

Applying BRDF Monotonicity for Refined Shading Normal Extraction from Multi-Light Image Collections

Ruggero Pintus¹, Antonio Zorcolo², Enrico Gobetti³

CRS4, Italy

Abstract

Multi-Light Image Collections (MLICs) are often transformed into geometric normals and BRDF normals for visual exploration under novel illumination. However, discrepancies between the chosen BRDF space and the complete optical behavior of objects, along with the possible presence of non-local lighting effects in measurements, often lead to sub-optimal visual outcomes even with the most accurate geometric normal recovery. In this paper, we introduce a modular component designed to convert the geometric normals into well-behaved shading normals, under the common and general assumption that the reflectance must be a monotonic function of the angle between the shading normal and the bisector of lighting and viewing directions. Since it does not require the coupling of shape and material estimation, the module allows seamless integration into existing reconstruction pipelines, supporting the mixing and matching of Photometric Stereo methods, BRDF models, and BRDF fitters. The performance and versatility of the approach are demonstrated through experiments.

CCS Concepts

• *Computing methodologies* → *Shape modeling*; *Computer graphics*; • *Applied computing* → *Arts and humanities*;

1. Introduction

Multi-Light Image Collections (MLICs), i.e., stacks of photos acquired with a fixed viewpoint and varying surface illumination, are one of the most widely used solutions for capturing information on the shape and appearance of objects in many fields, cultural heritage being the prevalent one [PDC*19]. In the most common configuration, data is acquired using a different point or directional light per image (using a dome or a moving light), leading to a pixel representation with a different color for each sampled light direction.

While classic virtual inspection solutions for such data were restricted to exploiting low-frequency analytical relighting representations fitting the acquired light-direction-dependent colors, recent work started targeting physically-based rendering using recovered normals and Bidirectional Reflectance Distribution Functions (BRDF) maps, since this representation can be easily distributed at low bitrate, produces a physically-reasonable result, and supports natural integration with standard high-quality and real-time rendering solutions (Sec. 2). However, despite the considerable advances in normal computation through Photometric Stereo and BRDF parameter fitting, typical renderings with such a model may sometimes fail to match the original measurements well. This issue arises not only when geometric normals are inaccurately recovered but also in cases of a "perfect" geometric reconstruction that is fully suitable for geometric inspection purposes. Several practical factors contribute to this problem. In addition to the consideration that input data may include non-local lighting effects that the employed

BRDF fitter may not account for, there often is an inherent discrepancy between the simplified normal+BRDF model and the complex optical behaviors of many real-world objects, including layering, where different materials are stacked and interact in complex ways or subsurface scattering, where light penetrates the surface, interacts with the material beneath, and then exits at a different location. These phenomena are challenging to model accurately with standard single normal+BRDF representations. In particular, the shape space of the reflectance function defined by the chosen BRDF model may not be capable of closely matching the input data given the local geometric normal. For this reason, several approaches try to expand the solution by jointly minimizing shape and appearance together in a single framework (e.g., through alternating minimization techniques) [GCHS05, AZK08, HS16]. Unfortunately, they increase the system's computational complexity, are less controllable than separate Photometric Stereo and BRDF fitting solutions, are more difficult to upgrade, and often do not ensure proper convergence [HS16].

This paper proposes an approach that does not break the modularity of classical shape and material characterization steps. This is done by decoupling the normals used for shading from the geometric normals, as done in many rendering frameworks. In our approach, the geometric normal coming from Photometric Stereo is slightly tweaked (i.e., moved by a few degrees), to ensure a better consistency of the behavior of the measured data with the output that can be generated with local shading using a generic BRDF model. Based on the common and general assumption that the recovered BRDF

should be a monotonic function of the angle between the shading normal and the bisector of lighting and viewing directions, we introduce a novel monotonicity metric and search in the neighborhood of the photometric stereo normal, for the normal that maximizes our metric. After that, we use this refined shading normal to fit the dense analytic pixel BRDF. Unlike methods that require coupled shape and material solvers, this approach integrates seamlessly into existing reconstruction pipelines, allowing flexible combinations of Photometric Stereo methods, BRDF models, and BRDF fitters. Users can choose to utilize our corrected normals as improved geometric normals, particularly when using simple normal recovery techniques, or, in the most general case, to use separate normal maps for geometric analysis and relighting using local shading.

2. Related work

Computing reflectance transformation or shape and material models from MLICs are well-researched subjects. We discuss here the most closely related approaches, referring the reader to established surveys [GGG*16, PDC*19, LBFS21, WRG*23] for broader coverage.

Relightable images directly encode, per-pixel, a continuous map from lighting parameters to rendered color, eliminating the need for explicitly determining a separation between shape, material, and lighting [PDC*19]. Polynomial Texture Mapping (PTM) [MGW01] and Hemi-Spherical Harmonics (HSH) [GKPB04] are formulations of widespread use. The compactness and low complexity of relightable images make them highly suitable for fast interactive relighting in both local and remote visualization, making them de-facto standard formats for relighting in Cultural Heritage application scenarios. However, these methods do not provide a physics-based characterization of the surface independent of the measurement illumination. The lack of separation between shape and material components makes integrating such representations into common rendering pipelines difficult. This makes it difficult to insert them into scenes or to simulate shape-dependent effects such as shadows or global illumination. For these reasons, practical solutions to reconstruct a separate surface geometry and optical response are the focus of much research [PDC*19]. Photometric Stereo and BRDF fitting are standard approaches to address those issues.

After the first photometric stereo methods tuned for Lambertian surfaces [Woo80], a large variety of solutions have been proposed to tackle complex models using conventional fitting (e.g., [Geo03, RK09, STMI13, HLHZ08]), example-based solutions (e.g., [HS05, HS16]) or deep-learning approaches (e.g., [CHS*19, SSS*17, JLX*24]). Our module is independent from the method used to estimate the geometric normal, but acts as a post-process to tweak it towards the direction that better explains the measured data using a monotonic BRDF function. In prior work, monotonicity has already been used for normal reconstruction from highlights or for specular lobes characterization, by considering monotonicity with respect to the cosine of the bisector between the view and light directions, and the surface normal [OKT*19, STMI12, AZK08, HMI10]. In these works the monotonicity metrics are simply computed by using the signum-based method [Gri06, Mat24], which locally checks the sign of the derivatives, and computes its normalized sum. We also employ monotonicity constraints but in the context of a refinement step that improves SV-BRDF fitting starting from geometric

normals rather than as a means to compute normals from measurements. In this work, we introduce a novel monotonicity metric and exploit it to compute a shading normal from a previously computed geometric normal.

A sub-class of methods that extract optical surface behavior in terms of SV-BRDF can accept MLIC data as input signal [GGG*16, PAZ*23]. This is a specific condition since MLIC represents only a slice of the bigger multi-dimensional BRDF space. While several solutions find the BRDF parameters that best match the measured data using a fixed photometric normals, others cope with normal refinement using an alternating minimization approach for normal and BRDF fitting [CNK16]. Although powerful, the downside of alternating minimization approaches lies in their computational expense and their significant reliance on obtaining a favorable initial solution, which is typically the starting normal. This challenge arises due to the non-convex nature of the underlying problem, which is fraught with local minima [HS16]. Moreover, it should be noted that, in real conditions, local color may often be affected by non-local illumination effects (e.g., inter-reflections), making this optimization even harder. Most solutions use local minimizers, and there is no guarantee that the starting solution is in the basin of convergence, especially when there is only very sparse sampling, as in MLIC conditions.

3. Method

We consider a generic pipeline that takes as input a MLIC of an object, accompanied by geometric and radiometric calibration metadata, and produces, as output, a normal and a BRDF map. In the classic approach, such a generic pipeline processes the MLIC through a photometric stereo module to estimate per-pixel normals, storing them into a normal map. These normals are then fed to an SV-BRDF fitting module to estimate per-pixel BRDF parameters. This work introduces a new module that modifies the photometric stereo normals into shading normals better suited for BRDF estimation using local illumination approaches.




Figure 1: Sampled BRDF dependency on normal estimation on a synthetic dataset of a sphere. The leftmost column shows an image obtained with the correct sphere normal, while the rightmost column is tilted by $\Delta\theta = 30^\circ$ degrees. The first row shows the dependency of $\hat{\rho}$ to θ_h , highlighting how the monotonicity is strongly violated with a wrong normal.

Our input data consists of the MLIC reflectance measurements (the point/pixel appearance profile), the corresponding series of light directions L_i and intensities I_i , the point normal N , and a single

viewing direction V . For a MLIC acquisition, the camera has a fixed position and viewing direction, so, for a general camera matrix, each pixel has its viewing vector which remains fixed within its entire appearance profile. We choose Rusinkiewicz’s representation [Rus98] to define a 2D domain to organize the BRDF values spatially. Based on the triplet (N, V, L_i) those values are projected on a domain that is defined by two axes. The horizontal axis represents the values of the angle θ_h between the surface normal and the half vector $H = \frac{V+L}{\|V+L\|}$, while the vertical axis is the angle θ_d between the half vector H and the viewing V (or the light L) vector. Given a general BRDF model $f(N, V, L, \dots)$ and a measured reflectance R , we choose the standard cosine-based formulation for the BRDF value normalization, where $R = f(N, V, L, \dots)(N \cdot L)I$. Each BRDF value associated with a measured reflectance R_i in the appearance profile will be $\rho_i = \frac{R_i}{I_i \cos(\theta_i)}$, where θ_i being the angle between the light direction L_i and the normal N . For a standard (and practical) MLIC, the BRDF representation results in a sparse image. Note that both the spatial organization of the sparse samples and their radiometric values will change with a variation of the surface point normal N . Fig. 1 shows an example of the same appearance profile of a specific pixel seen under different surface normal vectors, including a “correct” normal and a normal in the neighborhood (here, tilted by $\Delta\theta = 30^\circ$ degrees to better highlight the behavior change). Observe that monotonicity is preserved in the first image of Fig. 1, while it is strongly violated in the other. In this example, the normal is tilted by $\Delta\theta = 30^\circ$ degrees to better highlight the behavior change, but the deformation follows the same rules for smaller angles.

We devised an operator to compute the monotonicity-based metric, focusing on the relationship between the angle θ_h and the normalized BRDF value. Initially, it eliminates the dependency on θ_d by projecting all BRDF samples onto the horizontal axis. Then, it further projects the multi-spectral BRDF value onto the hyper-diagonal of the multi-spectral hyper-cube by computing $\tilde{\rho} = |\tilde{\rho}(\theta_h)|$, where $\tilde{\rho}$ represents the projected BRDF value and $\tilde{\rho}$ denotes its amplitude. The functions are plotted above the corresponding BRDF images in Fig. 1. Monotonicity of these functions is evaluated using a metric inspired by the standard *Signum Formula*, yet with several refinements: (1) consideration of the duality of diffuse and non-diffuse signal composition of the BRDF; (2) asymmetric normalization weighted by the gradient magnitude; (3) incorporation of a quick estimation of the albedo value to enhance metric robustness to tiny variations and oscillations in the BRDF value due to measurement noise; (4) definition within the interval $\{0, 1\}$, where 1 indicates decreasing monotonicity (typical BRDF behavior), 0.5 represents a constant function (e.g., perfectly Lambertian diffuse case), and 0 signifies perfectly increasing monotonicity.

For each pixel, we begin by setting the initial value of the monotonicity as $\mu = 1$. We analyze a curve represented by a series of data points, specifically the values of $\tilde{\rho}$. We then examine each pair of adjacent data points along this curve. When comparing each pair, we calculate the difference between the $\tilde{\rho}$ value of the current point i and that of the next point $i + 1$, referred to as δ . When δ is bigger than a tolerance value ϵ , there is an upward trend between the two points and thus a violation of the BRDF’s decreasing monotonicity. We thus update the monotonicity by reducing it according to the magnitude of δ . This update is computed as $\mu_{new} = \mu_{old} / (1 + \delta)$. Upon examining all data points in the curve, the last updated (current)

value of μ becomes the final monotonicity metric. The tolerance $\epsilon = \max(\epsilon_{abs}, \epsilon_{rel})$ takes into account both the minimum absolute numerical error ϵ_{abs} in the BRDF values, and a relative tolerance η with respect to the diffuse signal *albedo*, i.e., $\epsilon_{rel} = (1 + \eta) \times albedo$. Setting $\eta = 0.1$ and $\epsilon_{abs} = 10^{-4}$ is sufficient to mitigate the effects of noise. The shading normal \tilde{N} is then obtained by launching an optimization procedure that maximizes the monotonicity as

$$\tilde{N} = \operatorname{argmin}_N [1 - (\mu(N) + \mathbf{R}(N))] \quad (1)$$

where $\mathbf{R}(N)$ denotes a regularization term aiming to minimize the deviation between the initial normal and the final one, thereby preventing excessive divergence. The solution is not computed for every possible normal, but only for those in a small cone around the photometric normal N . In this paper, the search cone angle is $\pi/6$. We perform all computations in a canonical search space where the photometric normal is aligned with the z-axis. Thus, the regularization term simplifies to $\mathbf{R}(N) = \lambda N_z$. We found that a small value of $\lambda = 10^{-4}$ works well in all our tests. Since the function may have local minima due to the nature of the response of the monotonicity field to changes in the normal orientation and the sampling of the BRDF, where it is common to have most values in the diffuse area, traditional gradient-based minimizers are not appropriate. For this reason, we adopted a bounded global search, calculating all monotonicity values for a dense set of normals sampled in the search cone and choosing the one with the smallest cost.

4. Results

The newly proposed normal map refinement module has been integrated into an existing MLIC pipeline [PAZ*23], without any modification to the photometric stereo and BRDF reconstruction components. In this paper, we first present results on a synthetic dataset for which ground truth is available (Sec. 4.1). We then illustrate the results obtained on a free-form acquisition of a painting (Sec. 4.2). This work briefly presents illustrative results that will be expanded in a future journal publication.

4.1. Synthetic evaluation and modularity test

We have evaluated our method using a variety of configurations. Here, we summarize the results obtained with two different photometric stereo solutions to estimate geometric normals and two different BRDF formulations. The first photometric stereo method is the classic *Trimmed PS* [WGT*05], which performs a Lambertian fitting after removing the darkest (shadow) and lightest (highlight) observations, and *CNN-PS* [Ike18], a Convolutional Neural Network-based Photometric Stereo that has demonstrated strong performance on the standard *DiLiGenT* benchmark [SWM*16]. The BRDF models we employ are the dichromatic isotropic Ward model [Wal05] characterized by a diffuse color, a specular color, and a monochromatic glossiness, and the stretched Blinn-Phong model [NNSK99], a compact empirical model with a physically-realistic behavior for dielectric and metallic materials.

Synthetic tests conducted on rendered models from manually defined maps allow for evaluating the methods in a fully controlled environment where ground truth is available. We selected a synthetic model from the *EveryTexture* database [Eve24], as it closely

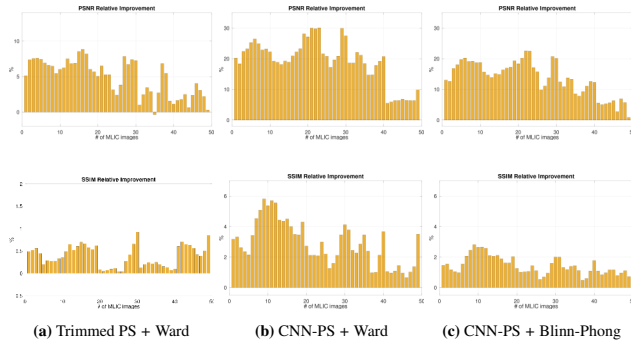


Figure 2: The performance of all relit images for the Paint-Texture-01 dataset is assessed. We present the percentage relative improvement for both the PSNR in the first row and the SSIM in the second row. We present the results for two pipelines that utilize combinations of Trimmed PS or CNN-PS as the Photometric Stereo method, and either Ward or Blinn-Phong formulation for the BRDF fitting step.

Synthetic Dataset	Method	FLIP Statistics			
		Mean	Weighted Median	1st Quartile	3rd Quartile
	Trimmed PS + Ward	0.09	0.16	0.08	0.24
	Trimmed PS + Ours + Ward	0.08	0.11	0.06	0.20
	CNN-PS + Ward	0.11	0.15	0.09	0.28
	CNN-PS + Ours + Ward	0.05	0.09	0.05	0.16
	CNN-PS + Blinn-Phong	0.15	0.19	0.13	0.31
	CNN-PS + Ours + Blinn-Phong	0.11	0.16	0.10	0.21

Table 1: FLIP statistics for one image of the Paint-Texture-01 dataset for two pipelines that utilize CNN-PS as the Photometric Stereo method, and either Ward or Blinn-Phong formulation for the BRDF fitting step.

resembles the detailed shape and appearance of real objects that are typically measured with MLICs, such as paintings. Utilizing the provided *Diffuse*, *Bump*, and *Normal* maps, we retained the *Diffuse* map as is, generated the *Gloss* component by converting the *Bump* map to monochrome and rescaling it, and assigned a constant highlight color. These maps were then cropped to generate synthetic MLICs using a fixed camera and 52 directional lights by applying the Ward model. The data presented here is for the *Paint-Texture-01* dataset (see the first column of Tab. 1 for an example image).


We applied the different combinations of methods and measured the change in accuracy obtained when using or not using our normal refinement module. To synthetically present results, we use percentage relative improvement, defined as by $\Gamma = 100 \frac{|m_1 - m_0|}{m_0}$, where m_0 denotes the metric value (PSNR or SSIM) representing the error between the original and the relighted image obtained with the non-refined normal, and m_1 the metric computed with the refined normal. Fig. 2 summarizes the PSNR and SSIM for this dataset. It can be noted that all figures are positive, showing that our method improves over standard solutions for every tested light direction. Additionally, Tab. 1 reports the improvements in FLIP [ANA*20] statistics for a sample relighted image in this dataset. This test provides insight into how our method consistently delivers positive performance when seamlessly integrated into general shape and appearance characterization frameworks.

4.2. Evaluation on a real painting

To assess the results obtained by our method in a typical cultural heritage setting, we have analyzed its behavior on the common handheld moving light source setup, also known as free-form MLIC



(a) A panel of the Retable of S. Bernardino



(b) Retablo Crop #0 (c) Retablo Crop #1 (d) Retablo Crop #2

Figure 3: A panel from the Retable of Saint Bernardino captured using a free-form RTI acquisition setup. Three image excerpts showcasing various material and geometric properties are employed to evaluate the effectiveness of the proposed normal refinement technique.

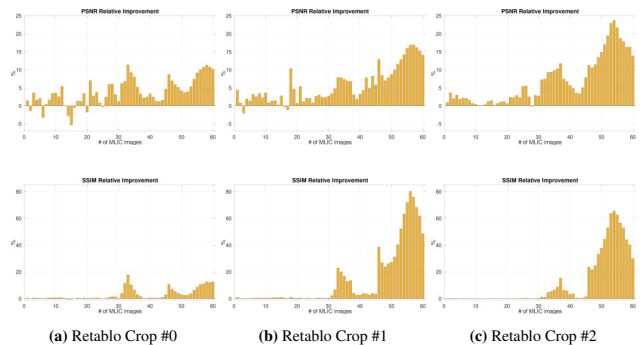


Figure 4: The performance of all relit images for the three Retablo crops is assessed. We present the percentage relative improvement for both the PSNR in the first row and the SSIM in the second row. Noticeably better results are observed in images with frontal illumination (higher image numbers) compared to those with raking light conditions (lower image numbers), which primarily display the surface's diffuse signal.

acquisition. In this work, we present the outcomes achieved in the capture, reconstruction, and relighting processes of a panel belonging to the retable of St. Bernardino (1455). This polyptych, originally housed in the chapel of St. Bernardino within the St. Francesco church in Cagliari, Italy, is currently held and exhibited at the Pinacoteca Nazionale in Cagliari. The panel (Fig. 3) measures 34x25cm and is executed in oil on a wooden support, portraying the prophet Daniel. The painting was acquired in its pre-restoration state using a free-form setup that comprises a 36.3 Megapixel DSLR FX Nikon D810 Camera equipped with a 50 AF Nikkor Lens and a handheld white LED (5500K) covering the entire visible spectrum. Sixty images were captured for the MLIC. The acquired data underwent

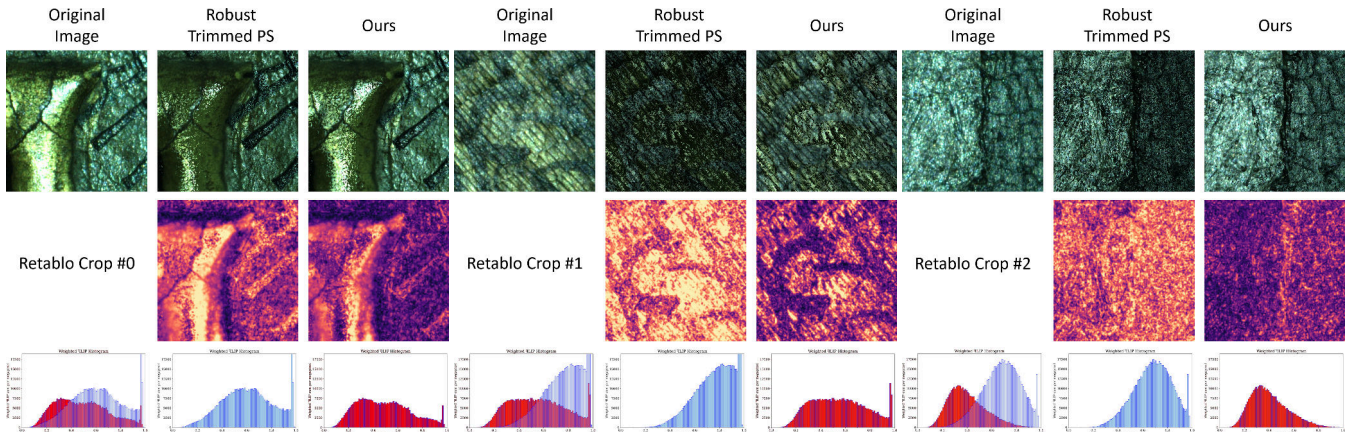


Figure 5: We juxtapose selected original images of the three Retablo’s crops with their digitally relit versions generated from the SV-BRDF calculated using a normal map from a standard Trimmed PS and our refined normal map. The final two rows showcase the respective \mathcal{A} LIP maps and associated statistics. The leftmost plot in the third row of each group of images overlays the two sets of statistics to highlight the improvement more effectively.

calibration using four glossy spheres (for light direction) and a gray frame positioned around the object, employing the camera and light calibration method introduced by Pintus et al. [PJZ*21].

Here, our attention is directed towards three specific image sub-regions/crops of the first panel (Fig. 3), chosen for their distinct characteristics. *Crop #0* encompasses a gold-colored relief featuring several cracks and a flat area with black writing. Typically, dark signals induce distortions in both the normal map and the final re-lighted image. *Crop #1* depicts a decorative detail situated within a predominantly flat area with cracks, while *Crop #2* focuses on a facial detail within a highly glossy region. In most original images, the highlights are so pronounced that the painted face is obscured and overshadowed by the glossy signal, with evident paint mixing and various visible damages.

Synthetic Dataset	Method	\mathcal{A} LIP Statistics			
		Mean	Weighted Median	1st Quartile	3rd Quartile
Retablo Crop #0	<i>Trimmed PS + Ward</i>	0.53	0.61	0.46	0.76
	<i>Trimmed PS + Ours + Ward</i>	0.39	0.47	0.32	0.63
Retablo Crop #1	<i>Trimmed PS + Ward</i>	0.75	0.82	0.68	0.94
	<i>Trimmed PS + Ours + Ward</i>	0.45	0.54	0.37	0.71
Retablo Crop #2	<i>Trimmed PS + Ward</i>	0.64	0.69	0.58	0.79
	<i>Trimmed PS + Ours + Ward</i>	0.35	0.38	0.29	0.50

Table 2: \mathcal{A} LIP statistics for the three retablo images in Fig. 5.

All evaluation metrics applied to these examples consistently demonstrate performance that exceeds previous results (Fig. 4). Some images exhibit a relative improvement in PSNR of over 20%, while SSIM performance shows even greater enhancements, reaching up to an 80% improvement. As observed in prior tests, the improvement is particularly noticeable in front-illuminated images, where highlights or other glossy effects are prevalent. Fig. 5 depicts the rendered images alongside the \mathcal{A} LIP maps and associated statistics. Additionally, Tab. 2 is provided containing the numerical values of the mean, weighted median, 1st, and 3rd quartile of the error. In *Crop #0*, the improvement is evident as the gloss gold region is restored with greater accuracy, while the perceptual error in the flat region, particularly in the darker parts of the writing, is notably reduced. Similar results are observed in *Crop #1* and *Crop #2*, where both the error images and the plots illustrate how the normal refinement significantly enhances the quality of surface characterization and the corresponding virtual relighting.

5. Conclusions

Recovering normal maps and BRDF maps from MLICs permits both object characterization and flexible re-illumination. However, fitting an analytical BRDF to recovered geometric normals may often lead to sub-optimal outcomes for visual inspection in typical relighting applications.

Building on the common and general assumption that the reflectance in a typical normal+BRDF model is a monotonic function of the angle between the normal and the bisector of lighting and viewing directions, we have proposed a solution that corrects the orientation of the geometric normals into shading normals where this assumption is not met. Differently from solutions that force using a coupled shape and material solver, our method integrates seamlessly into existing reconstruction pipelines, supporting a free mixing and matching of Photometric Stereo methods, BRDF models, and BRDF fitters. Users can decide whether to use shading normals also as improved geometric normals, especially when using the most common normal recovery techniques (e.g., based on trimming), or, more commonly, separate the normal maps used for geometric analysis and visualization.

In this paper, we have provided a description of the method and the initial tests. A more comprehensive description and evaluation will be included in a forthcoming journal publication. We are also currently working on further enhancing the approach, in particular by reducing its computation time by exploiting the similarity of nearby pixels, improving the global optimization solver, and streamlining the computation of the cost function.

Acknowledgments We thank the National Archaeological Museum and the National Art Gallery in Cagliari for access to artworks for digitization and fruitful collaboration. This work received funding from Sardinian regional authorities under the XDATA project (art 9 L.R. 20/2015). We also acknowledge the contribution of project “REFLEX – REFlectance EXploration: improving the acquisition, distribution, and exploration of multi-light image collections for surface characterization and analysis” (PRIN2022) funded by the EU Next-GenerationEU PNRR, component M4C2-1.1.

References

- [ANA*20] ANDERSSON P., NILSSON J., AKENINE-MÖLLER T., OSKARSSON M., ÅSTRÖM K., FAIRCHILD M. D.: FLIP: A Difference Evaluator for Alternating Images. *Proc. ACM SIGGRAPH* 3, 2 (2020), 15:1–15:23. doi:10.1145/3406183.4
- [AZK08] ALLDRIN N., ZICKLER T., KRIEGMAN D.: Photometric stereo with non-parametric and spatially-varying reflectance. In *Proc. CVPR* (2008), IEEE, pp. 1–8. doi:10.1109/CVPR.2008.4587656. 1, 2
- [CHS*19] CHEN G., HAN K., SHI B., MATSUSHITA Y., WONG K.-Y. K.: Self-calibrating deep photometric stereo networks. In *Proc. CVPR* (2019), IEEE, pp. 8739–8747. doi:10.48550/arXiv.1903.07366. 2
- [CNK16] CHOE G., NARASIMHAN S. G., KWEON I. S.: Simultaneous estimation of Near IR BRDF and fine-scale surface geometry. In *Proc. CVPR* (2016), IEEE, pp. 2452–2460. doi:10.1109/CVPR.2016.269. 2
- [Eve24] EVERY TEXTURE: 3D texture database, 2024. [Online; accessed-July-22-2024]. URL: <https://everytexture.com/>. 3
- [GCHS05] GOLDMAN D., CURLESS B., HERTZMANN A., SEITZ S.: Shape and spatially-varying BRDFs from photometric stereo. In *Proc. ICCV* (2005), IEEE, pp. 341–348 Vol. 1. doi:10.1109/ICCV.2005.219. 1
- [Geo03] GEORGHIADES A.: Incorporating the Torrance and Sparrow model of reflectance in uncalibrated photometric stereo. In *Proc. ICCV* (2003), IEEE, pp. 816–823. doi:10.1109/ICCV.2003.1238432. 2
- [GGG*16] GUARNERA D., GUARNERA G. C., GHOSH A., DENK C., GLENCROSS M.: BRDF representation and acquisition. *Computer Graphics Forum* 35, 2 (2016), 625–650. doi:10.1111/cgf.12867. 2
- [GKPB04] GAUTRON P., KRIVÁNEK J., PATTANAİK S. N., BOUATOUCH K.: A novel hemispherical basis for accurate and efficient rendering. *Rendering Techniques* (2004), 321–330. doi:10.2312/EGWR/EGSR04/321-330. 2
- [Gri06] GRIMMETT G.: *The random-cluster model*, vol. 333. Springer, 2006. doi:10.1007/978-3-540-32891-9. 2
- [HLHZ08] HOLROYD M., LAWRENCE J., HUMPHREYS G., ZICKLER T.: A photometric approach for estimating normals and tangents. *ACM TOG* 27, 5 (2008), 1–9. doi:10.1145/1457515.1409086. 2
- [HMI10] HIGO T., MATSUSHITA Y., IKEUCHI K.: Consensus photometric stereo. In *Proc. CVPR* (2010), IEEE, pp. 1157–1164. doi:10.1109/CVPR.2010.5540084. 2
- [HS05] HERTZMANN A., SEITZ S. M.: Example-based photometric stereo: Shape reconstruction with general, varying BRDFs. *IEEE TPAMI* 27, 8 (2005), 1254–1264. doi:10.1109/TPAMI.2005.158. 2
- [HS16] HUI Z., SANKARANARAYANAN A. C.: Shape and spatially-varying reflectance estimation from virtual exemplars. *IEEE TPAMI* 39, 10 (2016), 2060–2073. doi:10.1109/TPAMI.2016.2623613. 1, 2
- [Ike18] IKEHATA S.: CNN-PS: CNN-based photometric stereo for general non-convex surfaces. In *Proc. ECCV* (2018), Springer, pp. 3–18. doi:10.1007/978-3-030-01267-0_1. 3
- [JLX*24] JU Y., LAM K.-M., XIE W., ZHOU H., DONG J., SHI B.: Deep learning methods for calibrated photometric stereo and beyond: A survey. *arXiv preprint arXiv:2212.08414* (2024), 1–19. doi:10.48550/arXiv.2212.08414. 2
- [LBFS21] LAVOUÉ G., BONNEEL N., FARRUGIA J.-P., SOLER C.: Perceptual quality of BRDF approximations: dataset and metrics. *Computer Graphics Forum* 40, 2 (2021), 327–338. doi:10.1111/cgf.142636. 2
- [Mat24] MATHWORKS: monotonicity, Quantify monotonic trend in condition indicators, 2024. [Online; accessed-January-15-2024]. URL: <https://www.mathworks.com/help/predmaint/ref/monotonicity.html>. 2
- [MGW01] MALZBENDER T., GELB D., WOLTERS H.: Polynomial Texture Maps. In *Proc. SIGGRAPH* (2001), ACM, pp. 519–528. doi:10.1145/383259.383320. 2
- [NNSK99] NEUMANN L., NEUMANN A., SZIRMAY-KALOS L.: Compact metallic reflectance models. *Computer Graphics Forum* 18, 3 (1999), 161–172. doi:10.1111/1467-8659.00337. 3
- [OKT*19] ONO T., KUBO H., TANAKA K., FUNATOMI T., MUKAIGAWA Y.: Practical BRDF reconstruction using reliable geometric regions from multi-view stereo. *Computational Visual Media* 5 (2019), 325–336. doi:10.1007/s41095-019-0150-3. 2
- [PAZ*23] PINTUS R., AHSAN M., ZORCOLO A., BETTIO F., MARTON F., GOBBETTI E.: Exploiting local shape and material similarity for effective SV-BRDF reconstruction from sparse Multi-Light Image Collections. *ACM JOCCH* 16, 2 (2023), 39:1–39:31. doi:10.1145/3593428.2. 3
- [PDC*19] PINTUS R., DULECHA T. G., CIORTAN I., GOBBETTI E., GIACHETTI A.: State-of-the-art in Multi-Light Image Collections for surface visualization and analysis. *Computer Graphics Forum* 38, 3 (2019), 909–934. doi:10.1111/cgf.13732. 1, 2
- [PJZ*21] PINTUS R., JASPE VILLANUEVA A., ZORCOLO A., HADWIGER M., GOBBETTI E.: A practical and efficient model for intensity calibration of Multi-Light Image Collections. *The Visual Computer* 37, 9 (September 2021), 2755–2767. doi:10.1007/s00371-021-02172-9. 5
- [RK09] RUITERS R., KLEIN R.: Heightfield and spatially varying BRDF reconstruction for materials with interreflections. *Computer Graphics Forum* 28, 2 (2009), 513–522. doi:10.1111/j.1467-8659.2009.01390.x. 2
- [Rus98] RUSINKIEWICZ S. M.: A new change of variables for efficient BRDF representation. In *Rendering Techniques* (1998), Springer, pp. 11–22. doi:10.1007/978-3-7091-6453-2_2. 3
- [SSS*17] SANTO H., SAMEJIMA M., SUGANO Y., SHI B., MATSUSHITA Y.: Deep photometric stereo network. In *Proc. ICCV workshops* (2017), IEEE, pp. 501–509. doi:10.1109/ICCVW.2017.66. 2
- [STMI12] SHI B., TAN P., MATSUSHITA Y., IKEUCHI K.: Elevation angle from reflectance monotonicity: Photometric stereo for general isotropic reflectances. In *Proc. ECCV* (2012), Springer, pp. 455–468. doi:10.1007/978-3-642-33712-3_33. 2
- [STMI13] SHI B., TAN P., MATSUSHITA Y., IKEUCHI K.: Bi-polynomial modeling of low-frequency reflectances. *IEEE TPAMI* 36, 6 (2013), 1078–1091. doi:10.1109/TPAMI.2013.196. 2
- [SWM*16] SHI B., WU Z., MO Z., DUAN D., YEUNG S.-K., TAN P.: A benchmark dataset and evaluation for non-Lambertian and uncalibrated photometric stereo. In *Proc. CVPR* (2016), IEEE, pp. 3707–3716. doi:10.1109/TPAMI.2018.2799222. 3
- [Wal05] WALTER B.: *Notes on the Ward BRDF*. Tech. Rep. PCG-05, Program of Computer Graphics, Cornell University, 2005. 3
- [WGT*05] WENGER A., GARDNER A., TCHOU C., UNGER J., HAWKINS T., DEBEVEC P.: Performance relighting and reflectance transformation with time-multiplexed illumination. *ACM TOG* 24, 3 (2005), 756–764. doi:10.1145/1073204.1073258. 3
- [Woo80] WOODHAM R. J.: Photometric method for determining surface orientation from multiple images. *Optical engineering* 19, 1 (1980), 139–144. doi:10.1117/12.7972479. 2
- [WRG*23] WANG F., REN J., GUO H., REN M., SHI B.: DiLiGenT-Pi: Photometric stereo for planar surfaces with rich details-benchmark dataset and beyond. In *Proc. ICCV* (2023), IEEE, pp. 9477–9487. doi:10.1109/ICCV51070.2023.00869. 2