

Image-based Rendering from a Sparse Set of Images

Todd Zickler*
Harvard University

Sebastian Enrique
Columbia University

Ravi Ramamoorthi
Columbia University

Peter Belhumeur†
Columbia University

When the shape of an object is known, its appearance is determined by the spatially-varying bi-directional reflectance distribution function (SBRDF) defined on its surface. We present a method for recovering the SBRDF of a surface with known geometry from a sparse set of images. Unlike existing parametric methods (e.g., [Sato et al. 1997]) this approach does not require a choice of an analytic BRDF model, and is therefore capable of capturing arbitrary, complex reflectance effects. It is also different from existing non-parametric methods (e.g. [Matusik et al. 2002]) that are purely data-driven and can require thousands of images; unlike these methods, we exploit spatial coherence, which can drastically reduce the number of required input images.

Approach

For a curved surface with smooth spatial variation, each image provides a dense 2D slice of data in the higher-dimensional (mixed spatial and angular) SBRDF domain, and we frame SBRDF estimation as a scattered-data interpolation problem in this domain. Unlike existing methods, we interpolate simultaneously in the spatial and angular dimensions, enabling a controlled exchange between spatial and angular information, effectively giving up some of the spatial resolution in order to fill the holes between sparsely observed view and illumination conditions.

The scattered-data interpolation problem is solved by introducing: i) a new parameterization of the isotropic BRDF domain, and ii) a non-parametric representation of reflectance based on radial basis functions (RBFs). Our isotropic BRDF parameterization is designed to exploit the fact that a typical BRDF is quite compressible, meaning that it varies slowly over much of its angular domain. At the same time, it is designed to be suitable for general interpolation methods in \mathbb{R}^n .

Radial basis functions are used because they enable interpolation with few restrictions on the locations of the sample points. In particular, the use of radial basis functions means that, unlike existing BRDF representations based on spherical harmonics, Zernike polynomials, wavelets, and the basis of Matusik et al. [2003], we do not require a local preprocessing step to resample the image-based reflectance data at regular intervals. Radial basis functions are also useful because the cost of computing an RBF interpolant is dimension independent; and for a fixed number of samples, the size of the resulting representation does grow substantially as the dimension increases. As a result, essentially the same representation and method can be used to recover reflectance in three different cases, each of which requires interpolation in a different dimension:

1. homogeneous, isotropic BRDF data (\mathbb{R}^3)
2. spatially-varying reflectance from a single input image (\mathbb{R}^4)
3. spatially-varying reflectance from multiple images (\mathbb{R}^5)

Evaluation and Results

The method is applied to each of the three cases above, and it is evaluated both quantitatively and qualitatively by comparing it to existing parametric and non-parametric methods using synthetic data. In the case of homogeneous, isotropic BRDF data, the RBF representation is found to accurately represent both specular and retro-reflective reflectance, including Fresnel effects, and it provides improved accuracy when compared to a variety of parametric

models fit to the same data. In the case of smooth spatially-varying reflectance using multiple images, it is compared to angular-only interpolation techniques (like that of Matusik et al. [2002]), and in some cases it is found to reduce the number of required input images by more than an order of magnitude.

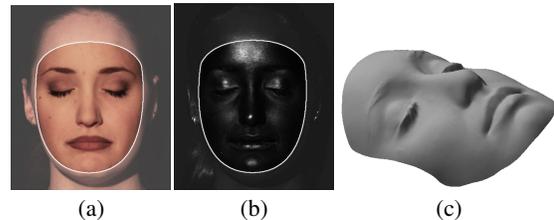


Figure 1: (a,b) Specular and diffuse components of a single input image. (c) Geometry used for SBRDF recovery and rendering.

In addition to this numerical evaluation, the method is applied to the problem of image-based modeling of a human face from one or more images. In the single-image case, the input is the surface geometry as well as one image captured under known viewpoint and known, point-source illumination (Fig. 1). The image is decomposed into its diffuse and specular components using polarizing filters. The diffuse component of the SBRDF is modeled as Lambertian and the RGB diffuse albedo at each surface point is estimated from the diffuse reflectance samples available in Fig. 1(a). The specular component is represented using RBFs, which effectively interpolate the specular SBRDF samples available in Fig. 1(b). The recovered SBRDF can be used to predict the appearance of the surface under novel view and illumination conditions. For example, Fig. 2 shows the predicted appearance under novel complex illumination.



Figure 2: Images synthesized using the SBRDF representation recovered from the single (decomposed) image shown in Fig. 1.

References

- MATUSIK, W., PFISTER, H., BRAND, M., AND MCMILLAN, L. 2002. Image-based 3D photography using opacity hulls. *ACM Transactions on Graphics (Proc. ACM SIGGRAPH)* 21, 3, 427–437.
- MATUSIK, W., PFISTER, H., BRAND, M., AND MCMILLAN, L. 2003. A data-driven reflectance model. *ACM Transactions on Graphics (Proc. ACM SIGGRAPH)* 22, 3, 759–769.
- SATO, Y., WHEELER, M. D., AND IKEUCHI, K. 1997. Object shape and reflectance modeling from observation. In *Proceedings of ACM SIGGRAPH*, 379–387.

*zickler@eecs.harvard.edu

†{senrique, ravir, belhumeur}@cs.columbia.edu