGB = (255,255,255) to be a reference white and (0,0,0) to be reference black
Cannot represent colors with requiring negative R,G,B values

Alternative parameterization: HSV

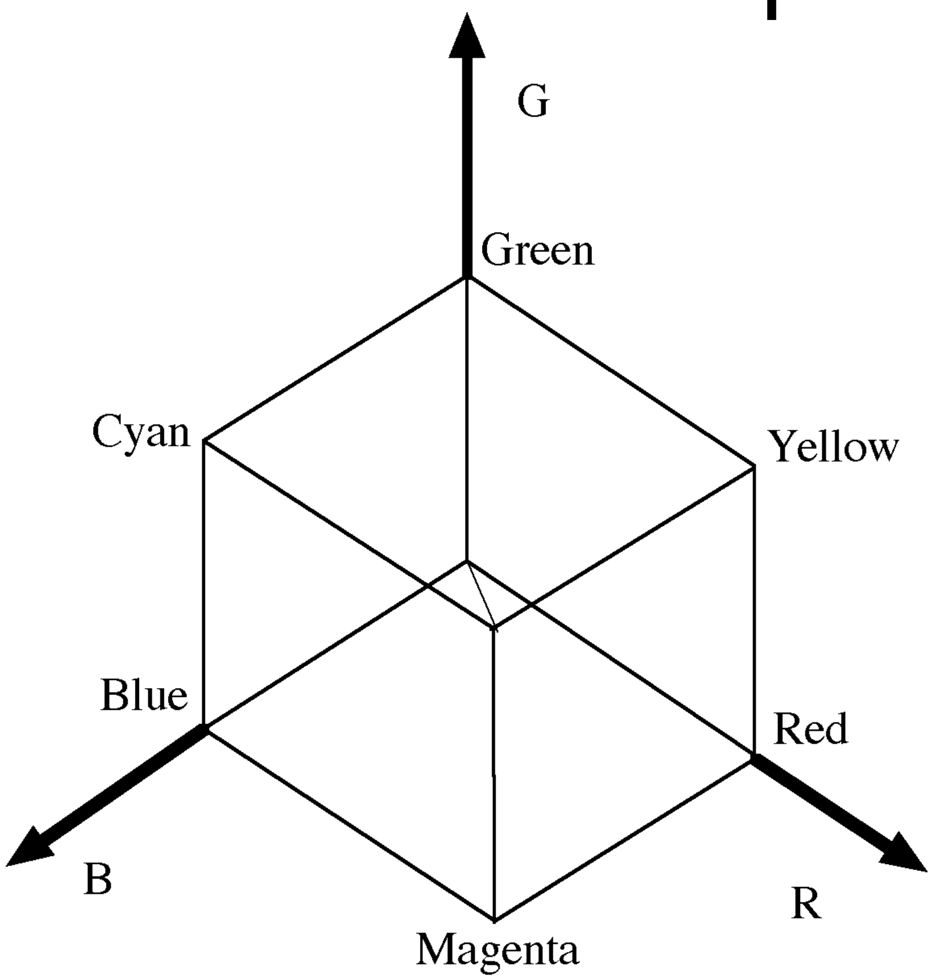

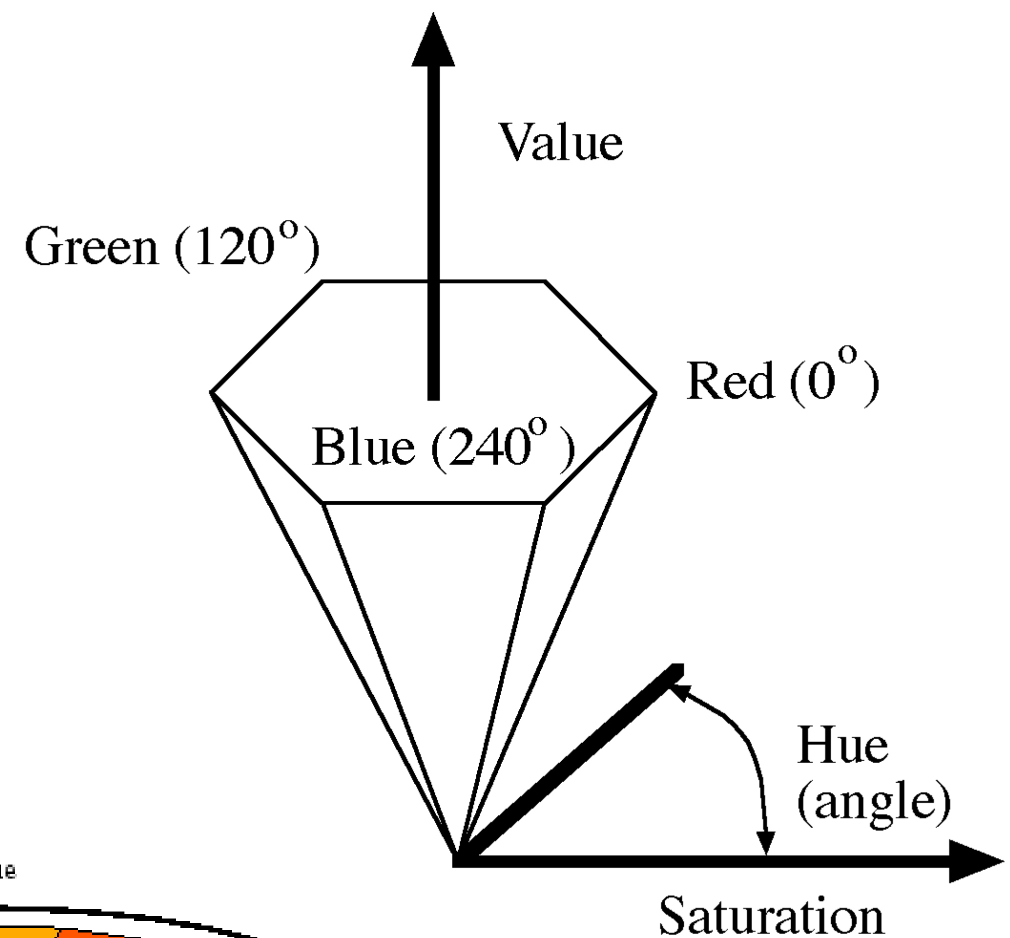

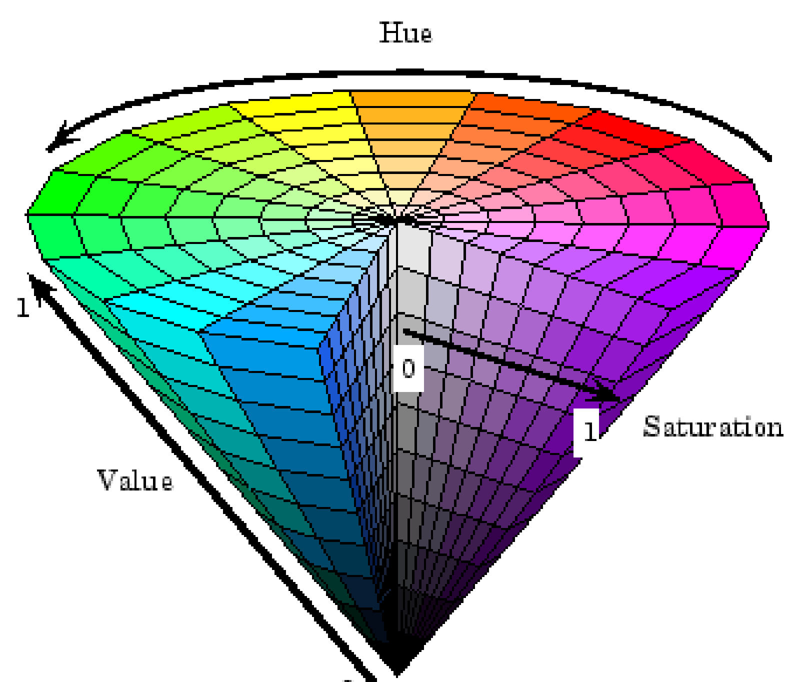

Separate out luminance (brightness) versus chrominance (color)

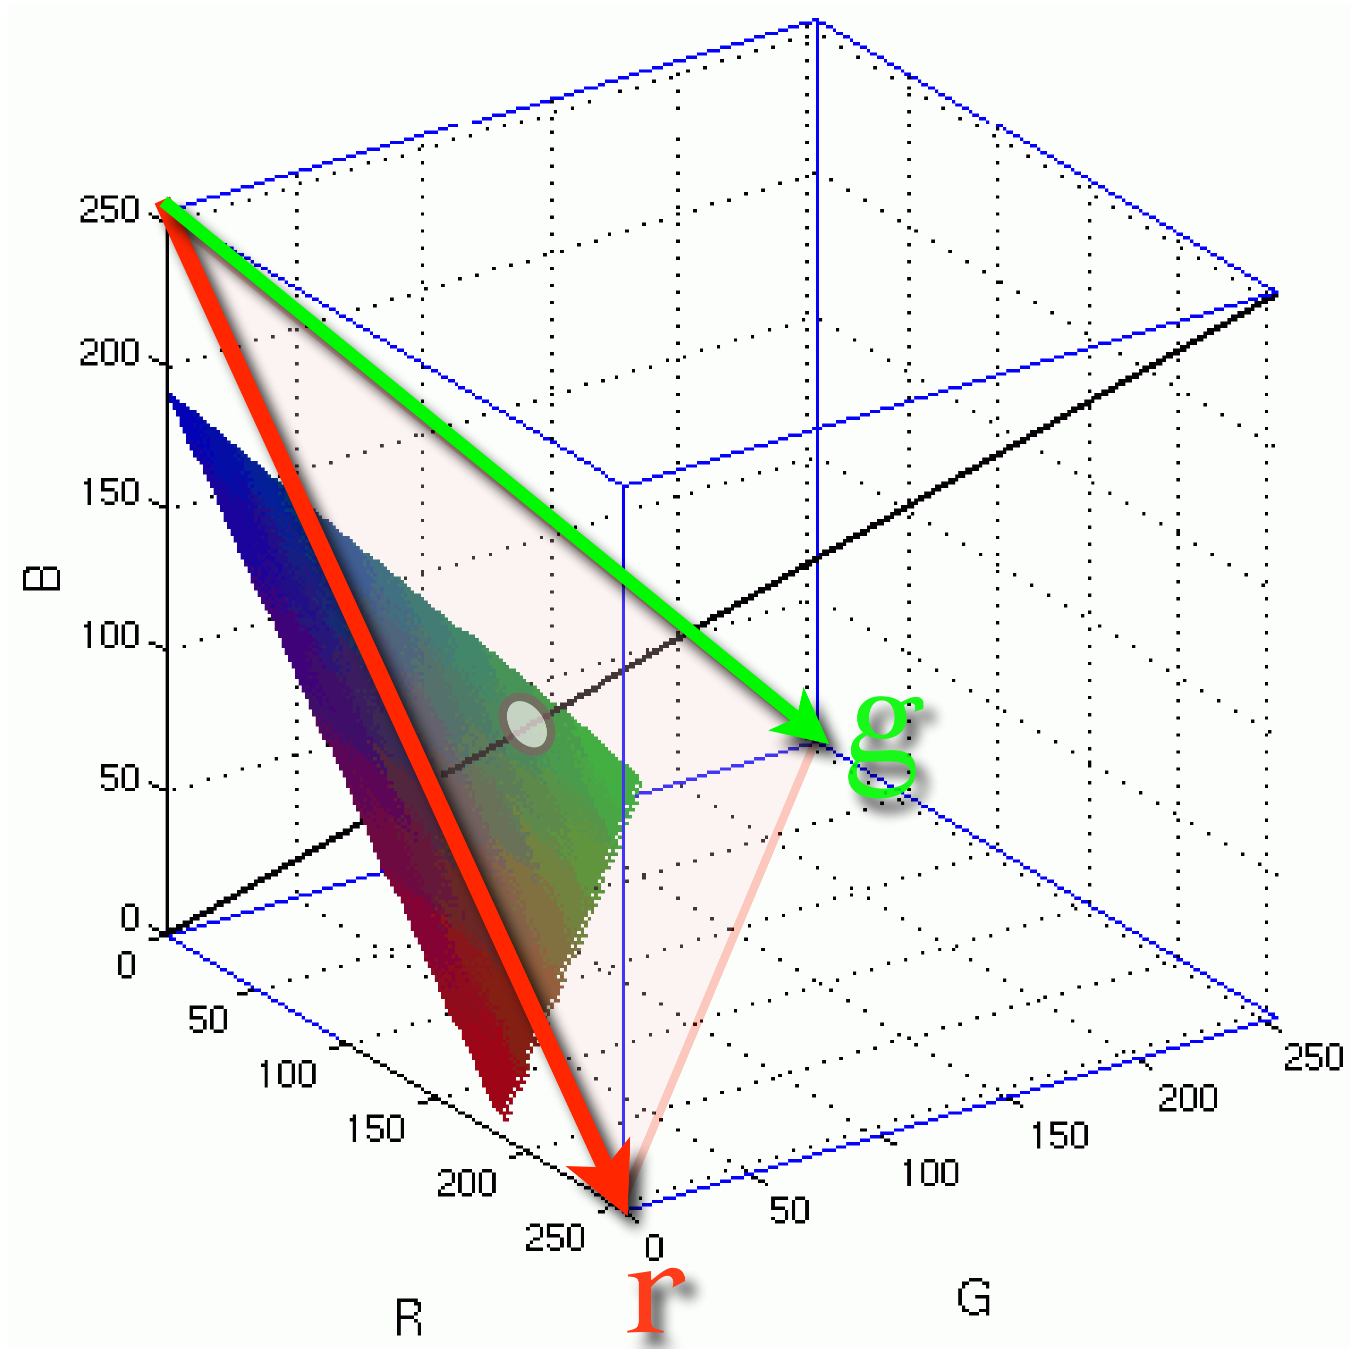

$$r = \frac{R}{R+G+B}$$

$$g = \frac{G}{R+G+B}$$

Points in the RGB cube having the same chromaticity but varying brightness lie along lines emanating from black (0,0,0)

Chrominance view of spectral locus

$$r = \frac{R}{R+G+B}$$

$$g = \frac{G}{R+G+B}$$

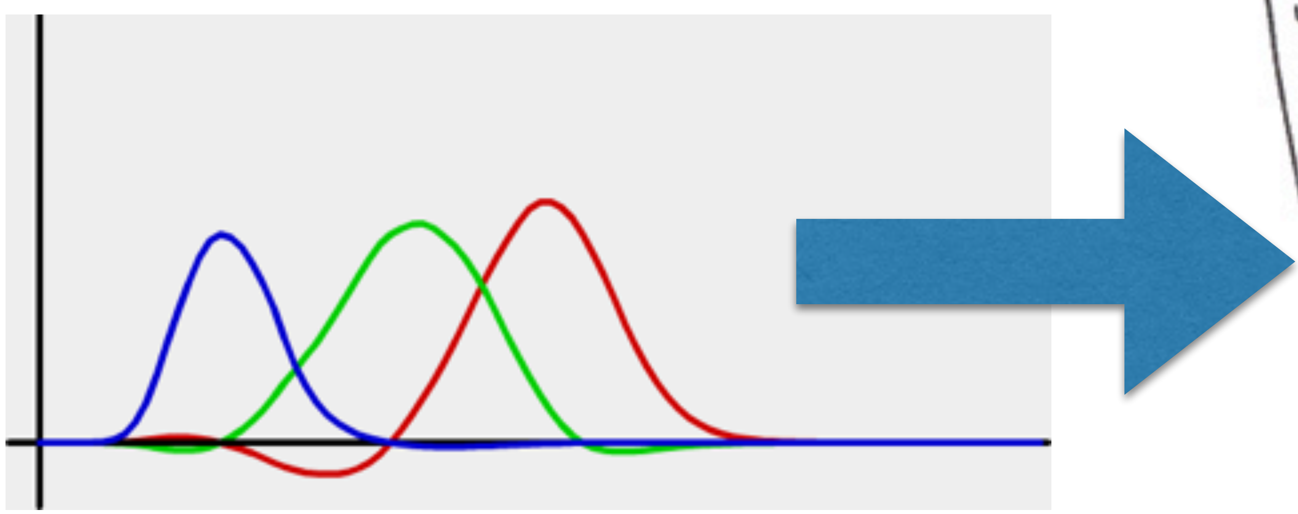

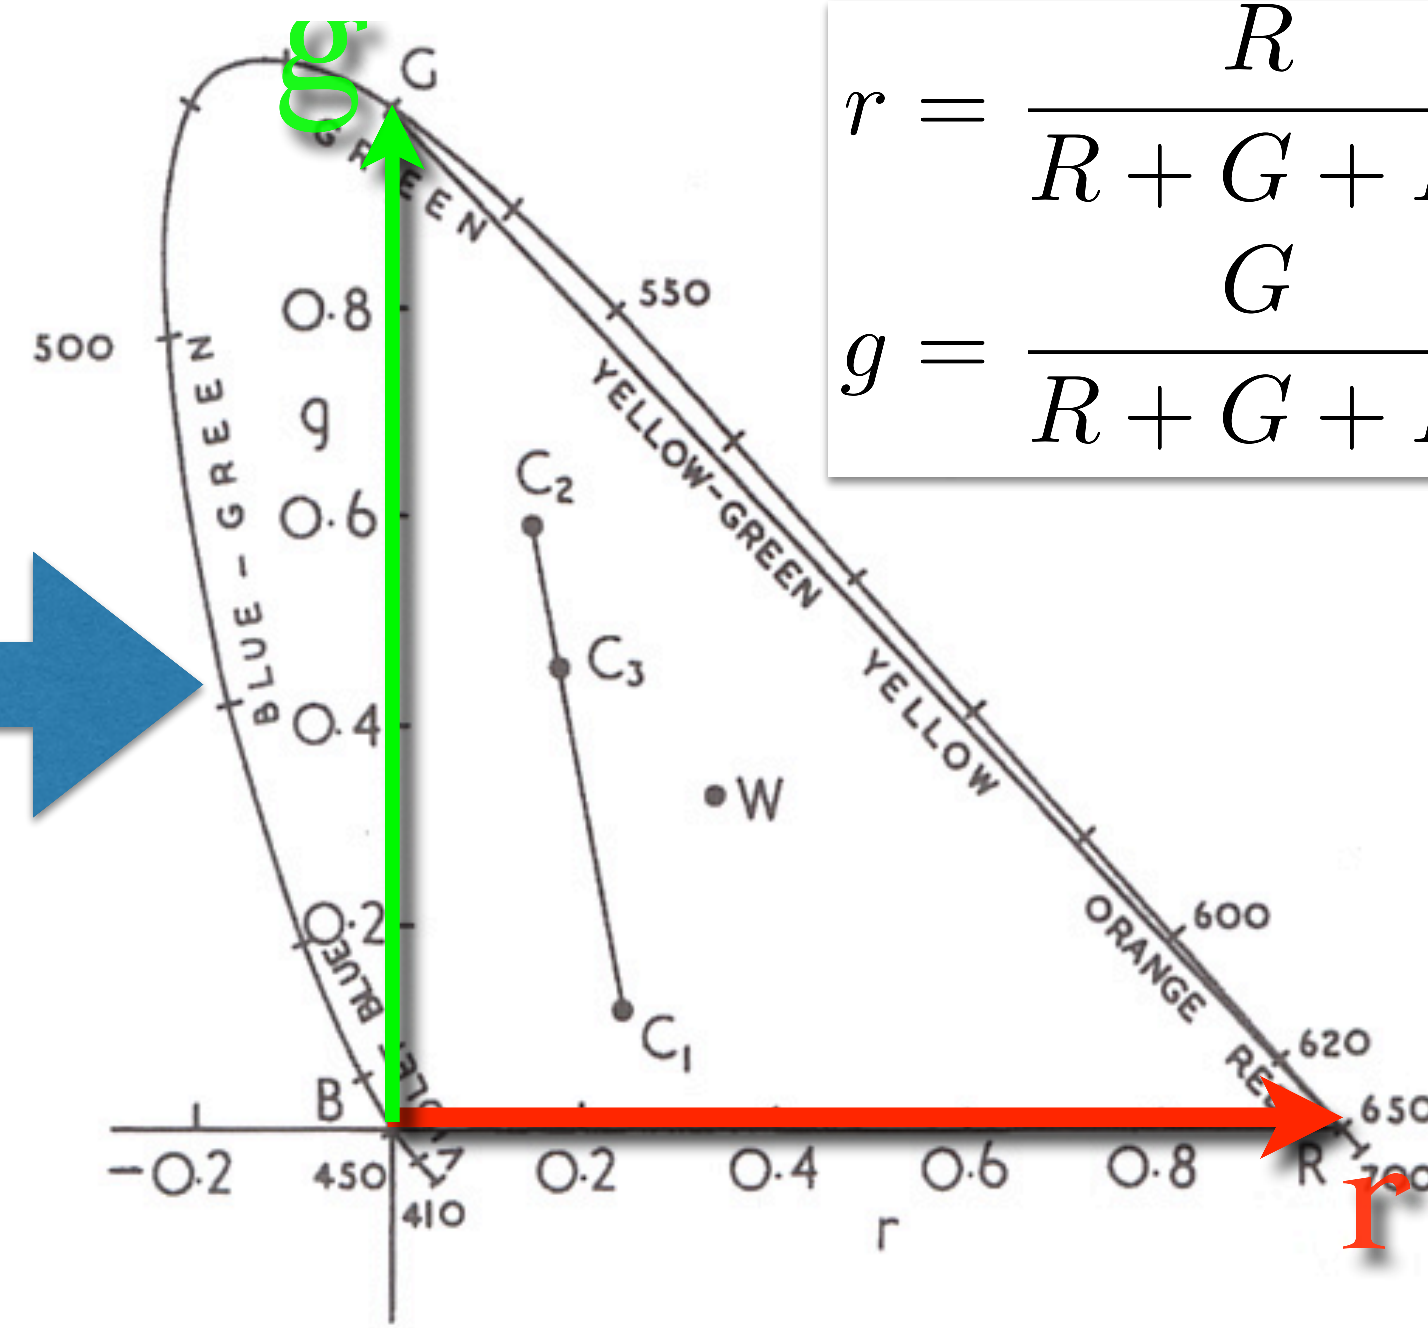

neutral (black to white) = (1/3,1/3)

r>0,g>0 does not enclose spectral locus

http://graphics.stanford.edu/courses/cs178/applets/threedgamut.html

CIE XYZ color space

CIE = “International Commission on Illumination”

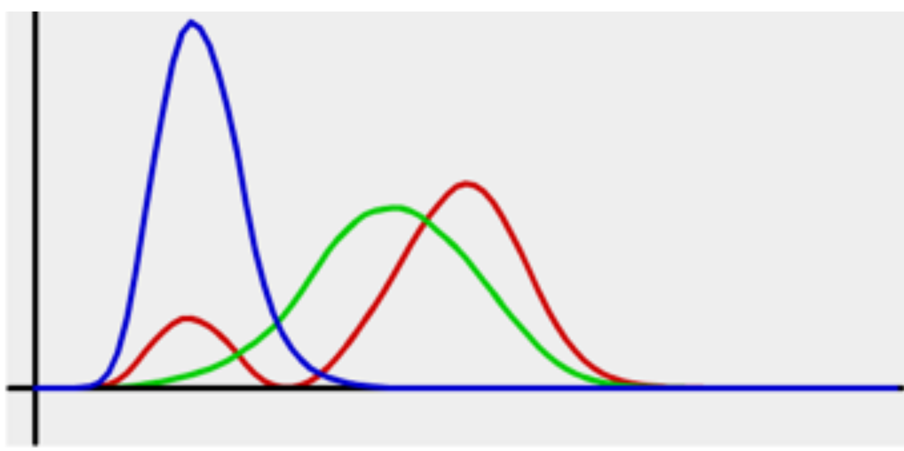

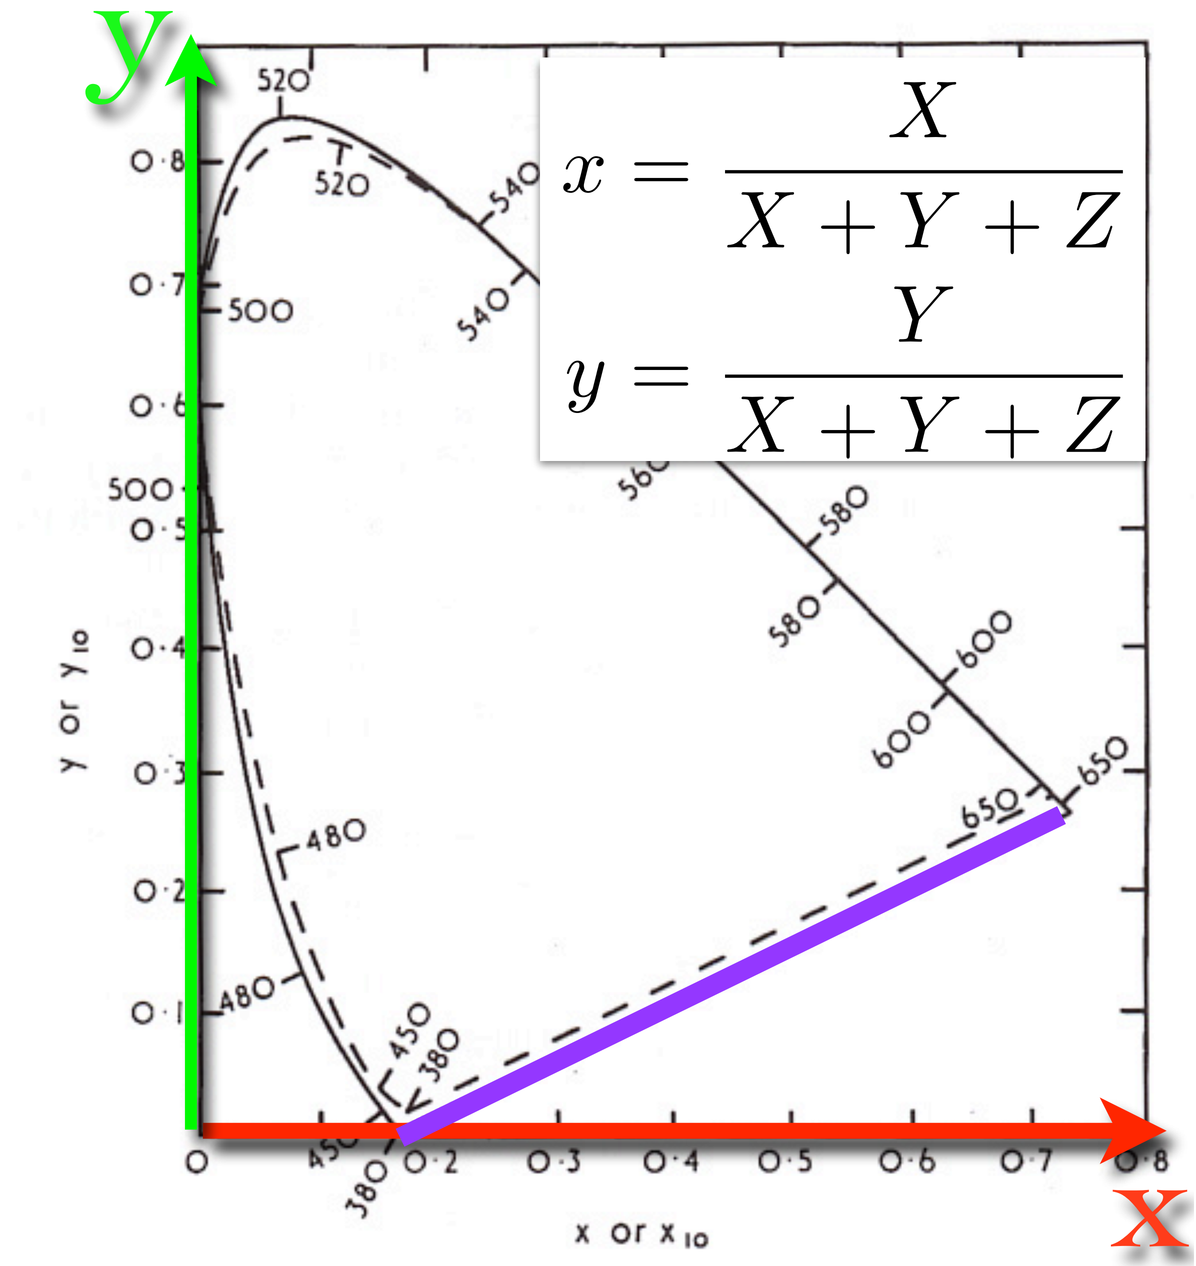

$$x = \frac{X}{X+Y+Z}$$

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

Define positive color matching functions (no real primaries)

x>0,y>0 does include all spectral colors

http://graphics.stanford.edu/courses/cs178/applets/threedgamut.html

CIE xy

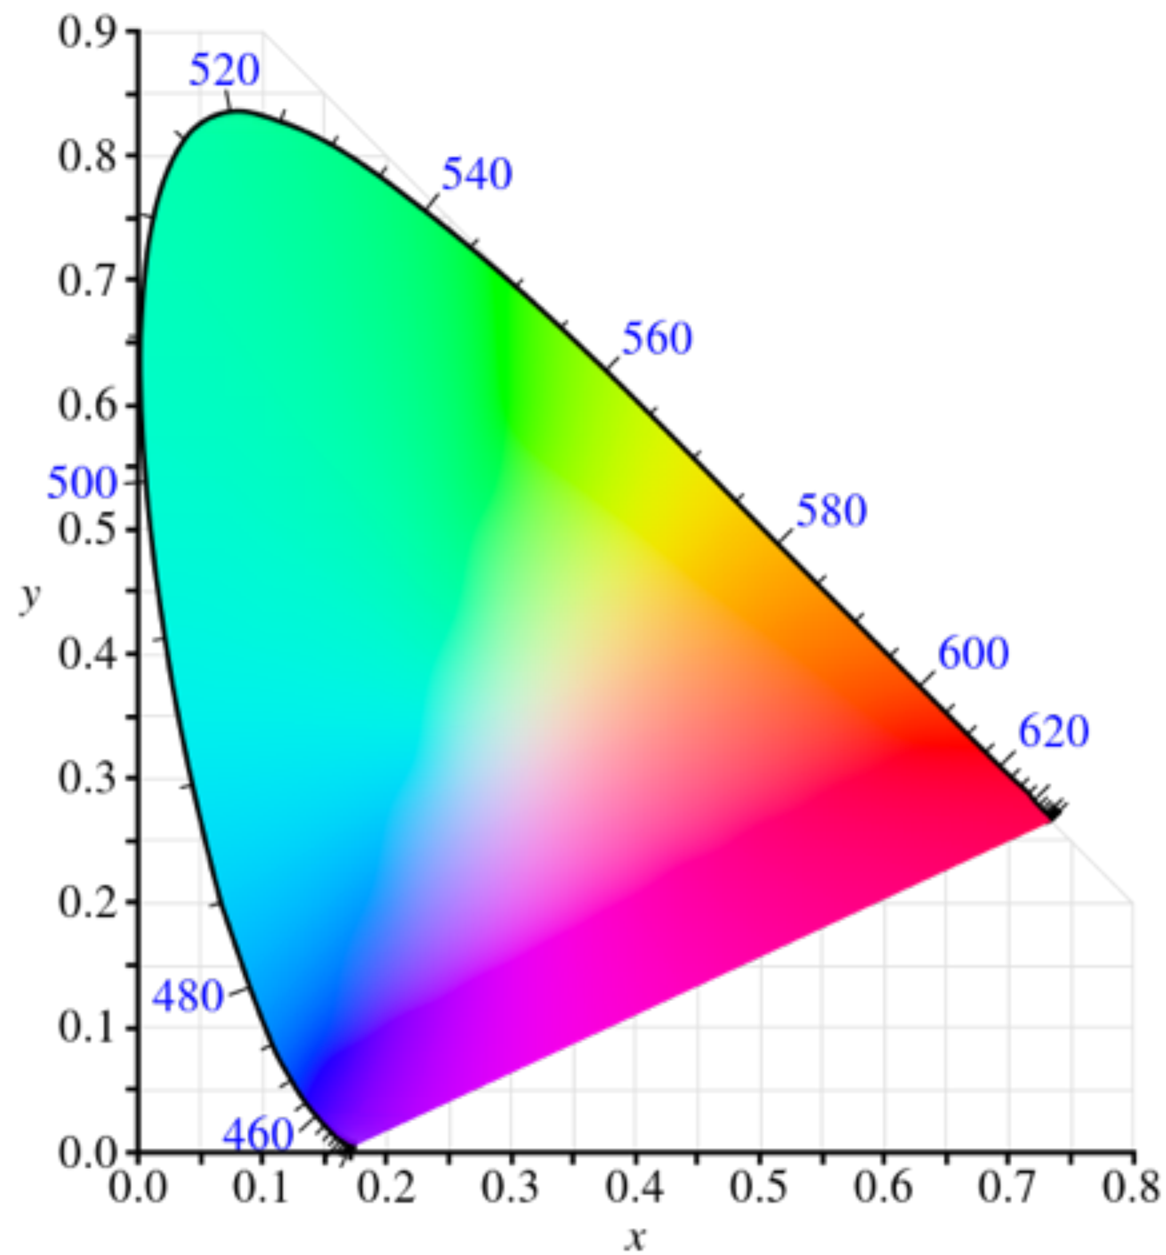

0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0.0
y
0.0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
x
520
540
560
580
600
620
500
480
460

CIE Luv (or Lab)

XYZ

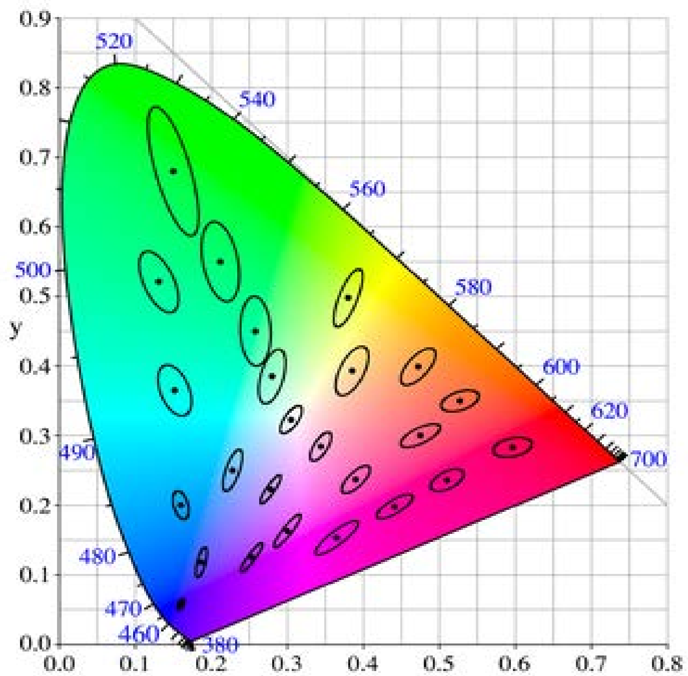

McAdam ellipses:
just-noticeable differences in color

Nonlinear (projective)
transformation

Lab

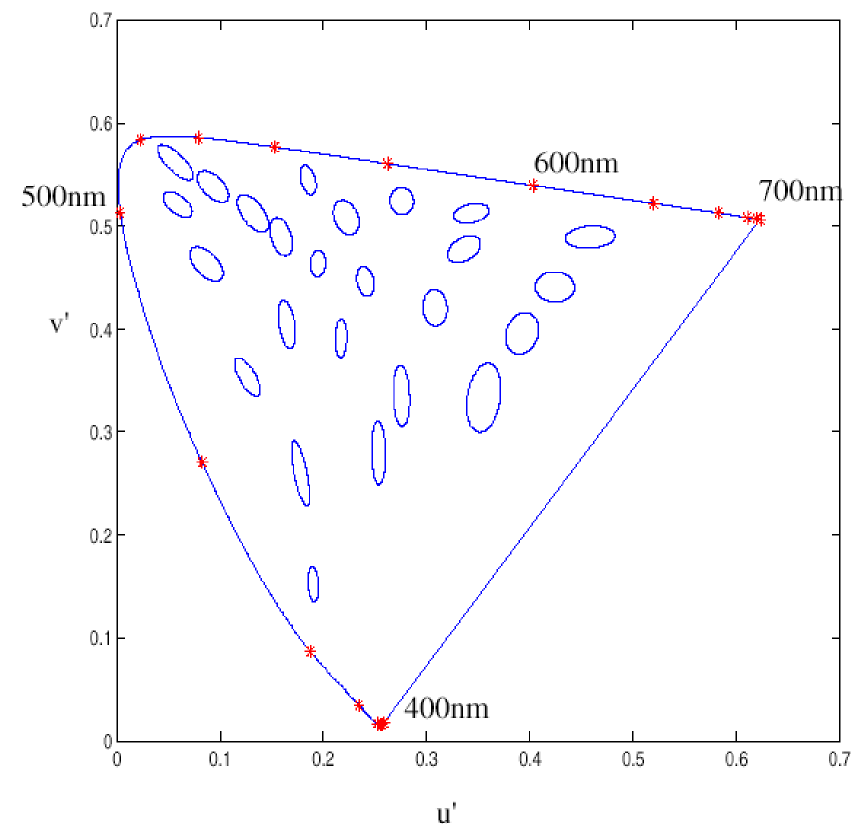

Ellipses are more circular

Luv,Lab is a nice color space to work with when measuring pixel differences (e.g., SSD)

L measures luminance (brightness) and u,v (or a,b) measures chrominance

Color constancy

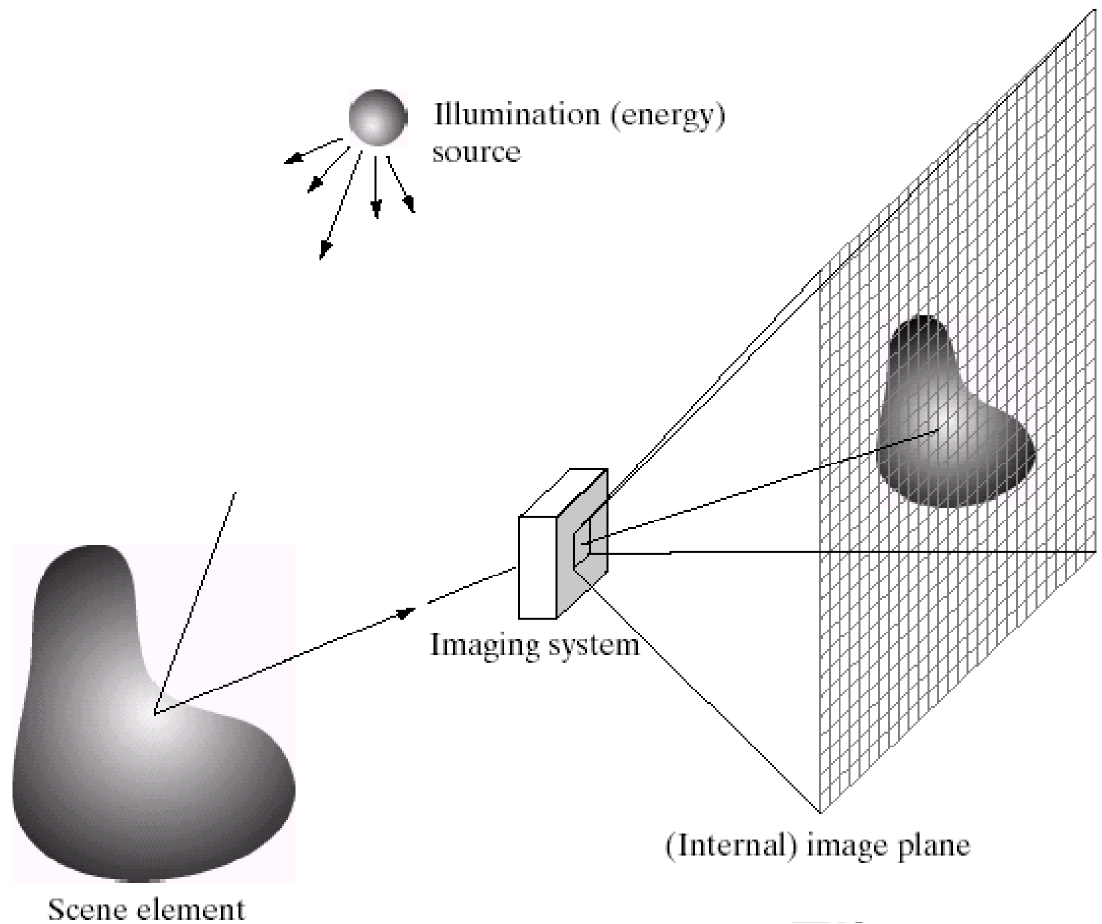

Appearance of surface = product of surface reflectance + illumination

Color constancy: an illumination-independent representation of the scene - what would it look like under white light?

Human color constancy

- Human visual system changes its sensitivity depending on the prevailing colors in the visual field
- Iris changes size (walking into a building from sunshine)
- If scene has lots of red light, red sensitivity of receptors decrease until scene looks white

Color constancy algorithms

Von Kries adaption model: multiple each RGB channel by gain factor

Easy approach:

1. Take a picture of a reference object (white or grey)
2. If object color is (r_w, g_w, b_w), use gains of $(1/r_w, 1/g_w, 1/b_w)$,

Color constancy algorithms

Statistically infer grey value from image

1. Grey-world assumption: assume average RGB color is grey
2. Assume brightest pixel value is a highlight with illuminant color
3. Gamut mapping: transform convex hull of all pixels in an image to gamut under “typical” white light

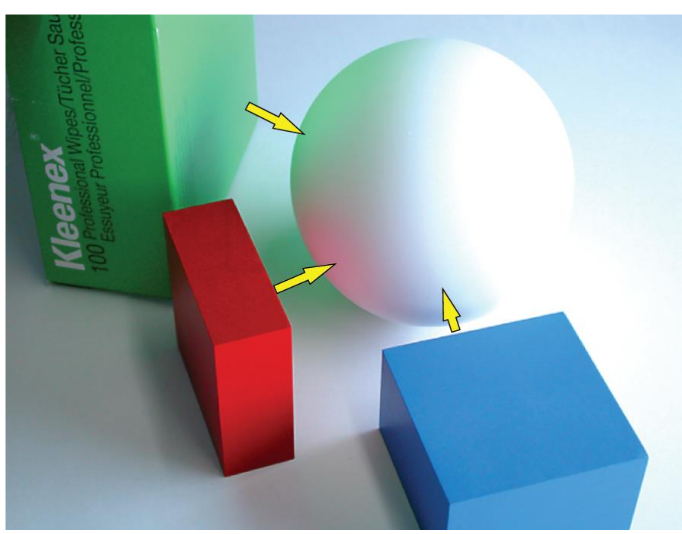

grey value = moon

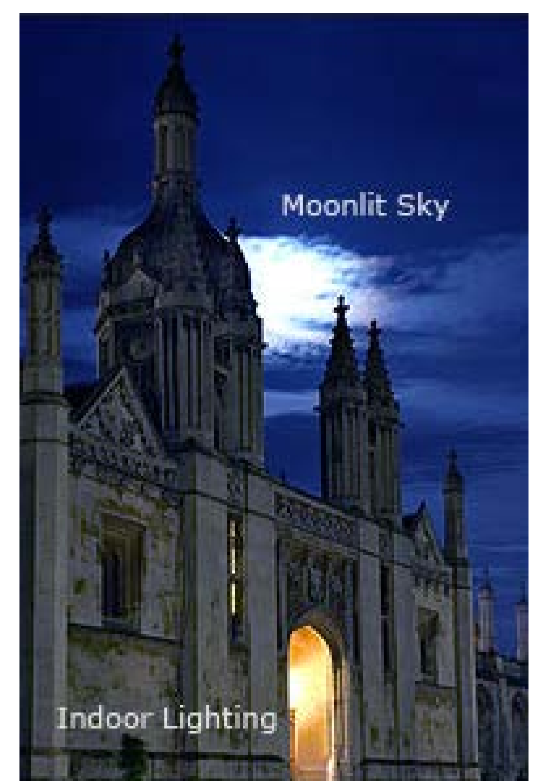

grey value = stone

Color as a cue for recognition

Usually gets clobbered by within-class variation

Color histograms for instance-level matching

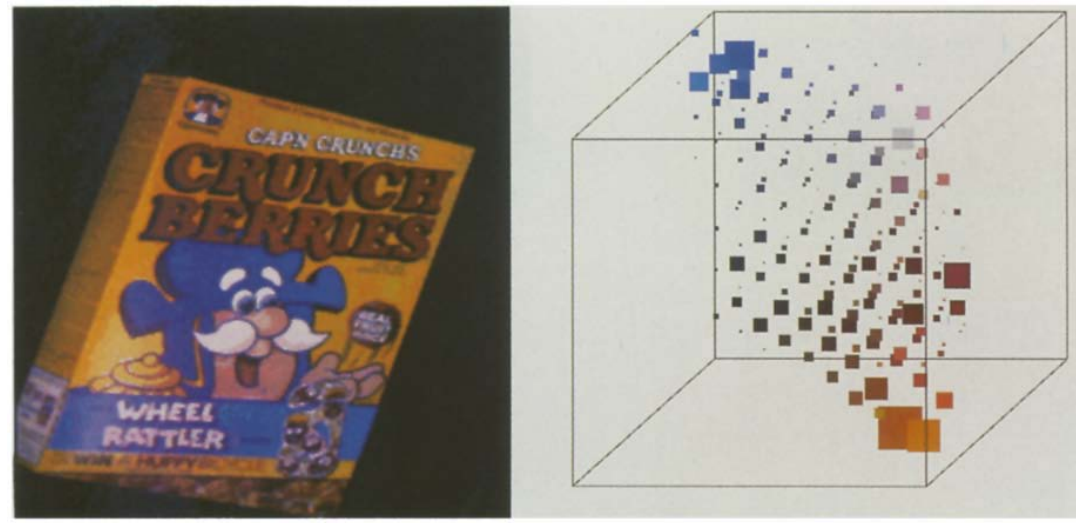

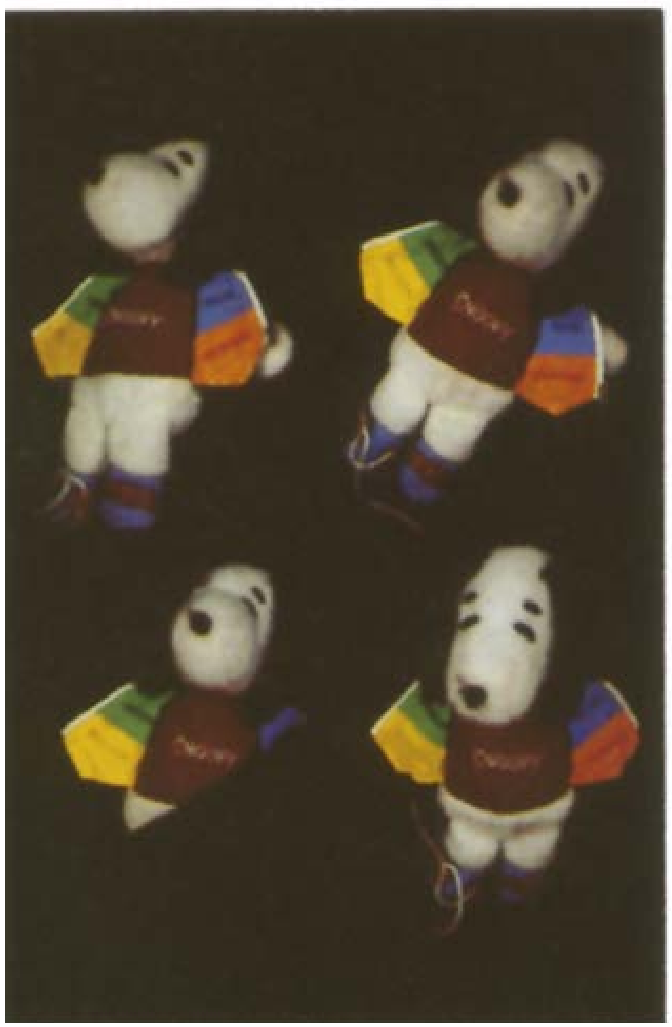

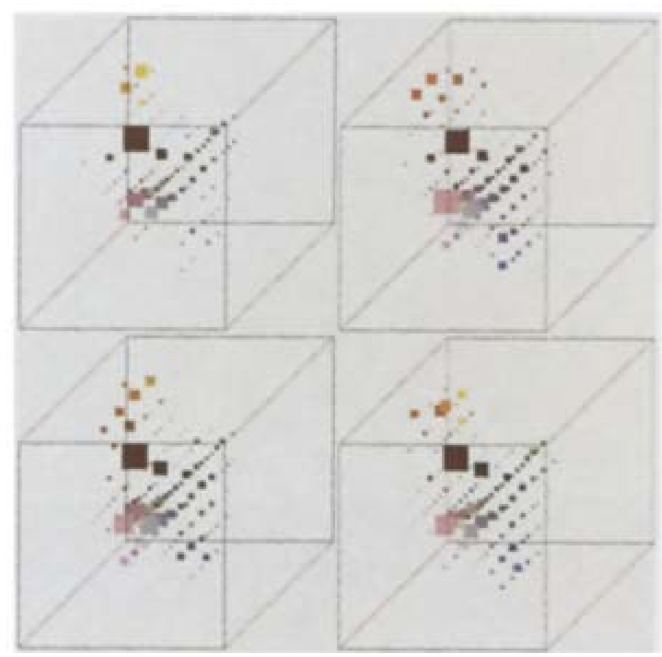

Swain & Ballard, “Color Indexing” IJCV 91

Skin detection

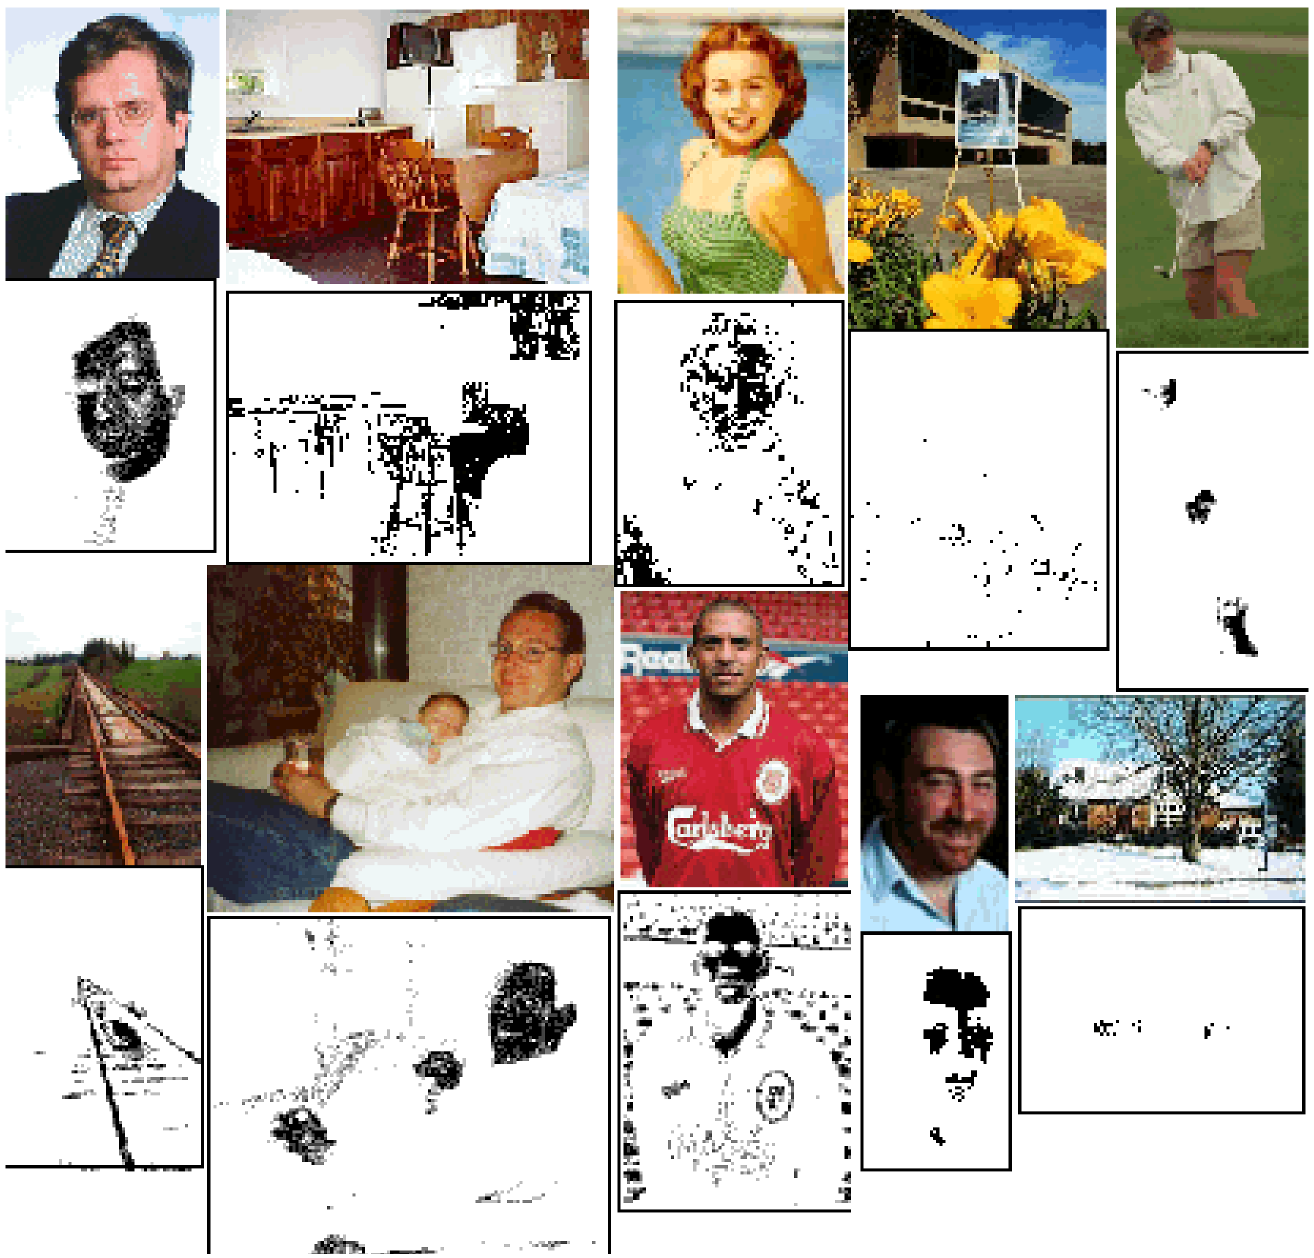

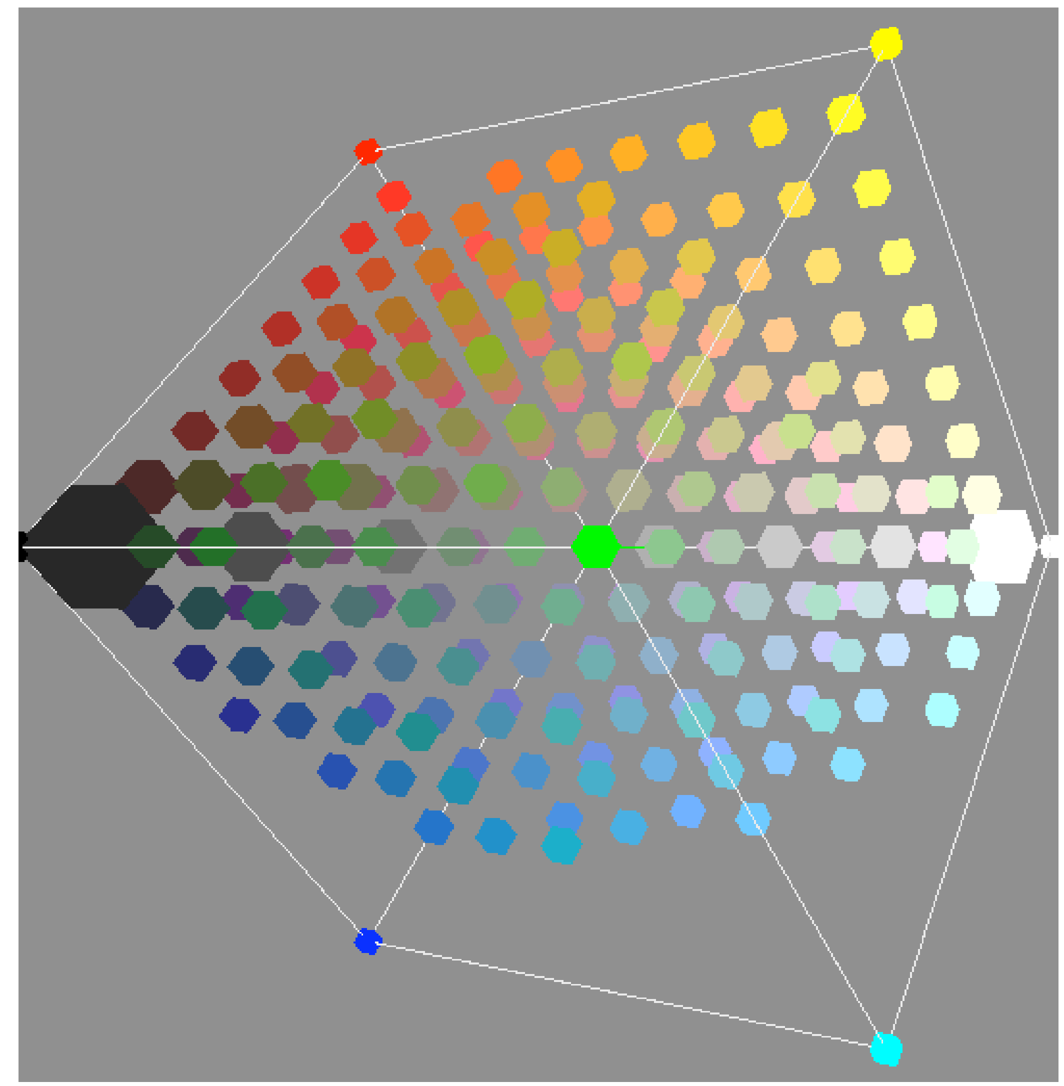

Jones & Rehg “Statistical Color Models with Application to Skin Detection” IJCV 02

Color histogram on RGB. Why not Lab?

SSD matching

Matching accuracy

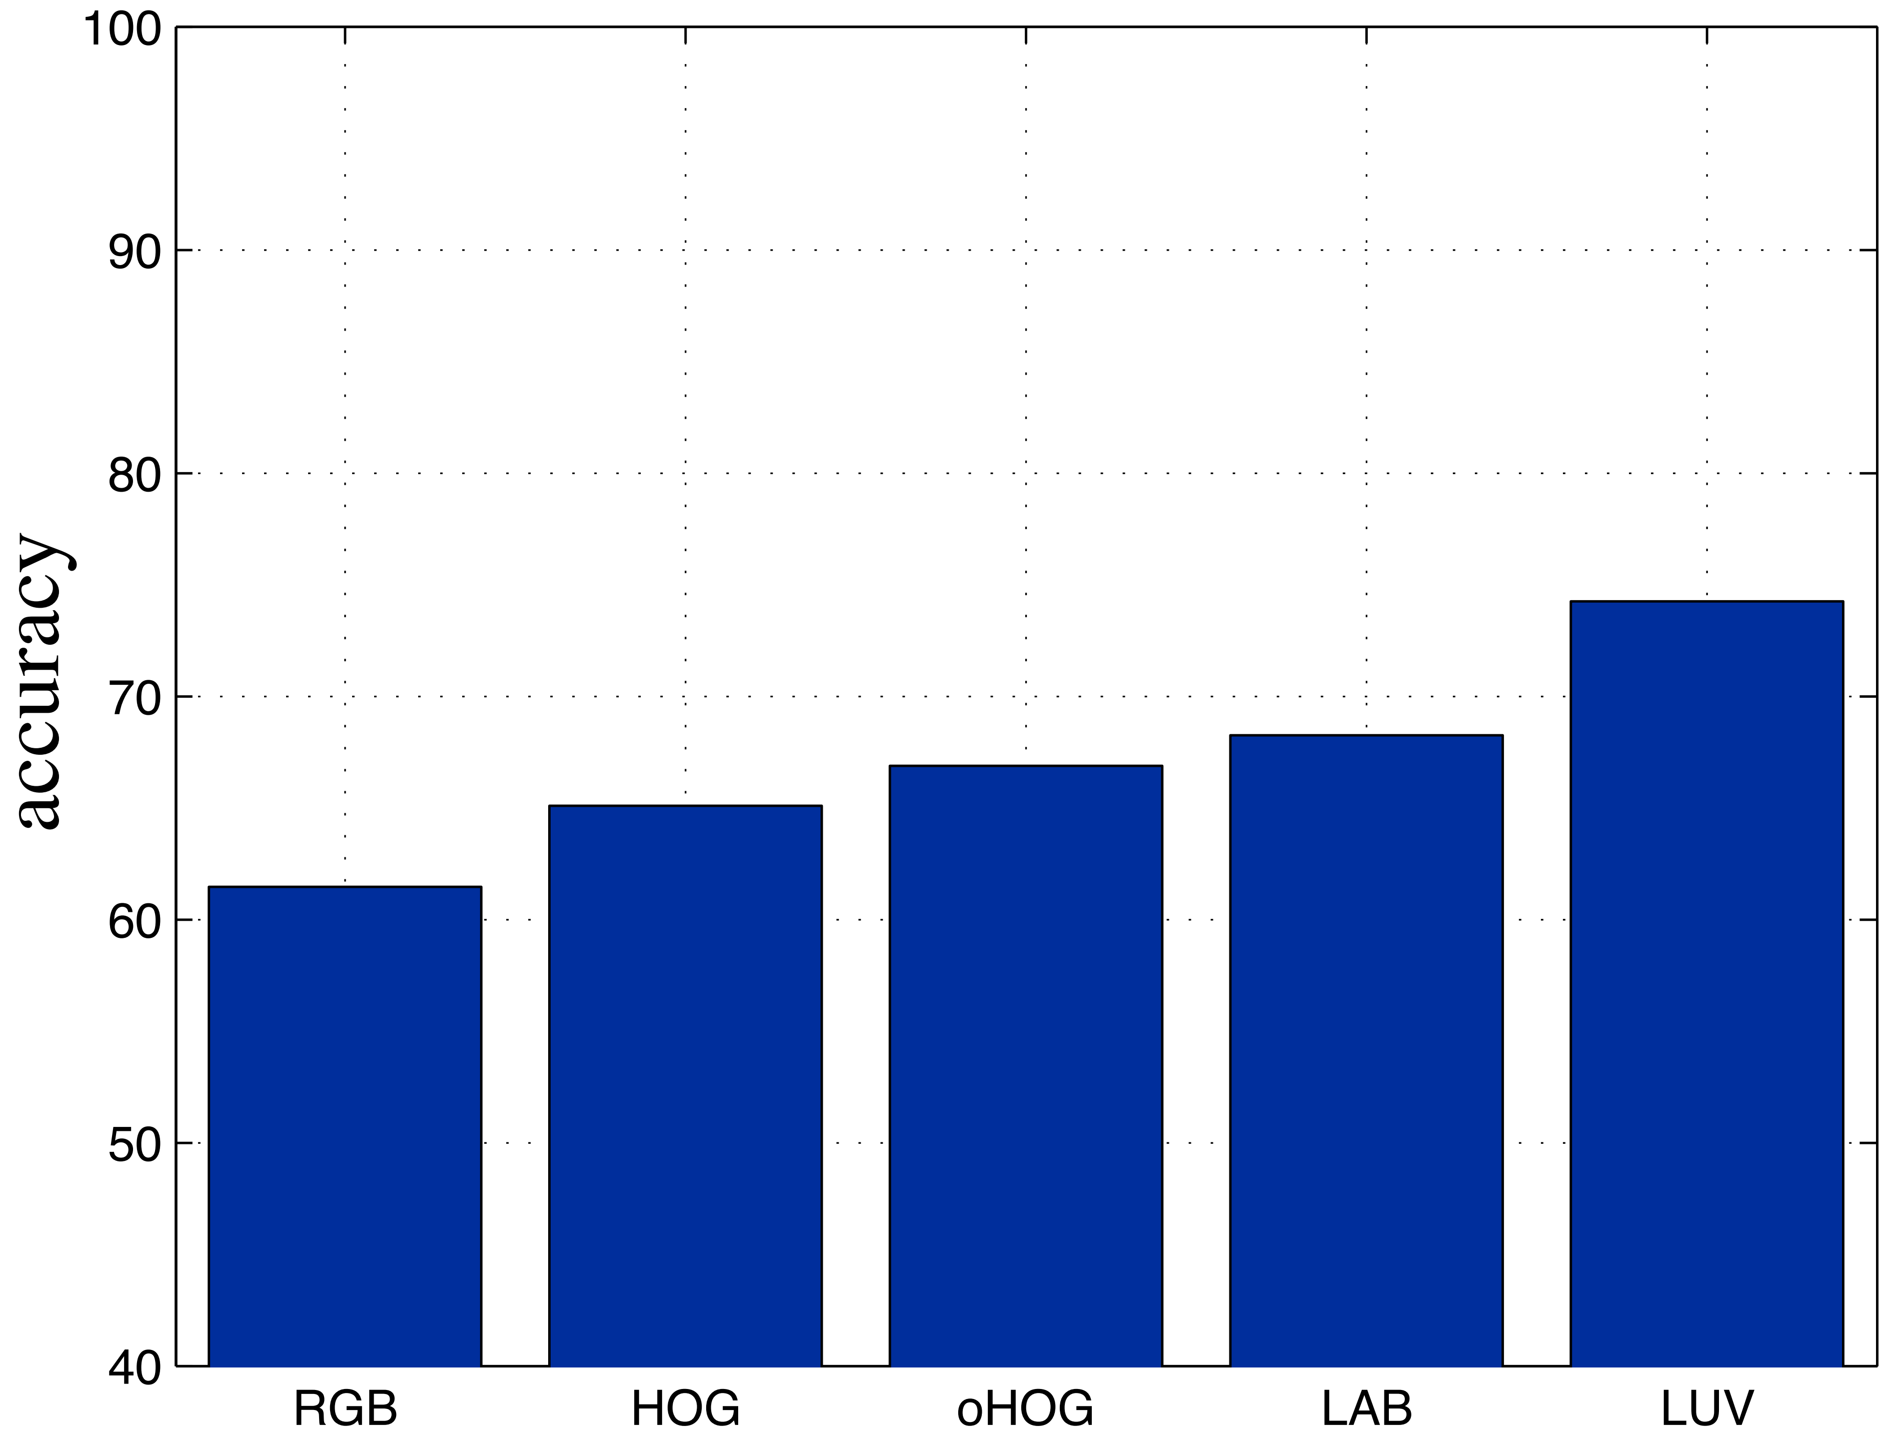

Park & Ramanan, 2015

Take-home points for color

- True “color” is represented as a continuous distribution over wavelengths

- Humans perceive color by projecting wavelength onto 3 color receptors types (implying metamers do exist, and they depend on the person/organism).

- As a consequence, color responses behave linearly. One convenient representation is a (trichromatic) color matching function.

- Lab,Luv are better color spaces for measuring euclidean distance