

Université de Montréal

Theory and Numerical Integration of Subsurface Light Transport

par
David Milaenen

Département d'informatique et de recherche opérationnelle
Faculté des arts et des sciences

Mémoire présenté à la Faculté des études supérieures et postdoctorales
en vue de l'obtention du grade de Maître ès sciences (M.Sc.)
en informatique

Août, 2015

© David Milaenen, 2015

Résumé

En synthèse d'images, reproduire les effets complexes de la lumière sur des matériaux transluminescents, tels que la cire, le marbre ou la peau, contribue grandement au réalisme d'une image. Malheureusement, ce réalisme supplémentaire est couteux en temps de calcul. Les modèles basés sur la théorie de la diffusion visent à réduire ce coût en simulant le comportement physique du transport de la lumière sous surfacique tout en imposant des contraintes de variation sur la lumière incidente et sortante. Une composante importante de ces modèles est leur application à évaluer hiérarchiquement l'intégrale numérique de l'illumination sur la surface d'un objet.

Cette thèse révisé en premier lieu la littérature actuelle sur la simulation réaliste de la transluminescence, avant d'investiguer plus en profondeur leur application et les extensions des modèles de diffusion en synthèse d'images. Ainsi, nous proposons et évaluons une nouvelle technique d'intégration numérique hiérarchique utilisant une nouvelle analyse fréquentielle de la lumière sortante et incidente pour adapter efficacement le taux d'échantillonnage pendant l'intégration. Nous appliquons cette théorie à plusieurs modèles qui correspondent à l'état de l'art en diffusion, octroyant une amélioration possible à leur efficacité et précision.

Fourier, traitement du signal, transluminescence, BSSRDF, structure d'accélération, illumination globale, synthèse d'images

Abstract

In image synthesis, reproducing the complex appearance of objects with subsurface light scattering, such as wax, marble and skin, greatly contributes to the realism of an image. Unfortunately, this added realism comes at a high computational cost. Models based on diffusion theory aim to reduce this computational cost by simulating the physical behaviour of subsurface light scattering while imposing smoothness constraints on the incident and outgoing light fields. An important component of these models is how they are employed to hierarchically evaluate the numerical integral of lighting over the surface of an object.

This thesis will first review the existing literature on realistic subsurface lighting simulation, before investigating in more depth the application and extension of modern diffusion models in image synthesis. In doing so, we propose and evaluate a new hierarchical numerical integration technique that uses a novel frequency analysis of the incident and outgoing light fields to reliably adapt the sampling rate during integration. We realize our resulting theory in the context of several state-of-the-art diffusion models, providing a marked improvement in their efficiency and accuracy.

Fourier analysis, signal processing, subsurface scattering, BSSRDF, acceleration data structure, global illumination, image synthesis

Contents

Résumé	ii
Abstract	iii
Contents	iv
List of Tables	viii
List of Figures	ix
Acknowledgments	xiv
Chapter 1: Introduction	1
1.1 Motivation	1
1.2 Objectives	2
1.2.1 Outline	3
Chapter 2: Light Transport	5
2.1 Radiometry	5
2.2 Rendering Equation	7
2.3 Algorithms	9
2.3.1 Path Tracing	10

2.3.2	Photon Mapping	11
2.4	Radiative Transport Equation	12
2.5	Volume Rendering Equation	14
2.6	Algorithms	15
2.6.1	Path Tracing	15
2.6.2	Photon Mapping	16
Chapter 3:	State of the Art	17
3.1	Bidirectional Surface Scattering Reflectance Distribution Function	17
3.1.1	Decomposition of S	18
3.1.2	Diffusion Theory	18
3.1.3	Searchlight Problem	21
3.1.4	Dipole	22
3.1.5	Quadpole	23
3.1.6	Multipole	24
3.1.7	Improving the Accuracy of Dipole Approximations	26
3.1.8	Quantized Diffusion	27
3.1.9	Directional Dipole	30
3.1.10	Related Techniques	32
3.2	Numerical Integration for Realistic Image Synthesis	33
3.2.1	Monte Carlo	33

3.2.2	Two-pass Hierarchical Technique	33
3.2.3	Single-pass Hierarchical Technique	35
3.3	Frequency Analysis	36
Chapter 4: A Frequency Analysis and Dual Hierarchy for Efficient Rendering of Subsurface Scattering		39
4.0	Abstract	39
4.1	Introduction	40
4.2	Previous Work	41
4.2.1	Overview	43
4.3	Fourier Analysis	44
4.3.1	Fourier Transform of a BSSRDF	44
4.3.2	Outgoing Radiance Bandwidth Computation	47
4.4	Hierarchical Approach	49
4.4.1	Hierarchical Surface Integration	49
4.4.2	Dual Hierarchy for Pixel-Surface Integration.	50
4.5	Implementation	51
4.6	Results and Discussions	52
4.7	Conclusion	56
Chapter 5: Conclusion		58

Bibliography 60

List of Tables

1.I	Nomenclature used throughout this thesis.	4
-----	---	---

List of Figures

1.1	Many techniques are used to produce realistic images of materials that undergo subsurface scattering, such as (in reading order) marble [DJ05], skin, wax [HCJ13b], and milk [DI11]. Properties like the (small) size of a character [JB02] can also be depicted with translucence in non-photorealistic images.	2
2.1	The planar angle (left) is the length of the arc on a unit circle and the solid angle (right) is the area on a unit sphere, both with respect to a central point and from a projected object.	5
2.2	The BRDF mathematically expresses the relation between the irradiance and the reflected radiance at a given point on an object. Conceptually, the BRDF describes the reflection behaviour of a surface.	9
2.3	Path tracing (left) computes the color of a pixel by averaging the contribution of many light carrying paths through the scene, shot from the point of view. Photons are accumulated on the surfaces in photon mapping (right) and used as a measure of the incident radiance at rendering time.	10
2.4	The <i>Cornell box</i> scene rendered with path tracing (left) and photon mapping (right). Notice the high frequency noise in the shadowy parts of the first image and the constant low frequency noise on the rectangular prisms' faces in the second one.	11
2.5	Ecosystem scene rendered without (left) and with (right) participating media. Additional effect makes the image look substantially more realistic. From <i>pbrt</i> [PH10].	12

3.1	As opposed to the BRDF (Figure 2.2), the BSSRDF does not assume that light enters and leaves at the same surface point. It describes the relation between the irradiance at an incident point and the reflected radiance at the exit point on a surface bounding a volumetric scattering medium. . . .	18
3.2	The searchlight problem consists of an infinitesimal beam of light striking a semi-infinite slab at normal incidence. The scattered distribution of light is a 1D radially symmetric function, denoted R	21
3.3	The dipole method consists of two point light sources, a positive (bottom) and a negative (top), placed in such a way that the boundary condition (<i>e.g.</i> , Equation 3.15) is satisfied.	22
3.4	The basic dipole configuration is mirrored in the case of a $\pi/2$ angle with an adjacent face, forming a quadpole, to satisfy the boundary condition. . .	24
3.5	The multipole configuration is an extension to regular dipole that may apply to thin slabs and multilayered media.	25
3.6	Comparison of quantized diffusion with other diffusion-based approaches for a semi-infinite material with a white light at normal incidence. The QD model retains high-frequency details present in a ground truth Monte Carlo simulation. From d'Eon and Irving [DI11].	30
3.7	Ray sources replace the point sources in this augmented diffusion configuration. The refracted direction ω_r , which stands as the positive source, is mirrored by a modified tangent plane to obtain the negative ray source ω_v . This model handles directionality and relaxes the assumption of flat boundaries.	30
3.8	Poisson disk samples on a complex 3D object. From Bowers et al. [BWWM10].	34

3.9	In reading order: (1) the naive method where all the irradiance samples from all the lights are computed per shading point e , (2) the method from Jensen and Buhler [JB02], (3) clustering the point lights and using the method from Jensen and Buhler [JB02], and (4) the links between the light clusters and the irradiance samples clusters are also clustered, allowing a faster evaluation. From Arbree et al. [AWB08].	36
3.10	The Fourier transform of a signal captures its frequency variations. Low spatial frequency regions of the image are concentrated around the origin of the Fourier domain while high spatial frequency regions may span the entire Fourier domain. From Belcour [Bel12].	37
4.1	We introduce a hierarchical method to accelerate the rendering of multiple scattering with BSSRDFs (IV). We overview our approach in the PICNIK scene, above: our frequency analysis of BSSRDFs allows us to predict the screen-space sampling rates (II) which are used to devise bounds on the variation of outgoing radiance. These bounds allow us to efficiently integrate the BSSRDF using a dual hierarchy over clustered illumination samples (I) and shading points (i.e., pixels; III).	40
4.2	We sample incident illumination over the object (a) according to its subsurface scattering properties and construct two spatial acceleration structures: one over these samples (c) and one over pixels (d). To render, we simultaneously traverse the trees (e), using our outgoing radiance bandwidth estimate s_p (b) to stop the tree traversal and shade super-pixels of area A	43

4.3	Assuming that the incoming light-field has infinite bandwidth, we conservatively estimate the bandwidth of the outgoing light-field $[B_s, B_\theta]$ as the bandwidth of the BSSRDF along the outgoing spatial positions and directions (a). The interaction with the material limits the spectrum of the local light-field by the BSSRDF spatial and angular bandwidth (b). To estimate the bandwidth at the camera position, we first shear spatially the spectrum to account for curvature (c). Then, we scale spatially to account for foreshortening (d) and finally shear angularly the spectrum to account for transport (e).	45
4.4	First row: The sampling rate s_p computed from the screen-space bandwidth estimation. Second row: Pixel areas from which the sampling rate predicts an adequate approximation of the outgoing radiance variation.	48
4.5	We compare our approach (red) to Jensen and Bulher [JB02] (blue) for different settings of ϵ . We highlight the $\epsilon \in [0.01, 0.2]$ values and consistently reach equal-quality (measured in RMSE; y -axis) in less render time (in seconds; x -axis).	52
4.6	The TOAD scene has a bumpy geometry with detailed textures. We compare the difference images of the multiple scattering term against the ground truth for an equal rendering time (196s). The difference images are scaled by 50 for visualization. Our approach achieves more accurate estimation than the single-tree in the same rendering time.	53
4.7	The BUNNY scene. We compare the difference images of the multiple scattering term against the ground truth for an equal rendering time (60s). The difference images are scaled by 200 for visualization. In this example, our approach removes artifacts under the tail and reduces moiré patterns present in the single-tree approach.	54

4.8	The PICNIK scene challenges the assumptions of our work, and we only obtain equal-quality benefit at lower rendering times (albeit enough for visual convergence).	55
-----	--	----

Acknowledgments

First of all, I want to thank my supervisor Derek Nowrouzezahrai for his guidance and support all the way through my Master's degree. I also want to thank my unofficial co-supervisor Pierre Poulin and his open door, (almost) always willing to give a few minutes of his time. I am grateful to the members of my jury, Max Mignotte, Neil Stewart, and Derek Nowrouzezahrai. I also want to thank Laurent Belcour for all the knowledge he provided and the time he gave to help me solve my numerous problems. The co-authors of my research paper all deserve a special mention: Laurent Belcour, Jean-Philippe Guertin, Toshiya Hachisuka and Derek Nowrouzezahrai. Derek and I graciously acknowledge funding from the *Fonds de recherche sur la nature et les technologies*, through their *Établissement de nouveaux chercheurs universitaires* program (2014-NC-173734).

These two years would not have been the same without the help of my colleagues and friends Gilles-Philippe, David, Adrien, Jean-Philippe, Cynthia, Olivier, Chaitanya, Aude, and Kim: thank you for the CG discussions, the tips, and the constant encouragements. I will never forget the lab events and the great time I had with you and all my wonderful colleagues of the LIGUM: Joël, Étienne, Sonia, Yangyang, Cihan, Melino, Alexandre, Antoine, Arnaud, Nicolas, Florian, Mahdi, Guofu, Jonathan, Omar, Nicolas, Mohamed, and Dorian. Finally, I want to thank my parents, Ginette and André, who were always a good moral support. This work would not have been the same without all of you!

Chapter 1

Introduction

The interaction of light in our environment produces many marvelous effects. Among these effects is the scattering of light through liquids, precious gems, wax, leaves, flesh, and skin. The properties of these materials are such that the light enters an object and does not exit it before scattering irregularly within its volume, which complicates the simulation of these effects in computer graphics (CG). To sidestep this challenge, CG has traditionally ignored these effects for artistic and/or computational reasons: as traditional rendering methods only consider light interactions at one surface point at a time, a more general function is necessary to describe the way light scatters inside a volumetric medium. This area of interest gained attention in the past 15 years as the field has matured to consider more complex and realistic light transport effects.

1.1 Motivation

Accurately reproducing light's interaction with non-opaque objects is fundamental for creating realistic images. It adds pleasant details to objects, as observed in the real world. The appearance of skin, for instance, is primarily due to subsurface scattering: only 6% of light is reflected directly off the skin, whereas the remaining 94% is the result of subsurface scattering [KB04]. Needless to say, this effect is very important for any idealistically modelled character. It is this quest for realism that has pushed scientists to develop and improve simulation, models, data structures, and techniques over the recent years, aiming for accuracy and efficiency.

Accounting for the transparency of materials augments the dimensionality of the underlying physically-based simulation of light transport, making it impractical to explicitly compute all light interactions in the environment. In order to remain practical and efficient, it is necessary to use adaptive methods. Doing so reduces the computational cost of generating images by restraining the simulation to the regions of interest. Our work will aim to more accurately determine these regions of interest during adaptive integration.



Figure 1.1 – Many techniques are used to produce realistic images of materials that undergo subsurface scattering, such as (in reading order) marble [DJ05], skin, wax [HCJ13b], and milk [DI11]. Properties like the (small) size of a character [JB02] can also be depicted with translucence in non-photorealistic images.

1.2 Objectives

The growing theoretical understanding of light transport in dense scattering media has led to the development of mathematical models dedicated to such materials. This thesis presents the following contributions:

1. An extension of the frequency analysis of Durand et al. [DHS⁺05], Bagher et al. [BSS⁺12], and Belcour et al. [BBS14] suited to dense volumetric media.

2. A numerical technique to estimate the frequency spectrum of scattering light in such media, which we use to determine the spatial and angular variation of outgoing radiance over the surface of a translucent object, and which is capable of supporting any underlying dipole model.
3. The application of a dual-tree structure to accelerate the spatial integration step in joint image- and object-space, adaptively evaluating the translucent model for the final rendering.

1.2.1 Outline

Prior to detailing our technical contributions, we will review the theory of light transport and how it is modelled in the context of image synthesis. This thesis is structured as follows:

Introduction to Light Transport: We will introduce concepts and equations of light transport, based on a set of radiometric formulations describing visible light. Given the simplified formulation for light reflected directly off surfaces, we will detail the more involved simulation of light scattering through a volumetric media.

Introduction to Computer Rendering Algorithms: The evolution from mathematical models to practical application in computer graphics is subject to multiple restrictions (physical material, complexity, lack of information in the model). The rendering simulation process is complex and we will introduce its algorithmic details along with the theoretical aspects of light transport concepts.

Subsurface Scattering and Adaptive Rendering: The state of the art in rendering dense translucent media will be thoroughly detailed, identifying the pros and cons of current approaches. An alternative space, the Fourier domain, will also be studied in order to efficiently predict the variation of light. We will show how this prediction in that space can be used as an oracle to perform adaptive rendering in a hierarchical integration scheme, which will drive the contributions presented in this dissertation.

Symbol	Description
\mathbf{x}	Position
\mathbf{n}	Normalized surface normal
$\boldsymbol{\omega}$	Normalized direction, always pointing away from the surface, $\boldsymbol{\omega} \cdot \mathbf{n} > 0$
$\int_{2\pi}$	Integral over the hemisphere of the surface normal
\mathcal{A}	Surface area
L_e	Emitted radiance
L_r	Reflected radiance
L_i	Incident radiance
L_o	Outgoing radiance
$f_r(\mathbf{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o)$	BRDF
$S(\mathbf{x}_i, \boldsymbol{\omega}_i, \mathbf{x}_o, \boldsymbol{\omega}_o)$	BSSRDF
$p(\mathbf{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o)$	Phase function
σ_a	Absorption coefficient
σ_s	Scattering coefficient
σ_t	Extinction coefficient $\sigma_t = \sigma_s + \sigma_a$
α	Albedo $\alpha = \sigma_s / \sigma_t$
g	Mean cosine of the scattering angle
σ'_s	Reduced scattering coefficient $\sigma'_s = (1 - g)\sigma_s$
σ'_t	Reduced extinction coefficient $\sigma'_t = \sigma'_s + \sigma_a$
α'	Reduced albedo $\alpha' = \sigma'_s / \sigma'_t$
D	Diffusion constant $D = 1/3\sigma'_t$
σ_{tr}	Effective transport coefficient $\sigma_{tr} = \sqrt{\sigma_a/D}$
C_1, C_2	Approximations of the Fresnel reflectance, define in d'Eon [d'E12]
C_ϕ	$C_\phi = \frac{1}{4}(1 - 2C_1)$
C_E	$C_E = \frac{1}{2}(1 - 3C_2)$

Table 1.I – Nomenclature used throughout this thesis.

Chapter 2

Light Transport

2.1 Radiometry

Radiometry provides a set of mathematical formalisms and tools to describe the evolution of electromagnetic radiation. Visible light being one of these radiations, we will use a radiometric formulation that encompasses light transport. This theory is required for further study of the light transport process and its application in computer graphics.

Radiant Flux: Also referred to as *power* or simply *flux*, the *radiant flux* represents the amount of energy, measured in *Watts* [$W = J \cdot s^{-1}$], passing through a surface per unit time. It includes the energy emitted by a light source as well as the energy reflected, transmitted or received by the surface. It is denoted by the symbol Φ .

Solid Angle: The extension of the two-dimensional *planar angle* to the third dimension is called the *solid angle*. The planar angle, measured in *radians* [rad], is the length of the arc subtended by the projection of an object onto the unit circle. In the same way, the solid angle, illustrated in Figure 2.1 and measured in *steradians* [sr], is the area subtended by an object projected onto the unit sphere, through the sphere's center.




Figure 2.1 – The planar angle (left) is the length of the arc on a unit circle and the solid angle (right) is the area on a unit sphere, both with respect to a central point and from a projected object.

Radiant Intensity: The flux per unit solid angle from a single direction is called *radiant intensity* or simply *intensity*. It is measured in Watts per steradian [$W \cdot sr^{-1}$] and is denoted by the symbol I . It is expressed in terms of flux as

$$I(\boldsymbol{\omega}) = \frac{d\Phi(\boldsymbol{\omega})}{d\boldsymbol{\omega}}. \quad (2.1)$$

Since a unit sphere's total surface area is 4π , an isotropic point light source would have a uniform radiant intensity of $I = \Phi/4\pi$.

Radiance: Denoted by the symbol L and measured in Watts per steradian per square meter [$W \cdot sr^{-1} \cdot m^{-2}$], the *radiance* is the radiant intensity per unit projected area along its direction. In terms of flux, it is expressed as

$$L(\mathbf{x}, \boldsymbol{\omega}) = \frac{d^2\Phi(\mathbf{x}, \boldsymbol{\omega})}{d\boldsymbol{\omega}d\mathcal{A}(\mathbf{n} \cdot \boldsymbol{\omega})}, \quad (2.2)$$

with the normal \mathbf{n} of the surface subjected to the projection. We can also define flux in terms of the radiance by integrating over the surface area \mathcal{A} and the hemisphere as

$$\Phi = \int_{\mathcal{A}} \int_{2\pi} L(\mathbf{x}, \boldsymbol{\omega})(\mathbf{n} \cdot \boldsymbol{\omega})d\boldsymbol{\omega}d\mathcal{A}. \quad (2.3)$$

Irradiance and Radiant Exitance: The flux per unit area arriving at a surface is called *irradiance*. It is denoted by the symbol E and is measured in Watts per square meter [$W \cdot m^{-2}$]. It can be expressed in terms of flux as

$$E(\mathbf{x}) = \frac{d\Phi(\mathbf{x})}{d\mathcal{A}(\mathbf{x})}. \quad (2.4)$$

The *radiant exitance* represents the flux per unit area leaving a surface and is denoted by the symbol M , with the same units as the irradiance. In order to differentiate the radiance arriving at a surface and the radiance leaving a surface, we will denote L_i as the *incident* radiance and L_o as the *exitant*, or outgoing, radiance. Thus, in terms of the

radiance, the irradiance and radiant exitance can be distinguished by

$$E(\mathbf{x}, \boldsymbol{\omega}) = \int_{2\pi} L_i(\mathbf{x}, \boldsymbol{\omega})(\mathbf{n} \cdot \boldsymbol{\omega})d\boldsymbol{\omega} \quad (2.5)$$

$$M(\mathbf{x}, \boldsymbol{\omega}) = \int_{2\pi} L_o(\mathbf{x}, \boldsymbol{\omega})(\mathbf{n} \cdot \boldsymbol{\omega})d\boldsymbol{\omega}. \quad (2.6)$$

Spectroradiometry: The previous definitions can be extended to express variations as a function of wavelength. For instance, the radiance per wavelength is called *spectral radiance*. In image synthesis, we are interested in the range of wavelengths visible to the human eye.

Luminance: The relative sensitivity of the human eye to wavelengths is described as *luminance*. It accounts for the fact that, for instance, an amount of energy in the green wavelengths will appear brighter to humans than the same amount of energy in the blue ones.

Polygon Mesh: A polygon mesh is a collection of vertices, edges and faces which defines the shape of a polyhedral object. The sub-field of computer graphics that studies polygon meshes is large, but we need to be conscious that our 3D objects, and more specifically their surfaces, are modelled with such constructions of meshes.

Pixel: A *picture element*, or *pixel*, is a point (square) in an image where data is normally represented as values on a regular grid. Physically, it is the smallest point of color visible on a monitor. At a high level, rendering is the process of passing from a three-dimensional domain to the pixel domain, *i.e.*, an image. Thus, the appropriate colour for each pixel has to be determined.

2.2 Rendering Equation

In computer graphics, and especially in realistic rendering, real world behaviours are the main source of information in producing accurate images. When realism is the primary goal, physical laws form the foundation of rendering convincing images. This section defines the *rendering*

equation [Kaj86], which models the interaction of light with surfaces in a three-dimensional scene, and presents several rendering techniques to numerically solve this equation.

The outgoing radiance at a surface location \mathbf{x} is expressed by the rendering equation as the sum of the emitted (L_e) and reflected (L_r) radiance,

$$L_o(\mathbf{x}, \boldsymbol{\omega}) = L_e(\mathbf{x}, \boldsymbol{\omega}) + L_r(\mathbf{x}, \boldsymbol{\omega}). \quad (2.7)$$

Without loss of generality, this radiance determines the colour perceived when observed from a point of view $\boldsymbol{\omega}$ with normal \mathbf{n} . The relation between the irradiance and the reflected radiance at point \mathbf{x} of normal \mathbf{n} , illustrated in Figure 2.2, is described by the *bidirectional reflectance distribution function* or *BRDF* [NRH⁺92]. Often denoted as f_r , it can be expressed in terms of the irradiance (Equation 2.5) as

$$f_r(\mathbf{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o) = \frac{dL_r(\mathbf{x}, \boldsymbol{\omega}_o)}{dE(\mathbf{x}, \boldsymbol{\omega}_i)} = \frac{dL_r(\mathbf{x}, \boldsymbol{\omega}_o)}{L_i(\mathbf{x}, \boldsymbol{\omega}_i)(\mathbf{n} \cdot \boldsymbol{\omega}_i)d\boldsymbol{\omega}_i}, \quad (2.8)$$

and so the reflected radiance can be written as

$$L_r(\mathbf{x}, \boldsymbol{\omega}_o) = \int_{2\pi} f_r(\mathbf{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o) L_i(\mathbf{x}, \boldsymbol{\omega}_i) (\mathbf{n} \cdot \boldsymbol{\omega}_i) d\boldsymbol{\omega}_i. \quad (2.9)$$

The rendering equation (Equation 2.7) can now be rewritten as

$$L_o(\mathbf{x}, \boldsymbol{\omega}) = L_e(\mathbf{x}, \boldsymbol{\omega}) + \int_{2\pi} f_r(\mathbf{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o) L_i(\mathbf{x}, \boldsymbol{\omega}_i) (\mathbf{n} \cdot \boldsymbol{\omega}_i) d\boldsymbol{\omega}_i. \quad (2.10)$$

This means that, in order to compute the outgoing radiance at \mathbf{x} in viewing direction $\boldsymbol{\omega}_o$, we need to integrate the incoming radiance from all incoming directions $\boldsymbol{\omega}_i$. This is not an easy task and there are numerous ways of approaching it.

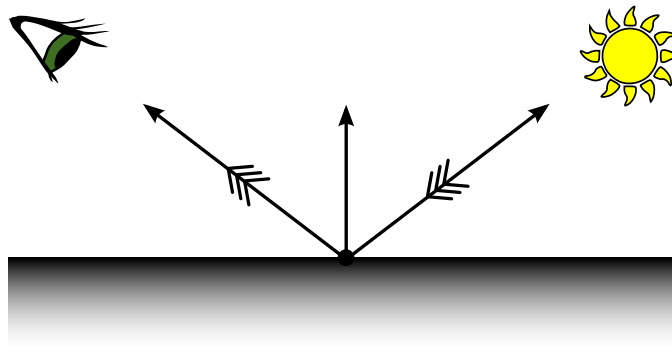


Figure 2.2 – The BRDF mathematically expresses the relation between the irradiance and the reflected radiance at a given point on an object. Conceptually, the BRDF describes the reflection behaviour of a surface.

2.3 Algorithms

As mentioned previously, analytical solutions to Equation 2.10 are only available in some specific, simple scenarios. Instead, in general, we resort to numerical integration to solve the rendering equation. There are several algorithms available, which can be grouped in three main categories: *sampling*, *density estimation*, and *finite element* methods.

The main application of finite element methods for solving the rendering equation is *radiosity* [GTGB84]. The idea is to model and compute the irradiance exchanged between discrete surfaces. The three-dimensional scene is subdivided into small patches that will contain the final light distribution, *i.e.*, the corresponding radiant exitance, related via a set of linear equations. Initially, this technique was used for scenes containing only diffuse materials, which are independent of the viewpoint. Subdividing a scene appropriately may also be a challenge due to numerous known difficulties in the domain of geometry, making radiosity algorithms impractical for complex scenes. We will therefore focus on sampling and density estimation techniques.

2.3.1 Path Tracing

It would be impractical to compute all the light interactions for an entire scene. Tracing infinitesimal *beams*, starting from the pixels, allows a simulation limited to a domain visible in an image. This technique is referred to as *ray tracing* [App68, Whi80] and these reversed rays of “light” are good at handling mirror reflections and illumination that comes from the direct incidence of a light source. In order to account for more effects, ray tracing has been extended with Monte Carlo stochastic techniques to form the more comprehensive *path tracing* [KVH84, CPC84] algorithm. The concept of rays is extended to paths, as light energy is transported and accumulated along such paths. Here, several jittered rays can be averaged for each pixel, according to the BRDFs, allowing an estimation of the integral over all light paths.

Path tracing additionally allows the simulation of realistic effects such as depth of field, glossy reflection, color bleeding, and motion blur. However, this technique requires a large number of samples to resolve the error due to numerical variance. The error, manifested in the form of noise in a Monte Carlo integral estimate is proportional to $1/\sqrt{N}$ for N samples. Four times the number of samples only reduces noise in an image by half, which leads to the problem that high quality rendering with path tracing is extremely time consuming.




Figure 2.3 – Path tracing (left) computes the color of a pixel by averaging the contribution of many light carrying paths through the scene, shot from the point of view. Photons are accumulated on the surfaces in photon mapping (right) and used as a measure of the incident radiance at rendering time.

2.3.2 Photon Mapping

Because a scene can be non-rigid and is likely to be modified, it is important to decouple illumination from geometry. *Photon mapping* [JC94, Jen96, JC98, Jen01] techniques propose to accumulate light as packetized particles, *i.e.*, photons, onto the surfaces. The photons are traced from the light sources and deposited onto surfaces in the scene. The radiance is then computed as a function of the number and power of the photons near an intersection point of the surface for a given point of view, with each photon weighted by the BRDF. This is a two-step algorithm, where the photons are accumulated in a first pass before the final rendering pass. This decoupling allows the reutilization of lighting information on objects at rendering time. Photon mapping is well suited for glossy surfaces as well as in generating realistic effects that result from the formation of complex light paths, such as caustics. Even though this technique is only subject to low frequency errors instead of high frequency noise in the rendered image, it is a biased (but consistent) method, which requires many photons to converge and thus consumes a large amount of memory.




Figure 2.4 – The *Cornell box* scene rendered with path tracing (left) and photon mapping (right). Notice the high frequency noise in the shadowy parts of the first image and the constant low frequency noise on the rectangular prisms’ faces in the second one.

2.4 Radiative Transport Equation

As mentioned in Section 2.2, Equation 2.10 expresses the equilibrium energy balance of light at *surface points* in a scene, and it is based on radiometry in order to respect physical laws. However, this equation assumes that light travels from surfaces to surfaces freely, which is often not the case in reality. Space is generally occupied by media like fog, dust, smoke, or even fire, and these are referred to as *participating media*, since they are composed of small particles that participate in the light transport process. When the light travels inside a volume that is not a vacuum, specific equations are required to model the behaviour of light within it.



Figure 2.5 – Ecosystem scene rendered without (left) and with (right) participating media. Additional effect makes the image look substantially more realistic. From *pbrt* [PH10].

The change in radiance along a ray in a participating medium is influenced by four types of interactions, which together must respect the law of conservation of energy: *absorption*, *emission*, *outscattering*, and *inscattering*. Considering the differential change of radiance $(\boldsymbol{\omega} \cdot \Delta)L(\mathbf{x}, \boldsymbol{\omega})$, this section derives the *RTE*, or *radiative transport equation* [Cha60], as the sum of these four quantities. They will be expressed term by term through Equations 2.11 to 2.15.

Absorption: When light enters a participating medium, it has a chance of being absorbed and converted into other forms of energy (*e.g.*, heat) that can be omitted from the rendering process, as they are invisible to humans. The denser the media, the more light it absorbs.

This interaction is parameterized by an *absorption coefficient* σ_a and is described in mathematical terms as

$$(\boldsymbol{\omega} \cdot \nabla)L_o(\mathbf{x}, \boldsymbol{\omega}) = -\sigma_a L_i(\mathbf{x}, \boldsymbol{\omega}). \quad (2.11)$$

Emission: As in the rendering equation for surfaces, participating media can also emit light (flames, chemical reactions). It is similarly described as

$$(\boldsymbol{\omega} \cdot \nabla)L_o(\mathbf{x}, \boldsymbol{\omega}) = L_e(\mathbf{x}, \boldsymbol{\omega}). \quad (2.12)$$

Outscattering: When light enters a participating media, it also has a chance of being scattered in a direction other than its initial $\boldsymbol{\omega}$. The likelihood of this event is also parameterized by a *scattering coefficient* σ_s and is described as

$$(\boldsymbol{\omega} \cdot \nabla)L_o(\mathbf{x}, \boldsymbol{\omega}) = -\sigma_s L_i(\mathbf{x}, \boldsymbol{\omega}). \quad (2.13)$$

It is common to combine the effects of both the absorption and scattering with an *extinction coefficient* $\sigma_t = \sigma_a + \sigma_s$, and so

$$(\boldsymbol{\omega} \cdot \nabla)L_o(\mathbf{x}, \boldsymbol{\omega}) = -\sigma_t L_i(\mathbf{x}, \boldsymbol{\omega}). \quad (2.14)$$

Inscattering: Just as light can be scattered away from the original direction $\boldsymbol{\omega}$, light from other directions can be scattered into direction $\boldsymbol{\omega}$, increasing the radiance along the ray. The inscattering contribution is computed by integrating over the sphere of possible directions, weighted by the scattering coefficient σ_s and a *phase function* p . Similarly to how the BRDF describes the relation between the irradiance and the reflected radiance for a given point of view, the phase function describes the angular distribution of light intensity scattered from incoming directions $\boldsymbol{\omega}_i$ into the outgoing direction $\boldsymbol{\omega}_o$. The resulting equation for the inscattered component of light along a ray in a participating media is

$$(\boldsymbol{\omega} \cdot \nabla)L_o(\mathbf{x}, \boldsymbol{\omega}) = \sigma_s \int_{4\pi} p(\mathbf{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o) L_i(\mathbf{x}, \boldsymbol{\omega}_i) d\boldsymbol{\omega}_i. \quad (2.15)$$

The *mean cosine of the scattering angle*, denoted g , is

$$g = \int_{4\pi} (\boldsymbol{\omega}_i \cdot \boldsymbol{\omega}_o) p(\mathbf{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o) d\boldsymbol{\omega}_i. \quad (2.16)$$

A positive g means the phase function is forward scattering, whereas a negative g corresponds to a backward scattering function, and a constant phase function of $g = 0$ denotes isotropic scattering.

By combining the effects of absorption, emission, outscattering and inscattering, we can fully describe the differential change of radiance along a ray, which leads to the final expression for the radiative transport equation

$$\begin{aligned} (\boldsymbol{\omega} \cdot \nabla) L_o(\mathbf{x}, \boldsymbol{\omega}) &= L_e(\mathbf{x}, \boldsymbol{\omega}) \\ &- \sigma_t L_i(\mathbf{x}, \boldsymbol{\omega}) \\ &+ \sigma_s \int_{4\pi} p(\mathbf{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o) L_i(\mathbf{x}, \boldsymbol{\omega}_i) d\boldsymbol{\omega}_i. \end{aligned} \quad (2.17)$$

2.5 Volume Rendering Equation

In practice, we are interested in obtaining the radiance arriving from a ray instead of the change in radiance along it. In order to find this radiance, Equation 2.17 has to be integrated on both sides. To do so, we introduce the *transmittance*, which describes the fraction of light passing through a medium for a travelled distance. Starting with Equation 2.14, integrating between \mathbf{x} and the point entering the medium \mathbf{x}' along a ray of light gives

$$L_o(\mathbf{x}, \boldsymbol{\omega}) = e^{-\tau(\mathbf{x}, \mathbf{x}')} L_i(\mathbf{x}', \boldsymbol{\omega}) \quad (2.18)$$

where

$$\tau(\mathbf{x}, \mathbf{x}') = \int_{\mathbf{x}}^{\mathbf{x}'} \sigma_t d\mathbf{x}_t, \quad (2.19)$$

for 3D points \mathbf{x}_t on the straight line between \mathbf{x} and \mathbf{x}' . The transmittance is $e^{-\tau}$ and τ is called the *optical depth*. In a homogeneous medium, σ_t is constant and the optical depth is simply

$$\tau(\mathbf{x}, \mathbf{x}') = \sigma_t \|\mathbf{x}' - \mathbf{x}\|, \quad (2.20)$$

where $\|\cdot\|$ is the Euclidean norm. Integrating the complete RTE yields the *volume rendering equation*

$$\begin{aligned} L_o(\mathbf{x}, \boldsymbol{\omega}) &= \int_{\mathbf{x}}^{\mathbf{x}'} e^{-\tau(\mathbf{x}, \mathbf{x}')} L_e(\mathbf{x}_t, \boldsymbol{\omega}) d\mathbf{x}_t \\ &\quad - e^{-\tau(\mathbf{x}, \mathbf{x}')} L_i(\mathbf{x}', \boldsymbol{\omega}) \\ &\quad + \int_{\mathbf{x}}^{\mathbf{x}'} e^{-\tau(\mathbf{x}, \mathbf{x}')} \sigma_s \int_{4\pi} p(\mathbf{x}_t, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o) L_i(\mathbf{x}_t, \boldsymbol{\omega}_i) d\boldsymbol{\omega}_i d\mathbf{x}_t. \end{aligned} \quad (2.21)$$

2.6 Algorithms

Again, analytic solutions to Equation 2.21 can only be derived for simple scenarios. In practice, it has to be evaluated numerically. We review popular techniques used to solve the RTE in the context of volumetric rendering, each of which are extensions of those introduced in Section 2.3.

2.6.1 Path Tracing

Extending path tracing to handle participating media is fairly straightforward. When a ray passes through a participating medium, and before its intersection with an object, a random point is selected along the ray and a stochastic scattering event is simulated, changing the trajectory of the ray in space. If this point falls outside the participating medium, the initial path is continued until it hits a surface and the standard surface rendering process is applied. Otherwise, it is scattered in another direction with respect to its phase function. This modification is easy to implement and introduces no bias. However, this process is subject to noise which further increases the already high sampling rate requirements for path-tracing-based numerical integration solutions.

2.6.2 Photon Mapping

Photons can be stored inside a volume in the same way as they are on the surface (Section 2.3.2). The transmittance affects the concentration of photons within the medium. Photon mapping is good at handling heterogeneous media, where the transmittance varies along the ray. The energy of these photons is finally weighted by the phase function. Much like ray tracing, storing photons in a first pass is trivial and faster than regular path tracing. The density can be estimated at rendering time by regularly sampling along a ray that passes through the volume.

Chapter 3

State of the Art

The previously described numerical methods used to solve the RTE suffer from noise and efficiency issues. In practice, this is a major problem since it requires trading precision in order to have smooth (but approximated) renderings at affordable costs. This section introduces the concept of diffusion based *subsurface scattering* and motivates the use of frequency analysis and adaptive rendering, to solve the RTE more efficiently and accurately in certain important practical circumstances.

3.1 Bidirectional Surface Scattering Reflectance Distribution Function

The bidirectional reflectance distribution function was introduced in Section 2.2, relating the ratio of the radiance leaving the surface, towards a point of view, to the incident radiance, at a given point. This concept can be generalized to *bidirectional surface scattering reflectance distribution function*, or *BSSRDF* [NRH⁺92], for volumetric media with a surface boundary. The BSSRDF describes the radiance leaving the surface, in the presence of subsurface volumetric scattering underneath the boundary, as in Figure 3.1. As opposed to the BRDF, we cannot assume that light enters and exits the surface boundary at the same point. Often denoted as S , the BSSRDF can be written as

$$S(\mathbf{x}_i, \boldsymbol{\omega}_i, \mathbf{x}_o, \boldsymbol{\omega}_o) = \frac{dL_r(\mathbf{x}_o, \boldsymbol{\omega}_o)}{d\Phi_i(\mathbf{x}_i, \boldsymbol{\omega}_i)}, \quad (3.1)$$

and so the reflected radiance expressed in terms of the BSSRDF is obtained from Equation 2.3,

$$L_r(\mathbf{x}_o, \boldsymbol{\omega}_o) = \int_{\mathcal{A}} \int_{2\pi} S(\mathbf{x}_i, \boldsymbol{\omega}_i, \mathbf{x}_o, \boldsymbol{\omega}_o) L_i(\mathbf{x}_i, \boldsymbol{\omega}_i) (\mathbf{n} \cdot \boldsymbol{\omega}_i) d\boldsymbol{\omega}_i d\mathcal{A}, \quad (3.2)$$

for each patch of surface area \mathcal{A} made up of points \mathbf{x}_i . This modification turns the 2D scattering equation into a substantially more complex 4D integral. This implies major changes in traditional rendering algorithms.




Figure 3.1 – As opposed to the BRDF (Figure 2.2), the BSSRDF does not assume that light enters and leaves at the same surface point. It describes the relation between the irradiance at an incident point and the reflected radiance at the exit point on a surface bounding a volumetric scattering medium.

3.1.1 Decomposition of S

It is common to decompose S , in order to make use of specialized algorithms, as a sum of three terms,

$$S = S^{(0)} + S^{(1)} + S_d, \quad (3.3)$$

where $S^{(0)}$, the *reduced radiance*, accounts for the unscattered extinct (absorbed) radiance, $S^{(1)}$, the *single scattering* term, accounts for the radiance scattered only once before leaving the surface, and S_d , the *multiple scattering* or *diffusion* term, accounts for the radiance scattered more than once.

3.1.2 Diffusion Theory

Equation 3.3 allows more specialized, efficient, and practical evaluation of the BSSRDF. The quantities $S^{(0)}$ and $S^{(1)}$ are usually inexpensive to compute using Monte Carlo methods, as they converge fast. Our main application of classical diffusion theory will be the calculation of S_d , since using common algorithms to compute it gives noisy results and converges slowly.

In a highly scattering media, the numerous scattering events tend to blur the light towards a uniform distribution. Based on this observation, the *diffuse reflectance distribution profile*, R_d , is approximated based on the diffusion approximation. The diffusion approximation considers the first two angular moments of the radiance distribution in a medium, respectively denoted as *fluence*, ϕ , and *flux*, \mathbf{E} :

$$\phi(\mathbf{x}, \boldsymbol{\omega}) = \int_{4\pi} L(\mathbf{x}, \boldsymbol{\omega}) d\boldsymbol{\omega}, \quad (3.4)$$

$$\mathbf{E}(\mathbf{x}, \boldsymbol{\omega}) = \int_{4\pi} L(\mathbf{x}, \boldsymbol{\omega}) \boldsymbol{\omega} d\boldsymbol{\omega}. \quad (3.5)$$

The radiance is then expanded according to this first order approximation [MS67, RH01] and normalized in order to maintain conservation of energy:

$$L(\mathbf{x}, \boldsymbol{\omega}) \approx \frac{1}{4\pi} \phi(\mathbf{x}, \boldsymbol{\omega}) + \frac{3}{4\pi} \mathbf{E}(\mathbf{x}, \boldsymbol{\omega}) \cdot \boldsymbol{\omega}. \quad (3.6)$$

Fick's first law of diffusion postulates that, under the assumption of steady state, the vector flux can be used to define fluence as $\mathbf{E}(\mathbf{x}) = -D\nabla\phi(\mathbf{x})$. By substituting Equation 3.6 into the RTE and applying Fick's law, the *diffusion equation* can be obtained by integrating over $\boldsymbol{\omega}$ [Ish78], which yields

$$-D\nabla^2\phi(\mathbf{x}, \boldsymbol{\omega}) + \sigma_a\phi(\mathbf{x}, \boldsymbol{\omega}) = Q(\mathbf{x}, \boldsymbol{\omega}), \quad (3.7)$$

where Q is a light source function, $D = 1/3\sigma'_t$ is the diffusion constant and σ'_t the reduced extinction coefficient. In Equation 2.16, this reduces the anisotropic problem ($g \neq 0$) to an approximated, simpler problem with isotropic scattering using modified scattering and extinction coefficients,

$$\sigma'_s = \sigma_s(1 - g), \quad (3.8)$$

$$\sigma'_t = \sigma'_s + \sigma_a. \quad (3.9)$$

Assuming an isotropic unit point light source Q , or *monopole*, and an infinite medium, the

diffusion equation has a simple solution, the *diffusion Green's function*:

$$\phi^m(\mathbf{x}) = \frac{1}{4\pi D} \frac{e^{-\sigma_{tr}r}}{r}, \quad (3.10)$$

where r is the distance to the monopole m at \mathbf{x} and $\sigma_{tr} = \sqrt{\sigma_a/D}$ is the transport coefficient.

When the medium is constrained in a finite region of space (*e.g.*, when a surface bounds the volume as in Figure 3.1), the diffusion approximation becomes subject to boundary conditions, namely that the inward diffuse flux must be zero at each point \mathbf{x}_s on the surface,

$$\int_{2\pi_-} L(\mathbf{x}_s, \boldsymbol{\omega})(\boldsymbol{\omega} \cdot \mathbf{n})d\boldsymbol{\omega} = 0, \quad (3.11)$$

where $2\pi_-$ is the inward hemisphere of directions. For the case with mismatched indices of refraction between the two media, there is an important reflection phenomenon at the surface that must be taken into account. The *average diffuse Fresnel reflectance*, F_{dr} , from the *Fresnel formula* for the reflectance at a dielectric surface, F_r , is

$$F_{dr} = \int_{2\pi} F_r(\eta, \mathbf{n} \cdot \boldsymbol{\omega}_o)(\mathbf{n} \cdot \boldsymbol{\omega}_o)d\boldsymbol{\omega}_o, \quad (3.12)$$

for a relative index of refraction η . It can be computed analytically [KL69], but a rational approximation is often used instead [EH79]

$$F_{dr} = \begin{cases} \frac{0.0636}{\eta^3} - \frac{0.3319}{\eta^2} + \frac{0.7099}{\eta} - 0.4399, & \eta < 1 \\ -\frac{1.4399}{\eta^2} + \frac{0.7099}{\eta} + 0.6681 + 0.0636\eta, & \eta > 1 \end{cases}. \quad (3.13)$$

This results in the following boundary conditions:

$$-\int_{2\pi_-} L(\mathbf{x}, \boldsymbol{\omega})(\boldsymbol{\omega} \cdot \mathbf{n}_-)d\boldsymbol{\omega} = F_{dr} \int_{2\pi_+} L(\mathbf{x}, \boldsymbol{\omega})(\boldsymbol{\omega} \cdot \mathbf{n}_+)d\boldsymbol{\omega}, \quad (3.14)$$

where the - and + subscripts denote inward and outward directions respectively, explaining the difference in signs between the two sides of the equation. Once again employing a first-order

approximation of the radiance distribution, the boundary conditions can be reordered as

$$\phi(\mathbf{x}_s) - 2AD(\mathbf{n} \cdot \nabla)\phi(\mathbf{x}_s) = 0, \quad (3.15)$$

where

$$A = \frac{1 + F_{dr}}{1 - F_{dr}}. \quad (3.16)$$

The challenge now lies in solving this boundary condition for S_d .

3.1.3 Searchlight Problem

A simplified version of the diffusion approximation problem has been proposed in medical physics and astrophysics [Cha58], known as the *searchlight problem*. As illustrated in Figure 3.2, it consists of an infinitesimal *pencil beam* of light striking a semi-infinite slab at normal incidence. Photons are then distributed as they scatter through the medium and leave the surface. This distribution, denoted R , is called a *reflectance distribution profile*. Because the medium is homogeneous and the pencil beam is at normal incidence, this profile is 1D radially symmetric, which means

$$R(\mathbf{x}_o - \mathbf{x}_i) = R(\|\mathbf{x}_o - \mathbf{x}_i\|). \quad (3.17)$$




Figure 3.2 – The searchlight problem consists of an infinitesimal beam of light striking a semi-infinite slab at normal incidence. The scattered distribution of light is a 1D radially symmetric function, denoted R .

3.1.4 Dipole

The previously defined boundary conditions come from the observation that the incoming radiance on the surface can be treated as a light source inside the medium, considering the first-order scattering events. A well-known solution to solving these boundary conditions is the *method of images*. This method extrapolates an analogous problem where a negative light source is mirrored with respect to a mirror plane, convenient for subsurface reflection which is often modeled as a semi-infinite planar medium.




Figure 3.3 – The dipole method consists of two point light sources, a positive (bottom) and a negative (top), placed in such a way that the boundary condition (e.g., Equation 3.15) is satisfied.

Jensen et al. [JMLH01] proposed a practical *dipole* model, based on Farrell et al. [FPW⁺92]’s method, suitable for simulating the appearance of subsurface scattering in computer graphics. By placing one positive real point light source at distance z_r beneath the surface and one negative virtual point light source at distance $z_v = z_r + z_b$, where $z_b = 4AD$, the boundary conditions can be satisfied. The resulting fluence along the surface boundary is

$$\phi(\mathbf{x}) = \frac{1}{4\pi D} \left(\frac{e^{-\sigma_{tr}d_r}}{d_r} - \frac{e^{-\sigma_{tr}d_v}}{d_v} \right), \quad (3.18)$$

where $d_r = \|\mathbf{x} - \mathbf{x}_r\|$ is the distance from the surface point \mathbf{x} to the real light source and

$d_v = \|\mathbf{x} - \mathbf{x}_v\|$ the distance to the virtual light source. The real light source is located at a distance of one *mean-free path*, $z_r = 1/\sigma'_t$, which is the average distance travelled by a photon underneath the surface. With Fick's law, we are now able to compute the diffuse reflectance,

$$R_d(\|\mathbf{x}_o - \mathbf{x}_i\|) = \frac{1}{4\pi} \left(\frac{z_r(1 + \sigma_{tr}d_r)e^{-\sigma_{tr}d_r}}{d_r^3} - \frac{z_v(1 + \sigma_{tr}d_v)e^{-\sigma_{tr}d_v}}{d_v^3} \right). \quad (3.19)$$

Finally, by taking into account the Fresnel reflection for both incoming and outgoing radiance, we can obtain the multiple scattering term of the BSSRDF,

$$S_d(\mathbf{x}_i, \boldsymbol{\omega}_i, \mathbf{x}_o, \boldsymbol{\omega}_o) = \frac{1}{\pi} F_t(\eta, \boldsymbol{\omega}_i) R_d(\|\mathbf{x}_o - \mathbf{x}_i\|) F_t(\eta, \boldsymbol{\omega}_o) \quad (3.20)$$

3.1.5 Quadpole

Assuming a flat surface with infinite extent on arbitrary geometry may lead to considerable error in the diffuse reflectance. Commonly modelled with polygon meshes, objects are likely to have sharp corners, which contradict the semi-infinite media assumption made earlier. Donner and Jensen [DJ05] extended the dipole model to account for more complex geometry, resulting in the *quadpole* method.

Assuming a geometry with convexity (*e.g.*, the corner of a cube), as shown in Figure 3.4, the initial real and virtual light sources are respectively placed as usual at a depth z_r and z_v . This time, the infinite surface is replaced by a corner which adds another boundary surface to the problem. This new boundary condition is solved again with the method of images by mirroring the dipole around the side face at a distance $2z_b$. This defines the fluence along the surface as

$$\phi(\mathbf{x}) = \frac{1}{4\pi D} \left(\frac{e^{-\sigma_{tr}d_r}}{d_r} - \frac{e^{-\sigma_{tr}d_v}}{d_v} + \frac{e^{-\sigma_{tr}d_{rm}}}{d_{rm}} - \frac{e^{-\sigma_{tr}d_{vm}}}{d_{vm}} \right), \quad (3.21)$$

where d_{rm} and d_{vm} are the distance from the surface point \mathbf{x} to the new real and virtual light

sources, respectively. The diffuse reflectance is then computed in the same way as the dipole, as

$$R_d(\|\mathbf{x}_o - \mathbf{x}_i\|) = \frac{1}{4\pi} \left(\frac{z_r(1 + \sigma_{tr}d_r)e^{-\sigma_{tr}d_r}}{d_r^3} - \frac{z_v(1 + \sigma_{tr}d_v)e^{-\sigma_{tr}d_v}}{d_v^3} + \frac{x_r(1 + \sigma_{tr}d_{rm})e^{-\sigma_{tr}d_{rm}}}{d_{rm}^3} - \frac{x_v(1 + \sigma_{tr}d_{vm})e^{-\sigma_{tr}d_{vm}}}{d_{vm}^3} \right), \quad (3.22)$$

where x_r and x_v are the distances to the side face for the initial dipole pair and the new one, respectively, as in Figure 3.4.

Donner and Jensen [DJ05] suggest that, for non-trivial meshes, an interpolation could be used between the dipole, quadpole, and the next extended model, the *multipole*.

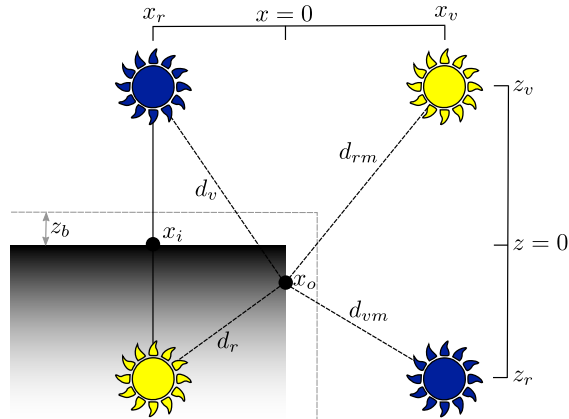


Figure 3.4 – The basic dipole configuration is mirrored in the case of a $\pi/2$ angle with an adjacent face, forming a quadpole, to satisfy the boundary condition.

3.1.6 Multipole

Another dangerous assumption that arises from the searchlight problems is the semi-infinity of the medium. This prevents heterogeneous multilayered media to be handled or for a thin layer of material to receive any backlit contribution of the light. Again, Donner and Jensen [DJ05, DJ08] propose an extension of the dipole model to support such effects.

The new boundary condition is handled by mirroring the dipole analogously at the bottom side

of the material slab and by doing the same for the upper side. By repeating the process a few times, until the error gets negligible, the boundary condition is solved over the surface. This setup is illustrated in Figure 3.5, where the poles are placed at depths

$$\begin{aligned} z_{r,i} &= 2i(d + 2z_b) + l \\ z_{v,i} &= 2i(d + 2z_b) - l + 2z_b, i = -n, \dots, n \end{aligned}, \quad (3.23)$$

where $2n + 1$ is the number of dipoles, l the depth of the original real source and d the slab thickness. The corresponding fluence in this configuration is

$$\phi(\mathbf{x}) = \frac{1}{4\pi D} \sum_{i=-n}^n \left(\frac{e^{-\sigma_{tr}d_{r,i}}}{d_{r,i}} - \frac{e^{-\sigma_{tr}d_{v,i}}}{d_{v,i}} \right). \quad (3.24)$$

The diffuse reflectance can finally be computed as

$$R_d(\|\mathbf{x}_o - \mathbf{x}_i\|) = \frac{1}{4\pi} \sum_{i=-n}^n \left(\frac{z_{r,i}(1 + \sigma_{tr}d_{r,i})e^{-\sigma_{tr}d_{r,i}}}{d_{r,i}^3} - \frac{z_{v,i}(1 + \sigma_{tr}d_{v,i})e^{-\sigma_{tr}d_{v,i}}}{d_{v,i}^3} \right). \quad (3.25)$$

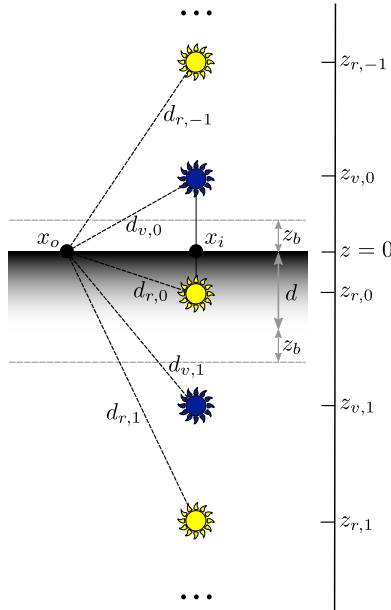


Figure 3.5 – The multipole configuration is an extension to regular dipole that may apply to thin slabs and multilayered media.

3.1.7 Improving the Accuracy of Dipole Approximations

d'Eon and Irving [DI11, d'E12] suggested modifications to the diffusion equation and introduced new parameters, relying on Grosjean's work. Because the classical dipole model is easy to implement, their goals were more accurate exitant radiance calculation and consistent boundary conditions, all the while keeping the simplicity and ease of implementation. This new technique requires negligible additional computation.

Grosjean [Gro51, Gro56, Gro59] introduced an alternative solution to the Green's function approximation of the fluence of Equation 3.10, still assuming a monopole and an infinite medium, namely

$$\phi^m(\mathbf{x}) = \frac{e^{-\sigma'_t r}}{4\pi r^2} + \frac{\alpha'}{4\pi D} \frac{e^{-\sigma_{tr} r}}{r}, \quad (3.26)$$

where α' is the *reduced albedo*, or *reduced reflection coefficient*, which represents the diffuse reflectivity, and D an improved diffusion coefficient,

$$D = \frac{2\sigma_a + \sigma'_s}{3\sigma_t'^2} = \frac{1}{3\sigma_t'} + \frac{\sigma_a}{3\sigma_t'^2}, \quad (3.27)$$

with $\sigma_{tr} = \sqrt{\sigma_a/D}$ defined with this new coefficient. Equation 3.26 is the sum of the exact single scattering and an approximate multiple scattering. This approximation was found to be more accurate for all absorption levels and distance than other proposed approximations [WWW58, KI98, GR01]. The diffusion theory focusses on multiple scattering: the first term of the equation is left for other methods to handle. This results in a modified diffusion equation:

$$-D\nabla^2\phi(\mathbf{x}, \boldsymbol{\omega}) + \sigma_a\phi(\mathbf{x}, \boldsymbol{\omega}) = \alpha'Q(\mathbf{x}, \boldsymbol{\omega}). \quad (3.28)$$

The boundary condition forcing fluence to be zero at distance $z_b = 2AD$ remains the same, but with the modified D and an improved reflection parameter,

$$A \approx \frac{1 + 3C_2}{1 - 2C_1}, \quad (3.29)$$

where C_1 and C_2 are defined by d'Eon [d'E12]. With these changes, the resulting fluence distribution along the boundary surface is

$$\phi(\mathbf{x}) = \frac{\alpha'}{4\pi D} \left(\frac{e^{-\sigma_{tr}d_r}}{d_r} - \frac{e^{-\sigma_{tr}d_v}}{d_v} \right). \quad (3.30)$$

Instead of using Fick's law, which relies on the vector flux to define fluence, d'Eon and Irving [DI11] used a Robin boundary condition for the exitant radiance, which is a linear combination of the fluence and the vector flux,

$$R_d(\mathbf{x}) = C_\phi \phi(\mathbf{x}) - C_E D (\mathbf{n} \cdot \nabla) \phi(\mathbf{x}), \quad (3.31)$$

where $C_\phi = \frac{1}{4}(1 - 2C_1)$ and $C_E = \frac{1}{2}(1 - 3C_2)$. From this, the diffuse reflectance can be rewritten

$$R_d(\mathbf{x}) = R_d^\phi(\mathbf{x}) + R_d^E(\mathbf{x}), \quad (3.32)$$

with

$$R_d^\phi(\mathbf{x}) = C_\phi \frac{\alpha'}{4\pi D} \left(\frac{e^{-\sigma_{tr}d_r}}{d_r} - \frac{e^{-\sigma_{tr}d_v}}{d_v} \right), \quad (3.33)$$

$$R_d^E(\mathbf{x}) = C_E \frac{\alpha'}{4\pi} \left(\frac{z_r(1 + \sigma_{tr}d_r)e^{-\sigma_{tr}d_r}}{d_r^3} - \frac{z_v(1 + \sigma_{tr}d_v)e^{-\sigma_{tr}d_v}}{d_v^3} \right). \quad (3.34)$$

This reflectance profile can still be directly derived for the quadpole and multipole configurations.

This improved theory consistently results in more accurate profiles, with negligible additional computation, making it suitable for CG purposes. Habel et al.'s [HCJ13a] technical report provides an excellent, comprehensive survey of classical and improved diffusion theory.

3.1.8 Quantized Diffusion

d'Eon and Irving [DI11] proposed to approximate the diffuse reflectance using a sum of Gaussians. One of the main objectives to quantized diffusion is to handle thin- and multi-layered media, several orders of magnitude thinner than previously possible. Instead of placing all energy at

a depth of one mean-free path, which is more or less accurate as the distance to the source is reduced, they apply the light source function $Q(l) = \alpha' \sigma'_t e^{-\sigma'_t l}$ to the 3D-normalized Gaussian

$$G_{3D}(v, r) = \frac{1}{(2\pi v)^{3/2}} e^{-r^2/(2v)} \quad (3.35)$$

of variance v . This allows the construction of a planar surface Gaussian

$$\int_0^\infty G_{3D}(v, \sqrt{r^2 + l^2}) Q(l) dl = \frac{1}{2} \sigma'_s f(\sigma'^2_t v/2) G_{2D}(v, r), \quad (3.36)$$

where G_{2D} is the 2D-normalized Gaussian

$$G_{2D}(v, r) = \frac{1}{2\pi v} e^{-r^2/(2v)} \quad (3.37)$$

and $f(x) = e^x \operatorname{erf}(\sqrt{x})$ contains the complementary error function erf .

Photons leaving a point source and undergoing diffusion produce Gaussian distributions with a mean displacement proportional to \sqrt{tD} where t is the time since the emission. This is related to Equation 3.10 by integrating over contributions from all emission times t in the past,

$$\phi^m(\mathbf{x}) = \frac{1}{4\pi D} \frac{e^{-\sigma_a r}}{r} = \int_0^\infty \frac{c}{(4\pi D c t)^{3/2}} e^{\sigma_a c t} e^{-r^2/(4D c t)} dt, \quad (3.38)$$

where c is the speed of light in the medium. In terms of the 3D-normalized Gaussian, it is expressed as

$$\int_0^\infty \frac{c}{(4\pi D c t)^{3/2}} e^{\sigma_a c t} e^{-r^2/(4D c t)} dt = \int_0^\infty e^{-\tau \sigma_a} G_{3D}(2D\tau, r) d\tau \quad (3.39)$$

where $\tau = ct$, measured in meters. The distribution of photons at time τ grows as a 3D Gaussian of variance $v = 2D\tau$, and all photons have travelled a distance of $\tau = ct$, weighting the Gaussian distribution by an absorption of $e^{-\tau \sigma_a}$. They quantize the interval $\tau \in [0, \infty]$ with $k + 1$ discrete values τ_i with $\tau_0 = 0$ and $\tau_k = \infty$, weighting the Gaussian to approximate the distribution of

photons as

$$\phi^m(\mathbf{x}) = \int_0^\infty e^{-\tau\sigma_a} G_{3D}(2D\tau, r) d\tau \approx \sum_{i=0}^{k-1} w_i G_{3D}(v_i, r), \quad (3.40)$$

where

$$w_i = \int_{\tau_i}^{\tau_{i+1}} e^{-\tau\sigma_a} d\tau = \frac{e^{-\tau_1\sigma_a} - e^{-\tau_2\sigma_a}}{\sigma_a}. \quad (3.41)$$

They showed that $\tau_i = s^{i-1}\tau_1$ with $s \in [1.5, 2.0]$ produces an accurate and sparse set of Gaussians. The variance v_i is chosen using $v = 2D\tau$ and taking the average shape of the Gaussian in the interval $[\tau_i, \tau_{i+1}]$ as $v_i = D(\tau_i + \tau_{i+1})$.

From the method of images, the fluence and flux weights for a Gaussian are

$$\begin{aligned} w_\phi(v, z) &= \int_0^\infty \frac{e^{-\frac{(-l+z)^2}{2v}}}{\sqrt{2\pi v}} \alpha' \sigma_t e^{\sigma_t l} dl \\ &= \frac{\alpha' \sigma_t}{2} e^{z\sigma_t + \frac{\sigma_t^2 v}{2}} \left(1 - \operatorname{erf} \left[\frac{z + \sigma_t v}{\sqrt{2v}} \right] \right) \end{aligned} \quad (3.42)$$

and

$$w_{\mathbf{E}}(v, z) = D\sigma_t \left(-w_\phi(v, z) + \frac{\alpha' e^{-\frac{z^2}{2v}}}{\sqrt{2\pi v}} \right), \quad (3.43)$$

for a total reflectance weight of

$$w_R(i) = C_\phi \left(w_\phi(v_i, z_r) - w_\phi(v_i, -z_v) \right) + C_{\mathbf{E}} \left(w_{\mathbf{E}}(v_i, z_r) + w_{\mathbf{E}}(v_i, -z_v) \right). \quad (3.44)$$

This leads to the diffuse reflectance approximated by the sum of Gaussians

$$R_d(\mathbf{x}) \approx \alpha' \sum_{i=0}^{k-1} w_R(i) w_i G_{2D}(v_i, r). \quad (3.45)$$

Note that the equations in this section are simplified to match a semi-infinite material configuration. Indeed, one of the drawbacks of this technique is the high complexity of the mathematical background. Extensive equations are available in d'Eon and Irving [DI11], based on the multipole configuration of Section 3.1.6.




Figure 3.6 – Comparison of quantized diffusion with other diffusion-based approaches for a semi-infinite material with a white light at normal incidence. The QD model retains high-frequency details present in a ground truth Monte Carlo simulation. From d’Eon and Irving [DI11].

3.1.9 Directional Dipole

Frisvad et al. [FHK14] proposed to replace point light sources with ray sources. Figure 3.6 shows results of diffusion-based approaches with light forced at normal incidence. Indeed, previous models were simplified that way, ignoring the angle of incidence of light. Using a ray source allows taking that direction into account, which better matches the actual behaviour of light refracting through media.

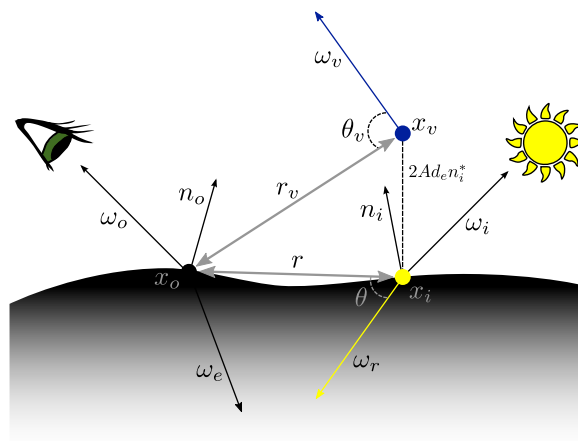


Figure 3.7 – Ray sources replace the point sources in this augmented diffusion configuration. The refracted direction ω_r , which stands as the positive source, is mirrored by a modified tangent plane to obtain the negative ray source ω_v . This model handles directionality and relaxes the assumption of flat boundaries.

The first step is to replace the isotropic light source in the diffusion equation by a ray of light in an infinite medium. Based on recent solutions to the diffusion equation for a ray of light in an infinite medium [MSG05a, MSG05b], a directional source term can be inserted into Equation 3.10 to retrieve the fluence as a function of the incident ray angle θ for a monopole,

$$\phi^m(\mathbf{x}, \theta) = \frac{1}{4\pi D} \frac{e^{-\sigma_{tr}r}}{r} \left(1 + 3D \frac{1 + \sigma_{tr}r}{r} \cos \theta \right). \quad (3.46)$$

Note that for a perpendicular ray of light, $\cos \theta = 0$ and we obtain the solution of Equation 3.10 for a point light source. As opposed to previous work, the gradient of the previous equation is computed,

$$\begin{aligned} \nabla \phi^m = \frac{1}{4\pi D} \frac{e^{-\sigma_{tr}r}}{r^3} & \left(\boldsymbol{\omega}_r 3D(1 + \sigma_{tr}r) - (\mathbf{x}_o - \mathbf{x}_i)(1 + \sigma_{tr}r) \right. \\ & \left. - (\mathbf{x}_o - \mathbf{x}_i) 3D \frac{3(1 + \sigma_{tr}r) + (\sigma_{tr}r)^2}{r} \cos \theta \right), \end{aligned} \quad (3.47)$$

where $\boldsymbol{\omega}_r$ is the refracted incident ray at \mathbf{x}_i . This allows the direct application of Equation 3.31 in order to obtain the exitant radiance of a directional monopole,

$$\begin{aligned} R_d^m(\mathbf{x}_i, \boldsymbol{\omega}_r, \mathbf{x}_o) = \frac{1}{4\pi^2} \frac{e^{-\sigma_{tr}r}}{r^3} & \left[C_\phi \left(\frac{r^2}{D} + 3(1 + \sigma_{tr}r)r \cos \theta \right) \right. \\ & - C_E \left((\boldsymbol{\omega}_r \cdot \mathbf{n}_o) 3D(1 + \sigma_{tr}r) - (\mathbf{x}_o - \mathbf{x}_i) \right. \\ & \left. \left. \cdot \mathbf{n}_o \left((1 + \sigma_{tr}r) + 3D \frac{3(1 + \sigma_{tr}r) + (\sigma_{tr}r)^2}{r} \cos \theta \right) \right) \right]. \end{aligned} \quad (3.48)$$

Again, the boundary condition that arises from the surface is treated using the method of images. However, for the fluence to vanish at the boundary, as illustrated in Figure 3.7, the direction has to be mirrored too. Mirroring the directional monopole through the tangent plane defined by \mathbf{n}_i is numerically unstable on a rough detailed surface, and so a modified tangent plane, defined by \mathbf{n}_i^* , is used to mirror the source:

$$\mathbf{n}_i^* = \frac{\mathbf{x}_o - \mathbf{x}_i}{|\mathbf{x}_o - \mathbf{x}_i|} \times \frac{\mathbf{n}_i \times (\mathbf{x}_o - \mathbf{x}_i)}{|\mathbf{n}_i \times (\mathbf{x}_o - \mathbf{x}_i)|}. \quad (3.49)$$

As opposed to previous methods, the real light source is not displaced one mean free path underneath the surface. This choice was made since the light source corresponds more accurately to the actual light ray, but it requires clamping the distance to $r = |\mathbf{x}_o - \mathbf{x}_i|$ for numerical stability where r tends towards zero, and for regions where the assumption of uniform emergent radiance is invalid due to directionality. The clamping is performed as follows:

$$r = \begin{cases} \max(|\mathbf{x}_o - \mathbf{x}_i|, d_e), & \boldsymbol{\omega}_r \cdot \boldsymbol{\omega}_e < 0 \\ \max(|\mathbf{x}_o - \mathbf{x}_i|, D), & \boldsymbol{\omega}_r \cdot \boldsymbol{\omega}_e \geq 0 \end{cases}, \quad (3.50)$$

where $d_e = 0.7104/\sigma'_s = 2.121D/\alpha'$ is the distance at which the fluence vanishes for an extrapolated boundary [GE52, Ish78].

With this configuration, depicted in Figure 3.7, the diffuse reflectance is computed as

$$R_d(\mathbf{x}_i, \boldsymbol{\omega}, \mathbf{x}_o) = \frac{R_d^m(\mathbf{x}_i, \boldsymbol{\omega}_r, \mathbf{x}_o) - R_d^m(\mathbf{x}_v, \boldsymbol{\omega}_v, \mathbf{x}_o)}{4C_\phi(1/\eta)}, \quad (3.51)$$

where the mirror positions are $\mathbf{x}_v = \mathbf{x}_i + 2Ad_e\mathbf{n}_i^*$ and $\boldsymbol{\omega}_v = \boldsymbol{\omega}_r - 2(\boldsymbol{\omega}_r \cdot \mathbf{n}_i^*)\mathbf{n}_i^*$.

This representation of light as ray sources increases realism, but is more complex and requires more computation time due to the many light directions required to produce visually pleasing results. This model, unlike the standard dipole, is not reciprocal (*i.e.* swapping incident and emergent points), as it does not depend only on relative distance.

3.1.10 Related Techniques

In addition to the advances in diffusion theory, several related topics of interest to realistic rendering of subsurface scattering have been explored. For instance, real-world measurement techniques [OM85, Ant00, NGD⁺06] allow better understanding, comparison and image generation of highly scattering materials. These measurements can also be directly edited by artists to manipulate the appearance of selected portions of an object, while maintaining a smooth appearance [WZT⁺08, CP⁺08, STPP09]. Textures may also be applied to subsurface scattering,

and this is often done in the context of human skin [Her05]. These alternative approaches have different performances/accuracy requirements than the dipole-based techniques we base our work on, such as interactivity for artists and extremely high precision for material information gathering.

3.2 Numerical Integration for Realistic Image Synthesis

Rendering dense media with Monte Carlo path tracing is costly due to the growing number of scattering events occurring in the material before the light leaves the surface. This is the main issue that promoted the development of the aforementioned faster analytical approaches. However, efficient integration remains an important part in making even these analytic approaches practical. For every shading point \mathbf{x}_o , it is necessary to integrate the incoming light over the entire object’s surface area in order to compute the final color, even with the use of fast analytic approximations.

3.2.1 Monte Carlo

Instead of computing the contribution of millions of surface samples, Monte Carlo approaches may be used to stochastically sample the surface with respect to the object mean free path and the distance to the shading point. This allows for a faster evaluation. Variations of this strategy include *Russian roulette acceptance* [Khi89] and *dart throwing* sample generation [CJW⁺09] for creating points with uniform distribution, still based on distance and density, as well as importance sampling [HCJ13b]. Unfortunately, these methods tend to introduce noise, something we often want to avoid when reproducing the smooth appearance of BSSRDFs.

3.2.2 Two-pass Hierarchical Technique

To extend the practical application of their dipole model, Jensen and Buhler [JB02] introduced a technique to approximate the dipole evaluation, with minimal additional bias, as they exploit

the exponential-falloff-with-distance property of the diffuse reflectance. This method consists of clustering the illumination points that are distant to the shading point in order to accelerate the evaluation, allowing each shading point to only treat a subset of the dipole sources over the object. It is decoupled into two passes: a sampling and an evaluation pass.



Figure 3.8 – Poisson disk samples on a complex 3D object. From Bowers et al. [BWWM10].

The first step caches the incoming irradiance at various surface locations. Such a set of points P is referred to as a *point cloud*. For the basic algorithm, each sample p has an associated position \mathbf{p}_x , normal \mathbf{p}_n , and surface area p_a . For an object of total surface area \mathcal{A} , $\sum_{p \in P} p_a = \mathcal{A}$ must hold. To ensure an accurate evaluation, the required number of points depends on multiple factors including the geometry, the variation of the light, the scattering properties, and the integration technique. Jensen and Buhler [JB02] suggest limiting the distance between the points to a maximum of one mean free path, setting the approximate number of points required for a given object as $|P| = \mathcal{A} \frac{\sigma_t'}{\pi}$. Various methods are available to generate these points, such as *point repulsion* [Tur92], dart throwing [CJW⁺09] or *Poisson disks* [BWWM10]. These methods share a common property: they ensure a minimal distance between all neighboring points, and the resulting point cloud often satisfies a *blue noise* distribution constraint. It is also possible to generate a non-uniform distribution, since surface points are not required to have the same p_a . For instance, the sampling density could be increased around discontinuities in the geometry and decreased on flat surfaces. Once the samples are generated, the irradiance can be evaluated using common techniques such as path tracing or photon mapping.

From this set of irradiance points, it is possible to evaluate the contribution of all samples for a given shading point \mathbf{x}_o , using one of the methods detailed earlier (dipole, directional dipole, quantized diffusion). However, it would be impractical to do so since it is too expensive to evaluate the exact contribution for all the points generated during the pre-pass. Indeed, this pass can generate hundreds of thousands of points. The second pass adaptively evaluates the contribution of the point cloud, based on the exponentially decreasing contribution of the light samples with distance, as in Figure 3.2. Various hierarchical data structures can be used to accelerate this process, such as an octree or a kd-tree. At each node, the average irradiance, the total surface area, and the irradiance-weighted average location of its children are stored. At rendering time, this tree is traversed until a node is determined to be close enough to the actual contribution of its children. There are two major criteria used to make this decision: \mathbf{x}_o must not be in the node’s bounding box, and the estimation of the solid angle subtended by the irradiance samples, $\Delta\omega = p_a / \|\mathbf{x}_o - \mathbf{p}_x\|^2$, must be larger than a user-defined quality threshold ϵ . Algorithm 1 describes this tree-traversal scheme.

This technique can apply to most of the previously studied BSSRDF models, with a few modifications. For instance, Frisvad et al. [FHK14] replace the irradiance sample with a list of differential irradiance samples, and d’Eon and Irving [DI11] accumulate the irradiance into a 1D radial binning data structure.

3.2.3 Single-pass Hierarchical Technique

The gathering of illumination in a separate pass may become a significant expense in large, complex scenes. Arbree et al. [AWB08] address this problem using a single-pass algorithm based on Jensen et al. [JB02] and making use of lightcuts [WABG06].

They unify the evaluation of the surface irradiance and the subsurface transport by clustering not only sample illumination points, but their entire light paths. To do so, along with the surface being discretized with point samples, the lights in the scene are also discretized. The discretized lights are then clustered similarly to the surface samples. This setup could remain a two-pass summation of the contributions; however, they suggest using a single evaluation of

the clusters' contribution as illustrated in Figure 3.9, refining the node with the highest error at each step, based on a heuristic. The contributions are stored in order to avoid recomputing the light contribution, which was an important advantage of the original method of Jensen et al. [JB02].




Figure 3.9 – In reading order: (1) the naive method where all the irradiance samples from all the lights are computed per shading point e , (2) the method from Jensen and Buhler [JB02], (3) clustering the point lights and using the method from Jensen and Buhler [JB02], and (4) the links between the light clusters and the irradiance samples clusters are also clustered, allowing a faster evaluation. From Arbree et al. [AWB08].

3.3 Frequency Analysis

One of the properties of participating media is their tendency to blur out details and decrease contrast, typically requiring higher computational cost due to numerous scattering events that must be simulated. This also applies to subsurface scattering, where the diffusion process results in a smooth appearance. *Frequency analysis of light transport* studies how light propagation may tend to increase or decrease the final color variation of an image.

In terms of signal processing, an image subject to a blurring from a band-limiting kernel contains less relevant frequency content than its unblurred version. Thus, such an image may require less information to be reconstructed correctly, if we are mindful of the content of its *Fourier spectrum*. Using this approach, the rendering time could be reduced while producing a perceptually equivalent image. We will study this variation in the *Fourier domain*. The *Fourier transform* expresses a signal in terms of its frequency variation rather than its variation in the

primal domain:

$$\mathcal{F}[f] = \int_{\mathbf{x} \in \mathbb{R}^N} f(\mathbf{x}) e^{-2\pi i \boldsymbol{\xi}^T \mathbf{x}} d\mathbf{x}. \quad (3.52)$$

The Fourier transform calculates the frequency spectrum of f . This complex signal s is presented under the form $s(\boldsymbol{\xi}) = A(\boldsymbol{\xi}) e^{i\phi(\boldsymbol{\xi})}$, where A describes the amplitude of the energy associated with frequency $\boldsymbol{\xi}$, and ϕ , the phase shift associated for that frequency. Intuitively, the amplitude corresponds to the projected energy of the input signal to a given sinusoid with the associated frequency and shifted by the associated phase. For instance, a slowly varying signal will result in a tight spectrum around the origin of the Fourier domain, producing identifiable frequency amplitude (with the possibility of identifying a direction of variation). On the contrary, a highly varying signal will be spread in the Fourier domain in a hardly categorizable frequency content. Figure 3.10 depicts the behaviour of the Fourier transform for various input signal images.



Figure 3.10 – The Fourier transform of a signal captures its frequency variations. Low spatial frequency regions of the image are concentrated around the origin of the Fourier domain while high spatial frequency regions may span the entire Fourier domain. From Belcour [Bel12].

Durand et al. [DHS⁺05] presented a frequency-analysis framework for computer graphics, by establishing correspondences between light phenomena, such as shading, occlusion, and transport, to equivalent operations in the Fourier domain. This resulted in an approach to accurately predict the frequency content of a rendered image, which they used as an oracle to determine a

sampling rate for rendering, such as the number of rays required in path tracing. This theory was later used by Bagher et al. [BSS⁺12] to accelerate the rendering of diffuse and specular materials.

More recently, Belcour et al. [BBS14] extended this analysis to the behavior of light along light paths, thus improving the convergence of several existing methods in participating media. They decompose the RTE in order to analyse it in the Fourier domain. Their approach works well when the media is not optically thick, as the number of scattering events is not as high as with dense media. We saw that, in the case of diffusion-based subsurface scattering, the equations are simplified by various hypotheses, such as normal incidence and medium semi-infinity. In our work, we wish to extend the Fourier analysis of light transport to these diffusion-based BSSRDF models, which requires an adapted treatment. Indeed, no such analysis exists for BSSRDFs in rendering and, in light of previously mentioned integration algorithms, the need for efficient adaptive integration algorithms is another important requirement that we will address.

In this context where multiple BSSRDF models have been presented along with integration techniques, we present the following contributions:

1. An extension of the frequency analysis of Durand et al. [DHS⁺05], Bagher et al. [BSS⁺12], and Belcour et al. [BBS14] that apply to most of the previously studied BSSRDF models;
2. A numerical technique to estimate the frequency spectrum of scattering light for translucent media, which we use to determine the spatial and angular variation of outgoing radiance over the surface;
3. An extension to the integration technique from Jensen and Buhler [JB02] under the form of a dual-tree. The spatial integration step is done in joint image- and object-space, adaptively evaluating the underlying BSSRDF model for the final rendering.

Chapter 4

A Frequency Analysis and Dual Hierarchy for Efficient Rendering of Subsurface Scattering

This chapter presents our paper submitted to the *Computer Graphics Forum*, formatted to fit this thesis. The authors are David Milaenen, Laurent Belcour, Jean-Philippe Guertin, Toshiya Hachisuka and Derek Nowrouzezahrai.

4.0 Abstract

BSSRDFs are commonly used to model subsurface light transport in highly scattering media such as skin and marble. Rendering with BSSRDFs requires an additional spatial integration, which can be significantly more expensive than surface-only rendering with BRDFs. We introduce a novel hierarchical rendering method that can mitigate this additional spatial integration cost. Our method has two key components: a novel frequency analysis of subsurface light transport, and a dual hierarchy over shading and illumination samples. Our frequency analysis predicts the spatial and angular variation of outgoing radiance due to a BSSRDF. We use this analysis to drive adaptive spatial BSSRDF integration with sparse image and illumination samples. We propose the use of a dual-tree structure that allows us to simultaneously traverse a tree of shade points (i.e., pixels) and a tree of object-space illumination samples. Our dual-tree approach generalizes existing single-tree accelerations. Both our frequency analysis and the dual-tree structure are compatible with most existing BSSRDF models, and we show that our method improves rendering times compared to the state of the art method of Jensen and Buhler [JB02].

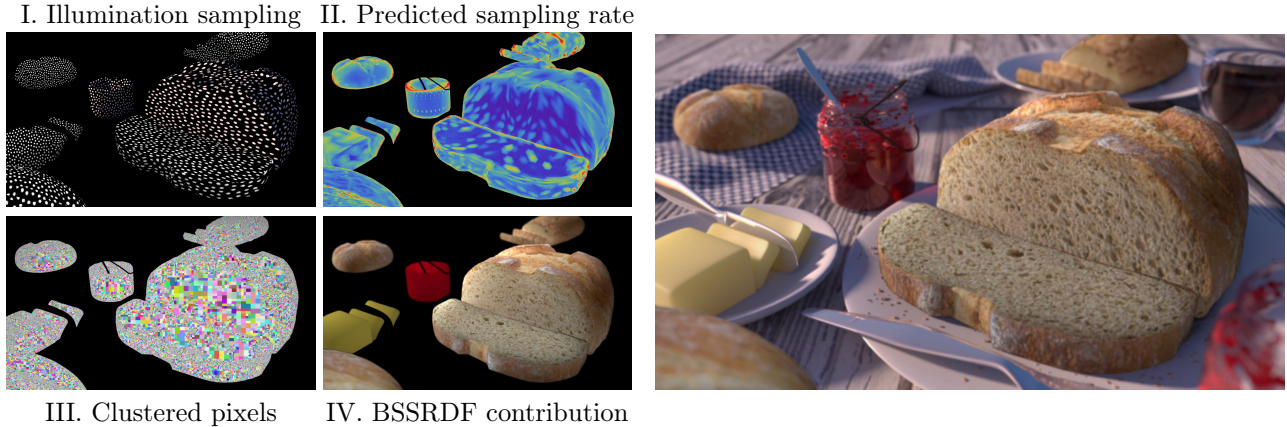


Figure 4.1 – We introduce a hierarchical method to accelerate the rendering of multiple scattering with BSSRDFs (IV). We overview our approach in the PICNIK scene, above: our frequency analysis of BSSRDFs allows us to predict the screen-space sampling rates (II) which are used to devise bounds on the variation of outgoing radiance. These bounds allow us to efficiently integrate the BSSRDF using a dual hierarchy over clustered illumination samples (I) and shading points (i.e., pixels; III).

4.1 Introduction

Including subsurface scattering effects in virtual scenes can significantly increase the realism of rendered images. Since many real-world materials exhibit subsurface scattering effects, modeling and simulating them remains an important problem in realistic image synthesis.

Accurate light transport in highly absorbing media can be modeled mathematically with the Bidirectional Scattering Surface Reflectance Distribution Function (BSSRDF). Many BSSRDF models exist, with varying degrees of accuracy: classical dipole models [JMLH01, d’E12] and quantized diffusion [DI11] do not account for the angular variation of incident radiance, however more recent models do [FHK14, HCJ13b, d’E14]. Unlike BRDFs, BSSRDFs describe light transport between two *different* locations on an object. As such, an additional spatial integration (over the surface) is required in order to render objects with BSSRDFs. Jensen and Buhler [JB02] introduced an adaptive hierarchical integration method to amortize the cost of this spatial integration using clusters of spatial illumination samples. While this approach has been successfully used in many applications, it does not take the smoothness of the resulting outgoing radiance (i.e., in screen-space) into account.

We propose a novel integration method that clusters **both** pixels and illumination points as illustrated in Figure 4.1. We conduct a frequency analysis of subsurface scattering that is agnostic to the underlying BSSRDF model. Specifically, we study the frequency content of the spatial and angular variation of radiance after its BSSRDF interaction. This leads us to a theoretically sound criterion for sparse sampling and adaptive integration. Using this criterion, we leverage a dual hierarchical data structure to accelerate the final evaluation of the multiple scattering term. Our hierarchical evaluation is motivated by the existing tree-based approach of Jensen and Buhler [JB02]; our dual-tree structure, however, amortizes computation cost across both pixels *and* illumination points. We are able to generate higher-quality results in less rendering time compared to the single tree method of Jensen and Buhler [JB02]. Concretely, we propose:

- a frequency analysis of shading with BSSRDFs,
- a numerical approach for estimating the BSSRDF spectra, which we use to determine the variation of outgoing radiance over the surface of a translucent object, capable of supporting any underlying dipole model, and
- the application of a dual-tree structure to the problem of BSSRDF estimation in joint image- and object-space, directly leveraging our frequency analysis to adaptively traverse the structure and accelerate the final rendering.

4.2 Previous Work

We focus on work that most closely aligns with our approach: specifically, we review integration schemes for BSSRDF models, and frequency analyses of light transport.

BSSRDF Integration Techniques. In all cases, the bottleneck of dipole-like techniques remains the numerical evaluation of the spatial-angular integration in Equation 4.1. Jensen and Buhler [JB02] compute an approximate evaluation of this contribution from sparsely sampled

irradiance samples distributed over a translucent object’s surface. Here, the outgoing radiance at any shade point is computed by traversing a tree over the irradiance samples and terminating traversal according to a quality criterion. This two-pass approach introduces a controllable bias and has remained compatible (often without modification) with many of the newer dipole models we discussed in Section 4.1: notably, Frisvad et al. need only substitute the (diffuse) irradiance samples with a vector of differential irradiance samples, and d’Eon and Irving use a supplemental 1D radial directional radiance bin.

Arbree et al. [AWB08] propose a scalable approach to rendering large translucent scenes based on multidimensional lightcuts [WABG06]. They aggregate the computation of irradiance samples by simultaneously clustering light sources and irradiance samples such that the resulting contribution to a given shade point can be well approximated. While this work also uses two trees, it treats each pixel independently without taking the resulting image smoothness into account (see Figure 4 and Section 4.1 of [AWB08]). We do not consider the evaluation cost of (ir)radiance samples, but we do cluster evaluation over pixels.

The idea of applying a doubly-adaptive traversal originates from the particle simulation literature [GR87], and the implementation of d’Eon and Irving’s quantized diffusion model [DI11] in Pixar’s RenderMan implicitly leverages a similar principle (i.e., with REYES’ adaptive micropolygon evaluation). We are similarly motivated by concurrent work that applies dual-tree structures to density estimation problems in realistic image synthesis (an anonymized manuscript of this concurrent work in submission is included in our supplemental material, for reference) [Ano15]. One of our contributions is a well-founded oracle to terminate shading tree traversal based on our BSSRDF frequency analysis.

Frequency Analyses of Light Transport. Durand et al. [DHS⁺05] presented the first comprehensive Fourier analysis of light transport in scenes with opaque surfaces, and a proof-of-concept adaptive image space sampling approach to reconstruct noise-free images at super-pixel sampling rates. Bagher et al. [BSS⁺12] derived atomic operators for bandwidth estimation in order to study environmental reflection with acquired BRDFs. Belcour et al. [BBS14]

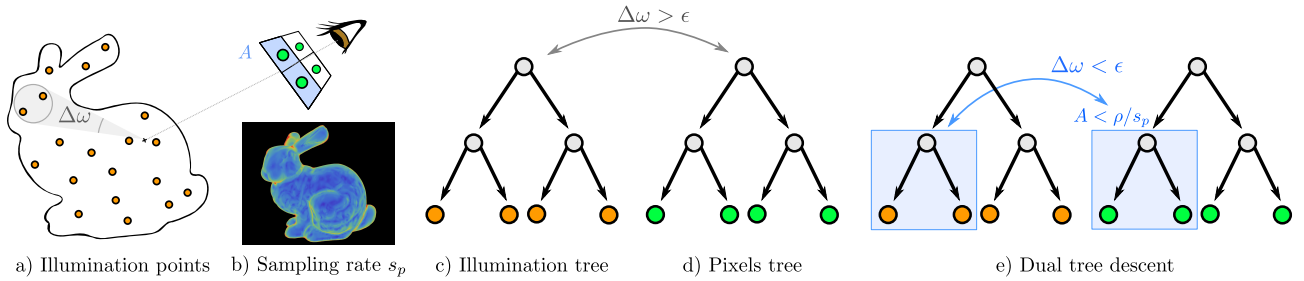


Figure 4.2 – We sample incident illumination over the object (a) according to its subsurface scattering properties and construct two spatial acceleration structures: one over these samples (c) and one over pixels (d). To render, we simultaneously traverse the trees (e), using our outgoing radiance bandwidth estimate s_p (b) to stop the tree traversal and shade super-pixels of area A .

extend these frameworks to incorporate the study of scattering in arbitrary participating media, however their analysis is not suited for dense media and BSSRDFs. We bridge this gap with a frequency analysis of scattering in dense media, similarly leveraging matrix-vector formulations of frequency-space bandwidth operators.

4.2.1 Overview

Figure 4.2 overviews our approach: after sparsely evaluating incident radiance on the surface of each translucent object (Figure 4.2a), we compute a per-pixel bandwidth estimate of the multiply-scattered outgoing radiance (Figure 4.2b). We build two spatial acceleration structures, one over illumination samples (Figure 4.2c) and another over pixels (Figure 4.2d). In order to compute the object’s final shading, we simultaneously traverse both trees, hierarchically accumulating the contribution of groups of illumination samples to groups of pixels (Figure 4.2e). We use the frequency bandwidth of the outgoing radiance predicted by our theory (Section 4.3) to terminate traversal along each tree, significantly reducing the number of BSSRDF evaluations necessary to compute the final image without introducing visible artifacts.

We present our BSSRDF frequency analysis theory, as well as its numerical realization for computing image-space radiance bandwidths in Section 4.3. We introduce our variant of the dual tree construction and explain how the bandwidth predictions are used during hierarchical

traversal in Section 4.4. Finally, we discuss our implementation details in Section 4.5 and compare our method to the state of the art in Section 4.6.

4.3 Fourier Analysis

We will derive conservative, numerical estimates of the frequency bandwidth of the outgoing radiance in image space, taking into account the effects of curvature, foreshortening, transport and multiple scattering on the incident light field’s frequency content. We will show that the BSSRDF acts as a band-limiting filter on the incident radiance distribution, and we will derive a conservative expression of the resulting spatio-angular bandwidth of the outgoing radiance spectrum (Section 4.3). We will use these bandwidth estimates, combined with the formulation of Bagher et al. [BSS⁺12], to predict the variation of outgoing radiance in image space (Section 4.3.2), which will in turn drive our hierarchical dual tree traversal and integration (Section 4.4).

4.3.1 Fourier Transform of a BSSRDF

Given a BSSRDF model $S(\mathbf{x}_i, \boldsymbol{\omega}_i, \mathbf{x}_o, \boldsymbol{\omega}_o)$, the outgoing radiance at the object surface L_o in direction $\boldsymbol{\omega}_o$ and at position \mathbf{x}_o is expressed as:

$$L_o(\mathbf{x}_o, \boldsymbol{\omega}_o) = \iint_{\mathcal{A} \times \mathcal{H}} S(\mathbf{x}_i, \boldsymbol{\omega}_i, \mathbf{x}_o, \boldsymbol{\omega}_o) L_i(\mathbf{x}_i, \boldsymbol{\omega}_i) d\boldsymbol{\omega}_i^\perp d\mathbf{x}_i, \quad (4.1)$$

where \mathcal{A} is the object’s surface area, \mathcal{H} is the set of (hemispherical) incident directions, L_i is the incident radiance, and $d\boldsymbol{\omega}_i^\perp = \cos\theta_i d\boldsymbol{\omega}_i$ is the projected solid angle.

If we apply a Fourier transform to Equation 4.1, converting products in the primal domain to convolutions in the frequency domain and integration in the primal domain to DC evaluation in

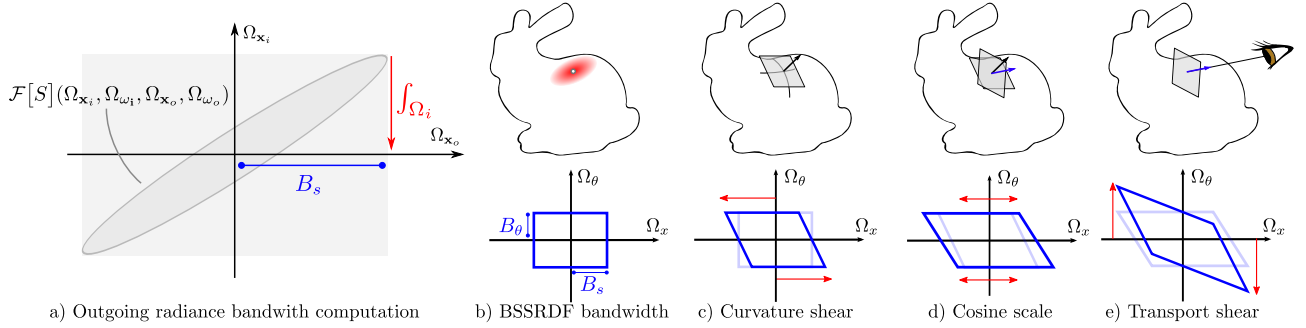


Figure 4.3 – Assuming that the incoming light-field has infinite bandwidth, we conservatively estimate the bandwidth of the outgoing light-field $[B_s, B_\theta]$ as the bandwidth of the BSSRDF along the outgoing spatial positions and directions (a). The interaction with the material limits the spectrum of the local light-field by the BSSRDF spatial and angular bandwidth (b). To estimate the bandwidth at the camera position, we first shear spatially the spectrum to account for curvature (c). Then, we scale spatially to account for foreshortening (d) and finally shear angularly the spectrum to account for transport (e).

the frequency domain, we obtain:

$$\mathcal{F}[L_o](\Omega_{\mathbf{x}_o}, \Omega_{\boldsymbol{\omega}_o}) = \left[\mathcal{F}[\widehat{S}] \circ \mathcal{F}[L_i] \right] (0, 0, \Omega_{\mathbf{x}_o}, \Omega_{\boldsymbol{\omega}_o}), \quad (4.2)$$

where $\mathcal{F}[f]$ is the Fourier transform of f , \circ the convolution operator, and Ω_x the frequency variation of x . Concretely, the outgoing radiance’s spatial-angular frequency spectrum $\mathcal{F}[L_o](\Omega_{\mathbf{x}_o}, \Omega_{\boldsymbol{\omega}_o})$ results from evaluating the convolution of the Fourier transform of the cosine-weighted BSSRDF $\mathcal{F}[\widehat{S}] = \mathcal{F}[S(\mathbf{x}_i, \boldsymbol{\omega}_i, \mathbf{x}_o, \boldsymbol{\omega}_o) \cos(\theta_i)]$ with the Fourier transform of the incident light $\mathcal{F}[L_i]$ at the incoming spatial and directional DC frequencies $(\Omega_{\mathbf{x}_i}, \Omega_{\boldsymbol{\omega}_i}) = (0, 0)$.

Assuming that $\mathcal{F}[L_i]$ contains all-frequency content, the resulting outgoing bandwidth (along $\Omega_{\mathbf{x}_o}$ and $\Omega_{\boldsymbol{\omega}_o}$) after convolution against the spectrum of the cosine-weighted BSSRDF $\mathcal{F}[\widehat{S}]$ will match the bandlimit of $\mathcal{F}[\widehat{S}]$ (see Figure 4.3a). We will discuss how to compute the spatial and angular bandwidths $\{B_o, B_\theta\}$ of the cosine-weighted BSSRDF given its local orientation.

Spatial Bandwidth. We compute the cosine-weighted BSSRDF’s spatial bandwidth numerically by sampling and projecting $S(\mathbf{x}_i, \boldsymbol{\omega}_i, \mathbf{x}_o, \boldsymbol{\omega}_o) \cos(\theta_i)$ into the frequency domain, across its different dimensions. Depending on the underlying BSSRDF model, the cosine-weighted

BSSRDF may depend on the viewing direction, the incident lighting direction, and the distance and angle between \mathbf{x}_o and \mathbf{x}_i .

For instance, the dipole model has a separable form:

$$\mathcal{F}[\widehat{S}] = \mathcal{F} [R_d(\|\mathbf{x}_i - \mathbf{x}_o\|) F_i(\theta_i) \cos(\theta_i) F_o(\theta_o)] ,$$

where R_d is the diffuse reflectance, and F_i and F_o are the incident and outgoing Fresnel terms [JMLH01, Equation 5]. Here, we take advantage of the separability of the model (w.r.t. θ_i and θ_o) to express its Fourier transform as

$$\mathcal{F}[\widehat{S}] = \underbrace{\mathcal{F} [R_d(\|\mathbf{x}_i - \mathbf{x}_o\|)] \mathcal{F} [F_i(\theta_i) \cos \theta_i]}_{\mathcal{F}[\widehat{S}_i](\Omega_{\mathbf{x}_i}, \Omega_{\mathbf{x}_o}, \Omega_{\boldsymbol{\omega}_i})} \mathcal{F} [F_o(\theta_o)] .$$

Since we are only concerned with the DC $[\Omega_{\mathbf{x}_i}, \Omega_{\boldsymbol{\omega}_i}] = [0, 0]$ hyperplane, the spatial bandwidth is computed with the 1D diffuse reflectance spectrum $\mathcal{F} [R_d] (\Omega_{x_o})$. We discuss the outgoing term $\mathcal{F} [F_o] (\Omega_{\omega_o})$ below.

In contrast, the directional dipole [FHK14] additionally takes $\boldsymbol{\omega}_i$ and the direction between \mathbf{x}_i and \mathbf{x}_o into account:

$$\mathcal{F}[\widehat{S}_i] = \mathcal{F} \left[\frac{e^{-\sigma_{tr}\|\mathbf{x}_i - \mathbf{x}_o\|}}{4\pi^2\|\mathbf{x}_i - \mathbf{x}_o\|^3} M(\mathbf{x}_i - \mathbf{x}_o, \boldsymbol{\omega}_{12}) F_i(\theta_i) \cos \theta_i \right] ,$$

where $M(\mathbf{x}_i - \mathbf{x}_o, \boldsymbol{\omega}_{12})$ models the spatial-directional scattering distribution and $\boldsymbol{\omega}_{12}$ is the refraction of $\boldsymbol{\omega}_i$ at \mathbf{x}_i [FHK14, Equation 17]. We extract the outgoing spatial bandwidth by taking the maximum 1D bandwidth for various angles between $\mathbf{x}_i - \mathbf{x}_o$, the normal at \mathbf{x}_i and the refracted ray $\boldsymbol{\omega}_{12}$.

In all instances, we compute a conservative estimate of the outgoing spatial and directional frequency *bandwidths*, B_s and B_θ , as the values required to retain 95% of the energy of the discrete power Fourier spectrum.

Angular Variation. The angular variation of the BSSRDF is modulated by the outgoing Fresnel term above, and we use a windowed Fourier transform to compute the bandwidth of $\mathcal{F}[F_o](\Omega_{\omega_o})$, again as the 95th energy percentile spectrum value. We tabulate these bandwidths as a function of θ_o , and use them to modulate B_θ ; this is particularly important at grazing angles, where the effects of the spectrum of the outgoing Fresnel term can significantly impact the angular bandwidth of the outgoing radiance.

4.3.2 Outgoing Radiance Bandwidth Computation

Given the spatial-angular bandwidth of the outgoing radiance at a shade point, conservatively estimated as the BSSRDF bandwidth, we need to compute the associated pixel frequency bandwidth. To do so, we are motivated by Bagher et al.’s [BSS⁺12] bandwidth tracking approach, applying bandwidth evolution operators defined by Durand et al. [DHS⁺05] to the bandwidth vector $[B_s, B_\theta]^T$. Figure 4.3 (c – d) illustrates the transport operators in the following order:

1. we transform from local shade point coordinates to global coordinates by projecting the outgoing spectrum onto the shade point’s tangent plane, which amounts to a shear in the spatial frequency according to the local curvature k ,
2. we take the foreshortening towards the viewpoint due to $\cos \theta_o$ into account, stretching the spectrum spatially, and
3. we evaluate the spectrum at the sensor, after transport through free-space with a distance d , by applying an angular shear to the spectrum.

These operations can be compactly expressed as matrix operators, if we act directly on frequency bandwidths instead of the full spectra [BSS⁺12], as:

$$\mathbf{T}_d = \begin{bmatrix} 1 & 0 \\ d & 1 \end{bmatrix}, \quad \mathbf{P}_x = \begin{bmatrix} 1/\cos \theta_x & 0 \\ 0 & 1 \end{bmatrix}, \quad \text{and} \quad \mathbf{C}_k = \begin{bmatrix} 1 & k \\ 0 & 1 \end{bmatrix}.$$




Figure 4.4 – First row: The sampling rate s_p computed from the screen-space bandwidth estimation. Second row: Pixel areas from which the sampling rate predicts an adequate approximation of the outgoing radiance variation.

We apply these operators, in order, to the outgoing radiance bandwidth (i.e., the BSSRDF bandwidth $[B_s, B_\theta]$), to predict the final screen space bandwidth vector for a pixel as:

$$\begin{bmatrix} B_p & B_a \end{bmatrix}^T = \mathbf{T}_d \mathbf{P}_x \mathbf{C}_k \begin{bmatrix} B_s & B_\theta \end{bmatrix}^T. \quad (4.3)$$

Isolating the screen space angular bandwidth B_a above,

$$B_a = B_\theta + d (B_s + kB_\theta) / \cos \theta, \quad (4.4)$$

and applying the Nyquist criterion, we arrive at the pixel sampling rate s_p (in units of pixel^{-1}) as twice the angular screen space bandwidth,

$$s_p = 2 B_a \max (f_x/W, f_y/H), \quad (4.5)$$

for a $W \times H$ image resolution and a horizontal and vertical field of view of f_x and f_y . Figure 4.4

visualizes the screen space sampling rate for the scenes we render. We discuss how we use the pixel sampling rate to drive our new dual hierarchical BSSRDF integration approach below.

4.4 Hierarchical Approach

We explain how to utilize our bandwidth estimation in order to accelerate rendering with BSSRDFs. We first review the single hierarchy approach of Jensen and Buhler [JB02], then explain how we can use a dual hierarchy to adaptively cluster both illumination samples and pixels simultaneously.

4.4.1 Hierarchical Surface Integration

Jensen and Buhler [JB02] pointed out that we can cluster illumination samples over the surface in order to reduce the cost of BSSRDF evaluations. The underlying observation is that we can aggregate contributions from illumination samples that are distant from a given shading point. We can thus evaluate the BSSRDF only once for a cluster of such illumination samples, resulting in fewer BSSRDF evaluations.

This approach has two passes. In the first pass, pre-integrated illumination samples are inserted into a tree data structure where each inner node i represents the aggregated information of its children. For example, each node stores the average illumination, the total surface area A_i , and the irradiance-weighted average location \mathbf{p}_i of its children. In the second pass, we traverse this tree until the current node accurately represents all the contributions of its children to a given shading point. If the shading point is in the bounding volume of the current node, we keep traversing the tree and consider contributions from the children nodes. Otherwise, we traverse to the child nodes only if the conservative estimate of the solid angle subtended by the illumination samples, $\Delta\omega = A_i / \|\mathbf{x}_o - \mathbf{p}_i\|^2$, is larger than the user-defined quality threshold ϵ (Algorithm 1). While this approach significantly reduces the cost of integration over the surface, it is repeated for each shading point without considering the smoothness of resulting pixels values in screen-space.

Algorithm 1 Single-hierarchy tree traversal: \mathbf{x}_o is the shading point/pixel, with I_L and I_R as children of the active node.

```

procedure SINGLE( $\mathbf{x}_o, I$ )
  if  $I$  is leaf or ( $\Delta\omega < \epsilon$  and  $\mathbf{x}_o \notin \text{BBOX}(I)$ ) then
     $c \leftarrow$  contribution of  $I$  to  $\mathbf{x}_o$ 
    add  $c$  to  $\mathbf{x}_o$ 
  else
    SINGLE( $\mathbf{x}_o, I_L$ ), SINGLE( $\mathbf{x}_o, I_R$ )

```

4.4.2 Dual Hierarchy for Pixel-Surface Integration.

We leverage a dual hierarchy to avoid traversing the illumination tree at *every* pixel. Similar to the spatial hierarchy of illumination samples in the previous approach, we also cluster pixels in the screen space and traverse two trees simultaneously. Each node in our pixel-tree stores the average world-space position \mathbf{p}_o corresponding to the pixel group, its bounding box, the average normal direction, the average view direction, and the list of pixels covered by the node. This dual-tree approach allows us to evaluate the contribution from a cluster of illumination samples to a cluster of pixels. Algorithm 2 is a pseudocode of our dual-tree approach.

The key difference from the single tree approach is that, at each traversal step, we have a choice of refining the pixel and/or illumination point clusters. For refining clusters of illumination samples, we use a criterion similar to the single tree approach. We always traverse down the tree if bounding volumes of pixels and illumination samples intersect. Otherwise, we decide if we want to keep traversing the tree based on the extended solid angle measure, $\Delta\omega = A_i / \|\mathbf{p}_o - \mathbf{p}_i\|^2$, which uses the average position \mathbf{p}_o of clustered pixels.

Criterion to Refine Pixel Clusters. To refine pixel clusters, we use our frequency analysis to conservatively predict the potential variation in pixels. Given a pixel sampling rate $s_p[i]$ for the i^{th} pixel in a pixel tree node, a conservative estimate of a screen-space filter extent, centered about the node, is

$$P = \rho / \max_i (s_p[i]), \quad (4.6)$$

where ρ is a user-defined parameter that intuitively corresponds to the fraction of captured outgoing radiance required to avoid discontinuity artifacts. For all our examples, we found that using $\rho = 0.75$ produces good results. The ρ setting influences pixel cluster refinement during traversal.

We refine the cluster only if our criterion predicts a high variation of outgoing radiance in the parent node’s pixels (the SHADE routine in Algorithm 2). During shading (SHADE procedure) we do not adaptively refine the illumination cluster and conservatively assume that $\Delta\omega < \epsilon$ is satisfied for all the children nodes. We could alternatively continue refining along the illumination tree for sub-nodes of the pixel tree. However, not refining results in better performance without any noticeable visual artifacts.

Algorithm 2 Dual-hierarchy traversal: S and I are the root nodes of the shading point and illumination trees, with $S_{\{L|R\}}$ and $I_{\{L|R\}}$ their respective left and right children.

```

procedure DUAL( $S, I$ )
  if  $\Delta\omega < \epsilon$  and  $\text{BBOX}(S) \cap \text{BBOX}(I) = \emptyset$  then
    SHADE( $S, I$ )
  else
    if  $S$  is leaf and  $I$  is leaf then
      SHADE( $S, I$ )
    else if  $S$  is leaf then
      DUAL( $S, I_L$ ), DUAL( $S, I_R$ )
    else if  $I$  is leaf then
      DUAL( $S_L, I$ ), DUAL( $S_R, I$ )
    else
      DUAL( $S_L, I_L$ ), DUAL( $S_R, I_L$ )
      DUAL( $S_L, I_R$ ), DUAL( $S_R, I_R$ )

procedure SHADE( $S, I$ )
  if  $\text{LENGTH}(S) < \rho / \text{BANDWIDTH}(S)$  then
    SHADE( $S_L, I$ ), SHADE( $S_R, I$ )
  else
     $c \leftarrow$  contribution of  $I$  to  $S$ 
    add  $c$  to all pixels  $\mathbf{x}$  in  $S$ 

```

4.5 Implementation

We implemented our approach in the G3D Innovation Engine [MMO14] and our results were measured on a 2.8 GHz Intel Core i7 930 with 12 GB of RAM. Both our illumination and pixels




Figure 4.5 – We compare our approach (red) to Jensen and Bulher [JB02] (blue) for different settings of ϵ . We highlight the $\epsilon \in [0.01, 0.2]$ values and consistently reach equal-quality (measured in RMSE; y -axis) in less render time (in seconds; x -axis).

hierarchies are kd-trees, split along the largest bounding volume dimension. Our single- and dual-tree implementations use the same underlying kd-tree structure.

We uniformly sample points on translucent objects with Bowers et al. [BWWM10] blue noise approach, and image-space curvature values are interpolated from object-space values precomputed with the robust curvature estimator of Kalogerakis et al. [KSNS]. In Section 4.3.1 we compute BSSRDF bandwidths as the 95th percentile of the discrete spectrum, where this setting balances numerical stability and accuracy. We use $\rho = 0.75$ (Equation 4.6) in all our scenes and plots, as we found this value avoids discontinuity artifacts while providing good performance. We discuss the performance vs. accuracy trade-offs of ρ and ϵ in Section 4.6.

4.6 Results and Discussions

We have tested our approach on objects with a range of scattering parameters, as well as adapting our frequency analysis to support *several* BSSRDF models: the standard dipole [JMLH01], the “better dipole” [d’E12], and the directional dipole [FHK14]. We use three scenes of increasing radiometric complexity: BUNNY, TOAD, and PICNIK (Figures 4.7, 4.6, and 4.1). TOAD uses the directional dipole, and the remaining scenes use the better dipole.

We compare root mean square error (RMSE) of our technique to the single hierarchy of Jensen and Buhler [JB02], for *total* render time, on the BUNNY and TOAD scenes (Figure 4.5). We sampled ϵ to generate the plots, and our approach consistently reaches equal quality in less time.

Comparisons in the BUNNY scene (Figure 4.7) illustrate our scalability with pixel coverage: the performance discrepancy between the full-view (Figure 4.5a) and zoom-in (Figure 4.5b) renderings is due to the total number of pixels present in the pixel hierarchy. As expected, the benefit of our approach increases with the number of translucent pixels: one can expect our approach to scale sub-linearly here, which is particularly favorable given recent trends towards higher resolution renderings and higher pixel supersampling rates.



Figure 4.6 – The TOAD scene has a bumpy geometry with detailed textures. We compare the difference images of the multiple scattering term against the ground truth for an equal rendering time (196s). The difference images are scaled by 50 for visualization. Our approach achieves more accurate estimation than the single-tree in the same rendering time.




Figure 4.7 – The BUNNY scene. We compare the difference images of the multiple scattering term against the ground truth for an equal rendering time (60s). The difference images are scaled by 200 for visualization. In this example, our approach removes artifacts under the tail and reduces moiré patterns present in the single-tree approach.

Our screen space adaptive sampling rate accounts both distance, local curvature, foreshortening and BSSRDF properties from first principles. Moreover, it properly explains (and subsumes) most of the previously used heuristics in the literature, e.g., depth and normal min/max methods [NW09]. Our sampling rate formulation (Equation 4.4) is simple and only requires the

precomputation of two values (B_θ, B_s) *per material*. We do not require an additional pass to aggregate min/max statistics over the G-buffer.

Limitations. The PICNIK scene (Figure 4.8) is a “failure” case: specifically, our current implementation creates a separate dual tree per object, and since the PICNIK scene includes several (smaller) translucent objects, we only obtain a benefit for a sub-region of the quality/performance range. Moreover, the solid angles $\Delta\omega$ spanned by pixel-tree nodes are more sensitive to errors for small objects and small BSSRDF scales. Since our technique approximates $\Delta\omega$ for a group of pixels, it is sensitive to these scenarios and we plan to address this issue in the future by devising more appropriate $\Delta\omega$ estimates. Overall, the fact that the additional tree construction time is amortized over fewer pixels, and the nature of our non-conservative $\Delta\omega$ estimate in the presence of smaller objects (in image-space), contribute to the suboptimal performance profile in this scene. This also explains the reduced error reduction rate for small ϵ .




Figure 4.8 – The PICNIK scene challenges the assumptions of our work, and we only obtain equal-quality benefit at lower rendering times (albeit enough for visual convergence).

In some difficult scenarios, high frequencies may be missed due to pixel discretization: for instance, a worst-case scenario would involve a camera facing an object with staggered depth discontinuities, which may miss small depth changes due to pixel aliasing. Here, we would group pixels that should not have been grouped.

4.7 Conclusion

We presented a new frequency analysis of BSSRDFs in order to predict the variation of outgoing radiance for multiple subsurface scattered light. We build and traverse a dual hierarchy over illumination samples and pixels using a well-founded refinement strategy that leverages our frequency bandwidth estimates. This yields an adaptive rendering strategy that almost consistently outperforms the state-of-the-art. Moreover, our frequency analysis and bandwidth estimates apply to a variety of existing BSSRDF models with negligible precomputation, our rendering technique scales positively with shading resolution, and all without introducing any additional approximation error.

Our approach leads to several interesting open questions:

1. An interesting avenue would be to combine our work and the one of Arbre et al. [AWB08]. They cluster both light source positions and illumination points at the surface of the object while we cluster both illumination points and shading points. Based on the same multiple cluster idea, it should be possible to build a tree that accounts for those three components during rendering;
2. Our frequency analysis does not account for surface global illumination transport: we ignore visibility, and the use of spatial illumination samples ignores incident radiance variation. Modeling this behavior more accurately could lead to less conservative bandwidth estimates and traversal criteria;
3. There are no reasons why our theory and implementation could not support other existing diffusion models (e.g., *quantized diffusion* [DI11]), and so implementing these models under our framework is interesting even if only for the sake of completeness;
4. Investigating how increases in ϵ should affect our choice of ρ , and vice versa, leads to the interesting question of whether an “optimal” parameter setting for both these values could be computed automatically;

5. Lastly, we are exploring the effects on performance and accuracy of replacing our position-based solid angle approximation with the actual projected solid angle of the underlying surface elements.

Chapter 5

Conclusion

Beginning with a foundational view of modern radiometry in the context of realistic image synthesis, we introduced the mathematical formalisms necessary to understand the behaviour of light transport in dense participating media. Specifically, we first presented the rendering equation that models the *global* interaction of light in scenes comprising solid, opaque reflectors. Here, the *local* interaction of light incident on different reflecting materials (*e.g.*, metals, matte reflectors) is modelled by the bidirectional reflectance distribution function (BRDF). Since analytic solutions to the rendering equation are only possible in simple geometries, we also detailed the most prevalent numerical techniques used to solve the rendering equation in the most general setting. These techniques can roughly be categorized as either sampling-based or density estimation-based solutions.

Since, in the most general case, free-space is occupied by media like fog, dust or smoke, the assumption that light travels unimpeded from surfaces to surfaces no longer holds. To model the behaviour of light in these more representative scenarios that include participating volumetric media, we presented the radiative transport equation (RTE) that models four additional types of interactions in the media: absorption, emission, outscattering, and inscattering. We showed how the differential form of the RTE can be reformulated as an integration problem and, similarly to the surface-only rendering equation, we presented numerical solutions to the RTE integration problem in scene with participating media.

These aforementioned numerical methods suffer from noise and efficiency issues. To address these problems, we presented the concept of diffusion based subsurface scattering to solve the RTE more efficiently and accurately, under important practical circumstances. We have described the generalization of BRDF to bidirectional surface scattering reflectance distribution function (BSSRDF), which does not assume that light enters and exits the surface at the same point (Figure 3.1). This imposes major changes in traditional rendering algorithms, for which we

have explored the state of the art of diffusion theory: the searchlight problem, dipole, quadpole, multipole, quantized diffusion, directional dipole, and several related techniques.

The efficient numerical spatial integration necessary for these various analytic approaches is an important part of making them practical. Since standard sampling approaches are prone to noise, we have reviewed the two-pass hierarchical approach of Jensen and Buhler [JB02], along with its single-pass extension by Arbree et al. [AWB08], which are the most common techniques used to address this spatial integration noise.

Starting from the seminal frequency analysis of light transport by Durand et al. [DHS⁺05], we have devised a novel extension to enable a similar frequency analysis of light transport in the context of subsurface volumetric scattering with BSSRDF models. We have contributed a novel numerical technique to estimate the frequency spectrum of scattered light in translucent media. We demonstrated the practical benefits of our formulation and adapted a novel dual-tree data structure to adaptively evaluate the underlying BSSRDF in joint image- and object-space.

Our frequency analysis has limitations, however: high frequency details may be missed due to pixel discretization limitations. In its practical application, our cluster approximation is sensitive to the size of objects on small BSSRDF scales. Further investigation of our solid angle approximation of the cluster could lead to a more robust solution. Another interesting direction of future work could aim for interactivity with progressive rendering, where pixel clusters could be assigned importance values based on their predicted intensity.

Bibliography

- [Ano15] Anonymous. Dual-tree for adaptive density estimation, 2015.
- [Ant00] V.S. Antyufeev. *Monte Carlo Method for Solving Inverse Problems of Radiation Transfer*. Inverse and ill-posed problems series. VSP, 2000.
- [App68] Arthur Appel. Some techniques for shading machine renderings of solids. In *Proceedings of the April 30–May 2, 1968, Spring Joint Computer Conference, AFIPS '68 (Spring)*, pages 37–45, New York, NY, USA, 1968. ACM.
- [AWB08] Adam Arbree, Bruce Walter, and Kavita Bala. Single-pass scalable subsurface rendering with lightcuts. In *Computer Graphics Forum*, volume 27, pages 507–516. Wiley Online Library, 2008.
- [BBS14] Laurent Belcour, Kavita Bala, and Cyril Soler. A local frequency analysis of light scattering and absorption. *ACM Trans. on Graph.*, 33(5):163:1–163:17, September 2014.
- [Bel12] Laurent Belcour. *A Frequency Analysis of Light Transport: from Theory to Implementation*. PhD thesis, Grenoble Université, 2012.
- [BSS⁺12] Mahdi M. Bagher, Cyril Soler, Kartic Subr, Laurent Belcour, and Nicolas Holzschuch. Interactive rendering of acquired materials on dynamic geometry using bandwidth prediction. In *ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*, pages 127–134, 2012.
- [BWWM10] John Bowers, Rui Wang, Li-Yi Wei, and David Maletz. Parallel poisson disk sampling with spectrum analysis on surfaces. *ACM Trans. on Graph.*, 29(6):166:1–166:10, December 2010.
- [Cha58] Subrahmanyan Chandrasekha. *Proceedings of the National Academy of Sciences of the United States of America*. Number 44. National Academy of Sciences, 1958.

- [Cha60] Subrahmanyan Chandrasekhar. *Radiative Transfer*. Dover Books on Intermediate and Advanced Mathematics. Dover Publications, 1960.
- [CJW⁺09] D. Cline, S. Jeschke, K. White, A. Razdan, and P. Wonka. Dart throwing on surfaces. In *Proceedings of the Twentieth Eurographics Conference on Rendering*, EGSR'09, pages 1217–1226, Aire-la-Ville, Switzerland, Switzerland, 2009. Eurographics Association.
- [CP⁺08] Rui Wang 0004, Ewen Cheslack-Postava, Rui Wang 0003, David P. Luebke, Qianyong Chen, Wei Hua, Qunsheng Peng, and Hujun Bao. Real-time editing and relighting of homogeneous translucent materials. *The Visual Computer*, 24(7-9):565–575, 2008.
- [CPC84] Robert L. Cook, Thomas Porter, and Loren Carpenter. Distributed ray tracing. *SIGGRAPH Comput. Graph.*, 18(3):137–145, January 1984.
- [d'E12] Eugene d'Eon. A better dipole. <http://www.eugenedeon.com/project/a-better-dipole/>, Nov 2012. [accessed 7-April-2015].
- [d'E14] Eugene d'Eon. A dual-beam 3d searchlight bssrdf. In *ACM SIGGRAPH 2014 Talks*, pages 65:1–65:1, 2014.
- [DHS⁺05] Frédo Durand, Nicolas Holzschuch, Cyril Soler, Eric Chan, and François X. Sillion. A frequency analysis of light transport. *ACM Trans. Graph.*, 24(3):1115–1126, July 2005.
- [DI11] Eugene D'Eon and Geoffrey Irving. A quantized-diffusion model for rendering translucent materials. *ACM Trans. Graph.*, 30(4):56:1–56:14, July 2011.
- [DJ05] Craig Donner and Henrik Wann Jensen. Light diffusion in multi-layered translucent materials. *ACM Trans. Graph.*, 24(3):1032–1039, July 2005.
- [DJ08] Craig Donner and Henrik Wann Jensen. Rendering translucent materials using photon diffusion. In *ACM SIGGRAPH 2008 Classes*, SIGGRAPH '08, pages 4:1–4:9, New York, NY, USA, 2008. ACM.

- [EH79] W.G. Egan and T.W. Hilgeman. *Optical properties of inhomogeneous materials: applications to geology, astronomy, chemistry, and engineering*. Academic Press, 1979.
- [FHK14] Jeppe Revall Frisvad, Toshiya Hachisuka, and Thomas Kim Kjeldsen. Directional dipole model for subsurface scattering. *ACM Trans. Graph.*, 34(1):5:1–5:12, December 2014.
- [FPW⁺92] Thomas J Farrell, Michael S Patterson, Brian Wilson, et al. A diffusion theory model of spatially resolved, steady-state diffuse reflectance for the noninvasive determination of tissue optical properties in vivo. *Med. Phys.*, 19(4):879–888, 1992.
- [GE52] S. Glasstone and M.C. Edlund. *The Elements of Nuclear Reactor Theory*. Van Nostrand, 1952.
- [GR87] L. Greengard and V. Rokhlin. A fast algorithm for particle simulations. *J. Comput. Phys.*, (2):325–348, December 1987.
- [GR01] Reindert Graaff and Kees Rinzema. Practical improvements on photon diffusion theory: application to isotropic scattering. *Physics in medicine and biology*, 46(11):3043, 2001.
- [Gro51] C. C. Grosjean. The exact mathematical theory of multiple scattering of particles in an infinite medium. *Memoirs Kon. Vl. Ac. Wetensch*, 36(13), 1951.
- [Gro56] C.C. Grosjean. A high accuracy approximation for solving multiple scattering problems in infinite homogeneous media. *Il Nuovo Cimento*, 3(6):1262–1275, 1956.
- [Gro59] C.C. Grosjean. *Multiple isotropic scattering in convex homogeneous media bounded by vacuum*. Oct 1959.
- [GTGB84] Cindy M. Goral, Kenneth E. Torrance, Donald P. Greenberg, and Bennett Battaile. Modeling the interaction of light between diffuse surfaces. *SIGGRAPH Comput. Graph.*, 18(3):213–222, January 1984.

- [HCJ13a] Ralf Habel, Per H. Christensen, and Wojciech Jarosz. Classical and improved diffusion theory for subsurface scattering. Technical report, Disney Research Zurich, June 2013.
- [HCJ13b] Ralf Habel, Per H. Christensen, and Wojciech Jarosz. Photon beam diffusion: A hybrid monte carlo method for subsurface scattering. In *Eurographics Symposium on Rendering*, pages 27–37, 2013.
- [Her05] Christophe Hery. Implementing a skin bssrdf: (or several...). In *ACM SIGGRAPH 2005 Courses*, SIGGRAPH '05, New York, NY, USA, 2005. ACM.
- [Ish78] Akira Ishimaru. *Wave Propagation and Scattering in Random Media*. IEEE Press-Oxford University, 1978.
- [JB02] Henrik Wann Jensen and Juan Buhler. A rapid hierarchical rendering technique for translucent materials. *ACM Trans. Graph.*, 21(3):576–581, July 2002.
- [JC94] Henrik Wann Jensen and Niels Jørgen Christensen. Photon maps in bidirectional monte carlo ray tracing of complex objects, 1994.
- [JC98] Henrik Wann Jensen and Per H. Christensen. Efficient simulation of light transport in scenes with participating media using photon maps. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '98, pages 311–320, New York, NY, USA, 1998. ACM.
- [Jen96] Henrik Wann Jensen. Global illumination using photon maps. In *Proceedings of the Eurographics Workshop on Rendering Techniques '96*, pages 21–30, London, UK, UK, 1996. Springer-Verlag.
- [Jen01] Henrik Wann Jensen. *Realistic Image Synthesis Using Photon Mapping*. A. K. Peters, Ltd., Natick, MA, USA, 2001.
- [JMLH01] Henrik Wann Jensen, Stephen R. Marschner, Marc Levoy, and Pat Hanrahan. A practical model for subsurface light transport. In *ACM SIGGRAPH*, pages 511–518, 2001.

- [Kaj86] James T. Kajiya. The rendering equation. *SIGGRAPH Comput. Graph.*, 20(4):143–150, August 1986.
- [KB04] Aravind Krishnaswamy and Gladimir V. G. Baranoski. A biophysically-based spectral model of light interaction with human skin. *Comput. Graph. Forum*, 23(3):331–340, 2004.
- [Khi89] A. I. Khisamutdinov. Choosing the “russian roulette and splitting” parameters for monte carlo computation of radiation transfer. *USSR Comput. Math. Math. Phys.*, 29(1):198–203, August 1989.
- [KI98] Arnold D Kim and Akira Ishimaru. Optical diffusion of continuous-wave, pulsed, and density waves in scattering media and comparisons with radiative transfer. *Applied optics*, 37(22):5313–5319, 1998.
- [KL69] G. Kortum and J.E. Lohr. *Reflectance Spectroscopy*. Springer London, Limited, 1969.
- [KSNS] Evangelos Kalogerakis, Patricio Simari, Derek Nowrouzezahrai, and Karan Singh. Robust statistical estimation of curvature on discretized surfaces. In *Eurographics Symposium on Geometry Processing*, pages 13–22.
- [KVH84] James T. Kajiya and Brian P Von Herzen. Ray tracing volume densities. *SIGGRAPH Comput. Graph.*, 18(3):165–174, January 1984.
- [MMO14] Morgan McGuire, Michael Mara, and Others. G3D Innovation Engine, 2014. <http://g3d.sourceforge.net/>.
- [MS67] Thomas Murray MacRobert and Ian Naismith Sneddon. Spherical harmonics : an elementary treatise on harmonic functions, with applications, 1967.
- [MSG05a] S. Menon, Q. Su, and R. Grobe. Determination of g and μ using multiply scattered light in turbid media. *Phys. Rev. Lett.*, 94:153904, Apr 2005.
- [MSG05b] S. Menon, Q. Su, and R. Grobe. Generalized diffusion solution for light scattering from anisotropic sources. *Opt. Lett.*, 30(12):1542–1544, Jun 2005.

- [NGD⁺06] Srinivasa Narasimhan, Mohit Gupta, Craig Donner, Ravi Ramamoorthi, Shree K. Nayar, and Henrik Wann Jensen. Acquiring scattering properties of participating media by dilution. *ACM Transactions on Graphics / SIGGRAPH*, 25(3), August 2006.
- [NRH⁺92] F. E. Nicodemus, J. C. Richmond, J. J. Hsia, I. W. Ginsberg, and T. Limperis. Radiometry. chapter Geometrical Considerations and Nomenclature for Reflectance, pages 94–145. Jones and Bartlett Publishers, Inc., USA, 1992.
- [NW09] Greg Nichols and Chris Wyman. Multiresolution splatting for indirect illumination. In *ACM I3D*, pages 83–90, 2009.
- [OM85] J. C. Oelund and N. J. McCormick. Sensitivity of multiple-scattering inverse transport methods to measurement errors. *J. Opt. Soc. Am. A*, 2(11):1972–1978, Nov 1985.
- [PH10] Matt Pharr and Greg Humphreys. *Physically Based Rendering, Second Edition: From Theory To Implementation*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2nd edition, 2010.
- [RH01] Ravi Ramamoorthi and Pat Hanrahan. An efficient representation for irradiance environment maps. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '01, pages 497–500, New York, NY, USA, 2001. ACM.
- [STPP09] Ying Song, Xin Tong, Fabio Pellacini, and Pieter Peers. SubEdit: A Representation for Editing Measured Heterogeneous Subsurface Scattering. *ACM Transactions on Graphics*, 28(3), August 2009.
- [Tur92] Greg Turk. Re-tiling polygonal surfaces. *SIGGRAPH Comput. Graph.*, 26(2):55–64, July 1992.
- [WABG06] Bruce Walter, Adam Arbree, Kavita Bala, and Donald P Greenberg. Multidimensional lightcuts. *ACM Transactions on Graphics (TOG)*, 25(3):1081–1088, 2006.

- [Whi80] Turner Whitted. An improved illumination model for shaded display. *Commun. ACM*, 23(6):343–349, June 1980.
- [WWW58] A.M. Weinberg, P. Wigner, and E.K. Wigner. *The physical theory of neutron chain reactors*. University of Chicago Press, 1958.
- [WZT⁺08] Jiaping Wang, Shuang Zhao, Xin Tong, Stephen Lin, Zhouchen Lin, Yue Dong, Baining Guo, and Heung-Yeung Shum. Modeling and rendering of heterogeneous translucent materials using the diffusion equation. *ACM Trans. Graph.*, 27(1):9:1–9:18, March 2008.