



HAL
open science

Apparence matérielle : représentation et rendu photo-réaliste

Mahdi Mohammadbagher

► **To cite this version:**

Mahdi Mohammadbagher. Apparence matérielle : représentation et rendu photo-réaliste. Mathématiques générales [math.GM]. Université de Grenoble, 2012. Français. NNT : 2012GRENM063 . tel-00770181

HAL Id: tel-00770181

<https://theses.hal.science/tel-00770181v1>

Submitted on 4 Jan 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

THÈSE

Pour obtenir le grade de

DOCTEUR DE L'UNIVERSITÉ DE GRENOBLE

Spécialité : **Mathématiques Appliquées, Informatique**

Arrêté ministériel :

Présentée par

Mahdi M. Bagher

Thèse dirigée par **Nicolas Holzschuch**
et codirigée par **Cyril Soler**

préparée au sein du **Laboratoire Jean Kuntzmann**
et de **École Doctorale Mathématiques, Sciences et Technologies de l'Information, Informatique de Grenoble**

Material Appearance

Photo-realistic representation and rendering

Thèse soutenue publiquement le ,
devant le jury composé de :

Steve Marschner

Professeur à l'Université Cornell, Ithaca, Rapporteur

Mathias Paulin

Professeur à l'Université Paul Sabatier, Toulouse, Rapporteur

Jaakko Lehtinen

Senior research scientist at NVIDIA Research, Helsinki, Examineur

Florence Forbes

Directeur de recherche, HDR, à INRIA, Grenoble, Examineur

Nicolas Holzschuch

Directeur de recherche, HDR, à INRIA, Grenoble, Directeur de thèse

Cyril Soler

Chargé de recherche, CR1, à INRIA, Grenoble, Co-Directeur de thèse



Material Appearance: photo-realistic representation and rendering

Abstract:

This thesis presents some advances in efficient representation of material appearance in a lighting simulation. The scope of this thesis is two-fold: an interactive shading algorithm to render measured reflectance with dynamic geometry using frequency analysis of light transport and hierarchical shading and up-sampling in deferred shading context, and a new normal distribution function for the Cook-Torrance micro-facet BRDF model, along with a new shadowing and masking function and a generalization of Schlick's approximation of the Fresnel term.

In the first part, we introduce a real-time frequency analysis of light transport framework that allows us to estimate the bandwidth and variance of the shading integrand. The bandwidth and variance are a function of frequencies in the illumination, distance traveled by light, BRDF and texture, and the geometry configuration (curvature). We use this information to under-sample the image, and also use an adaptive number of samples for shading. We devise a single-pass hierarchical shading and up-sampling scheme to assemble an image out of the sparsely shaded image pixels. We extend our interactive technique to use pre-convolved shading for real-time performance. We also take advantage of the bandwidth information to perform multi-sample anti-aliasing in deferred shading by sub-sampling only a small portion of image pixels whose bandwidth is smaller than 1 pixel^{-1} .

In the second part, we propose a new distribution function for the Cook-Torrance micro-facet BRDF, based on our observations on the reflectance measurements. We isolate the distribution components of the reflectance data and directly observe that existing distribution functions are insufficient. Then we devise the Shifted Gamma Distribution (SGD) fitting more accurately to the data. We derive the shadowing and masking function from the distribution. We observe that not all materials have the Fresnel behavior expected by Schlick's approximation. Hence, we generalize the Schlick's approximation to more accurately fit the model to the measurements. We introduce a two-step fitting approach, that fits each RGB channel separately — accounting for wave-length dependent effects. We show that our shading model outperforms existing models and accurately represents a wider range of materials from diffuse to glossy and highly specular materials.

Keywords: Material appearance, Reflectance, BRDF, Frequency analysis, Micro-facet, Cook-Torrance

Apparence Matérielle : représentation et rendu photo-réaliste

Résumé :

Cette thèse présente quelques avancées sur la représentation efficace de l'apparence matérielle dans une simulation de l'éclairage. Nous présentons deux contributions : un algorithme pratique de simulation interactive pour rendre la réflectance mesurée avec une géométrie dynamique en utilisant une analyse fréquentielle du transport de l'énergie lumineuse et le shading hiérarchique et sur-échantillonnage dans un contexte deferred shading, et une nouvelle fonction de distribution pour le modèle de BRDF de Cook-Torrance.

Dans la première partie, nous présentons une analyse fréquentielle de transport de l'éclairage en temps réel. La bande passante et la variance sont fonction de l'éclairage incident, de la distance parcourue par la lumière, de la BRDF et de la texture, et de la configuration de la géométrie (la courbure). Nous utilisons ces informations pour sous-échantillonner l'image en utilisant un nombre adaptatif d'échantillons. Nous calculons l'éclairage de façon hiérarchique, en un seul passage. Notre algorithme est implémenté dans un cadre de *deferred shading*, et fonctionne avec des fonctions de réflectance quelconques, y compris mesurées. Nous proposons deux extensions : pré-convolution de l'éclairage incident pour plus d'efficacité, et anti-aliasing utilisant l'information de fréquence.

Dans la deuxième partie, nous nous intéressons aux fonction de réflectance a base de micro-facette, comme le modèle de Cook-Torrance. En nous basant sur les réflectances mesurées, nous proposons une nouvelle distribution des micro-facettes. Cette distribution, Shifted Gamma Distribution, s'adapte aux donnée avec plus de précision. Nous montrons également comment calculer la fonction d'ombrage et de masquage pour cette distribution. Dans un deuxième temps, nous observons que pour certains matériaux, le coefficient de Fresnel ne suit pas l'approximation de Schlick. Nous proposons une généralisation de cette approximation qui correspond mieux aux données mesurées. Nous proposons par ailleurs une nouvelle technique d'optimisation, canal par canal, en deux étapes. Notre modèle est plus précis que les modèles existants, du diffus au spéculaire.

Keywords : Apparence matérielle, réflectance, BRDF, analyse fréquentielle, micro-facette, Cook-Torrance

Acknowledgments

It is often said that writing a Ph.D. is a work of many, and for this thesis, this is also the case. My sincere thanks to my advisors Nicolas Holzschuch and Cyril Soler for their immense support and invaluable guidance, as well as my co-authors Kartic Subr and Laurent Belcour for their assistance. I am grateful to the people at Maverick/Artis team for creating such great work environment. My sincere gratitude to my family for their unconditional love and care, and to my friends who supported me in one way or another in the past three years.

Contents

I Introduction	1
I.1 Motivation	1
I.2 List of publications	3
I.3 Organization	3
II Introduction (français)	5
II.1 Motivation	5
II.2 Liste des publications	7
II.3 Organization	7
III Background and previous work	9
III.1 Introduction (français)	10
III.2 Introduction	10
III.3 Light transport	10
III.3.1 Radiometric terms	11
III.3.1.1 Radiant power (flux)	11
III.3.1.2 Radiance (L)	11
III.3.1.3 Irradiance (E) and Radiosity (B)	12
III.3.1.4 Radiant intensity (I)	12
III.3.2 The rendering equation	12
III.3.3 The Monte Carlo method	13
III.3.4 Monte Carlo importance sampling	14
III.3.5 Hardware-accelerated rendering	14
III.3.5.1 The rendering pipeline	14
III.3.5.2 Forward vs. deferred shading	15
III.3.6 Multi-resolution screen-space rendering	15
III.3.7 Light field	17
III.3.7.1 Light field propagation	18
III.3.7.2 Light field parameterization	18
III.3.8 Frequency analysis of light transport	18
III.3.8.1 Transport in free space	19
III.3.8.2 Reflection	19
III.3.8.3 Shading	20
III.3.8.4 Texture mapping	20
III.3.8.5 Occlusion	21
III.3.8.6 Depth of field	21
III.3.8.7 Motion blur	21
III.4 Material appearance	21
III.4.1 Definitions	21
III.4.1.1 Reflectance	21

III.4.1.2	BRDF	22
III.4.2	Taxonomy of appearance	23
III.4.3	Reflectance acquisition	23
III.4.3.1	Acquisition techniques	25
III.4.3.2	Isotropic BRDF databases	26
III.4.4	Reflectance representation	27
III.4.4.1	Tabulated BRDF	27
III.4.4.2	BRDF re-parameterization	28
III.4.4.3	Dimensionality reduction and compression techniques	29
III.4.4.4	Analytical models	30
III.4.4.5	Comparison between existing analytical models	37
III.5	Summary and conclusion	38
IV	Interactive rendering of reflectance using frequency analysis	41
IV.1	Introduction (français)	41
IV.2	Introduction	43
IV.2.1	Overview	44
IV.3	Real-time frequency analysis	45
IV.3.1	Illumination, textures, and BRDF bandwidth	45
IV.3.2	Computing one-bounce 2D bandwidth	49
IV.3.3	Image-space bandwidth and sampling rate	51
IV.3.4	Adaptive sampling for shading	51
IV.3.5	Adaptive sampling for pre-convolved shading	52
IV.3.6	Implementation roadmap	53
IV.4	Hierarchical shading algorithm	53
IV.4.1	Bandwidth buffer initialization	54
IV.4.2	Shading and up-sampling	55
IV.4.3	Shading computation	55
IV.5	Results and discussion	56
IV.5.1	Behavior of our algorithm	56
IV.5.2	Comparison with related work	59
IV.5.3	Rendering results	59
IV.5.4	Discussion	61
IV.6	Conclusion	62
V	A new distribution for Cook-Torrance BRDF	65
V.1	Introduction (français)	65
V.2	Introduction	66
V.3	Our reflectance model	67
V.3.1	Data observation	67
V.3.2	The SGD distribution	69
V.3.3	Shadowing and masking	70
V.3.4	The Fresnel term	71
V.3.5	Importance sampling	72

V.4	Fitting algorithm	73
V.4.1	Fitting RGB channels separately	73
V.4.2	Non-linear optimization	73
V.5	Results and Comparison	74
V.5.1	Visual comparison	74
V.5.2	BRDF Lobes comparison	77
V.5.3	Quantitative Error Measure	77
V.5.4	Influence of D , G , and F parameters in our BRDF model	80
V.5.5	Fitting RGB channels altogether for D parameters	81
V.5.6	Timings	83
V.5.7	Limitations	83
V.6	Conclusion	85
VI	Conclusion	87
VII	Conclusion (français)	89
A	Appendix A: SGD fitting results	91
	Bibliography	93

List of Figures

I.1	Examples of rendering applications.	2
III.1	The geometry of radiance.	11
III.2	The rendering equation.	13
III.3	The GPU rendering pipeline for OpenGL 4 and DirectX 11.	16
III.4	The low frequency nature of indirect illumination.	17
III.5	Light transport.	18
III.6	Two popular 4D light field parameterizations.	19
III.7	Local shading for a curved receiver with arbitrary BRDF.	20
III.8	BRDF.	22
III.9	Taxonomy of appearance.	24
III.10	Gonio-reflectometer for isotropic BRDF measurement.	25
III.11	Image-based isotropic BRDF measurement.	26
III.12	Gonio-reflectometer for anisotropic BRDF measurement.	27
III.13	The MERL-MIT isotropic BRDF database.	28
III.14	The Cornell BRDF database.	28
III.15	The half/difference-angle parameterization.	29
III.16	A 2D factorization example.	30
III.17	The local coordinate system	30
III.18	The geometry of reflection.	31
III.19	The vectors and angles for the albedo corrected Ward-Dür model.	32
III.20	Vectors for the Lafortune model.	33
III.21	Micro-facets for Cook-Torrance BRDF.	34
III.22	Comparison between several analytical models for <i>beige-fabric</i>	38
III.23	Comparison between several analytical models for <i>fruitwood-241</i>	39
III.24	Comparison between several analytical models for <i>chrome</i>	39
IV.1	Flatland illustration of local bandwidth propagation.	45
IV.2	4D and 2D local light-field parametrization.	46
IV.3	Review of spectral operations.	46
IV.4	Most wavelets have a band-limited spectrum.	47
IV.5	Instant 2D bandwidth estimation using wavelets.	47
IV.6	Input environment map and its local angular bandwidth.	48
IV.7	area-light source examples and their bandwidth.	48
IV.8	Input BRDF (<i>gold-paint</i>) and its local angular bandwidth.	49
IV.9	Matrix operators on 2D bandwidth.	50
IV.10	Combining bandwidth estimates from sampled incident directions.	52
IV.11	Our rendering pipeline.	54
IV.12	Up-sampling interpolation scheme.	55
IV.13	Three identical objects with different materials.	56

IV.14	Rendering times.	57
IV.15	Computation times.	57
IV.16	Validation of predicted variance and bandwidth.	58
IV.17	Effect of the number of sampled directions on time and quality.	59
IV.18	Comparison of our algorithm with a reference.	60
IV.19	Comparison with related work.	61
IV.20	Adaptive sampling for pre-convolved shading.	62
IV.21	area-light source rendering examples.	63
IV.22	Application of our algorithm to dynamic geometry.	64
IV.23	Adaptive multi-sample anti-aliasing using bandwidth information.	64
V.1	Cook-Torrance BRDF using Beckmann and SGD distributions.	66
V.2	The micro-facet probability distribution for <i>brass</i>	67
V.3	The SGD distribution is an accurate fit for a wide range of materials.	68
V.4	The influence of parameters α and p	69
V.5	The approximation to shadowing term G_1 for SGD.	70
V.6	The Fresnel term extracted from the measured data.	71
V.7	Our Fresnel approximation.	72
V.8	Uniform and importance sampling for <i>specular-violet-phenolic</i>	73
V.9	Comparison between BRDF models for <i>aventurine</i>	75
V.10	Comparison between BRDF models for <i>colonial-maple-223</i>	76
V.11	BRDF lobes with Beckmann, GGX and SGD distributions for <i>hematite</i>	78
V.12	Normalized average fitting errors (logarithmic scale).	79
V.13	Pixel-by-pixel comparison for <i>aventurine</i>	80
V.14	The influence of D , G , and F parameters for <i>beige-fabric</i>	82
V.15	The influence of D , G , and F parameters for <i>natural-209</i>	82
V.16	The influence of D , G , and F parameters for <i>specular-violet-phenolic</i>	83
V.17	Fitting errors of the SGD distribution with different fitting methods.	84
V.18	An example of a two-layer material.	85
V.19	A failure case for our fitting algorithm.	85

List of Tables

III.1 Radiometric terms.	12
III.2 P_{22} functions corresponding to classical micro-facet distributions, with $x = \tan^2 \theta$	35

Introduction

Contents

I.1 Motivation	1
I.2 List of publications	3
I.3 Organization	3

In this chapter we give a brief introduction about the main goals of this thesis. It lists the publications presented in this thesis and concludes with a short overview of the chapters that will follow.

I.1 Motivation

Computers, as powerful calculators, can take descriptions — such as 3D shapes, materials and lights — of an environment as well as rules about how — physically or artistically — light and matter interact and generate pictures. This process is called *image synthesis* or *rendering* and is a major area of research in the field of *computer graphics*. Such synthesized images can be seen as visual stimuli causing a perceptual response in the viewer, as if the viewer is immersed in that virtual environment.

In the case of photo-realistic rendering, accurate representation of a virtual scene — geometry and appearance descriptions — as well as physically correct illumination simulation are very important. Hence, computer graphics is a multi-disciplinary area of research, borrowing theories and techniques from optics — physical and geometrical —, biology, psychology, art, mathematics, and computer and engineering sciences.

Computers, no matter how powerful, are limited in their computational resources. And, graphics simulations are extremely costly. Therefore, computation time in rendering is of major concern. Efficient rendering algorithms maximize the use of computation resources and minimize calculations to deliver images as fast as possible to the viewer, allowing interactivity. Typically, interactive rates are achieved at the cost of losing some degree of accuracy by approximations to the real-world phenomena being simulated.

The geometric description of a scene can be either mathematically expressed or explicitly defined by surface representations such as polygonal meshes. The appearance description of a surface defines its material characteristics mathematically in a way that a computer program can predict how light interacts with it. This involves defining how much of the incident light is reflected, refracted, and absorbed when light is hit by a surface. The

rendering algorithm is concerned with rules about how these descriptions can be efficiently combined in a lighting simulation to generate pictures.

Image synthesis has a wide range of applications such as visual effects, video games, virtual prototyping (e.g. automotive design), flight simulations, astrology, archeology, cartography, education, scientific visualizations, user-interfaces, and digital arts, just to name a few. Some applications require real-time interactive feedback such as video games and flight simulations, which implies fast rendering times. And, other applications require extreme accuracy and realism such as visual effects and automotive design. See Figure I.1 for two examples.



(a) Real-time video game example

(From *Call of Duty: Black Ops II*)



(b) Realistic automotive design example

(From *Author's personal portfolio*)

Figure I.1: Examples of rendering applications.

Existing geometric descriptions and rendering algorithms are sophisticated enough to produce photo-realistic images in finite time with current computation resources. The limiting factor, for the time being, is the lack of an efficient way to accurately transfer real-world material appearance to digital models. *Digital appearance modeling* is about capturing and simulating how material and light interact within an environment for a given surface point. That depends on the nature of illumination, material characteristics, and the geometric profile of the surface. A deep insight to material appearance is not only beneficial in image synthesis but also in inverse problems such as object recognition and 3D reconstruction from video or photographs.

Material appearance can be very complex, e.g. non-homogeneous, spatially varying and multi-layered. Real-world materials are also capable of exhibiting complex behavior such as sub-surface scattering and time or wavelength-dependent effects. We perceptually distinguish materials from their color and texture, as well as various degrees of gloss, haziness, luster, matte, anisotropy, translucency, and transparency. Unfortunately, these qualitative factors are difficult to measure quantitatively for general material types. And, it is unknown how many perceptual factors are involved in the qualitative perception of appearance. Therefore, to numerically capture a material's appearance, one should measure how much of the incident light is scattered for various wavelengths and times, for every possible incident and viewing direction, for every point on the surface. This is a multi-dimensional quantity, 3 to 12 dimensions, requiring very sophisticated measurement

devices for acquisition, and special treatment for rendering with such vast amount of acquired data. Current acquisition techniques are subject to inaccuracy and noise. Dense acquisition of a single point on a single material sample may take hours to several days. The data size for a single material acquisition can be as small as 33 MB for an isotropic homogeneous material — three dimensional quantity — to several gigabytes for capturing higher dimensional scattering functions. The current appearance representation models are subject to inaccuracies and they lack generality — they don't efficiently support a wide range of material types.

In this thesis, we attack this problem by two different approaches. First, we look into how we can render material appearance in an interactive environment using acquired reflectance data for accurate shading, while avoiding pre-computations based on geometry. This requires a framework to assess the light transport from the light source to the viewer and predict the optimal number of samples required for shading. We use frequency analysis to account for many factors affecting light transport including illumination properties, material properties, and geometric features. Second, we investigate a more accurate representation of measured reflectance using a shifted gamma micro-facet distribution. This reflectance model outperforms all the existing parametric models for isotropic homogeneous materials.

I.2 List of publications

Most of the contributions presented in this thesis have been presented in the following publications:

- 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 (I3D), pages 127–134, Costa Mesa, United States, March 2012. [M. Bagher 2012b]
- Mahdi M. Bagher, Cyril Soler, Kartic Subr, Laurent Belcour and Nicolas Holzschuch. *Interactive rendering of acquired materials on dynamic geometry using frequency analysis*. IEEE Transactions on Visualization and Computer Graphics, 2013. Invited paper. [M. Bagher 2013]
- Mahdi M. Bagher, Cyril Soler and Nicolas Holzschuch. *Accurate fitting of measured reflectances using a Shifted Gamma micro-facet distribution*. Computer Graphics Forum, vol. 31, no. 4, June 2012. [M. Bagher 2012a]

I.3 Organization

In the Chapter III, we are going to briefly review the existing body of work related to this thesis. Then, in Chapter IV, we present an interactive technique for rendering acquired materials using frequency analysis. In Chapter V, we introduce a new micro-facet distribution for accurate shading of measured reflectance using a shifted gamma distribution. Finally, we conclude our contributions in Chapter VI.

Introduction (français)

Contents

II.1 Motivation	5
II.2 Liste des publications	7
II.3 Organization	7

Dans ce chapitre, nous donnons une brève introduction sur les objectifs principaux de cette thèse. Il dresse la liste des publications présentées dans cette thèse et se termine par un bref aperçu des chapitres qui vont suivre.

II.1 Motivation

Outils de calcul puissants, les ordinateurs peuvent à partir de formes 3D, de matériaux et de lumières — d'un environnement, ainsi que des règles sur la façon dont — physiquement ou artistiquement — lumière et la matière interagissent et générer des images. Ce processus est appelé synthèse d'image qui est un secteur important de recherche dans le domaine de l'infographie. Ces images de synthèse peuvent être considérées comme des stimuli visuels provoquant une réponse perceptive chez le spectateur, comme si celui-ci était immergé dans cet environnement virtuel.

Dans le cas de rendu photo-réaliste, une représentation précise d'une scène virtuelle — géométrie et les descriptions d'apparence — ainsi que la simulation d'éclairage physiquement correcte sont très importants. Par conséquent, l'infographie est un domaine multidisciplinaire de recherche, qui emprunte des théories et des techniques au domaine de l'optique — physique et géométrique —, la biologie, la psychologie, l'art, les mathématiques et informatique et sciences de l'ingénieur.

Aussi puissant, qu'ils puissent être, les ordinateurs sont limités dans leurs ressources de calcul. Et, les simulations graphiques sont extrêmement coûteuses. Par conséquent, le temps de calcul de rendu est un critère majeur. Algorithmes de rendu efficace maximiser l'utilisation des ressources calcul et minimiser les calculs de délivrer des images le plus rapidement possible pour le spectateur, ce qui permet l'interactivité. En règle générale, les temps interactifs sont obtenus au prix d'une certaine perte de précision par des approximations des phénomènes du monde réel simulés.

La description géométrique d'une scène peut être soit exprimée mathématiquement ou explicitement définie par les représentations de surface tels que des maillages polygonaux.

La description de l'apparence d'une surface définir des caractéristiques matérielles mathématiquement de manière qu'un programme d'ordinateur puisse prédire comment la lumière interagit avec celui-ci. Il s'agit de définir la quantité de lumière incidente qui est réfléchie, réfractée, et absorbée lorsque la lumière rencontre par une surface. L'algorithme de rendu porte sur la façon dont ces descriptions peuvent être efficacement combinées dans une simulation de l'éclairage pour générer des images.

La synthèse d'image a une large gamme d'applications telles que les effets spéciaux, les jeux vidéo, le prototypage virtuel (par exemple, la conception automobile), des simulateurs de vol, l'astronomie, l'archéologie, la cartographie, l'éducation, la visualisation scientifique, les interfaces utilisateur et les arts numériques, pour n'en nommer que quelques-uns. Certaines applications nécessitent une rétroaction temps réel, tels que les jeux vidéo et les simulateurs de vol, ce qui implique un rendu rapide. Et, d'autres applications requièrent précision extrême et un réalisme élevé comme les effets visuels et le design automobile.

Les descriptions géométriques existantes et les algorithmes de rendu sont suffisamment sophistiqués pour produire des images photo-réalistes en temps raisonnable avec les ressources de calcul actuels. Le facteur limitatif, pour le moment, est le manque de moyen efficace de transférer l'aspect matériel du monde réel avec précision à des modèles numériques. La modélisation numérique de l'apparence consiste en la capture et la simulation de la lumière interagissent dans un environnement d'un point de surface donnée. Cette dépend de la nature de l'éclairage, des caractéristiques des matériaux, et du profil géométrique de la surface. Une compréhension profonde de l'aspect matériel est non seulement bénéfique en synthèse d'images, mais aussi dans les problèmes inverses tels que la reconnaissance d'objets et la reconstruction 3D à partir de vidéo ou photos.

L'apparence matérielle peut être très complexe, par exemple non-homogène, variant dans l'espace et à plusieurs couches. Les matériaux du monde réel sont également susceptibles d'avoir un comportement complexe tel que la diffusion sous-surface, et montre une dépendance en temps ou en longueur d'onde. En tant qu'humains, nous pouvons distinguer les matières, leur couleur et leur texture, ainsi que divers degrés de brillance, le flou, le lustre, le mat, l'anisotropie, la translucidité et la transparence. Malheureusement, ces facteurs qualitatifs sont difficiles à mesurer quantitativement pour les types de matériaux généraux. Et, il est inconnu combien de facteurs perceptifs sont impliqués dans la perception qualitative de l'apparence. Par conséquent, pour capturer numériquement l'apparence d'un matériau, il convient de mesurer quelle quantité de lumière incidente est diffusée pour différentes longueurs d'onde, pour chaque direction incidente et chaque direction du regard possible, pour chaque point à la surface. C'est une quantité multi-dimensionnelle, 3 à 12 dimensions, nécessitant des appareils de mesure très sophistiqués pour l'acquisition et le traitement mais également pour le rendu avec une telle grande quantité de données acquises. Les techniques d'acquisition courantes sont également soumises à des inexactitudes et à du bruit. L'acquisition dense d'un point unique sur un échantillon de certains matériaux peut prendre des heures voire plusieurs jours. Le volume des données pour une acquisition d'un seul matériau peut être aussi faible que 33 Mo pour un matériau homogène et isotrope — quantité trois dimensionnelle — à plusieurs giga-octets pour la saisie des fonctions de diffusion. Les modèles actuels de représentation d'apparence sont sujettes à des inexactitudes et ils manquent de généralité — ils n'ont pas la possibilité de représenter efficacement

un large éventail de types de matériaux.

Dans cette thèse, nous attaquons ce problème par deux approches différentes. Premièrement, nous examinons la façon dont nous pouvons rendre l'aspect matériel dans un environnement interactif en utilisant les données de réflectance acquises pour l'ombrage précis, en évitant des pré-calculs basés sur la géométrie. Cela nécessite un cadre pour évaluer le transport de la lumière de la source lumineuse à l'observateur et de prédire le nombre optimal d'échantillons requis pour l'ombrage. Nous utilisons l'analyse fréquentielle pour tenir compte des nombreux facteurs affectant transport de la lumière y compris les propriétés d'éclairage, les propriétés des matériaux et les caractéristiques géométriques. Deuxièmement, nous étudions une représentation plus précise de la réflectance mesurée à l'aide d'une *shifted gamma micro-facet distribution*. Ce modèle de réflectance surpasse tous les modèles paramétriques existants pour des matériaux homogènes isotropes.

II.2 Liste des publications

Majeure partie des contributions présentées dans cette thèse ont été présentées dans les publications suivantes :

- 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 (I3D), pages 127–134, Costa Mesa, United States, March 2012. [M. Bagher 2012b]
- Mahdi M. Bagher, Cyril Soler, Kartic Subr, Laurent Belcour and Nicolas Holzschuch. *Interactive rendering of acquired materials on dynamic geometry using frequency analysis*. IEEE Transactions on Visualization and Computer Graphics, 2013. Invited paper. [M. Bagher 2013]
- Mahdi M. Bagher, Cyril Soler and Nicolas Holzschuch. *Accurate fitting of measured reflectances using a Shifted Gamma micro-facet distribution*. Computer Graphics Forum, vol. 31, no. 4, June 2012. [M. Bagher 2012a]

II.3 Organization

Dans le Chapitre III, nous allons d'examiner brièvement le corpus existant de travaux relatifs à cette thèse. Puis, en Chapitre IV, nous présentons une technique interactive pour rendre des matériaux acquis en utilisant une analyse fréquentielle. Dans le Chapitre V, nous introduisons une nouvelle distribution de micro-facettes pour l'estimation précise d'une réflectance mesurée, à l'aide d'une *shifted gamma distribution*. Enfin, nous concluons notre contribution au Chapitre VI.

Background and previous work

Contents

III.1 Introduction (français)	10
III.2 Introduction	10
III.3 Light transport	10
III.3.1 Radiometric terms	11
III.3.1.1 Radiant power (flux)	11
III.3.1.2 Radiance (L)	11
III.3.1.3 Irradiance (E) and Radiosity (B)	12
III.3.1.4 Radiant intensity (I)	12
III.3.2 The rendering equation	12
III.3.3 The Monte Carlo method	13
III.3.4 Monte Carlo importance sampling	14
III.3.5 Hardware-accelerated rendering	14
III.3.5.1 The rendering pipeline	14
III.3.5.2 Forward vs. deferred shading	15
III.3.6 Multi-resolution screen-space rendering	15
III.3.7 Light field	17
III.3.7.1 Light field propagation	18
III.3.7.2 Light field parameterization	18
III.3.8 Frequency analysis of light transport	18
III.3.8.1 Transport in free space	19
III.3.8.2 Reflection	19
III.3.8.3 Shading	20
III.3.8.4 Texture mapping	20
III.3.8.5 Occlusion	21
III.3.8.6 Depth of field	21
III.3.8.7 Motion blur	21
III.4 Material appearance	21
III.4.1 Definitions	21
III.4.1.1 Reflectance	21
III.4.1.2 BRDF	22
III.4.2 Taxonomy of appearance	23

III.4.3	Reflectance acquisition	23
III.4.3.1	Acquisition techniques	25
III.4.3.2	Isotropic BRDF databases	26
III.4.4	Reflectance representation	27
III.4.4.1	Tabulated BRDF	27
III.4.4.2	BRDF re-parameterization	28
III.4.4.3	Dimensionality reduction and compression techniques	29
III.4.4.4	Analytical models	30
III.4.4.5	Comparison between existing analytical models	37
III.5	Summary and conclusion	38

III.1 Introduction (français)

Un riche corpus de travaux existe déjà sur la simulation de l'éclairage et des images réalistes. Dans ce chapitre, nous allons explorer brièvement les travaux existants et donner les définitions des concepts que nous allons traiter dans les chapitres suivant. Plus précisément, nous allons expliquer comment la simulation de l'éclairage se fait dans un environnement synthétique. Nous allons brièvement expliquer les phénomènes de transport de la lumière et ses implications. Et enfin, nous allons exposer les concepts liés à l'aspect matériel, l'acquisition de réflexion, et de la représentation.

III.2 Introduction

A rich body of work about lighting simulation and realistic imagery already exists. In this chapter, we are going to briefly explore the existing related work and also give definitions of the concepts we will be dealing with in the up-coming chapters. More specifically, we are going to explain how lighting simulation is done within a synthetic environment. We will briefly explain the light transport phenomena and its implications. And finally, we will explore concepts related to material appearance, reflectance acquisition, and representation.

III.3 Light transport

According to *quantum mechanics*, light exhibits a dual behavior by being both a particle and a wave [Feynman 1988]. On the one hand, *particle optics* models light as particles called *photons*. Each photon carries some energy as it is traveling through the scene. Tracing photons, being shot from a light source, is a fundamental solution to lighting simulation. On the other hand, *wave optics* characterizes light as electromagnetic waves to describe light's interaction with matter. Apart from the above mentioned, *geometrical optics* or *ray optics* abstracts light propagation in terms of *rays*. We can think of a ray as a stream of photons traveling along the same path. The over-simplification in ray optics results in ignoring




Figure III.1: The geometry of radiance.

wave-like optical effects such as interference and diffraction, e.g. the rainbow colors on CDs. Since such wave-like effects rarely happen in everyday life, it is usually ignored or treated separately in computer graphics. [Born 1998] is an excellent reference to the main principles of optics.

III.3.1 Radiometric terms

Radiometric terms describe units of measurement for physical quantities; light in our case. Radiometric terms are a function of wave-length, but for the sake of simplicity we omit the wavelength dependence. Light as an electromagnetic energy is measured in joules (J) and is called the *radiant energy* (Q). Each particle carries some joules of energy.

III.3.1.1 Radiant power (flux)

To measure the rate of flow of energy we define the *radiant power* or *flux* (Φ), which is defined as the energy per time or Watt (W).

$$\Phi = \frac{dQ}{dt} \quad \left[\text{Watt}(W) \equiv \frac{J}{s} \right]$$

III.3.1.2 Radiance (L)

To describe how much light is traveling in a particular direction (solid angle), we define *radiance* (L) in $\frac{W}{m^2 sr}$ as the flux per unit area per unit solid angle.

$$L(x, \omega) = \frac{d^2\Phi}{\cos\theta d\omega dA} \quad \left[\frac{W}{m^2 sr} \right]$$

Figure III.1 depicts the geometry of radiance.

What we perceive as brightness in a scene is initially due to radiance. In other words, radiance measures how much flux is received by a viewer — the eye or a camera — depending on the solid angle subtended by the optical system.

Definition	Name	Unit
$\Phi = \frac{dQ}{dt}$	Radiant power (Flux)	$Watt(W) \equiv \frac{J}{s}$
$L(\mathbf{x}, \omega) = \frac{d^2\Phi}{\cos\Theta d\omega dA}$	Radiance	$\frac{W}{m^2 sr}$
$E(\mathbf{x}) = \frac{d\Phi_i}{dA}$	Irradiance (radiant emittance)	$\frac{W}{m^2}$
$B(\mathbf{x}) = \frac{d\Phi_o}{dA}$	Radiosity (radiant exitance)	$\frac{W}{m^2}$
$I(\omega) = \frac{d\Phi}{d\omega}$	Radiant intensity	$\frac{W}{sr}$

Table III.1: Radiometric terms.

III.3.1.3 Irradiance (E) and Radiosity (B)

We also need to measure the flow of energy toward or away from a surface patch. It can be measured by integrating the incident or outgoing flux over all directions per unit area of the surface.

Irradiance (E) measures the total flux entering a surface point per unit area, $\frac{W}{m^2}$.

$$E(\mathbf{x}) = \frac{d\Phi_i}{dA} \left[\frac{W}{m^2} \right]$$

Radiosity (B) measures the total flux exiting a surface point per unit area, $\frac{W}{m^2}$.

$$B(\mathbf{x}) = \frac{d\Phi_o}{dA} \left[\frac{W}{m^2} \right]$$

III.3.1.4 Radiant intensity (I)

In contrast to irradiance (E) and radiosity (B), *radiant intensity (I)* is the quantity describing the flux entering or existing a surface point with respect to a solid angle instead of an area.

$$I(\omega) = \frac{d\Phi}{d\omega} \left[\frac{W}{sr} \right]$$

A summary of radiometric terms and units appears in Table III.1. Please refer to [Glassner 1994] for more details.

III.3.2 The rendering equation

To estimate the radiance leaving an opaque surface point, [Kajiya 1986] introduced the rendering equation (See Figure III.2):

$$L_o(x, \omega, \lambda, t) = L_e(x, \omega, \lambda, t) + \int_{\Omega^+} f_r(x, \omega', \omega, \lambda, t) L_i(x, \omega', \lambda, t) (-\omega' \cdot \mathbf{n}) d\omega'$$

This equation is an integral over the visible hemisphere for a given surface point x , in viewing direction (ω), given a wavelength (λ) and time (t). L_e is the emitted light. f_r is the




Figure III.2: The rendering equation describes the total amount of light emitted from a point x along a particular viewing direction, given a function for incoming light and a BRDF. (From *Wikipedia*)

bidirectional reflectance distribution function (BRDF), specifying the ratio of the reflected radiance exiting in direction ω to incident irradiance from direction ω' . And finally, $(-\omega' \cdot \mathbf{n})$ is the attenuation factor due to the incident angle — the cosine term.

III.3.3 The Monte Carlo method

The Monte Carlo method [Metropolis 1949], was developed to numerically solve quantitative problems using repetitive stochastic sampling. It has many applications in various fields of science such as physics, mathematics, and engineering — from fluid simulation and oil discovery to business risk calculation. For example, in mathematics, Monte Carlo can be used to evaluate multidimensional definite integrals with complicated boundary conditions. Please refer to [Kalos 1986, Spanier 1969, Hammersley 1965] for an introduction to Monte Carlo methods. In computer graphics, it's widely used to estimate the value of high dimensional definite integrals such as in the rendering equation. Using the Monte Carlo method to estimate an integral, we repeatedly evaluate the integral using random samples and average the results until it converges. Monte Carlo ray-tracing [Kajiya 1986, Veach 1995, Ward 1988, Ward 1992b], bidirectional path tracing [Lafortune 1996], stochastic radiosity [Bekaert 1999] and photon mapping [Jensen 1996] are examples of the Monte Carlo method's applications in computer graphics.

As mentioned by [Veach 1998], Monte Carlo integration is about estimating a definite integral of a real valued function using random sampling:

$$I = \int_{\Omega} f(x) d\mu(x)$$

If we take N samples $[X_1, \dots, X_N]$ of a real-valued random variable x according to a probability density function $p(x)$, then the Monte Carlo estimator for the integral I will be:

$$F_N = \frac{1}{N} \sum_{i=1}^N \frac{f(X_i)}{p(X_i)}$$

The convergence rate of the Monte Carlo integration depends on the variance of the estimator:

$$\text{Var} [F_N] = \text{Var} \left[\frac{1}{N} \sum_{i=1}^N \frac{f(X_i)}{p(X_i)} \right] = \frac{1}{N^2} \sum_{i=1}^N \text{Var} \left[\frac{f(X_i)}{p(X_i)} \right] = \frac{1}{N} \text{Var} \left[\frac{f}{p} \right]$$

The larger the variance, the more samples we need for the integration to converge.

III.3.4 Monte Carlo importance sampling

The larger the variance, the more noise is in the rendered images using a fixed number of samples. Several variance reduction techniques exist. They include increasing the number of samples, using stratified sampling instead of uniform sampling, or choosing a density function $p(x)$ close to the integrand $f(x)$.

The latter is called Monte Carlo importance sampling and is widely used in computer graphics. It helps the numerical estimation to converge with a smaller sample set, and therefore faster convergence, without introducing any bias. It is important to choose a $p(x)$ with a convenient method of generating samples from it. The importance function can simply be an approximation of the integrand or some factors of it. Using the exact integrand as the importance function will result in zero variance. But in most cases, the integrand is unknown.

In the case of the rendering equation, ideally, the importance function should be the product of the BRDF, the incident radiance, and the cosine factor. but unfortunately, this product function is unknown. For diffuse shading with low-frequency illumination, uniform sampling or importance sampling based on the cosine term is sufficient. However, for glossy and specular materials with all-frequency illumination, it requires a very large number of samples to converge.

Since the value of the product with respect to the parameters are not known beforehand, it is possible to importance sample based on either the BRDF or incident illumination to reduce the variance of the integrand to some extent. If the BRDF and the incident illumination are both high frequency, the number of samples for the integration to converge will still be very high. Multiple importance sampling [Veach 1998] is a workaround to combine multiple sampling strategies to reduce variance without introducing any bias.

III.3.5 Hardware-accelerated rendering

Solving the rendering equation with Monte Carlo sampling for every pixel in a 2D image is extremely expensive and therefore slow. Fortunately, this computation is highly parallelizable — calculations can be carried out simultaneously —, since the same operations have to be executed millions of times. A *graphics processing unit (GPU)* is a many-cores dedicated hardware for parallel graphics processing. Commodity graphics hardware are widely available on any consumer desktop computer these days.

III.3.5.1 The rendering pipeline

GPUs have programmable rendering pipelines. The rendering pipeline consists of taking a 3D scene, made of *polygons* for example, and outputting a 2D rendered view of the scene. Graphics programming APIs such as OpenGL and Direct3D help programmers program graphics hardware without having to write in assembly language. Figure IV.11 shows an overview of the rendering pipeline for an state of the art GPU supporting OpenGL 4 and DirectX 11. In this pipeline, some but not all stages are programmable. The three

main programmable stages are vertex, geometry, and fragment/pixel programs operating on vertices, primitives, and pixels accordingly.

III.3.5.2 Forward vs. deferred shading

In a classical *forward shading* approach, all geometry must be rendered independently for each light source in the scene including the hidden surfaces. This is somewhat wasteful and therefore slow for complex shading and lighting. *Deferred shading* [Deering 1988, Saito 1990] is a practical solution that decouples scene geometry from lighting. It defers the lighting until all the geometry is rendered into geometry buffers called the *G-buffers*. The G-buffers contain lighting properties such as the position and normal for each pixel. Hence, lighting can be applied as a 2D post process, allowing to render many lights without significant performance loss. The G-buffers can also be used for post-processing effects such as glow, auto-exposure, fog, HDR tone-mapping, edge-smoothing, distortion, etc.

One major disadvantage of deferred shading is its incompatibility with hardware anti-aliasing [Fatahalian 2010], due to the separation between the geometry and lighting stages. Furthermore, deferred shading can be scaled to multiple materials at the cost of an extra G-buffer term for material ID. In addition, transparency is not supported. One possible workaround is to use depth peeling [Mammen 1989, Everitt 2001, Bavoil 2008] at the cost of extra G-buffers and computation.

III.3.6 Multi-resolution screen-space rendering

Lighting effects such as global illumination, in contrast to local illumination, are very difficult to handle using GPUs. For each surface point, indirect illumination must be gathered from all the other visible surface points. This is ill-suited for the GPU, because it is difficult to parallelize and make the memory access coherent. Sophisticated rendering algorithms are required to achieve interactive high quality global illumination on the GPU [McGuire 2009, Wang 2009b, Gautron 2008, Barsi 2005].

One way to overcome this architectural limitation of GPUs is to use multi-resolution screen-space rendering. These techniques take advantage of the fact that indirect illumination from diffuse and glossy surfaces are low-frequency and smooth in nature. See Figure III.4 for an example. Therefore, shading every single pixel can be wasteful, while it can be approximated by rendering indirect illumination at lower resolutions and up-sampling them.

A family of interactive rendering techniques heuristically shade pixels at various levels of coarseness and up-sample them to assemble a fully shaded 2D image within a deferred shading context. These algorithms splat shading in a screen-space hierarchical buffer such as a mip-map. After the splatting is done, they recurse through all the levels and interpolate values from coarser levels using bi-lateral up-sampling to preserve edges. Doing so saves computation power since only a subset of pixels will be shaded after all. In chapter IV, we draw inspiration from some recent multi-pass multi-resolution screen-space rendering algorithms [Nichols 2009, Shopf 2009, Nichols 2010b, Nichols 2010a, Soler 2010, Segovia 2006, Ritschel 2009] and introduce a single-pass multi-resolution rendering tech-

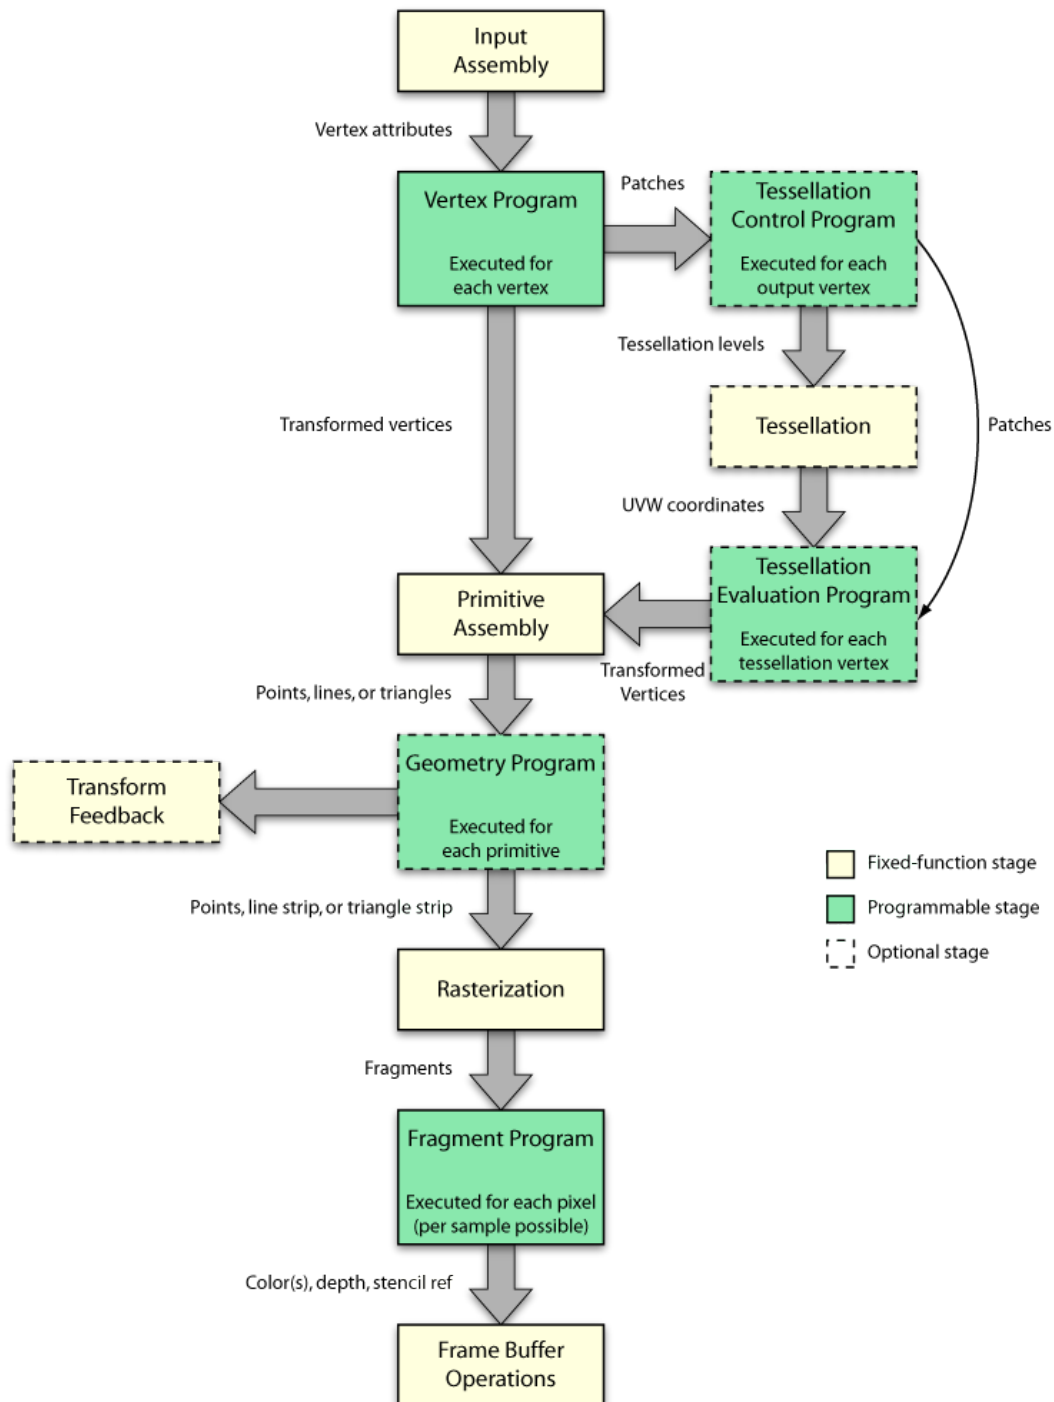


Figure III.3: The GPU rendering pipeline for OpenGL 4 and DirectX 11. (From *Wikipedia*)



Figure III.4: The low frequency nature of indirect illumination: Direct light only (left); indirect light (center); the combined image (right). (From [Nichols 2009])

nique for all-frequency direct illumination, that performs the shading and up-sampling in parallel at the same time.

Multi-resolution splatting for indirect illumination [Nichols 2009] splats illumination in image-space from a series of *virtual point lights (VPLs)* to a multi-resolution buffer. It extends the idea of *reflective shadow maps (RSM)* [Dachsbacher 2005] which is essentially a light-space G-buffer (camera depth, world position, normal, and reflected flux), using shadow mapping hardware to generate VPLs. RSM ignores the visibility of indirect illumination rays, assuming that the errors due to incorrect visibility is unnoticeable. Storing maximum and minimum depth and normal values for each pixel in a *min-max mip-map*, they start from 64 splats representing the whole image. They subdivide each splat into 4 new sub-splats if it contains discontinuities in depth or normal values. After the refinement is done, they up-sample the splatted indirect illumination from the coarsest level to the finest, making sure that empty pixels remain empty. Finally the indirect illumination is combined with direct illumination to build the final image. *Interactive Indirect Illumination Using Adaptive Multi-resolution Splatting* [Nichols 2010b] reduces the rendering costs further by grouping the sub-splats that potentially receive the same radiance.

Hierarchical Image-Space Radiosity for Interactive Global Illumination [Shopf 2009] replaces splatting with gathering. Instead of using a multi-resolution hierarchical buffer and iterative refinement, they flatten the hierarchy into a single resolution texture map with stencil buffering to gather indirect illumination. They also extended this technique to rendering global illumination for dynamic area lighting supporting visibility [Nichols 2010a].

III.3.7 Light field

Michael Faraday, in a lecture entitled "Thoughts on Ray Vibrations" in 1846, suggested that light should be considered as a field, much like the magnetic field. Years after, Alexander Gershun coined the phrase *Light field*, in a paper on radiometric properties of light in 3D space in 1936.

Light field describes the amount of light traveling in every direction through any point in space. This is a 5D plenoptic function: three dimensions to determine the position (x, y, z) and two dimensions to determine the direction (θ, ϕ) . If we consider a ray outside the convex hull of an object or in an empty region of space, the fact that radiance




Figure III.5: Light transport: light goes through several transformations before reaching our eyes. (From [Durand 2005])

along any ray remains constant reduces the 5D domain to 4D. In computer graphics, the radiance along rays in empty space is called the *4D light field* [Levoy 1996] or *Lumigraph* [Gortler 1996].

III.3.7.1 Light field propagation

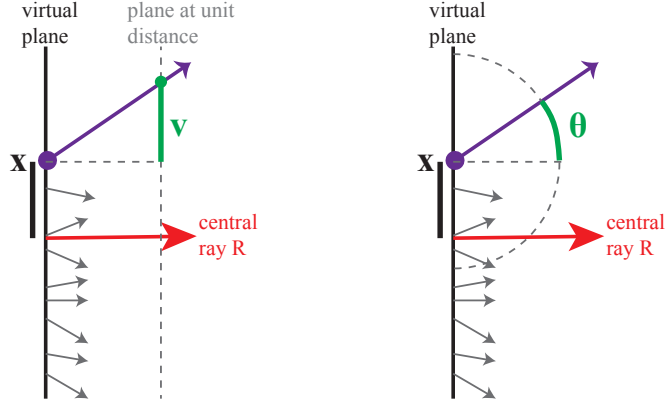
In ray optics, we assume photons are continuously traveling along light rays carrying energy. As light propagates through the scene, it interacts with the environment around it. Phenomena such as transport in free space, shading, and occlusion modify the local light field around a central light ray. Figure III.5 depicts various ways of light field propagation.

III.3.7.2 Light field parameterization

There are two popular parameterizations for the 4D light field at a virtual plane orthogonal to a central ray (Figure III.6). The *two-plane parameterization* [Chai 2000] uses the intersection v with a parallel plane at unit distance expressed in the local frame of x (Figure III.6-a). And, the *plane-sphere parameterization* [Camahort 1998] uses the angle θ with the central direction (Figure III.6-b). These two parameterizations are linked by $v = \tan\theta$ and are equivalent around the origin, as mentioned in [Durand 2005].

III.3.8 Frequency analysis of light transport

Durand et al. [Durand 2005] introduced a comprehensive study of the frequency content of a local light field, characterized by a 4D slice of radiance in the neighborhood of a central ray, and how it is altered by phenomena such as transport, shading, and occlusion. Such signal processing framework can be used to predict the frequency content of the local light



(a) two-plane parameterization (b) plane-sphere parameterization

Figure III.6: Two popular 4D light field parameterizations: (a) two-plane parameterization and (b) plane-sphere parameterization

field as it travels through the scene to reach the viewer. This information can then be used to control sampling rates for rendering.

We define the Fourier spectrum of the radiance field ℓ_R by $\widehat{\ell}_R$ in two-plane parameterization as:

$$\widehat{\ell}_R(\Omega_x, \Omega_v) = \int_{x=-\infty}^{\infty} \int_{v=-\infty}^{\infty} \ell_R(x, v) e^{-2i\pi\Omega_x x} e^{-2i\pi\Omega_v v} dx dv$$

In a simple scenario, light is emitted from the light source, transported in free space, optionally occluded by obstacles, transported again, reflected by one or more surfaces, and finally transported to the viewer. See Figure III.5.

Durand et al. [Durand 2005] studied the local light field propagation assuming a pinhole camera. Soler et al. [Soler 2009] extended the frequency analysis of light transport to study the local light field for a photographic lens with finite-sized aperture to support depth of field with adaptive sampling.

Now we briefly go through each operation:

III.3.8.1 Transport in free space

Transport in free space is a shear of the local light field in the Fourier domain.

$$\widehat{\ell}_{R'}(\Omega_x, \Omega_v) = \widehat{\ell}_R(\Omega_x, \Omega_v + d\Omega_x)$$

Please refer to [Durand 2005] for the derivations.

III.3.8.2 Reflection

Reflection is described by two scale transforms due to the incoming and outgoing angles and two shears due to the surface curvature k .




Figure III.7: Local shading for a curved receiver with arbitrary BRDF. (From [Durand 2005])

$$\widehat{\ell}_i(\Omega_x, \Omega_\theta) = e^{-i\Omega_\theta \theta_0} / |\cos \theta_0| \widehat{\ell}(-\Omega_x \cos \theta_0, \Omega_\theta)$$

$$\widehat{\ell}'_i(\Omega'_x, \Omega'_\theta) = \widehat{\ell}_i(\Omega'_x + k\Omega'_\theta, \Omega'_\theta)$$

$$\widehat{\ell}'_o(\Omega'_x, \Omega'_\theta) = \widehat{\ell}'_o(\Omega'_x - k\Omega'_\theta, \Omega'_\theta)$$

$$\widehat{\ell}''_o(\Omega_x, \Omega_\theta) = e^{i\Omega_\theta \theta_1} / |\cos \theta_1| \widehat{\ell}_o(\Omega_x \cos \theta_1, \Omega_\theta)$$

III.3.8.3 Shading

Shading is a convolution with a small kernel corresponding to the spectrum of the clamped cosine term followed by a clamping by the BRDF angular bandwidth.

$$\widehat{\ell}'_o(\Omega'_x, \Omega'_\theta) = (\widehat{\ell}'_i(\Omega_x, \Omega_\theta) \otimes \widehat{\text{cos}}_+(\Omega_\theta) \delta_{\Omega_x=0}) \widehat{\rho}'(\Omega'_\theta)$$

The important shading effects due to curvature, cosine term, and the BRDF are summarized in Figure III.7.

III.3.8.4 Texture mapping

Texture mapping is a multiplication of radiance in the primal domain and therefore a convolution in the Fourier space.

$$\widehat{\ell}_T(\Omega_x, \Omega_\theta) = \widehat{T}(\Omega_x) \delta_{\Omega_\theta=0} \otimes \widehat{\ell}_o(\Omega_x, \Omega_\theta)$$

III.3.8.5 Occlusion

Occlusion corresponds to a convolution by the spectrum of the blockers.

$$\widehat{\ell}_R(\Omega_x, \Omega_v) = \widehat{\ell}_R(\Omega_x, \Omega_v) \otimes \widehat{V}(\Omega_x, \Omega_v)$$

V is the visibility function ranging from 0 for blocked to 1 for visible.

III.3.8.6 Depth of field

At a *finite-sized aperture lens*, the spectrum of local light field is band-limited by the spectrum of the aperture response function.

$$\widehat{\ell}_{L_+}(\Omega_x, \Omega_v) = \widehat{\ell}_{L_-}(\Omega_x, \Omega_v) \widehat{a}(\Omega_x, \Omega_v)$$

L_+ (respectively L_-) represents the light field after (respectively before) the lens, and \widehat{a} is the Fourier transform of the indicator function of the set of rays not blocked by the aperture.

III.3.8.7 Motion blur

Egan et al. [Egan 2009] extended the frequency analysis of light transport even further to render motion blur using adaptive space-time sampling. They generalized light transport for moving objects, BRDF reflection, and moving shadows to the time domain.

III.4 Material appearance

Material appearance refers to the visual impression of a certain material. Digital appearance modeling is the data structures and algorithms we use to mathematically express and simulate a material from any view in any environment. [Dorsey 2008] is an excellent reference to digital modeling of material appearance.

To simulate the material appearance, we need to define how light interacts with matter from any incident direction and any point of view, under any illumination conditions. Therefore, we are going to define what we mean by reflectance and BRDF. Then we will look at the taxonomy of appearance and how by making assumptions we can reduce the number of dimensions of a general reflectance function. Next, we are going to introduce the basic concepts of reflectance acquisition. And finally, we will describe various ways of representing reflectance.

III.4.1 Definitions

III.4.1.1 Reflectance

According to Nicodemus et al. [Nicodemus 1992], *reflection* is the process by which electromagnetic flux (power), incident on a stationary surface or medium, leaves the surface or medium from the incident side without change in frequency; *reflectance* is the fraction of the incident flux that is reflected.




Figure III.8: BRDF describing how the radiance leaving the surface is related to the radiance incident on the surface.

In this thesis, we skip the concepts of *transmittance* and *absorption* for the sake of simplicity and only look into reflectance.

III.4.1.2 BRDF

Bidirectional reflectance distribution function (BRDF) indicates the ratio of the radiance reflected in a particular direction Θ_r as a result of the total incident flux per unit area from another direction Θ_i . This ratio is shown in Figure III.8 and expressed below:

$$f_r(\lambda, \mathbf{x}, \Theta_i \rightarrow \Theta_r) = \frac{dL_r(\lambda, \mathbf{x} \rightarrow \Theta_r)}{L_i(\lambda, \mathbf{x} \leftarrow \Theta_i) \cos \theta_i d\omega_i}$$

A BRDF must be always positive, energy conserving, and symmetric (reciprocal).

III.4.1.2.1 BRDF and energy conservation

The energy conservation means that the amount of energy reflected to all directions from light, incident from a single direction, must be between 0 and 1.

$$\int_{\omega_r} f_r(\lambda, \mathbf{x}, \Theta_i \rightarrow \Theta_r) \cos \theta_r d\omega_r \leq 1$$

III.4.1.2.2 BRDF and Helmholtz reciprocity

The reciprocity constraint basically means that light paths are reversible. In other terms, the BRDF must remain the same when the direction of incidence and reflectance are reversed.

$$f_r(\lambda, \mathbf{x}, \Theta_i \rightarrow \Theta_r) = f_r(\lambda, \mathbf{x}, \Theta_r \rightarrow \Theta_i)$$

III.4.2 Taxonomy of appearance

Figure III.9 depicts the whole taxonomy of appearance. In the following, we explain each function:

- General scattering function (12D)
Starting from the general scattering function, it's a 12 dimensional function depending on incoming and reflected positions on the surface, angles, wavelengths and times.
- BSSRDF (8D)
Assuming we do not have phosphorescence and fluorescence, we can reduce the dimensions to 8D and get *Bidirectional Scattering Surface Reflectance Distribution Function (BSSRDF)*.
- SVBRDF (6D)
Ignoring sub-surface scattering, the scattering function reduces to 6 dimensions. *Spatially-varying Bidirectional Reflectance Distribution Function (SVBRDF)* is a BRDF that varies with the position. It can also be seen as a texture map that varies with illumination and viewing direction, called *Bidirectional Texture Function (BTF)*.
- BSSDF (6D)
On the other hand, if we assume that the spatial variation only depends on $\mathbf{x}_r - \mathbf{x}_i$ and $\mathbf{y}_r - \mathbf{y}_i$, it results in *Bidirectional Sub-surface Scattering Distribution Function (BSSDF)*.
- BRDF (4D)
Ignoring variation across the surface results in a 4D function called *Bidirectional Reflectance Distribution Function (BRDF)*.
- Isotropic BRDF (3D)
Ignoring anisotropy reduces BRDF from 4D to 3D for *isotropic BRDFs*.
- Light Field (4D)
Light Fields and their variants such as *Lumigraphs* and *Surface Light Fields (SLF)* are 4D functions resulted from ignoring dependence on incident light direction.
- Diffuse texture (2D)
Finally, *diffuse textures* are 2D maps not dependent on light direction at all.

III.4.3 Reflectance acquisition

Reflectance of real world materials can be acquired using special measurement devices. To choose the right acquisition technique, the precision, cost, acquisition time, and range of materials that can be acquired must be considered.

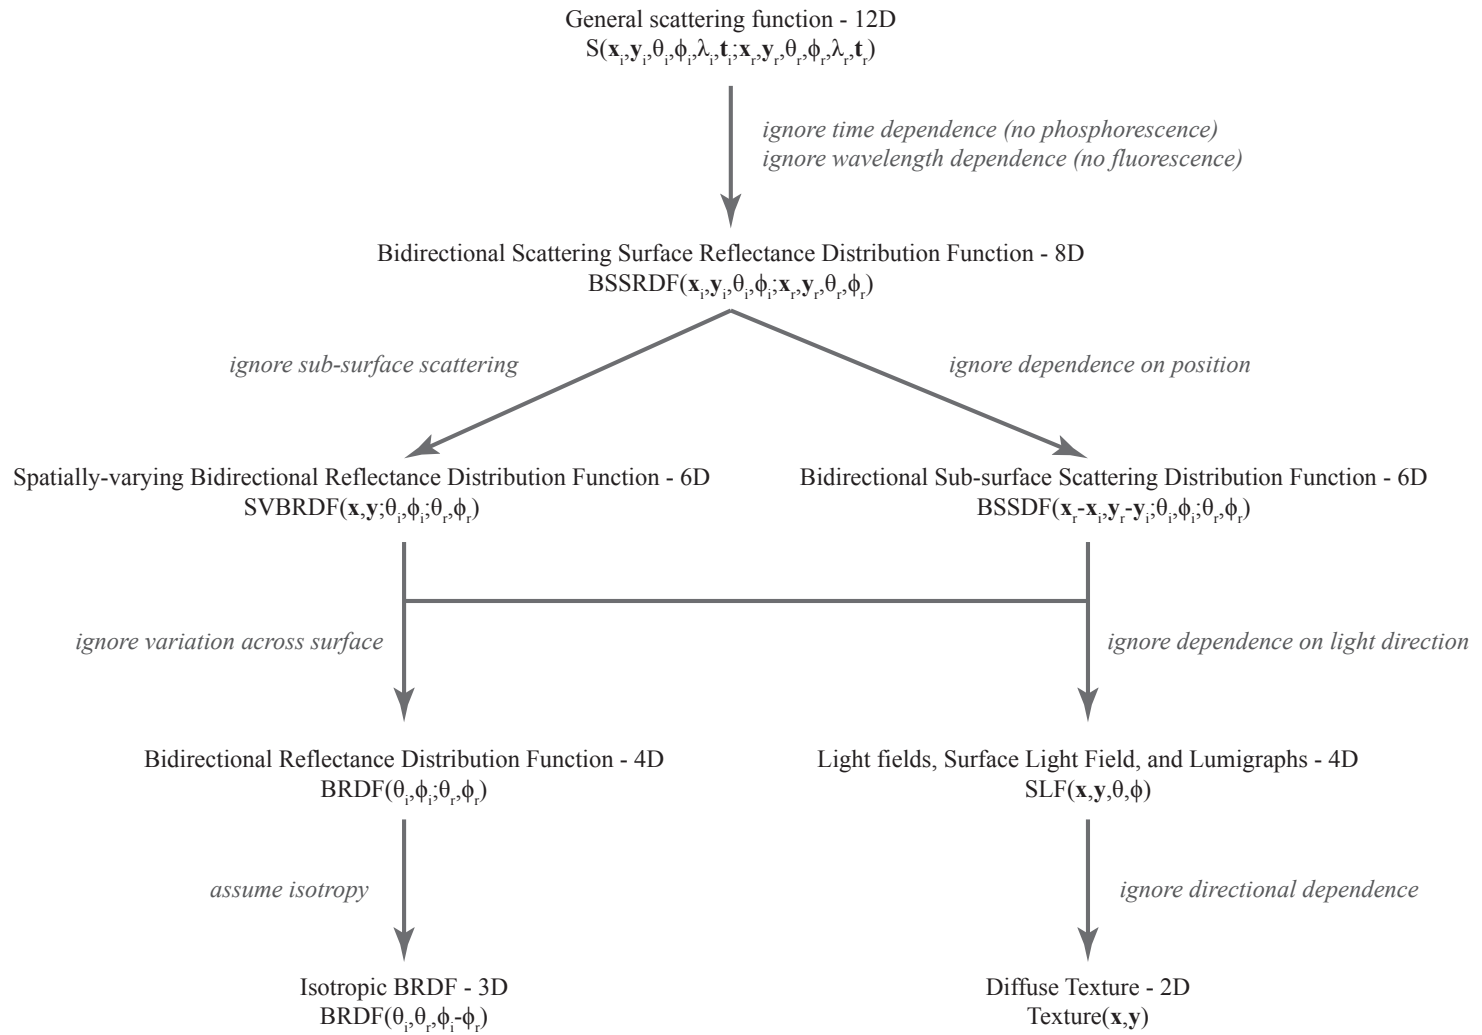


Figure III.9: Taxonomy of appearance. (From Szymon Rusinkiewicz)




Figure III.10: Gonio-reflectometer for isotropic BRDF measurement. (From [Li 2005])

III.4.3.1 Acquisition techniques

The classic device for measuring BRDFs is a *gonio-reflectometer* [Murray-Coleman 1990, nis 2010, cor 2005]. It involves a light source, a detector, and a material sample. Each component can be positioned by stepper motors or robotic arms. If a flat material sample is used, the detector can take only one sample at a time. Therefore, the measurement process is very lengthy, taking days to measure a single material. However, it can be carefully calibrated and multi-spectral high dynamic range detectors can be used.

Li et al. [Li 2005] setup a gonio-reflectometer with only three degrees of freedom for measuring isotropic BRDFs. This technique uses two degrees of freedom in the orientation of a planar sample, and one degree of freedom in the angular position of the light source. See Figure III.10.

Ward [Ward 1992a] accelerated this process by capturing the entire hemisphere of reflected directions simultaneously. He used a half-silvered hemispherical mirror, a CCD detector with a fish-eye lens, and a flat sample in the center. Since only the light direction should vary, it is a much faster acquisition process. Due to the size and shape of the reflecting hemisphere and the size of the sample, the angular resolution near the grazing angles is limited.

Dana et al. [Dana 2004] proposed a gonio-reflectometer for capturing spatially-varying materials. It uses a concave parabolic mirror, a CCD camera, a beam splitter, and translational stages.

Marschner [Marschner 1999, Marschner 2000] and similarly Matusik et al. [Matusik 2003] used a spherical material sample in a way that all surface normal variations would be visible in a single image. See Figure III.11. Multi exposure photography was used to capture the high dynamic range. This setup is limited to homogeneous isotropic materials only.

Ngan et al. [Ngan 2005] setup a gonio-reflectometer for measuring anisotropic BRDF, using a set of cylindrical stripes of an anisotropic material. Figure III.12 depicts this acquisition setup. The cylinder tilt, surface normal variation, light position and strip orientation each contributes to one degree of freedom for acquiring a 4D BRDF. However, the resolu-




Figure III.11: Image-based isotropic BRDF measurement. The surface normals of the spherical sample provide two degrees of freedom, the rotating light provides the third for the 3D isotropic BRDF. (From [Ngan 2006])

tion of the acquired BRDF is mainly limited by the number of material stripes mounted on the cylinder.

More recently, a large body of work has been done for BRDF acquisition either using a more sophisticated setup or inferring reflectance from a single image. The contributions in this thesis are not related to material acquisition, and therefore the reader is referred to the following references for further reading. [Ren 2011, Romeiro 2008, Romeiro 2010, Debevec 2000, Dong 2010, Ghosh 2010a, Ghosh 2010b, Ben-Ezra 2008, Munoz 2011, Hullin 2010]

III.4.3.2 Isotropic BRDF databases

In this thesis we have used the MERL-MIT homogeneous opaque isotropic BRDF database from [Matusik 2003]. This database contains 100 densely measured isotropic materials from diffuse to glossy and specular. We chose this database because the measurements are dense and no interpolation is needed, and also it contains a wide range of materials. The resolution is 1,458,000 combinations of viewing and illumination directions. See Figure III.13.

Cornell University [cor 2001] has a few multi-spectral measured paints (automotive, spray, and house paints). Some materials are acquired using a gonio-reflectometer, and some other using an image-based technique. The resolution is low (1439 points) and samples should be interpolated for use in rendering. And, there is not much variety. See Figure III.14.

CUReT database [cur 1999] contains BRDF measurements for about 60 different real-world surfaces. The resolution is as low as 200 combinations of viewing and illumination directions.

One major problem with acquired materials in general is that the signal-to-noise ratio is usually very low at grazing angles.




Figure III.12: Gonio-reflectometer for anisotropic BRDF measurement. The cylinder tilt, surface normal variation, light position, and strip orientation each contribute to one degree of freedom for acquiring a 4D BRDF. (From [Ngan 2005])

III.4.4 Reflectance representation

Reflectance representation can be derived from actual physical measurements (data driven reflectance), or can be artistically created, or a mix of both. The choice of what representation technique to employ mainly depends on the application and the costs of using it. Artists like to have control over the overall look of a material while it may not necessarily end up being physically correct, e.g. energy conserving. Editing reflectance requires meaningful parameters to tweak, which can be derived from fitting the parameters of an analytical expression to the measurements.

III.4.4.1 Tabulated BRDF

One possible way of representing material appearance is to use the reflectance measurements directly. This can be done by tabulating the measured values in the memory for on-the-fly evaluation at render time. However, this has several disadvantages. Large memory consumption and access make this approach almost impractical for densely sampled acquisitions. For example, each isotropic BRDF in the MERL-MIT database is 33 MB. An anisotropic material, such as the ones measured by [Ngan 2005], takes about 750 MB of memory. The lack of parameters for tweaking makes it extremely difficult to edit the reflectance of a material directly from the measurements. And finally, there is no practical expression for importance sampling of acquired materials. For BRDF importance sampling, one should pre-compute the importance samples by inverting the cumulative distribution function (CDF) of the BRDF [Lawrence 2005] and tabulate it, which is not very practical.

In chapter IV, we show how an efficient sampling prediction framework can be used to achieve interactive frame-rates when rendering using tabulated BRDF measurements.




Figure III.13: The MERL-MIT isotropic BRDF database. (From [Matusik 2003])




Figure III.14: The Cornell BRDF database. (From [cor 2001])

III.4.4.2 BRDF re-parameterization

One simple way to more efficiently represent isotropic BRDFs is the change of variables proposed by Rusinkiewicz [Rusinkiewicz 1998]. He re-parameterizes the BRDF in terms of the half-way vector (the vector half-way between incoming and outgoing vectors) and a difference vector (incoming vector in a frame of reference in which the half-way vector is at the north pole), instead of the incoming and outgoing angles. Figure III.15 depicts the half/difference-angle parameterization of BRDFs. This re-parameterization aligns the features of the BRDF such as the specular peak with the new coordinate axes. Change of basis using this re-parameterization guaranties less non-zero coefficients, due to the weak dependency on combinations of both axes. Please note that the MERL-MIT BRDF database is already parametrized by the half-angle parametrization.




Figure III.15: The half/difference-angle parameterization: The BRDF as function of (θ_i, ϕ_i) and (θ_o, ϕ_o) is shown on the left, and the re-parameterization as a function of the half-angle (θ_h, ϕ_h) and a difference-angle (θ_d, ϕ_d) is shown on the right. (From [Rusinkiewicz 1998])

III.4.4.3 Dimensionality reduction and compression techniques

BRDF measurement data, as a multi-dimensional quantity, can be reduced and compressed for better rendering performance. Mathematical techniques [Fodor 2002] such as principal component analysis (PCA) and factor analysis can be used to compress BRDFs.

III.4.4.3.1 Decomposition into basis functions

A solution is to take advantage of the sparsity of the signal after a change of basis. Decomposition into basis functions such as tensor products of the spherical harmonics [Westin 1992], Zernike polynomials [Koenderink 1996], spherical wavelets [Schröder 1995, Lalonde 1997, Claustres 2007], or spherical Gaussians [Wang 2009a] are used for BRDF compression.

III.4.4.3.2 Factorization and separable approximation

Factorization techniques and separable approximations [Kautz 1999, McCool 2001, Latta 2002] factor the 4D BRDF into two or more 2D sets of projections for conventional storage such as hardware 2D texture. The idea is to decompose the high dimensional BRDF into products of two or more factors of lower dimensionality, using singular value decomposition (SVD) for example.

$$f(\hat{\omega}_o, \hat{\omega}_i) = \sum_{n=1}^N u_n(\pi_u(\hat{\omega}_o, \hat{\omega}_i)) v_n(\pi_v(\hat{\omega}_o, \hat{\omega}_i))$$

where u_n and v_n are two-dimensional functions (suitable for 2D texture storage) and π_u and π_v are projection functions associated with each map.

Figure III.16 depicts an example of factorizing a 2D input signal into two and three 1D factors.




Figure III.16: A 2D factorization example: the 2D input signal (left); separable approximation with two 1D factors (center); separable approximation with three 1D factors (right). (From [McCool 2001])




Figure III.17: The local coordinate system defined by the normal (\hat{n}), tangent (\hat{t}), and bi-normal (\hat{b}) vectors. (From [Wikipedia](#))

III.4.4.4 Analytical models

Empirical models qualitatively emulate material appearance while ignoring physical properties of materials, while physically-based models apply first principles of physics to a surface's microscopic structure to predict how a material interacts with light.

One major advantage of analytical models is that the parameters can be used for fitting to the measurements. Evaluation at render time is fast and efficient, because instead of tabulating the whole BRDF we only need to store a few coefficients related to the parameters in the model. For some models analytical importance sampling expressions exist which makes rendering fast and efficient. But on the other hand, not all analytical models are physically plausible. And, the fitting is subject to numerical inaccuracies, specially for fitting to highly specular materials, which degrades the appearance quality of a rendered material.

Now, we are going to look at some of the existing parametric models for reflectance. Please note that BRDF is always defined in local coordinate system, defined by the normal, tangent and bi-normal. See [Figure III.17](#).




Figure III.18: The geometry of reflection: L is the light direction, N is the surface normal at point P , V is the view direction, R is the ideal specular reflection direction, and H is the half-way vector between L and V .

III.4.4.4.1 Empirical models

Phong BRDF We start with the *Phong BRDF* [Phong 1975].

$$\rho_{\text{phong}}(L, V) = \rho_d + \rho_s \frac{(V \cdot R)^n}{N \cdot L}$$

ρ_d and ρ_s are the diffuse and specular coefficients, which can take on different values for different RGB channels, and n is the shininess.

See Figure III.18 for the vectors L , V , N , and R .

The diffuse and specular components of a BRDF model describe two different phenomena. The diffuse term is the result of reflecting light almost equally in all directions, due to absorption and sub-surface scattering. The specular term describes the rather directional component of reflection. It describes the direction and fall-off of the specular peak.

The Phong BRDF model is not physically correct since it does not conserve energy. More specifically, the specular term of Phong does not conserve energy at sufficiently large incident angles [Lewis 1994].

Lewis [Lewis 1994] added a correction factor to make Phong loosely energy conserving.

$$\rho_{\text{lewis-phong}}(L, V) = \frac{\rho_d}{\pi} + \rho_s \cdot \frac{n+2}{2\pi} (V \cdot R)^n$$

Blinn [Blinn 1977] re-parameterized Phong based on the half-way vector (H) instead of the mirror specular direction (R).

$$\rho_{\text{blinn-phong}}(L, V) = \frac{\rho_d}{\pi} + \rho_s \cdot \frac{n+2}{2\pi} (H \cdot N)^n$$

Ward BRDF Ward [Ward 1992a] replaced the cosine term to a power in the Phong model with an exponential. He extended the model to anisotropic specular reflection using an elliptical Gaussian.




Figure III.19: The definition of vectors and angles for the albedo corrected Ward-Dür model. (From [Geisler-Moroder 2010])

$$\rho_{ward}(L, V) = \frac{\rho_d}{\pi} + \rho_s \cdot \frac{1}{\sqrt{(N \cdot L)(N \cdot V)}} \cdot \frac{\exp[-\tan^2 \theta_h (\cos^2 \phi_h / \alpha^2 + \sin^2 \phi_h / \beta^2)]}{4\pi\alpha\beta}$$

This model is reciprocal but not strongly energy conserving.

Dür [Dür 2006] corrected the normalization factor for the Ward model.

$$\rho_{ward-dur}(L, V) = \frac{\rho_d}{\pi} + \rho_s \cdot \frac{1}{(N \cdot L)(N \cdot V)} \cdot \frac{\exp[-\tan^2 \theta_h (\cos^2 \phi_h / \alpha^2 + \sin^2 \phi_h / \beta^2)]}{4\pi\alpha\beta}$$

Still, both Ward and Ward-Dür models do not meet energy balance at grazing angles [Neumann 1999]. Geisler-Moroder and Dür [Geisler-Moroder 2010], recently, introduced a new ward model with bounded albedo to ensure energy conservation.

$$\rho_{ward-dur-gm}(L, V) = \frac{1}{\pi\alpha\beta} \cdot \left(-\tan^2 \delta \left(\frac{\cos^2 \phi}{\alpha^2} + \frac{\sin^2 \phi}{\beta^2} \right) \right) \cdot \frac{2(1 + \cos \theta_l \cos \theta_v + \sin \theta_l \sin \theta_v \cos(\phi_v - \phi_l))}{(\cos \theta_l + \cos \theta_v)^4}$$

See also Figure III.19 for the definition of vectors and angles associated with the equation.

Lafortune BRDF The Lafortune BRDF [Lafortune 1997] is a cosine lobe model as a generalization of the Phong model. It allows the definition of multiple lobes (possibly anisotropic) around arbitrary axes with respect to the surface.

$$\rho_{lafortune}(u, v) = \rho_s \left[C_x u_x v_x + C_y u_y v_y + C_z u_z v_z \right]^n$$

where u and v are incoming and outgoing vectors in the local coordinate system at the surface, where the z axis is the normal, as illustrated in Figure III.20. n controls the sharpness of the lobe. Multiple lobes of this form can be summed up to form a BRDF.




Figure III.20: Vectors for the Lafortune model: u is the incoming vector and v is the outgoing vector in the local coordinate system, where z is the normal at the surface. (From [Lafortune 1997])

Ashikhmin-Shirley BRDF The Ashikhmin-Shirley BRDF [Ashikhmin 2000] is another extension to the Phong model. This model is anisotropic and physically plausible and accounts for Fresnel effects. The formulas are not included here due to their complexity, and the reader is referred to their paper.

III.4.4.4.2 Physically-based models

Physically-based models infer the visual appearance of a material from its microscopic structure, based on physics first principles.

The micro-facet theory Torrance and Sparrow [Torrance 1967] introduced a new geometric description of the surface to model light reflection from rough surfaces, that resulted in the idea of *micro-facet BRDF* models. The Torrance-Sparrow model assumes that a rough surface is made of many tiny V-shaped grooves composed of smooth mirror-like facets called the *micro-facets*. This model predicts reflection as a sum of specular reflection from these facets plus a diffuse reflection due to multiple reflections and/or internal scattering. Because the micro-facets are perfectly specular, only those with a normal equal to the half-angle vector cause specular reflection.

The Cook-Torrance reflectance model Cook and Torrance [Cook 1981] introduced the *Cook-Torrance BRDF* model based on the micro-facet theory. The full BRDF is:

$$\rho(\mathbf{i}, \mathbf{o}) = \frac{\rho_d}{\pi} + \frac{\rho_s}{\pi} \frac{F(\mathbf{i} \cdot \mathbf{h})D(\theta_h)G(\mathbf{i}, \mathbf{o})}{\cos(\theta_i) \cos(\theta_o)} \quad (\text{III.1})$$

where \mathbf{i} is the incoming direction, \mathbf{o} is the outgoing direction, \mathbf{h} is the half-vector ($\mathbf{h} = \mathbf{i} + \mathbf{o}$, normalized). θ_i , θ_o and θ_h are the respective angles between these vectors and the surface normal \mathbf{n} . ρ_d and ρ_s are the diffuse and specular coefficients, respectively (see notations on Figure III.21).

D is the probability distribution function for micro-facet normals. G is the shadowing and masking coefficient, expressing the amount of light that is blocked by other micro-facets before reaching the point (shadowing), or after reflection (masking). For isotropic




Figure III.21: For Cook-Torrance BRDF, the material micro-geometry is a random distribution of specular micro-facets. Angles are relative to the normal \mathbf{n} of the macro-geometry.

materials, G is approximated by a product of one-dimensional functions:

$$G(\mathbf{i}, \mathbf{o}) \approx G_1(\mathbf{i})G_1(\mathbf{o})$$

and G_1 is computed from D through a double integration [Smith 1967, Brown 1980, Bourlier 2002, Walter 2007].

F is the Fresnel term, depending on the refraction index η of the material. It is usually close to constant at normal incidence, and increases to unity at grazing angles.

The normal probability distribution function (D) The micro-facets are randomly oriented and their distribution can be statically modeled as a *normal probability density function*. It is a statistical distribution of slopes to the mirror-like facets on the surface.

Given a direction $\omega = (\theta, \phi)$ and an infinitesimal solid angle $d\omega$ centered around this direction, the probability for the normal of a micro-facet to be inside the cone $(\omega, d\omega)$ is $D(\omega) (\omega \cdot \mathbf{n}) d\omega$.

In order to be physically plausible, D should follow several requirements: it should be positive, the projected area of the micro-facets in a direction should be equal to the projected area of the macro-surface and in particular, $\int D(\omega) (\omega \cdot \mathbf{n}) d\omega = 1$. D is usually expressed as:

$$D(\theta) = \frac{\chi_{[0, \frac{\pi}{2}]}(\theta)}{\cos^4 \theta} P_{22}(\tan^2 \theta) \quad (\text{III.2})$$

where $P_{22}(x)$ is a positive function of the variable $x \in [0, \infty)$ and $\chi_{[0, \frac{\pi}{2}]}(\theta)$ ensures sidedness: it is equal to 1 if $\theta < \frac{\pi}{2}$, and 0 otherwise. With this expression,

$$D(\theta) \cos \theta d\omega = \frac{1}{2} P_{22}(\tan^2 \theta) d(\tan^2 \theta) \quad (\text{III.3})$$

Since $\int D(\theta) (\omega \cdot \mathbf{n}) d\omega = 1$, we have $\int_0^\infty P_{22}(x) dx = 1/\pi$.

Importance sampling based on the distribution We use Equation III.3 for importance sampling [Walter 2007]: for a given incoming direction \mathbf{i} , we first pick a random micro-facet normal \mathbf{m} with probability $D(\theta_m) \cos \theta_m$, then compute the reflected direction \mathbf{o} so that \mathbf{m} is the half-vector. Finally, we compute the sampling weight as the BRDF divided

	Beckmann	Exponential	TR/GGX
$P_{22}(x)$	$\frac{1}{\pi\alpha^2} e^{-x/\alpha^2}$	$\frac{1}{2\pi\alpha^2} e^{-\sqrt{x/\alpha^2}}$	$\frac{\alpha^2}{\pi(\alpha^2+x)^2}$

Table III.2: P_{22} functions corresponding to classical micro-facet distributions, with $x = \tan^2 \theta$.

by the probability, multiplied by the Jacobian of \boldsymbol{o} :

$$w_{\boldsymbol{o}} = \frac{\rho(\boldsymbol{i}, \boldsymbol{o}) \cos \theta_o}{D(\theta_m) \cos \theta_m} \left\| \frac{\partial \omega_o}{\partial \omega_m} \right\|$$

To pick \boldsymbol{m} with the probability $D(\theta_m) \cos \theta_m$, we need the associated Cumulative Distribution Function F :

$$F(\theta_m, \phi_m) = \int_0^{\phi_m} \int_0^{\theta_m} D(\theta) \cos \theta \, d\omega = \frac{\phi_m}{2} \int_0^{\tan^2 \theta_m} P_{22}(u) \, du$$

F is separable: $F = F_\phi F_\theta$, with F_ϕ and F_θ from 0 to 1:

$$\begin{aligned} F_\phi(\phi_m) &= \frac{\phi_m}{2\pi} \\ F_\theta(\theta_m) &= \pi \int_0^{\tan^2 \theta_m} P_{22}(u) \, du \end{aligned}$$

Assuming we have two uniform random variables u_1 and u_2 in $[0, 1)$, we get the sampling equations:

$$\begin{aligned} \phi_m &= F_\phi^{-1}(u_1) = 2\pi u_1 \\ \theta_m &= F_\theta^{-1}(u_2) \end{aligned}$$

For any micro-facet distribution D following Equation III.2, importance sampling reduces to computing the integral of P_{22} and inverting it.

Existing normal probability distributions Table III.2 lists the P_{22} functions corresponding to existing micro-facet distributions. Beckmann is frequently used [Cook 1982]; it is a Gaussian in $\tan \theta$, and is also very close to the Phong distribution [Walter 2007]. The exponential distribution is mostly used for the propagation of radio waves [Bourlier 2002, Brown 1980], and has a sharper peak and larger tails than Beckmann. The Trowbridge-Reitz [Trowbridge 1975] and GGX [Walter 2007] are independently formed but identical distributions. They have a sharper peak and larger tails than Beckmann and Exponential distributions. Please note that α is the root-mean-square slope of the micro-facets, and it captures surface roughness.

Recently, Löw et al. [Löw 2012] proposed to use a condensed version of the ABC model as the distribution in the Cook-Torrance BRDF. Please refer to their article for full details.

The shadowing and masking attenuation factor (G) Some micro-facets may block the light before arriving at a particular point causing *shadowing*, or light may get blocked by other micro-facets after it's reflected off the surface causing *masking*. These geometrical effects (shadowing and masking) are grouped under a unified *geometric term* (G).

In the original Cook-Torrance BRDF model [Cook 1981] the shadowing function G was independent from the micro-facet distribution.

$$G = \min \left\{ 1, \frac{2(N \cdot H)(N \cdot V)}{(V \cdot H)}, \frac{2(N \cdot H)(N \cdot L)}{(V \cdot H)} \right\}$$

Further research [Smith 1967, Brown 1980, Bourlier 2002] proved that we get a more accurate shadowing term by integration from the micro-facet distribution D . G is approximated as the product of two one-dimensional functions: $G(\mathbf{i}, \boldsymbol{\theta}) \approx G_1(\theta_i)G_1(\theta_o)$. G_1 is the Smith shadowing function [Smith 1967] and is computed from D through a double integration:

$$\begin{aligned} P_2(r) &= \int_{-\infty}^{\infty} P_{22}(r^2 + q^2) dq & \text{(III.4)} \\ \Lambda(\theta) &= \int_{\frac{1}{\tan \theta}}^{\infty} (r \tan \theta_v - 1) P_2(r) dr \\ G_1(\theta) &= \frac{\chi_{[0, \frac{\pi}{2}]}(\theta)}{1 + \Lambda(\theta)} \end{aligned}$$

For further details, please refer to the original papers or the short summary in [Walter 2007].

The Fresnel term (F) The *Fresnel term* (F) defines the ratio of light being reflected by the well-oriented micro-facets. It depends on the wavelength of the incoming light, the geometry of the surface and the light, and the incoming angle. Schlick's approximation to the Fresnel term [Schlick 1994] is widely used in computer graphics.

$$F_\lambda(\theta) = F_\lambda(0) + (1 - F_\lambda(0))(1 - \cos \theta)^5$$

Lazányi et al. [Lazányi 2005] introduced a more accurate Fresnel approximation for metals.

$$F_\lambda(n, k, \cos \theta) = \frac{(n - 1)^2 + 4n(1 - \cos \theta)^5 + k^2}{(n + 1)^2 + k^2}$$

Where for $k = 0$ we get back to Schlick's Fresnel approximation.

The Oren-Nayar reflectance model The Oren-Nayar model [Oren 1994] is a generalization of Lambert's perfect diffuse reflection model. They take the micro-facet approach, but contrary to Cook-Torrance model, they assume the surface is made of Lambertian micro-facets. Because they account for shadowing and masking, although the facets

are Lambertian, the overall BRDF is not Lambertian anymore. The Oren-Nayar model simulates back-scattering but not specular reflection.

$$\rho_{oren-nayar}(L, V) = \frac{\rho}{\pi} (A + B \max[0, \cos(\phi_v - \phi_l)] \sin(\max[\theta_v, \theta_l]) \tan(\min[\theta_l, \theta_v]))$$

Where

$$A = 1.0 - 0.5 \frac{\alpha^2}{\alpha^2 + 0.33}$$

$$B = 0.45 \frac{\alpha^2}{\alpha^2 + 0.09}$$

This is an approximation to the full solution for Lambertian facets with reflectance ρ and Gaussian slope distribution with the standard deviation of α .

The He-Torrance reflectance model The He-Torrance model [He 1991] is a very comprehensive analytical model accounting for shadowing and masking, anisotropy, as well as polarization effects. This model is able to predict the appearance of a nearly smooth surface, producing variations in sharp and blurred reflections not predicted by other models. Due to the high complexity of this analytical model, we do not explain this model here. Please refer to the original paper for details.

III.4.4.5 Comparison between existing analytical models

Many different analytical models for reflectance have been proposed in the past decades. Hence, there is a need for quantitative and qualitative comparison between all these models compared to a ground truth such as physical reflectance measurements.

Ngan et al. [Ngan 2005] fitted the previously mentioned analytical BRDF models to most of the materials in the MERL-MIT database. The isotropic models compared in their work include: Ward [Ward 1992a], Ward-Dür [Dür 2006], Blinn-Phong [Blinn 1977], Lafortune [Lafortune 1997], Cook-Torrance [Cook 1981], He-Torrance (also known as He) [He 1991], and Ashikhmin-Shirely [Ashikhmin 2000]. They fitted the materials with a single lobe. Adding more lobes to the BRDF model makes it more accurate, but it increases the instability of the fitting process. 3 lobes are usually considered to be the limit for automatic fitting.

Figures III.22, III.23, III.24 show examples of the comparison between these models for *beige-fabric* as a diffuse material, *fruitwood-241* as a glossy material, and *chrome* as a specular material respectively.

Ngan et al. [Ngan 2005] reported that physically-based models like Cook-Torrance and He-Torrance, as well as the Ashikhmin-Shirley model often produce the best results. They further explained that the He model does not produce noticeably superior visual results to the Cook-Torrance model. The Lafortune model is not a good candidate for glossy materials due to the anisotropic blurriness near grazing angles. The Blinn-Phong, Ward, and Ward-Dür models are only suitable to a sub-class of materials, since they have near constant reflection power independent of the incident angle.




Figure III.22: Comparison between several analytical models for *beige-fabric* from MERL-MIT database — figures from left to right: reference, Ward, Ward-Dür, Blinn-Phong, Lafortune, Cook-Torrance, He, and Ashikhmin-Shirely. The numbers below each figure represent the L^2 metric error value. (From [Ngan 2005] supplemental material)

III.5 Summary and conclusion

In this chapter, we briefly explored the simulation of light transport within a synthetic environment. We defined the radiometric terms involved in illumination simulation, the rendering equation and its estimation using Monte Carlo sampling. We looked at hardware-accelerated rendering, its pipeline, and the difference between forward and deferred shading. We briefly reviewed some multi-resolution screen-space rendering techniques that implement interactive global illumination on graphics hardware in a more efficient manner. We defined light fields, its propagation and parameterization. We also introduced the concept of analyzing light transport in frequency space using Fourier theory.

Then, we talked about material appearance, what we mean by reflectance and BRDF. We explored the taxonomy of appearance, and discussed various techniques for reflectance acquisition using special optical measurement devices. We looked at various ways to represent reflectance using the measurements directly, re-parameterization, dimensionality reduction and compression techniques, and also analytical models. finally, we compared existing analytical models.

In the next chapter, we will introduce a frequency analysis framework for interactive rendering of measured reflectance.




Figure III.23: Comparison between several analytical models for *fruitwood-241* from MERL-MIT database — figures from left to right: reference, Ward, Ward-Dür, Blinn-Phong, Lafortune, Cook-Torrance, He, and Ashikhmin-Shirely. The numbers below each figure represent the L^2 metric error value. (From [Ngan 2005] supplemental material)




Figure III.24: Comparison between several analytical models for *chrome* from MERL-MIT database — figures from left to right: reference, Ward, Ward-Dür, Blinn-Phong, Lafortune, Cook-Torrance, He, and Ashikhmin-Shirely. The numbers below each figure represent the L^2 metric error value. (From [Ngan 2005] supplemental material)

Interactive rendering of measured reflectance using frequency analysis ¹

Contents

IV.1 Introduction (français)	41
IV.2 Introduction	43
IV.2.1 Overview	44
IV.3 Real-time frequency analysis	45
IV.3.1 Illumination, textures, and BRDF bandwidth	45
IV.3.2 Computing one-bounce 2D bandwidth	49
IV.3.3 Image-space bandwidth and sampling rate	51
IV.3.4 Adaptive sampling for shading	51
IV.3.5 Adaptive sampling for pre-convolved shading	52
IV.3.6 Implementation roadmap	53
IV.4 Hierarchical shading algorithm	53
IV.4.1 Bandwidth buffer initialization	54
IV.4.2 Shading and up-sampling	55
IV.4.3 Shading computation	55
IV.5 Results and discussion	56
IV.5.1 Behavior of our algorithm	56
IV.5.2 Comparison with related work	59
IV.5.3 Rendering results	59
IV.5.4 Discussion	61
IV.6 Conclusion	62

IV.1 Introduction (français)

Comme nous l'avons vu dans le chapitre précédent, la synthèse d'images photo-réaliste est très coûteuse en raison de la nécessité d'une intégration numérique par pixel sur plusieurs directions incidentes. L'utilisation de BRDF mesurée assure un maximum de réalisme à la

¹This chapter is based on the publication [M. Bagher 2012b] and is extended with new features and more discussions.

précision des mesures, mais rend le calcul plus coûteux du fait d'un accès mémoire élevé et des coûts de stockage prohibitifs.

Pour certaines applications, telles que la conception de formes, un feedback interactif d'apparence matérielle est souhaitable. La nécessité de permettre de changer dynamiquement la géométrie pose le premier défi a calcul interactif. En outre, il est important que le résultat soit réaliste et compatible avec son aspect final après post-conception hors ligne, par exemple a l'aide d'un modèle physique de rendu.

Une gamme d'approches parvient partiellement à résoudre ce problème. À une extrémité, des algorithmes rapides se concentrent sur la géométrie modifiable avec les modèles de matériaux simples. D'autres algorithmes s'efforcent de décrire une variété d'effets tels que l'illumination globale. Cette dernière approche nécessite généralement le pré-calcul du transfert de rayonnement (comprenant la visibilité) et empêche l'édition de la géométrie [Sloan 2002]. Le rendu réaliste interactif d'une géométrie dynamique avec des matériaux complexes tels que des mesurées BRDFs acquis est encore un problème de recherche ouvert et est l'objet de ce chapitre. Nous réalisons cette combinaison au prix de l'illumination globale et la visibilité.

Simuler l'apparence de matériaux sous éclairage haut fréquence nécessite l'estimation d'une intégrale de l'éclairage incident à chaque pixel modulée par la fonction de réflectance du matériau. Les intégrands sont généralement échantillonnées, et le taux d'échantillonnage dépend de la matière : les matériaux diffus nécessitent de nombreux échantillons sur les directions incidentes, mais présentent une faible variation entre les pixels voisins ; les matériaux spéculaires nécessitent moins d'échantillons sur les directions incidentes mais provoquent de grandes variations entre les pixels voisins.

Nous nous appuyons sur les théories de l'analyse fréquentielle du transport de la lumière [Durand 2005] (cf. III.3.8) pour exploiter systématiquement la relation entre l'échantillonnage à l'image de l'espace (reconstruction) et d'échantillonnage pour le shading (intégration). Pour la reconstruction, nous proposons un nouvel algorithme multirésolution en une passe. Pour l'intégration, nous estimons le *nombre d'échantillons*. Notre prédiction peut être utilisée en conjonction avec une stratégie d'échantillonnage pour l'intégration numérique. Le gain dû à notre prédiction est complémentaire au gain éventuel de stratégies d'échantillonnage efficaces.

Dans ce chapitre, nous introduisons le concept de calcul et de stockage des fréquences maximales locales du champ de radiance. Nous proposons une représentation concrète de la bande passante, ainsi que qu'une approximation de la variance de la fonction à intégrer, avec un algorithme rapide pour la calculer. Nous utilisons cette information pour adapter les taux d'échantillonnage pour la reconstruction et l'intégration pendant le rendu. Notre algorithme de rendu se compose de deux étapes principales. Tout d'abord, pour chaque pixel, nous estimons la bande passante en espace-écran et la variance. Cette information est stockée hiérarchiquement, dans un tampon ayant la même taille que l'image générée. Ensuite, nous utilisons cette information pour échantillonner l'image. On calcul moins de pixels dans les zones variant de façon régulière, et d'adapte le nombre d'échantillons en fonction de la variance prédite. Nous rendons l'image finale en interpolant les pixels calculés, en prenant en compte les discontinuités.

IV.2 Introduction

As we saw in the previous chapter, photo-realistic image synthesis is very costly due to the need for per pixel numerical integration over multiple incident directions. Using tabulated BRDF measurements ensures maximum realism, up to the accuracy of the measurements, but it makes shading even more computationally expensive due to high memory access and storage cost.

For some applications, such as shape design, interactive feedback of material appearance is desirable. The need to allow dynamically changing geometry poses the first challenge to interactive shading. In addition, it is important that the shading is realistic and consistent with its final appearance after post-design off-line rendering, e.g. using a physically-based off-line renderer.

A gamut of approaches partially address this problem. At one end, fast algorithms focus on editable geometry with simple material models. Other algorithms strive to depict a variety of effects such as global illumination. The latter approach typically requires the pre-computation of radiance transfer (including visibility) and prevents geometry editing [Sloan 2002]. The interactive rendering of editable geometry and realistically portraying complex materials such as acquired BRDFs is still an open research problem and is the focus of this chapter. We achieve this combination at the cost of global illumination and visibility.

Simulating material appearance under all-frequency illumination requires the estimation of an integral of the incident illumination at each pixel modulated by the material's reflectance function. The integrand is typically sampled, and the sampling rate depends on the material: diffuse materials require many samples over incident directions, but exhibit low variation between neighboring pixels; specular materials require fewer samples over incident directions but cause large variation across nearby pixels.

We leverage theories in frequency analysis of light transport [Durand 2005] (refer to III.3.8) to systematically exploit the relationship between sampling in image-space (reconstruction) and sampling for shading (integration). For reconstruction, we propose a new multi-resolution algorithm. For integration, we predict the required *number of samples*. Our prediction may be used in conjunction with any sampling strategy for numerical integration. The gain due to our prediction is complementary to the benefit from efficient sampling strategies.

In this chapter, we introduce the concept of computing and storing the maximum local frequencies of the radiance field. We propose a practical representation of —*local bandwidth*—, as well as an approximation of the —*variance*— of the shading integrand, along with a fast algorithm to compute it. We use this information to adapt sampling rates for reconstruction and integration during rendering. Our rendering algorithm consists of two major steps. First, for each pixel, we estimate the screen-space bandwidth and variance. This information is stored, hierarchically, in a buffer having the same size as the picture generated. Next, we use this information to sample the image. We shade fewer pixels in smoothly varying areas, and adapt the number of samples according to the predicted variance. We render the final image using the scattered shaded pixels by edge-aware up-sampling.

This chapter introduces the following contributions:

1. *Rapid bandwidth and variance computation*: we quickly (about 8ms) predict local variation in the image due to reflected illumination, and expected variance for each pixel.
2. *Multi-resolution shading*: our multi-resolution deferred shading algorithm uses the local frequency information for efficient sampling. We adaptively sample for reconstruction (shading only some pixels) and for integration (number of light samples for each shaded pixel)
3. *Adaptive multi-sample anti-aliasing*: we only compute sub-pixel shading for those pixels where the predicted image-space frequency is greater than 1 pixel^{-1} .
4. *Reflectance bandwidth estimation*: we estimate local bandwidth of arbitrary reflectance functions using wavelets.
5. *Adaptive sampling for pre-convolved shading*: Reflectance usually filters out high frequencies in the illumination. We take advantage of this fact by taking less samples for integration using pre-convolved illumination.
6. *Local light sources*: We show applications of our technique with local light sources such as an area-light source.

IV.2.1 Overview

The radiance arriving at each pixel p after one-bounce direct reflection at a point x (ignoring visibility) is

$$L_p = \int_{\Omega_x} L_i(\omega) \rho(x, \omega, \omega_{x \rightarrow p}) \omega \cdot \mathbf{n}(x) \, d\omega. \quad (\text{IV.1})$$

Here $\omega_{x \rightarrow p}$ denotes the direction from x to the eye through pixel p , $\mathbf{n}(x)$ is the normal at x , L_i is radiance from illumination, Ω_x is the set of incident directions on the hemisphere above the local tangent plane, and ρ is the reflectance function. This integral is typically estimated using Monte Carlo estimators as an average of N_p illumination samples:

$$L_p \approx \frac{G}{N_p} \sum_{i=1}^{N_p} \frac{L_i(\omega_i)}{g(\omega_i)} \rho(x, \omega_i, \omega_o) \omega_i \cdot \mathbf{n}(x) \quad (\text{IV.2})$$

where the $\omega_i \in \mathcal{S}^2$ are random incidence directions distributed according to the importance function $g(\omega_i)$ and G is the importance function integrated over Ω_x .

We accelerate rendering by first, *avoiding shading all pixels*: we compute the integral for an adaptively sparse set of pixels depending on local variations, and up-sample from neighboring pixels for the others (Section IV.4.2). Second, for each pixel p where we estimate the integral, we *adaptively choose* N_p according to the predicted variance of the shading integrand (Section IV.3.4). In Section IV.4, we present a multi-resolution shading algorithm that implements this two-fold strategy.




Figure IV.1: Flatland illustration of local bandwidth propagation. Our idea is to only propagate local bandwidth information (dotted rectangles). Then, using local bandwidth along a few (typically 16) incident directions at R , we estimate the local image variation at p and use it to determine image-space sampling rates (see Eq. IV.7).

IV.3 Real-time frequency analysis

We only propagate maximum local frequencies (bandwidth) about light paths. See Figure IV.1 for an example of bandwidth propagation in Flatland.

2D bandwidth: We analyze the local lightfield using the parametrization of Durand et al. [Durand 2005]. Their parametrization is in 4D and we define the bandwidth of the local light-field as a 2D vector with the maximum non-zero Fourier frequencies in space and angle. For robustness we use a quantile (the 95th percentile) of the power spectra rather than the absolute maximum. For non-band-limited signals, we store an arbitrarily large value until the final calculation in image-space, where we clamp to the maximum representable frequency, which depends on the extent of anti-aliasing chosen. We denote the bandwidth using $\nu \equiv [\nu_s \ \nu_a]^T$ so that the rectangle with opposite corners $(-\nu_s, -\nu_a)$ and (ν_s, ν_a) contains the 2D spatio-angular spectrum of the local light field around a central ray (Figure IV.2).

From [Durand 2005], we derive simple linear transformations undergone by ν for each step of the transport process (see Figure IV.3). We describe how to derive sampling rates using the bandwidth information. Finally, we explain how to estimate the variance of the shading integrand for adaptive sampling.

IV.3.1 Illumination, textures, and BRDF bandwidth

We present in this section a general technique we used to estimate local bandwidth on 2D signals. We used it for textures, distant illumination, BRDF and local light sources, since they all can be easily converted to 2D images—or sets of 2D images.

The most naive way to estimate 2D bandwidth would be to compute a windowed Fourier transform of the input signal (2D image) and measure the bandwidth of the spec-




Figure IV.2: *Left*: 4D local light-field parametrization adopted by Durand et al. [Durand 2005]. *Right*: 2D parametrization introduced by Soler et al. [Soler 2009]. We define local bandwidth $\nu \equiv [\nu_s \ \nu_a]^T$ (black dot) so that 95% of the spectral energy lies in the dotted rectangle.

	<i>4D ray space</i>	<i>Fourier domain</i>
<i>Transport (free space)</i>	spatial shear	angular shear
<i>Occlusion</i>	product	convolution (spatial)
<i>Curvature</i>	angular shear	spatial shear
<i>BRDF</i>	convolution (angular)	product
<i>Texture</i>	product	convolution (spatial)

Figure IV.3: Review of spectral operations from [Durand 2005].

trum. This is a costly computation, and at the same time it does not provide a usable bandwidth estimate, because it implies a compromise between locality and the range of possibly measured frequencies.

Instead, we use 2D wavelets. Wavelets are very localized both in space and frequency, as proven by theory [Mallat 2008] and illustrated in Figure IV.4. Consequently, the set of wavelets that contribute to the image value at a given pixel may provide a good estimate of the local spectrum in the image. Of course, power spectra are not additive because doing so neglects the phase of the different wavelet bands. This overestimates the image bandwidth.

More formally, for a given point x in the image signal s , we have:

$$s(x) = \sum_i \beta_i \phi_i(x) + \sum_i \sum_j \lambda_{i,j} \frac{1}{2^i} \psi\left(\frac{x - 2^i j}{2^i}\right).$$

(where ϕ is the scale function and ψ is the mother wavelet). Wavelets of the same scale have identical bandwidths, so we compute the maximum wavelet coefficient $\lambda_i^{max} = \max_j |\lambda_{i,j}|$ per frequency band, and estimate the wavelet level as:

$$I_w = \operatorname{argmin}_i \sum_{k=i}^n \lambda_k^{max} < \varepsilon \max_k \lambda_k^{max} \quad (\text{IV.3})$$

The result is independent of ε as long as it is small enough. We used $\varepsilon = 0.01$ in all our experiments. Figure IV.5 illustrates this process.




Figure IV.4: Most wavelets have a band-limited spectrum. We use Daubechies wavelets of order 4, which are drawn at the *Left* for various detail levels (bands 2,3,5 and 6 respectively). *Right*: power spectra of the same wavelets. Most of the spectrum energy is concentrated around frequency 2^{n+1} where n is the wavelet band.




Figure IV.5: Instant 2D bandwidth estimation using wavelets. The source image is first converted into wavelets using a 2D transform, alternating horizontal and vertical pyramid steps (*left*). Then for each pixel, we build a histogram of wavelet energy for each frequency band (*middle*). We extract the 95 percentile of this energy to get the maximum order of significant wavelets bands for that pixel. This number is converted into bandwidth thanks to the very limited spectrum of each wavelet.

We experimented with various methods based on wavelets. An alternative is to apply a 1D wavelet transform to image lines and columns separately and apply the percentile extraction to the spectrum of each line. This gives for each pixel a horizontal and a vertical bandwidth. Taking the maximum value of the two gives a result practically similar to the estimate given by the presented method.

Distant illumination: The local 2D bandwidth of distant illumination along ω is purely angular:

$$v(\omega) = [0, v_a(\omega)]^T$$

The distant illumination is first converted into an angular map. Then, for each pixel, the image wavelet hierarchy level I_w from Eq. IV.3 is converted back into angular bandwidth using

$$v_a(\omega) = \frac{2\pi}{2^{I_w} \lambda_{max}}$$





Figure IV.6: Input environment map and its local angular bandwidth computed using a 2D wavelet decomposition.



(a) area-light source (b) spatial bandwidth (c) area-light source (d) spatial bandwidth

Figure IV.7: (a) A finite-sized solid white square-shaped area-light source, and (b) its spatial bandwidth. It contains high bandwidth at the boundaries and zero bandwidth in the middle. (c) A textured area-light source and (d) its spatial bandwidth.

where λ_{max} is the maximum eigenvalue of the Jacobian for the 2D mapping of spherical coordinates onto the image plane. This approach allows to compute instant angular bandwidth in real time on GPU for environment maps. See Figure IV.6, for an example of computing the local bandwidth of an environment map as distant illumination.

Local light sources (spatial): The bandwidth of an area-light source is purely spatial, and is computed using the above method. The area-light source can have any planar shape and can be textured.

To correctly account for the boundary effects of a finite-sized area-light source, we pad the area-light source image with zeros, so that its size is twice the original. We compute the bandwidth of the zero-padded image and crop it back to the original size afterwards. Figure IV.7 shows an un-cropped solid white square-shaped area-light source as well as a textured area-light source example.

Texture (spatial): We extract the spatial bandwidth using the same approach, this time accounting for the Jacobian of the mapping onto the surface so that the bandwidth is correctly expressed in inverse meters.

Reflectance (angular): We only demonstrate separable reflectance: an spatially homoge-

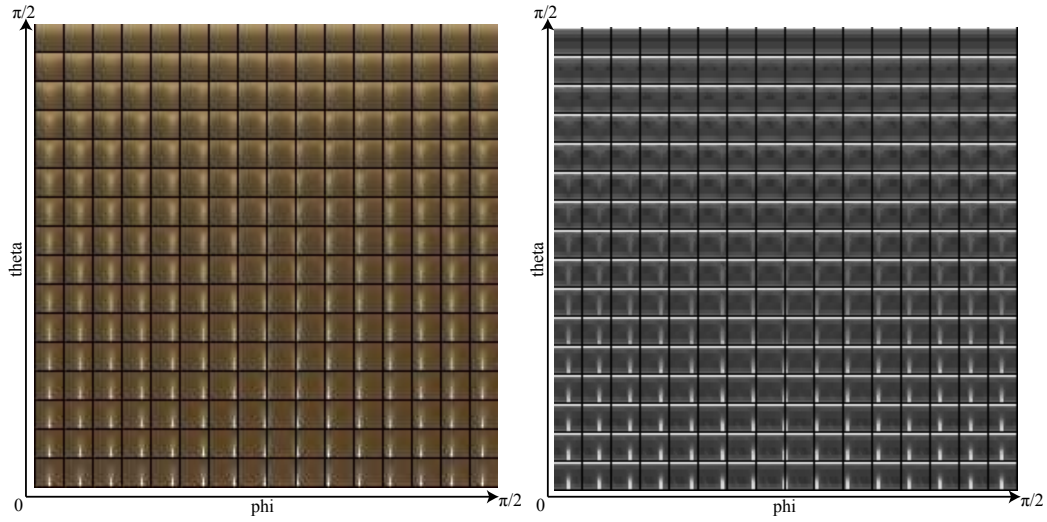


Figure IV.8: Input BRDF (*gold-paint*) and its local angular bandwidth computed using a 2D wavelet decomposition. Each square block represents a single slice of the BRDF.

neous angular reflectance distribution along with a texture. However, all the derivations for bandwidth hold for spatially-varying BRDFs.

For each incident direction in the local tangent frame, we compute the angular bandwidth map of the outgoing BRDF lobe. We use the same technique as for distant illumination and apply 2D wavelet transforms on the slices. We store the result for each lobe of the BRDF in a large texture. For the general case of 4D reflectance data we use 16×16 input directions and a 16×16 image for each reflectance lobe, packed into a 256^2 texture. Since the maximum expressible bandwidth depends on resolution, we compute the bandwidth for higher-resolution angular slices and reduce it to 16×16 . See Figure IV.8 for an example. We tried compression of the BRDF bandwidth with principal component analysis but did not observe any improvement.

IV.3.2 Computing one-bounce 2D bandwidth

At the light: the bandwidth of the local light field leaving light sources depends on the geometry and emission of the light sources. For distant illumination, ν_s is zero and ν_a is directly computed from the environment map. For local light sources at finite distance, such as an area-light source, ν_a is zero and ν_s is computed directly from the light source.

Transport through free space: since transport through free space results in an angular shear of the local light field's spectrum [Durand 2005], the transported bandwidths can be written as $T_d \nu$ for transport by a distance of d (see Figure IV.1), where T_d is defined in Figure IV.9.

Reflection: in the frequency analysis framework, reflection is realized in four steps [Durand 2005]:

1. Re-parametrization of the incident light-field into the frame of the reflecting surface.

$$\begin{aligned}
T_d &= \begin{bmatrix} 1 & 0 \\ -d & 1 \end{bmatrix} & P_x &= \begin{bmatrix} \frac{1}{\cos \theta_x} & 0 \\ 0 & 1 \end{bmatrix} & P_i &= \begin{bmatrix} \cos \theta_i & 0 \\ 0 & 1 \end{bmatrix} \\
S &= \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} & C_c &= \begin{bmatrix} 1 & c \\ 0 & 1 \end{bmatrix} & \mathcal{B}_{t,\rho} &\equiv \begin{bmatrix} \nu_s + t \\ \min(\rho, \nu_a) \end{bmatrix}
\end{aligned}$$

Figure IV.9: Matrix operators on 2D bandwidth.

This is a spatial scale by $\cos \theta_i$ of the spectrum, followed by a spatial curvature shear of length c in the frequency domain (c is the Gaussian curvature of the surface expressed in m^{-1});

2. The product with the incident cosine, which is an angular convolution in Fourier space with a Bessel function;
3. The angular convolution of the light field with the BRDF is a band-limiting product in the frequency domain, while the spatial product by the texture is a convolution by the spectrum of the texture in the Fourier domain.
4. Re-parametrization along the outgoing direction. This is a mirror reflection in the spatial domain, followed by a spatial curvature shear of length $-c$ and a spatial scale of $1/\cos \theta_x$.

We translate these into matrix operations applicable to the bandwidth vector ν of the incident local light field (see Figure IV.9). The re-parametrization (first and last steps) are simply scaling (P_i and P_x). The curvature causes a shear of ν by matrices C_c and C_{-c} . The mirror re-parametrization is a multiplication by matrix S . The reflectance function band-limits angular frequencies based on its own angular bandwidth ρ while the convolution with local texture augments the spatial bandwidth by the bandwidth t of the texture. We denote this using the operator $\mathcal{B}_{t,\rho}$. We neglect the product by the incident cosine, which only adds a small constant to the angular frequency.

We quickly pre-compute angular and spatial bandwidths of the reflectance distribution (and texture). This computation is applicable to any type of reflectance function (analytical BRDFs, acquired BRDFs or artistic shaders).

The overall transformation undergone by incident bandwidths during reflection can thus be represented by a reflection operator \mathcal{R} over the bandwidth vector:

$$\mathcal{R} = P_x C_{-c} S \mathcal{B}_{t,\rho} C_c P_i \quad (\text{IV.4})$$

The bandwidth around a light path arriving at pixel p after one-bounce of a single ray from the light is

$$\nu^i = T_{d'} \mathcal{R} T_d \nu_l^i \quad (\text{IV.5})$$

Here d is the distance from the light source to the bouncing point on the surface, d' is the distance from the surface to the image plane and ν_l^i is the bandwidth while originating at the light source along direction ω_i .

IV.3.3 Image-space bandwidth and sampling rate

The bandwidth at pixel p depends on the choice of ω_i sampled at x . That is the 2D bandwidth ν at pixel p is a combination of the individual bandwidths ν^i along the sampled directions.

We compute the 2D bandwidth at each pixel ν as a weighted average of the sampled incident illumination $L_i(\omega_i)$ at x , reflectance and the 2D bandwidths of the associated one-bounce paths ν^i :

$$\nu = \begin{bmatrix} \nu_s \\ \nu_a \end{bmatrix} = \frac{1}{\sum_{j=1}^{n_b} \lambda_j} \sum_{i=1}^{n_b} \lambda_i \nu^i \quad (\text{IV.6})$$

with

$$\lambda_i = L_i(\omega_i) \rho(\omega_i, \omega_{x \rightarrow p}) \omega_i \cdot \mathbf{n}_x$$

Although the bandwidth at each pixel is estimated using multiple samples, a small choice of n_b is sufficient (see Figure IV.17).

Using a max operation in place of the sum in Eq. IV.6 would be the only conservative choice, but it does not capture view-dependent effects. For instance bandwidth after reflection from a specular sphere would be equally high regardless of viewing or light direction (see Figure IV.10). Eq. IV.6 is a heuristic approximation to the actual variance of the shading integrand that is bounded to the range of bandwidths from contributing light paths while giving more credit to light paths with larger energy (If all light paths contribute the same bandwidth, the approximation is exact). This way, we account for the material reflectance, relative orientation of illumination and view, and local geometry.

The required sampling rates at the image plane are twice the local image-space bandwidth (Nyquist criterion) b_p (in pixel^{-1}):

$$b_p = \nu_a \max \left[\frac{f_x}{W}, \frac{f_y}{H} \right] \quad (\text{IV.7})$$

where f_x and f_y are the horizontal and vertical fields of view, and the rendered image is $W \times H$ pixels.

IV.3.4 Adaptive sampling for shading

In this section, we estimate the variance of the Monte-Carlo estimator L_p of Eq. IV.1 in order to determine a suitable number N_p of integration directions for each pixel p . Because there is no analytical formula giving the variance $V(L_p)$ of L_p using importance sampling, we conservatively overestimate it assuming uniform sampling:

$$V(L_p) \leq \frac{4\pi^2}{N_p} \sigma_p^2$$

The variance σ_p^2 of the shading integrand about a single illumination direction ω_i , at a point x that projects to pixel p , is

$$\begin{aligned} \sigma_p^2 &= \text{E}(\lambda_i^2) - \text{E}(\lambda_i)^2 \\ &\leq \text{E}(\chi_i^2) \end{aligned}$$




Figure IV.10: Combining bandwidth estimates from sampled incident directions. *Middle*: Applying a *max* overestimates sampling rates (1 pixel^{-1} almost everywhere on the sphere). *Right*: Our approach. Eq. IV.6 predicts view-dependent sampling rates. *Left*: Final result.

where $\chi_i = L_i(\omega_i)\rho(\omega_i, \omega_{x \rightarrow p})$, and $E(X)$ denotes the expected value of X . Finally,

$$E(\chi_i^2) = E(\hat{L}_i^2 \otimes \hat{\rho}^2)$$

where \hat{f} denotes the Fourier transform of f and \otimes denotes convolution. This equality is a consequence of Parseval's theorem (see, e.g. [Oppenheim 1975]). The convolution is in the angular domain. We have:

$$E(\hat{L}_i^2 \otimes \hat{\rho}^2) \leq (v_a^i + v_{\rho a}^i) \chi_i^2 \quad (\text{IV.8})$$

where v_a^i is the angular component of $v^i = T_d v_l^i$, and $v_{\rho a}^i$ is the local angular bandwidth of the reflectance function. Therefore, we have:

$$V(L_p) \leq \frac{4\pi^2}{N_p} \sum_{i=1}^{n_b} (v_a^i + v_{\rho a}^i) \chi_i^2 \quad (\text{IV.9})$$

To keep the variance of the shading globally constant, we need to keep N_p proportional to the sum of the bandwidths, weighted by the illumination and reflectance values, along sampled directions ω_i . The sum is a conservative approximation of the variance of the integrand. $n_b = 16$ provides acceptable quality (see Figure IV.17).

The summations over incident directions — Eq. IV.6 and Eq. IV.9 — indicate that we implicitly account for the relative alignment (phase) of the illumination and reflectance. Previous approaches that neglect phase cannot predict variation due to view-dependent effects.

IV.3.5 Adaptive sampling for pre-convolved shading

As an approximation, we propose a faster integration scheme, that would need far less directional samples.

Starting again from Eq. IV.1, we can write the illumination to be the sum of a low frequency component L_i^l and a high frequency component L_i^h . The high frequency component

is going to cancel out due to the low-pass filter characteristics of the reflectance:

$$\begin{aligned} L_p &= \int_{\Omega_x} (L_i^l(\omega) + L_i^h(\omega)) \rho(x, \omega, \omega_{x \rightarrow p}) \mathbf{n} \cdot \omega d\omega \\ &\approx \int_{\Omega_x} L_i^l(\omega) \rho(x, \omega, \omega_{x \rightarrow p}) \mathbf{n} \cdot \omega d\omega \end{aligned}$$

Doing this, we can rewrite Eq. IV.9, replacing the bandwidth of the illumination by the bandwidth of the filtered illumination, which in turn is equal to the bandwidth of the BRDF:

$$V(L_p) \leq \frac{4\pi^2}{N_p} \sum_{i=1}^{n_b} 2\nu_{\rho a}^i \chi_i^2. \quad (\text{IV.10})$$

To compute the shading integral, we pick samples in the pre-filtered illumination L^l (e.g a pre-filtered environment map), at the level given by the maximum frequency of the BRDF. The number of samples needed to compute the shading integral is therefore greatly reduced.

The calculation we performed is based on the same principle used by pre-convolved environment map shading [Kautz 2000]. The main difference is that we do not project the BRDF nor the illumination onto a basis function. We only filter out parts of the illumination that are out of the support of the spectrum of the BRDF. We still perform a Monte-Carlo integration to compute L_p , but using less samples. Section IV.5.3 gives examples of using this technique.

IV.3.6 Implementation roadmap

The practical implementation is simply as follows: for each pixel, we compute the image bandwidth using Eq. IV.6 and Eq. IV.7, and the number of samples for the shading integrand using Eq. IV.9. Both are summations over $n_b = 16$ incoming directions ω_i . For each ω_i , the local 2D bandwidth ν^i is given by Eq. IV.4 and Eq. IV.5, using the matrices listed in Figure IV.9. These matrices require the curvature c , normal n , the incident and outgoing angles (θ_i and θ_o), the pre-computed spatial and angular bandwidth of the material (resp. t and ρ) for the current pixel and direction ω_i , and the pre-computed bandwidth of the light source ν_l^i in direction ω_i .

IV.4 Hierarchical shading algorithm

Our rendering algorithm consists of four steps (see Figure IV.11):

1. load the pre-computed bandwidth for BRDF(s) and compute illumination bandwidth on the fly.
2. a geometry pass that renders G-buffers such as depth, position, normals and material IDs.
3. we fill the bandwidth buffer with image-space bandwidth, the number of integration samples to use per pixel and the screen-space bandwidth.

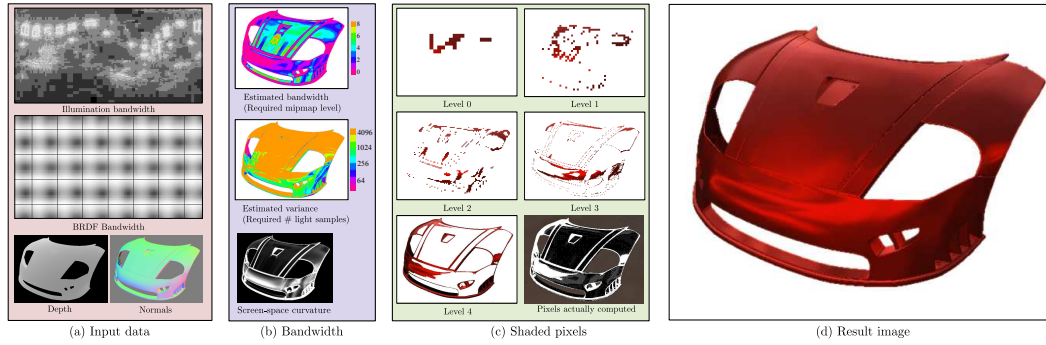


Figure IV.11: Our rendering pipeline: (a) at each frame, we first render G-Buffers. From these, we compute an additional *bandwidth buffer* (b) that stores image space bandwidth and shader integrand variance maps as well as screen-space curvature. The former is stored in a multi-resolution pyramid. During rendering of the final image from coarse to fine scale, depending on our bandwidth and variance predictions, pixels are either explicitly shaded (numerical integration) or up-sampled from parent pixels (c). Once we reach the finest level, the image is fully rendered (d).

4. a single-pass multi-resolution shading step, interleaved with up-sampling.

Rendering G-Buffers is a classical geometry pass where we store normals, position, depths and material IDs into a set of screen-space buffers. Note that G-Buffers do not need to be hierarchically built (mip-mapped) in our method; we build a multi-resolution pyramid only for the bandwidth buffer.

IV.4.1 Bandwidth buffer initialization

The bandwidth buffer contains two important values: the local image-space bandwidth and the number of samples to be used for shading each pixel. These are computed using the G-Buffers ((IV.7) and (IV.9)). Although these estimations involve numerical integration, they are several orders of magnitude faster than the actual shading, since a coarse sampling is sufficient (Figure IV.17). Rather than storing b_p (see (IV.7)) in the bandwidth buffer, we store

$$\min(\lfloor \log_2 \frac{1}{b_p} \rfloor, \min(\log_2(W), \log_2(H))) \quad (\text{IV.11})$$

which is the pyramid resolution at which pixel p needs to be shaded, accounting for the local variation at p . The *floor* operation ensures that the Nyquist sampling rate is respected. Storing b_p directly in the bandwidth buffer leads to identical results; our optimization simplifies tests for deciding the pyramid resolution while shading each pixel.

The bandwidth is mip-mapped using a min filter, so that at a given level in the hierarchy, the value for a pixel conservatively tells us whether sub-pixels should be computed at this level. We do the same for the variance estimate using a max filter.

In addition, we estimate screen-space curvature on-the-fly and use it in the bandwidth estimation, without the need for storing it for later use.

```

for all points  $p$  at level  $L$  do
  if  $b_w(p) < L$  then
    compute  $w_0, \dots, w_3$ 
     $c(p) \leftarrow \sum_k w_k c(p_k)$ 
  else if  $b_w(p) == L$  then
    shade( $p$ )

```




Figure IV.12: Up-sampling interpolation scheme. *Left*: Pseudo-code for the computation of one level. *Right*: relative weights α_i for parent pixels of pixel p at the next level.

IV.4.2 Shading and up-sampling

We render the image hierarchically, progressively from coarse to fine. At a given resolution (say $2^k \times 2^k$), we examine the bandwidth buffer and shade the pixels for which the bandwidth buffer pyramid contains the current coarseness resolution k . For pixels whose bandwidth buffer entries are less than the current resolution (i.e. $< k$), we bilaterally up-sample from neighbors at the preceding level of coarseness ($2^{k-1} \times 2^{k-1}$), only accounting for pixels that are already computed. The parents' values are averaged with coefficients

$$w_i = g_z(z - z_i) g_a(p - p_i) \alpha_i$$

where z (resp. z_i) are the depths of the shaded (resp. parent) pixels, and g_z is a Gaussian that cancels out pixels of irrelevant depth, and α_i are bilinear weights (Figure IV.12). The last term g_a is an anisotropic 2D Gaussian defined as

$$g_a(\mathbf{v}) = e^{-\mathbf{v}^T M \mathbf{v}} \quad \text{with} \quad M = R_\phi^T \begin{bmatrix} 1 & 0 \\ 0 & \frac{1}{\cos^2 \theta_x} \end{bmatrix} R_\phi$$

where ϕ is the angle of the screen-projected normal at the surface, and θ_x the angle between the normal and the view direction. This enables efficient anisotropic filtering aligned with the highest and lowest screen-space frequencies, since $\frac{b_p}{\cos \theta_x}$ and b_p estimate the minimum and maximum directional screen-space bandwidth around current pixel.

We continue this process over successive levels, until we reach the finest resolution where we shade all remaining pixels. Pseudo-code for the algorithm is presented in Figure IV.12.

IV.4.3 Shading computation

For each shaded pixel we read the number of samples N_p from the bandwidth buffer. We estimate reflected radiance (IV.2), unless stated otherwise, by importance sampling the BRDF lobe for the current view direction. In our implementation, we read N_p samples randomly, from multiple pre-computed vectors of importance samples that are stored in a texture. Our algorithm is compatible with any importance sampling strategy. In practice, we importance sample the reflectance by numerical inversion of its cumulative distribution.




Figure IV.13: Three identical objects with different materials (`color-changing-paint3`, `red-metallic-paint` and `gold-paint`), and a marble textured ground floor with `white-marble` material. We shade more pixels per area for specular materials (red arrows) but we use a larger number of samples to compute shading on diffuse materials (black arrows).

We always shade with depth and normal values at the finest level, since the G-Buffers are not mip-mapped. This is possible because the bandwidth buffer predicts whether, for any sub-pixel of the current level pixel, the computation will yield similar estimates despite potential variation in the depths, normals and illumination.

IV.5 Results and discussion

All timings reported here were measured on an NVIDIA GeForce GTX 560 Ti graphics card with 1GB of memory.

Figure IV.13 shows an example of a scene with three identical objects but different materials and a marble textured ground floor. This is a clear example of how our fast frequency analysis helps us spend the samples where they are most needed.

IV.5.1 Behavior of our algorithm

Computation time: the computation time for our algorithm scales linearly with the total number of shading samples (Figure IV.14, left). The total number of shading samples required depends on the desired image quality, the material and the illumination. Figure IV.15 tabulates the computation times for our algorithm to obtain equal quality as ground truth, for several scenes. It details the cost of individual steps: the cost of bandwidth computation is independent of the scene and the material, and negligible compared to the overall cost (8 ms, or less than 0.33%). Shading estimation consumes most of the total time (up to 90 %).

Memory cost: The memory footprint of our algorithm on the GPU is approximately




Figure IV.14: *Left*: Rendering times are linear in the number of shading samples for various models. *Right*: Rendering time against number of samples for anti-aliasing (blue) in log scale. Our algorithm scales sub-linearly (with an exponent of 0.7) for anti-aliasing.

Model	Car body	Bunny	Buddha	Dragon
Reference	2056 ms	1628 ms	2634 ms	1739 ms
Ours: Total	223 ms	398 ms	207 ms	190 ms
Bandwidth calculation	7 ms	8 ms	7 ms	8 ms
Shading integration	216 ms	390 ms	200 ms	182 ms

Figure IV.15: Computation times (in ms) at 512×512 screen resolution: Car body (Fig. IV.13), Bunny (Fig. IV.17), Buddha (Fig. IV.16) and Dragon (Fig. IV.23). Our algorithm provides a 4 to 10 \times speedup compared to a forward-shaded reference of comparable quality, depending on the material and screen occupancy. Bandwidth computation is fast (< 8 ms) for all scenes.

432 MB at a resolution of 512×512 pixels: 17 MB for the G-buffers (4 RGBA buffers for position, normal, tangent, material ID and the depth buffer); 2 mip-mapped buffers (5.5 MB each) for the bandwidth buffer and for the shading computations with up-sampling; 2 RGB buffers of 6 MB for the environment map and its bandwidth; and 2 buffers of 196 MB for the raw BRDF data and its importance samples. Increasing the picture resolution only increases the cost of the G-buffers and mip-mapped buffers: 112 MB (instead of 22.5) for 1024×1024 and 448 MB at 2048×2048 , which is currently the maximum for our algorithm.

Validation: Figure IV.16 compares our predictions with reference variance and image-space bandwidth. Our predictions are similar in spatial distribution, and of the same order of magnitude. Our variance estimate is conservative, as explained in Section IV.3.4. We computed the reference bandwidth² using a windowed Fourier transform over the image,

²Reference local frequencies cannot be computed exactly using the windowed Fourier transform because

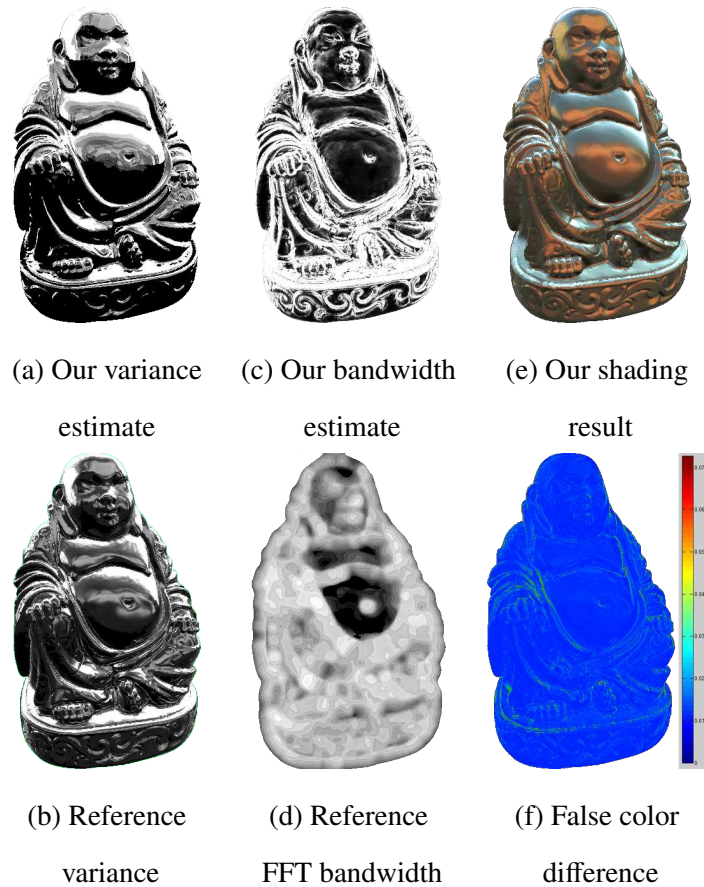


Figure IV.16: Validation of predicted variance and bandwidth. (a) our estimate for the variance of the shading integrand (Eq. IV.9), (b) reference variance computed using brute-force sampling, (c) our estimate of the image bandwidth, (d) reference image bandwidth (windowed FFT of reference image), (e) our result, and (f) relative error of (e) with respect to a path-traced reference. Material: `color-changing-paint3`

with a window size of 32×32 pixels. We computed the reference variance using extensive sampling.

Influence of parameters: the main parameter for our algorithm is the number of samples we use for the bandwidth estimation, n_b (see Eq. IV.6). Figure IV.17 shows the influence of varying this parameter. The results are indistinguishable even in the zoomed-in insets and the Peak Signal-Noise Ratio stays almost constant for all values of n_b . The rendering time has a small dependency on n_b . We used a small value, $n_b = 16$, for all results in this work. This makes sense as n_b is only used to estimate the bandwidth and not for the actual illumination computations.

of the uncertainty principle. However, using our wavelet estimation technique here would not constitute a fair test either.




Figure IV.17: Effect of the number of sampled directions n_b for bandwidth computation on time (ms) and quality (PSNR). The data points in the plots are at $n_b = 16, 64, 256, 1024$. The enlarged insets are virtually indistinguishable for the different n_b . The plots depict that fast bandwidth computation ($n_b = 16$) is sufficient. Resolution: 512×512 . Material: color-changing-paint3.

IV.5.2 Comparison with related work

Brute-force rendering: Figure IV.18 compares our result with a forward shading reference computed using importance sampling and a fixed number of shading samples per pixel. For this scene, we achieve a $2.5\times$ speedup due to our adaptive sampling. The extent of our speed-up depends on the material, the environment map and the area occupied by the object on screen. Figure IV.15 tabulates rendering times for our algorithm and brute-force rendering for several scenes. We observe a speed-up of $4\times$ to $10\times$.

Spherical Gaussian approximation: Figure IV.19 compares the results of Wang *et al.* [Wang 2009a] with ours and ground truth computed using path-tracing. The authors kindly provided us with their best images for the materials. Our algorithm accurately shades acquired materials. The fast algorithm [Wang 2009a] is visibly different from the ground truth.

IV.5.3 Rendering results

Adaptive sampling for pre-convolved shading: To use adaptive sampling for pre-convolved shading (ASFPCS), we pre-convolve the environment map with Phong lobes of various exponent and assemble them into an array of textures. At the time of rendering, we estimate the number of adaptive samples using Eq. IV.10) and look-up pre-filtered illumination for

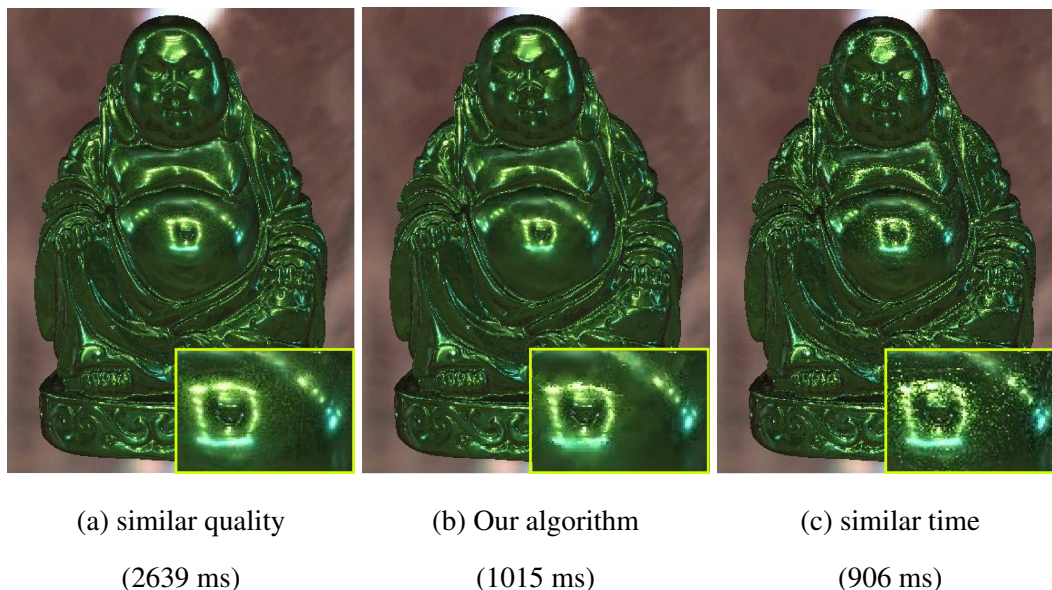


Figure IV.18: Comparison of our algorithm (b) with a reference for equal quality (a) and equal time (c). For this scene, we achieve a $2.5\times$ speedup (without anti-aliasing). Forward-shaded references use BRDF-importance sampling and a fixed number of shading samples per pixel. Material: green-metallic-paint.

each sample according to the maximum frequency of the BRDF. figure IV.20 depicts examples of using adaptive sampling for pre-convolved shading (ASFPCS). We achieved a significant speed-up (10 to 20 times depending on the material and scene coverage).

Area-light source rendering: To render area-light sources, we propagate bandwidth through the scene from the area-light source towards the camera. Once we know the adaptive number of samples for shading by bandwidth propagation, we sample the area-light source directly (no BRDF importance sampling) and shade in the hierarchical buffer. Figure IV.21 depicts two scenes rendered with the area-light sources from Figure IV.7.

Adaptive multisample anti-aliasing: Our bandwidth prediction reduces the cost of multi-sample anti-aliasing by adaptive sampling. Standard deferred shading evaluates shaders at every sample: 16 samples per pixel costs 16 times more. Our algorithm scales sub-linearly in the number of samples (see Figure IV.14, right). We render the G-buffers at the higher resolution (4 or 16 times the number of pixels), but compute shading at the appropriate level in the pyramid, depending on the predicted bandwidth. Anti-aliasing only requires a few extra shader evaluations at the finest levels (blue pixels in Figure IV.23).

Dynamic geometry, normal and displacement mapping: We compute bandwidth and variance estimates using only the information from the G-buffers (normals, geometry and curvature). Our algorithm handles dynamic geometry, displacement mapping and normal




Figure IV.19: Comparison with related work [Wang 2009a]. Our algorithm results in pictures that are identical to ground truth, while [Wang 2009a] result in clear differences.

maps seamlessly. Figure IV.22 shows an example of our bandwidth estimation algorithm on a dynamic geometry, running at 15 fps.

IV.5.4 Discussion

Best and worst case: The rendering cost for our algorithm depends on the total number of shading samples predicted. The best case is when the predicted variance is low (a specular material) with low frequency illumination or when the spatial variation is low (a smooth diffuse surface). We do not save time when both angular variance and spatial frequencies are high (a bumpy surface with high frequency illumination).

Conservative bandwidth estimate: We conservatively predict bandwidth, choosing sub-optimal (excessive) sampling over artifacts from insufficient sampling. The min operator of Figure IV.9 is larger than the real bandwidth of the product of the BRDF and illumination; we also over-estimate variance.

Bandwidth estimate and importance sampling: Since we base all our bandwidth predictions on uniform sampling and render the result using BRDF importance sampling, our predictions for variance are sometimes very conservative, and it happens that an acceptable result can still be achieved with less samples. It would be interesting to derive acceptable approximations of the variance for importance sampling, but we keep this as a future avenue.

Visibility and global illumination: we focus on accurate depiction of material appearance rather than scenes. We ignore effects such as visibility for shadows and global illumination.

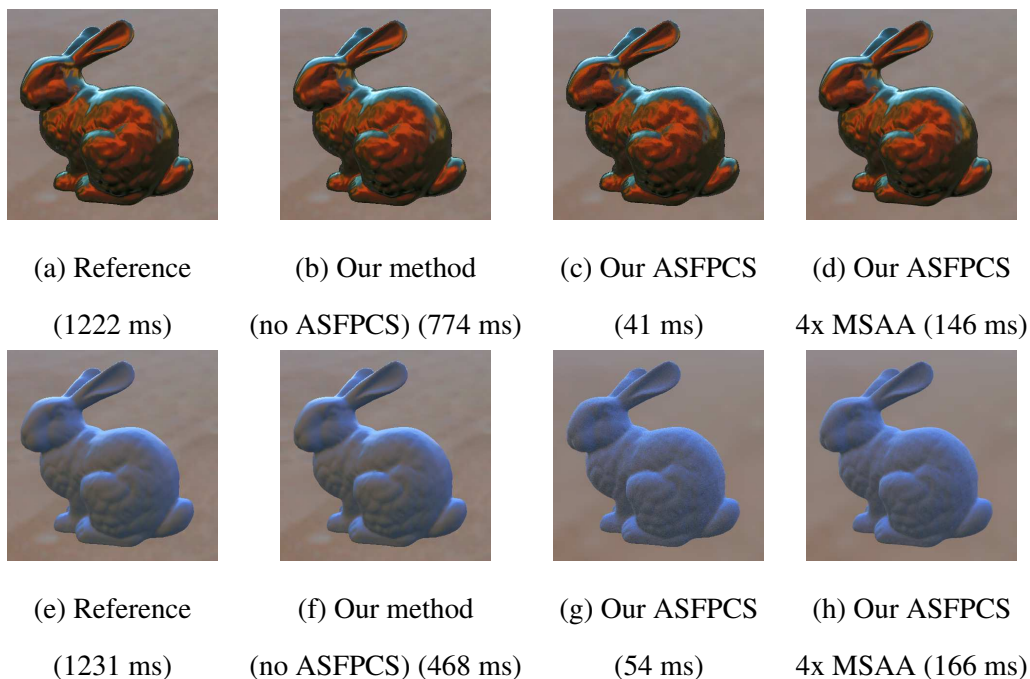


Figure IV.20: Adaptive sampling for pre-convolved shading (ASFPCS) examples (c, g) for `color-changing-paint3` and `blue-rubber`, compared to the reference (a,e) and our adaptive sampling method without pre-convolved shading (b,f). Our ASFPCS technique with 4x multi-sample anti-aliasing is still faster than our method without pre-convolved shading. The speed-up is 10 to 20 times.

Spatially-varying BRDFs: we have only used homogeneous (non spatially-varying) materials, modulated by a texture. We pre-computed local bandwidths separately for reflectance ($4D$) and texture ($2D$). Extending our algorithm to fully spatially varying BRDFs ($6D$) is possible at the extra cost of $6D$ bandwidth pre-computation.

IV.6 Conclusion

We have introduced an algorithm for interactively and accurately shading dynamic geometry with acquired materials. Our contribution is to:

- only shade a small fraction of pixels where the local bandwidth is predicted to be large.
- adaptively sample shading integrals based on the predicted variances.

We achieve these predictions by quickly computing local bandwidth information from standard G-buffers. We have introduced the concept of a bandwidth buffer, to store this infor-




Figure IV.21: *Left*: Buddha rendered with a solid white square area-light source and red-metallic-paint material (513 ms). The area-light source is positioned at the top of the object and is shown as inset. *Right*: Dragon rendered with a textured area-light source and color-changing-paint3 material (376 ms). The area-light source is positioned at the top of the object and is shown as inset.

mation, which is exploited to sub-linearly scale multi-sample anti-aliasing with deferred shading.

Our work can be extended in many ways. We would like to apply this technique to indirect lighting and visibility effects as well. This constitutes a real challenge: although extending our theory to these situations seems feasible, the resulting calculations needed by visibility will increase the computation cost significantly. Accounting for anisotropy in the bandwidth estimate is another interesting avenue that will require significantly more calculations.




Figure IV.22: Application of our algorithm to dynamic geometry, with red-metallic-paint. Since our algorithm depends only on data from the G-buffers, it can handle dynamic geometry, displacement mapping and normal maps. Rendered images as well as pixels for which we performed shading computations (white pixels). The screen-space bandwidth adapts to the curvature (Rendering time: 66 ms).




Figure IV.23: Adaptive multi-sample anti-aliasing using bandwidth information. Aliasing artifacts are visible in the top row of enlarged insets (55 ms). Our multi-resolution algorithm handles anti-aliasing seamlessly by adding extra levels to the pyramid (one level for 4 \times and two levels for 16 \times). We only compute shading at the finest level for the blue pixels in the rightmost column, where the predicted bandwidth is high. The cost of anti-aliasing is thus reduced: 141 ms for 4 \times , 390 ms for 16 \times . Material:gold-paint.

A new distribution for Cook-Torrance BRDF ¹

Contents

V.1	Introduction (français)	65
V.2	Introduction	66
V.3	Our reflectance model	67
V.3.1	Data observation	67
V.3.2	The SGD distribution	69
V.3.3	Shadowing and masking	70
V.3.4	The Fresnel term	71
V.3.5	Importance sampling	72
V.4	Fitting algorithm	73
V.4.1	Fitting RGB channels separately	73
V.4.2	Non-linear optimization	73
V.5	Results and Comparison	74
V.5.1	Visual comparison	74
V.5.2	BRDF Lobes comparison	77
V.5.3	Quantitative Error Measure	77
V.5.4	Influence of D , G , and F parameters in our BRDF model	80
V.5.5	Fitting RGB channels altogether for D parameters	81
V.5.6	Timings	83
V.5.7	Limitations	83
V.6	Conclusion	85

V.1 Introduction (français)

Dans ce chapitre, nous présentons une nouvelle distribution de micro-facettes, la fonction d'ombrage et de masquage, aussi que le term de Fresnel pour la BRDF de Cook-Torrance. Nous introduisons également un ajustement précis et robuste pour les BRDF isotropes (dont les métaux) avec un seul-lobe spéculaire (en plus de la composante diffuse).

¹This chapter is based on the publication [M. Bagher 2012a]




Figure V.1: *gold-metallic-paint* material from the MERL database. Comparison between measured data (c) and Cook-Torrance BRDF using (b) Beckmann distribution or (d) SGD distribution. (a) and (e): difference measured in Lab color space between ground truth and (b) and (d) respectively.

La BRDF de Cook-Torrance [Torrance 1967] simule la géométrie d’un micro Matériau spéculaire composé de micro-facettes, et tire la BRDF mathématiquement à partir de la distribution de probabilité de micro-facettes —voire III.4.4.4.2 pour une introduction. Bien qu’il soit généralement utilisé avec la distribution Beckmann, le modèle peut être utilisé avec de nombreuses autres distributions. Notre étude de BRDFs mesurées [Matusik 2003] montre que les matériaux brillants ne suivent pas la distribution Beckmann : les micro-facettes sont plus susceptibles d’être alignés avec la surface, entraînant un pic dans la distribution, compatible avec l’apparence lisse.

Nous introduisons densité de probabilité SGD pour les micro-facettes. Cette répartition permet une optimisation plus précise pour tous les matériaux mesurés, en comparant à des modèles existants de BRDF, qu’ils soient brillants au diffus Nous utilisons la BRDF de Cook-Torrance avec la distribution SGD pour se rapprocher des matériaux mesurés. Il en résulte en une représentation compacte et précise. Comme la distribution SGD se rapproche très bien de données mesurées, nous sommes capables de concevoir une méthode d’optimization plus simple. Cette méthode converge rapidement (2,5 minutes en moyenne) et utilise seulement deux tranches de données mesurées pour s’adapter à des paramètres inconnus.

V.2 Introduction

In this chapter, we present a new micro-facet distribution, shadowing and masking function, and Fresnel term for the Cook-Torrance BRDF. We also introduce an accurate and robust fitting for any isotropic BRDF (including metals) with a single specular lobe (in addition to the diffuse component).

The Cook-Torrance BRDF [Torrance 1967] simulates the micro geometry of a material as specular micro-facets and derives the BRDF mathematically from the probability distribution of micro-facets — refer to III.4.4.4.2 for an introduction. Although it is typically used with the Beckmann distribution, the model can be used with many other distributions. Our study of measured BRDFs [Matusik 2003] shows that shiny materials do not follow the




Figure V.2: The micro-facet probability distribution for *brass* extracted from the measured data. The SGD distribution (Right) fits the data accurately both at the peak and the tail, unlike existing distributions (Left)

Beckmann distribution: micro-facets are more likely to be aligned with the surface normal, resulting in a sharper peak, consistent with the smooth appearance.

We introduce the SGD probability distribution function for micro-facets. This distribution provides a more accurate fitting for most measured materials in the MERL-MIT database, comparing to existing BRDF models, from highly shiny to diffuse. See Figure V.2. We use the Cook-Torrance BRDF with SGD distribution to approximate measured materials resulting in a compact and accurate representation. Since the SGD distribution approximates measured data very well, we were able to design a simpler fitting method. This method converges quickly (2.5 minutes on average) and it only uses two slices of the measured data to fit the unknown parameters.

This chapter is organized as follows: in the next section, we describe the SGD probability distribution function, and in section V.4 our fitting method. Section V.5 presents our results and compares them with previous work. In section V.6, we conclude and discuss potential avenues for future work.

V.3 Our reflectance model

V.3.1 Data observation

Given the Cook-Torrance BRDF:

$$\rho(\mathbf{i}, \mathbf{o}) = \frac{\rho_d}{\pi} + \frac{\rho_s}{\pi} \frac{F(\mathbf{i} \cdot \mathbf{h})D(\theta_h)G(\mathbf{i}, \mathbf{o})}{\cos(\theta_i) \cos(\theta_o)}$$

and Smith shadowing functions G_1 [Smith 1967] directly computed from D :

$$G(\mathbf{i}, \mathbf{o}) \approx G_1(\theta_i)G_1(\theta_o)$$

, we are interested in observing the behavior of each component of the Cook-Torrance BRDF directly from the measurements. However, it is hard to compare measured data




Figure V.3: The SGD distribution is an accurate fit for a wide range of materials. The distribution is scaled by ρ_d and ρ_s to fit the red curve.

with BRDF models, since models depend on position and several functions. For simple comparisons, we take $\mathbf{i} = \mathbf{o}$. As a consequence, $\mathbf{h} = \mathbf{i}$: we get an equation that only depends on θ_h , with two unknown functions D and G_1 [Ashikhmin 2007]:

$$\rho(\theta_h) = \frac{\rho_d}{\pi} + \frac{\rho_s}{\pi} \frac{F(0)D(\theta_h)G_1(\theta_h)^2}{\cos^2(\theta_h)} \quad (\text{V.1})$$

For small values of θ_h and specular materials, G_1 is almost constant, equal to 1 (see Figure V.5). The measured data varies only with D , and we can observe directly the micro-facet distribution.

Figure V.2 compares the measured distribution D for *brass* from the MERL-MIT database [Matusk 2003] with approximations using different distributions. The peak of the measured distribution is much sharper than the predictions, while the tail is larger. If we fit the tail accurately (for example with the Beckmann distribution, $\alpha = 7 \times 10^{-2}$), then we miss the peak. On the other hand, if we fit the peak well, we miss the tail. To fit this data with existing distributions we would need several lobes. This observation holds for many materials in the MERL-MIT database: metals, metallic paints and shiny plastics.

At large angles, the measured distribution seem to decrease exponentially. But for small angles, the decrease is sharper than an exponential. Based on this observation, we suggest a slope in x^{-p} , where $x = \tan^2 \theta_h$ and p is a parameter that depends on the model, multiplied by an exponential factor:

$$P_{22}(x, \alpha, p) = \frac{K_{\alpha,p} e^{-\frac{x}{\alpha}}}{(x + \alpha^2)^p} \quad (\text{V.2})$$

where $K_{\alpha,p}$ is a normalization factor. With this distribution, we fit the measured data accurately (see Figure V.2(b)), both at the peak and the tail.

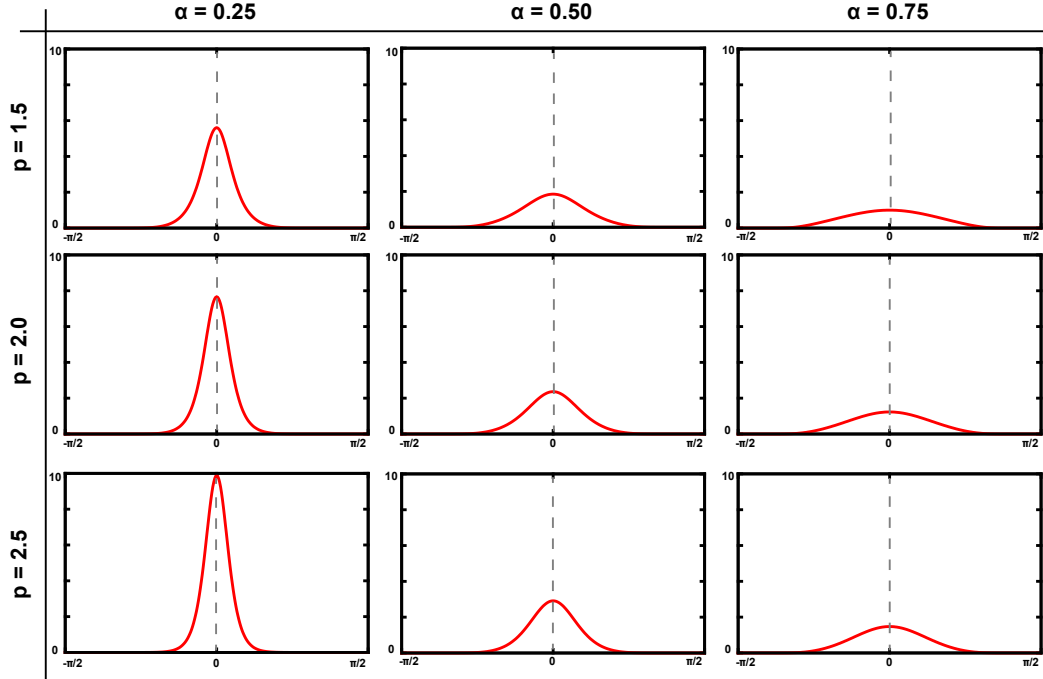


Figure V.4: The influence of parameters α and p on the shape of the distribution.

V.3.2 The SGD distribution

We compute the normalization factor $K_{\alpha,p}$ so that $\int_0^\infty P_{22}(x) dx = 1/\pi$. The result is a shifted gamma distribution (SGD):

$$P_{22}(x) = \frac{\alpha^{p-1}}{\Gamma(1-p, \alpha)} \frac{e^{-\frac{\alpha^2+x}{\alpha}}}{(\alpha^2+x)^p}$$

$$D(\theta_m) = \frac{\chi_{[0, \frac{\pi}{2}]}(\theta_m)}{\pi \cos^4 \theta_m} P_{22}(\tan^2 \theta_m)$$

where Γ is the incomplete Gamma function: $\Gamma(s, x) = \int_x^\infty t^{s-1} e^{-t} dt$

Figure V.4 illustrates the influence of parameters α and p on the shape of the distribution. Both parameters affect the peak and fall off of the curve simultaneously.

The SGD fits accurately a large range of materials in the MERL-MIT database, from diffuse to specular (see Figures V.2(b), V.3 and Appendix A).




Figure V.5: The shadowing term G_1 is expensive for SGD. Our approximation is very accurate; the L_2 error is below 10^{-5} .

V.3.3 Shadowing and masking

Once we have D , we can compute the shadowing and masking function, G_1 , using Equation III.4:

$$\begin{aligned}
 P_2(r) &= \frac{1}{\sqrt{\pi\alpha}\Gamma(1-p, \alpha)} U\left(p, p + \frac{1}{2}, \alpha + \frac{r^2}{\alpha}\right) e^{-\alpha - \frac{r^2}{\alpha}} \\
 \Lambda(\theta) &= \int_{\frac{1}{\tan\theta}}^{\infty} (r \tan\theta - 1) P_2(r) dr \\
 G_1(\theta) &= \frac{\chi_{[0, \frac{\pi}{2}]}(\theta)}{1 + \Lambda(\theta)}
 \end{aligned}$$

where U is the confluent hyper-geometric function of the second kind. We do not have a closed form expression for G_1 . We either pre-compute G_1 using Gauss integration and store its values for varying θ or use the following approximation:

$$\begin{aligned}
 G_1(\theta) &\approx \begin{cases} 1 + \lambda(1 - e^{c(\theta - \theta_0)^k}) & \text{if } \theta > \theta_0 \\ 1 & \text{if } \theta \leq \theta_0 \end{cases} \quad (\text{V.3}) \\
 \theta_0 &= \frac{\pi}{2} - \left(\frac{\log\left(1 + \frac{1}{\lambda}\right)}{c}\right)^{\frac{1}{k}}
 \end{aligned}$$

θ_0 ensures that $G_1(\pi/2) = 0$. We get the (λ, c, k) parameters by fitting the approximation over the precomputed values, using Levenberg-Marquadt [Lourakis 2004]. This approximation is a very good fit for G_1 for all materials (the L_2 error is negligible, mostly below 10^{-5} , see Figure V.5).

The choice between using stored values for G_1 and the approximation depends on the processing power, memory bandwidth and storage capacities. For offline rendering, where storage is not an issue, we use precomputed values. For GPU rendering, we use the approximation.

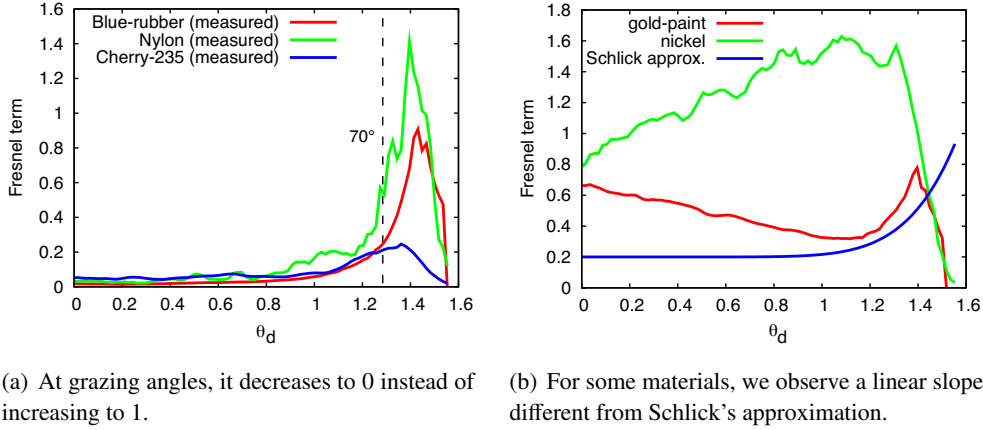


Figure V.6: The Fresnel term extracted from the measured data does not behave as predicted by the theory

V.3.4 The Fresnel term

Since we have an expression for D and G_1 , we can express the Fresnel term using the measured data:

$$F(\mathbf{i} \cdot \mathbf{h}) = \left(m(\mathbf{i}, \mathbf{o}) - \frac{\rho_d}{\pi} \right) \frac{\pi \cos \theta_i \cos \theta_o}{\rho_s D(\theta_h) G_1(\theta_i) G_1(\theta_o)} \quad (\text{V.4})$$

where $m(\mathbf{i}, \mathbf{o})$ is the measured data. For simplicity, we will note $\theta_d = \arccos(\mathbf{i} \cdot \mathbf{h})$. Looking at the data for the Fresnel term, we make two important observations:

- The measured data behaves strangely at grazing angles. In theory, the Fresnel term should be increasing with θ_d up to 1 when $\theta_d = \pi/2$. The Fresnel term extracted from the measured data starts decreasing after a certain angle (from 70° to 80°), and is equal to 0 for $\theta_d = \pi/2$ (see Figure V.6(a)). This is probably caused by the acquisition process: at grazing angles, the sensor is saturated by direct illumination from the light source². We discard data for $\theta_d > 70^\circ$.
- For some materials such as *gold-paint* or *nickel*, the Fresnel term does not behave as predicted by the Schlick approximation (see Figure V.6(b)). Instead of a plateau followed by a sharp increase, we observe linear slope, followed by an increase. To model these materials, we suggest the following generalization of Schlick approximation:

$$F(\theta_d) = F_0 + (1 - F_0)(1 - \cos \theta_d)^5 - F_1 \cos \theta_d \quad (\text{V.5})$$

For $F_1 = 0$, we get the original approximation. Figure V.7 shows how well our generalization fits to the measured Fresnel.

²F. Durand, personal communication.




Figure V.7: Our generalization to Schlick’s approximation of the Fresnel term fits well to all the materials in the MERL-MIT database.

V.3.5 Importance sampling

Efficient Monte-Carlo integration requires importance sampling. For this, we need the inverse of the integral of P_{22} (see Equation III.4). We easily have F_θ :

$$F_\theta = 1 - \frac{\Gamma\left(1 - p, \alpha + \frac{\tan^2 \theta}{\alpha}\right)}{\Gamma(1 - p, \alpha)}$$

but there is no closed form expression for the inverse for F_θ , so we can’t use it for exact importance sampling. As an approximation, we importance sample the GGX distribution [Walter 2007], and compute weights, w , to fit our model:

$$\begin{aligned} \theta_m &= \arctan\left(\frac{\alpha \sqrt{u_2}}{\sqrt{1 - u_2}}\right) \\ \phi_m &= 2\pi u_1 \\ \mathbf{m} &= (\cos \phi_m \sin \theta_m, \sin \phi_m \sin \theta_m, \cos \theta_m) \\ \mathbf{o} &= 2|\mathbf{i} \cdot \mathbf{m}| \mathbf{m} - \mathbf{i} \end{aligned}$$

where u_1 and u_2 are realizations of a uniform random variable in $[0, 1)$. Our BRDF model contains a diffuse and a specular component. We alternatively sample either component according to the following probabilities:

$$\begin{aligned} p_{\text{specular}} &= \frac{\rho_s}{\rho_d + \rho_s} \\ p_{\text{diffuse}} &= 1 - p_{\text{specular}} \end{aligned}$$

The importance sampling weights for our distribution is:

$$\begin{aligned} w_{\text{specular}} &= \frac{4\rho_s}{p_{\text{specular}}} F(\mathbf{i} \cdot \mathbf{h}) D(\theta_h) G(\mathbf{i}, \mathbf{o}) \cos^3 \theta_m \frac{(\alpha^2 + \tan^2 \theta)^2}{\alpha^2} \\ w_{\text{diffuse}} &= \frac{\rho_d}{p_{\text{diffuse}}} \end{aligned}$$

This expression accounts for the jacobian of sampling \mathbf{m} instead of \mathbf{o} [Walter 2007]. Figure V.8 shows the same material rendered with and without importance sampling.




Figure V.8: Uniform and importance sampling for *specular-violet-phenolic* with the same rendering time (6 mn).

V.4 Fitting algorithm

V.4.1 Fitting RGB channels separately

Fitting BRDF models with measured materials is usually done in a two-step approach [Ngan 2005]: first compute the diffuse and specular colors using a linear least square optimization, then compute the BRDF parameters (α , F_0) through a second optimization.

This approach places constraints on the second optimization process: the parameters must fit all RGB channels. We used a different approach: we compute the BRDF parameters separately for each channel, red, green and blue. Although we compute and store more data for each material, the fitting process is faster and more efficient, as we only fit real-valued functions, with less constraints for each fitting.

We also get wavelength dependent effects. The Fresnel term should change with the wavelength [Lazányi 2005] and we do find different values for F_0 and F_1 for the red, green and blue channels, for all materials in the MERL-MIT database (see Appendix A). Often, we get similar values for the α and p parameters for the micro-facet distribution.

We do not take any explicit step to conserve the color of the material, since we treat each channel independently. We still get BRDF models that match very well the color of the measured materials (see Figures 1, V.9, and V.10 and the Appendix A). We attribute this to the very good match between our distribution function and the data behavior.

V.4.2 Non-linear optimization

Given the measured data for a single channel, we find the optimal value for the parameters ($\rho_d, \rho_s, \alpha, p, F_0, F_1$) so that the BRDF model is as close as possible to the measured data.

We could compute the L_2 norm of the difference between the BRDF model and the measured BRDF over the whole hemisphere, as in [Ngan 2005]. But regular sampling in

θ_i and θ_o results in under-sampling near the specular peak, resulting in poor quality fits. Instead we fit in two steps:

1. We fit ρ_d and the probability distribution parameters (α, p) to the measured data for varying θ_h , with $\theta_d = 0$.
2. We fit ρ_s and the Fresnel term parameters (F_0, F_1) to the measured data for varying θ_d , with $\theta_h = 0$.

Each step is the fitting of a real-valued function of several parameters, which is simple and stable. Although we fit the parameters using only two slices of the measured reflectance, we still capture the full behavior of the BRDF.

We define the following quantities:

- $m(\mathbf{i}, \mathbf{o})$: the full measured reflectance, as extracted from the database,
- $m_h(\theta_h) = m(\mathbf{h}, \mathbf{h})$: the slice of the BRDF corresponding to $\theta_d = 0$,
- $m_d(\theta_d) = m(\mathbf{i}, \text{reflection}(\mathbf{i}))$: the slice of the BRDF corresponding to $\theta_h = 0$.

In a first step, we fit $m_h(\theta_h)$ with:

$$\rho_h(\theta_h) = \frac{\rho_d}{\pi} + \frac{F_r}{\pi} \frac{D(\theta_h)G_1^2(\theta_h)}{\cos^2(\theta_h)}$$

The function ρ_h depends on 4 parameters: ρ_d , F_r , α and p . Using Levenberg-Marquadt optimization [Lourakis 2004], we find the value of these parameters that minimizes the error function:

$$E_1(\rho_d, F_r, \alpha, p) = \sum_i w_i (m_h(\theta_h) - \rho_h(\theta_h))^2 \quad (\text{V.6})$$

At the end of this step, we have the value for 3 parameters: ρ_d , α and p . The last coefficient, F_r is the product of ρ_s and the Fresnel term for $\theta_d = 0$.

In the second step, we first extract the specular coefficient times the Fresnel term from $m_d(\theta_d)$ using the values ρ_d , α and p we just computed:

$$\rho_s F_D(\theta_d) = \left(m_d(\theta_d) - \frac{\rho_d}{\pi} \right) \frac{\pi \cos^2(\theta_d)}{D(0)G_1^2(\theta_d)}$$

We fit this function with our Fresnel term approximation (Equation V.5), using Levenberg-Marquadt optimization. The data gives us $\rho_s F_0$ and $\rho_s F_1$. Using $F_D(\pi/2) = 1$, we get ρ_s , F_0 and F_1 .

V.5 Results and Comparison

V.5.1 Visual comparison

Figures V.9 and V.10 shows a side-by-side comparison for *aventurine* and *colonial-maple-223* between images rendered using the measured reflectance from the MERL-MIT database

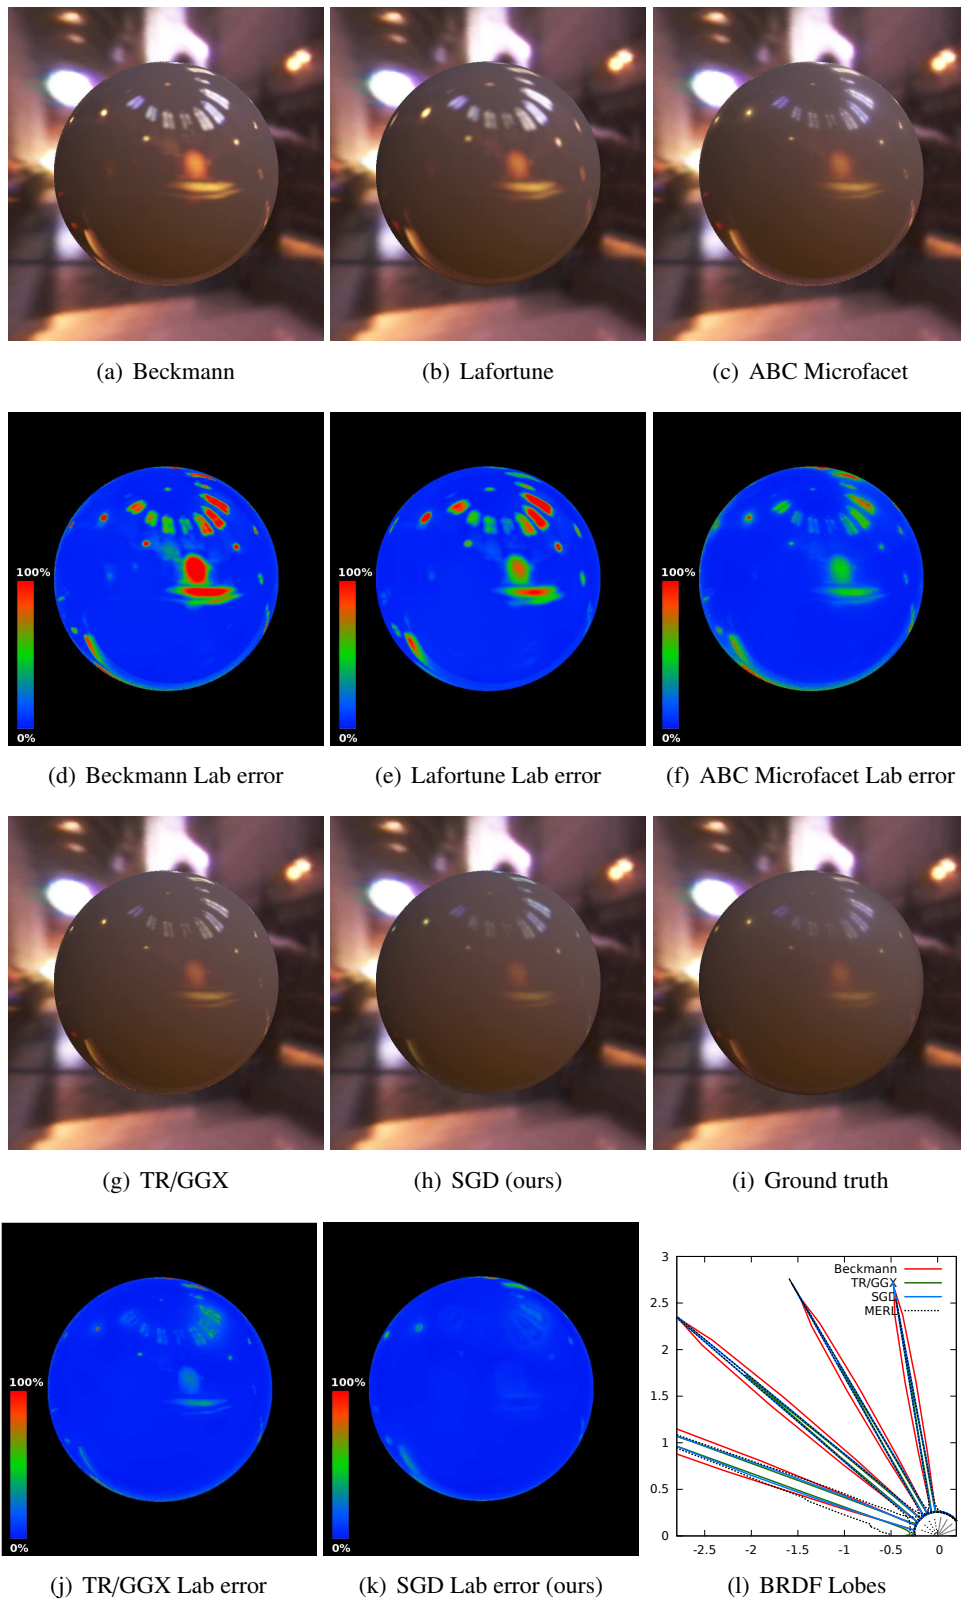


Figure V.9: Comparison between the Cook-Torrance model with Beckmann, TR/GGX and SGD distributions, Lafortune, ABC microfacet model [Löw 2012] and ground truth for *aventurine* from the MERL-MIT database (spelled *aventurnine* in [Ngan 2005]).

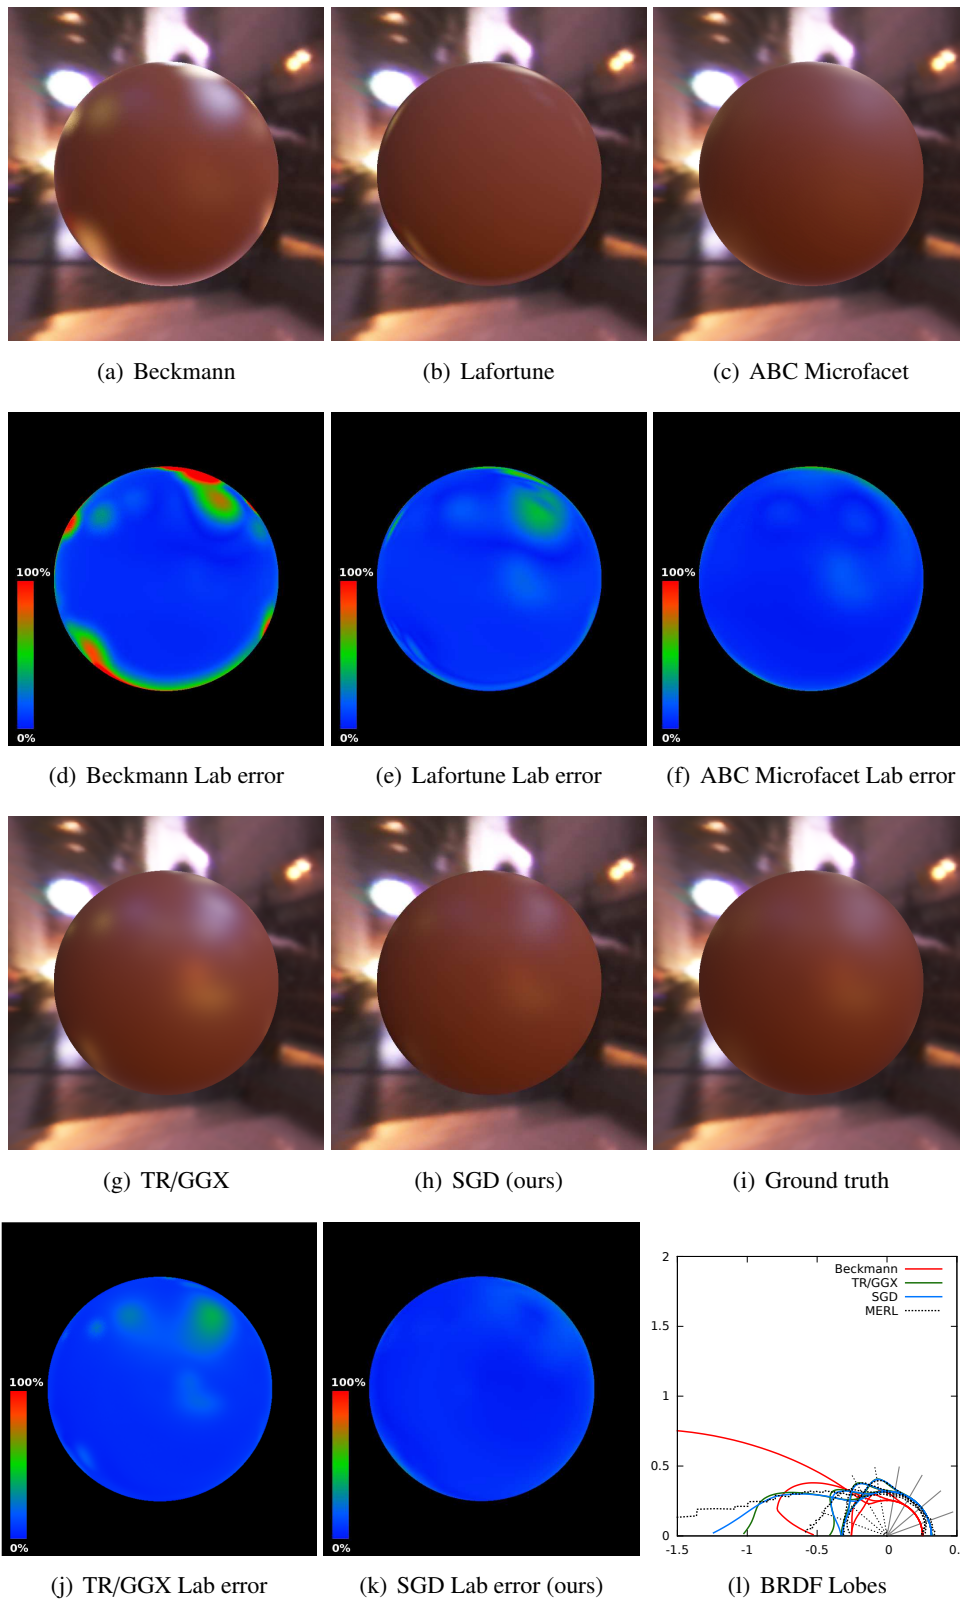


Figure V.10: Comparison between the Cook-Torrance model with Beckmann, TR/GGX and SGD distributions, Lafortune, ABC microfacet model [L w 2012] and ground truth for *colonial-maple-223* from the MERL-MIT database.

and several BRDF models: Lafortune, ABC microfacet model [Löw 2012], and Cook-Torrance with the Beckmann, TR/GGX and SGD distributions. For each model, we show the difference with the reference image, measured in Lab space. We chose the Lab color space because it is perceptually uniform. Qualitatively, our BRDF model provides the best approximation to the measured data.

In Appendix A, we provide a similar comparison for all 100 materials in the MERL-MIT database: reference image and Cook-Torrance BRDF with Beckmann and SGD distributions, along with differences in Lab space. For 13 materials in the database, we couldn't find the corresponding parameters in the supplemental of [Ngan 2005], so we only provide the comparison between the measured material and our distribution. Visually, the Cook-Torrance BRDF with the SGD distribution looks very close to the measured reflectance.

V.5.2 BRDF Lobes comparison

Figures V.9(I), V.10(I) and V.11, plot the BRDF lobes for *aventurine*, *colonial-maple-223*, *hematite*, *black-obsidian*, *nickel*, and *black-oxidized-steel*, with the Beckmann, GGX and SGD distributions compared to the measurements, for four incoming directions (10° , 30° , 50° , and 70°).

The SGD distribution closely fits the measured data, while the Beckmann distribution either misses the tail by trying to fit the peak of the lobe, or vice-versa. This is particularly evident for *black-obsidian* where the Beckmann distribution underestimates the peak. In the *hematite*, *aventurine*, *nickel*, *black-oxidized-steel*, and *colonial-maple-223* cases, which are ranging from highly-specular to glossy-specular and glossy-diffuse, the Beckmann distribution significantly overestimates the lobes. The GGX distribution performs better than Beckmann and SGD is very close to the actual shape of the lobes.

V.5.3 Quantitative Error Measure

For each material, we define the fitting error as the L_2 norm of the difference between measured data and the Cook-Torrance BRDF:

$$E = \left(\int (m(\mathbf{i}, \mathbf{o}) - \rho(\mathbf{i}, \mathbf{o}))^2 \cos \theta_i \cos \theta_o d\omega_i \right)^{\frac{1}{2}}$$

To ensure a fair comparison between bright and dark materials, we normalize this error by dividing it with the maximum albedo of the BRDF, as in [Ngan 2005] (the Appendix A has the un-normalized error value). Figure V.12 compares the fitting errors for the Cook-Torrance BRDF using the SGD and Beckmann distributions (parameters for Beckmann from [Ngan 2005]). We sorted the materials by increasing α roughness, going from highly specular to diffuse. For all but 3 of the materials (*black-fabric*, *yellow-matte-plastic*, and *white-paint*), the SGD distribution provides the best approximation. For many materials, we improve the accuracy by one order of magnitude, and for some, such as *hematite*, by two orders of magnitude.

For the materials where we do not improve over Beckmann distribution, we think that the problem comes from the fitting algorithm rather than the SGD distribution. Manually

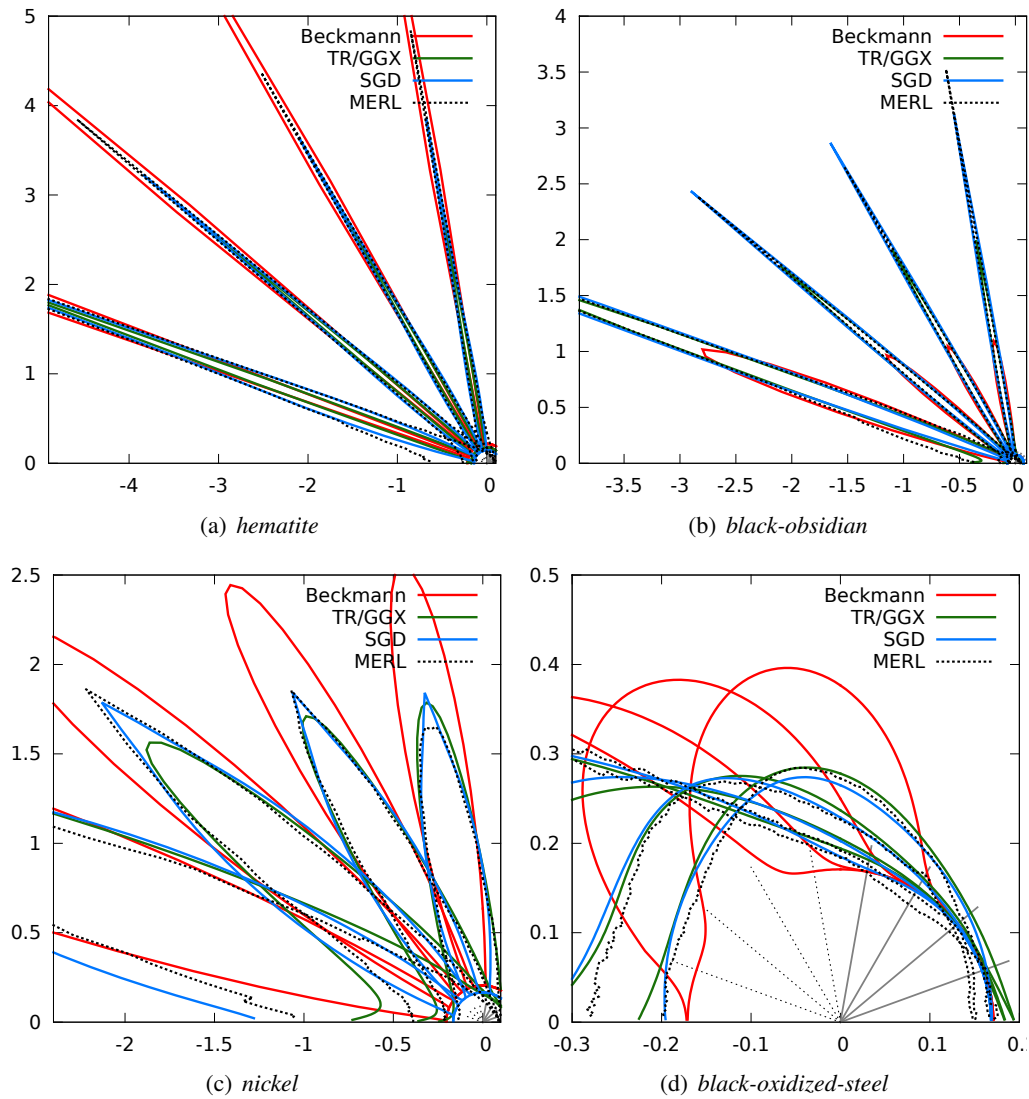


Figure V.11: BRDF lobes (cubic root applied) with Beckmann, GGX and SGD distributions for *hematite*, *black-obsidian*, *nickel*, and *black-oxidized-steel*. The SGD distribution provides a much closer fit to measured reflectance than the Beckmann and GGX distribution.




Figure V.12: Normalized average fitting errors (logarithmic scale) of the Cook-Torrance model with Beckmann — with Ngan parameters and also with our fitting —, TR/GGX and SGD distributions. In most cases, SGD provides a much better fit. Only 85 materials are listed — sorted by increasing alpha —, since Ngan et al. has provided parameters for 85 materials only.

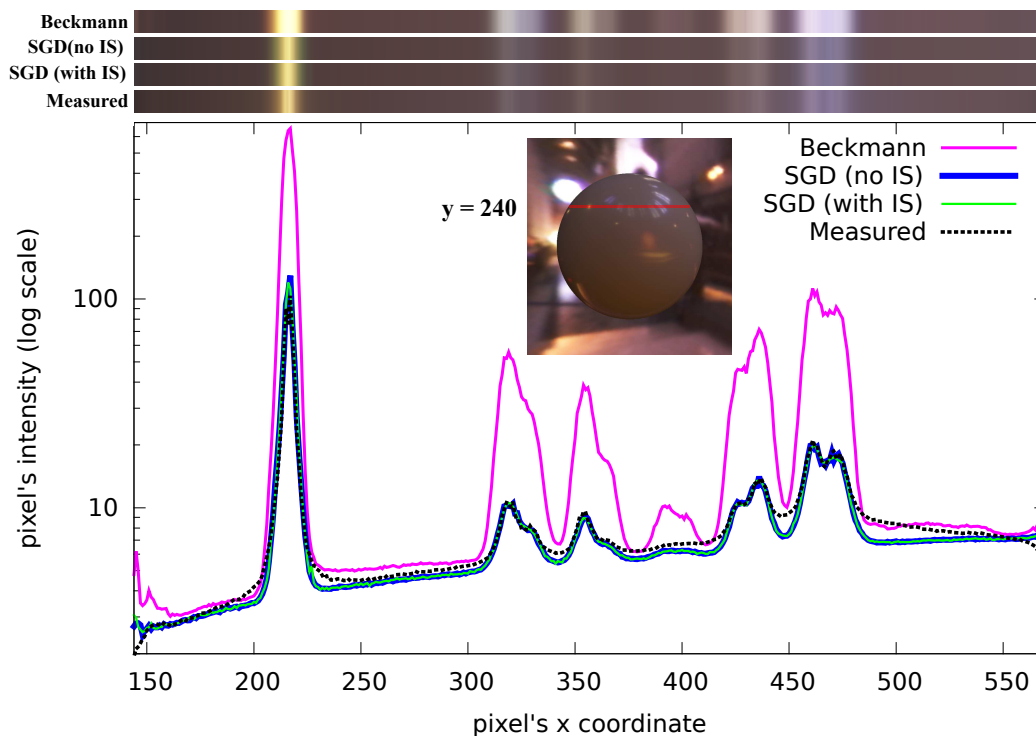


Figure V.13: Pixel-by-pixel comparison between the reference, Beckmann and SGD (with and without importance sampling) for *aventurine* from MERL-MIT database.

setting the parameters for these materials improves the accuracy (all the parameters we use here are computed in a single run of our optimization process, without any tweaking for specific materials).

To evaluate our fitting method, we used it to compute parameters for the Beckmann and TR/GGX distributions, and reported the error for each material in Figure V.12. For almost all materials, the error is higher than with Beckmann distribution and parameters from [Ngan 2005]. Our fitting method works better if the shape of the distribution matches the measured data.

Figure V.13 provides a pixel-by-pixel comparison for a horizontal scanline of rendered image between: measured reflectance, Cook-Torrance with Beckmann distribution and SGD distribution, with and without importance sampling. The Beckmann distribution overestimates specular lobes; SGD remains very close to the reference. The figure also validates that using GGX importance sampling for SGD distribution gives the same results, but reduces the noise.

V.5.4 Influence of D , G , and F parameters in our BRDF model

The Cook-Torrance micro-facet model contains three main components, namely:

- Normal distribution function (D)

- Shadowing and masking function (G)
- Fresnel term (F)

We are going to study the influence of each of these parameters in detail for a diffuse, a glossy, and a specular material as representatives of all the materials in the MERL-MIT database. Figures V.14, V.15, and V.16 shows the result of our experiments with *beige-fabric* (diffuse), *natural-209* (glossy), and *specular-violet-phenolic* (specular). Changing the α parameter of the distribution significantly changes the appearance of the materials. Typically, smaller α values are associated with highly specular materials and the larger values are associated with diffuse ones. For example, the average α and p values for these three materials are as following:

- *beige-fabric*: $\alpha = 0.21, p = 0.1$
- *natural-209*: $\alpha = 0.017, p = 0.088$
- *specular-violet-phenolic*: $\alpha = 0.0006, p = 0.62$

In contrast to the distribution parameters, keeping the shadowing and masking function constant and equal to unity does not affect the visual appearance of the materials noticeably. For most materials, specially the highly specular ones, the value of the G function is almost always equal to one. Cancelling the Fresnel variation by setting the Fresnel term to be constant and equal to $F(\theta_d = 0)$ does not make a very noticeable difference either. The error gradually increases as the specularity increases.

Based on the results of this experiment, we conclude that the normal distribution function (D) is the most important component of the micro-facet model. The NDF defines the portion of the micro-facets aligned with the half-angle vector causing specular reflection.

V.5.5 Fitting RGB channels altogether for D parameters

The results from the fitting algorithm showed that we get similar values for distribution (D) parameters — α and p — in some cases. The Fresnel coefficients are expected to be different, since Fresnel is a wavelength dependent effect. Therefore, as an experiment, we modify the two-step fitting algorithm, described in V.4, to fit the distribution parameters (α , and p) to the RGB channels at once — resulting in a single α , and p instead of separate values for each channel. In the first step, we fit $\rho_d^R, \rho_d^G, \rho_d^B, \alpha$, and p to the measured data considering all the channels at once. The second step is like before, fitting ρ_s, F_0 , and F_1 for each RGB channel separately — with identical α , and p values for all channels. We compute the error and compare it with the old fitting method. See Figure V.17 for the error curves. The materials are sorted by the error in the fitting method explained in V.4. The results show that the fitting is more unstable and the error is generally higher than the case we fit each channel independently. Specially, the algorithm failed to fit *gold-metallic-paint2*, *alum-bronze*, *two-layer-silver*, *violet-phenolic*, and *white-acrylic* — and the error is significantly high. The Figure V.17 also sorts the 100 materials in the MERL-MIT database in order of fitting errors using our algorithm with our BRDF model.

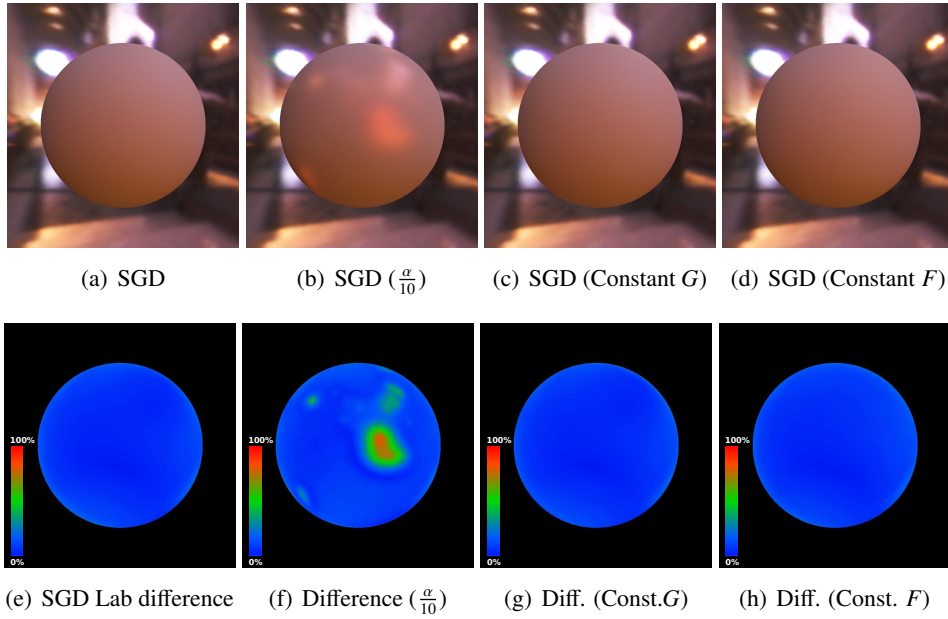


Figure V.14: Top row: The influence of D , G , and F in our model for *beige-fabric* from the MERL-MIT database as a diffuse material. Bottom row: The associated Lab difference figures compared to the ground truth.

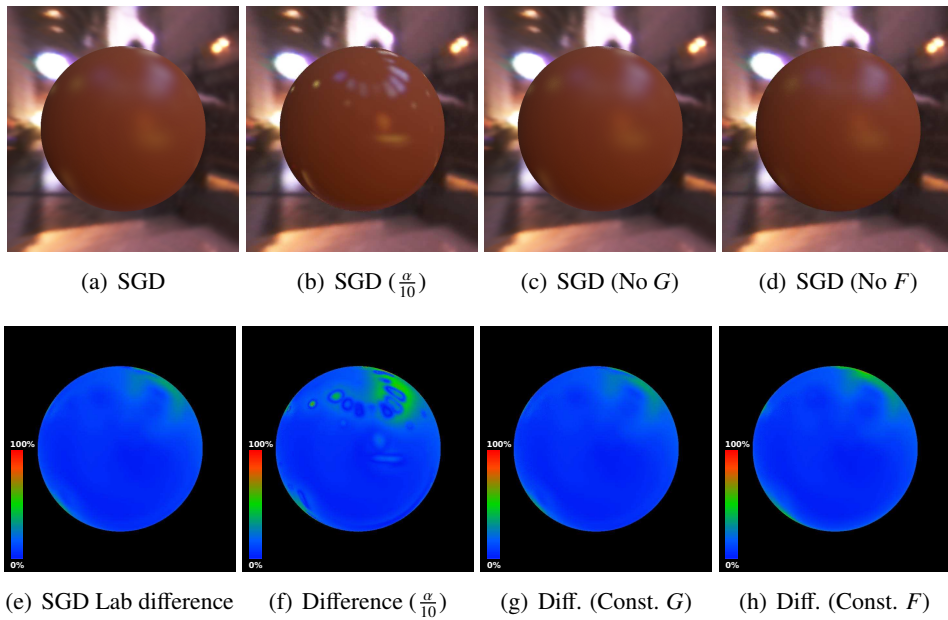


Figure V.15: Top row: The influence of D , G , and F in our model for *natural-209* from the MERL-MIT database as a glossy material. Bottom row: The associated Lab difference figures compared to the ground truth.




Figure V.16: Top row: The influence of D , G , and F in our model for *specular-violet-phenolic* from the MERL-MIT database as a specular material. Bottom row: The associated Lab difference figures compared to the ground truth.

V.5.6 Timings

Since our algorithm only uses two slices of the BRDF measurement for fitting the parameters, it is extremely fast: 2.5 minutes on average to fit all the parameters for a single material on a single core 2.57 GHz Intel CPU. For 65 % of the materials, the fitting process took less than 1 minute.

We also measured the cost of evaluating our reflectance model inside a ray-tracer, using valgrind profiling tool. A single evaluation takes 3044 CPU cycles, compared to 1416 with Beckmann and 551 with Lafortune. Our model is more expensive because we estimate the BRDF independently for each color channel.

V.5.7 Limitations

For some materials in the MERL-MIT database, even with the SGD distribution, we see evidence of a multi-lobe behavior. This is quite obvious for *two-layer-gold* and *two-layer-silver* (see Figure V.18), and it appears (but it's less obvious) for *alum-bronze*. For these materials, adding a second lobe would improve the accuracy. However, the number of materials with multi-lobe behavior is 3 out of 100.

Our fitting method is based on only two slices of the material and on the assumption that the material follows the Cook-Torrance model. It will miss behavior that are outside of the model, for example *color-changing-paint*.

The *white-paint* material illustrates a failure case of our fitting algorithm (see Figure V.19). The error with SGD is larger than with Beckmann. The difference is clearly




Figure V.17: Normalized average fitting errors (logarithmic scale) of the Cook-Torrance model with the SGD distribution with channels fitted separately — different α and p for each channel — (red curve) and altogether — the same α and p for all channels — (green curve). Fitting RGB channels separately provides a better fit. All 100 materials in the MERL-MIT database are listed here, sorted by increasing error.

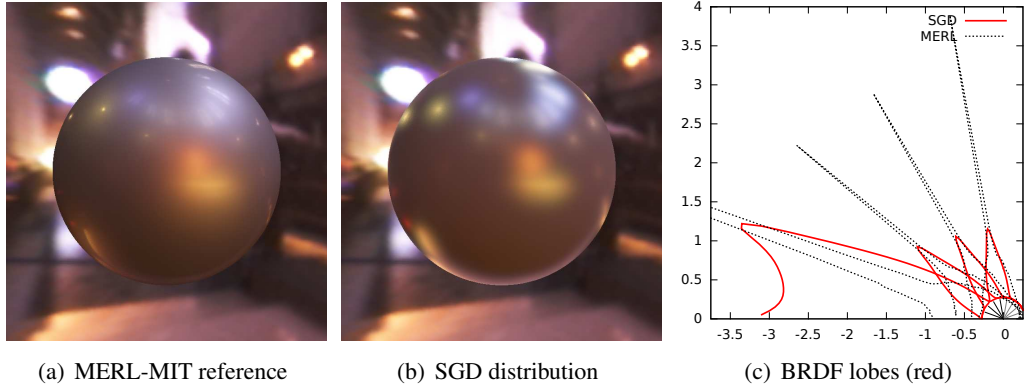


Figure V.18: An example of a two-layer material, *two-layer-silver*, that requires two lobes for accurate fitting.

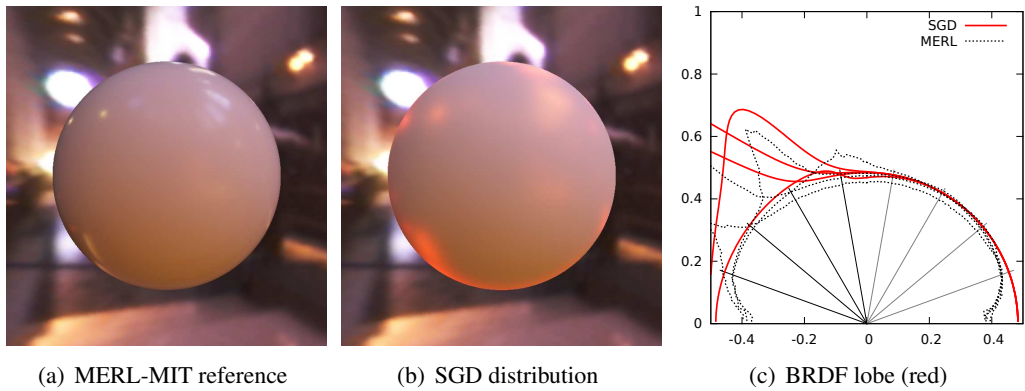


Figure V.19: *white-paint*, a failure case for our fitting algorithm.

visible in the rendered pictures (Figure V.19(b)). A possible explanation is visible in the lobes (Figure V.19(c)): the Fresnel term is almost null at normal incidence, and our fitting algorithm does not pick the right shape for the specular lobe.

An interesting observation here is that the normalized average (numerical) error value is different from the visual (qualitative) error. For example, many materials such as *steel* have higher average error than *white-paint* – see Figure V.17 —, but they visually look closer to the ground truth than *white-paint*. This is because the error for *steel*, as an example, is distributed across the image, while the error for *white-paint* is concentrated in some sub-parts of the image — becoming more visible. Please refer Appendix A for rendered images of *steel*.

V.6 Conclusion

We have presented the SGD micro-facet probability distribution function for the Cook-Torrance BRDF. This distribution provides an accurate fit for most measured materials in the MERL-MIT database. We derived the shadowing and masking function from the SGD

distribution. We proposed a generalized Fresnel approximation that fits better to the measurements. We have also presented a new fitting algorithm, where we fit each color channel independently (instead of pre-computing the diffuse and specular colors), and we fit using only two slices of the BRDF. This fitting algorithm is very efficient when combined with the SGD distribution, taking 2.5 mn to compute the parameters for a measured distribution.

In the Appendix A, we provide parameters for the SGD distribution for all materials in the MERL-MIT database. We hope that this will help researchers and software developers in using the SGD distribution.

In future work, we want to continue working on micro-facet distributions, finding a distribution that provides the same accurate fit to measured data while being easier to compute and integrating into nicer functions. We also want to target material acquisition: we describe an acquired material accurately with only 18 coefficients. Can we use this property for faster acquisition of new materials?

Conclusion

In this thesis, we focused on presenting new ways to represent and render reflectance of arbitrary isotropic opaque materials.

We introduced a real-time framework for frequency analysis of light transport. The estimation of bandwidth and variance of the shading integrand enabled us to adaptively sample both the angular and the image space while shading. We presented a single pass hierarchical shading and up-sampling algorithm to accelerate the shading process. This made it possible to render tabulated BRDFs, evaluation of which on the GPU is extremely expensive, at interactive frame rates. We also proposed a pre-convolved shading technique to render materials in real-time. With the help of frequency analysis, we were able to implement multi-sample anti-aliasing in a deferred shading pipeline. This technique could be employed for applications in which accurate shading of materials, dynamic geometry, and interactivity are essentially required.

Then, we introduced a new distribution function for the Cook-Torrance micro-facet BRDF, namely the Shifted Gamma Distribution (SGD). We derived the shadowing and masking function from the distribution, and generalized the Schlick's Fresnel approximation to fit to the measurements with higher accuracy. We advised a solution for importance sampling using this shading model. We proposed a straightforward two-step fitting algorithm that fit each RGB channel separately and accurately for a wide range of materials from diffuse to glossy and highly specular in the MERL-MIT database.

We developed the above mentioned solutions to advance the state of the art towards finding a unified and straightforward approach to appearance modeling. Although we focused on representation of reflectance of arbitrary isotropic opaque materials, many of our contributions can be extended to more general materials — e.g. anisotropic and non-opaque materials. The frequency analysis of light transport theory can be further developed to predict the sampling rate for rendering anisotropic reflectance or refraction for example. We can account for visibility and multiple reflections to extend our bandwidth estimation technique to global illumination. The SGD distribution can be used as the reflection term of a more general, e.g. non-opaque, material.

One important aspect of digital appearance modeling is appearance capture. Although we have not explored this area in this thesis, our new reflectance distribution suggests that using much less measurement samples we can infer appearance of an isotropic opaque material with high fidelity.

Conclusion (français)

Dans cette thèse, nous nous sommes concentrés sur la présentation de nouvelles façons de représenter et de rendre la réflectance de matériaux opaques isotropes.

Nous avons introduit un cadre temps réel pour l'analyse de la fréquence de l'opérateur de transport de lumière. L'estimation du débit et de la variance de l'intégrande nous a permis d'échantillonner de manière adaptative à la fois en angle et en espace. Nous avons présenté un algorithme de rendu hiérarchique en une passe, basé sur un sur-échantillonnage. Cela a permis de rendre des BRDFs tabulées, dont l'évaluation sur le GPU est extrêmement coûteux, à des cadences interactives. Nous avons également proposé une technique de pré-convolution pour rendre les matériaux en temps réel. Avec l'aide de l'analyse fréquentielle, nous avons permis l'utilisation de l'anti-aliasing dans un pipeline en deferred shading. Cette technique pourrait être utilisée pour les applications où l'éclairage le rendu précis et interactif des matériaux, avec une géométrie dynamique sont nécessaires.

Ensuite, nous avons introduit une nouvelle fonction de distribution des micro-facettes pour le modèle de BRDF de Cook-Torrance, que nous avons appelée la Shifted Gamma Distribution (SGD). Nous avons calculé la fonction d'ombrage et de masquage de la distribution, et généralisé l'approximation de Schlick pour le terme de Fresnel, pour s'adapter aux mesures effectuées avec une plus grande précision. Nous avons proposé une solution pour l'échantillonnage adaptatif. Nous avons proposé un algorithme de regretion simple en deux étapes sur chaque canal RGB séparément, qui l'est révélé très précis pour un large éventail de matériaux diffus, brillants et hautement spéculaires.

Nous avons développé les solutions mentionnées ci-dessus pour faire avancer l'état de l'art en vue de trouver une approche unifiée et simple à la modélisation de l'apparence. Bien que nous nous soyons concentrés sur la représentation de la réflectance de matériaux isotropes opaques, beaucoup de nos contributions peut être étendues à d'autres matériaux, par exemple anisotropes et non opaques. la théorie de l'analyse fréquentielle de le transport de lumière peut être développée afin de prédire le taux d'échantillonnage pour le rendu de réflexion ou de réfraction anisotrope par exemple. Nous pouvons rendre compte de la visibilité et de réflexions multiples et étendre notre technique d'estimation de débit à l'illumination globale. La distribution SGD peut être utilisée comme fonction de réflectance pour des matériaux plus généraux, par exemple non opaques.

Un aspect important de la modélisation numérique est la capture de l'apparence. Même si nous n'avons pas exploré cette voie dans cette thèse, notre nouveau modèle de réflectance suggère qu'il est possible de mesurer l'apparence de beaucoup de matériaux en utilisant peu d'échantillons.

Appendix A: SGD fitting results

Due to the large size of this appendix, for more convenience, it is prepared as a separate file.

Bibliography

- [Ashikhmin 2000] Michael Ashikhmin and Peter Shirley. *An Anisotropic Phong BRDF Model*. Journal of Graphics Tools, vol. 5, pages 25–32, 2000. (Cited on pages 33 and 37.)
- [Ashikhmin 2007] Michael Ashikhmin and Simon Premoze. *Distribution-based BRDFs*. University of Utah, <http://www.cs.utah.edu/~premoze/dbrdf/>, 2007. (Cited on page 68.)
- [Barsi 2005] Attila Barsi, László Szirmay-Kalos and Gábor Szijártó. *Stochastic glossy global illumination on the GPU*. In Proceedings of the 21st Spring Conference on Computer Graphics, SCCG '05, pages 187–193, New York, NY, USA, 2005. ACM. (Cited on page 15.)
- [Bavoil 2008] Louis Bavoil and Kevin Myers. *Order independent transparency with dual depth peeling*. Rapport technique, NVIDIA, 2008. http://developer.download.nvidia.com/SDK/10.5/opengl/src/dual_depth_peeling/doc/DualDepthPeeling.pdf. (Cited on page 15.)
- [Bekaert 1999] Philippe Bekaert. *Hierarchical and stochastic algorithms for radiosity*. PhD thesis, Department of Computer Science, K. U. Leuven, December 1999. (Cited on page 13.)
- [Ben-Ezra 2008] Moshe Ben-Ezra, Jiaping Wang, Bennett Wilburn, Xiaoyang Li and Le Ma. *An LED-only BRDF measurement device*. In Computer Vision and Pattern Recognition, 2008. (Cited on page 26.)
- [Blinn 1977] James F. Blinn. *Models of light reflection for computer synthesized pictures*. ACM SIGGRAPH Computer Graphics, vol. 11, no. 2, pages 192–198, July 1977. (Cited on pages 31 and 37.)
- [Born 1998] M. Born and E. Wolf. Principles of Optics. Cambridge University Press, 1998, 1998. (Cited on page 11.)
- [Bourlier 2002] C. Bourlier, G. Berginc and J. Saillard. *One- and two-dimensional shadowing functions for any height and slope stationary uncorrelated surface in the monostatic and bistatic configurations*. IEEE Transactions on Antennas and Propagation, vol. 50, no. 3, pages 312–324, March 2002. (Cited on pages 34, 35 and 36.)
- [Brown 1980] G. Brown. *Shadowing by non-Gaussian random surfaces*. IEEE Transactions on Antennas and Propagation, vol. 28, no. 6, pages 788–790, November 1980. (Cited on pages 34, 35 and 36.)

- [Camahort 1998] Emilio Camahort, Apostolos Lerios and Don Fussell. *Uniformly Sampled Light Fields*. Rapport technique, Austin, TX, USA, 1998. www.cs.utexas.edu/ftp/techreports/tr98-09.pdf. (Cited on page 18.)
- [Chai 2000] Jin-Xiang Chai, Xin Tong, Shing-Chow Chan and Heung-Yeung Shum. *Plenoptic sampling*. In Proceedings of the 27th annual conference on Computer graphics and interactive techniques, SIGGRAPH '00, pages 307–318, New York, NY, USA, 2000. ACM Press/Addison-Wesley Publishing Co. (Cited on page 18.)
- [Claustres 2007] Luc Claustres, Loïc Barthe and Mathias Paulin. *Wavelet encoding of BRDFs for real-time rendering*. In Proceedings of Graphics Interface 2007, GI '07, pages 169–176, New York, NY, USA, 2007. ACM. (Cited on page 29.)
- [Cook 1981] Robert L. Cook and Kenneth E. Torrance. *A reflectance model for computer graphics*. ACM SIGGRAPH Computer Graphics, vol. 15, no. 3, pages 307–316, August 1981. (Cited on pages 33, 36 and 37.)
- [Cook 1982] R. L. Cook and Kenneth E. Torrance. *A Reflectance Model for Computer Graphics*. ACM Transactions on Graphics, vol. 1, no. 1, pages 7–24, 1982. (Cited on page 35.)
- [cor 2001] *Cornell Reflectance Data*. <http://www.graphics.cornell.edu/online/measurements/reflectance/>, 2001. (Cited on pages 26 and 28.)
- [cor 2005] *Cornell Light Measurement Laboratory*. <http://www.graphics.cornell.edu/research/measure/>, 2005. (Cited on page 25.)
- [cur 1999] *Columbia-Utrecht Reflectance and Texture Database*. <http://www.cs.columbia.edu/CAVE/software/curet/>, 1999. (Cited on page 26.)
- [Dachsbacher 2005] Carsten Dachsbacher and Marc Stamminger. *Reflective shadow maps*. In Proceedings of the 2005 symposium on Interactive 3D graphics and games, I3D '05, pages 203–231, New York, NY, USA, 2005. ACM. (Cited on page 17.)
- [Dana 2004] Kristin J. Dana and Jing Wang. *Device for convenient measurement of spatially varying bidirectional reflectance*. Journal of the Optical Society of America, vol. 21, no. 1, pages 1–12, Jan 2004. (Cited on page 25.)
- [Debevec 2000] Paul Debevec, Tim Hawkins, Chris Tchou, Haarm-Pieter Duiker, Wesley Sarokin and Mark Sagar. *Acquiring the reflectance field of a human face*. In Proceedings of the 27th annual conference on Computer graphics and interactive techniques, SIGGRAPH '00, pages 145–156, New York, NY, USA, 2000. ACM Press/Addison-Wesley Publishing Co. (Cited on page 26.)
- [Deering 1988] Michael Deering, Stephanie Winner, Bic Schediwy, Chris Duffy and Neil Hunt. *The triangle processor and normal vector shader: a VLSI system for high performance graphics*. SIGGRAPH Computer Graphics, vol. 22, no. 4, pages 21–30, June 1988. (Cited on page 15.)

- [Dong 2010] Yue Dong, Jiaping Wang, Xin Tong, John Snyder, Yanxiang Lan, Moshe Ben-Ezra and Baining Guo. *Manifold bootstrapping for SVBRDF capture*. ACM Transactions on Graphics, vol. 29, no. 4, pages 98:1–98:10, July 2010. (Cited on page 26.)
- [Dorsey 2008] Julie Dorsey, Holly Rushmeier and François Sillion. Digital modeling of material appearance. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2008. (Cited on page 21.)
- [Dür 2006] Arne Dür. *An Improved Normalization for the Ward Reflectance Model*. Journal of Graphics Tools, vol. 11, no. 1, pages 51–59, 2006. (Cited on pages 32 and 37.)
- [Durand 2005] Frédo Durand, Nicolas Holzschuch, Cyril Soler, Eric Chan and François X. Sillion. *A frequency analysis of light transport*. ACM Transactions on Graphics, vol. 24, no. 3, pages 1115–1126, 2005. (Cited on pages 18, 19, 20, 42, 43, 45, 46 and 49.)
- [Egan 2009] Kevin Egan, Yu-Ting Tseng, Nicolas Holzschuch, Frédo Durand and Ravi Ramamoorthi. *Frequency analysis and sheared reconstruction for rendering motion blur*. ACM Transactions on Graphics, vol. 28, no. 3, pages 93:1–93:13, July 2009. (Cited on page 21.)
- [Everitt 2001] Cass Everitt. *Interactive Order-Independent Transparency*. Rapport technique, NVIDIA, 2001. http://developer.download.nvidia.com/SDK/10/opengl/src/dual_depth_peeling/doc/DualDepthPeeling.pdf. (Cited on page 15.)
- [Fatahalian 2010] Kayvon Fatahalian, Solomon Boulos, James Hegarty, Kurt Akeley, William R. Mark, Henry Moreton and Pat Hanrahan. *Reducing shading on GPUs using quad-fragment merging*. ACM Transactions on Graphics, vol. 29, no. 4, pages 67:1–67:8, 2010. (Cited on page 15.)
- [Feynman 1988] Richard P. Feynman. QED: The Strange Theory of Light and Matter. Princeton University Press, October 1988. (Cited on page 10.)
- [Fodor 2002] Imola Fodor. *A Survey of Dimension Reduction Techniques*. <http://www.llnl.gov/CASC/sapphire/pubs/148494.pdf>, 2002. (Cited on page 29.)
- [Gautron 2008] Pascal Gautron, Jaroslav Křivánek, Kadi Bouatouch and Sumanta Patanaik. *Radiance cache splatting: a GPU-friendly global illumination algorithm*. In ACM SIGGRAPH 2008 classes, SIGGRAPH '08, pages 78:1–78:10, New York, NY, USA, 2008. ACM. (Cited on page 15.)
- [Geisler-Moroder 2010] David Geisler-Moroder and Arne Dür. *A New Ward BRDF Model with Bounded Albedo*. Computer Graphics Forum, vol. 29, no. 4, pages 1391–1398, 2010. (Cited on page 32.)

- [Ghosh 2010a] Abhijeet Ghosh, Tongbo Chen, Pieter Peers, Cyrus A. Wilson and Paul Debevec. *Circularly polarized spherical illumination reflectometry*. ACM Transactions on Graphics, vol. 29, no. 6, pages 162:1–162:12, December 2010. (Cited on page 26.)
- [Ghosh 2010b] Abhijeet Ghosh, Wolfgang Heidrich, Shruthi Achutha and Matthew O’Toole. *A Basis Illumination Approach to BRDF Measurement*. Int. J. Comput. Vision, vol. 90, no. 2, pages 183–197, November 2010. (Cited on page 26.)
- [Glassner 1994] Andrew S. Glassner. Principles of digital image synthesis. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1994. (Cited on page 12.)
- [Gortler 1996] Steven J. Gortler, Radek Grzeszczuk, Richard Szeliski and Michael F. Cohen. *The lumigraph*. In Proceedings of the 23rd annual conference on Computer graphics and interactive techniques, SIGGRAPH ’96, pages 43–54, New York, NY, USA, 1996. ACM. (Cited on page 18.)
- [Hammersley 1965] J. M. Hammersley and D. C. Handscomb. Monte Carlo Methods (Methuen’s Monographs on Applied Probability & Statistics). Methuen, January 1965. (Cited on page 13.)
- [He 1991] Xiao D. He, Kenneth E. Torrance, François X. Sillion and Donald P. Greenberg. *A comprehensive physical model for light reflection*. ACM SIGGRAPH Computer Graphics, vol. 25, no. 4, pages 175–186, July 1991. (Cited on page 37.)
- [Hullin 2010] Matthias B. Hullin, Johannes Hanika, Boris Ajdin, Hans-Peter Seidel, Jan Kautz and Hendrik P. A. Lensch. *Acquisition and analysis of bispectral bidirectional reflectance and reradiation distribution functions*. ACM Transactions on Graphics, vol. 29, no. 4, pages 97:1–97:7, July 2010. (Cited on page 26.)
- [Jensen 1996] Henrik Wann Jensen. *Global illumination using photon maps*. In Eurographics Workshop on Rendering, pages 21–30, London, UK, UK, 1996. Springer-Verlag. (Cited on page 13.)
- [Kajiya 1986] James T. Kajiya. *The rendering equation*. SIGGRAPH Computer Graphics, vol. 20, no. 4, pages 143–150, August 1986. (Cited on pages 12 and 13.)
- [Kalos 1986] Malvin H. Kalos and Paula A. Whitlock. Monte Carlo methods. J. Wiley & Sons, 1st édition, October 1986. (Cited on page 13.)
- [Kautz 1999] Jan Kautz and Michael D. McCool. *Interactive rendering with arbitrary BRDFs using separable approximations*. In ACM SIGGRAPH 99 Conference abstracts and applications, SIGGRAPH ’99, pages 253–, New York, NY, USA, 1999. ACM. (Cited on page 29.)
- [Kautz 2000] Jan Kautz and Michael D. McCool. *Approximation of Glossy Reflection with Prefiltered Environment Maps*. In Proceedings Graphics Interface 2000, pages 119–126, May 2000. (Cited on page 53.)

- [Koenderink 1996] Jan J. Koenderink, Andrea J. Van Doorn, Kristin J. Dana and Shree Nayar. *Bidirectional reflection distribution function expressed in terms of surface scattering modes*. International Journal of Computer Vision, vol. 31, pages 28–39, 1996. (Cited on page 29.)
- [Lafortune 1996] Eric P. Lafortune, Eric P. Lafortune, Eric P. Lafortune, Yves D. Willems and Yves D. Willems. *Rendering Participating Media with Bidirectional Path Tracing*. In Eurographics Workshop on Rendering, pages 91–100. Springer-Verlag/Wien, 1996. (Cited on page 13.)
- [Lafortune 1997] Eric P. F. Lafortune, Sing-Choong Foo, Kenneth E. Torrance and Donald P. Greenberg. *Non-linear approximation of reflectance functions*. In Proceedings of the 24th annual conference on Computer graphics and interactive techniques, SIGGRAPH '97, pages 117–126, New York, NY, USA, 1997. ACM Press/Addison-Wesley Publishing Co. (Cited on pages 32, 33 and 37.)
- [Lalonde 1997] Paul Lalonde and Alain Fournier. *Filtered Local Shading in the Wavelet Domain*. In Eurographics Workshop on Rendering, pages 163–174, London, UK, UK, 1997. Springer-Verlag. (Cited on page 29.)
- [Latta 2002] Lutz Latta and Andreas Kolb. *Homomorphic factorization of BRDF-based lighting computation*. In Proceedings of the 29th annual conference on Computer graphics and interactive techniques, SIGGRAPH '02, pages 509–516, New York, NY, USA, 2002. ACM. (Cited on page 29.)
- [Lawrence 2005] Jason Lawrence, Szymon Rusinkiewicz and Ravi Ramamoorthi. *Adaptive Numerical Cumulative Distribution Functions for Efficient Importance Sampling*. In Eurographics Symposium on Rendering, June 2005. (Cited on page 27.)
- [Lazányi 2005] István Lazányi and László Szirmay-Kalos. *Fresnel Term Approximations for Metals*. In Winter School of Computer Graphics (WSCG), Short papers, 2005. (Cited on pages 36 and 73.)
- [Levoy 1996] Marc Levoy and Pat Hanrahan. *Light field rendering*. In Proceedings of the 23rd annual conference on Computer graphics and interactive techniques, SIGGRAPH '96, pages 31–42, New York, NY, USA, 1996. ACM. (Cited on page 18.)
- [Lewis 1994] Robert R. Lewis. *Making Shaders More Physically Plausible*. Computer Graphics Forum, vol. 13, no. 2, pages 109–120, 1994. (Cited on page 31.)
- [Li 2005] Hongsong Li, Sing Choong Foo, Kenneth E. Torrance and Stephen H. Westin. *Automated three-axis gonioreflectometer for computer graphics applications*. In Optical Engineering, page 2006, 2005. (Cited on page 25.)
- [Lourakis 2004] M. I. A. Lourakis. *levmar: Levenberg-Marquardt nonlinear least squares algorithms in C/C++*. <http://www.ics.forth.gr/~lourakis/levmar/>, Jul. 2004. (Cited on pages 70 and 74.)

- [Löw 2012] Joakim Löw, Joel Kronander, Anders Ynnerman and Jonas Unger. *BRDF models for accurate and efficient rendering of glossy surfaces*. ACM Transactions on Graphics, vol. 31, no. 1, pages 9:1–9:14, February 2012. (Cited on pages 35, 75, 76 and 77.)
- [M. Bagher 2012a] Mahdi M. Bagher, Cyril Soler and Nicolas Holzschuch. *Accurate fitting of measured reflectances using a Shifted Gamma micro-facet distribution*. Computer Graphics Forum, vol. 31, no. 4, June 2012. (Cited on pages 3, 7 and 65.)
- [M. Bagher 2012b] Mahdi M. Bagher, Cyril Soler, Kartic Subr, Laurent Belcour and Nicolas Holzschuch. *Interactive rendering of acquired materials on dynamic geometry using bandwidth prediction*. In Michael Garland and Rui Wang, editors, ACM Siggraph Symposium on Interactive 3D Graphics and Games (I3D), pages 127–134, Costa Mesa, United States, March 2012. ACM Siggraph, ACM. (Cited on pages 3, 7 and 41.)
- [M. Bagher 2013] Mahdi M. Bagher, Cyril Soler, Kartic Subr, Laurent Belcour and Nicolas Holzschuch. *Interactive rendering of acquired materials on dynamic geometry using frequency analysis*. IEEE Transactions on Visualization and Computer Graphics, 2013. Invited paper. (Cited on pages 3 and 7.)
- [Mallat 2008] Stphane Mallat. *A wavelet tour of signal processing, third edition: The sparse way*. Academic Press, 3rd édition, 2008. (Cited on page 46.)
- [Mammen 1989] Abraham Mammen. *Transparency and Antialiasing Algorithms Implemented with the Virtual Pixel Maps Technique*. IEEE Computer Graphics and Applications, vol. 9, no. 4, pages 43–55, July 1989. (Cited on page 15.)
- [Marschner 1999] Steve R. Marschner, Stephen H. Westin, Eric P. F. Lafortune, Kenneth E. Torrance and Donald P. Greenberg. *Image-Based BRDF Measurement Including Human Skin*. In Eurographics Symposium on Rendering, pages 131–144, 1999. (Cited on page 25.)
- [Marschner 2000] Steve R. Marschner, Stephen H. Westin, Eric P. F. Lafortune and Kenneth E. Torrance. *Image-Based Bidirectional Reflectance Distribution Function Measurement*. Applied Optics, vol. 39, no. 16, pages 2592–2600, Jun 2000. (Cited on page 25.)
- [Matusik 2003] Wojciech Matusik, Hanspeter Pfister, Matt Brand and Leonard McMillan. *A data-driven reflectance model*. ACM Transactions on Graphics, vol. 22, no. 3, pages 759–769, July 2003. (Cited on pages 25, 26, 28, 66 and 68.)
- [McCool 2001] Michael D. McCool, Jason Ang and Anis Ahmad. *Homomorphic factorization of BRDFs for high-performance rendering*. In Proceedings of the 28th annual conference on Computer graphics and interactive techniques, SIGGRAPH '01, pages 171–178, New York, NY, USA, 2001. ACM. (Cited on pages 29 and 30.)

- [McGuire 2009] Morgan McGuire and David Luebke. *Hardware-Accelerated Global Illumination by Image Space Photon Mapping*. In Proceedings of the 2009 ACM SIGGRAPH/EuroGraphics conference on High Performance Graphics, New York, NY, USA, August 2009. ACM. (Cited on page 15.)
- [Metropolis 1949] Nicholas Metropolis and Stan Ulam. *The Monte Carlo Method*. Journal of the American Statistical Association, vol. 44, no. 247, pages 335–341, September 1949. (Cited on page 13.)
- [Munoz 2011] Adolfo Munoz, Jose I. Echevarria, Francisco Seron, Jorge Lopez-Moreno, Mashhuda Glencross and Diego Gutierrez. *BSSRDF Estimation from Single Images*. Computer Graphics Forum, vol. 30, pages 455–464, 2011. (Cited on page 26.)
- [Murray-Coleman 1990] J.F. Murray-Coleman and A.M. Smith. *The Automated Measurement of BRDFs and their Application to Luminaire Modeling*. Journal of the Illuminating Engineering Society, pages 87–99, 1990. (Cited on page 25.)
- [Neumann 1999] László Neumann, Attila Neumann and László Szirmay-Kalos. *Compact Metallic Reflectance Models*. Computer Graphics Forum, vol. 18, no. 3, pages 161–172, 1999. (Cited on page 32.)
- [Ngan 2005] Addy Ngan, Frédo Durand and Wojciech Matusik. *Experimental Analysis of BRDF Models*. In Eurographics Workshop on Rendering, pages 117–126, June 2005. (Cited on pages 25, 27, 37, 38, 39, 73, 75, 77 and 80.)
- [Ngan 2006] Wai Kit Addy Ngan. *Acquisition and modeling of material appearance*. PhD thesis, Cambridge, MA, USA, 2006. (Cited on page 26.)
- [Nichols 2009] Greg Nichols and Chris Wyman. *Multiresolution splatting for indirect illumination*. In Symposium on Interactive 3D graphics and games (I3D), pages 83–90, 2009. (Cited on pages 15 and 17.)
- [Nichols 2010a] Greg Nichols, Rajeev Penmatsa and Chris Wyman. *Interactive, Multiresolution Image-Space Rendering for Dynamic Area Lighting*. Computer Graphics Forum, vol. 29, no. 4, pages 1279 – 1288, 2010. (Cited on pages 15 and 17.)
- [Nichols 2010b] Greg Nichols and Chris Wyman. *Interactive Indirect Illumination Using Adaptive Multiresolution Splatting*. IEEE Transactions on Visualization and Computer Graphics, vol. 16, no. 5, pages 729–741, 2010. (Cited on pages 15 and 17.)
- [Nicodemus 1992] F. E. Nicodemus, J. C. Richmond, J. J. Hsia, I. W. Ginsberg and T. Limperis. *Radiometry*. chapitre Geometrical considerations and nomenclature for reflectance, pages 94–145. Jones and Bartlett Publishers, Inc., USA, 1992. (Cited on page 21.)
- [nis 2010] *NIST reference reflectometer: STARR facility*. <http://physics.nist.gov/>, 2010. (Cited on page 25.)

- [Oppenheim 1975] Alan V. Oppenheim and Ronald W. Schaffer. Digital signal processing. Prentice–Hall, 1975. (Cited on page 52.)
- [Oren 1994] Michael Oren and Shree K. Nayar. *Generalization of Lambert’s reflectance model*. In Proceedings of the 21st annual conference on Computer graphics and interactive techniques, SIGGRAPH ’94, pages 239–246, New York, NY, USA, 1994. ACM. (Cited on page 36.)
- [Phong 1975] Bui Tuong Phong. *Illumination for computer generated pictures*. Communications of the ACM, vol. 18, no. 6, pages 311–317, June 1975. (Cited on page 31.)
- [Ren 2011] Peiran Ren, Jiaping Wang, John Snyder, Xin Tong and Baining Guo. *Pocket reflectometry*. ACM Transactions on Graphics, vol. 30, no. 4, pages 45:1–45:10, July 2011. (Cited on page 26.)
- [Ritschel 2009] T. Ritschel, T. Engelhardt, T. Grosch, H.-P. Seidel, J. Kautz and C. Dachsbacher. *Micro-rendering for scalable, parallel final gathering*. ACM Transactions on Graphics, vol. 28, no. 5, pages 132:1–132:8, 2009. (Cited on page 15.)
- [Romeiro 2008] Fabiano Romeiro, Yuriy Vasilyev and Todd Zickler. *Passive Reflectometry*. In Proceedings of the 10th European Conference on Computer Vision: Part IV, ECCV ’08, pages 859–872, Berlin, Heidelberg, 2008. Springer-Verlag. (Cited on page 26.)
- [Romeiro 2010] Fabiano Romeiro and Todd Zickler. *Blind reflectometry*. In Proceedings of the 11th European conference on Computer vision: Part I, ECCV’10, pages 45–58, Berlin, Heidelberg, 2010. Springer-Verlag. (Cited on page 26.)
- [Rusinkiewicz 1998] Szymon Rusinkiewicz. *A New Change of Variables for Efficient BRDF Representation*. In Eurographics Workshop on Rendering, June 1998. (Cited on pages 28 and 29.)
- [Saito 1990] Takafumi Saito and Tokiichiro Takahashi. *Comprehensible rendering of 3-D shapes*. SIGGRAPH Computer Graphics, vol. 24, no. 4, pages 197–206, September 1990. (Cited on page 15.)
- [Schlick 1994] Christophe Schlick. *An Inexpensive BRDF Model for Physically-based Rendering*. Computer Graphics Forum, vol. 13, pages 233–246, 1994. (Cited on page 36.)
- [Schröder 1995] Peter Schröder and Wim Sweldens. *Spherical wavelets: efficiently representing functions on the sphere*. In Proceedings of the 22nd annual conference on Computer graphics and interactive techniques, SIGGRAPH ’95, pages 161–172, New York, NY, USA, 1995. ACM. (Cited on page 29.)
- [Segovia 2006] B. Segovia, J. C. Iehl, R. Mitanchey and B. Péroche. *Non-interleaved deferred shading of interleaved sample patterns*. In Symposium on Graphics Hardware (GH), pages 53–60, 2006. (Cited on page 15.)

- [Shopf 2009] Jeremy Shopf, Greg Nichols and Chris Wyman. *Hierarchical Image-Space Radiosity for Interactive Global Illumination*. Computer Graphics Forum, vol. 28, no. 4, pages 1141–1149, 2009. (Cited on pages 15 and 17.)
- [Sloan 2002] Peter-Pike Sloan, Jan Kautz and John Snyder. *Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments*. ACM Transactions on Graphics, vol. 21, no. 3, pages 527–536, July 2002. (Cited on pages 42 and 43.)
- [Smith 1967] B. Smith. *Geometrical shadowing of a random rough surface*. IEEE Transactions on Antennas and Propagation, vol. 15, no. 5, pages 668–671, September 1967. (Cited on pages 34, 36 and 67.)
- [Soler 2009] Cyril Soler, Kartic Subr, Frédo Durand, Nicolas Holzschuch and François Sillion. *Fourier depth of field*. ACM Transactions on Graphics, vol. 28, no. 2, pages 18:1–18:12, May 2009. (Cited on pages 19 and 46.)
- [Soler 2010] Cyril Soler, Olivier Hoel and Frank Rochet. *A Deferred Shading Algorithm for Real-Time Indirect Illumination*. In ACM SIGGRAPH Talks, page 18:1, 2010. (Cited on page 15.)
- [Spanier 1969] J. Spanier and E. M. Gelbard. Monte Carlo principles and neutron transport problems. Addison-Wesley series in computer science and information processing. Addison-Wesley Pub. Co., 1969. (Cited on page 13.)
- [Torrance 1967] K. E. Torrance and E. M. Sparrow. *Theory for Off-Specular Reflection From Roughened Surfaces*. Journal of the Optical Society of America, vol. 57, no. 9, pages 1105–1112, September 1967. (Cited on pages 33 and 66.)
- [Trowbridge 1975] T. S. Trowbridge and K. P. Reitz. *Average irregularity representation of a rough surface for ray reflection*. Journal of the Optical Society of America, vol. 65, no. 5, pages 531–536, May 1975. (Cited on page 35.)
- [Veach 1995] Eric Veach and Leonidas J. Guibas. *Optimally combining sampling techniques for Monte Carlo rendering*. In SIGGRAPH 95, pages 419–428, 1995. (Cited on page 13.)
- [Veach 1998] Eric Veach. *Robust monte carlo methods for light transport simulation*. PhD thesis, Stanford, CA, USA, 1998. http://graphics.stanford.edu/papers/veach_thesis/. (Cited on pages 13 and 14.)
- [Walter 2007] B. Walter, S. Marschner, H. Li and K. E. Torrance. *Microfacet models for refraction through rough surfaces*. In Eurographics Symposium on Rendering, 2007. (Cited on pages 34, 35, 36 and 72.)
- [Wang 2009a] Jiaping Wang, Peiran Ren, Minmin Gong, John Snyder and Baining Guo. *All-frequency rendering of dynamic, spatially-varying reflectance*. ACM Transactions on Graphics, vol. 28, no. 5, pages 133:1–133:10, December 2009. (Cited on pages 29, 59 and 61.)

-
- [Wang 2009b] Rui Wang, Rui Wang, Kun Zhou, Minghao Pan and Hujun Bao. *An efficient GPU-based approach for interactive global illumination*. ACM Transactions on Graphics, vol. 28, no. 3, pages 91:1–91:8, July 2009. (Cited on page 15.)
- [Ward 1988] Gregory J. Ward, Francis M. Rubinstein and Robert D. Clear. *A ray tracing solution for diffuse interreflection*. SIGGRAPH Computer Graphics, vol. 22, no. 4, pages 85–92, June 1988. (Cited on page 13.)
- [Ward 1992a] Gregory J. Ward. *Measuring and modeling anisotropic reflection*. ACM SIGGRAPH Computer Graphics, vol. 26, no. 2, pages 265–272, July 1992. (Cited on pages 25, 31 and 37.)
- [Ward 1992b] Gregory J. Ward and Paul S. Heckbert. *Irradiance Gradients*. pages 85–98, 1992. (Cited on page 13.)
- [Westin 1992] Stephen H. Westin, James R. Arvo and Kenneth E. Torrance. *Predicting reflectance functions from complex surfaces*. ACM SIGGRAPH Computer Graphics, vol. 26, no. 2, pages 255–264, July 1992. (Cited on page 29.)