

Efficient Reflectance and Visibility Approximations for Environment Map Rendering

Paul Green¹

Jan Kautz²

Frédo Durand¹

¹Massachusetts Institute of Technology, Computer Science and Artificial Intelligence Laboratory

²University College London

Abstract

We present a technique for approximating isotropic BRDFs and precomputed self-occlusion that enables accurate and efficient prefiltered environment map rendering. Our approach uses a nonlinear approximation of the BRDF as a weighted sum of isotropic Gaussian functions. Our representation requires a minimal amount of storage, can accurately represent BRDFs of arbitrary sharpness, and is above all, efficient to render. We precompute visibility due to self-occlusion and store a low-frequency approximation suitable for glossy reflections. We demonstrate our method by fitting our representation to measured BRDF data, yielding high visual quality at real-time frame rates.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation; Bitmap and frame buffer operations; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Color, Shading, Shadowing and Texture

1. Introduction

There have been many techniques developed for interactive and real-time rendering of static objects illuminated by distant environment lighting [NRH03, LSS04, WTL04]. These techniques separate the evaluation of the rendering equation into an expensive precomputation stage and an efficient on-line rendering stage. The trend has been to increase the richness of simulated effects at the cost of increased complexity in preprocessing and storage. In contrast, our work emphasize visual faithfulness rather than physical accuracy and, above all, it focuses on simplicity of implementation.

We address the above issues and introduce a technique that requires a small memory footprint (about 1KB per BRDF), and yet produces realistic high-frequency reflections. Our method represents an isotropic Bidirectional Reflectance Distribution Function (BRDF) using a sum of Gaussian functions of arbitrary bandwidths for a fixed set of view elevation angles. The explicit use of scale (in the form of variance) enables reflections of arbitrary frequencies. In contrast to analytical models of reflectance based on the half-angle formulation, our representation is well suited for efficient prefiltered environment mapping. Additionally, BRDF parameters can be changed at run-time, without reprocessing any data.

Furthermore, we make the observation that correct shadowing of glossy reflections is perceptually not crucial, leading to a simplified form of precomputed radiance transfer (PRT). Our simplified PRT ignores multiple-bounce specular reflection and is much cheaper than previous methods [NRH04, LSS04, GKMD06], but still visually effective. In our work we focus only on the specular component of a BRDF. The diffuse component can be efficiently rendered with precomputed radiance transfer [SKS02] or irradiance environment maps [RH01b]. Our contributions can be summarized as follows:

- We introduce an isotropic BRDF representation that accurately fits measured BRDF data and requires very little storage.
- Our representation is ideal for integration with prefiltered environment maps, enabling high-frequency reflections at real-time rates, evaluated entirely on the GPU.
- We introduce a light-weight visibility approximation for shadowed reflections, where only large-scale accuracy is important for visually realistic renderings.

2. Related Work

Blinn and Newell [BN76] first introduced *environment mapping* for mirror-like reflections on curved objects. Prefilter-

ing an environment map can produce diffuse [MH84] and Phong-like reflections [Gre86] on arbitrary geometries from an environment, without the need to recompute the integral between the BRDF and the environment map. This is especially beneficial in a real-time context [HS99], as prefiltering can be performed in an offline step, albeit restricting the lighting to be static. Commonly, only glossy reflections using the Phong model [Pho75] or Phong-like BRDF approximations [KM00, MLH02] are supported, which keeps the resulting prefiltered environment map two-dimensional. Our method differs from these techniques, as it allows for arbitrary isotropic BRDFs, but still keeps the memory requirements to a minimum.

Kautz and McCool [KM00] proposed using a single- or multi-lobe approximation of isotropic BRDFs for a fixed set of azimuthal view angles, where the shape of each rotationally-symmetric lobe is allowed to vary. This results in a 3D environment map, consisting of a stack of 2D prefiltered environment maps, one for each view angle in the case of a single-lobe approximation, or several in the case of multi-lobe approximations. This method requires considerable amounts of memory — a 3D environment map for each BRDF in a scene — but can handle the common phenomenon of increased highlight sharpness at grazing angles. By reparameterizing the BRDF, and using a nonlinear parametric approximation, our method is able to capture all of the same effects, while reducing the memory requirements to a handful of coefficients per BRDF and a single mip-mapped environment map (which is used for all BRDFs). An added benefit of our approach is the ability to change the BRDF at run-time, without the need to recompute filtered environment maps.

Cabral et al. [CON99] proposed an image-based technique that stores prefiltered environment maps for a fixed set of possible view-directions. They reparameterize the environment maps by the central reflection direction of a BRDF to enable better interpolation when generating in-between views; additional compression [LK02] reduces memory consumption. While these methods support arbitrary BRDFs, they are also limited to static incident lighting and fixed BRDFs due to the expensive prefiltering and compression steps. Our method makes use of the same reparameterization to achieve better approximation results.

Ngan et al. [NDM05] have demonstrated that analytic BRDF models based on the half-vector parameterization reproduce real BRDFs more faithfully than models based on the mirror direction. Unfortunately, we observe that the half-angle definition does not lend itself to prefiltered environment mapping because it results in kernels that are not rotationally symmetric. To avoid this asymmetry, we follow previous work [RH01b] and base our representation on the mirror angle. However, we use a sum of Gaussian functions to improve the expressiveness of the mirror-angle represen-

tation and to better reproduce the lobe anisotropy observed in real materials.

Diffuse reflections can be accurately and efficiently rendered using Spherical Harmonics (SH) [RH01b]. Spherical Harmonics have also been used to render reflections with arbitrary isotropic [RH02] and anisotropic [KSS02] BRDFs, but are limited to low-frequency reflections. Kautz et al. [KVHS00] demonstrated sharp reflections from dynamic lighting, if the material is restricted to Phong-like BRDFs. In contrast, our method is not limited to low-frequency reflections and is suitable for dynamic reflections with arbitrary high-frequency isotropic BRDFs.

None of the above techniques consider self-shadowing or inter-reflections. Precomputed radiance transfer incorporates reflections, inter-reflections, and self-shadowing from dynamic environment map lighting [SKS02, KSS02]. Although the original techniques were restricted to low-frequency lighting environments, this limitation has been overcome using nonlinear wavelet approximation and clustering methods [NRH03, LSSS04, WTL04]. In this paper, we make the important observation that accurate visibility for glossy reflections is visually less important than for diffuse surfaces, where shadows play an important role. Based on this observation, we use a low-frequency representation of visibility, which allows us to render visually plausible, high-frequency reflections at real-time rates.

Recently, Tsai and Shih [TS06] have proposed a technique that uses spherical radial basis functions to represent light transport. This method produces accurate renderings at near real-time rates. However it requires several orders of magnitude more storage (several MBs as compared to ~1KB for our method) and hours of precomputation time. In addition, their rendering frame rates are asymmetric with regard to changes in viewpoint (~50fps) and changes in lighting (~15fps). Our method does not suffer from this limitation.

Green et al. [GKMD06] approximate the transport function with a sum of Gaussians for a set of 92 fixed views. The transport function contains shadowing effects, making fitting complicated and expensive (takes several hours to precompute on a cluster of 20 PCs). We build on their work but significantly simplify both the representation and the preprocessing. We propose fitting a reparameterized version of the BRDF with a sum of Gaussian functions and to handle the visibility separately. This dramatically simplifies the fitting stage and the representation, and achieves surprisingly high visual faithfulness.

Standard ambient occlusion techniques [Bun04] weight shading by calculating the percentage of the visible hemisphere that is unoccluded. We use a low-order Spherical Harmonic representation of visibility, in which the DC term encodes the ambient occlusion, and the higher-order terms add view-dependence. Spherical Harmonics have recently been used to encode visibility in dynamic scenes [RWS*06].




Figure 1: Reparameterization of $\vec{\omega}_i$ w.r.t. to the reflection direction \mathbf{r} . Outgoing view V is parameterized by elevation angle θ_o w.r.t. to the surface normal \mathbf{n}_x , and incoming light direction L is parameterized by a spherical displacement δ_i w.r.t. the reflection direction \mathbf{r} .

This technique is limited to shading diffuse surfaces in low-frequency lighting to render soft shadows, whereas we focus on visibility for glossy reflections. It may be possible to extend the method of Ren *et al.* [RWS*06] to allow dynamic scenes.

3. Reflectance Approximation

In this section we introduce our approximation to the rendering equation [Kaj86] and discuss the assumptions and reparameterizations necessary. We first discuss unshadowed rendering, and then describe extensions for shadowed rendering.

3.1. Unshadowed Reflectance

We wish to calculate the outgoing radiance $L_o(\mathbf{x}, \vec{\omega}_o)$ along viewing direction $\vec{\omega}_o$ at a surface point \mathbf{x} . Unshadowed outgoing radiance $L_o(\mathbf{x}, \vec{\omega}_o)$ can be described by the integration of the distant incident radiance L_i attenuated by the geometric term $(\mathbf{n}_x \cdot \vec{\omega}_i)$ and the reflectance function $f_x(\vec{\omega}_o, \vec{\omega}_i)$:

$$L_o(\mathbf{x}, \vec{\omega}_o) = \int_{\Omega} f_x(\vec{\omega}_o, \vec{\omega}_i) (\mathbf{n}_x \cdot \vec{\omega}_i) L_i(\vec{\omega}_i) d\vec{\omega}_i. \quad (1)$$

We define a transfer function $\mathbf{T}_{\mathbf{x}, \vec{\omega}_o}(\vec{\omega}_i) = f_x(\vec{\omega}_o, \vec{\omega}_i) \cos \theta_i$ (where $\vec{\omega}_i = (\theta_i, \phi_i)$ is the incident lighting direction expressed in spherical coordinates, and similarly for $\vec{\omega}_o$). For a fixed outgoing direction $\vec{\omega}_o$, $\mathbf{T}_{\mathbf{x}, \vec{\omega}_o}$ is a cosine weighted 2D slice of the BRDF, parameterized by incoming direction $\vec{\omega}_i$. Note, we subscript $\mathbf{T}_{\mathbf{x}, \vec{\omega}_o}$ with $\vec{\omega}_o$ to emphasize that $\vec{\omega}_o$ is held constant.

As noted in [RH02] it is beneficial to reparameterize the transfer function to align prominent features. In particular, when reparameterized, isotropic BRDFs reduce in dimensionality from 4D to 3D. In addition to improving compression, the alignment of features is key to the high quality interpolation of the transfer function across elevation angles. We use a reparameterization that is similar in spirit to the reflection direction parameterization discussed in [RH02].

Let \mathbf{r} be the reflection direction given direction $\vec{\omega}_o$ and

surface normal \mathbf{n}_x . We can reparameterize $\vec{\omega}_i$ by a spherical displacement $\vec{\delta}_i$ w.r.t. the reflection direction \mathbf{r} (see Figure 1). Using the change of variables $\vec{\omega}_i \rightarrow \mathbf{r} + \vec{\delta}_i$ and assuming an isotropic BRDF — allowing us to drop the azimuthal dependence, ϕ_o , of $\vec{\omega}_o$ — we can rewrite (1) as:

$$L_o(\mathbf{x}, \vec{\omega}_o) = \int_{\Omega} \mathbf{T}_{\mathbf{x}, \theta_o}(\vec{\delta}_i) L_i(\mathbf{r} + \vec{\delta}_i) d\vec{\delta}_i. \quad (2)$$

Our goal is to find a compact representation of $\mathbf{T}_{\mathbf{x}, \theta_o}$ which enables an accurate and efficient integration of (2). To this end, we approximate $\mathbf{T}_{\mathbf{x}, \theta_o}$ as a sum of K isotropic Gaussian functions:

$$\mathbf{T}_{\mathbf{x}, \theta_o}(\vec{\delta}_i) \approx \tilde{\mathbf{T}}_{\mathbf{x}, \theta_o}(\vec{\delta}_i) = \sum_k w_{k, \theta_o} G_{\sigma_{k, \theta_o}}(\mu_{k, \theta_o} - \vec{\delta}_i) \quad (3)$$

where μ_{k, θ_o} , σ_{k, θ_o} and w_{k, θ_o} are the mean, standard deviation and weight respectively of the k^{th} Gaussian for elevation angle θ_o (for clarity and convenience we will omit any further subscripting of θ_o on the Gaussian parameters). Typically we use $K = 3$ Gaussian lobes. Substituting our approximation (3) into (2) yields:

$$\begin{aligned} L_o(\mathbf{x}, \vec{\omega}_o) &\approx \int_{\Omega} \tilde{\mathbf{T}}_{\mathbf{x}, \theta_o}(\vec{\delta}_i) L_i(\mathbf{r} + \vec{\delta}_i) d\vec{\delta}_i \\ &= \int_{\Omega} \sum_k w_k G_{\sigma_k}(\mu_k - \vec{\delta}_i) L_i(\mathbf{r} + \vec{\delta}_i) d\vec{\delta}_i \\ &= \sum_k w_k \int_{\Omega} G_{\sigma_k}(\mu_k - \vec{\delta}_i) L_i(\mathbf{r} + \vec{\delta}_i) d\vec{\delta}_i \quad (4) \\ &= \sum_k w_k (G_{\sigma_k} * L_i)(\mathbf{r} + \mu_k) \quad (5) \end{aligned}$$

where $G_{\sigma_k} * L_i$ denotes convolution of L_i with a Gaussian of standard deviation σ_k . By preconvolving and storing $G_{\sigma_k} * L_i$, evaluating (2) becomes trivial.

3.2. Arbitrary Elevations

The analysis described in the previous section assumed that for a particular outgoing direction $\vec{\omega}_o = (\theta_o, \phi_o)$ we have an approximated transfer function $\tilde{\mathbf{T}}_{\mathbf{x}, \theta_o}$. However, instead of storing $\tilde{\mathbf{T}}_{\mathbf{x}, \theta_o}$ at many finely sampled elevation angles θ_o , our strategy is to coarsely sample θ_o and interpolate the parameters of our model to in-between views [GKMD06]. In practice, we typically store 15 $\tilde{\mathbf{T}}_{\mathbf{x}, \theta_o}$ samples at equally spaced elevation angles from 0 to $\pi/2$ (see Figure 3 for sample data).

3.3. Adding Shadowing

In the previous section we derived an efficient method for evaluating the unshadowed reflectance integral. To extend this framework to handle shadowing we include a piecewise constant approximation to the visibility function. Using the same $(\theta_o, \vec{\delta}_i)$ reflection space parameterization, we can

rewrite (4), now including the visibility function $V_{\mathbf{x}}$:

$$L_o(\mathbf{x}, \vec{\omega}_o) \approx \sum_k^K w_k \int_{\Omega} G_{\sigma_k}(\mu_k - \vec{\delta}_i) V_{\mathbf{x}}(\mathbf{r} + \vec{\delta}_i) L_i(\mathbf{r} + \vec{\delta}_i) d\vec{\delta}_i. \quad (6)$$

Instead of evaluating the full visibility function, we make the approximation that visibility is constant over the entire support of each Gaussian function G_{σ_k} used in our reflectance approximation. We use the visibility value evaluated at the mean direction $V_{\mathbf{x}}(\mathbf{r} + \mu_k)$ as a piecewise constant approximation to the true visibility over the whole range of G_{σ_k} . This separability approximation of visibility allows us to factor $V_{\mathbf{x}}$ out of the integral, and once again outgoing radiance can be described by a convolution:

$$L_o(\mathbf{x}, \vec{\omega}_o) \approx \sum_k^K w_k V_{\mathbf{x}}(\mathbf{r} + \mu_k) (G_{\sigma_k} * L_i)(\mathbf{r} + \mu_k). \quad (7)$$

Representing $V_{\mathbf{x}}(\mathbf{r} + \mu_k)$ accurately is difficult, because visibility contains high-frequencies and varies for each \mathbf{x} . However it is beneficial if our approximation captures low-frequency, large-scale visibility variation which is important to avoid perceptually troublesome artifacts such as bright highlights in occluded regions (see figures 5 and 7). We therefore represent the visibility with low-order spherical harmonics (SH), similar to PRT [SKS02].

3.3.1. Discussion

If spherical harmonics based diffuse PRT [SKS02] is used in conjunction with our technique, the same PRT data can be used for the visibility term of the glossy component. In diffuse PRT, the visibility times the cosine — i.e., $V_{\mathbf{x}}(\omega) \cos \omega$ — is represented using SH. By dividing the cosine out, the same data can be reused for the visibility of the glossy component.

4. Model Fitting

In this section we describe our technique for estimating the parameters of the transfer function $\tilde{\mathbf{T}}_{\mathbf{x}, \theta_o}(\vec{\delta}_i)$. We want to estimate parameters μ_{k, θ_o} , σ_{k, θ_o} and w_{k, θ_o} for each Gaussian and θ_o that minimize the L^2 distance between the actual transfer function and our approximate transfer function:

$$\min_{\theta_o} \sum_i \left\| \tilde{\mathbf{T}}_{\mathbf{x}, \theta_o}(\vec{\delta}_i) - \sum_k^K w_{k, \theta_o} G_{\sigma_{k, \theta_o}}(\mu_{k, \theta_o} - \vec{\delta}_i) \right\|^2 \quad (8)$$

We use nonlinear optimization to find the best set of parameters [CL96]. Initially, we experimented with additional terms in the objective function to ensure consistency across θ_o (similar to Green et al. [GKMD06]). However, we found that they were unnecessary. In fact, not including smoothness terms makes the optimization faster, less fickle to initial conditions, and achieved more accurate fits. For $K=3$ Gaussians and $N=15$ views, it takes about 10 minutes to fit our




Figure 2: Our fit (bottom rows) to two measured brdf data (top rows). We use $K=3$ Gaussians. Elevation angle increases from $\theta_o = 0$ on the left to $\approx \pi/2$ on the right. Using more Gaussians can achieve better fits.

model to measured BRDF data [MPBM03] using MATLAB optimization routines. See Figure 2 for two example fits.

Visibility. We precompute visibility using ray-casting at every vertex. The visibility is then converted into 4-th order SH. This is identical to PRT precomputation.

5. Rendering

Rendering glossy shadowed reflections with our technique can be done entirely on the GPU, achieving real-time frame rates. Rendering entails evaluating Equation 7 on the GPU.

For now, let us assume we have an appropriately preconvolved lighting environment ($G_{\sigma_j} * L_i$) stored in a cube map, where mip-map level j is filtered with a Gaussian of size σ_j . We arrange the means μ_k and standard deviations σ_k in one texture map, and the RGB-weights w_k in another texture map, as can be seen in Figure 3. In case visibility is used, we have additional texture maps (or per-vertex data) storing the SH coefficients.

In a pixel shader we compute the elevation angle θ_o of the view direction in the local tangent frame, as well as the reflection direction \mathbf{r} . Based on θ_o , we query the parameters μ_k , σ_k , and w_k for all K Gaussians (3 or 4 in practice) from the two texture maps. We then perform a lookup at the appropriate level of the filtered cube map, based on σ_k and the direction $\mathbf{r} + \mu_k$. This value is then scaled by the weights w_k and the visibility term, which is queried by directly evaluating the 4-th order SH along direction $\mathbf{r} + \mu_k$.



Figure 3: The complete BRDF data for the green-metallic-paint measured BRDF, using a K -term Gaussian approximation and an elevation sampling of N views. The data (means, standard deviations, weights) is stored as two texture maps of size $K \times N$ texels. Typically we use $K=3$ Gaussians and $N=15$ views.

The basic shader using 3 Gaussians is only 54 ARB instructions long. When visibility is used we evaluate our SH representation of visibility in the shader. The length of the shader for second, third, and fourth-order SH grows to 107, 140, and 181 ARB instructions respectively.

5.1. Spherical Convolution

As stated above, our technique uses preconvolved environment maps in order to efficiently integrate the BRDF with the lighting. As has been done before [MLH02, AG02], we build a set of progressively convolved lighting environments, which are stored in a single mip-mapped cube map. Each level of the mip-map is computed by convolving the environment map with a spherical Gaussian that has twice the standard deviation of the previous level, i.e., $\sigma_{j+1} = 2\sigma_j$; appropriate sub-sampling is applied in order to store the result in a cube map.

Accurate convolution of the lighting environment is usually considered too costly, as it is done in the inconvenient spherical domain. However, accurate filtering can be done quickly in the spherical Fourier domain (Spherical Harmonics), where convolution becomes a product. We transform the lighting into SH by applying a variant of the Fast Fourier Transform [HKRM03] (code available from <http://www.cs.dartmouth.edu/~geelong/sphere>). A fast convolution between the transformed SH lighting and an analytical SH approximation of a spherical Gaussian [RH01a] of bandwidth G_σ can be performed by multiplying their Fourier representations.

Using this method, it only takes 4 seconds to perform an accurate convolution with a spherical Gaussian with $\sigma = 1$ degree (required SH-order: 325), as opposed to hours when using a naive SH implementation. Convolutions with larger σ 's are much faster, as only a lower-order SH representation is needed. The results are virtually free from ringing artifacts, as the SH-order can be chosen sufficiently high without much performance penalty. Alternatively, real-time frame rates can be achieved by performing the convolution




Figure 4: Percent L^2 error of 5 different fits to measured BRDF [MPBM03] data using our model. These fits use $K=3$ Gaussians and $N=15$ views. Grazing angles are more difficult to fit due elongated highlights (see Figure 2).

using graphics hardware [KVHS00], but at the cost of lower quality and less accurate filtering.

6. Results

In Figure 2 we compare original measured BRDF data with our fits. As can be seen, using only three Gaussians introduces some error, especially at grazing angles. Using more Gaussians will decrease the error at the expense of increased rendering times. Figure 4 shows the percent L^2 error of 5 different BRDFs as a function of view elevation angle. Our method has a peak percent L^2 error of 18% at severely grazing angles, and average error below 5% for normal views.

Figure 5 compares rendering glossy reflections with and without shadowing. Note that small-scale visibility events (e.g., in the wrinkles on the arm) are visually not important and in these areas both models are nearly indistinguishable. However, our low-frequency representation is well-suited for the large shadows that are quite noticeable (see the base of the statue). Figure 7 shows a comparison where visibility is approximated by the percentage of the occluded hemisphere [Bun04]. The renderings are similar, however there are regions where ambient visibility produces unwanted highlights in shadowed regions (e.g., the foot).

Figure 6 shows a model rendered with five different BRDFs. The approximation uses $K=3$ Gaussians and $N=15$ elevation angles. We represent and render a wide range of materials of varying degrees of specularity, all with high visual quality. When moving the lighting or changing the view, the reflections remain temporally coherent (see the accompanying video), even though the representation is sampled at only 15 views. For this particular model, the rendering speed was about 1000 frames per second (NVIDIA 8800 756MB RAM and 2.8Ghz P4). Timings are for offscreen rendering of shadowed glossy reflections only, and do not include the diffuse PRT component, or drawing to the screen. To our knowledge, this is the fastest method for rendering shadowed glossy reflections.



Figure 5: Visibility comparison. The same model is rendered with and without shadowing. Notice how small-scale shadow events are perceptually not important; at first glance both models look identical (upper half of model). However, larger shadows are quite noticeable. Our visibility representation exploits this observation. Model contains 100K vertices and glossy reflections render at 350Hz on a NVidia 8800.




Figure 6: The bunny model (35K vertices) rendered with five different BRDFs ($K=3$ Gaussians and $N=15$ views). A wide variety of materials can be captured with our technique. Glossy reflections are rendered at 1KHz with a NVidia 8800 card.

7. Discussion

Our technique is very effective for direct glossy reflections. It even allows on-the-fly changes of the BRDF without expensive reprocessing. However, we do not support inter-reflections. While glossy inter-reflections are not an impor-

tant effect on many models, it can be noticeable on adjacent planar surfaces.

We propose using a light-weight method for adding visibility. By moving the visibility component outside the integral, we simplify rendering dramatically. This approximation introduces an error, especially for blurrier reflections.

8. Conclusions

We have presented a real-time method for rendering realistic glossy reflections from environment maps. Our method is based on a nonlinear approximation of the BRDF based on a weighted sums of Gaussian functions. This representation is compact, captures arbitrary frequencies, and is efficient for rendering. In addition, we introduced an effective approximation to shadowed glossy reflections.

In the future, we would like to investigate reflections from bump mapped surfaces and the accompanying aliasing issues when viewed from afar. Our current model is a step in the right direction, and we are hopeful that it may help to prevent aliasing. We also hope to investigate methods to extend our technique to anisotropic BRDFs.

9. Acknowledgments

The authors thank the MIT Computer Graphics Group and in particular Tom Mertens for his insightful comments and Thomas Yeo for many discussions about SH convolution and the sphere. This work was supported by NSF CAREER award 0447561 *Transient Signal Processing for Realistic Imagery* and a Ford Foundation predoctoral fellowship. Jan Kautz acknowledges support from an Emmy-Noether fellowship from the German Research Foundation. Frédo Durand acknowledges a Microsoft Research New Faculty Fellowship and a Sloan fellowship.

References

- [AG02] ASHIKHMIN M., GHOSH A.: Simple blurry reflections with environment maps. *Journal of Graphics Tools* 7, 4 (2002), 3–8.
- [BN76] BLINN J., NEWELL M.: Texture and Reflection in Computer Generated Images. *Communications of the ACM* 19 (1976), 542–546.
- [Bun04] BUNNELL M.: Dynamic ambient occlusion and indirect lighting. In *GPU Gems 2: Programming Techniques for High-Performance Graphics and General-Purpose Computation* (2004), Addison-Wesley, pp. 223–233.
- [CL96] COLEMAN T., LI Y.: An Interior, Trust Region Approach for Nonlinear Minimization Subject to Bounds. *SIAM Journal on Optimization* 6, 2 (1996), 418–445.
- [CON99] CABRAL B., OLANO M., NEMEC P.: Reflection Space Image Based Rendering. In *Proceedings SIGGRAPH* (Los Angeles, California, August 1999), pp. 165–170.
- [GKMD06] GREEN P., KAUTZ J., MATUSIK W., DURAND F.: View-dependent precomputed light transport using nonlinear gaussian function approximations. In *SI3D '06: Proceedings of the 2006 symposium on Interactive 3D graphics and games* (New York, NY, USA, 2006), ACM Press, pp. 7–14.
- [Gre86] GREENE N.: Environment Mapping and Other Applications of World Projections. *IEEE Computer Graphics & Applications* 6, 11 (November 1986), 21–29.
- [HKRM03] HEALY D. M., KOSTELEK P., ROCKMORE D., MOORE S.: FFTs for the 2-Sphere-Improvements and Variations. *Journal of Fourier Analysis and Applications* 9, 4 (July 2003), 341–385.
- [HS99] HEIDRICH W., SEIDEL H.: Realistic, Hardware-accelerated Shading and Lighting. In *Proceedings SIGGRAPH* (Aug. 1999), pp. 171–178.
- [Kaj86] KAJIYA J.: The Rendering Equation. In *Proceedings SIGGRAPH* (August 1986), pp. 143–150.
- [KM00] KAUTZ J., MCCOOL M.: Approximation of Glossy Reflection with Prefiltered Environment Maps. In *Proceedings Graphics Interface* (May 2000), pp. 119–126.
- [KSS02] KAUTZ J., SLOAN P.-P., SNYDER J.: Fast, Arbitrary BRDF Shading for Low-Frequency Lighting Using Spherical Harmonics. In *13th Eurographics Workshop on Rendering* (June 2002), pp. 301–308.
- [KVHS00] KAUTZ J., VÁZQUEZ P.-P., HEIDRICH W., SEIDEL H.-P.: A Unified Approach to Prefiltered Environment Maps. In *Eleventh Eurographics Workshop on Rendering* (June 2000), pp. 185–196.
- [LK02] LATTA L., KOLB A.: Homomorphic Factorization of BRDF-based Lighting Computation. In *Proceedings SIGGRAPH* (July 2002), pp. 509–516.
- [LSS04] LIU X., SLOAN P.-P., SHUM H.-Y., SNYDER J.: All-Frequency Precomputed Radiance Transfer for Glossy Objects. In *Proceedings Eurographics Symposium on Rendering 2004* (June 2004), pp. 337–344.
- [MH84] MILLER G., HOFFMAN R.: Illumination and Reflection Maps: Simulated Objects in Simulated and Real Environments. In *SIGGRAPH Course Notes – Advanced Computer Graphics Animation* (July 1984).
- [MLH02] MCALLISTER D., LASTRA A., HEIDRICH W.: Efficient Rendering of Spatial Bi-directional Reflectance Distribution Functions. In *Proceedings Graphics Hardware* (September 2002), pp. 79–88.
- [MPBM03] MATUSIK W., PFISTER H., BRAND M., MCMILLAN L.: A data-driven reflectance model. *ACM Transactions on Graphics* 22, 3 (July 2003), 759–769.
- [NDM05] NGAN A., DURAND F., MATUSIK W.: Experimental analysis of brdf models. In *Rendering Techniques 2005: 16th Eurographics Workshop on Rendering* (June 2005), pp. 117–126.
- [NRH03] NG R., RAMAMOORTHY R., HANRAHAN P.: All-Frequency Shadows Using Non-linear Wavelet Lighting Approximation. *ACM Transactions on Graphics* 22, 3 (July 2003), 376–381.
- [NRH04] NG R., RAMAMOORTHY R., HANRAHAN P.: Triple Product Wavelet Integrals for All-Frequency Relighting. *ACM Transactions on Graphics* 23, 3 (Aug. 2004), 477–487.
- [Pho75] PHONG B.-T.: Illumination for Computer Generated Pictures. *Communications of the ACM* 18, 6 (June 1975), 311–317.
- [RH01a] RAMAMOORTHY R., HANRAHAN P.: A Signal-Processing Framework for Inverse Rendering. In *Proceedings SIGGRAPH* (August 2001), pp. 117–128.
- [RH01b] RAMAMOORTHY R., HANRAHAN P.: An Efficient Representation for Irradiance Environment Maps. In *Proceedings SIGGRAPH* (August 2001), pp. 497–500.



Figure 7: A Dragon model rendered with our real-time environment mapping method. Top row: visibility approximated as the percentage of the occluded hemisphere (Ambient Visibility). Bottom row: Our Spherical Harmonic visibility approximation. The left column shows full renderings and while the right column shows the visibility weights per pixel. Our approximate visibility technique eliminates highlights that occur in occluded regions (e.g. see foot and inset images) when rendered using ambient visibility.

[RH02] RAMAMOORTHY R., HANRAHAN P.: Frequency Space Environment Map Rendering. In *Proceedings SIGGRAPH* (July 2002), pp. 517–526.

[RWS*06] REN Z., WANG R., SNYDER J., ZHOU K., LIU X., SUN B., SLOAN P.-P., BAO H., PENG Q., GUO B.: Real-time soft shadows in dynamic scenes using spherical harmonic exponentiation. In *SIGGRAPH '06: ACM SIGGRAPH 2006 Papers* (New York, NY, USA, 2006), ACM Press, pp. 977–986.

[SKS02] SLOAN P.-P., KAUTZ J., SNYDER J.: Precomputed Radiance Transfer for Real-Time Rendering in Dynamic, Low-Frequency Lighting Environments. In *Proceedings SIGGRAPH* (July 2002), pp. 527–536.

[TS06] TSAI Y.-T., SHIH Z.-C.: All-frequency precomputed radiance transfer using spherical radial basis functions and clustered tensor approximation. *ACM Trans. Graph.* 25, 3 (2006), 967–976.

[WTL04] WANG R., TRAN J., LUEBKE D.: All-Frequency Relighting of Non-Diffuse Objects using Separable BRDF Approximation. In *Proceedings Eurographics Symposium on Rendering 2004* (June 2004), pp. 345–354.