CMPT 888 Assignment 3: Photometric Stereo

Colin J Brown, Instructor: Ping Tan

September 17, 2015

1 Introduction and Method

When a static scene is photographed with varying lighting conditions, information about the surfaces of objects in the scene can be extracted [1]. The purpose of this assignment was to extract the surface normals and albedo in scenes with a single, but changing, directional light source and matte and mirrored calibration spheres (see Figure 1a) .

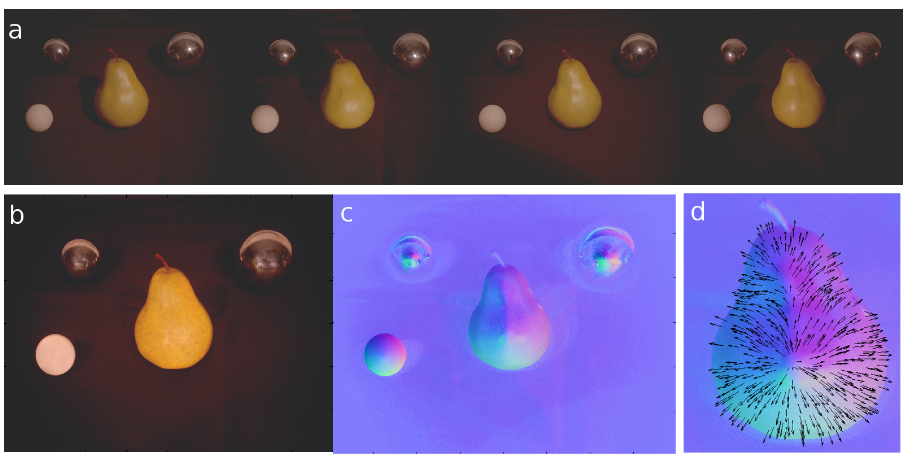

Figure 1: a) Four sample images from the pear dataset with different lighting angle in each image. b) Extracted albedo color map. Note that the pear and matte sphere appear flat since shading is removed and only color and texture remain. c) Extracted normal map. d) Normal vectors sampled on pear.

The light direction in each scene was first approximated using the mirror spheres (for which binary masks were given). For a mirrored surface, reflecting the view vector (assumed WOLOG to be [0,0,1]) about the normal, $\vec{n}$, at the brightest point gives the light direction. This normal was calculated using the known geometry of a sphere. Since two mirrored spheres were provided in each scene, the light direction vectors from each image were taken as an average of the light directions calculated from each sphere. Light intensity was calculated using the matte sphere (for which, also, a mask was given). Since matte materials scatter light in all directions evenly, the brightest point on the matte sphere gives the light intensity.

With lighting calculated, for each image, the scene objects' surface albedo, ρ, and normals, n, could be computed at each pixel. For a given pixel, let I represent the column vector of pixel intensities across each image. Then for L, a matrix containing a lighting row vector ($\vec{l}$, encoding lighting direction and intensity) for each image, we have,

$$L^{-1}I = \rho\vec{n} \tag{1}$$

Note that $\vec{n}$, is a unit vector, so $\rho = ||\rho\vec{n}||$. Since we are using a Lambertian model here and not taking light occlusion into account, our model cannot handle specular highlights or shadows. Thus, for each pixel, we remove the brightest 25% and darkest 25% of pixels before computing ρ and $\vec{n}$. The algorithm was applied on each colour channel independently and then surface normals were averaged across the 3 channels.

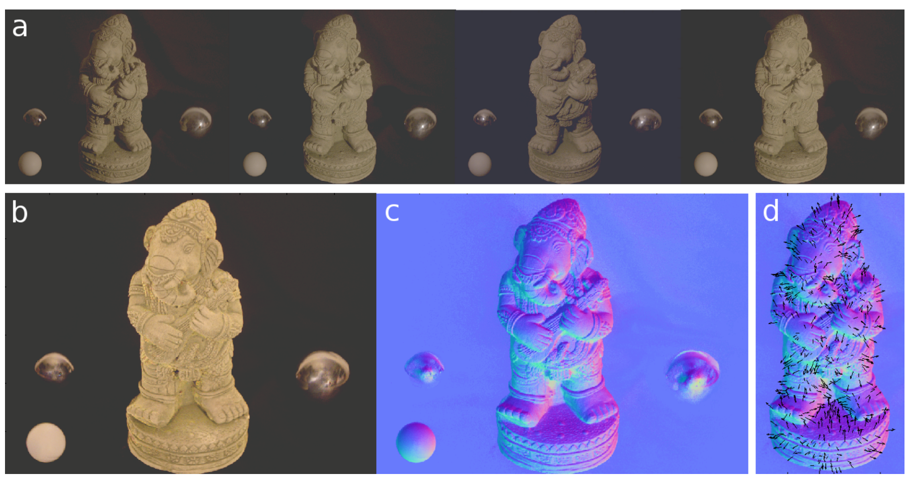

Figure 2: a) Four sample images from the elephant statue dataset with different lighting angle in each image. b) Extracted albedo color map. c) Extracted normal map. d) Normal vectors sampled on statue.

2 Experiments and Discussion

The method was tested on two datasets, a pear scene (Figure 1) and an elephant statue scene (Figure 2), each containing 21 (linearized) images of static scenes with changing lighting. It is clear that despite removal of specular highlights, this algorithm produces better results on objects with matte materials. For example, whereas the normal map for the elephant statue (Figure 2c) appears correct despite the high level of detail, the pear, which has a shinier surface, appears incorrectly to have a strong ridge down the middle. Furthermore, it is evident in both scenes that normal maps of the mirrored spheres are completely incorrect, whereas the normal maps of the matte spheres appear very accurate.

Shadows also cause artifacts when using this approach, despite measures to exclude pixels likely to be in shadow. For example, in Figure 1b, faint shadows can be seen to the right of the right and left of each object even though this albedo map should be independent of any lighting information. Similar artifacts can be seen in the normal map (Figure 1c). While more agressive culling of the brightest and darkest pixels may mitigate the effects of highlights and shadows, it would also decrease the stability of the solution and begin to introduce noise. Instead, perhaps use of a model which explicitly takes these phenomena into account would be a better solution.

Once the albedo and normal maps were generated, the scenes were re-rendered under novel lighting directions (see Figures 3 and 4). Despite a lack of depth information, the synthetic re-lighting of the scenes is surprisingly realistic.

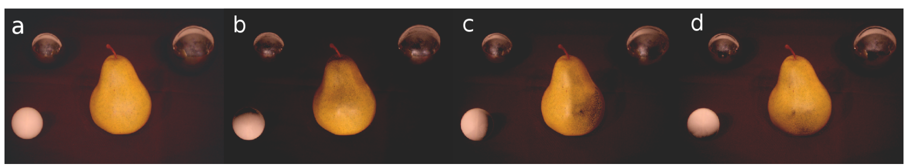

Figure 3: Renderings of the pear scene with light direction a) from camera, b) from below, c) from left and d) from above.

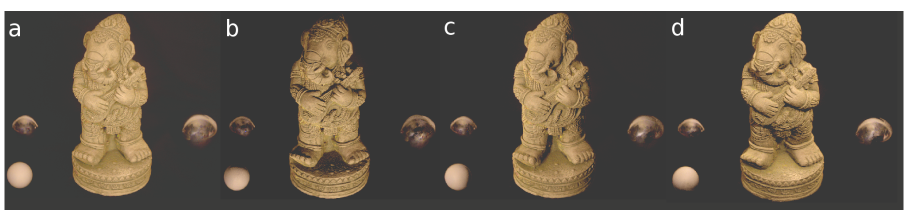

Figure 4: Renderings of the elephant statue scene with light direction a) from camera, b) from below, c) from left and d) from above.

However, again since specular highlights and shadows are not modelled, they are conspicuously missing from the rendered images. Specular highlights could be added using a different rendering model (e.g. Phong) and selecting reasonable specularity parameters for each material. By using the normal maps to reconstruct each object in 3D, shadows could be re-introduced via raytracing.

References

[1] Robert J Woodham. Photometric stereo: A reflectance map technique for determining surface orientation from image intensity. In *22nd Annual Technical Symposium*, pages 136–143.

International Society for Optics and Photonics, 1979.