Photometric Stereo

Introduction to Computer Vision
CSE 152
Lecture 15

Midterm Exam Results

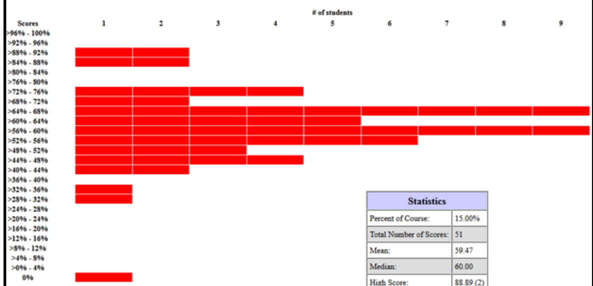

Statistics	
Percent of Course:	15.00%
Total Number of Scores:	51
Mean:	59.47
Median:	60.00
High Score:	88.89 (2)
Low Score:	31.11 (1)
Standard Deviation:	15.28

Announcements

- **Final exam will be a take home exam**
- Homework 2 has been graded
- Homework 3 is due May 29, 11:59 PM
- Homework 4 will be assigned this week

- Reading:
  - Section 12.1.1 Shape from shading and photometric stereo

Shading reveals 3-D surface geometry

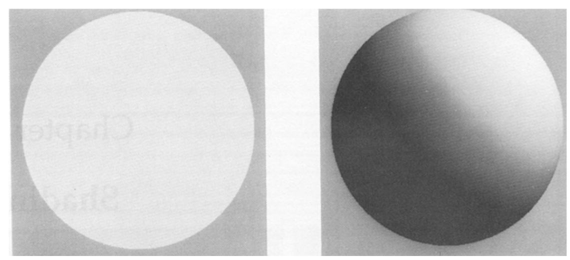

Two shape-from-X methods that use shading

- Shape-from-shading: Use just one image to recover shape. Requires knowledge of light source direction and BRDF everywhere. Too restrictive to be useful.

- Photometric stereo: Single viewpoint, multiple images under different lighting.

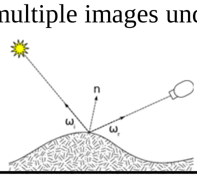

BRDF
(four dimensional function)

Photometric Stereo Rigs: One viewpoint, changing lighting

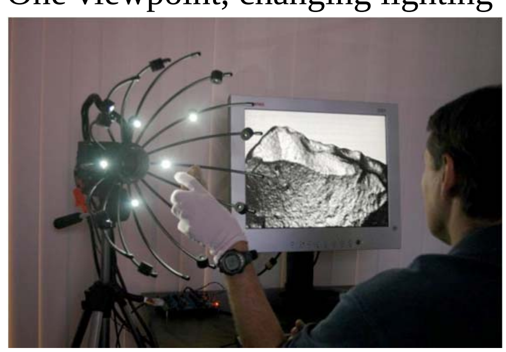

An example of photometric stereo

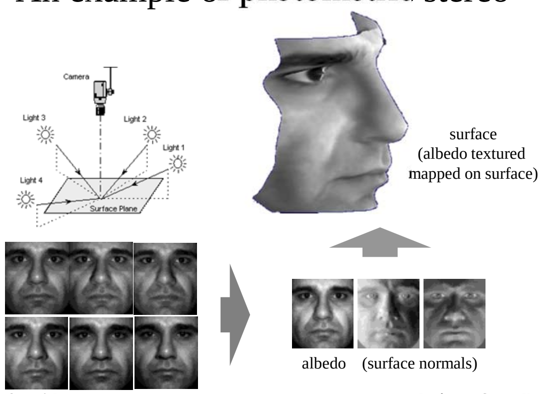

Multi-view stereo vs. Photometric Stereo: Assumptions

- Multi-view Stereo
  - Multiple images
  - Dynamic scene
  - Multiple viewpoints
  - Fixed lighting
- Photometric Stereo
  - Multiple images
  - Static scene
  - Fixed viewpoint
  - Multiple lighting conditions

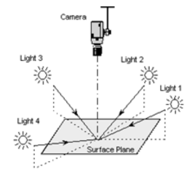

Photometric stereo

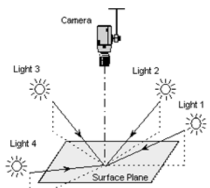

- Single viewpoint, multiple images under different lighting.
  1. Arbitrary known BRDF, known lighting
  2. Lambertian BRDF, known lighting
  3. Lambertian BRDF, unknown lighting.

I. Photometric Stereo: General BRDF and Reflectance Map

BRDF

- Bi-directional Reflectance Distribution Function
  $\rho(\theta_{in}, \phi_{in} ; \theta_{out}, \phi_{out})$
- Function of
  - Incoming light direction:
    θ_{in} , ϕ_{in}
  - Outgoing light direction:
    θ_{out} , ϕ_{out}
- Ratio of incident irradiance to emitted radiance

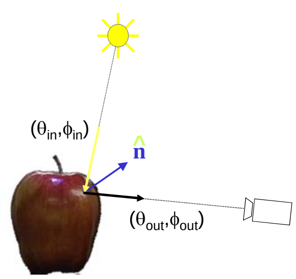

Coordinate system

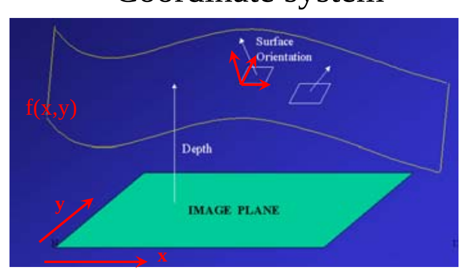

Surface: $\mathbf{s}(x,y) = (x, y, f(x,y))$

Tangent vectors: $\frac{\partial s(x,y)}{\partial x} = \left(1, 0, \frac{\partial f}{\partial x}\right)$

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

Normal vector

$\mathbf{n} = \frac{\partial s}{\partial x} \times \frac{\partial s}{\partial y}$

$= \left(\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y}, -1\right)$

Gradient Space (p,q)

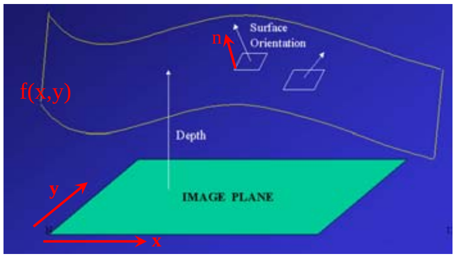

Gradient Space : (p,q)

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

Normal vector

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

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

Image Formation

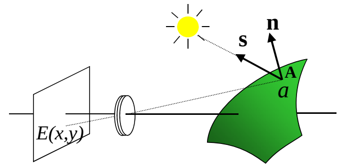

For a given point A on the surface, the image irradiance E(x,y) is a function of

1. The BRDF at A
2. The surface normal at A
3. The direction of the light source

Reflectance Map

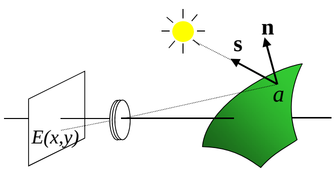

Let the BRDF be the same at all points on the surface, and let the light direction **s** be a constant.

1. Then image irradiance is a function of only the direction of the surface normal.
2. In gradient space, we have E(p,q).

Example Reflectance Map: Lambertian surface

E(p,q)

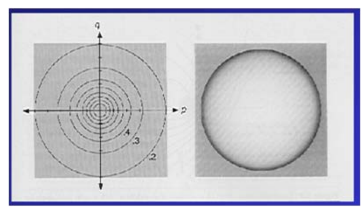

For lighting from front

LAMBERTIAN REFLECTANCE MAP

$$E = L\rho \frac{1 + p p_s + q q_s}{\sqrt{1+p^2+q^2}\sqrt{1+p_s^2+q_s^2}}$$

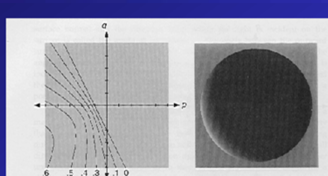

$p_s = -2 \quad q_s = -1$

Light Source Direction, expressed in gradient space.

Reflectance Map of Lambertian Surface

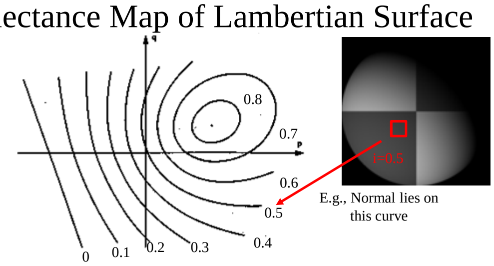

What does the intensity (Irradiance) of one pixel in one image tell us?

It constrains the surface normal projecting to that point to a curve

Two Light Sources
Two reflectance maps

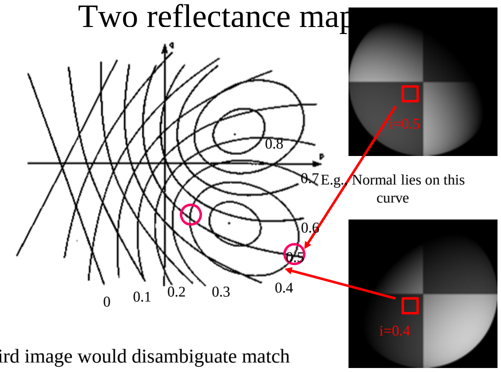

A third image would disambiguate match

Three Source Photometric stereo: Step1

Offline:

Using source directions & BRDF, construct reflectance map for each light source direction. $R_1(p,q)$, $R_2(p,q)$, $R_3(p,q)$

Online:

1. Acquire three images with known light source directions. $E_1(x,y)$, $E_2(x,y)$, $E_3(x,y)$
2. For each pixel location (x,y), find (p,q) as the intersection of the three curves
   $R_1(p,q)=E_1(x,y)$
   $R_2(p,q)=E_2(x,y)$
   $R_3(p,q)=E_3(x,y)$
3. This is the surface normal at pixel (x,y). Over image, the normal field is estimated

Normal Field


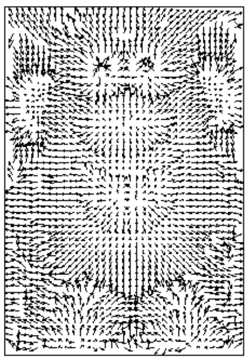

Plastic Baby Doll: Normal Field


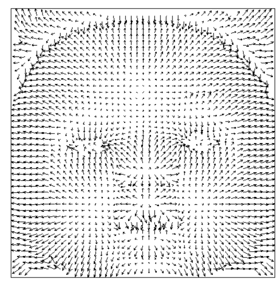

Next step:
Go from normal field to surface


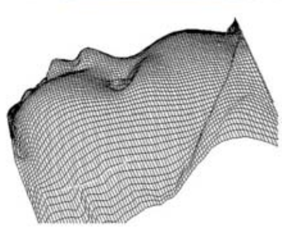

Recovering the surface f(x,y)

Many methods: Simplest approach

1. From estimate $\mathbf{n} = (n_x, n_y, n_z)$, $p = n_x/n_z$, $q = n_y/n_z$
2. Integrate p=df/dx along a row (x,0) to get f(x,0)
3. Then integrate q=df/dy along each column starting with value of the first row

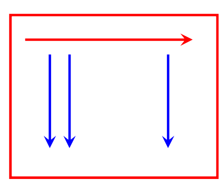

f(x,0)

II. Photometeric Stereo: Lambertian Surface, Known Lighting

Lambertian Surface

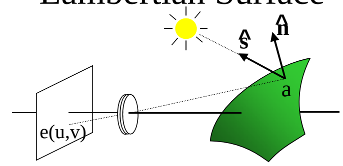

At image location (u,v), the intensity of a pixel x(u,v) is:

$$e(u,v) = [a(u,v)\hat{n}(u,v)] \cdot [s_0\hat{\mathbf{s}}] = \mathbf{b}(u,v) \cdot \mathbf{s}$$

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.
- $\hat{\mathbf{s}}$ is the direction to the light source.

Lambertian Photometric stereo

- If the light sources $\mathbf{s}_1$, $\mathbf{s}_2$, and $\mathbf{s}_3$ are **known**, then we **can** recover b from as few as three images. (Photometric Stereo: Silver 80, Woodham81).

$$[e_1 \;\; e_2 \;\; e_3] = \mathbf{b}^T[\mathbf{s}_1 \;\; \mathbf{s}_2 \;\; \mathbf{s}_3]$$

- i.e., we measure e_1, e_2, and e_3 and we know $\mathbf{s}_1$, $\mathbf{s}_2$, and $\mathbf{s}_3$. We can then solve for **b** by solving a linear system.

$$\mathbf{b}^T = [e_1 \;\; e_2 \;\; e_3][\mathbf{s}_1 \;\; \mathbf{s}_2 \;\; \mathbf{s}_3]^{-1}$$

- Normal is: **n** = **b**/|**b**|, albedo is: |**b**|

What if we have more than 3 Images? Linear Least Squares

$$[e_1 \;\; e_2 \;\; e_3 \ldots e_n] = \mathbf{b}^T[\mathbf{s}_1 \;\; \mathbf{s}_2 \;\; \mathbf{s}_3 \ldots \mathbf{s}_n]$$

Rewrite as

$$\mathbf{e} = \mathbf{Sb}$$

where

e is n by 1
b is 3 by 1
S is n by 3

Let the residual be

r=e-Sb

Squaring this:

$$r^2 = r^T r = (e-Sb)^T(e-Sb) = e^Te - 2b^TS^Te + b^TS^TSb$$

$(r^2)_b=0$ - zero derivative is a necessary condition for a minimum, or

$-2S^Te+2S^TSb=0;$

Solving for b gives

$$\mathbf{b} = (\mathbf{S}^T\mathbf{S})^{-1}\mathbf{S}^T\mathbf{e}$$

Input Images

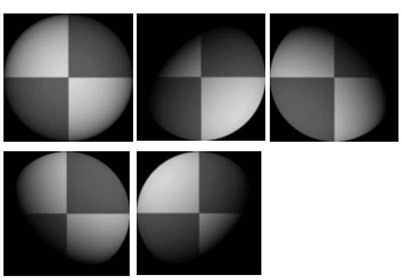

Recovered albedo

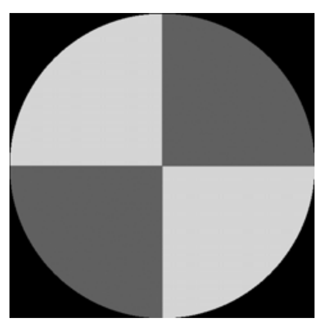

Recovered normal field

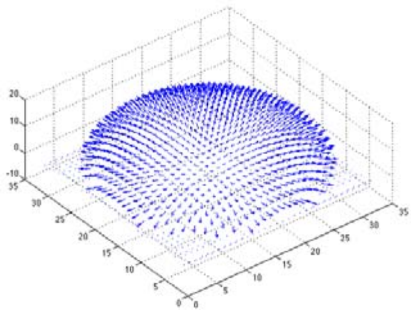

Surface recovered by integration

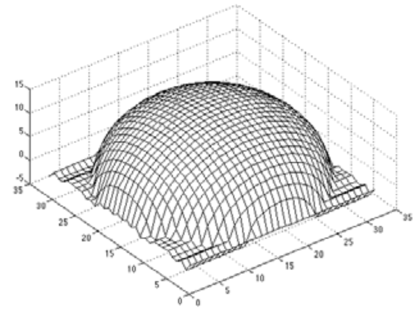

An example of photometric stereo

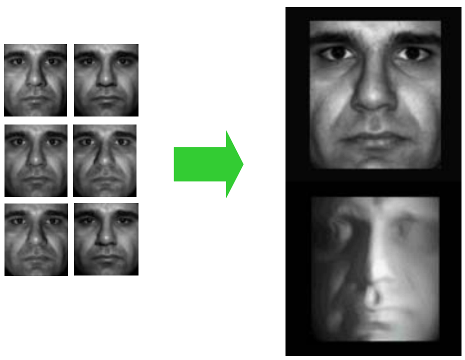

III. Photometric Stereo with unknown lighting and Lambertian surfaces

How do you construct subspace?

$[\mathbf{E}_1\ \mathbf{E}_2\ \mathbf{E}_3] = \mathbf{B}^T[\mathbf{s}_1\ \mathbf{s}_2\ \mathbf{s}_3]$

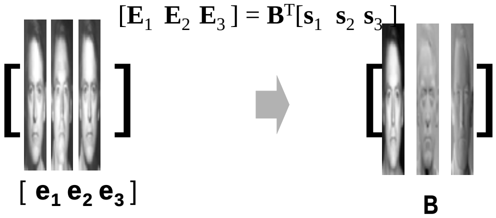

- Given three or more images $\mathbf{E}_1 \ldots \mathbf{E}_n$, estimate $\mathbf{B}$ and $\mathbf{s}_i$.
- How? Given images in form of $E=[E_1\ E_2 \ldots]$, Compute $[U,S,V] = SVD(E)$ and B^* is the n by 3 matrix formed by first 3 columns of U.

Face Basis

Original Images

Basis Images

Render Image from Basis Images

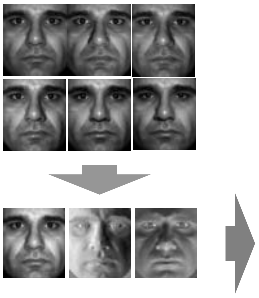

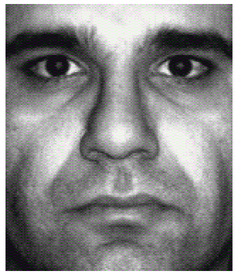

Rendered Image
(Single Light Source)

Do Ambiguities Exist? Yes

- Is B unique?
- For any invertible matrix **A**, **B* = BA** also a solution
- For any image of **B** produced with light source **S**, the same image can be produced by lighting **B*=BA** with **S*=A**$^{-1}$**S** because
  $$\mathbf{X} = \mathbf{B^*S^*} = \mathbf{B\,AA^{-1}S} = \mathbf{BS}$$
- When we estimate B using Singular Value Decomposition (SVD), the rows are NOT generally the normal times the albedo.

GBR Transformation

Only **Generalized Bas-Relief** transformations satisfy the integrability constraint:

$$\mathbf{A} = \mathbf{G}^T = \begin{bmatrix} \lambda & 0 & -\mu \\ 0 & \lambda & -\nu \\ 0 & 0 & 1 \end{bmatrix}^T$$

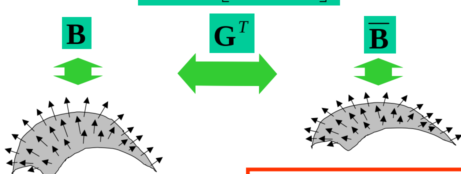

$f(x,y)$

$$\bar{f}(x,y) = \lambda f(x,y) + \mu x + \nu y$$

Generalized Bas-Relief Transformations

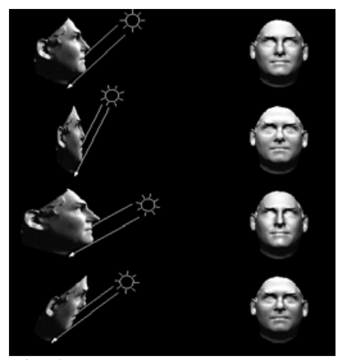

Without knowledge of light source location, one can only recover surfaces up to GBR transformations.

$$\bar{f}(x,y) = \lambda f(x,y) + \mu x + \nu y$$

f: true depth

$\bar{f}$: GBR transform of depth

Uncalibrated photometric stereo

1. Take n images as input without knowledge of light directions or strengths
2. Perform SVD to compute B*.
3. Find some A such that B*A is close to integrable.
4. Integrate resulting gradient field to obtain height function f*(x,y).

Comments:
- f*(x,y) differs from f(x,y) by a GBR.
- Can use specularities to resolve GBR for non-Lambertian surface.