

Photometric Stereo

Introduction to Computer Vision
CSE152
Lecture 16

CS152, Spring 2012

Intro Computer Vision


Announcements

- HW 3 due on Tuesday

CS152, Spring 2012

Intro Computer Vision

Shading reveals 3-D surface geometry



CS152, Spring 2012

Intro Computer Vision

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.
 1. Arbitrary known BRDF
 2. Lambertian BRDF, known lighting
 3. Lambertian BRDF, unknown lighting.

CS152, Spring 2012

Intro Computer Vision

Photometric Stereo Rigs: One viewpoint, changing lighting



CS152, Spring 2012







Intro Computer Vision

Multi-view stereo vs. Photometric Stereo: Assumptions


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

CS152, Spring 2012

Intro Computer Vision



Gradient Space (p,q)



Gradient Space : (p,q)

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

Normal vector


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

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

CS152, Spring 2012

Intro Computer Vision

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

CS152, Spring 2012

Intro Computer Vision

Reflectance Map




Let the BRDF be the same at all points on the surface, and let the light direction $\mathbf{s}$ be constant, and viewing direction constant.

1. Then image irradiance is a function of only the direction of the surface normal.
2. In gradient space, we can write $E(p,q)$
3. We can measure $E(p,q)$ by taking an image of a sphere made of a single material under distant lighting

CS152, Spring 2012

Intro Computer Vision

Example Reflectance Map: Lambertian surface




For lighting from front

CS152, Spring 2012

Intro Computer Vision

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.

10


Reflectance Map of Lambertian + Specular Surface



CS152, Spring 2012

Intro Computer Vision

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

CS152, Spring 2012

Intro Computer Vision

Two Light Sources Two reflectance maps



A third image would disambiguate match

CS152, Spring 2012

Intro Computer Vision

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

CS152, Spring 2012

Intro Computer Vision


An Example Normal Field



CS152, Spring 2012

Intro Computer Vision


Plastic Baby Doll: Normal Field



CS152, Spring 2012

Intro Computer Vision

Next step: Go from normal field to surface




CS152, Spring 2012

Intro Computer Vision

Recovering the surface $f(x,y)$

Many methods: Simplest approach


1. From normal field $\mathbf{n} = (n_x, n_y, n_z)$, $p = -n_x/n_z$, $q = -n_y/n_z$
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



CS152, Spring 2012

Intro Computer Vision

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, one expects to get back to the starting value.
- Might not happen because of noisy estimates of (p, q)

CS152, Spring 2012

Intro Computer Vision

What might go wrong?

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

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

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




CS152, Spring 2012

Intro Computer Vision

II. Photometric 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) .
- $\mathbf{n}(u, v)$ is the direction of the surface normal.
- s_0 is the light source intensity.
- $\mathbf{s}$ is the direction to the light source.

CS152, Spring 2012

Intro Computer Vision

Lambertian Photometric stereo

- If the light sources $\mathbf{s}_1$, $\mathbf{s}_2$, and $\mathbf{s}_3$ are **known**, then we **can** recover $\mathbf{b}$ 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 for each pixel location and we know $\mathbf{s}_1$, $\mathbf{s}_2$, and $\mathbf{s}_3$. We can then solve for $\mathbf{b}$ 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: $\hat{\mathbf{n}} = \mathbf{b}/|\mathbf{b}|$, albedo is: $|\mathbf{b}|$

CS152, Spring 2012

Intro Computer Vision

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

$$[e_1 \ e_2 \ e_3 \dots e_n] = \mathbf{b}^T [s_1 \ s_2 \ s_3 \dots s_n]$$

Rewrite as

$$\mathbf{e} = \mathbf{S}\mathbf{b}$$

where

$\mathbf{e}$ is n by 1

$\mathbf{b}$ is 3 by 1

$\mathbf{S}$ is n by 3

Let the residual be

$$\mathbf{r} = \mathbf{e} - \mathbf{S}\mathbf{b}$$

Squaring this:

$$\begin{aligned} r^2 &= \mathbf{r}^T \mathbf{r} = (\mathbf{e} - \mathbf{S}\mathbf{b})^T (\mathbf{e} - \mathbf{S}\mathbf{b}) \\ &= \mathbf{e}^T \mathbf{e} - 2\mathbf{b}^T \mathbf{S}^T \mathbf{e} + \mathbf{b}^T \mathbf{S}^T \mathbf{S} \mathbf{b} \end{aligned}$$

$\nabla_{\mathbf{b}}(r^2) = \mathbf{0}$ - zero derivative is a necessary condition for a minimum, or
 $-2\mathbf{S}^T \mathbf{e} + 2\mathbf{S}^T \mathbf{S} \mathbf{b} = \mathbf{0}$;


Solving for $\mathbf{b}$ gives

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

CS152, Spring 2012

Intro Computer Vision


Input Images



CS152, Spring 2012

Intro Computer Vision


Recovered albedo



CS152, Spring 2012

Intro Computer Vision


Recovered normal field



CS152, Spring 2012

Intro Computer Vision


Surface recovered by integration



CS152, Spring 2012

Intro Computer Vision

Lambertian Photometric Stereo



CS152, Spring 2012

Intro Computer Vision

Reconstruction with albedo map



CS152, Spring 2012

Intro Computer Vision


Without the albedo map



CS152, Spring 2012

Intro Computer Vision

Another person



CS152, Spring 2012

Intro Computer Vision

No Albedo map



CS152, Spring 2012

Intro Computer Vision

III. Photometric Stereo with unknown lighting and Lambertian surfaces

CS152, Spring 2012

Intro Computer Vision

How do you construct subspace?

$$[e_1 \ e_2 \ e_3] = B^T [s_1 \ s_2 \ s_3]$$

$[e_1 \ e_2 \ e_3]$ B

- Given three or more images $e_1 \dots e_n$, estimate B and s_i .
- How? Given images in form of $E = [e_1 \ e_2 \ \dots]$, Compute SVD(E) and B^* is n by 3 matrix formed by first 3 singular values.

CS152, Spring 2012

Intro Computer Vision

GBR Transformation

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

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

B

$f(x,y)$

G^T

B̄

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

CS152, Spring 2012 Intro Computer Vision

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

CS152, Spring 2012 Intro Computer Vision

Uncalibrated photometric stereo

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

Comments:

- $\bar{f}^*(x,y)$ differs from $f(x,y)$ by a GBR.
- Can use specularities to resolve GBR for non-Lambertian surface.

CS152, Spring 2012 Intro Computer Vision