Scattering Parameters and Surface Normals from Homogeneous Translucent Materials using Photometric Stereo

Bo Dong   Kathleen D. Moore   Weiyi Zhang   Pieter Peers
College of William & Mary

Abstract

This paper proposes a novel photometric stereo solution to jointly estimate surface normals and scattering parameters from a globally planar, homogeneous, translucent object. Similar to classic photometric stereo, our method only requires as few as three observations of the translucent object under directional lighting. Naively applying classic photometric stereo results in blurred photometric normals. We develop a novel blind deconvolution algorithm based on inverse rendering for recovering the sharp surface normals and the material properties. We demonstrate our method on a variety of translucent objects.

1. Introduction

A commonly used lightweight method for shape acquisition is photometric stereo [26]. Photometric stereo exploits the direct relation between the incident irradiance and the outgoing observed radiance at a surface point to infer the surface normal from as little as three observations under directional lighting. While originally developed for Lambertian materials only, it has been extended to handle more general surface reflectance functions. Photometric stereo is also commonly applied to translucent materials, such as human skin. However, the resulting photometric normals are "blurred" due to the indirect relation between incident irradiance and the observed radiance at a surface point [21].

The two most common applications of photometric stereo are shape recovery and surface detail recovery (typically followed by an embossing step onto a global shape [22]). In this paper, we focus on this second application for which the impact of translucency on the accuracy of photometric stereo is particularly striking. We propose a novel method that computes "deblurred" normals of homogeneous, translucent, globally planar objects that exhibit local surface normal variation. In addition, our method also obtains estimates of the reduced albedo and mean free path that characterize the appearance of the translucent material. Notably, our method only requires the same input as "traditional" photometric stereo (*i.e.*, three or more photographs

under directional lighting). We formulate our method as a blind deconvolution problem and rely on inverse rendering to gauge the fitness of the estimated parameters. We validate the effectiveness of our method on a variety of measured materials, and compare the estimated parameters and normals to baseline measurements.

2. Related Work

Photometric Stereo Photometric stereo [26] is a lightweight method for estimating surface normals of Lambertian objects from observations under as few as three directional light sources. Since its inception, many improvements and variants have been proposed that extend to non-Lambertian surface reflectance (*e.g.*, [1, 4, 10, 14, 25]), minimize the impact of shadows [2], and jointly estimate normals and lighting (*e.g.*, [3, 8, 19, 24]). However, all these methods assume that the observed radiance at a surface point is directly related to the incident irradiance at the *same* surface point only. Ma *et al.* [20] recognize the potential issue of subsurface light transport on photometric stereo, and propose a photometric stereo variant that relies on polarized spherical gradient illumination to infer normals from specular reflections (isolated via polarization differencing), thus avoiding the bias introduced by subsurface scattering. Recently, Gu *et al.* [11] showed that 7 high-frequency multiplexed illumination patterns are sufficient to practically recover surface normals from translucent objects. However, both methods require more complex hardware and more measurements than classic photometric stereo. In contrast, the proposed method removes the effects of translucency in post-processing. This allows us to maintain the lightweight acquisition procedure and setup of classic photometric stereo.

Subsurface Scattering The Bidirectional Subsurface Scattering Distribution Function (BSSRDF) [23], an 8D function, describes the appearance of translucent materials. The BSSRDF is often split into two components: a component that models single scattering, and a component that models multiple scattering. The former can be effectively modeled by a 4D BRDF (*i.e.* Bidirectional Reflectance Distribution Function) approximation [13]. For optically thick

materials dominated by scattering, the diffusion approximation holds for multiple scattered light, (*i.e.*, the dependence between incident and exitant lighting directions is lost). Several compact analytical models have been proposed to model the diffusion approximation of multiple scattering (also called *diffuse BSSRDF*) in terms of index of refractions, and *reduced scattering* and *absorption coefficients*. Jensen *et al.* [16] model multiple scattering using a dipole approximation. Donner *et al.* [6] extend the dipole model to a multipole model to better model translucent slabs of finite depth. d'Eon and Irvin [5] propose the quantized diffusion approximation to better address the excessive blurring of the dipole approximation and approximate the diffusion profile by a sum of 64 Gaussians. Recently, Haber *et al.* [12] introduce the photon beam diffusion approximation that is computationally more convenient than quantized diffusion (*i.e.*, it only requires summing as few as 5 dipoles) . Furthermore, the photon beam diffusion approximation also accounts for the effects of oblique incident light angles.

Reduced scattering and absorption parameters are acquired from physical material samples by fitting to the observed diffusion profiles (*e.g.*, [9, 16]) under (spatially) high-frequency illumination. Variations in surface normals are ignored, potentially introducing incorrect estimates. In contrast, the proposed method uses directional lighting without spatial variation and relies on surface normal variation to infer scattering and absorption coefficients. The method of Donner *et al.* also employs uniform lighting and relies on multispectral observations to estimate the reduced scattering and absorption coefficients. However, their method is limited to skin only. Similar to us, Zhu *et al.* [27] also exploit local surface variations (*i.e.*, curvature) to infer the BSSRDF parameters under spherical gradient illumination. However, their method requires knowledge of the surface geometry, and uses a large and complex setup. In contrast, the proposed method requires as few as three lights and jointly estimates the BSSRDF parameters and normals.
Image Deblurring The proposed method is also related to blind image deconvolution (*e.g.*, [17, 18]). Instead of deblurring a single natural image, we deblur normal maps/irradiance images which exhibit different statistics. Furthermore, our "blur" kernel is not sparse. We do not claim any technical contributions to blind deconvolution, but rather rely on the developments in this field.

3. Method

In the following, we assume that the surface is globally planar and boundless with localized surface normal variations. Furthermore, we assume that the surface consists of an optically thick, homogeneous translucent material with a known index of refraction. We will consider only multiple scattering, ignoring single scattering and direct surface reflections.

3.1. Background

The observed radiance $L(x_o, \omega_o)$ at a position x_o in a direction ω_o of a translucent material under a directional light source $L(\omega_i)$ is given by [16]:

$$L(x_o, \omega_o) = \frac{1}{\pi} \int_A F_r(\eta, (n(x_o) \cdot \omega_o)) R_d(x_i, \omega_i, x_o, \omega_o)$$
$$F_r(\eta, (n(x_i) \cdot \omega_i)) L(\omega_i)(n(x_i) \cdot \omega_i) dx_i,$$
$$(1)$$

where F_r is the Fresnel transmission at the incident and exitant surface point with index of refraction η, $n(x)$ is the surface normal at x, and R_d is the diffuse BSSRDF. Several models exist for modeling R_d, ranging from the classical dipole-diffusion model [16] to the quantized diffusion model [5]. In this paper we employ the photon beam diffusion model [12] which accounts for oblique incident light angles and which is fast to evaluate: $R_d(||x_i - x_o||, \omega_i)$. We parameterize R_d in terms of reduced albedo $\alpha' = \frac{\sigma'_s}{\sigma'_t}$ and diffuse mean free path $l_d = \frac{1}{\sigma'_t}$, which provides a more intuitive description than a parameterization in terms of reduced scattering σ'_s and reduced extinction σ'_t – a change in α' roughly corresponds to a change in "profile-height", while a change in l_d roughly corresponds to a change in "profile-width" or extent [7].

As shown in [21], ignoring the subsurface light transport and directly applying photometric stereo [26] to the observed radiance yields a "blurred" normal estimate. This can be clearly seen by rewriting Eq. (1) as

$$L'_{\omega_i} = R_{d\omega_i} * I_{\omega_i}, \qquad (2)$$

where '$*$' is the convolution operator, L' is the ratio between exitant radiance and outgoing Fresnel transmittance, and I is the product between foreshortening and incident Fresnel. Note that I depends only on the dot product between the surface normal $n(x_i)$ and the incident lighting direction ω_i. To better understand the impact of this 'blurring' of the incident irradiance, consider the following didactic example were $\eta = 1.0$ (*i.e.*, $F_r = 1$), and $R_{d\omega_i}$ is independent of ω_i, the "blurred" normal $\overline{n}(x_o)$ can then be expressed as

$$\overline{n} \sim R_d * n. \qquad (3)$$

This suggests that the "sharp" normal field n can be recovered by "deblurring" $\overline{n}$ with R_d. However, directly deblurring the normal field $\overline{n}$ is difficult due to the unit-length constraint on both $\overline{n}$ and n. Furthermore, the index of refraction is unlikely to be 1.0, and the shape of the diffuse BSSRDF R_d depends on α' and l_d, which are unknown. Instead of directly working on the normal field, we propose to work on the observed radiance L. We will employ a blind deconvolution approach to recover both the sharp normal field as well as the reduced albedo α' and mean free path l_d of the translucent material.

3.2. Example

Before developing our algorithm, we first consider the following example to gain further intuitive insight. Consider the limit case of a flat surface without surface variation, *i.e.*, $n(x_i) = n$, and $\eta = 1.0$. In this case Eq. (2) reduces to:

$$L'_{\omega_i} = I_{\omega_i} \int_A R_{d\omega_i} dx_i. \tag{4}$$

It trivially follows that $\overline{n} = n$. Note, however, that I_{ω_i} is equivalent to the surface reflectance of a Lambertian surface with unit albedo. Thus, when computing the surface normal from the observed radiance values, the normalization length of the computed normal equals $\int_A R_{d\omega_i} dx_i$, the total diffuse reflectance. However, this total diffuse reflectance only depends on the reduced albedo α', and is independent of the diffuse mean free path l_d [16]. This illustrates that it is possible to estimate the reduced albedo from typical photometric stereo observations of a homogeneous translucent material (*i.e.*, it is a function of the normalization length). Furthermore, it also illustrates that without surface variations, we cannot infer the diffuse mean free path l_d.

Now consider the case where there is sufficient surface variation, and the reduced albedo α' is known. Deblurring with an incorrect kernel (*i.e.*, incorrect mean free path l_d, but correct α') results in spatially varying artifacts such as ringing. Consequently, even though the reduced albedo is correct, the normalization of estimated normals will also vary spatially, and deviate from the ideal unit length.

The above intuitive analysis shows that the normalization length of the computed normals is indicative of the accuracy of the estimated reduced albedo α' and mean free path l_d. However, we found this normalization length to be very sensitive to calibration inaccuracies and measurement noise. Instead of directly relying on this cue, we will follow an inverse rendering approach to estimate both the sharp surface normals as well as the BSSRDF parameters.

3.3. Algorithm

To estimate BSSRDF parameters and sharp surface normals, we search for the reduced albedo α' and mean free path l_d that best explain the observations. We formulate this as a non-linear optimization in α' and l_d (using the method of Hooke and Jeeves [15] that does not rely on derivatives and thus is less sensitive to measurement noise). In each optimization step, we first remove the effects of the outgoing Fresnel transmission from the observed radiance (Algorithm 1, line 2). Next, we deconvolve each of the radiance images with the proposed BSSRDF kernel (line 3). We model the BSSRDF kernel using the photon beam model [12], and employ a preconditioned conjugate gradient method for computing the deconvolution (*i.e.*, the images are stacked in a vector, and the convolution with R_d is explicitly written as multiplication with a square matrix).

Algorithm 1 Fitness of a proposed BSSRDF (α' and l_d) given observations under directional lighting.

1: **function** FITNESS($\alpha'^{(k)}$, $l_d^{(k)}$, $n^{(k-1)}$, $e^{(k-1)}$, $\{L_{\omega_i}\}_{\omega_i}, \eta$)
2:      $L'_{\omega_i} = L_{\omega_i} / F_r(\eta, (n^{(k-1)} \cdot \omega_o)$
3:      $I_{\omega_i} = L'_{\omega_i} *^{-1} R_{d\omega_i}(\alpha'^{(k)}, l_d^{(k)})$
4:      $\overline{n} = \text{PhotometricStereo}(\{I_{\omega_i}\}_{\omega_i}, \eta)$
5:      Compute $\overline{L}_{\omega_i}$ using Eq. (1) and $\overline{n}$
6:      $e = \sum_{\omega_i} ||L_{\omega_i} - \overline{L}_{\omega_i}||_2$
7:      **if** $e < e^{(k-1)}$ **then**
8:          $n^{(k)} \leftarrow \overline{n}$
9:          $e^{(k)} \leftarrow e$
10:      **else**
11:          $n^{(k)} \leftarrow n^{(k-1)}$
12:          $e^{(k)} \leftarrow e^{(k-1)}$
13:      **end if**
14:      **return** e
15: **end function**

Given the deblurred irradiance values, we estimate (normalized) surface normals (line 4). To account for incoming Fresnel transmission, we formulate this as a non-linear optimization using again the method of Hooke and Jeeves [15]. Next, using these estimated surface normals, we regenerate the observed radiance images by evaluating Eq. (1) (line 5). The fitness of the proposed BSSRDF parameters is then evaluated by the RMS error of the regenerated radiance images and the observed radiance images (line 6). Note, that the first step (removal of the outgoing Fresnel transmission) requires knowledge of the unknown surface normals. We circumvent this problem by keeping track of the best-so-far estimated normals (lines 7-13), and initialize it with the "blurred" normals.

Validation We validated the proposed deconvolution approach on synthetic datasets. Figure 1 shows one such result. We simulated the acquisition using Eq. (1), with $\eta = 1.3$, $\alpha' = 0.999$, and $l_d = 0.436$, which correspond to the properties of apple [16]. We distributed 3 light sources on a 45 degree cone around the virtual camera. As can be seen, our method is capable of recovering the original normal maps, as well as the BSSRDF parameters.

In addition, we validated that the error landscape exhibits a single global minimum by plotting the errors obtained from a systematic parameter scan (Figure 2). In general, the error landscape exhibits a long and narrow valley in (α', l_d) space, being narrow and steepest in terms of reduced albedo α'. While measurement noise can affect the location of the global minimum in the error landscape, its effect is greatly reduced by the 'reblurring' step in our method (step (5) in Algorithm 1). Consequently, we found that we can still obtain good estimates of α' and l_d in the presence of moderate

Blurred
Error: 3.1 degrees

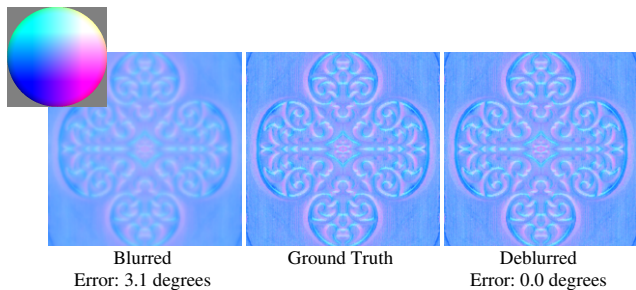Ground Truth

Deblurred
Error: 0.0 degrees

Figure 1. **Inset:** false color encoding exemplar. **Left:** naively applying photometric stereo on a (simulated) translucent material yields blurred surface normal estimates. **Middle:** ground truth normals. **Right:** the surface normals estimated by the proposed method are a closer match to the ground truth normals (note the recovery of the fine details in the 'flat' areas). Furthermore, the estimated BSSRDF parameters ($\alpha' = 0.998689, l_d = 0.436094$) are an almost exact match to the ground truth BSSRDF parameters ($\alpha' = 0.998691, l_d = 0.436110$).

amounts of measurement noise. Note that noise amplification during deblurring can potentially impact the quality of the estimated sharp normals. The effects of noise amplification can be reduced by adding an appropriate noise-suppression term to the deconvolution algorithm. However, we did not explore this option.

We found that our method is not sensitive to the initial values of α' and l_d, and use the randomly chosen $\alpha' = 0.9987$ and $l_d = 0.4342$ as starting point in our implementation.

Experimental Validation Our capture setup consists of a single DSLR camera (Nikon D700) with 4 synchronized flashes mounted approximately 1m from the sample. We currently capture 4 photographs (1 per flash), one more than strictly necessary, for added robustness – we validated that our method also works well with the minimum number of measurements. We cross-polarize the light sources and the camera to cancel out direct surface reflections and single scattering. To radiometrically calibrate our setup, we capture RAW images of a planar Spectralon sample (an ideal diffuse reflector with +99% reflectivity) and normalize these images by the foreshortening of the central pixel. During acquisition, each captured image of a globally planar translucent object is then divided by the respective normalized Spectralon image. This simple calibration procedure ensures that camera and light source vignetting, as well as differences in light source power, are compensated. We assume an index of refraction of 1.3 which is a good approximation for most organic materials – Moore *et al.* [21] showed that the differences in impact on photometric stereo for translucent materials is small when $\eta = [1.2, 1.5]$.

We validate our method by comparing both surface normals as well as BSSRDF parameters with ground truth estimates. We selected soap as a ground truth material as it can


Figure 2. Typical error landscape of the surface normals computed from deblurred (simulated) measurements. The global minimum (highlighted) occurs at $\alpha' = 0.9987$, and $l_d = 0.436$, and is equal to the simulation parameters.

be easily melted and cast into a desired shape. We place an opaque target surface normal sample in a container and submerge it with molten soap. After letting the soap solidify, we carefully remove the mold, yielding a translucent material with the desired (inverse) surface variation. We obtain the ground truth BSSRDF parameters by fitting them to the observations of illuminating a single surface point (with a cross-polarized projector). Figure 3 clearly illustrates the effects of translucency on the accuracy of the estimated photometric surface normals (ground truth (a) versus blurred (b-d)). The different degrees of blurring per color channel are due to the wavelength dependence of BSSRDF parameters α' and l_d. The surface normals, per color channel, obtained with our method (f-h) are closer to the the ground truth normals. Furthermore, the estimated BSSRDF parameters ($\alpha' = [0.999, 0.991, 0.991]$ and $l_d = [0.103, 0.269, 0.239]$, plotted in green in (j-l)) are a close match to the estimated ground truth parameters ($\alpha' = [0.999, 0.992, 0.987]$ and $l_d = [0.286, 0.221, 0.207]$, plotted in red in (j-l)).

3.4. Multiple Color Channels

Ideally, the above blind deconvolution algorithm should result in identical normal maps for each color channel. However, in practice, the different normal maps will differ slightly, due to measurement noise and/or deviations from the theoretical BSSRDF models.

The above algorithm can be easily extended to a joint multi-channel blind deconvolution. We alter the non-linear optimization to search over all color channels' BSSRDF parameters, and change the deconvolution (step (3) in Algorithm 1) to deconvolve all channels simultaneously. The latter is achieved by stacking all the L'_{ω_i} of each color channel in a vector, and concatenating the $n \times n$ $R_{d\omega_i}$ matrices into a single $n \times (n \times \#channels)$ matrix. We again employ preconditioned conjugate gradients (with normal equations) to solve for the sharp normals in a least squares sense.

However, the above non-linear optimization now needs to perform a search in a $\#channels \times 2$ dimensional space

Figure 3. **(a)** Ground truth surface normals captured directly from the mold used to cast the normal variation on the pink soap material (i). **(b-d)** Naively applying photometric stereo directly on the observations yields blurred normals. **(f-g)** Result from our method computed on each channel seperately. **(e)** Result from joint optimization on all color channels simultaneously. **(j-l)** Corresponding ground truth and recovered BSSRDF profiles.

(as opposed to a $2D$ space), which has a slower convergence rate. To speed up computations, we first perform the search on each channel separately, and use the obtained estimates as a start point for the full joint optimization.

3.5. Results & Discussion

Implementation We implemented our algorithm on the GPU in CUDA. Nevertheless, due to the large size of the kernels, the computational cost is still high. We therefore optimize for the BSSRDF parameters on a scaled-down subset, and deconvolve the full image with the obtained parameters. In general we first downsample the photographs to approximately 1.6 pixels/mm, and then crop an 80×80 pixel area. Computing the BSSRDF parameters for a single color channel takes 1 to 2 hours on an Intel E5630 at 2.53Ghz with an NVidia GTX 660 Ti, and 12 hours for a joint three-color channel optimization.

Results Figure 3-(e) shows the jointly optimized surface normal map of the fabricated soap example. The respective BSSRDF parameters, shown in blue in Figure 3-(j-l), are $\alpha' = [0.999, 0.991, 0.991]$ and $l_d = [0.333, 0.284, 0.235]$, which are closer to the ground truth parameters ($\alpha' = [0.999, 0.992, 0.987]$ and $l_d = [0.288, 0.221, 0.207]$) than

the results of optimizing each channel separately.

Due to the high frequency normal variations, it is difficult to make a per-pixel RMSE comparison between the ground truth (*i.e.*, mold) and the estimated normals. Furthermore, the estimates of the "ground truth" BSSRDF parameters are potentially affected by inaccuracies in the PSF estimate of the projector, and due to differences in spectrum between the flashes used for photometric stereo and the projector used for estimating the ground truth parameters. Nevertheless, both are still good indicators of the reconstruction quality (albeit not absolute measures).

We tested our blind deconvolution method on several homogeneous translucent materials (all assuming $\eta = 1.3$). Table 1 summarizes our results in terms of estimated BSSRDF parameters. In general, the error on diffuse mean free path is larger than the error on reduced albedo. However, in terms of BSSRDF shape, the differences are small. Figures 4,5, and 6 show additional results for white soap, cheddar cheese, and potato (albeit without ground truth normals).

Limitations The frequency spectrum of BSSRDF profiles contain very few zeros, indicating that we should be able to cover most of the surface details. However, as is typical
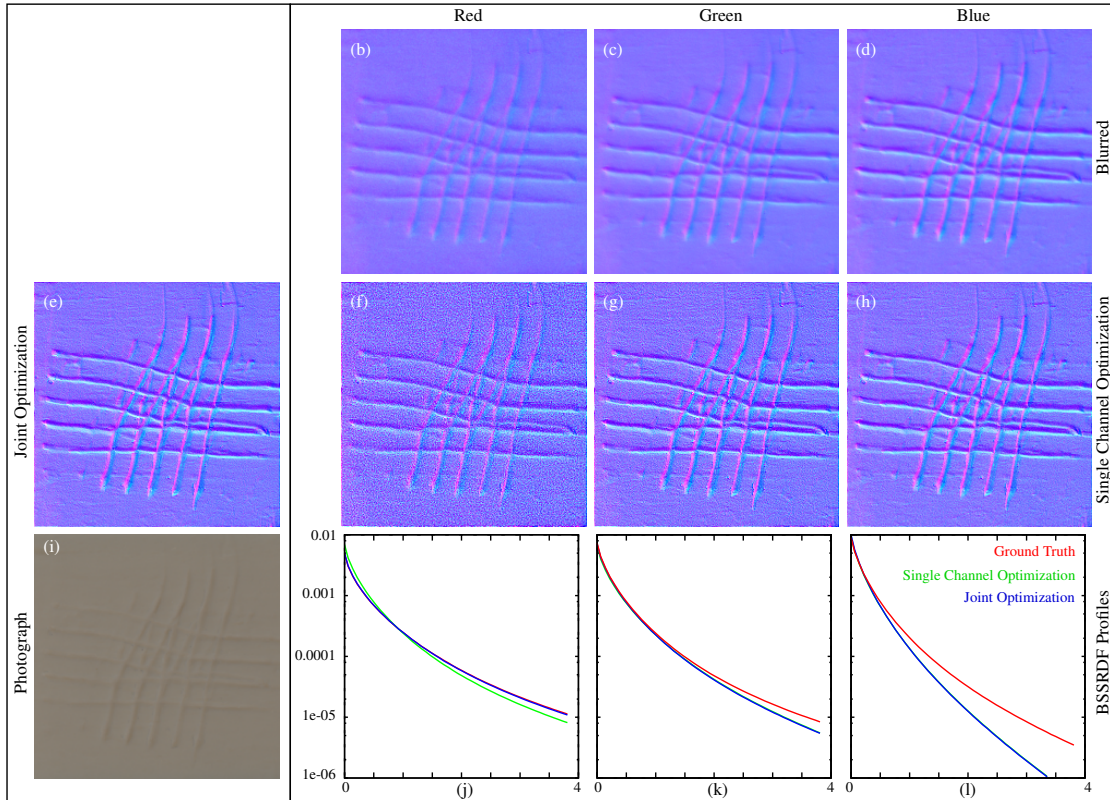
Figure 4. Result for a white soap material (i). For this example, no ground truth normals are available. **(b-d)** Naively applying photometric stereo directly on the observations yields blurred normals. **(f-g)** Result from our method computed on each channel seperately. **(e)** Result from joint optimization on all color channels simultaneously. **(j-l)** Corresponding ground truth and recovered BSSRDF profiles.

Table 1. BSSRDF parameters estimated using the joint multi-channel deconvolution versus ground truth parameters.

Material	Parameter	R	G	B
Pink Soap	est. α'	0.999	0.991	0.991
	g.t. α'	0.999	0.992	0.987
	est. l_d	0.333	0.284	0.235
	g.t. l_d	0.288	0.221	0.207
White Soap	est. α'	0.999	0.998	0.991
	g.t. α'	0.999	0.999	0.998
	est. l_d	0.302	0.223	0.156
	g.t. l_d	0.309	0.219	0.179
Cheddar Cheese	est. α'	0.999	0.968	0.790
	g.t. α'	0.999	0.977	0.875
	est. l_d	0.329	0.289	0.168
	g.t. l_d	0.431	0.414	0.305
Potato	est. α'	0.967	0.997	0.967
	g.t. α'	0.999	0.996	0.956
	est. l_d	1.027	1.057	1.037
	g.t. l_d	1.029	0.947	1.025

for deconvolution, our method suffers from noisy amplification, and thus the accuracy of the recovered normals greatly depends on the measurement quality of the blurred signal. Using the joint multi-channel optimization can help in regularizing the estimation of BSSRDF parameters from noisy color channels. However, in some cases the noise dominates the signal from the other color channels, resulting in a degradation of the estimated normals and BSSRDF parameters. We currently omit such noisy channels manually from the joint optimization. However, it would be interesting to investigate automatic weighting schemes and/or noise suppression to handle such cases.

Our current implementation is limited to globally planar surfaces only. Surface curvature can deform the kernel, effectively yielding a spatially varying kernel that depends on the "deblurred" surface normal. However, we found that we still obtain good estimates for slightly curved surfaces (*e.g.*, the potato example in Fig. 6 is slightly curved). Similarly, our method also assumes a homogeneous translucent material. Subsurface light transport in heterogeneous translucent materials is difficult to describe by simple analytical models, requiring a full data-driven scattering model.

Finally, we currently rely on cross-polarization to isolate the reflectance from multiple scattering. However, for certain materials, such as marble, this separation is not perfect at thin structures (*e.g.*, close to edges), resulting in artifacts in the separation.
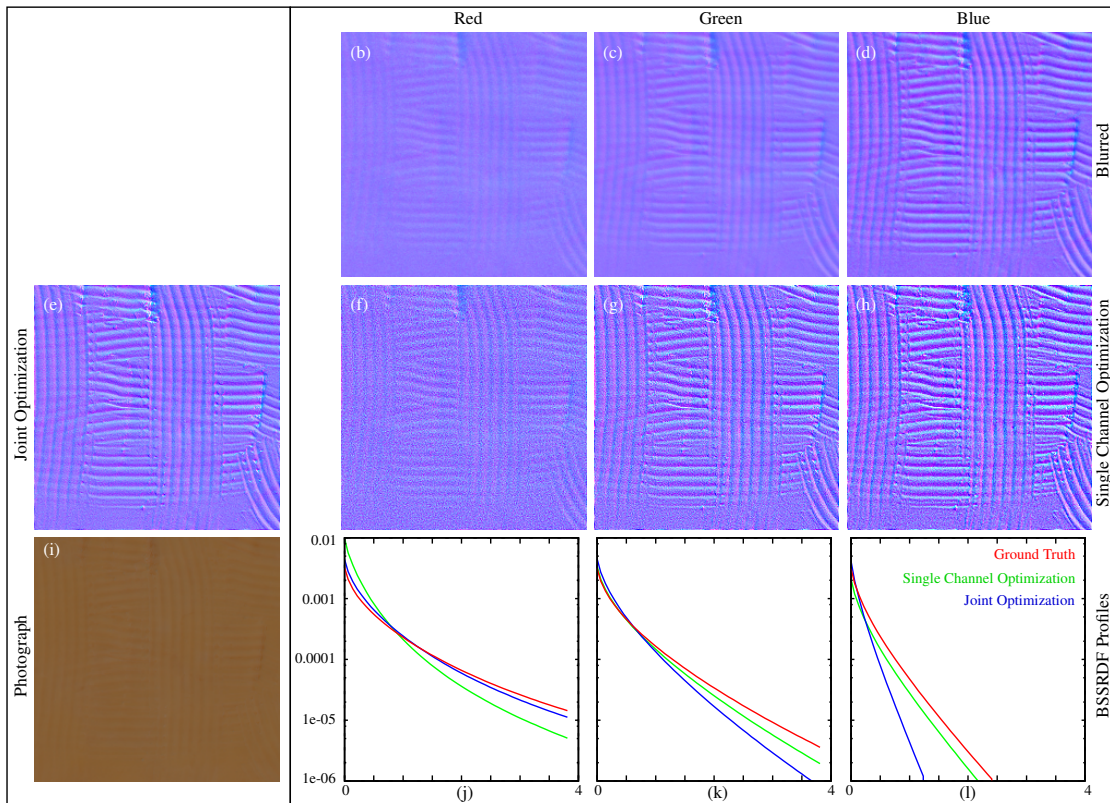
Figure 5. Result for Cheddar Cheese (i). For this example, no ground truth normals are available. **(b-d)** Naively applying photometric stereo directly on the observations yields blurred normals. **(f-g)** Result from our method computed on each channel seperately. **(e)** Result from joint optimization on all color channels simultaneously. **(j-l)** Corresponding ground truth and recovered BSSRDF profiles.

4. Conclusion

We proposed a novel photometric stereo approach for estimating surface normals and scattering parameters of translucent objects based on blind deconvolution and inverse rendering, and demonstrated our method on a variety of translucent materials. For future work, we would like to incorporate advanced deconvolution algorithms to improve computation time and quality. In addition, we would like to extend our method to curved surfaces.

Acknowledgments

References

[1] N. Alldrin and D. Kriegman. Toward reconstructing surfaces with arbitrary isotropic reflectance: A stratified photometric stereo approach. In *ICCV*, pages 1–8, 2007.

[2] S. Barsky and M. Petrou. The 4-source photometric stereo technique for three-dimensional surfaces in the presence of highlights and shadows. *IEEE Trans. Pattern Anal. Mach. Intell.*, 25(10):1239–1252, 2003.

[3] R. Basri and D. Jacobs. Photometric stereo with general, unknown lighting. In *IEEE CVPR*, pages 374–381, 2001.

[4] M. Chandraker, J. Bai, and R. Ramamoorthi. A theory of differential photometric stereo for unknown BRDFs. In *IEEE CVPR*, pages 2505–2512, 2011.

[5] E. D'Eon and G. Irving. A quantized-diffusion model for rendering translucent materials. *ACM Trans. Graph*, 30(4):56:1–56:14, July 2011.

[6] C. Donner and H. W. Jensen. Light diffusion in multi-layered translucent materials. *ACM Trans. Graph*, 24(3):1032–1039, 2005.

[7] C. Donner, J. Lawrence, R. Ramamoorthi, T. Hachisuka, H. W. Jensen, and S. K. Nayar. An empirical bssrdf model. *ACM Trans. Graph.*, 28(3), 2009.

[8] P. Favaro and T. Papadhimitri. A closed-form solution to uncalibrated photometric stereo via diffuse maxima. In *IEE CVPR*, pages 821–828, 2012.

[9] M. Goesele, H. P. A. Lensch, J. Lang, C. Fuchs, and H.-P. Seidel. Disco: acquisition of translucent objects. *ACM Trans. Graph.*, 23(3):835–844, 2004.

[10] D. B. Goldman, B. Curless, A. Hertzmann, and S. M. Seitz. Shape and spatially-varying brdfs from photometric stereo. In *IEEE CVPR*, pages 341–348, 2005.

[11] J. Gu, T. Kobayashi, M. Gupta, and S. Nayar. Multiplexed Illumination for Scene Recovery in the Presence of Global
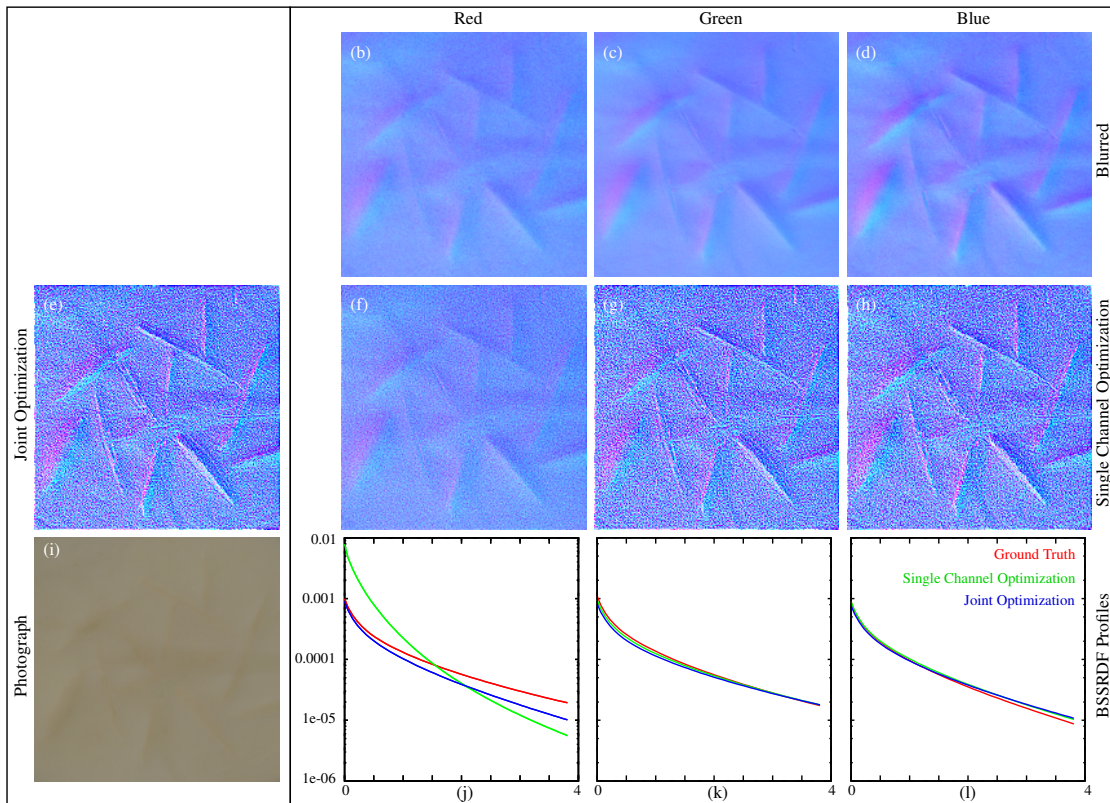
Figure 6. Result for a Potato (i). For this example, no ground truth normals are available. **(b-d)** Naively applying photometric stereo directly on the observations yields blurred normals. **(f-g)** Result from our method computed on each channel seperately. **(e)** Result from joint optimization on all color channels simultaneously. **(j-l)** Corresponding ground truth and recovered BSSRDF profiles.

Illumination. In *ICCV*, pages 1–8, Nov 2011.

[12] R. Habel, P. H. Christensen, and W. Jarosz. Photon beam diffusion: A hybrid monte carlo method for subsurface scattering. *Comput. Graph. Forum*, 32(4):27–37, 2013.

[13] P. Hanrahan and W. Krueger. Reflection from layered surfaces due to subsurface scattering. In *Proc. SIGGRAPH 93*, pages 165–174, 1993.

[14] T. Higo, Y. Matsushita, and K. Ikeuchi. Consensus photometric stereo. In *IEEE CVPR*, pages 1157–1164, 2010.

[15] R. Hooke and T. A. Jeeves. " direct search" solution of numerical and statistical problems. *J. ACM*, 8(2):212–229, Apr. 1961.

[16] H. W. Jensen, S. Marschner, M. Levoy, and P. Hanrahan. A practical model for subsurface light transport. In *Proc. SIGGRAPH 2001*, pages 511–518, 2001.

[17] D. Kundur and D. Hatzinakos. Blind image deconvolution. *IEEE Signal Processing Magazine*, 13(3):43–64, 1996.

[18] A. Levin, Y. Weiss, F. Durand, and W. T. Freeman. Understanding blind deconvolution algorithms. *IEEE Trans. Pattern Anal. Mach. Intell.*, 33(12):2354–2367, 2011.

[19] F. Lu, Y. Matsushita, I. Sato, T. Okabe, and Y. Sato. Uncalibrated photometric stereo for unknown isotropic reflectances. In *IEEE CVPR*, pages 1490–1497, 2013.

[20] W. Ma, T. Hawkins, P. Peers, C. Chabert, M. Weiss, and P. Debevec. Rapid acquisition of specular and diffuse normal maps from polarized spherical gradient illumination. In *Proc. Eurographics Symposium on Rendering*, 2007.

[21] K. D. Moore and P. Peers. An empirical study on the effects of translucency on photometric stereo. *The Visual Computer*, 29(6-8):817–824, 2013.

[22] D. Nehab, S. Rusinkiewicz, J. Davis, and R. Ramamoorthi. Efficiently combining positions and normals for precise 3d geometry. *ACM Trans. Graph.*, 24(3):536–543, 2005.

[23] F. E. Nicodemus, J. C. Richmond, J. J. Hsia, I. W. Ginsberg, and T. Limperis. Geometric considerations and nomenclature for reflectance. *National Bureau of Standards Monograph 160*, 1977.

[24] T. Papadhimitri and P. Favaro. A new perspective on uncalibrated photometric stereo. In *IEEE CVPR*, pages 1474–1481, 2013.

[25] B. Shi, P. Tan, Y. Matsushita, and K. Ikeuchi. Elevation angle from reflectance monotonicity: Photometric stereo for general isotropic reflectances. In *ECCV*, pages 455–468, 2012.

[26] R. J. Woodham. Photometric method for determining surface orientation from multiple images. *Optical Engineering*, 19(1):3050–3068, 1980.

[27] Y. Zhu, P. Garigipati, P. Peers, P. Debevec, and A. Ghosh. Estimating diffusion parameters from polarized spherical gradient illumination. *IEEE CG&A*, 99, 2013.