

Light Propagation in Digital Material Structures

Dissertation

zur

Erlangung des Doktorgrades (Dr. rer. nat.)

der

Mathematisch-Naturwissenschaftlichen Fakultät

der

Rheinischen Friedrich-Wilhelms-Universität Bonn

vorgelegt von

M.Sc. Zdravko Velinov

aus

Moskau, Russland

Bonn, März, 2019

Angefertigt mit Genehmigung der Mathematisch-Naturwissenschaftlichen
Fakultät der Rheinischen Friedrich-Wilhelms-Universität Bonn

1. Gutachter: Prof. Dr. Reinhard Klein
2. Gutachter: Prof. Dr. Matthias Hullin
3. Gutachter: Prof. Dr. Thomas Schultz
4. Gutachter: Prof. Dr. Christoph Bourauel

Tag der Promotion: 20.03.2019

Erscheinungsjahr: 2019

Contents

Abstract	v
Zusammenfassung	vii
Acknowledgements	ix
1 Introduction	1
1.1 Publications	3
1.2 Contributions and outline	4
1.3 Relation to previous results	5
2 Preliminaries	9
2.1 Rendering	9
2.1.1 Light propagation	10
2.1.2 Surface scattering	12
2.1.3 Volume scattering	14
2.1.4 Wave optics	16
2.2 Acquisition	18
2.2.1 Appearance acquisition devices	18
2.2.2 Perspective camera model	19
2.2.3 Geometric calibration	20
2.2.4 Light source registration	20
2.2.5 Light source intensities estimation	22
2.2.6 Photometric calibration	22
2.3 Inverse methods	22

3	An Interactive Appearance Model for Microscopic Fiber Surfaces	25
3.1	Introduction	26
3.2	Background	28
3.3	Related Work	30
3.4	Proposed Approach	31
3.4.1	Reflectance model	31
3.4.2	Parameter fitting	36
3.4.3	Post processing	37
3.5	Implementation and Results	38
3.6	Discussion	42
3.7	Conclusions	44
4	Real-Time Rendering of Wave-Optical Effects on Scratched Surfaces	47
4.1	Introduction	49
4.2	Related work	51
4.2.1	Wave-optical shading models	51
4.2.2	Rendering scratch-like features	52
4.2.3	Extended light sources	52
4.3	Theory	54
4.3.1	Model	54
4.3.2	Integration of the spatially varying reflectance over the camera pixel footprint	58
4.3.3	Scratch area density	60
4.3.4	Integration of the reflectance in angular domain over the solid angle subtended by area light sources	61
4.4	Implementation	69
4.5	Results	70
4.6	Discussion	73
	Appendices	77
4.A	Additional performance evaluation	77
5	Appearance Capture and Modeling of Human Teeth	85
5.1	Introduction	87
5.2	Related Work	89
5.2.1	Teeth Capture	91

CONTENTS

5.2.2	Facial Appearance Capture	92
5.2.3	Volumetric Appearance Capture	92
5.2.4	Rendering Translucent Materials	93
5.3	Method Overview	94
5.4	Data Acquisition	94
5.4.1	Image Capture	94
5.4.2	Geometry Capture	95
5.5	Teeth Appearance Model	97
5.5.1	Enamel	97
5.5.2	Dentin	98
5.5.3	Gingiva and oral cavity	99
5.6	Model Parameter Optimization	100
5.6.1	Enamel and Gums	101
5.6.2	Dentin	105
5.7	Results	107
5.8	Discussion and Future Work	110
Appendices		117
5.A	Derivative of distance with respect to vertex translation	117
5.B	Optimization process statistics	119
5.C	Optimization results	124
6	Closure	135
6.1	Summary of contributions	135
6.2	Limitations, future work and outlook	137
6.2.1	Fabrics	137
6.2.2	Worn surfaces and wave optics	138
6.2.3	Inverse rendering of participating media	140
6.2.4	Outlook	141
A	Publikationen	143
List of Symbols		145
List of Figures		153
Index		157
Bibliography		159

Abstract

Light takes many paths throughout the physical environment, defined by a multitude of interactions with matter, before reaching an imaging sensor or the human eye and thus forming the visual perception of the world. Recreating these interactions in virtual environments requires to consider both gradual scattering by volumes or sudden change introduced by surface boundaries between materials. In many cases, the properties defining these phenomena are not constant and can be affected by external forces. Two main paradigms have arisen in the field of physics and computer graphics to model the complexity of these interactions: geometrical and wave optics. On macroscopic scale the wave nature of light is often neglected as the statistics of the process of incoherent interference reduces the impact of this component. An exception to this rule are many worn smooth surfaces. The notion of scale is violated by the microscopic structures and their intricate iridescent appearance can be clearly observed by the naked eye. Finally, deep structures within translucent objects can significantly influence their overall appearance. They invalidate the common assumption employed in computer graphics that the appearance of objects can be modeled as optically thick and fully defined by the surface reflectance. We thus outline three cases where the existing approaches within the computer graphics literature have difficulties explaining the appearance of everyday objects and provide specific solutions for acquiring and recreating the digital material appearance. They are presented in the form of an article thesis (cumulative thesis) discussing recent publications by the author, regarding the following topics:

- Appearance of semi-unstructured assemblies of fibers interwoven into pile fabrics and their characteristic appearance dependent on the history of tactile interactions. In Chapter 3 we outline a complete end-to-end pipeline capable of acquiring the appearance of this class of

materials and modeling the dynamic change of reflectance based on virtual interaction.

- Efficient rendering of wave-optical effects appearing on many everyday objects with smooth worn surface. In Chapter 4 we derive equations and develop algorithms for illumination and anti-aliasing the appearance of these surfaces with real-time performance on modern commodity hardware.
- Modeling and acquisition of the appearance of human teeth and oral cavity. In Chapter 5 we develop a framework that pairs path space derivative estimation with stochastic gradient descent to acquire jointly inner geometry and scattering parameters from photographs captured in a conventional multiple camera and light source setup.

Zusammenfassung

Licht nimmt viele Wege durch die physische Umgebung, die durch eine Vielzahl von Wechselwirkungen mit Materie definiert wird, bevor es einen bildgebenden Sensor oder das menschliche Auge erreicht und somit die visuelle Wahrnehmung der Welt formt. Das Nachbilden dieser Wechselwirkungen in virtuellen Umgebungen erfordert sowohl die Betrachtung gradueller Streuprozesse in ausgedehnten Medien als auch die Einbeziehung abrupter Änderungen an Grenzflächen zwischen verschiedenen Materialien. In vielen Fällen sind die Eigenschaften, die diese Phänomene definieren, nicht konstant und können durch externe Kräfte beeinflusst werden. Auf dem Gebiet der Physik und der Computergrafik haben sich zwei Hauptparadigmen herauskristallisiert, um die Komplexität dieser Wechselwirkungen zu modellieren: die geometrische und die Wellenoptik. Auf makroskopischer Ebene wird die Wellennatur von Licht oft vernachlässigt, da deren Einfluss durch die statistischen Eigenschaften der stattfindenden inkohärenten Interferenz reduziert wird. Eine Ausnahme von dieser Regel sind viele abgenutzte glatte Oberflächen. Die Annahme der Skaleneigenschaft wird hier durch die mikroskopischen Strukturen und deren komplizierte irisierende Erscheinung Schillerfarben verletzt; Ein Effekt, der mit bloßem Auge beobachtbar ist. Schließlich können tiefe Strukturen innerhalb durchscheinender Objekte ihren Gesamteindruck wesentlich beeinflussen. Sie entkräften die in der Computergrafik übliche Annahme, dass Objekte als optisch dicht modelliert und, als solche, vollständig durch die Eigenschaften ihrer Oberfläche definiert werden können. Wir skizzieren daher drei Fälle, in denen die bestehenden Ansätze in der Computergrafik-Literatur Schwierigkeiten haben, das Erscheinungsbild alltäglicher Objekte zu erklären und präsentieren spezifische Lösungen für deren Messung und Simulation. Diese Methoden sind im Folgenden in der Form einer kumulativen Arbeit beschrieben und befassen sich mit den jüngsten Publikationen des Autors, aufgeteilt in die folgenden Themengebiete:

te:

- Das Aussehen von semi-unstrukturierten Faserverbänden und daraus erstellten Geweben sowie deren charakteristisches Aussehen in Abhängigkeit von vorangegangenen taktilen Wechselwirkungen. In Kapitel 3 skizzieren wir eine komplette End-to-End-Pipeline, die in der Lage ist, das Erscheinungsbild dieser Materialklasse zu reproduzieren und dynamische Änderungen aufgrund virtueller Interaktion zu modellieren.
- Effizientes Rendern wellenoptischer Effekte, die oftmals auf Alltagsgegenständen mit glatter Oberfläche auftreten. Kapitel 4 führt durch die Herleitung und Entwicklung von Algorithmen für die Berechnung des Aussehens solcher Oberflächen. Dabei wird besonderes Augenmerk auf die effiziente Berechnung unter komplexer Beleuchtung und die Verwendung von Antialiasing-Techniken gelegt. Die Implementierung erlaubt Renderings in Echtzeit unter Verwendung aktueller handelsüblicher Hardware.
- Modellierung und Messung des Aussehens von menschlichen Zähnen und des Mundraums. In Kapitel 5 entwickeln wir ein Framework zur Kombination der Abschätzung von Weg-/Raum-Gradienten mit stochastischen Gradientenverfahren zur gemeinsamen Rekonstruktion der Geometrie- und Streuparameter. Der dafür verwendete Datensatz wurde mit einer herkömmlichen Mehrfachkamera aufgenommen wurden.

Acknowledgements

Nowadays the field of computer graphics has progressed to ever growing engineering and scientific complexity and parts of this thesis were developed as a team effort. It is thus essential to thank everyone that contributed to the projects described in the main body of this text. First I would like to thank my academic advisor Prof. Matthias Hullin who created a nice scholarly environment where scientific problems can be freely discussed and solved within group and one-on-one discussions. Further thanks go to Prof. Reinhard Klein for his work in directing the Grad School on Digital Material Appearance. Sebastian Werner is the person that I would like to thank the most for his extensive contribution on multiple projects and many discussions that we had on various topics over the years.

The 'Capture and Effects' and 'Rendering' groups at Disney Research Zurich were instrumental in the teeth appearance project and thus I would like to thank their respective lead researchers Thabo Beeler and Jan Novak. I am indebted to Marios Papas for his hard work on parts of the pipeline used in this project. I would also like to thank Derek Bradley, Paulo Gotardo, Parsa Mirdehghan, Steve Marschner who additionally contributed to the project. Early programming work on the optimization pipeline was done by Lifan Wu. Dr. med. dent. Lukas Stucky, Raphael Laue, Dr. med. dent. Sven Mühlemann and the rest of the team at the Center of Dental Medicine at the University of Zurich performed the high accuracy intra-oral geometry scans. I thank our volunteers Gaspard Zoss and Jan Wezel who participated in both our intra-oral scanning and appearance capture sessions. The content from this project is reprinted with permission by Disney Research.

Prof. Wenzel Jakob was a longtime contributor to the 'Scratch Iridescence' project whom I would like to thank for the in-depth discussions. Rodrigo Martín assisted with proofreading and discussing some of the publications about this project.

Two of the works included in the main body of this thesis were funded by the X-Rite Chair and Graduate School on Digital Material Appearance. Funding on the teeth appearance project was provided by Disney Research.

Multiple appearance capture sessions were done with the Dome II at University of Bonn. Heinz-Christian Steinhausen, Sebastian Merzbach and Dennis den Brok assisted with using the device and the respective BTF processing pipeline.

Finally, I would like to thank my family for providing moral support during my PhD studies and the students at University of Bonn who organized many fun events which kept me a sane and happy member of the society.

CHAPTER 1

Introduction

Our planet is bathed each day by waves of light created within the great nuclear furnace in the depths of our Sun. They propagate through the atmosphere to reach every directly exposed object on the surface and reflect or transmit through its volume until they finally imping the eye of a biological being or the sensor of a synthetic device. The projected image is a combination of many light paths, but ultimately contains visual cues about the closest observed object. The field of computer graphics strives to recreate the entire richness of interaction between light and matter while considering other modalities that might affect the perception of the physical world. Over the years many algorithms were developed to solve specific subsets of this problem which find application in the entertainment industry to create movies and interactive experiences, in fabrication for designing entire products before even a single unit gets produced or in cultural heritage to preserve objects which might decay in the physical world.

Digital materials describe the relationship between the incident and outgoing light at each point in space through statistical models. They are derived from physics first principles or by data-driven approaches aiming to produce faithful compact representation. Historically, computer graphics treated the phenomena of light propagating through scattering media separately from light scattered at surface boundaries, giving rise to a different set of equations describing the process of light transport and its respective parameters. Recent works showed that these concepts can be expressed in a similar framework [Heitz et al. 2016b; Dupuy et al. 2016; Velinov and Hullin 2016] by considering the limit case of semi-infinite opaque surface and further

experimental validation of some of these findings is available in this thesis.

Both the automotive and movie industry are interested in digitizing objects for the purpose of interactively editing and seamlessly blending them back in photographs or movie scenes. To achieve this task, elaborate measurement setups are built that are composed of multiple light sources and cameras [Guarnera et al. 2016; Klehm et al. 2015] to capture multiple photographs of static objects or living humans. Many of them consider only surface reflectance for modeling and acquisition simplicity. However, many everyday objects and living tissue (plants, fruits, skin, teeth, nails) are in fact translucent and they can further contain complex inner structure. Similarly to surface and volume rendering, the problem of appearance acquisition for different media and surfaces was considered separately by many previous works. One example covered by this work is modeling and acquisition of hidden media within human teeth. In fact, similar phenomenon can be observed on skin where individual veins can be distinguished beneath the skin or seeds that are visible in sliced fruits. These structures have to be efficiently expressed for the purpose of light transport simulation. Triangular meshes are the most common shape representation used in graphics. The necessary framework for inverse rendering inner structures with these primitives in path space is derived in this work. Complete light transport simulation in these specific cases is also required to account correctly for inter-reflection. Differentiable path tracing is the approach cover by this work to fully solve the dependence of appearance on light transport effects. Another often neglected fact about surface appearance that is not considered by acquisition techniques is that appearance can be affected by tactile interaction and lead to change of appearance. One solution of this problem is further provided in this work.

Geometrical optics were successfully used to recreate scenes from the real world in virtual environments. However, smooth surfaces can reflect close to mirror-like sources of light spread throughout the environment and thus violate the notion of scale. Furthermore, microscopic defects comparable in size to the wavelength on these surfaces show iridescent behavior. Techniques for modeling diffraction [Werner et al. 2017] were successfully used in the past to explain this phenomenon. However, efficient modeling of surfaces requires to consider both the angular extent of light sources and the projected area of a camera pixel footprint. While in offline rendering it is enough to solve this problem numerically, the time budget in real-time graphics is much more limited and requires deriving efficient approximations of these integral

equations. The mirror-like structure also prevents the use of punctual light sources to approximate extended light sources as discontinuities become quite apparent. The specialized solutions of this problem are discussed in the main body of the text.

By considering many facets of appearance modeling and acquisition, this work provides a portfolio of approaches that can be used to recreate objects and digital doubles for the purpose of object prototyping and visual effects. Multiple topics regarding appearance modeling are discussed, ranging from surface and volume based modeling and acquisition, and efficient approaches for light computation of iridescent worn surfaces. These are supplemented with novel techniques for solving the core of their respective problem. Many of the approaches include the necessary tools and acquisition pipelines to apply them in production environments.

1.1 Publications

The following publications are included within this article thesis (cumulative dissertation):

- Zdravko Velinov and Matthias B. Hullin [Oct. 2016]. “An Interactive Appearance Model for Microscopic Fiber Surfaces.” In: *Vision, Modeling, and Visualization 2016*
- Zdravko Velinov, Sebastian Werner, and Matthias B. Hullin [May 2018b]. “Real-Time Rendering of Wave-Optical Effects on Scratched Surfaces.” In: *Computer Graphics Forum* 37.2. ISSN: 1467-8659. DOI: 10.1111/cgf.13347
- Zdravko Velinov, Marios Papas, Derek Bradley, Paulo Gotardo, Parsa Mirdehghan, Steve Marschner, Jan Novák, and Thabo Beeler [Dec. 2018a]. “Appearance Capture and Modeling of Human Teeth.” In: *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 37.6, 207:1–207:13. ISSN: 0730-0301. DOI: 10.1145/3272127.3275098. URL: <http://doi.acm.org/10.1145/3272127.3275098>

Velinov et al. [2018b] builds upon theory that was previously published by the author of this thesis as a second author with equal contribution,

- Sebastian Werner, Zdravko Velinov, Wenzel Jakob, and Matthias B. Hullin [Nov. 2017]. “Scratch Iridescence: Wave-optical Rendering of Diffractive Surface Structure.” In: *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 36.6, 207:1–207:14. ISSN: 0730-0301. DOI: 10.1145/3130800.3130840. URL: <http://doi.acm.org/10.1145/3130800.3130840>

The author of this thesis built the interactive tools and data structures that were used to generate results in the publication. Further help with streamlining different derivations led ultimately to the development of the main scientific and technical contributions described in Velinov et al. [2018b]. Chapter 2 briefly introduces the main digital material modeling advancements brought by Werner et al. [2017] and the theory upon which they are built. The main contributions of that work are excluded from the list of novel techniques and results detailed by this thesis. Further in-depth information can be acquired by reading the original work.

1.2 Contributions and outline

Preliminary information introducing basic concepts regarding light propagation modeling and acquisition is provided in Chapter 2. The main body of the text covers three cases of digital materials requiring special treatment to accurately model their structure. They are paired with novel techniques for modeling the interaction with light and/or acquisition of appearance from photographs.

Pile fabrics have dynamic appearance dependent on the history of interactions. Chapter 3 contributes to the appearance modeling literature:

- a model capable of simulating the reflectance of unstructured random pile fabrics,
- a complete end-to-end acquisition pipeline capable of estimating the parameters of the model,
- a rendering approach for recreating the tactile interaction.

Worn surfaces often contain structures comparable to the wavelength. Long concentric iridescent highlights can be observed around the mirror reflection of light sources proportional to their angular extent. Chapter 4 provides specific derivations and algorithms for efficient rendering in real time:

- an efficient approximation of the integral over the projected area of spherical and polygonal light sources,
- an efficient approximation of the integral over the projected camera footprint,
- a set of optimized for rendering data structures to store the individual scratches appearing on the surface.

Recreating the appearance of *teeth and the oral cavity* requires considering the complete light transport to account for inter-reflections and subsurface light propagation. Chapter 5 provides:

- a model capable of accounting for the phenomena observed on the surface and interior of teeth,
- a complete acquisition and parameter estimation framework,
- specific equations and techniques to acquire the hidden inner dentin surface boundary.

The works are presented as close as possible to their previously published form. Minor changes involve ensuring unique math notation to clearly separate concepts introduced in different chapters and typographic or misprint corrections. All symbols are outlined in a combined chapter at the end of the thesis. Chapter 4 includes an appendix (cf. Sec. 4.A) comprised of a supplemental paper published along with the proceedings of the conference in electronic form. Chapter 5 combines the original appendix (cf. Sec. 5.A) and supplemental data published along with the paper in electronic form (cf. Sec. 5.B, 5.C).

1.3 Relation to previous results

Pile fabrics appearance modeling and acquisition. Modeling surfaces as a Bidirectional Scattering Distribution Function (BSDF) is one of the most common approaches employed in computer graphics and it received extensive research over the years [Guarnera et al. 2016]. Unstructured surfaces, such as pile fabrics, are more suitable for regular distribution based modeling than using specialized techniques that encode the weave pattern directly into the BSDF [Sadeghi et al. 2013; Irawan and Marschner 2012]. Recent advances

in volume rendering brought multiple techniques capable of modeling complex fabrics over different scales [Jakob et al. 2010; Zhao et al. 2011; Heitz et al. 2015]. These techniques require evaluating an infinitely dimensional integral expressed as Neumann series. Therefore they are unpractical for real-time rendering where existing hardware acceleration focuses on single reflection modeling and shadow casting. Instead, the model described within this work takes inspiration from these volume rendering approaches and derives a Bidirectional Reflectance Distribution Function (BRDF) in the limit case of semi-infinite surface and describes the respective pipeline for parameter estimation. The model bears some similarities to isotropic BSDF models [Seeliger 1888; Chandrasekhar 1960], however the anisotropic nature of the microflake theory allows expressing both rough and directionally smooth surfaces, which makes it suitable for fabrics modeling. A parallel work focusing on theory [Dupuy et al. 2016] noted the similarity of this kind of models to the microfacet theory with the only difference being how different sides of the flakes are treated. The main contributions over this work is a complete pipeline covering acquisition of appearance and rendering of tactile interaction.

Worn surfaces appearance modeling and illumination. Microscopic defects appearing on worn surfaces were studied by multiple recent works [Yan et al. 2014; Yan et al. 2016; Raymond et al. 2016]. These works consider only efficient integration from geometrical optics perspective. Werner et al. [2017] showed that in fact these surfaces contain observable wave-optical phenomenon. Building upon this work that was co-authored with the author of this thesis, an efficient set of algorithms is derived for rendering these surfaces under extended light source illumination and integration over the camera pixel footprint in real time. The equations expressed within this thesis are based on approximations derived directly from the integrals in angular and spatial domain. Other recent works on area-light modeling fit distributions to these integrals for microfacet models [Heitz et al. 2016a; Dupuy et al. 2017; Heitz and Hill 2017] or they provide exact solution only for the Phong model [Arvo 1995; Lecocq et al. 2016]. Further contributions included in the text cover efficient data structures for storing the individual scratches.

Teeth and oral cavity appearance modeling and acquisition. In dental literature multiple destructive measurements were performed to identify the average scattering parameters of teeth [Fried et al. 1995; Pop-Ciutrila

et al. 2015; Spitzer and Ten Bosch 1975]. Teeth as any other living tissue have very different parameters when they are alive or dead and their parameters can vary quite noticeably between each individual person based on genetic, age or hygiene related reasons. Thus capturing these parameters in vivo is more desirable. Existing dental scanners (3Shape TRIOS[®]) extract diffuse albedo and very similar approaches are used by teeth geometry reconstruction techniques [Wu et al. 2016; Kawai et al. 2014]. Low-dimensional model based approaches are often applied to extract the appearance of teeth [Buchillard et al. 2007; Farag et al. 2013; Mehl and Blanz 2005; Mostafa et al. 2014], however teeth are, in practice, optically thin and forward scattering which requires considering both their surface model and internal structure. Path tracing approaches are able to realistically recreate the complete light transport within the tooth interior and model inter-reflection of light from other elements of the oral cavity. Previous works demonstrated the feasibility of building path tracing based inverse renderers [Gkioulekas et al. 2013; Khungurn et al. 2015; Gkioulekas et al. 2016]. The main extension provided in this work regards using triangular meshes with this framework to estimate the inner structure from photographs. Additionally, a complete end-to-end pipeline for capture, parameter estimation and rendering is presented.

CHAPTER 2

Preliminaries

Digital doubles of humans and objects are created through a process covering many steps. Rendering plays a central role in the process of bringing virtual environments to life. Different models were developed to express physical phenomena and some of them are even tailored towards artists. While synthetically generating virtual representations was done successfully for decades, it is a very laborious task to exactly match reality. Acquisition and inverse methods promise to automate many of these tasks. They have already been used in both the movie and automotive industry to digitize humans and materials, respectively. Multiple aspects of this process are discussed in the following, including rendering (cf. Sec. 2.1), acquisition (cf. Sec. 2.2) and inverse methods (cf. Sec. 2.3).

2.1 Rendering

Rendering covers the process of converting parametric representation of a virtual environment to visual imagery. Light propagation algorithms lie at its foundation and they can vary in their simulation accuracy and completeness. Surfaces are often represented by triangular meshes. They determine reflectance and transmission at a boundary in space. Their behavior is mainly guided by the rendering equation [Kajiya 1986] and geometrical optics; however, wave optics have been successfully incorporated to further bring details on the order of the wavelength. On the other hand, volumes take a step further and consider the subsurface behavior of materials. Multiple subsections in the following discuss these aspects of rendering. Fundamental concepts

regarding light propagation are first discussed in Sec. 2.1.1. Then surfaces and volumes are discussed with their respective models in Sec. 2.1.2, 2.1.3. The discussion ends in Sec. 2.1.4 with wave optics as an alternative approach for light propagation and surface reflectance modeling.

2.1.1 Light propagation

Light is defined by the wave-particle duality. Computer graphics models the interaction of light with matter on scales observable by the human eye and conventional cinema and photography cameras. The interaction of light on cosmic scale with the curvature of space-time is rarely considered. Similarly, the statistic nature of the process of incoherent interference to a large extent reduces the cases when wave-optical effects can be observed on the scales of interest. Geometrical optics is the preferred paradigm for expressing light propagation on macroscopic scale in graphics. It considers the radiance propagated along a ray under the assumption that steady state equilibrium is reached and change over time is not observable. Polarization of the electromagnetic field is not considered. Scattering effects are assumed to completely explain the statistics of light transport and fluorescence or phosphorescence are ignored. Geometrical optics expresses interaction of light as a superposition of the radiance of each individual path and interference between waves of light is ignored. The main equation that guides most modern rendering systems which considers all scattering phenomena within a medium is the Radiative Transfer Equation (RTE). Commonly in physics and graphics literature, it is expressed in a differential form [Chandrasekhar 1960],

$$(\boldsymbol{\omega} \cdot \nabla)L(\mathbf{x}, \boldsymbol{\omega}) = -\sigma_t(\mathbf{x}, \boldsymbol{\omega})L(\mathbf{x}, \boldsymbol{\omega}) + \sigma_s(\mathbf{x}, \boldsymbol{\omega})L_s(\mathbf{x}, \boldsymbol{\omega}) + \sigma_a(\mathbf{x}, \boldsymbol{\omega})L_e(\mathbf{x}, \boldsymbol{\omega}), \quad (2.1)$$

where out-scattering combined with absorption are modeled by the extinction coefficient $\sigma_t(\mathbf{x}, \boldsymbol{\omega}) = \sigma_s(\mathbf{x}, \boldsymbol{\omega}) + \sigma_a(\mathbf{x}, \boldsymbol{\omega})$, light emitted within the volume $L_e(\mathbf{x}, \boldsymbol{\omega})$ is absorbed according to the absorption coefficient $\sigma_a(\mathbf{x}, \boldsymbol{\omega})$, radiance $L_s(\mathbf{x}, \boldsymbol{\omega})$ is in-scattered in the cross section according to the scattering coefficient $\sigma_s(\mathbf{x}, \boldsymbol{\omega})$,

$$L_s(\mathbf{x}, \boldsymbol{\omega}_o) = \int_{\mathcal{S}^2} f_p(\boldsymbol{\omega}_i, \boldsymbol{\omega}_o) L_i(\boldsymbol{\omega}_i) d\boldsymbol{\omega}_i, \quad (2.2)$$

with the phase function $f_p(\boldsymbol{\omega}_i, \boldsymbol{\omega}_o)$ characterizing scattering of the incident radiance $L_i(\boldsymbol{\omega}_i)$ at a point in space. Everything is parameterized according to the direction of propagation $\boldsymbol{\omega}$ and the spatial position \mathbf{x} . RTE is

commonly solved by considering light transmitted between two points in the direction of propagation, thus both sides are integrated with respect to the light propagation direction ω ,

$$L(\mathbf{x}, \omega) = \int_0^\infty \mathcal{T}_t(\mathbf{x}, t) [\sigma_s(\mathbf{x}, \omega) L_s(\mathbf{x}, \omega) + \sigma_a(\mathbf{x}, \omega) L_e(\mathbf{x}, \omega)] dt. \quad (2.3)$$

In the classical theory of radiative transfer the scattering events are uncorrelated and thus the transmittance is according to an exponential law w.r.t. distance along the ray t ,

$$\mathcal{T}_t(\mathbf{x}, t) = e^{-\int_0^t \sigma_t(\mathbf{x}-t\omega, \omega) dt}. \quad (2.4)$$

This relation is also known as the Beer-Lambert law in the case of isotropic media where the scattering events are independent of the direction of light propagation.

The transition between media can significantly alter the distribution of light propagated throughout a scene. The light transport equation (LTE), which is also known as the rendering equation [Kajiya 1986], gives this relation,

$$L(\mathbf{x}, \omega_o) = L_e(\mathbf{x}, \omega_o) + \int_{\mathcal{S}^2} f(\mathbf{x}, \omega_i, \omega_o) L_i(\mathbf{x}, \omega_i) |\mathbf{n}(\mathbf{x}) \cdot \omega_i| d\omega_i. \quad (2.5)$$

It bears similarity to the RTE, with the Bidirectional Scattering Distribution Function (BSDF) $f(\mathbf{x}, \omega_i, \omega_o)$ defining the amount of incident light $L_i(\mathbf{x}, \omega_i)$ irradiated over the projected area $|\mathbf{n}(\mathbf{x}) \cdot \omega_i|$, that is scattered towards the outgoing direction ω_o , and combined with light directly emitted from the light source $L_e(\mathbf{x}, \omega_o)$.

Both the RTE and LTE can be expressed as Neumann series. The result of the respective equation at a point in space is applied as incident radiance at another point to compute the propagated radiance, while also taking into account transmittance along the path segment connecting the two path-vertices. Integrating this infinite series for most practical purposes is not possible in closed form. Many modern rendering systems instead solve this problem iteratively during the image rendering process or precompute parts of the integral in an optimized data structure and spherical basis functions. The most common way, employed in the offline rendering community, to solve this equation is by Monte Carlo numerical methods. In essence, integration is evaluated based on a set of random samples drawn from a set of distributions and weighted respectively to avoid introducing bias [Veach 1998]. The final estimate approaches the true value of the integral proportionally to the number of samples.

2.1.2 Surface scattering

The differential form of the light transport equation (LTE) can be used to derive the definition of a BSDF,

$$f(\mathbf{x}, \omega_i, \omega_o) = \frac{dL(\mathbf{x}, \omega_o)}{dE(\mathbf{x}, \omega_i)} = \frac{dL(\mathbf{x}, \omega_o)}{L_i(\mathbf{x}, \omega_i) |\mathbf{n}(\mathbf{x}) \cdot \omega_i| d\omega_i}, \quad (2.6)$$

which is geometrically illustrated in Fig. 2.1. The earliest known source of this relationship is by Nicodemus [Nicodemus 1970]. His work examines this relation from the perspective of reflectance. Later works redefine it for transmissive materials [Dereniak et al. 1982] and finally combine it in a single-scattering function [Bartell et al. 1981]. The practicality of employing a certain BSDF in rendering frameworks is often evaluated on its reciprocity, energy conservation and availability of efficient sampling techniques.




Figure 2.1: Spatially Varying Bidirectional Reflectance Distribution Function (SV-BRDF) parameterization on the hemisphere according to the spatial position \mathbf{x} and the infinitesimal incident $d\omega_i$ and outgoing $d\omega_o$ light cones.

Reciprocity Helmholtz is often attributed for this principle [Helmholtz 1856]. In the case of reflection, the incident and outgoing light direction can be interchanged in the definition of BSDF without altering the value $f_r(\omega_i, \omega_o) = f_r(\omega_o, \omega_i)$. Refractive media requires an additional factor to be applied to compensate for ray foreshortening based on the refractive index $f_t(\omega_i, \omega_o)/\eta_o^2 = f_t(\omega_o, \omega_i)/\eta_i^2$. Finally, in the case of the phase function

in radiative transfer, reciprocity of anisotropic media is achieved by applying the appropriate scattering coefficient $\sigma_s(\omega_i) f_p(\omega_i, \omega_o) = \sigma_s(\omega_o) f_p(\omega_i, \omega_o)$. Taking into account these relations enables interchangeably sampling from distributions according to the camera pupil and light source surface and successfully combining them in a single integrator [Veach 1998]. Certain data-driven models might violate this requirement because they encode many phenomena within their structure and instead focus on a faithful and compact representation.

Energy conservation When accounting for all possible paths, it is crucial to guarantee that no new energy is introduced by the BRDF function ($\int_{\Omega^+} f_r(\omega_i, \omega_o) |\mathbf{n}(\mathbf{x}) \cdot \omega_i| d\omega_i \leq 1$). Complete energy preservation is not guaranteed by this requirement and many single-scattering BRDF models suffer from high loss of energy when certain parameter values are used. This leads to albedo coefficients higher than one when fit to real-world data. The correct way to account for the lost energy is to simulate the complete light transport on the surface [Heitz et al. 2016b]. However, the performance of this approach makes it unpractical for most applications. One heuristic approach to alleviate the issue is to apply additional energy compensations terms, such as the matte term by Kelemen and Szirmay-Kalos [2001].

Sampling Generating samples according to the BSDF function is one of the Monte Carlo algorithms often employed to reduce variance. It is most efficient for peaky scattering functions. For many close to diffuse and glossy surfaces it can be more beneficial to sample according to the light source geometry or the environment [Shirley et al. 1996]. Numerically solving direct illumination is highly unpractical for real-time and interactive rendering applications because it often requires a significant number of samples to reach noise-free images. Instead, closed-form approximation of the integrals over the geometry of the light source are used, as shown in Fig. 2.2. They can be evaluated more efficiently in constant time [Arvo 1995; Lecocq et al. 2016; Heitz et al. 2016a; Heitz and Hill 2017; Dupuy et al. 2017; Velinov et al. 2018b]. Similarly, the problem can be inverted to account for the spatial integral over the projected camera pixel footprint [Yan et al. 2014; Yan et al. 2016; Velinov et al. 2018b]. *Ray differentials* [Igehy 1999] are used to compute the pixel footprint by casting additional rays to determine the projected ellipse on the surface.




Figure 2.2: BRDF can be modified to account for the angular extent of the light source and the spatial extent of the projected camera pixel footprint under (a) the far field approximation and separability assumption. The exact solution would cover (b) the paired integral over the angular extent of the camera and light source with respect to the covered surface, which does not have any known analytic solutions in the general case.

2.1.3 Volume scattering

The amount of light in-scattered within a cross section of a volume is defined by the scattering coefficient and the phase function (scattering kernel) $\sigma_s(\mathbf{x}, \omega_i) f_p(\omega_i, \omega_o)$. Similarly to the definition of a BRDF, phase functions are required to obey energy conservation,

$$\int_{\mathcal{S}^2} f_p(\omega_i, \omega_o) d\omega_i = 1. \quad (2.7)$$

They define the angular relation between the incident radiance and the radiance scattered within a cross section of a volume. Usually the functions are defined in rotation-invariant fashion where no dependence on the global coordinate frame exists,

$$f_p(\omega_i, \omega_o) = \frac{dL_s(\mathbf{x}, \omega_o)}{dL_i(\omega_i)}. \quad (2.8)$$

In computer graphics one popular choice is to use the Henyey-Greenstein phase function [Henyey and Greenstein 1941] which is defined by a single *mean cosine* parameter g ,

$$f_{\text{HG}}(\omega_i, \omega_o) = \frac{1}{4\pi} \frac{1 - g^2}{(1 + g^2 + 2g(\omega_i \cdot \omega_o))^{3/2}}. \quad (2.9)$$




Figure 2.3: Participating media composed of microscopic flakes with (a) isotropic and (b) anisotropic distribution.

Recent works in the graphics community introduced models for participating media composed of anisotropic particles, also known as microflake models [Jakob et al. 2010; Zhao et al. 2011; Heitz et al. 2015]. An example of such medium is shown in Fig. 2.3. Microflake models are defined with respect to the microscopic geometry and cover all volumetric parameters that are redefined to be angularly dependent. Both scattering and extinction are defined with respect to the projected area of the microflake $\sigma(\omega)$,

$$\sigma_t(\mathbf{x}, \omega) = \rho \sigma(\mathbf{x}, \omega) \quad (2.10)$$

$$\sigma_s(\mathbf{x}, \omega) = \alpha \rho \sigma(\mathbf{x}, \omega), \quad (2.11)$$

where the volume density ρ affects both scattering and extinction, and the single-scattering albedo α defines the amount of light scattered at each event. The actual projected area is defined based on the distribution of normals (NDF) $D(\omega_h)$,

$$\sigma(\omega) = \int_{\Omega} D(\omega_h) \langle \omega \cdot \omega_h \rangle d\omega_h. \quad (2.12)$$

The phase function of the microflake distribution is derived with respect to the phase function of an individual microflake surface element $p(\omega_i, \omega_o)$,

$$f_p(\omega_i, \omega_o) = \frac{1}{\sigma(\omega_i)} \int_{\Omega} p(\omega_i, \omega_o) D(\omega_h) \langle \omega_i \cdot \omega_h \rangle d\omega_h, \quad (2.13)$$

where $\langle \omega_i \cdot \omega_h \rangle$ expresses the clamped to positive values cosine between the incident light vector ω_i and microscopic normal vectors,

$$\omega_h = \frac{\omega_i + \omega_o}{|\omega_i + \omega_o|}. \quad (2.14)$$

For most practical purposes, it is assumed that microflakes are mirror-like $p(\omega_i, \omega_o) = \frac{\delta_{\omega_h}(\omega_h)}{4|\omega_i \cdot \omega_h|}$, where the relationship is defined by a Dirac delta function $\delta_{\omega_h}(\omega_h)$ that is equal to $+\infty$ at 0. One example of a microflake distribution which provides analytic expressions for all volumetric parameters is the Symmetric GGX (SGGX) distribution [Heitz et al. 2015],

$$f_p(\omega_i, \omega_o) = \frac{D(\omega_h)}{4\sigma(\omega_i)} \quad (2.15)$$

$$D(\omega_h) = \frac{1}{\pi \sqrt{|S_{\text{GGX}}|} (\omega_h^T S_{\text{GGX}}^{-1} \omega_h)^2}, \quad (2.16)$$

where S_{GGX} is a covariance matrix incorporating a basis function and the respective standard deviations along each axes. The well defined relationship between geometry and appearance enabled the original paper to derive efficient ways for downsampling volumes. In Chapter 3 this relationship is redefined as a BRDF and used to fit the reflectance of pile fabrics.

2.1.4 Wave optics

The wave based modeling of light traces its origin to the treatise *Traite de la Lumiere* by Huygens in 1690. Its essence lies in the idea that light propagates as waves where each point on the wavefront can act as a source of spherical waves. This principle explained the diffraction phenomenon observed by earlier scholars who documented light forming periodic patterns behind a set of apertures comparable in size to the wavelength [Cecchini and Pelosi 1990; Hooke 2005; Marsh 2002]. Later, Fresnel and Kirchhoff proposed specific theoretic equations that represent this phenomenon. In modern physics, the most complete equation that is able to express the emerging electric field from a small aperture both in far-field and near-field is the general form of the Rayleigh-Sommerfeld formula [Sommerfeld 1954],

$$U(\mathbf{x}) = \frac{1}{i\lambda} \int_{\mathbb{R}^2} U_i(\mathbf{x}') \left(\frac{1}{kq} - i \right) \frac{\exp(ikq)}{q} (\mathbf{n} \cdot \omega) d\mathbf{x}', \quad (2.17)$$

where the relationship between the electric field emerging from the aperture $U_i(\mathbf{x}')$ and the outgoing electric field $U(\mathbf{x})$ is expressed with respect to the

wavenumber $k = 2\pi/\lambda$, the distance between the points ($q = |\mathbf{x} - \mathbf{x}'|$) on the diffraction grating \mathbf{x}' and the illuminated screen \mathbf{x} , and the incident light vector $\omega = \frac{\mathbf{x} - \mathbf{x}'}{|\mathbf{x} - \mathbf{x}'|}$. The main assumptions introduced by this equation are that light is passing through a thin opaque screen. Methods for computing diffraction behind thick screens do exist in the field of physics [Feit and Fleck 1978; Harrington 1993] and even more exact methods based on discretizing the Maxwell equations [Musbach et al. 2013] were recently explored in the field of graphics. However, the practicality of these approaches is quite limited when considering integrating them in existing image rendering systems. We will continue our discussion by considering only equations related to Eq. 2.17 and the scalar diffraction theory which are proven to integrate easily within existing Monte Carlo path tracing frameworks. One simplification of Eq. 2.17 that is often employed is assuming distant illumination $q \gg \lambda$,

$$U(\mathbf{x}) = \frac{1}{i\lambda} \int_{\mathbb{R}^2} U_i(\mathbf{x}') \frac{\exp(ikq)}{q} (\mathbf{n} \cdot \omega) d\mathbf{x}'. \quad (2.18)$$

This relation can be re-expressed as a Fourier transform $\mathcal{F}\{\cdot\}_\xi$ of the diffraction grating [Harvey et al. 2000] to relate the emerging electric field with the *diffracted radiance*,

$$L(\xi) = \frac{\lambda^2}{A_c} |\mathcal{F}\{U_i(\hat{\mathbf{x}}') \exp(i2\pi(\xi \cdot \mathbf{x}'))\}_\xi|^2 (\mathbf{n} \cdot \omega), \quad (2.19)$$

where the equation is re-expressed in direction cosine domain ξ by considering a scaled coordinate system $\hat{\mathbf{x}} = \mathbf{x}/\lambda$. In his work, Werner et al. [2017] redefined this expression w.r.t. a transfer function $\mathcal{T}_s(\mathbf{x})$ by considering uniform illumination of the diffracting aperture and derived a BRDF model for modeling microscopic defects on worn surfaces,

$$f_r(\xi) = \frac{1}{A_c \lambda^2} |\mathcal{F}\{\mathcal{T}_s(\mathbf{x})\}_\xi|^2. \quad (2.20)$$

Motivated by the Cittert–Zernike theorem [Goodman 1996], the BRDF was further modified to account for the spatial coherence, when combining contributions from individual spatial structures, by applying a spatial weighting centered around the shaded position on the surface \mathbf{x}_0 ,

$$\mathcal{G}(\mathbf{x}) = \exp\left(-\frac{|\mathbf{x} - \mathbf{x}_0|^2}{2\sigma^2}\right) \quad (2.21)$$

$$f_r(\xi) = \frac{1}{A_c \lambda^2} |\mathcal{F}\{\mathcal{T}_s(\mathbf{x}) \mathcal{G}(\mathbf{x})\}_\xi|^2, \quad (2.22)$$

where the standard deviation of the coherence area was chosen based on values from literature, representing the average of common illumination scenarios $\sigma = 10 \mu\text{m}$ [Divitt and Novotny 2015; Mandel and Wolf 1995]. The final area affected by interference is proportional to the radius of coherence $A_c = \pi\sigma^2$. Scratches in general cover only parts of the surface. To seamlessly combine response from the smooth material ($\mathcal{T}_{\text{base}}(\mathbf{x})$) and individual scratch elements, parts of the transfer function must be masked out ($\mathcal{T}_{\text{mask}}^{(k)}(\mathbf{x})$) and replaced by the scratch transfer function ($\mathcal{T}_{\text{scratch}}^{(k)}(\mathbf{x})$),

$$\mathcal{T}_s(\mathbf{x}) = \mathcal{T}_{\text{base}}(\mathbf{x}) - \sum_k \left[\mathcal{T}_{\text{mask}}^{(k)}(\mathbf{x}) + \mathcal{T}_{\text{scratch}}^{(k)}(\mathbf{x}) \right]. \quad (2.23)$$

Following multiple Fourier transform rewritings, which are going to be omitted for brevity, the BRDF takes its final form,

$$f_r(\mathbf{x}, \xi) = \frac{F}{A_c \lambda^2} |\mathcal{B}(\xi) - \mathcal{S}(\xi)|^2, \quad (2.24)$$

where the relationship further incorporates the Fresnel term F and it is split in a base response $\mathcal{B}(\xi)$ and scratch response $\mathcal{S}(\xi)$ function. Note that the cosine term is correctly omitted compared to the original paper to match the definition of BRDF used in computer graphics; even though, many renderers fold it inside the implementation of the BRDF evaluation function. Further details regarding this particular BRDF are discussed in Chapter 4.

2.2 Acquisition

Acquiring the appearance of materials covers the choice of a proper device and calibration techniques required to establish the spatial relationship between objects in the setup and radiometric calibration of the imaging device and light source intensities.

2.2.1 Appearance acquisition devices

Multiple devices were proposed over the years to measure the appearance of materials, starting with gonireflectometers whose measurement principles were described by Nicodemus [1970]. The basic premise is that a single camera and light source are moved relative to a sample. Various setups have

been developed with different elements of the setup fixed and sensors that are able to measure only homogeneous BRDF [Hsia and Richmond 1976] or Bidirectional Texture Function (BTF) representation of reflectance [Dana et al. 1999]. An evolution of this setup is to incorporate a hemispherical gantry with multiple light sources [Malzbender et al. 2001] and cameras [Müller et al. 2004]. One such device [Schwartz et al. 2013] is used to capture a BTF in Chapter 3. In the case of face acquisition light domes are constructed to surround an entire person and sometimes they are close to a spherical shape. Historically, these devices were first used for relighting a person for video compositing purposes [Wenger et al. 2005] and later for skin reflectance acquisition [Weyrich et al. 2006]. We used a similar light dome [Beeler 2012] to acquire the appearance in Chapter 5 with external cameras mounted on tripods. This decoupling allowed greater degree of freedom when setting up the cameras to point towards the person. More details regarding other types of devices are included in the surveys by Guarnera et al. [2016] and Klehm et al. [2015].

2.2.2 Perspective camera model

Most calibration techniques do not consider the lens geometry and instead use a very simple perspective projection model that can be optionally augmented to account for distortions. Note that this kind of camera cannot be built in the physical world because of diffraction limitations. The closest physically possible device is *camera obscura*, where light enters a chamber through a finite aperture and it is subsequently projected on a screen. *Pinhole camera* is another term used to describe this model. The projection operation is defined as follows:

$$\mathcal{P}(\mathbf{X}) = \left(\begin{bmatrix} \mathbf{X}^{(x)} \\ \mathbf{X}^{(y)} \end{bmatrix} \cdot \mathbf{F} \right) \frac{1}{\mathbf{X}^{(z)}} + \mathbf{C}, \quad (2.25)$$

where the point in world space \mathbf{X} is multiplied by the focal lengths \mathbf{F} and divided by the depth $\mathbf{X}^{(z)}$ to account for perspective distortion, and the optical centers \mathbf{C} are applied to offset the image in screen space. Accounting for radial distortion can be done by using the following function,

$$\mathcal{P}_{\text{radial}}(\mathbf{x}_{\mathcal{P}}) = \mathbf{x}_{\mathcal{P}} \left(1 + \mathbf{K} \cdot \begin{bmatrix} |\mathbf{x}_{\mathcal{P}}|^2 \\ |\mathbf{x}_{\mathcal{P}}|^4 \\ |\mathbf{x}_{\mathcal{P}}|^6 \end{bmatrix} \right), \quad (2.26)$$

where \mathbf{K} is a vector of camera distortion coefficients applied to the screen coordinates \mathbf{x}_p .

2.2.3 Geometric calibration

The acquisition setup must be calibrated before capturing images. Both the position and geometry of the light sources and cameras must be determined with respect to the subject. *Camera resectioning* or *geometric camera calibration* are the two technical terms used for describing the process of determining the camera position and orientation in space and its intrinsic parameters: focal length and lens distortion. The common way of extracting these parameters is by imaging *fiducial markers* with a known pattern and detecting individual points from different camera perspectives to jointly estimate all parameters. Two approaches are used to calibrate the setups used within this thesis. The first method [Lourakis and Argyros 2009] operates on an unknown set of 3D points and solve the problem in least squares sense by minimizing the sum of the reprojection errors,

$$\mathbf{X}^+ = \arg \min_{\mathbf{X}} \sum_i \sum_j |\mathbf{x}_p^{(i,j)} - \mathcal{P}(\mathbf{X}^{(i)})|^2, \quad (2.27)$$

where the optimized points \mathbf{X} are projected ($\mathcal{P}(\mathbf{X})$) in screen space and compared against the points detected in the images $\mathbf{x}_p^{(i,j)}$. Numerical solution is computed using the algorithm by Levenberg-Marquardt (cf. Sec. 2.3). The second method [Garrido-Jurado et al. 2014] further includes a sophisticated technique for generating *fiducial markers* that ensure correct detection of individual features across images, even when parts of the markers are occluded.

2.2.4 Light source registration

Most hemispherical gantries have a fixed geometry that is known in advance or it can be easily measured by an external distance measurement device. By imaging a mirror sphere the reflections of each individual light source can be established by switching them on in a sequence. Establishing the connections between a point on the screen and a reflection vector can be accomplished after detecting the sphere geometry based on least squares optimization of the projected compared to the imaged circular outline. Reflected vectors are computed from each camera view according to the normal vector \mathbf{n} at each

highlight on the sphere,

$$\omega_i = \omega_o - 2\mathbf{n}(\mathbf{n} \cdot \omega_o). \quad (2.28)$$

They are used to solve a linear system to find the position of the light sources $\mathbf{P}_{\text{light}}^{(i)}$ in world space in least squares sense,

$$\begin{aligned} \mathbf{M}_M &= \sum_j (\mathbf{I} - \omega_i^{(j)} \otimes \omega_i^{(j)}) \\ \mathbf{q} &= \sum_j (\mathbf{I} - \omega_i^{(j)} \otimes \omega_i^{(j)}) \mathbf{p}_s^{(j)} \\ \mathbf{P}_{\text{light}}^{(i)} &= (\mathbf{M}_M \mathbf{M}_M^T)^{-1} \mathbf{M}_M^T \mathbf{q}, \end{aligned} \quad (2.29)$$

where the individual incident light vectors $\omega_i^{(j)}$ trace lines through the point on the sphere $\mathbf{p}_s^{(j)}$ in world space for each camera view j . More details regarding this derivation can be found in Traa [2013]. The positions derived from this linear system can be very imprecise when a low number of cameras are used. Therefore, computing the rigid transformed of the prior-known light source positions to the ones computed with the above-stated equation is preferred. Computing the rigid transform in world space can be similarly achieved by solving a linear system by first centering the positions,

$$\begin{aligned} \hat{\mathbf{P}}_{\text{light}} &= \frac{1}{N_{\text{light}}} \sum_i \mathbf{P}_{\text{light}}^{(i)} \\ \hat{\mathbf{P}}_{\text{meas}} &= \frac{1}{N_{\text{light}}} \sum_i \mathbf{P}'_{\text{meas}}^{(i)}, \end{aligned} \quad (2.30)$$

and then applying SVD [Golub and Kahan 1965] to extract the rotation and translation transformations,

$$\begin{aligned} \mathbf{H} &= (\mathbf{P}'_{\text{meas}} - \hat{\mathbf{P}}_{\text{meas}}) \otimes (\mathbf{P}_{\text{light}}^{(i)} - \hat{\mathbf{P}}_{\text{light}}) \\ [\mathbf{U}, \mathbf{S}, \mathbf{V}] &= \text{SVD}(\mathbf{H}) \\ \mathbf{R} &= \mathbf{V} \mathbf{U}^T \end{aligned} \quad (2.31)$$

$$\mathbf{T} = -\mathbf{R} \hat{\mathbf{P}}_{\text{meas}} + \hat{\mathbf{P}}_{\text{light}}. \quad (2.32)$$

The resulting transformation can be further refined by optimizing the re-projection error of the light source positions projected on the sphere and removing the outliers to get closer fit to the prior-known points.

2.2.5 Light source intensities estimation

To ensure that all images are illuminated equally in the final images, a reflectance standard (chrome or spectralon sphere) has to be imaged. These are often manufactured to resemble mirror-like ($f_r = \frac{\delta(\omega_h \cdot \mathbf{n})}{|\omega_i \cdot \omega_o|}$) or diffuse Lambertian surface ($f_r = 1/\pi$). The true light intensities are directly extracted for mirror-like surfaces and the cosine to the surface is divided out for Lambertian surface. Further corrective terms account for the Fresnel term and albedo of the reflectance standard. The resulting values can be used to correct the final acquired images by dividing out their RGB values. In the virtual scenes that are later used for inverse rendering, light sources are specified with intensity equal to the squared distance.

2.2.6 Photometric calibration

The final acquired images must be radiometrically linear to match the rendering process. Eliminating the camera response curve is achieved by imaging a color checker in exposure series with times that are doubled between each shot. Additionally, dark frames with the lens cap put on must be taken to subtract the effects of dark current (sensor noise). Two algorithms are most often used to estimate response curves: Debevec and Malik [1997] and Robertson et al. [2003]. All cameras also need to have consistent colors. By using the combined High Dynamic Range (HDR) images of the color checker, the correspondence with the real color values can be established. Methods to transform the color space vary between simple color balance with an additional bias term [Ilie and Welch 2005] and solving a linear system or polynomial that gives the corrective terms.

2.3 Inverse methods

Inverse methods are used to recover parameters, given a reference image. *Inverse rendering* solves this problem by minimizing the distance between a rendered image and a reference photograph. Distance metrics in most cases are defined by the l^p -norm,

$$|\mathbf{x}|_p := \left(\sum_i |x_i|^p \right)^{1/p}, \quad (2.33)$$

defined with respect to the vector of values $\mathbf{x} = [x_1 \dots x_n]$ and $p \geq 1$. The Euclidean norm ($p = 2$) and the taxicab norm ($p = 1$) are the most popular metrics employed to solve inverse problems. Each individual value x_i is called a *residual* and *cost function* $\psi(\chi)$ by some authors. The distance metric is also called a *loss function* $\rho(\psi)$. Distance metrics can be applied additionally in parameter space to penalize certain behavior based on prior knowledge about the problem. These constraining terms are known as *regularization*. Given these two functions, an optimization problem can be defined to estimate the parameters of a distribution,

$$\chi^+ = \arg \min_x \rho(\psi(\chi)). \quad (2.34)$$

For most non-linear equations a closed-form solution does not exist. Instead, numerical algorithms are used to iteratively solve the equation. The most basic approach is Newton's method that is based on iteratively following the gradient of the loss function,

$$\chi_{n+1} = \chi_n - \gamma \nabla \rho(\chi_n), \quad (2.35)$$

where γ represents a small number that can be derived by a simple line search along the gradient for a value that improves the *loss function* or a more sophisticated algorithm. Another algorithm for solving only l^2 -norm problems is Gauss-Newton,

$$\chi_{n+1} = \chi_n - (\mathbf{J}_\chi^T \mathbf{J}_\chi)^{-1} \mathbf{J}_\chi^T \psi(\chi), \quad (2.36)$$

where each step is defined by the pseudoinverse of the Jacobian matrix $\mathbf{J}_\chi = \frac{d\psi(\chi)}{d\chi}$. Inverting the matrix is not always possible. Levenberg-Marquardt [1944; 1963] combines these two algorithm by using a damping term $\gamma \mathbf{I}$ to derive a next step at each iteration,

$$\chi_{n+1} = \chi_n - (\mathbf{J}_\chi^T \mathbf{J}_\chi + \gamma \mathbf{I})^{-1} \mathbf{J}_\chi^T \psi(\chi), \quad (2.37)$$

where γ can be successively increased until an inverse matrix is found or a more sophisticated method can be used. More in-depth discussion of numerical optimization algorithms is provided in the book by Nocedal and Wright [2006].

Another set of numerical optimization algorithms was developed to deal with noisy problems. Many of them gained popularity in the machine learning community. One such algorithm is Adam [Kingma and Ba 2014]. The

basic premise is that a running average is applied on the gradient and squared gradient which roughly corresponds to the first \hat{m} and second \hat{v} moment estimation,

$$\begin{aligned}\hat{m}_n &= b_1 \cdot \hat{m}_{n-1} + (1 - b_1) \cdot \nabla \rho \\ \hat{v}_n &= b_2 \cdot \hat{v}_{n-1} + (1 - b_2) \cdot \nabla \rho^2,\end{aligned}\tag{2.38}$$

with $b_1, b_2 \in [0, 1)$ being a set of coefficients controlling the exponential decay rate of these moving averages.

Since the algorithm has initialization bias in its initial form, additional corrective terms are applied on the first $\hat{m}'_n = \hat{m}_n / (1 - b_1^n)$ and second $\hat{v}'_n = \hat{v}_n / (1 - b_2^n)$ moment. Finally, the algorithm applies the update step,

$$\chi_n = \chi_{n-1} - \alpha_l \frac{\hat{m}_n}{\sqrt{\hat{v}_n + \epsilon}},\tag{2.39}$$

where a small value ϵ is used to prevent singularities and the learning rate is determined by a coefficient α_l . Choosing an optimal set of hyperparameters (b_1, b_2, α_l) can be done manually or by a hyperparameter optimization algorithm [Bergstra et al. 2011].

Later Chapters 3 and 5 provide examples of using numerical optimization algorithms for estimating optical properties of surfaces and volumes.

CHAPTER 3

An Interactive Appearance Model for Microscopic Fiber Surfaces

Preface Surface reflectance in computer graphics is often represented by statistical models that express the relation between the incident and outgoing light. These models are built around the notion of local flatness and distant illumination. Acquisition of their appearance is done with devices that measure the reflectance over the upper hemisphere by ignoring the effects of light transmission.

Light reflectance of pile fabrics is the subject of this chapter. The presence of loose fibers in these weave assemblies makes them disobey the notion that materials can be model by a simple acquisition of reflectance by flat samples in a single state. Their surface reflectance on macroscopic scale is affected by the orientation of each individual fiber. The main contribution described by this chapter is an acquisition and visualization techniques for modeling the tactile interaction with this class of fabrics. A low-dimensional modeling approach is used that can be integrated in real-time rendering systems with minimum cost. Furthermore, the runtime for parameter estimation is relatively low and it can better generalize to novel viewpoints without heavy regularization during the optimization process. These simplifications, however, impact certain aspects of the digital reproduction and its applicability. High-frequency surface detail is removed by this parameterization and mesoscopic scale inter-reflection and shadowing are encoded in the BRDF; thus making

them unsuitable for fabricating surfaces with the same optical parameters. Macroscopic light interaction defined by global illumination is not handled. Subsurface and translucency cannot be directly acquired by this kind of acquisition setup.

A later chapter in the thesis discusses the importance of subsurface modeling and a more elaborate framework that can extract the appearance of materials composed of multiple homogeneous layers in the presence of strong global illumination effects. The information provided in this chapter is as published in Velinov and Hullin [2016].

Abstract Modeling and rendering the appearance of fabrics remains one of the most challenging research topics in computer graphics. Today’s most advanced models are based on volumetric fiber distributions, obtained from advanced measurement techniques like micro-CT, and only renderable through expensive volume rendering approaches. In this paper, we propose an analytical BRDF model for pile fabrics, i.e., textiles like velvet, plush or Alcantara, that are characterized by open-ended strands of yarn. A fascinating property of many such materials is that they “memorize” tactile interaction and that their appearance depends on the last direction in which the surface was brushed. Our reflectance model, inspired by recent microflake models, links the microscopic structure of a material to the interaction that caused that structure, and to the resulting macroscopic reflectance behavior. We present an end-to-end pipeline for fitting the parameters of our model to measured reflectance data, for manipulating fiber distributions using tactile input, and rendering the resulting spatially varying surface appearance in real time.

3.1 Introduction

Textile materials are a central ingredient in industries as diverse as clothing, furniture or automotive. Capturing and faithfully recreating the characteristic visual appearance of fabrics is therefore an important goal and an active field of research in computer graphics. So far, most research has focused on woven or knitted fabrics whose structure is tightly held together and hence static. This property enables relatively straightforward appearance acquisition by capturing images of a material sample under different lighting and viewing conditions.

In contrast, pile fabrics (velvet, alcantara and plush) contain loose fibers which causes their appearance to depend on the history of interaction. Such

materials are visually intriguing, because their tilted fibers can exhibit complex specular reflectance that is not aligned with the macroscopic cloth surface. More familiar and fascinating, however, is the experience of being able to “paint” on a carpet, pillow or car seat using one’s fingers. The goal of this work is to recreate this experience, which implies the need for three major ingredients: firstly, a reflectance model capable of recreating the interaction-dependent appearance of pile fabrics; secondly, a method to derive the model parameters from physical samples; and finally a rendering algorithm, along with a user interface, for displaying and manipulating virtual objects in real time.

The need for a real-time solution limits the feasible techniques that can be applied to solve the problem. Spatially varying bidirectional reflectance distribution functions (SVBRDFs) based on microfacet distributions are considered a physically correct way of modeling the appearance of opaque surfaces. They are built under certain assumptions about the surface geometry; most importantly, that the surface is comprised of microscopic mirrors distributed relative to the surface normal.

Like all textile materials, however, pile fabrics are not opaque surfaces. Rather, light that interacts with them is scattered most intricately by a volume filled by microscopic fibers. Consequently, the state-of-the-art in fabric rendering is held by sophisticated rendering schemes that treat the material as a volume of scattering microflakes [Jakob et al. 2010] or even explicit fibers. However, these techniques are computationally expensive, and modeling the interaction between external forces and every fiber would be unfeasible. Rather, we borrow ideas from these models and use them to derive a novel SVBRDF representation that approximates not only the reflectance caused by microflake distributions, but also the connection between these distributions and the way that the material has been interacted with.

To validate our approach, we measure the appearance of an alcantara sample and fit the proposed model to the obtained data. We use hemispherical bidirectional texture function (BTF) measurements, which also include non-local effects like interreflection, parallax or shadowing. To fit the model to the measured BTF data we perform a principal component analysis (PCA) on the surface distribution of normals (NDF) to derive an initial guess which is then refined using non-linear least squares optimization.

The main contributions of our work are thus threefold:

- We introduce a SVBRDF model for the local appearance of pile fabrics

and the way that this appearance is affected by tactile interaction. The shading model is based on a distribution of microscopic fibers that is in general not aligned with the macroscopic material surface.

- We develop a complete pipeline for measuring, fitting and simulating materials by example.
- We demonstrate our model in an application that allows the user to interact with an object made of virtual pile fabric in real time.

The rest of the paper is structured as follows. In Section 3.2 we discuss the class of pile fabrics and how to approach the problem of modeling them. We discuss related work in Section 3.3. Then we introduce our pipeline in Section 3.4 by making a breakdown of each of its elements. We apply our approach to measured data and compare the results in Section 3.5. Section 3.6 discusses the strengths and weaknesses of the proposed approach. We conclude the paper with an outline of the achieved results in Section 3.7.

3.2 Background




Figure 3.1: The surface of Alcantara is comprised of microscopic fibers whose orientation is dependent on the history of previous tactile interactions. The distribution of fibers gives its characteristic anisotropic appearance.

Velvet is woven on a special loom that weaves two thicknesses of the material at the same time. They are cut afterwards to create the characteristic




Figure 3.2: Our SVBRDF fitting pipeline. We use hemispherical bidirectional texture functions (BTF) measured by a camera dome as input to our pipeline. Each operation is performed per pixel in the main parameter estimation phase. Post processing is performed on the whole texture based on information about the neighboring similar regions. The output of the pipeline is a set of texture files containing SVBRDF parameters.

pile effect on each separate piece. Plush is produced in similar manner as velvet. In some weaving patterns the pile fibers might be pulled from the base fabric without being cut. Alcantara is made of fine polyester microfibers. The polyester fibers are cut into small segments and then processed by a carding machine into a soft felt. The cloth-like appearance is achieved by a process of needle punching and treatment with special agents. The most characteristic property of these materials is that they quickly restore their geometry to a stationary state close to the initial after tactile interaction. However, most of the fibers are still slightly bent towards the direction of interaction (See Figure 3.1). The ability to reach quickly stationary state is highly dependent on the fiber length. Longer fibers as rule of thumb take more time to restore their shape. We limit our study to fabrics with short fibers which allows us to capture the appearance of different brushing patterns without requiring the acquisition of parameters used for dynamics simulation.

To reproduce the characteristic look of these materials we have devised a complete pipeline (See Figure 3.2). Every stage is covered - from appearance acquisition to material modeling and interactive simulation. We capture BTF of material sample brushed in eight different directions in a camera dome. We process it by applying PCA and consequently a nonlinear least squares optimization to fit the measured reflectance in microflake distribution adapted for surface rendering. Additional post processing is done to eliminate misalignment between neighboring pixels of the SVBRDF. Then we remap the differently brushed regions to certain states of the fabric. Afterwards we simulate it in a custom real-time ray tracing software running on the GPU. The simulation relies on the geometric information encoded within

the microflake distribution to perform accurate interpolation.

Our approach is suitable for material showcases which enable the user to display different objects made of this kind of materials. It enhances the perceptual qualities of the modeled material by providing direct feedback to the user who can interact with the fabric.

3.3 Related Work

We identify several strands of prior work that are closely related to our problem and approach.

Time-dependent appearance. In our application, the interaction with a material leads to a change in appearance over time, a problem that has already been explored in the context of paint drying and dust accumulation [Sun et al. 2007]. In this work, a robotic arm was used to acquire gonio-reflectometric data in regular time intervals, which was then used to fit the parameters of analytical BRDF models. The time dependency was then investigated in more detail by analyzing temporal series of measured BRDFs. In contrast, in our work we are not only interested in a time series but the more complex connection between tactile interaction with a material and its appearance. This leads to a higher-dimensional fitting problem. In that sense, our work is related to weathering simulation [Miller 1994; Chen et al. 2005; Wang et al. 2006]. The main difference here is that instead of parameters like the “exposedness” of points on the surface that affect weathering and wear, our materials respond to user interaction. To our knowledge, this effect has not been treated in computer graphics literature to date.

Appearance of fabric. Textile materials play an important role in our every-day lives and have received considerable attention in the graphics community. The models proposed in literature span a wide range. On one end are analytical BRDFs based on microfacet models [Ashikmin et al. 2000] or fiber geometry [Irawan and Marschner 2012; Sadeghi et al. 2013]. On the other end, purely image-based models, like the bidirectional texture function (BTF) [Dana et al. 1999], capture the view- and light-dependent reflectance of surface samples. The characteristic properties of fabric make it particularly suitable for volumetric scattering models that can either be isotropic [Schröder et al. 2011] or anisotropic like microflake representations [Jakob

et al. 2010; Zhao et al. 2011]. Today’s microflake models can be represented and rendered efficiently [Heitz et al. 2015]. They have been shown to produce similar rendering quality as explicit fiber models [Khungurn et al. 2015], although such models are still several orders of magnitude too expensive for real-time rendering. In our work, we use a hybrid approach. Using a BTF as input data, we fit the parameters of a custom spatially varying BRDF which, in turn, is derived from a single-layer microflake model. We derive our model from linear transport theory similarly to Hanrahan and Krueger [1993]. However, this early work relied on isotropic scattering and the empirical Henyey-Greenstein phase function while we use the microflake theory to derive an anisotropic BRDF. Unlike most existing analytical BRDF models [Ashikmin et al. 2000; Irawan and Marschner 2012; Sadeghi et al. 2013], ours allows for asymmetric or tilted microflake distributions. The only model derived under the assumptions of the microfacet theory for the purpose of simulating metallic surfaces that shares these characteristics is the ellipsoid NDF [Dong et al. 2015]. Our model decouples material parameters (gloss and color) from interaction parameters (fiber directionality), enabling the plausible manipulation of fiber distributions without leaving the appearance space of the prescribed material. Similar to other BRDF models our technique trades-off accuracy for performance, mainly, by ignoring parallax, inter-reflection and assuming simplified representations of the microgeometry. The implications of these properties are that our model would deviate from the ground-truth BTF rendering. However, our model has smaller memory footprint, requires less computation and therefore is more widely applicable.

3.4 Proposed Approach

Our solution is comprised of three major parts: reflectance model, parameter fitting to measured data and interactive simulation. We discuss each one of them in the following subsections.

3.4.1 Reflectance model

Physically correct modeling of pile fabrics under tactile interaction requires parametrization of the appearance based on the microscopic structure. Materials like Alcantara are comprised of many loose fibers that are bound together into a single fabric. This characteristic feature makes them deviate

from the assumptions of microfacet theory of small mirror-reflecting surfaces. The microflake theory is a newer approach of looking at the process of modeling the microscopic geometry. The key difference is that surfaces are replaced by volumes of oriented microscopic particles. We take a similar approach to Heitz et al. [2015] who introduced the classical Trowbridge-Reitz (GGX) model [Trowbridge and Reitz 1975] to microflake theory by extending its distribution of normals to the whole sphere. The preferred framework of microflake models up to this point was to use volume rendering techniques. This method of rendering is unsuitable for real-time rendering because it introduces a random walk step which severely reduces the convergence rate. We propose a different approach which can enable regular surface rendering techniques to be applied.

We start our derivation from the radiative transfer equation limited to single scattering. It is a fundamental assumption made when deriving closed-form analytical BRDF models in the microfacet theory. Therefore, our approach would be comparable in quality to these models. We further assume optically thick semi-infinite medium per surface element. This assumption removes the need of performing a random walk on the surface. The reflected radiance at a given point on the surface $L(\mathbf{x}, \omega_o)$ under these assumptions can be expressed as:

$$L(\mathbf{x}, \omega_o) = \int_0^\infty e^{-\sigma_t(\omega_o)t} \sigma_s(\omega_o) f_p(\omega_i, \omega_o) L_i(\omega_o, \omega_i, t) dt. \quad (3.1)$$

Light is scattered along the outgoing direction from the light source proportional to the scattering coefficient $\sigma_s(\omega_o) = k_s \rho \sigma(\omega_o)$. The $\sigma(\omega_o)$ and ρ are the projected area of the microscopic flake and the volume density, respectively. The albedo k_s determines the specular contribution in our model. Afterwards, light is exponentially attenuated along distance t at rate proportional to $\sigma_t(\omega_o) = \rho \sigma(\omega_o)$. $f_r(\omega_i, \omega_o)$ represents the phase function which in our case is a specular reflectance function using the GGX distribution of normals extended on the sphere (SGGX). $L_i(\omega_o, \omega_i, t)$ represents the in-scattered radiance. We further assume far field illumination which lets us approximate the light transport as parallel rays scattered towards the outgoing direction. Complete picture of the assumptions that we apply to compute the radiance is shown in Figure 3.3. The resulting equation for the in-scattered radiance is:

$$L_i(\omega_o, \omega_i, t) = L_i(\omega_i) \exp\left(-\sigma_t(\omega_i) \frac{\cos \theta_o}{\cos \theta_i} t\right). \quad (3.2)$$

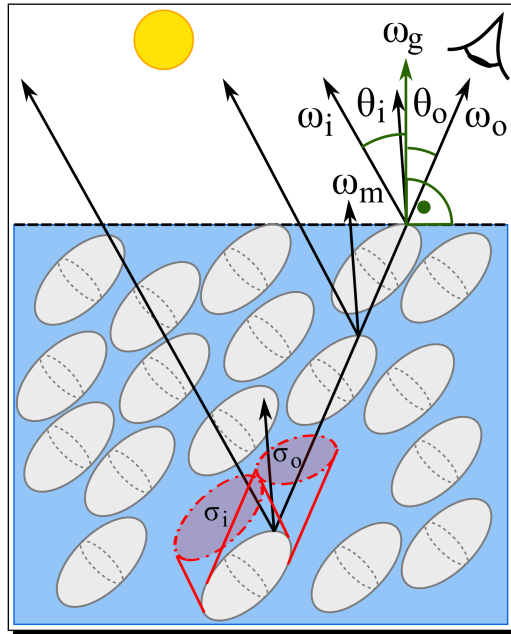


Figure 3.3: We assume homogeneous semi-infinite volume per surface pixel and far-field illumination. It enables us to use parallel incident lights along the scattering direction.

Both cosine factors are computed relative to the surface normal. They express the path difference between the incident and outgoing light. Inserting equation Eq. 3.2 into Eq. 3.1 lets us integrate the RTE under the above-stated assumptions:

$$L(\mathbf{x}, \omega_o) = k_s \underbrace{\frac{f_p(\omega_i, \omega_o) \sigma(\omega_o)}{\sigma(\omega_o) + \sigma(\omega_i) \frac{\cos \theta_o}{\cos \theta_i}}}_{f_r} L_i(\omega_i). \quad (3.3)$$

Note that the volume density ρ cancels out because it is part of the numerator and denominator. The main implication of this property is that in semi-infinite medium the appearance is unaffected by scaling of the average free path length.

The phase function that we use is the mirror-like function proposed by Heitz et al. [2015],

$$f_p(\omega_i, \omega_o) = \frac{D(\omega_h)}{4\sigma(\omega_o)}. \quad (3.4)$$

The SGGX distribution of normals $D(\omega_h)$ is dependent on the microscopic normal vector ω_h defined as the half-way vector between the incident and outgoing light directions. We refer to the original paper for more information. Note that we use the reciprocal parameters in our notation which is consistent with Jakob et al. [2010]. $D(\omega_h)$ is defined as follows:

$$D(\omega_h) = \frac{1}{\pi \sqrt{|S_{\text{GGX}}|} (\omega_h^T S_{\text{GGX}}^{-1} \omega_h)^2} \quad (3.5)$$

The covariance matrix S_{GGX} is defined as a squared product of a rotation M and scale matrix S_σ .

$$S_{\text{GGX}} = M S_\sigma S_\sigma^T M^T \quad (3.6)$$

The scale matrix is expressed in our case as a diagonal matrix composed of the projected areas along each major axis. For convenience we will express the elements along the diagonal as S_v .

$$S_\sigma := (\delta_{kn} S_v^{(k)})_{kn} \quad (3.7)$$

$$S_v^{(k)} = \sigma(\omega_k) \quad (3.8)$$

We have experimented with adding a Fresnel term based on measured values of the refractive index; however, it did not result in improvement of the overall quality compared to measured data.

The model is also extendable to multiple scattering by evaluating the complete light transport. Similar approach was applied by Heitz et al. [2016b] in the context of the microfacet theory and Hanrahan and Krueger [1993] for simulating layered materials. However, computing multiple bounces along material’s surface can be quite costly. Therefore we add a regular isotropic diffuse term to compensate for the lost energy and the residual due to the complex microscopic structure. Our final model which can be substituted in the rendering equation is as follows:

$$f_r(\omega_i, \omega_o) = k_s \frac{D(\omega_h)}{4 \cos \theta_i \cos \theta_o \left(\frac{\sigma(\omega_o)}{\cos \theta_o} + \frac{\sigma(\omega_i)}{\cos \theta_i} \right)} + \frac{k_d}{\pi}. \quad (3.9)$$

The additional $\cos \theta_i$ factor accounts for the change of normals used by RTE simulation. The k_d represents the spectrum of the reflected light by the diffuse component. It is visible that the BRDF fulfills the requirement for reciprocity, i.e. $f_r(\omega_i, \omega_o) = f_r(\omega_o, \omega_i)$.

In the special case of isotropic BRDF where $\sigma(\omega_o) = 1$, $\sigma(\omega_i) = 1$ and $D(\omega_h) = \frac{1}{\pi}$ the model reduces to the Seeliger’s Law [Seeliger 1888]. This property can be shown by rewriting equation Eq. 3.3 in the following form:

$$L(\mathbf{x}, \omega_o) = \frac{D(\omega_h) \cos \theta_i}{4(\sigma(\omega_o) \cos \theta_i + \sigma(\omega_i) \cos \theta_o)} L_i(\omega_i). \quad (3.10)$$

Similarly, if we apply the same constants in equation Eq. 3.9 we would derive the isotropic Chandrasekhar’s BRDF [Chandrasekhar 1960].

Importance sampling is essential to ensure fast convergence in ray tracing applications. The ideal importance sampling would be exactly according to the BRDF $f_r(\omega_i, \omega_o)$. However, that’s rarely possible. The most common approximation is to sample according to the distribution of normals. Our model is derived from tilted GGX distribution. It is one of the similarities of our model with the ellipsoid NDF derived in the supplemental material of Dong et al. [2015]. Ignoring the other terms lets us sample our model in similar fashion. We refer to the original work for more details.

Modeling time-varying fabrics with this model requires the interpolation of material parameters. The geometry related parameters are contained within the distribution of normals (NDF). We use the original SGGX distribution which contains a covariance matrix decomposed into rotation and standard deviation part. Interpolation can be performed linearly on each of these parameters.

3.4.2 Parameter fitting

To validate the applicability of our model for approximating the appearance of real world materials we devised a pipeline which fits parameters to measurement Bidirectional Texture Functions (BTF). Our model has two main components: diffuse and specular. Fitting the specular term can be broken down into two parts: initial guess estimation and nonlinear least squares fitting. The initial guess step is required to prevent reaching suboptimal local minimum when the next step gets applied. First the BTF is sampled at a specific point on the surface and complete slice of all available pairs of light and view directions is taken with their corresponding reflected radiance to irradiance ratio. Direct inversion of the NDF with the geometry contribution would be unfeasible for the initial guess that we would like to estimate. Therefore we extract the NDF for a perfect mirror reflection by using the following relation:

$$D_{\text{BTF}}(\omega_h) = \frac{1}{2\pi} \int_{\Omega^+} f_{\text{BTF}}(\omega_i, 2(\omega_h \cdot \omega_i)\omega_h - \omega_i) d\omega_i. \quad (3.11)$$

This equation can be solved through numerical integration. We use stratified Monte Carlo integrator for this purpose. The resulting image is stored in parabolic map with the following coordinate mapping:

$$\begin{bmatrix} x' & y' \end{bmatrix} = \begin{bmatrix} 0.5 + 0.5\frac{x}{z+1} & 0.5 + 0.5\frac{y}{z+1} \end{bmatrix} \quad (3.12)$$

Proper normalization is not essential because the SGGX NDF is scale invariant.

We extract the major axes by applying PCA on the resulting NDF to determine the rotation of the SGGX covariance matrix. Given $D_{\text{BTF}}(\omega_h) : \Omega^+ \mapsto R$ we perform a search for axes which satisfy the equation:

$$\omega_2 = \arg \max_{\omega_h \in \Omega^+} (D_{\text{BTF}}(\omega_h)) \quad (3.13)$$

$$\omega_0 = \arg \max_{\omega_h \in \Omega^+ : \omega_2 \cdot \omega_h = 0} (D_{\text{BTF}}(\omega_h)) \quad (3.14)$$

$$\omega_1 = \omega_0 \times \omega_2 \quad (3.15)$$

The only component left is the standard deviation which is proportional to the projected area on each axis. It can be extracted through a simple integration:

$$M = [\omega_0 \quad \omega_1 \quad \omega_2]$$

$$S_v = \int_{\Omega^+} (M^T \omega_h^T) D_{\text{BTF}}(\omega_h) d\omega_h. \quad (3.16)$$

Having determined the major axes and the projected area we have a good enough initial guess on the specular contribution. To derive optimal parameters for the complete BRDF we take the estimated specular term and apply nonlinear least squares optimization. We compute the residual as the difference of the discrete samples $W_{\omega_i, \omega_o}^{(k)}$ contained within the BTF and our BRDF function leading us to the following equations representing the residual:

$$g^{(k)}(W_{\omega_i, \omega_o}^{(k)}) = f_{\text{BTF}}(W_{\omega_i, \omega_o}^{(k)}) - f_r(\omega_i^{(k)}, \omega_o^{(k)}), \quad (3.17)$$

The sum of squares is expressed as:

$$\rho(W_{\omega_i, \omega_o}^{(k)}) = \frac{1}{2} \sum_{k=1}^{N_L N_V} (g^{(k)}(W_{\omega_i, \omega_o}^{(k)}))^2. \quad (3.18)$$

Given $\rho : \mathbb{R}^n \mapsto \mathbb{R}$. We must find global minimizer:

$$W_{\omega_i, \omega_o}^{(k)+} = \arg \min \{ \rho(W_{\omega_i, \omega_o}^{(k)}) \} \quad (3.19)$$

This is a popular problem in mathematical optimization and many solutions exist to solve this problem. We used Levenberg-Marquardt method for this purpose by applying it to the specular and diffuse components of the BRDF.

3.4.3 Post processing

The resulting SVBRDF from the fitting procedure has many areas where the tangent directions of neighboring pixels are pointing towards opposite directions. It is essentially caused by symmetry of the BRDF as opposite bases would generate the same appearance. Interpolating this kind of bases would result in visible artifacts. The first step to alleviate this issue is to sort the basis according to standard deviation. By applying this method we are exploiting the mostly homogeneous structure of the material. However, due to symmetry in the BRDF it is still possible to have misalignment. Therefore, we further detect the outlines of areas where this behavior occurs and label the individual regions. Afterwards, we reorient each individual region towards the dominant direction (the region covering most area). It can be further realigned, so that the dominant direction points towards the preferred direction, ensuring smooth interpolation during interaction. All of these operations don't change the overall appearance as we are exploiting symmetries in the BRDF function.

3.5 Implementation and Results

We have implemented the parameter fitting pipeline in C++. We have experimented with CUDA implementation for some parts of the pipeline. However, it was mostly memory bandwidth limited, so the overall performance improvements were fairly modest. We use separate tools for the per pixel parameter estimation and the post processing step for faster development iteration when testing different experimental techniques.

We work with BTF that was captured with 30° rotations of the camera dome turntable and 11 cameras. It was rectified as a set of images with resolution 1024×1024 and resampled with 151×151 light-view vector pairs. It was compressed in Full Matrix Factorization (FMF) format [Koudelka et al. 2003]. Our SVBRDF fitting procedure takes about 150 minutes per 200×200 patch on an Intel i7-5820k system running at 3.30 GHz. We work with chunks of 64×64 randomly selected light-view vector pairs to ensure low error. It is the leading factor affecting the current execution times. The main bottleneck is the nonlinear least squares optimization followed by the Monte Carlo integration used for estimating the NDF. We have experimented with some photometric normals techniques for estimating the initial guess of the SGGX basis axes; however, they did not yield better results than the NDF estimation technique that we proposed in this paper. We use a configuration of 4 dispatcher threads which manage information for 4 different pixels at the same time and a shared thread pool that processes the bulk of the data. Dispatcher threads assist in the data processing while they wait for completion of related tasks. This architecture ensures full utilization of the CPU.

After the initial per pixel parameter step is complete we get above 20 dB peak signal-to-noise (PSNR) ratio in the angular domain for most of the samples as can be seen in Figure 3.4. It guarantees us below 10% error for most of the surface. Increasing this ratio would require further fitting of multiple lobes and taking into account light transport related effects. It is possible to apply multiple lobes in the parameter estimation procedure; however, it would be more unstable. It should be noted that our results are based on a plane estimation and subsequent rectification of photographs which results in different quality between the samples in the angular domain. The main reasons for this quality disparity is the difference in resolution between top and grazing view angles and parallax inside the measured data. Our technique does not take parallax into account which is required to further improve the




Figure 3.4: PSNR in the angular domain for each pixel of the SVBRDF produced for Alcantara material sample. The original BTF lit from top is shown as a background. Some of the areas show distinct striped patterns which can be explained by inter-reflectance and shadowing effects.

visual quality. Standard solution to this problem is to measure geometry by using structured lighting techniques. Applying structured lighting techniques would be challenging as most of the surface detail is close to the error caused by the instability of projectors. There are also inter-reflection and subsurface scattering effects which impede exact geometry acquisition. These problems can be bypassed by working only with the top-down view of the BTF material. Example images are provided in Figure 3.8. Ground truth rendering of the original BTF with the same scene is provided in Figure 3.9.

Post processing the data is crucial to guarantee smooth interpolation across the whole surface. We have chosen eight patches which are brushed in the same direction. Essentially they represent eight different states of the material. The objective of the post processing stage is to also align them so that we can guarantee smooth transitions when tactile interaction is applied. Overall the resulting vector field is smooth after our post processing step




Figure 3.5: Orientation of tangent vectors of the SGGX basis after postprocessing. It is visible that the appearance is highly dependent on direction of material manipulation.

which is confirmed by Figure 3.5. Further confirmation of the alignment of the SGGX basis with the brushing directions is shown in Figure 3.6. The post processing step as we apply it works only on the SGGX basis vectors and the standard deviation and does not introduce any errors as it performs swapping and axis reversal operations.




Figure 3.6: Ellipsoid distributions fit to the original BTF representation. The distributions are leaning towards the brushing direction connecting appearance and microscopic structure.

We have developed an application that demonstrates the capability of our SVBRDF technique to simulate tactile interaction. We provide images rendered by our framework in Figure 3.7. We have used the real-time GPU ray tracing backend in our custom rendering framework. However, it is directly applicable in regular GPU rasterization based deferred shading by storing information about tangent space and one additional specular power component. Therefore we consider our technique as widely applicable for rendering in many scenarios. It is visible that the change of appearance when interacting with the material is handled by our approach. Our painting solution is based on software rasterization on the GPU of two-dimensional capsules with linear fall-off along the edges. Interpolation is performed linearly with spherical linear interpolation for the SGGX basis. Editing can be performed by taking into account that smoothness corresponds to lower standard deviation values and anisotropy can be boosted by reducing the value of the

smallest ellipsoid axis.




Figure 3.7: Our model can be manipulated in real time and it contains parameters that are intuitive for editing. The first image on the left is a surface whose SVBRDF was generated by our parameter fitting pipeline. The second image displays tactile interaction with the material. The third image presents simple recoloring of the SVBRDF to green color. The last image on the right presents recoloring and boosting the anisotropy which is controlled by the standard deviation.

3.6 Discussion

We are working with SVBRDF representation of a measured material. This class of techniques shares similar set of simplification related to handling parallax, inter-reflection and microgeometry.

Parallax is usually handled by geometry acquisition. In our case we work with a material with fine fiber structure which is challenging for the currently existing structured light techniques. They are negatively affected by the instability of the projected structured patterns on the surface. Inaccuracies in the rotations of the turn-table sample holder can affect the final quality of the reconstruction. Also, they are negatively affected by inter-reflection and




Figure 3.8: Higher level of detail can be achieved if the top-down view of the BTF is taken which is not affected by parallax and rectification errors. The first image on the left is a surface whose SVBRDF was generated by our parameter fitting pipeline. The second image displays tactile interaction with the material. The third presents simple recoloring of the SVBRDF to green color. The last image presents recoloring and boosting the anisotropy which is controlled by the standard deviation.




Figure 3.9: Ground truth BTF rendering. Note that due to the rectification error the effective resolution of the BTF decreases at lower viewing angles, making the appearance smoother. The problem is further emphasized by intermixing caused by the parallax error.

subsurface scattering effects. All of these factors make the exact acquisition of reconstructed geometry extremely challenging. Our research does not focus on resolving these issues; therefore, this is left as future work.

Inter-reflection can be handled by our model. Full light transport simulation is prohibitively expensive for real-time rendering of complex scenes. Another concern is that the parameter estimation would take much longer time. Therefore, within this paper we demonstrated only single scattering data fits. It is left as future study to determine the quality change when multiple scattering is applied for high quality offline rendering.

Complex microgeometry can be handled through multiple lobe simulation. The main concerns when applying multiple lobes is that the parameter estimation becomes unstable as we are working with ill-posed problem. Another concern is that applying interpolation becomes more ambiguous when simulating tactile interactions. These are legitimate concerns and we are looking forward to address in future work.

3.7 Conclusions

We introduced a SVBRDF model based on homogeneous infinite volume approximation of pile fabrics which can simulate the change of appearance based on tactile interactions. Our model is simple, efficient and easy to implement both in ray tracing and traditional rasterization graphics pipelines.

We developed a complete pipeline for measuring, fitting and simulating materials with microscopic fiber surfaces. We validated our pipeline against

measurements of Alcantara material and proven that it provides good quality for materials simulated with our low-dimensional SVBRDF model. Parallax and inter-reflection still remain a challenging problem. However, they would require introducing sophisticated methods for geometry acquisition, which are out of the scope of our research.

We demonstrated the applicability of our SVBRDF model for modeling the tactile interaction between human hand and object made of this class of materials in real time. We further showed that our model is easy to edit which makes it applicable for generating synthetic materials for the purpose of designing objects made of pile fabrics.

CHAPTER 4

Real-Time Rendering of Wave-Optical Effects on Scratched Surfaces

Preface Illumination plays a significant role in the appearance formation of materials. For example, hidden structures can be revealed by observing the reflection of extended light sources. The strife of computer graphics is to create virtual representation of greater and greater visual fidelity. Thus, the problem of proper evaluation of the reflectance with respect to the light source geometry is important both in offline and real-time rendering applications. BRDF modeling approaches often summarize the overall reflective properties of surfaces, however thin mirror-like features can violate the notion of scale that is often employed to derive these models. The scales on which these features are present are often close to the wavelength and geometrical optics fails to reproduce their iridescent behavior.

Building upon the initial work on BRDF modeling of iridescent worn surfaces by Werner et al. [2017], this chapter extends the theory to real-time applications where evaluation of all integrated terms must be done with a single sample per pixel both in spatial and angular domain. The discussion covers efficient algorithms for analytic evaluation of spherical and polygonal lights illumination of iridescent scratches that are separately anti-aliased. Both evaluation of extended light source illumination

and anti-aliasing vectorial wave-optical SVBRDFs was not covered by previous works and it is the main claim of contribution to the real-time rendering literature.

Multiple simplifications were introduced to make integration tractable. The Fresnel term is excluded from the approximations. It can be added in an artificial way by taking an average value along line segment, but it would definitely introduce errors. Shadowing is not implemented, but sampling approaches can be applied, such as Heitz et al. [2018]. Only rectangular profiles are derived analytically, however triangular are possible by pre-integrating them and fitting a curve that best represents the reflectance behavior.

Close to exact evaluation of the illumination by extended light sources has potential use both for generating synthetic imagery and inverse rendering. The content in this chapter corresponds to Velinov et al. [2018b].




Figure 4.1: We introduce a diffraction model for fine surface scratches which seamlessly integrates in modern real-time rendering pipelines. Left: Our approach reaches 30fps at 720p resolution on a notebook GPU (GTX 970M) with 10000 scratches and two light sources: a sphere and a rectangle approximating the highest intensity parts of the environment. Right: Our technique provides anti-aliased scratches for multi-scale rendering without multi-sampling. Shown here is the reflection of a triangular light source on a planar target.

Abstract The visual appearance of real-world materials is characterized by surface features across many scales and has received significant attention by the graphics community for decades. Yet, even the most advanced microfacet models have difficulties faithfully recreating materials like snow, sand, brushed metal or hair that feature scale-violating glints and speckles and defy any traditional notion of filtering and level of detail. In this work, we address

an important subset of such materials, namely metal and dielectric surfaces that are covered with microscopic scratches, e.g., from polishing processes or surface wear. The appearance of such surfaces features fine-scale spatial detail and iridescent colors caused by diffraction, and has only recently been successfully recreated. We adopt the scratch iridescence model, which is known for plausible results in offline Monte Carlo settings but unsuitable for real-time applications where extensive illumination sampling is prohibitively expensive. In this paper, we introduce an efficient technique for incoherently integrating the contributions of individual scratches, as well as closed-form solutions for modeling spherical and polygonal area light sources, and for the first time bring scratch iridescence within reach of real-time applications.

4.1 Introduction

Real-world surfaces feature geometric detail across a wide range of scales, which results in a characteristic visual appearance that computer graphics has long attempted to faithfully model, capture and reproduce. The state-of-the-art in digital material appearance is dominated by models that represent surface microgeometry by statistics like micro-facet or micro-flake distributions. This class of models is capable of closely approximating the far-field scattering properties of most real-world surfaces [Guarnera et al. 2016], and it has recently even been extended to include wave-optical effects [Holzschuch and Pacanowski 2017; Belcour and Barla 2017]. Yet, significant effort is needed to enable the recreation of fine-scale reflectance detail such as isolated glints [Yan et al. 2014; Yan et al. 2016; Jakob et al. 2014b]. In recent work, Werner et al. [2017] introduced a spatially varying bidirectional reflectance distribution function (SVBRDF) for a sub-class of microstructured materials, namely metallic and dielectric surfaces with fine enough scratches to produce iridescent reflections. Their method, to our knowledge, is the first to bridge the gap between non-paraxial wave-optical scattering theory and texture required for spatial detail in computer graphics applications. Neglected in prior wave-optical scattering models, their model correctly handles spatial variation by combining coherence theory with a vector-graphics representation for discrete scratches on a surface. In this paper, we address what we consider the most important drawback of the proposed technique: its need for extensive sampling of area light sources and pixel footprints. While the model can be evaluated with a single light sample for point sources,




Figure 4.2: Scratch iridescence model by Werner et al., rendered under point light source (1 light sample per pixel).

the resulting appearance is highly unrealistic (Fig. 4.2). Multi-sampling, on the other hand, produces well-filtered renderings in offline settings but is far too costly for real-time applications. Where precomputed radiance transfer delivers high-quality results for ordinary reflectance distributions under environment lighting [Sloan et al. 2002], highly specular glints as well as micrometer-level spatial detail prohibit the use of such approaches for the class of materials at hand. To mitigate this problem, we propose a set of analytical pre-integrations and approximations that build upon the model of Werner et al. and enable it to achieve real-time performance under area and environment lighting. In particular, our technical contributions include

- a closed-form solution of the integral over the pixel footprint visible from the camera,
- closed-form solutions for integrals over polygonal and spherical area light sources, and
- an optimized pipeline and data structure to enable efficient lookup of discrete scratch particles on the GPU.

With these improvements, our model achieves a speed-up of at least three orders of magnitude, enabling the rendering of scratch iridescence at HD resolution in real time. Our work is a promising step towards introducing these effects in current real-time rendering pipelines, and it is immediately applicable for visualization purposes in offline rendering production.

4.2 Related work

Our proposed method builds upon Werner et al.’s theory and model for wave-optical effects on worn surfaces [Werner et al. 2017], which heavily relies on Monte Carlo integration for the evaluation of the reflectance function in both spatial and angular domain. Multiple samples are cast on the footprint of each pixel in order to mitigate aliasing effects in the spatial domain; in angular domain, area and environment lights are evaluated through costly importance sampling. Here, our goal is to alleviate both these requirements. By deriving analytical pre-filtered versions of the reflectance function, we reduce the effort of evaluating the model from thousands of samples to one sample per pixel, effectively making it capable of real-time operation. In the following, we briefly review the state-of-the-art in wave-optical shading models, and provide an overview of real-time rendering techniques for extended light sources. For a historic study of general BRDF modeling, we refer the interested reader to the survey by Guarnera et al. [2016]. Further information about prefiltering different components of microfacet BRDFs can be found in the survey by Bruneton and Neyret [2012].

4.2.1 Wave-optical shading models

A conclusive overview of wave-optical scattering models for different surface classes can be found in Krywonos [2006], and fundamental principles are discussed in Goodman [1996]. Based on such models of light-surface interaction, the computer graphics community has developed a range of techniques accounting for diffraction and interference in ray-based frameworks. Stam [Stam 1999], as one of the first, proposed a diffraction model for rough surfaces. The concept of spatial coherence, which is critical for spatially resolved structures, was used by Levin et al. [2013] and Dhillon et al. [2014]. In contrast to Levin, who predicted BRDFs of lithographically structured surfaces, Dhillon used this concept to efficiently simulate diffraction effects of measured biological real-world surfaces. More recently, Belcour and Barla [2017] proposed a modified microfacet model to recreate interference from thin layers. Dong et al. [2015] derived far-field BRDF models from Kirchhoff scattering and microfacet theory with support for tilted surfaces. Holzschuch and Pacanowski [2017] proposed a two-scale microfacet model based on generalized Harvey-Shack theory to combine reflection and diffraction. Toisoul and Ghosh [2017] introduced a model suitable for sim-

ulating the iridescent behavior of surfaces with periodic structure for small angles. Their technique is based on a tabulated pre-convolution lookup of the reflected radiance by periodic diffraction gratings. However, the memory requirements prohibit the use of their approach for real-time rendering of worn surfaces, which are composed of many scratches with randomized parameters. This approach relates to many previously studied techniques for precomputation described in the survey by Ramamoorthi [Ramamoorthi 2009], and shares the same limitations.

4.2.2 Rendering scratch-like features

Rendering scratches and line-like features such as fibers [d'Eon et al. 2013] or hair [d'Eon et al. 2011; Yan et al. 2017] are closely related due to the common geometrical properties. Recent models for reflection from worn surfaces are limited to geometrical optics and represent scratches using scratch profiles made of aligned surfaces with associated procedural BRDFs, their positions described by curves on the surface [Bosch et al. 2004]. More recently, Raymond et al. [2016] proposed a multi-scale SVBRDF model from a stack of scratch layers. Yan et al. [2016] utilize optimized integration of the normal distribution function on normal-mapped surfaces to recreate glint-like reflectance behavior.

4.2.3 Extended light sources

The appearance of materials strongly depends on the distribution of light surrounding them. The graphics community therefore has invested major efforts to develop efficient approximate methods for shading with extended light sources. In opposition to standard brute-force path tracing where luminaires are sampled, many techniques have been studied to efficiently sample area sources of different geometry [Shirley et al. 1996]; in particular spherical triangles [Arvo 1995], spherical rectangles [Ureña et al. 2013], and spherical ellipses [Guillén et al. 2017]. When applied to strongly specular material models with high angular bandwidth, these integration strategies require a high number of samples for converged, noise-free images, making them unsuitable for most real-time rendering purposes.

In concurrent work, researchers have attempted to reduce area lights to representative point sources [Picott 1992], causing a severe loss in quality unless sets of many virtual point lights (VPLs) are used [Keller 1997; Dachs-

bacher et al. 2014]. The number of VPLs needed to obtain high-quality renderings increases with the bandwidth of the material, so that for specular materials real-time evaluation becomes more and more challenging. Techniques that project the illumination into spherical basis functions for efficient pre-computation [Sloan et al. 2002] suffer from the same problem: high-frequency illumination in combination with high-frequency reflectance requires impractically high orders of spherical harmonics to achieve output of sufficient quality.

A third strand of research aims to approximate area light sources using a single sample, ranging from the well-known computation of form factors in the radiosity method for diffuse surfaces [Goral et al. 1984; Baum et al. 1989] to extensions of the representative point method that approximate the combined effect of specular lobe and area sources [Drobot 2014; Karis 2013; Carpentier and Ishiyama 2017]. In the context of our work, the most relevant class of methods are those that derive analytical solutions for the illumination integral over the solid angle subtended by area light sources. For glossy materials, for instance, the problem has been addressed by several methods based on polynomial expansions of the two-dimensional (2D) illumination integral [Bao and Peng 1993], on multiple one-dimensional (1D) integrals [Poulin and Amanatides 1991; Tanaka and Takahashi 1997], or a single contour integral using irradiance tensors for polygonal area sources [Arvo 1995; Snyder 1996; Lecocq et al. 2016]. More recently, Heitz et al. [2016a] proposed to transform microfacet distributions into linearly transformed cosine distributions (LTCs) to approximate the reflection of polygonal and linear [Heitz and Hill 2017] light sources. Finally, Dupuy et al. [2017] used spherical pivot transformed distributions (SPTDs) to efficiently approximate the integral over the solid angle subtended by spherical light sources for microfacet models.

The appearance model that serves as the basis for the proposed technique [Werner et al. 2017] specializes on microscopic iridescent scratches on worn surfaces, which gives rise to extremely detailed and specular reflections. Applied naïvely to real-time rendering, their model misses important visual properties such as elongated scratches, but instead produces point-like features (cf. Fig. 4.2). We aim to mitigate these drawbacks while using only a single sample per pixel to achieve real-time capability.

4.3 Theory

In this work, we seek to extend the scratch iridescence model by Werner et al. [2017] for use in real-time applications, and in particular for rendering well-filtered scratches under extended light sources with a single sample per pixel. To this end, we provide approximations and techniques that are required to efficiently evaluate core elements of the model. In particular, we consider the missing pieces necessary to achieve real-time capability to be (1) anti-aliasing and (2) lighting by area light sources, which need to be realized using only one sample per pixel. To better understand the properties of the original model, we will in the following discuss its main elements and how to exploit them for our purpose.

4.3.1 Model

In its original form, the scratch iridescence model is given by the SVBRDF

$$f_r(\xi, \mathbf{x}) = \frac{1}{\pi\sigma^2\lambda^2} |\mathcal{B}(\xi) - \mathcal{S}(\xi, \mathbf{x})|^2, \quad (4.1)$$

where ξ is the vector of direction cosines [Krywonos 2006] and \mathbf{x} the shaded surface position. In essence, a base response function \mathcal{B} and a scratch response function \mathcal{S} model the respective contributions of the unscratched surface and the iridescent scratches. We base our approximation of the scratch response \mathcal{S} on a vector graphics representation (discrete lines) and sum over the contributions $\mathcal{S}^{(m)}$ of individual scratch primitives:

$$\begin{aligned} \mathcal{S}(\xi, \mathbf{x}) &= \sum_m \mathcal{S}^{(m)}(\xi^{(m)}, \mathbf{x}) \\ &= \sum_m \mathcal{W}^{(m)} \mathcal{D}^{(m)} \eta^{(m)}(\xi^{(m)}, \mathbf{x}). \end{aligned} \quad (4.2)$$

Assuming a rectangular scratch cross-section, the individual factors for the m^{th} scratch are:

$$\mathcal{W}^{(m)}(\xi^{(m)}) = W_s^{(m)} \text{sinc}(\pi W_s^{(m)} \xi_2 / \lambda), \quad (4.3)$$

$$\mathcal{D}^{(m)} = 1 - e^{i4\pi D_s^{(m)} / \lambda}, \quad \text{and} \quad (4.4)$$

$$\eta^{(m)}(\xi^{(m)}, \mathbf{x} | l_1, l_2) = d \text{erf} \left(\left((s_t^{(m)} + h) / \sigma + i\sigma k \xi_1 \right) / \sqrt{2} \right) \Big|_{h=l_1}^{l_2}. \quad (4.5)$$

$\mathcal{W}^{(m)}$ denotes the spectral term which, in turn, depends on the *scratch width* $W_s^{(m)}$. $\mathcal{D}^{(m)}$ denotes the *depth*-related term, and $\eta^{(m)}$ encodes the spatial phases with respect to \mathbf{x} . $\mathcal{D}^{(m)}$ is approximated for small angles to enable analytic integration for area light sources in the following. To evaluate this deviation from the original work by Werner et al., we provide a qualitative comparison of this approximation in Fig. 4.3 which shows that the overall shape of the scattering function is well preserved when only separate scratches are considered. This approximation is provided in the public implementation of the work [Werner et al. 2017] without being discussed in the main text.

We further provide the main elements of the model for completeness and easier implementation:

$$\begin{aligned}\xi^{(m)} &= (\xi_1^{(m)}, \xi_2^{(m)}) = (\omega_i + \omega_o) \cdot (\mathbf{t}^{(m)}, \mathbf{b}^{(m)}) \\ \mathbf{s}'^{(m)} &= \mathbf{s}^{(m)} - \mathbf{x} = s_t^{(m)} \mathbf{t}^{(m)} + s_b^{(m)} \mathbf{b}^{(m)} \\ k &= \frac{2\pi}{\lambda} \\ d &= \sqrt{\pi/2} \sigma \exp\left(-\frac{1}{2\sigma^2} s_b^{(m)2} - \frac{\sigma^2}{2} k^2 \xi_1^{(m)2} - ik s_b^{(m)} \xi_2^{(m)}\right),\end{aligned}$$

where $\xi^{(m)}$ is the vector of direction cosines, $(\mathbf{t}^{(m)}, \mathbf{b}^{(m)})$ the tangent and bitangent of a scratch, $\mathbf{s}'^{(m)}$ its position with respect to \mathbf{x} which is further decomposed in tangential ($s_t^{(m)}$) and bitangential ($s_b^{(m)}$) components, $\sigma = 60 \mu\text{m}$ the radius of the coherence area and $k = \frac{2\pi}{\lambda}$ the wavenumber with respect to wavelength λ .

The base material \mathcal{B} can be any physically based material model under the assumption of negligible scratch-surface interference. More formally,

$$\begin{aligned}f_r(\xi, \mathbf{x}) &= \frac{1}{\pi\sigma^2\lambda^2} (|\mathcal{B}|^2 + |\mathcal{S}|^2 - 2\mathcal{B} * \mathcal{S}) \\ &\approx \frac{1}{A_c\lambda^2} ((1 - \rho_c)|\mathcal{B}|^2 + |\mathcal{S}|^2),\end{aligned}\quad (4.6)$$

where the cross-term related to scratch-surface interaction is neglected and $A_c = \pi\sigma^2$ is the coherence area covered by the spatial filter. To ensure energy conservation, a simple alpha blending term based on the *scratch area density* ρ_c is introduced. We focus on using microfacet materials to define the base contribution as they are prevalent in the current rendering systems and provide overall easier control over the appearance of the surface through intuitive parameters. For efficient rendering of microfacet surfaces lit by extended polygonal, tube or spherical light sources we refer to the recent




Figure 4.3: Comparison of diffracted radiance for a single scratch with rectangular profile. The result obtained with our model (black line; no angle dependence, $\mathcal{D}^{(m)} \propto \exp(i4\pi D_s^{(m)}/\lambda)$) underestimates the diffraction of side lobes and neglects re-distribution of energy due to the angular dependence of the complex phase in the original model (blue dashed line; $\mathcal{D}^{(m)} \propto \exp(i2\pi(\cos\theta_i + \cos\theta_o)D_s^{(m)}/\lambda)$). However, the overall appearance is conserved as we closely preserve the envelope shape of the scattering function ($\lambda = 700$ nm).

works of Heitz et al. [2016a], Heitz and Hill [2017] and Dupuy et al. [2017]. We use these distributions fit to the GGX model and directly superimpose the scratches as illustrated by Eq. 4.6.

The solution for incoherent superposition additionally assumes that mutual interference of scratch ensembles, expressed by cross-terms appearing in $|\mathcal{S}(\xi)|^2$, can be neglected such that

$$|\mathcal{S}(\xi)|^2 = \left| \sum_m \mathcal{S}^{(m)}(\xi^{(m)}) \right|^2 \approx \sum_m \left| \mathcal{S}^{(m)}(\xi^{(m)}) \right|^2. \quad (4.7)$$

All previously expressed formulations within this section follow directly from the work by Werner et al. and its publicly available implementation. Based on them, we are now able to separately compute the radiance scattered by each single scratch and to derive more efficient formulas for certain individual parts.




Figure 4.4: Footprint anti-aliasing evaluation. Without spatial integration over the footprint (Left), scratches appear segmented. Footprint integration (Middle) generates connected scratches. Anti-aliasing is only performed along scratches which reflects the assumption that the Gaussian filter extent is much smaller than the scratch width. Monte Carlo integration (Right) additionally is able to anti-alias across scratches.

It is generally assumed that the scratches fall completely within the coherence area. Consequently $\eta^{(m)}$ can be simplified by taking the limit of an infinite scratch:

$$\eta^{(m)} = \sqrt{2\pi}\sigma \exp\left(-\frac{1}{2\sigma^2}s_b^{(m)2} - \frac{\sigma^2}{2}k^2\xi_1^2 - ik s_b^{(m)}\xi_2\right). \quad (4.8)$$

The complex part of the exponential is completely canceled when we take the absolute value squared, leading to a very compact form of the term for the reflectance along the scratch:

$$|\eta^{(m)}|^2 = 2A_C \exp\left(-\frac{1}{\sigma^2}s_b^{(m)2} - \sigma^2k^2\xi_1^2\right), \quad (4.9)$$

which separates into an angular and spatial part, respectively

$$\begin{aligned} |\eta_a^{(m)}|^2 &= \exp(-\sigma^2k^2\xi_1^2); \\ |\eta_s^{(m)}|^2 &= 2A_C \exp\left(-\frac{s_b^{(m)2}}{\sigma^2}\right). \end{aligned} \quad (4.10)$$

The separability enables us to derive the integral over a pixel footprint required for correct anti-aliasing in spatial domain and the integral over solid angle subtended by area light sources. Evaluation of the absolute value squared of $\mathcal{D}^{(m)}$ additionally yields

$$|\mathcal{D}^{(m)}|^2 = (2 - 2\cos(4\pi D_s^{(m)}/\lambda)), \quad (4.11)$$




Figure 4.5: Comparison of numeric (reference) vs analytic (approximation) results of the integration of scratch area density $\rho_{\mathcal{P}}$ and spatial phases $|\eta_{\mathcal{P}}^{(m)}|^2$, respectively given by the integral of ρ_c and $|\eta_s^{(m)}|^2$ over the pixel footprint. For footprint sizes smaller than the coherence area ($A_{\mathcal{P}}/A_{\mathcal{C}} < 1$, depicted by the dashed red line) our solution for $|\eta_{\mathcal{P}}^{(m)}|^2$ (Left) fails as a single pixel only observes a small fraction of the coherence area, an unlikely case which is not covered by our approximation. On the other hand, our solution for $\rho_{\mathcal{P}}$ (Middle, $\xi_1 = 0$) overestimates the area covered by scratches for smaller ratios $A_{\mathcal{P}}/A_{\mathcal{C}}$ but closely resembles the complex dependency. For increasing area ratios our assumptions are matched and the error becomes smaller. The dashed vertical lines correspond to ratios of footprint area to coherence area $A_{\mathcal{P}}/A_{\mathcal{C}} \in [1, 10, 20]$. The corresponding pixel footprints are depicted as color matching ellipses on the surface used for our simulations (Right).

which is a well known result that can be traced back to the derivation of the equation proposed by Stam [Stam 1999] for modeling the appearance of CDs with predefined periodic structure. These results already improve the performance of the model for the specific case of incoherent surfaces.

4.3.2 Integration of the spatially varying reflectance over the camera pixel footprint

The full flux received by a pixel is determined by the irradiance impinging on its spatial footprint, which we approximate with an ellipse in the tangent plane with primary and secondary axes a_1 and a_2 respectively. In a wave-optical context with spatial coherence taken into account, we can think of the pixel footprint as being subsampled with coherent samples (i.e. Fourier transform of spatially weighted surface features). The coherent subsamples are converted to radiance (via BRDF) and then averaged over the surface

area visible in the pixel, which corresponds to an incoherent superposition. Following Eq. 4.6 and Eq. 4.7 we obtain

$$\int_{\mathbf{x} \in A_{\mathcal{P}}} f_r(\xi, \mathbf{x}) d\mathbf{x} \approx \frac{1}{A_{\mathcal{P}} A_c \lambda^2} \int_{\mathbf{x} \in A_{\mathcal{P}}} \left[(1 - \rho_c) |\mathcal{B}|^2 + \sum_m |\mathcal{S}^{(m)}|^2 \right] d\mathbf{x}, \quad (4.12)$$

where $\mathbf{x} = (x, y)^T$ is the point of intersection, $A_{\mathcal{P}}$ is the pixel footprint and $A_{\mathcal{P}} = a_1 a_2 \pi$ is a factor which normalizes the integral to the covered area. In the case of incoherent scratches (and scratches extending over the coherence area) the only term dependent on \mathbf{x} is $|\eta_s^{(m)}|^2$ (cf. Eq. 4.10). What is left to solve is then

$$\frac{1}{A_{\mathcal{P}}} \int_{\mathbf{x} \in A_{\mathcal{P}}} |\eta_s^{(m)}|^2 d\mathbf{x} = \frac{2A_c}{A_{\mathcal{P}}} \int_{\mathbf{x} \in A_{\mathcal{P}}} \exp\left(-\frac{s_b^{(m)2}}{\sigma^2}\right) d\mathbf{x}. \quad (4.13)$$

If this integral needs to be performed over an arbitrary region, interdependencies between the integration bounds for tangential and bitangential direction might occur. However, for our case it should be safe to assume that the footprint is much larger than the coherence area so that we can extend the integration bounds of the bitangential direction to

$$\begin{aligned} \int_{\mathbf{x} \in A_{\mathcal{P}}} \exp\left(-\frac{s_b^{(m)2}}{\sigma^2}\right) d\mathbf{x} &\approx \int_{-\infty}^{\infty} \exp\left(-\frac{(s_y^{(m)} - y)^2}{\sigma^2}\right) dy \int_{s_1}^{s_2} dx \\ &= \sigma \sqrt{\pi} l_{\text{contained}}^{(m)}, \end{aligned} \quad (4.14)$$

where s_1 and s_2 are the intersection points of the scratch with the pixel footprint and thus the second integral yields the length fraction $l_{\text{contained}}^{(m)}$ of the scratch that is contained in the footprint. To ensure that we do not overestimate $l_{\text{contained}}^{(m)}$, we perform a line-ellipse intersection test and retrieve the accurate contained geometric length. We can then approximate the integral over spatial phases by

$$|\eta_{\mathcal{P}}^{(m)}|^2 := \int_{\mathbf{x} \in A_{\mathcal{P}}} |\eta_s^{(m)}|^2 d\mathbf{x} \approx 2 \frac{\sigma^3}{a_1 a_2} \sqrt{\pi} \cdot l_{\text{contained}}^{(m)}. \quad (4.15)$$

We expect this approximation to be valid for larger footprints and to overestimate the integral for smaller ones. We therefore propose a correction term for small footprint-to-coherence area ratios which linearly interpolates between two limit cases of the integral. For the lower bound we assume

the footprint size to be negligible, the upper bound is given by the previous formulation in Eq. 4.15. We thus obtain a corrected estimate of

$$|\eta_{\mathcal{P}}^{(m)}|^2 \approx \frac{2\sigma^2}{A_{\mathcal{P}}} \left[\alpha_{\mathcal{P}} \sigma \sqrt{\pi^3} \cdot l_{\text{contained}}^{(m)} + (1 - \alpha_{\mathcal{P}}) A_{\mathcal{P}} e^{-s_b^{(m)2}/\sigma^2} \right] \quad (4.16)$$

with the blending factor $\alpha_{\mathcal{P}} = \min(\frac{1}{2} \frac{A_{\mathcal{P}}}{A_{\mathcal{C}}}, 1.0)$. With this correction term, we obtain low errors for the standard case of footprint areas about 7 times larger than the coherence area (cf. Fig. 4.5, middle). Although we overestimate the scratch area density with our approximation which lets scratches appear wider than expected, our method still yields visually close results in comparison to reference Monte Carlo renderings visible in Fig. 4.4.

4.3.3 Scratch area density

To allow combination between the incoherent model and arbitrary base BRDFs, Werner et al. use a modified alpha-blending step that takes into account the scratch area contained within the coherence area. More formally, the resulting area density ρ_c is given by the Gaussian-weighted scratch length and its width. For scratches extending over the coherence area this yields

$$\begin{aligned} \rho_c &= \frac{1}{A_{\mathcal{C}}} \sum_m W_s^{(m)} \int_{s_1}^{s_2} ds \exp\left(-\frac{|\mathbf{s}''^{(m)}(s)|^2}{2\sigma^2}\right) \\ &\approx \frac{1}{A_{\mathcal{C}}} \sum_m W_s^{(m)} \sqrt{2\pi}\sigma \exp\left(-\frac{s_b^{(m)2}}{2\sigma^2}\right) \end{aligned} \quad (4.17)$$

with $\mathbf{s}''^{(m)}(s) = \mathbf{s}'^{(m)} + \mathbf{s}\mathbf{t}^{(m)}$, m denoting the m^{th} scratch. For integration over the pixel footprint, we have to modify the area density to account for *subsampling* of the footprint with samples of the size of the coherence area:

$$\begin{aligned} \rho_{\mathcal{P}} &= \frac{1}{A_{\mathcal{P}}} \int_{\mathbf{x} \in A_{\mathcal{P}}} \rho_c(\mathbf{x}) d\mathbf{x} \\ &= \frac{1}{A_{\mathcal{P}} A_{\mathcal{C}}} \sqrt{2\pi}\sigma \sum_m W_s^{(m)} \int_{\mathbf{x} \in A_{\mathcal{P}}} \exp\left(-\frac{s_b^{(m)2}}{2\sigma^2}\right) d\mathbf{x} \\ &\approx \frac{2}{A_{\mathcal{P}}} \sum_m W_s^{(m)} l_{\text{contained}}^{(m)}, \end{aligned} \quad (4.18)$$

where we used the solution of the integral found in Eq. 4.15. Eq. 4.18 now yields the mixing factor necessary for the alpha-blending between our scratch

response and (arbitrary) microfacet base BRDFs which replaces ρ_c within the context of Eq. 4.6. Fig. 4.5 (top) reveals that for small footprint size, a high deviation between reference and approximation is obtained, which corresponds to unlikely cases of pixels covering only a tiny fraction of the coherence area. For increasing size however, the results agree well and only small errors are reported. Our footprint integration scheme relies on the extension of integration limits to $(-\infty, \infty)$ (cf. Eq. 4.14). This enables us to develop anti-aliasing along scratches, but disables smoothing across scratches which becomes visible in Fig. 4.4, a result comparable to standard anti-aliasing techniques based on weighted area sampling (e.g. [Gupta and Sproull 1981]), albeit with an anisotropic filter kernel. With this we are able to significantly improve performance and recreate connected scratches with the minor disadvantage of non-perfect anti-aliasing.

4.3.4 Integration of the reflectance in angular domain over the solid angle subtended by area light sources

The angular dependence of the shading model can be fully expressed by Eq. 4.3, 4.10. Both depend on the direction cosine along a specific direction. To derive an approximation of the integral over the solid angle subtended we must first rewrite the rendering equation, so that it is in direction cosine space. We start from the classical rendering equation first introduced to graphics community by Kajiya [Kajiya 1986]:

$$L = \iint f(\phi_i, \phi_o, \theta_i, \theta_o) L_i \cos \theta_i \sin \theta_i d\theta_i d\phi_i. \quad (4.19)$$

The direction cosines irrespective of features in the tangent plane can be expressed as:

$$\begin{aligned} \omega_o^\perp &= \omega_o \cdot (\mathbf{t}, \mathbf{b}) = (\omega_{o,1}^\perp, \omega_{o,2}^\perp) \\ \xi_1 &= \cos \phi_i \sin \theta_i + \omega_{o,1}^\perp \end{aligned} \quad (4.20)$$

$$\xi_2 = \sin \phi_i \sin \theta_i + \omega_{o,2}^\perp. \quad (4.21)$$

Rewriting these equations according to the azimuth and altitude angles yields

$$\begin{aligned} \theta_i &= \text{asin}|\xi - \omega_o^\perp| \\ \phi_i &= \text{atan} \frac{\xi_2 - \omega_{o,2}^\perp}{\xi_1 - \omega_{o,1}^\perp}, \end{aligned} \quad (4.22)$$

which we use to derive the Jacobian determinant for a transformation of the integral into direction cosine space:

$$|\mathbf{J}| = \left| \frac{\delta(\theta_i, \phi_i)}{\delta(\xi_1, \xi_2)} \right| = \left| \begin{array}{cc} \frac{\xi_1 - \omega_{o,1}^\perp}{\cos \theta_i \sin \theta_i} & \frac{\xi_2 - \omega_{o,2}^\perp}{\cos \theta_i \sin \theta_i} \\ -\frac{\xi_2 - \omega_{o,2}^\perp}{|\xi - \omega_o^\perp|^2} & \frac{\xi_1 - \omega_{o,1}^\perp}{|\xi - \omega_o^\perp|^2} \end{array} \right| = \frac{1}{\cos \theta_i \sin \theta_i}, \quad (4.23)$$

with which the rendering equation can be expressed as:

$$L = \iint f(\xi_1, \xi_2) L_i d\xi_1 d\xi_2. \quad (4.24)$$

The derived relationship signifies that the cosine between the outgoing direction and the surface normals acts only as an offset of the projected light source and thus simplifies further derivations.

The BRDF that we are studying is separable in angular domain and changes according to the following relation:

$$|\mathcal{S}^{(m)}(\xi^{(m)})|^2 \propto \text{sinc}^2\left(\pi \frac{W_s^{(m)}}{\lambda} \xi^{(m)}\right) \exp\left(-\sigma^2 k^2 \xi_1^{(m)2}\right). \quad (4.25)$$

The key observation is that the standard deviation of the Gaussian function is relatively low. We can establish an upper bound of the standard deviation by taking the highest wavelength in the visible spectrum $\lambda_{\max} = 750 \mu\text{m}$:

$$s_{\max} = \frac{\lambda_{\max}}{2\sqrt{2}\pi\sigma} = 1.40674 \cdot 10^{-3}. \quad (4.26)$$

A common empirical rule for upper bounds is to use three times the standard deviation; from this it becomes clear that for most practical purposes the area with the highest contribution would be significantly smaller than the projection of most light sources. The main benefit of this observation is that we can limit the integration to a small band around the origin in direction cosine space as shown in Fig. 4.6. Furthermore, we can ignore the curvature around the light source's edges such that integration can be performed separately for each axis. The integral of a Gaussian distribution at infinity is well known while the integral of $\text{sinc}^2(x)$ can be derived through integration by parts. The final equation has the following form:

$$\begin{aligned} S(\xi|\zeta_0, \zeta_1)^2 &= W_s^{(m)2} \mathcal{D}^{(m)2} \eta_s^{(m)2} (\mathcal{N}^{(m)}(\xi, \zeta_1) - \mathcal{N}^{(m)}(\xi, \zeta_0)) \\ \mathcal{N}^{(m)}(\xi, \zeta) &= \frac{\sqrt{\pi}}{\sigma k} \left(2 \frac{\text{Si}(kW_s^{(m)}\zeta)}{kW_s^{(m)}} - 4 \frac{\sin^2(kW_s^{(m)}\zeta/2)}{k^2 W_s^{(m)2} \zeta} \right) \end{aligned} \quad (4.27)$$




Figure 4.6: Approximation of the integral over the solid subtended by triangular light source in direction cosine domain. We limit the integration over a small band along the bitangential direction cosine, starting at the origin $(-\omega_o^\perp)$.

where ζ_0 are ζ_1 are the integration bounds according to ξ_2 which represents the azimuthal reflectance, and $\text{Si}(x)$ is the trigonometric integral

$$\text{Si}(x) = \int_0^x \text{sinc}(t) dt = \int_0^x \frac{\sin(t)}{t} dt. \quad (4.28)$$

As $\text{Si}(x)$ is a nonelementary integral and not available on the GPU, we approximate it as the superposition of a sigmoid and a hyperbolically decaying oscillation. The following equation is used by our code:

$$\text{Si}(x) \approx \frac{\pi}{2} \tanh(x) - \frac{a_0 x}{1 + (x - \text{sgn}(x)b_0)^2} + \text{sgn}(x) \frac{1 - \cos(x)}{1 + |x|} \quad (4.29)$$

with $a_0 = 0.66142739$ and $b_0 = 0.38272292$ obtained via least-squares fit of the middle term to the residual in range $x \in [-10, 10]$ to prevent overestimation of the oscillations. We provide approximation error analysis in Fig. 4.7.




Figure 4.7: Comparison between the approximation of Si and the exact function (left), as well as the corresponding residual (right).

Spherical light sources

We compute the boundaries in direction cosine domain by using the normalized equation of a disk that can be expressed as:

$$\mathbf{G}_E = \frac{\mathbf{p} + r_t \cos \theta \mathbf{t} + r_b \sin \theta \mathbf{b}}{|\mathbf{p} + r_t \cos \theta \mathbf{t} + r_b \sin \theta \mathbf{b}|}. \quad (4.30)$$

The sphere is just a special case of this equation where the radius of the




Figure 4.8: Comparison of the specular lobe of (a) the ground truth Monte Carlo sampling of light source at 16k spp, (b) our spherical light source approximation and (c) the same approximation with exact $\text{Si}(x)$. Equirectangular mapping of the viewing directions is used with radiance boosted ten times. The light source is placed close to the origin of the polar coordinate system and it is mirror reflected π radians around the axis while the scratch is perpendicular to the main axis.

disk is expressed as $r = r_t = r_b$ and the major axes are perpendicular to the relative position \mathbf{p} ($(\mathbf{t} \cdot \mathbf{p}) = (\mathbf{b} \cdot \mathbf{p}) = 0$), leading to

$$\mathbf{G}_S = \frac{\mathbf{p} + r \cos \theta \mathbf{t} + r \sin \theta \mathbf{b}}{\sqrt{r^2 + |\mathbf{p}|^2}}. \quad (4.31)$$

Constructing the orthogonal basis around \mathbf{p} is achieved by using the branch-less version of the code proposed by Duff et al. [2017]. To find the intersection points with the line passing through the origin we must solve the following equation:

$$t_{\xi_1} r \cos \theta + b_{\xi_1} r \sin \theta + p_{\xi_1} + \omega_{o,\xi_1} \sqrt{r^2 + |\mathbf{p}|^2} = 0 \quad (4.32)$$

Using the trigonometric identities turns the equation into quadratic form:

$$\nu = \tan \frac{\gamma}{2}; \quad \sin \gamma = \frac{2\nu}{1 + \nu^2}; \quad \cos \gamma = \frac{1 - \nu^2}{1 + \nu^2} \quad (4.33)$$

We use the method stated in Numerical Recipes [Press et al. 1992] for solving quadratic equations in a numerically stable fashion:

$$D_q = b_{\xi_1}^2 r^2 - \left(\left(p_{\xi_1} + \omega_{o,\xi_1} \sqrt{r^2 + |\mathbf{p}|^2} \right)^2 - t_{\xi_1}^2 r^2 \right) \quad (4.34)$$

$$\nu = \left(\frac{-(b_{\xi_1} r + \text{sgn}(b_{\xi_1}) \sqrt{D_q})}{p_{\xi_1} + \omega_{o,\xi_1} \sqrt{r^2 + |\mathbf{p}|^2} - t_{\xi_1} r} \right)^{\pm 1}, \quad (4.35)$$

and find $\sin \gamma$ and $\cos \gamma$ by back-substituting ν into the respective trigonometric identities (cf. Eq. 4.33).

The complete solution requires clipping of the spherical disk to limit the contribution to the upper hemisphere. We perform clipping by making a check whether the intersection point is below the surface by substituting only the normal direction cosine component z in the numerator of Eq. 4.31. In case it is below the surface, we snap the point to the arc defined by the intersection of the disk with the plane that splits the circle into two hemispheres. The equation of an arc $\mathbf{C}^{(m)}$ can be defined as the circular interpolation in the plane formed by the projected vertex on the hemisphere $\mathbf{U}^{(m)}$ and the orthogonal vector $\hat{\Gamma}^{(m)}$:

$$\begin{aligned}
\mathbf{U}^{(m)} &= \frac{\mathbf{V}^{(m)} - \mathbf{x}}{|\mathbf{V}^{(m)} - \mathbf{x}|} \\
\mathbf{U}^{(m)} &= \left(\xi^{(m)} - \omega_o^\perp, \sqrt{1 - |\xi^{(m)} - \omega_o^\perp|^2} \right) \\
\Gamma^{(m)} &= \mathbf{U}^{(m+1)} - \mathbf{U}^{(m)} (\mathbf{U}^{(m)} \cdot \mathbf{U}^{(m+1)}) \\
\hat{\Gamma}^{(m)} &= \frac{\Gamma^{(m)}}{|\Gamma^{(m)}|} \\
\mathbf{C}^{(m)} &= \mathbf{U}^{(m)} \cos \gamma + \hat{\Gamma}^{(m)} \sin \gamma
\end{aligned} \tag{4.36}$$

where the angle on the circle is defined in the range $\gamma \in [0, \text{acos}(\mathbf{U}^{(m)} \cdot \mathbf{U}^{(m+1)})]$. To find the intersection points with the plane between the two hemispheres we must solve the equation:

$$t_z r \cos \theta + b_z r \sin \theta + p_z = 0 \tag{4.37}$$

It is solved in analogy to Eq. 4.32, however, two corner cases have to be considered: The first is caused as the sphere approaches grazing angles and the projected ellipse of decreasing size eventually turns into a line. We detect this case and perform a simple intersection between a line and the arc to find the bounds of the integral. The second corresponds to a sphere fully below the horizon. We test whether the sphere center is further than one radius under the horizon and discard shading of the respective point upon validation.

To evaluate our derivation we rendered the response in angular domain according to different viewing directions by using our approximation and Monte Carlo sampling of the light source (Fig. 4.8). Apart from minor discrepancies due to our approximation of the $\text{Si}(x)$ -function, we closely match the ground truth. Thus our approach can be used without performing extra-normalization.

Polygonal light sources

Polygonal lights can be represented by a collection of spherical triangles whose edges are defined as parts of three great circles. Each edge can be expressed as an arc via Eq. 4.36. To find the intersections with each edge the following equation must be solved:

$$\omega_{o,\xi_1} + U_{\xi_1}^{(m)} \cos \gamma + \hat{\Gamma}_{\xi_1}^{(m)} \sin \gamma = 0 \tag{4.38}$$




Figure 4.9: Comparison of the specular lobe of (a) the ground truth Monte Carlo sampling, (b) our triangular light source approximation and (c) the same approximation with exact $\text{Si}(x)$ computation. Equirectangular mapping of the viewing directions is used with radiance boosted ten times. The coordinate system follows the same principles that we outlined for the spherical light sources.

Analogously to Eq. 4.31 we must solve a quadratic equation that has the following roots:

$$D = \hat{\Gamma}_{\xi_1}^{(m)2} + U_{\xi_1}^{(m)2} - \omega_{o,\xi_1}^2 \quad (4.39)$$

$$\nu = \left(\frac{-\left(\hat{\Gamma}_{\xi_1}^{(m)} + \text{sgn}(b_{\xi_1})\sqrt{D}\right)}{\omega_{o,\xi_1} - U_{\xi_1}^{(m)}} \right)^{\pm 1} \quad (4.40)$$

In most cases, the line intersects exactly one edge. For the corner case when a point lies on two arcs we determine whether to keep if there is a difference of the derivatives according to x :

$$\frac{\delta y}{\delta x} = \frac{-U_{\xi_2}^{(m)} \sin \gamma + \Gamma_{\xi_2}^{(m)} \cos \gamma}{-U_{\xi_1}^{(m)} \sin \gamma + \Gamma_{\xi_1}^{(m)} \cos \gamma}. \quad (4.41)$$

After we have determined all intersection points we must form line segments. For this purpose we sort the roots by using bitonic sort [Batcher 1968] by exploiting swizzling and the min and max operations supported by the GPU. We form a line segment out of each pair and compute Eq. 4.27. We clip the polygons by the plane that splits the hemisphere and form a new triangle when necessary.

Similar to the sphere implementation we validate the response in angular domain based on different viewing directions in Fig. 4.9. Analogously, minor discrepancies become visible that are related to the $\text{Si}(x)$ approximation.

4.4 Implementation

We implemented our technique inside a real-time rendering framework following the principles expressed by Burns and Hunt [2013]; we use rasterization to cast primary rays in a scene and directly sample the geometry. The main benefit of this technique is that our approach is evaluated only for the visible pixels. Another similar approach that was widely used in the real-time rendering community was G-buffer based deferred shading, which traces its origin to the technique proposed by Saito and Takahashi [1990]. This technique is not directly applicable to our approach as it would require to store the intersected scratches. Also, storing high quality normals requires additional bandwidth; more traditional ways of encoding the normal in 8-bit or 16-bit variables would lead to severe artifacts caused by the highly specular phenomenon that we are modeling. The shading pipeline used by our

renderer follows the widely adopted principles within production real-time renderers [Hill et al. 2014].

We adopt two basic data structures to store the scratches within our pipeline. We use the basic structure covered by Werner et al. for the case when the footprint integration is enabled, as it requires complete knowledge about the scratches covered by a single pixel. For near-field rendering the visibility buffer structure can be exploited to optimize the intersection checks. We store the scratches associated with a single triangle that fall within the coherence area, which enables faster linear time lookup. It works very efficiently on the GPU for high polygon meshes at the expense of bigger storage requirements. Each node associated with a polygon stores a 32-bit index into a scratch reference buffer and a 32-bit number of scratches within a polygon. To obtain the scratches within the polygon we perform cylinder-line intersection tests where the radius is proportional to three times the coherence area. Additional experiments with a hierarchical version of this technique did not show significant improvement in any of the scenes that we tested, which we attribute to a more divergent execution profile.

We compute the projected camera pixel footprint by first applying the partial derivative functions ($dFdx$, $dFdy$) on the world coordinates of the intersection points, which is supported by the GPU. Then we construct an orthogonal basis on the surface and pass only the tangent and the extent of the pixel footprint ellipse to the functions that evaluate the spatial response.

We perform the computation individually for each light source. Clustering is also applicable to our rendering algorithms; however, we left the exploration of these techniques for future work. For the specific case of polygonal lights we evaluate the contribution of each triangle within the inner loop of the scratch intersection code.

4.5 Results

We performed a validation of our technique against classic Monte Carlo integration according to the publication by Werner et al. which is shown in Fig. 4.10. As expected from the angular domain plots (cf. Fig. 4.8, Fig. 4.9) our results match well. Main discrepancies are color differences of some of the scratches, most likely caused by our $Si(x)$ approximation and the shape of the reflection by the base material. Note that we use the techniques described by Heitz and Dupuy et al. [2016a; 2017] for computing the integral over the




Figure 4.10: Comparison between Monte Carlo integration over the solid angle subtended by (a) spherical and (c) polygonal light source against our approximation of the same integral for (b) spherical and (d) polygonal light source. The difference in the base material highlight is due to the properties of the LTC and SPTD distributions.




Figure 4.11: Recreation of the dining table scene from the original work on iridescent scratches with finer scratches $\sigma = 10nm$ and 200000 individual scratches split between the spoon and the fork. Recreating one of the bright windows in the scene enhances the realism (b) compared to (a) a single light. We provide additional analysis of this many-lights approach in the supplemental material (Sec. 4.A).

solid angle subtended of the base material by spherical and polygonal light source. These are known to fail to match the shape of the highlight at grazing angles. The spherical light source is situated around a meter away from the spoon and has a radius of 5 mm. We use a star shaped light composed of eight individual polygons to illustrate that our technique performs well with polygonal light sources of any shape. In both cases we use scratch parameters (width and depth) uniformly distributed between 0.8 and 6 microns. We show a faithful recreation (Fig. 4.11) of one of the high quality scenes from the original work by Werner et al. Furthermore, we include validation of our technique against one of the ground truth scenes in the supplemental material (Sec. 4.A).

A benchmark of our technique on a notebook GPU (NVIDIA GTX 970M) is provided in Table 4.1; our technique works at real-time performance for spherical light sources. Computation cost is mostly dominated by intersection tests. Except for extreme close-ups, we achieve close to real-time performance which is most likely related to the more complex code for intersecting the projection of a triangle light source. Our technique that reuses the triangle intersections by building per triangle arrays noticeably improves performance. We attribute the low performance on the GPU on divergence caused by branching in the code. Applying our anti-aliasing technique increases the shading cost as we must intersect and evaluate shading for noticeably more scratches. We tested our technique on a previous generation desktop GPU (NVIDIA GTX 980) and it performs at least a factor of two faster in comparison to the notebook GPU. Further benchmarks against high quality scenes are provided in the supplemental material (Sec. 4.A). We expect that our technique would perform much better on current and next generation GPUs making it a useful tool for visualization of iridescent scratches in real time.

4.6 Discussion

We developed a set of techniques which allow the direct application of iridescent scratches within real-time rendering frameworks. Our techniques achieve several orders of magnitude of performance improvement over previous Monte Carlo sampling based techniques. However, to reach truly practical implementation quality for production code bases, several avenues for improvement remain, which we will discuss in the following.

Improved storage and lookup. We used direct segmentation of scratches

	AA disabled					
Type	BVH			Per Triangle Array		
Zoom-In	0.5x	1x	2x	0.5x	1x	2x
Intersection Only	4.1ms	4.4ms	5.5ms	2.9ms	3.5ms	3.8ms
Sphere	4.5ms	6.3ms	12.8ms	3.9ms	4.7ms	8.8ms
Triangle	13.9ms	26.5ms	60.6ms	12.5ms	21.1ms	43.4ms
	AA enabled					
Type	BVH					
Zoom-In	0.5x	1x	2x			
Intersection Only	18.1ms	17.5ms	15.5ms			
Sphere	27.3ms	31.7ms	36.3ms			
Triangle	48.5ms	60.3ms	94.1ms			

Table 4.1: Benchmark of the scene presented in Fig. 4.10 (b) with 10000 scratches split into 411810 segments with and without anti-aliasing enabled. AA is performed using camera footprint integration. The values were capture on a notebook GTX 970M GPU at 720p image resolution. We provide the respective images under different zoom levels in the supplemental material (Sec. 4.A).

according to the polygon boundaries of the shaded mesh. Another possibility that would improve storage efficiency is to store the extent of scratch particles as box-like structures. This technique is known as *deferred decals* in the real-time rendering community. It would greatly reduce the memory footprint of the data structures and reduce lookup times as more shallow data structures are traversed. Another possibility is to apply procedural models to greatly reduce storage. We aim to pursue these and other ideas that improve the memory usage.

More GPU friendly implementation. Currently, the intersection code is heavily reliant on branching to handle corner cases. We are interested in exploring other possible techniques that can handle this problem in a more GPU friendly branchless fashion.

Modeling improvements. We employ the simple Harvey model for our computations as doing an analytic solution of this model is tractable. We would like to explore possible solutions to handle the Generalized Harvey-Shack model [Krywonos 2006]. Another open problem is the handling of more complex scratch profiles as well as their variation along the scratch,

as they cannot be integrated in closed form and necessitate more advanced anti-aliasing techniques. Although many possibilities exist for decomposing the 1D profile functions into terms that are easier to integrate, most of them are not particularly efficient memory-wise or performance-wise. We are also interested in deriving efficient techniques for approximating the contribution of many wavelengths as we limited the results shown in this publication to RGB. Currently, it can be approximated by evaluating the functions for many wavelengths in spectral rendering fashion and converting them to RGB. We provide performance evaluation of this approach in the supplemental material (Sec. 4.A). Another limitation that we share with other area lighting techniques is handling of the Fresnel term which is only approximately supported for dielectrics by the distribution based approaches that fit the Schlick approximation. We are interested in exploring this problem in the future.

Handling of more complex light sources. We currently handle only spherical and polygonal light sources. However, the approach of handling the intersection in direction cosine space should fairly easily generalize to other geometry shapes. An often neglected but still very important open problem is the handling of shading by complex concave meshes. From a theoretic standpoint it is possible to clip polygons so that only visible faces are considered in the computation of the solid angle subtended, which, however, would be a very computation intensive task. We are interested in looking for solutions of this problem, especially with respect to the special case of rendering scratched surfaces and in the more generic case of microfacet theory based modeling. Rendering of environment light probes is also another possible direction for future research.

Shadow mapping. The closed-form solution of shadow mapping by area light sources is considered an open problem in the real-time rendering community. Some special case solutions exist as the ones shown by Dupuy et al. [2017] for shadows cast by spheres. However, the generic case of arbitrary geometry of light sources and shadow casters is unsolved.

Appendix

4.A Additional performance evaluation

To demonstrate the applicability of the techniques we are going to provide a set of renderings of the dining table high quality scene from the original work by Wernel et al. Similar to the original work we pick a coherence area radius of $\sigma = 10\mu m$ which yields thinner scratches in practice. The scenes have approximately 200000 scratches split between the spoon and the fork, resulting in densely worn surfaces which better resemble reality. The result is a very faithful recreation of the scenes provided in the original work. To better understand the sources of error compared to the original model we prepared a table of assumptions (Table 4.3). The scenes as rendered by our framework use precomputed irradiance in light maps to approximate the global illumination and are tone mapped with the Reinhard operator similarly to the original work.

We split the results in two categories: images with 16 spectral bands (Figure 4.13a, 4.13b, 4.13c, 4.13d) and regular RGB (Figure 4.12a, 4.12b, 4.12c, 4.12d). The result of using more spectral bands is slightly more muted colors. Note that we do not consider scratches without anti-aliasing in this case as the smaller coherence area results in much greater subpixel detail. The performance impact is shown in Table 4.2. Adding more spectral bands does not result in integer increase of the rendering times as the computation of the response is the less expensive part of the evaluation. On the other hand, adding polygonal lights is quite expensive as it was shown in the main body of the work. The performance penalty of more polygons is close to linear. The slight deviation from linear dependence of the performance based on the number of polygons might be explained by scheduling effects on the GPU. However, this kind of detailed analysis would require an advanced profiling

Light Polygons	RGB		16 Spectral Bands	
	Figure	Render Time	Figure	Render Time
None	Figure 4.12a	42ms	Figure 4.13a	58ms
2	Figure 4.12b	408ms	Figure 4.13b	753ms
4	Figure 4.12c	562ms	Figure 4.13c	995ms
6	Figure 4.12d	799ms	Figure 4.13d	1649ms

Table 4.2: Performance with different combinations of rectangles approximating key parts of the environment. Benchmark was performed on a laptop with NVIDIA GTX 970M GPU.

(simulation) software for the GPU which is not freely available at this point.

Modeling bright light passing through windows is overall justified as it can be seen by comparing Figure 4.12a and 4.12b. The same observation is valid for the multiple spectral bands renderings in Figure 4.13a and 4.13b. However, the step towards capturing even more windows in the scene as in Figure 4.12c and 4.13c may not result in a change for this particular view. This observation points towards possible future investigation of culling the light sources that are unlikely to contribute. Approximating parts of the environment, such as the walls in Figure 4.12d and 4.13d is less justified as it results in fairly uniform reflections that do not contribute significantly outside of the perfect mirror reflection. One possible topic of future investigation in this case is polygons with gradients to better match the environment.

We further validated our technique against one of the ground truth scenes in the original work by Werner et al. in Figure 4.18. Note that the scratches might appear slightly brighter in our renderings compared to the photographs as we are working with a pinhole camera that perfectly resolves them.

We created an artificial benchmark within the main body of this chapter to outline the cost of having scratches lit by different light sources. Here, we provide the images without comment. Note that the coherence area radius is $\sigma = 60\mu m$ in this case as it makes the scratches visible even without anti-aliasing. We provide separate figures for sphere light sources with anti-aliasing enabled (Figure 4.14) and disabled (Figure 4.15). Similarly we provide the same figures for a light source made of a single triangle with anti-aliasing enabled (Figure 4.16) and disabled (Figure 4.17)




Figure 4.12: Images in RGB of increasing complexity: (a) no extra rectangles, (b) one rectangle approximating the window with the biggest solid angle, (c) approximating also the next big window, (d) adds the reflected wall. 200000 scratches spread equally between the spoon and the fork.




Figure 4.13: Images with 16 spectral bands of increasing complexity: (a) no extra rectangles, (b) one rectangle approximating the window with the biggest solid angle, (c) approximating also the next big window, (d) adds the reflected wall. 200000 scratches spread equally between the spoon and the fork.




Figure 4.14: Sphere light source with 10000 scratches and anti-aliasing enabled at different zoom levels: (a) 0.5x, (b) 1x, (c) 2x.




Figure 4.15: Sphere light source with 10000 scratches and anti-aliasing disabled at different zoom levels: (a) 0.5x, (b) 1x, (c) 2x.




Figure 4.16: Triangle light source with 10000 scratches and anti-aliasing enabled at different zoom levels: (a) 0.5x, (b) 1x, (c) 2x.




Figure 4.17: Triangle light source with 10000 scratches and anti-aliasing disabled at different zoom levels: (a) 0.5x, (b) 1x, (c) 2x.

Assumption	Usecase/Effect	Reference
Small angle phase dependence	$(\mathcal{D}^{(m)} \propto \exp(i4\pi D^{(m)}/\lambda))$; Analytic area-light integration	Sec. 4.3.1 Fig. 4.3
Rectangular profiles only	Separation of angular and spatial terms	Sec. 4.3.1
Scratches extend over coherence area	Separation of angular and spatial terms	Sec. 4.3.1 Eq. 4.9, 4.10
No scratch-surface interaction (Incoherence)	Base reflectance decoupled from scratch reflectance	Sec. 4.3.1 Eq. 4.6
No scratch-scratch interaction (Incoherence)	Separate reflectance computation for each single scratch	Sec. 4.3.1 Eq. 4.7
Elliptic pixel footprint	-	Sec. 4.3.2
Coherence area much smaller than: - Pixel footprint - Projected light source	- Separation of integrals, limit-case solutions for $ \eta_{\mathcal{P}}^{(m)} ^2$, $\rho_{\mathcal{P}}$ Separation of integrals over solid subtended by light source	Sec. 4.3.1-4.3.3 Sec. 4.3.4

Table 4.3: Convenience list of the assumptions that are used by our model as well as a reference to its first introduction.

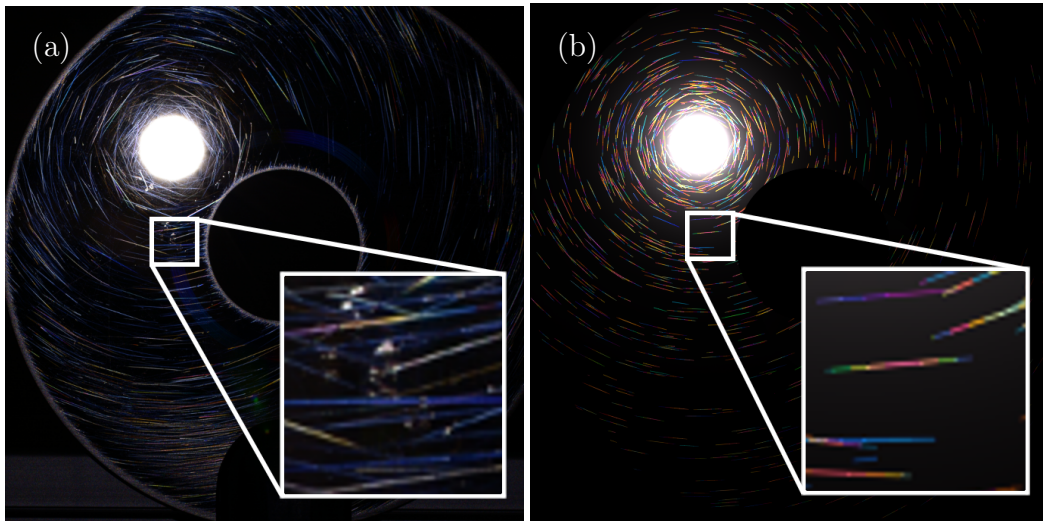


Figure 4.18: Comparison between (a) ground truth photo of a disc and (b) rendering with our sphere light approximation. Variation is achieved by splitting the scratches into multiple segments. The phenomenological connection with the ground truth is preserved by our approximation.

CHAPTER 5

Appearance Capture and Modeling of Human Teeth

Preface Transparent and translucent materials are optically thin and light easily enters their inner volume. Deep structures within them can noticeably affect the process of light propagation and consequently their overall appearance. This phenomenon is modeled by the theory of radiative transfer which simulates the physical interaction of light with matter through the statistical models of extinction, scattering and phase function. Human teeth are characterized by the same phenomenon and they require careful modeling of the subsurface light interaction. Light passes through their thin parts which can be observed in many back-lit scenarios. Therefore, a framework capable of simulating the radiative transfer in its entirety is required to accurately render the overall material appearance. Global illumination and occlusion phenomena are further observed within the oral cavity, thus reinforcing the need for a highly accurate light transport simulation.

In this chapter a complete pipeline for acquisition and inverse rendering is proposed. Inverse rendering is based on differentiable path tracing that drives a joint numerical optimization of inner geometry and optical properties. The main contribution to the literature on differentiable path tracing frameworks is a complete set of algorithms for recovering inner triangular mesh boundaries between index-matched participating media. Transferring the knowledge acquired during this project to modeling of

other materials has the potential to further enrich the visual fidelity of digital doubles of humans and translucent objects.

The main simplifications that keep the problem tractable are the use of homogeneous layers of material, index-matched boundaries and a thin smooth surface material that does not bend light. Anisotropy is modeled only within the bounds of the phase function. The main impact of smoothness is the presence of highly specular highlights that are automatically masked to avoid negatively impacting the optimization process. Further masking is done over the parts of the scene that are not covered by the model (face, lips and retractor). Bending of light (refraction) affects geometrically thin areas where the overall reflectance would deviate from index-matched propagation. Even with these deficiencies, the model is able to recreate faithfully the appearance of teeth within the directly visible and unmasked areas of the scene. Other parts of the scene require regularization to keep their values within the physically plausible limits.

This chapter is based on Velinov et al. [2018a].



Figure 5.1: Various comparisons of our recovered teeth appearance parameters rendered next to corresponding reference images, from different view-points and lighting conditions. Reference images are on the left for individual teeth comparisons, and above for full teeth rows.

Abstract Recreating the appearance of humans in virtual environments for the purpose of movie, video game, or other types of production involves the acquisition of a geometric representation of the human body and its scattering parameters which express the interaction between the geometry and light propagated throughout the scene. Teeth appearance is defined not only by the light and surface interaction, but also by its internal geometry and the intra-oral environment, posing its own unique set of challenges. Therefore, we present a system specifically designed for capturing the optical properties of live human teeth such that they can be realistically re-rendered in computer

graphics. We acquire our data in vivo in a conventional multiple camera and light source setup and use exact geometry segmented from intra-oral scans. To simulate the complex interaction of light in the oral cavity during inverse rendering we employ a novel pipeline based on derivative path tracing with respect to both optical properties and geometry of the inner dentin surface. The resulting estimates of the global derivatives are used to extract parameters in a joint numerical optimization. The final appearance faithfully recreates the acquired data and can be directly used in conventional path tracing frameworks for rendering virtual humans.

5.1 Introduction

A beautiful smile is often considered key to success and the key to a beautiful smile is a nice set of teeth. The ideal teeth appearance depends on the social and cultural context of the subject, as well as on their age since teeth appearance changes over time. In addition to the surface properties of teeth, the appearance also heavily depends on the geometric and optical properties of internal organic structures. These can be influenced by many factors, e.g. the recurrent use of substances like caffeine or tobacco, neglected teeth hygiene, illness, in particular at a young age, or a dysfunction or complete absence of the tooth nerve. All of these can lead to a gradual change of appearance and contribute to the unique look of each single tooth. Conversely, the appearance of teeth may disclose the habits, age, or socio-economic background of their owner.

When creating CG characters, for example for visual effects in motion pictures, it is therefore essential that their teeth match their background story. Similarly, when creating digital doubles of real actors, it is paramount to faithfully reproduce their teeth appearance. While there has been substantial work, in particular in the medical dental field, on capturing the shape and positioning of teeth in vivo, acquisition of their optical properties has been less actively pursued. While the most recent intra-oral scanners do capture color, they offer only very basic appearance acquisition and do not account for the complex optical properties of teeth, ending up with an overall unnatural look, as illustrated by the dental scan using a *3Shape TRIOS[®]* device shown in *Fig. 5.2*. Previous works dedicated to capturing the reflectance of teeth limited themselves to simple models [Buchallard et al. 2007] or derived synthetic models which partially approximate light transport [Larsen



Figure 5.2: **Dental scanning.** Current dental scanners extract a diffuse texture. The diffuse albedo is incorrectly darker in the areas where the tooth is optically thinner, and glossy highlights are also incorrectly encoded, leading to brighter appearance than expected.

et al. 2012], in both cases delivering lower visual quality compared to a more complete light transport simulation.

The scattering properties of a tooth are determined by its two outermost layers composed of materials with very different optical properties (Fig. 5.3). The inner core is made out of dentin, which has the biggest influence on its chroma and is also responsible for the fluorescent nature of teeth. A layer of enamel with varying thickness covers the dentin. It is typically thinner at the gums and thicker towards the tip, causing a chroma gradient. Enamel is a structurally dense material, but thinner than dentin. It has a high degree of translucency, which depends substantially on its mineralization and hydration. At the microscopic level, enamel is formed by rods that run from the dentin towards the surface and cause strongly anisotropic light propagation. Translucent enamel displays the characteristics of opalescence, causing the tooth to backscatter in the blue spectra and hence appear red-yellowish when lit from the other side. Finally, surface variations, cracks, stains and other imperfections add yet another layer of visual richness to teeth appearance.

This work presents a step towards overcoming the limitations of present acquisition systems by capturing the optical properties of teeth such that they can be realistically re-rendered in virtual environments. Our work complements existing shape reconstruction methods, such as intra-oral scanning or extra-oral fitting. In fact we use geometry acquired through an intra-oral scan and segment it according to Wu et al. [2016]. We acquire the appearance data by imaging a person’s mouth area from multiple views and under multiple lighting conditions.

To model the complex interaction of light within the oral cavity and

tooth volume, we use a modified path tracing framework capable of estimating derivatives with respect to optimized parameters while performing regular rendering. Essentially, our method iteratively solves for the derivatives of the entire radiative transfer integral over the entire visible scene. Furthermore, we can perform optimization of not just scattering parameters (albedo, density, surface texture) but also geometric parameters (dentin surface shape). We model dentin shape as a set of deformations from the base enamel shape, which we parametrically control in the optimization loop. Thus we are able to predict the approximate shape and contribution of the hidden dentin surface on the appearance. Our approach represents a step forward in visual fidelity compared to previous works on tooth appearance capture. To the best of our knowledge, the proposed use of morphable shapes for representing hidden structures has not been explored by previous methods and represents a major contribution to the field of differentiable path tracing. We believe that our framework is general enough to be successfully employed in other applications, such as appearance acquisition of organic tissues (skin, fruits, plants), translucent gemstones, and other man-made objects.

We validate our technique by capturing a dead tooth extracted from a mouth cavity, where the ground truth inner dentin and outer enamel geometry is available. We further apply our method on in-vivo data acquired in an appearance capture session of human volunteers (Fig. 5.1). In both cases our approach closely recreates the overall appearance. Since the main framework employed within the optimization loop is based on path tracing, the resulting set of parameters is readily available for re-rendering the teeth with very minor modification within production environments.

5.2 Related Work

In the following we discuss other methods that capture properties of teeth in the context of computer graphics and beyond. Our appearance capture approach is also related to methods for reconstructing the appearance of real skin, in particular faces. Since the ultimate goal is to re-render the captured teeth in CG, we outline relevant work in volumetric rendering, particularly those dedicated to teeth. Finally, as inverse rendering is a key ingredient to our optimization pipeline, we end by discussing how this concept has been explored in other contexts.




Figure 5.3: **Anatomy.** a) A tooth is made of two different materials with very different optical properties. The inner core is formed by dentin, which has the biggest influence on the tooth's chroma, and is covered by a layer of optically denser enamel. Finally, the tooth is bounded to the sides by gingiva tissue (gums), which scatters skin tinted light into the tooth. b) At a microscopic level, enamel is formed by rods that run from the dentin towards the surface and cause strong anisotropic light propagation.

5.2.1 Teeth Capture

Reconstructing the 3D shape of teeth is important in the field of dentistry, in particular for diagnosing medical problems or planning restoration procedures. Common approaches include plaster casting or CT scanning [Omachi et al. 2007; Yanagisawa et al. 2014]. More lightweight intra-oral scanners have also been developed (e.g. *3MTM True Definition*, *iTero*, *3Shape TRIOS[®]*, [Ahn et al. 2017]), which typically use structured light to recover teeth shape and, in some instances, a basic surface color. Detailed survey of these devices from dental practitioner standpoint is provided by Mangano et al. [2017]. In the areas of computer graphics and vision, photometric reconstruction methods such as shape-from-shading [Carter et al. 2010] and feature-based reconstruction [Zheng et al. 2011] have been proposed. Because teeth generally have few identifiable features, model-based fitting approaches have gained popularity [Mehl and Blanz 2005; Buchaillard et al. 2007; Farag et al. 2013; Mostafa et al. 2014]. The current state-of-the-art in teeth shape capture is the recent model-based method of Wu et al. [2016], who build a statistical model of teeth shape from dental scans and then fit the model to silhouette features, presenting high-quality results for teeth reconstruction from passive imagery. Most shape reconstruction papers limit their appearance models to color extraction from images, as it is the case in Wu et al.; Kawai et al. [2016; 2014]. They are thus severely limited in their ability to faithfully re-render teeth from novel viewpoints or in different lighting conditions. Our method for detailed appearance capture of teeth would naturally complement these techniques for teeth shape reconstruction.

In medical dental literature, several experiments have measured appearance properties of teeth, in particular the scattering and absorption values of both dentin and enamel [Spitzer and Ten Bosch 1975; Fried et al. 1995; Pop-Ciutrla et al. 2015]. These studies are motivated by different applications in dentistry, for example the creation of false teeth substitutes that are identical to real teeth [Pop-Ciutrla et al. 2015], or the study of how lasers designed for soft tissue surgical procedures interact with hard dental tissue [Fried et al. 1995]. In all cases the measurement process is a destructive one, where individual teeth are sliced into thin layers for measurement. In our work we design a method to measure the appearance properties of teeth in vivo, with the goal of re-rendering for computer graphics. The optical properties reported by dentistry literature, in particular the average measured absorption and scattering properties of enamel [Spitzer and Ten Bosch

1975] and dentin [Pop-Ciutrila et al. 2015], are employed as initial values in our optimization procedure, described in Sec. 5.6.

To the best of our knowledge, the only existing specialized model for rendering teeth is presented by Larsen et al. [2012] which makes several compromises in order to enable real-time performance. In contrast, we are interested in higher fidelity reproduction of the real teeth appearance for the application of high-quality digital humans.

5.2.2 Facial Appearance Capture

Teeth appearance capture can be considered part of the larger field of facial appearance modeling and capture, which has received a lot of attention over the past decades. An in-depth overview can be found in the state-of-the-art report of Klehm et al. [2015] and is beyond the scope of this paper. The relevant work in the area of facial appearance capture is concerned with the estimation of scattering properties of skin, which falls in the wider area of volumetric appearance capture discussed in the next section.

5.2.3 Volumetric Appearance Capture

A number of methods have been proposed for acquiring surface-based definitions of translucency (e.g. the methods by Debevec et al.; Tong et al.; Donner et al. [2000; 2005; 2008]). We take a different approach and explicitly model the volumetric optical properties of teeth, which falls into the wider area of volumetric appearance capture. The most common approach to estimate the absorption and reduced scattering coefficients is to inject light at a specific point and image the diffusion profile on the medium. Approaches employ single beams [Jensen et al. 2001; Weyrich et al. 2006], stripes [Tariq et al. 2006], and dots [Ghosh et al. 2008]. Wang et al. [2008] employed a polygrid diffusion algorithm for solving the diffusion equation. Munoz et al. [2011] proposed a pipeline for predicting the appearance of optically thick homogenous translucent objects based on the diffusion theory. Papas et al. [2013] measured a database of diffusion profiles of pigments used in a mixture estimation pipeline capable of reproducing the appearance of translucent materials. In contrast, Gkioulekas et al. [2013] invert the radiative transfer equation [Chandrasekhar 1960] and solve in operator-centric fashion for optical parameters using a combination of Monte Carlo rendering, stochastic gradient descent, and material dictionaries.

More generally our method builds upon the family of papers that compute analytic path contribution gradients which are numerically integrated on a sensor. Initially Hašan and Ramamoorthi [2013] provided analytic gradients of paths with respect to albedo to facilitate interactive albedo editing. This was later generalized to handle all Bidirectional Scattering Distribution Function (BSDF) and volume related properties [Khungurn et al. 2015] and more general lighting/viewing conditions [Gkioulekas et al. 2016] for the purpose of inverse rendering. Our paper extends this line of work by allowing the joint optimization of optical properties and geometry, which is essential for optimizing the dentin geometry in the case of teeth.

5.2.4 Rendering Translucent Materials

The concept of BSDF was introduced in physics to model transparent surfaces [Bartell et al. 1981; Dereniak et al. 1982]. In order to render materials with a significant degree of internal scattering, Stam [1995] proposed a grid-based method for spatially resolving the diffusion of light. To render effects due to subsurface scattering, Jensen et al. [2001] utilized the diffusion dipole, which was later extended to handle layered materials [Donner and Jensen 2005; Donner et al. 2008] and edges and corners [Donner and Jensen 2007]. It was demonstrated that the quality of diffusion-based approaches can be improved by quantization [d’Eon and Irving 2011] and combination with Monte Carlo simulations [Li et al. 2005; Habel et al. 2013]. Frisvad et al. [2014] proposed a directional dipole that better handles transport at oblique angles. Frederickx and Dutré [2017] introduced a forward scattering dipole model based on the theory of functional integrals for radiative transfer. Despite all of these improvements, the poor handling of the boundary leads to inaccurate results when the object is thin and (locally) highly translucent, such as teeth.

Approaches based on Monte Carlo integration, such as volumetric (bidirectional) path tracing [Kajiya and Von Herzen 1984; Lafortune and Willems 1996], metropolis light transport [Pauly et al. 2000], or photon mapping [Jensen and Christensen 1998; Donner and Jensen 2007] are computationally more intensive, but provide high-accuracy results. Since we strive for visual fidelity and can afford render times on the order of minutes, we use volumetric path tracing during inverse rendering and to render our validation results.

5.3 Method Overview

Light takes many paths when interacting with teeth. First, it interacts with the surface, characterized by a rough dielectric layer, and is either reflected or refracted into the interior. The interactions in the interior are mainly defined by the two layers of a tooth: enamel and dentin. We thus identify three sets of parameters defining the appearance of teeth: surface reflectance and transmission, sub-surface scattering parameters, and inner dentin shape. Our approach aims to derive a plausible set of model parameters for the purpose of high-quality rendering in a scenario where teeth are illuminated externally with respect to the oral cavity. We split our discussion into a general definition of the appearance model in Sec. 5.5, our data acquisition methodology in Sec. 5.4, and the respective parameter estimation pipeline in Sec. 5.6 followed by qualitative evaluations of the proposed model and optimization on real-world measurements in Sec. 5.7.

5.4 Data Acquisition

We fit our appearance model described in Sec. 5.5 to real teeth measurements captured for three volunteers. In this section we describe the data capture and the pre-processing and post-processing steps.

5.4.1 Image Capture

In order to obtain accurate appearance parameters, we require a dense sampling of incoming and outgoing light directions. For this reason, we capture each volunteer in a light dome [Beeler 2012] consisting of 155 lights (each one a RGB LED triplet), where each light is approximately one meter away from the teeth. Due to occlusions of the head, we record imagery for 54 of the lights, which accounts for all visible incoming directions. We capture multi-view imagery from three color cameras (Ximea CB200CG), positioned around the edge of the light dome with long lenses to maximize teeth resolution, recording 20 megapixel images, synchronized with individual light activations. Please refer to Fig. 5.4 (a) for an image of the light dome and example captured imagery.

Calibration We geometrically calibrate the cameras using a planar fiducial pattern [Garrido-Jurado et al. 2014]. The positions of the lights are measured using a Leica DISTOTM S910 laser distance measuring device, and then aligned to the camera coordinate system by imaging a mirror sphere. Color calibration is performed by recording the X-Rite ColorChecker[®] chart with all lights turned on to extract corrective terms similar to Ilie and Welch [2005]. The individual light intensities are measured by recording a diffuse spherical spectralon and considering it as ideal Lambertian surface. In order to account for spurious incoming light, a black (no light) image is recorded and subtracted from all measurements.

Volunteer Preparation Recording the 54 light directions requires approximately 10-15 seconds, during which time the subject must remain still. To facilitate this, we provide a stabilizing neck brace, a head strap and a rigid object to bite on, which also helps to separate the upper and lower teeth. To minimize occlusions by the face, we retract the lips and cheeks using an OptraGate lip retractor from Ivoclar Vivadent, often used in dentistry.

Post-Processing Despite stabilizing the volunteer’s head during acquisition, there is still a minimal amount of residual motion left, which we remove as a postprocessing step by aligning all frames to the first image using optical flow [Brox et al. 2004]. Our imagery contains shadows cast by the lips and lip retractor, which will confuse the appearance optimization since we do not model the geometry of the face or the retractor. Additionally, we will not attempt to match specular highlights since we are more interested in the volumetric light transport. Therefore, we manually mask the imagery to remove shadows, the gums and inner mouth and employ the method of Mallick et al. [2005] to automatically detect and mask specular highlights, keeping only valid teeth pixels that will inform the optimization. These tasks do not form an integral part of our algorithm, and are easy to perform by non-experts. They could also be further automated for use in a production pipeline when the complete facial appearance and geometry is available.

5.4.2 Geometry Capture

The captured imagery alone is not sufficient to perform appearance optimization, as we require knowledge of the teeth geometry as well. While lightweight teeth reconstruction methods exist (e.g. Wu et al. [2016]), they

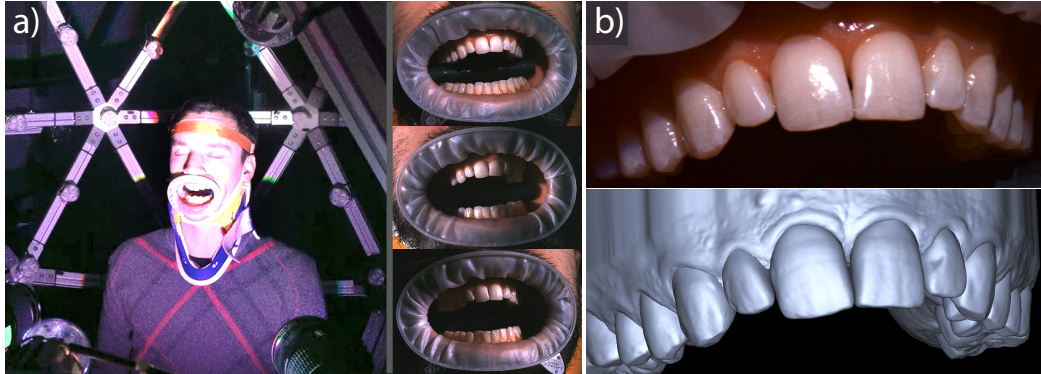


Figure 5.4: **Data Acquisition.** (a) We capture the teeth from three view-points under 54 different light directions using a light dome. Here we show three different light directions for the frontal camera. (b) The teeth are scanned and we align the geometry to the imagery, in preparation for the appearance optimization.

fail to recover the accurate shape of the back side of the teeth and only hypothesize the teeth shape where image coverage is low. Since we need to evaluate volumetric light paths, we require accurate teeth boundaries on all sides. Therefore, we obtain a professional teeth scan for the volunteer using a *3Shape TRIOS*[®] intra-oral scanner. In order to segment the teeth in the scan and separate them from the gums we align template teeth meshes to the scan geometry using the semi-automatic template fitting approach described in Wu et al. [2016]. Since the appearance of teeth is tightly coupled with the gums, we also model the gums geometry artistically based off the scans. Finally, we align the resulting per-tooth geometry and gums model to the camera coordinate frame by manually selecting a sparse set of corresponding landmarks and solving for a rigid transformation (see Fig. 5.4 (b)).

The final result of the data acquisition is, per volunteer, a total of 162 images (54 individual light directions for three views) masked to identify valid teeth pixels, with precise teeth and gums geometry aligned to the imagery. Part of this dataset will be used in the model parameter optimization method and for validation.

5.5 Teeth Appearance Model

The appearance of a tooth is highly influenced by the shape and optical properties of its enamel and dentin parts (Fig. 5.3). We represent both materials as bounded volumes and introduce a number of key simplifying assumptions to accelerate the inverse rendering.

5.5.1 Enamel

Boundary Geometry and BSDF The enamel boundary is the outer boundary of the tooth, and is modeled as a triangle mesh. In order to model the interaction of light with the air-enamel boundary, we use a modified version of the smooth-dielectric BSDF, which adopts the concept of separating the air and enamel using a “thin” dielectric layer.

The BSDF consists of a reflective and a transmissive component, $f(\chi, \omega_i, \omega_o) = f_r(\chi, \omega_i, \omega_o) + f_t(\chi, \omega_i, \omega_o)$, where ω_i and ω_o correspond to the incident and outgoing light directions.

The reflection component is defined in a standard way,

$$f_r(\chi, \omega_i, \omega_o) = \frac{\delta(\omega_h, \mathbf{n})}{|\omega_i \cdot \mathbf{n}|} F(\omega_i \cdot \mathbf{n}, \eta), \quad (5.1)$$

where F is the Fresnel term that quantifies the amount of light reflected off the dielectric boundary (see e.g. [Pharr et al. 2016] for the definition), χ is the spatial position, ω_h is the half-way vector between ω_i and ω_o , \mathbf{n} is the surface normal, and $\delta(a, b)$ is a Dirac delta function that equals $+\infty$ when $a = b$ and 0 otherwise. We use a fixed refractive index from dentist literature $\eta = 1.631$ [Meng et al. 2009] for computing the Fresnel term.

The transmission component,

$$f_t(\chi, \omega_i, \omega_o) = a_t(\chi) \frac{\delta(\omega_i, -\omega_o)}{|\omega_i \cdot \mathbf{n}|} (1 - F(\omega_i \cdot \mathbf{n}, \eta)), \quad (5.2)$$

introduces the simplification that transmitted light does not bend, i.e. it continues along a straight line, which greatly simplifies the light-transport simulation. The spatially varying $a_t(\chi)$ modulates the absorption of the thin enamel/air boundary layer. During inverse rendering we store $a_t(\chi)$ in a texture which is indexed by a UV parametrization of the outer enamel mesh.

The BSDF can be interpreted as an infinitesimally thin dielectric layer, and, while physically inaccurate, the model is expressive enough to capture

specular highlights and stains, which are essential to faithfully represent the appearance. Ignoring the bending of light rays during transmission greatly accelerates our Monte Carlo simulations during inverse rendering as direct illumination can be sampled using shadow rays that connect to lights through the boundary. We avoid highlights during optimization by masking them and reintroduce them in the final images by using a mixture of this BSDF representing the transmission term and a rough dielectric representing the reflection term according to the GGX distribution [Walter et al. 2007].

Volumetric Properties Once transmitted through the boundary, light will start interacting with the inner material. To model the propagation through the volume, we rely on the radiative transfer equation [Chandrasekhar 1960],

$$(\boldsymbol{\omega} \cdot \nabla L) = -\sigma_t L(\boldsymbol{\chi}, \boldsymbol{\omega}) + \alpha \sigma_t \int f_p(\boldsymbol{\omega}_i, \boldsymbol{\omega}_o) L(\boldsymbol{\chi}, \boldsymbol{\omega}_i) d\boldsymbol{\omega}_i, \quad (5.3)$$

which describes the change of radiance L due to absorption and scattering, here quantified by the extinction coefficient σ_t and single-scattering albedo α . To reproduce the characteristic wavelength-dependent scattering and absorption of teeth, these quantities are RGB vector values.

In order to model the directional distribution of scattered light, we use a two-lobe phase function:

$$f_p(\boldsymbol{\omega}_i, \boldsymbol{\omega}_o) = \beta f_{\text{HG}}(\boldsymbol{\omega}_i \cdot \boldsymbol{\omega}_o, g_1) + (1 - \beta) f_{\text{HG}}(\boldsymbol{\omega}_i \cdot \boldsymbol{\omega}_o, g_2), \quad (5.4)$$

where β balances the influence of the two lobes and f_{HG} is the Henyey-Greenstein [1941] phase function parameterized by the mean cosine g of the scattering angle.

5.5.2 Dentin

Boundary Geometry and BSDF The thickness of the enamel layer largely affects the amount of light reaching the dentin core of the tooth and has therefore significant impact on the overall appearance. We allow varying the thickness by modeling the enamel-dentin interface using a blend-shape model defined by a set of deformations:

$$\mathbf{V} = \mathbf{B}_0 + \mathbf{c} \cdot (\mathbf{B} - \mathbf{B}_0)^{\text{T}}. \quad (5.5)$$




Figure 5.5: **Dentin deformation modes.** The fully deformed (a) dentin blend shape of an incisor serves as a boundary defining enamel thickness. For optimization purposes it is further decomposed in three deformation modes along the major axes extracted according to the canonical tooth space such that they represent the following directions: (a) buccal, (b) proximal, (c) occlusal.

We denote the resulting shape \mathbf{V} , its individual vertices \mathbf{v}_i , the meshes representing the deformation by each mode \mathbf{B} , and the vector weights \mathbf{c} . We chose to represent the deformations with respect to an initial mesh \mathbf{B}_0 as it enables intuitive visualization of each mode. Through experimentation we came to the conclusion that three modes are enough to capture most of the visible changes in scattering introduced by the difference of enamel thickness. The modes represent the deformation along the three major axes in a canonical tooth space; see the illustration for a single tooth in Fig. 5.5.

We use a highly simplified BSDF to model the interactions of light with the enamel-dentin interface, which assumes an *index-matched* boundary, i.e. light passes through the interface unaltered.

Volumetric Properties The directional scattering is modeled analogously to enamel, i.e. using a two-lobe phase function.

5.5.3 Gingiva and oral cavity

Gingiva or gums and the oral cavity in general do have an indirect impact on the appearance of teeth since they scatter skin tinted light into the tooth. In order to explain that stray light and prevent the optimization from predicting a tinted tooth we have to model and include the gums and a substitute for the oral cavity walls.

The gums are modeled as a triangle mesh wrapping around the outer boundary of the root of the teeth. As we can directly observe the gums, we use the same set of parameters as the enamel to extract their appearance through numerical optimization with fixed refractive index $\eta = 1.33$.

The oral cavity is highly occluded and non-trivial to capture, however it has a great impact on the appearance as light bounces from its walls and illuminates areas in shadow. We use a substitute mesh with a diffuse reflectance set as free parameter. Its starting point is derived by using similarity theory and average values derived from previous optimization of the scattering parameters of gingiva.

5.6 Model Parameter Optimization

We formulate the problem of extracting the parameters that describe the appearance of teeth as a non-linear optimization problem. Assuming we have M_p pixels in each of the rendered and objective images, and a vector \mathbf{x} with P_p parameters to be optimized according to a loss function $\rho(\psi) : \mathbb{R}^{M_p} \rightarrow \mathbb{R}$ and a cost function $\psi(\chi) : \mathbb{R}^{P_p} \rightarrow \mathbb{R}^{M_p}$, we formally define the set of optimal parameters

$$\chi^+ = \arg \min_{\chi} \rho(\psi(\chi)). \quad (5.6)$$

Since light transport in participating media is generally difficult to evaluate in closed form, the cost function cannot be computed analytically. Instead, we take the route of differentiable path tracing by computing analytic derivatives associated with each *light path* and accumulating them to compute approximate global derivatives to guide a nonlinear optimization algorithm, thus iteratively solving Eq. 5.6. Since the number of light paths is finite—we use 64 samples per pixel—we obtain a noisy estimate of the true gradient. As such, we take the first-order stochastic gradient descend (SGD) approach proposed by Kingma and Ba [2014], which was demonstrated to be sufficiently reliable for solving problems with such characteristics. A similar decision was made by Gkioulekas et al. [2016] who used AdaDelta [Zeiler 2012] in their Monte Carlo based SGD approach for inverse rendering of participating media.

The derivative for each pixel with respect to a free parameter is expressed starting from the loss function $\rho(\psi)$, then applying the change according to image-reconstruction filter F used when sampling the camera pixel footprint, and finally applying the individual gradients $\nabla \psi_{\chi}$ according to the respective

parameters,

$$\nabla\rho = \frac{d\rho}{dF} \frac{dF}{d\psi_\chi} \nabla\psi_\chi. \quad (5.7)$$

We accumulate the first order derivatives at each interaction with a medium, surface or shape containing free parameters separately by exploiting the generalized product rule. Furthermore, we assume a pinhole camera with box filtering function dependent only on the number of samples N_s and the L2 loss function, resulting in the final equation:

$$\begin{aligned} \nabla\rho &= \frac{1}{N_s} (C - C_{\text{ref}}) \cdot \left(\prod_{i=0}^K \psi_\chi^{(i)} \right) \left(\sum_{i=0}^K \frac{\nabla\psi_\chi^{(i)}}{\psi_\chi^{(i)}} \right) \\ &= \frac{1}{N_s} (C - C_{\text{ref}}) \cdot \mathcal{T} \left(\sum_{i=0}^K \frac{\nabla\psi_\chi^{(i)}}{\psi_\chi^{(i)}} \right), \end{aligned} \quad (5.8)$$

where C and C_{ref} are vectors holding linearized rendered and reference RGB images, respectively, and N_s denotes the number of samples and K_b the maximum number of bounces. The complete throughput is expressed as a matrix \mathcal{T} , thus separating the specific change introduced by each interaction from the complete computation of light propagation.

Three groups of parameters are treated within our optimization pipeline: the dentin shape as well as the surface and subsurface parameters of dentin, enamel, and gums. We discuss them in their own separate sections and show them in Fig. 5.6. For the purpose of optimization we define the respective block of gradients

$$\nabla\psi_\chi = \nabla\psi_{\mathbf{B}} \parallel \nabla\psi_{\sigma_t} \parallel \nabla\psi_\alpha \parallel \nabla\psi_{f_p} \parallel \nabla\psi_f \quad (5.9)$$

as the concatenation of the gradients w.r.t. the deformable dentin shape model ($\nabla\psi_{\mathbf{B}}$), volume extinction ($\nabla\psi_{\sigma_t}$), single-scattering albedo ($\nabla\psi_\alpha$), the phase function ($\nabla\psi_{f_p}$), and the BSDF surface model ($\nabla\psi_f$). The gradients of subsurface parameters ($\nabla\psi_{\sigma_t}$, $\nabla\psi_\alpha$, $\nabla\psi_{f_p}$) are further composed of separate groups of RGB components associated with dentin and enamel.

5.6.1 Enamel and Gums

Enamel and gums are treated similarly by the optimization as they are both modeled to share the same types of optical properties such as a homogeneous volume and an absorption layer. For brevity we discuss the case of enamel. The gums are treated analogously. The only difference is that enamel encloses another internal volume (dentin) with free optical and geometrical properties.

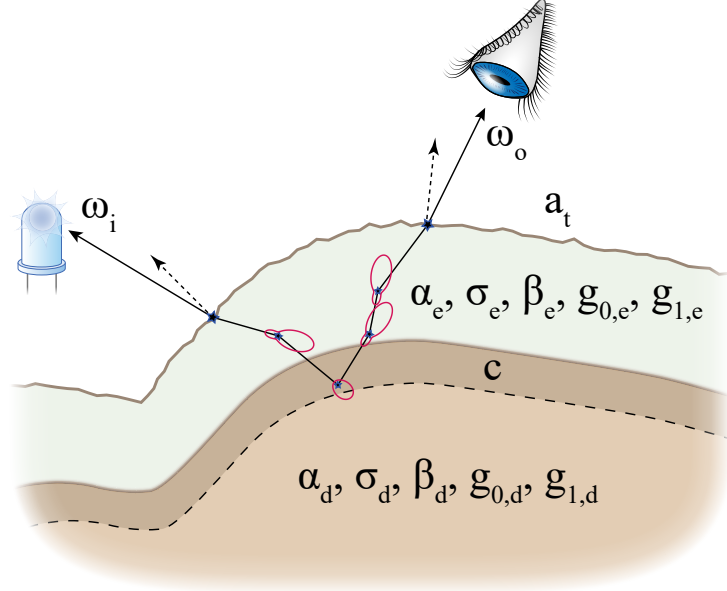


Figure 5.6: **Model parameters.** We jointly optimize surface absorption (a_t), enamel scattering parameters ($\alpha_e, \sigma_e, \beta_e, g_{1,e}, g_{2,e}$), optical and geometry parameters of dentin ($\alpha_d, \sigma_d, \beta_d, g_{1,d}, g_{2,d}, \mathbf{c}$). We further create a relation between enamel and dentin parameters by smoothing the border between media which establishes correlation between scattering and geometry parameters.

Air-Enamel boundary The BSDF model introduces a linear dependency on its transmissive absorption coefficient. Excluding the reflectance term, the derivative at each transmission is completely defined by the light transport equation, which leads to a very simple equation for the accumulated term:

$$\frac{\nabla \psi_f^{(i)}}{\psi_f^{(i)}} = \frac{1}{a_t}. \quad (5.10)$$

Phase function We use a dual-phase function leading to three sets of parameters with the following accumulated term:

$$\frac{\nabla \psi_{f_p}}{\psi_{f_p}} = \frac{1}{f_p(\omega_i, \omega_o)} \left(\frac{df_p}{dg_1} \parallel \frac{df_p}{dg_2} \parallel \frac{df_p}{d\beta} \right). \quad (5.11)$$

We use a linear blend between the phase functions leading to the following relation for the analytic gradient with respect to the blend factor:

$$\frac{df_p}{d\beta} = f_p(\omega_i \cdot \omega_o, g_1) - f_p(\omega_i \cdot \omega_o, g_2). \quad (5.12)$$

The derivative with respect to the mean cosine g_j of the j -th lobe is

$$\frac{df_p}{dg_j} = -b(j) \left(\frac{2g_j}{1-g_j^2} + \frac{3g_j + 3\cos\theta}{1+g_j^2+2g_j\cos\theta} \right) f_{\text{HG}}(\omega_i, \omega_o, g_j), \text{ where} \quad (5.13)$$

$$b(j) = \begin{cases} \beta, & \text{for } j = 1, \\ 1 - \beta, & \text{for } j = 2. \end{cases} \quad (5.14)$$

These derivatives need to be further modified to account for correlation with other parameters. We employ a blend shape model to approximate the geometry of dentin. Due to discontinuity introduced between the two media by the mesh, we cannot compute derivatives expressing correlation between scattering and geometry parameters in path space. To alleviate this issue we interpolate the phase function parameters of dentin and enamel according to a smooth function in spatial domain based on distance q to the dentin shape. We use a logistic sigmoid,

$$l = \frac{1}{1 + e^{-q/s_l}}, \quad (5.15)$$

$$\frac{dl}{dq} = \frac{e^{q/s_l}}{s_l(1 + e^{q/s_l})^2}, \quad (5.16)$$

where we define the spread factor s_l such that the major influence is 1mm around the border. We use this weight to linearly interpolate between enamel and dentin phase function parameters,

$$\beta_{\text{mix}} = \beta_e l + (1 - l)\beta_d \quad (5.17)$$

$$g_{\text{mix}} = \mathbf{g}_e l + (1 - l)\mathbf{g}_d, \quad (5.18)$$

We are going to re-express the derivative according to phase function parameters with respect to blend weights from the perspective of a point inside enamel,

$$\begin{aligned} \frac{\nabla\psi_{f_p, \mathbf{c}}}{\psi_{f_p, \mathbf{c}}} &= \frac{1}{f_p(\omega_i, \omega_o)} \left(\beta_{\text{mix}}(g_{1,e} - g_{1,d}) \frac{df_p}{dg_{0,\text{mix}}} + \right. \\ &\quad \left. (1 - \beta_{\text{mix}})(g_{2,e} - g_{2,d}) \frac{df_p}{dg_{1,\text{mix}}} + \right. \\ &\quad \left. (\beta_e - \beta_d)(f_{\text{HG}}(g_{1,e}) - f_{\text{HG}}(g_{2,e})) \frac{dl}{dq} \frac{dq}{d\mathbf{c}} \right). \end{aligned} \quad (5.19)$$

This relationship can be expressed symmetrically for dentin by exchanging the respective mean cosine coefficients g . Both derivatives with respect to

enamel and dentin parameters are computed at the same time by applying their respective weight l and $(1-l)$. We refer to Appendix 5.A for more details about the derivative according to distance and the dedicated Section 5.6.2 on dentin geometry for specific details about transforming from derivative according to distance to derivative according to blend weights \mathbf{c} .

Single-scattering albedo Assuming a piece-wise constant medium, the light gets attenuated by the scattering coefficient $\sigma_s = \alpha\sigma_t$. The derivative is analogous to the one of the BSDF function, given as

$$\frac{\nabla\psi_\alpha^{(i)}}{\psi_\alpha^{(i)}} = \frac{1}{\alpha}. \quad (5.20)$$

Therefore the modified derivative can be expressed with respect to the the blend shape weights analogously to the phase function parameters,

$$\frac{\nabla\psi_{\alpha_{\text{mix}}}^{(i)}}{\psi_{\alpha_{\text{mix}}}^{(i)}} = \frac{1}{\alpha_{\text{mix}}}(\alpha_e - \alpha_d)\frac{dl}{dq}\frac{dq}{d\mathbf{c}}, \quad (5.21)$$

The difference between the high albedo values is much larger, resulting in poor scaling of the problem. To alleviate this issue we re-expressed the parameter in a dual domain, $\alpha = 1 - 10^\kappa$, and alter the respective derivative,

$$\frac{\nabla\psi_\kappa^{(i)}}{\psi_\kappa^{(i)}} = \frac{\nabla\psi_{\alpha_{\text{mix}}}^{(i)}}{\psi_{\alpha_{\text{mix}}}^{(i)}}10^{-\kappa}\log 10. \quad (5.22)$$

Extinction coefficient The density (precisely the extinction coefficient) acts on both transmission and scattering. In the simple case of a homogeneous medium, the derivative depends on two terms. The first term ($\nabla\psi_{\sigma_t,s}/\psi_{\sigma_t,s}$) stems from the scattering coefficient applied at each scattering event. The second term accumulated over each path segment ($\nabla\psi_{\sigma_t,t}/\psi_{\sigma_t,t}$) derives from the Beer-Lambert law and the path-length dependence. The collective derivative is

$$\frac{\nabla\psi_{\sigma_t}}{\psi_{\sigma_t}} = \frac{\nabla\psi_{\sigma_t,s}}{\psi_{\sigma_t,s}} + \frac{\nabla\psi_{\sigma_t,t}}{\psi_{\sigma_t,t}} = \frac{N_s - 1}{\sigma_t} - \sum_{k=0}^{N_s} t_k, \quad (5.23)$$

where N_s is the number of path segments and t_k is the length of the k -th path segment. We further reparameterize in a dual domain, $\sigma_t = e^\xi$, yielding

$$\frac{\nabla\psi_\xi}{\psi_\xi} = e^\xi \left(\frac{N_s - 1}{\sigma_t} - \sum_{k=0}^{N_s} t_k \right). \quad (5.24)$$

5.6.2 Dentin

Similarly to the enamel and the gums we optimize the phase function, density, and albedo for dentin. We also leverage the blend shape model described in Sec. 5.5.2 to enable optimizing the geometry of dentin, which has a crucial impact on the appearance.

Dentin geometry Modeling deformable surfaces within the context of differentiable path tracing requires the exact computation of analytic gradients to establish the connection between deformation parameters and the effect on light transmittance through participating medium. We are going to discuss the specific case of index-matched boundary between two homogeneous participating media. The transmittance in this case is a composition of three operations: deformation ($\psi_{\mathbf{G}}$), intersection distance (ψ_d), and transmittance (ψ_t). Collectively, this subset of the light transport evaluation can be expressed as function composition,

$$\psi_B(\mathbf{c}) = (\psi_t \circ \psi_d \circ \psi_{\mathbf{G}})(\mathbf{c}). \quad (5.25)$$

We will break down the set of operations, starting from the very last one and build the respective transformations required to reach the final solution. We express the *deformation* model as already described in Eq. 5.5. To simplify the application of the chain rule against the derivative of the per triangle intersection we are going to restructure the matrix of vertices $\mathbf{v}^{(m)} = [\mathbf{v}_0^{(m)} \ \mathbf{v}_1^{(m)} \ \mathbf{v}_2^{(m)}]$ into a tensor of vertex triplets associated with each ray intersection $\mathbf{G}_T = [\mathbf{v}^{(0)} \ \dots \ \mathbf{v}^{(n)}]$, with n being the number of intersections. The change according to the model parameters $\mathbf{c} = [c_0 \ \dots \ c_k]$, which are k in total, in this case is expressed by the change of shape $\nabla\psi_{\mathbf{G}} : \mathbb{R}^k \rightarrow \mathbb{R}^{k \times n \times 3 \times 3}$ derived from Eq. 5.5 as

$$\nabla\psi_{\mathbf{G}} = \mathbf{B} - \mathbf{B}_0 = \Delta\mathbf{B}. \quad (5.26)$$

The next relation that is crucial to derive is the change of *intersection distance* w.r.t. the model. It is governed by the ray-triangle intersection. The intersection point $\mathbf{p}_r^{(m)}$ is expressed as a the distance $t^{(m)}$ traversed

along the ray in direction \mathbf{r}_d starting from point \mathbf{r}_o :

$$\mathbf{p}_r^{(m)} = \mathbf{r}_o + \mathbf{r}_d t^{(m)}, \quad (5.27)$$

$$t^{(m)} = \frac{\begin{vmatrix} \mathbf{v}_2^{(m)} - \mathbf{v}_0^{(m)} \\ \mathbf{r}_o - \mathbf{v}_0^{(m)} \\ \mathbf{v}_1^{(m)} - \mathbf{v}_0^{(m)} \end{vmatrix} \begin{vmatrix} \mathbf{v}_1^{(m)} - \mathbf{v}_0^{(m)} \\ \mathbf{r}_d \\ \mathbf{v}_2^{(m)} - \mathbf{v}_0^{(m)} \end{vmatrix}^{-1}}{|\mathbf{D}_1| |\mathbf{D}_2|^{-1}} = |\mathbf{D}_1| |\mathbf{D}_2|^{-1}. \quad (5.28)$$

The transformation to distance space involves the computation of a Jacobian matrix which represents a linear transformation such that $\nabla\psi_d : \mathbb{R}^{k \times n \times 3 \times 3} \rightarrow \mathbb{R}^{k \times n}$.

For convenience we define the vector columns of the adjugate matrix as $\mathbf{h}_m^{(k)} \equiv \text{col}_m \text{adj}(\mathbf{D}_k)$. By using the definition of the derivative of a determinant we express the columns of the Jacobian ($\nabla\psi_d^{(m)} \equiv \text{col}_m \nabla\psi_d$) with the following relation:

$$\nabla\psi_d^{(m)} = t^{(m)} \begin{bmatrix} -\mathbf{h}_0^{(1)} - \mathbf{h}_1^{(1)} - \mathbf{h}_2^{(1)} \\ \mathbf{h}_2^{(1)} \\ \mathbf{h}_0^{(1)} \end{bmatrix} - \frac{t^{(m)}}{|\mathbf{D}_2|} \begin{bmatrix} -\mathbf{h}_0^{(2)} - \mathbf{h}_2^{(2)} \\ \mathbf{h}_0^{(2)} \\ \mathbf{h}_2^{(2)} \end{bmatrix}. \quad (5.29)$$

The final part is computing the transformation to transmittance space. As outlined in Fig. 5.7 the goal is to compute the effect a change of shape $\Delta\mathbf{B}$ has on the light transport along a segment t . Lets assume the segment crosses a boundary, exiting from a set of media whose optical properties are given by \mathbf{S}_e and entering into another one characterized by \mathbf{S}_i . Then the transmittance at the border (excluding the endpoints) is expressed as

$$\psi_t = \exp((\mathbf{S}_i - \mathbf{S}_e) \cdot \mathbf{t}), \quad (5.30)$$

where $\mathbf{S} = [\sigma_{t1} \dots \sigma_{tn}]$ is a vector of the extinction coefficients which is further split into a vector of entered (\mathbf{S}_i) and exited media (\mathbf{S}_e). These coefficients are appended to \mathbf{S}_i and \mathbf{S}_e at each successive crossing between media. The Jacobian transformation $\nabla\psi_t : \mathbb{R}^{k \times n} \rightarrow \mathbb{R}^k$ is

$$\nabla\psi_t = (\mathbf{S}_i - \mathbf{S}_e) \exp((\mathbf{S}_i - \mathbf{S}_e) \cdot \mathbf{t}). \quad (5.31)$$

We are now able to express the entire computation of the accumulated term by also taking into account the common terms with the complete light transport:

$$\frac{\nabla\psi_{\mathbf{B}}(\mathbf{c})}{\psi_{\mathbf{B}}(\mathbf{c})} = \frac{1}{\psi_{\mathbf{B}}(\mathbf{c})} \nabla\psi_d \nabla\psi_t \nabla\psi_{\mathbf{G}} = (\mathbf{S}_i - \mathbf{S}_e) \nabla\psi_d \nabla\psi_{\mathbf{G}}. \quad (5.32)$$




Figure 5.7: A path segment going from one medium into another will be influenced when changing the shape \mathbf{B} as it will pass more through one of the media and less through the other.

This relation shows that when entering an optically thicker medium with triangles translated in the positive direction w.r.t. to the ray direction, the transmittance increases as the ray traverses less distance through the medium with higher extinction (Fig. 5.7).

5.7 Results

We implemented the optimization procedure in Python by using the scientific computing package Numpy as a building block of our custom implementation of Adam. We implemented custom volumetric integrator in Mitsuba [Jakob 2010] which enabled path tracing of derivatives with respect to the optimized parameters and developed the respective BSDF, blend shape and data structures used to store the parameters. We used the Python bindings to interface with the renderer and exchange free parameters χ , the resulting derivative blocks $\nabla\psi_\chi$, and rendered image C . We split derivative estimation and image rendering in two independently seeded rendering passes. Our optimization and rendering runs on a dual CPU machine equipped with Intel Xeon E5-2680v3 processors. We run the optimization at very low resolution (approximately a fifth of the captured resolution across each dimension) to keep the execution times feasible.

We setup the virtual optimization scene based on calibration data of teeth geometry, camera, and light positions and optimize against our stabilized images. We associate free parameters with the appearance properties of all objects in the scene: teeth, gingiva and oral cavity. We use a point light representation and the recovered positions to simulate our calibrated

LED light sources. For the case of a single tooth scene we derive the volume parameters starting point by densely exploring the space spanned by the values reported in literature [Spitzer and Ten Bosch 1975; Fried et al. 1995; Pop-Ciutnila et al. 2015] for both dentin and enamel and comparing the rendered results with our reference measurements. This step is necessary since there is a wide range of values reported in literature mainly due to the various pre-processing procedures performed on the measured sample surfaces which affect roughness and consecutively the appearance. We tabulated the three dimensional domain of albedo, density, and phase function anisotropy into a set of 50000 points and evaluated a scale invariant error of a virtual homogeneous tooth against reference measurements of a masked tooth illuminated from various directions and viewed by different cameras. For each channel we find the point in the domain with the smallest error and use it as our initialization. We do not perform highlight or high frequency normal map fitting. For the case of a complete set of teeth we initialize the appearance of all teeth with the converged values derived by the stochastic single tooth optimization.

We first explore the capability of our technique to estimate scattering parameters of teeth and re-render them in virtual environments by using isolated dead teeth. We acquired micro-CT scans to extract the exact geometry of enamel and dentin. We only mask highlights for this scene since the tooth is static. In Fig. 5.8 we provide qualitative comparisons of our training set for a single view (left) and additional extrapolated appearance renderings (middle, right) on a testing set with novel viewpoints corresponding to images captured by other cameras but were not visible to the optimization. The procedure is initialized with average values derived from the dental literature. Overall our method provides plausible fits in the case of an isolated dead tooth. Dead teeth contain details which are less pronounced in live teeth. One example is dead tissue and cracks which cast volumetric shadows within the tooth volume that our model does not explicitly support. Other discrepancies in the highlights are caused by the reflection of the metal frame of our dome. These preliminary results validated the plausibility of our technique to estimate parameters of live teeth.

We proceed with capture of live teeth in vivo of a few human volunteers. We provide separate fits by progressively enabling more parameters as shown in Fig. 5.9. Optimizing with a homogeneous enamel-only volume results in very smooth appearance which lacks detail. To capture the translucency of teeth around the edges we add an inner layer to model dentin. Finally, we

model the missing surface details by enabling a textured stains-absorption layer on the enamel outer boundary. We provide the complete comparison of fitting results in Fig. 5.10, including validation from novel viewpoints and reference images. Overall our method is able to match the appearance characteristics of a live tooth. Note that when a sample location is masked from the camera then the optimization might arrive to sub-optimal results (due to over-fitting), such as the recovered diffuse albedo of the mouth.

The results of our method on the more challenging complete sets of teeth are shown in Fig. 5.11, Fig. 5.12, Fig. 5.13 on three different volunteers. Additional data regarding the initialization values, textures, masks and loss function values is provided in the supplemental material. All scenes use independent parameters for each tooth that are jointly optimized. With these scenes we test the capability of our model to capture realistic teeth appearance in the presence of gums and inter-reflections from neighboring teeth. We would like to point out that the measurements with the left camera shown in Fig. 5.13 were not used in the optimization. Our model is optimized using only the information of center and right cameras. The left camera comparison is an evaluation of the ability of our method to extrapolate to novel view points. Overall our method performs better for teeth in the front part of the mouth where shadows from the oral cavity occur less frequently than the teeth in the back. In the case of a shadow boundary in our measurements we mask the nearby affected teeth since we do not have an accurate reconstruction of the oral cavity and thus we can not simulate the shadow accurately. Most of the molars are usually masked out in our measurements and consequently some of them have false appearance. An accurate reconstruction of the entire oral cavity and the lip retractor might be required to achieve better fits on these particular teeth along with a regularization term. Regularizing the optical properties of these teeth to the average values of the frontal teeth is one option that we tested and partially alleviates the issue. We note that the un-textured reflectance of the oral cavity proxy in this case reaches a compromise between the inner color and the white lip retractor, resulting in ambient illumination which is specularly reflected by some teeth in shadowed areas. The ambient light reflected by the outside environment and light dome is not modeled, similarly to the single tooth experiment. It leads to further discrepancies with the real physical properties that are encoded in the current set of parameters (optical thickness and texture). Capturing a light probe along with the actor can be used in the future to account for the missing light. We experimented with different loss functions (L1, rela-

tive L1, L2, relative L2) and came to the conclusion that L2 loss function produces higher quality results as it applies a significantly lower weight on dark and shadowed regions where most of the error due to our measurement methodology is concentrated.

5.8 Discussion and Future Work

We presented a system for capturing bulk properties that are key for digitally reproducing the appearance of teeth. Our approach is the first to consider the internal structure of each tooth when reverse-engineering its optical properties and while the results match the photographs well, we would like to point out a number of limitations of our system and how they could be addressed in future work.

Our acquisition setup is very similar to those used for facial-appearance capture, so suitable measurements could likely be made with capture rigs already used in industry. Furthermore, a dense capture setup is actually not required, as Fig. 5.10 shows an optimized fit using only 3 cameras and 5 light sources.

Our parametric model currently does allow for capturing craze lines and cracks in enamel, and assumes the tooth is free of cavities and fillings. These can potentially impact the appearance (e.g. amalgam fillings tend to significantly darken the entire tooth which can be observed on the right molar in Fig. 5.13) but we believe they need to be encoded in the geometric representation of the tooth and would require a joint optimization of the appearance and the outer tooth geometry. Such optimization is significantly more challenging than inferring the outer geometry and appearance independently. Another option is to consider dental filling and cracks as a morphable mesh similarly to the dentin inner shape and perform joint optimization in similar manner. In this work, we focused on making the first step towards capturing the appearance parameters and the dentin geometry, but we believe the joint optimization represents an interesting avenue for extending our work.

We heavily simplified the simulation of light transport at dielectric boundaries, completely disregarding the effect of light bending upon refraction. While this purely performance-motivated simplification did not cause visually significant differences in our experiments, we would like to note that implementing search-based methods [Walter et al. 2009; Holzschuch 2015] or manifold walks [Jakob and Marschner 2012] would permit simulating light

transport more accurately at acceptable render times. We also did not attempt to capture high-frequency glossy highlights. These could be added by subtracting the main diffuse reflectance and performing standard BRDF fitting separately according to one of the existing techniques in the graphics literature [Guarnera et al. 2016].

The volumetric model can be further improved by considering anisotropic particle distributions [Jakob et al. 2010] that would better represent the tubular structure of enamel. Another option is to enable heterogeneity and spatial variation of the scattering parameters of enamel and dentin. In this work, we opted for homogeneous approximations with anisotropic phase functions to avoid excessively long rendering times. Another danger of adding more parameters to the model is over-fitting and it may require regularization and priors to constrain the optimization. A careful analysis of potential gains in visual fidelity is to be carried out first before committing to more sophisticated models; our experiments did not reveal such needs but we did not attempt to model heavily stained and damaged teeth where additional parameters could potentially help.

The choice to use per-path analytic gradients for each parameter enables the cancellation or reuse of certain terms within the implementation. We further optimize the layout of the block-sparse gradients and avoid a significant amount of computation. Another alternative is to use automatic differentiation which would need to be applied to a substantial part of the rendering system, including the carefully optimized ray/triangle intersection code, potentially entailing additional performance and storage costs.




Figure 5.8: **Dead Tooth Validation.** We validate our appearance estimation fit on a dead incisor with known dentin geometry. Left: 9 light directions from front (top to bottom) are shown from the front camera. We optimize the appearance using the 9 reference images as targets, initialized with default appearance parameters as shown. Our fit result closely matches the reference images. Center: Extrapolating the fit parameters to a side view validates that the volumetric appearance is reasonably close to the real tooth. Right: We extrapolate further to 9 different side light directions, as seen from the front camera view.




Figure 5.9: **Model Validation.** All three components of our appearance model are required to generate realistic teeth appearance. Fitting only a single homogeneous enamel volume appears too opaque (left). Fitting both enamel and our dentin layer allows spatially-varying translucency (left-center). Adding a stain texture layer on top provides the best fit (right-center) to the reference image (right).

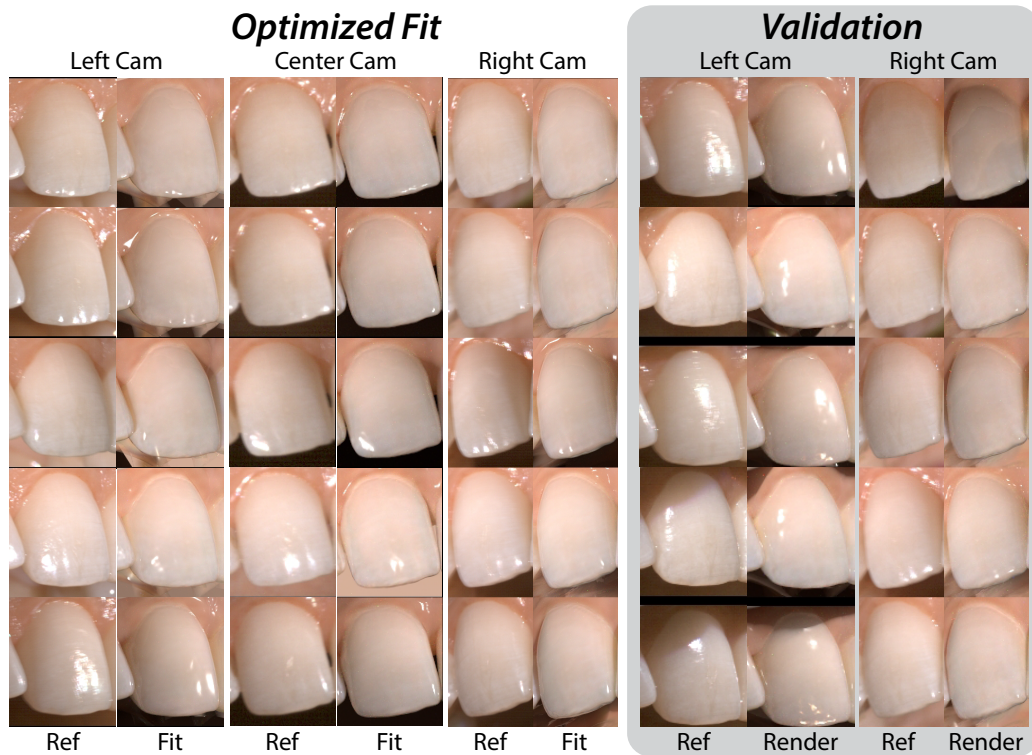


Figure 5.10: **Single Tooth Fitting and Validation.** Here we show the accuracy of our method on a single tooth. Left: The optimized appearance fit to three cameras (columns) and five lighting conditions (rows). Our fit tooth appearance closely matches the reference photo. Right: We validate the recovered appearance properties by rendering under five new lighting directions that were not part of the optimization. Our model extrapolates well to these unseen conditions.




Figure 5.11: **Upper Jaw Fitting: Volunteer 1.** We compare the optimized result of our model and optimization on a complete upper jaw consisting of 14 teeth. We show three different viewpoints of the same person’s teeth next to corresponding reference images at the same resolution.

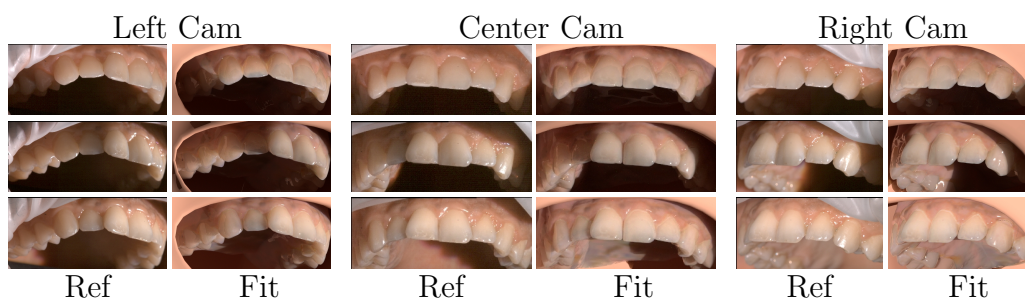


Figure 5.12: **Upper Jaw Fitting: Volunteer 2.** Fitting results on the entire upper jaw of another volunteer. We show three different viewpoints of the same person’s teeth and comparisons to corresponding reference images.




Figure 5.13: **Upper Jaw Fitting: Volunteer 3.** Fitting results on the entire upper jaw of another volunteer. In this experiment we only use the center and right cameras during the optimization. The left camera image comparison is a qualitative, modeling power evaluation of our method when used to extrapolate to novel views.

Appendix

5.A Derivative of distance with respect to vertex translation

Computing the derivative of distance to a point w.r.t. translation of individual vertices is essential when computing the contribution of spatially varying parameters defined by a blend shape model. We will consider these three primitives being closest to a point: a point, an edge or a plane. These are three cases when computing the distance between a point and a triangle according to the method described by Eberly [2015]. Note that we are going to assume that within infinitesimal translations the closest primitive does not change.

Derivative of the distance w.r.t. a triangle vertex. The distance,

$$q_{\mathbf{p}} = |\mathbf{p}_s - \mathbf{v}|, \quad (5.33)$$

is defined by the spatial position of a point \mathbf{p}_s , and the closest triangle vertex \mathbf{v} . The derivative is computed by differentiating according to each respective coordinate of the vertex and populating the respective column of the triangle derivative matrix $\frac{dq_{\mathbf{p}}}{d\mathbf{G}_T}$,

$$\frac{dq_{\mathbf{p}}}{d\mathbf{v}} = -\frac{\mathbf{p}_s - \mathbf{v}}{|\mathbf{p}_s - \mathbf{v}|} \quad (5.34)$$

Derivative of the distance to a triangle edge. The distance to a triangle edge is expressed according to the vector orthogonal to the line and

connecting the target point,

$$\mathbf{t} = \frac{\mathbf{v}_1 - \mathbf{v}_0}{|\mathbf{v}_1 - \mathbf{v}_0|} \quad (5.35)$$

$$\mathbf{q}_p = (\mathbf{p}_s - \mathbf{v}_0) - \mathbf{t}(\mathbf{t} \cdot (\mathbf{p}_s - \mathbf{v}_0)) \quad (5.36)$$

$$q_p = |\mathbf{q}_p|, \quad (5.37)$$

where we use the projection on the tangent \mathbf{t} determined by the two vertices \mathbf{v}_0 and \mathbf{v}_1 to derive the final equation. Differentiating this equation is more involved and we provide just the final set of equations for reference,

$$\mathbf{A}_{\mathbf{v}_0} = \left(\frac{\mathbf{t} \cdot (\mathbf{p}_s - \mathbf{v}_0)}{|\mathbf{v}_1 - \mathbf{v}_0|q_p} - 1 \right) \mathbf{q}_p \quad (5.38)$$

$$\mathbf{A}_{\mathbf{v}_1} = -\frac{\mathbf{t} \cdot (\mathbf{p}_s - \mathbf{v}_0)}{|\mathbf{v}_1 - \mathbf{v}_0|q_p} \mathbf{q}_p. \quad (5.39)$$

The individual derivatives according to each vertex ($\mathbf{A}_{\mathbf{v}_0}, \mathbf{A}_{\mathbf{v}_1}$) are further assigned to the correct corresponding columns in the triangle derivative matrix $\frac{dq_p}{d\mathbf{G}_T}$.

Derivative of the distance to a triangle interior plane. The distance to a triangle interior point can be expressed by taking into account the normal vector \mathbf{n} defined by the two edge vectors $\mathbf{u}_0 = \mathbf{v}_1 - \mathbf{v}_0$ and $\mathbf{u}_1 = \mathbf{v}_2 - \mathbf{v}_0$,

$$\mathbf{n} = \frac{\mathbf{u}_0 \times \mathbf{u}_1}{|\mathbf{u}_0 \times \mathbf{u}_1|},$$

and projecting the vector connecting a triangle vertex \mathbf{v}_0 and the target point \mathbf{p}_s ,

$$q_p = \mathbf{n} \cdot (\mathbf{p}_s - \mathbf{v}_0). \quad (5.40)$$

We are going to use the columns of the identity matrix $\mathbf{I}_i = \text{col}_i \mathbf{I}$ to derive a compact representation of the derivative,

$$\begin{aligned} \mathbf{A}_{ki} &= -\mathbf{I}_i \times \mathbf{u}_k \\ \mathbf{A}_{\text{cross}} &= \begin{bmatrix} \mathbf{A}_{00} + \mathbf{A}_{10} & \mathbf{A}_{01} + \mathbf{A}_{11} & \mathbf{A}_{02} + \mathbf{A}_{12} \\ -\mathbf{A}_{00} & -\mathbf{A}_{01} & -\mathbf{A}_{02} \\ -\mathbf{A}_{10} & -\mathbf{A}_{11} & -\mathbf{A}_{12} \end{bmatrix} \\ \frac{dq_p^{(i,j)}}{d(\mathbf{v}_0, \mathbf{v}_1, \mathbf{v}_2)} &= \text{sgn}(q_p) \mathbf{A}_{\text{cross}}^{(i,j)} \cdot \left((\mathbf{p}_s - \mathbf{v}_0) + q_p \frac{\mathbf{n}}{|\mathbf{u}_0 \times \mathbf{u}_1|} \right). \end{aligned} \quad (5.41)$$

5.B Optimization process statistics

We show the statistics of the optimization process for our three volunteers in Table 5.1, 5.2, 5.3. Image resolutions that are used in the numerical optimization are shown separately for each camera. These are cropped regions around the mouth and they can contain masked pixels that are not rendered. We provide the number of optimization iterations and optimized parameters per channel and optimization phase. If a two step processed is used both phases must execute, otherwise only the textured phase is used. Our experiments show that a single "textured" phase in this case can be enough. The final renderings that are included in the main body of the paper are with five times smaller resolution across each dimension compared to the captured photographs. The resolution used for rendering and optimization is the same, but with higher number of sample used to render the final images. Wall clock timing is further provided in these tables for the average time per iteration and the total time to complete the optimization process. Additional fine-grain timings are provided for each individual phase and channel, separately. Multiple stages of the pipeline are presented to evaluate their impact on the overall computation. Scene setup involves rebuilding Bounding Volume Hierarchies (BVH) and passing the optimized set of parameters to the renderer. BVH rebuilds are done on a single thread in Mitsuba. A future implementation can apply multithreading to these parts of the code to improve its performance. Rendering the gradients and the reference images for this set of parameters has very similar timings. To reduce the performance impact of this stage we used block-sparse gradients. Finally, the gradients are processed by the optimization algorithm which has the lowest performance impact. Note that all timings are done on a single machine equipped with two Intel Xeon E5-2680v3 processors. Computations can be trivially spread between channels and images across an entire cluster of machines.

The loss function per channel is shown in Fig. 5.14, 5.15, 5.16. We keep track of the lowest loss value to avoid the fact that Adam does not always converge to the lowest minimum while traversing null space.




Figure 5.14: Loss function during the optimization process of the first volunteer's teeth

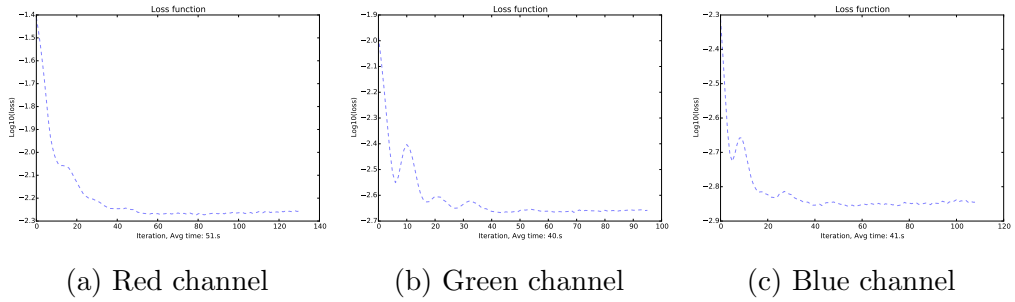


Figure 5.15: Loss function during the optimization process of the second volunteer's teeth

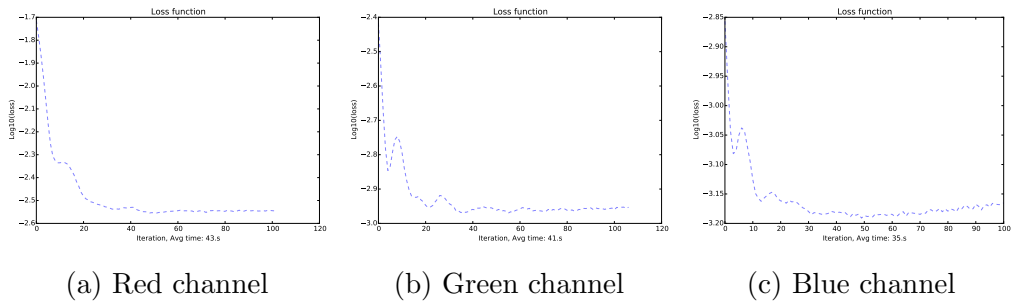


Figure 5.16: Loss function during the optimization process of the third volunteer's teeth

		First volunteer		
Operation		Red	Green	Blue
”Homogeneous” Phase	Parameter count	220	220	220
	Resolution (width \times height)	288 \times 117	288 \times 117	288 \times 117
		385 \times 195	385 \times 195	385 \times 195
		532 \times 233	532 \times 233	532 \times 233
	Iteration count	130	173	185
	Scene setup (average)	4.469 s	4.551 s	4.480 s
	Render gradients (average)	98.256 s	73.765 s	67.763 s
	Render reference (average)	96.558 s	66.230 s	59.496 s
	Process gradients (average)	0.221 s	0.234 s	0.221 s
	Scene setup (total)	580.997 s	787.307 s	828.819 s
	Render gradients (total)	12 773.298 s	12 761.410 s	12 536.103 s
	Render reference (total)	12 552.557 s	11 457.776 s	11 006.741 s
	Process gradients (total)	28.724 s	40.437 s	40.828 s
Total time per channel	25 935.575 s	25 046.929 s	24 412.490 s	
Total time per phase	75 394.994 s			
”Textured” Phase	Parameter count	196810	196810	196810
	Resolution (width \times height)	288 \times 117	288 \times 117	288 \times 117
		385 \times 195	385 \times 195	385 \times 195
		532 \times 233	532 \times 233	532 \times 233
	Iteration count	90	139	173
	Scene setup (average)	7.077 s	7.145 s	5.148 s
	Render gradients (average)	113.477 s	91.518 s	75.169 s
	Render reference (average)	130.821 s	81.202 s	59.586 s
	Process gradients (average)	0.442 s	0.455 s	0.390 s
	Scene setup (total)	636.890 s	993.163 s	890.574 s
	Render gradients (total)	10 212.956 s	12 720.953 s	13 004.201 s
	Render reference (total)	11 773.932 s	11 287.045 s	10 308.423 s
	Process gradients (total)	39.758 s	63.255 s	67.477 s
Total time per channel	22 663.536 s	25 064.415 s	24 270.674 s	
Total time per phase	71 998.625 s			
Total time	147 393.619 s			

Table 5.1: Execution timing of the optimization process of the first volunteer’s teeth

		Second volunteer		
		Red	Green	Blue
”Homogeneous” Phase	Operation			
	Parameter count	205	205	205
	Resolution (width \times height)	311×109	311×109	311×109
		394×164	394×164	394×164
		491×183	491×183	491×183
	Iteration count	114	114	140
	Scene setup (average)	5.712 s	5.590 s	4.171 s
	Render gradients (average)	73.617 s	64.339 s	48.766 s
	Render reference (average)	98.723 s	92.849 s	63.987 s
	Process gradients (average)	0.158 s	0.157 s	0.138 s
	Scene setup (total)	651.171 s	637.272 s	583.920 s
	Render gradients (total)	8392.352 s	7334.702 s	6827.173 s
	Render reference (total)	11 254.390 s	10 584.754 s	8958.113 s
	Process gradients (total)	18.042 s	17.845 s	19.388 s
Total time per channel	20 315.955 s	18 574.573 s	16 388.593 s	
Total time per phase	55 279.121 s			
”Textured” Phase	Parameter count	196795	196795	196795
	Resolution (width \times height)	311×109	311×109	311×109
		394×164	394×164	394×164
		491×183	491×183	491×183
	Iteration count	131	98	105
	Scene setup (average)	4.807 s	4.788 s	4.793 s
	Render gradients (average)	81.862 s	74.103 s	65.001 s
	Render reference (average)	99.146 s	79.609 s	78.463 s
	Process gradients (average)	0.305 s	0.304 s	0.305 s
	Scene setup (total)	629.759 s	469.253 s	503.281 s
	Render gradients (total)	10 723.969 s	7262.116 s	6825.057 s
	Render reference (total)	12 988.100 s	7801.665 s	8238.585 s
	Process gradients (total)	39.938 s	29.837 s	32.052 s
	Total time per channel	24 381.767 s	15 562.871 s	15 598.975 s
Total time per phase	55 543.613 s			
Total time	110 822.734 s			

Table 5.2: Execution timing of the optimization process of the second volunteer’s teeth

		Third volunteer		
Operation		Red	Green	Blue
”Homogeneous” Phase	Parameter count	220	220	220
	Resolution (width \times height)	290×114	290×114	290×114
		524×233	524×233	524×233
		437×211	437×211	437×211
	Iteration count	100	119	119
	Scene setup (average)	6.093 s	6.072 s	4.301 s
	Render gradients (average)	77.520 s	62.845 s	48.050 s
	Render reference (average)	91.295 s	64.934 s	49.017 s
	Process gradients (average)	0.140 s	0.138 s	0.125 s
	Scene setup (total)	609.305 s	722.575 s	511.833 s
	Render gradients (total)	7752.013 s	7478.595 s	5717.914 s
	Render reference (total)	9129.522 s	7727.129 s	5833.075 s
	Process gradients (total)	14.013 s	16.413 s	14.869 s
	Total time per channel	17 504.853 s	15 944.712 s	12 077.692 s
Total time per phase	45 527.256 s			
”Textured” Phase	Parameter count	196810	196810	196810
	Resolution (width \times height)	290×114	290×114	290×114
		524×233	524×233	524×233
		437×211	437×211	437×211
	Iteration count	101	90	99
	Scene setup (average)	5.090 s	7.014 s	7.066 s
	Render gradients (average)	75.960 s	67.092 s	58.403 s
	Render reference (average)	88.083 s	73.596 s	58.436 s
	Process gradients (average)	0.294 s	0.333 s	0.343 s
	Scene setup (total)	514.078 s	631.240 s	699.538 s
	Render gradients (total)	7671.958 s	6038.248 s	5781.920 s
	Render reference (total)	8896.363 s	6623.611 s	5785.133 s
	Process gradients (total)	29.738 s	29.970 s	33.908 s
	Total time per channel	17 112.137 s	13 323.070 s	12 300.499 s
Total time per phase	42 735.707 s			
Total time	88 262.963 s			

Table 5.3: Execution timing of the optimization process of the third volunteer’s teeth

5.C Optimization results

The images that show results from the optimization process are split in three major categories. First, the reference images are as captured in the light dome with masked version additionally provided next to them. Masking is done automatically to remove highlights and manually on parts of the scene that are occluded by elements that are not modeled (lips, oral cavity). Next, the results of the optimization process are shown with the initial values next to them. The initial images are with all parameters picked to represent close to average values from the literature over all wavelengths and thus they appear grey. The final images are shown last with their L2 loss function multiplied by a factor of a hundred. The estimated textures are provided separately for each volunteer. As a rule of thumb, only the directly visible areas have close to correct transmission parameters. Surfaces on the border of gums might contain smooth transitions and underneath the gums they can be with a similar tint. Similarly between the teeth, they can contain spurious shadows.

Our first volunteer is a male in his mid-twenties, smoker with regular diet. His teeth are perfectly healthy and smoking stains are not observable. The first data fits are focused on his right central incisor from multiple camera views (cf. Fig. 5.17 to 5.19). Also, data fits of an entire tooth row from multiple views are shown in Fig. 5.20 to 5.22. Textures are provided in Fig. 5.29.

Our second volunteer is a male in his mid-thirties, smoker with regular diet. His teeth are stained and darker than the other volunteers. Furthermore, his frontal right incisor was chipped and repaired multiple times. The effect of this defect is captured by the stains layer of our model. Data fits of an entire tooth row from multiple views are shown in Fig. 5.23 to 5.25. Textures are provided in Fig. 5.30.

Our third volunteer is a male in his mid-thirties, non-smoker with regular diet. His left molar was repaired with amalgam which is an opaque metallic substance that cannot be explicitly modeled by our approach. The overall effect is a homogeneous change of the volumetric scattering parameters of that tooth since the optimization process tries to compensate for the volumetric shadows cast inside its interior. Data fits of an entire tooth row from multiple views are shown in Fig. 5.26 to 5.28. The first camera is not used in the optimization procedure and therefore it is completely masked-out

(cf. Fig. 5.26). Textures are provided in Fig. 5.31.




Figure 5.17: Results for the right central incisor of the first volunteer from the viewpoint of the first camera

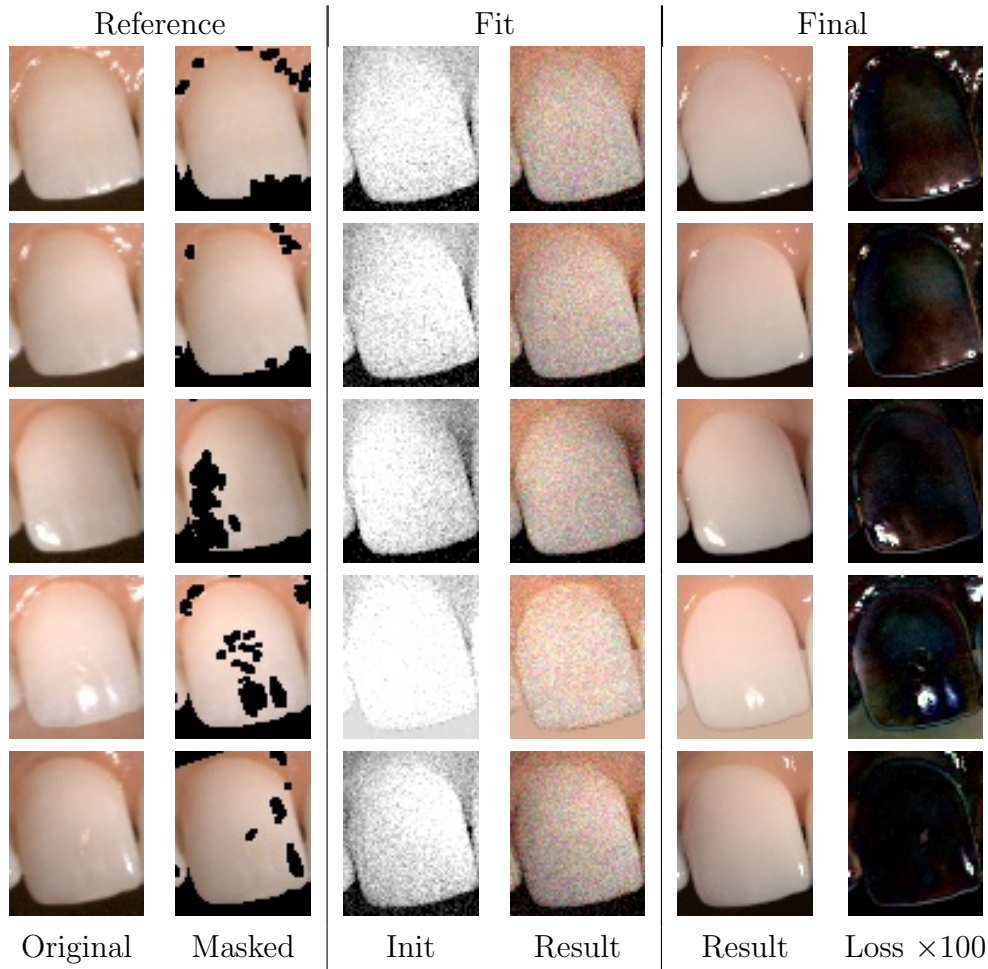


Figure 5.18: Results for the right central incisor of the first volunteer from the viewpoint of the second camera




Figure 5.19: Results for the right central incisor of the first volunteer from the viewpoint of the third camera

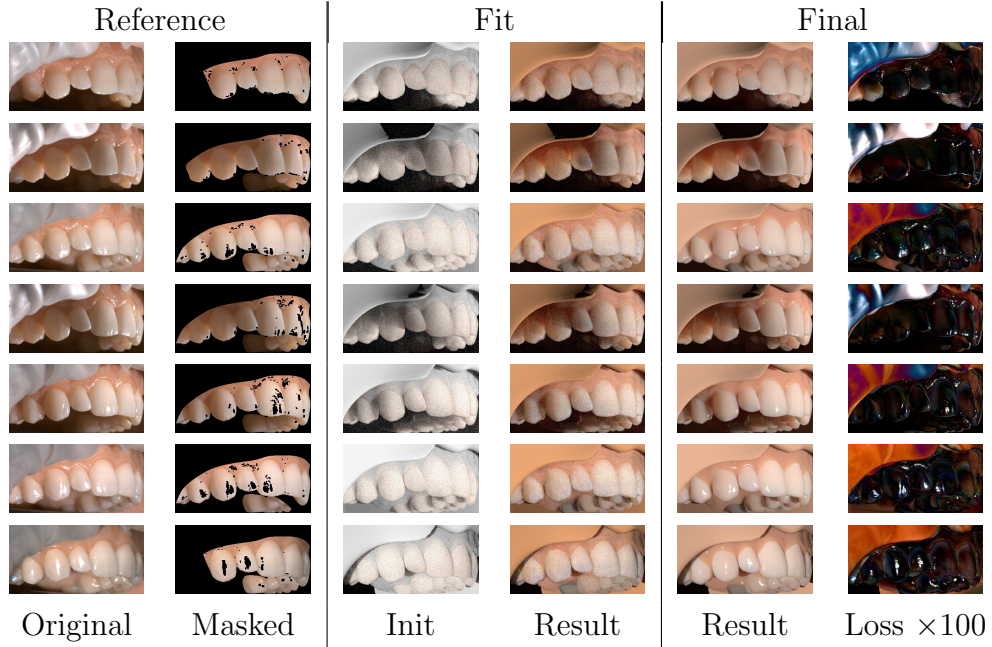


Figure 5.20: Results for the entire tooth row of the first volunteer from the viewpoint of the first camera

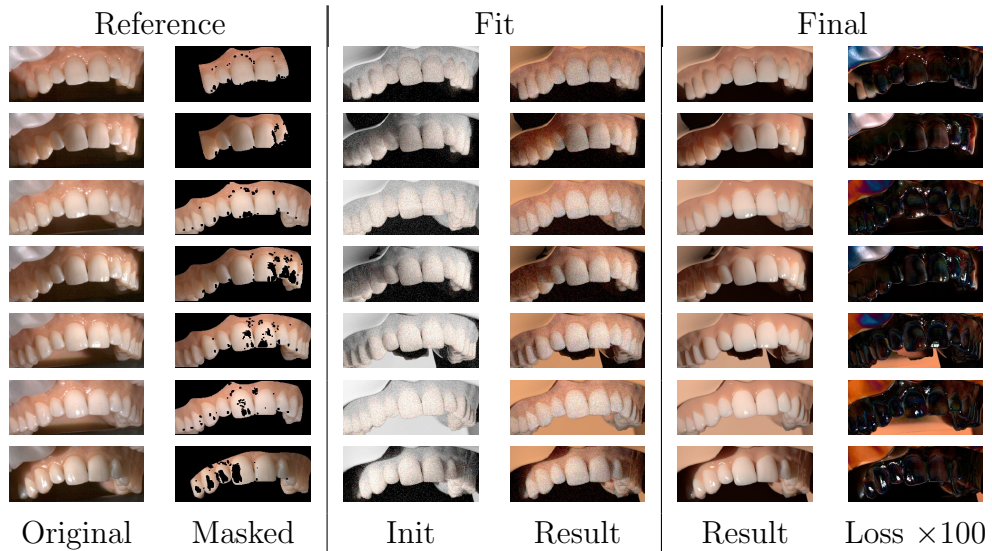


Figure 5.21: Results for the entire tooth row of the first volunteer from the viewpoint of the second camera

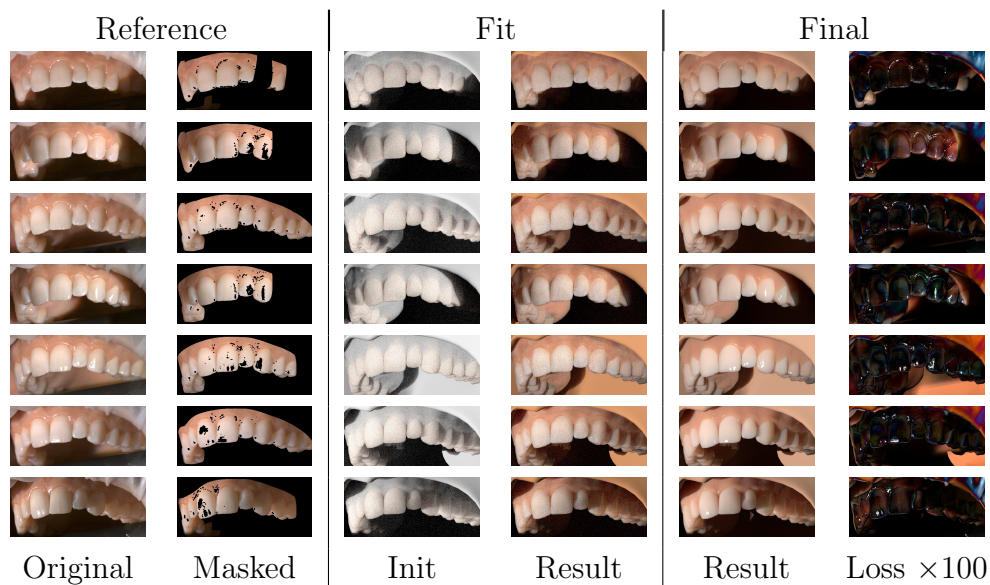


Figure 5.22: Results for the entire tooth row of the first volunteer from the viewpoint of the third camera

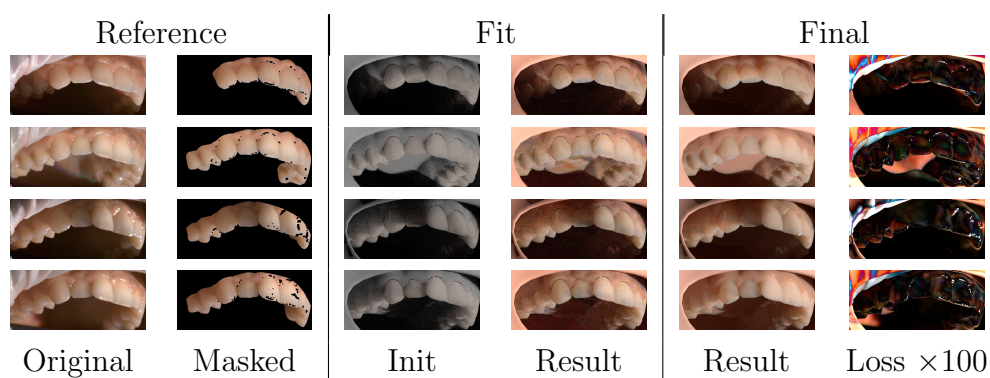


Figure 5.23: Results for the entire tooth row of the second volunteer from the viewpoint of the first camera

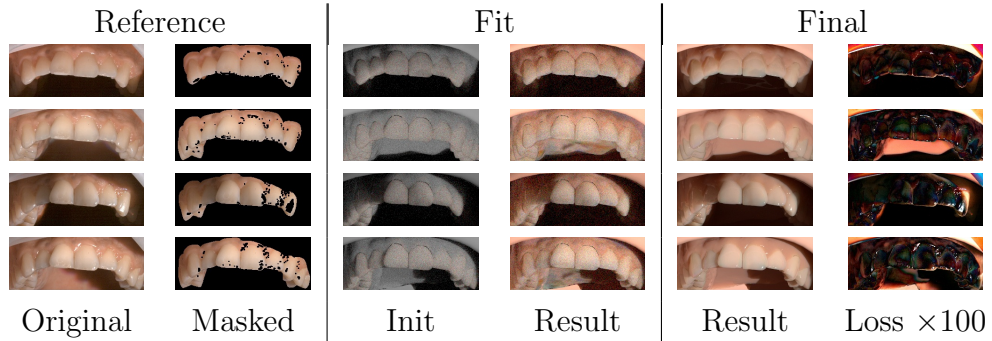


Figure 5.24: Results for the entire tooth row of the second volunteer from the viewpoint of the second camera

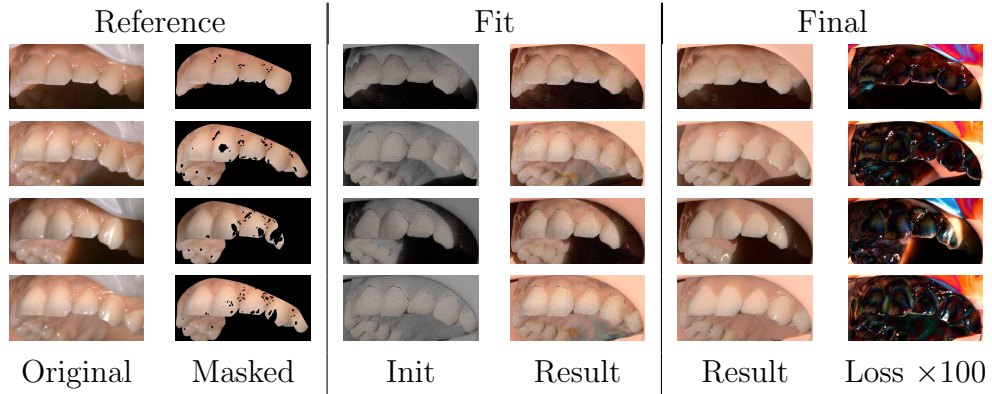


Figure 5.25: Results for the entire tooth row of the second volunteer from the viewpoint of the third camera

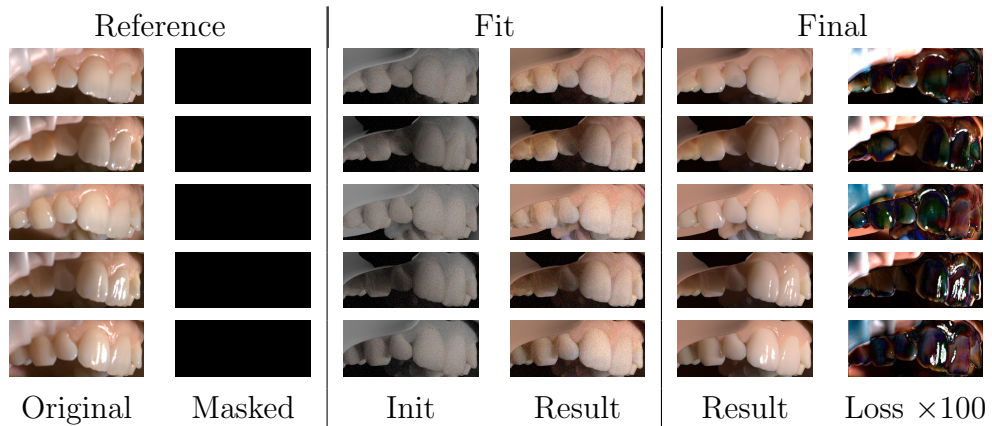


Figure 5.26: Results for the entire tooth row of the third volunteer from the viewpoint of the first camera

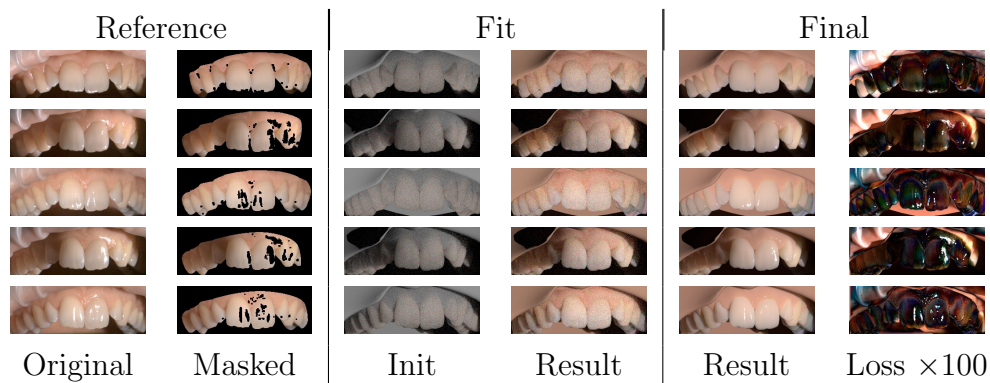


Figure 5.27: Results for the entire tooth row of the third volunteer from the viewpoint of the second camera

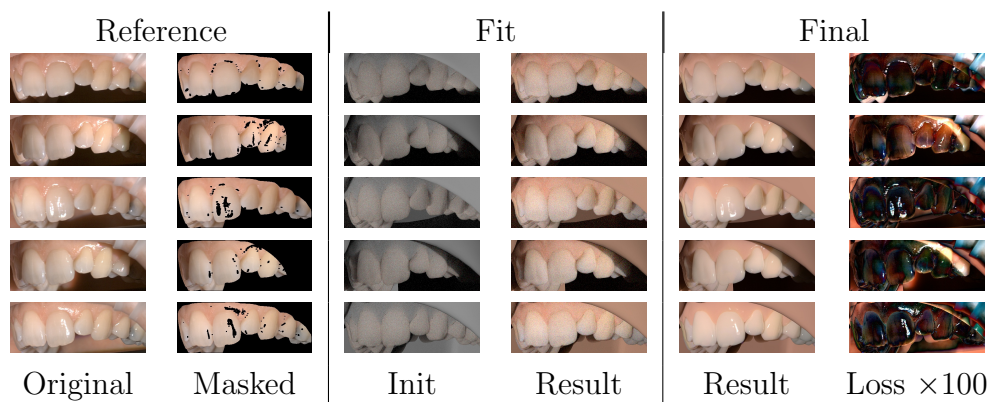


Figure 5.28: Results for the entire tooth row of the third volunteer from the viewpoint of the third camera


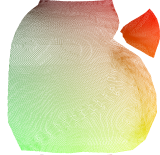
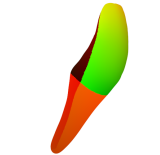

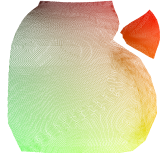
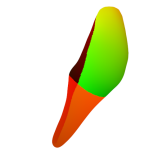
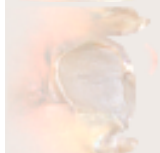
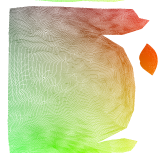

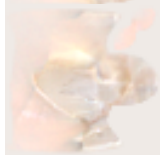
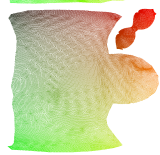
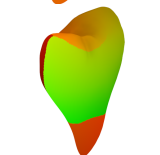

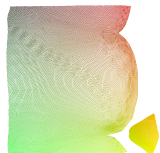





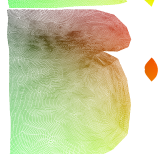
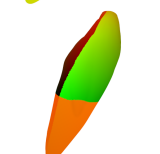

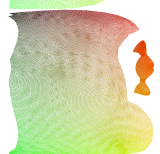




Name	Texture	Coordinates	Meshes
Right Central Incisor			
Right Lateral Incisor			
Right Canine			
Right First Premolar			
Left Central Incisor			
Left Lateral Incisor			
Left Canine			
Left First Premolar			
Gingiva (Gums)			

Figure 5.29: Textures estimated for our first volunteer

Name	Texture	Coordinates	Meshes
Right Central Incisor			
Right Lateral Incisor			
Right Canine			
Right First Premolar			
Left Central Incisor			
Left Lateral Incisor			
Left Canine			
Left First Premolar			
Gingiva (Gums)			

Figure 5.30: Textures estimated for our second volunteer


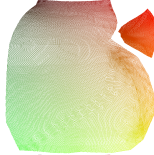


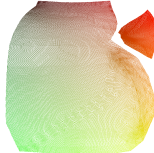


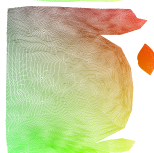


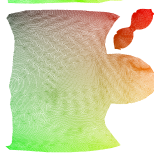
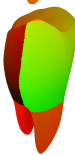

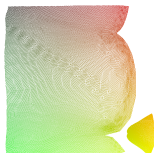





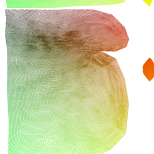


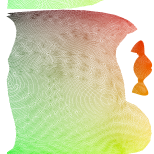


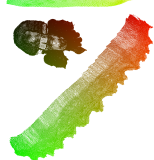
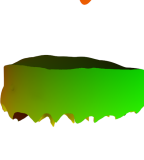
Name	Texture	Coordinates	Meshes
Right Central Incisor			
Right Lateral Incisor			
Right Canine			
Right First Premolar			
Left Central Incisor			
Left Lateral Incisor			
Left Canine			
Left First Premolar			
Gingiva (Gums)			

Figure 5.31: Textures estimated for our third volunteer

CHAPTER 6

Closure

The multitude of physical phenomena observed on everyday objects prevents the development of one truly universal approach that is both efficient and able to capture every possible reflectance behavior and thus giving rise to specialized techniques. They can be more suitable to solve the problem at hand, given the prior knowledge about the structure. Within this work, we demonstrated multiple techniques that are specifically tailored to a given problem that violates simplifications that are often employed in graphics, such as static unchanging material properties, geometrical optics modeling of light propagation or assuming that objects are opaque and their internal structure contributes insignificantly to the final appearance. This work ends with a discussion of the main novel results and techniques that were derived in the main text and how they fit within the overall field of digital material appearance modeling and acquisition. To avoid repetition, some of the discussion that was already provided in the respective chapters regarding the individual problems is omitted and replaced by a more high-level overview.

6.1 Summary of contributions

The main contributions are split across three distinct topics that are part of the digital material theory. They are going to be discussed separately in the following paragraphs.

Pile fabrics modeling and acquisition. Chapter 3 introduced a pipeline for extracting the appearance of pile fabrics by employing a low-dimensional

parametric model. Even with relatively small number of parameters, the appearance still matches the average reflectance behavior and the distribution based approach gives connection with the underlying geometry structure of fabrics. The interesting consequence of including a distribution with tilt is that reflection is well defined underneath the horizon compared to regular normal map based approaches, but it might require multi-scattering evaluation to avoid energy loss. Computing the reflectance relationship is simple and it can be done in real time on any existing graphics processor, even if real-time ray tracing (RTRT) is employed as the main rendering framework, similarly to the one developed for this project. All parameters are easy to modify and intuitive, as it is shown by the demo developed to demonstrate the capabilities of this technique. The approach has potential to be employed for synthesis and acquisition of surface appearance on a tight rendering budget.

Worn surfaces illumination. The highly reflective and uniform nature of smooth surfaces guarantees that defects close in size to the wavelength can be observed and therefore violate the notion of scale, which is discussed in Chapter 4. Building upon the existing literature of modeling iridescent scratches, a set of algorithms for solving illumination by extended light sources and anti-aliasing was proposed. Extended light sources are reflected by scratches and their mirror-like highlight forms the concentric structures appearing on the surface. Therefore, it is essential to properly evaluate illumination when these models are used in real-time rendering scenario. It is not enough to break light sources into many punctual lights because the discontinuities are going to be immediately apparent. Furthermore, anti-aliasing keeps the scratches connected and visible on the screen, even though they are essentially subpixel detail. These problems are solved with novel algorithms and further evaluated with different optimized data structures to validate that these techniques truly have the potential to work with real-time performance. Bringing microscopic detail to computer graphics is becoming increasingly important with the rising demand for higher rendering resolution and greater level of detail in movies, video games and design applications.

Human teeth and oral cavity modeling and acquisition. Digital doubles are increasingly used in the visual effects industry and video games to act as stunt doubles, virtual representation of living or long passed-away people, or to enhance existing footage. In Chapter 5 a framework for appearance

acquisition of teeth is described. Teeth vary greatly between individuals based on their physiological characteristics, therefore a single model with fixed parameters cannot reproduce all possible variation. Teeth have to be captured while still alive which comes with its own set of challenges. Data is acquired by intra-oral scanning of the exact geometry and capturing photographs in a light dome with optical flow stabilization. Existing models for surface and volume rendering are employed to extract a compact representation that seamlessly integrates in existing frameworks. The main difficulty lies in estimating the parameters of these models which requires considering direct and indirect illumination, and internal scattering within the structure of teeth. The main novel contribution is extending stochastic inverse rendering techniques to account for the inner geometry (dentin, in the case of teeth). This framework has the potential to be applied to other living tissue and man-made objects to estimate inner geometric structures purely based on photographs of the final appearance.

6.2 Limitations, future work and outlook

As the scientific progress continue, previous limitations of research works get lifted, but others are left as future work. Here these points are discussed within the context of individual works and how they can influence each other to bring even more accurate solutions of modeling and acquisition of digital materials.

6.2.1 Fabrics

The technique of using semi-infinite medium lands itself nicely to the problem of modeling opaque surfaces. However, certain fabrics might be translucent or even let light through gaps in the weave pattern. The problem as stated can be solved through inverse volume rendering, similarly to the triangular mesh modeling used for modeling teeth or by considering the radiative transport through regular grids [Gkioulekas et al. 2016]. One difficulty that arises is finding derivatives with respect to interpolated anisotropic media parameters which in earlier works on forward rendering was achieved through a power iteration [Jakob et al. 2010; Zhao et al. 2011]. It might be numerically unstable or inefficient to differentiate. Another difficulty is that delta tracking (also known as Woodcock tracking) is discrete and it is thus not dif-

ferentiable. Therefore integration has to fallback to less efficient algorithms, such as ray marching or regular tracking. Another problem revealed by the study on efficient integration of distributions of normals [Yan et al. 2014; Yan et al. 2016] is that subpixel detail can still influence significantly the final image. Similar problems can be expected to appear on fabrics where multiple parallel fibers can reflect light in mirror-like fashion. Of course, fabrics made of straight fibers are rarely the case, but the problem still exists and it is often ignored by works that downsample media with a smooth distribution [Heitz et al. 2015]. Hypothetically, it is possible to exploit the existence of this scale violation to extract individual fibers and even do super resolution of fabrics, but it would require further research. Fabrics appearance up to this point was investigated only within the context of isotropic and anisotropic media, however the correlation between individual fiber positions in weave assemblies implies that more sophisticated models can be built to simulate radiative transfer. One early work on random discrete media was by Moon et al. [2007] who exploited precomputation to model these surfaces. A more general approach of modeling these surfaces was proposed recently that integrates more seamlessly in existing path tracing frameworks Jarabo et al. [2018].

Tactile interaction can be further improved to better represent long fibers spanning multiple pixel or including complex physics simulation. In fact, during the development of this technique physics simulation was incorporated, but a significant effect was never found to justify using it. Velvet seems to be the only exception where long fibers manage to restore their initial orientation. Acquiring this behavior is more difficult because parts of the scene are occluded and the restoring motion can be close to instant. High-speed camera imaging paired with long macro lenses can be potentially used to study the behavior more in depth.

Another problem is the applicability of wave-optical approaches to this kind of problem that is going to be discussed in the next section.

6.2.2 Worn surfaces and wave optics

Modeling scratches as line segments might be limiting because real surfaces can contain arbitrary vectorial shapes with varying thickness, depth or length. More accurate models in the future might incorporate even more variation options. Furthermore, the entire scalar diffraction theory that is built around the Rayleigh-Sommerfeld equation and its various simplifica-

tions relies on local flatness which is also the case for any work relying on ray differentials [Igehy 1999] or partial derivative computation to perform camera pixel footprint integration. These are also single-scattering BRDFs because the theory is developed around thin screens and it does not consider inter-reflections. Lifting this limitation would require doing more accurate modeling of light propagation, which would be prohibitively expensive with the currently existing frameworks [Musbach et al. 2013]. Even more, it might not be justified to apply them on the complete rendered scene. Only limited number of surfaces show these effects and the reflected patterns deteriorate when they strike a glossy surface.

Wave-optical phenomena can be observed on items with complex microscopic structure (butterfly wings) and diffraction gratings form on many man-made objects with repeating structure (cell phone monitors, iridescent foil, finely-weaved fabrics). There is great potential for deriving efficient techniques to model many of these everyday items in efficient manner. The current optimization structures are probably going to be the limiting factor because every scratch element is stored separately in the scratch iridescence model. Newer seamless texturing approaches have to be developed to use distributions of diffraction gratings instead. Another interesting problem is propagation within fibers and the effects of transmission on wave optics. Recent works proposed frameworks for modeling layered material [Jakob et al. 2014a; Zeltner and Jakob 2018; Belcour 2018], but none of them considered the effects of the wave nature of light.

The lack of efficient techniques for rendering soft shadows was recently lifted by the advent of ratio estimator techniques for rendering extended light sources [Heitz et al. 2018] and the availability of ray tracing hardware. More efficient sampling algorithms are required to incorporate these features with our iridescent scratches techniques which can be an avenue for future research. The use of footprint integration in Chapter 4 removed the requirement for stochastic integration in spatial domain, but the fact that scratches have a very narrow reflection band across one dimension and a quite wide across another means that the usual approach of alternating between two sampling strategies [Veach and Guibas 1997] is wasting significant number of samples.

Acquisition of worn surfaces is another topic that was not fully explored by our works on scratch iridescence. Both the earlier work [Werner et al. 2017] and the one described within this thesis (cf. Chapter 4) showed phenomenological match with samples prepared in laboratory conditions. How-

ever, a fully automated acquisition approach was not proposed. The main challenge lies in the strictness of the specular condition which would require dense sampling of the angular and spatial domain. Another problem that has to be overcome is that real-world surfaces can be additionally contaminated by dust and grease. Even further, they can be composite surfaces, requiring more complex modeling approaches. These problems are a good subject for future work.

The use of footprint integration is biased by the fact that we use approximations and the assumption of local flatness stemming from the definition of ray differentials [Igehy 1999].

Recently, new methods for efficiently integrating microscopic structures that show wave-optical effects were introduced [Yan et al. 2018]. However, it is unclear how to pair efficiently these techniques with area or environment light illumination. Also, the memory footprint of these substitute distribution based approaches is not much more favorable than our technique. Another paper considered integration over arbitrary diffraction apertures [Scandolo et al. 2018] by using axis-aligned rectangular decomposition. This approach is not very suitable for surfaces with arbitrarily oriented unstructured features. Our technique fits better the profile of this problem.

6.2.3 Inverse rendering of participating media

Derivative path tracing is limited to functions that are continuous in path space. Lifting this limitation can be artificially achieved by smoothing borders between media and using inefficient integration techniques, as already mentioned (cf. 5.8). In a parallel work, Li et al. [2018] outlined how to solve accurately the case of subpixel dependence on deformable opaque geometry with an additional edge sampling strategy. However, even they note that it is a much more difficult problem in the case of volumes that would further involve accounting for the effects of scattering and extinction. These problems need to be first solved before opaque discontinuous medium can be incorporated within these differentiable frameworks. There is quite a lot of future research that must be done to make these techniques universally applicable.

Teeth modeling can be significantly improved. Bending of light can be incorporated, however it makes the problem much more difficult to differentiate if considering deformable boundaries between media. Additionally, saliva can act as thin film which requires considering modeling iridescence [Belcour

and Barla 2017]. The problem becomes even more difficult if dirty teeth are considered since they essentially become multi-layered material. Furthermore, they can have small scratches appearing on them that would require considering filtering subpixel detail. The appearance can further change over time and fluid simulation of saliva must be considered. Other elements of the oral cavity might change their optical properties based on how much blood flows through their veins. There is quite a lot of potential for future improvements of the visual fidelity.

More limitations were already outlined in the discussion about fabrics (cf. Sec. 6.2.1). One significant limitation of local and momentum based optimization methods is their sensitivity to the initial values which often leads them to local minimum in null space. In the case of teeth this problem is avoided by using a lot of prior knowledge, but in the general case that might not be the case.

6.2.4 Outlook

Overall the techniques described in this dissertation show great potential for future research. Wave optics are gaining popularity in the graphics community. Multiple works can be expected in the following years that exploit different aspects of them to bring greater level of detail or fabricate objects with specific reflectance properties. There is also growing interest in developing differentiable graphics pipelines, especially with the recent success of machine learning techniques. The potential of incorporating models and techniques from graphics can bring further improvements to computer vision algorithms. Finally, some of the works included in this thesis have already inspired other researchers in the field of computer graphics who cited them in their works at major venues which is a real testament of their impact on the community.

APPENDIX A

Publikationen

- Zdravko Velinov, Marios Papas, Derek Bradley, Paulo Gotardo, Parsa Mirdehghan, Steve Marschner, Jan Novák, and Thabo Beeler [Dec. 2018a]. “Appearance Capture and Modeling of Human Teeth.” In: *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 37.6, 207:1–207:13. ISSN: 0730-0301. DOI: 10.1145/3272127.3275098. URL: <http://doi.acm.org/10.1145/3272127.3275098>
- Zdravko Velinov, Sebastian Werner, and Matthias B. Hullin [May 2018b]. “Real-Time Rendering of Wave-Optical Effects on Scratched Surfaces.” In: *Computer Graphics Forum* 37.2. ISSN: 1467-8659. DOI: 10.1111/cgf.13347
- Sebastian Werner, Zdravko Velinov, Wenzel Jakob, and Matthias B. Hullin [Nov. 2017]. “Scratch Iridescence: Wave-optical Rendering of Diffractive Surface Structure.” In: *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 36.6, 207:1–207:14. ISSN: 0730-0301. DOI: 10.1145/3130800.3130840. URL: <http://doi.acm.org/10.1145/3130800.3130840>
- Zdravko Velinov and Matthias B. Hullin [Oct. 2016]. “An Interactive Appearance Model for Microscopic Fiber Surfaces.” In: *Vision, Modeling, and Visualization 2016*

List of Symbols

\mathcal{G}	Gaussian filter function
\mathcal{T}_s	Transfer function
$\mathcal{T}_{\text{base}}$	Base material transfer function
$\mathcal{T}_{\text{mask}}$	Masked-out transfer function
$\mathcal{T}_{\text{scratch}}$	Scratch element transfer function
\mathcal{F}	Fourier transform
F	Fresnel term
\mathbf{t}	Surface tangent vector
\mathbf{b}	Surface bitangent vector (orthogonal to \mathbf{t})
\mathbf{n}	Surface normal vector
ω	Light direction vector
ω_i	Incident light vector
ω_o	Outgoing light vector
ϕ_i	Azimuth incident light angle
θ_i	Zenith incident light angle
θ_o	Zenith outgoing light angle
\mathcal{S}^2	Unit sphere domain
Ω^+	Positive hemisphere
η	Refractive index
ρ	Volume density
b	Hyper-parameter
\hat{m}	First moment estimate of Adam
\hat{v}	Second moment estimate of Adam
α_l	Learning rate
p	Microflake phase function
U	Electric field
$\mathbf{x}_{\mathcal{P}}$	Position in screen coordinates

$\mathbf{x}_{\mathcal{P}}^{(i,j)}$	Position in screen coordinates of a pixel with coordinates (i,j)
\mathcal{P}	Projection transformation
\mathbf{X}	Matrix of positions in space
$\mathbf{P}_{\text{light}}^{(i)}$	Computed spatial position of a light source
$\mathbf{P}_{\text{meas}}^{(i)}$	Measured spatial position of a light source
$\hat{\mathbf{P}}_{\text{light}}$	Mean of computed spatial positions of light sources
$\hat{\mathbf{P}}_{\text{meas}}$	Mean of measured spatial positions of light sources
N_{light}	Number of light sources
\mathbf{F}	Focal length vector
\mathbf{C}	Vector of camera optical centers in pixel coordinates
\mathbf{K}	Vector of camera optical centers in pixel coordinates
k_s	Specular albedo
k_d	Diffuse albedo
f_{BTF}	Bidirectional Texture Function (BTF)
D	Distribution of normals
D_{BTF}	Distribution of normals approximately computed from a BTF
L_i	Radiance transported along the incident direction
S_v	Vector of projected areas
S_σ	Diagonal matrix of projected areas
S_{GGX}	Covariance matrix of the Spherical GGX distribution (SGGX)
$W_{\omega_i, \omega_o}^{(k)}$	index of a discrete radiance sample encoded in a BTF
$g^{(k)}$	Residual
ω_0	First axis of SGGX distribution
ω_1	Second axis of SGGX distribution
ω_2	Third axis of SGGX distribution
ω_k	k^{th} Axis of SGGX distribution
M	Basis matrix of SGGX distribution
$\mathbf{s}^{(m)}$	Position of a scratch m in world coordinates

LIST OF SYMBOLS

\mathbf{x}	Shaded position in world coordinates
ξ	Vector of direction cosines
ξ_1	Tangent direction cosine
ξ_2	Bitangent direction cosine
$\xi_1^{(m)}$	Tangent direction cosine to the m^{th} scratch
$\xi_2^{(m)}$	Bitangent direction cosine to the m^{th} scratch
$\omega_{o,1}^\perp$	Tangential component of outgoing light direction projected on the unit disk
$\omega_{o,2}^\perp$	Bitangential component of outgoing light direction projected on the unit disk
$\xi^{(m)}$	Vector of direction cosines w.r.t. the basis vector of scratch m
σ	Coherence area radius (60 μm)
A_C	Coherence area ($A_C = \pi\sigma^2$)
λ	Wavelength
k	Wave number
\mathcal{B}	Base response function
\mathcal{S}	Scratch response function
$\mathcal{S}^{(m)}$	Response function of the m^{th} scratch
$\mathcal{W}^{(m)}$	Response function w.r.t. width of the m^{th} scratch
$\mathcal{D}^{(m)}$	Response function w.r.t. depth of the m^{th} scratch
$W_s^{(m)}$	Scratch profile width parameter
$D_s^{(m)}$	Scratch profile depth parameter
$\mathbf{s}^{(m)}$	Position of scratch origin
$\mathbf{s}'^{(m)}$	Relative position to the scratch origin
$\mathbf{s}''^{(m)}(s)$	Scratch line equation
s	Point along scratch line
ζ	Boundary of the integral over the projected light source area
s_1	Integration lower boundary along scratch line
s_2	Integration upper boundary along scratch line
$A_{\mathcal{P}}$	Area of the projected camera pixel footprint
$s_t^{(m)}$	Tangential component of the relative position to the scratch origin

$s_b^{(m)}$	Bitangential component of the relative position to the scratch origin
$\eta^{(m)}$	Spatial phases of the m^{th} scratch
$\eta_{\mathcal{P}}^{(m)}$	Spatial phases integrated over a pixel footprint of the m^{th} scratch
d	Exponential component of the spatial phases of the m^{th} scratch
ρ_c	Scratch area density
$\eta_a^{(m)}$	Angular component of spatial phases of the m^{th} scratch
$\eta_s^{(m)}$	Spatial component of spatial phases of the m^{th} scratch
$\mathbf{t}^{(m)}$	Tangent of m^{th} scratch
$\mathbf{b}^{(m)}$	Bitangent of m^{th} scratch
ω_o^\perp	Outgoing light direction projected on the unit disk
a	Extent of the projected footprint ellipse (a_1, a_2)
$l_{\text{contained}}^{(m)}$	Scratch length contained within pixel footprint
l_1	Scratch lower boundary
l_2	Scratch upper boundary
$\rho_{\mathcal{P}}$	Scratch area density integrated over a pixel profile
$\alpha_{\mathcal{P}}$	Footprint corrective term blend factor
$\mathbf{V}^{(m)}$	Vertex m of polygonal light source
$q^{(m)}$	Intersection with projected light source
\mathbf{J}	Jacobian matrix
\mathbf{G}_E	Parametric equation over an ellipse projected on the hemisphere
\mathbf{G}_S	Parametric equation over a circle projected on the hemisphere
f	Bidirectional Scattering Distribution Function (BSDF)
f_r	Bidirectional Reflectance Distribution Function (BRDF)

LIST OF SYMBOLS

f_t	Bidirectional Transmission Distribution Function (BTDF)
ω_h	Half-way vector between outgoing and incident light vector (micro normal)
χ	Parameter block
α	Volume albedo
κ	Volume albedo reparameterized in dual domain
α_e	Enamel volume albedo
α_d	Dentin volume albedo
α_{mix}	Smoothed albedo between enamel and dentin
σ_t	Volume extinction coefficient
σ	Projected area of ellipsoid along direction ω
σ_a	Volume absorption
σ_{td}	Dentin volume extinction coefficient
σ_{te}	Enamel volume extinction coefficient
σ_s	Volume scattering coefficient
f_p	Phase function
f_{HG}	Phase function according to the Henyey-Greenstein model (HG)
β	Blend factor of the Dual-phase Henyey-Greenstein model (DHG)
β_e	Enamel blend factor of the Dual-phase Henyey-Greenstein model (DHG)
β_d	Dentin blend factor of the Dual-phase Henyey-Greenstein model (DHG)
β_{mix}	Smoothed blend factor between enamel and dentin of the Dual-phase Henyey-Greenstein model (DHG)
g	Mean cosine parameter of the Henyey-Greenstein model (HG)
$g_{1,e}$	Mean cosine parameter of the first phase of enamel Dual-phase Henyey-Greenstein model (DHG)
$g_{2,e}$	Mean cosine parameter of the second phase of enamel Dual-phase Henyey-Greenstein model (DHG)

$g_{1,d}$	Mean cosine parameter of the first phase of dentin Dual-phase Henyey-Greenstein model (DHG)
$g_{2,d}$	Mean cosine parameter of the second phase of dentin Dual-phase Henyey-Greenstein model (DHG)
\mathbf{g}_e	Mean cosine parameter vector of enamel Dual-phase Henyey-Greenstein model (DHG)
\mathbf{g}_d	Mean cosine parameter vector of dentin Dual-phase Henyey-Greenstein model (DHG)
g_{mix}	Smoothed mean cosine parameter between enamel and dentin
a_t	Absorption parameter of the tooth stain layer
M_p	Number of pixels in optimized image
P_p	Number of optimized parameters
N_s	Number of samples per pixel in the optimized image
K_b	Number of bounces considered in the optimized image
\mathbf{v}	Triangle vertex
q	Distance to surface
q_p	Distance to the the closest point on a surface
\mathbf{q}_p	Vector from the the closest point on a surface to a spatial position
\mathbf{u}	Edge vector
l	Logistic sigmoid blend weight
s_l	Logistic sigmoid spread parameter
F	Image reconstruction filtering function
\mathbf{G}_T	Tensor representation of triangle mesh
\mathbf{V}	Matrix of triangle vertices
\mathbf{B}	Dentin blend mode matrix of vertices
\mathbf{c}	Dentin blend shape factors
\mathbf{B}_0	Enamel matrix of vertices
\mathbf{D}	Determinant matrix
\mathbf{r}_o	Ray origin
\mathbf{r}_d	Ray direction
t	Intersection distance along a ray

LIST OF SYMBOLS

\mathbf{p}_r	Intersection position on a surface mesh in world space
\mathbf{p}_s	Position in space
\mathbf{h}	Vector element of a determinant in the ray intersection equation
\mathbf{S}	Vector of volume densities along a path
\mathbf{S}_e	Vector of volume densities of the exited volumes
\mathbf{S}_i	Vector of volume densities of the entered volumes
\mathbb{R}	Real number
L	Radiance
E	Irradiance
\mathcal{T}_t	Transmittance along a segment
ρ	Loss function
ψ	Cost function or residual between measured and optimized image
ψ_χ	Partial cost function w.r.t. an individual parameter
$\psi_{\mathbf{G}}$	Partial cost function w.r.t. triangle mesh geometry
$\psi_{\mathbf{B}}$	Partial cost function w.r.t. the blend shape geometry
ψ_{σ_t}	Partial cost function w.r.t. extinction
$\psi_{\sigma_t,s}$	Partial cost function w.r.t. extinction - scattering term
$\psi_{\sigma_t,t}$	Partial cost function w.r.t. extinction - extinction term
ψ_d	Partial cost function w.r.t. ray intersection distance
ψ_α	Partial cost function w.r.t. albedo
ψ_κ	Partial cost function w.r.t. albedo parameterized in dual domain
$\psi_{\alpha_{\text{mix}}}$	Partial cost function w.r.t. albedo smoothed between enamel and dentin
ψ_{f_p}	Partial cost function w.r.t. phase function
ψ_f	Partial cost function w.r.t. Bidirectional Scattering Distribution Function (BSDF)

ψ_t Partial cost function w.r.t. transmittance

List of Figures

2.1	Spatially Varying Bidirectional Reflectance Distribution Function (SV-BRDF) parameterization on the hemisphere	12
2.2	BRDF modified to account for the angular extent of the light source and the spatial extent of the projected camera pixel footprint	14
2.3	Participating media composed of microscopic flakes with isotropic and anisotropic distribution	15
3.1	The surface of Alcantara	28
3.2	Our microscopic fibers SVBRDF fitting pipeline	29
3.3	Microscopic fiber surface model	33
3.4	PSNR in the angular domain for each pixel of the SVBRDF produced for Alcantara material sample	39
3.5	Orientation of tangent vectors of the SGGX basis after post-processing	40
3.6	Ellipsoid distributions fit to the original BTF representation	41
3.7	Model tactile interaction and editing in real time	42
3.8	Top-down view fit of the BTF reflectance	43
3.9	Ground truth BTF rendering	44
4.1	Teaser: Illumination of worn surfaces that show wave-optical effects	48
4.2	Scratch iridescence model rendered under point light source (1spp)	50
4.3	Comparison of diffracted radiance for a single scratch with rectangular profile	56
4.4	Footprint anti-aliasing evaluation	57

4.5	Comparison of numeric vs analytic results of the integration of scratch area density and spatial phases terms	58
4.6	Approximation of the integral over the solid subtended by triangular light source in direction cosine domain	63
4.7	Comparison between the approximation of S_i and the exact function	64
4.8	Comparison of the specular lobe of the ground truth Monte Carlo sampling, our spherical light approximation with exact $S_i(x)$ and fitted curve	65
4.9	Comparison of the specular lobe of the ground truth Monte Carlo sampling, our triangular light approximation with exact $S_i(x)$ and fitted curve	68
4.10	Comparison between Monte Carlo integration over the solid angle subtended by a spherical and polygonal light source against our approximations	71
4.11	Recreation of the dining table scene from the original work on iridescent scratches	72
4.12	Images in RGB of increasing number of environment elements approximated by rectangles	79
4.13	Images with 16 spectral bands of increasing number of environment elements approximated by rectangles	80
4.14	Sphere light source with anti-aliasing enabled at different zoom levels	81
4.15	Sphere light source with anti-aliasing disabled at different zoom levels	81
4.16	Triangle light source with anti-aliasing enabled at different zoom levels	81
4.17	Triangle light source with anti-aliasing disabled at different zoom levels	81
4.18	Comparison between ground truth photo of a disc and rendering with our sphere light approximation	83
5.1	Teaser: Comparisons of our recovered teeth appearance parameters rendered next to corresponding reference images . . .	86
5.2	Dental scanning example	88
5.3	Tooth anatomy	90
5.4	Teeth data acquisition setup	96
5.5	Dentin deformation modes	99

LIST OF FIGURES

5.6	Tooth model parameters	102
5.7	Derivative of extinction with respect to blend shape deformation	107
5.8	Dead tooth validation results	112
5.9	Tooth model validation	113
5.10	Single tooth fitting and validation	114
5.11	Upper jaw fitting of first volunteer	115
5.12	Upper jaw fitting of second volunteer	115
5.13	Upper jaw fitting of third volunteer	115
5.14	Loss function during the optimization process of the first volunteer's teeth	120
5.15	Loss function during the optimization process of the second volunteer's teeth	120
5.16	Loss function during the optimization process of the third volunteer's teeth	120
5.17	Results for the right central incisor of the first volunteer from the viewpoint of the first camera	125
5.18	Results for the right central incisor of the first volunteer from the viewpoint of the second camera	126
5.19	Results for the right central incisor of the first volunteer from the viewpoint of the third camera	127
5.20	Results for the entire tooth row of the first volunteer from the viewpoint of the first camera	128
5.21	Results for the entire tooth row of the first volunteer from the viewpoint of the second camera	128
5.22	Results for the entire tooth row of the first volunteer from the viewpoint of the third camera	129
5.23	Results for the entire tooth row of the second volunteer from the viewpoint of the first camera	129
5.24	Results for the entire tooth row of the second volunteer from the viewpoint of the second camera	130
5.25	Results for the entire tooth row of the second volunteer from the viewpoint of the third camera	130
5.26	Results for the entire tooth row of the third volunteer from the viewpoint of the first camera	130
5.27	Results for the entire tooth row of the third volunteer from the viewpoint of the second camera	131
5.28	Results for the entire tooth row of the third volunteer from the viewpoint of the third camera	131

5.29 Textures estimated for our first volunteer 132
5.30 Textures estimated for our second volunteer 133
5.31 Textures estimated for our third volunteer 134

Index

- appearance acquisition
 - devices, 18
 - gonioreflectometer, 18
 - hemispherical gantry, 19
 - light dome, 19
- blend-shape, 98
- BRDF
 - Chandrasekhar's BRDF, 35
 - incoherent superposition, 56
 - scratch iridescence model, 54
 - Seeliger's Law, 35
 - SVBRDF, 54
- BSDF, 11
 - sampling, 13
 - thin dielectric, 97
- BTF, 36
- calibration
 - camera resectioning, 20
 - fiducial markers, 20
 - geometrical, 20
 - light source registration, 20
 - photometric, 22
 - spectralon, 22, 95
- Dirac delta function, 16, 97
- direction
 - buccal, 99
 - occlusal, 99
 - proximal, 99
- energy conservation, 13, 14
- fabrics
 - Alcantara, 29
 - pile fabrics, 28
 - velvet, 28
- generalized product rule, 101
- geometrical optics, 9
 - Beer-Lambert law, 11
 - camera obscura, 19
 - LTE, 11, 12
 - microflake models, 15
 - SGGX, 16, 34
 - NDF, 15
 - phase function, 14
 - Henyey-Greenstein, 14
 - pinhole camera, 19
 - rendering equation, 11, 61
 - RTE, 10, 34, 98
 - scattering kernel, 14
 - volume density, 15
 - volume scattering, 14
- gingiva, 99
- HDR, 22
- inverse

- rendering, 22
- methods, 22

- Neumann series, 11
- numerical optimization
 - Adam, 23
 - cost function, 23
 - hyperparameter, 24
 - Levenberg-Marquardt, 23
 - loss function, 23
 - non-linear optimization, 100
 - regularization, 23
 - residual, 23

- oral cavity, 99

- PCA, 36
- Perspective camera model, 19

- ray differentials, 13
 - projected camera pixel footprint, 14

- reciprocity, 12
- reprojection error, 20

- tooth anatomy, 90
 - dentin, 90
 - enamel, 90

- wave optics, 16
 - diffracted radiance, 17
 - diffraction, 16
 - electric field, 16
 - Rayleigh-Sommerfeld formula, 16
 - wavefront, 16

Bibliography

- Ahn, Jae Sung, Anjin Park, Ju Wan Kim, Byeong Ha Lee, and Joo Beom Eom (2017). “Development of Three-Dimensional Dental Scanning Apparatus Using Structured Illumination.” In: *Sensors*.
- Arvo, James Richard (1995). “Analytic Methods for Simulated Light Transport.” PhD thesis. New Haven, CT, USA: Yale University.
- Ashikmin, Michael, Simon Premože, and Peter Shirley (2000). “A Microfacet-based BRDF Generator.” In: *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH '00. New York, NY, USA: ACM Press / Addison-Wesley Publishing Co., pp. 65–74. ISBN: 1-58113-208-5. DOI: 10.1145/344779.344814. URL: <http://dx.doi.org/10.1145/344779.344814>.
- Bao, Hujun and Qunsheng Peng (1993). “Shading models for linear and area light sources.” In: *Computers & Graphics* 17.2, pp. 137–145. ISSN: 0097-8493. DOI: [http://dx.doi.org/10.1016/0097-8493\(93\)90097-S](http://dx.doi.org/10.1016/0097-8493(93)90097-S). URL: <http://www.sciencedirect.com/science/article/pii/009784939390097S>.
- Bartell, F. O., E. L. Dereniak, and W. L. Wolfe (Mar. 1981). “The theory and measurement of bidirectional reflectance distribution function BRDF and bidirectional transmittance distribution function BTDF.” In: *Radiation scattering in optical systems*. Ed. by G. H. Hunt. Vol. 257, pp. 154–160. DOI: 10.1117/12.959611.
- Batcher, K. E. (1968). “Sorting Networks and Their Applications.” In: *Proc. April 30–May 2, 1968, Spring Joint Computer Conference*. AFIPS '68 (Spring). Atlantic City, New Jersey: ACM, pp. 307–314. DOI: 10.1145/1468075.1468121. URL: <http://doi.acm.org/10.1145/1468075.1468121>.
- Baum, D. R., H. E. Rushmeier, and J. M. Winget (1989). “Improving Radiosity Solutions Through the Use of Analytically Determined Form-factors.”

- In: *Proc. 16th Ann. Conf. on Computer Graphics and Interactive Techniques*. SIGGRAPH '89. ACM, pp. 325–334. ISBN: 0-89791-312-4. DOI: 10.1145/74333.74367. URL: <http://doi.acm.org/10.1145/74333.74367>.
- Beeler, Thabo (2012). “Passive Spatio-Temporal Geometry Reconstruction of Human Faces at Very High Fidelity.” PhD thesis. ETH Zurich / Disney Research Zurich.
- Belcour, Laurent (2018). “Efficient Rendering of Layered Materials using an Atomic Decomposition with Statistical Operators.” In: *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 37.4, p. 1. DOI: 10.1145/3197517.3201289. URL: <https://hal.archives-ouvertes.fr/hal-01785457>.
- Belcour, Laurent and Pascal Barla (2017). “A Practical Extension to Microfacet Theory for the Modeling of Varying Iridescence.” In: *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 36.4, p. 65. DOI: 10.1145/3072959.3073620. URL: <https://hal.archives-ouvertes.fr/hal-01518344>.
- Bergstra, James, Rémi Bardenet, Yoshua Bengio, and Balázs Kégl (2011). “Algorithms for Hyper-parameter Optimization.” In: *Proceedings of the 24th International Conference on Neural Information Processing Systems*. NIPS'11. Granada, Spain: Curran Associates Inc., pp. 2546–2554. ISBN: 978-1-61839-599-3. URL: <http://dl.acm.org/citation.cfm?id=2986459.2986743>.
- Bosch, Carles, Xavier Pueyo, Stéphane Mérillou, and Djamchid Ghazanfarpour (2004). “A Physically-Based Model for Rendering Realistic Scratches.” In: *Computer Graphics Forum* 23.3, pp. 361–370. ISSN: 1467-8659. DOI: 10.1111/j.1467-8659.2004.00767.x. URL: <http://dx.doi.org/10.1111/j.1467-8659.2004.00767.x>.
- Brox, T., A. Bruhn, N. Papenbergh, and J. Weickert (2004). “High accuracy optical flow estimation based on a theory for warping.” In: *ECCV*, pp. 25–36.
- Bruneton, Eric and Fabrice Neyret (2012). “A Survey of Nonlinear Prefiltering Methods for Efficient and Accurate Surface Shading.” In: *IEEE Transactions on Visualization and Computer Graphics* 18.2, pp. 242–260. ISSN: 1077-2626. DOI: 10.1109/TVCG.2011.81. URL: <http://dx.doi.org/10.1109/TVCG.2011.81>.

BIBLIOGRAPHY

- Buchillard, Stéphanie I., S. H. Ong, Yohan Payan, and Kelvin Foong (2007). “3D Statistical Models for Tooth Surface Reconstruction.” In: *Comput. Biol. Med.* 37.10, pp. 1461–1471.
- Burns, Christopher A. and Warren A. Hunt (2013). “The Visibility Buffer: A Cache-Friendly Approach to Deferred Shading.” In: *J. Computer Graphics Techniques (JCGT)* 2.2, pp. 55–69. ISSN: 2331-7418. URL: <http://jcggt.org/published/0002/02/04/>.
- Carpentier, Giliam de and Kohei Ishiyama (2017). “Decima: Advances in Lighting and AA.” In: *Advances in Realtime Rendering, Part II (SIGGRAPH 2017 Course)*. Ed. by Natalya Tatarchuk et al. URL: <http://advances.realtimerendering.com/s2017/DecimaSiggraph2017.pdf>.
- Carter, C.N., R.J. Pusateri, Dongqing Chen, A.H. Ahmed, and A.A. Farag (2010). “Shape from shading for hybrid surfaces as applied to tooth reconstruction.” In: *IEEE ICIP*, pp. 4049–4052.
- Cecchini, R. and G. Pelosi (1990). “Diffraction: the first recorded observation.” In: *IEEE Antennas and Propagation Magazine* 32.2, pp. 27–30. ISSN: 1045-9243. DOI: 10.1109/74.80496.
- Chandrasekhar, S. (1960). *Radiative Transfer*. Dover books on advanced mathematics. Dover Publications.
- Chen, Yanyun, Lin Xia, Tien-Tsin Wong, Xin Tong, Hujun Bao, Baining Guo, and Heung-Yeung Shum (July 2005). “Visual Simulation of Weathering by γ -ton Tracing.” In: *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 24.3, pp. 1127–1133. ISSN: 0730-0301. DOI: 10.1145/1073204.1073321. URL: <http://doi.acm.org/10.1145/1073204.1073321>.
- d’Eon, Eugene, Guillaume Francois, Martin Hill, Joe Letteri, and Jean-Marie Aubry (2011). “An Energy-Conserving Hair Reflectance Model.” In: *Computer Graphics Forum* 30.4. ISSN: 1467-8659.
- d’Eon, Eugene and Geoffrey Irving (July 2011). “A Quantized-diffusion Model for Rendering Translucent Materials.” In: *ACM Transactions on Graphics* 30.4, 56:1–56:14. ISSN: 0730-0301. DOI: 10.1145/2010324.1964951. URL: <http://doi.acm.org/10.1145/2010324.1964951>.
- d’Eon, Eugene, Steve Marschner, and Johannes Hanika (2013). “Importance Sampling for Physically-based Hair Fiber Models.” In: *SIGGRAPH Asia 2013 Technical Briefs*. SA ’13. Hong Kong, Hong Kong: ACM, 25:1–25:4. ISBN: 978-1-4503-2629-2. DOI: 10.1145/2542355.2542386. URL: <http://doi.acm.org/10.1145/2542355.2542386>.

- Dachsbacher, Carsten, Jaroslav Krivánek, Miloš Hašan, Adam Arbree, Bruce Walter, and Jan Novák (2014). “Scalable Realistic Rendering with Many-Light Methods.” In: *Computer Graphics Forum* 33.1, pp. 88–104. ISSN: 1467-8659.
- Dana, Kristin J., Bram van Ginneken, Shree K. Nayar, and Jan J. Koenderink (Jan. 1999). “Reflectance and Texture of Real-world Surfaces.” In: *ACM Transactions on Graphics* 18.1, pp. 1–34. ISSN: 0730-0301. DOI: 10.1145/300776.300778. URL: <http://doi.acm.org/10.1145/300776.300778>.
- Debevec, Paul E. and Jitendra Malik (1997). “Recovering High Dynamic Range Radiance Maps from Photographs.” In: *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques. SIGGRAPH '97*. New York, NY, USA: ACM Press / Addison-Wesley Publishing Co., pp. 369–378. ISBN: 0-89791-896-7. DOI: 10.1145/258734.258884. URL: <https://doi.org/10.1145/258734.258884>.
- Debevec, Paul, Tim Hawkins, Chris Tchou, Haarm-Pieter Duiker, Westley Sarokin, and Mark Sagar (2000). “Acquiring the Reflectance Field of a Human Face.” In: *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques. SIGGRAPH '00*. New York, NY, USA: ACM Press / Addison-Wesley Publishing Co., pp. 145–156. ISBN: 1-58113-208-5. DOI: 10.1145/344779.344855. URL: <http://dx.doi.org/10.1145/344779.344855>.
- Dereniak, Eustace L., Langford G. Brod, and John E. Hubbs (1982). “Bidirectional transmittance distribution function measurements on ZnSe.” In: *Appl. Opt.* 21.24, pp. 4421–4425. DOI: 10.1364/AO.21.004421. URL: <http://ao.osa.org/abstract.cfm?URI=ao-21-24-4421>.
- Dhillon, D.S., J. Teyssier, M. Single, I. Gaponenko, M.C. Milinkovitch, and M. Zwicker (2014). “Interactive Diffraction from Biological Nanostructures.” In: *Computer Graphics Forum* 33.8, pp. 177–188. ISSN: 0167-7055. DOI: 10.1111/cgf.12425. URL: <http://dx.doi.org/10.1111/cgf.12425>.
- Divitt, Shawn and Lukas Novotny (Feb. 2015). “Spatial coherence of sunlight and its implications for light management in photovoltaics.” In: *Optica* 2.2, pp. 95–103. DOI: 10.1364/OPTICA.2.000095. URL: <http://www.osapublishing.org/optica/abstract.cfm?URI=optica-2-2-95>.
- Dong, Zhao, Bruce Walter, Steve Marschner, and Donald P. Greenberg (2015). “Predicting Appearance from Measured Microgeometry of Metal Surfaces.” In: *ACM Transactions on Graphics* 35.1, 9:1–9:13. ISSN: 0730-

BIBLIOGRAPHY

0301. DOI: 10.1145/2815618. URL: <http://doi.acm.org/10.1145/2815618>.
- Donner, Craig and Henrik Wann Jensen (July 2005). “Light Diffusion in Multi-layered Translucent Materials.” In: *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 24.3, pp. 1032–1039. ISSN: 0730-0301. DOI: 10.1145/1073204.1073308. URL: <http://doi.acm.org/10.1145/1073204.1073308>.
- Donner, Craig and Henrik Wann Jensen (2007). “Rendering Translucent Materials Using Photon Diffusion.” In: *Proceedings of the 18th Eurographics Conference on Rendering Techniques*. EGSR’07. Grenoble, France: Eurographics Association, pp. 243–251. ISBN: 978-3-905673-52-4. DOI: 10.2312/EGWR/EGSR07/243-251. URL: <http://dx.doi.org/10.2312/EGWR/EGSR07/243-251>.
- Donner, Craig, Tim Weyrich, Eugene d’Eon, Ravi Ramamoorthi, and Szymon Rusinkiewicz (Dec. 2008). “A Layered, Heterogeneous Reflectance Model for Acquiring and Rendering Human Skin.” In: *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 27.5, 140:1–140:12. ISSN: 0730-0301. DOI: 10.1145/1409060.1409093. URL: <http://doi.acm.org/10.1145/1409060.1409093>.
- Drobot, Michal (2014). “Physically based area lights.” In: *GPU Pro 5*. Ed. by Wolfgang Engel. CRC Press, pp. 67–100.
- Duff, Tom, James Burgess, Per Christensen, Christophe Hery, Andrew Kensler, Max Liani, and Ryusuke Villemin (2017). “Building an Orthonormal Basis, Revisited.” In: *J. Computer Graphics Techniques (JCGT)* 6.1, pp. 1–8. ISSN: 2331-7418. URL: <http://jcgt.org/published/0006/01/01/>.
- Dupuy, Jonathan, Eric Heitz, and Laurent Belcour (July 2017). “A Spherical Cap Preserving Parameterization for Spherical Distributions.” In: *ACM Transactions on Graphics* 36.4, 139:1–139:12. ISSN: 0730-0301. DOI: 10.1145/3072959.3073694. URL: <http://doi.acm.org/10.1145/3072959.3073694>.
- Dupuy, Jonathan, Eric Heitz, and Eugene d’Eon (2016). “Additional Progress Towards the Unification of Microfacet and Microflake Theories.” In: *Proceedings of the Eurographics Symposium on Rendering: Experimental Ideas & Implementations*. EGSR ’16. Dublin, Ireland: Eurographics Association, pp. 55–63. ISBN: 978-3-03868-019-2. DOI: 10.2312/sre.20161210. URL: <https://doi.org/10.2312/sre.20161210>.
- Eberly, David (2015). *Distance Between Point and Triangle in 3D*. Tech. rep.

- Farag, Aly, Shireen Elhabian, Aly Abdelrehim, Wael Aboelmaaty, Allan Farman, and David Tasman (2013). “Model-Based Human Teeth Shape Recovery from a Single Optical Image with Unknown Illumination.” In: *Medical Computer Vision: Recognition Techniques and Applications in Medical Imaging (MCV '12)*, pp. 263–272.
- Feit, M. D. and J. A. Fleck (1978). “Light propagation in graded-index optical fibers.” In: *Appl. Opt.* 17.24, pp. 3990–3998. DOI: 10.1364/AO.17.003990. URL: <http://ao.osa.org/abstract.cfm?URI=ao-17-24-3990>.
- Frederickx, Roald and Philip Dutré (July 2017). “A Forward Scattering Dipole Model from a Functional Integral Approximation.” In: *ACM Transactions on Graphics* 36.4, 109:1–109:13. ISSN: 0730-0301. DOI: 10.1145/3072959.3073681. URL: <http://doi.acm.org/10.1145/3072959.3073681>.
- Fried, Daniel, Richard E. Glens, John D. B. Featherstone, and Wolf Seka (1995). “Nature of light scattering in dental enamel and dentin at visible and near-infrared wavelengths.” In: *Journal of Applied Optics* 34.7, pp. 1278–1285.
- Frisvad, Jeppe Revall, Toshiya Hachisuka, and Thomas Kim Kjeldsen (Dec. 2014). “Directional Dipole Model for Subsurface Scattering.” In: *ACM Transactions on Graphics* 34.1, 5:1–5:12. ISSN: 0730-0301. DOI: 10.1145/2682629. URL: <http://doi.acm.org/10.1145/2682629>.
- Garrido-Jurado, S., R. Muñoz-Salinas, F.J. Madrid-Cuevas, and M.J. Marín-Jiménez (2014). “Automatic generation and detection of highly reliable fiducial markers under occlusion.” In: *Pattern Recognition* 47.6, pp. 2280–2292. ISSN: 0031-3203. DOI: <http://dx.doi.org/10.1016/j.patcog.2014.01.005>. URL: <http://www.sciencedirect.com/science/article/pii/S0031320314000235>.
- Ghosh, Abhijeet, Tim Hawkins, Pieter Peers, Sune Frederiksen, and Paul Debevec (2008). “Practical modeling and acquisition of layered facial reflectance.” In: *ACM Transactions on Graphics (TOG)*. Vol. 27. 5. ACM, p. 139.
- Gkioulekas, Ioannis, Anat Levin, and Todd Zickler (2016). “An Evaluation of Computational Imaging Techniques for Heterogeneous Inverse Scattering.” In: *Computer Vision – ECCV 2016*. Ed. by Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling. Cham: Springer International Publishing, pp. 685–701. ISBN: 978-3-319-46487-9.
- Gkioulekas, Ioannis, Shuang Zhao, Kavita Bala, Todd Zickler, and Anat Levin (Nov. 2013). “Inverse Volume Rendering with Material Dictionar-

BIBLIOGRAPHY

- ies.” In: *ACM Transactions on Graphics* 32.6, 162:1–162:13. ISSN: 0730-0301. DOI: 10.1145/2508363.2508377. URL: <http://doi.acm.org/10.1145/2508363.2508377>.
- Golub, G. and W. Kahan (1965). “Calculating the Singular Values and Pseudo-Inverse of a Matrix.” In: *Journal of the Society for Industrial and Applied Mathematics Series B Numerical Analysis* 2.2, pp. 205–224. DOI: 10.1137/0702016. eprint: <https://doi.org/10.1137/0702016>. URL: <https://doi.org/10.1137/0702016>.
- Goodman, J.W. (1996). *Introduction to Fourier Optics*. McGraw-Hill Series in Electrical and Computer Engineering: Communications and Signal Processing. McGraw-Hill. ISBN: 9780070242548. URL: <https://books.google.de/books?id=Q11RAAAAMAAJ>.
- Goral, Cindy, Kenneth Torrance, Donald Greenberg, and Bennett Battaile (1984). “Modeling the Interaction of Light Between Diffuse Surfaces.” In: *Proc. 11th Ann. Conf. on Computer Graphics and Interactive Techniques*. SIGGRAPH ’84. ACM, pp. 213–222. ISBN: 0-89791-138-5. DOI: 10.1145/800031.808601. URL: <http://doi.acm.org/10.1145/800031.808601>.
- Guarnera, D., G.C. Guarnera, A. Ghosh, C. Denk, and M. Glencross (2016). “BRDF Representation and Acquisition.” In: *Computer Graphics Forum* 35.2, pp. 625–650. ISSN: 0167-7055. DOI: 10.1111/cgf.12867. URL: <https://doi.org/10.1111/cgf.12867>.
- Guillén, Ibón, Carlos Ureña, Alan King, Marcos Fajardo, Iliyan Georgiev, Jorge López-Moreno, and Adrian Jarabo (2017). “Area-Preserving Parameterizations for Spherical Ellipses.” In: *Computer Graphics Forum (Proceedings of EGSR)* 36.4.
- Gupta, Satish and Robert F. Sproull (1981). “Filtering Edges for Gray-scale Displays.” In: *Proc. 8th Ann. Conf. on Computer Graphics and Interactive Techniques*. SIGGRAPH ’81. Dallas, Texas, USA: ACM, pp. 1–5. ISBN: 0-89791-045-1. DOI: 10.1145/800224.806783. URL: <http://doi.acm.org/10.1145/800224.806783>.
- Habel, Ralf, Per H. Christensen, and Wojciech Jarosz (June 2013). “Photon Beam Diffusion: A Hybrid Monte Carlo Method for Subsurface Scattering.” In: *Computer Graphics Forum (Proceedings of the Eurographics Symposium on Rendering)* 32.4. DOI: 10.1111/cgf.12148.
- Hanrahan, Pat and Wolfgang Krueger (1993). “Reflection from Layered Surfaces Due to Subsurface Scattering.” In: *Proceedings of the 20th Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH ’93. Anaheim, CA: ACM, pp. 165–174. ISBN: 0-89791-601-8. DOI:

- 10.1145/166117.166139. URL: <http://doi.acm.org/10.1145/166117.166139>.
- Harrington, Roger F. (1993). *Field Computation by Moment Methods*. Wiley-IEEE Press. ISBN: 0780310144.
- Harvey, James E., Cynthia L. Vernold, Andrey Krywonos, and Patrick L. Thompson (2000). “Diffracted radiance: a fundamental quantity in non-paraxial scalar diffraction theory: errata.” In: *Appl. Opt.* 39.34, pp. 6374–6375. DOI: 10.1364/AO.39.006374. URL: <http://ao.osa.org/abstract.cfm?URI=ao-39-34-6374>.
- Hašan, Milovš and Ravi Ramamoorthi (Apr. 2013). “Interactive Albedo Editing in Path-traced Volumetric Materials.” In: *ACM Transactions on Graphics* 32.2, 11:1–11:11. ISSN: 0730-0301. DOI: 10.1145/2451236.2451237. URL: <http://doi.acm.org/10.1145/2451236.2451237>.
- Heitz, Eric, Jonathan Dupuy, Cyril Crassin, and Carsten Dachsbacher (July 2015). “The SGGX Microflake Distribution.” In: *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 34.4, 48:1–48:11. ISSN: 0730-0301. DOI: 10.1145/2766988. URL: <http://doi.acm.org/10.1145/2766988>.
- Heitz, Eric, Jonathan Dupuy, Stephen Hill, and David Neubelt (2016a). “Real-time Polygonal-light Shading with Linearly Transformed Cosines.” In: *ACM Transactions on Graphics* 35.4, 41:1–41:8. ISSN: 0730-0301. DOI: 10.1145/2897824.2925895. URL: <http://doi.acm.org/10.1145/2897824.2925895>.
- Heitz, Eric, Johannes Hanika, Eugene d’Eon, and Carsten Dachsbacher (July 2016b). “Multiple-scattering Microfacet BSDFs with the Smith Model.” In: *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 35.4, 58:1–58:14. ISSN: 0730-0301. DOI: 10.1145/2897824.2925943. URL: <http://doi.acm.org/10.1145/2897824.2925943>.
- Heitz, Eric and Stephen Hill (2017). “Linear-Light Shading with Linearly Transformed Cosines.” In: *GPU Zen: Advanced Rendering Techniques*. Ed. by Wolfgang Engel. Bowker Identifier Services.
- Heitz, Eric, Stephen Hill, and Morgan McGuire (2018). “Combining Analytic Direct Illumination and Stochastic Shadows.” In: *Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*. I3D ’18. Montreal, Quebec, Canada: ACM, 2:1–2:11. ISBN: 978-1-4503-5705-0. DOI: 10.1145/3190834.3190852. URL: <http://doi.acm.org/10.1145/3190834.3190852>.
- Helmholtz, H. von (1856). *Handbuch der physiologischen Optik*. Vol. 1. Handbuch der physiologischen Optik. Leipzig: Leopold Voss, p. 169.

BIBLIOGRAPHY

- Henye, L. G. and J. L. Greenstein (1941). “Diffuse radiation in the Galaxy.” In: *Astrophys. J.* 93, pp. 70–83. DOI: 10.1086/144246.
- Hill, Stephen, Stephen McAuley, Jonathan Dupuy, Yoshiharu Gotanda, Eric Heitz, Naty Hoffman, Sébastien Lagarde, Anders Langlands, Ian Megibben, Farhez Rayani, and Charles de Rousiers (2014). “Physically Based Shading in Theory and Practice.” In: *ACM SIGGRAPH 2014 Courses*. SIGGRAPH ’14. Vancouver, Canada: ACM, 23:1–23:8. ISBN: 978-1-4503-2962-0. DOI: 10.1145/2614028.2615431. URL: <http://doi.acm.org/10.1145/2614028.2615431>.
- Holzschuch, Nicolas (Sept. 2015). “Accurate computation of single scattering in participating media with refractive boundaries.” In: *Computer Graphics Forum* 34.6, pp. 48–59. DOI: 10.1111/cgf.12517. URL: <https://hal.inria.fr/hal-01083246>.
- Holzschuch, Nicolas and Romain Pacanowski (July 2017). “A Two-scale Microfacet Reflectance Model Combining Reflection and Diffraction.” In: *ACM Transactions on Graphics* 36.4, 66:1–66:12. ISSN: 0730-0301. DOI: 10.1145/3072959.3073621. URL: <http://doi.acm.org/10.1145/3072959.3073621>.
- Hooke, R. (2005). *Micrographia*. Project Gutenberg (eBook). URL: <http://www.gutenberg.org/etext/15491>.
- Hsia, Jack J. and Joseph C. Richmond (1976). “Bidirectional reflectometry. part I: A high resolution laser bidirectional reflectometer with results on several optical coatings.” In: *Journal of Research of the National Bureau of Standards-A. Physics and Chemistry* 80A, pp. 189–205.
- Igehy, Homan (1999). “Tracing Ray Differentials.” In: *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH ’99. New York, NY, USA: ACM Press / Addison-Wesley Publishing Co., pp. 179–186. ISBN: 0-201-48560-5. DOI: 10.1145/311535.311555. URL: <http://dx.doi.org/10.1145/311535.311555>.
- Ilie, A. and G. Welch (2005). “Ensuring color consistency across multiple cameras.” In: *Tenth IEEE International Conference on Computer Vision (ICCV’05)*. Vol. 2, pp. 1268–1275. DOI: 10.1109/ICCV.2005.88.
- Irawan, Piti and Steve Marschner (Feb. 2012). “Specular Reflection from Woven Cloth.” In: *ACM Transactions on Graphics* 31.1, 11:1–11:20. ISSN: 0730-0301. DOI: 10.1145/2077341.2077352. URL: <http://doi.acm.org/10.1145/2077341.2077352>.
- Jakob, Wenzel (2010). *Mitsuba renderer*. <http://mitsuba-renderer.org>.

- Jakob, Wenzel, Adam Arbree, Jonathan T. Moon, Kavita Bala, and Steve Marschner (July 2010). “A Radiative Transfer Framework for Rendering Materials with Anisotropic Structure.” In: *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 29.4, 53:1–53:13. ISSN: 0730-0301. DOI: 10.1145/1778765.1778790. URL: <http://doi.acm.org/10.1145/1778765.1778790>.
- Jakob, Wenzel, Eugene d’Eon, Otto Jakob, and Steve Marschner (July 2014a). “A Comprehensive Framework for Rendering Layered Materials.” In: *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 33.4, 118:1–118:14. ISSN: 0730-0301. DOI: 10.1145/2601097.2601139. URL: <http://doi.acm.org/10.1145/2601097.2601139>.
- Jakob, Wenzel, Miloš Hašan, Ling-Qi Yan, Jason Lawrence, Ravi Ramamoorthi, and Steve Marschner (2014b). “Discrete Stochastic Microfacet Models.” In: *ACM Transactions on Graphics* 33.4, 115:1–115:10. ISSN: 0730-0301. DOI: 10.1145/2601097.2601186. URL: <http://doi.acm.org/10.1145/2601097.2601186>.
- Jakob, Wenzel and Steve Marschner (July 2012). “Manifold Exploration: A Markov Chain Monte Carlo Technique for Rendering Scenes with Difficult Specular Transport.” In: *ACM Transactions on Graphics* 31.4, 58:1–58:13. ISSN: 0730-0301. DOI: 10.1145/2185520.2185554. URL: <http://doi.acm.org/10.1145/2185520.2185554>.
- Jarabo, Adrian, Carlos Aliaga, and Diego Gutierrez (July 2018). “A Radiative Transfer Framework for Spatially-correlated Materials.” In: *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 37.4, 83:1–83:13. ISSN: 0730-0301. DOI: 10.1145/3197517.3201282. URL: <http://doi.acm.org/10.1145/3197517.3201282>.
- Jensen, Henrik Wann and Per H. Christensen (1998). “Efficient Simulation of Light Transport in Scenes with Participating Media Using Photon Maps.” In: *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH ’98. New York, NY, USA: ACM, pp. 311–320. ISBN: 0-89791-999-8. DOI: 10.1145/280814.280925. URL: <http://doi.acm.org/10.1145/280814.280925>.
- Jensen, Henrik Wann, Stephen R. Marschner, Marc Levoy, and Pat Hanrahan (2001). “A Practical Model for Subsurface Light Transport.” In: *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH ’01. New York, NY, USA: ACM, pp. 511–518. ISBN: 1-58113-374-X. DOI: 10.1145/383259.383319. URL: <http://doi.acm.org/10.1145/383259.383319>.

BIBLIOGRAPHY

- Kajiya, James T. (1986). “The Rendering Equation.” In: *Proc. 13th Ann. Conf. on Computer Graphics and Interactive Techniques*. SIGGRAPH ’86. ACM, pp. 143–150. ISBN: 0-89791-196-2. DOI: 10.1145/15922.15902. URL: <http://doi.acm.org/10.1145/15922.15902>.
- Kajiya, James T. and Brian P Von Herzen (Jan. 1984). “Ray Tracing Volume Densities.” In: *SIGGRAPH Comput. Graph.* 18.3, pp. 165–174. ISSN: 0097-8930. DOI: 10.1145/964965.808594. URL: <http://doi.acm.org/10.1145/964965.808594>.
- Karis, Brian (2013). “Real shading in unreal engine 4.” In: *SIGGRAPH ’13: ACM SIGGRAPH 2013 Courses*. Anaheim, California: ACM. ISBN: 978-1-4503-2339-0.
- Kawai, Masahide, Tomoyori Iwao, Akinobu Maejima, and Shigeo Morishima (Dec. 2014). *Automatic Photorealistic 3D Inner Mouth Restoration from Frontal Images*.
- Kelemen, Csaba and Laszlo Szirmay-Kalos (2001). *A Microfacet Based Coupled Specular-Matte BRDF Model with Importance Sampling*.
- Keller, Alexander (1997). “Instant Radiosity.” In: *Proc. 24th Ann. Conf. on Computer Graphics and Interactive Techniques*. SIGGRAPH ’97. ACM Press / Addison-Wesley Publishing Co., pp. 49–56. ISBN: 0-89791-896-7. DOI: 10.1145/258734.258769. URL: <http://dx.doi.org/10.1145/258734.258769>.
- Khungurn, Pramook, Daniel Schroeder, Shuang Zhao, Kavita Bala, and Steve Marschner (Dec. 2015). “Matching Real Fabrics with Micro-Appearance Models.” In: *ACM Transactions on Graphics* 35.1, 1:1–1:26. ISSN: 0730-0301. DOI: 10.1145/2818648. URL: <http://doi.acm.org/10.1145/2818648>.
- Kingma, Diederik P. and Jimmy Ba (2014). “Adam: A Method for Stochastic Optimization.” In: *CoRR* abs/1412.6980. URL: <http://arxiv.org/abs/1412.6980>.
- Klehm, Oliver, Fabrice Rousselle, Marios Papas, Derek Bradley, Christophe Hery, Bernd Bickel, Wojciech Jarosz, and Thabo Beeler (May 2015). “Recent Advances in Facial Appearance Capture.” In: *Computer Graphics Forum* 34.2, pp. 709–733.
- Koudelka, Melissa L., Sebastian Magda, Peter N. Belhumeur, and David J. Kriegman (2003). “Acquisition, compression, and synthesis of bidirectional texture functions.” In: *In ICCV 03 Workshop on Texture Analysis and Synthesis*.

- Krywonos, Andrey (2006). *Predicting Surface Scatter Using a Linear Systems Formulation of Non-paraxial Scalar Diffraction*. University of Central Florida. ISBN: 9780542975868. URL: http://etd.fcla.edu/CF/CFE0001446/Krywonos_Andrey_200612_PhD.pdf.
- Lafortune, Eric P. and Yves D. Willems (Aug. 1996). “Rendering participating media with bidirectional path tracing.” In: *Proceedings of the Eurographics Workshop on Rendering Techniques*. Porto, Portugal: Springer-Verlag, pp. 91–100.
- Larsen, Christian Thode, Jeppe Revall Frisvad, Peter Dahl Ejby Jensen, and Jakob Andreas Bærentzen (2012). “Real-Time Rendering of Teeth with No Preprocessing.” In: *Advances in Visual Computing*. Ed. by George Bebis, Richard Boyle, Bahram Parvin, Darko Koracin, Charless Fowlkes, Sen Wang, Min-Hyung Choi, Stephan Mantler, Jürgen Schulze, Daniel Acevedo, Klaus Mueller, and Michael Papka. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 334–345. ISBN: 978-3-642-33191-6.
- Lecocq, Pascal, Arthur Dufay, Gaël Sourimant, and Jean-Eudes Marvie (2016). “Accurate Analytic Approximations for Real-time Specular Area Lighting.” In: *Proc. 20th ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*. I3D ’16. Redmond, Washington: ACM, pp. 113–120. ISBN: 978-1-4503-4043-4. DOI: 10.1145/2856400.2856403. URL: <http://doi.acm.org/10.1145/2856400.2856403>.
- Levenberg, K. (1944). “A method for the solution of certain problems in least squares.” In: *Quarterly Journal on Applied Mathematics* 2, pp. 164–168.
- Levin, Anat, Daniel Glasner, Ying Xiong, Frédo Durand, William Freeman, Wojciech Matusik, and Todd Zickler (2013). “Fabricating BRDFs at High Spatial Resolution Using Wave Optics.” In: *ACM Transactions on Graphics* 32.4, 144:1–144:14. ISSN: 0730-0301. DOI: 10.1145/2461912.2461981. URL: <http://doi.acm.org/10.1145/2461912.2461981>.
- Li, Hongsong, Fabio Pellacini, and Kenneth E. Torrance (2005). “A Hybrid Monte Carlo Method for Accurate and Efficient Subsurface Scattering.” In: *Rendering Techniques (Proceedings of the Eurographics Symposium on Rendering)*, pp. 283–290.
- Li, Tzu-Mao, Miika Aittala, and Jaakko Lehtinen Frédo Durand (Nov. 2018). “Differentiable Monte Carlo Ray Tracing through Edge Sampling.” In: *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 37.6. to appear. DOI: 10.1145/3272127.3275098.
- Lourakis, Manolis I. A. and Antonis A. Argyros (Mar. 2009). “SBA: A Software Package for Generic Sparse Bundle Adjustment.” In: *ACM Trans.*

BIBLIOGRAPHY

- Math. Softw.* 36.1, 2:1–2:30. ISSN: 0098-3500. DOI: 10.1145/1486525.1486527. URL: <http://doi.acm.org/10.1145/1486525.1486527>.
- Mallick, Satya P, Todd E Zickler, David J Kriegman, and Peter N Belhumeur (2005). “Beyond Lambert: Reconstructing specular surfaces using color.” In: *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*. Vol. 2. IEEE, pp. 619–626.
- Malzbender, Tom, Dan Gelb, and Hans Wolters (2001). “Polynomial Texture Maps.” In: *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH '01. New York, NY, USA: ACM, pp. 519–528. ISBN: 1-58113-374-X. DOI: 10.1145/383259.383320. URL: <http://doi.acm.org/10.1145/383259.383320>.
- Mandel, Leonard and Emil Wolf (1995). *Optical Coherence and Quantum Optics*. Cambridge University Press. ISBN: 9780521417112.
- Mangano, Francesco Guido, Andrea Gandolfi, Giuseppe Luongo, and Silvia Logozzo (2017). “Intraoral scanners in dentistry: a review of the current literature.” In: *BMC oral health*.
- Marquardt, Donald W. (1963). “An Algorithm for Least-Squares Estimation of Nonlinear Parameters.” In: *Journal of the Society for Industrial and Applied Mathematics* 11, pp. 431–441.
- Marsh, J. H. (2002). “Optical sciences in Scotland.” In: *The 15th Annual Meeting of the IEEE Lasers and Electro-Optics Society*. Vol. 1, 9–10 vol.1. DOI: 10.1109/LEOS.2002.1133892.
- Mehl, Albert and Volker Blanz (2005). “A new approach for automatic reconstruction of occlusal surfaces with the biogeneric tooth model.” In: *Int. J. Comput. Dent.* 8, pp. 13–25.
- Meng, Zhuo, Steve Yao, Hui Yao, Yan Liang, Tiegeng Liu, Yanni Li, Guanhua Wang, and Shoufeng Lan (2009). “Measurement of the refractive index of human teeth by optical coherence tomography.” In: *Journal of Biomedical Optics* 14. DOI: 10.1117/1.3130322. URL: <http://dx.doi.org/10.1117/1.3130322>.
- Miller, Gavin (1994). “Efficient Algorithms for Local and Global Accessibility Shading.” In: *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH '94. New York, NY, USA: ACM, pp. 319–326. ISBN: 0-89791-667-0. DOI: 10.1145/192161.192244. URL: <http://doi.acm.org/10.1145/192161.192244>.
- Moon, Jonathan T., Bruce Walter, and Stephen R. Marschner (2007). “Rendering Discrete Random Media Using Precomputed Scattering Solutions.” In: *Proceedings of the 18th Eurographics Conference on Rendering Tech-*

- niques*. EGSR'07. Grenoble, France: Eurographics Association, pp. 231–242. ISBN: 978-3-905673-52-4. DOI: 10.2312/EGWR/EGSR07/231-242. URL: <http://dx.doi.org/10.2312/EGWR/EGSR07/231-242>.
- Mostafa, E., S. Elhabian, A. Abdelrahim, S. Elshazly, and A. Farag (2014). “Statistical morphable model for human teeth restoration.” In: *IEEE ICIP*, pp. 4285–4288.
- Müller, Gero, Jan Meseth, Mirko Sattler, Ralf Sarlette, and Reinhard Klein (Sept. 2004). “Acquisition, Synthesis and Rendering of Bidirectional Texture Functions.” In: *Eurographics 2004, State of the Art Reports*. Ed. by Christophe Schlick and Werner Purgathofer. INRIA and Eurographics Association, pp. 69–94.
- Munoz, Adolfo, Jose I. Echevarria, Francisco J. Seron, Jorge Lopez-Moreno, Mashhuda Glencross, and Diego Gutierrez (2011). “BSSRDF Estimation from Single Images.” In: *Computer Graphics Forum* 30.2, pp. 455–464. DOI: 10.1111/j.1467-8659.2011.01873.x. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1467-8659.2011.01873.x>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-8659.2011.01873.x>.
- Musbach, A., G. W. Meyer, F. Reitich, and S. H. Oh (2013). “Full Wave Modelling of Light Propagation and Reflection.” In: *Computer Graphics Forum* 32.6, pp. 24–37. ISSN: 1467-8659. DOI: 10.1111/cgf.12012. URL: <http://dx.doi.org/10.1111/cgf.12012>.
- Nicodemus, F. E. (1970). “Reflectance Nomenclature and Directional Reflectance and Emissivity.” In: *Appl. Opt.* 9.6, pp. 1474–1475. DOI: 10.1364/AO.9.001474. URL: <http://ao.osa.org/abstract.cfm?URI=ao-9-6-1474>.
- Nocedal, Jorge and Stephen J. Wright (2006). *Numerical Optimization*. 2nd. New York, NY, USA: Springer.
- Omachi, S., K. Saito, H. Aso, S. Kasahara, S. Yamada, and K. Kimura (2007). “Tooth shape reconstruction from ct images using spline Curves.” In: *Wavelet Analysis and Pattern Recognition*. Vol. 1, pp. 393–396.
- Papas, Marios, Christian Regg, Wojciech Jarosz, Bernd Bickel, Philip Jackson, Wojciech Matusik, Steve Marschner, and Markus Gross (July 2013). “Fabricating Translucent Materials Using Continuous Pigment Mixtures.” In: *ACM Transactions on Graphics* 32.4, 146:1–146:12. ISSN: 0730-0301. DOI: 10.1145/2461912.2461974. URL: <http://doi.acm.org/10.1145/2461912.2461974>.

BIBLIOGRAPHY

- Pauly, Mark, Thomas Kollig, and Alexander Keller (2000). “Metropolis light transport for participating media.” In: *Proceedings of the Eurographics Workshop on Rendering Techniques*. London, UK: Springer-Verlag, pp. 11–22.
- Pharr, Matt, Wenzel Jakob, and Greg Humphreys (2016). *Physically Based Rendering: From Theory to Implementation*. 3rd. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc. ISBN: 978-0-12-800645-0.
- Picott, Kevin P. (1992). “Extensions of the Linear and Area Lighting Models.” In: *IEEE Comput. Graph. Appl.* 12.2, pp. 31–38. ISSN: 0272-1716. DOI: 10.1109/38.124286. URL: <http://dx.doi.org/10.1109/38.124286>.
- Pop-Ciutrla, Ioana-Sofia, Razvan Ghinea, Maria del Mar Perez Gomez, Horatiu Alexandru Colosi, Diana Dudea, and Mandra Badea (2015). “Dentine scattering, absorption, transmittance and light reflectivity in human incisors, canines and molars.” In: *Journal of Dentistry* 43, pp. 1116–1124.
- Poulin, Pierre and John Amanatides (1991). “Shading and shadowing with linear light sources.” In: *Computers & Graphics* 15.2, pp. 259–265. ISSN: 0097-8493. DOI: [http://dx.doi.org/10.1016/0097-8493\(91\)90079-W](http://dx.doi.org/10.1016/0097-8493(91)90079-W). URL: <http://www.sciencedirect.com/science/article/pii/009784939190079W>.
- Press, W. H., S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery (1992). *Numerical Recipes in Fortran 77: The Art of Scientific Computing*. 2nd. Cambridge University Press, p. 178.
- Ramamoorthi, Ravi (Apr. 2009). “Precomputation-Based Rendering.” In: *Found. Trends. Comput. Graph. Vis.* 3.4, pp. 281–369. ISSN: 1572-2740. DOI: 10.1561/06000000021. URL: <http://dx.doi.org/10.1561/06000000021>.
- Raymond, Boris, Gaël Guennebaud, and Pascal Barla (2016). “Multi-scale Rendering of Scratched Materials Using a Structured SV-BRDF Model.” In: *ACM Transactions on Graphics* 35.4, 57:1–57:11. ISSN: 0730-0301. DOI: 10.1145/2897824.2925945. URL: <http://doi.acm.org/10.1145/2897824.2925945>.
- Robertson, M. A., S. Borman, and R. L. Stevenson (Apr. 2003). “Estimation-theoretic approach to dynamic range enhancement using multiple exposures.” In: *Journal of Electronic Imaging* 12, pp. 219–228. DOI: 10.1117/1.1557695.
- Sadeghi, Iman, Oleg Bisker, Joachim De Deken, and Henrik Wann Jensen (Apr. 2013). “A Practical Microcylinder Appearance Model for Cloth

- Rendering.” In: *ACM Transactions on Graphics* 32.2, 14:1–14:12. ISSN: 0730-0301. DOI: 10.1145/2451236.2451240. URL: <http://doi.acm.org/10.1145/2451236.2451240>.
- Saito, Takafumi and Tokiichiro Takahashi (1990). “Comprehensible Rendering of 3-D Shapes.” In: *Proc. 17th Ann. Conf. on Computer Graphics and Interactive Techniques*. SIGGRAPH ’90. Dallas, TX, USA: ACM, pp. 197–206. ISBN: 0-89791-344-2. DOI: 10.1145/97879.97901. URL: <http://doi.acm.org/10.1145/97879.97901>.
- Scandolo, Leonardo, Sungkil Lee, and Elmar Eisemann (2018). “Quad-Based Fourier Transform for Efficient Diffraction Synthesis.” In: *Computer Graphics Forum (Proceedings of the Eurographics Symposium on Rendering)* 37.4, pp. 167–176. DOI: 10.1111/cgf.13484. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.13484>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.13484>.
- Schröder, Kai, Reinhard Klein, and Arno Zinke (July 2011). “A Volumetric Approach to Predictive Rendering of Fabrics.” In: *Computer Graphics Forum (Proceedings of EGSR 2011)* 30.4, pp. 1277–1286.
- Schwartz, Christopher, Ralf Sarlette, Michael Weinmann, and Reinhard Klein (June 2013). “DOME II: A Parallelized BTF Acquisition System.” In: *Eurographics Workshop on Material Appearance Modeling: Issues and Acquisition*. Ed. by Holly Rushmeier and Reinhard Klein. Zaragoza, Spain: Eurographics Association, pp. 25–31. ISBN: 978-3-905674-48-4. DOI: 10.2312/MAM.MAM2013.025-031. URL: <http://diglib.eg.org/EG/DL/WS/MAM/MAM2013/025-031.pdf>.
- Seeliger, R. (1888). *Munch. Akad. II. Kl. Sitzungsber 18*.
- Shirley, Peter, Changyaw Wang, and Kurt Zimmerman (1996). “Monte Carlo Techniques for Direct Lighting Calculations.” In: *ACM Transactions on Graphics* 15.1, pp. 1–36. ISSN: 0730-0301. DOI: 10.1145/226150.226151. URL: <http://doi.acm.org/10.1145/226150.226151>.
- Sloan, Peter-Pike, Jan Kautz, and John Snyder (2002). “Precomputed Radiance Transfer for Real-time Rendering in Dynamic, Low-frequency Lighting Environments.” In: *Proc. 29th Ann. Conf. on Computer Graphics and Interactive Techniques*. SIGGRAPH ’02. San Antonio, Texas: ACM, pp. 527–536. ISBN: 1-58113-521-1. DOI: 10.1145/566570.566612. URL: <http://doi.acm.org/10.1145/566570.566612>.
- Snyder, J. (1996). *Area light sources for real-time graphics*. Tech. rep. MSR-TR-96-11. Microsoft Research.

BIBLIOGRAPHY

- Sommerfeld, A. (1954). *Optics. Lectures on Theoretical Physics*. New York: Academic Press.
- Spitzer, D. and J. J. Ten Bosch (1975). “The absorption and scattering of light in bovine and human dental enamel.” In: *Calcified Tissue Research* 17.2, pp. 129–137.
- Stam, Jos (1995). “Multiple scattering as a diffusion process.” In: *Rendering Techniques (Proceedings of the Eurographics Workshop on Rendering)*, pp. 41–50.
- Stam, Jos (1999). “Diffraction Shaders.” In: *Proc. 26th Ann. Conf. on Computer Graphics and Interactive Techniques*. SIGGRAPH ’99. ACM Press / Addison-Wesley Publishing Co., pp. 101–110. ISBN: 0-201-48560-5. DOI: 10.1145/311535.311546. URL: <http://dx.doi.org/10.1145/311535.311546>.
- Sun, Bo, Kalyan Sunkavalli, Ravi Ramamoorthi, Peter Belhumeur, and Shree Nayar (2007). “Time-Varying BRDFs.” In: *IEEE Transactions on Visualization and Computer Graphics*.
- Tanaka, Toshimitsu and Tokiichiro Takahashi (1997). “Fast Analytic Shading and Shadowing for Area Light Sources.” In: *Computer Graphics Forum* 16, pp. C231–C240. ISSN: 1467-8659. DOI: 10.1111/1467-8659.16.3conferenceissue.24. URL: <http://dx.doi.org/10.1111/1467-8659.16.3conferenceissue.24>.
- Tariq, Sarah, Andrew Gardner, Ignacio Llamas, Andrew Jones, Paul Debevec, and Greg Turk (2006). “Efficient estimation of spatially varying subsurface scattering parameters.” In: *Vision, Modeling, and Visualization (VMV2006)*, pp. 129–136.
- Toisoul, Antoine and Abhijeet Ghosh (July 2017). “Practical Acquisition and Rendering of Diffraction Effects in Surface Reflectance.” In: *ACM Transactions on Graphics* 36.4. ISSN: 0730-0301. DOI: 10.1145/3072959.3012001. URL: <http://doi.acm.org/10.1145/3072959.3012001>.
- Tong, Xin, Jiaping Wang, Stephen Lin, Baining Guo, and Heung-Yeung Shum (July 2005). “Modeling and Rendering of Quasi-homogeneous Materials.” In: *ACM Transactions on Graphics* 24.3, pp. 1054–1061. ISSN: 0730-0301. DOI: 10.1145/1073204.1073311. URL: <http://doi.acm.org/10.1145/1073204.1073311>.
- Traa, Johannes (2013). *Least-Squares Intersection of Lines*. Tech. rep. URL: http://cal.cs.illinois.edu/~johannes/research/LS_line_intersect.pdf.

- Trowbridge, T. S. and K. P. Reitz (May 1975). “Average irregularity representation of a rough surface for ray reflection.” In: *J. Opt. Soc. Am.* 65.5, pp. 531–536. DOI: 10.1364/JOSA.65.000531. URL: <http://www.osapublishing.org/abstract.cfm?URI=josa-65-5-531>.
- Ureña, Carlos, Marcos Fajardo, and Alan King (2013). “An Area-preserving Parametrization for Spherical Rectangles.” In: *Proc. Eurographics Symposium on Rendering*. EGSR ’13. Zaragoza, Spain: Eurographics Association, pp. 59–66. DOI: 10.1111/cgf.12151. URL: <http://dx.doi.org/10.1111/cgf.12151>.
- Veach, Eric (1998). “Robust Monte Carlo Methods for Light Transport Simulation.” AAI9837162. PhD thesis. Stanford, CA, USA. ISBN: 0-591-90780-1.
- Veach, Eric and Leonidas J. Guibas (1997). “Metropolis Light Transport.” In: *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH ’97. New York, NY, USA: ACM Press / Addison-Wesley Publishing Co., pp. 65–76. ISBN: 0-89791-896-7. DOI: 10.1145/258734.258775. URL: <http://dx.doi.org/10.1145/258734.258775>.
- Velinov, Zdravko and Matthias B. Hullin (Oct. 2016). “An Interactive Appearance Model for Microscopic Fiber Surfaces.” In: *Vision, Modeling, and Visualization 2016*.
- Velinov, Zdravko, Marios Papas, Derek Bradley, Paulo Gotardo, Parsa Mirdelghan, Steve Marschner, Jan Novák, and Thabo Beeler (Dec. 2018a). “Appearance Capture and Modeling of Human Teeth.” In: *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 37.6, 207:1–207:13. ISSN: 0730-0301. DOI: 10.1145/3272127.3275098. URL: <http://doi.acm.org/10.1145/3272127.3275098>.
- Velinov, Zdravko, Sebastian Werner, and Matthias B. Hullin (May 2018b). “Real-Time Rendering of Wave-Optical Effects on Scratched Surfaces.” In: *Computer Graphics Forum* 37.2. ISSN: 1467-8659. DOI: 10.1111/cgf.13347.
- Walter, Bruce, Stephen R. Marschner, Hongsong Li, and Kenneth E. Torrance (2007). “Microfacet Models for Refraction Through Rough Surfaces.” In: *Proceedings of the 18th Eurographics Conference on Rendering Techniques*. EGSR’07. Grenoble, France: Eurographics Association, pp. 195–206. ISBN: 978-3-905673-52-4. DOI: 10.2312/EGWR/EGSR07/195-206. URL: <http://dx.doi.org/10.2312/EGWR/EGSR07/195-206>.

BIBLIOGRAPHY

- Walter, Bruce, Shuang Zhao, Nicolas Holzschuch, and Kavita Bala (Aug. 2009). “Single Scattering in Refractive Media with Triangle Mesh Boundaries.” In: *ACM Transactions on Graphics* 28.3, 92:1–8. DOI: 10.1145/1531326.1531398. URL: <https://hal.inria.fr/inria-00388481>.
- Wang, Jiaping, Xin Tong, Stephen Lin, Minghao Pan, Chao Wang, Hujun Bao, Baining Guo, and Heung-Yeung Shum (July 2006). “Appearance Manifolds for Modeling Time-variant Appearance of Materials.” In: *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 25.3, pp. 754–761. ISSN: 0730-0301. DOI: 10.1145/1141911.1141951. URL: <http://doi.acm.org/10.1145/1141911.1141951>.
- Wang, Jiaping, Shuang Zhao, Xin Tong, Stephen Lin, Zhouchen Lin, Yue Dong, Baining Guo, and Heung-Yeung Shum (Mar. 2008). “Modeling and Rendering of Heterogeneous Translucent Materials Using the Diffusion Equation.” In: *ACM Transactions on Graphics* 27.1, 9:1–9:18. ISSN: 0730-0301. DOI: 10.1145/1330511.1330520. URL: <http://doi.acm.org/10.1145/1330511.1330520>.
- Wenger, Andreas, Andrew Gardner, Chris Tchou, Jonas Unger, Tim Hawkins, and Paul Debevec (July 2005). “Performance Relighting and Reflectance Transformation with Time-multiplexed Illumination.” In: *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 24.3, pp. 756–764. ISSN: 0730-0301. DOI: 10.1145/1073204.1073258. URL: <http://doi.acm.org/10.1145/1073204.1073258>.
- Werner, Sebastian, Zdravko Velinov, Wenzel Jakob, and Matthias B. Hullin (Nov. 2017). “Scratch Iridescence: Wave-optical Rendering of Diffractive Surface Structure.” In: *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 36.6, 207:1–207:14. ISSN: 0730-0301. DOI: 10.1145/3130800.3130840. URL: <http://doi.acm.org/10.1145/3130800.3130840>.
- Weyrich, Tim, Wojciech Matusik, Hanspeter Pfister, Bernd Bickel, Craig Donner, Chien Tu, Janet McAndless, Jinho Lee, Addy Ngan, Henrik Wann Jensen, et al. (2006). “Analysis of human faces using a measurement-based skin reflectance model.” In: *ACM Transactions on Graphics (TOG)*. Vol. 25. 3. ACM, pp. 1013–1024.
- Wu, Chenglei, Derek Bradley, Pablo Garrido, Michael Zollhöfer, Christian Theobalt, Markus Gross, and Thabo Beeler (Nov. 2016). “Model-based Teeth Reconstruction.” In: *ACM Transactions on Graphics* 35.6, 220:1–220:13. ISSN: 0730-0301. DOI: 10.1145/2980179.2980233. URL: <http://doi.acm.org/10.1145/2980179.2980233>.

- Yan, Ling-Qi, Miloš Hašan, Wenzel Jakob, Jason Lawrence, Steve Marschner, and Ravi Ramamoorthi (2014). “Rendering Glints on High-resolution Normal-mapped Specular Surfaces.” In: *ACM Transactions on Graphics* 33.4, 116:1–116:9. ISSN: 0730-0301. DOI: 10.1145/2601097.2601155. URL: <http://doi.acm.org/10.1145/2601097.2601155>.
- Yan, Ling-Qi, Miloš Hašan, Steve Marschner, and Ravi Ramamoorthi (2016). “Position-normal Distributions for Efficient Rendering of Specular Microstructure.” In: *ACM Transactions on Graphics* 35.4, 56:1–56:9. ISSN: 0730-0301. DOI: 10.1145/2897824.2925915. URL: <http://doi.acm.org/10.1145/2897824.2925915>.
- Yan, Ling-Qi, Miloš Hašan, Bruce Walter, Steve Marschner, and Ravi Ramamoorthi (July 2018). “Rendering Specular Microgeometry with Wave Optics.” In: *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 37.4, 75:1–75:10. ISSN: 0730-0301. DOI: 10.1145/3197517.3201351. URL: <http://doi.acm.org/10.1145/3197517.3201351>.
- Yan, Ling-Qi, Henrik Wann Jensen, and Ravi Ramamoorthi (July 2017). “An Efficient and Practical Near and Far Field Fur Reflectance Model.” In: *ACM Transactions on Graphics* 36.4, 67:1–67:13. ISSN: 0730-0301. DOI: 10.1145/3072959.3073600. URL: <http://doi.acm.org/10.1145/3072959.3073600>.
- Yanagisawa, R, Y Sugaya, S Kasahara, and S Omachi (2014). “Tooth shape reconstruction from dental CT images with the region-growing method.” In: *Dentomaxillofacial Radiology* 43.6, p. 20140080.
- Zeiler, M. D. (Dec. 2012). “ADADELTA: An Adaptive Learning Rate Method.” In: *ArXiv e-prints*. arXiv: 1212.5701 [cs.LG].
- Zeltner, Tizian and Wenzel Jakob (July 2018). “The Layer Laboratory: A Calculus for Additive and Subtractive Composition of Anisotropic Surface Reflectance.” In: *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 37.4, 74:1–74:14. ISSN: 0730-0301. DOI: 10.1145/3197517.3201321. URL: <http://doi.acm.org/10.1145/3197517.3201321>.
- Zhao, Shuang, Wenzel Jakob, Steve Marschner, and Kavita Bala (2011). “Building Volumetric Appearance Models of Fabric Using Micro CT Imaging.” In: *ACM SIGGRAPH 2011 Papers*. SIGGRAPH ’11. Vancouver, British Columbia, Canada: ACM, 44:1–44:10. ISBN: 978-1-4503-0943-1. DOI: 10.1145/1964921.1964939. URL: <http://doi.acm.org/10.1145/1964921.1964939>.

BIBLIOGRAPHY

Zheng, Shu-Xian, Jia Li, and Qing-Feng Sun (2011). “A Novel 3D Morphing Approach for Tooth Occlusal Surface Reconstruction.” In: *Comput. Aided Des.* 43.3, pp. 293–302.