2025
204

Dario Lanza

Modeling, Perception and Editing of Volumetric Materials

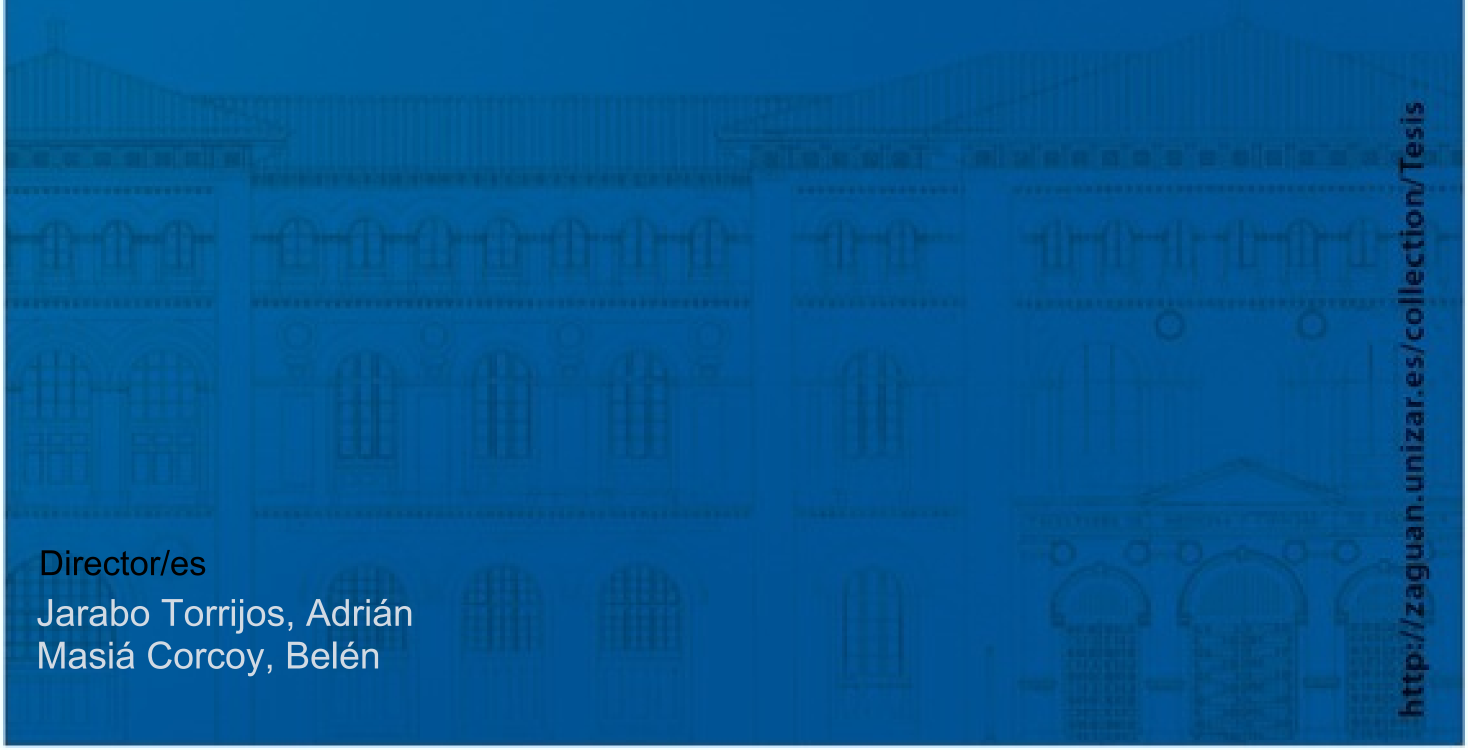

Director/es

Jarabo Torrijos, Adrián
Masiá Corcoy, Belén

ISSN 2254-7606

Prensas de la Universidad
Universidad Zaragoza

Universidad de Zaragoza
Servicio de Publicaciones

ISSN 2254-7606

Tesis Doctoral

MODELING, PERCEPTION AND EDITING OF VOLUMETRIC MATERIALS

Autor

Dario Lanza

Director/es

Jarabo Torrijos, Adrián
Masiá Corcoy, Belén

UNIVERSIDAD DE ZARAGOZA
Escuela de Doctorado

2025

Tesis Doctoral

Modeling, Perception and Editing of Volumetric Materials

Autor

Dario Lanza

Director/es

Dra. Dª. Masiá Corcoy, Belén
Dr. D. Jarabo Torrijos, Adrián

Universidad de Zaragoza
Escuela de Doctorado
2024

MODELING, PERCEPTION AND EDITING OF VOLUMETRIC MATERIALS

WRITTEN BY:

DARIO LANZA

SUPERVISED BY:

BELEN MASIÁ AND ADRIÁN JARABO

Tesis Doctoral - Ingeniería De Sistemas E Informática
Departamento De Informatíca E Ingeniería De Sistemas
Graphics and Imaging Lab
Universidad de Zaragoza

December 2024

To my parents, my brother, my family and my friends

ABSTRACT

This thesis focuses on complex volumetric materials for computer graphics, from their physically-based modeling to their perception and intuitive editing. Volumetric materials are ubiquitous in nature and manufactured objects, and range from organic tissues such as skin, teeth, food, wood, or leaves; to marble, jade, or ice; to aggregates of complex materials including hair, fur, and granular materials. As opposed to local light-matter interactions in opaque materials, in volumetric materials light penetrates inside the object, interacting with the matter inside it multiple times before leaving the object. As these interactions can be modeled at different scales (micro, meso and macro), a common approach is to first study the phenomena that happens at the micro-scale, and from then, through statistical aggregation, describe the final macro-scale behavior. The aggregated micro-scale phenomena are described by considering the ensemble averaged *optical properties* of matter in a small fraction of volume. In this thesis, we apply this statistical approach to *model* volumetric materials such as foundation cosmetics and pennaceous feathers. For cosmetics, inspired by scanning electron microscope (SEM) images, we model the optical properties of the material by considering flake-like or diffuse particles. For feathers, we model the aggregated barckreflection radiance of barb's medullas, as an ideal lambertian diffuser. Additionally we consider the aggregated effect for all the feather with a masking analytical term derived from the regular structure of barbs and barbules. The final appearance of a material, however, is not only influenced by the light-matter interactions (and thus the aforementioned optical properties), but also by how we, humans, *perceive* it. External factors, such as geometry and lighting, can alter this perception, acting as confounding factors. The way we perceive does not only affect modeling and simulation of these materials, but also their editing, which is fundamental for content creation. For volumetric materials, editing is particularly challenging since manipulating parameters often leads to unintuitive results. The second part of this thesis focuses on the perception and editing of volumetric materials, specifically, we examine aspects related to the *perception* of translucent materials, a class of materials that can be modeled as volumetric materials. We investigate whether dynamic lighting can alleviate the effects caused by the confounding factor of light, as well as analyze the performance of different similarity metrics for translucent materials. We then leverage the best performing similarity metric to improve on the *editing* of translucent materials, by introducing a two-dimensional manifold of translucent appearances over which we build a novel hybrid interface.

MEASURABLE CONTRIBUTIONS

This thesis has lead to the following publications, more details can be found in Section 1.4:

- Three journal publications in Computer Graphics Forum [126, 127, 173].
- One peer-reviewed publication at the international conference ACM Symposium on Applied Perception (SAP) [125].

Two research stays (totaling four months), one at Denmark Technical University (DTU) hosted by Prof. Jeppe Revall Frisvad (two months), and another one at Maxon Computer GmbH hosted by Dr. Enrico Rinaldi and Sebastian Häfele (two months).

RESUMEN

Esta tesis se centra en los materiales volumétricos complejos para gráficos por ordenador, desde su modelado basado en la física hasta su percepción y edición intuitiva. Los materiales volumétricos son omnipresentes en la naturaleza y los objetos manufacturados, y cubren desde tejidos orgánicos como piel, dientes, alimentos, madera u hojas; hasta el mármol, el jade o el hielo; pasando por agregados de materiales complejos como el pelo (animal y humano) y los materiales granulares como la sal o la arena. A diferencia de los materiales opacos, en los materiales volumétricos la luz entra en el interior del objeto, interactuando con la materia de su interior varias veces antes de salir del objeto. Dado que estas interacciones se pueden modelar a diferentes escalas, un enfoque común consiste en estudiar primero los fenómenos que ocurren a microescala y, a partir de ahí, mediante agregación estadística, describir el comportamiento final a macroescala.. En esta tesis, aplicamos este marco teórico para modelar la apariencia de dos materiales volumétricos:cosméticos de base y las plumas penáceas. Para los cosméticos, inspiradosen las imágenes del microscopio electrónico, modelamos las propiedades ópticas del material considerando partículas en forma de plaquetas o difusas. En el caso de las plumas, modelamos estadísticamente la reflectancia de barbas y bárbulos, , y derivamos su comportamiento agregado considerando la estructura mesoscópica de la pluma.

Sin embargo, la apariencia de un material no sólo depende de las interacciones ópticas luz-materia, sino también de cómo lo percibimos los seres humanos. La forma en que percibimos no sólo afecta al modelado y la simulación de estos materiales, sino también a su edición, fundamental para la creación de contenidos. En el caso de los materiales volumétricos, la edición es especialmente difícil, ya que los parámetros ópticos son espacios de alta dimensionalidad, y poco intuitivos para usuarios noveles. La segunda parte de esta tesis se centra en la percepción y edición de materiales volumétricos, y en concreto de materiales translúcidos. Para ello, investigamos el efecto de introducir iluminación dinámica en la percepción de objetos translúcidos, y analizamos el rendimiento de diferentes métricas de similitud para materiales translúcidos. A continuación, aprovechamos esta métrica de similitud para mejorar la edición de materiales translúcidos, introduciendo un espacio bidimensional de apariencias translúcidas, sobre el que podemos navegar para editar la apariencia de objetos usando una interfaz intuitiva de edición.

CONCLUSIONES

En la primera parte de esta tesis, nos hemos centradoen el modelado apariencia de materiales volumétricos desde un enfoque basado en física. En concreto hemos propuesto dos modelos de reflectancia para bases cosméticas y plumas penáceas. En el Capítulo 3 hemos modeladolas bases cosméticas como un material volumétrico compuesto por dos tipos de partículas: difusores y plaquetas. En el Capítulo 4 hemos modelado la reflectancia de las barbas y bárbulos que forman la pluma a escala microscópica, y la hemos combinado en un modelo de reflectancia teniendo en cuenta la oclusión debido a la estructura de la superficie de la pluma.

En la segunda parte de la tesis, nos centramos en un subconjunto de materiales volumétricos, concretamente en los materiales translúcidos, y estudiamos su percepción y edición. El Capítulo 5 investiga la percepción de materiales translúcidos en condiciones de iluminación dinámica. Mientras que los trabajos anteriores se centraban en la percepción de materiales translúcidos con estímulos estáticos, nosotros estudiamos cómo se ve afectada la percepción cuando se utilizan estímulos dinámicos. En el Capítulo 6 investigamos nuevas interfaces de edición trabajando en un espacio de dimensiones reducidas que correla con la percepción humana de la translucencia, en lugar de en el espacio de los parámetros ópticos. Para ello, hemos construido un espacio de apariencia translúcida utilizando una métrica de distancia que correla con la percepción humana, tal y como evaluamos mediante un estudio de usuarios.

ACKNOWLEDGEMENTS

Almost four years have passed, since I began this interesting journey in the research field of computer graphics. Now that I am close to finish this journey, it only seems fit to thank all the people that, in some way or other, made this journey easier, lighter and, overall, awesome.

Adrian and Belen, I will always be grateful for having you two as supervisors. I had a good time during this PhD and is largely due to you. You have taught me many things, both at a technical and personal level that I will treasure in the future years. Thank you very much also for the infinite patience that you had with me, and for all the help that you offered me during these years.

Thanks to all the co-authors with which whom I have shared some projects: *Alina*, thank you for your availability for measurements and your unwavering positivity about the cosmetic project. I would like to thank also *Adolfo*, for his help and advice and support (knowing that you are swamped with duties from all sides), and *Jeppe*, for hosting me at DTU and for being a source of balance. *Juan Raul*, we began this journey together, sharing both the highs and lows along the way. Knowing I could count on you during tight deadlines made the challenges easier to face. Although we often have different points of view on many topics, I always appreciated your professionalism, and overall your will to grow by always going outside of your comfort zone.

Thanks to the people at the Graphics and Imaging Lab for being such an amazing research lab. Thanks to *Jorge García, Sergio, Jorge Pina, Óscar and Maria* for being always such a group of proactive people, always ready to discuss new theories and with a great sense of humor. Thanks also to *Daniel Martín, Ana and Julio* for always being willing to help and a great source of humor. Special thanks to: *Edurne*, for listening to me, often spilling wise suggestions, and being a source of interesting discussion during coffee breaks. *Nestor*, you are truly one of the most coherent people I have ever met: Your resolution in acting following your ideals, even if might harm you, is truly commendable. Thanks for all the evenings spent climbing, the discussions about life, the books, and the various beers. *Julia*, you are always smiling and laughing, and yet, having deep conversations, about life and the future, always felt natural. *Martinxo*, for always having a direct and honest opinion, without too many filters. *Pablo*, for being an endless source of jokes, queer discussions about improbable situations, and in general for making the atmosphere lighter in the lab. *Diego Royo*, it has been great pleasure to meet you "capo". You have an astonishing calm and positive attitude towards whole life.

I would like to thank the crew at Maxon Computer Gmbh, even if it had been for a few months I remember having a fun time with you. *Matteo, Jonas* I will look forward to another dinner in a fancy restaurant after climbing. *Sebastian*, thank you for welcoming me in your team, even for few months. *Enrico*, Thank you for tutoring me at Maxon and always finding the time to answer my numerous questions.

I would also like to thank few people that I had the luck to met in Zaragoza: *David Morilla*, I have never met anyone that get so genuily excited about many themes, stay near to you has contagious effects on the mood, never lose it! *Lucia and Emi*, thank you for all the endless discussions in front of a beer about life, work and all the dark jokes. *Jorge*, for all the Sunday's breakfasts and all the various discussions about food and the future. *Miguel*, for alwyas sharing interesting discussions about philosophy and the capitalistic system.

Now, to the old (but always present) friends: *Simone and Eleonora*, we have been friends for more than ten years now, which makes it for almost a third of our lifetime. Even if now is harder and harder to be in touch, you always find time to listen to me and share the good and bad times of life, I truly appreciate it. Special thanks for a special friend *Lucione*! Jokes aside I am really glad I can count you as one of my best friend, your brutal honesty mixed with your sense of humor is something that will always make me laugh hard. *Federico*, still glad that we went through the Poli together and that we met there. No matter how much time it passes I always know that it will be always as if we spoke the day before. Thanks to *Domenico, Carla,* and *Roberto* for their stories about faraway countries that always pushed me to explore this vast world and go outside of my house and my comfort zone.
To my brother, Sergio, for being present and for being an amazing brother, always ready to laugh at some dumb joke and always ready to catch up on some very obscure (but fun) film reference. *To my parents*, for whom I do not have enough words, thank you for always believing in me (even when I did not), for the love, and for the skills that you taught me.

This thesis has been funded by the European Union, through the Horizon 2020 research and innovation programme, under the Marie Sklodowska-Curie grant agreement No 956585 (PRIME).

CONTENTS

Part I

INTRODUCTION

1

INTRODUCTION

Simplicity is the ultimate sophistication

— Leonardo da Vinci

In the past decades, the world has experienced a continuous and steady digitalization. The worldwide number of Internet users has almost duplicated in 10 years, from 2.8 billion in 2014 up to 5.3 billion users in 2024 [117], showing a continuous increase in people's use of digital content and digital media. While computer graphics have often been used for entertainment and CAD (Computer Aided Design), the pervasive digitalization of the world might present new opportunities and challenges for this field. Within this transformation, digital twins or replicas have been introduced, as *reliable* digital representations of objects or physical processes, which can offer a wide range of applications. In the context of this thesis we will focus on *visual* digital replicas, i.e., the faithful reproduction of the appearance of a real object.

A visual digital replica can be useful in the manufacturing sector, where technicians can use real-time visualization to monitor the status of ongoing processes, or it can be used to previsualize new products, reducing the costs and time associated with the prototyping phase. Visual digital replicas are also used in retail as well, where most of products are sold online: To attract more clients, companies now provide a *faithful* pre-visualization of their online products, to give a realistic idea of the products that are selling. Beyond digital replicas, computer graphics has valuable applications in scientific visualization and education, or in historical and paleontological reconstructions.

To realistically mimick reality, it is fundamental to simulate not only the light-matter interaction, that happens at the surface level, but also the interactions happening within an object's volume. Materials that account for these internal interactions are called *volumetric materials* and are ubiquitous in our daily lives: From organic materials to inorganic ones, up to aggregates of complex materials (see Figure 1.1).

To efficiently generate volumetric materials, a deeper understanding of how an object's final appearance is generated is essential, this is a challenging problem that involves different scientific disciplines. On one side, it requires insights into the physical processes that characterize how light interacts with matter. On the other side, it requires an understanding of how the light that reaches our retina is interpreted by our visual system, thus involving the fields of physiology and psychology. The correct understanding and modeling of both fields is required in order to have faithful digital reproductions of objects: Without the correct description of the light transport, we would simulate objects that do not look consistently realistic, while without considering the perceptual part we might simulate phenomena that have minor visual impact, thus wasting resources.

Figure 1.1: Examples of volumetric materials, that can be found in our daily-lives. In the top row we show pictures of volumetric materials (from left to right: fish, jade and sand). On the bottom row we show examples of rendered volumetric materials from left to right: skin (courtesy of [114]), textiles (courtesy of [256]) and ice (courtesy of [108]).

1.1 MODELING THE COMPLEXITY OF APPEARANCE

Correctly simulating the light that reaches our eyes, when observing an object under specific lighting conditions, is a challenging task that requires the simulation of the interaction of light with matter at different scales, involving wave, electromagnetic, and quantum effects. As it often happens, the explicit simulation of all these phenomena (i.e., simulating each quantum interaction) using actual technology is unfeasible. To overcome this problem, it is often common to study first the phenomena that happen at the micro-scale. Then, through statistical aggregation, the multiple occurrences of micro-scale phenomena are described, leading to the final macro-scale appearance. This allows us to derive macro-scale appearance, the one that can be observed with bare eyes, from known micro-scale effects in a first-principled manner.

This, however is a challenging task, since there are many phenomena, that happens at different levels (micro or meso) that affect the final macro-scale appearance. Therefore, it becomes important to decompose the final macro-scale into a hierarchy structure of physical phenomena that take place at different levels. Consider the appearance of the Moon: Its surface can be decomposed into multiple levels, starting with mountains and craters at a higher level, boulders and rocks at a middle level, and ending with regolith at a finer level (see Figure 1.2). This layered classification highlights the hierarchical structure of macro-scale appearance, where each level contributes to the overall effect: Altering the finer level, by replacing the Moon's regolith with Mars's regolith, will shift the Moon's color from white to red, while changing the disposition of mountains and craters affect the shadows that is possible to see with a bare eye.

Once the levels of this hierarchical structure have been identified is important to determine which optical phenomena, when statistically aggregated, leads to the same macro-appearance that we observe from the Earth. At higher levels, we can consider the relative masking and shadowing caused by mountains and craters. Similarly, the middle level composed of boulders and rocks, can be modeled through masking and shadowing terms, showing an example of a fractal-like pattern (i.e., the repetition of the same geometry at different scales). However, when reaching the microscopic level of regoliths other effects, such as coherent interference caused by the extreme roughness and compactness of particles, may become relevant. Through statistical aggregation, we aim to model all the multi-scale effects that generate the emerging appearance. In the case of the Moon's surface several models have been proposed to efficiently account for the aggregated appearance without the need to explicitly account for the finer levels of detail [80, 81, 171].

Depending on the required precision of the simulation, it might become important to not only characterize the surface alone of the object but also its internal structure. This has a major visual impact when simulating structures such as organic tissues like human skin, leaves, or fur when viewed closely. A convenient approximation of this volume-like aggregate behavior is to assume that the light-matter interaction inside the material can be modeled statistically, as a set of suspended independent scatterers, via the Radiative Equation, which we describe in more depth later (see Section 2.1.4). The suspended scatterers encode the ensemble average optical properties (e.g., properties that define how much dense, gloss, or transparent a material is) of the materials that is characterized and are key to describing volumetric materials with very different appearances. Simulating the light transport in volumes is computationally expensive, therefore, methods have been devised to volumetrize surfaces as well, aiming to simulate the scattering of light inside the object without the need to explicitly instantiate volumes. This often revolves around assuming that the underneath medium can be considered infinite in one or more directions [111, 260], or that the lateral displacement caused by the internal scattering is so small that can be ignored [76]. These hypotheses, in return, allow to reduce the dimensionality required to calculate the internal scattering of light from a three-dimensional volume to two or one-dimensional, simplifying the required computational time. In this thesis, we investigate the use of such volumetric surfaces for efficiently simulating foundation cosmetics in Chapter 3 and for describing the scattering that happens inside pennaceous feathers in Chapter 4. While the correct modeling of the light transport is a requisite to generate believable models, being aware of the part that perception plays in the final appearance is key to efficiently modeling and editing these materials.

1.2 PERCEPTUALLY-BASED AUTHORING TOOLS

While physically-based models are required for obtaining realistic images, understanding the processes within the human brain is equally important. Such knowledge can help determine when complex, computationally intensive models are necessary and when simpler, less demanding models can be effectively employed. An example of this can be viewed in the application involving volumetric materials: While computing the

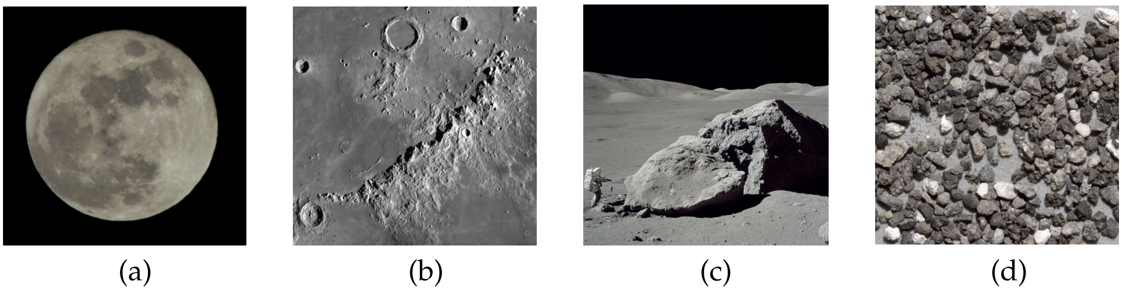

Figure 1.2: Different scales of appearance. (a) Moon viewed from a telescope. (b) Mountains and craters on the lunar surface. (c) Boulder on the Moon's surface. (d) Regoliths present on the moon's surface. The final appearance of (a) is caused by the aggregated effects of (b),(c) and (d).

light transport that happens beneath a surface is more physically based, many surface models are still used. The reason is that for opaque or metallic materials, most of the energy that a light beam carries beneath the surface will be absorbed quickly. Therefore, the final appearance of these materials can be well approximated by considering the few scattering events that happen close to the surface levels, thus making surface-based models a valid choice. Describing the light transport using volumes would be more physically correct, but less efficient in terms of computational power, and with little *perceived* benefits. Perceptual models have also applications in the manufacturing world, where they can be employed in quality control of screen productions [145, 224], among other fields.

Understanding the mechanism of human perception can have applications also when *editing* volumetric materials. Physically-based models often rely on parameters (i.e., optical properties) derived from micro-phenomena, which can produce unpredictable final appearances. This is due to the non-linear behavior resulting from the repeated occurrences of the same micro-phenomena. Moreover, many parameters are interdependent, leading to similar final appearances obtained through different optical parameters combinations. Finally, the number of optical parameters is often relativily high, thus increasing the number of axis (or dimensions) that needs to be explored during the editing phase. These three issues, non-linear behavior, paramter interplay, and number of dimensions, creates a barrier to the widespread diffusion of appearance models to the masses, and only specialized professionals can specify appearance in a way that is plausible for tasks such as previsualization and predictions. A solution to these challenges is the research of alternative editing spaces: Instead of the conventional bottom-up approach, where users have to fine-tune material properties to achieve a desired final appearance, these novel editing spaces should follow a top-down methodology. This approach allows users to specify the desired final appearance directly, from which one or more corresponding sets of optical properties should be derived automatically.

Additionally, while the optical properties of a material heavily influence the appearance of an object, they are not the only contributors. Confounding factors, i.e., factors that can alter the perception of an object but are not inherently related to the optical properties of the material (e.g., the object's shape or the lighting condition), can further hinder the editing task. For example, if a material that has been carefully edited appears drastically different when viewed under a different light condition or applied to a different shape, than it needs to be edited again. We delve more into the topic of confounding factors in Section 2.2.

The existence of confounding factors show also a limitation of working using optical parameters, that can be measured and quantified, but do not account for external factors such as lighting or geometry. To overcome such limitation, an alternative paradigm is to work in image space which better captures perceptual changes that could not be otherwise registered working solely in optical parameters space. Both issues related to the editing task, the difficulty of defining the final macro-behavior in the space of optical parameters and the difficulty caused by confounding factors, should be considered in the development of novel, intuitive, authoring tools. In this thesis, we investigate aspects of perception focusing on the editing task of translucent materials, a class of materials that can be modeled as volumetric materials (see Section 2.1.5). Specifically, we investigate if light motion can be used to mitigate the confounding effect of light position in Chapter 5, as well as analyze the performance of similarity metrics for translucent materials in Chapter 6. Successively, we investigate over the editing of translucent materials. We employ the best-performing similarity metrics to develop an intuitive manifold of translucent appearances over which we build a novel interface that can facilitate the editing task of translucent materials.

1.3 GOAL AND OVERVIEW

Due to the wide range of materials that is possible to describe, this thesis focuses on the advances in the digital representation of volumetric materials. As an object's appearance is the result of both light transport and human perception, in this dissertation we investigate both fields, with the goal of improving the future digital representation of volumetric materials.

- Part II is devoted to characterizing new reflectance models involving volumetric materials. Chapter 3, inspired by scanning electron microscope (SEM) images, propose to model the internal volume as an ensemble of flake-like or diffuse scatterers. Such modeling, allows us to obtain dewy and matte appearances, along with other known effects, such as the increase in reflectance when stacking multiple layers of cosmetics, which we validate through objective and subjective data. Chapter 4, on the other side, investigates the use of volumetric materials to model the scattering of barb and barbules, two types of fibers that are fundamental to model pennaceous feathers. Since we model both types of fibers as volumetrized primitives (i.e., we model the internal scattering of infinitely long cylinders), we generate an analytical model that efficiently accounts for the relative occlusion between barb and barbules.
- Part III is devoted to the perception and editing of translucent materials. Chapter 5 investigate the effects of light motion on the perception of translucent materials. In Chapter 6 we leverage the limits of the human visual system to develop a perceptually-based manifold of translucent appearances. We then employ such

manifold in a user study case finding that it leads to better qualitative and objective improvements than traditional editing approaches.

While I am the leading first author in most of the presented works, some have been carried out in collaboration with other researchers. Therefore, at the beginning of each chapter, a section called "About this chapter" briefly introduces and contextualizes the work, lists its technical contribution, and states my contribution to the project.

1.4 CONTRIBUTIONS AND MEASURABLE RESULTS

In this section, we list all the publications and work done during this thesis. Specifically, this thesis led to three publications in journals (Computer Graphics Forum), one publication in a conference (Symposium on Applied Perception), and two posters at SIGGRAPH.

- Chapter 3: Practical Appearance Model for Foundation Cosmetics
  - Dario Lanza, Juan Raúl Padrón-Griffe, Alina Pranovich, Adolfo Muñoz, Jeppe Revall Frisvad, Adrián Jarabo.
  - The paper has been published at Computer Graphics Forum (IF 2.7, JCR 44/131 Q2, Computer Science, Software Engineering). The paper has been presented at the Eurographics Symposium on Rendering (EGSR) Conference 2024.
- Chapter 4: A Surface-based Appearance Model for Pennaceous Feathers
  - Juan Raúl Padrón-Griffe, Dario Lanza, Adrián Jarabo, Adolfo Muñoz.
  - The paper has been published at Computer Graphics Forum (IF 2.7, JCR 44/131 Q2, Computer Science, Software Engineering). The paper has been presented at the the Pacific Conference on Computer Graphics and Applications (Pacific Graphics 2024).
  - The project was presented as a poster at SIGGRAPH 2024.
- Chapter 5: On the Influence of Dynamic Illumination in the Perception of Translucency
  - Dario Lanza, Adrián Jarabo and Belen Masia.
  - This paper has been accepted at the Symposium on Applied Perception (SAP 2022).
  - The project was presented as a poster at SIGGRAPH 2022 and won third place in the Student Research Competition (SRC) - Graduate Students.
- Chapter 6: Navigating the Manifold of Translucent Appearance
  - Dario Lanza, Belen Masia and Adrián Jarabo.
  - The paper has been published at Computer Graphics Forum (IF 2.7, JCR 44/131 Q2, Computer Science, Software Engineering). The paper has been presented at the EUROGRAPHICS Conference 2024.

1.4.1 *Research Internships*

Two research internships, for a total of four months, have been carried out during this PhD thesis:

- June - July 2022: two months research stay at Denmark Technical University. Host: Prof. Jeppe Revall Frisvad. This internship has led to the development of work presented in Chapter Chapter 3.
- January - February 2024: two months research stay at Maxon Computer Gmbh Hosts: Dr. Enrico Rinaldi and Sebastian Häfele.
- During my stay I had the opportunity to give an internal technical talk on Landscape Erosion Simulation in Cinema 4D to the engineering department of Maxon Computer GmbH.

1.4.2 *Outeach and Reviews*

During my PhD I had the opportunity to also give back to the commounity:

- I have been reviewer for the Eurographics conference, EAI ArtsIT conference and for the Perception Journal.
- I have participated to the SIGGRAPH Podcast SRC Student Research Competition winners.
- I have helped in writing one chaprter of the non-peer reviewed Unidigital IASAC project [178].

2

BACKGROUND AND RELATED WORK

In this chapter, we introduce basic concepts that are used throughout the rest of the thesis. In Section 2.1 we review mathematical and physical concepts relevant to rendering and appearance modeling in the context of this thesis. Section 2.2, then, introduces basic concepts of visual perception, discusses perceptual material constancy, and reviews previous work on the perception of translucent materials, over which we build the work presented in Part III.

2.1 PHYSICAL BACKGROUND

Visual appearance is a complex phenomenon that involves the interaction between light and matter at different scales. Depending on the scale at which we would like to characterize the behavior of light we can distinguish between three macro-areas: geometric, wave, and quantum optics. If we consider the interaction between light and matter on a scale larger than the wavelength of visible light (between 380 and 700 nm), geometric optics allows the modeling of many effects such as reflection, absorption, refraction and emission. Under these assumptions, light can be efficiently characterized as an infinitesimally small beacon of ray (sometimes called pencil ray) that propagates in a straight line in a medium that has a constant index of refraction.

Simulating light transport with geometric optics allows to model many of the most salient features that we normally perceive in our daily lives, with a comparatively low computational cost, making geometric optics the standard framework in computer graphics. Although efficient, geometric optics is not capable of modeling many of the small subtleties present in the world. To overcome this gap, it is key to incorporate wave-based phenomena.

When the scale of the interaction between matter and light is comparable to the wavelength of visible light, the light-matter interaction can be described using wave optics. At such small scale, most of the natural light is coherent, thus allowing the creation of interference or diffraction effects. Although there are many solvers capable of simulating these interactions, they are often computationally expensive and do not scale well with the number of elements in the scene. Recent work [198–200] aims to characterize wave optics using the geometric optics formalism, but the computation time required by these solutions is still high. Therefore, it is more common to describe the wave nature of light with ad-hoc solutions, aiming to simulate effects such as those caused by thin-film diffraction [50, 70, 74, 193], or by the structure of the microgeometry [43, 96, 231, 238, 240].

When the scale of interaction is less than five nanometers [97], the quantized nature of light and matter should be considered using quantum electrodynamics. Quantum optics, which is derived from quantum electrodynamics, aims to model such discrete nature, allowing to characterize and predict effects such as as phosphorescence, fluorescence, or blackbody radiation. These phenomena arise from changes in the quantum state of matter and as such they cannot be correctly explained by classical wave optics alone. These three phenomena generate visual effects that can be commonly found in our daily lives. Although the model of blackbody radiation is an abstract one, it is often used to describe the radiance distribution emitted by an object heated at a specific temperature (e.g., metals that are being heated or the Sun). On the other hand, fluorescence and phosphorescence are visual effects that can be commonly found in our daily lives (e.g., chlorophyll, among many other organic substances, shows a degree of fluorescence; while phosphorescence is more common in deep-sea animals). An explicit simulation of the quantum effects that explain these three phenomena is too expensive for the time budgets of standard computer graphics applications. Still, it is possible to efficiently reproduce the visual effects of these phenomena, using spectral renderers (i.e., render engines that account for all the wavelengths in the visible range during light transport, rather than RGB renderers that consider only three specific wavelengths). Blackbody radiation can be simulated through a closed-formula expression [191], which has been extended to non-blackbody [232], while fluorescence and phosphorescence are simulated through the so called re-radiation functions [69, 77, 109], that describes the absorption and re-emission of light at either a different wavelength (fluorescence) or at a different time (phosphorescence). The key idea of re-radiation functions is to describe the relationship between incoming and outgoing light vectors, treating these phenomena as "black boxes".

In this dissertation, we will work under the geometric optics assumption, since it offers a good trade-off between computation time and accuracy of the visual effects that we are investigating.

2.1.1 *Radiometric magnitudes*

RADIANT FLUX The radiance ϕ, or radiant power, is defined as

$$\Phi = \frac{dQ}{dt}, \tag{2.1}$$

where dQ is the infinitesimal energy (expressed in joules) that a ray beam carries over a portion of differential time dt. The radiant flux is expressed in Watts (W) defined as one joule per second ($[Js^{-1}]$).

IRRADIANCE The irradiance is defined as the radiant flux that passes through a surface per unit area. More formally, irradiance E is defined as

$$E = \frac{d\Phi}{dA}, \tag{2.2}$$

where dA is the infinitesimal surface over which the flux is passing. Irradiance is measured in $[Wm^{-2}]$.

RADIANCE While irradiance accounts for the radiant flux that passes through a surface, in all directions, it lacks any notion of directionality. To overcome this problem it is useful to define the radiant flux L over surface area per unit solid angle as

$$L = \frac{d\Phi}{d\omega dA^{\perp}}, \tag{2.3}$$

where $dA^{\perp}$ describes the projection of the infinitesimal surface dA over the differential solid angle $d\omega$. Radiance is measured in Watts per square meter per steradian ($[W sr^{-1} m^{-2}]$).

2.1.2 *The Scattering Function*

When light hits matter its interaction is defined by the optical properties of the material, which might absorb or scatter light. Such behavior is characterized by the bidirectional scattering distribution function (BSDF), defined as:

$$f_o\left(x_i, \omega_i, \omega_o\right) = \frac{dL_o\left(x_i, \omega_o\right)}{dE_i\left(x_i, \omega_i\right)}, \tag{2.4}$$

where L_o is the radiance observed at point x_i from direction ω_o and E_i is the infinitesimal irradiance that is observed at point x_i when the infinitesimal area dA is lit from direction ω_i. The BSDF is measured in $[sr^{-1}]$. Intuitively, BSDFs are a general way of describing the ratio between reflected and received light, with no further insights. In this sense, BSDFs can be seen as black boxes that allow the reproduction of known measurements, but with no intuition into the physical mechanisms that are taking place at the microscopic (or nanoscopic) level, thus limiting their capability to extrapolate to novel conditions (e.g., novel light conditions or shape).

2.1.3 *The Scattering Equation*

By solving Equation (2.4) for the outgoing radiance dL_o we obtain

$$dL_o\left(x_i, \omega_o\right) = f_o\left(x_i, \omega_i, \omega_o\right) dE \quad = f_o\left(x_i, \omega_i, \omega_o\right) dL_i\left(x_i, \omega_i\right) |\cos\theta_i| d\omega_i, \tag{2.5}$$

where we made use of the fact that the irradiance E can be described in terms of the radiance L by noting that $dA^{\perp} = dA|\cos\theta_i|$, where θ_i is the azimuthal angle between the normal of the infinitesimal surface dA and the direction of projection, $d\omega$.

Finally, we can integrate Equation (2.5) over the unitary spherical domain of directions Ω, obtaining

$$L_o\left(x_i, \omega_o\right) = \int_{\Omega} f_o\left(x_i, \omega_i, \omega_o\right) L_i\left(x_i, \omega_i\right) |\cos\theta_i| \, d\omega_i. \tag{2.6}$$

Equation (2.6) is also known as the rendering equation [115]. The rendering equation is used in computer graphics to calculate the radiance that reaches an infinitesimally small point in space x_i with direction ω_o.

2.1.4 *The Radiative Transfer Equation*

While Equation (2.4) allows us to model the scattering behavior of different materials, it only captures what happens on the surface level, with no insights into the physical processes that are taking place beneath the surface. Similarly, Equation (2.6) describes the light transport between macro surfaces, ignoring how light is affected by particles dispersed in a volume (e.g., water droplets dispersed in the air generate clouds). To overcome both problems, it is useful to describe the light transport inside a volume filled with a medium, characterized by particles with specific optical properties. Such characterization allows us to describe both the interaction between light and particles dispersed inside a portion of space (e.g., clouds) and the light transport that happens beneath the surface of an object since it can be also modeled as a collection of particles dispersed inside a volume. Explicitly describing all the particles inside a volume can quickly become inefficient (a single mol of an element contains 6.02210^{23} particles); therefore, it is more efficient to model the medium as a volume filled with scatterers, particles that encode the *average* optical properties of the constituent particles. Assuming that particles are independently distributed, i.e., they do not form clusters, nor that they emit energy, the behavior of light can described by the Radiative Transfer Equation (RTE) [23]:

$$\omega_o \cdot \nabla L(x, \omega_o) = -\sigma_t L(x, \omega_o) + \sigma_s \int_{\Omega} f_p(\mu) L(x, \omega') d\omega', \tag{2.7}$$

which describes the infinitesimal change in radiance at point x viewed from direction ω_o inside the volume filled with scatterers. The first term, $-\sigma_t L(x, \omega_o)$, describes how radiance is attenuated by the extinction coefficient σ_t, defined as $\sigma_t = \sigma_s + \sigma_a$, where σ_s is the scattering coefficient and represents the differential probability of a light beam being scattered from its initial direction per unit distance traversed inside the medium and σ_a, the absorption coefficient, that describes the probability density of a light beam to be fully absorbed in unit distance traversed inside the medium. The second term of Equation (2.7) accounts for all the radiance that can be scattered towards x with direction $-\omega_o$ that happens during the traversal of the infinitesimal volume. This additional incoming radiance is defined as the product between the infinitesimal scattering probability, σ_s, and the total incoming extra radiance defined as the spherical integral in the solid angle domain of the phase function, $f_p(\mu)$, and the radiance observed at position x viewed from ω', the integration variable. The phase function describes the directional behavior of scattering $f_p(\mu)$ [sr^{-1}], with $\mu = \omega_o \cdot \omega'$ and $\cdot \cdot$. the dot product. Similar to the definition of BSDF in Equation (2.4), the phase function is defined as

$$f_p(\omega_i, \omega_o) = \frac{d\Phi_s(\omega_o)}{d\Phi_i(\omega_i)} \tag{2.8}$$

where $\Phi_i(\omega)$ is the incoming radiant flux with direction ω over the infinitesimal volume dV and Φ_s is the scattered flux in direction ω_o over the the infinitesimal volume dV. Both Equation (2.4) and Equation (2.8) describe similar concepts, and in fact they are both measured in sr^{-1}. This is no coincidence, as both Equation (2.7) and Equation (2.6) describe the transport of light. In the case of Equation (2.6) the light transport only accounts for surfaces using the definition of BSDF given in Equation (2.4). In the case of Equation (2.7) the light transport accounts for the presence of scatterers inside the volume that is traversed, accounting for the directional behavior of scatterers through the phase function (see Equation (2.8)). Additionally, Equation (2.6) describes the boundary conditions for the RTE, which in this case are the surfaces that delimit the volume considered by the RTE.

As can be seen from Equation (2.7), the radiance of a light beam that passes through a medium is influenced by three parameters: σ_t, σ_s and f_p, which are also known as *bulk optical parameters*, sometimes called optical parameters. The bulk optical parameters represent the ensemble average optical properties of matter in the differential volume.

Equation (2.7), assumes that the σ_t and σ_s are independent of direction, as well as that the phase function only depends on the relative angle between ω and ω'. The formulation of Equation (2.7) might not be sufficient to express materials that show a certain anisotropy, such as clothes, thus requiring a more general formulation that also accounts for the average particle geometry; following [103] we can define:

$$\omega_o \cdot \nabla L(\omega_o) = -\sigma_t(\omega_o) L(\omega_o) + \sigma_s(\omega_o) \int_{\Omega} f_p(\omega_o \to \omega') L_i(\omega)\, d\omega', \tag{2.9}$$

with $\sigma_t(\omega_o)$ and $\sigma_s(\omega_o)$ the direction-dependent extinction and scattering coefficients and $f_p(\omega_o \to \omega)$ the directional-dependent phase function, in which we make explicit, through the $\to$ symbol, the relationship between the ω_o and ω', and not only their dot product as in Equation (2.7). This, however, does not impose any symmetry over the phase function, making it potentially not reciprocal, meaning that $f_p(\omega_o \to \omega) \neq f_p(\omega \to \omega_o)$; reciprocity can be achieved by imposing that $f_p(\omega_o \to \omega)\sigma(\omega_o) = f_p(\omega \to \omega_o)\sigma(\omega)$.

In this thesis, we will often use two coefficients that are related to σ_t. The first one is the single-scattering albedo coefficient defined as $\alpha = \frac{\sigma_s}{\sigma_t}$, which describes the probability, between 0 and 1, of scattering (instead of absorbing) at scattering events. The second coefficient is the mean-free path *mfp* defined as $mfp = \frac{1}{\sigma_t}$, which describes the average distance that a light beam might traverse inside the medium without incurring in any particle.
As a final remark, it is interesting to note that although Equation (2.7) has been initially phenomenologically derived [23], only 40 years later has been derived from first-principles [159].

2.1.5 *Translucent Materials in Computer Graphics*

By describing the scattering of light inside volumes, Equation (2.7) is more general than Equation (2.6), since it allows the simulation of the entire light transport inside an object, up to the point that volumes can be used to model surface models [39, 88]. This becomes particularly important to model light-permeable materials, also known as translucent materials, such as skin, food, or jade (see Figure 2.1, left, for examples of translucent materials).

From a perceptual point of view, this category of materials lies in between completely opaque and transparent objects. In computer graphics, it is common to describe an object made of a translucent material as an enclosing surface filled with a volumetric medium. This simplification allows the decoupling of refraction and reflection on the surface from the multiple scattering events inside the volume. For the rest of this dissertation, the medium inside will be characterized through three types of coefficients: the attenuation coefficient σ_t, the single scattering albedo α, and the phase function f_p. To simulate the surface, sometimes called dielectric interface, we will use Smith's microfacet [192, 210] model following the GGX distribution of microfacets. We will characterize the surface interface through ρ, the roughness of the microfacets, and η, the index of refraction. This approach efficiently decouples the refraction and reflection of light using the dielectric interface, and the internal scattering, as we show in Figure 2.1, right. In this dissertation, we will sometimes refer to *optically thick* and *optically thin* translucent materials; with these terms, we intend to distinguish between translucent materials that are either particularly dense, with a comparatively high average number of scattering events required to fully traverse the medium, and translucent materials that are particularly thin (i.e., with very few scattering events required to fully traverse the medium). Often optically thick materials end up generating appearances that share more features in common with opaque materials (i.e., background not directly visible), while optically thin materials share visual features with transparent materials (i.e., background is visible although blurred, strong presence of caustics).

2.2 MATERIAL PERCEPTION AND CONSTANCY

Being able to fully describe the way in which the human brain visually perceives and processes what surrounds us is still an open problem. In this section, we will briefly discuss some basic concepts of visual appearance perception that will be later used in Part III. When discussing perceptual studies, it is common to distinguish between distal stimulus–the *real* object that is the focus of our attention–and the proximal stimulus–the input that reaches our sensors–. In the case of vision, the proximal stimulus is the radiance that reaches the cones and rods in our retina.

Despite some failing cases, our visual system can robustly recognize the same distal stimuli even when the proximal stimuli change. As an example, consider a sphere made of metal: we can perceive the object as being made of metal independently of whether we observe it at sunset or when the sun is at the zenith and yet, the radiance distribution that reaches our retina is very different in the two cases. The ability of our system to overcome these changes and still be able to recognize is called perceptual constancy [41, 218]. Specifically, in this thesis, we will discuss *material* constancy, which is the capability of our system to consistently recognize (or not) the same material when the proximal stimulus is changed.

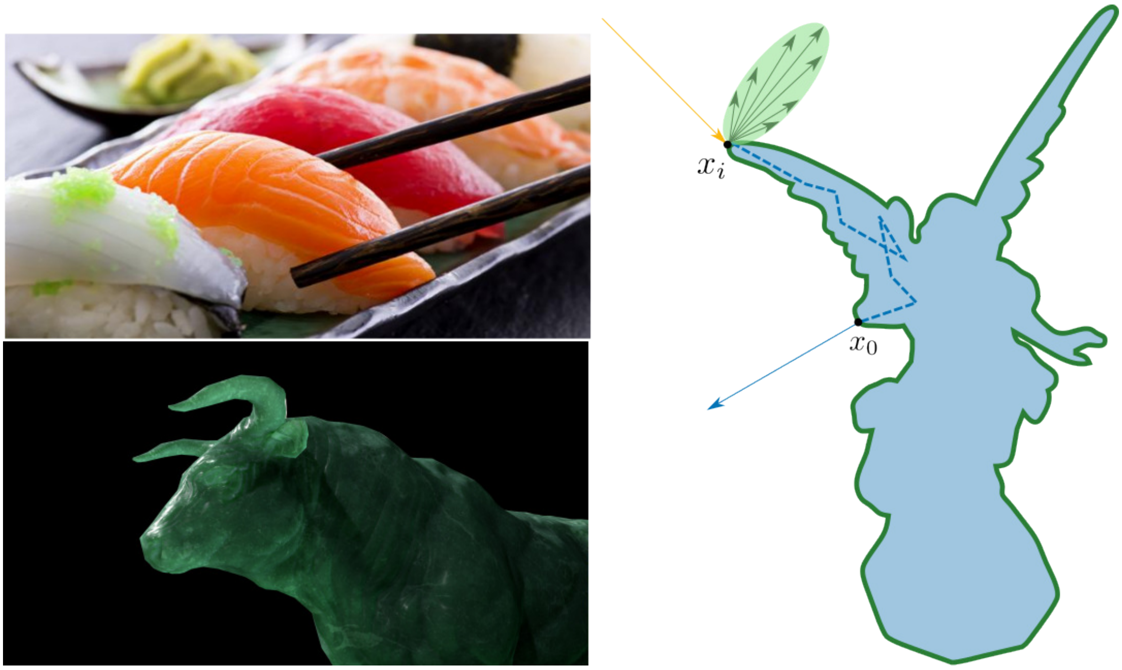

Figure 2.1: *Left*: Due to their internal scattering, translucent materials like fish or jade are characterized by a "glowing" effect. *Right*: Schematic of how translucent materials are simulated in computer graphics. Part of the incoming light is reflected back at point x_i (green lobe), or refracted inside (blue path). The refracted path gets scattered multiple times until it finally emerges back to the surface at point x_o.

Achieving material constancy is not a trivial task; the proximal stimulus that reaches our eye is a combination of multiple factors: recalling the example from before, the proximal stimulus is a combination of shape (the spherical shape of the object), the optical properties (the fact that the object is made of metal), the light conditions (whether we are observing the object at sunset or at zenith time) and the viewing conditions. The factors that can alter our perception of the material without being linked to the optical properties of the material itself are called *confounding factors*. In this context, several works have discussed the effect of various confounding factors, such as geometry and illumination [124, 185, 214], motion [146, 213], or distance [45, 110]. As an example, see Figure 2.2 for examples of successes and failures of material constancy.

Traditionally, the main theory that had been proposed on how material constancy is achieved assumed that the human visual system solved an inverse physics problem, essentially estimating the optical properties of the material, to solve the material recognition task [156]. Currently, a widely accepted theory asserts that our visual system relies on extracting meaningful features from the proximal stimulus. During our lifetime, our brain learns how to correctly associate the features extracted from the proximal stimulus with the distal stimulus [47]. This theory has reached now a certain consensus [22] through different psychophysical experiments. Further support has also recently emerged through the use of Deep Neural Networks (DNN). An example in this direction is the work of Storr et al [201], in which the authors show how training a neural network on simple images of glossy heightfields not only achieves human

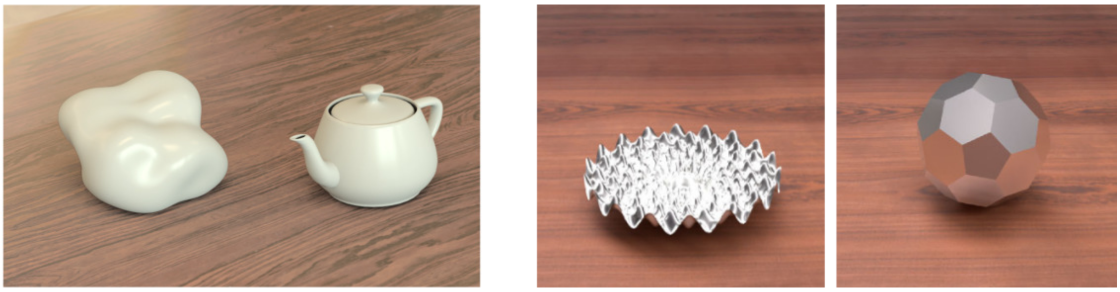

Figure 2.2: *Left*: Example of material constancy. Even if the shapes are different the two objects are perceived to be made of the same material (courtesy of [214]). *Right*: Example of material constancy failure. Both objects are made of the same material lit under the same light conditions. However, the object on the left is perceived as brighter than the one on the left (courtesy of [124]).

performance in glossy estimation but also leads to similar downfalls, suggesting, that humans, similarly to DNN, learn from patterns.

If our visual system relies on features extracted from the proximal stimulus, it becomes clear that material properties alone are not sufficient to describe how an object will be perceived. Consequently, the space of optical properties cannot solely resolve the final appearance of an object; confounding factors also need to be accounted for. This has a major impact on material editing, required, for instance, to generate reliable product pre-visualizations. This challenge arises because the parameters exposed by the models are physically based, but ignore the effects of confounding factors. As a result, users must further refine the optical parameters to account for the effects of confounding factors to achieve the desired appearance.

A solution to overcome this is to generate new authoring tools that take into account not only the material properties but also the final perceived appearance. This often resolves in authoring tools that work in image space [32, 123, 186, 188, 204, 233], as it already expresses the interplay between confounding factors and optical properties, rather than working in the optical parameter space.

Moreover, editing directly using the final perceived appearance allows users to avoid tweaking low-level optical parameters, which can be a daunting task. In this thesis, we investigate the research of novel editing spaces for translucent materials. In Chapter 6 we develop a manifold of translucent appearances working in image space, allowing us to account for perceptual distances rather than optical ones. However, we also maintain the link between appearance and optical properties, through inverse rendering, allowing us to develop a novel editing interface for editing translucent materials.

2.2.1 *Perception of Translucent Materials*

In this section, we briefly revise previous work on the perception of translucency, over which we build the work presented in Part III; for a more comprehensive review, we refer the reader to the work of Gigilashvili et al. [60]. Translucent materials allow light to pass through them and get scattered in multiple directions. Such internal scattering of light redistributes light in areas that would be otherwise darker [65], creating a characteristic glow. Traditionally, the study of the perception of translucent materials had always been associated with the study of transparent stimuli [156], accepting that the human visual system (HVS), categorized material based on inverse optics. With the advancement of Computer Graphics technology in the early 2000s, realistic rendering of translucent materials became possible, allowing Fleming and colleagues [49] to study the perception of translucent materials using computer-generated stimuli. Through the results of this study, Fleming and colleagues propose a new theory that claims that the HVS extracts some low-level statistics from the proximal stimulus, and then learns the association between the low-level statics and the material's property. Other authors have proposed the local contrast of the non-specular zone as a possible low-level cue used by the HVS to infer properties of translucent materials [161]. More recently, the work of Marlow et al. [147, 148] suggested that the covariance between surface orientation and light intensity is used as a clue in translucency perception. This is in accordance with the work of Kiyokawa et al. [122] that studies the relationship between highlights regions and diffuse regions in translucency perceptions.

These works correlate with the fact that *geometry*, can also alter the perception of translucent materials, with objects with smooth edges being perceived as more translucent than those with sharp edges, as shown by Xiao et al. [240]. The reverse has been also found true: Bumpy surfaces might look less bumpy depending on how much light is scattered inside the object [29].

These findings relate to the fact that we use the intensity gradient to estimate surface curvature: If we have a smooth gradient between a bright and a dark area, we usually assume that the geometry is smooth. While this is generally true for smooth opaque objects, for translucent materials such association between gradient and geometry does not hold particularly well. The internal surface scattering causes a non-linear re-distribution of the energy inside the object [66] that breaks this relationship. It is likely that when facing front-lit translucent materials we tend to reconstruct the geometry from surface highlights, as they can give a more reliable cue on surface orientation, and then use other metrics or features to assess material's property [147].

Highlights and diffuse regions are not the only areas used for assessing translucency: thin areas (i.e., areas which are particularly thin and therefore scatter less the light) tend to be observed by users when assessing the level of translucency of objects [62, 63]. Similar results are explored by Chowdhury and colleagues [29], who report that removing silhouette edges from the stimuli alters the observers' perception.

Lighting conditions have also been proven to be another confounding factor when observing translucency [239]: Objects that were lit from behind appear optically thinner, while the same object frontally lit is perceived to be more optically thicker. These results, however, were dependent on the phase functions of the scattering material: As expected, light direction was more relevant in forward scattering media where light penetrates deeper in the material while having little effect in isotropic media. Behavioral experiments point out how humans prefer to put the object between the light source and the eye when assessing the translucent nature of an object [59].

The absence of *chroma* can also alter the perception of translucency [135], although it is not sufficient to elicit a sense of translucency in materials [49].

Translucency perception is a complex phenomenon that is far from being completely understood. There are numerous studies that point out how translucency assessment is also entangled with the three-dimensional estimation of the object [122, 147, 148], however, these studies have focused on the perception of translucent materials that are optically thick (i.e., it is not possible to see through the object) when frontally lit. However, it is not clear if this also applies to objects that are backlit, or for optically thin materials (i.e., it is possible to see a blurred background through the object). It is likely that the HVS relies on different strategies depending on the lighting conditions (whether the object is lit frontally or from behind) and depending on the optical parameters (whether the object is optically thin or optically thick). In this context, the work of Xiao et al [239] shows that light conditions can lead to the same material being perceived differently. We investigate in a similar direction by studying if dynamic lighting can alleviate such lack of constancy in Chapter 5.

Part II

REFLECTANCE MODELS

In this part we delve into appearance modeling by developing novel reflectance models. The first chapter of this part, Chapter 3, is focused on the reproduction of foundation cosmetics. The main contributions are the development of a new reflectance capable of reproducing different ranges of appearance, and its validation, both visually and through ground-truth data. The second chapter of this part Chapter 4, aimg to simulate the far-field appearance of feathers. The main contributions are accounting for characteristics specific of feathers, such as non-cylindrical cross-sections, and the development of a new masking term that accounts for the relative occlusion of barbs and barbules, two key components of feathers.

3

A PRACTICAL APPEARANCE MODEL FOR COSMETICS FOUNDATIONS

ABOUT THIS CHAPTER

In this work, we introduce a novel reflectance model for foundation cosmetics that replicates a wide range of appearances—from glossy to matte to velvety—using a multilayered BSDF designed to simulate the stacking of multiple cosmetic layers. We take inspiration from microscopic particulates used in cosmetics, and characterize each individual layer as a stochastic participating medium with two types of scatterers that mimic the most prominent visual features of cosmetics: spherical diffusers, resulting in a uniform distribution of radiance; and platelets, responsible for the glossy look of certain cosmetics. We validate our model against measured reflectance data, and demonstrate the versatility and expressiveness of our model by thoroughly exploring the range of appearances that it can produce. As the leading author, I have led the modeling and implementation of the reflectance model, worked on the virtual scenes used in the various experiments in the paper, and contributed to the optimization phase. I have also led the writing of the manuscript. The work introduced in this chapter has been accepted to Computer Graphics Forum and presented at the Eurographics 2024 conference.

D. Lanza, J.R., Padrón, A. Pranovich, A. Muñoz, J.R., Frisvad & A. Jarabo
A Practical Appearance Model For Cosmetics Foundations

Computer Graphics Forum (EGSR 2024)

3.1 INTRODUCTION

Since ancient times, decorative cosmetics have been used in virtually all existing cultures around the globe. From the prehistoric mineral-based pigments [227] to modern sophisticated chemical makeup, cosmetics have been used for enhancing appearance and hiding imperfections, for ritual painting, theatrical purposes, or most recently, visual effects. However, despite their ubiquity, cosmetics have not been explored much in graphics, most likely because they are generally baked into the look-development process for virtual humans. Rendering of cosmetics has received little attention, limited to image-space [166, 242, 245] or texture-space manipulation [184, 248], or simplistic physical models [93, 131, 132]. Existing models either directly manipulate an image or describe how cosmetics modify the parameters of conventional analytic shading models. Our approach models the light-scattering particles in the cosmetic layer and computes the effect of such a layer using position-free Monte Carlo simulation.

Figure 3.1: We propose a practical appearance model for foundation cosmetics that can be stacked on top of human skin.*Left:* Rendering of a white female character skin without makeup. *Middle:* We apply a *foundation* layer with a matte finish, which reduces the specular highlights, especially visible in the eye region and next to the nose, and slightly changes the skin color. *Right:* We apply an additional *shinier* layer adding a reddish tint to the cheek.

Several types of cosmetics exist with different areas of application and intended goals. In the case of facial makeup, one or more relatively thin layers of cosmetics are applied on top of the skin, resulting in a multilayered structure. Facial makeup cosmetics, including *foundation*, *concealers*, *rougers* (*blusher*), *bronzers* and *highlighters*, consist of a combination of microscopic colored diffusers and platelets, either suspended in a water-based host medium, or presented as powder that sticks to the outermost sebum layer of the skin. At a macroscopic level, the diffusers and platelets interact with light as a scattering medium.

In this work, we focus on the base layer or *foundation*. A foundation aims to provide a uniform skin hue over the whole face upon which other makeup layers can be stacked. A key visual attribute of a foundation layer is its *finish*: Different foundation products exhibit distinct visual characteristics ranging from a *matte* finish that results in a non-shiny, velvety appearance; to a *dewy* finish that aims for a more natural appearance making the skin brighter and shinier to evoke an impression of a healthy skin.

Based on the observed types of appearance existing in foundation cosmetics, we propose an intuitive empirical appearance model for foundation makeup, where we model each cosmetic layer as a finite volumetric scattering layer, and where the bulk scattering properties are derived from a mixture of spherical-like and platelet-like scatterers. This enables us to stack an arbitrary number of makeup layers, resulting in a single bidirectional scattering distribution function (BSDF). We implement our model using the position-free Monte Carlo formalism [76], allowing for multiple scattering and arbitrary distributions.

Our model is able to reproduce a wide variety of appearances, from dewy to matte. We compare our model against measurements of real-world cosmetics, and show that our model closely fits such measurements, while reproducing the main visual features of cosmetics when applied on top of digital human skin. In addition, we show that our model can be used to model other types of cosmetics (e.g., blusher) and enables the stacking of multiple makeup layers as exemplified in Figure 3.1.

3.2 RELATED WORK

VOLUMETRIC MATERIALS. As discussed in Section 2.1.4 light scattering in a volume can be modeled using Equation (2.7) [23], or its variants for anisotropic (Equation (2.9)) [104] or correlated media [12, 108]. Based on this theoretical framework, numerous appearance models have been proposed for materials such as skin [3, 37, 100, 195], paper [175], cloth [254, 256], leaves [6], special pigments [74], ice [53], wood [137, 150], or granular media [154, 163]. Our work follows this line of work and models foundation makeup as a combination of isotropic and anisotropic scatterers inside a scattering medium.

MULTILAYERED MATERIALS. Early work on rendering of multilayered materials simulated the reflection and transmission by the multilayered structure of skin and leaves using subsurface scattering [36, 79], being limited to diffuse reflectance. Later work proposed using the adding-doubling method on spatio-angular scattering representations of thin slabs for stacks of isotropic [107] and anisotropic layers [250]. These methods require expensive precomputation and large storage requirements. On the other hand, approximate models have been successful in practice, due to their efficiency and simplicity [10, 228, 229]; these methods follow the microfacets formalism, which makes them very suitable for integration in modern offline and real-time renderers, at the cost of reduced fidelity of the underlying light transport. A more accurate precomputation-free solution was proposed by Guo et al. [76] based on Monte Carlo random-walks inside the layered material; they leveraged a position-free formulation of light transport for providing bidirectional estimators of the BSDF, leveraging variance reduction via multiple importance sampling. This work was later extended to support more advanced sampling estimates [54, 237]. Previous work [221] proposed a different formulation for multilayered materials, by proposing index-matched layers. This significantly simplified the layer stacking, allowing for closed-form solutions for the single-scattering BRDF, and allowing simple learning-based multiple scattering approximations. This learning-based modeling of layered materials was further explored recently for all components of the BRDF [75]. In this work, we leverage multilayer material formulations for the modeling of our makeup BSDF. Specifically, we implement our model on top of position-free Monte Carlo [76], following the index-matched simplifications that previous work proposed [221] in the context of layered materials.

VISUAL REPRODUCTION OF COSMETICS. A considerable body of work on visual reproduction of cosmetics has focused on image-space makeup transfer [24, 33, 113, 134, 138, 139, 206, 220, 242, 245, 247]. Such works do not aim at characterizing the underlying properties of cosmetics but rather transfers an example of makeup from an input photograph to a target image, as in style-transfer methods. The facial appearance of a 3D model is usually described by textures. One way to encode makeup is to model how texture layers like diffuse, specular, and scattering albedo change when the cosmetic is applied [184, 248]. While these methods are useful for setting parameters in an analytic spatially-varying BRDF, like the conventional microfacet model [210], they do not enable us to compute the appearance of an applied cosmetic based on the constituents of the makeup material. With reflectance measurements from a gonioreflectometer one can fit a microfacet model to the data and use this to analyze the appearances of different cosmetics. This has been done for liquid foundation cosmetics [160, 209], and the researchers found that a data-driven model based on principal component analysis (PCA) provided a better representation of the data than the Torrance-Sparrow model. However, a data-driven BSDF does not lend itself to editing and representation of related materials. Closer to our work, Huang et al. [93] and Li et al. [131, 132] reproduced the appearance of a liquid foundation by modeling a cosmetic layer using Kubelka-Munk theory for the diffuse reflectance and transmittance of the layer and a microfacet model for its specular reflectance. However, as found in the PCA study, the conventional analytic microfacet BRDF model cannot accurately capture the directional distribution of the scattered light and the Kubelka-Munk theory is diffuse, i.e., there is no directional dependency.

In contrast, we model the appearance of foundation makeup by computing the directional distribution of the light due to scattering by the different types of particles inside the layer.

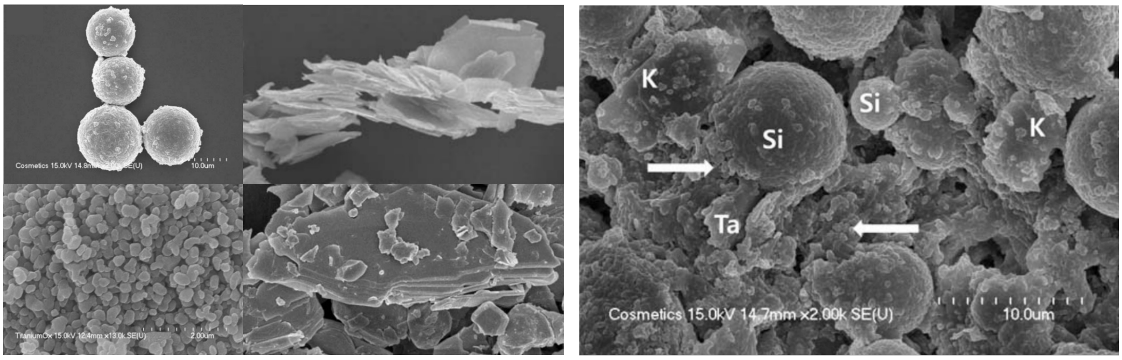

Natural State

Product Powder Mixed

Figure 3.2: *Left:* Scanning electron microscope (SEM) pictures of different minerals in their natural state (in clockwise order): Silica, Talc, Kaolin and Titanium Dioxide. Note how Talc and Kaolin have a polygonal and plate-like geometry, while Silica and Titanium Dioxide have a more round or cubic shape. *Right:* SEM image of the product powder mixed with the sebum stratum. Courtesy of Jeon and Chang [112].

3.3 FOUNDATION COSMETICS

Foundation cosmetics, as well as other types of cosmetics (e.g., eyeshadows, lipsticks or primer), are generally made from a combination of colorants, minerals, vitamins and other chemicals included for skin health reasons [226]. A fundamental distinction between foundations is the base used for dispersing these constituents. *Liquid-based foundations* typically disperse these particles within a host medium such as water or silicone. As the majority of the particles are not soluble in the host medium (talc is not water-soluble, for example) the final system is colloidal, meaning that tiny particles are dispersed within the host medium. In contrast, *powder-based foundations* use the sebum stratum to bind these materials together, forming a layer in which cosmetic particles are mixed with the sebum [112]. Common colorants include minerals such as talc, kaolin, silica, or titanium dioxide, which are present as small powder-like particles with different shapes [112]. While minerals like kaolin and talc have flake-like polygonal geometry, titanium dioxide or silica tend to have spherical or cuboidal shapes (Figure 3.2). This difference in shape results in radically different scattering behavior, leading to a different *finish* of the makeup. The color and finish are thus the two most important visual attributes for a correct reproduction of the appearance of a foundation layer.

COLORATION Since foundation layers are used to create an even layer, the color and hue of the foundation have to roughly match the underlying skin hue, making colorants a key aspect of a foundation layer. Colorants can be categorized into two types: dyers, derived synthetically, and pigments, which have a biological origin. As colorants can have an impact on human health, efforts have been intensified to develop new technologies to ensure a stringent control of colorant concentration in cosmetic products [71, 72]. The main source of coloration is scattering and absorption, since regulations [21] do not allow coloration through photoluminescence or other chemical reactions. In addition, diffractive particles are allowed to be used as colorants in effect

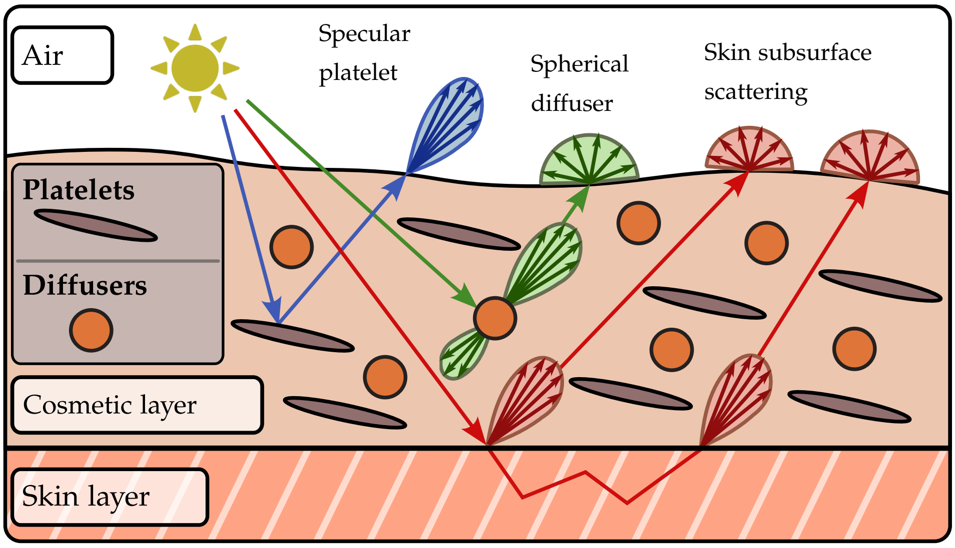

Figure 3.3: We model a single foundation layer as a medium filled with two types of particles: platelets and spherical diffusers. Platelets generate specular lobes, using a more anisotropic phase function. In contrast, spherical diffusers generate rougher finishes, due to a more isotropic phase function. We account also for the subsurface scattering due to the underlying skin model.

pigments [143], though these are not common due to the need for running extra safety studies. We focus on the most common scattering-based colorants.

FINISH Different foundations exhibit distinct visual characteristics ranging from a *matte* finish, resulting in a non-shiny, velvety appearance, to a *dewy* finish, aiming to achieve a more natural appearance that makes the skin brighter and shinier to evoke an impression of a healthy skin. The finish of a foundation layer is mainly caused by the concentration and shape of the different particles that constitute the bulk properties of the material, which directly affects the scattering behavior. While flake-like particles scatter light more coherently, spherical particles scatter light in a more isotropic fashion, leading to a more matte appearance.

SUMMARY The appearance of foundation cosmetics is the result of multiple scattering by a thin scattering medium, modeled by the combination of two types of scattering behavior: On one hand, we can assume that diffusers are responsible for the diffusive look and incoherent backscattering in matte appearances, while more directional scatterers (such as platelets) model the more coherent scattering reflection responsible for the highlights prominent in dewy foundations. These are somewhat aligned with different types of particles used to produce foundation cosmetics. Based on these different types of observed type of scattering, in the following we present our model, defined using radiative transfer theory as a mixture of different scatterers in a scattering medium.

3.4 A BSDF FOR FOUNDATION COSMETICS

As described in section 3.3, foundation cosmetics are applied as a layer on top of the skin, and its appearance is the result of the volumetric scattering in the layer. We model the layer statistically, as the combination of two uncorrelated scatterers uniformly distributed in the medium: spherical diffusers and anisotropic reflective platelets. Figure 3.3 shows a diagram of our model for a single foundation layer over the skin.

We assume elastic scattering, i.e., both the energy and scatterers numbers are conserved, and negligible wave-optical coherence in the multiple scattering component. In addition, given the small thickness of foundation layers, we assume no horizonal scattering, and therefore that light exits at the same position as it enters. Thus, we represent the foundation layer using a BSDF, defined as the result of all light-matter interactions occurring in the cosmetic for a ray of incident light at direction ω_i and exiting the surface at direction ω_o. Following Guo et al. [76] we model $f(\omega_i, \omega_o)$ using a position-free variant of the path integral as

$$f(\omega_i, \omega_o) = \int_{\Omega(\omega_i, \omega_o)} \Phi(\mathbf{x}) d\mu(\mathbf{x}), \tag{3.1}$$

with $\Omega(\omega_i, \omega_o)$ the space of light paths entering and exiting the differential surface at directions ω_i and ω_i respectively, $\Phi(\mathbf{x})$ the contribution of the path $\mathbf{x} \in \Omega(\omega_i, \omega_o)$ as the sequence of scattering and absorption events inside the foundation volumetric layer, and $\mu(\mathbf{x})$ the measure of path $\mathbf{x}$. By construction, the BSDF modeled by Equation (3.1) is energy conserving and reciprocal only if the scattering events conserve energy and are reciprocal.

MODELING THE OPTICAL PARAMETERS We model the foundation layer as a statistical volumetric scattering plane-parallel layer with thickness t. To account for specular-like scattering, we describe light transport using the anisotropic radiative transfer framework [104] (Equation (2.9)). The anisotropic RTE is characterized by the extinction and scattering coefficients, $\sigma_t(\omega_o)$ and $\sigma_s(\omega_o)$, and the phase function $f_p(\omega \to \omega_o)$. We assume that the host medium has negligible effect on these optical properties, and thus only spherical diffusers and platelets are responsible for the appearance. This is reasonable since the thickness of foundation layers is usually small, and the layer is in a semi-dry state. Most of the appearance is thus the result of the interaction with only the minerals and colorants. In Section 3.5 we validate this choice. Building on top of the spatial uncorrelation assumption, we model the extinction coefficient as the combined probability of extinction for both diffusers $\sigma_d(\omega_o)$ and platelets $\sigma_p(\omega_o)$, following

$$\sigma_t(\omega) = \sigma_d(\omega) + \sigma_p(\omega) = N_d\, C_d + N_p\, C_p(\omega), \tag{3.2}$$

with N_d and N_p the number density of spherical diffusers and platelets in m^{-3}, respectively, and C_d and $C_p(\omega)$ their extinction cross sections in m^2. Note that the extinction cross-section of the spherical diffusers has no dependency on the direction. To find practical user parameters, we model extinction as a function of a base extinction

$$\sigma_{\text{base}} = \sigma_d + N_p \max_{\omega} C_p(\omega) \tag{3.3}$$

representing the non-directional part of the extinction coefficient for all the scattering particles in the material. To include the directional dependency, we use

$$\sigma_t(\omega) = \sigma_{\text{base}}\left(c_d + (1 - c_d)\,\hat{C}_p(\omega)\right), \tag{3.4}$$

with $\hat{C}_p$ the normalized version of the extinction cross-section of the platelets and c_d the relative concentration of spherical diffusers out of the total concentration of particles:

$$\hat{C}_p(\omega) = \frac{C_p(\omega)}{\max_\omega C_p(\omega)} \quad \text{and} \quad c_d = \frac{\sigma_d}{\sigma_{\text{base}}}. \tag{3.5}$$

Since we assume that all chromatic effects arise only from absorption and that wave-dependent extinction cross-sections do not contribute, our $C_p(\omega)$ shows no spectral dependence.

To introduce practical parameters for specifying the part not being absorbed, we use non-directional spectral single scattering albedos for the spherical diffusers α_d and the platelets α_p. The single scattering albedo is defined by $\alpha(\omega) = \sigma_s(\omega)/\sigma_t(\omega)$, so we use the factors for diffuse and directional extinction in Equation 3.4 to model this:

$$\alpha(\omega) = \frac{\alpha_d\, c_d + \alpha_p\,(1 - c_d)\,\hat{C}_p(\omega)}{c_d + (1 - c_d)\,\hat{C}_p(\omega)}. \tag{3.6}$$

Finally, we build our phase function as a linear blend of the phase functions of diffusers and platelets $f_{p,d}$ and $f_{p,p}$, weighted by the total scattered light by each of the particles following

$$\begin{aligned} &f_p(\omega \to \omega_o) \\ &\quad = \frac{\alpha_d\, c_d\, f_{p,d}(\omega \to \omega_o) + \alpha_p\,(1 - c_d)\,\hat{C}_p(\omega)\, f_{p,p}(\omega \to \omega_o)}{\alpha_d\, c_d + \alpha_p\,(1 - c_d)\,\hat{C}_p(\omega)}. \end{aligned} \tag{3.7}$$

This phase function is normalized if both $f_{p,d}$ and $f_{p,p}$ are normalized. As discussed in Section 2.1.4 the phase function of Equation (2.9) is not reciprocal, but reciprocity is achieved due to $f(\omega \to \omega_o)\sigma_s(\omega) = f(\omega_o \to \omega)\sigma_s(\omega_o)$. Thus, our BSDF (3.1) is energy-conserving and reciprocal. In the following, we detail the scattering behavior of both diffusers and platelets, and the implementation of our model.

3.4.1 *Scattering by Diffusers*

Following the measurements by Wang et al. [222], who analyze colloidal systems made of nanoscopic titanium dioxide particles, we use a mixture of phase functions to model the scattering of diffusers. This is common practice in computer graphics and related areas [38, 67, 259]. In particular, we model the scattering from diffusers using a mixture of two Henyey-Greenstein lobes $p_{f,HG}(\omega \to \omega_o|g)$ each parameterized by the mean cosine of the scattering angle g [89]. This is similar to previous work modeling the scattering of sunscreen lotions [168]. The resulting phase function is

$$f_d(\omega \to \omega_o) = w_g p_{f,HG}(\omega \to \omega_o|g_1) + (1 - w_g) p_{f,HG}(\omega \to \omega_o|g_2), \tag{3.8}$$

with $w_g \in [0, 1]$ the blending parameter between the two lobes.

3.4.2 *Scattering by Platelets*

To reproduce the glossy appearance required by a dewy finish, we use platelet particles, i.e., flat particles suspended in the medium that generate a glossy reflection. Platelets are purely reflective micro-flakes [104] suspended in the medium, following the SGGX distribution of normals $D_{SGGX}(\omega_m)$ [88], with ω_m the microflake normal.

We assume a disk-like distribution of platelets, and parametrize $D_{SGGX}(\omega_m)$ by a roughness Λ_p parameter along the distribution's mean direction, and a rotation angle θ_p that rotates the mean direction of the distribution with respect to the tangent direction of the surface, so that the distribution does not need to be aligned with the surface normal.

The distribution of normals $D_{SGGX}(\omega_m)$ directly enables us to compute the projected area of the platelets, which we, as in previous work [87, 104], assume equal to $\hat{C}_p(\omega)$. As phase function, we opt for a specular SGGX phase function without Fresnel effects, so that the phase function for platelets is defined by

$$f_p(\omega \to \omega_o) = \frac{D_{SGGX}(\omega_h)}{4\hat{C}_p(\omega)}, \tag{3.9}$$

with $\omega_h = (\omega + \omega_o)/\|\omega + \omega_o\|$ the half vector.

3.4.3 *Implementation & Properties*

We implemented our model in PBRT v4 [177] on top of the position-free Monte Carlo framework from Guo et al. [76]. This allows us to compute multiple scattering as an average of random walks inside the layer with next-event estimation for variance reduction, and to easily stack multiple layers of cosmetics. The source code, scenes and reflectance measurements are available on the project page.

We compute Equation (3.1) using a Monte Carlo estimate, by randomly sampling paths starting at direction ω_i. We build the random walk by using exponential mean free path sampling with $\text{pdf}(\omega, s) = \sigma_t(\omega)\exp(-\sigma_t(\omega)\, s)$. A collision is found as long as the sampled distance s remains inside the layer. Otherwise, we move to the contiguous layer (the skin or outside). At each interaction inside the medium, we select a lobe (platelet or one of the two Henyey-Greenstein lobes) using Shirley's remapping [190] (with probabilities driven by c_d and w_g). Then, platelets are sampled using SGGX visible normals sampling [87], while Henyey-Greenstein phase functions are sampled using the usual CDF-inversion-based routine. Absorption is handled by multiplying the throughput by the single scattering albedo (Equation 3.6) at each bounce. To reduce variance, we only use Russian Roulette as a path termination criterion after 128 bounces. The sampling routines and PDF functions are evaluated following PFMC [76], using a forward path sampling inside the medium, with PDF the probability density of generating the path.

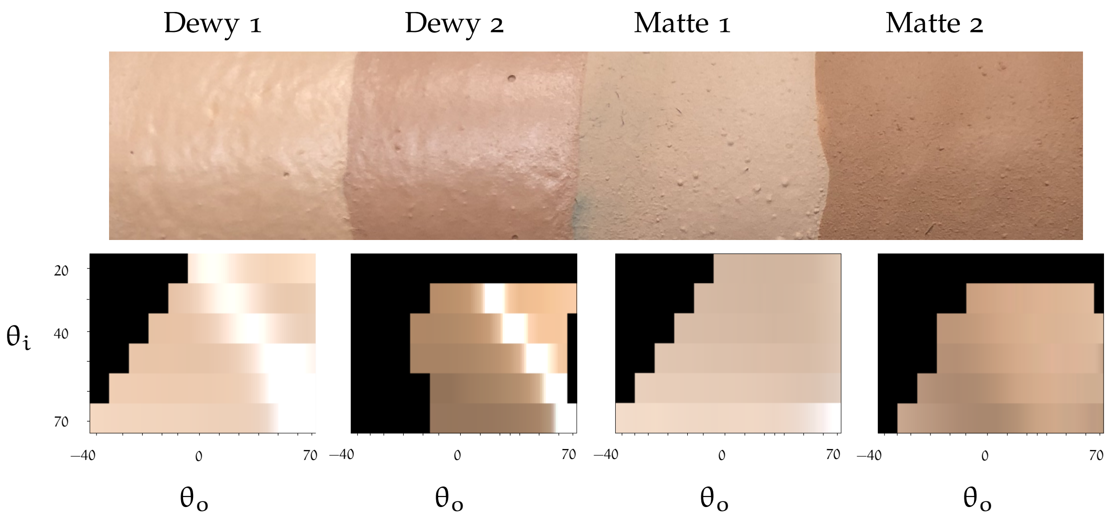

Figure 3.4: Measured samples (top row), wrapped around a cylinder for demonstration purposes. Reflectance profile measured with our setup (bottom row), with angles of observation θ_o on the horizontal axis and angles of incidence θ_i on the vertical axis. Negative angles of observation represent retro-reflectance configurations. Black pixel areas indicate that no information was available.

3.5 ANALYSIS AND EVALUATION

In this section, we validate our model against measurements of cosmetics reflectance. We capture the reflectance from four liquid-based foundations from *Clarins*, with varying colors and finishes, from dewy to matte. In particular, we select *Skin Illusions 105 Nude* (Dewy 1), *Skin Illusions 112 Nude* (Dewy 2), *Skin Illusions Velvet 103 Nude* (Matte 1) and *Skin Illusions Velvet 108* (Matte 2). We show samples for each of them in Figure 3.4, top row.

3.5.1 *Sample Preparation*

We captured the reflectances using a home-built setup, consisting of a gonioreflectometer, a light source (Xenon lamp emitting in about 6200K temperature), and an Ocean Optics spectrometer with fibre input and an attached collecting lens, as shown in Figure 3.5. We directed a collimated beam with a diameter of 4 mm onto the sample. The direction of incidence was explored by rotating the sample holder, and the direction of observation by independently rotating the receiver's arm. Reflectance values were estimated by dividing the received spectral values by the signal reflected from a white reference, a 99% spectralon by Labsphere. Calibration with spectralon additionally compensates for the geometrical foreshortening [225]. We collected our measurements on a flat black slab made of matte polymer, which we found easier to calibrate than synthetic skin, and that simplified the layer beneath the cosmetic, which was useful for optimization. We could not directly measure the thickness of the samples, however, we performed several measurements on a silicon plate (mirror) to validate that we were in the multiple scattering regime. We added thin layers of the cosmetic product until we found no difference in measurements with respect to the number of layers.

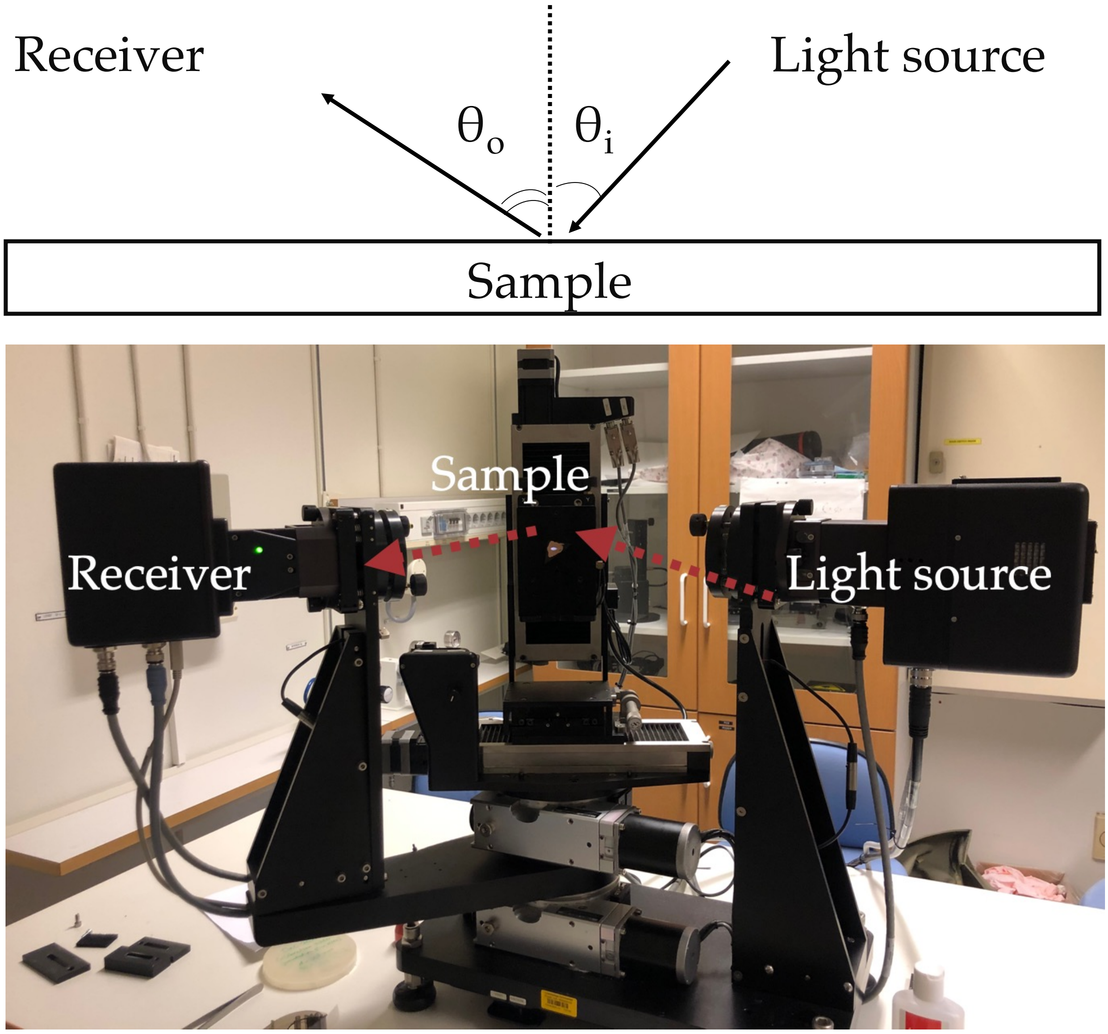

Figure 3.5: Experimental setup for measuring reflectance. We modify a commercial gonioreflectometer to collect spectral reflectance in non-specular directions.

In preparation for our samples, we applied a thick layer of cosmetic product on a black substrate. This is different from previous work [160], where measured cosmetics were applied on top of synthetic skin that might interfere with the cosmetic's reflectance data. We measure our samples in a semi-dry state, leaving them resting for an hour to allow the formation of a homogeneous layer. This resembles better the condition of a product when it is applied since usually a small quantity of product is spread over a comparatively large surface. Air exposition is likely to evaporate the product into a semi-dry state. This mechanism might also impact the hue of the product [246].

3.5.2 *Captured Data*

Figure 3.4 shows the captured data for our four liquid foundations. We plot measured data as sRGB- colored patches to better visualize the color and highlight shape changes. As we sparsely gather measurements we linearly interpolate between measurements and leave black pixels to represent lack of data. We measure both forward and backscattering for all products, although our setup does not allow us to measure for all light-view configurations, in some cases due to low a signal-to-noise ratio.

We only capture the inclination angle and not azimuthal data as our setup did not allow it. However, we believe that this is not a major issue since cosmetics are often applied on the skin without following any precise stroke direction, but rather in a circular manner, precisely to hide any anisotropy. Therefore, it is reasonable to assume that a small amount of reflectance changes occurs in the azimuthal plane and that most of the changes in appearance happen along the inclination angle.

As seen in Figure 3.4, the selected samples exhibit different reflectance behaviors: Dewy 1 and Dewy 2 displays a dewy finish, characterized by sharp highlights. Matte 1 can be considered an intermediate foundation, similar to a velvet finish, showing highlights for only grazing angles. In contrast, Matte 2, is clearly matte (absence of highlights). In all cases, we found some backscattering, in agreement with previous work [160].

3.5.3 *Comparison between model and captured data*

To analyze how accurately our model can represent the actual appearance of cosmetics, we fit our model to the measured data for each of the four samples, by solving

$$\operatorname*{argmin}_{\boldsymbol{\pi}_i} \sum_{(\theta_i,\theta_o)\in M} d\left(I\left(\theta_i,\theta_o\right), f_s\left(\theta_i,\theta_o,\boldsymbol{\pi}_i\right)\right), \tag{3.10}$$

where $\boldsymbol{\pi}_i$ is the appearance model parameters (see Section 3.4), $I(\theta_i, \theta_o)$ is the measured data in linear RGB space at incoming angle θ_i and outgoing angle θ_o, $f_s\left(\theta_i,\theta_o,\boldsymbol{\pi}_i\right)$ is our BSDF, and $d(\cdot,\cdot)$ is the L_2 distance metric. Our model is evaluated using Monte Carlo integration as described in Section 3.4. We use a derivative-free optimizer (modified Powell algorithms from the SciPy library); since this approach is sensitive to initialization we first perform a manual selection of the parameters and then use the optimizer to refine the solution. We use the L2 distance function in linear RGB space, although we show the results in sRGB space. We prepared the digital replica by setting the underlying layer as a black Lambertian surface and set the foundation layer thickness to 16 optical depths to enforce the multiple scattering conditions in which samples were taken. Computation-wise, we limit the number of bounces to 512, we then sample each direction 128 times. We evaluated this setup with preliminary experiments and found that it offered a good trade-off between time of execution and energy loss. We run the optimization described, allowing the model to optimize also for the IoR and surface roughness parameters (using the Trowbridge-Reitz distribution). At the end of the optimization cycle, we found values close to 1.0 for the index of refraction, thus supporting the validity of our interface-free simplification of the medium, which allows us to further reduce the number of parameters and reduces the computational cost by removing an interface from the computations.

Figure 3.6 shows the results of the optimization process for the four samples (last row), among the error measured under the L2 loss function. Our model captures the general behavior of the measured data in terms of reflected energy, color and highlight shapes. Some small differences are present in the intensity gradients of the diffuse part of the darker samples at grazing angles (Matte 2 and Dewy 2). We discuss these differences in Section 3.7.

ABLATION STUDY. We ran an ablation study of our model to examine the effect of its various components. Using the approach discussed above, we optimized three incremental versions of our proposed model: One composed of spherical diffusers only (Diffusers Only), one that is composed of platelets only (Platelets Only), and one that uses both particles but with a single lobe for the spherical diffuser particles (Full, Single Lobe). We initialize the optimization of each variation using the fitted parameters of our full model (Full, Two Lobes). Figure 3.6 shows the results of this ablation study:

It demonstrates that a model with only one particle type, either diffusers or platelets, is unable to properly reproduce the highlights of glossy appearances (Dewy 1 and Dewy 2). The reason is that one type cannot simultaneously generate highlights at grazing angles and in near-specular directions. In contrast, the combination of both particles with single-lobe diffusers generates results that are comparable with our full model, except for minor hue differences in Dewy 2. This suggests that a model with single-lobe diffusers could be sufficient to represent some real-world cosmetics, at the cost of losing expressivity, as we show in the next section.

3.6 RESULTS

In this section, we visually analyze our model in realistic use cases. We implemented our model as a layered material in PBRT v4 [177]. For the base skin, we use the random-walk-based default PBRT skin model. We first demonstrate our model using the realistic materials captured in the previous section; then we analyze the appearance space defined by the parameters of our model.

We report rendering times and sample counts, in addition to the user parameters obtained through optimization of the measured foundations, in Appendix C.1.

CAPTURED DATA Figure 3.7 shows our model with the parameters obtained from captured data, applied on top of the PBRT skin model. As expected, while the two dewy foundations enhance the highlights, matte foundations remove the highlights. Note how both dewy and matte foundations also slightly reduce the effect of skin subsurface scattering.

In Figure 3.8, we qualitatively evaluate our model against pictures of cosmetics applied on real skin. Our model reproduces the trends shown in the pictures: an increase in reflection when Dewy 1 is applied to the skin, and a more matte appearance, especially at grazing angles, for Matte 2.

APPEARANCE EXPLORATION Figure 3.9 (a) explores the effect of the concentration of diffusers c_d and the thickness t (in mean-free-path units). In this experiment, we use a single phase function for the spherical diffusers. Increasing the thickness reduces the effect on the appearance of the underneath layer, while increasing the concentration of

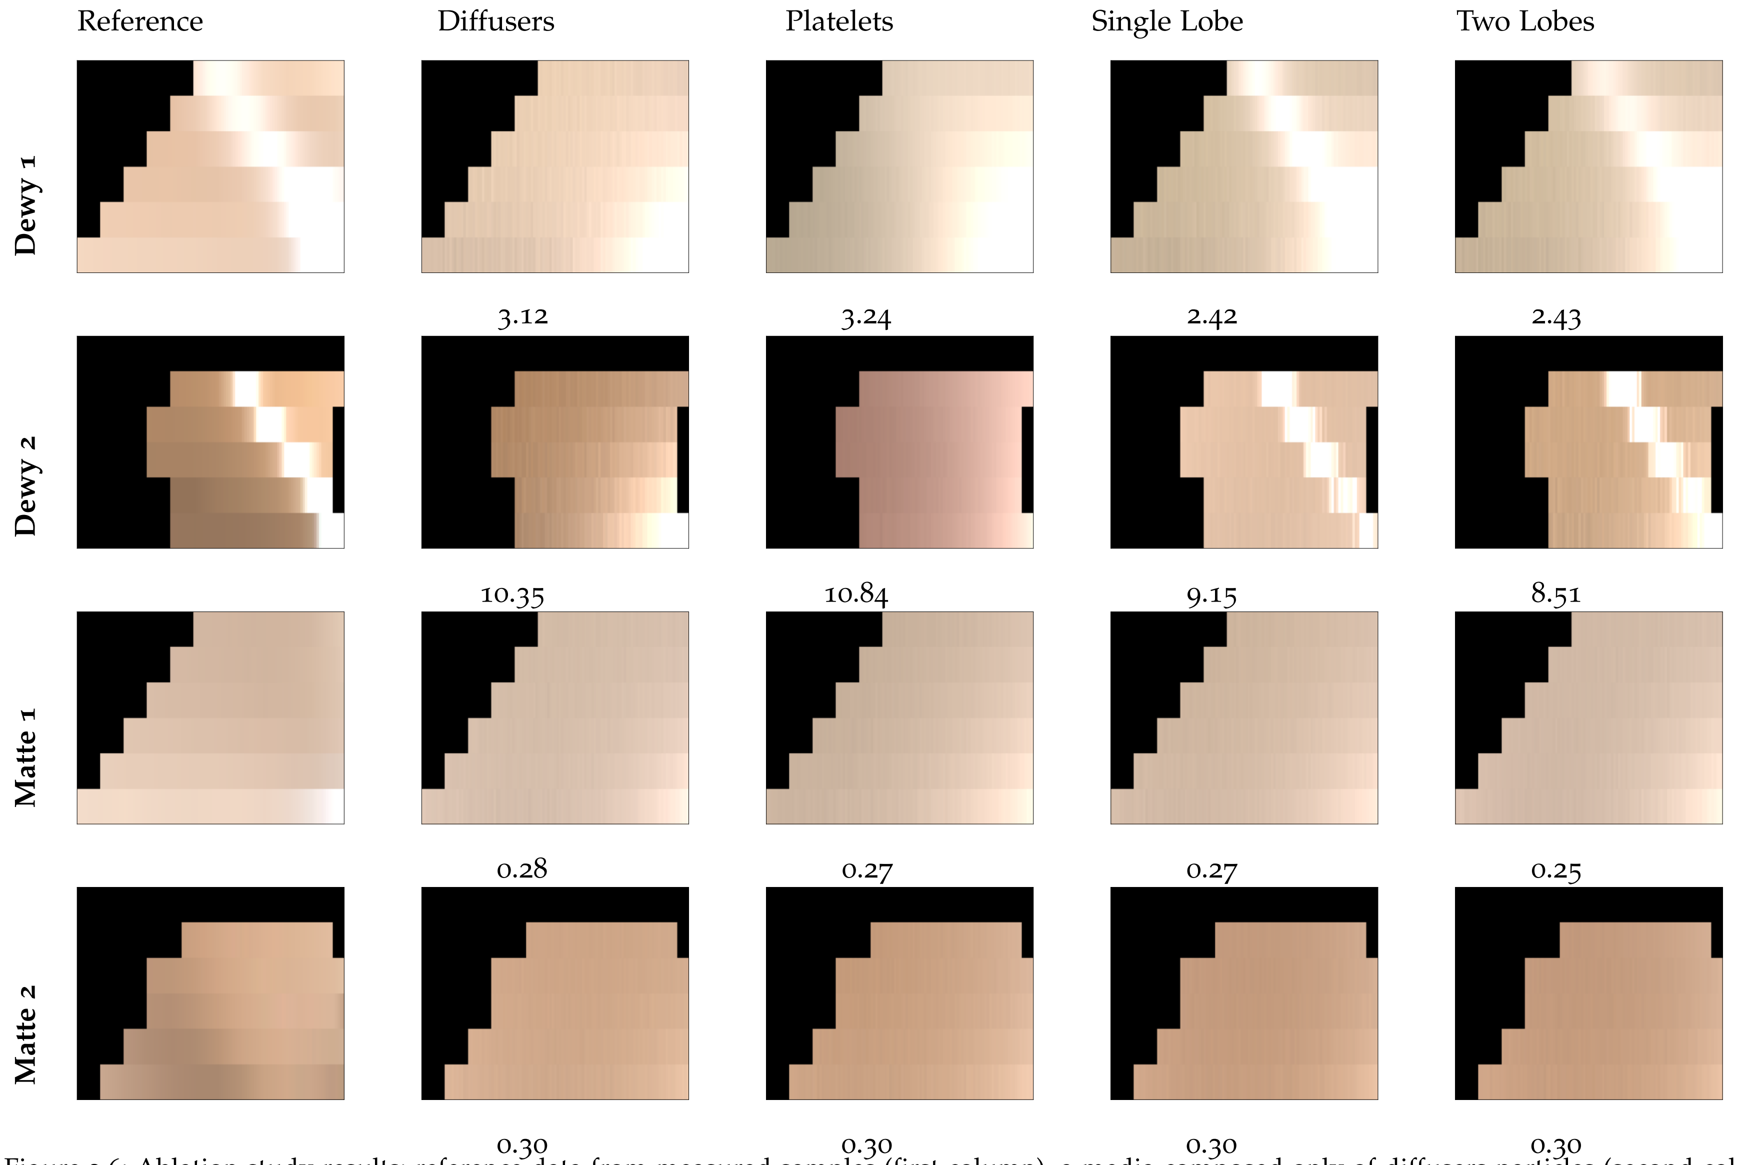

Figure 3.6: Ablation study results: reference data from measured samples (first column), a media composed only of diffusers particles (second column), a media made of only platelet particles (third column), a media with both types of particles (diffusers and platelets) with a single lobe phase function and our full model (last column). For each configuration, we report the error under the L2 distance. Notice how the models using only one type of particle cannot approximate well both types of appearances (matte and dewy) and the full model in overall exhibits the lowest error across all the products.

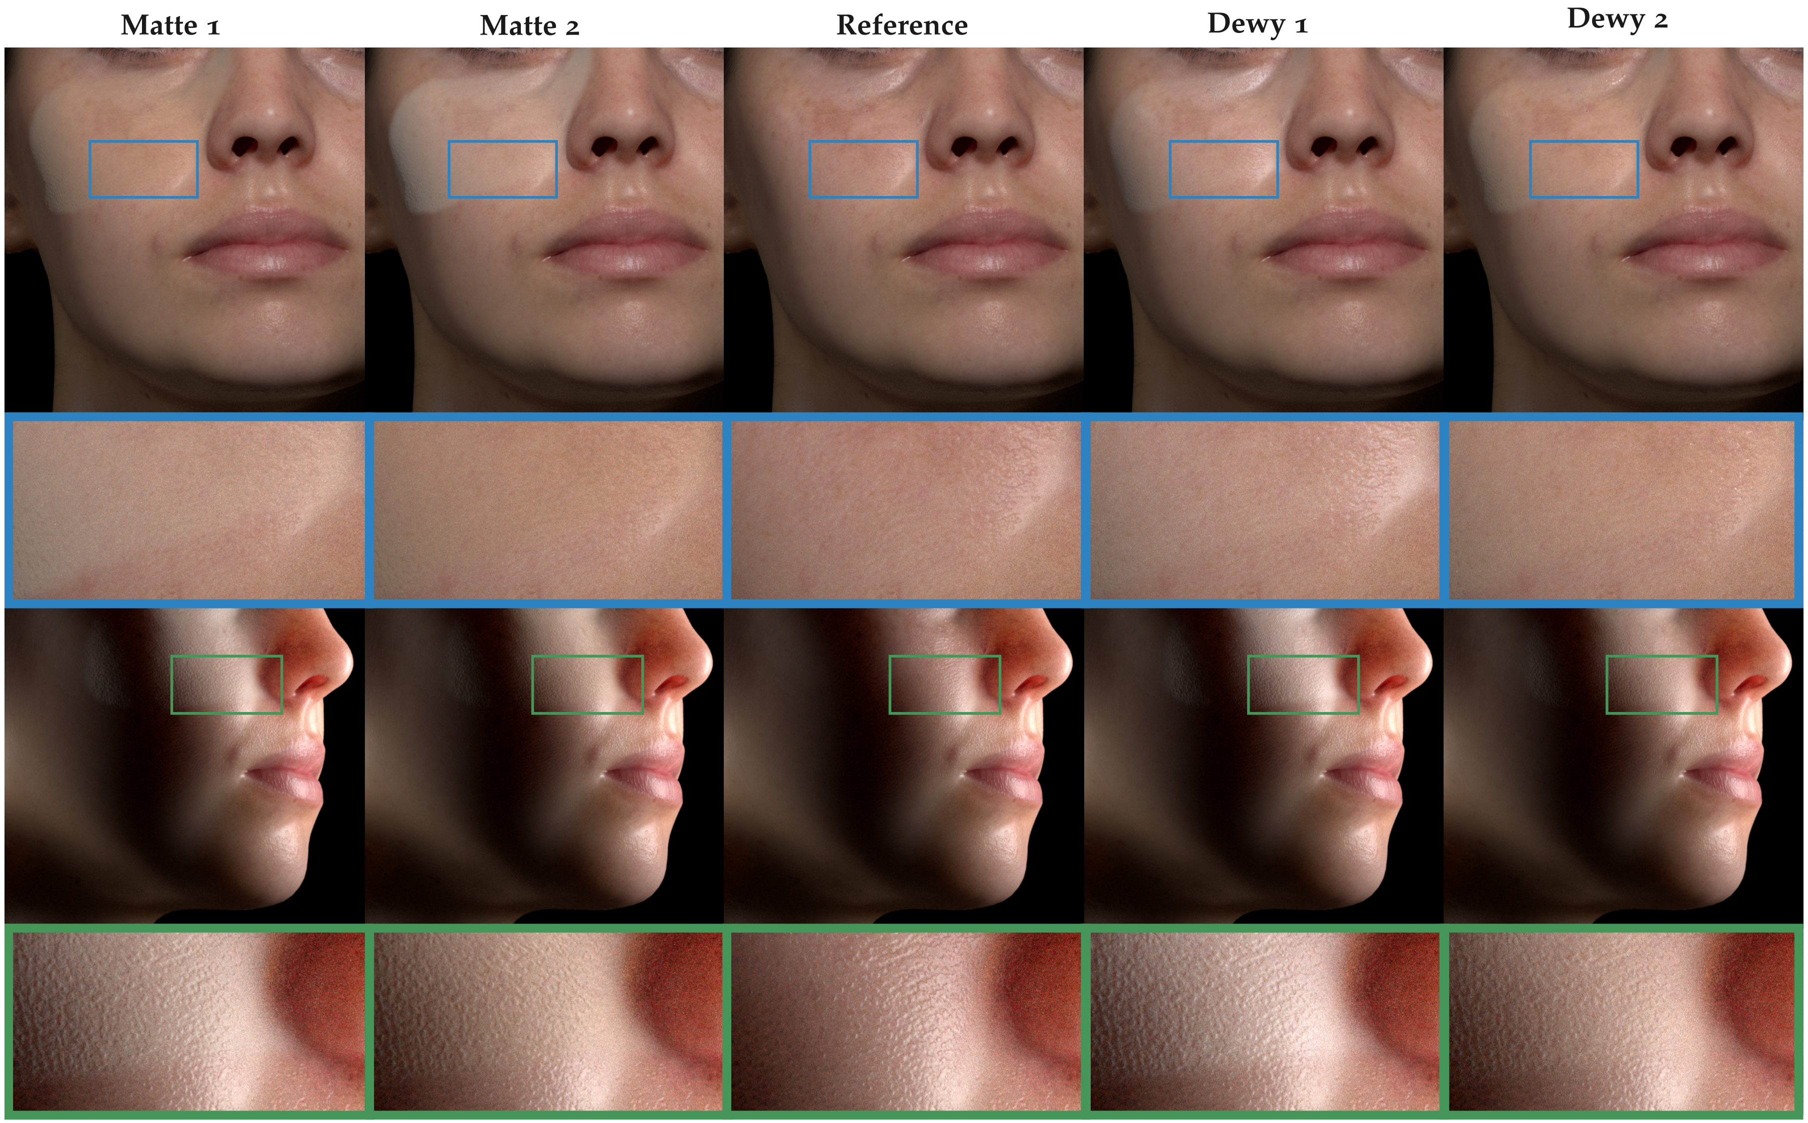

Figure 3.7: Analysis of our model applied over realistic skin. In the first row, we show the effect that cosmetics have on highlights. As can be seen in the insets (blue boxes), the matte cosmetics tend to reduce the brightness of highlights, while the two dewy products enhance them. In the second row, we show how a layer of cosmetics affects the subsurface scattering of the underlying skin. In each column, we apply the optical parameters retrieved for each product.

Figure 3.8: We compare renderings obtained with our model (right) with photographs (left) of cosmetics applied over bare skin on the forehead. In both renderings and pictures, we apply the cosmetic on the left side of the forehead, leaving the right side with bare skin for reference. Our model captures the subtle increase in reflectance when Dewy 1 is applied (central area of the forehead). On the other hand, the model is also capable of generating a matte appearance for grazing angles when Matte 2 is applied.

platelets (smaller c_d) increases the glossiness of the skin. In contrast, when increasing c_d the appearance shifts to matte, as expected. The reason for brighter highlights with increasing thickness in dewy foundations was observed by Yoshida et al. [249], where the authors explained that multiple applications of cosmetics lead to an increase in the reflected light. We have to clarify that this is not necessarily a general characteristic of all cosmetics, but rather a behavior that is shown by a certain subset of products.

In Figure 3.9 (b), we investigate the influence of c_d and Λ_p, the concentration of diffusers and the roughness of the platelets. We see that both parameters alter the perception of glossiness, although in different ways. While c_d affects the intensity of the highlights, Λ_p alters the shape of the highlights, going from sharper to broader highlights. Similar to before, these phenomena can be explained by observing that increasing c_d only decreases the chances of hitting a platelet, without altering the scattering behavior of the particle itself. On the other hand, Λ_p affects the scattering behavior of the particle, with rougher appearances obtained for bigger particles. We believe that these two parameters, c_d, and Λ_p can be helpful during editing, as they allow artists to edit two different dimensions of gloss perception, contrast, and sharpness of the highlight [176].

We investigate the influence of using one or two-lobed phase functions in Figure 3.9(c) to see how t and c_d alter the appearance when diffusers are modeled with different phase functions. We observe that when diffusers are defined with a two-lobed phase function, the same parameters yield darker and more saturated colors. This can be explained by noting that when having two strongly anisotropic lobes, light tends to be either forwarded or backscattered, thus resulting in a darker appearance, since less light gets scattered horizontally.

As expected, the influence of the skin layer is greater for small thickness values, while for higher values the appearance converges to a similar result, regardless of the appearance of the skin.

In Figure 3.10, we explore the capabilities of our model to reproduce the adequate hue for different skin types, for different product types. In each row, we apply four cosmetics strips, with only one of them matching the hue of the underlying skin type.

This illustrates how a cosmetic product that does not match the hue of the underlying skin can be easily spotted. We run this experiment for both dewy and matte foundations, showing that our model is capable of generating an adequate hue for the different types of products. Matte cosmetics, as expected, provide a closer color match, while dewy foundations provide more desaturated colors in exchange for increased highlights

We also observe that the effect of glossiness is more prominent for darker skin tones, as observed in previous work on glossiness perception [176]. This is a key aspect that has to be considered for different skin types.

Finally in Figure 3.11, we explore the results of applying layers of foundations characterized by different final appearances. We run this experiment by taking dewy and matte foundation cosmetics and observing how they interact with each other depending on the order in which they are applied. We observe that the topmost layer dominates the final appearance. This is expected as the brightness of highlights is dominated by the first bounce, therefore the first few bounces are key in determining the final appearance.

3.7 CONCLUSIONS

We proposed a scattering model for characterizing cosmetic material using colored volumetric spherical diffusers and platelet particles. We demonstrated that this model is capable of reproducing the main characteristics of foundation layers by fitting the model to measured bidirectional reflectance data of sample products. We investigated the appearance space that our model spans, confirming that is possible to achieve believable results with our model. As opposed to previous work, we offer a model that can be easily edited and used to predict different appearances without necessarily being entangled with the skin type on which it was applied. We believe that our model offers a solid base but additional work is required to capture all the subtle nuances of cosmetic materials.

Strictly speaking, our model focuses on foundation makeup, and we have only validated it against this type of cosmetics. However, as shown in Figure 3.1, we can use our model to replicate the look of other types of makeup, which shows the versatility of our approach. Nevertheless, a more thorough analysis against measurements would be required for validating other ranges of cosmetics. Note also that the foundation samples are specific for light skin; however, as shown in Figure 3.10 our model is able

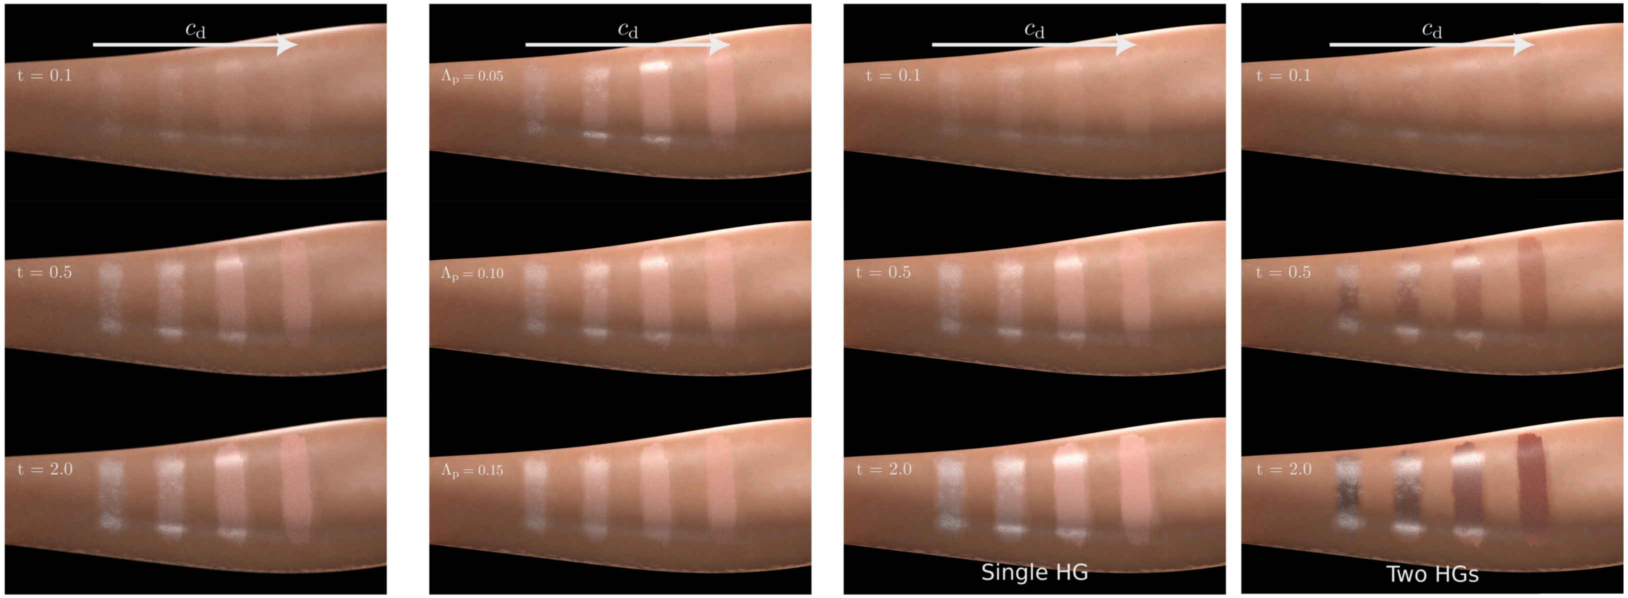

(a) Thickness vs Concentration

(b) Roughness vs Concentration

(c) single lobe vs two lobed

Figure 3.9: Appearance exploration - we investigate the influence of the parameters of our model on the final appearance. In (a) we show the influence of thickness t and concentration of diffusers c_d. Thickness has a major impact on perceived reflectance. In (b) we explore how Λ_p, the roughness of platelets, with c_d affect the shape and brightness of highlights. Λ_p mainly affects the shape of the highlights, while c_d alters their intensity. In (c) we study the effect of using a single or a two-lobed phase function for diffusers. Two-lobed phase functions generate darker and more saturated appearances.

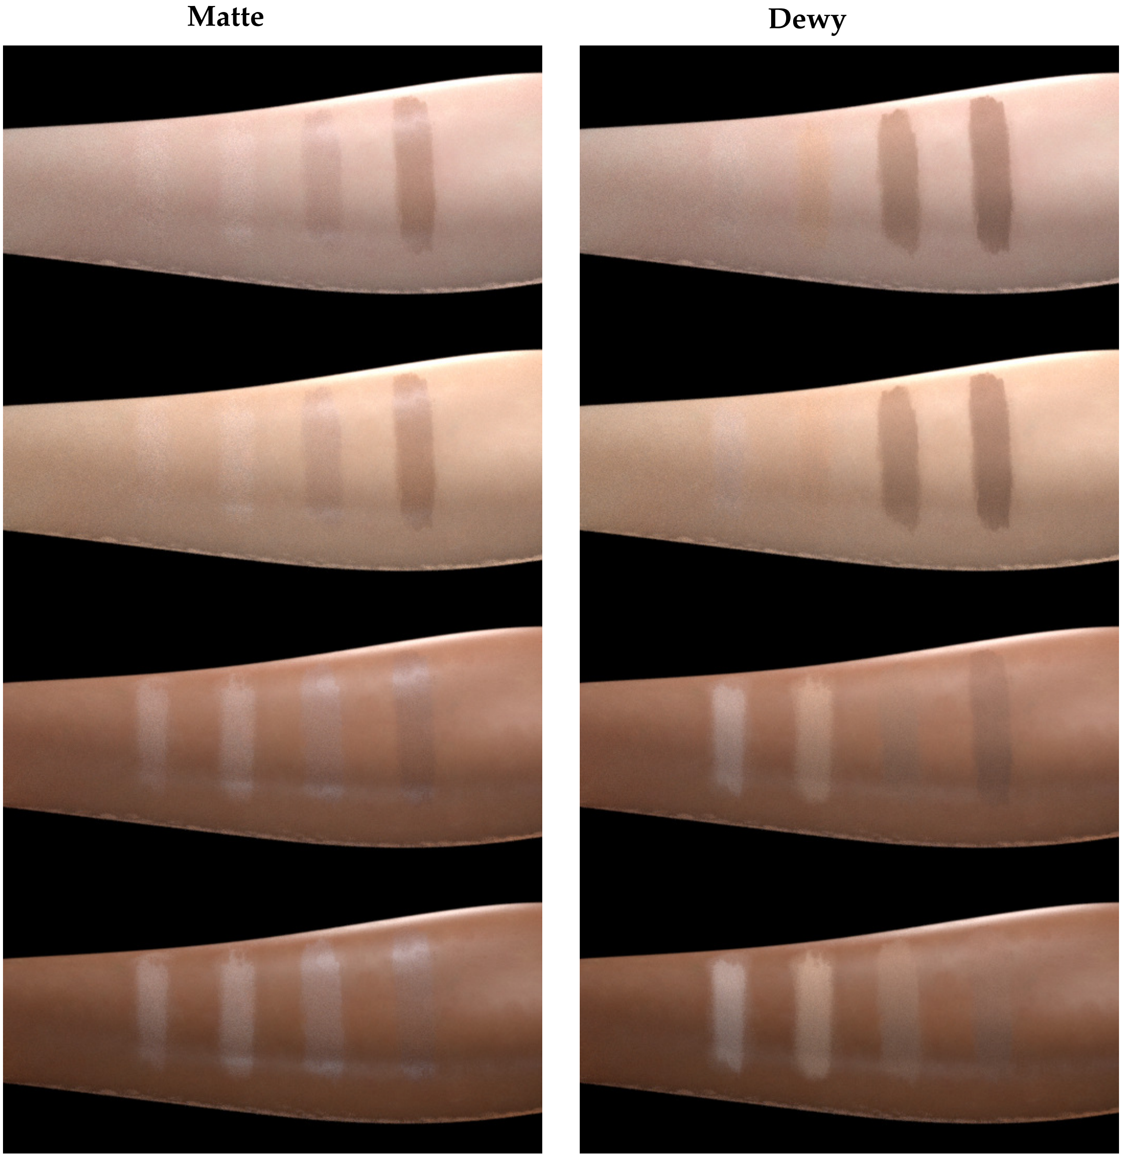

Figure 3.10: Experiment matching the hue of different skin types. In each row, we simulate different cosmetic strips aimed to match the hue of different skin types. We apply the same cosmetics to all rows changing the underneath skin type. Note how incorrect hues can be easily spotted if not applied to the correct skin type. We repeat this experiment for the two categories of matte and dewy cosmetics. All results are obtained using $t = 0.5$

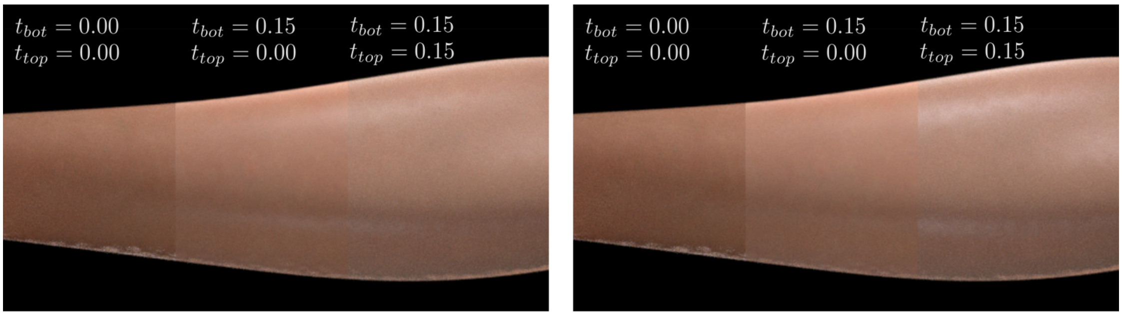

(a) bottom: dewy - top: matte

(b) bottom: matte - top: dewy

Figure 3.11: Effect of stacking two layers of foundation. (a) the bottom layer uses a dewy foundation, while the top layer is composed of a matte foundation. In (b) the bottom layer uses a matte foundation, while the top layer is a dewy foundation. Note how the top-most material dominates the final appearance.

to qualitatively capture the appearance of cosmetics targeting other skin types. Particularly interesting for future work is to investigate reflectance models for foundation layers that incorporate glittering effects. However, measuring this class of materials can prove to be a challenging task, as it is not clear how to measure the reflectance of a material that glitters if not by taking the aggregated behavior. An additional limitation is that our model builds upon pure ray optics, and thus it ignores diffraction effects occurring due to the powder-like nature of the scatterers. That might explain why our model is not able to fully reproduce the reflectance behavior exhibited by Dewy 2, which we hypothesize is caused by the interference caused at grazing angles observed in previous work for different materials [42, 140]. Adding a diffraction reflection lobe [92] could help our model better fit the captured data and appearance behavior at grazing angles, and we think it is an interesting research direction for future work.

Studying the effect of heterogeneous and uncorrelated scatterers is an interesting avenue for future work as concurrent work shows that varying pigments radius and mass fraction can have a significant impact on the colored appearance of cosmetics [212]. Nevertheless, we found that using the classical RTE formulation for uncorrelated media fits well with measurements, while it also reduces the additional complexity of specifying non-exponential mean-free-paths.

In terms of implementation, our model is based on the PFMC framework, which requires several samples to converge to a noise-free solution. Implementing our model using faster position-free approaches [13] or using the SpongeCake model [221] should be trivial, and would likely reduce significantly the overhead of our method. Finally, our model focuses on the local microscopic appearance of cosmetics, and thus does not account for the mesoscopic effects of cosmetics masking small crevices or pores which should change the normal or displacement mapping modeling of the skin mesogeometry.

4

A SURFACE-BASED APPEARANCE MODEL FOR PENNACEOUS FEATHERS

ABOUT THIS CHAPTER

In this work, we present a far-field model for pennaceous feathers. We represent the far-field appearance of feathers using a BSDF that implicitly represents the light scattering from the main biological structures of a feather such as the shaft, barb, and barbules. Our model accounts for the particular characteristics of feather barbs such as the non-cylindrical cross-sections and the scattering media via a numerically-based BCSDF. To model the relative visibility between barbs and barbules, we derive a masking term for the differential projected areas of the different components of the feather's microgeometry, which allows to analytically compute the masking between barbs and barbules. While I am not the leading author I contributed to the implementation and revision of the reflectance model, the development of the scenes used in the experiments, and in running of the experiments. For completeness, we report here the full project. The work discussed in this chapter has been accepted to Computer Graphics Forum and presented at the Pacific Graphics 2024 conference.

J.R. Padrón-Griffe, D. Lanza, A. Jarabo & A. Muñoz
A Surface-based Appearance Model for Pennaceous Feathers

Computer Graphics Forum (Pacific Graphics 2024)

4.1 INTRODUCTION

Feathers are the distinctive characteristic of birds and are a unique structure not shared with any other animal. They are crucial for flying but also allow birds to control their body temperature, camouflage, and they are fundamental for communication in mating, aggression, and dominance [90].

The appearance of feathers is varied and rich, different among species, and is partially explained by the geometrical complexity of their structure at several levels: Each feather is composed of a central shaft, called the rachis, with its base (the calamus) is inserted to the skin. Serial fiber-like branches (barbs) emerge from both sides of the rachis. A second level of branching emerges from the barbs (the barbules), which in pennaceous feathers become attached to adjacent barbs, forming a flattened surface (the vane). Depending on the structure formed by barbules, feathers can have both pennaceous and plumaceous sections. When barbules are interlocked to adjacent barbs, they create a tense and semi-rigid surface, characteristic of pennaceous areas. When barbules are loose, they form a fluffy, irregular volume typical of plumaceous sections. Each of the barbs and barbules components possesses a scattering microstructure and nanostructure that greatly affects appearance: Different pigments (carotenoids and melanin) produce coloration through spectral absorption, while barbs have a mul-

Figure 4.1: A rendering of the wing of an *Amazon parrot*. *Left:* Bird model with a baked texture from a reference photograph of an Amazon parrot. Notice how the baked texture looks flat. **Center:** A surface-based rendering of the wing feathers, using a hair model [149] for representing the scattering from the microscopic structures of the wing, similar to previous approaches [7, 78]. *Right:* The same scene with feathers rendered using our model accounting for accurate masking and complex medulla using our BCSDF. Our far-field reflectance model represents more accurately the feather substructures and coloration mechanisms, given a closer appearance match with the baked texture overall appearance including features such as the diffuse look of feathers, and the subtle goniochromatism due to visibility changes between barbs and barbules.

tilayered structure with an inner scattering medulla which in some species have a quasi-ordered nanoscopic structure producing diffuse structural coloration. It is only the aggregated behavior of light-matter interactions at these scales that produce the complex appearance of feathers. Beyond direct applications for traditional computer graphics applications, the appearance of feathers has also been studied for ornithology [90], paleontology for color reconstruction of extinct dinosaur feathers [251], or fabrication of new biomimetic materials [44].

Compared to other biological appearances such as skin [37, 195], hair [149] or fur [244], rendering of feathers, and in particular of pennaceous feathers, is a relatively unexplored area in computer graphics, with some notable exceptions that either oversimplify the appearance [202], bake it in a bidirectional texture function [27], focus on specific bird species [95], or use expensive curve-based representations for the barbs with simplified fiber scattering functions [7]. While modeling the barbs as curves [7, 95] is flexible, allows for a very fine-detailed representation of the feathers, and explicitly accounts for geometric effects such as visibility, it might become very expensive when representing many feathers, quickly becoming impractical in most applications.

In this work we propose a far-field surface-based appearance model for pennaceous feathers, that encodes the geometric complexity of the feather by using lightweight textures, and that it is able to correctly predict the geometric attenuation that so far could only be modeled with explicit curve-based feathers, see Figure 4.1. At the core of our model is a new bidirectional scattering distribution function (BSDF) that accounts for the scattering of barbs and barbules stochastically, based on their density and orientation, as well as their relative visibility. Inspired by previous works [7, 95, 261], we model the individual scattering of both barbs and barbules using a fiber-based bidirectional curve scattering distribution function (BCSDF). However, as opposed to these previous works, our BCSDF accounts for the ellipticity of the fibers, and the effect of the internal non-centered scattering medulla, which are crucial for the appearance of most feathers, specially bright colored feathers.

We demonstrate our work representing a wide baseline of feather appearances, including pure white feathers, blue feathers due to diffuse structural coloration, or feathers with color resulting from hybrid coloration. None of these can be represented with previous models. In particular, our contributions are:

- A flexible BSDF that models the far-field appearance of pennaceous feathers,
- an elliptical BCSDF that generalizes previous curve scattering models by accounting for both elliptical cross-section and the presence of a non-centered scattering medulla,
- and an analytic masking term that combines the relative contribution between barbs, barbules and transparency, based on their differential projects areas.

4.2 RELATED WORK

FEATHER APPEARANCE The modeling of the appearance of feathers remains a relatively unexplored area of research. Most of previous works focus on the modeling of the geometry of the feather, defining parametrical models that simulate the geometrical structure of each feather. Rachis and barbs have been modeled using NURBS and structured in feathers using L-Systems [27], and Bezier curves both for the structure of feathers and for individual barbs [27, 202]. Data-driven approaches based on images of real feathers have been able to model not only rachis and barbs but also barbules [8, 9]. Lately, a generative algorithm based on the biological growth of feathers has been able to model realistic feathers with three-dimensional features [252]. These works are orthogonal to our work, and could be used for generating our surface-based representation described in Section 4.6.2.

In terms of appearance modeling, Chen et al. [27] proposed to use data-driven appearance based on BTFs, which require capturing the feather and has limited angular resolution. Leaning and Fagnou [129] and Haapaoja and Genzwürker [78] modeled the scattering from barbs using a hair BCSDF [149], but ignored the scattering from barbules. Harvey and colleagues [82] proposed a data-driven model for the African Emerald Cuckoo, based on measured data and expensive wave simulations. Closest to our work, Huang et al. [95] and Baron et al. [7] accounted for the scattering of both barbs and barbules, with the former modeled statistically. In particular, Huang and colleagues [95] focused on the iridiscent appearance of plumaceous rock dove neck feathers, proposing a BSDF for explicit barbs that simulates barbule scattering using a microfacet-like model with thin-film iridiscense, while accounting for inter-barbule masking. In contrast, our work focuses on general pennaceous feathers. Baron et al. [7] focus on general feathers, representing barbs as curves where the BCSDF combines the scattering of both barbs and barbules, using Marschner's hair model [149]. In contrast, our model is a general, compact surface-based representation for pennaceous feathers, where the scattering of both barbs and barbules is modeled using a new BCSDF that accounts for elliptical cross-sections and inner scattering medulla allowing to represent a wider range of light-matter interactions occurring at barb and barbule scale.

MICROSTRUCTURE-BASED SURFACE SCATTERING As discussed in Section 1.1 statistical aggregated representations of the detail at microscopic level are the common choice for representing the scattering of small features. For surface models, the most common approach are microfacet models [31, 105, 219] that represent rough surfaces as statistical aggregates of tiny planar surfaces, generally assuming uncorrelation for computing the geometric attenuation [86]. Other approaches assume other microscopic scatterer shapes, including spherical scatterers [260], bumps and cavities [155], scratches [16], or micrograins [141], which allow computing specific geometric attenuation terms. Closer to our work, far-field models for woven cloth [101, 182] assume a particular structure between perpendicularly interwoven yarns with fiber-like scattering; this structure allows to derive a closed-form masking functions between the different yarns. Our model also models the scattering of each component using a BCSDF, but incorporates the additional complexity of the multiscale structure of barbs and barbules.

SCATTERING FROM FIBERS Cylinder-based scattering models based on the BCSDF have been proposed for realistic human hair for circular [28, 149, 261] and elliptical fiber cross sections [11, 94, 120]. Follow-up work extended hair models to fur, by adding

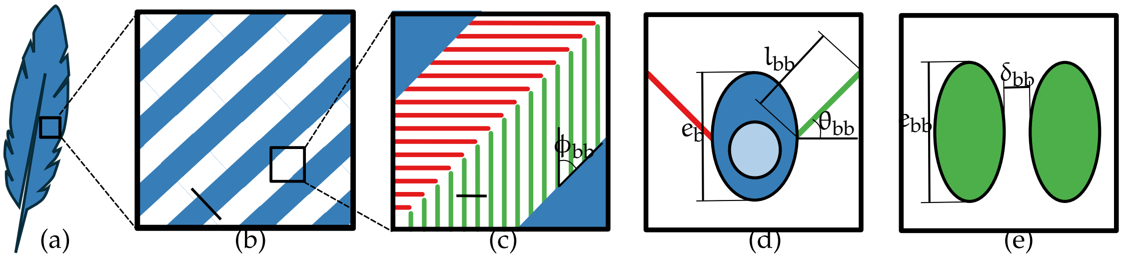

Figure 4.2: (a) The appearance of the feather's vane (a) depends on its underlying structure. (b) Parallel barbs emerge from the rachis. (c) Each barb branches into two sets of barbules, proximal and distal. (d) Barbs are modeled as infinite cylinders with elliptical cross section, with an inner medulla with elliptical cross section. Barbules' thickness is negligible so at this level they are modeled as the plane in which they lie. (e) At a smaller scale, barbules are also cylinders with elliptical cross secion, and can oclude each other. They have some spacing between them that at a larger level is treated as partial transparency that is considered at a larger (barb) scale. The microgeometrical parameters of our model are also represented.

a scattering medulla crucial for the soft look of fur [243, 244]. Cloth fibers and yarns have also been represented using the BCSDF, either measured with goniorreflectometers [182] or simulated from the microgeometry of the fibers [2, 236]. To reduce the cost of rendering individual fibers, aggregation techniques have been proposed to approximate the aggregated appearance in fur [257] and cloth fibers [119].

Our model for fiber scattering builds upon these models, but generalized to the specific structure of the barbs and barbules that compose the feather, combining elliptical cross-sections with a non-centered scattering medulla.

4.3 ON THE APPEARANCE OF FEATHERS

The appearance of feathers is the result of a complex light-matter interactions at multiple scales. It depends on both internal structure and chemical composition. The structure and type of feather affect its appearance at a high level, while the chemical composition and micro-nanoscale structure fundamentally change its coloration, opacity and reflectance.

Structure

Although feathers are incredibly diverse in shape and size, they are all composed primarily of keratin, a protein present in all epidermal structures of all vertebrates. Unlike the epidermal structures of mammals (e.g. hair or nails) whose main component is *α-keratin*, the structures of birds and reptiles are based on *β-keratin*, which makes them more rigid. This protein is present in all the epidermal structures of birds (beak, claws, feathers ...) in different ways depending on the corresponding biological functions. The microstructure of *β-keratin* on feathers is, biomechanically, more flexible, providing unique properties required for flight [174].

Feather structure

The main components of a typical feather consist of a long and central axis in which two parts are considered: the lower and wider part, which is inserted on the skin (calamus) and the rest of the axis (rachis). On both sides of the rachis, the vane grows as a lamina and is formed by barbs, which grow along the entire rachis in a branching manner. Each barb consists of a central axis of tens of microns (ramus), which branches out on both sides in rows of hundreds of microscopic-sized barbules (proximal and distal), coupled by hooklets [82, 205]. Barbules are up to an order of magnitude smaller than barbs [95], and branch from their respective barb, often in a $\phi_{bb} = 45^\circ$ angle to maximize their overlap, favoring the grip of the hooklets [40]. This ramification results in features at multiple scales (see Figure 4.2). Previous work [9] collected data from ornithological studies including barb density, barb angles or barbule length. There is an important variety of geometric structures and configurations for barb and barbules depending on the bird species and even sex [152]. For instance, some cross sections such as the barbule cross-section shapes of a peacock tail feather (see Fig. 3 [51]) or a Lawes' parotia breast feather (see Fig. 3 [196]) with ridged forms. Nevertheless, the cross sections of these internal structures tend to be elliptical (see Fig. 4 [152]) as we show in the Scanning electron microscopic (SEM) images of the kingfisher feather Figure 4.3 (c), with eccentricity up to 6 depending on the specie [99]. For this reason, as previous work [95], we approximate these cross sections as ellipses.

Feather taxonomy

Feathers can be classified according to size, shape, and type [203]. One of the main features is the type of vane: plumulaceous (fuffy, diffuse and soft) or pennaceous (mostly flat, firm and rigid, due to interlocked barbules). In this work, we focus on pennaceous feathers, as these ones can be modeled with a BSDF over a surface that represents the vain. A surface representation allows us to encode the feathers as texture, a representation quite effective especially for far field scenarios. On the other hand, plumulaceous feathers require the modeling of individual barbs over a volume. For a further discussion about the feather taxonomy and potential representations, we refer the reader to previous work specialized on feather geometry [8, 203].

Coloration

Feathers combine pigmentation and structural coloration to create extraordinary and diverse colors, more than for any mammal and most vertebrates [91]. The colors of the bird feathers play an essential role in a variety of biological functions such as camouflage, visual signaling or mating [180].

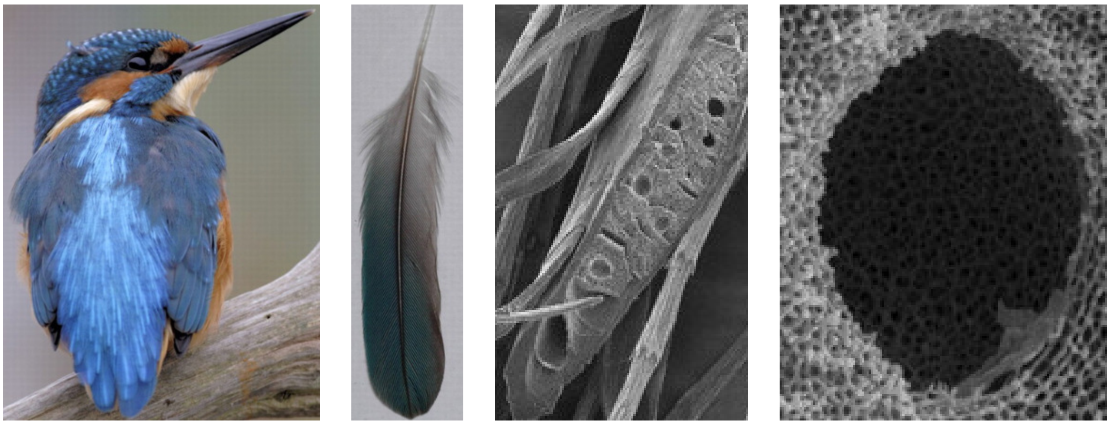

(a) (b) (c) (d)

Figure 4.3: (a) Kingfisher (*Alcedo atthis*). (b) Blue tail feather of a kingfisher. (c) Top-view SEM of a sectioned barb of the feather with several barbules (50 μm). (d) Cross section SEM of a cut vacuole and the surrounding spongy structures (2 μm). Notice the elliptical shape of the barb cross section and the presence of the medulla with an internal spongy structure (Source [197]).

Pigmentation

Pigment-based coloration is the most common in vertebrates and is caused by the absorption of certain wavelengths as light traverses the pigments (Figure 4.4, left). Pigments are molecules that absorb and reflect each of the wavelengths of light, hence they can be associated with a particular color. These pigments are deposited between the keratin sheets that make up the barbs, barbules and rachis [180]. The most prevalent pigments in feathers are based on melanin, which is responsible for a wide range of colors, usually brown, gray and black tones [187] or carotene, that produce more vivid colors such as reds, yellows, and oranges [153, 180, 187, 217].

Structural Coloration

The structural coloration of bird feathers is caused by the scattering and interference of light with the feather nanoscopic structure [52]. This coloration does not depend entirely on absorption, but on the way the structure of the feathers disperses certain wavelengths. One source of such structural coloration are specialized keratin structures in the medulla of the barbs, which create a pseudo-ordered matrix of keratin and air bubbles, known as the spongy layer (Figure 4.3, d). These lead to non-iridescent tones of blue, violet and UV [179], as illustrated in Figure 4.4, although in some cases they generate partial iridescence for highly directional illuminations [197], which disappears with lower frequency illuminations. We account for this effect with a diffuse medulla, that approximates the backscattering given by constructive interference in such quasi-ordered structures [169].

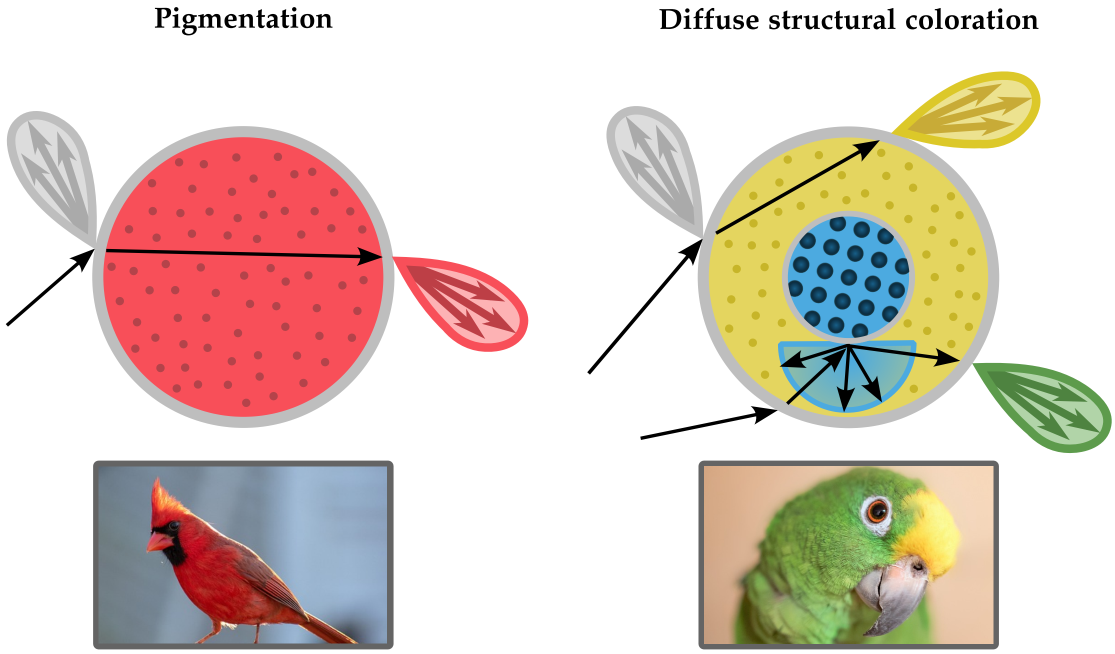

Figure 4.4: Cross-section fiber schematic of the coloration mechanisms in feathers supported by our appearance model. *Left*: Pigments-based coloration where some wavelengths are absorbed by the pigment granules located in the keratin matrix. *Right*: Diffuse structural coloration, where in addition to pigmentation, the medulla scatters colored light due to interference.

Hybrid coloration

Some colors such as blues, greens and purples are created by the combination of the structure and the pigments. For instance, the green color is usually a combination of yellow pigmentation with a blue structural color [153] (Figure 4.4, right). Feather goniochromatism can occur from the partial occlusions between barbs and barbules of different colors at different view orientations, which we model through a masking expression.

4.4 OVERVIEW

The scattering of each barb and barbule is modeled as a *bidirectional curve scattering distribution function* (BCSDF) [258], that assumes that each fiber is an infinite locally-straight cylinder (Section 4.5). The aggregated appearance comes from the combination of such BCSDFs in their local frame, as obtained from the tangent map and the geometrical parameters of the BSDF. The scattering radiance from each BSDF is weighted based on the respective projected visible area as defined by our masking term (Section 4.6). Masking is computed on the fly when our feather BSDF is evaluated or sampled. The structural and optical parameters of our model can affect either particular BCSDFs (for barbs or barbules), the global BCSDF or both. These are summarized in Table 4.1 and some of them are illustrated on Figure 4.2. The feather's geometry is encoded in a 2D texture applied on a plane or curved plane: The red channel encodes the barb orientation and the blue channel stores a flag identifying rachis, vane, and background. We manually designed the texture to roughly match the feathers' silhouette based on feather photographs. More details about the authoring process can be found in the supplemental document.

Parameter	Definition
β_m	Longitudinal roughness of cortex
β_n	Azimuthal roughness of cortex
η_c	Refractive index of cortex
a_c	Cortex semi-major axis
b_c	Cortex semi-minor axis
d_m	Diffuse reflectance of medulla
η_m	Refractive index of medulla
a_m	Medulla semi-major axis
b_m	Medulla semi-minor axis
$(c_{m,a}, c_{m,b})$	Medulla center (w.r.t. to cortex center)
ϕ_b	Azimuthal barb angle
ϕ_{bb}	Azimuthal barbule angle
θ_{bb}	Longitudinal barbule angle
δ_{bb}	Barbule separation
l_{bb}	Barbule length

Table 4.1: Parameters of our Feather BSDF. The angles (ϕ_b, ϕ_{bb}, θ_{bb}) are expressed in degrees. The sub-index b refers to barbs, while the bb refers to barbules. l_{bb} and δ_{bb} are measured in number of barb and barbules per length unit respectively. The eccentricity for e_b and barbules cross section e_{bb} are computed from the cortex axes a_c and b_c of the corresponding internal structure.

We introduce the following assumptions:

1. No wave-optical effects are present between barbs and barbules, and all scattering can be considered in the ray-optics regime. Wave optics is particularly significant for iridescent feathers such as the hummingbird and peacock feathers. In this work, we focus on pennaceous non-iridescent feathers which extend a large group of feathers.
2. As previous work [95], we assume an elliptical cross-section of both barbs and barbules. While this is mostly accurate for barbs, barbules comes in different shapes which is specially relevant for iridiscent feathers [82].
3. Multiple scattering only occurs inside the fibers (accounted by the BCSDF), but not among barbs and barbules on the same surface. When rendering (path tracing), multiple scattering emerges from the interactions among different surfaces.

4. We assume that shadowing has a negligible effect, which allows our model to compensate for energy loss due to the lack of multiple scattering among barbs and barbules.
5. The microstructure of the feather's vane is locally regular: In a differential surface patch, barbs are parallel, proximal barbules are parallel and distal barbules are parallel. This removes the need of explicitly modeling the geometry of all barbs and barbules.

In the following, we first define our new BCSDF that supports some of the observed structural properties of both barbs and barbules; then, we define our surface BSDF that leverages the BCSDFs for modeling the aggregated scattering at the feather's vane.

4.5 SCATTERING FROM BARBS AND BARBULES

Here we model the scattering of the main individual components of the feather as a BCSDF. Then, we describe our structural model for barbs and barbules; then, we describe the BCSDF used for modeling the three of them. Following the formulation from [149], we define the far-field BCSDF in polar coordinates, as

$$S(\phi_i, \theta_i, \phi_o, \theta_o, h^-, h^+) = \frac{D\cos(\theta_i)^{-2}}{h^+ - h^-} \int_{h^-}^{h^+} S(\phi_i, \theta_i, \phi_o, \theta_o, h) dh, \tag{4.1}$$

where (ϕ_i, θ_i) and (ϕ_o, θ_o) parametrize in polar coordinates the incident and outgoing direction respectively, D is the fiber width, $\cos(\theta_i)^{-2}$ accounts for the projected solid angle of the specular cone, h^- and h^+ parametrize the limits of the visible fiber diameter (for a fully visible fiber, $h^- = -1$ and $h^+ = 1$), and $S(\phi_i, \theta_i, \phi_o, \theta_o, h)$ is

$$\begin{aligned} S(\phi_i, \theta_i, \phi_o, \theta_o, h) = \iint & \mathcal{S}(\phi_i, \theta_i, \Phi, \Theta, h) G_M(\Theta - \theta_o|\beta_m) \\ & G_N(\Phi - \phi_o|\beta_n) \cos(\Theta) d\Theta d\Phi, \end{aligned} \tag{4.2}$$

where $G_M(\Delta\theta|\beta_m)$ and $G_N(\Delta\phi|\beta_n)$ are longitudinal and azimuthal Gaussian detector functions respectively [261], that regularize the scattered field accounting for surface roughness (parametrized by β_m and β_n, respectively). Finally, $\mathcal{S}(\phi_i, \theta_i, \Phi, \Theta, h)$ models the transfer function inside the fiber, defined as the integral of paths starting at (ϕ_i, θ_i, h) and outgoing at direction (Φ, Θ) following the path integral [215]

$$\mathcal{S}(\phi_i, \theta_i, \Phi, \Theta, h) = \int_\Omega f(\mathbf{x}) d\mu(\mathbf{x}), \tag{4.3}$$

with Ω the space of paths $\mathbf{x}$ that start at (ϕ_i, θ_i, h) and end at (Φ, Θ), and $\mu(\mathbf{x})$ the measure of the integral. For purely absorbing fiber, the space of paths starting at (ϕ_i, θ_i, h) is singular specular path, and thus $\mathcal{S}(\phi_i, \theta_i, \Phi, \Theta, h)$ is a sum of impulse functions, one per bounce inside the fiber, with amplitude the attenuation at each bounce.

4.5.1 *Fiber models*

We do not model the scattering from the rachis as a BCSDF; instead, we assume that the diffuse scattering medulla is dominant and model it as a colored diffuse surface.

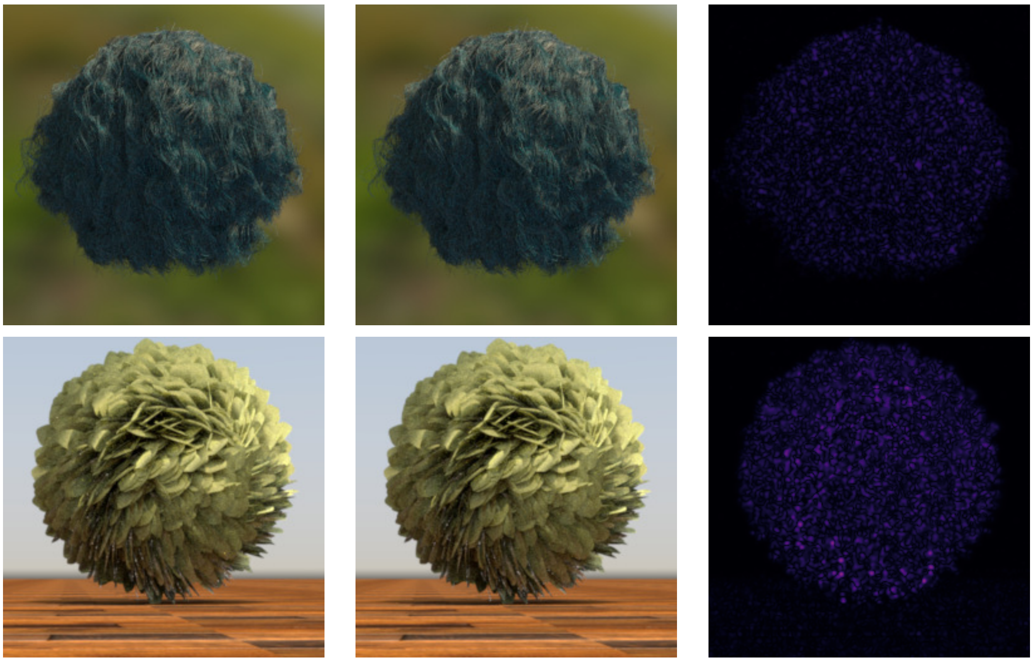

1000 lobes 5 lobes Error

Figure 4.5: Energy conservation validation of the BCSDF. For a five-bounces evaluation of our BCSDF we observe minimal differences with respect to a 1000-bounces counterpart, in both fiber-only, single scattering test (top) and inside our feathers BSDF with multiple scattering (bottom). Error is computed using the FLIP metric [4].

We model barbs and barbules as cylindrical β-keratin cortex with elliptical cross section, with axes a_c and b_c. The interface is a rough dielectric, with roughness modeled as azimuthal and longitudinal 1D Gaussian detectors as described above), and index of refraction η_c. The β-keratin hosting medium is filled with absorbing pigments, leading to an exponential transmittance as predicted by Beer-Lambert law.

Additionally, the barb contains a medulla. As opposed to hair and fur models [244], the medulla is generally not aligned with the cortex center, nor have the same eccentricity. We thus do not model cortex and medulla as concentric cylinders, but instead allow the medulla to position freely inside the cortex, with axes a_m and b_m and center $(c_{m,a}, c_{m,b})$. As anticipated before, we approximate the reflectance of the medulla as a diffuse surface. The albedo ranges from white (when there is no spongy layer) to blue, representing the non-iridiscent coloration of the spongy layer (as described before). We use measured reflectances for the blue tint, from the work by Noh et al. [169].

4.5.2 *Rendering*

Our model for barbs and barbules involves elliptical cross sections and a diffuse medulla (for the case of barbs), so our solution is not simple enough for a closed-form expression for Equation (4.1), specially as we account for higher-order bounces inside the fiber, so we need to solve the integral numerically. Inspired on the work of

Chiang et al. [28] we leverage the stochastic nature of modern renderers and compute the nested integrals in Equation (4.1) with a Monte Carlo estimate.

We implement a stochastic Monte-Carlo based path tracer from which we obtain a set of fixed number of *lobes*. In our experiments, we observed that five lobes seem to be accurate enough as we shown in Figure 4.5. Each of these lobes consists of an outgoing radiance value (throughput) associated to an outgoing direction. Each interaction with each dielectric interface (cortex) deterministically generates two rays: one exiting the cortex, which we store as a lobe, and one entering the fiber, which keeps interacting and generating lobes until the final number is reached. The stochastic exploration of those sets of lobes comes from the integration variable of Equation (4.1) h and, in the case of barbs, the interactions with the medulla, from which we generate new rays using cosine sampling.

For evaluating the BCSDF, we go through all the lobes, evaluating the Gaussian detector functions $G_M(\Theta - \theta_o|\beta_m)$ and $G_N(\Phi - \phi_o|\beta_n)$ (see Equation (4.2)) for the direction of the lobe (Θ, Φ), and $\mathcal{S}(\phi_i, \theta_i, \Phi, \Theta, h)$ is approximated to the lobe's throughput. For sampling the BCSDF, we follow a strategy similar to previous works [261] where we build a discrete pdf where the accumulated throughput of each lobe corresponds to the probability of such lobe.

4.6 OUR SURFACE APPEARANCE MODEL

As opposed to previous works [95] that explicitly model the rachis and barbs as curves, we represent feathers using a surface-based representation, where the geometry explicitly model the rachis and vanes, while barbs and barbules in the vanes are modeled as microgeometry, using a BSDF defined in the local coordinate system of the barbs. We encode the feather parameters in texture space, using a mask to distinguish between rachis and vanes, and defining the local tangent direction on the surface using a rotation angle. This removes the need of modeling explicit geometry of the feather, resulting in a compact representation.

Our BSDF $f_s(\omega_i, \omega_o)$ is a linear combination of four different components: barb BCSDF S_b and proximal and distal barbules BCSDFs S_{bb} modeled using Equation (4.2) (below in Cartesian coordinates instead of polar), and transmittance through the vane. The BSDF is defined in the local coordinates of the barb, defined by the normal of the surface and the barb's tangent direction, following

$$\begin{aligned} f_s(\omega_i,\omega_o) = \Big[& w_b(\omega_o)\, S_b\left(\omega_i, \omega_o, h_b^-(\omega_o), h_b^+(\omega_o)\right) G(\omega_i) \\ & + w_{bp}(\omega_o)\, S_{bb}\left(\mathbf{T}_{bp}\omega_i, \mathbf{T}_{bp}\omega_o, h_{bp}^-(\omega_o), h_{bp}^+(\omega_o)\right) G(\mathbf{T}_{bp}\omega_i) \\ & + w_{bd}(\omega_o)\, S_{bb}\left(\mathbf{T}_{bd}\omega_i, \mathbf{T}_{bd}\omega_o, h_{bd}^-(\omega_o), h_{bd}^+(\omega_o)\right) G(\mathbf{T}_{bd}\omega_i) \\ & + w_t(\omega_o)\, \delta(1 - \omega_i \cdot \omega_o)\Big] |\omega_i \cdot n|^{-1}, \end{aligned} \tag{4.4}$$

where n is the surface normal, $w_b(\omega_o)$, $w_{bp}(\omega_o)$ and $w_{bd}(\omega_o)$ model the projected area of the barb, proximal barbule and distal barbule as viewed from ω_o, respectively, $w_t(\omega_o) = 1 - (w_b(\omega_o) + w_{bp}(\omega_o) + w_{bd}(\omega_o))$ is the amount of transparency, $\delta(\cdot)$ is the Dirac delta distribution, $\mathbf{T}_{bp}$ and $\mathbf{T}_{bd}$ transform the coordinate system to the frame of proximal and distal barbules, according to their rotation $\phi_{bb} = 45°$ and inclination θ_{bb} (parameter of our model), and $G(\omega_i) = \cos(\theta_i)$ the foreshortening over the fiber. The integration ranges for the corresponding BCSDFs are $[h_b^-(\omega_o), h_b^+(\omega_o)]$, $[h_{bp}^-(\omega_o), h_{bp}^+(\omega_o)]$ and $[h_{bd}^-(\omega_o), h_{bd}^+(\omega_o)]$. Both the projected areas $w(\omega_o)$ and the integration ranges (h^-, h^+) for barb and barbules are given by our masking term (Section 4.6.1). Finally, note that for all BCSDFs we set the fiber width $D = 1$, since it is accounted by the projected areas w_b, w_{bp} and w_{bd}, respectively.

4.6.1 *Masking*

We do not explicitly model this microgeometry, but devise an analytical masking expression that analyzes the visibility among its components. Our masking expression works at two different scales: Barbs and barbules. The key insight for developing our masking components is that, at their respective local coordinate systems, all the microgeometrical elements are either ellipses or segments. Both are governed by implicit and parametric equations from which we can derive both projected areas within a parameter subrange. Furthermore, instead of stochastically explore the visibility, as the work by [95] does for barbules, we find analytical points that correspond to visibility discontinuities by tracing 2D rays at the analytical boundaries, and from the arcs/segments between the corresponding intersections we calculate projected areas.

We name $\omega_o' = \{\omega_{ox}', \omega_{oy}'\}$ the two-dimensional direction in the local coordinate space (the longitudinal dimension is ignored) after transforming ω_o. For an ellipse (representing barb or barbule) with center $\mathbf{c} = \{c_x, c_y\}$ and axes a_x and a_y the visibility discontinuities as given by that ellipse are its tangent points, obtained as:

$$\mathbf{p}_{t0} = \begin{pmatrix} c_x + a_x\cos(\psi) \\ c_y + a_y\sin(\psi) \end{pmatrix} \qquad \mathbf{p}_{t1} = \begin{pmatrix} c_x + a_x\cos(\psi+\pi) \\ c_y + a_y\sin(\psi+\pi) \end{pmatrix} \tag{4.5}$$

where $\psi = \tan^{-1}\left(\frac{-a_y\omega_{ox}'}{a_x\omega_{oy}'}\right)$.

For barbs, the corner among barbule segments $\mathbf{p}_c$ represents a visibility discontinuity, too. These visibility boundary points are the origin of 2D rays traced towards the rest of the elements of the microgeometry at their specific scale. Once the visibility among the different geometrical elements, projected areas towards ω_o' are calculated through a simple cosine for the case of segments and for the case of ellipses, its projected area A is:

$$A(u_0, u_1) = \omega_{ox}' a_y(\sin(u_0) - \sin(u_1)) + \omega_{oy}' a_x(\cos(u_1) - \cos(u_0)) \tag{4.6}$$

where u_0 and u_1 are the starting and ending parameters from the ellipse's parametric equation, obtained by inverting the parametric equation for the intersection (visibility boundaries) points. The derivation of this inversion and of Equation (D.13) are on the supplemental material.

Barbule masking At a smaller scale, which we call *barbule masking* we align the reference frame at the local barbule frame. Given the cross section of barbules (ellipses) and a separation among them, we obtain the ratio between barbule visibility and trans-

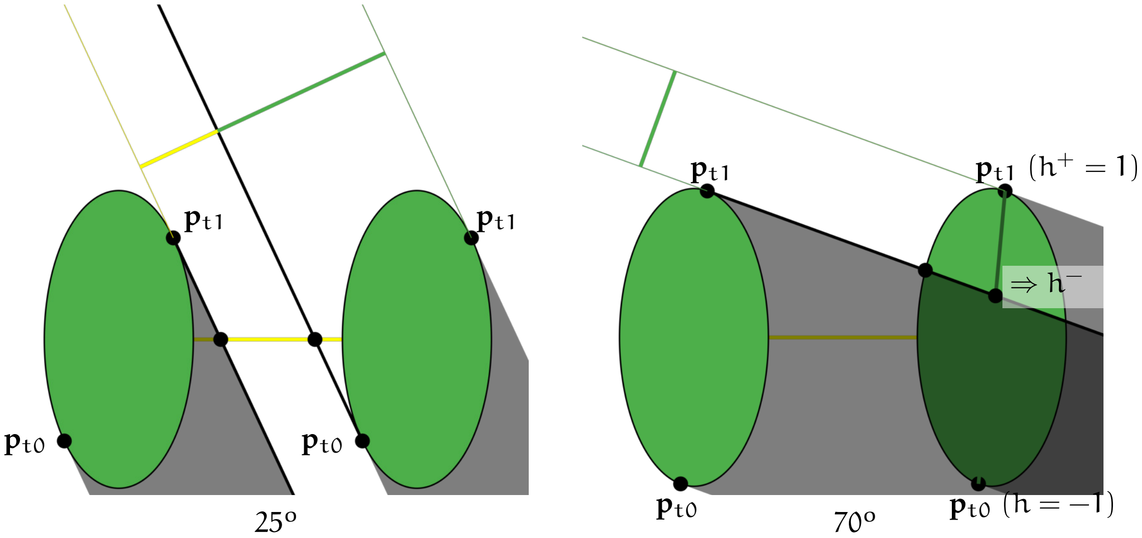

Figure 4.6: Masking between two barbules at different view inclinations. Barbules (and their projected area) are marked in green. The barbule separation is modeled as the yellow segment. The traced rays are the black lines. We first trace a ray from the tangent point $\mathbf{p}_{t1}$, generating two cases. If it intersects the other ellipse (70° case) we need to intersect the ellipse's diameter to identify $\mathbf{p}_d$ and therefore h^-. The separation (transmittance) cannot be seen ($w'_{bb} = 1$). If it does not intersect the other ellipse (25° case) then we need to intersect the separation segment with two different rays. We then calculate the local barbule weight w'_{bb} from the projected areas a_{bb} and a_s.

mittance, as well as the casting segment range for the barbule (see Figure 4.2 and Figure D.1). This is done at the local coordinate system of proximal and distal barbules, and the result is used for the larger-scale barb masking. Note that previous work [95] models this masking term stochastically, and not analytically.

We explicitly model two barbules as ellipses while the separation between them is a segment (as represented in Figure D.1). The barbule model parameters are, for reducing the number of parameters, the barbule axis ratio e_{bb} and the relative barbule separation δ_{bb}. We trace 2D rays from a tangent point at each ellipse towards each other and the separation segment. If there is occlusion between ellipses (ray has intersected) the lower limit of the integration range h^- ($h^-_{bp}(\omega_o)$ or $h^-_{bd}(\omega_o)$) require tracing the same ray again towards the segment between $\mathbf{p}_{t0}$ and $\mathbf{p}_{t1}$ that represents the diameter of the ellipse. h^- is the parameter of the parametric equation of this diameter at the intersection point $\mathbf{p}_d$ (if there is no intersection, $h^- = 1$). The other integration range limit is $h^+ = 1$. These geometrical elements can be visualized on Figure D.1. This barbule masking term is applied twice, for both proximal and distal barbules, obtainig h^-_{bp}, h^+_{bp}, h^-_{bd}, h^+_{bd} and, from the projected areas, local weights w'_{bp} and w'_{bd}. As barbules are separated, there is also the possibility of local transparency with weights $1 - w'_{bp}$ and $1 - w'_{bd}$, respectively .

Barb masking at the upper scale, *barb masking*, we align the reference frame to the barb local frame, so the cross section of the barb is an ellipse, and the cross sections of barbules are two segments (distal and proximal), with partial transparency obtained from the barbule masking component. Here, we compute the visibility ratios between barbs, barbules, and transmittance (see Figure D.2 and Figure 4.2, d). Note that previ-

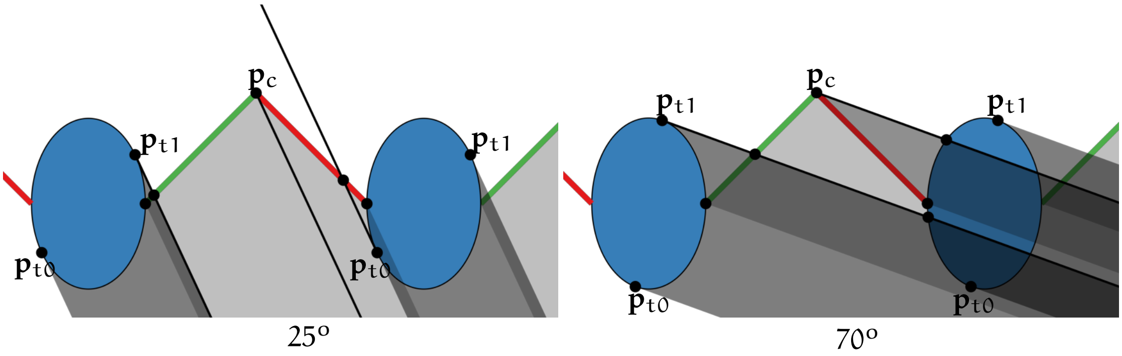

Figure 4.7: Cross section of barbs, representing the masking between barbs at two different view inclinations. Barbules (in red and green, respectively) are partially transmitting, depending on the view direction at their particular local coordinates (see Figure D.1), while barbs (in blue) are considered to be opaque. Depending on the view direction, each element (barbs and barbules) totally or partially occludes the rest. The limits of such occlusions are identifying by tracing 2D rays (marked in black). By considering all the particular intersection ranges we obtain the weights.

ous work [95] does not provide a masking term that accounts for barbs and need to explicitly model the geometry of each individual barb.

Parameters are also relative with respect to the horizontal size of the barb, and include axis ratio e_b, barbule length l_{bb} and barbule inclination θ_{bb}. From these we obtain the a representation of the microgeometry with two ellipses (barbs) and two pairs of segments (barbules). We trace rays form the visiblity discontinuity points $\mathbf{p}_{t0}$, $\mathbf{p}_{t1}$ and $\mathbf{p}_c$ towards the rest of geometrical elements (segments and ellipses) and, from the ranges between intersection points, calculate the projected areas as before. Each interval correspond to total or partial oclussions (at this scale, barbules are partially transparent) so each projected area must be multiplied by $1 - w'_{bp}$ if it is occluded by the proximal barbule and by $1 - w'_{bd}$ if it is occluded by the distal barbule. We combine all projected areas and obtain w_b, w_{bp} and w_{bd}. We also obtain h_b^- and h_b^+ using the same procedure than for barbules. A detailed step-by-step procedure for obtaining these masking terms can be found as supplemental material.

4.6.2 *Rendering*

For evaluating the BSDF as described by Equation (4.4), we simply calculate the weights from the masking expression and then evaluate the corresponding BCSDFs at their local coordinate systems. For sampling it, we use a discrete distribution, similar to previous work [7] with four possible events (hit barb, hit proximal/distal barbules and delta transmittance) whose probabilities come from weights obtained from the analytical masking expression. The selected event is then sampled, deterministically if it is a delta transmittance or by using the BCSDF sampling routine of the specific fiber otherwise.

Figure	# Feathers	Resolution	Spp	Time
Figure 4.8	2500	512 × 256	1024	2.2 min
Figure 4.9	82	730 × 420	256	4.4 min
Figure 4.13	496	512 × 512	256	1.5 min
Figure 4.11	1	256 × 512	1024	1.5 min

Table 4.2: Rendering time for different scene complexities: single feather, feather pelt, feather wing and feather ball.

4.7 ANALYSIS AND RESULTS

In this section, we perform several experiments to show the capability of our feather BSDF, formal validation of our analytical masking term and comparison to previous work and photographs. We author each feather manually, by encoding the feathers geometry, including shaft, vane and barb orientation, on a 2D texture on top of a relatively simple geometry. Exploring more effective procedural tools for authoring (e.g., [8]) is left as an interesting future research direction.

We implemented our model in Mitsuba 0.6 [102] as a new BSDF. All our renders have been computed on an Intel Core i9-10900KF CPU with 20 cores. Table 4.2 reports the rendering time of all Figures including scene with different complexity in terms of geometry (single feather, feather pelt, feather wing and feather balls), material (brown, black, blue, red and green feathers) and lighting conditions. Since our is surface-based, it is practical in both terms of rendering time and memory for fully feathered assets, which would be very expensive is curve-based feathers. We plan to release the source code and the scenes to try to encourage future work in this topic.

4.7.1 *Model analysis*

We perform ablation studies to analyze the impact of each component of our model. Figure 4.8 shows how increasing the complexity of the fiber cross section affects the final appearance produced by our fiber BCSDF in a significant manner. As previous work had already shown [120], an elliptical cross section with higher eccentricity increases the brightness of the specular lobes (see renders) and introduces new patterns as the elliptical blobs (see latlong reflectance plots). As we introduce a medulla, reflectance distribution gets more uniform and appearance gets smoother. When introducing a non-concentric medulla, reflectance distribution becomes asymmetric and the change in appearance with respect to the concentric case is important, showing a greater variation on appearance with respect to feather orientation.

Figure 4.9 shows the effect of both our masking and our BCSDF representing feathers, compared against using only hair barbs [78, 129] (first column), and against a more sophisticated feather structure with barbs and barbules modeled with a hair model (second column) and combined with a masking term, similar to the one proposed by Baron and colleagues [7]. As shown in the four examples, the medulla (third column) is

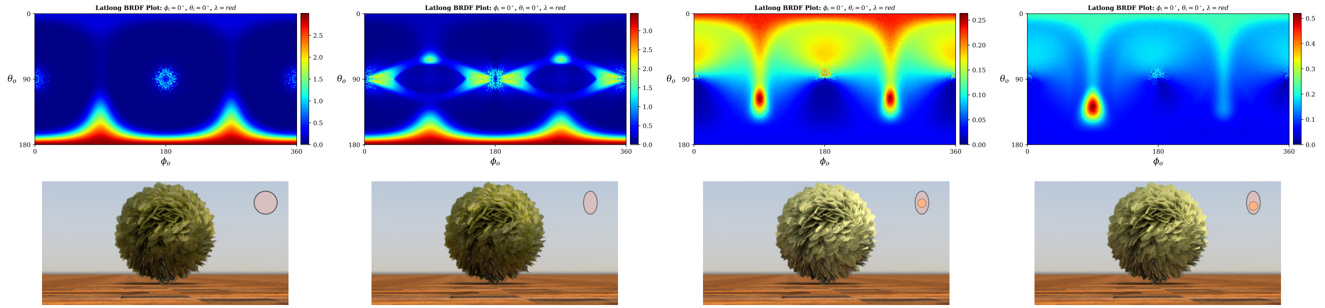

Circular Cortex Elliptical Cortex Concentric Medulla Non-Concentric Medulla

Figure 4.8: Ablation studies of the Fiber BCSDF at (ϕ_i = 0°, θ_i = 45°): single cortex cylinder (first column), elliptical cortex cylinder (second column), nested cylinders with concentric medulla (third column) and nested cylinders with non-concentric medulla (last column). We show the latlong reflectance plots for the red wavelength. Fiber BCSDF parameters: $\sigma_{c,a} = (0.133, 1.967, 2.1)$, $\eta_c = \eta_m = 1.53$, $\beta_m = 0.13$, $\beta_n = 0.12$. The elliptical cross section is $e_b = 1.6$. Medulla is circular with radius 0.6 ($a_m = 0.6, b_m = 0.6$) and the non-concentric medulla is located at $c_x = -0.1$ and $c_y = 0.1$. Latlong reflectance plots are obtained with 10000 spp. Notice how the elliptical cortex increases the reflectance of the render, while the presence of the medulla makes the appearance softer. The latlong reflectance plots also reveals how each component affects the lobe patterns in a significant manner.

essential to produce softer and non-iridescent structural colors such as green and blue tones. Notice that a closer appearance matching would require to match the lighting conditions and feather geometries more carefully. Nevertheless, the overall appearance produced by our full model better reproduces the real feathers' appearance than previous works.

Finally, in Figure 4.10 we validate our analytical masking term by comparison against a brute-force 100spp Monte-Carlo simulation of masking over explicit geometry (consisting of 100 barbs with 100 barbules each). Our masking expression yields similar weights to the reference, but without any noise and in constant time.

4.7.2 *Appearance exploration*

In Figure 4.13 we analyze the feather expressivity of our appearance model by exploring the range of appearances that it can achieve. In particular: 1) Increasing the relative barbule separation (δ_{bb}) increases feather transparency as less transmitted rays are blocked by barbules. 2) Increasing barb axis ratio (e_b) increases the frequency of the highlights. Finally, 3) increasing the medulla size changes the hue of the feather, as well as the distribution of reflectance towards a more diffuse one. We provide a more thorough parameter exploration in Appendix D.3.

4.7.3 *Appearance matching with photographs*

We photograph two feathers using a conventional smartphone under different light conditions, and create an scene to roughly match lighting conditions, appearance and shape of the captures, while ignoring the plumaceous parts of the feather.

Figure 4.11 shows our results with the Amazon parrot, demonstrating that our masking is able to predict the view-dependent changes (top near frontal view, bottom rotated feather) of the final color due to different coloration between green barbs and yellow barbules.

Figure 4.12 shows a black goose feather lighted from behind. Similar to before, we capture the feather under varying rotation: Without masking, the transparency is independent of the rotation, while our masking predicts the loss on transparency as the feather rotates. In addition, this figure demonstrates again the importance of the medulla to predict the blue tint of the feather due to structural coloration of feathers, which is not physically possible with simpler hair BCSDF as in previous works.

Note that we do not expect to perfectly match the photographs as they were been taken under uncontrolled illumination conditions and the structural and optical parameters of the feather were unknown. In addition, near-field details such as visible individual barbs are not accounted by our model. Still, we show that the combination of our BCSDF and our masking term predict challenging appearance features observed in feathers.

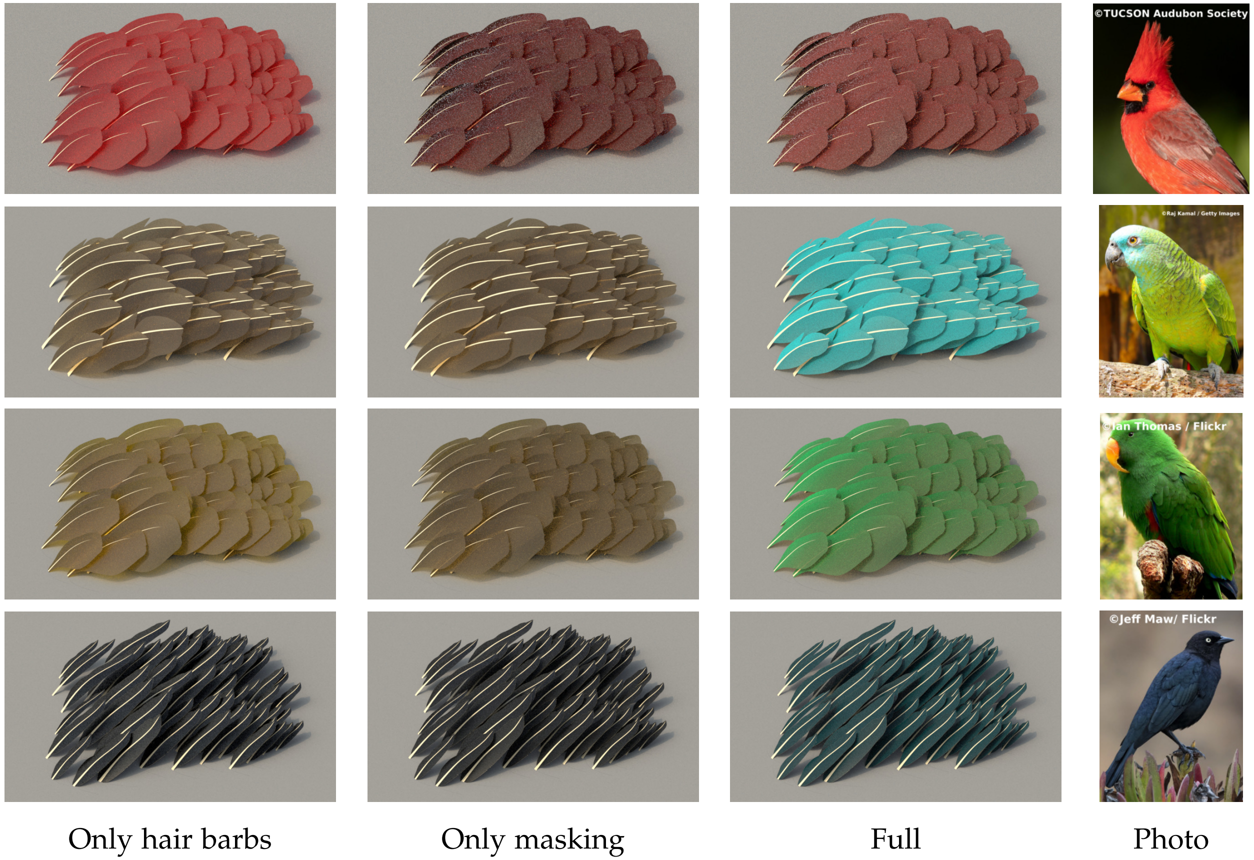

Figure 4.9: Ablation studies of our feather BSDF for a feather pelt scene. [Only hair barbs]: Only barb with hair BCSDF [149]; [Only masking]: Barb and barbules with Hair BCSDF combined with our masking term similar to [7]; and [Full]: Adding a diffuse medulla inside the Barb BCSDF. Notice how our full model is capable of producing a wide variety of complex appearances that resembles closer the appearance of the photographs than previous works for several birds: From top to bottom, northern cardinal, blue-fronted amazon parrot (we reproduce the blue feathers in the face), electus parrot and Brewer's blackbird. In particular, the diffuse reflectance of the medulla is critical to achieve similar tones to the photographs, while the masking is important for goniochromatic and occlusions effects.

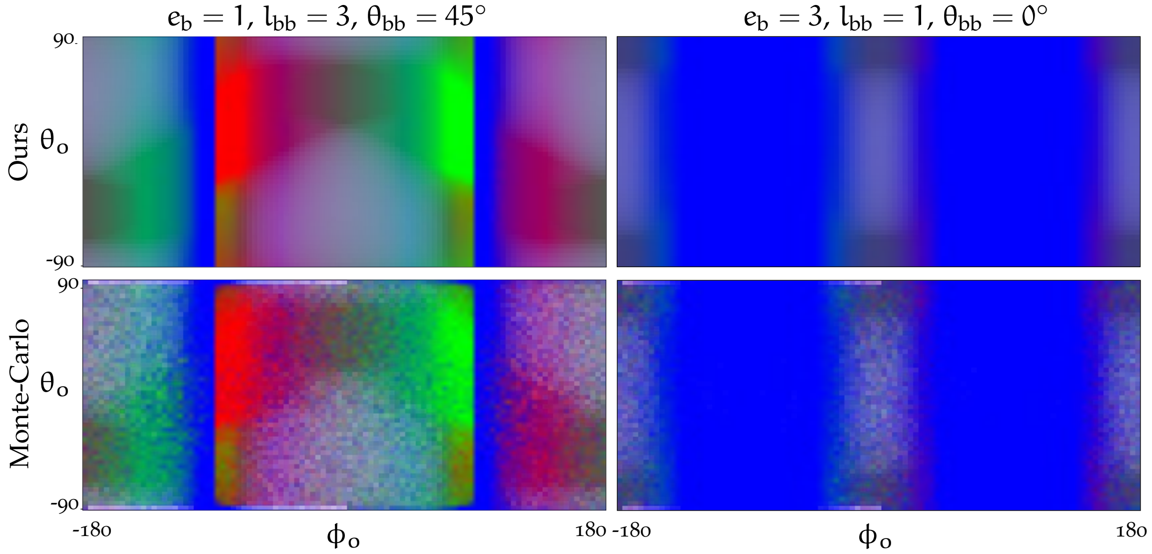

Figure 4.10: Comparison between our analytical masking expression (top) and a Monte-Carlo simulation (100spp) with ray tracing towards a full modeled microgeometry for the same parameters (bottom). Two different masking parametrizations have been explored (left and right columns), and the rest of the parameters are the same between the two: $e_{bb} = 1$, $\delta_{bb} = 1$. The horizontal and vertical axes per plot correspond to the view direction ω_o in polar coordinates (θ_o, ϕ_o). The colormap represents the weights of the different microgeometrical components: w_b, w_{bp} and w_{bd} are mapped to the blue, green, and red channels, respectively. Note how our analytical masking expression is able to obtain a noiseless result in a minimal fraction of the time (constant) while being accurate with respect to a full modeled microgeometry.

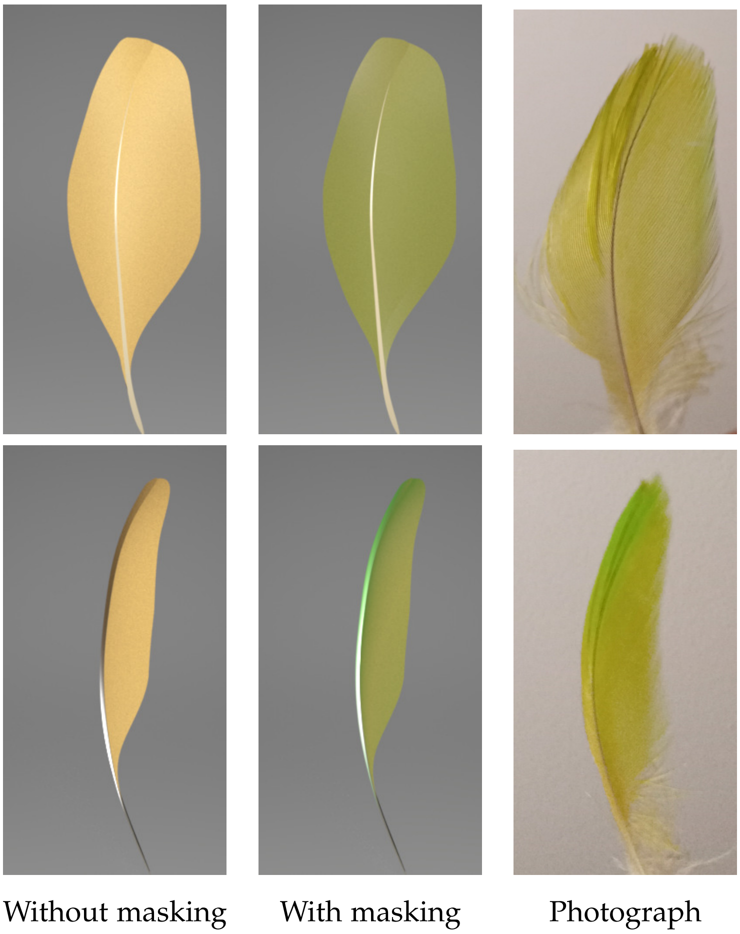

Figure 4.11: Qualitative appearance matching on an *Amazon parrot* feather, for a frontal (top) and lateral (bottom) views. As the feather rotates, view dependent changes on the feather's color become apparent: These are produced by visibility changes between barbs (yellow) and barbules (green). Our masking model is roughly able to predict these changes. Please see the supplemental video for a dynamic example of this goniochromism effect.

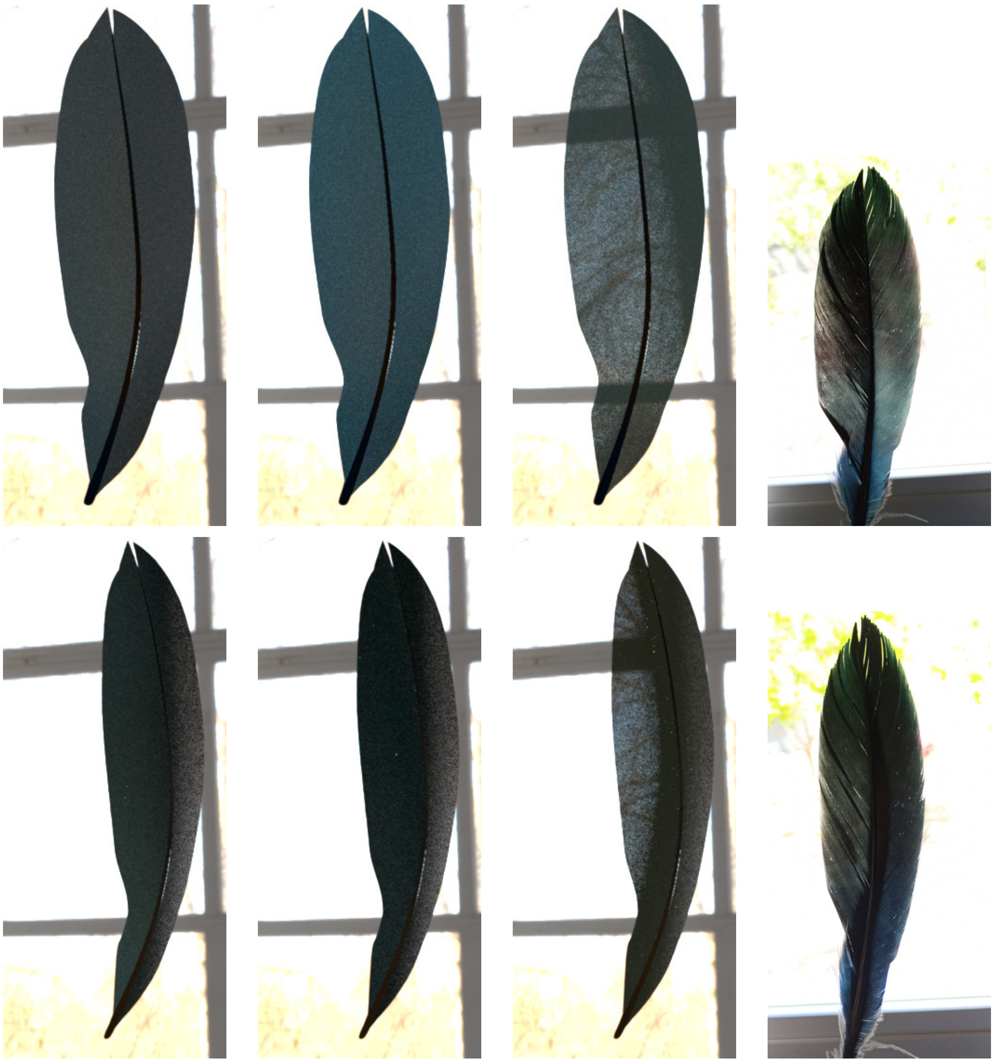

Figure 4.12: Qualitative appearance matching on a *black goose* feather, for a frontal (top) and lateral (bottom) views, under strong back-lighting. Our masking term can reproduce view-dependent transparency. Note that the bluish appearance can only be achieved by the inclusion of a blue medulla.

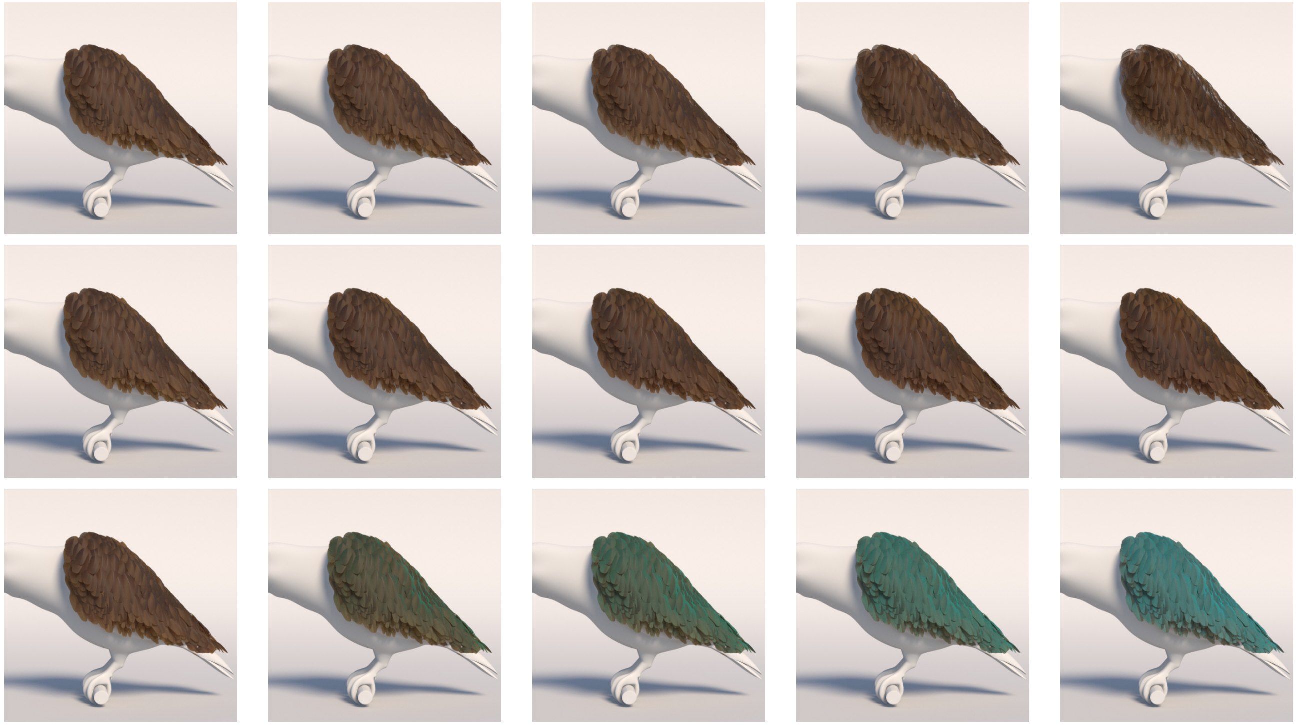

Figure 4.13: Parameter exploration of our model. *Top:* We vary the $\delta_{bb} = [0.1, 0.3, 0.5, 1.0, 3.0]$. Larger separation between the barbules increased the transmittance with high values showing prominent transparency effects. *Mid:* We show the effect of increasing the barb axis ratio $e_b = [1.0, 1.2, 1.4, 1.6, 1.8]$, progressively making the central highlight sharper. *Bottom:* We increase the size of the medulla $b_m = [0.0, 0.2, 0.4, 0.6, 0.8]$, for higher values, the color is mostly dominated by the color of the medulla and the appearance becomes more diffuse.

4.8 DISCUSSION

FAR-FIELD VS NEAR-FIELD Our model assumes far-field rendering of feathers and thus is not able to capture the high-frequency details of visible barbs, which are naturally handled by curve-based methods, which in the far field suffer aliasing. Extending our method to near-field appearance would require more sophisticated authored maps, and redefining the masking terms to support per-point masking. Deriving a proper filtering technique for such a near-field model is also left as future work.

SHADOWING To simplify our BSDF we decided to omit shadowing, which intuitively have little effect in natural lighting conditions. This allows us to use a far-field integration of the BCSDF, instead of requiring integrating both the input and outgoing visible surfaces of the fibers, at the cost of removing some physical plausibility.

ENERGY CONSERVATION Given the lack of shadowing, our BSDF is energy preserving without the need of multiple scattering simulations, except for the potential energy loss in the BCSDF, which might be significant when introducing elliptical cross-sections and an inner scattering medulla. However, as shown in Figure 4.5, we found that limiting light transport inside the fiber to five lobes result in a minimal energy loss.

WAVE OPTICS Our model omits important wave optical behavior, except for the structural coloration of the medulla, which we approximate as a diffuse reflection. Other effects such as the iridiscense due to the grating-like structure of barbs and barbules, or the thin-film-like structure in barbules [95] is omitted. Incorporating these effects, as well as a more principled definition of the medulla's diffuse structural coloration is left as future work.

OTHER TYPES OF FEATHERS An important limitation of our surface-based approach is that is not particularly suitable for plumaceous feathers, where the vane is not as structured and are therefore more suitable for curve-based representations. Still, our BCSDF could be directly applied to curve-based representations [7], and hybrid approaches where curve-based barbs encode flyaways and plumaceous feathers, are probably the best choice for a general model covering the whole spectrum of feather appearances.

NON-ELLIPTICAL BARBS AND BARBULES An additional assumption of our work is the use of elliptical cross-sections for barbs and barbules, which is a coarse approximation for both, specially barbules that appear in a large variety of cross sections. We made this choice for limiting the space of parameters of our model into something manageable; however, our methodology for computing the BCSDF based on path tracing is very general, so fibers with arbitrary cross-section could be used, in the same spirit as the work of Aliaga et al. [2] for cloth fibers.

PARAMETER SPACE Our model is parametrized by a total of 23 parameters including the masking parameters (5 geometrical parameters), BCSDF barb parameters (12 parameters for cortex and medulla) and BCSDF barbule parameters (6 parameters for the cortex), with absorption coefficients and diffuse reflectance defined by RGB values. The model is expressive as shown in our experiments and the parameter space might seem challenging.However, these parameters all have an intuitive meaning, either structurally or appearance-wise, and the general appearance features of feathers (goniochromatism, view-dependant transmittance, highlights...) emerge naturally. We selected some of the parameter values based on the structural and optical properties observed on feathers in previous works [99, 152, 208] and we found the rest empirically. However, targeting a specific feather's appearance required several iterations of parameter editing in our experience. A more example-based automatic parameter selection would definitely be an interesting next step on this research path. Introducing a length-dependent barb radius would probably produce more realistic results.

CONCLUSION We have presented a practical surface-based far-field appearance model for feathers, in which we model the complex microgeometry of a feather as a lightweight texture and an analytical masking term that accounts for the angular dependent visibility conditions, which previous curve-based models handled explicitly via curve visibility. On the core of our model is a new fiber BCSDF that supports elliptical cross sections and (potentially-colored) scattering medulla, which is crucial for representing a wide variety of real-world feathers.

Part III

PERCEPTION AND EDITING OF TRANSLUCENT MATERIALS

The first chapter of this part (Chapter 5) is devoted to understanding how the perception of translucency is affected by dynamic lighting. The main contributions are the design of a user study aimed at understanding the effects of light motion on the perception of translucency and its successive analysis. The second chapter (Chapter 6), instead investigates the *editing* of translucent materials. The key contributions are a user study focused on determining an adequate similarly metric for translucent materials, the construction of a perceptually meaningful manifold for homogeneous translucent materials, and the development of an editing interface that leverages the manifold, validated through a second user study.

5

ON THE INFLUENCE OF DYNAMIC ILLUMINATION IN THE PERCEPTION OF TRANSLUCENCY

ABOUT THIS CHAPTER

Previous work, has shown that different light conditions can result in constancy failure for translucent materials [239]. However, the stimuli used were static images, which can be a potential limitation, since dynamic motion has been proven, e.g., to improve glossiness perception [35, 183, 230]. In this work, we study the effect of dynamic lighting on the perception of translucent materials and see if it can alleviate such constancy failure. We design and run a psychophysiological experiment and conduct the subsequent statistical analysis. Surprisingly, we find no statistical difference between using static or dynamic stimuli, suggesting that light motion might not impact the perceived densitiy of a translucent material. As the leading author, I have led the design, development, and analysis of the user study. I have also contributed significantly to the writing of the manuscript. The work introduced in this chapter has been presented as a poster to SIGGRAPH 2022, where it was awarded with the third place in the Student Research competition (SRC) - Graduate Students. Later on, the full work was accepted at the Symposium on Applied Perception 2022.

D. Lanza, A. Jarabo & B. Masia

On the Influence of Dynamic Illumination in the Perception of Translucency

Symposium on Applied Perception (SAP 2022)

5.1 INTRODUCTION

As discussed in Section 2.2.1, numerous works have been devoted to understanding how confounding factors affect our perception of translucent objects. However, existing previous literature analyzed the influence of confounding factors by running most of the experiments using static viewing and illumination conditions. This is a rather strong limitation, since it limits the viewing condition to a single snapshot, and reduces the potential information encoded in the temporal domain, which human observers are known to leverage for gathering additional data. For example, it has been shown that non-static conditions favor accurate glossiness perception [35, 183, 230]. In this work, inspired by this line of papers, we analyze the perception of translucent objects under *dynamic illumination conditions*. In particular, we aim to understand whether dynamic illumination can or not reduce the problem of constancy failure, which has potential applications in authoring or previsualization, where the common setup assumes static illumination. Such static setup has been largely explored, including illumination configurations targeting better shape [181, 216] or material [17] understanding. However, even a carefully fine-tuned appearance might break when the illumination changes, due to the aforementioned problem of constancy failure. Our goal is to complement

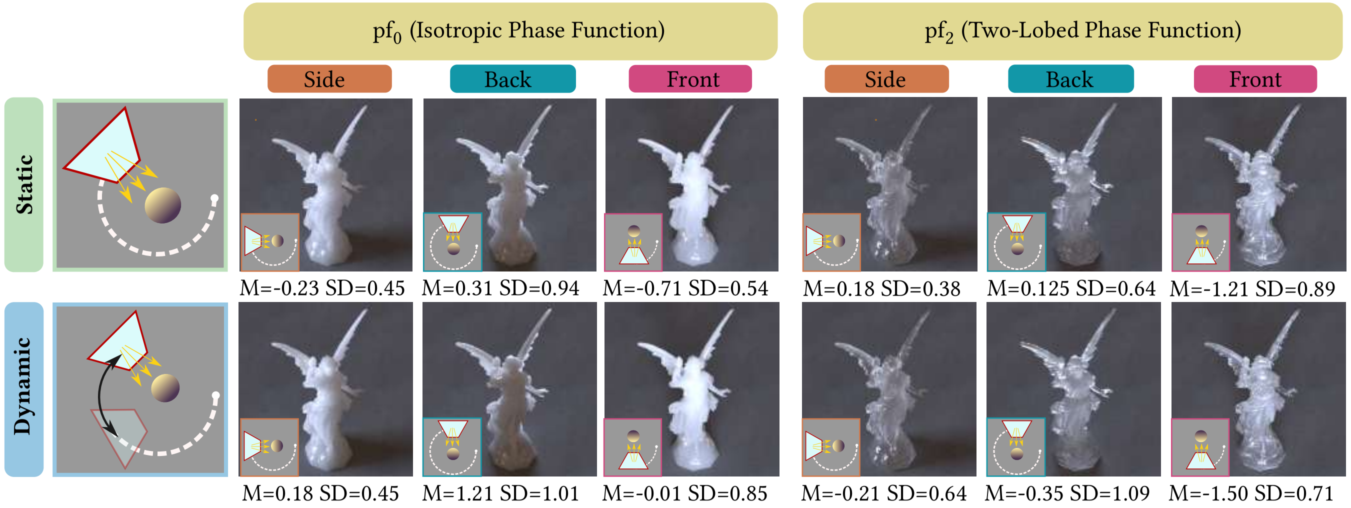

Figure 5.1: We analyze the effect of dynamic natural illumination on the perception of translucent materials. We design a matching experiment where participants estimate the optical density (the extinction coefficient σ_t) of a reference object with *static* (top row) or *dynamic* (bottom row) illumination. We do this under three different light directions in the match stimulus (*side*, *back* and *front*, depicted in the inset of each image), for different optical properties of the translucent medium. The images depict our test object, rendered with the average density estimated by participants for each condition, for a fixed reference density $\sigma_t = 4.0$ (M and SD represent the mean and standard deviation of the estimation error). Surprisingly, there are no significant differences between static and dynamic illumination when estimating the reference (ground-truth) density.

Figure 5.2: *Left:* Selected frames of a dynamic reference stimulus (video) used in our experiment, illustrating the changes in appearance as the illumination moves. *Right:* Example of a static reference stimulus used. From left to right, the light rotation used to render the stimuli is 135°, 180°, and 225° for the frames of the dynamic Reference, and 210° for the static Reference.

previous behavioral studies that noted how users tend to use viewpoint motion [59], exploring how the temporal evolution of shading can improve this problem. Specifically, we investigate if showing a stimulus in a scenario with a dynamic lighting can alleviate or not users' translucency constancy failure, similar to previous studies that explored novel ways to facilitate shape [181, 216] or material [17] understanding with the help of specific lighting settings. To do that, we compare participants' performance when determining the density of translucent objects under both static and dynamic lighting conditions (an example is shown in Fig. 5.2). We perform a series of guided tasks, where participants estimate material properties under a variety of lighting conditions and optical parameters (see Fig. 5.1). The results of our experiments confirm, as discussed in previous works, that both the direction of the illumination and the scattering directionality are key factors determining the perception of translucency. However, quite surprisingly, we observe that humans perform equally well when assessing the optical properties of translucent materials in both static and dynamic scenarios, suggesting that we do not leverage the extra information encoded in the temporal dimension in this particular scenario.

5.2 EXPERIMENT DESIGN

The goal of our study is assessing the effect of having dynamic illumination on the perception of translucent materials. More specifically, evaluating whether dynamic lighting leads to a more accurate perception of translucent appearance as compared to a static counterpart. We do this through an asymmetric matching task, which has been used successfully in similar contexts [239]. In a matching task, participants need to adjust the density of the scattering medium of a Match stimulus to resemble that of the Reference stimulus. In our experiment, this is done in two different conditions: with a *static* and with a *dynamic* illumination in the Reference stimulus (in the latter case the Reference stimulus is thus a video). Each condition is evaluated for a variety of illumination directions, and material properties, in particular the phase function and the density (the medium extinction coefficient σ_t). To avoid direct comparison, we use an

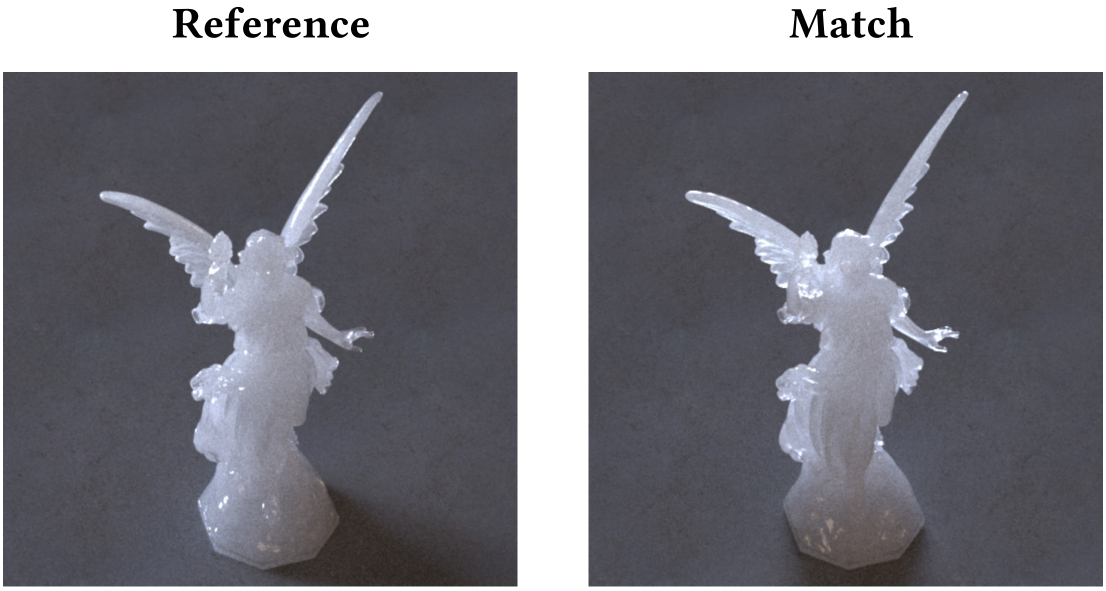

Figure 5.3: Experiment design. We show the user two images, or a video and an image, side-by-side. The user is asked to edit the Match image density (right) until it visually matches the Reference (left).

asymmetric matching task experiment; in our work, this is by done by having different lighting condition of the Target and Match stimuli, always differently (see Figure 5.3).

5.2.1 *Stimuli*

Sample stimuli used in the experiment can be seen in Figure 5.3. The stimuli are rendered images featuring the *Lucy* statue from the Stanford 3D Repository [130], which is often used to evaluate the perception of translucency due to its combination of thick and thin features [67, 239]. All stimuli were generated using the Mitsuba 0.6 physically-based renderer [102].

As explained in Section 2.1.5, translucent objects are modeled by a scattering medium inside their volume and a dielectric interface at the surface of the object. We model the interface as a smooth dielectric ($\rho = 0$), so that it presents highlights that are associated with translucent materials [49, 67, 147, 148, 161]. For the scattering medium, we fix the index of refraction to $\eta = 1.5$, which is a common value in translucent objects such as wax and glass [58], and the single scattering albedo to $\alpha = 0.99$, following previous work [67], so scattering dominates over absorption. These parameters are fixed for all the stimuli shown in the study, while we analyze the role of the remaining two parameters that model the behavior of the scattering medium: the density (or extinction coefficient) σ_t and the phase function f_p.

DENSITY The Reference stimuli are thus rendered using four different density levels, ranging from 3 to 6 in logarithmic scale, such that $\sigma_t = \exp(d)\ m^{-1}$, with $d = [3..6]$. In the Match stimuli, the density is the parameter that the participants in the study will need to adjust so that the material looks like that of the Reference; participants will be able to adjust it within the range $d = [0..10]$, with a step size of 0.25. For reference, the height of the *Lucy* is set to 5.23m.

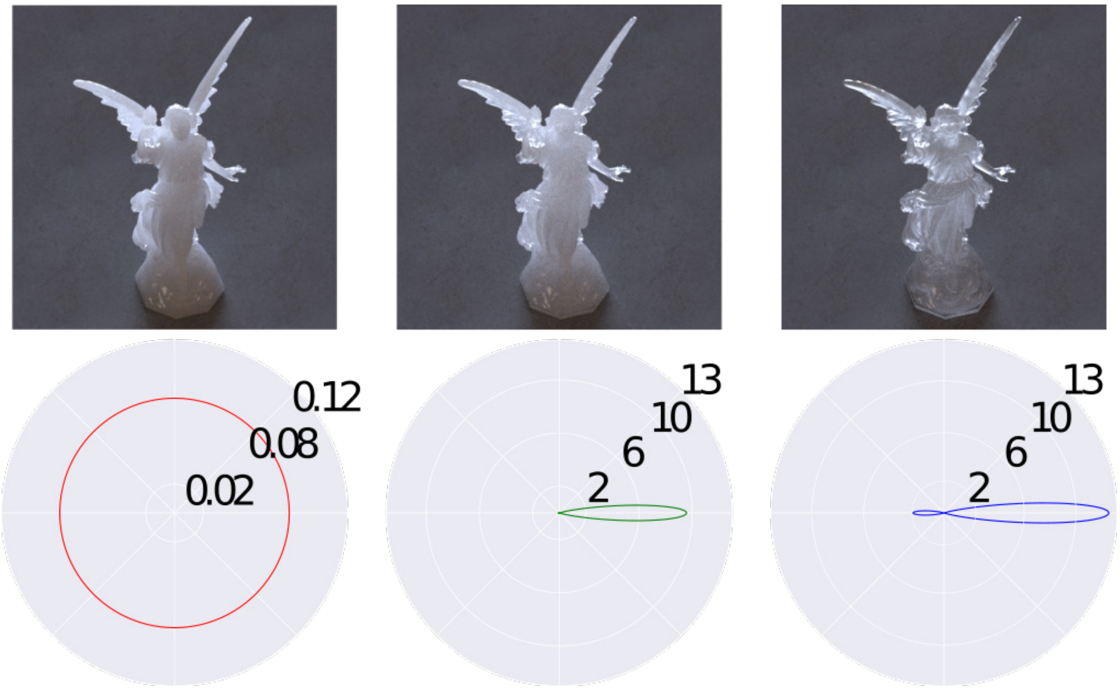

(a) $f_{p,0}$ (b) $f_{p,1}$ (c) $f_{p,2}$

Figure 5.4: Polar plots illustrating the shape of our three phase functions (bottom), together with their effect on the appearance of a translucent object (top) (all other parameters are kept fixed). Note the large magnitude difference in the primary lobes of $f_{p,1}$ and $f_{p,2}$, and with the isotropic $f_{p,0}$.

PHASE FUNCTION We render the stimuli with three different phase functions that yield perceptually distant appearances according to Gkioulekas et al.[67] (see Figure 5.4). The first phase function, $f_{p,0}$, is a purely isotropic one (i.e., $f_{p,0}(\mu) = 1/4\pi$). The other two phase functions, $f_{p,1}$ and $f_{p,2}$, are two-lobed (backward and forward) phase functions defined using two von Mises–Fisher distributions as

$$pf_{1,2}(\mu) = w\, f_{vMF}(\mu; k_1) + (1 - w)\, f_{vMF}(\mu; k_2)\,, \tag{5.1}$$

where the exact definition of $f_{vMF}(\mu; k)$ can be found in Appendix A.1. Phase function $f_{p,2}$ is aggressively forward scattering, while $f_{p,1}$ presents a less prominent forward lobe than $f_{p,2}$ (see Table A.1 for details at A.1).

ILLUMINATION Each trial in the matching task is composed by a Reference and a Match stimulus; the illumination in both is never the same, to create an asymmetry in the matching task and avoid pixel-by-pixel comparisons. We illuminate the scene using a captured environment map, in contrast with most previous works, which employ different forms of synthetic lighting [49, 239]. This has the advantage that natural

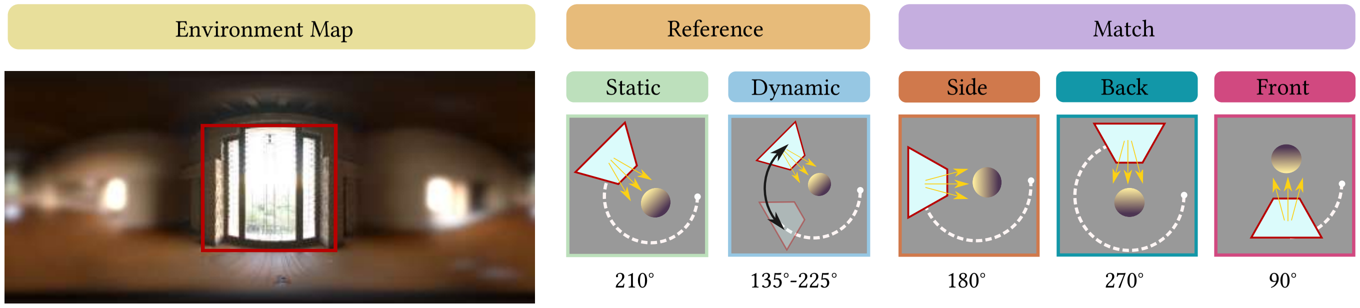

Figure 5.5: *Left: Ennis* environment map used to render our stimuli. We highlight in a red box the window used as a reference for the rotation. Middle: Schematic representation of the two different illumination conditions for the Reference, where in the dynamic case the illumination rotates around the object. *Right:* Schematic representation of the three illumination setups for the Match images. Below each, we report the rotation angle along the azimuthal direction.

environment maps are more realistic, as well as usually preferred over synthetic illuminations when comparing surface reflectance of materials [47, 48, 124]. When selecting the environment map, we seek a tradeoff between naturalness and controllability, to be able to assess the influence of lighting direction. We use the *Ennis* environment map, shown in Figure 5.5. In it, most of the irradiance comes from one main region and two other smaller areas, making the illumination highly directional while retaining a natural setup. Having such directional lighting has the additional benefit of creating more dramatic changes in the appearance when rotating the environment map. We use the large area light (marked in red in Figure 5.5) as a reference for the direction of the illumination. We further slightly blur areas outside it to reduce high-frequency reflections, following initial pilot studies that showed that the movement of the highlights across the surface when rotating the environment map could be distracting for participants, whereas our focus lies on the properties of the medium. Blurring out the details alters the original, captured environment map; a similar result could have been obtained by making the surface rougher, but this would have increased the rendering time to generate the stimuli. Slightly blurring the environment map thus offered a good compromise between realism, efficiency, and placing the focus on the scattering medium.

In the *static* condition, the Reference stimulus is rendered with a back-side illumination, in which the main light is at an azimuthal angle of 210° (back and side illuminations have been shown to be more informative than front ones [239]). In the *dynamic* condition, we render 45 frames, spanning from azimuthal = 135° to azimuthal = 225°, and the frames are then played in a bounce-loop. We set the video's frame rate at 30 frames per second (fps), and a rotation speed of 2º per frame. From preliminary tests, this speed seemed a good fit between showing the stimuli with enough speed to avoid an almost static Reference and slow enough to show the evolution of the light patterns as the illumination moved. To account for the influence of the illumination direction, the Match stimulus is rendered under three different lighting directions: *side* (main light at 180°), *front* (main light at 90°) and *back* (main light at 270°). Figure 5.5 illustrates the lighting scheme for both the Reference and Match stimuli, in the static and dynamic conditions.

5.2.2 *Procedure*

The experiment is carried out in two separate sessions, one for the *static* condition, and another one for the *dynamic* condition. The two sessions are separated by at least 24 hours, to avoid fatigue and learning effects. The order of the two conditions is randomized, and its potential effect is checked in the subsequent analysis. Each session consists of 36 randomized matching trials (4 densities × 3 phase functions × 3 light directions). Before starting the experiment, the participant is asked to fill in an anonymous questionnaire including basic demographics and questions about previous knowledge in computer graphics and art. During the experiment, we record the responses of the participants to the matching tasks.

5.2.3 *Participants and Apparatus*

The experiment protocol is in accordance with the Declaration of Helsinki and was approved by the institution's Research Ethics Committee. Twelve participants (4 female, 8 male), with an average age of 30.3 years (±8.42) and normal or corrected-to-normal vision, took part in the experiment. All participants completed the experiment (they could withdraw at any point).

The experiments took place in a room with a controlled, constant illumination using an office light. The stimuli were shown on an ASUS Vz239he 24″ LCD display that had been previously color calibrated and the viewer was asked to sit at a distance of 60cm. We tone-mapped the HDR renderings to sRGB using a simple exposure-gamma tonemapper with $\gamma = 2.2$ and fixed exposure.

5.3 DATA ANALYSIS

Our experiment has three within-subjects factors, described in detail in Sec. 5.2: light motion (2 levels, *static* or *dynamic*), light direction (3 levels, *front*, *side* and *back*), and phase function (3 levels, $f_{p,0}$, $f_{p,1}$and $f_{p,2}$). Since it is a matching task, our dependent variable is the error in the estimated density adjusted by the participants, $\sigma_{t,est} - \sigma_{t,real}$, where $\sigma_{t,est}$ is the density estimated by the user and $\sigma_{t,real}$ is the actual density present in the Reference. We analyze our data using repeated measures ANOVA, with post-hoc pairwise comparisons with Bonferroni correction when applicable. We set the threshold ϵ to consider a factor statistically significant at $\epsilon = 0.05$. No data was discarded due to outlier rejection. We also measured and analyzed the time to completion of each trial, but we found only a first-order interaction between time spent and the light direction. However, the successive Post-Hoc analysis did not highlighted any common pattern. While we discuss here the main significant effects and interactions, the full results of the analysis can be found in Appendix A.2.

5.3.1 *Does a dynamic lighting setup improve density estimation?*

The first question we sought to answer was whether having a dynamic lighting setup led to improved density estimation with respect to the static counterpart. Our analysis revealed no significant effect of the light motion factor on the error incurred by participants ($p = 0.1016$). Fig. 5.6 plots the estimated density against the real density for each condition, and illustrates how the trends are similar for both the static (top row) and dynamic (bottom row) cases.

There is, however, a significant first-order interaction between light motion and phase function ($p < 0.001$). The type of phase function (see Fig. 5.4) has a significant effect on the ability of participants to estimate density in the case of the dynamic setup ($p < 0.001$), but not in the case of the static setup ($p = 0.3889$): The error remains roughly constant across phase functions for the static case (marginal means $M_{pf_0} = -0.16$, $M_{pf_1} = -0.23$, $M_{pf_2} = -0.33$), while in the dynamic case a large increase in

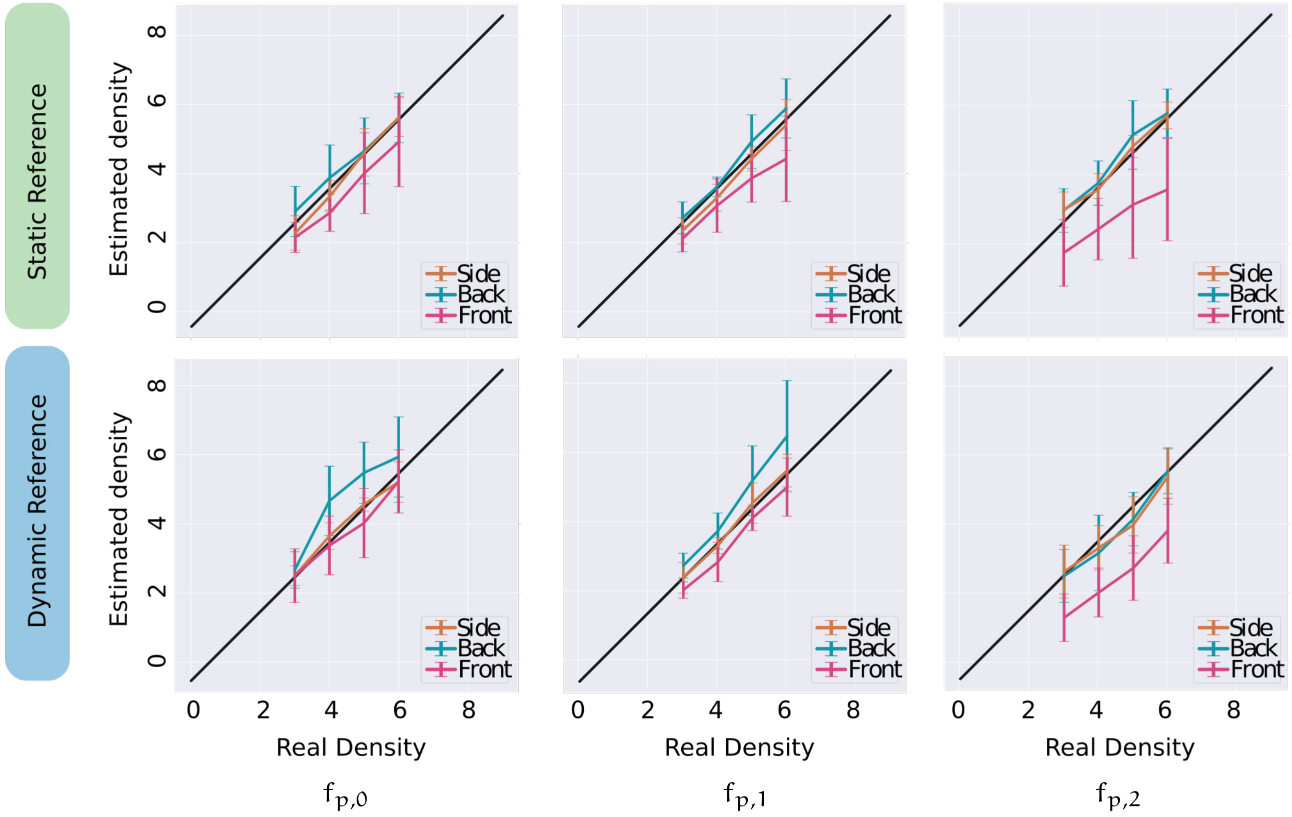

Figure 5.6: Results, averaged across participants, showing the estimated density with respect to the real density, for the static (top row) and dynamic (bottom row) cases. We report also one standard deviation of the estimated density as error bar. Columns correspond to the three phase functions tested, and line colors to the three different light directions used (*side*, *back* and *front*). The black line marks the ideal (ground truth) response.

error is observed for the most forward scattering phase function, $f_{p,2}$ (marginal means $M_{pf_0} = 0.19$, $M_{pf_1} = 0.10$, $M_{pf_2} = -0.65$).

5.3.2 *Is there a more favorable light direction for density estimation?*

Light direction has a significant effect on the accuracy of the estimation, both in the static and dynamic cases ($p < 0.001$). Further, there is an interaction between the phase function and the light direction ($p < 0.01$). Post-hoc analyses reveal the nature of this effect, which is shown in Fig. 5.7 (bottom). The *front* light direction is consistently harder to estimate than the *side* and *back* conditions (marginal means $M_{side} = -0.05$,$M_{back} = 0.22$, $M_{front} = -0.90$), and the effect is aggravated in the case of phase function $f_{p,2}$, the most forward scattering one (see Fig. 5.7). The higher error when viewing the Match front lit might also be caused by the experimental setup, since the *front* (90°) light condition is the farthest from the static Reference (210°), while the *side* (180°) and *back* (270°) conditions are visually closer to the static Reference. We further discuss this topic in Section 5.4.

Interestingly, in the dynamic lighting setup, there is no significant effect of the light direction for $f_{p,0}$, whereas in the static case the *front* condition is performing worse than the *side* one, and worse than its dynamic counterpart. Although a weak effect, this suggests that the dynamic lighting may be aiding in the estimation of the density for the challenging *front* case. This only occurs, however, for the isotropic $f_{p,0}$, and not for the more complex $f_{p,1}$ and $f_{p,2}$, where the error remains similar between the static and dynamic conditions.

5.4 DISCUSSION

We have designed a matching task experiment where participants estimated the density of a translucent object under different lighting conditions and material properties. We discuss here our main findings, and contextualize them with respect to previous work.

THE INFLUENCE OF LIGHT DIRECTION In contrast to previous works investigating the effect of light direction [239], where the stimuli were illuminated with synthetic lighting (e.g., spherical harmonics), here we employ an environment map captured from a real scene, and therefore more representative of a real-world illumination.

Both our findings and Xiao et al.'s [239] reveal the influence of light direction and its interaction with the phase function. Like Xiao et al., we observe that the *front* lighting condition leads to higher inaccuracies in density estimation than side or back lighting, statistically significant in non-isotropic, forward-scattering phase functions. They report an overestimation of the Match density when the *Reference* is frontally lit, and we find an (equivalent) underestimation of the Match density when the *Match* is frontally lit. Frontal lighting has the effect of "flattening" and decreasing local contrast of non-specular surfaces [29, 162], which can hinder the ability to discriminate between similar appearances. Moreover, we note that $f_{p,2}$ performs differently than the other two, when

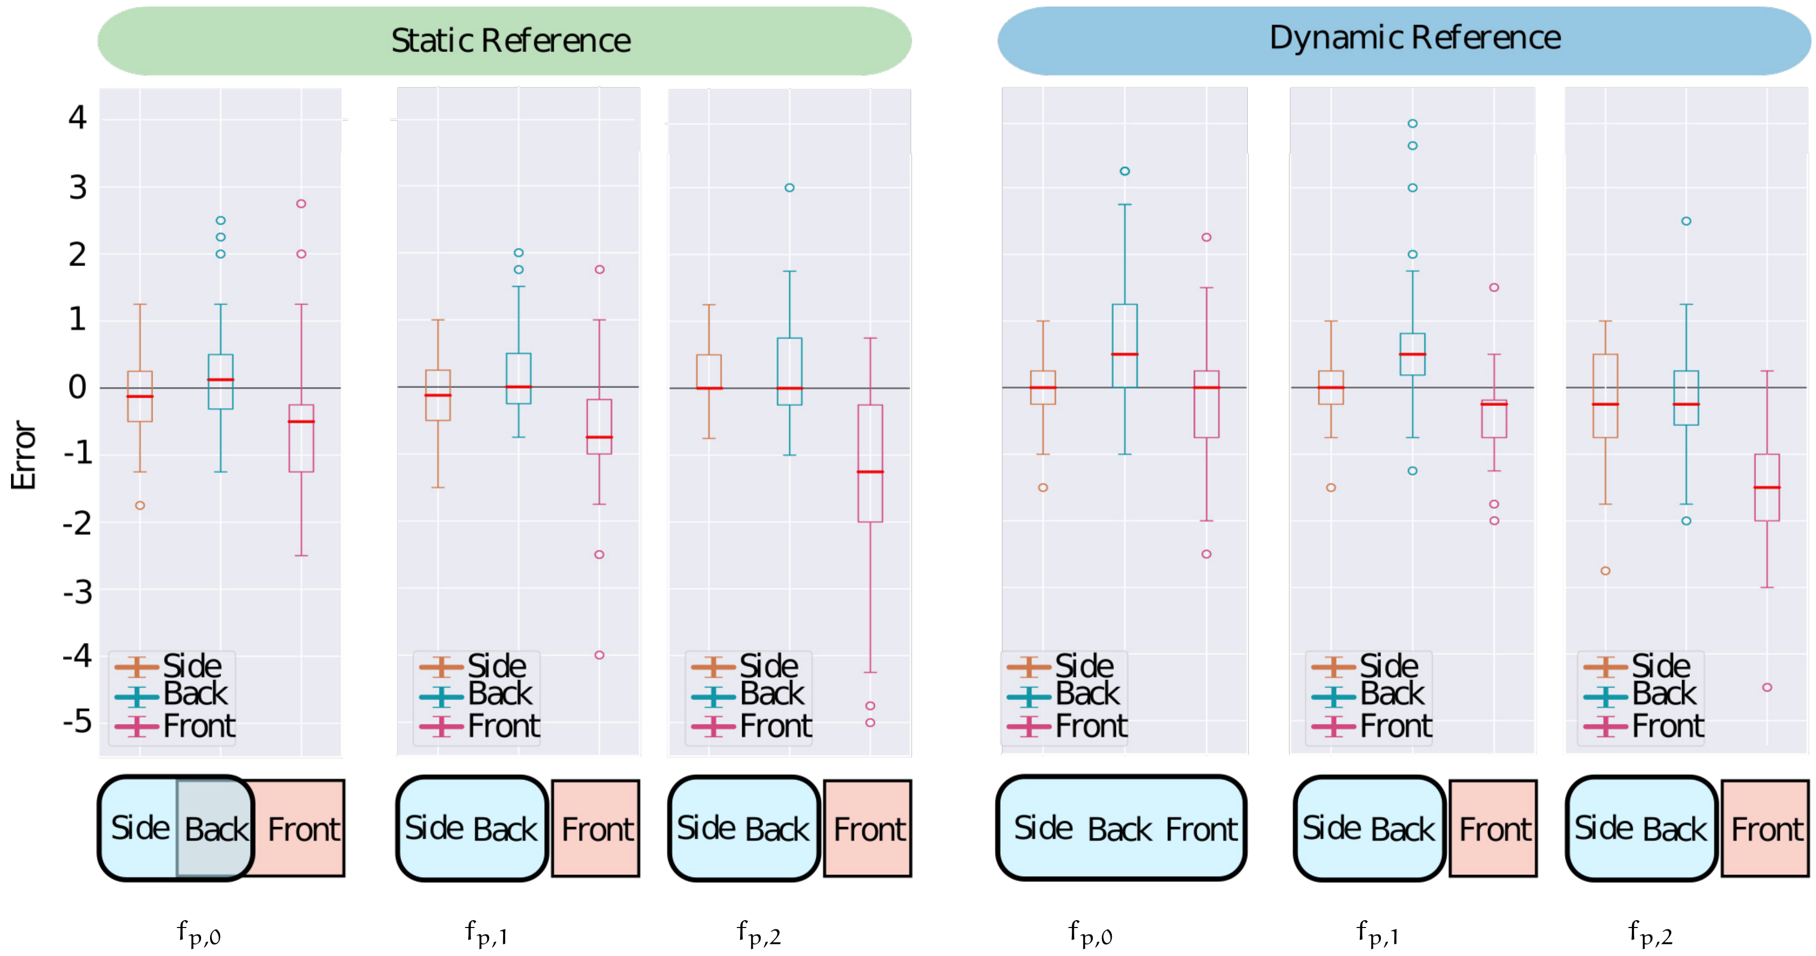

Figure 5.7: Error in the estimated density for each combination of light condition and phase function, for both the static and dynamic experiments. Below them, we show the results of the post-hoc analysis on the significant differences between light directions for each phase function. Please refer to the text for details.

frontally lit. We believe that this is caused by the presence of the strong backward and forward scattering peaks, that increases the sharpness of details by forward-scattering light inside the object in all but some contour areas, where backward scattering increases brightness. However, while the trends resulting from light direction are similar for the same phase functions, Xiao et al. find a larger error in the estimated density than in our case. This may be caused by the different types of illumination that we used to render the stimuli: As opposed to Xiao et al.´s low-frequency synthetic illumination, we use a more natural light with a sharp, dominating high-frequency, directional light. As in previous work [240] we hypothesize that the increased directionality of the light emphasizes the translucency specially in thin areas, in which scattering differences are better perceived by our HVS [61]. This correlates with the perceived effect of illumination on opaque material perception [124]. Nevertheless, this is just a hypothesis which would require further testing in equivalent conditions.

THE INFLUENCE OF LIGHT MOTION Previous studies have shown that observers, when offered the possibility, tend to use motion clues when investigating translucent objects, by either moving the head or by directly rotating the object [61]. Our experiment looks at a slightly different scenario, in which the light source moves around the object. Although similar, these two setups convey different information. By moving the object, or the head, the observer can see the same object from a different perspective, recovering information about the geometry of the object. In our scenario, the participant has extra information about the evolution of light patterns, which could help in the estimation of translucency properties. However, our data do not reveal a more accurate translucency perception in the dynamic with respect to the static case. It seems that participants were not able to leverage the extra information provided by these light patterns. Given the strong influence of light direction on the perception of translucency observed in both our and previous works, this was a surprising result, which requires further investigation. Moreover, analyzing the time that users spent on each trial did not show significant differences between the static and dynamic conditions.

VISUAL EQUIVALENCE While care was taken so that in all cases the Match and Reference were different pixel-wise, to avoid pixel matching, , our lighting setup is such that some Match light directions are closer to the Reference than others. For instance, the *side* (180º) and *back* (270º) conditions are closer to the static Reference (210º) than the *front* one (90º). This could have led to participants more easily estimating the "correct" optical density when the Match was lit from the side or from behind than when it was lit frontally, since these two conditions might present visual clues that are also present in the static Reference. In this regard, it is interesting to note that in the dynamic Reference condition, in which the illumination spans from 135º to 225º (thus closer to the match when frontally lit), users did not improve their performance. This seems to further support the idea that people do not use dynamic illumination cues to assess the nature of translucent materials when rotating the illumination. In any case, exploring the visual equivalence in translucent materials, for both static and dynamic scenarios, remains an open problem.

LIMITATIONS AND FUTURE WORK As with any study of this kind, our findings are strictly valid only for the conditions here tested. We limited our experiment to one shape (the *Lucy* model) to keep the experiment size tractable, and to focus on the effect of the illumination. Other works have looked into the effect of geometry, and shown differences between simpler (sphere, torus) and more complex geometries [58], so extrapolation of our findings to a variety of geometries should be done with caution.

An important constraint of our experiment is that participants were not able to control the direction or speed of movement, since they were shown a fixed video, in a loop. We made this choice in order to ensure control of the stimuli and consistency between participants over the stimuli viewed, and because in certain application scenarios free exploration of the object would not be possible. It remains as future work to test whether this would have an effect on our findings.

Finding that we may be unable to leverage the extra information offered by a dynamic illumination for translucent density estimation can have implications for computational design and editing of materials, both in rendering and fabrication scenarios. We hope this work serves as yet another step towards our understanding of material appearance, and specifically of translucent appearances.

As a future work, it would be interesting to investigate the impact of motion translucent materials with a rougher surface interface. As noted in previous work, there is an interaction between translucency and glossiness perception [58]. In particular, it seems that the surface halo created by rough materials generates areas with low local contrast, similar to what was noted in previous work [162]. We argue that this constancy failure might be alleviated by a dynamic reference since the halo created by surface glossiness would likely vary more sharply than translucency. Still, since translucency is a global (not local) effect, understanding the tight interaction between glossiness and translucency under dynamic lighting remains as future work.

SUBSEQUENT STUDIES

In this paragraph, we briefly discuss and relate the work described in this chapter with a recent work from Gigilashvili et al. [57], which was published a year after the work presented in this chapter. In this work, the authors run two user studies: The first aims to investigate if movement alone is enough to elicit a sense of translucency, finding out that it does not seem the be the case. In the second study, the authors run an experiment similar to the one discussed in this chapter. They also use an asymmetric matching task, studying the impact that motion can cause on density estimation of translucency, using stimuli different than ours. In this case the stimuli featured different shape (the so-called "spiky sphere"), lighting condition (a surface area emitting a D65 spectrum, positioned in front of the shape), and type of motion (with both shape and light moving). Interestingly, using these different stimuli, Gigilashvili et al. [57] report results similar to ours, corroborating the thesis proposed in this chapter that motion does not affect the way users estimate the density of translucent materials.

6

NAVIGATING THE MANIFOLD OF TRANSLUCENT APPEARANCE

ABOUT THIS CHAPTER

While the previous chapter focused on the *perception* of translucent materials, in the following chapter we investigate the *editing* of translucent materials. The main contributions include a user study to identify an appropriate similarity measure for translucent materials; the creation of a perceptually meaningful manifold for homogeneous translucent materials, leveraging this measure; and the development of an editing interface based on this manifold, validated through a second user study. As the leading author, I have led the design, development, and analysis of both user studies; I have also been the main developer of the manifold and contributed significantly to the writing of the manuscript. The work introduced in this chapter has been accepted at the Computer Graphics Forum and presented at the Eurographics 2024 conference.

D. Lanza, B. Masia & A. Jarabo
Navigating the Manifold of Translucent Appearance

Computer Graphics Forum (Eurographics 2024)

6.1 INTRODUCTION

Materials that exhibit some degree of translucency are ubiquitous in the real world, ranging from organic materials, such as milk or wax, to inorganic materials like jade or glass. Light transport algorithms capable of simulating them with great accuracy, and in a computationally efficient manner, are well-developed [170]. When it comes to *editing* these material models, however, existing techniques are considerably less developed.

Translucent materials are typically modeled using the bidirectional scattering surface reflectance distribution function (BSSRDF), and manipulating their appearance through tuning of the optical parameters of the BSSRDF is a daunting challenge. These parameters span a high-dimensional space, with dimensions that correlate poorly with human perception, and complex, non-linear interactions between them, leading to distant optical parameters resulting in similar appearances [255]. This problem is aggravated in the case of novice users, who are not acquainted with the individual impact of each of these parameters on the final appearance, not to mention the intricate manner in which they interact.

Bridging the gap between physically-based parameters, used in analytical or measured models of appearance, and high-level attributes that humans can better understand and control is a long-standing problem in the field of material appearance modeling. As discussed in 1.2, this often involves finding low-dimensional perceptual manifolds, which enable a more intuitive navigation of the targeted range of appearance [186, 189, 211], and in some cases even provide a relationship between the dimensions of these manifolds and the parameters of reflectance models [176, 233]. Unfortunately, all these works focus on building a perceptually-based appearance manifold for opaque BRDFs.

Despite –or owing to– their higher complexity, translucent appearance models have received less attention in this area. Most efforts have focused on optically-thick materials, where editing can be done by manipulating the diffusion profile [19, 194, 223], thus limiting material editing to a particular set of translucent materials. In addition, the perception of translucency, and its interaction with factors like lighting and geometry, is not fully understood [46, 49, 61, 65, 122, 125], further hindering the task.

Our goal is to explore intuitive methods for navigating and editing translucent appearance, allowing users to abstract themselves from the low-level optical properties defining appearance and focus only on the appearance itself. We do this by building a perceptually-meaningful continuous manifold of translucent appearance (Figure 6.2). This manifold is built under the premise that perceptually similar appearances should be closer together, so global exploration of our manifold will allow drastic appearance changes, and local navigation will result in fine-tuning of appearance. We focus on homogeneous, achromatic translucent appearance, with the goal of evaluating whether an editing paradigm based on image navigation of such a manifold can be practical for users and outperform more traditional, slider-based editing approaches, thus justifying further exploration of its applicability to the vast space of translucent appearance.

We build our manifold leveraging an objective distance measure, which we select by conducting a perceptual study. Since our ultimate objective is appearance editing, we validate the usefulness of such space by proposing an interface that builds on top of it; we evaluate this interface with a user study, comparing it to a standard, slider-based interface Figure 6.1. We show that novice users have better objective performance in editing tasks using our prototype interface, and at the same time have a more satisfactory editing experience.

Specifically, our contributions are:

- A perceptual study to determine an adequate objective distance measure for translucent appearance that correlates with human perception of translucency over a wide range of optical parameters.
- A perceptually-meaningful continuous manifold of the space of translucent appearance, suitable for editing homogeneous translucent materials.
- An editing interface that leverages this manifold, and which we validate through a user study, showing its effectiveness and superior performance in comparison to a standard, slider-based approach.

6.2 BACKGROUND AND RELATED WORK

In this section, we briefly summarize the related work on translucency editing, perceptual spaces of appearance, and material design and editing interfaces.

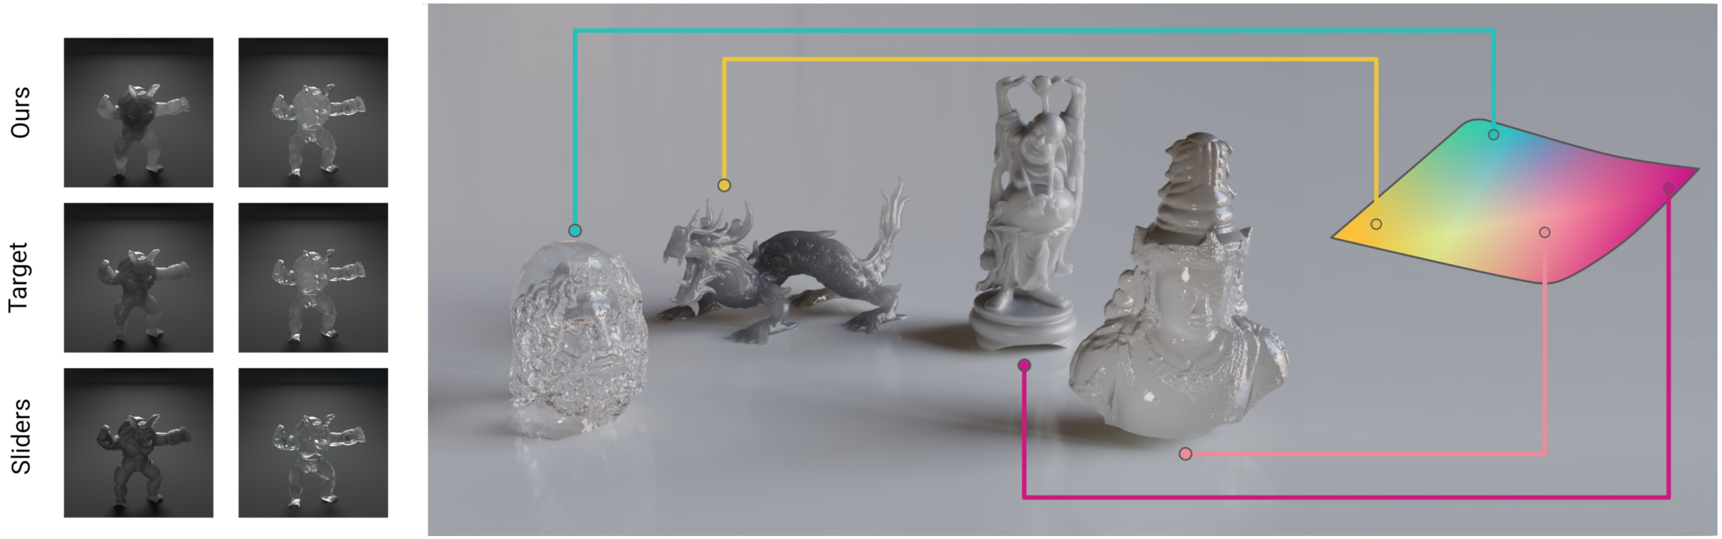

Figure 6.1: We propose a continuous, perceptually-motivated manifold for translucent appearance (illustrated on the right of the right panel, color-coding is for visualization purposes only). The manifold covers a wide range of materials, as shown by the objects on the right panel, whose appearance has been selected from the corresponding points selected in the manifold. We also propose an interactive editing interface that leverages this manifold, and allows users to navigate the space of translucent appearance more effectively than when using traditional slider-based interfaces. We evaluate this by means of a comprehensive user study including two different editing tasks. An example of the results obtained in this study with the two interfaces (ours and a slider-based one) for a certain target image is shown on the left panel. The user study shows that participants were able to better match the target image using our interface, with higher level of satisfaction.

6.2.1 *Light Transport in Translucent Materials*

As described in Section 2.1.5, we model translucent materials as a volumetric medium enveloped in a surface. The medium is characterized by the extinction coefficient σ_t, the single scattering albedo α, and the phase function describing the angular scattering $f_p(\mu)$, with $\mu = \omega \cdot \omega'$. We restrict our work to homogeneous achromatic media, avoiding spatial, angular, and spectral dependence of these parameters. For the phase function we use the common Henyey-Greenstein model [89], parameterized by the directionality parameter g. Additionally, the surface interface is characterized by ρ the roughness of the microfacet distribution and η the index of refraction.

6.2.2 *Editing Translucency*

As seen in Section 2.1.5 a vast section of previous work has been devoted to studying the perception of translucent materials. However, *editing* of translucent materials in an intuitive way is a relatively unexplored topic. Most existing works have focused on optically thick materials, where the light is at the diffusion regime: In this case, editing is generally done via modifying the diffusion profile [19, 114, 194, 241], or via local manipulation of the single scattering albedo in volumes [83]. For a broader range of optical thickness, the general approach falls back to manipulating optical parameters via sliders, either directly or via simplified controls [234]. Our work targets all ranges of optical thickness and poses editing as a low-dimensional intuitive navigation through the potential range of appearance, rather than low-level properties. Inverse rendering of participating media has proved a useful tool for manipulating translucent appearance. This includes optically thin [164] and thick materials [34, 111], as well as more general materials using fully-differentiable volume path tracing [64, 68]. Recently, neural-based approaches [25, 133] have demonstrated good performance on inverting the BSSRDF. While we leverage inverse rendering to obtain an estimate of material properties, our focus is on providing an improved method to intuitively edit translucency.

6.2.3 *Perceptual Spaces for Appearance Editing*

A considerable body of work has been dedicated to deriving low-dimensional perceptual manifolds of material appearance in which materials are organized along meaningful, interpretable dimensions. Often, a key motivation behind these is to facilitate editing and control of appearance, shielding the user from the high-dimensional, nonlinear spaces spanned by the optical or physically-based parameters of BSDF models.

A seminal work in this area is that of Pellacini et al. [176], where they derive a two-dimensional perceptual space for gloss. The space is built from psychophysical data, making use of multidimensional scaling, and they establish a relationship between their two main perceptual dimensions. A close goal is pursued by Wills et al. [233], who build a two-dimensional perceptual space for gloss for measured BRDFs. More recently, Toscani et al. [211] used a similar methodology with a wider range of appearance, and with the addition of an experiment to interpret and label the three main dimensions of their derived space.

Instead of focusing on a single attribute (e.g., gloss) and deriving a perceptual model for it, other works have broadened the scope, often at the same time lifting the restriction of having a perceptually-linear space, where a linear change in space corresponds to a linear change in appearance, and settling, instead, for *intuitive*, perceptually-meaningful spaces, in which items are spatially sorted following perceptual criteria, but without any extra relationship between space in the manifold and perceptual space. Matusik [151] evaluated different dimensionality reduction techniques for measured BRDF data, and defined a set of 15 attributes (which they call perceptual traits) which are then used to navigate the space of appearance. Inspired by this, Serrano et al. [186] generated intuitive manifolds for a carefully-curated list of attributes, building on the dimensionality reduction proposed by Nielsen et al. [167]. Shi et al. [189] follow a similar approach, but do not establish specific attributes a priori, and instead extract dimensions that explain perceived differences in appearance. They also propose an editing interface, discussed in Section 6.2.4.

Akin to these works, we seek to build a low-dimensional space that is linked to the optical properties of the material appearance model in a manner that allows for appearance editing. We draw inspiration from their methods, but while their focus is on opaque BRDFs, we deal with translucent materials, a much less explored domain.

In the context of translucent materials, the closest work to ours, which identifies a two-dimensional perceptually-meaningful manifold for a set of physical scattering parameters, is that of Gkioulekas et al. [67]. They specifically explore the impact of the phase function on translucent appearance, and relate the two axes of their space to moments of the phase function. We build on their findings, including the distance metric they use, but we tackle a wider range of appearance, and focus on providing an intuitive space for editing it. Recently, Liao et al. [136] proposed an unsupervised learning-based model that identifies perceptually relevant dimensions for translucent material appearance from images, and show that manipulations in the latent space of the model can lead to modifying the appearance of the object. While they work in image space, we seek to retain the mapping to optical parameters of the material, required for physically-based rendering.

6.2.4 *Material Design Interfaces*

Different interfaces and interaction paradigms have been proposed in the context of material appearance design and editing. They mainly rely on one of two paradigms: those based on a set of sliders that allows modifying a series of parameters, or those based on navigation of a structured collection of images. Both were analyzed by Kerr and Pellacini [118] in their comparison of material design interface paradigms for novice users. A third category relies on directly painting with brush tools, and has been used for opaque BRDFs, particularly highlights [30, 172], but requiring custom material models.

Slider-based models can rely on *physical* sliders, that expose and control the optical properties of the underlying BSDF model (as used by many commercial, off-the-shelf modeling tools [5]), or on *perceptual* sliders, allowing editing along the dimensions of perceptually-meaningful spaces of appearance like the ones introduced in Section 6.2.3. Some BSDF models, like Disney's Principled BSDF [18] are already designed to provide an optimal combination of controllability, expressiveness, and robustness.

Hybrid approaches that combine image navigation with subsequent tuning of parameters are commonly found in everyday applications such as Adobe Filter gallery [1] or Microsoft PowerPoint template gallery [157]. They offer good balance between a rapid selection of the desired result, and the fine-tuning precision given by sliders. In the context of material appearance, the recent work by Shi et al. [189] proposed to use such a hybrid paradigm for authoring opaque metallic-like BRDFs. We draw inspiration from this approach, and explore the use of such hybrid editing interfaces for translucency, which is a higher-dimensional problem with a more complex interaction between the low-level parameters used to define appearance.

6.3 A DISTANCE MEASURE FOR TRANSLUCENT APPEARANCE

Measuring the perceptual distance between the appearance of two materials is an open problem, and multiple measures have been proposed [128]. These measures can operate in material space or in image space. The former have the obvious advantage that they do not require rendering the images, but they have shown to correlate poorly with perception, especially when dealing with suprathreshold differences [165, 186]. Image-based metrics allow to factor in the influence of the geometry and the illumination, critical in material appearance [123], and in particular in the case of translucency [65, 239]. In this section, we briefly describe our choice of distance measure: we curate a set of metrics based on the literature (Section 6.3.1), and then conduct a perceptual study to select the best one for our particular type of stimuli (Section 6.3.2). For a more detailed description of the metrics and the experiment, we refer the reader to Appendix B.

6.3.1 *Image-based Metrics*

While the literature in the area of image-based appearance metrics is vast, the specific case of translucent materials is not as well studied as that of opaque BRDFs. Studies range from analyzing existing metrics proposed for BRDFs for their use in translucent stimuli, to the proposal of models specifically designed to predict human responses to those translucent stimuli. We explore the use of representative metrics of both, as well as other commonly used image-based metrics. In particular, we evaluate: the *cubic root metric* [165], which Gkioulekas et al. [67] found to correlate well with the perception of translucency; the *anisoshading ratio* proposed by Kiyokawa [122] for translucency; and a set of image-based metrics including the well-known *MS-SSIM*, widely used for image comparison, the recent *FovVideoVDP* [145], and the learning-based metric *LPIPS* [253]. Details on these metrics can be found in the Appendix B.1.

6.3.2 *Perceptual Study*

We seek to find, for the type of stimuli that we work with, which metric correlates best with human judgements. As mentioned, previous work has analyzed the performance of different image-based metrics on the perception of translucency [67]. However, they did so with a set of stimuli that focused on studying the role of the phase function. Since their focus is on densely sampling the space of phase functions, they fix the other scattering parameters to isolate the effect of the former (σ_a and σ_t take the values of the green channel of marble [111], with perturbations of them around these default values; roughness is also fixed), leading to a more reduced range of translucent appearance (in terms of σ_a, σ_t and roughness coverage) than the one we cover in the present work. Therefore, we build on their work, and assess through a perceptual study whether their proposal of the cubic root metric applies to a larger range of appearance, even when compared to more recent and sophisticated image-based metrics.

6.3.2.1 *Stimuli*

Our stimuli are images of a translucent object against a uniform background; we use a fixed geometry, and vary the material's optical parameters and the lighting conditions. Images are rendered with Mitsuba 0.6 [102], using the default volumetric path tracer with no limit to the number of bounces. We fix the geometry to the *Lucy* statue, since it has both thick and thin parts and shows a balance of medium and high-frequency details, which are helpful cues [61]. We add a background horizontal plane, to project caustic patterns that are also used to assess translucency [56].

Since the perception of a translucent material is heavily affected by where the light is positioned with respect to the object, we render our stimuli with three different variations in lighting: *front*, *side* and *back* [239]. In all three cases, we use a strong area light to generate a directional light source from the specific direction, and add a filling environment map (*Ennis* [98]).

We model translucent materials as a medium enclosed by a surface, yielding five optical parameters that determine material properties (see Section 2.1.5). We fix the index of refraction (IOR) to $\eta = 1.5$, since most translucent materials have an IOR close to this number [61]. We explore the whole range of phase function asymmetry $g \in [-1, 1]$ and single scattering albedo $\alpha \in [0, 1]$, while we bound both surface roughness and extinction to $\rho \in [0.01, 0.6]$ and $\sigma_t \in [0.01, 10000]$, respectively. Given that the parameter space $\boldsymbol{\pi} = [\rho, \sigma_t, \alpha, g]$ is four-dimensional, a grid-based sampling would be impractical; instead, we sample the space using a low-discrepancy quasi-random distribution (Halton). The use of a low-discrepancy sampler allows to have a uniform coverage of the sampling space and scales well in multidimensional spaces. Even within our scheme, the optical parameters sampled have a highly non-linear relationship with translucent appearance. Based on findings from previous work on translucency perception [63], we choose a power-like distribution for the single scattering albedo α and extinction σ_t and empirically set the parameters of the power function to provide a balanced coverage of appearances. Specifically, we follow $\alpha = \xi^5$ and $\sigma_t = \xi^{-5}$, with ξ a random number, skewing the sampling towards lower extinction values and including optically thinner appearances that would otherwise be underrepresented.

6.3.2.2 *Procedure*

We avoid ranking methods and resort to a two-alternative forced choice (2AFC) experimental paradigm. This is common practice in complex appearance evaluation tasks, where the measured effect may not follow a linear scale. As in standard 2AFC experiments, the participants need to select which of the two *candidate* images is most similar to a *reference* one.

We sample sixteen sets of material properties, under three different lighting conditions. This yields a total of 5,040 triplets, each evaluated by at least five different participants, leading to a total of 25,200 trials. Given the large number of trials, we use Amazon Mechanical Turk, implementing control trials to discard unreliable participants (slightly under 25% of participants were discarded).

6.3.2.3 *Results*

Following previous works [55, 123, 233], we compute the *accuracy* of each image-based metric when predicting human responses using both its *raw* and *majority* modalities. The raw accuracy considers each response as a correct answer for each triplet, while the majority accuracy considers that only the majority's response for each triplet is correct. Differences between both modalities exist when there is low agreement between participant responses for a given triplet. Table 6.1 shows the accuracy results for each of the five metrics when predicting human responses, in both its *raw* and *majority* modalities.

Image quality metrics do not provide a good prediction in this scenario, even in the case of those designed to detect supra-threshold differences. On the other hand, we see how the cubic root metric initially developed as a metric for BRDF comparison [165] offers the best results, consistent with what Gkioulekas et al. [67] found in their exploration of the phase function. These are complementary findings to the work of Gkioulekas et al., since we span a wider range of optical parameters and appearance, going from crystal-like to darker materials (see Appendix B for representative subset of the materials used). Therefore, the cubic root metric, initially proposed for BRDFs, will be our distance metric of choice.

6.4 A PERCEPTUALLY MEANINGFUL SPACE FOR TRANSLUCENT APPEARANCE

Editing material appearance is generally done by moving along non-orthogonal nonlinear dimensions via sliders. In our case, that means a four-dimensional space, which is a challenging task for novice (or even experienced) users.

Instead, our goal is to provide a more natural navigation along the manifold of translucent appearance. That imposes a set of design constraints for our manifold: 1) The manifold should be perceptually meaningful, so that similar appearances are close together; 2) the manifold should be continuous, to allow navigation along the space of

	Majority			Raw		
Metric	Side	Back	Front	Side	Back	Front
Cubic root metric	**0.77**	**0.81**	**0.84**	**0.70**	**0.73**	**0.75**
Anisoshading ratio	0.61	0.65	0.69	0.50	0.50	0.63
MS-SSIM	0.29	0.27	0.25	0.22	0.19	0.18
FovVideoVDP	0.29	0.27	0.25	0.23	0.20	0.17
LPIPS	0.75	0.77	0.77	0.69	0.71	0.72

Table 6.1: Accuracy (both majority and raw) of the various image-based metrics when predicting human similarity judgements, for the three lighting conditions: side, back, and front. Performance is generally consistent across lighting conditions. The cubic root metric outperforms the rest, and is thus our distance metric of choice.

possible translucent appearances; and 3) the manifold should be two-dimensional, so that it can be easily navigated by the user.

Mathematically, we want to find a two-dimensional manifold $\mathcal{M} \subseteq \mathcal{R}^2$ that is a convex hull of appearances. Each point $y_i \in \mathcal{M}$ maps to a set of optical parameters $\boldsymbol{\pi}_i$, thus establishing a relationship $m(y_i) = \boldsymbol{\pi}_i$. For that manifold to be perceptually meaningful, we impose that for a pair of points $y_i, y_j \in \mathcal{M}$, it holds that

$$|y_i - y_j| \propto d(\mathbf{i}_i, \mathbf{i}_j), \tag{6.1}$$

with $d(\cdot,\cdot)$ being the cubic root metric, and $\mathbf{i}_i = \text{Render}(\boldsymbol{\pi}_i)$ the generated appearance with optical parameters $m(y_i) = \boldsymbol{\pi}_i$. Note that as we describe later, $\mathbf{i}_i$ is not necessarily a single image; in fact, in the construction of our manifold and its navigation, we use a triplet of three images with the same optical parameters and under different illumination conditions. In the following, we describe our manifold construction from a discrete set of samples, which we navigate by using an interpolation scheme.

6.4.1 *Manifold Construction*

In order to build our manifold, we leverage dimensionality reduction techniques. This allows us to leverage the cubic root metric, which we showed that correlates best with human perception in our scenario (Section 6.3), so that materials with similar appearance are closer together, and different materials are pushed further away. A benefit of this approach is that similarity relations between optical parameters are implicitly handled, as well as non-linearities between optical parameters and resulting appearance.

We sample $M = 100$ tuples of optical parameters $\hat{\pi}$ using the sampling procedure described in Section 6.3.2, and render the *Lucy* statue under three lighting conditions (front, back and side). Then we compute a $M \times M$ distance matrix relating the perceptual distance between each pair of appearances, using the average distance of all three lighting conditions. By associating each point to three different lighting conditions we enforce the manifold to place closer together points that carry optical parameters that generate similar appearances under different lighting conditions, making the manifold more robust.

Finally, we feed the IsoMap algorithm [207] with the distance matrix, and set the number of neighbors to five and the dimensionality of the output space to two. With these settings, we build a two-dimensional manifold where each image triplet $\overline{\mathbf{i}_i}$ generated with optical parameters $\overline{\boldsymbol{\pi}_i}$ has a coordinate $\overline{y_i} \in \overline{\mathcal{M}}$ in the manifold.

DISCUSSION While other algorithms for dimensionality reduction exist (e.g., multidimensional scaling or kernel principal component analysis) we chose the IsoMap algorithm since it works well with arbitrary distance matrices and handles well non-linear spaces [207]. Other alternatives using neural networks, following previous works on perceptual spaces for material appearance [32, 123, 204], were considered, but we found that IsoMap works well in our case, and produces stable and meaningful manifolds that can be easily navigated, while the latent spaces of neural-based methods might lead to unintuitive spaces and would require re-training with extensive labeled data for translucent appearance.

6.4.2 *Manifold Navigation*

The manifold defined above provides a convex hull of appearances for translucent materials. However, the discrete set of samples it is built with fill the 2D manifold unevenly, and do not allow for smooth navigation. To build a continuous manifold inside the convex hull of the manifold from a discrete set of samples, we follow an approach similar to Wills et al. [233], and estimate the optical parameters $\boldsymbol{\pi}_i$ from the finite set $\overline{\mathcal{M}}$ of data points $(\overline{\mathbf{i}_j}, \overline{\boldsymbol{\pi}_j}, \overline{y_j})$ used for building the manifold.

Computing $\boldsymbol{\pi}_i$ directly from the optical parameters associated with the finite set of data points $\overline{\boldsymbol{\pi}_j}$ does not necessarily result into a perceptually meaningful appearance manifold. The reason is the complex interactions between optical parameters and final appearance, arising due to the non-linear effect of each different parameter, as well as appearance metamerism due to well-known similarity relations between optical parameters [255], where very different optical parameters might result in a similar appearance. Instead, we compute the interpolation between datapoints in appearance space (i.e., in image space), and compute $\mathbf{i}_i$ as

$$\mathbf{i}_i = \frac{1}{K_N} \sum_{j=1}^{N} K(|y_i - \overline{y_j}|)\, \overline{\mathbf{i}_j}, \tag{6.2}$$

with $\overline{y_j}$ and $\overline{\mathbf{i}_j}$ the position and appearance of the N nearest data points to y_i (we set $N = 5$), $K(\cdot)$ a radially symmetric kernel (we use a triangle kernel), and $K_N = \sum_{j=1}^{N} K(|y_i - \overline{y_j}|)$ the normalization factor. This approach has two benefits: First, it makes navigation inside the manifold smooth and predictive. Second, it imposes that the relationship in Equation 6.1 holds for the generated appearance $\mathbf{i}_i$ at y_i and the rest of the manifold, assuming a dense enough set of datapoints generating a manifold. Figure 6.2 shows the manifold for the *back* light condition.

Unfortunately, Equation 6.2 results into a relatively simple image-space interpolation, with no direct relationship between y_i and $\boldsymbol{\pi}_i$. Instead, for building the map between y_i and $\boldsymbol{\pi}_i$ we need to compute $\mathbf{i}_i$, and from that compute $\boldsymbol{\pi}_i = \text{Render}^{-1}(\mathbf{i}_i)$, as we describe in the next subsection.

DISCUSSION While this interpolation has some similarities to the one performed by Wills et al. [233], note that they interpolate the optical parameters in the parameter space, which is only feasible in the low-dimensional space of gloss in opaque BRDFs. In the case of translucent materials, this approach leads to suboptimal results, and causes non-smooth transitions inside the manifold. Additionally, while the interpolation in Equation 6.2 does not guarantee that there is a set of parameters $\boldsymbol{\pi}_i$ that would generate the appearance in $\mathbf{i}_i$ (i.e., $\mathbf{i}_i$ might not be physically meaningful), we found that there is always a close match between $\boldsymbol{\pi}_i$ and $\mathbf{i}_i$. Both cases are shown in Section 6.4.4.1 and Figure 6.3.

6.4.3 *Optical Parameters Retrieval*

We leverage physically-based inverse rendering for computing the inverse problem $\boldsymbol{\pi}_i = \text{Render}^{-1}(\mathbf{i}_i)$, by solving

$$\min_{\boldsymbol{\pi}_i} d\left(\mathbf{i}_i, \text{Render}(\boldsymbol{\pi}_i)\right), \tag{6.3}$$

with $d(\cdot,\cdot)$ our distance metric (cubic root metric), averaged over the three different illuminations. We solve for $\boldsymbol{\pi}_i$ using the Adam optimizer [121] (learning rate $lr = 0.025$, 100 iterations), with gradients computed using the physically-based differentiable rendering in Mitsuba 3 [106]. We initialize the optimization by interpolating the optical parameters following

$$\boldsymbol{\pi}_{i,0} = \frac{1}{K_N} \sum_{j=1}^{N} K(|y_i - \overline{y_j}|)\, \overline{\boldsymbol{\pi}_j}, \tag{6.4}$$

which we empirically found works well in our case. To make our problem tractable, we inverse rendered images with a resolution of 128×128 pixels, as it was a good balance between the quality of results and the time to convergence. In our tests, we found that increasing the resolution for inverse rendering did not have a major impact on the reconstruction error, but resulted in significantly longer reconstruction times. In particular, we found that performing the optimization for a native resolution (512×512) and the lower-resolution one (128×128) provided similar error values in both the recovered optical parameters (MAE of 0.16 vs 0.14, respectively) and error in image space (cubic root metric of 0.037 vs 0.048, respectively). We further analyze the quality of the inverse rendering stage in Section 6.4.4.

6.4.4 *Analysis of the Space*

The resulting two-dimensional manifold can be seen in Figure 6.2. Roughly, the manifold shows two main axes of appearance: From top-left to bottom-right, we can see that the appearances are distributed from glass-like (top-left) to a more opaque diffuse-like appearance (bottom-right) similar to, e.g., marble or jade, which agrees with the findings by Gkioulekas et al. [67]. On the other hand, the horizontal axis distributes appearances from dark (left) to bright (right). Since the main purpose of our generated manifold is to enable navigation of the space of translucent appearance, we evaluate this through a user study by integrating our manifold into an appearance editing interface in Section 6.5. Prior to this, in the remainder of this section, we analyze the resulting manifold in terms of *validity*, *stability* and *generality*. Finally, we discuss the relationship of this manifold with the underlying optical parameters.

6.4.4.1 *Validity*

To assess our approach to build the continuous manifold (Sections 6.4.2 and 6.4.3), we compare it to the case in which the manifold is generated by simply interpolating the optical parameters from the initial set of samples. Our measure of error is the cubic root metric, and results are shown in Figure 6.3. The error using inverse rendering is lower, and also shows that the image-based interpolation results in physically plausible appearance.

6.4.4.2 *Stability*

To analyze the dependence of our manifold on individual material samples, we perform a stability analysis [233]. We do this by building M manifolds following a leave-one-out procedure, with $M-1$ manifold samples being used to build each manifold, leaving one material sample out each time a manifold is built (manifold construction is as described in Section 6.4.1). For each of these M manifolds, we compute its Procrustes distance [15, 116] to our resulting, final manifold. This is done by computing the Euclidean transformation (rotation, reflection, scaling, and translation) between the two embeddings that minimizes the distance between them, i.e., the sum of squared differences between the points in both manifolds. The average normalized Procrustes distance across our M manifolds is $d_P = 0.017$; with distances in the range $[0, 1]$, this value is indicative of a stable manifold.

6.4.4.3 *Generality*

Since we build our manifold based on a single geometry (*Lucy*), we would like to assess whether such manifold remains valid for other geometries. For that, we compare our manifold with the manifolds generated for a set of additional geometries that show different ranges of spatial frequencies and distributions of thickness, which are relevant geometric features for both reflection and transmittance. These additional manifolds are built for the *Armadillo* and *Bunny* from the Stanford Repository, and a *sphere* and a *cube*. All these geometries were re-scaled to have a thickness similar to the pedestal in the *Lucy* statue, so that the level of translucency is comparable. The manifold for each of these geometries is generated using the same procedure as with *Lucy*, described in Section 6.4.1.

We analyze the generality of our approach by performing pairwise comparisons between the manifolds generated for all five different geometries, analyzing the cor-

Figure 6.2: Depiction of our continuous manifold for the *Lucy* geometry under the *back* lighting condition. The convex hull of the manifold is depicted by the dark-red line. Note that the manifold smoothly interpolates between different translucent appearances, naturally following two main high-level dimensions of translucency: glass-likeness (bottom-right to top-left) as described by Gkioulekas et al. [67], and brightness (left to right). These axes are shown in the bottom-right inset. We refer to Appendix B for the depiction of the manifold under other lighting conditions.

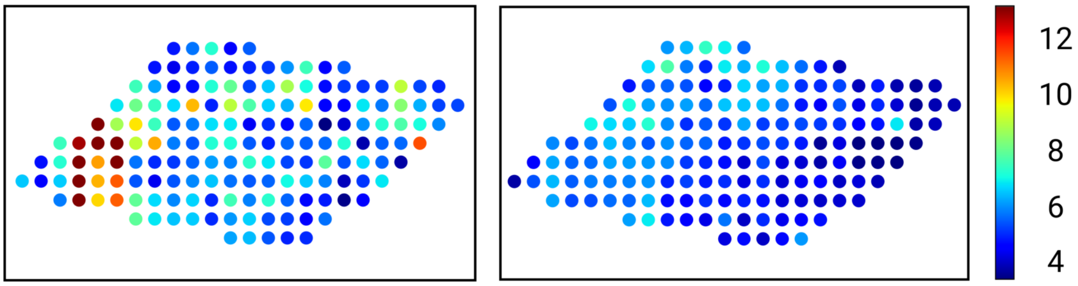

Figure 6.3: Error (cubic root metric) between the interpolated image, $\mathbf{i}_i$, and the rendered one, Render($\boldsymbol{\pi}_i$), for a set of uniformly distributed points along the manifold. We show the error when optical parameters $\boldsymbol{\pi}_i$ are computed through parameter interpolation *(left)*, and by means of inverse rendering *(right)*.

	Armadillo	Bunny	Sphere	Cube
Lucy	0.96	0.94	0.88	0.82

Table 6.2: Pairwise correlation between manifolds generated for different geometries and the *Lucy* statue, computed using the Mantel test. Values correspond to the Pearson product-moment correlation coefficient, which ranges from -1 to 1, with 1 indicating a strong positive correlation and -1 a strong negative one. Results shown are statistically significant ($p < 0.05$).

relation of the distance matrices using the non-parametric Mantel test [144]. Results are shown in Table 6.2, with 1.0 being perfect correlation: Even for very different geometries all manifolds exhibit great consistency. As expected, simple geometries (i.e., sphere and cube) show smaller correlation, in line with previous work [58] that noted that simple geometries are perceived differently from more complex geometries. However, as this complexity increases, the manifolds become extremely consistent, even for geometries with very different distributions of thickness and spatial frequency (e.g., *Bunny* vs *Lucy*).

6.4.4.4 *Underlying Optical Parameters*

Here we analyze the relationship between the underlying optical parameters and the perceptual manifold. These parameters are retrieved from the continuous manifold following the procedure described in Section 6.4.3. Given the complex interplay of optical parameters resulting in strong similarity relationships between albedo, extinction and phase function [67, 255], we opt for analyzing the reduced coefficients (i.e., first-order similarity relationships [235]). Figure 6.4 shows the distribution of the reduced extinction σ_t^* and reduced single scattering albedo α^* inside the manifold: As we can observe, there is a trend that relates the reduced extinction and single scattering albedo with the glass-like appearance (diagonal axis, from top-left to bottom-right), so that lower σ_t^* and α^* increases the glass-like look, and higher values increase the diffuseness, as expected. The brightness (horizontal axis of the manifold), on the other hand, is mostly related with the single scattering albedo. While these trends seem consistent, further exploration would be required to find a complete description of the high-level percep-

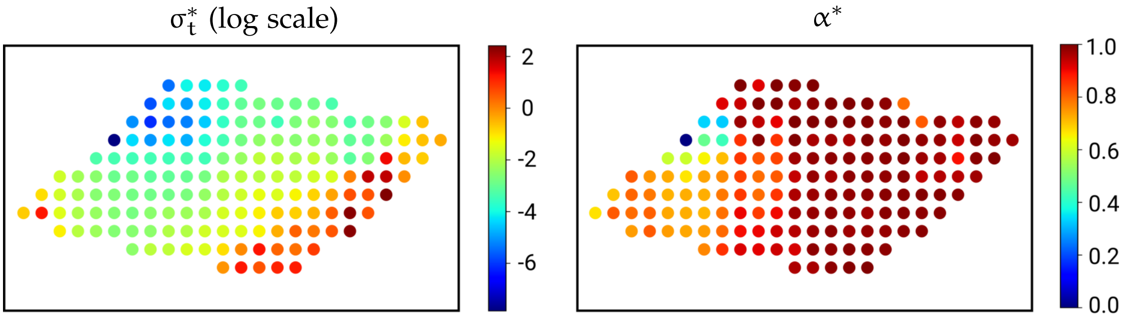

Figure 6.4: Visualization of the reduced optical parameters in our continuous manifold, obtained via inverse rendering. We show the reduced extinction coefficient σ_t^* (*left*, log scale) and reduced single scattering albedo α^* (*right*, linear scale), computed using first-order similarity. Please refer to the main text for an analysis of the observed trends.

tual axes within the manifold, and their potential high-order relationships with the underlying optical parameters.

6.5 AN INTERFACE FOR EDITING TRANSLUCENT APPEARANCE

This section seeks to assess whether our perceptually meaningful space can indeed be effectively used to edit translucent appearance. To do that, we build a prototype editing interface that leverages our manifold, and compare it to a standard editing method (based on tuning sliders for each of the optical parameters) in a user study comprising two editing tasks on a diverse set of materials and geometries.

6.5.1 *Interface Overview*

We build a prototype material editing interface that leverages our proposed translucency manifold. Our interface is integrated into the widespread, open-source modeling software Blender [14]. The main goal of the manifold is to enable users to rapidly obtain a reasonable approximation to the desired final appearance. This approximation can be very hard to reach through individual tuning of four (or more) optical parameters that interact in complex ways. At the same time, it has been shown that sliders perform better than image navigation for *precise adjustments* in matching tasks [118].

Taking all of this into account, we propose a hybrid interface that combines image navigation through a zoomable version of our manifold, and subsequent fine-tuning of the optical parameters using a slider-based interface. Image navigation facilitates an initial good approximation to the desired appearance (which is hard to obtain with sliders alone), and fine-tuning through sliders enables refining this approximation to obtain an optimal match.

Our interface can be seen in Figure 6.5. It has three different areas: one shows the preview rendering of the object whose appearance is being edited, as in most editing interfaces (we use the Cycles built-in progressive renderer for display); another

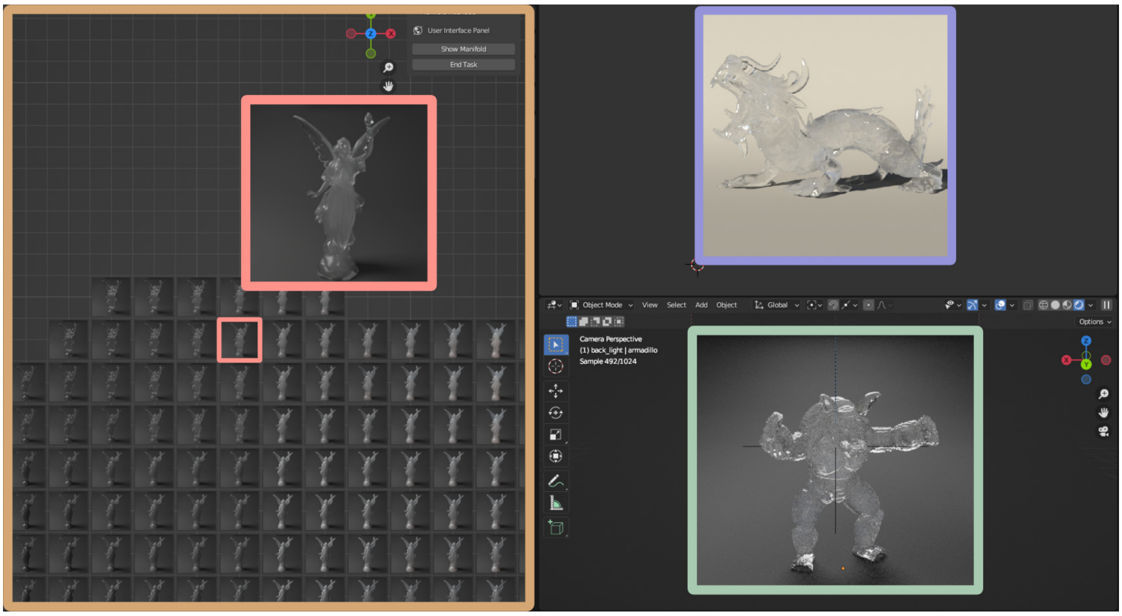

Figure 6.5: Screenshot of our proposed editing interface, as displayed in our user studies, while in image navigation mode. There are three distinct areas. *Left:* Manifold navigation area (orange box), which can be zoomed and panned, and includes a preview of the selected appearance at each instant (pink). *Top right:* Target image area, showing the target image (purple); this area is included for the purpose of our user studies, but can be hidden for normal interface usage. *Bottom right:* Preview area, showing the object being edited, rendered interactively as its appearance is changed (green).

shows the manifold (or a part of it, depending on the selected scale), together with an inset that shows a large version of the selected manifold sample; and a third area is devoted to showing the target image that will be used in the user studies. The third area is thus optional, and could be hidden resulting in a split screen-type interface. The arrow keys allow the user to traverse the manifold, whose scale can be changed to show a coarser or finer view of it. Finally, our interface allows to toggle between the three lighting conditions used to build the manifold, i.e., displaying the manifold with the *front*, *side* or *back* lighting condition (please refer to the supplementary video in our webpage `https://graphics.unizar.es/projects/translucent_manifold/`, for a demonstration of the interface).

6.5.2 *User Study*

We conduct a user study to validate the use of our manifold in an editing application scenario. Participants in the study perform editing tasks on a variety of translucent materials. Each editing task involves a *target* image and a *match* image, whose material appearance will be edited by the participants to match that of the target appearance. They conduct each task with two different interfaces: *ours* (described in Section 6.5.1), and a standard *slider-based* interface, in which the four optical parameters (α, σ_t, g and ρ) are exposed by means of sliders that the participants can tune to edit the material's appearance. The slider-based interface follows the same layout as ours (Figure 6.5), with the four sliders occupying the area reserved for the manifold in our interface.

TASKS The study features two types of tasks. The first type is a *matching task*, in which sixteen participants are asked to match the material of an object under fixed lighting to a target image of the *same* object and lighting. This provides a constrained task that facilitates quantitatively measuring users' performance [118]. In the second type, the *natural task*, ten participants are asked to match the material of an object under fixed lighting to a target image of a *different* object and under a *different* lighting. This is designed to mimic a look-dev use case scenario, in which the user aims to match a certain appearance (e.g., from a photograph or any Internet image) on a new geometry and lighting.

PROCEDURE For each interface and task, participants need to conduct a series of trials, one per target image that needs to be matched. We test four different materials per task; further, in the matching task, we test two different lighting conditions. This results in a total of eight target images in the *matching task* and four in the *natural task*, which need to be matched with each interface. We divide each experiment into two blocks, one per interface; the order of the blocks is randomized for each participant. Each block consists of four trials (i.e., matching four target images), one for each of the four target materials. Lighting conditions are pseudo-randomized so that each participant sees two trials with each lighting condition.

Participants start the experiment conducting a short demographics questionnaire, and then proceed to a training session. In the training session, they were asked to use the two interfaces to match the appearance of a glass-like material applied on the Stanford *Bunny*. Then, they perform one block, followed by a 5-minutes break, and later the second block. After each trial, participants answer a question regarding their level of satisfaction with the result, on a 1 to 7 Likert scale. In addition to that, we log time to completion, optical parameters of the final result, and toggling between lighting conditions events. There is no time limit to complete the tasks, and participants underwent a short questionnaire and debriefing session at the end of the study.

PARTICIPANTS Sixteen subjects participated in the matching task (four female, twelve male, none identifying themselves as non-binary, not listed, or preferring not to say). The average age was 24.8 years, with a standard deviation of 3.58 years. For the natural task, we recruited ten participants (four female, and six male, none identifying themselves as non-binary, not listed, or preferring not to say) among the sixteen initial subjects; in this way, participants were already familiar with both interfaces.

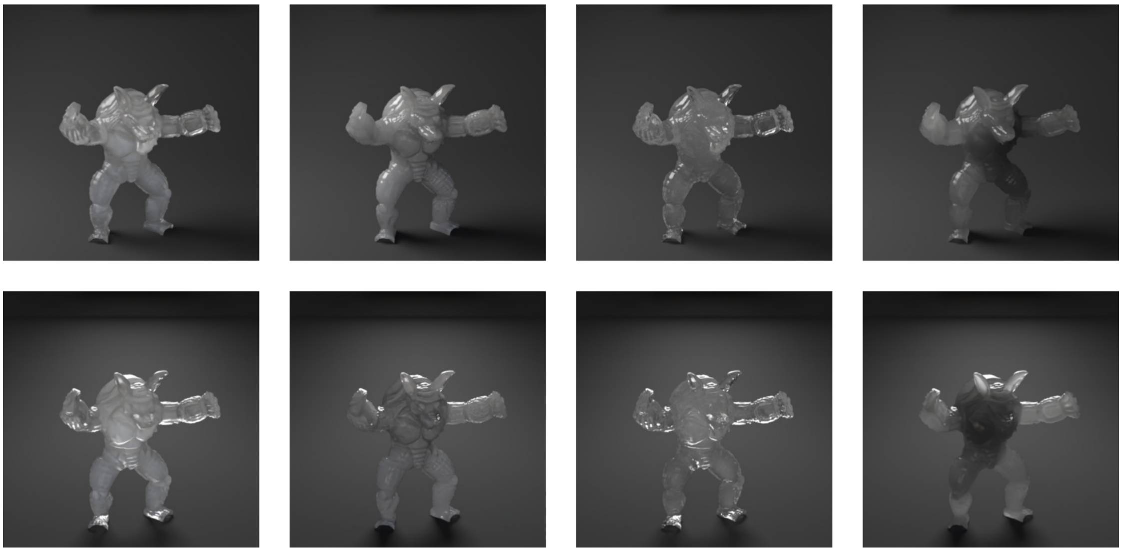

Figure 6.6: Target images used in the *matching task* of the user study. We use the *Armadillo* statue with four different materials (*columns*), and each of these materials is rendered under two different light conditions: back (*top row*) and side (*bottom row*) lighting. Materials are selected such that they cover different regions of the manifold.

6.5.2.1 *Matching Task*

STIMULI We use the *Armadillo* statue for both the target image and match scene for this task. With this choice, we not only test the interface but also its ability to generalize to other geometries, since the manifold features the *Lucy* statue. We choose two lighting conditions: the *side* and *back* conditions described in Section 6.3.2.1, since they convey translucent appearance better than the *front* one. Stimuli are rendered with four target materials, selected to span a range of translucent appearance. Figure 6.6 shows the eight target images for the matching task (4 materials × 2 lighting conditions).

MEASURED VARIABLES Since it is a matching task, we can compute the perceptual *error* (cubic root metric) between the target and the result obtained by each participant. We also analyze the *time* to completion of each trial, as well as the level of *satisfaction* provided by participants at the end of the trial (value in a 1 to 7 Likert scale).

RESULTS We use repeated measures ANOVA to test whether the interface used (two levels, *slider-based* vs. *ours*) had a significant effect on the perceptual error or the time to completion. We use Bonferroni correction, and set the threshold for significance at $\alpha = 0.05$. In the case of the satisfaction level, which is a rating on a 1 to 7 Likert scale, we use the non-parametric Wilcoxon signed-rank test to evaluate statistical significance.

Figure 6.7 shows all three measured variables for both interfaces tested. We can see that the results obtained are better in the case of our proposed interface than for the slider-based one, as shown by both the error and the level of satisfaction. Satisfaction (Figure 6.7, right) is not only significantly higher with our interface ($Z = 204.00$, $p-value < 0.0001$) but also has a lower standard deviation, while the error

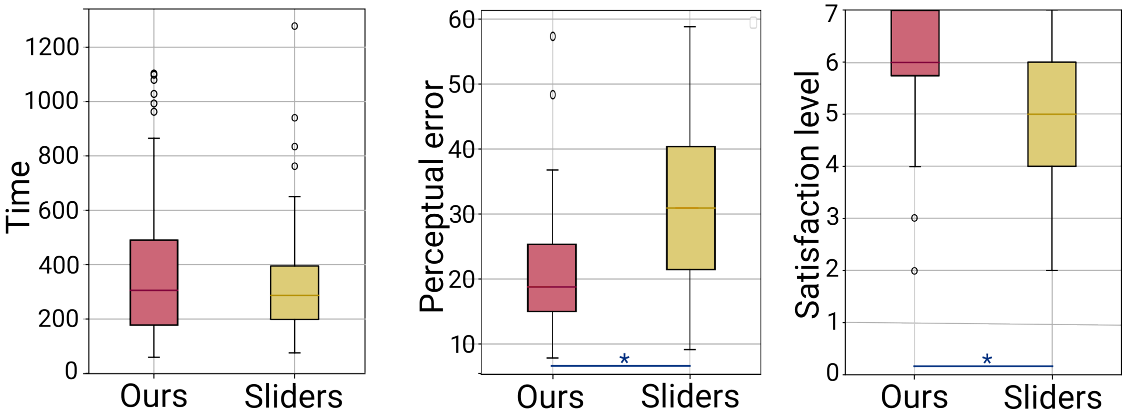

Figure 6.7: Comparison between interfaces for the matching task of the user study. We show, for each interface, the average time per trial (*left*), the error in image space computed using the cubic root metric (*center*), and the reported level of satisfaction with the result (*right*). Users are able to obtain better results, as shown by the increased satisfaction level and decreased error, using similar time than with the slider-based interface.

with respect to the target (Figure 6.7, center) is significantly lower for our interface ($F(1,15) = 12.28$, $p - value = 0.003$). While the average editing time is longer with our interface, differences between both interfaces in terms of time to completion are not statistically significant; despite this, participants are more satisfied with our interface, possibly owing to the increased quality of the results.

We show qualitative results obtained by the participants in our study in Figure 6.8. For three different target images, we show the final appearance obtained by two participants, with both the slider-based and our interface. Within the complexity of the task, our interface enables more closely matching the target, while in the case of the slider-based interface, participants can really struggle to reach even a fairly similar appearance. Interestingly, when using our interface, in 42% of the trials participants did not fine-tune their results, and kept the result obtained through image navigation of the manifold. However, in a number of cases the fine-tuning step can play a significant role, as shown in the example in Figure 6.9.

Overall, our study indicates that, while sliders can be successfully tuned for fine edits (as shown in Figure 6.9), navigating the full space of translucent appearance with sliders is a daunting task for users. Meanwhile, image navigation provides a usable alternative, either to reach a good match, or to provide a good initial approximation in more challenging cases.

6.5.2.2 *Natural Task*

STIMULI The target images feature the *XYZ RGB Dragon* lit by an environment map, and are shown in Figure 6.10. The match scene is the same as the one in the previous task (see Section 6.5.2.1), therefore participants can select between three conditions for

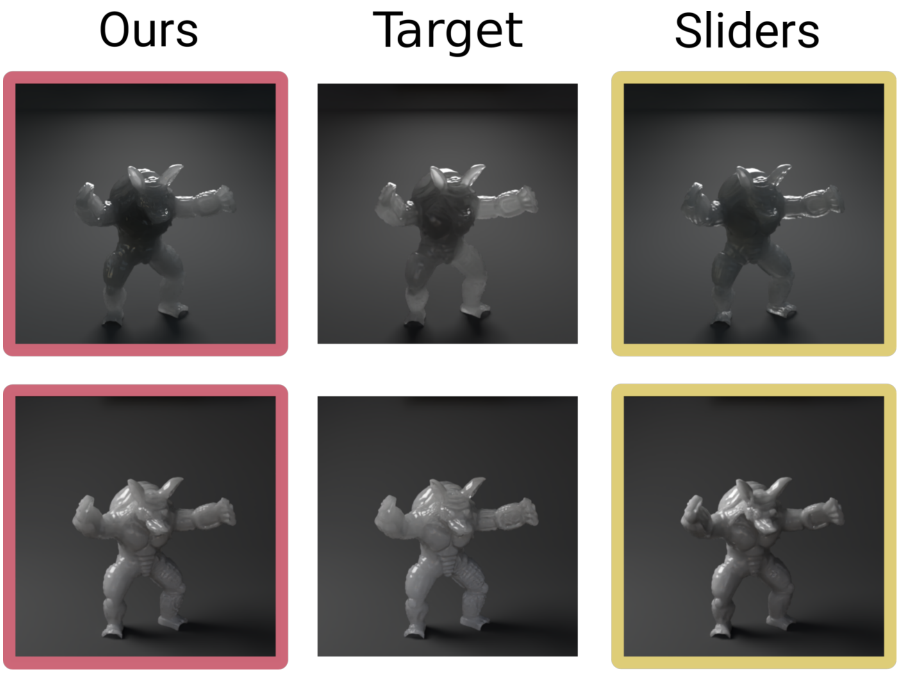

Figure 6.8: Qualitative results obtained in the matching task with the interfaces tested. For two target images, we show the results obtained with each interface (*slider-based* and *ours*) by two different participants.

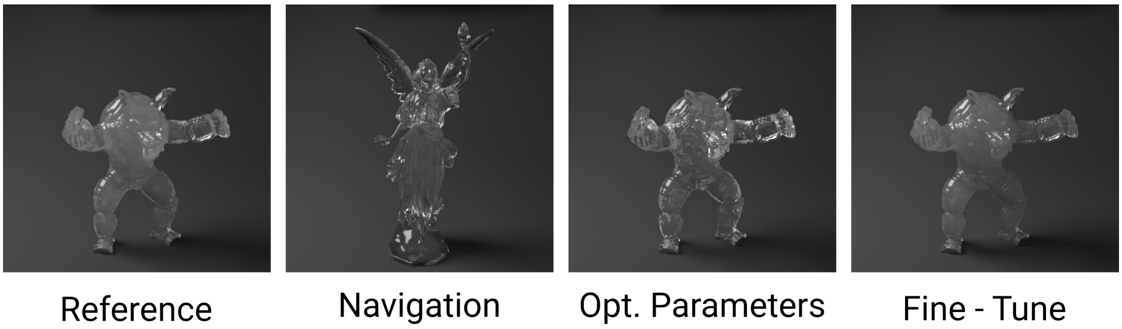

Figure 6.9: Fine-tuning with our interface. Given a target image (*left*), we show the final image selected in the manifold during image navigation (*center left*) and the match scene rendered with the corresponding optical parameters (*center right*). Its appearance can then be fine-tuned via sliders to reach a final result (*right*) closer to the target. Please refer to the text for more details.

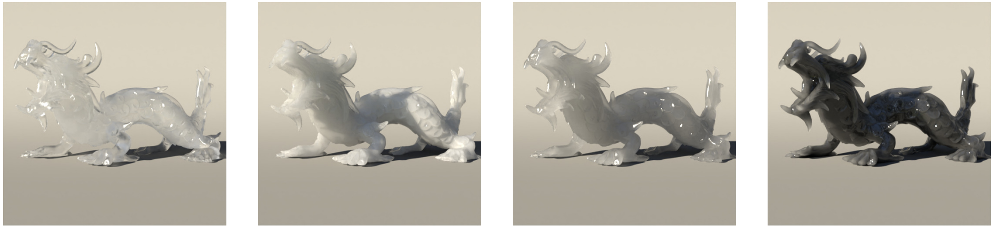

Figure 6.10: Target images used in the natural task of the user study. We employ the *XYZ RGB Dragon* geometry with four different materials, rendered with an environment map.

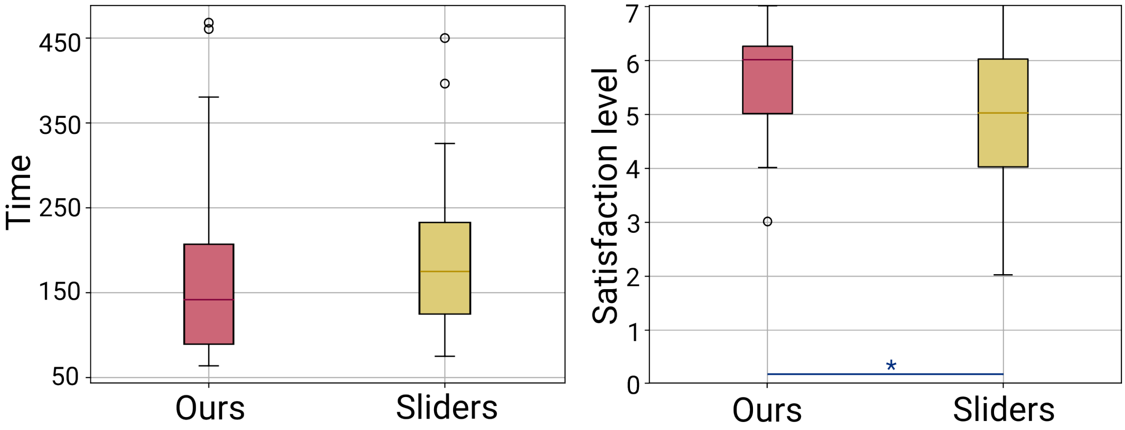

Figure 6.11: Comparison between interfaces for the natural task of the user study. We show, for each interface, the average time per trial (*left*), and the reported level of satisfaction with the result (*right*).

lighting the scene (side, back, and front). Again, target images are rendered with four materials covering a wide range of appearance.

MEASURED VARIABLES In this case, in the absence of an objective measure of error (since the target and match images have different geometry and illumination), our measured variables are the *time* to completion of each trial, and the level of *satisfaction* provided by participants at the end of the trial.

RESULTS We perform the same analyses as in the previous task for the two measured variables, time and level of satisfaction. Figure 6.11 shows these measures for both interfaces tested. While participants spent on average less time editing with our interface than with the slider-based one, this difference was not statistically significant. However, our interface yields a higher level of satisfaction with the result ($Z = 107.5, p-value = 0.01$). Both measures (especially level of satisfaction) exhibit less variance in the case of our interface. We also observe a drop in the time to complete each trial with respect to that of the matching task, with only a slight decrease in the level of satisfaction. Feedback during the debriefing sessions suggests that participants attempted to *approximate* the appearance, rather than exactly matching it. Sample qualitative results can be seen in Figure 6.12.

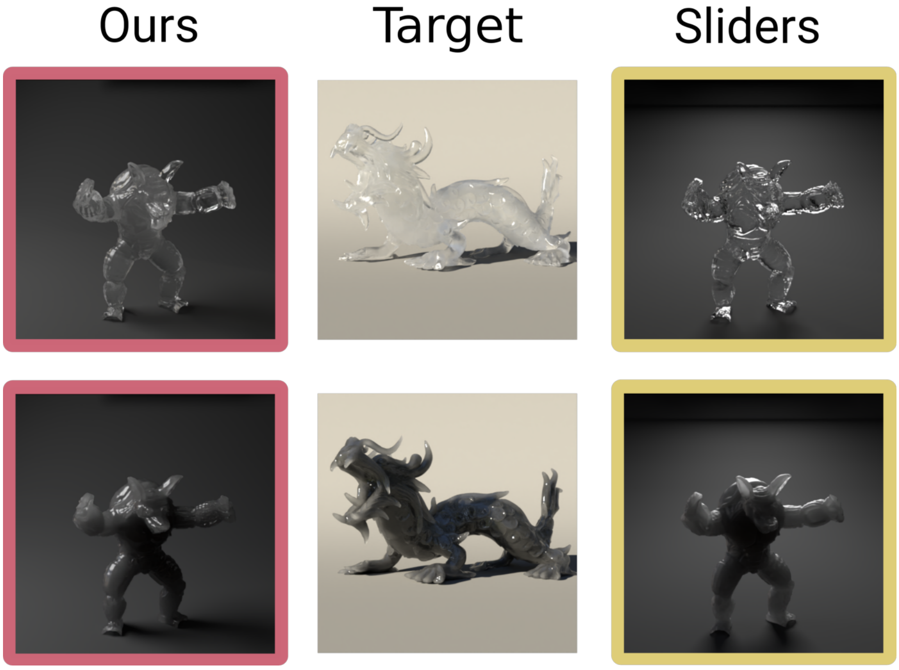

Figure 6.12: Qualitative results obtained in the natural task with the interfaces tested. In this task, participants had to mimic the appearance of a target image on a match scene that featured a different geometry and a different illumination. Despite the difficulty of the task, our interface led to significantly higher levels of satisfaction. The times to complete the task with our interface were 90s *(top)* and 218s *(bottom)*, while for the slider-based interface were 325s and 265s, respectively.

6.6 DISCUSSION

We propose a novel way of exploring and editing translucent materials, via a continuous perceptually-motivated manifold of translucent appearance. To validate the benefits of this approach we have built an interface that leverages our manifold in an image navigation paradigm, followed by a fine-tuning step. An evaluation of its applicability is done by means of a user study, in which we compare our proposed interface to a standard, slider-based one, in two editing tasks and for varied materials and geometries. Our interface enables non-expert users to reach a better result (based on both objective and subjective measures) with differences between both interfaces being statistically significant.

In particular, participants report higher levels of satisfaction, rating our interface as more intuitive and less frustrating than the traditional slider-based approach. Furthermore, for matching tasks, our navigation (plus slider-based fine-tuning) allowed a closer match to the reference, as well as a faster workflow for general look-dev tasks. In terms of workflow, the usual trend was to find a rough estimate of the appearance through navigation, and then locally explore this estimate, first using navigation, and then (sometimes) by fine-tuning using sliders. This is in line with the "block and refine" strategy found by Kerr and Pellacini [118]. At the same time, our results somewhat contradict the findings of Kerr and Pellacini, where a navigation-like paradigm was rated as the least effective, in that case for opaque reflectance modeling. This might be explained by the discrete nature of their image-based navigation through the reflectance parameters, as opposed to the continuous navigation on a perceptually-motivated manifold in our case, as well as by the fact that we focus on a different type of material.

The image-based navigation of our manifold requires computing the optical parameters via inverse rendering, which can lead to differences between the image in the manifold and the actual result. In most cases this was not noticeable, likely because translucency constancy is not good enough to perceive these small differences across different lighting or geometry. The Appendix B material compiles detailed participants' results illustrating this point.

The perception of translucency depends on geometry and lighting, and thus strictly speaking our analysis is only valid for the geometry and lighting considered. We address this by (i) including three lighting directions (commonly used in translucency perception); (ii) selecting a geometry with thin and thick parts, fine details and sharp edges; and (iii) studying generalization to other geometries, including both complex and simpler ones (Section 6.4.4.3). We choose directional lighting because it better reveals the translucent nature of the material, and our natural task shows that users can leverage this manifold to edit complex geometries lit under an environment map. An analysis of the effect of other lighting conditions, including, e.g., more diffuse ones or other factors, could be done following the methodology laid out in this work.

LIMITATIONS The main limitation of image navigation leveraging our manifold, also pointed out by users, is its limited ability for very fine adjustments. This is particularly relevant for the matching task. This motivates the addition of a slider-based fine-tuning step complementing the image-based navigation, but it could potentially be addressed by a fully continuous navigation interface.

In terms of appearances, we have explored homogeneous, achromatic materials, and we have also limited our study to classical exponential media, to keep the problem tractable. The effect of colored albedo deserves further exploration, and the additional degrees of freedom of non-exponential transmittance [12, 108] have been proven useful in look-dev tasks in production [234], at the cost of further increasing the complexity of editing. Having shown here the benefits that a perceptually-motivated manifold can have for editing translucent appearance, we hope these additional aspects will be explored in future work.

Part IV

CONCLUSIONS

7

CONCLUSIONS

In this thesis, we have presented several contributions to the study of volumetric materials. As the appearance of a material depends both on by the interaction between light transport and the human visual system, we explored both in the context of volumetric materials. While each chapter – each contribution presented – has already dwelled on the conclusions and future work on that particular topic, here we do a final summary and outline some of the key challenges that lie ahead in each area.

7.1 REFLECTANCE MODELS

In the first part of this thesis, we focus on light transport in volumetric materials and discuss novel reflectance models, specifically for cosmetics foundations and pennaceous feathers. In Chapter 3 we model foundation cosmetics as a volumetric material composed of two types of particles: diffusers and platelets. In Chapter 4 we model the coherent back-reflection of medullas as a diffuse material and provide a masking term for the self-occlusion caused by barbs and barbules. In both chapters, the key insights stem from understanding the internal structure of the studied material. This highlights the necessity and importance of fundamental research aimed at better understanding the optical mechanisms that cause certain macro-appearances. Gaining more knowledge about the internal micro- and meso-structures of materials will enhance our capability to better define the problem of macro-appearance, in a first-principled manner.

In this thesis we have used physically-based models; however, other approaches can be leveraged to solve the problem of reproducing the appearance of a material. For instance, data-driven models are becoming increasingly popular. Physically-based models, while not always achieving a perfect match with ground-truth data, exhibit good generalization capabilities for unexpected scenarios. In contrast, fully data-driven models excel at mimicking ground-truth data but often struggle when confronted with novel situations. Physically-based models typically require less data for testing but depend on acquiring highly specific information, such as details about the internal structure of the material. In contrast, data-driven models demand larger amounts of data, though the type of data, such as reflected radiance, are often easier to obtain.

A potential advancement for both approaches lies in leveraging repeated measurements of the same sample and a more ample analysis of both captured and simulated data. On one side, it is important to quantify and report the error incurred by measurement tools, which require multiple measurements of the same sample. On the other, having multiple measurements over the same sample would allow for a statical interpretation between ground-truth data and the error caused by the model (whether physically-based or data-driven). In this case, using a statistical approach, in which the capabilities of the model are not compared over a single function (i.e., a possible realization of a statistical distribution), but over a *distribution* of functions, might be a more general and robust way to compare real data with synthetic data. Such type of analysis can have applications, e.g., when considering dust-based cosmetics, or, more generally, aggregates of complex materials (e.g., textiles or hairs), in which two different samples might yield two drastically different reflectance behaviors.

7.2 PERCEPTION AND EDITING OF TRANSLUCENT MATERIALS

In the second part of the thesis, we focus on a subset of volumetric materials, specifically, translucent materials, and study their perception and editing. Chapter 5 investigates the perception of translucent materials under dynamic lighting conditions. While previous work focused on the perception of translucent materials with static stimuli, we look at how perception is affected when using dynamic stimuli. While we find that dynamic lighting does not generate a significative improvement in density estimation, it might be worth studying what lighting conditions, or shapes, do help with material constancy, similar to the work of Havran et al. [84] for gloss, whose findings were later confirmed by subsequent research demonstrating improved user performance in material constancy [26, 124, 185].

In Chapter 6 we investigate novel editing spaces working in image space, rather than in optical parameter space. We build a manifold of translucent appearance using a distance metric that correlates well with human perception, as we evaluate it through a user study. This manifold, which is two-dimensional, enables a more intuitive navigation of the appearance space than working in optical parameter space. We then develop a hybrid user interface that leverages such manifold and proves its efficiency in two editing tasks, finding objective and subjective improvements over a standard, slider-based interface.

While in this work we have worked in image space, rather than in the optical parameter space, allowing users to visually explore the different nuances of translucency, other editing paradigms should investigated. With the advances of large language models, it is becoming increasingly common to use natural language as a possible input to solve the cumbersome editing task. Using natural language allows for very intuitive editing, through text-based descriptions of what the users desire. However, it is important to consider that natural language alone is not sufficient to describe all the subtleties of reality. For instance, we can utilize non-verbal languages, such as *visual* language, as an alternative medium of conveying meaning, particularly in the context of visual editing faced in this thesis.

As an example, consider colors: While many words can be used to describe different hue colors, they are not sufficient to clearly define all the color nuances, which are best expressed using visual language. To overcome this issue recent work aimed to solve this issue by allowing input of both visual and natural language [20, 73]. In the context of this thesis, we believe that similar concepts can be applied to translucent, and in more in general, to volumetric materials. This becomes clear when we ponder on the possible meaning of expressions such as “more translucent” or “less translucent”.

If we consider translucency as a property of a material that lies between “transparency” and “opaqueness” it becomes hard to define what “more translucent” might mean: Do we refer to an object that has more internal scattering? But if this is true then the phrasing “more translucent” would refer to an object that is opaque rather than translucent, thus leading to personal interpretation of the meaning of such expressions. Still, it is common to describe translucent materials through analogies, using expressions such as “marble-like” or “milk-like” to convey specific visual characteristics. In future work, it would be interesting to integrate both visual and natural language. Such a system would allow users to visually explore all the ranges of “marble-like” materials, while retaining the ability of editing, using the visual language, the proposed materials.

PERSONAL CONCLUSIONS

On the technical side, I am really glad that I had the opportunity to work on topics involving human perception, coming from a background in computer graphics, this has been a refreshing and interesting field to discover. On a personal note, I am happy that this thesis allowed me to mature skills that extend beyond research and can be applied to various aspects of work. First and foremost is how to give and build effective presentations: whether it is teaching a class, presenting your results to your supervisors, or showing a project to your peers, I believe that presenting and explaining your point of view to an audience is an important skill in this world. I always found presenting one of the most fun and thrilling aspects of my work as a researcher. Incidental to this, I have also learned the value of paying attention to detail. From ensuring consistency in capitalization to the careful use of one word instead of another, everything matters and conveys a message. This lesson has fundamentally impacted the way I approach documents, papers, and books. For these lessons, and many others that I do not have the time to recall, I am really glad that I undertook this journey. This thesis has not been exactly a stroll in the park, but rather a quite tortuous hike, but I guess this is what makes a hike more memorable than a Sunday stroll in the park. I had fun and the years spent at the Graphics Imaging Lab in Zaragoza will always be remembered in a bright light.

Part V

APPENDICES

A

ON THE INFLUENCE OF DYNAMIC ILLUMINATION IN THE PERCEPTION OF TRANSLUCENCY - ADDITIONAL DETAILS

We provide here additional details on the definition of the phase function used and extra statistical analysis.

A.1 VON MISHES-FISHER PHASE FUNCTION

We use the phase function proposed by Gkioulekas and colleagues [67] to model anisotropic scattering. It is based on the *von Mises-Fisher (vMF)* distribution, defined as

$$f_{\mathrm{vMF}}(\mu; k) = \frac{k}{k\pi \sinh k} e^{k\,\mu}. \tag{A.1}$$

Each lobe of the phase function is governed by a single parameter $k \in [-100, 100]$, which controls the spread (anisotropy) of the scattering. Table A.1 compiles the values used to render our stimuli.

	k_1	k_2	w
$f_{p,1}$	100	-0.95	0.6
$f_{p,2}$	100	-75	0.9

Table A.1: Parameters for $f_{p,1}$and $f_{p,2}$, the phase functions defined as a mixture of two von Mishes-Fisher distributions (Eq. (5.1)).

A.2 ADDITIONAL ANALYSIS RESULTS

We report here the full results of the repeated-measures ANOVA. Table A.2 reports the results for the case of the error ($\sigma_{t,est} - \sigma_{t,real}$), and Table A.3 for time spent per trial. Please refer to the main text for further details.

Factor	F-Statistic	DF	p-value
light direction	47.3906	(2,22)	0.0000
light motion	3.1918	(1,11)	0.1016
phase function	17.1836	(2,22)	0.0000
direction*motion	0.9743	(2,22)	0.3932
direction*phase function	8.5727	(4,44)	0.0000
motion*phase function	11.1971	(2,22)	0.0004
direction*motion*phase function	1.0834	(4,44)	0.3763

Table A.2: Results of the repeated-measures ANOVA for the error in the estimated density.

Factor	F-Statistic	DF	p-value
light direction	9.0612	(2,22)	0.0013
light motion	0.0000	(1,11)	0.9976
phase function	2.2318	(2,22)	0.1311
direction*motion	0.3799	(2,22)	0.6883
direction*phase function	1.2806	(4,44)	0.2921
motion*phase function	0.9068	(2,22)	0.4184
direction*motion*phase function	0.7840	(4,44)	0.5418

Table A.3: Results of the repeated-measures ANOVA for the time spent per trial.

B

NAVIGATING THE MANIFOLD OF TRANSLUCENT APPEARANCE - ADDITIONAL DETAILS AND RESULTS

In this chapter, we provide additional details and information on the following topics

- Section B.1 Additional details on the perceptual study of distance metrics for translucent appearance
- Section B.2 Details on stimuli generation
- Section B.3 Additional results of the matching task
- Section B.4 Additional results of the natural task

B.1 USER STUDY

B.1.1 *Metrics Definition*

In the following, we provide a small introduction to the various metrics evaluated during the perceptual study for an image-based metric.

CUBIC ROOT METRIC One of the most thorough studies of translucent appearance, by Gkioulekas et al. [67], tested different simple image-based metrics over a variety of geometries and material properties (a total of 6,777 rendered images), and showed that the so-called *cubic root metric* performed best in terms of consistency between geometries, and correlated well with human perception. The cubic root metric was introduced by Ngan et al. [165] in the context of opaque BRDFs, and is defined as

$$d(\mathbf{i}_i, \mathbf{i}_j) = \sqrt{\sum_{k=1}^{N} \left(\sqrt[3]{i_k} - \sqrt[3]{j_k} \right)^2}, \tag{B.1}$$

where i_k, j_k are pixels k of images $\mathbf{i}_i$ and $\mathbf{i}_j$ respectively, in linear color space, with N is the total number of pixels.

ANISOSHADING METRIC In our analysis, we also include a recently developed model for translucent materials [122]. The key idea behind this model is that translucency perception is influenced by the incongruence in 3D shape information provided by specular and non-specular shading patterns. Through psychophysical studies, they show how the difference in the anisotropy of the shading between specular and diffuse image components can be a plausible predictor of perceived translucency.

IMAGE QUALITY METRICS In addition to these two, we also include in our analysis the well-known MS-SSIM, widely used for image comparison, and the recent FovVideoVDP metric [145]. While FovVideoVDP is designed as a video metric, with emphasis on the temporal and peripheral aspects of vision, benchmarking results show that it is among the best performing also on conventional SDR and HDR images. We choose these two metrics because, in the recent Unified Photometric Image Quality (UPIQ) dataset [158], their predictions yield the highest (FovVideoVDP) and second highest (MS-SSIM) Spearman rank order correlation with human judgements [145]. Finally, we also include the deep learning-based LPIPS metrics [253] in our work, as one of the most used metrics in perception studies.

B.1.2 *Perceptual Study*

To increase the robustness of the subjective judgments gathered, we resort to a two-alternative forced choice (2AFC) experimental paradigm [67, 123, 233]. In each experimental trial, the participant is shown a triplet of images: one *reference* and two *candidates*. The participant is then asked to select the candidate stimulus that is more similar to the reference. Figure B.1 shows the graphic user interface used by participants in our perceptual study. In Figure B.2 we show a subset of the stimuli used in our user study, differently from previous work [67] we focus on a different set of appearances encompassing crystal-like, matte, and dark lustrous appearances.

Our stimuli, comprising sixteen sets of material properties, yield 1,680 possible triplets for each light condition. We show a subset of the stimuli used in Fig B.2. Since we have three light conditions, a total of 5,040 triplets need to be evaluated (all three images in a triplet feature the same lighting). In the interest of robustness of the subjective data, we seek that each triplet be answered by at least five different participants, leading to 25,200 trials. We resort to crowdsourcing (through the Amazon Mechanical Turk platform) to gather the perceptual data, which has been shown to be effective in similar studies [85, 110, 123, 128, 185], and include a training stage before the experiment, as well as control triplets to discard unreliable participants.

Amazon Mechanical Turk defines a single trial as a HIT (Human Intelligent Task), in our user study, each HIT is composed of 100 triplets, along which we also add ten control triplets. The purpose of a control triplet is to check whether a participant can be trusted or not by posing a triplet with a clear answer, in this case, the reference stimulus is used also as one of the two candidates. Following pilot studies, in which the control triplets were always passed easily, we settled to discard all the data from an HIT if more than two control triplets are failed. For each successful HIT, we pay a single participant 1.88 euros, since we estimated the completion time of a single HIT to be around fifteen minutes, following the minimum wage of 7.5 euro per hour.

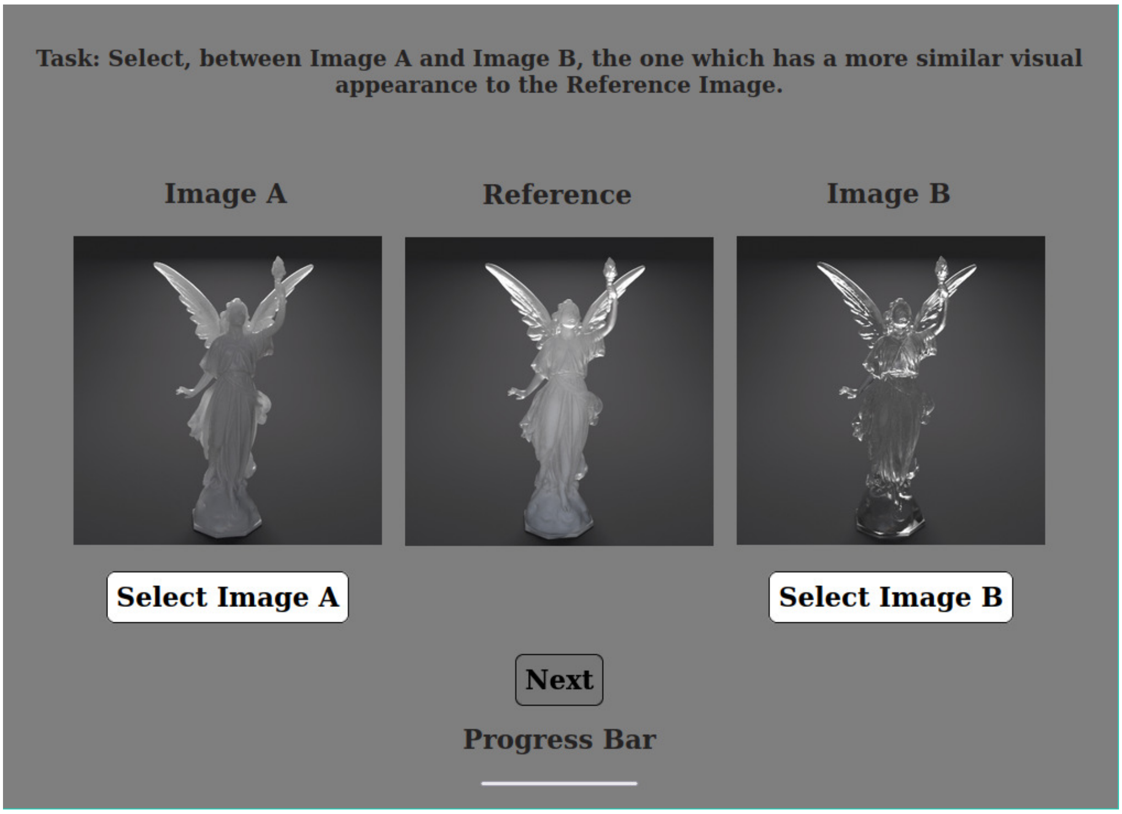

Task: Select, between Image A and Image B, the one which has a more similar visual appearance to the Reference Image.

Figure B.1: User interface used for the perceptual study analyzing the perceptual distance between translucent materials. Participants have to select the image that most closely resembles the Reference stimulus (in the middle). Once participants select and confirm their choice, they can pass to the next triplet.

B.2 MANIFOLD - ADDITIONAL RESULTS

In this section, we showcase additional results when we generate the manifold. Figures B.3 and B.4 show the same manifold presented in the main paper with a side or a front lit illumination.

Figure B.2: We show a subset of the various appearances covered in our user study. Each column shows the same material lit under different light conditions: back (first row), side (second row), and front (third row). The triplet showed in the user interface uses different materials with the same light conditions.

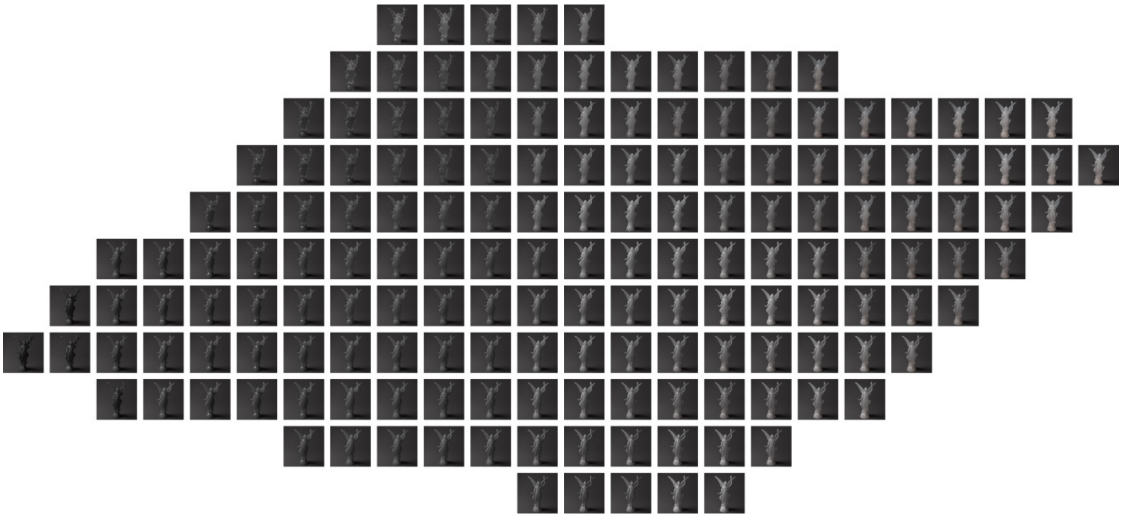

Figure B.3: Manifold obtained using side-lit stimuli

B.3 MATCHING TASK: USERS' RESULTS

In this section, we report the results of our user study. Each table covers all the results a single user obtained throughout the experiment. For each table, we report the Reference target (Ref) the appearance selected during the navigation phase (Ours - Selected) and the renderings obtained using the inverse rendering parameters (Ours - IR). We then show the result of the fine-tuning phase (Ours-Tune). This step was optional and therefore not all users desired to fine-tune their results, for these cases, we report the rendered image obtained using the optical parameters using inverse rendering. The last column shows results using the traditional slider interface (Slider). It is possible

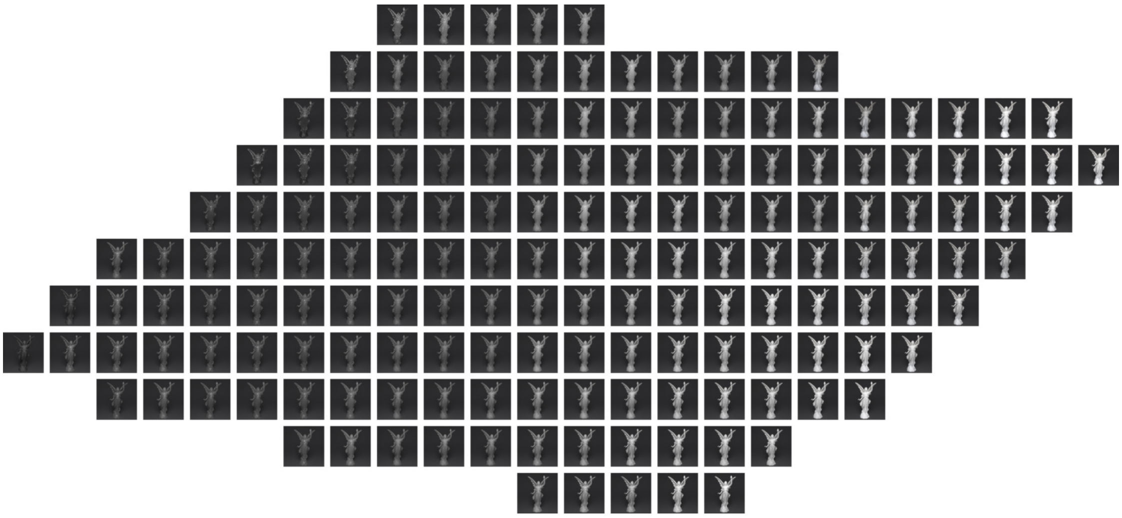

Figure B.4: Manifold obtained using front-lit stimuli

to see some inconsistencies between Ours-Sel and Ours-Manif, these inconsistencies happen due to the inverse rendering step that is not always capable of finding the optical parameters that can reproduce a similar appearance to the one previewed in the second column.

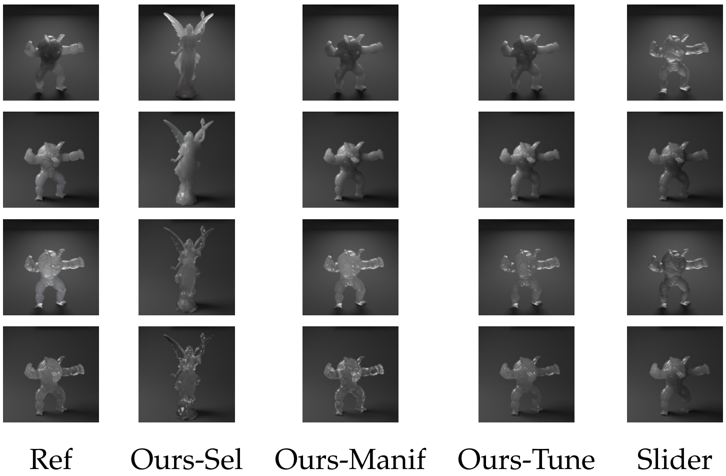

Figure B.5: Results for the *Matching Task* for User 1

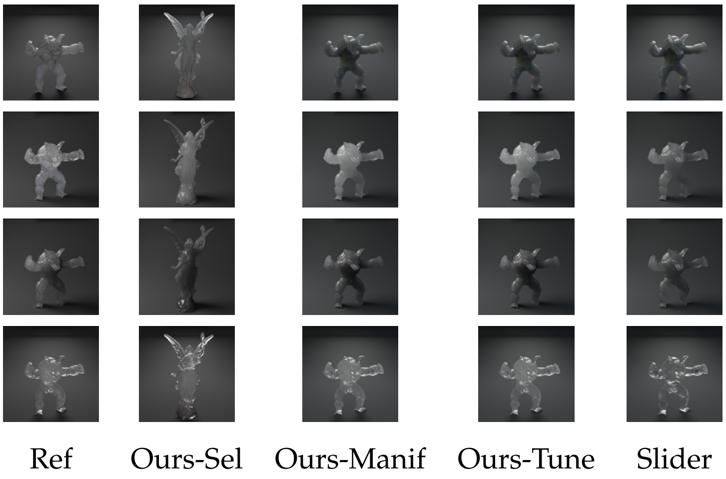

Figure B.6: Results for the *Matching Task* for User 2

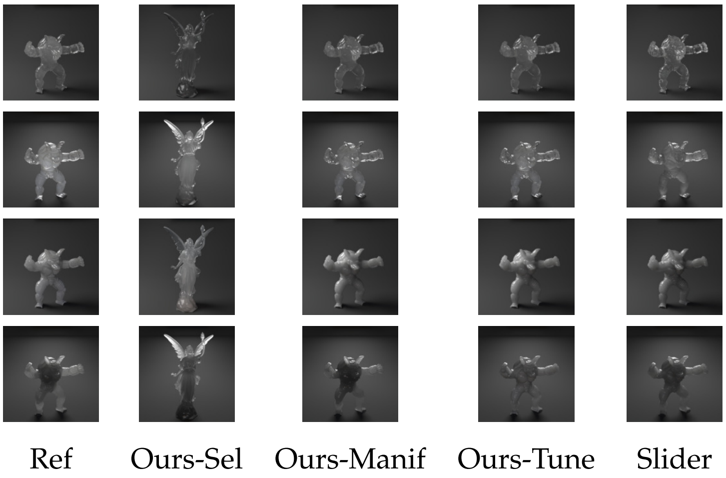

Figure B.7: Results for the *Matching Task* for User 3

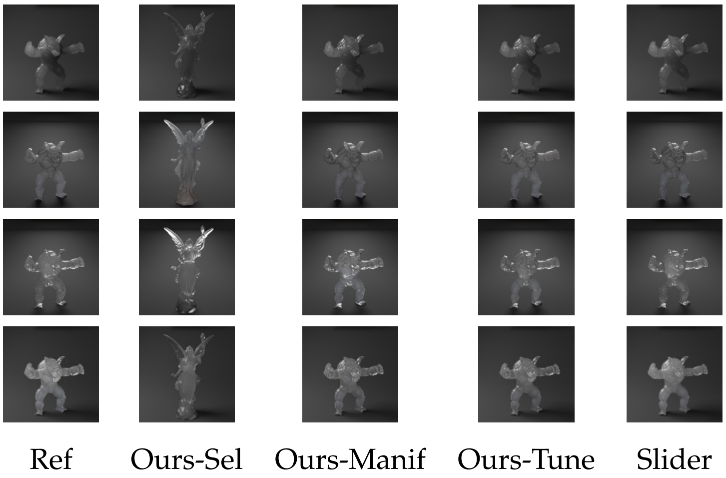

Figure B.8: Results for the *Matching Task* for User 4

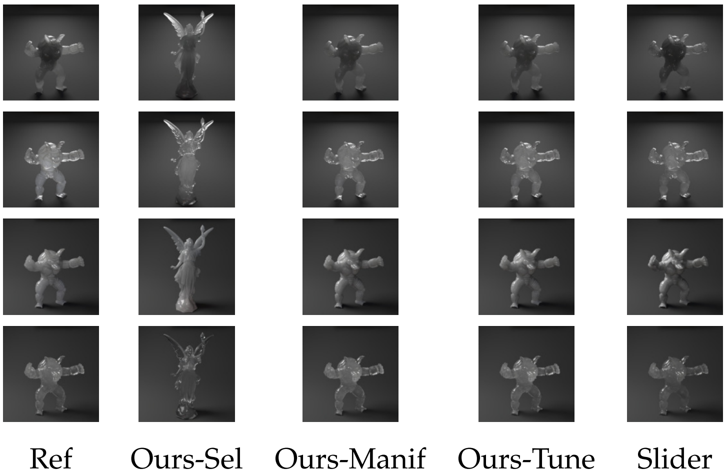

Figure B.9: Results for the *Matching Task* for User 5

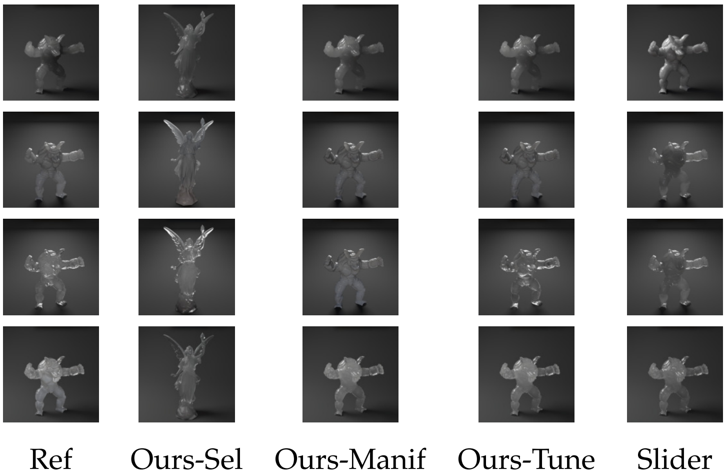

Figure B.10: Results for the *Matching Task* for User 6

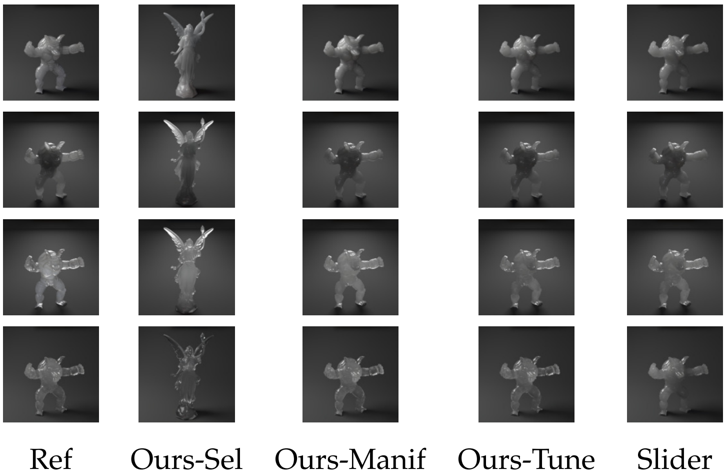

Figure B.11: Results for the *Matching Task* for User 7

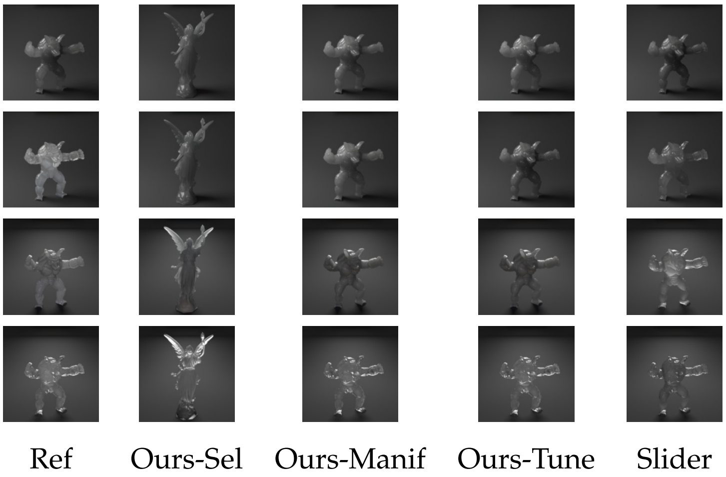

Figure B.12: Results for the *Matching Task* for User 8

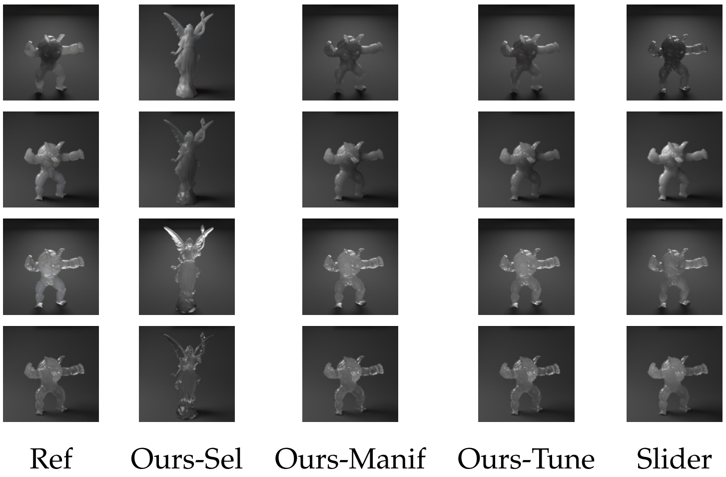

Figure B.13: Results for the *Matching Task* for User 9

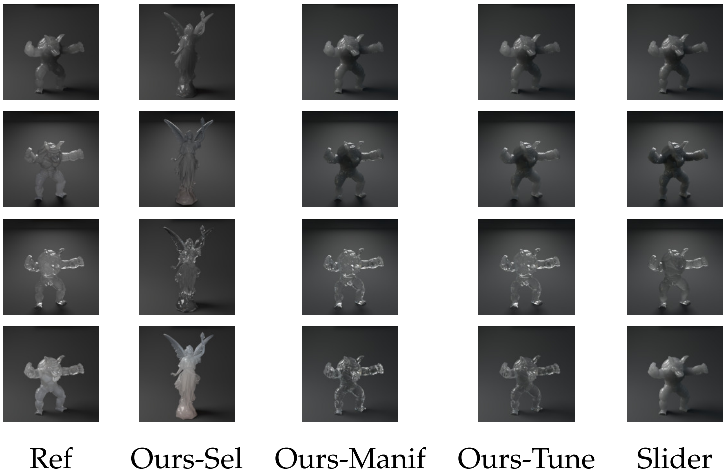

Figure B.14: Results for the *Matching Task* for User 10

Ref Ours-Sel Ours-Manif Ours-Tune Slider

Figure B.15: Results for the *Matching Task* for User 11

Ref Ours-Sel Ours-Manif Ours-Tune Slider

Figure B.16: Results for the *Matching Task* for User 12

Ref Ours-Sel Ours-Manif Ours-Tune Slider

Figure B.17: Results for the *Matching Task* for User 13

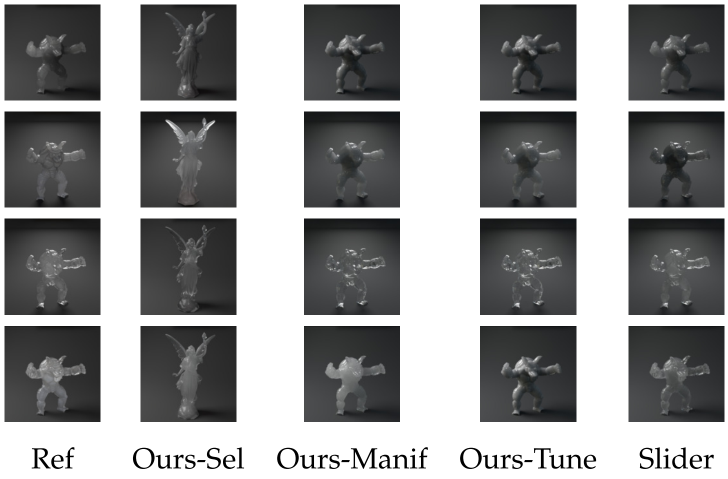

Figure B.18: Results for the *Matching Task* for User 14

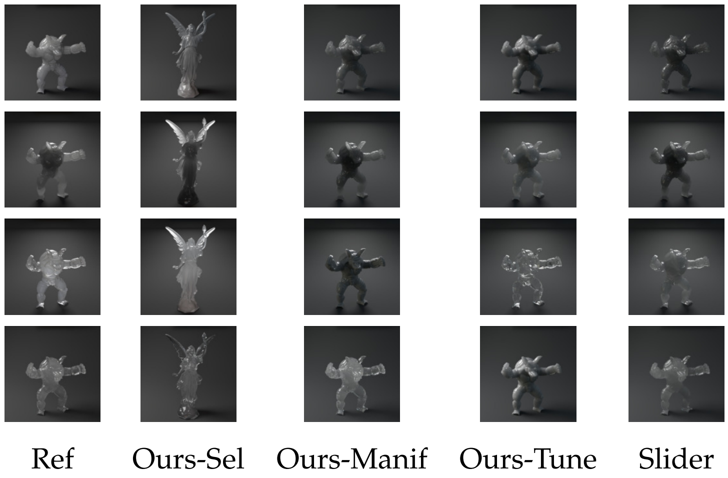

Figure B.19: Results for the *Matching Task* for User 15

b.4 *natural task*: users' results

In this section, we report the results of our user study. Each table covers all the results a single user obtained throughout the experiment. For each table, we report the Reference target (Ref) inverse rendering parameters obtained by navigating the manifold (Ours - Manif). We then show the result of the fine-tuning phase (Ours-Tune). This step was optional and therefore not all users desired to fine-tune their results, for these cases, we report the rendered image obtained using the optical parameters using inverse rendering. The last column shows results using the traditional slider interface (Slider). The task was very challenging, especially for novice users, and in many instances, the results of using the two interfaces are equally distant from the reference.

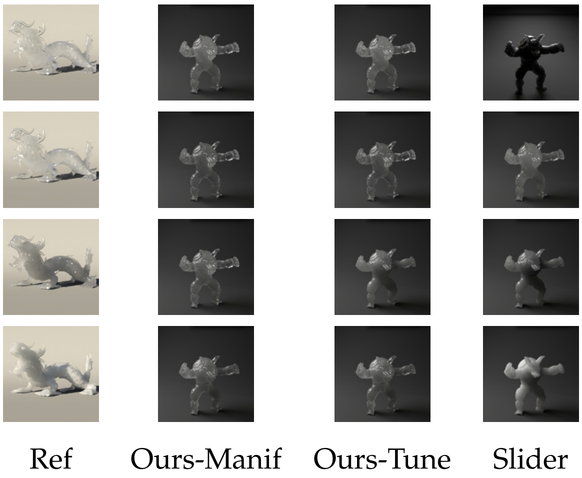

Figure B.20: Results for the *Natural Task* for User 1

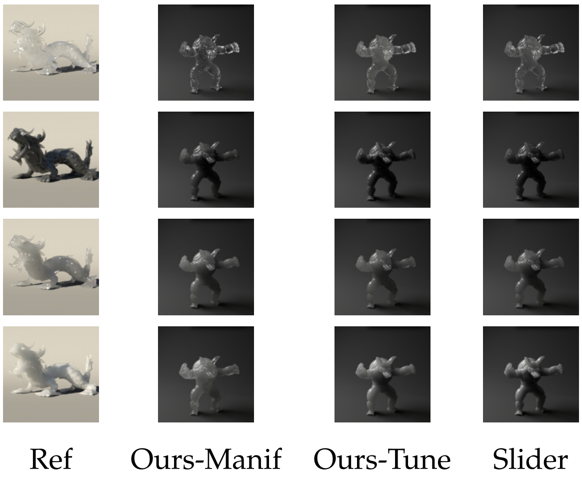

Figure B.21: Results for the *Natural Task* for User 2

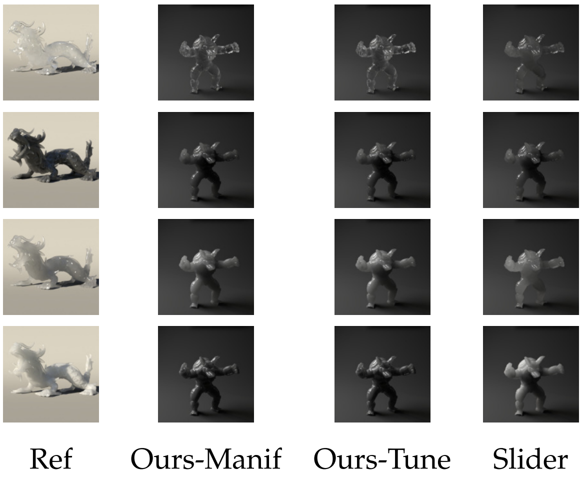

Figure B.22: Results for the *Natural Task* for User 3

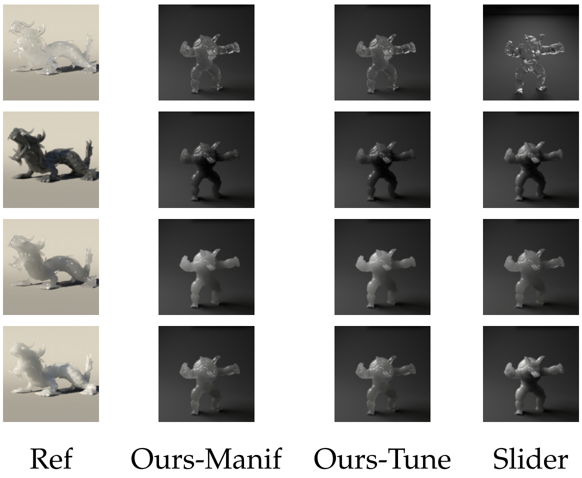

Figure B.23: Results for the *Natural Task* for User 4

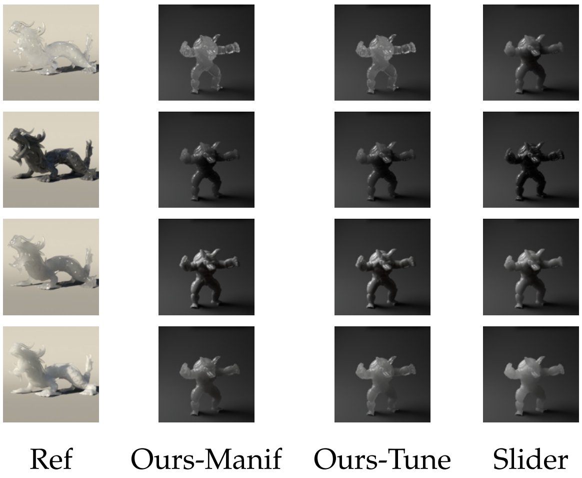

Figure B.24: Results for the *Natural Task* for User 5

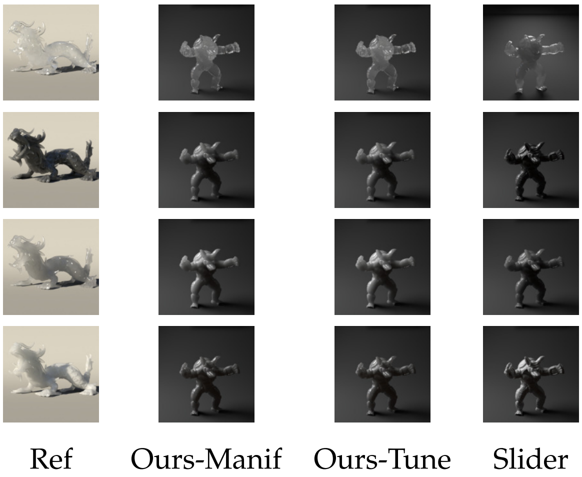

Figure B.25: Results for the *Natural Task* for User 6

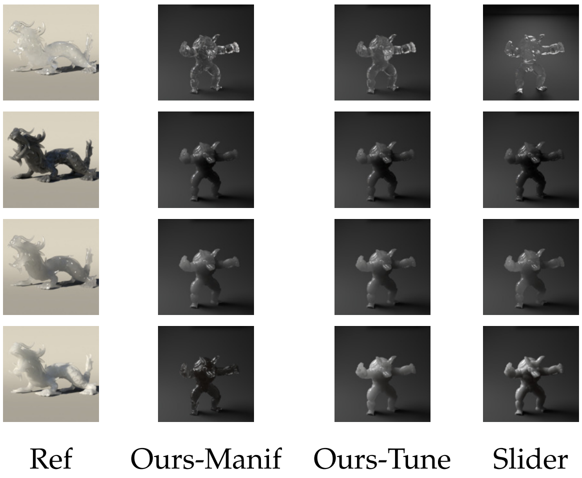

Figure B.26: Results for the *Natural Task* for User 7

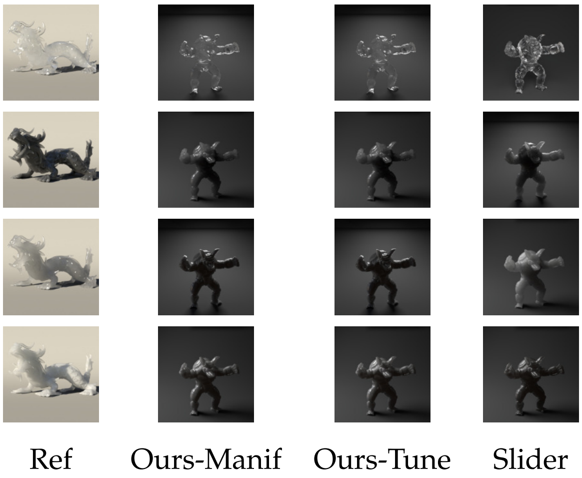

Figure B.27: Results for the *Natural Task* for User 8

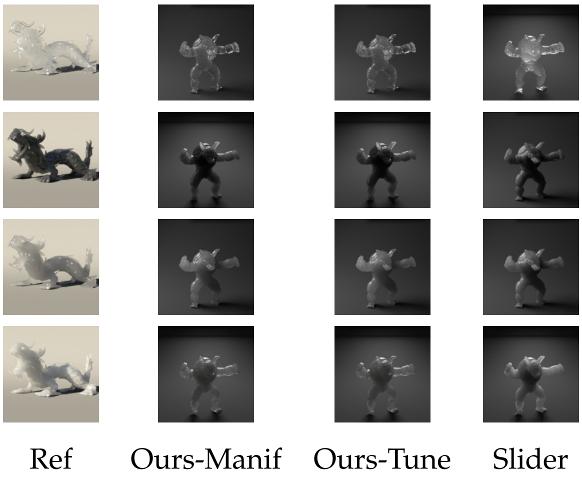

Figure B.28: Results for the *Natural Task* for User 9

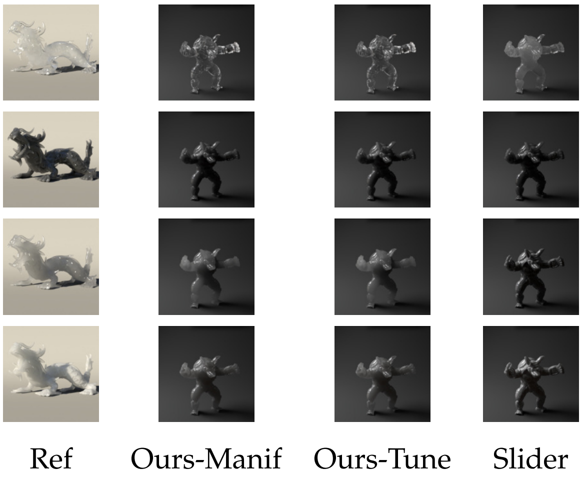

Figure B.29: Results for the *Natural Task* for User 10

C

PRACTICAL APPEARANCE MODEL FOR FOUNDATION COSMETICS - ADDITIONAL DETAILS

C.1 OPTIMIZED PARAMETERS AND RENDERING TIME

In the following, we report the optimized parameters (see Table C.1) for the four foundation samples measured and renderings time for some of the figures in Table C.2. Notice how our cosmetic layer only introduces a small overhead in terms of rendering time despite our work being more focused on rendering accuracy than computational performance and we do not perform any sophisticated performance optimization strategies.

Parameter	Dewy 1	Dewy 2	Matte 1	Matte 2
α_p	0.13	0.002	0.15	0.49
θp	1.47°	5.5°	2.01°	52.2°
α_p	(0.75,0.85,0.99)	(0.8,0.8,0.8)	(0.91, 0.89,0.85)	(0.93,0.83,0.71)
c_d	0.90	0.99	0.89	0.29
α_d	(0.99,0.98,0.95)	(0.97,0.93,0.87)	(0.95,0.94,0.91)	(0.9,0.9,0.84)
g_1	0.55	0.77	0.37	0.24
g_2	0.09	-0.23	-0.25	-0.22
w_g	1.0	0.79	0.62	0.51
t	16	16	16	16

Table C.1: Optimized parameters fit from our measured reflectance data using our cosmetic appearance model for several cosmetic products.

Figure	Resolution	SPP	Time	Thickness
7 (First row - Matte 1)	1223 × 1269	128	11.4 min	0.5
7 (First row - Matte 2)	1223 × 1269	128	11.0 min	0.5
7 (First row - Reference)	1223 × 1269	128	9.25 min	0.5
7 (First row - Dewy 1)	1223 × 1269	128	9.95 min	0.5
7 (First row - Dewy 2)	1223 × 1269	128	10.95 min	0.5
7 (Second row - Matte 1)	1332 × 1187	512	25.3 min	0.35
7 (Second row - Matte 2)	1332 × 1187	512	24.8 min	0.35
7 (Second row - Reference)	1332 × 1187	512	24.8 min	0.35
7 (Second row - Dewy 1)	1332 × 1187	512	24.2 min	0.35
7 (Second row - Dewy 2)	1332 × 1187	512	25.3 min	0.35
11 (First row - no foundation)	1332 × 1187	512	25.3 min	0.0
11 (First row - one layer)	1332 × 1187	512	25.3 min	0.25
11 (First row - two layers)	1332 × 1187	512	25.3 min	1.25

Table C.2: Rendering times for the the figures in the paper.

D

A SURFACE-BASED APPEARANCE MODEL FOR PENNACEOUS FEATHERS - ADDITIONAL DETAILS AND RESULTS

D.1 ELLIPSE EQUATIONS

In this Section, we explore the different geometrical representations of the ellipse, as well as the properties and derivations of those representations that our work uses for ray intersections, tangent point detection and projected area integrals.

An ellipse is characterized by its center $\mathbf{c} = c_x, c_y$ and its two axes a_x and a_y.

Its **implicit equation** for a two-dimensional point $\mathbf{p} = p_x, p_y$ is:

$$\frac{(p_x - c_x)^2}{a_x^2} + \frac{(p_y - c_y)^2}{a_y^2} - 1 = 0 \qquad . \tag{D.1}$$

We can obtain the **normal** $\mathbf{n} = n_x, n_y$ for each point from the gradient of the left part of the Eq. D.1, which yields:

$$\mathbf{n} = \begin{pmatrix} n_x \\ n_y \end{pmatrix} = \begin{pmatrix} 2\frac{p_x - c_x}{a_x^2} \\ 2\frac{p_y - c_y}{a_y^2} \end{pmatrix} \tag{D.2}$$

which should be normalized afterwards.

The **parametric equation** of this same ellipse is based on an angular parameter $\theta \in [-\pi, \pi)$:

$$\mathbf{p} = \begin{pmatrix} p_x \\ p_y \end{pmatrix} = \begin{pmatrix} c_x + a_x \cos\theta \\ c_y + a_y \sin\theta \end{pmatrix} \tag{D.3}$$

We can also obtain an expression for the **normal direction** according to this parametric version as described in Eq. D.3 by calculating the direction perpendicular to the tangent (derivative), as:

$$\mathbf{n} = \begin{pmatrix} n_x \\ n_y \end{pmatrix} = \begin{pmatrix} -a_y \cos\theta \\ -a_x \sin\theta \end{pmatrix} \tag{D.4}$$

which should be normalized afterwards.

We can also invert Equation D.3 in order to **obtain the parameter** θ from a point in the ellipse $\mathbf{p}$ as:

$$\theta = \operatorname{atan2}\left(\frac{p_y - c_y}{a_y}, \frac{p_x - c_x}{a_x}\right) \tag{D.5}$$

Another operation on ellipses we will need to do (for casing shadows / masking) is to find the points in the ellipse for which the tangent is a specific direction $\mathbf{d} = d_x, d_y$, which is equivalent to saying that the normal $\mathbf{n}$ is perpendicular to $\mathbf{d}$, which combined from Equation (D.4) yields:

$$\begin{aligned} \mathbf{n} \cdot \mathbf{d} &= 0 \\ -a_y \cos\theta_t d_x - a_x \sin\theta_t d_y &= 0 \\ a_y d_x + a_x \tan\theta_t d_y &= 0 \\ \tan\theta_t &= \frac{-a_y d_x}{a_x d_y} \\ \theta_t &= \tan^{-1}\frac{-a_y d_x}{a_x d_y} + k\pi \qquad k \in \mathbb{Z} \end{aligned} \tag{D.6}$$

For the purpose of casting, the only solutions that we need for θ_t in Equation (D.6) are for $k \in \{0, 1\}$ because other solutions generate the same set of points from the parametric ellipse, as described in Eq. D.3. Those two casting points are:

$$\begin{aligned} \mathbf{p_0} &= \begin{pmatrix} c_x + a_x \cos\left(\tan^{-1}\frac{-a_y d_x}{a_x d_y}\right) \\ c_y + a_y \sin\left(\tan^{-1}\frac{-a_y d_x}{a_x d_y}\right) \end{pmatrix} \\ \mathbf{p_1} &= \begin{pmatrix} c_x + a_x \cos\left(\tan^{-1}\frac{-a_y d_x}{a_x d_y} + \pi\right) \\ c_y + a_y \sin\left(\tan^{-1}\frac{-a_y d_x}{a_x d_y} + \pi\right) \end{pmatrix} \end{aligned} \tag{D.7}$$

Also, it is common for fiber-based microgeometry modeling to use a projected diameter according to a casting direction $\mathbf{d}$, a rectilinear segment parametrized by $h \in (-1, 1)$ that generates ray origin points $\mathbf{o}$:

$$\mathbf{o}(h) = \mathbf{p_0}\frac{h-1}{2} + \mathbf{p_1}\frac{h+1}{2} - l\mathbf{d} \tag{D.8}$$

where l is a distance large enough so that the segment does not intersect with the ellipse. In our case we use $l = 10\max(a_x, a_y)/|\mathbf{d}|$. For the inverse operation (finding the parameter h of the projection of a given point in the ellipse) we trace a ray from the point towards direction $-\mathbf{d}$ and estimate the parameter h within the segment by inverting Equation (D.8).

In order to calculate the **intersection with a 2d ray** we need to start from the implicit equation as described in Equation (D.1) and define a system of equations with the equation of the 2d ray with origin $\mathbf{o}$ and direction $\mathbf{d}$: $\mathbf{p} = \mathbf{o} + t\mathbf{d}$. By substituting the point in the implicit equation with the equation of the ray we get:

$$\begin{aligned} \frac{(o_x + td_x - c_x)^2}{a_x^2} + \frac{(o_y + td_y - c_y)^2}{a_y^2} - 1 &= 0 \\ \frac{t^2 d_x^2 + 2td_x(o_x - c_x) + (o_x - c_x)^2}{a_x^2} + \frac{t^2 d_y^2 + 2td_y(o_y - c_y) + (o_y - c_y)^2}{a_y^2} - 1 &= 0 \\ t^2\left(\frac{d_x^2}{a_x^2} + \frac{d_y^2}{a_y^2}\right) + t\left(\frac{2d_x(o_x - c_x)}{a_x^2} + \frac{2d_y(o_y - c_y)}{a_y^2}\right) + \left(\frac{(o_x - c_x)^2}{a_x^2} + \frac{(o_y - c_y)^2}{a_y^2} - 1\right) &= 0 \end{aligned} \tag{D.9}$$

which is a second degree equation which obtains all the t parameters (if any) of the intersections.

The last interesting operation to do with an ellipse is to calculate the integral of the **projected "area"** from a direction **d** within a parameter range $[\theta_0, \theta_1]$, which is with respect to the unnormalized normal direction (which corresponds to the area differential) in Equation (D.4):

$$\begin{aligned}\int_{\theta_0}^{\theta_1} \mathbf{n}(\theta) \cdot \mathbf{d} d\theta &= \int_{\theta_0}^{\theta_1} -d_x a_y \cos\theta - d_y a_x \sin\theta d\theta \\ &= d_x a_y (\sin\theta_0 - \sin\theta_1) + d_y a_x (\cos\theta_1 - \cos\theta_0) \qquad \text{(D.10)}\end{aligned}$$

D.2 MASKING

In this supplemental Section we describe the mathematical and algorithmic details of our masking analytical expression at both levels (barbule and barb level), that were left out from the main paper due to their level of detail. They are included here for reproducibility purposes.

D.2.1 *Barbule masking*

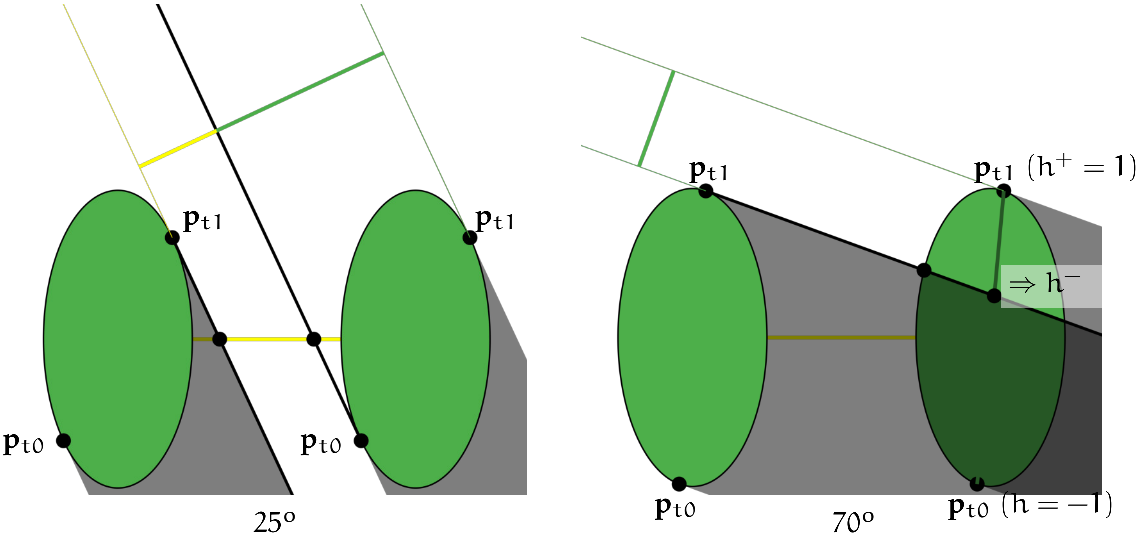

Figure D.1: **Masking between two barbules at different view inclinations.** Barbules (and their projected area) are marked in green. The barbule separation is modeled as the yellow segment. The traced rays are the black lines. We first trace a ray from the tangent point $\mathbf{p}_{t1}$, generating two cases. If it intersects the other ellipse (70° case) we need to intersect the ellipse's diameter to identify $\mathbf{p}_d$ and therefore h^-. The separation (transmittance) cannot be seen ($w'_{bb} = 1$). If it does not intersect the other ellipse (25° case) then we need to intersect the separation segment with two different rays. We then calculate the local barbule weight w'_{bb} from the projected areas a_{bb} and a_s.

The barbule masking expression is computed twice, for both proximal and distal barbules, at their respective local coordinate systems given by the transform matrices $\mathbf{T}_{bp}$ and $\mathbf{T}_{bd}$, respectively. We will name $\omega'_o = \{\omega'_{ox}, \omega'_{oy}\}$ the two-dimensional direction in the local coordinate space (the longitudinal dimension is ignored) after transforming ω_o. From this expression we will obtain the integration range limits for each barbule BCSDF ($h^-_{bp}(\omega_o)$,$h^+_{bp}(\omega_o)$, $h^-_{bd}(\omega_o)$ and $h^+_{bd}(\omega_o)$) and local barbule (w'_{bp} and w'_{bd}, respectively) between 0 and 1, that are later used at the barb masking component. The local transmittances (visibility through barbules) are $1 - w'_{bp}$ and $1 - w'_{bd}$, respectively.

We explicitly model two barbules as ellipses while the separation between them is a segment (as represented in Figure D.1). The barbule model parameters are, for reducing the number of parameters, the barbule axis ratio e_{bb} and the relative barbule separation δ_{bb}.

For identifying the visibility boundaries among the two barbules and the separation segment, we first cast a two-dimensional ray, following the local masking direction ω'_o. The origin point for such ray is the visibility boundary of the ellipse, that is, the point in the ellipse that is tangent to ω'_o. For an ellipse with center $\mathbf{c} = \{c_x, c_y\}$ and axes a_x and a_y, there are two tangent points:

$$\mathbf{p}_{t0} = \begin{pmatrix} c_x + a_x \cos(\psi) \\ c_y + a_y \sin(\psi) \end{pmatrix} \qquad \mathbf{p}_{t1} = \begin{pmatrix} c_x + a_x \cos(\psi + \pi) \\ c_y + a_y \sin(\psi + \pi) \end{pmatrix} \tag{D.11}$$

where $\psi = \tan^{-1}\left(\frac{-a_y \omega'_{ox}}{a_x \omega'_{oy}}\right)$.

We choose which ellipse casts to the other one and which of the two points we use as ray origin depending on the quadrant of ω'_o. If this ray intersects the other ellipse at point $\mathbf{p}_{bb}$, there is no transmittance through the barbules: $w'_{bb} = 1$. The lower limit of the integration range h^- ($h^-_{bp}(\omega_o)$ or $h^-_{bd}(\omega_o)$) require tracing a the same ray again towards the segment between $\mathbf{p}_{t0}$ and $\mathbf{p}_{t1}$ that represents the diameter of the ellipse. h^- is the parameter of the parametric equation of this diameter at the intersection point $\mathbf{p}_d$. The other integration range limit is $h^+ = 1$. These geometrical elements can be visualized on Figure D.1 (left).

If the ray does not intersect the other ellipse, then the barbule is fully visible ($h^- = -1$ and $h^+ = 1$) and we intersect this ray and the separation segment between both ellipses, obtaining an intersection point $\mathbf{p}_{s0}$. We also trace a new ray from the other tangent point of the other ellipse towards the same separation segment and find another intersection $\mathbf{p}_{s1}$. Then, the projected areas (a_{bb} for the barbule and a_s for the separation segment) are:

$$a_s = \omega'_o \cdot (\mathbf{p}_{s1} - \mathbf{p}_{s0})^{\perp} \qquad a_{bb} = A\left(\psi, \psi + \pi\right) \tag{D.12}$$

where $A\left(\beta_0, \beta_1\right)$ is the projected area of the barbule between two angles of its parametric form, calculated as:

$$A\left(\beta_0, \beta_1\right) = \omega'_{ox} a_y (\sin(\beta_0) - \sin(\beta_1)) + \omega'_{oy} a_x (\cos(\beta_1) - \cos(\beta_0)) \tag{D.13}$$

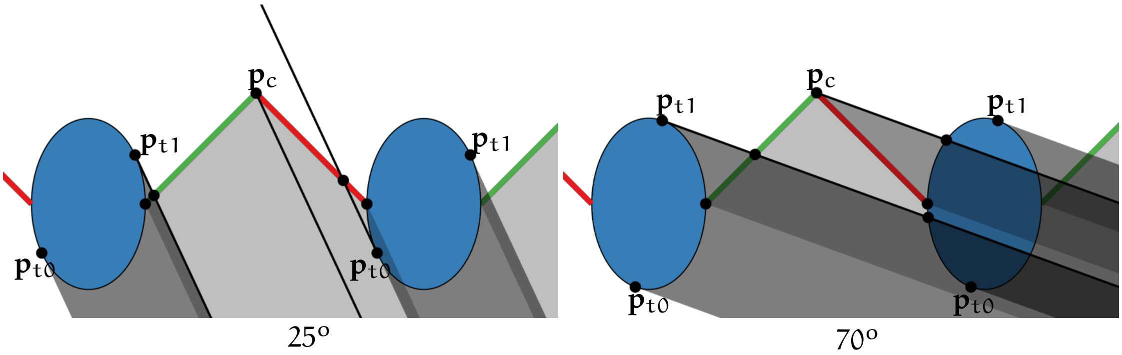

Figure D.2: Cross section of barbs, representing the masking between barbs at two different view inclinations. Barbules (in red and green, respectivelly) are partially transmittant , depending on the view direction at their particular local coordinates (see Figure D.1), while barbs (in blue) are considered to be opaque. Depending on the view direction, each element (barbs and barbules) totally or partially occludes the rest. The limits of such oclussions are identifying by tracing 2D rays (marked in black). By considering all the particular cases at each element (barb to barb oclussion, barb to barbules, barbules to barb and barbules to barbules) we obtain the rates of each element.

and the local barbule weight is then calculated as the relative projected area: $w'_{bb} = a_{bb}/(a_{bb} + a_s)$. A diagram that represents this geometrical components is shown in Figure D.1 (right).

Ellipse intersections, tangent points and projected areas are derived from the implicit and parametric equations of the ellipse. Such derivations can be found on Section D.1

D.2.2 *Barb masking*

The barb masking term is similar to the barbule term, albeit slightly more complex. We model two barbs as ellipses, as well as two sets of barbules as segments, symmetric with a specific inclination θ_{bb} (so we are actually considering the plane at which the barbules lie, no matter their orientation). By doing so, we are assuming that barbules' thickness is negligible compared to barbs thickness, which is often the case: they tend to be one order of magnitude thinner than barbs [142]. As parameter models, and similar to barbules, all distance parameters are relative with respect to the horizontal size of barbs, including barb axis ratio e_b and barbule length l_{bb} and inclination θ_{bb} (with respect to the horizontal plane). local barbule weights w'_{bp} and w'_{bd} from the barbule masking expression. Barbule segments are aligned with the horizontal axis of the plane, so we can obtain three significant points: the corner between barbule segments ($\mathbf{p}_c$), the limit points of the proximal and distal barbules at the barb ($\mathbf{p}_{bp}$ and $\mathbf{p}_{bd}$, respectively).

Similar to the barbule masking term, we cast rays from one of the tangent points of both ellipses, obtained as described by Equation (D.11), towards the contiguous ellipse and both barbule segments. This gives the location of the boundary of the full occlusion for barbs ($\mathbf{p}_{b\to b}$) and barbules ($\mathbf{p}_{b\to bp}$ and $\mathbf{p}_{b\to bd}$, respectively). We also cast a ray from the corner $\mathbf{p}_c$ to the next barb and next proximal barbule, obtaining the corresponding intersection points $\mathbf{p}_{c\to b}$ and $\mathbf{p}_{c\to bd}$, in order to identify the boundaries of partial occlusion: as there is separation between barbules (as obtained by the barbule masking expression) it is possible that there is some partial visibility through the barbules. All these points are shown for diferent cases on Figure D.2.

Then we obtain projected areas of the barb (a_b), proximal barbule (a_{bp}), distal barbule (a_{bd}) and transmittance (a_t, that represents visibility through the vane). There are multiple combinations of cases regarding partial or global oclussion, from one or both barbules to the barb or to each other of the barbules, depending on whether each of the three different rays intersect each of the microgeometrical components.

If the ray from the barb does not intersect the proximal barbule, then $\mathbf{p}_{b\to bp}$ does not exists and it is fully occluded ($a_{bp} = 0$). Otherwise:

$$a_{bp} = w'_{bp} \max\left(0, \left(\omega_o \cdot (\mathbf{p}_c - \mathbf{p}_{b\to bp})^{\perp}\right)\right) \tag{D.14}$$

where $(\mathbf{p}_c - \mathbf{p}_{b\to bp})^{\perp}$ represents the perpendicular (unnormalized) direction of the visible part of the proximal barbule.

The distal barbule can be directly visible or visible through the proximal barbule depending on their orientation with respect to the view direction ω_o, so it is calculated as:

$$a_{bd} = \begin{cases} w'_{bd} \max\left(0, (\omega_o \cdot (\mathbf{p}_{bd} - \mathbf{p}_{b\to bd})^{\perp})\right) & \text{if } (\omega_o \cdot (\mathbf{p}_{bd} - \mathbf{p}_c)^{\perp}) > 0 \\ (1 - w'_{bp}) w'_{bd} \max\left(0, (\omega_o \cdot (\mathbf{p}_c - \mathbf{p}_{b\to bd})^{\perp})\right) & \text{otherwise} \end{cases} \tag{D.15}$$

where the max function helps identifying whether the distal barbule is fully occluded by the barb ($\mathbf{p}_{b\to bd}$ is outside the segment that represents the distal barbule).

For the transmittance, if barbs occlude each other ($\mathbf{p}_{b\to b}$ exists) then $a_t = 0$. If not, if the distal barbule is visible, $\left(\omega_o \cdot (\mathbf{p}_{bd} - \mathbf{p}_c)^{\perp}\right) > 0$, then we trace a new ray for the other tangent point of the other ellipse towards the distal barbule, obtaining $\mathbf{p}_t$. Then:

$$\begin{aligned} a_t = & (1 - w'_{bp}) \max\left(0, \left(\omega_o \cdot (\mathbf{p}_c - \mathbf{p}_{b\to bp})^{\perp}\right)\right) \\ & + (1 - w'_{bd}) \max\left(0, \left(\omega_o \cdot (\mathbf{p}_t - \mathbf{p}_{b\to bd})^{\perp}\right)\right) \end{aligned} \tag{D.16}$$

Otherwise, only the proximal barbule cotributes to transmittance, so this ray is traced towards the distal barbule, obtaining $\mathbf{p}_t$. Then:

$$a_t = (1 - w'_{bp}) \max\left(0, \left(\omega_o \cdot (\mathbf{p}_t - \mathbf{p}_{b\to bp})^{\perp}\right)\right) \tag{D.17}$$

For the barbs, we obtain h^- following the same procedure than with the barbules (ignoring barbule occlusion). The projected area is obtained by sorting all relevant points on the ellipse: $\mathbf{p}_{t0}$, $\mathbf{p}_{b\to b}$, $\mathbf{p}_{bd}$, $\mathbf{p}_{c\to b}$ and $\mathbf{p}_{t1}$ according to their parametrical angle (using the inverse of the parametric function). Depending on the ordering this leads to up to three different parameter ranges (ignoring the part of the barb fully occluded by the other barb below $\mathbf{p}_{b\to b}$). The projected area of each of those ranges is calculated using Equation (D.13), multiplied by a factor depending on the potential partial occlusions, $(1-w'_{bp}$ or $(1-w'_{bd})$, if the range is occluded by the proximal or distal barbule, respectively, $(1-w'_{bp})(1-w'_{bd})$ if occluded by both and 1 if unoccluded). The contribution of all those ranges is added into a_b.

The final weights for the BSDF ($w_b(\omega_o)$, $w_{bp}(\omega_o)$, $w_{bd}(\omega_o)$ and $w_t(\omega_o)$) are obtained by normalizing all the projected areas. For any particular weight i:

$$w_i(\omega_o) = \frac{a_i}{\sum_j a_j} \tag{D.18}$$

There is a particular corner case that rarely happens, that has been omitted for the sake of simplicity and efficiency. When viewed at grazing angles, barbules could partially occlude not only the barb but the next barbules (one after the other) leading to an infinite sequence of partial occlusions that is approximated by a single occlusion in our case, with minimal effect on the final appearance.

D.3 APPEARANCE EXPLORATION

In this section, we explore some of the potential appearances of the most interesting parameters of our appearance model for feathers. For these experiment, we create a matrix of renders for a feather wing under sunlight conditions, where we vary only two parameters we are interested and let the rest of parameters fixed. In particular, we include experiments of the following configurations: (1) barb eccentricity vs barbule longitudinal angle, (2) barbule length vs barbule separation, (3) barb roughness vs barb eccentrity, (4) barb roughness vs barbule rougness and (5) Barb eumelanin vs pheomelanin.

In Figure D.3 we explore the relationship in terms of appearance between the barbule inclination and the barb eccentricity. As we increase the barb eccentricity the feather becomes slightly brighter in some regions (close to the rachis) and darker in others (upper right vane). On the other hand, a larger barbule longitudinal angle (barbule inclination) results into a larger influence of the barbule BCSDF for some regions of the barn depending of the barb angle given by the geometry texture.

In Figure D.4 we analyze how the different configurations between barbule length and barbule separation affect the masking term in terms of appearance. The larger the barbule separation, the more transparent the feather becomes. On the other hand, the larger the barbule length is, the more prominent the barbule appearance is.

In Figure D.5, the barb roughness parameter controls the transition between a specular surface and a matte surface, while higher barb eccentricity values increase the intensity at some specific regions and darker it at others.

In Figure D.6 we vary both the barb roughness and the barbule roughness and evaluate how affect the final appearance of a single feather. The larger the roughness, the more matte and softer the final appearance becomes. At low roughness, the specular highlights from the barb and barbules are clearly noticeable and they can be a helpful setup to debug the contribution of each feather component in more detail.

In Figure D.7 we perform an appearance parameter exploration for the two main pigments of our barb BCSDF. Please notice how the appearance of the single feather

Figure D.3: Appearance exploration of a feather wing for two free model variables: the barb eccentricity parameter (rows) $b_c = [1.0, 1.3, 1.6, 1.8]$ and the barbule longitudinal angle parameter (columns) $\theta_{bb} = [0°, 10°, 15°, 30°, 45°]$.

change from white and yellowish appearances for low eumelanin and pheomelanin values to dark brown and black appearances for higher values.

Figure D.4: Appearance exploration of a feather wing for two free model variables under front point light source: the barbule length parameter (rows) $l_{bb} = [0.01, 0.1, 1.0, 2.0, 10.0]$ and the barbule separation parameter (columns) $\delta_{bb} = [0.01, 0.1, 0.5, 1.0, 2.0]$.

Figure D.5: Appearance exploration of a feather wing for two free model variables: the barb eccentricity (rows) $b_c = [1.0, 1.3, 1.6, 1.8]$ and the barb roughness parameter (rows) $\beta_b^m = [0.01, 0.05, 0.1, 0.5, 1.0]$.

Figure D.6: Appearance exploration of a feather wing for two free model variables: the barb roughness parameter (columns) $\beta_{b}^{m} = [0.01, 0.1, 0.5, 1.0, 2.0]$ and the barbule roughness parameter (rows) $\beta_{bb}^{m} = [0.01, 0.1, 0.5, 1.0, 2.0]$.

Figure D.7: Appearance exploration of a feather wing for two free model variables: the barb eumelanin parameter (rows) $\rho_{b,e} = [0.1, 0.5, 0.75, 1.0, 2.0, 3.0]$ and the barb pheomelanin (columns) $\rho_{b,p} = [0.1, 0.5, 0.75, 1.0, 2.0, 3.0]$.

BIBLIOGRAPHY

[1] Adobe. *Photoshop*. `https://www.adobe.com/products/photoshop.html`. 2023.

[2] Carlos Aliaga, Carlos Castillo, Diego Gutierrez, Miguel A. Otaduy, Jorge Lopez-Moreno, and Adrian Jarabo. "An Appearance Model for Textile Fibers." In: *Computer Graphics Forum (Proc. EGSR 2017)* 36.4 (2017).

[3] Carlos Aliaga, Menqi Xia, Xao Xie, Adrian Jarabo, Gustav Braun, and Christophe Hery. "A Hyperspectral Space of Skin Tones for Inverse Rendering of Biophysical Skin Properties." In: *Computer Graphics Forum* 42.4 (2023), e14887.

[4] Pontus Andersson, Jim Nilsson, Tomas Akenine-Möller, Magnus Oskarsson, Kalle Åström, and Mark D. Fairchild. "FLIP: A Difference Evaluator for Alternating Images." In: *Proc. of the ACM on Computer Graphics and Interactive Techniques* 3.2 (2020), 15:1–15:23.

[5] Autodesk. *Maya*. `https://www.autodesk.com/`. 2023.

[6] Gladimir V. G. Baranoski and Jon G. Rokne. "Efficiently simulating scattering of light by leaves." In: *The Visual Computer* 17.8 (2001), pp. 491–505.

[7] Jessica Baron, Daljit Singh Dhillon, N. Adam Smith, and Eric Patterson. "Microstructure-based appearance rendering for feathers." In: *Computers & Graphics* 102 (2022), pp. 452–459.

[8] Jessica Baron and Eric Patterson. "Procedurally Generating Biologically Driven Feathers." In: *Advances in Computer Graphics* 11542 (2019), pp. 342–348.

[9] Jessica Baron and Eric Patterson. "Computer-Aided Visual Analysis of Feathers." In: *ICPR Workshop on Visual observation and analysis of Vertebrate And Insect Behavior (VAIB)*. 2021.

[10] Laurent Belcour. "Efficient rendering of layered materials using an atomic decomposition with statistical operators." In: *ACM Trans. on Graphics* 37.4 (2018), 73:1–73:15.

[11] Alexis Benamira and Sumanta Pattanaik. "A combined scattering and diffraction model for elliptical hair rendering." In: *Computer Graphics Forum* 40.4 (2021), pp. 163–175.

[12] Benedikt Bitterli, Srinath Ravichandran, Thomas Müller, Magnus Wrenninge, Jan Novák, Steve Marschner, and Wojciech Jarosz. "A radiative transfer framework for non-exponential media." In: *Trans. on Graphics* (2018).

[13] Benedikt Bitterli and Eugene d'Eon. "A position-free path integral for homogeneous slabs and multiple scattering on Smith microfacets." In: *Computer Graphics Forum* 41.4 (2022), pp. 93–104.

[14] Blender Foundation. *Blender*. `https://www.blender.org/`. 2023.

[15] Fred L Bookstein. *Morphometric tools for landmark data*. Cambridge University Press, 1997.

[16] Carles Bosch, Xavier Pueyo, Stéphane Mérillou, and Djamchid Ghazanfarpour. "A Physically-Based Model for Rendering Realistic Scratches." In: *Computer Graphics Forum* 23.3 (2004), pp. 361–370.

[17] Adrien Bousseau, Emmanuelle Chapoulie, Ravi Ramamoorthi, and Maneesh Agrawala. "Optimizing environment maps for material depiction." In: *Computer graphics forum*. Vol. 30. 4. Wiley Online Library. 2011, pp. 1171–1180.

[18] Brent Burley. *Practical Physically-Based Shading in Film and Game Production: Physically Based Shading at Disney*. SIGGRAPH Courses (2012). 2012.

[19] Brent Burley. *Physically Based Shading in Theory and Practice : Extending the Disney BRDF to a BSDF with integrated subsurface scattering*. SIGGRAPH Courses (2015). 2015.

[20] Muhammad Atif Butt, Kai Wang, Javier Vazquez-Corral, and Joost van de Weijer. "Colorpeel: Color prompt learning with diffusion models via color and shape disentanglement." In: *Proc. European Conference on Computer Vision (ECCV)*. Springer. 2025, pp. 456–472.

[21] J. Buzek and B. Ask. *Regulation (EC) No 1223/2009 of the European parliament and of the council of 30 November 2009 on cosmetic products*. `https://www.eumonitor.eu/9353000/1/j9vvik7m1c3gyxp/vibn2mp7slr0`. 2009.

[22] Alice C Chadwick and Robert William Kentridge. "The perception of gloss: A review." In: *Vision research* 109 (2015), pp. 221–235.

[23] S Chandrasekhar. *Radiative Transfer*. Dover Publications, 1960.

[24] Huiwen Chang, Jingwan Lu, Fisher Yu, and Adam Finkelstein. "PairedCycleGAN: asymmetric style transfer for applying and removing makeup." In: *Proc. Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2018, pp. 40–48.

[25] Chengqian Che, Fujun Luan, Shuang Zhao, Kavita Bala, and Ioannis Gkioulekas. "Towards Learning-based Inverse Subsurface Scattering." In: *Proc. International Conference on Computational Photography (ICCP)* (2020), pp. 1–12.

[26] Bin Chen, Chao Wang, Michal Piovarči, Hans-Peter Seidel, Piotr Didyk, Karol Myszkowski, and Ana Serrano. "The effect of geometry and illumination on appearance perception of different material categories." In: *The Visual Computer* 37.12 (2021), pp. 2975–2987.

[27] Yanyun Chen, Yingqing Xu, Baining Guo, and Heung-Yeung Shum. "Modeling and rendering of realistic feathers." In: *ACM Trans. on Graphics* 21.3 (2002), pp. 630–636.

[28] Matt Jen-Yuan Chiang, Benedikt Bitterli, Chuck Tappan, and Brent Burley. "A Practical and Controllable Hair and Fur Model for Production Path Tracing." In: 35.2 (2016).

[29] Nahian S Chowdhury, Phillip J Marlow, and Juno Kim. "Translucency and the perception of shape." In: *Journal of vision* 17.3 (2017), pp. 17–17.

[30] Mark Colbert, Sumanta Pattanaik, and Jaroslav Krivanek. "BRDF-Shop: Creating Physically Correct Bidirectional Reflectance Distribution Functions." In: *IEEE Comput. Graph. Appl.* 26.1 (2006), 30–36.

[31] Robert L. Cook and Kenneth E. Torrance. "A reflectance model for computer graphics." In: *SIGGRAPH Comput. Graph.* 15.3 (1981), 307–316.

[32] Johanna Delanoy, Manuel Lagunas, J Condor, Diego Gutierrez, and Belén Masia. "A Generative Framework for Image-based Editing of Material Appearance using Perceptual Attributes." In: *Computer Graphics Forum*. Vol. 41. 1. Wiley Online Library. 2022, pp. 453–464.

[33] Han Deng, Chu Han, Hongmin Cai, Guoqiang Han, and Shengfeng He. "Spatially-invariant style-codes controlled makeup transfer." In: *Proc. Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2021, pp. 6549–6557.

[34] Xi Deng, Fujun Luan, Bruce Walter, Kavita Bala, and Steve Marschner. "Reconstructing Translucent Objects using Differentiable Rendering." In: *Proc. SIGGRAPH* (2022).

[35] Katja Doerschner, Roland W Fleming, Ozgur Yilmaz, Paul R Schrater, Bruce Hartung, and Daniel Kersten. "Visual motion and the perception of surface material." In: *Current Biology* 21.23 (2011), pp. 2010–2016.

[36] Craig Donner and Henrik Wann Jensen. "Light diffusion in multi-layered translucent materials." In: *ACM Trans. on Graphics* 24.3 (2005), pp. 1032–1039.

[37] Craig Donner, Tim Weyrich, Eugene d'Eon, Ravi Ramamoorthi, and Szymon Rusinkiewicz. "A layered, heterogeneous reflectance model for acquiring and rendering human skin." In: *ACM Trans. on Graphics* 27.5 (2008), 140:1–140:12.

[38] Marc Droske, Johannes Hanika, Jiri Vorba, Andrea Weidlich, and Manuele Sabbadin. "Path tracing in production: the path of water." In: *SIGGRAPH Courses (2023)*. ACM, 2023.

[39] Jonathan Dupuy, Eric Heitz, and Eugene d'Eon. "Additional progress towards the unification of microfacet and microflake theories." In: *Eurographics Symposium on Rendering - Experimental Ideas & Implementations (EGSR EI&I)*. Eurographics Association, 2016, pp. 55–63.

[40] A Ennos, J Hickson, and ANNA Roberts. "Functional morphology of the vanes of the flight feathers of the pigeon Columba livia." In: *The Journal of experimental biology* 198.5 (1995), pp. 1219–1228.

[41] William Epstein, Gary Hatfield, and Gerard Muise. "Perceived shape at a slant as a function of processing time and processing load." In: *Journal of Experimental Psychology: Human Perception and Performance* 3.3 (1977), p. 473.

[42] Becca E. Ewing, Samuel D. Butler, and Michael A. Marciniak. "Improved grazing angle bidirectional reflectance distribution function model using Rayleigh–Rice polarization factor and adaptive microfacet distribution function." In: *Optical Engineering* 57.10 (2018), p. 105102.

[43] Viggo Falster, Adrian Jarabo, and Jeppe Revall Frisvad. "Computing the bidirectional scattering of a microstructure using scalar diffraction theory and path tracing." In: *Computer Graphics Forum*. Vol. 39. 7. Wiley Online Library. 2020, pp. 231–242.

[44] Li Feng, Fen Wang, Hongjie Luo, and Boxin Qiu. "Review of recent advancements in the biomimicry of structural colors." In: *Dyes and Pigments* 210 (2023), p. 111019.

[45] Jirí Filip, Mike J Chantler, Patrick R Green, and Michal Haindl. "A psychophysically validated metric for bidirectional texture data reduction." In: *Trans. on Graphics* 27.5 (2008), p. 138.

[46] Roland W Fleming and Heinrich H Bülthoff. "Low-level image cues in the perception of translucent materials." In: *ACM Trans. Appl. Percept.* 2.3 (2005), pp. 346–382.

[47] Roland W. Fleming, Ron O. Dror, and Edward H. Adelson. "How do Humans Determine Reflectance Properties under Unknown Illumination." In: *CVPR 2001*. 2001.

[48] Roland W Fleming, Ron O Dror, and Edward H Adelson. "Real-world illumination and the perception of surface reflectance properties." In: *Journal of vision* 3.5 (2003), pp. 3–3.

[49] Roland W Fleming, Henrik Wann Jensen, and Heinrich H Bülthoff. "Perceiving translucent materials." In: *Proc. of the ACM Symposium on Applied Perception*. 2004, pp. 127–134.

[50] G. Fourneau, R. Pacanowski, and P. Barla. "Interactive Exploration of Vivid Material Iridescence using Bragg Mirrors." In: *Computer Graphics Forum* (2024).

[51] Pascal Freyer and Doekele G. Stavenga. "Biophotonics of diversely coloured peacock tail feathers." In: *Faraday Discuss.* 223 (0 2020), pp. 49–62.

[52] Pascal Freyer, Bodo D Wilts, and Doekele G Stavenga. "Reflections on iridescent neck and breast feathers of the peacock, Pavo cristatus." In: *Journal of the Royal Society Interface Focus* 9.1 (2019), p. 20180043.

[53] Jeppe Revall Frisvad, Niels Jørgen Christensen, and Henrik Wann Jensen. "Computing the Scattering Properties of Participating Media Using Lorenz-Mie Theory." In: *ACM Trans. on Graphics* 26.3 (2007), 60:1–60:10.

[54] Luis E. Gamboa, Adrien Gruson, and Derek Nowrouzezahrai. "An efficient transport estimator for complex layered materials." In: *Computer Graphics Forum* 39.2 (2020), pp. 363–371.

[55] Elena Garces, Aseem Agarwala, Diego Gutierrez, and Aaron Hertzmann. "A Similarity Measure for Illustration Style." In: *Trans. on Graphics* 33.4 (2014).

[56] Davit Gigilashvili, Lucas Dubouchet, Jon Yngve Hardeberg, and Marius Pedersen. "Caustics and translucency perception." In: *Electronic Imaging* (2020), pp. 33–1.

[57] Davit Gigilashvili, David Norman Díaz Estrada, and Lakshay Jain. "Does Motion Increase Perceived Magnitude of Translucency?" In: *London Imaging Meeting*. Vol. 4. Society for Imaging Science and Technology. 2023, pp. 81–85.

[58] Davit Gigilashvili, Weiqi Shi, Zeyu Wang, Marius Pedersen, Jon Yngve Hardeberg, and Holly Rushmeier. "The Role of Subsurface Scattering in Glossiness Perception." In: *ACM Trans. Applied Perception* 18.3 (2021), pp. 1–26.

[59] Davit Gigilashvili, Jean-Baptiste Thomas, Jon Yngve Hardeberg, and Marius Pedersen. “Behavioral investigation of visual appearance assessment.” In: *Color and imaging conference*. Vol. 2018. 1. 2018, pp. 294–299.

[60] Davit Gigilashvili, Jean-Baptiste Thomas, Jon Yngve Hardeberg, and Marius Pedersen. “Translucency perception: A review.” In: *Journal of Vision* 21.8 (2021), pp. 4–4.

[61] Davit Gigilashvili, Jean-Baptiste Thomas, Marius Pedersen, and Jon Yngve Hardeberg. “On the appearance of objects and materials: Qualitative analysis of experimental observations.” In: *Journal of the International Colour Association* 27 (2021), pp. 26–55.

[62] Davit Gigilashvili, Philipp Urban, Jean-Baptiste Thomas, Jon Yngve Hardeberg, and Marius Pedersen. “Impact of shape on apparent translucency differences.” In: *Color and Imaging Conference*. Vol. 2019. 1. 2019, pp. 132–137.

[63] Davit Gigilashvili, Philipp Urban, Jean-Baptiste Thomas, Marius Pedersen, and Jon Yngve Hardeberg. “The Impact of Optical and Geometrical Thickness on Perceived Translucency Differences.” In: *Journal of Perceptual Imaging* 5 (2022).

[64] Ioannis Gkioulekas, Anat Levin, and Todd Zickler. “An evaluation of computational imaging techniques for heterogeneous inverse scattering.” In: *Proc. ECCV*. Springer. 2016, pp. 685–701.

[65] Ioannis Gkioulekas, Bruce Walter, Edward H Adelson, Kavita Bala, and Todd Zickler. “On the appearance of translucent edges.” In: *Proc. CVPR*. 2015, pp. 5528–5536.

[66] Ioannis Gkioulekas, Bruce Walter, Edward H Adelson, Kavita Bala, and Todd Zickler. “On the appearance of translucent edges.” In: *Proc. Computer Vision and Pattern Recognition (CVPR)*. 2015, pp. 5528–5536.

[67] Ioannis Gkioulekas, Bei Xiao, Shuang Zhao, Edward H. Adelson, Todd Zickler, and Kavita Bala. “Understanding the Role of Phase Function in Translucent Appearance.” In: *Trans. on Graphics* 32.5 (2013).

[68] Ioannis Gkioulekas, Shuang Zhao, Kavita Bala, Todd Zickler, and Anat Levin. “Inverse volume rendering with material dictionaries.” In: *Trans. on Graphics* 32.6 (2013), pp. 1–13.

[69] Andrew S Glassner. “A model for fluorescence and phosphorescence.” In: *Photorealistic Rendering Techniques*. Springer, 1995, pp. 60–70.

[70] Jay S. Gondek, Gary W. Meyer, and Jonathan G. Newman. “Wavelength dependent reflectance functions.” In: *Proc. Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH ’94. ACM, 1994, 213–220.

[71] Eugenia Guerra, Maria Celeiro, J. Pablo Lamas, Maria Llompart, and Carmen Garcia-Jares. “Determination of dyes in cosmetic products by micro-matrix solid phase dispersion and liquid chromatography coupled to tandem mass spectrometry.” In: *Journal of Chromatography A* 1415 (2015), pp. 27–37.

[72] Eugenia Guerra, Maria Llompart, and Carmen Garcia-Jares. "Miniaturized matrix solid-phase dispersion followed by liquid chromatography-tandem mass spectrometry for the quantification of synthetic dyes in cosmetics and foodstuffs used or consumed by children." In: *Journal of Chromatography A* 1529 (2017), pp. 29–38.

[73] Julia Guerrero-Viu, Milos Hasan, Arthur Roullier, Midhun Harikumar, Yiwei Hu, Paul Guerrero, Diego Gutierrez, Belen Masia, and Valentin Deschaintre. "Texsliders: Diffusion-based texture editing in clip space." In: *ACM SIGGRAPH 2024 Conference Papers*. 2024, pp. 1–11.

[74] Ibón Guillén, Julio Marco, Diego Gutierrez, Wenzel Jakob, and Adrian Jarabo. "A general framework for pearlescent materials." In: *ACM Trans. on Graphics* 39.6 (2020), 253:1–253:15.

[75] Jie Guo, Zeru Li, Xueyan He, Beibei Wang, Wenbin Li, Yanwen Guo, and Ling-Qi Yan. "MetaLayer: a meta-learned BSDF model for layered materials." In: *ACM Trans. on Graphics* 42.6 (2023), 222:1–222:15.

[76] Yu Guo, Miloš Hašan, and Shuang Zhao. "Position-free Monte Carlo simulation for arbitrary layered BSDFs." In: *ACM Trans. on Graphics* 37 (6 2018), pp. 1–14.

[77] Diego Gutierrez, Francisco J Seron, Adolfo Munoz, and Oscar Anson. "Visualizing underwater ocean optics." In: *Computer Graphics Forum*. Vol. 27. 2. Wiley Online Library. 2008, pp. 547–556.

[78] Rasmus Haapaoja and Christoph Genzwürker. "Mesh-driven generation and animation of groomed feathers." In: *ACM SIGGRAPH 2019 Talks*. SIGGRAPH '19. Los Angeles, California: ACM, 2019.

[79] Pat Hanrahan and Wolfgang Krueger. "Reflection from layered surfaces due to subsurface scattering." In: *SIGGRAPH '93*. ACM, 1993, pp. 165–174.

[80] Bruce Hapke. "Bidirectional reflectance spectroscopy: 1. Theory." In: *Journal of Geophysical Research: Solid Earth* 86.B4 (1981), pp. 3039–3054.

[81] Bruce Hapke. *Theory of Reflectance and Emittance Spectroscopy*. 2nd ed. Cambridge University Press, 2012.

[82] Todd Alan Harvey, Kimberly S. Bostwick, and Steve Marschner. "Directional reflectance and milli-scale feather morphology of the African Emerald Cuckoo, Chrysococcyx cupreus." In: *Journal of the Royal Society Interface* (2013).

[83] Miloš Hašan and Ravi Ramamoorthi. "Interactive albedo editing in path-traced volumetric materials." In: *Trans. on Graphics* 32.2 (2013), pp. 1–11.

[84] Vlastimil Havran, Jiri Filip, and Karol Myszkowski. "Perceptually motivated BRDF comparison using single image." In: *Computer graphics forum*. Vol. 35. 4. Wiley Online Library. 2016, pp. 1–12.

[85] Jeffrey Heer and Michael Bostock. "Crowdsourcing Graphical Perception: Using Mechanical Turk to Assess Visualization Design." In: *Proc. of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '10. Atlanta, Georgia, USA: ACM, 2010, 203–212.

[86] Eric Heitz. “Understanding the Masking-Shadowing Function in Microfacet-Based BRDFs.” In: *Journal of Computer Graphics Techniques (JCGT)* 3.2 (2014), pp. 48–107.

[87] Eric Heitz, Jonathan Dupuy, Cyril Crassin, and Carsten Dachsbacher. “The SGGX microflake distribution.” In: *ACM Trans. on Graphics* 34 (4 2015), 48:1–48:11.

[88] Eric Heitz, Johannes Hanika, Eugene d’Eon, and Carsten Dachsbacher. “Multiple-scattering microfacet BSDFs with the Smith model.” In: *Trans. on Graphics* 35 (2016), pp. 1 –14.

[89] Louis G. Henyey and Jesse Leonard Greenstein. “Diffuse radiation in the galaxy.” In: *The Astrophysical Journal* 93 (1941), pp. 70–83.

[90] Geoffrey E Hill and Kevin J McGraw. *Bird coloration: function and evolution*. Vol. 2. Harvard University Press, 2006.

[91] Geoffrey E Hill and Kevin J McGraw. *Bird coloration: mechanisms and measurements*. Vol. 1. Harvard University Press, 2006.

[92] Nicolas Holzschuch and Romain Pacanowski. “A Two-Scale Microfacet Reflectance Model Combining Reflection and Diffraction.” In: *Trans. on Graphics* 36.4 (2017).

[93] Cheng-Guo Huang, Tsung-Shian Huang, Wen-Chieh Lin, and Jung-Hong Chuang. “Physically based cosmetic rendering.” In: *Computer Animation and Virtual Worlds* 24.3-4 (2013), pp. 275–283.

[94] W. Huang, M. B. Hullin, and J. Hanika. “A Microfacet-based Hair Scattering Model.” In: *Computer Graphics Forum* 41.4 (2022), pp. 79–91.

[95] Weizhen Huang, Sebastian Merzbach, Clara Callenberg, Doekele Stavenga, and Matthias Hullin. “Rendering Iridescent Rock Dove Neck Feathers.” In: *Proc. ACM SIGGRAPH 2022*. SIGGRAPH ’22. Vancouver, BC, Canada: ACM, 2022.

[96] Weizhen Huang, Sebastian Merzbach, Clara Callenberg, Doekele Stavenga, and Matthias Hullin. “Rendering iridescent rock dove neck feathers.” In: *Proc. ACM SIGGRAPH 2022*. 2022, pp. 1–8.

[97] Donald R. Huffman. “THE APPLICABILITY OF BULK OPTICAL CONSTANTS TO SMALL PARTICLES.” In: *Optical Effects Associated with Small Particles*. 1988, pp. 277–324.

[98] ICT Vision & Graphics Lab. *High-Resolution Light Probe Image Gallery*. `https://vgl.ict.usc.edu/Data/HighResProbes/`. 2010.

[99] Branislav Igic, Liliana D’Alba, and Matthew D Shawkey. “Fifty shades of white: how white feather brightness differs among species.” In: *The Science of Nature* 105.3 (2018), pp. 1–15.

[100] Jose A. Iglesias-Guitian, Carlos Aliaga, Adrian Jarabo, and Diego Gutierrez. “A Biophysically-Based Model of the Optical Properties of Skin Aging.” In: *Computer Graphics Forum* 34.2 (2015), pp. 45–55.

[101] Piti Irawan and Steve Marschner. “Specular reflection from woven cloth.” In: *Trans. on Graphics* 31.1 (2012).

[102] Wenzel Jakob. *Mitsuba renderer*. `http://mitsuba-renderer.org/`. 2010.

[103] Wenzel Jakob, Adam Arbree, Jonathan T Moon, Kavita Bala, and Steve Marschner. "A radiative transfer framework for rendering materials with anisotropic structure." In: *Trans. on Graphics* 29.4 (2010), pp. 1–13.

[104] Wenzel Jakob, Adam Arbree, Jonathan T Moon, Kavita Bala, and Steve Marschner. "A radiative transfer framework for rendering materials with anisotropic structure." In: *ACM Trans. on Graphics* 29.4 (2010), 53:1–53:13.

[105] Wenzel Jakob, Miloš Hašan, Ling-Qi Yan, Ravi Ramamoorthi, and Steve Marschner. "Discrete Stochastic Microfacet Models." In: *ACM Trans. on Graphics (Proceedings of SIGGRAPH)* 33.4 (July 2014), 115:1–115:10.

[106] Wenzel Jakob, Sébastien Speierer, Nicolas Roussel, Merlin Nimier-David, Delio Vicini, Tizian Zeltner, Baptiste Nicolet, Miguel Crespo, Vincent Leroy, and Ziyi Zhang. *Mitsuba 3 renderer*. Version 3.0.1. https://mitsuba-renderer.org. 2022.

[107] Wenzel Jakob, Eugene d'Eon, Otto Jakob, and Steve Marschner. "A comprehensive framework for rendering layered materials." In: *ACM Trans. on Graphics* 33.4 (2014), pp. 1–14.

[108] Adrian Jarabo, Carlos Aliaga, and Diego Gutierrez. "A radiative transfer framework for spatially-correlated materials." In: *Trans. on Graphics* 37.4 (2018), pp. 1–13.

[109] Adrian Jarabo and Victor Arellano. "Bidirectional rendering of vector light transport." In: *Computer Graphics Forum*. Vol. 37. 6. Wiley Online Library. 2018, pp. 96–105.

[110] Adrian Jarabo, Hongzhi Wu, Julie Dorsey, Holly Rushmeier, and Diego Gutierrez. "Effects of approximate filtering on the appearance of bidirectional texture functions." In: *IEEE Trans. Visualization and Computer Graphics* 20.6 (2014), pp. 880–892.

[111] Henrik Wann Jensen, Stephen R Marschner, Marc Levoy, and Pat Hanrahan. "A practical model for subsurface light transport." In: *Proc. CGIT*. 2001, pp. 511–518.

[112] Myung-Ok Jeon and Byung-Soo Chang. "Property of silica and fine structure of cosmetic white powders." In: *Applied Microscopy* 42 (2012), pp. 87–93.

[113] Wentao Jiang, Si Liu, Chen Gao, Jie Cao, Ran He, Jiashi Feng, and Shuicheng Yan. "PSGAN: pose and expression robust spatial-aware GAN for customizable makeup transfer." In: *Proc. Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2020, pp. 5194–5202.

[114] Jorge Jimenez, Károly Zsolnai, Adrian Jarabo, Christian Freude, Thomas Auzinger, Xian-Chun Wu, Javier von der Pahlen, Michael Wimmer, and Diego Gutierrez. "Separable subsurface scattering." In: *Computer Graphics Forum* 34.6 (2015), pp. 188–197.

[115] James T. Kajiya. "The Rendering Equation." In: *SIGGRAPH Comput. Graph.* 20.4 (1986), 143–150.

[116] David G Kendall. "A survey of the statistical theory of shape." In: *Statistical Science* 4.2 (1989), pp. 87–99.

[117] Kepios Pte. Ltd. *DataReportal*. `https://datareportal.com/reports/digital-2024-global-overview-report/`. 2024.

[118] William B. Kerr and Fabio Pellacini. “Toward evaluating material design interface paradigms for novice users.” In: 2010.

[119] Apoorv Khattar, Junqui Zhu, Emiliano Padovani, Jean-Marie Aurby, Marc Droske, Ling-Qi Yan, and Zahra Montazeri. “A Multi-scale Yarn Appearance Model with Fiber Details.” In: *Computational Visual Media*. 2024.

[120] Pramook Khungurn and Steve Marschner. “Azimuthal Scattering from Elliptical Hair Fibers.” In: *Trans. on Graphics* 36.2 (2017).

[121] Diederik P Kingma and Jimmy Ba. “Adam: A Method for Stochastic Optimization.” In: *arXiv e-prints* (2014), arXiv–1412.

[122] Hiroaki Kiyokawa, Takehiro Nagai, Yasuki Yamauchi, and Juno Kim. “The perception of translucency from surface gloss.” In: *Vision Research* (2022), p. 108140.

[123] Manuel Lagunas, Sandra Malpica, Ana Serrano, Elena Garces, Diego Gutierrez, and Belen Masia. “A Similarity Measure for Material Appearance.” In: *Trans. on Graphics* 38.4 (2019).

[124] Manuel Lagunas, Ana Serrano, Diego Gutierrez, and Belen Masia. “The joint role of geometry and illumination on material recognition.” In: *Journal of Vision* 21.2 (2021), pp. 2–2.

[125] Dario Lanza, Adrian Jarabo, and Belen Masia. “On the Influence of Dynamic Illumination in the Perception of Translucency.” In: *Proc. of the ACM Symposium on Applied Perception*. 2022, pp. 1–9.

[126] Dario Lanza, Belen Masia, and Adrian Jarabo. “Navigating the Manifold of Translucent Appearance.” In: *Computer Graphics Forum* (2024).

[127] Dario Lanza, Juan Raúl Padrón-Griffe, Alina Pranovich, Adolfo Muñoz, Jeppe Revall Frisvad, and Adrián Jarabo. “Practical Appearance Model for Foundation Cosmetics.” In: *Computer Graphics Forum* 43.4 (2024), e15148.

[128] Guillaume Lavoué, Nicolas Bonneel, Jean-Philippe Farrugia, and Cyril Soler. “Perceptual quality of BRDF approximations: dataset and metrics.” In: *Computer Graphics Forum* 40 (2021).

[129] James Leaning and Damien Fagnou. “Feathers for mystical creatures: creating pegasus for Clash of the Titans.” In: *ACM SIGGRAPH 2010 Talks*. SIGGRAPH ’10. Los Angeles, California: ACM, 2010.

[130] Marc Levoy, J Gerth, Brian Curless, and Kari Pulli. *The Stanford 3D scanning repository*. `http://www-graphics.stanford.edu/data/3dscanrep`. Accessed: 2022-06-09. 2005.

[131] Chen Li, Kun Zhou, and Stephen Lin. “Simulating makeup through physics-based manipulation of intrinsic image layers.” In: *Proc. Computer Vision and Pattern Recognition (CVPR)*. 2015, pp. 4621–4629.

[132] Chen Li, Kun Zhou, Hsiang-Tao Wu, and Stephen Lin. “Physically-based simulation of cosmetics via intrinsic image decomposition with facial priors.” In: *IEEE Trans. on Pattern Analysis and Machine Intelligence* 41.6 (2019), pp. 1455–1469.

[133] Chenhao Li, Trung Thanh Ngo, and Hajime Nagahara. "Inverse Rendering of Translucent Objects using Physical and Neural Renderers." In: *Proc. CVPR*. 2023, pp. 12510–12520.

[134] Tingting Li, Ruihe Qian, Chao Dong, Si Liu, Qiong Yan, Wenwu Zhu, and Liang Lin. "BeautyGAN: instance-level facial makeup transfer with deep generative adversarial network." In: *Multimedia (MM '18)*. ACM, 2018, 645—653.

[135] Chenxi Liao, Masataka Sawayama, and Bei Xiao. "Crystal or jelly? Effect of color on the perception of translucent materials with photographs of real-world objects." In: *Journal of Vision* 22.2 (2022), pp. 6–6.

[136] Chenxi Liao, Masataka Sawayama, and Bei Xiao. "Unsupervised learning reveals interpretable latent representations for translucency perception." In: *PLOS Computational Biology* 19 (2022).

[137] Albert Julius Liu, Zhao Dong, Miloš Hašan, and Steve Marschner. "Simulating the structure and texture of solid wood." In: *ACM Trans. on Graphics* 35.6 (2016), 170:1–170:11.

[138] Si Liu, Wentao Jiang, Chen Gao, Ran He, Jiashi Feng, Bo Li, and Shuicheng Yan. "PSGAN++: robust detail-preserving makeup transfer and removal." In: *IEEE Trans. on Pattern Analysis and Machine Intelligence* 44.11 (2022), pp. 8538–8551.

[139] Si Liu, Xinyu Ou, Ruihe Qian, Wei Wang, and Xiaochun Cao. "Makeup like a superstar: deep localized makeup transfer network." In: *International Joint Conference on Artificial Intelligence (IJCAI)*. AAAI Press, 2016, pp. 2568–2575.

[140] Joakim Löw, Joel Kronander, Anders Ynnerman, and Jonas Unger. "BRDF Models for Accurate and Efficient Rendering of Glossy Surfaces." In: *Trans. on Graphics* 31.1 (2012).

[141] Simon Lucas, Mickael Ribardiere, Romain Pacanowski, and Pascal Barla. "A Micrograin BSDF Model for the Rendering of Porous Layers." In: *Proc. SIGGRAPH Asia 2023*. SA '23. Sydney, Australia: ACM, 2023.

[142] Rafael Maia, Liliana D'Alba, and Matthew D Shawkey. "What makes a feather shine? A nanostructural basis for glossy black colours in feathers." In: *Proc. of the Royal Society B: Biological Sciences* 278.1714 (2011), pp. 1973–1980.

[143] Frank J. Maile, Gerhard Pfaff, and Peter Reynders. "Effect pigments—past, present and future." In: *Progress in Organic Coatings* 54.3 (2005), pp. 150–163.

[144] Nathan Mantel. "The detection of disease clustering and a generalized regression approach." In: *Cancer research* 27 2 (1967), pp. 209–20.

[145] Rafał K. Mantiuk, Alexandre Chapiro, Gizem Rufo, Trisha Lian, Rafał K. Mantiuk, Gyorgy Denes, Alexandre Chapiro, and Anton Kaplanyan. "FovVideoVDP: a visible difference predictor for wide field-of-view video." In: *Trans. on Graphics* 40 (2021), 49:1–49:19.

[146] Ruiquan Mao, Manuel Lagunas, Belen Masia, and Diego Gutierrez. "The Effect of Motion on the Perception of Material Appearance." In: 2019, 16:1–16:9.

[147] Phillip J. Marlow and Barton L. Anderson. "The cospecification of the shape and material properties of light permeable materials." In: *Proc. of the National Academy of Sciences* 118 (2021).

[148] Phillip J Marlow, Juno Kim, and Barton L Anderson. “Perception and misperception of surface opacity.” In: *Proc. of the National Academy of Sciences* 114.52 (2017), pp. 13840–13845.

[149] Stephen R. Marschner, Henrik Wann Jensen, Mike Cammarano, Steve Worley, and Pat Hanrahan. “Light scattering from human hair fibers.” In: *Trans. on Graphics* 22.3 (2003), 780–791.

[150] Stephen R. Marschner, Stephen H. Westin, Adam Arbree, and Jonathan T. Moon. “Measuring and modeling the appearance of finished wood.” In: *ACM Trans. on Graphics* 24.3 (2005), pp. 727–734.

[151] Wojciech Matusik. “A data-driven reflectance model.” PhD thesis. Massachusetts Institute of Technology, 2003.

[152] Dakota McCoy, Allison Shultz, Charles Vidoudez, Emma Heide, Jacqueline Dall, Sunia Trauger, and David Haig. “Microstructures amplify carotenoid plumage signals in tanagers.” In: *Scientific Reports* 11 (2021), p. 8582.

[153] Maria M Mendes-Pinto, Amy M LaFountain, Mary Caswell Stoddard, Richard O Prum, Harry A Frank, and Bruno Robert. “Variation in carotenoid–protein interaction in bird feathers produces novel plumage coloration.” In: *Journal of the Royal Society Interface* 9.77 (2012), pp. 3338–3350.

[154] Johannes Meng, Marios Papas, Ralf Habel, Carsten Dachsbacher, Steve Marschner, Markus Gross, and Wojciech Jarosz. “Multi-Scale Modeling and Rendering of Granular Materials.” In: *ACM Trans. on Graphics* 34.4 (2015), 49:1–49:13.

[155] Stéphane Mérillou, Jean-Michel Dischler, and Djamchid Ghazanfarpour. “A BRDF Postprocess to Integrate Porosity on Rendered Surfaces.” In: *IEEE Trans. on Visualization and Computer Graphics* 6.4 (2000), 306–318.

[156] Fabio Metelli. “The perception of transparency.” In: *Scientific American* 230.4 (1974), pp. 90–99.

[157] Microsoft. *Microsoft PowerPoint*. https://www.microsoft.com/es-es/microsoft-365/powerpoint. 2023.

[158] Aliaksei Mikhailiuk, María Pérez-Ortiz, Dingcheng Yue, Wilson Suen, and Rafał K. Mantiuk. “Consolidated Dataset and Metrics for High-Dynamic-Range Image Quality.” In: *IEEE Trans. on Multimedia* 24 (2020), pp. 2125–2138.

[159] Michael I. Mishchenko. “Vector radiative transfer equation for arbitrarily shaped and arbitrarily oriented particles: a microphysical derivation from statistical electromagnetics.” In: *Appl. Opt.* 41.33 (2002), pp. 7114–7134.

[160] Yusuke Moriuchi, Shoji Tominaga, and Takahiko Horiuchi. “Precise Analysis of Spectral Reflectance Properties of Cosmetic Foundation.” In: vol. 5575 LNCS. 2009, pp. 138–148.

[161] Isamu Motoyoshi. “Highlight–shading relationship as a cue for the perception of translucent and transparent materials.” In: *Journal of Vision* 10.9 (2010), pp. 6–6.

[162] Isamu Motoyoshi. “Highlight–shading relationship as a cue for the perception of translucent and transparent materials.” In: *Journal of Vision* 10.9 (2010), pp. 6–6.

[163] Thomas Müller, Marios Papas, Markus Gross, Wojciech Jarosz, and Jan Novák. "Efficient rendering of heterogeneous polydisperse granular media." In: *ACM Trans. on Graphics* 35.6 (2016), 168:1–168:14.

[164] Srinivasa G Narasimhan, Mohit Gupta, Craig Donner, Ravi Ramamoorthi, Shree K Nayar, and Henrik Wann Jensen. "Acquiring scattering properties of participating media by dilution." In: *Proc. SIGGRAPH*. 2006, pp. 1003–1012.

[165] Addy Ngan, Frédo Durand, and Wojciech Matusik. "Image-driven Navigation of Analytical BRDF Models." In: *Symposium on Rendering*. Ed. by Tomas Akenine-Moeller and Wolfgang Heidrich. 2006.

[166] Tam V. Nguyen and Luoqi Liu. "Smart mirror: intelligent makeup recommendation and synthesis." In: *Multimedia (MM '17)*. ACM, 2017, pp. 1253–1254.

[167] Jannik Boll Nielsen, Henrik Wann Jensen, and Ravi Ramamoorthi. "On Optimal, Minimal BRDF Sampling for Reflectance Acquisition." In: *Trans. on Graphics* 34.6 (2015).

[168] Gunnar A. Niklasson, Sara L. Niklasson, Celina Notfors, Junxin Wang, Maria Strømme, and Cecilia Århammar. *Optics of sunscreen lotions: Preliminary results on scattering and absorption coefficients*. arXiv:2204.13507 [physics.med-ph]. 2022.

[169] Heeso Noh, Seng Fatt Liew, Vinodkumar Saranathan, Simon G. J. Mochrie, Richard O. Prum, Eric R. Dufresne, and Hui Cao. "How Noniridescent Colors Are Generated by Quasi-ordered Structures of Bird Feathers." In: *Advanced Materials* 22.26-27 (2010), pp. 2871–2880.

[170] Jan Novák, Iliyan Georgiev, Johannes Hanika, and Wojciech Jarosz. "Monte Carlo Methods for Volumetric Light Transport Simulation." In: *Computer Graphics Forum* 37.2 (May 2018).

[171] Michael Oren and Shree K Nayar. "Generalization of Lambert's reflectance model." In: *Proc. of the 21st annual conference on Computer graphics and interactive techniques*. 1994, pp. 239–246.

[172] Romain Pacanowski, Xavier Granier, Christophe Schlick, and Pierre Poulin. "Sketch and Paint-based Interface for Highlight Modeling." In: *Eurographics Workshop on Sketch-Based Interfaces and Modeling*. Ed. by Christine Alvarado and Marie-Paule Cani. 2008.

[173] Juan Raúl Padrón-Griffe, Dario Lanza, Adrián Jarabo, and Adolfo Muñoz. "A Surface-based Appearance Model for Pennaceous Feathers." In: *Computer Graphics Forum* 43.7 (2024).

[174] Yanhong Pan et al. "The molecular evolution of feathers with direct evidence from fossils." In: *Proc. of the National Academy of Sciences* 116.8 (2019), pp. 3018–3023.

[175] Marios Papas, Krystle de Mesa, and Henrik Wann Jensen. "A Physically-Based BSDF for Modeling the Appearance of Paper." In: *Computer Graphics Forum* 33.4 (2014), pp. 133–142.

[176] Fabio Pellacini, James A Ferwerda, and Donald P Greenberg. "Toward a psychophysically-based light reflection model for image synthesis." In: *Proc. SIGGRAPH*. 2000, pp. 55–64.

[177] Matt Pharr, Wenzel Jakob, and Greg Humphreys. *Physically Based Rendering: From Theory to Implementation*. MIT Press, 2023.

[178] *Proyecto Unidigital IASAC Inteligencia Artificial y Sistemas Autónomos Cognitivos*. `https://unidigitaliasac.unizar.es/`. 2024.

[179] Richard O Prum. "The anatomy and physics of avian structural colours." In: *Proc. Int. Ornithol. Congr*. Vol. 22. 1999, pp. 1633–1653.

[180] R. Riedler, C. Pesme, J. Druzik, M. Gleeson, and E. Pearlstein. "A review of color-producing mechanisms in feathers and their influence on preventive conservation strategies." In: *Journal of the American Institute for Conservation* 53.1 (2014), pp. 44–65.

[181] Szymon Rusinkiewicz, Michael Burns, and Doug DeCarlo. "Exaggerated shading for depicting shape and detail." In: *Trans. on Graphics* 25.3 (2006), pp. 1199–1205.

[182] Iman Sadeghi, Oleg Bisker, Joachim De Deken, and Henrik Wann Jensen. "A Practical Microcylinder Appearance Model for Cloth Rendering." In: *Trans. on Graphics* 32.2 (2013).

[183] Yuichi Sakano and Hiroshi Ando. "Effects of head motion and stereo viewing on perceived glossiness." In: *Journal of Vision* 10.9 (2010), pp. 15–15.

[184] Kristina Scherbaum, Tobias Ritschel, Matthias Hullin, Thorsten Thormählen, Volker Blanz, and Hans-Peter Seidel. "Computer-suggested facial makeup." In: *Computer Graphics Forum* 30.2 (2011), pp. 485–492.

[185] Ana Serrano, Bin Chen, Chao Wang, Michal Piovarci, Hans-Peter Seidel, Piotr Didyk, and Karol Myszkowski. "The effect of shape and illumination on material perception: model and applications." In: *ACM Trans. on Graph.* 40.4 (2021).

[186] Ana Serrano, Diego Gutierrez, Karol Myszkowski, Hans-Peter Seidel, and Belen Masia. "An intuitive control space for material appearance." In: *Trans. on Graphics* 35.6 (2016).

[187] Matthew D Shawkey and Liliana D'Alba. "Interactions between colour-producing mechanisms and their effects on the integumentary colour palette." In: *Philosophical Trans. of the Royal Society B: Biological Sciences* 372.1724 (2017), p. 20160536.

[188] Weiqi Shi, Zeyu Wang, Cyril Soler, and Holly E. Rushmeier. "A Low-Dimensional Perceptual Space for Intuitive BRDF Editing." In: *EGSR*. 2021.

[189] Weiqi Shi, Zeyu Wang, Cyril Soler, and Holly Rushmeier. "A Low-Dimensional Perceptual Space for Intuitive BRDF Editing." In: *Proc. EGSR*. 2021.

[190] Peter Shirley, Samuli Laine, David Hart, Matt Pharr, Petrik Clarberg, Eric Haines, Matthias Raab, and David Cline. "Sampling Transformations Zoo." In: *Ray Tracing Gems: High-Quality and Real-Time Rendering with DXR and Other APIs*. Ed. by Eric Haines and Tomas Akenine-Möller. Apress, 2019, pp. 223–246.

[191] Robert Siegel and John R Howell. *Thermal Radiation Heat Transfer*. Taylor & Francis, 2001.

[192] Bruce Smith. "Geometrical shadowing of a random rough surface." In: *IEEE Trans. on antennas and propagation* 15.5 (1967), pp. 668–671.

[193] Brian E. Smits and Gary W. Meyer. "Newton's Colors: Simulating Interference Phenomena in Realistic Image Synthesis." In: *Photorealism in Computer Graphics*. Ed. by Kadi Bouatouch and Christian Bouville. Springer Berlin Heidelberg, 1992, pp. 185–194.

[194] Ying Song, Xin Tong, Fabio Pellacini, and Pieter Peers. "SubEdit: A Representation for Editing Measured Heterogeneous Subsurface Scattering." In: *Trans. on Graphics* 28.3 (2009).

[195] Jos Stam. "An illumination model for a skin layer bounded by rough surfaces." In: *Rendering Techniques 2001 (EGWR 2001)*. Springer, 2001, pp. 39–52.

[196] Doekele G Stavenga, Hein L Leertouwer, N Justin Marshall, and Daniel Osorio. "Dramatic colour changes in a bird of paradise caused by uniquely structured breast feather barbules." In: *Proc. of the Royal Society B: Biological Sciences* 278.1715 (2011), pp. 2098–2104.

[197] Doekele G. Stavenga, Jan Tinbergen, Hein L. Leertouwer, and Bodo D. Wilts. "Kingfisher feathers – colouration by pigments, spongy nanostructures and thin films." In: *Journal of Experimental Biology* 214.23 (2011), pp. 3960–3967.

[198] Shlomi Steinberg, Pradeep Sen, and Ling-Qi Yan. "Towards Practical Physical-Optics Rendering." In: *ACM Trans. on Graphics* 41.4 (2022), 1–13.

[199] Shlomi Steinberg and Ling-Qi Yan. "A generic framework for physical light transport." In: *ACM Trans. on Graphics* 40.4 (2021), pp. 1–20.

[200] Shlomi Steinberg and Ling-Qi Yan. "Physical light-matter interaction in hermite-gauss space." In: *ACM Trans. on Graphics* 40.6 (2021), pp. 1–17.

[201] Katherine R Storrs, Barton L Anderson, and Roland W Fleming. "Unsupervised learning predicts human perception and misperception of gloss." In: *Nature human behaviour* 5.10 (2021), pp. 1402–1417.

[202] Lisa Marie Streit. "Modelling of Feather Coat Morphogenesis for Computer Graphics." PhD thesis. 2003, pp. 9–24.

[203] Lisa Streit and Wolfgang Heidrich. "A Biologically-Parameterized Feather Model." In: *Computer Graphics Forum* 21.3 (2002), pp. 565–573.

[204] J Daniel Subias and Manuel Lagunas. "In-the-wild Material Appearance Editing using Perceptual Attributes." In: *arXiv preprint arXiv:2302.03619* (2023).

[205] Tarah N. Sullivan, Bin Wang, Horacio D. Espinosa, and Marc A. Meyers. "Extreme lightweight structures: avian feathers and bones." In: *Materials Today* 20.7 (2017), pp. 377–391.

[206] Zhaoyang Sun, Feng Liu, Wen Liu, Shengwu Xiong, and Wenxuan Liu. "Local facial makeup transfer via disentangled representation." In: *Proc. Asian Conference on Computer Vision (ACCV)*. Springer, 2020, pp. 8538–8551.

[207] Joshua B. Tenenbaum, Vin de Silva, and John C. Langford. "A Global Geometric Framework for Nonlinear Dimensionality Reduction." In: *Science* 290.5500 (2000), pp. 2319–2323.

[208] Jan Tinbergen, Bodo D. Wilts, and Doekele G. Stavenga. "Spectral tuning of Amazon parrot feather coloration by psittacofulvin pigments and spongy structures." In: *Journal of Experimental Biology* 216.23 (2013), pp. 4358–4364.

[209] Shoji Tominaga and Yusuke Moriuchi. “Principal component analysis-based reflectance analysis/synthesis of cosmetic foundation.” In: *Journal of Imaging Science and Technology* 53.6 (2009), pp. 60403–1–60403–8.

[210] Kenneth E Torrance and Ephraim M Sparrow. “Theory for off-specular reflection from roughened surfaces.” In: *Journal of the Optical Society of America* 57.9 (1967), pp. 1105–1114.

[211] Matteo Toscani, Dar’ya Guarnera, Giuseppe Claudio Guarnera, Jon Yngve Hardeberg, and Karl R. Gegenfurtner. “Three Perceptual Dimensions for Specular and Diffuse Reflection.” In: *ACM Trans. Appl. Percept.* 17.2 (2020).

[212] Loïc Tran, Philippe Lalanne, Benjamin Askenazi, and Kevin Vynck. “Physically based modeling of the colored appearance of complex media for cosmetics.” In: *Proc. of the 33rd IFSCC Congress*. 2023.

[213] Peter Vangorp, Timothy S Condon, James A Ferwerda, Kavita Bala, Roeland Schoukens, and Philip Dutré. “Visual equivalence in dynamic scenes.” In: (2009).

[214] Peter Vangorp, Jurgen Laurijssen, and Philip Dutré. “The influence of shape on the perception of material reflectance.” In: *Trans. on Graphics* (2007), 77–es.

[215] Eric Veach. *Robust Monte Carlo Methods for Light Transport Simulation*. Vol. 1610. Stanford University PhD thesis, 1997.

[216] Romain Vergne, Romain Pacanowski, Pascal Barla, Xavier Granier, and Christophe Schlick. “Light warping for enhanced surface depiction.” In: *ACM SIGGRAPH 2009 papers*. 2009, pp. 1–8.

[217] N Walsh, James Dale, Kevin J McGraw, MA Pointer, and Nicholas I Mundy. “Candidate genes for carotenoid coloration in vertebrates and their expression profiles in the carotenoid-containing plumage and bill of a wild bird.” In: *Proc. of the Royal Society B: Biological Sciences* 279.1726 (2012), pp. 58–66.

[218] Vincent Walsh and Janusz Kulikowski. *Perceptual constancy: Why things look as they do*. Cambridge University Press, 1998.

[219] Bruce Walter, Stephen R Marschner, Hongsong Li, and Kenneth E Torrance. “Microfacet Models for Refraction through Rough Surfaces.” In: vol. 2007. 2007.

[220] Zhaoyi Wan, Haoran Chen, Jie An, Wentao Jiang, Cong Yao, and Jiebo Luo. “Facial attribute transformers for precise and robust makeup transfer.” In: *Proc. Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 2022, pp. 1717–1726.

[221] Beibei Wang, Wenhua Jin, Miloš Hašan, and Ling-Qi Yan. “SpongeCake: A Layered Microflake Surface Appearance Model.” In: *Trans. on Graphics* 42.1 (2022).

[222] Junxin Wang, Changgang Xu, Annica M. Nilsson, Daniel Luis Abreu Fernandes, and Gunnar A. Niklasson. “A novel phase function describing light scattering of layers containing colloidal nanospheres.” In: *Nanoscale* 11.15 (2019), pp. 7404–7413.

[223] Rui Wang, Ewen Cheslack-Postava, Rui Wang, David Luebke, Qianyong Chen, Wei Hua, Qunsheng Peng, and Hujun Bao. “Real-Time Editing and Relighting of Homogeneous Translucent Materials.” In: *Vis. Comput.* 24.7 (2008), 565–575.

[224] Zhou Wang, Eero P Simoncelli, and Alan C Bovik. "Multiscale structural similarity for image quality assessment." In: *The Thrity-Seventh Asilomar Conference on Signals, Systems & Computers, 2003*. Vol. 2. Ieee. 2003, pp. 1398–1402.

[225] Gregory J. Ward. "Measuring and modeling anisotropic reflection." In: *Computer Graphics (SIGGRAPH '92)* 26.2 (1992), pp. 265–272.

[226] Eliza Wargala, Martyna Sławska, Agnieszka Zalewska, and Magdalena Toporowska. "Health effects of dyes, minerals, and vitamins used in cosmetics." In: *Women* 1.4 (2021).

[227] Ian Watts. "Red ochre, body painting, and language: interpreting the Blombos ochre." In: *The cradle of language* 2 (2009), pp. 93–129.

[228] Andrea Weidlich and Alexander Wilkie. "Arbitrarily layered micro-facet surfaces." In: *GRAPHITE*. ACM, 2007, pp. 171–178.

[229] Philippe Weier and Laurent Belcour. "Rendering Layered Materials with Anisotropic Interfaces." In: *Journal of Computer Graphics Techniques* 9.2 (2020), pp. 37–57.

[230] Gunnar Wendt, Franz Faul, Vebjørn Ekroll, and Rainer Mausfeld. "Disparity, motion, and color information improve gloss constancy performance." In: *Journal of vision* 10.9 (2010), pp. 7–7.

[231] Sebastian Werner, Zdravko Velinov, Wenzel Jakob, and Matthias B Hullin. "Scratch iridescence: Wave-optical rendering of diffractive surface structure." In: *ACM Trans. on Graphics* 36.6 (2017), pp. 1–14.

[232] Alexander Wilkie and Andrea Weidlich. "A physically plausible model for light emission from glowing solid objects." In: *Computer Graphics Forum*. Vol. 30. 4. Wiley Online Library. 2011, pp. 1269–1276.

[233] Josh Wills, Sameer Agarwal, David Kriegman, and Serge Belongie. "Toward a perceptual space for gloss." In: *Trans. on Graphics* 28.4 (2009), pp. 1–15.

[234] Magnus Wrenninge, Ryusuke Villemin, and Christophe Hery. "Path traced subsurface scattering using anisotropic phase functions and non-exponential free flights." In: *Tech. Rep.* Pixar Inc., 2017.

[235] Douglas R Wyman, Michael S Patterson, and Brian C Wilson. "Similarity relations for the interaction parameters in radiation transport." In: *Applied optics* 28.24 (1989), pp. 5243–5249.

[236] Mengqi (Mandy) Xia, Bruce Walter, Eric Michielssen, David Bindel, and Steve Marschner. "A Wave Optics Based Fiber Scattering Model." In: *Trans. on Graphics* 39.6 (2020).

[237] Mengqi Xia, Bruce Walter, Christophe Hery, and Steve Marschner. "Gaussian product sampling for rendering layered materials." In: *Computer Graphics Forum* 39.1 (2020), pp. 420–435.

[238] Mengqi Xia, Bruce Walter, Christophe Hery, Olivier Maury, Eric Michielssen, and Steve Marschner. "A Practical Wave Optics Reflection Model for Hair and Fur." In: *ACM Trans. on Graphics* 42.4 (2023), pp. 1–15.

[239] Bei Xiao, Bruce Walter, Ioannis Gkioulekas, Todd Zickler, Edward Adelson, and Kavita Bala. "Looking against the light: How perception of translucency depends on lighting direction." In: *Journal of vision* 14.3 (2014), pp. 17–17.

[240] Bei Xiao, Shuang Zhao, Ioannis Gkioulekas, Wenyan Bi, and Kavita Bala. "Effect of geometric sharpness on translucent material perception." In: *Journal of vision* 20.7 (2020), pp. 10–10.

[241] Kun Xu, Yue Gao, Yong Li, Tao Ju, and Shi-Min Hu. "Real-time homogeneous translucent material editing." In: *Computer Graphics Forum* (2007).

[242] Lin Xu, Yangzhou Du, and Yimin Zhang. "An automatic framework for example-based virtual makeup." In: *Proc. International Conference on Image Processing (ICIP)*. IEEE, 2013, pp. 3206–3210.

[243] Ling-Qi Yan, Henrik Wann Jensen, and Ravi Ramamoorthi. "An Efficient and Practical near and Far Field Fur Reflectance Model." In: *Trans. on Graphics* 36.4 (2017).

[244] Ling-Qi Yan, Chi-Wei Tseng, Henrik Wann Jensen, and Ravi Ramamoorthi. "Physically-Accurate Fur Reflectance: Modeling, Measurement and Rendering." In: *Trans. on Graphics* 34.6 (2015).

[245] Qixin Yan, Chunle Guo, Jixin Zhao, Yuekun Dai, Chen Change Loy, and Chongyi Li. "BeautyREC: robust, efficient, and component-specific makeup transfer." In: *Proc. Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2023, pp. 1102–1110.

[246] Yuchun Yan, Juhyun Lee, Jongin Hong, and Hyeon-Jeong Suk. "Measuring and describing the discoloration of liquid foundation." In: *Color Research and Application* 46.2 (2020), pp. 362–375.

[247] Chenyu Yang, Wanrong He, Yingqing Xu, and Yang Gao. "EleGANt: Exquisite and locally editable GAN for makeup transfer." In: *Proc. European Conference on Computer Vision (ECCV)*. Springer, 2022, pp. 737–754.

[248] Xingchao Yang, Takafumi Taketomi, and Yoshihiro Kanamori. "Makeup extraction of 3D representation via illumination-aware image decomposition." In: *Computer Graphics Forum* 42.2 (2023), pp. 293–307.

[249] Kenichiro Yoshida and Natsuko Okiyama. "Estimation of reflectance, transmittance, and absorbance of cosmetic foundation layer on skin using translucency of skin." In: *Opt. Express* 29.24 (2021), pp. 40038–40050.

[250] Tizian Zeltner and Wenzel Jakob. "The layer laboratory: a calculus for additive and subtractive composition of anisotropic surface reflectance." In: *ACM Trans. on Graphics* 37.4 (2018), 74:1–74:14.

[251] Fucheng Zhang, Stuart Kearns, Patrick Orr, Michael Benton, Zhonghe Zhou, Diane Johnson, Xing Xu, and Xiaolin Wang. "Fossilized melanosomes and the colour of Cretaceous dinosaurs and birds." In: *Nature* 463 (2010), pp. 1075–8.

[252] Jiajun Zhang and Takashi Kanai. "Biological Modeling of Feathers by Morphogenesis Simulation." In: *Proc. International Conference on Cyberworlds (CW)*. 2020, pp. 63–70.

[253] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. "The Unreasonable Effectiveness of Deep Features as a Perceptual Metric." In: *Proc. Computer Vision and Pattern Recognition (CVPR)* (2018), pp. 586–595.