New methods for the computational fabrication of appearance

Thiago Siqueira Pereira

A Dissertation
Presented to the Faculty
of Princeton University
in Candidacy for the Degree
of Doctor of Philosophy

Recommended for Acceptance
by the Department of
Computer Science
Adviser: Szymon Rusinkiewicz

June 2015

© Copyright by Thiago Siqueira Pereira, 2015.

All rights reserved.

Abstract

3D printing has been advancing rapidly with new machines becoming available each year. They can already accurately reproduce an object's shape. However, they are very limited when reproducing the object's appearance. Computational fabrication of appearance is an interesting research direction which seeks to extend the appearance reproduction capabilities of current devices and also to manage their limitations. It can have great impact in a number of different fields including product prototyping and design, realistic prosthesis and watermarks in security. This thesis presents three appearance fabrication works: a similarity metric, a light routing algorithm and reflectance fabrication process.

First, recent spatially varying reflectance (svBRDF) printing systems can reproduce an input document as a combination of matte, glossy and metallic inks. Due to the limited number of inks, this reproduction process incurs some distortion. To preserve a material's perceived variation with lighting and view, we introduce an improved BRDF similarity metric that builds on both experimental results on reflectance perception and on the statistics of natural lighting environments. We validate it quantitatively as well as through a perceptual study. We also show how to adapt traditional color gamut mapping methods to svBRDFs to preserve textures and edges.

Second, we use multi-material 3D printing to fabricate objects with embedded optical fibers, exploiting total internal reflection to guide light inside an object. We introduce automatic fiber design algorithms together with new manufacturing techniques to route light between two arbitrary surfaces. Our implicit algorithm optimizes light transmission by minimizing fiber curvature and maximizing fiber separation while respecting manufacturing constraints. Our methods enables new applications in sensing and display such as surface displays of arbitrary shape.

Third, existing BRDF fabrication methods are restricted to using pigments with isotropic light scattering. We propose the use of magnetic reflective pigments such that we can control their orientation by applying a magnetic field. We show how dynamic magnetic fields let us control not only off-specular lobes direction but also lobe width and anisotropy. We show how this magnetic control can be coupled with a projector to fabricate spatially-varying anisotropic BRDFs.

Acknowledgements

I would first like to thank my advisor, Szymon Rusinkiewicz, for his aid with all aspects of this work and infinite patience during its development .

Extremely invaluable were Wojciech Matusik, Steve Marshner, Tom Funkhouser and Adam Finkelstein, for all their ideas, advices and suggestions on how to keep improving this work.

Thanks to Luiz Velho for getting me interested in Computer Graphics and guiding my first steps in research.

The authors acknowledge the help and suggestions of Samuel Muff, Xavier Snelgrove, Moira Forberg, Shinjiro Sueda, Hao Li for the work in Chapter 4.

Much of the BRDF metric work originated in conversations with Wojciech Matusik and Fabio Pellacini, and we thank them for many helpful suggestions. We also thank Wojciech Matusik, Boris Ajdin, Jason Lawrence, and Paul Debevec for the BRDF, svBRDF, and environment-map datasets used in Chapter 2.

The work in Chapter 3 was made easier by the help and suggestions in questions of electronics and chemistry from numerous people including Sema Berkiten, Carolina Araujo, Linguang Zhang and Xinyi Fan.

Thanks to all Tiggraph reviewers and members of the Princeton Graphics Group who have read preliminary version of this work and provided great suggestions.

I thank the NSF grants CCF-1012147, CCF-1027962 and IIS-1116296, the DARPA grant N66001-12-1-4242, the Intel Science and Technology Center for Visual Computing for their support.

To my parents.

Contents

Abstract iii

Acknowledgements v

1 Introduction **1**

1.1 Previous work on reproduction of appearance 3

1.2 Previous work on non-appearance displays 5

1.3 Research contributions 6

2 Gamut Mapping Spatially Varying Reflectance with an Improved BRDF Similarity Metric **8**

2.1 Introduction 8

2.2 Related Work 10

2.3 Reproduction Framework 12

2.4 BRDF Similarity Metric 13

2.4.1 Synthetic Environment 14

2.4.2 Expression in Half-Angle Coordinates 16

2.4.3 Color and Image Comparison 17

2.4.4 BRDF Mapping Results 19

2.4.5 Perceptual Studies 22

2.5 svBRDF Mapping 25

2.6 svBRDF Results 27

2.7 Implementation Details 31
2.8 Conclusion and Future Work 32

3 Printing BRDFs with magnetically aligned pigments 34

3.1 Introduction 34
3.2 Related Work 35
3.3 Overview 38
3.4 Pigments and composition 39
3.4.1 Magnetic alignment 41
3.4.2 Concentration 45
3.5 Printing setup 46
3.5.1 Magnetic setup 47
3.5.2 Projector curing 51
3.6 BRDFs from dynamic magnetic fields 53
3.6.1 Random field 54
3.6.2 Circular field 55
3.6.3 Spiral field 56
3.6.4 Line field 56
3.6.5 Frequency of fields 57
3.7 Measurements 58
3.7.1 Circular field 58
3.7.2 Spiral field 62
3.7.3 Line field 62
3.8 Printed results 65
3.9 Conclusion 66

4 Computational Light Routing 67

4.1 Introduction 67

4.2 Related work . . . 69
4.3 Fiber fabrication . . . 71
4.3.1 Choice of materials . . . 71
4.3.2 Geometric factors . . . 74
4.3.3 Field of view . . . 75
4.3.4 Bending loss . . . 76
4.4 Fiber routing . . . 79
4.4.1 Curvature term . . . 82
4.4.2 Compression term . . . 83
4.4.3 Base layout . . . 85
4.4.4 Path constraints . . . 90
4.4.5 Implementation details . . . 91
4.5 Applications . . . 94
4.5.1 Display . . . 94
4.5.2 Sensing . . . 98
4.6 Limitations . . . 100
4.7 Conclusion . . . 101

5 Conclusion **102**
5.1 Future Work . . . 103

Bibliography **105**

Chapter 1

Introduction

3D printing has been advancing rapidly with new machines becoming available each year. The resolution of printers keeps increasing which allows them to accurately reproduce an object's shape. However, one important component is missing: object appearance. Cheap printers today, like some Makerbots, can only fabricate objects with two colors and even very expensive printers, like the Object Connex, are limited in the simultaneous color or materials available (Figure 1.1).

Current 3D color printers are already enabling compelling applications such as the 3D selfie. However, these small scale human replicas still look far from realistic (Figure 1.2). The shoes do not look like leather, the mug does not look ceramic. With printer/computer systems that allow careful control of material appearance, these replicas could be made more realistic. We could have hair that actually looks like hair, pants that look like velvet and even realistic skin.

This ultimate general-purpose machine is still far in the future, but progress in fabrication of objects with more complex appearance can have applications in many fields today. These fields include product design and prototyping, realistic medical prosthesis, new watermarks for security and also applications in displays and photography where a fabricated object can process light.



Figure 1.1: Current 3D printing systems are limited in the richness of material appearance they can reproduce. Cheap printers, like the Makerbot on the left, can only fabricate two colors and even expensive printers, like this Objet Connex on the right, can achieve full color but are limited in the materials it can use.



Figure 1.2: A 3D selfie is a compelling application of 3D printing, however the whole object looks like it was made of a single material.

Creating objects that interact with light in a user prescribed way is a challenging problem. The reason is most fabrication processes are very limited in how many simultaneous materials that can be used and in the resolution of the objects they create. Handling these limited resources in order to best achieve a prescribed light interaction goal is naturally a computational problem.

The field of computational fabrication of appearance aims at using algorithms to circumvent all the limitations of current fabrication process. Figure 1.3 shows different areas that come together to best solve problems in appearance fabrication. First, new optimization algorithms can be useful for designing 2D prints or 3D ob-

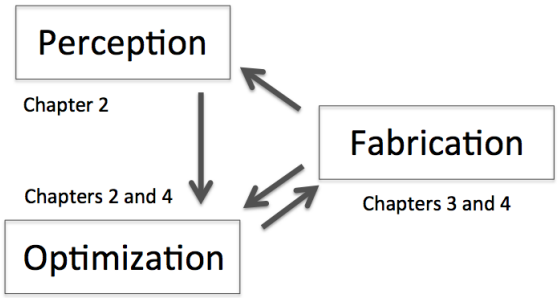


Figure 1.3: This thesis makes contributions in diverse areas such as material perception, optimization for design and appearance fabrication. All of these come together to advance our capabilities in designing and fabricating objects that interact with light in a prescribed way.

jects where the objective function is a light interaction goal while taking into account all the manufacturing limitations as constraints. Chapter 4 is an example of this approach where we design a set of 3D printable optical fibers that maximize light transmission and transport light between two surfaces of complex shape. Second, a better understanding of human perception of light and materials can be used to focus these optimizations in the aspects that are relevant to the end user. Chapter 2 proposed a material similarity function that can be used as a proxy for human perception in a spatially varying BRDF approximation process where printing constraints are considered. Finally, it is also possible to conceive of new fabrication processes that extend the capabilities of current machines. In Chapter 3, we investigate the use of magnetically aligned reflective flakes for printing spatially varying anisotropic BRDFs.

1.1 Previous work on reproduction of appearance

A number of works aim at creating physical instances that reproduce a given appearance. These appearance representations can be increasingly complex: uniform

BRDFs, spatially varying BRDFs, anisotropic BRDFs, subsurface scattering and even reflectance fields.

Some methods are restricted to a single BRDF and cannot achive spatial variation. Weyrich et al. [66] uses digital design and fabrication to create a reflective surface that achieves an input BRDF. They calculate a surface whose normal distribution matches the desired light scattering and mill it with aluminium. This lets them fabricate custom highlights of complex shape such as a teapot. Hullin et al. [24] achives a more limited set of BRDFs (Gaussian lobes), but allows for it to be changed dynamically. They excite water waves to achive an average scattering distribution.

Many other methods can achieve spatially varying appearance, usually at the cost of generality in the scattering distributions. Some of these works have been restricted to inks with isotropic appearance. Hersch et al. [23] models the appearance of 2D color printing integrated with metallic inks. Matusik et al. [35] uses a larger set of inks to print spatially varying BRDFs.

More recent works can fabricate anisotropic appearance. One approach is to use very small scale details and model diffraction with wave optics to achieve target anisotropic svBRDFs [32, 69, 18]. In a different aproach, Lan et al. [29] uses a two step process together with a 3D printer to create anisotropic bumps and print isotropic glossy ink over it. In Chapter 3, we describe a framework to print anisotropic appearance without using a 3D printer. We instead use reflective pigments anisotropically aligned by a magnetic field.

Some previous work can reproduce subsurface scattering. Papas et al. [43] fabricated objects of a uniform material using a mixture of pigments. These objects achieve a prescribed subsurface diffusion profile (BSSRDF). Other approaches [9, 22] use a 3D printer to fabricate a spatial arrangement of isotropic scattering material that approximates a spatially-varying BSSRDF.

Finally, some fabrication works can reproduce reflectance fields. These can exhibit even non-local light interactions such as cast shadows. Fuchs et al. [17] uses an arrangement of lenslets to flatten an input light distribution and modulates it with an attenuation layer before displaying the outgoing image or light field. Unfortunately their resolution is limited. Malzbender et al. [34] achieves higher resolution by fixing the viewpoint and using an array of curved mirrors with a transmissive layer on top.

Chapters 2 and 3 of this thesis also address the problem of reproducing appearance.

1.2 Previous work on non-appearance displays

We now discuss works that aim at creating an output light distribution, but whose purpose is not to reproduce a given physical object or scene. Many of these physical objects might be refered to as types of displays, but they all share the same challenges with the works previously described.

Weyrich et al. [65] explore the milling of diffuse surfaces for reproducing the shading of a scene. Alexa and Matusik [2] extend this technique to display distinct images when lit from different directions.

Many works explore refraction on curved surfaces as means of creating a target image. Finckh [14] proposed a method to compute a smooth surface that generates a prescribed image with its caustics. Papas et al. [42] designed and fabricated such objects. More recently, Kiser et al. [56] improved on the previous methods and showed high-quality caustic images. Finally, Papas et al. [41] showed an application of these refractive surfaces for steganography.

All these works fabricate objects with a single transparent material. More recently, Willis et al. [68] has shown how multi-material 3D printing can be used for printing optical fibers. These can be used to transport an image from one location to another. Brockmeyer et al. [5] showed how to design fibers whose end points lie on a sphere.

Pereira et al. [47] presented a design algorithm for routing printed fibers which enables the fabrication of objects of more complex shape. We describe this work in Chapter 4.

1.3 Research contributions

This thesis describes three works in increasingly more complex scenarios of light interacting with fabricated objects.

The first two chapters consider 2D prints where light reflects locally. Chapter 2 deals with 2D printing of spatially varying isotropic BRDFs. Chapter 3 deals with 2D printing of spatially varying anisotropic BRDFs. Finally, Chapter 4 deals with 3D printing objects where light can be transported globally between distant points on the objects surface.

Chapter 2 presents an svBRDF gamut mapping algorithm that minimizes distortions in the angular and spatial domains. To preserve a material's perceived variation with the direction of lighting and view, we introduce an improved BRDF similarity metric that builds on both experimental results on reflectance perception and on the statistics of natural lighting environments. Our experiments show better preservation of object color and highlights, as validated quantitatively as well as through a perceptual study. As for the spatial domain, we show how to adapt traditional color gamut mapping methods to svBRDFs. Our solution takes into account the contrast between regions, achieving better preservation of textures and edges. This work was previously published in [46].

In Chapter 3, we propose to use magnetic reflective pigments whose orientation can be controlled by applying a magnetic field. We identified existing magnetic reflective particles and measured their reflectance properties when oriented by a number of different input fields. Our setup uses electromagnets controlled by a microcontroller

to create time-varying magnetic fields. This magnet setup is integrated with a projector curing-based printer that can create spatially-varying reflectance by alternating between magnetic passes and resin curing passes. We show applications including printing anisotropic BRDFs and bump maps.

In Chapter 4, we use multi-material 3D printing to fabricate objects with embedded optical fibers, exploiting total internal reflection to guide light *inside* an object. We introduce automatic fiber design algorithms together with new manufacturing techniques to route light between two arbitrary surfaces. Our implicit algorithm optimizes light transmission by minimizing fiber curvature and maximizing fiber separation while respecting constraints such as fiber arrival angle. We also discuss the influence of different printable materials and fiber geometry on light propagation in the volume and the light angular distribution when exiting the fiber. Our methods enable new applications such as surface displays of arbitrary shape, touch-based painting of surfaces and sensing a hemispherical light distribution in a single shot. This work was previously published in [47].

Chapter 2

Gamut Mapping Spatially Varying Reflectance with an Improved BRDF Similarity Metric

2.1 Introduction

While most real-world materials exhibit a variety of appearances ranging from diffuse to glossy to metallic, traditional printing methods can only reproduce grayscale or color images. Motivated by advances in printing technologies, recent work has shown how to create physical reproductions of materials with angular dependent effects. These include opaque materials [35], but also translucent ones [9, 22]. In these works, the authors note that practical applications require goal-based printing pipelines, i.e. the user should only specify the output appearance instead of how much ink or material is to be used in each place. For instance, Matusik et al. [35] describe a reflectance printing system in which the user inputs a *spatially varying bidirectional reflectance distribution function* (svBRDF) [40], and the system finds the proper combination of inks of different reflectances through halftoning.

However, one challenge in traditional printing is also present in goal-based reflectance printing: gamut mapping [38]. Any printer has only a few different inks available, and it is not possible to achieve exact reproduction of input documents. These distortions may be in the angular dimension (e.g. highlights are not broad enough) or in the spatial dimension (e.g. edges have reduced contrast). In this work, we address the problem of svBRDF gamut mapping: finding the best possible approximation to an input svBRDF in the reproducible set of the printer, i.e. the printer's BRDF gamut.

While a large literature exists for image gamut mapping [38, 6, 30, 27], these algorithms cannot handle the angular effects of svBRDFs. To minimize angular domain distortion, it is possible to use a BRDF similarity metric [44, 39, 45]. Matusik et al. [35] used a metric that optimizes reproduction for point light visualization, which does not necessarily lead to good reproductions under natural environments. The distortion can be significant, specially for metals and specular materials. Inspired by the statistics of lighting environments [10], we propose a metric based on a new synthetic environment that correlates well with natural environments. We show that using this metric reduces perceptual distortion (Figure 2.1).

In addition, previous work has ignored the spatial arrangement of BRDFs, which may lead to loss of contrast in edges and texture. Finding a method that preserves the contrast between BRDFs and scales to the size of svBRDF datasets is a challenge. We have adapted recent gamut mapping approaches [27, 30] and show how they perform on svBRDF datasets.

Our method has applications beyond reproduction. For instance, it could fit parametric models to captured svBRDFs. In addition, our metric could be used to achieve more perceptually accurate svBRDF decompositions [31] and interactive edits [45].

Our main contributions are:

- An improved perceptual BRDF similarity metric based on a new synthetic lighting environment that correlates with natural environments (Section 2.4). We validate it through gamut mapping experiments and a perceptual user study.

- An adaptation of an image gamut mapping algorithm for svBRDFs. Our solution builds on our metric for clustering and optimal projections, but also takes the spatial arrangement of the BRDFs into account (Section 2.5). Our experiments show this method better preserves textures and edges (Section 2.6).

2.2 Related Work

BRDF similarity metric: An important part of an svBRDF gamut mapping system is its BRDF metric. Pellacini et al. [44] proposed a perceptually uniform reparametrization of the low-dimensional monochrome Ward BRDF space based on a psychophysical study. Generalizing this work to the high-dimensional space of real BRDFs, however, would be impractical. This has led to the use of algorithms, rather than explicit perceptual measurements, for determining BRDF similarity.

One approach are metrics that have an analytical expression, such as the L^2-norm and the metric of Lawrence and Pellacini [45], even though neither is perceptually inspired. Ngan et al. proposed a perceptual image-based metric[39] inspired by studies showing that humans are best at judging reflectance when seen under natural environments [15]. Therefore, they propose to compare BRDFs by instead comparing environment-mapped rendered images of spheres made of the materials (although [62] argues that more complex shapes are better for human perception).

However, Ngan et al. do not recommend a specific environment to be used as metric. Our experiments show that this choice can make a big difference. Matusik et al. [35] used a single point light as environment in their svBRDF reproduction system. While it leads to a simple metric, they claim that the point light does not corre-

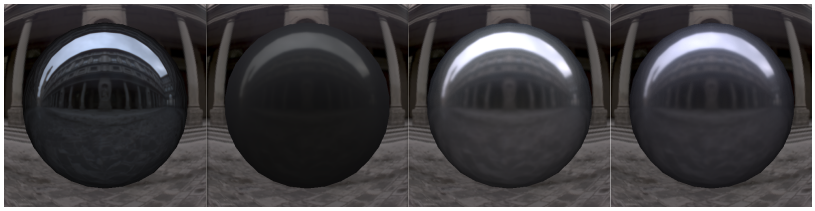


Figure 2.1: We address the problem of mapping a BRDF to a constrained gamut, such that it is close according to some similarity metric. Our metric (right) achieves better reproductions of the target material (left) when seen inside natural lighting environments compared to Matusik et al. (middle left) [35] and Pellacini et al. (middle right) [45].

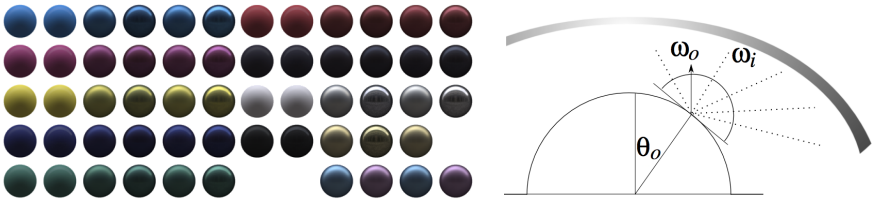


Figure 2.2: a) Reflectance inks from [35] b) Rendered sphere in our environment and its coordinate systems.

late well with natural environments when testing on large datasets of BRDFs. Our environment BRDF metric also follows this image-based approach, leveraging its corresponding perceptual results, but in addition we introduce a synthetic environment for comparing BRDFs that is consistent with the statistics of natural environments [10].

While a BRDF metric should be consistent with high-level reflectance perception, it should also build on low-level perception of color. For this, we use the CIELAB metric [12], even though other metrics could be used.

svBRDF gamut mapping: When mapping svBRDFs, it is not enough to maximize similarity per pixel. Instead of a reproduction that approximates the original in an absolute sense, we should rather aim at *relative* reproduction of the svBRDF, which aims at preserving edges and textures.

A good survey on color gamut mapping can be found in [38]. However, its unclear how to extend many color-specific concepts that are central to these algorithms, such

as hue preservation, luminance remapping and black point compensation. More recent developments, which we extend to svBRDFs, are spatial gamut mapping algorithms. They allow the same color to be mapped differently depending on its spatial position by using signal processing [70, 13] or optimization methods [27, 30].

In addition to spatial domain similarity, Guthe et al. [21] also consider the angular domain and develop a metric for Bidirectional Texture Functions. However, their metric predicts just-noticeable differences, while we focus on larger differences. The first work that focuses on large-scale changes in the angular domain for svBRDF gamut mapping was Matusik et al. [35]. In this work, the authors preserve spatial details by mapping material bases in the svBRDF convex hull and preserving combination weights. In their work, a BRDF is always mapped the same, independently of its position. In contrast, we allow it to map differently in order to preserve contrast.

Hersch et al. [23] describe a reproduction system with color and metallic inks and Stollnitz et al. [59] with multiple color inks. However, both works mainly focus on predicting the appearance of a combination of inks while we focus on minimizing perceptual BRDF distortion.

Other works gamut map materials with subsurface scattering appearance [9, 22]. Due to the more limited availability of scattering inks, these works focus on reproducing the achromatic characteristics of the materials.

2.3 Reproduction Framework

Matusik et al. [35] present a reflectance printing system that receives an input svBRDF and maps it to the printer gamut. We follow many of their assumptions about the printing process. First, they capture the BRDFs of a variety of ink stacks (e.g. a stack of cyan, yellow and silver foil ink would yield a metallic yellow ink). In this work, whenever we talk about inks, we refer to these stacked composites. They

also discuss how through halftoning a printer can generate any convex combination of its basis inks. In short, in this work, we assume the printer gamut to be any convex combination of their composite basis inks (Figure 2.2). As a source gamut we use the MERL database [36], containing a representative set of 100 BRDFs, and we consider several svBRDFs captured in previous work [31].

We represent all these BRDFs as $\rho(\theta_h)$ curves: sampled one-dimensional functions of the half-angle θ_h, defined to be the angle between the surface normal and the bisector of view and light direction [55]. This bisector gives the direction a microfacet would need to be oriented for perfect mirror reflection between light and observer. Since our input svBRDF maps and the output inks all have isotropic reflectance, this representation captures the main visual features of these BRDFs, namely the color and shape of the highlights. Its major disadvantage is that it does not model phenomena such as retro-reflection and grazing-angle effects. We find that previously proposed BRDF metrics performed well for some θ_h curves, but also disagreed widely with human perception for others. For this reason, we present an improved BRDF metric in the next section.

2.4 BRDF Similarity Metric

In this section, we present an improved environment-based BRDF metric. Our metric builds on the idea of comparing two BRDFs by comparing rendered images of objects having the two BRDFs under natural environments [39]. While it has been shown that humans can better perceive reflectance when seeing complex shapes [62], we follow Ngan et al. [39] and use simply the sphere shape, which leads to a simpler metric. After rendering, we convert each pixel's color to the CIELAB color space and compare the resulting images, taking an L^p difference pixel-wise. By doing so, we model human color perception more accurately.

However, we find the result of this metric to depend on the environment chosen. We designed a synthetic environment (subsection 2.4.1) that predicts well many natural environments because it is inspired by experimental analysis of natural environments. We propose its use as a reference for comparing BRDFs. We also show (subsection 2.4.2) how to find an analytical expression for our metric. In addition, we discuss two important features of our metric: use of CIELAB and choice of L^p-norm (subsection 2.4.3). Finally, we validate our metric by applying it to the gamut mapping problem and through perceptual user studies (subsections 2.4.4 and 2.4.5).

2.4.1 Synthetic Environment

In this subsection, we describe our proposed synthetic environment, but first we present a comparison of multiple environment metrics for the gamut mapping problem. We compare the metrics induced by several reference environment maps (beach, Grace, Uffizi, kitchen, St Peters) on three sets of BRDFs: the APLS printer inks, the MERL database and a set of synthetic Ward BRDFs [64] with ks and kd ranging from 0 to 1 (ks + kd $\leq$ 1) and roughness ranging up to 0.25. For each set, we compute pairwise similarities between all pairs of BRDFs under one environment, then compare the distances to those obtained using a different environment map. Ideally, we would like the BRDF similarity values to be consistent across different environments, in other words to have perfect correlation. Indeed, on the inks dataset, the correlation is quite high. However, for the MERL and Ward datasets there are substantial differences: though there is still a clear correlation between the results (Figure 2.3) using different environment maps, relative distances could vary by well over a factor of 2. It is interesting to notice that the Uffizi environment was most "typical" in the sense of agreeing best on average with other environments. Grace was most atypical. In conclusion, the environment-based metric is dependent on the environment that is used.

Let us consider next the metrics induced by two synthetic environments: a point light source at the camera and our new proposed environment. The latter also has a singularity at the pole, but has a heavy tail (Figure 2.3), with energy distributed as $1/\tan(\theta/2)$. It is motivated by the observation that real-world environments tend to have total energy per frequency varying roughly as $1/f$ [10] (sum of $2l+1$ spherical harmonics coefficients of energy $1/l^2$). This is achieved with an environment having energy distributed roughly as $1/\theta$. We actually choose $1/\tan(\theta/2)$, so that the function goes to zero as theta goes to π. We refer to this as the $1/f$ environment and its corresponding metric as the **1/f metric**.

In our experiments with the nine Debevec environments, we found significant greater variance in energy distribution than is acknowledged in previous work [15, 10]. While we also find the mean energy per spherical harmonic coefficient to be $1/l^2$, the exponents range betwen 1.5, 2 (St. Peters, Galileo), 2.1 ($1/f$ environment) up to 3 and 4 (Uffizi, Beach). Any environment that we pick is a compromise and we must evaluate how it agrees with others as a metric.

Comparing the BRDF metrics induced by point and $1/f$ to those of real environments, we find that they all have high correlation on inks. On Ward and MERL, the point light source is worse than any real environment, while the $1/f$ environment is comparable (Figure 2.3). In particular, it has high agreement with beach. This is easy to understand, since beach has essentially a major source of light, i.e. the sun, but also a radial falloff as light scatters off the sky. The conclusion is that the analytic $1/f$ environment does about as well as any real environment, plus it is radially symmetric (leading to faster evaluation).

For gamut mapping applications, we can replace any metric by its composition with an increasing function. Therefore, one could argue that correlation between metrics is not an appropriate measure of similarity between metrics. For this reason, we additionally ran all our correlations experiments using Spearman's correlation. This

is a correlation function that is invariant to composition with increasing functions. While some details do change, our conclusions are essentially the same.

One important advantage of the $1/f$ environment is its symmetry. In the next section, we show how it can be used to find an analytical expression for our metric.

2.4.2 Expression in Half-Angle Coordinates

In this section, we exploit the symmetry of our proposed environment and the symmetry of our chosen BRDF representation of θ_h curves to find an analytical expression for our metric, which can be used as an alternative to actually rendering the images. This expression is easier to plug in optimization methods. The process of computing the $1/f$ metric on θ_h curves is complex, since it requires a full spherical convolution. Let us now derive our metric equation and show how to precompute these convolution weights.

We start from the illumination equation to calculate the image we would obtain by rendering a sphere of constant BRDF described by a θ_h curve under any radially symmetric environment. Given two different BRDFs, we can integrate the L^p difference of their images. We assume that both the viewer and the environment's symmetry center are in the up direction (Figure 2.2). Parameterizing the sphere with θ_o, ϕ_o, we can see that radiance arriving at the eye from the sphere only depends on θ_o. Using the area element of the projected sphere in this parametrization and simplifying:

$$\frac{d(\rho_1, \rho_2)}{2\pi} = \int_0^{\pi/2} d_c(I_1(\theta_o), I_2(\theta_o))^p \sin\theta_o \cos\theta_o \mathrm{d}\theta_o, \tag{2.1}$$

where d_c is any color space metric. At this point, we look into the spherical convolution that results in the rendered image $I(\theta_o)$. This means we now integrate over incident light directions ω_i for a fixed θ_o value. We denote the viewer direction in the incident hemisphere's coordinate system ω_o. The outgoing light in this direction can

be calculated by the following integral $I(\theta_o) = \int f(\omega_i, \omega_o) E(\omega_i) \cos\theta_i \mathrm{d}\omega_i$. Integrating instead in the θ_h, ϕ_h coordinates [55] where the BRDF is a function of a single variable $\rho(\theta_h)$:

$$I(\theta_o) = \int \rho(\theta_h) E(\omega_i) \cos\theta_i \left| \frac{\partial \omega_i}{\partial \theta_h} \times \frac{\partial \omega_i}{\partial \phi_h} \right| \mathrm{d}\phi_h \mathrm{d}\theta_h. \tag{2.2}$$

We can precalculate all that does not depend on the BRDF: $I(\theta_o) = \int A(\theta_o, \theta_h) \rho(\theta_h) \mathrm{d}\theta_h$. We now make the expression of the function A more explicit. For our setup, we know that ω_o is in the z-direction (Figure 2.2) and we rewrite the environment $E(\omega_i) = E(\angle(\omega_i, \omega_o)) = E(2\theta_d)$. We also know that $\cos\theta_i = \omega_i^z$ and from the definition of h: $\cos\theta_i = 2\cos\theta_d \cos\theta_h - cos\theta_o$. We also need the θ_h, ϕ_h area element [53]. In addition, by solving $\cos\theta_i > 0$, we find the range of integration of ϕ_h to be $[-cos^{-1}K, cos^{-1}K]$, where $K = -\cot\theta_o \cot 2\theta_h$. Substituting, we obtain $A(\theta_o, \theta_h) =$

$$4 \sin\theta_h \int E(2\theta_d)(2\cos\theta_d \cos\theta_h - \cos\theta_o) \cos\theta_d \mathrm{d}\phi_h.$$

We believe that it is not possible to obtain a closed-form expression for this integral for our environment. For this reason, we discretize and precalculate $A(\theta_o, \theta_h)$ numerically. In conclusion, these weights let us compute $I(\theta_o)$ as a matrix-vector multiplication, which is easier to discretize and optimize. This formulation, lets us evaluate our metric in 3 ms in MATLAB. Our equivalent implementation rendering using a modern graphics card and BRDF importance sampling, but no symmetry, takes 120 ms. This speed up becomes crucial when processing svBRDFs.

2.4.3 Color and Image Comparison

We use the CIELAB color metric (D50 illuminant) since it builds on perceptual experiments on human color perception. While the RGB Euclidean metric could be used, our experiments find this solution to be inappropriate. In many cases, the RGB

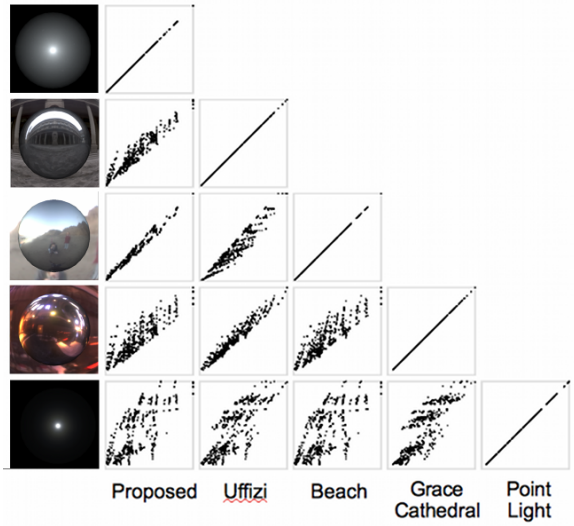
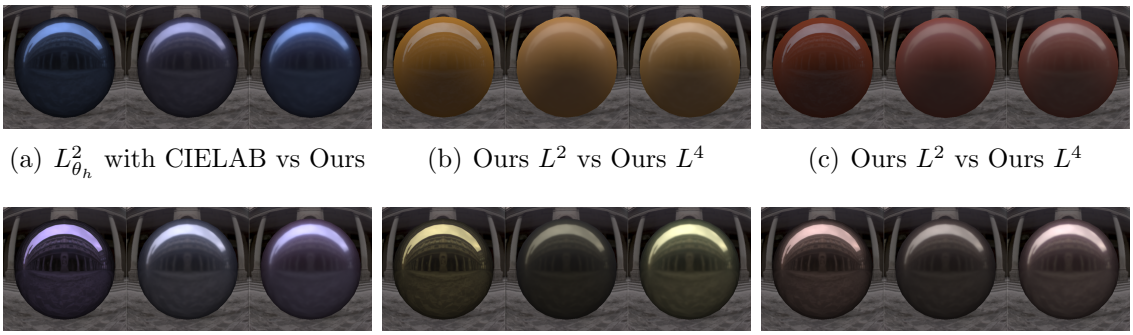


Figure 2.3: Scatterplots of correlations between BRDF similarity under different environments. Each plot considers a pair of environments, with each point corresponding to a pair of BRDFs. Its x, y coordinates are the distances between the materials as measured in both environments (so that points closer to the diagonal indicate better agreement). Our synthetic $1/f$ environment (leftmost column) has high correlation with the natural environments.



(a) $L^2_{\theta_h}$ with CIELAB vs Ours (b) Ours L^2 vs Ours L^4 (c) Ours L^2 vs Ours L^4

(d) Ours RGB vs Ours CIELAB (e) Point CIELAB vs Ours (f) Point CIELAB vs Ours

Figure 2.4: These comparisons display the target material (left) and the result with our metric (right). Our BRDF metric is built on four main ideas. First, that BRDFs should be compared through their rendered images in lighting environments. Image a) shows a reproduction using L^2-norm in θ_h-space with CIELAB color comparison. Our result preserves the hue. Second, that color comparisons should be performed using a perceptual color metric. Image d) middle shows the results of our metric using RGB space. Third, that to preserve highlights, the integration over angle should be performed as an L^4 norm. Middle of images b) and c) show our metric using L^2 instead, which results in blurred highlights. Finally, that BRDFs should be compared in the $1/f$ environment. For comparison, images e) and f) middle present the result using the point light metric [35].

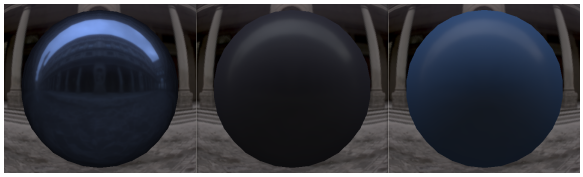


Figure 2.5: Restricted to a diffuse gamut, our metric (right) preserved the color, as opposed to the point metric (middle).

metric leads to incorrect hue, which can happen for both diffuse and glossy materials (Figure 2.4-d).

The integral in Equation 2.1 of our metric is essentially a simple image comparison metric. While we could have used more complex image metrics [63], the simpler choice of an L^p norm was favored. We experimentally chose $p = 4$. Compared to L^2, the L^4 norm gives less weight to small pixel errors and more weight to large errors. This is consistent with image perception by humans, since we are very tolerant of small changes in the mean intensity of an image.

Experimenting with gamut mapping, we find the highlights to be particularly sensitive to changes in p. Figure 2.4-b,c shows a comparison of reproductions using $p = 2, 4$. The highlights are sharper with L^4, while the overall colors barely changed. While highlights continue to improve for higher values of p, this leads to deviations in diffuse color.

2.4.4 BRDF Mapping Results

Having fully described the proposed BRDF metric, we now present comparison to previously proposed metrics using gamut mapping experiments. The images shown next are the result of solving a gamut mapping optimization problem. The mapping

of a BRDF y onto the gamut of a set of inks W using the metric d can be written as:

$$\begin{aligned} \min_{x} \quad & d(Wx, y) \\ \text{s.t.} \quad & x \geq 0, \sum x_i = 1. \end{aligned}$$

where Wx is a convex combination of the columns of W, i.e, Wx is any ink in the gamut defined by the convex hull of our basis inks. This is a non-linear optimization problem when using the CIELAB metric, but is a simpler quadratic problem in the RGB case. While most previous methods have not used the CIELAB metric for color comparison, we chose to implement them with CIELAB for a more fair comparison. We next show a comparison between ours and the $L^2_{\theta_h}$, point and cosine metrics. We discuss how they compare under different environments and gamuts.

We begin by showing gamut mapping results using the $L^2_{\theta_h}$ norm:

$$\int_0^{\pi/2} d_c(\rho_1(\theta_h), \rho_2(\theta_h))^2 d\theta_h$$

where d_c is the CIELAB metric. As can be seen in Figure 2.4-a, even when using the CIELAB metric, this simple $L^2_{\theta_h}$ metric fails to reproduce hue in many cases. The main drawback of this metric is that it does not compare BRDFs under any kind of lighting environment. The extreme mistakes above can be avoided by using an environment metric such as the point light. However, as we show below, our proposed metric can still achieve better results compared to the **point light metric** (Figure 2.6). The reason behind these results, as we saw earlier, is the low correlation between natural environments and the point light. In Figure 2.6, we show the exact same materials illuminated under a point light. The point metric indeed results in more similar materials when seen under point lighting.

We show many different cases where our metric achieves improved results. In Figure 2.4, our main advantage is the correct overall color. In other cases, as in

Figure 2.6, the improvement is in the highlight color and sharpness: notice how our reproduction shows the reflection of the buildings.

We also compare our metric to the weighted **cosine metric** [45]. We implemented it in RGB space because it led to a simpler quadratic optimization problem. Therefore, it is only fair to compare to our metric for gray materials. Our experiments show that this metric behaves similar to the point CIELAB metric. They both give too much weight to the highlights and tend to miss the overall color (Figure 2.1).

Overall, for our mapping experiments on the MERL database, our metric leads to less perceptual distortion when the materials are seen under natural environments like beach, kitchen and Uffizi. As for Grace and St Peters, the point light metric results in less perceptual distortion on average (Figure 2.6). This can be understood because these environments are composed of a large collection of distinct point lights. However, in regions where the environment is an area light, our metric more closely matches the original (Figure 2.6).

These results show that our metric can achieve good reproductions of the MERL materials on the inks gamut. It appears that most noticeable artifacts still left are a result of gamut limitations, e.g. the inks gamut does not include very diffuse materials, very glossy materials or very dark materials. We push these limitations further by removing the basis BRDFs that used silver or gold foil inks in their composition (columns 4,6,10 and 12 in Figure 2.2), which are necessary to approximate some metallic materials. We show results when mapping aluminium, but we observe similar behavior with many metallic materials. In Figure 2.7, we can see how our metric prefers a reconstruction with a broad highlight, which is certainly far from the target due to gamut limitations. The other metrics prefer a very dark reconstruction in order to have a sharper highlight. In these extreme trade-offs it is not so clear what is desirable, but, in our opinion, our reproduction was more faithful. Even only using diffuse inks, our metric can still create a reasonable reproduction (Figure 2.5).

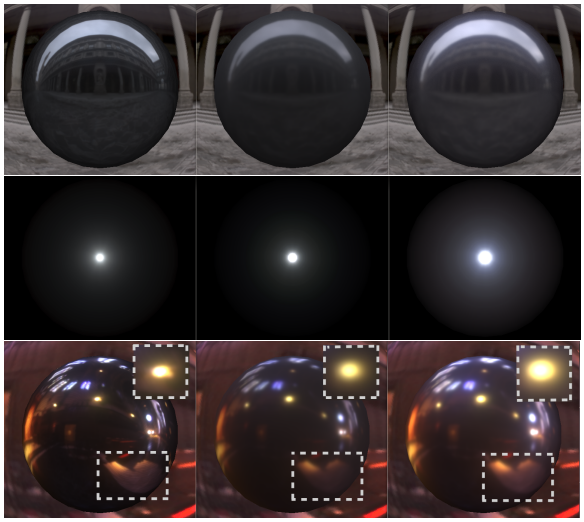


Figure 2.6: We achieve better reproductions under natural environments like Uffizi and beach, while the point metric is better under point and Grace. However, a detailed analysis (dashed lines) of Grace shows that in regions dominated by an area light, our proposed metric is better. Within each set: target (left), point (middle) and our metric (right). From top to bottom: visualization under Uffizi, point and Grace. We clip high intensity pixels.

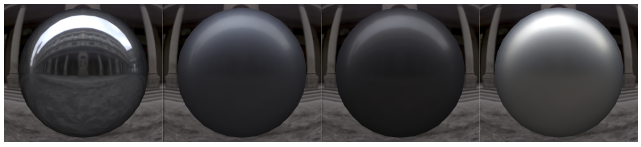


Figure 2.7: If we do not employ foil inks, we cannot reproduce aluminium (left). The $L^2_{\theta_h}$ and point light metrics (middle) preferred dark results in order to have a sharper highlight. Our metric (right) preferred a broader highlight.

Our comparisons to previous metrics under different environments and gamuts have shown the improved perceptual quality of our metric. To further validate it, we ran two perceptual user studies.

2.4.5 Perceptual Studies

We designed two perceptual studies to compare the performance of our solution to two other metrics in the gamut mapping problem. For each study, we used a different selection of materials from the MERL database. Our selected gamut was the

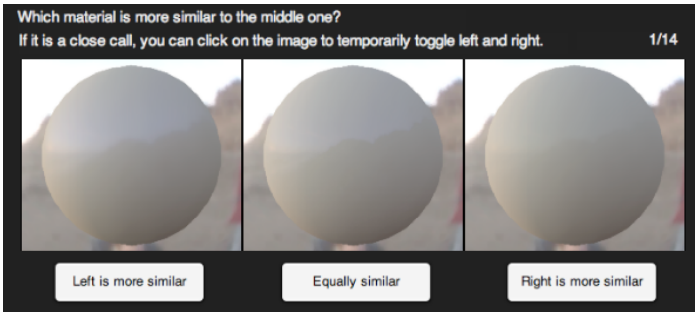


Figure 2.8: User interface used in our perceptual user studies.

same set of BRDF inks discussed in previous sections. We chose three representative environments based on our previous correlations analysis: beach, Uffizi and Grace.

Both studies use an interface (Figure 2.8) where the target material is presented in the middle and mappings with two different metrics are presented at the sides. The subjects are asked to select which of the mappings is more similar to the middle one. In addition, we also give them the option "Equally similar". By clicking and holding the mouse button, the user can temporarily observe the left and right images swapped. This overlaying makes it faster to observe some of the appearance differences. We also randomize left and right. This kind of perceptual study retains the disadvantage of the study of Pellacini et al. [44], i.e. only images are compared instead of the actual physical material.

We recruited 16 subjects between 21 and 37 years old. There were 11 males and 5 females. Some of the subjects participated in both studies. The average session took 25 minutes with 9 seconds per question. All the sessions were performed on the same calibrated display. Each sphere displayed occupied roughly 3.5 degrees of visual angle.

In our first experiment, we compare our metric to the point light CIELAB metric. A share of the MERL database lies inside the inks gamut. Since, for these cases, both metrics yield essentially the same results, we decided to focus on out of gamut materials. For this purpose, for each environment, we calculate CIELAB pixel differences

Ours preferred over	Uffizi and Beach	Grace
point CIELAB	73%	41%
cosine RGB	74%	67%

Table 2.1: Frequency with which the $1/f$ metric is preferred. Statistical significance p ¡ 0.05 for ±5% confidence.

between the two mappings. For beach/Uffizi/Grace, we discard material environment pairs where the average pixel color difference is less that 1/1/1 JND (just noticeable difference) and maximum less than 3/3/6 JND. We believe these thresholds are conservative because 50% of the materials remain and subjects find that many materials are still similar. A total of 153 questions are equally distributed among environments.

We observe that in this setting only 14% of the responses were 'Equally similar'. This means that our subjects found significant enough differences between the two metrics to justify a selection. We also observe that our subject population was rather consistent among themselves. On average, only 14.5% of the responses were the opposite of the majority (e.g. majority prefers metric A and response prefers B). In addition, we found this majority is usually significant: on average 75% of the subjects agree on a choice. The consistency of this population leads us to believe that the observed preferences are generalizable to a larger population.

For the beach and Uffizi environments, we find that in most cases where subjects have a preference they prefer our metric compared to the point metric. It was chosen 73% ± 2% (p ¡ 0.05). For the Grace environment, they prefer our metric in only 41% ± 3.5% (p ¡ 0.05). These results are consistent with our discussion in the previous section.

In our second experiment, we compare our metric to an RGB implementation of the cosine metric [45]. Since we are comparing against an RGB implementation, we restricted our study to only approximately monochrome materials in the MERL database, not necessarily out of gamut. This study consisted of 35 materials for a total of 105 questions equally distributed among the three environments. Since

many materials are in gamut, we observed a higher rate of 'Equally similar' responses 31%. We again observed a consistent population, only 8% of the responses are the opposite of the majority opinion. Even though the target materials are monochrome, the mappings with the cosine metric often result in some chroma. This led to users preferring our metric in 74%±4% (p ¡ 0.05) of the questions for Uffizi and beach and 67%±6% (p ¡ 0.05) for Grace.

All these results demonstrate a frequent preference for our metric. This preference is also considerable, which is attested by the non-forced choice nature of our study.

2.5 svBRDF Mapping

After finding improved BRDF mappings, we now consider multiple BRDFs per image, i.e. an svBRDF. Possibly the simplest approach to svBRDF mapping is to consider each pixel's BRDF in isolation and map it to the closest in gamut BRDF. This approach is known as **clipping**. Another approach is the **convex compression** solution presented in Matusik et al. [35]. In their work, they represent each BRDF in the document by a convex combination of basis BRDFs. They choose the basis near the convex hull of the source gamut, so that, when these are mapped, all the other BRDFs are compressed inside the destination gamut as well.

Both of these methods have advantages and drawbacks. Clipping has the property that it does not change materials that are already in gamut. However, it can lose spatial details in regions of the svBRDFs where all pixels map to the same in-gamut color. Convex compression can introduce large changes even to in-gamut colors, which leads to a loss of global contrast. In cases where multiple basis clip to the same point, compression also leads to loss of details.

To overcome these limitations, we have adapted two algorithms [27, 30] from the spatial gamut mapping literature. Both techniques use optimization to preserve

point-wise BRDF similarity and their spatial differences. They can be written in the following general form:

$$\begin{aligned} \min_{x_p} \quad & \sum_{p \in V} d(Wx_p, y_p) + \alpha \sum_{(u,v) \in E} d(Wx_u - Wx_v, y_u - y_v) \\ \text{s.t.} \quad & x_p \geq 0, \sum x_p^i = 1, \forall p \in V. \end{aligned}$$

where x_p^i are all the ink weights associated with vertex p, y_p is the target BRDF at vertex p and W is the gamut matrix as described in the previous section. Our objective function is non-linear because similarity is measured using our environment metric d. In fact, clipping can also be written in this general form by setting $\alpha = 0$.

The major difference between these two algorithms is how to define the sets of vertices V and edges E. The first is the **gradient-based** method [27]. It aims at preserving spatial gradients (i.e. difference between neighboring pixels). In other words, choose V to be the set of pixels and E to be the edges in all four neighborhoods. While this solution does manage to achieve a balance between clipping and compression algorithms, often having the advantages of both, it also suffers from two major drawbacks. It creates halo artifacts around strong edges (Figure 2.9). Even though halo could be improved with sparse gradient norms, this solution is also very slow, on the order of an hour for a 30 by 30 image. Compared to gradient reconstruction methods, this problem is harder because it is non-linear and constrained, which renders common speed-up techniques inapplicable.

A solution to these two problems is preserving the contrast between *regions* instead of pixels. This fixes halo because it focuses the optimization at preserving significant contrasts, as opposed to all pixel differences. In addition, there is a major speed-up since the number of variables becomes much smaller, proportional to the number of regions.

This method is similar to the optimization based algorithms of Lau et al [30]. It consists of four parts: clustering, optimization, interpolation and clipping. First, it starts by clustering the pixels into regions defined by BRDF and spatial distance with k-means. Clustering is performed in the perceptual space implied by our metric appended by the two spatial coordinates. In other words, we use the $I(\theta_o)$ curves in CIELAB space weighted by $\sqrt{\cos\theta_o \sin\theta_o}$ as our feature vectors. Second, the same optimization above is applied by taking V to be the cluster BRDF means and E to be neighboring clusters. This procedures preserves the differences between adjacent regions. Third, we add back the details. The simplest way would be to add the displacement between the mapped and source cluster material to all the other pixels in this cluster, but this may lead to artifacts at clustering boundaries. To avoid this, we follow the approach of Lau et al. [30], in which the authors interpolate the displacement vectors $Wx_p - y_p$ based on the inverse distance to each cluster center. Finally, to ensure that the image is in gamut, we run a final clipping step. This whole process has the advantage of good global contrast as a result of the optimization but also good preservation of details. We will refer to this method as the **cluster-based** solution.

In conclusion, clipping and convex compression fail to preserve contrast and edges. The gradient and cluster-based methods do a good job at preserving contrast, but only the cluster-based scales to the size of svBRDF datasets. Next section shows results that support these claims.

2.6 svBRDF Results

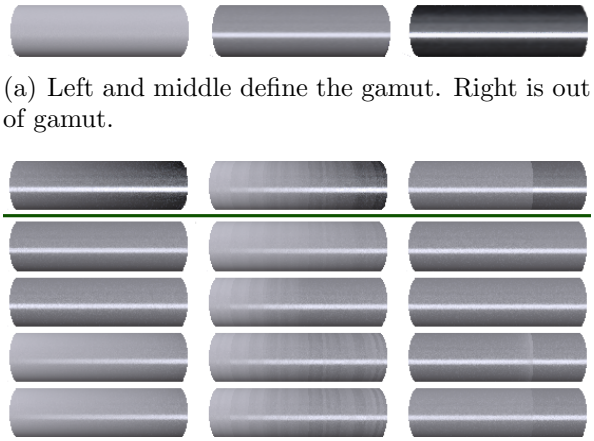
To compare these methods, we present three simple but challenging use cases. We use the same gamut in all cases: the convex combinations of the diffuse and specular inks shown in Figure 2.9-a. The metallic ink also shown in this figure is out of gamut and

is mapped to the specular ink by our metric. Each svBRDF is a horizontal gradient of two of these three materials (Figure 2.9). We choose to visualize them wrapped around cylinders because they are isometric to the document plane, but let us see multiple orientations in one single image. As a result, in a single cylinder any left to right variation is due to material change while top to bottom variations are caused by illumination. The number of clusters and α are inputs to the algorithm (values used are shown in captions).

In the first column, the input cylinder (top) is completely out of gamut: a gradient between specular and metallic. Both convex mapping and clipping lose all spatial variation. Gradient and cluster preservation preserve the spatial variation by using some diffuse ink on the left. The second input cylinder is only half out of gamut: a gradient with noise added to the combination weights. Convex compression preserves the gradient variations, but compresses the details. Clipping perfectly reproduces the left half of the cylinder, but loses all variations in the right half. Again, gradient and cluster preservation achieve good reproductions. Our third case is a simple edge between out of gamut materials. Convex compression and clipping lose the edge. Gradient preservation preserves the edge but leads to halo artifacts. The cluster-based solution preserves the edge without halo. Overall, the cluster-based method gives good results in all cases.

We also ran experiments on the svBRDFs from Lawrence et al. [31]. We found the full set of 57 inks to result in very good reproductions for this svBRDF dataset. To create more challenging cases, many of the experiments presented include gamuts with a reduced number of inks, gamuts with darker inks or svBRDFs where we retained the spatial variations but replaced the basis materials.

In Figure 2.11, all inks were scaled by 90% and we also mixed 25% diffuse white into all pixels to move the image more out of gamut. In the middle, we show how clipping can lead to total loss of edges. On the right, the cluster-based method can



(a) Left and middle define the gamut. Right is out of gamut.

Figure 2.9: Comparison of four svBRDF gamut mapping algorithms. Each column shows the mappings of a different cylinder. The target (above green line) is mapped using convex compression (first below green line), clipping (second), gradient preservation (third) and the cluster-based solution (bottom). Only the cluster-based solution works in all cases.

recover the contrast and make the edges visible again at the cost of worse matching of specularity.

As another example, we replaced the basis materials for the wheel. Our target appearance is made of a convex combination of brass, mapped brass and perfect black (Figure 2.13). Combinations of mapped brass and brass are projected to the same material. This leads to a strong loss of texture when using the simple clipping method. For both the wheel and the season greetings card, the convex compression method of Matusik et al. [35] would lead to similar results to clipping because these are cases where multiple basis materials are mapped to the same place.

The target in Figure 2.10 is a combination of copper and brass, both out of gamut. The clipping output is closer to the target but loses the contrast between red and yellow regions. The cluster-based method leads to a darker reproduction, but recovers the contrast. It preserves the specularity of the input, which is only possible by using the gold and silver foil inks. This explains the reproduction's lack of red.

As a final example, we studied the behavior of our metric with progressively restricted gamuts (Figure 2.12) using the clipping method. The target dove (first

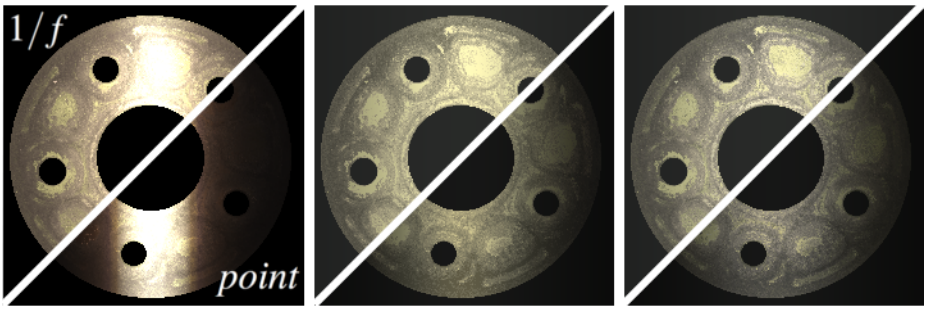


Figure 2.10: The clipping output is closer to the target (left) but loses global contrast. The cluster-based result (right) is darker in the red region enhancing the contrast. Visualized with our environment (top triangle) and point light (bottom). Generated with $\alpha = 50$ and 25 clusters.



Figure 2.11: Clipping (middle) can lead to total loss of edges from the target (left). In this case, the cluster method can recover the contrast and make the edges visible (right) at the cost of worse matching of specularity. Visualized with our environment (top triangle) and point light (bottom). Generated with $\alpha = 250$ and 10 clusters.

column) is well reproduced by the full gamut (second column). In the third column, we removed the metallic inks. This led the system to use the foil inks, which are more specular than desired. For the fourth column, we removed both metallic and foil inks. The result still shows a very narrow highlight from the inks that include a finish layer. The fifth column displays the svBRDF clipped using only the diffuse inks. Some specularity can still be seen since these are not perfectly diffuse.

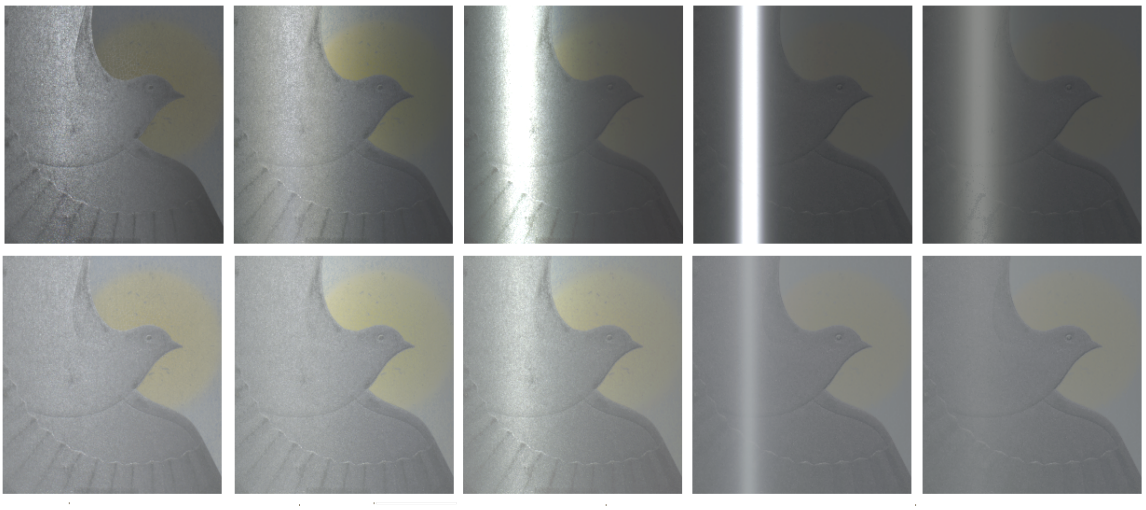


Figure 2.12: On the left we show the target svBRDF, which is well reproduced by the inks (second column). In the third column, we show the result using a gamut without metallic inks. In the fourth column, we further remove the foil inks. Finally, we show the projection on a diffuse gamut. Top row shows the dove visualized with the point light environment, bottom row with a $1/f$ environment.



Figure 2.13: The cluster-based method (right) preserves texture, while clipping destroys it (middle). Generated with $\alpha = 1.25$ and 25 clusters.

2.7 Implementation Details

Discrete metric: Discretizing our metric is straightforward. All functions of θ_h and θ_o become vectors by sampling the angles uniformly and transforming the integrals into summations. Equation 2.2 defines a linear mapping between the space of θ_h curves and the space of θ_o curves, which we write in vector notation as: $I = A\rho$. Each of its entries is precalculated by numerical integration. Our metric can be

computed by applying this matrix to the reflectance vectors, transforming the result to CIELAB and replacing the image space integral by a weighted summation.

Optimization: We implemented the BRDF gamut mapping process using an interior point method available through MATLAB's `fmincon` function. Because of the CIELAB non-linearity, our metric is a non-convex function. However, in practice, we did not observe local minima. We initialize the algorithm with a uniform combination of all inks. All the svBRDF algorithms were also implemented using the interior-point method in MATLAB. All are initialized with uniform inks. Most of the implementation details of the cluster-based method can be found in Lau et al. [30], but there are two major differences that we find necessary to make this solution practical for svBRDFs. First, in their work, they formulate this problem with a quadratic objective function subject to non-linear gamut constraints. Instead, we chose to formulate the non-linear objective function but linearly constrained problem. Second, even though projecting a single BRDF to a gamut of 57 inks takes only 2 seconds, we find it computationally prohibitive to simply project all pixels in our final svBRDF clipping step. Instead we quantize the input svBRDF with a large number of clusters (on average 400) using k-means. We only project the centers of these clusters. We visually inspected all images to make sure this quantization step is introducing negligible distortion.

2.8 Conclusion and Future Work

We presented a new gamut mapping algorithm for svBRDFs. We have shown how a synthetic environment can lead to a simple metric, but still agree with most natural environments resulting in perceptually accurate reproductions, including material color and highlights. We have also adapted some existing image gamut mapping methods to the svBRDF context. We show how the cluster-based solution leads to

good preservation of textures and edges, avoids halo artifacts and scales well to size of the svBRDFs.

One limitation of our approach is that we restricted the BRDF to depend only on θ_h. Therefore, we cannot represent retro-reflection, grazing-angle and anisotropic appearance. An interesting question is how to approximate anisotropic by isotropic BRDFs. In addition, our method is limited to svBRDFs, in which all interaction happens at the surface. We would like to extend our metric to scattering materials. Another future direction is extending our methods to svBRDFs on height maps or even 3D surfaces, allowing its application in 3D printing.

Chapter 3

Printing BRDFs with magnetically aligned pigments

3.1 Introduction

Fabricating custom bidirectional reflectance distribution functions (BRDFs) is a problem that has attracted the attention of the graphics community. It could have great impact in product design, manufacturing and security printing, yet existing approaches have many limitations. Most of them are restricted to isotropic reflectance [35, 43, 9, 22] Even more recent methods that fabricate anisotropic reflectance [29, 32, 69] cannot print BRDFs that give the impression of a bump map.

In this work[1], we propose to print appearance using ink with reflective magnetic particles, so that their orientation can be controlled by applying a magnetic field [48]. These particles are embedded in a carrier liquid that lets them rotate freely to align with the magnetic field lines until the liquid solidifies, the particles' orientation becomes fixed. We investigated the use of orientable particles for printing spatially-varying surface appearance (svBRDF).

[1]This chapter describes research performed in collaboration with Steve Marschner.

For svBRDF fabrication, anisotropic scattering particles can be seen as reflection from a discontinuous surface 3.1. This let us achieve large reflection angles with a completely flat print. It also let us avoid the strong dependencies between nearby pixels of a continuous surface. This greatly extends the range of manufacturable appearance using our approach compared to previous work. For example, minimizing the error in the approximation of a normal distribution by a continuous surface was an important concern in [66]. It also reduces problems caused by masking and self-shadowing.

The interesting appearance of magnetic flakes have been used for applications such as security printing and cosmetics, but their BRDFs have not been studied. In addition, existing fabrication techniques use static magnetic fields which do not allow for enough BRDF control. We have prototyped a fabrication setup that applies a dynamic magnetic field while selectively hardening pixels using projected UV light. This let us fabricate spatially varying complex BRDFs.

Our main contributions in this work are:

- We identified existing magnetic reflective particles and measured their reflectance properties when oriented by a number of different input magnetic fields
- We built a projector curing-based printer that can create spatially-varying anisotropic reflectance by applying a dynamic magnetic field synchronized with the curing of individual pixels

3.2 Related Work

Outside graphics. Using a magnetic field to orient reflective particles in ink was first proposed in [48]. The authors use a uniform rotating magnetic field created by solenoids to align reflective flakes horizontally and increase the inks reflectance in

the specular direction. The change in appearance with viewpoint has been used in security printing [26, 8]. For example, some bank notes may include some characters with magnetically oriented flakes. Matching the reflectance of these characters is impossible with standard printer pigments and is still hard without access to the exact magnetic field that created it. Magnetic flakes have also been used in the cosmetics industry such as in magnetic nail polish [60]. After a coating is applied to a nail, a custom shaped magnet is applied and quickly creates spatial variation of color. For this application, the spatial variation is the major aspect although the result is light/view dependent appearance.

Neither of these applications aim at automatically synthesizing a target appearance. The methods developed in this research may have great impact in both settings. For cosmetics, measuring and predicting the scattering of pigments will enable previewing the resulting pattern given a magnet shape. For both cosmetics and security printing, our work shows more complex BRDFs can be fabricated which might be useful for their applications.

Different mechanisms for creating images from orienting magnetic particles have been proposed in the literature using custom shaped magnets by either cutting [26] or milling [8]. Their main limitation is that the magnetic field does not change with time, which limits the range and reproducibility of the achieved BRDFs. In addition, the resulting orientation field is continuous. This can be seen in Figures 3.3 and 3.6. Continuity stems directly from the continuity of the magnetic field and as a result no hard edges can be created.

In the context of UV-curable magnetic compositions, the use of multiple static magnetic fields coupled with multiple light masks to create spatial variation has been proposed [54]. Our proposal for spatial variation builds on this idea. Instead of static fields we use dynamic fields created with the electromagnets described above. Instead

of physical masks for each printing pass, we propose to use a projector powered by a UV-led to cure different sets of pixels at a time.

Fabrication. Previous research in goal-based appearance fabrication can be divided in two approaches: those that use scattering pigments and those that use continuous surfaces.

Papas et al. [43] fabricate objects of a uniform material that achieve a prescribed subsurface light diffusion profile (BSSRDF) using a mixture of pigments of isotropic scattering. Other approaches [9, 22] use a 3D printer to fabricate a spatial arrangement of isotropic scattering material that approximates a spatially-varying BSSRDF. All these works are focused on the spatial dimension of the BSSRDF and ignore it's angular dimension. Matusik et al. [35] uses a dithered combination of glossy and diffuse inks to reproduce a target spatially-varying reflectance function. Due to the use of glossy inks, their approach allows for controlling specular reflection. However, all these works restrict themselves to unoriented pigments. In our work, we use oriented pigments to achieve anisotropic results. This is made possible by our use of a magnetic field to orient particles.

Surface based approaches [66, 14, 42] are more flexible in creating anisotropic reflectance (e.g. even teapot shaped highlights can be achieved) but usually use all their resolution to control angular reflection. As a consequence no spatial variation is achieved. This is largely a consequence of surface continuity which restricts their methods to smooth surfaces causing an effective loss in resolution. Two recent works use light diffraction to achieve higher resolution. The work of Levin et al. [32] does not allow non-symmetric reflectances such as we can accomplish with off-specular peaks, they are limited to symmetric specular peaks of complex shapes. Ye et al. [69] present off-specular peaks but with small bending away from the specular direction. They also only fabricate spatially varying results with a few BRDFs since it can take a few hours to compute a diffraction pattern for each BRDF used.

Finally, a recent approach takes a hybrid road. Lan et al. [29], in a two step process, use a 3D printer to create anisotropic bumps and print glossy inks over it. Even if their inks are isotropic, they can achieve anisotropic results with the underlying geometry. Our method can create anisotropic reflectance using only paper and ink, no 3D object is necessary. For cases where a 3D object is desired, their approach can be used together with our oriented pigments. This would extend the range of manufacturable appearance in their method even further.

3.3 Overview

Current spatially-varying reflectance printing methods only use inks with isotropic reflectance. These inks are isotropic because even if they are sometimes composed of non-spherically symmetric particles, these particles are usually present in random orientations and so are their average reflections. In our work, we propose to create anisotropic reflectance by orienting reflective particles with a magnetic field. Each printed pixel will have a different BRDF depending on the magnetic field that has been applied to it. Section 3.4 will describe the appearance of these pigments under static magnetic fields, will discuss different media in which they can be used and at which concentrations.

To gain some intuition on the effect of orientation on the inks reflectance, imagine each flake as a perfect mirror (Figure 3.1). In practice, each particle actually has a scattering profile. Randomly distributed mirrors will result in more diffuse like appearance. After alignment with the magnetic field, most of these mirrors will share a commom direction. Reflection will be stronger when the half-vector between the viewer and the light is in the same direction as the micro mirror's normals. Static magnetic fields are very limited in the BRDFs they can generate. We propose the use of time-varying fields. Section 3.5 describes our electromagnet setup that gives

us control of the 3 components of the field. Section 3.6 describes different fields and the BRDFs they generate. Section 3.7 shows our measurements of these BRDFs.

Magnetically orientable flakes are used both in security printing and for cosmetics. While there are many variations in their pipelines, they all share three common steps and so does our method. First the magnetic ink is applied on a substrate. Next a magnetic field is applied to orient the particles (Figure 3.2). Finally, the ink hardens while the magnetic field is still being applied. This hardening process freezes the particle orientations. We describe our fabrication setup for spatially varying BRDFs in Section 3.5 and show printed results in Section 3.8.

3.4 Pigments and composition

In this section, we describe our experiments with orienting magnetic pigments in static magnetic fields, how the appearance changes with concentration and the details of the final composition we use.

Most pigments used in metallic/pearlescent paints are flakes (small facets) that reflect light. However, most of these are not magnetic. In our work, we propose the use of magnetically orientable reflective particles that have different scattering distributions when oriented in different ways. Fortunately, many pigments with these characteristics have been developed for applications such as security printing and cosmetics. In these settings these particles are usually flat and have a layered material structure where at least one of them is magnetic [50]. These pigments are usually composed of a flat substrate of mica flakes, coated with layers of magnetic iron oxides and titanium dioxides. Their reflective properties including color can be controlled by varying these layers thickness which influences the interference of light between all these dielectric layers.

We have experimented with magnetic pigments using two different compositions: a solvent based and a uv-curable resin. First with pigments in a solvent-based composition: magnetic nail polish [60] which is commercially available under multiple brand names. In its expected use, after a coating is applied to a nail, a custom shaped magnet is brought near the nail which quickly creates spatial variation. After waiting a few seconds the coating dries, fixing the particles orientation. Figure 3.3 shows two samples of this product and the pattern that it can generate on a nail. Besides the very inconvenient smell, solvent-based compositions have another big disadvantage. It is hard to control how fast it dries which limits their use to orientation processes with a static magnetic field.

We also experimented with magnetic pigments in a UV-curable composition. In this case, the medium only hardens when exposed to UV light. While UV-curable magnetic nail polish are commercially available, the products we found did not seem to have the same visual quality compared to the magnetic nail polish due either to low concentration or low quality of pigments. Therefore, we identified a supplier of the raw pigments (Colorbridge) and mixed them in a UV-curable medium. Their magnetic pigments are available in two sizes, 5-25 microns or 10-60 microns. It is possible to see a glittery appearance for both sizes of pigments, but it is more pronounced with the larger pigments. Both sizes have the same qualitative behaviour and we have experimented with both of them. The measurements in Section 3.7 are for the larger pigments.

The same visual effect as the solvent-based magnetic nail polish can be achieved using this UV-curable composition. We have used two UV-curable mediums: a 3D printer resin and a nail polish gel. We chose to use the UV gel for convenience, but the same results apply to the printer resin. The UV curable medium has a number of advantages. First, by controlling the UV light intensity, we can dynamically control the drying time. This extra degree of freedom is especially useful when orienting

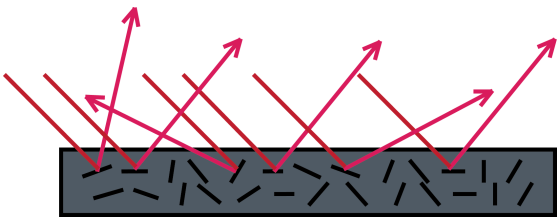


Figure 3.1: Parallel rays of light incident on oriented particles will scatter according to the designed distribution of facets. For simplicity, we depicted each reflection as perfect mirror reflection. In practice, each particle has a scattering profile that depends on its shape, size and composition. Refraction omitted for simplicity.

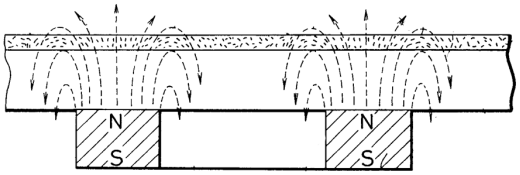


Figure 3.2: A spatially varying magnetic field (dashed line) created by two magnets is used to orient reflective particles (top layer). Image taken from [26].

pigments in a dynamic magnetic field. Second, it allows us to selectively harden some pixels by exposing them to light, while other pixels remain free to rotate. This feature will be later explored to create a spatially varying BRDF.

Next we present a small discussion of the appearance of these pigments under static magnetic fields, followed by our experiments with different concentrations of pigments and then we describe the final composition used in the rest of the paper.

3.4.1 Magnetic alignment

In this section, we introduce a qualitative discussion of the appearance of magnetic pigments in static magnetic fields. We discuss experiments using magnetic nail polish, but the results are essentially the same with the UV-curable composition. In section 3.7, we quantitatively support these conclusions with BRDF measurements of these pigments.



Figure 3.3: Comercially available magnetic nail polish is used to create patterns in nails. We used it for our preliminary experiments. A magnet hidden in the lid is brought close to the nail reorienting the flakes and thus creating spatial color variation.

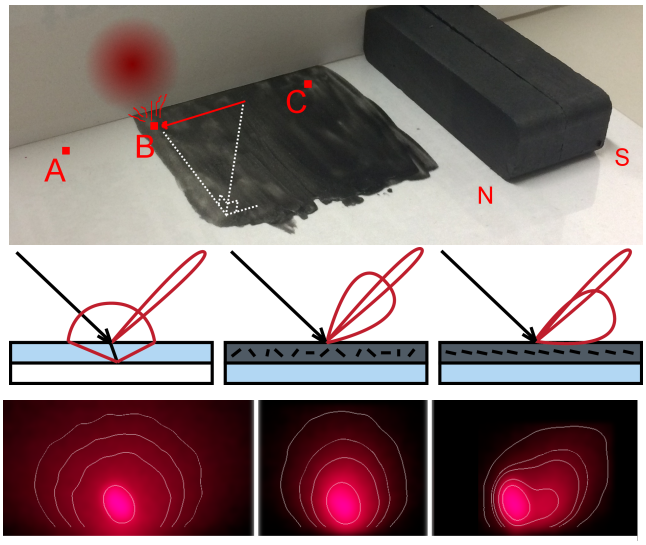


Figure 3.4: Top: a picture of our setup where a laser is used to create an incident ray at three different points. Point A (left) is a transparency on top of white paper shown here only for comparison. Point B (middle) is unoriented ink. Point C (right) is oriented ink. The second row shows three diagrams depicting each point's reflective lobes. The light that exits each point hits a vertical white sheet of paper. The third row shows the image formed on paper. By orienting the ink, we can create an off-specular lobe in a controllable direction (right).

These static field results are already used in security printing. We show them here only as examples to provide more intuition on the appearance of these particles in a spatial arrangement. In a first experiment shown in Figure 3.4, we used a bar magnet to orient pigments applied on a transparency (Figure 3.4). This experiment shows how the BRDF depends on the direction of the magnetic field. For this setup, the painted surface is mostly near the north pole of the magnet. While we are not depicting the magnetic field lines, their effect is the tilting of the microflakes' normals towards the magnet. They tilt by some amount at the point labeled C and are mostly unaffected at the point B due to the distance to the magnet. We use a laser with an elevation angle of roughly 45 degree to shine light in the selected points. The light distribution exiting each point hits a vertical sheet of paper of which we take a photo. We show both diagrams and pictures of the reflectance of each point. First, point A (left) is a transparency sheet overlaid on white paper. Its resulting reflectance is compose of mirror reflection at the air-transparency interface plus a diffuse lobe given by the paper. The diffuse lobe is shown for comparison. It is interesting to notice that points A, B and C all show a very similar mirror lobe given by either the air-transparency or the air-ink interface. When the paint is present there is no wide diffuse lobe, a large amount of the light that goes into the ink is reflected by the microflakes. For point B (middle), this creates a much smaller glossy lobe surrounding the mirror reflection. Finally point C has oriented microflakes, they create an off-specular lobe in a position that can be controlled by orienting the microflakes. When images are created using these pigments, it is this off-specular lobe that we mainly concerned with since most view directions are not aligned with the ideal specular direction. Yet, we remind the reader that these materials have traditional specular reflection from their top surface.

In Figure 3.5, we painted circles with nail polish. The left image shows unoriented appearance. While the reflectance changes with the elevation angle of the light (as any isotropic specular material), it does not depend on which side the light is coming

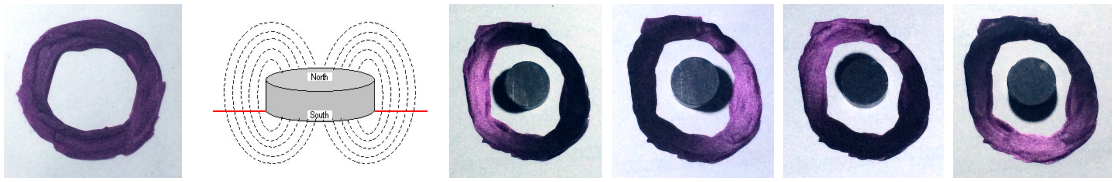


Figure 3.5: From left to right: unoriented nail polish, magnetic field lines arriving at an angle at the paper level (shown in red), pictures of oriented nail polish with difference light directions (see the magnet's shadow).

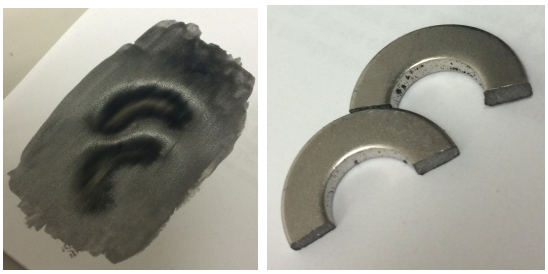


Figure 3.6: A spatially varying magnetic field can be used to create spatially varying reflectance. A coating of nail polish was applied on paper while the two magnets on the right were place 3mm below the paper. The resulting flake orientations gives the impression of a bump map. We remind the reader that this print is completely flat.

from. We used a disk magnet and painted a circle around it. The microflakes tend to bend towards the center. We show it lit from many light positions. Positions on the circle whose microflakes are oriented towards the light reflect more light and appear purple. Positions where normals are oriented away from the light appear very dark. We remind the reader that the actual normals of the surface are pointing up and are not responsible for the appearance variation.

We can also place the magnet underneath (or above) the paper. Results produced using this method are shown in the nail in Figure 3.3 and also in Figure 3.6 where we placed two arc shaped magnets 3 mm below the paper while applying the nail polish. Once again, the paint is completely flat. In both figures, the result resembles a bump map with a specular material.

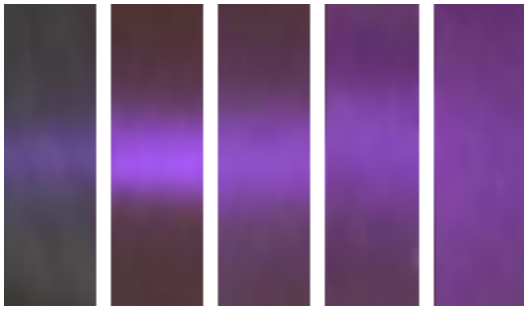


Figure 3.7: From left to right, increasing pigment concentration in a UV-curable composition. The color variation in each sample is caused by a static spatially varying field. The normals of the flakes at the center of each image are pointing towards the viewer, while the normals at the top and bottom are increasingly tilted away from the viewer.

3.4.2 Concentration

We performed an experiment to compare the BRDFs of different concentrations of pigments in the UV-curable nail polish gel. In this example, we used compositions at 1%, 6%,12%,18% and 30%, where the concentration is measured in mass of pigment divided by total mass of composition. These five samples have the same thickness as determined by the offset of a pair of glass slides.

A visual comparison quickly shows the lower concentrations being more transparent, as expected, where reflection is weak at 1% and transmission dominates the visible image. For the other concentrations, pigment reflection dominates. We can also see how higher concentrations required stronger magnetic fields to achieve complete alignment. This is potentially a good thing if we can then use the field strength to create large BRDF variations.

Figure 3.7 shows these five concentrations with the lower one on the left. These samples were created in the spatially varying magnetic field of a bar magnet. This field varies from top to bottom. These are not highlights created by uniform materials on a curved surface. For these images, the normals of the flakes at the center of the image are pointing towards the viewer, while the normals at the top and bottom are

increasingly tilted away from the viewer. The pigment orientation varying from top to bottom makes it look like a cylinder. The same spatially varying magnetic field was used for the 5 samples shown. The light source and the camera are in roughly the same orientation perpendicular to the sample plane. These images suggest different levels of pigment alignment at different concentrations. Pigments at 12% and 18% look aligned but with a bigger lobe width then 6%. With this field intensity, the 30% sample looks unaligned, yet this 30% composition can be perfectly aligned with an even stronger field.

In the next section, we will describe an electromagnet setup to orient these pigments. This setup can create fields which are much weaker than a bar magnet. Therefore, we chose to use the 6% concentration that does not require strong fields and achieves very good alignment. If we had a stronger setup, we expect a composition like 12% to be a better choice, since in that case field strength can be used to achieve an ever wider range of BRDF lobe widths.

Finally, we also add a 10% isopropyl alcohol by mass to the composition. Even a small percentage like this, already lowers the viscosity which also makes it easier for the pigments to align. We experimented with higher concentrations of alcohol, but they interfere with the curing process.

3.5 Printing setup

In this section, we propose a setup that allows fabrication of spatially varying BRDFs. The main idea is to use the UV-curable composition and a sequence of projected patterns with UV light. We can apply a dynamic magnetic field to change the BRDFs and then harden only pixels that are lit with the projector. This process locks the pigments at those locations, while others are still free to move. For example, to print a checkerboard with two BRDFs, we would apply the first magnetic field and then

illuminate the white squares in the checkerboard, completing the first pass. Then we would apply the second magnetic field and illuminate the black squares of the checkerboard, completing the second pass.

We first detail how to create a dynamic magnetic field. Then we detail the fabrication process.

3.5.1 Magnetic setup

We chose to design and build a setup of magnets that achieves the following requirements: 3 degrees of freedom (x,y,z), fast switching between fields, strong intensity, built from off the shelf components, easy to integrate with a projector. This setup was designed to be general and allow for many experiments with complex dynamic magnetic fields. However, these requirements come at a cost: limited spatial range. As we will see in this section, our setup gives us a field that is uniform in a small square of around 1x1cm.

In order to have strong fields from off the shelf components, we chose to use electromagnets with an iron core. We used both medium magnets with a 4 cm diameter and large magnets with a 5 cm diameter. Both have resistance around 20Ω. A guiding principle in our design of our magnetic setup was to keep the magnets as close as possible to each other to maximize field intensity. Magnetic field intensity falls with $1/r^3$ at a distance from the magnet.

We started our experiments using three perpendicular electromagnets (Figure 3.9-left) equally distant from a center point. While this gives us a 3 degree of freedom control, it creates a very non-uniform field. A second option would be to use 3 pairs of magnets, each pair aligned with one cartesian axis. While this setup would give us the most uniform field, it would be harder to integrate with the projector and its lens. Increasing the distance between the magnets and the printing area, would open space for the light path but decrease intensity. We chose to use only 5 magnets: two

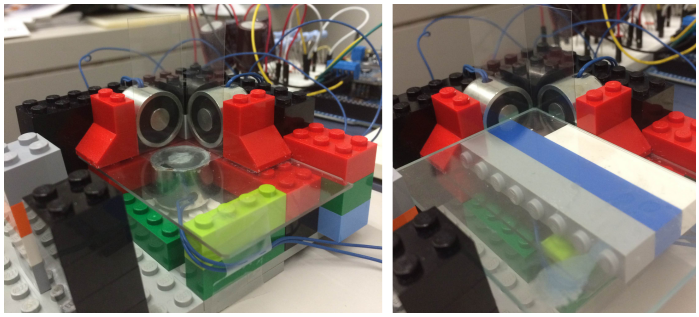


Figure 3.8: Our first prototype used three perpendicular electromagnets (left image) equally distant from a center point. This setup let us apply different currents to each of the magnets and achieve any 3d desired field at this point. Its main problem is the fields are very non-uniform at the center point.



Figure 3.9: Our final magnetic setup includes 5 magnets. Two pairs to achieve a uniform field in the x,y axis and z magnet (not shown) pointing up.

pairs for each of x and y axis, together with a single magnet for the z axis. These magnets were positioned in a tightly packed configuration on a machine milled frame (Figure 3.9), this means the four horizontal magnets are at a distance of 2 cm from the center point (the radius of the medium magnet). Notice that we used two larger magnets for the y axis.

This setup let us apply different currents to each of the magnets and achieve any 3d desired field at this point. We can achieve fast switching magnetic fields by programming the output currents using an Arduino microcontroller. The output of



Figure 3.10: We use a projector powered by a UV-led to selectively harden individual pixels. For BRDF fabrication, we apply a magnetic field (or a field schedule) and then harden all pixels compatible with the applied field. We then switch the field and harden other pixels. Printing proceeds in multiple passes.

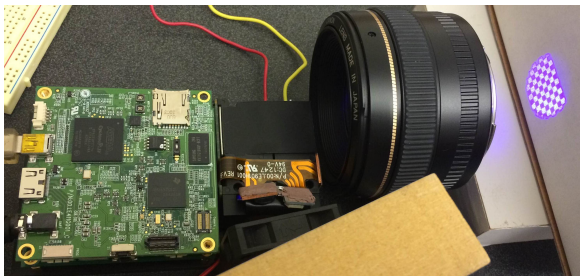


Figure 3.11: We used a lens in reverse to focus the light. This resulted in a small image but large enough given our target area of only 1x1cm.

the arduino is a PWM wave (a digital square wave) and it has to be passed through an RC filter before reaching the magnets. The 5V output of the microcontroller is not nearly enough to align these magnetic flakes. We observed that using an external 12V power source with a motor shield was enough to orient the nail polish composition, but still not enough for the UV-curable composition. Our final design uses a 24V power source instead, resulting in currents around 700 mA due to the resistances of the RC filter and the magnets themselves.

There are two problems created by having only one magnet in the z axis. First, at equal distances, it would give us a z field which is half as strong as x and y. This is easy to fix by bringing the z magnet a little closer to the printing surface. We placed

it at a distance of 1.6 cm from the printing plane which was found to be enough to increase the intensity. The second problem is the z field is not spatially uniform. Our solution to this problem was to use the other four magnets to compensate for the non-uniformity of z. At a high level, it works by running both x magnets (same for y) with opposite currents, such that they cancel each other at the exact center but still create a non-zero field around the center that can be used to cancel deviations in z.

In order to implement this solution, we need to find the exact currents to run through each magnet. We have measured the responses of all five magnets at the printing plane. Then for each target magnetic field, we solve for five currents that make the field at multiple x, y positions equal to the target. Aiming for multiple positions creates a more uniform end result. This process can be written as a linear system. Since scaling all the currents scales the magnetic field, we only solve the linear system for target fields on a unit hemisphere. The other side of the sphere can be found by simply scaling all currents by -1.

All this happens as a pre-computation. In order to have these results accessible by the microcontroller code, we gather the five currents calculated by the linear system and fit a fourth degree polynomial in two dimensions u, v. We chose u, v to be a stereographic mapping of the hemisphere since it introduces lower distortion.

Using the original measurements of each magnet and the optimized currents, we can simulate the field combinations and calculate the average (over all possible fields and positions) angle between the desired field and the achieved field. We calculated this angle error to have a mean of 3.4 and a maximum of 5.1 degrees. To validate this process, we re-measured the field for 5 different magnetic field targets at 4 different locations in the printing plane in a 8×8 mm square. Over this smaller set of 20 measurements, we observed a mean error of 2.5 and max of 7.3 degrees. These numbers contrast with the much larger errors if we do not use 5 magnets simultaneously.

The naive solution gives a mean measured error of 5.8 degrees and maximum error of 18 degrees. We scale all down all currents uniformly in order to have the maximum current calculated with the linear system be the maximum current delivered by the system. This uniform scaling results in a roughly 40% loss in total intensity. Our final system achives fields of 9000 uT, which is around 200× the earth's magnetic field.

All measurements were performed by moving a magnetic sensor with two linear stages. The response of each magnet is measured multiple times with different input currents and we use the slope of the resulting line which is more robust to noise.

3.5.2 Projector curing

Our full setup consists of a projector along the z axis together with the five magnets (Figure 3.10). We used the DLP Lightcrafter from Texas Instruments as it was easy to disassemble. This allowed us to replace one of its color LEDs with a high-power 385nm UV LED (from LED Engin model number LZ110UA00-U4). We power this LED with a current of 800 mA. Since our magnetic field is only uniform in a small area (1×1 mm), it was enough to project an image in a small area as well. We used a lens in reverse to focus the projectors image at a size of 2×2 mm (Figure 3.11).

During our printing process, we project each image during only 20 seconds. This is enough to partially cure the resin and stop the pigments from realigning with new fields. At the end of the printing process, we project a full white frame and cure all pixels for an additional 30 seconds. This means curing is not the bottleneck in our printing process. Given the current strength of the magnetic field we use, it takes around 45 seconds to align the pigments for each projected image. In our experiments with stronger magnetic fields, we find that it is possible to bring this alignment time to a few seconds, which would make curing the bottleneck.

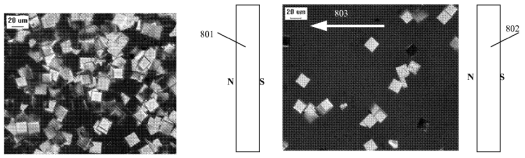


Figure 3.12: Microscopic image of magnetic flakes before (left) and after (right) applying a uniform magnetic field. The facets align along their diagonal. Images taken from [49]. The second image is of a different sample ink with lower density of flakes.

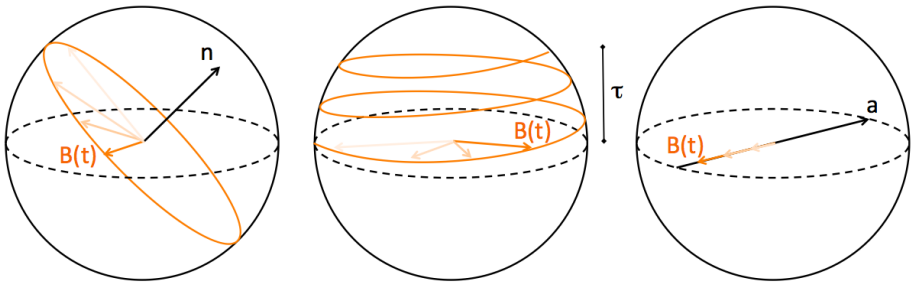


Figure 3.13: Time varying magnetic fields used to orient flakes. From left to right: circular field, spiral field and linear field.

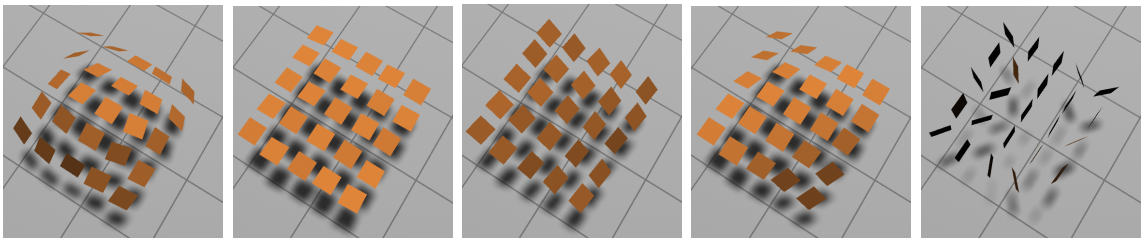


Figure 3.14: Depiction of flake orientations. From left to right: unoriented/spiral, circular pointing up, circular tilted right, anisotropic, spiral all the way to vertical.

3.6 BRDFs from dynamic magnetic fields

This section describes our experiments with time-varying magnetic field schedules (Figure 3.13) and our observations of the BRDFs they generate (Figure 3.14). The qualitative results discussed here are supported by measurements in the next section.

Raksha et al. [49] have taken microscopic images of flakes in an unoriented state and also of aligned flakes when applying a field along the x axis (Figure 3.12). This image shows how most of the flakes align along their diagonals. These images show that the alignment of diagonals with the field is imperfect, although quite good, but most interesting they show how each particle is still free to rotate around the x axis. This effect would cause greater light spread in one direction (y axis) than in the other (x axis). We will refer to this as the extra degree of freedom when aligning magnetic pigments in a static magnetic field. Since, pigments are free to move in this direction, their macro appearance is also much more glittery.

This degree of freedom is a great opportunity since it enables achieving anisotropic reflectance distributions, but it is also a great challenge since the simple process of aligning flakes with a static field will lead to different results depending on the starting configuration of flakes. We can think of the alignment process as being state dependent. This has two practical consequences: initial randomness and dependency on history.

First, whenever the composition with pigments is deposited, the pigments will have some random initial orientation. This orientation is usually not complety random and it does depend on how the composition was applied. For example, applying these pigments with a nail polish brush already causes some anisotropy aligned with the brushing direction. Whichever is the deposition method, a system that aims at good repeatability of BRDFs needs to clean this initial randomness of the flakes. We will discuss this cleaning process in more details below. In fact, this cleaning process needs to be applied to remove the effects of previous print passes as well.

Second, this dependency of magnetic field history means we have much more degrees of freedom: different sequences of magnetic fields will be useful as long as the sequence generates a different repeatable BRDF. If history did not matter, then all the degrees of freedom we would have would be field intensity and direction which would limit the range of achievable BRDFs. Next, we describe some time varying magnetic field schedules that we have experimented with and the BRDFs they generate. Experimenting with other schedules could lead to interesting future work.

Understanding the degree of alignment of the flakes is important because it will translate into the width of the BRDF's specular lobe. This width will be determined by three different factors. First, it is influenced by the roughness of the particle surface and scattering caused by it. Second, and mainly, it is influence by the orientation distribution of the flakes. Finally, this lobe increases further due to scattering at the air-medium interface. Notice that the incoming light is scattered twice by the interface (which has low but non-negligible roughness) on its way in and on its way out.

All BRDFs described in this work also have a specular lobe at this interface. Since this lobe is independent of the magnetic flakes orientation we will ignore for most of our discussion.

3.6.1 Random field

Our first attempt at a cleaning process was applying a sequence of random fields. This sequence does clean any deposition marks, however it had an undesirable side effect. It usually introduced some tilt in the flake distribution. In other words, the peak of the BRDF lobe was not in the direction of the normal, but instead got stuck in some other arbitrary direction. The role of cleaning will be left to the circular field.

It is unfortunate, that random fields had this side effect, since random would also provide a great solution to the problem of randomizing the flakes orientations,

or unaligning them. We will later discuss how the spiral field can perform this role instead.

3.6.2 Circular field

In the first work with alignment of magnetic pigments, Pratt and Salzberg [48] had showed how one way to control this extra degree of freedom. They used two electromagnets creating two horizontal magnetic fields in the x and y direction. By simply switching each field on and off, alternating at high frequency between the x and the y field, they obtained improved aligment of each facets normal direction along the z axis, resulting in a more specular sample.

Since our setup has x,y,z control and also supports digitally controlled continuous current levels (instead of on and off), we use a natural extension of their alternating field: the circular field. The circular field takes as input a normal vector n, direction of the BRDF lobe peak and therefore the perceived normal, and another input is the field intensity I. We calculate a basis u, v orthogonal to n. We then apply a circular field contained in this plane spanned by u, v:

$$B_c(t) = I(u \sin wt + v \cos wt)$$

, where I is the intensity of the input field and $w = 2\pi f$ with frequency f of 2 Hz.

As in their experiment, this also leads to improved alignment of the flake distribution, but lets us control the lobe direction. The parameter I can be used to control this degree of alignment where bigger intensities decrease the lobe width making the material more specular.

The circular field works as a very consistent cleaning process that can be used to erase any effect of previous fields and bring the flakes to a known distribution. We use it to "reset" the initial deposition of pigments.

3.6.3 Spiral field

With the circular field, there is no way to increase the lobe width, this question brings us to the next field: spiral. This field has constant maximum intensity (it is on the surface of a sphere). It starts as a circular field but then the z component slowly increases to a maximum value τ and the new field is normalized. Increasing τ increases the lobe width which also makes the appearance darker. This happens because many pigments try to align with the input field (which has a positive z component) and lose their alignment with the sample's normal.

More specifically, we use a thresholded Gaussian function $f(t) = \max[g(t), 1 - \tau]$ for interpolation between the circular field $B_c(t)$ and the vertical field $B_v = (0, 0, 1)$. So the spiral field is the normalized version of:

$$B_s(t) = f(t)B_c(t) + (1 - f(t))B_v$$

This field does not suffer from the drawback of the random field. The lobes that are generated still have their center in the direction of the sample's normal.

We have not experimented with more general spirals that start from an arbitrary plane spanned by u, v and move towards a normal direction n, but we expect it to work just the same.

3.6.4 Line field

Finally, the last field we discuss is the line field. We use this field to create anisotropic BRDFs. This field takes as input a tangent vector a in the x, y plane. The field then oscillates in the direction of this vector:

$$B_a(t) = a \sin wt$$

If this field is applied right after the circular field, nothing happens. The reason being all pigments would already be planar and no longer exhibit the extra degree of freedom. We propose to use this field after the spiral instead. As such, the lobe width first grows with the spiral and then the line field orients the pigments but still leaves the extra degree of freedom. In this way, the final result is anisotropic.

We have experimented with tangent vectors a outside of the horizontal plane, but have not measured their appearance. Visual inspection suggests they work in the expected way. Another interesting future direction is controlling the anisotropic ratio, which this line field does not allow.

3.6.5 Frequency of fields

As described above, our circular oscillations happen at 2 Hz. Higher frequencies (e.g. 8 Hz) can be used for faster alignment, but would also introduce more dampening of the effective intensity. We remind the reader that the square wave output from the microcontroller passes through a filter before reaching the magnets.

We arrived at this chosen frequency through testing alignment at higher and lower frequencies. We performed a preliminary experiment with a random field in a unit sphere. Using fields at a frequency of around .5 Hz, the pigments had enough time to align and the resulting BRDF seemed very close to the last random field having very little influence from it's history. At around 10 Hz, we obtained a pulsating appearance at a lower frequency than 10. This suggests some aliasing phenomena is already happening.

All these reported frequency values are very dependent on the media. In early experiments with nail polish which has lower viscosity, we observed pulsating appearance only at higher frequencies like 15 Hz. In addition, even 3 Hz is not enough to avoid the influence of the last field.

3.7 Measurements

We have measured the BRDFs of magnetic pigments oriented using three of the dynamic magnetic field schedules: circular field, spiral field and line field. These measurements quantitatively show how BRDF parameters such as lobe width, direction and anisotropy vary as a function of parameters of the magnetic schedules such as direction and intensity.

We used a gonioreflectometer to measure these BRDFs. During acquisition the camera is fixed in the z direction, while the light direction is free to move in the hemisphere. This gives us a 2D slice of the BRDF which is enough to characterize the major features of our flake distributions. The angle between the light and the camera varies between -75 and 75 degrees in increments of 5o, the three samples at angles -5,0,5 degrees are not acquired since the light source blocks the camera view direction. The highlight from the sample's surface is in this direction, but this is not a problem since we are most interested in measuring the off-specular lobe caused by the pigments.

Since we measured multiple painted samples in parallel, the camera view direction slightly varies for each samples in the same image. We have measured the position of each sample and used it to estimate a precise view direction. Some plots shown have BRDF not as a function of the light angle, but instead as a function of the half-angle which already takes the variation in view direction into account. Roughly speaking the half-angle equals half the light angle in this setup.

3.7.1 Circular field

We started by measuring the BRDF generated by circular fields to create pigment normal distributions where the average normal makes an angle α with the painted sample's normal, where α ranges from 0 to 30 in increments of 5 degrees. The

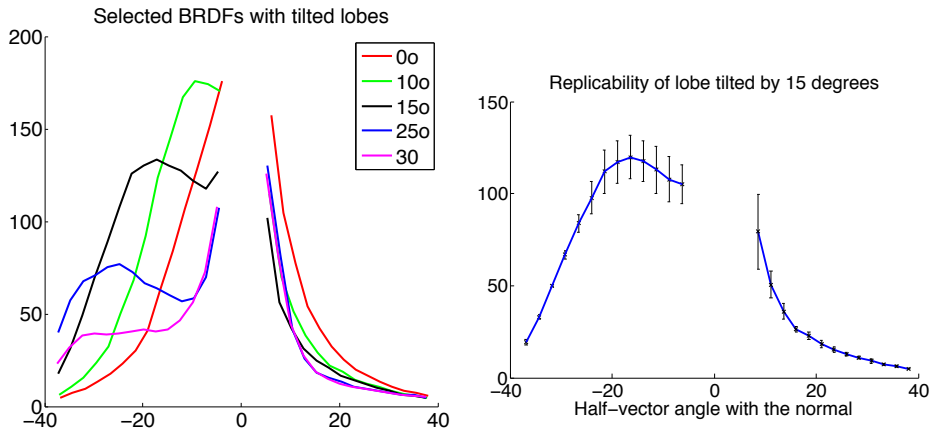


Figure 3.15: Left: this figure shows the major half-angle slice for selected raw measurements for $\alpha = 0, 10, 15, 25, 30$. It shows the surface's specular highlight centered at 0 and also the off-specular peaks centered at increasing angles. Right: The $\alpha = 15°$ sample repeated 3 times shows good repeatability of our orientation process.

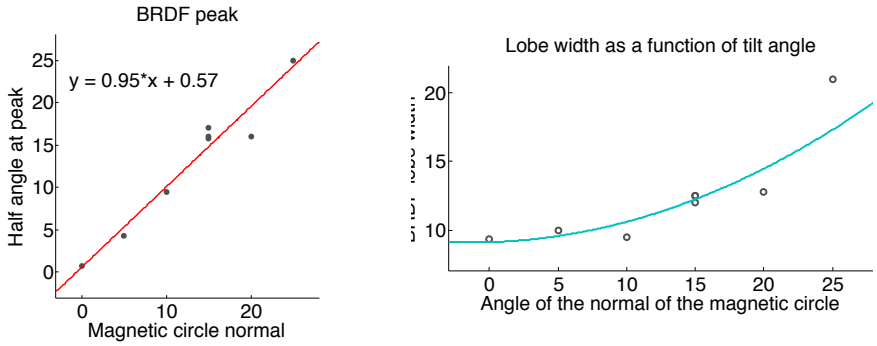


Figure 3.16: Effect of circular field's normal direction. Right, lobe width increases when orienting at higher angles. Left, the center of the lobe has a linear relationship with the circular field's normal.

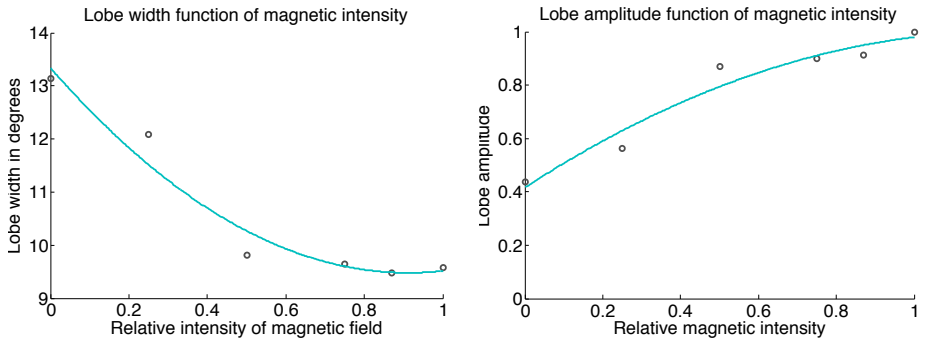


Figure 3.17: Effect of magnetic field strength on the BRDF. On the left, the lobe width decreases with increasing intensity. On the right, amplitude increases.

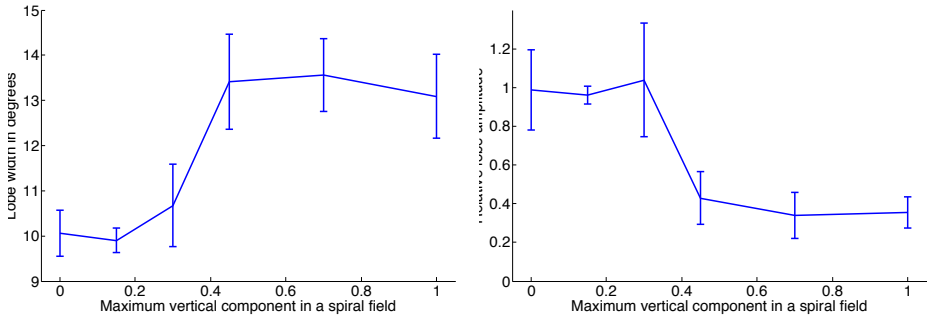


Figure 3.18: By varying the maximum z component in the spiral field we can control lobe amplitude and width.

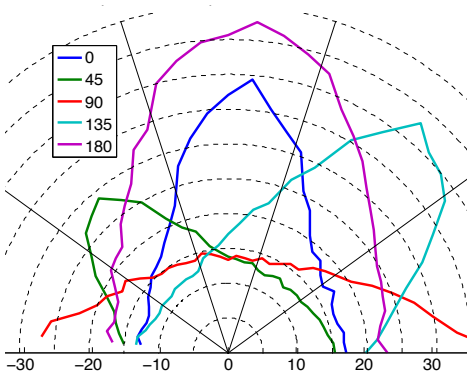


Figure 3.19: Polar plot of anisotropic BRDFs observed at 45° elevation angle.

BRDF at 30 degrees is already close to reflecting light only at grazing angles. In Figure 3.15-left, we show the major half-angle slice for selected raw measurements for $\alpha = 0, 10, 15, 25, 30$. The plot shows the surface's specular highlight centered at 0 and also the off-specular peaks centered at increasing angles. We painted the $\alpha = 15°$ sample 3 times, to explore the repeatability of our orientation process. Figure 3.15-right shows the mean BRDF in blue and the corresponding error bars.

In order to interpret our data, we fit a parametric model to slices of the off-specular lobes with center, width and amplitude as free parameters. For low α values such as 0 and 5, a simple laplacian distribution ($A\exp(-|x-u|/b)$ fits the data well. In log space this is an absolute value function where A is the amplitude, u is the center and b is the lobe width (or deviation). However for higher α values the off-specular peaks tend to be more round. We add a single additional radius parameter that controls a transition from Gaussian to Laplacian and results in very good fits of the rounded peaks.

With increasing α, we obtain as expected an increased peak center. Figure 3.16-left, shows how the center of the lobe is exactly at half-angle α. The lobe width also increases with α (Figure 3.16-right) while the amplitude decreases (not shown).

We also oriented the samples using a series of horizontal circular fields with increasing intensity. All these BRDFs are centered (lobes at 0 degrees), but they show variation in the specular lobe. Figure 3.17 shows how the lobe width decreases with increasing field strength. At intensity 0, we can see the lobe width of unoriented pigments is around 13 degrees. It decreases to around 9.5. It also shows how the amplitude increases. These are two sides of the same coin, with more pigments aligned at the center there is more light being reflected in the center direction. We expect the lobe width to decrease even further with even stronger fields.

3.7.2 Spiral field

We now show the effect of the spiral field. As described in the previous section, this field has a parameter which is the maximum z component which it converges to. The most visible effect is that the samples get darker. Figure 3.18-right shows the decrease in lobe amplitude with an increase of the maximum z value. At the same time, it can increase the lobe width as shown in Figure 3.18-left. A very interesting feature of both these transitions is that they happen very suddenly with a transition threshold around 0.3.

3.7.3 Line field

In this experiment, we demonstrate the anisotropic BRDFs generated with the line magnetic field. Each sample is painted separately as a way of starting with unoriented pigments with a wider lobe. The magnetic field line direction takes the value $(0, 45, 90, 135, 180)$. Figure 3.19 shows the BRDF in a polar plot at a fixed elevation angle of 45 degrees. The major axis of each BRDF clearly aligns with the magnetic field direction. We observed a mean ratio between major and minor axis of 2.5 for these five samples in the plot. The ratio ranged between 2.3 and 2.9.

Starting the experiment from the unoriented pigments reduces the consistency of this experiment as can be seen in the different sizes of lobes that were generated.

We will repeat this experiment using the spiral field followed by the line field (as described in the previous section). This would reduce the dependency on the starting state.

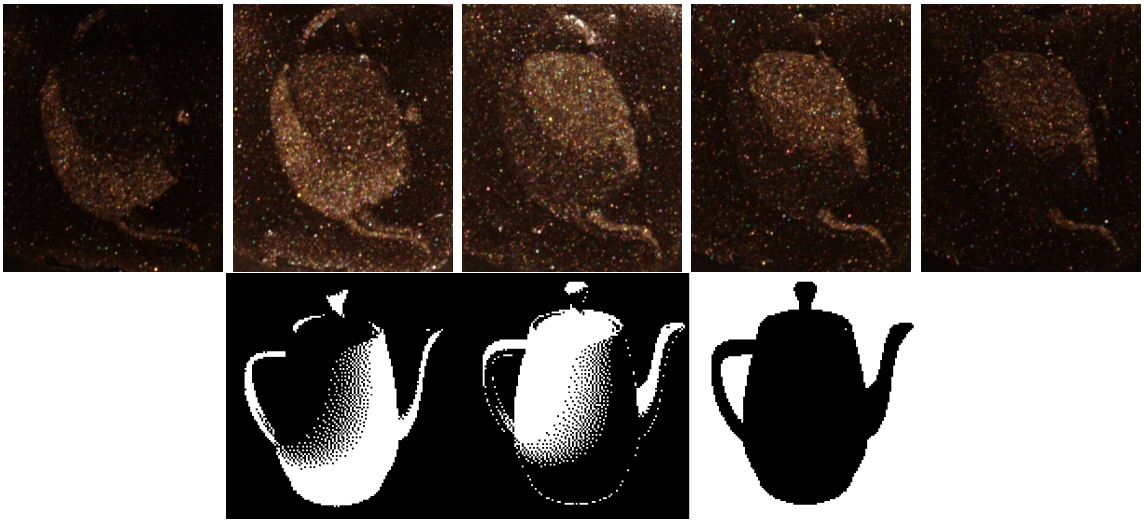


Figure 3.20: We printed a teapot with three different passes. We show it lit from different elevation angles (-55,-25,20,35,50). The masks had a different aspect ratio due to projector requirements.

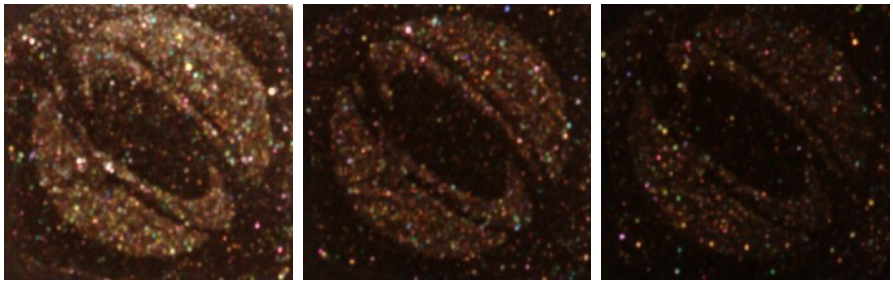


Figure 3.21: Siggraph logo printed with the spiral field.

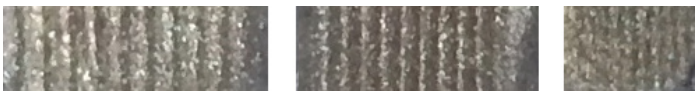


Figure 3.22: Resolution. From left to right, stripes of width 300 um, 250 um and 200 um. Our system shows good contrast at 250 um, but poor contrast at 200 um.

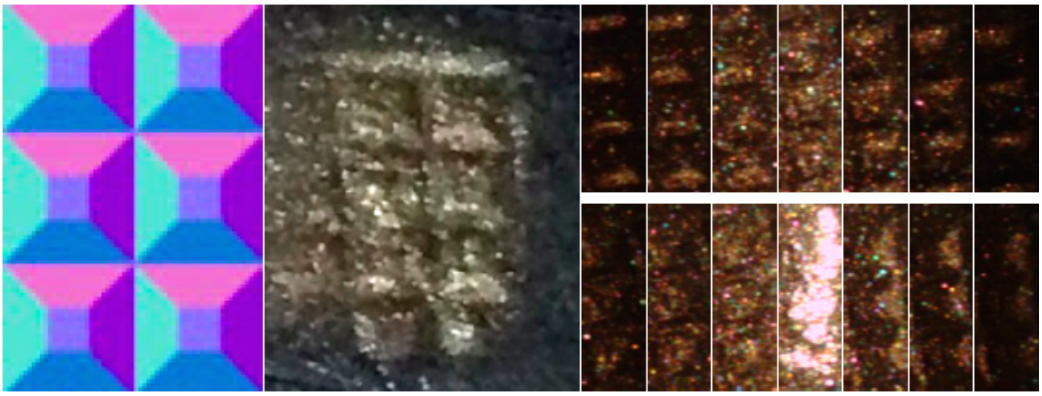


Figure 3.23: Six small pyramids. Left: seen under natural lighting. Bottom: lit from different elevation angles along the x axis (40,25,10,-20,-35,-50,-65). Top: lit from different elevation angles along the y axis (55,40,25,-20,-35,-50,-65)

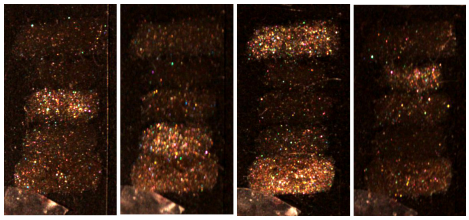


Figure 3.24: Five anisotropic BRDFs with tangents from top to bottom 0,45,90,135,180. We light them at an elevation angle of 45° and with azymuthal angle varying from left to right: 90, 135, 0, 45 degrees. These five samples were not printed with the projector. Each pixel was painted by hand and cured.

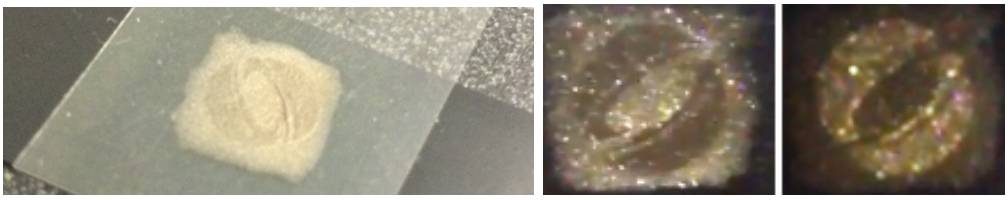


Figure 3.25: Siggraph logo where foreground and background have different anisotropic tangents. This print was created using the spiral field to randomize the flakes orientations followed by the line field to achieve the anisotropic distribution. This process was repeated for foreground and background. On the left, lit under a natural environment. On the middle and right: lit with a point light from different azimuthal angles at constant elevation.

3.8 Printed results

In this section, we show BRDF prints using different magnetic field schedules. We printed a Siggraph logo (Figure 3.21) using two different masks, where the logo has a specular BRDF and the background is made darker with the spiral field.

We printed a teapot (Figure 3.20) with three different masks. For the background we used the spiral field and two foreground masks were oriented with circular fields oriented up and down. We used dithering between the masks in this result which creates a smoother transition, but less smooth then expected. The curing process of white pixels seems to be leaking to neighboring black pixels.

To test this leaking process and the effective resolution of our system, we printed vertical stripes (Figure 3.22) where pigments are tilted to the left and to the right. We printed the stripes with different widths of 300 um, 250 um and 200 um. Our current setup shows good contrast at 250 um, but poor contrast at 200 um. This loss in resolution could be happening for multiple reasons. First, the image in the UV wavelength (385 nm) might be less focused than the one with pure visible light. Second, when projecting a white stripe, the chemical reaction might be propagating spatially to neighboring black pixels. Finally, the flakes themselves might be causing some light scattering to neighboring pixels.

We also printed a bump map with five different masks, one for each normal. Figure 3.23 shows them under natural lighting and also under controlled lighting.

We painted by hand five anisotropic BRDFs (Figure 3.24) with tangents from top to bottom 0,45,90,135,180. We light them at an elevation angle of 45° and with azymuthal angle varying from left to right: 90, 135, 0, 45 degrees.

We printed a Siggraph logo with different anisotropic tangents. Figure 3.25 shows the sample under natural illumination as well as lit from two different azimuthal angles at constant elevation. For these samples, we used the spiral field ($\tau = 0.45$) followed by the line field for each of background and foreground.

3.9 Conclusion

In this work, we have shown how magnetic reflective particles can be used to create spatially-varying anisotropic reflectance using a number of time-varying magnetic fields. We prototyped a projector-based printer that exposes the flakes to a time-varying magnetic field and then hardens them with UV-light. This setup was used to fabricate spatially-varying BRDFs.

One important direction of future work is investigating setups that use a print head with magnetic control instead of a projector. First, it would allow multiple layers of magnetic pigments. Specially at lower concentrations this would allow fabricating anisotropic subsurface scattering phenomena.

Second, this could be one direction towards 3D printing with rich appearance using magnetic pigments. Even using the pigments as a bump map on a 3D printed object could already allow for higher resolution details, just like bump maps in computer graphics.

In the context of DLP projector 3D printing, as in our setup, an interesting problem would be to actuate magnetic fields over a larger area (our setup only achieves uniform fields over 1×1 mm). This would not only allow for larger printed BRDF images, but also could enable grayscale DLP 3D printing.

Chapter 4

Computational Light Routing

4.1 Introduction

Despite recent advances there are still few fabrication techniques and algorithms that let us control how light propagates inside a solid object. Existing methods design surfaces that reflect [67, 35] and refract light [42, 14] or restrict themselves to reproducing light diffusion in solid objects [9, 22]. We present automatic object design algorithms that, coupled with 3D printing, let us fabricate complex objects with embedded optical fibers. These fibers let us control light propagation in objects, enabling novel display and sensing applications.

Printing optical fibers is made possible by modern multi-material 3D printers. We print two materials with different indices of refraction (Figure 4.1). The core material, where light propagates, has a higher index. A low-index cladding material surrounds the core. The difference in indices causes total internal reflection, allowing light to propagate with low loss even after multiple bounces.

Recently Willis et al. [68] have shown that many optical components can be custom printed including optical fibers. Using printed fibers they designed applications such as tangible displays in the form of chess pieces, custom sensors of mechanical

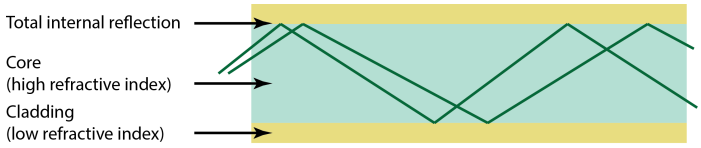


Figure 4.1: Total internal reflection happens because of the higher refractive index in the core. This allows good propagation of light inside an optical fiber.

motion and toys with dynamic eyes. However, the internal complexity of fiber volumes enabled by 3D printing, coupled with manufacturing constraints, led to a hard manual design problem.

We propose an automatic fiber routing algorithm that accounts for these constraints. It routes light between *arbitrary surfaces* enabling new applications of printed optical fibers. In entertainment applications, we created backlit face displays [33]. In scientific visualization, we created a brain-shaped display that let us visualize MRI data in context. In both cases, light from a micro-projector enters the fiber volume on a flat interface and is routed by the fibers to points on the surface (Figure 4.2). In addition, we prototyped sensing applications where fibers enable non-flat imaging surfaces. These include a touch sensitive display and a fiber hemisphere for BRDF acquisition. Finally, our algorithm could be used in photography applications in a design such as in Ford et al. [16] to optimize the coupling between the curved surface of a lens and a sensor.

In this work, we characterize the capabilities and limitations of printing optical fibers (Section 4.3) including the advantages of using different materials, how shape influences light propagation in the fibers and how to control the exiting light distribution. Next, we describe algorithms (Section 4.4) that take two arbitrary surfaces as input and design a volume of optical fibers to route light between them. Using our method, we design (Section 4.5) complex objects such as arbitrary shaped touch-sensitive displays and a hemispherical light distribution sensing component.

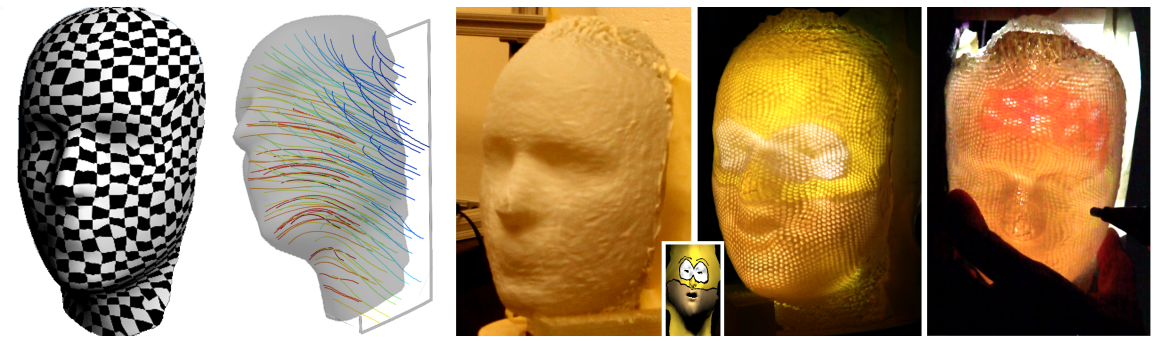


Figure 4.2: We use 3D printing to fabricate objects with embedded optical fibers that route light between two interfaces. We use this pipeline it to create displays of arbitrary shape, such as this animated face. Given a parameterized output surface (left), our algorithm automatically designs the fibers (middle-left) to maximize light transmission. We use a micro-projector to input an image (inset) on the printed object's (middle) flat interface, and it is routed to the surface (middle-right). We also present a painting application in which fibers are used for sensing and display. The light from a touch-sensitive infrared pen (right) is routed through the object to a camera.

4.2 Related work

Fabrication. Fabrication has introduced the graphics community to problems of how to control the interaction of light with physical objects. Most recent work focuses on the *surface* of objects, including reproduction of reflectance [35, 67] or refraction [42, 14]. Other work designs objects with a custom *interior* structure, but this has been restricted to reproducing subsurface scattering [9, 22]. Our printed fiber volumes provide control of light propagation through total internal reflection in the interior of objects.

Recently, 3D printed optical components [68] including optical fibers were shown to have many applications in human computer interaction such as motion sensors, volumetric displays and toys with custom flat displays. We present new applications of printed fibers enabled by using automatic design algorithms instead of manual design, such as display and sensing applications with non-flat surfaces. On the manufacturing

end, we discuss the light transmission trade-offs involved in using different materials and techniques to control the angular distribution of exiting light.

Traditional manufacturing. Fiber imaging applications depend on fiber bundles which are traditionally manufactured in two steps. First each individual fiber is made in isolation. In some methods, they are constrained to constant cross-section. In other methods, they are even constrained to cylindrically symmetric cross-sections [57]. Second these fibers are packed into a bundle which is then shaped. The process works by placing fibers together and applying heat to fuse them [28]. At this point a simple parallel bundle is obtained, its shape can then be changed by the simultaneous application of both heat and external forces. For example to create a taper, it's possible to stretch the bundle at its end points resulting in thinning at the middle where the object can be cut. To create an image inverter, a rotation can be applied to the end points resulting in fibers in a helix shape.

Printed fibers avoid this global deformation step, by placing material voxels with different refractive indices directly in their final position. Printing enables bundles of much more complex shape, such as bundles with the shape of a face. While it is possible to start with a traditional bundle and grind it to any shape, it would be hard to control the fiber arrival angle. This is made easy with printing. Printing also has no limitation regarding cross-sections which can easily vary along the fiber. Currently these benefits come at the cost of impurities in the fibers, but most important not perfectly flat interfaces between core and cladding resulting from voxel quantization which introduces light loss.

Inter-surface mapping. Given input and output surfaces, our routing algorithm creates fibers to connect them. There are a variety of methods for morphing surfaces [20]. While these morphing techniques create a sequence of intermediate steps through *time*, our approach can be seen as morphing through *space*. Our work differs from previous morphing methods because we are mostly interested in proper-

ties of the fibers themselves, such as curvature. In addition, we have to handle spatial manufacturing constraints such as keeping fibers from getting too close.

We solve this 3D geometric problem using variational methods [11]. This approach has been used for other geometric problems including surface design [37] and deformation [58]. [61] use thin-plate implicit functions to create a smooth morphing between two surfaces. Techniques more similar to our work include [4], who used triharmonic functionals for smooth shape deformation, and [25], who also used higher-order functionals for shape deformation. In our work, we use triharmonic functionals to have better control and continuity of curvature at the boundary conditions. In addition, we propose other objective terms to represent manufacturing limitations such as the minimum fiber spacing.

4.3 Fiber fabrication

In this section, we describe different fiber designs with which we have experimented. These include the use of different materials (Section 4.3.1) and cross-section geometries (Section 4.3.2). We also include an analysis of the fiber's field of view and how to improve it (Section 4.3.3). Finally, we measure how transmission varies with length and curvature (Section 4.3.4). The capabilities and limitations of fiber printing will guide the design of our algorithm (Section 4.4).

4.3.1 Choice of materials

An optical fiber is composed of two parts with different materials (Figure 4.1). The core is chosen to have a high refractive index, while the cladding has a low index. Most of the light propagation occurs in the core, with the light bouncing back when it hits the cladding due to total internal reflection. When designing optical fiber volumes for 3D printing, an important question is the choice of materials for core

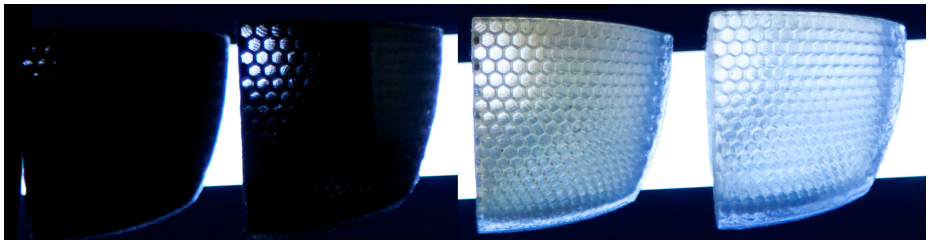


Figure 4.3: Comparison of fiber designs with different materials. For all designs, light is strongest when the camera direction is aligned with the surface normal. On the left, VeroBlack cladding transmits no light. On the right, TangoPlus as cladding shows good transmission in the fiber direction but also lets light pass even from other directions since it is transparent. This is a problem since leaked light from one fiber may be captured by other fibers. Support material as cladding (middle-right) absorbs some leaked light which reduces cross-talk. Even lower cross-talk can be achieved with a 3 material design (middle-left) by surrounding support cladding with VeroBlack. While no leaking occurs, the resulting transmission is very low. On the back, a laptop screen feeds white light into all these fiber bundles.

and cladding. In this subsection, we detail our experiments with different material combinations. For all experiments and applications in this work, we fabricated fibers using a Objet Connex 500 multi-material 3D printer.

We classify printed fiber designs according to different properties. The ones that depend on material selection are effective light transmission, cross talk between fibers and numerical aperture. The effective transmission depends on how absorbent are the core and the cladding and how much light is internally reflected. Cross-talk occurs when light crosses the fiber walls into a neighboring fiber. It is caused by irregularities on the fabricated fiber walls, but also depends on how much of the leaked light is absorbed by the cladding. The numerical aperture is the range of incident angles that lead to total internal reflection. Incident rays that are outside this acceptance angle range will mostly cross the fiber wall and leak. Numerical aperture depends mostly on the two refractive indices n_{co}, n_{cl} and is given by the following formula: $\sin(\theta_{max}) = \sqrt{n_{co}^2 - n_{cl}^2}$ where θ_{max} is the angle of incidence with respect to the fiber axis. The acceptance angle is identical to the exiting angle of light from a fiber, which we will also call its field of view.

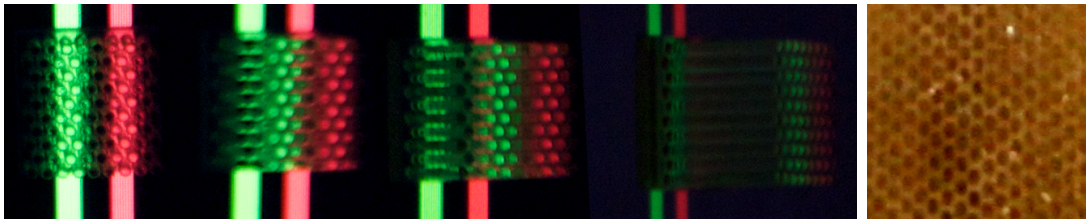


Figure 4.4: Air cladding fibers have 180 degree field of view and excellent transmission, but they are impractical since support material is added between the fibers during printing. This support material is very hard to clean for the intricate fiber geometries used in this work. From left to right, we see light imaged from increasingly grazing angles. On the back, a laptop screen feeds light into the bundle. At extreme right, we show the hexagonal cross-section at the base of our fibers. This shape improves tiling of the input image plane with no observable loss in fiber transmission.

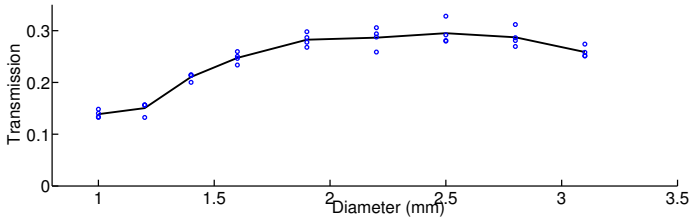


Figure 4.5: Transmission as a function of cross-section diameter (mm) measured for a fiber of 8 cm of bending radius. Multiple conflicting diameter-related factors influence transmission including number of light bounces and incidence angle resulting in a sweet spot.

Similar to [68], we chose VeroClear as the core. This is a transparent and colorless material with a high refractive index (1.47). We also experimented with FullCure 720, which has similar refractive index but has a yellowish color. The cladding material should have a lower refractive index. We experimented with air as cladding (Figure 4.4) and also with different printer materials (Figure 4.3). Both TangoPlus (a clear flexible material with index 1.42) and the printer's support material work well as cladding. Support material absorbs more light, leading to fibers that are more resilient to leaking. For this reason, we follow [68] and choose support as cladding. Both TangoPlus and support material have similar numerical aperture equivalent to a θ_{max} of 25 degrees.

4.3.2 Geometric factors

In this section, we describe our experiments with cross-section geometry and core/cladding diameter ratio. We experimented with three different cross-sections: circle, square and hexagon. The square cross-section does not work. To our surprise, both circle and hexagonal fibers showed similar transmission. We chose to use hexagonal fibers because the hexagon tiles the plane well, leading to a larger area dedicated to the core of fibers. Circular sections always have extra cladding spaces due to tiling limitations. In addition, hexagonal fibers reduce the number of mesh triangles needed to represent the geometry. This is a mundane, but practical, concern since printer software cannot handle models with very large triangle count. Besides cross-section geometry, the usable area also depends on the core/cladding ratio. Ideally, we would like to make the core very small and the cladding even smaller, but we are limited by the printer resolution. Most of the results shown in this work have a core diameter around 1mm and cladding spacing (distance between neighboring fibers) around 0.15 mm. Our experiments with cladding spacing down to 0.05 mm failed because this size is very close to the printer resolution (0.042 mm in the x or y direction).

We have measured how fiber transmission varies with cross-section diameter (Figure 4.5). For this experiment, we restricted ourselves to circular cross-sections and printed fibers in an s-shape with 8 cm bending radius in each arc (more setup details in subsection 4.3.4). We printed this fiber ranging in diameter from 1 to 3.1 mm. Multiple conflicting factors influence transmission. A larger diameter results in fewer light bounces and therefore less light loss. However, a larger diameter often results in bounces at a larger incident angle which brings more loss. This is because the fiber curvature becomes non-negligible. Overall, transmission seems to vary only up to a factor of 2 due to diameter which is a much smaller impact than the one we observe from fiber bending and length. Yet these measurements reveal a sweet spot around

2.5 mm. We leave as future work determining how this sweet spot changes with fiber curvature and length.

4.3.3 Field of view

As seen in Figure 4.3, the limited exiting angle of a fiber has large impact on its use for display purposes. It constrains the field of view to the directions which are approximately aligned with a pixel's surface normal. In this section, we discuss different ways to control the exiting light distribution with angle by making changes to the fiber ending. A solution to this field of view problem for flat displays is to use planar diffuser sheets, but this solution does not extend to arbitrarily shaped objects. After a few experiments, we chose to use a layer of diffuse white paint, which results in a 180 degree field of view. We experimented with applying white diffuse paint using spray cans and an air-brush. The air-brush provided finer control of the thickness of the applied layer of paint. This is important because the thicker the paint the more diffuse the appearance, but also the less light is transmitted.

Our first attempt at a diffuser was adding a layer of support material to the fiber endings. We also attempted to increase scattering by appending heterogeneous endings: small dithered combinations of support and VeroClear or TangoPlus and VeroClear. All these solutions failed to increase the field of view by a significant amount. Next, we printed small lenses at the end of the fibers. The lenses do extend the field of view by some small angle because rays that arrive outside the accepted angle range are mapped inside. However, some of the rays that were already accepted are bent to angles outside this range and therefore not transmitted. The net result is that from larger viewing angles only half a pixel is lit. We also experimented with using a layer of a white diffuse printer material: VeroWhite. While it averages the angular light distribution well, it also introduces some spatial blur due to subsurface scattering. After all these experiments, we settled on using white paint for diffuser.

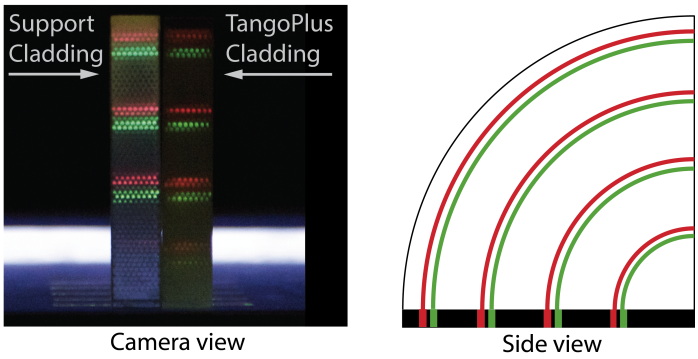


Figure 4.6: Comparison of the transmission of 90° turn fibers with varying curvature radii for both support and TangoPlus cladding. The radius of curvature of the top fibers is 5 cm.

In conclusion, we manufacture our fibers using the following design: VeroClear for core and support for cladding, hexagonal cross-section and white paint as diffuser.

4.3.4 Bending loss

In this section, we present an experimental evaluation of printed fibers. Light leaking is worse the higher the curvature. In our experience, short fibers with curvature radius of 5 cm or larger can transmit incident light very well while fibers with curvature radius smaller than 1 cm carry almost no light. Figure 4.6 shows the transmission of 90 degree turn fibers with varying curvature radii. Even though the length of the fibers is larger the larger the radius, we can see that loss due to curvature dominates in this example.

To measure fiber attenuation, we printed fibers in s-shapes (Figure 4.7). We also extend them by a straight segment so that many fibers have the same height. We can then focus a camera at this height making simultaneous measurements of multiple fibers with a single photo. Both arcs of the s-shape have the same radius. For each radius, we vary the arc angle resulting in varying total length but keeping the height fixed.

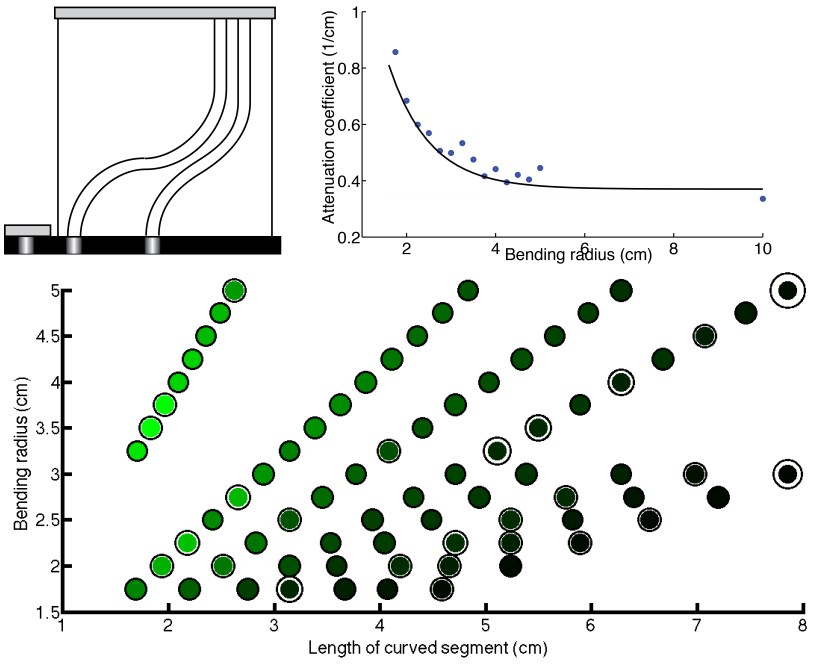


Figure 4.7: We measure how fiber transmission varies as a function of curvature and length. Our setup consists of multiple fibers in an s-shape of constant curvature that continue into a straight segment (top-left). We fit an offset exponential (top-right) to describe the fiber attenuation coefficient (1/cm) as a function of the bending radius (cm). On the bottom, we show our measured transmission values (shown in green) as a function of bending radius and length of its s-segment. The ratio of the radius of each two concentric circles is the measurement's relative fitting error.

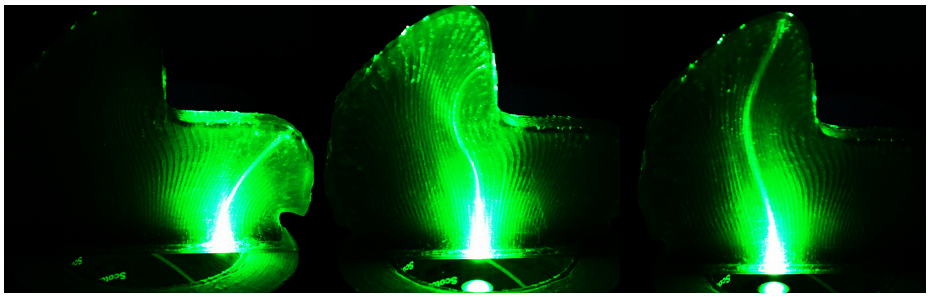


Figure 4.8: Light propagation inside a poorly designed object cross-section. Since we are imaging from the side what we actually observe is scattering along the volume. While some light arrived at its destination even for complex routes, much light is leaking and scattering through the volume. The middle image also shows how more leaking happened at a high curvature region. We see how many light rays escape along the tangent direction.

We measured two of these sets of fibers. The first set had a height of 6 cm and 6 radii values ranging uniformly between 1.75 to 3 cm. For each of these radii, we measure 7 different arc angles $(28°, 36°, 45°, 51°, 60°, 67°, 75°)$. The second set had a height of 7 cm and 8 radii values ranging uniformly between 3.25 to 5 cm. For each of these radius, we measure 4 different arc angles $(15°, 28°, 36°, 45°)$. Notice that low angle values were avoided to guarantee that there is no line of sight between the pixel on the laptop screen and the camera.

We perform the transmission measurements using an LCD screen as input. The two fiber volumes are placed on the laptop screen which is all black except for gaussian blobs at the fiber positions (Figure 4.7) top-left. These gaussians reduce aliasing caused from alignment problems. More light exits the fiber along its axis direction, this can cause problems due to slight off axis measurements. We use a transparent diffuser sheet (drawn in gray) to smooth the angular light distribution before measuring. We manually select the center and radius of each fiber in a high-dynamic range image and add the contribution of all pixels to achieve the final intensity value. The same process is performed for the screen pixels without any fibers on them (Figure 4.7 top-right). We simply place the diffuser sheet directly on the screen to get a measurement of the input light.

We assume the attenuation coefficient only depends on the bending radius R. $i_{out} = i_{in} \exp(-\alpha(R)l)$ where l is the length of a segment of constant radius in cm. Each measurement is not only influenced by the coefficient associated with its bending radius, but also by the straight segment loss. We performed two fittings to this data. First we assume $\alpha(R) = a + b \exp(-cR)$. This exponential decrease of attenuation has been commonly observed both in theory and practice with traditional optical fibers for both fibers with diameter comparable and larger than light's wavelength [19, 57]. We find $\alpha(R) = 0.37 + 2.5 \exp(-1.1R)$, where attenuation is measured in cm^{-1} and 0.37 corresponds to the straight fiber attenuation. This gives us a mean relative error

of 15% and 68 of the 74 data points have relative error less than 33%. While this error is non-negligible, these fitting results support the exponential attenuation model with length and the exponential decrease of α. Figure 4.7 shows the transmission of the curved part of our measured fibers where brighter green means higher transmission. These are compensated measurements where the loss of the straight segment was estimated. The length axis measures only the curved part of each fiber. We plot two concentric circles for each measurement, the ratio of their radius is our model's relative fitting error.

Second, we drop any hypothesis on $\alpha(R)$ and fit 15 bending attenuation values. Mean relative error decreases to 10% and now 68 of the 74 data points have error less than 20%. Figure 4.7 shows both the exponential model in black and the 15 attenuation coefficients in blue. We chose to plot the straight fiber attenuation coefficient as a point at radius 10 cm; this value is actually at infinite radius.

All these measurements confirm that transmission is greatly reduced as the bending radius decreases which motivates our algorithmic decisions in the next section.

4.4 Fiber routing

In theory, fibers carry light with no loss because of total internal reflection. In practice, however, limited printer resolution introduces a minimum spacing between neighboring fibers and printer voxel quantization introduces surface irregularities that cause light to leak while propagating (Figure 4.8).

In this section, we present an algorithm to design fibers to route light to a user provided surface. We present different objective functions and design constraints that enable the automatic design of complex bundles to route light. Our algorithm receives two parameterized surfaces as input and creates fibers to connect them. It maximizes light transmission by minimizing fiber curvature. In addition, it maximizes fiber

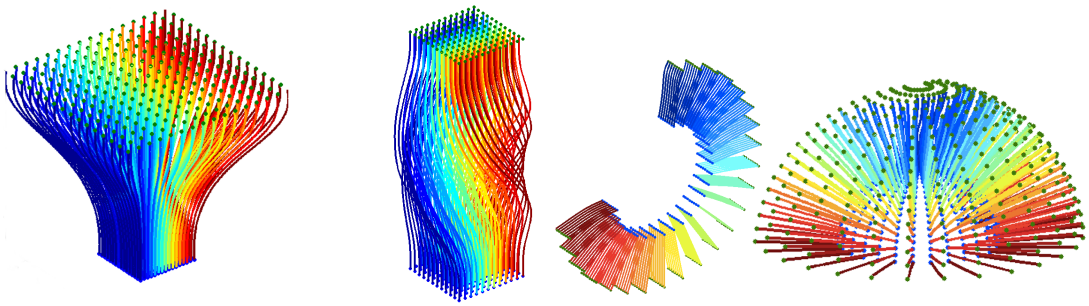


Figure 4.9: Sample results show curvature optimized routes while respecting user provided parametrization constraints. The input and output surfaces can be arbitrary as shown in these cylinder and sphere routings.

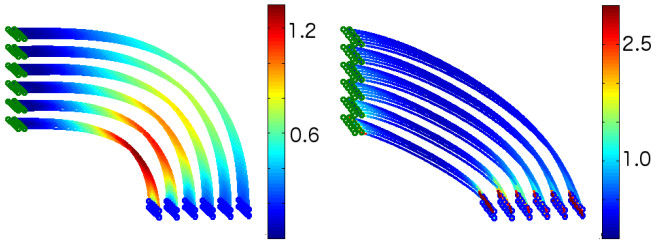


Figure 4.10: On the right, fibers generated by minimizing the thin-plate energy resulting in higher curvature in concentrated regions. On the left, minimizing the third derivative energy which results in more uniform curvature. Plots display color coded curvature at different scales.

separation and constrains fibers to arrive orthogonal to the input/output surfaces and conform to the user provided parameterization. We also show how to incorporate additional degrees of freedom into our optimization by automatically selecting a parameterization of the volume's flat interface (Subsection 4.4.3).

We propose an implicit formulation for the routing problem. Our algorithm receives as input both an input and an output surface, together with their u, v parameterizations. It then calculates u, v coordinates for every point in space by solving a variational problem. Each fiber can be seen as the set of points in space that have a given u_0, v_0 coordinate — in other words, a level set. In our current formulation, we solve for both u and v as separate optimization problems, so from now on we will only discuss u.

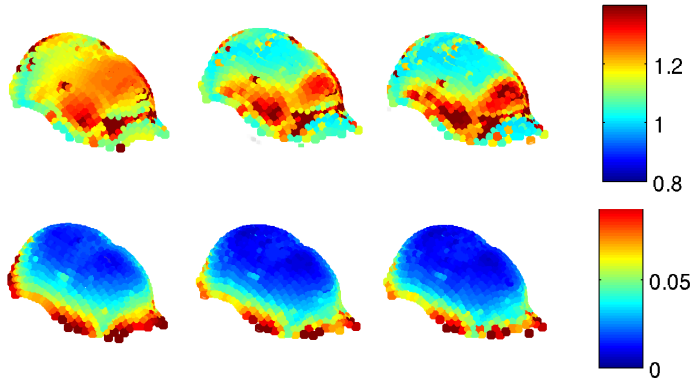


Figure 4.11: Effect of compression weight. In this example, we route fibers to a face. In the first row, we color-coded fiber endings with the max compression over the fiber path. In the second row, we show the mean curvature over the fiber path. From left to right, compression weights of 0, 0.1 and 1. Not only, mean compression, but also mean curvature decrease. The drawback was that maximum compression and maximum curvature both increased.

Figure 4.9 shows sample results of our algorithm for motivation. The green and blue points represent the input and output surfaces which can be arbitrary. We will denote the base by B and the target surface by Q. Each input point has an associated u coordinate, which we denote by $g(x)$. The algorithm extends these coordinates to all free space (displayed here as color).

The input parameterization serves as a hard constraint during optimization. We impose the additional hard constraint that the fibers arrive orthogonal to the surfaces. Every time light crosses between different media, there is a reflected and a transmitted ray. We constrain the arrival angle to maximize the power of the transmitted ray. As an extreme case, rays arriving at very high angles would be internally reflected and no light would exit the desired surface. This angle constraint can be achieved by constraining the derivative of $u(x)$ in the direction of the normal $n(x)$. Notice that this normal need not be the true normal of the surface. For instance, when routing complex models, we replaced it by smoothed normals, which still keep the arrival angle low but are less constraining.

Our objective function includes two different volumetric design goals: curvature $C(u)$ and compression $K(u)$. We formulate these objectives in the following quadratic program:

$$\begin{aligned} \underset{u}{\text{minimize}} \quad & C(u) + w_k K(u) \\ \text{subject to} \quad & u(x) = g(x),\ x \in Q \cup B, \\ & \nabla u(x) \cdot n(x) = 0,\ x \in Q \cup B., \end{aligned} \tag{4.1}$$

4.4.1 Curvature term

As we have shown in the previous section, both fiber curvature and length influence transmission. Yet we observe, while solving our arbitrary surface routing problem, that high curvature is often the main cause of high loss (Figure 4.6). This is partly due to the fiber end points being hard constraints in our formulation. For simplicity, we chose to minimize fiber curvature and we leave as future work optimizing actual fiber transmission.

We would like to obtain a set of fibers that has minimum curvature. However, the expression for curvature is non-linear in u. It involves the product of second and first derivatives. We solve instead linear proxies to curvature that depend on derivatives. We experimented with minimizing second derivatives (thin-plate energy):

$$\int u_{xx}^2 + u_{yy}^2 + u_{zz}^2 + 2u_{xy}^2 + 2u_{xz}^2 + 2u_{yz}^2 \tag{4.2}$$

but chose to use third derivatives instead:

$$C(u) = \int u_{xxx}^2 + u_{yyy}^2 + u_{zzz}^2 + 3u_{xxy}^2 + 3u_{xxz}^2 + \cdots. \tag{4.3}$$

Here the subscripts denote partial derivatives. The thin-plate energy resulted in fibers that were mostly straight but had high curvature in concentrated regions. As shown in Figure 4.10, the third derivative energy better distributes curvature along

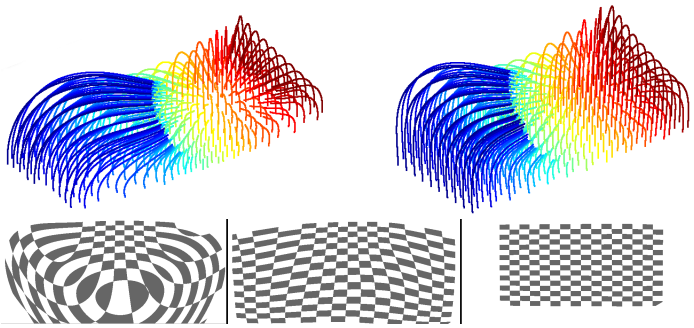


Figure 4.12: We add a stretch energy term to keep the base from growing too large in cases where that minimizes curvature. On the top, we show the results with stretch weight 0.0001 and 0.001. The lower area weight result's width was 40% larger. On the bottom, we show the base parameterizations for weight 0, 0.0001 and 0.001.

the length of the fiber, avoiding these high curvature values. All results in this paper were generated using the third-derivative energy.

4.4.2 Compression term

If two fibers get too close to each other, the resulting volume might not be manufacturable. The reason is that there is a minimum width of cladding a printer can actually print. We use the term *compression* to refer to many fibers coming together in a small area. In this section, we present an objective function to minimize average compression, or in other words, to maximize fiber spacing.

Let us start by considering the 2D case. In our implicit formulation, compression can be written as $|\nabla u|$. For example, when all fibers are going up then u_y is zero and this constraint reduces to $|\frac{du}{dx}|$. This is exactly the number of fibers du per unit of space dx. We write our quadratic objective as $K(u) = \int u_x^2 + u_y^2$.

How to extend it to 3D space? One might consider maximizing the area of the fiber cross-section $|\nabla u \times \nabla v|$. However, this could still lead to thin walls between fibers in one direction. We simply minimize both $|\nabla u|^2$ and $|\nabla v|^2$ independently. This also maintains the separable structure of our optimization problem. We can still solve two entirely separate problems one for u and one for v.

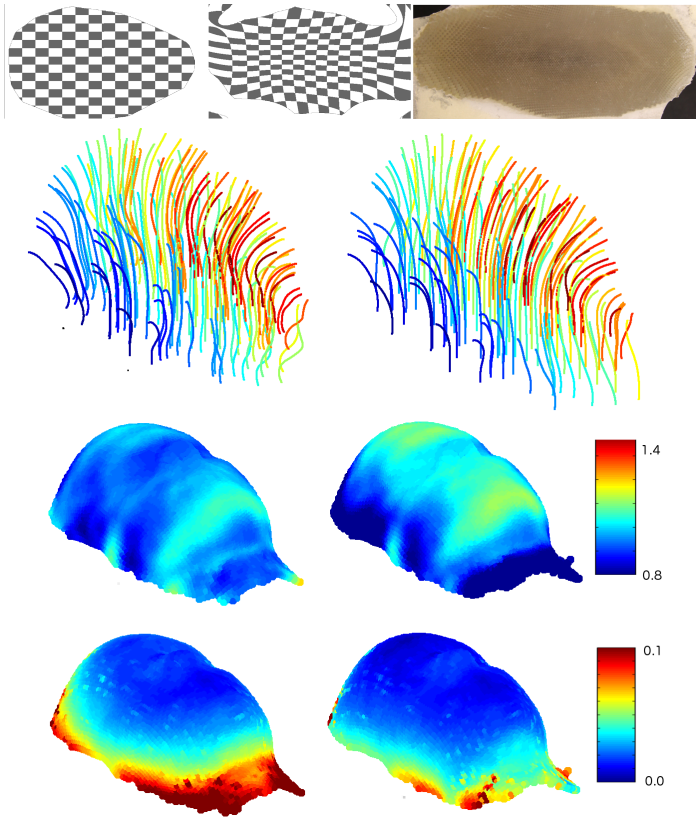


Figure 4.13: By optimizing the base parameterization jointly with fiber routing (right column), we obtain fibers with much less curvature (last row) and only slightly more compression (third row). On the top right, we show the base of the printed face model with its optimized fiber placement.

As an example of a result generated using our compression objective see Figure 4.11, where we route light to a face. In the first row, we color-coded fiber endings with the max compression over the fiber path. In the second row, we show the mean curvature over the fiber path. As expected, we see that the mean curvature decreases with higher compression weights (more green). Surprisingly, the mean curvature also decreases from 0.038 to 0.026. This happened for all models we experimented with and can be understood by remembering we are only optimizing an approximation to curvature. However, there are drawbacks. Maximum compression and maximum curvature both increase from 1.65 to 1.70 and from 0.54 to 0.65 respectively. As a

consequence, we usually keep low weights on compression to benefit from the decrease in mean values without hurting the maximum values much.

While adding these compression objectives gives us a useful design tool to balance between compression and curvature, we would also like to add more degrees of freedom to the algorithm. In the next section we discuss how to jointly optimize the volumetric routing and the base parameterization.

4.4.3 Base layout

So far we assumed the user provided parametric coordinates for the surface and the base. For some applications such as shape displays, the user is usually not interested in any particular base parameterization. This means that ideally the algorithm should be free to change this parameterization if this change results in more efficient routing. In this section, we describe how to incorporate this extra degree of freedom to improve routing. With this variation, our algorithm can automatically decide where to place the base and with which scale. It also allows for non-linear stretching of the base. The major drawback of changing the fiber cross-section non-uniformly at the base is decreasing the input energy of some fibers and therefore the display's dynamic range. For all points on the base surface, we replace the constraint $u(x, y) = g(x, y)$, where g was fixed, by a linear combination of basis functions:

$$u(x, y) = \sum_i \alpha_i \, h_i(x, y). \tag{4.4}$$

Fortunately, this is still a linear constraint. The only variables are the values of u and the coefficients α_i. In our current implementation, we chose a radial thin-plate spline basis $h_i(r) = r^2 \log(r)$ centered at a grid of control points. In addition to these splines, we include the affine basis $h_1(x, y) = 1, h_2(x, y) = x, h_3(x, y) = y$.

Figure 4.13 shows the effect of using base layout. In the first row, we see a visualization of the unoptimized and optimized base parameterization. In the second row, notice how the fibers on the right are more well behaved after their base positions are optimized.

The median of the maximum compression of each fiber increases slightly by 4.6%, while curvature showed large decreases. For example, the median of the mean curvatures decreased 39% and the median of the maximum curvatures decreased 29%. The mean compression of each fiber is visualized on the third row and the mean curvature on the last.

This method proved effective at finding smooth base parameterizations. Moreover, we observed that these converge quickly when the number of control points is increased. For grids of 8×8 points, we could usually get no more improvement. Therefore, all results generated in this paper used a grid of 4×4 points, which helps in keeping the linear system sparse. In the future, we would like to experiment with other families of warps in order to capture high frequencies in the optimized base parameterization.

As we chose a family of thin-plate warps of the base parameterization, there is no guarantee it is a bijection. As happened in this example, the optimized base parameterization may assign the same u, v coordinate to multiple points on the base plane. This was a not a big problem since the resulting parameterization tends to be well-behaved around positions corresponding to fibers. The only step of the algorithm where not having a bijection was a problem was while sampling the base to grow fibers. Currently, we uniformly sample in u, v and solve for the corresponding x, y position to start the fiber. When there are multiple answers we found it adequate simply to choose the one nearest to the projected center of the mesh.

Base layout introduced an undesirable side effect when designing an inward looking hemisphere that routes light from a plane (Figure 4.14, right). Both our energies force

the base to grow very large, since that reduces both curvature and compression (Figure 4.12, left). Since it is impractical to make these very large objects, we added an extra objective term to keep stretch low. This term works as a weak quadratic prior that pulls u_x, u_y, v_x, v_y towards their mean values on the target surface. Figure 4.12 shows how this term provides control over stretch. Both curvature and compression are volumetric terms, while stretch is an area term. We normalize all energies by volume and area respectively before adding them.

After the addition of the base layout constraints and the stretch energy term, our optimization problem is written below.

$$\begin{aligned} \underset{u}{\text{minimize}} \quad & C(u) + w_k K(u) + w_s S(u) \\ \text{subject to} \quad & u(x) = \sum_i \alpha_i h_i(x),\ x \in B, \\ & u(x) = g(x),\ x \in Q, \\ & \nabla u(x) \cdot n(x) = 0,\ x \in Q \cup B., \end{aligned} \tag{4.5}$$

Using this algorithm, we routed and printed a few different surfaces (Figure 4.14). The parameters used and some summary statistics of routing quality including curvature and compression are shown in Table 4.1. Execution time, number of fibers and voxels can be found in Table 4.2. In the next section, we show how these routings behave when carrying light.

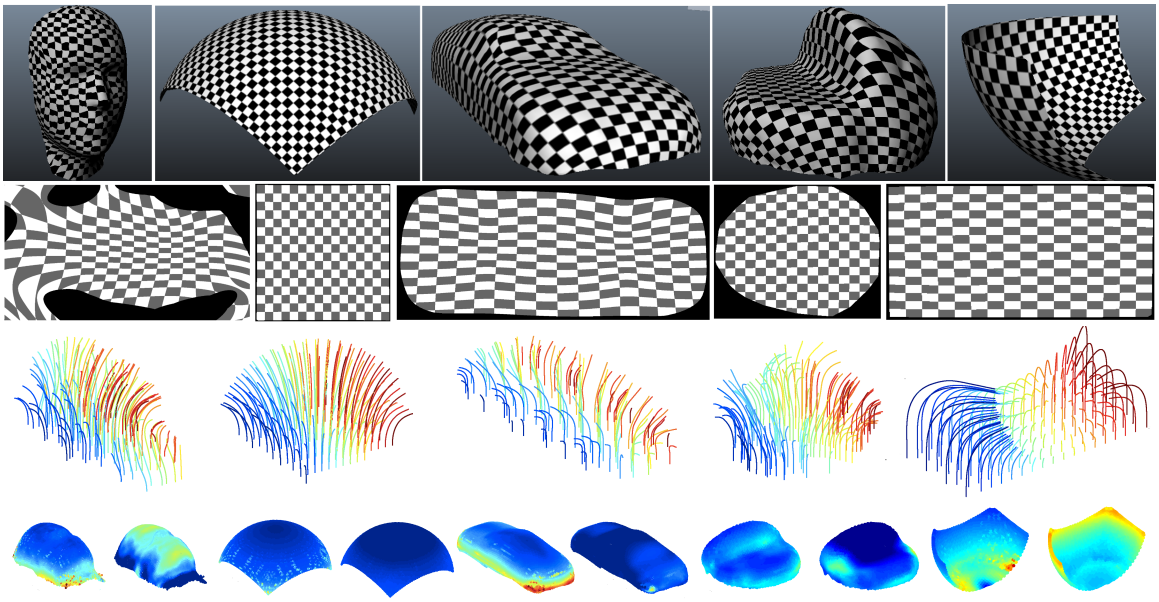


Figure 4.14: Table of routing results. First row, the input parameterized mesh. Second row, the optimized base parameterization. Third row, randomly selected fibers. Fourth row, for each model we show mean curvature and mean compression per fiber.

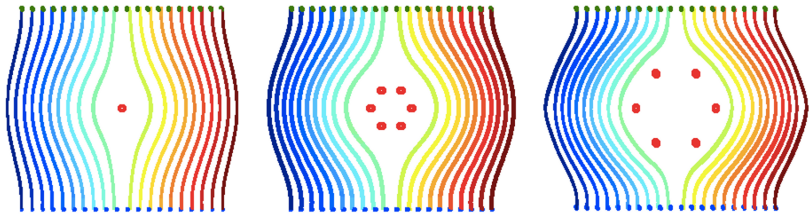


Figure 4.15: Routing around objects.

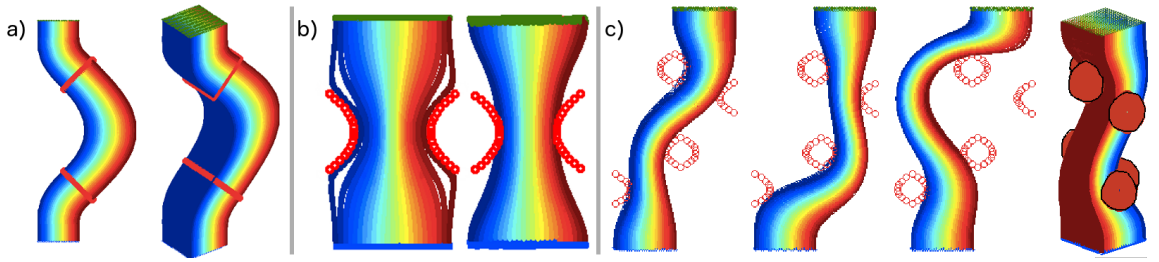


Figure 4.16: Routing with constraints. Relay points (a) can be added to constrain the path of individual fibers using equality constraints. Bottleneck constraints can be implemented with equality constraints but this may lead to low quality paths (b-left). We used inequality constraints to achieve a low curvature routing through a bottleneck (b-right). Multiple obstacles (c) can be specified simultaneously in both 2D and 3D.

Table 4.1: Routing statistics. Each provided statistic is first a mean or max operator over each fiber's path, followed by a median or max over all fibers.

Bundle	Compression / Area Weight	Median of Mean Curv	Max of Mean Curv	Median of Max Curv	Max of Max Curv	Median of Max Compr	Max of Max Compr
Face	0 / 0	0.022	0.072	0.073	1.05	1.16	1.55
Car	0.001 / 0.001	0.029	0.079	0.067	0.369	0.88	1.48
Brain	0.001 / 0.001	0.027	0.046	0.073	0.191	1.0	1.4
Inward	0 / 0.001	0.029	0.058	0.071	0.349	1.5	2.5
Sphere	0.01 / 0.001	0.013	0.034	0.022	0.838	0.94	1.40

Table 4.2: Timings and sizes.

Bundle	Voxels	Fibers	Optimize	Meshing
Face	45k	5k	335s	289s
Car	13k	7k	61s	391s
Brain	11k	5k	43s	300s
Inward	63k	2k	391s	97s
Sphere small	48k	2k	960s	62s
Sphere medium	48k	3.5k	960s	160s
Sphere	48k	7k	960s	473s

4.4.4 Path constraints

All routing examples discussed so far considered as constraints only the input and output surfaces. In fact, many important applications of printed fibers [68] require additional constraints on the fiber's path. For example, we may want fibers to guide light *inside* a character's body to its eye or head. Alternatively, some sensing applications of fibers may require avoiding other components such as buttons.

In this section, we show how additional linear equality or inequality constraints can be added to our optimization formulation to represent these path constraints. They provide more routing control to the user letting him specify relay points, bottlenecks, obstacles and even route around objects.

The simplest of these are relay points. If the user desires that a specific fiber u_0, v_0 pass through a certain point in space x, we can simply identify which voxel contains the point and add the linear constraint $u(x) = u_0, v(x) = v_0$. Figure 4.16-a shows a case where relay points were used to bend the fiber 45° to the right. This same case highlights limitations of our implicit method: for bending angles larger than 60° the resulting red fibers extended horizontally all the way to infinity, failing to connect input and output surfaces.

Equality constraints can also be used to route fibers around an object (Figure 4.15). This was accomplished by identifying the coordinates of the two middle fibers, let's say u_0, u_1, and adding the linear constraint $u(x) = .5(u_0 + u_1)$ to on or more middle points.

While equalities are useful, we found them to be rather limited. The problem is that in other settings, we just don't know a priori the coordinate value of a point in space. For example, we experimented with routing through a bottleneck (Figure 4.16-b) using equality constraints by setting the left wall's coordinate to 0 and the right's to 1. This led to very high curvature fibers since the leftmost and rightmost fibers stuck to the walls. We found inequality constraints to be much more expressive.

By adding wall constraints $u \leq 0$ to the left and $u \geq 0$ to the right, we can achieve a low curvature fiber that fits in the bottleneck.

As a final example, we route fibers though multiple obstacles (Figure 4.16-c). Ideally, we would have a forbidden region constraint. However, this is essentially an 'or' constraint. At the obstacle points, u should be either less than 0 or larger than 1 which is not a convex constraint. Instead, we request as input for each obstacle if the fiber bundle should go to its right or left. Different combinations give the user control over the routing. These obstacles also work in 3D. In Figure 4.16-c, we apply these constraints independently to u and v. Each obstacle can constrain u, v separately or both at the same time.

We attempted solving simple versions of routing using this non-convex constraint and a local optimization technique (interior point method). However, being a local method it fails to jump over obstacles and it simply gets stuck to whichever left/right combination it was initialized with. As such a local method is not better than simply fixing sides.

4.4.5 Implementation details

We solve our optimization problem by discretizing u into a grid and using finite differences to approximate all derivatives. This solution has one variable for each grid cell. Fortunately, since we are searching for very smooth fiber routings, small grids are sufficient to capture these low frequencies. Our problem is a quadratic program with linear equality constraints. Therefore, it is efficiently solved with a sparse linear system obtained using Lagrange multipliers. The only exception is when we use inequalities for path constraints. In this case, we use an interior point method. All our implementations were performed in Matlab.

We grow the fibers by integrating the curves from their initial positions in the input surface along the tangent field $t = \nabla u \times \nabla v$. We chose integrating fibers for simplicity

of implementation. A more robust solution would be to extract the isocurves using interpolation of the grid data. In practice, we did not observe significant drifting even with this simple integration solution.

To create the final mesh, we create hexagonal cylinders for each fiber. The remaining question is choosing the cross-section orientation and shape at each point. Since $\nabla u, \nabla v$ might not have the same norm and might not be orthogonal, our fibers have anisotropic cross-sections. We observed very small loss in transmission for small deviations from uniform cross-section. This justifies not imposing additional cross-section uniformity constraints which would hurt other objectives.

In addition, when printing hexagonal fiber cross-sections, it is important to have a smoothly changing basis of the plane orthogonal to the fiber. We use ∇u to define one axis and choose the other one as $\nabla u \times t$, where t is the tangent. In fact, even when printing circular cross-sections, this smooth basis helps in connecting adjacent sections into triangles. For a given point on a fiber, after determining the local frame, we search for the 6 closest points, one in each of the adjacent fibers. With knowledge of these points we can create the local anisotropic cross-section without intersecting the neighboring cylinders.

Table 4.2 shows our algorithm runtime split in two stages: optimize and meshing. Optimization time is the solution of the linear system and depends only on the number of voxels. Meshing includes time for fiber growth and mesh generation. The time it takes depends only on the number of fibers. This is a big advantage of our implicit algorithm: the same optimized implicit function can be used to route fibers at multiple resolutions (spheres in Table 4.2).

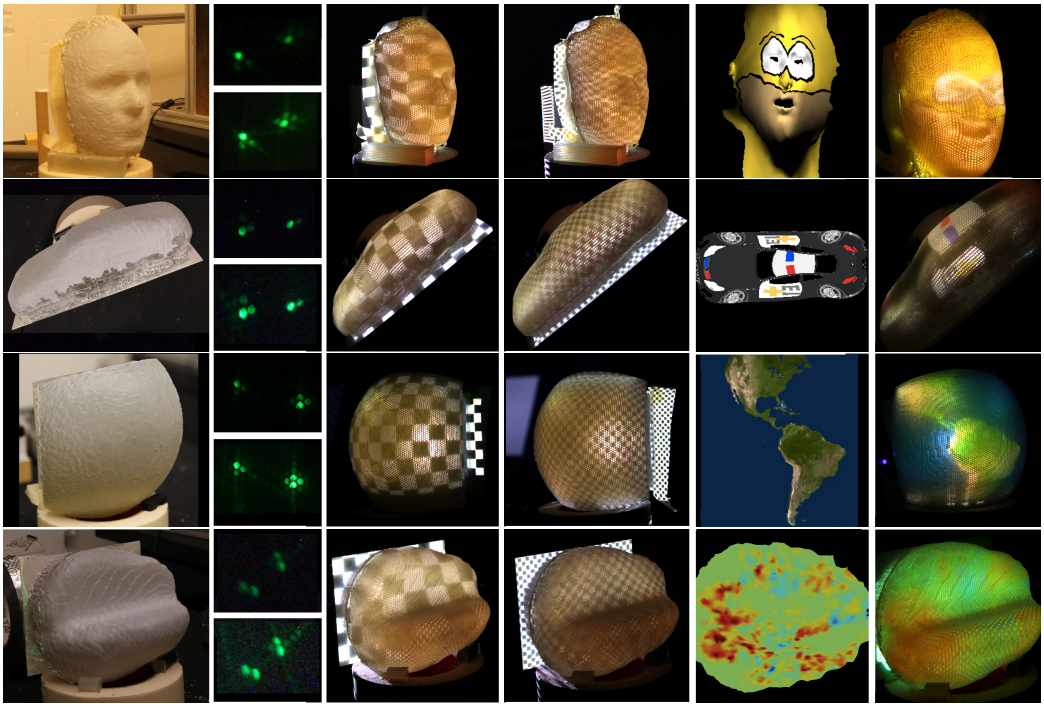


Figure 4.17: Table of displays. First column, our printed displays. Second column, representative point spread functions imaged by lighting a few single fibers in the back. We show two versions of each PSF picture: a non saturated and a saturated version. The saturated image shows the dark tail of these distributions. Third and fourth columns, checkerboard images show how these objects can display sharp images but with low contrast. Fifth column shows the projected textures and the sixth column the result as seen on the surface.

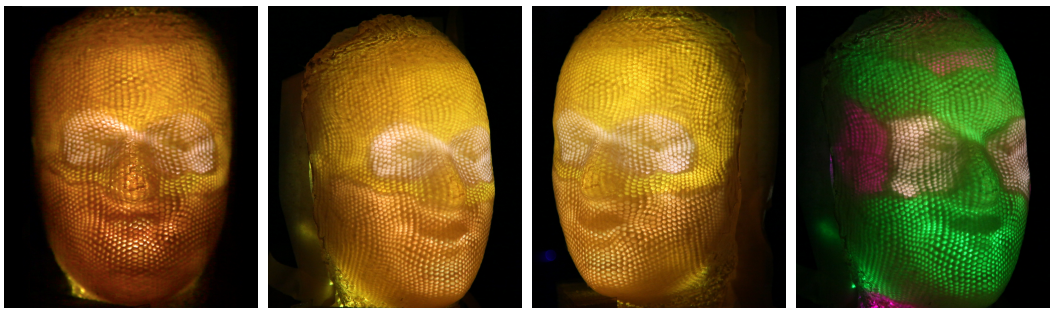


Figure 4.18: Face seen from different viewpoints and animation frames including open and closed mouths.

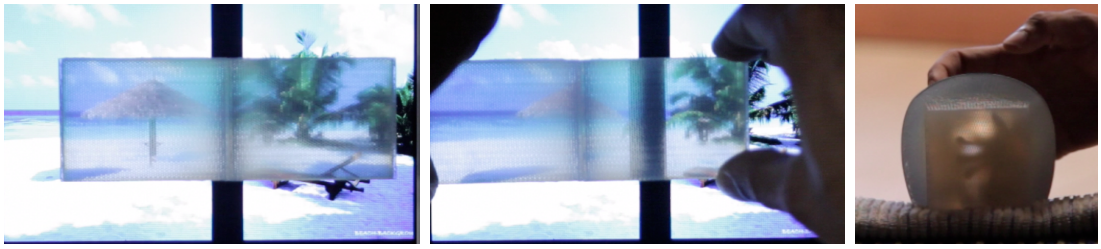


Figure 4.19: The left image shows a tiling of the two input images. The black band is not seen from this viewpoint. It is currently not seamless because the relayed image is also darker. In the middle image, we move the tiling component. Notice how the black band is now very visible and the tent pole is not. The right image shows how a half-sphere fiber volume with no diffuser can be used as a fish-eye lens. It can implement complex projections; shown here is a stereographic mapping. Notice how nearby objects (a hand) are imaged through these fibers.

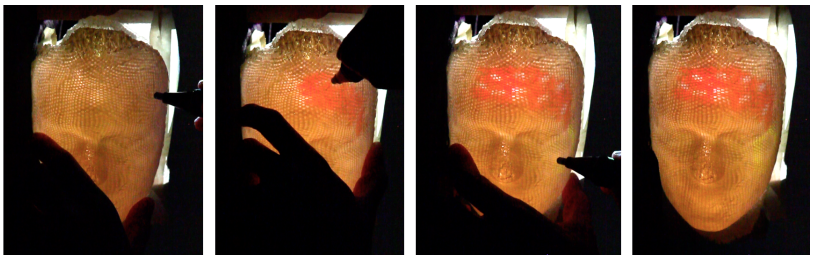


Figure 4.20: Painting on a face using a touch-sensitive IR LED pen.

4.5 Applications

4.5.1 Display

We can print displays of arbitrary shape by connecting projector pixels to surface elements. We use a micro-projector to input an image on the fiber bundle's flat interface, and route light to the surface where it meets the diffuse paint for a complete field of view. Previous work on displaying animated content on surfaces can be divided in those that project from the outside [51, 52] and those that project from the inside [33]. When projecting from outside, a single projector is usually not enough for complex objects. Because not all points may be visible from a single viewpoint, multiple projectors are often used [51]. This projection technique has a long history

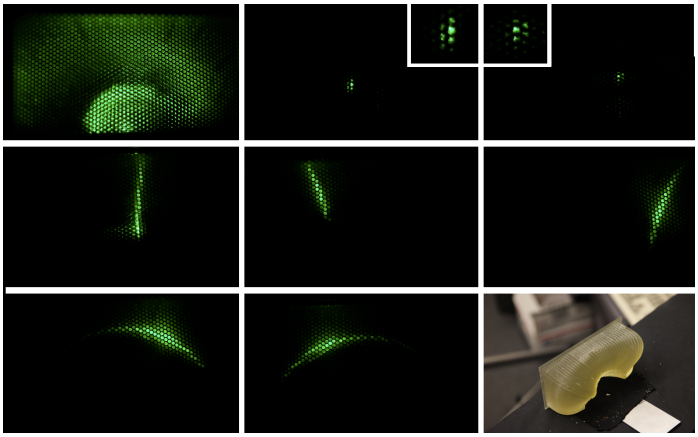


Figure 4.21: Image of the hemispherical light distribution exiting different small material samples when illuminated by a directional light source. We use black cloth to mask the input material leaving only a small sample located at the center of the hemisphere (lower right). The angular light distribution is routed by the fibers to the camera's focal plane. In the first row, white diffuse paper and white specular paper illuminated from two different directions (insets show enlarged images of the captured highlights). The next five images show specular highlights of anisotropic metal.

in entertainment including the singing heads at Disney's Haunted Mansion and light shows on architectural monuments. The major drawback of all these techniques is that they require unobstructed clear air between the projectors and the object. For indoor settings, this prevents the viewers from touching or even getting close to the display. For outdoor settings, scattering of the light rays caused by fog may reveal the technique. To avoid these limitations, [33] project from the inside of the object using a wide-angle lens. Lenses restrict the maximum bending angle of light and also restrict the display geometry to star-shaped objects. Our displays use fibers to carry light all the way to the surface pixel, enabling more light bending and more complex display shapes.

For each result in this section, the input parameterized mesh and the associated routing can be seen in Figure 4.14. The printed displays and an evaluation of their point spread function and contrast are shown in Figure 4.17. This figure also shows

the warped input image we project on the flat side and how it is seen on the display side. Table 4.2 shows the number of fibers for each display.

We printed a face display to project animated characters making different facial expressions. Figure 4.18 shows its appearance with varying content. It also shows the display seen from multiple viewpoints, which shows how the diffuse layer of paint worked well. In particular, from a fixed viewpoint, even points on the surface contour are visible. This display shows light bending by 90 degrees (side of face). It shows how we can achieve very concentrated point spread functions, which lead to sharp edges and adequate contrast. In fact, this face display shows how the contrast on the side of the face is actually better than in the front. The first reason is some of the light that was to be routed to the side escaped through the tangent and arrived at the front of the face (see checkerboard images). Another reason is that some light enters the cladding instead of the core. These rays traverse the object and are not routed, eventually hitting the front of the face. We plan to investigate the use of masks to block this incoming light.

We also printed a half-sphere display that turns light by almost 90 degrees. We projected a rotating globe animation. For this sphere, the layer of applied paint was not thick enough. As a result, more light exists the surface in the normal direction than at grazing angles. This can be seen by the increased brightness in the middle of the globe picture. Notice that this cannot be fixed by calibration, since it is a view dependent effect; it could only be fixed by a thicker layer of diffuse paint.

We printed a car display that allows us to dynamically change its content. We projected cars of different colors and blinking lights. While the image at the middle of this car display is very sharp (even for the 90 degree turns at the door), the front and back of the car are blurry. This happened because the car was the most elongated display we printed (16 cm). When we aim the projector at the center of the flat interface, light rays with destination at the front and back enter the fibers with high

arrival angles. The image at these fibers was sharper when centering the projection there. This could be improved by placing the projector at a larger distance or even by printing displays with non-flat input interfaces.

Finally, our last display has the shape of a brain with half a lobe removed to reveal its inside. We use it to display volumetric f-MRI data. Each point on the surface of our display is assigned the color of the corresponding point in the MRI volume. This shows how our technique can handle even very concave display surfaces. However, this comes at the expense of dynamic range, since transmission of different fibers varies considerably. We show uncalibrated checkerboard images and a manually calibrated MRI image which slightly improved the result. There are two reasons for the varying dynamic range of fibers. First, because of the concave region, the surface parameterization of the brain mesh is much less uniform than the other models (compare the pixel sizes on the left and right lobes). Second, fibers that arrive at the cutting plane took much more complex curved paths than other fibers (see Figure 4.14). The brain is the only result in this paper where we used intensity calibration.

We also printed simpler fiber volumes where both interfaces are flat. Using fibers, we prototyped merging two images where we route the light from the physical displays to the final display plane (Figure 4.19). The top image shows a tiling of the two input images. The black band cannot be seen from this viewpoint. It is currently not seamless because the relayed image is also darker and blurrier. In future work, we would like to use fibers to route light around an object making it 'invisible' from a given view direction.

The content for the face, sphere and car displays was generated by positioning the display surface mesh near a moving or deformable mesh and transferring the shading attributes at each frame. This process creates an animated texture that can be projected at the object's base. Notice that the parameterization used is not the user provided one, but instead the new mesh parameterization induced by the base

layout optimization. Both the car and face textures highlight these distortions. For the brain, we simply intersect the display surface with the volumetric image.

4.5.2 Sensing

While fibers can be used to carry light from a projector to a display surface, they can also be used to carry light in reverse. We designed two sensing applications of fibers: touch-sensitive painting on surfaces and hemispherical light distribution acquisition.

Touch-sensitive painting. Simultaneous sensing and display through the same fiber can be challenging. Projected light is usually much stronger, which compromises the signal to noise ratio of the sensing measurements. For interactive painting, we deal with this problem by sensing infrared light. It is important to use near-infrared, since our experiments showed our printed fibers are not capable of carrying wave-lengths farther from the visible spectrum.

For our setup, we used a pen that has an LED at its tip. It emits light when pressure is applied to the tip, such as when touching the surface. The input light is then carried in the reverse direction and lights the appropriate pixel on the object's flat side. While there is some leaking onto neighboring fibers, this was not significant. We have a camera placed near the projector with an infrared filter to detect which fiber was lit. The infrared filter worked well and the projector emitted almost no infrared light; as a result, the signal to noise ratio was extremely good. Both the camera and projector have to stay close to each other such that both fall inside the numerical aperture of all fibers. This was not an issue in our setup with the face as a display. Figure 4.20 shows this system in action. We start with a skin color canvas and draw a red and white head band together with a big yellow earring. Notice how the pen is detected even for fibers that curve 90 degrees. Painting on physical objects was previously explored with a head mounted display and a separate haptic interface

for detecting position of the pen [1]. Our setup allows displaying and sensing through the same object and only requires projectors and cameras.

Hemispherical light distribution acquisition. We designed a set of fibers to route light between a plane and an inward looking hemisphere. As a proof of concept, we applied it to capturing the angular light distribution coming from a material sample. This is an important component in BRDF acquisition. Our setup and results are shown in Figure 4.21. Black cloth is used to isolate a sample of three different materials including white diffuse paper, white specular paper and anisotropic metal. We light the material with a laser pointer from multiple directions and image the response on the flat side. No diffuser material is used for this application, since we do not need an extended field of view. In fact, we would rather have a shorter field of view. The images acquired by our setup can capture the shape of the different material highlights and their orientation. Our assembly provides good angular coverage at the cost of introducing some blur. We believe these qualitative results are promising. In future work, our transmission model could be used to calibrate the hemisphere allowing quantitative evaluation.

Besides sensing light, our first experiments suggest that it is also possible to use such a hemisphere to illuminate the material sample. We leave as future work a setup that can illuminate from any direction on one hemisphere and capture from any direction in the other hemisphere, thus enabling full BRDF capture. This would build on prior work in no-moving-parts BRDF acquisition including systems such as that of Dana [7], which combined a parabolic mirror and beam splitter to allow BRDF acquisition without requiring either camera or light source to move around the sample. However, this system was limited in its angular coverage. Ben-Ezra et al. [3] built a BRDF measurement system based on LEDs placed on a hemisphere, using them for both illumination and sensing. While this allowed for BRDF acquisition with no moving parts and full angular coverage, the sampling rate was limited because of the

need to place discrete LEDs on a hemispherical dome. Our 3D printed hemisphere-to-plane optical fiber assembly holds the potential for no-moving-parts, full-coverage BRDF acquisition, with an angular sampling rate limited only by the ratio of fiber width to overall printed hemisphere size.

Printed fibers should have additional applications in photography. As a very simple example, an outward looking sphere behaves as a fish-eye lens. It can implement complex projections with a passive setup. Shown here, a stereographic mapping (Figure 4.19, right).

4.6 Limitations

For all our applications, the major limitation is resolution. For applications where material cost is not a problem, bigger object sizes can be used to gain more resolution. We also expect printer resolution to increase in the future. Contrast is also currently a limitation, since some light leaks from the fibers. As more printable materials become available this should improve contrast in two ways, by increasing the difference in refractive index and reducing leaking and by blocking leaked light with highly absorbing materials.

A limitation of our implicit algorithm is that there is no guarantee that a fiber that starts in the input surface will connect to the output surface. In our experience this is a rare event, but it is possible that a fiber leaving the bottom surface simply loops back on itself or extends to infinity. For future work, we could consider an explicit algorithm that represents fibers as spline curves, we could then easily constrain it's end points and its topology, but things such as compression might be harder to optimize. After some experiments, we found another limitation: our algorithm does not handle discontinuous u, v mappings between the two surfaces. Continuity imposes the restriction that a fiber has to end near its neighbors at the base. Finally, we

would like to investigate how to minimize actual fiber curvature or even maximize fiber transmission directly.

4.7 Conclusion

In this work, we showed how a combination of automatic design algorithms and new manufacturing techniques enable new applications of optical fibers. Our algorithm employs an implicit representation of fibers to minimize curvature and compression while conforming to the user provided parameterization and constraining fiber arrival angles. We showed new applications of fiber printing such as touch-sensitive displays of arbitrary shape and acquisition of a hemispherical light distribution.

In future work, we would like to investigate other applications of fibers. For instance, it should be possible to couple printed lenslets and fibers to route an input light field. Beyond relaying light, we could design optical fiber networks where fibers split and merge enabling linear optical computation.

Chapter 5

Conclusion

This thesis has presented several novel methods related to the computational fabrication of appearance. These methods highlight how algorithms can help solve a number of different problems when fabricating objects that interact with light in a specified way. These included modeling perception, object optimization and extending fabrication capabilities.

Chapter 2 showed how computational proxies of human *perception* can guide ink selection when printing BRDFs. We introduced an improved BRDF similarity metric that builds on both experimental results on reflectance perception and on the statistics of natural lighting environments. We validated it quantitatively as well as through a perceptual study. In addition, we also showed how to adapt traditional color gamut mapping methods to svBRDFs to preserve spatial details like textures and edges.

Chapter 3 showed how using magnetic pigments can extend the capabilities of apperance *fabrication* enabling anisotropic BRDFs. We used dynamic magnetic fields to achieve finer control of printed BRDFs and included a qualitative and quantitative discussion of the BRDFs generated by different fields. We proposed a fabrication setup that uses electromagnets together with UV light projection to create spatially varying BRDFs.

Chapter 4 showed an example of the use of *optimization* to design objects while respecting manufacturing constraints. We introduced an automatic fiber design algorithms together with new manufacturing techniques to route light between two arbitrary surfaces. Our methods enables new applications in sensing and display such as surface displays of arbitrary shape, interactive touch-based painting and hemispherical light sensing.

5.1 Future Work

The ultimate goal of appearance fabrication is to create a digital fabrication system that can create a copy of any given object. Achieving this goal will require not only more general machines but also a more general gamut mapping process. This general process should be able to approximate a desired object taking into account the machine limitations but at the same time achieving an indistinguisable reproduction of the input object. The works in this thesis suggest very interesting research directions that could bring us closer to this ultimate goal.

Chapters 2 and 3 proposed perception similarity and a fabrication method for BRDFs. Extensions of these to include subsurface scattering (BSSRDF) would be a promising direction. A multi-layer sample of lower density anisotropic magnetic pigments should be a very flexible primitive to create very general subsurface scattering effects. More advanced BSSRDF fabrication pipelines coupled with a better understanding of perceptual similarity might lead to advances in realistic digital fabrication of skin, a very important application.

Our algorithm for light routing with printed fibers (Chapter 4) also suggests much more general future light transport algorithms. Given multiple fabrication primitives like lenses, mirrors and fibers fabricated into a single object, a challenge would be

how to computationally combine them in order to fabricate an object that achieves a prescribed light transport.

Moving beyond appearance, extensions of our routing algorithm could be used to route printed wires for applications in printed electronics. In addition, the use of different magnetic pigments coupled with magnetization during the printing process could point to ways of printing 3D objects that have customized magnetic domains thoughout their structure.

Bibliography

[1] Maneesh Agrawala, Andrew C. Beers, and Marc Levoy. 3d painting on scanned surfaces. In *Proceedings of the 1995 symposium on Interactive 3D graphics*, I3D '95, pages 145–ff., New York, NY, USA, 1995. ACM.

[2] Marc Alexa and Wojciech Matusik. Reliefs as images. In *ACM SIGGRAPH 2010 Papers*, SIGGRAPH '10, pages 60:1–60:7, New York, NY, USA, 2010. ACM.

[3] Moshe Ben-Ezra, Jiaping Wang, Bennett Wilburn, Xiaoyang Li, and Le Ma. An LED-only BRDF measurement device. In *Proc. CVPR*, 2008.

[4] Mario Botsch and Leif Kobbelt. Real-time shape editing using radial basis functions. In *Computer Graphics Forum*, pages 611–621, 2005.

[5] Eric Brockmeyer, Ivan Poupyrev, Moshe Mahler, Joanne Dauner, and James Krahe. Papillon: Expressive eyes for interactive characters. In *ACM SIGGRAPH 2013 Emerging Technologies*, SIGGRAPH '13, pages 12:1–12:1, New York, NY, USA, 2013. ACM.

[6] CIE. Guidelines for the evaluation of gamut mapping algorithms. In *Technical Report CIE 156:2004*, 2004.

[7] Kristin J. Dana. BRDF/BTF measurement device. In *Proc. ICCV*, 2001.

[8] Despland, Degott, and Muller. Method and means for producing a magnetically induced design in a coating containing magnetic particles. *European Patent*, 2005.

[9] Yue Dong, Jiaping Wang, Fabio Pellacini, Xin Tong, and Baining Guo. Fabricating spatially-varying subsurface scattering. *ACM Trans. Graph.*, 29:62:1–62:10, July 2010.

[10] R.O. Dror, T.K. Leung, E.H. Adelson, and A.S. Willsky. Statistics of real-world illumination. In *Computer Vision and Pattern Recognition.*, 2001.

[11] Jean Duchon. Splines minimizing rotation-invariant semi-norms in sobolev spaces. In *Constructive Theory of Functions of Several Variables*, Lecture Notes in Mathematics. Springer Berlin / Heidelberg, 1977.

[12] Mark D. Fairchild. Color appearance models. John Wiley and Sons, 2005.

[13] I. Farup, C. Gatta, and A. Rizzi. A multiscale framework for spatial gamut mapping. *Image Processing, IEEE Transactions on*, 16(10):2423 –2435, oct. 2007.

[14] Manuel Finckh, Holger Dammertz, and Hendrik P. A. Lensch. Geometry construction from caustic images. ECCV'10. Springer-Verlag, 2010.

[15] Roland Fleming, Ron O. Dror, and Edward H. Adelson. How do humans determine reflectance properties under unknown illumination. In *In Proceedings of CVPR Workshop on Identifying Objects Across Variations in Lighting: Psychophysics and Computation*, pages 347–368, 2001.

[16] Joseph Ford, Igor Stamenov, Stephen J. Olivas, Glenn Schuster, Nojan Motamedi, Ilya P. Agurok, Ron Stack, Adam Johnson, and Rick Morrison. Fiber-coupled monocentric lens imaging. In *Imaging and Applied Optics*, page CW4C.2. Optical Society of America, 2013.

[17] Martin Fuchs, Ramesh Raskar, Hans-Peter Seidel, and Hendrik P. A. Lensch. Towards passive 6d reflectance field displays. *ACM Trans. Graph.*, 27(3):58:1–58:8, August 2008.

[18] Daniel Glasner, Todd Zickler, and Anat Levin. A reflectance display. *ACM Trans. Graph.*, 33(4):61:1–61:12, July 2014.

[19] D Gloge. Bending loss in multimode fibers with graded and ungraded core index. *Applied optics*, 11(11):2506–2513, 1972.

[20] Jonas Gomes, Lucia Darsa, Bruno Costa, and Luiz Velho. *Warping and morphing of graphical objects*. Morgan Kaufmann Publishers Inc., 1998.

[21] Michael Guthe, Gero Müller, Martin Schneider, and Reinhard Klein. Btf-cielab: A perceptual difference measure for quality assessment and compression of btfs. *Computer Graphics Forum*, 28(1):101–113, February 2009.

[22] Miloš Hašan, Martin Fuchs, Wojciech Matusik, Hanspeter Pfister, and Szymon Rusinkiewicz. Physical reproduction of materials with specified subsurface scattering. *ACM Trans. Graph.*, 2010.

[23] Roger D. Hersch, Fabien Collaud, and Patrick Emmel. Reproducing color images with embedded metallic patterns. *ACM Trans. Graph.*, 2003.

[24] Matthias B. Hullin, Hendrik P. A. Lensch, Ramesh Raskar, Hans-Peter Seidel, and Ivo Ihrke. Dynamic display of BRDFs. In Oliver Deussen and Min Chen, editors, *Computer Graphics Forum (Proc. EUROGRAPHICS)*, pages 475–483, Llandudno, UK, 2011. Eurographics, Blackwell.

[25] Alec Jacobson, Tino Weinkauf, and Olga Sorkine. Smooth shape-aware functions with controlled extrema. *Comp. Graph. Forum*, 31(5), 2012.

[26] Kashiwagi and Tamura. Painting with magnetically formed pattern and painted product with magnetically formed pattern. *US Patent*, 1994.

[27] R. Kimmel, D. Shaked, M. Elad, and I. Sobel. Space-dependent color gamut mapping: a variational approach. *Image Processing, IEEE Transactions on*, 2005.

[28] Longin Kociszewski, Dariusz Pysz, and Ryszard Stepien. Double crucible method in the fiber optic image guides (tapers) manufacturing. In *Video Communications and Fiber Optic Networks*, pages 206–219. International Society for Optics and Photonics, 1993.

[29] Yanxiang Lan, Yue Dong, Fabio Pellacini, and Xin Tong. Bi-scale appearance fabrication. *ACM Trans. Graph.*, 32(4):145:1–145:12, July 2013.

[30] Cheryl Lau, Wolfgang Heidrich, and Rafal Mantiuk. Cluster-based color space optimizations. In *Proceedings of International Conference on Computer Vision (ICCV 2011*, 2011.

[31] Jason Lawrence, Aner Ben-Artzi, Christopher DeCoro, Wojciech Matusik, Hanspeter Pfister, Ravi Ramamoorthi, and Szymon Rusinkiewicz. Inverse shade trees for non-parametric material representation and editing. *ACM Trans. Graph.*, 2006.

[32] Anat Levin, Daniel Glasner, Ying Xiong, Frédo Durand, William Freeman, Wojciech Matusik, and Todd Zickler. Fabricating brdfs at high spatial resolution using wave optics. *ACM Trans. Graph.*, 32(4):144:1–144:14, July 2013.

[33] Gordon E. Liljegren and Eugene L. Foster. Figure with back projected image using fiber optics. In *US Patent Number 4978216*. Walt Disney Company, 1990.

[34] Tom Malzbender, Ramin Samadani, Steven Scher, Adam Crume, Douglas Dunn, and James Davis. Printing reflectance functions. *ACM Trans. Graph.*, 31(3):20:1–20:11, June 2012.

[35] Wojciech Matusik, Boris Ajdin, Jinwei Gu, Jason Lawrence, Hendrik P. A. Lensch, Fabio Pellacini, and Szymon Rusinkiewicz. Printing spatially-varying reflectance. SIGGRAPH Asia, 2009.

[36] Wojciech Matusik, Hanspeter Pfister, Matt Br, and Leonard Mcmillan. A data-driven reflectance model. *ACM Transactions on Graphics*, 22:759–769, 2003.

[37] Henry P. Moreton and Carlo H. Séquin. Functional optimization for fair surface design. *ACM SIGGRAPH Comput. Graph.*, 26(2):167–176, July 1992.

[38] Jan Morovic. Color gamut mapping. Wiley, 2008.

[39] Addy Ngan, Fredo Durand, and Wojciech Matusik. Image-driven navigation of analytical brdf models. In *Eurographics Symposium on Rendering*, 2006.

[40] F. E. Nicodemus, J. C. Richmond, J. J. Hsia, I. W. Ginsberg, and T. Limperis. Radiometry. geometrical considerations and nomenclature for reflectance. 1992.

[41] Marios Papas, Thomas Houit, Derek Nowrouzezahrai, Markus Gross, and Wojciech Jarosz. The magic lens: Refractive steganography. *ACM Trans. Graph.*, 31(6):186:1–186:10, November 2012.

[42] Marios Papas, Wojciech Jarosz, Wenzel Jakob, Szymon Rusinkiewicz, Wojciech Matusik, and Tim Weyrich. Goal-based caustics. *Comput. Graph. Forum*, 2011.

[43] Marios Papas, Christian Regg, Wojciech Jarosz, Bernd Bickel, Philip Jackson, Wojciech Matusik, Steve Marschner, and Markus Gross. Fabricating translucent materials using continuous pigment mixtures. *ACM Trans. Graph.*, 32(4):146:1–146:12, July 2013.

[44] Fabio Pellacini, James A. Ferwerda, and Donald P. Greenberg. Toward a psychophysically-based light reflection model for image synthesis. SIGGRAPH, 2000.

[45] Fabio Pellacini and Jason Lawrence. Appwand: editing measured materials using appearance-driven optimization. *ACM Trans. Graph.*, 26, July 2007.

[46] Thiago Pereira and Szymon Rusinkiewicz. Gamut mapping spatially varying reflectance with an improved brdf similarity metric. *Comp. Graph. Forum*, 31(4):1557–1566, June 2012.

[47] Thiago Pereira, Szymon Rusinkiewicz, and Wojciech Matusik. Computational light routing: 3d printed optical fibers for sensing and display. *ACM Trans. Graph.*, 33(3):24:1–24:13, June 2014.

[48] Pratt and Salzberg. Process for orienting ferromagnetic flakes in paint films. *US Patent*, 1947.

[49] Raksha, Markantes, Coombs, Phillips, Kohmann, Argoitia, and Teitelbaum. Anisotropic magnetic flakes. *US Patent*, 2009.

[50] Coombs Chu Raksha, Markantes. Multilayer magnetic reflecting pigment flakes and foils. *European Patent*, 2011.

[51] Ramesh Raskar, Greg Welch, Kok-Lim Low, and Deepak Bandyopadhyay. Shader lamps: Animating real objects with image-based illumination. In *Proceedings of the 12th Eurographics Workshop on Rendering Techniques*, pages 89–102, London, UK, UK, 2001. Springer-Verlag.

[52] Ramesh Raskar, Remo Ziegler, and Thomas Willwacher. Cartoon dioramas in motion. In *Proceedings of the 2nd international symposium on Non-photorealistic animation and rendering*, NPAR '02. ACM, 2002.

[53] William A. Rense. Polarization studies of light diffusely reflected from ground and etched glass surfaces. *J. Opt. Soc. Am.*, 40(1):55–56, Jan 1950.

[54] Charles T. Markantes Paul G. Coombs Matthew R. Witzman Roger W. Phillips, Charlotte R. LeGallee. Methods for producing imaged coated articles by using magnetic pigments. *United States Patent US 6,759,097 B2*, 2004.

[55] Szymon M. Rusinkiewicz. A new change of variables for efficient brdf representation. In *In Eurographics Workshop on Rendering*, pages 11–22, 1998.

[56] Yuliy Schwartzburg, Romain Testuz, Andrea Tagliasacchi, and Mark Pauly. High-contrast computational caustic design. *ACM Trans. Graph.*, 33(4):74:1–74:11, July 2014.

[57] John M. Senior. *Optical fiber communications (2nd ed.): principles and practice*. Prentice Hall International (UK) Ltd., 1992.

[58] O. Sorkine, D. Cohen-Or, Y. Lipman, M. Alexa, C. Rössl, and H.-P. Seidel. Laplacian surface editing. In *Proceedings of the 2004 Eurographics/ACM SIGGRAPH symposium on Geometry processing*, SGP '04, pages 175–184, New York, NY, USA, 2004. ACM.

[59] Eric J. Stollnitz, Victor Ostromoukhov, and David H. Salesin. Reproducing color images using custom inks. SIGGRAPH, New York, NY, USA, 1998.

[60] Thevenet. Method of applying makeup to surface by means of a magnetic composition including reflective particles having metallic luster. *US Patent*, 2007.

[61] Greg Turk and James F. O'Brien. Shape transformation using variational implicit functions. ACM SIGGRAPH '99, 1999.

[62] Peter Vangorp, Jurgen Laurijssen, and Philip Dutré. The influence of shape on the perception of material reflectance. SIGGRAPH '07, New York, NY, USA, 2007. ACM.

[63] Zhou Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *Image Processing, IEEE Transactions on*, 2004.

[64] Gregory J. Ward. Measuring and modeling anisotropic reflection. In *Proceedings of the 19th annual conference on Computer graphics and interactive techniques*, SIGGRAPH '92, pages 265–272, New York, NY, USA, 1992. ACM.

[65] Tim Weyrich, Jia Deng, Connelly Barnes, Szymon Rusinkiewicz, and Adam Finkelstein. Digital bas-relief from 3d scenes. In *ACM SIGGRAPH 2007 Papers*, SIGGRAPH '07, New York, NY, USA, 2007. ACM.

[66] Tim Weyrich, Pieter Peers, Wojciech Matusik, and Szymon Rusinkiewicz. Fabricating microgeometry for custom surface reflectance. In *ACM SIGGRAPH 2009 papers*, SIGGRAPH '09, pages 32:1–32:6, New York, NY, USA, 2009. ACM.

[67] Tim Weyrich, Pieter Peers, Wojciech Matusik, and Szymon Rusinkiewicz. Fabricating microgeometry for custom surface reflectance. *ACM Transactions on Graphics*, 28(3):1–6, 2009.

[68] Karl Willis, Eric Brockmeyer, Scott Hudson, and Ivan Poupyrev. Printed optics: 3d printing of embedded optical elements for interactive devices. ACM UIST '12, New York, NY, USA, 2012. ACM.

[69] Genzhi Ye, Sundeep Jolly, V. Micheal Bove, Qionghai Dai, Ramesh Raskar, and Gordon Wetzstein. Toward BxDF Display using Multilayer Diffraction. *ACM Trans. Graph. (SIGGRAPH Asia)*, 33(6), 2014.

[70] P. Zolliker and K. Simon. Retaining local image information in gamut mapping algorithms. *Image Processing, IEEE Transactions on*, 16(3):664 –672, march 2007.