



Shading For Fourier Volume Rendering

Alireza Entezari¹, Randy Scoggins², Torsten Möller¹, Raghu Machiraju³

¹Graphics, Usability, and Visualization Lab, Simon Fraser University

²Randy Scoggins, U.S. Army Engineer Research and Development Center

³Computer and Information Science, The Ohio State University

Abstract

The work presented here describes two methods to incorporate viable illumination models into Fourier Volume Rendering (FVR). The lack of adequate illumination has been one of the impediments for the wide spread acceptance of FVR. Our first method adapts the Gamma Corrected Hemispherical Shading (GCHS) proposed by Scoggins et al. [11] for FVR. We achieve *interactive* rendering for constant diffusive light sources. Our second method operates on data transformed by spherical harmonic functions. This latter approach allows for illumination under varying light directions. It should be noted that we only consider diffuse lighting in this paper. We demonstrate and compare the effect of these two new models on the rendered image and document speed and accuracy improvements.

Keywords: *Fourier Volume Rendering, Shading, Spherical Harmonics, Fourier Transform*

1. INTRODUCTION

Fourier domain volume rendering (FVR), introduced by Malzbender [6], and extended by Totsuka and Levoy [12], is a direct volume rendering algorithm that is significantly different from all commonly used ones (for an overview of volume rendering algorithms see [7]). FVR is based on the Fourier projection-slice theorem [2]. Therefore the algorithmic complexity for rendering an N^3 volumetric dataset is $O(N^2 \log N)$ after incurring a $O(N^3 \log N)$ pre-processing step. All other rendering algorithms are required to traverse the entire volume and hence have an algorithmic complexity of $O(N^3)$. By suitably pre-multiplying the frequency data, one can also achieve depth cues and effects similar to diffuse lighting.

There are two impediments to widespread acceptance of Fourier volume rendering. The first one is the inability to deal with occlusion. At best, the images produced by FVR resemble X-Ray images. The Radon transform which is the basis for the existence of the projection-slice theorem is essentially a summation of the function values along the rays through the object. Unless fundamental advances are brought to include opacity weighted sums into

the back-projection process, occlusion will remain a major impediment for widespread acceptance of FVR.

Another major impediment has been the lack of illumination models in the Fourier domain. Moreover, every time that lighting attributes are altered the Fourier transform has to be re-computed and several expensive frequency space operations have to be completed. Lighting is essentially non-linear. Consider diffuse illumination. A surface or voxel point is illuminated by evaluating the operator $max(0, \mathbf{N} \cdot \mathbf{L})$, where \mathbf{N} is the surface normal at the point under discussion and \mathbf{L} is the direction of light pointing on the surface. This is a non-linear operator since $max(0, \mathbf{N} \cdot (\mathbf{L}_1 + \mathbf{L}_2))$ is not equal to the sum of $max(0, \mathbf{N} \cdot \mathbf{L}_1)$ and $max(0, \mathbf{N} \cdot \mathbf{L}_2)$. The non-linearity of this operator precludes its projection into the Fourier basis. The achievement of Totsuka and Levoy [12] was the replacement of the non-linear illumination with the hemispherical illumination operator. This operator is proportional to the quantity $(1 + \mathbf{N} \cdot \mathbf{L})$. Now, the Fourier transform can be applied to all sample points hemispherically shaded. However, the approximation results in fairly poor images (see Fig. 4). Scoggins et al. [11] suggested a better approximation by raising the hemispherical reflection term to the cube. They called it the Gamma-Corrected Hemispherical Shading (GCHS). Their insight was that culling back-lit surfaces is equivalent to using the maximum operator. Hence, if an approximate function treats front-facing surfaces more favorably than back-facing ones, our goal of shading in the Fourier domain would be more realizable. However, it should be noted that the GCHS operators are not linear. On the other hand, one could still project the polynomial functions onto the Fourier basis. In this paper, we show how one could achieve the implementation of common illumination operators in the frequency domain and increase the effectiveness of FVR. The results obtained are of higher quality than those obtained in earlier work.

An advantage of the approach of Scoggins [11] is that re-lighting can be conducted for surfaces. Additional storage is certainly needed. However, it should be noted that for volumetric data the expense can be staggering. Re-lighting and re-viewing can easily be achieved if the function being subjected to a Fourier transform does not depend on the light and observer attributes. Light and view directions are often described in a global co-ordinate system. Spherical harmonic transforms allow the expression of illumination on a sphere as a linear combination of harmonic basis functions parametrized by the two orientation angles in a polar space. Through adequate rotations it can be shown that the illumination is

a linear combination of the transformed light directions. This allows for quick re-lighting. It should be noted that although there do not exist explicit surfaces, the spherical harmonic transform will have an impact on the volumetric-gradient implied surfaces.

Thus, a comprehensive framework for incorporating lighting in the frequency domain requires the application of both the spherical harmonics and the Fourier transform. The Fourier back projection can still be used for quick renderings. However, there is a cost that is associated with transforming shading operators into harmonic domains. Memory consumption can be high given the expression of intermediate operations as convolutions. This work reported here is an exploration of the research questions that address illumination operators in the frequency domain. The work in our opinion is in no way complete.

In Section 2 we discuss previous work in Fourier volume rendering and the use of spherical harmonics towards surface lighting problems. Section 3 describes the Gamma Corrected Hemispherical shading operator, while Section 4 describes the projection of this operator onto the Fourier basis. In Section 5 we discuss the use of spherical harmonics for volumetric shading and explain how re-lighting can be realized easily. Where necessary, we include appropriate images to drive home some observations and claims about our methods. Finally, in Section 6 we offer a summary and point to the future.

2. PREVIOUS WORK

Fourier Volume Rendering was first introduced by Malzbender [6] and Totsuka and Levoy [12] in 1993. Totsuka and Levoy extended the ideas of Malzbender to include diffuse and depth shading into Fourier Volume Rendering.

The idea of Fourier volume rendering was combined with the wavelet transform by Gross et al. [4] as a multi-resolution acceleration to the algorithm. Their method also benefits from the fact that the Fourier transforms of the wavelets and scaling functions can be computed analytically. Westenberg et al. [13] propose Fourier Wavelet Volume Rendering. This allows a direct computation of the wavelet decomposition of the rendered image, which leads to quick progressive refinement implementations of the volume rendered images. However, all of the Fourier volume rendering algorithms are limited to parallel projections and X-ray type rendering. Since the accuracy of the slicing operation has tremendous effects on the image quality, hardware accelerated slicing may cause unwanted artefacts. However, there do exist methods that exploit suitable hardware including texture maps [3].

Spherical harmonics have been applied to diverse problems related to general reflection functions [14], reflection variations due to bump-mapped surface [3], re-lighting images in image-based rendering systems [15], and to the acceleration of global illumination calculations [10]. Spherical harmonics are the analog to Fourier functions in the spherical coordinate system, with properties similar to the Fourier basis functions. Signals can be approximated with a linear equation, but many coefficients are often necessary to accurately reconstruct signals containing high frequency components. Recent work [1][8] has shown that a large number of coefficients in a spherical harmonic series can be set to zero if the reflection function being approximated is evaluated in a preferential local coordinate system.

3. APPROXIMATE SHADING

Images produced by normal diffuse shading employ a non-linear function, the \max function, to account for visibility for illumination. Typically the diffuse term is computed by

$$\max(0, N \cdot L), \quad (1)$$

where N is the normalized gradient vector and L is the normalized light vector. This requires repeating linear transformations, such as the Fourier transform, whenever lighting changes. If visibility is ignored, negative image intensity values result. Alternative shading is afforded by hemispherical formulation and the corresponding operator is computed as:

$$E = \frac{1}{2}(1 + N \cdot L) \quad (2)$$

Hemispherical shading has been used [5][12] as a substitute diffuse shading function since linear transforms may be applied to the surface normal data once and re-evaluated in transformed space as lighting changes. Hemispherical shading is similar to that seen in an overcast outdoor setting and produces reduced image contrast. To minimize the visual difference between hemispherical and diffuse shading, one can use gamma correction and scaling given by:

$$E = \left(\frac{1 + N \cdot L}{2} \right)^g = \sum_{i=1}^G a_i V_i \quad (3)$$

The Gamma Corrected Hemispherical Shading (GCHS) has been used [11] to modify hemispherical image histograms so that the image approximates diffuse shading with the \max function. If an integer gamma correction factor g is selected, an exact shading polynomial can be obtained in which coefficients a_i , $1 \leq i \leq G$, are derived from a light vector only, and the same number of geometric terms V_i from unit surface normals only. Therefore, linear transforms may be applied to the V_i and re-evaluated for changing lights, through a_i without performing a new transform.

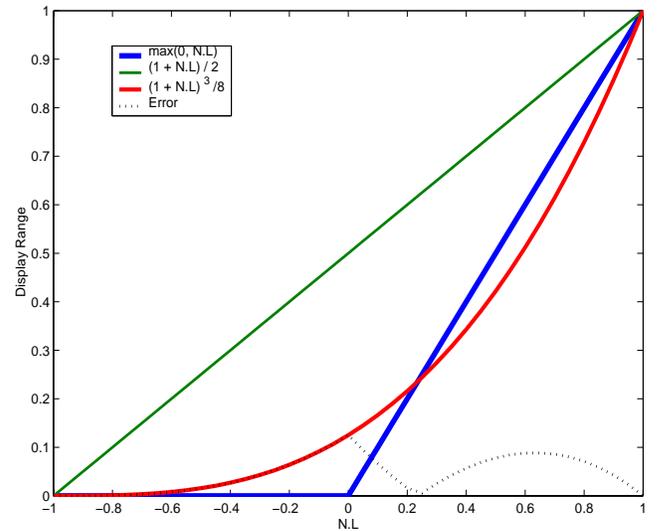


FIGURE 1. Differences between Linear and Cubic approximations to Max function

To illustrate how gamma correction can significantly reduce visual differences between hemispherical and diffuse shading consider Fig. 1. Values on the abscissa measure the *inner product* of the surface normal vectors, N , and the light vector L . The ordinate values represent the device display range normalized from zero (0) to one (1). The curve labeled $\max(N \cdot L, 0)$ characterizes the commonly used mapping in which negative *dot product* results are clamped to zero (0) and positive results are mapped to the entire display range of the device. The curve labeled $(1 + N \cdot L)/2$ is the hemispherical shading function in which the full range of the *dot product* $N \cdot L$ is mapped to the display range. This produces positive display values even for back-facing surfaces which would be set to zero (0) by the *max* function. Note that, prior to dividing $(1 + N \cdot L)$ by two (2) to scale to the display range, positive values of $N \cdot L$ correspond to values greater than or equal to one (1) in $(1 + N \cdot L)$ and negative $N \cdot L$ values correspond to those less than one (1)

Applying a gamma correction of greater than one (1) to $(1 + N \cdot L)$ and then scaling can therefore greatly reduce the display intensities that correspond to a negative $N \cdot L$. While values greater than one (1) are also effected, the error relative to $N \cdot L$ can be made to be relatively small by selecting an appropriate gamma exponent. In this manner, we eliminate the necessity of employing the *max* operator for visibility without introducing negative intensity, and produce images that are visually almost indistinguishable from normal shading (see Fig. 2). The smallest possible integer exponent is most useful since fewer terms are produced in a polynomial expansion of the equation, as in [11]. An exponent of three (3) yields the curve $(1 + N \cdot L)^3/8$ in Fig. 1, has been found to be optimal since it is the smallest exponent that produce an inflection point, allowing the curve to pass through both 0 and 1 . This results in the *Error* curve of Fig. 1. A more detailed explanation of the GCHS method and its' salient features are available in [11].

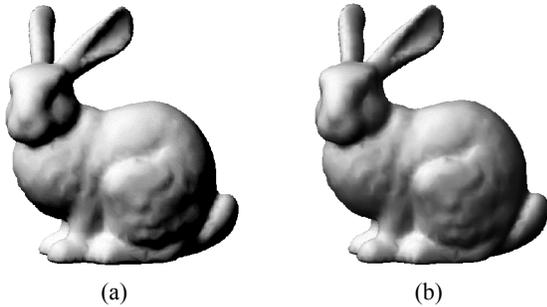


FIGURE 2. (a) Diffuse shading, (b) Gamma corrected hemispherical shading.

4. FOURIER DOMAIN APPROXIMATE SHADING

The utility of the gamma corrected shading is to approximate the actual diffuse shading term $\max(N \cdot L, 0)$ with the better approximation namely, $(1 + N \cdot L)^g$. Since this term is also not linear, a direct implementation is susceptible to changes in light-source direction. It is necessary to subject the shaded version of the data to a Fourier Transform. An alternative is the following: for changes in light direction, the shading is computed through a linear combination of various convolutions of versions of the data in the

frequency domain. In this section we show how this can be realized. The operation necessary for gamma-corrected shading is:

$$(1 + N \cdot L)^3 = \left(1 + \frac{\nabla f \cdot L}{\|\nabla f\|}\right)^3 = \left(1 + 3 \frac{\nabla f \cdot L}{\|\nabla f\|} + 3 \frac{(\nabla f \cdot L)^2}{\|\nabla f\|^2} + \frac{(\nabla f \cdot L)^3}{\|\nabla f\|^3}\right) \quad (4)$$

The above computation can be implemented in the Fourier domain by considering each of the terms separately in the following way.

4.1 Linear ($\nabla f \cdot L$)

This has already been shown by Totsuka and Levoy [12]. We simply remember that a derivative in the spatial domain is synonymous with a linear ramp function in the Fourier domain. We first express the term under consideration as:

$$\nabla f \cdot L = f_x l_1 + f_y l_2 + f_z l_3 \quad (5)$$

Here, we denote the derivative components with a subscript, i.e.

$\frac{df}{da} = f_a$. The components of the light vector are denoted by $L = (l_1, l_2, l_3)$. Now we can conclude that the Fourier transform (denoted by $T\{\cdot\}$) is:

$$\begin{aligned} T\{\nabla f \cdot L\} &= i\omega_x F l_1 + i\omega_y F l_2 + i\omega_z F l_3 \\ &= (i\omega_x l_1 + i\omega_y l_2 + i\omega_z l_3) F \\ &= (i\Omega \cdot L) F \end{aligned} \quad (6)$$

It should be marked that we use the notation $T\{f\} = F$. This particular term can be computed in a very straightforward manner [12].

4.2 Quadratic ($(\nabla f \cdot L)^2$)

A brute-force evaluation would lead to

$$\begin{aligned} (\nabla f \cdot L)^2 &= (f_x l_1 + f_y l_2 + f_z l_3)^2 \\ &= l_1^2 f_x^2 + l_2^2 f_y^2 + l_3^2 f_z^2 + 2l_1 l_2 f_x f_y + 2l_1 l_3 f_x f_z + 2l_2 l_3 f_y f_z \end{aligned} \quad (7)$$

We can re-write any combination of two derivatives f_a and f_b in the following way:

$$f_a f_b = \frac{1}{2} (f_a^2 + f_b^2) - \frac{1}{2} (f_a - f_b)^2 \quad (8)$$

Hence, we can re-write Equation 7 in the following way in the frequency domain:

$$T\{(\nabla f \cdot L)^2\} = \frac{1}{2} (i\Omega \cdot L)^2 (F \otimes F) - (F \otimes ((i\Omega \cdot L)^2 F)) \quad (9)$$

Note that convolution is denoted by the symbol \otimes .

4.3 Cubic ($(\nabla f \cdot L)^3$)

The cubic term can be computed similarly. Here we use the fact, that

$$f_a f_b f_c = \frac{1}{6}(f^3)_{abc} - \frac{1}{2}f(f^2)_{abc} + \frac{1}{2}f^2 f_{abc} \quad (10)$$

From this we conclude

$$\begin{aligned} T\{(\nabla f \cdot L)^3\} &= \frac{1}{6}(i\Omega \cdot L)^3(F \otimes F \otimes F) \\ &- \frac{1}{2}(F \otimes ((i\Omega \cdot L)^3(F \otimes F))) + \frac{1}{2}(F \otimes F \otimes ((i\Omega \cdot L)^3 F)) \end{aligned} \quad (11)$$

Equation 6, Equation 9, and Equation 11 show how the components of Equation 4 can be implemented in the Fourier domain. Please note that the normalization with $1/\|\nabla f\|^k$ can be precomputed as well, just as Totsuka and Levoy suggested [12]. However, a problem for efficient implementation remains. The equations above require a convolution with a full volume. This can only be precomputed for a constant light source for real-time rendering. A change in the light direction will require that the terms in Equation 9 and 11 be recomputed. On the other hand no transformation of the data is required. Thus, our first method will allow us to change the light sources albeit at a higher cost. However, for constant light sources this method can allow for expedient rendering as the observer position changes. Our second proposed method, on the other hand, allows for relighting.

4.4 Results

To demonstrate our methods, we have rendered the data sets shown in Fig. 3 and Fig. 4. The first example compares linear and cubic illumination for a synthetic dataset of size 128^3 . The dataset consists of a sphere with high intensity value inside a bigger sphere with a lower intensity value. The images in the top row use the original FVR, which does not incorporate shading information. Of these two images, the right image uses depth cueing to give higher intensities to the surfaces closer to the viewer. The middle row adds linear (hemispheric) shading to the model. These images demonstrate the added advantage of this illumination model over basic FVR images. The light intensity diminishes as the surface normal is away from the light direction. The images in the bottom row incorporate the GCHS illumination model to allow a more accurate approximation of the *max* function. In the images obtained from the application of this illumination model surfaces on the right hemisphere are less lit than the left hemisphere. This variation of intensity is what one expects to observe on diffusively lit surfaces. Another noteworthy observation is that the images from linear shading have significantly higher erroneous energy comparing to the GCHS version (especially visible in Fig. 4). This result confirms the theory as expounded in Section 3 and characterized by various plots in Fig. 1.

The second set of images in Fig. 4 are obtained from rendering the UNC brain dataset of size 256^3 . As in the spherical dataset, the images in the top row are not influenced by the light in the scene. The middle row incorporates linear shading and the images exhibit higher luminosities since a linear approximation of the *max* function is employed. The images in the bottom row do not exhibit this excess intensity. As a result, the surfaces are illuminated more accurately. This has a dramatic effect on the surfaces facing away from the light, as displayed in Fig. 4. It should be noted that the light source is placed to the left of the dataset.

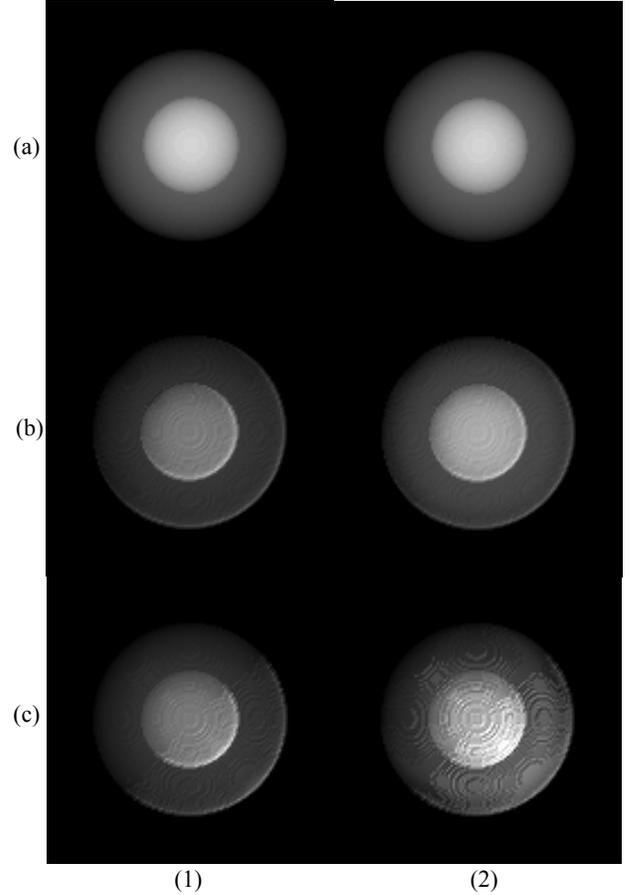


FIGURE 3. Rendering of a synthetic sphere. (a) traditional FVR without any shading. (b) using Totsuka and Levoy’s hemispherical reflections. (c) Our cubic shading term. Column (1) is without depth cueing. Column (2) includes depth cueing. In the lit images, light is coming from (1,1,1).

The rendering speed for the sphere (UNC brain) for original FVR is 116 frames per second (fps) (27fps) without depth cueing and 79fps (17fps) using depth cueing. For linear and cubic illumination two volumes are sliced instead of one volume in the original FVR (the original data set as well as its gradient magnitude volume). Hence the resulting frame-rates are somewhat lower - 98fps (21fps) without the depth cueing and 64fps (14fps) with depth cueing. (The depth cueing requires even more computations, which explains the lower frame rates.) All our experiments were conducted on a Pentium 4 2.2GHz processor equipped with 1.5GB main memory.

5. VOLUME SHADING WITH SPHERICAL HARMONICS

A spherical harmonic approximation allows the diffuse shading of surfaces without repeated 3D Fourier transforms. We will first introduce spherical harmonics briefly and then explain how they can be exploited for FVR.

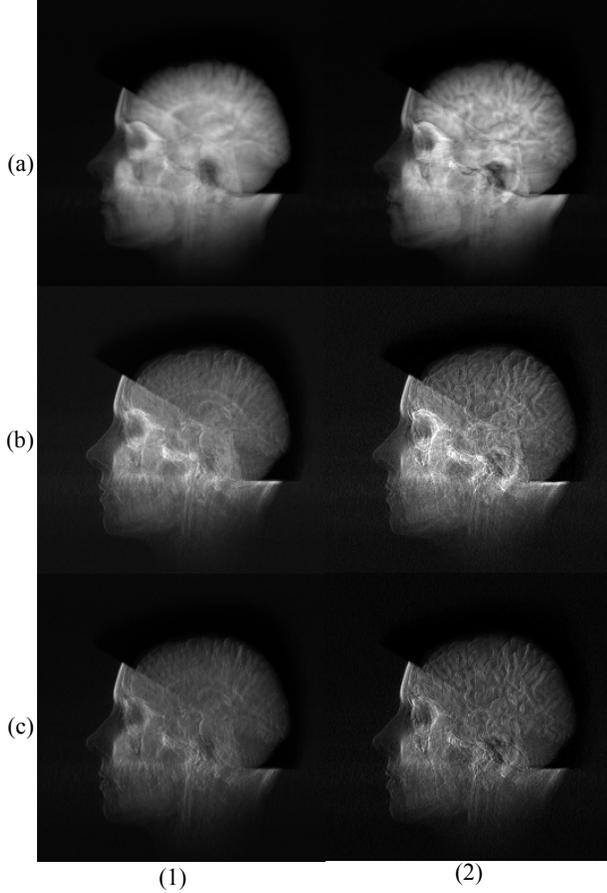


FIGURE 4. Rendering of the UNC brain data set. (a) traditional FVR without any shading. (b) using Totsuka and Levoy's hemispherical reflections. (c) Our cubic shading term. Column (1) is without depth cueing. Column (2) includes depth cueing. In the lit images light is coming from left

5.1 Spherical Harmonic and Surface Shading

Spherical harmonics are a group of functions that form an orthonormal basis on the unit sphere. While they are often given in complex form, a real version of the transform exists and is suitable for images and scalar which consist of integers and real numbers. Real spherical harmonics are computed as:

$$Y_{lm} = \sqrt{\frac{2l+1(l-|m|)!}{4\pi(l+|m|)!}} P_m^l(\cos\theta) \begin{cases} \cos m\phi & \text{for } m \geq 0 \\ \sin |m|\phi & \text{for } m < 0 \end{cases} \quad (12)$$

where P_m^l are the associated Legendre polynomials, $0 \leq l \leq \infty$, and $-l \leq m \leq l$. Spherical harmonics are the analog of the Fourier basis functions in spherical coordinates, and the transform has properties similar to that of the Fourier transform. A function f may be approximated to any degree of accuracy by the expansion:

$$f(\theta, \phi) = \sum_l \sum_m f_{lm} Y_{lm}(\theta, \phi) \quad (13)$$

where

$$f_{lm} = \int_0^{2\pi} \int_0^\pi f(\theta, \phi) Y_{lm}(\theta, \phi) \sin\theta(d\theta)d\phi \quad (14)$$

The series is exact for $l = \infty$ but in many cases very good approximations may be obtained with only a few terms [1][8]. For instance, the diffuse shading equation can be approximated well with low order harmonics.

Significantly, the effect of the *max* function can be considered by setting $f(\theta, \phi) = \max(\theta)$, or equivalently not integrating over the lower hemisphere in Equation 14. This leads to a linear shading equation that prevents significant back-face lighting using the approximation in Equation 13. The most compact and accurate approximation is obtained when a surface's normal corresponds to the z-axis of the global coordinate system. Though this is generally not true, it is possible to rotate the spherical harmonic basis to combine surface normals and lighting direction in a single coordinate system. This, and the interpreting of shading in terms of convolution in spherical coordinates, leads to an elegant diffuse shading equation that both incorporates the *max* function and is linear.

First, we express diffuse shading in terms of the general illumination equation. The radiance due to diffuse reflection originating from direction ω that is reflected into direction ω' is given by:

$$E = \int_{\Omega} L(\omega') \rho(\omega', \omega) \max(\cos\theta, 0) d\omega' \quad (15)$$

As described in [9], this equation may be interpreted as a convolution in spherical coordinates. For diffuse shading, ρ is a constant (set to $1/\pi$) and reflection from surfaces facing away from the light source is handled by applying the *max* function to the area projection term $\cos\theta$. If a local coordinate system is used where the surface normal is the z axis, then $N \cdot L = \cos\theta$. Using spherical coordinates for direction, this equation may be expressed as a *spherical harmonic* series by substituting the spherical harmonic forms of L and $\max(\cos\theta, 0)$ into Equation 15. The light and *max* expansions are given by:

$$L(\theta_L, \phi_L) = |L| \sum_{l=0}^M \sum_{m=-l}^l L_{lm} Y_{lm}(\theta_L, \phi_L) \quad (16)$$

and

$$\max(\cos\theta, 0) = \sum_{l=0}^M \sum_{m=-l}^l c_{lm} Y_{lm}(\theta', \phi') \quad (17)$$

respectively. Note that under the constraint imposed above -- the surface normal coincides with the z axis --- for isotropic reflection, there is no azimuthal dependence. Equation 17 reduces to:

$$\max(\cos\theta, 0) = \sum_{l=0}^M c_{l0} Y_{l0}(\theta', \phi') \quad (18)$$

Equation 15 may now be written as:

$$E = \frac{|L|}{\pi} \sum_l \sum_m \sum_p L_{lm} c_{p0} \int_0^{2\pi} \int_0^\pi Y_{lm}(\theta_L, \phi_L) Y_{p0}(\theta', \phi') d\theta' d\phi' \quad (19)$$

However, the integration is performed in the local coordinate system in which the reflection function is evaluated, indicated by θ' and ϕ' , while the light direction is expressed in global un-primed coordinates θ_L and ϕ_L . To perform the integration, both coordinate systems must be expressed in the same coordinate system. Spherical harmonics are closed under rotation so Y_{lm} in global coordinates may be expressed in terms of a series of other spherical harmonics evaluated in a local coordinate system with normal equal to the z axis. For a normal vector with global spherical coordinates (α, β) , the rotation is given by:

$$Y_{lm}(R_{\alpha, \beta}(\theta', \phi')) = \sum_{m'=-l}^l D_{mm'}^l(\alpha) e^{im'\beta} Y_{lm}(\theta', \phi') \quad (20)$$

where $R_{\alpha, \beta}(\theta', \phi')$ represents a rotation of local spherical coordinates to their values in the global coordinate system. Following [8], Equation 20 is substituted in Equation 19 and the integration performed. The orthogonality property of spherical harmonics results in a non-zero integral only for $p=l$. Also, since c_{pq} is zero except for $q=0$, Equation 15 simplifies to:

$$E = \frac{|L|}{\pi} \sum_l \sum_m L_{lm} c_{l0} D_{m0}^l(\alpha) e^{im\beta} \quad (21)$$

where the rotation coefficient is given by [8]:

$$D_{m0}^l(\alpha) e^{im\beta} = \sqrt{\frac{4\pi}{2l+1}} Y_{lm}(\alpha, \beta) \quad (22)$$

Equation 21 may be simplified further for the case of a directional light source since $L_{lm} = Y_{lm}(\theta_L, \phi_L)$. Substituting the directional light term and Equation 22 in Equation 21 yields the final form of the spherical harmonic rendering equation as:

$$E = \frac{|L|}{\pi} \sum_l \sum_m \sqrt{\frac{4\pi}{2l+1}} Y_{lm}(\theta_L, \phi_L) c_{l0} Y_{lm}(\alpha, \beta) \quad (23)$$

Equation 23, reveals shading as a convolution in spherical coordinates since it involves the summation of spherical harmonic coefficients for the light and reflection functions. As is the case for the Fourier Transform, this corresponds to a signal space convolution. Having presented a linear shading equation which includes the effect of the max function, we now apply this result to shading volumes.

5.2 Spherical Harmonic Volumes

Volume data does not have explicit surfaces, however the gradient operator is often employed to enhance discontinuity which indicate surfaces, and to compute per-voxel surface normal vectors for shading. The normal at a volume coordinate (x, y, z) computed in this manner is given by:

$$N(x, y, z) = \frac{\nabla f(x, y, z)}{|\nabla f(x, y, z)|} = n_x \hat{i} + n_y \hat{j} + n_z \hat{k} \quad (24)$$

Equation 24 provides an orientation that may be used to evaluate Equation 23 for each voxel. Rather than utilizing the complex form of the spherical harmonic series, we employ the real form, as described in [1][8] for two reasons: images are real, and using real data allows a factor of two savings in memory and speed when performing the Fourier transform on the volume. Additionally, the

real spherical harmonic series may be expressed as products of unit surface normal vector components [8] obtained from Equation 24 so that computing and storing voxel normals in spherical coordinates is not necessary.

Expressing Equation 23 in terms of the real spherical harmonics which are functions of normal vector components, and the light in terms of a unit direction vector (x_L, y_L, z_L) results in a linear equation for shading the volume:

$$E(x, y, z) = \frac{|L|}{\pi} \sum_l \sum_m c'_{l0} Y_{lm}(x_L, y_L, z_L) Y_{lm}(n_x, n_y, n_z) \quad (25)$$

where $c'_{l0} = \sqrt{(4\pi)/(2l+1)} c_{l0}$. The 3D Fourier transform of the shaded volume may be directly computed as:

$$T(E) = \frac{|L|}{\pi} \sum_l \sum_m Y_{lm}(x_L, y_L, z_L) c'_{l0} T(Y_{lm}(n_x, n_y, n_z)) \quad (26)$$

Note that the terms in Equation 25 which undergo the Fourier transform are completely independent of light direction and surface reflection function. Therefore, the transform may be computed once for a given set of volume unit normal vectors and then used to shade the volume for any directional light source. The terms which change as a function of light direction may be quickly computed.

Shading the entire volume is inefficient since only a slice contributes to the final image. The individual harmonic Fourier terms for all l and m planes corresponding to the desired 2D slice may first be reconstructed and then used to produce the 2D shaded slice for inverse Fourier transforming. While Equation 26 allows shading in the Fourier domain, it does require volumes for each Y_{lm} used in the truncated spherical harmonic series. Fortunately, over 98 percent of the energy in the image may be obtained with only 9 spherical harmonic planes, for $l=0, 1, 2$ and $-l \leq m \leq l$ [1][8]. However, this is a significant burden in memory and reconstruction time. Possible optimizations and methods to avoid storing the full set of harmonics are described in the future work section to follow.

5.3 Results

We have again implemented the spherical (synthetic) dataset as well as the UNC brain from Section 4.4. The results are shown in Fig. 5 and Fig. 6.

For purposes of comparison, we have included the linear and cubic illumination techniques from Fig. 3 and Fig. 4. The images show, that the method based on spherical harmonics (bottom row) creates similar results to the cubic illumination (middle row). The overall appearances are greatly improved compared to the hemispherical illumination model. Moreover, the spherical harmonics method allows us to change the light source interactively. We first included an ambient term of 0.3 in all of our earlier visualizations. In Fig. 7 we have removed this influence. The accurate characterization of diffuse illumination operators results in higher quality images.

Whether we change the light source or the viewpoint - rendering times are the same. For the sphere (UNC brain) we achieve 21fps (5fps) for the spherical harmonics without depth cueing. Using depth cueing this decreases to 13fps (3fps). These frame-rates are what we expected. The Fourier Transform of ten (10) volumes are required in memory to implement spherical harmonic shading, cor-

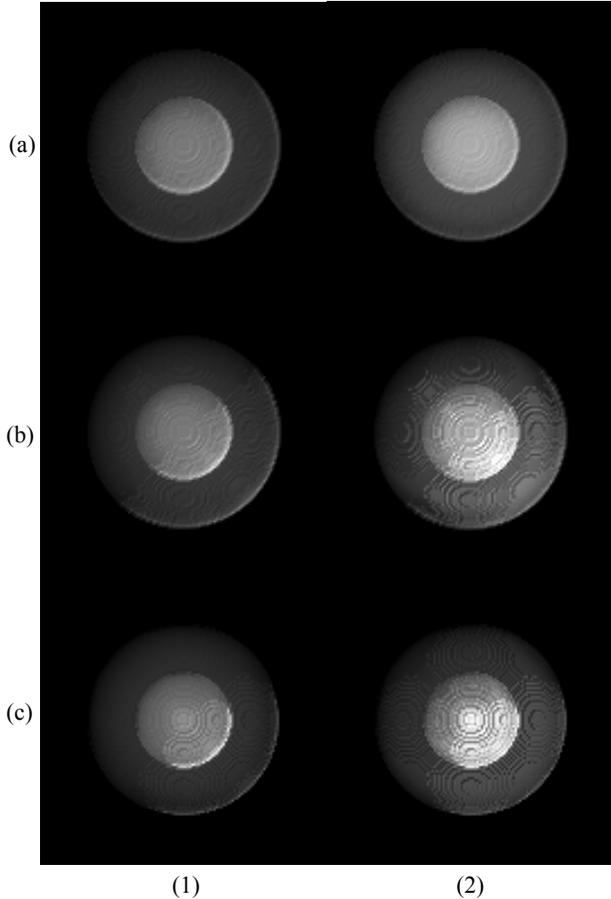


FIGURE 5. Rendering of a synthetic sphere. (a) using Totsuka and Levoy’s hemispherical reflections. (b) Our cubic shading term. (c) Our spherical harmonic implementation. Column (1) is without depth cueing. Column (2) includes depth cueing. In these images light is coming from (1,1,1)

responding to the nine $Y_{lm}(n_x, n_y, n_z)$ terms in Equation 26 for diffuse shading plus the data volume itself for ambient shading. Thus memory cost is significantly higher than FVR without shading and somewhat higher than the GCSH shading described above. Slice reconstruction within each of these volumes must also be performed, then multiplying by $L_{lm}c'_{l0}$ and summing to assemble one image. The other methods discussed only have to slice two volumes. Hence an increase on the order of five-fold in rendering time was expected. Further each floating point, complex Fourier volume needs 8 times the storage of the original byte-size data set.

6. CONCLUSIONS AND FUTURE WORK

The work presented in this paper attempts to make Fourier Volume Rendering (FVR) a more viable alternative for the visualization of volumetric datasets. It has long been shown that the superior complexity of $O(N^2 \log N)$ vs. $O(N^3)$ of all other rendering algorithms will have more of an impact as our data sets grow. Lighting is one important cue used to derive shape and hence insight from visualizations and in this paper we presented two ways for improved illumination in the Fourier Domain. One approach uses an improved hemispherical reflection term. We showed how this term, although not linear, can be implemented in the Fourier Domain. Further we have successfully employed spherical harmonics to cope with the

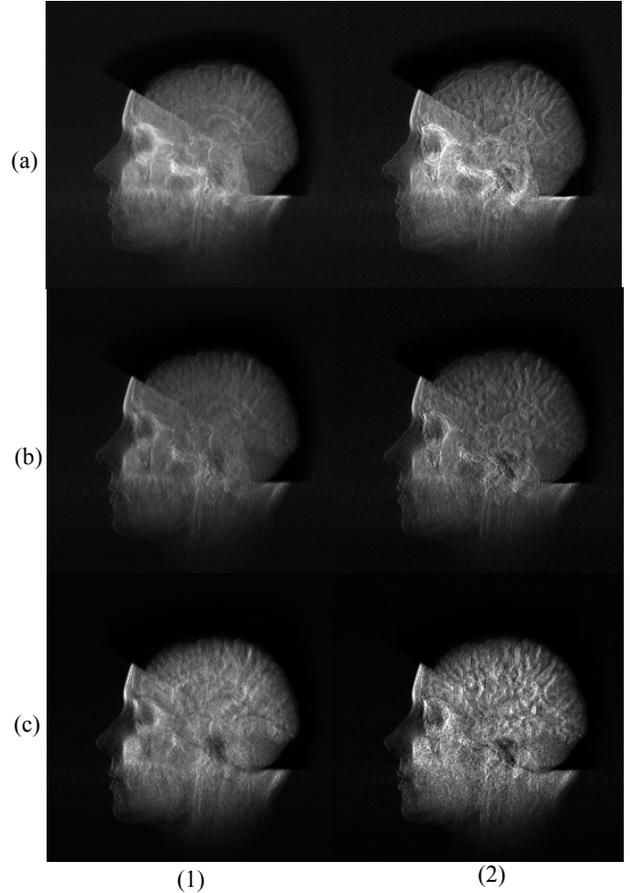


FIGURE 6. Rendering of the UNC brain data set. (a) using Totsuka and Levoy’s hemispherical reflections. (b) Our cubic shading term. (c) Our spherical harmonic implementation. Column (1) is without depth cueing. Column (2) includes depth cueing. In these images light is coming from left.

non-linearity of the shading operation in both the original and approximate versions. An efficient implementation led to real-time manipulation (21fps) of the data when both viewpoint and light source were altered. For constant light sources, we can use the first method, which allows a nearly five-fold improvement and hence rendering rates of 116fps.

Novel techniques for incorporating realistic shading into FVR visualization have overcome the visual differences inherent in hemispherical shading at the cost of increased computations and memory use. A valuable insight, from a systems viewpoint, would be obtained by performing comparison to other volume rendering methods, such as ray-casting. Due to time and space limitations, meaningful quantitative comparisons to other methods were not possible. However, we felt the significant contribution of this work is the description of the shading techniques. Future work in this area will develop optimized methods to reduce memory use and reconstruction time at which point cost comparisons will be useful and necessary.

Our algorithms, although showing real-time performance, can be further improved. Some ideas include the folding of 2 volumes into the real and imaginary part of the signal, which should afford one a speedup by a factor of two. Similar results can be achieved by applying the Hartley transform (the real version of the Fourier

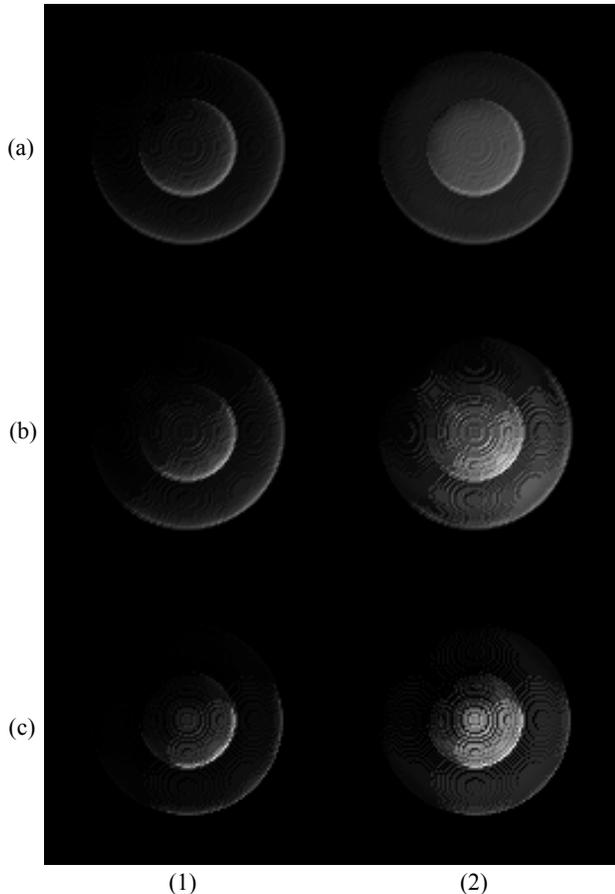


FIGURE 7. Rendering of a synthetic sphere without an ambient term. (a) using Totsuka and Levoy's hemispherical reflections. (b) Our cubic shading term. (c) Our spherical harmonic implementation. Column (1) is without depth cueing. Column (2) includes depth cueing. In these images light is coming from (1,1,1)

Transform). Other optimizations may be realized by manipulating the normal vector component terms that are the basis of the real spherical harmonic transform. Since these are derived from a gradient operation, evaluation of the partial derivative in Fourier space may allow fewer volumes to be maintained at the expense of multiple convolutions in the pre-processing step.

7. ACKNOWLEDGEMENTS

The second author wishes to acknowledge support from the NSF Engineering Research Center at Mississippi State University, and from the U.S. Army ERDC through the ILIR program. The fourth author wishes to acknowledge the support of the NSF through CAREER award (ACI 9734483).

The first and the third authors would like to thank the National Science and Engineering Research Council of Canada (NSERC) as well as the Advanced Systems Institute of British Columbia (ASI BC) for partial funding of this work

8. REFERENCES

- [1] R. Basri, D. Jacobs, "Lambertian reflectance and linear subspaces," *Proceedings of International Conference on Computer Vision'01*, pp. 383--390, 2001.
- [2] R.N. Bracewell, *Two Dimensional Imaging*, Prentice Hall Inc., Englewoods Cliffs, NJ, 1995.
- [3] B. Cabral, N. Max, and R. Springmeyer, "Bidirectional reflection functions from surface bump maps," *Computer Graphics.*, 21(4), pp. 273--281 1987.
- [4] M.H. Gross, L. Lippert, R. Dittich, S. Häring, "Two Methods for Wavelet-Based Volume Rendering", *Computers & Graphics*, 21(2), pp. 237-252, March 1997.
- [5] M. Levoy, "Volume rendering using the Fourier projection-slice theorem," *Graphics Interface '92*. pp. 61-69, 1992.
- [6] T. Malzbender, "Fourier Volume Rendering," *ACM Transactions on Graphics*, 12 (3), pp. 233-250, July 1993.
- [7] M. Meißner, J. Huang, D. Bartz, K. Mueller, R. Crawfis, "A Practical Comparison of Popular Volume Rendering Algorithms", *Proceedings of the 2000 Symposium on Volume Visualization*, pp. 81-90, October 2000.
- [8] R. Ramamoorthi, P. Hanrahan, "An Efficient Representation for Irradiance Environment Maps" *Proceedings of ACM SIGGRAPH 2001*, pp. 497-500, August 2001.
- [9] R. Ramamoorthi, P. Hanrahan, "On the Relationship Between Radiance and Irradiance: Determining the Illumination from Images of a Convex Lambertian Object," *J. Opt. Soc.*, Vol.18, No.10, October 2001.
- [10] F. X. Sillion, J. R. Arvo, S. H. Westin, D. P. Greenberg., "A Global Illumination Solution for General Reflectance Distributions," *Computer Graphics*, 25(4), pp. 187-196, July 1991.
- [11] Scoggins R., Machiraju R., Moorhead R., "Approximate Shading for Re-Illumination of Synthetic Images," *Proceedings of IEEE Visualization 2001*, San Diego, CA, 2001, pp. 379-386.
- [12] T. Totsuka, M. Levoy, "Frequency domain volume rendering", *Computer Graphics*, 27:(4), pp.271-278, August 1993.
- [13] M.A. Westenberg, J.B.T.M. Roerdink, "Frequency Domain Volume Rendering by the Wavelet X-ray Transform", *IEEE Transactions On Image Processing*, 9 (7), pp. 1249-1261, July 2000.
- [14] S. H. Westin, J. R. Arvo, K. E. Torrance, "Predicting Reflectance Functions From Complex Surfaces," *Computer Graphics*, 26 (2), pp. 255-264, July 1992.
- [15] T-T. Wong, P-A. Heng, S-H. Or, W-Y. Ng, "Image-based Rendering with Controllable Illumination," *Eurographics Rendering Workshop 1997*, pp. 13-22, June 1997.