

A BRDF Database Employing the Beard-Maxwell Reflection Model

Harold B. Westlund^a

Gary W. Meyer^b

Department of Computer and Information Science
University of Oregon · Eugene, Oregon · U.S.A.

Abstract

The Beard-Maxwell reflection model is presented as a new local reflection model for use in realistic image synthesis. The model is important because there is a public domain database of surface reflection parameters, the Nonconventional Exploitation Factors Data System (NEFDS), that utilizes a modified form of the Beard-Maxwell model. Additional surface reflection parameters for the database can be determined because a measurement protocol, using existing radiometric instruments, has been specified. The Beard-Maxwell model is also of historical significance because it predates many computer graphics reflection models and because it includes several features that are incorporated into existing local reflection models. The NEFDS is described and a special shader is developed for use with NEFDS. The shader makes use of the alias method for determining random variates from discrete probability distributions. Realistic images are synthesized from the existing database and from samples that were characterized using the measurement protocol.

Key words: BRDF, local illumination

1 Introduction

Computer graphics has made great progress in the simulation of object appearance. Reflection models have been developed that characterize both the spectral and the spatial distributions of light reflected from an object's surface. Image synthesis systems have been constructed that use Monte Carlo techniques to evaluate these reflection models and simulate the objects in the context of globally illuminated environments. With sufficient processing time, these programs are capable of producing individual photorealistic images. Recent developments in the area of real-time shading have made it possible to utilize sophisticated reflection models in interactive programs [18, 20]. These advances in real-time rendering are

likely to bring renewed attention to the subject of surface reflection modeling.

Unfortunately, the accurate synthesis of object appearance is currently limited by the small amount of readily available surface reflectance data. Many sophisticated computer graphics reflection models have been proposed. The best of these models characterize both the spectral and the spatial distributions of reflected light and are therefore appropriate for modeling object appearance. These models have a sound theoretical foundation and often contain measurable surface reflection parameters. Regrettably, measurement protocols and actual measured data are seldom available for the models. This situation is in contrast to object shape information where polygon data is now commercially available.

This paper presents the Beard-Maxwell [30] reflection model which has not been discussed in the computer graphics literature even though a large database of measurements exists for the model. The model is of historical importance because it predates many of the current computer graphics reflection models and foreshadows several of the features that have been incorporated into these models. The model has practical significance because the public domain Nonconventional Exploitation Factors Data System (NEFDS) [1] utilizes a modified version of the Beard-Maxwell model and contains over 400 materials for which the parameters of the modified Beard-Maxwell model have been measured. Furthermore, the NEFDS can be extended because a measurement protocol exists, using standard radiometric instruments, to acquire model parameters for additional materials.

The paper also describes a rendering system that has been developed for use with NEFDS. The Radiance software package [2] has been modified so that it can be employed to make pictures with data taken from NEFDS. Since the Beard-Maxwell model is not an invertible function, this necessitated the use of a probability mass function (the discrete counterpart of the probability density function) and the alias method for generating random variates. The research discussed in this paper also includes an attempt to add additional material to the database. The measurements that were necessary to ac-

^aCurrent address: Radical Entertainment, Vancouver, BC, Canada. hwestlund@radical.ca

^bCurrent address: Department of Computer Science and Engineering, University of Minnesota, Minneapolis, MN, U.S.A. meyer@cs.umn.edu

comply with this are described, and the quality of the approximation provided by Beard-Maxwell and NEFDS is evaluated.

This paper is divided into four major sections. In the next section, a brief review is done of existing computer graphics reflection models. The following section introduces the complete Beard-Maxwell model and shows how it has been modified for use in the NEFDS. An overview of the NEFDS is given in the penultimate section and example renderings that were made using the data are discussed in the last section. These pictures were made using a modified version of a public domain rendering program and data from both NEFDS and from new measurements made using the NEFDS measurement protocol. The paper concludes with a summary of the work and suggestions for further research.

2 Background

2.1 BRDF

In order to create a realistic image of an object, the light reflection properties of the object's surface must be specified. The most common means of quantifying surface reflection of light is by utilizing the bi-directional reflectance distribution function (BRDF). The BRDF, ρ , is defined as the ratio of differential reflected radiance to differential incident irradiance:

$$\begin{aligned} \rho(\Theta_i; \Theta_r; \lambda) &= \frac{dL_r(\Theta_i; \Theta_r; \lambda)}{dE_i(\Theta_i; \lambda)} \\ &= \frac{dL_r(\Theta_i; \Theta_r; \lambda)}{dL_i(\Theta_i; \lambda) \cos \theta_i d\omega_i} \end{aligned} \quad (1)$$

where the subscripts i and r denote incident and reflected respectively, $\Theta = (\theta, \phi)$ is the direction of light propagation, λ is the wavelength of light, L is radiance, E is irradiance, and $d\omega$ is an element of solid angle [31]. This surface reflection geometry is shown in Figure 1.

2.2 Physics Based BRDFs

Over the past century, a wide variety of analytical reflection models have been created to represent the physics of light reflection off surfaces. They have been developed in the fields of physics and engineering accompanying the development of radio transmission, radar, heat transfer theory, remote sensing, astronomy, and many other areas of research. A few relevant models are presented here.

In 1963, Beckman and Spizzichino [6] developed a model of reflection based upon wave optics. Their reflection model was derived through the use of the Kirchhoff approximation of the Helmholtz boundary conditions. Also in 1963, Hapke [15] developed a reflection model which predicted the light reflection off the lunar surface using scattering theory. He was able to account




Figure 1: Light reflection geometry

for both forward and backward scatter through his selection of the scattering function.

Torrance and Sparrow [37], in 1967, created an analytical reflection model based on a surface composed of small, randomly distributed, mirror-like facets. These micro-facets were distributed using a Gaussian based distribution, which was in correspondence to the measured surface reflection. Their model was novel in its inclusion of a term for representing the shadowing and masking of neighboring facets which helped to predict the measured off-specular peak.

2.3 Physics Based BRDFs in Computer Graphics

The work of Torrance and Sparrow was first introduced to computer graphics by Blinn [7] who extended it by utilizing the facet distribution proposed by Trowbridge and Reitz [38]. This distribution is built on the assumption that the facets are ellipsoids of revolution. Another adaptation of Torrance and Sparrow's reflection model to computer graphics was offered by Cook and Torrance [10]. They incorporated the Fresnel term into the model to capture the wavelength dependency of first surface reflection. Recently Ashikhmin et al. [5] have shown how to produce BRDFs from very general micro-facet distributions by using a model with a simple shadowing term.

Blinn [8] utilized the work of Hapke to develop a reflection model for use in the simulation of the rings of Saturn. He utilized several functions to represent the phase function of light off the scattering particles, including the Henyey-Greenstein [19] function and a function of Rayleigh scattering. Blinn also discussed the extension of the model to account for shadowing and multiple scattering.

He et al. [16] utilized the work of Beckman and Spizzichino in addition to the work of Torrance and Sparrow to generate a model which accounted for many physics derived attributes. Their model utilized polariza-

tion of light, directional Fresnel effects and a more complex distribution of micro-facets. It was the first model to provided a unified means of predicting scatter from surfaces with autocorrelation distances much smaller, on the same order, or much greater than the wavelength of incident light. Beyond He et al., Stam [36] is one of the few to have attempted the development of more general physics based BRDFs in computer graphics.

2.4 Empirical BRDFs in Computer Graphics

As pointed out by Ward [29] the very complexity that gives these physics based BRDFs validity also hampers their general use in computer graphics. In Ward’s paper he presented a reflection model using an elliptical Gaussian based formula which is ideally suited for fitting to measured data. Ward demonstrated its validity through measurement of several surface BRDFs which were fit to his reflection model.

Lafortune et al. [26] presented a BRDF model built upon sums of cosine lobes specifically constructed to fit measured BRDF data. The number, orientation, and shape of the cosine lobes are variable, allowing for a convincing fit to a wide variety of BRDF distributions.

Dana et al. [11] fit measured data to both the Oren-Nayer model [32] and the Koenderink et al. decomposition [22]. Oren and Nayer built a model of diffuse reflection based on direct and indirect reflections off diffuse micro-facets. Koenderink et al. decomposed the BRDF into a variable-order vector of coefficients, fit to data.

2.5 BRDF Storage Representations in Computer Graphics

Fitting measured data to a BRDF model is avoided if the data rather than the model parameters is stored. Gondek et al. [14] utilized this technique with an adaptive subdivision scheme over the sampled reflection hemisphere. Schröder and Sweldens [35] as well as Lalonde and Fournier [27] offer an efficient method of storing and utilizing BRDFs using spherical wavelets. Cabral et al. [9] and Westin et al. [47] offered the same using spherical harmonics.

3 A New Reflection Model

The modified Beard-Maxwell model presented in this paper incorporates the best attributes of physical and empirical BRDFs. It is a physics based model capturing subtle BRDF characteristics required in realistic image synthesis. Its empirical measurements are themselves built upon the physical model, providing a means to set the required parameters accurately with relatively little effort. Since only a minimal set of defining parameters are used, the model requires minimal storage space.

3.1 Beard-Maxwell Reflection Model

The Beard-Maxwell model presented by Maxwell et al. [30] was originally used to describe the reflection properties of rough, painted surfaces displaying Fresnel effects and later was applied to a wider variety of surfaces. Their model is built on the assumption that the material surface is a three dimensional terrain of micro-facets of varying orientation. In this model, reflected light is the result of only two physical occurrences. Light is reflected off one of the micro-facets (first surface reflectance) and light is scattered out of the surface after having first entered the sub-surface medium (volumetric reflectance). The Beard-Maxwell reflectance model thus takes the form

$$\rho(\Theta_i; \Theta_r) = \rho_{fs}(\Theta_i; \Theta_r) + \rho_{vol}(\Theta_i; \Theta_r) \quad (2)$$

where ρ_{fs} and ρ_{vol} are the first surface and volumetric reflectance functions respectively. For the sake of notational clarity the wavelength of this wavelength dependent reflectance function is not listed as an explicit parameter. We will take this same liberty to hide other functional parameters which are tangential to current discussions.

First surface reflection causes light to be reflected in the specular direction (i.e., mirror reflection) off each individual micro-facet as determined by the micro-facet’s normal rather than the macro-surface normal. Therefore the distribution of the first surface reflectance is determined by the distribution of the micro-facet normals which in turn is driven by the density function $\Xi(\Theta)$, the relative density of micro-facet normals (per steradian) in direction Θ . Maxwell et al. calculated the first surface reflectance to be

$$\rho_{fs}(\Theta_i, \Theta_r) = \frac{R(\beta)\Xi(\hat{H})}{4 \cos \theta_i \cos \theta_r} SO \quad (3)$$

where \hat{H} is the half angle vector, β is the bistatic angle (i.e., the angle between either the incident or reflected direction and the half angle vector), $R(\beta)$ is the Fresnel reflectance and SO is a shadowing and obscuration¹ term. Figure 2 is a diagram of the geometry used by the Beard-Maxwell model.

The shadowing and obscuration term in (3), SO, accounts for the height distribution of the micro-facets. Shadowing and obscuration are due to intersections between the other surface facets and the incident and reflected light rays respectively. Torrance and Sparrow [37] accounted for the light ray intersections with a purely analytical function derived from theory. Blinn [7] and Cook

¹Obscuration is often termed masking in the literature, but we will use obscuration here to follow the original discussion in both Maxwell et al. [30] and the NEFDS documentation [4].




Figure 2: Beard-Maxwell BRDF model geometry

and Torrance [10], who were the first to introduce shadowing and obscuration in computer graphics, also used a purely analytical function. Beard and Maxwell chose to create a function with free parameters set using empirical data. This function, which has a maximum value of one for both specular reflection ($\theta_{\hat{H}} = 0$) and pure backscatter ($\beta = 0$), is defined as

$$SO = \frac{1 + \frac{\theta_{\hat{H}}}{\Omega} e^{-2\beta/\tau}}{1 + \frac{\theta_{\hat{H}}}{\Omega}} \left(\frac{1}{1 + \frac{\phi_n \theta_i}{\Omega}} \right) \quad (4)$$

where the free parameters Ω and τ modify the falloff of SO in the forward-scattered and back-scattered directions respectively, and ϕ_n is a parameter computed from geometry which affects the rate of change of SO [30]. The process of gathering empirical data used to set the equivalent free parameters in the modified Beard-Maxwell model is discussed below in Section 3.2.

Rather than attempting to measure Ξ of (3) directly, Maxwell et al. replaced it with the measured zero-bistatic ($\beta = 0$) first surface reflectance at the half angle, $\rho_{fs}(\hat{H}, \hat{H})$. $SO = 1$ when $\beta = 0$, so (3) can be rewritten as

$$\Xi(\hat{H}) = \frac{4\rho_{fs}(\hat{H}, \hat{H}) \cos^2 \theta_{\hat{H}}}{R(0)}. \quad (5)$$

This is simplified further by their assumption that the surface is isotropic. In that case to generate reflectance values for all incident and reflected directions, $\rho_{fs}(\hat{H}, \hat{H})$ needs only be sampled at the angles $0 \leq \theta_{\hat{H}} \leq \frac{\pi}{2}$ and $\phi = 0$.

Light reflected from the first surface is assumed to maintain its original polarization (i.e., incident and reflected light are of like polarization) while any light reflected through volumetric scattering is assumed to be totally depolarized. In this way, separation of the first surface reflectance from the volumetric reflectance can easily be performed by using a polarized illuminant and

then measuring the like and cross polarized light. The first surface reflectance can be expressed in measurable quantities by using the four combinations of polarization, $\rho_{\perp\perp}$, $\rho_{\perp\parallel}$, $\rho_{\parallel\perp}$, and $\rho_{\parallel\parallel}$ ²:

$$\rho_{fs} = (\rho_{\perp\perp} - \rho_{\perp\parallel}) + (\rho_{\parallel\parallel} - \rho_{\parallel\perp}). \quad (6)$$

Similarly, the volumetric reflectance can be expressed in measurable quantities:

$$\rho_{vol} = 2\rho_{\perp\parallel} + 2\rho_{\parallel\perp} \quad (7)$$

The volumetric reflectance of (2), ρ_{vol} is represented by either Lambertian reflectance, ρ_D , or directional diffuse reflectance, ρ_{ddif} :

$$\rho_{vol} = \begin{cases} \rho_D \\ \rho_{ddif} = \frac{2\rho_v f(\beta)g(\theta_{\hat{N}})}{\cos \theta_i + \cos \theta_r} \end{cases} \quad (8)$$

where ρ_v is the base directional diffuse reflectance when both θ_i and θ_r are zero and where $f(\beta)$ and $g(\theta_{\hat{N}})$ are empirically determined functions that characterize differential subsurface scattering and substrate reflection respectively. The directional diffuse component accounts for any subsurface directional scattering and is selected only if there is any directional reflection unaccounted for by the first surface reflectance. The Lambertian reflectance is selected if the first surface reflectance accounts for all directional reflectance. The directional diffuse component of (8) is based on a model of a subsurface of finite thickness built on a reflective base. Maxwell et al. found that in practice surfaces were well represented by using a constant valued scattering function and by modeling the surface as having infinite thickness. With these simplifications, f and g can be both set to unity and the volumetric reflectance becomes:

$$\rho_{vol} = \begin{cases} \rho_D \\ \rho_{ddif} = \frac{2\rho_v}{\cos \theta_i + \cos \theta_r} \end{cases} \quad (9)$$

3.2 Modification to Beard-Maxwell

The NEF uses a modified form of the Beard-Maxwell reflection model to characterize the reflection properties of surfaces in the NEF database. The NEF Beard-Maxwell (NEF-BM) model includes changes in an attempt to more accurately represent surfaces other than paint, the original target of the Beard-Maxwell model. Additionally, fixed radiometric measurements are introduced which can be

²We follow the same notation as in Maxwell et al. to denote polarized reflectance by $\rho_{\alpha_i \alpha_r}$, where the orientations of α_i and α_r indicate the polarization of the incident and reflected light respectively. The polarization state of the incident light is either parallel (\parallel) or perpendicular (\perp) to the plane of incidence. The polarization state of the reflected light is described in the same way, but with respect to the reflectance plane.

used to determine appropriate values for the model's free parameters.

In the Beard-Maxwell model, only one volumetric component was utilized, either the Lambertian reflectance or the directional diffuse reflectance. In the NEF-BM model *both* are used. The volumetric Lambertian reflectance is referred to as simply diffuse reflectance and is attributed to multiple first surface reflections. The volumetric directional diffuse reflectance is itself referred to as volumetric reflectance, ρ_{vol} , and is attributed to all the reflectance due to subsurface light interaction.

The NEF-BM model uses a simplified form of (4) for its shadowing and obscuration function:

$$SO = \frac{1 + \frac{\theta_{\hat{N}}}{\Omega} e^{-2\beta/\tau}}{1 + \frac{\theta_{\hat{N}}}{\Omega}}.$$

The model also uses the simplified form of directional subsurface scattering from (9). So the complete NEF-BM BRDF can be written

$$\rho(\Theta_i; \Theta_r) = \frac{R(\beta)}{R(0)} \frac{\rho_{fs}(\theta_{\hat{H}}, 0; \theta_{\hat{H}}, 0) \cos^2 \theta_{\hat{H}}}{\cos \theta_i \cos \theta_r} \left(\frac{1 + \frac{\theta_{\hat{H}}}{\Omega} e^{-2\beta/\tau}}{1 + \frac{\theta_{\hat{H}}}{\Omega}} \right) + \rho_D + \frac{2\rho_v}{\cos \theta_i + \cos \theta_r} \quad (10)$$

As in the BM model, the NEF-BM performs the zero-bistatic BRDF measurements at all four combinations of polarization using (6) and (7) to generate the first surface reflection parameters as well as the diffuse and volumetric coefficients. Additionally the NEF-BM utilizes two other series of radiometric measurements.

The first additional set of measurements is a series of specular BRDF measurements for parallel polarized source and receiver over $10 \text{ deg} \leq \theta_i \leq 80 \text{ deg}$ with $\phi_i = 0 \text{ deg}$ and $\phi_r = 180 \text{ deg}$. For each sampled incident direction, the receiver direction is varied about the mirror direction in the plane of incidence by $\theta_i - 5 \text{ deg} \leq \theta_r \leq \theta_i + 5 \text{ deg}$. Data obtained from these measurements provides a means of calculating the complex index of refraction components, n and k which are required for computing the Fresnel reflection coefficient, R . Determination of n and k is performed by utilizing Brewster's angle, the angle of minimum reflectance of parallel-polarized light [17, 4].

The second series of measurements uses a fixed source direction and varying receiver direction, in the plane of incidence, over $-90 \text{ deg} \leq \theta_r \leq 90 \text{ deg}$. These measurements offer means to select the shadowing and obscuration parameters, τ and Ω , in order to reduce the error away from the forward-scattered and back-scattered directions. The values of τ and Ω are set using a

least square fit of the modeled BRDF, with shadowing and obscuration, to this measured data. This second series of measurements also provides a good verification of the previously selected model parameters.

We now introduce the dependency of the BRDF on the wavelength of light which has hitherto been considered constant. The naive method of characterizing spectral dependency would be to perform all required measurements at a sufficiently dense sampling of wavelengths. However this would undermine one of the primary motivations of the NEF-BM model which is to represent the BRDF using relatively few measurements.

The solution was found in the observation that for many surfaces, the BRDF at one wavelength is similar in shape, but not necessarily magnitude, to the BRDF of the surface at nearby wavelengths. To utilize this information, the measurements detailed above are only performed at key reference wavelengths and intermediate values are linearly interpolated. The interpolated values are then scaled by directional hemispherical reflectance (DHR) [31] values, $\rho_d(\lambda)$, measured at a dense sampling of wavelengths. The full spectral dependent BRDF then is:

$$\rho(\bar{\Theta}; \lambda) = \rho_d(\lambda) \left[\frac{\rho(\bar{\Theta}; \lambda_j)}{\int \int \rho(\bar{\Theta}_{DHR}; \lambda_j) \cos \theta_i d\theta_i d\phi_i} \left(\frac{\lambda_k - \lambda}{\lambda_k - \lambda_j} \right) + \frac{\rho(\bar{\Theta}; \lambda_k)}{\int \int \rho(\bar{\Theta}_{DHR}; \lambda_k) \cos \theta_i d\theta_i d\phi_i} \left(\frac{\lambda - \lambda_j}{\lambda_k - \lambda_j} \right) \right] \quad (11)$$

where j and k are the indices of the reference wavelengths bounding λ , $\bar{\Theta}$ is the arbitrary geometry $(\Theta_i; \Theta_r)$, and $\bar{\Theta}_{DHR}$ is the geometry of DHR measurement $(\Theta_i; \Theta_{DHR})$.

This method offers the additional benefit of increased accuracy. DHR measured values are typically more accurate than the corresponding integrated values of measured BRDFs. Scaling the BRDFs to ensure their integrated values match the DHR values corrects propagated measurement errors.

3.3 Limitations of the Model

These measurements generate a reflection model which faithfully represents the actual reflection behavior of many target surfaces. However, a model is an approximation; there are limitations to the types of BRDFs it can accurately represent. Some surface BRDFs simply cannot be accurately approximated with the NEF-BM model using any combination of parameter values. For a model which was initially constructed to represent simple paint reflection, it covers a broad range quite accurately, but care must be used in applying it to new BRDFs.

Another limitation of the NEF-BM model is its inaccuracy at grazing angles. The model was created primarily for remote sensing. Because remote sensing involves interpreting imaged data from aircraft or satellites, accuracy at grazing angles was not a large concern. For use in computer graphics, however, grazing angles are to be expected in every scene.

The problem lies in the fact that the NEF-BM BRDF at grazing incidence, Θ_{gi} , integrated over the reflected hemisphere, results in values greater than one:

$$\begin{aligned} \int_{\Omega_r} \rho(\Theta_{gi}; \Theta_r) d\Omega_r &= \int_{\Omega_r} \rho_{fs}(\Theta_{gi}; \Theta_r) d\Omega_r \\ &+ \int_{\Omega_r} \rho_D d\Omega_r + \int_{\Omega_r} \frac{2\rho_v}{\cos \theta_{gi} + \cos \theta_r} d\Omega_r \\ &= 1 + \pi\rho_D + 4\pi\rho_v \end{aligned}$$

Consequently, light incident at grazing angles will result in more energy reflected than was incident. In practice this has not been a problem mainly because of the foreshortening of incident energy at grazing angles.

4 NEFDS

The NEFDS allows those working in the area of remote sensing to perform complex radiometric calculations while taking into consideration the properties of the target material, the material background, the measuring sensors, and the atmospheric environment. At the heart of this system, and of particular interest to researchers in the field of computer graphics, is the BRDF database. This database contains the BRDFs of hundreds of surfaces from a dozen different categories. In this section, we will present a discussion of the BRDFs available and offer an overview of the relevant portions of the NEFDS.

4.1 Available Material Types

There are currently over 400 materials available in the NEFDS. These materials correspond to a wide variety of objects ranging from dirt to tree canopies. The materials fall into 12 different categories: asphalt, brick, camouflage, composite, concrete, fabric, water, metal, paint, rubber, soil and wood.

Although the number of materials is large and the variety is wide, the selection is limited by two key conditions. One requirement for the inclusion of a material in the database is that the BRDF of the material must be well represented by the NEF-BM BRDF model. For example, the BRDF should not be characterized by extreme anisotropy since the NEF-BM BRDF model only works with isotropic data. (There are in fact some materials included in the NEFDS which are anisotropic. To represent anisotropic material using the NEF-BM model, the material's BRDF is averaged to isotropy.)

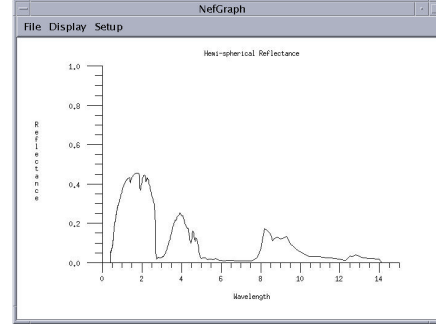


Figure 3: Spectral reflectance function for material 0221, red dirt

The second condition is a result of the data's origin. As mentioned before, the main use of the NEFDS has been in the area of remote sensing. Because of this, the materials in the database were selected on the basis of applicability to that field of work. They include objects which would be viewed from a remote sensor (e.g., a satellite). This still offers a potential wealth of BRDF data for use in computer graphics but is currently constrained to this one area. This however does not restrict future measurement and the inclusion of other materials having more direct applicability to computer graphics.

4.2 Operational Modes

NEFDS can be accessed either through the XWindows interface, *NefMenu*, or by using command line control. A discussion of the full control offered by NEFDS is beyond the scope of this paper (see the NEF Users Guide for detailed discussion [1]). Instead a brief overview will be given.

The model parameter data files used to specify the NEF-BM models are located in the material and groups files [3]. These data files can be viewed interactively through *NefMenu*, providing fully titled and organized fields. *NefMenu* also allows a convenient means of browsing the materials, organized within intuitive groups.

Most interactive programs in NEFDS are available both through the XWindows interface and by command line control. For example, the spectral distribution of hemispherical reflectance can be graphed either through *NefMenu* or through the program *NefPlot*. Such a plot of material number 0221, red dirt, is shown in Figure 3.

The NEFDS may be queried for BRDF values of materials or material groups at ranges of wavelengths for any given geometry. The XWindows control offers interactive control over these parameter settings allowing the user to quickly determine key reflection attributes. The command line counterpart, *BRDF* allows BRDF values to be obtained through a batch process. This is the technique

utilized for the BRDF sampling used to generate BRDF tables for rendering discussed in the next section.

5 Rendering data

In this section, BRDF data obtained from NEFDS was used to render synthetic images. NEFDS was queried and the discretely sampled BRDFs were stored in data files. These files were then used as lookup tables to determine BRDF data during the execution of the rendering program. The program was selected based on the physically accurate global illumination solution that it produces. In addition, the program was extended to efficiently execute Monte Carlo integration over the sampled BRDFs.

5.1 The Alias Method

The inverse method is used in some rendering systems to generate random variates of the probability density function (PDF) corresponding to the BRDF. In general analytically generating random variates of arbitrary PDFs is not possible. However, there are techniques available to create the random variates of an arbitrary probability mass function (PMF), the discretized counterpart to the PDF. Gentle [13] and Knuth [21] both offer a nice overview and discussion of the many methods available. For this article, the alias method proposed by Walker [40, 41, 42] was selected because it generates random variates in constant time.

Many variations of Walker's original model have been proposed offering advantages in one way or another [12, 23, 24, 25, 33, 39]. The modification by Vose [39] was chosen because it allows for initialization in $\mathcal{O}(n)$ time versus the original paper's time of $\mathcal{O}(n \log n)$. Vose also presents several optimizations for memory and execution time although they are not used here.

The alias method is most easily viewed as a form of the rejection method in which the rejected values are recycled into usable data. It requires an initial setup which need only be performed once. After the setup is complete, random variates are created by transforming variates of a uniform distribution.

Consider the random variate \mathbf{X} which can take on any of the t values,

$$X = \{x_1, x_2, \dots, x_t\} \quad (12)$$

with corresponding probability

$$P = \{p_1, p_2, \dots, p_t\}. \quad (13)$$

Ensure that the probabilities form a valid PMF by requiring that $\sum_{c=1}^t p_c = 1$.

In the setup process, the alias method creates two new lists, the rejection list and the alias list. The rejection list is a list of probabilities,

$$R = \{r_1, r_2, \dots, r_t\}, \quad (14)$$

whose elements r_i determine whether or not an alias is to be used. Each probability r_i in the rejection list forms its own PMF together with its complement, $1 - r_i$. The alias list,

$$A = \{a_1, a_2, \dots, a_t\}, \quad (15)$$

consists of t indices which each may take on integer values 1 through t , representing the index of the alias. The setup process is performed only once.

After the setup has been completed, selection of the random variate \mathbf{X} is performed by choosing two uniformly distributed variates, $\mathbf{u} \sim U(0, 1)$ and $\mathbf{i} \sim U_d(1, t)$, where $U_d(1, t)$ is a discrete uniform distribution over the integers 1 through t . \mathbf{X} is given the value x_i if $\mathbf{u} \leq r_i$ (\mathbf{u} was not rejected). \mathbf{X} is given the value x_{a_i} if $\mathbf{u} > r_i$ (\mathbf{u} was rejected so the alias was used). The time required to generate this random variate is equal to the total time required to generate \mathbf{u} , perform a comparison, and then lookup the final value. Assuming these three actions can be performed in constant time, so can the generating step.

Consider a BRDF discretely sampled at s reflected directions $\{\Theta_{r,1}, \Theta_{r,2}, \dots, \Theta_{r,s}\}$, and for each reflected direction, t incident directions, $\{\Theta_{i,1}, \Theta_{i,2}, \dots, \Theta_{i,t}\}$, for a total of st samples. The samples each have a representative solid angle $\Delta\omega_{i,n}$ and a representative projected solid angle $\Delta\Omega_{i,n} = (\Theta_{i,n} \cdot \hat{N})\Delta\omega_{i,n}$ corresponding to $\Theta_{i,n}$, where $1 \leq n \leq t$.

The continuous PDF, $h_{\Theta_{r,m}}$, associated with the reflected direction $\Theta_{r,m}$ then becomes the PMF, $g'_{\Theta_{r,m}}$, by

$$\begin{aligned} g'_{\Theta_{r,m}}(\Theta_{i,n}) &= \int_{\Delta\Omega_{i,n}} h_{\Theta_{r,m}}(\Theta_i) d\Omega_i \\ &\approx h_{\Theta_{r,m}}(\Theta_{i,n}) \Delta\Omega_{i,n} \end{aligned}$$

where $1 \leq m \leq s$ and $1 \leq n \leq t$. The approximation is valid for sufficiently small $\Delta\Omega_{i,n}$ and will be used in generating the random variates in the following section. However, it does require a renormalization to ensure that the sum of each PMF is 1. When renormalized and applied to the BRDF, the PMF becomes

$$\begin{aligned} g_{\Theta_{r,m}}(\Theta_{i,n}) &= \frac{h_{\Theta_{r,m}}(\Theta_{i,n}) \Delta\Omega_{i,n}}{\sum_{c=1}^t h_{\Theta_{r,m}}(\Theta_{i,c}) \Delta\Omega_{i,c}} \\ &= \frac{\rho(\Theta_{i,n}; \Theta_{r,m}) \Delta\Omega_{i,n}}{\sum_{c=1}^t \rho(\Theta_{i,c}; \Theta_{r,m}) \Delta\Omega_{i,c}} \quad (16) \end{aligned}$$

We are now ready to define the lists from (12) and (13) required to generate random variates of the discretely sampled BRDF. For generating random variates of the incident directions corresponding to reflected direction $\Theta_{r,m}$ use

$$X = \{\Theta_{i,1}, \Theta_{i,2}, \dots, \Theta_{i,t}\}$$

and

$$P = \{g_{\Theta_r,m}(\Theta_{i,1}), g_{\Theta_r,m}(\Theta_{i,2}), \dots, g_{\Theta_r,m}(\Theta_{i,t})\}$$

5.2 The Rendering Program

The Radiance Lighting Simulation and Rendering System (Radiance) [28, 44] was used to generate the synthetic images in this work. Radiance is a suite of programs built around an advanced distributed raytracer designed for realistic image synthesis. It was selected because it is a physics based rendering system designed to accurately model the light behavior of a scene using physical units [44]. Using such a system reinforces the validity of the results obtained by this physics based BRDF. Additionally, the source code to Radiance is publicly available [2] and the program is currently in wide use, aiding future work.

Radiance is a distributed ray tracer which utilizes Monte Carlo importance sampling. The direction of the ray is stochastically selected so that the distribution matches the corresponding distribution function associated with the reflection model. This requires an inversion of the distribution function to map a uniformly random number to the desired distribution. Ward’s implementation uses the relatively straightforward implementation of the Gaussian inversion to achieve this. That is, Ward’s reflection model is based on the Gaussian function which provides easy inversion and thus a straightforward method of Monte Carlo integration. This method, however, is not compatible with the BRDFs generated by NEF-BM which are not of a Gaussian nature.

A new shader,³ called iBRDF was built for Radiance using the alias method described in the preceding section. This provides a means for discretized versions of arbitrary isotropic BRDFs to be used in the creation of synthetic images.

Uniform sampling is used to discretize the BRDF. The BRDF is separated into M, N, and R divisions of the reflected θ , incident θ , and incident ϕ respectively. Since the BRDF is isotropic we define the coordinate system so that $\phi_r = 0$. For most of the BRDFs generated using NEFDS, uniform sampling was found to be sufficient. Good results were obtained by sampling at 40 values of θ_i and θ_r between 0 and $\pi/2$ and at 80 values of ϕ_r between 0 and π ⁴. For some BRDFs containing greater variation in the reflectance lobe, a larger number of samples was required.

The iBRDF data is accessible via two different functions. The first is a direct lookup of the BRDF values

given incident and reflected directions. This lookup requires only constant time since the sample grid consist of fixed divisions of θ_i and ϕ_i . For this function, tri-linear interpolation is used to interpolate between the three sampled axes. The second function creates random variates distributed as g_{Θ_i} in (16) using the alias method. Interpolation is performed using a stochastic version of linear interpolation for both the selection of the sampled ϕ_r and the selection of the incident direction within the selected $\Delta\omega_i$.

Radiance separates the BRDF into two components, diffuse and specular. The diffuse component refers only to pure Lambertian while specular refers to the remainder of the BRDF. The advantage of this separation is that the diffusely reflected radiance is slowly varying across the surface. Radiance utilizes this by caching the diffusely reflected radiance values and interpolating between those cached results whenever possible [46, 45]. Additionally, since the diffuse reflectance is pure Lambertian, the cached values are valid for all reflected angles. The above two functions provide the required BRDF information for specular reflectance. The diffuse reflectance is easily obtained by finding the minimum reflectance value of the BRDF for each Θ_r and using this as the diffuse portion of the BRDF, ρ_d . The specular portion of the BRDF used in the alias method is then the difference

$$\rho_{specular}(\Theta_i; \Theta_r) = \rho(\Theta_i; \Theta_r) - \rho_d(\Theta_r).$$

Radiance requires that all spectra be reduced to an RGB tristimulus equivalent, and it performs separate illumination calculations for each of the R, G, and B components. This limits the circumstances under which the color calculations will be exact to those where light reaching the viewing plane has struck only a single colored object (or that object and other spectrally neutral surfaces in the scene). Three separate iBRDF data sets are produced, one for each of the three axes of the RGB color space used in Radiance.

5.3 Rendering NEFDS material

Sampled BRDFs (at $\lambda = 550\text{nm}$) of several NEFDS materials were used to generate images in the modified rendering program. Figure 4 shows three vases modeled with the NEFDS materials Bare Construction Lumber, Gloss Paint on Metal, and Scored Aluminum. Figure 5 is rendered using the same materials and additionally the NEFDS materials Cement, and Weathered Concrete along with the addition of a texture map.

The variance in the reflected light over these five sampled BRDFs is shown even clearer in viewing the BRDFs directly. Figures 6 and 7 show the BRDFs of cement and lumber respectively. Notice the significant difference in geometry that can be characterized by the NEF-BM

³Radiance uses the term material type rather than shader.

⁴Another consequence of isotropic reflection is the need to only sample half the hemisphere. The other half (i.e., $\pi < \phi_r < 2\pi$) can be determined using the laws of reciprocity.



Figure 4: Vases with NEF materials (left to right) Bare Construction Lumber, Gloss Paint on Metal, and Scored Aluminum. Image from [43].




Figure 5: Cubes with textured NEF materials (top row) Cement, Gloss Paint on Metal, (bottom row) Bare Construction Lumber, Scored Aluminum, and Weathered Concrete. Image from [43].

model. This strength (the ability to capture a wide variety of BRDF distributions) coupled with the systematic method of measurement make for a very powerful rendering tool.

5.4 Fitting New Models to NEF-BM

Not only does NEFDS offer hundreds of pre-existing materials, it also allows new materials to be added. The measured BRDF values of two metallic paint samples were used to create two new modeled BRDFs using the NEF Beard-Maxwell model. The two samples, termed fine and coarse, were respectively dominated by small and large metallic flakes. The paint with a greater number of fine flakes had a larger diffuse component due to more edge scattering. This is easily captured by the NEF-BM model as can be seen in the rendered image of Figure 8. Measurements were performed at the National Institute of Standards and Technology (NIST) using the reference spectrogoniophotometer, the Spectral Tri-function Automated Reference Reflectometer (STARR) [34].

The first measurement⁵ discussed here is the near zero-bistatic BRDF measurement. This measurement is used to derive the distribution of micro-facets, which are responsible for the first surface reflection in the Beard-Maxwell model. Measurement of the near zero-bistatic BRDF is performed in a single plane, from the surface normal to grazing angles. Source and detector are held at a fixed six degree separation (near zero) in this plane of




Figure 6: BRDF of Cement

measurement.

Figure 9 shows the measured versus modeled BRDF values at the near zero-bistatic angles for the fine and coarse metallic paint samples. As can be seen there is a very good fit to the measured data at the points outside the specular peak. The crossover point at about 18 degrees is matched for both measured and modeled data.

In addition to the near zero-bistatic measurements, the specular BRDF measurements were performed using parallel polarized source and receiver over $10 \text{ deg} \leq \theta_i \leq 80 \text{ deg}$ with $\phi_i = 0 \text{ deg}$ and $\phi_r = 180 \text{ deg}$. As mentioned earlier, for each sampled incident direction, the receiver direction was varied about the mirror direction in the plane of incidence by $\theta_i - 5 \text{ deg} \leq \theta_r \leq \theta_i + 5 \text{ deg}$. This provided the means by which the complex index of refraction was determined, used in determining the Fresnel reflection. Lastly, the shadowing and obscuration

⁵Measurements were performed in the four polarization states, $\rho_{\perp\perp}$, $\rho_{\perp\parallel}$, $\rho_{\parallel\perp}$, and $\rho_{\parallel\parallel}$. Unless noted otherwise, the measured values listed are the calculated unpolarized values, computed from these measured polarized values.




Figure 7: BRDF of Bare Construction Lumber




Figure 8: Coarse and fine metallic paint on vases




Figure 9: Comparison of measured and modeled six degree bistatic angle BRDF values

Table 1: Comparison of measured versus modeled out-of-plane BRDF values for coarse metallic flake sample.

θ_i (degs.)	θ_r (degs.)	ϕ_r (degs.)	Coarse metallic flake BRDF	
			Measured (1/sr)	Modeled (1/sr)
-30	45	45	0.269	0.255
-30	45	90	0.0635	0.0698
-30	45	135	0.0389	0.0405
-45	45	45	0.207	0.223
-45	45	90	0.0469	0.0499
-45	45	135	0.0331	0.0332
-60	45	45	0.150	0.173
-60	45	90	0.0401	0.0469
-60	45	135	0.0325	0.0316

Table 2: Comparison of measured versus modeled out-of-plane BRDF values for fine metallic flake sample.

θ_i (degs.)	θ_r (degs.)	ϕ_r (degs.)	Fine metallic flake BRDF	
			Measured (1/sr)	Modeled (1/sr)
-30	45	45	0.254	0.321
-30	45	90	0.1023	0.1001
-30	45	135	0.0627	0.0622
-45	45	45	0.210	0.295
-45	45	90	0.0727	0.0769
-45	45	135	0.0479	0.0454
-60	45	45	0.179	0.240
-60	45	90	0.0581	0.0641
-60	45	135	0.0386	0.0377

terms, τ and Ω , were determined with the fixed source, varying receiver measurements.

The model parameters for the NEF Beard-Maxwell reflection model are selected on the basis of in-plane BRDF measurements only. In order to verify that the NEF Beard-Maxwell model is in fact an appropriate model to represent the BRDF of the surface, validation of out-of-plane BRDF values should be performed. Out-of-plane measurements were made at a variety of incident and reflected directions for both coarse and fine metallic flake samples—the results are listed in Tables 1 and 2. As can be seen from these figures, there is a very good correspondence between the measured and modeled BRDF values, indicating that the NEF Beard-Maxwell model is appropriate for representing these metallic surfaces.

6 Conclusion

The NEFDS provides a collection of surface reflection data that can be useful for work in realistic image synthesis. Although the database was developed for application in the remote sensing field, it contains materials of importance to computer graphics researchers. The database has a good user interface, it is well documented, and it can be extended with additional measurements. The protocol for performing these measurements is well defined and it makes use of existing radiometric instruments.

The NEFDS incorporates a modified form of the Beard-Maxwell reflection model. This is a physics based reflection model that can be fit to measured data. This gives it the advantage of a first principles physics derivation coupled with a measurement scheme for determining the model parameters. The key insight in the model is that all first surface reflectances can be calculated by measuring zero-bistatic first surface reflectances. The principal limitation of the model is its inaccuracy at grazing angles.

The acceptance of sophisticated reflectance models in computer graphics depends upon the availability of data necessary to use of the models. Rendering systems such as Radiance are capable of producing photo realistic pictures. With the addition of real-time shading techniques, even personal computers can make acceptable interactive pictures of surfaces with complex reflectance properties. Just as the need for geometric data has spawned a 3D modeling industry, databases containing the reflectance characteristics of common materials will be of increasing importance in computer graphics. NEFDS is an example of such a database.

7 Acknowledgments

The authors would like to acknowledge the significant contributions made by one of the original co-developers of the NEF-BM model and the NEF database, Michael Metzler of ISciences LLC. Metzler wrote software that allowed us to directly access the data in the NEF database without going through the user interface. Metzler was also instrumental in specifying the protocol that was employed to make the measurements as well as performing the parameter estimation from the measurement data described in Section 5.4. Maria Nadal of the National Institute of Standards and Technology (NIST) acquired the data, some of which is shown in Figure 9, that allowed us to synthesize pictures of the metal flake samples. The scenes shown in Figures 4 and 5 were modeled and rendered by Peter Walker. Funding for this research was provided by NIST under the program administration of Fern Hunt.

8 References

- [1] *NEF User Guide*, 9.1 edition. <http://math.nist.gov/~FHunt/appearance/nefds.html>, Accessed March 20, 2002.
- [2] Radiance home page. <http://radsite.lbl.gov/radiance>. Accessed March 20, 2001.
- [3] *NEF Specifications*, July 1996. ORD 258-96, <http://math.nist.gov/~FHunt/appearance/nefds.html>, Accessed March 20, 2002.
- [4] *Nonconventional Exploitation Factors (NEF) Modeling*, August 1996. ORD 257-96, <http://math.nist.gov/~FHunt/appearance/nefds.html>, Accessed March 20, 2002.
- [5] Michael Ashikhmin, Simon Premoze, and Peter S. Shirley. A microfacet-based brdf generator. *Computer Graphics (Proceedings of SIGGRAPH 2000)*, pages 65–74, July 2000.
- [6] Petr Beckmann and André Spizzichino. *The Scattering of Electromagnetic Waves from Rough Surfaces*. Pergamon Press, Oxford, England, 1963.
- [7] James F. Blinn. Models of light reflection for computer synthesized pictures. *Computer Graphics (Proceedings of SIGGRAPH 77)*, 11(2):192–198, July 1977.
- [8] James F. Blinn. Light reflection functions for simulation of clouds and dusty surfaces. *Computer Graphics (Proceedings of SIGGRAPH 82)*, 16(3):21–29, July 1982.
- [9] Brian Cabral, Nelson Max, and Rebecca Springmeyer. Bidirectional reflection functions from surface bump maps. *Computer Graphics (Proceedings of SIGGRAPH 87)*, 21(4):273–281, July 1987.
- [10] Robert L. Cook and Kenneth E. Torrance. A reflectance model for computer graphics. *Computer Graphics (Proceedings of SIGGRAPH 81)*, 15(3):307–316, August 1981.
- [11] Kristin J. Dana, Bram van Ginneken, Shree K. Nayar, and Jan J. Koenderink. Reflectance and texture of real-world surfaces. *ACM Transactions on Graphics*, 18(1):1–34, January 1999.
- [12] George S. Fishman and L. Stephen Yarberr. Generating a sample from a k -cell table with changing probabilities in $O(\log_2 k)$ time. *ACM Transactions on Mathematical Software*, 19(2):257–261, 1993.
- [13] James E. Gentle. *Random Number Generation and Monte Carlo Methods*. Springer-Verlag New York, Inc., New York, U.S.A., 1998.
- [14] Jay S. Gondek, Gary W. Meyer, and Jonathan G. Newman. Wavelength dependent reflectance functions. *Computer Graphics (Proceedings of SIGGRAPH 94)*, pages 213–220, July 1994.
- [15] Bruce W. Hapke. A theoretical photometric function for the lunar surface. *Journal of Geophysical Research*, 68(15):4571–4586, 1963.
- [16] Xiao D. He, Kenneth E. Torrance, François X. Sillion, and Donald P. Greenberg. A comprehensive physical model for light reflection. *Computer Graphics (Proceedings of SIGGRAPH 91)*, 25(4):175–186, July 1991.
- [17] Eugene Hecht. *Optics*, chapter 4: The Propagation of Light, pages 85–147. Addison Wesley Longman, Inc., third edition, 1998.
- [18] Wolfgang Heidrich and Hans-Peter Seidel. Realistic, hardware-accelerated shading and lighting. *Computer Graphics (Proceedings of SIGGRAPH 99)*, pages 171–178, August 1999.
- [19] L. Henyey and J. Greenstein. Diffuse reflection in the galaxy. *Journal of Astrophysics*, 93:70–77, 1941.

- [20] Jan Kautz and Michael D. McCool. Interactive rendering with arbitrary brdfs using separable approximations. *Tenth Eurographics Workshop on Rendering*, pages 281–292, June 1999.
- [21] Donald E. Knuth. *Seminumerical Algorithms*, volume 2 of *The Art of Computer Programming*. Addison-Wesley Longman, Reading, Massachusetts, U.S.A., second edition, 1998.
- [22] J. J. Koenderink, A. J. van Doorn, and M. Stavridi. Bidirectional reflection distribution function expressed in terms of surface scattering modes. *European Conference on Computer Vision*, pages 28–29, 1996.
- [23] R. A. Kronmal and A. V. Peterson. The alias and alias-rejection-mixture methods for generating random variables from probability distributions. In *Proceedings of the 1979 Winter Simulation Conference*, pages 269–280. Institute of Electrical and Electronics Engineers, 1979.
- [24] R. A. Kronmal and A. V. Peterson. On the alias method for generating random variables from a discrete distribution. *The American Statistician*, 33:214–218, 1979.
- [25] R. A. Kronmal and A. V. Peterson. A variant of the acceptance-rejection method for computer generation of random variables. *Journal of the American Statistical Association*, 76:446–451, 1981.
- [26] Eric P. F. Lafortune, Sing-Choong Foo, Kenneth E. Torrance, and Donald P. Greenberg. Non-linear approximation of reflectance functions. *Computer Graphics (Proceedings of SIGGRAPH 97)*, pages 117–126, August 1997.
- [27] Paul Lalonde and Alain Fournier. A wavelet representation of reflectance functions. *IEEE Transactions on Visualization and Computer Graphics*, 3(4):329–336, October–December 1997.
- [28] Greg Ward Larson and Rob Shakespeare. *Rendering with Radiance, The Art and Science of Lighting Visualization*. Morgan Kaufmann Publishers, Inc., San Francisco, U.S.A., 1998.
- [29] Gregory J. Ward Larson. Measuring and modeling anisotropic reflection. *Computer Graphics (Proceedings of SIGGRAPH 92)*, 26(2):265–272, July 1992.
- [30] J. R. Maxwell, J. Beard, S. Weiner, and D. Ladd. Bidirectional reflectance model validation and utilization. Technical Report AFAL–TR–73–303, Environmental Research Institute of Michigan (ERIM), October 1973.
- [31] F. E. Nicodemus, J. C. Richmond, J. J. Hsia, I. W. Ginsberg, and T. Limperis. Geometric considerations and nomenclature for reflectance. Technical Report MN-160, U.S. Department of Commerce, National Bureau of Standards, October 1977.
- [32] Michael Oren and Shree K. Nayar. Generalization of the lambertian model and implications for machine vision. *International Journal of Computer Vision*, 14:227–251, 1995.
- [33] A. V. Peterson and R. A. Kronmal. On mixture methods for the computer generation of random variables. *The American Statistician*, 36:184–191, 1982.
- [34] James E. Proctor and P. Yvonne Barnes. Nist high accuracy reference reflectometer-spectrophotometer. *Journal of Research of the National Institute of Standards and Technology*, 101(5):619–627, September–October 1996.
- [35] Peter Schröder and Wim Sweldens. Spherical wavelets: Efficiently representing functions on the sphere. *Computer Graphics (Proceedings of SIGGRAPH 95)*, pages 161–172, August 1995.
- [36] Jos Stam. Diffraction shaders. *Computer Graphics (Proceedings of SIGGRAPH 99)*, pages 101–110, August 1999.
- [37] K. E. Torrance and E. M. Sparrow. Theory for off-specular reflection from roughened surfaces. *Journal of Optical Society of America*, 57(9):1105–1114, 1967.
- [38] T. Trowbridge and K. Reitz. Average irregularity representation of roughened surfaces. *Journal of the Optical Society of America*, 65(5):531–536, 1975.
- [39] Michael D. Vose. A linear algorithm for generating random numbers with a given distribution. *IEEE Transactions on Software Engineering*, 17:972–975, 1991.
- [40] A. J. Walker. Fast generation of uniformly distributed pseudorandom numbers with floating point representation. *Electronics Letters*, 10(41):553–554, 1974.
- [41] A. J. Walker. New fast method for generating discrete random numbers with arbitrary frequency distribution. *Electronics Letters*, 10(8):127–128, 1974.
- [42] A. J. Walker. An efficient method for generating discrete random variables with general distributions. *ACM Transactions on Mathematical Software*, 3:253–256, 1977.
- [43] Peter Andrew Walker. A visualization system for bidirectional reflectance distribution functions. Master’s thesis, University of Oregon, 1999.
- [44] Gregory J. Ward. The radiance lighting simulation and rendering system. *Computer Graphics (Proceedings of SIGGRAPH 94)*, pages 459–472, July 1994.
- [45] Gregory J. Ward and Paul S. Heckbert. Irradiance gradients. In *Proceedings of the Third Annual Eurographics Workshop on Rendering*, pages 85–98, Bristol, U.K., 1992. Springer-Verlag.
- [46] Gregory J. Ward, Francis M. Rubinstein, and Robert D. Clear. A ray tracing solution for diffuse interreflection. *Computer Graphics (Proceedings of SIGGRAPH 88)*, 22(4):85–92, August 1988.
- [47] Stephen H. Westin, James R. Arvo, and Kenneth E. Torrance. Predicting reflectance functions from complex surfaces. *Computer Graphics (Proceedings of SIGGRAPH 92)*, 26(2):255–264, July 1992.