CSE252a – Computer Vision – Assignment 2
Instructor: Prof. David Kriegman.
Revision 2

Instructions:

- This assignment should be solved, *and written up* in groups of 2. Work alone *only* if you can not find a partner. No groups of 3 are allowed.
- Attempt all questions
- Submit your assignment electronically by email to obeijbom@cs.ucsd.edu with the subject line *CSE252 Assignment 2*. The email should have two files attached.
  1. A pdf file with your writeup. This should have all code attached in the appendix. Name this file: CSE_252_hw2_writeup_lastname1_lastname2.pdf.
  2. A compressed archive with all your matlab code files. Name this file: CSE_252_hw2_code_lastname1_lastname2.zip.

  The code is thus be attached *both* as text in the writeup appendix and as m-files in the compressed archive.
- Please make this a proper report, with methods, thoughts, comments and discussions. All code should be tucked away in an appendix.
- No physical hand-in for this assignment.
- You may do problems on pen an paper, just scan and include in the writeup pdf file.
- In general, MATLAB code does not have to be efficient. Focus on clarity, correctness and function here, and we can worry about speed in another course.

Irradiance 1 [3 pts]

Consider a rectangular surface with vertices $(-3, -2, 1), (3, -2, 1), (-3, 2, 1),$ and $(3, 2, 1)$. If the radiance on the surface is $L \cdot (x^2 + y^2 + 3)$ (with L constant) what is the irradiance arriving at position $(0, 0, 0)$ with normal vector $(0, 0, 1)$?

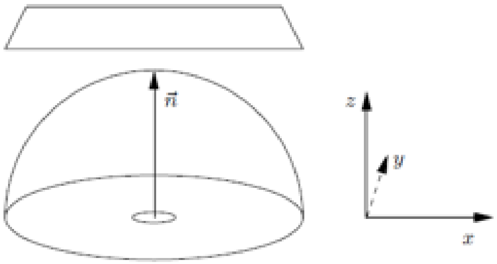

Figure 1: Problem 1 Setup

Irradiance 2 [3 pts]

a. Consider a cylinder with radius r and height h whose base is centered at $z = 0$ along the xy-plane. If the walls of the cylinder have constant radiance L and the top of the cylinder has constant radiance $2L$, what is the irradiance E at the point $(0, 0, 0)$ assuming that the surface at $(0, 0, 0)$ has a normal vector of $(0, 0, 1)$?

b. What is the irradiance if the radiance of the top is now Ld^2 where d is the distance to the center of the top ($d \in [0, r]$)?

Express your answers in terms of the given quantities (e.g. convert any angles you might have introduced to the given quantities).

Steradians [1 pts]

How many steradians are in a quarter sphere?

Lambertian surfaces [2 pts]

A Lambertian surface is one that appears equally bright from all viewing direction. In other words, the emitted radiance from a Lambertian surface is not a function of outgoing direction. Assume that we have an ideal Lambertian surface which also reflects all incident lights (absorbing none), the BRDF $\rho(\theta_{in}, \psi_{in}, \theta_{out}, \psi_{out})$ of such a surface will be a constant. What is that constant?

Hint: If we integrate the radiance of the surface over all directions, the total emitted energy is equal to the incoming energy.

Photometric Stereo and Specularity Removal [10pts]

The goal of this part of the assignment is to implement a couple of different algorithms that reconstruct a surface using the concept of photometric stereo. Additionally, you will implement the specular removal technique of Mallick et al., which enables photometric stereo reconstruction of certain non-Lambertian materials. You can assume a Lambertian reflectance function once specularties are removed, but the albedo is unknown and non-constant in the images. Your program will take in multiple images as input along with the light source direction (and color when necessary) for each image. You will also implement a second example-based photometric stereo algorithm which is based on simultaneously imaging two objects of the same material, one of which has known structure.

Part 1

Implement the photometric stereo technique described in Forsyth and Ponce 5.4 and the lecture notes. Your program should have two parts:

a) Read in the images and corresponding light source directions, and estimate the surface normals and albedo map.

b) Reconstruct the depth map from the normals. You can first try the naive scanline-based shape by integration method described in the book. If this does not work well on real images, you can use the matlab implementation of the Horn integration technique at the bottom of this assignment.

Try this out on the synthetic dataset (synthetic_data.mat) with three subsets of images:

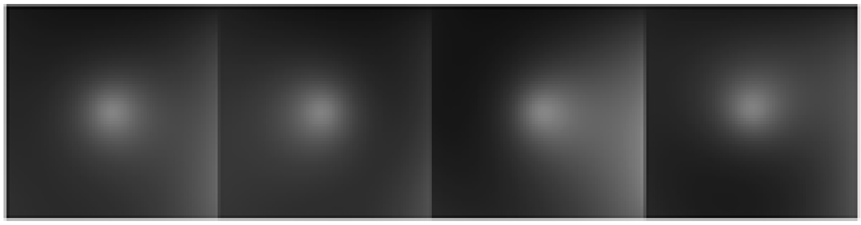

Figure 2: Synthetic Data

a) im1, im2, im4

b) all four images (Most accurate)

What to include in your report:

a) Estimated albedo map

b) Estimated surface normals by either showing

- Needle map (you will need to subsample the image to get a needle map which can be displayed. You can use the matlab functions meshgrid() and quiver3())
- Three images showing three components of surface normal

c) A wireframe of a depth map (you can use surf() in matlab).

Part 2

Implement the specularity removal technique described in Beyond Lambert: Reconstructing Specular Surfaces Using Color (by Mallick, Zickler, Kriegman, and Belhumeur; CVPR 2005). Your program should input an RGB image and light source color and output the corresponding SUV image. Try this out first with the specular sphere images and then with the pear images. What to include in your report: For each specular sphere and pear images.

a) The original image (in RGB colorspace).

b) The recovered S channel of the image.

c) The recovered diffuse part of the image - Use $G = \sqrt{U^2 + V^2}$ to represent the diffuse part.

Part 3

Combine parts 1 and 2 by running your photometric stereo code on the diffuse components of the specular sphere and pear images. For comparison, run your photometric stereo code on the original images (converted to grayscale) as well. You should notice erroneous ”bumps” in the resulting reconstructions - the result of violating the Lambertian assumption. What to include in your report: For each specular sphere and pear images.

a) The recovered diffuse images ($G = \sqrt{U^2 + V^2}$)

b) Estimated albedo map (original and diffuse images)

c) Estimated surface normals (original and diffuse images) by either showing

- Needle map (you will need to subsample the image to get a needle map which can be displayed. You can use the matlab functions meshgrid() and quiver3())
- Three images showing three components of surface normal

d) A wireframe (original and diffuse images) of a depth map (you can use surf() in matlab).

Part 4

Implement the photometric stereo algorithm described in section 2 of Shape and Materials by Example: A Photometric Stereo Approach (by Hertzmann and Seitz). Use the algorithm to reconstruct the 3D shape of the bottle using the set of corresponding bottle and sphere images. You may find it useful to use an algorithm to speedup the nearest-neighbor search, such as a kd-tree. What are the differences in assumptions made (e.g., with respect to light source, BRDF, etc.) for this algorithm compared to the algorithm from Part 3? What to include in your report:

a) Estimated surface normals of the bottle by either showing

- Needle map (you will need to subsample the image to get a needle map which can be displayed. You can use the matlab functions meshgrid() and quiver3())
- Three images showing three components of surface normal

b) A wireframe of a depth map (you can use surf() in matlab).

Resources

Data

Synthetic Images, Specular Sphere Images, Pear Images : Available in a Matlab *.mat file (graciously provided by Satya Mallick) which contains

a) im1, im2, im3, im4 ... images.

b) l1, l2, l3 , l4 ... light source directions

c) c (when required) color of light source.

Bottle and Sphere Images provided from Shape and Materials by Example: A Photometric Stereo Approach (by Hertzmann and Seitz).

All files are on the course webpage.

Integration Code

While a naive integration routine will probably work ok for synthetic data, it might have problems with real images. Here we provide a matlab implementation of the Horn integration technique (summarized here by Satya Mallick). You still need to write your own integration code, but if your algorithm doesn't handle real images well, use this routine instead.

Good luck!