

EventPSR: Surface Normal and Reflectance Estimation from Photometric Stereo Using an Event Camera

Supplementary Material

Bohan Yu^{1,2†} Jin Han^{3,4} Boxin Shi^{1,2*} Imari Sato^{3,4}

¹ State Key Laboratory for Multimedia Information Processing, School of Computer Science, Peking University

² National Engineering Research Center of Visual Technology, School of Computer Science, Peking University

³ The University of Tokyo

⁴ National Institute of Informatics

ybh1998@pku.edu.cn, jinhan@nii.ac.jp, shiboxin@pku.edu.cn, imarik@nii.ac.jp

6. Details for the parameterized BRDF model

The detailed formula of D and G in Eq.(3) are as follows:

$$D = \frac{1}{\pi R (\mathbf{N} \cdot \mathbf{H})^4} \cdot \exp \left(\frac{(\mathbf{N} \cdot \mathbf{H})^2 - 1}{R(\mathbf{N} \cdot \mathbf{H})^2} \right),$$

$$G = \min \left(1, \frac{2 (\mathbf{H} \cdot \mathbf{N}) (\mathbf{w}_i \cdot \mathbf{N})}{\mathbf{w}_i \cdot \mathbf{H}}, \frac{2 (\mathbf{H} \cdot \mathbf{N}) (\mathbf{w}_r \cdot \mathbf{N})}{\mathbf{w}_i \cdot \mathbf{H}} \right), \quad (12)$$

$$\mathbf{H} = \frac{\mathbf{w}_i + \mathbf{w}_r}{|\mathbf{w}_i + \mathbf{w}_r|_2},$$

where \mathbf{N} is the surface normal, and R is the surface metallic. Both of them have the same definition as Eq.(4), with the pixel coordinates (x, y) omitted. \mathbf{w}_i and \mathbf{w}_r are the incident and reflection direction, respectively. The specular distribution term D describes the microfacet normal direction distribution. Here we use Beckmann distribution [16] as D . The geometry term describes the cast shadow of the microfacet model. We use the same formula in the Disney BRDF model [4] and Torrance-Sparrow model [26] as G .

7. Event Integral Drifting issue

We show a case of event integral curve in Fig. 9. Due to non-ideal triggering of real event camera, the integral of events (orange line) can be constantly rising or falling, causing mismatch compared with rendered database (green line). However, as shown in Eq. (9), we directly compare each individual pair of events with the rendered database without the event integral, to avoid the event integral drifting issue.




Figure 9. Demonstration of the event integral drifting issue. The integral of the event signal (orange line) is constantly rising in this example, while the actual intensity (green line) is returning to the initial value.

8. Device Coverage

The monitors in our data acquisition platform surround the objects horizontally, leaving two uncovered areas at the top and bottom. For a mirror ball, the coverage map, shown in Fig. 10, reveals these uncovered regions at the center and border of the object. Within the coverage region, surface normal and material properties are accurately estimated, while in the uncovered areas, they are interpolated during the second gradient-tuning stage. We achieved a coverage area of 69.25%, using five monitors in our data acquisition platform. This coverage can be further improved with a more compact monitor arrangement.

9. Roughness and Metallic Standardization

The definitions of roughness and metallic vary across different models used by comparison methods. To ensure consistency, we normalize the roughness and metallic outputs for all results in the logarithmic domain. Let the original output of a method on the entire synthetic dataset be represented as a vector \mathbf{X} (where \mathbf{X} can denote either roughness or metallic), and let the ground truth values be represented

*Corresponding author: Boxin Shi

†This work was conducted while the first author, Bohan Yu, was doing internship at National Institute of Informatics.




Figure 10. The coverage map on data acquisition platform. The cyan on the sphere represents the area that can be covered by the light scanning pattern. The red area is the uncovered area.

as $\hat{\mathbf{X}}$. The normalized method output \mathbf{X}^* is defined as:

$$S = \frac{\text{std}(\log(\hat{\mathbf{X}}) - \text{mean}(\log(\hat{\mathbf{X}})))}{\text{std}(\log(\mathbf{X}) - \text{mean}(\log(\mathbf{X})))},$$

$$T = \text{mean}(\log(\hat{\mathbf{X}})) - \text{mean}(\log(\mathbf{X})) * S,$$

$$\mathbf{X}^* = \exp(\log(\mathbf{X}) * S + T).$$
(13)

This process is similar to applying batch normalization to the logarithmic roughness and metallic values across the entire synthetic dataset. Consequently, the mean and standard deviation of the log-transformed roughness and metallic are aligned with those of the ground truth. By normalization in this manner, discrepancies in scale and power across different material property definitions are eliminated, ensuring fair and consistent quantitative comparisons.

10. Robustness on Various Materials

To demonstrate the robustness of EventPSR in recovering various materials, we evaluate the performance of the three methods on a cow object with 6 different surface materials, from ideal Lambertian to pure mirror, as shown in Fig. 11. Since the PS-based method [13] is designed from the Lambertian reflection model, the error of normal estimation is becoming larger as the surface material is becoming shiny. The shiny materials are also challenging for NeRF-based method [31], which even failed to reconstruct the object after the “rough metallic” material in Fig. 11. Because highly specular and metallic materials are too challenging for the geometry initialization stage in NeLF++ [31].

	Lambertian	Weak specular	Strong specular	Rough metallic	Shiny metallic	Mirror
Ours Data Ratio	0.388	0.392	0.411	0.448	0.446	0.361
Ours Normal	5.155	5.593	3.345	2.093	1.362	1.401
NeLF Normal	3.221	2.975	2.738	Failed	Failed	Failed
SDM Normal	8.153	2.513	2.590	14.899	21.099	24.685
Ours Roughness	N/A	0.010	0.010	0.009	0.000	0.000
NeLF Roughness	N/A	0.030	0.027	Failed	Failed	Failed
SDM Roughness	N/A	0.008	0.006	0.063	0.010	0.001
Ours Metallic	N/A	0.027	0.007	0.008	0.000	0.001
NeLF Metallic	N/A	0.067	0.032	Failed	Failed	Failed
SDM Metallic	N/A	0.448	0.123	0.134	0.118	0.386

Figure 11. Case study on robustness. The materials of objects in the first row vary from ideal Lambertian to pure mirror. EventPSR performs consistently for different materials, while the comparing methods fail or have an increasing error as the specular part gets more and more dominant.

11. Additional Results on Synthetic Data

To evaluate the performance of EventPSR, we conduct extensive experiments on synthetic dataset, as shown in Fig. 13. The synthetic dataset includes objects with varying levels of roughness, metallicity, and geometric complexity to test the robustness and versatility of our method.

12. Ablation study

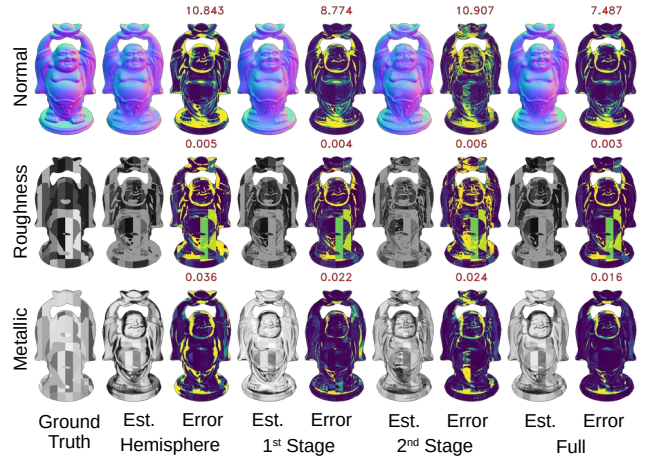


Figure 12. Ablation study shows that, high illumination coverage ratio and the full two-stage pipeline solution are both essential to achieve high accuracy.

We conduct ablation experiments on synthetic data with reduced light coverage ratio using a hemisphere (50% coverage ratio) and evaluate our two-stage estimation pipeline. As shown in Fig. 12, hemisphere illumination yields higher errors in bottom and boundary areas, likely due to specular

reflection in uncovered regions. The first-stage-only results exhibit high overall error, likely due to low parameter resolution, while the second-stage-only results show incorrect patches, possibly stuck in local minima. The full pipeline’s superior accuracy underscores the importance of both high light coverage and our proposed two-stage solution.

13. Discussion and Potential Improvements.

The accuracy of our current solution is bottlenecked by non-Lambertian effects, and the speed is limited by the illumination device. We propose several potential extensions: First, to enhance robustness against cast shadows and inter-reflections, explicit modeling methods such as DANI-Net [18] or deep learning-based PS [5, 14] can be employed. Besides, applying a multi-view system can provide more comprehensive coverage of specular regions under limited illumination. For the diffuse albedo recovery, we can integrate a frame-based camera to create a hybrid camera system, or leverage alternative intensity estimation techniques like [11]. For the illumination system, we utilize multiple OLED panels for broad light coverage, uniform intensity, and mostly white color. A different illumination setup may have non-uniform intensity & color, or position misalignment. These artifacts can be reduced by the end-to-end calibration process. For real-time capturing, our high event efficiency design depicted in Fig. 2 (c) suggests that the event camera can accommodate illumination with increased speed. To achieve real-time processing, we can replace the grid search with more efficient algorithms, such as heuristic search or locality-sensitive hashing.

14. Light Pattern and Relighting Result Video

We provide a video including the working data acquisition platform, and the light scanning pattern illuminated by the monitors. We also render the object relighting results with white uniform albedo in the supplementary video.




Figure 13. Additional evaluation results on the synthetic dataset.