

Enhanced Photometric Stereo with Multispectral Images

Tsuyoshi Takatani^{*†}

Yasuyuki Matsushita[‡]

Stephen Lin[‡]

Yasuhiro Mukaigawa[†]

Yasushi Yagi[†]

[†]Osaka University

[‡]Microsoft Research Asia

{takatani,mukaigaw,yagi}@am.sanken.osaka-u.ac.jp

{yasumat,stevelin}@microsoft.com

Abstract

We introduce a technique based on multispectral images aimed at improving Lambertian photometric stereo. Many photometric stereo algorithms assume Lambertian reflectance, but deviations from this ideal produce errors in shape reconstruction. To alleviate this problem, we exploit the wavelength-dependence of material reflectance. Based on the observation that reflectance at certain wavelengths for a given object is more Lambertian than at others, we propose a method for identifying such wavelengths by a matrix rank analysis, and use them to achieve more accurate photometric stereo. We merge reconstruction results from different wavelengths to produce the final surface normal map. Experimental results on synthetic and real data demonstrate the greater accuracy of this method compared to conventional photometric stereo based on brightness images.

1 Introduction

Photometric stereo, a technique originally proposed by Woodham [1] and Silver [2], is a traditional method for estimating surface orientations from multiple images taken under different lighting conditions. A key factor in this process is the surface reflectance, which describes how the shading at each surface point changes in relation to the lighting direction and the surface normal. By varying the light direction, the corresponding changes in the shading are used to infer the surface normals according to the reflectance model.

Surface reflectance is normally assumed to follow Lambert's law, whereby the intensity I of the reflected light is proportional to the inner product of the lighting vector \mathbf{l} and the surface normal vector \mathbf{n} :

$$I(x) = \rho(x)\mathbf{n}^T\mathbf{l}, \quad (1)$$

where $\rho(x)$ denotes the albedo, or intrinsic color, of the surface at point x . This Lambertian model of reflectance is frequently used in photometric stereo because of its simplicity and convenience. However, the small deviations from the Lambertian model displayed by many materials produces inaccuracies in shape reconstruction.

Various methods have been proposed for dealing with non-Lambertian reflectance. Non-Lambertian highlights are treated as outliers within a Lambertian photometric stereo framework [3], or removed from images [4]. Some instances of photometric stereo are formulated for specific parametric reflectance models,

^{*}This work has been done while the author was a research intern at Microsoft Research Asia.

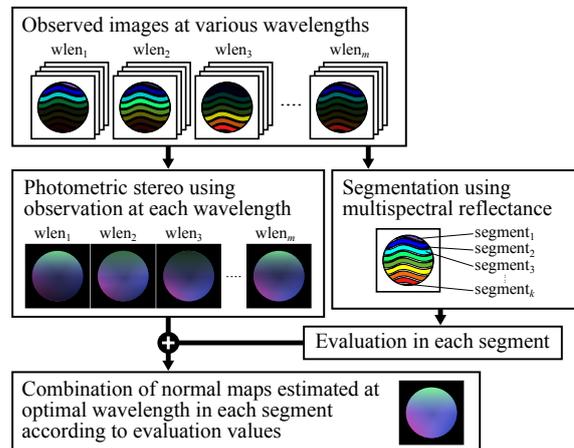


Figure 1. Overview of our algorithm for enhancing the performance of photometric stereo using multispectral images.

such as the Torrance-Sparrow model [5], a combination of isotropic Ward models [6], and m -lobed reflectance maps [7]. Several works have been presented for an arbitrary form for the reflectance [8, 9]. Some methods analyze color for photometric stereo, but mainly to minimize image acquisition with differently colored illuminations [10], to improve robustness by incorporating additional measurements [11], or to aid in removing non-Lambertian reflectance components [4, 12].

Our algorithm is outlined in Fig. 1. Our input is a set of images captured at various wavelengths and for different lighting directions. For each wavelength, we apply Lambertian photometric stereo. In addition, the multispectral image constructed from the set of wavelengths is segmented into different multispectral regions. In each of these, we evaluate the photometric stereo results for each wavelength, using matrix rank analysis to identify which wavelength is most consistent with the Lambertian model. The surface normal map computed for each wavelength are then merged by selecting the optimal wavelength for each region. We thus obtain the final normal map. This approach is validated with simulated images. Experiments on real images support this multispectral approach, showing visible improvements in surface normal estimation over conventional photometric stereo based on brightness images.

2 Photometric Stereo using Multispectral Images

2.1 Wavelength-dependent Reflectance

We first examine the wavelength-dependence of reflectance and its effects on photometric stereo recon-

struction. When accounting for a continuous range of wavelengths, the Lambertian reflectance model in Eq. (1) can be expressed as

$$I(\lambda, x) = \sigma(\lambda)\rho(\lambda, x)e(\lambda)\mathbf{n}(x)^T\mathbf{l}, \quad (2)$$

where λ ranges over the measurable wavelengths of the sensor, σ is the spectral sensitivity of the camera, ρ is the spectral reflectance at a given surface point, and e is the spectrum of the light source.

The diffuse reflection that is conventionally used for photometric stereo arises from a physical process in which light penetrates an object surface, scatters off material particles, and is emitted out of the surface. According to the Lambertian model, this outgoing light is distributed uniformly in all directions, with an intensity given by Eq. (2).

In reality, the directional distribution of the outgoing light is determined by how the light is scattered within the material. This subsurface scattering is physically determined with respect to the optical properties of the material, virtually all of which are wavelength-dependent. For example, the wavelength can affect the number of scattering events. For details on the wavelength-dependence of optical properties in materials, we refer the reader to [13].

This wavelength-dependence of optical properties suggests that surface reflectance is also spectrum-sensitive, thereby affecting the reconstruction quality of photometric stereo. We demonstrate this in Fig. 2 with a simulated sphere rendered using real reflectance data for red plastic, obtained from the MERL BRDF dataset [14]. Forty images with different lighting directions were used for photometric stereo, and the color spectrum was coarsely sampled in red, green, and blue channels. The angular error maps in the reconstructions reveal different degrees of accuracy in the different color channels.

2.2 Region-based Identification of Optimal Wavelengths

Given a set of multispectral images captured with a stationary camera and different lighting directions, our algorithm first computes Lambertian photometric stereo, performs a region segmentation, and identifies the optimal wavelength for each region, which is then used to merge the normal maps computed with different wavelengths. The details of the process are as follows.

Region segmentation As mentioned previously, the optimal wavelength for Lambertian photometric stereo may vary across the image because of varying BRDFs. In particular, our investigations have shown that regions of different color often have different optimal wavelengths. To improve the reconstruction accuracy, we therefore segment the scene by normalized multispectral information, and determine the optimal wavelength for each of the regions. We apply k -means clustering for the segmentation, using an image where the object is fully illuminated (*i.e.*, no shadows).

Optimal wavelength identification in each region To find the optimal wavelength in each segmented region, we evaluate the data in each region.

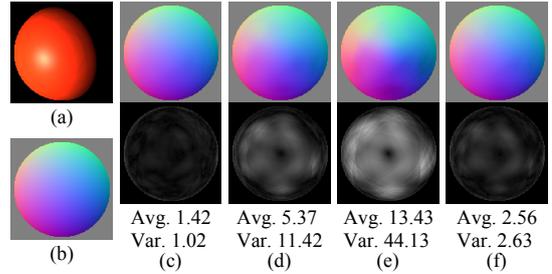


Figure 2. Effect of different wavelengths on photometric stereo. (a) One of 40 images rendered with different lighting directions; (b) ground truth of the normal map; normal maps and angular error maps using the (c) red, (d) green, and (e) blue color channels for Lambertian photometric stereo; and (f) results with brightness images.

For a region with N pixels and M lighting conditions, we can express Eq. (2) in matrix form as

$$\mathbf{I} = g(\lambda)\mathbf{N}\mathbf{L}, \quad (3)$$

where $g(\lambda) = \sigma(\lambda)\rho(\lambda)e(\lambda)$, $\mathbf{I} \in \mathbb{R}^{N \times M}$ is a matrix of image intensities, $\mathbf{N} \in \mathbb{R}^{N \times 3}$ is a matrix of surface normals, and $\mathbf{L} \in \mathbb{R}^{3 \times M}$ is a matrix of lighting vectors.

As seen in Eq. (3), the Lambertian model predicts that the rank of the observation matrix \mathbf{I} must equal that of the normal vector matrix \mathbf{N} , given that the sampled lighting directions are not co-planar (*i.e.*, $\text{rank}(\mathbf{L}) = 3$). As $\text{rank}(\mathbf{I})$ deviates further from $\text{rank}(\mathbf{N})$, the data becomes more poorly fitted by the Lambertian model and then Lambertian photometric stereo should produce less accurate results as discussed in [15].

In view of this effect, we evaluate the wavelengths by determining how close the rank of its observation matrix is to $\text{rank}(\mathbf{N})$. The evaluation function used in this work is

$$E = \frac{\sigma_{k+1}}{\sigma_k}, \quad (4)$$

where $k = \text{rank}(\mathbf{N})$ and σ_i denotes the i -th eigenvalue of \mathbf{I} . This function becomes large when the rank of \mathbf{I} falls below k (*i.e.*, σ_k close to 0) or exceeds k (*i.e.*, σ_{k+1} becomes large). Thus, the wavelength that minimizes E is that for which the rank of \mathbf{I} is closest to $\text{rank}(\mathbf{N})$. A small additive term may be included in the denominator to avoid division by zero, but in practice we have found σ_k never to be zero.

Assuming the rank of \mathbf{N} to be three, Eq. (4) is evaluated with $k = 3$. Corresponding values of k should be used in degenerate cases of lower rank. For each segmented region, the optimal wavelength minimizes E , and to enhance Lambertian photometric stereo, we estimate the final normal map by combining the normal maps of each region that correspond to its optimal wavelength.

3 Experimental Results

3.1 Simulation images

We first validate our method on a simulated scene for which the ground truth surface normals are known. Our scene consists of an egg-shaped surface rendered with the reflectances of seven different materials from the MERL BRDF dataset [14]. Forty images under

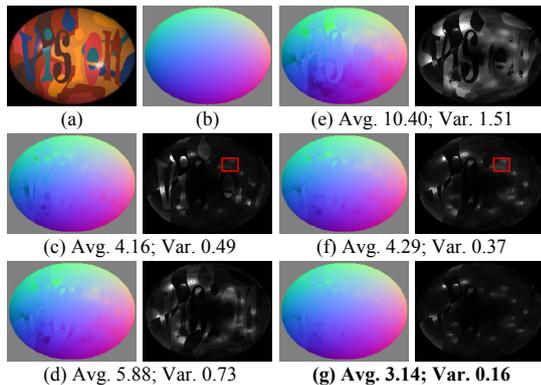


Figure 3. Normal map estimation for simulation images. (a) One of 40 images rendered under different light directions; (b) ground truth of the normal map; normal maps and angular error maps using the (c) red, (d) green, and (e) blue color channels; (f) with brightness images; (g) with our method.

Table 1. Evaluation function values.

Material name	Red	Green	Blue
blue-acrylic	0.960	0.985	0.642
green-plastic	0.840	0.831	0.885
light-brown-fabric	0.687	0.852	0.968
orange-paint	0.767	0.812	0.897
purple-paint	0.348	0.837	0.847
red-phenolic	0.403	0.932	0.900
yellow-matte-plastic	0.434	0.604	0.898

different lighting directions are generated. One of these is shown in Fig. 3(a). The ground truth is exhibited in Fig. 3(b). We note that the MERL BRDF dataset provides reflectance data in terms of only red, green, and blue, so our multispectral technique is performed with three color channels in this instance.

Photometric stereo is first computed for each of the color channels separately, and also for the brightness images obtained by summing the three color channels. The estimated normal maps, angular error maps, and error statistics are shown in Fig. 3(c-f). Overall, the brightness images yield better results than the images of individual color channels, probably because of a higher signal-to-noise ratio (SNR) resulting from summing multiple measurements. However, in some regions, using an individual color channel yields a better estimate than the brightness images (see the red rectangles in Fig. 3 for example).

The result of combining normal map regions from different color channels is shown in Fig. 3(g). Each region from the three color channel images is evaluated using Eq. (4), giving the values shown in Table 1. Pixels with sharp highlights detected by intensity thresholding were discarded from the evaluation by removing their rows from the observation matrix. With this multispectral technique, we achieve reconstructions that surpass those derived from brightness images, despite our technique not using summed measurements to improve the SNR.

3.2 Real images

We conducted an experiment with a real scene captured using a fixed camera (Point Grey Research Grasshopper), a halogen lamp (OSRAM), and nine

narrow band filters at 450, 488, 580, 650, 694, 730, 780, 880 and 940 nm (Edmond Optics). A translation stage (Sigma) was used to switch the narrow band filters automatically. Images were recorded under twelve different lighting directions. We applied photometric stereo to the images under perspective projection [16] to estimate the surface normals.

The target object in this experiment is a plushie shown in Fig. 4, made of fabric of different colors.

Normal estimation using each wavelength The normal maps of the plushie estimated from different wavelengths are shown in Fig. 4. Subtle but visible differences are apparent in the results obtained from the different wavelengths. Long wavelengths tend to yield smoother reconstruction results, possibly because they penetrate deeper into the surface. By contrast, shorter wavelengths often yield sharper surface details. At certain wavelengths, such as 650 nm, surface normal discontinuities may appear, coinciding with sharp color gradients on the surface. These effects may indicate problems in fitting the Lambertian model to the data at these wavelengths.

From the surface normal reconstructions for the different wavelengths, we next determine which wavelengths are most consistent within the Lambertian model with our segmentation and evaluation scheme.

Segmentation and region evaluation Segmentation is performed using the nine color channels of our multispectral images. The segmentation result of the target object is shown in Fig. 5(a). The plushie is divided into nine different region types. Disconnected areas on the image with the same region index are treated collectively as one region.

After segmentation, Eq. (4) is used to determine the optimal wavelength for each region. The values of the evaluation function for the different wavelengths are shown in Fig. 5(c) for each region. The evaluation curves tend to be even over the color spectrum, possibly because the fabric material is less transparent, and hence more similar degrees of light penetration into the surface at different regions and for different wavelengths.

Combination of normal maps using evaluation function values The evaluation function minima in Fig. 5(c) yield the optimal wavelength for each region, displayed in the label map in Fig. 5(b). Long wavelengths in the visible range are often selected.

A reasonable supposition is that the optimal wavelength corresponds to the surface color of a region, since these wavelengths produce the brightest image and a high SNR. However, our results suggest otherwise. For example, the optimal wavelength is in the red spectral range even for a green surface, because of the subsurface scattering properties of this region.

The normal map computed by our method from the label map and the normal maps at different wavelengths is shown in Fig. 4. Though the ground truth is unavailable for this object, our normal map appears qualitatively more correct than the normal map computed from brightness images. With our method, the normal map is less sensitive to texture.

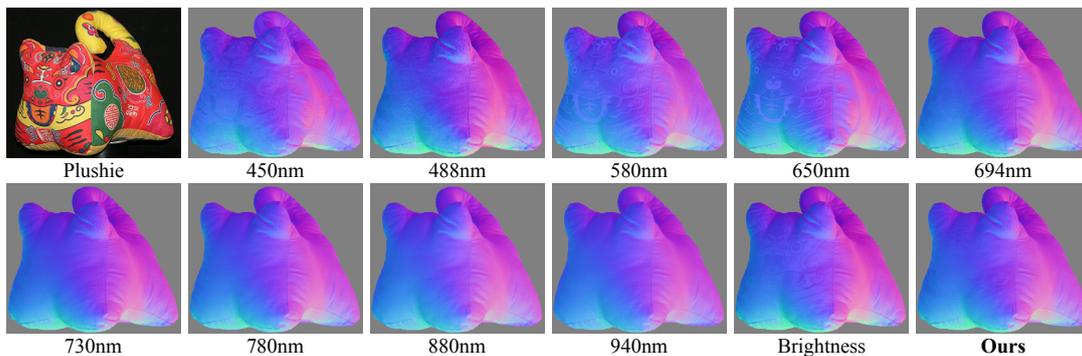


Figure 4. Picture and normal maps estimated using different wavelengths, brightness images, by applying our method.

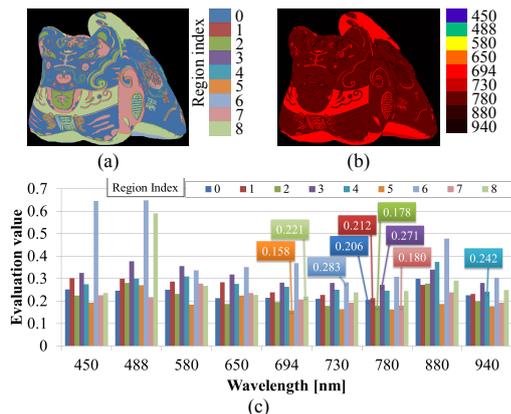


Figure 5. Our algorithm illustrated. (a) Nine regions segmented by k -means clustering in multispectral space. (b) Optimal wavelength for each region identified by Eq. (4). (c) Evaluation function values versus wavelength for each region.

4 Conclusion

In this paper, we proposed a new technique based on multispectral images to improve Lambertian photometric stereo. The resulting normal map is obtained by combining different normal maps computed at an optimal wavelength determined in each region of the image. Optimal wavelength determination does not require ground truth geometry, but is based instead on matrix rank analysis. Our experiments demonstrated significant wavelength-dependence of various materials in photometric stereo reconstruction. Our method demonstrably shows improvements over traditional photometric stereo based on brightness images.

Some materials have reflectance properties that differ significantly from the Lambertian model for any wavelength. To broaden the applicability of our technique, future work will extend our approach to other parametric reflectance models used in photometric stereo. In addition, we intend to investigate the wavelength-dependent reflectance of various objects for the purpose of material recognition.

Acknowledgments

This research is granted by the Japan Society for the Promotion of Science (JSPS) through the “Funding Program for Next Generation World-Leading Researchers (NEXT Program),” initiated by the Council for Science and Technology Policy (CSTP).

References

[1] R. J. Woodham, “Photometric method for determining surface orientation from multiple images,” *Optical*

Engineering, vol. 19, no. 1, pp. 139–144, 1980.

[2] W. Silver, “Determining shape and reflectance using multiple images,” tech. rep., Master thesis, MIT, 1980.

[3] S. Barsky and M. Petrou, “The 4-source photometric stereo technique for three-dimensional surfaces in the presence of highlights and shadows,” *IEEE Trans. on Pattern Anal. Mach. Intell.*, vol. 25, no. 10, pp. 1239–1252, 2003.

[4] S. P. Mallick, T. E. Zickler, D. J. Kriegman, and P. N. Belhumeur, “Beyond lambert: Reconstructing specular surfaces using color,” in *Proc. CVPR*, vol. 2, pp. 619–626, IEEE, 2005.

[5] A. Georgiades, “Incorporating the torrance and sparrows model of reflectance in uncalibrated photometric stereo,” in *Proc. ICCV*, pp. 816–823, IEEE, 2003.

[6] D. B. Goldman, B. Curless, and A. H. S. M. Seitz, “Shape and spatially-varying brdfs from photometric stereo,” in *Proc. ICCV*, vol. 1, pp. 341–348, IEEE, 2005.

[7] H. Tagare and R. de Figueiredo, “A theory of photometric stereo for a class of diffuse non-lambertian surfaces,” *IEEE Trans. on Pattern Anal. Mach. Intell.*, vol. 13, no. 2, pp. 133–152, 1991.

[8] N. Alldrin and D. Kriegman, “Photometric stereo with non-parametric and spatially-varying reflectance,” in *Proc. CVPR*, IEEE, 2008.

[9] T. Higo, Y. Matsushita, and K. Ikeuchi, “Consensus photometric stereo,” in *Proc. CVPR*, pp. 1157–1164, IEEE, 2010.

[10] R. Anderson, B. Stenger, and R. Cipolla, “Color photometric stereo for multicolored surfaces,” in *Proc. ICCV*, IEEE, 2011.

[11] P. H. Christensen and L. G. Shapiro, “Three-dimensional shape from color photometric stereo,” *International Journal of Computer Vision*, vol. 13, no. 2, pp. 213–227, 1994.

[12] M. Liao, X. Huang, and R. Yang, “Interreflection removal for photometric stereo by using spectrum-dependent albedo,” in *Proc. CVPR*, pp. 689–696, IEEE, 2011.

[13] M. Bass, ed., *Handbook of Optics: Optical properties of materials, nonlinear optics, quantum optics*. Optical Society of America, 2009.

[14] W. Matusik, H. Pfister, M. Brand, and L. McMillan, “A data-driven reflectance model,” *ACM Trans. on Graph.*, vol. 22, no. 3, pp. 759–769, 2003.

[15] L. Wu, A. Ganesh, B. Shi, Y. Matsushita, Y. Wang, and Y. Ma, “Robust photometric stereo via low-rank matrix completion and recovery,” in *Proc. ACCV*, 2010.

[16] A. Tankus and N. Kiryati, “Photometric stereo under perspective projection,” in *Proc. ICCV*, vol. 1, pp. 611–616, IEEE, 2005.