

A PRACTICAL FRAMEWORK FOR MEASURING  
AND MODELING THE APPEARANCE OF  
STRONGLY ANISOTROPIC MATERIALS

A Thesis

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Master of Science

by

Nicolas Savva

January 2017

© 2017 Nicolas Savva  
ALL RIGHTS RESERVED

## ABSTRACT

We present a practical sparse measurement technique and a novel parameter fitting approach for the appearance of strongly anisotropic materials, with application to finished wood. Our approach makes use of bilateral symmetry arguments to reduce the amount of input data required to capture a spatially varying BRDF. This significantly decreases the necessary acquisition and computation time to recover the model parameters, with an observed speedup close to an order of magnitude over previous work, while achieving significantly improved results. We validate the quality of the rendered results from the new approach using additional dense ground truth measurements obtained using a 4-DoF spherical gantry. We also demonstrate a field measurement system using a portable hoop with individually addressable LEDs. The device is inexpensive, simple to build, fast in operation, and fully compatible with the proposed acquisition technique. We provide a database of wood BRDFs and an implementation of the model and optimization fitting pipeline under Mitsuba [24] to demonstrate the results.

## BIOGRAPHICAL SKETCH

Nicolas Savva was born in Plovdiv, Bulgaria before moving to spend the next two decades on the sunny Mediterranean island of Cyprus. From 2007 to 2011 he studied at Cornell University, on a Fulbright scholarship, where he received a Bachelor degree in Physics and Computer Science. Following that, Nicolas remained in Ithaca where he joined the Program of Computer Graphics, in the summer of 2011, as a member of the Sustain team. This marked the beginning of his graduate studies and the period of further exploration of several topics of interest under the domain of physically-based rendering and simulation. Throughout the next several years Nicolas was heavily involved in pedagogy while furthering his professional development as a frequent course assistant and occasional instructor for numerous courses, ranging from introductory Computer Science classes to some more advanced courses in Computer Graphics and Scientific Computing. Nicolas is graduating with a research MS in Computer Science with concentrations in Computer Graphics (CG), Artificial Intelligence (AI), and Computational Science and Engineering (CSE). Nicolas is now part of the Autodesk Cloud Renderer team, in San Francisco, working on the next generation of physically-based rendering technology.

This work is dedicated to my family  
In loving memory of my grandfather Nikolai

## ACKNOWLEDGEMENTS

This work would not have been possible without the patience and support of my graduate committee: Many thanks to Steve Marschner, for all the wisdom (hopefully) imparted through numerous discussions, and Bruce Walter for helping me see this project through. Bruce has been instrumental in helping with the development of the framework, offering his rendering expertise and assisting with the appearance measurements. Thank you Kavita for being a sounding board and my academic godmother all these years. David Bindel has also spent many hours with me sharing both his pedagogical and scientific computing views and expertise. I wish to acknowledge Don Greenberg for taking me into his lab, his passion for Computer Graphics, and the generous initial support. Last but not least I want to thank Doug James, for acting as my undergraduate advisor, and encouraging me to pursue this academic path many years ago.

There are numerous other people I am indebted to over the years for being my mentors, patrons, colleagues and friends... Sparing some trees (and bytes), I decided to maintain a longer online text that I may occasionally update through my personal website.

Please check:

*[www.nicolassavva.com/acknowledgeCornell/](http://www.nicolassavva.com/acknowledgeCornell/)*

You have my gratitude - thank you all!

## TABLE OF CONTENTS

Biographical Sketch . . . . .	iii
Dedication . . . . .	iv
Acknowledgements . . . . .	v
Table of Contents . . . . .	vi
List of Tables . . . . .	vii
List of Figures . . . . .	viii
<b>1 Introduction</b>	<b>1</b>
<b>2 Background</b>	<b>4</b>
2.1 Mathematical Preliminaries . . . . .	5
2.2 Prior work . . . . .	8
<b>3 Method and Fitting Approach</b>	<b>20</b>
3.1 Acquisition Pipeline Overview . . . . .	20
3.2 Finished Wood Appearance Model . . . . .	22
3.3 Data Acquisition . . . . .	27
3.4 Parameter Fitting Procedure . . . . .	34
<b>4 Results and Analysis</b>	<b>43</b>
4.1 Overview . . . . .	43
4.2 Ground Truth Validation . . . . .	43
4.3 Comparison to Previous Work . . . . .	51
4.4 IRLS Convergence, Fit Quality Analysis and Parameter Sensitivity	53
<b>5 Conclusion</b>	<b>65</b>
5.1 Limitations and Future Work . . . . .	65
5.2 Closing Remarks . . . . .	67
<b>A Fitting Approach Pseudocode</b>	<b>68</b>
<b>Bibliography</b>	<b>71</b>

## LIST OF TABLES

3.1	Fit Comparison: We show results for two different sample locations. The top row shows the output of our approach compared to previous work (middle row) and the measured validation data (bottom row) . . . . .	42
A.1	Parameters of the wood BRDF model . . . . .	70



## LIST OF FIGURES

1.1	Anisotropic Materials in our lives: Wood and fabrics are commonplace materials in our environment that exhibit intriguing behavior when interacting with light (source: [ <a href="https://www.pinterest.com/andrewgball/house-ideas/">https://www.pinterest.com/andrewgball/house-ideas/</a> ]) . . .	1
2.1	Wood table: A periodic table arrangement depicting the many different kinds of woods from around the world. It becomes immediately obvious that wood is a rather diverse and broad category (source: [ <a href="https://www.pinterest.com/andrewgball/house-ideas/">https://www.pinterest.com/andrewgball/house-ideas/</a> ]) . . . . .	4
2.2	Illustration of several different ways in which light scatters from a material surface (source: [ <a href="http://www.mitsuba-renderer.org/">http://www.mitsuba-renderer.org/</a> ]) . . . . .	6
2.3	BRDF and SV-BRDF taxonomy (source: [16]) . . . . .	8
2.4	Photographs of a selection of the material spheres found in the MERL dataset [33] . . . . .	12
2.5	Prior Work Devices: Prior techniques of comparable quality require long measurements times or equipment confined to the laboratory setting . . . . .	19
3.1	Pipeline Overview: An end-to-end schematic view of our acquisition and parameter fitting framework . . . . .	21
3.2	Wood Parameters on a 1D orbit: Sparse sampling of the hemisphere plotted as a 1D function to illustrate relevant parameter discriminative features . . . . .	22
3.3	Orbit Rings: An illustration of the range of possible projected BRDF for different locations of the surface of our strongly anisotropic material samples. Note the 1D dotted gray orbit intercepting the strongly anisotropic signal colored in orange . . .	23
3.4	Wood Structure: Wood contains rough cylindrical air tubes (source: NC Brown Center for Ultrastructure Studies, SUNY, Syracuse, NY) . . . . .	24
3.5	Many ways of cutting wood: Different cuts expose different fiber orientations (source: [Beals and Davis 1977]). . . . .	25
3.6	Grain Patterns: Many different wood distributions depending on the wood species (source: [Beals and Davis 1977]). . . . .	25
3.7	Illustration of Subsurface highlights from wood surface (source: Marschner 2005) . . . . .	26
3.8	Scattering from cylindrical geometry. Reflection from interfaces parallel to an axis end up preserving the inclination (source: Marschner 2005) . . . . .	27
3.9	Wood BRDF parameters illustration . . . . .	28
3.10	Three dimensional perspective on the left and topdown view on the right used in subsequent illustrations (such as in Fig 3.11) . .	29

3.11	The Subsurface highlight distributions: Visualization of the strongly Anisotropic BRDF . . . . .	29
3.12	Sparse Samples: Hemisphere of Illumination Direction with a highlighted single Orbit of Measurements . . . . .	30
3.13	Illustration of the Capture Setup: The camera is placed directly above and images are captured with a sequence of light positions around a ring . . . . .	31
3.14	Caliber software package [28]: rigid constraint optimization for localization of the camera and light source . . . . .	32
3.15	Irradiance Correction: the captured data is scaled using the incident illumination to recover measured BRDF values . . . . .	33
3.16	LED Loop: A photo of our portable addressable LED loop device (left) and the evolution of the prototype (right) . . . . .	33
3.17	Illustration of the subsurface highlight cone with different elevation for the wood fiber (Top images yellow fiber orientation nearly parallel to the surface; bottom images yellow fiber orientation significantly slanted out of the plane) . . . . .	34
3.18	Symmetry Projection . . . . .	37
3.19	Labeled parameters on the projected gaussian and respective logarithmic space parabola . . . . .	37
3.20	IRLS objective function and parameter minimization (source: [18])	39
3.21	Gaussian Transform: Parabola Parameters in Log Space . . . . .	40
3.22	Parameter Map result of the wood model for the oak sample using our proposed fitting approach . . . . .	41
4.1	Checkerboard Comparison to Measurement Data: Interleaved blocks of acquired real data and synthesized fitted model output . . . . .	45
4.2	Checkerboard Comparison of Fitted Result to Measurement Data (top image oak; bottom image mahogany). Note the grid boundaries are barely visible indicating a close match . . . . .	46
4.3	Progressively zooming into our captured sample surface to view a small region of the Spatially-Varying wood BRDF) . . . . .	47
4.4	Closeup of Fig 4.3 for the area inside the red rectangle. Top image showing the dense validation BRDF data while the bottom image shows the ring subset data used by our fitting algorithm . . . . .	48
4.5	Our symmetry model fit of a diffuse location on the surface . . . . .	49
4.6	Our symmetry model fit of a fiber with a large elevation angle (truncated gaussian) . . . . .	49
4.7	Our symmetry model fit of a fiber parallel to the surface . . . . .	50
4.8	The SRM05 [31] fitting approach . . . . .	51

4.9	Comparison of BRDF fit for a region against ground truth and a prior wood appearance model (Marschner et al SG05). The center column consists of ground truth measurements for a region of the walnut sample at different magnification levels (rows). The left and right columns correspond to the results generated by the new symmetry aware model and the dense measurements plane fitting approach respectively. Note that the rendered results do not include the surface specular component. Blue regions denote areas of missing data in the ground truth. . . . .	56
4.10	Compare Paper Target Diffuse Fit (SRM05 left; our approach right)	57
4.11	Performance Comparison to [Marschner 2005] . . . . .	57
4.12	Comparison of Oak Sample Pixels . . . . .	58
4.13	Closeup Comparison of Oak Sample Pixels . . . . .	58
4.14	Comparison of the fiber parameter maps output from our approach and the prior work [31] . . . . .	59
4.15	Comparison of 1D signal Fit: Reasonable fit for both fitting approaches . . . . .	60
4.16	Comparison of 1D signal Fit: Our symmetry based approach matches while the prior approach was unable to converge . . . .	60
4.17	$L_2$ distance plot for several wood samples. A cumulative histogram showing the population of pixels within a certain bucket of $L_2$ error from prior work and measurements. . . . .	61
4.18	Extracted Parameter Maps for Several Wood Samples. Columns (left to right): Diffuse Color, Fiber Color, Highlight Width, Fiber Direction (Ours), Fiber Direction (SRM05) . . . . .	62
4.19	The approach is robust and can converge even when there is significant noise in the signal while not being sensitive to the initialization (top figure shows convergence to ground truth parameter). The IRLS algorithm exhibits fast convergence to a fixed point within a few iterations steps as can be seen from the above plots (bottom figure shows the diminishing parameter gradients). . . . .	63
4.20	Pseudocolored clusters of closely matching BRDF pixels (top image) and one white pixel representatives of each unique bin of pixels (bottom image). Less than 10% of signals are unique in sufficiently large samples . . . . .	64

# CHAPTER 1

## INTRODUCTION



Figure 1.1: Anisotropic Materials in our lives: Wood and fabrics are commonplace materials in our environment that exhibit intriguing behavior when interacting with light (source: [<https://www.pinterest.com/andrewgball/house-ideas/>])

### Motivation

Finished wood is an important material in interiors and is used in a wide variety of manufactured products, from furniture to musical instruments. It is a uniquely attractive material due to its complex, anisotropic subsurface structure (see Fig 3.4), and realistically modeling its appearance requires multi-channel parameter maps that encode the 3D orientation of fibers [31].

While previous methods for measuring these parameter maps produce good results, they require making large numbers of measurements over the whole

hemisphere, then fitting a BRDF model to them [31]. This requires long measurements with expensive 2-DoF motion equipment, which produce large amounts of data that are slow to process. This has very much hampered the adoption of anisotropic parameter maps, so that the practical state of the art is still to use color textures with a microfacet layer for the finish [50].

This work introduces a new measurement technique that uses a 1D family of light source positions along a circular path, requiring only around 100 images for high quality results. This means the capture can be done with simple motions or with a 1D array of sources, making it much faster and less expensive. To process this data, we propose a new fitting technique for highly anisotropic materials, which makes anisotropy into an asset rather than a liability and leverages the symmetry inherent in the material to robustly and efficiently fit a multi-parameter wood BRDF to these sparse measurements. The technique has two phases: for each pixel it first extracts an axis of bilateral symmetry, then performs a fast 1D Gaussian fit to the data. Both are simple and robust, leading to high quality results.

## **Chapter Overview**

In the following sections we begin with a summary of relevant mathematical concepts and a survey of the state of the art (Sec 2), followed by a presentation of the fitting method in detail (Sec 3). We then show results from two different measurement setups: a laboratory setup using a 20 megapixel SLR and a moving light source, and a field-deployable setup using a high-definition video camera and a ring of LEDs. Results are validated against full-hemisphere measure-

ments and compared to previous work that uses the full hemisphere data (Sec 4), demonstrating that the new method produces detailed, artifact-free maps that are superior to the results of previous work, at a fraction of the time, storage, and apparatus complexity. The conclusion and future work chapter (Sec 5) discusses the limitations of the current approach and outlines a proof of concept system to act as the next step for its evolution - showing how we can increase measurement automation and overall practicality. Relevant software, data and additional results will be made available and can be found through the project website [43].

## CHAPTER 2 BACKGROUND



Figure 2.1: Wood table: A periodic table arrangement depicting the many different kinds of woods from around the world. It becomes immediately obvious that wood is a rather diverse and broad category (source: <https://www.pinterest.com/andrewgball/house-ideas/>)

The appearance of materials, as perceived by humans, depends on the light that is absorbed, transmitted and reflected until it reaches the eyes of the viewer. The final appearance of an object can vary significantly, even in the case of a single category, as illustrated in the figure above (Fig 2.1) using a wide range of wood samples. Many parameters and properties come into play: the viewing and lighting directions, surface roughness, surface finish and other coatings and pigments, even the internal structures in proximity to the viewing point inside

the sample. Accurate measurement and modeling of such materials still remains a major challenge in Computer Graphics though significant progress has been made by researchers over the last couple of decades. There is still no universal model that can closely represent most classes and even the non-standardized specialized solutions that do exist (such as the illustrated finished wood examples) require a disabling amount of labor, money and storage to achieve high quality results. This work focuses on alleviating some of these shortcomings for strongly anisotropic materials; we focus on finished wood, but the same ideas could also be adapted to some woven fabric and brushed metal surfaces. The following two subsections introduce some relevant mathematical concepts and survey the related state of the art in material acquisition and representation.

## 2.1 Mathematical Preliminaries

The Bidirectional Reflectance Distribution Function (BRDF) defines how light is reflected at an interface of a surface. This function is effectively a four parameter equation that describes how incident energy is redirected in a hemisphere above a surface point. First introduced by Nicodemus [37] as a photorealistic representation of opaque surfaces, it is a simplified reflectance representation compared to the BSSRDF [25] where light can be incident at one point of the surface propagate through the surface and emerge at a different surface location.

Formally the BRDF is defined as the ratio of the outgoing radiance to the incoming irradiance (please see the symbols table for definitions):

$$f_r(v_i, v_r) = \frac{L_o(v_r)}{E_i(v_i)} = \frac{L_r(v_r)}{L_i(v_i)\cos(\theta_i)d\omega_i} \quad \left[ \frac{1}{sr} \right]$$



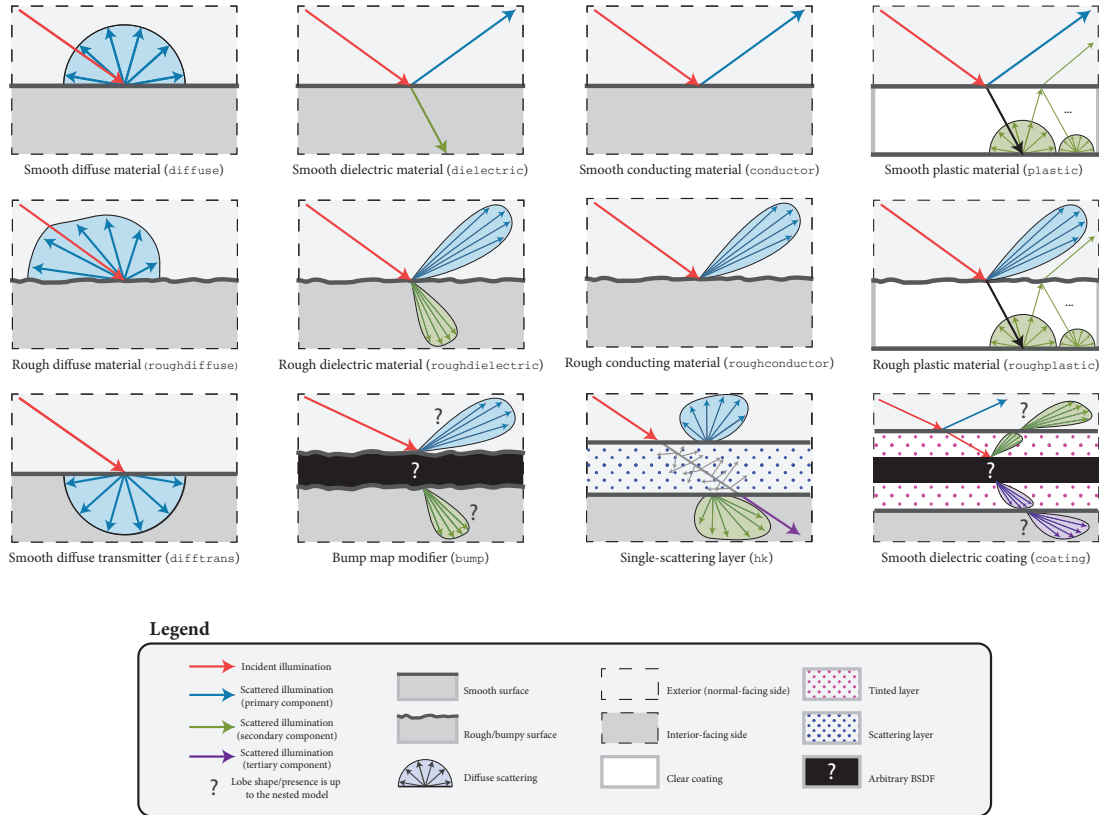


Figure 2.2: Illustration of several different ways in which light scatters from a material surface (source: [http://www.mitsuba-renderer.org/])

The particular function aims to represent the reflectance characteristics of homogeneous materials. The logical extension of the BRDF to non homogeneous spatially varying surfaces is given by the SV-BRDF [19] which defines a BRDF for each position on the material surface by adding the surface coordinates as function parameters  $f_r(x, y, v_i, v_r)$ . This particular generalization works well for flat and opaque surfaces though it makes it much harder to capture materials which may now require a large number of samples to be taken both in the angular ( $\omega_i, \omega_r$ ) and spatial domains ( $x, y$ ). This usually results in long measurements and processing times as well as expensive specialized equipment to capture the appearance data. When the variation in both the reflectance and the

small-scale geometry is significant, effects such as masking, occlusion and inter-reflection can occur at the mesoscopic scale, when the surface interacts with the neighboring local variations. A different reflectance function, the Bidirectional Texture Function (BTF) [7] can describe the rough surfaces and is preferable to the SVBRDF for surfaces with coarse scale variations (see [9] for a recent BTF state of the art survey). The Bidirectional Transmittance Distribution Function (BTDF) [47] is the equivalent of a BRDF for light passing through a surface. Combining the BRDF and BTDF together we end up with a scattering function characterizing both transmission and reflection. The above taxonomy of scattering functions is important to keep in mind when considering the scope of a proposed appearance model and the tradeoff between practicality and complexity. For the purposes of this project we will focus on flat surfaces and use a SV-BRDF. There are several BRDF properties that we expect will hold true. The light paths should be reversible so that for any pair of outgoing and incident directions the following property holds:  $f_r(v_i, v_r) = f_r(v_r, v_i)$  (Helmholtz reciprocity). We also expect that the BRDF is non-negative for any pair of directions  $f_r(v_r, v_i) \geq 0$ . Conservation of energy is also required for a physically-based model and hence the sum of the reflected energy must be less than or equal to the incident amount. Additionally, we make an effort to model Fresnel effects [44] as we strive for a physically-based model accounting for the surface refraction. We will be dealing with an anisotropic BRDF and expect that the reflection will change with respect to the surface rotation around the normal to the surface (as opposed to isotropic BRDFs where the reflectance distribution is invariant to planar rotations of the surface).

## 2.2 Prior work

There are several classes of BRDF models – we will briefly cover some of the more relevant to our work in the following subsections: **(1) Physically-based models** tend to model a rough surface using real world optics and usually result in mathematical models with accurate formulas and tunable parameters, **(2) Phenomenological models** fit an analytical formula to reflectance data and reproduce approximate reflectance characteristics without attempting to model the actual behavior of the real world material, **(3) Data-driven BRDF models** do not attempt to model the data and instead use a grid or table to lookup and interpolate the appearance of the material.

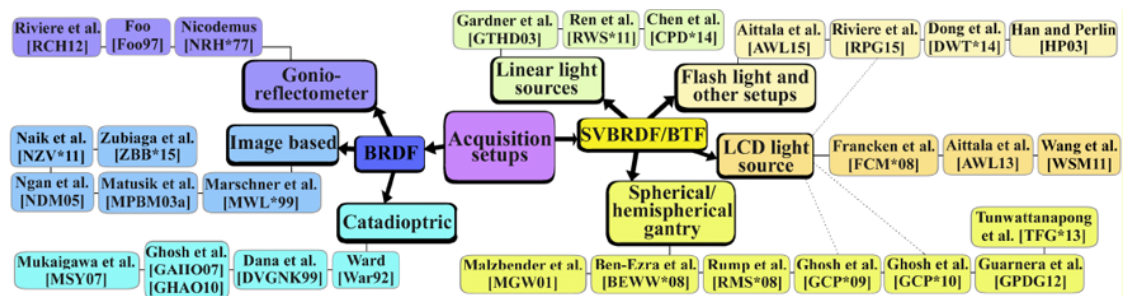


Figure 2.3: BRDF and SV-BRDF taxonomy (source: [16])

### Acquisition Devices and Setups

The major bottleneck in the appearance modeling process is usually the acquisition or characterization through capture or observation step. Researchers over the years have developed numerous approaches which we attempt to summarize in the upcoming subsections. Typically a light source is used to uniformly illuminate an area of the sample surface and a detector measures a small area

within this illuminated region. The different systems offer a range of accuracies, costs, and speeds for measuring the reflectance functions. We get a reasonable overview of this body of work through the taxonomy shown in Fig 2.3. Significant work has also been done on quantifying the error, for both the physical and the perceptual distance, to assess the quality of acquisition setups and the relevant appearance models.

### **Image-based Measurements**

Image-based acquisition techniques typically only require commonplace consumer equipment and make use of photographs, effectively reducing the cost of the setup. A series of pictures can be taken and such images can quickly capture the reflected light with a range of surface to camera orientations. A drawback of such an approach tends to be that more time and additional consideration is usually required when capturing the wavelength spectrum BRDFs (typically done using removable filters in front of the camera) [30]. Marschner et al. [32] proposed an accurate and rapid isotropic BRDF capture approach that works for a broad range of homogeneous materials and results in high accuracy and resolution supporting a large range of reflection and illumination directions. A CCD sensor handheld camera with a set of color filters and an electronic flash light source are sufficient to measure surfaces with simple geometric shapes (for which analytic formulations exist) and can also be adapted for irregular geometries given a 3d model of the surface. The characterized camera moves from near the light source to the opposite side from the light source, thus measuring configurations from close to a retro-reflection all the way to grazing angle reflections. Additional photos are taken to give the intensity and location of

the light source as well as the camera and sample poses. The authors determined that about thirty images from different position are enough to cover the 3D BRDF domain. Every pixel in the images is then used as a single sample in the BRDF domain using bundle adjustment. Similar to the above work, Matusik et al. [33] uses such an approach to measure 100 isotropic materials and build a large reliable BRDF database (see Fig 2.4). Ngan et al. [36] later collaborated with Matusik to put forward an anisotropic BRDF acquisition setup that works for flexible and flat samples. Strips of the material with different orientations are taken from the flat sample and wrapped around a cylinder to deal with anisotropy. The cylinder is then tilted using a motor allowing for the acquisition of the missing degrees of freedom. The light source rotates around the cylinder with a static camera which is used to capture the cylindrical target. A set of eight pictures with different exposures is taken at each position to reconstruct a high dynamic range image. The primary limitation resolution-wise is the number of material strips that can be attached to the cylinder where the number of light positions and tiling strips can be changed to achieve a certain BRDF measurement resolution. Naik et al. [35] uses space-time images captured with a time of flight camera with two different setups. Two known Lambertian materials, a source and a receiver are used in a three-bounce scattering setup. In the first of the configurations, the laser illuminates the surface and the camera views the receiver indirectly measuring the patch. In the other configuration, the patch is not directly visible from the camera and the source and receiver are the same surface. This is based on around-the-corner viewing where the light is multiplexed along different paths, some of which are of the same length, arriving from multiple paths to the same point at the same time. The measurements can be decoded by solving a sparse underdetermined system

by using the halfway vector parametrization, recovering the parameters of the Ashikmin-Premoze model [3]. Analyzing the streak images is enough to find the specular peak. The setup can be used to take many BRDF measurements simultaneously but it does require costly high-speed cameras and usually has low signal to noise ratio brought on by the size of patches and multiple scattering bounces. More recent work by Zubiaga et al. [51] works locally in Fourier space and analyses how up to 2nd order BRDF moments induce blurring, warping and coloring of the reflected radiance to further understand how the properties of BRDF influence the observed appearance. Around forty unimodal materials from the MERL [33] database are used with a heuristic method for diffuse and specular separation of two dimensional slices from those BRDFs.

### **Gonioreflectometers**

Gonioreflectometers measure the spectral reflectance of surfaces and can handle both specular and diffuse scenarios. Nicodemus outlines the construction of such a device, which was further developed and constructed by several others: Torrance and Sparrow [45], He et al. [21], and Blinn [5]. Hsia and Richmond [22] had the following setup: A sample is placed in a holder mounted on a turntable that can rotate around the vertical axis, a laser beam light source is used to illuminate the surface and a detector can capture reflected light from the sample. The sample holder is mounted on an arm attached to the turntable and placed in front of the detector. Barium-sulfate coated averaging spheres are used to measure any incident light. Foo [10] made use of a three axis gonioreflectometer design which had two degrees of freedom. It consisted of a stationary detector, a light source that could rotate around the sample and a folding mirror.




Figure 2.4: Photographs of a selection of the material spheres found in the MERL dataset [33]

This system was capable of measuring reflection at large grazing angles and was also capable of high dynamic range captures with rather precise measurements. The drawback of the system was the speed of acquisition and the fact that it could only measure isotropic BRDFs. Li et al. [27] also proposed a similar system to that of Foo. Rivera et al. [41] makes use of an in-plane polarized multispectral radiometer. The polarized detection system takes a Fresnel equation approach to identify the polarization axis and has a setup consisting of three lin-

early polarized laser light sources. It can sample at zero lighting angles and is fully calibrated for multispectral and polarized in-plane BRDF measurements. The reconstructed BRDF data proves to be relatively robust to noise and makes use of a proposed inversion algorithm for the high angular measurements of BRDFs.

### **Spherical Gantries and Catadioptric Setups**

Catadioptric systems use refracted and reflected light to reduce optical aberrations, usually resulting in generally efficient image-based BRDF acquisition setups with devices that lack moving parts. Ward [49] measures anisotropic surfaces under various orientations with repeat measurements while capturing the entire hemisphere of the reflected and refracted illumination at the same time. A half-silvered hemisphere and a fisheye lens is used to achieve this; however grazing angles and sharp specular peaks cannot be acquired with the given device. Dana et al. [7] uses a robot arm which can hold and rotate the material sample, along with a Fresnel lens, a video camera and a halogen bulb. The light is fixed in position and the camera is moved around to acquire measurements from seven different locations. For anisotropic materials the sample is rotated around the  $z$ -axis and the above procedure is repeated. This setup is unable to capture fine-scale texture variations and noise within measurements tends to be an issue. Ghosh et al. [13, 15] proposes a device without moving parts. A camera focusing on a zone of reflected directions, a light source with a beam splitter, a mirror dome and mirror parabola are used for this setup. The focus of the illumination beam is on the mirrored components that the beam reflects back to the origin. It can be used for a specially designed orthonormal zonal basis function



illumination which allows for fast BRDF acquisition and good signal to noise ratio. The measurements can then be fitted to an analytical reflection model or projected into a spherical harmonics basis. Mukaigawa et al. [34] designed a device for anisotropic BRDF capture which makes use of a projector as the light source. The projector is placed at the focal point of an ellipsoidal mirror. A beam splitter is required since the camera and projector cannot be at the same spatial point. The number of images needed depends on the wanted sampling for lighting and viewing directions and can be scaled based on the desired accuracy. The authors go on to fit the acquired data to the Ward anisotropic model. Malzbender et al. [29] makes use of a hemispherical device which has 50 strobe light sources with the camera placed at the apex of the sphere to acquire pictures of nearly flat samples on the floor illuminated with a single source at a time. Polynomial Texture Maps are used to represent the measured data encoded as surface luminance at the local texture coordinates for the sample point and normalized projected light vector. Ben-Ezra [4] also uses a hemispherical device to capture anisotropic BRDFs provided an accurate geometric and radiometric calibration. LEDs are used as both the light sources and detectors without use of a conventional camera or any moving parts in the setup. This results in a fast acquisition system. The described device uses 84 LEDs pointed towards the center of a hemisphere. Each LED is switched one at a time while the rest act as detectors. Multiplexing illumination can be used to improve the signal to noise ratio and multispectral data can easily be acquired using colored LEDs. Due to the small number of LEDs used the BRDF resolution is limited and retroreflection cannot be captured since an individual LED cannot be an emitter and detector at the same time. Ramp [42] built a hemispherical gantry with 151 fixed cameras uniformly distributed. The cameras are equipped with flashes which act as

the light sources. All the cameras capture an image for every flash and multiple images are taken for HDR reconstruction. Ghosh et al. [14] put forward three setups to estimate spatially varying BRDFs using polarised second order spherical gradient illumination patterns. The device can capture both isotropic and anisotropic materials used to recover specular reflections, specular roughness, specular albedo and the reflection vector. The first setup makes use of 150 linearly polarized LED lights with the object at the center of the sphere. The second setup is suitable for flat objects and uses an LCD screen as a light source which is placed close to the subject. The third setup uses a roughly specular sphere, that can reflect onto the subject the light emitted by a projector, with the subject again at the center of the hemisphere. The camera is placed to observe the subject from the apex of the hemisphere and in this fashion dense sampling can be achieved. Ghosh also goes on to use circularly polarized spherical illumination to split specular and diffuse albedo and estimate the index of refraction and specular roughness for isotropic SVBRDFs given a known surface orientation. The device configuration is similar to the one outlined above with the difference of using circular polarizers at the light sources. Four pictures can then be used with three differently oriented linear polarizers and a circular polarizer placed in front of the camera to acquire the required data to be used in the recovery of the Stokes field. Guarnera et al. [17] uses the same approach and extends it to cover unpolarized illumination to also obtain the perpixel surface normal estimate from the same input. More recently Tunwattanapong [46] built a device with a spinning semi-spherical illumination arc consisting of 105 LEDs pointing towards the center and spinning around the sphere to sweep out a continuous spherical harmonic illumination pattern. The performed experiments show that approximately forty images are sufficient to estimate the captured anisotropic

SVBRDFs as well as recover geometry of highly diffuse or highly specular objects. Gardner et al. [12] also built a low cost linear light source device to capture flat samples making use of a fixed camera and a structured light diode. A 50cm long neon tube is translated horizontally over the surface of the sample along with the camera. Given the camera and light position for each frame the data is used to fit the Ward isotropic model. A laser projects a stripe on the surface which is deformed by the surface and used to recover the geometry. The device is used to recover diffuse and specular colors, surface normals and specular roughness. Chen [6] refined this approach for capture of anisotropic surface reflectance by using a microfacet model to model the anisotropy. Ren [39] later used a handheld linear light source alongside a BRDF chart to obtain isotropic SV-BRDFs from video using a mobile device. The BRDF chart was made of 24 flat tile square samples of known BRDFs. The tiles were made out of specular material except for a single diffuse one which was used for camera calibration. A 40cm fluorescent tube was slowly translated by hand over the sample after being closely placed to the BRDF chart.

### **Flash Illumination and LCD Light Sources**

Francken et al. [11] demonstrates that a common SLR camera and an LCD display can be used to recover detailed normal maps of specular objects using a halfway vector formulation. A gray code lighting pattern is used to quickly estimate the illumination directions. Aittala [1] is similar to the above setup and relies on the design of the image formation model. It uses a Fourier basis for the measurements where Bayesian inference is used for the reconstruction of isotropic BRDFs. Wang et al. [48] makes use of a conventional LCD as an area

light source coupled with a vision camera to rapidly capture isotropic glossy and bumpy surfaces. Two images are used for calibration and the pose with respect to the camera is computed using a target placed on the surface. The surface is illuminated with a half-white/half-black image with a vertical edge and the overall roughness is estimated by fitting a gaussian filter that blurs the step-edge to produce the one observed. The information is then used to estimate the microscale roughness and works well for highly glossy surfaces such as paints, plastics and metals. Recently Riviere et al. [40] used a mobile device LCD as an extended light source mounted at normal incidence half a meter from an isotropic planar material sample in a dark room. The inherent linear polarization of the LCD panel is used to separate the specular and diffuse contribution with the use of two pictures of the sample, making use of a differently oriented plastic sheet linear polarizer placed in front of the camera. The surface normal, specular roughness and albedo are also estimated by using different illumination patterns. Riviera also goes on to demonstrate that mobile devices with collocated camera and light flash can be used to capture the backscatter surface reflectance which can then be fitted to a microfacet BRDF model and used to approximate such a distribution. A video is acquired using the handheld mobile device in a dimly lit room where data is captured over the upper hemisphere above the sample. The diffuse gray squares of a ColorChecker are used for reflectance calibration and the top view frame of the sample at normal incidence is used as the reference frame to register the rest of the frames. The mobile device magnetometer and accelerometers give an estimate of the lighting and viewing directions. Aittala et al. [2] later went on to use a mobile device measurement setup for stationary materials to obtain a flash-no-flash image pair of stationary textured material with known size. A full anisotropic SV-BRDF is reconstructed

using the measurements where the flash image gives rough retroreflection data for each pixel and the no flash ambient image is used to find other locations on the surface with similar local reflectance properties. Tiles are used to approximate a repeating texture pattern and a high frequency detail transfer approach is used together with a crude SV-BRDF fit in a nonlinear optimization to fit an analytic SV-BRDF model which is then reverse propagated to the full image. A different measurement device presented by Han [20] makes use of a kaleidoscope image which gives information from many different sample viewpoint. A 45 degree beamsplitter is used to share the optical path between the camera and a projector used for illumination. A series of pictures is takes with a sequence of known incoming illumination directions. The lack of moving parts and registration issues make this a fast approach for measurement of BSSRDF and BTFs.

For those who wish to read further on the topics covered in this chapter, recent survey papers and short courses such as [16] and [50] do a much more comprehensive and excellent job of covering the state of the art in BRDF representation, acquisition and the relevant devices. The following books by Filip [19], Pouli [38], and Dorsey [8] are also excellent resources.




Figure 2.5: Prior Work Devices: Prior techniques of comparable quality require long measurements times or equipment confined to the laboratory setting

## CHAPTER 3

### METHOD AND FITTING APPROACH

#### 3.1 Acquisition Pipeline Overview

The main idea motivating our work is the observation that the extreme anisotropy of wood makes it easier to measure than most surfaces, because its linear highlight always shows up near the edges of the BRDF hemispherical domain (see Fig 3.3). We will always find the features we need to identify the fiber direction and the roughness even when only making observations near the edge (as is the case with the subset of samples denoted by the dotted lines in the previous figure). Our approach leverages this to measure the highlight cone features with a sparse orbit of illumination directions, which allows us to deduce the orientation of the physical fiber structures. The bilateral symmetry of the observed sparse signal (see Fig 3.2) helps us identify the fiber direction from the data so that we can efficiently fit the appearance model parameters.

The overall aim is to use the simplest model that can fully represent the appearance of each supported anisotropic material. Finished wood [31] is used as an illustrative example here, while as previously mentioned the technique can be adapted for other strongly anisotropic materials (as shown by additional examples with an appropriate model for specular reflections from some types of woven fabric [23]). We design our end-to-end approach in a way that minimizes the amount of input data in an attempt to streamline the capture process so as to reduce the total time that goes both into acquisition as well as model fitting process compared to previous work (see Fig 3.1 for an end-to-end pipeline overview). A symmetry argument around the axis of the fiber due to the under-

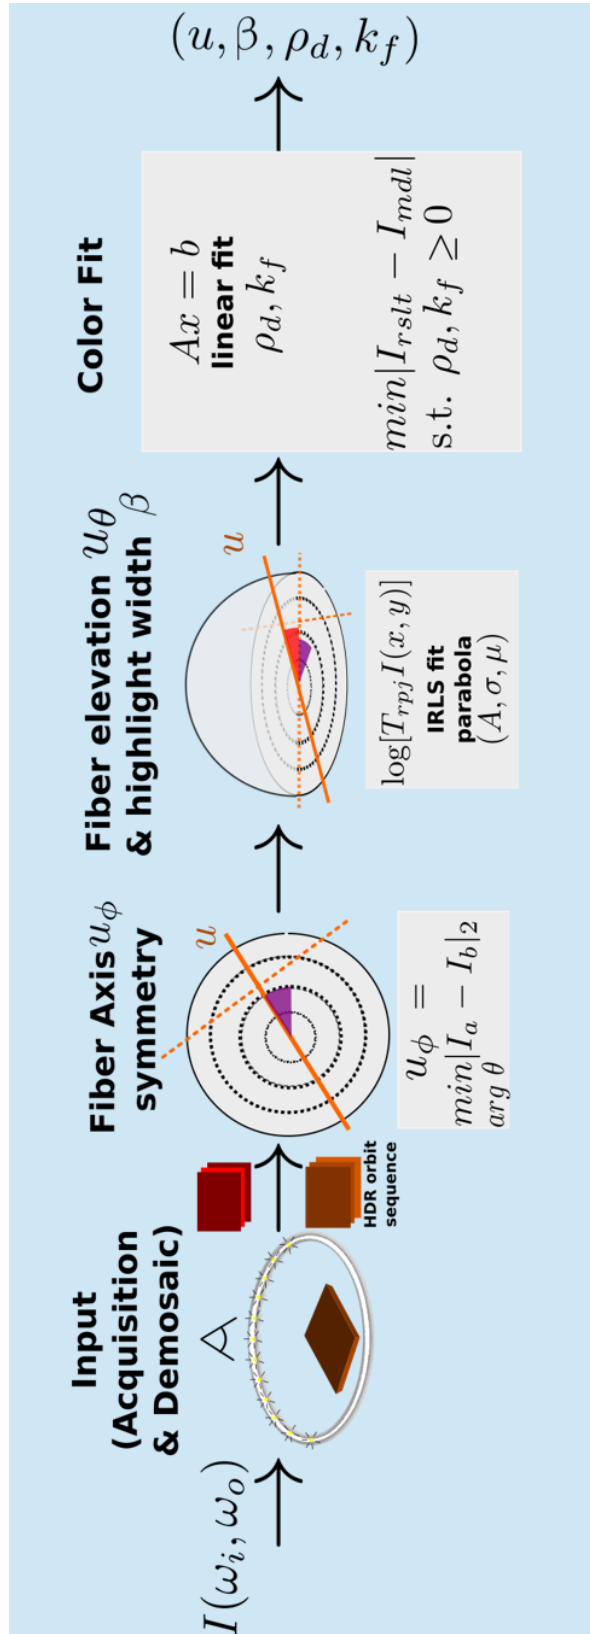


Figure 3.1: Pipeline Overview: An end-to-end schematic view of our acquisition and parameter fitting framework



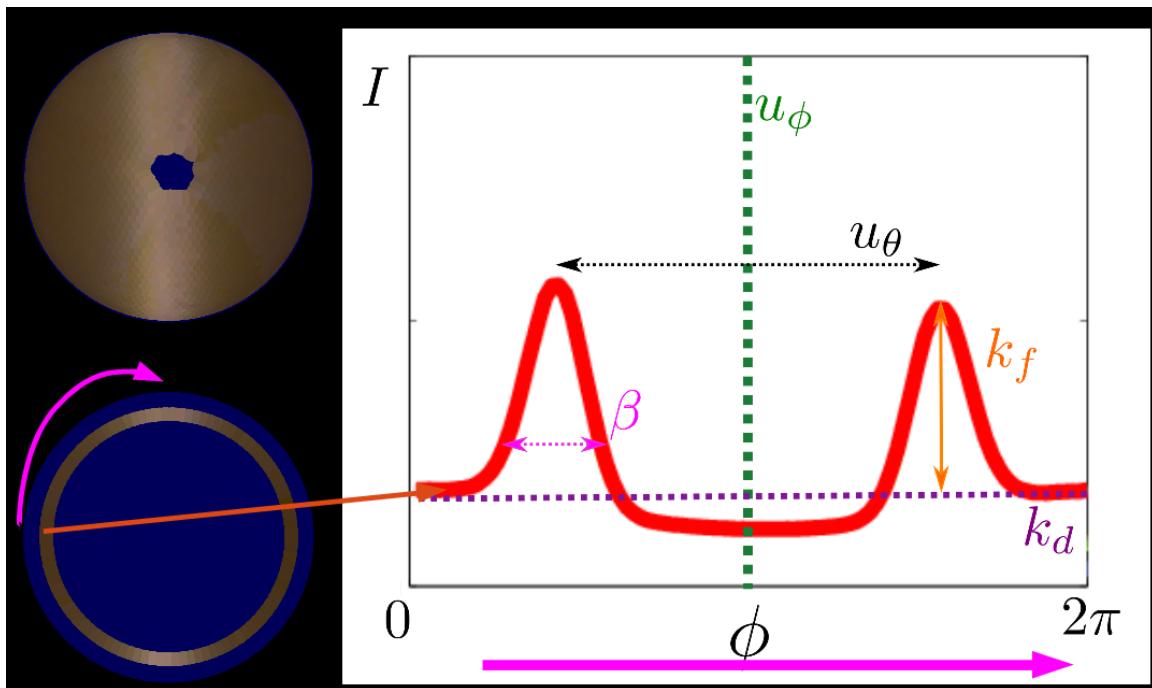


Figure 3.2: Wood Parameters on a 1D orbit: Sparse sampling of the hemisphere plotted as a 1D function to illustrate relevant parameter discriminative features

lying BRDF model allows us to significantly improve on both of these aspects. Fitting the model is performed in two steps: We first determine the fiber direction and highlight width using the bilateral symmetry of the signal and then proceed to use those values to fit the diffuse and fiber colors.

### 3.2 Finished Wood Appearance Model

Empirical evidence suggests that the model put forward by [31] successfully represents the appearance of finished wood, so we adopt it in its existing form. We fit the model parameters according to the same wood BRDF as outlined below, which we reproduce here for the sake of completeness, and illustrate in Fig




Figure 3.3: Orbit Rings: An illustration of the range of possible projected BRDF for different locations of the surface of our strongly anisotropic material samples. Note the 1D dotted gray orbit intercepting the strongly anisotropic signal colored in orange




Figure 3.4: Wood Structure: Wood contains rough cylindrical air tubes (source: NC Brown Center for Ultrastructure Studies, SUNY, Syracuse, NY)

3.9 (readers should see [31] for more details). The following four model parameters need to be determined: the diffuse color ( $k_d$ ), fiber color ( $k_s$ ), fiber direction ( $u$ ), and highlight width ( $\beta$ ) satisfying the wood BRDF equation:

$$f_r(v_i, v_r) = f_s(v_i, v_r) + T_i T_r (\rho_d + f_f(u, v_i, v_r))$$

where:

$T_i$  and  $T_r$  are the Fresnel transmission and reflection coefficients respectively

$$f_f(u, v_i, v_r) = k_f \frac{g(\beta, \psi_h)}{0.5 \cos^2(\psi_d)} \quad (3.1)$$




Figure 3.5: Many ways of cutting wood: Different cuts expose different fiber orientations (source: [Beals and Davis 1977]).

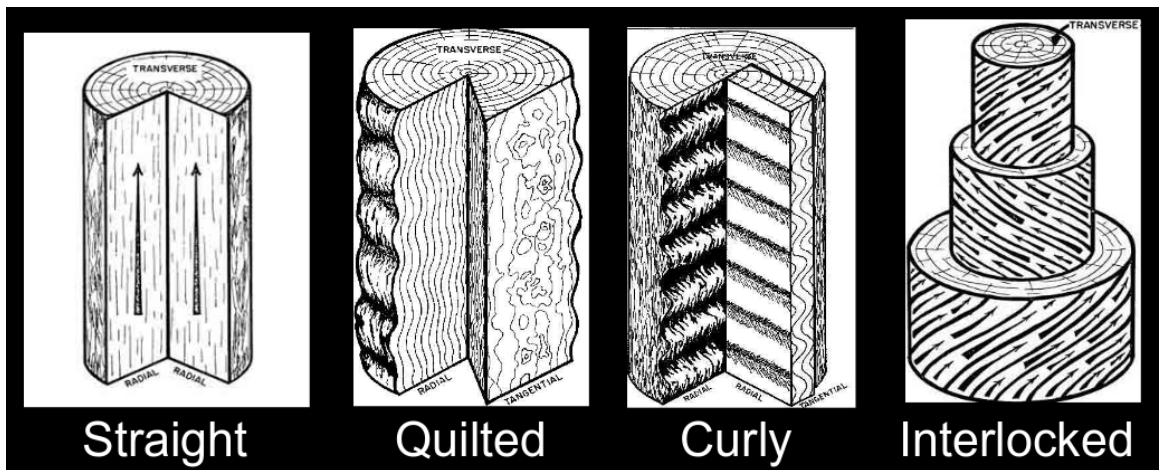


Figure 3.6: Grain Patterns: Many different wood distributions depending on the wood species (source: [Beals and Davis 1977]).




Figure 3.7: Illustration of Subsurface highlights from wood surface (source: Marschner 2005)

$$g(\beta, \psi_h) = \frac{1}{\beta \sqrt{2\pi}} e^{-\frac{\psi_h^2}{2\beta^2}} \quad (3.2)$$

$$\psi_d = \psi_r - \psi_i \quad (3.3)$$

$$\psi_h = \psi_r + \psi_i \quad (3.4)$$

$$\psi_r = \sin^{-1}(s(v_r) \cdot u) \quad (3.5)$$

$$\psi_i = \sin^{-1}(s(v_i) \cdot u) \quad (3.6)$$

(see Table A.1 for other symbol definitions)

We observe significant improvements in the results by fitting the highlight width parameter for every pixel instead of just using a user determined constant value for the whole surface.




Figure 3.8: Scattering from cylindrical geometry. Reflection from interfaces parallel to an axis end up preserving the inclination (source: Marschner 2005)

### 3.3 Data Acquisition

We use two different setups for the measurements: (a) The validation measurements are performed on a 2-DoF spherical motion device [26]. The camera (Canon EOS 70D) is fixed looking perpendicular to the surface. Our camera is used in combination with a macro lens which gives us a spatial resolution of  $10\mu\text{m}$  per pixel. For orbit measurements the light source is moved to equally spaced positions,  $3.6^\circ$  apart for a total of 100 measurement positions, in an orbit around the sample using a fixed incidence angle, of  $60^\circ$  from the normal (see Fig 3.13). For each light source position we capture multiple exposures (usually 2-3 images are sufficient to capture the dynamic range) which are combined

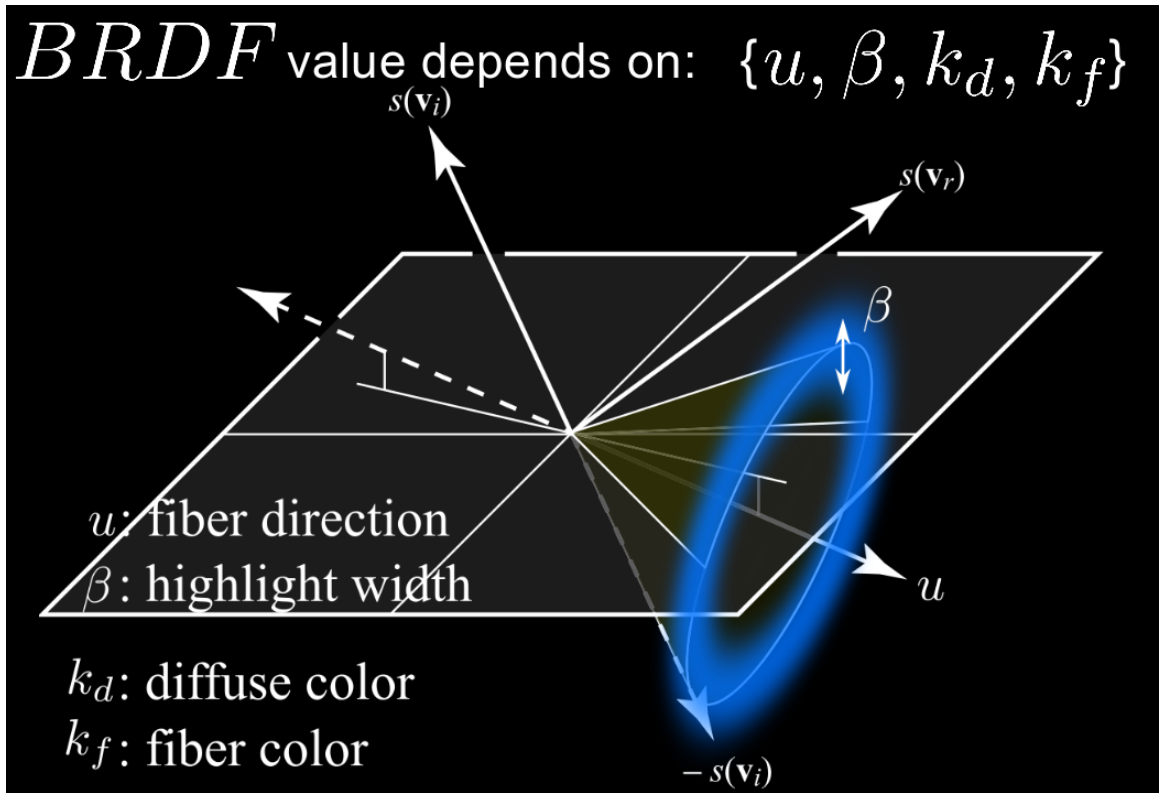


Figure 3.9: Wood BRDF parameters illustration

using HDR techniques to reconstruct the linear response signal. Dense measurements for the validation set use an equi-solid-angle distribution covering the hemisphere with approximately 1800 samples. (b) The prototype version of our measurement device makes use of a battery operated addressable LED strip attached to a PVC pipe loop with similar dimensions to the gantry setup (see Fig 3.16). A second version of the prototype is motorized and makes use of lego modular bricks loop. We attempt to make the surface of the material sample as flat and smooth as possible since the goal is to measure the subsurface behavior and not any surface highlight effects that may arise from the geometry of the interface. Planar wood samples are prepared by first sanding the surface which is followed by several applications of a lacquer or varnish finish. The camera setup




Figure 3.10: Three dimensional perspective on the left and topdown view on the right used in subsequent illustrations (such as in Fig 3.11)

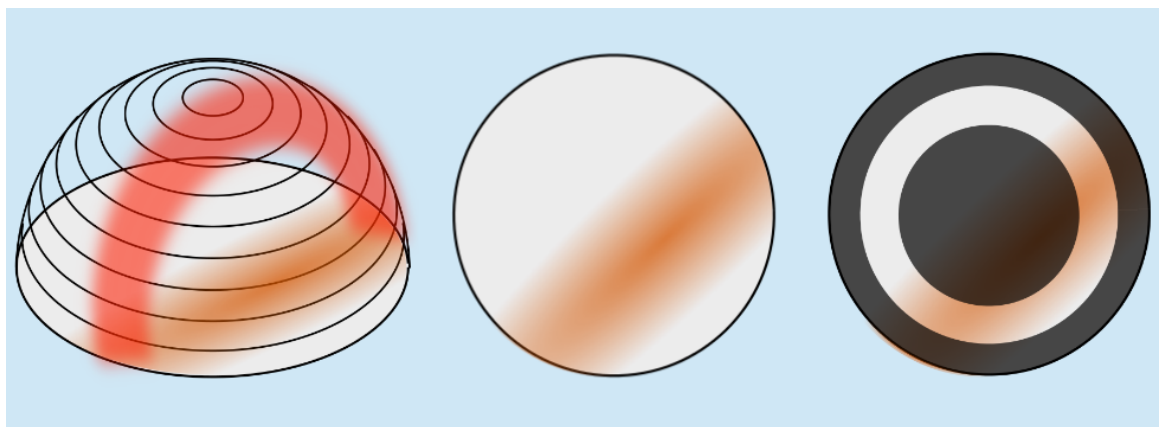


Figure 3.11: The Subsurface highlight distributions: Visualization of the strongly Anisotropic BRDF

is calibrated using images of a reference checkerboard. Light source positions are imaged using a mirror and used to build a calibrated model of the gantry



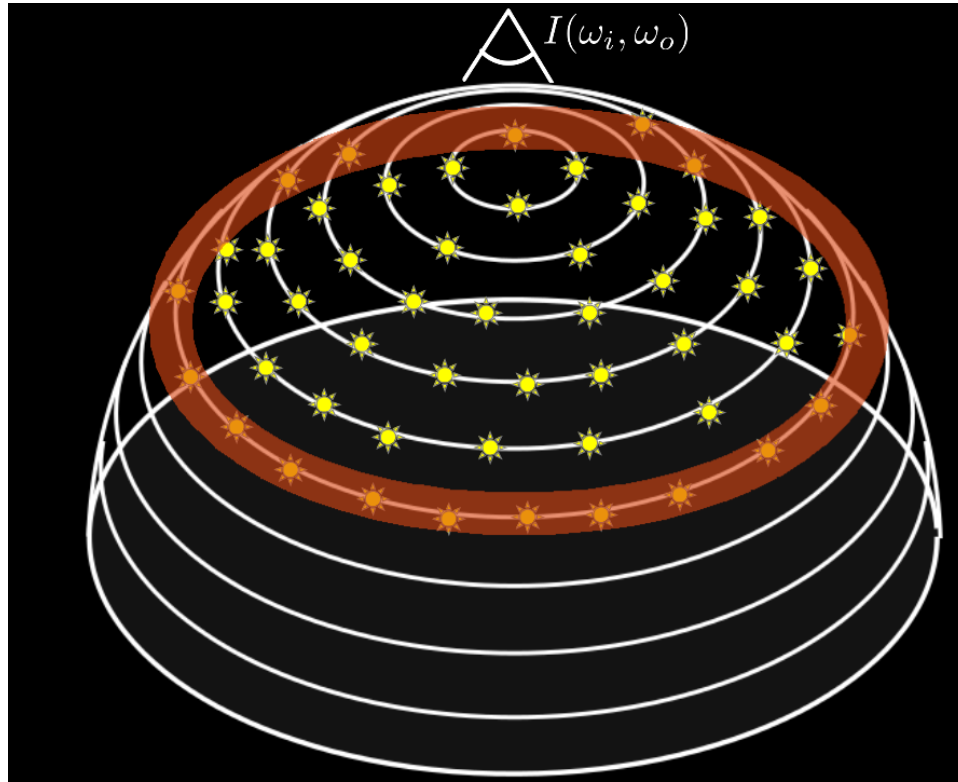


Figure 3.12: Sparse Samples: Hemisphere of Illumination Direction with a highlighted single Orbit of Measurements

using the Caliber software package [28]. The known light positions are used to correct for irradiance variations across the sample plane, due to distance and angle (see Fig 3.15). Additional care has to be taken with the stock Canon camera to correct for shutter speed and aperture variations. During measurement the multiple exposures are merged to a high dynamic range image, demosaiced and then downsampled to produce each measured image. White balancing is applied based on measurements of a Colorchecker calibration target and optical tracking of embedded targets is used to check sample alignment between measurements.




Figure 3.13: Illustration of the Capture Setup: The camera is placed directly above and images are captured with a sequence of light positions around a ring

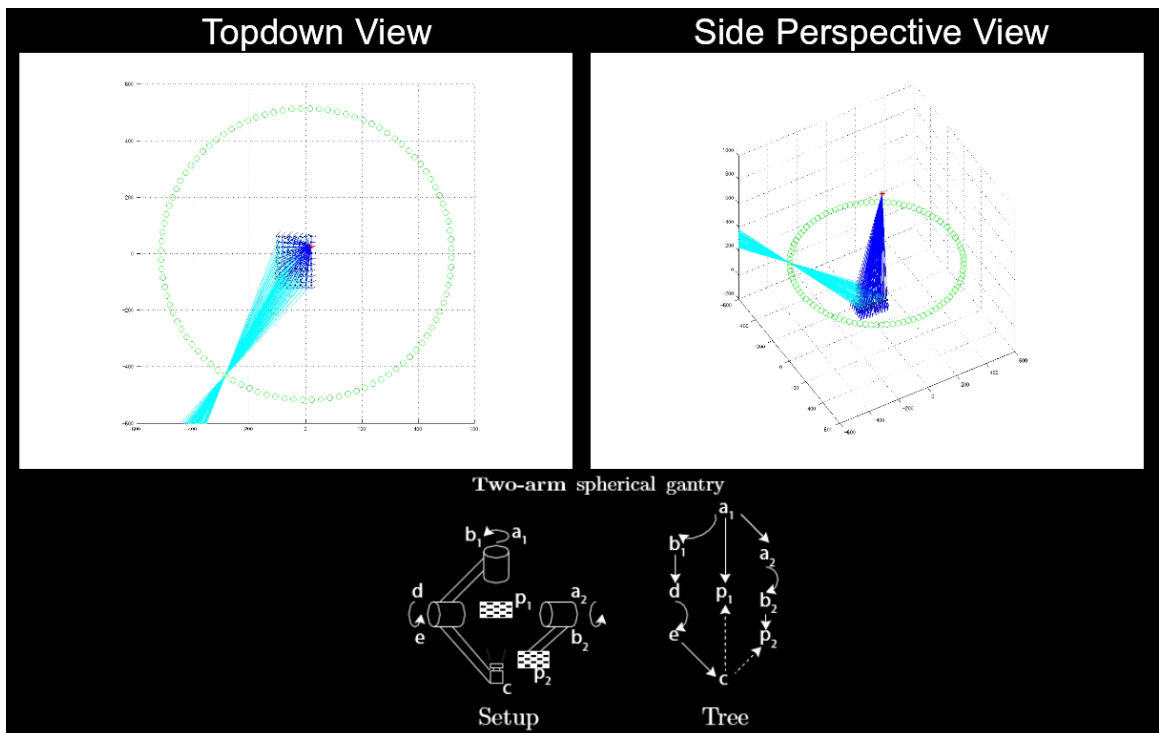


Figure 3.14: Caliber software package [28]: rigid constraint optimization for localization of the camera and light source



Figure 3.15: Irradiance Correction: the captured data is scaled using the incident illumination to recover measured BRDF values




Figure 3.16: LED Loop: A photo of our portable addressable LED loop device (left) and the evolution of the prototype (right)

### 3.4 Parameter Fitting Procedure

The fitting technique we are about to discuss is simple, scalable and outperforms blackbox general purpose optimizers. Our goal in parameter fitting is to take advantage of the special form of the function we are fitting. We observed that the highlight is linear, which means it can be observed adequately just from a ring. But more specifically, the subsurface highlight is a 1D function: it depends only on the projection of the light direction onto the fiber axis. This means we can expect to see symmetry across that axis, and once we identify it then we only have a 1D function to fit (Fig 3.2).




Figure 3.17: Illustration of the subsurface highlight cone with different elevation for the wood fiber (Top images yellow fiber orientation nearly parallel to the surface; bottom images yellow fiber orientation significantly slanted out of the plane)

Our measurement setup examines just the reflection from subsurface fibers,

as a function of incident direction. When we look at this function we expect to see a bright reflection when the half vector (between the refracted light and view directions) is perpendicular to the fibers. This condition is met when the light direction is on a cone centered on the fiber direction: on the hemisphere, a circle. By observing this cone, we need to find the fiber direction, the width of the reflection band around the cone, and the specular and diffuse colors. Previous work identified the cone by measuring the whole hemisphere and fitting a plane, but our method makes the observation that the whole hemisphere is not needed; just the data from one incidence angle determines the cone, and the other parameters follow. To gain some intuition it is useful to think about the case where the fiber direction is close to parallel: we will observe two peaks where the fixed-incidence-angle path crosses the cone; they are diametrically opposite when the fiber is parallel to the surface and move closer together as the fiber direction lifts out of the surface, and they are symmetrically arranged around the fiber direction. Our fitting approach does not work exactly this way, but the intuition of a pair of peaks is a relevant idealization.

We fit the measured data to the model in two steps (as outlined in Algorithm 1; see Fig 3.1 for an illustration of the steps in our fitting approach). Each pixel of the planar sample is processed independently. For each location  $I(x, y, \phi)$  in the input HDR images, where  $\phi$  denotes the  $n$  light positions forming a constant elevation orbit around the sample and  $x, y$  correspond to pixel coordinates in image space, we observe a potentially truncated gaussian signal profile that slices through the subsurface reflection cone (see Fig 4.4, Fig 3.3). There are three possible cases: the set of observations at our data points pass through the highlight cone once, twice or not at all. We consequently expect to observe a projected gaussian lobe in our signal at each respective transit.

The whole signal, and therefore our slice of the signal, is expected to be bilaterally symmetric across the plane containing the fiber direction and its projection into the surface. Identifying the symmetry plane tells us the azimuth of the fiber direction, and it is easy to do robustly as a first step. Furthermore, the highlight cone will be symmetric around the projected direction of the fiber (elongated feature). Note that we make the assumption that there is a single fiber direction and a highlight width which does not depend on the color channel.

## Fiber Orientation and Highlight Width

Our first operation is the conversion of the RGB- $\phi$  signal for the pixel into a grayscale signal that we use for the  $u$  and  $\beta$  fitting step. An approximate projected fiber symmetry axis in the sample plane is determined using a brute force iteration over the 50 ( $n/2$ ) possible positions at which we can fold the 1D signal and subtract it from itself searching for the bilateral symmetry. The score is scaled by a penalty factor that discourages setting the slicing plane at locations where the highlight intensity is high (this is needed since it is possible for the signal to have multiple axes of symmetry; in particular when the fiber is nearly planar there is an additional symmetry with a  $90^\circ$  rotation). The  $L_2$  minimum value is further refined by supersampling the signal and using bisection to resolve a somewhat more precise axis of symmetry. Once we know the fiber symmetry axis  $u_\phi$  we project the 1D grayscale signal into the plane of symmetry. The resulting signal has a noisy truncated gaussian profile.

At this point we know the azimuth of the fiber direction and need to find




Figure 3.18: Symmetry Projection

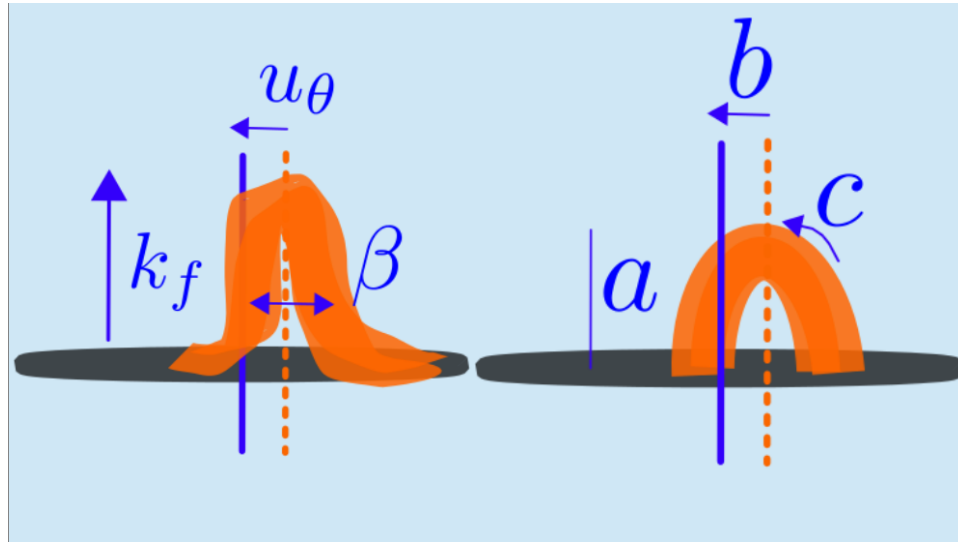


Figure 3.19: Labeled parameters on the projected gaussian and respective logarithmic space parabola

the elevation and width. The key observation behind this next step is that the signal is supposed to be a function of the angle between the half vector and the fiber, and under a small angle assumption this is simply a function of the horizontal component of the projection of the half vector into the symmetry plane. Thus all we really have to do is consider our data as a set of (angle, reflectance) pairs and fit a gaussian to it. An Iterative-Reweighed Least Squares (IRLS) approach is used to fit the above gaussian profile. The weights are initialized to the value of the signal at each datapoint (the stronger the signal, the higher the initial weight). We empirically chose to subtract the bottom 6% percentile



as the base signal to account for the diffuse offset of the gaussian signal and to suppress outliers at the minimum value. Subtracting the 6th percentile is a heuristic diffuse/specular separation, which will be followed later by a more principled estimate in the linear fit.

We observe that our model will have the following gaussian form as a 1D profile which turns into a parabola once transformed to log-space (Fig 3.21):

$$\rho_d + Ae^{\frac{(x-\mu)^2}{2\sigma^2}}$$

$$\ln(A) + \frac{(x-\mu)^2}{2\sigma^2}$$

$$f(x) = cx^2 + bx + a$$

The following IRLS iteration with the corresponding objective function (Fig 3.20) can be applied:

$$\begin{bmatrix} \Sigma y_{(k-1)}^2 & \Sigma xy_{(k-1)}^2 & \Sigma x^2 y_{(k-1)}^2 \\ \Sigma xy_{(k-1)}^2 & \Sigma x^2 y_{(k-1)}^2 & \Sigma x^3 y_{(k-1)}^2 \\ \Sigma x^2 y_{(k-1)}^2 & \Sigma x^3 y_{(k-1)}^2 & \Sigma x^4 y_{(k-1)}^2 \end{bmatrix} \begin{bmatrix} a^{(k)} \\ b^{(k)} \\ c^{(k)} \end{bmatrix} = \begin{bmatrix} \Sigma y_{(k-1)}^2 \ln y \\ \Sigma xy_{(k-1)}^2 \ln y \\ \Sigma x^2 y_{(k-1)}^2 \ln y \end{bmatrix}$$

The signal is transformed to logarithmic space where we now fit a parabola through our data [18]. Note that we need to use a constrained least square approach since the value of  $c$  needs to be negative corresponding to positive (physical) highlight widths. We first solve the unconstrained problem; if the minimizing solution ends up with a positive  $c$  parameter we solve again for the other parameters with  $c = \kappa$  where  $\kappa$  is within the admissible domain and is chosen to be a small negative value (typically  $\kappa = -0.001$ ). After each iteration the residual is computed using the new values of the parameters for the ex-

$$\min \|diag(f^{prev})(\ln[f^{new}] - \ln[y])\|^2$$

where  $f_j^k = A \exp \frac{(x_j - \mu)^2}{\sigma^2}$  at step k and  
 $y =$  measured signal

$$y^{(k)} = \begin{cases} \hat{y} & \text{for } k = 0 \\ e^{a^{(k)} + b^{(k)}x + c^{(k)}x^2} & \text{for } k > 0 \end{cases}$$

Figure 3.20: IRLS objective function and parameter minimization (source: [18])

pected gaussian model. The value  $y_{k-1}$  of the model from each iteration is used as a weight in estimating the parameters  $a_k, b_k, c_k$  for the next iteration. The IRLS process terminates if we are within  $\xi$  of the observed gaussian signal, the partial derivatives of the parabola parameters have reached a certain threshold or if 10 iteration steps have been reached. The value of  $\xi$  is chosen to be 1% since prior studies have demonstrated that small differences in intensity values (below around a 2% threshold) are not perceivable by human eyes. The linear system shown above summarizes the iteration step. The resulting parameters are used to determine the  $u_\theta$  of the fiber direction using the  $\mu$  parameter of the fit and the highlight width of the gaussian lobe that is computed from the final  $\sigma$  parameter.




Figure 3.21: Gaussian Transform: Parabola Parameters in Log Space

## Diffuse and Subsurface Fiber Colors

So far we have used the proposed bilateral symmetry approach and fitted the projected gaussian signal to determine the fiber direction and highlight width, now we just need to calculate the remaining two parameters of the model for the two colors (one for the diffuse component  $k_d$  and one for the subsurface specular component  $k_s$ ). Having determined the fiber direction and highlight width we can go on to perform the second step in the fitting approach to compute the diffuse color and subsurface color components of the BRDF model. This is a linear problem.

We set up the system of equations as shown in line 24 of Algorithm 1 using the gaussian BRDF model (Sec 3.2). We multiply each color channel of the measured data with the Moore-Penrose pseudoinverse of the above system of equations to compute the corresponding color component parameters, effectively splitting the measured signal into the diffuse component and subsurface component for each color channel.




Figure 3.22: Parameter Map result of the wood model for the oak sample using our proposed fitting approach

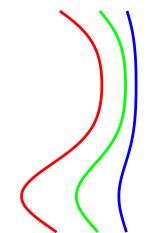
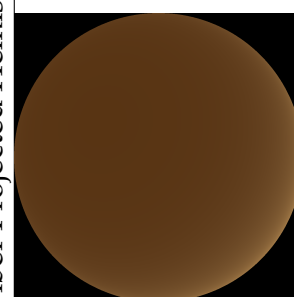
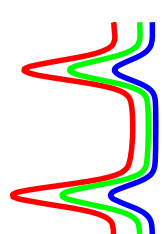
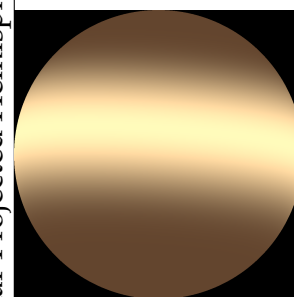
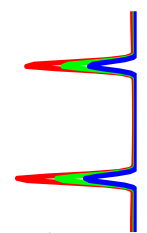
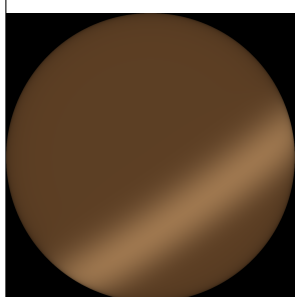
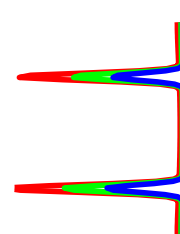
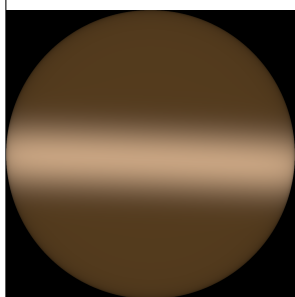
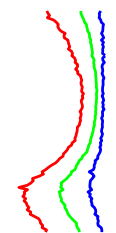
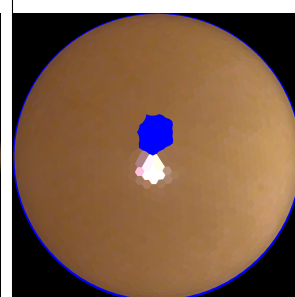
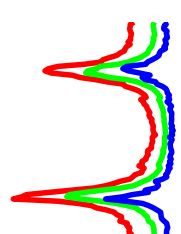
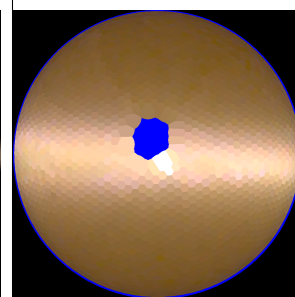
	Steep Fiber [ $\theta = 60^\circ$ ]	Steep Fiber Projected Hemisphere	Planar Fiber [ $\theta = 60^\circ$ ]	Planar Projected Hemisphere
SYM-Orbit				
Wood-SRM05				
Ground Truth				

Table 3.1: Fit Comparison: We show results for two different sample locations. The top row shows the output of our approach compared to previous work (middle row) and the measured validation data (bottom row)

## CHAPTER 4

### RESULTS AND ANALYSIS

#### 4.1 Overview

The symmetry-aware parameter fitting approach described in the previous chapter (Sec 3.4) produces parameter maps of comparable quality to prior work at a fraction of the previous end-to-end time as shown in Fig 3.22 and Fig 4.18. Sec 4.2 compares the rendered results to the dense ground truth measurements in more detail for several representative samples. In Sec 4.3 we contrast the results of the proposed model against the equivalent fitting approach found in [31] (and [23]) showcasing that a variable highlight width map is critical for a better parameter fit at the used  $10\mu m/px$  resolution. The subsequent section (Sec 4.4) gives an analysis of the numerical behavior of our IRLS based fitting approach which is seen to converge as well as remain robust to small perturbations in the input data. For additional visuals, presentation slides and video results please visit the project website [43].

#### 4.2 Ground Truth Validation

We have qualitatively tested the synthesized sequences using the wood parameter maps obtained by our approach. The results are perceptually close and the overall absolute difference between the raw measurements and the generated sequences is relatively small. Readers are invited to visit the project website where they can view higher resolution image content as well as video sequences

and assess the results firsthand. Checkerboard sequences are generated by alternating blocks of the measured data and synthesized sequences (see Fig 4.1) which result in almost seamless animations consequently demonstrating that our appearance model can produce almost indistinguishable results when compared to the ground truth Fig 4.2. Furthermore, visualization of the absolute pixel difference between the measurement and generated results from the fit only results in relatively dark frames with small numbers of sparkling pixels. The difference could most probably be primarily explained by accounting for the deviations from perfectly flat sample surfaces to model the grooves and dents of some of the roughness of the material surface that can result in the unaccounted specular highlights seen in the measured data.

In Fig 4.3 and Fig 4.4 we can see a typical target sample and progressively zoom into the surface where we visualize the captured BRDF data both for the full hemisphere validation set and the sparse equivalent sparse ring measurements used by our approach.

We compare the output result from the model evaluation, using the extracted parameter maps, with the observed measurements under both the input as well as under novel viewing and illumination configurations. A few surface points are also selected to assess the quality of the fit using BRDF plots for the reflected intensity. The overall visual match appears to be quite good for the range of the acquired samples, fitting a range of strongly anisotropic material samples, and can be evaluated further in the accompanying supplementary content and videos [43].

The orbit approach is able to fit significantly tilted fibers where only a truncated gaussian peak is observed. The projected hemisphere BRDF and 1D



Figure 4.1: Checkerboard Comparison to Measurement Data: Interleaved blocks of acquired real data and synthesized fitted model output

orbit signal plots (Fig 4.5-4.7) can also be used to confirm that our model output is closely matching the expected behavior so long as the observed data can be represented using the modular parametric BRDF model for the material which turns out to be true for the vast majority of sample pixels (more than 98% of the pixels have less than 15% relative error; see plot 4.17). Our approach produces almost indistinguishable results when compared to the captured photographs.





Figure 4.2: Checkerboard Comparison of Fitted Result to Measurement Data (top image oak; bottom image mahogany). Note the grid boundaries are barely visible indicating a close match




Figure 4.3: Progressively zooming into our captured sample surface to view a small region of the Spatially-Varying wood BRDF)

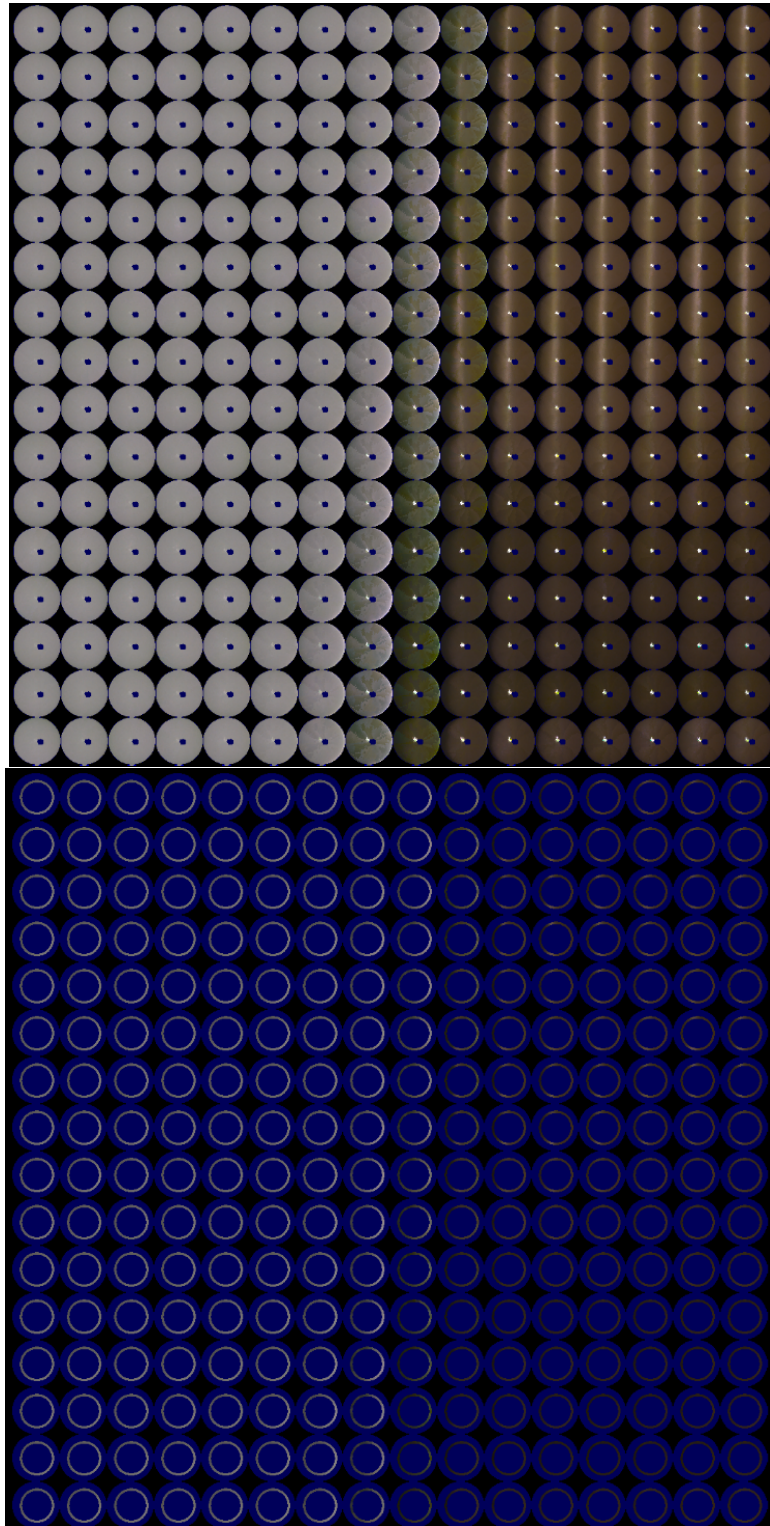


Figure 4.4: Closeup of Fig 4.3 for the area inside the red rectangle. Top image showing the dense validation BRDF data while the bottom image shows the ring subset data used by our fitting algorithm

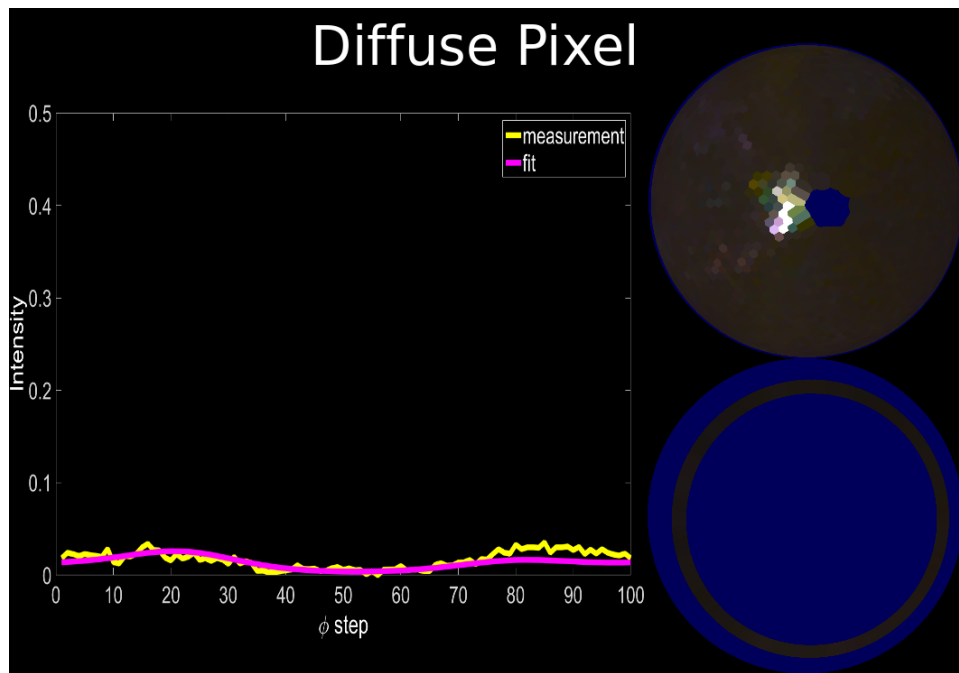


Figure 4.5: Our symmetry model fit of a diffuse location on the surface




Figure 4.6: Our symmetry model fit of a fiber with a large elevation angle (truncated gaussian)




Figure 4.7: Our symmetry model fit of a fiber parallel to the surface

### 4.3 Comparison to Previous Work




Figure 4.8: The SRM05 [31] fitting approach

Our results compare favorably when viewed against the output of the SRM05 [31] version of the fitting algorithm. We have improved upon the end-to-end time as well as the quality of the output with results that are closer to the ground truth data. The major reason for the closer match is that our approach automatically fits the highlight width instead of using a fixed value as is the case in the SRM05 approach (Fig 4.10). The fitting approach has been accelerated by two orders of magnitude using more sparse data and a more efficient nonlinear least squares optimization routine (Fig 4.11).

We compare our symmetry-aware parameter extraction performance and result quality against the models and approaches from Ward [49], [31] and [23] for the metal, finished wood and woven fabric models respectively. Note that the results for metal and fabric were omitted from this writeup where we focus on

the wood BRDF as an illustrative example. Table 4.11 shows that our technique significantly reduces the capture time and storage requirements with near interactive performance while maintaining the closest match to the measurements. Fig 4.8 and Fig 4.13 illustrate that even though the highlight width in [31] can be tuned to give plausible results it is highly unlikely to represent the actual variability in the particular parameter for most real life finished wood samples once we reach a certain level of resolution. The Ward model cannot represent the out of plane axis and consequently cannot reproduce the anisotropic behaviour of the wood samples.

We also see that the prior approach and our symmetry based approach disagree somewhat on the fiber map orientation (Fig 4.14). Both the quantitative and qualitative comparison to ground truth hint that we are closer to the measurements than the prior work. The prior approach was unable to fit diffuse pixels on the surface since the constant highlight width would have to be set to a large value in such a case effectively defaulting to not showing any anisotropy under such a scenario.

The SRM05 results in an incorrect fiber direction when there is a big mismatch between the observed highlight width and the user specified parameter. By that logic it is also unable to give accurate fiber orientations for input data with spatially-varying highlight widths and as such we cannot make a more meaningful comparison of the accuracy of our fiber orientation map to that of the prior work (see Fig 4.15-4.16). It follows that the output of this approach should be evaluated using primarily the measured validation data.

## 4.4 IRLS Convergence, Fit Quality Analysis and Parameter Sensitivity

In practice the proposed symmetry projection fitting approach and IRLS algorithm have a certain number of limitation and cases where we expect them not to produce the desirable results. We can argue that our observed data falls under three broad categories of signals: (a) the model fits the data and converges to the correct parameters, (b) the model fits the data but converges to the wrong values, and (c) the model does not fit the data. There is not much we can do for case (c) other than potentially detecting such a scenario and failing gracefully at such locations. As far as the other two cases go (a) is the desired outcome and we want to minimize the number of (b) cases as much as possible. We have performed a set of experiments where we use a range of gaussian signals spanning the whole domain of possible fiber orientations and highlight widths. We go on to degrade the signals with additive noise (zero mean) and observe how robust the fitting approach is to such input. Under ideal conditions, we expect to have a single global minimum since we are fitting a parabola in log-space, which is convex. Some crude experiments were performed to fit other strongly anisotropic materials, as well, such as is the case with specular highlight from the fibers of woven fabrics. As mentioned, other modular plugins can be developed in the future for any material that is compatible with our fitting approach so that a more comprehensive analysis can be performed with real world input data.

In an attempt to determine how sensitive the model is to small perturbation in the input, we experiment with moderate degradation in the input signal and



also observe how stable the fitting approach output parameters are under perturbation. The approach is most sensitive to the fiber direction though we can still tolerate up to 2 deg–4 deg change in the orientation without suffering perceivable changes in the output. Significant changes in the highlight width can result in noticeable changes for the generated animation sequences while not being immediately obvious in static images (a change of up to 10%–20% can be made without any visible effects). The diffuse and highlight colors are much less sensitive to small perturbations in the input.

The approach converges with ease even under high amounts of noise for fiber orientations that are effectively parallel to the surface, which exhibit the gaussian spikes from slicing the centered fuzzy highlight cone. The symmetry based approach also automatically behaves appropriately when it encounters diffuse signals (absence of gaussian spikes). The highlight width increases and the fiber color decreases resulting in a primarily matching diffuse signal. The plots in Fig 4.19 show that despite the use of an adaptive iterative approach we manage to achieve fast convergence and only need to take between 3 – 5 steps before the fit reaches a less than 2%  $L_2$  threshold while we also take an extra precaution and limit the maximum number of iterations to avoid spending an unbound amount of computation on signals that do not appear to converge.

Our cumulative relative error plot shows that we have a certain amount of error for almost all of the sample pixels. This is most likely due to the fact that we closely follow the 1D signal profile but do not exactly match it. There is a certain amount of noise in the signal which probably contributes 5%-10% of the relative error. Beyond 15% relative error almost 98% of our pixels match the ground truth and around 20% relative error virtually all of the pixels are match-

ing the measurements other than a small handful of outliers. The deviations can probably be explained by pixels with mixed signals due to registration errors or sample surface defects where the surface is not perfectly flat. Despite this numerical deviation the phenomenological behavior of the fitted result is remarkably similar to the measurement.

In practice, the proposed approach appears to extract high quality parameter maps for a wide variety of wood samples that also exhibit the whole spectrum of fiber orientation distributions.




Figure 4.9: Comparison of BRDF fit for a region against ground truth and a prior wood appearance model (Marschner et al SG05). The center column consists of ground truth measurements for a region of the walnut sample at different magnification levels (rows). The left and right columns correspond to the results generated by the new symmetry aware model and the dense measurements plane fitting approach respectively. Note that the rendered results do not include the surface specular component. Blue regions denote areas of missing data in the ground truth.




Figure 4.10: Compare Paper Target Diffuse Fit (SRM05 left; our approach right)

	Capture Time [min]	HDR Time [min]	Fitting Time [min]	Total Time [min]	Storage [Gb]
<b>SRM05</b>	~290	~60	~360	>700+	130.8
Hemi Fit					
<b>SYM</b>	~15	~5	~4	~24	7.4
Orbit Fit (ours)					

Full :  $57 \times 31 \times 4$  exposures = 1757 [4752 x 3158] images

Orbit :  $100 \times 4$  exposures = 100 [4752 x 3158] images

Figure 4.11: Performance Comparison to [Marschner 2005]




Figure 4.12: Comparison of Oak Sample Pixels

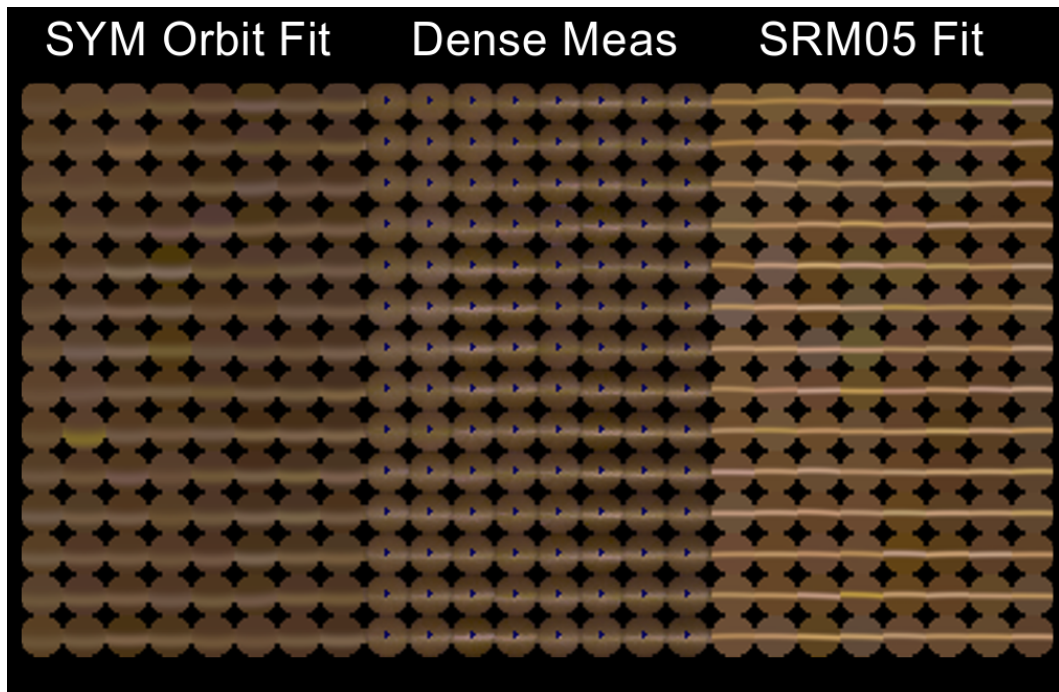


Figure 4.13: Closeup Comparison of Oak Sample Pixels

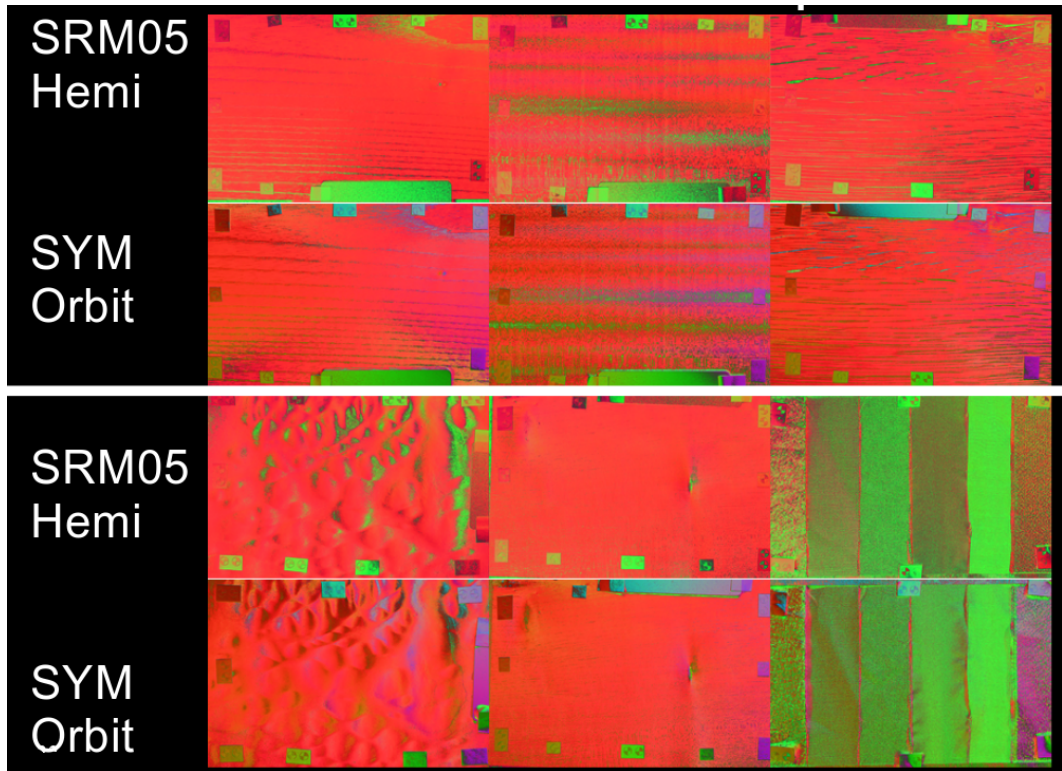


Figure 4.14: Comparison of the fiber parameter maps output from our approach and the prior work [31]




Figure 4.15: Comparison of 1D signal Fit: Reasonable fit for both fitting approaches

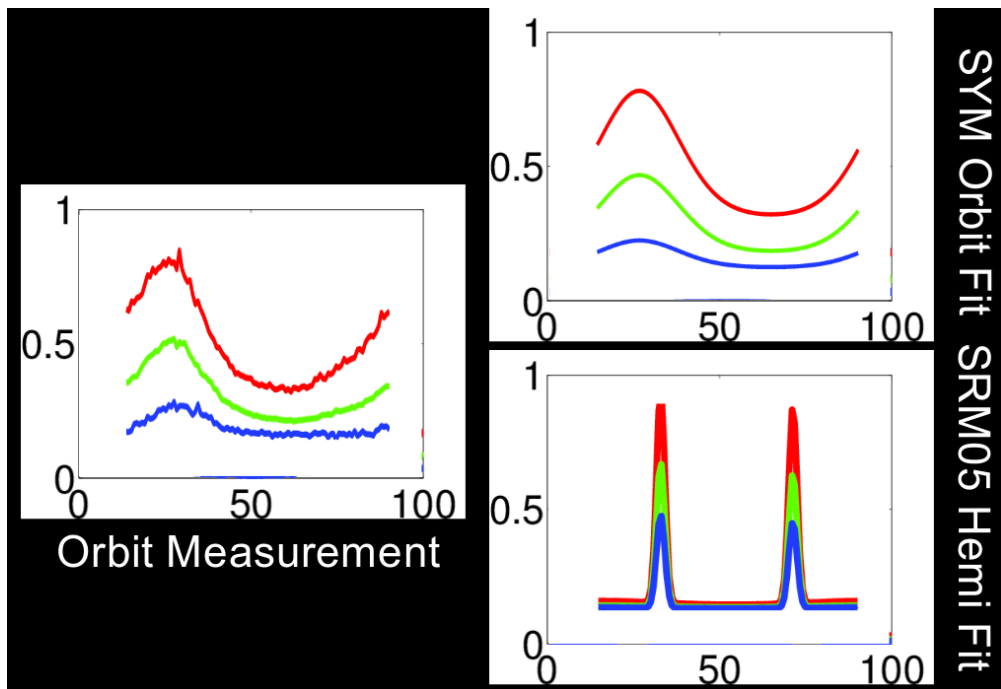


Figure 4.16: Comparison of 1D signal Fit: Our symmetry based approach matches while the prior approach was unable to converge

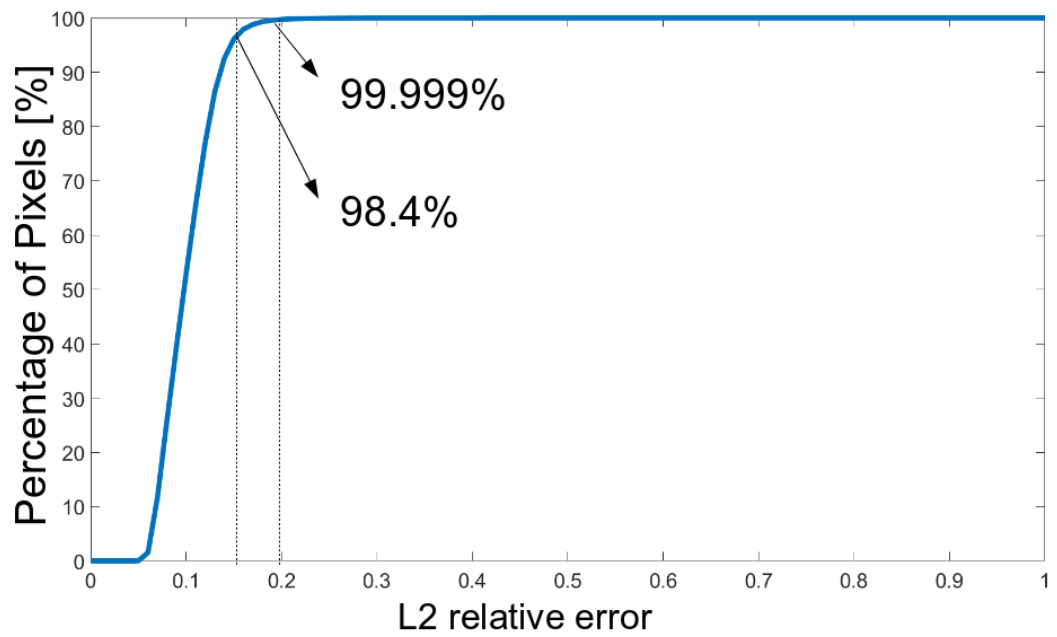


Figure 4.17:  $L_2$  distance plot for several wood samples. A cumulative histogram showing the population of pixels within a certain bucket of L2 error from prior work and measurements.






Figure 4.18: Extracted Parameter Maps for Several Wood Samples. Columns (left to right): Diffuse Color, Fiber Color, Highlight Width, Fiber Direction (Ours), Fiber Direction (SRM05)




Figure 4.19: The approach is robust and can converge even when there is significant noise in the signal while not being sensitive to the initialization (top figure shows convergence to ground truth parameter). The IRLS algorithm exhibits fast convergence to a fixed point within a few iterations steps as can be seen from the above plots (bottom figure shows the diminishing parameter gradients).




Figure 4.20: Pseudocolored clusters of closely matching BRDF pixels (top image) and one white pixel representatives of each unique bin of pixels (bottom image). Less than 10% of signals are unique in sufficiently large samples

## CHAPTER 5

### CONCLUSION

#### 5.1 Limitations and Future Work

We have shown that the sparse symmetry material acquisition approach can produce high quality parameter maps for a range of strongly anisotropic materials at a fraction of the time (close to an order of magnitude compared to prior work) while being reasonably robust to noise. The current formulation of our optimization approach is restricted to flat planar surfaces. The elevation angle of the illumination source from the surface normal and resulting half-vector limits the range of fiber elevations over which we can observe the strongly anisotropic signal. Despite this the output quality of the approach degrades gracefully and we can still get reasonable fits with limited highlight observations. We observe that a lot of pixels have very similar 1D signals. Future work will aim to alleviate this shortcoming using a refinement step where we utilize an ensemble of additional orbit illumination configuration measurements with translation in the plane parallel to the sample surface. A small set of such measurements within the immediate vicinity of the sample should theoretically be enough to sufficiently sample the hemisphere and observe the anisotropic highlight signal. We have taken steps towards this with our automated LED device illustrated in Fig 3.16. The proposed measurement technique can make use of the illumination loop mounted on an autonomous device. Such an approach would allow for additional automation where the rover device can move around the floor of a room and capture the appearance of materials under illumination from the LED loop device. The procedure can then be repeated so that

a larger area of the Spatially-Varying BRDF is stitched together incrementally. We also observe that a lot of the 1D signal profiles exhibit significant similarity. This can be exploited by quantizing the signals which we can consequently use as an encoding allowing us to cluster together large numbers of pixels and consequently avoiding excessive duplicate computation (see Fig 4.20).

## 5.2 Closing Remarks

Physically-based appearance fitting for strongly anisotropic materials is possible with a limited amount of captured data. We show that high quality parameter maps can be reconstructed in a fraction of the time compared to previous techniques using inexpensive equipment. We have demonstrated a bilateral symmetry approach that improved upon prior work and achieved an overall speedup close to an order of magnitude. Our prototype measurement device also helped alleviate the data acquisition bottleneck and several acceleration schemes were proposed to push the approach closer to near interactive performance.

APPENDIX A  
**FITTING APPROACH PSEUDOCODE**

---

### Algorithm 1: Per-pixel BRDF fitting approach

```

1: procedure FITSYMBRDF(orbit_pixel_values)
2:   hdr_reconstruction(raw_image_stack)
3:   color_correct(hdr_stack) / white_balance(hdr_stack)
4:   irradiance_flat_field_correction(hdr_stack)
5:   for each pixel in sample image do
6:     compute grayscale value for 1D orbit signal
7:     for  $\text{symm\_axis}_k = 1 \dots \frac{\text{len}(I_{\text{posn}})}{2}$  do
8:       shift 1D signal start idx by one sample
9:        $\text{score}_k = \|I(1 : \text{midpt}) - I(\text{end} : \frac{\text{midpt}}{2} + 1)\|_2$ 
10:    end for
11:     $u_{\text{azimuth}} = \text{angle2uv}(\min(\text{symm\_score}))$ 
12:    symm_signal = Project 1D orbit signal to  $u_{\text{azimuth}}$  line
13:     $\beta, \mu = \text{IR-LS}(\text{symm\_signal})$  ▷ Iterative-Reweighted Least Squares
14:
15:     $u_{\text{elevation}} = \text{func}(\mu, l_{\text{dir}})$ 
16:     $\rho_d, k_f = \text{FitColor}(u, \beta, l_{\text{dir}})$  ▷ Linear fit diffuse and subsurface colors
17:  end for
18:  return  $u, \beta, \rho_d, k_f$ 
19: end procedure
20: procedure IRLS(symm_signal)
21:  ▷ Initialize weights to intensity at each datapoint  ▷ Update weights using approximate fit
    parameters at given step ▷ Return if residual <  $\xi$  or iter >= 10
22:  return  $\mu, \sigma$ 
23: end procedure
24: procedure FITCOLOR( $u, \beta, l_{\text{dir}}$ )
25:
26:   $A = \begin{bmatrix} 1 & \text{gauss\_reconstruct}(u, \beta, l_{\text{dir}(1)}) \\ \vdots & \vdots \\ 1 & \text{gauss\_reconstruct}(u, \beta, l_{\text{dir}(n)}) \end{bmatrix}$ 
27:   $\rho_d, k_f = \text{pinv}(A) \times \mathbf{1D\_signal}_{\text{rgb}}$  ▷ Moore-Penrose pseudoinverse
28: end procedure

```

---



Table A.1: Parameters of the wood BRDF model

Symbol	Range	Description
$v_i$	$[0, \frac{\pi}{2}]$	incident light angle from normal to the surface
$v_r$	$[0, \frac{\pi}{2}]$	reflected light angle from normal to the surface
$u$	$[-1, 1] \in \mathbf{S}_2$	wood fiber direction
$\beta$	$(0, \infty] \in \mathbf{R}^+$	highlight width
$\eta$	1.55	refractive index of wood/cellulose
$k_d$	$[\mathbf{R}, \mathbf{G}, \mathbf{B}] \in [0, 1]^3$	diffuse color
$k_s$	$[\mathbf{R}, \mathbf{G}, \mathbf{B}] \in [0, 1]^3$	specular color
$k_f$	$[\mathbf{R}, \mathbf{G}, \mathbf{B}] \in [0, 1]^3$	subsurface color
$T_i$	$[0, 1]$	fresnel transmission coefficient
$T_r$	$[0, 1]$	fresnel reflection coefficient
$f_f(u, v_i, v_r)$	$[\mathbf{L}, \mathbf{R}]$	subsurface wood BRDF function component
$f_s(v_i, v_r)$	$[\mathbf{L}, \mathbf{R}]$	specular wood BRDF function component
$f_r(v_i, v_r)$	$[\mathbf{L}, \mathbf{R}]$	reflection wood BRDF function component
$\psi_d$	$[0, \pi]$	Difference between incident and reflected angle (inside the material)
$\psi_h$	$[0, \pi]$	Half-vector of incident and reflected angle (inside the material)
$\psi_r$	$[0, \pi]$	Reflected angle inside the material
$\psi_i$	$[0, \pi]$	Incident angle inside the material
$g(\beta, \psi_h)$	$[\mathbf{L}, \mathbf{R}]$	gaussian centered around subsurface highlight cone

## BIBLIOGRAPHY

- [1] Miika Aittala, Tim Weyrich, and Jaakko Lehtinen. Practical SVBRDF capture in the frequency domain. *ACM Trans. Graph.*, 32(4):110:1–110:12, July 2013.
- [2] Miika Aittala, Tim Weyrich, and Jaakko Lehtinen. Two-shot SVBRDF capture for stationary materials. *ACM Trans. Graph.*, 34(4):110:1–110:13, July 2015.
- [3] Michael Ashikhmin and Simon Premoze. Distribution-based BRDFs. *Technical Report, University of Utah*, 2:6, 2007.
- [4] Moshe Ben-Ezra, Jiaping Wang, Bennett Wilburn, Xiaoyang Li, and Le Ma. An LED-only BRDF measurement device. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1–8. IEEE, 2008.
- [5] James F Blinn. Models of light reflection for computer synthesized pictures. In *ACM SIGGRAPH Computer Graphics*, volume 11, pages 192–198. ACM, 1977.
- [6] Guojun Chen, Yue Dong, Pieter Peers, Jiawan Zhang, and Xin Tong. Reflectance scanning: Estimating shading frame and BRDF with generalized linear light sources. *ACM Trans. Graph.*, 33(4):117:1–117:11, July 2014.
- [7] Kristin J. Dana, Bram van Ginneken, Shree K. Nayar, and Jan J. Koenderink. Reflectance and texture of real-world surfaces. *ACM Trans. Graph.*, 18(1):1–34, January 1999.
- [8] Julie Dorsey, Holly Rushmeier, and Francois Sillion. *Digital Modeling of Material Appearance*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2008.
- [9] J. Filip and M. Haindl. Bidirectional texture function modeling: A state of the art survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(11):1921–1940, Nov 2009.
- [10] Sing Choong Foo. *A gonioreflectometer for measuring the bidirectional reflectance of material for use in illumination computation*. PhD thesis, Cornell University, 1997.

- [11] Yannick Francken, Tom Cuypers, Tom Mertens, Jo Gielis, and Philippe Bekaert. High quality mesostructure acquisition using specularities. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1–7. IEEE, 2008.
- [12] Andrew Gardner, Chris Tchou, Tim Hawkins, and Paul Debevec. Linear light source reflectometry. *ACM Trans. Graph.*, 22(3):749–758, July 2003.
- [13] Abhijeet Ghosh, Shruthi Achutha, Wolfgang Heidrich, and Matthew O’Toole. BRDF acquisition with basis illumination. In *2007 IEEE 11th International Conference on Computer Vision*, pages 1–8. IEEE, 2007.
- [14] Abhijeet Ghosh, Tongbo Chen, Pieter Peers, Cyrus A Wilson, and Paul Debevec. Estimating specular roughness and anisotropy from second order spherical gradient illumination. In *Computer Graphics Forum*, volume 28, pages 1161–1170. Wiley Online Library, 2009.
- [15] Abhijeet Ghosh, Wolfgang Heidrich, Shruthi Achutha, and Matthew O’Toole. A basis illumination approach to BRDF measurement. *International Journal of Computer Vision*, 90(2):183–197, 2010.
- [16] Dar’ya Guarnera, Giuseppe Claudio Guarnera, Abhijeet Ghosh, Cornelia Denk, and Mashhuda Glencross. BRDF Representation and Acquisition. *Computer Graphics Forum*, 2016.
- [17] Giuseppe Claudio Guarnera, Pieter Peers, Paul Debevec, and Abhijeet Ghosh. Estimating surface normals from spherical Stokes reflectance fields. In *European Conference on Computer Vision*, pages 340–349. Springer, 2012.
- [18] Hongwei Guo. A simple algorithm for fitting a gaussian function. In *Streamlining Digital Signal Processing*, pages 297–305. John Wiley and Sons, Inc., 2012.
- [19] Michal Haindl and Jirí Filip. *Visual Texture: Accurate Material Appearance Measurement, Representation and Modeling*. Advances in Computer Vision and Pattern Recognition. Springer, 2013.
- [20] Jefferson Y Han and Ken Perlin. Measuring bidirectional texture reflectance with a kaleidoscope. *ACM Transactions on Graphics (TOG)*, 22(3):741–748, 2003.
- [21] Xiao D. He, Kenneth E. Torrance, François X. Sillion, and Donald P. Green-

- berg. A comprehensive physical model for light reflection. In *Proceedings of the 18th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '91, pages 175–186, New York, NY, USA, 1991. ACM.
- [22] Jack J Hsia and Joseph C Richmond. Bidirectional reflectometry part i. a high resolution laser bi-directional reflectometer with results on several optical coatings. *Journal of Research of the National Bureau of Standards-A. Physics and Chemistry A*, 80:189–205, 1976.
- [23] Piti Irawan and Steve Marschner. Specular reflection from woven cloth. *ACM Trans. Graph.*, 31(1):11:1–11:20, February 2012.
- [24] Wenzel Jakob. Mitsuba renderer. <http://www.mitsuba-renderer.org>, 2010.
- [25] Henrik Wann Jensen, Stephen R. Marschner, Marc Levoy, and Pat Hanrahan. A practical model for subsurface light transport. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '01, pages 511–518, New York, NY, USA, 2001. ACM.
- [26] Mark Levoy. Stanford spherical gantry. <http://graphics.stanford.edu/projects/gantry/>, 2004.
- [27] Hongsong Li, Sing-Choong Foo, Kenneth E Torrance, and Stephen H Westin. Automated three-axis gonioreflectometer for computer graphics applications. *Optical Engineering*, 45(4):043605–043605, 2006.
- [28] Albert Liu, Steve Marschner, and Noah Snavely. Caliber: Camera localization and calibration using rigidity constraints. *International Journal of Computer Vision*, 118(1):1–21, 2016.
- [29] Tom Malzbender, Dan Gelb, and Hans Wolters. Polynomial texture maps. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 519–528. ACM, 2001.
- [30] Stephen R Marschner. *Inverse rendering for computer graphics*. PhD thesis, Cornell University, 1998.
- [31] Stephen R. Marschner, Stephen H. Westin, Adam Arbree, and Jonathan T. Moon. Measuring and modeling the appearance of finished wood. In *ACM SIGGRAPH 2005 Papers*, SIGGRAPH '05, pages 727–734, New York, NY, USA, 2005. ACM.

- [32] Stephen R. Marschner, Stephen H. Westin, Eric P. F. Lafortune, Kenneth E. Torrance, and Donald P. Greenberg. *Image-Based BRDF Measurement Including Human Skin*, pages 131–144. Springer Vienna, Vienna, 1999.
- [33] Wojciech Matusik, Hanspeter Pfister, Matt Brand, and Leonard McMillan. A data-driven reflectance model. In *ACM SIGGRAPH 2003 Papers, SIGGRAPH '03*, pages 759–769, New York, NY, USA, 2003. ACM.
- [34] Yasuhiro Mukaigawa, Kohei Sumino, and Yasushi Yagi. Multiplexed illumination for measuring BRDF using an ellipsoidal mirror and a projector. In *Asian Conference on Computer Vision*, pages 246–257. Springer, 2007.
- [35] Nikhil Naik, Shuang Zhao, Andreas Velten, Ramesh Raskar, and Kavita Bala. Single view reflectance capture using multiplexed scattering and time-of-flight imaging. In *Proceedings of the 2011 SIGGRAPH Asia Conference, SA '11*, pages 171:1–171:10, New York, NY, USA, 2011. ACM.
- [36] Addy Ngan, Frdo Durand, and Wojciech Matusik. Experimental Analysis of BRDF Models. In Kavita Bala and Philip Dutre, editors, *Eurographics Symposium on Rendering (2005)*. The Eurographics Association, 2005.
- [37] F. E. Nicodemus, J. C. Richmond, J. J. Hsia, I. W. Ginsberg, and T. Limperis. Radiometry. chapter Geometrical Considerations and Nomenclature for Reflectance, pages 94–145. Jones and Bartlett Publishers, Inc., USA, 1992.
- [38] Tania Pouli, Erik Reinhard, and Douglas W. Cunningham. *Image Statistics in Visual Computing*. A. K. Peters, Ltd., Natick, MA, USA, 1st edition, 2013.
- [39] Peiran Ren, Jiaping Wang, John Snyder, Xin Tong, and Baining Guo. Pocket reflectometry. *ACM Trans. Graph.*, 30(4):45:1–45:10, July 2011.
- [40] J. Riviere, P. Peers, and A. Ghosh. Mobile surface reflectometry. *Computer Graphics Forum*, 35(1):191–202, 2016.
- [41] Nicolas Riviere, Romain Ceolato, and Laurent Hespel. Multispectral polarized BRDF: design of a highly resolved reflectometer and development of a data inversion method. *Opt. Appl*, 42:15, 2012.
- [42] Martin Rump, Gero Mller, Ralf Sarlette, Dirk Koch, and Reinhard Klein. Photo-realistic rendering of metallic car paint from image-based measurements. *Computer Graphics Forum*, 27(2):527–536, 2008.

- [43] Nicolas Savva. Wood BRDF acquisition and fitting project website. [www.nicolassavva.com/MS-CU16/](http://www.nicolassavva.com/MS-CU16/), 2016.
- [44] Christophe Schlick. An inexpensive BRDF model for physically-based rendering. *Computer Graphics Forum*, 13:233–246, 1994.
- [45] K. E. Torrance and E. M. Sparrow. Theory for off-specular reflection from roughened surfaces\*. *J. Opt. Soc. Am.*, 57(9):1105–1114, Sep 1967.
- [46] Borom Tunwattanapong, Graham Fyffe, Paul Graham, Jay Busch, Xueming Yu, Abhijeet Ghosh, and Paul Debevec. Acquiring reflectance and shape from continuous spherical harmonic illumination. *ACM Trans. Graph.*, 32(4):109:1–109:12, July 2013.
- [47] Bruce Walter, Stephen R. Marschner, Hongsong Li, and Kenneth E. Torrance. Microfacet models for refraction through rough surfaces. In *Proceedings of the 18th Eurographics Conference on Rendering Techniques, EGSR'07*, pages 195–206, Aire-la-Ville, Switzerland, Switzerland, 2007. Eurographics Association.
- [48] Chun-Po Wang, Noah Snavely, and Steve Marschner. Estimating dual-scale properties of glossy surfaces from step-edge lighting. In *Proceedings of the 2011 SIGGRAPH Asia Conference, SA '11*, pages 172:1–172:12, New York, NY, USA, 2011. ACM.
- [49] Gregory J Ward. Measuring and modeling anisotropic reflection. In *ACM SIGGRAPH Computer Graphics*, volume 26, pages 265–272. ACM, 1992.
- [50] Michael Weinmann and Reinhard Klein. Advances in geometry and reflectance acquisition (course notes). In *SIGGRAPH Asia 2015 Courses, SA '15*, pages 1:1–1:71, New York, NY, USA, 2015. ACM.
- [51] Carlos J Zubiaga, Laurent Belcour, Carles Bosch, Adolfo Muñoz, and Pascal Barla. Statistical analysis of bidirectional reflectance distribution functions. In *SPIE/IS&T Electronic Imaging*, pages 939808–939808. International Society for Optics and Photonics, 2015.