LERPS: Lighting Estimation and Relighting For Photometric Stereo

Ashish Tiwari and Shanmuganathan Raman
Computer Vision, Imaging, and Graphics (CVIG) Lab, Indian Institute of Technology Gandhinagar, India
E-mail: {ashish.tiwari, shanmuga} @ iitgn.ac.in

Paper No. 9119

Introduction

- We propose a deep learning based framework called **Lighting Estimation and Relighting for Photometric Stereo (LERPS)** designed to jointly perform the following tasks.
  (a) Lighting estimation
  (b) Image relighting
  (c) Per-pixel surface normal estimation
- During training, the network uses multiple differently lit images of an object one at a time.
- Inference is performed using just a single image.

Method

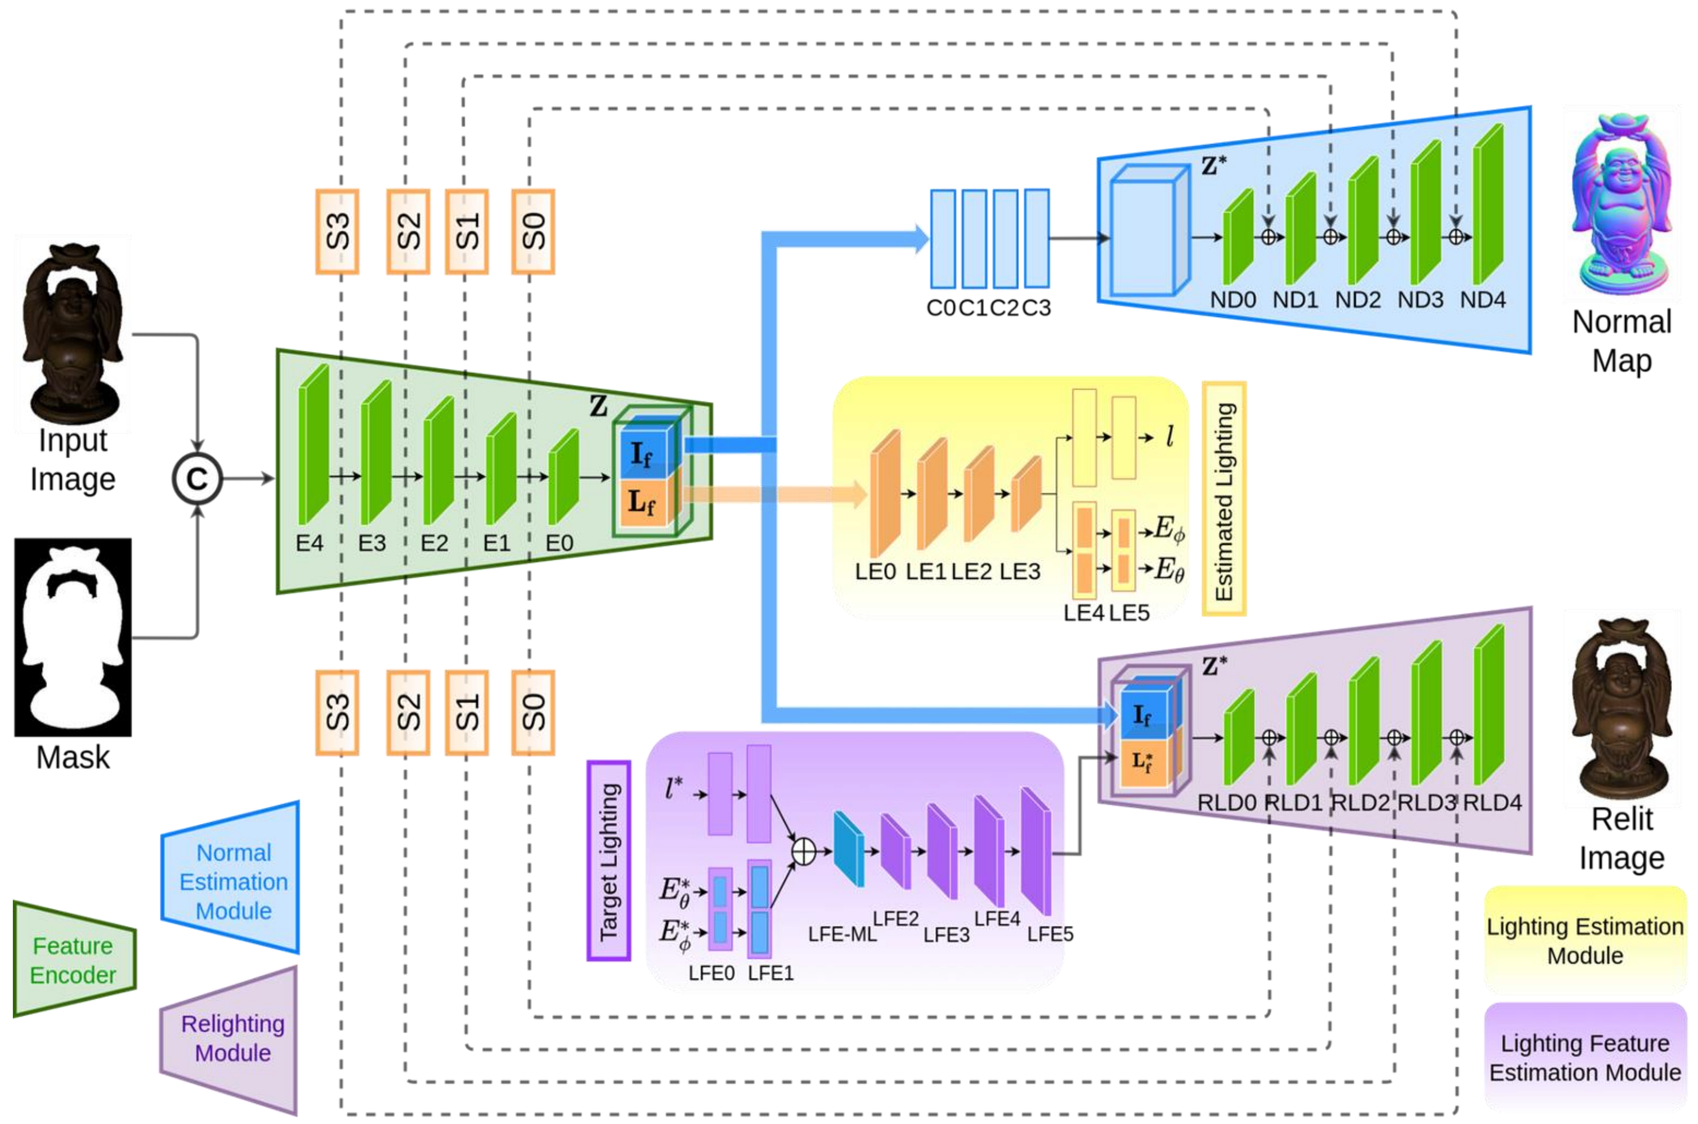

Fig. 1 Detailed architecture of LERPS framework.

- We use a combination of cross entropy loss, GAN loss, L1-loss (over image and gradient image), and feature matching loss for training the network.
- Feature Matching Loss:

$$L_F = \frac{1}{N}(\boldsymbol{I_{F1}} - \boldsymbol{I_{F2}})^2$$

I_{F1} and I_{F2} are geometric features of two input images I_1 and I_2 of the same object under different lightings.

Analysis

- The light space (the upper hemisphere) is characterized by azimuthal angle: $\phi \in [0°, 180°]$ and elevation angle: $\theta \in [-90°, 90°]$.
- We divide the light space into $K_d = 36$ bins.

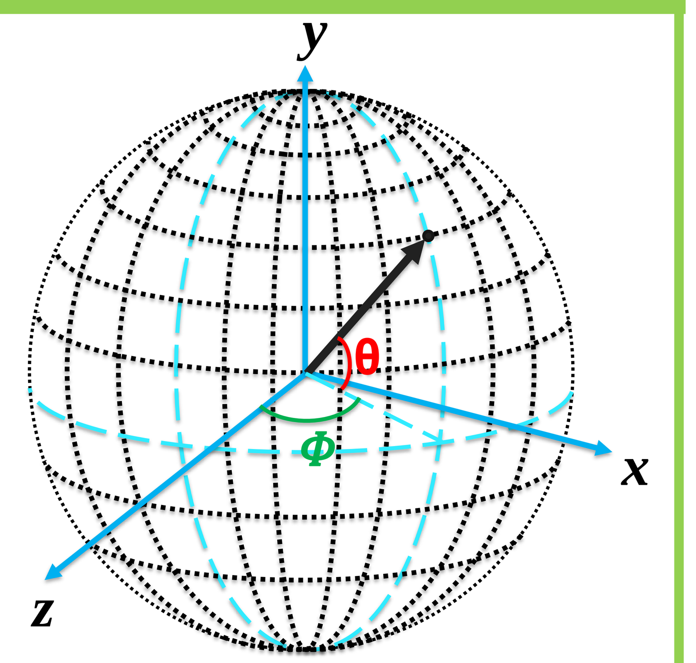

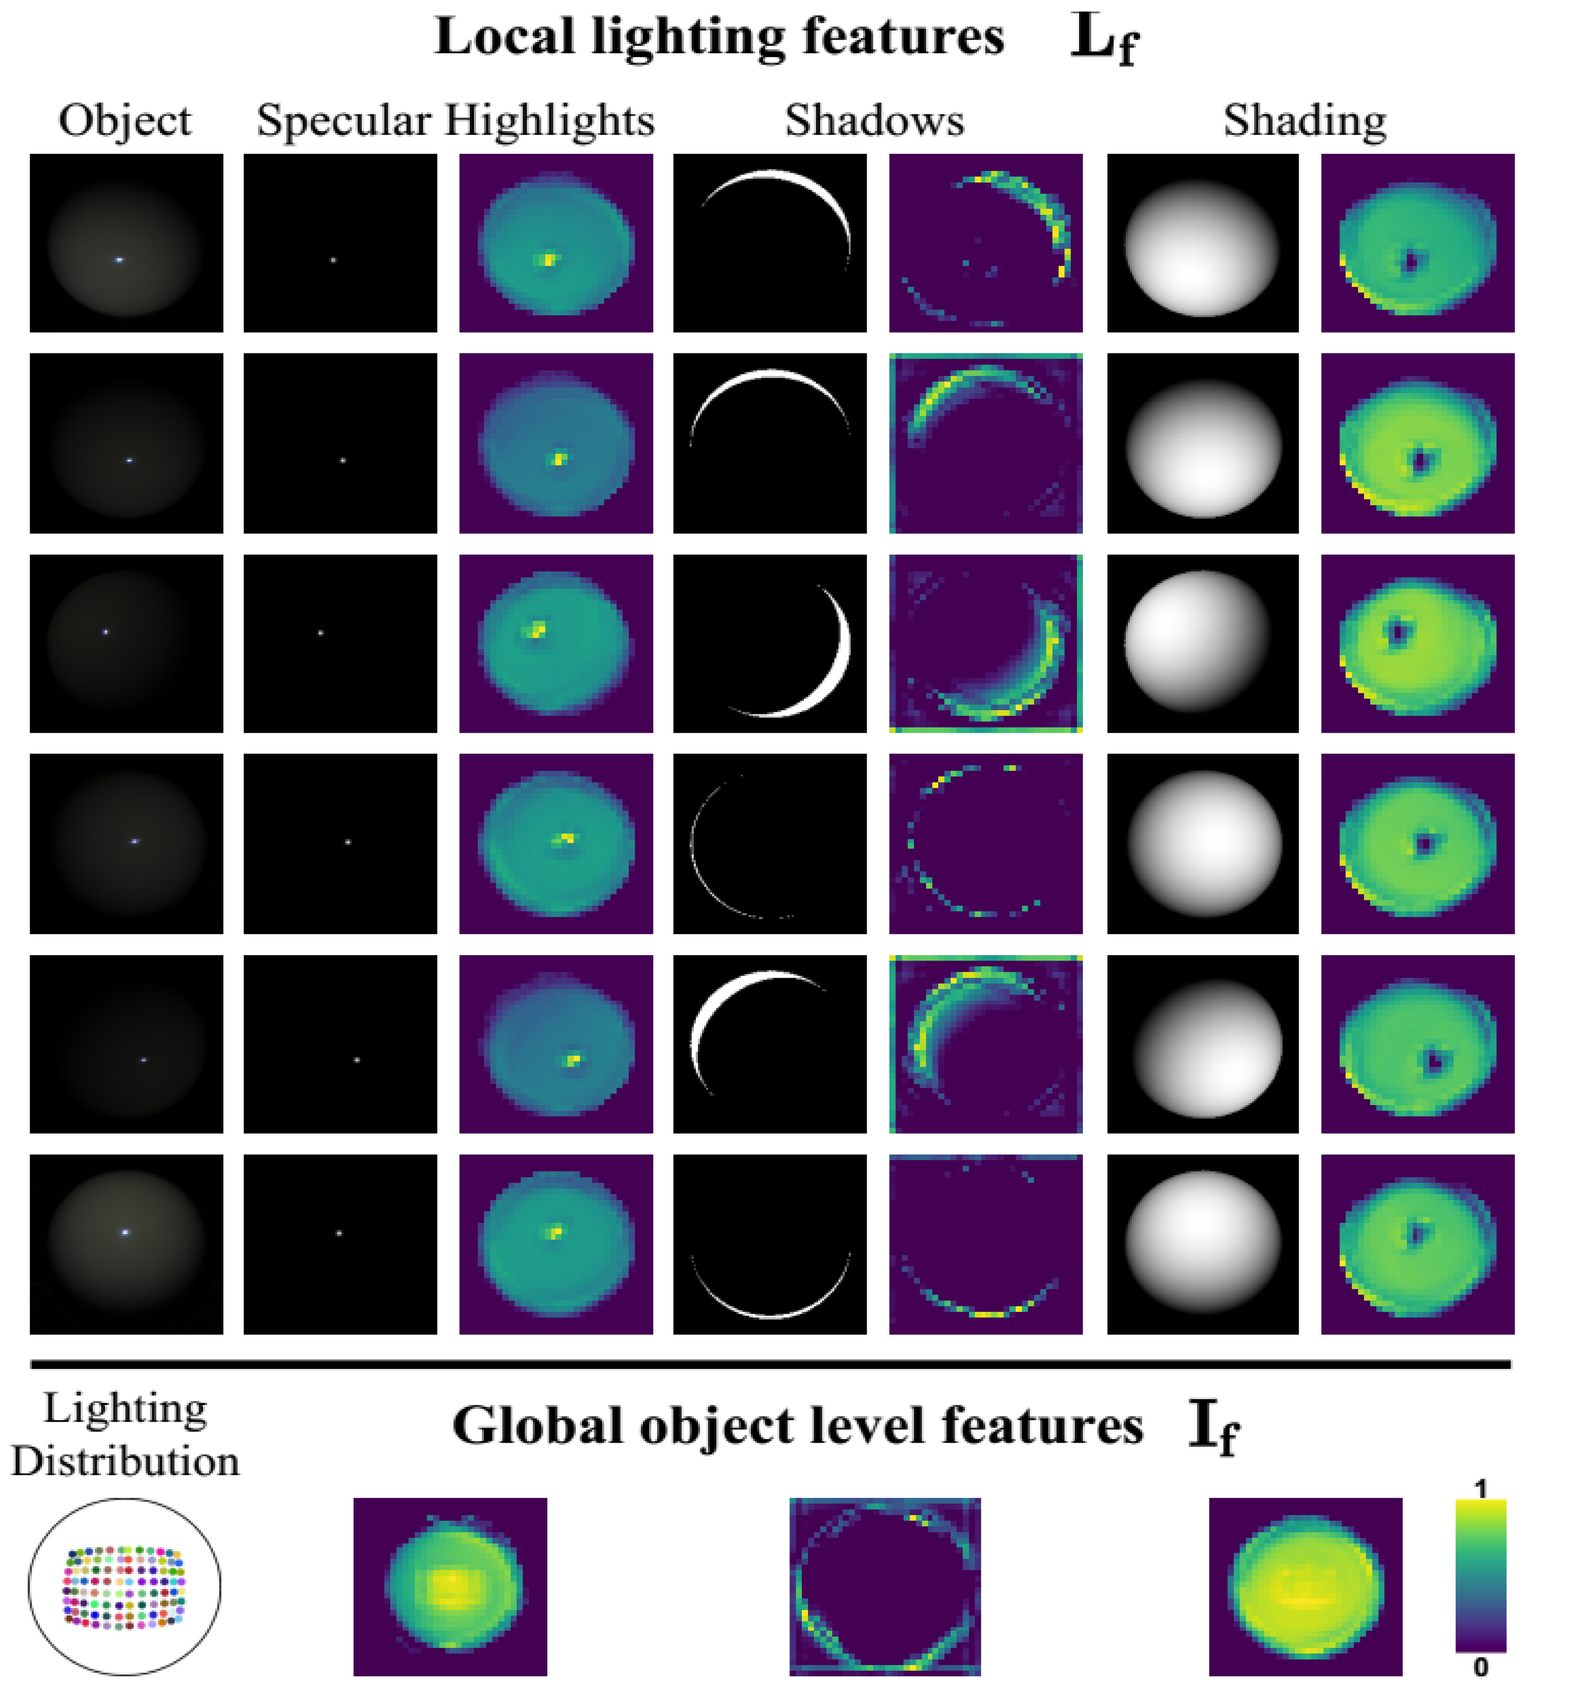

Fig. 2 Learned feature visualization.

- The network explicitly segregates global geometric features and local lightingspecific features of the object from a single image.
- The local features resemble attached shadows, shadings, and specular highlights, providing valuable lighting estimation and relighting cues.
- The global features capture the lighting-independent geometric attributes.

Results and Conclusion

- Our method ranks best with an average MAE of 7.89 and a standard deviation of 0.24 units for surface normal estimation over the DiLiGenT benchmark dataset.
- On an average, the relit images have over 95% structural similarity with the desired target images.
- LERPS captures and disentangles the global lighting-independent and the local lighting-specific features of the object.

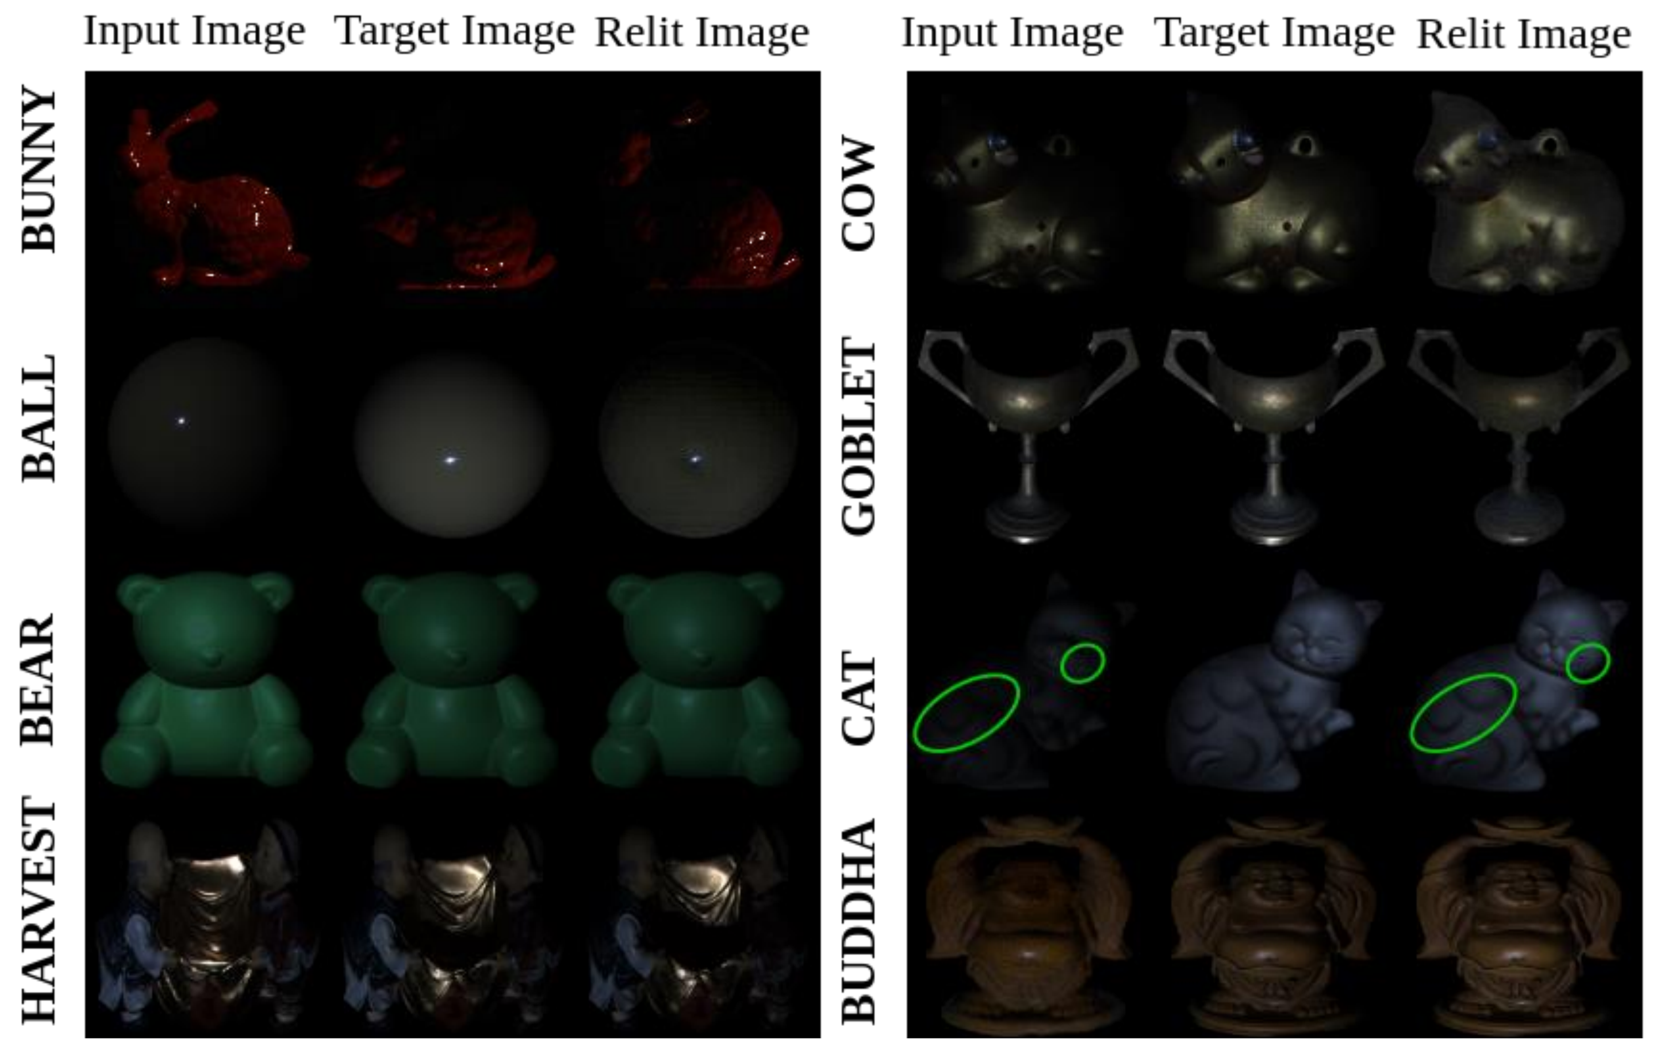

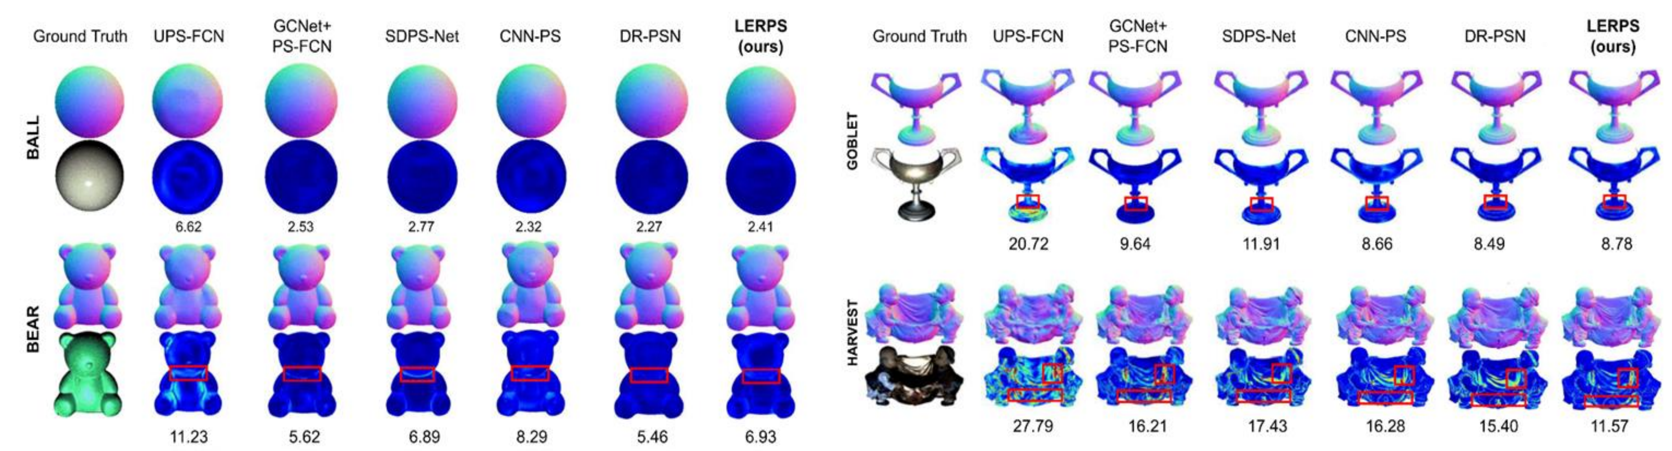

Key References

- Guanying Chen, Kai Han, Boxin Shi, Yasuyuki Matsushita, and Kwan-Yee K Wong, “Self-calibrating deep photometric stereo networks,” in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2019, pp. 8739–8747
- Boxin Shi, Zhe Wu, Zhipeng Mo, Dinglong Duan, Sai-Kit Yeung, and Ping Tan, “A benchmark dataset and evaluation for non-lambertian and uncalibrated photometric stereo,” in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2016, pp. 3707–3716.
- Guanying Chen, Michael Waechter, Boxin Shi, Kwan-Yee K Wong, and Yasuyuki Matsushita, “What is learned in deep uncalibrated photometric stereo?,” in European Conference on Computer Vision. Springer, 2020, pp. 745–762.