Vision, Modeling, and Visualization (2019)
H.-J. Schulz, M. Teschner, and M. Wimmer (Eds.)

Normal Map Bias Reduction for Many-Lights Multi-View Photometric Stereo

J. Gan[1], P. Bergen[1], T. Thormählen[1], P. Drescher[2], and R. Hagens[2]

[1]Philipps-Universität Marburg, Germany
[2]Beiersdorf AG, Germany

Abstract

In this paper, we improve upon an existing many-lights multi-view photometric stereo approach. Firstly, we show how to detect continuous regions for normal integration, which leads to a fully automatic reconstruction pipeline. Secondly, we compute per-pixel light source visibilities using an initial biased reconstruction in order to update the estimated normal map to a solution with reduced bias. Thirdly, to further improve the normal accuracy, we compensate for interreflections of light between surface locations. Our approach is evaluated on both synthetic and real-world data and it is shown that the normal accuracy is improved by around 50 percent.

CCS Concepts

• ***Computing methodologies*** → *Reconstruction;*

1. Introduction

Photometric stereo [Woo80] is a well-established technique in computer vision and is able to recover high-resolution normal maps of the object from multiple images taken under changing lighting conditions. Especially, many-lights approaches, such as the Light Stage approach by Debevec [Deb12], are commonly used in practice because the computation of the normal map is mathematically simple and fast, and produces visually convincing results.

On the other hand, low-frequency bias has been a well-known drawback of normal integration and approaches have been proposed to address the problem by combining photometric stereo with other triangulation techniques, such as multi-view stereo [NRDR05, EVC08].

While such hybrid approaches can achieve detailed normal maps and reconstructions of high visual quality, it is often overlooked that the normal estimates are not very accurate compared to the ground truth. Some research has focused on bias-reduction in photometric stereo in order to estimate accurate normal maps [WGS*10, IWMA14, QDW*18], but typically setups with few lights are considered.

In this paper, we propose an approach to reduce the bias of normal maps in the presence of many lights by considering occlusion and interreflections in an iterative process. We first estimate an initial biased normal map and then automatically select continuous regions for normal integration in the camera view, which results in an initial 3D model. Based on the initial 3D model, we compute per-pixel light source visibilities for our many-lights illumination setup (consisting of 2200 LEDs). Furthermore, we compensate for 1-bounce interreflections of light between different locations on the model. Once a refined solution is obtained, the process is repeated for a few times. Figure 1 gives an overview of our system, which extends the deformation-based approach adopted from [GWT*18].

The main contributions of this paper are:

- An approach to detect large discontinuities by normal integration and residual analysis
- An approach to reduce normal map bias in the presence of many light sources by taking per-pixel light visibilities into account
- An approach to reduce normal map bias by compensating for surface interreflections in a many-lights setup

2. Related Work

Discontinuities Detection In order to recover the underlying surface from the normal map, the integrability constraint must be satisfied. However, the constraint is typically violated for discontinuities, such as occlusion boundaries. Therefore, integrable regions are usually marked manually in most systems.

Some approaches are proposed to automatically find depth discontinuities. Raskar et al. [RTF*04] look for shape boundaries by using a camera that is flashed in different arranged positions. In [WBL*12], Wang et al. first apply edge detection on an intensity image and then identify discontinuities by two SVM classifiers, which are trained from hundreds of manually-labelled continuous and discontinuous edges. Other approaches try to compensate the defect of discontinuities by normal integration. Karaçali et al. [KS02] use a minimum norm solution to incorporate

DOI: 10.2312/vmv.20191314

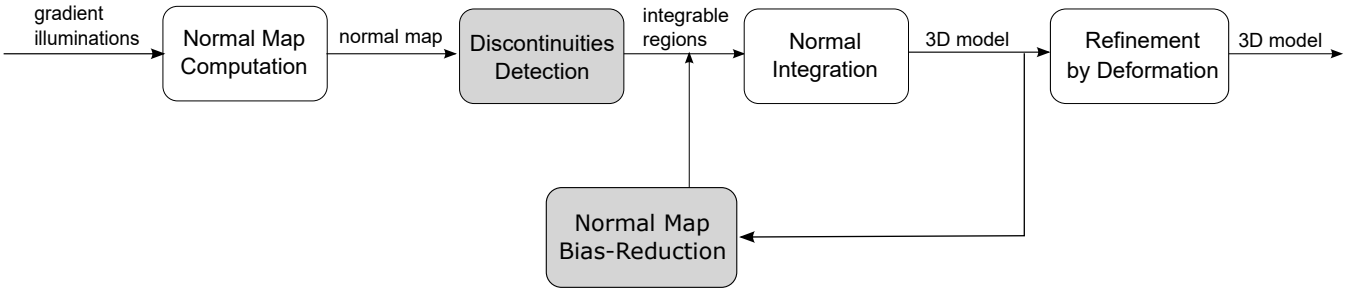

Figure 1: *The reconstruction pipeline of the proposed multi-view photometric stereo system. The steps marked in gray are main contributions of this paper.*

depth discontinuities. Wu et al. [WT06] present an EM algorithm to enforce discontinuity-preserving integrability for surface reconstruction. Quéau et al. [QDA18] give a thorough review of normal integration in the presence of discontinuities.

In contrast to the above work, our approach automatically avoids large discontinuities by performing region-based segmentation on the residual of the initial normal map integration and later either reconstructs each region separately or keeps only the largest region.

Occlusions and interreflections Dealing with occlusions and global illumination effects, such as interreflections, is a key challenge in photometric stereo. In order to address the problem, various approaches have been proposed. In [CAK07], Chandraker et al. propose a graph cut-based approach to remove shadows. They assume every pixel is visible to at least three light sources in order to achieve an initial shadow labelling. Instead, we use ray tracing to find the initial per-pixel visibilities for every light source using the initial (biased) reconstruction. In [EVC08], Hernández et al. use a RANSAC scheme to identify outliers as shadows. RANSAC is also used in [SZP10] by Sunkavalli et al., but visibility subspaces are estimated instead of per-pixel measurements to detect both attached and cast shadows in uncalibrated photometric stereo. Different from explicitly handling deviations, some other work by Wu et al. [WGS*10] and Ikehata et al. [IWMA14] focus on trying to detect any type of deviation of a Lambertian model and then apply photometric stereo only to the remaining intensities. Recently, deep learning based approaches [Ike18, TM18] have been introduced to photometric stereo that assume known directional lighting, but require pre-processing for the light calibration. In [CHS*19], Chen et al. handle the problem of uncalibrated photometric stereo by explicitly learning to estimate light directions.

However, all approaches above assume that each light source is turned on separately, which is not applicable for many-lights systems because far too many images would need to be captured. Typically, many-lights systems turn on a few lights at the same time and generate a certain illumination pattern, e.g., spherical gradient illumination patterns as proposed by Ma et al. [MHP*07]. Our approach can deal with the presence of many lights because it uses ray tracing to find per-pixel visibilities for every light source based on the initial mesh generated from normal integration.

3. Normal Map Bias Reduction

In this section, we present our algorithm for discontinuities detection as well as occlusions and interreflections bias reduction. Our many-lights illumination setup is a lighting sphere, which consists of 2200 LEDs. Objects are illuminated under six spherical gradient illumination patterns, which are axis parallel and complementary in pairs, and are captured synchronously by six calibrated digital SLR cameras from multiple viewpoints.

3.1. Discontinuities Detection

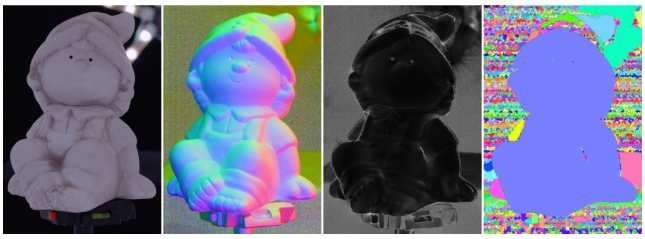

Figure 2: *Discontinuities detection for the dataset Dwarf. From left to right: intensity image, color-coded normal map, residual map, and watershed segmentation for the residual map.*

Input to our algorithm is a set of six gradient illumination images for each camera, which are denoted by $\{L^x, L^{-x}, L^y, L^{-y}, L^z, L^{-z}\}$. Based on the input images, we estimate photometric normals $\mathbf{n}$ following the approach by Wilson et al. [WGP*10]:

$$\mathbf{n} = \frac{(L^x - L^{-x}, L^y - L^{-y}, L^z - L^{-z})^\top}{\|(L^x - L^{-x}, L^y - L^{-y}, L^z - L^{-z})^\top\|}. \quad (1)$$

Given the normal map, we then integrate an initial depth map by the linear approach proposed by Grochulla and Thormählen [GT15]. Starting from an arbitrary point i on the normal map, the relation between the depth d_i at point i and the depth d_j at adjacent points j should be given by

$$d_j = \frac{\mathbf{l}_i \cdot \mathbf{n}_i}{\mathbf{l}_j \cdot \mathbf{n}_i} d_i \quad (2)$$

where $\mathbf{l} \in \mathbb{R}^3$ is the line of sight from the camera and $\mathbf{n}_i$ is the normal at point i.

Since effects such as image noise, depth discontinuities, shadows, interreflections, etc., are not accounted for in Eq. (1), the condition in Eq. (2) is typically not perfectly fulfilled. To quantitatively analyze the residual of Eq. (2), we formulate an deviation energy over the four neighbors of a point i:

$$\mathbf{e}_i = \sum_j (\mathbf{l}_j \cdot \mathbf{n}_j \, d_j - \mathbf{l}_i \cdot \mathbf{n}_i \, d_i)^2 \tag{3}$$

Using Eq. (3), we are able to compute a residual value for each pixel of the normal map. The residual is then used to heuristically determine reconstruction regions for the scene by applying watershed segmentation. Figure 2 shows the detection results for the *Dwarf* object. The far right image in Figure 2 shows the result of watershed segmentation for the residual map. When using the normal information directly, it is difficult to detect some large discontinuities such as occlusions (see Figure 5). This is because occlusions that cause a discontinuity in depth might not have a discontinuity in the normal map. In contrast, with the residual map generated by our approach, we are able to detect discontinuities and segment the scene into several large integrable regions. In our experiments, the maximum region is always selected for further processing. Furthermore, different integrable regions from multiple views can be merged into a complete reconstruction. The details of the region merging algorithm are beyond the scope of this paper.

3.2. Occlusions Bias Reduction

In most many-lights systems, the visibilities of light sources are not considered. For Eq. (1) it is assumed in [WGP*10] that the point on the surface has a hemispherical field of view but in reality this requirement is only met for very special objects, such as a perfect sphere. For more complex objects, self-shadowing typically occurs with respect to certain light sources. As shown in Figure 3, the visibility range for a surface point in a concavity is much less than a hemisphere.

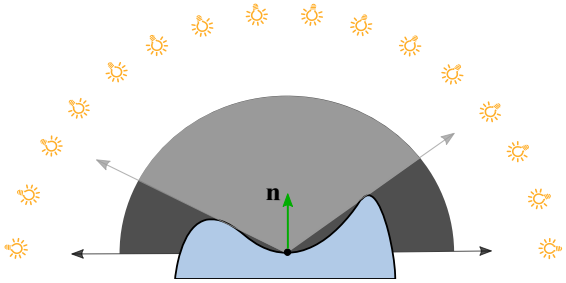

Figure 3: *True light source visibilities for a surface point in a concavity (light gray) vs. assumed full visibility of the hemisphere (dark gray).*

Given a surface point (x,y) in the input intensity image, for each light source with index k we denote the light direction as $\mathbf{d}_k$ and light intensity as L_k. The light visibility V_k is either 0 or 1. The observed intensity can be indicated as $I(x,y)$. We compute the initial per-pixel visibility of each light source by shooting shadow rays to all light sources from the surface point of the initial 3D model. Given the above information and the normal $\mathbf{n}$ from the normal map, we can formulate

$$I(x,y) = \rho \sum_k \mathbf{n}^T \mathbf{d}_k L_k V_k, \tag{4}$$

where ρ is the unknown surface albedo and the radiometric response function of the camera is assumed to be linear. Furthermore, we can extract the normal from the sum:

$$I(x,y) = \rho \mathbf{n}^T \underbrace{\sum_k \mathbf{d}_k L_k V_k}_{\mathbf{D}}, \tag{5}$$

where $\mathbf{D}$ is a 3×1 column vector and we get a different $\mathbf{D}_g$ for each gradient illumination g. With a total of six gradient illuminations we get an overdetermined system of linear equations, which can be solved for the three unknowns $\rho\mathbf{n} = (\rho n_x, \rho n_y, \rho n_z)^T = (u_x, u_y, u_z)^T$. Afterwards, we can use the non-linear constraint that the length of the normal $\mathbf{n}$ is 1.0 to eliminate the unknown albedo ρ:

$$\mathbf{n} = \frac{(u_x, u_y, u_z)^T}{||(u_x, u_y, u_z)||}. \tag{6}$$

Solving the overdetermined system of linear equations for each pixel gives us an updated normal map, which in turn can be used to compute updated light visibilities.

3.3. Interreflections Bias Reduction

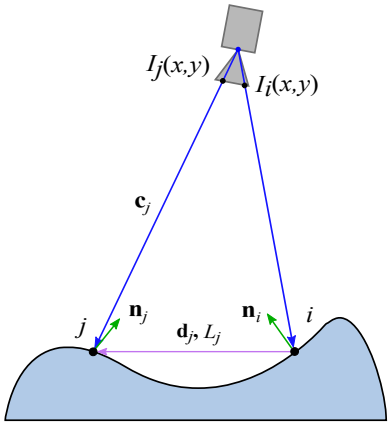

Figure 4: *Interreflections between two surface points.*

When light interacts with matter it can be absorbed, transmitted or reflected. In this section, we compensate for the reflected light between surface points, i.e. interreflections. Given two pixels i and j, where j is the emitting pixel, the geometric setup we are dealing with is shown in Figure 4. When considering interreflections, the observed intensity $I(x,y)$ for a pixel now not only depends on the reflected part of the sum of the direct light from all light sources k, but additionally on the sum of the reflected light that is received from all other surface locations j. Thus, Eq. (4) becomes

$$I(x,y) = \rho \sum_k \mathbf{n}^T \mathbf{d}_k L_k V_k + \rho \sum_j \mathbf{n}^T \mathbf{d}_j L_j V_j, \tag{7}$$

where $\mathbf{d}_j$ is the 3D direction between the two interacting 3D points (see Figure 4) and V_j is the visibility, which is either 0 or 1, and encodes if the two points i and j see each other (i.e., they are not occluded by other geometry). When modelling the patch at location j as a small Lambertian emitter, the light intensity L_j reaching location i from j depends on the normal $\mathbf{n}_j$ at location j and the distance $r_{\mathbf{d}_j}$ between surface points j and i. It can be approximated by:

$$L_j \approx \frac{-\mathbf{n}_j^T \mathbf{d}_j}{r_{\mathbf{d}_j}^2} L_{\mathbf{n}_j}, \tag{8}$$

where $L_{\mathbf{n}_j}$ is the light intensity in direction $\mathbf{n}_j$. Because a Lambertian emitter has a constant luminance, the light intensity $L_{\mathbf{n}_j}$ is proportional to the intensity I_j observed in the input image at pixel j. Introducing a global constant of proportionality α, Eq. (7) can be rewritten as follows:

$$I(x,y) = \rho \mathbf{n}^T \underbrace{\left(\sum_k \mathbf{d}_k L_k V_k + \alpha \sum_j \mathbf{d}_j \frac{(-\mathbf{n}_j^T \mathbf{d}_j)}{r_{\mathbf{d}_j}^2} I_j V_j \right)}_{\mathbf{D}}. \tag{9}$$

Similar as before, we solve for the normal $\mathbf{n}$ using the overdetermined system of linear equations. Normalization of $\mathbf{n}$ eliminates the unknown albedo ρ. All visibilities V_k and V_j, normals $\mathbf{n}_j$, directions $\mathbf{d}_j$, and distances $r_{\mathbf{d}_j}$ are determined using the previous reconstruction. These input values are kept constant during each iteration and the solution is independently computed in parallel for each normal map pixel i.

4. Results

In this section, we first present the results of the proposed discontinuities detection approach for two real-world objects: *Dwarf* and *Face1*, then evaluate our approach for normal map bias reduction on both synthetic and real-world image data. In order to quantitatively evaluate the approach, we use two synthetic objects, *Ref1* and *Face2*, in the experiments. We also test our approach on two real-world objects: *Ref2* and *Budda*. All objects are recovered by our approach and the one by Wilson et al. [WGP*10] (referred to as *baseline* approach in the following), which neglects the compensation of occlusions and interreflections. In all experiments we use six digital SLR cameras and each camera generates six images with a resolution of 5496×3670 pixels taken under gradient illuminations. For each object we perform 3 iterations to find a refined solution.

4.1. Discontinuities Detection

The results of the proposed discontinuities detection approach are shown in Figure 5. For comparison, Canny edge detection and watershed segmentation are applied on the intensity image, the normal map, and the residual map. The results show that for the intensity image and the normal map, fewer depth discontinuities are detected, e.g., the nose in the *Face1* example occludes parts of the right side of the face. In contrast, if the residual map is used as input, all large occlusions are found.

Table 1: *Comparison of the estimated to the ground truth normal map (RMSE)*

Methods	*Face2*	*Ref1*	*Ref2*
baseline	0.048	0.046	0.074
occlusion comp.	0.044	0.041	0.065
proposed	0.027	0.021	0.034

4.2. Bias Reduction on Synthetic Data

In this section, we generate six synthetic images of an object under simulated gradient illuminations and apply our approach and the *baseline* approach on the synthetic data to estimate normal maps of the object. The ground truth normal maps are known exactly for the synthetic objects, thus, we can quantitatively verify the accuracy of both approaches by computing the normal errors between the estimated normal map and the ground truth.

Figure 6 shows the ground truth normal map and the normal maps recovered by the different approaches. For the dataset *Ref1* the magnifications of three grooves with different shapes are shown. The normal map estimated by the *baseline* approach shows a bias that is significantly reduced by our approach. Similarly, for the synthetic *Face2* example the normals around the nose also show obvious bias for the *baseline* approach, which is reduced by our approach. For both approaches, the errors inside the nostrils are still present because almost complete shadowing occurs.

Table 1 shows the pixel-wise root-mean-square error (RMSE) of the estimated normal maps. The results show that the normal map errors are reduced by around 50% compared to the *baseline* approach.

4.3. Bias Reduction on Real-World Data

In this section, we test our approach on two real-world objects: *Ref2* and *Budda*. Dataset *Ref2* is an aluminum cuboid with two concave 45 degree grooves on the surface with a depth of 5.0 and 2.5 millimeters, which were cut with high accuracy by a CNC machine. The aluminum cuboid is covered with white chalk spray to achieve a diffuse surface. The results are shown in Figure 7. The RMSE of the recovered normal map of *Ref2* is given in Table 1. It shows that the normal bias is reduced by 54% with our approach compared to the *baseline* approach. For the *Budda* dataset, the ground truth of the normal map is unknown. We demonstrate the difference between the two models reconstructed by the baseline and our approach in Figure 8. The pseudo-color map (where blue corresponds to no difference and red to a difference of 0.29 mm) shows plausible refinements of the reconstruction by our approach, especially for the grooves.

4.4. Limitation

In our approach, we use ray tracing techniques to find per-pixel light source and surface-surface visibilities. This leads to a large computational effort. Using a custom CUDA ray-tracing implementation that employs a bounding volume hierarchy (BVH) as

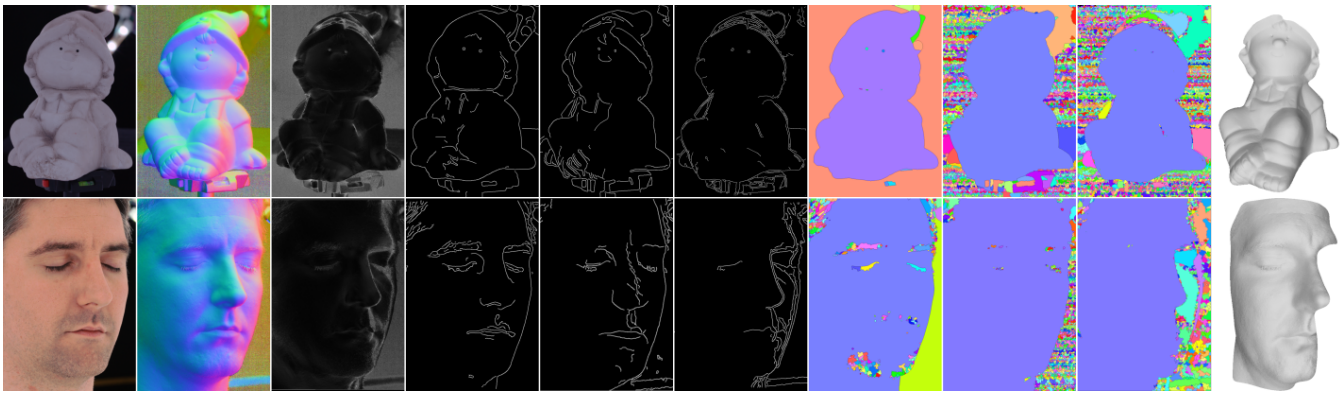

Figure 5: *Discontinuities detection for the datasets Dwarf and Face1. From left to right: intensity image, color-coded normal map, residual map, Canny edge detection and watershed segmentation for the three different input images, and 3D reconstruction of the largest region detected by our proposed discontinuities detection approach.*

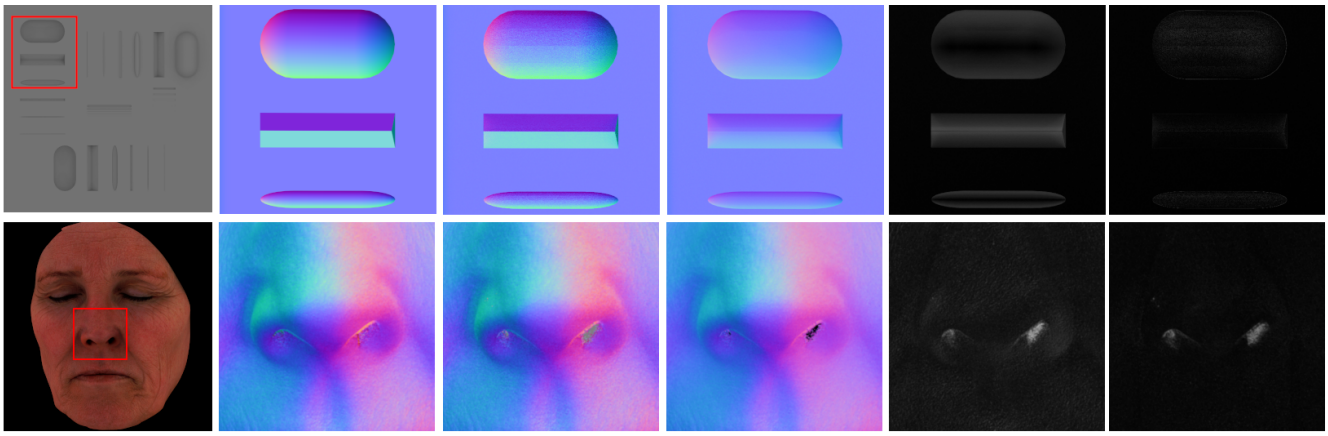

Figure 6: *Normal maps recovery and comparison for the datasets Ref1 and Face2. From left to right: intensity image, ground truth normal map, normal map estimated by our approach, normal map estimated by the baseline approach, difference between the baseline approach and ground truth, and difference between our approach and ground truth.*

an acceleration structure to compute ray-surface intersections, the computation time for the shown examples is reduced to approximately 2 to 5 hours depending on the target area and the graphics card. These large computation times make the approach less applicable in practice. Furthermore, as can be best observed in the *Ref1* example in Figure 6, our bias reduction approach produces a smaller RMSE compared to the baseline approach, but this comes at the price of an increased noise level (that is introduced by the employed ray tracing techniques).

We use an initial biased reconstruction to update the normals to a solution with less bias. We also have tested our data using the ground truth geometry as the reference to improve the normal maps and observed that similar RMSE improvements and noise amplifications occur. This indicates that using a biased initialization does not influence the output strongly. In general, there is no formal guarantee that our iterative process will converge. However, we encountered not a single convergence problem in our experiments.

Our approach assumes a surface of diffuse material. Handling of more complex material properties (specular components, subsurface scattering, etc.) is not considered in this work.

Futhermore, the proposed solution for the reduction of normal map bias is partly specific to systems that employ many-light illumination patterns (such as [WGP*10,GT15]). For other approaches (e.g., learning-based methods [MHP*19] or shape-from-shading approaches [LSHG16]) different model violations can be observed and must be addressed specifically.

The residual in the discontinuities detection might be caused by discontinuities and additional model violations (such as non-diffuse surface reflections or interreflections). Therefore, the approach is only applicable if the effect of the additional violations is below the threshold parameter of the segmentation algorithm.

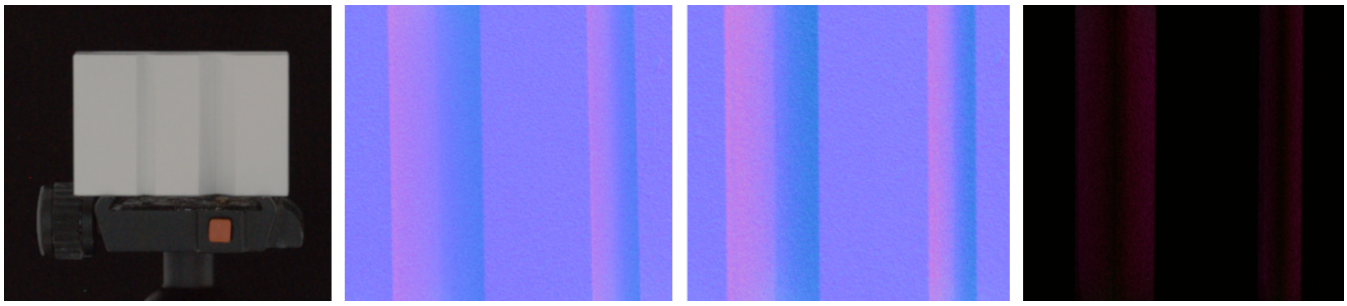

Figure 7: *Normal map recovery for the dataset* Ref2*. From left to right: intensity image, normal map estimated by the* baseline *approach, normal map estimated by our approach, and the difference between the two normal maps.*

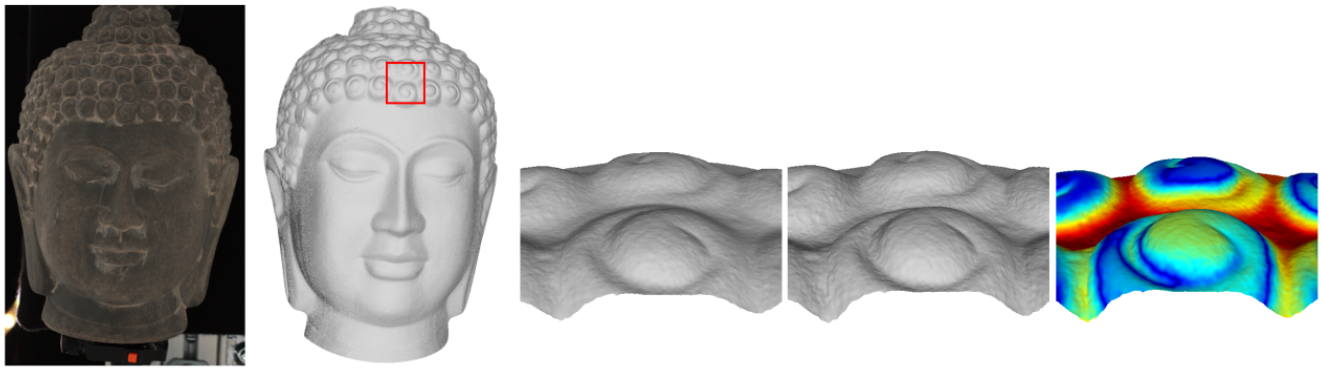

Figure 8: *Normal map recovery for the dataset* Budda*. From left to right: intensity image, initial reconstruction estimated by the* baseline *approach, detail magnifications of the 3D model estimated by the* baseline *approach and our approach, and difference between the two models in pseudo-color. Our approach especially improves the reconstruction in areas with concavities.*

4.5. Conclusion

This paper presents a multi-view photometric stereo approach for reducing normal map bias in many-lights systems. Known approaches for few lights are not applicable for many-lights systems because they typically require to capture a separate image for each light source, which is not feasible for a many-lights setup (e.g., our illumination hardware has 2200 lights). Compared with other many-lights approaches, we automatically detect continuous regions and take the occlusions and interreflections into account. Our approach is evaluated on both synthetic and real-world data and compared to a *baseline* approach. The RMSE is reduced by approximately 50%. The results on real-world data also show similar improvements in accuracy and plausible refinements, especially for areas with concavities.

Many-lights systems are commonly used in practice to generate 3D reconstructions for movies and games but when occlusions and interreflections are not considered the results might look visually plausible but are not close to the ground truth. Reducing the estimation bias is possible (as shown in this paper) but our approach requires a large computational effort of several hours even on modern GPUs. Thus, an interesting approach for future research would be to develop a many-lights acquisition setup that allows producing accurate results with less computation.

References

[CAK07] CHANDRAKER M. K., AGARWAL S., KRIEGMAN D. J.: Shadowcuts: Photometric stereo with shadows. In *2007 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2007), 18-23 June 2007, Minneapolis, Minnesota, USA* (2007). 2

[CHS*19] CHEN G., HAN K., SHI B., MATSUSHITA Y., WONG K. K.: Self-calibrating deep photometric stereo networks. *CoRR abs/1903.07366* (2019). arXiv:1903.07366. 2

[Deb12] DEBEVEC P.: The light stages and their applications to photoreal digital actors. In *SIGGRAPH Asia Technical Briefs* (2012). 1

[EVC08] ESTEBAN C. H., VOGIATZIS G., CIPOLLA R.: Multiview photometric stereo. *IEEE Trans. Pattern Anal. Mach. Intell. 30*, 3 (2008), 548–554. 1, 2

[GT15] GROCHULLA M., THORMÄHLEN T.: Combining photometric normals and multi-view stereo for 3d reconstruction. In *Proceedings of the 12th European Conference on Visual Media Production, London, United Kingdom, November 24-25, 2015* (2015), pp. 10:1–10:8. 2, 5

[GWT*18] GAN J., WILBERT A., THORMÄHLEN T., DRESCHER P., HAGENS R.: Multi-view photometric stereo using surface deformation. *The Visual Computer 34*, 11 (2018), 1551–1561. 1

[Ike18] IKEHATA S.: CNN-PS: cnn-based photometric stereo for general non-convex surfaces. In *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part XV* (2018), pp. 3–19. 2

[IWMA14] IKEHATA S., WIPF D. P., MATSUSHITA Y., AIZAWA K.: Photometric stereo using sparse bayesian regression for general diffuse

surfaces. *IEEE Trans. Pattern Anal. Mach. Intell. 36*, 9 (2014), 1816–1831. 1, 2

[KS02] KARAÇALI B., SNYDER W. E.: Partial integrability in surface reconstruction from a given gradient field. In *Proceedings of the 2002 International Conference on Image Processing, ICIP 2002, Rochester, New York, USA, September 22-25, 2002* (2002), pp. 525–528. 1

[LSHG16] LANGGUTH F., SUNKAVALLI K., HADAP S., GOESELE M.: Shading-aware multi-view stereo. In *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part III* (2016), pp. 469–485. 5

[MHP*07] MA W., HAWKINS T., PEERS P., CHABERT C., WEISS M., DEBEVEC P. E.: Rapid acquisition of specular and diffuse normal maps from polarized spherical gradient illumination. In *Proceedings of the Eurographics Symposium on Rendering Techniques, Grenoble, France, 2007* (2007), pp. 183–194. 2

[MHP*19] MEKA A., HAENE C., PANDEY R., ZOLLHOEFER M., FANELLO S., FYFFE G., KOWDLE A., YU X., BUSCH J., DOURGARIAN J., DENNY P., BOUAZIZ S., LINCOLN P., WHALEN M., HARVEY G., TAYLOR J., IZADI S., TAGLIASACCHI A., DEBEVEC P., THEOBALT C., VALENTIN J., RHEMANN C.: Deep reflectance fields - high-quality facial reflectance field inference from color gradient illumination. vol. 38. 5

[NRDR05] NEHAB D., RUSINKIEWICZ S., DAVIS J., RAMAMOORTHI R.: Efficiently combining positions and normals for precise 3d geometry. *ACM Trans. Graph. 24*, 3 (2005), 536–543. 1

[QDA18] QUÉAU Y., DUROU J., AUJOL J.: Variational methods for normal integration. *Journal of Mathematical Imaging and Vision 60*, 4 (2018), 609–632. 2

[QDW*18] QUÉAU Y., DURIX B., WU T., CREMERS D., LAUZE F., DUROU J.: Led-based photometric stereo: Modeling, calibration and numerical solution. *Journal of Mathematical Imaging and Vision 60*, 3 (2018), 313–340. 1

[RTF*04] RASKAR R., TAN K., FERIS R. S., YU J., TURK M.: Non-photorealistic camera: depth edge detection and stylized rendering using multi-flash imaging. *ACM Trans. Graph. 23*, 3 (2004), 679–688. 1

[SZP10] SUNKAVALLI K., ZICKLER T. E., PFISTER H.: Visibility subspaces: Uncalibrated photometric stereo with shadows. In *Computer Vision - ECCV 2010, 11th European Conference on Computer Vision, Heraklion, Crete, Greece, September 5-11, 2010, Proceedings, Part II* (2010), pp. 251–264. 2

[TM18] TANIAI T., MAEHARA T.: Neural inverse rendering for general reflectance photometric stereo. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018* (2018), pp. 4864–4873. 2

[WBL*12] WANG Y., BU J., LI N., SONG M., TAN P.: Detecting discontinuities for surface reconstruction. In *Proceedings of the 21st International Conference on Pattern Recognition, ICPR 2012, Tsukuba, Japan, November 11-15, 2012* (2012), pp. 2108–2111. 1

[WGP*10] WILSON C. A., GHOSH A., PEERS P., CHIANG J., BUSCH J., DEBEVEC P. E.: Temporal upsampling of performance geometry using photometric alignment. *ACM Trans. Graph. 29*, 2 (2010), 17:1–17:11. 2, 3, 4, 5

[WGS*10] WU L., GANESH A., SHI B., MATSUSHITA Y., WANG Y., MA Y.: Robust photometric stereo via low-rank matrix completion and recovery. In *Computer Vision - ACCV 2010 - 10th Asian Conference on Computer Vision, Queenstown, New Zealand, November 8-12, 2010, Revised Selected Papers, Part III* (2010), pp. 703–717. 1, 2

[Woo80] WOODHAM R. J.: Photometric method for determining surface orientation from multiple image. *Optical Engineering 19*, 1 (1980), 139–144. 1

[WT06] WU T., TANG C.: Visible surface reconstruction from normals with discontinuity consideration. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2006), 17-22 June 2006, New York, NY, USA* (2006), pp. 1793–1800. 2