Analytic Rendering of Multiple Scattering in Participating Media

Srinivasa G. Narasimhan, Ravi Ramamoorthi and Shree K. Nayar
Computer Science Department, Columbia University
New York, NY, USA
E-mail: {srinivas, ravir, nayar}@cs.columbia.edu

Additional Key Words and Phrases: multiple scattering, radiative transfer, atmospheric effects, participating media, physically-based glows, volume rendering

---

ABSTRACT

We consider the addition of physically-based weather effects like haze, fog and mist to images. Most computer graphics images are rendered under clear day or night conditions and little attention has been devoted to efficiently adding realistic weather effects. Multiple light scattering is dominant in a variety of atmospheric conditions (in general, most participating media) and is hard to model accurately using simple approximations such as single scattering and diffusion. In the past, accurate multiple scattering simulation has therefore required very expensive volumetric Monte Carlo methods. This paper focuses on multiple scattering from light sources immersed in participating media such as bad weather and fluids. We derive a new analytic formula for multiple scattering from a point light source in a medium. We extensively validate our model using monte carlo simulations as well as using controlled experiments with a light source immersed in milk. The model accurately predicts complex scattering effects produced by a wide range of concentrations of the medium. For rendering area sources of complex shapes and radiance distributions, we show that the analytic formula reduces to a simple depth-dependent convolution of the image. This leads to an efficient implementation of adding physically-based glows around sources. We demonstrate fast and accurate addition of weather effects to real photographs. Although the paper focuses on atmospheric effects, the results can be used to render sources through virtually any participating medium such as smoke, tissue and blood.

1. INTRODUCTION

Real-world imagery often includes weather effects like haze, fog and mist. However, most computer graphics images are rendered under clear atmospheric conditions, and simulation of weather effects has not received much attention. An important feature of images in bad weather is the glow around bright light sources, caused by multiple scattering in the atmosphere. The shape and extent of the glow depends on the particular type and density of particles constituting the weather condition. For instance, minute molecules in pure air do not produce glows, whereas larger particles such as aerosols, water droplets and impurities in haze, fog and mist can produce significant glows. While these effects are important in creating realistic images in bad weather, they are difficult to accurately and efficiently simulate using current computer graphics methods.

Volumetric Monte Carlo or finite element simulations, equivalent to volume ray tracing and radiosity, can give accurate results for general conditions and have been applied by a number of researchers [Pattanaik and Mudur 1993; Blasi et al. 1993; Languenou et al. 1994; Max 1994; Rushmeier and Torrance 1987; Antyufeev 2000]. These methods are based on numerically solving an integro-differential equation known as the radiative transfer equation [Chandrasekhar 1960; Ishimaru 1978], analogous in some ways to the rendering equation for surfaces. However, these simulations are very time consuming (unless the implementations are highly parallelized and hierarchical [Lecocq et al.]), leading us to look for alternative simple and efficient analytic models or approximations.

This paper presents a novel analytic solution for the scattering (including both single and multiple scattering) from a point source in a participating medium with arbitrary scattering properties (phase function), and a wide range of densities (optical thickness). The derivation is based on analyzing the spherical radiative transfer equations. To our knowledge, this represents the first such derivation of an analytic formula for general multiple scattering in the radiative transfer literature, generalizing previous work [Ambartsumian 1945; Marshak 1947; Elliott 1955; Chandrasekhar 1960]. Using our analytic formula, it is possible, for instance, to model the effects of small aerosols and large water drops constituting weather conditions like haze, fog and mist; the effects of impurities in smoke and water; as well as the effects of molecules and proteins in tissues, blood and milk. In simultaneous work, we have applied inverse methods to a simpler version of the model presented here to the computer vision problem of determining weather conditions from photographs [Narasimhan and Nayar 2003].

The closest previous analytic work in computer graphics is the application of the diffusion approximation for optically dense media by [Stam 1995]. Earlier, [Kajiya and Herzen 1984] simulated clouds by deriving an analytic formula in terms of spherical harmonics, but truncated the expansion after the first order harmonic, essentially obtaining a diffusion equation. Diffusion theory is based on the idea that multiple scattering within a dense medium eventually makes the angular radiance distribution almost uniform. Diffusion has been applied in plane parallel media such as clouds [Kajiya and Herzen 1984; Koenderink and van Doorn 2001] and subsurfaces of translucent materials [Jensen et al. 2001], where sources are outside the medium. However, it is not suitable for scattering in bad weather, because it cannot capture highly directional effects, where the sources and observer are both within the medium. Indeed, if applied to our problem, the glows around the light sources in figure 12 would fill the entire image, instead of having a finite extent with a rapid decay of intensity at large angles.

We validate our formula using monte carlo simulations as well as using real experiments with a spherical bowl of milk at different concentrations, that correspond closely to the conditions under which we derive our model. The experiments indicate that previous approximations like diffusion theory cannot fit the shape of the measured angular distributions, while our model gives accurate fits to within 3% with a small number of terms. Experiments on outdoor light sources (like street lamps) as well as experiments with a modified milk setup indicate that our model is reasonably accurate even in cases where the theoretical assumptions are not fully satisfied.

To efficiently apply our analysis to light sources of complex shapes and radiance distributions, we show that the glows can be computed by a depth-dependent image convolution. This makes implementation simple, requiring less than 25 lines of Matlab code. As a re-

sult, we have an interactive method to add physically correct glows to light sources seen through any participating medium. We demonstrate using three examples how our model can be applied to efficiently simulate weather effects in photographs. Further, we show that other approximations like single scattering [Ebert and Parent 1990; Sakas 1990; Max. 1986; E. Nakamae and Nishita 1986; Nishita and Nakamae 1987] qualitatively cannot capture the glows around light sources that are critical to hazy, foggy or misty appearances of the resulting images.

While this paper focuses on scattering of light from a source *within* a participating medium, we expect our model to be applied in the future to problems where the sources are *outside* the scattering media such as in the simulation of clouds [Kajiya and Herzen 1984; Nishita et al. 1996; Koenderink and van Doorn 2001] or subsurface scattering [Hanrahan and Krueger 1993; Dorsey et al. 1999; Pharr and Hanrahan 2000; Jensen et al. 2001]. We believe our analysis and formulae are likely to have broad impact in computer graphics applications ranging from rendering 3D models to image-based rendering, and in various fields like medical imaging, underwater imaging and digital entertainment.

The paper is organized as follows. Section 2 describes the analytic scattering model from a point light source in a spherical medium. In section 3, we describe the validation of our model using brute force monte carlo simulations. In section 4, we extensively validate the model using real experiments with milk. In section 5, we describe the effect of source visibility on multiple scattering and show that visibility is not a big problem in the presence of lamp covers, say, if we have a direct line of sight to the source. Sections 6 and 7 discuss issues related to rendering and present results of adding weather effects to photographs. Section 8 concludes with a description of future work.

2. ANALYTIC MULTIPLE SCATTERING MODEL

We present an analytical model to compute the light field (radiance at any spatial location in any direction) due to scattering from an isotropic point source, immersed in a homogeneous infinite scattering medium. This model includes the effects of both single and multiple scattering, which do not need to be handled separately. In section 7, we will extend this model to handle sources of complex shapes, sizes and radiance distributions.

Consider the homogeneous scattering medium shown in figure 1(a). An isotropic point source at the center illuminates the medium. The light field due to this source is said to exhibit *spherical symmetry*. In other words, the light field I in the medium only depends on the radial optical thickness $T = \sigma R$ (σ is the extinction coefficient of the medium and R is the radial distance from the source) and the angle $\theta = \cos^{-1}\mu$ from the radial direction. As we shall show, I can be expressed in terms of a Legendre polynomial series. We call the angular distribution of intensities $I(T,\mu)$ for any T as the 3D *scattering point spread function* (PSF). The main symbols used in this paper and their associated meanings are tabulated in Table I.

2.1 Model Derivation

We present the derivation of an analytic form for $I(T,\mu)$ in this subsection. The derivation is rather lengthy and involved. So, readers interested mainly in the implications of our model to graphics can skip this subsection.

We follow the physics-based theory of *radiative transfer* [Chandrasekhar 1960; Ishimaru 1978] for modeling light scattering within a spherical medium. The key idea in this approach is to investigate the difference between light incident on, and exiting from, an

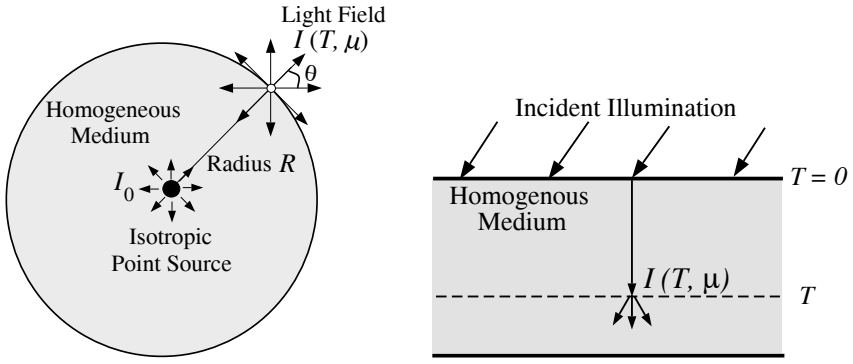

(a) Spherical Medium (b) Plane Parallel Medium

Fig. 1. *(a) An isotropic point source illuminating a spherical homogeneous scattering medium. The multiple scattered light field in the medium depends on radial optical thickness $T = \sigma R$ and the angle θ from the radial direction. The light field is said to exhibit spherical symmetry. (b) Plane parallel model used by most previous works to describe multiple scattering and diffusion. The source is generally collimated and is outside (infinitely far away from) the medium. Such a model can be used for scattering in clouds or subsurface scattering, for instance, but not in our case where the source is divergent and is within the participating medium.*

Symbol	Meaning
I	Light Field
I_0	Radiant Intensity of Isotropic Point Source
R	Radial distance in Spherical Medium
σ	Extinction Coefficient
$T = \sigma R$	Optical Thickness
W_0	Single scattering Albedo
P	Phase Function of Medium
L_k	Legendre Polynomial of order k
$\mu = \cos\theta$	Cosine of angle θ from the radial direction
q	Henyey Greenstein Phase function Parameter
c_m	Model Constants

Table I. *A table denoting the main symbols and their associated meanings.*

infinitesimal volume of the medium. Mathematically, the change in flux through a small volume is given by an integro-differential equation, called the *Radiative Transfer Equation (RTE)*. The directional intensity at any location in the atmosphere is then obtained by solving this equation. We begin by writing the spherical radiative transfer equation (RTE) [Chandrasekhar 1960] :

$$\mu \frac{\partial I}{\partial T} + \frac{1-\mu^2}{T}\frac{\partial I}{\partial \mu} = -I(T,\mu) + \frac{1}{4\pi}\int_0^{2\pi}\int_{-1}^{+1} P(\cos\alpha)\, I(T,\mu')\, d\mu'\, d\phi' \tag{1}$$

$$\cos\alpha = \mu\mu' + \sqrt{(1-\mu^2)(1-\mu'^2)}\cos(\phi-\phi')\,. \tag{2}$$

$P(\cos\alpha)$ is called the *phase function* of the particles in the medium and it denotes the angular distribution of light scattered by any particle. α is the angle between the incom-

ing irradiance direction (θ', ϕ') and the outgoing scattered radiance direction (θ, ϕ). For simplicity, the angles are written in terms of their cosines as $\mu = \cos\theta$, $\mu' = \cos\theta'$.

STEP 1: ELIMINATING PARTIAL DERIVATIVE $\frac{\partial I}{\partial \mu}$

In the first step, we eliminate the partial derivative of I with respect to μ by integrating the RTE with respect to μ over the range $[-1, +1]$. Most of the details of this step are adapted from [Chandrasekhar 1960]. However, a key difference in our approach to eliminating the partial derivative is to use exact integrals instead of approximating the integrals by linear summations [Chandrasekhar 1960].

By integrating P over the azimuth angle ϕ', Ishimaru and Chandrasekhar define a function that does not depend on the azimuth angle as,

$$P^{(0)}(\mu, \mu') = \frac{1}{2\pi} \int_0^{2\pi} P(\cos\alpha) d\phi' . \tag{3}$$

As we shall see, the use of $P^{(0)}$ simplifies the mathematics involved in modeling the multiple scattering around a point light source. Substituting into equation 1 we get,

$$\mu \frac{\partial I}{\partial T} + \frac{1-\mu^2}{T} \frac{\partial I}{\partial \mu} = -I(T, \mu) + \frac{1}{2} \int_{-1}^{+1} P^{(0)}(\mu, \mu') I(T, \mu') d\mu' . \tag{4}$$

Chandrasekhar defines a function $Q_m(\mu)$, for some $m > 0$, such that

$$L_m(\mu) = -\frac{d((1-\mu^2) Q_m(\mu))}{d\mu} \iff Q_m(\mu) = \frac{L'_m(\mu)}{m(m+1)} \tag{5}$$

Consider the integral

$$\int_{-1}^{+1} (1-\mu^2) Q_m \left(\frac{\partial I}{\partial \mu} \right) d\mu . \tag{6}$$

Integrating by parts and using 5, it has been shown that we can eliminate the partial derivative with respect to μ (see [Chandrasekhar 1960]),

$$\int_{-1}^{+1} (1-\mu^2) Q_m \left(\frac{\partial I}{\partial \mu} \right) d\mu = \int_{-1}^{+1} I(T, \mu) L_m d\mu . \tag{7}$$

When there is no confusion, we drop the parameters μ and T for brevity. Multiplying 4 by Q_m and integrating with respect to μ over $[-1, +1]$, we get,

$$\begin{aligned} &\int_{-1}^{+1} \mu Q_m \frac{\partial I}{\partial T} d\mu + \int_{-1}^{+1} \frac{1-\mu^2}{T} Q_m \frac{\partial I}{\partial \mu} d\mu \\ &= -\int_{-1}^{+1} Q_m I \, d\mu + \frac{1}{2} \int_{-1}^{+1} Q_m d\mu \int_{-1}^{+1} P^{(0)}(\mu, \mu') I(T, \mu') d\mu' , \end{aligned} \tag{8}$$

Substituting equation 7, we can rewrite the RTE as,

$$\int_{-1}^{+1} \mu Q_m \frac{\partial I}{\partial T} d\mu + \frac{1}{T} \int_{-1}^{+1} L_m I \, d\mu$$
$$= -\int_{-1}^{+1} Q_m I d\mu \; + \frac{1}{2} \int_{-1}^{+1} Q_m(\mu) d\mu \int_{-1}^{+1} P^{(0)}(\mu, \mu') I(T, \mu') d\mu' , \tag{9}$$

STEP 2: LEGENDRE POLYNOMIALS FOR $P(\cos\alpha)$

Most phase functions (both isotropic and anisotropic) can be expanded using Legendre polynomials L_k [Chandrasekhar 1960]:

$$P(\cos\alpha) \; = \; \sum_k W_k L_k(\cos\alpha) . \tag{10}$$

As an example, for the popular Henyey-Greenstein phase function [Henyey and Greenstein 1941] (for purely scattering media),

$$P(\cos\alpha) = \frac{1-q^2}{(1+q^2-2q\cos\alpha)^{3/2}} . \tag{11}$$

where, $q \in [0,1]$ is called the *forward scattering* parameter. Here, the coefficients of the Legendre polynomial expansion are shown to be:

$$W_k \; = \; (2k+1)\, q^k , \quad k \geq 1 , \tag{12}$$

Note that our derivation is not limited to any particular phase function. In other words, we may simply use different coefficients W_k for different phase functions. By using (2), it has been shown that [Ishimaru 1978]

$$P^{(0)}(\mu, \mu') = \sum_{k=0}^{\infty} W_k L_k(\mu) L_k(\mu') . \tag{13}$$

As we shall show in the appendix A, this expansion of $P^{(0)}$ greatly simplifies the mathematics in our derivation.

STEP 3: LEGENDRE POLYNOMIALS FOR $I(T, \mu)$

This is the key step in our derivation. We assume that a particular solution to the RTE is of the separable form:

$$I_m(T, \mu) = g_m(T) f_m(\mu) . \tag{14}$$

Since I_m does not depend on the azimuth angle ϕ, the angular dependence f_m can be represented using Legendre polynomials. Note that previous methods [Marshak 1947; Elliott 1955] that solve the RTE for isotropic phase functions have used spherical harmonics ex-

pansion of I_m which is not necessary for spherical symmetry. Substituting into 9, we get,

$$g'_m \int_{-1}^{+1} \mu Q_m f_m d\mu + \frac{g_m}{T} \int_{-1}^{+1} L_m f_m d\mu + g_m \int_{-1}^{+1} Q_m f_m d\mu - \frac{g_m}{2} \int_{-1}^{+1} Q_m d\mu \int_{-1}^{+1} P^{(0)}(\mu,\mu') f_m(\mu') d\mu' = 0\,. \tag{15}$$

It turns out that just equating f_m to the Legendre polynomial L_m does not produce meaningful non-trivial solutions to the RTE. Instead a simple extension where f_m is written as the sum of two consecutive Legendre polynomials allows us to compute meaningful and non-trivial solutions to the RTE:

$$\boxed{f_m(\mu) = L_{m-1}(\mu) + L_m(\mu)\,.} \tag{16}$$

Similarly, we shall expand $L'_k(\mu)$ and $\mu L'_k(\mu)$ using Legendre polynomial series [MacRobert 1967]:

$$\begin{aligned} L'_k(\mu) &= (2k-1)L_{k-1}(\mu) + (2k-5)L_{k-3}(\mu) + \ldots \\ \mu L'_k(\mu) &= kL_k(\mu) + (2k-3)L_{k-2}(\mu) + \ldots\,. \end{aligned} \tag{17}$$

We exploit the orthogonality of Legendre polynomials that appear in the f_m, $P(\cos\alpha)$ and $P^{(0)}$ terms to simplify the equation in 15. We substitute equations 5, 16, 13 and 17, into equation 15 and simplify each term using the orthogonality of Legendre polynomials:

$$\int_{-1}^{+1} L_i L_j d\mu = \begin{cases} \frac{2}{2n+1} & \text{if } i = j = n; \\ 0 & \text{otherwise} \end{cases} \tag{18}$$

The above property can be used to greatly simplify the mathematics and only the main results are listed below.

—Term 1 in Equation (15):

$$g'_m \int_{-1}^{+1} \mu Q_m f_m d\mu = g'_m \left(\frac{2}{(m+1)(2m+1)} \right). \tag{19}$$

—Term 2 in Equation (15)

$$\frac{g_m}{T} \int_{-1}^{+1} L_m f_m d\mu = \frac{g_m}{T} \left(\frac{2}{2m+1} \right). \tag{20}$$

—Term 3 in Equation (15)

$$g_m \int_{-1}^{+1} Q_m f_m d\mu = \frac{2g_m}{m(m+1)}\,. \tag{21}$$

—Term 4 in Equation (15)

$$\frac{g_m}{2}\int_{-1}^{+1} Q_m d\mu \int_{-1}^{+1} P^{(0)}(\mu,\mu')\, f_m(\mu') d\mu' = \frac{2g_m}{m(m+1)}\left(\frac{W_{m-1}}{2m-1}\right). \tag{22}$$

The fact that $P^{(0)}$ can be expressed as products of Legendre polynomials is clearly a considerable advantage in simplifying the terms of equation 15. Substituting the 4 terms from equations (19, 20, 21, 22) into equation (15), we obtain

$$g_m' + \frac{g_m}{T}\alpha_m + g_m\beta_m = 0 \tag{23}$$

The solution to (23) is,

$$g_m(T) = I_0 e^{-\beta_m T - \alpha_m \log T}, \tag{24}$$

where the following form for g_m along with the above equation for f_m satisfies the RTE:

$$\boxed{\begin{aligned} g_m(T) &= I_0\, e^{-\beta_m T - \alpha_m \log T} \\ \alpha_m &= m+1 \\ \beta_m &= \frac{2m+1}{m}\left(1 - \frac{W_{m-1}}{2m-1}\right), \end{aligned}} \tag{25}$$

where, I_0 is the radiant intensity of the point source. Note that the function $g_m(T)$ captures the attenuation of light in the medium, whereas the Legendre polynomial $L_m(\mu)$ explains the angular spread of the brightness observed due to multiple scattering. Note that the above equation automatically satisfies the boundary condition: $g_m(\infty) = 0$.

STEP 4: SUPERPOSING SOLUTIONS

In the final step, we simply superpose solutions corresponding to different values of m to get the final solution.

$$I(T,\mu) = \sum_{m=1}^{\infty} c_m g_m(T)\,\left(L_{m-1}(\mu) + L_m(\mu)\right). \tag{26}$$

For our analysis, we use the solution to the RTE given by all $c_m = 1$ and we show this produces accurate results in the paper.

Hence, we rewrite the above equation as,

$$\boxed{I(T,\mu) = \sum_{m=0}^{\infty}\left(g_m(T) + g_{m+1}(T)\right) L_m(\mu),} \tag{27}$$

with $g_0 = 0$. Note that this series solution is valid (converges) only for $T > 1$. However, for cases with small q, multiple scattering is minimal for $T \leq 1$, and we can simply use a single scattering approximation in that domain.

2.2 Notes on the Derivation and the Solution

We present in this section some observations about our derivation of the model as well as some features of the model itself.

2.2.1 *Features of the Model.* The model is valid for both **isotropic and anisotropic scattering** and thus describes multiple scattering within several participating media. For instance, in the Henyey Greenstein phase function, the parameter $q = 0$ corresponds to isotropic scattering, and $0 \leq |q| \leq 1$ corresponds to anisotropic scattering.

The model is valid for both **absorbing and purely scattering media**. The zeroth coefficient of the phase function, W_0, called the *single scattering albedo*, denotes the ratio of the scattering coefficient to the extinction (scattering + absorption) coefficient. When $W_0 = 1$, there is no absorption of light and the medium is said to be purely scattering. The Henyey-Greenstein phase function is modified to handle absorption by multiplying equations 11 and 12 by the factor W_0. Note that monte-carlo or other numerical simulation methods for scattering do not converge for the case of purely scattering (infinite) media. In the rest of the paper, wherever not explicitly mentioned, W_0 is assumed to be 1.

2.2.2 *Separating Reduced and Diffuse Intensity.* Note that a separation of total intensity into a sum of reduced intensity and diffuse scattered intensity is sometimes done in solving certain RTEs. See for example, the problems of "Diffuse Reflection and Transmission" based on the Plane parallel RTE. In this case, it has been shown by Chandrasekhar [Chandrasekhar 1960], Ishimaru [Ishimaru 1978] and Antyufeev [Antyufeev 2000] that a separation helps in solving the plane parallel RTE. However, the same texts mentioned before do not use this separation while describing spherical RTEs which is the focus of our work. In our derivation, we do not separate the reduced and diffuse intensities. Instead, we obtain an expression for the total intensity (reduced + diffuse) directly from the Spherical RTE.

2.2.3 *Normalizing the Solution.* A simple glance at the solution suggests that the flux can be arbitrarily large due to $T^2 = \sigma^2 R^2$ being in the denominator. However, this can be easily fixed by using the constants $c_m = \sigma^2$ instead of $c_m = 1$. In this way, I_0/R^2 can be factored out of the PSF. However, in many cases, where the exact values of I_0 and σ may not be known, the above issue can be handled by normalizing the maximum to be 1. This not only factors out I_0/R^2 but also removes the factor $c_m = \sigma^2$. Thus, effectively, $c_m = 1$ is still valid.

2.2.4 *Choice of Particular Solution $g_m(T)f_m(\mu)$.* Note that our particular solution is a separable product of two functions - $g_m(T)$, which is only a function of optical thickness T; and $f_m(\mu)$, which is only a function of cosine μ of the angle (see equation 14). At first glance, this form suggests that the angular distribution of the PSF independent of the optical thickness T. This is incorrect, since the above separable formulation is only a single term in the entire solution. The complete solution in fact is a DOT product of g_m terms and f_m terms for $m \geq 0$. Thus, the final form of the solution is such that the angular distribution of the PSF *does* depend on the optical thickness T. We wish to bring to the readers's attention to prior papers (for instance, [Ambartsumian 1945]) that have followed this approach of using separable functions in solving other RTEs and have shown this approach to be more accurate for larger values of T.

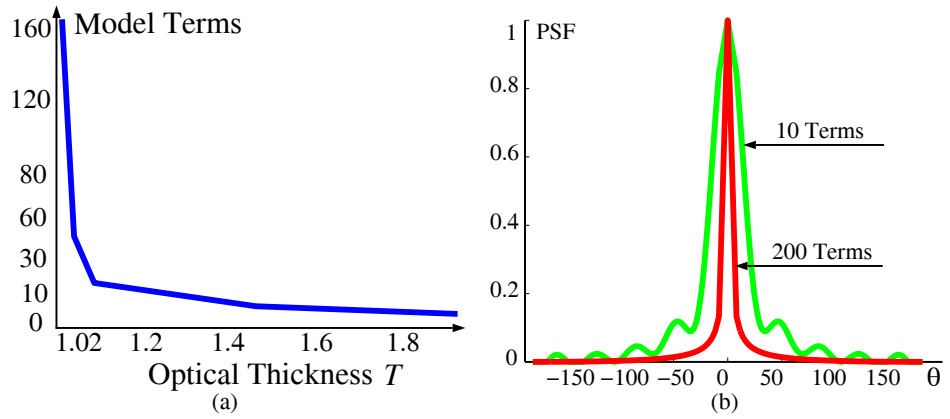

Fig. 2. *Number of coefficients needed in the model. Sharply peaked PSFs (small T) require a large number of terms in the series expansion (27), whereas less than* 10 *terms are needed for wide PSFs (large T). (b) Using small number of terms for narrow or peaked PSFs ($q = 0.9$, $T = 1.02$) produces ringing effects.*

2.2.5 *Number of Terms and Convergence.* The separability of the particular solution is precisely the reason for our solution to *not* converge for $T \leq 1$. Here we present a discussion about the convergence of our solution. Note that for $T \leq 1$, the attenuation function g_m does not decrease with each term m and hence the infinite series will not converge. However, for $T \geq 1$, it is simple to observe that each successive term g_m progressively attenuates the angular term L_m (which has an absolute upper bound of 1) and eventually reaches 0. For a particular error tolerance, larger the value of T, the faster the convergence. In the case of $T \leq 1$, the solution produces a delta function. However, the magnitude of the delta does not converge to a finite value. We must also analyze the convergence of the solution for different values of q for fixed T. Note that for large q (close to 1), the PSF should include more terms to be accurate and thus shows slow convergence. On the other hand, for smaller q the convergence is faster.

The **number of terms** needed in the model (eq. 27) depends on the values of T and q as well as the precision required. Note that g_m decreases rapidly for large optical thickness T or small q. Hence only a few terms (for $T > 2$, $m < 10$) are required for sufficient accuracy (less than 0.1 percent deviation from 500 terms). Higher order terms are needed when the optical thickness T or q is close to 1 (for $T \approx 1.02$, $m \approx 200$), and the medium exhibits significant anisotropy. Figure 2 gives the number of terms for an accurate approximation for different values of T.

2.2.6 *Initial/Boundary Conditions and the choice of $c_m = 1$.* There are several reasons for the particular choice of coefficients $c_m = 1$. Imagine that the constant of integration for a particular solution is k_m (see equation 24). Thus, different particular solutions may have different k_ms. However, note that the spherical RTE represents scattering from an isotropic point source. Hence, the complete solution must agree with the isotropy of the point source. In this sense, our derivation must satisfy an initial condition that the radiant intensity I_0 of the point source be the same for *all angles*. Thus, $k_m = I_0$ used as an integration constant for *a* particular solution must be the same for *all* particular solutions. This results in the radiant intensity being a global scale factor to the PSF as required. Thus when the particular solutions are combined, the choice of $c_m = 1$ allows for this condition.

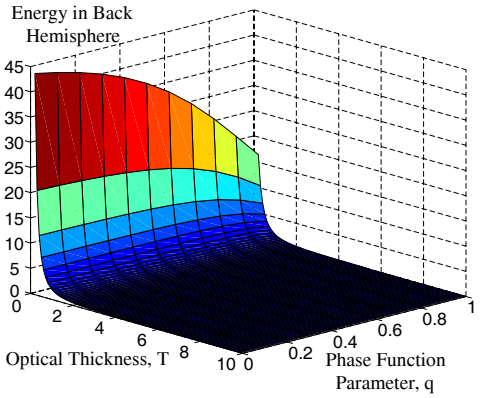

Fig. 3. Plot showing the total energy in the back hemisphere ($\mu <= 0$) of the multiple scattered intensity for a range of optical thicknesses and phase function parameter values. Theoretically, the backscattered diffuse energy goes to zero at infinity. However, the plot shows that the model can be used for media with finite extent if the optical thickness is reasonably large. A common boundary condition enforced is that there be no inward flux at the boundary of the medium. Since the back hemisphere energy is low for large optical thicknesses, this boundary condition is accurately satisfied by the model.

Furthermore, the choice of $c_m = 1$ must also satisfy boundary conditions. A desirable boundary conditions for forward scattering media is that the backscattered energy ($I(T,\mu)$ for $\mu <= 0$) is zero at the boundary of the medium. Note that the solution trivially satisfies this condition for an infinite medium. The plot in figure 3 shows the decay of the total energy in the back hemisphere ($\mu <= 0$) with distance. Theoretically, the backscattered diffuse energy goes to zero at infinity. However, the plot shows that the model can be used for media with finite extent if the optical thickness is reasonably large and is also evidenced by our real experiments (see next section). Note also that $I(T,-1) = 0$ giving a realistic PSF that is maximum for $\mu = 1$, stays positive and decreases rapidly to 0 at $\mu = -1$. This makes the angular profile a true point spread function, and avoids scattered light from completely filling the screen (washout), which is not the case in reality (for eg., mild fog). We have also shown that this choice is accurate through real experiments as well as numerical Monte Carlo simulations. See the next section on validation. Further boundary conditions may have to be applied depending on the problem at hand (for instance, handling visibility) and this is a topic of future work.

2.2.7 *Relation to Diffusion.* The popular diffusion model [Stam 1995; Koenderink and van Doorn 2001; Jensen et al. 2001] for highly dense media is simply a 2 term approximation to the full solution of the RTE. In other words, the angular distribution of diffusion is linear, i.e., $aL_0(\mu) + bL_1(\mu) = a + b\mu$. We will show that diffusion is inherently incapable of representing multiple scattering in media of various densities and that our model with higher order terms is much more accurate. An important distinction to be made between our model and the diffusion approximation is that our model is derived using the RTE for spherically symmetric media (figure 1(a)) whereas the popular diffusion model is derived from the RTE for plane parallel media (figure 1(b)). Hence, the forms and the boundary conditions are somewhat different from our model.

2.2.8 *Wavelength Dependence:*. Note that in general the parameters q and T of our model will vary with the wavelength of light radiated by the source. Therefore, if we are interested in multiple wavelengths (typically, 3 in the case of color images), then the light field due to multiple scattering can be computed from our model using the corresponding sets of values for the parameters q and T. The values of q and T for different wavelengths have been estimated and tabulated in the atmospheric optics literature [Kopeika 1981] and can be directly used with our model.

3. COMPARISON WITH MONTE CARLO

In this section, we compare our model to the results obtained using the standard Monte Carlo technique for simulating scattering. We use a brute force monte carlo renderer (Dali) [Jensen]. As expected, the Monte Carlo technique is very slow. In our experience, the simulations took anywhere between a few hours to days to obtain accurate results. The simulations were run in parallel on a 32 machine Pentium III PC cluster each with 256MB RAM. The large variation in the timing is due to the fact that monte carlo simulations are strongly effected by the parameters used to describe the participating medium (W_0, T, q). For instance, consider the single scattering albedo W_0. Note that monte carlo does not converge for pure scattering infinite media ($W_0 = 1$). For $W_0 < 1$, smaller the single scattering albedo faster the convergence. Similarly, the larger the optical thickness (T), the longer Monte Carlo takes to converge.

In figure 4, we show two plots obtained using monte carlo simulations of scattering and the corresponding PSFs obtained using our analytic model. Both the model as well as the corresponding Monte Carlo simulation were run on the same set of input parameters (W_0, T, q). The number of photon samples used in the Monte Carlo simulation was 250,000 with a maximum of 150 scattering events (bounces) simulated for each photon. The angles ($\cos^{-1}\mu$) were computed using an orthographic mapping (see figure 10). Thus our analytic model can be used accurately as well as instantaneously irrespective of the values of the parameter set describing the participating media. Due to the large simulation time required for the Monte Carlo technique, we validate our model more extensively using real experiments with milk described in the next section.

4. VALIDATION USING EXPERIMENTS WITH MILK

In this section, we describe the extensive experiments we performed to validate our model. The main challenge here is to carefully design the experiments under controlled settings that conform to the theory as closely as possible. We chose the scattering medium to be milk since it strongly scatters light and hence the multiple scattering effects are pronounced. Furthermore, we can dilute milk with water to create different optical thicknesses. To accurately simulate a spherical medium, we constructed a 40 cm diameter spherical container made of transparent plastic (Lucite) and filled it with milk. The apparatus used for the experiments is shown in figure 5. At the center of the container is a small spherical frosted glass bulb. We acquired images of the bulb from various angles and found that the bulb emits light more or less uniformly in all directions (except the narrow back region of the bulb holder) making it essentially isotropic. At the top of the container is a flat opening used to fill milk, to insert the light source and also to clean the bowl. Note that this is a fundamentally different experiment from that done by [Jensen et al. 2001]. We have the source *inside* a large *spherical* container, while [Jensen et al. 2001] have the

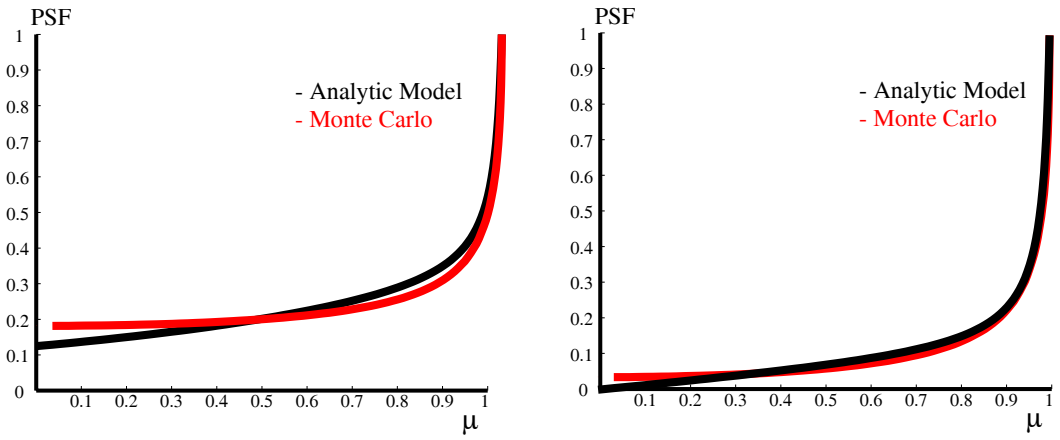

Fig. 4. *Comparison with Monte Carlo Simulations. Two plots show the PSFs obtained using a brute force Monte Carlo renderer (Dali) and our analytic model. While Monte Carlo takes a long time to converge (8 hrs on a 32 PC cluster), our model can be computed in real-time.*

source and observer *outside* a *semi-infinite plane-parallel* medium. We acquired images of the apparatus with a radiometrically calibrated high dynamic range (12 bits per pixel) Kodak DCS 760 camera. Fifteen different concentrations of milk obtained by diluting milk with different amounts of water were used in our experiments.

A mapping between the observed intensities in the image domain and the 3D scattering PSF is achieved using the geometry of figure 5(b). The mapping between the angle γ and μ is obtained by using simple trigonometry:

$$\frac{D}{\sin(\pi - \cos^{-1}\mu)} = \frac{R}{\sin\gamma}\ , \quad \Rightarrow \quad \mu = \sqrt{1 - \frac{D^2}{R^2}\sin^2\gamma}\ , \tag{28}$$

where, R is the radius of the spherical bowl and D is the distance from the pinhole to the center of the bowl. We must also handle refraction at the milk - air interface. We assume that that the refraction at the milk - glass and glass - air interfaces are minimal since the glass is very thin. The refractive index of milk was assumed to be 1.4 [Nallasamy 1984]. We do not handle inter-reflections or total internal reflections. We further assume that the fresnel coefficient of refraction is constant for all incident angles. This is generally true for most angles except the grazing angles. Then, we use Snell's Law to map the refracted angle μ in figure 5 to the incident angle μ_i as:

$$\frac{\sin\mu_i}{\sin\mu} = 1.4\,, \tag{29}$$

The angle μ_i is then used to fit our model. We verified the accuracy of our model by fitting the best PSF to the observed PSF (radial profile) for each image. The PSF model in eq. 27 has only three parameters W_0, T and q. We fit the observed PSF using a simple non-linear search tool in Matlab. The parameters of the PSF generally depend on wavelength of incident light. Ideally, we must constrain the model parameters according to their wavelength dependencies. However, since we lack this information in this experiment, we fit the model separately for each color channel.

A small representative set of 4 images (out of 15 milk concentrations) with the corresponding measured PSFs and the computed model fits is shown in figure 6. We see that in

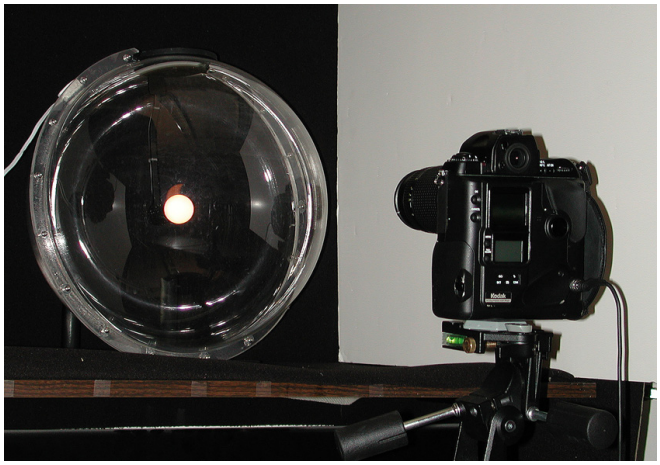

(a) Experimental Setup

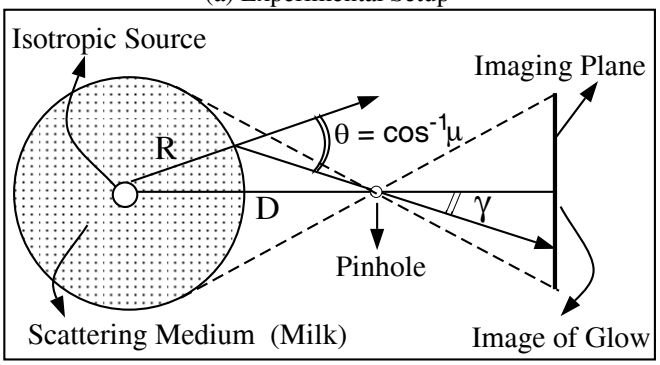

(b) Measurement Geometry

Fig. 5. *(a) Apparatus for measuring scattering within milk. On the left, a small bulb is positioned at the center of a spherical container made out of Lucite plastic. During experimentation, this container was filled with milk and placed in a black enclosure to avoid reflections from outside the container. The white background is shown only for clarity. The inter-reflections visible in the image were negligible when the container was filled with milk (see figure 6(a)). The small glass bulb is frosted which makes it roughly diffuse/isotropic. (b) The geometry shown on the right is used to measure multiple scattered intensities in different directions. The mapping between the image ray (angle γ) and the scattered ray (angle θ) is given in eq. 28.*

all cases, our model produces very accurate results (to within 3%). The parameters W_0 and q do not vary with concentration and hence we forced them to be constant for all the 15 experiments while fitting. This makes the validation much more strong. For comparison, we also empirically fit a 2-term diffusion-type angular dependence, i.e. $a + b\mu$ (the first two terms of our model). The diffusion approximation gives large errors ($20\% - 50\%$) in the left three plots. In fact, the real data is concave, while a function of the form $a + b\mu$ is only linear in μ. This clearly shows that we need higher order terms, and a diffusion-like model does not have the flexibility to capture the shapes of the angular distributions of interest. For dense media (large T), shown in the rightmost image, it is conceivable that a diffusion-type model could be fit to the data, although we again emphasize that our analytic formula is different from diffusion as we are solving a different problem (spherical RTE and not plane parallel RTE). These results show that our model works for a range of densities.

Note that we try to validate the model by fitting our model to observed milk data. We wish to clarify that we were not able to obtain the precise values for q, T and W_0 for the milk experiments we performed. Nallasamy [Nallasamy 1984] presents the scattering coefficient σ and forward scattering parameter q as a function of the sizes of the fat molecules in milk. Since we do not have information about the sizes of fat molecules used in our experiments, it is not possible to obtain the right scattering coefficients. Furthermore, the above paper does not give information about the absorption coefficient. Hence, we could not obtain the scattering albedo W_0 also.

In order to validate the model in a more strict sense, we performed the fitting in the following manner. First, we forced the parameters W_0 and q to be constant for all the 15 concentrations. Then, we searched for the best fit optical thickness T for each concentration. The values we obtained after fitting were approximately $W_0 = 0.8$, $q = 0.85$, and $T = \{1.1, 1.2, 1, 3, 1.5, 1.7, 1.9, 2.2, 2.3, 2.5, 2.8, 3.0, 3.2, 3.5, 3.8, 4.5\}$. The search space was $W_0 \in (0,1)$ in steps of 0.1; $q \in (0,1)$ in steps of 0.1 and $T \in (1,10)$ in steps of 0.05. The good fits suggest that the parameters are consistent with observations. The values shown are for the red channel. Similar values were obtained for the green and blue channels respectively. Note that we do not know how to constrain the relative ratios of parameters of the different color channels and hence we fit the model independently to each channel.

5. ACCURACY OF MODEL WITH REAL OUTDOOR LIGHT SOURCE

We also validated the model using a real light source in the outdoors. Imagine an outdoor light source in foggy conditions. The commonly appearing glow around the source is due to multiple scattering. Our model can be used to describe the glow around a light source in bad weather. We verified this model using images of distant sources and their glows with a high dynamic range (12-bits per pixel) Kodak digital camera. Weather data from a weather website was obtained at the time of image acquisition (rain, $q \approx 0.95$, 2.25 miles visibility). The source was about 1 km away from the sensor. The PSF measured from the image of the glow of one of the sources, and the PSF computed using our model are shown in figure 7. The angle $\cos^{-1}\mu$ is measured from the radial direction originating from the source. The comparison between the measured and computed PSFs shows the accuracy of our model.

6. EFFECT OF SOURCE VISIBILITY ON MULTIPLE SCATTERING

In section 2, the model was derived for an isotropic point source in open space without any occluders in the medium. In the previous section, we validated our model under controlled settings (spherical container, isotropic source, no visibility issues) that suited the theory. In this section, we describe the effect of blocking out different parts of the light source, on the observed scattering PSF.

If we have a direct line-of-sight to the source, occluders at large angles to the line-of-sight (including back covers of lamps) cause little problem since the scattering PSF falls off rapidly at large angles. To validate this claim, we conducted several carefully designed experiments using the setup in figure 5 to demonstrate the effect of source visibility on our model. In other words, we block different solid angles of the source and measure the scattered light field for several milk concentrations. The top row in figure 8 shows the various visibility configurations used in our experiments. To block different solid angles of the light bulb, we constructed a tube made of several filter adapter rings threaded together.

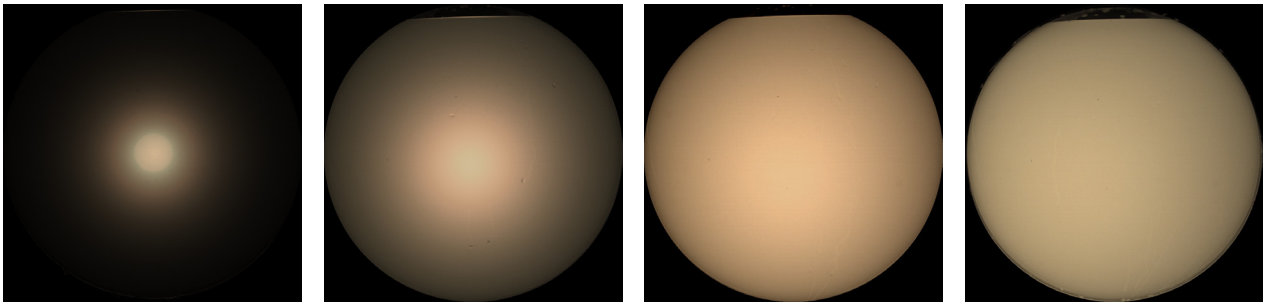

(a) High dynamic range photographs captured with different milk concentrations (increasing from left to right).

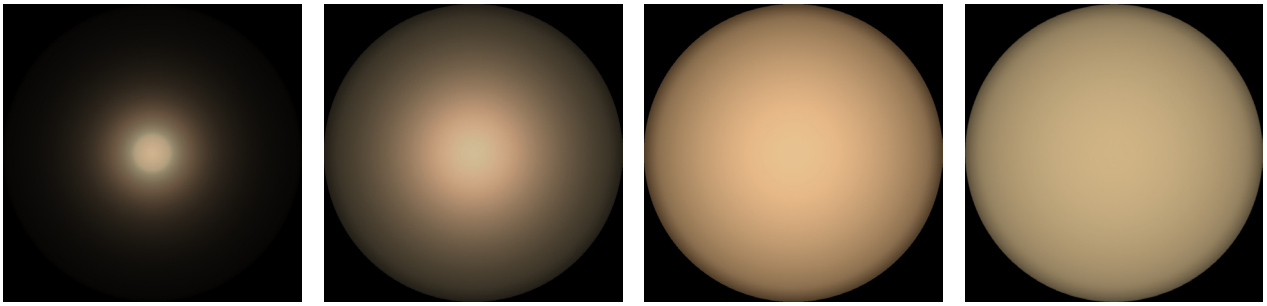

(b) Images rendered using our model.

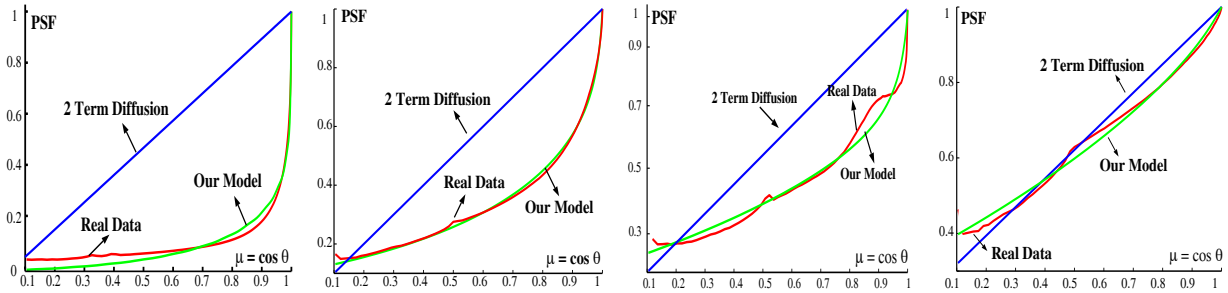

(c) Fits to the real data (red) with a 2-term diffusion like model (blue) and our model (green).

Fig. 6. *Validation with real data. (a) 4 out of 15 images of different concentrations of milk in the spherical container. Note that the definition of the source decreases as milk concentration increases (from left to right), showing that the PSFs are wider. The top parts of the images are flat since the container has a flat opening large enough to fill milk, insert the light source and also to clean the bowl. (b) Images rendered using our model in eq. 27. Some of the color variations are due to images being displayed at different radiance scales. (c) Plots of real radial angular profiles for selected experiments as well as fits with our model and an empirical 2-term diffusion-like model. The fits were performed for each color channel independently (for brevity, only the red channel fits are shown). Note that* $\mu = \cos\theta$ *(figure 5(b)). It is clear that our model is accurate while the 2-term fit fails to capture multiple scattering effects.*

The more the number of rings used, the more the blockage. The threading in the black rings are designed to trap incident light. Using the radii of the rings and the bulb, we analytically computed the solid angle that is blocked. A schematic representation of the visible solid angle of the source is shown in the second row of figure 8. For each source visibility configuration, we measured the scattered light field (PSF) for different milk concentrations.

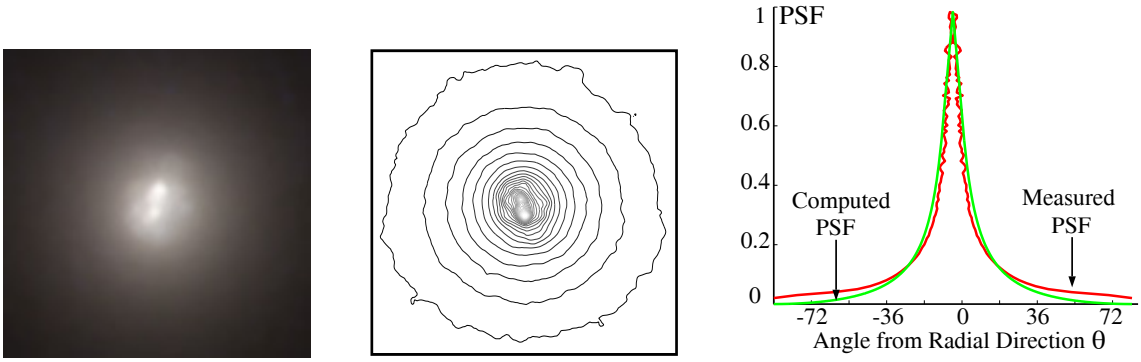

(a) (b) (c)

Fig. 7. *Verification of the analytic model using a distant light source as seen on a rainy night. (a) Image of a distant point source and its glow. (b) Iso-brightness contours of the image showing roughly concentric rings. (c) Comparison between measured PSF and PSF computed using the model.*

In total, we conducted experiments using seven different visibility configurations with each of five milk concentrations. A subset of these experiments are shown in figure 8. In these experiments, we observed that blocking parts of the back hemisphere of the light source did not significantly impact the angular distribution of the scattered light field. This can also be explained by noting that forward scattering is much greater than backward scattering (as predicted by our model too). The PSF curves for all the visibility configurations for each milk concentration is shown in figure 9. We observe that the PSFs for visibility angles 315^o, 270^o and 240^o are almost identical. Only as more and more of the solid angle is blocked (visibility angles $180^o, 90^o, 60^o, 30^o$), do we see a significant change in the observed PSFs. Since we fit our model to the visibility configuration of about 340^o in the previous section and found that the fit overlaps with the 315^o configuration as well (we do not show this curve for clarity), we conclude that partial occlusion of the back hemisphere of light sources does not effect our model appreciably.

7. ISSUES RELEVANT TO RENDERING

So far, we described multiple scattering from a point light source. In practice, however, there are three obstacles for rendering general sources and scenes by directly using eq. 27: (a) visibility issues in real scenes, (b) sources of complex shapes and radiances and (c) efficiency of the algorithm. We address these issues in this section.

7.1 Visibility Issues in Real Scenes

Previously we presented experiments with varying occlusions of the source when the source and the observer have a direct line of sight and observed that partial occlusions of the back hemisphere do not effect the PSF appreciably. Thus our model can be directly applied in these cases. In practice, however, real sources are blocked by complex scene geometry, say adjacent scene objects. It is not clear how to enforce these complex boundary conditions using our model. In this paper, we will not deal with these complex visibility issues. In future work, we wish to understand the precise limitations of our model under complex visibility conditions. However, our experiments on real light sources [Narasimhan and Nayar 2003] have shown that eq. 27 is usually accurate and suffices to produce visually compelling results as long as the occluders are far away from the source.

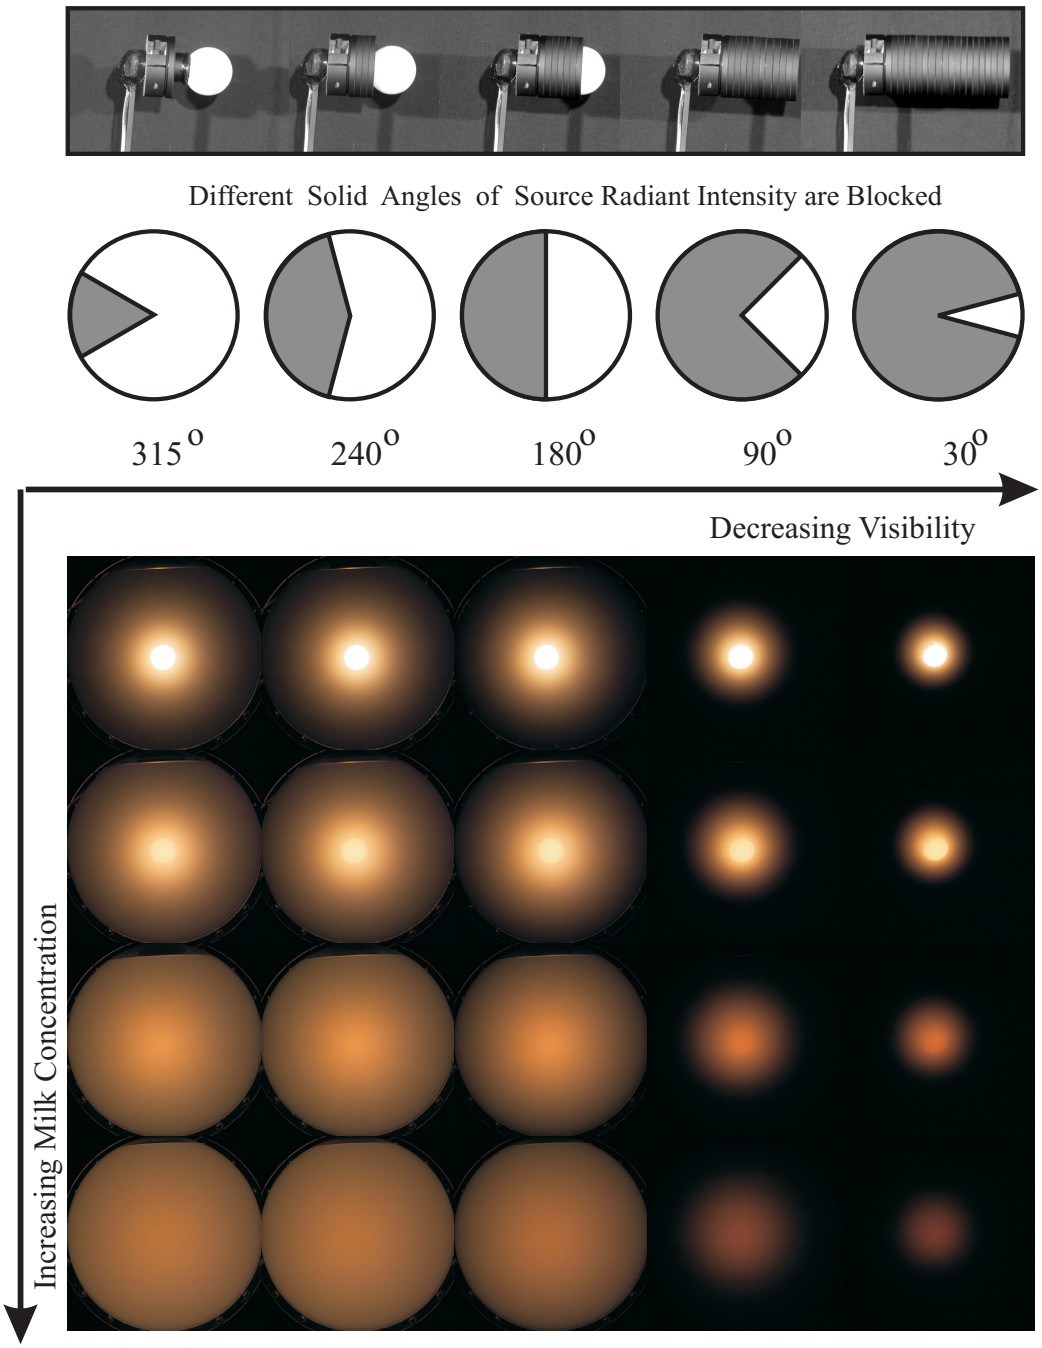

Fig. 8. *Effect of Light Source Visibility on Multiple Scattering. The first two images (on each row corresponding to the blocked angles of upto 120 degrees) are very similar, while we see only small differences in the image pertinent to the 180 degree case. This shows that our model (which is derived for 360 degree visibility) can reasonably approximate upto a back-hemisphere of light being blocked.*

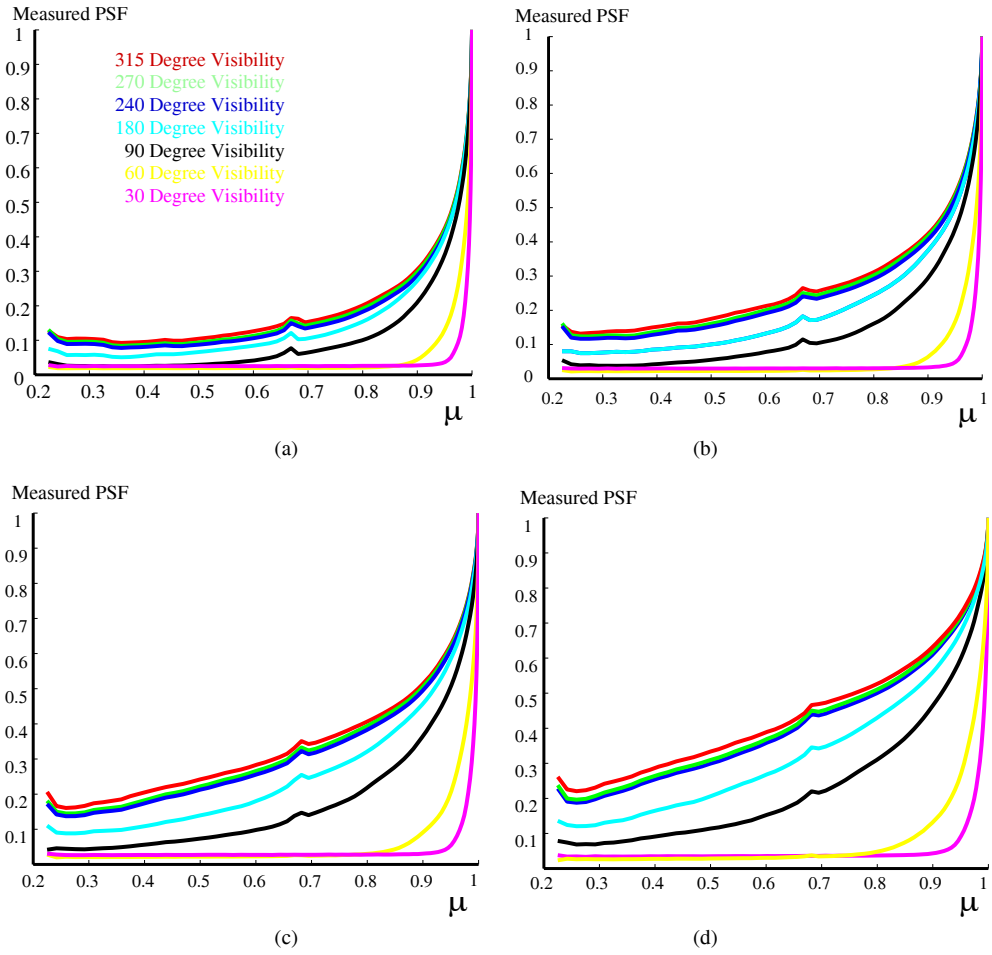

Fig. 9. *PSFs showing the effect of Light Source Visibility on Multiple Scattering. (a) - (d) PSFs measured using 4 increasing milk concentrations. In each plot, the observed PSFs for 7 different source visibility solid angles are shown. Notice that the Red, Green and Blue curves corresponding to partial occlusions of the back hemisphere are very similar. From this we conclude that occlusions of upto* 120° *do not show appreciable differences in the scattered intensities. Hence our model can be applied accurately even for these cases source occlusion.*

7.2 Sources with Complex Shapes and Radiances

First, we discuss how to apply the 3D PSF to simulate the glow around a point source in the image domain. Two schematics are shown in figure 10 to measure multiple scattered rays from different directions depending on whether the camera is a pinhole or an orthographic camera. Similarly, to simulate multiple scattering measurements in lens-based cameras with finite aperture, the appropriate incoming rays and their angles need to be considered. The 3D PSF is converted to a 2D PSF in the image domain using the projection angles ($\cos^{-1}\mu$) shown in figure 10. Similarly, for a camera that is placed outside the medium, we may use the inverse of the mapping function given in eq. 28. Since the 3D PSF is rotationally symmetric, so is the 2D PSF.

For an area source of arbitrary shape and radiance distribution, we subdivide the source

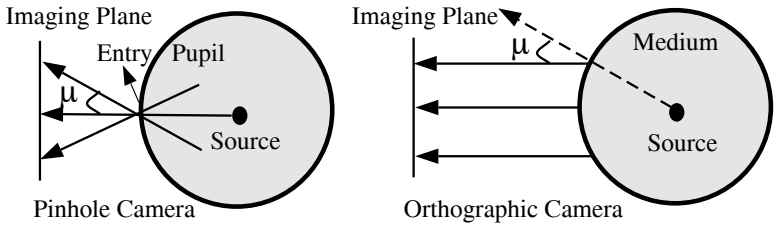

Fig. 10. *Measuring the 3D multiple scattering PSF in the image domain with a pinhole and an orthographic camera. The projection of the 3D PSF onto the image yields the 2D PSF.*

into source elements, apply the PSF to each source element and add the results. This follows directly from the linearity principle of light (assuming incoherence). Note that the PSF depends on the distance of each source element. If the entire area of the source is at the same depth from the observer (as is usually the case), then the PSFs corresponding to different source elements will be the same. Then, the multiple scattering from an area source is simply a convolution of the image of the source with the 2D PSF.

7.3 Efficient Algorithm to Simulate Glows

We described that the glows around a single light source can be implemented as a convolution. For rendering glows around multiple sources at different depths, the method essentially reduces to a spatially varying (depth dependent) convolution. This allows the use of a range of fast spatially varying filtering tools in image processing to simulate glows. To summarize, a step-by-step algorithm to add glows to light sources of arbitrary shapes and radiance distributions in an image is given below.

(1) Segment the image into regions (sources) of equal depths. For synthetic scenes, the depths come from computer models. For real scenes, depths may be estimated using computer vision techniques or range sensors. Note that we require only coarse depth information as opposed to a monte carlo simulation that requires precise depth information. This makes it simple to manually assign depths to real photographs and this is the approach we have taken in our experiments.

(2) For each image region at the same depth R, perform the following steps:
   (a) Input the model (PSF) parameters $T = \sigma R$ and q (or in general, any other phase function parameters). The PSF parameters can be specified interactively, or come from computer models.
   (b) Compute the 3D PSF using eq. 27. Then, compute the 2D PSF in the image domain using the image projection shown in figure 10, or using the geometry of figure 5(b) when the sensor is outside the medium.
   (c) Convolve the image region for a given depth with the 2D PSF obtained in the previous step.

The algorithm suggests that although the mathematical derivation of the analytic formula is complicated, implementation is straightforward and efficient. Our current implementation is about 25 lines of Matlab code.

7.4 General Implications for Rendering

Our model has broad applicability for efficient volume rendering of synthetic or real scenes and sources, as seen through any participating medium. For instance, consider the rendering of a computer generated 3D model of a scene illuminated by various light sources. To generate a foggy or hazy appearance of the scene, we need to simulate physically-based 3D glows around the sources and their brightness contributions at the scene points. Given the 3D geometry of the scene, we can assume that any scene point is at the edge of a spherical medium of radius equal to the distance of the scene point from the source. We can then compute the irradiance at the scene point due to the source using eq. 27. Since all of these computations are analytical, the implementation can be made very efficient for simulating realistic weather effects. Similarly, the model can be used for rendering in other application domains such as medical endoscopy or for underwater imaging.

The model can also be used in the context of adding multiple scattering effects (glows) to images using only rough depth estimates, which, in most cases, can be easily provided by manual segmentations of the photographs. While this is conceptually similar to the works of [Nakamae et al. 1990; Spencer et al. 1995; Beckman et al. 1994], those methods consider glare effects and diffraction in the cornea, a very different problem. Also, brute-force Monte Carlo simulation of glows is too computationally intensive to be tractable. In contrast, our method is efficient allowing us to interactively create images that appear hazy, foggy or misty. Potentially, this can be used in image-based rendering applications or simply as a photoshop-like tool to add weather effects to images.

8. ADDING WEATHER TO PHOTOGRAPHS

We now demonstrate the addition of weather effects to a single photograph using a manually provided rough depth map and some extensions like the attenuation and airlight models. We show results obtained using photographs of three different scenes shown in figure 11(a), figure 12(a) and figure 14(a). All the photographs were acquired using a 12 bits per pixel Kodak DCS 760 color camera. The camera was radiometrically calibrated (linear response function). Multiple exposures of the scene were captured and combined to obtain very high dynamic range photographs of the scene.

8.1 Simple Convolution

In our first example, we added 3 different weather conditions ($q = [0.9, 0.8, 0.5]$ and $T = [1.05, 1.8, 3.0]$ respectively) to the cluster of light sources shown in figure 11(a) at roughly the same depth from the observer. In this case, a single convolution of the original image with the 2D PSF was computed. Note that although the light sources occlude each other, this method still produces visually pleasing results.

8.2 Depth Dependent Convolution with Attenuation

The scene in the previous example was assumed to be at the same depth from the observer. Now, we present an example scene where several depths are visible and we apply depth dependent filters (PSFs with depth dependent values for T) at each depth separately. Consider the traffic scene shown in figure 12 acquired at night. The scene was hand segmented into rough regions of different depths. For instance, a simple ramp function was applied to the road. The resulting depth map is shown in figure 13. The light sources were identified using simple thresholding and the appropriate PSF was used to convolve sources at

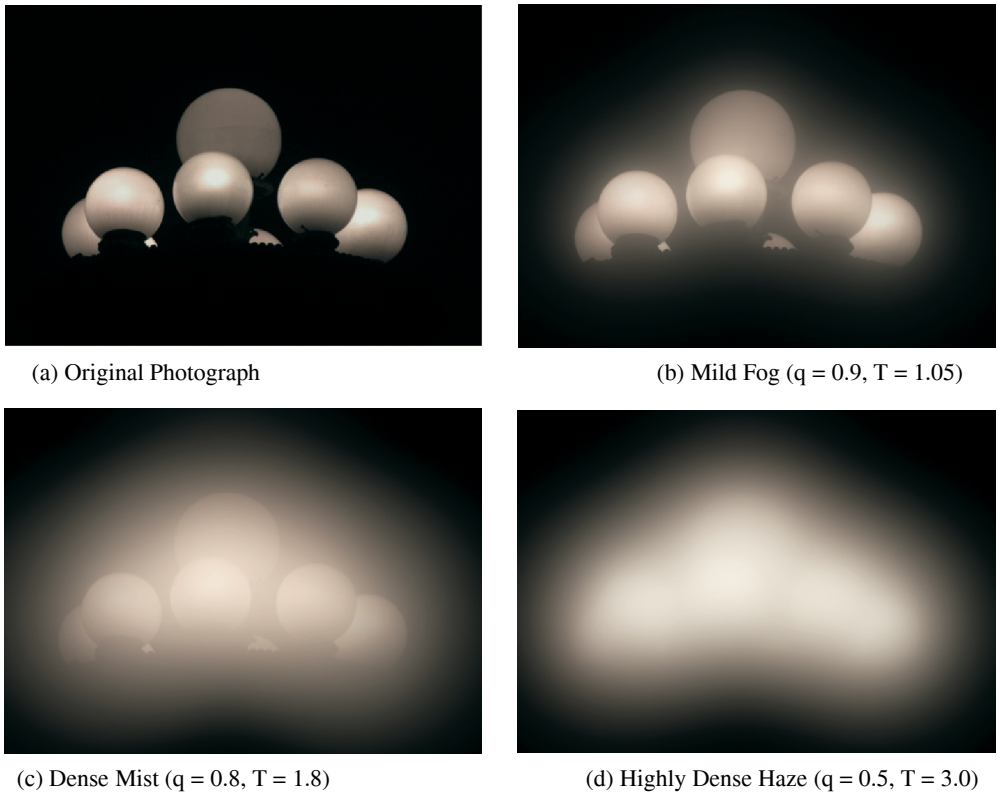

(a) Original Photograph (b) Mild Fog (q = 0.9, T = 1.05)

(c) Dense Mist (q = 0.8, T = 1.8) (d) Highly Dense Haze (q = 0.5, T = 3.0)

Fig. 11. *Light sources glowing under different atmospheric conditions. (a) Original image of a cluster of lamps. (b) - (d) Appearances under different atmospheric conditions and densities (or optical thickness values, T) using a simple convolution of the original photograph with the corresponding scattering PSF.*

each depth separately. For regions in the background, we just applied a simple exponential attenuation model analogous to that used in OpenGL. Thus the final expression for each pixel in the image is a linear combination of the attenuation and the multiple scattering model in eq. 27:

$$L_0 e^{-T}\,\delta(1-\mu) + I(T,\mu)\,, \tag{30}$$

where, L_0 is the radiance of the background scene point. Note that the attenuation model can only dim the intensities of light sources and cannot produce glows (angular spreads) around the sources. Hence, we multiply the attenuation model by a delta function in the head-on direction $\delta(1-\mu)$. Results on applying our model with two different optical thicknesses and with $q = [0.95, 0.75]$ are shown in figure 12(b) and (c).

Note that simple techniques like Gaussian blurring cannot produce *glowing* images shown in figure 13(b). Another simple approximation could be using a Gaussian blurred image added to the attenuated image. These approximations do not work since we cannot set the width of the blur filter according to the depth of scene point in a physically consistent manner. Ad hoc methods could be tried requiring time consuming human intervention, but note here that our method provides an accurate physics-based way of setting

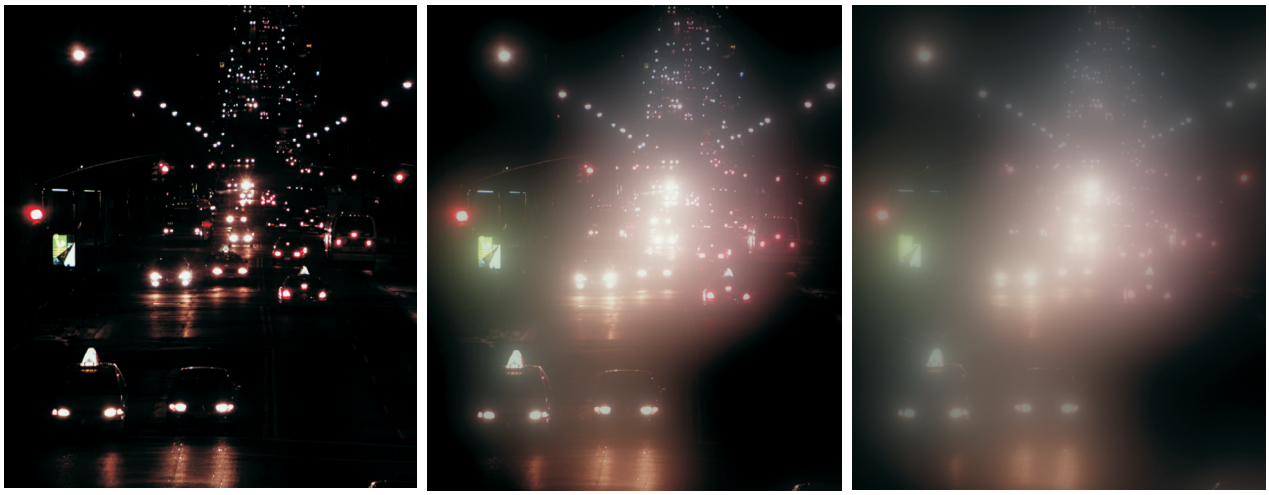

Fig. 12. *On the left is a real photograph of a traffic scene on a clear night. The next images show the effects of mist and fog (with different atmospheric visibilities) added interactively using our analytic model for multiple scattering from light sources. Notice the glows around the sources (street lamps, car headlights and advertisement signs) and that the brightnesses and extents of these glows depend on the depths of scene points. As we shall show, simple approximations like Gaussian blurs, single scattering or diffusion cannot accurately model these complex multiple scattering effects.*

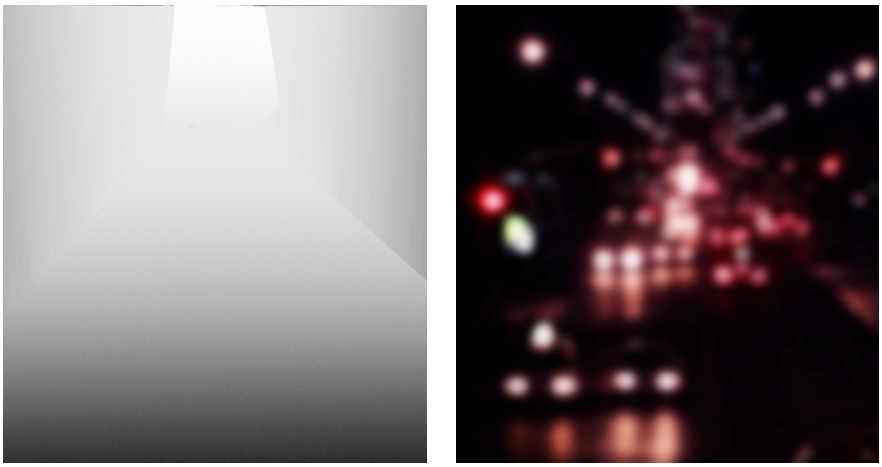

Fig. 13. *On the left, roughly segmented depth map of the scene in figure 12. The road is modeled using a simple ramp (linear) function. Brighter values denote farther scene points. On the right, Gaussian blurring with different widths (proportional to scene depth) applied to each depth separately. Note that such operations cannot produce realistic glows around light sources in photographs. Compare the Gaussian blurred image to the more realistic images in figures 12(b) and (c) generated using our model.*

parameters of the model taking into account the depths of light sources.

8.3 Depth Dependent Convolution with Attenuation and Airlight

The third scene we demonstrate is a scene photographed in the evening with decorative lights on trees. In this case, we took into account environmental illumination (due to the sky) in addition to the attenuation model and the multiple scattering model mentioned above. The sky was assumed to be overcast. The scattering due to skylight, called airlight [Koschmieder 1924], was assumed to be mainly single scattered. The simple attenuation plus airlight model was then applied to the original photograph according to a rough depth segmentation of the scene. The final expression used to render these images is as follows:

$$[L_0 e^{-T} + L_{sky}(1 - e^{-T})]\,\delta(1-\mu) + I(T,\mu)\,, \tag{31}$$

where L_{sky} is the horizon brightness [Narasimhan and Nayar 2002]. As in the second example, we multiply attenuation and airlight by a delta function in the head-on direction.

Two different amounts of mist and fog ($q = [0.8, 0.9]$, minimum $T = [1.05, 2.0]$) were added to the image in figure 14(a). These results are illustrated in figures 14(c) and (d). Note the glowing appearance of the trees and also that the extent of the glows vary with the distance of the trees from the observer. Further, compare our technique with only single scattering shown in figure 14(b). The results in figure 12 and figure 11 indicate that our technique suffices to produce high quality images without noticeable artifacts, and that single scattering on its own cannot produce any glows. Each simulation took less than a minute using Matlab potentially making the method interactive.

In summary, our method for adding weather effects to images is physically-based, fast (near real-time), produces much better results than current approximations like diffusion, blurring and single scattering, and at the same time, the implementation is very simple.

9. CONCLUSIONS AND FUTURE WORK

In this paper, we addressed multiple scattering in participating media like haze, fog and mist, with the practical motivation being addition of weather effects to images, capturing features like the glows around light sources. This is a different problem from that typically handled in plane-parallel diffusion, since we are considering strong angular dependence with the source and observer in the medium. In fact, neither the diffusion approximation, single scattering nor simple blurring can produce the results in this paper. We derived and applied a new analytic formula based on a Legendre polynomial expansion of the spherical radiative transfer equation. We extensively validated our formula using Monte Carlo simulations as well as real experiments with milk under controlled conditions. We applied the results to fast simulation of weather effects in real photographs using rough depth maps. Though the derivation is mathematically involved, the implementation is simple (essentially, spatially varying filtering) and is orders of magnitude faster than Monte Carlo simulation.

Current limitations of the work include practical issues like spatially varying media, non-isotropic sources and limitations of visibility. We believe this is a first important step, and there are clear ways they could be addressed in future work. Some initial directions are presented below.

—*Analytic Bases for Variations of the Spherical RTE:* In this work, we developed analytic bases (Legendre polynomials) for multiple scattering from a point light source placed at the center of a spherical medium. Future work includes the derivation of similar analytic

(a) Original Photograph

(f) Airlight + Attenuation model (minimum T = 1.05)

(c) Mild Mist (q = 0.8, minimum T = 1.05)

(d) Dense Fog (q = 0.9, minimum T = 2.0)

Fig. 14. *Decorative lights on trees glowing in fog and mist. (a) Original Photograph. (b) Foggy image rendered using single scattering (airlight + attenuation models). (c) and (d) Glows are added to lights on the trees. Note that multiple scattering effects due to light sources are significant as compared to single scattering effects.*

bases for simple variations on the spherical RTE configurations of medium and source geometry. For instance, what happens when the source is not at the center of the medium but rather at a different location inside or outside the medium? Then, the light field also depends on the azimuthal angle from the radius vector. In this case, we believe that using spherical harmonics can enable us to capture the multiple scattering within this volume.

—*Analytic approximation for Visibility of Sources:* In our analytic model, we assumed that an isotropic point source is immersed in a spherical medium with no other objects or occluders. In reality, the source could be occluded by nearby objects or could be non-isotropic. Areas of darkening may be seen in the medium where the occluders cast shadows. This is a hard problem to solve analytically. However, we believe that certain analytic approximations can be obtained. For instance, to darken the areas in the shadow region, we may place *negative* isotropic sources in the interior of the occluder so as to create the right area of darkening. To handle non-isotropic sources, we need to consider only a source obtained by blocking a finite solid angle of an isotropic source. Note

that an arbitrary non-isotropic source could be constructed from smaller solid angles of isotropic sources.

—*Guiding Monte Carlo Ray Tracers using Analytic Models:* One of the main problems of Monte Carlo approaches to rendering volumetric effects is that they are computationally very expensive. Their computation cost depends on the parameters of the medium (say, the volume of the medium and the density of particles in the medium). In other words, the time required to render an image of a scene in dense fog could be several orders of magnitude greater than the time required to render the same scene in mild fog. Also, the time required to render a larger (scaled) version of the same scene with the same particle density could be orders of magnitude higher. Thus, rendering scenes by interactively changing parameters of the scene (be it parameters of the medium or the objects in the scene itself) can be more or less ruled out. On the positive side, Monte Carlo methods can accurately take into account arbitrary source and medium configurations, complex visibility effects as well as arbitrary source radiance distributions.
A possible direction of future work could be to use of our analytic models and approximations to guide Monte Carlo techniques to faster and more accurate convergence. We believe that our analytic models will provide starting estimates that are very close to the final solution in most cases. Thus, the hard dependence of Monte Carlo on the parameters of the scene is alleviated. Also, we can take advantage of the flexibility in Monte Carlo to render subtle effects due to complex scenes, that are not captured using the analytic models. This hybrid method enjoys the advantages of both the techniques.

Using the above set of models and algorithms, it may be possible to render a wide variety of volumetric effects (atmospheric, underwater, subsurface) efficiently as well as accurately in general settings.

ACKNOWLEDGMENTS

This work was supported in parts by a DARPA/ONR HumanID Contract (N00014-00-1-0916) and an NSF Award (IIS-99-87979). The authors thank Prof. Henrik Wann Jensen for useful discussions on the topic and for his monte carlo renderer (Dali) [Jensen]. The authors also thank the Stanford Graphics Lab for allowing them to run monte carlo simulations on the Stanford Chromium PC cluster. Lastly, the authors thank Estuardo Rodas for building the “milk tank” setup.

REFERENCES

AMBARTSUMIAN, V. 1945. A point source of light within a scattering medium. *Bulletin of the Erevan Astronomical Observatory 6,* 3.

ANTYUFEEV, S. 2000. *Monte Carlo Method for Solving Inverse Problems of Radiative Transfer*. Inverse and Ill-Posed Problems Series, VSP Publishers.

BECKMAN, C., NILSSON, O., AND PAULSSON, L. 1994. Intraocular light scattering in vision, artistic painting, and photography. *Applied Optics 33,* 21.

BLASI, P., SAEC, B. L., AND SCHLICK, C. 1993. A rendering algorithm for discrete volume density objects. *Computer Graphics Forum 12,* 3, 201–210.

CHANDRASEKHAR, S. 1960. *Radiative Transfer*. Oxford Univ. Press.

DORSEY, J., EDELMAN, A., JENSEN, H., LEGAKIS, J., AND PEDERSEN, H. 1999. Modeling and rendering of weathered stone. In *SIGGRAPH 99*. 225–234.

E. NAKAMAE, K. HARADA, T. I. AND NISHITA, T. 1986. Montage : The overlaying of the computer generated image onto a background photograph. In *Computer Graphics*. Vol. 20. 207–214.

EBERT, D. AND PARENT, R. 1990. Rendering and animation of gaseous phenomena by combining fast volume and scanline a-buffer techniques. In *SIGGRAPH 90*. 357–366.

ELLIOTT, J. 1955. Milne's problem with a point-source. *Proc. of Royal Soc. of London, Series A, Mathematical and Physical Sciences 228,* 1174.

HANRAHAN, P. AND KRUEGER, W. 1993. Reflection from layered surfaces due to subsurface scattering. In *SIGGRAPH 93*. 165–174.

HENYEY, L. AND GREENSTEIN, J. 1941. Diffuse radiation in the galaxy. *Astrophysics Journal 93*, 70–83.

ISHIMARU, A. 1978. *Wave Propagation and Scattering in Random Media. Volume 1: Single Scattering and Transport Theory*. Academic Press.

JENSEN, H., MARSCHNER, S., LEVOY, M., AND HANRAHAN, P. 2001. A practical model for subsurface light transport. In *SIGGRAPH 01*. 511–518.

JENSEN, H. W. Dali rendering software.

KAJIYA, J. AND HERZEN, B. 1984. Ray tracing volume densities. In *SIGGRAPH 84*. 165–174.

KOENDERINK, J. J. AND VAN DOORN, A. 2001. Shading in the case of translucent objects. In *Human Vision and Electronic Imaging VI, B.E. Rogowitz, T.N. Pappas (eds.), SPIE*. 312–320.

KOPEIKA, N. 1981. General wavelength dependence of imaging through the atmosphere. *Applied Optics 20,* 9 (May).

KOSCHMIEDER, H. 1924. Theorie der horizontalen sichtweite. *12*, 33–53,171–181.

LANGUENOU, E., BOUATOUCH, K., AND CHELLE, M. 1994. Global illumination in presence of participation media with general properties. In *Eurographics Rendering Workshop*. 69–85.

LECOCQ, P., KEMENY, A., MICHELIN, S., AND ARQUES, D. Mathematical approximation for real-time lighting rendering through participating media. In *Pacific Graphics 2000*.

MACROBERT, T. 1967. *Spherical Harmonics: An Elementary Treatise on Harmonic Functions with Applications*. Pergamon Press.

MARSHAK, R. 1947. Note on the spherical harmonic method as applied to the milne problem for a sphere. *Physical Review 71,* 7.

MAX, N. 1994. Efficient light propagation for multiple anisotropic volume scattering. In *Eurographics Rendering Workshop*. 87–104.

MAX., N. L. 1986. Atmospheric illumination and shadows. In *SIGGRAPH 86*. 117–124.

NAKAMAE, E., KANEDA, K., OKAMOTO, T., AND NISHITA, T. 1990. A lighting model aiming at drive simulators. In *SIGGRAPH 90*. 395–404.

NALLASAMY, M. 1984. Mie scattering functions for milk fat globules. In *Rev. Sci. Instrum.* Vol. 55. 814.

NARASIMHAN, S. G. AND NAYAR, S. K. 2002. Vision and the atmosphere. *28,* 3 (August).

NARASIMHAN, S. G. AND NAYAR, S. K. 2003. Shedding light on the weather. In *To appear in CVPR 03*.

NISHITA, T. AND NAKAMAE, E. 1987. A shading model for atmosphere scattering considering luminous intensity distribution of light sources. In *Computer Graphics*. Vol. 21. 303–310.

NISHITA, T., Y.DOBASHI, AND E.NAKAMAE. 1996. Display of clouds taking into account multiple anisotropic scattering and sky light. In *SIGGRAPH 96*. 379–386.

PATTANAIK, S. AND MUDUR, S. 1993. Computation of global illumination in a participating medium by monte carlo simulation. *Journal of Visualization and Computer Animation 4,* 3, 133–152.

PHARR, M. AND HANRAHAN, P. 2000. Monte carlo evaluation of non-linear scattering equations for subsurface reflection. In *SIGGRAPH 00*. 75–84.

RUSHMEIER, H. AND TORRANCE, K. 1987. The zonal method for calculating light intensities in the presence of a participating medium. In *SIGGRAPH 87*. 293–302.

SAKAS, G. 1990. Fast rendering of arbitrary distributed volume densities. In *Eurographics 90*. 519–530.

SPENCER, G., SHIRLEY, P. S., ZIMMERMAN, K., AND GREENBERG, D. P. 1995. Physically-based glare effects for digital images. In *SIGGRAPH 95*. 325–334.

STAM, J. 1995. Multiple scattering as a diffusion process. In *Eurographics Rendering Workshop*. 41–50.

A. APPENDIX : RELATION TO DIFFUSION

In this section, we discuss how the final formula in eq. 27 relates to the popular diffusion model. We write down the first 2 terms of the series explicitly, taking $I_0 = 1$ for simplicity,

$$g_1(T) = \frac{e^{-3(1-W_0)T}}{T^2} \quad g_2(T) = \frac{e^{-\frac{5}{2}(1-q)T}}{T^3} \tag{32}$$
$$I(T,\mu) = \frac{e^{-3(1-W_0)T}}{T^2} L_0(\mu) + \left\{ \frac{e^{-3(1-W_0)T}}{T^2} + \frac{e^{-\frac{5}{2}(1-q)T}}{T^3} \right\} L_1(\mu) + \ldots$$

where these formulae include both single and multiple scattering. First, consider the case when W_0 is close to 1, so there is very little absorption. Since $\alpha_m = m+1$, and g_m depends on $\exp[-\alpha_m \log T]$, succeeding terms g_m decay as $1/T^{m+1}$. Thus, regardless of the exponential factors, for large optical thicknesses T, $g_1(T)$ will be the dominant term, and we will have a two term expansion $g_1(T)(L_0(\mu) + L_1(\mu))$. Substituting $L_0 = 1$ and $L_1 = \mu$, we get,

$$I(T,\mu) \approx \frac{e^{-3(1-W_0)T}}{T^2}(1+\mu). \tag{33}$$

Conceptually, this is very similar to the two term diffusion expansion, valid for very optically dense media (large T). As with standard diffusion theory, the dependence on the medium, as determined by q is weak, vanishing in the above formula. However, the standard diffusion equation is derived using radiative transfer in plane parallel media shown in figure 1(b) and its form is different.

We note that as for diffusion, the coefficient of μ decays as $1/T^2$, as it must for flux to be conserved (we will scale the solution by σ^2 as explained in the paper). Next, we consider the angular dependence $1+\mu$. This angular dependence ensures that the source intensity is maximum head on ($\mu = 1$), always remains positive, and eventually decays to 0 (for $\mu = -1$). It also ensures that the common boundary condition that there be no inward intensity (or the simpler condition that there be no inward flux), in case of a finite medium (for some extent T) is approximately satisfied, since $I(T,\mu)$ is relatively small for $\mu < 0$. Diffusion theory is consistent with similar boundary conditions, but they are not usually applied for a spherical finite medium of radius T, but for semi-infinite plane parallel atmospheres. The standard diffusion model for a point source has the coefficient of L_0 decaying as $1/T$ instead of $1/T^2$. This causes the field to become uniform with angle at large T. However, that would predict no angular dependence or uniform glows over the entire image, which is clearly unphysical in our scenario; standard diffusion theory cannot be applied to situations involving strong angular dependences of sources inside media.