Assignment A8: 3D from Stereo Image Pairs (adapted from Prof. Gerig, Spring 2013)

CS 6320
Spring 2014

Assigned: 19 February 2014

Due: 5 March 2014

For this problem, handin a lab report pdf (include name, date, assignment and class number in pdf) which studies camera calibration. You should handin the report pdf as well as the Matlab code used in the study. The code should conform to the style requested in the class materials.

In addition, please turn in a hardcopy of the report in class before the start of class on March 5, 2014.

The goal of this assignment is to implement an algorithm that reconstructs a surface using the concept of photometric stereo. You can assume a Lambertian reflectance function, but the albedo is unknown and nonconstant in the images. Your program will read multiple images as input along with the light source direction for each image. All data sets for this assignment are provided are provided on the class web page under the Data link (in subdir photometric-stereo). These synthetic images are called PS1.tif, PS2.tif, PS3.tif, and PS4.tif,

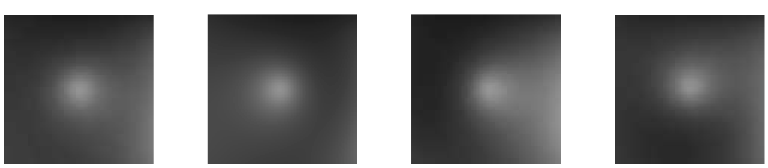

Figure 1: Photometric Stereo Input Set.

and have illumination directions $[0, 0, 1]^T$, $[0.2, 0, 1]^T$, $[-0.2, 0, 1]^T$, and $[0, 0.2, 1]^T$. Make sure they are scaled to the range 0 to 1.

Your program should have two parts:

- Read in the images and estimate the surface normals and the albedo map.
- Reconstruct the depth map from the normals via integration.

Since we have given you four images (size 100x100), you have the possibility for an exact solution (use 3 images only) or an overconstrained solution (all four images).

For the second part to get surface from normals, you should implement the naïve direct integration approach to integrate the partial derivatives and construct the depth map.

Shape from Shading from Real Images

Use the dog images in Figure 2; they have illumination directions $[16, 19, 30]^T$, $[13, 16, 30]^T$,

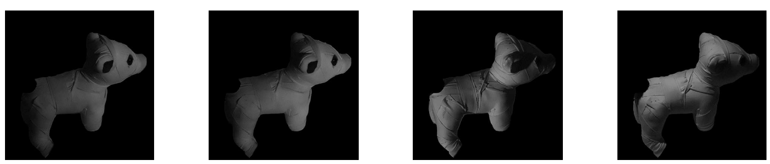

Figure 2: Photometric Stereo Real Input Set.

$[-17, 10.5, 26.5]^T$, and $[-9, 25, 4]^T$. Make sure they are scaled to the range 0 to 1.

Integration

Integration from normals to get the depth map is more difficult for noisy real images since integrability is violated and normals are only estimated approximately. You might first use a single integration step as shown in the slides and Forsyth and Ponce book. Alternatively, you might apply integration with different start-points and then average the resulting depth map.

Discussion

You should write up a report including your approach and results:

- Estimated albedo map
- Estimated surface normals by either showing:
  - A needle (vector) map (e.g., 2D vectors on the 100x100 grid), or
  - Three images showing the three components of surface normals (normalized normals),
- The reconstructed surface as:
  - Graylevel depth image where depth z is encoded as intensity, and
  - wireframe of a depth map (you can use surf() in matlab) to display Z(X,Y) similarly to figure 5.13 on page 86, and
  - shaded image generated from a user-dened light source position (given a light source direction and surface normals, you can easily calculate shaded images for arbitrary light source positions).

Commentary about any issues that arose, ways to improve your method, critical assessment of results, etc.