Image-Based Rendering of Synthetic Diffuse Objects in Natural Scenes

Mais Alnasser* Hassan Foroosh*
Computational Imaging Lab, University of Central Florida

Abstract

We present a method for solving the global illumination problem for synthetic diffuse objects. The approach generates realistic shading for applications where a synthetic object is to be inserted in a natural scene. Using only a few images of the surroundings of the object, we first build an environment map for representing the ambient light. We then model the global illumination integral using Chebyshev polynomials. We show that due to the orthogonality of 2D Chebyshev moments, the global illumination integral can reduce to the inner product of two vectors, representing the irradiance and the Bidirectional Reflectance Distribution Function (BRDF). The Chebyshev moments of these two functions are computed off-line and stored in the memory. The rendering of the object in the scene then becomes a simple problem of computing the inner product of the two vectors for each point.

1. Introduction

Accurate modeling and computation of light transport is one of the key elements to realistic rendering of synthetic objects. Several physically based global illumination techniques have been developed in the past two decades to accurately model and compute the light distribution [3,5,11,12]. Our paper contains two main contributions to this area:

- We extend the use of cube maps [1,9,10,13] to image-based computation of global illumination.
- We generate a uniform grid for storage and interpolation of global illumination in a preprocessing step prior to rendering. Although this idea was originally proposed by Greger et al. [4], the data structure and interpolation strategy used in our approach is different. In particular, we use the uniform grid data structure for transport and storage of irradiance in the computation of global illumination, whereas Greger et al. used the data structure simply for storage of illumination precomputed using an off the shelf global illumination algorithm. Furthermore, we use Chebyshev moments for compact representation of the radiance field at a set of sampled points, which as shown below reduces the solution of global illumination to a simple inner-product during the rendering time.

The latter idea is similar to the use of spherical harmonics that are currently considered as the state of the art for representing low-frequency lighting. However, we will show that the results produced by Chebyshev polynomials are superior to those obtained by spherical harmonics.

*e-mail: {nasserm,foroosh}@cs.ucf.edu

2 Background

Let p be a point in space, and let ω_i and ω_o denote the incoming light direction and outgoing (viewing) direction at point p. The radiance $L(p, \omega_o)$ is defined as the power exiting from p, per unit solid angle in the direction ω_o, per unit projected area orthogonal to that direction. Radiance L is measured in watt meter^{-2} steradian^{-1}. Under equilibrium the distribution of radiance satisfies the following Fredholm integral equation of the second kind known as the light transport equation:

$$L(p,\omega_o) = L_e(p,\omega_o) + \int_{\Omega} L(p,\omega_i) f(p,\omega_i,\omega_o) \max(0, \omega_i \cdot \mathbf{n}_p) d\omega_i \quad (1)$$

where $L(p, \omega_o)$ is the radiance exiting p in the direction ω_o, L_e is included to account for the case when p is emitting light, f is the BRDF, which represents the ratio of irradiance to radiance, $L(p, \omega_i)$ is the radiance incident in p along the incoming direction, and Ω is the hemisphere of directions around the normal $\mathbf{n}_p$ at point p.

Hereafter, for brevity of notation we use $L_{\text{in}} = L(p, \omega_i) \max(0, \omega_i \cdot \mathbf{n}_p)$. Also, we will denote the forshortened BRDF as $f_c = f(p, \omega_i, \omega_o)$.

3 Chebyshev Representation of Incoming Radiance and BRDF

We will model the incoming radiance using cube maps [8]. A cube map is basically a set of six texture maps arranged in a cross shape, which when folded over the object

The 18th International Conference on Pattern Recognition (ICPR'06)
0-7695-2521-0/06 $20.00 © 2006 IEEE

to be rendered can be used for approximating the radiosity. Figure 5 shows an example of a cube map that we generated using a standard off the shelf camera. Essentially, to construct a cube map, we placed ourselves in the environment that defines the lighting conditions, and took pictures in six mutually orthogonal directions corresponding to the faces of a cube. Using a cube map for global illumination has the advantage that it allows for image-based modeling of global illumination. The incoming radiance can then be approximated by taking samples in many directions from inside the cube where the object of interest will be rendered. These samples are used to reconstruct an approximation to the incoming radiance function.

Various approaches have been proposed in the literature for such approximation. For low-frequency lighting, the state of the art currently is considered to be the recently introduced spherical harmonics [8, 13], which are based on associated Legendre polynomials. Although, spherical harmonics provide a very good solution to the problem, we show in this paper that their performance is inferior to Chebyshev polynomials in terms of the trade off between quality, the order of approximation, and the computational complexity.

Chebyshev polynomials belong to a class of polynomials known to be orthogonal. They provide an orthogonal basis for approximating functions using the so called Chebyshev moments. Unlike continuous moments such as Zernike, and Legendre, discrete Chebyshev moments have been defined in the literature [7]. This avoids the requirement for discretization and the associated errors. On the other hand, continuous moments are typically defined over the unit sphere requiring a coordinate transformation, which the discrete Chebyshev moments do not require. There is, however, one disadvantage. Discrete chebyshev moments become unstable when the moment order becomes large. A solution to this problem is using the more stable orthonormal Chebyshev polynomials, which have been discussed in [6] and also in this paper.

For a given positive integer N, and for an integer value i in the range $[0, N-1]$, the orthonormal chebyshev discrete polynomials [6] $t_n(x)$ are defined as following:

$$t_n(x) = \alpha(2x+1-N)t_{n-1}(x) + \beta t_{n-2}(x)$$

where,

$$\begin{aligned} n &= 2, ..., N-1, \quad x = 0, ..., N-1 \\ \alpha &= \frac{\sqrt{4n^2-1}}{n\sqrt{N^2-n^2}}, \\ \beta &= -\frac{n-1\sqrt{2n+1}\sqrt{N^2-(n-1)^2}}{n\sqrt{2n-3}\sqrt{N^2-n^2}} \end{aligned} \quad (2)$$

and $t_0(x) = N^{-1/2}$ and $t_1(x) = \frac{\sqrt{3}(2x+1-N)}{\sqrt{n(N^2-1)}}$.

In our case, since we use images (cube maps) to represent the incoming radiance, N can be chosen equal to the image size (assuming that the images are $N \times N$). The above polynomials are orthonormal, which avoids the need for the normalization required for stability.

The Chebyshev moments of the incoming radiance are then given by

$$T^L_{mn} \simeq \sum_{i=0}^{s-1}\sum_{j=0}^{s-1} t_m(i)t_n(j)L_{\text{in}}(i,j) \quad (3)$$

where the series is truncated to the order s, and $0 \leq m, n \leq N-1$.

Note that L_{in} depends on the altitude and the azimuth angles mapped to integer values i and j. Also in practice, s is taken as the square root of the number of samples taken to approximate the function, i.e. in our case the number of samples taken from the cube map. Given the Chebyshev moments of the incoming radiance, we can reconstruct it using

$$L_{\text{in}}(i,j) \simeq \sum_{i=0}^{k-1}\sum_{j=0}^{k-1} T^L_{mn} t_m(i)t_n(j) \quad (4)$$

where k is square root of the number of coefficients required. The approximation is of course due to truncation, which in practice is applied to save in computational time. For low frequency light, this truncation error is in practice very small.

Similarly the BRDF can be expanded in terms of its Chebyshev moments using

$$f_c(i,j) \simeq \sum_{i=0}^{k-1}\sum_{j=0}^{k-1} T^f_{mn} t_m(i)t_n(j) \quad (5)$$

where

$$T^f_{mn} \simeq \sum_{i=0}^{s-1}\sum_{j=0}^{s-1} t_m(i)t_n(j)f_c(i,j) \quad (6)$$

The advantage of this representation lies in the fact that the discrete 2D Chebyshev polynomials are orthogonal. As a result, assuming that the point p is not emitting any light, we can write the global illumination equation for point p in (1) as

$$L_p = \sum_{m=0}^{N-1}\sum_{n=0}^{N-1} T^L_{mn}T^f_{mn} \quad (7)$$

Therefore the solution at any point p reduces to a simple inner product of the Chebyshev moments for the incoming radiance and the forshortened BRDF. The Chebyshev moments of both the incoming radiance and the forshortened BRDF can be pre-computed and stored in a structured grid. The rendering process can then be accelerated, since it would involve only an inner product per point. In particular, the process can also be mapped into the graphics hardware for real-time and interactive rendering.

We applied our approach to synthetic spheres in different cube map environments. The choice of sphere objects allows us to account for all possible normal directions. We generated some of these cube maps by taking pictures of the environment. To approximate the incoming radiance at a point inside the cube map, samples are taken along different altitude and azimuth directions. Using these samples the incoming radiance is represented by Chebyshev moments. Similarly the forshortened BRDF is sampled and projected into the Chebyshev space. The outgoing radiance is then simply given by the inner product of the two. This is similar to spherical harmonics, which is considered as the state of the art for low frequency light. Therefore we evaluated our results against spherical harmonics.

For this purpose, we generated a gold standard result for each environment using the Monte Carlo quadrature integration and importance sampling. In each case, we increased the number of samples in Monte Carlo integration until the gold standard result would remain practically invariant to further sample increases. We then generated the results for both Chebyshev and spherical harmonics, with different orders of approximations, and compared the performance of each method against the gold standard using Signal to Noise Ratio (SNR) and relative error. Results are plotted in Figure 2 for varying orders for both methods.

In practice, our approach is always outperforming spherical harmonics by a significant SNR difference. Typically, the SNR of the spherical harmonics approaches ours at much higher orders of approximation. Therefore Chebyshev moments can be used to both save in computational time and memory.

4 Discussion and Conclusion

We have presented a new image-based approach that would allow for interactive rendering of synthetic objects using the light captured from a natural scene. The proposed approach outperforms the spherical harmonics significantly, which are considered as the state of the art for dynamic low frequency lighting. The method can be readily combined with precomputed transfer methods to also incorporate self-shadowing and self-interreflection effects if required. Also, since the proposed method shares common algorithmic features with spherical harmonics, it is also possible to obtain real-time rendering in a dynamic lighting environment by mapping the algorithm to the graphics hardware as in [7].

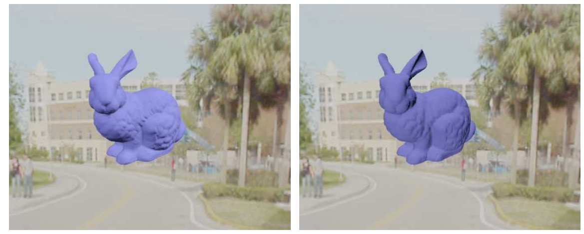

Figure 1. Results for bunny and a cube map from outdoor images. Left: spherical harmonics (4 coefficients) Right: Chebyshev (4 coefficients).

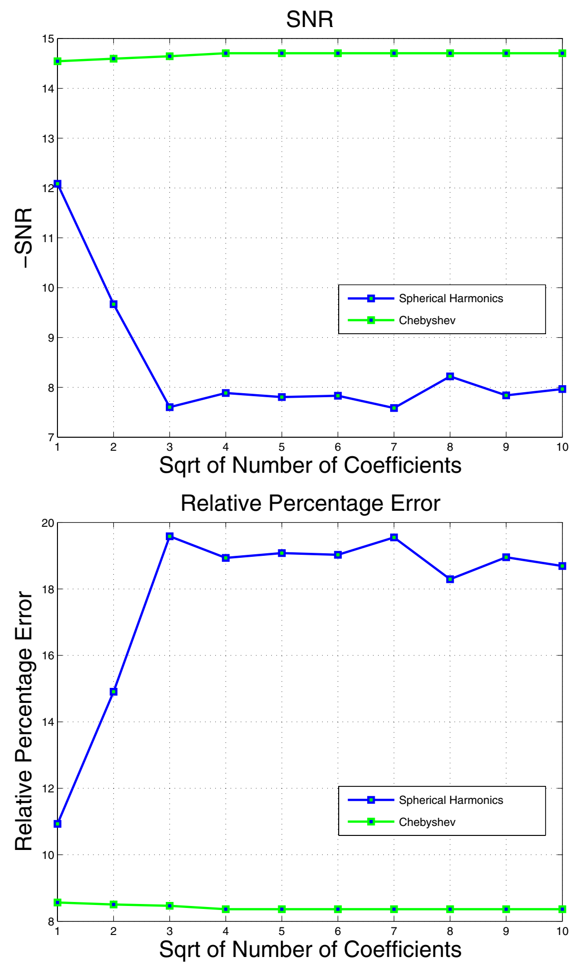

Figure 2. Comparison of Average SNR and Relative Error for Spheres in different environments

References

[1] J. Blinn and M. Newell. Texture and reflection in computer generated images. *Communications of the ACM*, 19:542–

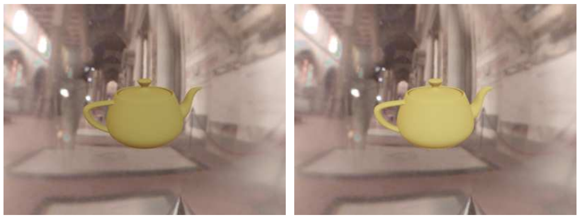

Figure 3. Results for Teapot and Gelileo cube map environment. Left: spherical harmonics (9 coefficients) Right: Chebyshev (9 coefficients).

Figure 4. Buddha rendered in the Galileo cube map environment using Chebyshev moments (9 coefficients)

Figure 5. An example of one of our cube maps

546, 1976.

[2] B. Cabral, N. Max, and P. Nemec. Bidirectional reflection functions from surface bump maps. In *SIGGRAPH*, pages 273–281, 1987.

[3] M. Cohen and J. Wallace. *Radiosity and Realistic Image Synthesis*. A.P. Professional, 1993.

[4] G. Greger, P. Shirley, P. Hubbard, and D. Greenberg. The irradiance volume. *IEEE Computer Graphics & Applications*, 18(2):32–43, 1998.

[5] H. Jensen. *Realistic Image Synthesis using Photon Mapping*. A.K Peters, 2001.

[6] R. Mukundan. Improving image reconstruction accuracy using discrete orthonormal moments. In *CISST*, pages 287–291, 2003.

[7] R. Mukundan and P. Lee. Image analysis by tchebichef moments. *IEEE Trans. on Image Processing*, 10(9):1357–1364, 2001.

[8] M. Nijasure, S. Pattanaik, and V. Goel. Real-time global illumination on gpu. *Journal of Graphics Tools*, 2004.

[9] R. Ramamoorthi and P. Hanrahan. An efficient reproduction for irradiance environment maps. In *SIGGRAPH*, pages 497–500, 2001.

[10] R. Ramamoorthi and P. Hanrahan. Frequency space environment map rendering. *ACM Transactions on Graphics*, 21(3):517–526, 2002.

[11] P. Shirley. *Realistic Ray Tracing*. A.K. Peters, 2000.

[12] F. Sillion and C. Puech. *Radiosity and Global Illumination*. Morgan Kaufman Publishers, 1994.

[13] P.-P. Sloan, J. Kautz, and J. Snyder. Fast arbitrary brdf shading for low-frequency lighting using spherical harmonics. In *SIGGRAPH*, pages 291–296, 2002.