

Reflectance Capture using Univariate Sampling of BRDFs (Supplementary Material)

Zhuo Hui,[†] Kalyan Sunkavalli,[‡] Joon-Young Lee,[‡] Sunil Hadap,[‡]
Jian Wang,[†] and Aswin C. Sankaranarayanan[†]
[†]Carnegie Mellon University [‡]Adobe Research

List of Figures

Figures 1 and 2: Identifiability analysis for Blinn-Phong and Isotropic Ward models.

Figure 3: Performance of surface normals and BRDF recovery for varying number of images.

Figure 4: Validation of BRDF reconstruction for different θ_d .

Figure 5: Performance of BRDF recovery for varying baseline between camera and flash, i.e., varying θ_d .

Figures 6 and 7: Visualization with light cascade for recovered BRDFs and ground truth materials in MERL database.

Figures 8 and 9: Visual comparisons of BRDF estimation using univariate sampling against state-of-the-art BRDF estimation/sampling methods on MERL database.

Figure 10: Example inputs for synthetic validation experiments.

Figure 11: Surface normal estimates on synthetic data.

Figure 12: Comparisons against Riviere et al. [4] on synthetic data.

Figures 13: Visual comparisons against ground truth images on synthetic data.

Figure 14: Visualization with hardware prototype for PointGrey camera.

Figure 15: Reflectance and shape estimation on real data captured with PointGrey.

Figure 16: Comparison of results obtained from PointGrey and iPhone 6S cameras.

Figure 17: Material trait analysis on real data captured with both PointGrey camera and iPhone 6s.

1. Identifiability under univariate sampling

In the main paper, we showed that parameters associated with the Cook-Torrance model can be identified from noiseless univariate samples. We provide similar results for other popular parametric BRDF models including the Blinn-Phong model, and the Isotropic Ward model and is extendable to the Ashikhmin-Shirley model.

Similar to the proof for the Cook-Torrance model, outlined in the main paper, the derivations underlying these results adhere to the following steps:

- First, we derive an expression for the BRDF purely as a function of θ_h , and with $\theta_d = \phi_d = 0$. Note that isotropic BRDFs are invariant to changes in ϕ_h .
- Second, since all models have a Lambertian diffuse term and a specular term, we show that the parameters of the Lambertian term can be identified by setting $\theta_h = \pi/2$.
- Third, we show that the parameters defining the specular lobe can be identified by observing the BRDF at few different values of θ_h .

Blinn-Phong model. BRDF measurements under the Blinn-Phong model and univariate sampling (i.e., $\theta_d = \phi_d = 0$) can be written as:

$$\rho(\theta_h) = \rho_d + \rho_s \frac{\beta + 2}{2\pi} \cos^\beta \theta_h,$$

where ρ_d, ρ_s and β are the parameters defining the model. Given ρ_d , we can write

$$\log(\rho(\theta_h) - \rho_d) = \log \rho_s + \log(\beta + 2) - \log(2\pi) + \beta \log \cos \theta_h.$$

Hence, if we plot $\log(\rho(\theta_h) - \rho_d)$ as a function $\log \cos \theta_h$, then the resulting plot is a straight line whose slope is β and intercept is $\log \rho_s + \log(\beta + 2) - \log(2\pi)$. Hence, we can recover all three parameters of the model.

In Figure 1, we show the estimated parameters via univariate measurements and compare against the ground truth Blinn-Phong model.

Isotropic Ward model. BRDF measurements under the Isotropic Ward model and univariate sampling (i.e., $\theta_d = \phi_d = 0$) can be written as:

$$\rho(\theta_h) = \rho_d + \frac{\rho_s}{\cos \theta_h} \frac{\exp(-\tan^2 \theta_h / \beta^2)}{4\pi\beta^2},$$

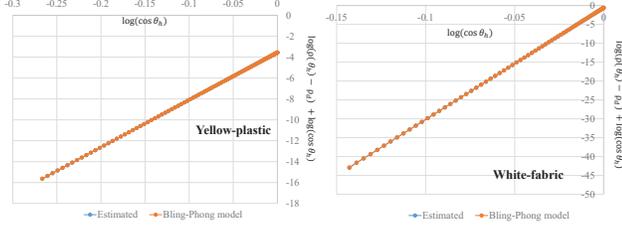


Figure 1. Shown are the plots for the value of $\log(\rho(\theta_d) - \rho_d)$ for different value of $\log(\cos(\theta_h))$ of two materials in MERL database. We compare the results against the ground truth Bling-Phong model. As can be seen here, the parameters can be determined from univariate measurements and very close to the ground truth parameters.

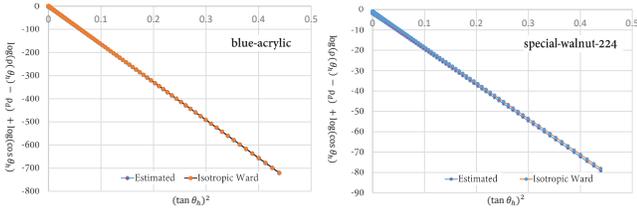


Figure 2. Shown are the plots for the value of $\log(\rho(\theta_d) - \rho_d) + \log(\cos \theta_h)$ for different value of $\tan^2 \theta_h$ of two materials in MERL database. We compare the results against the ground truth isotropic Ward model. As can be seen here, the parameters can be determined from univariate measurements and very close to the ground truth parameters.

where ρ_d, ρ_s and β are the parameters defining the model. First, we observe that $\rho_d = \rho(\pi/2)$. Second, given ρ_d , we can write

$$\log(\rho(\theta_h) - \rho_d) = \log\left(\frac{\rho_s}{4\pi\beta^2}\right) - \log(\cos \theta_h) - \frac{\tan^2 \theta_h}{\beta^2}.$$

Equivalently,

$$\log(\rho(\theta_h) - \rho_d) + \log(\cos \theta_h) = \log\left(\frac{\rho_s}{4\pi\beta^2}\right) - \frac{\tan^2 \theta_h}{\beta^2}.$$

If we plot the LHS expression as a function of $\tan^2(\theta_h)$, then we expect a straight line whose slope is $1/\beta^2$ and intercept is $\log \rho_s - \log(4\pi\beta^2)$, from which we can identify both β and ρ_s . More specifically, we can identify these parameters from values of $\rho(\theta_h)$ at two distinct values of θ_h .

In Figure 2, we show the estimated parameters via univariate measurements and compare against the ground truth Isotropic Ward model.

The proof for the Ashikhmin-Shirley model follows very closely the one we described above for the Cook-Torrance model.

2. Performance on MERL database

In this section, we empirically validate the performance of our dictionary-based method for BRDF reconstruction

from univariate samples on the MERL database. We start by analyzing its performance vis-a-vis the ground truth data, and then compare our reconstructions against those from other state-of-the-art BRDF modeling/sampling algorithms.

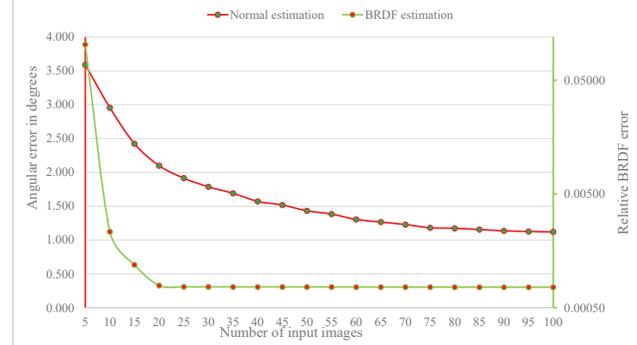


Figure 3. We estimate the angular errors as well as the relative BRDF errors for our proposed method on the MERL database. The plots were obtained by averaging across all 100 BRDFs in the MERL database and 100 randomly-generated normals per material.

Varying number of input images. Figure 3 characterizes the errors in surface normal as well as the BRDF estimation for varying number of input images, or equivalently, the number of samples for the half angles. For the normal estimation, we report the average angular error for the univariate sampling. The average angular error is computed by randomly generating 100 normals per material and varying across all 100 material BRDFs in the database. To characterize the performance for BRDF estimation, we assume the knowledge of true surface normal and report the average relative BRDF error for varying number of images. As noted from Figure 3, both the angular errors and relative BRDF errors degrade gracefully with a smaller number of images.

BRDF estimation against ground truth. In this experiment, we evaluate how accurately we can reconstruct a BRDF given only univariate samples using a dictionary-based prior. We assume a collocated setup, i.e. $\theta_d = 0^\circ$, set $\phi_h, \phi_d = 0^\circ$, and sample a chosen MERL BRDF at different values of θ_h to obtain the univariate measurements. In Figure 4, we visualize the original BRDF (parameterized by (θ_d, θ_h)) as well as 3 different 1-D slices corresponding to $\theta_d = (0^\circ, 15^\circ, 30^\circ)$. We then reconstruct the BRDF using a dictionary composed of the remaining 99 BRDFs in the MERL database, using only the samples corresponding to $\theta_d = 0^\circ$. Figure 4 also visualizes the reconstructed BRDFs, as well as 1-D slices of these reconstructions at $\theta_d = (0^\circ, 15^\circ, 30^\circ)$. As can be seen here, in spite of estimating the BRDF from only one slice of samples, the reconstructed results closely match the ground truth BRDFs even

at other angles as indicated by the measured RMSE values.

Varying baseline between camera and flash. Smartphones and tablets have different relative positioning of their cameras and flash units. Hence, it is important to characterize the stability of our reflectance estimation technique for varying baseline. When the object is at an approximately fixed distance from the device, changes in baseline can be modeled as changes in θ_d . Figure 5 showcases this by characterizing BRDF estimation errors as a function of θ_d . For each θ_d , we compute average error over 100 materials in MERL database. Note that the performance remains stable for θ_d less than 65 degree; this is sufficient to capture the operating scenario underlying a wide range of mobile devices. Beyond 65 degrees, the univariate sampling completely misses the specular lobe which results in poor performance in estimating the BRDF.

In addition, we ability of our method to handle a wide range of materials, by rendering spheres lit by point lights for different values of θ_d (see Figures 6 and 7). Renderings with our reconstructed BRDFs closely match those produced using ground truth BRDF data, demonstrating that our method is able to produce realistic results for a majority of isotropic materials in the MERL database.

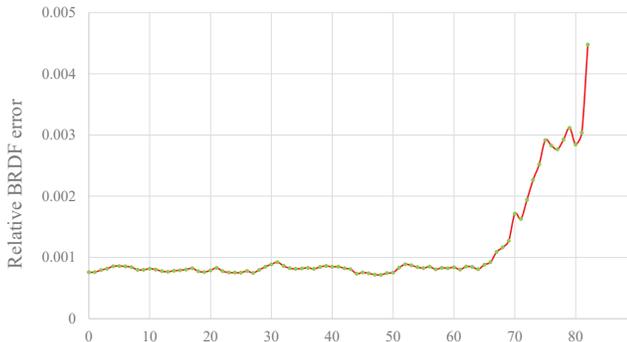


Figure 5. We plot the approximation accuracy in terms of relative BRDF errors by varying θ_d in degrees. For each θ_d , we compute average error over 100 materials in MERL database.

Comparisons. In this section, we compare our BRDF estimation method against different ways to approximate raw BRDF measurements – the bivariate model [5], Isotropic Ward, Cook-Torrance, as well as different sampling strategies – full isotropic sampling [1], bivariate sampling, and optimal sampling [6]. We follow the method of [3] to fit the parameters of the Cook-Torrance and Ward models to the full raw BRDF measurements. To compare against different sampling strategies, we use the same dictionary-based method to reconstruct BRDFs from different sets of BRDF samples. For Hui et al. [1], we randomly sample the full isotropic BRDF space for 100 samples of lighting/view

pairs. We also sample the bivariate BRDF space spanned by θ_h and θ_d with the same number of lighting/view combinations. To compare against Xu et al. [6], we sample the BRDF with 20 entries as indicated in their paper. Finally, we collect 100 univariate BRDF samples with collocated lighting and view directions for our method.

For evaluation, we isolate the testing material while leaving the remaining 99 materials in the dictionary. The visualization of the materials in MERL database for all these methods are shown in Figures 8 and 9, where we observe once again that univariate sampling is quite competitive to state-of-the-art models.

3. Performance on synthetic data

We evaluate our technique on synthetic data rendered using the Mitsuba physically-based renderer [2]. We simulate the same optical setup as our real capture prototype. In particular, we specify a perspective camera and a point light source collocated at the camera. We apply an arbitrary spatially-varying BRDF to a near-planar surface with a bump map. We also place checkerboard patterns in the corners of this sample, as is done in our real experiments. We render this scene under different viewpoints (and corresponding light directions) and create a set of 65 input images. Some example images from this dataset are shown in Figure 10.

We process this data using our reflectance and shape estimation method. The checkerboards are used to calibrate the camera and lights. Finally we use our dictionary-based method to estimate surface normals and BRDFs. We evaluate the accuracy of these estimates as well as compare novel renderings of our reconstructions to ground truth photographs.

Normal estimation. Figure 11 demonstrates how our surface normal estimates evolve over the course of 10 iterations (starting from a flat surface at the 0-th iteration). As can be seen here, the average angular error successively reduces and we get highly accurate normal estimates even for complex spatially-varying materials. Figure 12 compares our reconstructions against those from the method of Riviere et al. [4] on this synthetic data. For both the methods, we show the estimated surface normals as well as the integrated depth map against the ground truth depth. As demonstrated in Figure 12, our technique recovers more accurate gross shape as well as more fine scale structures than [4].

Reflectance estimation. In Figure 13, we characterize our reflectance and shape estimation on the synthetic data, by comparing renderings of our normal/BRDF estimates under two novel lighting and view directions against ground truth renderings. While our estimation is done on images

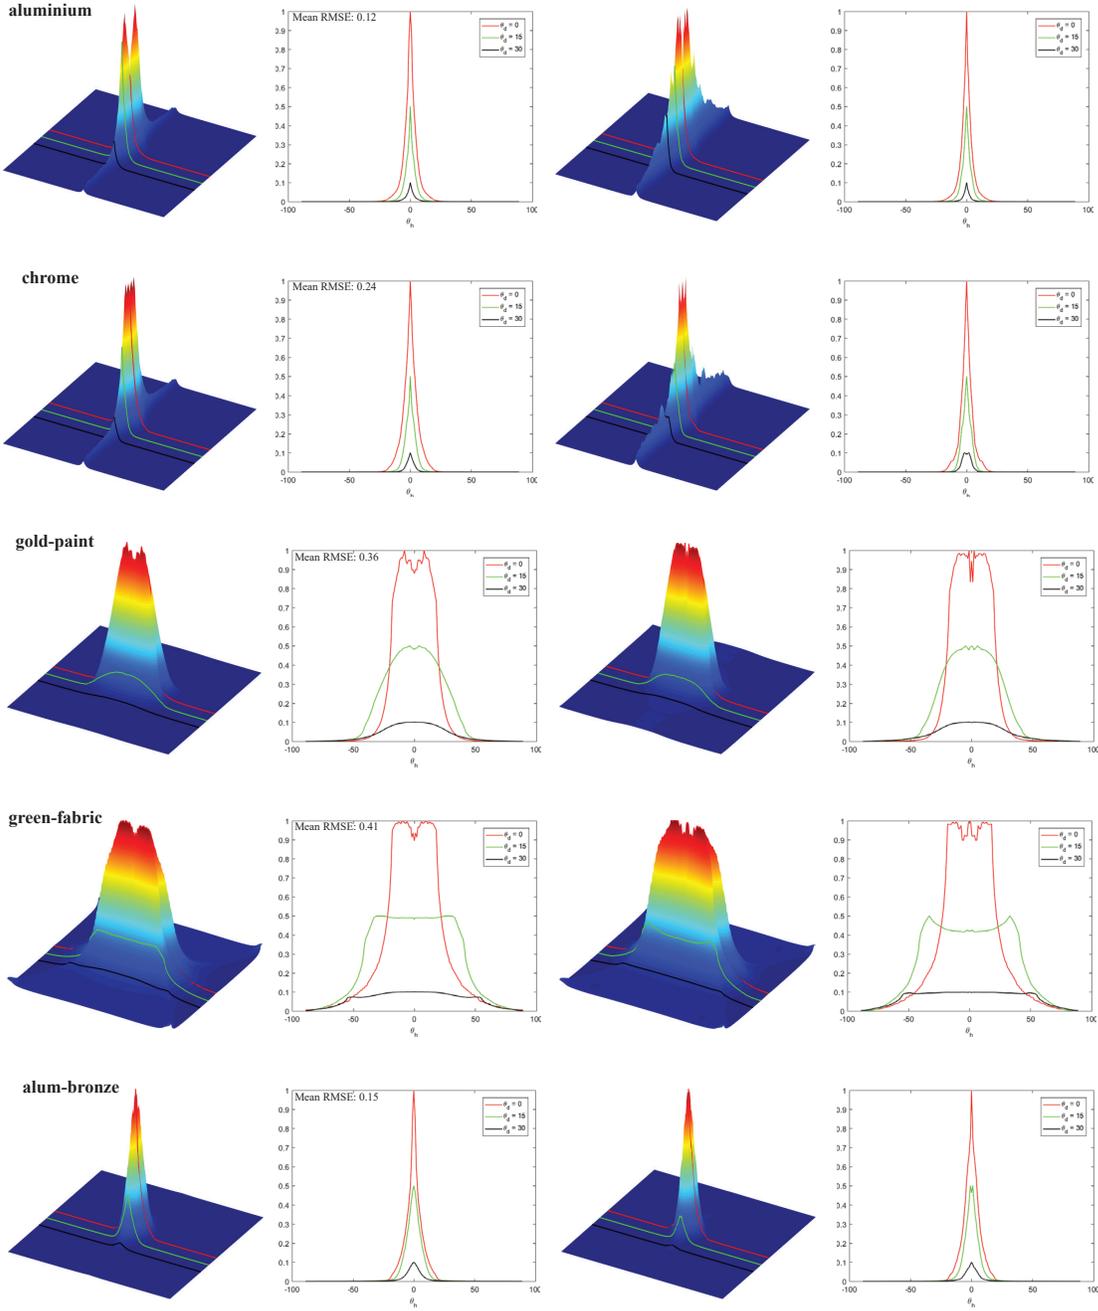
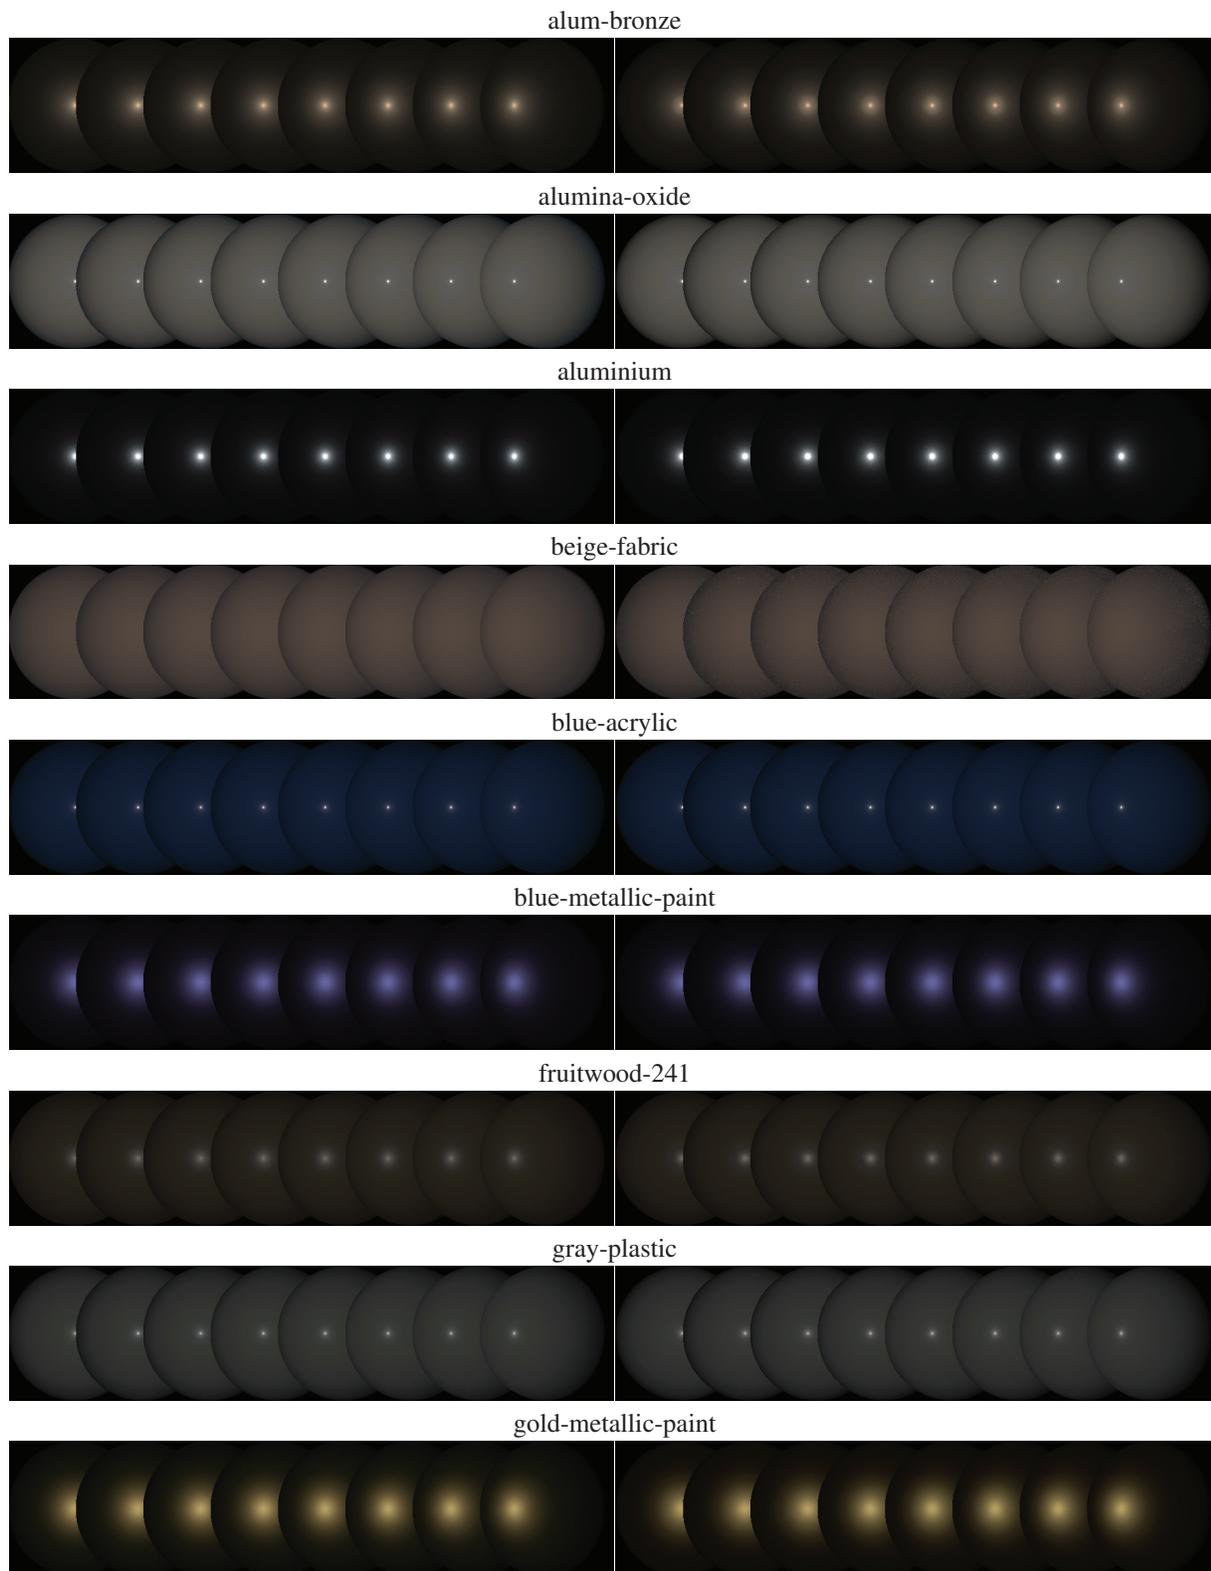


Figure 4. Evaluation on the MERL database. On the left, we visualize ground truth MERL BRDFs on the 2-D plane parameterized by (θ_d, θ_h) (for $\phi_d = 0^\circ$), as well as 3 1-D slices corresponding to $\theta_d = (0^\circ, 15^\circ, 30^\circ)$. We reconstruct the BRDF using only the collocated univariate samples, i.e., $\theta_d = 0^\circ$, and visualize it on the right. As can be seen here, by using the univariate measurements only, we can reconstruct the BRDF at other θ_d quite accurately, as indicated by the mean RMSE shown in the top-left of each plot. We normalize the BRDF ρ by $\hat{\rho} = \frac{\rho - \min(\rho)}{\max(\rho) - \min(\rho)}$ and plot the curve with different θ_d .

captured with a collocated setup, we are evaluating these renderings on non-collocated view/light pairs. This allows us to test how accurate our method is at reproducing the appearance of samples it has not seen. As can be seen from this

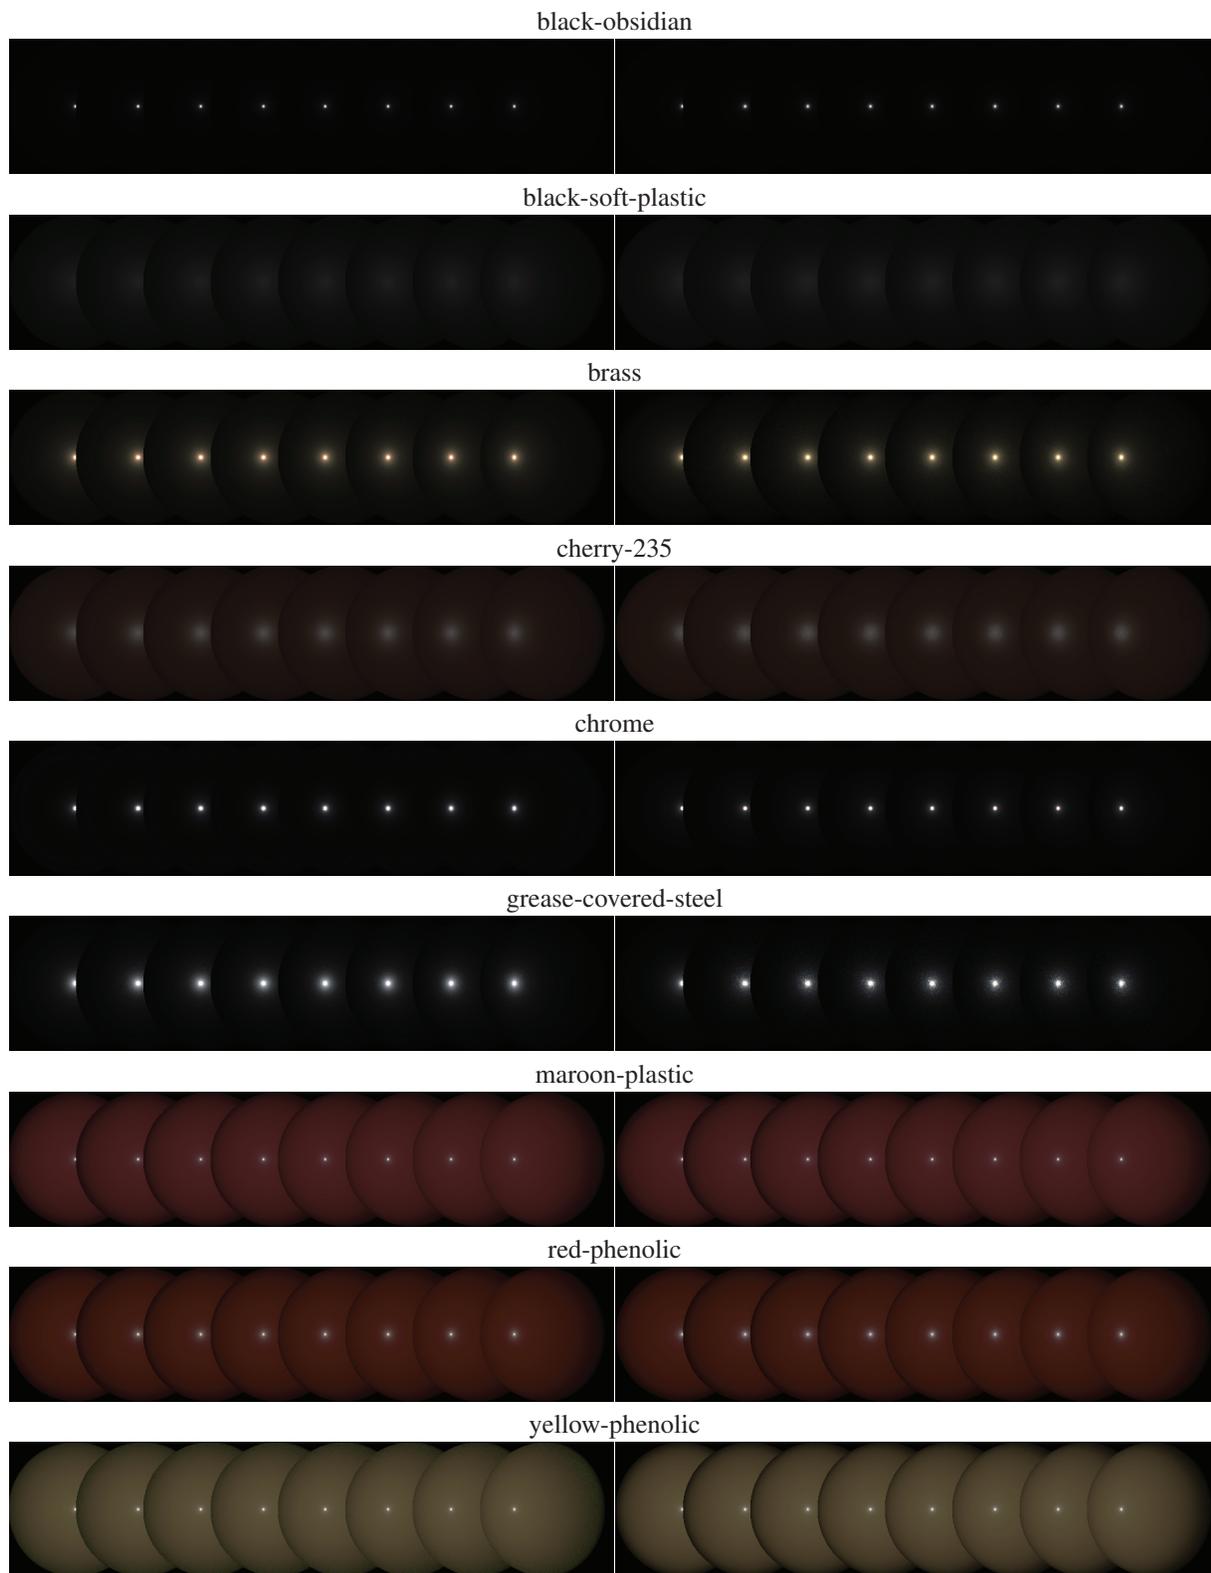
figure, our results closely resemble images rendered with the ground truth normals and BRDFs, indicating the robustness and accuracy of our method.



Estimated results

Ground truth

Figure 6. Evaluation on the MERL database. We visualize the ground truth MERL BRDF data and our univariate sample-based reconstruction by rendering these BRDFs on spheres lit by point light sources for different values of θ_d . Our rendered materials are visually almost identical to the ground truth BRDFs, indicating the accuracy of our reconstruction.

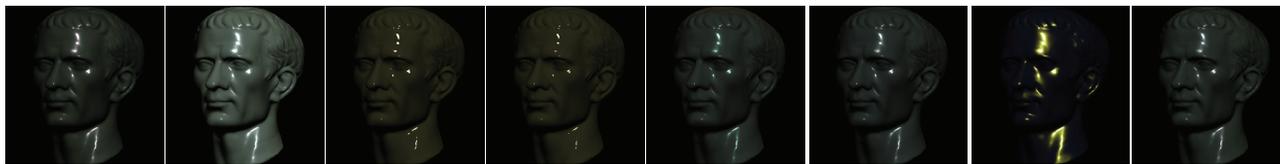


Estimated results

Ground truth

Figure 7. Evaluation on the MERL database. We visualize the ground truth MERL BRDF data and our univariate sample-based reconstruction by rendering these BRDFs on spheres lit by point light sources for different values of θ_d . Our rendered materials are visually almost identical to the ground truth BRDFs, indicating the accuracy of our reconstruction.

aventurine



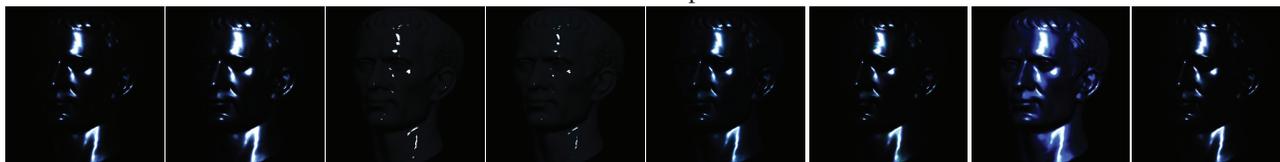
black-oxidized-steel



blue-metallic-paint



blue-metallic-paint2



brass



colonial-maple-223



gold-paint



green-latex

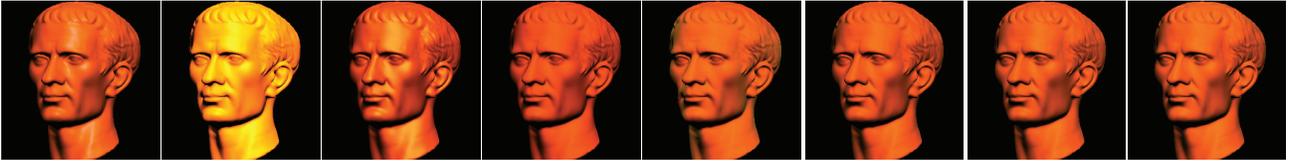


Hui et al. 2015 Bivariate model Cook-Torrance Isotropic Ward Bivariate sample Univariate samples Xu et al. 2016 Ground truth
Figure 8. We compare the visual performance for the state-of-the-art BRDF estimation methods for materials in MERL database.

green-metallic-paint



orange-paint



pearl-paint



silver-paint



pickled-oak-260



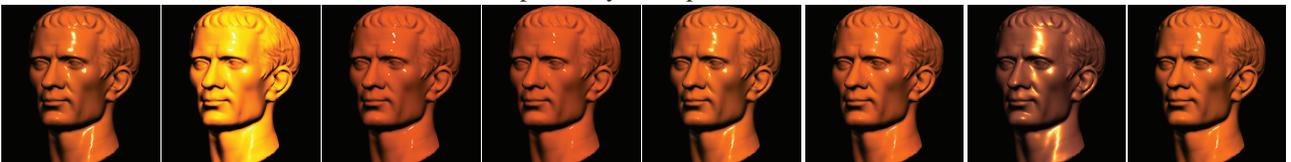
yellow-plastic



teflon



specular-yellow-phenolic



Hui et al. 2015 Bivariate model Cook-Torrance Isotropic Ward Bivariate sample Univariate samples Xu et al. 2016 Ground truth
Figure 9. We compare the visual performance for the state-of-the-art BRDF estimation methods for materials in MERL database.



Figure 10. Example images from our synthetic dataset.

4. Performance on real data

We evaluate our method on real data captured using a hardware prototype consisting of a PointGrey Grasshopper 3 with a light pipe. Our prototype acquisition hardware consists of a PointGrey Grasshopper 3 GS3-U3-23S6C-C camera and a light pipe. During image capture, we keep the camera and light source fixed and move the target sample, as shown in Figure 14. Figure 15 showcases our results – geometry, reflectance, as well as images rendered under novel lights – on datasets captured using the prototype. These six real world datasets have 113, 136, 105, 160, 92 and 110 input images, respectively. As shown here, images rendered using our estimated normals and BRDFs under novel lights closely resemble actual captured photographs, indicating the robustness and accuracy of our method.

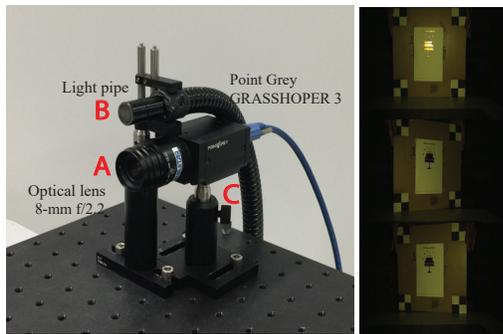


Figure 14. Our hardware prototype and example captured images.

Comparisons between PTGrey Grasshopper and iPhone 6s. Figure 16 compares the performance of the shape estimates by using PointGrey prototype. Our prototype acquisition hardware consists of a PointGrey Grasshopper 3 GS3-U3-23S6C-C camera and a light pipe. During image capture, we keep the camera and light source fixed and move the target sample, as shown in Figure 14.

The PointGrey Grasshopper camera has a pixel pitch of $6.45\mu\text{m}$ which allows us to collect high-quality images. For shiny objects with sharp specular lobes, we also collect HDR images for improved performance. In contrast, the iPhone 6s has a much smaller pixel pitch ($1.22\mu\text{m}$) and the app that we used to collect the dataset did not allow for HDR acquisition or changing the brightness of the flash. Hence,

it is expected that we see a drop in performance when using the iPhone 6s to capture shape and reflectance. Surprisingly, this difference is minimal as seen in the quality of results in Figure 16

Material trait analysis. Figure 17 demonstrates material trait analysis using our method on a dataset captured with both our PointGrey prototype and iPhone 6s. This is similar in spirit to the results in the main paper.

References

- [1] Z. Hui and A. C. Sankaranarayanan. A dictionary-based approach for estimating shape and spatially-varying reflectance. In *ICCP*, 2015. 3
- [2] W. Jakob. Mitsuba renderer, 2010. URL: <http://www.mitsuba-renderer.org>, 3:10, 2015. 3
- [3] A. Ngan, F. Durand, and W. Matusik. Experimental analysis of brdf models. *Rendering Techniques*, 2005(16th):2, 2005. 3
- [4] J. Riviere, P. Peers, and A. Ghosh. Mobile surface reflectometry. In *Computer Graphics Forum*, 2016. 1, 3, 11
- [5] F. Romeiro, Y. Vasilyev, and T. Zickler. Passive reflectometry. In *ECCV*, 2008. 3
- [6] Z. Xu, J. B. Nielsen, J. Yu, H. W. Jensen, and R. Ramamoorthi. Minimal brdf sampling for two-shot near-field reflectance acquisition. *TOG*, 2016. 3



Figure 11. We illustrate how our surface normal estimates evolve for a number of scenes with different geometry and materials. Starting from flat surfaces, our method refines the normals so that over the course of 10 iterations they converge to results that have a mean angular error of only 0.73° to 1.14°.

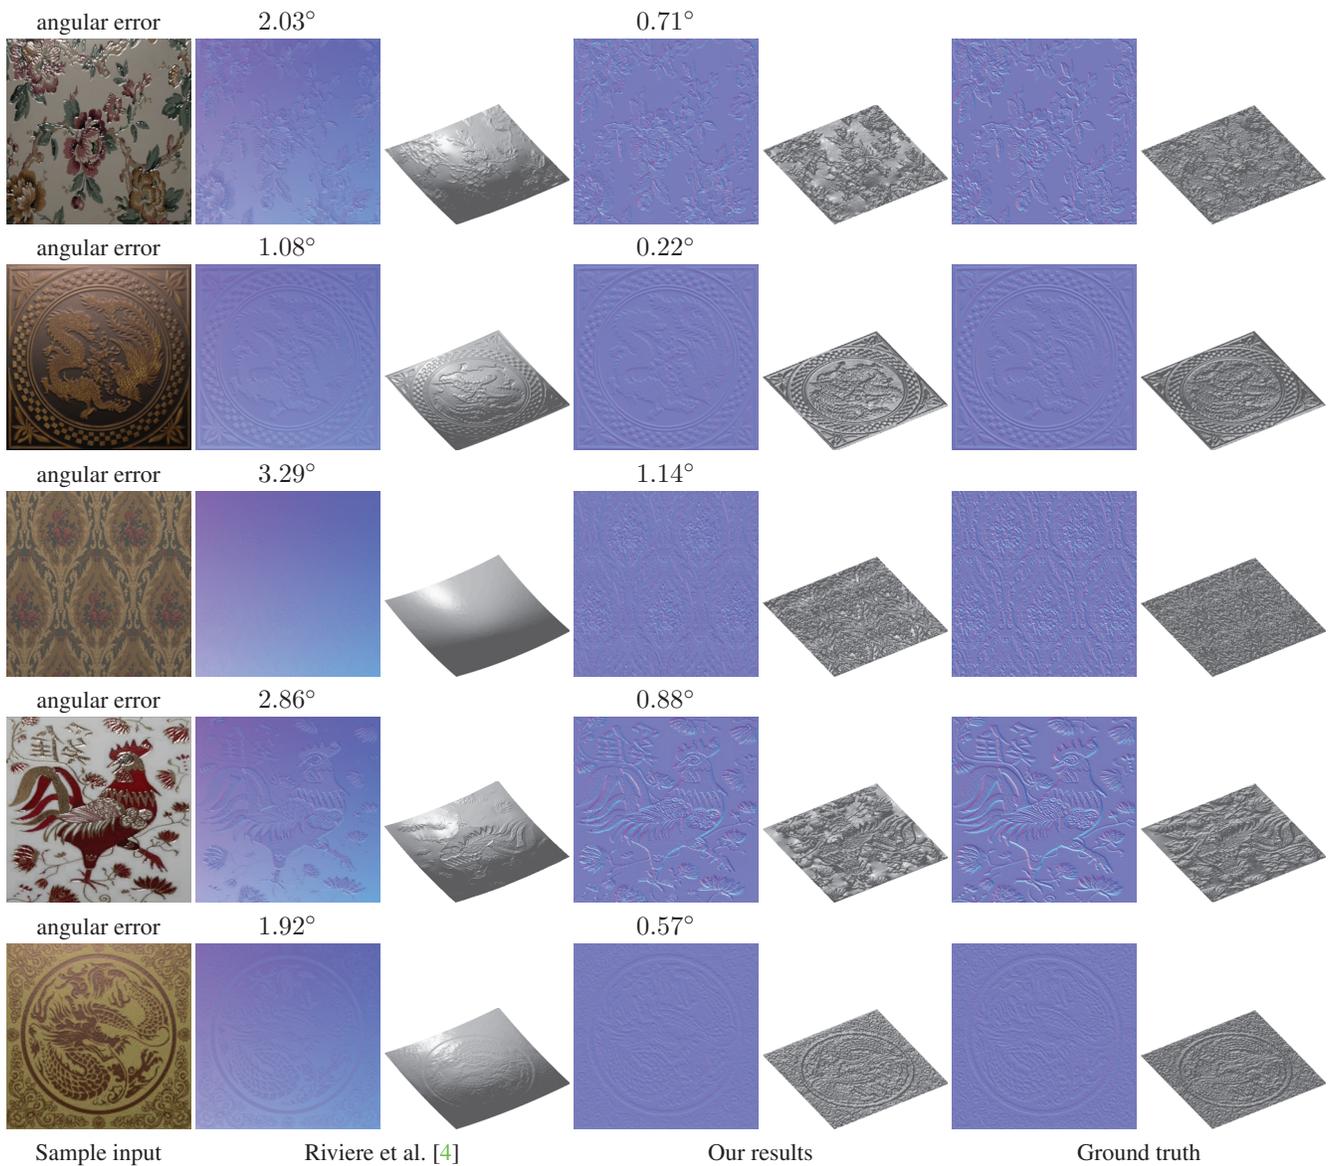
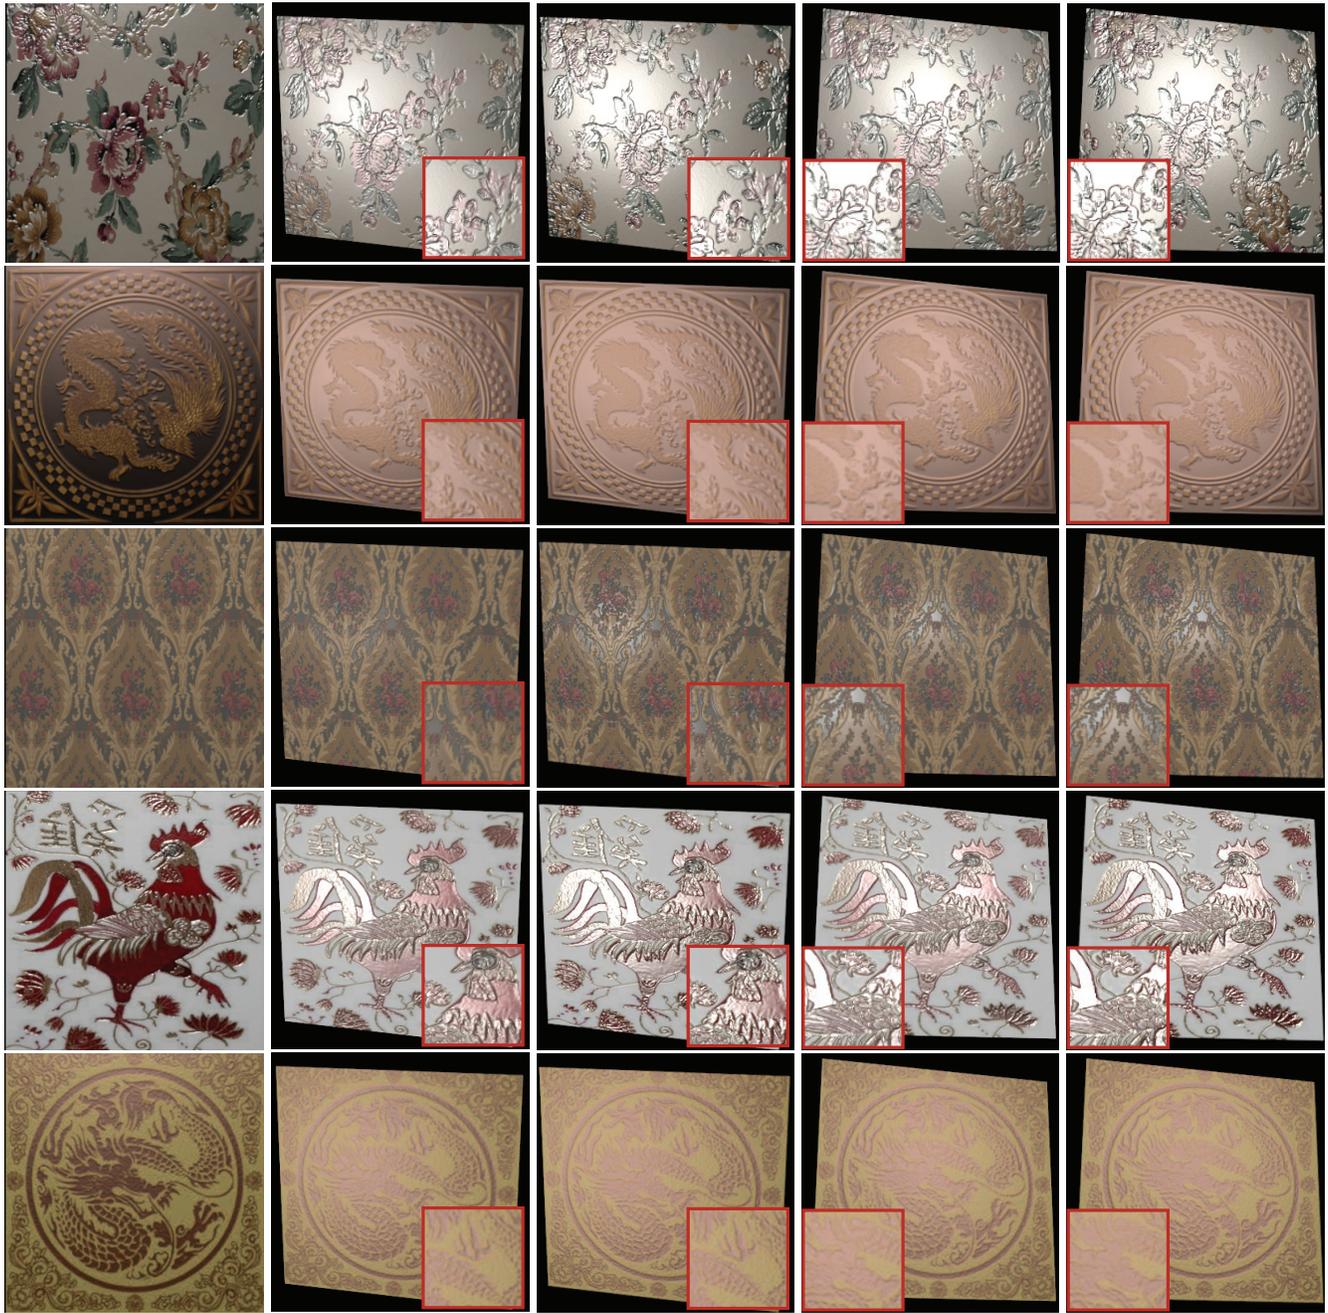


Figure 12. We compare our surface normal estimates against those from Riviere et al. [4] on our synthetic dataset. Shown are (left-right) one sample image, estimated normals and recovered 3D shape via Poisson reconstruction using Riviere et al., our method, and the ground truth. Please note that our reconstructions, like the actual samples, are close to planar and contain more fine-scale detail. In addition our average angular error is significantly lower than theirs.



Sample input

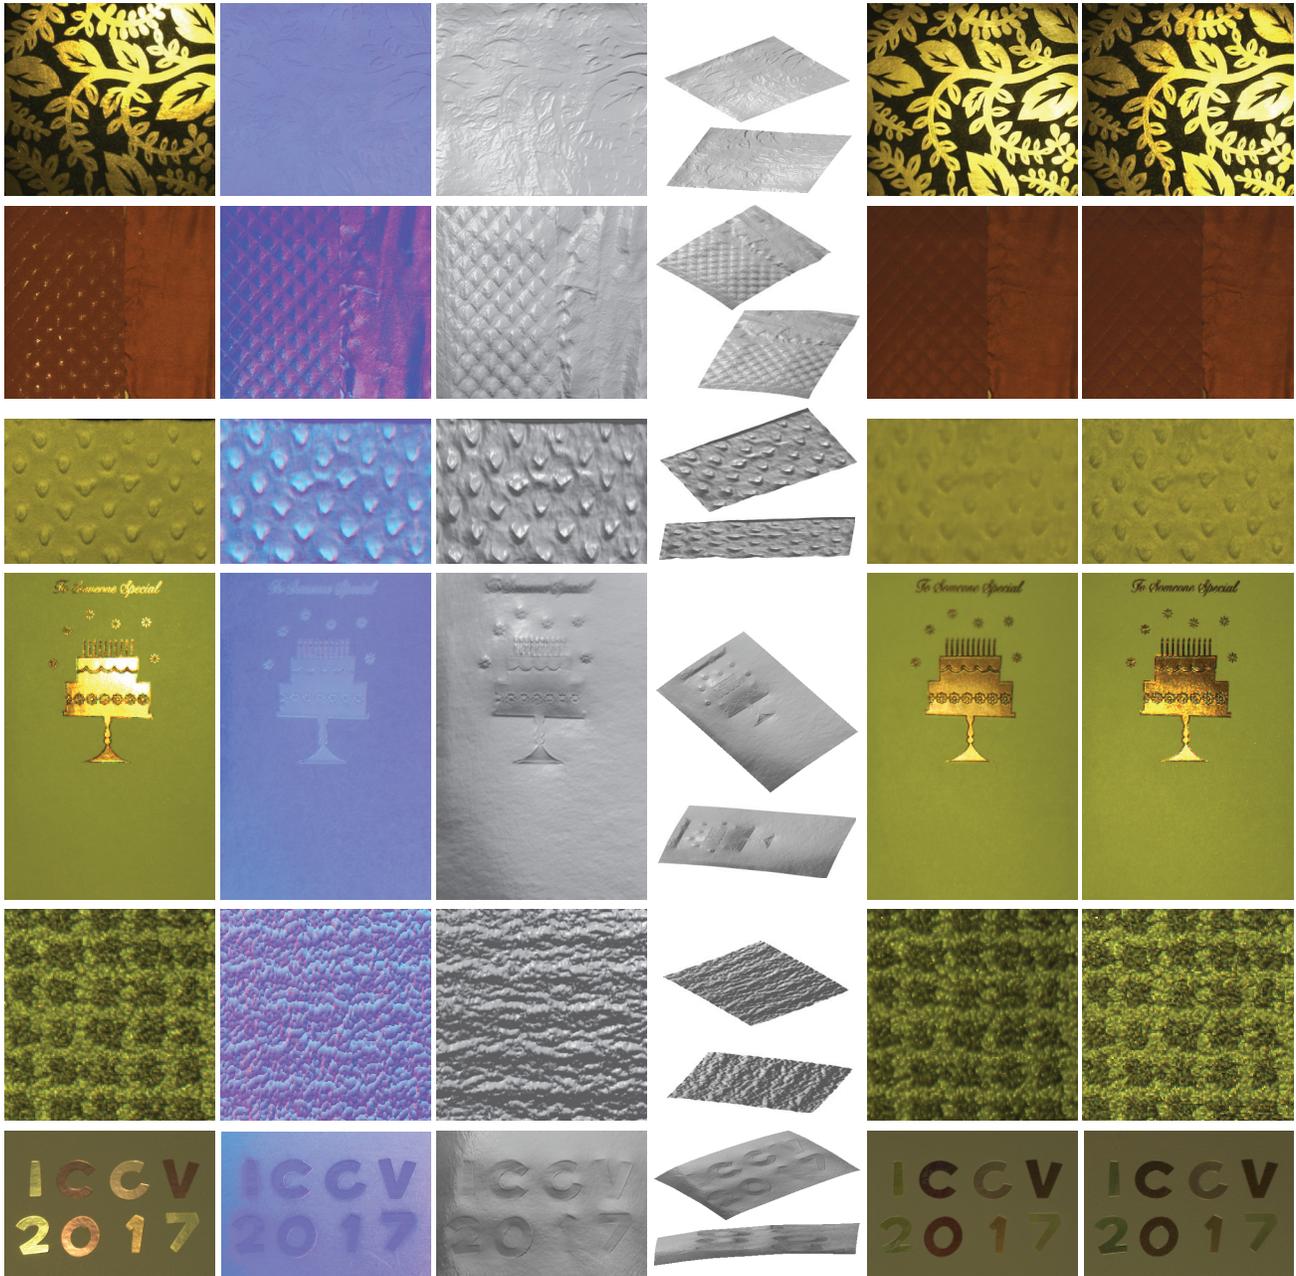
Rendering with
estimated normals/BRDF

Ground truth
rendering

Rendering with
estimated normals/BRDF

Ground truth
rendering

Figure 13. We visualize our estimated reflectance and normals under two novel view and lighting conditions. We also compare them against ground truth renderings under the same conditions. Our results closely match the ground truth images, demonstrating that while our method only observed collocated measurements, it generalizes well to novel views/lights.



(a) Input sample (b) Estimated normals (c) Recovered surface (d) Rendering (e) Photograph

Figure 15. We demonstrate shape and reflectance estimation on images captured using an iPhone 6S (a). We show the estimated normal map in false color (b) and depthmaps (c). We also compare rendered (d) results against actual captured photographs under novel lighting (e).

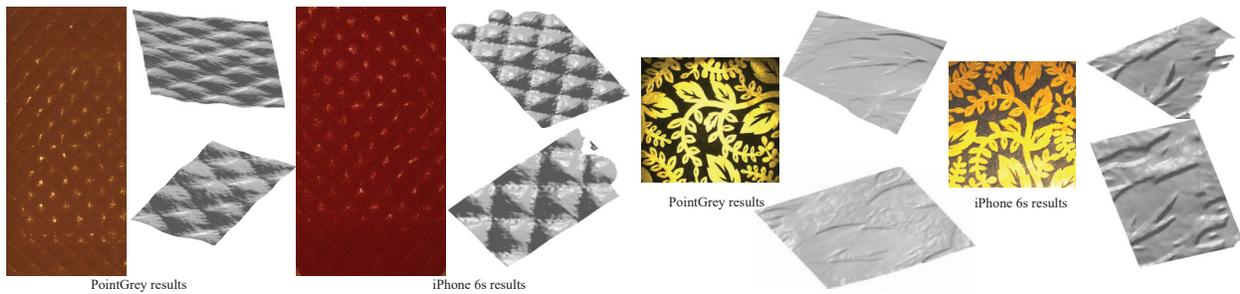


Figure 16. We demonstrate the performance of shape estimation by using both pointGrey and iPhone 6s. As can be seen here, we can achieve nearly the same quality by using either PointGrey or iPhone 6s, indicating the robustness of our technique.

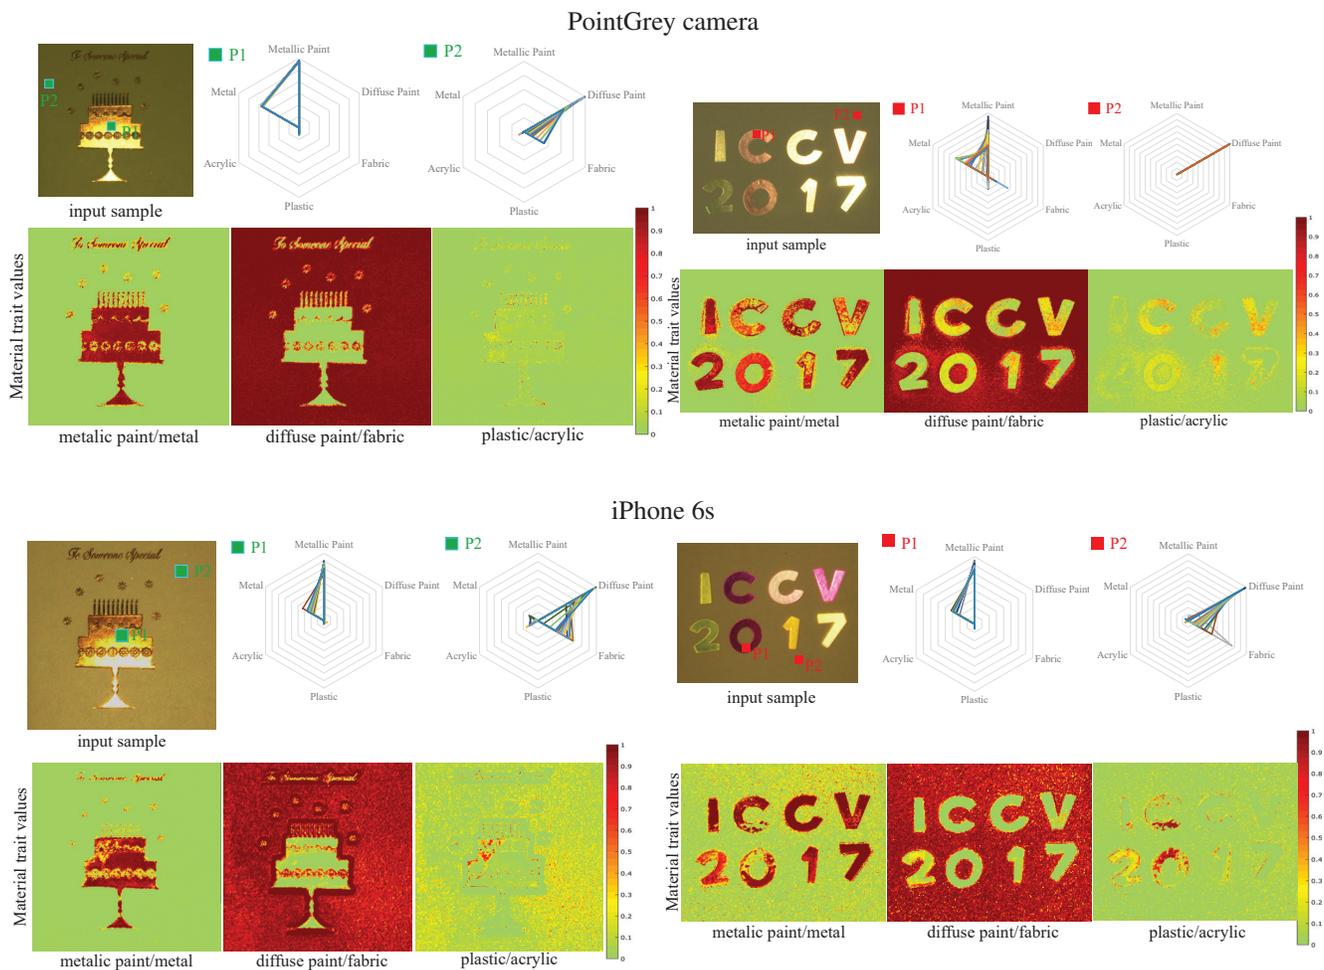


Figure 17. Material trait analysis on real capture data. (top) For two regions indicated by p_1 and p_2 , we plot the associated material trait values according to Section 5 in the main paper. Pixels (p_1) with metallic properties have large values in *metallic paint* and *metal* while pixels (p_2) with diffuse Lambertian-like materials show large values in *diffuse paint* and *fabric*. (bottom) We visualize per-pixel material trait values for three material groups — *metallic paint+metal*, *diffuse paint+fabric*, and *plastic+acrylic*. This leads to clean, consistent material segmentations.