

Shape and Spatially-Varying Reflectance Estimation From Virtual Exemplars (Supplementary Material)

Zhuo Hui, *Student Member, IEEE*, and Aswin C. Sankaranarayanan, *Member, IEEE*

Comparisons. Figure 1 showcases the performance of many photometric stereo techniques for two objects: a black-obsidian bunny and a gold-painted elephant. We used 253 input images for each object and recovered the 3D surface by using Poisson reconstruction on the estimated surface normals [1]. Photometric stereo under Lambertian model fails to recover precise normal maps indicating the presence of non-Lambertian components. The ℓ_1 -error model is the variant of classical Lambertian Photometric stereo by using ℓ_1 norm, instead of squared ℓ_2 norm, as the objective function. If the deviations from Lambertian models in the observed images are sparse (for example, when the reflectance is Lambertian+mirror), then this model is expected to perform well. However, ℓ_1 norm minimization fails to produce good results since the materials under consideration have no mirror-like lobes, making the model mismatch not sparse spatially. The robust PCA-based solver [2] produces better normal maps as compared to the traditional photometric stereo, however it produces overly smoothed estimates; this can be attributed to removal of non-Lambertian cues which are invaluable for precise normal estimation. The solution of Alldrin et al. [3] while significantly better than Lambertian technique produces errors greater than 1° . We simulated the performance of example-based photometric stereo which is identical to the proposed technique when applied to a two-material (white-diffuse and chrome) dictionary. As expected, having a larger dictionary of BRDFs as in the proposed technique does provide significant improvements in surface normal estimation. In contrast, coarse-to-fine, returns reliable normal estimates for both scenes indicating the robustness of the underlying solution. Finally, the gradient descent-based refinement of the coarse-to-fine scheme outperforms all the other techniques in both normal and surface reconstruction. The mean angular errors for both test objects can be reduced to nearly 0.1° , which is significantly smaller than the finest sampling used, $\theta_5 = 0.5^\circ$.

Spatially-invariant BRDF. Figure 2 shows the BRDF estimation errors for different materials in the database — rank-ordered from worst-to-best performance — both for the per-pixel as well as the rank-1 BRDF estimation techniques. We can observe that enforcing the rank-1 prior — indicating spatially-invariant BRDF — provides significantly better BRDF estimates since we can pool together information from all pixels. In Figure 3, we show relighted images and the visualization of BRDF slice for a subset of materials. From the performance, we see that the low-rank prior approach performs better. As expected, pooling information from multiple pixels does significantly help in restricting the solution space.

Spatially-varying BRDF. We report BRDF estimation results over an object whose SV-BRDF is piecewise constant. Specifically, we simulate an object (see Figure 4) with 6 distinct regions, each with its own BRDF. We estimate the SV-BRDF using both the per-pixel BRDF approach as well as the low-rank prior approach. For the low-rank approach, we control the rank of the solution by varying the penalty parameter associated with the nuclear norm. Figure 4 showcases the performance of both the per-pixel as well as the low-rank approaches. We present both qualitative, in terms of rendered images, as well as quantitative results, in terms of relative BRDF error, for varying values of the rank of the solution. We observe that the initial per-pixel estimates introduce specular points in the left forehead and cheek. The incorporating low-rank constraint, especially for $K \geq 6$, causes the rendered appearance to closely resemble to the ground truth. Quantitatively, we observe that the relative BRDF error plots reduces significantly when the rank of the solution is greater than $K = 6$. Note that we do not enforce any spatial clustering in our BRDF estimation results.

BRDF estimation. The BRDF estimates are obtained using the surface normals estimated from gradient descent method. In Figure 5, as a means of comparison to the ground truth, we showcase the rendering results for both per-pixel and low-rank approaches. Note that, in Figure 5, there are not just a wide variety of unique materials (the helmet, the breast-plate, the chain, the red scabbard, to name a few) but also significant modeling deviations (inter-reflections, cast-shadows). The per-pixel estimates show the robustness to handle objects with complex spatial variations, but at times produce incorrect rendering results due to the insufficient observations. In contrast, the method incorporating the low-rank prior returns a more faithful rendition of the scene, indicating the advantages gained by pooling the information across multiple pixels.

Comparisons on benchmark dataset. Figure 6 showcases the comparisons on the box-and-whisker plots for the proposed technique with the state-of-the-art methods in [4] and [5]. We adopt the same strategy to demonstrate the performance on the dataset as shown in [6]. We refer the reader to [6] for more comprehensive description of the dataset and the benchmark algorithms. We also show the reconstructed 3D surfaces using Poisson reconstruction for all the scenes from the database.

Analysis on the coefficients estimates. In Figures 7 and 8, we showcase the top few estimates atoms in the dictionary associated with largest coefficient values for 4 test materials. For each material, we simulated 253 lighting directions across 100 surface normals. The coefficients for each color channel were separately solved. While the specular lobes of the selected atoms closely resemble to the test material, we observe that atoms with large coefficients vary across color channels. Further, in spite of obtaining accurate BRDF estimates, the individual coefficients can be quite different since the atoms of the dictionary are highly correlated. For such correlated dictionaries, small perturbation of the output BRDF can lead to large differences in the support of the sparse coefficient vector. Hence, it is more reliable and accurate to compare the estimates to the ground truth in the BRDF space (as shown on the left of the plot) rather than the coefficients space.

REFERENCES

- [1] A. Agrawal, R. Raskar, and R. Chellappa, "What is the range of surface reconstructions from a gradient field?" in *ECCV*, 2006.
- [2] L. Wu, A. Ganesh, B. Shi, Y. Matsushita, Y. Wang, and Y. Ma, "Robust photometric stereo via low-rank matrix completion and recovery," in *ACCV*, 2011.
- [3] N. Aldrin, T. Zickler, and D. Kriegman, "Photometric stereo with non-parametric and spatially-varying reflectance," in *CVPR*, 2008.
- [4] B. Shi, P. Tan, Y. Matsushita, and K. Ikeuchi, "Bi-polynomial modeling of low-frequency reflectances," *PAMI*, vol. 36, no. 6, pp. 1078–1091, 2014.
- [5] S. Ikehata and K. Aizawa, "Photometric stereo using constrained bivariate regression for general isotropic surfaces," in *CVPR*, 2014.
- [6] B. Shi, Z. Wu, Z. Mo, D. Duan, and S.-K. Y. P. Tan, "A benchmark dataset and evaluation for non-lambertian and uncalibrated photometric stereo," in *CVPR*, 2016.
- [7] A. Hertzmann and S. Seitz, "Example-based photometric stereo: Shape reconstruction with general, varying BRDFs," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 27, no. 8, pp. 1254–1264, 2005.
- [8] W. Matusik, H. Pfister, M. Brand, and L. McMillan, "A data-driven reflectance model," *ACM Trans. Graphics*, vol. 22, no. 3, pp. 759–769, 2003.

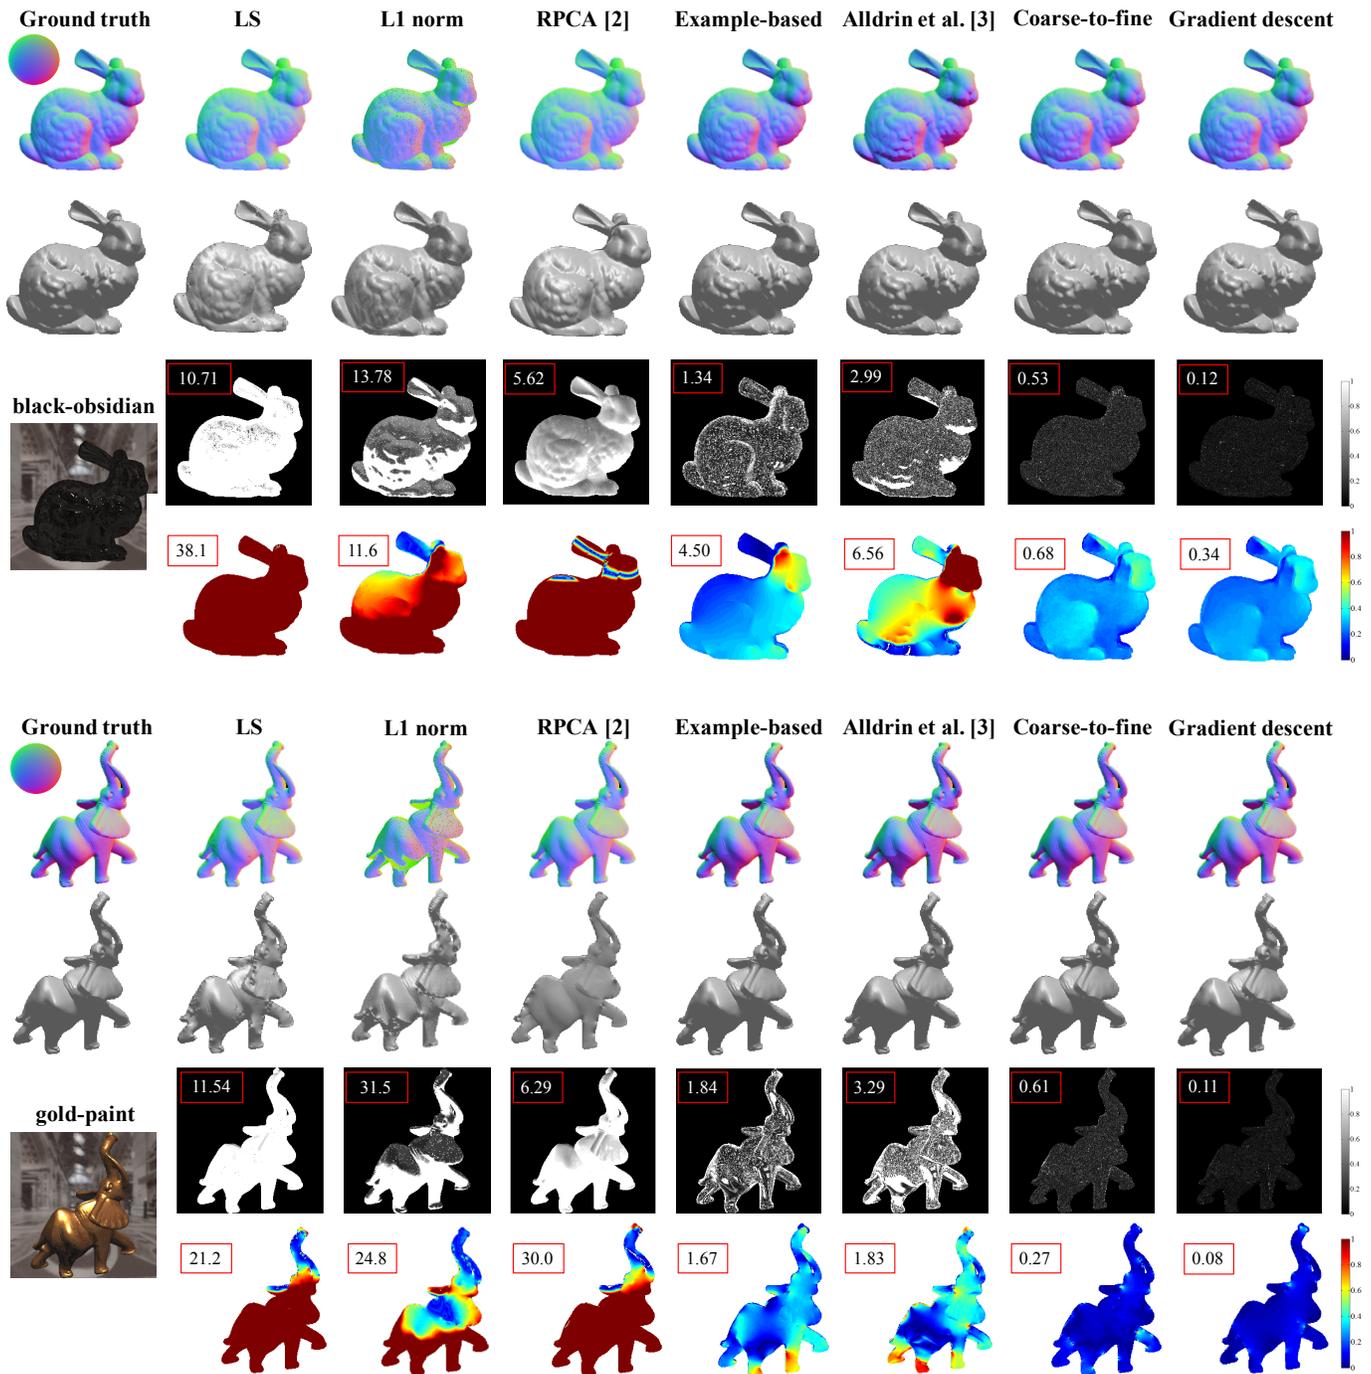


Fig. 1: **Normal estimation across algorithms.** We compare the performance of photometric stereo under Lambertian model (LS), an ℓ_1 -error model (L1), robust PCA-based approach [2] (RPCA), simulated example-based [7], Alldrin et al. [3], coarse-to-fine search, and gradient descent scheme on two objects using 253 images each. Shown are (top-bottom) the estimated surface normals, recovered 3D surface by using Poisson reconstruction, angular error in normal estimation in degrees and relative error in depth map based on different approaches. The insets in rows 3 and 4 for both objects are the average normal errors in degrees and the relative depth errors.

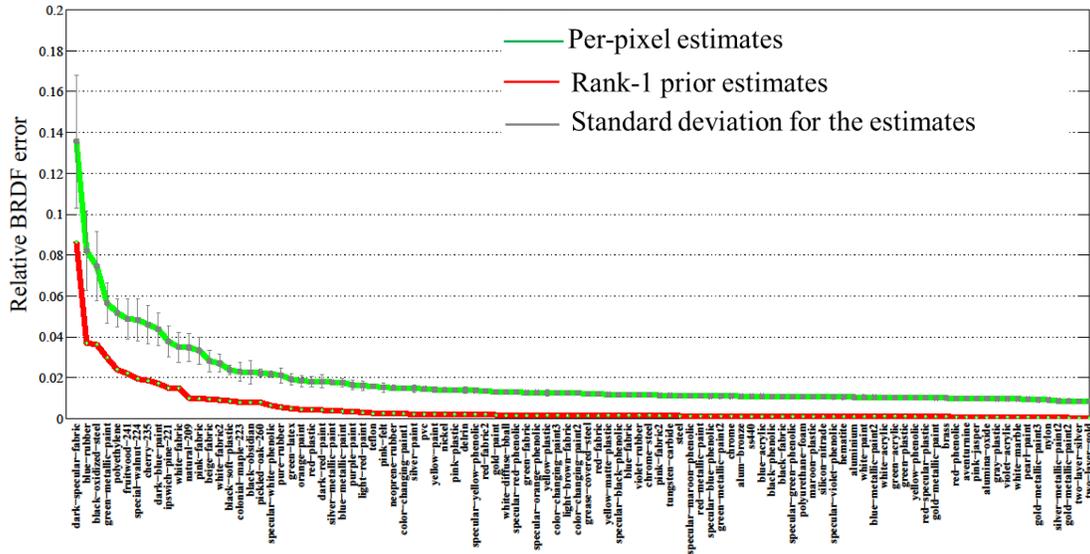


Fig. 2: **Quantitative BRDF evaluation.** Given 253 lighting directions, we evaluate accuracy of BRDF estimation across different materials. For each material, we generated 100 normals with random orientations and estimated the BRDF for each instance individually (per-pixel) as well as utilizing low-rank prior (rank-1). For the per-pixel estimates, we plot average and standard deviation of the errors.

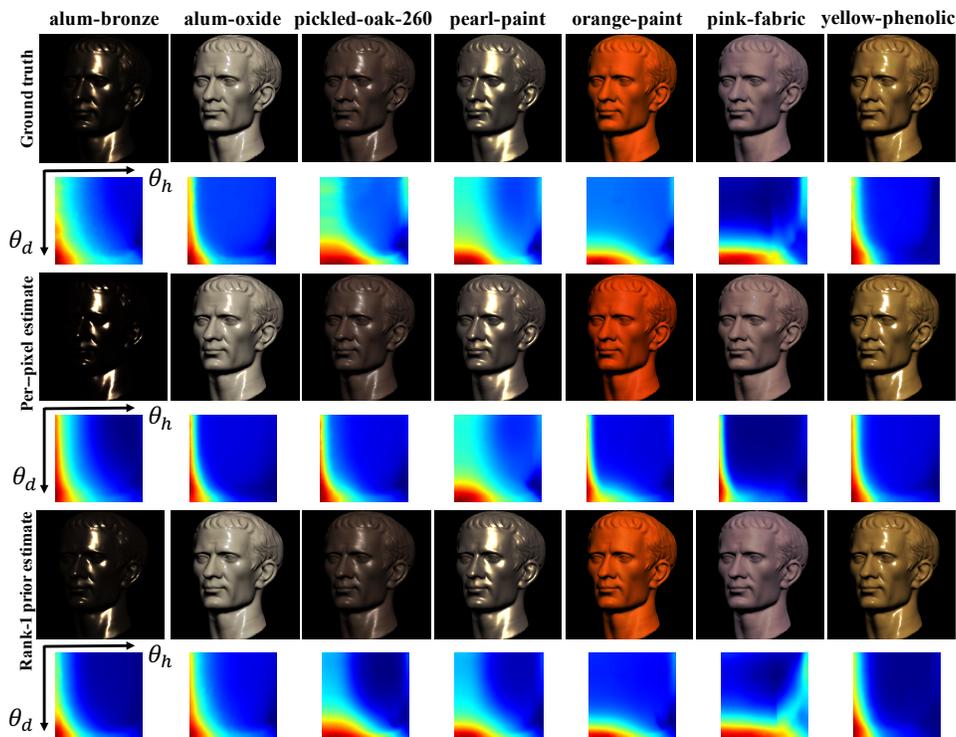


Fig. 3: **Qualitative BRDF evaluation.** Shown are rendered BRDF for the Caesar statue for a few select materials from MERL database [8]. (row 1) The rendered image based on ground truth BRDF; (rows 3 and 5) rendered images based on estimated averaged BRDF from per-pixel estimates and incorporating the low rank prior (rank-1) for 100 randomly generated normals, respectively; (rows 2, 4 and 6) the BRDF slices for $\phi_d = 90^\circ$. Note that how the low rank prior corrects the noisy specular peak (as shown in pink-fabric and orange-paint) by widening the lobe.

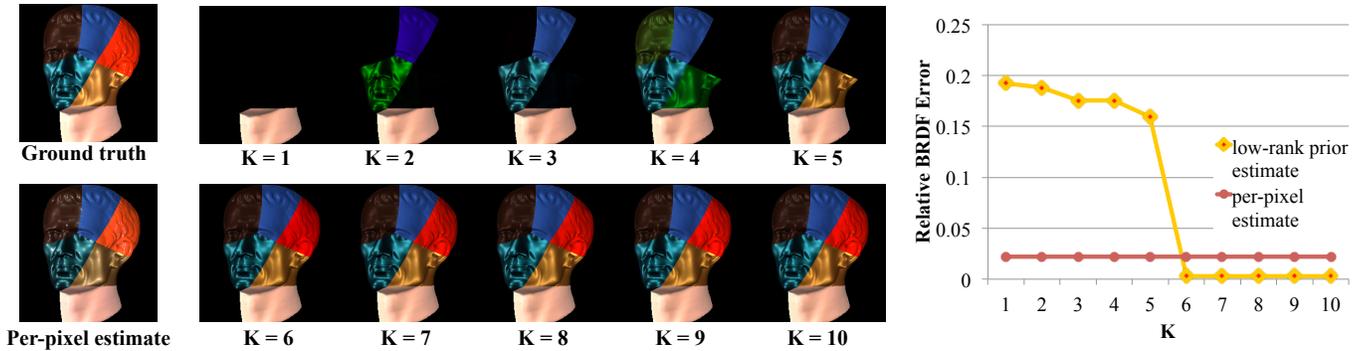


Fig. 4: **BRDF evaluation with low rank prior.** Shown are rendered images using the ground truth, the per-pixel estimate, and the low-rank prior estimate for different values of the rank of the solution, K . We also present a plot of the relative BRDF error as a function of the rank, K .

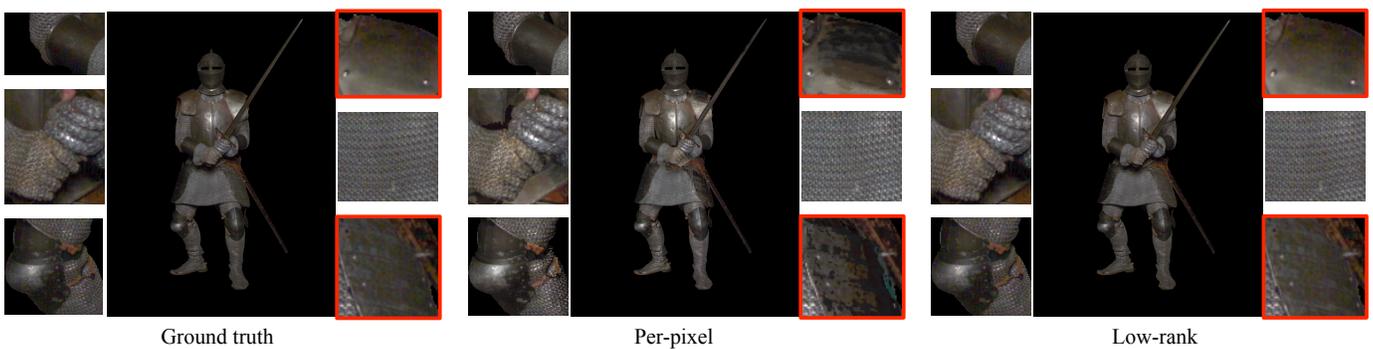


Fig. 5: **Relighting results on knight dataset.** Shown are the rendering results under a point light source for (from left to right) a ground truth image, per-pixel BRDF estimates, and estimates with the low-rank prior. We also show the close-up appearance for the relighting results.

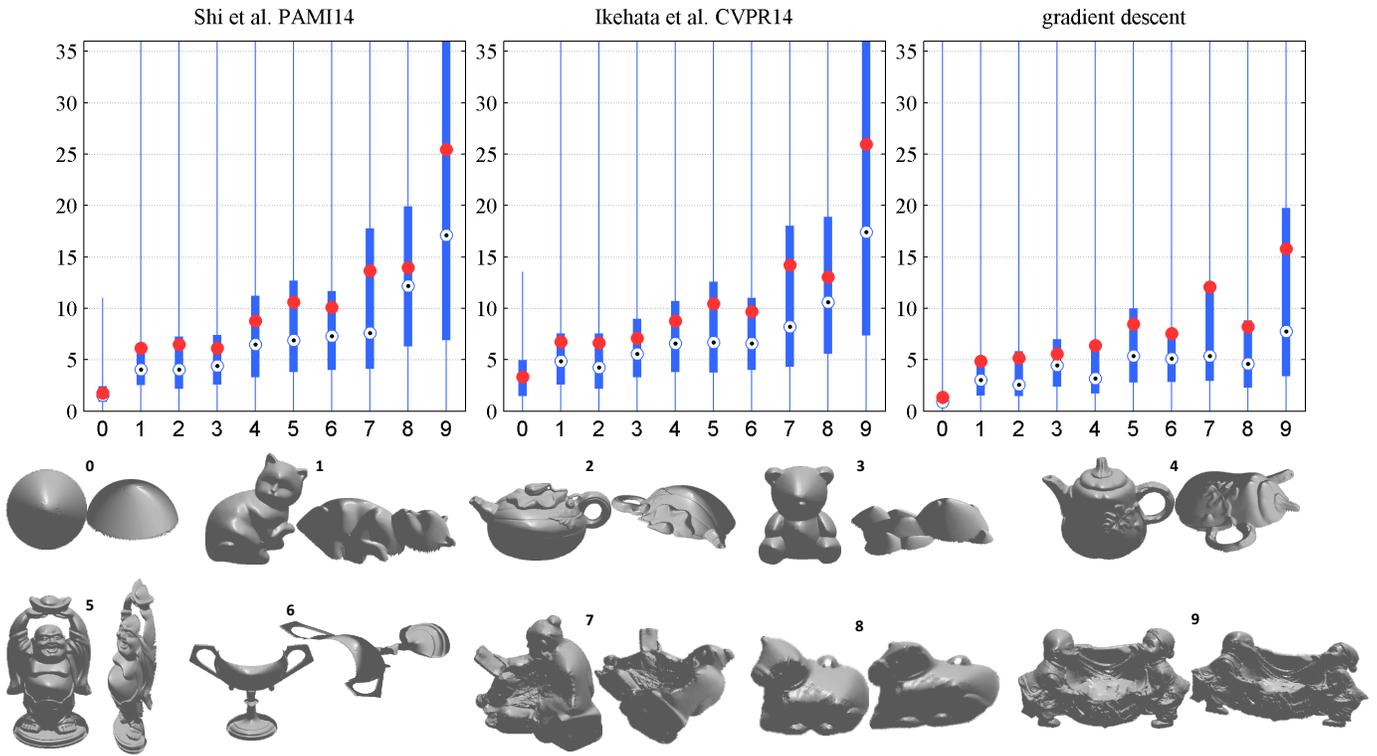


Fig. 6: **Evaluations on the benchmark photometric stereo dataset [6].** We compare the box-and-whisker plots for the proposed technique with the methods in [4] and [5]. We adopt the same notation fromas in [6] for these plots. In particular, the median and mean angular errors for each object (index number onshown as the horizontal axisindex) are denoted as the red and blue points, respectively. The length of the bar indicates the range between the first to the quartile values for the angular errors. We also show 3D surfaces recovered using Poisson reconstruction on the estimated surface normals of the gradient descent technique. In all, the proposed gradient descent technique has smaller angular errors as measured in mean, median, as well as maximum error.

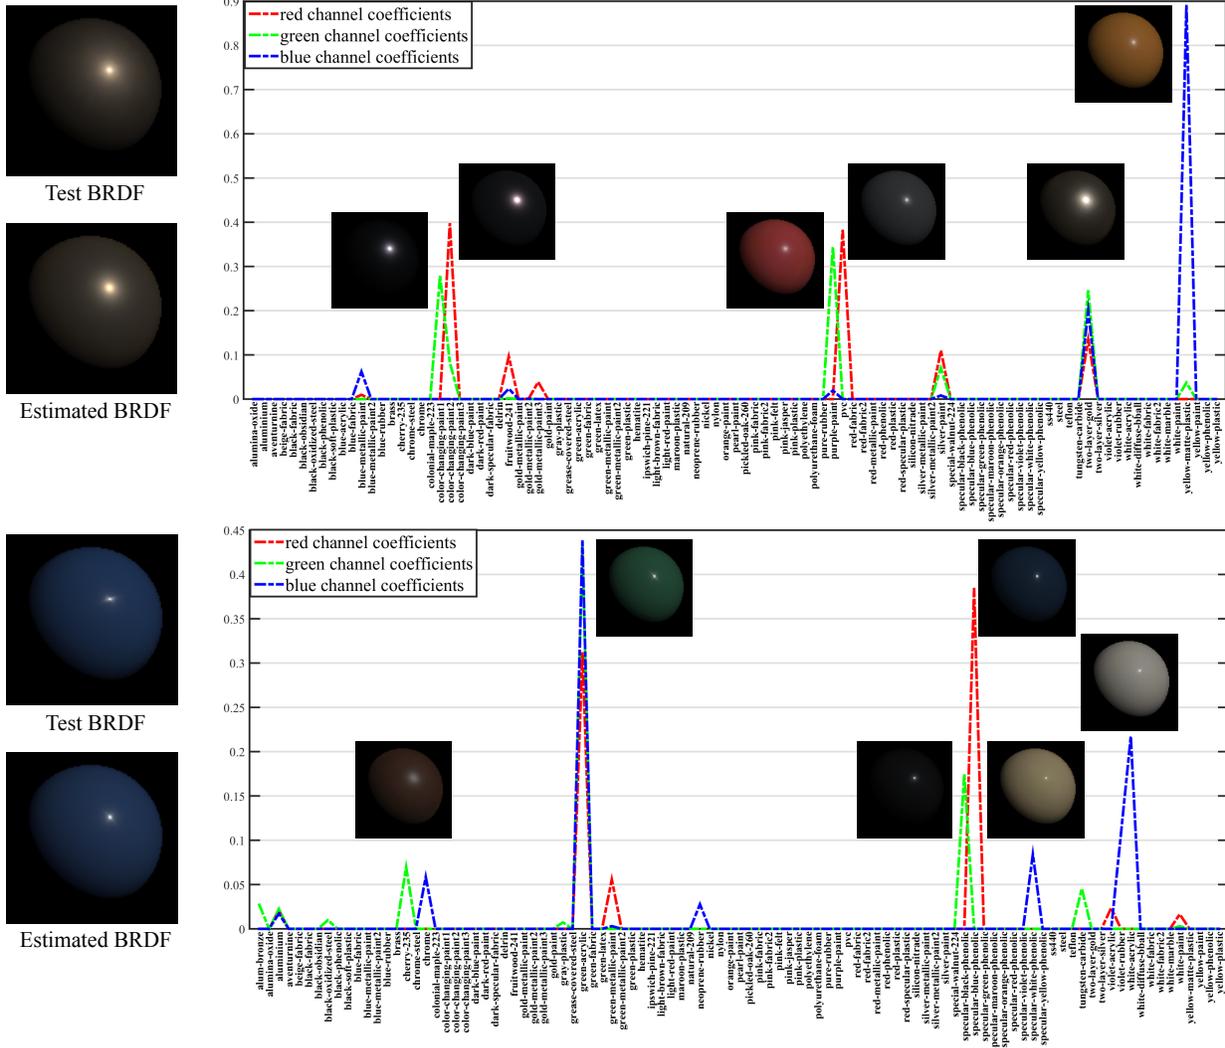


Fig. 7: **Analysis on the coefficients estimates.** We show the test and estimated material BRDFs as well as the atoms with largest coefficients value in each color channel (RGB). For each material, we plot the coefficients against the materials in the dictionary. The insets in the plots are the corresponding material BRDFs shown on the sphere. PI also refer to Figure 8 for more results.

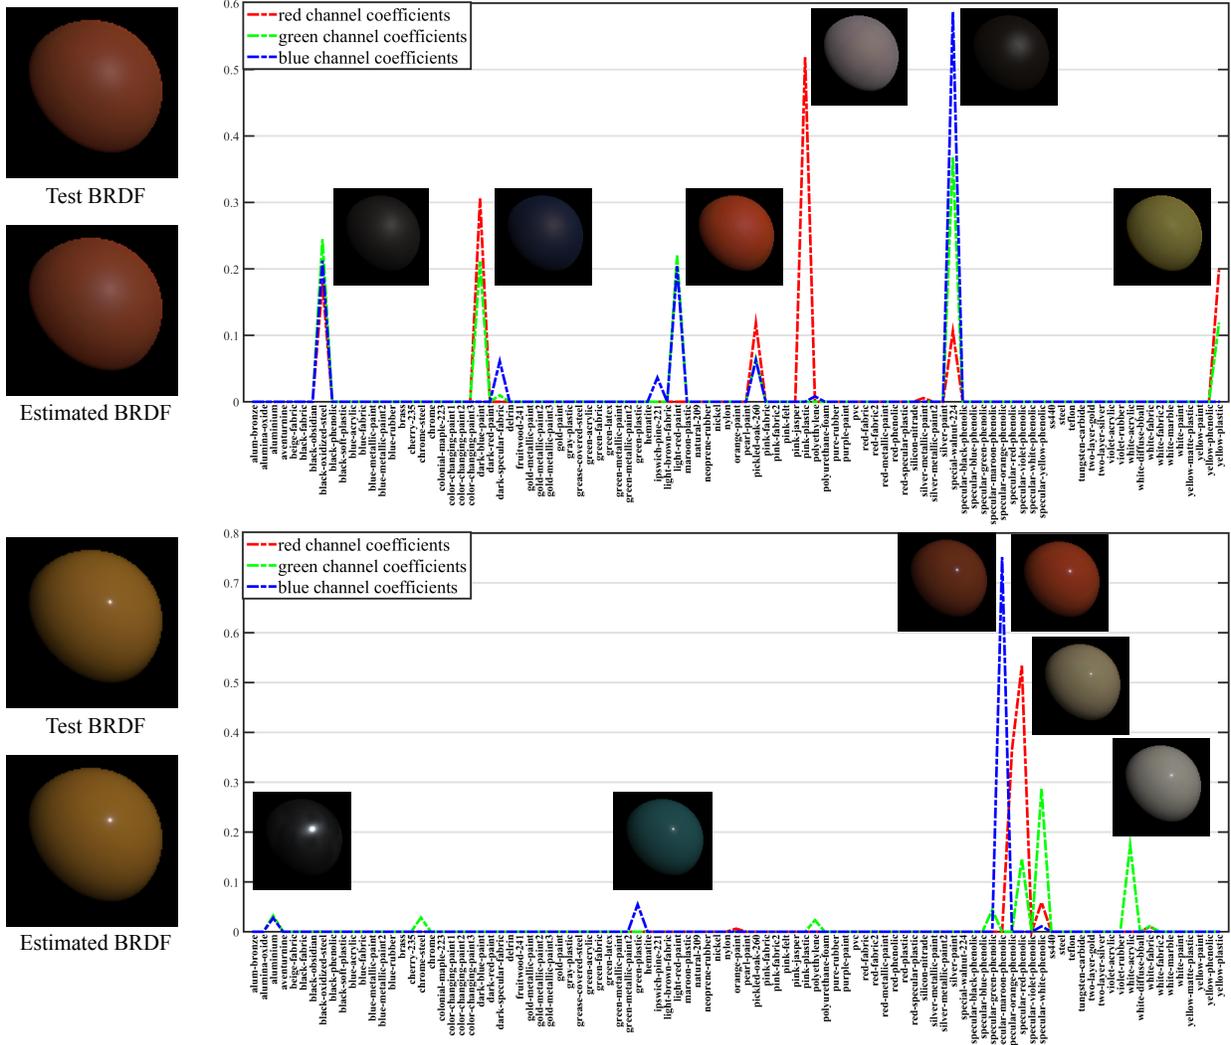


Fig. 8: **Analysis on the coefficients estimates.** We show the test and estimated material BRDFs as well as the atoms with largest coefficients value in each color channel (RGB). For each material, we plot the coefficients against the materials in the dictionary. The insets in the plots are the corresponding material BRDFs shown on the sphere.