

LUCES: A Dataset for Near-Field Point Light Source Photometric Stereo

Roberto Mecca¹
rmecca@crl.toshiba.co.uk

Fotios Logotheitis¹
flogotheitis@crl.toshiba.co.uk

Ignas Budvytis²
ib255@cam.ac.uk

Roberto Cipolla²
rc10001@cam.ac.uk

¹ Cambridge Research Laboratory,
Toshiba Europe, Cambridge, UK
Cambridge, UK

² University of Cambridge
Cambridge, UK

Abstract

This document provides supplementary material for the main publication. Section 1 provides additional details of the Photometric Stereo camera and LED setup. Section 2 provides a discussion on the inherent errors introduced in shape estimation from normals. Section 3 provides a complete qualitative comparison of all 5 methods evaluated in the main publication.

1 Photometric Stereo Setup

This section gives additional details on the PS setup.




Figure 1: Close-up of the setup used for acquiring Photometric Stereo images.

It consists of the following main components (see Figure 1):

- RGB camera FLIR BFS-U3-32S4C-C with 8mm lens
- 52 LED Golden Dragon OSRAM
- variable voltage for adjustable LED power
- Arduino Mega 2560

The Arduino Mega 2560 controls the LEDs that are turned on and off individually during the Photometric Stereo image capture. The PCB has been designed to accommodate a specific 8mm lens that allows a reasonable wide field of view on a target object placed few centimeters away from the camera. The type of OSRAM LEDs (LW W5SN) are capable of emitting 5600 Kelvin white light up to 191 lumen. The 120° of viewing angle of the LEDs allows a complete lighting of the scene. The LEDs have been distributed on 6 different circumferences (all centered in the camera centre) of radii 35, 45, 55, 65, 75 and 85 mm. Depending on the circumference, the LEDs have been positioned at variable angle of 30° and 60° .

In order to avoid out of focus areas, the aperture of the lens is kept to the minimum ($f/11$). Therefore, variable time of exposures (among objects) have been used to prevent saturations or too dark regions. Note that a separate set of calibration images (see main paper Section 3.1) was captured for each different configuration.

The complete LED parameters (positions, brightness, directions, angular dissipation, see main paper Section 3.1) are included in the dataset.

2 Ground Truth Meshes

This section provides supplementary information about the laser-scanned meshes which were used as ground truth for the evaluation of the competing PS approaches. Note that because of the different sizes of the objects and the characteristics of the surfaces, the number of triangles in each object mesh varies from 260K to 6.1M as shown in Figure 2. As the scans can be of questionable quality at some regions due to visibility and/or specular material (i.e. numbers on the Die), manual segmentation was performed on the image domain to only evaluate on the reliable regions.

Ground truth discontinuity and non-differentiability. Over the years, the majority of PS approaches has assumed that the surface can be described as a continuous and differentiable depth map. It has been acknowledged before [9] that this assumption is violated in practice and some care has been taken to include some robustness to it (e.g. L_1 loss in [10] and Cauchy estimator by [9]). However, we believe that the extent of this issue is underestimated and thus we attempt to quantify it here by computing the following two metrics. Firstly, we compute a normal map through numerical differentiation of the ground truth depth map [9] and compare it with the ground truth normal map. This is shown in Table 1 of the main submission and illustrated visually here in Figure 3. The mean per-pixel angular error of computed and ground truth normals is on average 3.34° with the maximum value being 9.19° on the complex geometry house object. We emphasise that this effect is completely independent of the actual uncertainty of the ‘ground truth’ meshes and it is solely caused by the projection operation (and so the effect would be identical in synthetic data). Note, the ground truth normal map is computed by rendering (i.e. projection, discretisation and occlusion) of the

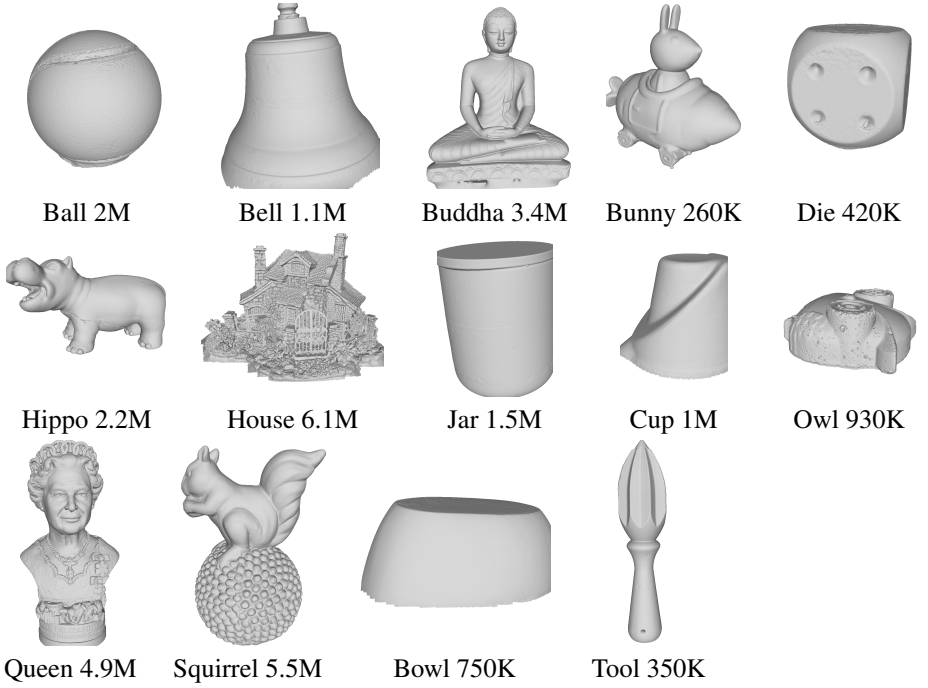


Figure 2: Laser-scanned meshes and their respective number of triangles.

surface normals into the image plane which is quite different than numerical differentiation of the ground truth depth map. Indeed, in Figure 3, the error is concentrated on boundaries.

In addition, we compute the average per-pixel error between the ground truth depth and the depth obtained by numerical integration (using [1]) of the ground truth normals (i.e. pseudo-depth). The observed error is not-negligible (1.67mm on average) and it propagates outwards from occlusion boundaries (as the numerical integration preserves the actual mean depth).


Finally, we note that the two error metrics explained above are likely to be close to the theoretical minimum (for normals and depth respectively) achievable by any approach that is reliant of the differentiable surface assumption. As these error bounds can be non-negligible, we motivate future research that avoids reliance of surface differentiability (e.g. direct depth regression).

3 Reconstructions


This section contains a complete qualitative comparison of all 5 methods L17 [1], Q18 [1], I18 [1], S20 [1] and L20 [1]. Estimated 3D surface view as well as depth Z error maps are provided in Figures 4 to 7. Note that errors of predicted normals are provided in the main publication.

References

- [1] S. Ikehata. Cnn-ps: Cnn-based photometric stereo for general non-convex surfaces. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 3–18, 2018.
- [2] F. Logothetis, R. Mecca, and R. Cipolla. Semi-calibrated near field photometric stereo. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 3, page 8, 2017.
- [3] F. Logothetis, I. Budvytis, R. Mecca, and R. Cipolla. A cnn based approach for the near-field photometric stereo problem. *BMVC*, 2020.
- [4] R. Mecca, Y. Quéau, F. Logothetis, and R. Cipolla. A single lobe photometric stereo approach for heterogeneous material. *SIAM Journal on Imaging Sciences*, 9(4):1858–1888, 2016.
- [5] Y. Quéau and J.-D. Durou. Edge-preserving integration of a normal field: Weighted least squares, TV and L1 approaches. In *SSVM*, 2015.
- [6] Y. Quéau, B. Durix, Tao Wu, D. Cremers, F. Lauze, and J.-D. Durou. Led-based photometric stereo: Modeling, calibration and numerical solution. *Journal of Mathematical Imaging and Vision (JMIV)*, 60(3):313–340, 2018.
- [7] H. Santo, M. Waechter, and Y. Matsushita. Deep near-light photometric stereo for spatially varying reflectances. In *European Conference on Computer Vision (ECCV)*, 2020.



Error maps between differentiated normals and ground truth normals.



Error maps between the integration of ground normals and the ground truth depth.

Figure 3: Differentiation/Integration errors.

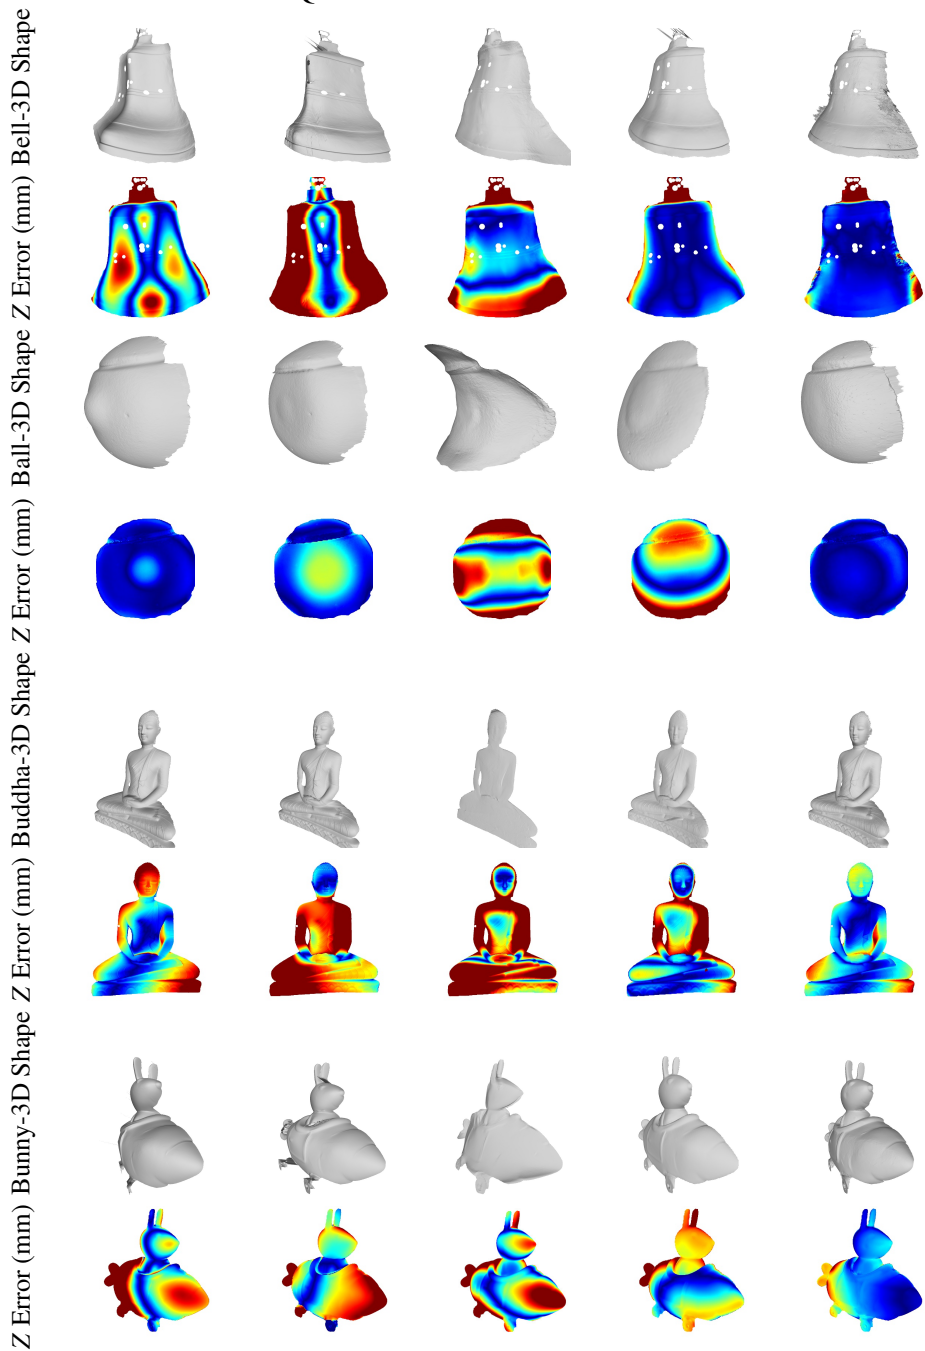


Figure 4: Evaluations 1-4/14

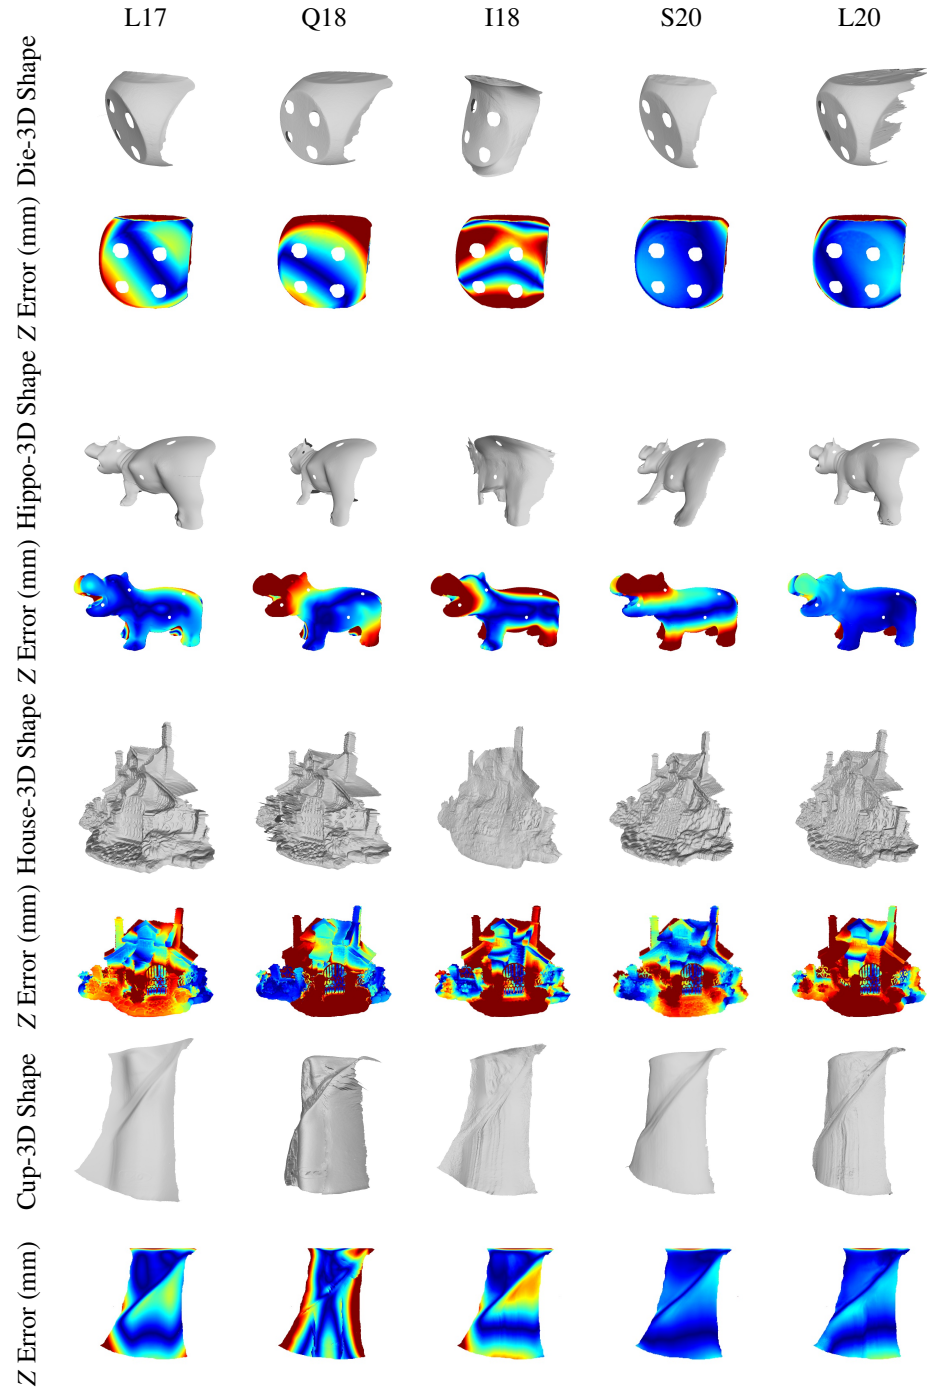


Figure 5: Evaluations 5-8/14

Z Error (mm) Squirrel-3D Shape Z Error (mm) Queen-3D Shape Z Error (mm) Jar-3D Shape Z Error (mm) Owl-3D Shape

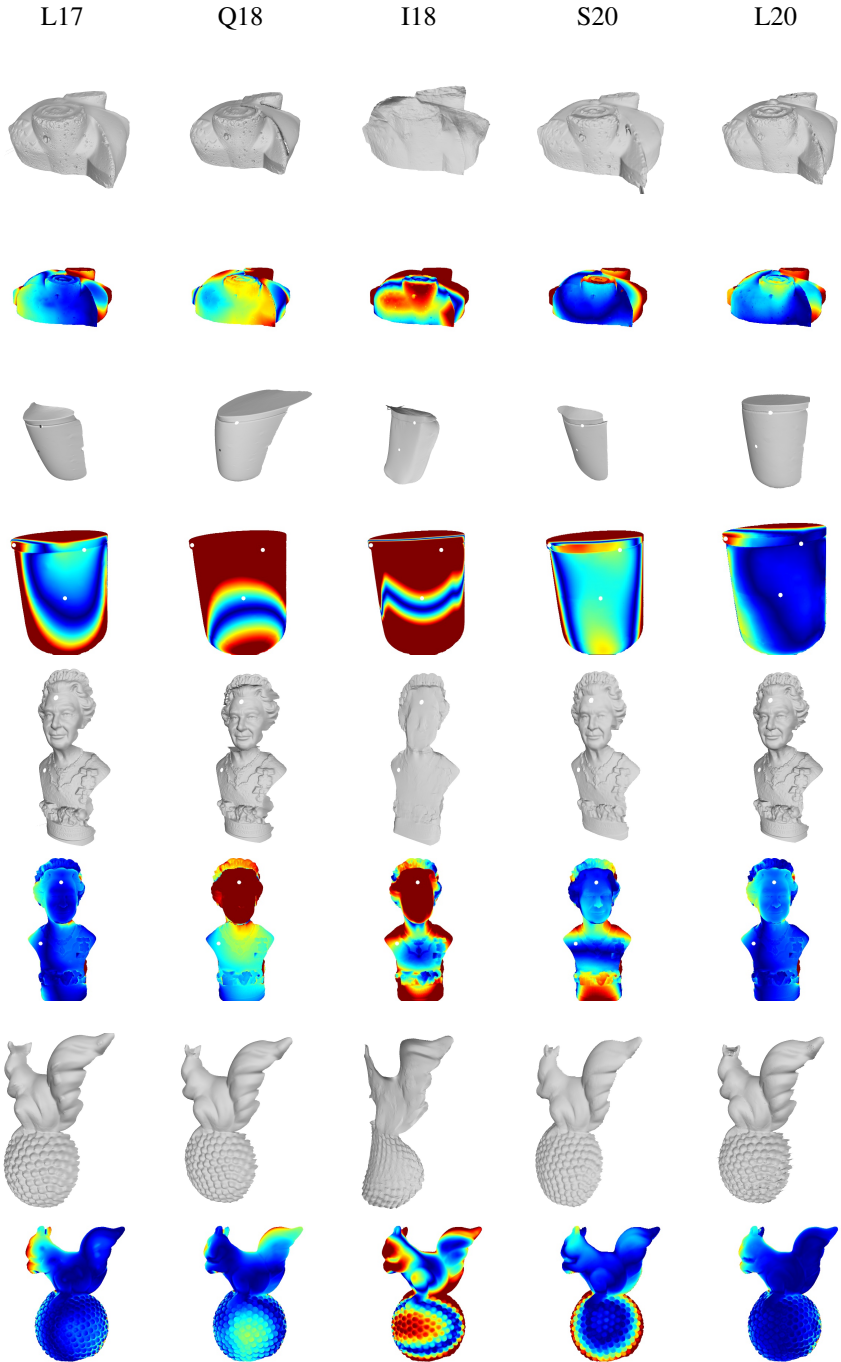


Figure 6: Evaluations 9-12/14




Figure 7: Evaluations 13-14/14