



Photometric Stereo, Shape from Shading SfS F&P Ch 5 (old) Ch 2 (new)

Guido Gerig

CS 6643, Spring 2016

Credits: M. Pollefev UNC CS256, Ohad Ben-Shahar CS BGU, Wolff JUN
(<http://www.cs.jhu.edu/~wolff/course600.461/week9.3/index.htm>)



Photometric Stereo




Depth from Shading?

First step: Surface
Normals from Shading


Second step:
Re-integration of
surface from Normals



Examples




<http://www.youtube.com/watch?v=sfCQ7f7PMbc&feature=related>



Simulated voyage over the surface of Neptune's large moon Triton

<http://www.youtube.com/watch?v=nwzVrC2GQXE>



<http://www.youtube.com/watch?v=KiTA6ftyQuY>



Shape from Shading

Inverting the image formation process




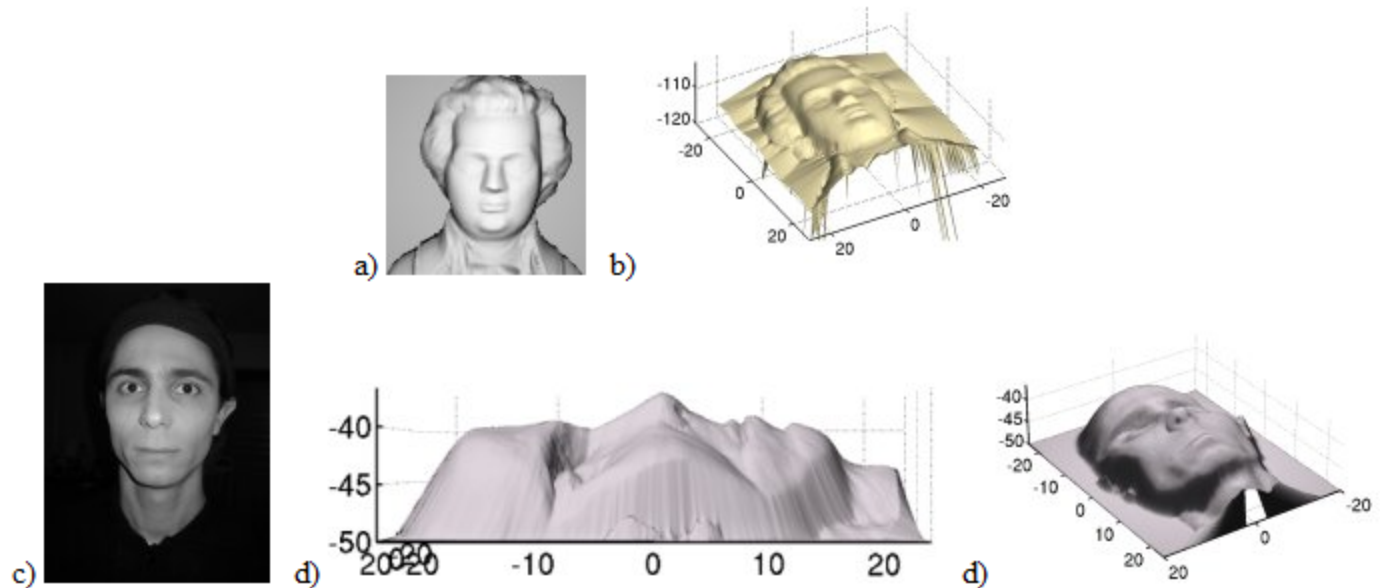
Image formation = “Shading from shape” (and light sources)



Shape from Shading

Authors: [Emmanuel Prados](#) and [Olivier Faugeras](#)

[CVPR'2005](#), International Conference on Computer Vision and Pattern Recognition, San Diego, CA, USA, June 2005.



a) Synthetic image generated from the classical Mozart's face [Zhang-Tsai-etal:99]; b) reconstructed surface from a) by new algorithm;
c) real image of a face; d)-e) reconstructed surface from c) by new algorithm.




Photometric Stereo

- Assume:
 - a local shading model
 - a set of point sources that are infinitely distant
 - a set of pictures of an object, obtained in exactly the same camera/object configuration but using different sources
 - A Lambertian object (or the specular component has been identified and removed)

Setting for Photometric Stereo




Multiple images with different lighting (vs binocular/geometric stereo)



Goal: 3D from One View and multiple Source positions



Input images





Usable Data
Mask





Scene Results

Needle Diagram:
Surface Normals

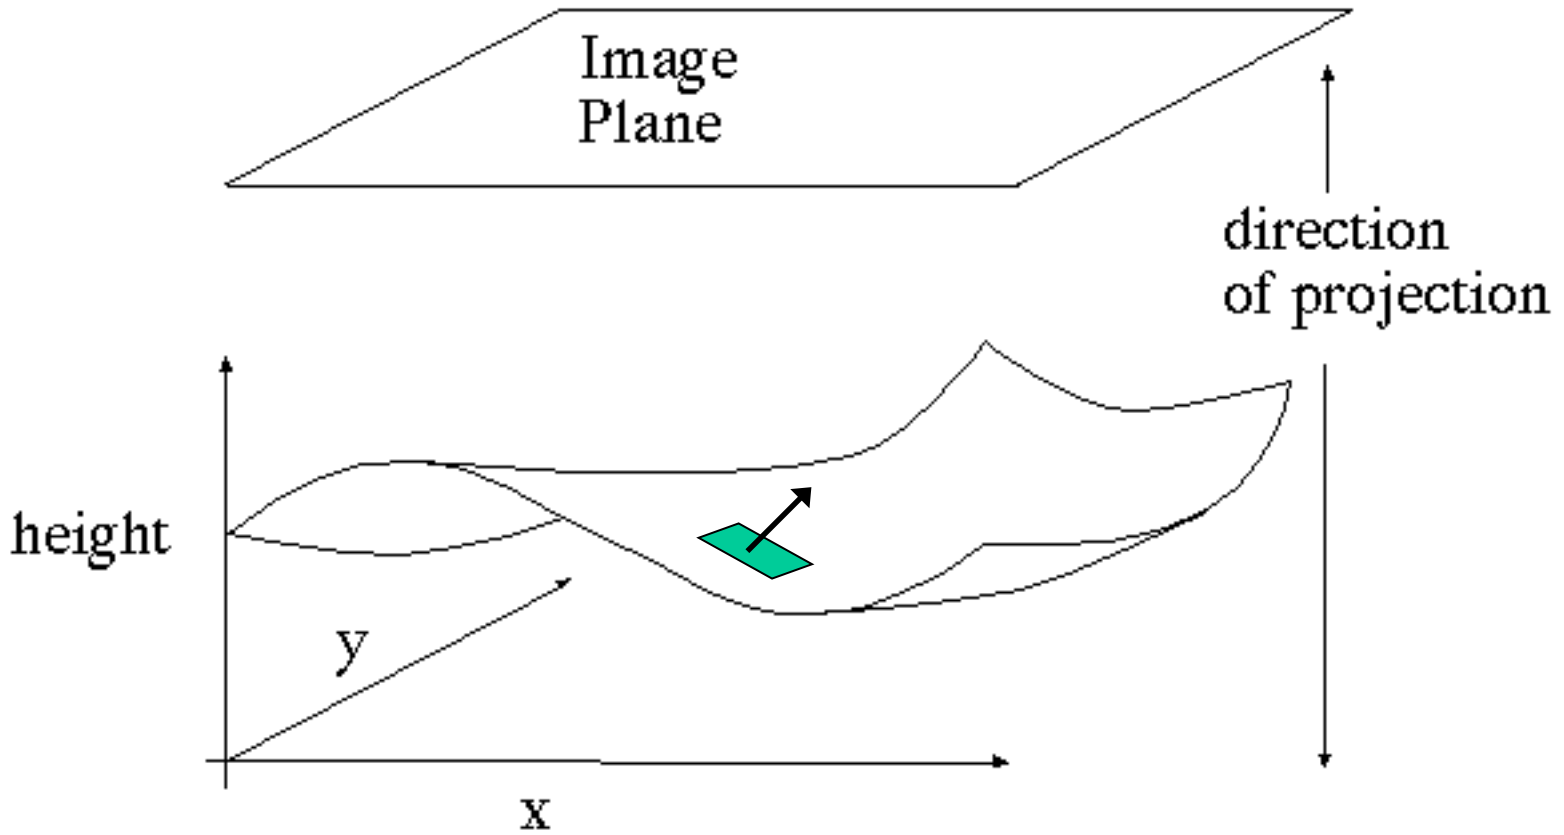


Albedo

Re-lit:



Projection model for surface recovery - usually called a Monge patch

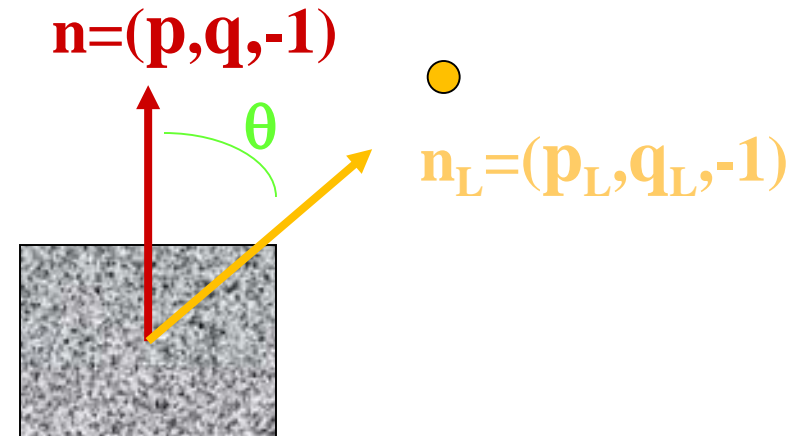





Lambertian Reflectance Map

LAMBERTIAN MODEL


$$E = \rho \langle \mathbf{n}, \mathbf{n}_L \rangle = \rho \cos \theta$$



$$\cos \theta = \frac{1 + pp_L + qq_L}{\sqrt{1 + p^2 + q^2} \sqrt{1 + p_L^2 + q_L^2}}$$

REFLECTANCE MAP IS A VIEWER-CENTERED REPRESENTATION OF REFLECTANCE

$$(f_x, f_y, -1) = (0, 1, f_x) \times (1, 0, f_y)$$



REFLECTANCE MAP IS A VIEWER-CENTERED REPRESENTATION OF REFLECTANCE

$$(-f_x, -f_y, 1) = (-p, -q, 1)$$

p, q comprise a **gradient** or **gradient space** representation for local surface orientation.


Reflectance map expresses the reflectance of a material directly in terms of viewer-centered representation of local surface orientation.



Reflectance Map

Shading on Lambertian surface – Overhead point source



$$I(x, y) = \rho(\hat{N} \cdot [0, 0, 1]) = \rho \frac{1}{\sqrt{p^2 + q^2 + 1}} = R(p, q)$$



Reflectance Map ($p_s=0, q_s=0$)

The Reflectance Map – Lambertian surface from overhead source position

$$R(p, q) = \frac{1}{\sqrt{p^2 + q^2 + 1}}$$



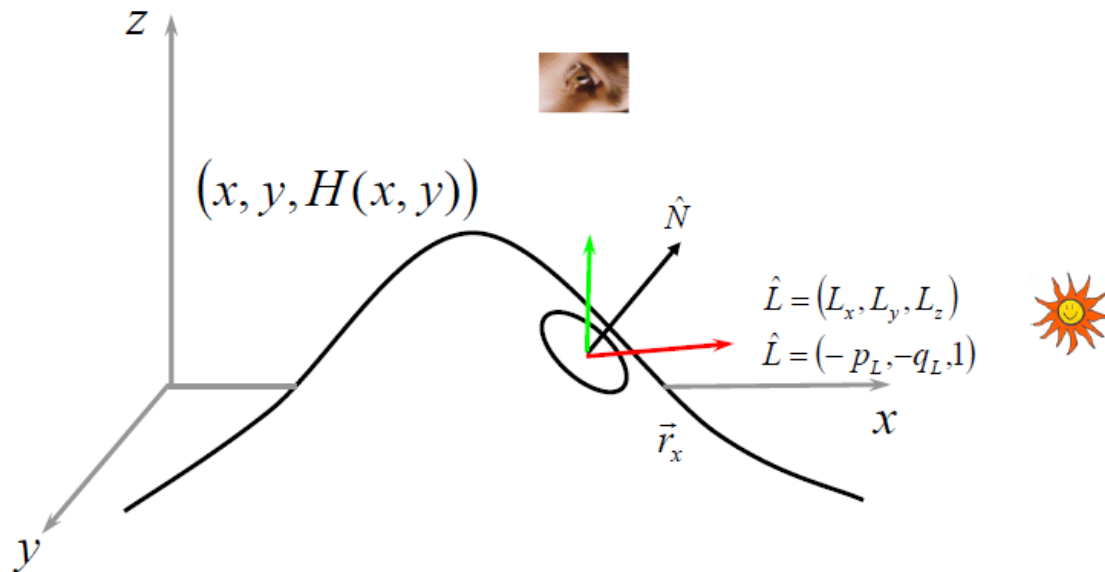


Reflectance Map

Shape from Shading

Shading on Lambertian surface – General point source

$$I = \rho(\hat{N} \cdot \hat{L}) = \rho \frac{-p \cdot L_x - q \cdot L_y + L_z}{\sqrt{p^2 + q^2 + 1} \sqrt{L_x^2 + L_y^2 + L_z^2}} = \rho \frac{p \cdot p_L + q \cdot q_L + 1}{\sqrt{p^2 + q^2 + 1} \sqrt{p_L^2 + q_L^2 + 1}}$$








Reflectance Map

The Reflectance Map – Lambertian surface from general source position

$$R(p, q) = \frac{p \cdot p_L + q \cdot q_L + 1}{\sqrt{p^2 + q^2 + 1} \sqrt{p_L^2 + q_L^2 + 1}}$$



Gradient point of maximum brightness



Reflectance Map (General)




Figure 10-13. The reflectance map is a plot of brightness as a function of surface orientation. Here it is shown as a contour map in gradient space. In the case of a Lambertian surface under point-source illumination, the contours turn out to be nested conic sections. The maximum of $R(p, q)$ occurs at the point $(p, q) = (p_s, q_s)$, found inside the nested conic sections, while $R(p, q) = 0$ all along the line on the left side of the contour map.



Reflectance Map



Given Intensity I in image, there are multiple (p, q) combinations (= surface orientations).

⇒ Use multiple images with different light source directions.

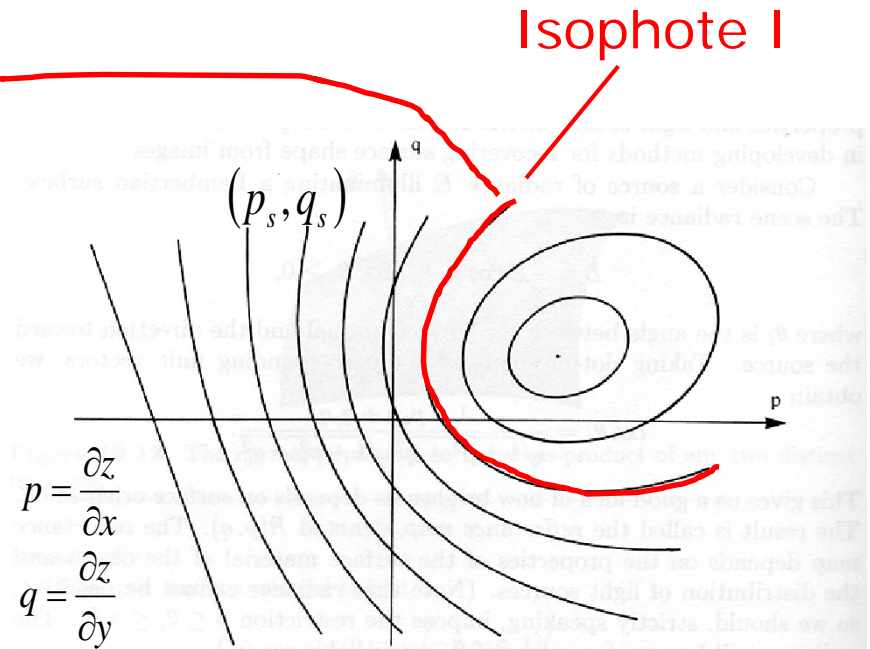


Figure 10-13. The reflectance map is a plot of brightness as a function of surface orientation. Here it is shown as a contour map in gradient space. In the case of a Lambertian surface under point-source illumination, the contours turn out to be nested conic sections. The maximum of $R(p, q)$ occurs at the point $(p, q) = (p_s, q_s)$, found inside the nested conic sections, while $R(p, q) = 0$ all along the line on the left side of the contour map.



Multiple Images = Multiple Maps

Can isolate p, q as contour intersection






Figure 10-21. In the case of a Lambertian surface illuminated successively by two different point sources, there are at most two surface orientations that produce a particular pair of brightness values. These are found at the intersection of the corresponding contours in two superimposed reflectance maps.



Example: Two Views

$$I_1(x, y) = R_1(p, q)$$

$$I_2(x, y) = R_2(p, q)$$




Photometric Stereo



$$p_s = q_s = 0$$


$$p_s = 0.5 \quad q_s = -0.0$$



Still not unique for certain intensity pairs.

Constant Albedo

Photometric Stereo



$$I_1 = \rho \mathbf{S}_1 \cdot \mathbf{N}$$

$$I_2 = \rho \mathbf{S}_2 \cdot \mathbf{N}$$

$$I_3 = \rho \mathbf{S}_3 \cdot \mathbf{N}$$

$$\begin{bmatrix} I_1 \\ I_2 \\ I_3 \end{bmatrix} = \begin{bmatrix} \mathbf{S}_1^T \\ \mathbf{S}_2^T \\ \mathbf{S}_3^T \end{bmatrix} \rho \mathbf{N}$$

$$\rho \mathbf{N} = \mathbf{S}^{-1} \mathbf{I}$$

Solve linear equation system to calculate $\bar{\mathbf{N}}$.



Varying Albedo

Solution Forsyth & Ponce:

For each point source, we know the source vector (by assumption). We assume we know the scaling constant of the linear camera (k). Fold the normal (\mathbf{N}) and the reflectance ($\rho(x, y)$) into one vector \mathbf{g} , and the scaling constant and source vector into another \mathbf{V}_j .

- Out of shadow:

$$\begin{aligned} I(x, y) &= kB(\mathbf{x}) \\ &= kB(x, y) \\ &= \underbrace{k\rho(x, y)\mathbf{N}(x, y)}_{\mathbf{g}(x, y)} \cdot \underbrace{\mathbf{S}_1}_{\mathbf{V}_1} \\ &= \mathbf{g}(x, y) \cdot \mathbf{V}_1 \end{aligned}$$

- In shadow:

$$I(x, y) = 0$$

where $\mathbf{g}(x, y) = \rho(x, y)\mathbf{N}(x, y)$ and $\mathbf{V}_1 = k\mathbf{S}_1$, where k is the constant connecting the camera response to the input radiance.



Multiple Images: Linear Least Squares Approach

- Combine albedo and normal
- Separate lighting parameters
- More than 3 images => overdetermined system

$$\mathcal{V} = \begin{pmatrix} \mathbf{V}_1^T \\ \mathbf{V}_2^T \\ \vdots \\ \mathbf{V}_n^T \end{pmatrix} \quad \mathbf{i}(x, y) = \{I_1(x, y), I_2(x, y), \dots, I_n(x, y)\}^T$$

$$\mathbf{i}(x, y) = \mathcal{V}\mathbf{g}(x, y)$$

\mathbf{g} is obtained by solving this linear system: $\bar{\mathbf{g}}(x, y) = \mathbf{V}^{-1}\mathbf{i}(x, y)$


- How to calculate albedo ρ and \bar{N} ?

$$\bar{\mathbf{g}}(x, y) = \rho(x, y)\bar{N}(x, y)$$

$$\rightarrow \bar{N} = \frac{\bar{\mathbf{g}}}{|\bar{\mathbf{g}}|}, \quad \rho(x, y) = |\bar{\mathbf{g}}|$$




Example LLS Input



Problem: Some regions in some images are in the shadow (no image intensity).

Dealing with Shadows (Missing Info)



- For each point source, we know the source vector (by assumption). We assume we know the scaling constant of the linear camera. Fold the normal and the reflectance into one vector \mathbf{g} , and the scaling constant and source vector into another \mathbf{V}_j

- Out of shadow:

$$\begin{aligned} I_j(x, y) &= kB(x, y) \\ &= k\rho(x, y)(\mathbf{N}(x, y) \bullet \mathbf{S}_j) \\ &= \mathbf{g}(x, y) \bullet \mathbf{V}_j \end{aligned}$$

- In shadow:

$$I_j(x, y) = 0$$

No partial shadow

Matrix Trick for Complete Shadows

- Matrix from Image Vector:

$$\mathcal{I}(x, y) = \begin{pmatrix} I_1(x, y) & \dots & 0 & 0 \\ 0 & I_2(x, y) & \dots & 0 \\ \dots & & & \\ 0 & 0 & \dots & I_n(x, y) \end{pmatrix}$$

- Multiply LHS and RHS with diag matrix

$$\mathcal{I}i = \mathcal{I}\mathcal{V}g(x, y)$$

$$\begin{pmatrix} I_1^2(x, y) \\ I_2^2(x, y) \\ \dots \\ I_n^2(x, y) \end{pmatrix} = \begin{pmatrix} I_1(x, y) & 0 & \dots & 0 \\ 0 & I_2(x, y) & \dots & \dots \\ \dots & \dots & \dots & 0 \\ 0 & \dots & 0 & I_n(x, y) \end{pmatrix} \begin{pmatrix} \mathbf{V}_1^T \\ \mathbf{V}_2^T \\ \dots \\ \mathbf{V}_n^T \end{pmatrix} \mathbf{g}(x, y)$$

Known
Known
Known
Unknown

⇒ Relevant elements of the left vector and the matrix are zero at points that are in shadow.






Obtaining Normal and Albedo

- Given sufficient sources, we can solve the previous equation (most likely need a least squares solution) for $\mathbf{g}(x, y)$.
- Recall that $\mathbf{N}(x, y)$ is the unit normal.
- This means that $\rho(x, y)$ is the magnitude of $\mathbf{g}(x, y)$.
- This yields a check
 - If the magnitude of $\mathbf{g}(x, y)$ is greater than 1, there's a problem.
- And $\mathbf{N}(x, y) = \mathbf{g}(x, y) / \rho(x, y)$.






Example LLS Input





Example LLS Result


- Reflectance / albedo:





Recap



- Obtain normal / orientation, no depth





Goal

Shape as surface with depth and normal,
so far only normal



Recovering a surface from normals - 1



- Recall the surface is written as

$$(x, y, f(x, y))$$

- This means the normal has the form:

$$N(x, y) = \left(\frac{1}{\sqrt{f_x^2 + f_y^2 + 1}} \right) \begin{pmatrix} -f_x \\ -f_y \\ 1 \end{pmatrix}$$

- If we write the known vector \mathbf{g} as

$$\mathbf{g}(x, y) = \begin{pmatrix} g_1(x, y) \\ g_2(x, y) \\ g_3(x, y) \end{pmatrix}$$

- Then we obtain values for the partial derivatives of the surface:

$$f_x(x, y) = (g_1(x, y) / g_3(x, y))$$

$$f_y(x, y) = (g_2(x, y) / g_3(x, y))$$

Recovering a surface from normals - 2



- Recall that mixed second partials are equal --- this gives us an integrability **check**. We must have:


$$\frac{\partial(g_1(x, y)/g_3(x, y))}{\partial y} = \frac{\partial(g_2(x, y)/g_3(x, y))}{\partial x}$$

- We can now recover the surface height at any point by integration along some path, e.g.

$$f(x, y) = \int_0^x f_x(s, y) ds + \int_0^y f_y(x, t) dt + c$$



Height Map from Integration



How to integrate?

Possible Solutions

- Engineering approach: Path integration (Forsyth & Ponce)
- In general: Calculus of Variation Approaches
- Horn: Characteristic Strip Method
- Kimmel, Siddiqi, Kimia, Bruckstein: Level set method
- Many others

Shape by Integation (Forsyth&Ponce)

- The partial derivative gives the change in surface height with a small step in either the x or the y direction
- We can get the surface by summing these changes in height along some path.

$$f(x, y) = \oint_C \left(\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y} \right) \cdot d\mathbf{l} + c$$

For example, we can reconstruct the surface at (u, v) by starting at $(0, 0)$, summing the y -derivative along the line $x = 0$ to the point $(0, v)$, and then summing the x -derivative along the line $y = v$ to the point (u, v)

$$f(u, v) = \int_0^v \frac{\partial f}{\partial y}(0, y) dy + \int_0^u \frac{\partial f}{\partial x}(x, v) dx + c$$

Obtain many images in a fixed view under different illuminants

Determine the matrix \mathcal{V} from source and camera information

Create arrays for albedo, normal (3 components),

p (measured value of $\frac{\partial f}{\partial x}$) and

q (measured value of $\frac{\partial f}{\partial y}$)

For each point in the image array

Stack image values into a vector i

Construct the diagonal matrix \mathcal{I}

Solve $\mathcal{I}\mathcal{V}\mathbf{g} = \mathcal{I}i$

to obtain \mathbf{g} for this point

albedo at this point is $|\mathbf{g}|$

normal at this point is $\frac{\mathbf{g}}{|\mathbf{g}|}$

p at this point is $\frac{N_1}{N_3}$

q at this point is $\frac{N_2}{N_3}$

end

Check: is $(\frac{\partial p}{\partial y} - \frac{\partial q}{\partial x})^2$ small everywhere?

top left corner of height map is zero

for each pixel in the left column of height map

height value=previous height value + corresponding q value

end

for each row

for each element of the row except for leftmost

height value = previous height value + corresponding p value

end

end

Simple Algorithm Forsyth & Ponce

Problem: Noise and numerical (in)accuracy are added up and result in distorted surface.

Solution: Choose several different integration paths, and build average height map.

Mathematical Property: Integrability

- Smooth, C2 continuous surface:

$$Z(x, y)_{xy} = Z(x, y)_{yx}$$

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

$$\Rightarrow \text{check if } \left(\frac{\partial p}{\partial y} - \frac{\partial q}{\partial x}\right)^2 \text{ is small}$$



SHAPE FROM SHADING

(Calculus of Variations Approach)

- First Attempt: Minimize error in agreement with Image Irradiance Equation over the region of interest:

$$\iint_{\text{object}} (I(x, y) - R(p, q))^2 dx dy$$

SHAPE FROM SHADING

(Calculus of Variations Approach)

- Better Attempt: Regularize the Minimization of error in agreement with Image Irradiance Equation over the region of interest:

$$\iint_{\text{object}} p_x^2 + p_y^2 + q_x^2 + q_y^2 + \lambda(I(x, y) - R(p, q))^2 dx dy$$

Linear Approaches for SFS

- Linear approaches reduce the non-linear problem into a linear through the linearization of the image irradiance equation .
- The idea is based on the assumption that **the lower order components in the reflectance map dominate**. Therefore, these algorithms only work well under this assumption.





Simple Scenario

- We will be concerned with the simplest scenario, where the following assumptions hold;
 - Camera; orthographic projection.
 - Surface reflectivity; Lambertian surface
 - Known/estimated illumination direction.
 - Known/estimated surface albedo/
 - The optical axis is the Z axis of the camera and the surface can be parameterized as $Z = Z(x, y)$.
- The image irradiance (amount of light received by the camera to which the gray-scale produced is directly proportional) can be defined as follows;

$$E(x, y) = R_{\rho, I}(p, q) = \rho \mathbf{I}^T \mathbf{n} = \frac{\rho}{\sqrt{1 + p^2 + q^2}} \mathbf{I}^T [-p, -q, 1]^T \quad (\text{A})$$

- Eq(A) is the typical starting point of many shape from shading techniques, yet it is of a great mathematical complexity, it is a non-linear partial differential equation in $p = p(x, y)$ and $q = q(x, y)$, which are the gradients of the unknown surface $Z = Z(x, y)$



Pentland's Approach

- Under the assumptions of :
 - Lambertian surface,
 - orthographic projections,
 - the surface being illuminated by distant light sources, and
 - the surface is not self-shadowing,
- Pentland defined the image irradiance equation as follows;




$$E(x, y) = R(p, q) = \frac{\rho(i_x p + i_y q - i_z)}{\sqrt{1 + p^2 + q^2}} = \frac{p \sin \sigma \cos \tau + q \sin \sigma \sin \tau + \cos \sigma}{\sqrt{1 + p^2 + q^2}}$$

Where light source direction is defined as:

$$I = [\sin \sigma \cos \tau, \sin \sigma \sin \tau, \cos \sigma]^T$$



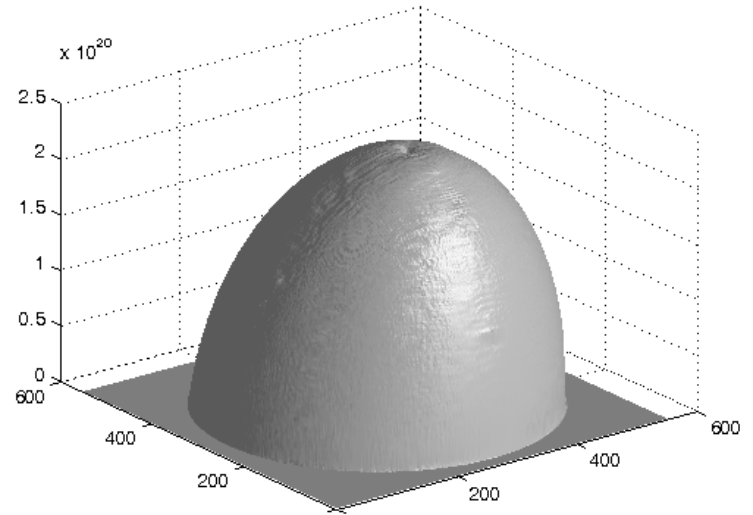



Pentland's Approach



Pentland, A., "Shape Information From Shading: A Theory About Human Perception," *Computer Vision., Second International Conference on* , vol., no., pp.404-413, 5-8 Dec 1988.

Shah's Approach




Deformable Models: SNAKES

Geodesic Snake
formulated as PDE

$$\frac{\partial c}{\partial t} = [\alpha]N$$




Curve evolves
over time



Speed

Normal direction
to curve



Deformable Models: SNAKES


Geodesic Snake:

$$\frac{\partial c}{\partial t} = [\kappa]N$$

Curvature (convex,
concave)

Curve evolves
over time

Normal direction
to curve



Mathematical
solution is circle

Deformable Models: SNAKES


Geodesic Snake:

$$\frac{\partial c}{\partial t} = [\kappa + \alpha]N$$

Curve evolves
over time

Speed

Normal direction
to curve



Another Solution to SFS: Kimmel, Siddiqi, Kimia, Bruckstein


Proposed Solution: Evolve a curve such that it tracks the height contours of $z(x, y)$.

[Kimmel *et al.*, IJCV95]

Height climbed while progressing a distance $|\Delta C|$ in the direction \hat{n} in the (x, y) plane is given by $|\Delta C| = |\Delta z| \cot(\alpha)$.

Let z denote time in the course of evolution, *i.e.*, $z = t$. Since $E = \rho\lambda \cos(\alpha)$, we have

$$\left| \frac{\Delta C}{\Delta t} \right| = \cot(\alpha) = \frac{E/\rho\lambda}{\sqrt{1 - (E/\rho\lambda)^2}} \quad (11)$$



$$\sin(\alpha) = \sqrt{1 - \left(\frac{E}{\delta\lambda}\right)^2}$$
$$\cos(\alpha) = \frac{E}{\delta\lambda}$$

[pdf document](#)

Kimmel, Siddiqi, Kimia, Bruckstein



Proposed Solution: Evolve a curve such that it tracks the height contours of $z(x, y)$.


[Kimmel *et al.*, IJCV95]

The curve evolution equation is:


$$\begin{cases} \frac{\partial \mathcal{C}}{\partial t} &= \frac{E/\rho\lambda}{\sqrt{1-E^2/(\rho\lambda)^2}} \cdot \hat{n}, \\ \mathcal{C}(s, 0) &= \mathcal{C}_0(s). \end{cases}$$

KimmeI, Siddiqi, Kimia, Bruckstein


Examples - Pyramids




shaded image



equal height contours




numerical solution




true surface

KimmeI, Siddiqi, Kimia, Bruckstein


Examples - Three Mountains




shaded image



equal height contours



numerical solution



true surface



Application Area: Geography






Abb. 13: Schräglichts-Schattrelieftopografie des Untersuchungsgebietes auf Basis des verbesserten DGM (Auflösung = 10m), überlagert mit den roten (=rot) und blauen (=blau) Strömnetzen und den Höhenlinien (=braun). Abbildungsmaßstab 1:25.000.



Application: Braille Code






Abbildung 3:

Oben links: Messanordnung mit einer Kamera und vier blauen LED-Leuchtfeldern.

Unten links: Ausschnitt einer Faltschachtel mit Blindenschrift-Prägung.

Rechts: 3D-Bild nach SFS-Analyse. Darunter ist ein Höhenprofil durch drei Braille-Punkte dargestellt.

Mars Rover Heads to a New Crater NYT Sept 22, 2008





Limitations

- Controlled lighting environment
 - Specular highlights?
 - Partial shadows?
 - Complex interreflections?
 - Fixed camera
 - Moving camera?
 - Multiple cameras?
- = > Another approach: binocular / geometric stereo