

Supplemental Material

All-Frequency Rendering with Dynamic, Spatially-Varying Reflectance




Figure 1: Comparison of SG, SH, and ZH bases for fitting the clamped cosine function. “Number of terms” refers to the order for the SH basis, and the number of lobes for the ZH and SG bases. (a) Error curves. (b) Fitted results.

1 Fitting the Clamped Cosine

Figure 1 compares our mixture of spherical Gaussians (SG) basis to spherical harmonics (SH) and zonal harmonics (ZH) [Sloan et al. 2005]. A simple 2-lobe fit with our basis provides good accuracy; error decreases further with more lobes. Asymptotic error decrease in our model compares favorably with increasing order of spherical harmonics and with increasing order and number of lobes in the ZH basis.

2 Fitting Isotropic Parametric Models

We tested our scheme on two parametric models for specular reflectance: Cook-Torrance (Gaussian lobe) and Blinn-Phong (cosine lobe). We also discuss the Ward model but because it is Gaussian-like and similar to Cook-Torrance, we show only the Cook-Torrance result. We separate each model into two factors: the normal distribution function (NDF) $D(\mathbf{h})$, and the rest of the factors $M_o(\mathbf{i})$. The NDF is fitted by SGs while M_o has an analytical form deduced in the following. Note that $\cos \theta = \mathbf{h} \cdot \mathbf{n}$. Parameters α, m, n are roughness constants for different models respectively. Also note that in canonical coordinates, \mathbf{n} is just the positive \mathbf{z} axis.

Cook-Torrance Model [Cook and Torrance 1981] has the form:

$$\rho_s(\mathbf{i}, \mathbf{o}) = \frac{F(\mathbf{i}, \mathbf{o})S(\mathbf{i}, \mathbf{o})}{\pi(\mathbf{n} \cdot \mathbf{i})(\mathbf{n} \cdot \mathbf{o})} \cdot e^{-(\theta/m)^2} \quad (1)$$

$$M_o(\mathbf{i}) = \frac{F(\mathbf{i}, \mathbf{o})S(\mathbf{i}, \mathbf{o})}{\pi(\mathbf{n} \cdot \mathbf{i})(\mathbf{n} \cdot \mathbf{o})} \quad (2)$$

$$D(\mathbf{h}) = e^{-\arccos^2(\mathbf{h} \cdot \mathbf{n})/m^2} \approx e^{-2(1-(\mathbf{h} \cdot \mathbf{n}))/m^2} \quad (3)$$




Figure 2: Single-lobe SG fit to isotropic parametric models. The solid curve represents the results of nonlinear optimization, while the dashed curve uses a simple analytic model. (a) Fitted λ related to Cook-Torrance’s roughness parameter m . (b) Fitted λ related to Blinn-Phong’s shininess parameter n .

Ward Model [Ward 1992] has the form:

$$\rho_s(\mathbf{i}, \mathbf{o}) = \frac{1}{\sqrt{(\mathbf{i} \cdot \mathbf{n})(\mathbf{o} \cdot \mathbf{n})}} \frac{e^{-\tan^2 \theta / \alpha^2}}{4\pi\alpha^2} \quad (4)$$

$$M_o(\mathbf{i}) = \frac{1}{4\pi\alpha^2 \sqrt{(\mathbf{i} \cdot \mathbf{n})(\mathbf{o} \cdot \mathbf{n})}} \quad (5)$$

$$D(\mathbf{h}) = e^{-\frac{1-(\mathbf{h} \cdot \mathbf{n})^2}{\alpha^2(\mathbf{h} \cdot \mathbf{n})^2}} \approx e^{-2(1-(\mathbf{h} \cdot \mathbf{n}))/\alpha^2} \quad (6)$$

Blinn-Phong Model [Blinn 1977] has the form:

$$\rho_s(\mathbf{i}, \mathbf{o}) = \frac{n+2}{2\pi} \cos^n \theta \quad (7)$$

$$M_o(\mathbf{i}) = \frac{n+2}{2\pi} \quad (8)$$

$$D(\mathbf{h}) = (\mathbf{h} \cdot \mathbf{n})^n \approx e^{-n(1-(\mathbf{h} \cdot \mathbf{n}))} \quad (9)$$

Figure 2 shows results from fitting a single-lobe SG to the NDF of the Cook-Torrance and Blinn-Phong models. Since the maximum value is 1 in both cases, we assign $\mu = 1$. To fit the lobe sharpness parameter, λ , we find that an analytic formula based on a simple two-term (linear) Taylor expansion of the exponentiated function matches the fitted curve very well, as illustrated in the figure. The formula is given by $\lambda = 2/m^2$ for the Cook-Torrance model, $\lambda = 2/\alpha^2$ for the Ward model, and $\lambda = n$ for the Blinn-Phong model.

Comparison with BRDF Factorization

The previously best-known method for representing BRDFs in PRT is to factorize the BRDF function into a light-dependent component and a view-dependent component [Wang et al. 2004; Liu et al. 2004; Wang et al. 2006], via:

$$\rho_s(\mathbf{i}, \mathbf{o}) = \sum_{k=1}^n \phi_k(\mathbf{i}) \psi_k(\mathbf{o}) \quad (10)$$

We compare reconstruction error from our SG mixture fitted result to that from BRDF factorization, using the following weighted

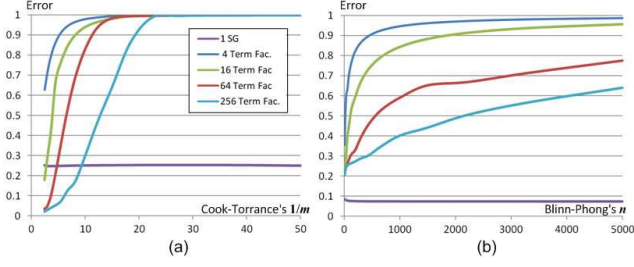


Figure 3: Error comparison of single-lobe SG fitted model with BRDF factorization. (a) Error related to Cook-Torrance's roughness parameter m . (b) Error related to Blinn-Phong's shininess parameter n .

sum of squares measure:

$$\frac{\sum_{\mathbf{i}, \mathbf{o}} (\rho_s(\mathbf{i}, \mathbf{o}) - \rho_s^R(\mathbf{i}, \mathbf{o}))^2 ((\mathbf{i} \cdot \mathbf{n})(\mathbf{o} \cdot \mathbf{n}))^2}{\sum_{\mathbf{i}, \mathbf{o}} (\rho_s(\mathbf{i}, \mathbf{o}) (\mathbf{i} \cdot \mathbf{n})(\mathbf{o} \cdot \mathbf{n}))^2}. \quad (11)$$

As shown in Figure 3, error from BRDF factorization increases with increasing specularity, while the error of our method is basically constant for a single-lobe model, regardless of its specularity.

To compare visual quality, we selected three Cook-Torrance BRDFs with different specular sharpness as shown in Figure 4. Even with 256 terms, a very heavyweight representation that can not be rendered interactively, BRDF factorization provides a poor approximation. On the other hand, our single-lobe SG model matches the ground truth very well.

3 Fitting Anisotropic Parametric Models

Anisotropic models require multiple-lobe SGs. All parameters of the SG mixture, including its lobe axes, are determined on the fly given the BRDF model parameters. The Ashikhmin-Shirley [Ashikhmin and Shirley 2000] anisotropic model, for example, has the form:

$$\rho_s(\mathbf{i}, \mathbf{o}) = \frac{\sqrt{(n_u + 1)(n_v + 1)} (\mathbf{n} \cdot \mathbf{h})^{n_u \cos^2 \phi_h + n_v \sin^2 \phi_h}}{8\pi (\mathbf{h} \cdot \mathbf{i}) \max(\mathbf{i} \cdot \mathbf{n}, \mathbf{o} \cdot \mathbf{n})} F(\mathbf{h} \cdot \mathbf{i})$$

$$M_{\mathbf{o}}(\mathbf{i}) = \frac{\sqrt{(n_u + 1)(n_v + 1)} F(\mathbf{h} \cdot \mathbf{i})}{8\pi (\mathbf{h} \cdot \mathbf{i}) \max(\mathbf{i} \cdot \mathbf{n}, \mathbf{o} \cdot \mathbf{n})} \quad (12)$$

$$D(\mathbf{h}) = (\mathbf{n} \cdot \mathbf{h})^{n_u \cos^2 \phi_h + n_v \sin^2 \phi_h} \quad (13)$$

where ϕ_h is the azimuthal angle formed by the half-angle vector \mathbf{h} and the Fresnel factor F is defined in equation (12) in the paper.

As in the case of isotropic models, we keep the analytic form of $M_{\mathbf{o}}$ and fit the NDF D using an SG mixture. Without loss of generality, we assume $n_u > n_v$; the opposite case follows by rotating the model $\pi/2$ azimuthally. Given n_u , n_v , and the number of SG lobes to be fit N (odd), we distribute the SGs along the longer length corresponding to the v tangent direction. We use a constant lobe sharpness $\lambda_i = n_u$ for all lobes.

As shown in Figure 5b, lobe axes \mathbf{p}_i are evenly distributed along the v direction using an angular interval θ_i defined by:

$$\theta_v = \arccos\left(\varepsilon^{\frac{1}{N+1}}\right), \quad \theta_i = \frac{2\theta_v}{N+1}. \quad (14)$$

The parameter θ_v represents the angular distance from the overall specular lobe center where D is 1 to where $D = \varepsilon$; the total angular extent of the lobe's longer axis is thus $2\theta_v$. We use $\varepsilon = 0.1$. Finally the lobe amplitudes μ_i are determined by sampling the NDF at the lobe center; i.e., $\mu_i = D(\mathbf{p}_i)$. Figure 5c compares error between



Figure 4: Rendering comparison with parametric, isotropic BRDFs: (a) ground truth, (b) single-lobe SG model, and BRDF factorization with (c) 256 terms, (d) 64 term, and (e) 16 terms.




Figure 5: Fitting anisotropic BRDF models with SG mixtures. (a) The anisotropic NDF. (b) SG mixture fit. (c) Fitting error and comparison with BRDF factorization.




Figure 6: Rendering comparison with parametric, anisotropic BRDFs: (a) ground truth, (b) 7-lobe SG model, BRDF factorizations with (c) 256 terms, and (d) 64 terms.

SG mixtures and BRDF factorization when fit to this anisotropic BRDF model. The error is calculated as in equation (11). Visual quality is compared in Figure 6.

4 Fitting Measured BRDFs

We also tested our model on a database of real, measured BRDFs from [Matusik et al. 2003]. For each BRDF, we first fit a microfacet BRDF model as in [Ngan et al. 2005]. We then fit an SG mixture to the derived NDF, while keeping analytic expressions for the other factors (shadowing and Fresnel). Figure 8 compares reconstruction error. We achieve lower error in all cases compared to a BRDF factorization with 256 terms.

5 SG Warping

Our approximation warps reflectance lobes in order to extract the BRDF view slice. Error arises in this warp because we represent the resulting lobe using a spherical Gaussian and thus assume it remains isotropic. We quantify this error based on a single-lobe fit to the Cook-Torrance BRDF, using the sum of squares measure in equation (11). The plot in Figure 9 shows that this error drops quickly as lobe sharpness increases. The visual error is small regardless of lobe sharpness, as can be seen in the rendered images of a sphere in the figure, for various values of roughness m .

6 SSDF Product and Compression

To investigate error in approximating products between SGs and SSDFs, we randomly selected visibility functions and their corresponding SSDFs, V_i^d , from N vertices in the scene shown in Figure 10, as well as M lobe directions \mathbf{p}_j . Given lobe sharpness λ , we calculated the squared error (vs. brute force integration with accurate visibility) for inner and vector products over all vertices and




Figure 7: SSDF product error.

lobe directions. Approximation error drops rapidly as a function of lobe sharpness, as shown in Figure 7.

Figure 10 compares our method with the method in [Ng et al. 2004] based on Haar wavelet compression of visibility cubemaps. Moderately tessellated geometry (13.3k vertices) and densely tessellated geometry (1M vertices) particularly for wavelet compression are used. The wavelet method uses the same cubemap resolution, $256 \times 256 \times 6$, that we used to sample visibility in our method, but requires denser geometry to obtain smooth shading result. To be conservative, we use 400 wavelet terms to compress the BRDF for the wavelet method, more than can be handled in a real-time implementation.

Our result is more accurate than the wavelet method, based on an equal compressed size (equal number of compression terms per vertex). Furthermore, our method supports better per-pixel interpolation, providing greater fidelity in the detailed highlight and shadow boundaries.

References

ASHIKHMIN, M., AND SHIRLEY, P. 2000. An anisotropic Phong BRDF model. *Journal of Graphics Tools* 5, 2, 25–32.

BLINN, J. F. 1977. Models of light reflection for computer synthesized pictures. In *Computer Graphics (Proceedings of SIGGRAPH 77)*, ACM, vol. 11, 192–198.

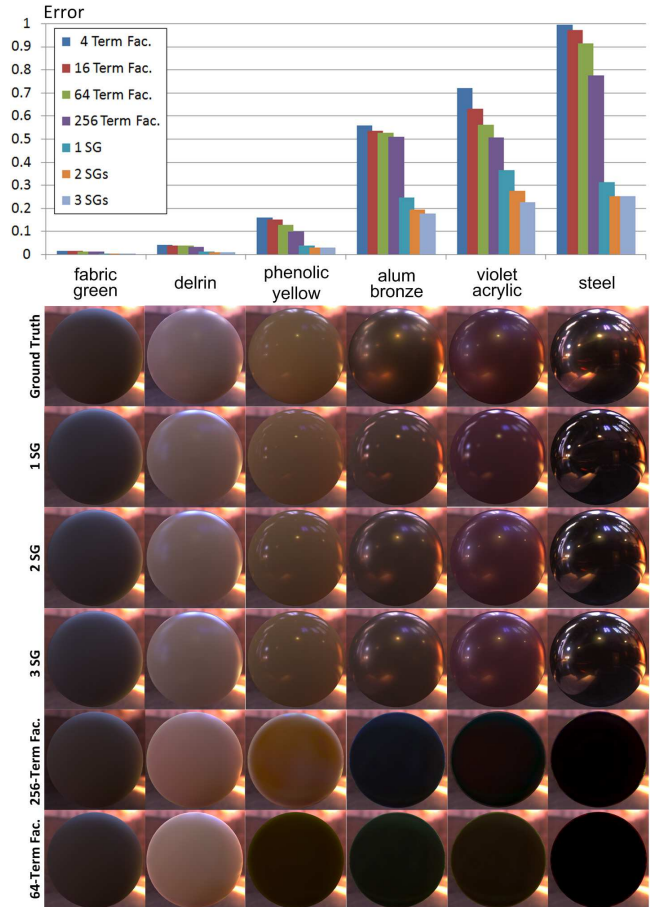


Figure 8: Fitting measured BRDFs with SG mixtures. Six BRDFs with different specularity were selected from the database in [Matusik et al. 2003].





Figure 9: SG warping error.

- COOK, R. L., AND TORRANCE, K. E. 1981. A reflectance model for computer graphics. In *Computer Graphics (Proceedings of SIGGRAPH 81)*, ACM, vol. 1, 307–316.
- LIU, X., SLOAN, P. P., SHUM, H. Y., AND SNYDER, J. 2004. All-frequency precomputed radiance transfer for glossy objects. In *Proceedings of the Eurographics Symposium on Rendering*, Eurographics Association, 337–344.
- MATUSIK, W., PFISTER, H., BRAND, M., AND MCMILLAN, L. 2003. A data-driven reflectance model. *ACM Transactions on Graphics* 22, 3, 759–769.
- NG, R., RAMAMOORTHY, R., AND HANRAHAN, P. 2004. Triple product wavelet integrals for all-frequency relighting. *ACM Transactions on Graphics* 23, 3, 477–487.
- NGAN, A., DURAND, F., AND MATUSIK, W. 2005. Experimental analysis of BRDF models. In *Rendering Techniques 2005: 16th Eurographics Workshop on Rendering*, 117–126.
- SLOAN, P.-P., LUNA, B., AND SNYDER, J. 2005. Local, deformable precomputed radiance transfer. *ACM Transactions on Graphics* 24, 3, 1216–1224.
- WANG, R., TRAN, J., AND LUEBKE, D. 2004. All-frequency relighting of non-diffuse objects using separable BRDF approximation. In *Rendering Techniques*, Eurographics Association, 345–354.
- WANG, R., TRAN, J., AND LUEBKE, D. 2006. All-frequency relighting of glossy objects. *ACM Transactions on Graphics* 25, 2, 293–318.
- WARD, G. J. 1992. Measuring and modeling anisotropic reflection.



Figure 10: PCA-compressed SSDFs vs. wavelets on visibility cubemaps.