Photometric Stereo

Introduction to Computer Vision
CSE 152
Lecture 4

Announcements

- Homework 1 is due today by 11:59 PM
- Homework 2 will be assigned today
  - Due Wed, Apr 26, 11:59 PM
  - To be completed individually, not in groups
- Reading:
  - Section 2.2.4: Photometric Stereo
    - Shape from Multiple Shaded Images

Shading reveals 3-D surface geometry

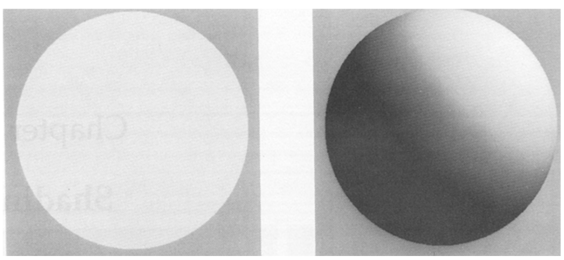

Two shape-from-X methods that use shading

- 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.

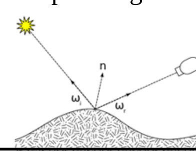

BRDF
(four dimensional function)

Photometric Stereo Rigs: One viewpoint, changing lighting

An example of photometric stereo

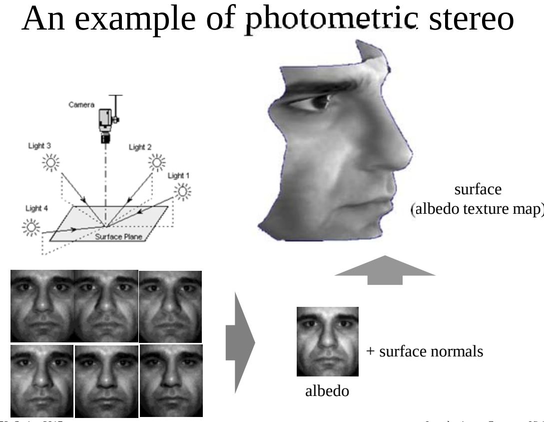

Photometric stereo

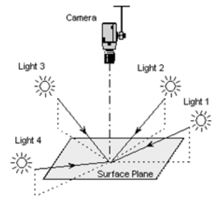

- Single viewpoint, multiple images under different lighting.
  1. Arbitrary known BRDF, known lighting
  2. Lambertian BRDF, known lighting
  3. Lambertian BRDF, unknown lighting.

1. 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

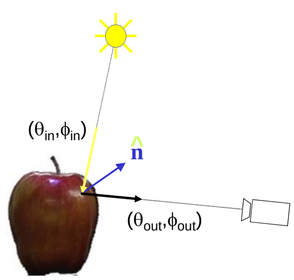

Coordinate system

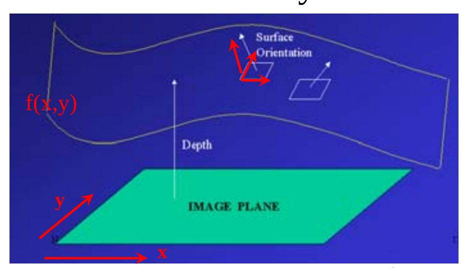

Surface: $\mathbf{s}(x,y) = (x, y, f(x,y))$

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)

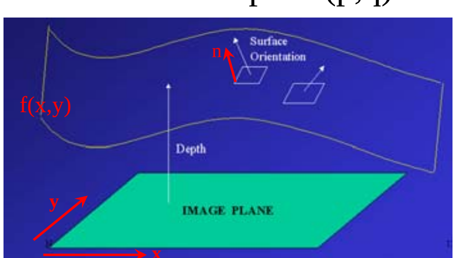

Gradient Space : (p,q)

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

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

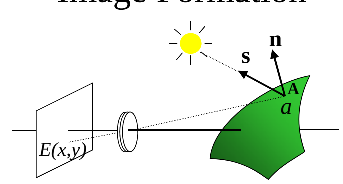

For a given point A on the surface *a*, the image irradiance E(x,y) is a function of

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

Reflectance Map

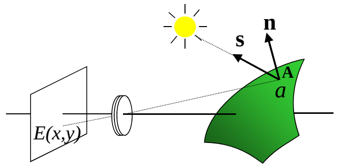

Let the BRDF be the same at all points on the surface, and let the light direction **s** be a constant.

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)

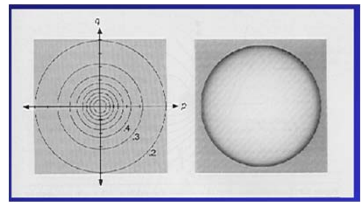

For lighting from front

LAMBERTIAN REFLECTANCE MAP

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

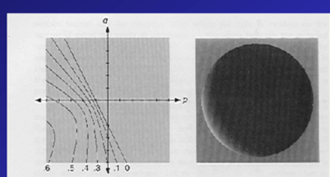

$p_s = -2 \quad q_s = -1$

Light Source Direction, expressed in gradient space.

Reflectance Map of Lambertian Surface

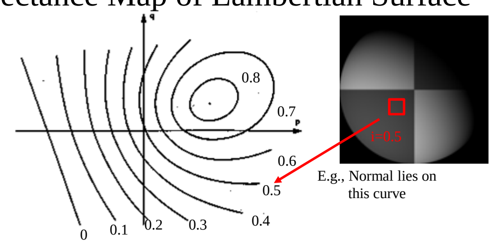

What does the intensity (irradiance) of one pixel in one image tell us?

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

Two Light Sources Two reflectance maps

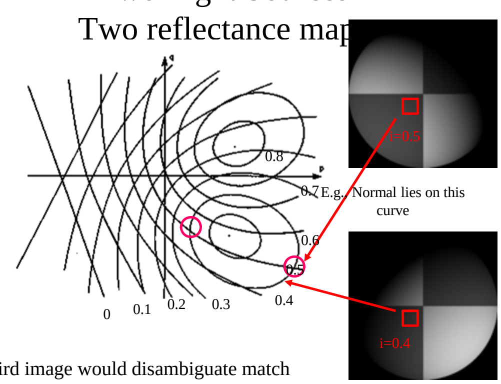

A third image would disambiguate match

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

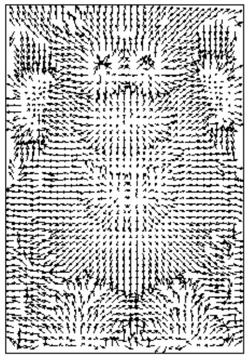

Plastic Baby Doll: Normal Field

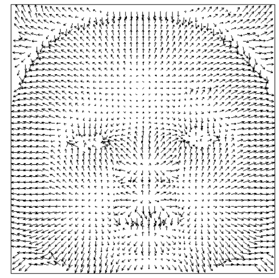

Next step: Go from normal field to surface

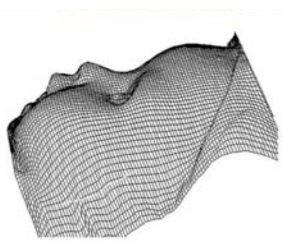

Recovering the surface f(x,y)

Many methods: Simplest approach

1. From estimate **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

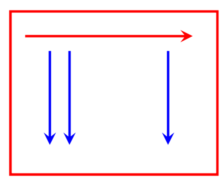

f(x,0)

What might go wrong?

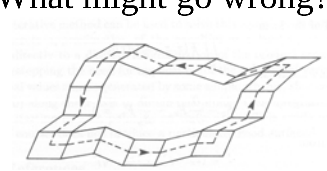

- Height z(x,y) is obtained by integration along a curve from (x_0, y_0).

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

- If one integrates the derivative field along any closed curve, on expects to get back to the starting value.
- Might not happen because of noisy estimates of (p,q)

What might go wrong?

Integrability. If f(x,y) is the height function, we expect that

$$\frac{\partial}{\partial y}\frac{\partial f}{\partial x} = \frac{\partial}{\partial x}\frac{\partial f}{\partial y}$$

In terms of estimated gradient space (p,q), this means:

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

But since p and q were estimated indpendently at each point as intersection of curves on three reflectance maps, equality is not going to exactly hold

Horn's Method

["Robot Vision, B.K.P. Horn, 1986]

- Formulate estimation of surface height z(x,y) from gradient field by minimizing cost functional:

$$\iint_{\text{Image}} (z_x - p)^2 + (z_y - q)^2 \, dxdy$$

where (p,q) are estimated components of the gradient while z_x and z_y are partial derivatives of best fit surface

- Solved using calculus of variations – iterative updating
- z(x,y) can be discrete or represented in terms of basis functions.
- Integrability is naturally satisfied.

2. Photometeric Stereo: Lambertian Surface, Known Lighting

Lambertian Surface

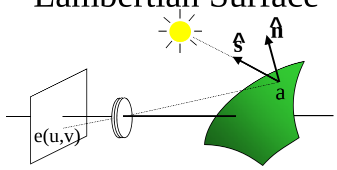

At image location (u,v), the intensity of a pixel x(u,v) is:

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

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.
- $\hat{s}$ 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.

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

- Normal $\hat{n} = \mathbf{b}/|\mathbf{b}|$ and albedo a = |**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

e = Sb

where

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

Solving for b gives

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

Let the residual be

r=e-Sb

Squaring this:

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

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

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

Input Images

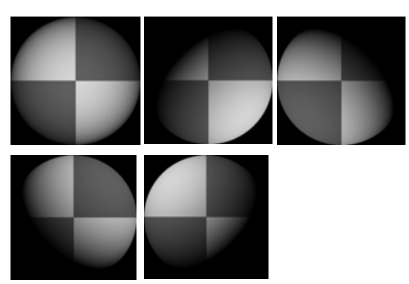

Recovered albedo

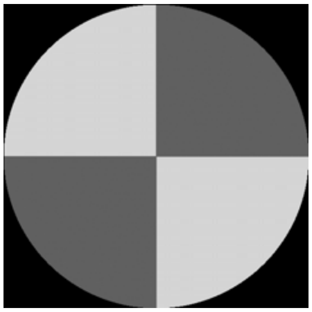

Recovered normal field

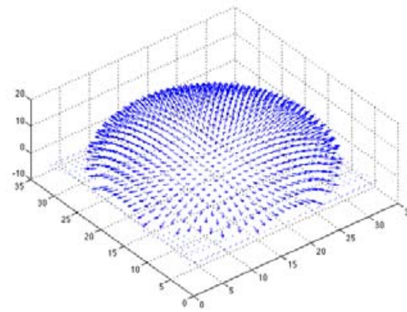

Surface recovered by integration

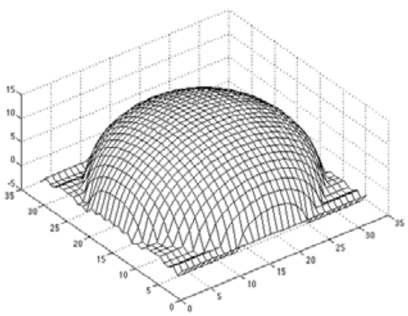

An example of photometric stereo

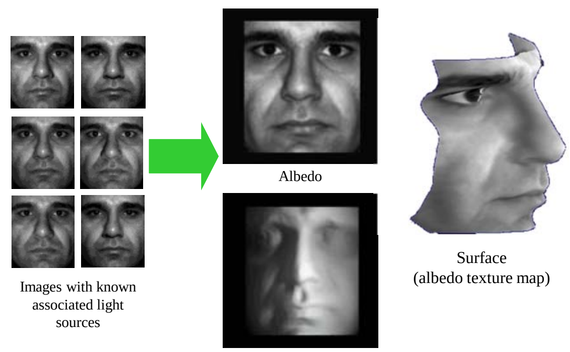

Next Lecture

- Binary image processing
- Reading:
  - Szeliski, Chapter 3 Image processing, Section 3.3 More neighborhood operators