

Intro to Computer Vision: HW 3

Alvin Lin

August 2018 - December 2018

Part A

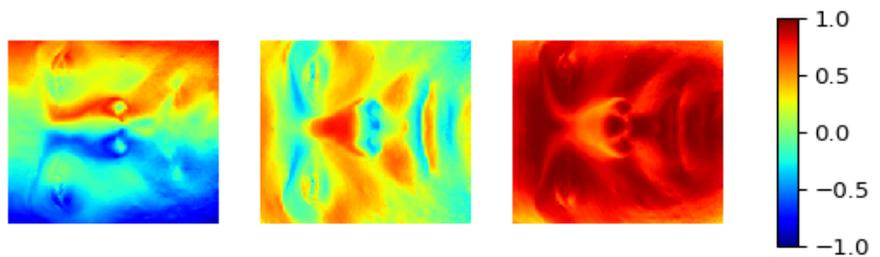
I implemented the `load_images()` function by using regular expressions to parse the filename for the azimuth and elevation and the `os` module to fetch all the files in the data directories. As a result, the `num_images` parameter ended up unused.

I implemented the pseudocode for photometric stereo as described in the assignment using many of the operations provided by `numpy`. The fast vector manipulations made many of the transpositions and calculations very easy. From the height map, you can see horizontal lines across the face as differences in shade. This is likely due to variations in illumination across the entire image dataset which become exaggerated by the integration process. The scale of height is also not properly captured by this process, since peaks such as the regions around the nose and the chin are not proportioned correctly.

Part B



(a) Albedo image for yaleB01



(b) Surface normals for yaleB01

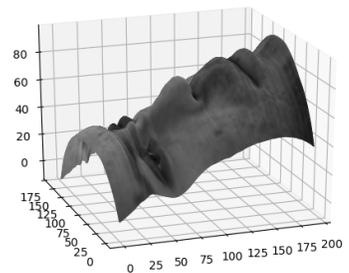
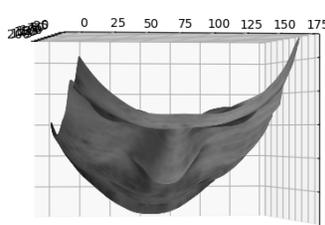
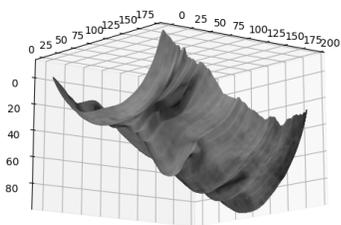


Figure 2: Height map of yaleB01

Part C

Many of the assumptions about illumination with regard to shape from shading are not true due to the nature of the surface. With shape from shading, we assume that the surface does not reflect light, but in reality, light bounces off of regions around our face (especially our nose) and contributes to the lighting of nearby features. The eye in particular tends to reflect a lot of light, which is evident from the fact that the area around the eyelid and eyeball are not correctly proportioned in our height map reconstruction. The eye position and the hair in all the images is also not uniform across the dataset, which can contribute to some error as seen in the examples below.



(a) Hair shadows are visible in this image



(b) Light is reflected off the person's forehead

If you have any questions, comments, or concerns, please contact me at alvin@omgimanerd.tech