

# Photometric stereo and shape from shading



# Course announcements

- Homework 3 has been posted and is due on March 9<sup>th</sup>.
  - Any questions about the homework?
  - How many of you have looked at/started/finished homework 3?
- Office hours for Yannis' this week: Wednesday 3-5 pm.
- Gaurav's office hours now happen on Smith 200.
- Many more talks this week:
  1. Manolis Savva, "Human-centric Understanding of 3D Environments," Wednesday March 7, 2:00 PM, NSH 3305.
  2. David Fouhey, "Recovering a Functional and Three Dimensional Understanding of Images," Thursday March 8, 4:00 PM, NSH 3305.

# Overview of today's lecture

- Light sources.
- Shape from shading.
- Photometric stereo.


# Slide credits

Most of these slides were adapted from:

- Srinivasa Narasimhan (16-385, Spring 2014).
- Todd Zickler (Harvard University).
- Steven Gortler (Harvard University).

Light sources

# “Physics-based” computer vision (a.k.a “inverse optics”)




**I**  $\longrightarrow$  shape, illumination, reflectance

# Lighting models: Plenoptic function

- Radiance as a function of position and direction
- Radiance as a function of position, direction, and time
- Spectral radiance as a function of position, direction, time and wavelength

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



**Fig.1.3**


The plenoptic function describes the information available to an observer at any point in space and time. Shown here are two schematic eyes-which one should consider to have punctate pupils-gathering pencils of light rays. A real observer cannot see the light rays coming from behind, but the plenoptic function does include these rays.

# Lighting models: far-field approximation

- Assume that, over the observed region of interest, all source of incoming flux are relatively far away

$$L(x, \omega, t, \lambda) \longrightarrow L(\omega, t, \lambda)$$

$$L(x, \omega) \longrightarrow L(\omega)$$






# Lighting models: far-field approximation

- Assume that, over the observed region of interest, all source of incoming flux are relatively far away

$$L(x, \omega, t, \lambda) \longrightarrow L(\omega, t, \lambda)$$

$$L(x, \omega) \longrightarrow L(\omega)$$




# Lighting models: far-field approximation


- Assume that, over the observed region of interest, all source of incoming flux are relatively far away

$$L(x, \omega, t, \lambda) \longrightarrow L(\omega, t, \lambda)$$



$$L(x, \omega) \longrightarrow L(\omega)$$



# Application: augmented reality



# Application: augmented reality



# Application: augmented reality



**(a)** Background photograph



**(b)** Camera calibration grid and light probe



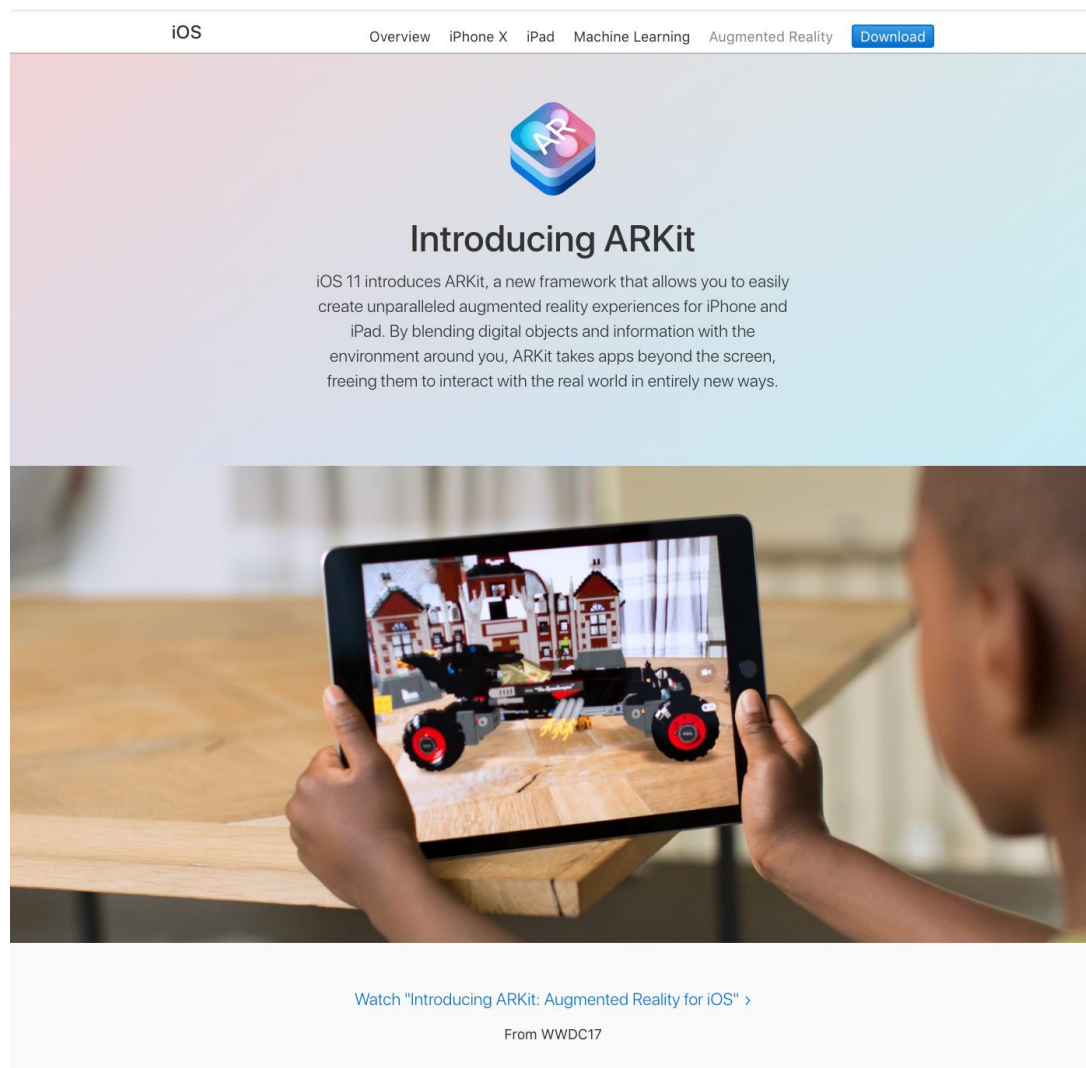
**(g)** Final result with differential rendering

# Application: augmented reality

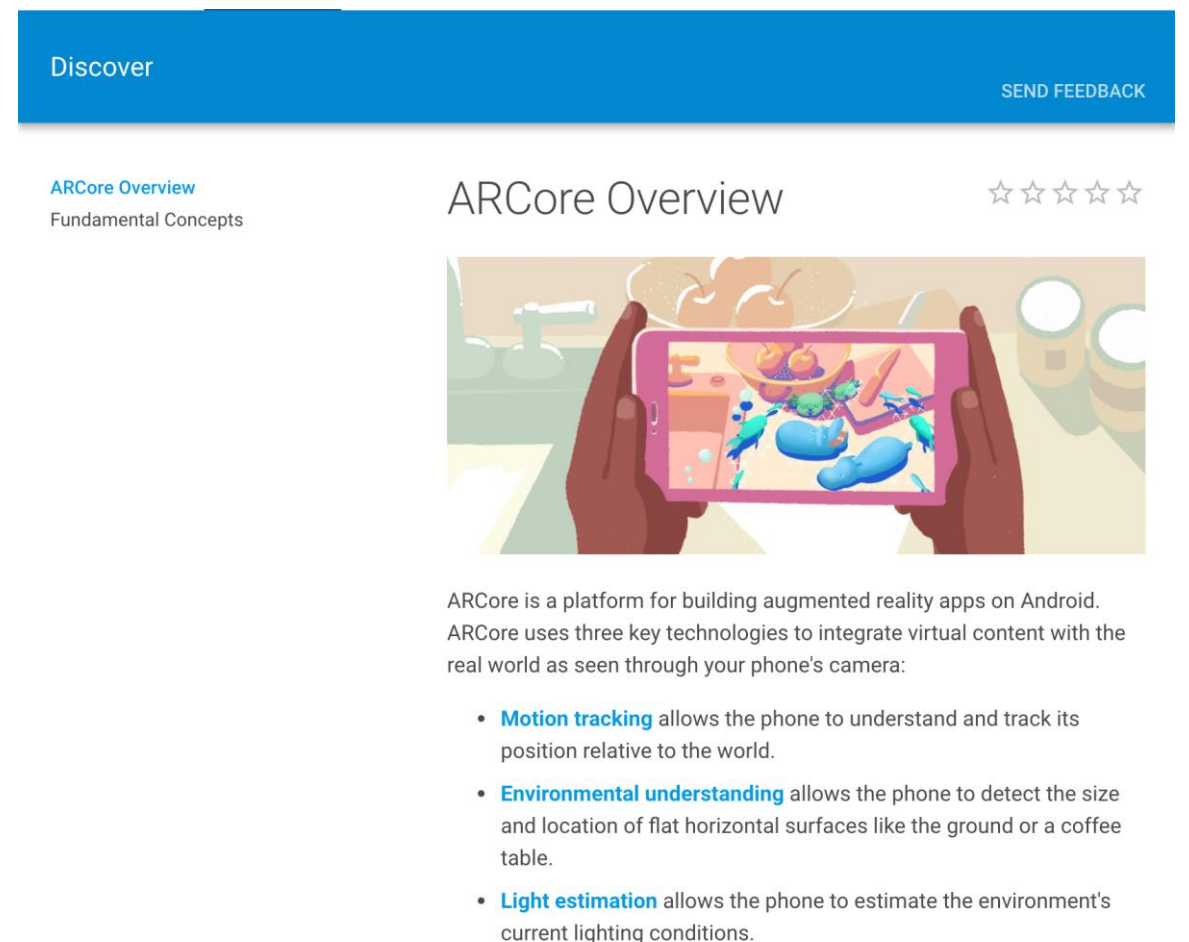




# Application: augmented reality



<https://developer.apple.com/arkit/>



<https://developers.google.com/ar/>

# Lighting models: far-field approximation

- One can download far-field lighting environments that have been captured by others

[\[http://gl.ict.usc.edu/Data/HighResProbes/\]](http://gl.ict.usc.edu/Data/HighResProbes/)

- A number of apps and software exist to help you capture your own environments using a light probe





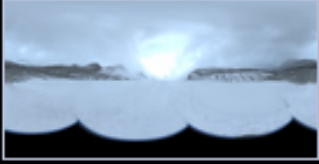

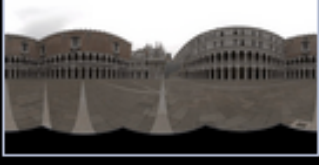



Figure 6. To produce the equal-area cylindrical projection of a spherical map, one projects each point on the surface of the sphere horizontally outward onto the cylinder, and then unwraps the cylinder to obtain a rectangular “panoramic” map.

**TABLE OF LIGHT PROBES:**

Image	Description	Interactive Preview	Download
<i>Uffizi Gallery, Italy</i>			
	Assembled from 18 14mm images taken using the Kodak DCS 520 camera	LDR panorama HDR panorama	HDR (7.3MB) EXR (7.9MB) Diffuse convolution
<i>Grace Cathedral, San Francisco, California</i>			
	Assembled from three 8mm fisheye images taken using the Canon EOS-1ds camera	LDR panorama HDR panorama	HDR (14MB) EXR (16MB) Diffuse convolution
<i>Dining room of the Ennis-Brown House, Los Angeles, California (website)</i>			
	Assembled from six 8mm fisheye images taken using the Canon d60 camera	LDR panorama HDR panorama	HDR (54MB) EXR (61MB) Diffuse convolution
<i>On a glacier in Banff National Forest, Canada</i>			
	Assembled from three 8mm fisheye images taken using the Canon EOS-1ds camera	LDR panorama HDR panorama	HDR (4.3MB) EXR (4.5MB) Diffuse convolution
<i>Pisa courtyard nearing sunset, Italy</i>			
	Assembled from three 8mm fisheye images taken using the Canon 5D camera	LDR panorama HDR panorama	HDR (20MB) EXR (22MB) Diffuse convolution
<i>Courtyard of the Doge's palace, Venice, Italy</i>			
	Assembled from five 8mm fisheye images taken using the Canon 5D camera	LDR panorama HDR panorama	HDR (22MB) EXR (19MB) Diffuse convolution




# Application: inferring outdoor illumination



From a single image (left), we estimate the most likely sky appearance (middle) and insert a 3-D object (right). Illumination estimation was done entirely automatically.

# A further simplification: Low-frequency illumination

$$L(\omega) = \sum_i a_i Y_i(\omega)$$



First nine basis  
functions are sufficient  
for re-creating  
Lambertian  
appearance

# Low-frequency illumination




Fig. 2. On the left, a white sphere illuminated by three directional (distant point) sources of light. All the lights are parallel to the image plane, one source illuminates the sphere from above and the two others illuminate the sphere from diagonal directions. In the middle, a cross-section of the lighting function with three peaks corresponding to the three light sources. On the right, a cross-section indicating how the sphere reflects light. We will make precise the intuition that the material acts as a low-pass filtering, smoothing the light as it reflects it.





Figure 3. Plot of spherical harmonic terms in Lambertian BRDF filter.

# Low-frequency illumination

$$L(\omega) = \sum_i a_i Y_i(\omega)$$



Truncate to first 9 terms


$$\vec{\ell} = (\ell_1, \dots, \ell_9)$$

# Application: Trivial rendering


Capture light probe



Rendering a (convex) diffuse object in this environment simply requires a lookup based on the surface normal at each pixel



Low-pass filter (truncate to first nine SHs)





# White-out: Snow and Overcast Skies



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



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

WHY?

# Diffuse Reflection from Uniform Sky

$$L^{surface}(\theta_r, \phi_r) = \int_{-\pi}^{\pi} \int_0^{\pi/2} L^{src}(\theta_i, \phi_i) f(\theta_i, \phi_i; \theta_r, \phi_r) \cos \theta_i \sin \theta_i d\theta_i d\phi_i$$

- Assume Lambertian Surface with Albedo = 1 (no absorption)

$$f(\theta_i, \phi_i; \theta_r, \phi_r) = \frac{1}{\pi}$$

- Assume Sky radiance is constant

$$L^{src}(\theta_i, \phi_i) = L^{sky}$$

- Substituting in above Equation:

$$L^{surface}(\theta_r, \phi_r) = L^{sky}$$

Radiance of any patch is the same as Sky radiance !! (white-out condition)

# Even simpler: Directional lighting

- Assume that, over the observed region of interest, all source of incoming flux is from one direction

$$L(x, \omega, t, \lambda) \longrightarrow L(x, t, \lambda) \longrightarrow s(t, \lambda) \delta(\omega = \omega_o(t))$$

$$L(x, \omega) \longrightarrow L(\omega) \longrightarrow s \delta(\omega = \omega_o)$$

- Convenient representation

$$\vec{\ell} = (\ell_x, \ell_y, \ell_z)$$


“light direction”  $\hat{\ell} = \frac{\vec{\ell}}{||\vec{\ell}||}$

“light strength”  $||\vec{\ell}||$



# Simple shading

ASSUMPTION 1:  
LAMBERTIAN





ASSUMPTION 2:  
DIRECTIONAL LIGHTING



$$L^{\text{out}}(\hat{\omega}) = \int_{\Omega_{\text{in}}} f(\hat{\omega}_{\text{in}}, \hat{\omega}_{\text{out}}) L^{\text{in}}(\hat{\omega}_{\text{in}}) \cos \theta_{\text{in}} d\hat{\omega}_{\text{in}}$$


$$I = a \hat{\mathbf{n}}^{\top} \vec{\ell}$$



[Prados, 2004]

# “N-dot-I” shading

ASSUMPTION 1:  
LAMBERTIAN





ASSUMPTION 2:  
DIRECTIONAL LIGHTING



$$L^{\text{out}}(\hat{\omega}) = \int_{\Omega_{\text{in}}} f(\hat{\omega}_{\text{in}}, \hat{\omega}_{\text{out}}) L^{\text{in}}(\hat{\omega}_{\text{in}}) \cos \theta_{\text{in}} d\hat{\omega}_{\text{in}}$$



$$I = a \hat{\mathbf{n}}^{\top} \vec{\ell}$$



[Prados, 2004]

# An ideal point light source

$$L(\mathbf{x}, \boldsymbol{\omega}) = \frac{s}{\|\mathbf{x} - \mathbf{x}_o\|^2} \delta \left( \boldsymbol{\omega} = \frac{\mathbf{x} - \mathbf{x}_o}{\|\mathbf{x} - \mathbf{x}_o\|} \right)$$



Think of this as a spatially-varying directional source where


1. the direction is away from  $\mathbf{x}_o$
2. the strength is proportional to  $1/(\text{distance})^2$

# Summary of some useful lighting models

- plenoptic function (function on 5D domain)
- far-field illumination (function on 2D domain)
- low-frequency far-field illumination (nine numbers)
- directional lighting (three numbers = direction and strength)
- point source (four numbers = location and strength)

Shape from shading

# Image Intensity and 3D Geometry



- Shading as a cue for shape reconstruction
- What is the relation between intensity and shape?
  - Reflectance Map



# Application: Detecting composite photos

---

Real photo




Fake photo



# “N-dot-I” shading

ASSUMPTION 1:  
LAMBERTIAN





ASSUMPTION 2:  
DIRECTIONAL LIGHTING



$$L^{\text{out}}(\hat{\omega}) = \int_{\Omega_{\text{in}}} f(\hat{\omega}_{\text{in}}, \hat{\omega}_{\text{out}}) L^{\text{in}}(\hat{\omega}_{\text{in}}) \cos \theta_{\text{in}} d\hat{\omega}_{\text{in}}$$


$$I = a \hat{\mathbf{n}}^{\top} \vec{\ell}$$



[Prados, 2004]



# Surface Normal



surface normal

$\mathbf{N}$

Equation of plane

or

$$Ax + By + Cz + D = 0$$

$$\frac{A}{C}x + \frac{B}{C}y + z + \frac{D}{C} = 0$$


Let

$$-\frac{\partial z}{\partial x} = \frac{A}{C} = p \quad -\frac{\partial z}{\partial y} = \frac{B}{C} = q$$

Surface normal

$$\mathbf{N} = \left( \frac{A}{C}, \frac{B}{C}, 1 \right) = (p, q, 1)$$

# Gradient Space



Normal vector

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

Source vector

$$\mathbf{s} = \frac{\mathbf{S}}{|\mathbf{S}|} = \frac{(p_s, q_s, 1)}{\sqrt{p_s^2 + q_s^2 + 1}}$$

$$\cos q_i = \mathbf{n} \cdot \mathbf{s} = \frac{(pp_s + qq_s + 1)}{\sqrt{p^2 + q^2 + 1} \sqrt{p_s^2 + q_s^2 + 1}}$$

$z = 1$  plane is called the Gradient Space ( $pq$  plane)

- Every point on it corresponds to a particular surface orientation

# Reflectance Map

- Relates image irradiance  $I(x,y)$  to surface orientation  $(p,q)$  for given source direction and surface reflectance
- Lambertian case:

$k$  : source brightness


$r$  : surface albedo (reflectance)

$c$  : constant (optical system)

Image irradiance:

$$I = \frac{r}{\rho} kc \cos q_i = \frac{r}{\rho} kc \mathbf{n} \times \mathbf{s}$$

Let  $\frac{r}{\rho} kc = 1$  then  $I = \cos q_i = \mathbf{n} \times \mathbf{s}$




# Reflectance Map

- Lambertian case

$$I = \cos q_i = \mathbf{n} \times \mathbf{s} = \frac{(pp_s + qq_s + 1)}{\sqrt{p^2 + q^2 + 1} \sqrt{p_s^2 + q_s^2 + 1}} = R(p, q)$$


Iso-brightness contour

Reflectance Map  
(Lambertian)



# Reflectance Map


- Lambertian case



Note:  $R(p, q)$  is maximum when  $(p, q) = (p_s, q_s)$


# Shape from a Single Image?

- 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

# Human Perception





# Human Perception




*2 possible illumination hypotheses*




**Examples of the classic bump/dent stimuli used to test lighting assumptions when judging shape from shading, with shading orientations (a)  $0^\circ$  and (b)  $180^\circ$  from the vertical.**

**a**



**b**



# Human Perception

- Our brain often perceives shape from shading.
- Mostly, it makes many assumptions to do so.
- For example:


Light is coming from above (sun).

Biased by occluding contours.

by V. Ramachandran

# Stereographic Projection


$(p, q)$ -space (gradient space)



Problem

$(p, q)$  can be infinite when  $q = 90^\circ$



$(f, g)$ -space



$$f = \frac{2p}{1 + \sqrt{1 + p^2 + q^2}} \quad g = \frac{2q}{1 + \sqrt{1 + p^2 + q^2}}$$

Redefine reflectance map as  $R(f, g)$

# Occluding Boundaries



$$\mathbf{n} \wedge \mathbf{e}, \quad \mathbf{n} \wedge \mathbf{v} \setminus \mathbf{n} = \mathbf{e} \times \mathbf{v} \quad \mathbf{e} \text{ and } \mathbf{v} \text{ are known}$$

The  $\mathbf{n}$  values on the occluding boundary can be used as the boundary condition for shape-from-shading

# Image Irradiance Constraint

- Image irradiance should match the reflectance map

Minimize

$$e_i = \iint_{\text{image}} (I(x, y) - R(f, g))^2 dx dy$$

(minimize errors in image irradiance in the image)

# Smoothness Constraint

- Used to constrain shape-from-shading
- Relates orientations  $(f, g)$  of neighboring surface points

Minimize

$$e_s = \iint_{\text{image}} \left( f_x^2 + f_y^2 \right) + \left( g_x^2 + g_y^2 \right) dx dy$$

$(f, g)$ : surface orientation under stereographic projection

$$f_x = \frac{\nabla f}{\|\nabla f\|}, f_y = \frac{\nabla f}{\|\nabla f\|}, g_x = \frac{\nabla g}{\|\nabla g\|}, g_y = \frac{\nabla g}{\|\nabla g\|}$$

(penalize rapid changes in surface orientation  $f$  and  $g$  over the image)

# Shape-from-Shading

- Find surface orientations  $(f, g)$  at all image points that minimize

$$e = e_s + \text{weight} / e_i$$

smoothness constraint      image irradiance error

Minimize

$$e = \iint_{\text{image}} \left( f_x^2 + f_y^2 \right) + \left( g_x^2 + g_y^2 \right) + \left( I(x, y) - R(f, g) \right)^2 dx dy$$



# Numerical Shape-from-Shading

- **Smoothness error** at image point  $(i,j)$

$$s_{i,j} = \frac{1}{4} \left( (f_{i+1,j} - f_{i,j})^2 + (f_{i,j+1} - f_{i,j})^2 + (g_{i+1,j} - g_{i,j})^2 + (g_{i,j+1} - g_{i,j})^2 \right)$$

Of course you can consider more neighbors (smoother results)

- **Image irradiance error** at image point  $(i,j)$



$$r_{i,j} = \left( I_{i,j} - R(f_{i,j}, g_{i,j}) \right)^2$$

Find  $\{f_{i,j}\}$  and  $\{g_{i,j}\}$  that minimize

$$e = \sum_i \sum_j (s_{i,j} + \lambda r_{i,j})$$


(Ikeuchi & Horn 89)

# Results





by Ikeuchi and Horn

# Results



Scanning Electron Microscope image  
(inverse intensity)

by Ikeuchi and Horn





# More modern results



Resolution: 640 x 500; Re-rendering Error: 0.0075.

---




Resolution: 590 x 690; Re-rendering Error: 0.0083.

---

# Single-lighting is ambiguous

ASSUMPTION 1:  
LAMBERTIAN





ASSUMPTION 2:  
DIRECTIONAL LIGHTING



$$L^{\text{out}}(\hat{\omega}) = \int_{\Omega_{\text{in}}} f(\hat{\omega}_{\text{in}}, \hat{\omega}_{\text{out}}) L^{\text{in}}(\hat{\omega}_{\text{in}}) \cos \theta_{\text{in}} d\hat{\omega}_{\text{in}}$$

$$I = a \hat{\mathbf{n}}^{\top} \vec{\ell}$$




[Prados, 2004]

Photometric stereo

# Lambertian photometric stereo


$$\begin{aligned} I_1 &= a \hat{\mathbf{n}}^\top \vec{\ell}_1 \\ I_2 &= a \hat{\mathbf{n}}^\top \vec{\ell}_2 \\ &\vdots \\ I_N &= a \hat{\mathbf{n}}^\top \vec{\ell}_N \end{aligned}$$





# Lambertian photometric stereo

$$\begin{aligned} I_1 &= a \hat{\mathbf{n}}^\top \vec{\ell}_1 \\ I_2 &= a \hat{\mathbf{n}}^\top \vec{\ell}_2 \\ &\vdots \\ I_N &= a \hat{\mathbf{n}}^\top \vec{\ell}_N \end{aligned}$$




define “pseudo-normal”  $\vec{\mathbf{b}} \triangleq a \hat{\mathbf{n}}$

solve linear system  
for pseudo-normal

$$\begin{bmatrix} I_1 \\ I_2 \\ \vdots \\ I_N \end{bmatrix}_{N \times 1} = \begin{bmatrix} \vec{\ell}_1^\top \\ \vec{\ell}_2^\top \\ \vdots \\ \vec{\ell}_N^\top \end{bmatrix}_{N \times 3} \begin{bmatrix} \vec{\mathbf{b}} \end{bmatrix}_{3 \times 1}$$

once system is solved,  
 $\mathbf{b}$  gives normal  
direction and albedo

# Photometric Stereo



# Solving the Equations

$$\underbrace{\begin{pmatrix} \hat{e}_1^T I_1 \hat{u} \\ \hat{e}_2^T I_2 \hat{u} \\ \hat{e}_3^T I_2 \hat{u} \end{pmatrix}}_{\mathbf{I}_{3 \times 1}} = \underbrace{\begin{pmatrix} \hat{e}_1^T \mathbf{s}_1^T \hat{u} \\ \hat{e}_2^T \mathbf{s}_2^T \hat{u} \\ \hat{e}_3^T \mathbf{s}_3^T \hat{u} \end{pmatrix}}_{\mathbf{S}_{3 \times 3}} \underbrace{\hat{u}}_{\tilde{\mathbf{n}}_{3 \times 1}} r \mathbf{n}$$

$$\tilde{\mathbf{n}} = \mathbf{S}^{-1} \mathbf{I}$$

inverse

$$r = |\tilde{\mathbf{n}}|$$

$$\mathbf{n} = \frac{\tilde{\mathbf{n}}}{|\tilde{\mathbf{n}}|} = \frac{\tilde{\mathbf{n}}}{r}$$

# More than Three Light Sources

- Get better results by using more lights

$$\begin{bmatrix} I_1 \\ \vdots \\ I_N \end{bmatrix} = \begin{bmatrix} \mathbf{s}^T \\ \vdots \\ \mathbf{s}_N^T \end{bmatrix} \mathbf{n}$$

- Least squares solution:

$$\mathbf{I} = \mathbf{S} \tilde{\mathbf{n}} \quad \longleftarrow \quad N \times 1 = (\underline{N \times 3})(3 \times 1)$$

$$\mathbf{S}^T \mathbf{I} = \mathbf{S}^T \mathbf{S} \tilde{\mathbf{n}}$$


$$\tilde{\mathbf{n}} = \left( \mathbf{S}^T \mathbf{S} \right)^{-1} \mathbf{S}^T \mathbf{I}$$

- Solve for  $\mathbf{r}, \mathbf{n}$  as before

Moore-Penrose pseudo inverse

# Computing light source directions


- Trick: place a chrome sphere in the scene



- the location of the highlight tells you the source direction

# Specular Reflection - Recap

- For a perfect mirror, light is reflected about **N**




$$R_e = \begin{cases} R_i & \text{if } \mathbf{v} = \mathbf{r} \\ 0 & \text{otherwise} \end{cases}$$

- We see a highlight when  $\mathbf{v} = \mathbf{r}$
- Then **S** is given as follows:

$$\mathbf{s} = 2(\mathbf{n} \times \mathbf{r})\mathbf{n} - \mathbf{r}$$

# Computing the Light Source Direction

Chrome sphere that has a highlight at position  $\mathbf{h}$  in the image



- Can compute  $\mathbf{N}$  by studying this figure
  - Hints:
    - use this equation:  $\|\mathbf{H} - \mathbf{C}\| = r$
    - can measure  $\mathbf{c}$ ,  $\mathbf{h}$ , and  $r$  in the image



# Limitations

- Big problems
  - Doesn't work for shiny things, semi-translucent things
  - Shadows, inter-reflections
- Smaller problems
  - Camera and lights have to be distant
  - Calibration requirements
    - measure light source directions, intensities
    - camera response function

# Trick for Handling Shadows

- Weight each equation by the pixel brightness:

$$I_i (I_i) = I_i (r \mathbf{n} \times \mathbf{s}_i)$$

- Gives weighted least-squares matrix equation:

$$\begin{pmatrix} I_1^2 \\ \vdots \\ I_N^2 \end{pmatrix} = \begin{pmatrix} I_1 \mathbf{s}_1^T \\ \vdots \\ I_N \mathbf{s}_N^T \end{pmatrix} r \mathbf{n}$$

- Solve for  $r, \mathbf{n}$  as before

# Depth from normals

---


- Solving the linear system per-pixel gives us an estimated surface normal for each pixel



Input photo



Estimated normals




Estimated normals  
(needle diagram)

- How can we compute depth from normals?
  - Normals are like the “derivative” of the true depth

# Normal Integration

---


- Integrating a set of derivatives is easy in 1D
  - (similar to Euler's method from diff. eq. class)



- Could just integrate normals in each column / row separately
- Instead, we formulate as a linear system and solve for depths that *best agree with the surface normals*

# Depth from normals

---



$$\begin{aligned} V_1 &= (x+1, y, z_{x+1,y}) - (x, y, z_{xy}) \\ &= (1, 0, z_{x+1,y} - z_{xy}) \end{aligned}$$

$$\begin{aligned} 0 &= N \cdot V_1 \\ &= (n_x, n_y, n_z) \cdot (1, 0, z_{x+1,y} - z_{xy}) \\ &= n_x + n_z(z_{x+1,y} - z_{xy}) \end{aligned}$$

Get a similar equation for  $V_2$


- Each normal gives us two linear constraints on  $z$
- compute  $z$  values by solving a matrix equation

# Results



1. Estimate light source directions
2. Compute surface normals
3. Compute albedo values
4. Estimate depth from surface normals
5. Relight the object (with original texture and uniform albedo)


# Results: Lambertian Sphere




Input Images



Needles are projections  
of surface normals on  
image plane



Estimated Surface Normals





Estimated Albedo





# Lambertain Mask



# Results – Albedo and Surface Normal



# Results – Shape of Mask






# Results: Lambertian Toy

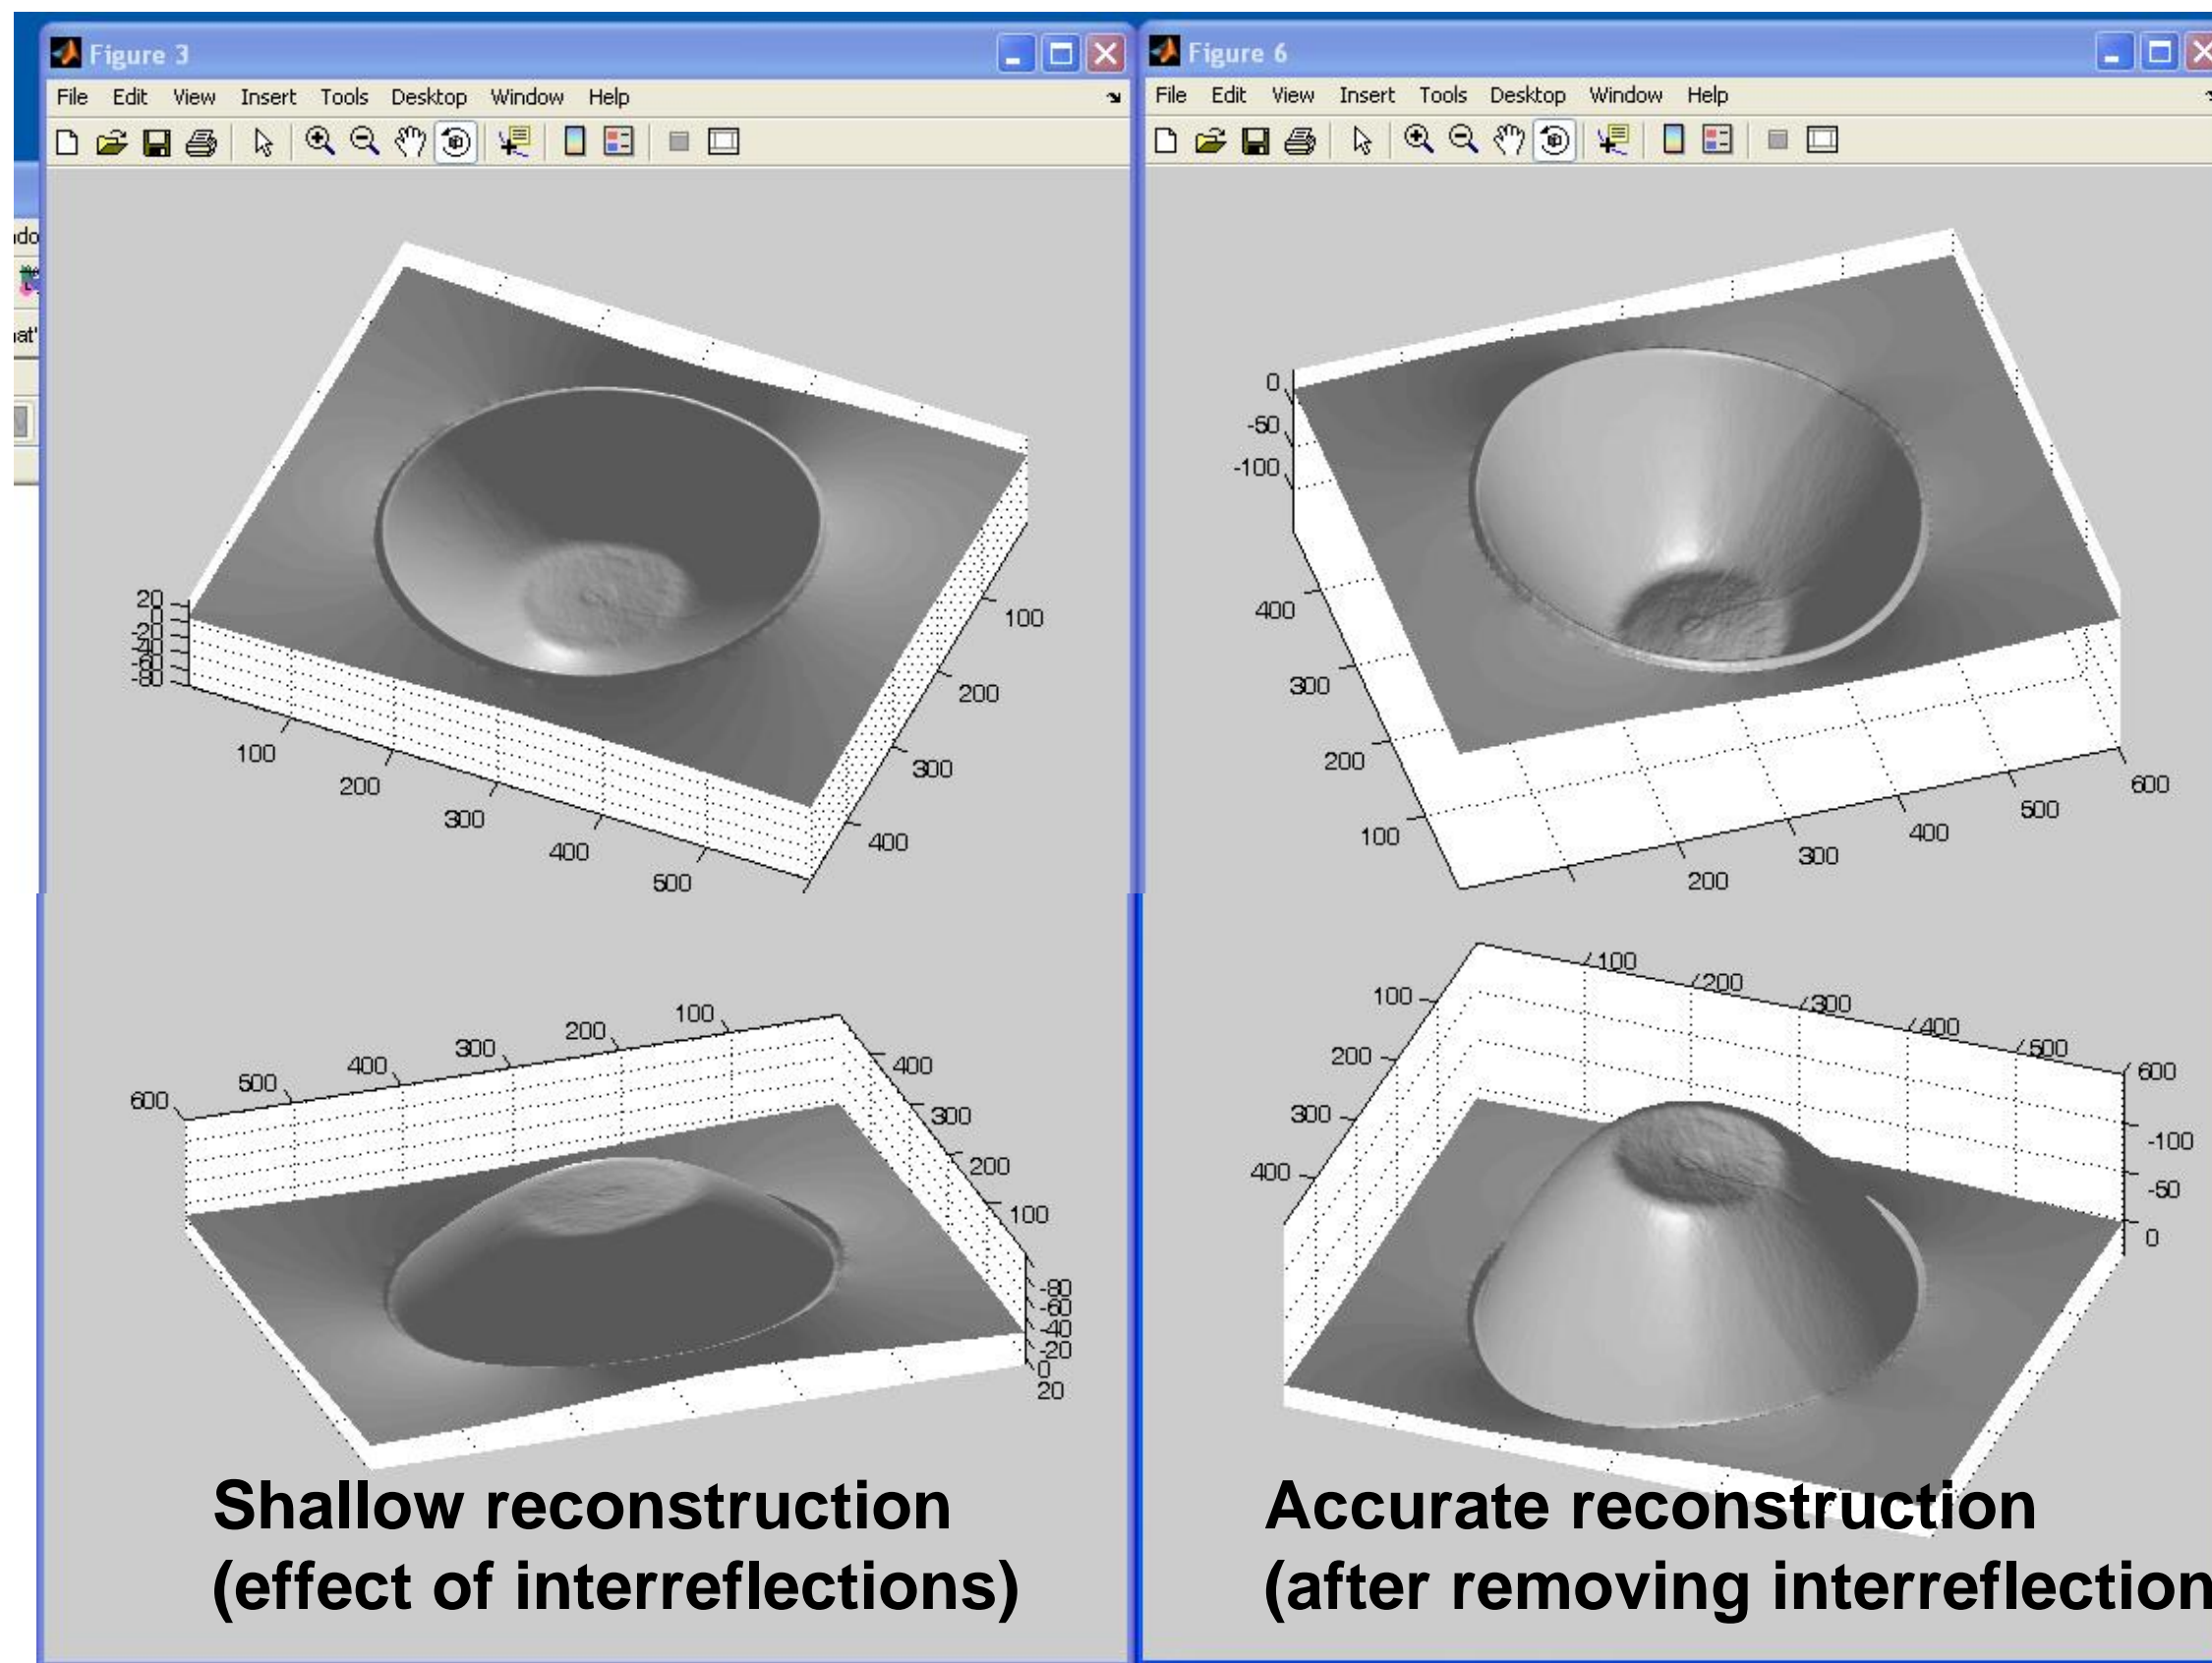




# Original Images



# Results - Shape





What if the light directions are unknown?



# What if the light directions are unknown?

$$\begin{aligned} I_1 &= a \hat{\mathbf{n}}^\top \vec{\ell}_1 \\ I_2 &= a \hat{\mathbf{n}}^\top \vec{\ell}_2 \\ &\vdots \\ I_N &= a \hat{\mathbf{n}}^\top \vec{\ell}_N \end{aligned}$$





define “pseudo-normal”  $\vec{\mathbf{b}} \triangleq a \hat{\mathbf{n}}$

solve linear system  
for pseudo-normal

$$\begin{bmatrix} I_1 \\ I_2 \\ \vdots \\ I_N \end{bmatrix}_{N \times 1} = \begin{bmatrix} \vec{\ell}_1^\top \\ \vec{\ell}_2^\top \\ \vdots \\ \vec{\ell}_N^\top \end{bmatrix}_{N \times 3} \begin{bmatrix} \vec{\mathbf{b}} \end{bmatrix}_{3 \times 1}$$

# What if the light directions are unknown?

$$\begin{aligned} I_1 &= a \hat{\mathbf{n}}^\top \vec{\ell}_1 \\ I_2 &= a \hat{\mathbf{n}}^\top \vec{\ell}_2 \\ &\vdots \\ I_N &= a \hat{\mathbf{n}}^\top \vec{\ell}_N \end{aligned}$$





define “pseudo-normal”  $\vec{\mathbf{b}} \triangleq a \hat{\mathbf{n}}$

solve linear system  
for pseudo-normal at  
each image pixel

$$\begin{bmatrix} I_1 \\ I_2 \\ \vdots \\ I_N \end{bmatrix}_{N \times M} = \begin{bmatrix} \vec{\ell}_1^\top \\ \vec{\ell}_2^\top \\ \vdots \\ \vec{\ell}_N^\top \end{bmatrix}_{N \times 3} \begin{bmatrix} B \end{bmatrix}_{3 \times M}$$

# What if the light directions are unknown?

$$\begin{aligned} I_1 &= a \hat{\mathbf{n}}^\top \vec{\ell}_1 \\ I_2 &= a \hat{\mathbf{n}}^\top \vec{\ell}_2 \\ &\vdots \\ I_N &= a \hat{\mathbf{n}}^\top \vec{\ell}_N \end{aligned}$$




define “pseudo-normal”  $\vec{b} \triangleq a \hat{\mathbf{n}}$

solve linear system  
for pseudo-normal at  
each image pixel

$$\begin{bmatrix} I_1 \\ I_2 \\ \vdots \\ I_N \end{bmatrix}_{N \times M} = \begin{bmatrix} \vec{\ell}_1^\top \\ \vec{\ell}_2^\top \\ \vdots \\ \vec{\ell}_N^\top \end{bmatrix}_{N \times 3} \begin{bmatrix} B \end{bmatrix}_{3 \times M}$$


How do we solve this  
system without  
knowing light matrix  $L$ ?

# Factorizing the measurement matrix




# Factorizing the measurement matrix

- Singular value decomposition:



To reduce to rank 3, we just need to set all the singular values to 0 except for the first 3



This decomposition minimizes  $|\mathbf{I} - \mathbf{L}\mathbf{B}|^2$

Are the results unique?

# Bas-relief ambiguity

$$\mathbf{I} = \mathbf{L} \mathbf{B} = (\mathbf{L} \mathbf{Q}^{-1}) (\mathbf{Q} \mathbf{B})$$



(Belhumeur et al., 1999)


What assumptions have we made for all this?



# What assumptions have we made for all this?

- Lambertian BRDF
- Directional lighting
- No interreflections or scattering

# Shape independent of BRDF via reciprocity: “Helmholtz Stereopsis”



$$I = f(\text{shape}, \text{illumination}, \text{reflectance})$$

$$f^{-1} =$$



# What assumptions have we made for all this?

- Lambertian BRDF
- Directional lighting
- No interreflections or scattering

# References

Basic reading:

- Szeliski, Section 2.2, 12.1.
- Gortler, Chapter 21.