Introduction to computer vision VI

Instructors: Jean Ponce and Matthew Trager
jean.ponce@inria.fr, matthew.trager@cims.nyu.edu

TAs: Jiachen (Jason) Zhu and Sahar Siddiqui
jiachen.zhu@nyu.edu, ss12414@nyu.edu

Slides will be available after class at:
https://mtrager.github.io/introCV-fall2019/

Light, shading, color and shodows

- Radiometry
  - Radiance
  - Irradiance
  - BRDF
  - Photometric stereo

- Color
  - Spectral radiometry
  - Source and surface colors
  - Trichromacies
  - The eye
  - Color spaces
  - The bichromatic model

Light and shadows

[These slides courtesy of S. Narasimhan, CMU.]

Reflections

Reflections and refractions

Refraction

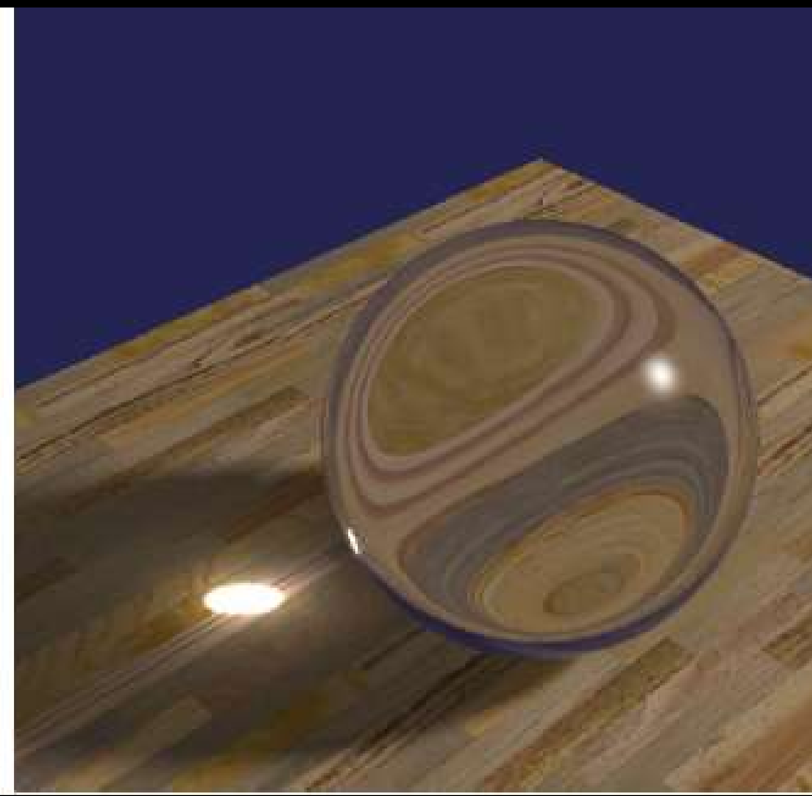

Caustics

Is glass really transparent?

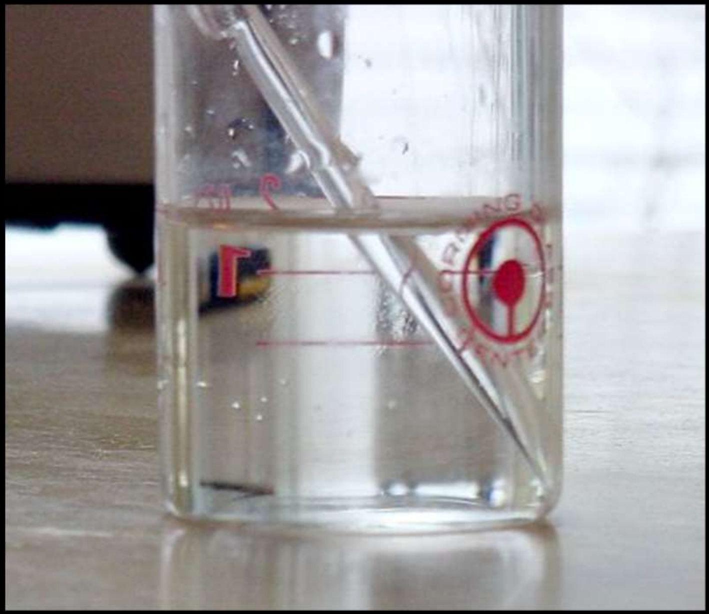

Interreflections

Scattering


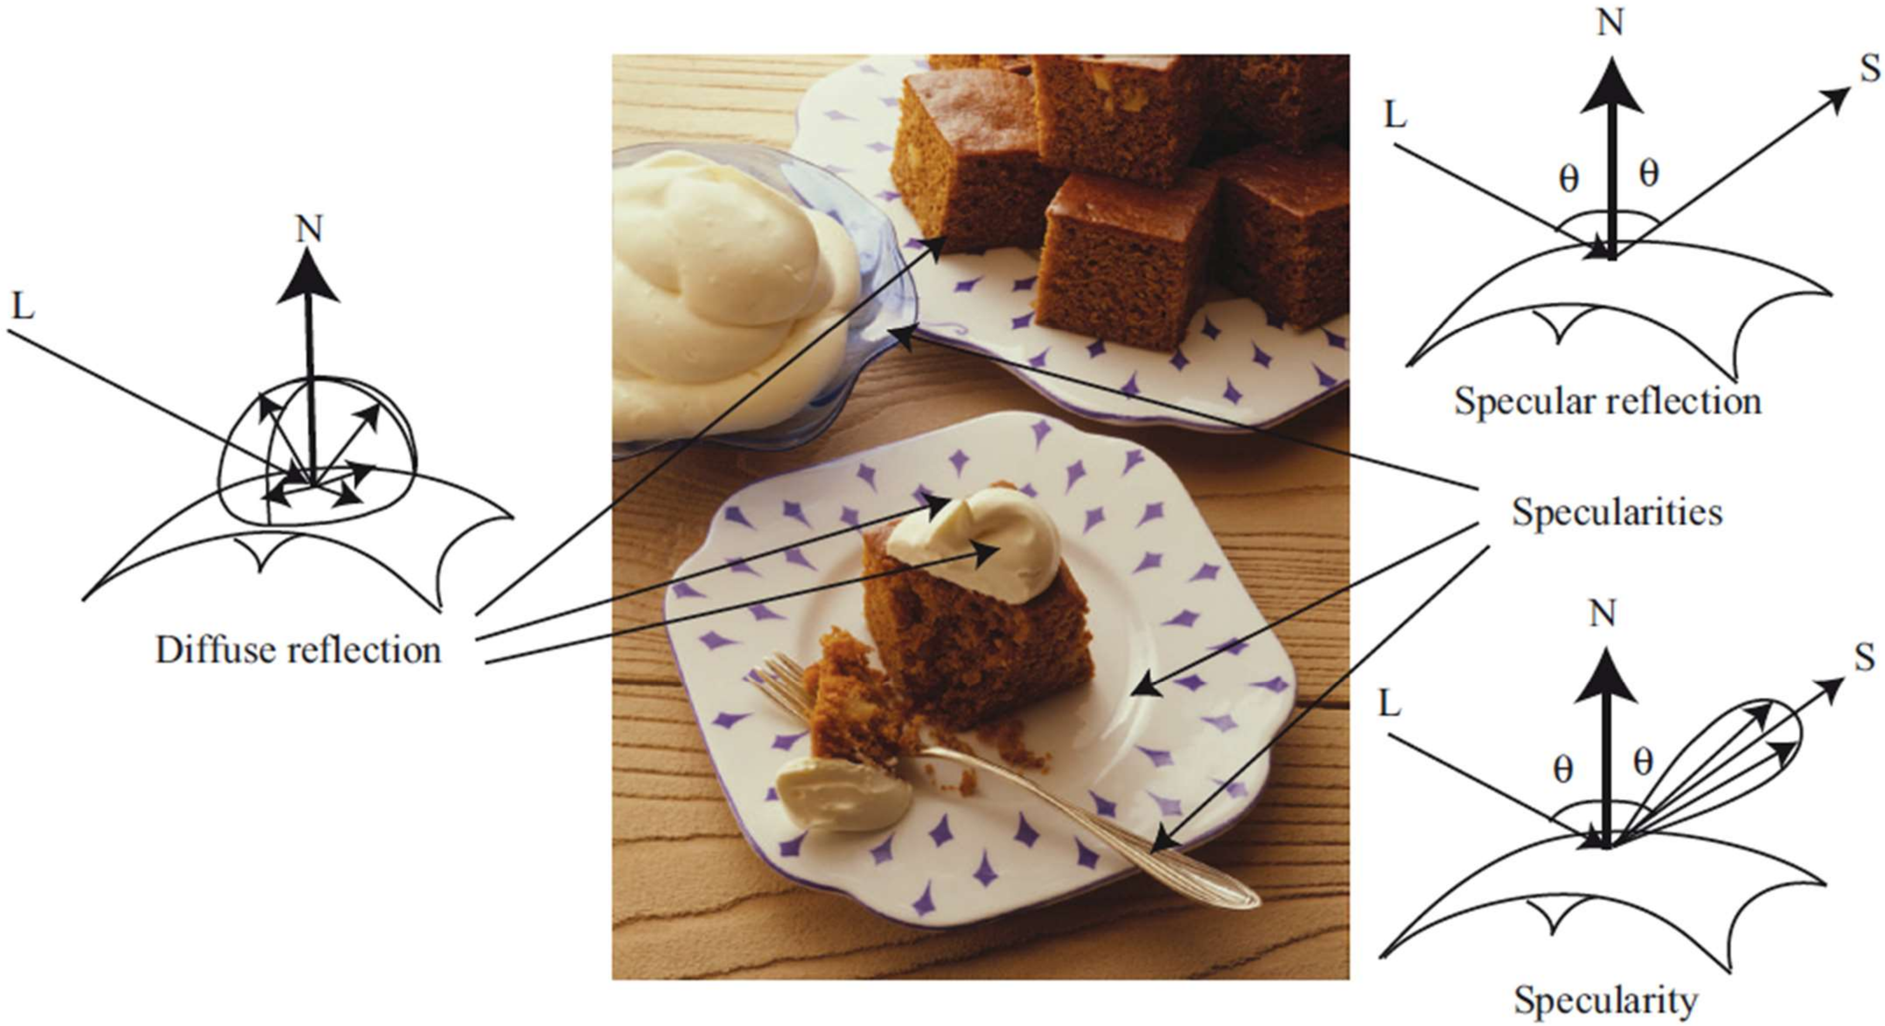

The two main types of reflection: diffuse and specular

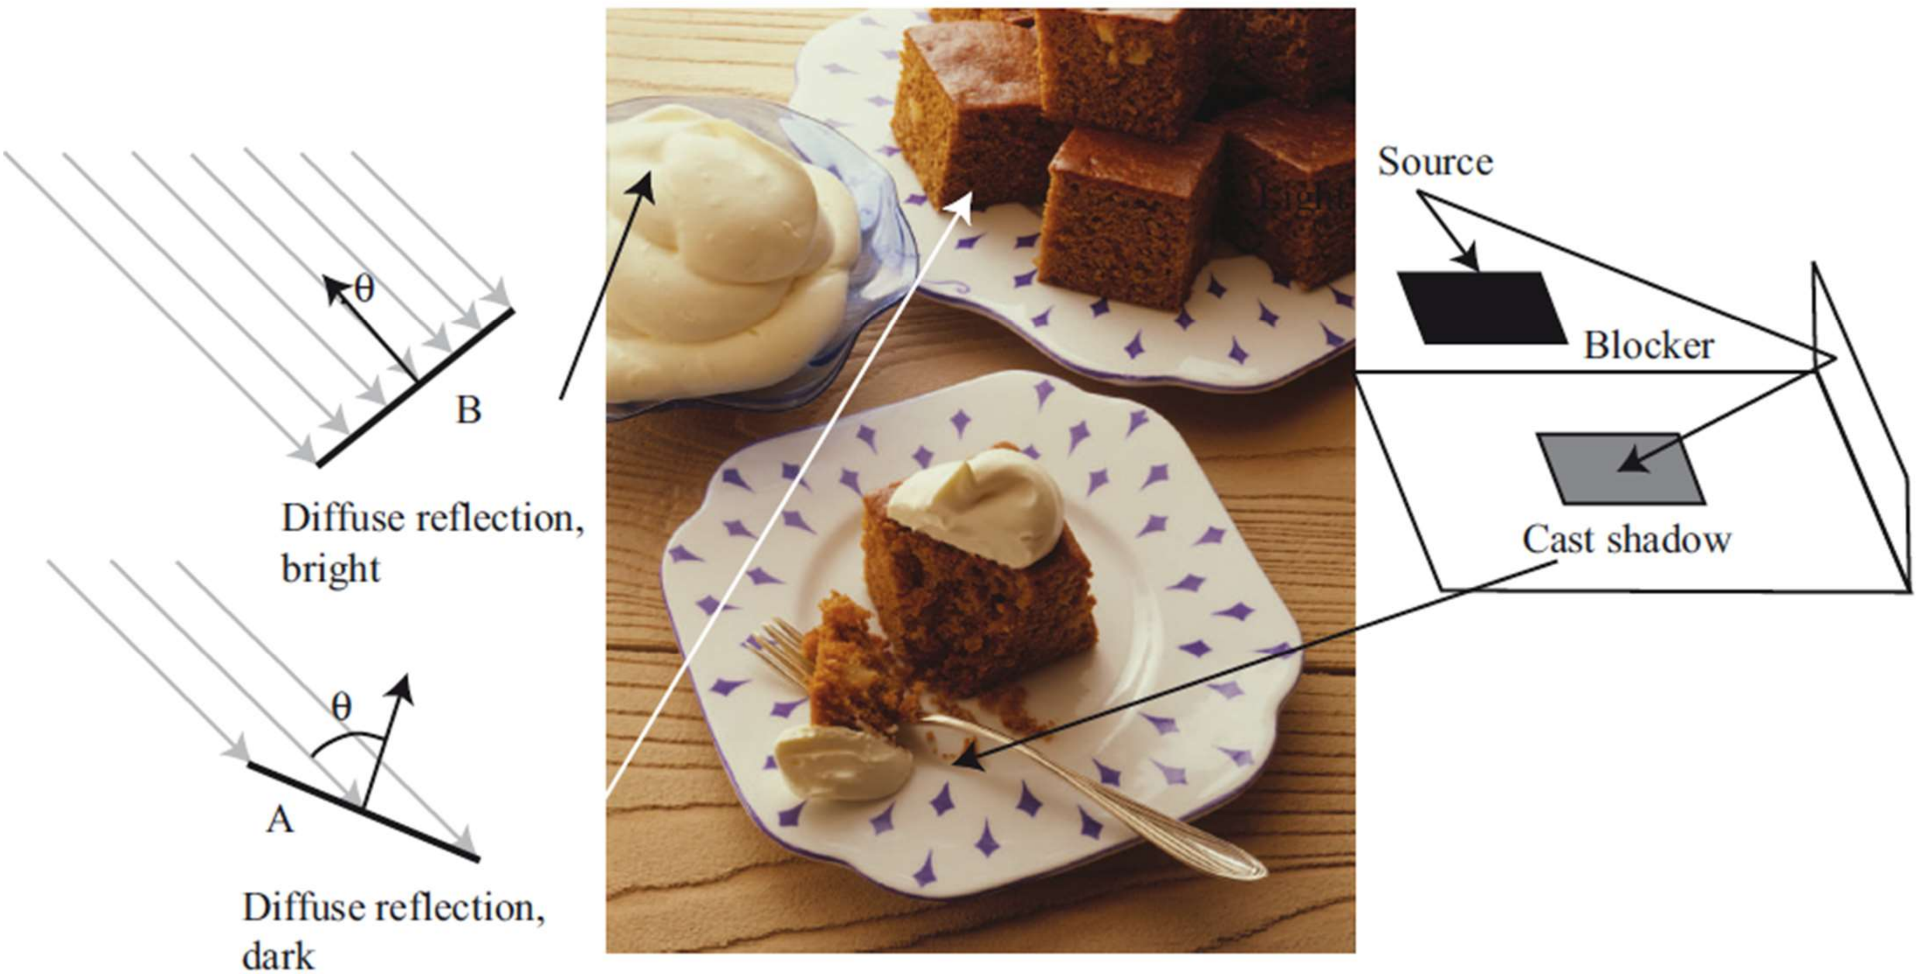

Shading and orientation (distant light sources)
Interreflections and soft shadows

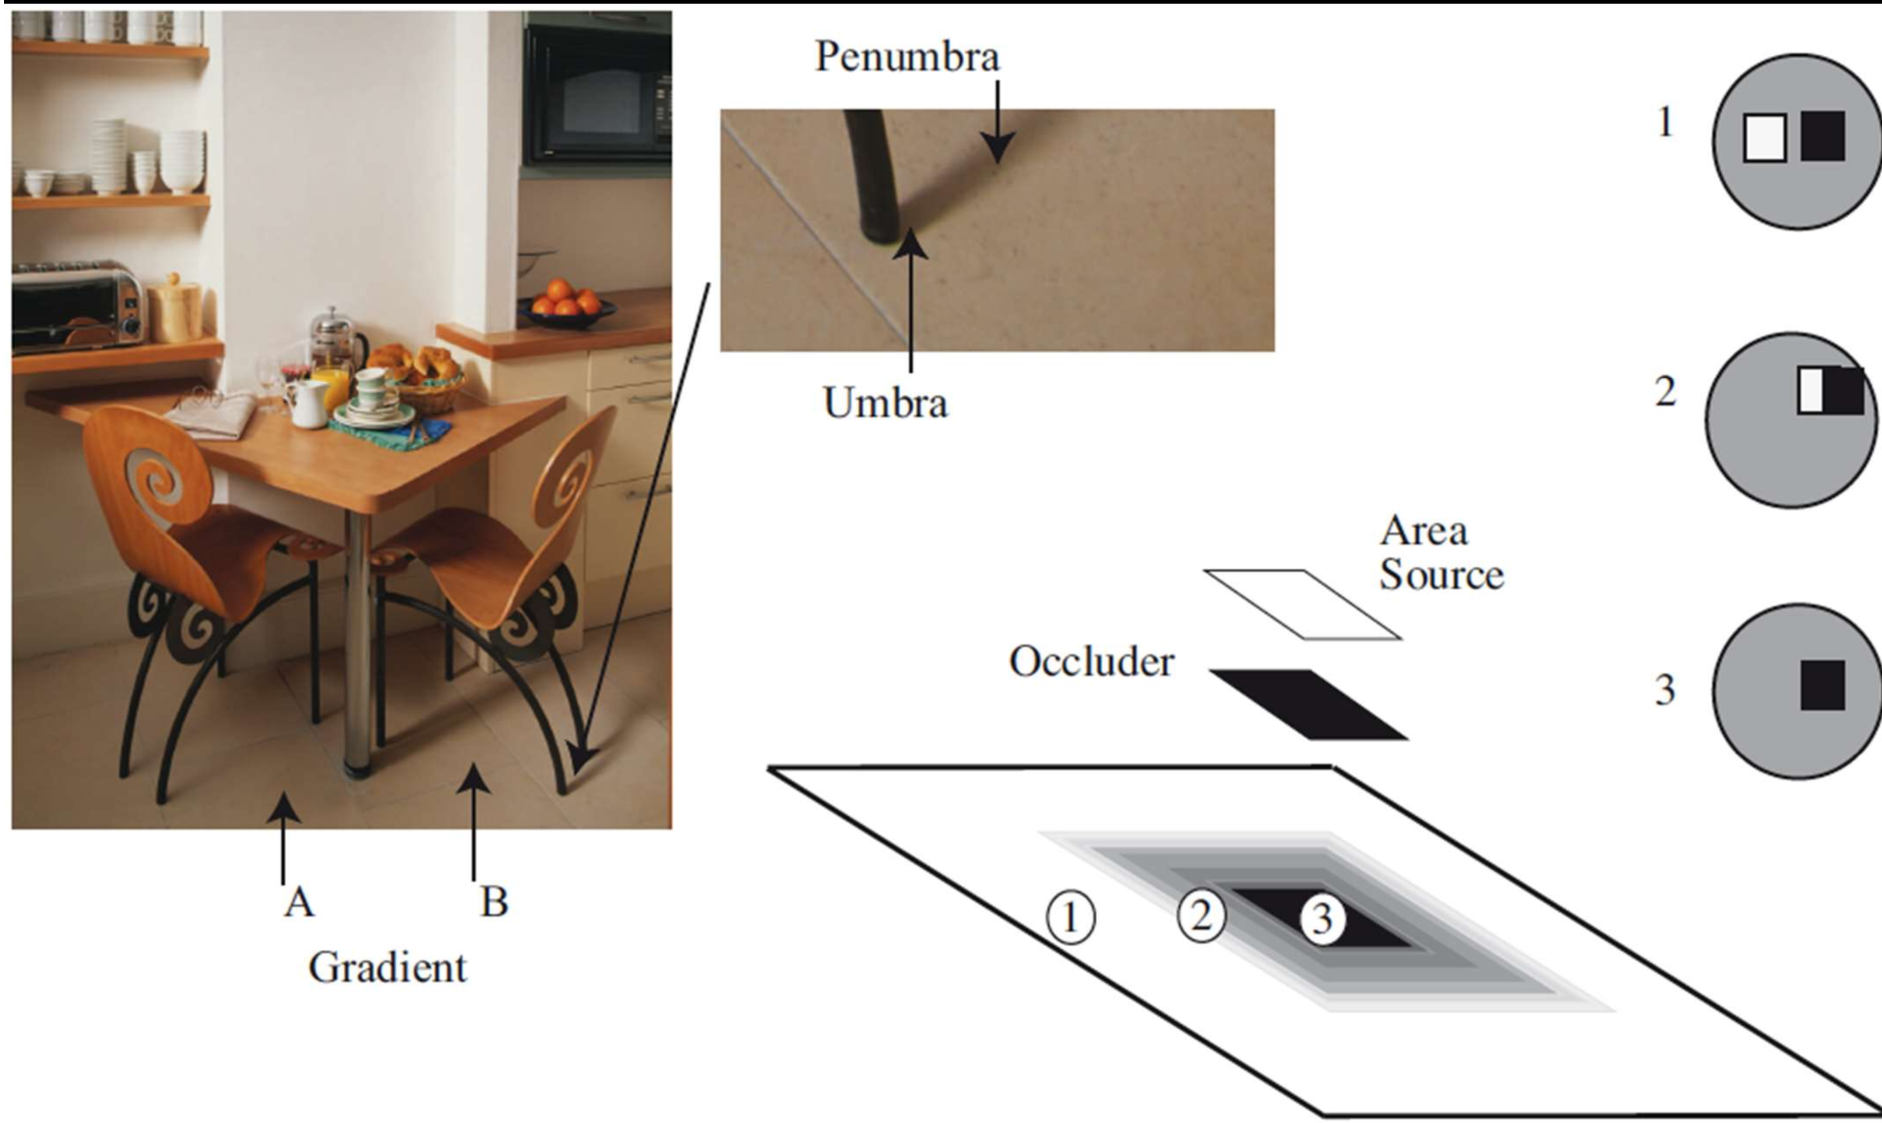

Area sources and soft shadows

Indoor scene with some directional componenent but almost no shadows: small occluders and large sources

Image Formation: Radiometry

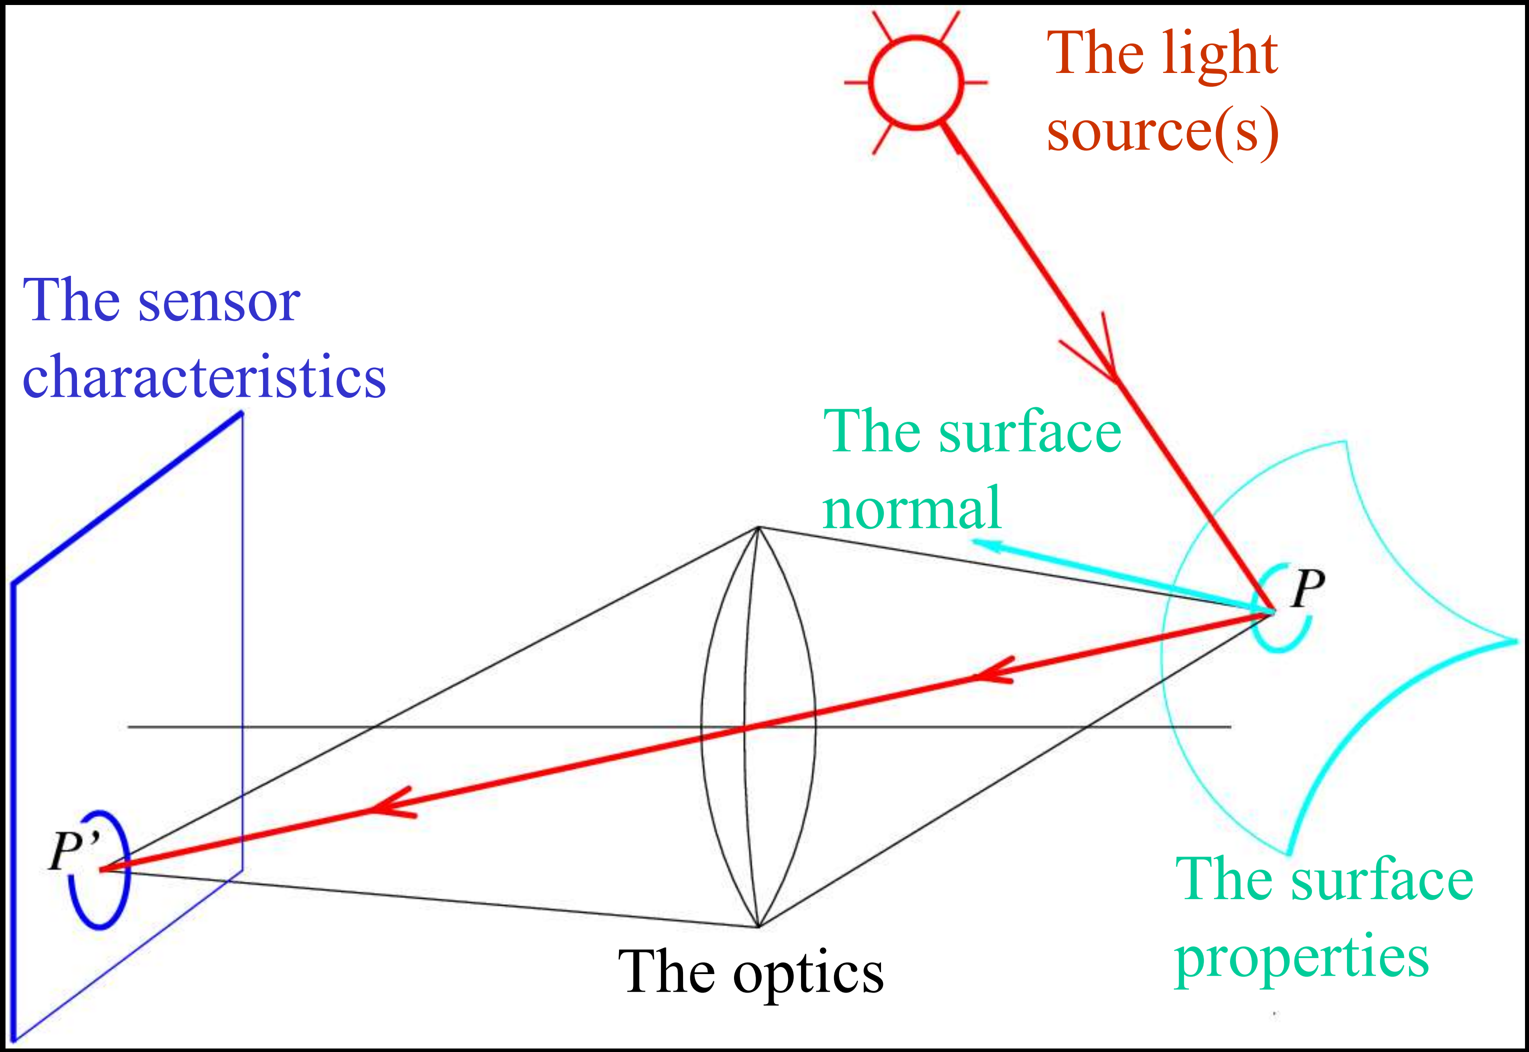

What determines the brightness of an image pixel?

DEFINITION: The radiance is the power traveling at some point in a given direction per unit area perpendicular to this direction, per unit solid angle.

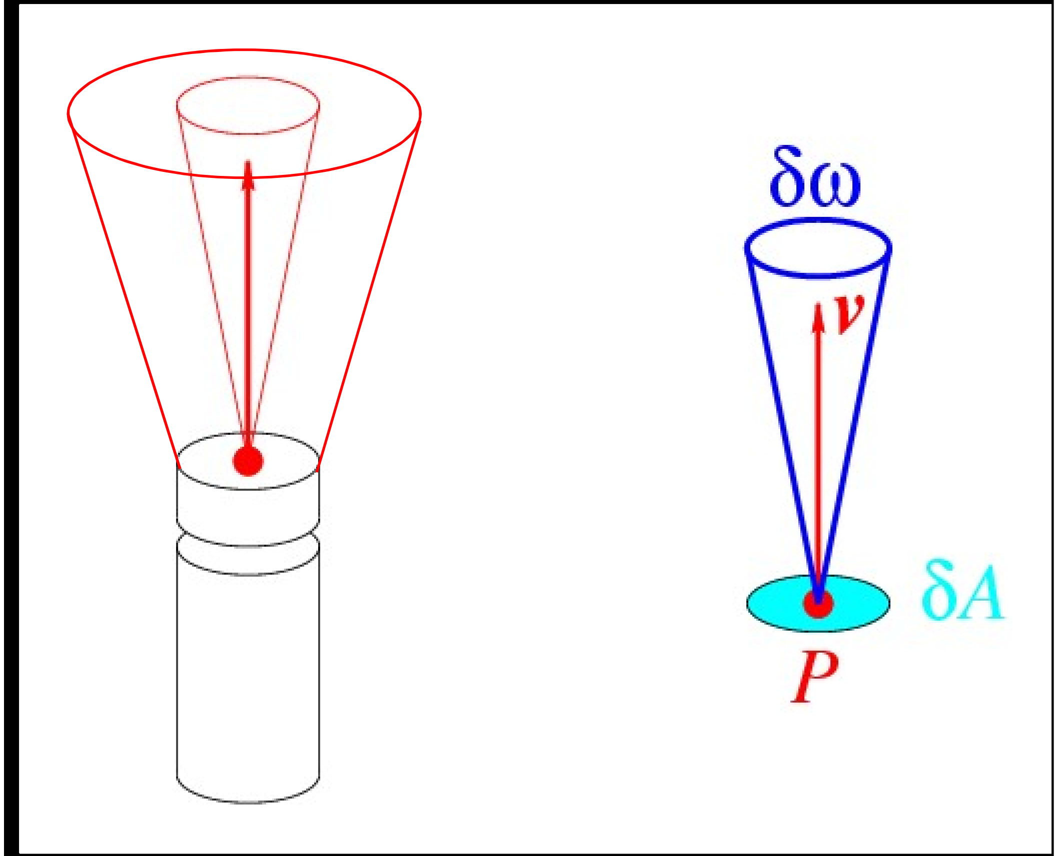

$$\boxed{\delta^2 P = L(P, v)\, \delta A\, \delta\omega}$$

DEFINITION: The radiance is the power traveling at some point in a given direction per unit area perpendicular to this direction, per unit solid angle.

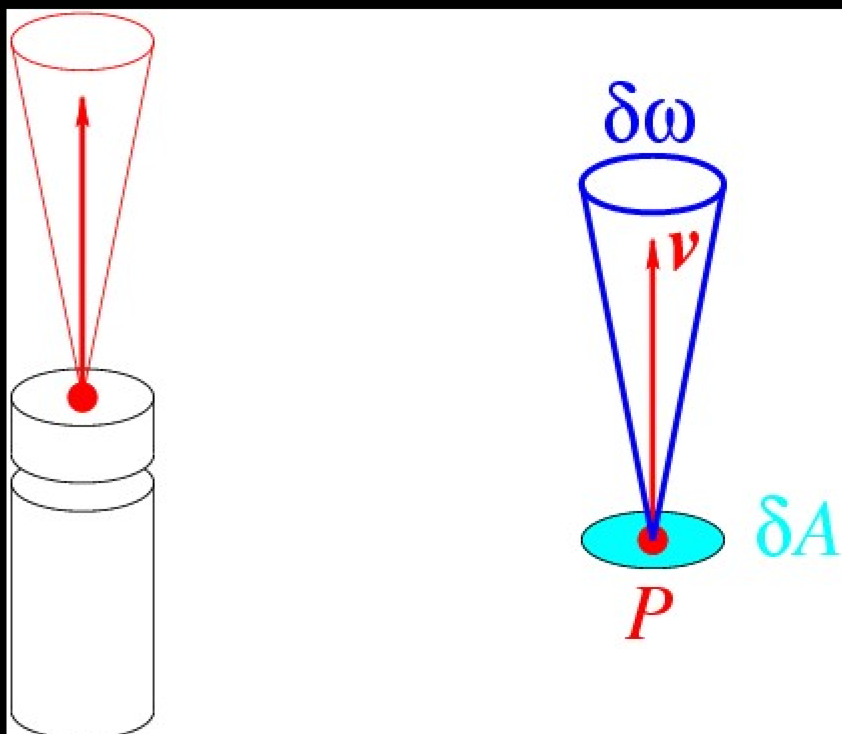

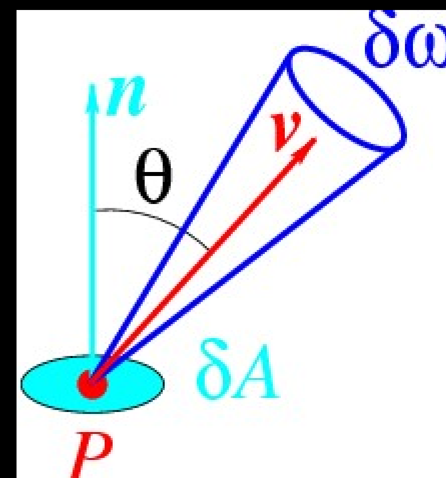

$$\boxed{\delta^2 P = L(\,P, \boldsymbol{v}\,)\ \delta A\ \delta\omega}$$

$$\boxed{\delta^2 P = L(\,P, \boldsymbol{v}\,)\ \cos\theta\ \delta A\ \delta\omega}$$

DEFINITION: Angles and Solid Angles

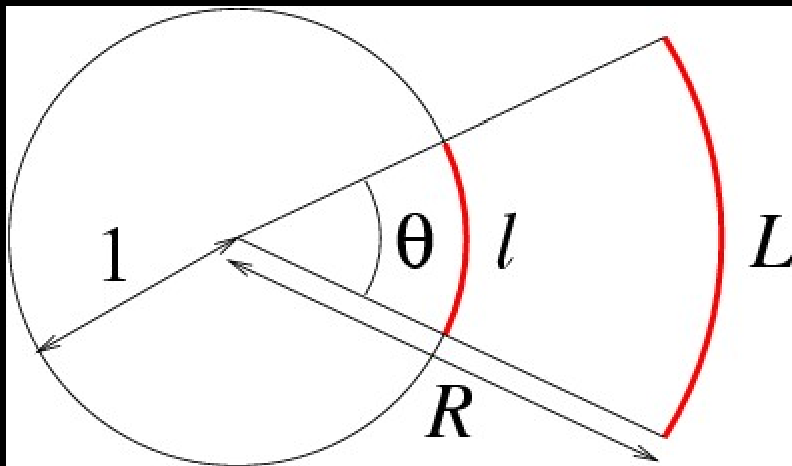

$$\boxed{\theta = l = \frac{L}{R}}$$ (radians)

DEFINITION: Angles and Solid Angles

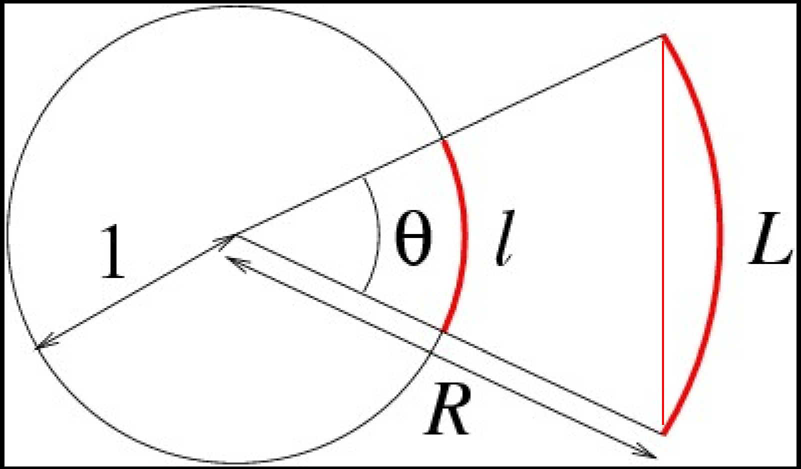

$$\boxed{\theta = l = \frac{L}{R}}$$ (radians)

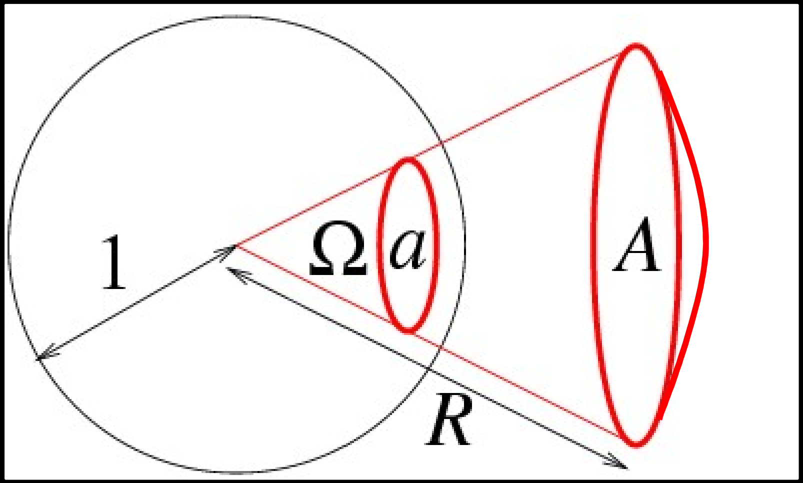

$$\boxed{\Omega = a = \frac{A}{R^2}}$$ (steradians)

PROPERTY: Radiance is constant along straight lines (in vacuum).

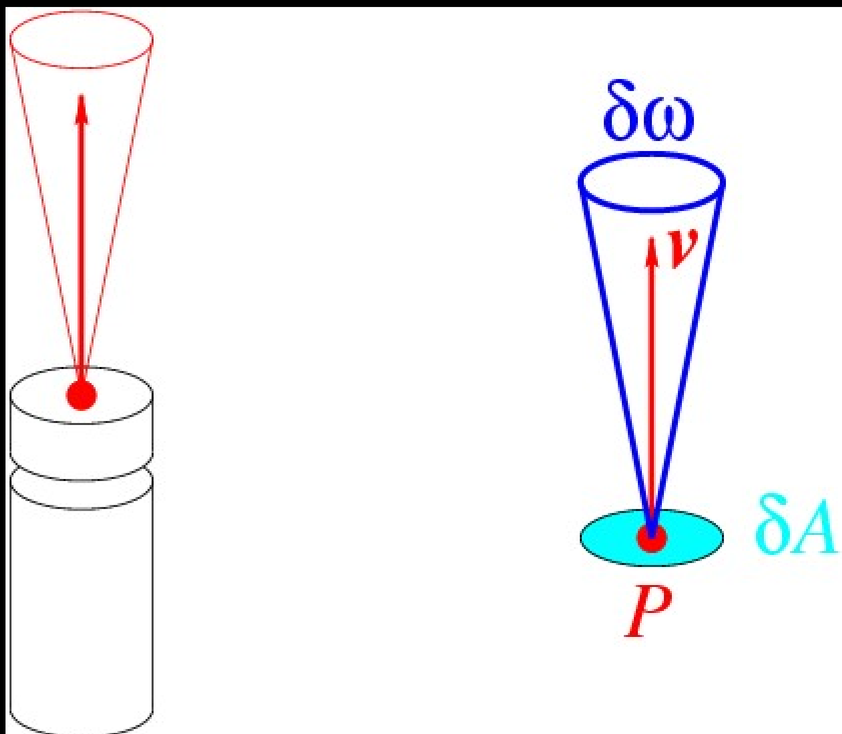

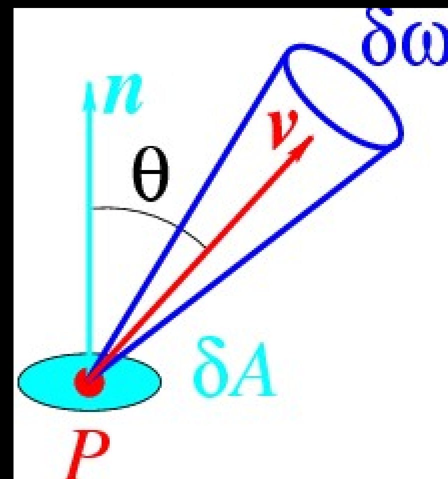

$$\boxed{\delta^2 P = L(P, \boldsymbol{v})\, \delta A\, \delta\omega} \Rightarrow \boxed{\delta^2 P = L(P, \boldsymbol{v}) \cos\theta\, \delta A\, \delta\omega}$$

DEFINITION: Irradiance

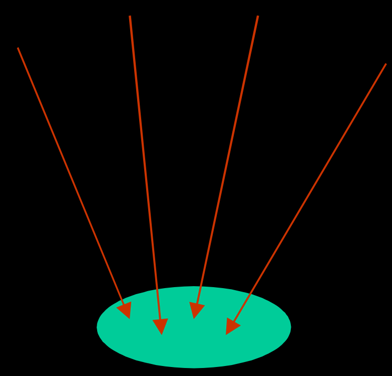

The irradiance is the power per unit area incident on a surface.

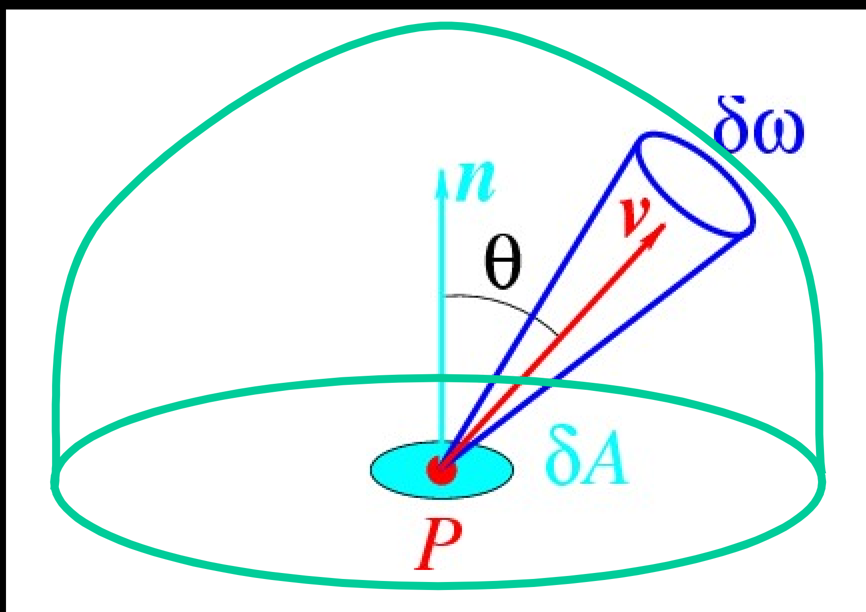

$$\delta^2 P = \delta E\ \delta A = L_i(P, \boldsymbol{v}_i) \cos\theta_i \delta\omega_i\ \delta A$$

$$\boxed{\delta E = L_i(P, \boldsymbol{v}_i) \cos\theta_i\ \delta\omega_i}$$

$$\boxed{E = \int_H L_i\ (P, v_i) \cos\theta_i\ d\omega_i}$$

Photometry

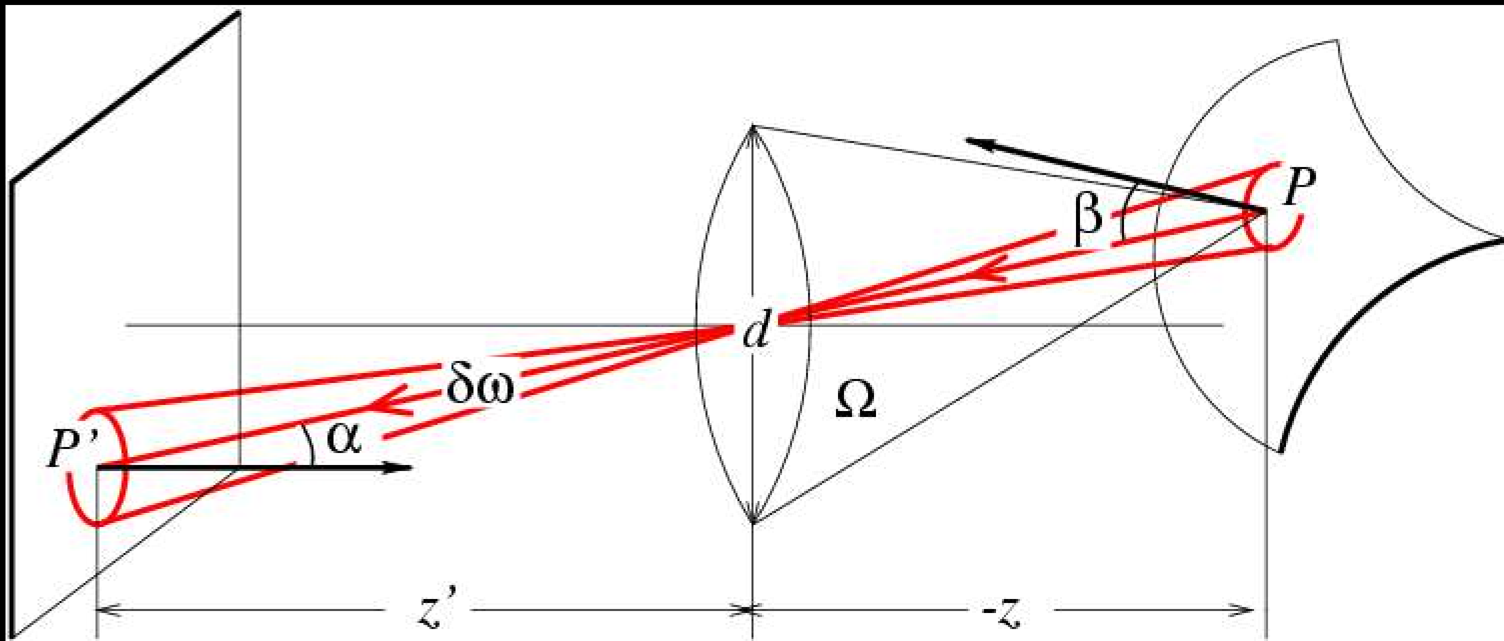

$$E = \left[\frac{\pi}{4} \left(\frac{d}{z'} \right)^2 \cos^4 \alpha \right] L$$

- L is the radiance.
- E is the irradiance.

DEFINITION: The Bidirectional Reflectance Distribution Function (BRDF)

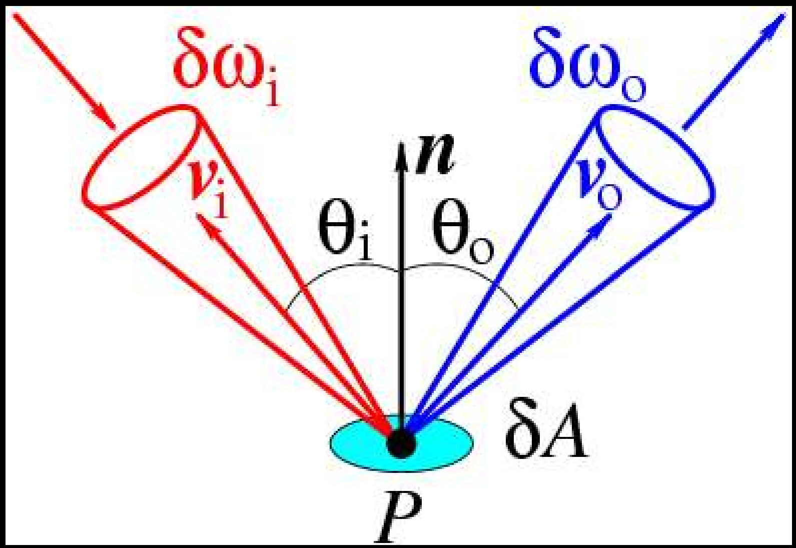

The BRDF is the ratio of the radiance in the outgoing direction to the incident irradiance (sr^{-1}).

$$\delta L_o (P, \boldsymbol{v}_o) = \rho_{BD} (P, \boldsymbol{v}_i , \boldsymbol{v}_o) \, \delta E_i (P, \boldsymbol{v}_i)$$

$$= \rho_{BD} (P, \boldsymbol{v}_i , \boldsymbol{v}_o) \, L_i (P, \boldsymbol{v}_i) \cos \theta_i \, \delta\omega_i$$

Helmoltz reciprocity law: $\rho_{BD} (P, \boldsymbol{v}_i , \boldsymbol{v}_o) = \rho_{BD} (P, \boldsymbol{v}_o , \boldsymbol{v}_i)$

DEFINITION: Specular Surfaces as Perfect or Rough Mirrors

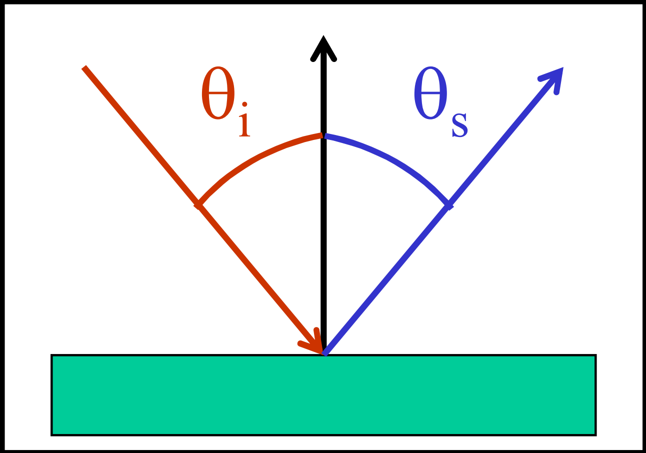

Perfect mirror

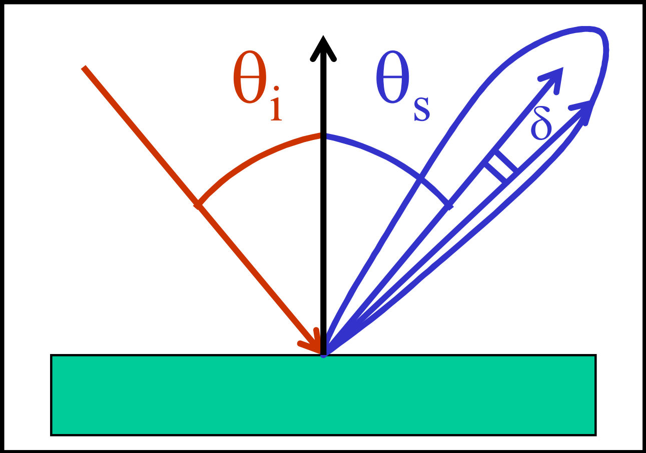

Rough mirror

Perfect mirror: $L_o\,(P,\boldsymbol{v}_s) = L_i\,(P,\boldsymbol{v}_i)$ [BRDF=double delta function]

Phong (non-physical model): $L_o\,(P,\boldsymbol{v}_o)= \rho_s L_i(P,\boldsymbol{v}_i)\cos^n \delta$

Hybrid model: $L_o\,(P,\boldsymbol{v}_o)= \rho_d \int_H L_i(P, \boldsymbol{v}_i)\cos\theta_i\, d\omega_i + \rho_s L_s(P,\boldsymbol{v}_i)\cos^n \delta$

DEFINITION: Radiosity

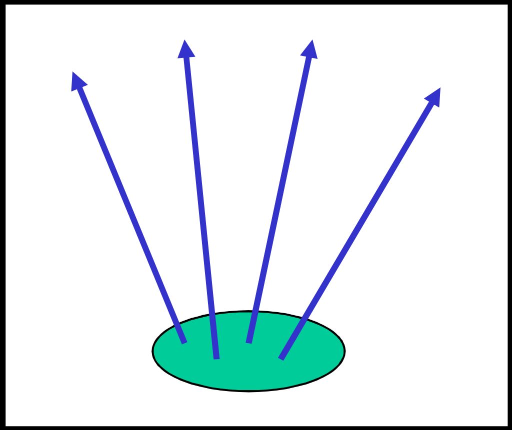

The radiosity is the total power Leaving a point on a surface per unit area (W x m^{-2}).

$$\boxed{B(P) = \int_H L_o (P , \boldsymbol{v}_o) \cos\theta_o \, d\omega}$$

Important case: L_o is independent of $\boldsymbol{v}_o$.

$$B(P) = \pi L_o(P)$$

DEFINITION: Lambertian (or Matte) Surfaces

A Lambertian surface is a surface whose BRDF is independent of the outgoing direction (and by reciprocity of the incoming direction as well).

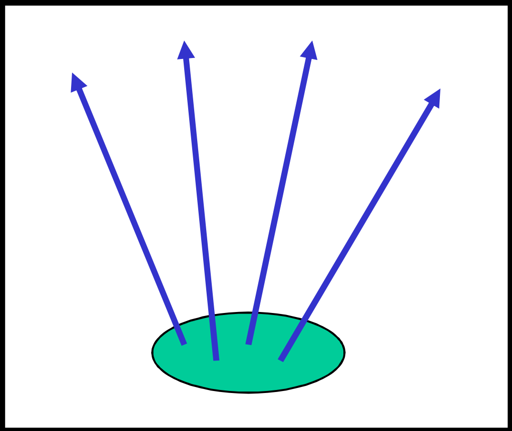

$$\rho_{BD}(\, \boldsymbol{v}_i \,, \boldsymbol{v}_o \,) = \rho_{BD} = \text{constant.}$$

Note: $1 \geq \Pi\, \rho_{BD}$ ➡ The albedo is $\rho_d = \Pi\, \rho_{BD}$.

DEFINITION: Point Light Sources

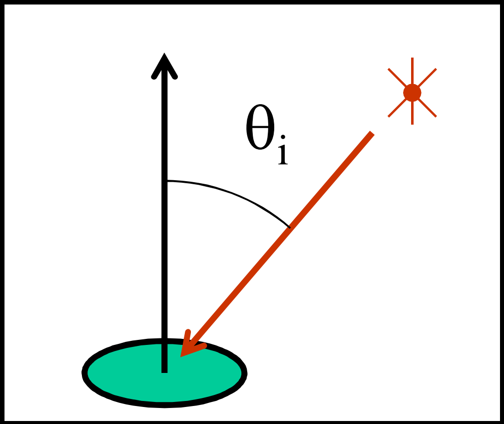

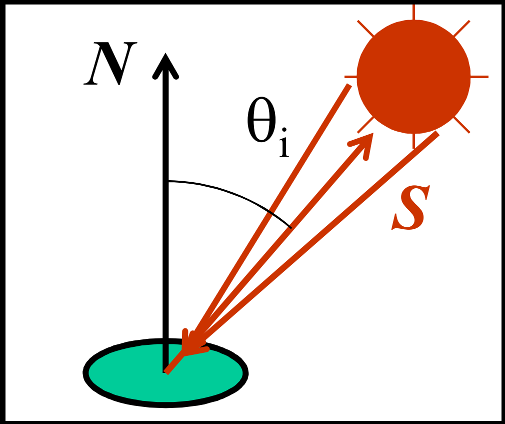

A point light source is an idealization of an emitting sphere with radius ε at distance R, with ε << R and uniform radiance L_e emitted in every direction.

For a Lambertian surface, the corresponding radiosity is

$$B(P) = \left[\rho_d(P)\ L_e \frac{\Pi \varepsilon^2}{R(P)^2} \right] \cos\theta_i \approx \rho_d(P) \frac{\mathbf{N}(P) \cdot \mathbf{S}(P)}{R(P)^2}$$

Lambertian sphere with the sun behind its observer

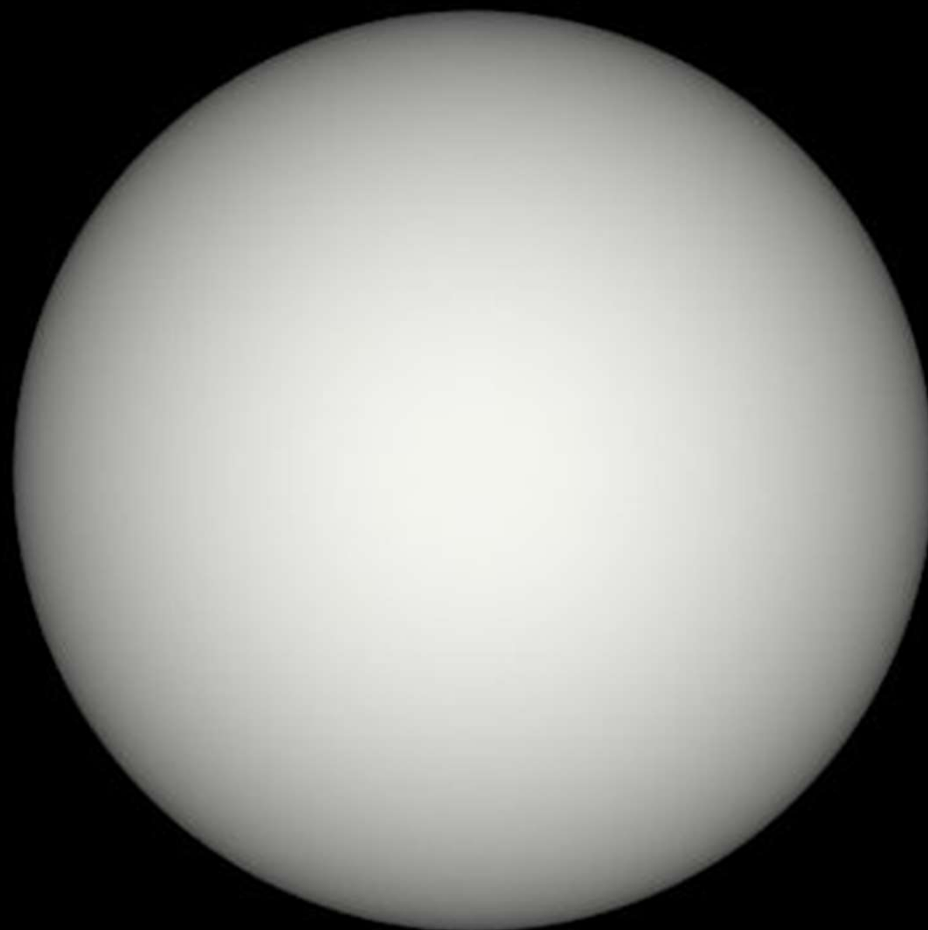

Is the moon Lambertian?

The moon "looks" matte but is not Lambertian

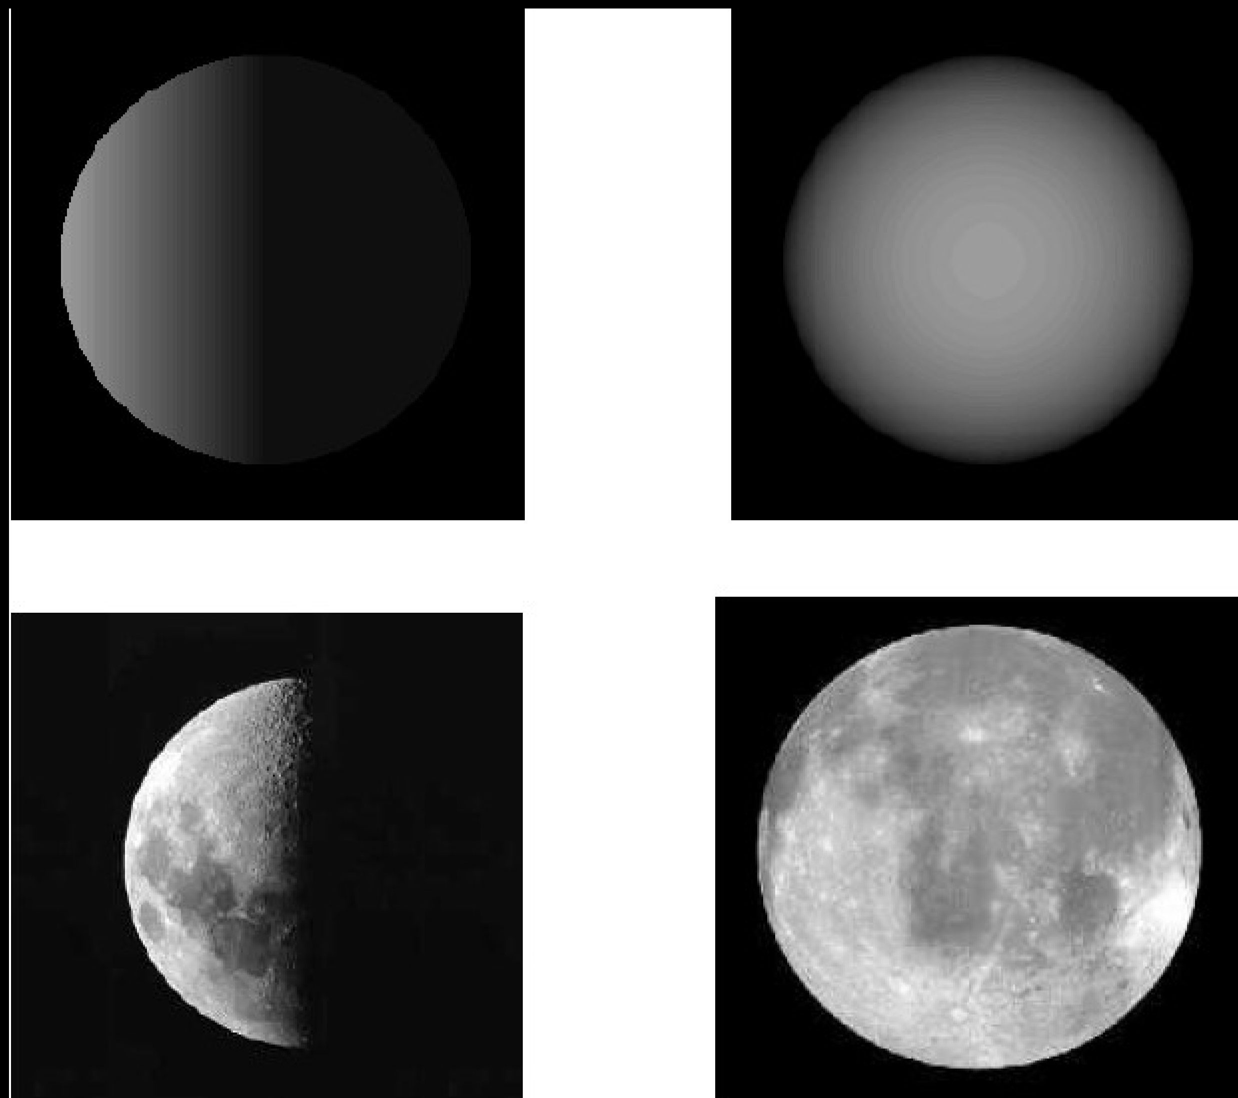

Local Shading Model

- Assume that the radiosity at a patch is the sum of the radiosities due to point light sources alone.

No interreflections.

- For point sources: $$B(P) = \sum_{\text{visible } s} \rho_d(P) \frac{\mathbf{N}(P) \cdot \mathbf{S}_s(P)}{R_s(P)^2}$$

- For point sources at infinity:

$$B(P) = \rho_d(P)\mathbf{N}(P) \cdot \sum_{\text{visible } s} \mathbf{S}_s(P)$$

Photometric Stereo (Woodham, 1979)

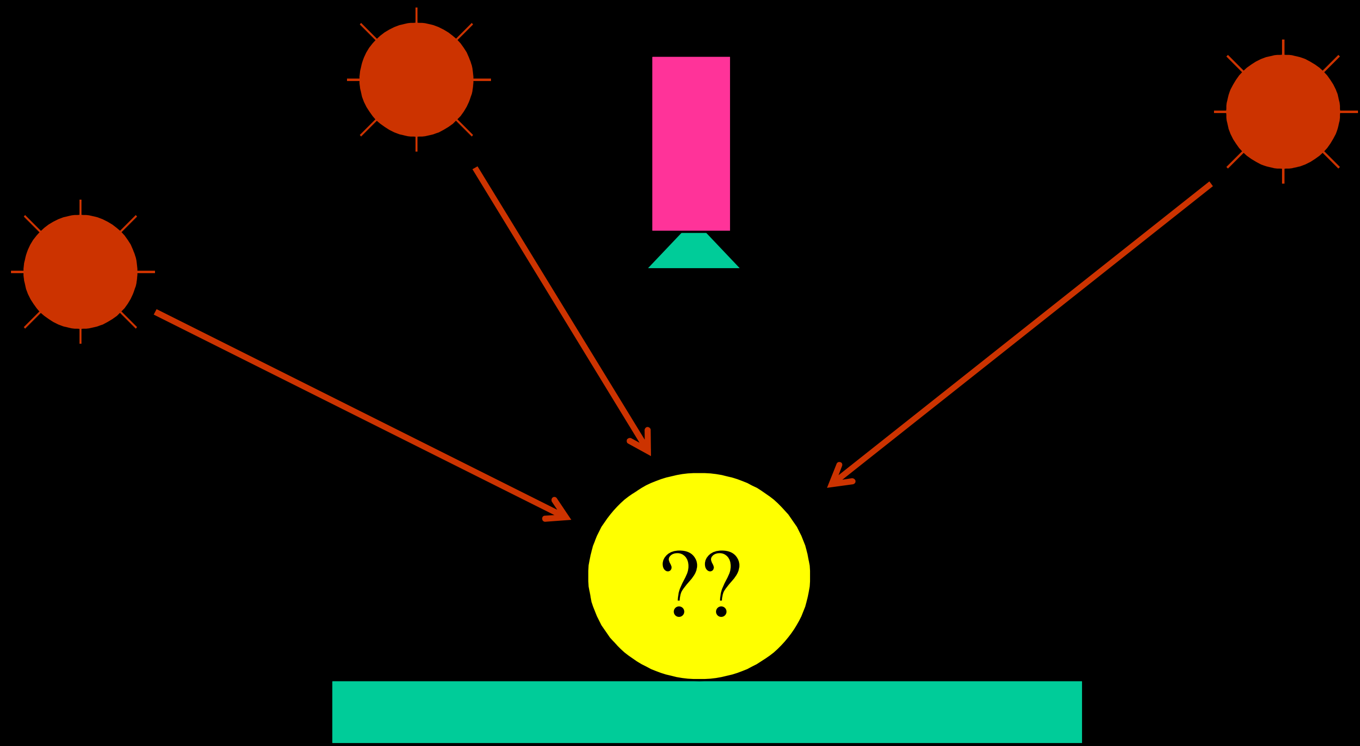

Problem: Given n images of an object, taken by a fixed camera under different (known) light sources, reconstruct the object shape.

$$I(P) = kB(P) = k\rho\, \mathbf{N}(P) \cdot \mathbf{S}$$

Photometric Stereo: Example (1)

- Assume a Lambertian surface and distant point light sources.

$$\boxed{I(P) = kB(P) = k\rho\, \mathbf{N}(P)\cdot \mathbf{S} = \mathbf{g}(P)\cdot \mathbf{V}}$$

with $\mathbf{g}(P)=\rho\, \mathbf{N}(P)$
and $\mathbf{V}= k\, \mathbf{S}$

- Given n images, we obtain n linear equations in ***g***:

$$\mathbf{i} = \begin{bmatrix} I_1 \\ I_2 \\ \ldots \\ I_n \end{bmatrix} = \begin{bmatrix} \mathbf{V}_1 \cdot \mathbf{g} \\ \mathbf{V}_2 \cdot \mathbf{g} \\ \ldots \\ \mathbf{V}_n \cdot \mathbf{g} \end{bmatrix} = \begin{bmatrix} \mathbf{V}_1^T \\ \mathbf{V}_2^T \\ \\ \mathbf{V}_n^T \end{bmatrix} \mathbf{g} \Rightarrow \mathbf{i} = \mathbf{V}\,\mathbf{g} \Rightarrow \boxed{\mathbf{g} = \mathbf{V}^{\dagger}\ \mathbf{i}}$$

Photometric Stereo: Example (2)

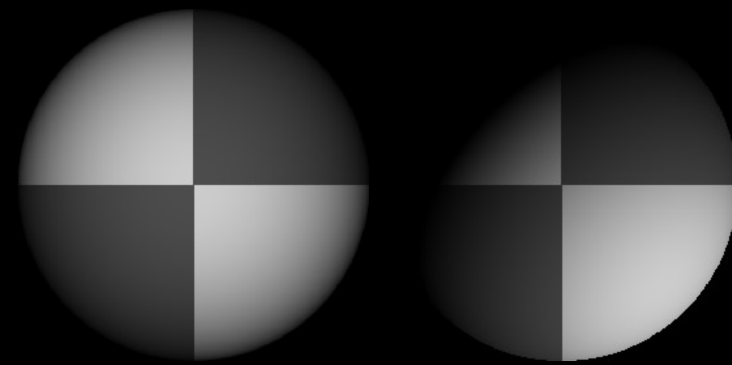

- What about shadows?
- Just skip the equations corresponding to zero-intensity pixels.
- Only works when there is no ambient illumination.

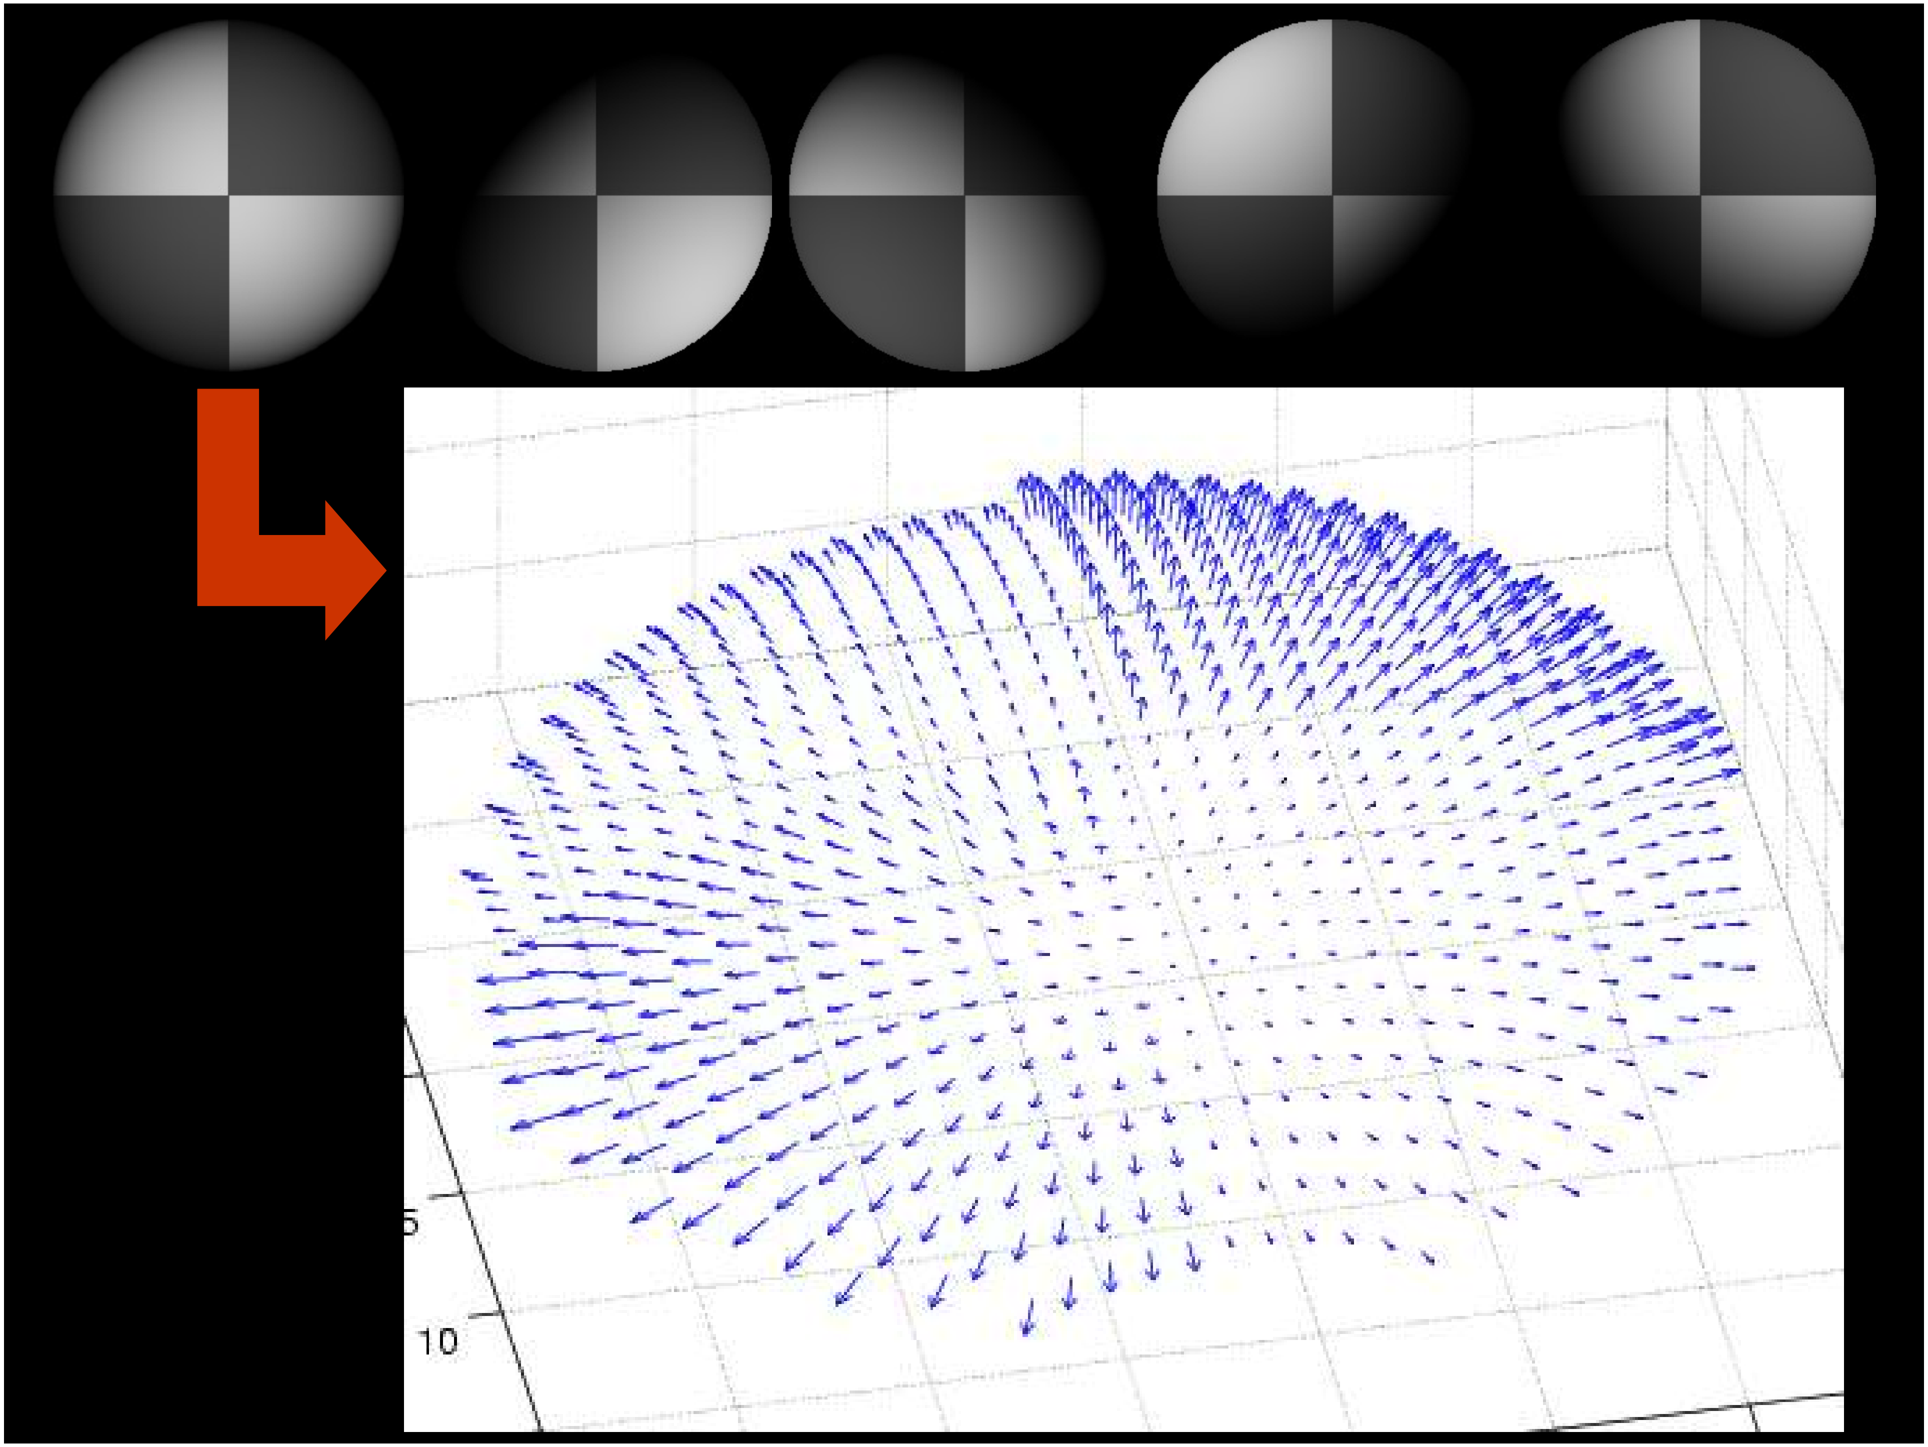
5
10

Photometric Stereo: Example (3)

$$\mathbf{g}(P)=\rho(P)\mathbf{N}(P)$$

$$\rho(P) = |\,\mathbf{g}(P)\,|$$

$$\boxed{\mathbf{N}(P) = \frac{1}{|\,\mathbf{g}(P)\,|}\mathbf{g}(P)}$$

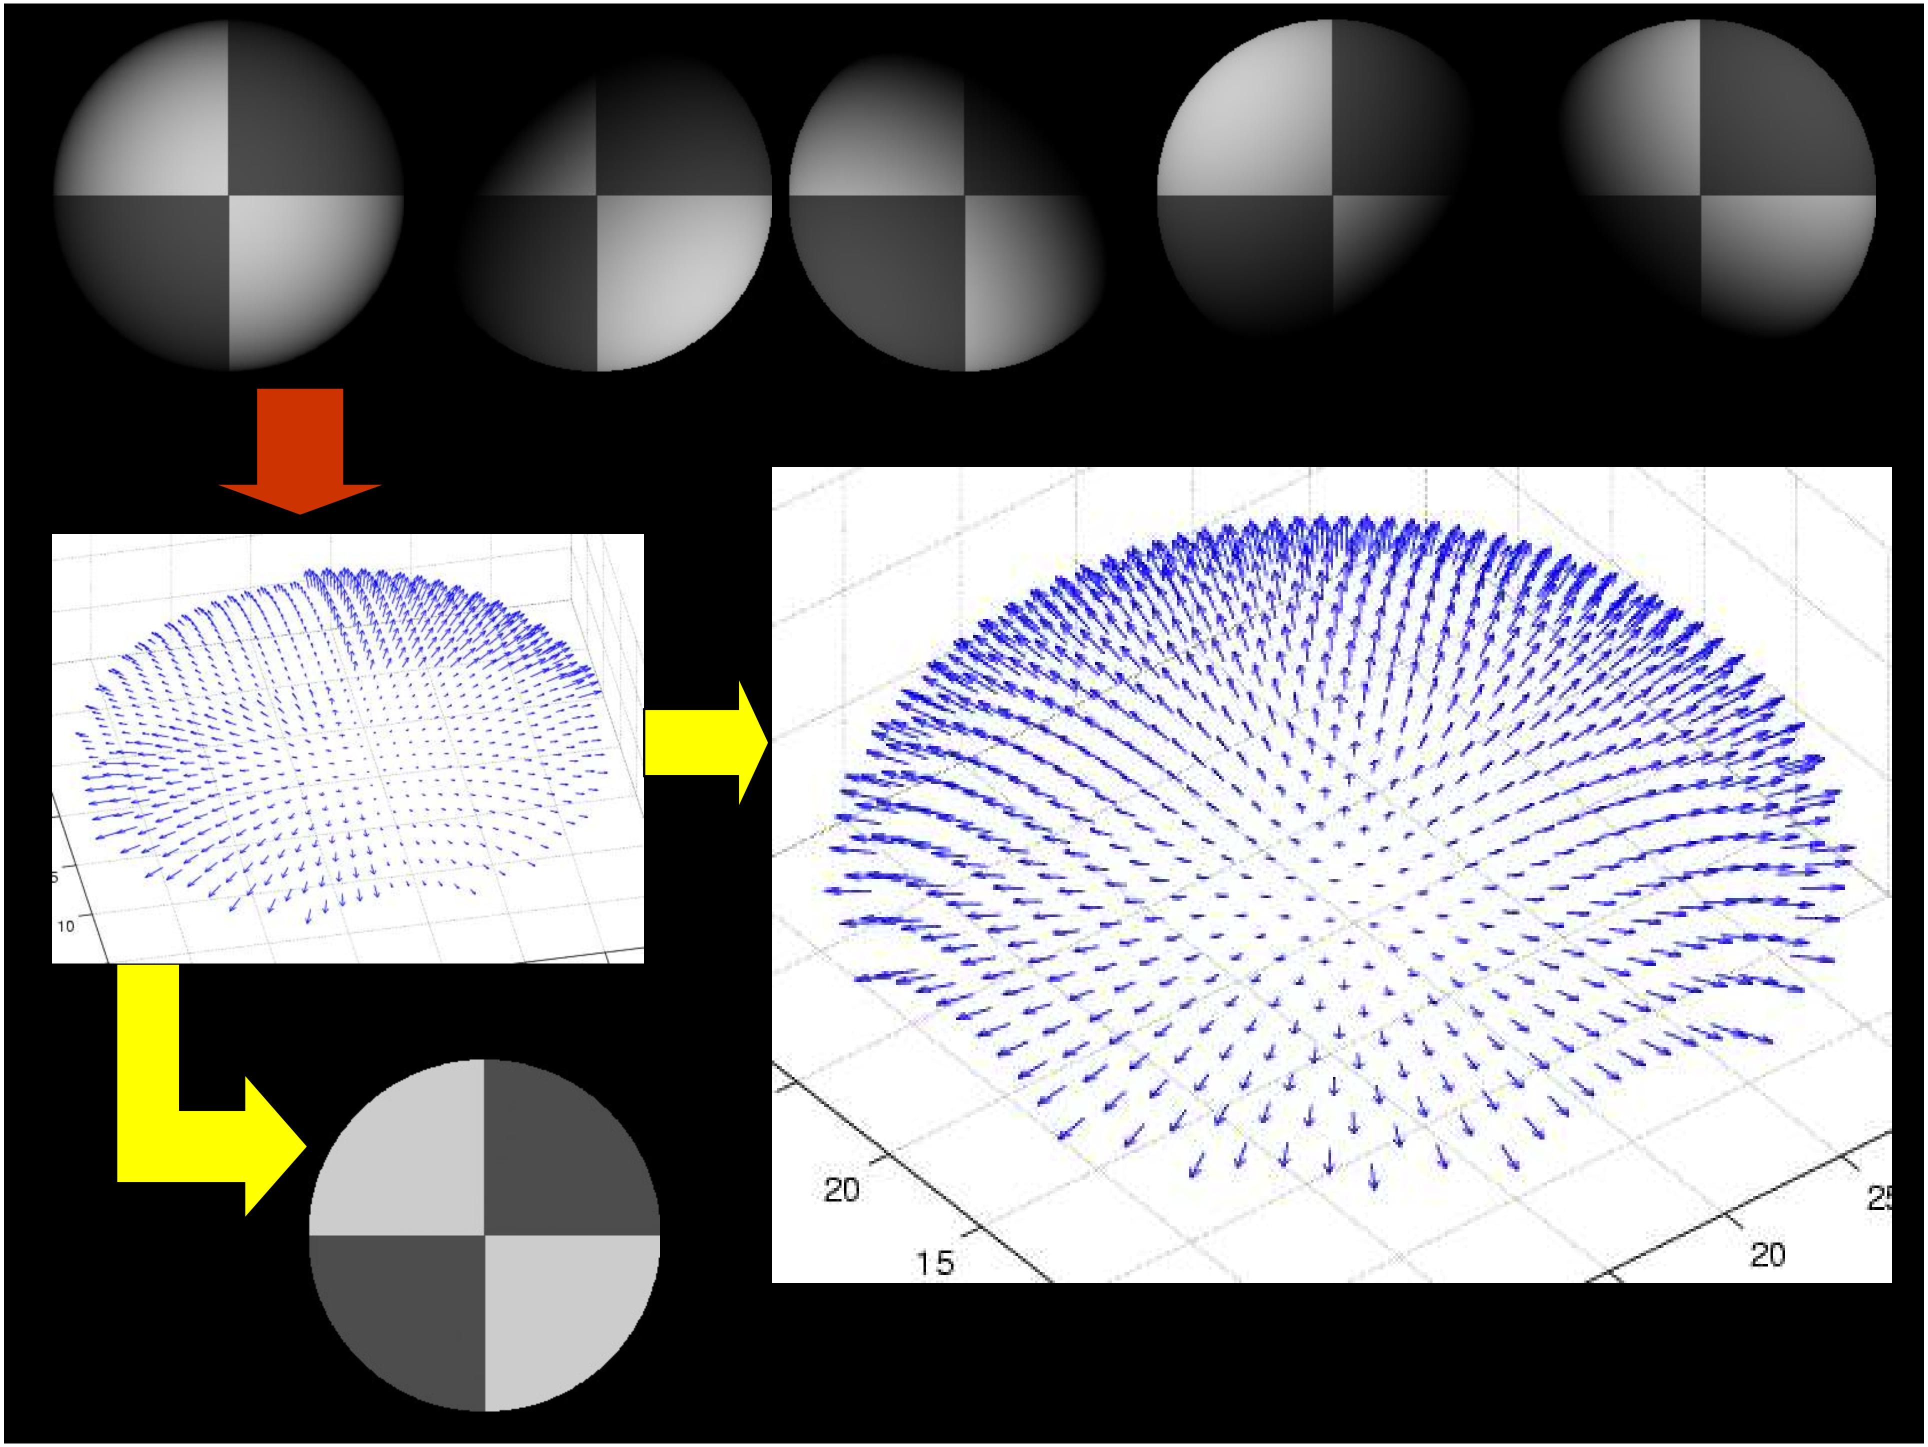
10
20
15
20

Photometric Stereo: Example (3)

Integrability!

$$\frac{\partial}{\partial y}\left(\frac{\partial z}{\partial x}\right) = \frac{\partial}{\partial x}\left(\frac{\partial z}{\partial y}\right)$$

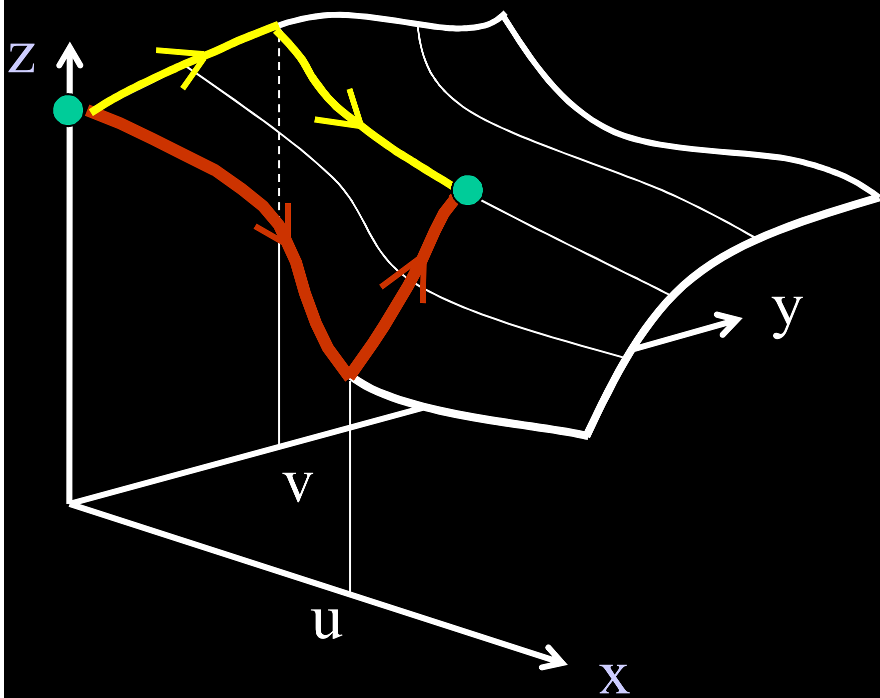

$$\mathbf{N} = \begin{bmatrix} a \\ b \\ c \end{bmatrix} \propto \begin{bmatrix} -\frac{\partial z}{\partial x} \\ -\frac{\partial z}{\partial y} \\ 1 \end{bmatrix} \Rightarrow \begin{cases} \frac{\partial z}{\partial x} = -\frac{a}{c} = -p \\ \frac{\partial z}{\partial y} = -\frac{b}{c} = -q \end{cases}$$

$$z(u,v) = \int_0^u \frac{\partial z}{\partial x}(x,0)\,dx + \int_0^v \frac{\partial z}{\partial y}(u,y)\,dy$$

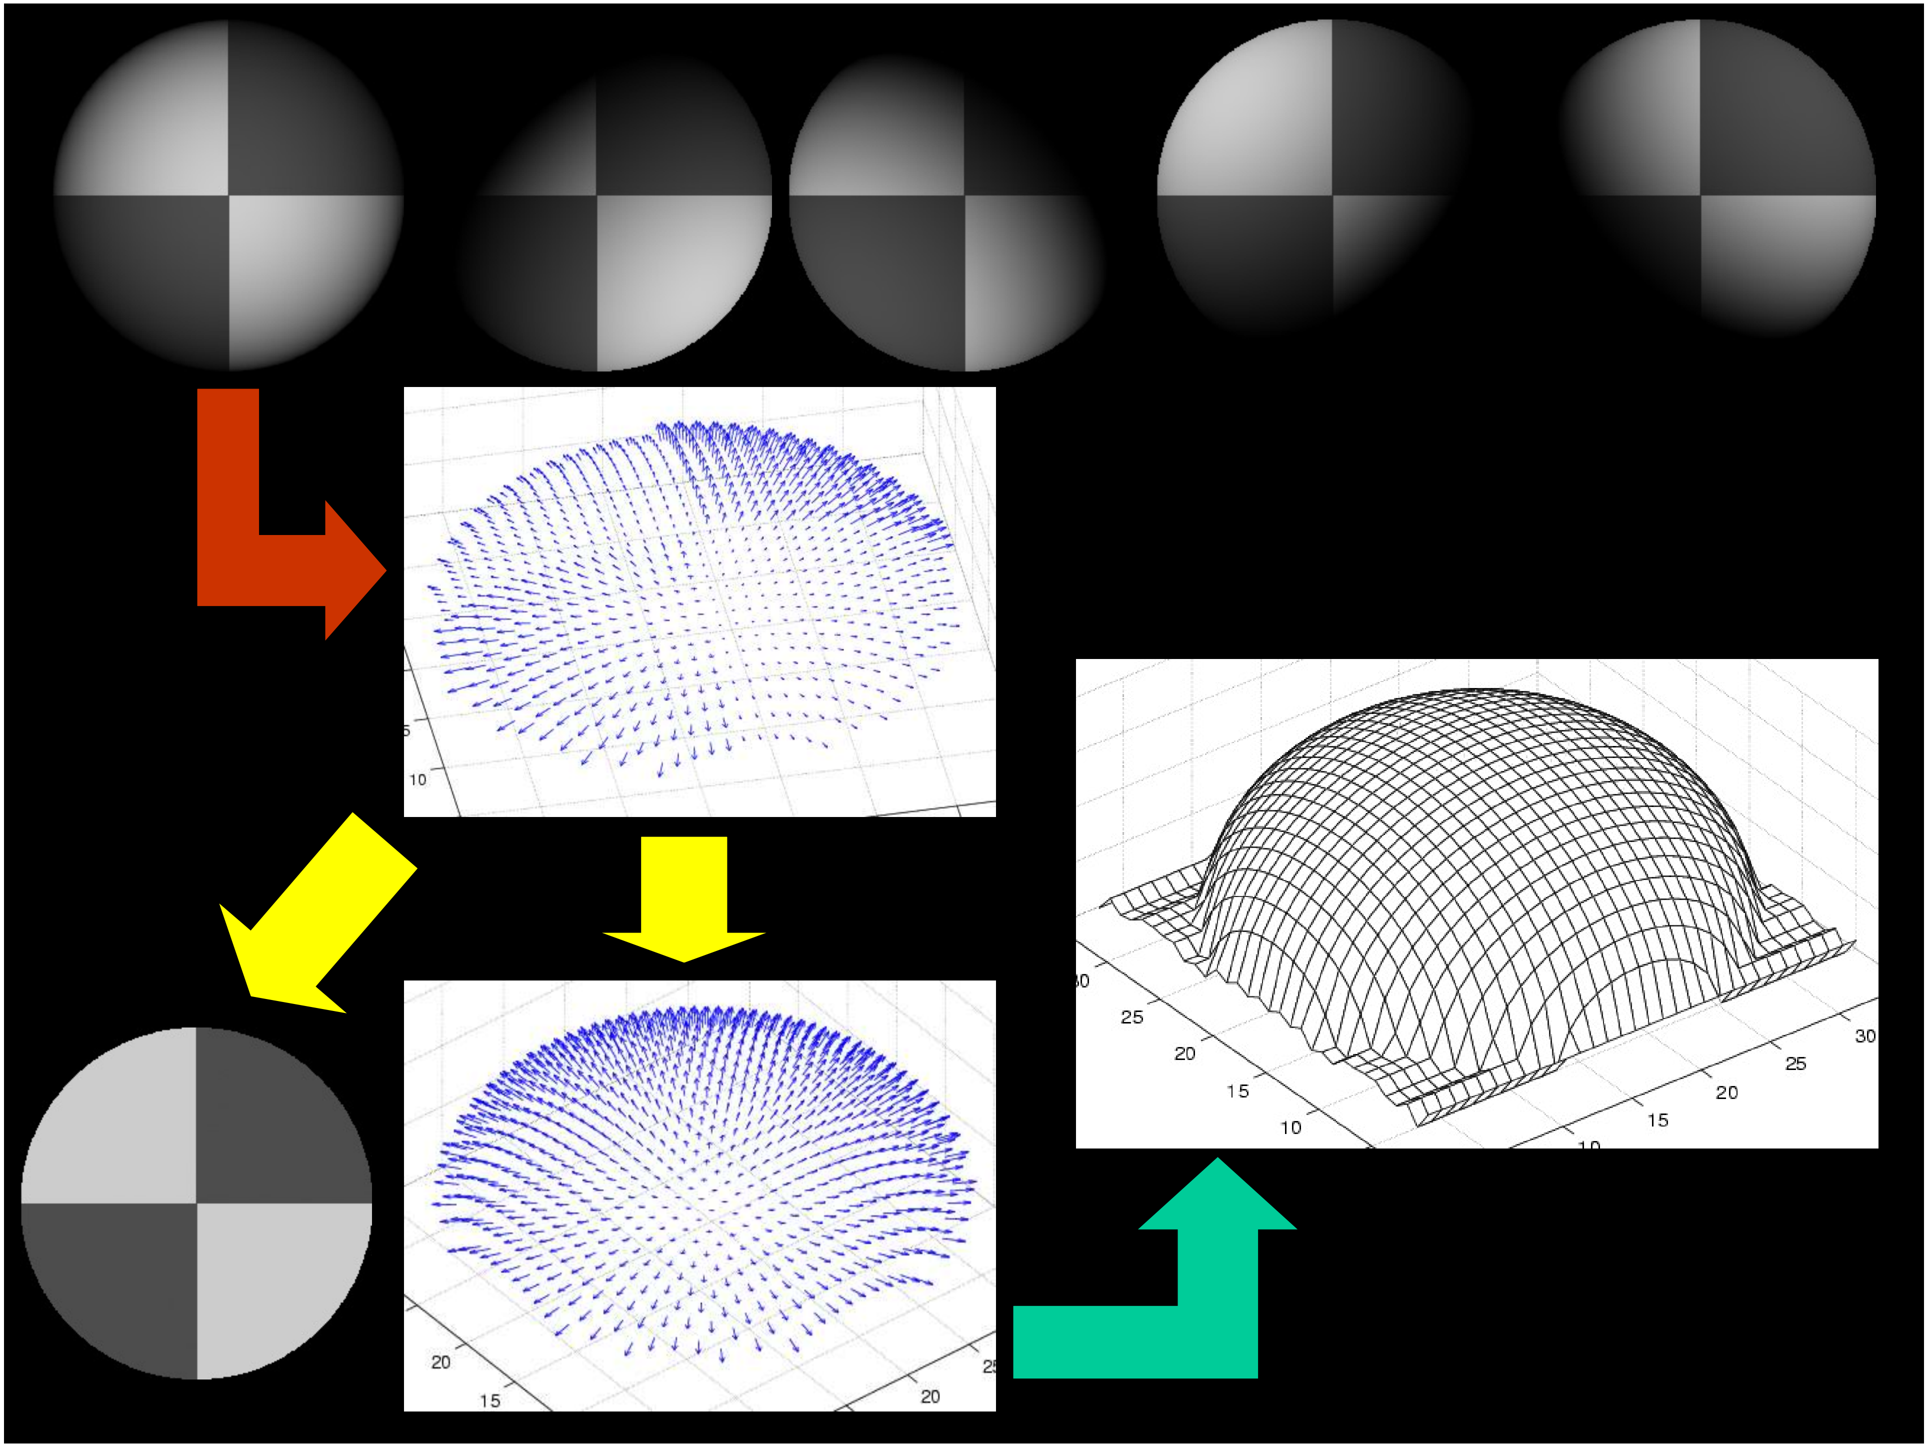
10
20
15
20
25
20
15
10
30
25
20
15

Reflectance map: known reflectance and light source

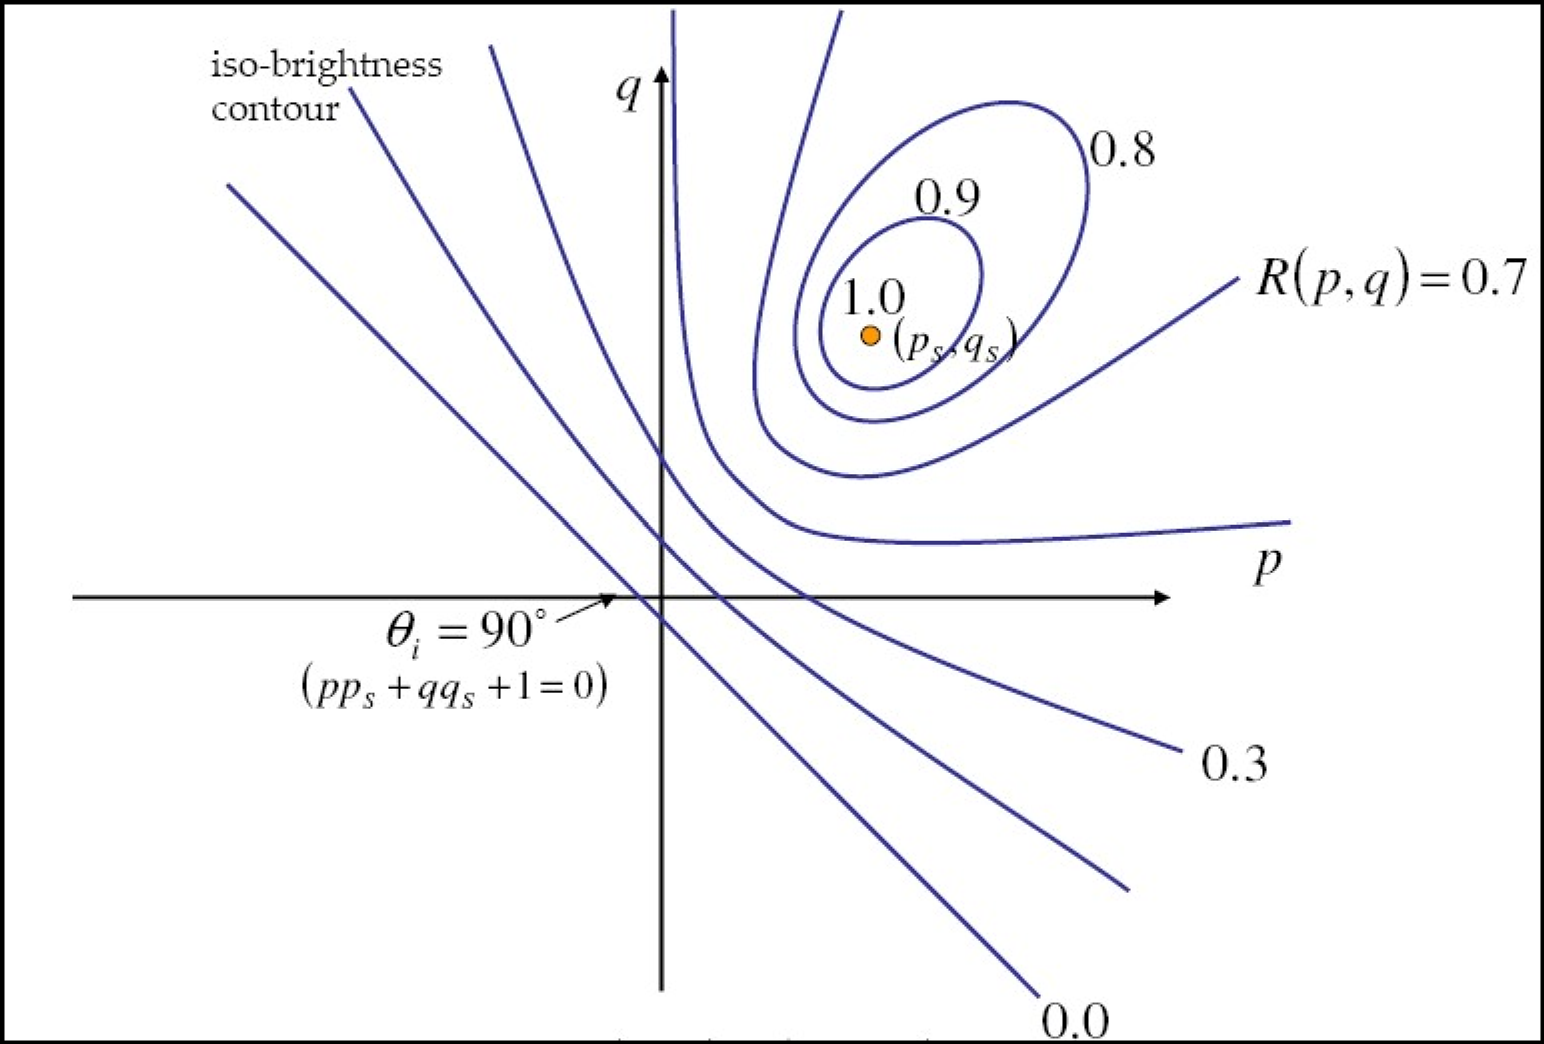

[Here: Lambertian case.]

- Given a single image of an object with known surface reflectance taken under a known light source, can we recover the shape of the object?
- Given *R(p,q)* (*(p_S,q_S)* and surface reflectance) can we determine *(p,q)* uniquely for each image point?

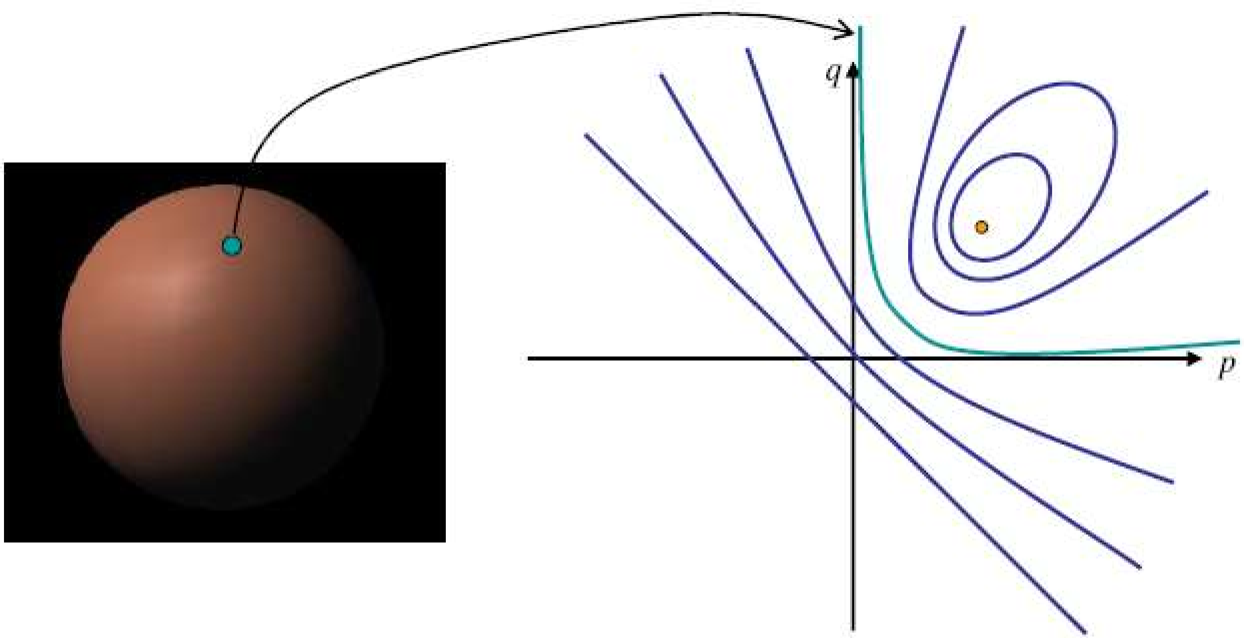

NO

Using more images: photometric stereo

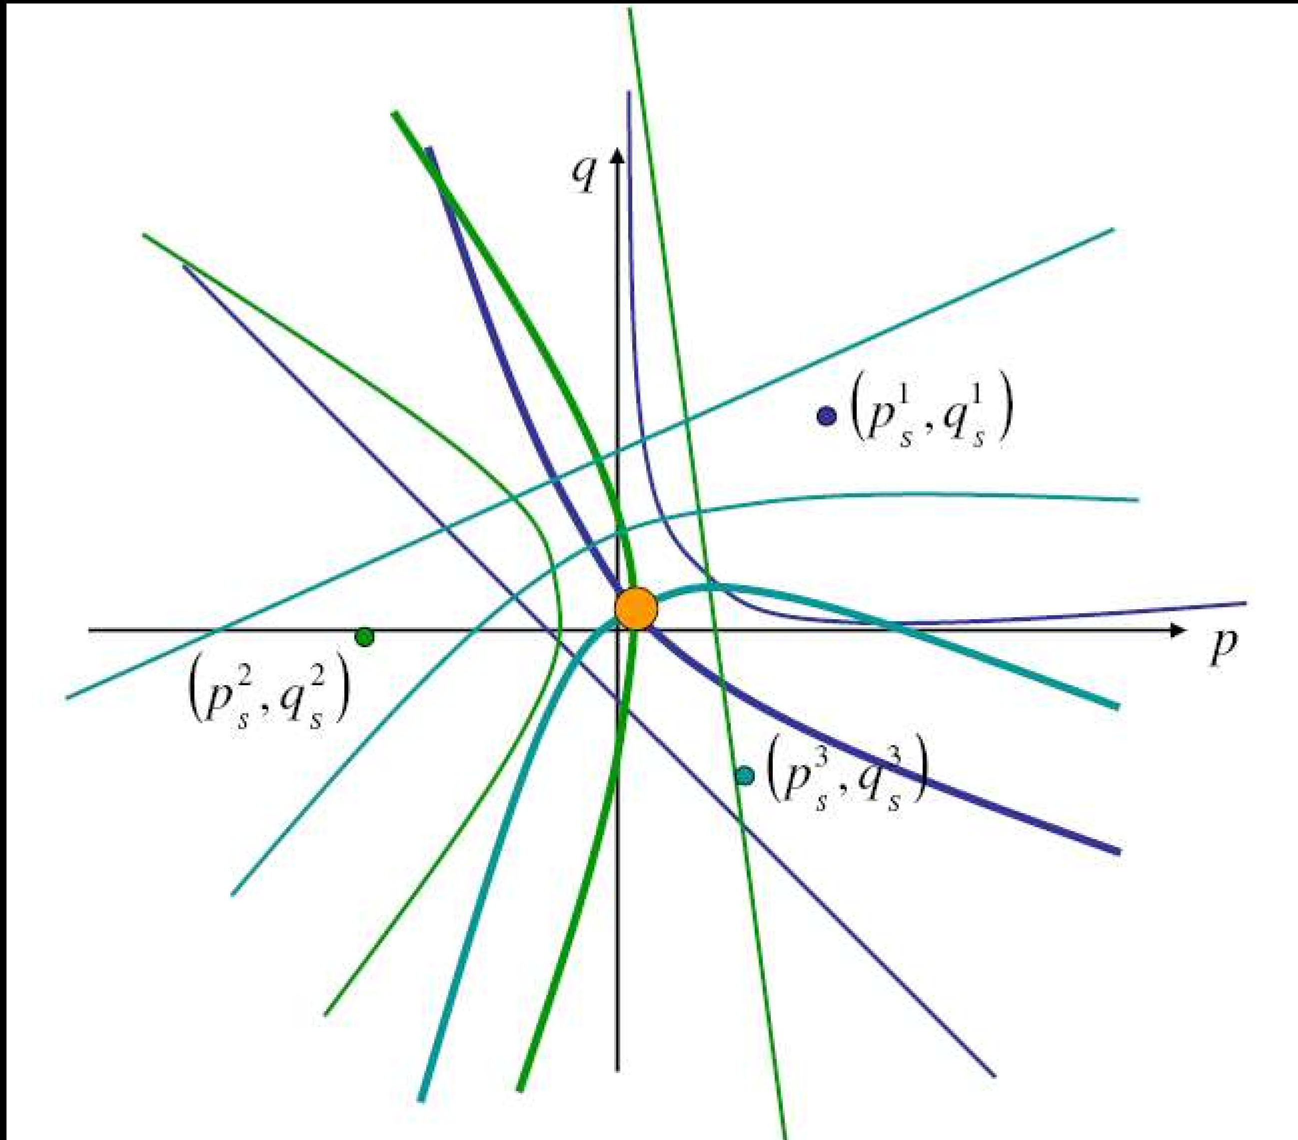

Example-based photometric stereo

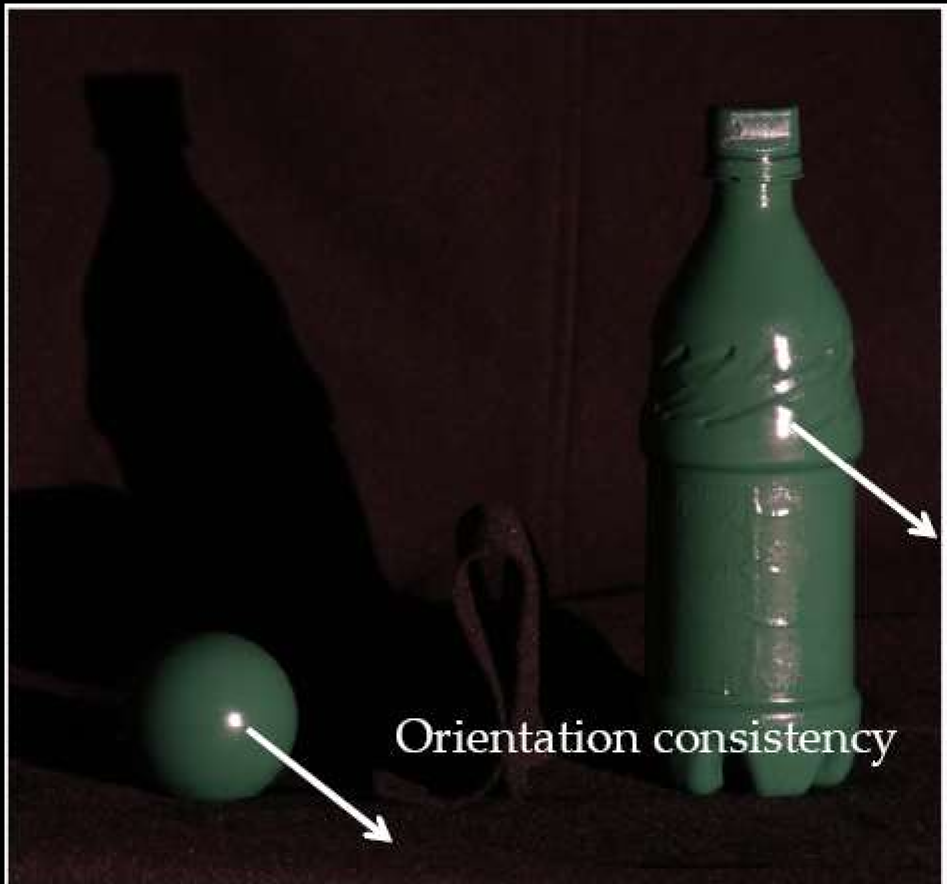

[Silver, 1980; Woodham, 1982, Seitz & Hertzmann, 2003]

Problem: Interreflections

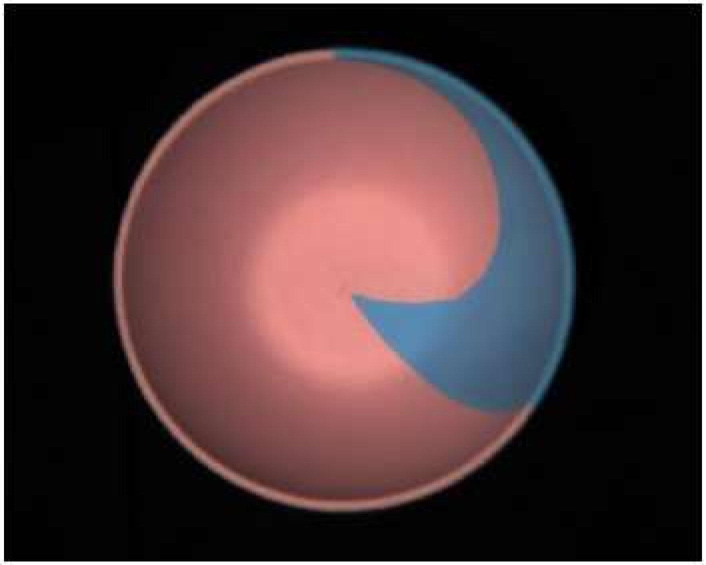

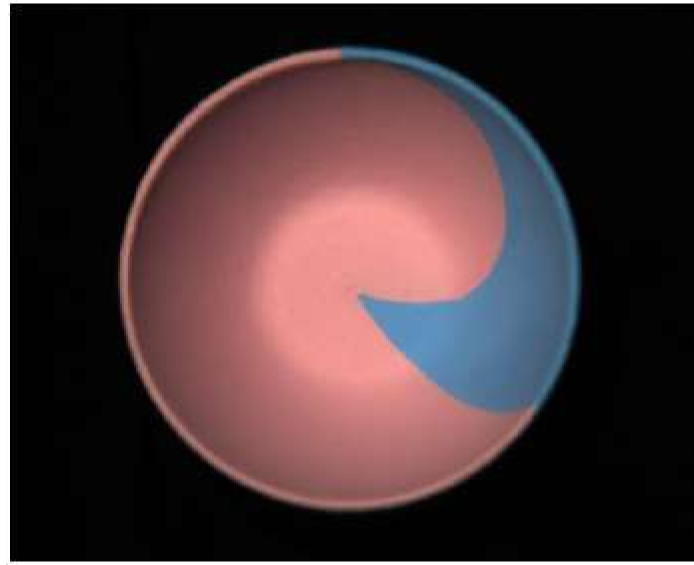

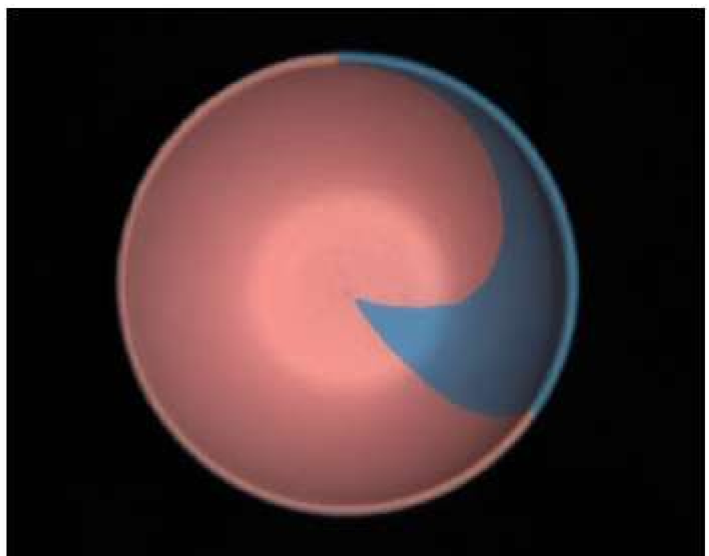

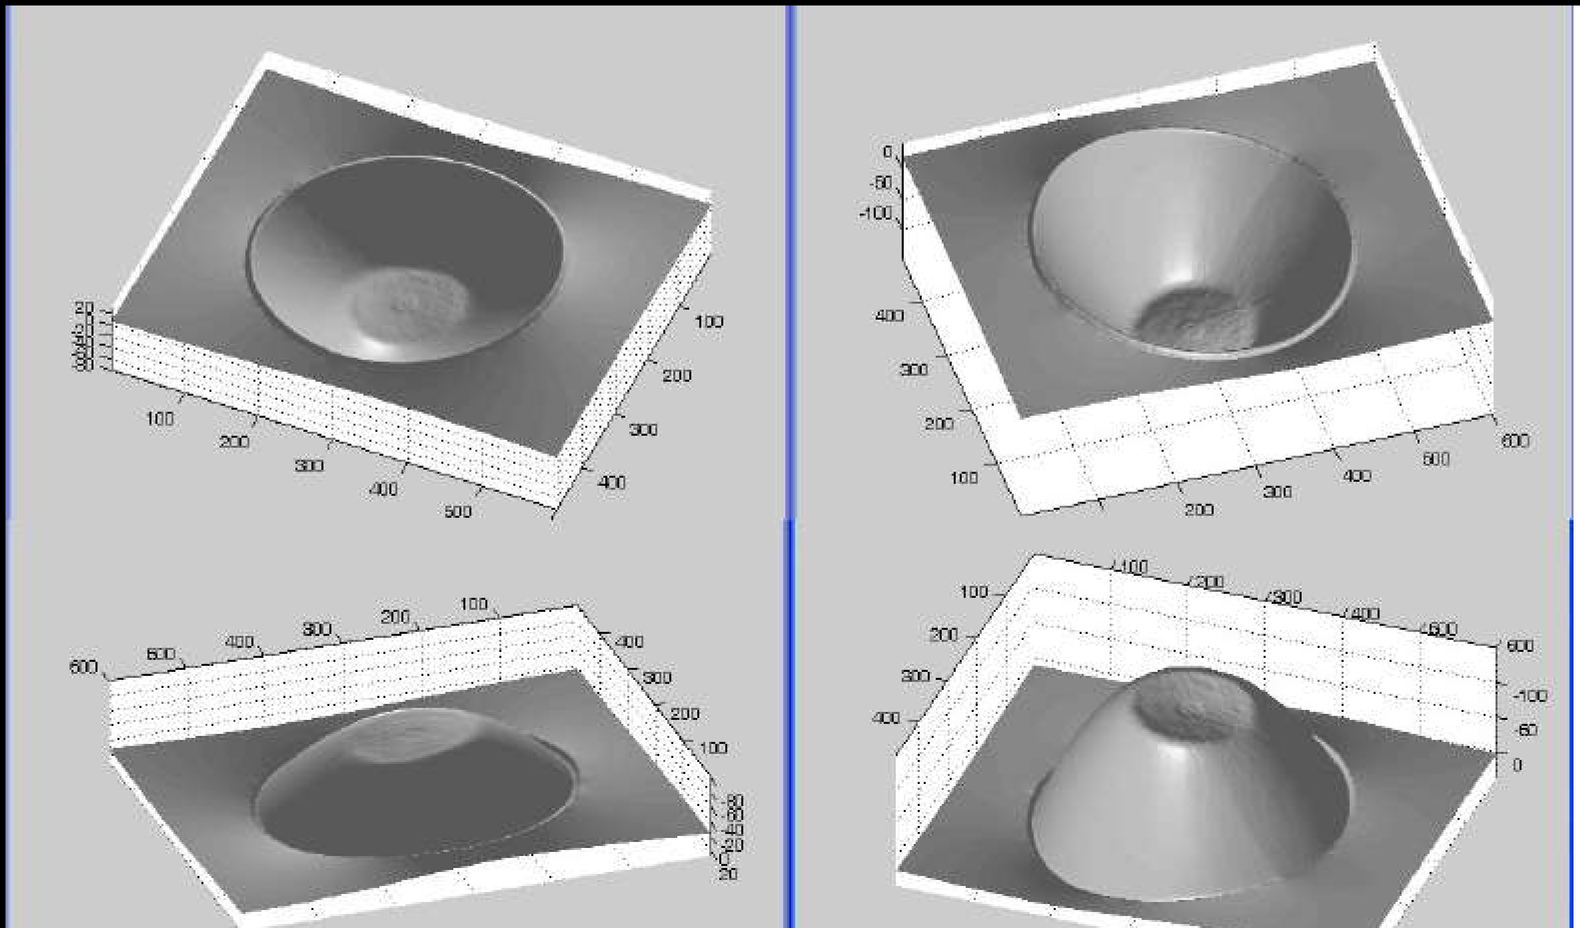

Shallow reconstruction
(effect of interreflections)

Accurate reconstruction
(after removing interreflections)

COLOR: Light energy usually depends on wavelength

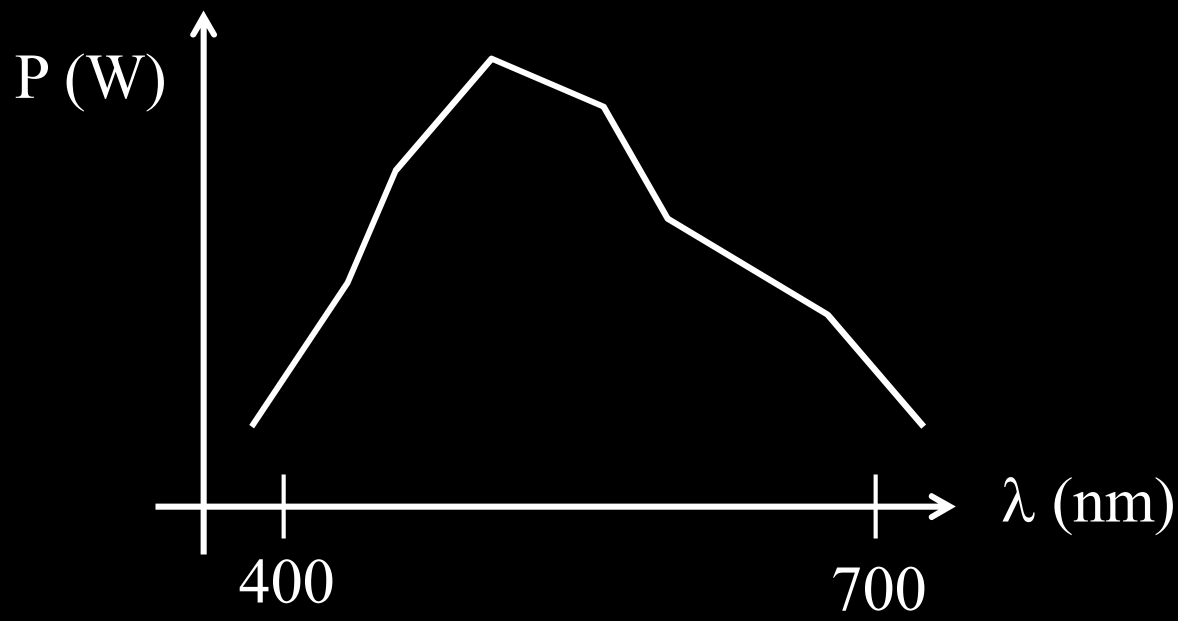

What causes the apparent color of a surface?

- the color and geometry of the light sources,
- the color and geometry of the surface,
- the spectral characteristics and geometry of the sensor.

Spectral Units

- Radiance L(P,v) – Spectral Radiance L^(P,v)

- The spectral radiance is the power traveling at some point in a given direction per unit area perpendicular to the direction of travel, per unit angle and per unit Wavelength (W $m^{-2}sr^{-1}$).

- The power emitted in the range $[\lambda,\lambda+\delta\lambda]$ is

$$\boxed{\delta P = L^{\lambda}(P,v) \cos\theta \ \delta\omega \ \delta A \ \delta\lambda.}$$

- Irradiance -> Spectral Irradiance
- BRDF -> Spectral BRDF
- Etc.

Sources

- Black body (Planck's Law):

$$E(\lambda)=\frac{1}{\lambda^5}\frac{1}{\exp(hc/\lambda kT)-1}$$

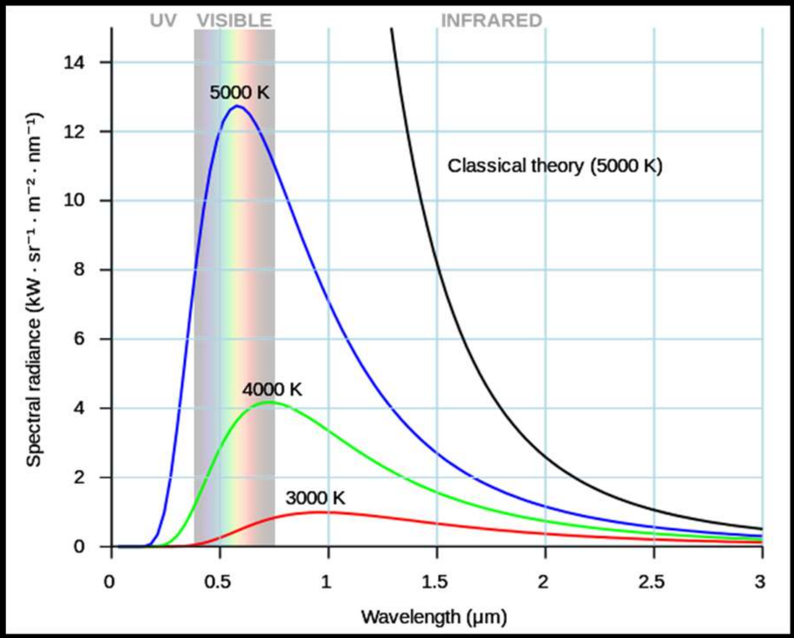

- Sun (daylight) and sky (airlight) [sky brighter at horizon]
- Incandescent lights
- Fluorescent lights
- LEDs
- Sodium or mercury arc lamps

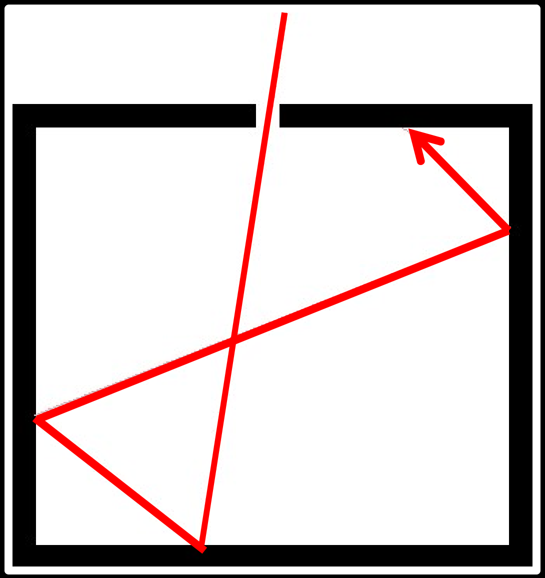

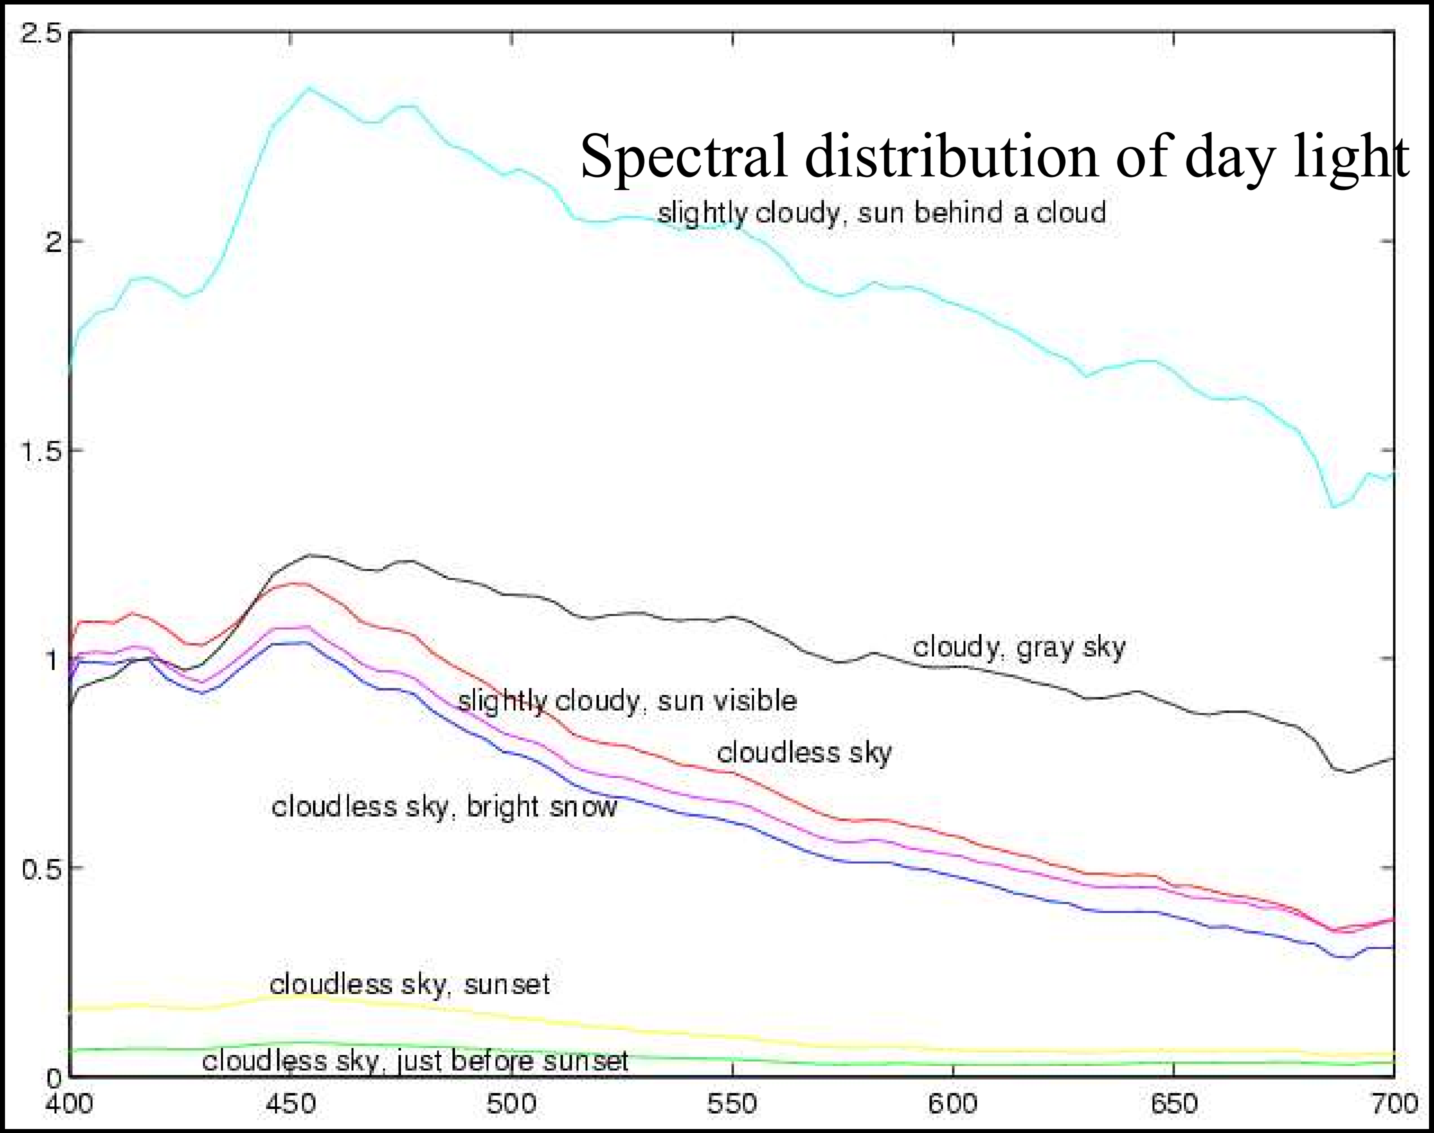

The sun looks yellow because shorter-wavelength frequencies (blue) are scattered out of sun rays faster than longer-wavelength frequencies (red) -> yellow sun and blue sky.

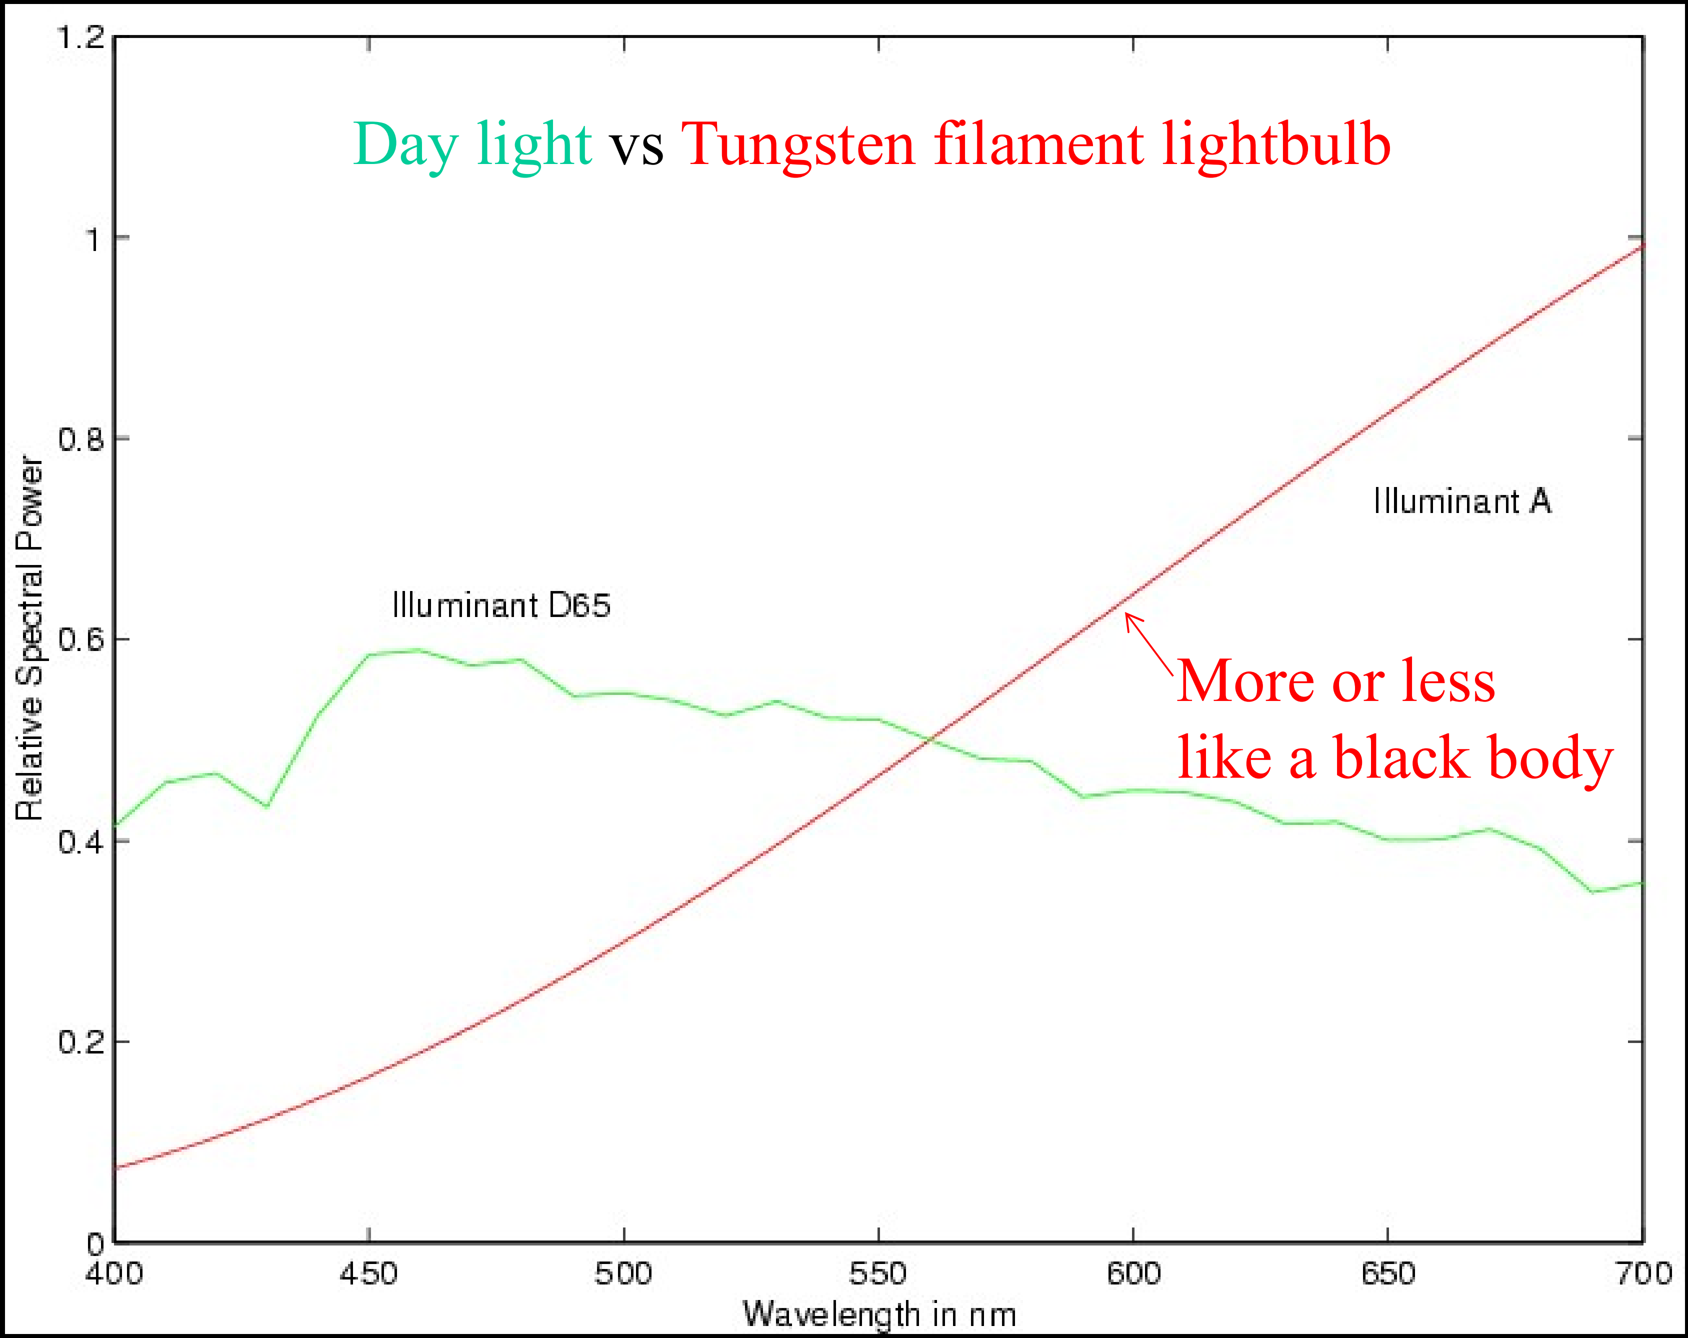
Day light vs Tungsten filament lightbulb
Illuminant D65
Illuminant A
More or less
like a black body
Relative Spectral Power
Wavelength in nm
0
0.2
0.4
0.6
0.8
1
1.2
400
450
500
550
600
650
700

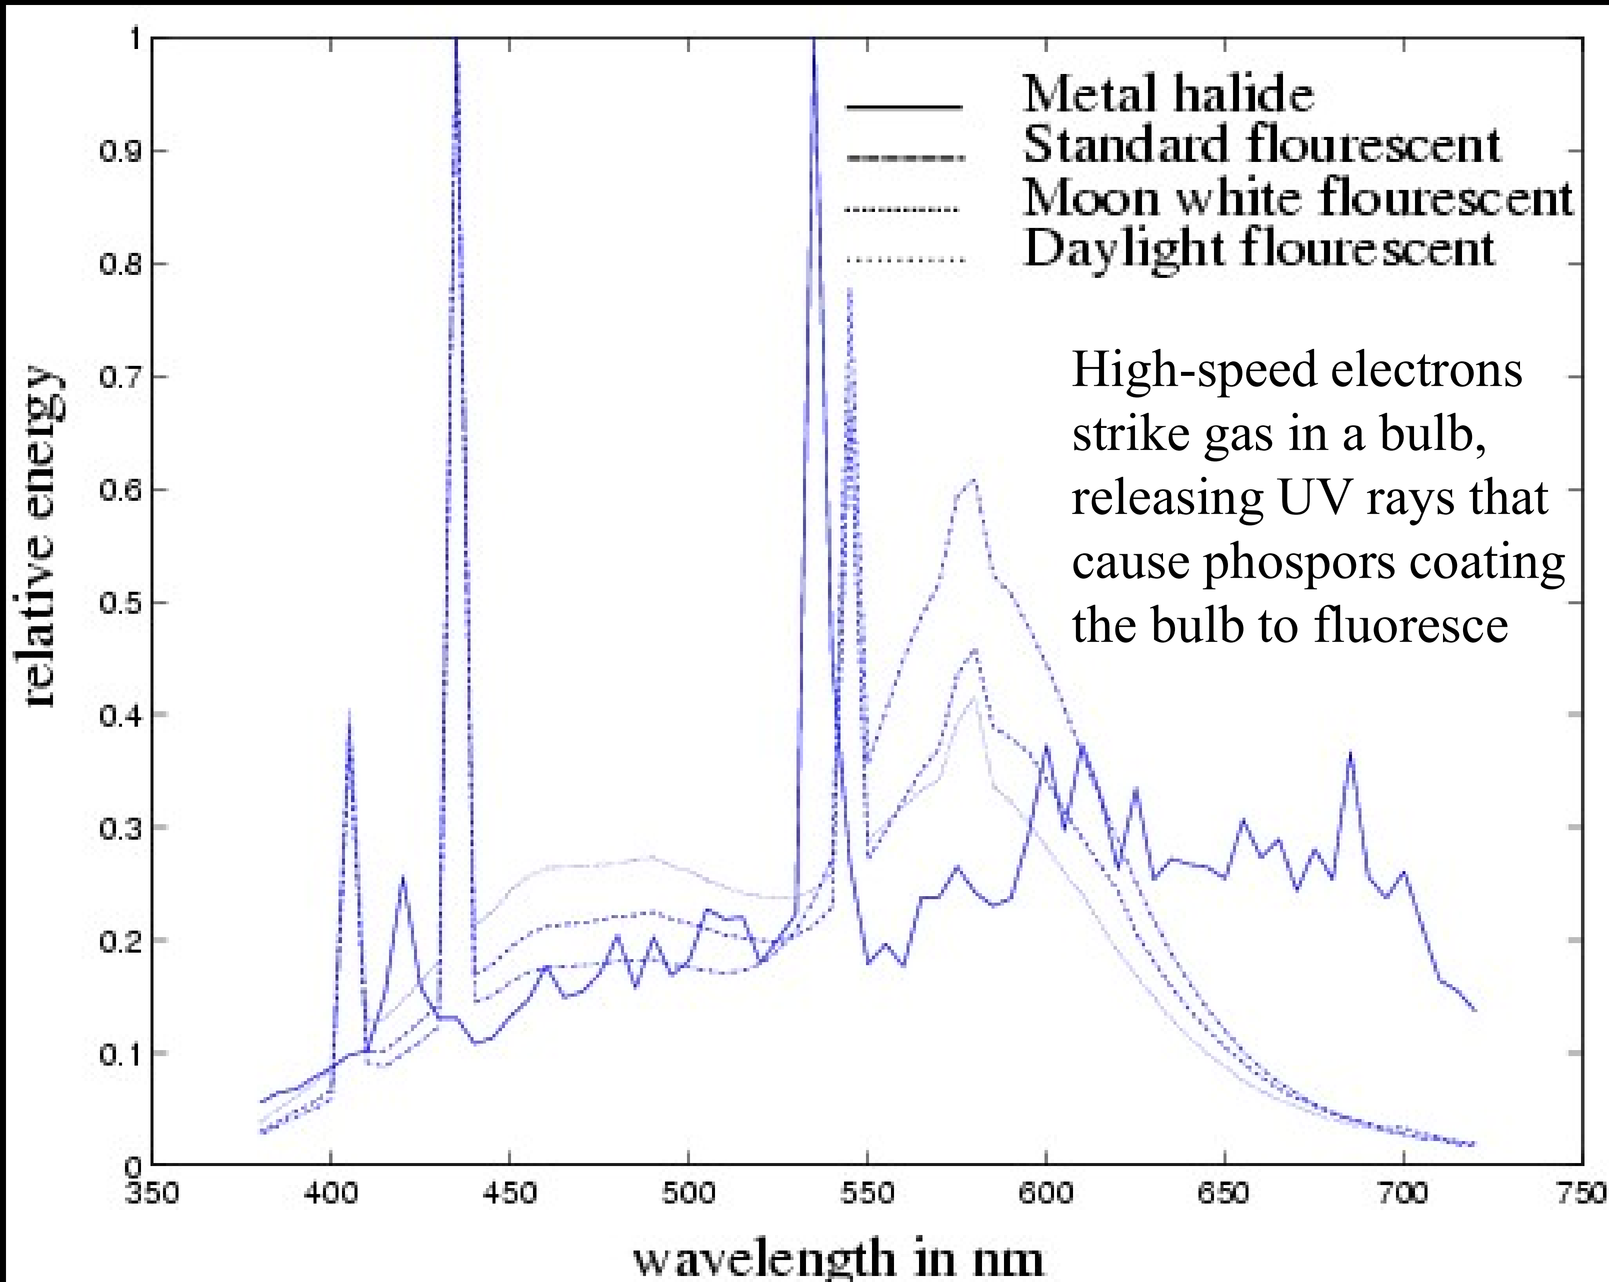
Metal halide
Standard flourescent
Moon white flourescent
Daylight flourescent
High-speed electrons strike gas in a bulb, releasing UV rays that cause phospors coating the bulb to fluoresce
relative energy
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
350
400
450
500
550
600
650
700
750
wavelength in nm

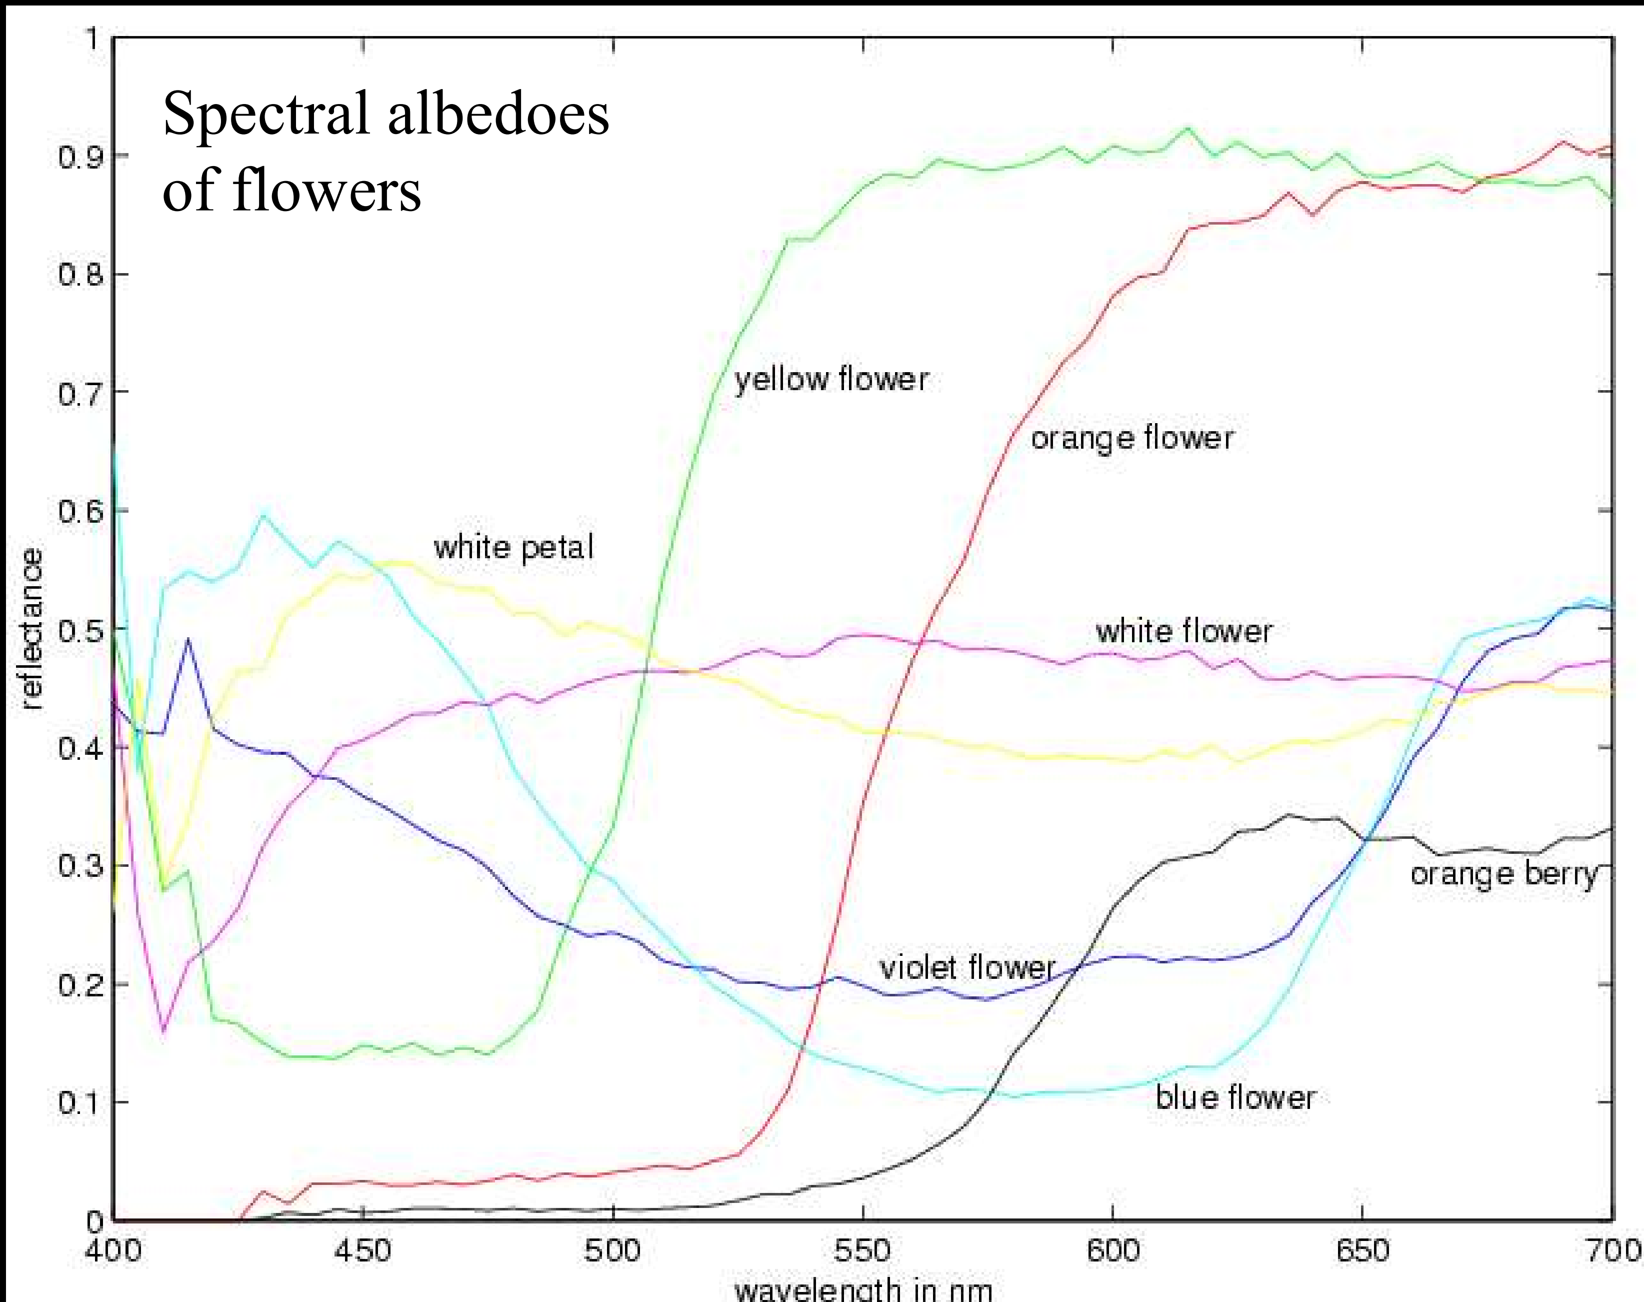
Spectral albedoes
of flowers
yellow flower
orange flower
white petal
white flower
orange berry
violet flower
blue flower
reflectance
wavelength in nm
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
400
450
500
550
600
650
700

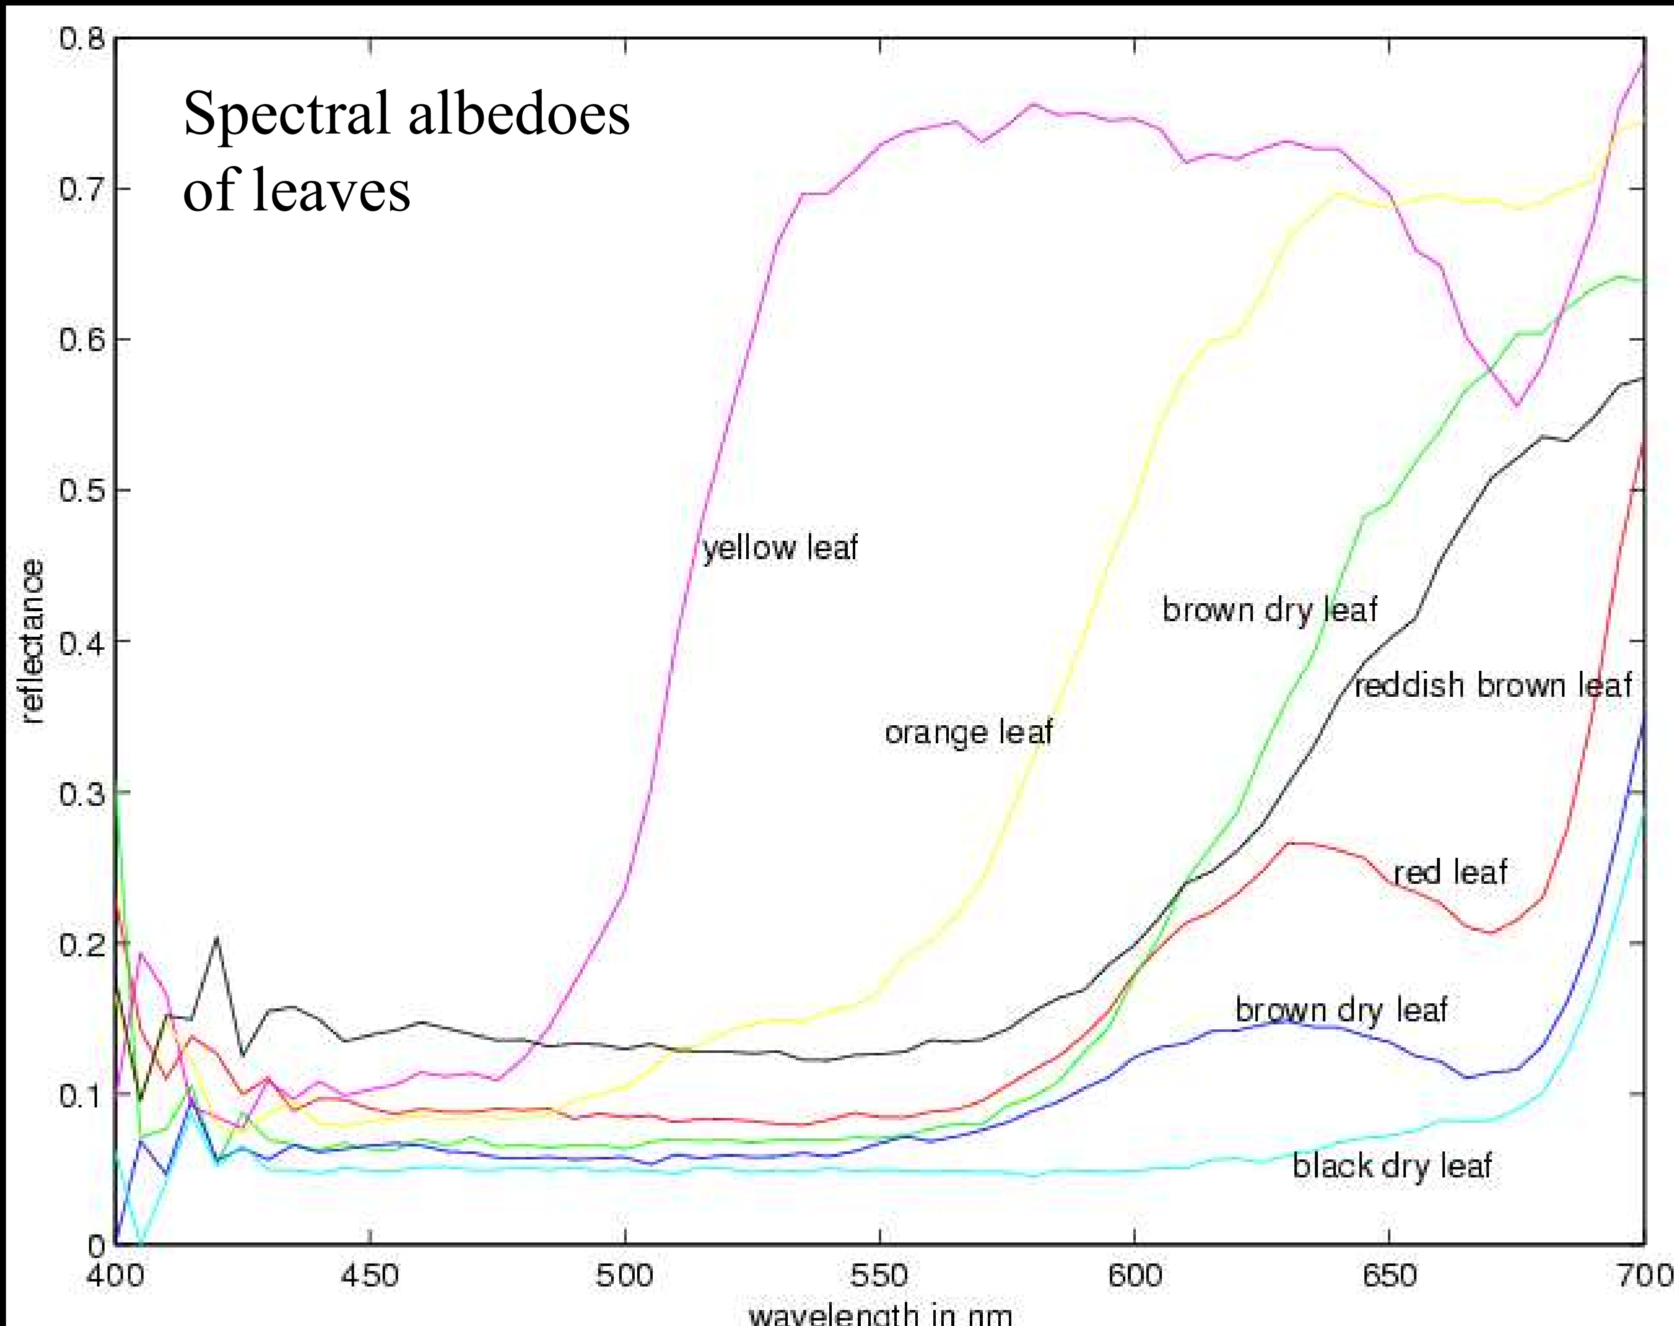
Spectral albedoes
of leaves
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
reflectance
400
450
500
550
600
650
700
wavelength in nm
yellow leaf
orange leaf
brown dry leaf
reddish brown leaf
red leaf
brown dry leaf
black dry leaf

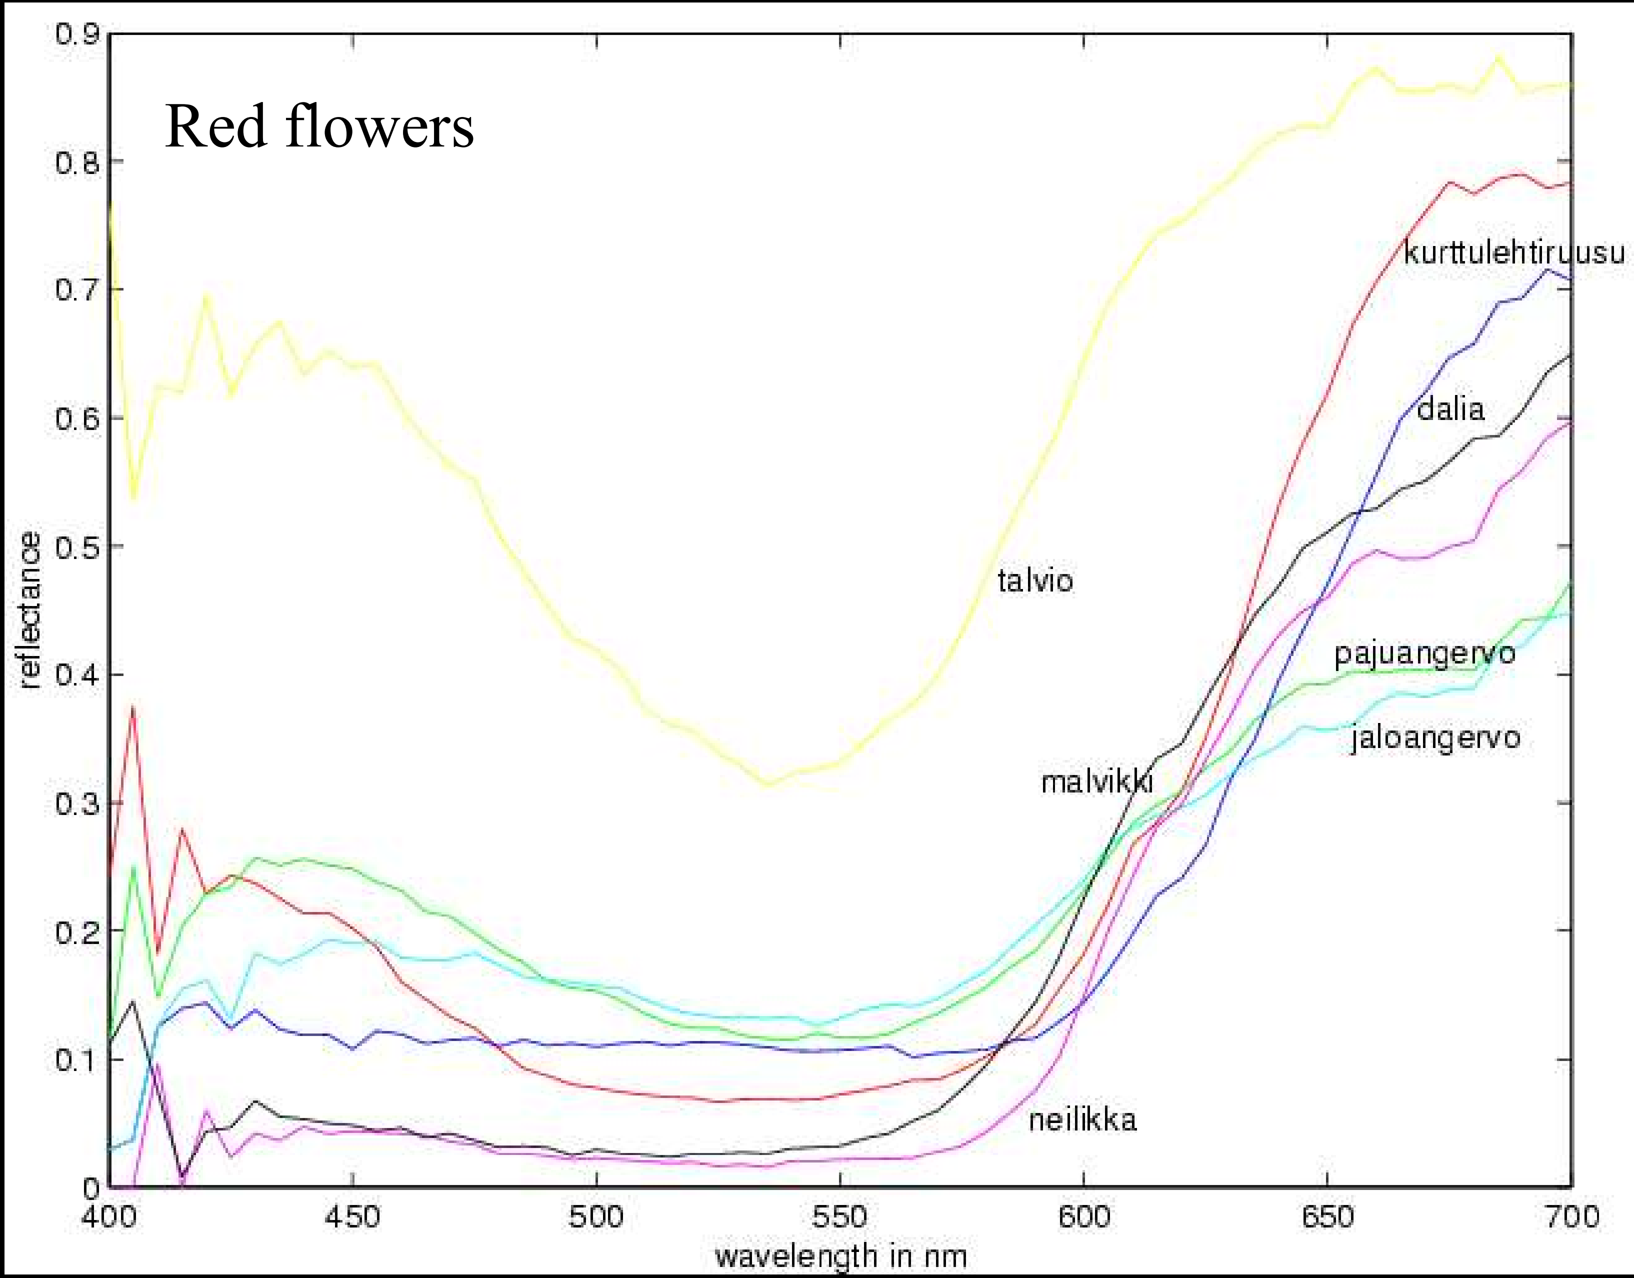

Red flowers
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
reflectance
400
450
500
550
600
650
700
wavelength in nm
kurttulehtiruusu
dalia
talvio
pajuangervo
jaloangervo
malvikki
neilikka

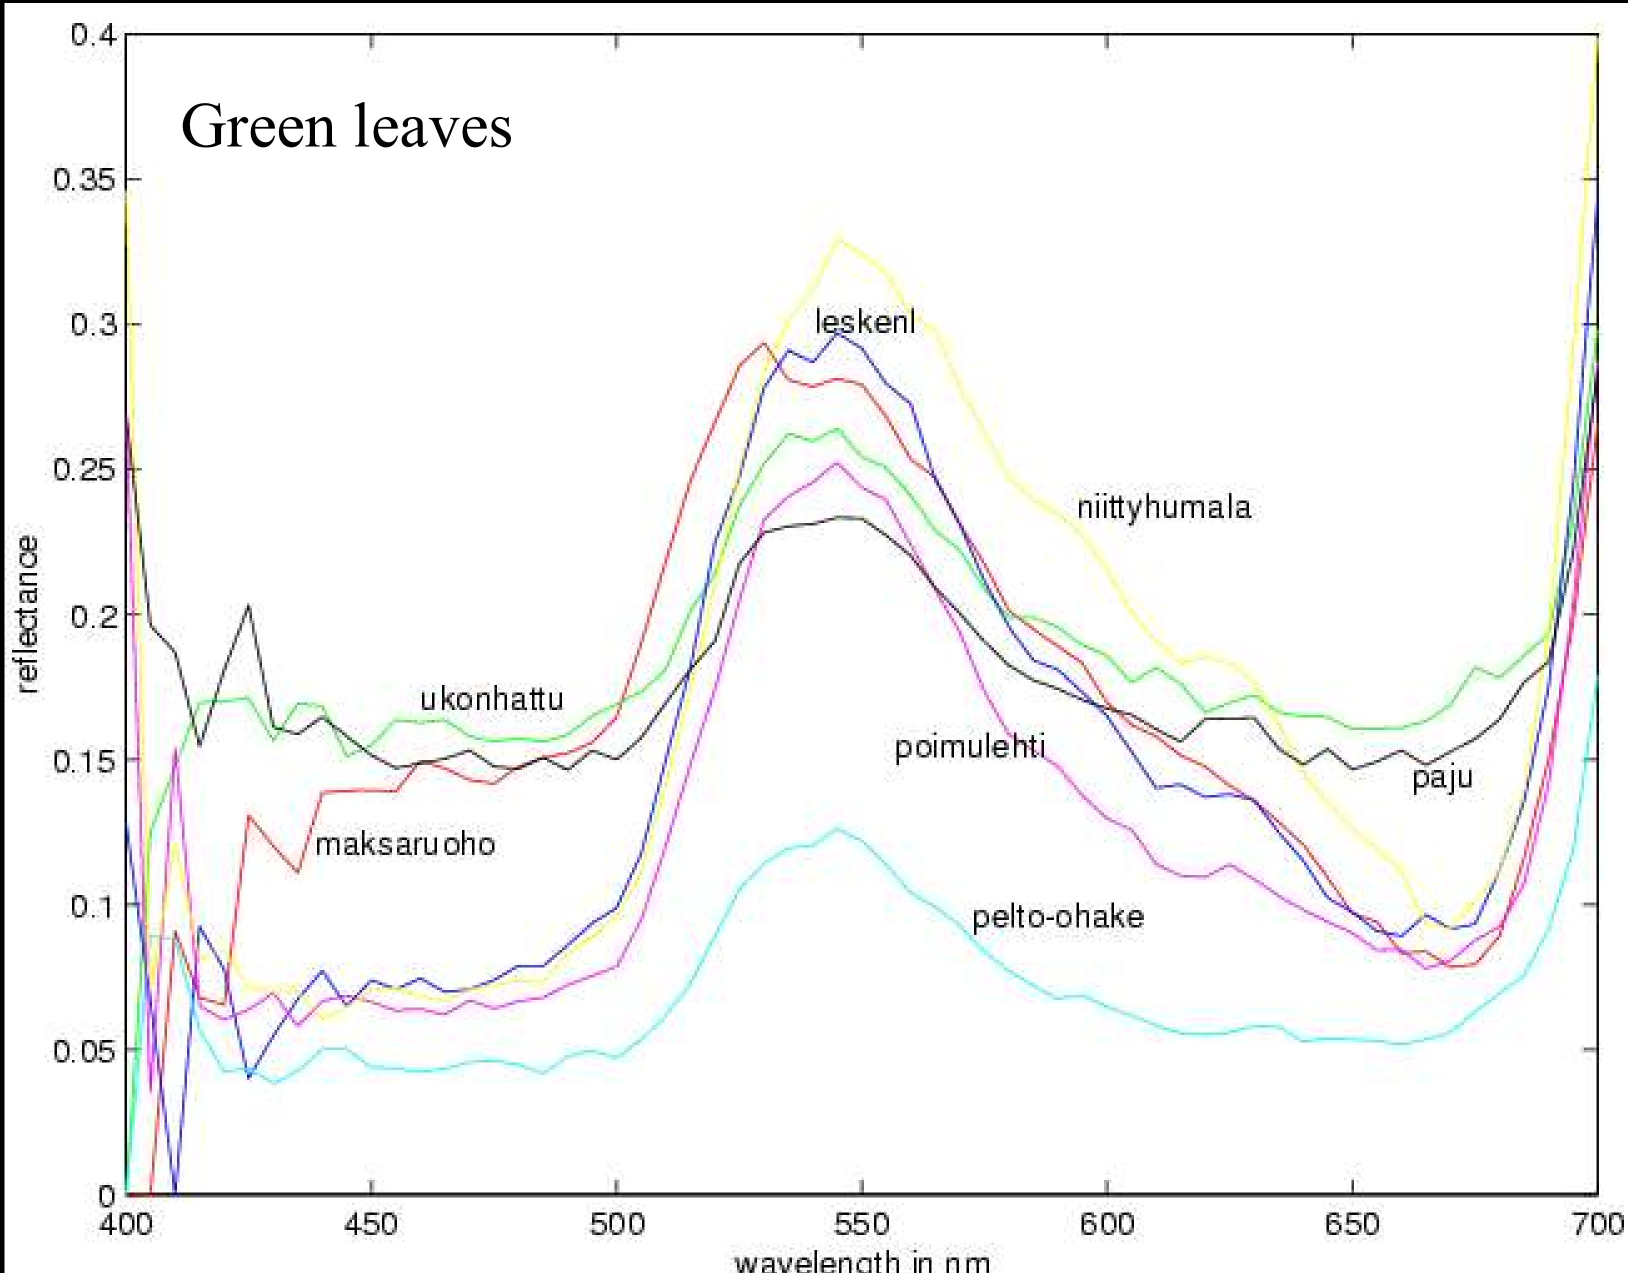
Green leaves
0.4
0.35
0.3
0.25
0.2
0.15
0.1
0.05
0
reflectance
400
450
500
550
600
650
700
wavelength in nm
leskenl
niittyhumala
ukonhattu
poimulehti
paju
maksaruoho
pelto-ohake

Color Matching Experiments

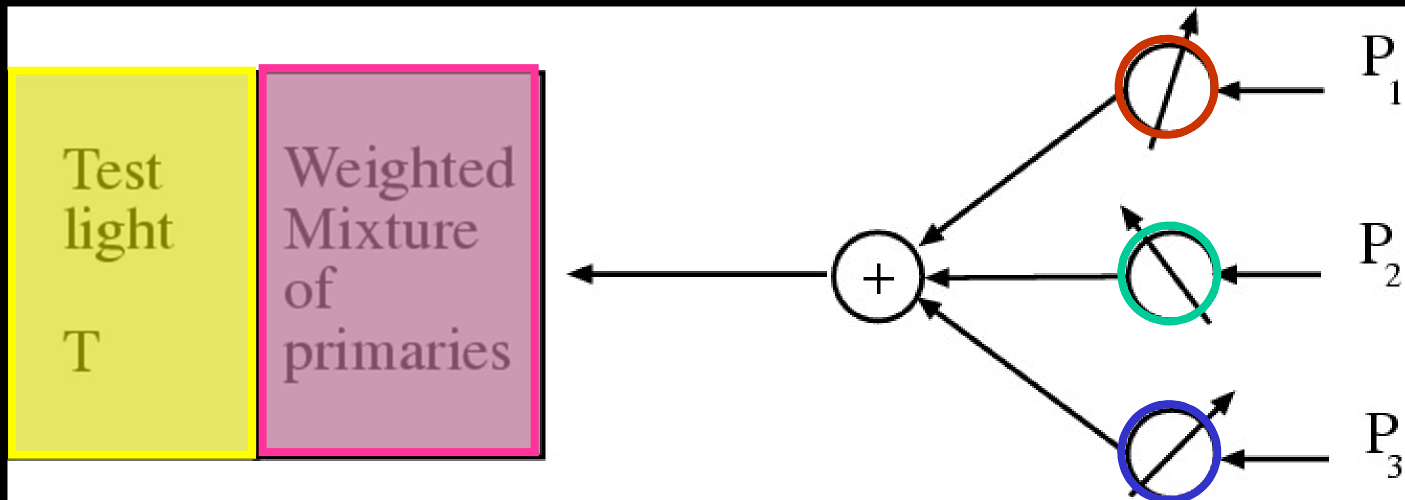

Adjust the knobs on the primaries until the split field looks uniform.

$$T = w_1P_1 + w_2P_2 + \ldots + w_kP_k$$

Notation only!

Trichromacy

$\boxed{\text{Three primaries are sufficient to match any test light!}}$

Note:

- This requires using independent primaries and allowing subtractive matching.
- For a given set of primaries and test light, people tend to pick the same weights.
- This is related to our three types of cones.

Color Matching Experiments

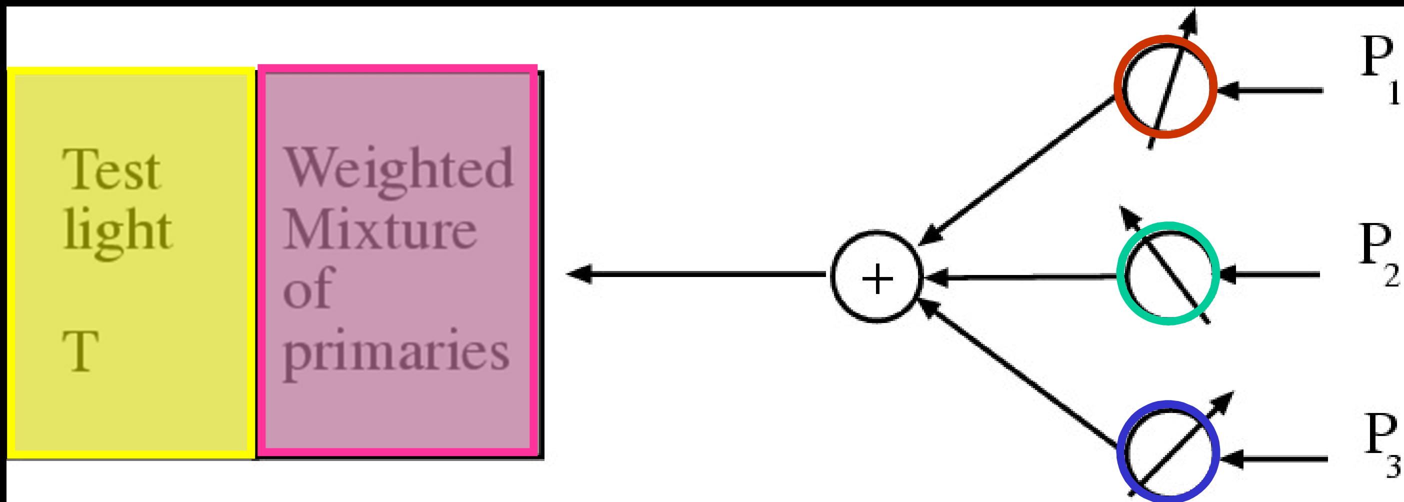

Adjust the knobs on the primaries until the split field looks uniform.

$$T = w_1 P_1 + w_2 P_2 + \ldots + w_k P_k$$

Notation only!

$$T + w_1 P_1 + \ldots + w_n P_n = w_{n+1} P_{n+1} + \ldots + w_k P_k$$

Subtractive matching

Linearity of Colour Matching (Grassman's Laws)

$$T_a = w_{a1}P_1 + w_{a2}P_2 + w_{a3}P_3 \quad \text{and} \quad T_b = w_{b1}P_1 + w_{b2}P_2 + w_{b3}P_3$$

$$\Downarrow$$

$$\lambda\, T_a + \mu\, T_b = (\lambda\, w_{a1} + \mu\, w_{b1})P_1 + (\lambda\, w_{a2} + \mu\, w_{b2})P_2 + (\lambda\, w_{a3} + \mu\, w_{b3})P_3$$

$$T_a = w_1P_1 + w_2P_2 + w_3P_3$$

and

$$T_b = w_1P_1 + w_2P_2 + w_3P_3$$

$$\Rightarrow \quad T_a = T_b$$

Note: this does not mean that T_a and T_b have the same spectrum.

Principle of Univariance

- The eye photoreceptors do not directly measure spectral radiance.

- Two lights will match if they produce the same photoreceptor responses even if they have quite different spectral radiances.

- Because of linearity, the response of each photoreceptor can be modeled as

$$P_k=\int_\Lambda \sigma_k(\lambda)E(\lambda)d\lambda.$$

- The sensitivity σ_k can be measured experimentally.

The human eye..

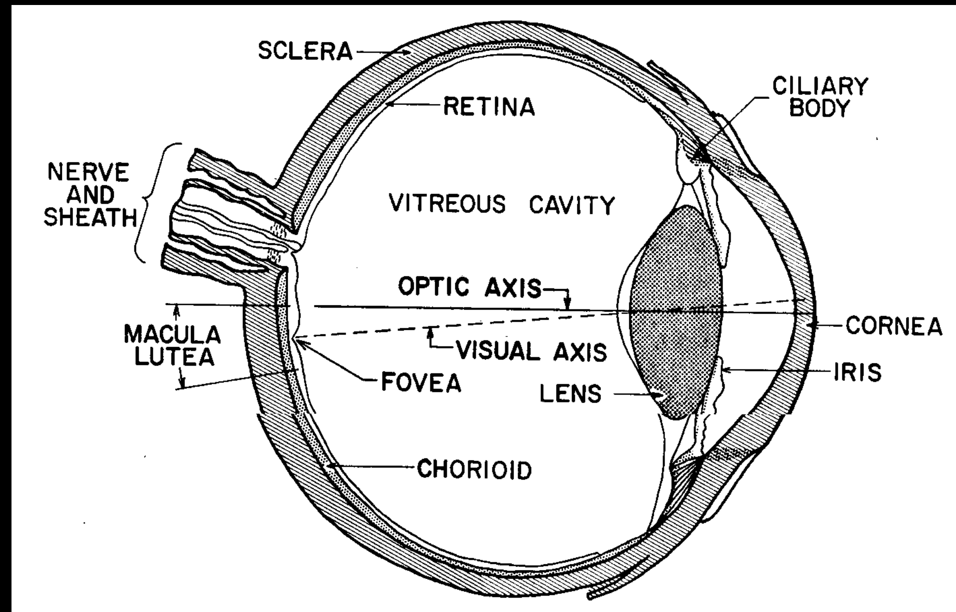

Cones in the fovea

Rods and cones in the periphery

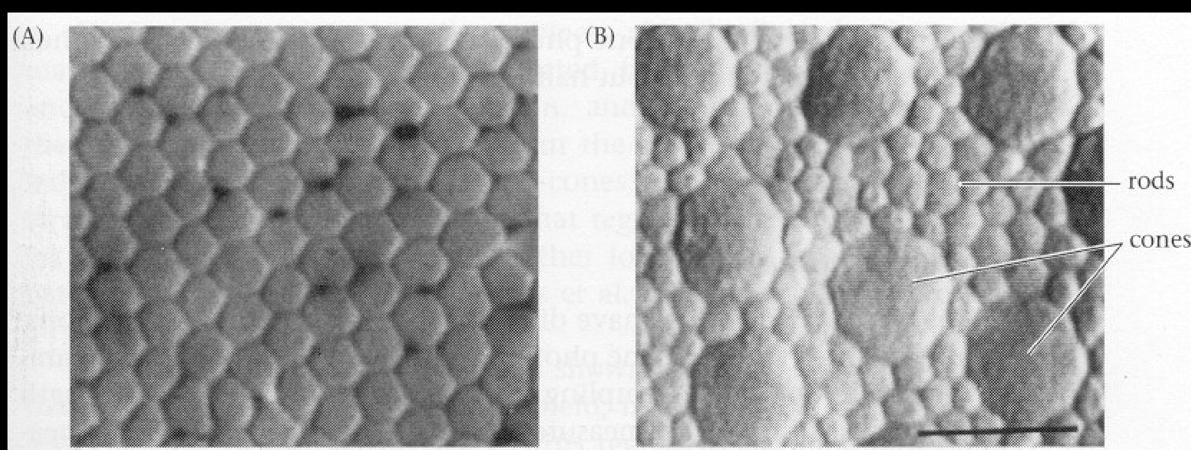

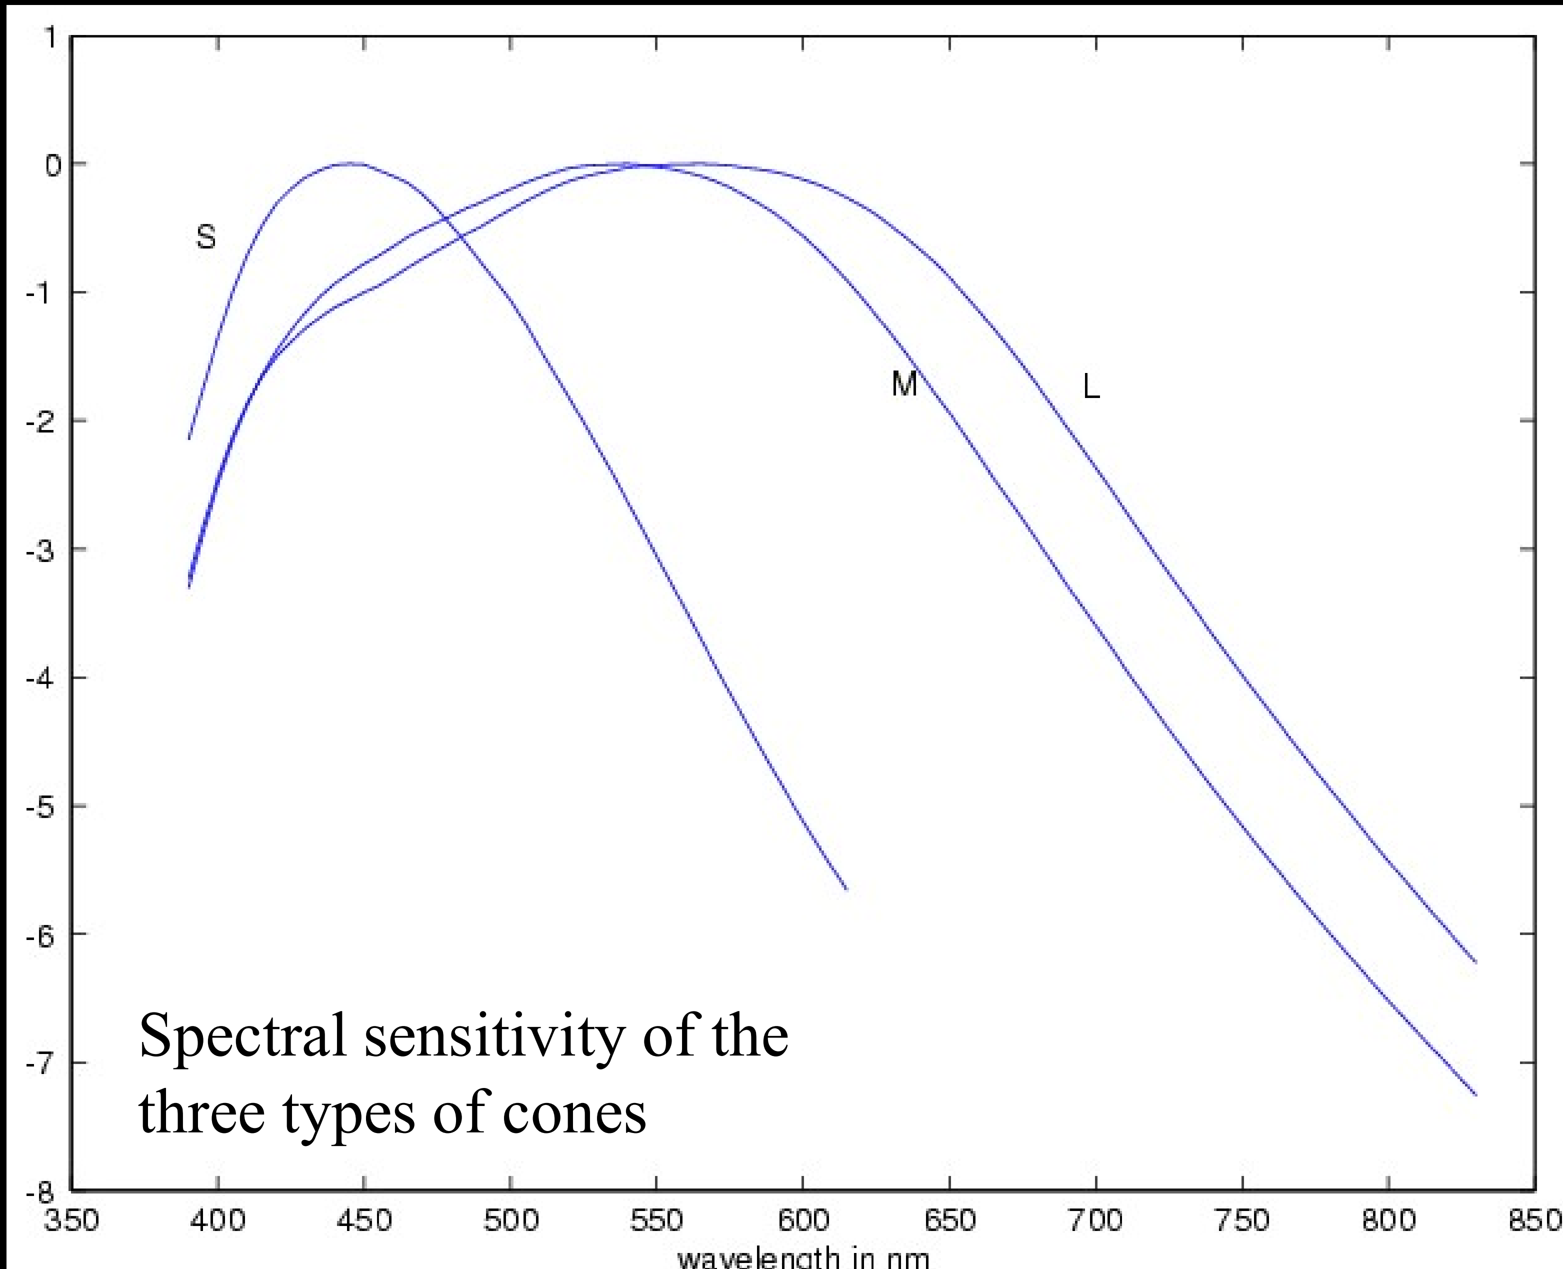

Spectral sensitivity of the three types of cones

Color Matching Functions

- Problem: given a set of primaries, what are the weights matching a given spectral radiance?

- Experiments:

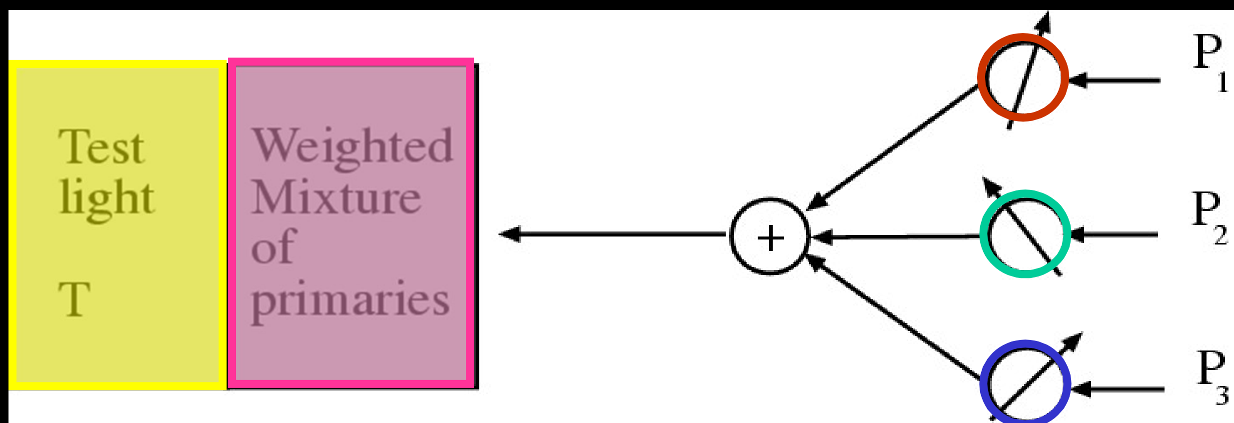

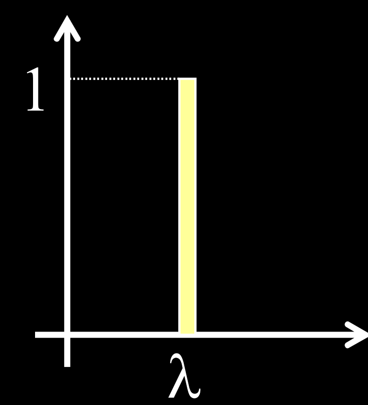

$$L(\lambda) = f_1(\lambda)P_1 + f_2(\lambda)P_2 + f_3(\lambda)P_3$$

Color matching functions

- To match S use linearity:

$$\boxed{S=\left(\int_\Lambda f_1(\lambda)S(\lambda)d\lambda\right)P_1+\left(\int_\Lambda f_2(\lambda)S(\lambda)d\lambda\right)P_2+\left(\int_\Lambda f_3(\lambda)S(\lambda)d\lambda\right)P_3}$$

RGB Color Matching Functions

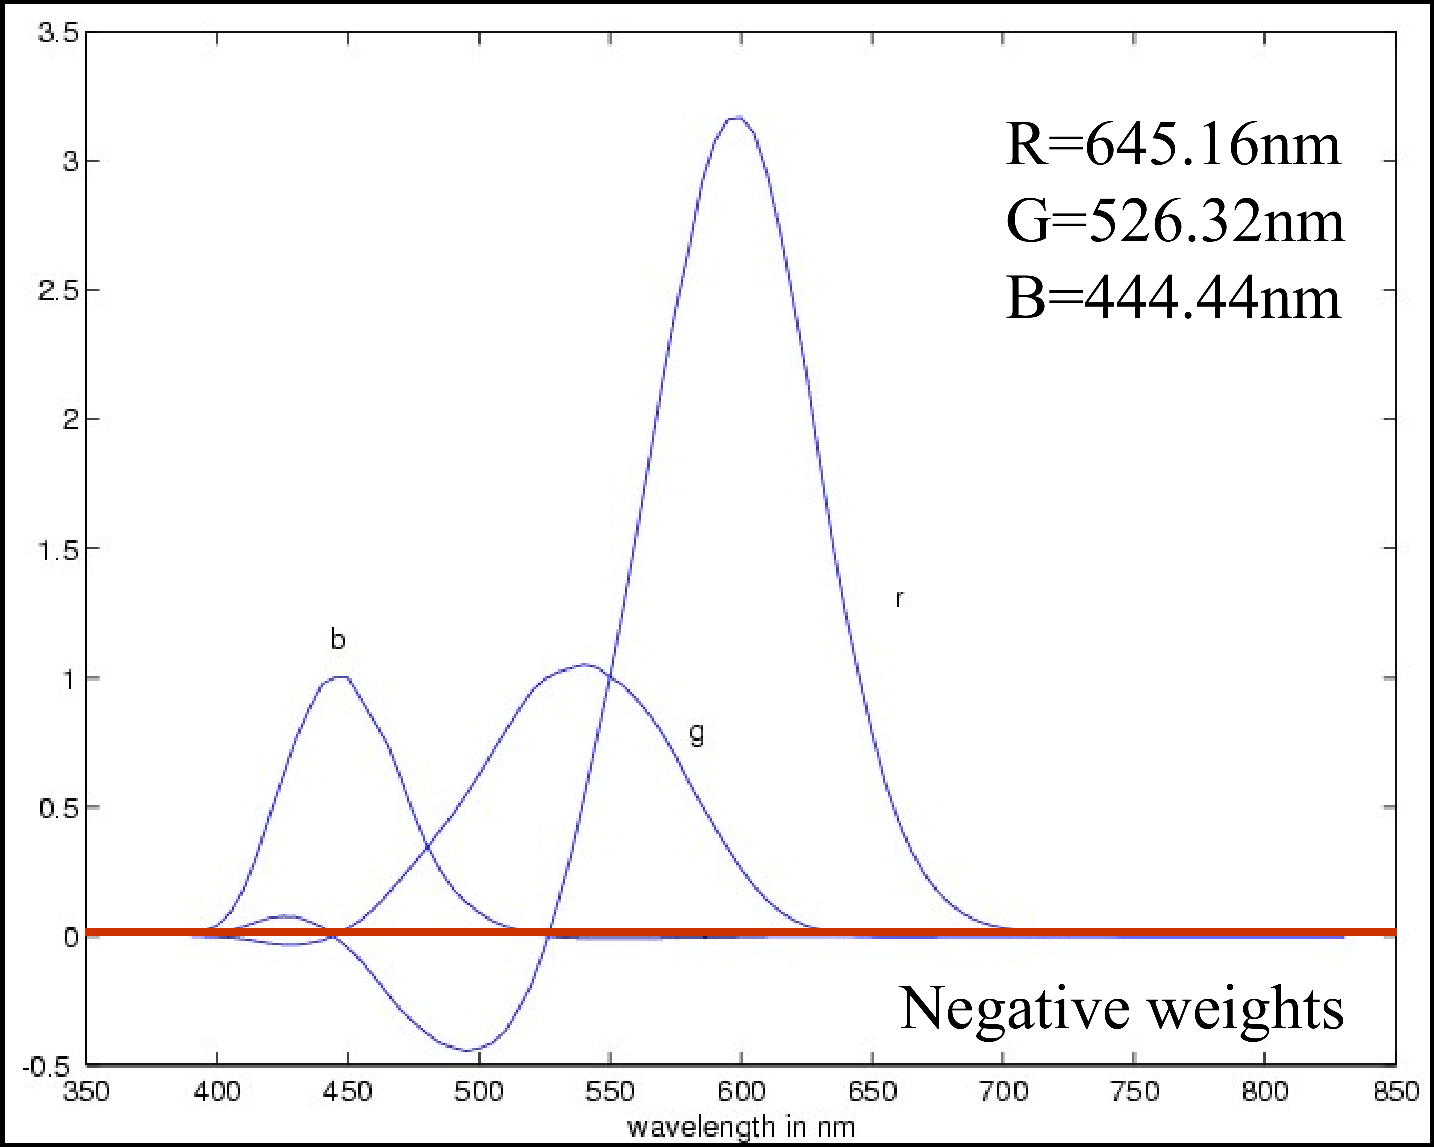

A rather poor reproduction of the RGB color cube..

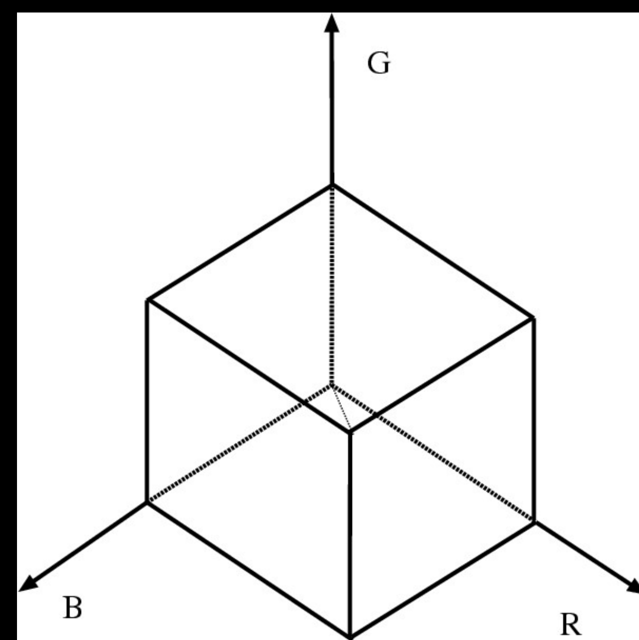

CIE XYZ Color Matching Functions

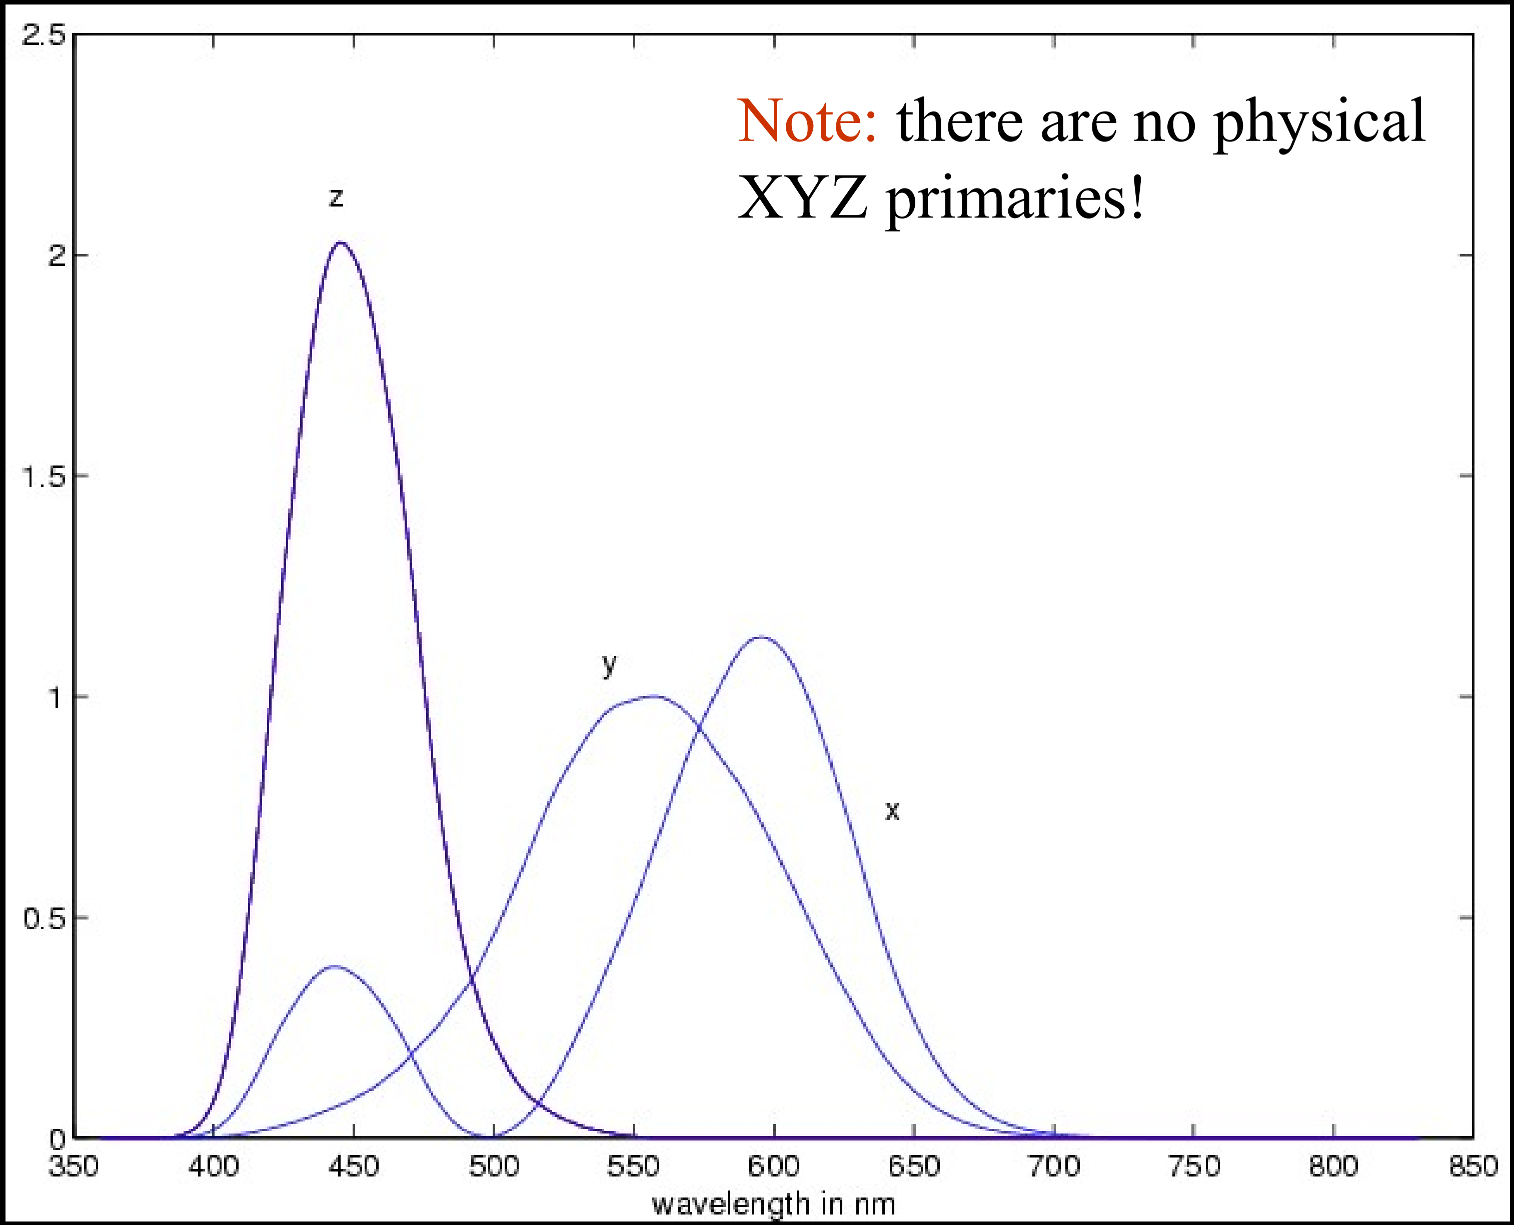

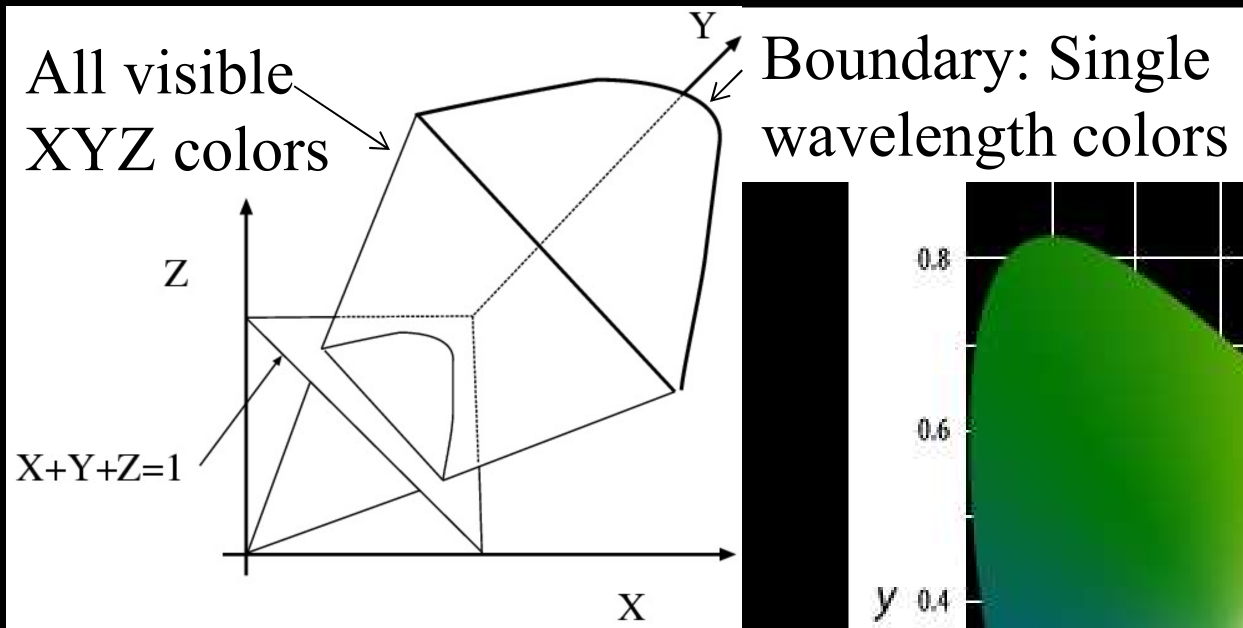

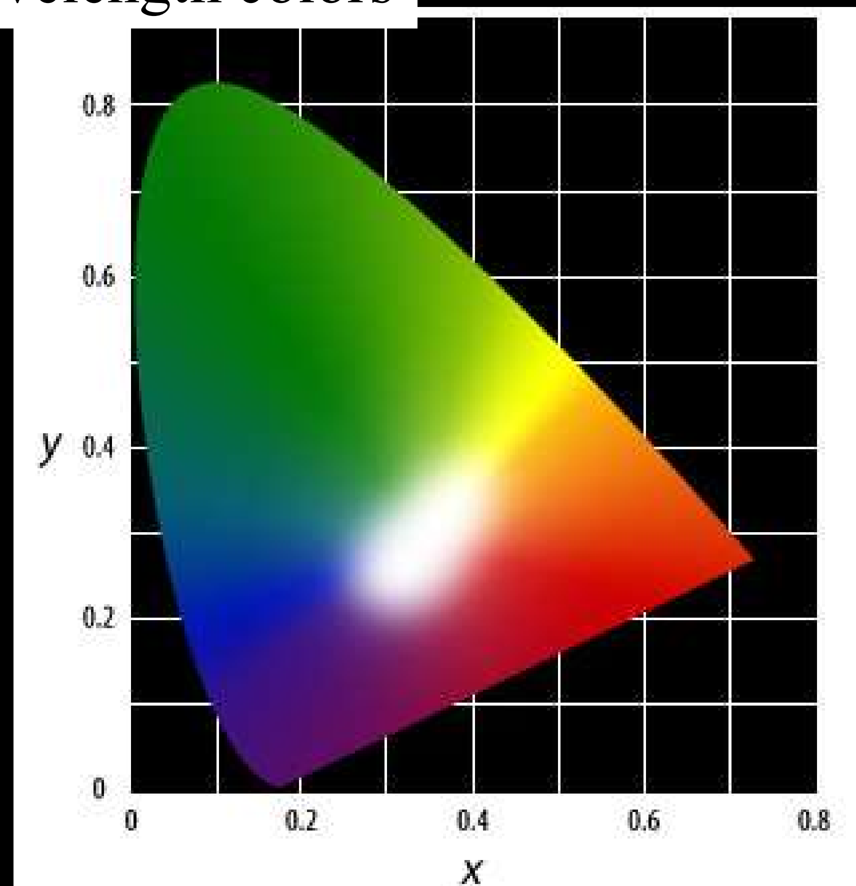

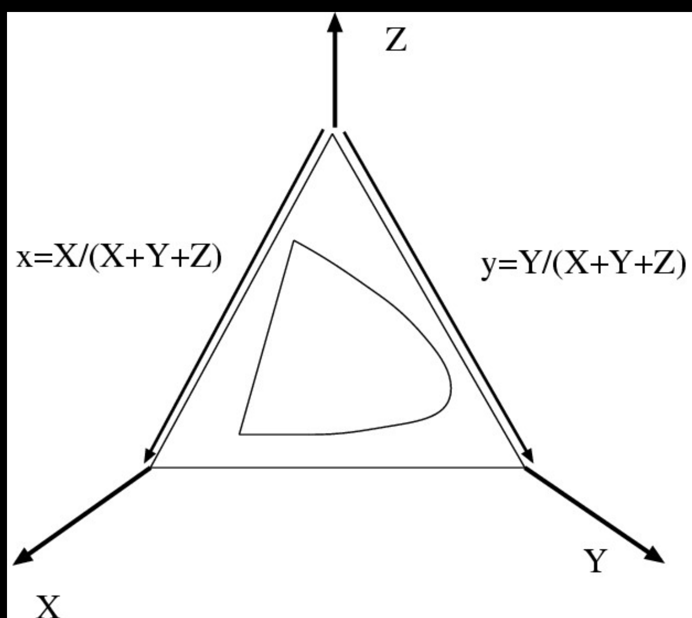

CIE XYZ and xy spaces