BRDF’s and Relighting

Topics in Image-Based Modeling and Rendering

CSE291 J00

Lecture 12

Environment Matte

Basic Assumption: Single ray into object, single ray out.

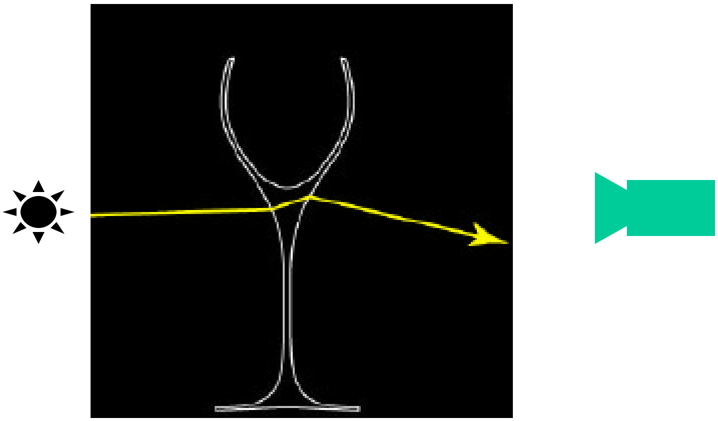

Background textures w/ foreground object

Background Textures

Note: Complex O(log k)

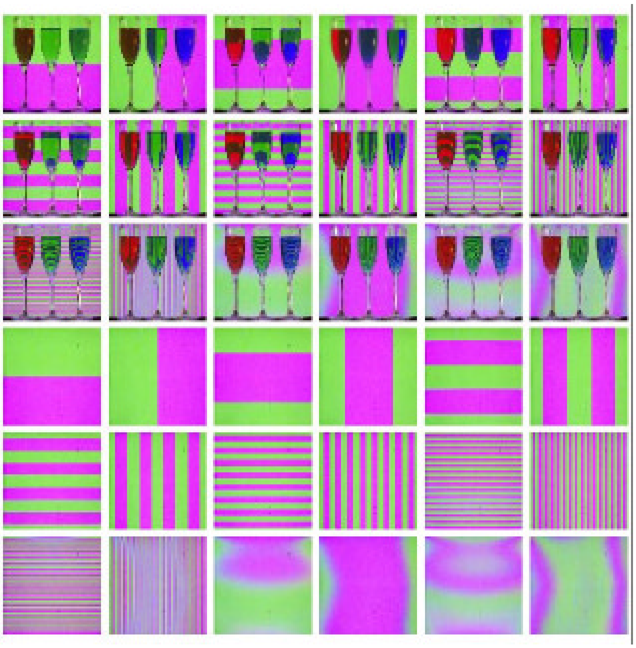

Explanation of Φ

$$F = \sum_{i=1}^{m} R_i M(T_i, A_i)$$

R_i: Reflectance coefficient

M: Texture mapping operator for axis-aligned rectangle (A_i) of texture (T_i)

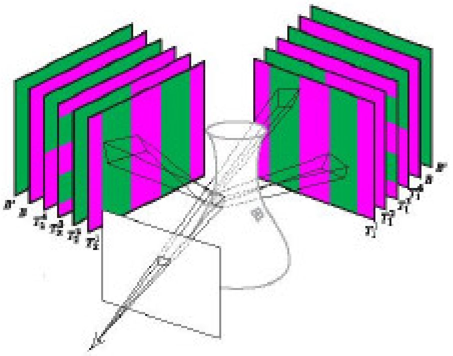

Environment Matte Example

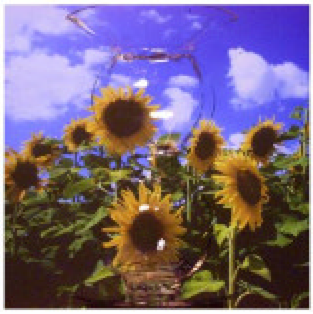

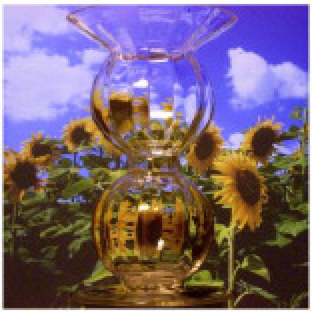

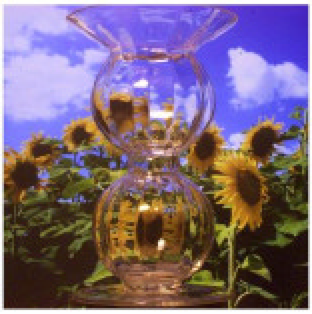

Alpha Matte **Environment Matte** **Photograph**

A composited object

BRDF

Surface Reflectance Models

Common Models

- Lambertian
- Phong
- Physics-based
  - Specular
    [Blinn 1977], [Cook-Torrance 1982], [Ward 1992]
  - Diffuse
    [Hanrahan, Kreuger 1993]
  - Generalized Lambertian
    [Oren, Nayar 1995]
  - Thoroughly Pitted Surfaces
    [Koenderink et al 1999]
- Phenomenological
  [Koenderink, Van Doorn 1996]

Arbitrary Reflectance

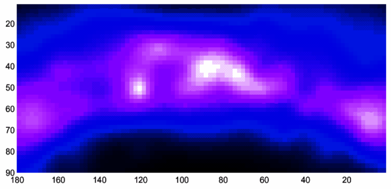

- Non-parametric model
- Anisotropic
- Non-uniform over surface
- BRDF Measurement *[Dana et al, 1999]*

BRDF

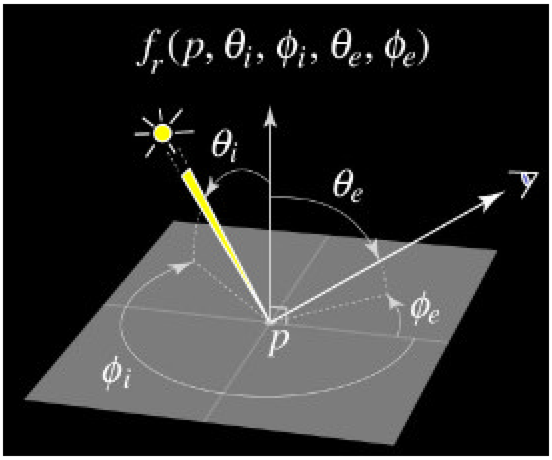

Adapted from Steve Marschner

Isotropic BRDF
Function of Three variables

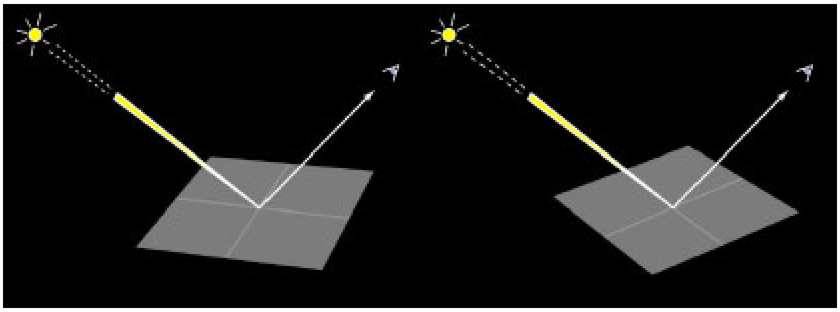

$f_r(\theta_i, \phi_\iota, \theta_e, \phi_e) = f_i(\theta_i, \theta_e, \phi_i - \phi_e)$

Adapted from Steve Marschner

BSSRDF

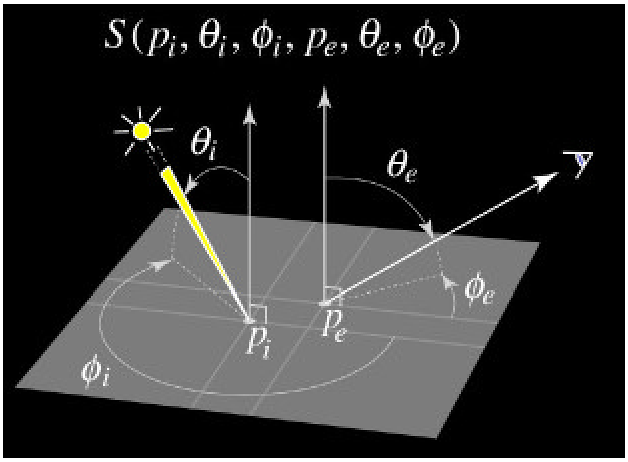

Bidirectional Subsurface Scattering Reflectance Distribution Function

Adapted from Steve Marschner

Off Specular Reflection

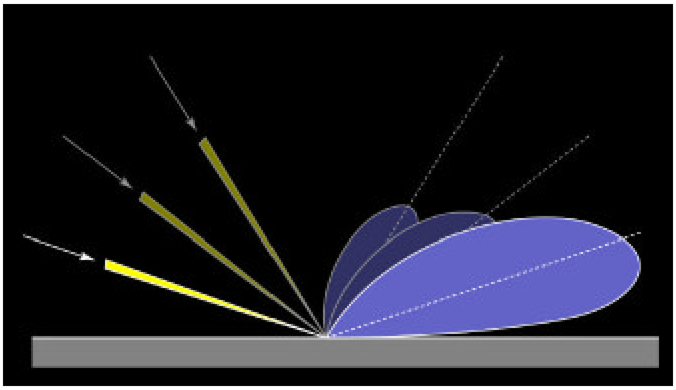

Adapted from Steve Marschner

Backscatter

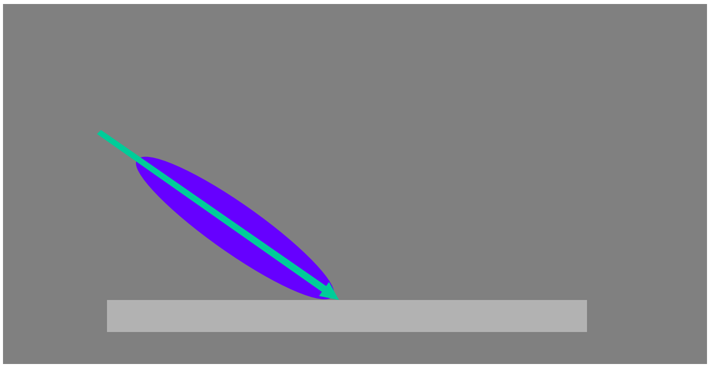

Adapted from Steve Marschner

BRDF

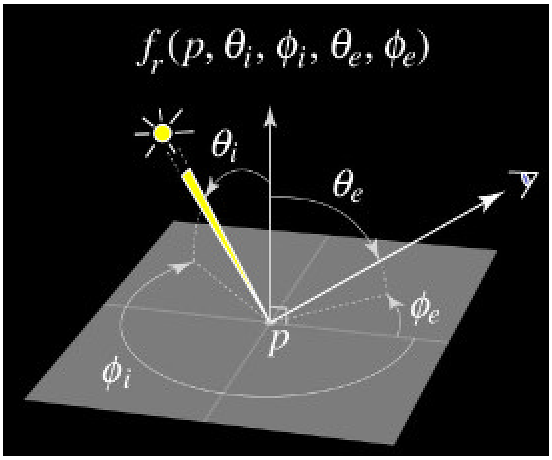

Adapted from Steve Marschner

BSSRDF

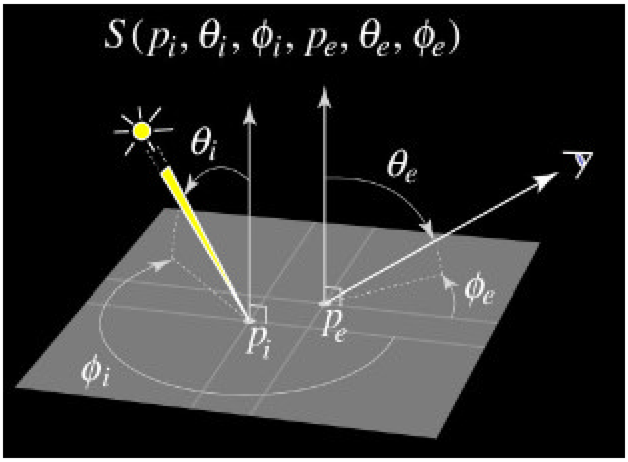

Bidirectional Subsurface Scattering Reflectance Distribution Function

Adapted from Steve Marschner

Materials: Conductor

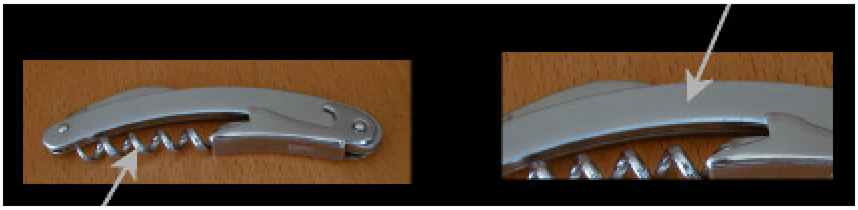

Conductor

Conductor + Microgeometry

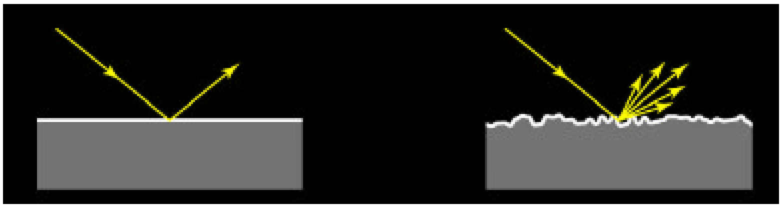

Material: Insulator

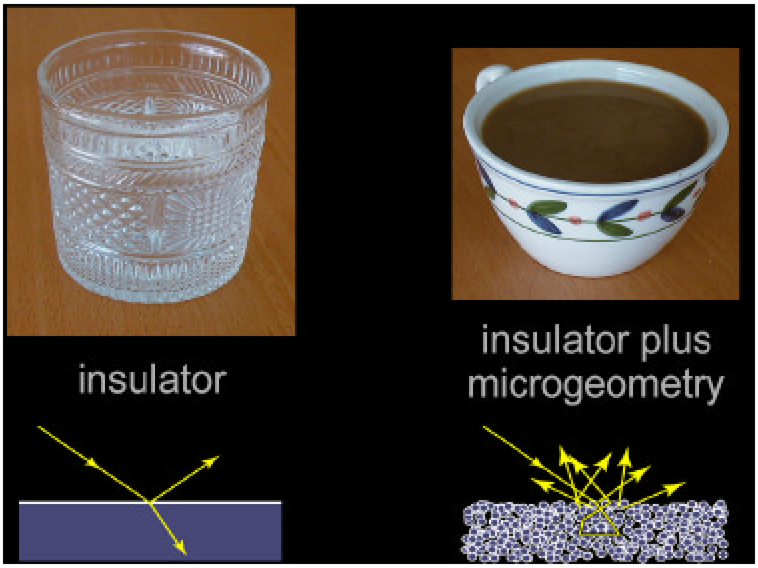

Adapted from Steve Marschner

Measured BRDFs

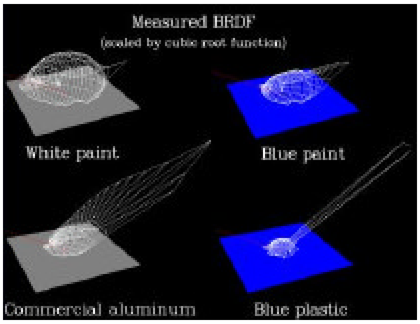

BRDF cross-sections

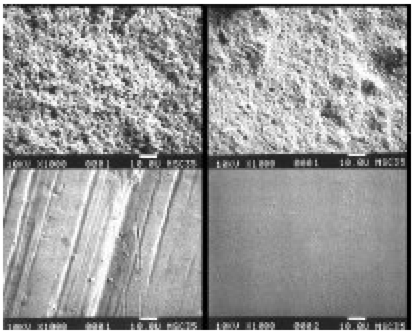

Surface microstructure

Ward reflectance model

- A physically realizable variant of the Phong model (satisfies energy conservation and reciprocity).

[redacted]

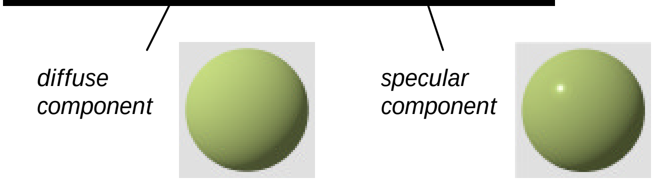

  - ρ_d: proportion of incident radiation reflected diffusely.
  - ρ_s: proportion of incident radiation reflected specularly.
  - α: surface roughness, or blur in specular component.

From Ron Dror's slides

Cook-Torrance Model (1982)

- Diffuse (Lambertian) and Specular and Fresnel reflection
- Microfacet model – surface is modeled as a collection of parallel symmetric V-groves called microfacets (facets are large w.r.t. wavelength, small w.r.t. pixel size).

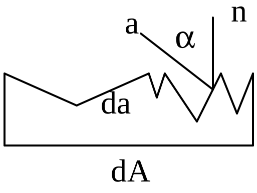

[redacted]

- Facet distribution is given by a specific distribution (e.g., Gaussian, $D = k\exp(\alpha/m)^2$
- Facets are purely specular

Geometric Attenuation: Masking and Shadowing

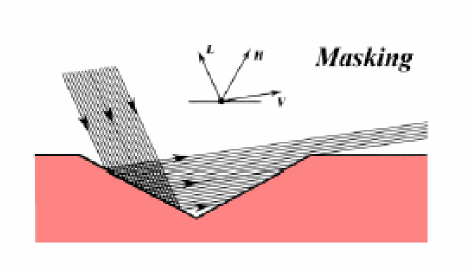

Geometric Attenuation

G=1-L1/L2

[redacted]

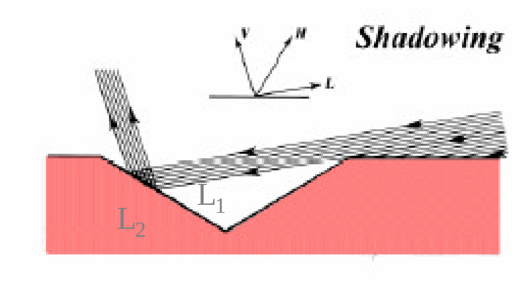

[redacted]

Fresnel Equation for Polished Copper

[redacted]

Reflectance as Function of Angle of Incidence for Copper

Blue Green Red

Generalized Lambertian Model (Oren, Nayar 1994)

- Like Torrance-Sparrow, but with Lambertian facets.
- Intensity doesn't fall of as quickly as function of incident illumination.

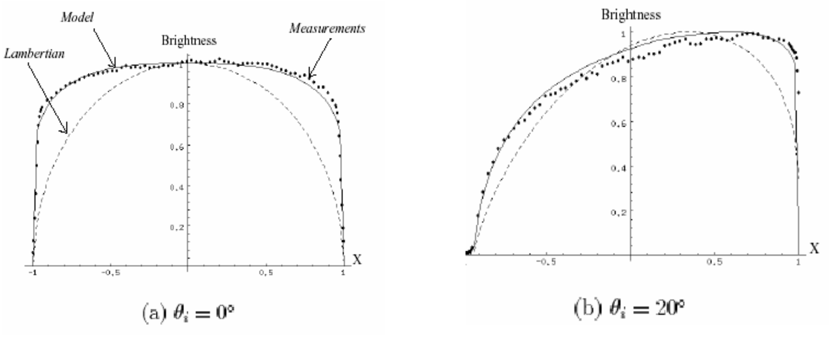

(a) $\theta_i = 0°$

(b) $\theta_i = 20°$

Rough Cylindrical Clay Vase

Velvet: A general BRDF

Portrait of Sir Thomas More, Hans Holbein the Younger, 1527

[After Koenderink et al, 1998]

Measuring Isotropic BRDF's

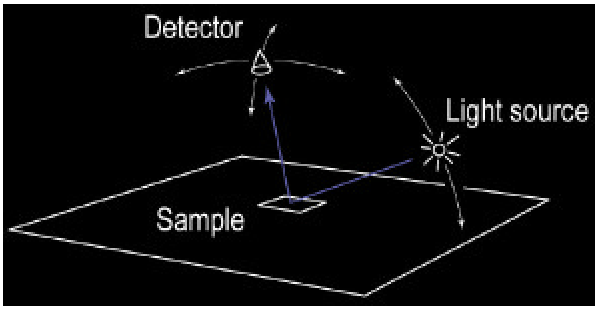

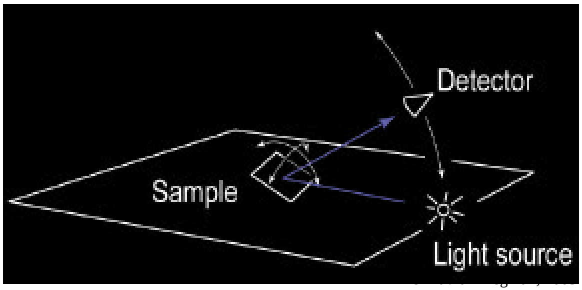

Adapted from Steve Marschner

Image-based: Marschner

- Known Geometry of Sample,
- From single image, one obtains 2-D slice of BRDF.

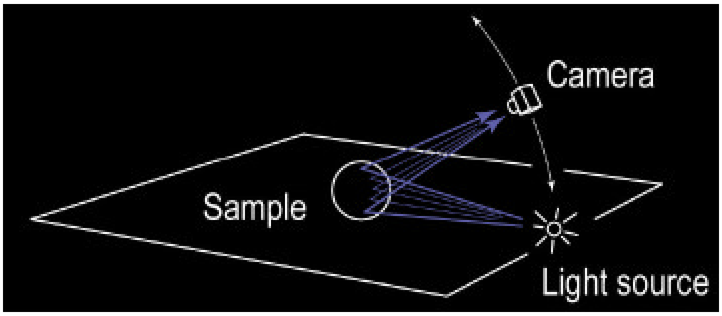

Adapted from Steve Marschner

Known illumination: inverse rendering

- If one assumes that illumination and surface geometry are known in advance, one can recover samples of the BRDF from an image (Marschner; Sato & Ikeuchi).

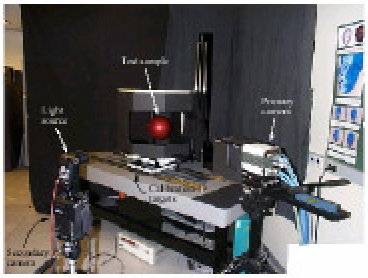

From Ron Dror's slides

Empirical BRDF’s

Consider a collection of basis functions b_j (θ_i, ϕ_i, θ_e, ϕ_e)

Represent an arbitrary BRDF as

$$r(q_i, f_i, q_r, f_r) = \sum_j w_j b_j(q_i, f_i, q_r, f_r)$$

Given measurements, estimate w_j to fit BRDF to data.

What is a good set of basis functions?

1. Product of spherical harmonics
2. Wavelets
3. Zernike polynomials

Spherical Harmonics

(Basri, Jacobs’01; Ramamoorthi, Hanrahan’01)

- A set of orthonormal basis functions defined on the unit sphere
- The first nine harmonics: $Y_{lm}(q, j)$

l=0

- [redacted]CS:

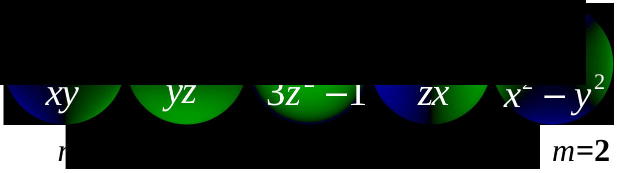

.

. [redacted] *m*=**2**

. (Borrowed from: Ramamoorthi, Hanrahan, SIGGRAPH’01)

Phenomenological BRDF model Zernike Polynomials

[Koenderink & van Doorn, 1996]

- A problem with spherical harmonics, half of sphere should be zero.
- General compact representation defined on disk
- Preserve Helmholtz Reciprocity
- Preserve reciprocity/isotropy if desired
- Domain is product of hemispheres
- Same topology as unit disk, adapt basis

Zernike Polynomials

- Optics, complete orthogonal basis on unit disk using polynomials of radius

$$\boxed{Z_n^m(r,f)=\sqrt{\frac{n+1}{p}}R_n^m(r)e^{imf}}$$

n-|m| even
|m|≤n

- R has terms of degree at least m. Even or odd depending on m even or odd
- Orthonormal, using measure ρdρdφ

Cool Demo: http://wyant.opt-sci.arizona.edu/zernikes/zernikes.htm

|m| —— 0 ■ 2

n
|

0

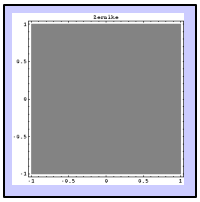

|m| ≤ n

n-|m| must
be even

|m| ≤ n

1

n-|m| must
be even

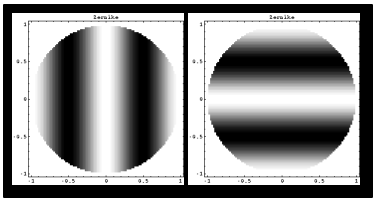

|m| ≤ n

n-|m| must
be even

2

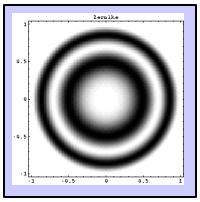

n-|m| must
be even

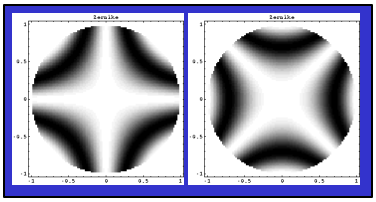

“Relighting”

1. Steerable lighting
2. Lambertian Surfaces and linear subspaces
3. Arbitrary BRDF, arbitrary lighting

Superposition of lighting: An important point

If $I_1 = R(S, L_1)$ is the image of a scene S under lighting L_1
and if $I_2 = R(S, L_2)$ is the image of a scene S under lighting L_2,

then the image of the scene under lighting L_1+L_2 is simply

$$I_1 + I_2$$

Basis functions: Example

Consider sinusoid of frequency f, we can specify any sinusoid as sum of two basis sinusoids cos(ft) and sin(ft) as:

$$a \cos(ft) + (1-a) \sin(ft)$$

Such basis function are sometimes called steerable functions – over some transformation of the parameter (e.g. t here, (x,y) for image plane), the function can be represented as the linear combination of a finite collection of basis functions.

Steerable lighting for relighting

1. Choose a steerable basis for the lighting – all rendering lighting conditions will defined as linear combinations of the basis lighting.
2. Gather (or synthesize) images of a scene under the basis lighting.
3. Render new images by taking linear combinations of basis images.

Why does it work? Superposition

Lambertian Surface: $\rho(\theta_{in}, \phi_{in}; \theta_{out}, \phi_{out})$=constant

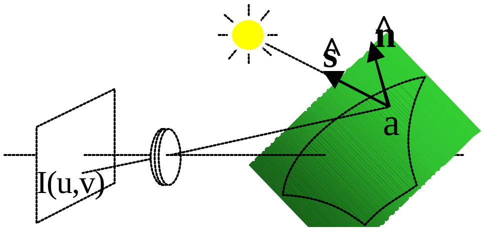

At image location (u,v), the intensity of a pixel I(u,v) is:

$$\begin{aligned} I(u,v) &= [a(u,v)\hat{\mathbf{n}}(u,v)]\cdot[s_0\hat{\mathbf{s}}] \\ &= \mathbf{b}(u,v)\cdot\mathbf{s} \end{aligned}$$

where

- **a(u,v) is the albedo of the surface projecting to (u,v).**
- **$\hat{n}$(u,v) is the direction of the surface normal.**
- **s_0 is the light source intensity.**
- **$\hat{s}$ is the direction to the light source.**

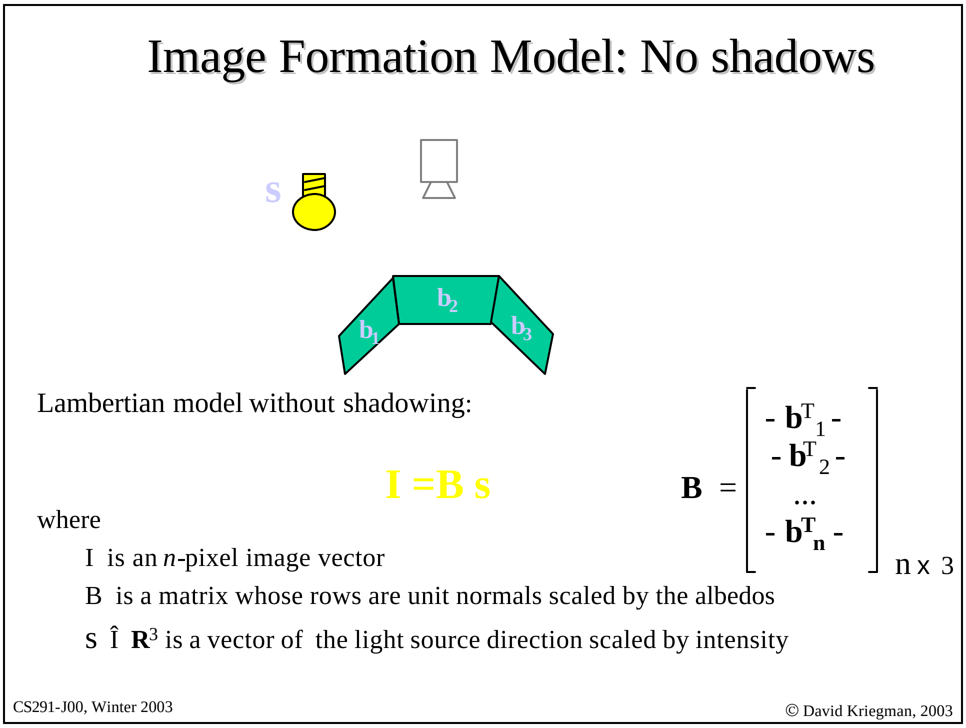

Image Formation Model: No shadows

Lambertian model without shadowing:

$$I = B\,s$$

$$B = \begin{bmatrix} - \mathbf{b}^T_1 - \\ - \mathbf{b}^T_2 - \\ \dots \\ - \mathbf{b}^T_n - \end{bmatrix}_{n \times 3}$$

where

I is an *n*-pixel image vector

B is a matrix whose rows are unit normals scaled by the albedos

s Î $\mathbf{R}^3$ is a vector of the light source direction scaled by intensity

CS291-J00, Winter 2003

© David Kriegman, 2003

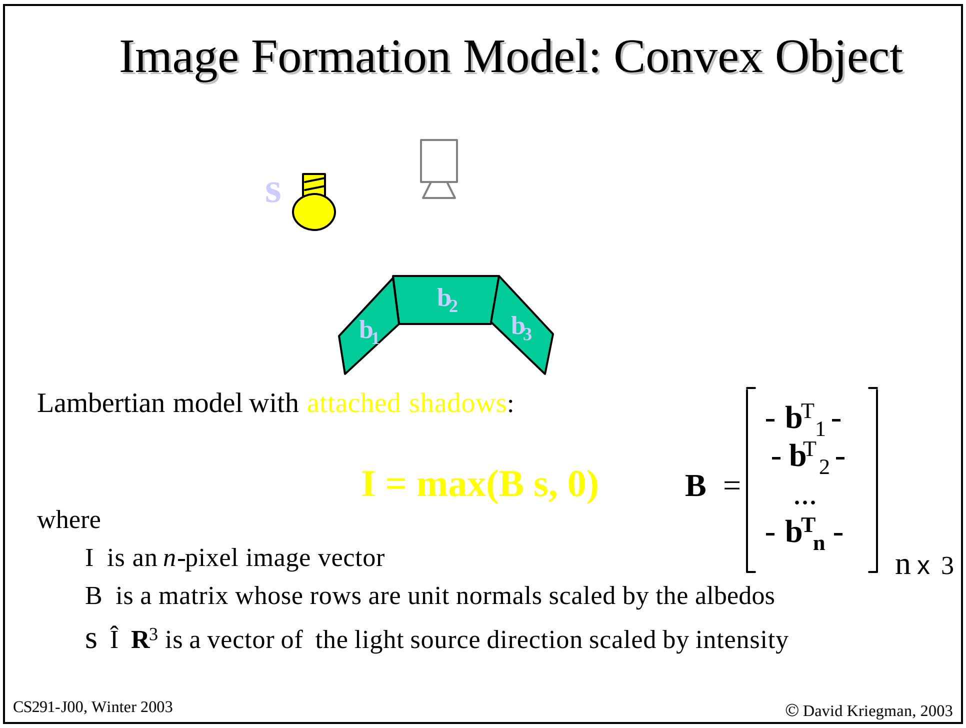

Image Formation Model: Convex Object

Lambertian model with attached shadows:

$$I = \max(B\,s, 0)$$

$$B = \begin{bmatrix} - \mathbf{b}^T_1 - \\ - \mathbf{b}^T_2 - \\ \dots \\ - \mathbf{b}^T_n - \end{bmatrix}_{n \times 3}$$

where

I is an *n*-pixel image vector

B is a matrix whose rows are unit normals scaled by the albedos

s Î $\mathbf{R}^3$ is a vector of the light source direction scaled by intensity

CS291-J00, Winter 2003

© David Kriegman, 2003

Illumination Subspace

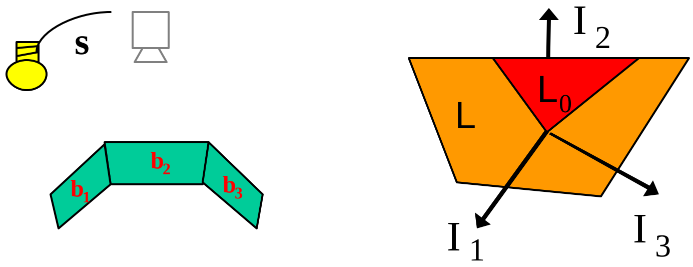

$$L = \{\mathbf{I} \mid \mathbf{I} = B\mathbf{s}, \text{ for all } \mathbf{s} \in R^3\}$$

- L is a 3-D linear subspace of image space, $\mathbf{R}^n$.
- L is spanned by 3 linearly independent images.
- Cone can be generated from L.
- See also [Woodham 81], [Shashua 92], [Hallinan 95], [Hayakawa 94], [Rosenholtz, Koenderink 96]

Multiple Sources: No shadows

- Consider two sources
- Light source 1: $I_1 = Bs_1$
- Light source 2: $I_2 = Bs_2$
- Image with both lights on:

$$I=I_1+ I_2 = Bs_1 + Bs_2 = B(s_1 + s_2)$$

So, what does this mean?

Computing **L**

For k images $X = [\ x_1,\ x_2, \ldots, x_k\ ]$ imaged under k unknown point light sources $S = [s_1, s_2, \ldots, s_k]$,

$$X = B\ S$$

Given $k \geq 3$ images we can compute B* that spans **L** with

1. singular value decomposition
2. or methods robust to outliers.

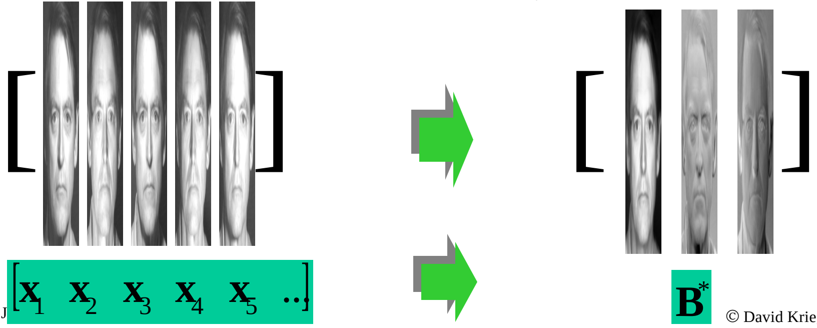

$[\mathbf{x}_1\ \ \mathbf{x}_2\ \ \mathbf{x}_3\ \ \mathbf{x}_4\ \ \mathbf{x}_5\ \ \ldots]$
$\mathbf{B}^*$

Still Life

Original Images

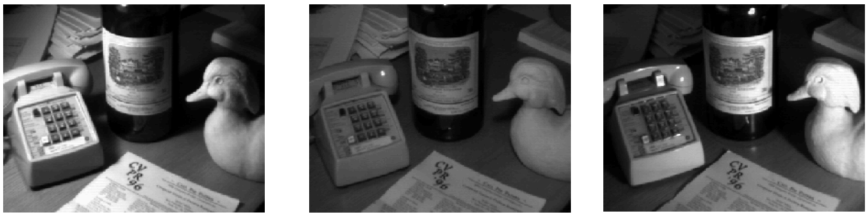

Basis Images

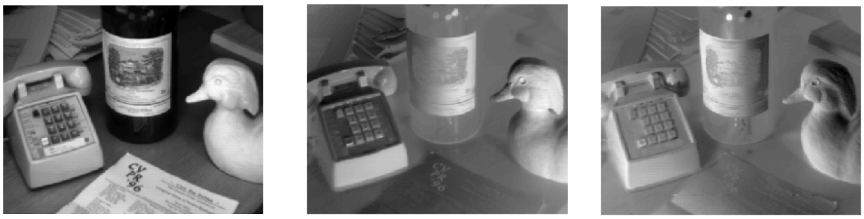

Samples from the Cone

1 Light **2 Lights** **3 Lights**

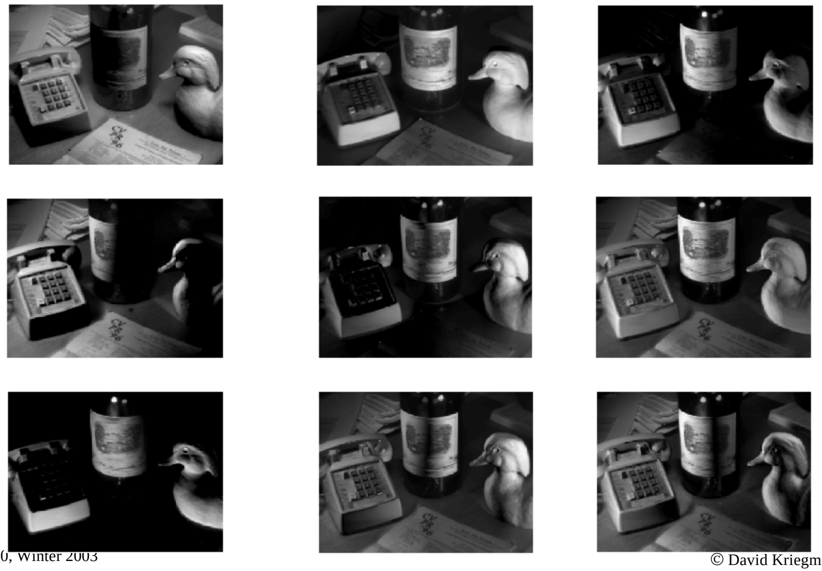

Face Basis

Original Images

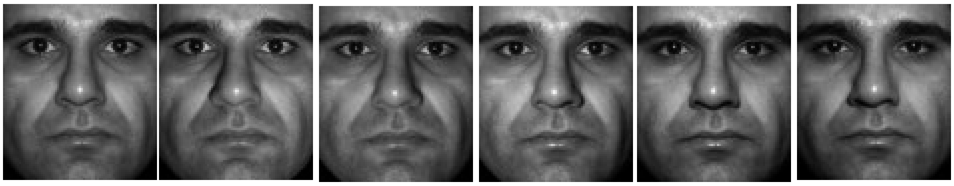

Basis Images spanning L

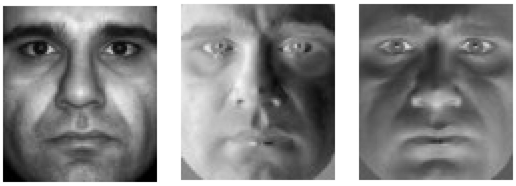

Image-Based Rendering: Attached Shadows

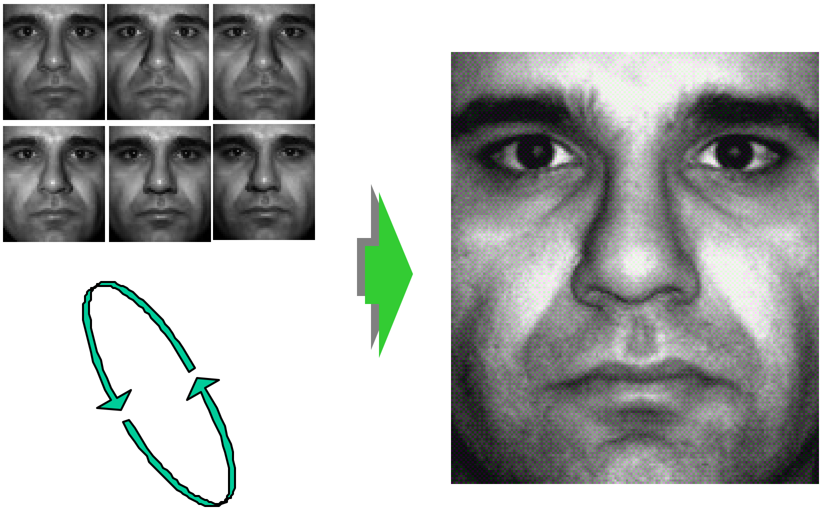

Single Light Source

Face Movie

Lumigraph/Light Field Rendering

[Levoy, Hanrahan, 1996]
[Gortler et al, 1996]

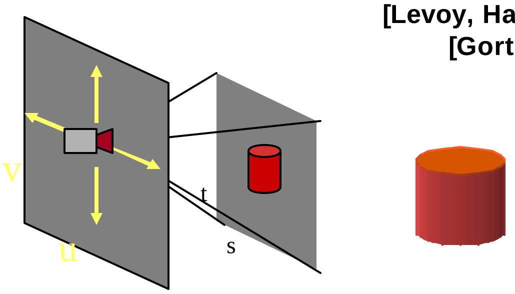

- Set of light rays is a 4-D manifold.
- Single camera image provides a 2-D sampling of the real scene's radiance.
- Moving the camera over a 2-D surface yields a sampling of 4-D radiance (light) field L(u,v,s,t).
- Assume lighting is fixed.

Lumigraph/Light Field Rendering

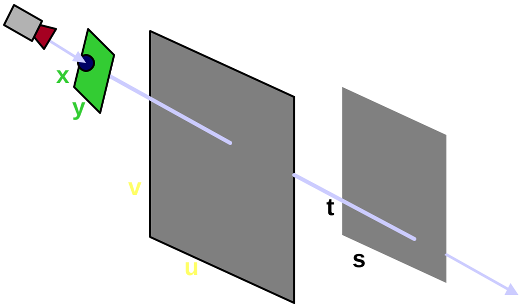

Image from new viewpoint but under same lighting is rendered by indexing each visual ray into the Lumigraph

$$I(x,y) = L(u,v,s,t)$$

where u,v,s,t are functions of (x,y).

A not so good idea

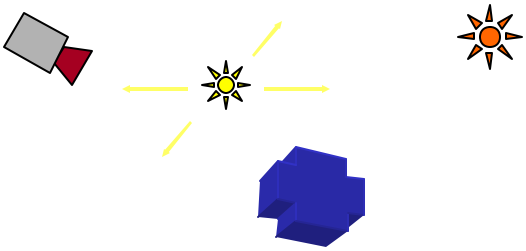

No and Yes

- Scene depth is needed.
- Depth reconstruction method.
- Rendering method.

Can we synthesize images under novel lighting by indexing into the image set?

Why is depth needed? Correspondence

- Isotropic point light source in known position
- Calibrated camera

- Which pixel corresponds to which light source ray?
- Which light source ray indirectly illuminates which pixel?

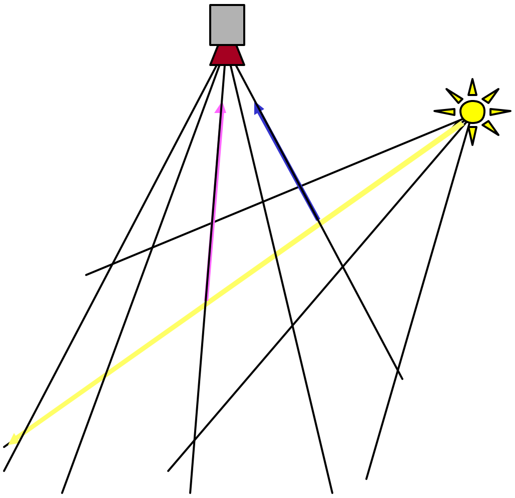

A 2-D schematic

Why is depth needed? Correspondence

- Isotropic point light source in known position.
- Calibrated camera

- Which pixel corresponds to which light source ray?
- Which light source ray indirectly illuminates which pixel?

Depth resolves this.

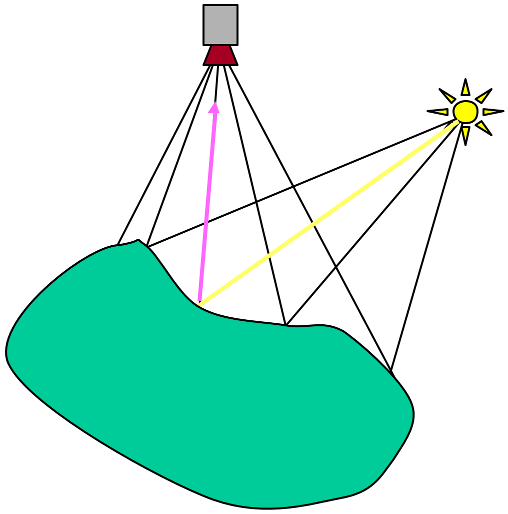

A 2-D schematic

Scene Depth

We would like to recover the depth l :

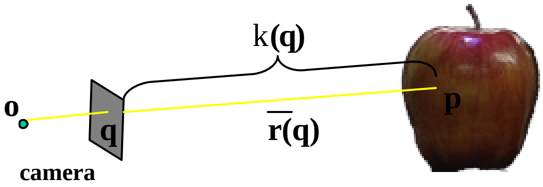

$$\mathbf{p}(k) = \mathbf{o} + l\,\overline{\mathbf{r}}$$

Reconstruction from the Illumination Field

1. **Consider fixed camera and point light source**

2. **Light moves over a star-shaped surface**

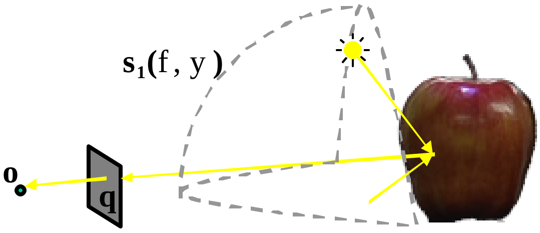

Intensity $I_1(f, y)$ as function of (q, y)

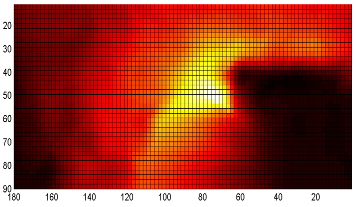

Intensity of One Pixel: $\mathbf{I_1}(f_1, y_1)$

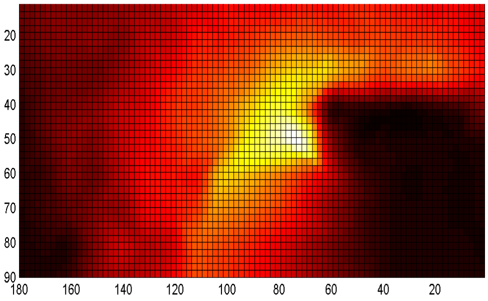

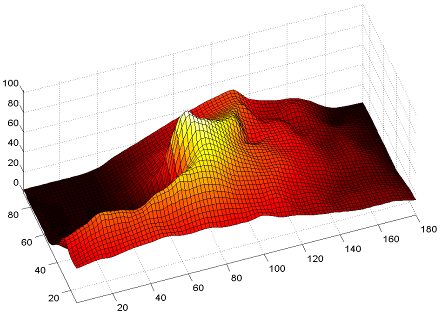

This is effectively a 2-D slice of a surface point's BRDF except for

- **Shadowing**
- **Variation in the distance between the sources and the surface point.**

Double Covering of the Illumination Field

Consider the effect of moving the light over a second surface:

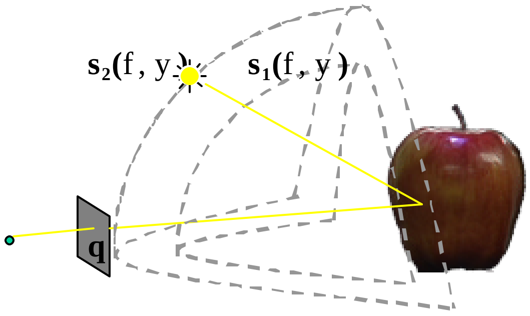

Intensity $I_2(f, y)$ as a function of (f, y)

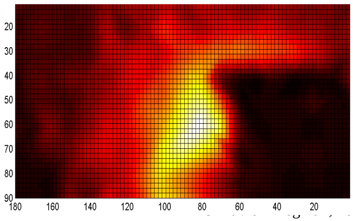

I_1(f $_1$, y $_1$) and I_2(f $_2$, y $_2$)

Inner Sphere: I_1(f $_1$, y $_1$)

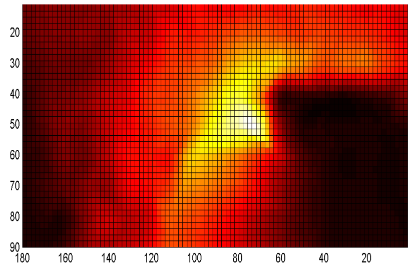

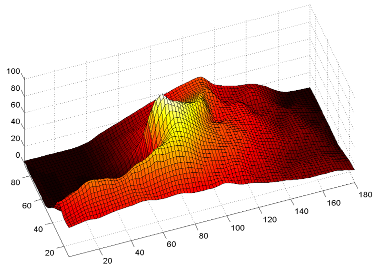

Outer Sphere: I_2(f $_2$, y $_2$)

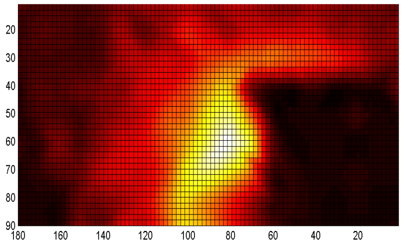

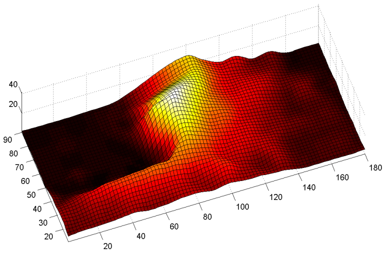

Image Acquisition

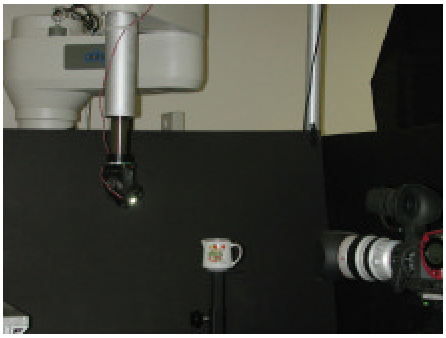

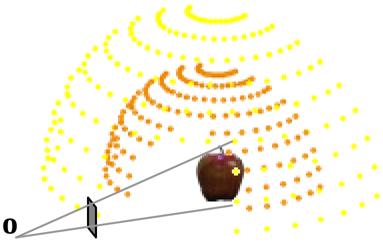

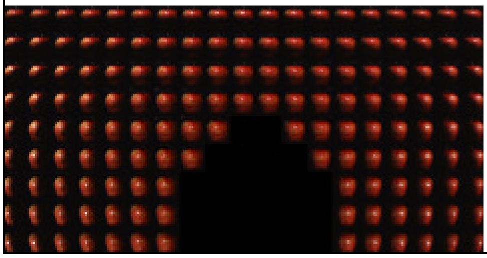

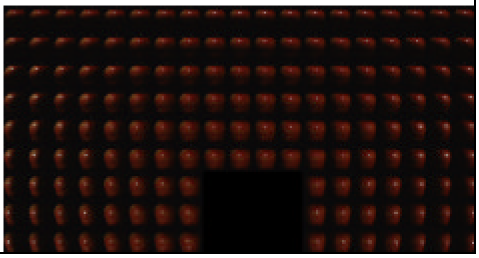

Relation Between Intensity Maps

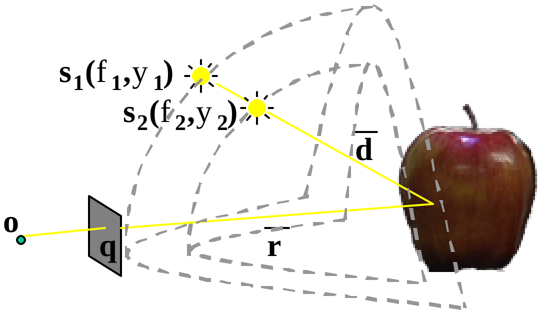

When the surface point p, s_1(f $_1$, y $_1$) and s_2(f $_2$, y $_2$) are collinear (in correspondence), the measured pixel intensities are simply related by the relative $1/r^2$ losses.

Depth Estimation

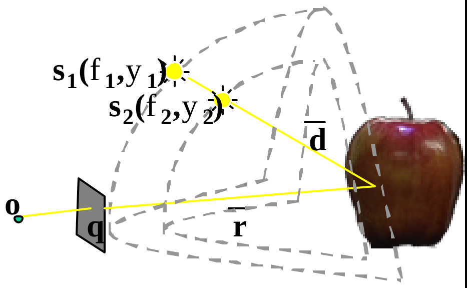

This correspondence can be expressed as a change of coordinates f $_2$(f $_1$, y $_1$; l) and y $_2$(f $_1$, y $_1$;l) parameterized by the depth l. We can then estimate l by minimizing

$$O(l)=\int\int[I_2(f_2(f_1,y_1;l),y_2(f_1,y_1;l))-d^2(\mathbf{p}(l))I_1(f_1,y\,)]^2\,df_1dy_1$$

where

$$d^2(\mathbf{p}(l))=\frac{\|\mathbf{p}(l)-\mathbf{s}_1\|^2}{\|\mathbf{p}(l)-\mathbf{s}_2\|^2}$$

An apple and its depth map

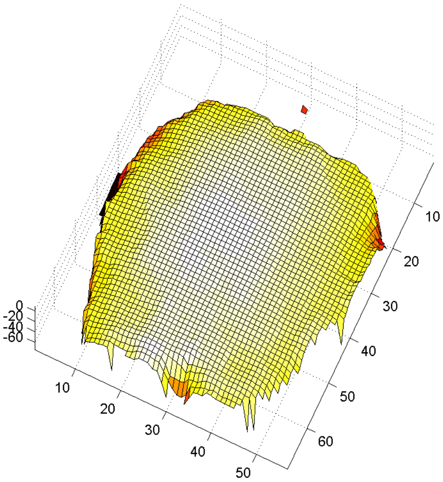

A Reconstructed Depth Map

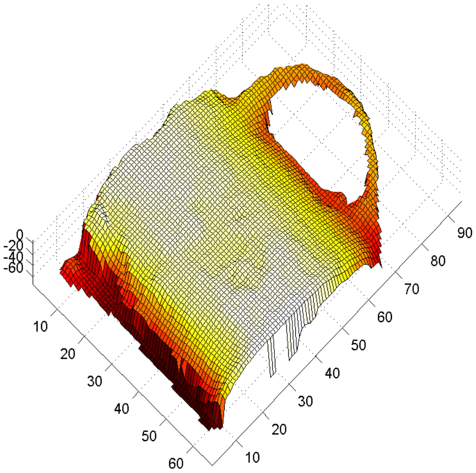

143 Images on each surface

Rendering Synthetic Images

Point Source Example

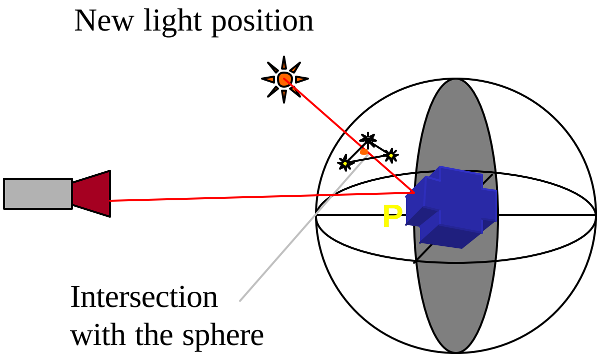
New light position

Intersection with the sphere

- For a given image point, there is a scene point: P
- Intersect light ray through P with sphere.
- Find triangle of light sources containing P.
- Interpolate pixel intensities of images corresponding to the triangle vertices & scale by $1/r^2$

Rendered Images

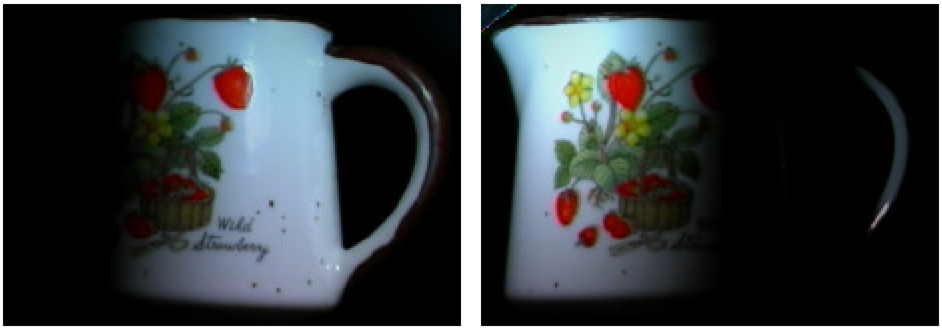

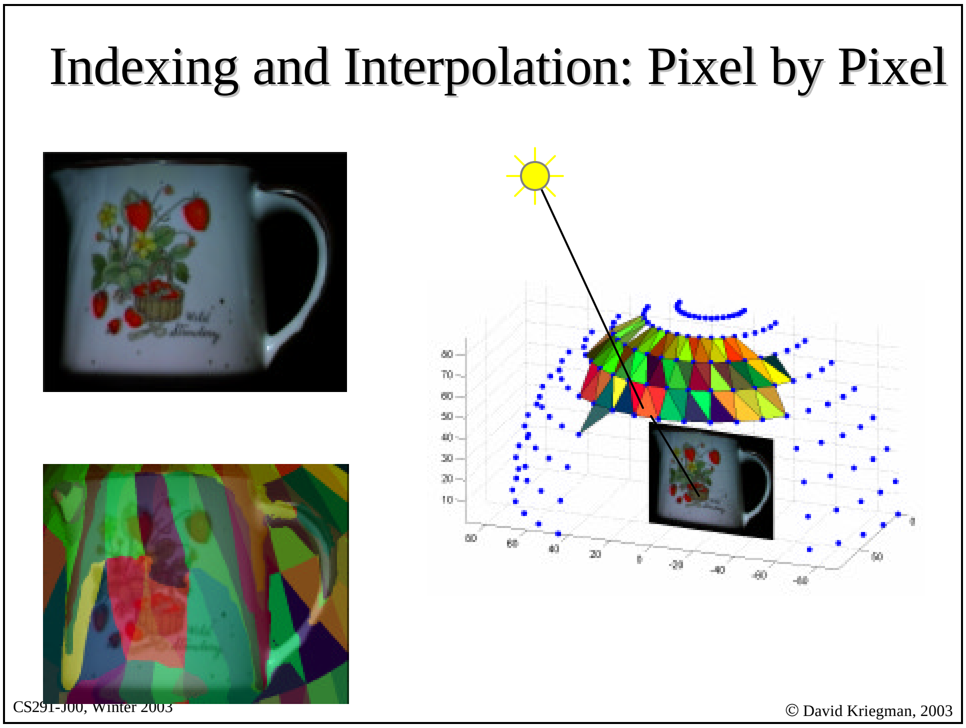

Indexing and Interpolation: Pixel by Pixel

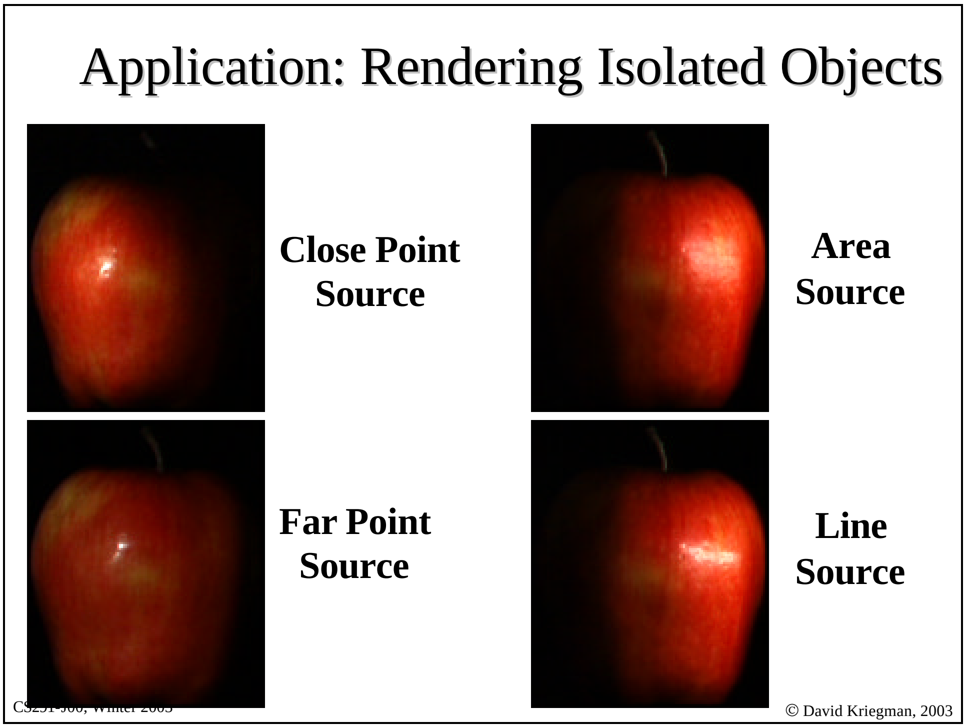

Application: Rendering Isolated Objects

Close Point Source

Area Source

Far Point Source

Line Source

Application: Rendering Isolated Objects

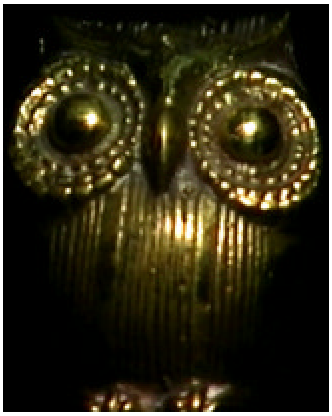

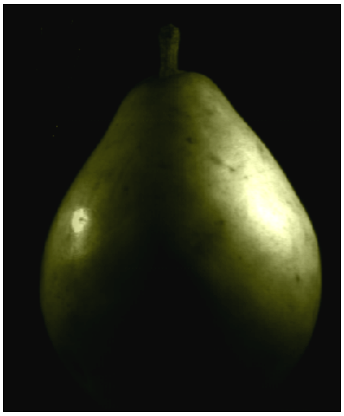

Point Source
(not on captured surface)

Multiple Sources -
Area Source

Rendered Image: A Sea Shell

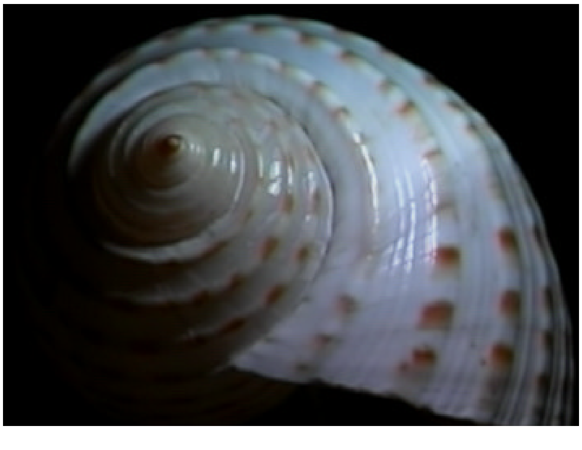

Isotropic point light source located between acquisition spheres.

Moving Light source

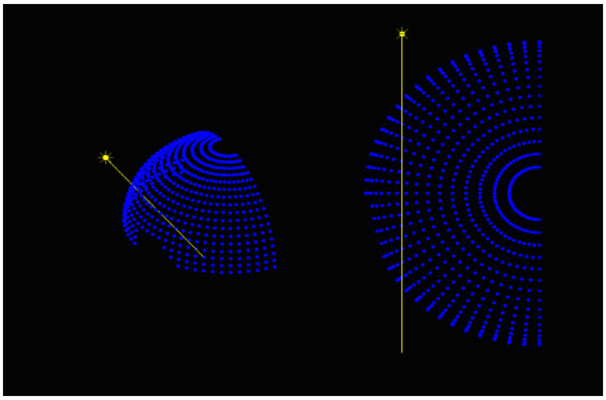

Moving Light source

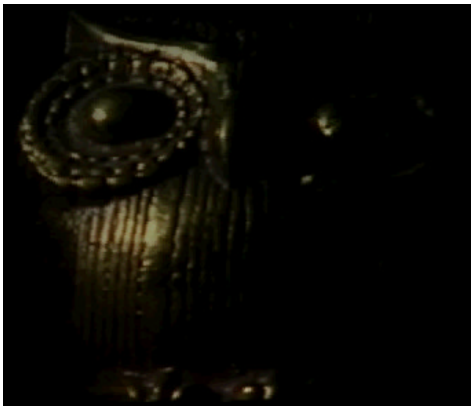

A Moving Light Source

A comparison

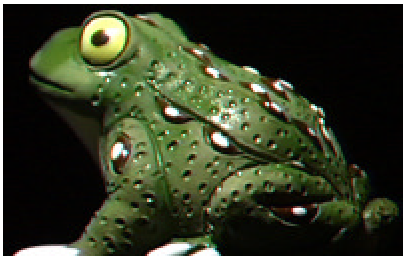

Real Image

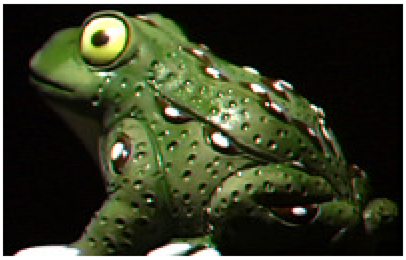

Rendered Image

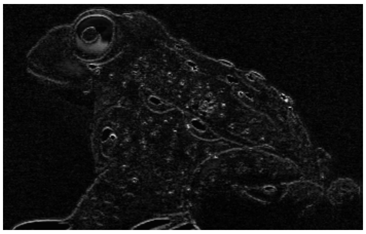

Magnitude of difference Image

Embedding Objects in Synthetic Scenes

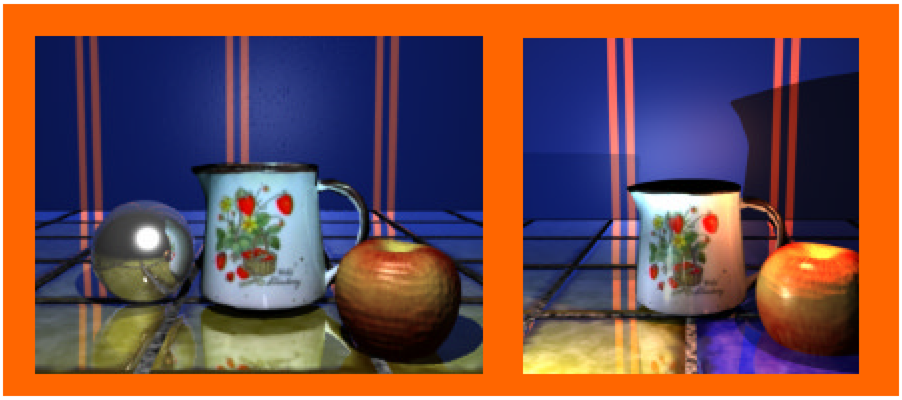

- Blue Moon Rendering Tools to render scene
- Custom surface shader to implement rendering method by indexing the illumination dataset