Image-based Separation of Diffuse and Specular Reflections using Environmental Structured Illumination

Bruce Lamond Pieter Peers Abhijeet Ghosh Paul Debevec
University of Southern California, Institute for Creative Technologies

Abstract

We present an image-based method for separating diffuse and specular reflections using environmental structured illumination. Two types of structured illumination are discussed: phase-shifted sine wave patterns, and phase-shifted binary stripe patterns. In both cases the low-pass filtering nature of diffuse reflections is utilized to separate the reflection components. We illustrate our method on a wide range of example scenes and applications.

1. Introduction

An important aspect in understanding a specific scene is how it interacts with incident illumination. The two most widely considered characteristics in reflectance behavior are diffuse and specular reflections. Being able to separate these two characteristics in a photographed scene, quickly and robustly, is of interest for numerous methods in computer graphics and computer vision, e.g., geometrical normal estimation, and reflectance measurement.

In this paper we consider the specific case where incident illumination originates from (multiple) distant sources. Often, such incident illumination is modeled in an image-based fashion by an environment map. Therefore, we call this kind of incident illumination *environmental incident illumination*. Unlike single light source conditions, specular highlights under environmental illumination conditions can cover large continuous areas of the object. Furthermore, different solid angles of incident illumination can significantly differ in color.

In the presented method, we employ a small set of high frequency structured environmental lighting conditions to separate diffuse and specular reflections from photographs of a scene taken from a single viewpoint. Our method is completely image-based, requires just a few photographs, and only makes mild assumptions about the frequency behavior of the reflectance properties in the scene. The core idea is that diffuse reflections act as a low-pass filter in the incident angular domain, and can be separated using a band-pass filter. Our method is aimed at separating both components with an equal degree of quality, as opposed to specular highlight removal methods which focus on the diffuse component only. This is useful for developing applications that estimate hybrid photometric normals that rely on either specular information such as [3, 5, 15, 30] or diffuse reflectance information [34], and for efficiently estimating reflectance characteristics.

To summarize the contributions of this work:

- First, it is to our knowledge the first separation method to use the *local* frequency behavior of the reflectance properties, and a theoretical justification is provided. It does not rely on color information or polarization states.
- Second, it is an image-based method where each pixel in the photograph is processed independently. No segmentation is required, nor is neighborhood information used. This allows straightforward component separation for textured objects.
- Finally, unlike most previous methods that only work for point light sources or constant area light sources, our method is able to separate reflection components under any (controlled) environmental illumination.

2. Previous Work

A large body of literature exists on separating reflectance components. These methods can be subdivided into four rough categories: polarization based methods [4, 15, 18, 33], color based methods [1, 9, 11, 13, 14, 16, 17, 27, 28, 29, 35], a combination of both [8, 19, 20, 32], or neighborhood (pixel or temporal) based methods [4, 7, 12, 24, 25, 26]. Discussing each technique in detail falls beyond the scope of this paper. Therefore, we will focus this discussion on the most relevant papers.

Shafer [27] introduced the dichromatic model: a colorless specular intensity and a diffuse chroma. To estimate the diffuse color, information from multiple pixels is used, and a global diffuse color is estimated. To deal with textured surfaces, a segmentation is performed, and for each

segment a diffuse color is computed. However, the need for segmentation limits the technique's effectiveness for highly textured objects. More recent work on color separation avoids segmentation by making additional assumptions such as known (or estimated) light source color(s) [16, 17]. Wolff and Boult [33] use polarization as an alternative tool to separate the reflection components. A key observation is that specular reflections become polarized, while diffuse reflections are basically unpolarized. Disadvantages are that multiple photographs are required (with different polarization directions), and that the amount of polarization is dependent on the angles of incidence and reflection. Nayar *et al.* [19] present the first system that combines color-space (for the diffuse component) and polarization based separation (for the specular component). However, their method only works well for dielectric specular reflections as opposed to metallic specular reflections.

None of the above methods are effective for separating reflection components under environmental illumination conditions which may be non uniformly colored. Due to their reliance on inter-pixel information or knowledge of the light source color, color-space methods will not work well when specular highlights cover large continuous areas of the object. These large highlit areas can cover large non uniformly colored solid angles and are unavoidable when using environmental illumination. Furthermore, color space methods generally fail in separating components with the same color. Polarization methods with unpolarized light sources require a large number of photographs (different polarization directions), and are not always equally effective for metallic specular reflections.

All of the methods above work under uncontrolled or semi-uncontrolled lighting conditions, making separation of diffuse and specular reflections a significantly under-constraint problem. Methods such as [4, 15, 21] use controlled illumination to separate diffuse or specular reflections. However, each of these methods rely on a (relative) sparse angular light source sampling, making them less suited for scenes containing sharp specular materials.

The method presented here does not rely on polarization states, or on inter-pixel relations, nor does it impose any restrictions on the color of the materials, light sources, or the viewpoint. It works equally well for metallic and dielectric specular reflections. It requires as few as three photographs of the scene under controlled illumination.

3. Problem Statement

Before detailing our separation method, we first formalize the problem and introduce necessary notations. Given diffuse and specular reflectance functions for each pixel $\mathbf{p}$ and incident lighting directions ω, denoted by $\mathbf{D}(\mathbf{p},\omega)$ and $\mathbf{S}(\mathbf{p},\omega)$ respectively, the goal is to compute the specular radiance $\rho_s(\mathbf{p})$, and diffuse radiance $\rho_d(\mathbf{p})$:

$$\rho_d(\mathbf{p}) = \int_\Gamma \mathbf{D}(\mathbf{p},\omega)\bar{\mathbf{I}}(\omega)\,d\omega, \quad (1)$$

$$\rho_s(\mathbf{p}) = \int_\Gamma \mathbf{S}(\mathbf{p},\omega)\bar{\mathbf{I}}(\omega)\,d\omega, \quad (2)$$

where $\bar{\mathbf{I}}(\omega)$ is the lighting environment for which the separation is desired, and Γ is the solid angle over which $\bar{\mathbf{I}}(\omega)$ is defined. Unless noted differently we will assume $\Gamma = \mathbf{\Omega}$ (the full sphere of incident lighting directions), and $\bar{\mathbf{I}}(\omega) = 1$.

In this paper we will compute $\rho_d(\mathbf{p})$, and $\rho_s(\mathbf{p})$ from a few observations $\mathbf{E}_i(\mathbf{p})$ of a scene taken from a fixed vantage point under different controlled distant environmental illumination conditions $\mathbf{I}_i$. Ignoring inter-reflections, and occlusions, the observed radiance $\mathbf{E}_i(\mathbf{p})$ under the i-th illumination condition $\mathbf{I}_i$ is governed by:

$$\mathbf{E}_i(\mathbf{p}) = \int_\Gamma \left(\mathbf{D}(\mathbf{p},\omega) + \mathbf{S}(\mathbf{p},\omega)\right)\,\mathbf{I}_i(\omega)\,d\omega, \quad (3)$$

The presented method will compute $\rho_d(\mathbf{p})$ and $\rho_s(\mathbf{p})$ for each pixel without using information from other pixels. Therefore, we drop the reference to a specific pixel $\mathbf{p}$ for notational simplicity.

We will now present two different sets of structured illumination conditions $\mathbf{I}_i$, and show how they can be used to extract the diffuse radiance ρ_d and specular radiance ρ_s from the observations $\mathbf{E}_i$.

4. Phase-shifted Sine Wave Patterns

It is well known that diffuse reflectance functions behave as a low-pass function on the incident illumination (e.g., [2, 23]). This allows us to design high frequency environmental illumination conditions that do not stimulate the diffuse reflectance function, and thus serve to separate diffuse and specular reflection components. A logical example of such an illumination condition is a (high frequency) sine wave pattern over the sphere of incident lighting directions.

4.1. Theory

Let us define different sine wave illumination patterns, each with a fixed frequency f, but with different phases φ_i: $\sin(f\phi + \varphi_i)$, with $\omega = (\theta,\phi)$, and $d\omega = \sin(\theta)d\theta d\phi$. The specific orientation of the azimuthal angle ϕ and polar angle θ does not matter, as long as they are consistently chosen. Note that these sine wave patterns are constant over θ. In order to obtain a physically realizable illumination condition, we add a constant offset, such that the complete illumination condition has positive values everywhere: $\mathbf{I}_i(\omega) = \sin(f\phi + \varphi_i) + 1$. Inserting this into Equation (3) gives:

$$\mathbf{E}_i = \int_\Gamma \left(\mathbf{D}(\omega) + \mathbf{S}(\omega)\right)\left(\sin(f\phi + \varphi_i) + 1\right)\, d\omega. \quad (4)$$

Since $\mathbf{I}_i$ is independent of θ, we include the integration over θ into the definition of $\mathbf{D}$, and $\mathbf{S}$, and focus on the integration over ϕ. This yields:

$$\begin{aligned}\mathbf{E}_i &= \rho_d + \rho_s + \int_\Gamma \mathbf{D}(\phi)\sin(f\phi + \varphi_i)d\phi \\ &+ \int_\Gamma \mathbf{S}(\phi)\sin(f\phi + \varphi_i)\, d\phi. \quad (5)\end{aligned}$$

If the frequency f is large enough, then the third term is approximately zero [23]. The fourth term can be written as:

$$\begin{aligned}&\textstyle\int_\Gamma \mathbf{S}(\phi)\sin(f\phi + \varphi_i)\, d\phi \\ &= \int_\Gamma \mathbf{S}(\phi)\left(\sin(f\phi)\cos(\varphi_i) + \cos(f\phi)\sin(\varphi_i)\right)\, d\phi \quad (6) \\ &= \left(\cos(\varphi_i)S + \sin(\varphi_i)C\right), \quad (7)\end{aligned}$$

where S and C are constants only dependent on the specific form of the specular reflectance function $\mathbf{S}(\cdot)$. Equation (7) defines a phase-shifted sine wave, with the phase directly related to the reflected view-direction and the surface normal. The amplitude of this sine wave corresponds to the magnitude of the f-th frequency component of $\mathbf{S}(\cdot)$. For sufficiently sharp specular reflections, and f low enough, the f-th frequency component is approximately ρ_s, and thus the sine wave defined by Equation (7) reaches its maximum amplitude ρ_s for some φ_{peak}.

Putting this all together yields:

$$\begin{aligned}\mathbf{E}_{peak} &\approx \rho_d + 2\rho_s, \\ \mathbf{E}_{peak+\pi} &\approx \rho_d. \quad (8)\end{aligned}$$

It is, however, unlikely that one of the phases φ_i equals exactly φ_{peak}, and thus $\mathbf{E}_{peak}$ is almost never directly observed. However, φ_{peak} can be computed from any collection of three phases φ_i, as described in the next subsection.

4.2. Implementation

Given three observations $\mathbf{E}_i$ of the scene under phase-shifted sine waves with phases $\varphi_i, i \in \{0, 1, 2\}$, we can compute ρ_s and ρ_d as follows. Each observation is of the following form:

$$\mathbf{E}_i = \rho_d + \rho_s + \left(\cos(\varphi_i)S + \sin(\varphi_i)C\right). \quad (9)$$

We can write this equation as a linear system with three unknowns A, S and C: $\mathbf{E}_i = A + \cos(\varphi_i)S + \sin(\varphi_i)C$, where $A = \rho_d + \rho_s$. If we have three or more linearly independent observations (e.g., with phases 0, 60 and 120 degrees), then all three unknowns can be computed by solving the linear system, or a linear least squares when more than three observations are made (e.g., for robustness). From this it can be easily shown that:

$$\begin{aligned}\rho_d &\approx A - \sqrt{S^2 + C^2}, \\ \rho_s &\approx \sqrt{S^2 + C^2}. \quad (10)\end{aligned}$$

Note that this approach is similar to phase shifting in structured light-based surface reconstruction.

The only question remaining is which frequency f to use. Ramamoorthi and Hanrahan [23] showed that the response of a diffuse reflectance function decays quadratically with respect to the frequency of the incident illumination. Therefore, to eliminate the effect of the diffuse reflections as much as possible, we recommend using $f > 10$ as a lower limit. The upper limit for the frequency f depends on the roughness of the specular reflectance function.

5. Phase-shifted Binary Stripe Patterns

A disadvantage of the phase-shifted sine wave patterns is that they require a radiometrically calibrated emitter capable of emitting a large number of intensity values. Often such a device is not available. Emission of binary patterns is, however, often easily achievable. An additional advantage of binary patterns is that no radiometric calibration (of the emitter) is required. Therefore, we investigate whether separation of diffuse and specular reflections can still be performed using only binary patterns: more specifically, phase-shifted binary stripe patterns. These binary stripe patterns can be viewed as a one-bit quantized version of the phase-shifted sine wave patterns of before.

Formally, define a set of phase-shifted binary stripe patterns as:

$$\mathbf{I}_i(\omega) = \Psi\left(\frac{f\phi + \varphi_i}{2\pi}\right) + 1, \quad (11)$$

with $\varphi_i = \frac{2\pi}{nf}i$, n the number of illumination patterns, and:

$$\Psi(x) = \begin{cases} +1 & \text{if } x - \lfloor x \rfloor \in [0, 0.5) \\ -1 & \text{if } x - \lfloor x \rfloor \in [0.5, 1). \end{cases} \quad (12)$$

Inserting these patterns into Equation (3), allows us to write the observed radiance similarly to Equation (5):

$$\begin{aligned}\mathbf{E}_i &= \rho_d + \rho_s + \int_\Gamma \mathbf{D}(\phi)\Psi\left(\frac{f\phi + \varphi_i}{2\pi}\right)d\phi \\ &+ \int_\Gamma \mathbf{S}(\phi)\Psi\left(\frac{f\phi + \varphi_i}{2\pi}\right)\, d\phi. \quad (13)\end{aligned}$$

We can follow a similar reasoning as with the phase-shifted sine wave patterns and equate the third term to zero.

Peak detection can then be used to obtain the following separation formula:

$$\begin{aligned} \rho_d &\approx \min_i \mathbf{E}_i, \\ \rho_s &\approx \tfrac{1}{2}\left(\max_i \mathbf{E}_i - \min_i \mathbf{E}_i\right). \end{aligned} \tag{14}$$

Please refer to [10] for an in depth derivation.

6. User-specified Environmental Illumination

Separating diffuse and specular reflections under uniform illumination is useful, however in many applications a separation under a particular environmental lighting condition is required (e.g., gradient illumination [15]). It can be shown that computing the maximum and minimum of the observations under phase-shifted binary stripe patterns modulated by the user-specified environmental illumination yields the desired separation.

Formally, given some user selected illumination condition $\bar{\mathbf{I}}$, and a set of phase-shifted binary stripe patterns $\mathbf{I'}_i, i \in \{0, ..., n-1\}$ with frequency f and phase φ_i, the desired separation is a diffuse component: $\rho_d = \int_\Gamma \mathbf{D}(\omega)\bar{\mathbf{I}}(\omega)\, d\omega$, and a specular component: $\rho_s = \int_\Gamma \mathbf{S}(\omega)\bar{\mathbf{I}}(\omega)\, d\omega$. Now define n novel illumination conditions $\mathbf{I}_i(\omega) = \bar{\mathbf{I}}(\omega)\mathbf{I'}_i(\omega)$. By restricting the selected illumination condition $\bar{\mathbf{I}}$ to not exhibit high frequency changes in intensity over a distance of $\frac{2\pi}{f}$ for every point in $\bar{\mathbf{I}}$, a successful separation can be obtained by computing both components as in Equations (3) and (14). Please refer to [10] for a more detailed derivation.

7. Separation Results

Our experiments use a setup similar to the device presented in [22], which covers approximately the front half of the sphere of incident lighting directions. The setup consists of a $1.5m$ diameter hemisphere coated with a rough specular surface, a radiometrically calibrated (i.e., gamma corrected) projector equipped with a fish-eye lens, and a radiometrically calibrated high resolution digital still camera. The projector is placed close to the focus of the hemisphere and aimed at the apex. The target object is placed near the lens of the projector. Finally, a camera is placed outside the dome and views the target object through a hole at the apex of the dome. Illumination patterns are emitted from the projector onto the dome, where they are subsequently reflected back towards the target object. The rough specular coating minimizes inter-reflections within the dome, while reflecting a maximum amount of light towards the target object. Note that our method is not limited to this acquisition device, and any extended controllable light source such as a CRT or TFT screen can be used.

A practical concern when projecting the phase-shifted sine-wave patterns, is to avoid the singularities at the poles. The first reason is that inaccuracies due to the limited resolution of the emitter can significantly impact accuracy in such a case. A second more fundamental issue is that while high frequencies are guaranteed to have a negligible impact on the diffuse reflections, there is no guarantee that it will have a non-zero impact on specular reflections. This is especially the case when the specular lobe is centered around the singularity. In order to avoid these problems, we directly emit a phase-shifted sine-wave pattern in the native parameterization of the emitter. While this does not exactly correspond to the theoretical patterns, we found that it still works well in practice, and avoids the singularity issues.

Figure 1 shows separations on four different objects. The leftmost example is obtained using phase-shifted sine wave patterns as in Section 4. The two middle examples are obtained using the phase-shifted binary stripe patterns as in Section 5. The rightmost example shows a separation result under user selected environmental illumination as in Section 6. The brightness of the individual components have been adapted for optimal viewing.

8. Discussion

The two presented methods have their advantages and disadvantages.

The phase-shifted sine wave patterns allow the peak phase to be computed from more than the minimum three measurements by solving a linear least squares system. This provides a convenient way to counteract potential measurement noise of the camera and brightness fluctuations due to the emitter. Because the computation of the components from the phase-shifted binary stripe patterns relies on the minimum and maximum observed intensities, the effect of measurement noise cannot be easily resolved by emitting more stripe patterns.

The binary stripe patterns, however, have the advantage that they do not require radiometric calibration of the emitter. Additionally, the computational simplicity and ease of implementation make this method very attractive. The cost of this simplicity is that this method is less forgiving to selecting an insufficient frequency range, although using more illumination patterns with different phases can partially resolve this. The phase-shifted sine wave patterns fail more gracefully due to their smoother appearance and analytical computation of the peak phase. This behavior is especially useful when separating objects that contain a wide range of specular reflectance functions (e.g., both dull and sharp reflectance functions in a single scene). Please refer to [10] for an in depth discussion and illustrations of the effects of using different frequency patterns.

Limitations: In our previous discussion we have ignored inter-reflections, and the effects of occlusions.

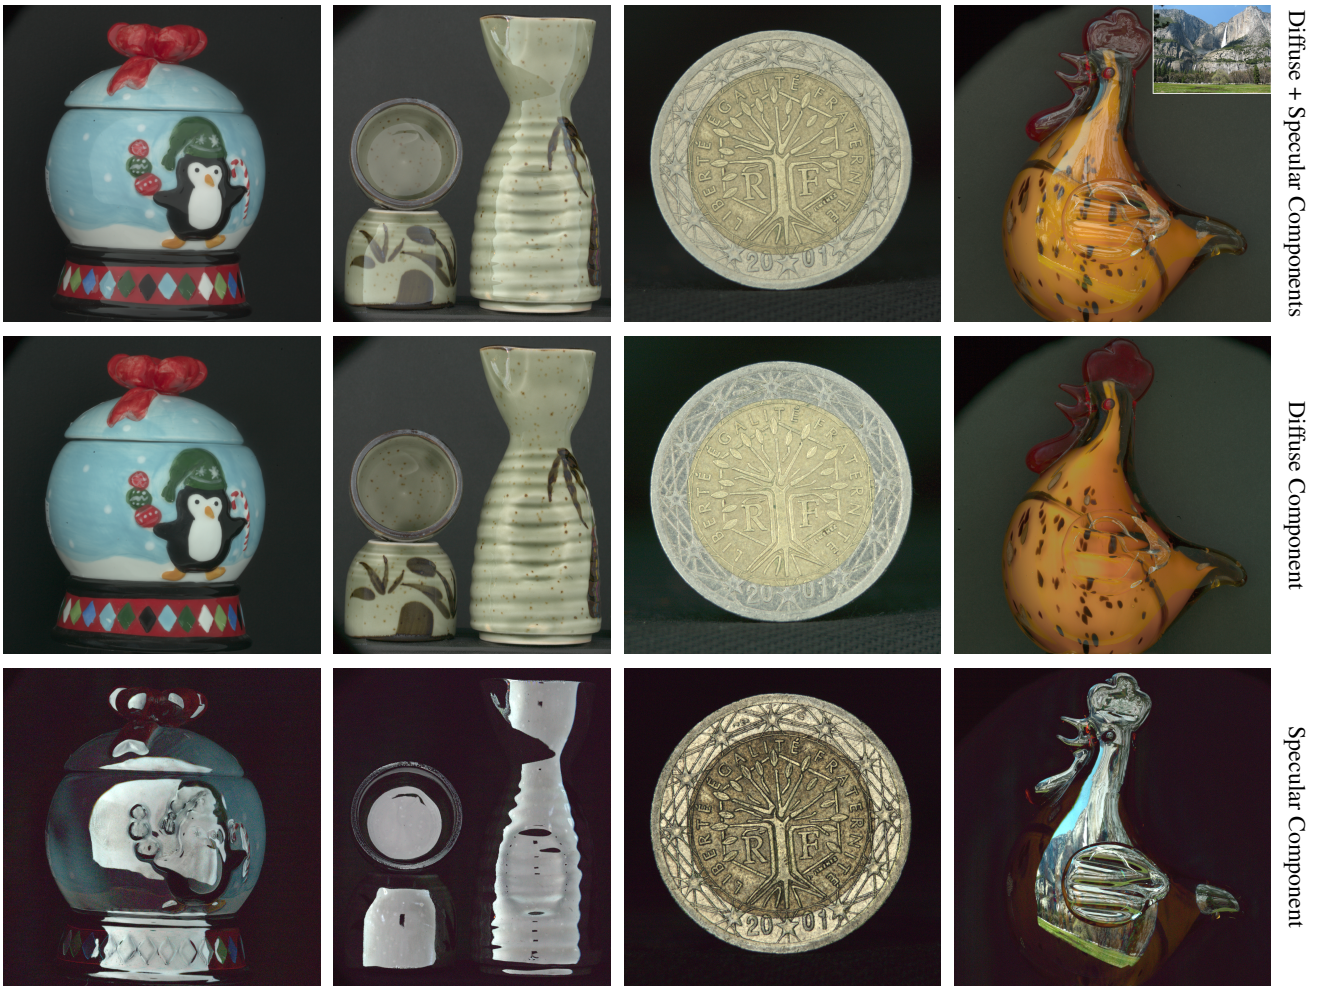

Figure 1. Results of our separation method. Left: separation using phase-shifted sine wave patterns. Middle (2): separation using phase-shifted stripe patterns. Right: separation under user-specified environmental illumination (shown in the inset).

Occlusions are light transport effects that change the frequency content of the reflectance functions. Hence, artifacts are introduced when separating. Increasing the frequency of the patterns helps to avoid these artifacts, but at the cost of potentially classifying parts of rough specular reflections as diffuse. Related to occlusions is the use of extended light sources that do not cover the whole sphere of incident lighting directions. Conceptually, lighting directions not covered by the light source can be considered occluded, and as such, similar artifacts occur.

Inter-reflections also influence the separation. Many previous methods only handle inter-reflections partially, or not at all. The presented method deals correctly with multiple specular reflections. Any diffuse inter-reflection, however, will result in an incorrect separation due to the low-pass behavior of diffuse reflections which removes any frequency information from the incident illumination.

Validation: We have compared the presented method to a polarization based separation approach. We covered a CRT monitor with a linear polarization sheet. A series of photographs with a linear polarizer rotated to a range of different orientations in steps of 6 degrees placed in front of the camera are captured of the object under constant illumination. Computing the minimum and maximum per pixel yields the cross polarized and parallel polarized photograph of the scene. Additionally, we also compute separations under phase-shifted stripe and sine wave patterns. We calibrated the patterns such that each period covers approximately the same solid angle when viewed from the object.

All three separations are comparable in quality. Differences were mainly due to non-specular polarization maintaining reflections, occlusions, and the limited coverage of the sphere of incident lighting directions by the CRT monitor. Please refer to [10] for a visual comparison.

Relation to Nayar *et al.* [21]: The presented method, especially the phase-shifted stripe patterns, bears similarities to [21] which presents a method to separate direct and indirect reflections from a point light source (e.g., a projector). However, this work differs fundamentally from our diffuse-specular separation method.

A important theoretical key difference is that we separate *local* reflection components (an interaction at a single surface point), while Nayar *et al.* separates *global* reflection effects (interactions between different surface points). Another practical difference is that Nayar *et al.* project illumination directly onto the scene from (approximately) a point light source. In our setup we use an area light source (indirectly via a hemispherical dome or the surface of a CRT monitor). Further, Nayar *et al.* consider the spatial frequency behavior of the incident illumination (from a single direction per camera pixel) on the scene, whereas our methods look at the angular frequency behavior of a single spatial location (i.e., pixel). Thus the two methods operate on different domains.

An interesting case is where the effects of local and global interactions overlap. Such a case occurs when the diffuse component is solely due to subsurface scattering. Specular reflections, on the other hand, are always first surface reflection effects, and thus correspond to directly reflected illumination. Consequently, for this case only, Nayar *et al.* are able to separate diffuse and specular reflections from a point light source. Correspondingly, our method is able to separate direct and indirect reflections from extended light sources. Generally, however, the separated components from these two techniques will not correspond to the same physical phenomena.

It should be noted that the projector resolution limits the effective resolution of the camera in Nayar *et al.*, and thus consequently the amount of detail that can be captured. The presented method on the other hand *requires* photographs of sufficiently high resolution such that the solid angle of a specular reflection through a pixel covers less then a single period of the phase-shifted patterns.

9. Applications

We envision a large number of applications that can benefit from using our separation method, which is especially suited for methods that require structured illumination to start with. Due to space restrictions, we will only discuss two applications: photometric normal estimation, and spatially-varying BRDF (SVBRDF) estimation. Our goal is to illustrate that our method is applicable in many situations. Comparing the quality to comparable alternative methods falls outside the scope of this paper.

Photometric Normal Estimation: A number of methods have been developed for estimating normals from diffuse and specular reflections. In our application, we use an evolution of example-based photometric stereo [6]. We compute the separated diffuse and specular components under three illumination conditions. Inspired by [15] we use linear gradient patterns. However, in order to avoid exact geometric calibration of the emitted pattern (which depends on the shape of the emitter and the relative location of the object), we also capture a diffuse and specular sphere under the same illumination conditions. For both components, we correspond for each pixel on the object its normalized intensity to the normalized intensity on the corresponding diffuse and specular sphere. The normalization ensures that we remove the effects of the albedo. The pixel location on the sphere of the matching normalized intensity, can then be directly translated into a normal direction.

Figure 2 shows false color images of the normals of the Sake set of Figure 1 computed from the separated diffuse and specular reflections. The inset shows the false color coding of normal directions on a sphere. Note that a low albedo aversely affects the diffuse normal estimation, while the specular normals do not suffer from this. Occlusions also impact the estimation of the normals differently for both components. In the diffuse case, normals are bent towards the occluder. Normals estimated from the specular component are affected less due to their more narrow angular support. However, the range of surface normals that can be estimated from specular reflections is currently limited by the angular extent of the hemispherical setup used.

SVBRDF Estimation: We can use our separation method to capture and fit SVBRDFs. For simplicity we assume that the SVBRDF is isotropic, and that the diffuse reflection is Lambertian. We use the separation under uniform illumination to compute the specular and diffuse albedo, and use the above photometric normal estimation technique to obtain surface normals. We fit a diffuse BRDF using the diffuse albedo for each pixel. Next, we sample the SVBRDF from 10 lighting directions and subtract for each pixel the effect of the diffuse BRDF. We further segment the sample manually into regions with different specular BRDFs. For each region a Torrance-Sparrow [31] lobe is fitted to the remaining specular reflections from the 10 sampled lighting directions. Figure 3 shows a visualization and a reference photograph of an SVBRDF. There are differences notable due to the fact that we use a directional light source for visualizing the SVBRDF, which underestimates the extent of the reference light source.

10. Conclusion

We have presented an image-based method that exploits specific frequency behaviors of reflectance functions for separating reflection components under environmental illumination. Our method does not rely on polarization, color-

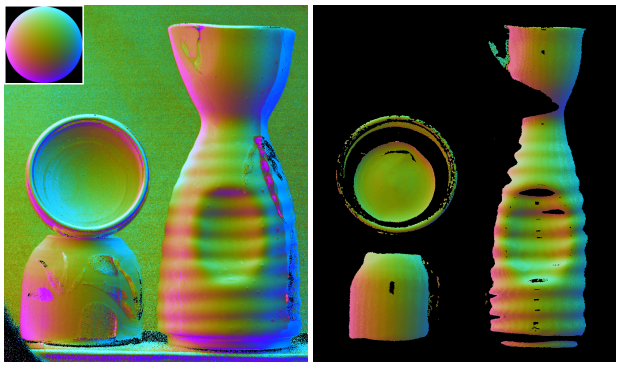

Figure 2. Normals computed from example-based photometric stereo based on the separate diffuse (left) and specular (right) components.

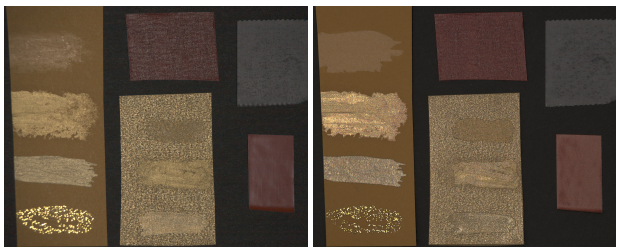

Figure 3. SVBRDF estimated with separated diffuse and specular albedos. **Left:** a reference photograph under a point light source. **Right:** a visualization of fitted SVBRDF parameters under similar lighting conditions.

space, or spatial/temporal neighborhood information. Two different sets of structured environmental illumination patterns were presented to achieve this separation. Finally, an extension was introduced to separate reflection components of a scene lit by a user-selected environmental illumination.

For future work we would like to improve the photometric stereo application presented in this paper. We also would like to investigate how to extend this method to uncontrolled incident light fields.

References

[1] R. Bajscy, S. Lee, and A. Leonardis. Detection of diffuse and specular interface reflections by color image segmentation. *IJCV*, 17(3), 1996. 1

[2] R. Basri and D. W. Jacobs. Lambertian reflectance and linear subspaces. *PAMI*, 25(2), 2003. 2

[3] T. Chen, M. Goesele, and H. P. Seidel. Mesostructure from specularities. In *CVPR*, pages 1825–1832, 2006. 1

[4] P. Debevec, T. Hawkins, C. Tchou, H.-P. Duiker, W. Sarokin, and M. Sagar. Acquiring the reflectance field of a human face. In *ACM SIGGRAPH*, 2000. 1, 2

[5] Y. Francken, T. Cuypers, T. Mertens, J. Gielis, and P. Bekaert. High quality mesostructure acquisition using specularities. In *CVPR*, 2008. 1

[6] A. Hertzmann and S. M. Seitz. Example-based photometric stereo: Shape reconstruction with general, varying brdfs. *PAMI*, 27(8):1254–1264, 2005. 6

[7] A. Jaklic and F. Solina. Separating diffuse and specular component of image irradiance by translating a camera. In *CAIP*, 1993. 1

[8] D. Kim and S. Lin. Variational specular separation using color and polarization. In *IAPR*, 2002. 1

[9] G. J. Klinker, S. A. Shafer, and T. Kanade. The measurement of highlights in color images. *IJCV*, 2, 1988. 1

[10] B. Lamond, P. Peers, A. Ghosh, and P. Debevec. Image-based separation of diffuse and specular reflections using environmental structured illumination, supplemental material. Technical report, USC, 2009. ICT-TR-01-2009. 4, 5

[11] S. Lee, H. Koo, N. Cho, and J. Park. Stochastic approach to separate diffuse and specular reflections. In *ICIP*, 2006. 1

[12] S. Lin and S. W. Lee. Estimation of diffuse and specular appearance. In *ICCV*, 1999. 1

[13] S. Lin, Y. Li, S. B. Kang, X. Tong, and H.-Y. Shum. Diffuse-specular separation and depth recovery from image sequences. In *ECCV*, 2002. 1

[14] S. Lin and H.-Y. Shum. Separation of diffuse and specular reflection in color images. *CVPR*, 2001. 1

[15] W.-C. Ma, T. Hawkins, P. Peers, C.-F. Chabert, M. Weiss, and P. Debevec. Rapid acquisition of specular and diffuse normal maps from polarized spherical gradient illumination. In *EGSR*, June 2007. 1, 2, 4, 6

[16] S. P. Mallick, T. Zickler, P. N. Belhumeur, and D. J. Kriegman. Specularity removal in images and videos: A pde approach. In *ECCV*, 2006. 1, 2

[17] S. P. Mallick, T. E. Zickler, D. J. Kriegman, and P. N. Belhumeur. Beyond lambert: Reconstructing specular surfaces using color. *CVPR*, pages 619–626, 2005. 1, 2

[18] V. Müller. Elimination of specular surface-reflectance using polarized and unpolarized light. In *ECCV*, 1996. 1

[19] S. Nayar, X. Fang, and T. Boult. Removal of Specularities using Color and Polarization. In *CVPR*, 1993. 1, 2

[20] S. K. Nayar, X.-S. Fang, and T. Boult. Separation of reflection components using color and polarization. *IJCV*, 21(3), 1997. 1

[21] S. K. Nayar, G. Krishnan, M. D. Grossberg, and R. Raskar. Fast separation of direct and global components of a scene using high frequency illumination. *ACM TOG*, 25(3), 2006. 2, 6

[22] P. Peers, T. Hawkins, and P. Debevec. A reflective light stage. Technical report, USC, 2006. ICT-TR-04-2006. 4

[23] R. Ramamoorthi and P. Hanrahan. The relationship between radiance and irradiance: Determining the illumination from images of a convex lambertian object. In *Journal Opt. Soc. Am.*, 2001. 2, 3

[24] R.Tan and K. Ikeuchi. Separating reflection components of textured surfaces using a single image. *PAMI*, 27(2), 2005. 1

[25] Y. Sato and K. Ikeuchi. Temporal-color space analysis of reflection. *Journal Opt. Soc. Am.*, 11, 1994. 1

[26] Y. Sato, M. D. Wheeler, and K. Ikeuchi. Object shape and reflectance modeling from observation. *Computer Graphics*, 31, 1997. 1

[27] S. Shafer. Using color to separate reflection components. *COLOR Research and Applications*, 10(4), 1985. 1

[28] R. Tan and K. Ikeuchi. Reflection components decomposition of textured surfaces using linear basis functions. In *CVPR*, 2005. 1

[29] R. T. Tan, K. Nishino, and K. Ikeuchi. Separating reflection components based on chromaticity and noise analysis. *PAMI*, 26(10), 2004. 1

[30] M. Tarini, H. P. A. Lensch, M. Goesele, and H.-P. Seidel. 3d acquisition of mirroring objects. *Graphical Models*, 67(4), 2005. 1

[31] K. E. Torrance and E. M. Sparrow. Theory of off-specular reflection from roughened surfaces. *J. Opt. Soc. Am.*, 57:1104–1114, 1967. 6

[32] S. Umeyama and G. Godin. Separation of diffuse and specular components of surface reflection by use of polarization and statistical analysis of images. *PAMI*, 26(5), 2004. 1

[33] L. Wolff and T. Boult. Constraining object features using polarization reflectance model. *PAMI*, 13(7), 1991. 1, 2

[34] R. J. Woodham, Y. Iwahori, and R. A. Barman. Photometric stereo: Lambertian reflectance and light sources with unknown direction and strength. Technical Report TR-91-18, 1991. 1

[35] K. Yoon, Y. Choi, and I. Kweon. Fast separation of reflection components using a specularity-invariant image representation. In *ICIP*, 2006. 1