Review

- Pinhole projection model
  - What are vanishing points and vanishing lines?
  - What is orthographic projection?
  - How can we approximate orthographic projection?
- Lenses
  - Why do we need lenses?
  - What is depth of field?
  - What controls depth of field?
  - What is field of view?
  - What controls field of view?
  - What are some kinds of lens aberrations?
- Digital cameras
  - What are the two major types of sensor technologies?
  - How can we capture color with a digital camera?

Assignment 1: Demosaicing

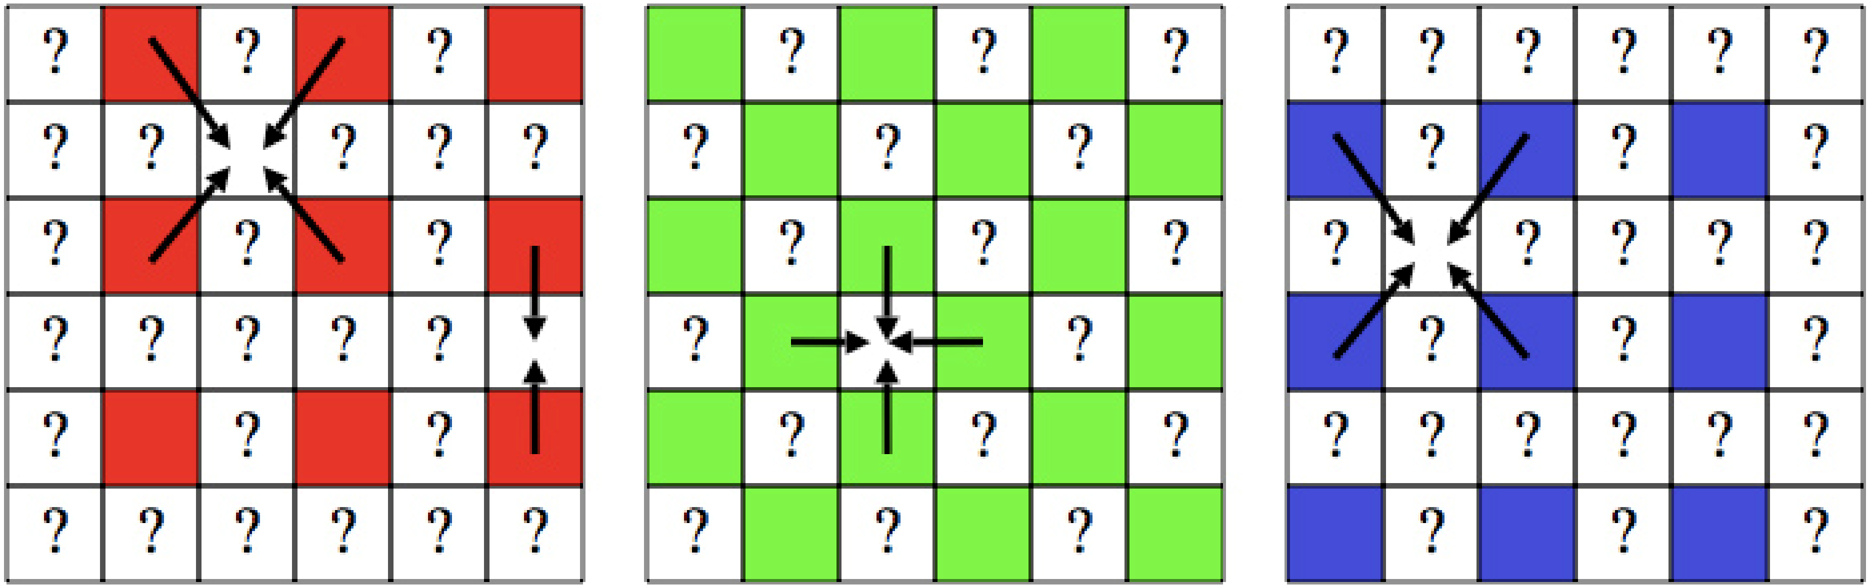

Historical context

- **Pinhole model:** Mozi (470-390 BCE), Aristotle (384-322 BCE)
- **Principles of optics (including lenses):** Alhacen (965-1039 CE)
- **Camera obscura:** Leonardo da Vinci (1452-1519), Johann Zahn (1631-1707)
- **First photo:** Joseph Nicephore Niepce (1822)
- **Daguerréotypes** (1839)
- **Photographic film:** Eastman (1889)
- **Cinema:** Lumière Brothers (1895)
- **Color Photography:** Lumière Brothers (1908)
- **Television:** Baird, Farnsworth, Zworykin (1920s)
- **First digitally scanned photograph:** Russell Kirsch, NIST (1957)
- **First consumer camera with CCD:** Sony Mavica (1981)
- **First fully digital camera:** Kodak DCS100 (1990)

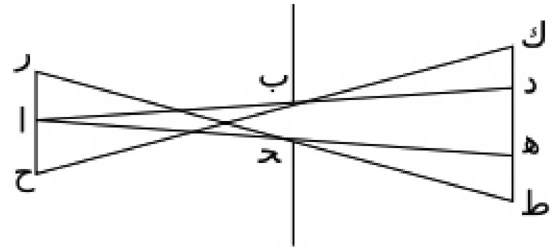

Alhacen's notes

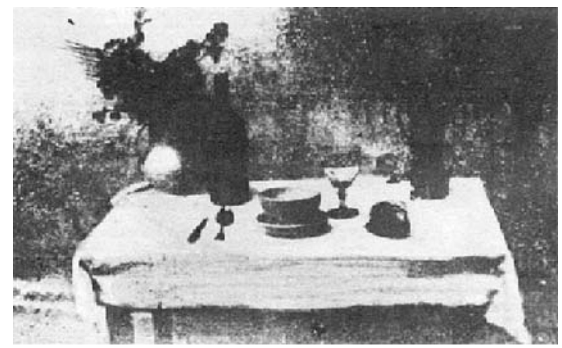

Niepce, "La Table Servie," 1822

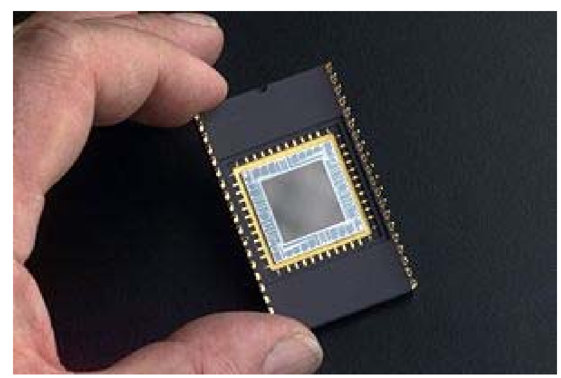

CCD chip

10 Early Firsts In Photography

http://listverse.com/history/top-10-incredible-early-firsts-in-photography/

Early color photography

Sergey Prokudin-Gorsky (1863-1944)

Photographs of the Russian empire (1909-1916)

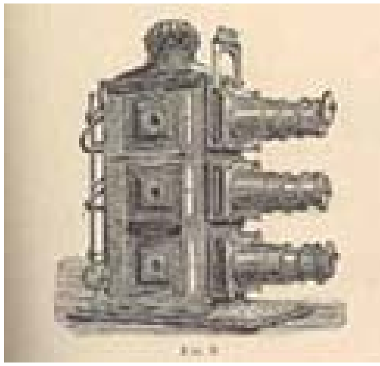

Lantern projector

http://en.wikipedia.org/wiki/Sergei_Mikhailovich_Prokudin-Gorskii

http://www.loc.gov/exhibits/empire/

"Fake miniatures"

Create your own fake miniatures: http://tiltshiftmaker.com/
http://tiltshiftmaker.com/tilt-shift-photo-gallery.php

Idea for class participation: if you find interesting (and relevant) links, send them to me or (better yet) to the class mailing list (comp776@cs.unc.edu).

Today: Capturing light

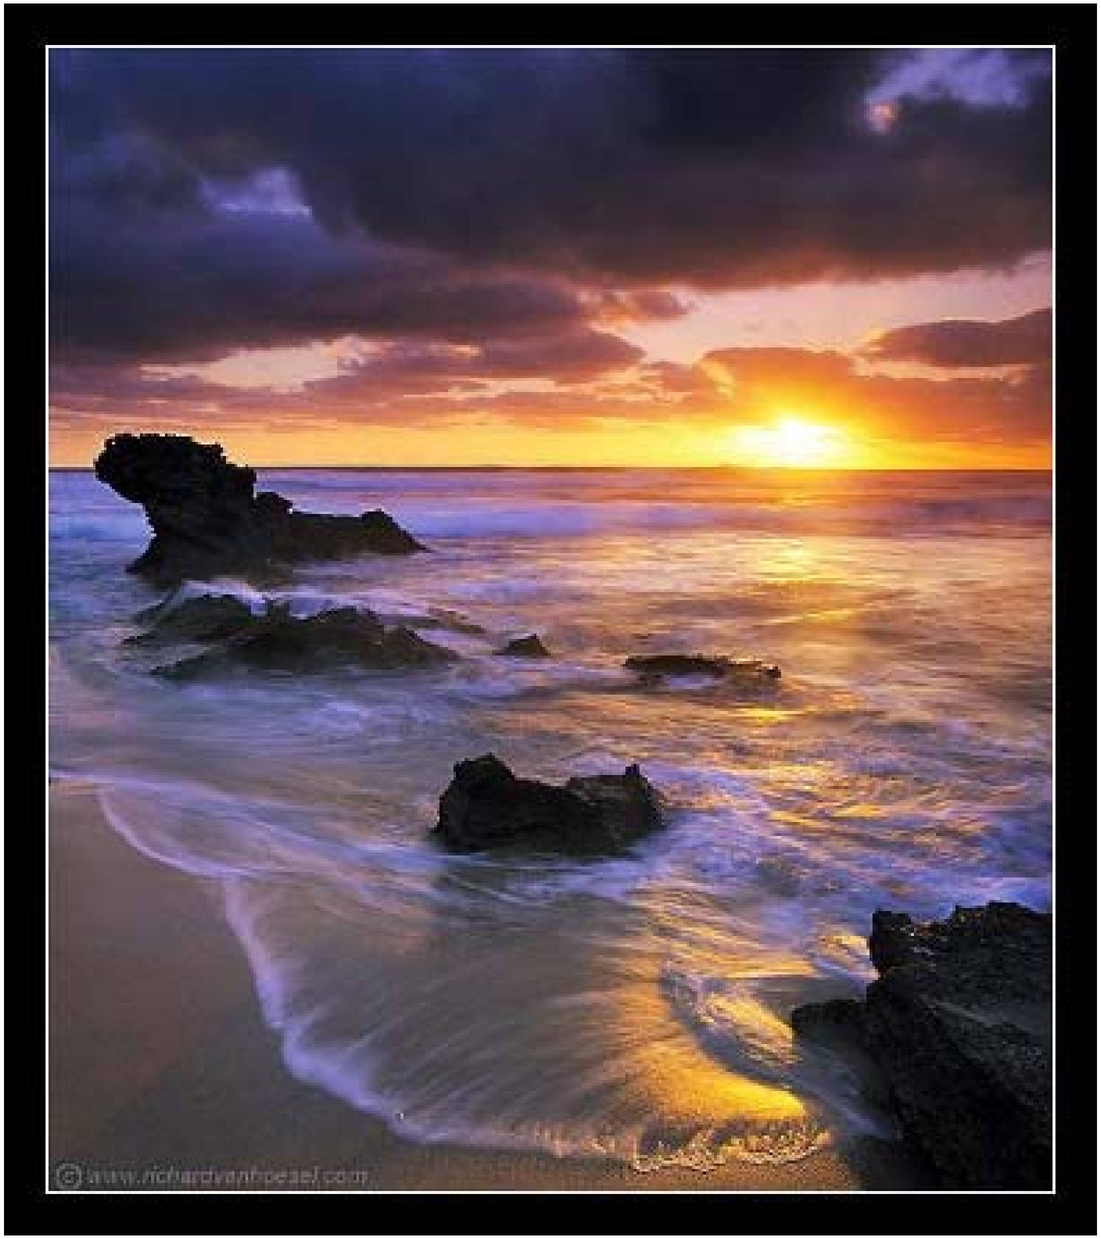

Source: A. Efros

Radiometry

What determines the brightness of an image pixel?

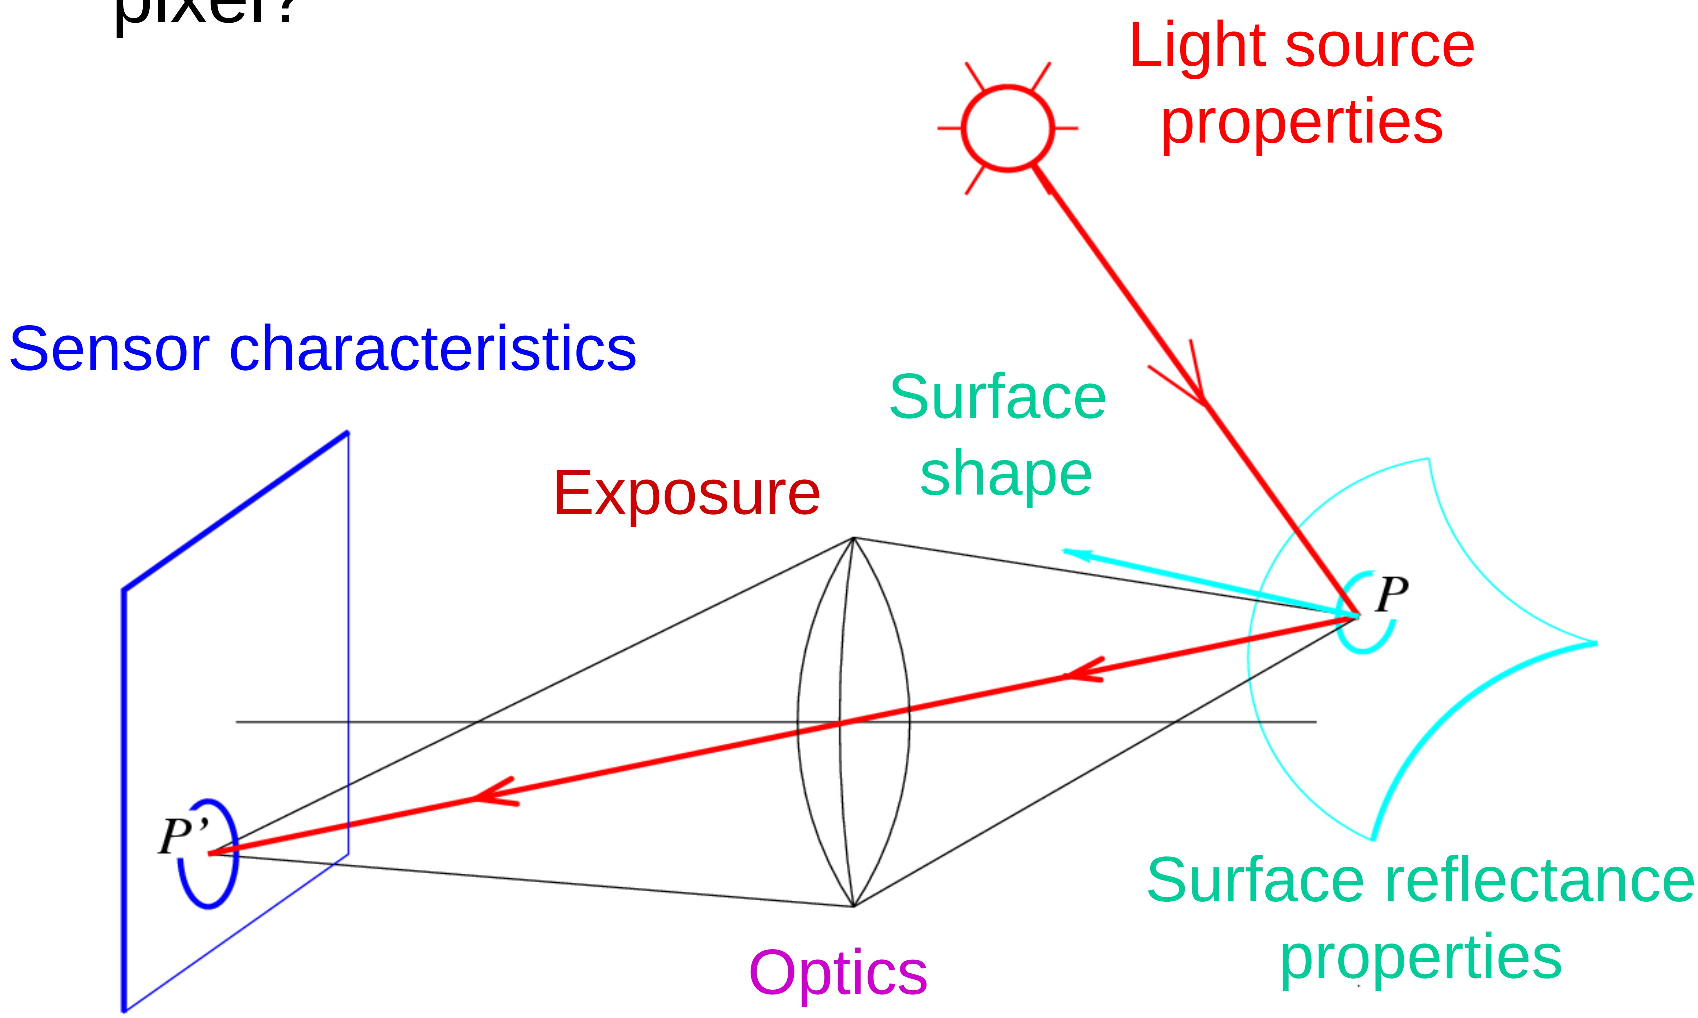

Slide by L. Fei-Fei

Solid Angle

- By analogy with angle (in radians), the solid angle subtended by a region at a point is the area projected on a unit sphere centered at that point
- The solid angle $d\omega$ subtended by a patch of area *dA* is given by:

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

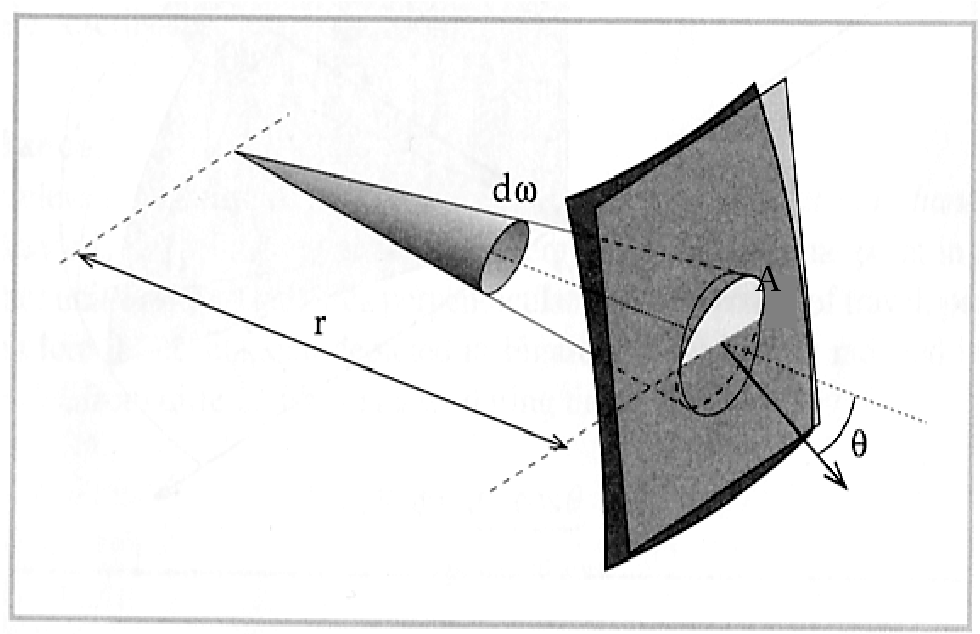

Radiometry

- Radiance (L): energy carried by a ray
  - Power per unit area perpendicular to the direction of travel, per unit solid angle
  - Units: Watts per square meter per steradian (W m^{-2} sr^{-1})
- Irradiance (E): energy arriving at a surface
  - Incident power in a given direction per unit area
  - Units: W m^{-2}
  - For a surface receiving radiance $L(x,\theta,\phi)$ coming in from dω the corresponding irradiance is

$$E(\theta,\phi) = L(\theta,\phi)\cos\theta \, d\omega$$

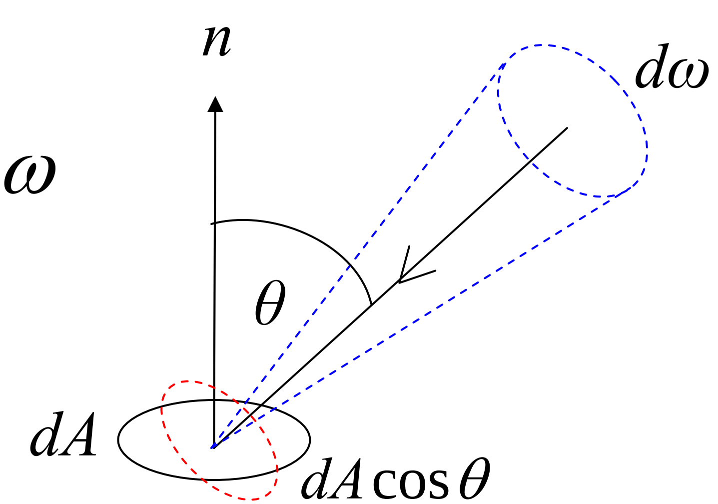

Radiometry of thin lenses

L: Radiance emitted from *P* toward *P*'

E: Irradiance falling on *P*' from the lens

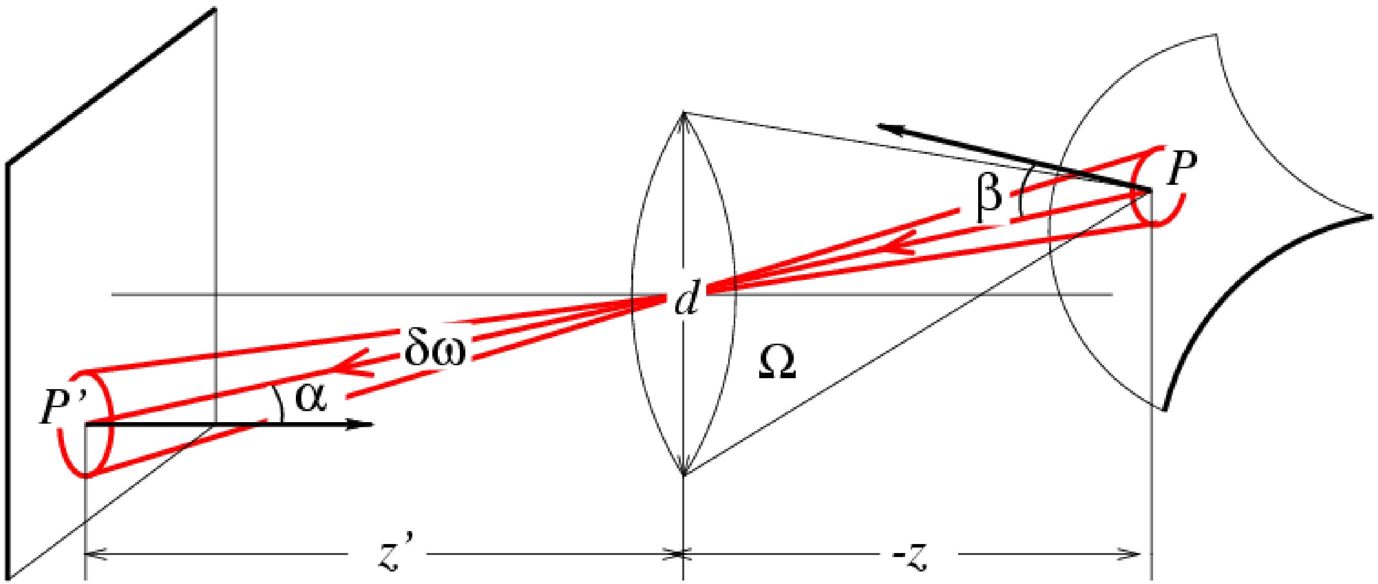

What is the relationship between *E* and *L*?

Forsyth & Ponce, Sec. 4.2.3

Example: Radiometry of thin lenses

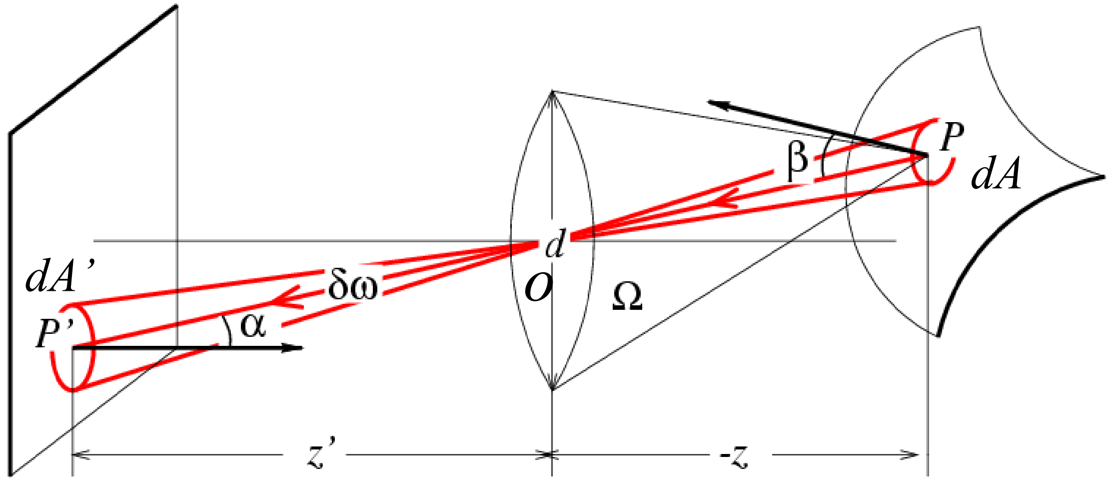

$$|OP| = \frac{z}{\cos\alpha}$$

$$|OP'| = \frac{z'}{\cos\alpha}$$

Area of the lens: $\frac{\pi d^2}{4}$

The power δP received by the lens from *P* is $\delta P = L\left(\frac{\pi d^2}{4}\right)\cos\alpha\,\delta\omega$

The radiance emitted from the lens towards dA’ is $\frac{\delta P}{\left(\frac{\pi d^2}{4}\right)\cos\alpha\,\delta\omega} = L$

The irradiance received at P’ is

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

Radiometry of thin lenses

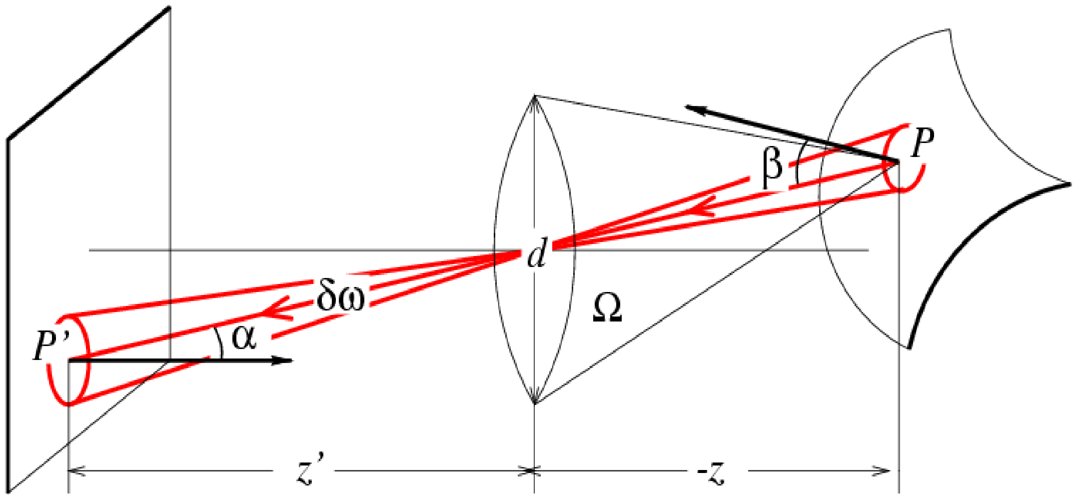

$$E = \left[\frac{\pi}{4} \left(\frac{d}{z'} \right)^2 \cos^4 \alpha \right] L$$

- Image irradiance is linearly related to scene radiance
- Irradiance is proportional to the area of the lens and inversely proportional to the squared distance between the lens and the image plane
- The irradiance falls off as the angle between the viewing ray and the optical axis increases

Forsyth & Ponce, Sec. 4.2.3

Radiometry of thin lenses

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

- Application:
  - S. B. Kang and R. Weiss, *Can we calibrate a camera using an image of a flat, textureless Lambertian surface?* ECCV 2000.

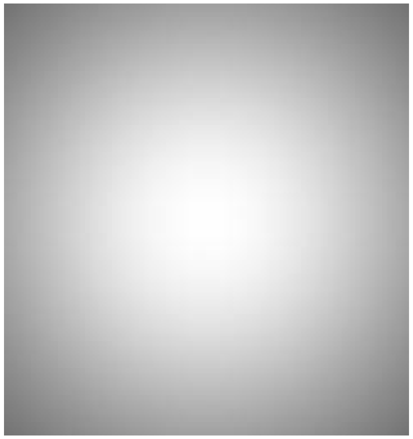

The journey of the light ray

$$X = E \cdot \Delta t$$

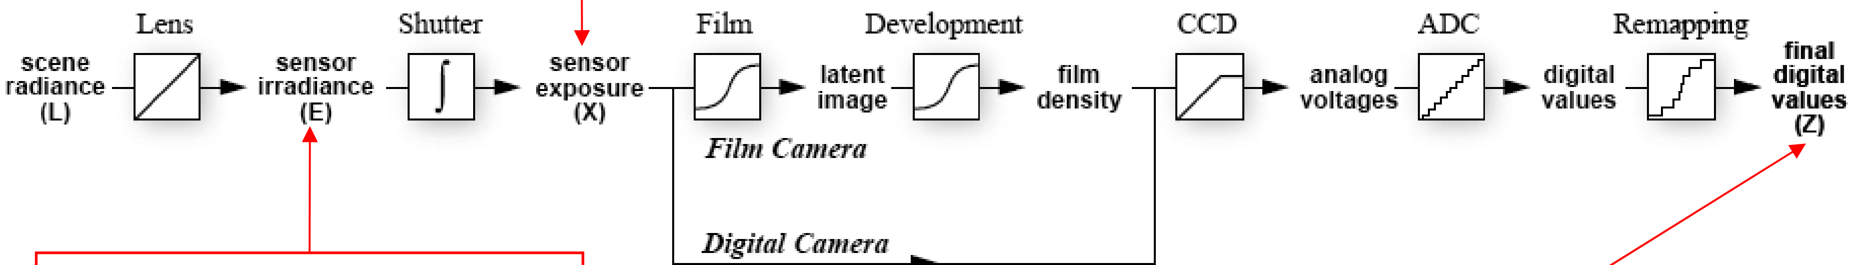

$$E = \left[\frac{\pi}{4} \left(\frac{d}{z'} \right)^2 \cos^4 \alpha \right] L$$

$$Z = f\left(E \cdot \Delta t\right)$$

- Camera response function: the mapping f from irradiance to pixel values
  - Useful if we want to estimate material properties
  - Enables us to create high dynamic range images

Source: S. Seitz, P. Debevec

The journey of the light ray

$$X = E \cdot \Delta t$$

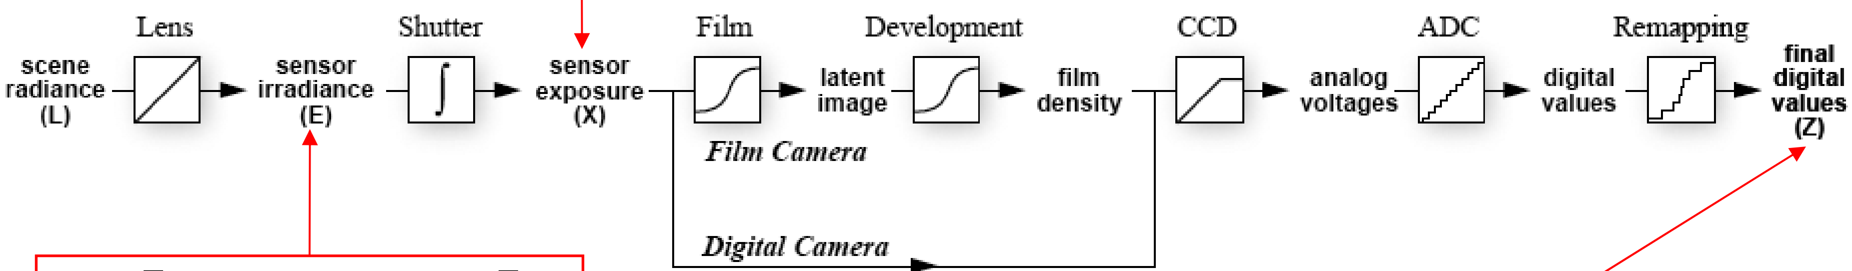

$$E = \left[\frac{\pi}{4} \left(\frac{d}{z'} \right)^2 \cos^4 \alpha \right] L$$

$$Z = f\left(E \cdot \Delta t\right)$$

- Camera response function: the mapping f from irradiance to pixel values

For more info

- P. E. Debevec and J. Malik. *Recovering High Dynamic Range Radiance Maps from Photographs*. In *SIGGRAPH 97*, August 1997

Source: S. Seitz, P. Debevec

The interaction of light and surfaces

What happens when a light ray hits a point on an object?

- Some of the light gets absorbed
  - converted to other forms of energy (e.g., heat)
- Some gets transmitted through the object
  - possibly bent, through “refraction”
- Some gets reflected
  - possibly in multiple directions at once
- Really complicated things can happen
  - fluorescence

Let’s consider the case of reflection in detail

- In the most general case, a single incoming ray could be reflected in all directions. How can we describe the amount of light reflected in each direction?

Slide by Steve Seitz

Bidirectional reflectance distribution function (BRDF)

- Model of local reflection that tells how bright a surface appears when viewed from one direction when light falls on it from another
- Definition: ratio of the radiance in the outgoing direction to irradiance in the incident direction

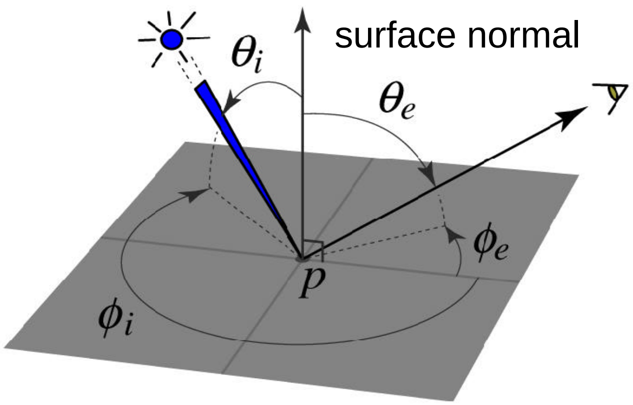

$$\rho(\theta_i,\phi_i,\theta_e,\phi_e) = \frac{L_e(\theta_e,\phi_e)}{E_i(\theta_i,\phi_i)} = \frac{L_e(\theta_e,\phi_e)}{L_i(\theta_i,\phi_i)\cos\theta_i d\omega}$$

- Radiance leaving a surface in a particular direction: add contributions from every incoming direction

$$\int_\Omega \rho(\theta_i,\phi_i,\theta_e,\phi_e,)L_i(\theta_i,\phi_i)\cos\theta_i d\omega_i$$

BRDF's can be incredibly complicated…

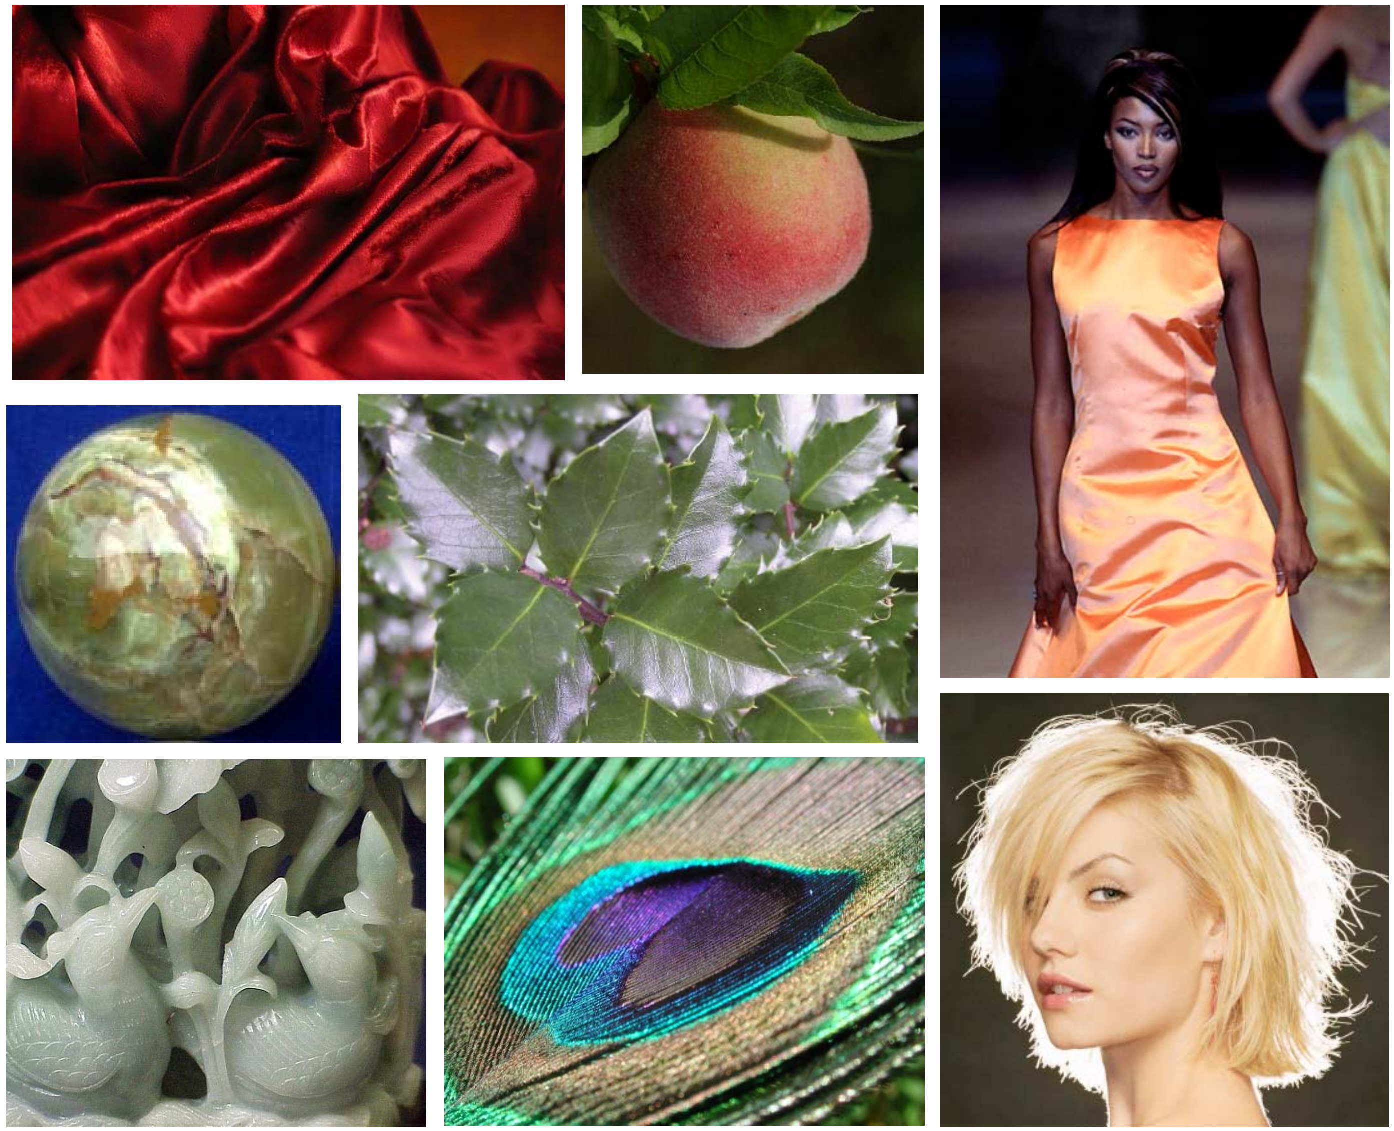

Diffuse reflection

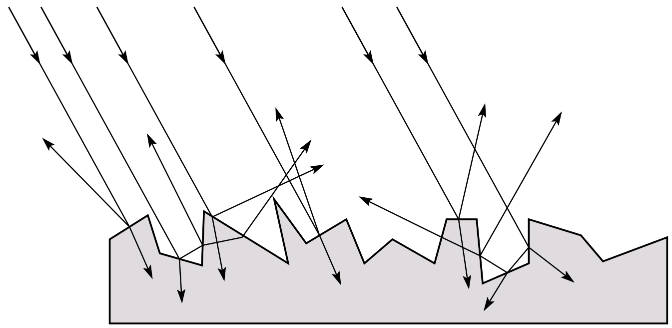

- Light is reflected equally in all directions: **BRDF is constant**
- Dull, matte surfaces like chalk or latex paint
- Microfacets scatter incoming light randomly
- *Albedo*: fraction of incident irradiance reflected by the surface
- *Radiosity:* total power leaving the surface per unit area (regardless of direction)

Diffuse reflection: Lambert's law

- Viewed brightness does not depend on viewing direction, but it *does* depend on direction of illumination

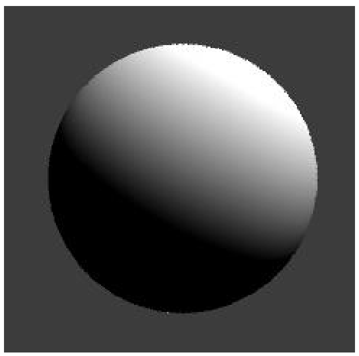

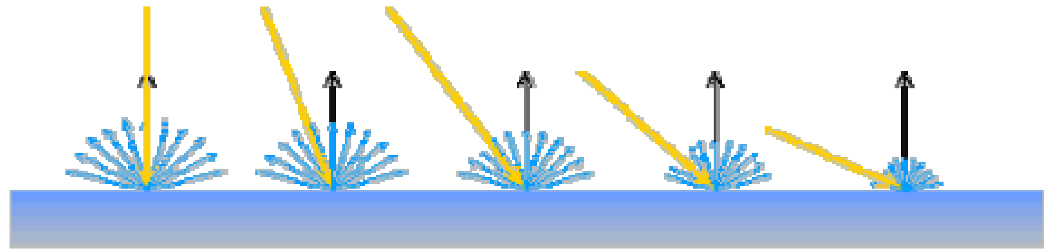

$$B(x) = \rho_d(x)\left(N(x) \cdot S_d(x)\right)$$

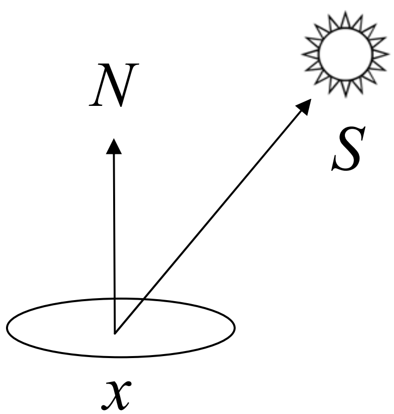

B: radiosity
ρ: albedo
N: unit normal
S: source vector (magnitude proportional to intensity of the source)

Specular reflection

- Radiation arriving along a source direction leaves along the specular direction (source direction reflected about normal)
- Some fraction is absorbed, some reflected
- On real surfaces, energy usually goes into a lobe of directions
- Phong model: reflected energy falls of with $\cos^n(\delta\theta)$
- Lambertian + specular model: sum of diffuse and specular term

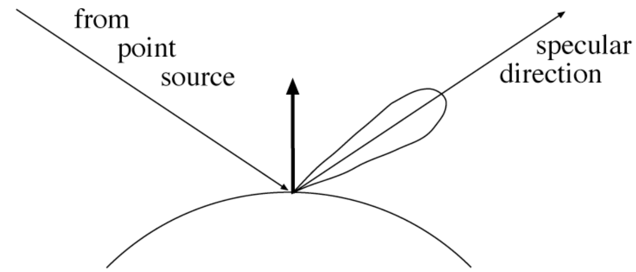

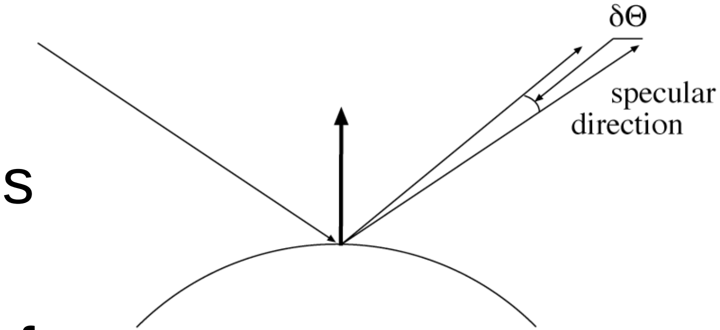

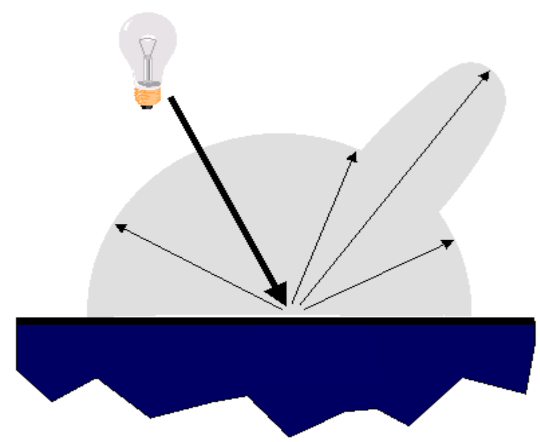

Specular reflection

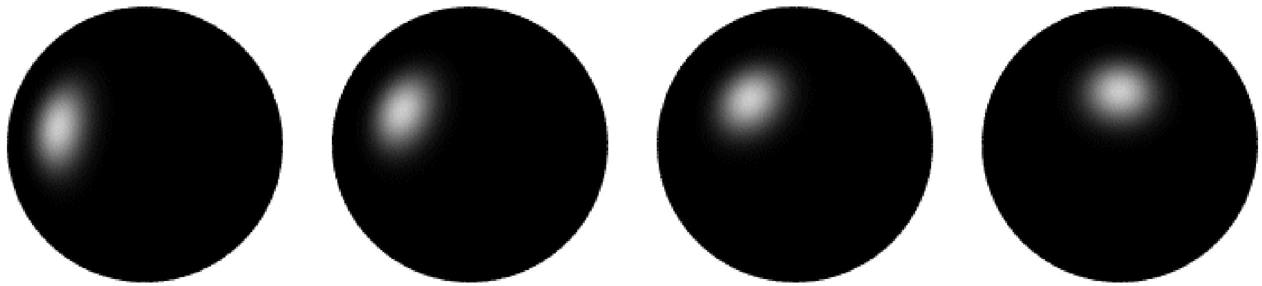

Moving the light source

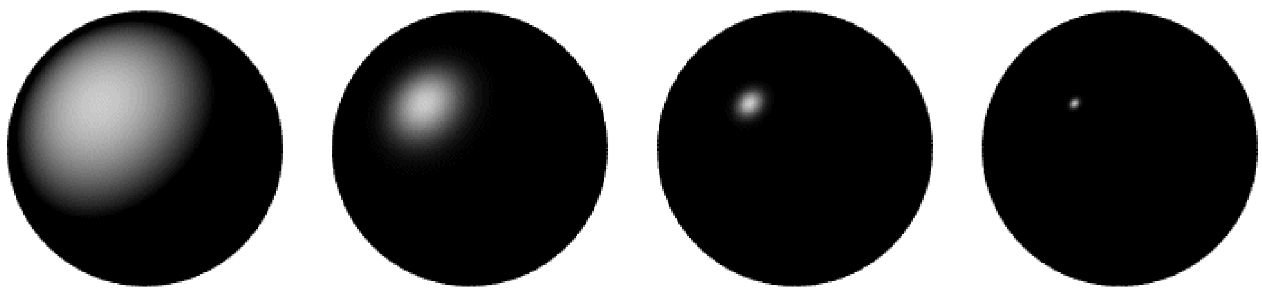

Changing the exponent

Photometric stereo

Assume:

- A Lambertian object
- A *local shading model* (each point on a surface receives light only from sources visible at that point)
- A set of *known* light source directions
- A set of pictures of an object, obtained in exactly the same camera/object configuration but using different sources
- Orthographic projection

Goal: reconstruct object shape and albedo

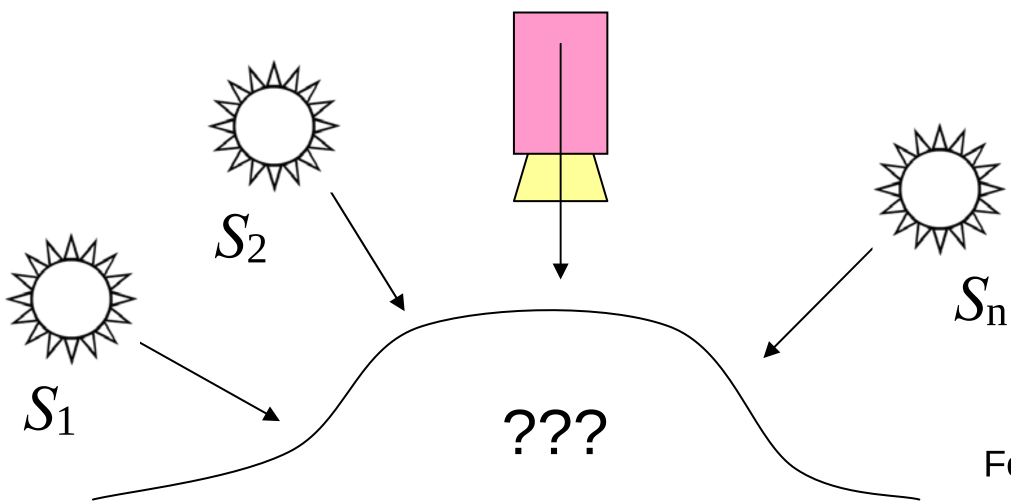

Forsyth & Ponce, Sec. 5.4

Surface model: Monge patch

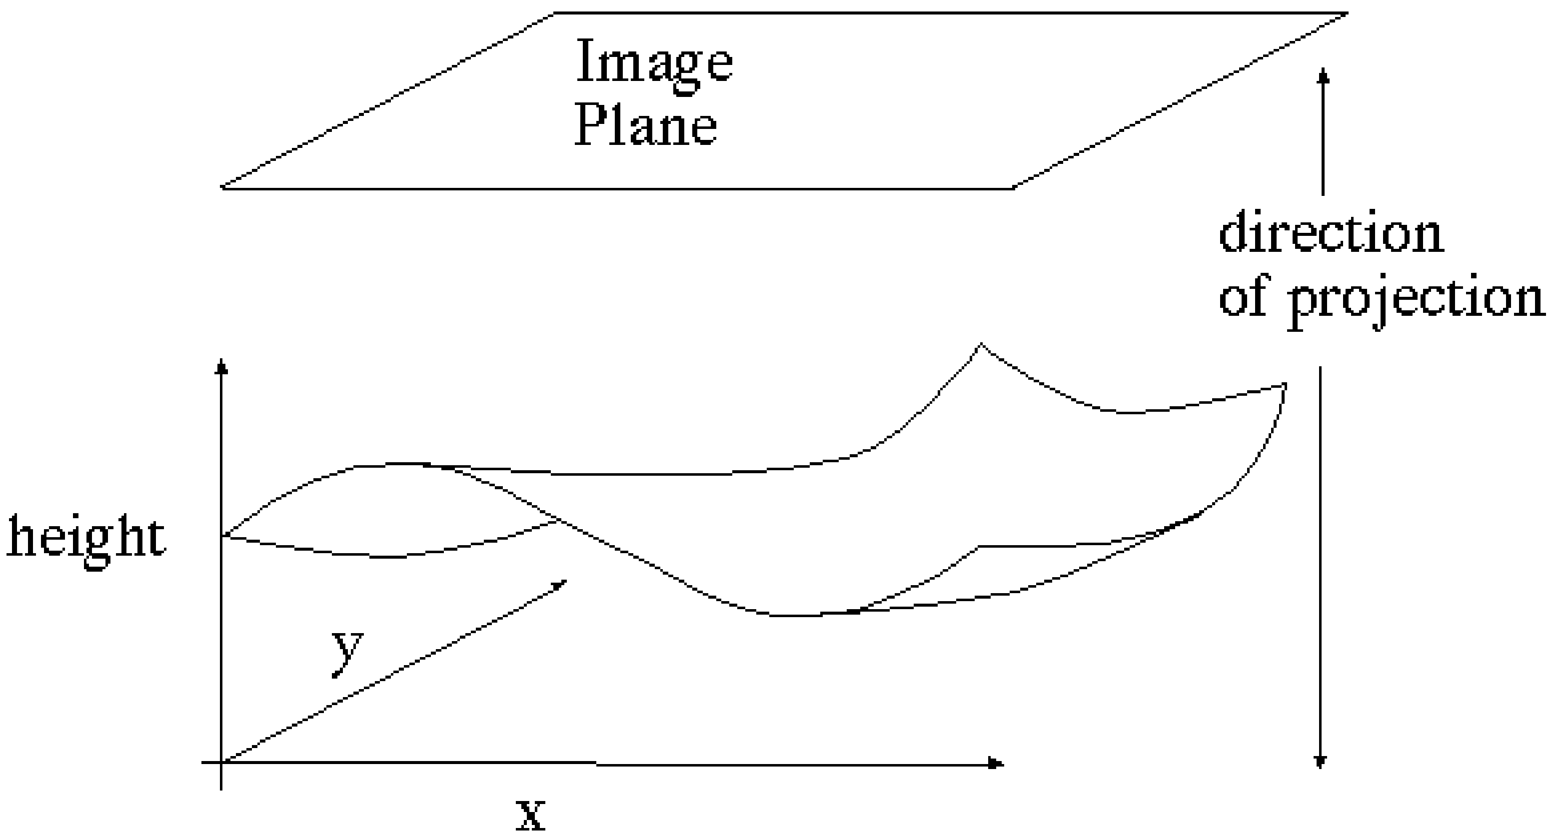

Forsyth & Ponce, Sec. 5.4

Image model

- Known: source vectors S_j and pixel values $I_j(x,y)$
- We also assume that the response function of the camera is a linear scaling by a factor of k
- Combine the unknown normal $N(x,y)$ and albedo $\rho(x,y)$ into one vector g, and the scaling constant k and source vectors S_j into another vector V_j:

$$
\begin{aligned}
I_j(x,y) &= k\,B(x,y) \\
&= k\,\rho(x,y)\big(N(x,y)\cdot S_j\big) \\
&= \big(\rho(x,y)N(x,y)\big)\cdot(k\,S_j) \\
&= g(x,y)\cdot V_j
\end{aligned}
$$

Forsyth & Ponce, Sec. 5.4

Least squares problem

- For each pixel, we obtain a linear system:

$$\begin{bmatrix} I_1(x,y) \\ I_2(x,y) \\ \vdots \\ I_n(x,y) \end{bmatrix} = \begin{bmatrix} V_1^T \\ V_2^T \\ \vdots \\ V_n^T \end{bmatrix} g(x,y)$$

($n \times 1$) known — ($n \times 3$) known — (3×1) unknown

- Obtain least-squares solution for $g(x,y)$
- Since $N(x,y)$ is the unit normal, $\rho(x,y)$ is given by the magnitude of $g(x,y)$ (and it should be less than 1)
- Finally, $N(x,y) = g(x,y) / \rho(x,y)$

Forsyth & Ponce, Sec. 5.4

Example

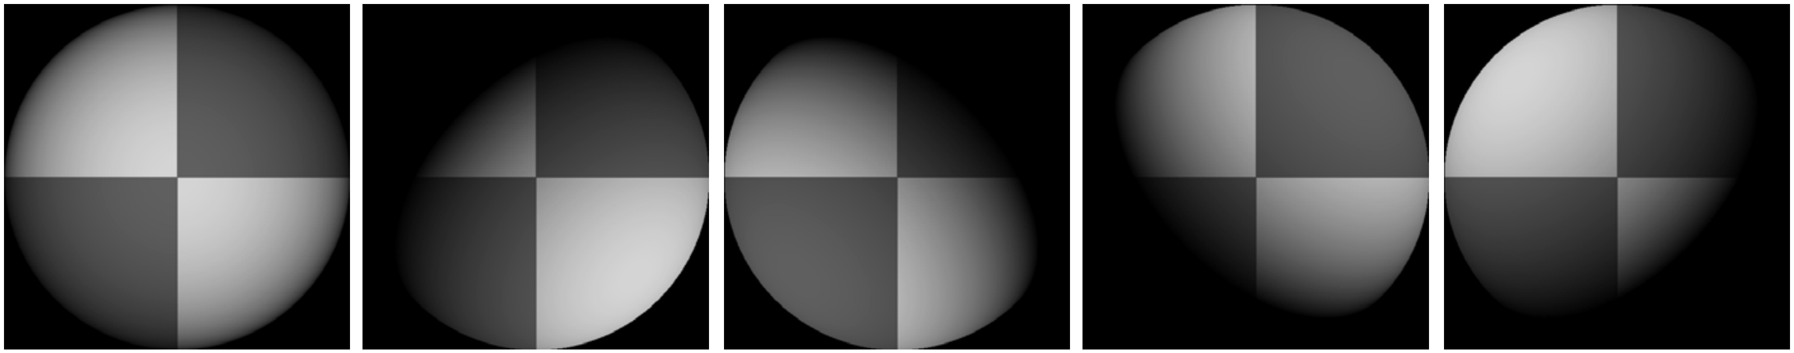

Recovered albedo

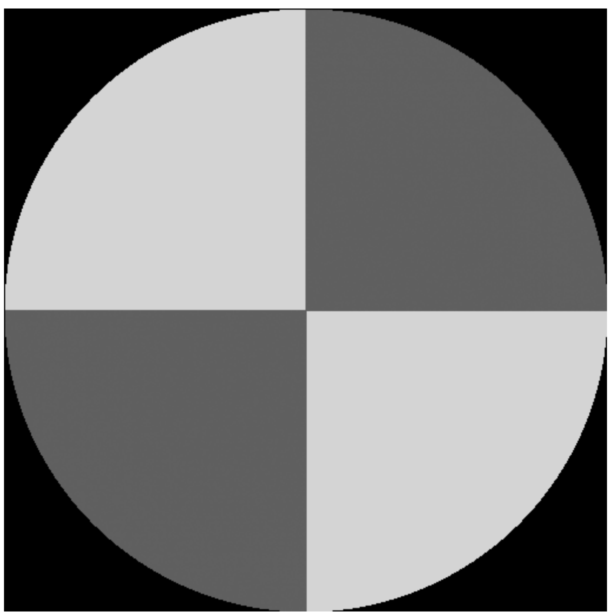

Recovered normal field

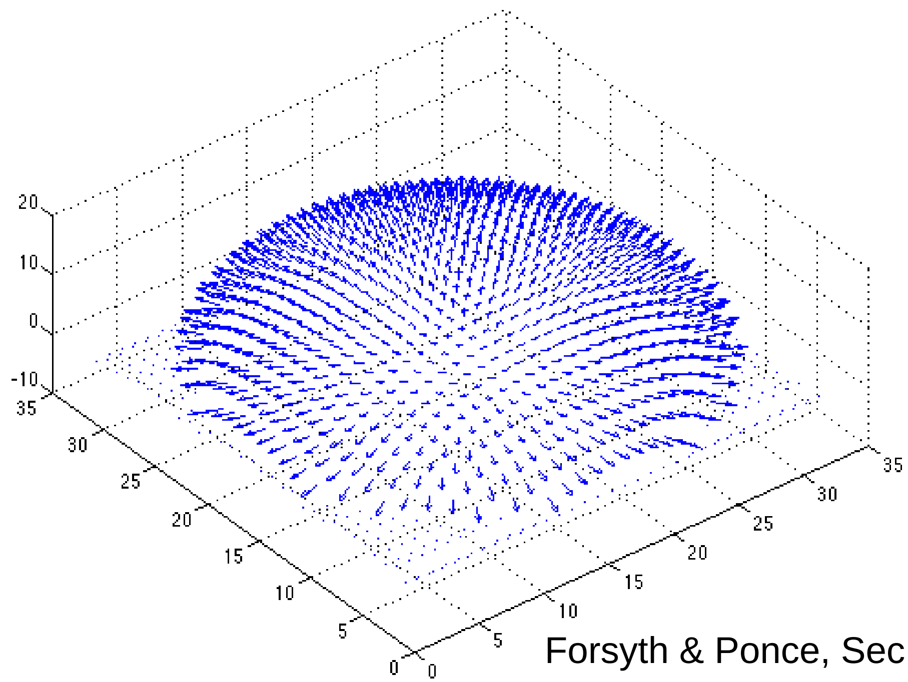

Forsyth & Ponce, Sec. 5.4

Recovering a surface from normals

Recall the surface is written as

$$(x, y, f(x, y))$$

This means the normal has the form:

$$N(x,y) = \left(\frac{1}{\sqrt{f_x^2 + f_y^2 + 1}} \right) \begin{pmatrix} -f_x \\ -f_y \\ 1 \end{pmatrix}$$

If we write the estimated vector *g* as

$$\mathbf{g}(x,y) = \begin{pmatrix} g_1(x,y) \\ g_2(x,y) \\ g_3(x,y) \end{pmatrix}$$

Then we obtain values for the partial derivatives of the surface:

$$f_x(x,y) = \left(g_1(x,y) / g_3(x,y) \right)$$

$$f_y(x,y) = \left(g_2(x,y) / g_3(x,y) \right)$$

Forsyth & Ponce, Sec. 5.4

Recovering a surface from normals

Integrability: for the surface f to exist, the mixed second partial derivatives must be equal:

$$\frac{\partial\left(g_1(x,y)/g_3(x,y)\right)}{\partial y} = \frac{\partial\left(g_2(x,y)/g_3(x,y)\right)}{\partial x}$$

(in practice, they should at least be similar)

We can now recover the surface height at any point by integration along some path, e.g.

$$f(x,y) = \int_0^x f_x(s,y)\,ds + \int_0^y f_y(x,t)\,dt + c$$

(for robustness, can take integrals over many different paths and average the results)

Forsyth & Ponce, Sec. 5.4

Surface recovered by integration

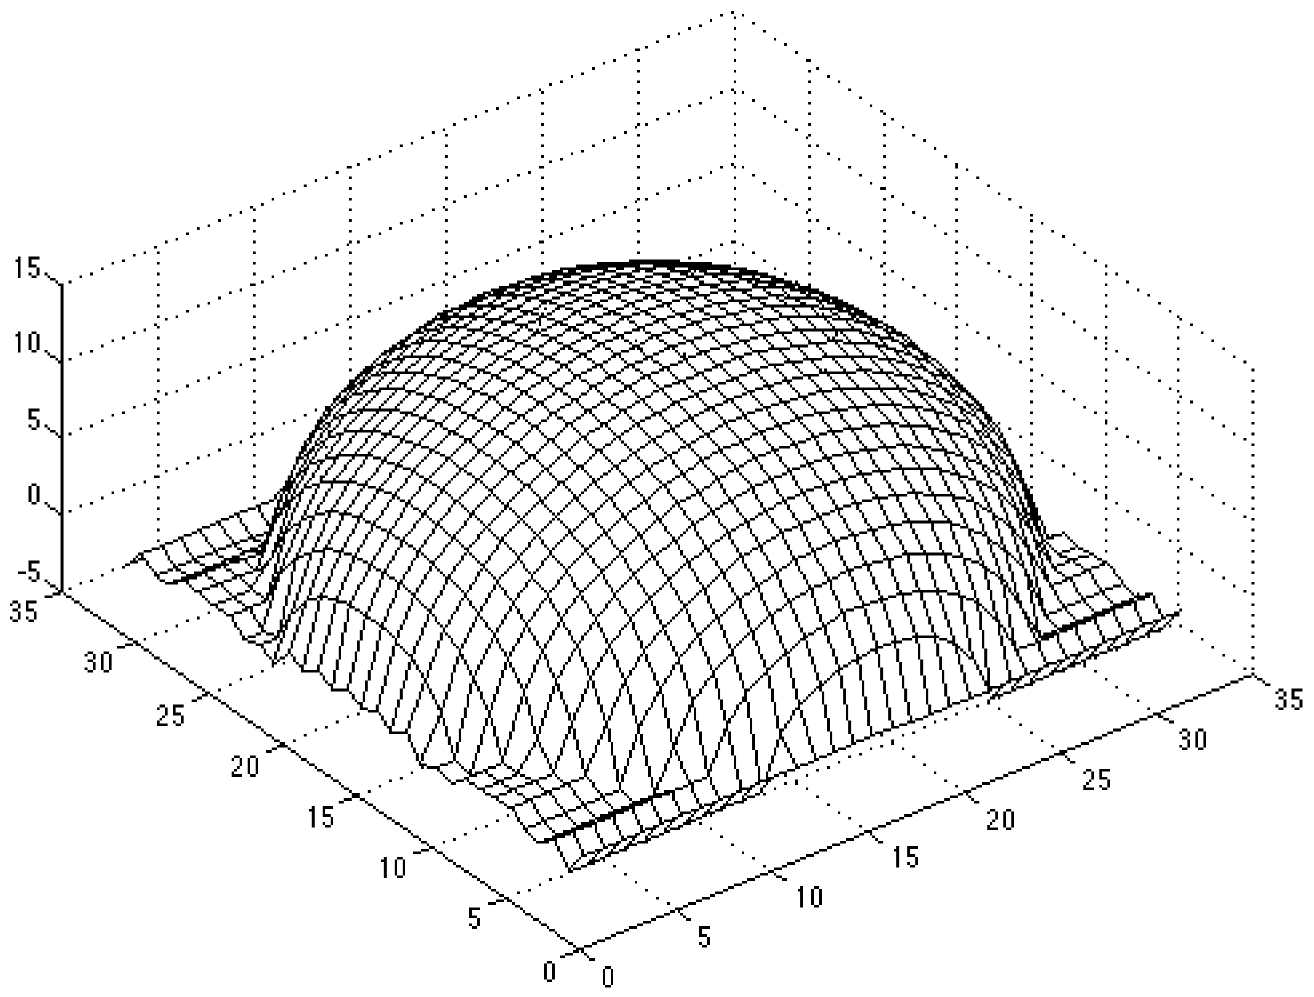

Forsyth & Ponce, Sec. 5.4

Limitations

- Orthographic camera model
- Simplistic reflectance and lighting model
- No shadows
- No interreflections
- No missing data
- Integration is tricky

Finding the direction of the light source

$$I(x,y) = N(x,y) \cdot S(x,y) + A$$

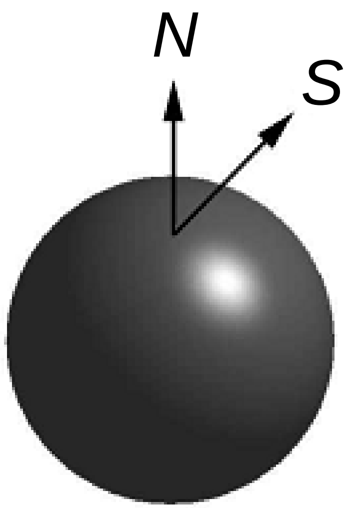

Full 3D case:

$$\begin{pmatrix} N_x(x_1,y_1) & N_y(x_1,y_1) & N_z(x_1,y_1) & 1 \\ N_x(x_2,y_2) & N_y(x_2,y_2) & N_z(x_2,y_2) & 1 \\ \vdots & \vdots & \vdots & \vdots \\ N_x(x_n,y_n) & N_y(x_n,y_n) & N_z(x_n,y_n) & 1 \end{pmatrix} \begin{pmatrix} S_x \\ S_y \\ S_z \\ A \end{pmatrix} = \begin{pmatrix} I(x_1,y_1) \\ I(x_2,y_2) \\ \vdots \\ I(x_n,y_n) \end{pmatrix}$$

For points on the *occluding contour*:

$$\begin{pmatrix} N_x(x_1,y_1) & N_y(x_1,y_1) & 1 \\ N_x(x_2,y_2) & N_y(x_2,y_2) & 1 \\ \vdots & \vdots & \vdots \\ N_x(x_n,y_n) & N_y(x_n,y_n) & 1 \end{pmatrix} \begin{pmatrix} S_x \\ S_y \\ A \end{pmatrix} = \begin{pmatrix} I(x_1,y_1) \\ I(x_2,y_2) \\ \vdots \\ I(x_n,y_n) \end{pmatrix}$$

P. Nillius and J.-O. Eklundh, “Automatic estimation of the projected light source direction,” CVPR 2001

Finding the direction of the light source

P. Nillius and J.-O. Eklundh, “Automatic estimation of the projected light source direction,” CVPR 2001

Application: Detecting composite photos

Real photo

Fake photo

Next time: Color

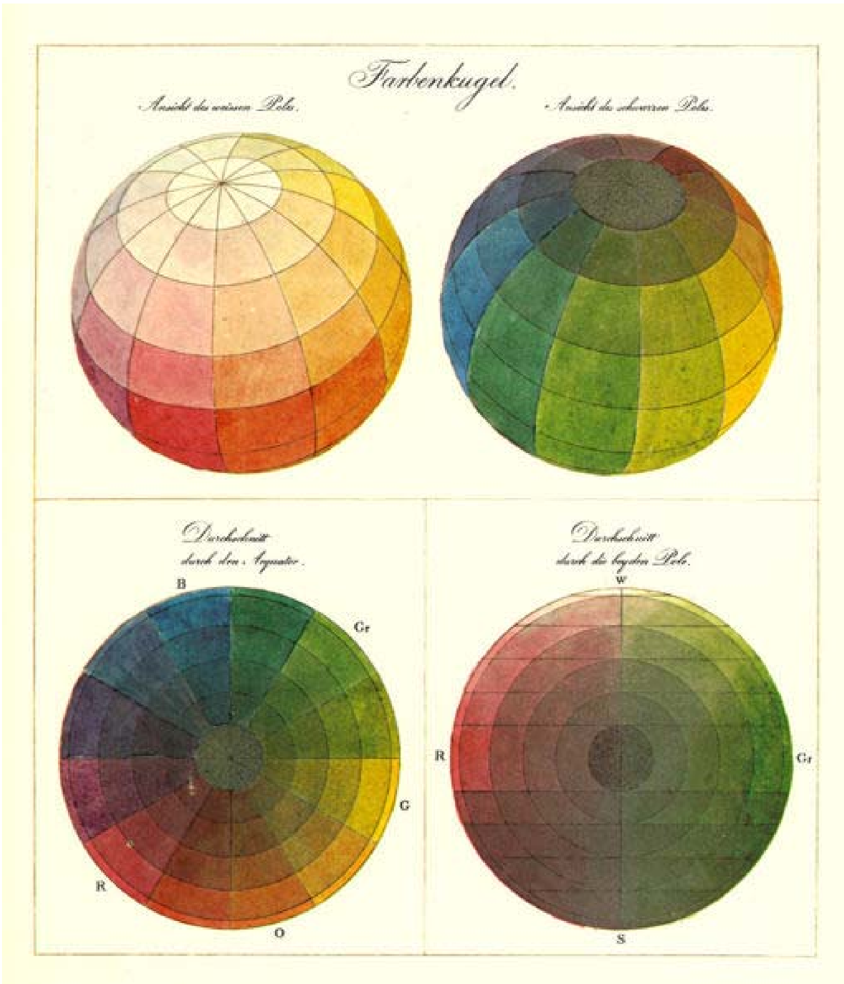

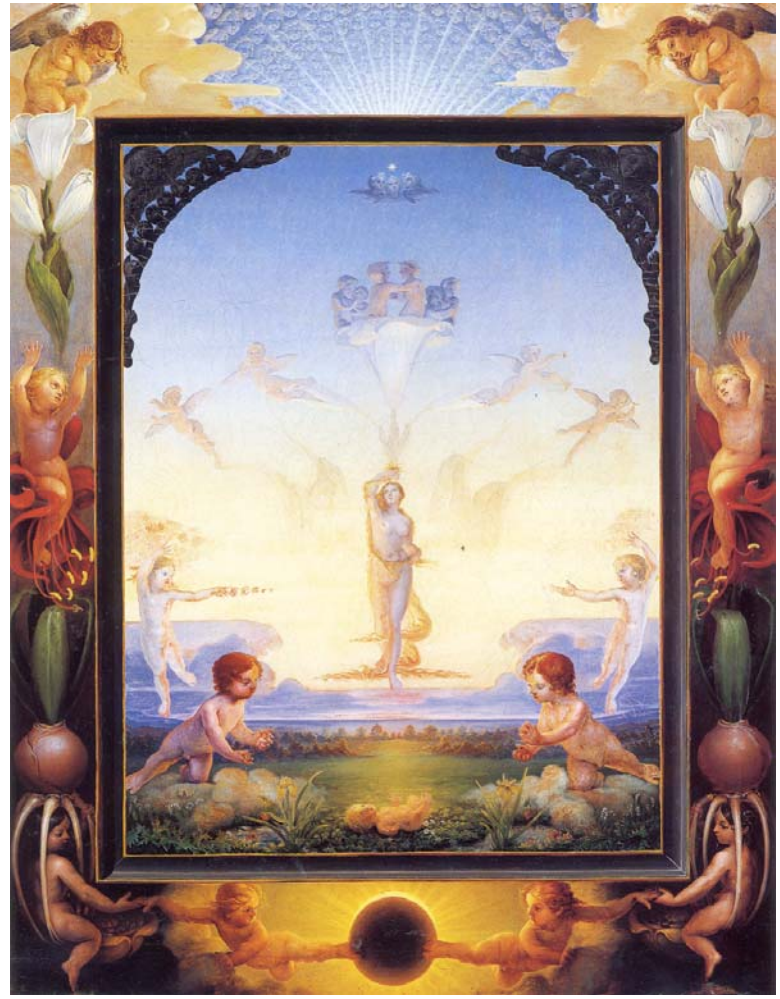

Phillip Otto Runge (1777-1810)