

UNIVERSITÄT
DES
SAARLANDES

Data-driven methods for high-fidelity material appearance reproduction in additive 3D printing

Dissertation zur Erlangung des Grades
des Doktors der Ingenieurwissenschaften
der Fakultät für Mathematik und Informatik
der Universität des Saarlandes

Vorgelegt von
Denis Sumin

Saarbrücken, 2024

Dean:	Prof. Dr. Roland Speicher
Date:	February 18th, 2025
Committee members	
Chair:	Prof. Dr. Thorsten Herfet
Reviewers:	Dr.-Ing. habil. Karol Myszkowski
	Prof. Dr. Hans-Peter Seidel
	Prof. Dr. Jeppe Rival Frisvad
Academic Assistant:	Dr. Mojtaba Bemana

Abstract

In this thesis, we address several key challenges in the realm of high-fidelity material-appearance reproduction. Despite significant recent advances in full-color 3D printing, reproducing many potential appearances remains a challenge due to the inherent translucency and scattering properties of commonly used print materials. This thesis develops a comprehensive 3D-printing preparation system tailored for multi-material inkjet 3D printers, focusing on overcoming these limitations.

Our work begins by highlighting the limitations of current methods, which often result in blurred details and color bleeding due to subsurface scattering. We introduce an innovative iterative optimization approach that refines the volumetric arrangement of print materials to counteract these effects, enabling crisp and accurate color texture reproduction even in highly translucent media.

To validate our method, we employ a general Monte Carlo simulation for light transport, supported by a practical calibration procedure to obtain scattering parameters of printer materials. Despite the medium's inherent translucency, objects fabricated using our system demonstrate superior detail preservation and color fidelity compared to native color texturing modes in commercial printers. The approach is applicable to general 3D geometries and effectively handles the challenging case of thin features, where material translucency and limited material deposition further complicate accurate reproduction.

We propose a data-driven approach using deep neural networks to predict light scattering within heterogeneous materials. This method offers a dramatic speedup, performing around two orders of magnitude faster than traditional Monte Carlo simulations while maintaining high-quality optimization results. This breakthrough enables full heterogeneous material optimization within practical time frames, making it feasible to execute within the actual printing duration.

Furthermore, we evaluate existing image quality metrics in their application to light field images to support future work on the fabrication of angularly and spatially varying appearances. We develop a new dense light-field dataset and derive subjective quality scaling through perceptual experiments, addressing typical artifacts arising from light-field processing. Our evaluation of existing quality metrics reveals the need for light-field-specific metrics to better predict perceived quality in complex tasks, and the developed dataset will aid in the future development of these specialized metrics.

Zusammenfassung

In dieser Dissertation befassen wir uns mit mehreren schwerwiegenden Komplikationen im Bereich der hochpräzisen Reproduktion vom Aussehen von Materialien. Trotz merklicher Fortschritte im Vollfarb-3D-Druck bleibt die Reproduktion vieler Erscheinungen aufgrund der inhärenten Transluzenz und Streueigenschaften der häufig verwendeten Druckmaterialien bisher ungelöst. Diese Arbeit entwickelt ein umfassendes Vorbereitungssystem für den 3D-Druck, das speziell auf 3D-Mehrmaterial-Tintenstrahldrucker zugeschnitten ist und das Ziel hat, diese Schwierigkeiten zu überwinden.

Unsere Arbeit beginnt mit dem Aufzeigen der Grenzen aktueller Methoden, die oft zu verschwommenen Details und Farbverläufen aufgrund von Streuung unterhalb der Oberfläche führen. Wir stellen einen innovativen iterativen Optimierungsansatz vor, der die volumetrische Anordnung der Druckmaterialien verfeinert um diesen Effekten entgegenzuwirken und eine klare und genaue Farbwiedergabe auch in hochtransluzenten Medien zu ermöglichen.

Zur Validierung unserer Methode verwenden wir eine allgemeine Monte-Carlo-Simulation für den Lichttransport, die durch ein praktisches Kalibrierungsverfahren zur Ermittlung der Streuparameter der Druckmaterialien unterstützt wird. Trotz der inhärenten Transluzenz des Mediums zeigen Objekte, die mit unserem System gefertigt wurden, höhere Detailgenauigkeit und bessere Farbtreue im Vergleich zu den nativen Farbtexturierungsmodi kommerzieller Drucker. Der Ansatz ist auf allgemeine 3D-Geometrien anwendbar und bewältigt sogar die Hürde dünner Strukturen, bei denen die Materialtransluzenz und der begrenzte Materialauftrag die genaue Reproduktion weiter erschweren.

Wir schlagen einen datenbasierten Ansatz unter Verwendung tiefer neuronaler Netzwerke vor, um die Lichtstreuung in heterogenen Materialien vorherzusagen. Diese Methode bietet eine erhebliche Beschleunigung und ist etwa zwei Größenordnungen schneller als herkömmliche Monte-Carlo-Simulationen, während sie gleichzeitig hochwertige Optimierungsergebnisse beibehält. Dieser Durchbruch ermöglicht eine vollständige heterogene Materialoptimierung in praktikablen Zeitrahmen, sogar während des 3D-Drucks.

Darüber hinaus bewerten wir bestehende Bildqualitätsmetriken in ihrer Anwendung auf Lichtfeldbilder, um zukünftige Arbeiten zur Herstellung von winklig und räumlich variierenden Erscheinungen zu unterstützen. Wir entwickeln einen neuen dichten Lichtfeld-Datensatz und leiten per menschlicher Wahrnehmung eine Qualitätsskalierung ab, die typische Bildfehler der Lichtfeldverarbeitung adressiert. Unsere Bewertung der bestehenden Qualitätsmetriken zeigt die Notwendigkeit lichtfeldspezifischer Metriken, um die wahrgenommene Qualität in komplexen Aufgaben besser vorherzusagen; der hier entwickelte Datensatz wird die zukünftige Entwicklung dieser Metriken unterstützen.

Acknowledgments

I would like to thank everyone who supported me during my Ph.D. studies and the process of writing this thesis.

First and foremost, I am deeply grateful to Karol Myszkowski for his visionary guidance and advice throughout the main phase of my research work at MPI, and for his infinite support during the long-lasting preparation of this thesis.

I want to thank Hans-Peter Seidel for creating the computer graphics group, which has been an excellent place to work, providing an inspirational and welcoming environment with numerous research opportunities. I am grateful to Ellen Fries and Sabine Budde, and everyone else at MPI and the Joint Administration for their help with any administrative tasks.

My sincere thanks go to the participants and organizers of the DISTRO project (Marie Skłodowska-Curie ITN) for the many collaboration opportunities and training events during the program. I am especially grateful to the DISTRO supervisors with whom I extensively collaborated on the research projects described in this thesis: Bernd Bickel, Jaroslav Křivánek, Tim Weyrich, and Alexander Wilkie. It was a great experience to work with and learn from you.

I want to extend my gratitude to my colleagues in the research projects and my co-authors, especially Oskar Elek and Tobias Rittig. Working with you was both exceptionally productive and enjoyable, and I am confident that our collaborative efforts have significantly elevated the quality of our results.

I also want to thank the members of D4, as well as the members of Bickel's group at IST Austria and the Charles University Computer Graphics Group, for many interesting discussions and joint leisure activities.

I would like to acknowledge the people who contributed to my development and education before my Ph.D. In particular, Tatiana Kravchuk, head of Moscow Lyceum of Information Technologies; Dmitry Vatolin, head of Moscow State University Video Group; and Alexey Voloboy, my supervisor at the Keldysh Institute of Applied Mathematics RAS. The education and personal support I received from you have been invaluable to my academic journey.

Finally, I want to express my heartfelt thanks to my parents, my wife, my brother, and the rest of my family for their great support and care during my Ph.D. and throughout my life.

Contents

Abstract iii

Zusammenfassung v

Acknowledgments vii

1 Introduction 1
- 1.1 Scattering-Aware Texture Reproduction . . . 3
- 1.2 Light-field Quality Metrics . . . 4
- 1.3 Thesis Contributions . . . 5
- 1.4 Publications . . . 7

2 Background and Related Work 9
- 2.1 Fabrication . . . 9
  - 2.1.1 Fabrication Methodology . . . 10
  - 2.1.2 Color 3D Printing . . . 10
  - 2.1.3 Translucency Fabrication . . . 11
  - 2.1.4 Subsurface Scattering Effects . . . 11
  - 2.1.5 Reflectance Fabrication . . . 12
  - 2.1.6 Anisotropic Appearance Fabrication . . . 12
- 2.2 Enhancement and Pre-correction . . . 13
- 2.3 Volume rendering . . . 14
  - 2.3.1 Forward Volume Rendering . . . 14
  - 2.3.2 Inverse Volume Rendering . . . 14
  - 2.3.3 Neural Rendering . . . 15
- 2.4 Material Characterization . . . 16
  - 2.4.1 Computational Fabrication . . . 16
  - 2.4.2 Other Contexts . . . 16
- 2.5 Light Fields . . . 17
  - 2.5.1 Light-field Datasets . . . 17
  - 2.5.2 Metrics and Experiments . . . 18
  - 2.5.3 Light-field Displays . . . 19

3 Scattering-aware Texture Reproduction for 3D Printing 21
- 3.1 Method Overview . . . 22
- 3.2 Object Appearance Model . . . 23
- 3.3 Material Calibration . . . 24
  - 3.3.1 Measurement . . . 25
  - 3.3.2 Model Fit . . . 26
  - 3.3.3 Results . . . 27
- 3.4 Material Mapping . . . 28
  - 3.4.1 Tonal Mapping . . . 29
  - 3.4.2 Tonal Quantization . . . 33

3.5 Structural Optimization . . . 33
3.5.1 Problem Definition . . . 33
3.5.2 Solution Concept . . . 34
3.5.3 Optimization . . . 34
3.5.4 Discussion . . . 37
3.6 Fabrication . . . 38
3.7 Results . . . 38
3.8 Discussion . . . 41
3.9 Conclusion . . . 43

4 Geometry-Aware Scattering Compensation for 3D Printing 45
4.1 Problem Exploration . . . 46
4.1.1 Full-volume Optimization . . . 46
4.1.2 Heuristic Optimization . . . 48
4.2 General Geometry Scattering Compensation . . . 49
4.2.1 Data Preparation . . . 49
4.2.2 Interior-to-Surface Points Mapping . . . 50
4.2.3 Halftoning . . . 51
4.2.4 Appearance Prediction . . . 52
4.2.5 Material Arrangement Refinement . . . 53
4.3 Content-aware Gamut Mapping . . . 56
4.3.1 Gamut of Achievable Color Pairs . . . 57
4.3.2 Mapping . . . 58
4.4 Results . . . 60
4.4.1 Baseline Comparison to Remapped Planar Slabs . . . 62
4.5 Conclusion . . . 64

5 Neural Acceleration of Scattering-Aware Color 3D Printing 65
5.1 Behavior of Volumetric Scattering . . . 66
5.2 A First Approach to Fast Print Preparation . . . 67
5.3 Appearance Prediction Network . . . 68
5.3.1 Training Dataset . . . 69
5.3.2 Feature Descriptor . . . 69
5.3.3 Network Architecture . . . 70
5.3.4 Training Protocol . . . 72
5.3.5 Inference . . . 73
5.4 Results . . . 74
5.4.1 Scattering Compensation . . . 74
5.4.2 Network Generalization . . . 78
5.4.3 Application: Spectral Prediction . . . 79
5.4.4 Differentiable Rendering Comparison . . . 81
5.5 Discussion . . . 82
5.6 Conclusion . . . 83

6 Towards a Quality Metric for Dense Light Fields 85
6.1 Data Collection . . . 86
6.1.1 Scenes . . . 87
6.1.2 Distortions . . . 87
6.2 Experiment . . . 88
6.3 Analysis of Subjective Data . . . 90
6.4 Evaluation of Quality Metrics . . . 91

6.4.1 Metric Performance Comparison 92
6.4.2 Sparse Light-field Reference Case 94
6.5 Conclusion 95

7 Conclusions and Future Work **97**

Bibliography - Own Work **99**

Bibliography **101**

Chapter 1

Introduction

During the last decades, digital content manipulation tools have undergone significant progress, offering increasingly sophisticated and accessible software applications for content creation, editing, and manipulation. Graphic design software, video editing tools, 3D modeling software, and other content manipulation tools have become more intuitive, feature-rich, and powerful. These advances have empowered individuals and professionals alike to create, modify, and enhance digital content with greater ease and efficiency. Moreover, this evolution has blurred the boundary between the digital and real worlds, often enabling seamless integration and interaction between virtual and physical environments.

Simultaneously, 3D fabrication technologies have experienced remarkable development, enabling precision, efficiency, and versatility in the creation of physical objects directly from computer-aided design (CAD) models. 3D fabrication methods can be divided into subtractive and additive manufacturing (AM). Subtractive manufacturing focuses on carving out the final product from a solid source. Computer numerical control (CNC) machining is a prominent example, where material is systematically removed to shape the desired object. Subtractive processes excel in achieving high precision and are often employed for detailed and intricate designs, making them suitable for a wide range of materials. Conversely, additive manufacturing, often referred to as 3D printing, involves the layer-by-layer construction of objects from digital models. This process enables the creation of complex geometries that might be challenging with traditional manufacturing methods. Additive technologies are particularly advantageous for rapid prototyping, customization, and reducing material waste.

AM processes provide the potential to create objects possessing various features, such as:

- shape with defined dimensional accuracy;
- specified mechanical properties, including strength, stiffness, and elasticity;
- physical attributes like toughness and thermal conductivity;
- specific functional properties, such as electrical conductivity, thermal insulation, chemical resistance, and biocompatibility;
- appearance, which entails the visual aesthetics of the 3D-printed objects, encompassing surface finish, color, texture, and overall visual appeal.

There are many additive manufacturing technologies, each with its strengths and limitations. Some processes may excel in achieving specific features, while others may prioritize cost-effectiveness and rapid production [Patterson et al. 2015]. Additionally, the choice of materials within a particular technology can significantly influence the capabilities of additive manufacturing in producing objects with the desired features. The field of additive manufacturing is extensive, offering a wide range of possible

applications on one end and numerous technology options on the other end. However, despite many impressive examples of AM advances, the technology is still in a phase of ongoing exploration and research. One of the tasks that has been challenging is high-fidelity appearance fabrication, which aims to reproduce the target object's appearance as closely as possible.

High-fidelity appearance reproduction ensures realistic visual representation, aids in design validation, and enhances product customization [Urban 2020]. In architecture, it enables lifelike models for design review and client visualization. In the medical field, it contributes to accurate anatomical models for surgical planning and educational purposes. In consumer products, marketing, and cultural heritage preservation, achieving accurate appearance reproduction is vital for showcasing designs, preserving artifacts, and creating attractive visual materials. This level of precision elevates the usability and applicability of 3D printing across diverse industries, making it a versatile tool for creating visually appealing and functionally accurate objects. Although notable progress has been made in multi-material printing, leading to significant advancements in achieving full-color 3D printing, accurate reproduction of detailed high-contrast and vivid, saturated color textures with sharp boundaries has remained a hurdle. This thesis is motivated by the gap between the advances in fabrication hardware and the limitations of the software in realizing their full potential.

Multi-color material jetting (MJ) 3D printing is the high-end AM process for achieving high-quality full-color 3D prints as it can currently produce full-color prints with a resolution in the order of 10 µm [Sitthi-Amorn et al. 2015; Stratasys 2017]. Even with the material-jetting AM process, inherent physical constraints of translucent printing materials restrict the range of reproducible appearances when using straightforward algorithms to control material disposition. The resins used as print materials in commercial photopolymerization printers are inherently translucent, i.e., exhibit significant sub-surface scattering. This characteristic aids in effective color mixing during full-color printing processes, allowing commercial printer drivers to achieve high-quality color reproduction. However, the associated light diffusion can cause excessive blurring and color bleeding, thereby limiting the saturation and sharpness of spatially-varying color textures.

We develop a data-driven print-preparation system which counteracts heterogeneous scattering to obtain the impression of a crisp albedo texture on top of the 3D print (Figure 1.1). This is achieved by optimizing the volumetric material distribution to accurately replicate the target appearance.

We first set up a scattering-aware 3D-printing framework for the limited case of planar slabs and develop a practical procedure to characterize printer primary materials. Then we extend the approach to general 3D geometries and propose solutions for extremely challenging cases of arbitrary 3D shapes such as thin structures where ideal texture reproduction may not be feasible due to material translucency and light scattering to another material side. Finally, we investigate the performance issue of our print-preparation system which heavily relies on accurate but slow light transport simulation. We adapt a deep neural network to predict the scattering within a highly heterogeneous medium to be fabricated and achieve practical timings of print-preparation tasks.

Appearance fabrication is not limited to the reproduction of texture; it also has an angular component: an object's appearance can vary (intentionally or not) depending on the viewing direction. Just as our print-preparation system is guided by discrepancy metrics in the texture domain, the analysis of printed objects' appearance in the angular domain requires specialized representation and quality metrics. A light-field representation (or simply light field) contains both angular and spatial information

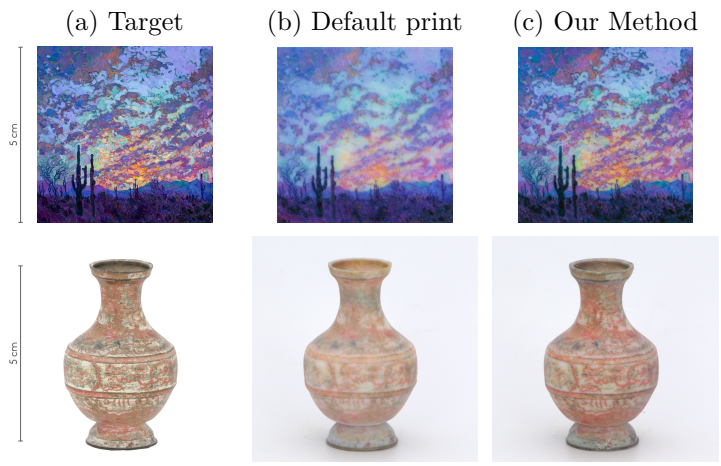


FIGURE 1.1: *Comparison of photographed 3D printouts to their rendered (a) counterparts (target models): (b) GrabCAD is the default printer software, and (c) Our sub-surface scattering and color-bleeding compensated solution. Both objects have size of approximately 5 cm. The planar slab in the top row is 1 cm thick.*

about the captured object or scene [Levoy et al. 1996]. Light field representation enables the processing of angularly varying effects. We evaluate existing quality metrics for light fields and propose a new light-field dataset. We perform a subjective experiment and derive the subjective quality scaling for various distortions that occur in light-field applications. This is an initial step towards a surface light-field metric that in future could guide the fabrication preparation and make the angular properties of 3D-printed objects more controllable.

This thesis consists of two main components:

1. A high-fidelity scattering-aware texture reproduction pipeline for 3D printing.
2. A survey on quality metrics applicable for measuring light fields quality, with a potential application to the challenging case of angular texture appearance changes.

Sections 1.1 and 1.2 further introduce these two parts of the thesis. Section 1.3 summarizes the main contributions of this thesis. In Section 1.4 we provide a list of publications covering the presented work.

1.1 Scattering-Aware Texture Reproduction

Color texture reproduction in 3D printing commonly ignores volumetric light transport between surface points on a 3D print. Such light diffusion leads to significant blur of details and color bleeding, and is particularly severe for highly translucent resin-based print materials. Given their widely varying scattering properties, this cross-talk between surface points strongly depends on the internal structure of the volume surrounding each surface point. Existing scattering-aware methods use simplified models for light diffusion, and often accept the visual blur as an immutable property of the print medium.

In Chapter 3, we address heterogeneous scattering in 3D-printed planar slabs. Our approach optimizes the volumetric material distribution to maintain a sharp and clear albedo texture, closely matching the intended appearance. Our method employs an efficient numerical optimizer on top of a general Monte-Carlo simulation of heterogeneous scattering, supported by a practical calibration procedure to obtain scattering parameters from a given set of printer materials. The presented system

is an end-to-end 3D-printing framework in which we control and thoroughly verify each component: acquisition of printing-materials' optical parameters, appearance prediction of an object to be fabricated, and volumetric deposition of specific printer material. Each part of the pipeline influences the others: appearance simulation requires correct optical parameters of the print materials, and volumetric optimization of materials' distribution relies on appearance prediction feedback. If any of these components is incorrect, the fabricated object would not match neither the target model nor the simulated preview of the printout. In this thesis, we present validation of each step in the pipeline and demonstrate results on real printouts.

In Chapter 4, we first revisit a previously made assumption that pushing the absorptive materials to the surface results in minimal volumetric light diffusion. We design a full-fledged optimization on a small domain for this task, where finding a global minimum is guaranteed, and confirm the previously proposed method. Then, we extend the above approach that was initially limited to color reproduction on planar surfaces, to arbitrary 3D shapes. In addition, we propose a content-aware gamut mapping that significantly improves color reproduction for the special case of thin surfaces, where lighting falling on one side is not fully attenuated inside material and affects the appearance of the another side. Despite the inherent translucency of the medium, we reproduce detailed surface textures on 3D prints; we evaluate our system using a commercial, five-color 3D print process and compare against the printer's native color texturing mode. We demonstrate how our method preserves high-fidelity in texture detail reproduction and minimizes errors in color reproduction that would otherwise lead to color desaturation due to light diffusion.

In Chapter 5 we dramatically speed up the print-preparation process by replacing the light transport simulation with a data-driven approach. Leveraging a deep neural network to predict the scattering within a highly heterogeneous medium, our method performs around two orders of magnitude faster than Monte Carlo rendering while yielding optimization results of similar quality level. The network is based on an established method from atmospheric cloud rendering, adapted to our domain and extended by a physically motivated weight sharing scheme that substantially reduces the network size. We analyze its performance in an end-to-end print preparation pipeline, compare quality and runtime to alternative approaches, and demonstrate its generalization to unseen geometry and material values. This enables full heterogeneous material optimization for 3D-print preparation within time frames in the order of the actual printing time.

1.2 Light-field Quality Metrics

So far we assumed that material reflectance properties are mostly Lambertian, where the reflected light radiance does not depend on the reflection direction. This assumption might not hold for a wide class of real-world materials and desired appearances, especially when considering the fabrication of 3D objects with angularly varying appearances, where the object's visual characteristics change when viewed from different perspectives. Angular light reflectance effects are specified by bidirectional reflectance distribution function (BRDF) and bidirectional texture function (BTF) for homogeneous and textured surfaces, respectively.

A high-quality 3D-print preparation system should consider the angular effects of 3D-objects' perception. A multi-view quality metric is essentially needed to incorporate the angular-dependent comparison of print specification with printout.

A light field is an ultimate generalization of a 2D image that describes angular light propagation in every direction through every point in the 3D space. Common 2D image is merely a single projection of the light-field 5D function. Light-field representation enables analysis of fabricated objects' angularly-varying appearance when seen from different viewpoints.

As the space of all possible light rays is huge, there practically exist several approaches to represent this space [Levoy et al. 1996]. Due to the enormous storage requirements, dense light fields are usually sparsely sampled in spatial and angular dimensions, stored using lossy compression, and reconstructed later. It is unclear how the distortions introduced on the way affect the perceived quality. While similar problems have been addressed for 2D images, only a few works consider dense light fields and lack systematic evaluation. In Chapter 6 we evaluate applicability of existing metrics to dense light fields and setup a ground for the future work towards specialized light-field discrepancy metrics.

1.3 Thesis Contributions

In this thesis we introduce major components of a 3D-print preparation system for multi-material jetting additive manufacturing process. We continue with the research of light-field quality metrics to support the future work on 3D-printing of objects with angular-dependent appearance.

The work presented in this thesis has been done as part of work of the "DISTRO" (Distributed 3D Object Design) Marie Skłodowska-Curie Innovative Training Network (ITN). Innovative Training Networks are designed to perform interdisciplinary and collaborative research covering wide topics. DISTRO was a collaboration between multiple laboratories in Visual Computing and 3D Computer Graphics research across Europe, comprising both academic and industrial partners.

Within DISTRO, we have been working on the distributed capture, editing, and fabrication of objects — from the real world to digital and back again. The overall goal was to enable users to casually capture objects, which can then be easily shared on the Web, customised in simple yet powerful ways, and finally fabricated again on a 3D printer.

Therefore, some of presented work is collaborative, but is still included for completeness to facilitate understanding of this thesis' contributions. Several papers which publish the presented work are attributed with "equal co-first authorship". All publications within the scope of this thesis are listed in Section 1.4. In this section we explicitly indicate the contributions of Denis Sumin.

Planar-slabs fabrication As a first step, we build a system capable of 3D-printing planar appearance on a bulk of white material. In Chapter 3 (published as [Elek et al. 2017; 2021]) we present:

- a complete end-to-end pipeline aimed at printing a sharp surface texture approaching opaque appearance qualities, subject to physical limitations;
- a custom nonlinear optimization method to compute a spatial material distribution that retains high-frequency details, compensating for the inherent translucency of the print materials;
- techniques of converting between different material parametrizations to facilitate the entire process up to the printing itself;
- a practical novel way to acquire the optical properties of the available printing materials.

This work was lead by Oskar Elek and Denis Sumin. Oskar Elek determined the initial project scope and lead the appearance prediction part of the system. Denis Sumin focused on the iterative material-arrangement method and developed the end-to-end pipeline; he also found a method to determine printing materials' optical parameters for channels with low signal-to-noise ratio (SNR) ratio. Ran Zhang contributed to the optical setup for calibration of printing materials.

Full-geometry objects fabrication We proceed with extending the print-preparation system to handle general 3D objects. In Chapter 4 (published as [Sumin et al. 2019]) we introduce:

- investigation of the solution space and validation of a previous material assignment heuristic using a full-volume conjugate-gradient optimization;
- design of a full-3D print-preparation pipeline that uses this heuristic for sharp, scattering-compensated texture reproduction on general 3D shapes;
- geometry-aware color gamut mapping where physical limits occur.

This work was done in a close collaboration by Denis Sumin and Tobias Rittig. Both first authors developed the full print-preparation pipeline; while Denis Sumin concentrated on the halftoning and material arrangement part, Tobias Rittig focused on the appearance prediction part. The solution space investigation was done by Tobias Rittig. The content-aware gamut mapping part was proposed and developed by Denis Sumin.

Neural-network acceleration of prediction system We conclude our print preparation system with a fast print appearance prediction to make the system practical. In Chapter 5 (published as [Rittig et al. 2021]) we contribute:

- fast appearance prediction for fully heterogeneous volumes using machine learning;
- up to 300× speedup with constant quality compared to existing 3D printing pipelines;
- a synthetic dataset of scattering volumes representing a wide gamut of material arrangements;
- vital architecture improvements that prepare the existing atmospheric-cloud neural-rendering framework for broader arrangements;
- a comparison of differentiable rendering and heuristic material arrangement refinement for 3D printing volumes.

This project was again a close collaboration of Tobias Rittig and Denis Sumin. The project's initial concept originated from Denis Sumin, who also conducted the initial experiments to validate the concept. The authors split the core work in a way that Tobias Rittig worked on the rendering part of the system while Denis Sumin focused on the data-processing and machine learning tasks.

Light-field quality metrics evaluation To enable fabrication of 3D objects with more general spatially and angularly varying appearance we perform evaluation of existing quality metrics for light-field images. In Chapter 6 (published as [Adhikarla et al. 2017]) we present:

- a publicly available dense light-fields dataset that is designed for training and evaluating quality metrics;

- a perceptual experiment that provides human quality judgments for several typical light-field distortions;
- an evaluation, analysis, and extensions of existing quality metrics in the context of light fields;
- a discussion of identified challenges in light field quality assessment, such as the need for a high-quality reference.

In this work, Denis Sumin was a key contributor to the project, with a focus on dataset implementation and the evaluation of light-field quality metrics.

1.4 Publications

Within the scope of the thesis, the following papers have been published:

Scattering-aware Texture Reproduction for 3D Printing [Elek et al. 2017]
Oskar Elek*, Denis Sumin*, Ran Zhang, Tim Weyrich, Karol Myszkowski, Bernd Bickel, Alexander Wilkie, Jaroslav Křivánek
ACM Transactions on Graphics (Proc. SIGGRAPH Asia) 36.6, 241:1–241:15. November 2017. ISSN: 0730-0301.

Robust and Practical Measurement of Volume Transport Parameters in Solid Photo-polymer Materials for 3D Printing [Elek et al. 2021]
Oskar Elek, Ran Zhang, Denis Sumin, Karol Myszkowski, Bernd Bickel, Alexander Wilkie, Jaroslav Křivánek, Tim Weyrich
Optics Express 29.5, pp. 7568–7588. February 2021. DOI: 10.1364/OE.406095.

Geometry-Aware Scattering Compensation for 3D Printing [Sumin et al. 2019]
Denis Sumin*, Tobias Rittig*, Vahid Babaei, Thomas Nindel, Alexander Wilkie, Piotr Didyk, Bernd Bickel, Jaroslav Křivánek, Karol Myszkowski, Tim Weyrich
ACM Transactions on Graphics (Proc. SIGGRAPH) 38.4. July 2019. ISSN: 0730-0301.

Neural Acceleration of Scattering-Aware Color 3D Printing [Rittig et al. 2021]
Tobias Rittig*, Denis Sumin*, Vahid Babaei, Piotr Didyk, Alexey Voloboy, Alexander Wilkie, Bernd Bickel, Karol Myszkowski, Tim Weyrich, Jaroslav Křivánek
Computer Graphics Forum (Proc. Eurographics) 40.2, pp. 205–219. May 2021. ISSN: 1467-8659. DOI: 10.1111/cgf.142626.

Towards a Quality Metric for Dense Light Fields [Adhikarla et al. 2017]
Vamsi Kiran Adhikarla, Marek Vinkler, Denis Sumin, Rafał Mantiuk, Karol Myszkowski, Hans-Peter Seidel, and Piotr Didyk
Proceedings of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR), 2017

** denotes shared co-first authorship of the works.*

Chapter 2

Background and Related Work

The primary focus of this work is appearance fabrication, an emerging field in computer graphics research. Its aim is to develop new methods and tools for the high-fidelity reproduction of 3D models in the physical world [Hullin et al. 2013]. As an applied, cross-disciplinary research topic, it draws upon and integrates advances from various sub-fields of computer science.

In this chapter, we provide a brief overview of background and related work, and establish connections between previous and our work.

We begin with related work in fabrication (Section 2.1) and color 3D printing, which are central to our work. The problem that our work addresses is quality degradation in existing 3D printing systems caused by subsurface scattering effects. As similar issues of signal degradation arise in other areas of computer graphics, we proceed with a literature review on pre-correction and image enhancement (Section 2.2).

Then we introduce the topics upon which we base our system and extend them to suit the requirements of our task, namely, multi-material 3D printing with photopolymer materials at resolutions in the order of 10 µm. Volume rendering (Section 2.3) serves as an appearance prediction engine for our print preparation system; moreover, the problem at hand can be categorized as an inverse volume rendering problem, as we aim to find a scattering volume that reproduces the target surface appearance. In this thesis, we first develop a solution for this task and then propose a custom neural rendering approach that dramatically reduces the time required for this search. Material characterization (Section 2.4) is essential for our system as it provides data for the appearance prediction component. In this work, we propose a practical method for acquiring absorption and scattering parameters of print materials and verify this approach by using it in the print-preparation system.

Finally, we review the literature on quality metrics for light fields (Section 2.5). Our 3D print preparation system utilizes conventional 2D image error metrics to guide the optimization of a scattering volume to be printed. This limits our method to fabricating objects with Lambertian surface appearance. We envision future development of our print preparation system by extending it to support angularly varying appearances, which would require calculating the difference between the target model and intermediate appearance predictions in a light field representation.

2.1 Fabrication

In this section, we review existing approaches to additive manufacturing in computer graphics. We discuss recent works in color 3D printing and translucency fabrication, which are closely related to the topic of our work. Subsurface scattering poses an inherent limitation of currently available printing materials, and our print preparation system must address these effects. Finally, we review recent efforts in reflectance and anisotropic appearance fabrication. While the system proposed in this thesis

is currently limited to Lambertian reflectance properties, we anticipate its future expansion to handle other reflectance models and produce accurate angularly varying appearances.

2.1.1 Fabrication Methodology

In computer graphics, at least two main paradigms of specifying fabricated objects are recognized: direct (forward) and constraint-driven (inverse). The OpenFab [Vidimče et al. 2013] and Spec2Fab [Chen et al. 2013a] systems are good respective examples of these two types of approaches. Both OpenFab and Spec2Fab can be seen as general frameworks which can accommodate various methods of specifying and/or optimizing for desired material properties, both visual and mechanical. Examples of these include custom BRDFs and other surface properties [Lan et al. 2013; Matusik et al. 2009; Rouiller et al. 2013; Weyrich et al. 2009], voxel micro-structure [Brunton et al. 2015; Doubrovski et al. 2015; Lou et al. 1998], continuous mixtures of inks [Papas et al. 2013], or other, technology-specific features [Pintus et al. 2010; Reiner et al. 2014].

Within that taxonomy, our work targets the inverse problem of additive fabrication of spatially varying subsurface scattering properties from a fixed set of translucent materials. As such, it falls into the same category as the works by Dong et al. [2010] and Hašan et al. [2010]. We observe, however, that once lateral subsurface scattering is taken into account, their goal of simultaneously controlling scattering *and* absorption at every point is an overdetermined problem. Instead, we forgo the control over the shape of scattering kernels and constrain diffuse albedo alone, providing the solver with the necessary degrees of freedom to match a target (albedo) texture despite lateral cross-talk through (heterogeneous) scattering. On the flip side, the lack of scattering control requires us to assume smoothly varying illumination of fabricated-object surface. As we demonstrate in Section 3.7, however, under many observing conditions this does not cause any practical issues.

2.1.2 Color 3D Printing

Closely related to our work are the recent efforts on color reproduction [Babaei et al. 2017; Brunton et al. 2015] using UV curable inkjet 3D printing. Although various color 3D printing techniques exist – using technologies such as paper lamination, powder-binder or fused filament fabrication – the inkjet technology gives arguably the highest output quality [Sitthi-Amorn et al. 2015; Stratasys 2017; 2020]. Using this technology, Brunton et al. [2015] extended error diffusion [Floyd et al. 1976] halftoning to 3D objects, while relying on a conventional color prediction model from 2D printing.

Babaei et al. [2017] showed how layering primaries with different thicknesses (contoning) can be used to control surface color and avoid halftoning artefacts. Their color reproduction relies on an absorption-only color prediction model, justified with the use of non-scattering materials. Shi et al. [2018] extended this approach to spectral color reproduction, using a deep neural network for the contone stack prediction. Morsy et al. [2022] addressed the problem of printed objects' surface quality and accuracy and proposed a method to counteract apparent quantization artifacts.

The above works ignore the important effects of lateral light scattering within the object, which induce unwanted color bleeding and texture blurring. Instead, they treat each small patch on the surface independently of its neighborhood. In our work, we specifically target this problem and propose a color reproduction method

that preserves texture detail by optimizing the full volumetric material arrangement so as to match the target surface appearance.

Although most of earlier works generate 3D results, either they do not consider high-frequency geometric features [Brunton et al. 2015] or propose special case solutions [Babaei et al. 2017].

Outside the realm of inkjet 3D printing, Morovič et al. [2019] simultaneously optimize for color and elasticity properties of objects fabricated by a multi-jet, powder-binding printer. Being constrained by extruder continuity in conventional fused filament fabrication (FFF) printers with multiple nozzles, Song et al. [2019] layer primary materials within one slice to achieve a color mixture.

Another noteworthy issue in the context of color 3D printing is material selection. In current fabrication devices, such as multi-material 3D printers, only a limited number of materials can be loaded into a printing machine at a time, despite the availability of many others. Users must determine which subset of available materials best suits a given print job. Several recent studies address this challenge and propose methods for its resolution [Ansari et al. 2020; 2022; Nindel et al. 2021; Piovarči et al. 2020].

2.1.3 Translucency Fabrication

Our current effort bears similarity to work on predefined translucency fabrication. Papas et al. [2013] determined pigment mixtures to reproduce homogeneous materials with a given color and translucency using silicon casts. Brunton et al. [2018] mixed primary cyan, magenta, yellow, key (black), and white (CMYKW) inks with a clear material to fabricate spatially-varying translucency. Our work addresses a different problem: surface texture color reproduction for 3D printers. Dong et al. [2010] and Hašan et al. [2010] fabricated heterogeneous subsurface scattering appearance by locally optimizing material stacks. Dong et al. [2010] furthermore employed inverse diffusion [Wang et al. 2008] to account for the interaction between different neighboring layers. In contrast to these works that reproduce scattering profiles explicitly, our method considers subsurface scattering implicitly by recreating the diffuse input color map while taking into account translucency of the printing materials.

2.1.4 Subsurface Scattering Effects

Established printing pipelines for color texture reproduction on 3D prints still follow strategies from 2D printing: their appearance formation models assume local mixture of primaries, without lateral cross-talk beyond *dot gain*, which also is routinely considered in 2D print production pipelines [Stollnitz et al. 1998].

Few works model the effect of lateral (subsurface) scattering in the material, and most assume it to be a spatially uniform blur that attenuates high frequencies across the surface. The resulting translucent appearance is taken as a welcomed effect [Brunton et al. 2015], kept to a minimum through specifically designed printing materials [Babaei et al. 2017], and/or is approximately compensated for by de-convolution [Babaei et al. 2017] or through an adaptation of unsharp-masking [Cignoni et al. 2008; Pintus et al. 2010].

As we will show, however, printing materials used in commercial 3D printers vary significantly in their effective scattering properties (Section 3.3). Such scattering heterogeneity implies a strong effect on the final appearance of a 3D printout. Hašan et al. [2010] and Dong et al. [2010] exploit this to control local scattering profiles by

suitable volumetric multiplexing of materials. Their models, however, approximate scattering only locally, ignoring heterogeneous scattering over larger lateral distances.

Bidirectional scattering-surface reflectance distribution functions (BSSRDFs) are a common approximation for translucent appearance in computer graphics [Christensen 2015; d'Eon et al. 2011; Donner et al. 2008; Jensen et al. 2002; 2001; Liang et al. 2024; Song et al. 2009]. We however argue (Section 3.2) that, in our application, any such model that considers scattering purely locally is bound to yield visible inaccuracies, as for most commercial printing materials the effective scattering distance is substantial (e.g., on average, light travels much farther than the print's resolvable feature size before contributing to the surface's appearance). Our work hence relies on a full Monte Carlo simulation of light transport in heterogeneous media to model appearance from a given material distribution. We offset the increased computational cost by employing an inverse-rendering algorithm with a high convergence rate. We furthermore speed our system up by employing a trained neural rendering approach instead of the slower Monte Carlo simulation.

2.1.5 Reflectance Fabrication

Earlier works aimed at fabricating custom surface reflectance using different hardware setups. Weyrich et al. [2009] designed microgeometries on a metallic surface to fabricate objects with predefined highlights. Matusik et al. [2009] created prints with spatially-varying BRDF using a 2D printer capable of printing with diffuse and metallic inks. Lan et al. [2013] used a bi-scale approach to fabricate anisotropic reflectances: they combined the appearance of a small-scale height field, manufactured using a 3D printer, with material BRDFs, printed using a flatbed printer. Piovarči et al. [2020] introduced novel printing hardware capable of jetting highly viscous varnishes to control the gloss of a given 2.5D object. In a recent study, Condor et al. [2023] demonstrate the dependence of an object's color perception on its gloss. Specifically, they observed that two samples with the same color but different surface gloss appear as if they have different colors. To address this issue, they propose a method to estimate this color shift and correct it, enabling the fabrication of objects with the same perceived color but different gloss levels. Our work forgoes accurate surface reflectance control and instead focuses on high-quality color and texture reproduction.

2.1.6 Anisotropic Appearance Fabrication

The fabrication of anisotropic, directionally varying appearance is a relatively unexplored and emerging area of research. While recent works have primarily focused on fabricating objects with specific appearance effects that can be categorized as anisotropic, to our knowledge, no general framework has been proposed thus far. Examples of recent studies in this field include: the fabrication of surfaces exhibiting a prescribed varying field of anisotropic appearance using FFF printers [Chermain et al. 2023]; 3D printing of goniochromatic effects (color change with viewing angle) on arbitrary surfaces using MJ printers [Abu Rmaileh et al. 2023]; and the fabrication of static 4D light fields [Zheng et al. 2020].

As part of future work, we envision the development of a general framework for fabricating anisotropic, angularly varying appearances. To integrate such framework into the optimization loop of our print preparation system, there is a need for a light field error metric that would guide the optimization. We take the initial step towards creating such a metric by studying existing light field quality metrics in this thesis.

2.2 Enhancement and Pre-correction

Enhancement is a very common operation in image processing [Gonzalez et al. 2007], typically aiming to improve the perception of images when presented with limited resolution and/or on devices with limited contrast. Similar problems are encountered in volume rendering and 3D visualization [Cai et al. 1996; Marchesin et al. 2007; Tao et al. 2009], or general rendering contexts [Ritschel et al. 2008]. The solutions are often variations of local edge enhancement and sharpening, e.g., via unsharp masking. Pre-correction generally refers to the modification of content to counteract subsequent degradations in the output. In general, we treat pre-correction and enhancement as two sides of the same coin, that is, as conceptual approaches to counteract or compensate for limitations of some parts in the reproduction pipeline. These include display, acquisition, or, in our case, fabrication. The essential component of designing each such method is taking into account specific limitations of the target platform, as ours and all of the referenced approaches do.

Our work is broadly related to the compensation techniques used in 2D printing for counterbalancing the dot gain [Rogers 1997]. The current best practice there is to account for the mechanical dot gain by adapting halftone dot areas [Hersch et al. 2005]. Moreover, the optical dot gain is taken into account using the Yule-Nielsen factor [Ruckdeschel et al. 1978]. These empirical models are practical but limited to a simple ink-on-paper setup and uniform-color halftones.

In the context of 3D printing, some works attempted to counteract the dot gain effects. Most closely related to our problem is the work by Cignoni et al. [2008], who recover the visibility of geometric details obfuscated by light diffusion, by countering the effect through modulation of the material albedo. Their key idea is to render the target object twice: once with simulated subsurface scattering and once without, in order to determine the shading difference that needs to be added back to maintain the appearance of a surface without scattering. Their method is elegant and effective but is limited to mid-tone target albedos to allow for sufficient freedom to darken and lighten the albedo as required. Furthermore, it does not consider the impact of light diffusion on the albedo texture itself, instead assuming artifact-free texturing of the 3D print. Babaei et al. [2017] used a similar approach except that they measured the difference and used it for deconvolution of the target texture.

Our method is inspired by these approaches, but tries to overcome these restrictions by accurately modeling crosstalk within the heterogeneous material while controlling material deposition volumetrically, rather than applying 2D corrections only.

Related to our thin feature treatment are works that enhance the shape to recover high-frequency geometric features which would otherwise be lost during the printing process [Herholz et al. 2017; Pintus et al. 2010]. While the geometry of our thin features is fully printable, it is *optically* too thin to reproduce arbitrary colors (consider a thin wall).

Other relevant works include pre-correction for light pollution (cross-talk) in display [Konrad et al. 2000] and projection systems [Baar et al. 2011], and improving visual acuity for observers with degraded vision [Montalto et al. 2015]. These problems conceptually compare to ours, but differ in aspects which make our problem more difficult to solve. Namely, the degradation due to light pollution [Baar et al. 2011] has a comparatively simple additive character, while blurring due to subsurface scattering in heterogeneous materials has a complex, nonlinear behavior. Also, pre-correcting for vision degradation [Montalto et al. 2015] can reasonably assume a constant point spread function (PSF), while in our case, the subsurface scattering kernel has a high, material-dependent spatial variance. In addition, in both above

cases, the degradation can be formulated as a convolution-like process with a relatively small kernel, while our problem deals with very wide-support 'kernels' with heavy tails that cannot be easily truncated.

2.3 Volume rendering

In the context of this work, volume rendering is a framework for modeling light transport in fabricated 3D objects with inherent heterogeneous scattering. We need to model light behavior inside printing materials to calibrate material parameters, obtain accurate predictions of printed objects (forward model), and find an arrangement of primary printing materials in the volume that results in a printout close to the target object specification (inverse model).

As our objective is to accurately model a physical process rather than solely producing visually appealing images, our approach is grounded in physically based methods for solving the volume rendering equation. Specifically, we rely on Monte Carlo path tracing, which represents the current state of the art for achieving realistic rendering results [Novák et al. 2018a; b]. Despite offering exceptional rendering quality, Monte Carlo path tracing is computationally intensive and slow. Consequently, in practical applications, it is common to substitute it with faster models that are validated against the accuracy of pure Monte Carlo physical rendering.

In this work, we first develop our core method for preparing 3D objects for printing using the base Monte Carlo path tracing. Then, we develop a specialized machine-learning-based model to enhance the performance of our solution.

2.3.1 Forward Volume Rendering

Numerous approximate methods to simulate scattering in optically dense materials have been developed [Christensen 2015; d'Eon et al. 2011; Donner et al. 2008; Jensen et al. 2001; Wang et al. 2008], their common deficiency being limited accuracy for non-planar geometries with heterogeneous scattering parameters – the very setup targeted by our method. Hašan et al. [2013] took advantage of light transport convexity for interactive re-rendering of arbitrarily shaped volumes. Their method only supports single-scattering albedo variations, while our printing materials additionally feature varying volume density. For this reason, our forward solution relies on full Monte Carlo simulation, an approach recently adopted also in the visual effects industry [Burley et al. 2018; Fascione et al. 2018].

2.3.2 Inverse Volume Rendering

Our material arrangement optimization is essentially an inverse volume rendering problem: the ultimate goal of our work is to find a scattering volume given an input surface appearance. The approach has previously been used to measure the scattering properties of *homogeneous* volumes [Gkioulekas et al. 2013b]. In a recent work Che et al. [2020] learn to predict *homogeneous* material coefficients from images using a renderer for training. Zhao et al. [2016] optimized the scattering parameters of a downsampled volume to match the appearance of an original high-resolution dataset. Finally, Gkioulekas et al. [2016] used time-resolved inverse rendering under structured lighting to recover the spatially-varying scattering properties of a heterogeneous medium. A part of their solution involves a full optimization similar to the

one we use to validate our heuristic approach. We employ machine learning to accelerate the forward volume rendering, which guides our heuristic method for the inverse reconstruction of a heterogeneous medium.

Differentiable volume renderers [Nimier-David et al. 2019; Zhang et al. 2019] are capable of reconstructing any volume parameter to match a target surface appearance using stochastic gradient descent (SGD). The recent publication of Radiative Backpropagation (RB) [Nimier-David et al. 2020] allows for practical usage of differentiable rendering without significant time or memory overheads. We show in Section 5.4 how such a general tool is applicable in the context of 3D print preparation and compare against the heuristic inverse method proposed in this work. Recently, differentiable renderers [Nimier-David et al. 2022] have become even more efficient. Vicini et al. [2021] and Nimier-David et al. [2022] proposed rendering algorithms with constant memory requirements and linear computation time in the number of scattering events which is particularly crucial for highly scattering mediums like printing materials in our work. Building upon these advancements and our own research, Nindel et al. [2021] applied differentiable volume rendering to full-color inkjet 3D printing and demonstrated the feasibility of this approach.

2.3.3 Neural Rendering

Machine learning has recently become a popular technique for image synthesis and has been applied to many problems in computer graphics. We refer the reader to the recent survey by Tewari et al. [2020] for a complete overview of general neural rendering. In Monte Carlo rendering, machine learning has been used for varying lighting in global illumination [Ren et al. 2013], path guiding in primary sample space [Müller et al. 2019; Zheng et al. 2018], importance sampling of first-bounce contribution [Bako et al. 2019; Huo et al. 2020], and image denoising [Bako et al. 2017; Chaitanya et al. 2017; Gharbi et al. 2019; Jiang et al. 2019; Kalantari et al. 2015; Vogels et al. 2018]. Here, we particularly focus on data-driven approaches in conjunction with volume rendering [Keller et al. 2018].

Neural networks have replaced some components, such as path computation, inside the rendering loop for compute-heavy problems like volume scattering. Vicini et al. [2019] propose a learned subsurface scattering model for homogeneous media that samples an exitant location over a polynomial approximation of the local shape and infers a corresponding path attenuation. Long light paths in heterogeneous, forward-scattering media like clouds are the motivation for Kallweit et al. [2017] to approximate multiple-scattering using a neural network. They learn the contribution of a directional light source towards a point inside the medium determined by Monte Carlo (MC) single-scattering. Recently, Panin et al. [2019] showed a 2.5× speed improvement over the latter work using spatial sharing of precomputed partial results.

Our approach is inspired by the work by Kallweit et al. [2017], but rather than evaluating the network for a single path sample, we replace MC integration altogether and predict the final radiance at once. Instead of directional lighting we assume the simpler case of smoother diffuse illumination, but with the added difficulty of a spatially-varying scattering albedo unlike the constant-albedo assumption of cloud modeling. Lastly, we improve the network architecture by a physically motivated weight sharing scheme that substantially reduces the network size.

2.4 Material Characterization

Our print preparation system is designed for real-world 3D printers and the available printing materials compatible with these printers. To accurately optimize the volumetric material arrangement to match the appearance of the target model, we must measure the properties of the printing materials and incorporate this data into the appearance prediction module. Given our focus on material characterization for use in physical reproduction, we begin our literature review starting from computational fabrication, followed by approaches used in other related fields.

2.4.1 Computational Fabrication

Several approaches suitable for translucent polymeric media have been developed in this context, with the aim to provide pragmatic means to characterize the involved printing materials [Dong et al. 2010; Hašan et al. 2010; Papas et al. 2013]. As such, their measurement approaches are typically adapted to the method-specific goals: Hašan et al. [2010] and Dong et al. [2010] both focus on reproducing spatially varying translucency effects, whereas Papas et al. [2013] aim at the reproduction of arbitrary homogeneous materials by mixing silicone pigments.

Consequently, common to these methods is the reliance on specialized data (implicitly [Dong et al. 2010] or explicitly [Hašan et al. 2010; Papas et al. 2013] represented scattering 'profiles') in conjunction with the *reduced* optical density and scattering albedo [Wyman et al. 1989]. Both of these representations carry the assumptions of smooth target signals and diffusive transport conditions, which is problematic as sharp material transitions are crucial for our work (cf. Figure 3.15). We also operate with strongly absorbing materials, which violate the conditions of diffusive transport [d'Eon et al. 2011; Jensen et al. 2001] – indeed, some of these methods [Dong et al. 2010; Papas et al. 2013] need to specifically compensate for that.

Several studies have recently been published for the same printing materials families as characterized and evaluated in our work. Iser et al. [2022] present the next generation of affordable measurements of print materials. They employ a simple spectrometer to obtain the absorption and scattering coefficients of a print material, as well as its phase function anisotropy factor. Chen et al. [2021, 2022] propose trainable deep learning models for optically characterizing 3D printers. They evaluate the proposed models on four state-of-the-art multimaterial 3D printers and demonstrate that such models can achieve accuracies sufficient for most applications.

2.4.2 Other Contexts

For bulk media, it is possible to obtain the transport parameters using, for instance, the inverse adding-doubling method [Leyre et al. 2014; Prahl et al. 1993]. In such conditions, this approach relates the total reflectance and transmittance of the medium to its intrinsic properties. It is also feasible to perform an isolated inference of some of the transport parameters (such as the scattering albedo α), followed by generalizing the relationship in a convenient closed mathematical form [Papas et al. 2013; Pharr et al. 2016]. These approaches often make the critical assumption of diffusive conditions [Jensen et al. 2001] and are limited by the low degrees of freedom of the measured quantities, meaning they do not generalize well for strongly absorbing and spatially varying materials (cf. [Correia et al. 2017]).

Media in specific conditions, or with properties that are known a priori or can be derived from first principles, do not need to be characterized by fully general measurement methods. A good example are fluid media, where dilution can enable

approximate isolation of single light-particle interactions [Narasimhan et al. 2006]. This is generally infeasible for most solids, since dissolving a solid material might alter its optical properties in an unknown way.

Specialized approaches exist for media with known components at unknown distributions, such as clouds [Derr 1980; Hu et al. 1993; Levis et al. 2017; Wu et al. 2009], biological tissues [Donner et al. 2008; Tuchin 1993], or milk [Frisvad et al. 2007]. In many cases including ours, knowledge like this is not available, precluding a specialized solution of any such kind.

To our knowledge, the most complete method has been proposed by Gkioulekas et al. [2013b] (who also provide an excellent literature review). This inverse-scattering approach measures the full set of optical parameters, including an approximate shape of the phase function (i.e., not only its anisotropy g used by our model). While successfully used to characterize liquids, gels, and suspensions with high accuracy, it is unclear whether this method would perform equally well on solid materials. Furthermore, it comes with strict calibration requirements for both lighting and geometry. This, in combination with the high material costs of the setup, has motivated us to develop an alternative methodology with emphasis on simplicity and affordability. Nevertheless, we draw inspiration from the work of Gkioulekas et al. through the use of *material dictionaries* to solve the inverse scattering problem.

2.5 Light Fields

The 3D print preparation system presented in this thesis requires an error metric for comparing the target model to the predicted appearance of a proposed print-material distribution. Currently, the system utilizes conventional error metrics. Specifically, we employ the root mean square error (RMSE) on linear values in the RGB color space and CIE dE76, which is the Euclidean distance in the perceptual CIELAB (or L*a*b*) color space. While effective for cases involving Lambertian-reflectance appearances, these metrics are unable to capture variations in the angular domain exhibited by a wide class of real-world materials and desired appearances. In such cases, objects may appear different when viewed from even slightly different angles. Extending the system to accommodate this variation would require the calculation of errors in a light field representation.

In this section, we provide an overview of existing datasets for light fields as well as the experiments that measure the perceived distortions in various types of content. For completeness, we also include references to works on light-field displays, as quality evaluations of light fields is usually conducted on specialized display devices.

2.5.1 Light-field Datasets

There are several publicly available light-field datasets. The most popular ones are: 4D light-field dataset [Wanner et al. 2013] containing seven synthetic scenes and five real-world scenes, Stanford archive [Computer Graphics Laboratory, Stanford University 2008] with twenty 4D light fields, and Disney 3D light-field dataset [Kim et al. 2013] containing five scenes. Although the first two datasets provide a good quality and reasonable number of light fields, they are captured over very narrow baselines that are insufficient for the new generation autostereoscopic displays. The Disney dataset provides high spatio-angular resolution light fields; however, they are few and do not have consistent spatial and angular resolution, which makes it difficult to use in quality evaluations. In the context of quality evaluation of 3D light fields, three real-world light fields are provided in the IRCCyN/IVC DIBR Images database

[Bosc et al. 2011]. These contain several scenes captured along a wide baseline at the cost of reduced angular resolution. Tamboli et al. [2016] provided 360° round table shots of three scenes that are used for quality evaluation on a 3D light-field display. These are rather simple scenes with single objects and the images contain a lot of noise. In our work, we provide first consistent dataset of dense, complex-scene light fields with large appearance variation. We use the dataset for training and evaluating quality metrics. The database can also serve as a ground truth for automultiscopic displays.

2.5.2 Metrics and Experiments

Because of their proven efficiency on 2D images, 2D objective metrics are viable candidates for evaluating light-field quality. Yasakethu et al. [2008] tested the suitability of objective measures – structural similarity index measure (SSIM) [Wang et al. 2004], peak signal-to-noise ratio (PSNR) and video quality model (VQM) [Pinson et al. 2004] for quality assessment for stereoscopic and 2D+Depth videos that are compressed at different bitrates. They carried out subjective experiments on an autostereoscopic display and showed that 2D metrics can be used separately on each view to assess 3D video quality. They used few sequences and studied only compression artifacts. Several metrics have been proposed to determine the quality of synthesized views from multiview images. Bosc et al. [2011] advocated two measures for assessing the quality of synthesized views. However, they did not conduct thorough subjective studies. Solh et al. [2009] presented a metric for quantifying the geometric and photometric distortions in multiview acquisition. Bosc et al. [2013] suggested a method to asses the quality of virtual synthesized views in the context of multiview video. Battisti et al. [2015] proposed more sophisticated framework for evaluating the quality of depth image based rendering techniques by comparing the statistical features of wavelet subbands and used image registration and skin detection steps for additional optimization. Sandic-Stankovic et al. [2016] exploited multi-scale pyramid decompositions with morphological filters for obtaining the quality of intermediate views and showed that they achieve significantly higher correlation with subjective scores. These methods form a class of metrics specific to view-interpolation artifacts, and 2D stimuli containing the interpolated views are used for subjective experiments.

Vangorp et al. [2011] ran a psychophysical study to account for the plausibility of visual artifacts associated with view interpolation methods. They considered such artifacts as a function of different number of input images; however, they limited their study to monocular viewing and Lambertian surfaces. An experiment was also performed for precomputed videos, so that the impact of user's interaction and dynamic aspects of free viewing could be judged. More recently, this work was extended to transitions between videos [Tompkin et al. 2013]. Similar studies were also performed in the context of panoramas [Morvan et al. 2009]. Tamboli et al. [2016] conducted subjective studies on a 3D light-field display. Users were asked to judge the quality as perceived from different viewing locations in front of the display and the scores were averaged over all locations. The user could rate the quality only from a certain viewing position. Moreover, they only considered three distinct scenes. We believe that, for inferring a light-field quality, all the views should be taken into account at the same time.

2.5.3 Light-field Displays

Our work focuses on wide-baseline 3D light fields which enable perfect simulation of stereoscopic viewing and continuous horizontal motion parallax crucial for new light-field displays. Although many light-field display designs exist [Masia et al. 2013], including more advanced ones that provide focus cues [Maimone et al. 2013], they suffer from several drawbacks such as limited field of view, discontinous motion parallax, visible crosstalk, and limited depth budget. Several strategies have been proposed to minimize these artifacts by filtering the content [Du et al. 2014; Zwicker et al. 2006] and manipulating depth [Didyk et al. 2012; Lang et al. 2010; Masia et al. 2013]. However, display designs that enable displaying reference light fields for quality measurements are still unavailable.

Chapter 3

Scattering-aware Texture Reproduction for 3D Printing

The past decade has witnessed an unprecedented technological development of 3D printing, both in breadth and depth: consumer-grade printers are now commonly available for a range of purposes, while increasingly advanced techniques allow us to fabricate novel shapes, mechanical properties, and appearances. Notably, the printers' capabilities have improved dramatically from printing single-material objects to producing detailed structures with per-voxel material variation. Such capabilities are most prominent in printers based on the photo-polymer jetting process, which can currently produce full-color prints with a resolution in the order of 10 µm [Sitthi-Amorn et al. 2015; Stratasys 2017].

In commercial photopolymerization printers, the print materials are inherently translucent, causing notable sub-surface scattering. While this property enhances color mixing and enables high-quality color reproduction in full-color print processes, it also results in light diffusion that can lead to over-blurring and potential color bleeding in spatially-varying color textures.

In this chapter (published as [Elek et al. 2017; 2021]), we describe our initial step to building a complete end-to-end pipeline aimed at enabling 3D-printing of high-fidelity color textures. In this initial step, we restrict the pipeline to printing only 2.5D (planar) slabs (see Figure 3.1), and we expand beyond this limitation in subsequent chapters.

This bears resemblance to works on fabrication of custom subsurface scattering [Dong et al. 2010; Hašan et al. 2010] and color in 3D prints [Brunton et al. 2015], which produce impressive physical renditions of naturally translucent objects, such as marble, wax and organic tissues, capitalizing on the translucency of the underlying materials. Such works, however, use approximate models for light diffusion, and often



FIGURE 3.1: *A still life photograph of our optimized printouts. The thickness of all the pictured samples is 1 cm.*

accept the apparent blur as an immutable property of the print medium; their models work best if the target appearance lacks high spatial frequencies, conveying a sense of translucency in proportion with the actual translucency of the materials.

In contrast, our work specifically targets the challenging case where the target appearance *does contain* high-frequency details, as in the common use case of assigning a detailed color texture to a 3D object, regardless of the translucency characteristics of the printer material. Our work counteracts heterogeneous scattering to obtain the impression of a crisp diffuse albedo texture on top of the 3D print, by optimizing for a fully volumetric material distribution that preserves the target appearance.

We use an efficient numerical optimizer alongside a general Monte Carlo simulation to model heterogeneous scattering. This simulation relies on scattering parameters of printer materials obtained through a practical calibration procedure.

3.1 Method Overview

Our method uses a physically-based Monte-Carlo scattering model to predict the appearance of a heterogeneous, 3D-printed volume (Section 3.2). To obtain the optical parameters for our print materials, we employ a lightweight calibration procedure outlined in Section 3.3.

Once the optical parameters for the available set of print materials are known, the core task of our work is to map a *surface* color specification to a *volumetric* distribution of materials that matches this color specification as closely as possible.

In a first approximation (presented in Section 3.4), we achieve this through a sequence of remapping steps, using the color and materials specifications outlined in Table 3.1. Given a color texture specification in (linear) RGB space $\mathcal{C}$, gamut-mapped to the printer's native gamut constraints and extruded into the volume, we determine a tonal volume that consists of idealized mixtures of the print materials in $\mathcal{M}$ (in our case CMYKW). In this step we assume a perfect, continuous tonal mixture and no lateral cross-talk, before quantizing values in $\mathcal{M}$ to a half-toned representation $\mathcal{M}'$ of realizable material assignments, ready for print.

The resulting object, however, would still be subject to excessive blur, due to lateral scattering not having been taken into account. We hence feed the predicted residual difference into an outer optimization loop (cf. Figure 3.2; detailed in Section 3.5) to obtain a refined RGB specification that, when mapped to the print materials, matches the input specification well. Key to effective convergence of that

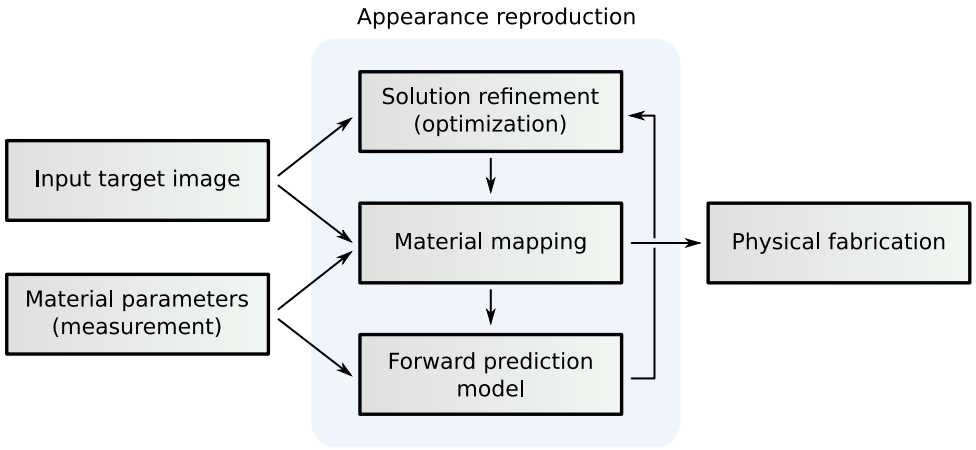


FIGURE 3.2: *High-level scheme of our pipeline, discussed in Section 3.1.*

TABLE 3.1: *Detailed summary of the spaces our method operates in (listed in the order of appearance in the text).*

Space	Type	Definition	Description
$\mathcal{C}$	continuous	$[0,1]^3$	Linear RGB colorspace describing the object appearance under a canonical 'white-sky' illumination. Under these conditions, we treat color as equivalent to observed reflectance, i.e., the relative portion of incident energy reflected by the object as a result of both surface and sub-surface light transport.
$\mathcal{V}$	continuous	$(\sigma_\mathrm{t}, \alpha, g, \eta)^3$	Space of intrinsic optical parameters of the (volumetric) printing materials. For each RGB channel, it defines four parameters: $\sigma_\mathrm{t} \in [0,\infty)$, the material's extinction coefficient (optical density); $\alpha \in [0,1]$, the single-scattering albedo, $g \in (-1,1)$, the scattering anisotropy; and $\eta \in (1,\infty)$, the refractive index.
$\mathcal{M}'$	discrete	$\{\mathrm{C},\mathrm{M},\mathrm{Y},\mathrm{K},\mathrm{W}\}$	Physical space of printer-specific materials, each of which has a unique set of optical parameters $\in \mathcal{V}$ used by our prediction techniques. Each voxel of the printed object must be assigned exactly one material.
$\mathcal{M}$	continuous	$[0,1]^5$	Virtual CMYKW tonal space of the printer. Given the properties of $\mathcal{M}'$, it is a *barycentric* space, i.e., for a given tonal mixture $w \in \mathcal{M}$ it must hold that $\sum w_i = 1$.

optimization is a carefully designed heuristic to drive the appearance residual into a volumetric RGB specification, prioritizing embeddings at different depths according to the *sign* of the residual texture. As such, it significantly differs from traditional error diffusion schemes (e.g., [Brunton et al. 2015]).

We evaluate our approach using a commercial poly-jetting printer (Section 3.6) and demonstrate its performance by assigning highly detailed target textures to solid slabs (Section 3.7), comparing to printouts using the printer's native color texturing mode. While our results show planar surfaces, there is no principal reason that would prevent a future extension to more general shapes. Please see Section 3.8 for a more detailed discussion.

3.2 Object Appearance Model

This section describes the forward model, which is used in our method's prediction step to obtain an appearance prediction from an object's parametric description.

In spite of the popularity of analytic BSSRDFs in translucent appearance fabrication [Dong et al. 2010; Hašan et al. 2010; Papas et al. 2013] and a progress in their *spatially varying* application [Sone et al. 2017; Song et al. 2009; 2013], we opt *not* to use an analytic BSSRDF as a forward prediction model. The main reason is that the current methods' accuracy decreases when the transport scale exceeds the size of the target appearance features (see Figure 3.3). This certainly happens in our conditions, as the target texture features are in the order of 10 µm, while our materials have mean free paths in the order of 0.1 mm (cf. Section 3.3), leading to visible scattering as far as 10 mm into the object. In addition, common spatially-varying BSSRDF models cannot fully account for volumetric transport due to lateral material heterogeneities, which is essential for our method.

Instead, we use standard path tracing (implementation provided by the Mitsuba renderer [Jakob 2010]) as the prediction model for the textured object's appearance. While computationally expensive, the accuracy of a Monte Carlo solution outweighs

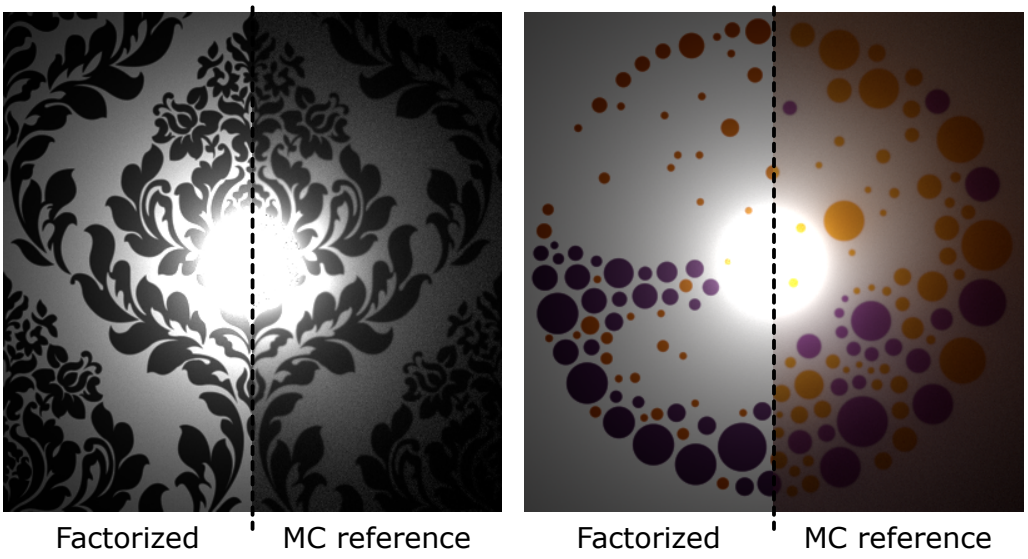


FIGURE 3.3: *Translucent transport in a material with a constant density (proportional to our printing materials) and laterally varying albedo, illuminated by a circular spotlight. We compare renderings using the factorized BSSRDF by Song et al. [2009] with a Monte Carlo path-traced reference. The BSSRDF tends to overestimate the transported energy and does not well reproduce subtle color bleeding from nearby features.*

the cost, as this is directly correlated with the quality of the resulting color and texture reproduction. Path tracing performs well under natural ambient illumination that we target, and can be sped up by disabling next-event estimation (from within the medium) which is inefficient for non-index-matched objects we deal with here.

Forward Prediction Model We seek to obtain the textured object's appearance, where we consider the object in a canonical position viewed from the top, illuminated by a normalized 'white sky' constant environment map. Under those conditions, the target appearance is defined as a diffuse RGB reflectance of the object in $\mathcal{C}$ (refer to Table 3.1 for an overview of our operating spaces). Formally we define the forward prediction model as

$$f\colon \mathcal{V} \to \mathcal{C}. \tag{3.1}$$

Equation (3.1) maps a volume (3D matrix) of the object's optical parameters $\boldsymbol{V} \in \mathcal{V}$ to a color image (2D matrix) $\boldsymbol{C} \in \mathcal{C}$.

Each CMYKW material $m \in \mathcal{M}'$ has a corresponding set of physical volume parameters $v_m \in \mathcal{V}$, which we measure via the procedure detailed in Section 3.3. Since every voxel can only be assigned a single material, the object volume $\boldsymbol{V} = \{\boldsymbol{v}_i\}$ is eventually quantized (half-toned): $\boldsymbol{v}_i \in \{v_m\}$. We detail this quantization in Section 3.4.2.

3.3 Material Calibration

Our prediction model described in Section 3.2 relies on the knowledge of the optical data $\in \mathcal{V}$ (Table 3.1) of our printing materials. While state-of-the-art measurement methods can deliver excellent accuracy (see [Gkioulekas et al. 2013b] and references therein), this can come at the cost of complex design and expensive equipment.

We instead propose a novel simple and affordable measurement method that well serves our need of obtaining data that yield *visually* accurate prediction. In this section we briefly explain the concept of our approach (Section 3.3.1), describe our fitting-based data extraction (Section 3.3.2) and present the resulting measurements (Section 3.3.3). A significantly expanded exposition is provided in the supplemental material.

3.3.1 Measurement

Since it is hard to measure the optical parameters of solid volumetric materials directly, we base our method on inverse rendering. Our approach is to acquire a complex, extrinsic interaction of light with the medium, and reproduce it in a simulation under equivalent conditions. By generating a sufficiently dense dictionary of virtual observations using different optical parametrizations $\in \mathcal{V}$, we obtain the measurements by finding a fit that matches the acquired physical observation [Gkioulekas et al. 2013b].

Our acquisition setup observes a simple printed step edge through a thin slice of the printing material, as sketched in Figure 3.4. The setup consists of six main components:

- *Sample.* A thin slice with the dimensions $40\times40\times0.5/1\,\text{mm}$, 3D-printed with one of the materials from $\mathcal{M}'$. Surface finish as smooth as possible.
- *Signal modulation.* Reflective surface containing two black and white regions separated by a straight step edge, placed under the sample.
- *Camera.* Located about 1 m above the sample, vertically aligned and laterally centered with the observed step edge, imaging the sample at 28 µm resolution.
- *Lighting.* Two diffuse, neutral illuminants, symmetrically placed around the sample. Positioned roughly along the bisector between the camera's optical axis and the sample plane, to minimize the influence of single-scattering events while still providing sufficiently strong and even illumination.
- *Thin glass slip.* To provide a flat, clear, smooth air interface with roughly known BRDF on top of the sample.

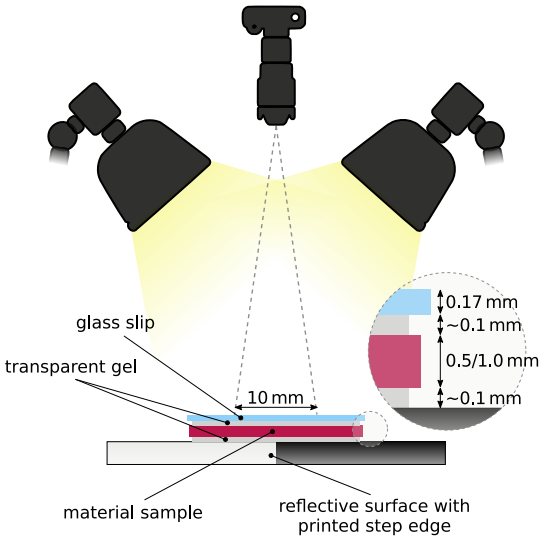


FIGURE 3.4: *Sketch of our acquisition setup (not to scale).*

- *Transparent gel.* For (approximate) index matching between the components, to suppress internal reflections and the influence of dust or residual roughness on the sample surface.

We built a physical version of this setup using a consumer DSLR camera (Canon EOS 700D body, EF-S 18–135 mm lens set to 135 mm), two 55 W, 5500 K-equivalent fluorescent lamps, thin microscope glass slides and a clear ultrasound transmission gel, chosen for its bubble-free applicability [Donner et al. 2008]. We printed the fitting edge on a Toughprint Waterproof paper using an office laser printer. The resulting edge appearances captured by our physical setup are shown in Figure 3.5, left.

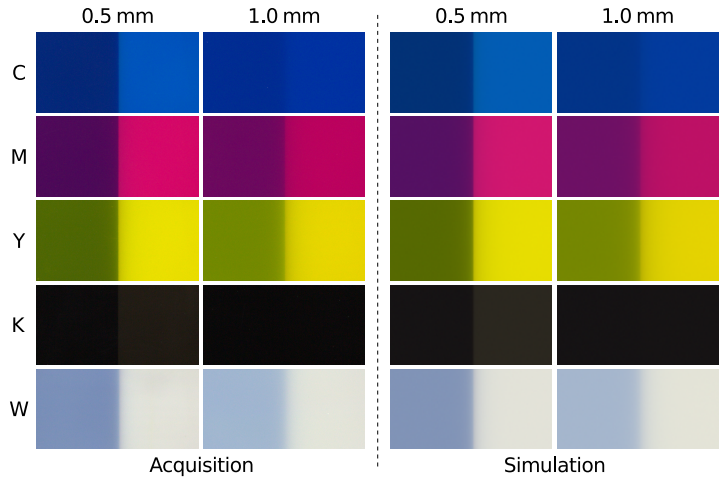


FIGURE 3.5: *Comparison of the step edges captured through the printing materials* $\mathcal{M} = \{\mathrm{C}, \mathrm{M}, \mathrm{Y}, \mathrm{K}, \mathrm{W}\}$ *by our physical (**left**) and virtual (**right**) setups, at two different sample thicknesses. The images cover an area of 2×1 cm.*

As for the virtual counterpart of this setup, we use the same environment as for the prediction (Section 3.2): a hand-modeled scene for the Mitsuba renderer, with path-tracing as a simulation framework. We generated the fitting dictionary for two different sample thicknesses – 0.5 mm and 1 mm – by sampling the space $\mathcal{V}$ as follows: $\sigma_\mathrm{t} \in \{1, \ldots, 30\}\,\mathrm{mm}^{-1}$, $\alpha \in \{0.05, \ldots, 0.99999\}$ and $g \in \{0, \ldots, 0.8\}$ (see the supplemental material for the complete enumeration). For the refractive index we used the value of $\eta = 1.5$ measured separately by optical ellipsometry. The resulting dictionary contains about 34k records and took two days on a 500-node CPU cluster to compute. We will make the dictionary publicly available, along with the detailed documentation of the data format.

A desired property of our setup is that it uses the same simulation framework (Monte Carlo path tracing) as a prediction model for the measurement as well as our optimization itself. This guarantees consistency of the obtained data and the resulting visual match.

3.3.2 Model Fit

Our fitting procedure operates in the space of 1D *edge profiles* $\boldsymbol{F} \in \mathcal{C}$ resulting from vertically averaging the captured edge appearances. We then find the *most similar* simulated edge profile $\boldsymbol{F}_\mathrm{S}$ for each profile $\boldsymbol{F}_\mathrm{A}$ acquired by our physical setup. This is done jointly over both acquired sample thicknesses $t \in T$ (in our case $T = \{0.5, 1\}$ mm).

We formally define the fitting as a minimization across the volume parameter space $\mathcal{V}$:

$$\underset{(\sigma_t,\alpha,g)\in\mathcal{V}}{\arg\min} \sum_{t\in T} \measuredangle\left(\mu \cdot \boldsymbol{F}_{A|t},\ \boldsymbol{F}_{S|t}(\sigma_t,\alpha,g)\right), \tag{3.2}$$

where μ is a scaling coefficient which we explain below. The function $\measuredangle$ is a distance metric in the edge profile space, and we define it as

$$\measuredangle(\boldsymbol{F}_A,\boldsymbol{F}_S) = \left\|\boldsymbol{F}_A - \boldsymbol{F}_S\right\|_2 + \lambda \cdot \left\| \left|\nabla \boldsymbol{F}_A\right| - \left|\nabla \boldsymbol{F}_S\right| \right\|_2, \tag{3.3}$$

with λ being a scalar parameter. Solving Equation (3.2) for every material—and for each RGB channel—leads to an optical characterization of the whole material space $\mathcal{M}'$. The solution itself is efficiently implemented as an exhaustive search in the material dictionary (Section 3.3.1), taking only a few minutes in Matlab on a desktop PC.

To obtain robust results, we regularize the fitting via the two scalar parameters μ and λ. While λ regularizes the fit locally by seeking a match in the profile gradients (in addition to their intensities), μ is a global regularizer that compensates for any possible global bias resulting from modeling mismatches between the physical and virtual setups. We determined $\lambda = 20$ empirically and $\mu = 0.913$ systematically – please refer to the supplemental material for details.

The simulated edge appearances of the best dictionary fits according to Equation (3.2) are shown in Figure 3.5, right.

3.3.3 Results

To our surprise, we found that the edge appearance profiles follow the predictions of the similarity theory [Wyman et al. 1989; Zhao et al. 2014] to a good accuracy, even in our thin-slice conditions. The main implication is that we can find a reliable fit of extinction coefficient σ_t and albedo α for *any* desired anisotropy g present in the dictionary. This benefits us greatly, because constraining g significantly simplifies the construction of our tonal mapping, as described later in Section 3.4.1.

We therefore performed a search across all g values in our dictionary and found that $g = 0.4$ yields the lowest overall fitting error. The resulting measurements are summarized in Table 3.2 and validated in Figure 3.5 for the thin slices and later in Section 3.7 for complex heterogeneous material compositions.

The main limitation of our setup is the measurement of strongly absorbing materials, as the resulting edge profiles have almost no discriminability in that part of $\mathcal{V}$. This turned out to be an issue for the R channel of cyan, G channel of magenta and B channel of the yellow material (marked bold in Table 3.2). We have resolved this issue by printing a color chart with patches of all possible *equal* combinations of

TABLE 3.2: *Measured optical parameters for the Stratasys Vero Opaque family of materials, constrained for scattering anisotropy $g = 0.4$.*

Material	Ext. coef. σ_t [mm^{-1}]			Scattering albedo α		
	R	G	B	R	G	B
Cyan	**9.0**	4.5	7.5	**0.05**	0.7	0.98
Magenta	2.5	**3.0**	10.0	0.98	**0.1**	0.9
Yellow	2.25	3.75	**19.0**	0.997	0.995	**0.15**
Black	5.0	5.5	6.5	0.35	0.35	0.35
White	6.0	9.0	24.0	0.9991	0.9997	0.999

the CMYKW materials and their subsets (shown later in Figure 3.8). We then ran a brute-force search in the problematic channels to identify their values, by seeking a best fit of our color predictions to the printed color chart, using our tonal mapping from Section 3.4.1 as a prediction model for efficiency sake.

3.4 Material Mapping

At the center of our method is the requirement to repeatedly transform between colorspace appearance definitions and material definitions, during the overarching optimization (Section 3.5). Here, we present efficient means to accurately convert between the different working spaces listed in Table 3.1.

For conceptual reasons, we describe our pipeline in terms of two modes of operation, where the former is a subset of the latter as illustrated in Figure 3.6:

- a direct, 'single-pass' appearance reproduction mode, agnostic of the global light transport within the printing materials, and
- an iterative, 'multi-pass' optimization-driven mode, which makes use of the direct mode to produce an improved match between material composition and the target appearance.

Qualitatively the direct mode is compatible with other state-of-the-art 3D color reproduction pipelines [Babaei et al. 2017; Brunton et al. 2015], although it needs be noted that our solution has a markedly different technical basis consistent within our scattering-aware optimization as a whole.

We now delineate the steps involved in our reproduction pipeline (also refer to Figure 3.6), followed by detailed descriptions of the more complex mapping procedures in Sections 3.4.1 and 3.4.2.

1. The target image $\boldsymbol{T}$ is vertically extruded down to a desired depth in the RGB working volume $\boldsymbol{X} \in \mathcal{C}$; the remainder of $\boldsymbol{X}$ is filled with white. The working volume $\boldsymbol{X}$ is then per-voxel transformed into the tonal space volume $\boldsymbol{M}$ using the $\mathcal{C} \rightarrow \mathcal{M}$ mapping detailed in Section 3.4.1. Note this also implicitly performs the gamut projection of the target image $\boldsymbol{T}$ itself, which then becomes the actual target in the algorithm.

2. The working volume $\boldsymbol{M} \in \mathcal{M}$ is quantized (half-toned) so that each voxel is exclusively assigned a single material from the working set $\mathcal{M}'$, as detailed in Section 3.4.2. This already constitutes a printable result, concluding the direct reproduction mode.

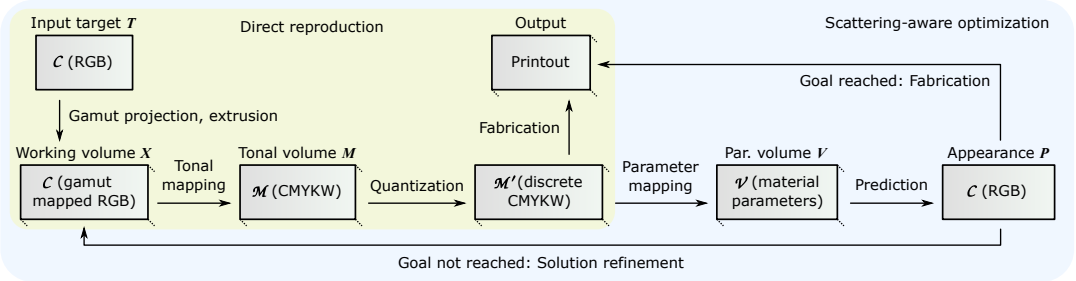


FIGURE 3.6: *Our reproduction pipeline, described in detail in Sections 3.4 and 3.5. Note the schematic distinction between 2D and 3D working data.*

3. Each voxel $\boldsymbol{m} \in \boldsymbol{M}$ is assigned three optical parameter vectors $v_m \in \mathcal{V}$ (one for each RGB channel) measured for its respective material $m \in \mathcal{M}'$.

4. The forward prediction model $f(\boldsymbol{V}) \to \boldsymbol{C}$ (described in Section 3.2) is used to compute the appearance of the working solution.

5. The predicted appearance image $\boldsymbol{C}$ is used to adjust the current solution in the optimization phase, which we describe in Section 3.5. This process is iteratively repeated until convergence, as sketched in Figure 3.6.

3.4.1 Tonal Mapping

Tonal mapping (also known as color separation) $\mathcal{C} \to \mathcal{M}$ is an important component in any appearance reproduction pipeline, since it amounts to a colorimetric characterization of the physical output device. Previous, dictionary-based approaches to establishing the mapping between the tonal space $\mathcal{M}$ and the input space $\mathcal{C}$ (e.g., [Babaei et al. 2016; Brunton et al. 2015]), however, require a high number of measurements, becoming impractical for higher-dimensional tonal spaces. In addition, as Brunton et al. explain, the translucent nature of our materials complicates the measurements. We therefore opt for a computational *predictive* approach instead.

We determine the mapping $\mathcal{C} \to \mathcal{M}$ through inversion of the mapping $\mathcal{M} \to \mathcal{C}$. Although this mapping corresponds to our forward appearance model (Section 3.2), its application would be very expensive for this purpose, for virtually the same reasons why the data-driven approaches are. We have therefore developed a tailored, analytic prediction model instead.

Albedo Mapping We can observe that for a semi-infinite, homogeneous medium with a smooth interface, its diffuse reflectance (i.e., color $\mathcal{C}$) depends on the refractive index η, scattering anisotropy g, and most importantly, the scattering albedo α. Crucially, the diffuse reflectance is *independent* on the optical density σ_t of the medium, as this only scales the sub-surface transport uniformly in all directions.

Consequently, by fixing the values of η and g we can obtain a bijective mapping $\alpha \leftrightarrow \mathcal{C}$. We use the measured $\eta = 1.5$, and $g = 0.4$ as justified in Section 3.3.3 (and further in the supplemental material).

Using the above canonical parametrization, we obtained the reflectance data via an analog brute-force Monte Carlo simulation. We computed 1024 data points (distributed uniformly in the α domain), and numerically fit them to an empirically chosen functional model, using simulated annealing under the L^2 norm. The resulting forward mapping $\alpha \to \mathcal{C}$ (plotted in Figure 3.7) is

$$\mathcal{C}(\alpha) = \mathcal{C}_s + (1 - \mathcal{C}_s) \cdot \sum_{k=1}^{K} a_k \cdot \alpha^{b_k}, \tag{3.4}$$

where $\mathcal{C}_s$ is the reflectance of the material's interface (i.e., the portion of energy that does not participate in the sub-surface transport) and equals $\mathcal{C}_s = 0.04526$ under the chosen parametrization.

For $K = 5$, we obtain the average absolute error of $3.9 \cdot 10^{-4}$ and a peak absolute error of $5.7 \cdot 10^{-3}$ with the parameters

$$a_k = \{0.065773,\ 0.201198,\ 0.279264,\ 0.251997,\ 0.201767\},$$
$$b_k = \{1.569383,\ 6.802855,\ 28.61815,\ 142.0079,\ 1393.165\}.$$

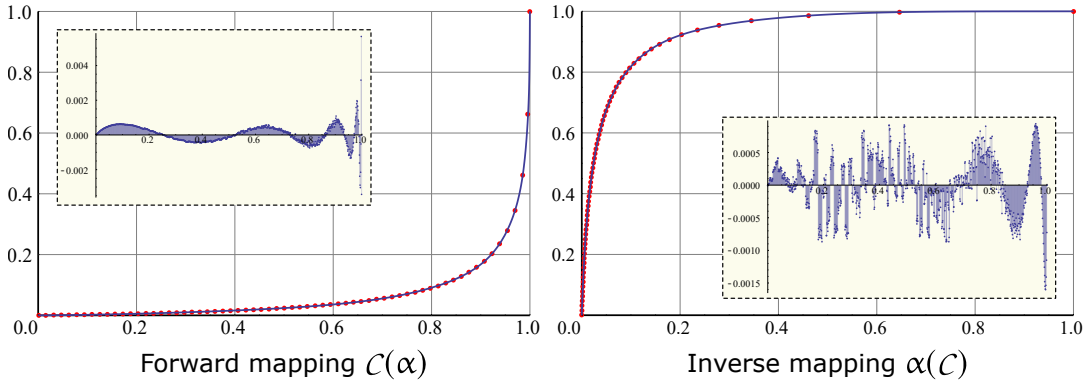


FIGURE 3.7: *Plots of the bi-directional mapping between scattering albedo α and color $\mathcal{C}$ (solid blue lines). The red points show the tabulated fitting data (sub-sampled by the factor of 16 for clarity). Note that the plots are normalized to [0,1], i.e., show only the volumetric portion of the reflected energy. The insets show the residual errors of the fits.*

For completeness we also computed the inverse mapping $\mathcal{C} \to \alpha$:

$$\alpha(\mathcal{C}) = \sum_{k=1}^{K} c_k \left(1 - \left(1 - \frac{\mathcal{C} - \mathcal{C}_\mathrm{s}}{1 - \mathcal{C}_\mathrm{s}}\right)^{d_k}\right). \tag{3.5}$$

For $K = 5$, we obtain the average absolute error of $3.5 \cdot 10^{-4}$ and the peak absolute error of $1.6 \cdot 10^{-3}$ with the parameters

$$c_k = \{0.163581,\ 0.391943,\ 0.029277,\ 0.316847,\ 0.098352\},$$
$$d_k = \{3.969542,\ 15.94272,\ 46.26871,\ 59.95706,\ 206.5716\}.$$

Note that Equation (3.5) is not a simple inversion of Equation (3.4); rather, a separate fit is necessary to obtain it. This inverse mapping is plotted alongside the forward one in Figure 3.7.

Our approach is conceptually similar to [Papas et al. 2013], who however rely on tabulated data for the mapping. We found our analytic fit to be more accurate, particularly in high-α regions where the non-linearity of the space is especially high. It is worth noting that Pharr et al. [2016] also obtain an analytic expression for the mapping, but use photon beam diffusion [Habel et al. 2013] to obtain the fitting data, which in itself is an approximate technique.

Mixture Mapping The mapping $\alpha \to \mathcal{C}$ (Equation (3.4)) enables us to define the forward mapping $\mathcal{M} \to \mathcal{C}$ by using the measured optical parameters $\{v_m\}$ of the printing materials $m \in \mathcal{M}'$. For this we convert the optical parameters of the continuous tonal mixture $w \in \mathcal{M}$ into the *linear space of absorption and scattering coefficients* σ_a and σ_s (which fundamentally express the respective particle concentrations) as

$$\sigma_{\mathrm{a}|w} = \sum_{m \in \mathcal{M}'} w_m \cdot (1 - \alpha_m) \cdot \sigma_{\mathrm{t}|m} \tag{3.6}$$

$$\sigma_{\mathrm{s}|w} = \sum_{m \in \mathcal{M}'} w_m \cdot \alpha_m \cdot \sigma_{\mathrm{t}|m}, \tag{3.7}$$

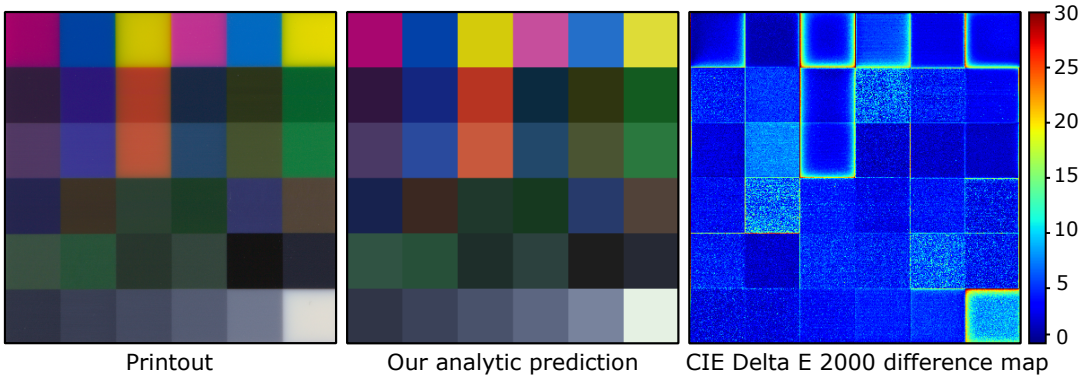


FIGURE 3.8: *Sampling of the CMYKW space with equal weights w_m for all material subsets, plus a grayscale palette. Our analytic prediction matches the printout well, especially when accounting for the different perception of the printout arising from the color bleeding between the patches.*

where $\alpha_m, \sigma_{\mathrm{t}|m} \in v_m$, i.e., the measured material albedos and extinction coefficients. The resulting mixture albedo is then

$$\alpha_w = \frac{\sigma_{\mathrm{s}|w}}{\sigma_{\mathrm{a}|w} + \sigma_{\mathrm{s}|w}}. \tag{3.8}$$

Applying Equation (3.4) yields the prediction for the mixture color, the accuracy of which is demonstrated in Figure 3.8.

Tonal Mapping The ability to efficiently predict the color for any continuous material mixture $w \in \mathcal{M}$ allows us to build the inverse mapping $\mathcal{C} \to \mathcal{M}$. We densely sample the 5D redundant space $\mathcal{M}$ by first sampling the weights for the CMYK materials and then computing the weight for the white material as $w_{\mathrm{W}} = 1 - \sum_{m \in \mathrm{CMYK}} w_m$, discarding all mixtures with $\sum w_m > 1$. For efficiency sake, we found it advantageous to bias the sampling of the CMYK dimensions towards lower concentrations, as the naïve uniform sampling yields a highly non-uniform distribution of the mapped points in $\mathcal{C}$ as a result of the strong non-linearity of the albedo space.

The mapping $\mathcal{C} \to \mathcal{M}$ then amounts to the application of Equation (3.8) and Equation (3.4) to the values $\in \mathcal{M}$ obtained by the above sampling procedure, and finding the closest point in the cloud of the resulting values $\in \mathcal{C}$. This step also implicitly performs the gamut mapping of the initial target image. The distance metric we use here is a Euclidean distance in the sRGB space (instead of the linear RGB space $\mathcal{C}$), due to its better perceptual uniformity. We cache this mapping in a 3D table with 256 bins in each of the RGB dimensions, and accelerate the point cloud search by building a kD-tree over it. A sparse sub-sampling of the mapped point cloud is shown in Figure 3.9.

The ability of our $\mathcal{C} \to \mathcal{M}$ mapping to evenly cover and reproduce the RGB color gamut is demonstrated in Figure 3.10.

As we discuss further in Section 3.5 the accuracy of color versus structure reproduction is, to a certain extent, a matter of tradeoff. This is directly visible in our results in Figure 3.10, but also in the corresponding numerical difference figures presented separately in Figure 3.11. In any case, both our direct-mapped and optimized results achieve a better color reproduction accuracy than the default print.

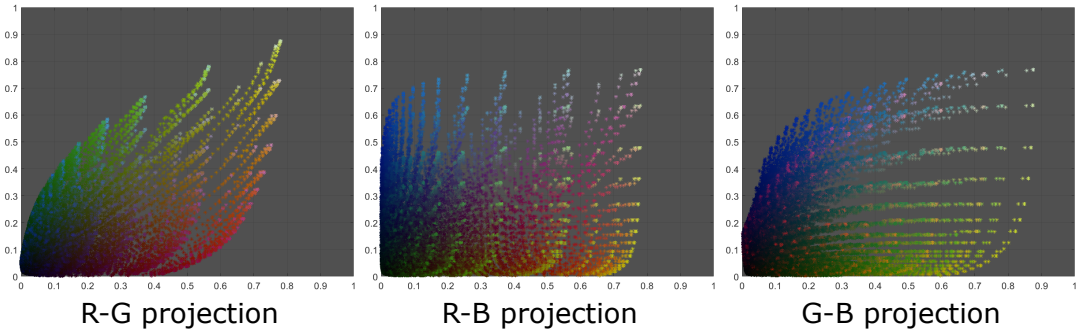


FIGURE 3.9: *Visualization of the CMYKW gamut in linear RGB (i.e., the mapping $\mathcal{M} \to \mathcal{C}$ described in Section 3.4.1). Shown is a small set of 30k points, out of the 32M points we used to construct the mapping.*

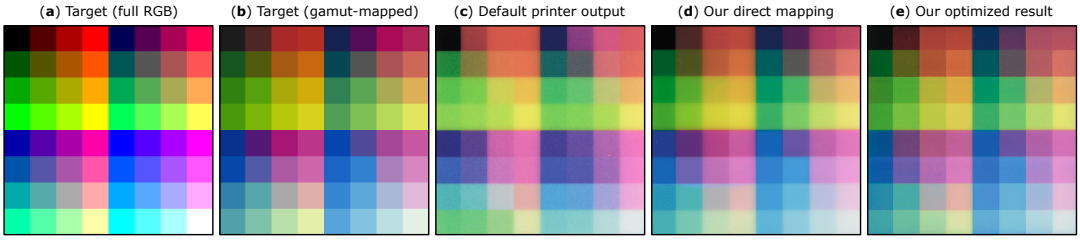


FIGURE 3.10: *Reproduction of the RGB gamut. We regularly sample the full RGB space $\mathcal{C}$ (**a**) and print it in the default printer color reproduction mode (**c**). We then print the gamut-mapped target (**b**) using both our direct reproduction method from Section 3.4 (**d**) and the optimized method from Section 3.5 (**e**). Our direct reproduction achieves a more consistent target match than the default print, further improved by the optimization in terms of both structure.*

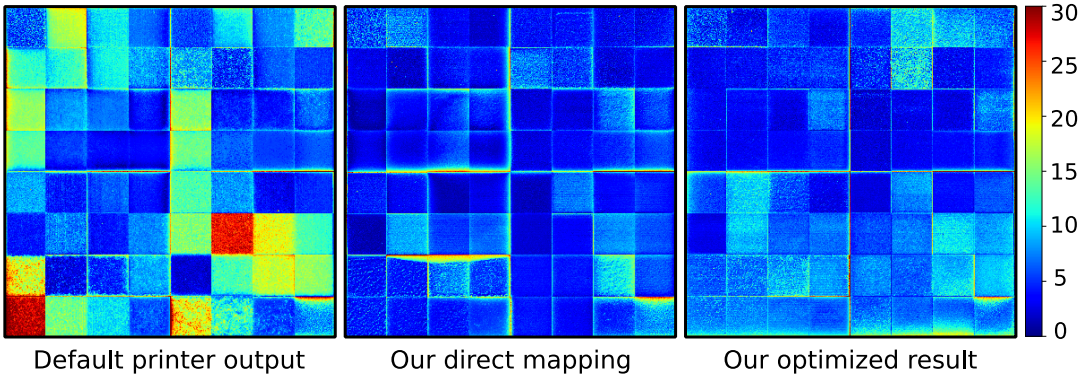


FIGURE 3.11: *Difference maps between the RGB color cube printouts (Figure 3.10, **c**, **d** and **e**) and the gamut-mapped target (Figure 3.10, **b**). We use the CIE Delta E 2000 color difference metric. The corresponding average color differences are 9.2, 4.1 and 5.7, respectively.*

3.4.2 Tonal Quantization

Tonal quantization essentially amounts to half-toning. We have compared the per-layer 2.5D approach of Brunton et al. [2015] with a full 3D half-toning [Lou et al. 1998]. We found that while the 3D half-toning approach leads to a slightly better structure preservation, the 2.5D variant leads to better colors and presents an overall better tradeoff in our specific case. We therefore use the method of Brunton et al. [2015] in all our results.

An interesting alternative would be combining our method with the recent contoning technique of Babaei et al. [2017], to suppress the quantization noise inherent to any half-toning. We see this as an interesting direction for future work.

3.5 Structural Optimization

In this section we describe the core, optimization-driven part of our pipeline (cf. Figure 3.6), which builds on the direct reproduction part detailed in Section 3.4. To recap, the purpose of this part is to iteratively refine the solution towards a better structural match with the target, by countering the distortions caused by sub-surface light transport.

3.5.1 Problem Definition

We seek a volumetric assignment of print materials $\boldsymbol{M} \in \mathcal{M}'$ that minimizes the distance from an opaque diffuse appearance goal represented by the target image $\boldsymbol{T}$ (Section 3.4):

$$\arg\min_{\boldsymbol{M}\in\mathcal{M}'} \mathcal{d}\left[\mathcal{f}(\boldsymbol{M} \to \boldsymbol{V}), \boldsymbol{T}\right], \tag{3.9}$$

where $\mathcal{d}$ is a distance metric. Because of the complex nature of the prediction operator $\mathcal{f}$ (Section 3.2), it is challenging to solve this problem using standard optimization techniques.

First, in spite of some similarities it is important to point out that our problem cannot be cast as a deconvolution. The principal difference is that deconvolution methods *remove* distortions that have already impacted the signal in question, while our solution has to embed the correction into the signal *a priori*. While these operations commute in the mathematical sense, they certainly do not in physical conditions where the signal is limited to a certain narrow gamut. A good demonstration of this problem has been given by Montalto et al. [2015] in the context of visual acuity enhancement. In comparison to them, we face two additional significant issues: our 'blurring kernel' is typically very asymmetric, and has a strong material-dependent spatial variation.

The complex non-linear nature of our appearance formation process, combined with the sheer size of our search space, also makes the application of generic optimization solvers very difficult. For instance, even a small 5×5 cm printout will have the lateral resolution of 600^2, with up to 50 layers modified by the optimization, corresponding to 18M voxels. This led us to develop a more specialized approach built around the properties of light transport in our conditions. Specifically, we seek physically justified heuristics that will significantly prune our search space yet leave enough freedom to find meaningful solutions.

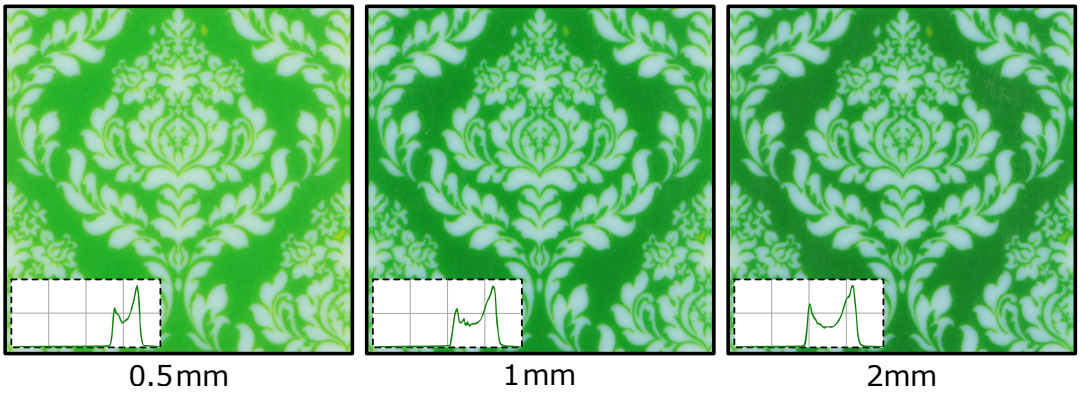


FIGURE 3.12: *The effect of varying the thickness of the colored layer (underlied by pure white material). Notice how the material is overall darkening and the undesirable color bleeding into white is increasing when more absorbing material is present. This manifests in shifting the green channel histogram towards lower values and spreading its width as shown in the insets.*

3.5.2 Solution Concept

Our solution's central idea comes from the observation that color fidelity and structural fidelity are actually competing requirements. This is demonstrated in Figure 3.12, showing a colored binary pattern vertically extruded into different depths (which we will refer to as *characteristic depth*). We can clearly see that increasing the characteristic depth trades structural acuity (which would peak in the extreme case of a single colored layer, as this would lead to the smallest possible color bleeding) for color saturation (which is directly proportional to the amount of applied colored material). The latter observation agrees with Brunton et al. [2015] who report an expansion of the color gamut with the characteristic depth increase.

The challenge lies in generalizing the above qualitative view to arbitrary target structures. It is clear that we need to consider a unique vertical material variation ('stacking') for every surface point. However, the appearance at every surface point is influenced by a large neighborhood, due to the relatively large diffusion distance of our materials. Consequently, we cannot optimize for each stacking in isolation (as was possible in previous approaches [Dong et al. 2010; Hašan et al. 2010] due to band-limited input specifications), but instead optimize *jointly* for the full 3D structure (Equation (3.9)) and iteratively refine it based on the global appearance prediction given by the operator f (Equation (3.1)).

3.5.3 Optimization

Based on the discussion in Section 3.5.2 the following guidelines apply while optimizing the 3D voxel structure of a printed object. First, we minimize the amount of absorbing material in each channel to reduce the undesired color bleeding and excessive material darkening (Figure 3.12). We refer to absorbing materials as 'dark', since they reduce the amount of light either selectively per channel (CMY materials) or overall (K material). Second, we concentrate the dark materials as close as possible to the object surface to minimize their impact on neighboring regions' colors. This effectively means that darker materials are always placed atop of lighter ones. Third, at edges between dark and light materials, we *erode* the dark material as a function of depth to further limit the lateral energy absorption (cf. Figure 3.13). As a result a tapered volume of the dark material is obtained, which helps retaining the contrast and symmetricity of the edge at the intended position.

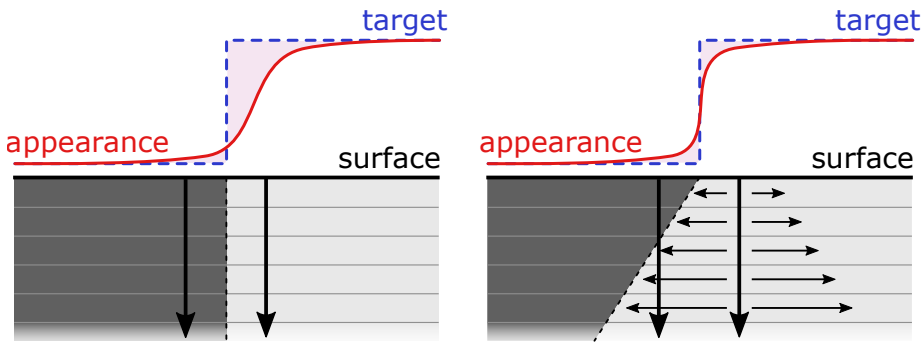


FIGURE 3.13: *Erosion during traversal. The benefit of progressively eroding the darker materials (**right**) compared to simple extrusion (**left**) is a better balance of the appearance 'profiles' around edges: the result overall matches the light side of the edge much better, at the cost of just a slightly worse match in the dark side.*

Outline Accounting for these guidelines, our solution starts from the top material layer and proceeds to subsequent layers as follows:

1. Assign a default material (typically plain white).
2. For a given iteration, express the current solution error as a signed energy difference between the target appearance and the prediction from the current iteration. If the error did not decrease since the last iteration, output the best solution so far.
3. Propagate the error vertically into the object, starting at the top layer. To account for the energy interchange between the neighboring stackings, diffuse the undeposited positive error by a small erosion kernel that increases with depth to achieve the tapered volume effect (Figure 3.13).
4. If a negative error cannot be deposited into the current layer ('too dark' case), and vice-versa for a positive error ('too light' case), proceed to the next lower layer.
5. If all error energy has been deposited, initiate a new iteration and continue to Item 2.

We now detail the key aspects of the procedure.

Representation We represent the working solution by an RGB proxy volume $\boldsymbol{X} \in \mathcal{C}$ (also see Figure 3.6). The advantage over representing it in the tonal space $\mathcal{M}$ (or even its discrete version $\mathcal{M}'$) is its linearity, and the fact that being expressed in the same space as both the target and the appearance prediction allows us to project the prediction gradient directly back into the solution. We store the proxy and hence compute the entire optimization in the native printer resolution.

Solution Refinement The prediction gradient (i.e., the residual difference between the prediction and the target) is used to adjust the solution. Since we represent the solution in the linear space $\mathcal{C}$, the adjustment is directly proportional to the gradient. The main challenge is the dimensionality mismatch: the gradient is a 2D image, while the solution a 3D volume.

In the refinement procedure (Algorithm 1), we divide the residual energy into its positive (the solution needs to be lightened) and negative (the solution needs to

Algorithm 1: Solution refinement (optimization core). All the operations are applied on entire matrices (images), i.e., separately for each channel and each vertical stack of voxels corresponding to a given surface pixel.

Input:
$\boldsymbol{X}$: current working (proxy) volume $\in \mathcal{C}$
$\boldsymbol{T}$: gamut-projected target image $\in \mathcal{C}$
$\boldsymbol{P}$: current appearance prediction image $\in \mathcal{C}$
Result:
$\boldsymbol{X}$: updated working volume $\in \mathcal{C}$
begin
  // Calculate positive and negative differences from target.
  $\boldsymbol{D} = \boldsymbol{T} - \boldsymbol{P}$
  $\boldsymbol{D}^{+} = \max(\boldsymbol{D}, 0)$
  $\boldsymbol{D}^{+}_{\text{res}} = 0$
  $\boldsymbol{D}^{-}_{\text{res}} = \min(\boldsymbol{D}, 0)$
  // Loop over all proxy volume layers, starting from the topmost one, down to the last layer Z.
  for $z = 0..Z, \boldsymbol{L}_z \in \boldsymbol{X}$ **do**
    // Conservatively darken the layer proportionally to c_{p}.
    $\boldsymbol{L}_z \mathrel{+}= c_{\text{p}} \cdot \boldsymbol{D}^{-}_{\text{res}}$
    // Lighten the layer in the full extent.
    $\boldsymbol{L}_z \mathrel{+}= \boldsymbol{D}^{+}$
    // Add any remaining residual energy from the above layers, diffusing it laterally with the 2D kernel κ_{e}.
    $\boldsymbol{L}_z \mathrel{+}= \kappa_{\text{e}}(\boldsymbol{D}^{+}_{\text{res}}, z)$
    // Calculate residuals for the next layer. $\boldsymbol{D}^{-}_{\text{res}}$ is the residual darkening energy from the current layer compensated for the absorbance of the above layers with c_{a}.
    The lower bound for $\boldsymbol{D}^{-}_{\text{res}}$ is chosen to enable enough energy propagation while avoiding divergence.
    $\boldsymbol{D}^{+}_{\text{res}} = \max(\boldsymbol{L}_z, 1) - 1$
    $\boldsymbol{D}^{-}_{\text{res}} = \text{clamp}(c_{\text{a}} \cdot \boldsymbol{L}_z, -2, 0)$
    // Ensure the layer stays in the gamut.
    $\boldsymbol{L}_z = \text{clamp}(\boldsymbol{L}_z, 0, 1)$
  end
end

be darkened) parts. We then traverse the solution vertically downwards, and add the positive residual aggressively to each layer, while the negative residual is applied conservatively and only up to the minimal necessary depth to minimize color bleeding (in line with the heuristics from the beginning of this section). We also slightly diffuse the positive residual during the traversal, which preserves the dark composition in homogeneous areas, but erodes it in the vicinity of structural gradients where an excess of dark material would cause the majority of color bleeding.

The described behavior is controlled by three empirical parameters (cf. Algorithm 1).

- $c_{\text{p}} \in (0, 1)$: proportionality coefficient. Conservatively tempers the darkening of the layering. We have experimentally determined $c_{\text{p}} = 0.5$ to yield good and stable results.

- $c_{\text{a}} \in [1, \text{inf})$: absorbance coefficient. Approximates the lower visibility of a given layer through the stack above and compensates for it. Since by design it cannot happen that a dark voxel is placed underneath a stack of lighter voxels, we

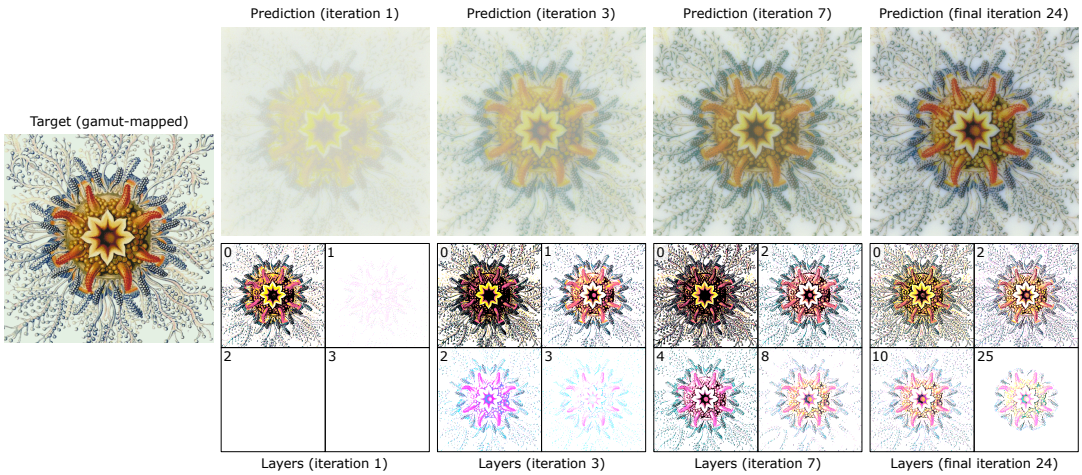


FIGURE 3.14: *Convergence of the optimization (Section 3.5) for a complex target. For a few selected iterations we visualize the distributions of the CMYKW materials in several selected layers under the surface (the z-index of each layer is shown in the upper-left corners). In spite of some residual color bleeding that could not have been compensated for, the optimization finds an overall good match to the target.*

apply this compensation to the negative energy only. We obtain the value of c_a by averaging the absorption coefficients of all significantly absorbing print materials $\in \mathcal{M}'$ and by computing the resulting transmittance of a single layer of such virtual material.

- k_e: 2D erosion kernel. Diffuses the positive residual energy to progressively suppress dark material assignment near structural gradients. For simplicity we use a small Gaussian kernel with the standard deviation equal to half the respective layer depth z.

We use the same set of parameters to generate all our results.

Stopping Criterion We stop the outer loop of the optimization if the value of our distance metric d (see Equation (3.9)) does not change significantly between two successive iterations. We use the structural similarity index (SSIM, [Wang et al. 2004]) as a metric, as its performance is usually not influenced by the Monte Carlo noise inherently present in the appearance prediction $\boldsymbol{P}$.

3.5.4 Discussion

Thanks to the simplicity of our core refinement procedure, it is very efficient. Since the radius of the erosion kernel k_e is bounded by the characteristic depth, the complexity of the refinement is in fact linear w.r.t. number of voxels in the working volume $\boldsymbol{X}$. The overall optimization process typically converges in about 20–25 iterations; we show an example of the convergence process for a complex input in Figure 3.14.

While we require the specification of some empirical parameters, they themselves have an intuitive meaning and, once set up, worked well for all our results without any modifications. We hypothesize their values could be automatically derived from the optical parameters of the printing materials (as we already do for c_a), but leave this to future investigation.

3.6 Fabrication

We printed all our results on the Stratasys J750 poly-jetting printer [Stratasys 2017] using a manual voxel printing mode provided by the manufacturer that allows us to address native output voxels directly. The tonal space $\mathcal{M}'$ corresponds to the CMYKW printing materials from the Vero Rigid Opaque family (specifically Black-Plus and PureWhite for KW). For the comparisons in Section 3.7, we additionally use the standard texture print mode of the machine, which can emulate up to 360k unique tonal mixtures (dubbed 'digital multi-materials' by the manufacturer). The lateral printing resolution was 300 DPI (i.e., 84.7 µm per voxel), and the vertical resolution 27 µm per layer. The total thickness of all our samples was 10 mm, which is the peak penetration distance we determined for our prints (by shining a white LED light through them, in a dark room).

While the manual voxel print gives us full control over the material composition, it produces raw printouts without a controlled surface finish. The surface is also polluted by the support material the printer automatically applies, which significantly distorts the appearance. To counteract this, we print a few extra transparent layers (using the Vero Clear material) on the sample surface, and scrub the residual support material with a grade-1000 sandpaper. We then spray each printout with a transparent lacquer to achieve a smooth and reproducible finish.

3.7 Results

We first list several parameters and statistics of our method.

- Our entire pipeline operates on the native printer resolution, which is 300 DPI laterally and 900 DPI vertically. The resulting voxel size is therefore $85 \times 85 \times 27$ µm.
- Our Monte Carlo prediction model takes about three minutes to compute for a 500×500 image with 500 samples per pixel, on a cluster with 100 quad-core Intel Xeon E5620 CPUs. The remaining steps of our pipeline take less than a minute to compute in a hybrid Python/C++ implementation. For a volume with $500 \times 500 \times 370$ voxels ($42 \times 42 \times 10$ mm) the optimization takes an hour to compute, on average, corresponding to approximately 20 iterations.
- Photographs of the printouts were taken using the same setup (camera, lens, lights) as for the calibration (Section 3.3). The white balance was set up using X-Rite ColorChecker Passport. We ensured consistency among all the printouts' photos, but the capturing system as a whole was not fully color calibrated, that is why the actual color perception in the images might differ slightly from the targets.

Prediction Accuracy As a first step we need to validate the accuracy of our prediction (Section 3.2) using the material data measured in Section 3.3. In Figure 3.15 we show several examples that we can indeed predict the appearance of our prints. A similar match is achieved for the remaining results in the paper.

Optimization Consistency Given that our target inputs cannot often be reproduced perfectly (i.e., there will always be some amount of color diffusion in translucent materials), we want to verify that our optimization can reach a goal that is guaranteed to be achievable. For this purpose we generate a very specific volumetric composition:



FIGURE 3.15: *Predictions of our printouts' appearance using the measured data from Table 3.2. The slight blurriness of the printouts' images is due to the band-limitedness of our optical capture system, plus any residual surface roughness which our surface treatment (Section 3.6) cannot eliminate entirely.*



FIGURE 3.16: *Validation of our optimization. Our method can reproduce target appearances arising from complex, randomly assigned material distributions, even if they by themselves are not explored by the optimization.*

randomized, colored characters with different thicknesses assigned up to 60 layers under the surface. We then render our prediction of this volume and use *this* image as a target for our optimization (which gives us the guarantee such appearance is in fact reproducible). Figure 3.16 shows that we can in fact achieve such appearances, even though they arise from a complex volumetric structure that by itself is not explored by our optimization.

Reproduction Evaluation The main results of our method are presented in Figure 3.17. We aimed to reproduce various different motifs. In the order of listing: two detailed drawings (*Flower* and *Marine*), two impressionist paintings with structured strokes (*Dusk* and *Tree*), and two highly heterogeneous natural materials (*Leaves* and

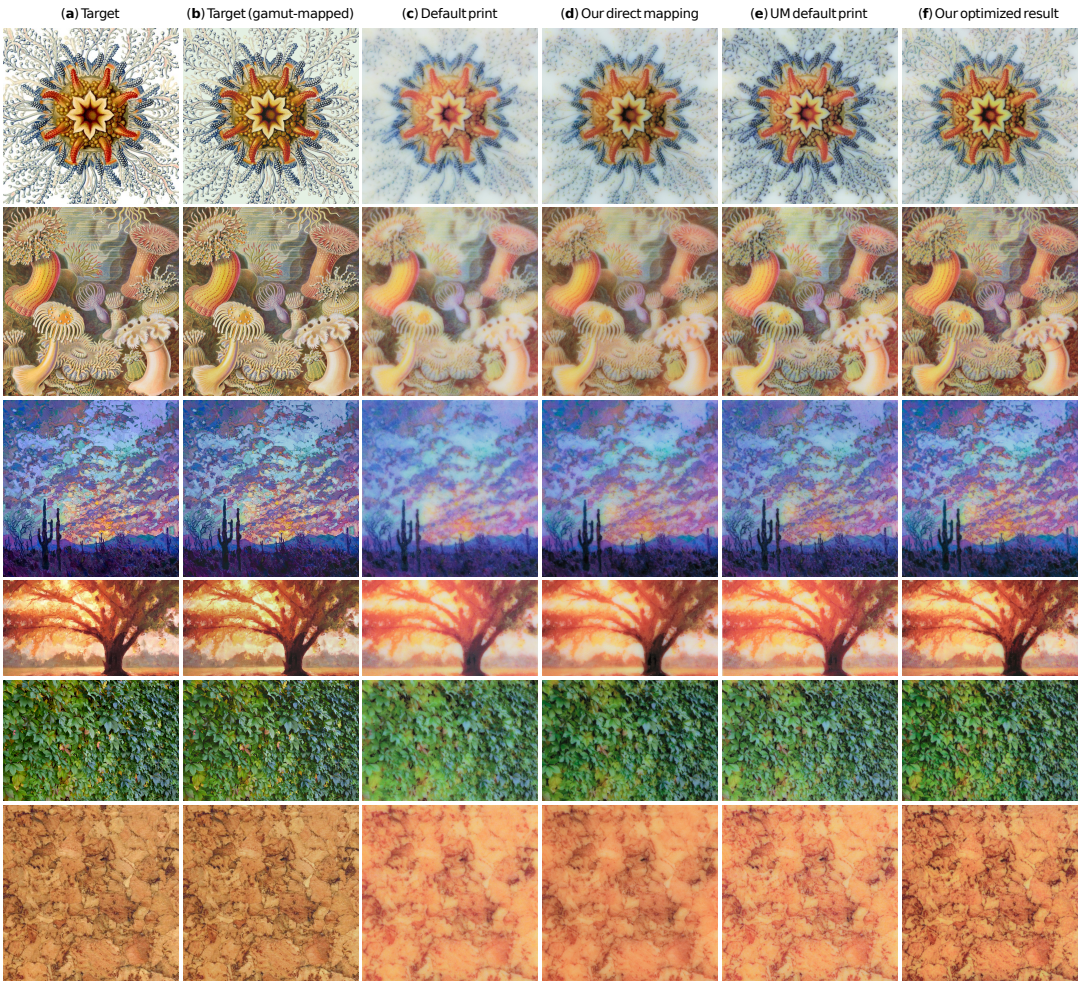


FIGURE 3.17: *Evaluation of the main results of our method. We compare the default printer output (**c**) and its sharpened version (**e**) which uses unsharp masking (UM), to the results of our method when using the direct reproduction mode (**d**) and the scattering-aware optimization (**f**). Results (**c**) and (**e**) were produced using the original target (**a**) as an input, while for the results (**d**) and (**f**) we used our gamut-mapped target (**b**).*

Cork). The reference method we compare against is the default texture printing mode of our machine, which functionally corresponds to the direct reproduction part of our pipeline (Section 3.4), i.e., focuses only on color reproduction. To provide a comparison for our optimized results, we use unsharp masking to enhance the targets and print them in the default texture print mode. This is compatible with the treatment of Babaei et al. [2017]; we also use an equivalent method to estimate the sharpening kernel.

We draw several conclusions here. First, our direct reproduction can achieve deeper colors compared to the default texture print. Part of the reason is that the default mode applies a relatively thin layer of the colored materials (usually below 1 mm, while we use 2.5 mm). In any case, our proposed material mapping generally reproduces the target colors at least as well as the default texture print.

Second, our optimization always leads to a better match of both structure and color than any of the compared methods. In some cases, the unsharp-masking solution yields almost comparable results (especially in lower-contrast prints where the color bleeding tends to be less noticeable). It is important to note however, that the



FIGURE 3.18: *A subset of our samples illuminated by direct sunlight on a clear day, comparing the unsharp-masked default prints (**top row**) with our optimized results (**bottom row**). The light is arriving from the top-right, with the zenith angle of approximately 50°. Although our results are not optimized for such directionally non-uniform light, their structure is preserved comparatively better (cf. respective targets in Figure 3.17).*

basic property of unsharp masking (and other such enhancement methods) is to increase contrast by adding dark material around edges. This in turn causes additional undesired color bleeding, which our method avoids by design (see, e.g., the first row in Figure 3.17).

Non-uniform illumination Although our method assumes a uniform 'white-sky' illumination during the prediction, we show that the quality of the obtained results holds well outside those conditions. First, in Figure 3.18 we show several samples illuminated by direct sunlight, which is strongly directionally non-uniform. Second, to visualize their local diffusion characteristics, we illuminated a subset of our samples with a laser line in Figure 3.19. In both scenarios, our results optimized for uniform lighting (Figure 3.17, column **f**) are superior to the unsharp-masked default prints (Figure 3.17, column **e**) in terms of structure preservation, agreeing with the main results.

Monochromatic Printing For grayscale targets, our method can sometimes produce chromatic quantization noise, especially for light tones. This is mainly because our tonal mapping (Section 3.4.1) can sometimes map grayscale tones to mixtures of CMY. We can get around this by explicitly constructing the tonal mapping for just the black and white materials, as demonstrated in Figure 3.20. A more general future solution would be based on a controlled black separation, akin to methods used in 2D printing.

3.8 Discussion

Convergence The relative simplicity of the proposed optimization procedure makes it particularly easy to implement. Thanks to the efficient design of our heuristics and the high accuracy of the Monte Carlo forward prediction (Section 3.2), the optimization also converges in a modest number of iterations (cf. Section 3.5.4).

The prediction itself is currently the biggest bottleneck of the pipeline (cf. timing figures in Section 3.7), but this is mostly due to the general CPU implementation of

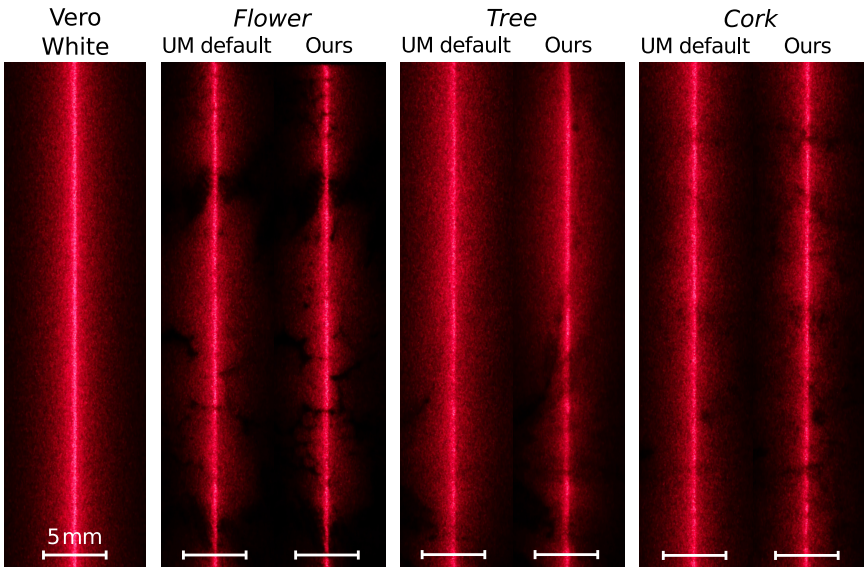


FIGURE 3.19: *A subset of our samples illuminated by a 650 nm laser line at matching locations, comparing the unsharp-masked default prints with our optimized results. For reference we also include a clear white slab in the comparison. Although our results are not optimized for this type of lighting, the lateral scattering is still inhibited to a greater degree. Please note that the graininess in the images arises from the camera's sensor noise.*

path tracing we use. Given the good parallelizability of path tracing, an optimized multi-GPU implementation could bring at least an order-of-magnitude speedup (or achieve comparable timings on a single high-end GPU). Orthogonally to that, a multi-resolution approach could converge in a comparable number of iterations, but could dramatically decrease the computation time for each of them in an amortized sense. Finally, given the high optical density of the print materials (Section 3.3), the simulated path lengths can easily reach thousands of scattering interactions. Since many of these occur in the underlying homogeneous white material, we hypothesize that applying the similarity theory [Wyman et al. 1989; Zhao et al. 2014] could also significantly decrease the simulation costs without compromising its accuracy.

General Geometry Because our solution primarily focuses on the input textures' complexity, we operate in the canonical setting of planar slabs. On the other hand, we believe a generalization of our results to full 3D geometry is perfectly plausible, and would involve two key steps: first, aligning the solution domain with the texture space of the input 3D object (including a correct handling of boundary conditions), and second, propagating the prediction gradient in an arbitrary direction in the object's voxel grid (similarly to the approach Brunton et al. [2015] chose for the 3D generalization of layered half-toning). That however still carries the assumption that the object's geometric features are at least as large as the materials' diffusion distance; if that were not the case, an additional redesign of our optimization heuristics would be necessary.

Volumetric Blockers The lateral transport could alternatively be prevented by placing opaque barriers near salient edges. We however do not have such an optically dense material at our disposal. More importantly, color bleeding influences all



FIGURE 3.20: *Comparison of our full CMYKW result and the KW (black and white) only, to the default printer output.*

features, from edges to smooth gradients, and such barriers cannot be placed everywhere as they would negatively influence the overall light transport and cause unwanted darkening.

Perceptual Considerations We considered implementing our optimization in a perceptually uniform colorspace (such as La*b*), but applying some of the optimization steps in such a space is not trivial. Orthogonally, a perceptually motivated global [Montalto et al. 2015] or, preferably, local contrast modification implemented via an adaptive tradeoff between structure and color (as hinted on in Section 3.5) could subjectively improve our results even further. In general, the perception of heterogeneous translucency is largely an open topic [Fleming et al. 2005; 2004; Gkioulekas et al. 2013a] and deserves a dedicated treatment.

3.9 Conclusion

This work presents a self-contained, end-to-end system to calibrate, model, and suitably compensate for lateral scattering in heterogeneous translucent 3D-printed volumes.

Our method involves a practical print material calibration, a hybrid material mapping scheme based on physical Monte-Carlo simulation and analytical prediction, and a carefully designed optimization loop that reliably leads to high-fidelity texture appearance.

We evaluate our system using a commercial, five-color 3D printing process where we compare favorably against the printer's native color texturing mode. Our method preserves high-frequency features well, without having to compromise on color gamut. While all of the results are demonstrated on planar surfaces, we are confident that our method can translate to more general shapes.

As a result, our system enables, for the first time, faithful reproduction of high-frequency color textures in 3D prints, despite the inherent blur of the translucent print materials.

Chapter 4

Geometry-Aware Scattering Compensation for 3D Printing

In Chapter 3 we described our first step toward developing a full-color scattering-aware 3D printing system. We introduced the problem of significant optical scattering exhibited by current printing materials, which often results in a blurry appearance of printed objects. To improve the print quality of target texture details, we formulated an inverse rendering problem and solved it for a limited case of planar slabs. Our approach involved an iterative optimization process aimed at compensating for volumetric light scattering by adjusting the material distribution within the printing volume. To ensure practical applicability to real 3D printers, we proposed and validated a practical material calibration method. As a result, we achieved a significantly better quality of texture reproduction compared to previously existing methods. However, that work was limited to printing planar objects only, which is a strong constraint in the realm of 3D printing. Furthermore, while our proposed optimization scheme yielded promising results, it lacked comprehensive verification, e.g. through a general full-volume optimization.

In this chapter (published as [Sumin et al. 2019]), we present a technique for optimizing color prints that accounts for volumetric scattering within general 3D shapes (Figure 4.1). We begin this chapter (Section 4.1) by describing a series of experiments that, for the first time, study the nature of *optimal* material distributions at the example of small test cases, using an accurate, full-volume conjugate-gradient optimization over a Monte-Carlo-simulated voxel volume. We validate that these solutions exhibit a trend consistent with a previously proposed (Chapter 3) heuristic that steers higher-absorbing material assignments toward the object surface.

Subsequently, we base our technique on that heuristic, applying it to the general 3D setting (Section 4.2) in a new iterative space traversal that relies on a custom view-independent Monte Carlo light transport simulation and accounts for arbitrarily



FIGURE 4.1: *A demonstration of 3D printouts obtained with our method. The Earth diameter is 5 cm.*

shaped objects. This ensures sharp texture reproduction – within physical limits – even around thin geometric features, where a limited material volume is available for optimization.

Our experimentation with the full-volume optimization also allowed us to derive a new geometry-aware color gamut (Section 4.3). This is a crucial component of our technique since it provides a new way of dealing with the dependency of the achievable color gamut on the local geometry in the presence of material translucency. The mapping enables reliable handling of the appearance in thin geometry areas and provides plausible color reproduction even where physics limits the achievable gamut. We compare our method to state-of-the-art techniques and demonstrate significant fidelity improvements.

4.1 Problem Exploration

The space of possible voxel color assignments within a 3D-printed object is vast, forcing any practical solution for material assignment in print preparation to make use of heuristics and approximations [Babaei et al. 2017; Brunton et al. 2015]. Before developing our own heuristic, however, we would like to learn about the general nature of *optimal* solutions. To that end, we conduct a full-volume optimization over all voxel assignments, for small test cases where finding such an optimal solution is feasible. Of particular interest are cases with strong texture gradients and competing targets on opposing surfaces. Subsequent Sections 4.2 and 4.3 then develop a robust method that builds upon the insights gained from these fully-optimized small-scale solutions.

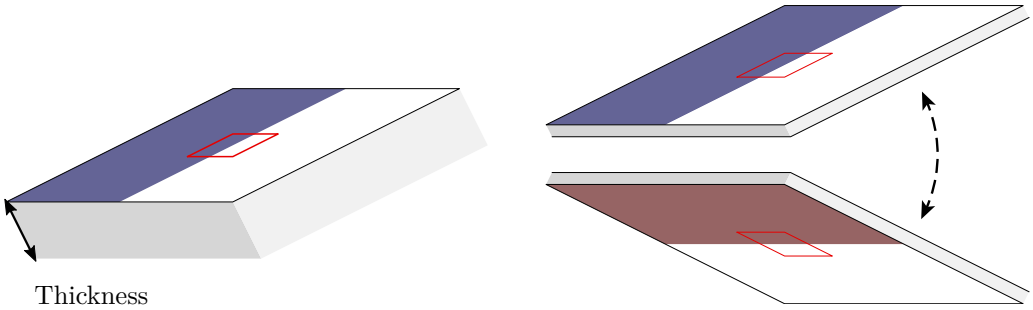
4.1.1 Full-volume Optimization

Our optimization operates on the continuous space of material concentrations at each voxel, aiming to find a (near) optimal material assignment that meets a target surface color specification for a small test geometry. Even for small cases, however, a full enumeration of possibilities is prohibitive. We hence chose a conjugate-gradient optimization to find a (potentially local) optimum. To hedge against non-convexity of the solution space, we run this optimization for various different initializations; robustness and generality are tested through different target specifications, material thicknesses, and error metrics.

Testbed At its core, the setup consists of a voxel volume, representing a slab of finite, parametric thickness and practically infinite lateral extent (i.e., big enough to avoid undesirable light transport from the sides). As depicted in Figure 4.2a, it is constructed from a texture patch centered at the region of interest and edge padding ("repeat" boundary extension) for expanding the volume laterally.

In order to keep the study general and independent of a particular set of printer materials, our testbed does not employ halftoning with its limiting discretization to obtain a heterogeneous medium. Instead, the linear RGB color at each voxel with applied Inverse Albedo Mapping (Equation (3.5)) is taken as the scattering albedo; density is set to a fixed uniform value comparable to existing printer materials.

The volume is assumed to have a smooth dielectric boundary ($\eta = 1.5$) and is illuminated using constant omnidirectional lighting. We modified a volume path tracer in Mitsuba [Jakob 2010] to compute alongside a radiance estimate also the partial derivatives of the latter w.r.t. each voxel's scattering albedo (i.e., the optimization



(A) *One-sided* (B) *Two-sided*

FIGURE 4.2: *Experimental Setup: One (a) or two (b) target textures are arranged into a centered region of interest (red square) on the front (or back) of a slab of parametric thickness. Edge-padding enlarges the volume to eliminate in-scattering of external lateral illumination.*

variable). That renderer is then used in the inner loop of a conjugate-gradient least-squares solver [Branch et al. 1999] (in the implementation of [SciPy 2008]) to find material assignments that minimize the difference between predicted and target surface appearance.

Additionally to the optimization of planar scattering, this setup can be employed in a two-sided mode, in which it is able to optimize for two surfaces on opposing sides concurrently (Figure 4.2b). By reducing the thickness of the slab, these two surfaces are less and less independent, but rather influence each other due to the light traveling through the object.

Findings In order to characterize the solution space for scattering compensation in 3D-print preparation, we ran a broad study over different input parameters, such as volume initializations (random, white, target extrusion, deep embedded confusion shape), slab thicknesses (in the range of $[0.5, 10]$ mm) and error metrics (RMS in linear RGB, CIE dE76 (CIELAB)).

The resolutions of the target textures and the actually optimized volume are set to 16×16 pixels/voxels in the tangential plane, and we experimented with 8, 16, and 32 depth layers, respectively. At that resolution, each iteration of the optimizer runs in 10 seconds, with convergence after typically 5 to 30 minutes.

From our experiments – see a collection of generated reports in the supplemental – we gathered the following findings.

1. As expected, perfect separation of front and back signal is not possible, given the materials' mean free path.

2. Voxels with high absorption occur closer to the surface. This effect is more pronounced the thinner the material slab is (see Figure 4.3, rows 2–3) but is still not able to compensate for the increased cross-talk between back and front.

3. Altering the metric for which we optimize does affect the internal distribution of the solution, but (a) the trend of higher absorption close to the surface persists, and (b) the visual differences in the corresponding outcomes are relatively small.

4. Initialization influences the solution structure, but resulting appearances exhibit only negligible visual differences.

A useful finding is that (2) qualitatively confirms a hypothesis underlying our approach in Chapter 3, which postulates that placing strong absorbers closer to the

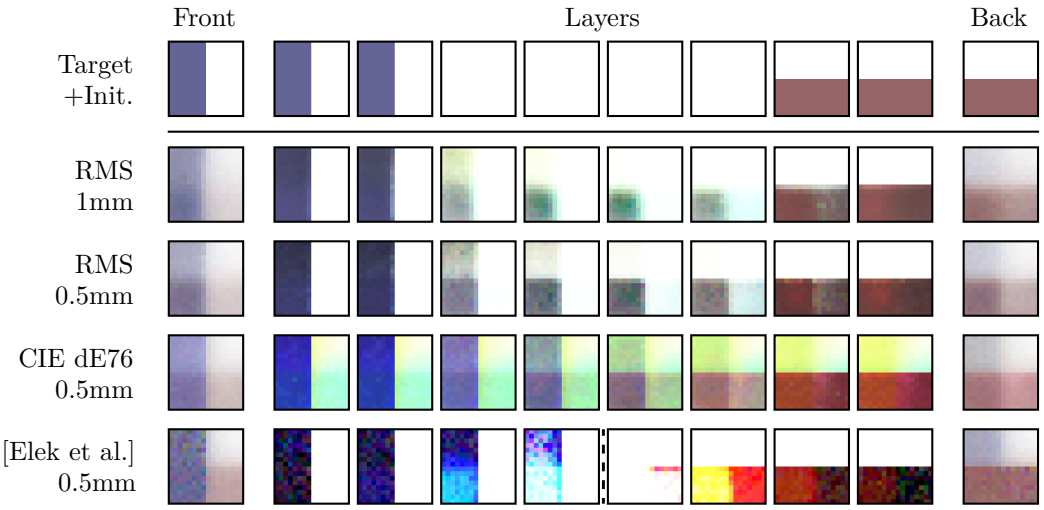


FIGURE 4.3: *Results of a conjugate-gradient optimization on volume assignments for a thin, two-sided slab. The columns show front and back predictions with the corresponding volume as horizontally scattered layers in between. Each row names the optimized difference metric and slab thickness.*

surface leads to better texture reproduction in the presence of strong subsurface scattering.

From (3) we further conclude that regardless of the error metric chosen, the outcome is dominated by the physical limitations. In the remainder, we hence use L2 in linear RGB. Nevertheless, an interesting difference remains that the use of CIE dE76, which puts more weight on color accuracy, leads to the deposition of complementary colors on the opposing side (Figure 4.3, row 4), albeit the resulting visual difference is minute. Lastly, (4) is suggestive of our results being very close to the global optimum.

4.1.2 Heuristic Optimization

Larger volumes will no longer be amenable to general optimization, due to the substantial overhead of dense partial-derivative computation. In order to scale to larger problem sizes, we require a more heuristic approach that approximates these types of solutions at a lower computational cost.

In Chapter 3 we proposed a heuristic whose volumetric assignments come close to what we observed in the full-volume optimization, specifically in how darker materials occur near the surface. In the remainder, we directly compare this method to the general optimization.

As the approach described in Chapter 3 is limited to planar, thick objects, without any effects from an opposite surface, we ran a global optimization under this exact geometric assumption. In a direct comparison (Figure 4.4), the solution of the general optimization resembles our heuristic solution. However, our heuristic solution seems to more aggressively push higher saturations toward the surface, producing slightly sharper contrast than the general optimization. This comes at the cost of a 27% higher RMS Error (General optimization: 0.120; our heuristic approach: 0.152).

Now, explore a straight-forward adaptation to two-sided cases, where we alternate two parallel iteration loops of their algorithm, each modifying voxels only half-way into the volume from opposite sides, but using Monte Carlo predictions that simulate the thin slab as a whole. As visible in the last row of Figure 4.3, method from Chapter 3 does not have a notion of opposing surfaces. It will simply ignore their

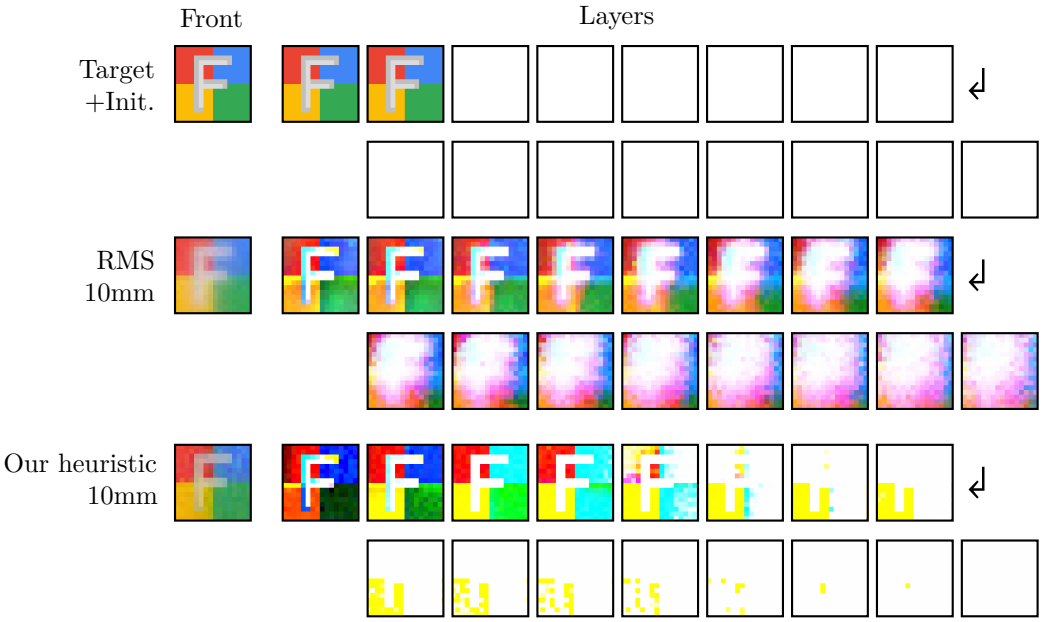


FIGURE 4.4: *Result comparison of a conjugate-gradient optimization on volume assignments with the heuristical approach (Chapter 3) for a thick, one-sided slab (10 mm of white with the top 1 mm optimized). The columns show the prediction with the corresponding volume as horizontally scattered layers next to it.*

demands resulting in sharper edges and better uniformity but at the cost of more visible cross-talk.

Overall, the qualitative agreement of the solutions led us to extend our heuristic approach to general 3D shapes, including situations with opposing surfaces near thin features.

4.2 General Geometry Scattering Compensation

In this section, we build a full general-geometry pipeline that can compensate for sub-surface scattering in 3D prints. Figure 4.5 gives and overview of the whole process, from the input appearance to the final printout. First, we convert ① the target shape with a diffuse texture to a voxel representation as described in Section 4.2.1. Our initial solution is created in step ②. We iterate over candidate solutions using our refinement loop, consisting of halftoning ③ (Section 4.2.3), prediction ④ (Section 4.2.4) and volume refinement ⑤ (Section 4.2.5). Finally, the halftoned volume is ready for fabrication ⑥.

4.2.1 Data Preparation

Using OpenVDB [Museth 2013], we convert the input 3D mesh into a distance transform on an isotropic voxel grid (step ①). As the resolution along the printer's three axes usually differs (600×300×940 DPI in our case), we choose to work with the coarsest of the three. Later, during halftoning (Section 4.2.3), we up-sample to the printer's native resolution. The coarser solution grid allows for significantly faster processing; conversely, optimization at the full resolution would have little benefit, as the need for halftoning reduces the effectively available printer resolution anyway.

We classify voxels as *interior*, *surface*, or *exterior*. Surface voxels encode the corresponding surface normal and gamut-mapped, target RGB color. We perform the

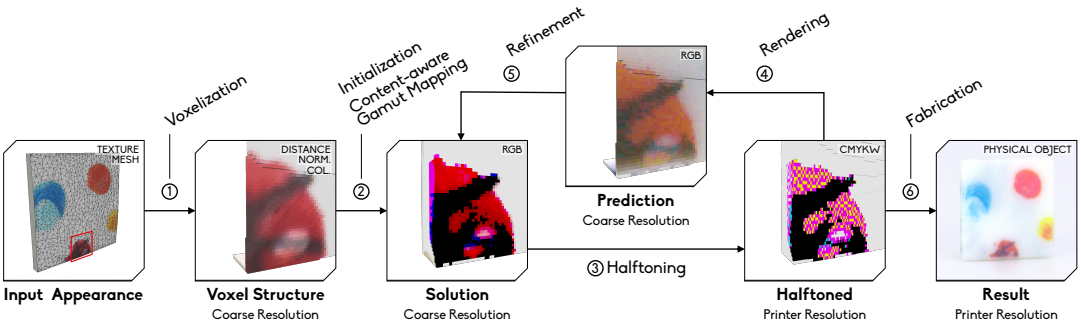


FIGURE 4.5: *Our pipeline takes a target in the form of a textured triangle mesh. It voxelizes it (1) at a uniform resolution with additional data stored on the surface. The inner loop iteratively optimizes an RGB volume by predicting the appearance (4) and propagating the residual error into the volume (5). Finally, we output discrete material assignments readable by the 3D printer (6).*

gamut mapping in CIELAB color space using a state-of-the-art algorithm [Rossier 2013]. Note that this is a conventional gamut mapping that ensures uniform RGB colors are within the printer gamut (as described in Chapter 3), different from our content-aware gamut mapping introduced in Section 4.3. The solution for the volumetric material arrangement is stored as linear RGB values representing the desired absorption at each non-exterior voxel of the coarse grid.

4.2.2 Interior-to-Surface Points Mapping

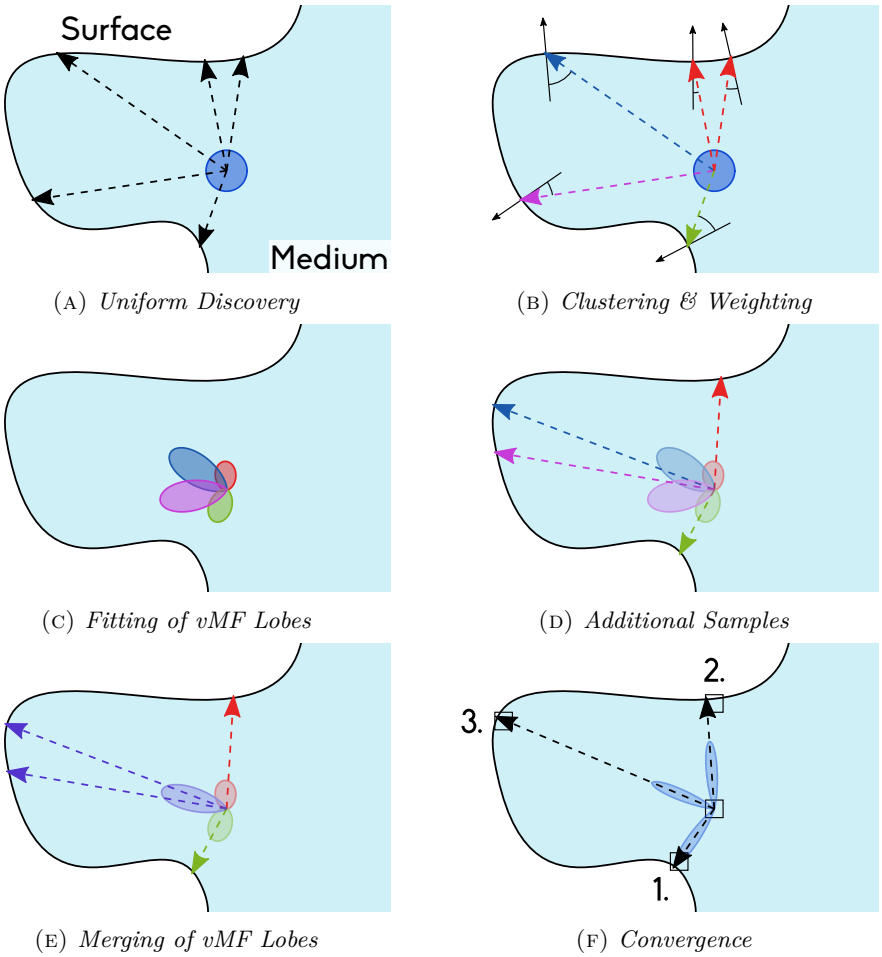
In planar scattering compensation (Chapter 3), each voxel layer corresponds to a distinct depth within the material, while columns group all voxels underneath a given surface point. For generally shaped 3D prints, there is no clear definition of voxel layers and columns. We propose three distinct changes to account for surface curvature and geometric thickness. First, we replace the notion of layers by partitioning the voxel object into *iso-shells* of equal (after rounding) distance to the surface. Second, lacking a precise definition of voxel columns in curved regions, we store for each non-exterior voxel its nearest surface voxel to associate voxels with a surface point.

While this mapping is sufficient in thick geometric cases, it becomes misleading for thin geometries where a considerable amount of light enters through adjacent surfaces (Section 4.1.1, Finding 1). Voxels thus affect the appearance of several surfaces, not only the closest one, which necessitates an appropriate representation for the refinement step (Section 4.2.5).

Accordingly, our third adjustment is a conceptual augmentation of the voxel-to-surface correspondence function (implicit in Chapter 3) to multiple connections.

Motivated by the energy distribution of a subsurface scattering kernel, the connections are concentrated around the surface normal direction and sorted by distance to model the major influence.

In practice, we compute additional surface connections for each inner voxel in thin geometries as shown in Figure 4.6. Through uniformly shooting rays over the sphere against the input geometry, connections are initially discovered (a). Discretized to voxel indices, the candidates are weighted according to their deviation from the surface's normal direction. We form clusters of direction vectors (b) and improve each cluster by shooting additional samples drawn from a von Mises-Fisher distribution fitted against the current best $\tau = 8$ candidates per cluster (c-d). By iteratively re-fitting the distributions, cluster directions can change and potentially merge (e). Typically after 8–10 iterations of 64 samples each, the algorithm converges to a list



(A) *Uniform Discovery* (B) *Clustering & Weighting*

(C) *Fitting of vMF Lobes* (D) *Additional Samples*

(E) *Merging of vMF Lobes* (F) *Convergence*

FIGURE 4.6: *Interior-to-Surface Mapping: Connections are initially discovered using uniform sampling (a) and further refined using fitted von Mises-Fisher (vMF) distributions (b-e) until convergence in surface normal direction (f). The result in (f) is sorted by distance and references discrete surface voxels.*

of surface connections that are aligned with the normal direction (f). Eventually, this list (sorted by distance and capped to $\lambda = 8$ items) contains only surface voxels whose directions are sufficiently separated in order to circumvent duplicate connections to the same surface.

4.2.3 Halftoning

Predicting the appearance of the solution at each refinement stage (and ultimately fabricating the model) requires both conversion into the operational grid of the printer and discretization into the material labels (CMYKW) as our target 3D printer can only place a single material at each voxel. We perform this as a single operation ③. For each point at the printer grid (higher resolution), we query the RGB values from the solution represented on the coarse grid using tri-linear interpolation. The RGB value is then converted to CMYKW tonal mixture using a static mapping model. Finally, we perform a per-slice error-diffusion halftoning using a serpentine-order with the Floyd-Steinberg kernel [Floyd et al. 1976]. Although more advanced 3D halftoning techniques are available [Brunton et al. 2015], for our examples, the difference was not visually significant.

4.2.4 Appearance Prediction

The material arrangement refinement ⑤ of our method relies on accurate prediction ④ of the surface appearance implied by the current solution. We employ Monte Carlo simulation for this purpose, using the same scattering parameters of the print materials as derived in Chapter 3. While time-consuming, this technique remains – to our knowledge – the only method capable of accurate predictions in the presence of complex geometric features (sharp edges, thin slabs, etc.), while also supporting heterogeneous media.

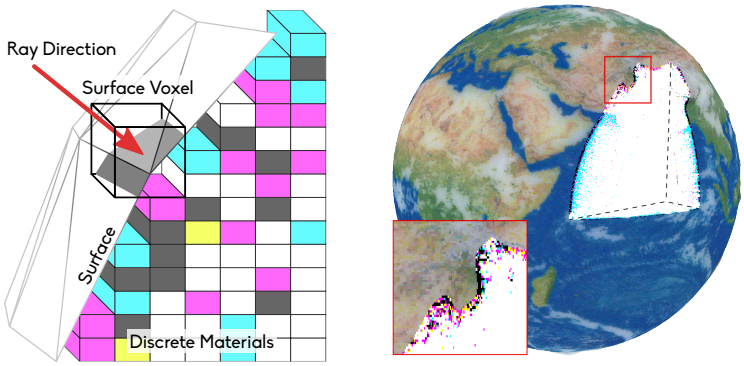
The volume grid input to the prediction stage consists of one of the five measured material parameters per voxel at the printer's native resolution. The lighting simulation sees the grid as block-shaped voxels, assuming a hard boundary between neighboring voxels, though in practice material mixing and mechanical dot gain may occur at voxel boundaries.

Light transport model and lighting The (virtual) lighting setup greatly influences the object appearance, as the print materials show a high degree of translucency. Optimizing the printout for specific lighting is feasible, however, it would reduce its generality. For this reason, we assume spatially and directionally constant, white illumination, factoring out all possible influences of the environment (such as self-shadowing or inter-surface reflections) from the optimization procedure. In fact, the prediction step is designed to only consider the subsurface part of light transport, which is the primary effect our method addresses. Specifically, our volumetric path tracer (implemented in Mitsuba [Jakob 2010]) stops tracing a path once it exits the medium through the surface into the air, and each such path brings unit radiance modulated by the Fresnel transmittance for a smooth dielectric boundary ($\eta = 1.5$).

Sensor Our goal is to match the appearance over the entire object surface, so predictions obtained with standard perspective or orthographic cameras would be insufficient. Performing the prediction directly in the UV texture space would require special consideration of texture distortion and discontinuities. Instead, we take advantage of our volumetric representation with the surface normal assigned to each surface voxel. This data is input to a custom sensor in Mitsuba, which renders for each 'pixel' the outgoing radiance of one surface voxel. As depicted in Figure 4.7, the 'viewing' direction for each voxel is aligned with its surface normal so that the transmission at the dielectric boundary is the same for all voxels.

While the light transport simulation runs on the printer-resolution grid, the predicted outgoing radiance is recorded at a set of surface samples, corresponding to the centers of the surface voxels of the coarser grid. This makes the prediction output ready for the subsequent stages of our pipeline, which work at the coarser resolution.

Preview setup For verification and demonstration purposes, the halftoned solution can be rendered with a conventional perspective sensor and omnidirectional illumination. An image obtained in such setup is compared in Figure 4.8 to the printout showing the quality of our prediction system. Figure 4.12 additionally offers a comparison for thin geometry (planar 0.5 mm slab). The *preview mode* is used in further figures and is explicitly marked with the term *rendered*.



(A) *Concept* (B) *Visualization*

FIGURE 4.7: *Rendering setup: For each surface voxel (in the solution resolution), we render point samples along the surface's normal direction (red arrow in (a)) and trace rays into the discretely heterogeneous medium (printer resolution). The visualization in (b) shows a typical surface prediction (stored in voxels) with the underlying halftoned solution as a cubic cross-section.*



(A) *Prediction* (B) *Printout*

FIGURE 4.8: *Prediction accuracy: (a) shows a* preview *prediction (perspective sensor, spherical illumination) of the Earth model. (b) is a photograph of the corresponding printout.*

4.2.5 Material Arrangement Refinement

Each iteration of our optimization loop (Figure 4.5) consists of rendering ④ the current (halftoned) solution and a refinement step ⑤, in which we propagate residual error into the volume. We initialize the solution by extruding the surface color into an initially fully white object up to a depth of $d = 0.1\,\text{mm}$, by assigning the target color of its closest surface point to each interior voxel in the depth range $[0, d]$. The thickness parameter d influences the overall convergence speed, as it allows our method to essentially skip a few early iterations if it is chosen well: but ultimately, it does not significantly affect the sharpness of the result.

Update step The rendering step provides an appearance prediction per surface point, from which a residual error compared to the target color is calculated. Positive residual values in a color channel imply that the current solution is too dark and needs lightening, while negative values indicate that darkening is required.

For each surface point, the refinement step distributes this residual into the set of voxels that correspond to this surface point (Section 4.2.2). Most of it is deposited into voxels which reference it as the *closest* surface point, while leftover lightening is

additionally also propagated to voxels which 'belong' to a different side of the object (i.e. that only hold a higher-order reference to this surface point).

We traverse all iso-distance layers, up to a maximum depth beyond which no noticeable light transport can occur – which in the case of our printer materials is 3 mm (corresponding to 18, 27, 72 mean-free-path lengths for the R, G, B channels of the white material respectively). In the refinement procedure (Algorithm 2), we first traverse the volume from the deepest iso-layers up to the surface and disperse any *positive* residual components, *lightening* the RGB values within the grid. Next, and similar to the planar compensation approach (Chapter 3), a *darkening* pass processes the volume from the surface iso-layers to the deepest ones, to disperse any *negative* components. Last, we again run multiple sparse lightening passes; this time depositing leftover positive residual only to voxels with an additional reference (i.e. *not* closest) to a surface point.

These references (derived in Section 4.2.2) are only available for voxels in thin regions and are sorted by ascending distance which in turn means by descending influence.

The last lightening pass enables 'negotiation' in cases when 'demands' of the two sides of a fine geometry are in conflict, i.e. one side should be darkened (in one of the R, G, B channels) while the other side should be lightened. The approach ensures that absorbing voxels are placed only near the darker surface and, at the same time, allows for processing the volume as a whole as opposed to processing two halves of the volume separately.

Recovery from overshooting We limit the fraction of residual dispersed in each update step, and propagate darkening and lightening only if the current iso-layer voxel can no longer receive any more darkening or lightening.

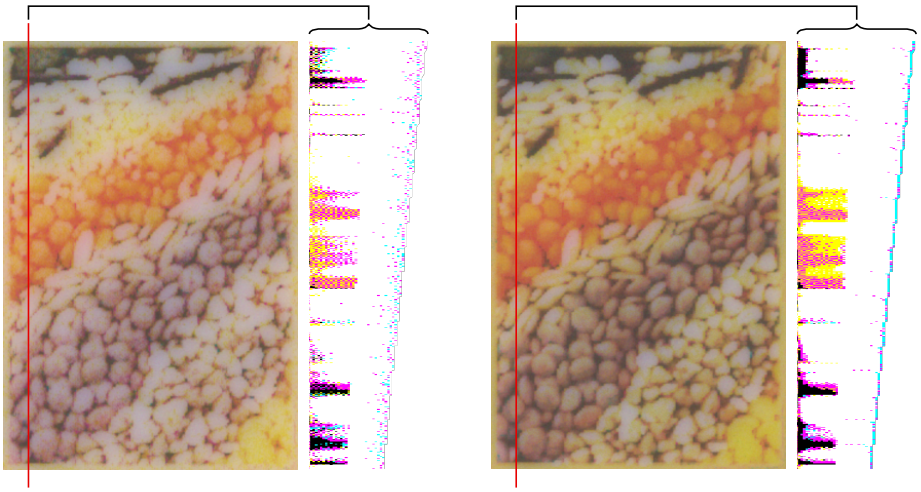
Even so, situations may still occur where the update leads to excessive darkening for individual surface points (for example due to a concurrent darkening of an opposing surface). In order to recover in these cases, in Chapter 3 we propose to uniformly lighten all voxels within the corresponding column whenever the residual is positive. However, such method of recovery disturbs the goal to minimize the number of absorbing voxels by placing them very close to the surface.

This problem is most severe when the target texture contains fine, dark details surrounded by larger areas of light material. Here, the dark surface voxels are back-lit by comparatively large amounts of in-scattered light. Any change of their concentration, such as darkening by changing the local absorption by a fixed amount, is amplified by the higher "back lighting" (absorption is multiplicative), increasing the risk of overshooting. In these cases, our previous uniform lightening implicitly draws the bulk of the darkening deeper into the volume, increasing dark color bleeding into the light surrounding and decreasing the globally achievable contrast. This effect is demonstrated in Figure 4.9, where, comparing the *rendered* predictions, a decrease in contrast is visible and the loss of concentration at the surface is revealed in the slice through the halftoned solution.

In contrast, in this work we rely on the new inside-out lightening pass described above to ensure that the darker voxels stay near the surface, mitigating potential blur. This *compressive* behavior is only stopped in the very proximity of the surface (0.25 mm in our system): shades of colors are better reproduced with 2 or 3 iso-layers than with a single iso-layer due to the small number of base materials (and, consequently, color combinations) in the subsequent Halftoning step ③. We observe that the new update step is effective in eliminating overshoot artifacts, to the extent that we are

Algorithm 2: Refinement (update step) ⑤

Input:
$\boldsymbol{X}$: current solution volume
$\boldsymbol{T}$: target-specification surface voxels
$\boldsymbol{P}$: current appearance-prediction surface voxels
Result:
$\boldsymbol{X}$: updated solution volume
begin
  // Calculate positive and negative differences from target
  // $\boldsymbol{D}[voxel, c]$ returns the c-th order voxel-to-surface correspondence
  $\boldsymbol{D} = \boldsymbol{T} - \boldsymbol{P}$; $\boldsymbol{D}^+ = \max(\boldsymbol{D}, 0)$; $\boldsymbol{D}^- = \min(\boldsymbol{D}, 0)$
  // Lightening and Darkening passes
  LIGHTEN($\boldsymbol{D}^+, 0$)
  DARKEN($\boldsymbol{D}^-$)
  // Pass over secondary voxel-to-surface correspondences to disperse the residual lightening to the rest of the volume; we use $\lambda = 8$
  for *correspondence* $= 1..\lambda$ **do**
    LIGHTEN($\boldsymbol{D}^+, correspondence$)
  end
end
procedure `Lighten(`$\boldsymbol{D}^+$`, correspondence):`
  // Loop over iso-layers $\boldsymbol{L}_z$, from the most distant one Z to the surface.
  for $z = Z..0, \boldsymbol{L}_z \in \boldsymbol{X}$ **do**
    // Loop over voxels belonging to the iso-layer.
    forall $v \in \boldsymbol{L}_z$ **do**
      v += $\boldsymbol{D}^+[v, correspondence]$
      // Update of the positive-residual for a given surface point ensures (unless close to the surface) the compression
      if $v.distance \geq 0.25mm$ **then**
        $\boldsymbol{D}^+[v, correspondence] = \text{clamp}(v - 1, 0, 1)$
      end
      $v = \text{clamp}(v, 0, 1)$
    end
  end
end
procedure `Darken(`$\boldsymbol{D}^-$`):`
  // Loop over iso-layers $\boldsymbol{L}_z$, from the surface to the most distant one Z.
  for $z = 0..Z, \boldsymbol{L}_z \in \boldsymbol{X}$ **do**
    forall $v \in \boldsymbol{L}_z$ **do**
      // Conservatively darken the voxel proportionally to c_p.
      // $c_\text{a} = 4.0$, $c_\text{p} = 0.5$ as in Chapter 3
      v += $c_\text{p} \cdot \boldsymbol{D}^-[v, 0]$
      $\boldsymbol{D}^-[v, 0] = \text{clamp}(c_\text{a} \cdot v, -2, 0)$
      $v = \text{clamp}(v, 0, 1)$
    end
  end
end



(A) *Update step in Chapter 3* (B) *Improved update step*

FIGURE 4.9: *Comparison of optimization strategies: The textured side (rendered) of planar slabs with varying thickness (3–4.3 mm) and uniform blue target color on the back side. Cross-sections show the arrangement of printing materials in the volume. Our strategy recovers well from over-darkening due to cross-talk and preserves absorption close to the surface.*

able to use a 3-times more aggressive value for darkening factor c_a, which results in faster convergence.

4.3 Content-aware Gamut Mapping

Smaller models often possess thin geometric features where one cannot assume enough volume for the printing materials to reproduce a range of colors. The colors on thin features are most severely affected by the colors on the opposite surface due to backlighting. This is a fundamental limitation of current color 3D printing technology and no solution comes without compromises. Our main insight for treating these cases is to adjust the color of input models to the capacity of geometry they rest on. This will reduce the risk of artifacts appearing due to insufficient volume and conflicting demands during our iterative refinement stage (Section 4.2.5). Note that the introduced content-aware gamut mapping is different from the conventional gamut mapping (Section 4.2.1) performed for all input colors. The latter is independent of the surrounding colors.

Babaei et al. [2017] pointed this problem out and offered a procedure in which they perform an adaptive gamut mapping where each color on the surface is mapped into the gamut afforded by the object thickness at that point. This approach, however, is limited, as it does not consider the actual colors in the vicinity. For a thin wall with nearly identical colors on both sides, it would drastically underestimate the practically achievable gamut, because it only considers half of the depth for each side. If colors on the two sides of a thin wall are contradicting, they would likely suppress each other.

We instead take into account both geometry and color *content* when remapping input colors in our preprocessing (step ② of our pipeline, see Figure 4.5). We take advantage of our full-volume conjugate-gradient optimization setup (Section 4.1.1) to investigate the range of achievable color pairs for different object thicknesses. Then, we use this data to process any pairs of conflicting colors and bring them into the gamut of realizable color combinations.



(A) *RGB Sampling* (B) *Example Pair*

FIGURE 4.10: *(a) The palette of target colors used to build the* convex hull *of content-aware gamuts. (b) One sample assignment for a two-sided wall (Figure 4.2b, thickness:* 0.5 *mm) with a target color on either side. Using conjugate-gradient optimization, we transform the target pair to an achievable pair for each wall thickness.*

4.3.1 Gamut of Achievable Color Pairs

In this section, we describe a one-time precomputation of the content-aware gamut. It consists of two steps: (1) collection of achievable color pairs for different thicknesses to understand the boundaries of the gamut in the color space of our prediction system (RGB), and (2) building a dense set of realizable color combinations (again, for different thicknesses) in a perceptual color space for performing the actual gamut mapping.

Collection of achievable RGB combinations Using the two-sided setup (Figure 4.2b), we run the conjugate-gradient optimization for thin walls of 10 different thicknesses from 0.25 to 3 mm. On the two sides, we place all combinations of uniform RGB colors where each channel is set to one of three values (0.0, 0.5, 1.0); this gives us 27 colors (Figure 4.10a) and 351 color combinations. The final result of each optimization is a pair of colors which are both nearest to the target pair and realizable at a specific thickness (Figure 4.10b). We also include the intermediate results obtained from each optimization iteration (around 20) as each rendering is an achievable pair on its own. Including these intermediate pairs helps to obtain a larger set of achievable color pairs (around 4000 for each thickness).

We use a per-channel *mean-difference* representation to process the calculated data: each pair of colors on the two sides – RGB and $R'G'B'$ – is transformed into:

$$\left(\frac{(R+R')}{2}, |R-R'|\right); \left(\frac{(G+G')}{2}, |G-G'|\right); \left(\frac{(B+B')}{2}, |B-B'|\right)$$

This representation captures the following features. First, the cross-talk can be estimated in each channel (R, G, B) separately; we did not find any inter-channel dependencies. Second, the achievable *difference* values are dependent on the *mean* values; e.g., the brighter the colors, the smaller the achievable difference. Third, the achievable differences are symmetrical with regard to the two sides.

Building the gamut In order to build a dense, discrete (content-aware) gamut we begin with computing the gamut boundary by finding convex hulls of 2D projections of the previously simulated color pairs. Figure 4.11 demonstrates the convex hulls of the collected data in the *mean-difference* representation for different wall thicknesses. Then, we tightly sample the CIELCh color space [Wyszecki et al. 1982] (collecting 163 million pairs) and verify if each *pair* of LCh colors is achievable at each of 10 thickness values. Our verification consists of checking if a given LCh pair is inside

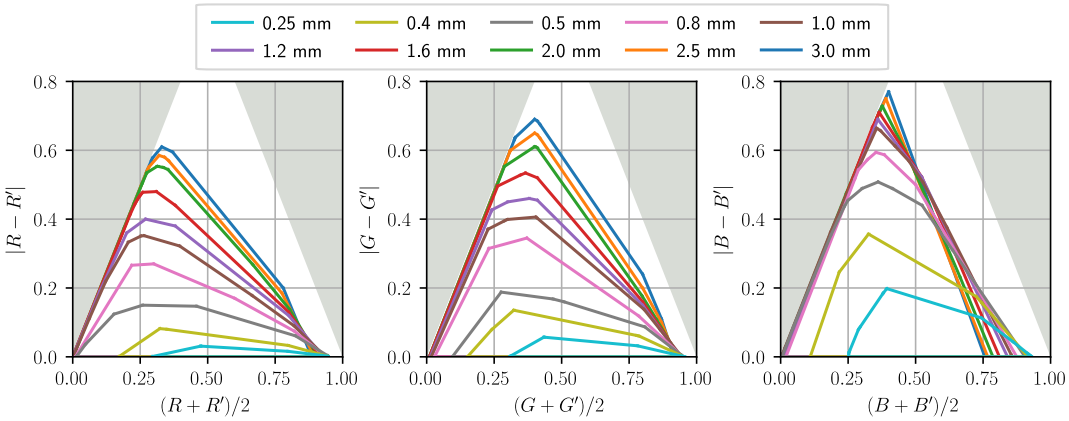


FIGURE 4.11: *Convex hulls of the achievable color pairs in the* mean-difference *representation for each channel and thickness of the object. Each convex hull represents around 4000 data points.*

the boundary of the set of achievable colors obtained in the previous step. Given a thickness, for each LCh pair, we compute its equivalent RGB pair and check if all 2D projections of this pair (RGB on one side, $R'G'B'$ on another) are within the convex hulls. If the test is valid for all projections (R/R', G/G', B/B'), we mark the color pair in-gamut.

4.3.2 Mapping

Upon having the *content-aware* gamut of achievable color pairs for a set of thicknesses, we – once for an object (Figure 4.5, step ②) – process the input colors depending on the nearby geometry and the close colors on the opposite side. This step ensures that conflicting colors on different sides of thin geometric features are reproducible. It is, in fact, an approximation to the solution which a full-volume conjugate-gradient optimization would produce if run on thin features, except the expensive optimization is replaced with gamut precomputation.

We experimentally observe that the backlit cross-talk is present when the object is thinner than 3 mm. Our content-aware gamut mapping, therefore, acts on parts of the model with thickness below 3 mm.

Before assigning conflicting pairs, we smooth the surface colors using a Gaussian filter. This way each color receives collective feedback from the other side, as we relate each color on one side to only a single conflicting color on the other side of a thin feature.

We iterate over all voxels and collect pairs of conflicting surface voxels that are (1) on opposite sides (see Section 4.2.2) and (2) closest to each other. The opposite-side condition is checked by analyzing the mutual arrangement of the two surface voxels and the connecting inner voxel. For example, in the volume demonstrated in Figure 4.6f, surface points 1. and 2. are marked as *conflicting* if the distance between them is less than 3 mm.

If any pair of conflicting colors is outside the gamut of achievable color pairs (Section 4.3.1), it will undergo an adjustment. This includes a simple lookup in the densely-sampled space of achievable LCh color pairs and finding the closest point using the CIEDE94 distances. The result of this preprocessing is a set of modified colors at thin features which are fed to the subsequent stages of our pipeline (Figure 4.5).

Content-aware gamut mapping limitations A major limitation of our method is that for each voxel we assign only a single conflicting color, lying on the opposite side. Although our smoothing step, indirectly, causes more than one voxels to influence the opposite side, one could argue that some neighboring voxels on the same side could have an effect as important as the opposite ones. In principle, one could incorporate more conflicts but should be careful that our method does not scale well with more colors as the number of projections grows unwieldy. Apart from this limitation, our nearest neighbor search is essentially a *gamut clipping* which is a less desirable category of gamut mapping techniques [Morovic et al. 2001]. Gamut clipping preserves the saturation but may lose image details. The more acceptable *gamut compression* category is however computationally intense in our 6D setup.

Content-aware gamut mapping results Figure 4.12 illustrates the effect of content-aware gamut mapping on thin planar slabs. We place two different textures (a) on the two sides of a slab. The bottom row shows the results we obtain using our normal procedure (Section 4.2), where the target images have gone only under a conventional gamut mapping. The results of both sides of the unprocessed targets exhibit artifacts at the thickness of 0.5 mm, along with hue shifts and loss of details. See, for example, how the window (marked with arrows) is visible on the opposite side of the printout or the objectionable green tint on the pillow in van Gogh's *Bedroom in Arles*. At the thickness of 1.0 mm, we see the gradual decrease of the cross-talk, and at 2.5 mm it is almost not visible. The middle row shows the results when using the gamut-mapped target specifications (b). As we can see, they not only achieve a much more consistent match when compared against their own target textures but also show good consistency with the original textures. Additionally, the images on



FIGURE 4.12: *Result of the content-aware gamut mapping demonstrated on planar slabs of different thicknesses (0.5, 1.0, 2.5 mm) with two images as target (a). (b) demonstrates the gamut-mapped target specification. The images on the gray background (c) are photographs of printouts as opposed to the rendered solutions in (d). Note the reduced cross-talk (i.e. the window under the arrow) and the reproduction of details and colors for the thinner cases.*



(A) *Target* (B) *No Content-aware Gamut Mapping* (C) *Content-aware Gamut Mapping*

FIGURE 4.13: *The effect of our content-aware gamut mapping demonstrated on an arbitrarily shaped object (rendered). The thickness of the geometry is* ≈1 mm. *Our proposed technique reduces cross-talk and improves the spatial consistency of the object.*

the gray background (c) demonstrate the printed 0.5-mm slabs. They are consistent with the renderings (d) and also demonstrate the achieved improvement.

Figure 4.13 demonstrates the effect of content-aware gamut mapping on a generally-shaped model. This object has a height of 3 cm and features approximately 1-mm thick 'walls'. The cross-talk is significantly reduced when using the content-aware gamut mapping. The proposed solution is more consistent with the target specification: the white parts of the model preserve the neutral color, the change of color on the opposite side of geometry is not visible.

4.4 Results

We fabricated the presented results using a Stratasys J750 PolyJet 3D printer and the Vero Rigid Opaque materials family. The printouts were cleaned using a grade-1000 sandpaper and sprayed with a transparent lacquer. We avoid scratching the absorbing voxels by adding a ≈0.12 mm shell of transparent material around the object. The pictures of the printouts were taken using a Canon EOS 700D camera with EF-S 18–135 lens set to 135 mm. We used two 55W, 5500 K fluorescent lamps and X-Rite ColorChecker Passport to set the white balance and the luminance level.

The refinement and appearance prediction were performed on isotropic voxel grids with resolution 300 DPI. The Monte Carlo simulation for appearance prediction was performed on a CPU cluster consisting of 100 quad-core Intel Xeon E5620 CPUs. It takes around 5 minutes to predict the appearance for a 4 cm object (i.e. red vase from Figure 4.14).

Target reproduction evaluation The main results are presented in Figure 4.14. For comparison, we fabricate the exact same models using our pipeline and two additional solutions: the built-in printer software (GrabCAD) and Cuttlefish, the commercially available software package that, to our knowledge, represents the approach described in [Brunton et al. 2015]. Our results are presented in the rightmost column.

In general, results produced using our method show significantly higher reproduction fidelity in terms of both color and spatial details. This is expected as the two other methods are agnostic to the mutual volumetric light transport between nearby surface points. The Cuttlefish printouts seem to be slightly color shifted, which may be explained by inter-machine variability in calibration data. Our thin results, especially the last row, show dramatic improvement when compared to the



FIGURE 4.14: *Comparison of photographed printouts to the rendered (a) target models: (b) GrabCAD is the default printer software, (c) Cuttlefish is based on [Brunton et al. 2015], and (d) Our scattering- and crosstalk-compensated solution. The last model is a thin planar slab of* 0.5 *mm thickness with two different textures on its front and back sides.*

(a) GrabCAD (b) Cuttlefish (c) Our method

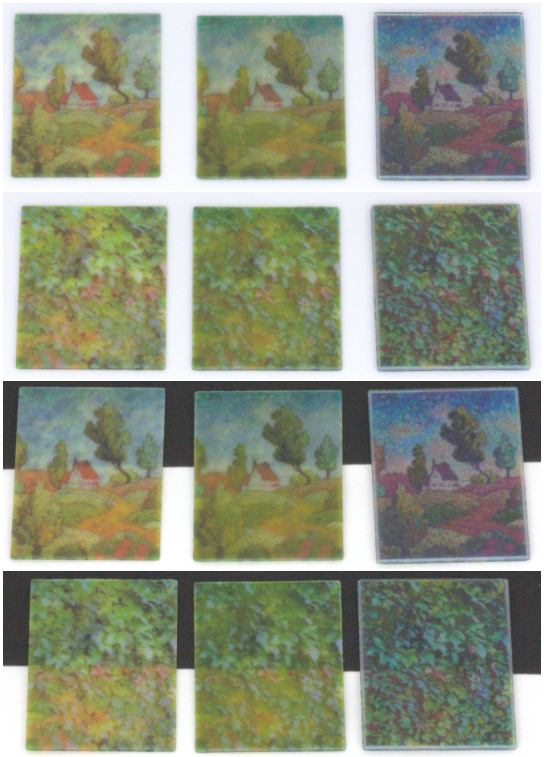


FIGURE 4.15: *Photos of the thin-slab objects (the 5th row of Figure 4.14) lying on a white table (rows 1–2) and on a black-white edge (rows 3–4).*

other methods: they are closer to the source image color-wise and suffer from much less backlit crosstalk.

One peculiar case is the color reproduction of the cheetah model. None of the methods is able to reproduce the orange segments of the skin pattern. We investigated this case and found out that both black and orange segments are reproducible when printed uniformly, isolated from each other. We then run a full-volume optimization including these two colors together and observe that even the full optimization is not able to reproduce them simultaneously. Although we are aware that this may be an extreme case (strong absorbing material next to a material with a large mean free path at high spatial resolution) where both our refinement and full-volume optimization fail to find a valid arrangement, it can plausibly show us the physical limits of this set of materials.

Figure 4.15 presents additional photos of the thin slabs printed with different methods (row 5 in the main results). Here, the prints are lying on a white table (rows 1–2) and on a black-white edge (row 3–4) in contrast to a vertical standing in the main-results figure. The absence of back-light changes the perceived appearance; the black-white edge demonstrates the extent of translucency we are dealing with.

4.4.1 Baseline Comparison to Remapped Planar Slabs

Lastly, we would like to validate whether a simpler extension of the pipeline described in Chapter 3 would have been possible. Intuitively, their method could also be generalized by optimizing a texture as-is in 2.5D and treating the resulting planar volume as a 3D texture that fills a 3D mesh. Such remapping would take into account the UV parametrization on the surface and the distance into the volume to select the correct

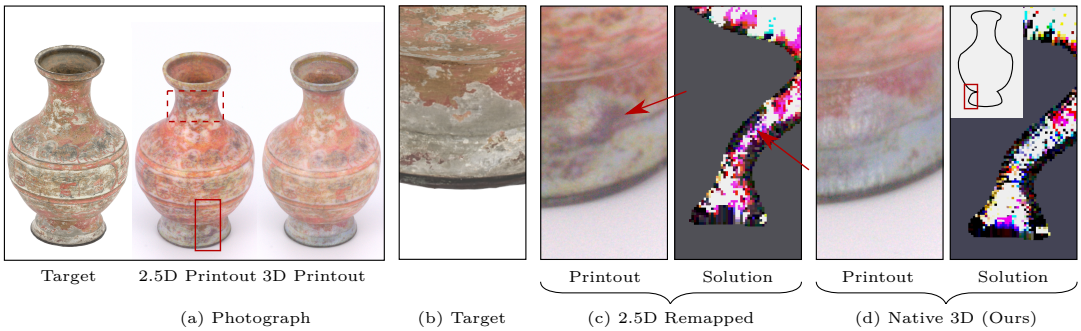


FIGURE 4.16: *Baseline Comparison (a) of results achieved with planar scattering compensation (remapped to 3D using UV + depth coordinates) and our native 3D pipeline without our refinement improvements (the new update step and the content-aware gamut mapping). (c) and (d) show cropouts of a region highlighting problems of the 2.5D method and cross-sections thereof in comparison to the target appearance (b).*

texel on the planar surface and layer from the stack respectively. A comparison of this idea to our native 3D treatment of the problem is described in the following.

We conduct this baseline experiment to evaluate the texture sharpness and color reproduction on 3D surfaces achievable by both methods. The remapping approach has disadvantages when implemented naïvely and requires special care to avoid pitfalls such as texture stretch or a mismatch in light scattering distance before and after the mapping. Please refer to the supplemental for a detailed description of our implementation.

Figure 4.16 compares printouts obtained by a remapped 2.5D solution to an inherent 3D solution. For a fair comparison, the native 3D version does not include our improved refinement step (Section 4.2.5) but employs the same heuristics as the 2.5D method. A comparison with all our features enabled and for all models in Figure 4.14 is provided in the supplemental.

We observe an overly saturated appearance across the surface of the 2.5D vase (Figure 4.16a). On thin geometries, such as the hollow neck (marked dashed) or foot (cropout), the surface exhibits unwanted blackening. The observations are in line with CIEDE2000 average (calculated on renderings) which is 15 for the 2.5D variant and 10 for the 3D one.

Conceptually, the columns in the 2.5D solution get expanded on concave geometry, whereas convex surfaces shrink the columns in depth.

In convex regions, the 2.5D mapping compresses the solution (especially in deeper iso-layers), reducing the fraction of white material in the volume. Since the light absorption is unrecoverable, the resulting appearance is darker and more saturated leading to the false impression of global contrast enhancement. In concave regions, the 2.5D mapping increases the absolute number of absorbing voxels (column expansion) with similar effect. Figure 4.17 shows how these two effects can even lead to severe color shifts on the cheetah's abdominal.

The third consequence of the remapping is a cutoff in depth due to a smaller thickness of the target geometry. Each side of a wall is by definition only assigned its half of the volume, leading to a cutoff in depth if the overall thickness is below the sum of filled depth on both sides. The planar solution consists of mostly black voxels near the surface, which remain after the remapping, while less absorptive layers underneath are discarded.

See examples of this blackening effect in Figure 4.16c (red arrow) and on the paws in Figure 4.17.



FIGURE 4.17: *Color shifts on the abdomen and blackening on the paws of a 2.5D remapped printout (left) compared to a native 3D printout (right). Similar setup as in Figure 4.16.*

This clipping cannot be solved by simply rescaling the composition in one column to the available depth, as this would induce unwanted brightness and even color shifts.

Our approach (described in Section 4.2.1) avoids the uncontrolled behavior of the 2.5D method, namely, inaccurate color reproduction and introduced artifacts by operating natively in the final geometric setup and by handling opposing surfaces seamlessly in the refinement step (Section 4.2.5).

4.5 Conclusion

We presented a technique that handles the inherent limitations of significantly translucent base materials and allows one to compensate lateral scattering in inkjet 3D prints when reproducing textured objects. In contrast to previous approaches, our technique is a genuine 3D approach and is capable of enhancing input textures on curved objects with arbitrary geometry. In the pathological case of very thin geometric features, we propose a preprocessing step that significantly improves consistency and accuracy.

Chapter 5

Neural Acceleration of Scattering-Aware Color 3D Printing

In the previous two chapters, we constructed a general print preparation system capable of producing high-quality full-color objects on multi-material inkjet 3D printers [Mimaki 2017; Stratasys 2017]. To address the main inherent limitation of UV-curable print materials, we developed an inverse rendering method that iteratively improves an arrangement of printing materials to be fabricated by a printer. During a forward pass, a prediction of the current solution's appearance is computed using a full MC simulation of light transport, which can take many hours for a small object on a high-end desktop system. In a backward pass, based on the difference between the target and predicted appearance of the current solution, a scattering-aware refinement procedure adjusts this solution. Although the proposed heuristic, which guided the solution adjustment, provided fast convergence of the method (in terms of the number of iterations), the total time required for preparing a single object was immense and utterly impractical when using a typical desktop computer. Alternatively, one could utilize a CPU cluster with numerous cores to obtain the results of our print preparation method within a reasonable timeframe. Nevertheless, despite arguably achieving the highest quality of target texture reproduction at the time of publication, the proposed system was not practically usable.

The reasons for this inefficiency of MC path tracing are manifold: with a high-albedo, forward-scattering material such as the white core of 3D printouts, light paths tend to grow in length while still contributing considerably. Thus, a high number of samples is required to reduce variance – we use 500 spp in Chapters 3 and 4. With increasing object volume, the MC estimator gets less efficient, as light paths grow longer and longer.

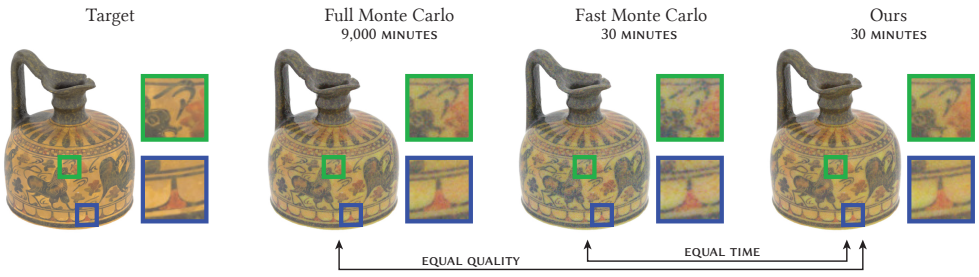


FIGURE 5.1: *We propose a neural scattering compensation for 3D color printing. Comparing to a method which uses noise-free Monte Carlo simulation our technique achieves 300× speedup in the above case while providing the same quality.*

In this chapter (published as [Rittig et al. 2021]), we address this significant bottleneck of the previously described color 3D printing techniques (Chapters 3 and 4) by proposing strategies (Figure 5.1) to speed up the included light transport simulation. We start (Section 5.2) by analyzing the impact of a reduced number of samples used in the MC simulation. Our experiments demonstrate the robustness of our refinement techniques to noisy appearance predictions and, consequently, significant speed benefits at only moderate quality loss. To further improve the quality of the achieved results, we leverage advances in machine learning and propose a data-driven forward appearance model using a deep neural network.

We rely on key insights from previous work in both MC-based color 3D printing and neural volume rendering. Based on a learning framework devised for approximating multiple-scattering in clouds [Kallweit et al. 2017], we build (Section 5.3) a full replacement for the MC forward prediction inside an appearance refinement loop. We depart from querying the learned predictor inside a MC estimator and obtain the final low-variance estimate from a single inference. Furthermore, we propose modifications to the architecture in the form of weight sharing, that relax the cloud-specific scope.

Our improvements offer two benefits. First, our solution generalizes better from simple geometries like cubes and spheres to more complex 3D geometries. Second, the network is able to predict a wider continuum of scattering parameters just from a sparse discrete training set. We train our neural network to predict light transport in heterogeneous volumes akin to composite multi-material prints with spatially varying albedo *and* density using data from MC renderings with complex volumetric textures. Our dataset allows the network to learn correct predictions for more general material arrangements than those typically found in 3D printouts. We argue that the dataset is critical for reliable execution of the iterative refinement as it encounters differing intermediate solutions. Moreover, as we will show, the neural predictor is also suitable for (spectral) soft-proofing during 3D print preparation.

With the new prediction integrated into the color 3D printing pipeline, and running on a single, GPU-equipped workstation, our approach outperforms previous work running on a multi-core workstation of similar caliber by around two orders of magnitude in speed. As a result, we enable the quality of previous methods at speeds suitable for more practical applications, without having to resort to massively parallel compute clusters. Finally, we also provide a comparison of heuristic backward refinement to recent differentiable rendering methods. While both perform on a comparable quality level, the heuristic shows a distinct convergence benefit.

5.1 Behavior of Volumetric Scattering

Volumetric light transport is typically a complex mapping between many volume points (voxels) and fewer surface points. While the space of possible volume assignments is wide in range, the gamut spanning their surface appearances is somewhat limited. As can be seen from the first two examples in Figure 5.2, vastly different volumes can exhibit comparable color saturation. That is because generally the influence of voxels decays with distance to the light source, i.e., the surface in this case. The deeper any absorptive material is embedded, the more it bleeds into the surrounding and affects local contrast, absolute brightness and texture sharpness on the surface, as can be observed across all examples in Figure 5.2. The overall absorptive power of a volume is physically limited by the geometric thickness and directly impacts the achievable color gamut.

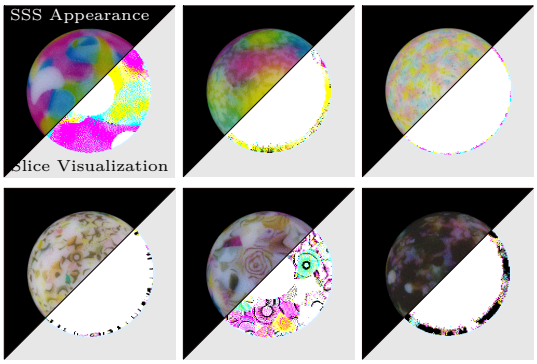


FIGURE 5.2: *Synthetic dataset examples: Monte Carlo predictions of generated volumes in a* 10 mm *sphere. The generator's output (Section 5.3.1) ranges from 3D Voronoi-like textures, to surface extrusions and thin surface coloring. Deeper voxels have less influence on color saturation but limit contrast and edge sharpness.*

5.2 A First Approach to Fast Print Preparation

The major bottleneck in our scattering-aware 3D printing pipeline is the computationally demanding Monte Carlo prediction. A simple approach to reduce the computation time of any MC-based algorithm is to reduce the number of samples at the cost of higher variance. We experimented with preparing a print with only four samples per surface point and found that, against expectations, our previously developed approach (Chapters 3 and 4) is robust against a low-quality forward prediction. In Figure 5.3, we compare this object with a preparation time of less than 30 core hours against a reference object optimized with 512 spp taking over 3,000 core hours. At first glance, it seems like the refinement loop averages out the noise over iterations and still converges to a reasonable result. However, compared to the high-sample MC refinement result, the quality degradation is visible; the color difference values against the target, in terms of mean CIE dE 2000, are also noticeably worse.

Also visually, a low sampling rate is not suitable for soft proofing purposes because it exhibits distracting (color) noise as demonstrated in Figure 5.3c. Image denoising methods could arguably be of help in that case, but they are unsuitable for predictions inside the refinement loop. Denoising images heavily relies on regularly sampled neighborhood information and thus cannot trivially be applied to our 3D voxel representations of *surface* radiance. Alternatively, volume path guiding [Herholz et al. 2019] appears promising to match the requirements for visual inspections



(A) *Optimized print MC 512 spp dE '00:* 8.43

(B) *Optimized print MC 4 spp dE '00:* 12.15

(C) *Volume preview MC 4 spp*

FIGURE 5.3: *High-sample MC previews showing the effect of reduced MC sampling-rate during refinement (a, b). A single 4 spp MC preview shows the level of quality the refinement is coping with (c).*

but ultimately does not promise enough speed-up for the print preparation to become practical.

With the aspiration of an approximative but fast appearance prediction to the high-dimensional problem of volume scattering, we shift our focus to data-driven methods. These have been used in the past to accelerate light transport simulation in homogeneous [Vicini et al. 2019] and even heterogeneous volumes [Kallweit et al. 2017]. The advantage of these methods is that they do not introduce high-frequency noise artifacts and have a constant cost per surface point, independent of scattering values and geometric features. Deviating from previous work, we do not employ (biased) machine learning inside an MC estimator. Instead, we eliminate the stochastic process completely and learn to predict the final, biased result.

5.3 Appearance Prediction Network

Our goal is to yield the full subsurface-scattering appearance of a surface point given a volumetric patch around that point using a Radiance Predicting Neural Network (RPNN) [Kallweit et al. 2017; Panin et al. 2019]. We assume uniform, omnidirectional incidence over the surface to account only for the subsurface part of light transport. The network ultimately replaces the MC forward prediction in an appearance optimization pipeline for 3D printing. Therefore, our learning ecosystem is targeted towards this application.

Problem Statement Formally speaking, we try to find a function $P\colon \mathcal{W}, \mathbb{R}_+^F \to \mathbb{R}_+$ with weights $\mathcal{W}$ that given a set of features $F_{\mathbf{x}_o}$ around point $\mathbf{x}_o$ results in the total in-scattered light from below the surface. This energy is described as the integral

$$L_o(\mathbf{x}_o, \omega_\mathrm{o}) = \int_A \int_{\mathcal{H}^2} S(\mathbf{x}_o, \omega_\mathrm{o}, \mathbf{x}_i, \omega_\mathrm{i}) L_i(\mathbf{x}_i, \omega_\mathrm{i}) |cos\,\theta_i|\ \mathrm{d}\omega_\mathrm{i}\ \mathrm{d}A\,, \tag{5.1}$$

where S is the bidirectional scattering surface distribution function (BSSRDF) between an incident and outgoing point $\mathbf{x}$, and direction ω respectively. A denotes the object's surface area, $\mathcal{H}^2$ the hemisphere above the surface and L_i stands for the incoming illumination. Our target application replicates diffuse appearances and thus we are only interested in the outgoing radiance along the surface normal $\mathbf{n}_\mathbf{x}$. Also, the assumption of diffuse illumination allows for the inner integral to be integrated out, leaving

$$L_o(\mathbf{x}_o, \mathbf{n}_{\mathbf{x}_o}) = \int_A S(\mathbf{x}_o, \mathbf{n}_{\mathbf{x}_o}, \mathbf{x}_i)\ \mathrm{d}A\,. \tag{5.2}$$

Learning an approximation of the remaining integral, we seek weights $\mathcal{W}$ of the network that minimize the L^2 distance between the RPNN prediction P and a ground truth Monte Carlo simulation L_o for all T training examples,

$$\arg\min_{\mathcal{W}} E(\mathcal{W}) = \frac{1}{T} \sum_{t=1}^{T} \left(L_o(\mathbf{x}_t, \mathbf{n}_{\mathbf{x}_t}) - P(\mathcal{W}, F_{\mathbf{x}_t})\right)^2\,. \tag{5.3}$$

In the following, we will first describe the construction of a dataset and how features F are derived from it. Next, we detail our contributions to RPNNs, and finally, we describe their deployment. We drop the location index $\mathbf{x}$ for brevity.

5.3.1 Training Dataset

For our supervised regression-learning task, the datasets consist of pairs of volumes and rendered points on their surface. The volumes store trichromatic scattering and absorption coefficients in each voxel. The target values represent the surface reflectance at a given point in direction of its normal. We choose to work with a synthetic dataset of basic shapes filled with generated volume assignments. We summarize the properties of the dataset here and refer the reader to our supplementary material for more details.

Shape Shape is one of the two major components affecting a volume's appearance. Our dataset includes basic convex geometry in the form of a cube, a rotated cube, a sphere and a thin cuboid. The former three objects are instantiated at different sizes between 5 mm to 20 mm without additional rotation. The thin slab includes thicknesses from 1 mm to 2.5 mm. Despite such sparse sampling of geometric variation, we notice comparably little artifacts on arbitrary geometry thanks to our weight sharing approach (Section 5.3.3).

Texture The volumes are highly heterogeneous mixtures of discrete base materials. We generate procedural volumetric textures and intersect them with geometry to produce a single volume entity. Based on spatial, deterministic OpenSimplex noise [KdotJPG 2014], the textures are assembled using a variety of parametric functions. Through thresholding and manifold parameter dependency we build rich procedural volumes that exhibit multiple scales of frequencies, a wide dynamic range and different levels of chromaticity. Refer to Figure 5.2 for examples of generated volumetric distributions and their surface appearance.

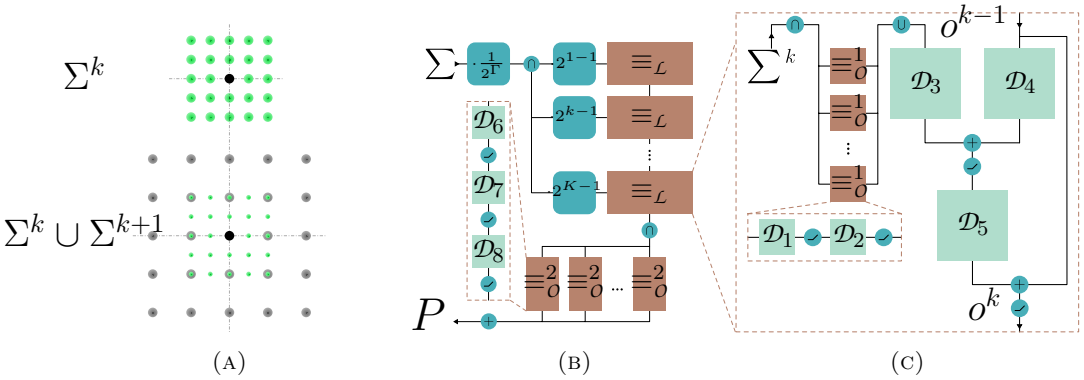
Values The values of the dataset describe scattering coefficient σ_s and absorption coefficient σ_a in inverse path length $[\mathrm{mm}^{-1}]$. These belong to five measured 3D printing materials (Chapter 3) each with three color channels. The measurements are fitted to a constant Henyey Greenstein phase function model with a forward scattering lobe of $g = 0.4$. We mix data from all channels together to train a single-channel network. For training stability, the values will also be rescaled as detailed in Section 5.3.4.

As the 3D printing materials vary in both extinction coefficient σ_t and albedo α, our approach is the first one to tackle fully heterogeneous volume scattering with machine learning.

5.3.2 Feature Descriptor

With light transport being an inherently global phenomenon, the outgoing radiance at each point on the object's surface depends on both local and global volumetric distribution of materials underneath. But in an machine learning (ML) setting, training on the full volumetric data becomes computationally intractable for larger objects and gigabytes of dense voxel grids. Motivated by these observations, Kallweit et al. [2017] introduced a sparse, hierarchical descriptor that extracts the most important features to describe scattering through a point. This descriptor captures the high-resolution details within the immediate vicinity of a considered surface point but reduces (exponentially) to a sparser representation for the distant volume.

Similar to Kallweit et al., our feature descriptor F extracts a set of volume *stencils* Σ^k at different levels $k \in \{1, \ldots, K\}$. Higher (coarser) levels represent a larger portion



(A) (B) (C)

FIGURE 5.4: *Network Architecture: (a) visualizes the spatial relationship between levels inside one stencil Σ and the central point $\mathbf{x}$. (b) Network overview where $\mathcal{D}$ blocks mark dense layers with trainable weights, $\equiv$ blocks indicate shared weights and round edges/circles signify arithmetic operations ($\cap$: slicing, $\cup$: concatenation) The factor $\frac{1}{2^\Gamma}$ keeps values in range for training stability. (c) Detail of one RPNN block.*

of the volume but at lower resolution. Obtaining a single stencil Σ^k involves sampling the neighborhood of surface point $\mathbf{x}$ and extracting a set of scattering and absorption coefficients at $|\Sigma^k| = I$ locations q_i^k,

$$\Sigma^k = \left\{\sigma_\mathrm{s}(q_1^k), \sigma_\mathrm{a}(q_1^k), \ldots, \sigma_\mathrm{s}(q_I^k), \sigma_\mathrm{a}(q_I^k)\right\}. \tag{5.4}$$

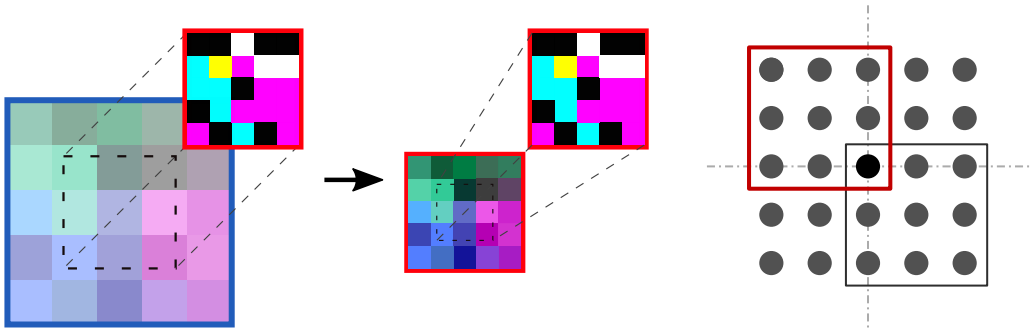
For each hierarchy *level*, the considered volume scales by 2^{k-1} with descriptor dimensions staying constant (I). With all stencils being centered around the same surface point $\mathbf{x}$ (Figure 5.4a), the resolution of the neighborhood effectively decreases exponentially. The final feature descriptor is a set of all stencils in the hierarchy $F = \Sigma = \bigcup_{k=1}^{K} \Sigma^k$. In practice, we use $K = 9$ levels and a stencil dimension of $I = 5 \times 3 \times 9$ with $[-2, -1, -4]$ and $[2, 1, 4]$ being corners of the index-space offsets from $\mathbf{x}$.

In contrast to Kallweit et al.'s requirement for an elongated and rotated stencil towards a directional light source, we take advantage of diffuse illumination in our setting and simplify the data lookups. We sample the volume with stencils axis-aligned to the voxel grid and of cubic world-space size. Due to asymmetry in the voxel resolutions of our 3D printer ($600 \times 300 \times 940$ DPI), the actual cell numbers still differ per axis. Our stencil of $5 \times 3 \times 9$ samples then covers a volume of $211\,\mu\mathrm{m} \times 254\,\mu\mathrm{m} \times 243\,\mu\mathrm{m} \cdot 2^{k-1}$ and reaches 65 mm of the surrounding on the coarsest level. Depending on channel and material, this distance is 146 to 1560 mean-free-path-lengths.

Values in the coarser levels $k > 1$ are averaged from the lowest level to reduce aliasing; volume partials outside the objects are treated as zeros. This is especially important with our high-frequency data for 3D printing as opposed to smoother value changes in clouds. Our filtering approach, detailed in Section 5.3.4, differs from mip-mapping in order to support perfect alignment to the central point $\mathbf{x}$ on coarser levels.

5.3.3 Network Architecture

Our general network architecture is based on the proposal of Kallweit et al. [2017]. Their architecture is based on a Multilayer Perceptron (MLP) and uses skip connections between consecutive blocks. One of their insights is, instead of passing all inputs from the top, each level of the descriptor Σ^k is passed into one RPNN block (Figure 5.4b). Each block consists of three dense layers ($\mathcal{D}_{[3,5]}$ in Figure 5.4c) in which



(A) *Unifying Level Size* (B) *Quadrants/Octants*

FIGURE 5.5: *Weight sharing. (a) Hierarchy levels are rescaled to the same size (with adjusted optical density) to enable processing with the same weights across levels ($\equiv_L$). (b) On each level, the stencil is split in half per axis and we share weights ($\equiv_O$) over correctly rotated quadrants (octants in three dimensions) to enforce symmetry following from rotational invariance around the black central point* $\mathbf{x}$.

residuals are mixed with knowledge from all previous levels (o^{k-1}). In the end, three dense layers $\mathcal{D}_{[6,8]}$ refine the gathered information and produce a single output value for P. Please refer to our supplemental for a complete specification.

In the remainder, this basic architecture is referred to as the *baseline* solution. In Section 5.4 we show that it already performs surprisingly well in our setting. This baseline explicitly differs from Kallweit et al.'s method in the number of levels $K$and the per-ray inference with explicit MC single-scattering simulation. As our domain features distinct media (density, heterogeneity) and illumination (diffuse) from the cloud setting, we include straightforward adapations to avoid computational overhead.

For our architecture we incorporate further domain knowledge of volumetric light transport and enhance the baseline by strategically placed weight sharing ($\equiv$) that implements two physical invariants: isometric scaling and rotational invariance.

Isometric Scaling Intuitively, the network should compute the same function on all K levels of the hierarchical volume representation Σ, that is, determine how much each voxel affects the central point of interest; the only difference is the geometric scale of the input representation. We enforce that identity by isometrically rescaling values according to the extent of each level of our multi-scale input.

In contrast to the well-known relationship for area and volume that scale with powers 2, and 3, respectively, scattering coefficients σ_s and σ_a scale with the power of -1 when the geometric reference frame changes. Figure 5.5a visualizes that shrinking a coarser level requires the optical density to be increased to keep the same appearance.

As shown in Figure 5.4b, our network scales all levels Σ^k to the size of the lowest (finest) level, increases optical density ($\cdot 2^{k-1}$) and processes them using the same weights ($\equiv_L$). While this weight sharing primarily reduces the network size, it also, implicitly, augments the discrete-valued training dataset in two ways:

1. Averaging of values for coarser levels ($k > 1$) creates more intermediate data points within the convex hull and mitigates the limitation to discrete materials.

2. Increasing the density then uniformly scales the convex hull of $(\sigma_\mathrm{s}, \sigma_\mathrm{a})$ to larger value ranges.

Figure 5.6 shows the convex hulls of 10^6 exemplary training points on different levels. Without the isometric scaling, the network would *only* see values from the lowest

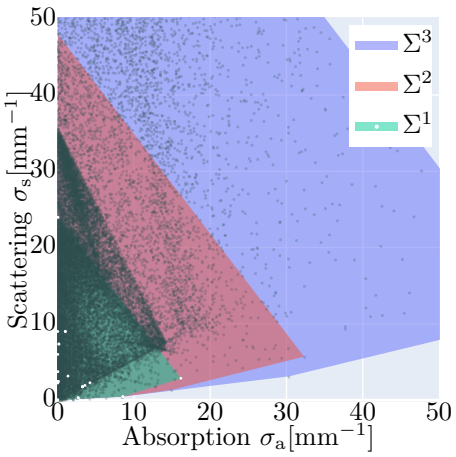


FIGURE 5.6: *Input values of three Σ levels during training: On the lowest level Σ^1 (white points), these are the 15 discrete values corresponding to CMYKW at three channels. Higher levels (with isometrically scaled density) also contain interpolated values according to the mixing ratios of base materials.*

level (green hull, white discrete points). Using our architecture, the network is forced to learn a consistent function for $\equiv_{\mathcal{L}}$ across a larger range of values. Note how averaging with non-scattering "air" on coarser levels also produces less dense data points towards the coordinate origin.

Rotational Invariance Object rotation against the grid axes should not impact the appearance. We can enforce this invariance by sharing weights between symmetrical parts within a grid around our central point of interest $\mathbf{x}$. As illustrated in Figure 5.5b, we split the data in all three dimensions in the middle (with potentially one row/column overlap), and process all eight parts (*octants*), after correct alignment, using the same filters ($\equiv_O$). We apply this concept twice: First, to each input stencil Σ^k within the already shared level $\equiv_{\mathcal{L}}$ as depicted in Figure 5.4c ($\equiv_O^1$). Second, at the end of the network ($\equiv_O^2$ in Figure 5.4b) to produce a single output value per octant. This way, one can weakly enforce rotational invariance without resampling data or using a grid-less learning framework. It also implicitly augments the data without explicitly storing duplicates.

Overall, our network consists of only eight trainable layers ($\mathcal{D}_{[1,8]}$) which greatly reduces the number of parameters and improves training speed. We exploit domain knowledge with our weight-sharing and support generalization of the network to unseen volume arrangements and material values.

5.3.4 Training Protocol

Our implementation is based on Tensorflow [Abadi et al. 2015] and is using the Adam optimization algorithm [Kingma et al. 2014] with a 0.001 learning rate. The networks are trained on dual-socket Intel Xeon CPU E5-2680 machines (40 threads) with 256 GB memory and a Nvidia Tesla K40m GPU (11 GB memory). Both, the implementation and the dataset to this paper will be published alongside the paper.

Dataset Loading For optimization with stochastic gradient estimation, the order of training samples needs to be shuffled. With the scale of volumetric data we are working with, the datasets typically do not fit into memory. Despite one stencil Σ being only 11 kB, with 200×10^6 surface points a dataset easily reaches into the order

of 2 TB. These are sizes that cannot even be streamed from storage quick enough for GPU processing. Thus, we employ an out-of-core method to process larger datasets using a sliding window approach over volumes within which we randomize samples. Channels of multi-channel (RGB) volumes are mixed to train a single-channel network as light transport is channel-independent.

Training on 340 volumes (sliding window of 110 volumes) took from 32 to 44 hours for the different networks variants. In our setup, the epoch size was set to 2M data points, and the whole training converged after 500 epochs.

Stencil Generation With memory consumption being a concern, we trade storage for computation and do not store explicit lower resolution representations of volumes. Instead, we repeatedly downscale particular neighborhoods of a requested sample point $\mathbf{x}$ using a box filter that doubles in footprint for each level as described in Section 5.3.2. Efficiently extracting spatial averages on the CPU is key to provide data points quickly for training on the GPU.

We chose to work with Summed Area Tables (SATs) that we implement in a two-level hierarchical way to avoid issues with numerical precision on larger volumes as well as reduce the memory consumption. A hierarchical split allows for storing a sparser representation of the SAT in constant areas by replacing their leaf-level tile with a constant value. Using SATs, it takes a constant-time lookup of eight values to calculate the average for any axis-aligned box in the volume. When querying for a training sample we build a hierarchy of volumetric patches Σ around a surface point $\mathbf{x}$. In contrast to a precomputed downsampling approach, this aligns perfectly for each $\mathbf{x}$ and avoids interpolation artifacts on coarser levels. Avoiding the exploding gradient problem, we rescale all input values to the network by a constant $\frac{1}{2^\Gamma}$ with $\Gamma = 6$ as visualized in Figure 5.4b.

5.3.5 Inference

Unlike previous work [Kallweit et al. 2017; Panin et al. 2019], our RPNN is not integrated with a path tracer. Instead, it is a full replacement of a renderer within a scattering-aware optimization loop (Chapter 4) for fabricating textured objects with a color 3D printer. The RPNN is able to predict the full subsurface light transport in one go without a separate single-scattering simulation as Kallweit et al. propose. We are able to do that, because the radiance function is smoother in our setup for two reasons. First, the lighting setup is uniform incidence over the surface, whereas Kallweit et al. [2017] account for sharp directional illumination. Second, our estimates are not sampling-based, i.e., integrated within the loop of a path tracing renderer as previous work [Kallweit et al. 2017; Panin et al. 2019; Vicini et al. 2019] did. Thus, the network is only asked to predict the averaged final appearance of the subsurface scattering, not individual paths.

On inference, we extract volumetric patches around each surface point as described in Section 5.3.4 and run them through the network's computational graph in parallel batches. Each color channel / spectral band is treated individually. This results in a linear runtime dependency on the number of surface points and color channels. Differing from MC rendering, the prediction quality stays constant without the need for further evaluations to improve the SNR.

5.4 Results

We evaluate our results on geometries and textures that have not been part of the training or validation dataset. These objects (shown in Figs. 5.8 to 5.11) are chosen to exhibit certain properties we deem important for our application.

Wedge (6 cm) has varying thickness from 0.5 mm to 6.5 mm. The mirrored texture in complementary colors on each side is designed to show cross-talk at different spatial frequencies and thicknesses.

Showcase (3 cm) features positive and negative curvature. The white side enables judgment of the scattering magnitude on thick geometry.

Octostar (3 cm) has constant thin geometry in various orientations. This object is particularly challenging as it contains loops.

Cork Bone (4 cm) has a high-frequency texture over a wide range of curvature, thickness and axis-orientation.

Yellow Vase (4 cm) has a rich texture and a thin handle.

Red Vase (5 cm) is partially hollow with average positive and negative curvature.

Cat (6 cm) features thin tubes in various orientations.

Thin Plane (3 cm) is a 0.5 mm slab with textures on both sides.

The volumetric arrangement of printing materials is achieved either by extruding a mixture of surface color into the depth or through iterative refinement (Chapter 4).

5.4.1 Scattering Compensation

We integrate our RPNN into the core of an optimization loop that refines volumetric distributions of materials for texture sharpness and contrast described in Chapter 4. Such scenario involves repeatedly predicting the surface appearance in order to gauge if adjustments to the volume have shown its desired effect. Any potential systematic errors would get amplified over more iterations and manifest themselves in the material arrangement.

In a typical run, it requires 8 to 12 iterations for a refinement to converge. Figure 5.7 plots the perceptual color differences between target and simulated objects in terms of CIE dE 2000 [Sharma et al. 2005] at each iteration. The figure compares

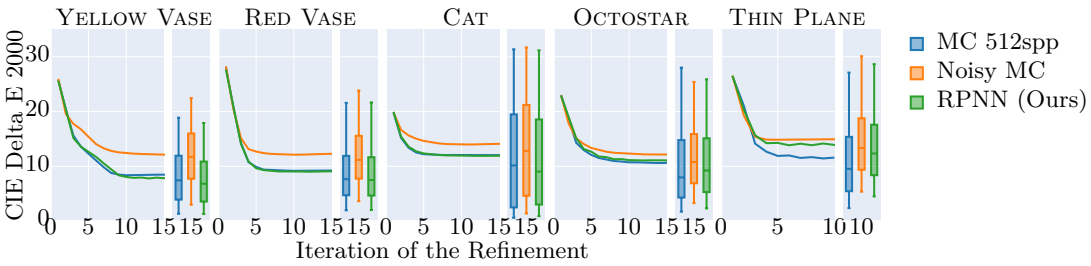


Figure 5.7: *Convergence plots for the five fabricated objects in Figure 5.8 using three prediction methods: MC-based refinement with 512 spp, our network-based solution and an equal-time MC prediction. Average CIE dE 2000 metric values against the target object are derived from a reference MC (re-)prediction (512 spp) for all methods. Boxplots show the 0.05, 0.25, 0.5, 0.75, and 0.95 quantiles of the last iteration.*

a reference refinement (MC 512 spp) to our RPNN-based method and an equal-time MC rendering (4 to 16 spp, depending on the object's surface area). For comparability, the volumes of all three methods have been re-rendered using MC 512 spp to obtain the difference values with equal noise-level. The comparably high level of CIE dE 2000 errors for each model reflects the fact that exact texture reproduction is often impossible in the presence of significant scattering for the given amount of high frequencies in the texture or the given geometry. Note that the error is calculated against a non-gamut-mapped target.

In most cases, our network-based refinement is on par with the established method of a 512 spp Monte Carlo rendering. For the challenging case of the Thin Plane our solution performs more comparably to a low-quality MC prediction. This is explainable from the training data, as the thickness (0.5 mm) lies outside the range of our training set (1 mm to 20 mm).

Figure 5.8 depicts MC preview renderings of the objects obtained from these refinement runs. Overall, reproductions with all three methods show acceptable performance. The results of our RPNN-based method are hard to distinguish from those of the MC 512 spp. Despite remarkable performance of the low-sample MC approach given the very noisy predictions, it shows systematically lower quality than the two other methods. For example, for this method, we observe noise patterns on Thin Plane and Yellow Vase, reduced sharpness on Red Vase, and color shifts (within yellowish patches) on Cat. Visually, the distances between methods on Thin Plane match the differences from the convergence graph in Figure 5.7.

The computational time required to run the compared methods is listed in Table 5.1. Our method was run on a desktop-grade workstation with a single Nvidia RTX 2070 GPU and 8-thread Intel CPU. Due to the heterogeneous hardware requirements, we are forced to convert the high-quality method's runtime, originally on a large compute cluster, to a comparable competitor. Again, low-quality MC predictions are based on equal time to RPNN-based prediction and thus not explicitly listed. Our method performs significantly faster with two orders of magnitude speedup for most objects. This ratio is depending on object geometry and grows with geometric volume. The RPNN predictor has a constant runtime per surface point, independent of underlying curvature, geometric volume or medium values. In contrast, with forward scattering materials and high albedo, a MC path-tracer's runtime is dependent on the volume underneath each surface point, as paths can scatter more and deeper into the medium. In the supplementary material of [Rittig et al. 2021], we provide additional speed comparisons with variants of Mitsuba2 [Nimier-David et al. 2019] including GPU path-tracing.

Table 5.1: *Prediction timings (per object, per iteration) for results in Figure 5.8 in the equal-quality setting. MC 512 spp is recalculated to a single 32-thread CPU system – NN is evaluated on a 8-thread CPU + GPU system. Low-spp MC is equal-time to RPNN.*

Object	Surface points	MC 512 spp	RPNN (Ours)	Speedup
Yellow Vase	1.089×10^6	38 745 s	128 s	303×
Red Vase	1.092×10^6	39 206 s	140 s	280×
Cat	0.637×10^6	20 987 s	75 s	280×
Octostar	1.833×10^6	29 059 s	236 s	123×
Thin Plane	0.348×10^6	428 s	16 s	27×



FIGURE 5.8: *Comparison of prediction methods inside a refinement loop: (a) target models (b) using 512 spp MC prediction, (c) using equal-time MC prediction with* 4 *to* 16 *spp, and (d) our proposed solution using RPNN-based predictions. Refinements with low-sampling MC prediction produce lower quality of texture reproduction compared to the other two. Timings are reported in Table 5.1.*

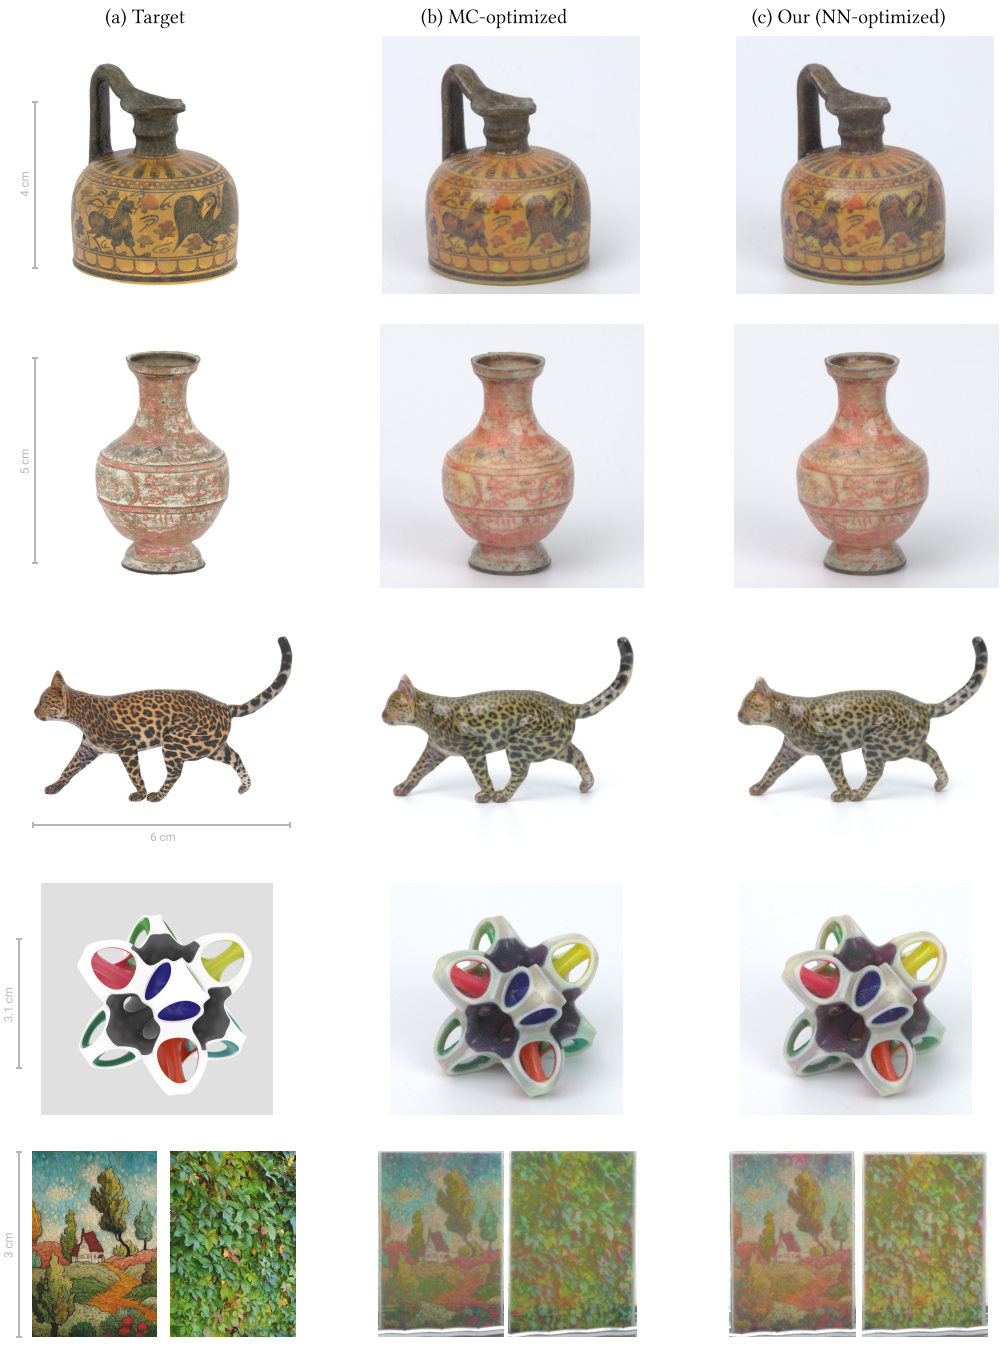
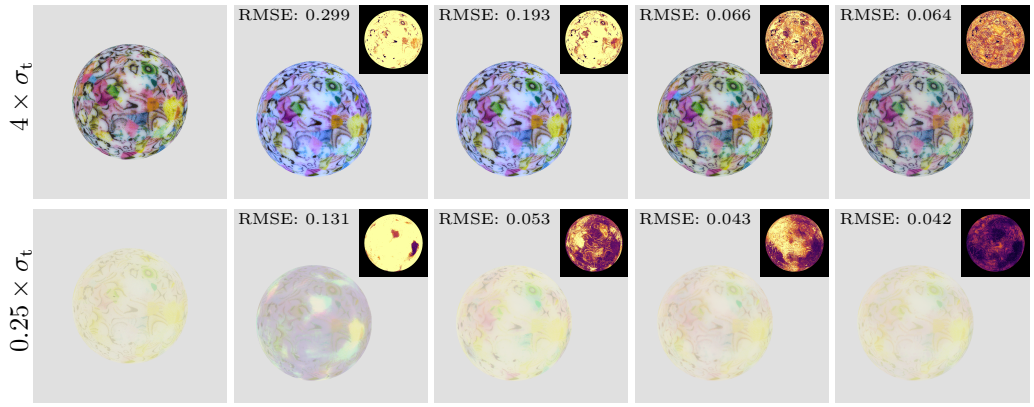


FIGURE 5.9: *Comparison of photographed printouts to the rendered (a) target models: (b) our refinement solution with MC-based predictions (Chapter 4), (c) our proposed solution with RPNN-based predictions. The bottom object is a thin planar slab (0.5 mm thick) with two different textures on its front and back sides.*



(A) *MC 512spp* (B) *Baseline* (C) *with* $\equiv_L$ (D) *with* $\equiv_O$ (E) *with* $\equiv_L + \equiv_O$ *(Ours)*

FIGURE 5.10: *Single forward predictions comparing the generalization of network architectures to new values. The* 15 mm *sphere is a training example with adjusted density. The top row is optically four times thicker, whereas the bottom row is four times thinner. Our weight-shared architecture handles values outside the training range with less artifacts. CIE dE 2000: 0*  *10*

Finally, we also show photographs of actual 3D printouts prepared by both high-quality methods in Figure 5.9. These are fabricated on a Stratasys J750 printer using the VeroOpaque material family, sanded and varnished with a clear-coat. On all fabricated samples, once again, the two high-quality methods are difficult to discern. Exception to this trend, as mentioned before, is the fabricated THIN PLANE.

5.4.2 Network Generalization

In the last subsection, we showed our predictor's performance averaged over many iterations. Here, we focus on single forward predictions and show raw network output values by visualizing the surface voxel values.

Evaluating our improvements to the Radiance Predicting Neural Network (RPNN) framework [Kallweit et al. 2017; Panin et al. 2019], we compare different architecture variants within our setting. Because of our weight sharing ($\equiv$) in multiple key locations, the number of trainable parameters, i.e., the learning capacity, varies between all four architectures as illustrated in Table 5.2. Against intuition, the variant with octant weight-sharing $\equiv_O$ increases the number of weights because it adds one overlapping column in each dimension (see Figure 5.5b) for every layer.

Values First, we compare generalization over new scattering and absorption values outside the range of the 15 discrete ones in the training dataset. We take a training example and rescale its density to simulate different scattering materials. This way, geometry and volume arrangement have been seen by the networks before, but the values are in a different range. Figure 5.10 depicts predictions for all four variants.

TABLE 5.2: *Number of trainable parameters for each architecture given the same stencil dimensions* $I = 5 \times 3 \times 9$.

Architecture	**Baseline**	**with** $\equiv_L$	**with** $\equiv_O$	**with** $\equiv_L + \equiv_O$ **(Ours)**
Parameters	2 046 871	365 851	3 767 341	706 861

For both rows, the baseline architecture exhibits mispredictions of the appearance, whereas ours handles this unfamiliar environment well. This demonstrates the benefits of our weight sharing over levels ($\equiv_{\mathcal{L}}$), as they implicitly augment the discrete training values to greater ranges as illustrated in Figure 5.6. In the baseline architecture, each level overfits to a distinct set of (interpolated) values in the lowest level's range.

Shape The results in Section 5.4.1 suggest that the network prediction performs well for arbitrary geometry. Here, we further verify this and dissect this generalization over shape per architecture. Figure 5.11 shows the quality level of predictions on general geometry. Despite training only on a very simplistic set, all architectures perform acceptably on concave, thin regions and arbitrary surface orientation against the stencil grid.

While RMSE error values are rather comparable over the architectures, the baseline exhibits distinct bias in multiple cases. On the white face of the Showcase object and around the blue+green cone where it intersects the coordinate axes, the baseline has banding artifacts. For the Wedge it fails to predict the crosstalk in the thin region correctly.

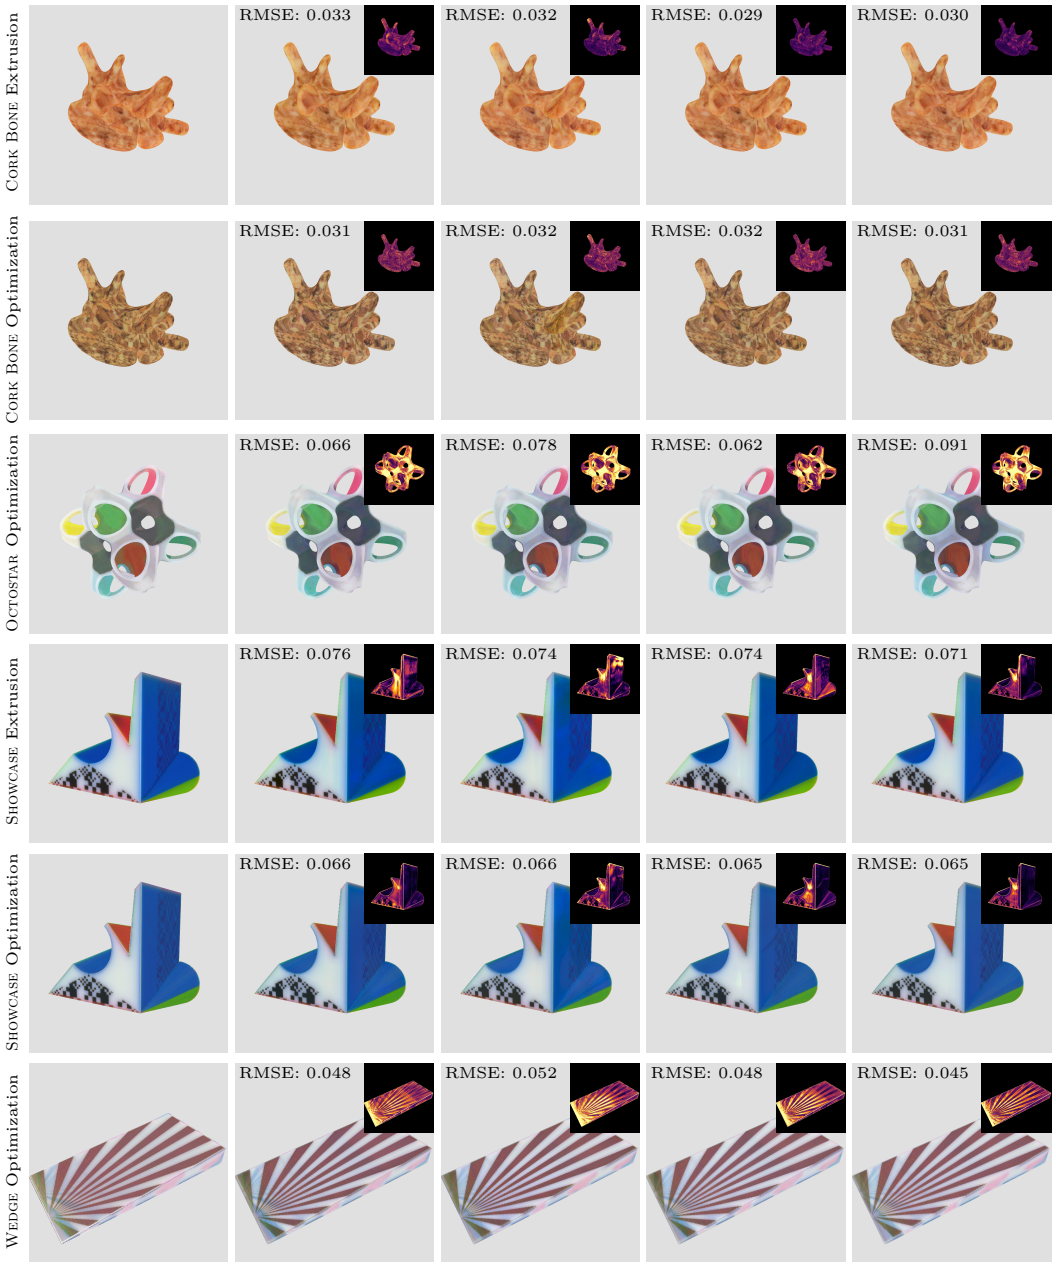
On the Cork Bone model the bias increases for the thinner parts and slightly less of the weight-shared architecture. This can be attributed to the octant weight sharing $\equiv_{\mathcal{O}}$, as it is less prone to overfitting on specific voxel arrangements in particular voxel positions. Because the weights are shared across all eight octants in the coordinate grid, training for individual voxel arrangements and occupancies in every voxel index is not necessary. Instead, training data from all eight directions contributes to the quality in all directions. Achieving true rotational invariance would require a spherical architecture and representation – a potential topic for future research.

On the challenging case of Octostar, the level sharing $\equiv_{\mathcal{L}}$ introduces considerable bias on the hollow bridge regions between colored corners. Architecture variants without this feature, such as the baseline, exhibit this problem less pronounced. We discuss this limitation in further detail in Section 5.5.

5.4.3 Application: Spectral Prediction

Our data-driven method can also be used for spectral volume predictions in a relatively straightforward manner. This application has important implications, as spectral 3D printing has shown promising results [Shi et al. 2018]. While, in this paper, we show only spectral predictions, spectral scattering compensation for 3D printing remains an exciting direction for future work.

Because the network is trained on a single color channel only, it can be repeatedly applied to predict individual spectral bands. In Figure 5.12, we demonstrate that our method matches the quality of an MC spectral rendering. This object is filled with *virtual* materials whose scattering properties are hand-crafted to mimic the spectral properties of a CMYKW ink set. The network predicts, just like the MC renderer, 31 spectral bins, from 400 to 700 nm, which are converted into RGB using standard procedures [Smith et al. 1931]. For predicting arbitrary spectral curves, a good generalization over scattering values is important. Our network architecture achieves that, as we demonstrate in the previous subsection. As with any spectral image, one can easily vary the diffuse illumination spectrum afterwards and predict the look under different illuminants. This result also highlights the benefits of using a network for soft proofing, because the runtime speedup we report above leaves room for more accurate colorimetric prediction within a practical timeframe.



(A) *MC 4,096 spp* (B) *Baseline* (C) *with* $\equiv_{\mathcal{L}}$ (D) *with* $\equiv_{\mathcal{O}}$ (E) *with* $\equiv_{\mathcal{L}} + \equiv_{\mathcal{O}}$ *(Ours)*

FIGURE 5.11: *Single forward predictions comparing four architecture variants over prediction error (RMSE) for general geometry. The baseline, a basic adaptation of Kallweit et al.'s RPNN, is compared against our weight sharing over levels* $\equiv_{\mathcal{L}}$ *and octants* $\equiv_{\mathcal{O}}$*. The objects' unique features are discussed in Section 5.4. CIE dE 2000: 0*  *10*



(A) *MC 128 spp* (B) *RPNN (Ours)* (C) *Difference*

FIGURE 5.12: *Spectral predictions for an object formed by virtual spectrally-defined materials. Our network generalizes well over scattering parameters from only 15 discrete values and can accurately predict continuous spectral curves. CIE dE 2000: 0* *10*

5.4.4 Differentiable Rendering Comparison

With the recent publication of Radiative Backpropagation (RB) [Nimier-David et al. 2020], time and memory costs of differentiable rendering methods become viable for inverse reconstruction of heterogeneous media. The paper contains experiments that hint at a potential application in the field of 3D print preparation. Naturally, the question follows how our heuristic-based iterative scheme compares to such a generic optimization framework. The available implementation of the experiments require us to recreate their specialized setup and compare both inverse methods in a more simplified environment. As the backward heuristic is independent of the used forward model (MC or RPNN) we chose MC for consistency.

The setup presented in [Nimier-David et al. 2020] differs from a traditional 3D printing setup in the following attributes:

- Constant density across the volume.
- Isotropic phase function.
- No printing material discretization (halftoning) in the iteration loop.

This is important to consider for speed (density) and quality (latter two) comparisons. A detailed description of the scenario, a discussion on timings and more results are presented in the supplemental materials of [Rittig et al. 2021]. In the next paragraphs we summarize the outcomes.

With stochastic gradient descent (SGD) being a local optimization method, the starting condition is important. Figure 5.13 compares cutouts of four optimizations. Choosing from several initializations, the best result for RB (a) compares in visual quality to the heuristic refinement (c) with (c) being slightly sharper. For the same initialization as (c), RB leaves texture details washed out (b). The most interesting result is (d) where RB is initialized with the result from (c). There, it minorly adjusts edge sharpness, but overall stays visually constant in the same local minimum.

Looking at the resulting volumetric distribution, slices are very different for each method. The heuristic (c) prioritizes compression of absorptive material close to the surface while RB finds different local minima depending on the initialization. Being designed for a particular initial condition the heuristic will not work stable with other initializations.

In conclusion, differentiable rendering is a general tool that requires further research to be applicable for practical 3D print preparation. Besides the required technical updates (proper volume modeling and material discretization), we see that the generic SGD optimization applied to the translucent printing materials produces a less crispy appearance. SGD's convergence behavior is sensitive to initialization and

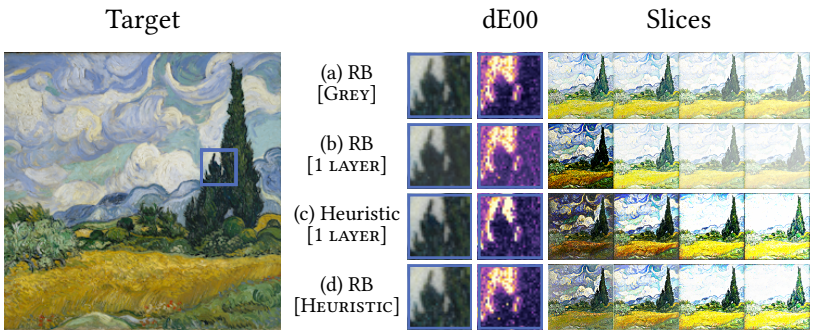


Figure 5.13: *Comparing Radiative Backpropagation (RB) [Nimier-David et al. 2020] with heuristic refinement (Chapter 4) on different initializations (a-d). The best result for RB (a) matches the quality of (c). On the same initialization (b) details are less pronounced. When continuing (c) with RB, the SGD optimization moves only slightly in the local minimum (d). CIE dE 2000: 0* *10*

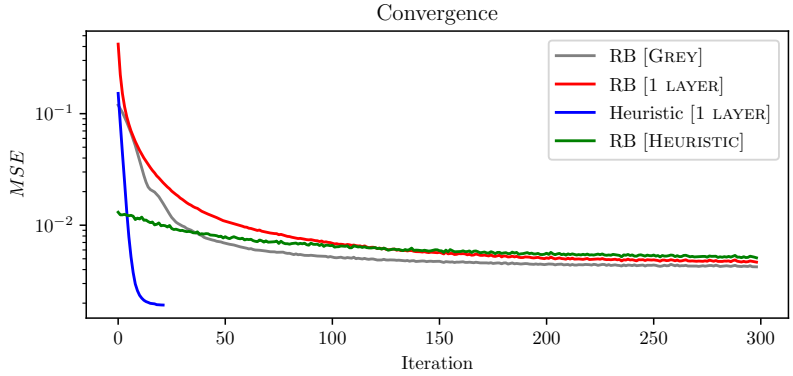


Figure 5.14: *Convergence behavior of Radiative Backpropagation (RB) [Nimier-David et al. 2020] and heuristic refinement (Chapter 4) for results shown in Figure 5.13. The heuristic converged after* 24 *iterations and visually matches the quality of RB [Grey] after* 300 *iterations.*

visual quality depends on the used error metric. A simple metric like MSE is not indicative of perceptual considerations and more perceptual metrics have yet to be developed for surfaces in voxel representation.

Most importantly, in the presence of long mean-free paths, the SGD optimization requires much more iterations to converge. As can be seen in Figure 5.14, the heuristic converges after 24 iterations whereas SGD optimization on MSE requires a long-tailed optimization to increase detail fidelity. When running both methods with the same forward predictor, which is the dominant factor, the absolute speed difference follows the same trend.

5.5 Discussion

Overall our results confirm that double (nested) weights sharing is a good way to regularize training, reduce weights, increase generalization and inject domain knowledge in an RPNN architecture.

With the generalization over values our published, pre-trained network can be useful even for future materials without re-rendering or re-training, given their measured scattering and absorption coefficients lie within the range illustrated in Figure 5.6. In case of vastly different value ranges one can still leverage the published dataset by re-rendering and subsequent re-training.

The results also show limitations of the method. When it comes to very thin objects like our Thin Plane, the network does not extrapolate well from the training data. More examples with reduced thickness would be a first solution to solve that. Also, the architecture exhibits artifacts with highly-concave and closed-loop geometry like our Octostar and parts on the Showcase. In an attempt to find the root-cause, we included similar objects to the training dataset, but the train- and test-time metrics didn't improve. We suspect the stencil representation being ambiguous between a completely solid object with air around and an open loop on higher levels.

Compared to MC rendering, our method provides faster evaluation times, which is independent of scene complexity and density of the medium. Besides practicality of the print preparation pipeline, it is also an important step for learning to generate a material distribution for a target surface appearance. Only with a fast forward predictor one can prepare enough training data to approach the backward step with machine learning.

Our training data is obtained by means of virtual simulations with an idealized voxel model. In reality, the printing process might induce partial voxel mixing, non-smooth interfaces and anisotropic scattering effects w.r.t. the layer orientation. Thus, a machine learning approach based on extensively measured data seems preferable, but requires a set of costly fabricated samples.

In a comparison with differentiable rendering, heuristic refinement performs on an equal quality level, partially with perceptually slightly "crisper" reproduction. But in terms of convergence, a naïve implementation requires ten times more iterations. An exciting trait of our RPNN forward model is that it is also differentiable and can deliver surface gradients w.r.t. volume patches. Untangling gradients shaped like our stencil Σ for a volumetric reconstruction is left to future work.

5.6 Conclusion

In this work we proposed a high-fidelity 3D printing pipeline that matches previous work in quality, while typically being 100 to 300 times faster. The key to this was combining the strengths of a state-of-the-art heuristic iterative optimizer for 3D color print preparation with a neural predictor for surface appearance from heterogeneous subsurface scattering.

As we have shown, despite the limited generality of any data-driven model, the network generalizes well to unseen geometry and material values. This robustness lends our solution for real-world deployment. In effect, this is the first time that full heterogenous material optimization for 3D-print preparation becomes possible within time frames in the order of the actual printing time.

Chapter 6

Towards a Quality Metric for Dense Light Fields

In the preceding chapters, we described a full-color 3D printing pipeline capable of producing high-fidelity objects, even in the presence of inherent light scattering within translucent printing materials. Our approach assumed uniform, omnidirectional light incidence across the surface, without self-shadowing or surface inter-reflections. Acting as a printer driver, the proposed software converts a user-specified target 3D model into a volumetric material arrangement suitable for 3D printing.

The task of determining this volumetric material arrangement was formulated as an inverse rendering problem and solved through an iterative refinement approach. Central to the optimization loop is the discrepancy between the predicted appearance of translucent materials and a given opaque target appearance. So far, we limited the scope of our work to creating objects with a Lambertian surface appearance, meaning they possess a uniform appearance irrespective of the viewing angle. In such setup, we used per-point difference metrics over the surface of the 3D objects to be fabricated, utilizing either L2 in the linear RGB color space or CIE dE76 (CIELAB). The latter operates within a perceptually uniform space, potentially enhancing the subjective quality of reproduction.

While our work primarily addresses materials and appearances that remain consistent across viewing angles, we recognize the importance of fabricating spatially- and angularly-varying appearances. To achieve this, a print-preparation system must compute target/prediction differences in a multi-view domain. Light field representation emerges as a promising solution, as it captures the entire light distribution in a scene [Levoy et al. 1996] and describes angular light propagation in every direction throughout 3D space. However, determining the discrepancy of light fields is not immediately evident.

Moreover, due to the enormous storage requirements, light fields are usually sparsely sampled in spatial and angular dimensions, stored using lossy compression, and reconstructed later. It remains unclear how the distortions introduced along the way affect perceived quality. In the context of new multiview and light-field displays [Masia et al. 2013; Wetzstein et al. 2012], and light-field 3D printing [Zheng et al. 2020] densely sampled light-fields spanning wide baselines are highly needed to produce accurate and sufficiently reliable reconstructions.

Similar problems have been addressed for 2D images, videos, and sparse multiview content. Many quality metrics have been designed to predict perceived differences between various versions of the same content [Aydin et al. 2010]. However, measuring quality for dense light fields still remains a complex task. While several works applied the existing metrics to such content [Dansereau et al. 2013; Higa et al. 2013], their performance has never been systematically evaluated in this context. One of the challenges is acquiring dense light-field data to validate a metric. Wide baselines

as in multi-camera rigs [Wilburn et al. 2005] need to be considered, and the reference light fields should be sufficiently dense to avoid uncontrolled visual artifacts. Obtaining human responses for light-field distortions is also difficult due to current display limitations. In this chapter (published as [Adhikarla et al. 2017]), we attempt to overcome these problems by first building a new dense light-field dataset which is suitable for testing quality metrics, and second, using a custom light-field viewing setup to obtain the quality judgments for this dataset. The collected subjective scores are used to evaluate the performance of existing metrics in the context of dense light fields.

We focus on light-field-specific angular effects akin to motion parallax, complex surface appearance, and binocular vision that arise in free viewing experience. To capture a rich variability over these effects and make quality scaling in our perceptual experiments tractable, we design (Section 6.1) fourteen real and synthetic scenes and introduce light-field distortions that are specific to light-field reconstruction, compression, and display. We then run a pair-wise comparison experiment (Section 6.2) over light-field pairs, and derive perceptual scaling of differences between original and distorted stimuli (Section 6.3). This allows us to investigate (Section 6.4) the suitability of a broad spectrum of existing image, video, and multiview quality metrics to predict such perceptual scaling. We also propose simple extensions of selected metrics to capture the angular aspects of light-field perception. While the original metrics are not meant for light fields, our results show that they can be used in this context, given a dense light field as the reference. We also demonstrate that the robustness of such metrics predictions drops when evaluating the quality between two distorted light fields.

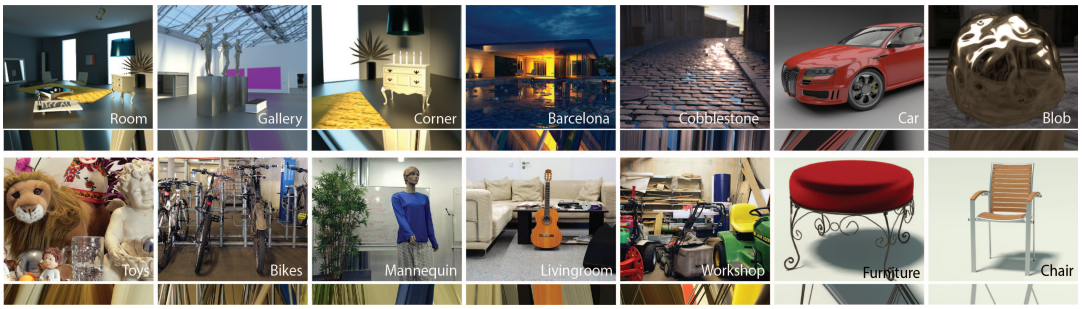
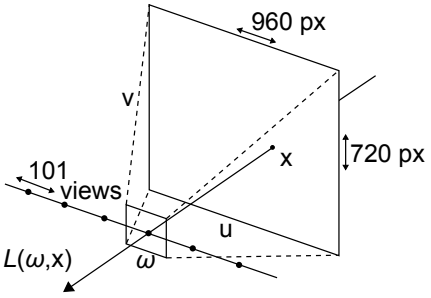


FIGURE 6.1: *Representative images of all light fields in our collection. Below each image representative epipolar-plane images (EPIs) are presented.*

6.1 Data Collection

Our dataset consist of light fields which are parameterized using two parallel planes [Levoy et al. 1996]. We consider only horizontal motion parallax that can be described using one plane and a line that is parallel to it.



More formally, we denote our light fields as $L(\omega, \mathbf{x}) \in (\mathbb{R} \times \mathbb{R}^2) \rightarrow \mathbb{R}^3$, where ω is a position on the line, and $\mathbf{x}$ is a position on the plane. We refer to them as *angular* and *spatial* coordinates, respectively. In practice, ω describes a position

of the viewer, and $\mathbf{x}$ is a coordinate of the observed image. Below, we describe the acquisition of our light fields.

6.1.1 Scenes

We designed and rendered nine synthetic and captured five real-world scenes (Figure 6.1). They span a large variety of different conditions, e.g., outdoor/indoor, daylight/night etc. They also contain objects with large range of different appearance properties. The scene objects distribution in depth is widely varied to study the artifacts resulting from disocclusions and depth discontinuities. For capturing real-world scenes, we used a one-meter long motorized linear stage with Canon EOS 5D Mark II camera and 50 mm and 28 mm lenses. After capturing all views, we performed lens distortion correction using PTLens [PTLens 2016], estimated the camera poses using Voodoo camera tracker [Viscoda 2016], and rectified all the images using the baseline drawn from the first to last camera using the approach in [Fusiello et al. 2000]. For rendered images, we used cameras with off-axis asymmetric frustums. For real-world scenes, the same effect was achieved by applying horizontal shift to the individual views. All the light fields are of identical spatial and angular resolution ($960\times720\times101$). The angular resolution was chosen high enough to avoid visible angular aliasing. This was achieved by assuring that the maximum on-screen disparities between consecutive views are around 1 pixel. To guarantee a comfortable viewing, the total disparity range during the presentation was limited to 0.2 visual degree [Shibata et al. 2011].

6.1.2 Distortions

We considered typical light-field distortions that are specific to transmission, reconstruction, and display. For each distortion, we generated multiple light fields by varying the distortion severity level. The exact levels were chosen to keep the differences between two consecutive levels small and similar. To this end, we conducted a small pilot study with 10 distortion levels, and then, selected the final levels manually.

Transmission To transmit the light-field data, an efficient data compression algorithm is highly required. We consider well-known 3D extension of HEVC encoder [Tech et al. 2016]. The light-field views are encoded into a bit stream at various quantization steps, and then, decoded back from the bit stream using the 3D-HEVC coder. We chose the following quantization steps: {25, 29, 33, 37, 41, 45}.

Reconstruction Light-field reconstruction techniques are used to recover a dense light field from sparse view samples. They interpolate the missing views using several techniques which alter the nature and appearance of the distortion. We chose the distortions resulting from linear (Linear) and nearest neighbor (NN) interpolation, as well as image warping using optical flow estimation (Opt). We also investigated the impact of using quantized depth maps (DQ). All the distortions are parametrized by the angular subsampling factor k (the distortion severity) that defines the angular resolution of the light field prior to applying the reconstruction technique. We considered $k \in \{2, 5, 8, 11, 18, 25\}$. The linear filter reconstructs dense light field by blending the reference views, and the NN method clones the closest reference view. For Opt method, we used the TV-L1 optical flow [Sánchez Pérez et al. 2013], and apply an image-warping technique [Brox et al. 2004] to synthesize in-between views. For DQ, we considered ground-truth depth map which is quantized using 8 discrete

levels. Then, we used the same image-warping technique [Brox et al. 2004] to reconstruct the light field. As this distortion requires ground-truth depth information, it is only applied to the synthetic scenes.

Display As an example of multiview autostereoscopic display artifacts, we chose a crosstalk between adjacent views, which can be modeled using a Gaussian blur in angular domain [Liu et al. 2015]. Consequently, we include such artifacts into our dataset (GAUSS). In particular, we considered the same angular subsampling parameters used in light-field reconstruction distortions and created hypothetical displays with corresponding number of views. The upsampling to higher resolution light field was achieved by using the display crosstalk model.

Four different distortions with all severity levels were applied to every scene. To all synthetic scenes we apply NN, LINEAR, OPT, and DQ. For all real-world scenes, we used NN, OPT, GAUSS, and HEVC. Including original light fields, our database consists of 350 different light fields and it is available online [MPI 2017]. The examples of resulting artifacts are presented in Figure 6.2. The whole light-field dataset is published as supplemental materials of [Adhikarla et al. 2017].

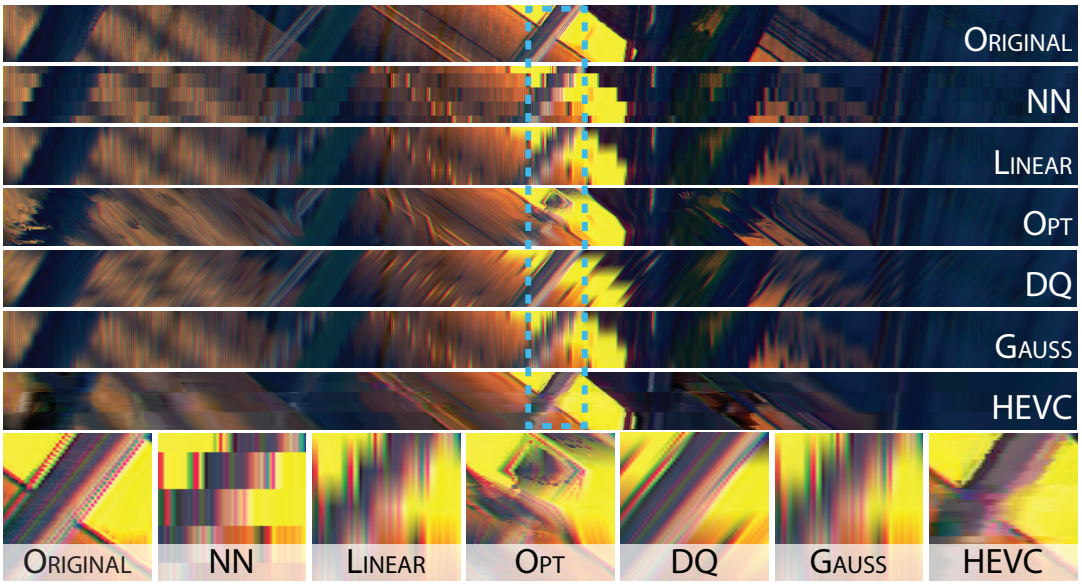


FIGURE 6.2: *Examples of distortions introduced to our light field for one of our scenes (*BARCELONA*). The images visualize central EPI of each of the distorted light fields and the enlarged portion of it is shown in the bottom row.*

6.2 Experiment

To acquire subjective quality scores that enable both training and testing different quality metrics, we performed a large-scale subjective experiment.

Equipment To simulate stereoscopic viewing with high-quality motion parallax, we used on our own setup (Figure 6.3) that consists of ASUS VG278 27″ Full HD 120 Hz LCD desktop monitor and NVIDIA 3D Vision 2 Kit for displaying stereoscopic images. Motion parallax was reproduced using a custom head tracking in which a small LED headlamp was tracked using a Logitech HD C920 Pro webcam (refer to the supplemental video). The head tracking allowed the participants to view light fields in an unconstrained manner. The viewing distance was approximately 60 cm,

and users could move their heads along a baseline of 20 cm in the direction parallel to the screen plane. The eye accommodation was fixed to the screen and did not change with eye vergence. The display was operated at the full brightness to minimize the effect of luminance on depth perception [Didyk et al. 2012].

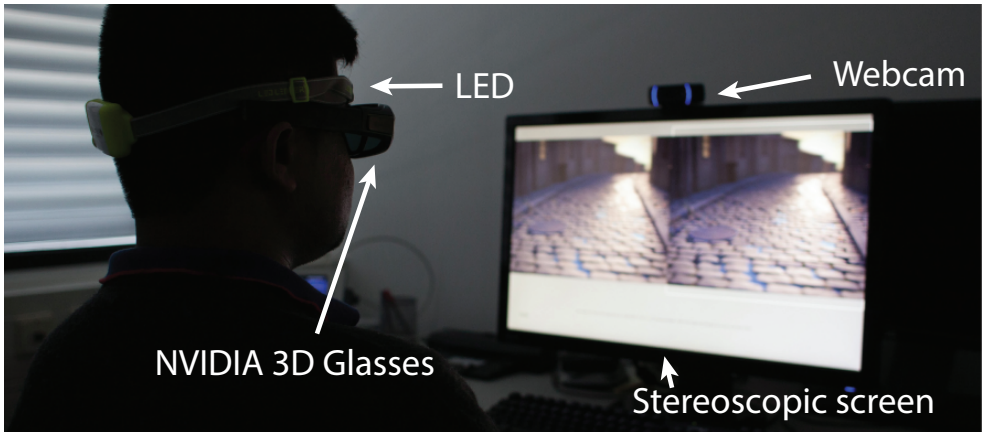


FIGURE 6.3: *Experiment session: viewer's position is tracked using a head lamp and a web-cam, a pair of NVIDIA 3D Vision 2 Kit active glasses provides stereoscopic viewing.*

Stimuli Each stimuli was a pair of light fields. As our scaling procedure used for obtaining quality scores (Section 6.3) can handle an incomplete set of comparisons and prefers when more comparisons are made for pairs of similar quality [Silverstein et al. 2001], each pair consisted of light fields with neighboring severity levels of the same distortion type. This results in 336 different stimuli which were presented stereoscopically.

Task We experimented with direct rating methods, such as ACR [ITU-T-P.910 2008], in order to measure Mean-Opinion-Scores of the distorted images. However, we found these methods to be insensitive to subtle but noticeable degradation of quality. Also participants found the direct rating task difficult. Therefore, we decided to use a more sensitive pair-wise comparison method with a two-alternative-forced-choice. In each trial, the participants were shown a pair of light fields side-by-side, and the task was to indicate the light field that a user "would prefer to see on a 3-dimensional display". Participants were given unlimited time to investigate the light fields, but they were allowed to give their response only after 80% of all perspective images were seen. The order of the light-fields pairs as well as their placement on the screen were randomized. Before each session, the participants were provided with a form summarizing the task, and a training session was conducted to familiarize participants with the experiment.

Participants Forty participants took part in the test, including both male (20) and female (20) aged 24–40 with normal or corrected-to-normal vision. Each subject performed the test in three sessions within one week. In one session, the participants saw 120–180 light-field pairs consisting of all the test conditions, but for a subset of the scenes. For a given subject, two test sessions were allowed during a single day, and these were separated by at least an hour of break.

6.3 Analysis of Subjective Data

The results of pair-wise comparison experiment are usually scaled in just-noticeable-differences (JNDs). We observed that considering measured differences as "noticeable" leads to incorrect interpretation of the experimental results. Two stimuli are 1 JND apart if 75% of observers can see the difference between them. However, our experimental question was not whether observers can tell if the light fields are different, but rather which one has higher quality. As shown in Figure 6.4, a pair of stimuli could be noticeably different from each other (JND>1), but they could appear to have the same quality. For that reason, we denote measured values as just-objectionable-differences (JODs). These units quantify the quality difference in relation to the perfect reference image. Note that the measure of JOD is more similar to visual equivalence [Ramanarayanan et al. 2007] or to the quality expressed as a difference-mean-opinion-score (DMOS) rather than to JNDs.

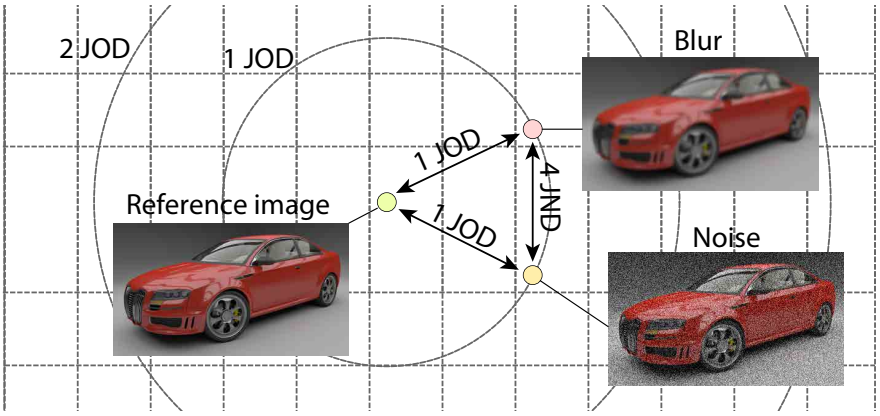


FIGURE 6.4: *Illustration of the difference between just-objectionable-differences (JODs) and just-noticeable-differences (JNDs). The image affected by blur and noise may appear to be similarly degraded in comparison to the reference image (the same JOD), but they are noticeably different and therefore several JNDs apart. The mapping between JODs and JNDs can be very complex and the relation shown in this plot using Cartesian and polar coordinates is just for illustration purposes.*

To scale the results in JOD units we used a Bayesian method based on the method of Silverstein and Farrell [Silverstein et al. 2001]. It employs a maximum-likelihood-estimator to maximize the probability that the collected data explains JOD-scaled quality scores under the Thurstone Case V assumptions [Rafał K. Mantiuk 2016]. The optimization procedure finds a quality value for each pair of light fields that maximizes the likelihood modeled by the binomial distribution. Unlike standard scaling procedures, the Bayesian approach robustly scales pairs of conditions for which there is unanimous agreement. Such pairs are common when a large number of conditions are compared. It can also scale the result of an incomplete and imbalanced pair-wise design, when not all the pairs are compared and some are compared more often. As the pair-wise comparisons provide relative quality information, the JOD values are relative. To maintain consistency across the scenes, we fix the starting point of the JOD scale at 0 for different distortions and thus the quality degradation results in negative JOD values.

The results of the subjective quality assessment experiment are shown in Figure 6.5. The error bars represent 95% confidence intervals, relative to the reference light field, computed by bootstrapping by sampling with replacement.

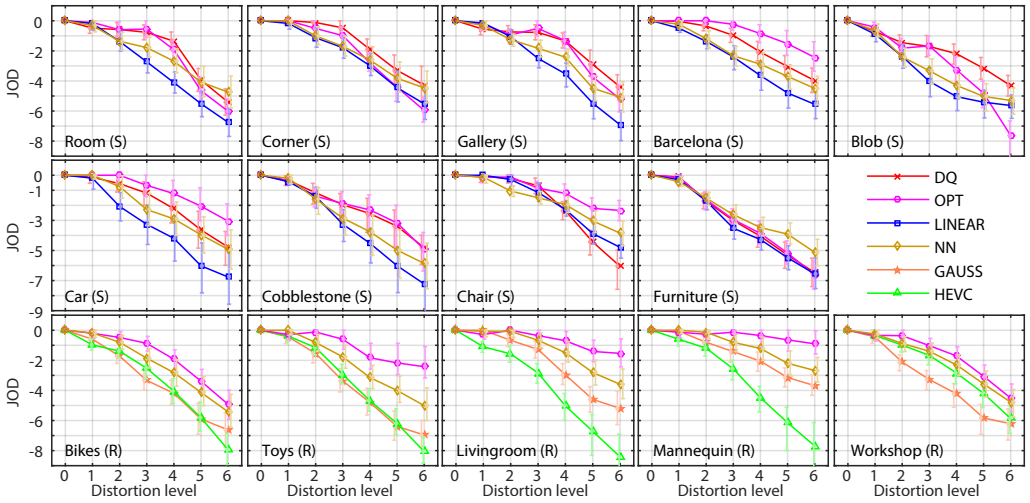


FIGURE 6.5: *The results of the subjective quality assessment experiment. The distortion level indicates the distortion severity: 0–reference 6–severest distortion level. JOD is the scaled subjective quality value. The error bars denote 95% confidence interval. The bars are horizontally displaced to avoid overlapping. The scene names are indicated in the corner of each plot. The character in parenthesis after the scene name indicates whether the scene is synthetic (S) or real-world (R).*

The results show interesting patterns in the objectionability of different distortions. OPT offers a consistent performance improvement over NN. The only exception is the *Furniture* scene featuring thin and irregularly shaped foreground objects, in which case all types of view interpolation are more objectionable than the selection of the nearest single view. The optical flow interpolation works better for real-world scenes as there are more features that can be detected. The LINEAR interpolation in most of the cases results in the worst performance, except for small distortion levels, which may indicate that visible blur due to this distortion is strongly objectionable. Similar findings have been reported by Vangorp et al. [2011] in their study on the visual performance of view interpolation methods in monocular vision. HEVC and GAUSS distortions are usually the easiest to detect as they induce significant amount of spatial distortion when compared to others. Overall the results show clearly that light-field quality is scene-dependent and successful quality metric must predict the effect of scene content on the visibility of light-field distortions.

6.4 Evaluation of Quality Metrics

We considered several popular image, video, stereo, and multiview quality metrics. We briefly describe the metrics and then show their individual performance on our dataset. For obtaining the quality of a light field using image quality metrics, we apply the metrics on individual light-field images and then average the scores over all images.

Quality metrics Although studies show that perceptual metrics perform better than an absolute difference (AD) [Lin et al. 2011], because of its significant usage in image quality assessment, we considered peak signal-to-noise ratio (PSNR). We also investigated $\text{SSIM}_{2\text{D}}$ [Wang et al. 2004], which is widely used on 2D images, and its extensions to angular domains – $\text{SSIM}_{2\text{D}\times1\text{D}}$ and $\text{SSIM}_{3\text{D}}$. $\text{SSIM}_{3\text{D}}$ computes the same statistics as standard $\text{SSIM}_{2\text{D}}$ but on 3D patches extracted from the light-field

volume. $\text{SSIM}_{2D\times1D}$ uses 2D×1D patch which contains a 2D window extracted from a particular view and a 1D row of pixels that extends from the center of the 2D window in the angular domain (see Figure 6.6). We applied the metrics to all light-field images without resampling and averaged the scores over all images. Although we experimented with various pooling strategies, we found that the average value performs best. Due to better performance, we chose the angular window sizes of 32 and 64 pixels for $\text{SSIM}_{2D\times1D}$ and SSIM_{3D} respectively. We also considered a multi-scale version of SSIM_{2D}– MS-SSIM [Wang et al. 2003] which extends SSIM_{2D} to compute differences on multiple levels. We also used GMSD [Xue et al. 2014] which provides good performance over a rich collection of image datasets. The most advanced 2D metric considered in our experiments was HDR-VDP-2 [Mantiuk et al. 2011] which stands out among perception-based quality metrics.

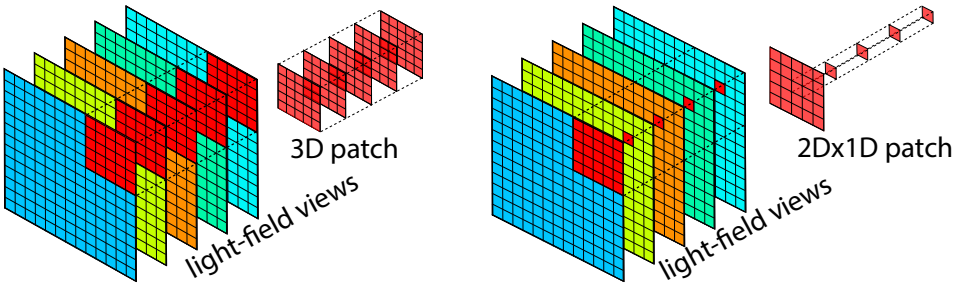


FIGURE 6.6: *Patches used in our extensions of* SSIM_{2D}.

We further considered the NTIA General Model – VQM [Pinson et al. 2004] which was standardized for video-signals evaluation (ANSI T1.801.03-2003). For this metric, light-field images are input in a form of video panning from the leftmost view to the rightmost view and back. We also chose the stereoscopic image quality metric – SIQM [Chen et al. 2013b] that is based on the concept of cyclopean image where, we averaged scores obtained from all stereo pairs shown in our experiment. To capture the full range of stereo quality metrics, we also included a stereoscopic video quality metric StSD_{LC} [Silva et al. 2013].

Finally, we chose metrics that address multiview data and account for interpolation artifacts. 3DSwIM proposed by Battisti et al. [2015] first shift-compensates blocks from the reference and distorted (interpolated) images. These matched blocks undergo the first level of Haar wavelet transform and histogram of the sub-band corresponding to horizontal details in the block is computed. Finally, the Kolmogorov-Smirnov distance of these histograms is taken as the metric prediction. Another metric for the multiview video is MP-PSNR [Sandic-Stankovic et al. 2016]. It computes the multi-resolution morphological pyramid decomposition on the reference and test images. Detail images of the top levels of these pyramids are then compared through the mean squared error. The resulting per pixel errors maps are then pooled and converted to a peak signal-to-noise ratio measure.

6.4.1 Metric Performance Comparison

The quality values predicted by each metric are expected to be related to JOD values, but this relation can be complex and non-linear. To account for this relation, we follow a common practice and fit a logistic function:

$$q(o) = a_1 \left\{ \frac{1}{2} - \frac{1}{1 + \exp\left[a_2\left(o - a_3\right)\right]} \right\} + a_4 o + a_5 \tag{6.1}$$

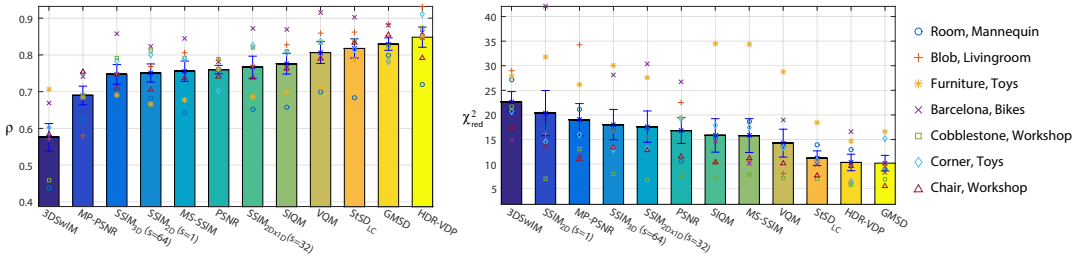


FIGURE 6.7: *The goodness-of-fit scores for the metrics expressed as Pearson Correlation Coefficient (ρ) and reduced chi-square (χ^2_{red}) after cross-validation. The results for each cross-validation fold are shown. $\chi^2_{red} = 1$ indicates that the goodness of fit between the metric predictions and the subjective data is in perfect agreement with the measured subjective variance and $\rho = 1$ indicates perfect positive linear relation between objective scores and JODs. The error bars represent standard error.*

where o is the output of a metric. The parameters $a_{1..5}$ are optimized to minimize a given goodness-of-fit measure. We computed several such measures, such as Spearman rank-order correlation, or MSE, which can be found in the supplementary materials. Here we report the reduced chi-squared statistic (χ^2_{red}) and Pearson correlation coefficient (ρ). χ^2_{red} is computed as a weighted average of the squared differences, in which weights are the inverse of sample variance. This accounts for the fact that larger JOD values are more uncertain (refer to Figure 6.5), and therefore, the accuracy of their prediction can be lower. For a fair comparison, we employed a seven fold cross-validation across different scenes. We measured the goodness-of-fit on two randomly chosen scenes in a cross-validation fold and averaged the results over all folds. The resulting Pearson correlation and χ^2_{red} values are shown in Figure 6.7. The performance of the metrics on individual distortions are shown in Figure 6.8. A more elaborate analysis including the evaluation on real-world and synthetic scenes separately is presented in the supplementary materials of [Adhikarla et al. 2017].

The results show good performance of 2D image and video quality metrics. This is unexpected as our dataset was meant to emphasize visibility of angular artefacts, which are not directly considered by these metrics. We observed, however, that angular distortions indirectly translate into the differences in spatial patterns, which could explain the good performance. We hypothesize that relatively better performance of HDR-VDP-2 and GMSD is achieved by detecting changes in contrast across multiple scales, which in case of HDR-VDP-2 is additionally backed by perceptual scaling of distortions and discarding of those that are invisible. A comparable performance of video (VQM) and stereoscopic (StSD_{LC}) metrics can be explained by their emphasis on the relation between neighboring views, which in some way captures angular aspects of light-field viewing. Figure 6.8 shows that some metrics are better at predicting some distortion types than the others. For example, HDR-VDP-2 consistently under-predicts quality for HEVC. Training such metrics for a particular distortion type could substantially boost their performance. Unexpectedly, our *ad hoc* attempts to extend the SSIM_{2D} metric by adding the angular dimension (SSIM_{3D}) or right away considering 3D patches ($\text{SSIM}_{\text{2D}\times\text{1D}}$) that should account for angular changes has led to significantly worse results. Clearly, there is a room for improvements and a suitable dataset, such as the one provided in this work, should help to develop a better metric in future.

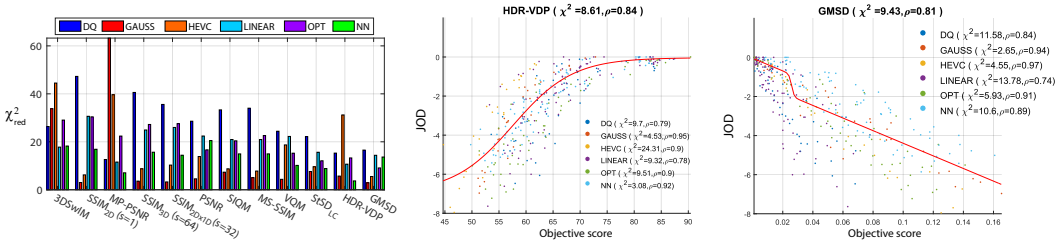


FIGURE 6.8: *Left: The prediction accuracy per-distortion reported as reduced chi-squared goodness of fit score. Middle and right: χ^2_{red}–fit for the metrics HDR-VDP and GMSD over all scenes. The prediction accuracy for individual distortions are shown inside the plots and the overall accuracy is indicated on the top of the plots.*

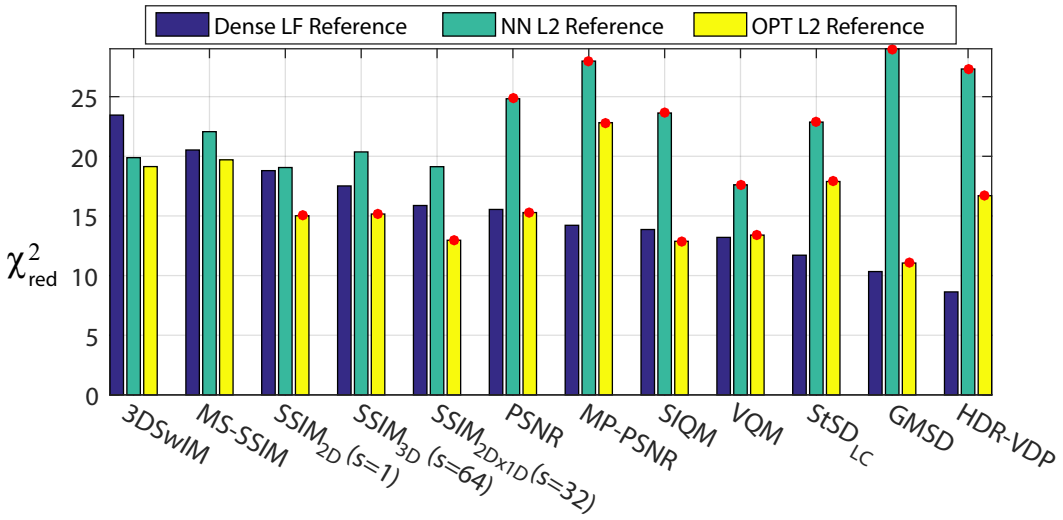


FIGURE 6.9: *The goodness-of-fit scores for the subset of the dataset when a dense LF is used as a reference (blue), when nearest-neighbour at the 2^{nd} distortion level is a reference (cyan), or when optical flow is used to up-sample the reference LFs. The dots at cyan bars mean that the value is statistically different from the dense LF case and the dots on the yellow bars that the values are statistically different from the* NN *case. The significance is computed by bootstrapping and running one-tailed test ($p = 0.05$).*

6.4.2 Sparse Light-field Reference Case

In all our tests, we provided a high quality, 101-view light field as a reference for the quality metrics. In practice, in most applications only sparsely sampled light field is available. When a sparse light field is used as a reference, a full-reference metric is given to compare two distorted light fields without a perfect reference. This is a task that such metrics were not designed for as they are intended to predict JODs relative to the perfect reference image, not JNDs relative to any other image (refer to Figure 6.4). This issue is potentially shared with other quality assessment tasks, for example when a metric is trained on 4K images, but it is used on much lower resolution images. However, this problem is exacerbated in case of light fields, where the reduction of angular resolution is often substantial.

To test whether the metrics can predict the quality of distorted light fields using sparse light fields as a reference, we measured the performance of the metrics on a subset of our dataset. As a reference, we chose light fields with distortion NN and severity level two, which correspond to original light fields subsampled to 21 angular views. For the testing light field, we considered all light fields with a higher distortion levels. For a fair comparison, we also ran the metrics on the same subset using full

101-view light fields as a reference. The results of these tests are shown as cyan and blue bars in Figure 6.9. The significant difference in goodness-of-fit scores (marked with dots) show that metrics predictions get worse if imperfect (sparse) reference is used. This suggests that the existing metrics must be provided with a high-quality reference light field to predict reliably the quality.

But if such high quality reference is not available, can it be approximated? Our subjective data shows that optical-flow interpolation (Opt) produces the highest quality results. Therefore, we used Opt to produce reference 101-view light fields from sparse 21-view light fields and reran the metrics on the subset. The results indicate that the predictions improved as compared to using sparse light field (yellow vs. cyan bars in Figure 6.9). This suggests that a potential solution to the problem of imperfect reference is to use high-quality interpolation method in order to generate reference.

6.5 Conclusion

We have established a new 3D dense light-field dataset together with the subjective quality scaling for various distortions that occur in light-field applications. Different methods in light-field processing lead to visual artifacts with quite different appearance, e.g., blur for Linear, ghosting for Opt, image flickering and jumping for NN. Our experiments reveal how these different artifacts affect perceived quality. Our subjective scores are derived from an interactive 3D light-field viewing setup and correspond precisely to overall quality of light fields rather than individual views. We have evaluated the potential of existing image, video, stereo, and multiview quality metrics in predicting the subjective scores. Our observations show that the metrics – HDR-VDP-2, GMSD, StSD_{LC} and VQM perform reasonably well when comparing a distorted light field to a dense reference, and can be used in applications requiring such comparisons. When dense light field is not available, which is the case in some applications, the usage of these metrics for quality assessment is not justified. The perceptually scaled data that we provide can be used for training and validating new light-field quality metrics. Of practical interest for such development is the problem identified in this work, where incomplete, sparse light fields must serve as the reference. Our results also reveal the quality of different light-field reconstruction method, which can directly guide the choice of the light-field reconstruction technique. In the current work, we did not consider aspects such as masking properties of the human visual system. It could be interesting to investigate how much the metrics gain by considering this effect rather than simple averaging of scores over all views. When creating our dataset, we did not consider focus cue. We are, however, not aware of any display setup that could be used to evaluate both motion parallax and focus cue quality. We also believe that the problems revealed in this work should be addressed before including additional cues.

Chapter 7

Conclusions and Future Work

In this thesis, we addressed several problems in the context of high-fidelity material-appearance reproduction.

The main part of the thesis (published as [Elek et al. 2017; 2021; Rittig et al. 2021; Sumin et al. 2019]) is devoted to the development of an end-to-end print-preparation system tailored for multi-material inkjet 3D printers, recognized as the most advanced devices for full-color 3D printing to date. We started from a practical problem: the inability to achieve sharp textures on 3D-printed objects using existing inherently translucent multi-color materials. We introduced an iterative optimization scheme designed to iteratively refine the volumetric arrangement of print materials, resulting in printouts closely resembling the user-specified 3D model. Implementing such an approach necessitated the definition of a discrepancy metric between the target model and intermediate optimization results. This, in turn, required establishing a practical material measurement method and developing an appearance prediction model. The proposed method can be categorized as inverse rendering as we search for such a scattering volume that, being rendered or fabricated, exhibits a certain surface appearance.

In the initial stages of our pipeline development, we adhered to accurate, albeit computationally expensive, physically-based Monte-Carlo estimation of light transport. At each step, we carefully validated that our virtual environment accurately reproduces the appearance of material-arrangement voxel grids that can then be sent to a printer. We then succeeded in replacing the accurate but slow Monte-Carlo-based appearance prediction with a fast predictor based on a trained neural network. This improvement enabled full heterogeneous material optimization for 3D-print preparation on a regular desktop computer within time frames comparable to the actual printing time.

A significant portion of our print-preparation system is data-driven, ensuring its applicability to any inkjet 3D printer and any set of print materials. Our approach involves collecting all necessary data independently, without relying on information from printer manufacturers. We extract print-material parameters from small samples of primary materials. Once these materials are characterized, our system is ready to prepare user-specified 3D objects for printing through an iterative method consisting of prediction and refinement steps. The proposed neural acceleration of light-transport computation is trained solely on synthetic data, demonstrating its ability to generalize to material parameters not encountered during the training phase.

The proposed 3D-printing pipeline effectively addresses hardware limitations associated with print materials from a software perspective. In certain scenarios, achieving the desired appearance may be challenging or even impossible. For instance, when dealing with thin geometries (below 2 mm in our experiment), material translucency significantly reduces the gamut of achievable colors, and saturated colors may not be reproducible in such thin areas. Moreover, the limited amount of material may

not suffice to achieve color separation on opposing surfaces of the geometry. In our work, we demonstrate that it is possible to account for such limitations and plausibly reproduce the target appearance.

In situations where perfect reproduction may not be feasible due to physical limitations, an optimization approach combined with a visually meaningful discrepancy metric allows for finding a solution as close to the user specification as possible. When working on the limits of what is physically reproducible, it is often beneficial to use perceptually uniform quality metrics, as this ensures the optimization result is closer to the user's expectations. We utilize the perceptually uniform CIELAB color space in our content-aware gamut mapping because it is the part of our system where perfect reproduction might not be possible.

We see significant potential in the recent advances in differentiable volume rendering. This field has developed actively alongside our work, and our inverse-rendering formulation for 3D-printing preparation aligns with these advancements. The recent work by Nindel et al. [2021] demonstrates the applicability of differentiable rendering to full-volume optimization of print materials and further extends our 3D-print preparation pipeline. The main advantage of the differentiable rendering approach lies in its generality: the error metric (or loss function) ultimately determines the path of the gradient-descent optimization process.

A prominent future direction for appearance 3D printing research is the fabrication of angularly- and spatially-varying appearances. Light field representation is well-suited for describing such visual features. To support future work in this area, we evaluate existing quality metrics for light field images [Adhikarla et al. 2017] and focus on their perceptual uniformity. We propose a new dense light-field dataset and derive subjective quality scaling through a perceptual experiment for various distortions that occur in light-field applications. This dataset, along with the corresponding perceptually scaled metric values, can be used for training and validating new light-field quality metrics, which would enable the high-fidelity fabrication of angularly- and spatially-varying appearances.

In general, we believe that the field of appearance fabrication is in its infancy, with many new works and developments yet to come. Fabrication of angularly- and spatially-varying appearances is just one promising direction for future research. New findings in the field of human perception of fabricated objects, implemented as error metrics, can be integrated into our or similar optimization-based 3D-print preparation systems. Future efforts could also explore the integration of dynamic and interactive elements within printed materials, paving the way for applications in various industries, including art, medical imaging, and augmented reality. Our work lays a strong foundation for these future explorations, providing a robust framework for addressing the complex challenges associated with high-fidelity material-appearance reproduction.

Bibliography - Own Work

Adhikarla, Vamsi Kiran, Marek Vinkler, Denis Sumin, Rafał Mantiuk, Karol Myszkowski, Hans-Peter Seidel, and Piotr Didyk (2017). "Towards a Quality Metric for Dense Light Fields". In: *Proceedings of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)* (cit. on pp. 6, 7, 86, 88, 93, 98).

Elek, Oskar, Denis Sumin, Ran Zhang, Tim Weyrich, Karol Myszkowski, Bernd Bickel, Alexander Wilkie, and Jaroslav Křivánek (Nov. 2017). "Scattering-aware Texture Reproduction for 3D Printing". In: *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)* 36.6, 241:1–241:15. ISSN: 0730-0301 (cit. on pp. 5, 7, 21, 97).

Elek, Oskar, Ran Zhang, Denis Sumin, Karol Myszkowski, Bernd Bickel, Alexander Wilkie, Jaroslav Křivánek, and Tim Weyrich (Feb. 2021). "Robust and Practical Measurement of Volume Transport Parameters in Solid Photo-polymer Materials for 3D Printing". In: *Optics Express* 29.5, pp. 7568–7588. DOI: 10.1364/OE.406095 (cit. on pp. 5, 7, 21, 97).

Rittig, Tobias, Denis Sumin, Vahid Babaei, Piotr Didyk, Alexey Voloboy, Alexander Wilkie, Bernd Bickel, Karol Myszkowski, Tim Weyrich, and Jaroslav Křivánek (May 2021). "Neural Acceleration of Scattering-Aware Color 3D Printing". In: *Computer Graphics Forum (Proc. Eurographics)* 40.2, pp. 205–219. ISSN: 1467-8659. DOI: 10.1111/cgf.142626 (cit. on pp. 6, 7, 66, 75, 81, 97).

Sumin, Denis, Tobias Rittig, Vahid Babaei, Thomas Nindel, Alexander Wilkie, Piotr Didyk, Bernd Bickel, Jaroslav Křivánek, Karol Myszkowski, and Tim Weyrich (July 2019). "Geometry-Aware Scattering Compensation for 3D Printing". In: *ACM Transactions on Graphics (Proc. SIGGRAPH)* 38.4. ISSN: 0730-0301 (cit. on pp. 6, 7, 45, 97).

Bibliography

Abadi, Martín, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng (2015). *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*. Software available from tensorflow.org. URL: https://www.tensorflow.org/ (cit. on p. 72).

Abu Rmaileh, Lubna and Alan Brunton (2023). "Meso-Facets for Goniochromatic 3D Printing". In: *ACM Trans. Graph.* 42.4. ISSN: 0730-0301. DOI: 10.1145/3592137. URL: https://doi.org/10.1145/3592137 (cit. on p. 12).

Ansari, Navid, Omid Alizadeh-Mousavi, Hans-Peter Seidel, and Vahid Babaei (2020). "Mixed integer ink selection for spectral reproduction". In: *ACM Trans. Graph.* 39.6. ISSN: 0730-0301. DOI: 10.1145/3414685.3417761. URL: https://doi.org/10.1145/3414685.3417761 (cit. on p. 11).

Ansari, Navid, Hans-Peter Seidel, and Vahid Babaei (2022). "Mixed integer neural inverse design". In: *ACM Trans. Graph.* 41.4. ISSN: 0730-0301. DOI: 10.1145/3528223.3530083. URL: https://doi.org/10.1145/3528223.3530083 (cit. on p. 11).

Aydin, Tunç Ozan, Martin Čadík, Karol Myszkowski, and Hans-Peter Seidel (2010). "Video quality assessment for computer graphics applications". In: *ACM Trans. Graph.* 29.6. ISSN: 0730-0301. DOI: 10.1145/1882261.1866187. URL: https://doi.org/10.1145/1882261.1866187 (cit. on p. 85).

Baar, Jeroen van, Steven Poulakos, Wojciech Jarosz, Derek Nowrouzezahrai, Rasmus Tamstorf, and Markus Gross (2011). "Perceptually-based compensation of light pollution in display systems". In: *Proc. of ACM Symposium on Applied Perception in Graphics and Visualization*. DOI: 10.1145/2077451.2077460 (cit. on p. 13).

Babaei, Vahid and Roger D. Hersch (2016). "N-ink printer characterization with barycentric subdivision". In: *IEEE Transactions on Image Processing* 25.7 (cit. on p. 29).

Babaei, Vahid, Kiril Vidimče, Michael Foshey, Alexandre Kaspar, Piotr Didyk, and Wojciech Matusik (2017). "Color contoning for 3D printing". In: *ACM Transactions on Graphics (Proc. SIGGRAPH)* 36 (to appear) (cit. on pp. 10, 11, 13, 28, 33, 40, 46, 56).

Bako, Steve, Mark Meyer, Tony DeRose, and Pradeep Sen (2019). "Offline Deep Importance Sampling for Monte Carlo Path Tracing". en. In: *Computer Graphics Forum* 38.7. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.13858, pp. 527–542. ISSN: 1467-8659. DOI: 10.1111/cgf.13858. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.13858 (visited on 04/30/2020) (cit. on p. 15).

Bako, Steve, Thijs Vogels, Brian Mcwilliams, Mark Meyer, Jan NováK, Alex Harvill, Pradeep Sen, Tony Derose, and Fabrice Rousselle (July 2017). "Kernel-predicting convolutional networks for denoising Monte Carlo renderings". In: *ACM Transactions on Graphics* 36.4, 97:1–97:14. ISSN: 0730-0301. DOI: 10.1145/3072959.3073708. URL: https://doi.org/10.1145/3072959.3073708 (visited on 05/15/2020) (cit. on p. 15).

Battisti, Federica, Emilie Bosc, Marco Carli, Patrick Le Callet, and Simone Perugia (2015). "Objective image quality assessment of 3D synthesized views". In: *Signal Processing: Image Communication* 30.C, pp. 78–88. ISSN: 0923-5965. DOI: http://dx.doi.org/10.1016/j.image.2014.10.005. URL: http://www.sciencedirect.com/science/article/pii/S0923596514001453 (cit. on pp. 18, 92).

Bosc, Emilie, Federica Battisti, Marco Carli, and Patrick Le Callet (2013). "A wavelet-based image quality metric for the assessment of 3D synthesized views". In: *Proc. SPIE*. Vol. 8648, 86481Z–86481Z–9. DOI: 10.1117/12.2002410. URL: http://dx.doi.org/10.1117/12.2002410 (cit. on p. 18).

Bosc, Emilie, Romuald Pépion, Patrick Le Callet, Martin Köppel, Patrick Ndjiki-Nya, Muriel Pressigout, and Luce Morin (Nov. 2011). "Towards a New Quality Metric for 3-D Synthesized View Assessment". In: *IEEE Journal on Selected Topics in Signal Processing*, J–STSP–ETVC–00048–2011. URL: https://hal.archives-ouvertes.fr/hal-00628070 (cit. on p. 18).

Branch, Mary Ann, Thomas F Coleman, and Yuying Li (1999). "A subspace, interior, and conjugate gradient method for large-scale bound-constrained minimization problems". In: *SIAM Journal on Scientific Computing* 21.1, pp. 1–23 (cit. on p. 47).

Brox, Thomas, Andrés Bruhn, Nils Papenberg, and Joachim Weickert (2004). "High accuracy optical flow estimation based on a theory for warping". In: *Comput. Vis.-ECCV 2004*. Springer, pp. 25–36 (cit. on pp. 87, 88).

Brunton, Alan, Can Ates Arikan, Tejas Madan Tanksale, and Philipp Urban (July 2018). "3D Printing Spatially Varying Color and Translucency". In: *ACM Transactions on Graphics* 37.4, 157:1–157:13. ISSN: 0730-0301 (cit. on p. 11).

Brunton, Alan, Can Ates Arikan, and Philipp Urban (Dec. 2015). "Pushing the limits of 3D color printing: error diffusion with translucent materials". In: *ACM Transactions on Graphics* 35.1, 4:1–4:13 (cit. on pp. 10, 11, 21, 23, 28, 29, 33, 34, 42, 46, 51, 60, 61).

Burley, Brent, David Adler, Matt Jen-Yuan Chiang, Hank Driskill, Ralf Habel, Patrick Kelly, Peter Kutz, Yining Karl Li, and Daniel Teece (July 2018). "The Design and Evolution of Disney's Hyperion Renderer". In: *ACM Transactions on Graphics* 37.3, 33:1–33:22. ISSN: 0730-0301 (cit. on p. 14).

Cai, Wenli, Tianzhou Chen, and Jiaoying Shi (1996). "Detail enhancement in volume rendering". In: *Proc. of SPIE*. Vol. 2644, pp. 267–72. DOI: 10.1117/12.235529 (cit. on p. 13).

Chaitanya, Chakravarty R. Alla, Anton S. Kaplanyan, Christoph Schied, Marco Salvi, Aaron Lefohn, Derek Nowrouzezahrai, and Timo Aila (July 2017). "Interactive reconstruction of Monte Carlo image sequences using a recurrent denoising autoencoder". In: *ACM Transactions on Graphics* 36.4, 98:1–98:12. ISSN: 0730-0301. DOI: 10.1145/3072959.3073601. URL: https://doi.org/10.1145/3072959.3073601 (visited on 05/15/2020) (cit. on p. 15).

Che, Chengqian, Fujun Luan, Shuang Zhao, Kavita Bala, and Ioannis Gkioulekas (2020). "Towards Learning-based Inverse Subsurface Scattering". en. In: *IEEE*

International Conf. on Computational Photography (ICCP). IEEE, pp. 1–12 (cit. on p. 14).

Chen, Danwu and Philipp Urban (2021). "Deep learning models for optically characterizing 3D printers". In: *Opt. Express* 29.2, pp. 615–631. DOI: 10.1364/OE.410796. URL: https://opg.optica.org/oe/abstract.cfm?URI=oe-29-2-615 (cit. on p. 16).

— (2022). "Inducing robustness and plausibility in deep learning optical 3D printer models". In: *Opt. Express* 30.11, pp. 18119–18133. DOI: 10.1364/OE.455115. URL: https://opg.optica.org/oe/abstract.cfm?URI=oe-30-11-18119 (cit. on p. 16).

Chen, Desai, David I. W. Levin, Piotr Didyk, Pitchaya Sitthi-Amorn, and Wojciech Matusik (2013a). "Spec2Fab: A reducer-tuner model for translating specifications to 3D prints". In: *ACM Transactions on Graphics (Proc. SIGGRAPH)* 32.4, 135:1–135:10 (cit. on p. 10).

Chen, Ming-Jun, Che-Chun Su, Do-Kyoung Kwon, Lawrence K. Cormack, and Alan C. Bovik (2013b). "Full-reference quality assessment of stereopairs accounting for rivalry". In: *Signal Processing: Image Communication* 28.9, pp. 1143–1155. ISSN: 0923-5965. DOI: http://dx.doi.org/10.1016/j.image.2013.05.006. URL: http://www.sciencedirect.com/science/article/pii/S0923596513000787 (cit. on p. 92).

Chermain, Xavier, Cédric Zanni, Jonàs Martínez, Pierre-Alexandre Hugron, and Sylvain Lefebvre (2023). "Orientable Dense Cyclic Infill for Anisotropic Appearance Fabrication". In: *ACM Trans. Graph.* 42.4. ISSN: 0730-0301. DOI: 10.1145/3592412. URL: https://doi.org/10.1145/3592412 (cit. on p. 12).

Christensen, Per H. (2015). "An approximate reflectance profile for efficient subsurface scattering". In: *ACM SIGGRAPH Talks*. Los Angeles, California: ACM. ISBN: 978-1-4503-3636-9 (cit. on pp. 12, 14).

Cignoni, Paolo, Enrico Gobbetti, Ruggero Pintus, and Roberto Scopigno (2008). "Color enhancement for rapid prototyping". In: *Proc. of International Symposium on Virtual Reality, Archaeology and Cultural Heritage*. Eurographics, pp. 9–16 (cit. on pp. 11, 13).

Computer Graphics Laboratory, Stanford University (2008). *The (New) Stanford Light Field Archive*. http://lightfield.stanford.edu/acq.html. Accessed: 2016-04-23 (cit. on p. 17).

Condor, Jorge, Michal Piovarci, Bernd Bickel, and Piotr Didyk (2023). "Gloss-Aware Color Correction for 3D Printing". In: *ACM SIGGRAPH 2023 Conference Proceedings*. SIGGRAPH '23. Los Angeles, CA, USA: Association for Computing Machinery. ISBN: 9798400701597. DOI: 10.1145/3588432.3591546. URL: https://doi.org/10.1145/3588432.3591546 (cit. on p. 12).

Correia, António, Peter Hanselaer, Hugo Cornelissen, and Youri Meuret (2017). "Radiance based method for accurate determination of volume scattering parameters using GPU-accelerated Monte Carlo". In: *Optics Express* 25.19, pp. 22575–86. DOI: 10.1364/OE.25.022575 (cit. on p. 16).

Dansereau, Donald G, Daniel L Bongiorno, Oscar Pizarro, and Stefan B Williams (2013). "Light field image denoising using a linear 4D frequency-hyperfan all-in-focus filter". In: *Proceedings of the SPIE Conference on Computational Imaging (SPIE'13)*. Vol. 8657 (cit. on p. 85).

d'Eon, Eugene and Geoffrey Irving (July 2011). "A Quantized-diffusion Model for Rendering Translucent Materials". In: *ACM Transactions on Graphics (Proc. SIGGRAPH)* 30.4, 56:1–56:14. ISSN: 0730-0301. DOI: 10.1145/2010324.1964951 (cit. on pp. 12, 14, 16).

Derr, V. E. (1980). "Estimation of the extinction coefficient of clouds from multi-wavelength lidar backscatter measurements". In: *Appl. Opt.* 19.14, pp. 2310–14. DOI: 10.1364/AO.19.002310 (cit. on p. 17).

Didyk, Piotr, Tobias Ritschel, Elmar Eisemann, Karol Myszkowski, Hans-Peter Seidel, and Wojciech Matusik (2012). "A luminance-contrast-aware disparity model and applications". In: *ACM Trans. Graph.* 31.6, p. 184 (cit. on pp. 19, 89).

Dong, Yue, Jiaping Wang, Fabio Pellacini, Xin Tong, and Baining Guo (2010). "Fabricating spatially-varying subsurface scattering". In: *ACM Transactions on Graphics (Proc. SIGGRAPH)* 29.4, 62:1–62:10. ISSN: 0730-0301. DOI: 10.1145/1778765.1778799 (cit. on pp. 10, 11, 16, 21, 23, 34).

Donner, Craig, Tim Weyrich, Eugene d'Eon, Ravi Ramamoorthi, and Szymon Rusinkiewicz (2008). "A layered, heterogeneous reflectance model for acquiring and rendering human skin". In: *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)* 27 (5), 140:1–140:12. ISSN: 0730-0301. DOI: http://doi.acm.org/10.1145/1409060.1409093 (cit. on pp. 12, 14, 17, 26).

Doubrovski, E. L., Elizabeth Yinling Tsai, Daniel Dikovsky, Jo M. P. Geraedts, Hugh Herr, and Neri Oxman (2015). "Voxel-based fabrication through material property mapping: A design method for bitmap printing". In: *Computer-Aided Design* 60. ISSN: 0010-4485 (cit. on p. 10).

Du, Song-Pei, Piotr Didyk, Frédo Durand, Shi-Min Hu, and Wojciech Matusik (2014). "Improving visual quality of view transitions in automultiscopic displays". In: *ACM Trans. Graph.* 33.6, p. 192 (cit. on p. 19).

Fascione, Luca, Johannes Hanika, Rob Pieké, Ryusuke Villemin, Christophe Hery, Manuel Gamito, Luke Emrose, and André Mazzone (2018). "Path tracing in production". In: *ACM SIGGRAPH 2018 Courses*. SIGGRAPH '18. Vancouver, British Columbia, Canada: Association for Computing Machinery. ISBN: 9781450358095. DOI: 10.1145/3214834.3214864. URL: https://doi.org/10.1145/3214834.3214864 (cit. on p. 14).

Fleming, Roland W. and Heinrich H. Bülthoff (2005). "Low-level image cues in the perception of translucent materials". In: *ACM Trans. Appl. Percept.* 2.3. ISSN: 1544-3558 (cit. on p. 43).

Fleming, Roland W., Henrik Wann Jensen, and Heinrich H Bülthoff (2004). "Perceiving translucent materials". In: *Proc. of ACM Symposium on Applied Perception in Graphics and Visualization*. Los Angeles, California, USA. ISBN: 1-58113-914-4 (cit. on p. 43).

Floyd, Robert W. and Louis Steinberg (1976). "An Adaptive Algorithm for Spatial Greyscale". In: *Proceedings of Society for Information Displays*. Vol. 17. 2, pp. 75–77 (cit. on pp. 10, 51).

Frisvad, Jeppe Revall, Niels Jørgen Christensen, and Henrik Wann Jensen (2007). "Computing the scattering properties of participating media using Lorenz-Mie theory". In: *ACM Trans. Graph.* 26, p. 60. DOI: http://doi.acm.org/10.1145/1276377.1276452 (cit. on p. 17).

Fusiello, Andrea, Emanuele Trucco, and Alessandro Verri (July 2000). "A Compact Algorithm for Rectification of Stereo Pairs". In: *Mach. Vision Appl.* 12.1, pp. 16–22. ISSN: 0932-8092. DOI: 10.1007/s001380050003. URL: http://dx.doi.org/10.1007/s001380050003 (cit. on p. 87).

Gharbi, Michaël, Tzu-Mao Li, Miika Aittala, Jaakko Lehtinen, and Frédo Durand (July 2019). "Sample-based Monte Carlo denoising using a kernel-splatting network". In: *ACM Transactions on Graphics* 38.4, 125:1–125:12. ISSN: 0730-0301. DOI: 10.1145/3306346.3322954. URL: https://doi.org/10.1145/3306346.3322954 (visited on 05/15/2020) (cit. on p. 15).

Gkioulekas, Ioannis, Anat Levin, and Todd E. Zickler (2016). "An Evaluation of Computational Imaging Techniques for Heterogeneous Inverse Scattering". In: *Proceedings of (ECCV) European Conference on Computer Vision*, pp. 685 –701 (cit. on p. 14).

Gkioulekas, Ioannis, Bei Xiao, Shuang Zhao, Edward Adelson, Todd Zickler, and Kavita Bala (2013a). "Understanding the role of phase function in translucent appearance". In: *ACM Transactions on Graphics* 32.5, 147:1–147:19 (cit. on p. 43).

Gkioulekas, Ioannis, Shuang Zhao, Kavita Bala, Todd Zickler, and Anat Levin (2013b). "Inverse volume rendering with material dictionaries". In: *ACM Transactions on Graphics* 32.6, 162:1–162:13. DOI: 10.1145/2508363.2508377 (cit. on pp. 14, 17, 24, 25).

Gonzalez, Rafael C. and Richard E. Woods (2007). *Digital image processing*. 3rd. Pearson (cit. on p. 13).

Habel, Ralf, Per H. Christensen, and Wojciech Jarosz (2013). "Photon beam diffusion: A hybrid Monte Carlo method for subsurface scattering". In: *Proc. of EGSR*, pp. 27–37 (cit. on p. 30).

Hašan, Miloš, Martin Fuchs, Wojciech Matusik, Hanspeter Pfister, and Szymon Rusinkiewicz (2010). "Physical reproduction of materials with specified subsurface scattering". In: *ACM Transactions on Graphics (Proc. SIGGRAPH)* 29.3, 61:1–61:10 (cit. on pp. 10, 11, 16, 21, 23, 34).

Hašan, Miloš and Ravi Ramamoorthi (2013). "Interactive albedo editing in path-traced volumetric materials". In: *ACM Transactions on Graphics* 32.2, 11:1–11:11 (cit. on p. 14).

Herholz, Philipp, Sebastian Koch, Tamy Boubekeur, and Marc Alexa (2017). "Unsharp masking geometry improves 3D prints". In: *Computers & Graphics* 66, pp. 135–142 (cit. on p. 13).

Herholz, Sebastian, Yangyang Zhao, Oskar Elek, Derek Nowrouzezahrai, Hendrik P. A. Lensch, and Jaroslav Křivánek (June 2019). "Volume Path Guiding Based on Zero-Variance Random Walk Theory". In: *ACM Transactions on Graphics* 38.3, 25:1–25:19. ISSN: 0730-0301. DOI: 10.1145/3230635. URL: http://doi.acm.org/10.1145/3230635 (visited on 11/06/2019) (cit. on p. 67).

Hersch, Roger David and Frédérique Crété (2005). "Improving the Yule-Nielsen modified Neugebauer model by dot surface coverages depending on the ink superposition conditions". In: *Electronic Imaging 2005*. International Society for Optics and Photonics, pp. 434–447 (cit. on p. 13).

Higa, Rogério Seiji, Roger Fredy Larico Chavez, Ricardo Barroso Leite, Rangel Arthur, and Yuzo Iano (2013). "Plenoptic image compression comparison between JPEG, JPEG2000 and SPITH". In: *Cyber Journals: JSAT* 3.6 (cit. on p. 85).

Hu, Y. X. and K. Stamnes (1993). "An accurate parameterization of the radiative properties of water clouds suitable for use in climate models". In: *J. Climate* 6, pp. 728–742. DOI: 10.1175/1520-0442(1993)006<0728:AAPOTR>2.0.CO;2 (cit. on p. 17).

Hullin, Matthias B., Ivo Ihrke, Wolfgang Heidrich, Tim Weyrich, Gerwin Damberg, and Martin Fuchs (2013). "Computational Fabrication and Display of Material Appearance". In: *Eurographics 2013 - State of the Art Reports*. Ed. by M. Sbert and L. Szirmay-Kalos. The Eurographics Association. DOI: 10.2312/conf/EG2013/stars/137-153 (cit. on p. 9).

Huo, Yuchi, Rui Wang, Ruzahng Zheng, Hualin Xu, Hujun Bao, and Sung-Eui Yoon (Jan. 2020). "Adaptive Incident Radiance Field Sampling and Reconstruction Using Deep Reinforcement Learning". In: *ACM Transactions on Graphics* 39.1,

6:1–6:17. ISSN: 0730-0301. DOI: 10.1145/3368313. URL: https://doi.org/10.1145/3368313 (visited on 04/30/2020) (cit. on p. 15).

Iser, Tomáš, Tobias Rittig, Emilie Nogué, Thomas Klaus Nindel, and Alexander Wilkie (2022). "Affordable Spectral Measurements of Translucent Materials". In: *ACM Trans. Graph.* 41.6. ISSN: 0730-0301. DOI: 10.1145/3550454.3555499. URL: https://doi.org/10.1145/3550454.3555499 (cit. on p. 16).

ITU-T-P.910 (2008). *Subjective audiovisual quality assessment methods for multimedia applications*. Tech. rep. (cit. on p. 89).

Jakob, Wenzel (2010). *Mitsuba renderer*. http://www.mitsuba-renderer.org (cit. on pp. 23, 46, 52).

Jensen, Henrik Wann and Juan Buhler (2002). "A rapid hierarchical rendering technique for translucent materials". In: *ACM Transactions on Graphics (Proc. SIGGRAPH)* 21.3. ISSN: 0730-0301 (cit. on p. 12).

Jensen, Henrik Wann, Stephen R. Marschner, Marc Levoy, and Pat Hanrahan (2001). "A practical model for subsurface light transport". In: *Proc. SIGGRAPH* (cit. on pp. 12, 14, 16).

Jiang, Giulio and Bernhard Kainz (Oct. 2019). "One Shot Radiance: Global Illumination Using Convolutional Autoencoders". In: *arXiv:1910.02480 [cs]*. arXiv: 1910.02480. URL: http://arxiv.org/abs/1910.02480 (visited on 04/30/2020) (cit. on p. 15).

Kalantari, Nima Khademi, Steve Bako, and Pradeep Sen (July 2015). "A machine learning approach for filtering Monte Carlo noise". In: *ACM Transactions on Graphics* 34.4, 122:1–122:12. ISSN: 0730-0301. DOI: 10.1145/2766977. URL: https://doi.org/10.1145/2766977 (visited on 05/15/2020) (cit. on p. 15).

Kallweit, Simon, Thomas Müller, Brian Mcwilliams, Markus Gross, and Jan Novák (2017). "Deep scattering: Rendering atmospheric clouds with radiance-predicting neural networks". In: *ACM Transactions on Graphics (TOG)* 36.6, pp. 1–11 (cit. on pp. 15, 66, 68–71, 73, 78, 80).

KdotJPG (Sept. 2014). *OpenSimplex Noise*. en. Accessed 2020-10-04. URL: https://uniblock.tumblr.com/post/97868843242/noise (visited on 05/04/2020) (cit. on p. 69).

Keller, Alexander, Jaroslav Křivánek, Jan Novák, Anton Kaplanyan, and Marco Salvi (Aug. 2018). "Machine learning and rendering". In: *ACM SIGGRAPH 2018 Courses*. Vancouver, British Columbia, Canada: Association for Computing Machinery, pp. 1–2. ISBN: 978-1-4503-5809-5. DOI: 10.1145/3214834.3214841. URL: https://doi.org/10.1145/3214834.3214841 (visited on 05/15/2020) (cit. on p. 15).

Kim, Changil, Henning Zimmer, Yael Pritch, Alexander Sorkine-Hornung, and Markus Gross (July 2013). "Scene Reconstruction from High Spatio-angular Resolution Light Fields". In: *ACM Trans. Graph.* 32.4, 73:1–73:12. ISSN: 0730-0301. DOI: 10.1145/2461912.2461926. URL: http://doi.acm.org/10.1145/2461912.2461926 (cit. on p. 17).

Kingma, Diederik P. and Jimmy Ba (Dec. 2014). "Adam: A Method for Stochastic Optimization". In: *arXiv:1412.6980 [cs]*. arXiv: 1412.6980. URL: http://arxiv.org/abs/1412.6980 (cit. on p. 72).

Konrad, J., B. Lacotte, and E. Dubois (2000). "Cancellation of image crosstalk in time-sequential displays of stereoscopic video". In: *IEEE Trans. Image Processing* 9.5. ISSN: 1057-7149 (cit. on p. 13).

Lan, Yanxiang, Yue Dong, Fabio Pellacini, and Xin Tong (2013). "Bi-scale appearance fabrication". In: *ACM Transactions on Graphics* 32.4. ISSN: 0730-0301 (cit. on pp. 10, 12).

Lang, Manuel, Alexander Hornung, Oliver Wang, Steven Poulakos, Aljoscha Smolic, and Markus Gross (July 2010). "Nonlinear Disparity Mapping for Stereoscopic 3D". In: *ACM Trans. Graph.* 29.4, 75:1–75:10. ISSN: 0730-0301. DOI: 10.1145/1778765.1778812. URL: http://doi.acm.org/10.1145/1778765.1778812 (cit. on p. 19).

Levis, A., Y. Y. Schechner, and A. B. Davis (2017). "Multiple-Scattering Microphysics Tomography". In: *Proc. IEEE Conf. Comp. Vision & Pat. Rec. (CVPR)*, pp. 5797–5806. DOI: 10.1109/CVPR.2017.614 (cit. on p. 17).

Levoy, Marc and Pat Hanrahan (1996). "Light Field Rendering". In: *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques*. New York, NY, USA: ACM, pp. 31–42. ISBN: 0-89791-746-4. DOI: 10.1145/237170.237199. URL: http://doi.acm.org/10.1145/237170.237199 (cit. on pp. 3, 5, 85, 86).

Leyre, Sven, Youri Meuret, Guy Durinck, Johan Hofkens, Geert Deconinck, and Peter Hanselaer (2014). "Estimation of the effective phase function of bulk diffusing materials with the inverse adding-doubling method". In: *Appl. Opt.* 53.10, pp. 2117–25. DOI: 10.1364/AO.53.002117 (cit. on p. 16).

Liang, Shiyu, Yang Gao, Chonghao Hu, Peng Zhou, Aimin Hao, Lili Wang, and Hong Qin (2024). "State of the Art in Efficient Translucent Material Rendering with BSSRDF". In: *Computer Graphics Forum* 43.1, e14998. DOI: https://doi.org/10.1111/cgf.14998. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.14998. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.14998 (cit. on p. 12).

Lin, Weisi and C.-C. Jay Kuo (2011). "Perceptual visual quality metrics: A survey". In: *J. Vis. Commun. Image Represent.* 22.4, pp. 297–312 (cit. on p. 91).

Liu, Jing, Tom Malzbender, Siyang Qin, Bipeng Zhang, Che-An Wu, and James Davis (2015). "Dynamic mapping for multiview autostereoscopic displays". In: *Proc. SPIE, vol. 9391*, 1I:1–1I:8. DOI: 10.1117/12.2185676. URL: http://dx.doi.org/10.1117/12.2185676 (cit. on p. 88).

Lou, Qun and Peter Stucki (1998). "Fundamentals of 3D halftoning". English. In: *Electronic Publishing, Artistic Imaging, and Digital Typography*. Ed. by RogerD. Hersch, Jacques André, and Heather Brown. Vol. 1375. Springer Berlin Heidelberg. ISBN: 978-3-540-64298-5 (cit. on pp. 10, 33).

Maimone, A., G. Wetzstein, D. Lanman, M. Hirsch, R. Raskar, and H. Fuchs (2013). "Focus 3D: Compressive Accommodation Display". In: *ACM Trans. Graph.* 32.5, pp. 1–13 (cit. on p. 19).

Mantiuk, Rafat, Kil Joong Kim, Allan G Rempel, and Wolfgang Heidrich (2011). "HDR-VDP-2: a calibrated visual metric for visibility and quality predictions in all luminance conditions". In: *ACM Trans. Graph.* 30.4, 40:1–40:12 (cit. on p. 92).

Marchesin, Stéphane, Jean-Michel Dischler, and Catherine Mongenet (2007). "Feature enhancement using locally adaptive volume rendering". In: *Proc. of IEEE VGTC Conference on Volume Graphics*. Prague, Czech Republic, pp. 41–48. ISBN: 978-3-905674-03-3. DOI: 10.2312/VG/VG07/041-048 (cit. on p. 13).

Masia, Belen, Gordon Wetzstein, Piotr Didyk, and Diego Gutierrez (2013). "A survey on computational displays: Pushing the boundaries of optics, computation, and perception". In: *Comput. Graph.* 37.8, pp. 1012–1038 (cit. on pp. 19, 85).

Matusik, Wojciech, Boris Ajdin, Jinwei Gu, Jason Lawrence, Hendrik P.A. Lensch, Fabio Pellacini, and Szymon Rusinkiewicz (2009). "Printing spatially-varying reflectance". In: *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)* 28.5, 40:1–40:7 (cit. on pp. 10, 12).

Mimaki (2017). *Mimaki 3D printer*. https://mimaki.com/special/3d_print/. [Online; Accessed 2019-04-25] (cit. on p. 65).

Montalto, Carlos, Ignacio Garcia-Dorado, Daniel Aliaga, Manuel M. Oliveira, and Feng Meng (May 2015). "A total variation approach for customizing imagery to improve visual acuity". In: *ACM Transactions on Graphics* 34.3. ISSN: 0730-0301 (cit. on pp. 13, 33, 43).

Morovic, Ján and M Ronnier Luo (2001). "The fundamentals of gamut mapping: A survey". In: *Journal of Imaging Science and Technology* 45.3, pp. 283–290 (cit. on p. 59).

Morovič, Peter, Ján Morovič, Ingeborg Tastl, Melanie Gottwals, and Gary Dispoto (Sept. 2019). "Co-optimization of color and mechanical properties by volumetric voxel control". en. In: *Structural and Multidisciplinary Optimization* 60.3, pp. 895–908. ISSN: 1615-1488. DOI: 10.1007/s00158-019-02240-8. URL: https://doi.org/10.1007/s00158-019-02240-8 (visited on 10/14/2019) (cit. on p. 11).

Morsy, Mostafa Morsy Abdelkader, Alan Brunton, and Philipp Urban (2022). "Shape dithering for 3D printing". In: *ACM Trans. Graph.* 41.4. ISSN: 0730-0301. DOI: 10.1145/3528223.3530129. URL: https://doi.org/10.1145/3528223.3530129 (cit. on p. 10).

Morvan, Y. and C. O'Sullivan (2009). "Handling Occluders in Transitions from Panoramic Images: A Perceptual Study". In: *ACM Trans. Appl. Percept.* 6.4, pp. 1–15 (cit. on p. 18).

MPI (2017). *Light-field archive*. http://lightfields.mpi-inf.mpg.de/Dataset.html. Accessed: 2017-07-04 (cit. on p. 88).

Museth, Ken (July 2013). "VDB: High-resolution Sparse Volumes with Dynamic Topology". In: *ACM Transactions on Graphics* 32.3, 27:1–27:22. ISSN: 0730-0301 (cit. on p. 49).

Müller, Thomas, Brian Mcwilliams, Fabrice Rousselle, Markus Gross, and Jan Novák (Oct. 2019). "Neural Importance Sampling". In: *ACM Transactions on Graphics* 38.5, 145:1–145:19. ISSN: 0730-0301. DOI: 10.1145/3341156. URL: https://doi.org/10.1145/3341156 (visited on 04/29/2020) (cit. on p. 15).

Narasimhan, Srinivasa G., Mohit Gupta, Craig Donner, Ravi Ramamoorthi, Shree K. Nayar, and Henrik Wann Jensen (July 2006). "Acquiring Scattering Properties of Participating Media by Dilution". In: *ACM Transactions on Graphics (Proc. SIGGRAPH)* 25.3, pp. 1003–12. ISSN: 0730-0301. DOI: 10.1145/1141911.1141986 (cit. on p. 17).

Nimier-David, Merlin, Thomas Müller, Alexander Keller, and Wenzel Jakob (July 2022). "Unbiased Inverse Volume Rendering with Differential Trackers". In: *ACM Trans. Graph.* 41.4, 44:1–44:20. DOI: 10.1145/3528223.3530073. URL: https://doi.org/10.1145/3528223.3530073 (cit. on p. 15).

Nimier-David, Merlin, Sébastien Speierer, Benoît Ruiz, and Wenzel Jakob (July 2020). "Radiative backpropagation: an adjoint method for lightning-fast differentiable rendering". In: *ACM Transactions on Graphics* 39.4, 146:146:1–146:146:15. ISSN: 0730-0301. DOI: 10.1145/3386569.3392406 (cit. on pp. 15, 81, 82).

Nimier-David, Merlin, Delio Vicini, Tizian Zeltner, and Wenzel Jakob (Nov. 2019). "Mitsuba 2: a retargetable forward and inverse renderer". In: *ACM Transactions on Graphics* 38.6, 203:1–203:17. ISSN: 0730-0301. DOI: 10.1145/3355089.3356498. URL: https://doi.org/10.1145/3355089.3356498 (visited on 04/29/2020) (cit. on pp. 15, 75).

Nindel, Thomas Klaus, Tomáš Iser, Tobias Rittig, Alexander Wilkie, and Jaroslav Křivánek (July 2021). "A Gradient-Based Framework for 3D Print Appearance Optimization". In: *ACM Trans. Graph.* 40.4. ISSN: 0730-0301. DOI: 10.1145/

3450626.3459844. URL: https://doi.org/10.1145/3450626.3459844 (cit. on pp. 11, 15, 98).

Novák, Jan, Iliyan Georgiev, Johannes Hanika, and Wojciech Jarosz (2018a). "Monte Carlo Methods for Volumetric Light Transport Simulation". In: *Computer Graphics Forum*. ISSN: 1467-8659. DOI: 10.1111/cgf.13383 (cit. on p. 14).

Novák, Jan, Iliyan Georgiev, Johannes Hanika, Jaroslav Křivánek, and Wojciech Jarosz (2018b). "Monte Carlo Methods for Physically Based Volume Rendering". In: *ACM SIGGRAPH 2018 Courses*. SIGGRAPH '18. Vancouver, British Columbia, Canada: ACM, 14:1–14:1. ISBN: 978-1-4503-5809-5. DOI: 10.1145/3214834.3214880. URL: http://doi.acm.org/10.1145/3214834.3214880 (cit. on p. 14).

Panin, Mikhail and Sergey Nikolenko (Nov. 2019). "Faster RPNN: Rendering Clouds with Latent Space Light Probes". In: *SIGGRAPH Asia 2019 Technical Briefs*. SA '19. Brisbane, QLD, Australia: Association for Computing Machinery, pp. 21–24. ISBN: 978-1-4503-6945-9. DOI: 10.1145/3355088.3365150. URL: https://doi.org/10.1145/3355088.3365150 (visited on 05/01/2020) (cit. on pp. 15, 68, 73, 78).

Papas, Marios, Christian Regg, Wojciech Jarosz, Bernd Bickel, Philip Jackson, Wojciech Matusik, Steve Marschner, and Markus Gross (July 2013). "Fabricating translucent materials using continuous pigment mixtures". In: *ACM Transactions on Graphics (Proc. SIGGRAPH)* 32.4, 146:1–146:12. ISSN: 0730-0301. DOI: 10.1145/2461912.2461974 (cit. on pp. 10, 11, 16, 23, 30).

Patterson, Albert, Paul Collopy, and Sherri Messimer (Aug. 2015). *State-of-the-Art Survey of Additive Manufacturing Technologies, Methods, and Materials*. DOI: 10.13140/RG.2.2.33894.47684/2 (cit. on p. 1).

Pharr, Matt, Wenzel Jakob, and Greg Humphreys (2016). *Physically based rendering: From theory to implementation*. 3rd. Morgan Kaufmann (cit. on pp. 16, 30).

Pinson, M. H. and S. Wolf (2004). "A new standardized method for objectively measuring video quality". In: *IEEE Transactions on Broadcasting* 50.3, pp. 312–322 (cit. on pp. 18, 92).

Pintus, R., E. Gobbetti, P Cignoni, and R. Scopigno (2010). "Shape enhancement for rapid prototyping". English. In: *The Visual Computer* 26.6–8. ISSN: 0178-2789 (cit. on pp. 10, 11, 13).

Piovarči, Michal, Michael Foshey, Vahid Babaei, Szymon Rusinkiewicz, Wojciech Matusik, and Piotr Didyk (2020). "Towards spatially varying gloss reproduction for 3D printing". In: *ACM Trans. Graph.* 39.6. ISSN: 0730-0301. DOI: 10.1145/3414685.3417850. URL: https://doi.org/10.1145/3414685.3417850 (cit. on pp. 11, 12).

Prahl, Scott A., Martin J. C. van Gemert, and Ashley J. Welch (1993). "Determining the optical properties of turbid media by using the adding-doubling method". In: *Appl. Opt.* 32.4, pp. 559–68. DOI: 10.1364/AO.32.000559 (cit. on p. 16).

PTLens (2016). *Lens distortion correction software*. http://www.epaperpress.com/ptlens/. Accessed: 2016-15-11 (cit. on p. 87).

Rafał K. Mantiuk (2016). *Thurstonian scaling for pair-wise comparison experiments*. https://github.com/mantiuk/pwcmp. Accessed: 2016-15-11 (cit. on p. 90).

Ramanarayanan, G, J Ferwerda, and B Walter (2007). "Visual equivalence: towards a new standard for image fidelity". In: *ACM Transactions on Graphics (TOG)* 26.3, p. 76. DOI: 10.1145/1276377.1276472. URL: http://portal.acm.org/citation.cfm?id=1276377.1276472 (cit. on p. 90).

Reiner, Tim, Nathan Carr, Radomír Měch, Ondřej Šťáva, Carsten Dachsbacher, and Gavin Miller (2014). "Dual-color mixing for fused deposition modeling printers". In: *Computer Graphics Forum (Proc. of Eurographics)* 33.2 (cit. on p. 10).

Ren, Peiran, Jiaping Wang, Minmin Gong, Stephen Lin, Xin Tong, and Baining Guo (July 2013). "Global illumination with radiance regression functions". In: *ACM Transactions on Graphics* 32.4, 130:1–130:12. ISSN: 0730-0301. DOI: 10.1145/2461912.2462009. URL: https://doi.org/10.1145/2461912.2462009 (visited on 04/30/2020) (cit. on p. 15).

Ritschel, Tobias, Kaleigh Smith, Matthias Ihrke, Thorsten Grosch, Karol Myszkowski, and Hans-Peter Seidel (2008). "3D unsharp masking for scene coherent enhancement". In: *ACM Transactions on Graphics (Proc. SIGGRAPH)* 27.4, 90:1–90:8. DOI: 10.1145/1360612.1360689 (cit. on p. 13).

Rogers, Geoffrey L (1997). "Optical dot gain in a halftone print". In: *Journal of Imaging Science and Technology* 41.6, pp. 643–656 (cit. on p. 13).

Rossier, Romain (2013). "Framework for printing with daylight fluorescent inks". In: (cit. on p. 50).

Rouiller, O., B. Bickel, J. Kautz, W. Matusik, and M. Alexa (2013). "3D-printing spatially varying BRDFs". In: *IEEE Computer Graphics and Applications* 33.6. ISSN: 0272-1716 (cit. on p. 10).

Ruckdeschel, FR and OG Hauser (1978). "Yule-Nielsen effect in printing: a physical analysis". In: *Applied Optics* 17.21, pp. 3376–3383 (cit. on p. 13).

Sánchez Pérez, Javier, Enric Meinhardt-Llopis, and Gabriele Facciolo (2013). "TV-L1 Optical Flow Estimation". In: *Image Processing On Line* 3, pp. 137–150. DOI: 10.5201/ipol.2013.26 (cit. on p. 87).

Sandic-Stankovic, D., D. Kukolj, and P. Le Callet (2016). "Multi-Scale Synthesized View Assessment Based on Morphological Pyramids". In: *Journal of Electrical Engineering* 67.1, pp. 3–11. ISSN: 1339-309X. DOI: http://dx.doi.org/10.1515/jee-2016-0001. URL: http://www.degruyter.com/view/j/jee.2016.67.issue-1/jee-2016-0001/jee-2016-0001.xml (cit. on pp. 18, 92).

SciPy (2008). *scipy.optimize.least_squares SciPy Reference Guide*. https://docs.scipy.org/doc/scipy/reference/generated/scipy.optimize.least_squares.html. [Online; Accessed 2019-01-10] (cit. on p. 47).

Sharma, Gaurav, Wencheng Wu, and Edul N Dalal (2005). "The CIEDE2000 color-difference formula: Implementation notes, supplementary test data, and mathematical observations". In: *Color Research & Application* 30.1, pp. 21–30 (cit. on p. 74).

Shi, Liang, Vahid Babaei, Changil Kim, Michael Foshey, Yuanming Hu, Pitchaya Sitthi-Amorn, Szymon Rusinkiewicz, and Wojciech Matusik (Dec. 2018). "Deep Multispectral Painting Reproduction via Multi-layer, Custom-ink Printing". In: *ACM Transactions on Graphics* 37.6, 271:1–271:15. ISSN: 0730-0301 (cit. on pp. 10, 79).

Shibata, Takashi, Joohwan Kim, David M Hoffman, and Martin S Banks (2011). "The zone of comfort: Predicting visual discomfort with stereo displays". In: *J. Vis.* 11.8, 11:1–11:29 (cit. on p. 87).

Silva, V. De, H. K. Arachchi, E. Ekmekcioglu, and A. Kondoz (Sept. 2013). "Toward an Impairment Metric for Stereoscopic Video: A Full-Reference Video Quality Metric to Assess Compressed Stereoscopic Video". In: *IEEE Transactions on Image Processing* 22.9, pp. 3392–3404. ISSN: 1057-7149. DOI: 10.1109/TIP.2013.2268422 (cit. on p. 92).

Silverstein, DA and JE Farrell (2001). "Efficient method for paired comparison". In: *J. Electron. Imaging* 10.2, pp. 394–398. DOI: 10.1117/1.1344187 (cit. on pp. 89, 90).

Sitthi-Amorn, Pitchaya, Javier E. Ramos, Yuwang Wangy, Joyce Kwan, Justin Lan, Wenshou Wang, and Wojciech Matusik (July 2015). "MultiFab: A Machine Vision Assisted Platform for Multi-material 3D Printing". In: *ACM Transactions on Graphics* 34.4, 129:1–129:11. ISSN: 0730-0301. DOI: 10.1145/2766962 (cit. on pp. 2, 10, 21).

Smith, T. and J. Guild (Jan. 1931). "The C.I.E. colorimetric standards and their use". In: *Transactions of the Optical Society* 33.3, 73–134. ISSN: 1475-4878. DOI: 10.1088/1475-4878/33/3/301 (cit. on p. 79).

Solh, Mashhour and Ghassan AlRegib (2009). "MIQM: A novel Multi-view Images Quality Measure". In: *2009 International Workshop on Quality of Multimedia Experience*, pp. 186–191. DOI: 10.1109/QOMEX.2009.5246953 (cit. on p. 18).

Sone, Hiroki, Toshiya Hachisuka, and Takafumi Koike (2017). "Parameter estimation of BSSRDF for heterogeneous materials". In: *Eurographics Short Papers* (cit. on p. 23).

Song, Haichuan, Jonàs Martínez, Pierre Bedell, Noémie Vennin, and Sylvain Lefebvre (June 2019). "Colored Fused Filament Fabrication". In: *ACM Transactions on Graphics* 38.5, 141:1–141:11. ISSN: 0730-0301. DOI: 10.1145/3183793. URL: https://doi.org/10.1145/3183793 (visited on 04/30/2020) (cit. on p. 11).

Song, Ying, Xin Tong, Fabio Pellacini, and Pieter Peers (July 2009). "SubEdit: A Representation for Editing Measured Heterogeneous Subsurface Scattering". In: *ACM Transactions on Graphics (Proc. SIGGRAPH)* 28.3, 31:1–31:10. ISSN: 0730-0301 (cit. on pp. 12, 23, 24).

Song, Ying and Wencheng Wang (2013). "A data-driven model for anisotropic heterogeneous subsurface scattering". In: *2013 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference*, pp. 1–7. DOI: 10.1109/APSIPA.2013.6694333 (cit. on p. 23).

Stollnitz, Eric J., Victor Ostromoukhov, and David H. Salesin (1998). "Reproducing Color Images Using Custom Inks". In: *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH '98. New York, NY, USA: ACM. ISBN: 0-89791-999-8 (cit. on p. 11).

Stratasys (2017). *J750 Printer*. http://www.stratasys.com/3d-printers/production-series/stratasys-j750. [Online; Accessed 15-09-2020] (cit. on pp. 2, 10, 21, 38, 65).

— (2020). *Stratasys ultimate full-color multi-material 3D printer*. https://www.stratasys.com/3d-printers/j8-series. [Online; Accessed 15-09-2020] (cit. on p. 10).

Tamboli, Roopak R., Balasubramanyam Appina, Sumohana Channappayya, and Soumya Jana (2016). "Super-multiview content with high angular resolution: 3D quality assessment on horizontal-parallax lightfield display". In: *Signal Processing: Image Communication* 47, pp. 42 –55. ISSN: 0923-5965. DOI: http://dx.doi.org/10.1016/j.image.2016.05.010. URL: http://www.sciencedirect.com/science/article/pii/S0923596516300674 (cit. on p. 18).

Tao, Yubo, Hai Lin, Hujun Bao, Feng Dong, and Gordon Clapworthy (2009). "Feature enhancement by volumetric unsharp masking". English. In: *The Visual Computer* 25.5–7, pp. 581–88. ISSN: 0178-2789. DOI: 10.1007/s00371-009-0328-2 (cit. on p. 13).

Tech, G., Y. Chen, K. MÜller, J. R. Ohm, A. Vetro, and Y. K. Wang (Jan. 2016). "Overview of the Multiview and 3D Extensions of High Efficiency Video Coding".

In: *IEEE Transactions on Circuits and Systems for Video Technology* 26.1, pp. 35–49. ISSN: 1051-8215. DOI: 10.1109/TCSVT.2015.2477935 (cit. on p. 87).

Tewari, Ayush, Ohad Fried, Justus Thies, Vincent Sitzmann, Stephen Lombardi, Kalyan Sunkavalli, Ricardo Martin-Brualla, Tomas Simon, Jason Saragih, Matthias Nießner, Rohit Pandey, Sean Fanello, Gordon Wetzstein, Jun-Yan Zhu, Christian Theobalt, Maneesh Agrawala, Eli Shechtman, Dan B. Goldman, and Michael Zollhöfer (Apr. 2020). "State of the Art on Neural Rendering". In: *arXiv:2004.03805 [cs]*. arXiv: 2004.03805. URL: http://arxiv.org/abs/2004.03805 (visited on 04/29/2020) (cit. on p. 15).

Tompkin, James, Min H. Kim, Kwang In Kim, Jan Kautz, and Christian Theobalt (2013). "Preference and artifact analysis for video transitions of places". In: *ACM Trans. Appl. Percept.* 10.3, 13:1–19 (cit. on p. 18).

Tuchin, Valery V. (1993). "Laser light scattering in biomedical diagnostics and therapy". In: *Journal of Laser Applications* 5.2, pp. 43–60. DOI: 10.2351/1.4745330 (cit. on p. 17).

Urban, Philipp (2020). "Graphical 3D Printing: Challenges, Solutions and Applications". In: *London Imaging Meeting* 1.1, pp. 87–90. DOI: 10.2352/issn.2694-118X.2020.LIM-35. URL: https://library.imaging.org/lim/articles/1/1/art00021 (cit. on p. 2).

Vangorp, Peter, Gaurav Chaurasia, P-Y Laffont, Roland W Fleming, and George Drettakis (2011). "Perception of Visual Artifacts in Image-Based Rendering of Façades". In: *Comput. Graph. Forum* 30.4, pp. 1241–1250 (cit. on pp. 18, 91).

Vicini, Delio, Vladlen Koltun, and Wenzel Jakob (July 2019). "A learned shape-adaptive subsurface scattering model". In: *ACM Transactions on Graphics* 38.4, 127:1–127:15. ISSN: 0730-0301. DOI: 10.1145/3306346.3322974. URL: https://doi.org/10.1145/3306346.3322974 (visited on 04/29/2020) (cit. on pp. 15, 68, 73).

Vicini, Delio, Sébastien Speierer, and Wenzel Jakob (Aug. 2021). "Path Replay Backpropagation: Differentiating Light Paths using Constant Memory and Linear Time". In: *Transactions on Graphics (Proceedings of SIGGRAPH)* 40.4, 108:1–108:14. DOI: 10.1145/3450626.3459804 (cit. on p. 15).

Vidimče, Kiril, Szu-Po Wang, Jonathan Ragan-Kelley, and Wojciech Matusik (July 2013). "OpenFab: a programmable pipeline for multi-material fabrication". In: *ACM Transactions on Graphics* 32.4. ISSN: 0730-0301. DOI: 10.1145/2461912.2461993 (cit. on p. 10).

Viscoda (2016). *Voodoo Camera Tracker*. http://www.viscoda.com/en/voodoo-download. Accessed: 2016-15-11 (cit. on p. 87).

Vogels, Thijs, Fabrice Rousselle, Brian Mcwilliams, Gerhard Röthlin, Alex Harvill, David Adler, Mark Meyer, and Jan Novák (July 2018). "Denoising with kernel prediction and asymmetric loss functions". In: *ACM Transactions on Graphics* 37.4, 124:1–124:15. ISSN: 0730-0301. DOI: 10.1145/3197517.3201388. URL: https://doi.org/10.1145/3197517.3201388 (visited on 05/15/2020) (cit. on p. 15).

Wang, Jiaping, Shuang Zhao, Xin Tong, Stephen Lin, Zhouchen Lin, Yue Dong, Baining Guo, and Heung-Yeung Shum (Mar. 2008). "Modeling and Rendering of Heterogeneous Translucent Materials Using the Diffusion Equation". In: *ACM Transactions on Graphics* 27.1, 9:1–9:18. ISSN: 0730-0301 (cit. on pp. 11, 14).

Wang, Z., E. P. Simoncelli, and A. C. Bovik (Nov. 2003). "Multiscale structural similarity for image quality assessment". In: *Signals, Systems and Computers, 2004. Conference Record of the Thirty-Seventh Asilomar Conference on*. Vol. 2, 1398–1402 Vol.2. DOI: 10.1109/ACSSC.2003.1292216 (cit. on p. 92).

Wang, Zhou, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli (2004). "Image quality assessment: from error visibility to structural similarity". In: *IEEE Transactions on Image Processing* 13.4, pp. 600–12 (cit. on pp. 18, 37, 91).

Wanner, S., S. Meister, and B. Goldluecke (2013). "Datasets and Benchmarks for Densely Sampled 4D Light Fields". In: *Proc. Vision, Modeling, and Visualization (VMV)* (cit. on p. 17).

Wetzstein, G., D. Lanman, M. Hirsch, and R. Raskar (2012). "Tensor Displays: Compressive Light Field Synthesis using Multilayer Displays with Directional Backlighting". In: *ACM Trans. Graph.* 31.4, pp. 1–11 (cit. on p. 85).

Weyrich, Tim, Pieter Peers, Wojciech Matusik, and Szymon Rusinkiewicz (2009). "Fabricating microgeometry for custom surface reflectance". In: *ACM Transactions on Graphics (Proc. SIGGRAPH)* 28, 32:1–32:6. ISSN: 0730-0301. DOI: 10.1145/1531326.1531338 (cit. on pp. 10, 12).

Wilburn, Bennett, Neel Joshi, Vaibhav Vaish, Eino-Ville Talvala, Emilio Antunez, Adam Barth, Andrew Adams, Mark Horowitz, and Marc Levoy (2005). "High performance imaging using large camera arrays". In: *ACM Trans. Graph.* 24.3, pp. 765–776 (cit. on p. 86).

Wu, Yonghua, Shuki Chaw, Barry Gross, Fred Moshary, and Sam Ahmed (2009). "Low and optically thin cloud measurements using a Raman-Mie lidar". In: *Appl. Opt.* 48.6, pp. 1218–27. DOI: 10.1364/AO.48.001218 (cit. on p. 17).

Wyman, Douglas R., Michael S. Patterson, and Brian C. Wilson (1989). "Similarity relations for the interaction parameters in radiation transport". In: *Applied Optics* 28.24 (cit. on pp. 16, 27, 42).

Wyszecki, Günter and Walter Stanley Stiles (1982). *Color science: concepts and methods, quantitative data and formulae*. Wiley classics library. Wiley, pp. 168–169. ISBN: 9780471021063 (cit. on p. 57).

Xue, W., L. Zhang, X. Mou, and A. C. Bovik (Feb. 2014). "Gradient Magnitude Similarity Deviation: A Highly Efficient Perceptual Image Quality Index". In: *IEEE Transactions on Image Processing* 23.2, pp. 684–695. ISSN: 1057-7149. DOI: 10.1109/TIP.2013.2293423 (cit. on p. 92).

Yasakethu, S. L. P., C. T. E. R. Hewage, W. A. C. Fernando, and A. M. Kondoz (Nov. 2008). "Quality analysis for 3D video using 2D video quality models". In: *IEEE Transactions on Consumer Electronics* 54.4, pp. 1969–1976. ISSN: 0098-3063. DOI: 10.1109/TCE.2008.4711260 (cit. on p. 18).

Zhang, Cheng, Lifan Wu, Changxi Zheng, Ioannis Gkioulekas, Ravi Ramamoorthi, and Shuang Zhao (Nov. 2019). "A differential theory of radiative transfer". In: *ACM Transactions on Graphics* 38.6, 227:1–227:16. ISSN: 0730-0301. DOI: 10.1145/3355089.3356522. URL: https://doi.org/10.1145/3355089.3356522 (visited on 05/14/2020) (cit. on p. 15).

Zhao, Shuang, Ravi Ramamoorthi, and Kavita Bala (July 2014). "High-order Similarity Relations in Radiative Transfer". In: *ACM Transactions on Graphics* 33.4. ISSN: 0730-0301 (cit. on pp. 27, 42).

Zhao, Shuang, Lifan Wu, Frédo Durand, and Ravi Ramamoorthi (Nov. 2016). "Downsampling Scattering Parameters for Rendering Anisotropic Media". In: *ACM Transactions on Graphics* 35.6, 166:1–166:11. ISSN: 0730-0301 (cit. on p. 14).

Zheng, Quan, Vahid Babaei, Gordon Wetzstein, Hans-Peter Seidel, Matthias Zwicker, and Gurprit Singh (2020). "Neural light field 3D printing". In: *ACM Trans. Graph.* 39.6. ISSN: 0730-0301. DOI: 10.1145/3414685.3417879. URL: https://doi.org/10.1145/3414685.3417879 (cit. on pp. 12, 85).

Zheng, Quan and Matthias Zwicker (Aug. 2018). "Learning to Importance Sample in Primary Sample Space". In: *arXiv:1808.07840 [cs, stat]*. arXiv: 1808.07840. URL: http://arxiv.org/abs/1808.07840 (visited on 04/29/2020) (cit. on p. 15).

Zwicker, Matthias, Wojciech Matusik, Frédo Durand, and Hanspeter Pfister (2006). "Antialiasing for automultiscopic 3D displays". In: *Proc. of EGSR*, pp. 73–82 (cit. on p. 19).