Dichromatic Separation: Specularity Removal and Editing

Satya P. Mallick
UC, San Diego

Todd Zickler
Harvard University

Peter Belhumeur
Columbia University

David Kriegman
UC, San Diego

The reflectance of a wide variety of materials (plastics, plant leaves, glazed ceramics, human skin, fruits and vegetables, paper, leather, etc.) can be described as a linear combination of specular and diffuse components, and for many applications we can benefit from separating an image in this way. Figures 2 and 3, for example, depict photo-editing and e-cosmetic applications in which visual effects are simulated by independently processing separated diffuse and specular image layers. Similarly, specular/diffuse separation plays an important role for image-based modeling applications in which diffuse (specular-free) texture maps are sought.

Separation of the two reflectance components in a single image is an ill-posed problem. In the past its solution has required the manual identification of highlight regions, the use of special acquisition systems (e.g., polarizing filters), or restrictive assumptions about the scene (e.g., untextured surfaces). Recently, we have introduced a method for specular/diffuse separation that overcomes many of these limitations [Mallick et al. 2006], and in this sketch, we build on this work, showing how it can be used for *dichromatic editing* — processing and recombining the two reflectance components for various visual effects. We present results on high-quality images and videos acquired in the laboratory in addition to images taken from the Internet. Results on the latter demonstrate robustness to low dynamic range, JPEG artifacts, and lack of knowledge of illuminant color.

Similar to most existing techniques for specular/diffuse separation, our approach is based on exploiting color differences between specular and diffuse reflections as described by Shafer's dichromatic model. According to this model, the color of the specular component at each surface point is the same as that of the illuminant (S), while the color of the diffuse component depends on the reflectance of the surface and can change from point to point.

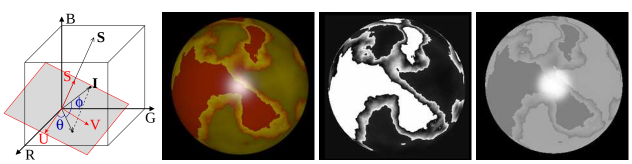

a b c d

Figure 1: (a) Partial separation obtained by rotation of color space and parameterization using angles θ and ϕ. (b) RGB image. (c) θ image: free of specularity. (d) ϕ image: contains an unknown mixture of diffuse and specular contributions.

We begin with a partial specular/diffuse separation obtained by a rotation and reparameterization of RGB color space. As shown in Fig. 1 (a), the rotation aligns one of the color axes (red, say) with S, the known source color. Given the RGB observation I at a pixel, we define the angular quantities θ and ϕ in this rotated space. There are two main advantages of this (θ,ϕ) representation of color. First, as shown in Fig 1 (c), θ is independent of both specular reflection and diffuse shading effects and therefore provides uncorrupted surface color (i.e., spectral reflectance) information. (We refer to this angle as "generalized hue" since it reduces to the standard definition of hue when the illuminant is white.) Second, the angle ϕ shown in Fig. 1 (d), is also independent of diffuse shading effects, and isolates the dimension of color space in which diffuse and specular mixing occurs.

Given this representation, the problem of specular/diffuse separation reduces to one of separating the specular and diffuse contributions to the ϕ-image. To accomplish this separation task, we share color information across the image through a series of local interactions. Intuitively, a reasonable estimate of the unknown diffuse contribution to ϕ at a point in a specular region (the center of the sphere in Fig. 1 (b), for example) is obtained from a nearby specular-free point with the same generalized hue. More formally, the diffuse contribution to each point of the ϕ-image is obtained by modeling the ϕ-image as a discrete approximation of a continuous 2D signal and evolving a non-linear PDE

$$\phi_t = -\left(\nabla\phi^\top \mathrm{M}(\theta)\nabla\phi\right)^{1/2},$$

that computes the multi-scale erosion (e.g., [Brockett and Maragos 1994]) of the 2D signal. We show that under reasonable assumptions, when using the input ϕ-image as an initial condition, the solution to this PDE yields the diffuse contribution to ϕ, and therefore uniquely determines the RGB diffuse component at each image point. Here, $\nabla\phi$ represents the spatial gradient, and the matrix $\mathrm{M}(\theta)$ determines the shape of the structuring set for the erosion process, enabling the sharing of diffuse color information without smoothing texture boundaries. In untextured regions, for example, diffuse color information can be shared in all directions and the required structuring set is disk-shaped. In textured regions, however, the required structuring set is linear and aligned with the iso-contours of θ. Finally, the same equation extends naturally to the case of videos, where erosion is performed along iso-surfaces of θ in three-dimensional space-time.

Figure 2: Dichromatic editing using a single input image. Left to right: input, specular-free, sharpened, re-lit, and 'wrinkled'.

a b c d e f

Figure 3: Dichromatic editing for e-cosmetics. Highlights on the skin of the input image (a) are removed while those on the lips are enhanced to obtain image shown in (b). Two regions of interest in (a) are shown in (c) and (e), and the corresponding regions in (b) are shown in (d) and (f).

Figures 2 and 3 show representative results. In each case, the specular/diffuse separation is computed as described above (without manual intervention), and these two layers are edited and recombined to produce a variety of visual effects. Additional results on images and videos are provided in the accompanying movie.

References

BROCKETT, R., AND MARAGOS, P. 1994. Evolution equations for continuous scale morphology. *IEEE transactions on Signal Processing 42*, 3377–3386.

MALLICK, S., ZICKLER, T., KRIEGMAN, D., AND BELHUMEUR, P. 2006. Specularity removal in images and videos: A PDE approach. In *Proc. European Conf. Computer Vision*, vol. 1, 550–563.