Diffuse-Specular Separation using Binary Spherical Gradient Illumination

Christos Kampouris

Stefanos Zafeiriou

Abhijeet Ghosh

Imperial College London

Abstract

We introduce a novel method for view-independent diffuse-specular separation of albedo and photometric normals without requiring polarization using binary spherical gradient illumination. The key idea is that with binary gradient illumination, a dielectric surface oriented towards the dark hemisphere exhibits pure diffuse reflectance while a surface oriented towards the bright hemisphere exhibits both diffuse and specular reflectance. We exploit this observation to formulate diffuse-specular separation based on color-space analysis of a surface's response to binary spherical gradients and their complements. The method does not impose restrictions on viewpoints and requires fewer photographs for multiview acquisition than polarized spherical gradient illumination. We further demonstrate an efficient two-shot capture using spectral multiplexing of the illumination that enables diffuse-specular separation of albedo and heuristic separation of photometric normals.

CCS Concepts

•Computing methodologies \rightarrow Computational photography; Reflectance modeling;

1. Introduction

Acquiring realistic models of real-word subjects including their shape and appearance has been a very active area of research in graphics and vision. While passive acquisition systems under uniform illumination have been the most common for multiview acquisition of 3D geometry and albedo texture, active illumination systems have usually been the preferred choice for reflectance acquisition. Recent research in active illumination has focused on fast acquisition for live subjects while still trying to obtain much of the appearance information required for realistic rendering of the acquired data. Specifically for facial capture, the active illumination technique of polarized spherical gradient illumination [MHP*07, GFT*11] remains the state of the art in terms of the quality of the acquired reflectance information including separated diffuse and specular albedo and surface normals, while allowing acquisition using a handful of photographs that can be rapidly acquired. The method usually requires a polarized LED sphere where every single light has a linear polarizer mounted on it in a specific orientation. The technique has also been recently extended to a single shot method using an even more sophisticated setup involving polarized RGB lights and multiple polarized cameras [FD15], with application in dynamic facial performance capture.

In this work, we present a novel variant of such active illumination for acquiring separated diffuse and specular reflectance using binary spherical gradient illumination. Similar to previous work, we employ an LED sphere to illuminate a subject with binary gradients and their complements. However, unlike previous work, we

© 2018 The Author(s) Eurographics Proceedings © 2018 The Eurographics Association. employ unpolarized illumination and instead rely on an inherent property of the binary gradients to enable diffuse-specular separation of acquired albedo and surface normals. In this respect, our method requires a simpler setup which does not require polarizing multiple (usually hundreds of) light sources, it requires fewer photographs and does not impose any restrictions on the acquisition viewpoint for diffuse-specular separation. Besides achieving albedo separation using a novel color-space analysis on the binary spherical gradients, we also demonstrate separation of photometric normals with robust estimation of a specular normal map from the acquired data (see Fig. 1). This is a key result of our method which we believe to also be somewhat surprising given the binary nature of the incident illumination for specular reflectance. We further demonstrate how spectral multiplexing can be exploited in conjunction with binary spherical gradient illumination to enable a very efficient two-shot acquisition approach, with potential future application in dynamic capture.

The rest of the paper is organized as follows: we first review some related work in Section 2 before describing our setup and acquisition in Section 3. We then describe our approach for diffusespecular separation of albedo (Section 4) and surface normals (Section 5) acquired using binary spherical gradients and their complements. We next present our spectral multiplexing approach for a two-shot acquisition in Section 6, before discussing a few additional results and comparisons in Section 7.

C. Kampouris, S. Zafeiriou, A. Ghosh / Diffuse-Specular Separation using Binary Spherical Gradient Illumination



Figure 1: Acquired images of a face lit with binary spherical gradient illumination and their complements (a - f), and estimated diffuse and specular albedo (g, h) and photometric normal maps (i, j) from the acquired data which enables realistic relighting with hybrid normal rendering (k).

2. Related Work

While there is significant literature on shape and appearance acquisition in vision and graphics, we will restrict the discussion here to active illumination methods.

Photometric acquisition: Woodham [Woo78] first introduced the technique of photometric stereo for estimating shape and reflectance (albedo) of a diffuse object using multiple (at least three) observations of the object lit with a point light source from a few independent directions. The approach was further extended to single shot acquisition of surface normals using RGB color-multiplexing, and this approach has been applied successfully for dynamic acquisition [HVB*07, KHE10]. The approach however assumes pure diffuse reflectance and usually restricts acquisition from a preferred frontal viewpoint which can see the response to all the light sources in a setup. Photometric stereo has been extended for non-lambertian objects using example based normal inference using a known (spherical) exemplar of similar material [GCHS05]. Mallick et al. [MZKB05] proposed a color-space separation technique for dielectric materials that linearly transforms RGB color into a formulated suv space where s aligns with white (color of light) and u, v align with chroma. They employ the u, v components for improved photometric stereo discarding specularities. We exploit such color-space analysis in this work for estimation of separated diffuse normals. Color-space analysis has also been employed by Lin&Shum [LS01] to remove specular reflectance for traditional point source photometric stereo assuming dielectric material reflectance. Their analysis is based on the observation of brighter measurements for a pixel containing some specular response compared to darker measurements which are assumed to be pure diffuse reflectance. We employ a similar observation in this work to photometric analysis of binary spherical gradient illumination.

Facial capture: Debevec et al. [DHT*00] first introduced a light stage apparatus for acquiring facial reflectance fields. Subsequently, LED spheres were employed for photometric stereo [WGT*05] and facial reflectance acquisition [WMP*06] under a dense set of incident lighting directions. Ma et al. [MHP*07] introduced polarized spherical gradient illumination (using an LED sphere) for efficient acquisition of the separated diffuse and specular albedos and photometric normals of a subject using just eight photographs, and demonstrated high quality rendering with the acquired data using the acquired hybrid normal maps. Their employed linear polarization pattern on the LED sphere however restricted acquisition to a single frontal viewpoint. Ghosh et al. [GFT*11] subsequently extended the view-dependent solution of Ma et al. for multi-view facial acquisition with polarized spherical gradient illumination. Their method switches the linear polarization state on the LED sphere (latitude vs longitue) instead of on the camera as required by Ma, providing a symmetric solution about the Y axis of the LED sphere. Additionally, Ghosh et al. employed both the gradients and their complements in the multiview acquisition process for greater SNR of the acquired data as well as for assistance in optical flow based motion alignment as proposed by Wilson et al. [WGP*10]. Fyffe et al. [FHW*11] applied the complimentary spherical gradient illumination based alignment of Wilson et al. in conjunction with high speed photography to acquire longer facial performance sequences. However, they acquired unpolarized spherical illumination and instead applied a heuristics based diffuse-specular separation on the acquired data to obtain separated reflectance maps for rendering. More recently, Fyffe & Debevec [FD15] have proposed spectral multiplexing with polarized spherical gradient illumination (using an RGB LED sphere) for single shot acquisition of separated diffuse and specular albedo and normals. Besides a polarized RGB LED sphere, their method requires multiple polarized cameras per viewpoint (parallel and cross-polarized) to obtain the desired reflectace and normal maps which makes it a somewhat complex setup. Another recent work employs single shot acquisition with twenty four cameras and six flashes in a multiview setup [FGT*16]. Here, a subset of cameras fire together with each flash in a sequence lasting a few milliseconds to obtain a dense set of multiview images lit with flashes from different directions. The data enables reconstruction of high resolution facial geometry and estimation of diffuse and specular reflectance maps using a combination of diffuse and specular photometric stereo and inverse rendering. While the approach is a middle ground between completely passive acquisition and multi-shot active illumination approaches, it requires more involved processing of acquired data. Researchers have also looked at more detailed separation of facial reflectance using a combination of polarized spherical illumination and structured lighting to estimate layered skin reflectance [GHP*08]. This approach has recently been reduced to single-shot acquisition using polarized light field imaging [KIG16], albeit with reduced spatial resolution of the acquired maps.

Frequency domain separation: Lamond et al. [LPGD09] have proposed employing high frequency phase shifting of incident environmental illumination as a mechanism for diffusespecular separation and employ a reflective hemispherical dome to illuminate objects with such structured illumination. Compared to a discrete LED sphere, their setup is suitable for highly specular objects due to the continuous nature of the incident illumination. Tunwattanapong et al. [TFG*13] instead employ a rotating semicircular LED arc to illuminate specular objects with continous spherical illumination. They further illuminate objects with higher order spherical harmonics using their setup for frequency domain diffuse-specular separation and reflectance estimation. Researchers have also investigated various types of illumination using an LCD panel to emit varying frequencies on planar samples including second order (polarized) gradients [GCP*09], gray codes [FCMB09], and band-limited Fourier patterns [AWL13]. Given the involved setups, these approaches are however not suited for acquisition of human subjects.

Binary spherical gradients: Closest to our work, Vlasic et al. [VPB*09] first proposed applying binary spherical gradient illumination for estimating surface normals for a full body performance capture application. Instead of directly estimating surface normals from the acquired data, they employed example based normal inference using a diffuse sphere due to the specific geometry of the LED sphere employed for the acquisition. Importantly, they assumed diffuse reflectance and did not investigate diffuse-specular separation with the acquired data. In this work, we employ unpolarized binary spherical gradient illumination



Figure 2: Acquisition setup consisting of an LED sphere (a), employed for binary spherical gradient illumination as seen on a mirror ball (b - g).

similar to [VPB*09], but demonstrate high quality diffuse-specular separation of albedo and surface normals for dielectric materials such as skin. We further demonstrate novel combination of the method with spectral multiplexing of the illumination which enables an efficient two-shot acquisition.

3. Acquistion

We employ an LED sphere for acquisition in this work consisting of 168 white LED lamps (ColorKinetics iW MR gen3) mounted on a 2.5m diameter spherical structure (see Fig. 2). The lamps are mounted on the structure along longitudes and latitudes. This helps simplify partitioning of the lights for each of the binary spherical gradient illumination conditions. Exactly half of the lights (84) participate in each lighting condition, with partitioning along two orthogonal longitudes for the X and Z gradients and their complements, and partition along the equatorial latitude for the Y gradient and its complement. We rapidly acquire six photographs of a subject under each of the binary gradients and their complements using a consumer DSLR camera (Canon EOS 750D with a 18-135mm lens) operating in burst mode. The LED sphere also has an additional set of 168 RGB LED lamps (ColorKinetics iColor MR gen3) mounted adjacent to the white LED lamps. We will later describe how we employ the RGB lamps for spectral multiplexing of the binary gradients for a two-shot acquisition in Section 6. In the following, we first describe how we employ the data acquired under white LED illumination for diffuse-specular separation.

C. Kampouris, S. Zafeiriou, A. Ghosh / Diffuse-Specular Separation using Binary Spherical Gradient Illumination

4. Albedo Separation

Having acquired the six images of a subject under binary spherical gradients and their complements, we first formulate separation of albedo under uniform illumination into diffuse and specular reflectance. Note that due to the nature of the complementary gradient lighting conditions, the sum of any two complementary lighting conditions equals the uniform (full-on) illumination condition which provides us with an estimate of the mixed albedo containing both the diffuse and specular reflectance. However, we make the observation that with the complementary gradients, a surface on a convex dielectric object exhibits pure diffuse reflectance when lit with the dark (off) hemisphere while exhibiting both diffuse and specular reflectance when lit with the bright (on) hemisphere of the binary gradients. This is due to the wide hemispherical integration lobe of diffuse reflectance (centered around the surface normal) which continues to integrate some light from the bright hemisphere even when a surface faces away towards the dark hemisphere while receiving no light only when facing exactly in the direction of the dark side. On the other hand, typical specular reflectance on a dielectric object has a much narrower lobe of reflection (centered around the reflection vector) which integrates light only when facing the bright side of the binary gradient. We employ this key observation in this work to formulate a novel color-space separation for the albedo under uniform illumination.

Our color-space separation proceeds as follows: for colored dielectric materials, specular reflectance is white in color when a surface is lit with white light. Hence, when a colored dielectric surface is lit with a binary spherical gradient and its complement, the reflected RGB color of the surface is more saturated in the darker (pure diffuse) lighting condition while exhibiting reduced saturation under the brighter (diffuse + specular) lighting condition due to mixing with some white reflection. Therefore, we first convert observed RGB colors of a surface under a binary spherical gradient and its complement $(R_g, G_g, B_g \text{ and } R_c, G_c, B_c \text{ respectively})$ into HSV space (H_g, S_g, V_g and H_c, S_c, V_c respectively). Assuming the gradient is brighter than its complement $(V_g > V_c)$, we formulate the amount of specular reflectance mixed in the bighter gradient to be equal to the amount $\boldsymbol{\delta}$ that needs to be subtracted from R_g, G_g, B_g such that its new saturation S'_g equals the saturation of the complementary gradient S_c . Formally, this can be expressed as:

$$S_c = S'_g = \frac{C_g}{V_g - \delta},\tag{1}$$

where C_g is the chroma (related to hue H_g) of the brighter (R_g, G_g, B_g) measurement. This leads to a closed-form solution for δ (Equation 2) which can be easily computed for every surface point seen under a binary spherical gradient and its complement.

$$\delta = V_g - \frac{C_g}{S_c}.$$
 (2)

Since we illuminate a subject with the three principal gradients (X, Y and Z) and their complements, we obtain three separate estimates of δ with this approach. These three separate estimates



(a) mixed albedo (b) spec. albedo

Figure 3: Mixed albedo (a), estimated from the sum of a gradient and its complement, contains both diffuse and specular reflectance. Specular reflectance (b) is estimated from the difference in color saturation in reflectance under a gradient compared to its complement. Diffuse albedo (c) is obtained by subtracting (b) from (a).

are useful to have since each estimate on its own exhibits a region of uncertainty at the boundary of the gradient and its complement where the estimate is noisy and imprecise due to insufficient RGB color differences. Fortunately, these uncertainties are in different regions over a sphere for each of the three principal gradients and we employ the median of the three estimates as our final robust estimate of observed specular reflectance. The estimate from a gradient-complement pair can be unstable when one of the measurements is too dark. We simply ignore any such outliers. Given that the bright side of the binary gradient is lit with uniform constant illumination, the separated specular reflectance δ is proportional to the specular albedo ρ_s modulated by Fresnel gain at grazing viewing angles. We simply subtract the estimated specular reflectance ρ_s from the mixed albedo (obtained from the sum of a gradient and its complement) to obtain our separated diffuse albedo ρ_d (see Fig. 3).

5. Separating Photometric Normals

Once we have separated the albedo observed under uniform (fullon) illumination into diffuse and specular components, we can proceed with separating photometric surface normals obtained using binary spherical gradients into pure diffuse and specular normals. Similar to the approach of Wilson et al. [WGP*10], the binary spherical gradients and their complements can be directly employed to compute photometric surface normals given as:

$$\vec{N}_{mixed} = \frac{(X - X', Y - Y', Z - Z')}{||(X - X', Y - Y', Z - Z')||}$$
(3)

5.1. Diffuse Normal

Consistent with the observations of Ma et al. [MHP*07], for a face we observe different amounts of surface details in these photometric normals computed from the red, green, and blue channels respectively due to the different amounts of subsurface scattering of light in skin at these wavelengths (see Fig. 4). However, these computed photometric normals \vec{N}_{mixed} contain a mix of both diffuse and specular reflectance and hence are not ideal for direct usage in

C. Kampouris, S. Zafeiriou, A. Ghosh / Diffuse-Specular Separation using Binary Spherical Gradient Illumination



Figure 4: Wavelength dependent mixed photometric normals estimated from the binary spherical gradients and their complements. For human skin, the surface details are increasingly visible for shorter wavelengths as seen in the crops around the eye.

rendering. Instead, we propose to separate the photometric normal into a pure diffuse photometric normal and a pure specular photometric normal for shading with a diffuse and a specular BRDF respectively. For this purpose, we exploit the color-space analysis of Mallick et al. [MZKB05] and transform the acquired binary spherical gradients and their complements from RGB to Mallick's proposed suv space. This transformation rotates the RGB color information to align any white component in the original RGB signal with the s component, while leaving the chroma in the u, v components. The chroma in colored dielectric materials is the result of absorption due to subsurface scattering and hence the u, v components can now be employed to recompute the photometric normals according to Equation 3 using the pure diffuse component of the original *RGB* signal given as $\sqrt{u^2 + v^2}$. This provides us with an estimate of a pure diffuse photometric normal which is very smooth in its appearance due to blurring of surface details with subsurface scattering (see Fig. 5, a).

5.2. Specular Normal

Subsequently, we propose to also estimate a specular photometric normal from the acquired data. We observe that surface details are clearly visible in the acquired mixed photometric normals, particularly in the shorter green and blue wavelengths. This implies that the mixed normals encode a mixture of a diffuse normal and some specular reflectance, with the contribution of the specular information in the mixture being strongest in the blue channel for skin. This enables us to formulate specular normal \vec{N}_{spec} estimation using the following simple mixture model:

$$\vec{N}_{mixed} \propto \alpha \cdot \vec{N}_{u,v} + (1-\alpha) \cdot \vec{R}_{spec}$$
 (4)

Equation 4 formulates the acquired mixed photometric normal \vec{N}_{mixed} to be the result of a linear combination of the pure diffuse





(a) *u*, *v* normals (b) init. spec. normals (c) corr. spec. normals

Figure 5: (a) Color-space separation and estimation of pure diffuse (u, v) normals from acquired binary spherical gradients and their complements. (b) Directly estimated specular normals using the acquired mixed (blue) photometric normal (Fig. 4, c) and the separated u, v diffuse normal. (c) Final synthesized specular normals with low-frequency correction.

u, *v* normal $\bar{N}_{u,v}$ (obtained using color-space separation) and an unknown reflection vector \vec{R}_{spec} . Here, the weights of the linear combination α is obtained from the relative strengths of the diffuse and specular reflectance ρ_d and ρ_s respectively: $\alpha = \frac{\rho_d}{\rho_d + \rho_s}$. α is computed using data from only one employed color channel for the mixed normal (e.g., blue). The unknown reflection vector can then be estimated as:

$$\vec{R}_{spec} = \frac{\vec{N}_{mixed} - \alpha \cdot \vec{N}_{u,v}}{||\vec{N}_{mixed} - \alpha \cdot \vec{N}_{u,v}||}.$$
(5)

Note that in Equation 5, the division by $(1 - \alpha)$ factor from Equation 4 is not required due to the normalization step. We found that omitting the division by $(1 - \alpha)$ factor is also beneficial for numerical stability. Finally, the specular normal \vec{N}_{spec} can be computed as the half-vector given the viewing direction \vec{V} the estimated reflection vector \vec{R}_{spec} :

$$\vec{N}_{spec} = \frac{\vec{V} + \vec{R}_{spec}}{||\vec{V} + \vec{R}_{spec}||}.$$
(6)



Figure 6: Specular normals directly estimated with binary spherical gradients exhibit mostly eight discrete spherical quadrants with a small gradient across the quadrants due to non-linear Fresnel gain. (a) Simulated dielectric sphere. (b) A real orange.

C. Kampouris, S. Zafeiriou, A. Ghosh / Diffuse-Specular Separation using Binary Spherical Gradient Illumination



(a) binary RGB grad. (b) binary RGB comp.

Figure 7: Spectrally multiplexed binary RGB gradients (a), and binary RGB complements (b) acquired using two-shots. Top-row: Mirror ball. Bottom-row: Face.

As can be seen in Fig. 5 (b), the specular normal \vec{N}_{spec} estimated using this simple procedure exhibits very fine scale skin surface details including fine scale wrinkles and pores, and scarring on the forehead that are not visible in the pure diffuse normal. However, we note that the specular normal directly obtained with the above procedure also suffers from significant low-frequency bias. This is because of the employment of binary gradient illumination for estimating the specular reflection vector which results in eight discrete quadrants over the sphere of directions (see Fig. 6, a). Employing Equation 5 to estimate the specular normal adds a little bit of smooth variation (gradient) across the discrete quadrants which can be more clearly seen on a real acquired spherical object - an orange (Fig. 6, b). This slight gradient across the eight primary directions is the result of non-linear Fresnel gain in the specular reflectance (as a function of $\vec{N} \cdot \vec{V}$) which is encoded both in the mixed normal \vec{N}_{mixed} and the mixing weight α (due to Fresnel being baked in ρ_s) in the mixture model. This Fresnel resultant gradient in conjunction with surface specular roughness helps create a desirable smooth variation in the high frequency of the estimated specular normal field. However, since these normals suffer from significant low-frequency bias, we cannot directly employ them as specular normals for rendering purposes.

Instead, we employ a practical solution to synthesize a specular normal with the correct low-frequency while borrowing the high frequency details from the directly estimated specular normal vector \vec{N}_{spec} . We do this by applying a high-pass filter to \vec{N}_{spec} and adding the high-pass to the low-frequency diffuse normal $\vec{N}_{u,v}$ obtained using color-space separation, and finally re-normalizing the resultant normal map. Fig. 5 (c) shows the final synthesized specular normal obtained using this procedure which is suitable for rendering.



(a) mixed albedo

(b) mixed normals

Figure 8: *Mixed albedo (a) and normals (b) computed from twoshots acquired with spectrally multiplexed binary RGB gradient illumination.*

6. Spectral Multiplexing

The previous sections described diffuse-specular separation with six separate measurements of binary spherical gradients under white illumination. We now describe how spectral multiplexing can be combined with binary spherical gradient illumination for an efficient two-shot acquisition approach. For this purpose, we employ the RGB LED lamps on our LED sphere to simultaneously emit the three principal binary spherical gradients along X, Y and Z multiplexed in individual RGB color channels. We encode the X, Y and Z gradients in G, B and R channels respectively. In a second shot, we acquire the complementary gradients X', Y' and Z' also encoded along G, B, and R color channels respectively (see Fig. 7). Note that we employ the proposed $XYZ \rightarrow GBR$ mapping because we found this to provide the highest surface details in the photometric normals with X and Y variation encoded in shorter wavelenghts with lower amounts of subsurface scattering and Z variation encoded in the longer wavelength (red) which exhibits higher amounts of scattering in skin.

The acquired data can now be employed to estimate albedo and photometric normals as follows: simply adding the two images acquired under RGB gradients and their complements yields an image acquired under full-on white illumination and is an estimate of the mixed albedo (Fig. 8, a). Similarly, doing a component-wise subtraction of the RGB complement from the RGB gradient followed by normalization yields the mixed photometric normals (Fig. 8, b).

Unlike with measurements under white illumination, we cannot apply the previous color-space (suv) separation to the photometric normals estimated with RGB illumination. Instead, we now apply a heuristic approach for separation of the mixed photometric normals. We simply blur the acquired photometric normals to obtain an approximation of a diffuse normal map (Fig. 9, a). Here, we empirically selected the blur kernel to produce a smooth normal map that is qualitatively similar to the u, v diffuse normals obtained with color-space separation. In order to now estimate a specular normal map, we simply subtract the blurred diffuse normals from the mixed photometric normals to isolate the high frequencies in the mixed normals (mostly due to skin pores and fine wrinkles). We then accentuate these high frequencies by scaling them up and adding back to the blurred diffuse normal to obtain a sharper pho-



(a) synth. diff. normals

nals (b) synth. spec. normals

Figure 9: Empirically synthesized diffuse (a) and specular (b) normals from the mixed photometric normals obtained using binary *RGB* gradient illumination.



(a) spec. albedo

(b) diff. albedo

Figure 10: Separated specular (a) and diffuse (b) albedo from solution to a linear system formulated from the measurements under binary RGB gradient illumination.

tometric normal map which we can employ as specular normals for rendering (Fig. 9, b). This heuristic approach, while not physically accurate, qualitatively mimics the desired photometric diffuse and specular normals obtained with measurements under white illumination (Fig. 5, Section 5).

We now describe how the spectrally multiplexed binary spherical gradient illumination can be employed for diffuse-specular separation of the mixed albedo shown in Fig. 8 (a). Instead of formulating the albedo separation with color-space analysis as previously, we formulate it now as a linear system. We treat the data from each color channel independently in the two acquired photographs to obtain two observations under every principal binary gradient G and its complement C. Without loss of generality, assuming a surface appears brighter under G and darker under C, we can formulate the diffuse-specular separation for that surface as a solution to the following linear system:

$$\begin{bmatrix} G \\ C \end{bmatrix} = \begin{bmatrix} N_d & 1 \\ (1 - N_d) & 0 \end{bmatrix} \begin{bmatrix} D \\ S \end{bmatrix}$$
(7)

Here, D is the diffuse component and S the specular component of the signal in each color channel, and N_d is the corresponding

© 2018 The Author(s) Eurographics Proceedings © 2018 The Eurographics Association. component of the diffuse normal (shifted to [0,1]) for that principal direction. We employ the synthesized diffuse normal obtained from blurring the mixed photometric normal (Fig. 9, a) for computing the weights N_d in the linear system. We separately solve Equation 7 three times, once for each color channel, to obtain three different estimates of D and S. Assuming monochromatic specular reflection, we can then just compute the final estimate of specular reflectance as ρ_s as the median of the three solutions for S. Fig. 10 (a) shows the specular albedo obtained from solving this linear system. We simply subtract the specular albedo from the mixed albedo to obtain an estimate of the separated diffuse albedo (Fig. 10, b).

7. Additional Results and Discussion

Fig. 11 presents a comparison of the separated diffuse and specular albedo and photometric normals acquired using binary spherical gradient illumination (top-row) vs. acqusition of such data using polarized spherical gradient illumination [GFT*11] (bottomrow) for an orange rubber ball and a face. We polarized the LED sphere with the latitude-longitude linear polarization pattern proposed by Ghosh et al. for this comparison, and acquired data for the binary spherical gradients with both polarization banks switched on to achieve unpolarized illumination. As can be seen, the separated albedo and photometric normals acquired using our proposed method using binary gradients are qualitatively quite comparable to those acquired using polarized spherical gradients. Furthermore, the data acquired requires half the number of photographs (six) compared to that required with polarized spherical gradients for multiview capture. However, the quality of albedo separation with binary gradients suffers from a bit more noise due to lower specular signal in the acquired data with unpolarized illumination compared to when employing polarized illumination. The albedo separation is also not as accurate over the white logo on the ball due to limitation of color-space analysis for white materials.

We also computed quantitative differences between the acquired maps of the orange ball using both the methods. These are as follows: diffuse albedo = 0.064 RMSE; specular albedo = 0.02 RMSE, diffuse normal = 8.93° angular RMSE, specular normal = 14.21° angular RMSE. Here, we compared the u, v diffuse normal obtained with our method against the red diffuse normal obtained with polarized spherical gradients because they are qualitatively the most similar in terms of blurring due to subsurface scattering. We note that the chroma based (u, v) diffuse normal computed with our method is qualitatively slightly softer than that obtained using cross-polarized spherical gradient illumination. In contrast, the specular normal obtained with our method exhibits slightly sharper surface details than that obtained using polarized spherical gradients. This is because any surface details missing in the u, v diffuse normal (due to subsurface scattering) is attributed instead to the specular normal due to Equation 5.

Fig. 12 presents rendering comparisons under a point light source for a face against a photograph (b) for acquisition with the binary spherical gradients (a). As can be seen, the hybrid normal rendering with the estimated diffuse and specular albedo and photometric normals closely approximate the facial appearance in the photograph, with the exception of self-occlusions around the nose and the subject's neck which are not modeled in the renderings.

C. Kampouris, S. Zafeiriou, A. Ghosh / Diffuse-Specular Separation using Binary Spherical Gradient Illumination



Figure 11: Comparison of acquired diffuse and specular albedo and photometric normals using binary spherical gradient illumination (top-row) vs. polarized spherical gradient illumination [*GFT*^{*}11] (bottom-row). (a - d) Orange rubber ball. (e - f) Face.



(a) rendering

(b) photograph

Figure 12: Hybrid normal rendering comparisons to photograph: (a) Rendering with data acquired with binary spherical gradients. (b) Photograph of subject lit with a point light source.

Fig. 13 presents some additional results of diffuse-specular separation of albedo and normals using the six binary spherical gradient illumination conditions and their complements (white illumination). Fig. 13 (top-row) presents results for the subject shown earlier in Fig. 1 acquired from a side viewpoint demonstrating similar high quality of diffuse-specular separation for multiple viewpoints. Fig. 13 also presents results for a subject with a darker skin-type simultaneously acquired from two different viewpoints (secondand third-rows), and a few additional non-face examples of convex dielectric objects such as a deep orange rubber ball (fourthrow), and two real fruits - an orange (fifth row), and a tangerine (sixth-row). The results clearly demonstrate that our method successfully achieves view-independent diffuse-specular separation of albedo and photometric normals for a wide range of dielectric materials and human skin types and is well suited for multiview acquisition.

7.1. Limitations

Our method assumes no subject motion during acquisition. However, given the complementary nature of the binary gradient patterns, the technique of Wilson et al. [WGP*10] could be employed for any motion compensation. While the binary spherical gradient illumination is able to separate diffuse and specular reflectance, the quality of albedo separation is not as high quality as that achieved with polarization. This is because of the weaker specular signal with unpolarized imaging of the binary gradients compared to polarized spherical gradient illumination. The achieved quality of albedo separation is still sufficient for estimating high quality separated diffuse and specular photometric normals which are very useful for realistic rendering. The spectrally multiplexed binary RGB gradients enable albedo separation but not true separation of the photometric normals for which we need to rely on our heuristic approach. The quality of albedo separation with RGB illumination suffers a bit compared to when employing white illumination due to fewer independent observations while solving for the unknowns. And spectrally matching the color of diffuse albedo estimated with RGB illumination requires additional color calibration steps.

7.2. Specular normals and binary illumination

It is worth pointing out that besides albedo separation, estimating a high quality specular normal using binary spherical gradient illumination is a key result of this work. However, it is also somewhat surprising given the binary nature of the incident illumination. Unlike previous work which has employed an explicit linear spherical gradient in the incident illumination [MHP*07,GFT*11], we exploit an implicit (non-linear) gradient due to Fresnel to estimate the high frequency component of the specular normal field. This observation that Fresnel can be exploited for this purpose even under constant (or binary) illumination is a novel contribution of this work. Besides being useful for realistic rendering of high frequency specular detail, the estimated specular normal can also be employed for embossing mesoscopic surface details on a base geometry [NRDR05, MHP*07].

8. Conclusion

We have presented a novel diffuse-specular separation technique using binary spherical gradient illumination that is well suited for reflectance separation of dielectric objects including human skin. C. Kampouris, S. Zafeiriou, A. Ghosh / Diffuse-Specular Separation using Binary Spherical Gradient Illumination

It requires fewer photographs and a simpler setup for multiview acquisition than the state-of-the-art method of polarized spherical gradient illumination. Besides albedo separation, the approach also enables robust estimate of separated diffuse and specular photometric normals for realistic rendering applications. We also present novel spectral multiplexing of the binary spherical gradients which yields an efficient two-shot acquisition with reflectance separation. Given the simple two-shot acquisition, we believe the binary RGB gradients can be applied for practical video-rate dynamic performance capture.

9. Acknowledgements

This work was supported by an EPSRC Early Career Fellowship (EP/N006259/1) and a Royal Society Wolfson Research Merit Award to Abhijeet Ghosh.

References

- [AWL13] AITTALA M., WEYRICH T., LEHTINEN J.: Practical SVBRDF capture in the frequency domain. *ACM Trans. on Graphics* (*Proc. SIGGRAPH*) 32, 4 (2013). 3
- [DHT*00] DEBEVEC P., HAWKINS T., TCHOU C., DUIKER H.-P., SAROKIN W., SAGAR M.: Acquiring the reflectance field of a human face. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques* (2000), ACM Press/Addison-Wesley Publishing Co., pp. 145–156. 2
- [FCMB09] FRANCKEN Y., CUYPERS T., MERTENS T., BEKAERT P.: Gloss and normal map acquisition of mesostructures using gray codes. In Proceedings of the 5th International Symposium on Advances in Visual Computing: Part II (Berlin, Heidelberg, 2009), ISVC '09, Springer-Verlag, pp. 788–798. 3
- [FD15] FYFFE G., DEBEVEC P.: Single-shot reflectance measurement from polarized color gradient illumination. In *International Conference* on Computational Photography (2015). 1, 3
- [FGT*16] FYFFE G., GRAHAM P., TUNWATTANAPONG B., GHOSH A., DEBEVEC P.: Near-instant capture of high-resolution facial geometry and reflectance. *Computer Graphics Forum* 35, 2 (2016), 353–363. URL: http://dx.doi.org/10.1111/cgf.12837, doi:10.1111/ cgf.12837.3
- [FHW*11] FYFFE G., HAWKINS T., WATTS C., MA W.-C., DEBEVEC P.: Comprehensive facial performance capture. *Computer Graphics Forum (Proc. Eurographics)* 30, 2 (2011). 2
- [GCHS05] GOLDMAN D. B., CURLESS B., HERTZMANN A., SEITZ S. M.: Shape and spatially-varying brdfs from photometric stereo. In *ICCV* (2005), pp. 341–348. 2
- [GCP*09] GHOSH A., CHEN T., PEERS P., WILSON C. A., DEBEVEC P. E.: Estimating specular roughness and anisotropy from second order spherical gradient illumination. *Comput. Graph. Forum* 28, 4 (2009), 1161–1170. 3
- [GFT*11] GHOSH A., FYFFE G., TUNWATTANAPONG B., BUSCH J., YU X., DEBEVEC P.: Multiview face capture using polarized spherical gradient illumination. *ACM Transactions on Graphics (TOG) 30*, 6 (2011), 129. 1, 2, 7, 8
- [GHP*08] GHOSH A., HAWKINS T., PEERS P., FREDERIKSEN S., DE-BEVEC P.: Practical modeling and acquisition of layered facial reflectance. ACM Trans. Graph. 27, 5 (Dec. 2008), 139:1–139:10. 3
- [HVB*07] HERNANDEZ C., VOGIATZIS G., BROSTOW G. J., STENGER B., CIPOLLA R.: Non-rigid photometric stereo with colored lights. In Proc. IEEE International Conference on Computer Vision (2007), pp. 1–8. 2

© 2018 The Author(s)

Eurographics Proceedings © 2018 The Eurographics Association.

- [KHE10] KLAUDINY M., HILTON A., EDGE J.: High-detail 3D capture of facial performance. In *International Symposium 3D Data Processing*, *Visualization and Transmission (3DPVT)* (2010). 2
- [KIG16] KIM J., IZADI S., GHOSH A.: Single-shot layered reflectance separation using a polarized light field camera. In *Proceedings of the Eurographics Symposium on Rendering: Experimental Ideas & Implementations* (2016), EGSR '16. 3
- [LPGD09] LAMOND B., PEERS P., GHOSH A., DEBEVEC P.: Imagebased separation of diffuse and specular reflections using environmental structured illumination. In *Proc. of International Conference on Computational Photography* (2009). 3
- [LS01] LIN S., SHUM H.-Y.: Separation of diffuse and specular reflection in color images. In Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2001). 2
- [MHP*07] MA W.-C., HAWKINS T., PEERS P., CHABERT C.-F., WEISS M., DEBEVEC P.: Rapid acquisition of specular and diffuse normal maps from polarized spherical gradient illumination. In *Proceedings* of the 18th Eurographics Conference on Rendering Techniques (2007), EGSR'07, Eurographics Association, pp. 183–194. 1, 2, 4, 8
- [MZKB05] MALLICK S. P., ZICKLER T. E., KRIEGMAN D. J., BEL-HUMEUR P. N.: Beyond lambert: Reconstructing specular surfaces using color. In CVPR (2005). 2, 5
- [NRDR05] NEHAB D., RUSINKIEWICZ S., DAVIS J., RAMAMOORTHI R.: Efficiently combining positions and normals for precise 3d geometry. *ACM Trans. Graph.* 24, 3 (July 2005), 536–543. 8
- [TFG*13] TUNWATTANAPONG B., FYFFE G., GRAHAM P., BUSCH J., YU X., GHOSH A., DEBEVEC P.: Acquiring reflectance and shape from continuous spherical harmonic illumination. ACM Trans. Graph. (Proc. SIGGRAPH) 32, 4 (2013). 3
- [VPB*09] VLASIC D., PEERS P., BARAN I., DEBEVEC P., POPOVIĆ J., RUSINKIEWICZ S., MATUSIK W.: Dynamic shape capture using multiview photometric stereo. ACM Transactions on Graphics 28, 5 (2009), 174. 3
- [WGP*10] WILSON C. A., GHOSH A., PEERS P., CHIANG J.-Y., BUSCH J., DEBEVEC P.: Temporal upsampling of performance geometry using photometric alignment. ACM Transactions on Graphics (TOG) 29, 2 (2010), 17. 2, 4, 8
- [WGT*05] WENGER A., GARDNER A., TCHOU C., UNGER J., HAWKINS T., DEBEVEC P.: Performance relighting and reflectance transformation with time-multiplexed illumination. 756–764. 2
- [WMP*06] WEYRICH T., MATUSIK W., PFISTER H., BICKEL B., DONNER C., TU C., MCANDLESS J., LEE J., NGAN A., JENSEN H. W., GROSS M.: Analysis of human faces using a measurementbased skin reflectance model. ACM Trans. Graph. 25, 3 (July 2006), 1013–1024. 2
- [Woo78] WOODHAM R. J.: Photometric stereo: A reflectance map technique for determining surface orientation from image intensity. In Proc. SPIE's 22nd Annual Technical Symposium (1978), vol. 155. 2

C. Kampouris, S. Zafeiriou, A. Ghosh / Diffuse-Specular Separation using Binary Spherical Gradient Illumination



Figure 13: Diffuse and specular albedo (b, c) and photometric normal maps (d, e) of various subjects acquired using binary spherical gradient illumination and rendering with accquired data (f).