

Diss. ETH No. 23031

A computational appearance fabrication framework and derived applications

A dissertation submitted to
ETH Zurich

for the Degree of
Doctor of Sciences

presented by

Marios Papas

MSc in Computer Science, UC San Diego, USA

born October 18, 1983

citizen of Cyprus

accepted on the recommendation of

Prof. Dr. Markus Gross, examiner

Prof. Dr. Wojciech Jarosz, co-examiner

Prof. Dr. Steve Marschner, co-examiner

2015

Abstract

Traditionally, control over the appearance of objects in the real world was performed manually. Understanding how some physical property of an object would affect its appearance was achieved primarily through trial and error. This procedure could be lengthy and cumbersome, depending on the complexity of the effect of physical properties on appearance and the duration of each fabrication cycle. Precise control of how light interacts with materials has many applications in arts, architecture, industrial design, and engineering. With the recent achievements in geometry retrieval and computational fabrication we are now able to precisely control and replicate the geometry of real-world objects. On the other hand, computational appearance fabrication is still in its infancy. In this thesis we lay the foundation for a general computational appearance fabrication framework, and we demonstrate a range of applications that benefit from it. We present various instances of our framework and detail the design of the corresponding components, such as: forward and backward appearance models, measurement, and fabrication. These framework instances help in understanding and controlling the appearance of three general classes of materials: homogeneous participating media (such as wax and milk), specular surfaces (such as lenses), and granular media (such as sugar and snow). More specifically we show how we can precisely measure, control, and fabricate the real-world appearance of homogeneous translucent materials, how to computationally design and fabricate steganographic lenses, and finally we present a fast appearance model for accurately simulating the appearance of granular media.

Zusammenfassung

Traditionell wurde Kontrolle über das Erscheinungsbild von Objekten der realen Welt manuell ausgeübt. Das Verständnis darüber, auf welche Art und Weise die Veränderung physikalischer Eigenschaften das Erscheinungsbild beeinflussen, wurde hauptsächlich durch systematisches Ausprobieren erlangt. Dieses Verfahren konnte jedoch langwierig und mühsam sein, abhängig von der Komplexität des Einflusses der physikalischen Eigenschaften auf das Erscheinungsbild, und der Dauer jedes Fertigungszykluses. Die präzise Steuerung von Lichtinteraktionen mit verschiedenen Materialien hat vielfältige Anwendungen in der Kunst, Architektur, sowie im Industriedesign und Ingenieurwesen. Mit den jüngsten Errungenschaften in der Geometrierekonstruktion sowie der computergesteuerten Fabrikation sind wir in der Lage, die Geometrie realer Objekte präzise zu steuern und zu replizieren. Andererseits steckt die computergesteuerte Herstellung von spezifischen Erscheinungsbildern immer noch in den Kinderschuhen. In dieser Dissertation legen wir die Grundlagen für ein Rahmenkonzept für allgemeine rechnergestützte Fabrikation von Erscheinungsbildern und zeigen eine Reihe von Anwendungen die davon profitieren. Wir präsentieren verschiedene Variationen unseres Konzepts und geben Details über das Design der entsprechenden Komponenten, wie etwa Forwärt- und Rückwärterscheinungsmodelle, Messung und Herstellung. Diese Instanzen unserer Konzepts erlauben uns das Erscheinungsbild dreier verschiedener allgemeinen Materialklassen - homogene partizipäre Medien (wie etwa Wachs oder Milch), spiegelnde Oberflächen (wie etwa Linsen) und granulare Medien (wie etwa Zucker oder Schnee) zu verstehen und zu kontrollieren. Insbe-

sondere zeigen wir wie man das Erscheinungsbild von homogenen lichtdurchlässigen Materialien präzise messen, kontrollieren und herstellen kann, wie man steganographischen Linsen computergestützt modellieren kann, und stellen letztlich ein effizientes Erscheinungsmodell für die akkurate Simulation des Aussehens granularer Medien vor.

for Georgia

Acknowledgments

First, I would like to thank my ETH advisor, Markus Gross, for his support, encouragement, and strategic vision throughout my doctoral studies. I thank my Disney Research advisor, Wojciech Jarosz, for his immense help and guidance at every step of this journey. I owe Wojciech a great debt of gratitude for teaching me on a daily basis how to think and perform as a researcher.

Furthermore, I would like to thank my mentors — Steve Marschner, Bernd Bickel, and Wojciech Matusik — for their advice and efforts during my studies. Their diverse skill set and knowledge were crucial components in the process of transforming ambitious ideas into successful publications.

I thank Wenzel Jakob for creating the Mitsuba Renderer and allowing me (and everyone in the community) to use his truly useful code-base. Almost every rendered image in this thesis was created in Mitsuba.

Most of the work presented in this thesis was made possible due to the contributions of my co-authors — Christian Regg, Johannes Meng, Phil Jackson, Ralf Habel, Derek Nowrouzezahrai, Carsten Dachsbacher, Markus Gross, Wojciech Jarosz, Steve Marschner, Bernd Bickel, and Wojciech Matusik.

I am indebted to Maurizio Nitti and Alessia Mara for their impactful contributions in our results and concept art. Their outstanding work allowed us to showcase the potential of our research.

I wish to express my gratitude to the people who engaged with me in fruitful scientific discussions during my studies — Amit Bermano, Olga Sorkine-Hornung, Daniele Panozzo, Iliyan Georgiev, Thabo Beeler, Derek Bradley, Fabrice Rousselle, Romain Prévost, Kaan Yücer, Jan Novák, Simon Heinzle, Fabio Zünd, Steven Poulakos, Nelson Max, Andrew Selle, Brent Burley, Chuck Tappan, Andreas Pitsillidis, and Klearchos Papas.

I owe a large amount of gratitude to the people who helped me directly with the writeup of this thesis. More specifically I thank Wojciech Jarosz, Steve Marschner, Jan Novák, Susanna Mineve, and Jodie Aberdeen for proofreading earlier versions of this document. I thank Henning Zimmer, Fabian Hahn, and Gosia Gubat for their help with the German abstract and Peter Kaufmann, Kaan Yücer, Yağiz Aksoy, Fabian Hahn, Amit Bermano, Alexandre Chapiro, Changil Kim, and Olga Diamanti for their help with L^AT_EX and the thesis structure.

I consider myself very fortunate to have the support of an amazing administration team including Stephan Veen, Andreas Baumann, Linda Breu, Martina Haefeli, Stefanie Lacher, Sarah Disch, Elena Habicher, Denise Spicher, and Markus Portmann. Thank you for your support, I could not have asked for a better working environment.

Finally, I would like to thank my parents, Androulla and Kyriakos for supporting me throughout my studies, my wife Georgia who stood by my side throughout this journey, and finally my brother Klearchos for showing me (by example) how interesting research can be.

Contents

Abstract	iii
Zusammenfassung	v
Acknowledgements	ix
Contents	xi
List of Figures	xv
List of Tables	xix
Introduction	1
1.1 Contributions	9
1.2 Organization	11
Related work	13
2.1 Appearance capture, modeling and fabrication	13
2.2 Translucency fabrication	19
2.3 Steganographic lens fabrication	22
2.4 Rendering granular media	26
2.5 Summary	29
Translucency fabrication	31
3.1 Introduction	34
3.2 Method overview	38

Contents

3.3	The forward appearance model	39
3.4	Measurement setup	41
3.5	Designing the backward model	46
3.5.1	Measurement database selection	47
3.5.2	Fitting the forward model to a single material	49
3.5.3	Global pigment parameter estimation	51
3.5.4	Mixture optimization	53
3.6	Local pigment parameter estimation	54
3.7	Appearance editing	56
3.7.1	Color reproduction	57
3.7.2	Color editing	58
3.7.3	Translucency editing	59
3.8	Fabrication	60
3.9	Results	62
3.9.1	Reproduction of simulated examples	62
3.9.2	Reproduction of real-world examples	64
3.10	Perceptual extension	68
3.11	Discussion and future work	71
	Steganographic lens fabrication	77
4.1	Introduction	79
4.2	Problem statement and goals	82
4.3	Generating source-optimized lenses	84
4.3.1	Input	84
4.3.2	Facet-patch matching	85
4.3.3	Facet orientation	86
4.3.4	Simulated annealing	87
4.3.5	Multi-scale height optimization	90
4.3.6	Facet-patch matching [optional]	91
4.3.7	Output	92
4.4	Generating universal steganographic lenses	93
4.4.1	Orientation initialization with dart-throwing	95
4.4.2	In-painting and texture synthesis	96
4.4.3	Iterative adjustment and convergence	96

4.5	Fabrication details and results	97
4.6	Discussion	100
4.7	Conclusion and future work	106
A fast, forward model for granular media		109
5.1	Introduction	110
5.2	Deriving diffusion parameters	112
5.3	Switching to diffusion	113
5.4	Multi-pole configuration	115
5.5	Results	117
5.6	Conclusion and discussion	121
Conclusion		125
Appendix A		133
Appendix B		141
References		147

List of Figures

1.1	The main components of the appearance fabrication framework proposed in this thesis	2
1.2	Interaction flowchart among the components of the proposed computational appearance fabrication framework	6
1.3	Sample results produced by the presented framework applications.	8
3.1	An instance of our framework depicting an example of translucency fabrication	32
3.2	A side by side comparison of real translucent materials next to their silicone replicas fabricated using our method	34
3.3	Overview of our system	37
3.4	An illustration of the setup used to measure the appearance of a target sample, and a resulting captured HDR image	42
3.5	A log plot of a measured bulk scattering profile as a function of distance	50
3.6	Per-wavelength-band scattering coefficients for the blue pigment dilution set as a function of concentrations	53
3.7	Screen-shots of our appearance editor	57
3.8	Out-of-gamut target materials and spectral reflectance measurements of our pigments	67
3.9	Measured target and fabricated replica HDR pairs for a variety of materials	74

List of Figures

3.10	Synthetic target and replica comparisons for the global method	74
3.11	Synthetic target and replica comparisons when using a simpler optimization that fits directly on parameters . .	74
3.12	Appearance prediction evaluation of our local method	75
4.1	An instance of our framework, used for source-optimized steganographic lens fabrication	78
4.2	A sketch of other potential applications for our steganographic lenses	80
4.3	Overview of source-optimized steganographic lens generation	85
4.4	Influence of simulated annealing on the lens geometry and rendering quality	87
4.5	Basic operation of simulated annealing	88
4.6	Height optimization displacement example	91
4.7	Heatmap of overlaps in a source image	93
4.8	Overview of universal steganographic lens generation .	94
4.9	Video sequence unscrambling example using a universal steganographic lens	95
4.10	Four source images with questions are warped to reveal pictographic answers with a single universal steganographic lens	97
4.11	A physically milled result of a source optimized lens .	99
4.12	A physically printed multi-view result	101
4.13	A photographic zoom-in of one of our 3D printed steganographic lenses	102
4.14	Inducing stable viewing by exploiting multi-view lens generation pipeline	104
4.15	Visualization of the warping behavior of a universal steganographic lens	105
4.16	A simulated result of a single source-optimized lens which encodes four hidden pictographic answers . . .	106

5.1	Diagonal comparison of proposed techniques with EPT on a sphere composed of snow grains	112
5.2	Overview of the full method presented in Meng et al. [2015], including our diffusion approximation extension	113
5.3	A selection of real granular materials illuminated from underneath with a cell phone’s flash light to reveal their translucency.	114
5.4	Rendering comparisons of the SANDCASTLE scene with 2 billion grains	118
5.5	Rendering comparisons of our SPICES scene showcasing flour, pink salt, brown sugar, and white sugar . . .	119
5.6	We visualize the individual contributions of each technique, which sum to our final image.	119
5.7	Rendering comparisons on a tall and a tiny SNOWMAN	121
A.1	Reflectance model investigation results	135
A.2	Profile comparisons between two rendering methods .	136
A.3	Target/Prediction/Replica comparison plots for the local method	137
A.4	White pigment dilution set measurements and model fits	138
A.5	Yellow pigment dilution set measurements and model fits	138
A.6	Red pigment dilution set measurements and model fits	139
A.7	Green pigment dilution set measurements and model fits	139
A.8	Blue pigment dilution set measurements and model fits	140
A.9	Black pigment dilution set measurements and model fits	140
B.1	Source optimized lens example with four source-target image pairs	142
B.2	Simulation of surface roughness effects on the target image	142
B.3	Input and output examples for a steganographic universal lens	143

List of Figures

B.4	More input and output examples for a steganographic universal lens	144
B.5	Image morphing example (in simulation) using a source optimized lens	145
B.6	More image morphing examples (in simulation) using a source optimized lens	145

List of Tables

3.1	Global estimate of reduced scattering and absorption coefficients for our pigments and base silicone	65
3.2	Concentration percentages prediction generated using our local method for the 9 targets	66
5.1	Time to unit variance comparisons (measured in seconds) for all methods	120
A.1	Error values for the 9 fabricated replicas using the local method with cosine concentration weights	134
A.2	Error values for the 9 fabricated replicas using the global method without local weights	134
A.3	Reduced scattering and absorption coefficients used for the synthetic database and targets	136

List of Tables

CHAPTER

1

Introduction

Controlling the appearance of objects in the real world has been a long standing goal of artists. This process involves modifying some physical properties of the object, which in turn have some indirect (and sometimes hard to understand) effect on appearance. Traditionally this has been performed manually and required a potentially lengthy trial-and-error procedure, depending on the complexity of the physical property's effect on appearance.

Precise control and understanding of appearance is desired not only by artists but is also useful for a plethora of other applications such as design of reflective/transmissive displays, calibration standards, lenses, headlights, retro-reflective gear, telescope lining [Brown *et al.* 2002], architecture, visible light communication, and more.

The appearance of an object is exposed to an observer when that object is subjected to some illumination conditions. Appearance can be observed or measured and it is dependent on the physical properties of the object, the illumination conditions, and the response of the observer. In this thesis we focus on




Figure 1.1: *The main components of the appearance fabrication framework proposed in this thesis.*

static appearance, ignoring dynamic appearance changes over time. We describe the important components (Figure 1.1) for a complete computational appearance fabrication framework.

Physical state. First, we need to specify what we define as the *physical state*. The physical state is described by three main components: the physical properties of the object, the illumination conditions and the response of the observer. We define physical properties, as the object properties that can potentially be controlled and may affect appearance. Some examples of physical properties are the composition of materials within the object (discrete or continuous), the shape or geometry of the object, surface roughness, layers and thickness thereof, and others. A requirement for the selected fabrication method should be to allow precise control over the desired appearance-relevant "knobs" of the physical state.

Forward appearance model. As a second step, we specify a *forward appearance model* that takes as input the physical state and simulates the observer's response. While a variety of physically based appearance models can be used to accurately simulate appearance, it is rather uncommon to have an appearance model with parameters that can be mapped exactly to an available physical state. In such cases we can compute a mapping from a physical state to forward model parameters, and consequently to the corresponding appearance. This mapping or training can be performed by an optimization which uses as input the forward model and a variety of appearance measurements with known corresponding physical states. The difficulty of this mapping and the number of required measurements depend on the expressiveness and accuracy of the forward model. Selection or design of an appropriate forward model for a target material class should be performed with care, such that the forward model can efficiently and accurately simulate a wide gamut of appearances, with as few parameters as possible.

Backward appearance model. Complimentary to the forward model, we also define a *backward appearance model*. This backward model uses as input a desired appearance and outputs fabrication specifications or a desired physical state. The backward model is tightly coupled with the fabrication method and should only output fabrication specifications that can be reproduced in the real world. It is common for a backward appearance model to utilize a copy of the forward appearance model and an optimization procedure for inverting appearance and converting it to fabrication specifications.

Appearance measurements. We would like to emphasize the need for accurate *appearance measurements* and include them

Introduction

as another important component within our framework. There are two main use cases where measurements are useful. The first use case is for training the forward model and mapping a physical state to the model parameters. In addition, measurements are also useful in the case of appearance modeling where a target appearance can be captured directly from the real world. When designing measurements it is important to have in mind the forward model and the main characteristics of the general appearance class to be measured. Measurements can be focused in parts of the domain where there is interesting appearance variation with respect to the model parameters. In addition, measurements could potentially be performed in simplified physical setups where the known components of the physical state are easy to reproduce in simulation as well as allowing for the forward model to be computationally efficient.

A motivating example. Using an implementation of the aforementioned components, we can examine a naive appearance replication example. Assume a simplified scenario where we would like to replicate the appearance of gray diffuse wall paint. We have a target paint sample that we would like to replicate. The goal of our automated process is to output a mixture of white and black paint (physical properties) such that the replica will have the same appearance under a specific illumination and observer. We need to define the physical state, the forward, and backward models. To define the physical state, we need to specify the illuminant, the observer, and the properties controlled by the fabrication method. The physical property which we can control in this case is the mixture of the replica. Since this is a low dimensional problem, the forward model can be a very simple 1D interpolation scheme of dense paint reflectance measurements with known mixtures.

The final component needed for replication is the backward model, which will provide corresponding physical properties for a given target appearance. In this case we can use a very naive backward model, which for a given target appearance finds the closest appearance measurement (e.g. CIEDE2000 [Sharma *et al.* 2005]) and returns the known physical properties (paint mixture) that correspond to it.

While this replication process is generic and does not require a complex appearance model, it does require dense sampling of the possible physical states. In the common case where the space of physical states is high dimensional, dense sampling might not be feasible. In such cases an expressive and low dimensional physically based appearance model can be useful. For example, the instance of our framework for mapping pigment concentrations to color and translucency (described in Chapter 3) only required a total of 21 training measurements with different pigment mixtures.

Appearance editor. We just described the necessary components of our framework to replicate a target appearance. In use cases where we would like to allow intuitive control over the appearance, we will add an additional procedure (*appearance editor*) between the forward model and the backward model. This procedure will transform a possibly measured appearance to a desired one, with the help of some meaningful and intuitive high-level edits. Then it is the responsibility of the backward model to find a physical state that, when fabricated, will match the novel appearance up to a user-specified threshold.

In summary, we highlight the main components and their interactions within our framework in Figure 1.2. Furthermore, the physical state encapsulates the physical parameters of the object, the illumination conditions and the observers' response.

Introduction




Figure 1.2: Interaction flowchart among the components of the proposed computational appearance fabrication framework. The appearance of a physical state (composed out of illumination conditions, physical properties, and an observer) can be depicted virtually with the use of a forward model. This virtual appearance can be edited and then passed as input to the backward model. Consecutively the backward model will translate the novel appearance to fabrication specifications that can be translated by the fabrication method to a new physical state with the desired appearance.

The forward model encapsulates the appearance model, training measurements, and a mapping from a physical state to appearance. The backward model is responsible for mapping appearance to fabrication specifications or a desired physical state. Finally the chosen fabrication method is tasked with translating the output of the backward model to a physical state in the real world.

In this thesis we show a range of applications that benefit from our computational appearance fabrication framework. We present various instances of our framework and detail the design of the corresponding components, such as forward and backward models, appearance editors, measurements design, and fabrication methods. Our framework that can be used for computational appearance fabrication within three general

classes of materials: homogeneous participating media, specular surfaces, and granular media.

Within these material classes we propose three applications which make use of our framework. For the case of homogeneous participating media we present a fully populated instance of our framework that can be used to fabricate replicas of translucent materials by using continuous mixtures of pigments with silicone. This method automatically deduces pigment concentrations that will match a target homogeneous translucent appearance. In the case of specular surfaces we present a method for generating the shape of steganographic lenses that can unscramble seemingly random patterns into desired images when the patterns are seen through the lens. This method proposes two different backward models (one for milling and one for 3D printing) along with a novel appearance editor that can be used in combination for translating desired scrambled and unscrambled image pairs into fabrication specifications. Finally, for the special case of granular media we only deal with appearance modeling and we provide a fast parametric forward appearance model for simulating their appearance. Example results of our applications are shown in Figure 1.3.

The current state of the art is at different stages for each material class and application. Thus, for some applications further research is needed in every part of the framework, and sometimes in a subset of the framework components. In the case of homogeneous participating media we have a set of forward appearance models that can be used to simulate their translucent appearance, but they are not suitable for fabrication. There is also work on fabricating spatially varying translucency using a discrete set of homogeneous voxels [Hasan *et al.* 2010] or layers [Tong *et al.* 2005]. We provide all the necessary components of our framework that can be used to solve an orthogonal problem: how to fabricate a desired homogeneous translucent appear-

Introduction




Figure 1.3: Sample results produced by the presented framework applications. On the left we provide a photograph of a glass filled with blue fabric softener side-by-side with its translucent silicone replica, generated by our automated method. In the middle we present a photograph of a computationally designed steganographic lens that can be used to unscramble patterns into predefined images. Finally, on the right we showcase a synthetic rendering of billions of snow grains rendered with our forward appearance model for granular media. This rendering converges $260\times$ faster than the reference.

ance by using continuous pigment mixtures. Our method can be used for completely homogeneous translucent objects but can also be potentially used to fabricate the primitives used by the aforementioned heterogeneous methods.

Similarly, for the case of controlling how light interacts with specular surfaces the current state of the art is also in a pretty complete state. For example we do have forward and backward models that can help us transition between physical properties and appearance. A common application is allow bending of directional light into images by computationally designing the shape of a lens. In contrast, our application aims to compute the geometry of a steganographic lens that unscrambles seemingly random looking patterns into desired images.

Contrary to homogeneous participating media and specular surfaces, the appearance of granular media is currently not well understood. We do not have parametric appearance models that can simulate their appearance at a reasonable level of

efficiency and accuracy. In this thesis we focus on solving this appearance modeling problem first, which is a prerequisite for attempting computational appearance fabrication of granular media.

1.1 Contributions

The main contributions of this thesis can be summarized as follows:

- A general framework for computational appearance fabrication.
- An application of our framework for fabricating homogeneous participating media.
- A multi-spectral measurement device for measuring reflectance and translucency.
- An application of our framework for fabricating steganographic lenses.
- A fast appearance model for granular media.

The aforementioned contributions have been published in the following international journal articles:

- Marios Papas, Christian Regg, Wojciech Jarosz, Bernd Bickel, Philip Jackson, Wojciech Matusik, Steve Marschner, Markus Gross. **Fabricating Translucent Materials using Continuous Pigment Mixtures**. In *ACM Transactions on Graphics (Proceedings of SIGGRAPH)*, 32(4), July 2013.
- Marios Papas, Thomas Houit, Derek Nowrouzezahrai, Markus Gross, Wojciech Jarosz. **The Magic Lens**:

Refractive Steganography. In *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)*, 31(6), November 2012.

- Johannes Meng, Marios Papas, Ralf Habel, Carsten Dachsbacher, Steve Marschner, Wojciech Jarosz, Markus Gross. **Multi-Scale Modeling and Rendering of Granular Materials**¹. In *ACM Transactions on Graphics (Proceedings of SIGGRAPH)*, 34(4), July 2015.

During the period of my doctoral studies, I have co-authored the following journal articles that are not directly related to this thesis:

- Marios Papas, Wojciech Jarosz, Wenzel Jakob, Szymon Rusinkiewicz, Wojciech Matusik, Tim Weyrich. **Goal-based Caustics.** In *Computer Graphics Forum (Proceedings of Eurographics)*, 30(2):503–511, June 2011.
- Marios Papas, Krystle De Mesa, Henrik Wann Jensen. **A Physically-Based BSDF for Modeling the Appearance of Paper.** In *Computer Graphics Forum (Proceedings of EGSR)*, 33(4):133-142, June 2014.
- Oliver Klehm, Fabrice Rousselle, Marios Papas, Derek Bradley, Christophe Hery, Bernd Bickel, Wojciech Jarosz, Thabo Beeler. **Recent Advances in Facial Appearance Capture.** In *Computer Graphics Forum (Proceedings of Eurographics)*, 34(2):709–733, May 2015.

¹From this paper we only consider the diffusion extension of the model as a contribution of this thesis.

1.2 Organization

In Chapter 2 we present an overview of the existing work relevant to this thesis. More specifically we summarize the current state of the art for each class of materials and highlight the work that has been done until now that is relevant to computational appearance fabrication.

In Chapter 3 we present an application of our full computational appearance fabrication framework for homogeneous participating media. We present a device for measuring the color and translucency of homogeneous translucent materials. By using this device, we measure various mixtures of silicone with pigments to deduce their scattering parameters. We use the same device to measure the appearance of some target materials, and we present a method for finding a pigment mixture which, when mixed, will have that target appearance. In addition, we present a method for editing appearance which allows the user to provide color and translucency edits while taking into account what can actually be fabricated by the set of available pigments.

In Chapter 4 we propose a method for computational fabrication of steganographic lenses. The main application is to scramble the incident light field on the array such that it can be used for steganographic purposes.

In Chapter 5 we present an extension to the rendering method of Meng et al. [2015] for simulating the appearance of granular media. Our key observation is that the appearance of granular media can be split into two main components: low-order and multiple scattering. The first observed effect is the structure in the appearance of granular media due to interactions at the outer part of the volume, which we attribute to low order scattering. The second effect is a more smoothly varying effect

Introduction

which we attribute to multiple scattering. In granular media with low absorption (e.g. cane sugar and snow) most of the contribution is due to multiple scattering. Multiple scattering is usually more difficult to simulate with traditional methods since it is the result of an extremely large number of interactions with the individual grains. Instead of simulating individual bounces of light, we propose to leverage a diffusion approach inspired by the work of Li et al. [2005] to efficiently summarize the effect of multiple scattering within granular materials with a single diffusion connection. This provides significant speedup in low absorbing granular media with minimal impact on accuracy.

The technical chapters of this thesis (Chapters 3, 4, 5, Appendix A and B) are largely word-for-word reproductions of published work that I have co-authored and use passages of text from these publications with explicit permission from the co-authors.

CHAPTER

2

Related work

In this section we present an overview of the relevant prior work in general appearance capture, modeling, editing, and fabrication (Section 2.1). We continue with an overview of the current state of the art relevant to our selected applications (Sections 2.2, 2.3 and 2.4).

2.1 Appearance capture, modeling and fabrication

Within the computer graphics community there is a vast amount of work on simulating and controlling how light interacts with materials from the real world. This is usually referred to as appearance modeling and facilitates the reproduction of real-world materials in virtual scenes. The main purpose of these appearance models is to simulate how light is scattered and absorbed as it interacts with a material. Surface interactions are modeled using a *Bidirectional Scattering Distribution Function* (BSDF) which encodes the ratio of outgoing radiance over incident irradiance as a function of the respective outgoing and

Related work

incident directions. This function can be reduced into a *Bidirectional Reflectance Distribution Function* (BRDF) when accounting only for reflectance (i.e. incident and outgoing directions lie on the same side of the surface), and into a *Bidirectional Transmittance Distribution Function* (BTDF) when accounting only for transmission (i.e. incident and outgoing directions lie on opposite surface sides).

BSDFs summarize surface interactions and are a key component of the *Rendering Equation* [Kajiya 1986]. Similarly phase functions (along with scattering and absorption coefficients) are used within the *Radiative Transfer Equation* [Chandrasekar 1960] to model interactions with participating media.

Many currently used physically-based renderers use Monte Carlo integration to solve these equations. Monte Carlo integration is commonly used to evaluate global illumination in research oriented renderers [Jakob 2010; Pharr and Humphreys 2010], but it is also heavily used in movie production (RenderMan[®], Maxwell Render[®], Hyperion, and Arnold among others). During this integration process, light paths originating from the sensor and/or the emitters are incrementally constructed as they interact with materials. In such interaction events, an outgoing direction and position are generated (sampled) and the appearance model is evaluated. The remaining path contribution is proportional to the value of the appearance model at each such interaction. Thus, we preferably want to generate the outgoing direction and position with a probability density function linearly related to the appearance model's function value. This is referred to as importance sampling and allows significant convergence improvements when used for Monte Carlo integration.

For this reason, we also include within the requirements of an appearance model the ability to efficiently sample and evaluate

2.1 Appearance capture, modeling and fabrication

light scattering behavior as a function of an incident direction and position. Another desired property of appearance models is being able to provide intuitive parameters that can be altered by the user for controlling various aspects of the appearance of a material such as glossiness, opaqueness or transparency, and color.

Physically-based appearance models designed with the goal of modeling the appearance of real-world materials often make use of measurements. These measurements can be used to discover and expose a characteristic appearance property and can also be used to validate the model based on how well it matches the real-world measurements.

Appearance models can be divided into two main categories: parametric and data-driven. Parametric models usually have some simplified interpretation of the underlying physical properties of the corresponding material. The parameters of these models can be either related to a physical property or to a desired appearance effect. When a more direct relation between the model parameters and the physical properties is required, measurements and an optimization procedure can be used to discover such a mapping. An advantageous property of parametric models with parameters related to physical properties, is that once the mapping is found then the model could be used to predict appearance even outside the convex hull of measurements.

Data-driven models are more heavily dependent on measurements and do not require a good understanding of the underlying physical properties. Oftentimes they can be better than parametric models at capturing the appearance of the target material, especially near the measured parts of the appearance domain. There are however some drawbacks and challenges with data-driven models. One drawback is that they potentially

Related work

require denser measurements than parametric models, depending on the variation and the dimensionality of the appearance within the domain. Unlike physically based parametric models, extrapolation outside the convex hull of measurements is more difficult. In addition, a summary of the measurements is usually needed in memory such that the model can be evaluated and sampled efficiently. Finally, another challenge with data-driven models is appearance editing. It is very difficult to devise methods that allow appearance edits that widely span the gamut of the represented material class.

There are many examples of specialized appearance models that can be used to simulate the appearance of real-world materials and objects. For materials with micro-surface roughness there is a wide variety of research on simulating the aggregate reflectance [Torrance and Sparrow 1967; Cook and Torrance 1981] and transmission [Walter *et al.* 2007] of specular micro-facets oriented according to an analytic micro-facet normal distribution function. These analytic parametric models were validated against measurements and can be efficiently evaluated and sampled. Oren and Nayar [1994] developed an analytic forward model for approximating the reflectance of diffuse micro-facets. Ashikmin *et al.* [2000] provided a method which allows an arbitrary micro-facet normal distribution function to be used for creating plausible energy preserving BRDFs. Ngan *et al.* [2005] performed BRDF measurements on a variety of materials and concluded that for isotropic BRDFs the existing analytic models perform well but in the case of anisotropic BRDFs a custom micro-facet normal distribution function was needed. They were able to produce a suitable distribution function by inverting the BRDF generator of Ashikmin *et al.* [2000] by fitting on measured data. Their resulting model was able to qualitatively reproduce the anisotropic BRDF measurements.

2.1 Appearance capture, modeling and fabrication

Another type of material whose appearance has attracted a lot of attention within the computer graphics community is hair fibers. The work by Kajiya and Kay [1989] on simulating the appearance of fur is considered the starting point. Marschner et al. [2003] presented a parametric model for simulating single scattering of light within hair volumes. Their hair model was able to better capture the characteristic highlights of dark hair than the model of Kajiya and Kay [1989]. A survey on hair appearance modeling can be found in the work of Ward et al. [2006]. More recently, researchers focused on hair capture [Paris et al. 2008; Jakob et al. 2009; Zinke et al. 2009] and efficient simulation of multiple scattering [Zinke et al. 2008] which is a prominent part of hair appearance, especially in the case of blond hair.

Other examples of appearance models include the work of Matusik et al. [2003] on designing a data driven model for interpolating the appearance of measured BRDFs, as well as the work of Marschner et al. [2005] for modeling the appearance of wood including anisotropic BRDF effects. Jakob et al. [2014] developed a method for consolidating the aggregate appearance from multiple thin layers with different BSDFs and media into a single BSDF model. Finally, Brady et al. [2014] presented a framework that can be used to synthesize analytical BRDF models able to describe real-world measured appearance, using genetic programming.

Appearance editing is another important procedure for the design of virtual scenes. Schmidt et al. [2014] performed a survey of the existing methods that facilitate artistic editing of appearance in virtual scenes. They categorized the existing methods according to the interface exposed to the user: direct, indirect, and goal-based. Direct user interfaces allow the user to directly edit the model parameters, they are easy to implement but usually not very intuitive. Indirect and goal-based interfaces

Related work

allow a more high-level specification of appearance and the underlying changes to the scene properties are performed by the system.

It is important to note, that even though many of the aforementioned appearance modeling and editing methods were able to capture and simulate the appearance of real-world materials, none of them addressed the challenges of computational appearance fabrication.

When simulating the appearance of real-world objects in a virtual scene, physically accurate appearance models are usually desirable but not required. On the other hand, when forward appearance models are used for fabrication, using a model that accurately simulates how light interacts with a material in the real world is essential.

In more recent years, researchers have started focusing on computational appearance fabrication. A notable example is the work of Weyrich et al. [2009b] on fabricating faceted surfaces that allow a customized angular reflectance distribution. Hullin et al. [2011] designed a prototype that can visualize a rough and anisotropic BRDF on a water surface by super-imposing multiple high frequency sine waves. Their method allows control of anisotropy by changing the frequency of the waves. Malzbender et al. [2012] combined lenses and blockers to achieve fabrication of spatially-varying BRDFs. Lan et al. [2013] presented a fabrication method that combines 3D-printing and milling for printing spatially-varying anisotropic BRDFs. Bermano et al. [2012] provided a method that takes as input a small set of up to four gray-scale images and designs a surface, that when fabricated, depicts one of these images (through self-shadowing) according to the illumination direction.

2.2 Translucency fabrication

Our goal with this application is to precisely control the appearance of homogeneous translucent materials in the real world by using continuous pigment mixtures with silicone. We allow appearance replication of measured translucent materials but also goal-based design and fabrication of novel translucent appearances, within the gamut of our pigments and silicone. In the remainder of this section we describe the work related to this application.

Color matching systems. Related to this application are the computerized commercial systems developed for color matching which have been successfully used in many industries (e.g., house paint, automotive etc.). They typically use a spectrophotometer to determine material reflection as a function of wavelength. Then they either determine the closest material, already in the database, or they output the combination of base pigments to obtain the best match. The most common systems include Sher-Color™ by Sherwin-Williams [Sherman and Simone 1989] and ChromaVision® by DuPont [Kelly 1987]. While some of these systems try to match sheen as well, the translucency of a material is never considered.

Forward appearance models. In the case of translucent materials, the main difference from the aforementioned surface scattering models (BSDFs) is that interactions are not only restricted to surfaces but can also happen anywhere within a participating medium. These interactions can be accurately simulated by using Monte Carlo methods to solve the radiative transfer equations such as volumetric path tracing. Computationally, this is an extremely expensive process even with the

Related work

current state-of-the-art efficiency improvements [Lafortune and Willems 1996; Walter *et al.* 2009; Jakob and Marschner 2012; Novák *et al.* 2012b; Novák *et al.* 2012a; Georgiev *et al.* 2013; Křivánek *et al.* 2014], especially in the case of optically thick and low absorbing media enclosed within a dielectric boundary.

The first practical model for approximating subsurface scattering in computer graphics was developed by Jensen *et al.* [2001]. An image-based measurement system was used to estimate model parameters for a variety of homogeneous translucent materials. This appearance model was later on extended by Donner and Jensen [2005]. In their work they were also able to approximate sub-surface scattering within and between thin layers. Their results include a rendering of a human face whose layered appearance was based on scattering parameters originating from the medical literature. Similarly, d'Eon and Irving [2011] introduced a BSSRDF model that accurately decouples single and multiple scattering. Their model as well as the model of Habel *et al.* [2013] approximates multiple scattered light from a linear light source which allows better approximation of light scattering from absorbing materials and thin translucent layers.

Capture. Hawkins *et al.* [2005] developed a method to estimate the phase function and albedo of aerosols. Narasimhan *et al.* [2006] created a method to estimate scattering parameters of translucent liquids by diluting low concentrations of a material in a clear solvent. Their method relied on the absence of multiple scattering from measurements and approximated the phase function with an analytic model [Henyey and Greenstein 1941]. More recently Gkioulekas *et al.* [2013] have presented an improved parameter estimation framework for liquids that could handle multiple scattering in measurements and provided scat-

tering parameters along with high-resolution tabulated phase functions.

Peers et al. [2006] developed a data-driven method for appearance capture and modeling of heterogeneous translucent materials by using a factored representation of the scattering function. The acquisition was performed using a projector-camera pair.

Weyrich et al. [2006] used a contact-based measurement device to estimate subsurface scattering parameters of human skin. For our measurements we use an improved version of this device to estimate reflectance and translucency in homogeneous translucent materials. Klehm et al. [2015] performed a survey on appearance capture of human faces and skin. An overview of the relevant forward appearance models was provided along with a categorization of the methods based on whether they could explicitly capture and model the inherent translucency of human skin.

Editing. There has been some recent work on authoring and editing the appearance of translucent materials [Xu *et al.* 2007; Song *et al.* 2009]. This work focused on editing and prescribing the subsurface scattering appearance in simulation, where inaccuracies in the forward appearance model and fabrication constraints can be completely ignored. We, on the other hand, are interested in editing the appearance of materials in the real world that can actually be fabricated by a mixture of our pigments. With our proposed editor we safely navigate the feasible gamut of translucent appearances based on the physical properties of our pigments and silicone.

Fabrication. In addition to the aforementioned appearance fabrication research, two closely related methods have addressed computational design and fabrication of materials with desired translucency [Hasan *et al.* 2010; Dong *et al.* 2010]. These approaches express the output as a material composed of discrete voxels or layers with known subsurface scattering properties. However, the materials for the layers/voxels are fixed. These methods can approximate materials with heterogeneous subsurface scattering as well. Nevertheless, they also have some disadvantages, in particular, materials with homogeneous subsurface scattering cannot be accurately reproduced since these methods rely on either voxel dithering or layering of discrete materials. Our proposed approach explores an orthogonal research direction that can be naturally combined with the discrete methods. Instead of manufacturing a compound multi-layer (or multi-voxel) material as in previous approaches, we fabricate materials using continuous pigment mixtures within a single homogeneous shape. Therefore, by default all our materials have homogeneous subsurface scattering properties.

2.3 Steganographic lens fabrication

Our goal with this application is to computationally design and fabricate lenses with steganographic properties. The main use-case is a user viewing a seemingly random looking pattern through the steganographic lens which optically unscrambles the pattern into the intended message or image. In the remainder of this section we describe the work related to this application.

Image obfuscation. Recent works in computer graphics encode hidden visual information into 2D images or 3D objects.

2.3 Steganographic lens fabrication

Different encoding and decoding schemes distinguish seemingly unrelated approaches that all share the same basic goal: encoding structured image information into potentially unstructured images or objects. Autostereograms [Tyler and Clarke 1990] encode a (depth) image that can be decoded by a viewer by controlling vergence, while camouflage and emerging images [Chu *et al.* 2010; Mitra *et al.* 2009] hide images decoded by the temporal delays of the human visual system. Classic “invisible ink” writing, using for instance lemon juice, is only revealed when exposed to a catalyst such as heat. More recent computational approaches use metallic inks to embed images that only appear under specular reflection [Hersch *et al.* 2003], or print with specialized inks that are only visible under UV light [Hersch *et al.* 2007]. Band Moiré images [Hersch and Chosson 2004] encode a hidden image in interfering Moiré patterns caused by superimposing transparent sheets. We also manufacture transparent surfaces, but instead, rely on a customized *refractive* lenslet-array to expose our hidden images and animations.

Patch matching. The shape of our lenslet-arrays is determined by matching regions between the source and target images. Many existing image feature and patch matching operations, commonly used in texture synthesis and image recognition [Lowe 1999; Pritchard and Heidrich 2003; Barnes *et al.* 2009; Barnes *et al.* 2010; Barnes *et al.* 2011; Barnes 2011], could be adapted to our task. However, our matching criteria is rather unique since we must constrain the smoothness of the resulting refractive lenslet-array. As such, we use a simple matching procedure specialized to our task.

3D displays. Simple auto-stereoscopic lenticular sheets [Lippmann 1908] “descramble” multiplexed left/right image pairs into distinct images seen from different viewpoints. Our approach can also be used to encode distinct images and, as such, can be seen as a special type of auto-stereoscopic display. We additionally encode target images completely unrelated to the structure of the source images by providing more degrees-of-freedom during lens optimization. With a directionally-dependent source image, our approach can also be interpreted as a way to optically warp one light field into another, opening up potential applications in light field display design [Gotoda 2010; Lanman *et al.* 2011; Wetzstein *et al.* 2011].

Hidden watermarks. Another family of methods generate digital watermarks, which can then be decoded using for example, lenticular sheets [Alasia 1976; Alasia 1998; Brosh and Wright 1994]. These methods require that the parameters of the lenticular sheets are known prior to the generation of the watermarks citeRenesse:2004:ODS.

Fabrication. Recent techniques can manufacture 3D objects, which, when lit in a controlled manner reveal hidden images. Niloy and Pauly [2009] designed 3D objects which, when lit from different directions, cast distinct shadow images, while Baran *et al.* [2012] constructed multilayer attenuators which cast several colored images depending on lighting conditions. Yue *et al.* [2012] allowed users to manually arrange refractive pixel-like sticks in a grid arrangement to generate projective pixel art. Relief images [Alexa and Matusik 2010] encoded two unique images in the diffuse shading of a height-field lit from two directions. Fuchs *et al.* [2008]

2.3 Steganographic lens fabrication

presented prototypes of various passive reflectance field displays where the outgoing light-field is dependent on the incident light direction. Related works [Finckh *et al.* 2010; Papas *et al.* 2011] optimized the geometry of refractive surfaces to deform incident light into desired caustic patterns. More recently, Schwartzburg *et al.* [2014] have presented a method for fabricating continuous caustic-generating lenses using an optimal transport map. Their method enforces the surface continuity as a hard constraint, whereas the method of Papas *et al.* treats continuity as a soft constraint. The benefit of the method of Schwartzburg *et al.* is the production of smooth lenses without discontinuities, which are easier to fabricate with the use of milling machines. A drawback of the method of Schwartzburg *et al.*, when considered for our steganographic application, is that the shape of the lens is highly correlated to the resulting caustic image. A warped version of the resulting caustic can be seen on the surface of the lens. In contrast, the faceted lenses generated by the method of Papas *et al.* breaks this correlation between the surface of the lens and the resulting caustic image, a property highly desirable for our application of steganographic lenses. Our goal is to fabricate steganographic lenses which is, in a sense, the optical dual of such caustic lens generating approaches. We apply a similar manufacturing process and optimization framework as Papas *et al.* [2011], but our optical system operates in reverse: we focus outgoing reflected or emitted light from a source image, using a lenslet-array, to form an image (or stereo pair, or animation) directly on an observer's eye(s).

2.4 Rendering granular media

Our goal with this application is to lay the foundation for future research on computational appearance fabrication of granular media. We identify as a necessary first step toward this goal, the design of a fast and accurate forward appearance model for simulating the appearance of granular media. This crucial component is currently missing from our computational appearance fabrication framework and is a necessary building block for future research in granular media fabrication. In the remainder of this section we describe the work related to our proposed forward appearance model.

Existing forward models. There is very little work on accurately simulating the appearance of granular media in the field of computer graphics. For this class of materials there has been work mainly on rendering by Moon et al. [2007] but their method is still prohibitively slow. Our goal is to create a fast multi-scale forward appearance model, inspired by Li et al.'s [2005] hybrid model that is able to render granular media accurately for the first few media interactions and then switch to a fast and accurate diffusion approximation. This will be a potentially useful tool for future research with the goal of controlling and fabricating the appearance of granular media.

Our work is a direct extension of the model developed by Meng et al. [2015]. They proposed a multi-scale modeling and rendering method that adapts to the structure of scattered light at different scales. They rely on explicit path tracing (EPT) of the individual grains only at the finest scale, and—by decoupling individual grains from their arrangement—they develop a modular approach for simulating longer-scale light transport. They model light interactions within and across grains as

separate processes and leverage this decomposition to derive parameters for classical radiative transport, including standard volumetric path tracing (VPT). This volumetric approximation can describe the large scale transport due to many grain interactions. They require only a one-time precomputation per exemplar grain, which they reuse for arbitrary aggregate shapes and a continuum of different packing rates and scales of grains.

Moon et al. [2007] precomputed scattering functions on spherical shells inside a random medium, accounting for the aggregate effect of transport within and between individual objects that make up the medium. A key difference in our work is that unlike Moon et al., we do not require any scene dependent global precomputations. Similarly to Meng et al. [2015] we only require a light-weight precomputation for a single constituent grain, and not an agglomeration of grains. This allows us to decouple shape, arrangement, scale, and packing of grains. Avoiding costly precomputations when changing these physical properties is essential when designing a backward appearance model and very useful when designing a fast forward appearance model.

Level-of-detail and prefiltering. Our approach combines several distinct representations and rendering methods to model the appearance and light transport within granular materials. Such multi-scale, level-of-detail approaches have a long history in graphics. Luebke et al. [2002] have given a thorough treatment of this early work, whereas Bruneton and Neyret [2012] have provided a recent survey of non-linear prefiltering techniques for smoothly transitioning between multi-scale representations of appearance.

Rendering granular materials has conceptual similarity to work in hair [Moon and Marschner 2006; Zinke and Weber 2006;

Related work

Zinke *et al.* 2008] and cloth rendering [Schröder *et al.* 2011; Zhao *et al.* 2013] where individual fibers are visible and therefore need to be modeled explicitly, but costly and smooth large-scale transport is approximated.

Methods outside graphics. Understanding the optical and heat transfer properties of densely packed media is of great importance for fields outside of graphics, including thermal engineering, atmospheric sciences, and nuclear reactor physics. Accurate solutions can be obtained using a full wave approach [Foldy 1945] and solving Maxwell’s equations [Durant *et al.* 2007], but are computationally infeasible and not necessary for graphics-related problems.

Since exact solutions considering each granule are impractical, a popular alternative is to make the “homogeneous phase approximation” [Randrianalisoa and Baillis 2009], where the densely packed medium (with statistically dependent scattering) is approximated using the standard RTE, valid at low volume fractions or for long-scale transport. Randrianalisoa and Baillis [2010] proposed a data-driven procedure to fit such “effective RTE properties” to the results of a Monte Carlo random walk through a discrete granular medium. This requires a separate, scene-dependent simulation, akin to Moon *et al.*, which we avoid.

Singh and Kaviany [1992] modeled dependent radiative transfer by pre-computing a scattering function with positional offsets for a large, smooth, dielectric spherical particle. Replacing the standard phase function in the Discrete Ordinates Method [Chandrasekar 1960] with this new scattering function allows them to account for intra-grain transport effects.

Donovan *et. al.* [2003] used a two-step process where they

sampled chord lengths for transport between grains from a statistical distribution and computed intra-grain transport by instantiating grain geometry. Our teleportation model is very similar, but uses a statistical representation for intra-grain transport.

2.5 Summary

To summarize, our findings show that there is a tremendous amount of work within the computer graphics community on transferring the appearance of real-world materials to the virtual world. Existing appearance models and editing methods often have the required properties for simulating the underlying material class. In many cases, though, they are not directly usable for computational appearance fabrication. The main requirements for fabrication-friendly forward appearance models are: they can correctly capture the real-world appearance of their material class, they have low dimensional parameters that can be mapped to a physical state, and finally are computationally efficient to evaluate, since they usually end up in an inner loop of an optimization procedure.

We also observe that the current state of the art with respect to computational appearance fabrication is at different stages for each class of materials. For the class of translucent materials the field is fairly advanced. Existing methods can already reproduce a limited gamut of translucent appearance in the real world with the use of 3D printers. For specular surfaces, computational appearance fabrication has also been demonstrated, and we provide an additional use-case of goal-based design of steganographic lenses. In the case of granular media, there is not much research available. There is some work on modeling the appearance of this material class, but there are

Related work

no suitable appearance models available which are suitable for computational granular media appearance fabrication. We provide the first step towards this future goal starting with a fast and accurate forward appearance model for modeling the appearance of granular media.

CHAPTER

3

Translucency fabrication

In this chapter we present a full instance of our framework (Figure 3.1) tailored for capture, design and physical reproduction of homogeneous translucent materials with desired color and translucency.

Our fabrication process allows us to precisely control pigment concentrations that can be suspended in a base material. Our backward appearance model is able to determine pigment concentrations that best reproduce the appearance of a given target material. In order to achieve this task, we need to train our forward appearance model such that it can be used to predict the appearance of a physical state. We first fabricate a collection of material samples composed of known mixtures of the available pigments with the base material. Their appearance is acquired using a custom-built measurement device. The same device is then used to measure the reflectance and translucency of a target material. Based on the forward model predictions which were trained on this database of mappings from pigment concentrations to reflectance profiles, we design a backward appearance model. This backward model uses an optimization process to compute a concentration of pigments that, when




Figure 3.1: An instance of our framework depicting an example of translucency fabrication. A physical state is translated into the virtual world with the use of our forward model. The appearance is then edited and passed as input to the backward model. Finally, the backward model uses an optimization to provide pigment mixtures that, when fabricated, will produce a new physical state with the desired appearance.

fabricated, will reproduce the target appearance. We demonstrate the practicality of our method by reproducing a variety of different translucent materials. We also present an appearance editing tool that allows the user to explore the appearance gamut defined by a fixed set of pigments and base silicone.

It is important to describe the three main components of the physical state at this point. First, the physical parameter of the object that can be controlled is the mixture composition of pigments and base silicone. In our case we use homogeneous mixtures of six pigments and base silicone at desired concentrations. We can additionally control the shape by using molds, but the computational aspect of geometry fabrication is outside the scope of our work.

The geometric illumination setup under which we simultaneously match appearance is divided into two cases. The first case

is when the surface of the object is uniformly illuminated by a distant light source (reflectance) and the second case when a beam of light is entering from a position on the surface and travels through the material (translucency). In addition, the light source emits light in five narrow spectral bands. The observer is defined as a gray-scale camera placed at some predefined distance on top of the sample with uniform sensitivity over the visible spectrum.

Due to the bidirectional nature of light transport and the fact that we can programatically control the emitters and the camera, we can also describe this setup as the equivalent of a pentachromat observer, viewing a material illuminated with an ideal "white" light source that emits radiant flux uniformly over the visible spectrum.

While matching appearance with such a sensitive observer allows for the fabricated appearance of an object to match for human observers under spectrally varying illumination, this drives the optimization in sub-optimal solutions when the appearance cannot be matched well in parts of the spectrum, to which the human visual system is not sensitive.

In practice, a more common scenario is to have a standard human observer and a fixed illuminant (e.g. Sunlight, LEDs and fluorescent lights). Under this fixed observer and fixed illuminant assumption, we can use perceptual metrics and thus focus appearance matching in parts of the spectrum where humans are more sensitive—effectively allowing our method to match appearance within a wider gamut.



Figure 3.2: *A side by side comparison of real translucent materials (right) next to their silicone replicas (left) fabricated using our method.*

3.1 Introduction

Most of the materials in our man-made environment are colored by dyes, pigments, or other colorants suspended in a scattering medium. Paints, plastics, papers, textiles, stained glass, ceramic glazes, candy—nearly all surfaces that are not metallic or completely transparent fall under this description. Many natural materials are also well approximated as colored scattering media—skin, leaves, flowers, foods—in which the colorants are naturally occurring. Because of the ubiquity of colored scattering materials, the technology of predicting and controlling their color is very mature, as epitomized by systems that automatically mix paints to match a given sample.

But color is not the only attribute of a colored scattering medium; pigmented media are, by their very nature, translucent. Some materials are so dense (wall paint, for instance) that the translucency can be ignored at macroscopic scales, but for others it is subtly (“opaque” plastic, skin) or obviously (translucent plastics, stained glass) part of the appearance. Translucency is a more complex phenomenon than diffuse color, and

currently the appearance of such materials is normally controlled by trial and error.

The goal of this chapter is to create the fundamental technology of controlling translucency as precisely as color can already be controlled, including an accurate forward model that can predict the appearance of translucent materials, automatically matching existing real or virtual materials, and synthetically adjusting mixtures with feedback about translucent appearance. Over the last dozen years, the field of computer graphics has developed an increasingly mature understanding of how to simulate [Jensen *et al.* 2001; d'Eon and Irving 2011], measure [Hawkins *et al.* 2005; Dorsey *et al.* 2008; Weyrich *et al.* 2009a], and manipulate [Xu *et al.* 2007; Song *et al.* 2009] the appearance of translucent objects in rendered scenes, and we believe this technology is becoming mature enough to be applied to the more demanding application of manipulating materials in the real world.

The resulting methods are directly useful in design applications involving pigmented translucent materials, such as industrial design of consumer products where appearance is important, or design of dental materials, prostheses, or animatronics that should match a given person's body—appearance in all these examples is currently matched manually by trial and error. Together with the ability to accurately render translucent materials, this will make translucent materials an integral part of realistic product pre-visualization, letting the user see on screen the exact appearance the manufactured product will have.

In addition, as multi-material 3D printers become sufficiently capable, the same fundamental technology of controlling the appearance of scattering materials can ultimately enable 3D printing of nearly arbitrary materials with precise control over appearance. Our work complements existing work on

spatial mixtures of 3D printed materials [Dong *et al.* 2010; Hasan *et al.* 2010] by examining how to control the properties of individual materials. Printers that can control both material properties and spatial arrangement precisely will enable *physical* appearance prototyping as well as the direct manufacture of products whose appearance is important.

In this research we have developed a system to generate a pigment mixture to match the subsurface scattering of a homogeneous target material. Our system is analogous to computerized paint matching systems that measure the spectral reflectance of a target material and then reproduce it by combining pigments in a scattering or clear base material. We use a similar concept, but instead of treating the material as having only diffuse surface reflection, we model it as a volume with corresponding subsurface scattering effects. This approach extends our system beyond just color matching to cover translucent materials with significant subsurface scattering in a way that can be easily integrated into existing systems. Figure 3.2 shows a sample of different materials that can be reproduced using our system. Our methods are applicable to any process where pigments are used to control the appearance of a material.

Our process for matching subsurface scattering of a given target material has the following steps. First, we manufacture a collection of samples with different concentrations of pigments mixed into a base material. We use a custom-built measurement system (Section 3.4) to acquire multi-spectral reflection profiles for each material sample. This allows us train our forward model and establish a mapping from pigment concentrations to reflection profiles that capture subsurface scattering properties. Using the same measurement device we acquire multi-spectral reflection profiles for a given target material. We design a backward model with an iterative optimization

procedure that computes the pigment concentrations required to reproduce the target material accurately (Section 3.5). We demonstrate the whole process by reproducing a number of organic and inorganic materials (Sections 3.8 and 3.9). In addition, we describe an appearance editing tool that enables interactive exploration of the range of translucent appearances that is achievable with a given set of pigments.

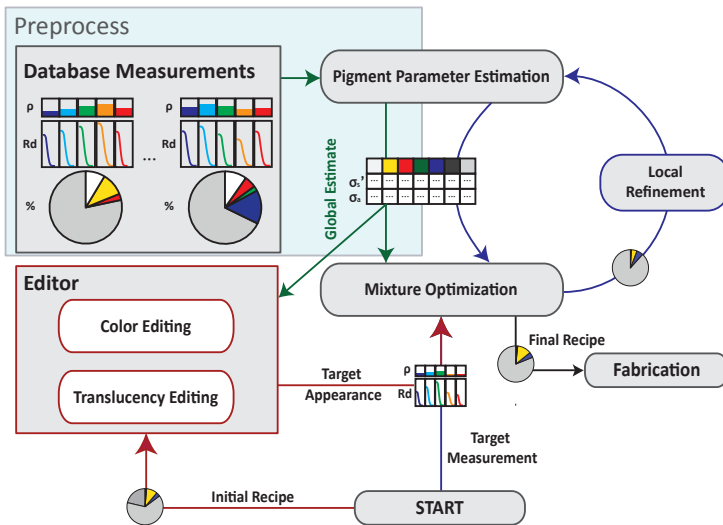


Figure 3.3: Overview of our system.

3.2 Method overview

At a high level, the goal of our system is to provide a recipe of how to mix pigments with a base material (clear silicone in our system) in order to reproduce a measured target material. Our process can be divided into five main stages (illustrated in Figure 3.3): database and target measurement, pigment parameter estimation, concentration estimation, and, finally, fabrication.

In a preprocess we fabricate and measure various mixtures of the available pigments with base material (Database Measurement). This collection of measurements becomes our appearance database which we then use to estimate a *global* set of pigment parameters for predicting the subsurface scattering appearance of a silicone mixture (Pigment Parameter Estimation).

To create a silicone replica, we first measure the diffuse reflectance and translucency of a target material (Target Measurement). We then perform an optimization that will estimate pigment concentrations for the target (Mixture Optimization), assuming global pigment parameters (Pigment Parameter Estimation). This process is refined by an iterative local pigment parameter estimation step, with the goal of enabling better results in regions of the domain where the forward model, and consequently the global set of pigment parameters, is no longer a good approximation.

In Section 3.3 we describe our chosen forward appearance model, which balances our needs for both accuracy and performance. We use a single, custom-built measurement device, which we describe in Section 3.4, to perform both the database and target measurements. In Section 3.5, we describe both the pigment parameter estimation (Section 3.5.3) and mixture optimization stages (Section 3.5.4) used for computing a pigment concentration recipe that will match a measured target

3.3 The forward appearance model

appearance. In Section 3.6 we describe an improvement to this method (Local Pigment Parameter Estimation) which accounts for inaccuracies in the forward model approximation. With this system in place, we enable users not only to replicate target materials, but also to edit and pre-visualize (Section 3.7) the appearance of a desired material before fabrication. The output of the appearance matching and appearance editing process is a recipe, which we fabricate by mixing pigments with silicone (Section 3.8). We compare the fabricated results of our system against a set of target translucent materials in Section 3.9. Finally, we propose an extension of our method (Section 3.10) that increases the appearance matching gamut in the case where we have a fixed illuminant and a human observer.

3.3 The forward appearance model

Before we can fabricate a replica, we need a computational model to predict the appearance of scattering materials which we can later invert to derive the scattering parameters necessary for fabrication. According to the theory of scattering media, a homogeneous material can be described by a phase function and two parameters, the absorption coefficient σ_a and the scattering coefficient σ_s , or equivalently by the extinction coefficient $\sigma_t = \sigma_a + \sigma_s$ and albedo $\alpha = \sigma_s/\sigma_t$. If any two of these parameters are known the other two can be computed. In highly scattering materials, the flow of light can be well modeled with a diffusion equation, which leads to approximate analytical models that are useful to describe translucent materials. In such materials, one can replace the scattering coefficient with the reduced scattering coefficient σ'_s , and then treat the phase function as isotropic (see Jensen and Buhler [2002] and references therein for details). In this thesis we will always use the reduced pa-

parameters σ'_s , σ'_t , and α' but will omit the customary primes for notational convenience.

Scattering Profile. From the many available diffusion models, we chose quantized diffusion [d'Eon and Irving 2011] to analytically express the subsurface reflectance profile since it remains relatively accurate even in the case of moderately absorbing materials. For the purposes of our task, the chosen diffusion model simply returns an analytic reflectance profile between two surface points x and y :

$$\frac{dL_r^\lambda(x)}{d\Phi_i^\lambda(y)} = R_d^{QD}(\alpha^\lambda, \sigma_t^\lambda, d; r), \quad (3.1)$$

as a function of their distance $r = \|x - y\|$, thickness d , and the reduced albedo α and reduced extinction coefficient σ_t per wavelength band λ . This also depends on η . In this thesis we always use a measured silicone index of refraction value of 1.41. The internal details of the diffusion model are largely unimportant for the rest of our pipeline, hence R_d can be treated as a black box.

Diffuse Reflectance. We also require a computational model for the diffuse reflectance ρ . We evaluated the accuracy of various analytical and numerical models for diffuse reflectance, but found that none matched Monte Carlo simulations well enough (see Appendix A for our detailed comparison analysis). We therefore created a dense tabulation of diffuse reflectance values from brute force Monte Carlo [Wang *et al.* 1995] simulations:

$$\rho(\alpha, \sigma_t, d) = 2\pi \int_0^\infty \frac{R_d^{MC}(\alpha, \sigma_t, d; r)}{\pi} r \, dr. \quad (3.2)$$

For our tabulation we parametrize ρ according to reduced albedo (α) and optical thickness ($\sigma_t d$, where d is the measured thickness of the sample) and store the values in a 2D table for lookup and interpolation.

Please note that our reflectance term assumes that the multiple scattering process (including internal reflection) is isotropic, and that ρ is the reflectance for incoming light inside the surface to outgoing light inside the surface.

3.4 Measurement setup

Using our fabrication process, we would like to reproduce a wide range of organic and inorganic materials that exhibit sub-surface scattering. To allow measuring of target materials, that cannot be diluted or modified in other ways, we use a non-invasive contact-based measurement device inspired by, but extended from, the design proposed by Weyrich et al. [2006]. We use this device to measure both the database silicone samples as well as the target materials we want to replicate.

Our setup incorporates a single housing to perform two distinct types of measurements. The first type of measurement is used to extract the **diffuse reflectance**, $\bar{\rho}^\lambda$, of the target material, and the second type is used to measure the **bulk scattering profile**, \bar{R}_d^λ , of the target material. In Figure 3.4 we illustrate a cross-section view of our enclosed device, which we fabricated using a 3D printer. The device has a height of 11 cm and a diameter of 25 cm. The sample is placed at the bottom of the measurement device, in contact with the base opening circle (A), and we place a monochromatic QImaging Retiga-2000R camera into the circular opening (B) at the top of the device. This non-invasive method only requires that a small circular patch of the target




Figure 3.4: *Top: An illustration of the setup used to measure the appearance of a target sample. The measurement device performs a diffuse reflectance measurement and as well as a bulk scattering profile measurement without any moving parts. Bottom Left: An example gray scale HDR image of a diffusion profile measurement. Bottom Right: An example gray scale HDR image of a diffuse reflectance measurement.*

material is fairly flat and can be brought in contact with our measurement device for a few seconds or minutes, depending on how dark the material is. All measurements are performed for five distinct wavelength-bands (indexed by λ) using two identical sets of five color LEDs, one for reflectance and one for profiles. This allows us to capture reflectance and profile

3.4 Measurement setup

measurements in sequence without having to move the sample or reconfigure the measurement device.

The camera used for our measurements was an actively cooled, monochromatic QImaging Retiga-2000R which captures 12-bit RAW images. After extensive testing we found that these RAW files have a linear radiance response and variations due to vignetting are negligible for the regions we measure. To create HDR images we sum up the unclipped pixels from each image and divide by the total exposure time. This approach effectively applies low weights to photographs with low exposure times, which is desirable since extremely low exposure photographs are prone to noise. For these measurements, an absolute radiance value is not required.

Diffuse Reflectance. The diffuse reflectance measurement is performed using an array of five fiber optic cables located in the top right edge of the device (C), aiming towards the center of the sample (A) as seen in Figure 3.4. The other ends of these fiber optic cables are mounted to an LED holder with five color LEDs. We place a 5 mm diameter opal glass diffuser at the outgoing end of the fiber optic cable to ensure a constant angular intensity distribution on the sample. Figure 3.4 shows a gray scale HDR diffuse reflectance capture (bottom-right) with one LED turned on.

In reality, we cannot easily observe $\bar{\rho}^\lambda$ directly. Instead, we observe values of reflected radiance, \bar{L}_s^λ , off the sample. However, we design the geometric configuration so that we can derive $\bar{\rho}^\lambda$ from our observed measurement under some reasonable assumptions. In particular, if we assume that the material is a homogeneous medium with a smooth Fresnel boundary and that single scattering is negligible, we have the following

general expression for the observed radiance:

$$\bar{L}_s^\lambda(\vec{\omega}_o) = 2\pi \int_0^\infty \int_\Omega \bar{L}_i^\lambda(\vec{\omega}_i) (\vec{\omega}_i \cdot \vec{\mathbf{n}}) \frac{S_d^\lambda(r, \vec{\omega}_i, \vec{\omega}_o)}{\pi} r \, d\vec{\omega}_i \, dr, \quad (3.3)$$

where $\vec{\omega}_i$ and $\vec{\omega}_o$ are the incident and outgoing directions respectively, $\vec{\mathbf{n}}$ is the surface normal, and $\bar{S}_d^\lambda(r, \vec{\omega}_i, \vec{\omega}_o) = F_t(\vec{\omega}_o) \bar{R}_d^\lambda(r) F_t(\vec{\omega}_i)$ is the BSSRDF with Fresnel reshaping.

In our case, we observe the sample from directly above, $\vec{\omega}_o = 0^\circ$, and we illuminate the sample from a single direction $\vec{\omega}_i$ at 45° to normal incidence (to avoid imaging direct specular reflection of the light). We measure the per-wavelength-band radiance of the sample, $\bar{L}_s^\lambda(0^\circ)$, by averaging an approximately 1 cm^2 square patch centered at (A). Assuming that the incident direction is constant at contributing regions, the observed radiance therefore simplifies to:

$$\bar{L}_s^\lambda(0^\circ) \approx \frac{I_i^\lambda}{t^2} \cos(45^\circ) F_t(0^\circ) F_t(45^\circ) 2\pi \int_0^\infty \frac{\bar{R}_d^\lambda(r)}{\pi} r \, dr. \quad (3.4)$$

Where I_i^λ is the intensity of the light source, located at distance t . To estimate the intensity we also perform a one-time measurement, $\bar{L}_c^\lambda(0^\circ)$, for a gray diffuse calibration target placed at (A). By assuming that the calibration target is perfectly Lambertian we have:

$$I_i^\lambda = \frac{t^2 \bar{L}_c^\lambda(0^\circ)}{\bar{\rho}_c^\lambda(0^\circ, 45^\circ) \cos(45^\circ)}, \quad (3.5)$$

where $\bar{\rho}_c^\lambda(0^\circ, 45^\circ)$ is the reported reflectance of the diffuse calibration target. Though we allocate a small warm-up time for stability within a measurement, we cannot assume that the intensity of the LEDs will be consistent across measurements over multiple days. To correct for fluctuations in LED intensity,

we additionally use four small reflectance standard patches (D) which are always visible in our measurements. We then scale I_i^λ by the average radiance ratio of these four patches in the material sample measurement and the calibration target measurement.

By combining Equations (3.4) and (3.5) and rearranging terms we obtain a simple expression for use as our model reflectance ρ (Equation 3.2):

$$\bar{\rho}^\lambda = 2\pi \int_0^\infty \frac{\bar{R}_d^\lambda(r)}{\pi} r \, dr \approx \frac{\bar{\rho}_c^\lambda(0^\circ, 45^\circ)}{F_t(0^\circ)F_t(45^\circ)} \frac{\bar{L}_s^\lambda(0^\circ)}{\bar{L}_c^\lambda(0^\circ)}. \quad (3.6)$$

Bulk Scattering Profile. Our contact measurement device also contains a second set of five LEDs and fiber optic cables originating from the left side (E).¹ These are in contact with the material sample and illuminate it at a location (F) which is not directly visible by the camera. The horizontal distance between the center of the 1 mm diameter fiber optic cable end and the first measurable location on the sample is 0.9 mm.

Light from the fiber optic cables propagates through the scattering material and into the field of view of the camera. We show an example gray scale HDR capture with one LED turned on in Figure 3.4 (bottom-left). We capture five such HDR images, one for each LED. We extract the horizontal scan-line, vertically aligned with the currently active light source, and use this as the bulk, per-wavelength-band, diffusion profile measurement, \bar{R}_d^λ . The measured profile has an inherent arbitrary scale factor, due to the unknown intensity of the LEDs. Additionally by

¹We measured the spectral distributions of all LEDs using a Photo Research SpectraScan PR 730 spectrometer to ensure there is negligible variability between corresponding LEDs across the two sets.

assuming that for our profile measurements the Fresnel transmission terms are not spatially varying, then the Fresnel terms can also be folded into this scale factor. We will later on show how knowledge of this unknown factor is not necessary for our purposes.

3.5 Designing the backward model

Once a target material is measured, our goal is to reproduce the material by mixing one or more pigments into the base material. To compute a recipe, or a vector containing a concentration for each available pigment, our forward model must be able to predict the appearance (both the reflectance and the scattering profile) that will result from any given set of concentrations. This mapping from recipe to appearance is then inverted in our backward model with an optimization process to determine the best recipe for matching the appearance of a target material.

Our key task is to derive a mapping from concentrations c_p for $p = 1, \dots, n_p$ to the observed appearance characteristics ρ^λ and R_d^λ for $\lambda = 1, \dots, n_\lambda$. We build this model from a number of example materials that are fabricated and measured ahead of time, forming what we call a *database* of appearance measurements (Section 3.5.1). Our approach leverages the approximate forward model from the previous section to intelligently interpolate among the database samples and to provide starting points for nonlinear optimization of pigment concentrations.

Given the measurements of \bar{R}_d^λ and $\bar{\rho}^\lambda$ for any particular material, the parameters of that material (α, σ_t) can be estimated by fitting to the forward model, seeking a match in the (relative)

diffusion profile and the (absolute) total reflectance:

$$(\alpha^\lambda, \sigma_t^\lambda) = \underset{\alpha, \sigma_t}{\operatorname{argmin}} F(R_d(\alpha, \sigma_t, d), \rho(\alpha, \sigma_t, d), \bar{R}_d^\lambda, \bar{\rho}^\lambda) \quad (3.7)$$

$$F(R_d, \rho, R'_d, \rho') = \left[E(R_d, R'_d) + (\rho - \rho')^2 \right]$$

where E is a profile difference measure we will describe in Section 3.5.2.

In principle, this fitting approach could be used to determine the material parameters of each of the training examples, from which the properties of each individual pigment could be derived. The parameters of a target material could then be determined in a second fit and used to find the pigment concentrations required. However, the diffusion approximation is not accurate enough to directly achieve a visual match using this simple approach, i.e. the model parameters are not linearly related with pigment concentrations. In the following subsections we describe our approach for finding a set of global pigment parameters, linearly related to pigment concentrations, that best predicts appearance.

3.5.1 Measurement database selection

The first step in estimating the pigment parameters is designing the input set, which we call the **Measurement Database**. Our goal is to use the methodology and machinery described in Section 3.4 to acquire the per-wavelength-band diffuse reflectance measurement, $\bar{\rho}^\lambda$, the per-wavelength-band bulk scattering profile, \bar{R}_d^λ , for a set of samples with known pigment concentrations, and then estimate α^λ and σ_t^λ for each pigment and for the base silicone using the methods of the previous section.

The main challenge, when the input set of pigments contains highly absorbing entries, is the design of database samples that will not violate the assumptions of diffusion theory which will be used to estimate their parameters. The two main assumptions for diffusion theory to hold are $\sigma_a^\lambda \ll \sigma_t^\lambda$ and that the multiple anisotropic scattering in the material can be well approximated using an approximately equivalent isotropic material with reduced scattering parameters.

We use a total of 6 pigments. With the exception of white and yellow pigments, the remaining (red, green, blue, and black) pigments are highly absorbing. We created for each pigment a set of database entries, which we call a dilution set. Each such dilution set consists of fabricated silicone samples with varying concentrations of that pigment, always mixed with white pigment at a concentration of 0.05%. For each dilution set, we used the minimum amount of pigment such that both the color and the profile differentiate enough from the appearance of white pigment at 0.05% concentration. The maximum concentration was chosen such that the smallest measurable profile is at least 3 mm long. We also include a dilution set with varying concentrations of only white pigment. This ensures that the scattering parameters of white pigment can be distinguished from the parameters of the base material. The physical size of our database samples is $10 \times 10 \times (2-4)$ cm, achieving a minimum optical thickness of about 10 along their minimum dimension. A set of plots showing the concentrations used for our dilution sets, along with the forward model fits, can be found in Appendix A.

3.5.2 Fitting the forward model to a single material

The simplest fitting operation is to fit the forward model to the measured appearance data \bar{R}_d^λ and $\bar{\rho}^\lambda$ for a single wavelength band of a single material.

Initial Guess. To find a starting point for the optimization, we use our diffuse reflectance lookup table to find α such that $\rho(\alpha, \sigma_t, d) = \bar{\rho}^\lambda$ initially assuming that the sample is semi-infinite ($d = \infty$). To obtain an initial guess for σ_t , we perform an asymptotic simplification of the quantized diffusion model, valid for $r \gg 1/\sigma_t$:

$$R_d(\alpha, \sigma_t, d; r) \approx k \frac{e^{-r\sqrt{\sigma_a/D}}}{r}, \quad (3.8)$$

where D is the diffusion coefficient and k is a constant. As shown in Figure 3.5, this asymptotic approximation states that, for large enough r , we can expect a plot of $\log(rR_d(r))$ against r to be a straight line with slope $-\sqrt{\sigma_a/D}$. Hence, by fitting a line to $\log(r\bar{R}_d^\lambda)$, we obtain $\sqrt{\sigma_a/D}$, from which we compute σ_t using the currently estimated value of α . We repeat these two steps (α and σ_t estimation), but for the following iterations we no longer assume a semi-infinite sample but instead we use the measured thickness of the sample, d . This process usually converges after 3–5 iterations.

Non-linear Optimization. Starting from these estimated values for σ_t and α , we use the Levenberg-Marquardt algorithm to compute the minimum of (3.7). To compute the difference between two scattering profiles we use the metric:

$$E(R_d, R'_d) = \frac{1}{r_1 - r_0} \int_{r_0}^{r_1} \left[(R_d(r)/\mu)^{\frac{1}{3}} - (R'_d(r)/\mu')^{\frac{1}{3}} \right]^2 dr \quad (3.9)$$




Figure 3.5: A log plot of a measured bulk scattering profile as a function of distance. Only the grey highlighted region is used for fitting with the diffusion model.

where we divide the profiles by their mean values (μ and μ') to account for the unknown intensity of the light source in the diffusion profile measurement. The interval $[r_0, r_1]$ is a range of distances over which the model is expected to fit well. This range is determined by shrinking the interval until a line fits within a given tolerance, and can be manually overridden to avoid any glitches in the measured profiles. An example from our measurements highlighting this range is shown in Figure 3.5.

As a final step in the single-material fitting process, we summarize the residual error of each database sample using a confidence:

$$z_m^\lambda = \min \left(\frac{\mu_d^\lambda}{d_m^\lambda - d_{75\%}^\lambda + \mu_d^\lambda}, 1 \right), \quad (3.10)$$

where d_m^λ is the residual (the minimum value of (3.7)) in wavelength band λ for the m^{th} sample, $d_{75\%}^\lambda$ is the 75th percentile

3.5 Designing the backward model

residual over the whole database for this wavelength band, and μ_d^λ is the median error for this wavelength band over the entire database. We use this confidence later as a weight in fitting pigment parameters.

This process of fitting to a single profile produces material parameters that correspond to the observed appearance, but because the model is only approximate, the best-fit parameters may not be close to the true parameters of the material. This is particularly problematic in the case of anisotropically-scattering materials for which the forward model is less accurate. To obtain more meaningful results we subsequently fit larger collections of samples at once, as described in the next section.

3.5.3 Global pigment parameter estimation

Once we have separately estimated the parameters of all the database samples (each of which consists of a known mixture of one or more pigments with the base material), we have material parameters for each sample, which describe that sample's appearance. However these pigment parameters are not linearly related to pigment concentrations as radiative transport theory predicts (see Figure 3.6). As a result, interpolating between the independently fit parameters can lead to poor prediction results.

To get more reliable predictions, we instead use the results of independent fitting to initialize a larger fitting problem that finds material parameters for each pigment that are globally consistent with all samples in the database, under the radiative transport theory assumption of a linear relationship between pigment concentrations and the parameters of the mixture. This linear relationship can be succinctly expressed using a matrix Σ ,

Translucency fabrication

which contains the properties of all samples in all wavelength bands:

$$\Sigma = \begin{bmatrix} \sigma_1 \\ \vdots \\ \sigma_{n_p} \end{bmatrix} \quad (3.11)$$

where

$$\sigma_p = \begin{bmatrix} \sigma_{s,p}^1 & \sigma_{t,p}^1 & \cdots & \sigma_{s,p}^{n_\lambda} & \sigma_{t,p}^{n_\lambda} \end{bmatrix} \quad (3.12)$$

are the material parameters of the p^{th} pigment, and a matrix C , which contains the known concentrations of all pigments in all database samples; entry c_{mp} is the concentration of pigment p in sample m . The $n_p \times 2n_\lambda$ matrix Σ has a row for each pigment (including the base material) and a column for each parameter in each wavelength band. The $n_m \times n_p$ matrix C has a row for each material in the database and a column for each pigment. With these definitions, the matrix $M = C\Sigma$ contains the material parameters of every material in the database.

To find globally consistent material parameters for the pigments, we fit the same objective function (3.7) independently for each wavelength band, except summed over all materials:

$$\sigma_{\text{global}}^\lambda = \underset{\Sigma^\lambda}{\operatorname{argmin}} \sum_{m=1}^{n_m} F(R_d(\alpha_m^\lambda, \sigma_{t,m}^\lambda, d), \rho(\alpha_m^\lambda, \sigma_{t,m}^\lambda, d), \bar{R}_d^\lambda, \bar{\rho}^\lambda) \quad (3.13)$$

where

$$\sigma_{s,m}^\lambda = \sum_{p=1}^{n_p} c_{mp} \sigma_{s,p}^\lambda, \quad \sigma_{t,m}^\lambda = \sum_{p=1}^{n_p} c_{mp} \sigma_{t,p}^\lambda, \quad \text{and} \quad \alpha_m^\lambda = \frac{\sigma_{s,m}^\lambda}{\sigma_{t,m}^\lambda}.$$

To ensure convergence to the global minimum, in this initial phase we estimate the properties of the pigments one at a time,

3.5 Designing the backward model

using the one- and two-pigment dilution sets described in Section 3.5.1. We begin with the white dilution set, optimizing (3.13), summing only over the materials in that set, for the properties of the base material and the white pigment. For each color dilution set, we then similarly optimize for the properties of the color pigment, holding the white and base materials fixed. We initialize these optimizations by fitting a line to the scattering parameters (from the previous step) of all entries in a dilution set.

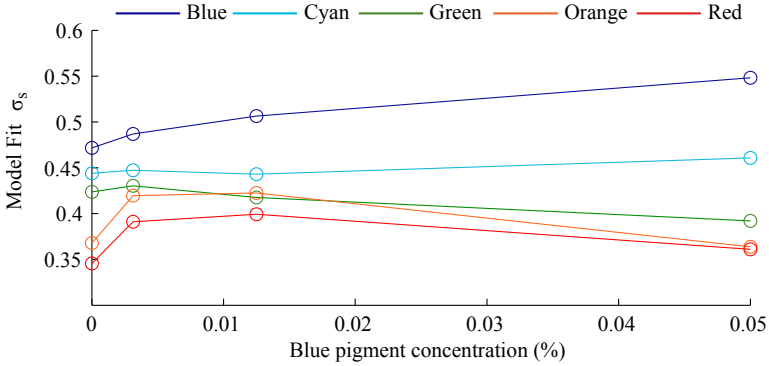


Figure 3.6: Per-wavelength-band scattering coefficients for the blue pigment dilution set as a function of concentrations. Note how σ_s is non-linear with respect to concentration.

3.5.4 Mixture optimization

Once we have parameters for each of the available pigments, we can compute a recipe to match a target material using the

same tools. Given the measured diffuse reflectance $\bar{\rho}^\lambda$ and scattering profile \bar{R}_d^λ for the target material, we first use the fitting process of Section 3.5.2 to estimate the $2n_\lambda$ -vector of scattering parameters $\hat{\Sigma} = [\hat{\sigma}_s^1 \cdots \hat{\sigma}_t^{n_\lambda}]^T$ for the target mixture. Then we solve the linear system $\mathbf{c}_i^T \Sigma = \hat{\Sigma}$ to get a p -vector of pigment concentrations $\mathbf{c}_i = [\hat{c}_1 \cdots \hat{c}_{n_p}]^T$. Using \mathbf{c}_i as an initial guess, we optimize the predicted appearance to the target:

$$\mathbf{c} = \underset{c_1, \dots, c_{n_p}}{\operatorname{argmin}} \sum_{\lambda=1}^{n_\lambda} F(R_d(\alpha^\lambda, \sigma_t^\lambda, d), \rho(\alpha^\lambda, \sigma_t^\lambda, d), \hat{R}_d^\lambda, \hat{\rho}^\lambda) \quad (3.14)$$

where α^λ and σ_t^λ are defined by

$$\sigma_s^\lambda = \sum_{p=1}^{n_p} c_p \sigma_{s,p}^\lambda \quad \text{and} \quad \sigma_t^\lambda = \sum_{p=1}^{n_p} c_p \sigma_{t,p}^\lambda.$$

The resulting vector \mathbf{c} is the recipe to replicate the appearance of the target material using the given pigment set.

3.6 Local pigment parameter estimation

The mapping of measurements to pigment concentrations as it is described until now assumes that the forward model can globally fit the entire database with a single set of pigment parameters. However, even though our forward model is relatively accurate, we cannot expect it to work well over the entire parameter range when using a single global set of pigment parameters. In particular, the forward model will be less accurate for low optical thickness, low albedo, or anisotropically scattering materials. This is the case, for instance, for Silicone Mixture 2 shown in Figure 3.9 which by design is within the gamut of our pigments, but is not replicated accurately.

3.6 Local pigment parameter estimation

To overcome inaccuracies in the forward model, we introduce a local refinement strategy to find a set of pigment parameters that locally fits the samples in the database that are most similar to the target. To accomplish this, we apply a higher weight on the neighbors, with respect to the pigment concentration, when estimating the “effective” pigment scattering parameters for finding the recipe.

This can be performed in an iterative procedure which interleaves the Parameter Estimation and Mixture Optimization stages with the difference that row weights are used in the Parameter Estimation stage to bias the error to be lower for neighboring mixtures already in the database. In practice, we found that the dot product of the normalized pigment concentration vectors between the currently predicted pigment concentrations and a database entry (excluding base silicone concentration) provides consistent results.

The procedure for the local pigment parameter estimation parallels the global optimization in Section 3.5.3; however, the objective function (3.13) is replaced by:

$$\sigma_{\text{local}}^\lambda(\mathbf{c}) = \underset{\tilde{\Sigma}^\lambda}{\operatorname{argmin}} \sum_{m=1}^{n_m} \quad (3.15)$$

$$\left[(w^\lambda(\mathbf{c}, \mathbf{c}_m) + k_{\text{reg}}) F(R_d(\tilde{\alpha}_m^\lambda, \tilde{\sigma}_{t,m}^\lambda, d), \rho(\tilde{\alpha}_m^\lambda, \tilde{\sigma}_{t,m}^\lambda, d), \bar{R}_d^\lambda, \bar{\rho}^\lambda) \right],$$

where the material parameters $\tilde{\alpha}$ and $\tilde{\sigma}_t$ are the local ones (derived from the optimization variable $\tilde{\Sigma}^\lambda$). The k_{reg} parameter regularizes the problem so that even when some pigments are not used by the nearby samples, their parameters stay close to the global parameters. The regularization parameter k_{reg} is set just high enough (around 10^{-4} relative to a unit maximum) to stabilize the optimization, while still low enough not to affect the quality of the local fit.

The weights w^λ are set to

$$w^\lambda(\mathbf{c}, \mathbf{c}_m) = z_m^\lambda D_{\text{mixture}}(\mathbf{c}, \mathbf{c}_m) \quad (3.16)$$

where

$$D_{\text{mixture}}(\mathbf{c}_1, \mathbf{c}_2) = \text{normalize}(\mathbf{c}_1) \cdot \text{normalize}(\mathbf{c}_2). \quad (3.17)$$

These weights cause the pigment estimation stage to find pigments parameters that fit well to database materials similar in composition to the target mixture, resulting in better prediction of appearance for the optimized recipe. After a new recipe has been found, we update the weights, and we repeatedly solve re-weighted systems until convergence or for a maximum of 5 iterations.

3.7 Appearance editing

In addition to matching the appearance of measured targets, we also developed an authoring tool which allows the user to pre-visualize recipes before fabrication and edit the desired color and translucency, while staying within the gamut imposed by the minimum and maximum possible pigment concentrations. Our approach provides an intuitive editing work-flow by decoupling edits to the color and the translucency of a desired material.

Overview. Our editing process starts with an initial recipe and sample thickness provided by the user. We then provide intuitive browsing control to fine tune the color followed by translucency. We took inspiration for our system from the “variations” control interface in Adobe Photoshop[®]. Screenshots from our editor are shown in Figure 3.7.




Figure 3.7: *Top: Screen-shot from the color selection phase of our editor. The variations along the CIELab primaries are shown above and below the current pigment recipe. Bottom: A screen-shot from the translucency selection phase. The user can navigate along the T,S,R coordinate system. The blacked out images show directions in translucency that are limited by our pigments' gamut.*

For editing color, our interface displays the diffuse reflectance color of the current recipe and six color-variations, estimated by moving in both the positive and negative directions of the three CIELab primaries. When the user selects a variation, the current recipe gets updated and the process repeats until the user has reached the desired color.

For editing translucency, our interface renders a synthetically fabricated recipe under step edge illumination to visualize the reflectance profile (similar to the captured photographs in Figure 3.9). We again show the current recipe and translucency-variations which can be selected to modify the shape of the diffusion profile without affecting the overall diffuse reflectance.

We detail the steps necessary to accomplish this below.

3.7.1 Color reproduction

To display the CIELab color predictions of a recipe, we require a mapping, $(L, a, b) = P_{\text{Lab}}(\mathbf{c}, d)$, to convert pigment concentra-

tions \mathbf{c} to CIE Lab color values, given the depth d of the sample. We use globally linear estimates of pigment scattering parameters from our database, and then linearly mix these parameters according to \mathbf{c} to get the resulting scattering and absorption coefficients of this mixture. These coefficients are then converted to a 5D reflectance vector using the interpolation method described in Section 3.5. We then convert our 5D spectral color values to spectral reflectance distributions by training a set of “eigen-spectral reflectance functions” from spectral reflectance measurements of our database [Park *et al.* 2007]. Once we have an approximate spectral distribution for the color, we convert to CIEXYZ and subsequently to CIE Lab and sRGB. The CIE Lab is used for computing the variation distances and sRGB for displaying the predicted reflectance values to the user.

3.7.2 Color editing

Our goal is to find variations of our recipes with a given distance in CIE Lab. For this, we first compute the partial derivatives $J_{\mathbf{c} \rightarrow \text{Lab}}$ of $P_{\text{Lab}}(\mathbf{c})$ with respect to pigment concentrations \mathbf{c} , using finite differences. This matrix represents the predicted color change with respect to pigment concentration changes. In order to find the pigment concentration changes that allows movement along the CIE Lab primaries, we first perform a singular value decomposition on $J_{\mathbf{c} \rightarrow \text{Lab}}$: $[U, \Sigma, V] = \text{svd}(J_{\mathbf{c} \rightarrow \text{Lab}})$ and compute the pseudo inverse of $J_{\mathbf{c} \rightarrow \text{Lab}}$ as $V_{\text{Lab}} = V \Sigma^{-1} U^T$.

Each column of the resulting 6×3 matrix V_{Lab} is now the derivative of the pigment vector with respect to the CIE Lab primaries. The difference in pigment concentrations, $\Delta_{\mathbf{c}}$, to achieve a desired change in CIE Lab, Δ_{Lab} , is given by $\Delta_{\mathbf{c}} = V_{\text{Lab}} \Delta_{\text{Lab}}$.

3.7.3 Translucency editing

Given the scattering and absorption coefficients for each wavelength band, we render a column image using the diffusion model described in Section 3.5. Each pixel representing the radiance of our 5 wavelength bands is transformed into CIE Lab using the process described in Section 3.7.1.

For navigating the translucency space without affecting the color, we reuse a part of the 6×6 matrix V from the singular value decomposition of the color derivatives matrix $J_{\mathbf{c} \rightarrow \text{Lab}}$. We split V into two sub-matrices $[V_c V_t] = V$, where V_c are the first 3 columns of V and V_t the remaining 3 columns. While the pigment vectors in V_c affect the color of the final recipe, the pigment vectors in V_t do not and can be used to influence just the translucency.

Unfortunately, the pigment vectors in V_t do not correspond to meaningful directions. One intuitive direction for controlling translucency is increasing the average reduced extinction coefficient $\langle \sigma_t^\lambda \rangle$, over all wavelength bands, by scaling the current pigment concentrations with the desired $s = \langle \sigma_t^\lambda \rangle$. For a semi-infinite slab, this directly corresponds to scaling all profiles by s , but for finite-slabs this can inadvertently affect the diffuse reflectance color. We can find the closest direction T in the space defined by V_t by solving the linear system $V_t T = C$. Now, given T , we find the orthonormal directions S and R that should not change color, nor the average reduced extinction coefficient $\langle \sigma_t^\lambda \rangle$. These vectors T , S , and R are used to define the variations in translucency that are presented to the user.

When editing translucency, it is often impossible to move in some directions, most often the S and R directions, because of the constraints on how much pigment can be used. For steps

that exceed the limits, we black out that direction to indicate to the user that they have reached the edge of the gamut.

3.8 Fabrication

Fabricating the generated recipes with a high degree of precision is critical to the correct evaluation of our method. The main challenges in fabrication are: ensuring that the correct amount of each pigment is added, avoiding air and other impurities, and finally ensuring that there is at least one side on the sample which appears near-specular. In this section we provide a brief description of the materials and machinery used for fabrication and then we describe the hierarchical dilution process we employ for improving the concentration accuracy, reducing waste and streamlining the fabrication process.

Silicone and Pigments. We use SortaClear40 Translucent, Addition Cure silicone rubber by Smooth-On. This silicone rubber cures at room temperature with a shrinkage of less than 0.01 %. The two-component silicone rubber requires a catalyst, mixed in with a ratio 1:10, to activate the curing process. The curing is roughly 24 hours. The silicone is completely cured after 7 days.

For manipulating the appearance of the silicone we use silicone pigments which are used for coloring silicone rubber. The pigment colors we use are: White (Pantone White C), Yellow (RAL 1018), Red (Pantone Red C), Green (Pantone 3292), Blue (Pantone 2757C), Black (Pantone Black C). These pigments are mainly absorbing except for White, Yellow and Red, for which scattering is significant.

Hierarchical Dilution. For fabricating the generated recipes as faithful as possible, we employ a hierarchical dilution scheme.

For each pigment we initially produce 1 kg master batches of 5% pigment concentration. This is the maximum ratio of pigment that still allows the silicone to cure. To achieve a target concentration, we then dilute some quantity of this master batch, in an iterative fashion, with base silicone. We start with some of the 5% mixture and repeatedly mix the current dilution with an equal amount of base silicone to halve the pigment concentration. Once the concentration is roughly twice the target concentration for our recipe, we mix with the exact ratio of base silicone needed to achieve the desired target concentration.

This dilution tree structure is automatically generated by a simple script using a bottom-up approach. Intermediate dilutions that are needed by many target recipes are merged and created only once. We ignore concentrations of pigments with absolute concentration less than 10^{-8} of the total sample weight, or less than 3 g for a typical 300 g sample. Our process also accounts for catalyst that will be added at the very end for curing the final samples.

We use a digital balance with accuracy of ± 0.05 g when mixing dilutions. To ensure high accuracy, we always use a minimum of 10 g of both dilution and base silicon when mixing. Though limiting the minimum weight increases the number of steps, we obtain a relative accuracy lower bound of 1% at each dilution step. This relative error is reduced with each dilution step by the mixing ratio.

Mix Preparation Process. During the fabrication process, preserving homogeneity is very important. Once all ingredi-

ents are added, we stir the mixture for several minutes until it is homogeneous. Since this process accumulates air, the mixture is placed in a vacuum chamber for about 10 minutes to further improve homogeneity. We then pour the mixture into a mold, constructed using acrylic plates. We use acrylic plates so the sample has smooth, near-specular surfaces for better compliance with our model assumptions and more angularly uniform reflectance for our $45^\circ/0^\circ$ measurement setup, due to the lack of micro-surface roughness. Finally when the samples cure, they are measured using the method described in Section 3.4.

3.9 Results

We validated our forward model, measurement setup, and optimization procedures first in simulation and also by fabricating physical replicas using silicone. We describe these here in turn.

3.9.1 Reproduction of simulated examples

To ensure that our method can properly measure and replicate real-world translucent materials, we first validated our entire pipeline in simulation by fitting on synthetic images generated using brute-force volumetric path tracing with the Mitsuba renderer [Jakob 2010]. We designed an analogous database of measurements and tested our method on synthetic targets. We created 6 synthetic pigments with prescribed RGB material parameters and isotropic phase functions. The database samples and targets are created by linearly blending the pigment parameters according to their concentrations. For our synthetic tests, we recreate the geometry of our physical setup: the camera and light sources are placed at the same positions and orientations, all the database samples had a thickness of 30 scene units and

our diffusion profile light source was modeled as a disk light touching the sample.

Our path traced measurements (which mimic our real measurements) omit the beginning of the profile, so to enable rendered validations in a realistic shadow edge configuration (which requires access to the entire profile) we performed an unconstrained fit of our forward model to these path traced measurements. Given the fitted profiles we use the same method as in Section 3.7 to render the shadow edge illumination seen in Figure 3.10.² The differences between the targets and the replicas in this synthetic test are barely noticeable. In fact, over this set of 14 synthetic targets the average relative reflectance error was 0.10% and the mean relative reduced mean free path error was 6.44%. This confirms that our measurement and optimization procedures are able to accurately replicate unknown translucent materials.

In Figure 3.11 we perform a similar synthetic evaluation, but using a simpler optimization strategy and error metric to demonstrate that our more involved approach is necessary to obtain accurate results. The only difference compared to Figure 3.10 is that the replicas are obtained by minimizing the MSE between the estimated/measured parameters for the target (reduced scattering coefficient and absorption coefficient) and the linearly blended global parameters of pigments of the tested recipe. This can be modeled as a simple linear system, similar to our initialization step in the Mixture Optimization section. The results show that, fitting on parameters directly does not perform as well as using our appearance distance term. This is due to two reasons: a) the material parameters recovered by a single target measurement are highly under-constrained;

²We confirm in Appendix A that these fitted profiles are indistinguishable from the 1D Monte Carlo profiles.

and b) our appearance distance term can leverage “appearance metamerism” (several recipes and pigment parameters can lead to indistinguishable translucent appearance) to tolerate ill conditioning in the mapping from optical properties to appearance.

3.9.2 Reproduction of real-world examples

In addition to validating on synthetic measurements, we also demonstrate the quality of our pipeline on real-world examples. We measured and replicated several homogeneous real-world materials such as strawberry drink, fabric softener, full-fat milk, low-fat milk, soap, and white chocolate. We computed the corresponding mixing recipes using our iterative optimization procedure and fabricated the samples in various shapes. Each replica has the same shape and thickness, d , as its corresponding target. The computation time for estimating a recipe is about 2 minutes with the global method and 30 minutes for the local method. The local method training database also includes the replicas fabricated using the global method.

For our real-world examples, we capture photographs of the original targets and the replicas using a Nikon D800 camera. We construct HDR images from 12-bit RAW images with 6 different exposure times. The RAW LDR images are processed with DCRAW [2004] and white balanced using a white reference card. We recover and account for the camera response curve using the technique of Mitsunaga and Nayar [1999]. This results in white balanced linear sRGB images which are then converted to HDRs using the aforementioned method. At the very end, we tone-map the HDR image to an 8-bit sRGB image by applying the sRGB gamma curve. We compare the target

material against the fabricated replicas under two different illumination configurations.

In Figure 3.2 we showcase side-by-side visual comparisons of various liquid and solid translucent targets next to their replicas, in their natural form under top and side illumination.

Figure 3.9 shows side-by-side shadow edge comparisons with an illuminant that approximates a color temperature of 4700K. For each triplet, the target is in the middle, and we show the results of the global and local methods to the left and right of the target respectively. Visually, even our global method is able to match both the translucency and reflectance of many of these materials quite well. For the global method, the appearance distance (defined in Equation (3.7)) averaged over all 9 targets is 0.00644, with a standard deviation of 0.01142. Our local refinement usually performs better, and sometimes significantly better, than predicting the appearance of the replicas using the global parameters (Table 3.1). In fact, the mean distance after local refinement decreases to 0.00276, with a standard deviation of 0.00213. The improvement is especially noticeable for the dark Silicone Mixture 2 seen in Figure 3.9, and some

Table 3.1: Global estimate of reduced scattering and absorption coefficients for pigments and base silicone used in our fabrication process. Units are in $1/\text{mm}$.

	Reduced scattering coefficient (σ_s)					Absorption coefficient (σ_a)				
	blue	cyan	green	orange	red	blue	cyan	green	orange	red
white	848.9	847.6	783.6	669.4	718.3	0.0	0.0	0.0	0.0	0.0
yellow	92.6	74.0	78.8	64.0	55.8	30.1	1.8	0.5	0.2	0.1
red	0.0	0.4	0.3	620.3	96.7	515.4	839.8	999.6	163.9	9.2
green	7.1	0.0	10.0	0.1	0.0	19.5	9.4	15.8	164.0	264.4
blue	92.3	9.1	0.0	0.1	0.2	41.6	90.2	181.1	1080.9	1045.9
black	25.3	21.8	22.9	25.9	26.6	321.6	331.4	330.1	323.7	325.4
base	0.078	0.053	0.061	0.054	0.021	0.002	0.002	0.001	0.001	0.001

Table 3.2: Concentration percentages prediction generated using our local method for the 9 targets used. The last 2 rows show the actual concentration of the fabricated silicone mixtures used as targets.

	white	yellow	red	green	blue	black
silicone mixture 1	0.20590%	0.20897%	0.00345%	0.00140%	0.00735%	0.00000%
blue fabric softener	0.16384%	0.00000%	0.00037%	0.00170%	0.01407%	0.00000%
low-fat milk	0.13572%	0.01291%	0.00000%	0.00012%	0.00000%	0.00000%
red soap	0.41249%	0.19489%	0.01169%	0.00247%	0.00000%	0.00000%
strawberry drink	0.25791%	0.06423%	0.00467%	0.00000%	0.00062%	0.00000%
white chocolate	0.32361%	0.24976%	0.00261%	0.00284%	0.00000%	0.00000%
full-fat milk	0.35230%	0.02702%	0.00000%	0.00000%	0.00000%	0.00012%
mocha drink	0.21805%	0.37069%	0.00953%	0.00543%	0.00124%	0.00000%
silicone mixture 2	0.00000%	0.19234%	0.00142%	0.00000%	0.00547%	0.00702%
silicone mixture 1 (actual)	0.20000%	0.20000%	0.00300%	0.00000%	0.00750%	0.00000%
silicone mixture 2 (actual)	0.00000%	0.20000%	0.00300%	0.00000%	0.00750%	0.00000%

improvement is also visible for the Blue Fabric Softener and Strawberry Drink samples.

We provide plots of the simulated, predicted, and measured diffuse reflectance and scattering profiles for our targets in Appendix A. In general, we observed a good prediction for materials with high reflectance and optical thickness. With a decrease of the reflectance and optical thickness, we notice an increase in matching error. This is not surprising, since diffusion theory is not a very good approximation for very dark materials with low optical thickness.

To isolate the performance of our optimization, we tested our pipeline with materials that are in theory exactly reproducible by our system. To ensure that our target material is within the gamut of reproducible materials, we fabricated target samples out of silicone mixed with pigments that are available to our fabrication process. We then measured these samples and used the diffuse reflectance and scattering profiles as input to our optimization process (these samples are not included in

3.10 Perceptual extension




Figure 3.8: *On the left we present the spectral reflectance values of three target materials whose appearance cannot be reproduced for our pentachromat observer. On the right we present spectral reflectance measurements of our pigments mixed with white. Both figures have the spectral intensities of our LEDs overlayed. The grayed out regions indicate spectral bands that are not sampled well with our LEDs. We observe that the three targets exhibit a significantly positive gradient in the region between 540 and 580 nm, which is not captured by our setup. In addition we can see on the right that our current set of pigments does not have a member that exhibits a significant positive gradient in this region, indicating that these targets are outside of the gamut of our pigments for a pentachromat observer.*

our measurement database). Table 3.2 shows the ground truth pigment concentrations of two samples and the estimated concentrations using our optimization process. In Figure 3.12, we compare the profiles of the original measured sample, the predicted appearance, and the measured fabricated replica. Note that the appearance of the replica closely matches the target, even though the pigment concentrations are different. This is an example of an “appearance metamer” where we cannot perceive the appearance difference between two different recipes. Such a replica would never be created when using a simpler parameter optimization as shown in Figure 3.11.

3.10 Perceptual extension

In this section we describe an extension of our method for improving color and translucency matching when the observer and illuminant are specified. The main idea is to match how humans perceive color and translucency, under a known illuminant, instead of fitting on the raw measurement data. In terms of our computational appearance fabrication framework this means that we fix the illumination conditions and observer to one of the standard human observers.

In Figure 3.8 (left), we present the spectral reflectance values of three target materials that are not suitable for our method. On the right we provide spectral reflectance measurements of our pigments mixed with white. Both figures have the spectral intensities of our LEDs overlaid. In the plot on the left (middle gray region) we observe three targets with a strong positive gradient in their spectral reflectance distribution, whereas, in the plots on the right we observe that our current set of pigments does not have a member that exhibits such a significant positive gradient in the corresponding spectral region. This is a strong indication that an exact reconstruction of the spectral reflectance distribution for these targets is not possible using our current set of pigments.

To alleviate this issue and also extend the gamut and accuracy of our method, we exploit the limitations of the human visual system. In particular we change the appearance distance metric (Equation (3.14)) which currently equally weights all spectral bands, to a perceptual distance which allows exploiting metamers. That makes use of the fact that two spectral distribution functions that are different, can actually be perceived as identical by humans. Our basic modification to Equation 3.14

is replacing the distance F with the perceptual distance \tilde{F} :

$$\tilde{c} = \underset{c_1, \dots, c_{n_p}}{\operatorname{argmin}} \tilde{F}(R_d(\tilde{\alpha}, \tilde{\sigma}_t, d), \rho(\tilde{\alpha}, \tilde{\sigma}_t, d), \tilde{R}_d, \tilde{\rho}), \quad (3.18)$$

where the raw measurements (\tilde{R}_d^λ and $\tilde{\rho}^\lambda$) are now projected into the perceptual color space CIELAB (denoted as \tilde{R}_d and $\tilde{\rho}$). As converting to this perceptual color space requires a reference illuminant, it is important to note that the resulting appearance of the replicas matches well only when they are illuminated with the specific illuminant we optimize for. We define this new perceptual distance (\tilde{F}) as a weighted sum of two terms:

$$\tilde{F}(R_d, \rho, \tilde{R}_d, \tilde{\rho}) = \left[w_r DE_{00}(\rho, \tilde{\rho})^2 + w_t \tilde{E}(R_d, \tilde{R}_d) \right], \quad (3.19)$$

a perceptual reflectance distance term on the left hand side of the addition, and a perceptual translucency distance (or profile shape) term on the right. We found that a relative weight $w_t/w_r = 25$ works well for our dataset.

Perceptual reflectance distance. The first step is to convert the 5D reflectance measurements ($\tilde{\rho}^\lambda$) into spectral reflectance distribution estimates. This process is described in the second half of Section 3.7.1. For replicas created using our pigments this method works well, but for arbitrary targets we found that the recovery of the spectral reflectance distribution function using spectral eigenvectors trained on our pigment database was not a good basis. Instead, we discovered that the method works better when the target reflectance distribution function is measured directly by using a spectrometer, or by using the spectral eigenvectors from Park et al. [2007]. In our case we used spectral measurements of our 4700K halogen light source as the reference illuminant. Now, given the spectral radiance

estimates of the target sample and the replica, we can convert the raw measurements ($\bar{\rho}^\lambda$) in CIELAB (denoted as $\bar{\rho}$), and estimate their color distance using the perceptual metric CIEDE2000 [Sharma *et al.* 2005].

Perceptual translucency distance. The next step is to convert Equation (3.9) to a perceptual distance metric for translucency. To do this we first perform an unconstrained fitting of the forward model on our reflectance partial profile measurements (\bar{R}_d^λ) and then use the model values to extrapolate the entire 5D profiles. The main reason for this optimization is that we miss the first few millimeters of the profiles, since they are occluded due to the measurement device wall during our measurement. We then convert the extrapolated 5D profiles to CIELab (denoted as \tilde{R}_d) and compute the integral of the CIEDE2000 squared distance for all measured profile locations:

$$\tilde{E}(R_d, \tilde{R}_d) = \frac{1}{r_1 - r_0} \int_{r_0}^{r_1} DE_{00}(R_d(r), \tilde{R}_d(r))^2 dr \quad (3.20)$$

Preliminary results. We implemented this new perceptual distance metric in our computational appearance framework. By converting our 5D multi-spectral measurements in this lower-dimensional 3D perceptual appearance space where the illuminant is our halogen lamp and the observer is defined as the standard human observer. We were able to compute mixtures for our pigments and silicone that, when viewed under the specified illuminant, provide a qualitatively better appearance match for the previously deemed out-of-gamut materials (carrot juice and purple fabric softener). We also observed that this appearance qualitatively matched better under sunlight which exhibits a similar spectral emission as the target illuminant. In contrast, the replicas produced with our perceptual

extension appeared different from the targets when we used a drastically different illuminant such as our green LED.

3.11 Discussion and future work

We presented a method to physically replicate homogeneous translucent materials by mixing pigments into a substrate. By using local fits to a diffusion model, we intelligently interpolate between measurements of known materials to predict the appearance resulting from novel mixtures. The method only considers homogeneous materials, and its accuracy drops off for highly forward scattering and highly absorbing materials, with low optical thickness, where diffusion theory is less accurate. We believe that there is room for improvement on the reproduction accuracy of our method.

It is important to discuss the work on fabricating spatially varying translucency using a discrete set of homogeneous voxels [Hasan *et al.* 2010] or layers [Tong *et al.* 2005]. These methods are a form of volumetric dithering where the primitive can be a homogeneous voxel or layer. They can achieve a desired heterogeneous translucent appearance by computationally specifying an arrangement of these primitives. These methods only control the shape and location of the primitives and have no control over other physical primitive properties such as albedo and opaqueness. Thus, the achievable range of different appearances is limited by the minimum size (or thickness) and the fact that the remaining physical properties of the primitives are fixed. Our method is able to solve an orthogonal problem: how to fabricate a desired homogeneous translucent appearance by using continuous pigment mixtures. Our method can be used to computationally design and fabricate the appearance of a complete homogeneous translucent object but can

also be potentially used to fabricate the primitives used by the aforementioned heterogeneous methods.

There are various options for expanding the reproducible appearance gamut of our method. One such option would be to add additional pigments with useful spectral features not present in our current set. With minimal changes to our method we can improve the spectral resolution of our measurement device by using our discrete set of LEDs (with the addition of an LED at the empty spectral region of 560 nm) and a multi-spectral camera.

There are many other directions in which our method can be further developed. Improved light transport models and/or better interpolation with larger databases can extend the gamut towards darker and less optically dense materials, ultimately covering the whole range from perfectly clear to opaque materials. We demonstrated the application of a perceptual distance metric that can be used to increase the appearance matching gamut of our method, for a human observer and a known illuminant. An interesting next step would be to find a new perceptual distance metric that can match appearance (possibly independently) under different illuminants such as sunlight, fluorescent light, and UV light.

By itself, the ability to control the appearance of homogeneous materials has implications for many industries where pigmented materials are used: plastics, food, prosthetics, and even dentistry, where the critical matching of appearance between natural teeth and artificial resins is still done painstakingly by eye. Translucent appearance matching is also important for fabricating more complex materials. Previous work on fabricating translucent appearance has been limited to spatial combinations of fixed materials. Our new method, with its ability to continuously tune material parameters, opens up

3.11 Discussion and future work

the possibility to more powerful new methods for inhomogeneous materials. The latter optimize both the spatial mixture of materials and also the properties of the individual materials themselves. Ultimately, these technologies will lead to future machinery that can automatically replicate the appearance of almost any translucent material.




Figure 3.9: Measured target and fabricated replica HDR pairs for a variety of materials. The mixing recipes for the fabricated replicas of each target (middle) were generated using the global method (left) and the local method(right) .



Figure 3.10: Synthetic target and replica comparisons for the global method. For each sample, we render the target (left), followed by the replica (right) under shadow edge illumination. Both renderings use a QD fit on ground-truth MC results.



Figure 3.11: Synthetic target and replica comparisons when using a simpler optimization that fits on parameters directly. Fitting on parameters directly does not perform as well as using our appearance distance term (cf. Figure 3.10).

3.11 Discussion and future work

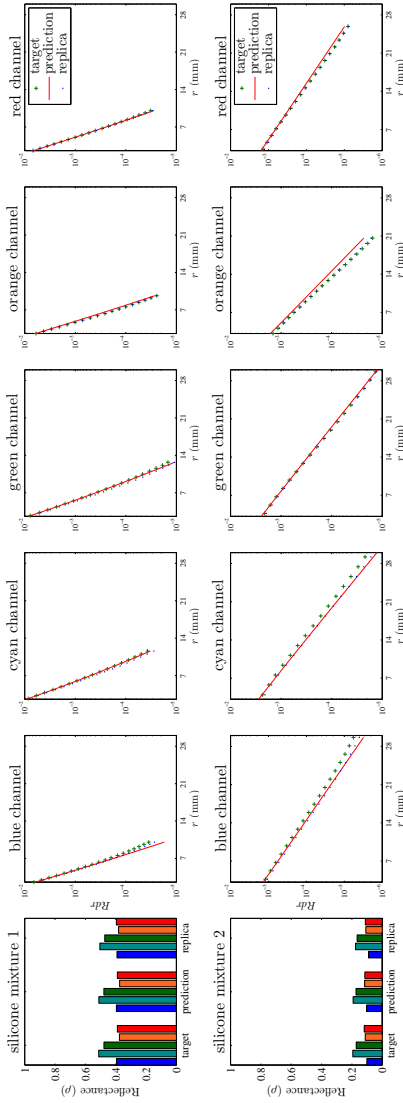


Figure 3.12: We evaluate the appearance prediction ability of our system, using the local method, by comparing the target measurements (a) with the model prediction (b) and finally the fabricated replica measurements (c). The first row shows comparisons of our system's prediction ability for Silicone Mixture 1, whereas the second row shows comparisons for Silicone Mixture 2. Both target samples are in gamut and their recipes can be found in Table 3.2. Note that even though Silicone Sample 2 has low optical thickness and reflectance values, the fabricated replica is a good match.

Translucency fabrication

CHAPTER

4

Steganographic lens fabrication

In this chapter we present another instance of our framework which provides an automatic approach to design and manufacture steganographic lenses. These lenses are composed of refractive lenslet arrays that reveal hidden images when placed over potentially unstructured printed or displayed source images. Our backward model determines the refractive geometry (or shape) of these surfaces by formulating and efficiently solving an inverse light transport problem, taking into account additional constraints imposed by our fabrication methods. The input of our backward model is a mapping of patches on the lens to regions on the source image, generated by our appearance editing (patch-matching) step. We fabricate several variants on the basic steganographic lens idea, including using a single source image to encode several hidden images. The latter are only revealed when the lens is placed at prescribed orientations on the source image or viewed from different angles. We also present an important special case, the *universal steganographic lens*, that forms an injection mapping from the lens surface to the source image grid, allowing it to be used

Steganographic lens fabrication

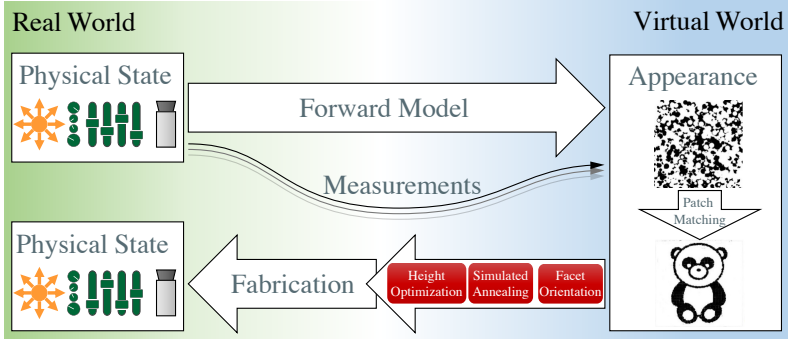


Figure 4.1: An instance of our framework, used for source-optimized steganographic lens fabrication. Please note our appearance editing step (right) which allows goal-based appearance edits by using user provided source and target image pairs. The appearance editing step provides a mapping to the backward model. Through the depicted optimization steps the backward model computes the appropriate geometry of the resulting lens, according to the desired fabrication method.

with arbitrary source images. We use this type of lens to decode hidden animation sequences.

An instance of our framework suitable for steganographic lens fabrication is shown in Figure 4.1. The illumination conditions are more relaxed in this application and they do not require rigorous radiometric and photometric calibration. Ideally we have some low frequency ambient illumination and/or a diffusely emissive display is used to depict the source image. For this application the physical parameter that we can computationally control is the lens geometry. We define the observer as a pinhole camera aimed at the center of the lens and positioned at a fixed distance.

By assuming that the surface of the lens is a perfect refractor,

the forward model is trivial and is mainly based on Fermat's principle and Snell's law of refraction.

One interesting aspect of this application is the design of an appearance editor which allows high-level appearance edits. In the case of source optimized lenses, the input of the editor is a combination of scrambled source images and decoded target images, whereas the output is a mapping of regions on the source image to patches on the lens surface. In the case of the universal steganographic lens, the user specifies some scrambling related parameters and the editor automatically outputs a mapping.

The goal of our backward model is to compute a valid lens geometry which guides the observer's view frustum from locations on the front side of the lens, to regions on the source image in the back of the lens, according to the mapping produced by the appearance editing step. We show variations of the backward model that accommodate the limitations of two common fabrication methods: 3D printing and milling.

4.1 Introduction

Steganographic techniques, from simple hidden message decoders to invisible inks and complex watermarking schemes, have led to active areas of research and have been applied in a wide variety of fields. Searching for and finding structure in unexpected places is also a fun and insightful process. Some common day examples of this expedition include the pursuits of a child armed with only a magnifying glass and their imagination, to a family huddled around a table, completing a jigsaw puzzle.

Steganographic lens fabrication




Figure 4.2: A sketch of other potential applications for our steganographic lenses. Left to right: architectural-scale steganographic lenses reflect hidden images to pedestrians as they walk past buildings; a teenager learns about dinosaurs by placing his steganographic lens over a museum exhibit pedestal; students find clever ways to exchange encoded images which are decoded with steganographic lenses; a \$100 bill has an embedded steganographic lens that, when placed over a scrambled verification pattern, confirms the validity of the bill; cinephiles view a movie billboard through specialized steganographic lens side-walk installations, revealing interesting facts about the movie.

We leverage and incite this sense of wonder, encountered when inanimate objects suddenly convey a unexpected message or reveal surprising behavior, by combining the ideas of steganography, hands-on physical user manipulation, and structure from unstructured patterns. We design and construct several different types of *steganographic lenses*, using a custom computational procedure, capable of warping both structured and unstructured image sequences into unexpected target images. Our steganographic lenses are composed of lenslets that, when placed atop an image/video and viewed from prescribed locations, warp the image through refraction to form the desired images specified during lens generation.

We pose secret image encoding as an inverse light transport problem and present a fully-automatic approach for designing and manufacturing various types of steganographic lenses (see Figure 3.2). We experiment with various use-cases, for example enabling multiple target images to be warped from a single source image depending on the viewing angle between the user and the lens, or depending on the relative rotation or

alignment of the lens and the source (see Section 4.5 for more results). In addition, while we experiment with two manufacturing processes to generate physical prototypes of hand-sized steganographic lenses, nothing about our technique precludes more exotic use-cases such as those depicted in Figure 4.2: e.g., replacing architectural fixtures with large-scale steganographic mirrors, revealing hidden messages for interactive and exploratory museum exhibitions, sending secret messages that can only be viewed with a user’s steganographic lens, or embedding thin, flexible steganographic lenses in paper currency as an anti-counterfeiting and validation measure.

We are motivated by recent work on computationally embedding images into physical material properties, classic steganographic techniques such as the Cardan grille, as well as “magic decoder rings” which reveal secret messages already present in the source image using masking or subtractive transmission. In contrast, our lenses use optical refraction (or reflection), and we require little relation between the input and output images as long as the original image contains all the colors of the target image. Furthermore, our approach can be passive, removing the need to carefully design or modify the source image (carrier signal) to encode the secret image. We also present an important special case of a steganographic lens called a *universal steganographic lens* (Section 4.4) that completely removes the dependence of the lens on the source image by generating an optical injection between the lens and an arbitrary source image grid.

Steganographic lenses are specialized passive display devices, related to light field displays, and we validate our simulation results with many real manufactured surfaces. Our approach is a first step towards realizing ideas such as those sketched in Figure 4.2. It has possible applications not just in art, education and entertainment (e.g. optical illusions, hidden message re-

trieval, optical decoder rings, holography), but also in banknote verification and security (see Section 4.7).

4.2 Problem statement and goals

Our goal is to design and manufacture passive micro-lenslet arrays, which we call *steganographic lenses*, capable of generating one or many *target images* when placed over one or many physically displayed *source images* (e.g., printed on paper or displayed on a monitor). The manufactured steganographic lenses we present (see Section 4.5) are all hand-sized, however the approaches we present are not limited in the size or form-factor of the lenses they can generate and one can imagine both larger-scale lens installations (in architectural settings) or miniaturized lenses (for banknote verification).

We will begin by describing the simplest case of warping a single source image to a single target image (i.e., viewed from a single position), and later address more interesting use-cases (Section 4.3.6).

Each steganographic lens is divided into an $N \times N$ grid of refractive *facets*¹ which we initially assume to be flat to simplify our exposition. Our target images are also divided into $N \times N$ tiles, which we wish to observe through the facets. We decompose source images, which are warped by our lenses, into a set of unstructured *image patches*, which are arbitrarily-sized rectangular regions in the source.

Each facet is constructed to refract light from the source image so that, when viewed from above, the image seen through each facet matches that of the corresponding image tile. In

¹A similar principle could also be applied to *reflective* facets.

4.2 Problem statement and goals

Section 4.3.5 we eliminate the flat-facet constraint using a multi-scale refinement, and so our finalized steganographic lenses are ultimately composed of a grid of $N \times N$ smooth lenslets. We will present two different types of steganographic lenses: *source-optimized* lenses and *universal* lenses.

Source-Optimized Lenses. In the simplest case, a source-optimized lens will “index” (through refraction) only the subset of source image patches necessary for it to reproduce its target image. As such, a source-optimized lens is “tied” to its input source image. This constraint allows us to better enforce smoothness and manufacturability, but at the cost of flexibility. This case is particularly useful when the source image must remain fixed (e.g. in the case of a well-known painting). We detail our approach for generating source-optimized steganographic lenses in Section 4.3.

Universal steganographic lenses. These lenses are an important special case and are constructed so that each lens facet indexes a unique source image patch, forming an *optically injective mapping* between the lens and the source (i.e., each pixel in the source image maps to *at most* a single facet). As such, a single universal steganographic lens can be used with an arbitrary number of source images which can be scrambled, according to the optical bijection, in order to generate a target image. This case is most useful when there is flexibility in manipulating a source image or when a single lens must be used with an undetermined set of images. We discuss the modified approach we use for universal steganographic lenses in Section 4.4.

4.3 Generating source-optimized lenses

Source-optimized lenses are most related to goal-based caustics [Papas *et al.* 2011], where a lenslet array warps incident light into a desired caustic image. Papas *et al.* form a bijection between facets and caustic image splats, however our search space is far larger as we are not constrained to the subset of solutions induced by one-to-one mappings. Instead, for source-optimized lenses, each facet has complete freedom to refract to an arbitrary source patch.

While having this larger space does expose a wider range of potential facet-patch pairings, special care must be taken to maintain a reasonable computation cost. We juggle these constraints using a multi-step and multi-pass approach: we ensure a baseline image quality using facet-patch matching (Section 4.3.2) *prior to* optimizing the geometric lens quality (Sections 4.3.3 and 4.3.5).

We decompose the procedure for generating source-optimized steganographic lenses into four mandatory (and one optional) steps (Sections 4.3.2 to 4.3.5): facet-patch matching, facet orientation, simulated annealing, and multi-scale height optimization. Optional facet-patch matching variants are detailed to handle special use-cases (Section 4.3.6) such as multi-source/multi-target warping, multi-view warping, and multi-rotation warping. These steps are repeated until a termination criterion is met. Figure 4.3 overviews the entire procedure.

4.3.1 Input

Source-optimized lenses require at least one input image and one or more target images to be specified. Focal distances from

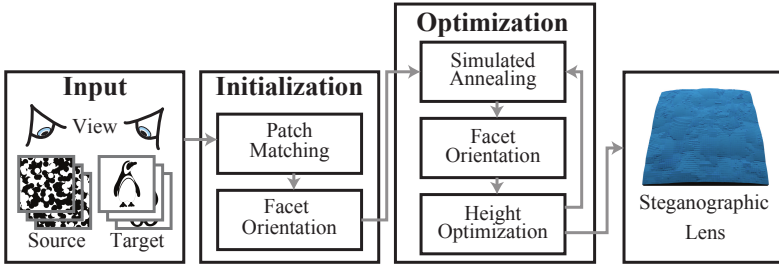


Figure 4.3: Overview of source-optimized steganographic lens generation.

the lens to the source, and from the viewer to the lens must also be specified.

4.3.2 Facet-patch matching

Each facet of a source-optimized lens should refract a portion of the source image that matches the facet's associate target image tile. The first step of surface generation is to find and assign suitable facet-to-patch matches.

We divide *target images* into grids of $N \times N$ tiles and map each tile to its facet on the lens. Now, for each facet, we must find and rank regions in the *source image* which can serve as potential matches (according to pixel-wise differences) to the target image tiles.

We then search image regions within a neighborhood $\{R_x, R_y\} \in [-m/2, m/2]^2$ of $m \times m$ pixels in the source image, starting directly underneath the facet (assuming the lens is aligned atop the source). We compute a *matching score* for every eligible region, $M(R_x, R_y) = M_d(R_x, R_y) + \lambda M_s(R_x, R_y)$,

Steganographic lens fabrication

where M_d is the image color matching term

$$M_d(R_x, R_y) = \sum_{i \in \mathbf{P}} \sum_{j \in \mathbf{P}} [\mathbb{T}(i, j) - \mathbb{S}(i + R_x, j + R_y)]^2 \quad (4.1)$$

and $M_s(R_x, R_y) = (m^2 - R_x^2 - R_y^2)^{-1} - m^{-2}$ is a smoothing term that penalizes distant matches which cause steep normals. Here, \mathbf{P} is the target patch region, \mathbb{S} indexes the source, \mathbb{T} indexes the target, and $\lambda = 255^2 \mathbf{P}^2$ is our regularization constant.

We compute M for every $R_x \times R_y$ offset and additionally compute matching scores at multiple scales S . We initially associate the best (lowest) match to the facet. Apart from a matching score, each potential match is assigned a *facet-patch transformation* (FPT) which consists of the match's (R_x, R_y) offset and the corresponding scale S , which map the facet to the associated region on the source. All matches with scores below an acceptance threshold are retained and sorted for future consideration, at each facet.

This marks the half-way point of the initialization phase for source-optimized lens generation.

4.3.3 Facet orientation

Once each facet has been associated its top-ranked FPT, we proceed with the final initialization phase: assigning an orientation to each facet in order to induce the refraction necessary to match its FPT's transformation. It is important to note that facet heights are not yet modified and thus, after these first two initialization steps, the surface is discontinuous and therefore not manufacturable.

4.3 Generating source-optimized lenses



Figure 4.4: Source/target images (left pair) are mapped with a source-optimized lens without (middle set) and with (right set) simulated annealing. Simulated results are indistinguishable but the lens generated with simulated annealing is smoother and manufacturable. The lens (128×128 facets with 11×11 micro-facets) is designed to be placed at 10 cm from the source and 40 cm from the viewer.

To determine facet normals we solve an inverse lighting problem: the normal of each facet is computed so that an eye ray passing through the facet center will refract to the matching source image patch. We define this mapping as $g_0 : (x, y, \Delta x, \Delta y) \rightarrow n_{xy}$, where (x, y) index the lens facets in the plane of the lens, $(\Delta x, \Delta y)$ are the offsets onto the source image from the facet (in the source image's parallel plane), and n_{xy} is the facet normal. With $\theta_b = \arctan(d/D)$, we can solve for θ_r using Snell's Law and the fact that $\theta_i = \theta_r - \theta_b$:

$$\theta_r = [\arctan(\eta_o - \sin \theta_b)] / [\eta_o \cos \theta_b - \eta_i] . \quad (4.2)$$

We solve for the surface normal that induces the unique refractive light path between the source patch and facet with a tailored version of Walter et al.'s [2009] technique, which simplifies the problem to discrete root finding solved using a handful of Newton iterations.

At this point, the initialization of the source-optimized lens generation is complete and we enter the optimization phases.

4.3.4 Simulated annealing

We find solutions in our large, non-linear search space using a stochastic simulated annealing search. This search will con-

Steganographic lens fabrication

sider FPTs across neighboring facets while increasing surface smoothness and maintaining high quality reproduction of the target image. Surface smoothness is a useful property for fabrication methods such as 3-axis milling where discontinuities are difficult to manufacture.

We induce smoothness between neighboring facet normals using stochastic search: we choose a facet at random and consider swapping its current FPT with another from its list (see Figure 4.5) if the new FPT assignment improves the lens smoothness of the lens (see below). We repeat this process at least $25 N^2$ times.




Figure 4.5: During simulated annealing, a facet is randomly selected (left) and its current FPT is swapped with another from the list of FPTs generated during facet-patch matching (right).

When considering the swap we compute the facet's new normal using facet orientation, and we compute a *facet energy* consisting of a geometric smoothness term E (see below). Prior to the first swap, the initial lens energy E_{initial} is the sum of E for every facet.

After computing the facet energy E , we perform an *energy validation* step to check whether the change in facet energy satisfies a

4.3 Generating source-optimized lenses

threshold $T = E_{\text{initial}} \cdot e^{0.01t} / (100 \cdot N^2)$, where t is the iteration number corresponding to simulated annealing's temperature metaphor. After repeating the swapping and energy validation steps, accepting FPT changes only if the energy change is below the threshold T , we compute the *total lens energy* as the sum of all facet energies. If this energy is lower than E_{initial} , we update the lens with the new facet settings and t is incremented.

This entire process is repeated until either the total energy of the lens does not decrease for 10 iterations in a row, or the reduction of energy is less than $10^{-1} \times E_{\text{initial}}$. An example of a lens before and after simulated annealing is shown in Figure 4.4 with its corresponding simulated results.

Smoothness Term for Normals. The per-facet geometric smoothness energy term

$$E = \alpha \|\Delta g_0\|^2 + \beta \|E_1\|^2 + \gamma \|E_2\|^2 \quad (4.3)$$

is computed every iteration at each facet. Here Δg_0 is a discrete geometric boundary gradient term

$$\begin{aligned} A &= \|g_0(X+1, Y) - g_0(X, Y)\|^2 = \|(a_x, a_y)\|^2 \\ B &= \|g_0(X-1, Y) - g_0(X, Y)\|^2 = \|(b_x, b_y)\|^2 \\ C &= \|g_0(X, Y+1) - g_0(X, Y)\|^2 = \|(c_x, c_y)\|^2 \\ D &= \|g_0(X, Y-1) - g_0(X, Y)\|^2 = \|(d_x, d_y)\|^2 \\ \Delta g_0 &= \sqrt{(A + B + C + D)/4}, \end{aligned} \quad (4.4)$$

measuring the deviation of the two normals in the x (A and B) and y (C and D) directions, and E_1 and E_2 are component-wise geometric error terms

$$E_1 = (a_y + b_y + c_x + d_y)/4, \quad E_2 = ((a_y - b_y) - (c_x - d_x))/8,$$

which correspond to deviations of the normal's y components in the x direction, and the normal's x components in the y direction (the change of the z component is corrected during height optimization; see Section 4.3.5) for E_1 , and a constraint on the integrability of the surface for E_2 . We normalize each term and set $\alpha = 0.125$, $\beta = 0.375$ and $\gamma = 0.5$ so that $E \in [0, 1]$.

4.3.5 Multi-scale height optimization

Once every facet has been assigned a FPT (and hence an orientation), we formulate the problem of optimizing heights for surface smoothness as a global optimization in the heights of the facets. This stage serves two important purposes. Firstly, facet heights are adjusted to maximize coincidence with neighboring facets along facet boundaries (see Figure 4.6), which increases manufacturability². Secondly, after facet heights are optimized, each facet is decomposed at a finer-scale into a grid of micro-facets which are optimized to approximate a smooth lenslet shape over the facet. Height optimization is equivalent to integrating a normal field into a height field and we use the over-constrained linear system of Papas et al. [2011] to solve this problem.

We repeat height optimization with facet orientation until convergence since: any height change may invalidate the selected FPT, which would require facet re-orientation; and, any facet re-orientation may reduce facet boundary smoothness, requiring further height optimization. After facet height and orientation are set, we dice facets into grids of micro-facets, performing a nested iteration of facet-orientation and height optimization on the micro-facets grids of each facet in order to generate smooth

²Smoothness is necessary for manufacturing with an engraving milling machine, but can be relaxed when using a 3D printer (Section 4.5).

4.3 Generating source-optimized lenses




Figure 4.6: Height optimization displaces facets with (non-geometric) normals (left) along z , enforcing smoothness along facet boundaries and inducing the correct geometric normal (right).

lenslets instead of planar facets for our final surface. Microfacets are generated once, at the end of the entire procedure.

4.3.6 Facet-patch matching [optional]

Several variants of the standard single-source/target steganographic lens can be realized by adjusting the facet-patch matching process. We detail three important examples below, but note that the application of these steps is optional depending on the desired use of the lens.

Rotating Lenses. One interesting use-case for source-optimized lenses is when a single lens produces several different target images depending on its relative rotation about the source image. To realize this type of lens, we can extend our matching scores to include scores (and FPTs) for rotated versions of the source image. We do so by augmenting the source sampling function S with a rotation parameter, and, for our square-shaped lenses, we have experimented with $0^\circ, 90^\circ, 180^\circ$, and 270° rotations.

Multiple Source-to-Target Pairings. Another use-case is when a single source-optimized lens is used to warp several source images to different target images. In this case the matching score is computed as the sum of matching scores for each source/target pair.

Multi-view Lenses. Multi-view source-optimized lenses are a more complex use-case: a lens that warps a single source image into different target images depending on the viewing angle of the user with the lens. This is a more complicated case since changes in the FPT for a given view direction induce changes in other view directions. We alternatively set each viewing direction as a “master” direction and optimized separately according to its FPT. The final matching score, as in the case of multiple source-to-target pairings, is the sum of the matching scores over all viewing direction conditions given the facet orientation for the master direction. After each direction has been selected, multiple matching scores corresponding to each master direction are sorted and only the top score which corresponds to the best master direction is kept for each FPT.

For each multi-target use-case, we require an additional optimization step after height optimization to re-introduce new FPTs that take the updated lens geometry into account. This occurs since height changes affect the regions each facet can “index” from the source (and thus invalidates the facet’s FPTs).

4.3.7 Output

The result of this process is a mesh of a steganographic lens surface which we then forward to one of two manufacturing pipelines for physical construction (see Section 4.5).




Figure 4.7: Heatmap for a source-optimized lens (left) with values from 0 (blue) to 18 (red) overlaps. Clearly, several facets are matched to the same source patch. The heatmap for a universal steganographic lens (right) illustrates the injective mapping: red corresponds to a single match and blue to no match.

4.4 Generating universal steganographic lenses

While source-optimized lenses are tied to the input source images, we can also devise a generalized patch-matching process that does *not* depend on the source image, allowing us to construct a steganographic lens that can be re-used with an arbitrary number of different sources. The dependence on the source image in source-optimized lenses arises from the fact that many lens facets may target the same region on the source image. Figure 4.7 visualizes the difference between the facet-patch mappings of a source-optimized lens and a universal steganographic lens, where we can clearly see that each patch on a source image is sampled by *at most* one facet on the universal steganographic lens.

Such a *universal steganographic lens* could be thought of as a “key” that can “unlock” hidden target images embedded in seemingly

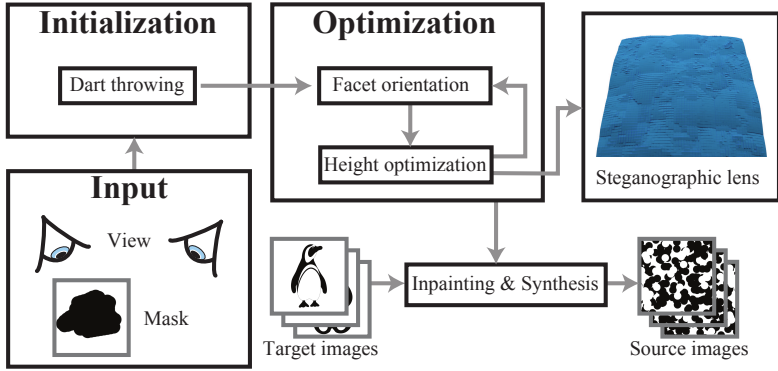


Figure 4.8: Overview of universal steganographic lens generation.

arbitrary (but specially authored) sources. An example use-case where this could be useful is the generation of encoded animations (see Figure 4.9).

The fundamental difference between universal and source-optimized lenses is that, once constructed, source images must abide by the optical injective mapping induced by the lens in order to be decoded by it. Figure 4.8 overviews the different steps of the universal steganographic lens generation process: assuming a single-view use-case again for simplicity, we substitute the facet-patch matching and simulated annealing steps of the source-optimized procedure with a customized dart-throwing approach that forms the injective mapping between our lens and the source. In addition, after repeating iterations of height optimization and facet orientation, as in the source-optimized lens procedure, we follow with an inpainting stage that embeds the target image into an arbitrary source image.

4.4 Generating universal steganographic lenses




Figure 4.9: *Unscrambling an unstructured video sequence with a universal steganographic lens with 16×16 facets. The middle row has the different (shrunk) source images, and the top and bottom rows are photographs of a manufactured lens placed over the source images as they animate.*

4.4.1 Orientation initialization with dart-throwing

We devise a modified dart-throwing approach which replaces facet-patch matching and facet orientation with: choosing a patch on the source image to match to facets on the lens, and determining the orientation of the facet in order to target the matched patch. For universal steganographic lenses, however, each source patch can only be matched to one facet at most.

We employ a multi-scale approach, dividing the source into $N \times N$ square patches, aligned with the facets of the lens. Then, each patch is divided into smaller square sub-regions in order to promote flexibility in finding a match without resulting in an identity mapping. At each facet, we then select one of the physically-reachable (via refraction) regions at random and mark it as used so it cannot be selected by any other facet. After

Steganographic lens fabrication

repeating for all the facets we are left with an injective mapping between the facets on the universal steganographic lens and the sub-regions (and patches) on the source.

We can optionally supply a source *mask* image to invalidate regions on the source from use in the dart-throwing selection procedure. The mask can be useful to reserve regions on the source images that are guaranteed to not perturb the output target image.

4.4.2 In-painting and texture synthesis

Given a universal steganographic lens, specified by the output of the dart-throwing process, and a target image, we take an empty (or arbitrary) source image and inpaint the sub-regions in order to match the target image, given the injective mapping.

Once in-painting is complete, some pixels of the source image (which are never indexed by any facets) are uninitialized (white). We optionally perform an additional texture synthesis step in order to fill the uninitialized regions with camouflage: we isolate un-mapped sub-regions in the immediate neighborhood surrounding mapped sub-regions, and synthesize an input pattern or a background image in these regions. This approach was employed in our *worm animation* example in Figure 4.9.

4.4.3 Iterative adjustment and convergence

After the execution of the dart-throwing injective initialization and the inpainting steps, the final universal steganographic lens is obtained by similarly repeating the height optimization and facet orientation stages until our stopping criterion is met.

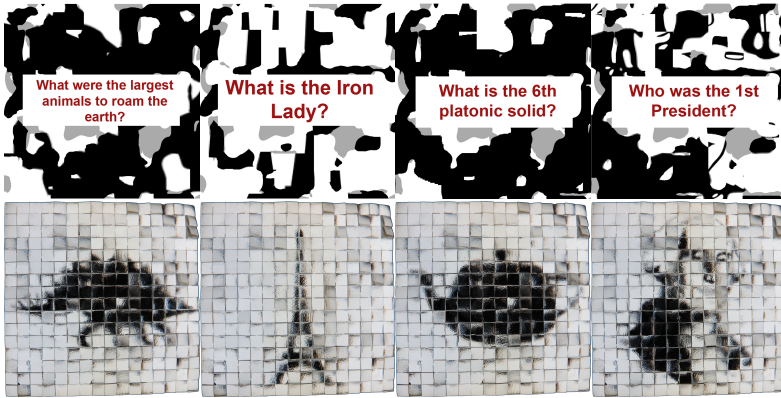


Figure 4.10: Four source images with questions (top; center zoom-ins) are warped to reveal pictographic answers (bottom) with a single universal steganographic lens (16×16 facets). Note that the region containing the question is not targeted by any of the facets. More examples and the full-sized sources are included in Appendix B.

As with source-optimized lenses, this results in a mesh surface which we forward to our manufacturing pipeline for physical construction (Section 4.5).

4.5 Fabrication details and results

We evaluate our lenses with simulated and fabricated results.

Fabricated Results. We physically manufacture our lenses using two separate processes: milling of acrylic blocks using a computer controlled engraving machine (a Roland EGX-600), and 3D printing with a clear material on a fast 3D prototyping machine (an Object Connex 350). These two manufacturing

Steganographic lens fabrication

modes have different physical constraints, which we can account for in our optimization. All lenses are 10×10 cm, with thickness varying between 2 to 4 cm. All lenses are placed directly on the source image, unless stated otherwise in the caption.

Figure 4.11 shows a source-optimized steganographic lens that transforms an image of random circles into a panda. We visualize the steganographic lens surface and show a simulated result as seen through the lens. We manufactured this lens using our milling approach — a photograph of the fabricated lens placed over the source image is also shown.

The precision of the milling machine is higher than our 3D printer, but the surfaces it produces have a frosted finish and need to be manually polished. This can be quite time consuming, and can also alter the shape of the lenslets, reducing quality. Because of this, and because 3D printed lenses do not need to be as continuous, we produce all our remaining results with the 3D printing method.

In Figure 4.12 we show another source-optimized steganographic lens with the panda as a target image but, when viewed at a different angle, the lens reveals a penguin image, all using the same random circle pattern as a source. Figure 3.2 also uses the same random circle source, but the source-optimized lens is capable of displaying four different targets (panda, bat, penguin, and whale) at four 90 degree rotation increments.

We also manufactured a universal steganographic lens that we place over a scrambled image sequence (Figure 4.9). When viewed through the physical lens, an animated sequence of a worm is revealed.

Steganographic lenses can also be used for more creative interactions. For instance, question-answer lenses can be designed




Figure 4.11: A physically milled result of a source optimized lens. Simulation (middle left) closely matches the target, however the smooth lens (middle right; 32×32 facets with 11×11 micro-facets) is degraded during milling and manual polishing, resulting in a lower quality physical result (right) than that of a 3D printer (see Figure 4.12).

to reveal pictorial answers when placed over source images posing questions. Figure 4.10 shows four different question images being answered using the same universal steganographic lens as in Figure 4.9. Since we use a universal steganographic lens, there is actually no limit on the number of randomized question-answer image pairs that can be used with this single lens (see Appendix B for more examples).

Simulated Results. In Figure 4.4 we illustrate, using simulation, the theoretical output that could be obtained with a higher quality manufacturing process. We simulate a high resolution lens (128×128 facets) that warps two images with fine-scale details. Our simulated annealing process is clearly capable of generating smooth, continuous lenses, but these lenses are beyond the manufacturing capabilities of our current milling or printing processes (especially at hand-held output sizes). In Figure 4.16 we show another example of a question-answer lens, this time source-optimized to the four question images, and simulated at a higher resolution (128×128 facets) than our current manufacturing processes allow.

4.6 Discussion

We have presented the idea of steganographic lenses as well as a general approach for creating various special case lenses. This is a significant first step towards realizing additional use-cases in the future, and we will discuss lessons learned and some considerations for extending our work in this section, before itemizing concrete ideas for future work in Section 4.7.

Manufacturing Quality and Considerations. Our proof-of-concept manufacturing process cannot yet yield lenses which match the high quality of our simulation results. In general, these discrepancies arise due to the differences between our idealized model of specular refraction used in simulation, and the additional surface roughness introduced during the manufacturing process (for both milling and 3D printing processes). In the case of milling, surface roughness is further diminished through a manual polishing post-process, which is necessary to render the lens smooth enough for refraction but which also degrades the surface quality and accuracy (when compared to the prescribed mesh output). Milling also imposes additional surface smoothness constraints, limiting the space of image patches each facet can index. To alleviate this constraint, we place milled lenses higher above the source and not directly atop it (as is the case for 3D printed lenses). Figure 4.11 is the only milled result we illustrate, and it is clearly of lower quality than the results generated using the 3D printer, due to the aforementioned issues. This lens is placed 10 cm from the source and is intended to be viewed at 40 cm above the lens.

Our fabricated lenses are currently thick and likely too bulky and expensive for immediate public use. Improved manufacturing processes in the future can permit us to shrink the thickness




Figure 4.12: A 3D printed multi-view result. Left to right: source image, lens (32×32 facets), and photographs of a manufactured lens viewed from two locations.

of the final lens object without compromising the warping capabilities of the lens, especially in the case of 3D printed surfaces where smoothness is less of an issue. In general, we do not see any reason why a higher quality (e.g., commercial) manufacturing process could not rectify many of the discrepancies between our simulation and manufactured results.

Physical Limitations. We consider the physical manufacturing limitations that must be taken into account when designing our lens optimization. We constrain the search range of each facet to 10% of the source image width to reduce the maximum steepness of facets and increase the lens surface regularity. When milling, lens thickness is limited to 4 cm by the milling machine and vertical transitions are limited in range due to the V-groove shape of the milling bit. On the other hand, two pragmatic limitations of the 3D printing process are its price and limited resolution. Transparent 3D printing material costs several orders of magnitude more than an equivalent volume of acrylic, and the 3D printer's precision is limited for our purposes, with significant surface roughness added as a side-effect of printing (Figure 4.13).

We found that, despite the additional undesired surface rough-



Figure 4.13: *A photographic zoom-in of one of our 3D printed steganographic lenses. Notice the easily visible z layers deposited by the printer and the semi-transparent nature of the printing material.*

ness introduced by 3D printing process, we can actually place the final lens objects *directly* on the source image mainly due to fewer constraints on inter-lenslet smoothness. By moving the lens closer to the source image we reduce blurriness and calibration sensitivity, resulting in increased usability. Reducing the distance between the lens and source often requires an increase in the facets' slopes. This increase introduces a non-negligible Fresnel reflection effect, especially when the angle between the average facet normal and the view direction is greater than 60° . We account for these effects during lens construction using the search distance threshold and using the physical lens in low (external) light settings.

Source Image Constraints and Manipulation Currently, the two types of lenses we produce place different constraints on the input. Source-optimized lens assumes a very rigid constraint on the source image (namely, that it is provided by the user as input and cannot be edited), whereas the universal steganographic lens allows complete freedom when generating

the source image(s). However, one could imagine a continuum of approaches between these two extremes: for instance, where a source image is provided, but the optimization is allowed to modify it for increased fidelity or smoothness. This type of approach could be very useful for our source-optimized examples which use an unstructured random pattern as the source. We have indeed experimented with applying variants of inpainting and texture synthesis procedure *during* the simulated annealing component of source-optimized lens generation. In doing so, we allow (un-masked) regions of the source image to be manipulated during lens optimization. Our initial results have proven promising and have permitted the generation of smoother and more accurate source-optimized lenses.

Use-case Customization. We have demonstrated that our approach can be readily and easily extended to handle more complex use-cases (see Section 4.7 for examples, as well as Figure 4.2), and it is our hope that this will promote future work in not only investigating interesting use-cases that we have not conceived, but also in using our approach for larger-scale applications.

Another example of this type of customization is presented in Figure 4.14, where we modify our multi-view lens generation to enhance the viewing stability of our manufactured lens: we generate a multi-view lens, with views arranged in a circular-cross layout, and with all the target images set to the same panda image. The apex of the inner viewing cone is 1.15° . The resulting multi-view lens affords a much larger “viewing zone”, in case a user slightly shifts their viewing position.

Cryptographic and Steganographic Strength. For certain applications steganographic or cryptographic strength is an

Steganographic lens fabrication

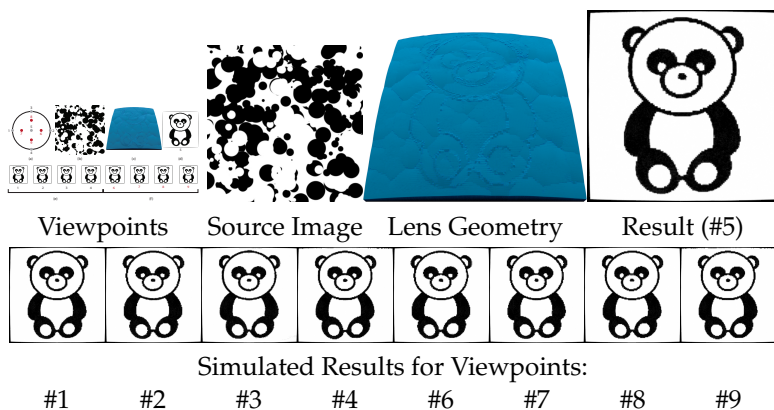


Figure 4.14: We induce a stable viewing zone by exploiting multi-view lens generation in a novel way: by using the **same** target images at each of the nine views, arranged in the layout on the left. The lens consists of 128×128 facets.

important consideration. For a useful discussion, “strength” needs to be better defined in our context and depends on the usage scenario. If an “attacker” gains access to the lens, the mapping can be easily inverted (i.e., it is weak). However, for image-specific lenses, if the attacker obtains the source image it is hard to deduce the target without the lens (i.e., it is strong).

For our physically-manufactured universal steganographic lenses the mapping is not cryptographically strong since parts of the target could be deduced from the source. However, this is a limitation of the manufacturing quality (using fewer/large facets), and not of our overall approach. With improved manufacturing, or in simulation, our universal steganographic lenses can be “steganographically stronger” in the limit of pixel-wide facets: e.g., by randomly scrambling the image pixels plus injecting arbitrary noise pixels to further evade detection.

A useful analogy of our universal steganographic lenses are pa-




Figure 4.15: Visualization of the warping behavior of a universal steganographic lens with 16×16 facets: Given the target image (b) and lens geometry (c), we compute the source image (a) using our inpainting method.

per shredders (cutting images into small slices) but our method shreds and scrambles in 2D while also allowing for zooming effects. A shredder that produces only a few slices (large facets) is less secure. Security is improved by increasing the number of slices (smaller facets) and mixing slices with other random shredded documents (inpainting with random pixels). We illustrate the mapping obtained from a universal steganographic lens that satisfies manufacturing constraints in Figure 4.15.

Visual Cryptography [Naor and Shamir 1994] can provide stronger cryptographic guarantees than our method. Our method does not explicitly optimize for cryptographic strength, though we believe this is an interesting avenue for future work.

3D Displays. Mass-produced 3D technology (e.g., lenticular sheets, parallax barriers) are less expensive to manufacture than our lenses, but these have significant limitations such as low signal to noise ratio, aliasing and crosstalk. Customizing 3D display optics to the expected displayed content is one possible direction which could overcome these issues. Our lenses




Figure 4.16: A single source-optimized lens (128×128 facets) encodes four hidden pictographic answers to four textual questions (simulated).

are an example of such customized optics and, although they are not yet cost-effective, fabrication quality and cost will only improve in time, opening the door to content-optimized display techniques, like ours, in this domain.

4.7 Conclusion and future work

We introduced *steganographic lenses*, a form of passive display device for exposing hidden messages and images from both seemingly random as well as structured source images. We are able to efficiently navigate a large search space of viable solutions using stochastic search and multi-scale height optimization and our process also addresses manufacturability constraints of the final output lens surfaces. We have manufactured physical prototypes using two processes: commodity 3D printing and engraving with a milling machine. We obtained high quality simulation results and also demonstrated promising

4.7 Conclusion and future work

manufactured results, despite using a proof-of-concept manufacturing process.

Future Work. In the case of colored source and target images, the source *must* contain all the colors desired in the target image. One can imagine using back-lit source projectors and modeling, or even inducing, *dispersion* effects in order to side-step this constraint.

As we currently assume perfect refractive surfaces, a better modeling of the surface roughness resulting from our manufacturing processes can be achieved by using a physically based BSDF such as the one introduced by Walter et al. [2007]. Renderings of steganographic lenses with the aforementioned rough glass BSDF can be found in Appendix B. This may lead to better agreement between simulated and manufactured results, however we believe that a higher-quality manufacturing process is a more important first step of investigation towards this end.

We similarly also assume perfect diffuse reflectance of the source image(s), and an interesting direction of future work would be to model the effects of glossy reflection or directional emission off of the source. In this case, the observed radiance of each facet will change not only as a function of location but also as a function of the facet's "viewing" direction. As discussed in Section 4.6, investigating methods for optimizing or manipulating a set of random source images to best reproduce a desired set of target images is also an interesting area of future work.

Our multi-view steganographic lenses may be suitable as a form of stereoscopic or light field display. We have experimented with the stereoscopic use case (using a 3D stereo pair as target images, and specifying view conditions that line up

Steganographic lens fabrication

with inter-eye distance), however with mixed results. We have yet to find a solution that consistently generates stereo image pairs that are comfortable to fuse.

As we feel that the general idea of steganographic lenses can be exploited for many different specialized applications, the most exciting areas of future work may in fact be the ones we have not yet conceived of. It is quite possible that tailored steganographic lens solutions will find their way into unique optical illusions (e.g., imagine an enhanced version of a magic mirrors room at an amusement park), cheap personal encryption, holography, banknote security, or architectural art (Figure 4.2).

CHAPTER

5

A fast, forward model for granular media

In this chapter we investigate the material class of granular media such as salt, sand and snow. A full computational fabrication pipeline for this class of materials remains undefined at the time of writing. In this chapter we present a critical component which is required for a full appearance fabrication pipeline— a fast forward model for simulating the appearance of granular media. This component could potentially be useful in the future for creating a full pipeline, since it fulfills the necessary requirements of the framework. This chapter is partially based on the paper by Meng et al. [2015] which presents a multi-scale model for rendering granular media. From that work we describe the diffusion extension which provides significant rendering speedups, and makes the model more compatible with the requirements of computational appearance fabrication.

5.1 Introduction

Treating each individual grain as explicit geometry and simulating global light transport using path tracing [Kajiya 1986] and its variants is a general solution for simulating the appearance of granular media, but is only practical for small collections of grains. At the other extreme, the aggregate object could be interpreted as a continuous medium. In this case, the smooth, large-scale appearance of the medium may be well expressed by rendering techniques [Cerezo *et al.* 2005] derived from the radiative transfer equation (RTE) [Chandrasekar 1960]. Methods based on the *diffusion approximation* [Stam 1995; Jensen *et al.* 2001] can further accelerate multiple scattering computation. Unfortunately, even if appropriate volumetric parameters could be determined for the continuous medium, such an approach cannot reproduce the fine-scale structure of individual visible grains (Figure 5.3).

Similar to Meng *et al.* [2015] we model a granular material as a procedurally defined, tiled assembly of objects representing individual grains. We augment the model by using an additional model for the transport of light in the material based on the *diffusion approximation* (DA). In Meng *et al.* [2015] they express the appearance of visible grains, using the most detailed model, which explicitly path-traces the grain geometry (*explicit path tracing*, or EPT). To more efficiently capture larger-scale transport above the scale of grains, they approximate the granular material as a continuous medium and render using *volumetric path tracing* (VPT). To avoid the need to trace long paths, for scales above the mean free path of the volumetric medium, we propose a diffusion-based approximation. We therefore implement a transition to a rendering technique based on a fast diffusion approximation [Stam 1995; Jensen *et al.* 2001], to account for large-scale transport due

to high-order multiple scattering. Diffusion allows us to effectively short-circuit the recursion of VPT and approximate long scattering paths directly. An overview of the full forward model of Meng et al. [2015] including our diffusion extension is shown in Figure 5.2. We refer the reader to that paper for the details of EPT and VPT, which are not part of this thesis.

We switch to diffusion, by sampling a location on the boundary mesh and estimating the diffusion transport. We use a technique inspired by the method of Li et al. [2005], though we use d’Eon and Irving’s [2011] improved diffusion model instead of the classical dipole [Jensen *et al.* 2001], and integrate its contribution using Monte Carlo [Habel *et al.* 2013]. We also propose a different virtual source placement procedure which in, our experiments, produces slightly better results than the approach described by Li et al. [2005]. Finally, we account for the finite thickness of the medium by using a multi-pole expansion [Donner and Jensen 2005]. We detail these changes in Section 5.4.

Using multiple levels of approximation in a practical rendering system requires solving two fundamental problems: consistency between approximations, and finding optimal switching criteria. In Meng et al. [2015] they obtained parameters for the first level of approximation (VPT), without any expensive scene-dependent precomputation that would preclude modifying the shape or composition of the aggregate object. They developed a stochastic “teleportation” model of light transport which accounts for *intra*-grain transport, and combined it with analytic estimates for *inter*-grain propagation imported from the physics literature [Torquato 2001]. After a one-time precomputation for an individual grain, they can reuse the results to derive parameters for arbitrary aggregate shapes and a continuum of different packing rates and scales. We detail the conversion of these VPT parameters to our next level of approximation in




Figure 5.1: *Diagonal comparison of proposed techniques (top left) with EPT (bottom right) on a sphere composed of snow grains. We compare four combinations of our methods: The first two combinations show how VPT with and without diffusion compare against EPT. The latter two show how preceding these methods with a few bounces of EPT can accurately recover high-frequency detail.*

Section 5.2. Second, we need to determine when to use each of the three approximations to obtain maximum efficiency gain while remaining visually accurate. We achieve this goal using carefully designed heuristics based on how paths diverge in the medium, and how deeply they penetrate the medium.

5.2 Deriving diffusion parameters

Meng et al. [2015] present a method to compute RTE parameters $\{\sigma_s, \sigma_t, f(\cos(\theta))\}$ for a classical homogeneous medium which approximately matches the statistical scattering behavior of the granular material. From these we then derive the necessary diffusion parameters $\{\sigma'_s, \sigma'_t\}$ by computing the mean scattering cosine g from $f(\cos(\theta))$, and from that the reduced scattering coefficient $\sigma'_s = (1 - g)\alpha_s\sigma_t$, which gives a reduced extinction coefficient $\sigma'_t = \sigma'_s + (\sigma_t - \sigma_s)$ and reduced albedo $\alpha'_s = \sigma'_s / \sigma'_t$.

Figure 5.1 compares EPT to VPT with and without DA for a more traditional rendering scenario. The results show that our derived diffusion parameters closely approximate the large-

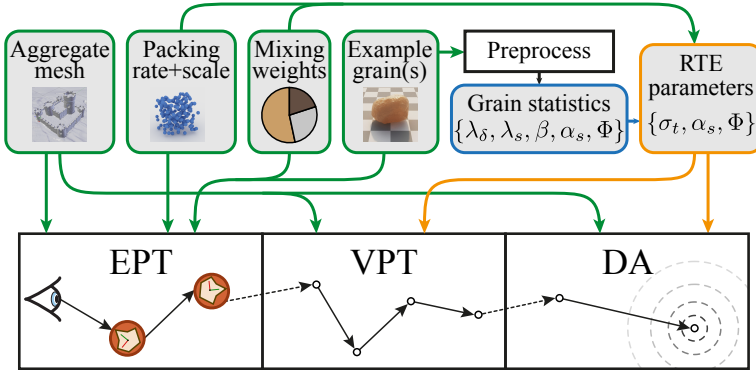


Figure 5.2: Overview of the full method presented in Meng et al. [2015], including our diffusion approximation extension. When rendering the granular material (bottom), primary rays from the eye start out path tracing (EPT) grains, then after enough scattering events take place, the path is continued using volumetric path tracing (VPT), and eventually terminated using a diffusion connection (DA) to approximate the contribution of all further interactions. The input (top, green) to EPT consists of an aggregate mesh, packing rate & scale of spheres, example grains, and mixing ratios of these grains. The grain scattering statistics (blue)—calculated in a one-time, per-grain-type preprocess—are combined with the packing rate and mixing ratios to obtain RTE parameters (orange) needed for the VPT and DA methods.

scale transport within granular materials. In isolation, however, neither VPT nor DA can retain the high-frequency details of discernible grains.

5.3 Switching to diffusion

While diffusion can estimate multiple scattering very efficiently, it also introduces a number of approximations and assumptions (most notably the assumption of isotropic scattering and



Figure 5.3: *A selection of real granular materials illuminated from underneath with a cell phone’s flash light to reveal their translucency. Granular materials can exhibit long-range light transport that still retains complex small-scale structure. Each HDR capture was individually tone-mapped for display.*

planar bounding geometry) that limit its accuracy in the general setting. To switch from VPT to DA, we adopt a criterion adapted from the work of Li et al. [2005] which aims to allow DA only when the approximations are not be too noticeable.

As the aforementioned work, our primary criterion for switching to diffusion is a minimum distance between the VPT path vertex \mathbf{x}_i^k and the surface of the boundary mesh. While Li et al. [2005] used a threshold of one reduced mean free path, we use a threshold of $d_{DA} = \min(1/\sigma'_t, 0.5/\sigma_{tr})$ —that is, we switch if \mathbf{x}_i^k is at least 1 reduced mean free path away from the boundary, or at least half a diffuse mean free path. This allows us to accept diffusion connections more frequently for lower albedo materials, which can be better represented using d’Eon and Irving’s improved diffusion model [2011].

Li et al. [2005] accelerate the distance computation using a kD-tree which stores uniformly sampled points on the surface. However, as the sampling distance needs to be smaller than the reduced mean free path, this approach is practically infeasible

5.4 Multi-pole configuration

for optically thick volumes, such as densely packed granular media. In our scenes, the reduced mean free path is often orders of magnitude less than the mean distance between aggregate mesh vertices.

For computing whether we are in a safe enough distance from the aggregate mesh, we reuse our aggregate mesh kD-tree with a new query. This query performs a depth-first search in the kD-tree nodes using a sphere centered at the candidate point for switching to diffusion \mathbf{x}_i^k , and a query radius d_{DA} . If this sphere intersects the bounding box of the current node, then the bounding boxes of the child nodes are computed and recursively queried. When the algorithm arrives at a leaf node, the minimum distance between all leaf node triangles and \mathbf{x}_i^k is computed and if it is smaller than d_{DA} we terminate, returning true as the answer to the query: "is there a surface with a distance smaller than d_{DA} from \mathbf{x}_i^k ?". When all intersecting nodes are processed and no triangles are found, we allow the switch from volumetric path tracing to diffusion. It should be noted that this query performs remarkably well, since only a small portion of the leaf nodes is actually processed, due to the small size of the sphere relative to the triangle size.

5.4 Multi-pole configuration

In general, the diffusion approximation is derived from the radiative transfer equation (RTE) which is solved by the VPT implementation in Meng et al. [2015]. The RTE assumes that the size of the scattering particles is negligible compared to the mean free path. The diffusion approximation further assumes a slab geometry and isotropic multiple scattering [Jensen *et al.* 2001]. We found that even under these heavy assumptions,

Additionally, instead of assuming a semi-infinite medium, we estimate a thickness $d = z_r + z_a$ for a finite slab, where z_r is the scalar projection of $(\mathbf{x}_i - \mathbf{x}_i^k)$ onto \mathbf{N}_i , and z_a the scalar projection of $(\mathbf{x}_a - \mathbf{x}_i^k)$ onto \mathbf{N}_a , obtained with an additional ray shot in direction $-\vec{\omega}_i$ from \mathbf{x}_i^k . We define $d_r = \sqrt{z_r^2 + r^2}$ and $d_v = \sqrt{z_v^2 + r^2}$, with $z_v = -z_r - 2z_b(0)$. We then derive the distances to the negative receiver and consecutive dipoles using the standard multi-pole expansion [Donner and Jensen 2005] but adopting the improved diffusion model [d'Eon and Irving 2011; Habel *et al.* 2013]. The definitions of the extrapolated boundaries at $z_b(0)$ and $z_b(d)$ can be found in [Donner and Jensen 2005]. We assume that the boundary of the granular medium is index matched (i.e. $\eta = 1$), since our grains are suspended in vacuum. Finally, we resume VPT from \mathbf{x}_i . The auxiliary point and normal $(\mathbf{x}_a, \mathbf{N}_a)$ are only used for computing the thickness of the approximate parallel slab and subsequent multi-pole mirroring.

Note, this is the dual and equivalent to computing the contribution from the dipole sources to the surface location.

5.5 Results

The diffusion approximation extension is implemented in Mitsuba [2010] as part of the granular media Integrator of Meng *et al.* [2015]. We render all results on a homogeneous cluster with nodes containing two 12-core Intel Xeon E5-2697v2 processors at 2.7 GHz with 64 GB RAM and report all render times in core-hours. For time related comparisons, we independently compute many low sample count images across the machines and average the resulting floating-point images for the final result. In all results, RGB channels are rendered separately and then combined and tone-mapped with an sRGB gamma curve.




Figure 5.4: *The SANDCASTLE contains about 2 billion grains, each composed of approximately 200k triangles. We report the high-order / total render times in hours and the variance in parentheses. Our approach (top half) renders the high-order scattering over $12 \times$ (50 vs. 628 hrs) faster than explicitly path tracing (EPT) the individual grains (bottom half) while providing visually indistinguishable results. The insets on the right provide equal time and equal variance comparisons.*

For all our figures we estimate variance (in parentheses) by dividing the *time to unit variance* (ttuv) by the core hours used for the specific scene and method. Time to unit variance for each method is an estimate of the time needed to achieve a variance of 1, assuming $1/N$ variance reduction. We compute ttuv by rendering low sample count versions of the image, and then multiplying the average sample variance over all pixels with the average time needed to render a low sample count image.

In Table 5.1 we provide a summary of render times and analysis of variance for the main scenes shown in Figures 5.4, 5.5, and 5.7. We quantify the efficiency of each algorithm using ttuv. On the top half of the table we report ttuv and the resulting speedup compared to EPT for the overall render time, and on




Figure 5.5: Rendering comparisons of our SPICES scene showcasing flour, pink salt, brown sugar, and white sugar. We provide approximate equal variance (top left) and equal-time comparisons (bottom left) between pure EPT and our successively faster techniques incorporating VPT and DA, reporting the high-order / total render time in hours and the (variance). Our full approach computes the high-order scattering over $120\times$ faster (34 vs. 4148 hrs) than EPT, resulting in an $8\times$ overall speedup for equal variance (top left). Note that all six images on the left include all light transport, and are just simulated using successive subsets of our full approach. Any differences are due to bias and/or variance.




Figure 5.6: We visualize the individual contributions of each technique, which sum to our final image.

Table 5.1: Time to unit variance comparisons (measured in seconds) for all methods. The value in parenthesis is the speedup relative to EPT alone. We report times for 100% acceptance probability (P_a) as well as our automatically computed optimal P_a . On the top we report times for rendering both low order and high order transport whereas on the bottom we report the timings and speedup only for the high order component.

Scene	Time to unit variance, combined (sec)				
	EPT	EPT+VPT		EPT+VPT+DA	
		$P_a = 1$	$P_a = \text{opt}$	$P_a = 1$	$P_a = \text{opt}$
SPICES	1617.3	400.3 (4.00 \times)	249.7 (6.50 \times)	219.8 (7.40 \times)	201.9 (8.00 \times)
SANDCASTLE	30.7	14.4 (2.10 \times)	13.7 (2.20 \times)	13.9 (2.20 \times)	13.7 (2.20 \times)
SNOWMAN LG.	4000.8	132.0 (30.3 \times)	39.6 (101. \times)	16.3 (246. \times)	15.4 (260. \times)
SNOWMAN SM.	2868.4	171.1 (16.8 \times)	52.7 (54.4 \times)	28.1 (102. \times)	25.2 (114. \times)
Time to unit variance, high order (sec)					
SPICES	688.8	5.3 (130.8 \times)	5.0 (137.8 \times)	0.3 (2554. \times)	0.4 (1968. \times)
SANDCASTLE	11.8	0.1 (206.8 \times)	0.1 (230.0 \times)	0.4 (26.50 \times)	0.1 (242.0 \times)
SNOWMAN LG.	3548.6	8.8 (404.6 \times)	9.2 (385.3 \times)	0.4 (8473. \times)	0.5 (7394. \times)
SNOWMAN SM.	2338.0	6.1 (381.5 \times)	8.6 (271.4 \times)	0.4 (6431. \times)	0.5 (4997. \times)

the bottom for only the high-order transport which we aim to accelerate with our approximations.

The EPT+VPT method of Meng et al. [2015] provides a significant overall speedup over EPT in all scenes, ranging from 2.1–30 \times . This is largely due to VPT’s ability to create shadow connections to the light sources at various depths inside the granular medium, whereas EPT must rely on random chance for any grain with a dielectric boundary. Enabling Meng et al.’s [2015] automatic acceptance rate calculation (P_a) provides further improvements, with speedups ranging from 2.2–101 \times . The details of the optimal acceptance rate are described in Meng et al. [2015], but the main goal of P_a is to balance the computation between EPT and higher-order methods in order to minimize ttuv. Enabling diffusion provides significant additional speedup in scenes with low-absorption grains and

5.6 Conclusion and discussion

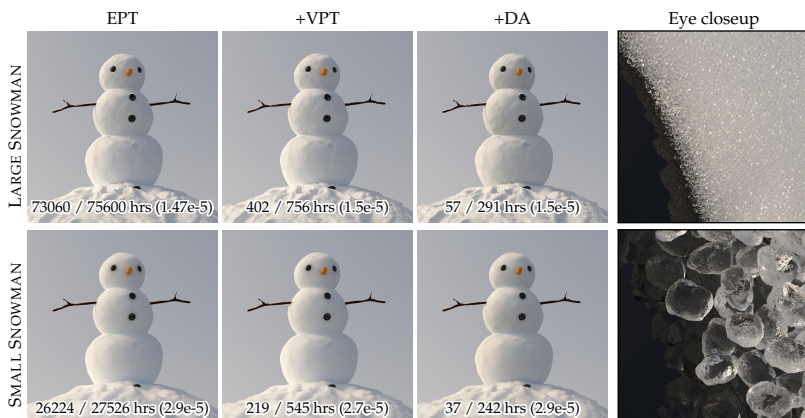


Figure 5.7: For the tall (1 m, top) and tiny (10 cm, bottom) SNOWMAN we report the high-order / total render time in hours and the (variance). All renderings in this figure simulate the full light transport. Any differences within the same row are due to bias.

long transport paths such as SNOWMAN (Figure 5.7, $259\times$ vs. $101\times$ of EPT+VPT [Meng *et al.* 2015]) at the cost of some visible bias. We attain these speedups in total render time primarily by reducing the computation time for the high-order scattering component with the use of diffusion. Measuring just the computation time spent on high-order scattering, our diffusion extension obtains dramatic speedups over EPT ranging from 241 – $7394\times$.

5.6 Conclusion and discussion

In this chapter we demonstrated how to approximate the large-scale behavior of a granular medium with a diffusion based model. Normally, path tracing a high-albedo granular medium like snow is completely impractical because very long, very

high-variance paths contribute significantly to its appearance. Our method uses diffusion approximation to reduce the length of paths, greatly reducing the time spent on high-order scattering. The result is that render times are no longer dominated by long paths; instead the computation is spent rendering the visible structure of grains due to low-order scattering. As shown by our results, this makes it feasible to render many practically relevant scenes that would otherwise take unacceptably long to converge.

The speedups of our method are limited by the continued need to path trace the low-order contributions that create visible grain structure; in many cases we succeed in reducing the cost of high-order paths essentially to zero. In order to obtain further speedup a new problem must be addressed, that of rendering the glittery, structured low-order contributions faster without smoothing out the appearance.

While our method allows mixtures of different grain types in a single medium, we currently assume the mixing ratios are homogeneous at the scale of the aggregate. For some scenes, mid-scale heterogeneity, in which the ratios of particle types vary spatially, is desirable. Our method could be extended to this case by modulating the particle ratios using a volumetric texture and computing correspondingly varying diffusion parameters. Once the spatially varying diffusion parameters are derived then we can leverage existing work on diffusion in heterogeneous media [Donner and Jensen 2005; Hery 2012; Arbree *et al.* 2011] for the implementation.

Leveraging diffusion solutions for non-exponential mean-free path distributions [d'Eon 2014] in our framework might further improve accuracy or allow us to switch to diffusion more quickly, effectively reducing the expensive computation currently needed on low-order.

We proposed a bottom-up approach for specifying the appearance of granular materials where the grain properties and their packing rate dictate the large-scale appearance. An interesting avenue for future work (similar to inverse bi-scale appearance design [Wu *et al.* 2013]) would be a top-down approach where the user specifies, or measures, the desired large-scale appearance. By using this desired appearance as input, the system then proposes the individual grain properties and their arrangement from a dictionary of possibilities. Editing one scale independently while maintaining a fixed appearance at other scales would be a valuable appearance design tool for scenes containing granular media.

In this chapter we displayed the capabilities and efficiency of our forward model as part of a physically-based renderer. We used our model to render with high resolution, complex scenes with intricate light transport to showcase the fidelity of our model when used as a regular appearance model. This is also the main reason for the long computation times reported. In practice the resolution, illumination conditions, and geometry can be significantly simplified in the case of real-world measurements, similar to the setup of Figure 5.3.

When measurements focus on capturing the large scale light transport of granular media—without aiming to extract small scale effects of individual grains—our simplified forward model (DA) can be evaluated quickly. The computation time is similar to the forward model used in Chapter 3 for translucency fabrication. More specifically we used the measurement device presented Chapter 3 along with our simplified forward model (DA) to recover the physical properties of real snow. The recovered physical properties were used in movie production to model the large-scale appearance of snow.

A fast, forward model for granular media

CHAPTER

6

Conclusion

In this thesis we presented a general framework for computational appearance fabrication and a variety of applications on three general classes of materials: homogeneous participating media, specular dielectrics and granular media. We conclude with future directions regarding the individual applications and present our outlook, on future work in the field of computational appearance fabrication.

Translucency fabrication. We presented a multi-spectral measurement device for measuring reflectance and translucency. With the help of this device we were able to train our novel forward and backward appearance models through very few (21) measurements. These models allowed us to simulate the measured appearance and also automatically generate recipes for our pigments and silicone. These recipes, when fabricated, reproduce the desired appearance. Additionally we presented an appearance editor which respects the gamut of the possible real-world physical states and at the same time allows intuitive control over the desired appearance.

Conclusion

We also presented a perceptual backward appearance model which for a fixed human observer and illuminant can increase the gamut of reproducible appearances compared to the published method [Papas *et al.* 2013]. An interesting extension to this work would be to investigate simultaneous appearance matching of multiple physical states with varying illumination conditions and observers.

The designed components of the framework bear some limitations. In our use case we only consider homogeneous translucent materials. In addition, the accuracy of the forward and backward models is reduced in parts of the domain where the underlying appearance model is not a good approximation. We can improve the accuracy with our local method or with a better forward model, but for the cases of highly absorbing materials the bulk-scattering profile measurement method needs to be re-designed. Even though our local method improves appearance replication accuracy, it does require additional target dependent sample generation which in some cases can be a significant drawback.

With the presented instance of our framework, we allow appearance control and fabrication of homogeneous participating media. This ability already has implications for many industries where pigmented materials are used, including plastics, foods, prosthetics, and even dentistry.

Our method can also be used to fabricate more complex materials. More specifically it can be combined with previous work which allows spatial combinations of fixed materials. An example would be using our continuous method to design and fabricate an expressive minimal set of primitives that can be used from the discrete methods [Dong *et al.* 2010; Hasan *et al.* 2010]. This minimal set of primitives can be computationally designed with the goal of spanning a wide ap-

pearance gamut. Moreover, the ability of continuously tuning material parameters opens the possibility to more powerful hybrid methods that can accurately control the real-world appearance of spatially varying translucent materials.

Steganographic lenses. We presented an application of our framework for computational fabrication of custom lens-let arrays. We were able to scramble the incident light field on the array such that it can be used for steganographic purposes. The resulting lenses expose hidden messages and images from both seemingly random as well as structured source images.

The current methods available for fabrication limit the minimum size of the patches relative to the size of the lens. They do not allow dithering or other color mixing methods that take advantage of the human visual system properties. With the invention of more accurate fabrication technologies our method could be extended to accommodate these type of effects, thus lifting the current limitations in color images and increasing the output color gamut.

Though our current manufacturing precision limits the resolution and number of hidden images we can encode, our simulations show that significantly more sophisticated results are possible with higher fidelity manufacturing. This opens the door to many possible applications for encoding and extracting hidden visual information, ranging from entertainment, to architecture and art.

Currently our forward model assumes that the fabrication method will reproduce the computer generated geometry perfectly. In reality the fabrication methods only approximate the desired output and exhibit systematic imperfections. Another interesting avenue for future work would be to model these

Conclusion

fabrication specific inaccuracies within the forward model and effectively minimize their impact on quality.

Granular media. We presented an extension to the appearance model of Meng et al. [2015] to simulate the appearance of homogeneous granular media with high computational efficiency. Our fast forward model approximates the multiple scattering component of light transport using diffusion theory. Our results show that our extension provides significant speedup, especially in media with extremely long paths with high contribution, with minimal impact on accuracy.

Our method uses a diffusion approximation to reduce the length of paths, greatly improving the time spent on high-order scattering. The result is that simulation times are no longer dominated by long paths. Most of the computation is spent rendering the visible structure of grains due to low-order scattering. Our forward model makes it feasible to render realistically complex scenes that would otherwise take unacceptably long to converge.

Currently there are just a handful of fabrication methods that use granular media. One example is the 3D Systems ProJet[®] 660 Pro 3D printer. This powder-based method was one of the first 3D printers able to reproduce full color. It uses a standard ink-jet printer with glue-infused inks. It applies a thin layer of homogeneous white powder and then performs a pass over it with the ink-jet printer heads to apply the pigmented glue and define the shape of the object. This process continues layer-by-layer until the object is completed. Then the surrounding powder that did not interact with the inks can be cleaned. This process allows limited control over the individual grain materials, such as modifying absorption at the grain boundaries. It also indirectly allows (spatially varying)

control over the optical properties of the background medium between grains. An interesting potential avenue for future work on granular media fabrication would be to use an instance of our framework to evaluate whether the appearance gamut of this method can be expanded by changing the optical thickness of the powder and inks or by adding new inks that can additionally scatter light.

A fully generic heterogeneous forward model would use explicit granular medium parameters such as per-grain physical properties: material, shape, size, and orientation, along with spatially varying pack rates and background medium optical properties. Our forward appearance model can already handle such an explicit parameter set but with globally homogeneous constraints on the grain physical properties, pack rate, and background medium.

According to our framework, the remaining steps needed for computational appearance fabrication of granular media are: a training-through-measurement procedure that can map the physical states to corresponding forward model parameters, and a backward model that can translate from a desired appearance to a physical state. Another remaining challenge and potential avenue for future work is the design of a measurement setup for mapping the explicit real-world granular medium physical parameters to appearance.

Outlook. We believe that our computational appearance framework will be useful for researchers who are interested in de-constructing, understanding, and controlling the appearance of real-world materials. We showed various powerful applications and instances of our framework and its use for different materials and problem specifications. We hope that the success and usefulness of these applications will inspire further

Conclusion

research on the appearance aspect of the currently developing field of computational fabrication.

Within the appearance modeling field we currently have an abundance of specialized forward models, each specific to a material class. One possible avenue for future work would be to use machine learning for creating a general "super" forward model that can identify appearance features from measurements and choose (or even create) an appropriate appearance model.

Most of the research in computer graphics community is focusing on capturing and simulating appearance within the visible spectrum. On the other hand there is an abundance of information that can be measured with existing devices outside of the visible spectrum. This "invisible" data can be potentially used in combination with visible data to augment existing appearance models. Some examples include exposing physical properties of materials such as internal structures inside surface boundaries or isolation of different compounds within an object that only interact with a specific range of the spectrum.

Recent advancements in nano-fabrication [Theocharous *et al.* 2014], show that large scale appearance can be affected by fabricated nano-scale structures. Inspired by these developments, we believe that future research can focus on simulating and modeling light transport efficiently outside of the simplified domain of geometric optics (e.g. wave and electromagnetic optics) to enable computational appearance fabrication at nano-scale.

Computational appearance fabrication is a promising new field. As more fabrication methods become available, there will also be more demand to match and control appearance accurately. I believe that these conditions currently taking shape will open up this growing field to a whole range of exciting new research,

which in turn will offer artists, engineers and hobbyists a real-life canvas, where they will be able to experiment and more fully express their creativity.

Conclusion

APPENDIX

A

In this appendix we provide supplemental material for the translucency fabrication chapter (Chapter 3).

Tables A.1 and A.2 provide error values for the local and global methods respectively. In Figure A.1 we provide plots of various appearance models that we considered when designing our forward model. In Table A.3 we provide the optical properties assigned to our synthetic database pigments. In Figure A.2 we evaluate Quantized Diffusion profile fits against path tracing. In Figure A.3 we compare the measured target appearance against the predicted and measured appearance of the replica. Finally, Figures A.4, A.5, A.6, A.7, A.8, and A.9 summarize the appearance measurements of our database and provide forward model fits on this data.

Table A.1: Error values for the 9 fabricated replicas using the local method with cosine concentration weights. The database includes the replicas generated using the global method. The first triplet of errors evaluates how well the fabricated replica matches the target. The second triplet evaluates how well our method believes it can match the targets appearance. The third triplet evaluates how well our method can predict the appearance of the replica given the local pigment parameters and the concentrations.

	Target-Replica Error			Target-Local Prediction Error			Local Prediction-Replica Error		
	Total	ρ	Shape	Total	ρ	Shape	Total	ρ	Shape
silicone mixture 1	0.00031	0.00017	0.00014	0.00011	0.00000	0.00011	0.00037	0.00020	0.00017
strawberry yogurt drink	0.00063	0.00049	0.00015	0.00044	0.00032	0.00012	0.00027	0.00015	0.00012
white chocolate	0.00169	0.00142	0.00026	0.00207	0.00184	0.00023	0.00032	0.00007	0.00025
pink soap	0.00172	0.00155	0.00017	0.00088	0.00077	0.00012	0.00060	0.00035	0.00024
full-fat milk	0.00175	0.00083	0.00093	0.00068	0.00052	0.00016	0.00108	0.00026	0.00082
mocca yogurt drink	0.00311	0.00294	0.00017	0.00254	0.00228	0.00026	0.00052	0.00032	0.00020
blue fabric softener	0.00349	0.00314	0.00035	0.00385	0.00367	0.00017	0.00059	0.00050	0.00009
low-fat milk	0.00604	0.00567	0.00037	0.00038	0.00010	0.00028	0.00601	0.00547	0.00054
silicone mixture 2	0.00611	0.00607	0.00004	0.00073	0.00048	0.00025	0.00638	0.00628	0.00011
mean	0.00276	0.00248	0.00029	0.00130	0.00111	0.00019	0.00179	0.00151	0.00028
std	0.00213	0.00217	0.00026	0.00125	0.00124	0.00007	0.00251	0.00249	0.00024

Table A.2: Error values for the 9 fabricated replicas using the global method without local weights. The database includes the replicas generated using the global method. The first triplet of errors evaluates how well the fabricated replica matches the target. The second triplet evaluates how well our method believes it can match the targets appearance. The third triplet evaluates how well our method can predict the appearance of the replica given the local pigment parameters and the concentrations.

	Target-Replica Error			Target-Global Prediction Error			Global Prediction-Replica Error		
	Total	ρ	Shape	Total	ρ	Shape	Total	ρ	Shape
silicone mixture 1	0.00090	0.00070	0.00019	0.00054	0.00005	0.00049	0.00140	0.00097	0.00043
strawberry yogurt drink	0.00098	0.00070	0.00029	0.00068	0.00049	0.00019	0.00078	0.00062	0.00016
pink soap	0.00146	0.00125	0.00021	0.00117	0.00105	0.00011	0.00037	0.00029	0.00007
white chocolate	0.00174	0.00150	0.00024	0.00262	0.00229	0.00033	0.00081	0.00023	0.00058
full-fat milk	0.00245	0.00113	0.00131	0.00032	0.00027	0.00004	0.00321	0.00091	0.00230
mocca yogurt drink	0.00298	0.00272	0.00026	0.00354	0.00292	0.00062	0.00149	0.00106	0.00043
blue fabric softener	0.00329	0.00299	0.00030	0.00409	0.00391	0.00018	0.00077	0.00064	0.00014
low-fat milk	0.00781	0.00752	0.00029	0.00022	0.00007	0.00015	0.00804	0.00739	0.00065
silicone mixture 2	0.03639	0.03565	0.00073	0.01747	0.01346	0.00402	0.03177	0.03004	0.00173
mean	0.00644	0.00602	0.00042	0.00341	0.00272	0.00068	0.00763	0.00691	0.00072
std	0.01142	0.01131	0.00037	0.00547	0.00425	0.00126	0.01672	0.01633	0.00077




Figure A.1: Reflectance model investigation results. On the left we see reflectance plots for a semi-infinite index-matched homogeneous medium for various reduced albedo values. In total 5 reflectance models were plotted against Monte Carlo. The model that performs the best is the Flock model with a mean relative reflectance error of 2.3% followed by Kubelka Munk with a relative error of 3.31%. On the right we see the performance of our reflectance model against Monte Carlo for varying reduced albedo and optical thickness. The relative index of refraction here was set to 1.41 (silicone). To evaluate we removed half of the reduced albedo and optical thickness samples, effectively reducing the total number of samples to interpolate from by 4. Then for those missing samples we used our interpolation method to estimate reflectance. The mean relative reflectance error was found at 0.03% over the entire set of missing samples.

Appendix A

Table A.3: *Reduced scattering and absorption coefficients used for the synthetic database and targets.*

	σ_s			σ_a		
	blue	green	red	blue	green	red
White	1100	900	800	0.001	0.002	0.002
Yellow	0	85	55	25	0.4	0
Red	2	0.01	170	500	1100	10
Green	0.01	220	0.001	25	20	270
Blue	50	0.01	0	40	140	1100
Black	0.1	0	3	350	360	340
Base	1e-4	2e-6	0	0.001	0.001	0.001




Figure A.2: *Profile comparisons between two rendering methods. For each sample, path tracing was used to render the profile on the left, whereas a Quantized Diffusion fit on the path traced profile was used to render the profile on the right. Ignoring Monte Carlo noise, we observe a good match between the two methods in this set of samples.*

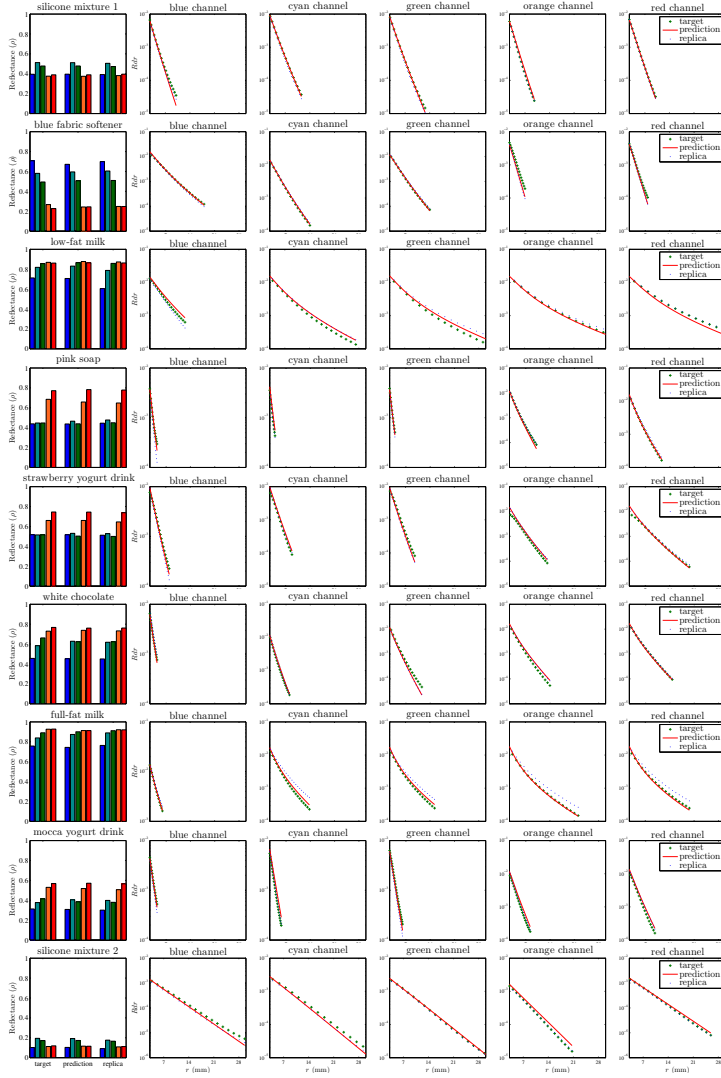


Figure A.3: Target/Prediction/Replica comparison plots for the local method.

Appendix A

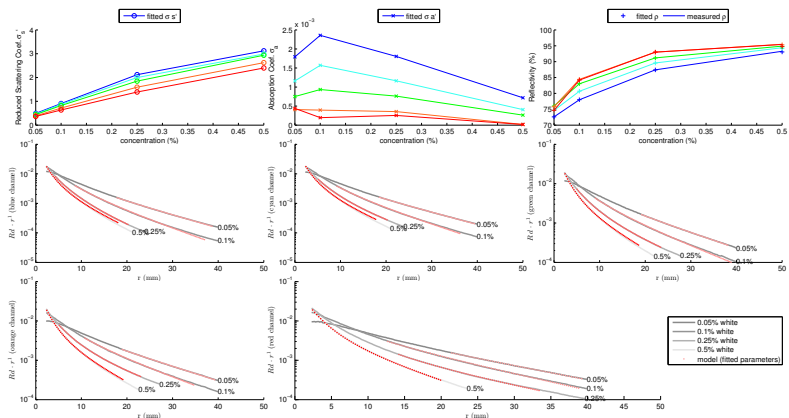


Figure A.4: White pigment dilution set measurements and forward model fits.

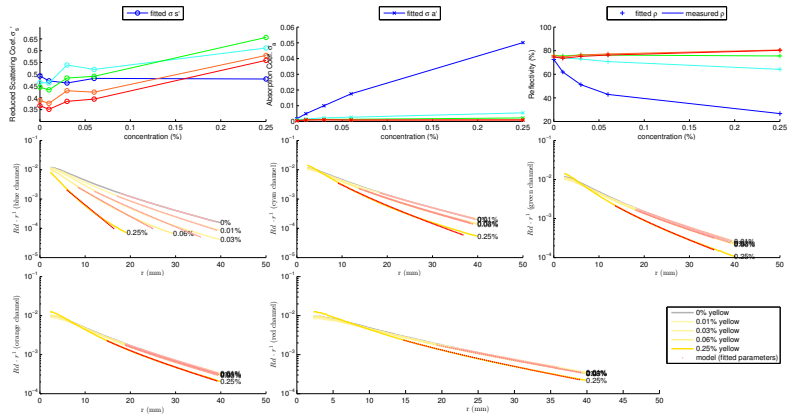


Figure A.5: Yellow pigment dilution set measurements and forward model fits.

Appendix A

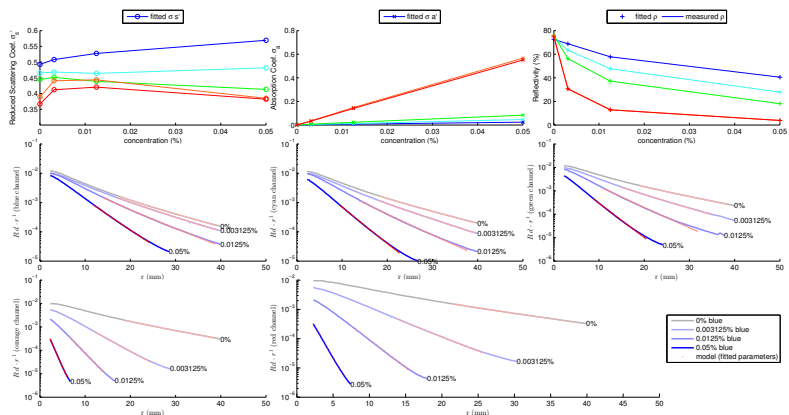


Figure A.8: Blue pigment dilution set measurements and forward model fits.

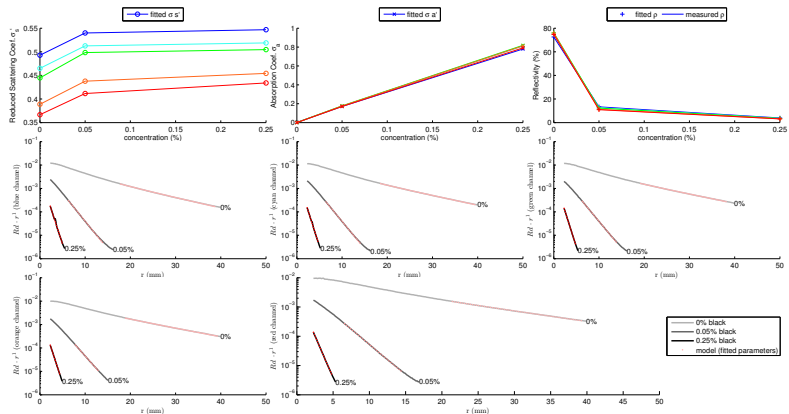


Figure A.9: Black pigment dilution set measurements and forward model fits.

APPENDIX

B

In this appendix we provide supplemental material for the steganographic lens fabrication chapter (Chapter 4).

In Figure B.1 we showcase a 3D printed source-optimized lens that can morph multiple source images to the desired targets. In Figure B.2 we simulate the effects of surface roughness using a forward appearance model. In Figures B.3 and B.4 we provide a variety of additional simulated and photographed results using our universal lens.

Finally in Figures B.5 and B.6 we provide simulated examples of image warping by using high resolution source-optimized lenses.




Figure B.1: A single manufactured lens (middle; 16×16 facets with 11×11 micro-facets) warps four source images (left) to four target images (right; photographs).




Figure B.2: We simulate the effects of surface roughness using Walter et al.'s model [2007]. We use the GGx distribution with increasing roughness (α_g) values from left to right. These simulations illustrate the type of degradation observed in our manufactured results.

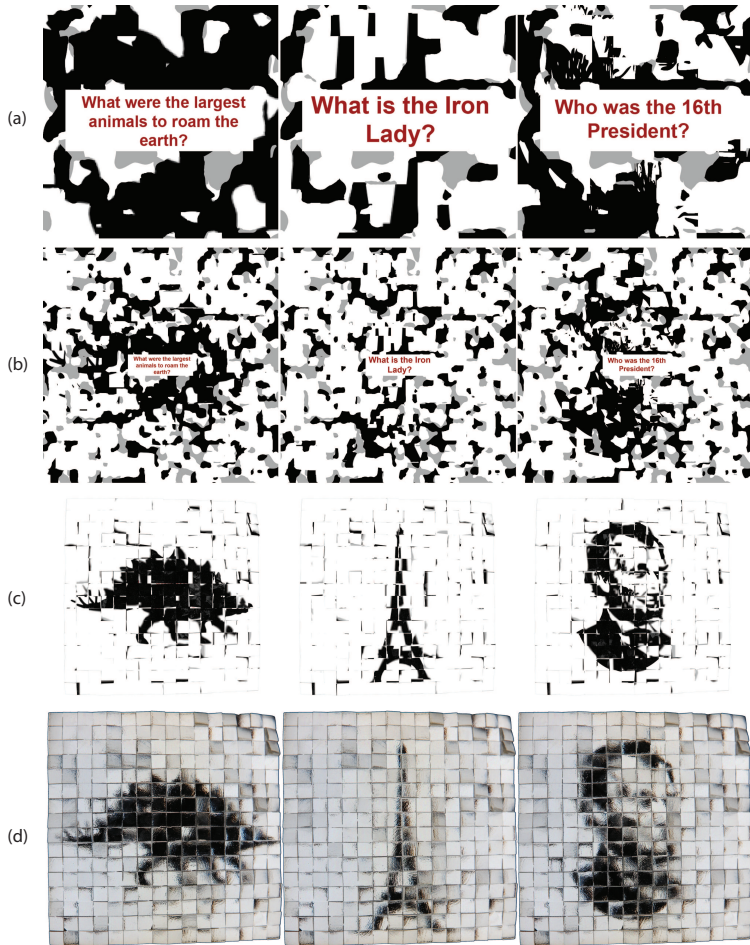


Figure B.3: A single steganographic lens warps the the source images in (b) (zoom-ins in (a)) to the target images (simulated results in (c); photos of a 3D printed lens in (d)), revealing pictographic answers to the questions in the source images. The region where there question text appears in the sources is masked and, thus, not accessed by the lens.

Appendix B

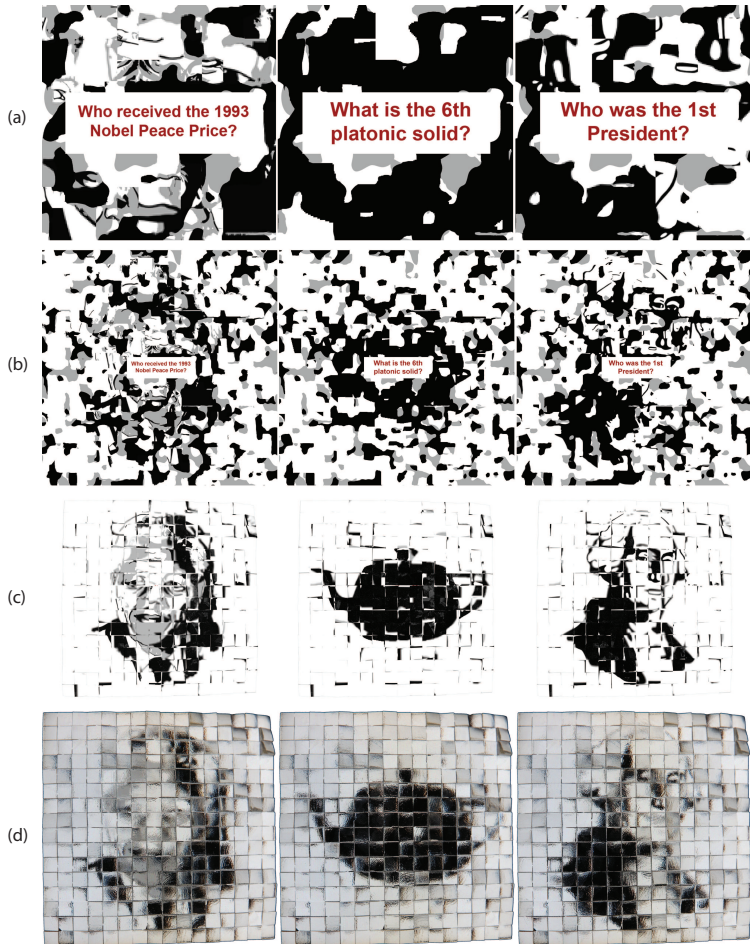


Figure B.4: A single steganographic lens warps the the source images in (b) (zoom-ins in (a)) to the target images (simulated results in (c)); photos of a 3D printed lens in (d), revealing pictographic answers to the questions in the source images. The region where there question text appears in the sources is masked and, thus, not accessed by the lens.




Figure B.5: *Morphing one image to another: with the source image (a) and target image (b) specified to our approach, the smooth lens (c) is generated and warps the source (a) to (d). This is a simulated result.*




Figure B.6: *Two more examples of morphing one image to another: for each row, the source image (a) and target image (b) are specified to our approach, the smooth lens (c) is generated and warps the source (a) to (d). These are also simulated results, however in this case the source and target images are of much higher-resolution. The resulting lens has 128×128 facets with 11×11 micro-facets, illustrating the scalability of our approach.*

Appendix B

References

- [Alasia 1976] Alfred V. Alasia. Process of coding indicia and product produced thereby. U.S. Patent Number 3937565, 1976. Filed Jun 3, 1974.
- [Alasia 1998] Alfred V. Alasia. Digital anti-counterfeiting software method and apparatus. U.S. Patent Number 5708717, 1998. Filed Nov 29, 1995.
- [Alexa and Matusik 2010] Marc Alexa and Wojciech Matusik. Reliefs as images. *ACM Trans. on Graphics*, 29(4):60:1–60:7, July 2010.
- [Arbree *et al.* 2011] A. Arbree, B. Walter, and K. Bala. Heterogeneous subsurface scattering using the finite element method. *IEEE Trans. on Visualization and Computer Graphics*, 17(7):956–969, July 2011.
- [Ashikmin *et al.* 2000] Michael Ashikmin, Simon Premože, and Peter Shirley. A microfacet-based BRDF generator. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '00, pages 65–74, New York, NY, USA, 2000. ACM Press/Addison-Wesley Publishing Co.
- [Baran *et al.* 2012] Ilya Baran, Philipp Keller, Derek Bradley, Stelian

References

- Coros, Wojciech Jarosz, Derek Nowrouzezahrai, and Markus Gross. Manufacturing layered attenuators for multiple prescribed shadow images. *Computer Graphics Forum (Proc. Eurographics)*, 31(2), 2012.
- [Barnes *et al.* 2009] Connelly Barnes, Eli Shechtman, Adam Finkelstein, and Dan B Goldman. PatchMatch: A randomized correspondence algorithm for structural image editing. *ACM Trans. on Graphics (Proc. SIGGRAPH)*, 28(3), August 2009.
- [Barnes *et al.* 2010] Connelly Barnes, Eli Shechtman, Dan B Goldman, and Adam Finkelstein. The generalized PatchMatch correspondence algorithm. In *European Conference on Computer Vision*, volume 6313 of *Lecture Notes in Computer Science*. Springer, September 2010.
- [Barnes *et al.* 2011] Connelly Barnes, Dan B Goldman, Eli Shechtman, and Adam Finkelstein. The PatchMatch randomized matching algorithm for image manipulation. *Communications of the ACM*, 54(11):103–110, November 2011.
- [Barnes 2011] Connelly Barnes. *PatchMatch: A Fast Randomized Matching Algorithm with Application to Image and Video*. PhD thesis, Princeton University, May 2011.
- [Bermano *et al.* 2012] Amit Bermano, Ilya Baran, Marc Alexa, and Wojciech Matusik. SHADOWPIX: Multiple Images from Self Shadowing. *Computer Graphics Forum*, 2012.
- [Brady *et al.* 2014] Adam Brady, Jason Lawrence, Pieter Peers, and Westley Weimer. genBRDF: Discovering new analytic BRDFs with genetic programming. *ACM Trans. on Graphics*, 33(4):114:1–114:11, July 2014.
- [Brosh and Wright 1994] Scott Brosh and Timothy Wright. Anti-counterfeiting process using lenticular optics and color masking. U.S. Patent Number 5303370, 1994. Filed Nov 13, 1992.

- [Brown *et al.* 2002] Richard JC Brown, Paul J Brewer, and Martin JT Milton. The physical and chemical properties of electroless nickel–phosphorus alloys and low reflectance nickel–phosphorus black surfaces. *Journal of Materials Chemistry*, 12(9):2749–2754, 2002.
- [Bruneton and Neyret 2012] Eric Bruneton and Fabrice Neyret. A survey of non-linear pre-filtering methods for efficient and accurate surface shading. *IEEE Trans. on Visualization and Computer Graphics*, 18(2):242–260, 2012.
- [Cerezo *et al.* 2005] Eva Cerezo, Frederic Pérez, Xavier Pueyo, Francisco J. Seron, and François X. Sillion. A survey on participating media rendering techniques. *The Visual Computer*, 21(5):303–328, 2005.
- [Chandrasekar 1960] S. Chandrasekar. *Radiative Transfer*. Dover Publications, 1960.
- [Christensen *et al.* 2012] Per H. Christensen, George Harker, Jonathan Shade, Brenden Schubert, and Dana Batali. Multiresolution radiosity caching for global illumination in movies. In *ACM SIGGRAPH Talks*, 2012.
- [Chu *et al.* 2010] Hung-Kuo Chu, Wei-Hsin Hsu, Niloy J. Mitra, Daniel Cohen-Or, Tien-Tsin Wong, and Tong-Yee Lee. Camouflage images. *ACM Trans. on Graphics*, 29(4):51:1–51:8, July 2010.
- [Coffin 2004] Dave Coffin. DCRAW, 2004. <http://www.cybercom.net/dcoffin/dcraw/>.
- [Cook and Torrance 1981] Robert L. Cook and Kenneth E. Torrance. A reflectance model for computer graphics. In *Computer Graphics (Proceedings of SIGGRAPH 81)*, pages 307–316, August 1981.
- [d’Eon and Irving 2011] Eugene d’Eon and Geoffrey Irving. A quantized-diffusion model for rendering translucent materials. *ACM Trans. on Graphics (Proc. SIGGRAPH)*, 30(4):56:1–56:14, 2011.

References

- [d'Eon 2014] Eugene d'Eon. Rigorous asymptotic and moment-preserving diffusion approximations for generalized linear Boltzmann transport in arbitrary dimensions. *Journal of Computational and Theoretical Transport*, page to appear, 2014.
- [Dong *et al.* 2010] Yue Dong, Jiaping Wang, Fabio Pellacini, Xin Tong, and Baining Guo. Fabricating spatially-varying subsurface scattering. *ACM Trans. on Graphics (Proc. SIGGRAPH)*, 29(4):62:1–62:10, July 2010.
- [Donner and Jensen 2005] Craig Donner and Henrik Wann Jensen. Light diffusion in multi-layered translucent materials. *ACM Trans. on Graphics (Proc. SIGGRAPH)*, 24(3):1032–1039, 2005.
- [Donovan *et al.* 2003] T.J. Donovan, T.M. Sutton, and Y. Danon. Implementation of chord length sampling for transport through a binary stochastic mixture. In *Nuclear Mathematical and Computational Sciences: A Century in Review, A Century Anew*, 2003.
- [Dorsey *et al.* 2008] Julie Dorsey, Holly Rushmeier, and François Sillion. *Digital Modeling of Material Appearance*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2008.
- [Durant *et al.* 2007] Stéphane Durant, Olivier Calvo-Perez, Nicolas Vukadinovic, and Jean-Jacques Greffet. Light scattering by a random distribution of particles embedded in absorbing media: full-wave Monte Carlo solutions of the extinction coefficient. *Journal of the Optical Society of America*, 24(9):2953–2962, 2007.
- [Finckh *et al.* 2010] Manuel Finckh, Holger Dammertz, and Hendrik Lensch. Geometry construction from caustic images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2010.
- [Foldy 1945] Leslie L. Foldy. The multiple scattering of waves. I. General theory of isotropic scattering by randomly distributed scatterers. *Physical Review*, 67:107–119, 1945.

- [Fuchs *et al.* 2008] Martin Fuchs, Ramesh Raskar, Hans-Peter Seidel, and Hendrik P. A. Lensch. Towards passive 6D reflectance field displays. *ACM Trans. on Graphics (Proc. SIGGRAPH)*, 27(3):58:1–58:8, August 2008.
- [Georgiev *et al.* 2013] Iliyan Georgiev, Jaroslav Krivánek, Toshiya Hachisuka, Derek Nowrouzezahrai, and Wojciech Jarosz. Joint importance sampling of low-order volumetric scattering. *ACM Trans. on Graphics (Proc. SIGGRAPH Asia)*, 32(6), November 2013.
- [Gkioulekas *et al.* 2013] Ioannis Gkioulekas, Shuang Zhao, Kavita Bala, Todd Zickler, and Anat Levin. Inverse volume rendering with material dictionaries. *ACM Trans. on Graphics*, 32(6):162:1–162:13, November 2013.
- [Gotoda 2010] Hironobu Gotoda. A multilayer liquid crystal display for autostereoscopic 3d viewing. *Stereoscopic Displays and Applications XXI*, 7524:1–8, 2010.
- [Habel *et al.* 2013] Ralf Habel, Per H. Christensen, and Wojciech Jarosz. Photon beam diffusion: A hybrid Monte Carlo method for subsurface scattering. *Computer Graphics Forum (Proc. Eurographics Symposium on Rendering)*, 32(4), 2013.
- [Hasan *et al.* 2010] Milos Hasan, Martin Fuchs, Wojciech Matusik, Hanspeter Pfister, and Szymon Rusinkiewicz. Physical reproduction of materials with specified subsurface scattering. *ACM Trans. on Graphics (Proc. SIGGRAPH)*, 29(4):61:1–61:10, July 2010.
- [Hawkins *et al.* 2005] Tim Hawkins, Per Einarsson, and Paul Debevec. Acquisition of time-varying participating media. *ACM Trans. on Graphics (Proc. SIGGRAPH)*, 24(3):812–815, August 2005.
- [Henyey and Greenstein 1941] L. G. Henyey and J. L. Greenstein. Diffuse radiation in the galaxy. *The Astrophysical Journal*, 93:70–83, 1941.

References

- [Hersch and Chosson 2004] Roger David Hersch and Sylvain Chosson. Band moiré images. *ACM Trans. on Graphics (Proc. SIGGRAPH)*, 23(3):239–247, August 2004.
- [Hersch *et al.* 2003] Roger D. Hersch, Fabien Collaud, and Patrick Emmel. Reproducing color images with embedded metallic patterns. *ACM Trans. on Graphics (Proc. SIGGRAPH)*, 22(3), July 2003.
- [Hersch *et al.* 2007] Roger D. Hersch, Philipp Donzé, and Sylvain Chosson. Color images visible under UV light. *ACM Trans. on Graphics*, 26(3):75:1–75:9, July 2007.
- [Hery 2012] Christophe Hery. Texture mapping for the better dipole model. Technical Report 12-11, Pixar, 2012.
- [Hullin *et al.* 2011] Matthias B. Hullin, Hendrik P. A. Lensch, Ramesh Raskar, Hans-Peter Seidel, and Ivo Ihrke. Dynamic display of BRDFs. In *Computer Graphics Forum (Proc. Eurographics)*, pages 475–483, 2011.
- [Jakob and Marschner 2012] Wenzel Jakob and Steve Marschner. Manifold exploration: A markov chain Monte Carlo technique for rendering scenes with difficult specular transport. *ACM Trans. on Graphics*, 31(4):58:1–58:13, July 2012.
- [Jakob *et al.* 2009] Wenzel Jakob, Jonathan T. Moon, and Steve Marschner. Capturing hair assemblies fiber by fiber. *ACM Trans. on Graphics*, 28(5):164:1–164:9, December 2009.
- [Jakob *et al.* 2014] Wenzel Jakob, Eugene d’Eon, Otto Jakob, and Steve Marschner. A comprehensive framework for rendering layered materials. *ACM Trans. on Graphics*, 33(4):118:1–118:14, July 2014.
- [Jakob 2010] Wenzel Jakob. Mitsuba renderer, 2010. <http://mitsuba-renderer.org>.

- [Jensen and Buhler 2002] Henrik Wann Jensen and Juan Buhler. A rapid hierarchical rendering technique for translucent materials. *ACM Trans. on Graphics (Proc. SIGGRAPH)*, 21(3):576–581, July 2002.
- [Jensen *et al.* 2001] Henrik Wann Jensen, Stephen R. Marschner, Marc Levoy, and Pat Hanrahan. A practical model for subsurface light transport. *Computer Graphics (Proc. SIGGRAPH)*, 35:511–518, 2001.
- [Kajiya and Kay 1989] James T. Kajiya and Timothy L. Kay. Rendering fur with three dimensional textures. In *Computer Graphics (Proc. SIGGRAPH)*, pages 271–280, 1989.
- [Kajiya 1986] James T. Kajiya. The rendering equation. *Computer Graphics (Proc. SIGGRAPH)*, 20:143–150, 1986.
- [Kelly 1987] Renee J. Kelly. Process for matching color of paint to a colored surface. U.S. Patent Number 4692481, 1987. Filed Sep 27, 1984.
- [Klehm *et al.* 2015] Oliver Klehm, Fabrice Rousselle, Marios Papas, Derek Bradley, Christophe Hery, Bernd Bickel, Wojciech Jarosz, and Thabo Beeler. Recent advances in facial appearance capture. *Computer Graphics Forum (Proc. Eurographics)*, 34(2):709–733, May 2015.
- [Křivánek *et al.* 2014] Jaroslav Křivánek, Iliyan Georgiev, Toshiya Hachisuka, Petr Vévoda, Martin Šik, Derek Nowrouzezahrai, and Wojciech Jarosz. Unifying points, beams, and paths in volumetric light transport simulation. *ACM Trans. on Graphics (Proc. SIGGRAPH)*, 33(4), aug 2014.
- [Lafortune and Willems 1996] Eric P. Lafortune and Yves D. Willems. Rendering participating media with bidirectional path tracing. In *Proc. Eurographics Workshop on Rendering*, pages 91–100, 1996.
- [Lan *et al.* 2013] Yanxiang Lan, Yue Dong, Fabio Pellacini, and Xin

References

- Tong. Bi-scale appearance fabrication. *ACM Trans. on Graphics*, 32(4):145:1–145:12, July 2013.
- [Lanman *et al.* 2011] Douglas Lanman, Gordon Wetzstein, Matthew Hirsch, Ramesh Raskar, and Wolfgang Heidrich. Polarization fields: Dynamic light field display using multi-layer LCDs. *ACM Trans. on Graphics (Proc. SIGGRAPH Asia)*, 30(6), December 2011.
- [Li *et al.* 2005] Hongsong Li, Fabio Pellacini, and Kenneth E. Torrance. A hybrid Monte Carlo method for accurate and efficient subsurface scattering. In *Rendering Techniques*, pages 283–290. Eurographics Association, 2005.
- [Lippmann 1908] Jonas Ferdinand Gabriel Lippmann. Epreuves reversibles donnant la sensation du relief. *Journal of Physics*, 1908.
- [Lowe 1999] David G. Lowe. Object recognition from local scale-invariant features. In *The Proceedings of the Seventh IEEE International Conference on Computer Vision*, volume 2, pages 1150–1157, 1999.
- [Luebke *et al.* 2002] David Luebke, Benjamin Watson, Jonathan D. Cohen, Martin Reddy, and Amitabh Varshney. *Level of Detail for 3D Graphics*. Elsevier Science Inc., 2002.
- [Malzbender *et al.* 2012] Tom Malzbender, Ramin Samadani, Steven Scher, Adam Crume, Douglas Dunn, and James Davis. Printing reflectance functions. *ACM Trans. on Graphics*, 31(3):20:1–20:11, June 2012.
- [Marschner *et al.* 2003] Stephen R. Marschner, Henrik Wann Jensen, Mike Cammarano, Steve Worley, and Pat Hanrahan. Light scattering from human hair fibers. *ACM Trans. on Graphics*, 22(3):780–791, July 2003.
- [Marschner *et al.* 2005] Stephen R. Marschner, Stephen H. Westin, Adam Arbree, and Jonathan T. Moon. Measuring and modeling the

- appearance of finished wood. *ACM Trans. on Graphics*, 24(3):727–734, July 2005.
- [Matusik *et al.* 2003] Wojciech Matusik, Hanspeter Pfister, Matt Brand, and Leonard McMillan. A data-driven reflectance model. *ACM Trans. on Graphics (Proc. SIGGRAPH)*, 22(3):759–769, 2003.
- [Meng *et al.* 2015] Johannes Meng, Marios Papas, Ralf Habel, Carsten Dachsbacher, Steve Marschner, Markus Gross, and Wojciech Jarosz. Multi-scale modeling and rendering of granular materials. *ACM Trans. on Graphics (Proc. SIGGRAPH)*, 34(4), July 2015.
- [Mitra and Pauly 2009] Niloy J. Mitra and Mark Pauly. Shadow art. *ACM Trans. on Graphics (Proc. SIGGRAPH Asia)*, 28(5), 2009.
- [Mitra *et al.* 2009] Niloy J. Mitra, Hung-Kuo Chu, Tong-Yee Lee, Lior Wolf, Hezy Yeshurun, and Daniel Cohen-Or. Emerging images. *ACM Trans. on Graphics*, 28(5):163:1–163:8, December 2009.
- [Mitsunaga and Nayar 1999] T. Mitsunaga and S.K. Nayar. Radiometric self calibration. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 1, pages 374–380, Jun 1999.
- [Moon and Marschner 2006] Jonathan T. Moon and Stephen R. Marschner. Simulating multiple scattering in hair using a photon mapping approach. *ACM Trans. on Graphics (Proc. SIGGRAPH)*, 25(3):1067–1074, 2006.
- [Moon *et al.* 2007] Jonathan T. Moon, Bruce Walter, and Stephen R. Marschner. Rendering discrete random media using precomputed scattering solutions. In *Proc. Eurographics Symposium on Rendering*, pages 231–242, 2007.
- [Naor and Shamir 1994] Moni Naor and Adi Shamir. Visual cryptography. In Alfredo De Santis, editor, *Advances in Cryptology, EUROCRYPT*, volume 950 of *Lecture Notes in Computer Science*, pages 1–12. Springer Berlin / Heidelberg, 1994.

References

- [Narasimhan *et al.* 2006] Srinivasa G. Narasimhan, Mohit Gupta, Craig Donner, Ravi Ramamoorthi, Shree K. Nayar, and Henrik Wann Jensen. Acquiring scattering properties of participating media by dilution. *ACM Trans. on Graphics (Proc. SIGGRAPH)*, 25(3):1003–1012, July 2006.
- [Ngan *et al.* 2005] Addy Ngan, Frédo Durand, and Wojciech Matusik. Experimental analysis of BRDF models. In *Proceedings of the Sixteenth Eurographics Conference on Rendering Techniques, EGSR '05*, pages 117–126, Aire-la-Ville, Switzerland, Switzerland, 2005. Eurographics Association.
- [Novák *et al.* 2012a] Jan Novák, Derek Nowrouzezahrai, Carsten Dachsbacher, and Wojciech Jarosz. Progressive virtual beam lights. *Computer Graphics Forum (Proc. Eurographics Symposium on Rendering)*, 31(4), June 2012.
- [Novák *et al.* 2012b] Jan Novák, Derek Nowrouzezahrai, Carsten Dachsbacher, and Wojciech Jarosz. Virtual ray lights for rendering scenes with participating media. *ACM Trans. on Graphics (Proc. SIGGRAPH)*, 31(4), July 2012.
- [Oren and Nayar 1994] Michael Oren and Shree K. Nayar. Generalization of lambert’s reflectance model. In *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '94*, pages 239–246, New York, NY, USA, 1994. ACM.
- [Papas *et al.* 2011] Marios Papas, Wojciech Jarosz, Wenzel Jakob, Szymon Rusinkiewicz, Wojciech Matusik, and Tim Weyrich. Goal-based caustics. *Computer Graphics Forum (Proc. Eurographics)*, 30(2), 2011.
- [Papas *et al.* 2013] Marios Papas, Christian Regg, Wojciech Jarosz, Bernd Bickel, Philip Jackson, Wojciech Matusik, Steve Marschner, and Markus Gross. Fabricating translucent materials using contin-

- uous pigment mixtures. *ACM Trans. on Graphics (Proc. SIGGRAPH)*, 32(4), July 2013.
- [Paris *et al.* 2008] Sylvain Paris, Will Chang, Oleg I. Kozhushnyan, Wojciech Jarosz, Wojciech Matusik, Matthias Zwicker, and Frédéric Durand. Hair photobooth: Geometric and photometric acquisition of real hairstyles. *ACM Trans. on Graphics (Proc. SIGGRAPH)*, 27(3):30:1–30:9, August 2008.
- [Park *et al.* 2007] Jong-Il Park, Moon-Hyun Lee, Michael D. Grossberg, and Shree K. Nayar. Multispectral imaging using multiplexed illumination. In *IEEE International Conference on Computer Vision (ICCV)*, pages 1–8, October 2007.
- [Peers *et al.* 2006] Pieter Peers, Karl vom Berge, Wojciech Matusik, Ravi Ramamoorthi, Jason Lawrence, Szymon Rusinkiewicz, and Philip Dutré. A compact factored representation of heterogeneous subsurface scattering. *ACM Trans. on Graphics*, 25(3):746–753, July 2006.
- [Pharr and Humphreys 2010] Matt Pharr and Greg Humphreys. *Physically Based Rendering, Second Edition: From Theory to Implementation*. Morgan Kaufmann Publishers Inc., 2nd edition, 2010.
- [Pritchard and Heidrich 2003] D. Pritchard and W. Heidrich. Cloth motion capture. *Computer Graphics Forum*, 22(3):263–271, September 2003.
- [Randrianalisoa and Baillis 2009] Jaona Randrianalisoa and Dominique Baillis. Radiative transfer in dispersed media: Comparison between homogeneous phase and multiphase approaches. *Journal of Heat Transfer*, 132(2):023405–023405, 2009.
- [Randrianalisoa and Baillis 2010] Jaona Randrianalisoa and Dominique Baillis. Radiative properties of densely packed spheres in

References

- semitransparent media: A new geometric optics approach. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 111(10):1372 – 1388, 2010.
- [Schmidt *et al.* 2014] Thorsten-Walther Schmidt, Fabio Pellacini, Derek Nowrouzezahrai, Wojciech Jarosz, and Carsten Dachsbacher. State of the art in artistic editing of appearance, lighting, and material. In *Eurographics 2014 - State of the Art Reports*, Strasbourg, France, April 2014. Eurographics Association.
- [Schröder *et al.* 2011] Kai Schröder, Reinhard Klein, and Arno Zinke. A volumetric approach to predictive rendering of fabrics. *Computer Graphics Forum (Proc. Eurographics Symposium on Rendering)*, 30(4):1277–1286, 2011.
- [Schwartzburg *et al.* 2014] Yuliy Schwartzburg, Romain Testuz, Andrea Tagliasacchi, and Mark Pauly. High-contrast computational caustic design. *ACM Trans. on Graphics*, 33(4):74:1–74:11, July 2014.
- [Sharma *et al.* 2005] Gaurav Sharma, Wencheng Wu, and Edul N. Dalal. The CIEDE2000 color-difference formula: Implementation notes, supplementary test data, and mathematical observations. *Color Research and Application*, 30(1):21–30, 2005.
- [Sherman and Simone 1989] Charles J. Sherman and Kenneth S. Simone. Process for manufacturing paints. U.S. Patent Number 4887217, 1989. Filed Jan 4, 1985.
- [Singh and Kaviany 1992] B.P. Singh and M. Kaviany. Modelling radiative heat transfer in packed beds. *International Journal of Heat and Mass Transfer*, 35(6):1397 – 1405, 1992.
- [Song *et al.* 2009] Ying Song, Xin Tong, Fabio Pellacini, and Pieter Peers. Subedit: A representation for editing measured heterogeneous subsurface scattering. *ACM Trans. on Graphics*, 28(3):31:1–31:10, July 2009.

- [Stam 1995] Jos Stam. Multiple scattering as a diffusion process. *Proc. Eurographics Workshop on Rendering*, pages 41–50, 1995.
- [Theocharous *et al.* 2014] Evangelos Theocharous, Christopher J. Chunnillall, Ryan Mole, David Gibbs, Nigel Fox, Naigui Shang, Guy Howlett, Ben Jensen, Rosie Taylor, Juan R. Reveles, Oliver B. Harris, and Naseer Ahmed. The partial space qualification of a vertically aligned carbon nanotube coating on aluminium substrates for EO applications. *Opt. Express*, 22(6):7290–7307, Mar 2014.
- [Tong *et al.* 2005] Xin Tong, Jiaping Wang, Stephen Lin, Baining Guo, and Heung-Yeung Shum. Modeling and rendering of quasi-homogeneous materials. *ACM Trans. on Graphics (Proc. SIGGRAPH)*, 24(3):1054–1061, 2005.
- [Torquato 2001] S. Torquato. *Random Heterogeneous Materials: Microstructure and Macroscopic Properties*. Interdisciplinary Applied Mathematics. Springer, 2001.
- [Torrance and Sparrow 1967] K. E. Torrance and E. M. Sparrow. Theory for off-specular reflection from roughened surfaces. *Journal of the Optical Society of America*, 57(9):1105–1112, 1967.
- [Tyler and Clarke 1990] Christopher W. Tyler and Maureen B. Clarke. The autostereogram. In *Proceedings of SPIE*, volume 1256 of *Stereoscopic Displays and Applications*, pages 182–196, 1990.
- [Walter *et al.* 2007] Bruce Walter, Stephen R. Marschner, Hongsong Li, and Kenneth E. Torrance. Microfacet models for refraction through rough surfaces. In *Rendering Techniques 2007: 18th Eurographics Workshop on Rendering*, pages 195–206, June 2007.
- [Walter *et al.* 2009] Bruce Walter, Shuang Zhao, Nicolas Holzschuch, and Kavita Bala. Single scattering in refractive media with triangle mesh boundaries. *ACM Trans. on Graphics*, 28(3), aug 2009.

References

- [Wang *et al.* 1995] Lihong Wang, Steven Jacques, and Liqiong Zheng. MCML: Monte Carlo modeling of light transport in multi-layered tissues. *Computer Methods and Programs in Biomedicine*, (8):313–371, July 1995.
- [Ward *et al.* 2006] Kelly Ward, Florence Bertails, Tae yong Kim, Stephen R. Marschner, Marie paule Cani, and Ming C. Lin. A survey on hair modeling: styling, simulation, and rendering. In *IEEE Trans. on Visualization and Computer Graphics*, pages 213–234, 2006.
- [Wetzstein *et al.* 2011] Gordon Wetzstein, Douglas Lanman, Wolfgang Heidrich, and Ramesh Raskar. Layered 3D: Tomographic image synthesis for attenuation-based light field and high dynamic range displays. *ACM Trans. on Graphics (Proc. SIGGRAPH)*, 30(4), 2011.
- [Weyrich *et al.* 2006] Tim Weyrich, Wojciech Matusik, Hanspeter Pfister, Bernd Bickel, Craig Donner, Chien Tu, Janet McAndless, Jinho Lee, Addy Ngan, Henrik Wann Jensen, and Markus Gross. Analysis of human faces using a measurement-based skin reflectance model. *ACM Trans. on Graphics (Proc. SIGGRAPH)*, 25(3):1013–1024, July 2006.
- [Weyrich *et al.* 2009a] Tim Weyrich, Jason Lawrence, Hendrik P.A. Lensch, Szymon Rusinkiewicz, and Todd Zickler. Principles of appearance acquisition and representation. *Foundations and Trends in Computer Graphics and Vision*, 4(2):75–191, October 2009.
- [Weyrich *et al.* 2009b] Tim Weyrich, Pieter Peers, Wojciech Matusik, and Szymon Rusinkiewicz. Fabricating microgeometry for custom surface reflectance. *ACM Trans. on Graphics (Proc. SIGGRAPH)*, 28(3):32:1–32:6, July 2009.
- [Wu *et al.* 2013] Hongzhi Wu, Julie Dorsey, and Holly Rushmeier. Inverse bi-scale material design. *ACM Trans. on Graphics (Proc. SIGGRAPH Asia)*, 32, 2013.

- [Xu *et al.* 2007] Kun Xu, Yue Gao, Yong Li, Tao Ju, and Shi-Min Hu. Real-time homogenous translucent material editing. *Computer Graphics Forum (Proc. Eurographics)*, 26(3):545–552, September 2007.
- [Yue *et al.* 2012] Yonghao Yue, Kei Iwasaki, Bing-Yu Chen, Yoshinori Dobashi, , and Tomoyuki Nishita. Pixel art with refracted light by rearrangeable sticks. *Computer Graphics Forum (Proc. Eurographics)*, 31(2), 2012.
- [Zhao *et al.* 2013] Shuang Zhao, Miloš Hašan, Ravi Ramamoorthi, and Kavita Bala. Modular flux transfer: efficient rendering of high-resolution volumes with repeated structures. *ACM Trans. on Graphics (Proc. SIGGRAPH)*, 32(4):131:1–131:12, 2013.
- [Zinke and Weber 2006] Arno Zinke and Andreas Weber. Global illumination for fiber based geometries. In *Proc. Ibero-American Symposium in Computer Graphics (SIACG)*, 2006.
- [Zinke *et al.* 2008] Arno Zinke, Cem Yuksel, Andreas Weber, and John Keyser. Dual scattering approximation for fast multiple scattering in hair. *ACM Trans. on Graphics (Proc. SIGGRAPH)*, 27(3):32:1–32:10, 2008.
- [Zinke *et al.* 2009] Arno Zinke, Martin Rump, Tomás Lay, Andreas Weber, Anton Andriyenko, and Reinhard Klein. A practical approach for photometric acquisition of hair color. *ACM Trans. on Graphics*, 28(5):165:1–165:9, December 2009.