Additional Technique Details: Multi-view Photometric Stereo with Spatially Varying Isotropic Materials

Zhenglong Zhou Zhe Wu Ping Tan
National University of Singapore

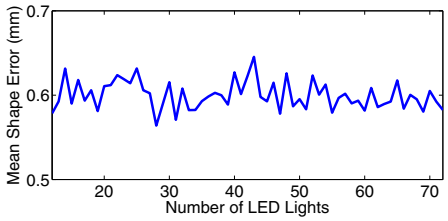

Figure 1. Mean shape error of the 'Buddha' example. This error does not change significantly with different number of LEDs.

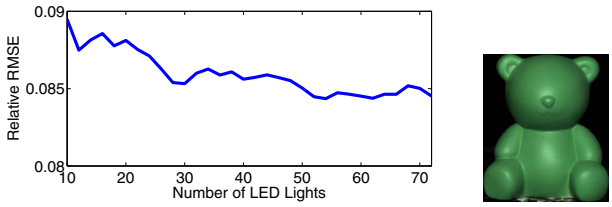

Figure 2. Relative RMSE BRDF error of a green paint. This error generally decreases with more LEDs.

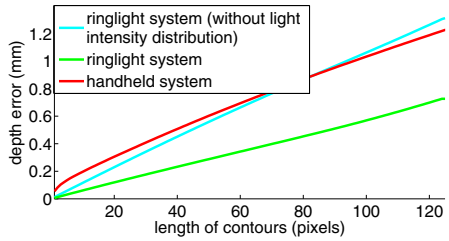

Figure 3. Average depth error of all iso-depth contours of the 'Cat' example. This error increases linearly with contour length.

1. Number of images at each viewpoint

We evaluate the accuracy of captured shape and BRDF with different number of input images from each viewpoint. We only perform this evaluation with the ring-light system for simplicity in data capture.

We first evaluated the shape accuracy on the 'Buddha' example shown in Figure 7 of the main paper. Figure 1 shows the mean shape reconstruction error (in millimeters) as a function of the number of LEDs in each viewpoint. We always chose equal number of uniformly distributed lights on both the outer and inner circles. Since our Fourier series fitting requires at least 5 LEDs from each viewpoint, we begin the plot from 10 lights (5 on each circle). We found the mean shape error did not change significantly for different number of LEDs.

We also evaluated the reflectance accuracy with different number of LEDs. We first measured the BRDF of a green paint after applying it to a sphere of known shape, and capturing the BRDF from images with calibrated directional lighting. We took this measurement as 'ground truth' and compared our result with it. This experiment is evaluated with another painted figurine shown on the right of Figure 2. The left of Figure 2 shows the relative root mean square error (RMSE) of our result with different number of LEDs from each viewpoint. We computed RMSE as the following

$$E_{rms} = \left(\sum_{\theta_h,\theta_d,\phi} \frac{(f(\theta_h,\theta_d,\phi) - \hat{f}(\theta_h,\theta_d,\phi))^2}{f(\theta_h,\theta_d,\phi)^2} \right)^{\frac{1}{2}}$$

Here, $f(\cdot)$, $\hat{f}(\cdot)$ are the 'ground truth' and recovered BRDFs respectively. This error converges to about 8.5% when about 30 LEDs are used. So in our experiments, we always used 30 LEDs for the ring-light system.

2. Errors in iso-depth contours

We also evaluated the accuracy of our obtained iso-depth contours with the 'Cat' example under both capture setups. Figure 3 shows the average depth error as a function of the contour length. The depth error was measured as the difference between the maximum and minimum 'ground truth' depths along an iso-depth contour. The red and green curves are the result from the handheld and ring-light systems respectively. The handheld system had larger error, since we only calibrated the overall light source intensity there, while we calibrated the light intensity distribution (see Section 3) for the ring-light system. The error of the ring-light system will become larger as shown by the cyan curve, if we only calibrate the overall intensity of each LED. Typically, the depth error increased linearly with the length of iso-depth

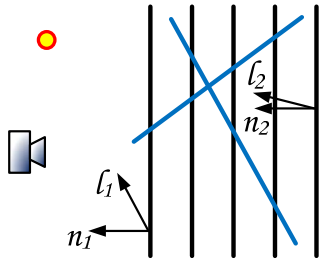

Figure 4. Top view of the calibration setup. We capture a diffuse board at several known positions (black lines) to calibrate camera vignetting and light intensity. Some additional boards (blue lines) are used to calibrate the angle θ_0.

contours. Though shorter contours had less depth fluctuation, longer ones required less iterations of depth propagation to cover the whole surface, and hence, accumulated less error. In our experiments, we always used iso-depth contours of 50-pixel length, which often had about 0.6 and 0.3 millimeter depth changes for the handheld and ring-light system respectively.

3. Calibrating the Ring-Light Setup

We calibrate the intensity distribution of each LED. Note that an LED is not an ideal point light source. Its intensities towards different directions are different. We then determine θ_0 to compute positions of LEDs in the local camera coordinate system.

We calibrate intensity distribution by capturing a diffuse board roughly parallel to the image plane at multiple depths as shown in Figure 4. The black lines indicate the positions of this board viewed from top-down. A check board calibration pattern is printed at the four corners of the board, such that its 3D position can be computed with the method described in [1]. Assume the board is Lambertian with unit albedo. At each point, the observed pixel intensity should be $I = \mathbf{n}^\top \mathbf{l} V$. Here $\mathbf{l}, \mathbf{n}$ are the local lighting and normal directions, and V is the light intensity. Hence, we can capture V at each point on the board as $I/\mathbf{n}^\top \mathbf{l}$. We linearly interpolate these captured values to obtain the result in a continuous 3D volume. During our experiments, we always divide an observed pixel intensity by the lighting intensity at its 3D position.

To calibrate the rotation angle θ_0, we capture the diffuse board at some slanted positions (indicated as blue lines in Figure 4) and compute the azimuth angle of the normal direction of the board. The computed angle should be $\theta_0 + \alpha$, where α is the true azimuth angle. α is known, since the 3D position of the board is known from its calibration pattern. Hence, we can obtain θ_0 by subtracting α from the estimated values.

References

[1] Z. Zhang. A flexible new technique for camera calibration. *IEEE Trans. Pattern Anal. Mach. Intell.*, 22:1330–1334, 2000. 2