

Discriminative Shape from Shading in Uncalibrated Illumination

– Supplementary Material –

Stephan R. Richter Stefan Roth
Department of Computer Science, TU Darmstadt

In this supplementary material we show additional results from the quantitative comparison to other methods in laboratory and natural illumination (*c.f.* Table 2(b) and Fig. 6 in the main paper). Note that all parameters are the same as for the experimental results in the main paper. The results on the data of [4] taken under laboratory illumination are shown in Fig. 2, including a comparison to their method. Instead of relying on smooth local context, our method combines local and global context in a learning framework, allowing for better reconstructions overall.

Fig. 3 shows the remaining three surfaces that form our new ground truth dataset of real objects in natural illumination, along with results. As can be seen, our method adapts well to unknown reflectance maps, since it does not rely on an illumination prior unlike the generative cross-scale method [1]. Additionally, our method better reconstructs fine detail, since it does not need to rely on strong smoothness priors.

Training dataset. The 3D models we used to render our training dataset (Fig. 1) are available at [2]. The model set consists of 67 comic characters, partially in different poses, totaling 100 models. While these models do not correspond to real-world object categories, they contain varying curvature, big and small parts, discontinuities, as well as a few planar regions, which appears ideal for learning shape variation.

References

- [1] J. T. Barron and J. Malik. Color constancy, intrinsic images, and shape estimation. In *ECCV*, 2012. 1
- [2] Dosch Design. http://www.doschdesign.com/products/3d/Comic_Characters_V2.html. 1
- [3] M. K. Johnson and E. H. Adelson. Shape estimation in natural illumination. In *CVPR*, 2011. 1
- [4] Y. Xiong, A. Chakrabarti, R. Basri, S. J. Gortler, D. W. Jacobs, and T. Zickler. From shading to local shape. *TPAMI*, 37(1):67–79, 2014. 1, 2




Figure 1: Sample models from our training dataset. We rotate each model into several views and extract patches from the rendered surface normals (top row). For illustration, we show a rendering under an illumination from [3] (bottom row).




Figure 2: Comparison on real images under laboratory illumination with median angular errors. We show two novel views of the reconstructed surfaces. The input images and the views of “local context” are taken from [4].

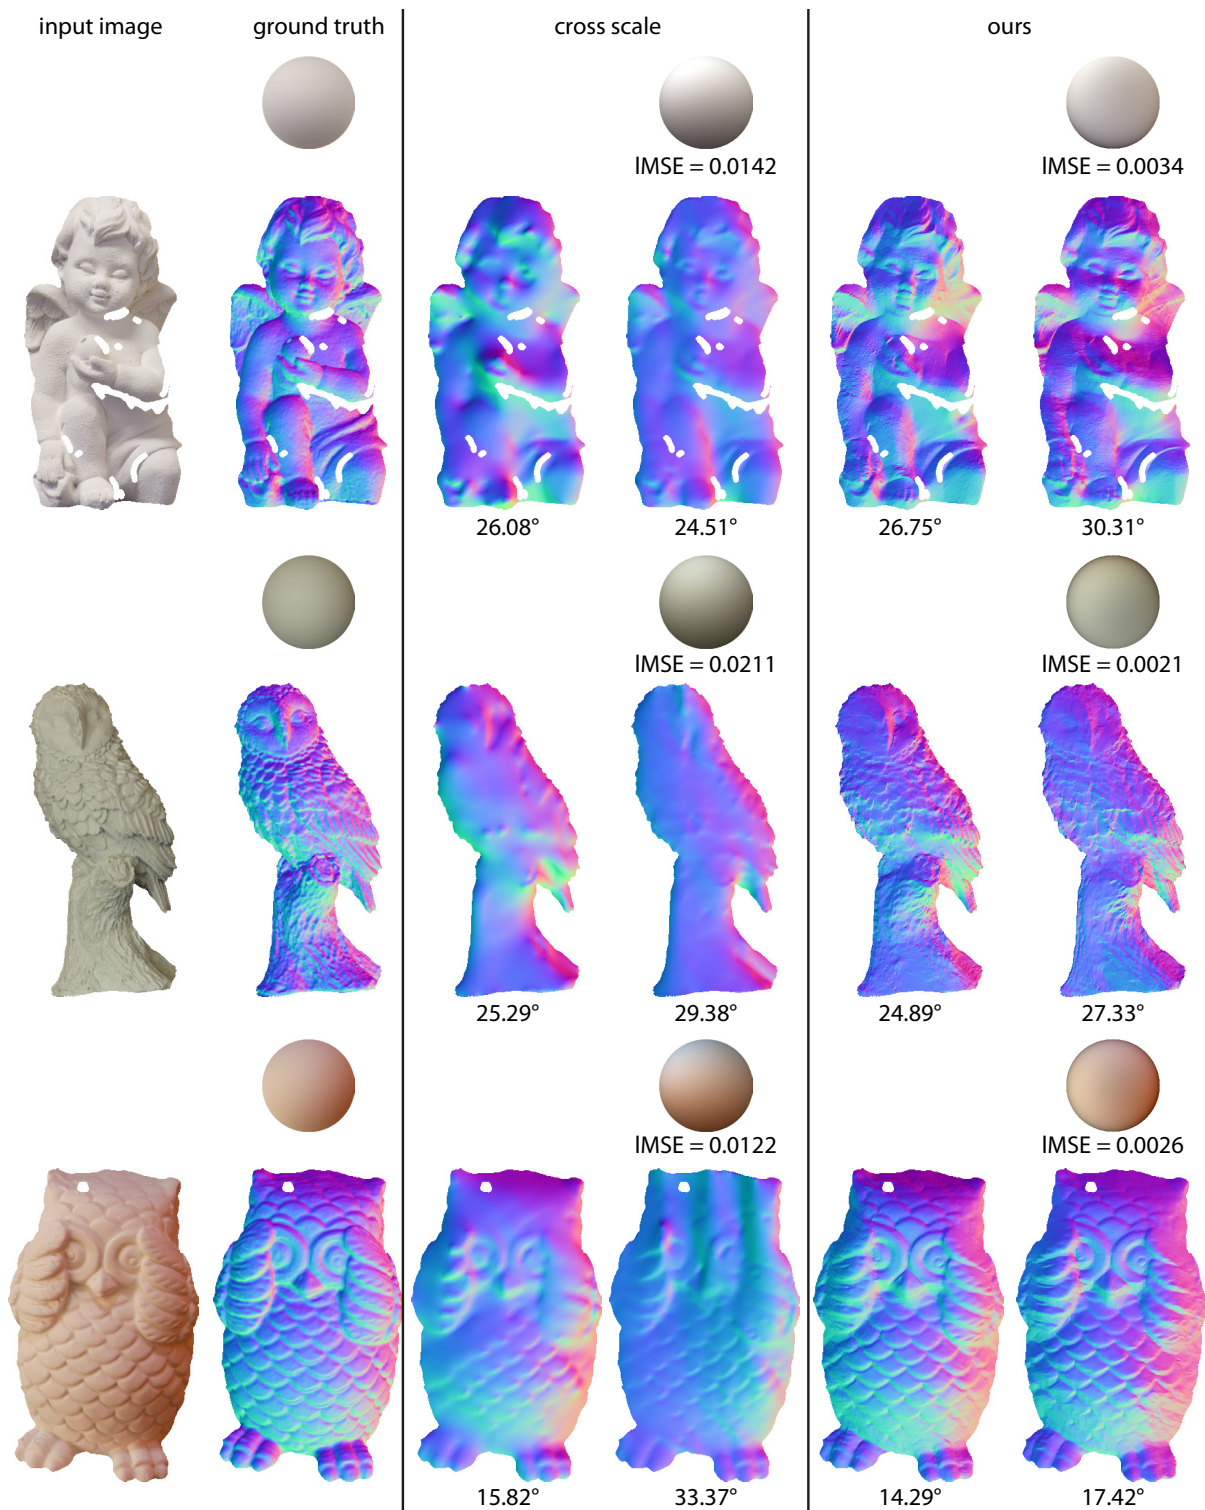


Figure 3: Comparison on real images under natural illumination with median angular errors. We show surface normal estimates for known (left) and unknown illumination (right). In the latter case, the predicted reflectance map with IMSE value is shown above the surface normals. We masked out parts of the surface that were invalid due to the multi-view stereo reconstruction of the ground truth data; these parts should not be evaluated against. These masks were only used in the evaluation and are overlayed on all surface normal maps shown.