


CS4670/5760: Computer Vision

Scott Wehrwein

Lecture 20: Photometric Stereo



Announcements

- P4 out Soon(TM)

Last Time: Two-View Stereo

Reading

Szeliski 12.1.1

Goals

- Understand the basics of modeling image formation:
 - Light sources (point, directional, area)
 - Materials (diffuse/Lambertian, specular, directional diffuse)
- Understand the intuition behind the cosine proportionality of image intensity to $\text{dot}(\mathbf{N}, \mathbf{L})$
- Be prepared to implement photometric stereo to solve for surface normals in a sequence of grayscale or color images with known light directions.
- Understand how to use matrix factorization (using the SVD) to solve for surface normals and light directions jointly.

Last Time: Two-View Stereo



Key Idea: use feature motion to understand shape

Today: Photometric Stereo



Key Idea: use pixel brightness to understand shape

Today: Photometric Stereo




Key Idea: use pixel brightness to understand shape

Photometric Stereo

What results can you get?



Input
(1 of 12)



Normals (RGB
colormap)



Normals (vectors)




Shaded 3D
rendering



Textured 3D
rendering


Photometric Stereo

What results can you get?



Radiometry

- What determines brightness of an image pixel?





Light Sources

Basic types of light sources:


- point source
- directional source
(a point source that's infinitely far away)
- area source
(a union of point sources)

What happens when light hits an object?


Materials - Three Forms




Ideal diffuse
(Lambertian)



Ideal
specular






Directional
diffuse



Ideal diffuse
(Lambertian)


Ideal
specular

Directional
diffuse




Ideal Diffuse Reflection

- Characteristic of multiple scattering materials
- An idealization but reasonable for matte surfaces





Directional Lighting



- Key property: all rays are parallel
- Equivalent to an infinitely distant point source

Lambertian Reflectance under directional lighting





$$I = N \cdot L$$

Image
intensity $=$ Surface
normal \cdot Light
direction

Image
intensity \propto $\cos(\text{angle between } N \text{ and } L)$


Lambertian Reflectance



1. **Incoming:** Reflected energy is proportional to cosine of the angle between L and N
2. **Outgoing:** Measured intensity is viewpoint-independent


Lambertian Reflectance: Incoming

1. **Incoming:** Reflected energy is proportional to cosine of angle between L and N




Lambertian Reflectance: Incoming

1. **Incoming:** Reflected energy is proportional to cosine of angle between L and N



Lambertian Reflectance: Incoming


1. **Incoming:** Reflected energy is proportional to cosine of angle between L and N



Light (energy) hitting surface is proportional to the cosine

Lambertian Reflectance: Outgoing



2. **Outgoing:** Measured intensity is viewpoint-independent



Lambertian Reflectance: Outgoing



2. Outgoing: Measured intensity is viewpoint-independent

LAMBERTIAN
DISTRIBUTION





Lambertian Reflectance: Outgoing

2. **Outgoing:** Measured intensity is viewpoint-independent



Lambertian Reflectance: Outgoing


2. **Outgoing:** Measured intensity is viewpoint-independent



Radiance $\propto B_0 \cos(\theta) \frac{1}{\cos(\theta)}$
(what the camera measures)

$$A \propto \frac{1}{\cos(\theta)}.$$


Image Formation Model: Final



$$I = k_d \mathbf{N} \cdot \mathbf{L}$$

1. **Diffuse albedo:** what fraction of incoming light is reflected?
 - Introduce scale factor k_d
2. **Light intensity:** how much light is arriving?
 - Compensate with camera exposure (global scale factor)
3. **Camera response function**
 - Assume pixel value is linearly proportional to incoming energy (perform radiometric calibration if not)

A Single Image: Shape from Shading



$$I = k_d \mathbf{N} \cdot \mathbf{L}$$


Assume k_d is 1 for now.

What can we measure from one image?

- $\cos^{-1}(I)$ is the angle between \mathbf{N} and \mathbf{L}
- Add assumptions:
 - A few known normals (e.g. silhouettes)
 - Smoothness of normals

In practice, SFS doesn't work very well: assumptions are too restrictive, too much ambiguity in nontrivial scenes.

Multiple Images: Photometric Stereo



$$I_1 = k_d \mathbf{N} \cdot \mathbf{L}_1$$

$$I_2 = k_d \mathbf{N} \cdot \mathbf{L}_2$$

$$I_3 = k_d \mathbf{N} \cdot \mathbf{L}_3$$

Write this as a matrix equation:

$$\begin{bmatrix} I_1 & I_2 & I_3 \end{bmatrix} = k_d \mathbf{N}^T \begin{bmatrix} \mathbf{L}_1 & \mathbf{L}_2 & \mathbf{L}_3 \end{bmatrix}$$

Solving the Equations

$$\underbrace{\begin{bmatrix} I_1 & I_2 & I_3 \end{bmatrix}}_{\substack{\mathbf{I} \\ 1 \times 3}} = k_d \underbrace{\mathbf{N}^T}_{\substack{\mathbf{G} \\ 1 \times 3}} \underbrace{\begin{bmatrix} \mathbf{L}_1 & \mathbf{L}_2 & \mathbf{L}_3 \end{bmatrix}}_{\substack{\mathcal{L} \\ 3 \times 3}}$$

$$\mathbf{G} = \mathbf{I}\mathbf{L}^{-1}$$

$$k_d = \|\mathbf{G}\|$$

$$\mathbf{N} = \frac{1}{k_d} \mathbf{G}$$

Solving the Equations

$$\underbrace{\begin{bmatrix} I_1 & I_2 & I_3 \end{bmatrix}}_{\substack{\mathbf{I} \\ 1 \times 3}} = k_d \underbrace{\mathbf{N}^T}_{\substack{\mathbf{G} \\ 1 \times 3}} \underbrace{\begin{bmatrix} \mathbf{L}_1 & \mathbf{L}_2 & \mathbf{L}_3 \end{bmatrix}}_{\substack{\mathcal{L} \\ 3 \times 3}}$$

$$\mathbf{G} = \mathbf{I}\mathbf{L}^{-1}$$

- When is \mathbf{L} nonsingular (invertible)?
 - All 3 light directions are linearly independent, or:
 - Light direction vectors are not coplanar.
- What if we have more than one pixel?
 - Stack them all into one big system.

More than Three Lights

$$\begin{bmatrix} I_1 & \dots & I_n \end{bmatrix} = k_d \mathbf{N}^T \begin{bmatrix} \mathbf{L}_1 & \dots & \mathbf{L}_n \end{bmatrix}$$

- Solve using least squares:

$$I = k_d N^T L$$

$$L^T k_d N = I^T$$

$$L^T (k_d N) = I^T$$

$$\mathbf{A} \quad \mathbf{x} \quad = \quad \mathbf{b}$$

$$k_d = ||x||$$

$$N = \frac{x}{||x||}$$

- Given G, solve for N and k_d as before.

More than one pixel

Previously:

$$\begin{array}{c} 1 \times n \\ \boxed{\mathbf{I}} \end{array} = \begin{array}{c} 1 \times 3 \\ \boxed{\mathbf{N}} \end{array} * \begin{array}{c} 3 \times n \\ \boxed{\mathbf{L}} \end{array} \quad n = \# \text{ images}$$

More than one pixel

Stack all pixels into one system:

$$\begin{array}{c} p \times n \\ \boxed{\text{I}} \end{array} = \begin{array}{c} p \times 3 \\ \boxed{\text{N}} \end{array} * \begin{array}{c} 3 \times n \\ \boxed{\text{L}} \end{array}$$

$p = \# \text{ pixels}$
 $n = \# \text{ images}$

Solve as before.

Color Images

- Now we have 3 equations for a pixel:

$$\mathbf{I}_R = k_{dR} \mathbf{L} \mathbf{N}$$

$$\mathbf{I}_G = k_{dG} \mathbf{L} \mathbf{N}$$


$$\mathbf{I}_B = k_{dB} \mathbf{L} \mathbf{N}$$

- Simple approach: solve for \mathbf{N} using grayscale or a single channel.
- Then fix \mathbf{N} and solve for each channel's k_d :

$$k_d = \frac{\sum_i I_i L_i N^T}{\sum_i (L_i N^T)^2}$$

Determining Light Directions


- Trick: Place a mirror ball in the scene.



- The location of the highlight is determined by the light source direction.

Determining Light Directions


- For a perfect mirror, the light is reflected across N :



$$I_e = \begin{cases} I_i & \text{if } \mathbf{V} = \mathbf{R} \\ 0 & \text{otherwise} \end{cases}$$


- So the light source direction is given by:

$$\mathbf{L} = 2(\mathbf{N} \cdot \mathbf{R})\mathbf{N} - \mathbf{R}$$



Determining Light Directions

- For a sphere with highlight at point H:



Compute N :


$$N_x = \frac{x_h - x_c}{r}$$

$$N_y = \frac{y_h - y_c}{r}$$

$$N_z = \sqrt{1 - x^2 - y^2}$$

- R = direction of the camera from $C = [0 \ 0 \ 1]^T$
- $L = 2(N \cdot R)N - R$

Results




from Athos Georghiades

Results



Input
(1 of 12)



Normals (RGB
colormap)



Normals (vectors)




Shaded 3D
rendering



Textured 3D
rendering

For (unfair) Comparison

- Multi-view stereo results on a similar object
- 47+ hrs compute time



Modern MVS
result



Ground truth

Taking Stock: Assumptions

Lighting	Materials	Geometry	Camera
directional	diffuse	convex / no shadows	linear
known direction	no inter- reflections		orthographic
> 2 nonplanar directions	no subsurface scattering		

Questions?

Unknown Lighting

- What we've seen so far: [Woodham 1980]
- Next up: Unknown light directions [Hayakawa 1994]

Unknown Lighting

Surface normals

Light directions

$$I = k N \cdot \ell L$$

Diffuse albedo

Light intensity

The diagram illustrates the components of the lighting equation $I = k N \cdot \ell L$. Arrows point from the labels to the corresponding variables: 'Surface normals' points to N , 'Light directions' points to ℓ , 'Diffuse albedo' points to k , and 'Light intensity' points to L .

Unknown Lighting

Surface normals,
scaled by albedo

Light directions,
scaled by intensity


$$I = N \cdot L$$

Unknown Lighting

Same as before, just transposed:

$p = \# \text{ pixels}$

$n = \# \text{ images}$


$$\begin{array}{|c|} \hline M \\ \hline \end{array} = \begin{array}{|c|} \hline L \\ \hline \end{array} * \begin{array}{|c|} \hline N \\ \hline \end{array}$$

Unknown Lighting

Measurements
(one image per row)


Light directions
(scaled by intensity)

Surface normals
(scaled by albedo)




M

=



L

*




N

Both L and N are now unknown!
This is a matrix factorization problem.

Unknown Lighting

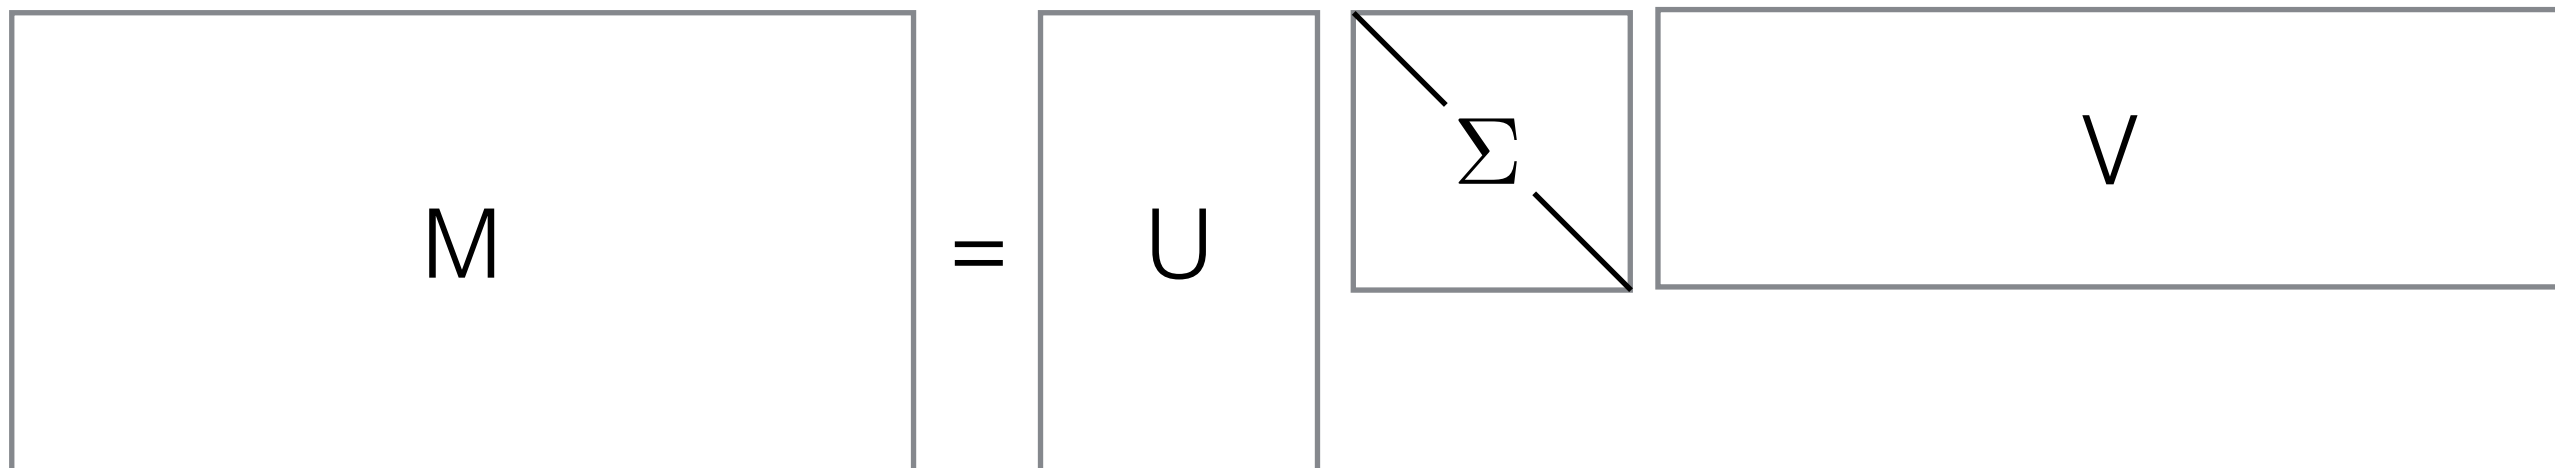
$$M_{ij} = L_i \cdot N_j$$



There's hope: We know that M is rank 3

Unknown Lighting

Use the SVD to decompose M:

$$M = U \Sigma V$$


SVD gives the best rank-3 approximation of a matrix.

Unknown Lighting

Use the SVD to decompose M:

$$M = U \Sigma V$$

What do we do with Σ ?

Unknown Lighting

Use the SVD to decompose M:

$$\boxed{M} = \boxed{U\sqrt{\Sigma}} \boxed{\sqrt{\Sigma}V}$$

What do we do with Σ ?

Unknown Lighting

Use the SVD to decompose M:

$$M = U\sqrt{\Sigma}V$$

Can we just do that?

Unknown Lighting

Use the SVD to decompose M:

$$M = U \sqrt{\Sigma} A^{-1} A \sqrt{\Sigma} V$$

Can we just do that? ...almost.

The decomposition is non-unique up to an invertible 3x3 A.

Unknown Lighting

Use the SVD to decompose M:

$$M = U\sqrt{\Sigma} \begin{matrix} A \\ A^{-1} \\ \sqrt{\Sigma}V \end{matrix}$$

$$L = U\sqrt{\Sigma}A$$

$$S = A^{-1}\sqrt{\Sigma}V$$

Unknown Lighting

Use the SVD to decompose M:

$$M = U\sqrt{\Sigma} \begin{matrix} A \\ A^{-1} \end{matrix} \sqrt{\Sigma} V$$


$$L = U\sqrt{\Sigma}A \quad S = A^{-1}\sqrt{\Sigma}V$$

You can find A if you know

- 6 points with the same reflectance, or
- 6 lights with the same intensity.

Unknown Lighting: Ambiguities

- Multiple combinations of lighting and geometry can produce the same sets of images.
- Add assumptions or prior knowledge about geometry or lighting, etc. to limit the ambiguity.




[Belhumeur et al.'97]

Questions?

Since 1994...


- Workarounds for many of the restrictive assumptions.
- Webcam photometric stereo:



Ackermann et al. 2012

Since 1994...


- Photometric stereo from unstructured photo collections (different cameras and viewpoints):




Shi et al, 2014

Since 1994...

- Non-Lambertian (shiny) materials:



Hertzmann and Seitz, 2005








Lights, camera, action

Sensor



Lights



Camera

