Photometric Stereo
Recap

Computer Vision I
CSE252A
Lecture 8a

Announcements

- HW1 due today

Coordinate system

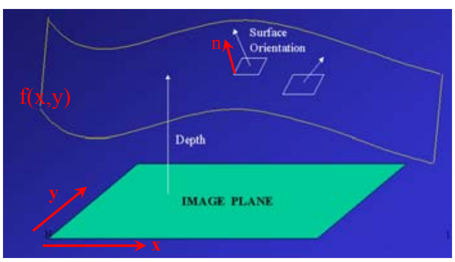

Gradient Space: (p,q)

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

Normal vector

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

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

Reflectance map

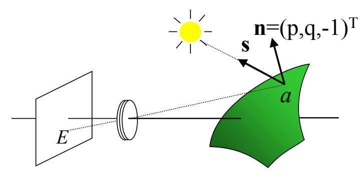

For known BRDF, fix light source direction/strength and projection direction

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

Reflectance Map of Lambertian Surface

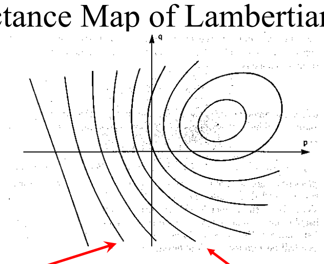

What does the value of one pixel in one image tell us?
It constrains normal to a curve

Two Light Sources
Two reflectance maps

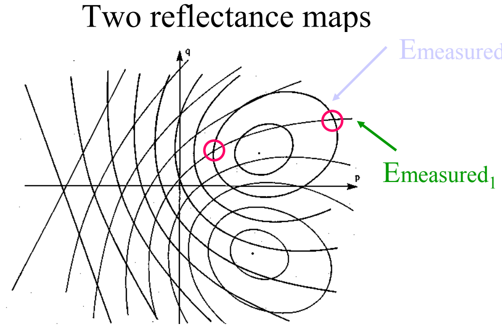

Third image would disambiguate match

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 first row

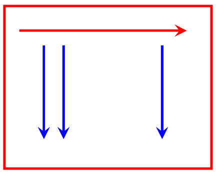

f(x,0)

What might go wrong?

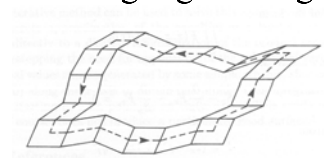

- Height z(x,y) is obtained by integration along a curve from (x_0, y_0).

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

- 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)

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:

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

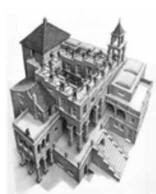

But since p and q were estimated independently at each point as intersection of curves on three reflectance maps, equality is not going to exactly hold

II. Photometeric Stereo: Lambertian Surface, Known Lighting

Lambertian Surface

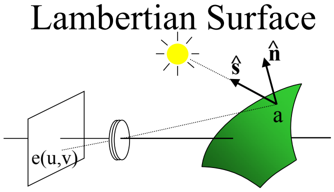

At image location (u,v), the intensity of a pixel x(u,v) is:

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

where

- a(u,v) is the albedo of the surface projecting to (u,v).
- **n**(u,v) is the direction of the surface normal.
- s_0 is the light source intensity.
- **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 is: $\mathbf{n} = \mathbf{b}/|\mathbf{b}|$, albedo is: $|\mathbf{b}|$

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

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

Rewrite as

$$\mathbf{e = Sb}$$

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:

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

$\frac{\partial \mathbf{r}^2}{\partial \mathbf{b}} = 0$ - zero derivative is a necessary condition for a minimum, or

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

Solving for **b** gives

$$\mathbf{b= (S^TS)^{-1}S^Te}$$

Input Images

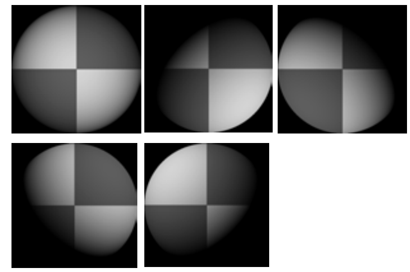

Recovered albedo

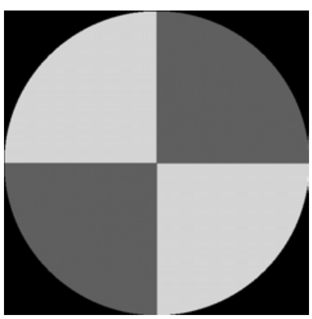

Recovered normal field

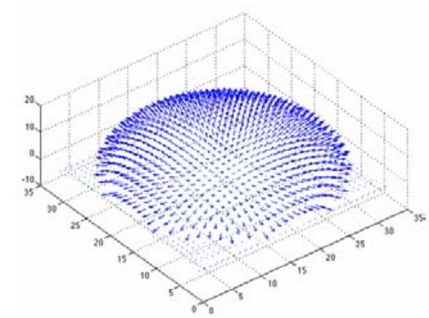

Surface recovered by integration