Image cues

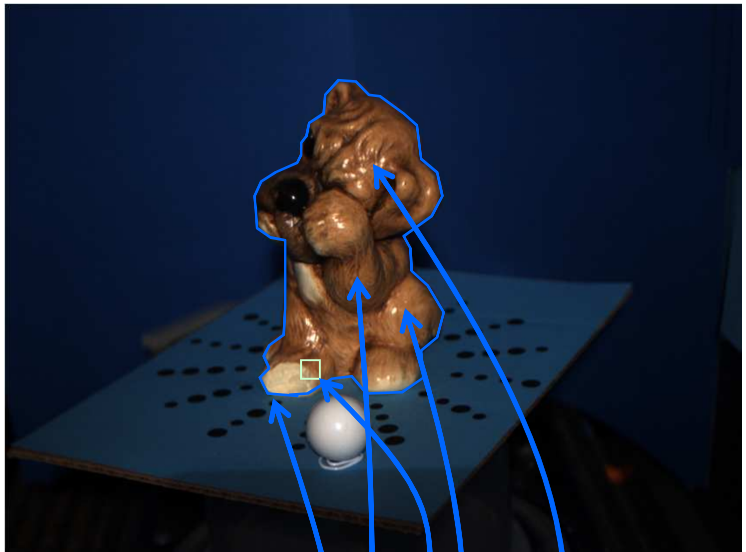

Color (texture)
Shading
Shadows
Specular highlights
Silhouette

Image cues

Shading [reconstructs normals]

shape from shading (SFS)

photometric stereo

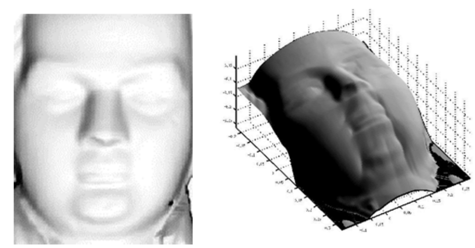

Specular highlights

[ignore, filtered]

[parametric BRDF]

Texture [reconstructs 3D]

stereo (relates two views)

Silhouette [reconstructs 3D]

shape from silhouette

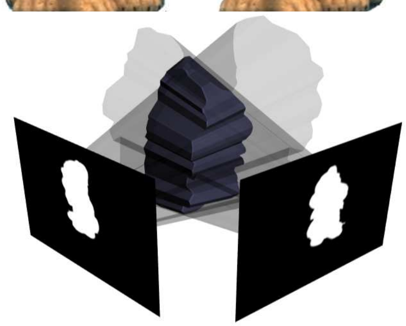

[Focus]

Geometry from shading

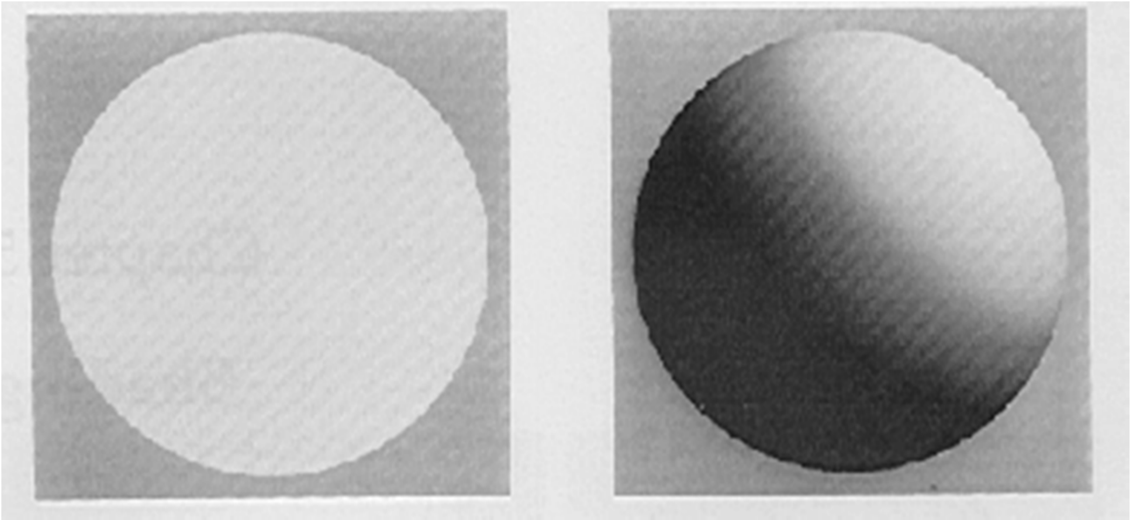

Shading reveals 3D shape geometry

Shape from Shading
One image
Known light direction
Known BRDF (unit albedo)
Ill-posed : additional constraints (intagrability …)

[Horn]

Photometric Stereo
Several images, different lights
Unknown Lambertian BRDF
1. Known lights
2. Unknown lights

[Silver 80, Woodman 81]

Reconstruct normals
Integrate surface

Shading

Lambertian reflectance

$$E(\mathbf{x}) = \rho L_i(\mathbf{x}, \theta_i, \phi_i)\cos\theta_i = \rho(\mathbf{n} \bullet \mathbf{l}_i)$$

albedo

normal

light dir

Fixing light, albedo, we can express reflectance only as function of normal.

Surface parametrization

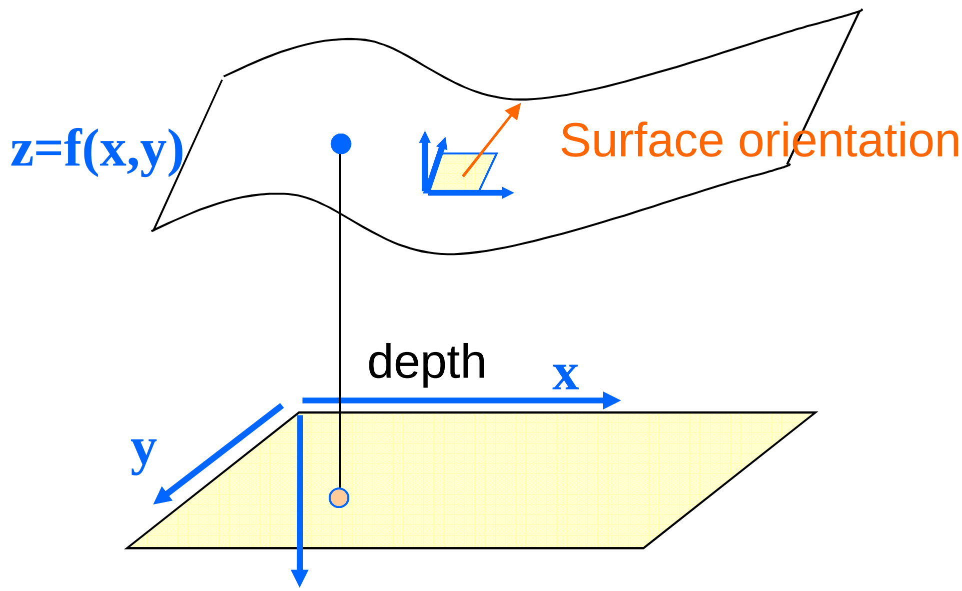

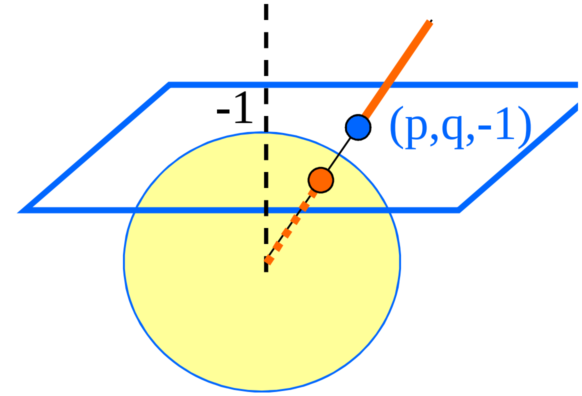

Surface

$$s(x,y) = (x, y, f(x,y))$$

Tangent plane

$$\frac{\partial s}{\partial x} = \left(1,0,\frac{\partial f}{\partial x}\right)^T \qquad \frac{\partial s}{\partial y} = \left(0,1,\frac{\partial f}{\partial y}\right)^T$$

Normal vector

$$\mathbf{n} = \frac{\partial s}{\partial x} \times \frac{\partial s}{\partial y} = \left(\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y}, -1\right)^T$$

Gradient space

$$p = \frac{\partial f}{\partial x} \qquad q = \frac{\partial f}{\partial y}$$

$$\mathbf{n} = (p, q, -1)$$

$$\hat{\mathbf{n}} = \frac{1}{\sqrt{p^2+q^2+1}}(p, q, -1)$$

Lambertian reflectance map

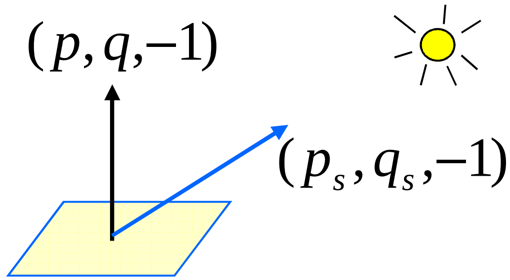

$$E(p,q) = L\rho \frac{1 + pp_s + qq_s}{\sqrt{1 + p^2 + q^2}\sqrt{1 + p_s{}^2 + q_s{}^2}}$$

Local surface orientation that produces equivalent intensities are quadratic conic sections contours in gradient space

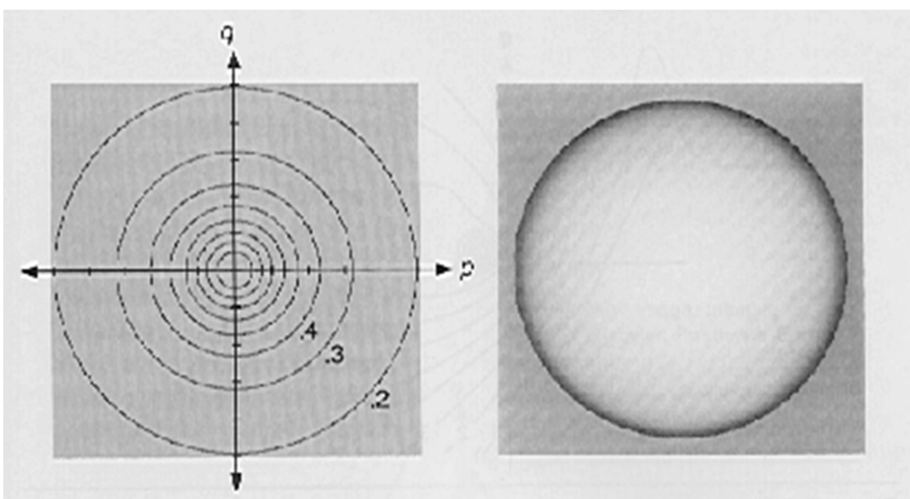

p_s=0,q_s=0

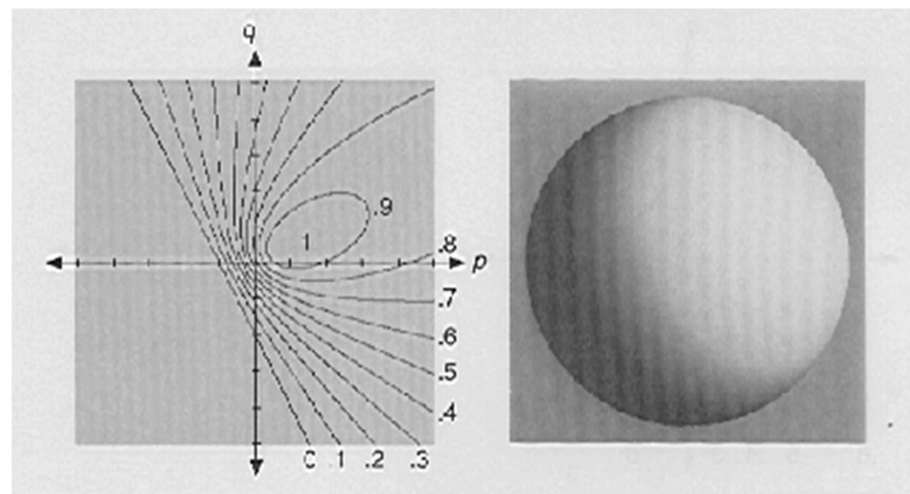

p_s=-2,q_s=-1

Photometric stereo

One image =one light direction

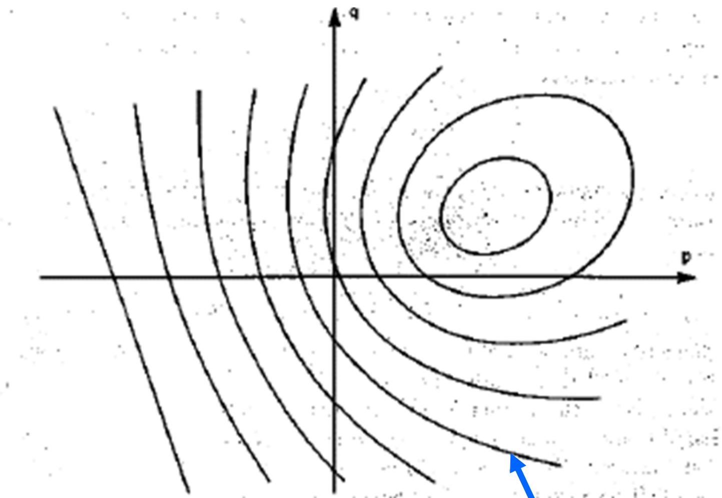

Radiance of one pixel constrains the normal to a curve

Two images = two light directions

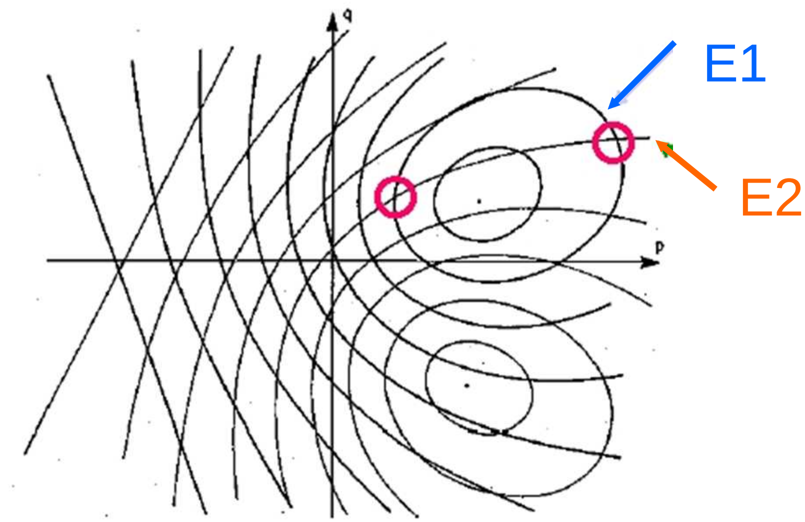

A third image disambiguates between the two.
Normal = intersection of 3 curves

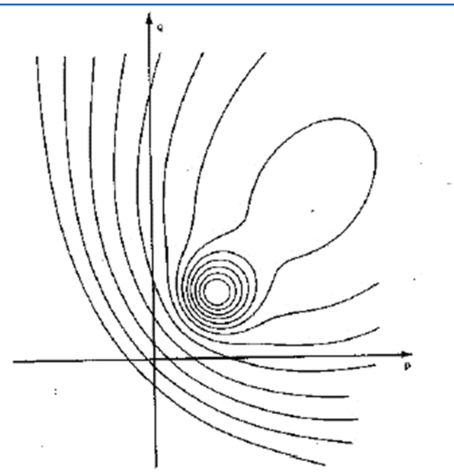

Specular reflectance

Photometric stereo

[Birkbeck]

One image, one light direction

$$I(\mathbf{x}) = B(\mathbf{x}) = \rho(\mathbf{x})\mathbf{n}(\mathbf{x}) \bullet \mathbf{l}_i$$

n images, n light directions

$$\begin{bmatrix} \mathbf{l}_1^T \\ \mathbf{l}_2^T \\ \vdots \\ \mathbf{l}_n^T \end{bmatrix} \rho(\mathbf{x})\mathbf{n}(\mathbf{x}) = \begin{bmatrix} I_1^T(\mathbf{x}) \\ I_2^T(\mathbf{x}) \\ \vdots \\ I_n^T(\mathbf{x}) \end{bmatrix}; \quad A\underbrace{\boxed{\rho(\mathbf{x})\mathbf{n}(\mathbf{x})}}_{\mathbf{b}(\mathbf{x})} = I(\mathbf{x})$$

<u>Given</u>: n>=3 images with different <u>known</u> light dir. (infinite light)

<u>Assume</u>: Lambertain object
orthograhic camera
ignore shadows, interreflections

<u>Recover</u> $\mathbf{b}(\mathbf{x}) = \rho(\mathbf{x})\mathbf{n}(\mathbf{x})$

<u>Albedo</u> = magnitude $|\mathbf{b}(\mathbf{x})|$

<u>Normal</u> = normalized $\dfrac{\mathbf{b}(\mathbf{x})}{|\mathbf{b}(\mathbf{x})|}$

Depth from normals (1)

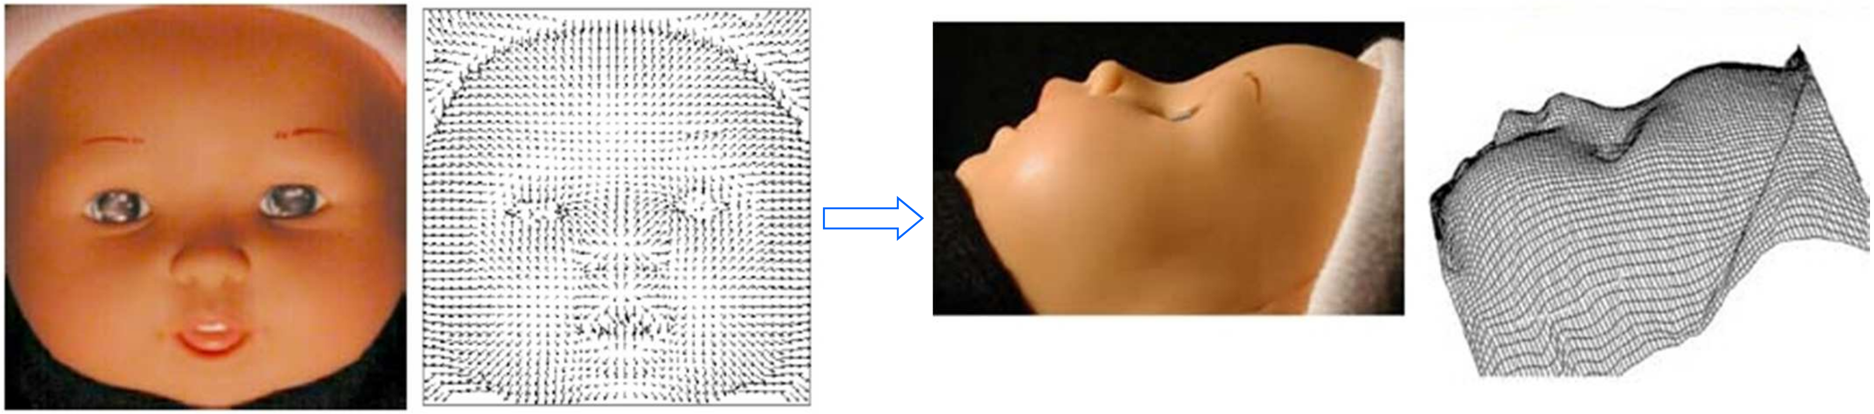

[D. Kriegman]

Integrate normal (gradients p,q) across the image

Simple approach – integrate along a curve from (x_0, y_0)

(x_0, y_0)

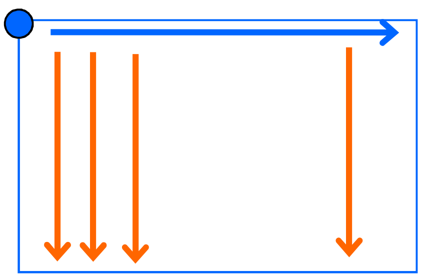

$f(x,0)$

1. From $\mathbf{n} = (n_x, n_y, n_z)$ $\quad p = n_x / n_z \quad q = n_y / n_z$
2. Integrate $p = \partial f / \partial x$ along $(x,0)$ to get $f(x,0)$
3. Integrate $q = \partial f / \partial y$ along each column

$$f(x,y) = f(x_0, y_0) + \int_{(x_0,y_0)}^{(x,y)} (p\,dx + q\,dy)$$

Depth from normals (2)

$$f(x,y) = f(x_0,y_0) + \int_{(x_0,y_0)}^{(x,y)} (p\,dx + q\,dy)$$

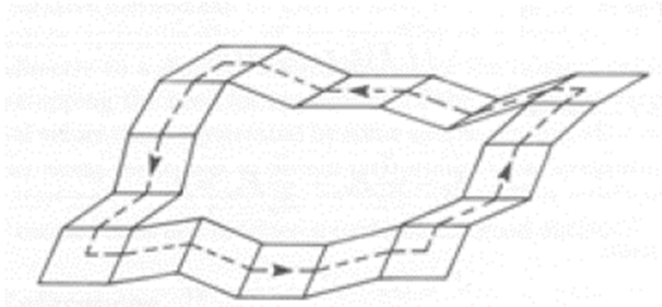

Integrate along a curve from (x_0, y_0)
Might not go back to the start because of noise – depth is not unique

[Escher] no integrability

Impose integrability

A normal map that produces a unique depth map is called integrable

Enforced by $\frac{\partial p}{\partial y} = \frac{\partial q}{\partial x}; \quad \frac{\partial}{\partial y}\frac{\partial f}{\partial x} = \frac{\partial}{\partial x}\frac{\partial f}{\partial y}$

Impose integrabilty

[Horn – Robot Vision 1986]

Solve f(x,y) from p,q by minimizing the cost functional

$$\iint_{image} (f_x - p)^2 + (f_y - q)^2 \, dxdy$$

- Iterative update using calculus of variation
- Integrability naturally satisfied
- F(x,y) can be discrete or represented in terms of basis functions

  Example : Fourier basis (DFT)-close form solution

[Frankfot, Chellappa
A method for enforcing integrability in SFS Alg.
PAMI 1998]

Example integrability

[Neil Birkbeck]

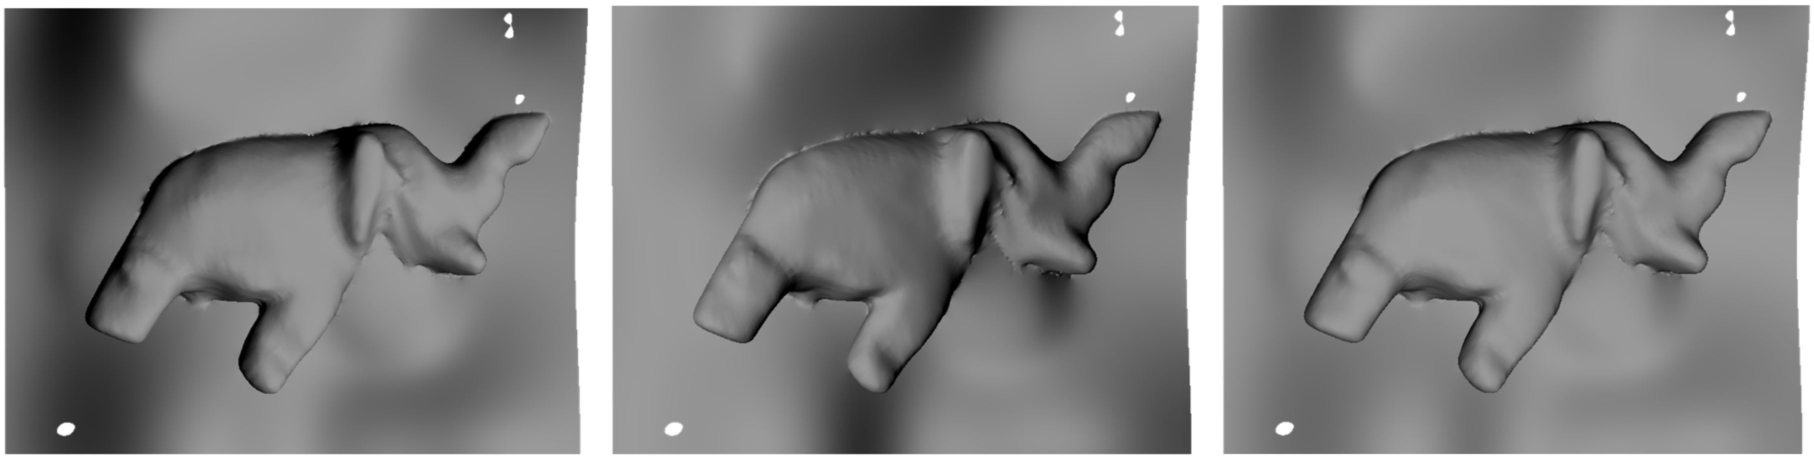

images with different light

normals

Integrated depth

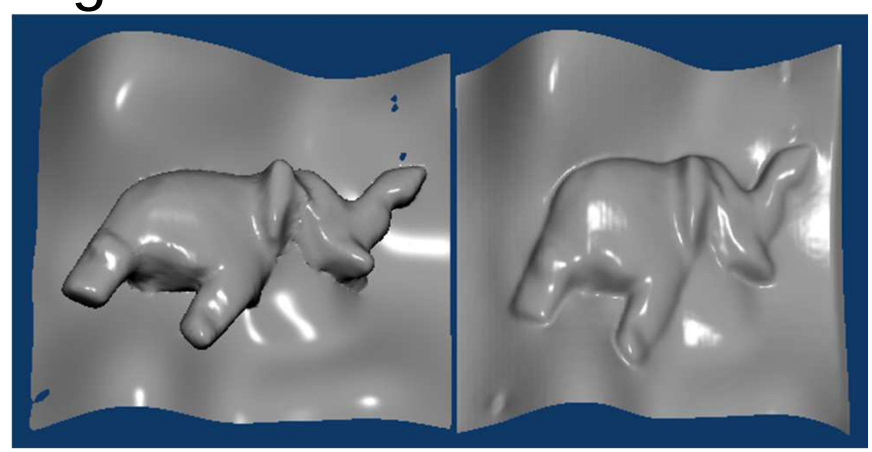

original surface

reconstructed

Image cues
Shading, Stereo, Specularities

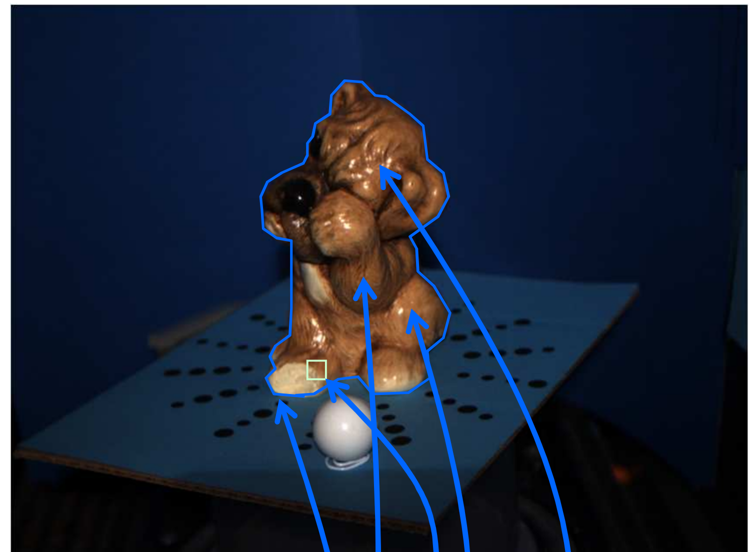

Color (texture)
Shading
Shadows
Specular highlights
Silhouette

Readings: See links on web page

Books: Szeliski 2.2, Ch 12

Forsythe Ch 4,5 (Lab related) .pdf on web site)

Upcoming

- Lab 4 due Mar31
- Exam 2: In-class Apr 2, same format as E1
  - Calculator and 4 sheets of your notes.
- Project presentations: Inclass Apr 9, 11
  - Present the motivation, related literature and libraries
  - Present your progress to-date
  - Prepare 5-10min presentation/person.
- Project report: Hand in at end of classes. (With an earlier hand-in I may have time to comment and you can polish it for a final hand-in)
- Project demos
  - Submit visual demos (videos and easily runnable code)
  - If you like can schedule an in-person demo.

All images

- Unknown lights and normals : It is possible to reconstruct the surface and light positions ?
- What is the set of images of an object under all possible light conditions ?

[Debevec et al]

Space of all images

Problem:

- Lambertian object
- Single view, orthographic camera
- Different illumination conditions (distant illumination)



1. 3D subspace:
[Moses 93][Nayar,Murase 96][Shashua 97]

+ convex obj
(no shadows)

2. Illumination cone:
[Belhumeur and Kriegman CVPR 1996]

3. Spherical harmonic representation:
[Ramamoorthi and Hanharan Siggraph 01]
[Barsi and Jacobs PAMI 2003]

3D subspace

Convex cone

Linear combination of harmonic imag.
(practical 9D basis)

3D Illumination subspace

<u>Lambertian reflection</u> : $I = \rho \mathbf{n} \bullet \mathbf{l} = \mathbf{b} \bullet \mathbf{l}$

(one image point $\mathbf{x}$)

<u>Whole image</u> : $I(:) = \mathbf{x} = \mathbf{B}\mathbf{l} \quad \mathbf{B} = \begin{bmatrix} \mathbf{b}^T{}_1 \\ \ldots \\ \mathbf{b}^T{}_n \end{bmatrix}$ $\boxed{n \times 3}$

(image as vector $\mathbf{I}$)

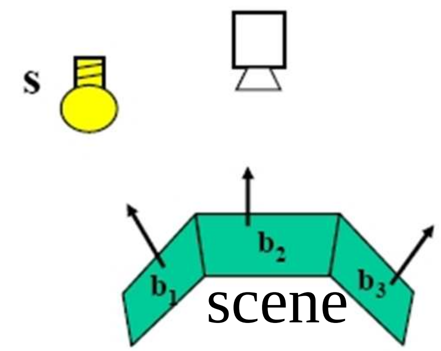

The set of images of a Lambertain scene surface with <u>no shadowing</u> is a subset of a 3D subspace. [Moses 93][Nayar,Murase 96][Shashua 97]

$$L = \{\mathbf{x} \mid \mathbf{x} = \mathbf{B}\mathbf{l}, \forall \mathbf{l} \in \Re^3\}$$

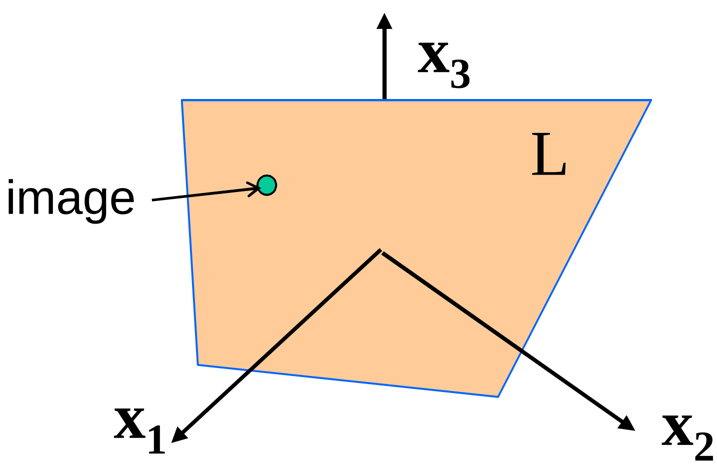

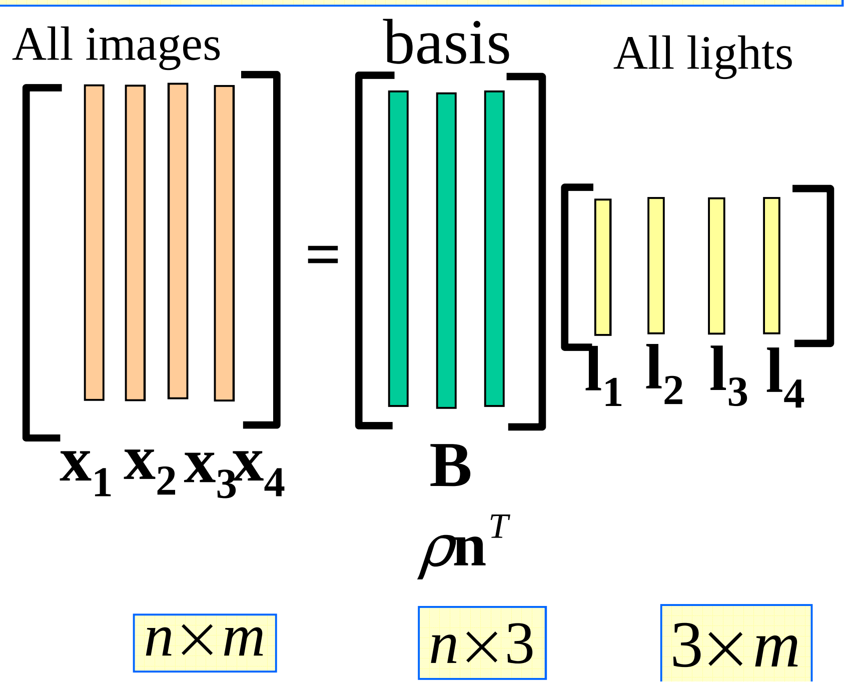

Reconstructing the basis

$$L = \{\mathbf{x} \mid \mathbf{x} = \mathbf{B}\mathbf{l}, \forall \mathbf{l} \in \Re^3\}$$

- Any three images without shadows span L.
- L – represented by an orthogonal basis B.
- How to extract B from images ?

PCA

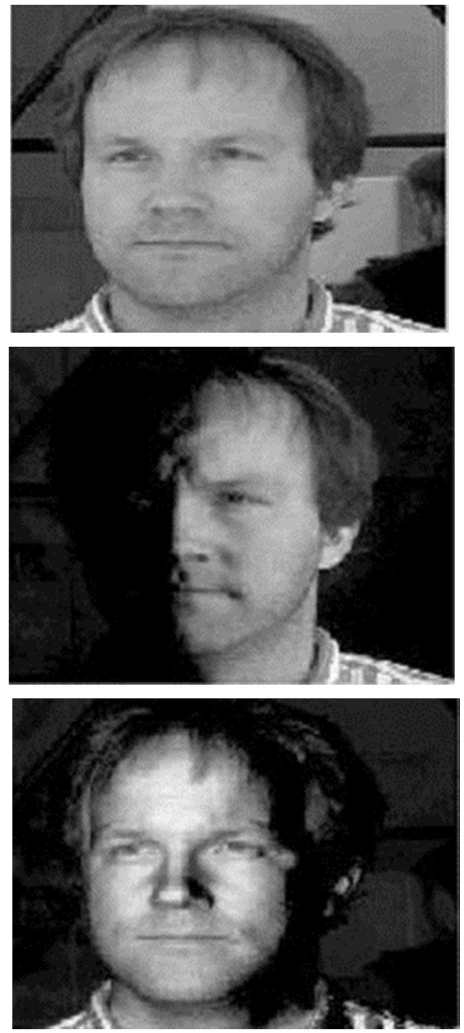

Shadows

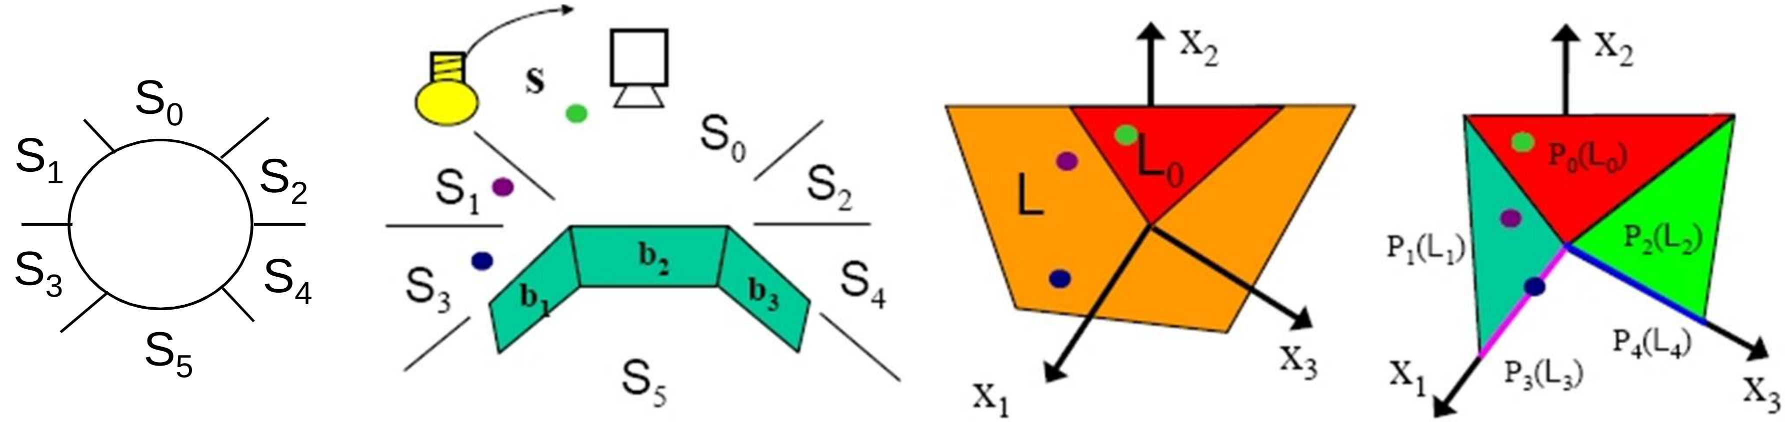

No shadows

Shadows

$$L = \{\mathbf{x} \mid \mathbf{x} = \mathbf{B}\mathbf{l}, \forall \mathbf{l} \in \Re^3\}$$

$$\mathbf{x} = \max(\mathbf{B}\mathbf{l}, 0)$$

Ex: images with all pixels illuminated

$$L_0 = L \cap \{\mathbf{x} \mid \mathbf{x} \in R^n, I_j \geq 0, \forall j\}$$

Single light source

- L_i intersection of L with an orthant i of R^n corresponding cell of light source directions S_i for which the same pixels are in shadow and the same pixels are illuminated.
- $P(L_i)$ projection of L_i that sets all negative components of L_i to 0 (convex cone)

The set of images of an object produces by a single light source is :

$$U = \{\mathbf{x} \mid \mathbf{x} = \max(\mathbf{B}\mathbf{l}, 0), \forall \mathbf{l} \in R^3\} = \bigcup_i P_i(L_i)$$

Shadows and multiple lights

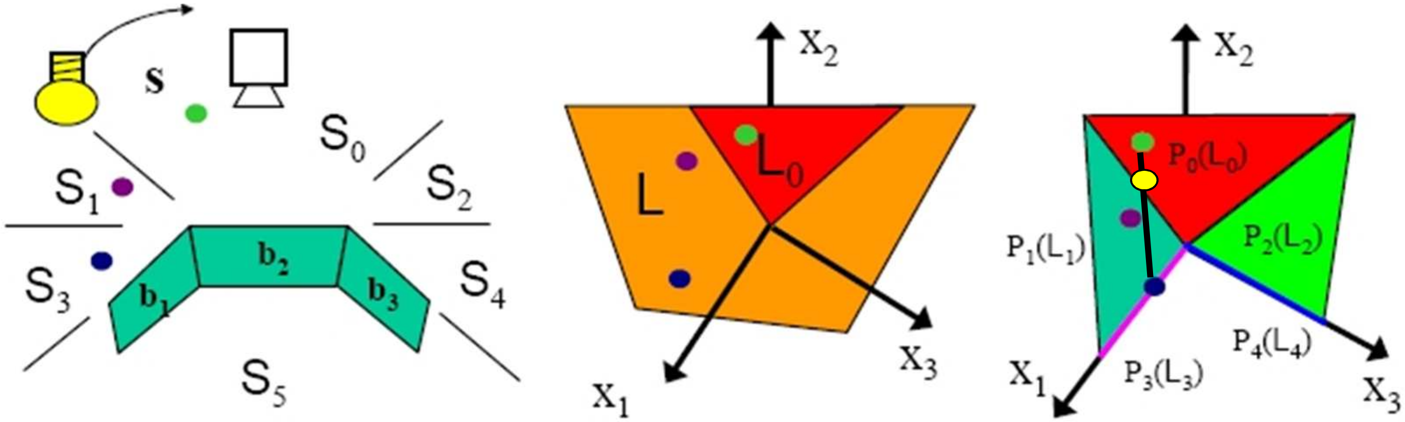

Shadows, multiple lights $\quad \mathbf{x} = \sum_i \max(\mathbf{Bl}_i, 0)$

The image illuminated with two light sources l_1, l_2, lies on the line between the images of x_1 and x_2.

The set of images of an object produces by an arbitrary number of lights is the convex hull of U = illumination cone C.

Illumination cone

The set of images of a any Lambertain object under all light conditions is a convex cone in the image space.
[Belhumeur,Kriegman: What is the set of images of an object under all possible light conditions ?, IJCV 98]

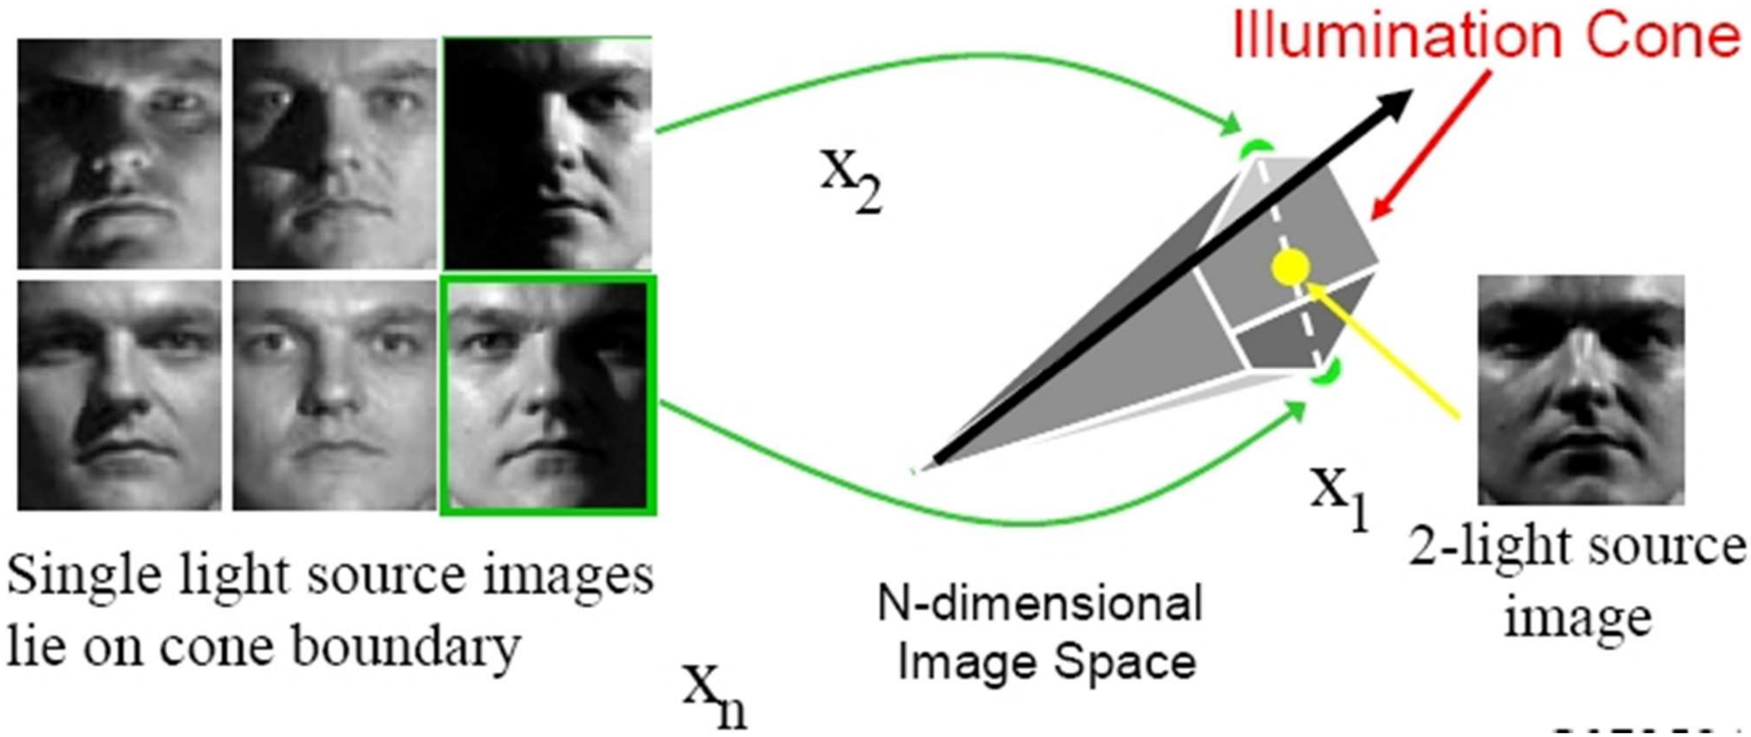

Do ambiguities exist ?

Can two different objects produce the same illumination cone ? YES “Bas-relief” ambiguity

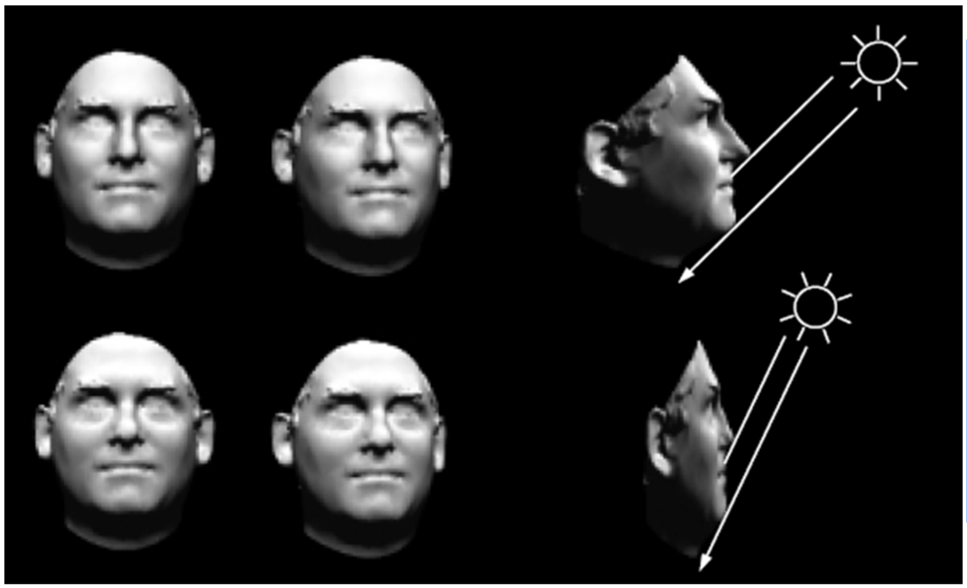

Convex object

- B span L
- Any $A \in GL(3)$, $B^* = BA$ span L
- $I = B^* S^* = (BA)(A^{-1}S) = BS$

  Same image B lighted with S and B^* lighted with S^*

When doing PCA the resulting basis is generally not normal*albedo

GBR transformation

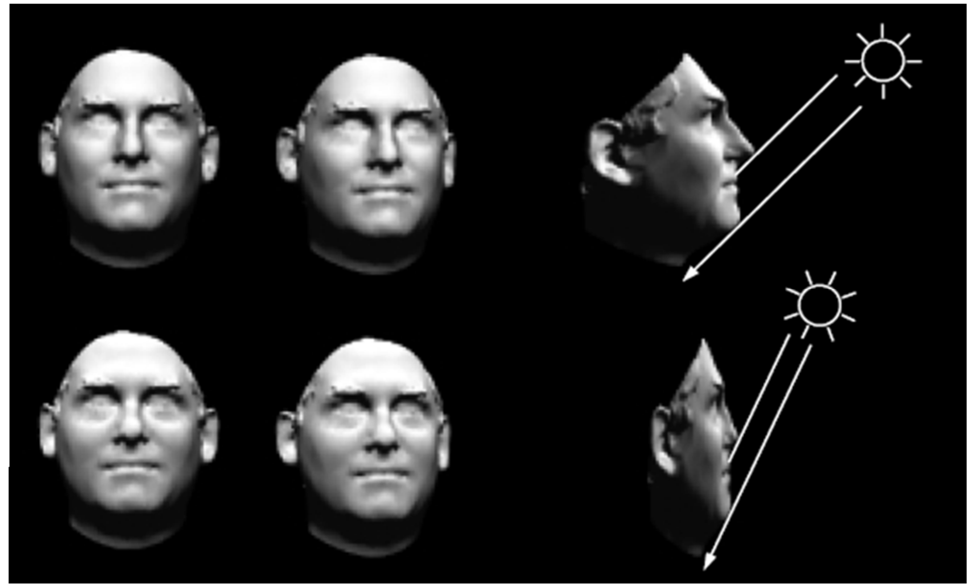

[Belhumeur et al: The bas-relief ambiguity IJCV 99]

<u>Surface integrability :</u>

Real B, transformed B*=BA is integrable only for General Bas Relief transformation.

$$A = G^T = \begin{bmatrix} \lambda & 0 & -\mu \\ 0 & \lambda & -\upsilon \\ 0 & 0 & 1 \end{bmatrix} \qquad \bar{f}(x,y) = \lambda f(x,y) + \mu x + \upsilon y$$

Uncalibrated photometric stereo

- Without knowing the light source positions, we can recover shape only up to a GBR ambiguity.

1. From n input images compute B^* (SVD)
2. Find A such that $B^* A$ close to integrable
3. Integrate normals to find depth.

Comments

- GBR preserves shadows [Kriegman, Belhumeur 2001]
- If albedo is known (or constant) the ambiguity G reduces to a binary subgroup [Belhumeur et al 99]
- Interreflections : resolve ambiguity [Kriegman CVPR05]

Spherical harmonic representation

Theory : infinite no of light directions
space of images infinite dimensional

[Illumination cone, Belhumeur and Kriegman 96]

Practice : (empirical) few bases are enough

[Hallinan 94, Epstein 95]

 = .2  +.3  +…

Simplification : Convex objects (no shadows, intereflections)

[**Ramamoorthi and Hanharan**: Analytic PCA construction for Theoretical analysis of Lighting variability in images of a Lambertian object: SIGGRAPH01]

[**Barsi and Jacobs**: Lambertain reflectance and linear subspaces: PAMI 2003]

Basis approximation

Spherical harmonics basis

- Sphere analog to the Fourier basis on the line or square
- Angular portion of the solution to Laplace equation in spherical coordinates $\nabla^2 \psi = 0$
- Orthonormal basis for the set of all functions on the surface of the sphere

$$Y_{lm}(\theta,\phi) = \underbrace{\sqrt{\frac{(2l+1)}{4\pi}\frac{(l-|m|)!}{(l+|m|)!}}}_{\text{Normalization factor}} \underbrace{P_{l|m|}(\cos\theta)}_{\text{Legendre functions}} \underbrace{e^{im\phi}}_{\text{Fourier basis}}$$

Illustration of SH

m

l

0

1

2

.

.

.

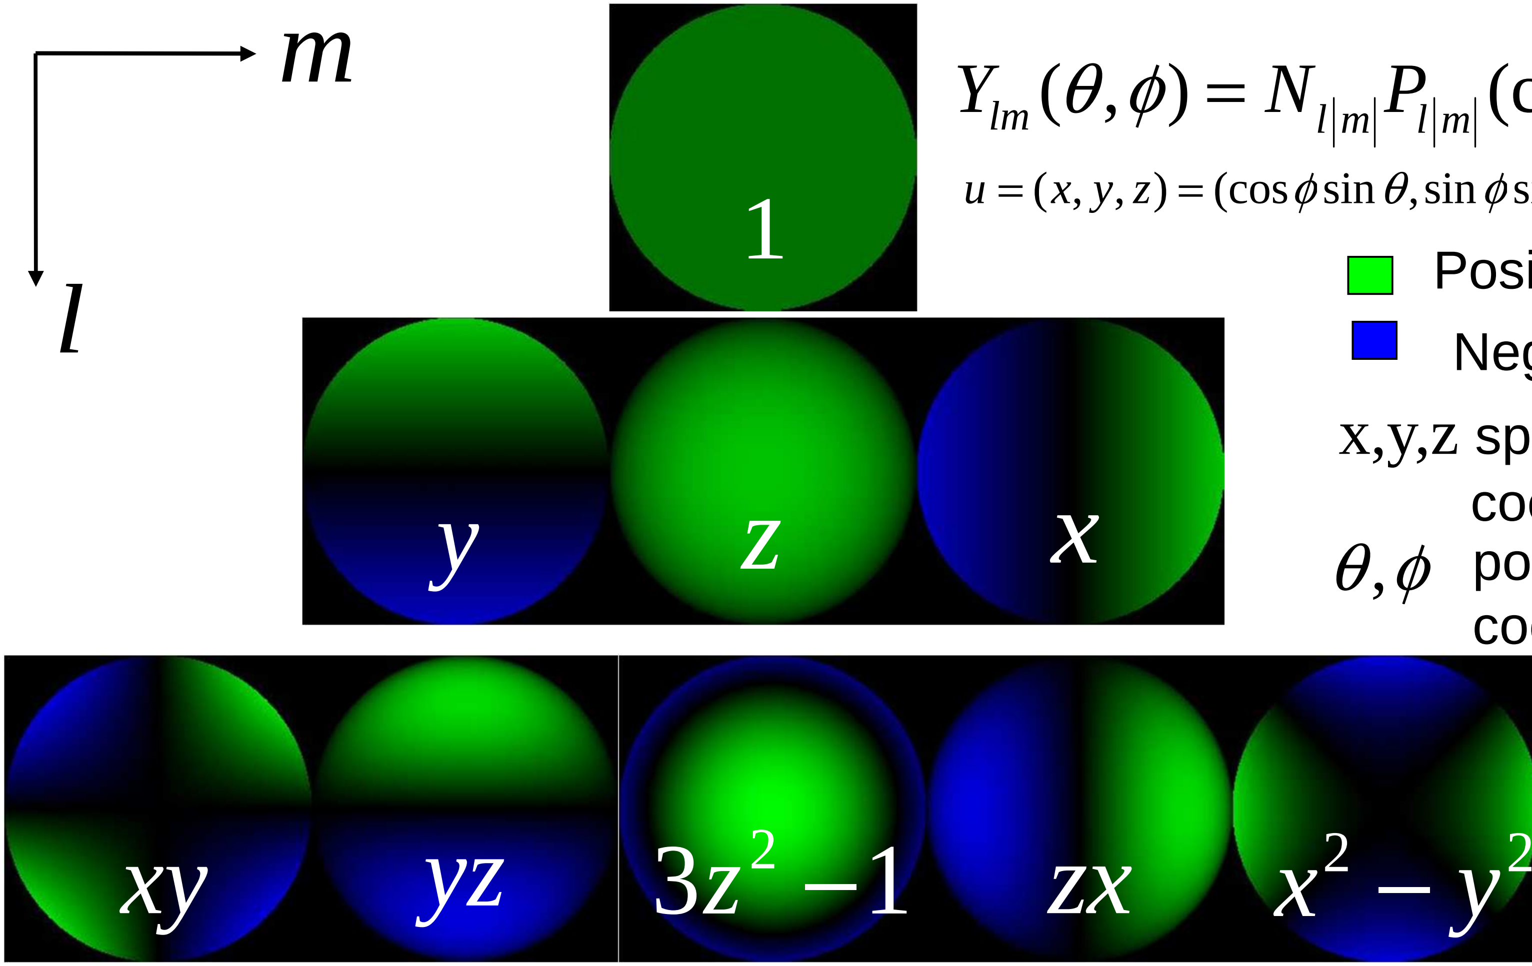

-2 -1 0 1 2

$$Y_{lm}(\theta,\phi) = N_{l|m|}P_{l|m|}(\cos\theta)e^{im\phi}$$

$$u = (x, y, z) = (\cos\phi\sin\theta, \sin\phi\sin\theta, \cos\theta)$$

Positive

Negative

x,y,z space coordinates;

θ,ϕ polar coordinates

odd components $Y_{lm} = Y_{l|m|}{}^{e} \pm iY_{l|m|}{}^{o}$ even components

Example of approximation

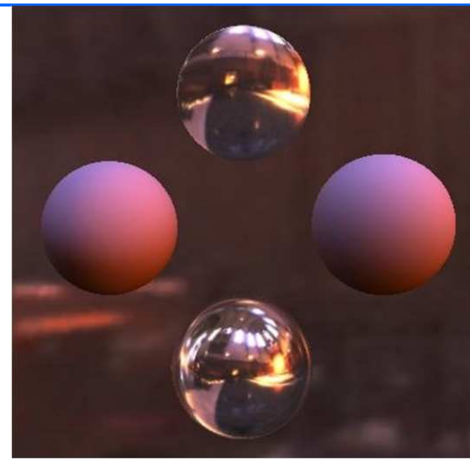

Efficient rendering
- known shape
- complex illumination (compressed)

Exact image

9 terms approximation

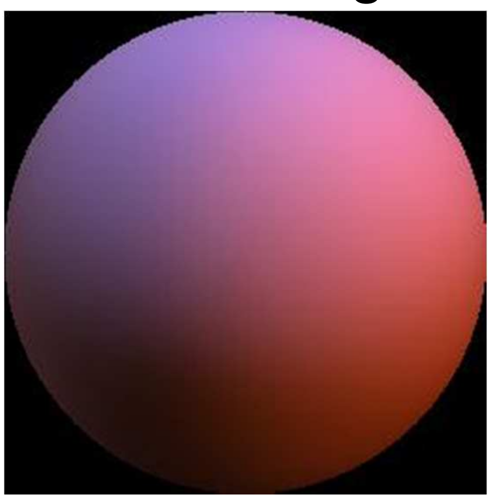

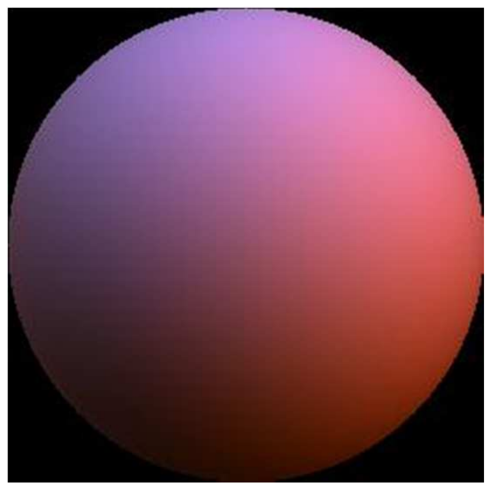

[Ramamoorthi and Hanharan: An efficient representation for irradiance enviromental map Siggraph 01]

Not good for hight frequency (sharp) effects ! (specularities)

Relation between SH and PCA

[Ramamoorthi PAMI 2002]

<u>Prediction</u>: 3 basis 91% variance
5 basis 97%

<u>Empirical</u>: 3 basis 90% variance
5 basis 94%

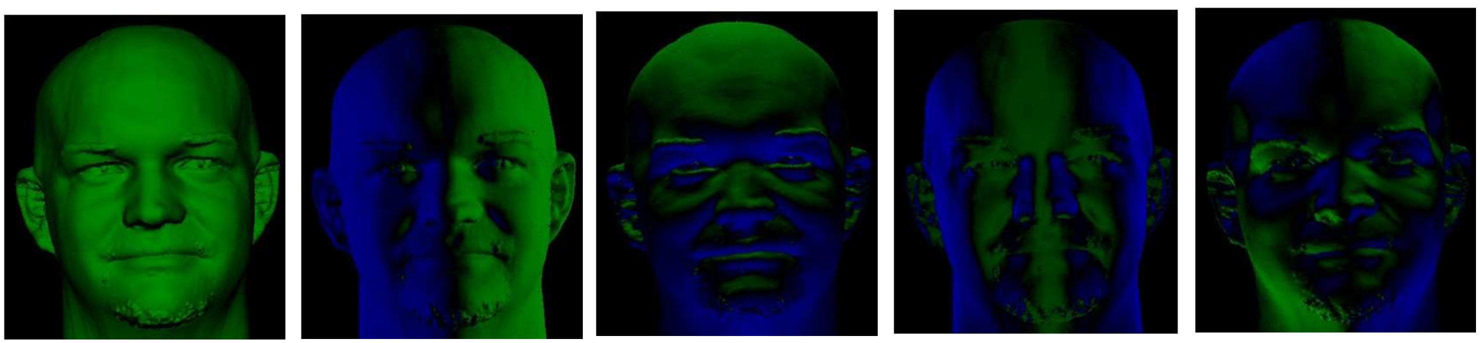

42% 33% 16% 4% 2%

Summary: Image cues

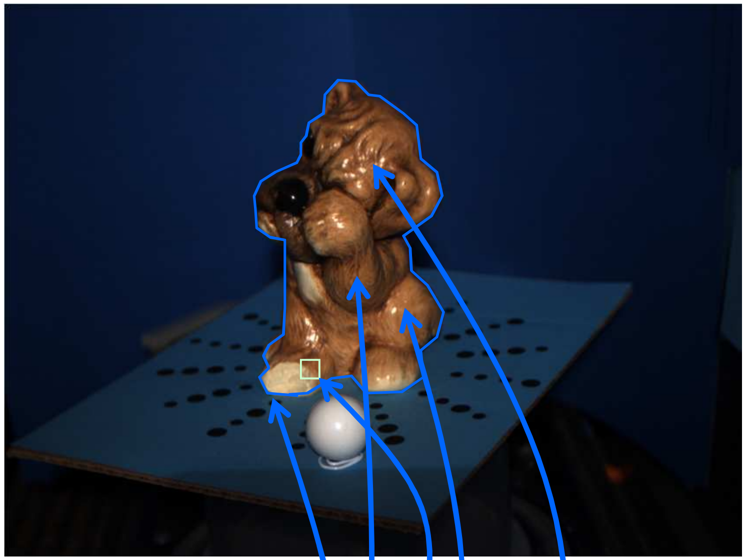

Color (texture)

Shading

Shadows

Specular highlights

Silhouette

Properties of SH

Function decomposition

f piecewise continuous function on the surface of the sphere

$$f(u) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} f_{lm} Y_{lm}(u)$$

where

$$f_{lm} = \int_{S^2} f(u) Y^{*}_{lm}(u) du$$

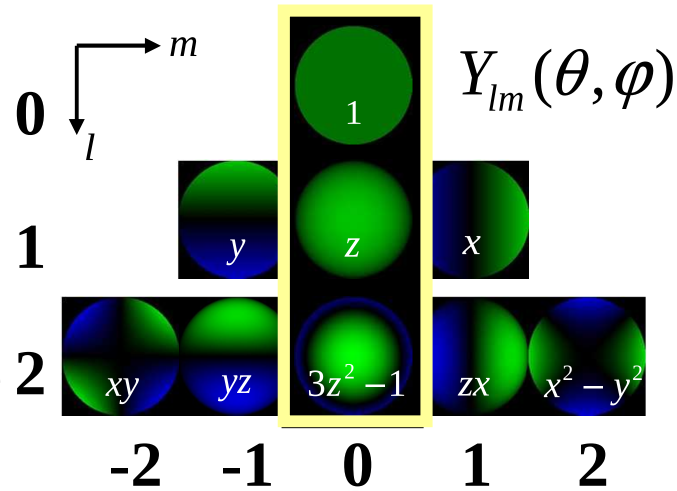

Rotational convolution on the sphere

Funk-Hecke theorem:

k circularly symmetric bounded integrable function on [-1,1] $\quad k(u) = \sum_{l=0}^{\infty} k_l Y_{l0}$

$$k * Y_{lm} = \alpha_l Y_{lm} \qquad \alpha_l = \sqrt{\frac{4\pi}{2l+1}} k_l$$

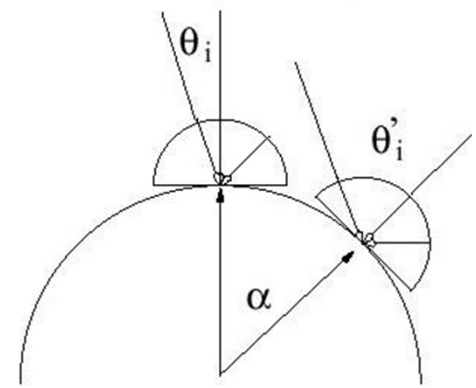

Reflectance as convolution

Lambertian reflectance

One light $$R(u') = l(u)\rho \max(0, u \bullet u')$$

Lambertian kernel $$k(u \bullet u') = \max(0, u \bullet u')$$

Integrated light $$R(u') = \int_{S^2} k(u \bullet u') l(u) du$$

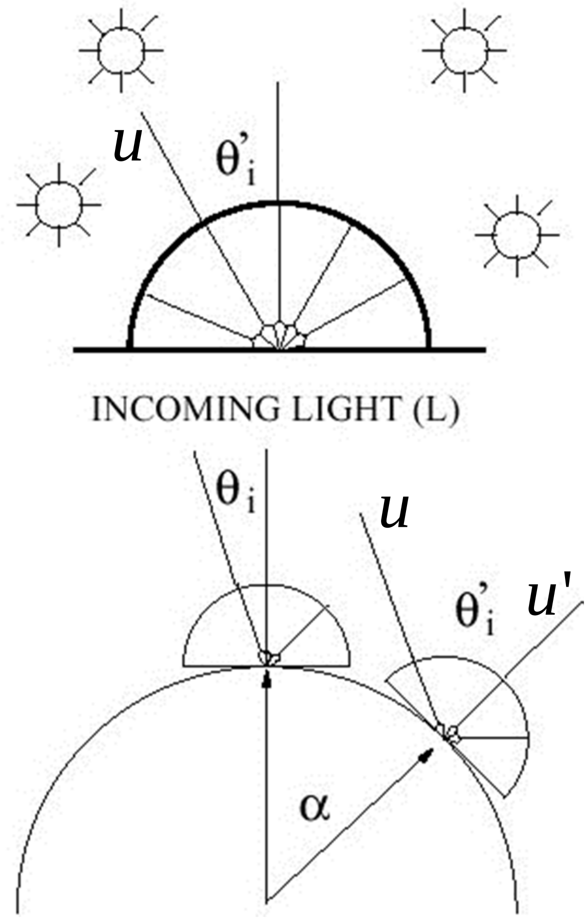

SH representation

light

$$l(u) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} l_{lm} Y_{lm}(u)$$

Lambertian kernel

$$k = \sum_{l=0}^{\infty} k_l Y_{l0}$$

Lambertian reflectance (convolution theorem)

$$R = k * l = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \left(\sqrt{\frac{4\pi}{2l+1}} k_l l_{lm}\right) Y_{lm} = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} r_{lm} Y_{lm}$$

Convolution kernel

<u>Lambertian kernel</u>

$$k(u \bullet u') = \max(0, u \bullet u')$$

$$k = \sum_{l=0}^{\infty} k_l Y_{l0}$$

$$k_l = \begin{cases} \dfrac{\sqrt{\pi}}{2} & n = 0 \\ \dfrac{\sqrt{\pi}}{3} & n = 1 \\ (-1)^{l/2+1} \dfrac{\sqrt{(2l+1)\pi}}{2^l (l-1)(l+2)} \begin{pmatrix} l \\ l/2 \end{pmatrix} & n \geq 2, \text{even} \\ 0 & n \geq 2, \text{odd} \end{cases}$$

Asymptotic behavior of k_l for large l

$$k_l \approx l^{-2} \qquad r_{lm} \approx l^{-5/2}$$

- Second order approximation accounts for 99% variability
- k like a low-pass filter

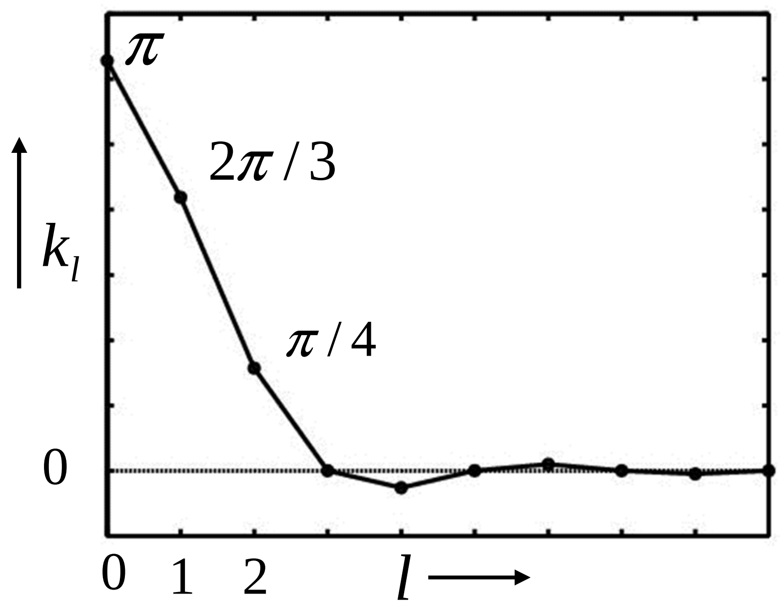

[Basri & Jacobs 01]
[Ramamoorthi & Hanrahan 01]

From reflectance to images

Unit sphere $\Rightarrow$ general shape

Rearrange normals on the sphere

Reflectance on a sphere $$R = k * l = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} r_{lm} Y_{lm}$$

Image point with normal n_i $$I_i = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \rho_i r_{lm} Y_{lm}(n_i)$$

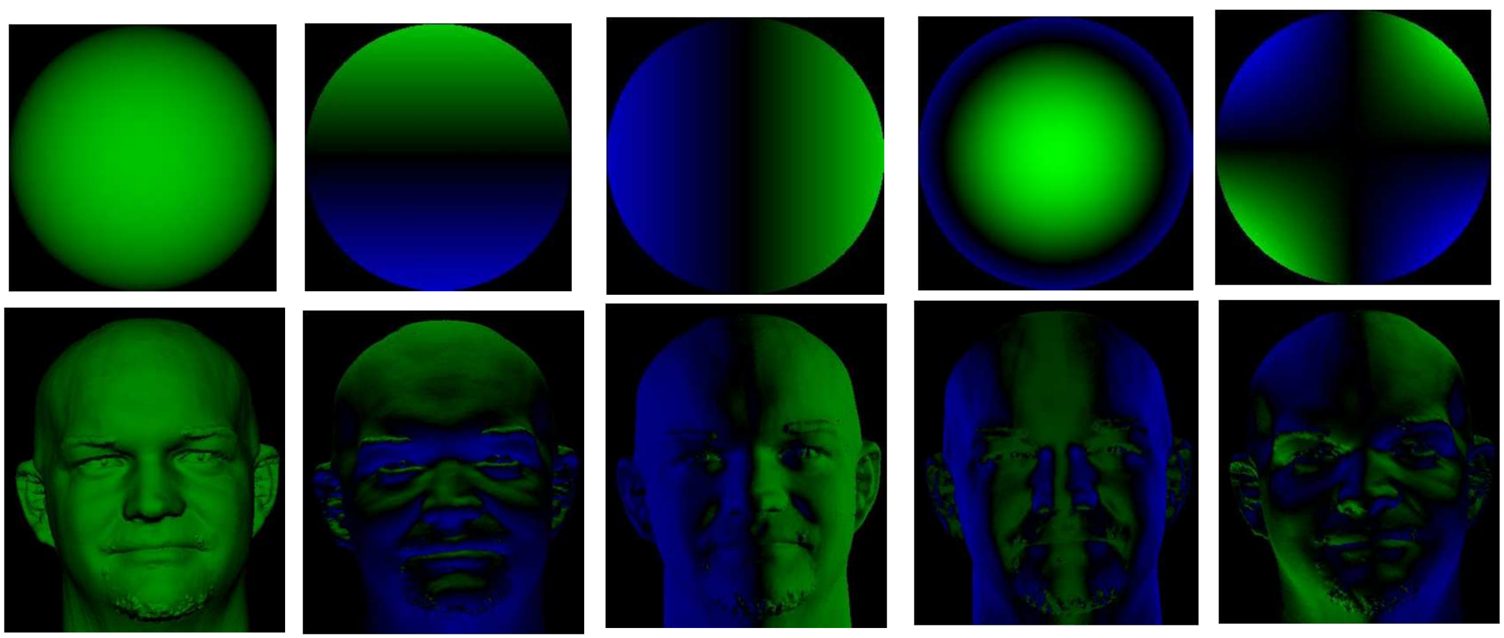

Shape from Shading

<u>Given</u>: **one** image of an object illuminated with a distant light source
<u>Assume</u>: Lambertian object, with known, or constant albedo (usually assumes 1)
orthograhic camera
known light direction
ignore shadows, interreflections
<u>Recover:</u> normals

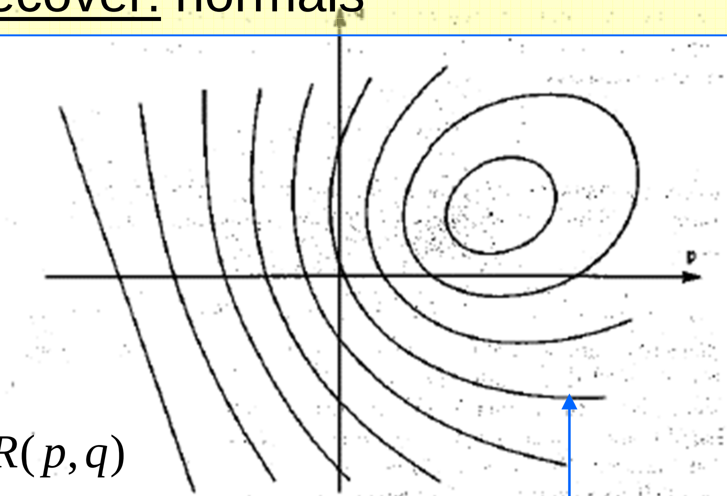

$R(p,q)$

Radiance of one pixel constrains the normal to a curve

ILL-POSED

Surface $s(x,y) = (x, y, f(x,y))$

Gradient space $p = \frac{\partial f}{\partial x}$ $q = \frac{\partial f}{\partial y}$

Normal $\mathbf{n} = (p, q, -1)$

$$\hat{\mathbf{n}} = \frac{1}{\sqrt{p^2 + q^2 + 1}}(p, q, -1)$$

Lambertian reflectance: depends only on n (p,q):

$$E(\mathbf{x}) = \cos(\mathbf{n}(\mathbf{x}), \mathbf{l}) = \frac{\mathbf{n}(\mathbf{x}) \bullet \mathbf{l}}{\|\mathbf{n}(\mathbf{x})\|}$$

Variational SFS

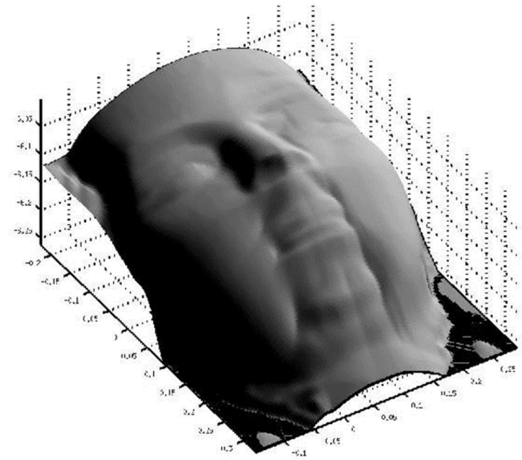

Image info shading **Recovers** Integrated normals

- Defined by Horn and others in the 70's.
- Variational formulation

$$\iint_{object} \left(I(x,y) - E(p,q)\right)^2 dxdy = \iint_{object} \left(I(x,y) - \frac{[p,q,-1]' \bullet \mathbf{l}}{\sqrt{p^2+q^2+1}} \right)^2 dxdy + \underbrace{\alpha \iint_{object} \left(\frac{\partial p}{\partial y} - \frac{\partial q}{\partial x} \right)^2 dxdy}_{regularization}$$

- Showed to be ill –posed [Brooks 92] (*ex . Ambiguity convex/concave*)
- Classical solution – add regularization, integrability constraints
- Most published algorithms are non-convergent [Duron and Maitre 96]

Examples of results

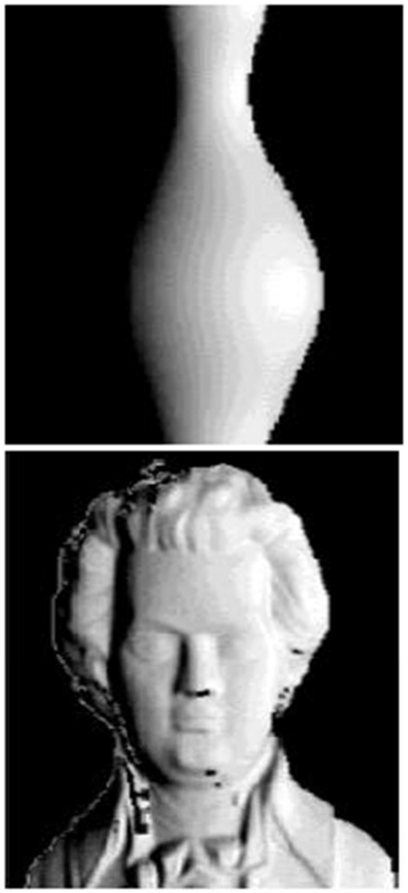

Synthetic images

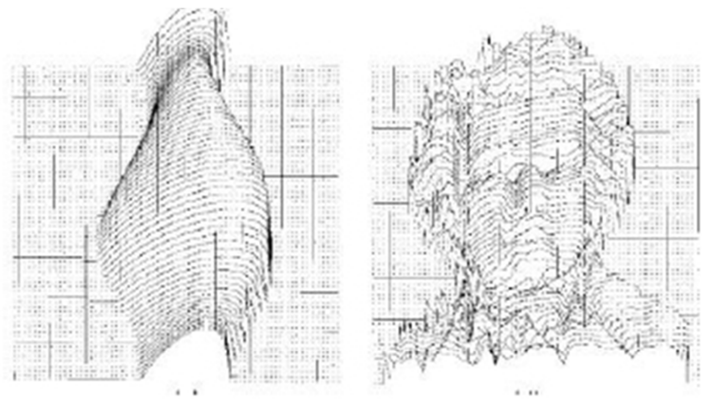

Tsai and Shah's method 1994

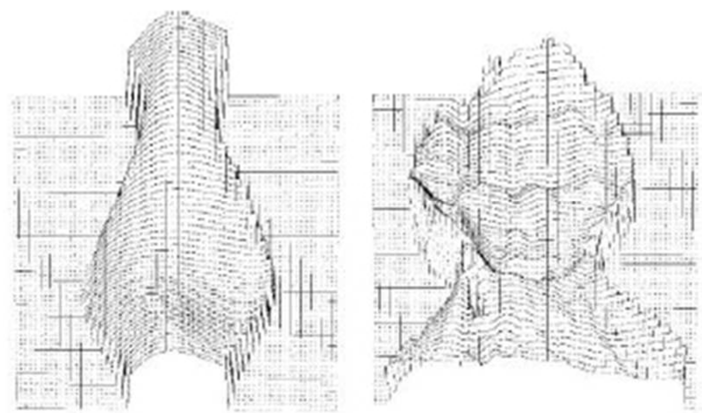

Pentland's method 1994

Well posed SFS

[Prados ICCV03, ECCV04**]** reformulated SFS as a well-posed problem

$$E(\mathbf{x}) = \cos(\mathbf{n}(\mathbf{x}), \mathbf{L}) = \frac{\mathbf{n}(\mathbf{x}) \cdot \mathbf{L}}{|\mathbf{n}(\mathbf{x})|}$$

Lambertian reflectance

Orthographic camera

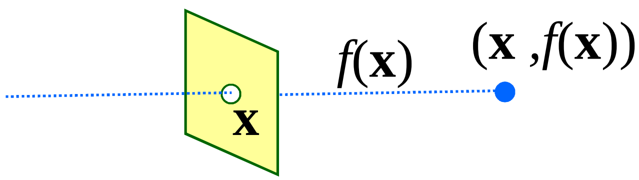

$$s = \{(\mathbf{x}, f(\mathbf{x})) \mid \mathbf{x} = (u, v) \in \Omega\}$$

$$\mathbf{n}(s(\mathbf{x})) = (\nabla f(\mathbf{x}), -1)$$

$$I(\mathbf{x}) = \frac{-\nabla f(\mathbf{x}) \cdot \mathbf{l} + c}{\sqrt{1 + |\nabla f(\mathbf{x})|^2}} \qquad \mathbf{L} = (\mathbf{l}, c)$$

Perspective camera

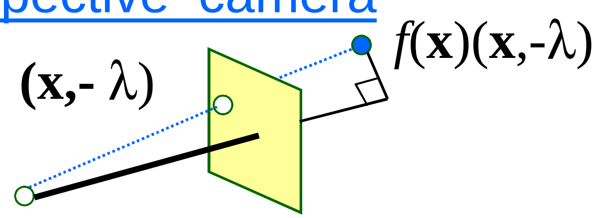

$$s = \{ f(\mathbf{x})(\mathbf{x}, -\lambda) \mid \mathbf{x} = (u, v) \in \Omega\}$$

$$\mathbf{n}(s(\mathbf{x})) = (\lambda \nabla f(\mathbf{x}), f(\mathbf{x}) + \mathbf{x} \cdot \nabla f(\mathbf{x}))$$

$$I(\mathbf{x}) = \frac{\lambda \mathbf{l} \cdot \nabla f(\mathbf{x}) + c(f(\mathbf{x}) + \mathbf{x} \cdot \nabla f(\mathbf{x}))}{\sqrt{\lambda^2 |\nabla f(\mathbf{x})|^2 + (f(\mathbf{x}) + \mathbf{x} \cdot \nabla f(\mathbf{x}))^2}}$$

Hamilton-Jacobi equations - no smooth solutions;

$$H(x, \nabla u) = 0$$

- require boundary conditions

Well-posed SFS (2)

Hamilton-Jacobi equations - no smooth solutions;

- require boundary conditions

Solution

1. **Impose smooth solutions** – not practical because of image noise
2. **Compute viscosity solutions** [Lions et al.93] (smooth almost everywhere)

   still require boundary conditions

E. Prados :general framework – characterization viscosity solutions.

(based on Dirichlet boundary condition)

efficient numerical schemes for orthogonal and perspective camera

showed that SFS is a well-posed for a finite light source

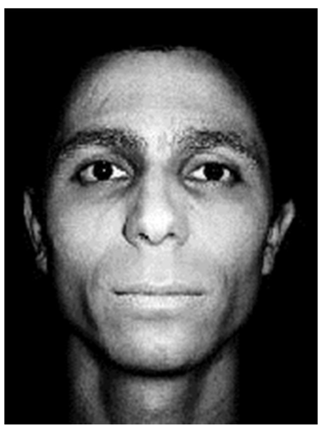

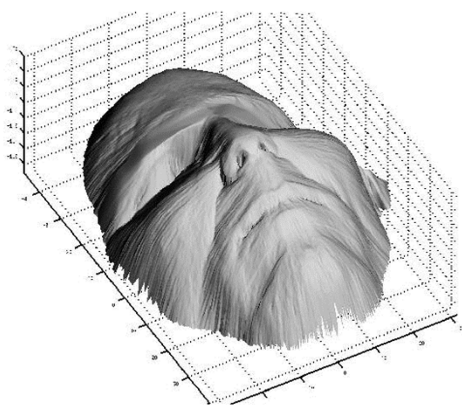

[Prados ECCV04]

Shading: Summary

Space of all images :	Lambertian object Distant illumination One view (orthographic)	
	+ Convex objects	
1. 3D subspace		3D subspace
2. Illumination cone:		Convex cone
2. Spherical harmonic representation:		Linear combination of harmonic imag. (practical 9D basis)

Reconstruction :	Single light source	
1. Shape from shading	One image Unit albedo Known light	Ill-posed + additional constraints
2. Photometric stereo	Multiple imag/1 view Arbitrary albedo Known light	
3. Uncalibrated photometric stereo	+ Unknown light	GBR ambiguity Family of solutions

Extension to multiple views

<u>Problem</u>: PS/SFS one view ⇨ incomplete object

<u>Solution</u> : extension to multiple views – rotating obj., light var.

<u>Problem</u>: we don’t know the pixel correspondence anymore

<u>Solution</u>: iterative estimation: normals/light – shape

initial surface from SFM or visual hull

Input images

Initial surface

Refined surface

1. Kriegman et al ICCV05; Zhang, Seitz … ICCV 03 — SFM
2. Cipolla, Vogiatzis ICCV05, CVPR06 — Visual hull

Multiview PS+ SFM points

[Kriegman et al ICCV05][Zhang, Seitz … ICCV 03]

1. SFM from corresponding points: camera & initial surface (Tomasi Kanade)

2. Iterate:

- factorize intensity matrix : light, normals, GBR ambiguity
- Integrate normals
- Correct GBR using SFM points (constrain surface to go close to points)

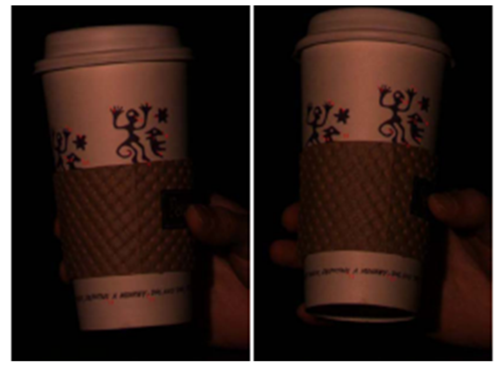

images

Initial surface

$$I = \mathbf{LN}$$

$$\sum_{xy} \left(\frac{\partial f}{\partial x} + \frac{n_x}{n_z} \right)^2 + \left(\frac{\partial f}{\partial y} + \frac{n_y}{n_z} \right)^2$$

Integrated surface

Rendered Final surface

Multiview PS + frontier points

[Cipolla, Vogiatzis ICCV05, CVPR06]

1. initial surface SFS

   visual hull – convex envelope of the object

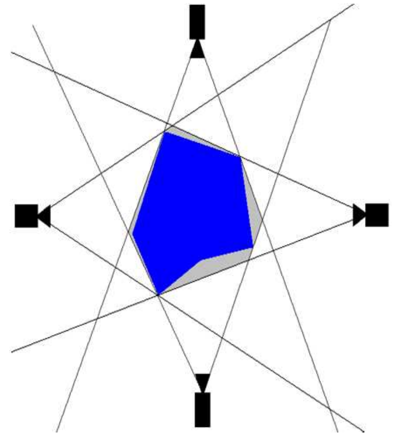

2. initial light positions from frontier points

   plane passing through the point and the camera center is tangent to the object > known normals

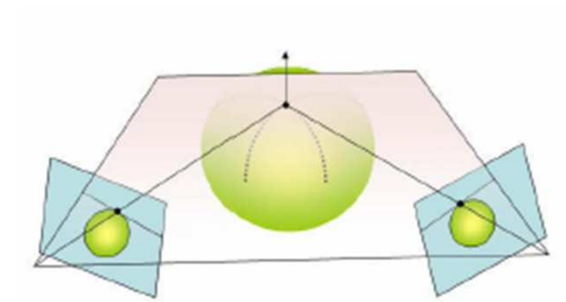

3. Alternate photom normals / surface (mesh)

   v *photom normals*

   n *surface normals – using the mesh*

   mesh –occlusions, correspondence in I

$$\sum_f \sum_i \left(l_i \bullet v_f - i_{fi} \right)^2$$

$$\sum_f \left| n_f - v_f \right|^2$$

Multiview PS + frontier points

(a) Input images.

(b) Visual hull reconstruction.

(c) Our results.

(a) Input images.

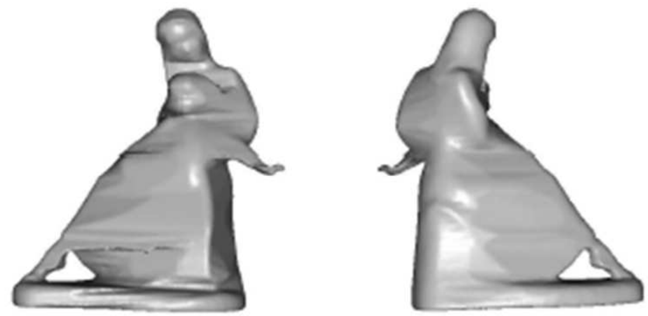

(b) Visual hull reconstruction.

(c) Our results.

Stereo

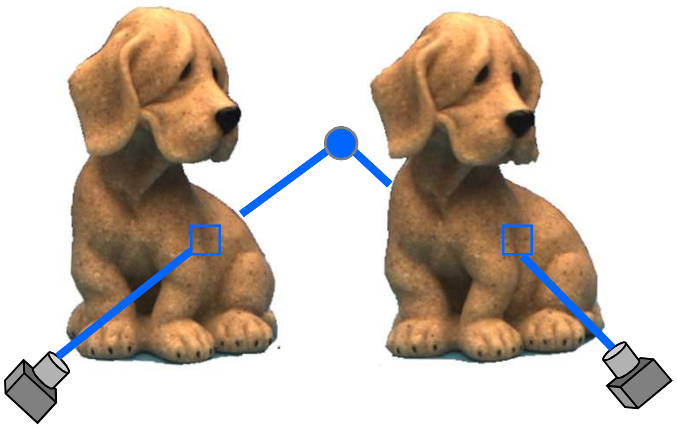

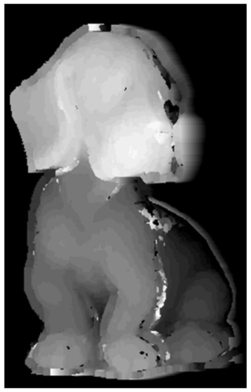

[Birkbeck]

[**Assumptions** two images
Lambertian reflectance
textured surfaces]

Image info texture

Recover per pixel depth

Approach triangulation of corresponding points

corresponding points

- recovered correlation of small parches around each point
- calibrated cameras – search along epipolar lines

Rectified images

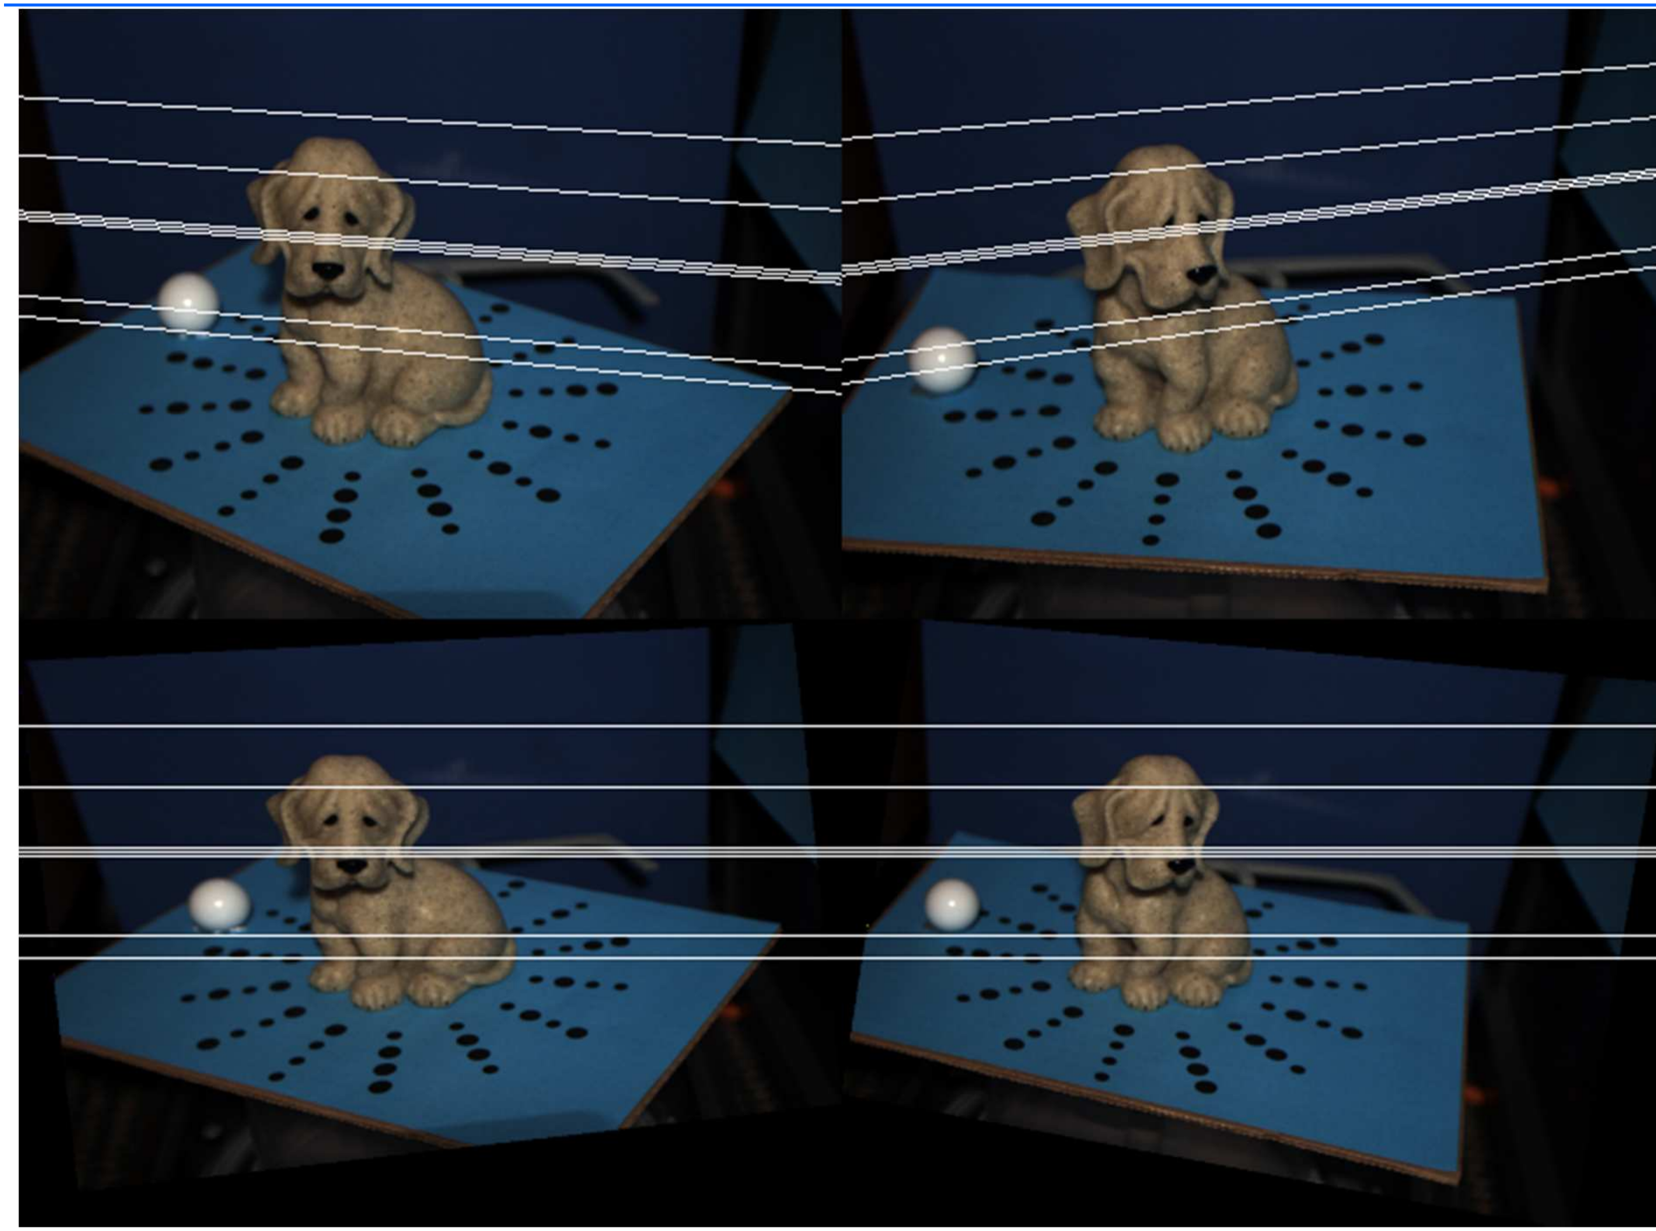

Disparity

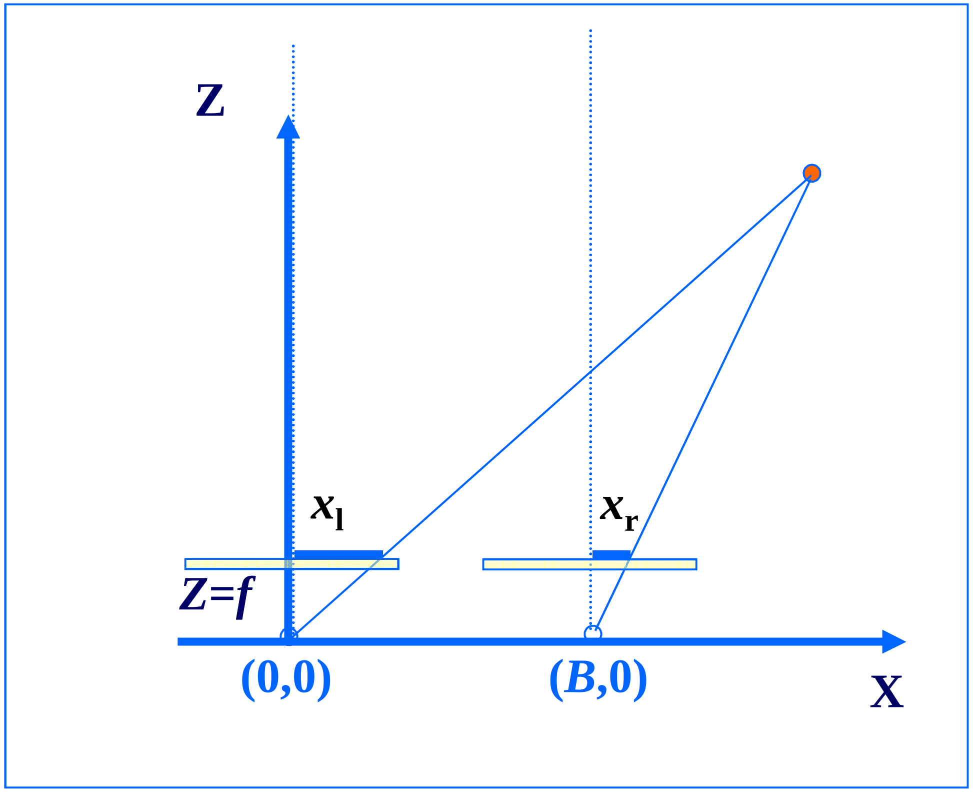

Disparity d

$$Z = \frac{f}{x_l} X = \frac{f}{x_r}(X - B)$$

$$Z = \frac{Bf}{\underbrace{\boxed{x_l - x_r}}_{d}} = \frac{Bf}{d}$$

Correlation scores

Point: $\mathbf{x} = (x, y, f(x, y))$ With respect to first image

Calibrated cameras: pixel in I_1 $\quad m_1 = P_1(\mathbf{x}) = (x, y)$

pixel in I_2 $\quad m_2 = P_2(\mathbf{x})$

Small planar patch: $N(x, y)$

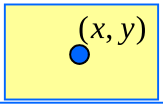

1. Plane parallel with image planes, no illumination variation

$$SAD_{12} = \int_{m \in N(m_1)} I_1(m_1 + m) - I_2(m_2 + m)dm$$

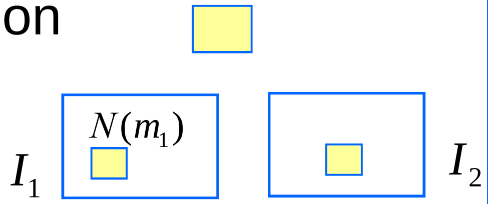

2. Compensate for illumination change

$$NCC_{12} = \int_{m \in N(m_1)} C_1(m)C_2(m)dm \quad C_i(m) = I_i(m_i + m) - \bar{I}_i(m_i) \quad \bar{I}_i \text{mean}$$

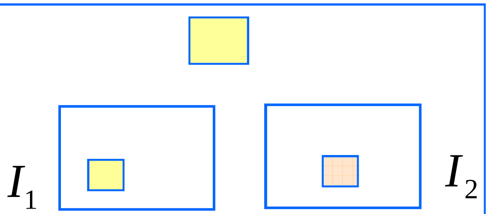

3. Arbitrary plane

$$\int_{m \in N(m_1)} I_1(m_1 + m) - I_2(H(m_1 + m))dm \quad H = R - \frac{\mathbf{t}\mathbf{N}^T}{d}$$

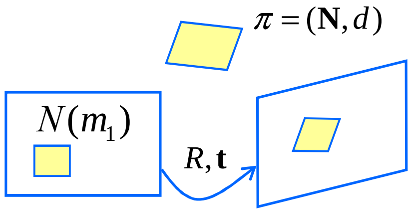

Specular surfaces

Reflectance equation

require: BRDF, light position

$$R_o = \rho(\theta_i, \phi_i, \theta_o, \phi_o) l(\theta_i, \phi_i) \cos(\theta_i)$$

Image info shading+specular highlights

Approaches

1. Filter specular highlights (*brightness, appear at sharp angles*)
2. Parametric reflectance
3. Non-parametric reflectance map (*discretization of BRDF*)
4. Account for general reflectance

Helmholz reciprocity [Magda et al ICCV 01, IJCV03]

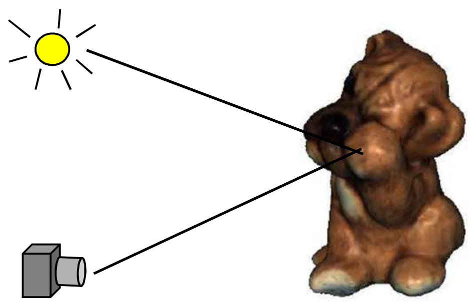
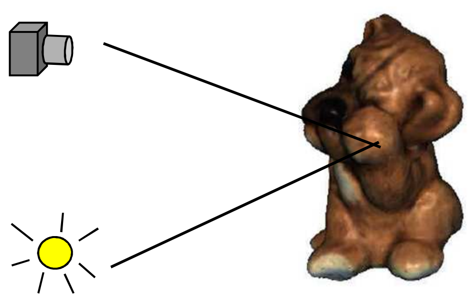

Shape and Materials by Example

[Hertzmann, Seitz CVPR 2003 PAMI 2005]

Reconstructs objects with general BRDF with no illumination info.
Idea : A reference object from the same material but with known geometry (sphere) is inserted into the scene.

Reference images

Multiple materials

Results

Summary of image cues

	Reflectance	Light	+	-
stereo	textured Lambertian	Constant **[SAD]**	Rec. texture Rec. depth discont. Complete obj	Needs texture Occlusions
		Varying **[NCC]**		
shading	uniform Lamb	Constant **[SFS]**		Uniform material Not robust Needs light pose
	unif/textured Lamb	Varying **[PS]**	Unif/varying albedo	Do not reconstr depth disc., one view