3D RECONSTRUCTION OF INTERREFLECTION-AFFECTED SURFACE CONCAVITIES USING PHOTOMETRIC STEREO

Steffen Herbort, Daniel Schugk, and Christian Wöhler
Image Analysis Group, TU Dortmund, Otto-Hahn-Straße 4, 44227 Dortmund, Germany
{steffen.herbort, daniel.schugk, christian.woehler}@tu-dortmund.de

Keywords: 3D surface reconstruction, Photometric Stereo, interreflections, surface concavities

Abstract: Image-based reconstruction of 3D shapes is inherently biased under the occurrence of interreflections, since the observed intensity at surface concavities consists of direct and global illumination components. This issue is commonly not considered in a Photometric Stereo (PS) framework. Under the usual assumption of only direct reflections, this corrupts the normal estimation process in concave regions and thus leads to inaccurate results. For this reason, global illumination effects need to be considered for the correct reconstruction of surfaces affected by interreflections. While there is ongoing research in the field of inverse lighting (i.e. separation of global and direct illumination components), the interreflection aspect remains oftentimes neglected in the field of 3D shape reconstruction. In this study, we present a computationally driven approach for iteratively solving that problem. Initially, we introduce a photometric stereo approach that roughly reconstructs a surface with at first unknown reflectance properties. Then, we show that the initial surface reconstruction result can be refined iteratively regarding non-distant light sources and, especially, interreflections. The benefit for the reconstruction accuracy is evaluated on real Lambertian surfaces using laser range scanner data as ground truth.

1 INTRODUCTION

The problem of image-based shape reconstruction has been examined by Horn (Horn, 1970) using a single image and regularization constraints for “Shape from Shading (SfS)". The Photometric Stereo (PS) principle evolved later on (Woodham, 1980), favorably removing the constraints, but needing at least three intensity measurements per pixel obtained under varying illumination directions. Both, the SfS and the PS setup, assume distant light sources with known illumination direction and intensity, which is nowadays termed “calibrated Photometric Stereo". Additionally, the observed surfaces are assumed to be strictly Lambertian (Lambert, 1760).

The case of *uncalibrated* PS has been examined extensively, starting with Hayakawa’s seminal work (Hayakawa, 1994), where the problem of an ambiguity in the determinable light source position from PS image data becomes exposed. That “Generalized Bas-Relief (GBR) ambiguity" (Belhumeur et al., 1999) was subject to extensive research and has been solved using e.g. the generic viewpoint constraint (Freeman, 1994; Yuille et al., 2000), minimization of the surface albedo distribution entropy (Alldrin et al., 2007), by modeling the illumination conditions using spherical harmonics (Basri et al., 2007), or with special relative light source arrangements (Zhou and Tan, 2010).

In parallel, non-Lambertian surfaces have been approached by different authors. (Ikeuchi, 1981) assume purely specular surfaces, (Nayar et al., 1988; Nayar et al., 1990a) assume hybrid surfaces with Lambertian and superimposed specular components. For these methods, the image formation and object reflection properties become relevant. Well-known reflectance function models are e.g. Lambert (Lambert, 1760), Oren-Nayar (Oren and Nayar, 1994), Blinn (Blinn, 1977), Phong (Phong, 1975), Lafortune (Lafortune et al., 1997), Beckmann-Spizzichino (Beckmann and Spizzichino, 1987), Torrance-Sparrow (Torrance and Sparrow, 1967), Ward (Ward, 1992), Ward-Dür-Geisler-Moroder (Geisler-Moroder and Dür, 2010).

Since PS works robustly on Lambertian surfaces, there has been research regarding the separation of diffuse and specular components, which then allows performing PS on the diffuse component only. These use e.g. polarization (Wolff, 1989; Wolff and Boult, 1991), color (Tan and Ikeuchi, 2005; Thomas and Sugimoto, 2010), or both (Nayar et al., 1997).

The model-based reflectance description has been relaxed by BRDF[1]-independent approaches like (Zickler et al., 2002; Alldrin et al., 2008).
A new formulation of the problem using a projective framework has been published by (Tan and Zickler, 2009; Tankus et al., 2005). Further variations of the PS framework regard e.g. exploiting the effects of non-distant light sources (Iwahori et al., 1990; Clark, 1992).
PS itself determines only pixelwise surface normals or gradients, while absolute depth determination requires a postprocessing step, which was elaborately examined by (Agrawal et al., 2006).
Due to the common assumption of single surface reflections (*direct* illumination) and disregarding higher orders (interreflections, a subset of *global* illumination), it is not possible to infer the underlying surface normal from the observed intensity correctly. Fig. 1 depicts an image separated into direct and global illumination components and illustrates the intensity aggregation in concavities.
Systems for synthetically computing global illumination exist in a large variety. One of the open source ray tracers that additionally employs physically plausible BRDFs, which is crucial for correct real world object representation, is PBRT[2] (Pharr and Humphreys, 2010).

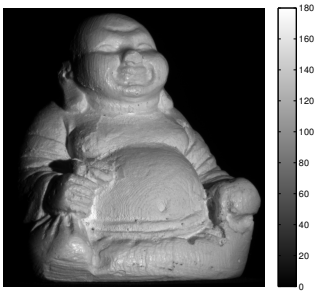

(a) camera image

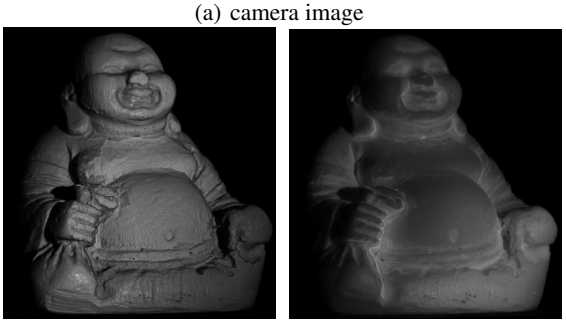

(b) direct component (c) global component

Figure 1: *Measured* illumination components of a camera image using the method by (Seitz et al., 2005) for (b) and a variant of that method with an inverse illumination pattern for (c). Note that all images have the same grayvalue scaling.

While extensive research has been carried out regarding (non-) Lambertian surfaces and uncalibrated PS, there is hardly any research regarding the image-based reconstruction of 3D shapes in the presence of interreflections. Basic thoughts have been published by (Nayar et al., 1990b) who iteratively refine concave v-shapes. The case of inversely computing global illumination with (at most) second order reflections along its path has been examined by (Yu et al., 1999), who use their data for realistic scene modeling and relighting. The effect of interreflections and the resulting inaccuracies on structured light laser scanning profiles have been discussed by (Gupta et al., 2011) and (Couture et al., 2011), but photometric aspects remain untreated therein.
The work by (Gu et al., 2011) stresses the influence of interreflections for PS. Their method is based on multiplexed illumination (Nayar et al., 2006) for the separation of global and direct illumination in the image acquisition stage. They apply high frequency patterns under different incident directions for estimating the two scene illumination components. The direct component is then used to estimate shapes with strongly increased accuracy using PS.
To the authors' best knowledge, there is no publication that actually deals with a method for the reconstruction of shapes with interreflection-affected concavities under illumination with simple point light sources (LEDs).

2 CONSIDERATION OF INTERREFLECTIONS

We propose a novel method that accounts for interreflections in a fully calibrated photometric stereo environment. The approach as such imposes no restrictions upon their reflective behavior and thus allows arbitrary and possibly even measured BRDFs. Furthermore, the light sources are not assumed to be distant, which is not a prerequisite for our algorithm to work but it improves the reconstruction accuracy. Calibration of the (point) light source positions and intensities is achieved using the method described by (Lenoch et al., 2012). Absolute surface depth data is computed using Agrawal's M-estimator approach (Agrawal et al., 2006). In the following, the concatenation of PS and Agrawal's method is denoted "PS+Z".
Our key to the consideration of interreflections lies in

[1]Bidirectional Reflectance Distribution Function

[2]Physically Based Ray-Tracer v2, `pbrt.org`

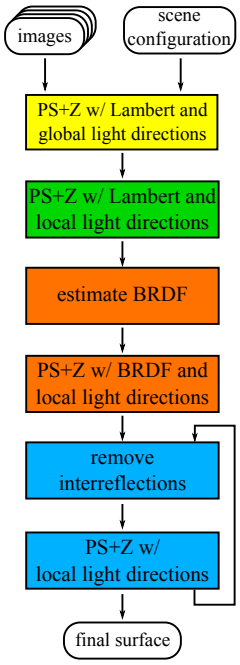

Figure 2: Algorithm overview. “PS+Z" denotes photometric normal estimation and followed by depth integration. Note that each group of components (green, orange, blue) can possibly be iterated.

an iterative scheme: (1) Our method initially reconstructs the surface based on distant light sources and Lambertian reflectance. (2) That result becomes refined assuming non-distant light sources. The resulting surface is afterwards used (3) to obtain the parameters of a parametric BRDF, which additionally refines the surface in another reconstruction step. Lastly (4), the consideration of interreflections is incorporated by determination of the amount of global illumination with a physically plausible rendering system. This allows us to remove a (significant) amount of interreflection-based radiance from the captured scene image data and thus to obtain a refined reconstruction result. Note that it may be beneficial to repeat some of the stages described above in order to ensure their convergence. Fig. 2 gives an overview of these elements.

We demonstrate the potential of our method on real data for a Lambertian and a non-Lambertian surface. The validation of the correctness of these results is achieved on the one hand by comparison with the surface data measured by an industry-standard laser scanner and on the other hand based on synthetically rendered images with a known underlying surface.

In contrast to related publications, we model point light sources and are thus not required to illuminate the scene with specialized projectors like (Gu et al., 2011); furthermore, there is no shape restriction like the v-shapes examined by (Nayar et al., 1990b).

3 ALGORITHM DETAILS

In stage 1 (yellow), the surface normals are computed for each pixel (u,v) based on the standard PS equation for Lambertian surfaces,

$$\vec{I}_{obs}(u,v) = I_0\ \rho(u,v)\ \mathbf{L}\ \vec{n}(u,v), \tag{1}$$

where u and v denote the pixel indices, $\vec{I}_{obs}(u,v) \in \mathbb{R}^{k\times 1}$ denotes the vector of k observed intensity samples under $i = 1...k$ varying (distant) normalized illumination directions $\vec{l}_i$ with $|\vec{l}_i|_2 = 1$. The light source relative radiance is referred to as I_0 and the light source directions are aligned row-wise in $\mathbf{L} = [l_{1x}, l_{1y}, l_{1z}; l_{2x}, l_{2y}, l_{2z}; ...]$. The surface normals are $\vec{n} = [n_x, n_y, n_z]^T$ with $|\vec{n}|_2 = 1$ and ρ denotes the surface albedo. The surface normal is thus determined in the least-mean-square (LMS) sense using

$$\vec{n}(u,v) = \frac{1}{I_0\ \rho(u,v)}\ \mathbf{L}^{-1}\ \vec{I}_{obs}(u,v) \tag{2}$$

with a pseudo-inverse $\mathbf{L}^{-1}$. Note, that locally varying albedos can be obtained when solving for $\vec{n}$ and ρ with known irradiance I_0:

$$\vec{n}^* = \mathbf{L}^{-1}\ \vec{I}_{obs}(u,v), \quad I_0\ \rho = |\vec{n}^*|_2, \quad \vec{n} = \frac{\vec{n}^*}{|\vec{n}^*|_2} \tag{3}$$

The obtained surface normal field is then integrated to absolute depth values using the M-estimator approach described by (Agrawal et al., 2006). For the transformation into a metric depth space, the algorithm is given the approximate object distance from the camera, which suffices for the determination of pixel extents and thus metric depth scaling.

In stage 2 (green), the light sources are regarded as non-distant. This alters the image intensity equation such that

$$\vec{I}_{obs}(u,v) = \frac{I_0\ \rho}{r_1^2(u,v)\ r_2^2(u,v)}\ \mathbf{L}(u,v)\ \vec{n}. \tag{4}$$

Note that we assume locally varying illumination directions $\vec{l}_i(u,v)$ for each pixel (u,v) and each light source $i = 1...k$, which is due to the scene being illuminated by non-distant point light sources. Additionally, each pixel is assigned the light travel distance from the light source to the object (r_1) and from the object to the camera sensor (r_2). Even on this small-scaled framework (object height $\approx 50\,mm$, mean$(r_1) \approx 300\,mm$ and mean$(r_2) \approx 310\,mm$), this shows a significant effect. Up to this point, the surface has been assumed to be perfectly Lambertian and

free of global illumination.
In stage 3 (orange), the BRDF of the underlying material is estimated globally based on the extracted surface from stage 2 and the obtained images. This is achieved by minimizing the RMSE of the intensity data $I = I(u,v)$ and the rendered intensities $R = R(\vec{P},u,v)$ of the surface with respect to the parameter $\vec{P}$ of the chosen BRDF:

$$\vec{P}_{\text{opt}} = \underset{\vec{P}}{\text{argmin}} \sqrt{\frac{1}{U \cdot V} \sum_{u=1}^{U} \sum_{v=1}^{V} [I-R]^2}. \qquad (5)$$

For Lambertian surfaces, the only parameter is the surface albedo ρ, which can be computed by determination of I_0 using (Lenoch et al., 2012) and thus having white balance data available.
In stage 4 (blue), the interreflection component is (iteratively) compensated before starting the PS+Z algorithm: Two separate images of the last stage's PS+Z surface result are rendered under (1) direct $R_{d,i}$ and (2) global $R_{g,i}$ illumination for each light source $i = 1...k$. Afterwards, the global component is determined as the difference between $R_{g,i}$ and $R_{d,i}$, which is used to compensate these in the real camera images I_i:

$$I_{\text{new},i} = I_i - \underbrace{(R_{g,i} - R_{d,i})}_{\text{interreflections}} \qquad i = 1...k \qquad (6)$$

As a last step for this stage, the PS+Z algorithm is started on the new image data basis $I_{\text{new},i}$.

4 EXPERIMENTS AND RESULTS

Figure 3 gives an overview of the setup used for data acquisition. Due to using the same camera for image acquisition and ground truth depth measurement (structured light laser scanning system), it is possible to capture pixel synchronous depth and image data, which facilitates the evaluation process significantly. The description and discussion of our experiments is structured as follows: Initially, we analyze the progressing change in PS+Z depth with increasing algorithm iterations (Sec. 4.1). Afterwards (Sec. 4.2), we show the beneficial effect of our approach qualitatively based on an analysis of the reflectance maps with global illumination.
Sec. 4.3 then shows the comparison of the PS+Z depth result with independently measured scanner depth data z_{RS} for a region of the surface that is especially affected by interreflections. To conclude the evaluation, the aspect of algorithm convergence is discussed in Sec. 4.4 along with suitable criteria that terminate the execution at a reasonable state.

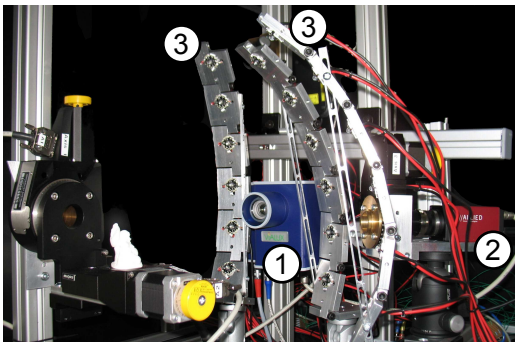

Figure 3: Overview of the experimental setup. The object has a height of about 50 mm, the distance between object and camera (2) / object and illumination (3) amounts to approximately 250 mm. Note that the laser range scanner (1) is used for validation only.

4.1 Progressing change in reconstructed depth

Fig. 4 shows the differences between the naive PS+Z depth $z_{PS,1}$ (no interreflection compensation) and its iterations $z_{PS,2...5}$ using image data with compensated interreflections. It becomes clear that the compensation of the global (interreflection) component produces deeper concavities compared to the uncompensated case. The results show that less deep concavities on the one hand (e.g. neck and ear region) exhibit only a change in depth over the early iterations

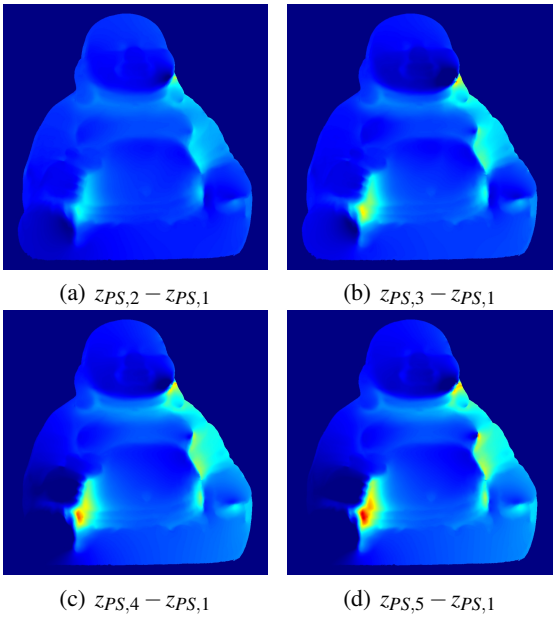
(a) $z_{PS,2} - z_{PS,1}$ (b) $z_{PS,3} - z_{PS,1}$
(c) $z_{PS,4} - z_{PS,1}$ (d) $z_{PS,5} - z_{PS,1}$
Figure 4: Difference between the naive PS depth $z_{PS,1}$ (no interreflection compensation) and the iterations using image data with compensated interreflections $z_{PS,2...5}$. Scale: $-0.5\,mm$ (dark blue) to $2.5\,mm$ (red)

of the algorithm, but remain static on the late iterations. Deep concavities on the other hand progress in their depth over several iterations. This already gives a hint that the algorithm converges suitably, since regions with shallow and deeper concavities converge independently.

4.2 Qualitative evaluation

Fig. 5 shows reflectance maps from different stages of the algorithm. Note that at iteration 5 (Fig. 5(c)), the object surface appears more similar to the physical image (Fig. 5(a)), especially in the interreflection-affected regions. For a human observer, it is obvious

(a) image

(b) stage 4, iteration 1 rendering result (global illumination)

(c) stage 4, iteration 5 rendering result (global illumination)

Figure 5: Rendered PS+Z results. Note the appearance difference at surface concavities.

that the underlying surface resembles the one of the physical image more closely due to the perceived intensities and due to the plausibility of where and how shadows are cast. In summary, the qualitative assessment result of the surface is very favorable, due to fine surface details being visible and surface concavities being shaped as expected from the real world object image.

4.3 Quantitative evaluation

Fig. 6 shows a comparison of the reconstructed depth $z_{\mathrm{PS,n}}$ and quantitatively measured depth data obtained using a laser range scanner z_{RS} for the iterations $n = 1...8$. The registration of both (already pixelsynchronous) depth profiles has been achieved by aligning the absolute depth of a convex surface part in the “hand"-region. This ensures that the PS+Z depth $z_{\mathrm{PS,n}}$ is unaffected by interreflections in that area, and thus allows unbiased registration. With naive photometric reconstruction (no compensation of interreflections, $z_{PS,1}$, Fig. 6(a)), there are very strong differences visible between the reconstructed surface $z_{PS,1}$ and the ground truth scanner depth z_{RS}, which are in the range of some *mm*. The differences initially reach

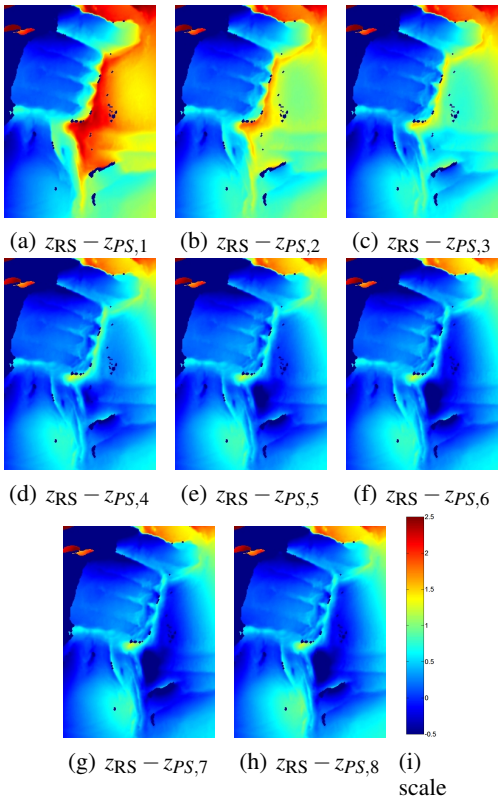

(a) $z_{\mathrm{RS}} - z_{PS,1}$ (b) $z_{\mathrm{RS}} - z_{PS,2}$ (c) $z_{\mathrm{RS}} - z_{PS,3}$

(d) $z_{\mathrm{RS}} - z_{PS,4}$ (e) $z_{\mathrm{RS}} - z_{PS,5}$ (f) $z_{\mathrm{RS}} - z_{PS,6}$

(g) $z_{\mathrm{RS}} - z_{PS,7}$ (h) $z_{\mathrm{RS}} - z_{PS,8}$ (i) scale

Figure 6: Ground truth validation. Scale: $-0.5\,mm$ (dark blue) to $2.5\,mm$ (red)

more than $2\,mm$ (Fig. 6(a)), which is unacceptable in terms of measurement accuracy for e.g. industrial applications. Tn the course of the iterations, the differences decrease quickly and already after the second iteration with interreflection-compensated image data ($z_{PS,3}$, Fig. 6(c)), the error has decreased over a wide range of the concavity.

4.4 Convergence criterion

Although it seems to be of less importance when to stop iterating due to e.g. shallow concavities remaining unchanged while deeper concavities still deepen, it makes sense to include another criterion for robust-

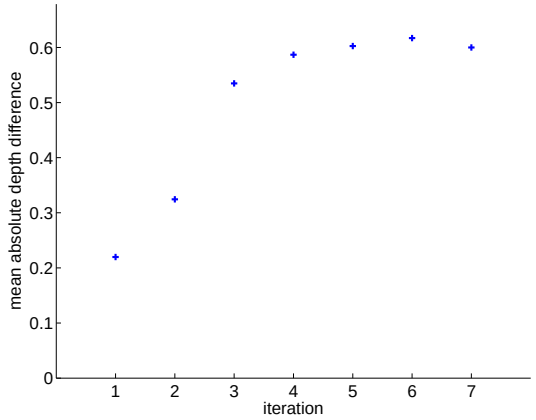

Figure 7: Average change in depth $\Delta z_n = \text{mean}(|z_{PS,n} - z_{PS,1}|)$ per iteration in the "bag"-section

ness and for saving computational time. We found it particularly useful to examine the mean absolute change in depth of the photometrically reconstructed depth data of iteration n (i.e. $z_{PS,n}$) compared to the initial PS+Z result without compensation of interreflections (i.e. $z_{PS,1}$) such that $\Delta z_n = \text{mean}(|z_{PS,n} - z_{PS,1}|)$. Once that change is lower than a certain threshold (here: $\Delta z_n < 0.02\,\text{mm}$), the computation is regarded as complete. In our case, convergence is then reached after the 4^{th} iteration.

5 SUMMARY AND CONCLUSION

We have presented a novel, iterative approach for dealing with interreflections in the context of PS based shape reconstruction. The algorithm initially reconstructs the depth profile without consideration of interreflections using a PS approach and subsequent absolute depth estimation based on the PS gradient fields. Since that profile mainly bares errors in regions affected by excessive light emission due to interreflections, we have introduced an iterative scheme that computes the intensity of the global illumination at surface concavities and then removed these from the physical image data before using that improved data for PS and absolute depth reconstruction.
The evaluation showed a qualitatively and quantitatively beneficial effect for the extracted surface. The results show furthermore, that convergence is reached already after ≈ 4 iterations (for the given surface). Furthermore, a criterion (mean absolute depth change) has been presented that allows to determine when the computation has finished.
In future work, there is mainly the need to perform experiments with non Lambertian surfaces, which are significantly more common in the real world than the applied Lambertian surfaces. Conceptually, this requires the algorithm to handle the BRDF estimation and absolute depth estimation step with greater robustness, but the basic functionality and algorithm principle will remain unchanged.

REFERENCES

Agrawal, A., Raskar, R., and Chellappa, R. (2006). What is the range of surface reconstructions from a gradient field? *Proceedings of the European Conference on Computer Vision (ECCV 2006)*, 1(TR2006-021):578–591.

Alldrin, N., Zickler, T., and Kriegman, D. (2008). Photometric stereo with non-parametric and spatially-varying reflectance. *2008 Conference on Computer Vision and Pattern Recognition (CVPR2008)*.

Alldrin, N. G., Mallick, S. P., and Kriegman, D. J. (2007). Resolving the generalized bas-relief ambiguity by entropy minimization. *2007 Conference on Computer Vision and Pattern Recognition (CVPR)*.

Basri, R., Jacobs, D. W., and Kemelmacher, I. (2007). Photometric stereo with general, unknown lighting. *International Journal of Computer Vision (IJCV)*, 72(3):239–257.

Beckmann, P. and Spizzichino, A. (1987). *The Scattering of Electromagnetic Waves from Rough Surfaces*. Number ISBN-13: 987-0890062382. Artech House Radar Library.

Belhumeur, P. N., Kriegman, D. J., and Yuille, A. L. (1999). The bas-relief ambiguity. *International Journal of Computer Vision (IJCV)*, 35(1):1040–1046.

Blinn, J. F. (1977). Models of light reflection for computer synthesized pictures. *ACM SIGGRAPH Computer Graphics*, 11(2):192–198.

Clark, J. J. (1992). Active photometric stereo. *Proceedings of the 1992 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'92)*, pages 29–34.

Couture, V., Martin, N., and Roy, S. (2011). Unstructured light scanning to overcome interreflections. *ICCV'2011*, pages 1–8.

Freeman, W. T. (1994). The generic viewpoint assumption in a framework for visual perception. *Nature*, 368:542–545.

Geisler-Moroder, D. and Dür, A. (2010). A new ward brdf model with bounded albedo. *Eurographics Symposium on Rendering 2010*, 29:1391–1398.

Gu, J., Kobayashi, T., Gupta, M., and Nayar, S. K. (2011). Multiplexed illumination for scene recovery in the presence of global illumination. *ICCV 2011*, pages 1–8.

Gupta, M., Agrawal, A., Veeraraghavan, A., and Narasimhan, S. G. (2011). A practical approach to 3d scanning in the presence of interreflections, subsurface scattering and defocus. *CVPR'2011, IJCV'2012*, pages 1–24.

Hayakawa, H. (1994). Photometric stereo under a light source with arbitrary motion. *Journal of Optical Society of America A (JOSA A)*, 11:3079–3089.

Horn, B. K. P. (1970). Shape from shading: A method for obtaining the shape of a smooth opaque object from one view. Technical Report 232, Messachusets Institute of Technology.

Ikeuchi, K. (1981). Determining surface orientations of specular surfaces by using the photometric stereo method. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 3(6):661–669.

Iwahori, Y., Sugie, H., and Ishii, N. (1990). Reconstructing shape from shading images under point light source illumination. *Proceedings of IEEE 10th International Conference on Pattern Recognition (ICPR'90)*, 1:83–87.

Lafortune, E. P. F., Foo, S.-C., Torrance, K. E., and Greenberg, D. P. (1997). Non-linear approximation of reflectance functions. *SIGGRAPH'97*, pages 117–126.

Lambert, J.-H. (1760). *Photometria, sive de mensura et gradibus luminis, colorum et umbrae*. Vidae Eberhardi Klett.

Lenoch, M., Herbort, S., and Wöhler, C. (2012). Robust and accurate light source calibration using a diffuse spherical calibration object. *Oldenburger 3D Tage 2012*, 11:1–8.

Nayar, S. K., Fang, X.-S., and Boult, T. (1997). Separation of reflection components using color and polarization. *International Journal of Computer Vision*, 21(3):163–186.

Nayar, S. K., Ikeuchi, K., and Kanade, T. (1988). Extracting shape and reflectance of lambertian, specular and hybrid surfaces. Technical Report CMU-FU-TR-88-14, The Robotics Institute, Carnegie Mellon University.

Nayar, S. K., Ikeuchi, K., and Kanade, T. (1990a). Determining shape and reflectance of hybrid surfaces by photometric sampling. *IEEE Transactions on Robotics and Automation*, 6(1):418–431.

Nayar, S. K., Ikeuchi, K., and Kanade, T. (1990b). Shape from interreflections. Technical Report CMU-RI-TR-90-14, Carnegie-Mellon University of Pittsburgh, PA, Robotics Institute.

Nayar, S. K., Krishnan, G., Grossberg, M. D., and Raskar, R. (2006). Fast separation of direct and global components of a scene using high frequency illumination. *ACM Transactions on Graphics (TOG2006), Proceedings of ACM SIGGRAPH 2006*, 25(3):935–944.

Oren, M. and Nayar, S. K. (1994). Generalization of lambert's reflectance model. *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH 1994)*, pages 239–246.

Pharr, M. and Humphreys, G. (2010). *Physically Based Rendering - From Theory to Implementation*. Morgan Kaufmann (Elsvier).

Phong, B. T. (1975). Illumination for computer generated pictures. *Communications of the ACM*, 18(6):311 –17.

Seitz, S. M., Matasushita, Y., and Kutulakos, K. N. (2005). A theory of inverse light transport. *ICCV 2005*, pages 1440 — 1447.

Tan, P. and Zickler, T. (2009). A projective framework for radiometric image analysis. *CVPR 2009*, pages 2977–2984.

Tan, R. T. and Ikeuchi, K. (2005). Separating reflection components of textured surfaces using a single image. *PAMI'05*, 27(2):179–193.

Tankus, A., Sochen, N., and Yeshurun, Y. (2005). Shape-from-shading under perspective projection. *International Journal of Computer Vision*, 63(1):21–43.

Thomas, D. and Sugimoto, A. (2010). Range image registration of specular objects under complex illumination. *Fifth International Symposium on 3D Data Processing, Visualization and Transmission (3DPVT2010)*.

Torrance, K. E. and Sparrow, E. M. (1967). Theory for off-specular reflection from roughened surfaces. *Journal of the Optical Society of America A (JOSA A)*, 57(9):1105–1114.

Ward, G. J. (1992). Measuring and modeling anisotropic reflection. *ACM SIGGRAPH Computer Graphics*, 26(2):265–272.

Wolff, L. B. (1989). Using polarization to separate reflection components. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'89)*, 1(1):363–369.

Wolff, L. B. and Boult, T. E. (1991). Constraining object features using a polarization reflectance model. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(7):635–657.

Woodham, R. J. (1980). Photometric method for determining surface orientation from multiple images. *Optical Engineering*, 19(1):139–144.

Yu, Y., Debevec, P., Malik, J., and Hawkins, T. (1999). Inverse global illumination: Recovering reflectance models of real scenes from photographs. *Association for Computing Machinery, Special Interest Group on Computer Graphics and Interactive Techniques (ACM SIGGRAPH1999)*, pages 215–224.

Yuille, A. L., Coughlan, J. M., and Konishi, S. (2000). The generic viewpoint constraint resolves the generalized bas relief ambiguity. *Conference on Information Science and Systems*.

Zhou, Z. and Tan, P. (2010). Ring-light photometric stereo. *Proceedings of the 11th European Conference on Computer Vision (ECCV'10)*, pages 1–14.

Zickler, T., Belhumeur, P. N., and Kriegman, D. J. (2002). Helmholtz stereopsis: Exploiting reciprocity for surface reconstruction. *Procedings of the European Conference on Computer Vision 2002*, 3:869–884.