

BxDF Material Acquisition, Representation, and Rendering for VR and Design

Giuseppe Claudio Guarnera
University of York
Norwegian University of Science and
Technology (NTNU)
giuseppe.guarnera@ntnu.no

Dar'ya Guarnera
Norwegian University of Science and
Technology (NTNU)
darya.guarnera@ntnu.no

Gregory J. Ward
Dolby Laboratories
Anywhere Software
gward@lmi.net

Mashhuda Glencross
The University of Queensland
mashhuda@gmail.com

Ian Hall
Yulio Technologies
ihall@yulio.com




Figure 1: (a) Surface-based fabric appearance model [Guarnera et al. 2017] for commercial interior design applications. (b) demonstrates data-driven BSDFs and the use of proxy geometry for rendering [Ward et al. 2012, 2014]. (c) Example of BRDF parameter remapping [Guarnera et al. 2019]: source models (on the left), remapped appearance (at the center), error map (on the right).

ABSTRACT

Photorealistic and physically-based rendering of real-world environments with high fidelity materials is important to a range of applications, including special effects, architectural modelling, cultural heritage, computer games, automotive design, and virtual reality (VR). Our perception of the world depends on lighting and surface material characteristics, which determine how the light is reflected, scattered, and absorbed. In order to reproduce appearance, we must therefore understand all the ways objects interact with light, and the acquisition and representation of materials has thus been an important part of computer graphics from early days. Nevertheless, no material model nor acquisition setup is without limitations in terms of the variety of materials represented, and different approaches vary widely in terms of compatibility and ease of use.

In this course, we describe the state of the art in material appearance acquisition and modelling, ranging from mathematical BSDFs to data-driven capture and representation of anisotropic materials, and volumetric/thread models for patterned fabrics. We further

address the problem of material appearance constancy across different rendering platforms. We present two case studies in architectural and interior design. The first study demonstrates Yulio, a new platform for the creation, delivery, and visualization of acquired material models and reverse engineered cloth models in immersive VR experiences. The second study shows an end-to-end process of capture and data-driven BSDF representation using the physically-based Radiance system for lighting simulation and rendering.

CCS CONCEPTS

• **Computing methodologies** → **Reflectance modeling**; *Image and video acquisition; Camera calibration; Computational photography*;

KEYWORDS

BRDF, BSDF, design, remapping, Cloth, Volumetric Models, Reflectance functions, material representation, acquisition, Virtual Reality

ACM Reference format:

Giuseppe Claudio Guarnera, Dar'ya Guarnera, Gregory J. Ward, Mashhuda Glencross, and Ian Hall. 2019. BxDF Material Acquisition, Representation, and Rendering for VR and Design. In *Proceedings of SIGGRAPH Asia '19, Brisbane, Australia, 17-20 November 2019*, 17 pages. <https://doi.org/10.1145/3355047.3362092>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SIGGRAPH Asia '19, 17-20 November 2019, Brisbane, Australia

© 2019 Copyright held by the owner/author(s).

ACM 978-1-4503-6941-1/19/11.

10.1145/3355047.3362092

1 INTRODUCTION

Humans have the capability to identify which material an object is made of, simply by looking at it, even though the appearance of materials may vary significantly depending on a wide range of properties such as colour, gloss, translucency, and texture, that interact with each other in a complex way. Moreover, the geometry of an object, the angle from which the material is viewed and the lighting directions influence Human perception of materials. A major challenge in computer graphics is how to simply and accurately measure the appearance of material characteristics from real-world objects and implement practical editable synthetic materials accurately matching the appearance of the original.

A variety of rendering systems are used in the software pipeline, resulting in a need for optimised material representations. This demands both a flexible acquisition process and representation methods. Unfortunately, up to now there is no widely adopted solution good enough for a wide range of commercial applications without significant labour and money. There is no standardised material model format for material capture and acquired material models have a high memory footprint. Finally, there is little standardisation across renderers, and material models are hard to edit by artists. Moreover, Virtual Reality (VR) is raising the bar for efficient, realistic and cost effective material modelling since it poses additional issues, related to the scale versatility, the need of tiling artifact reduction and the computational limitations of mobile devices.

This course introduces a broad range of material capture and representation methods. Characteristics of mathematical BSDF models as well as more complex, data-driven techniques for physically-based rendering are presented. Material appearance constancy across different rendering platforms is addressed, and practical applications in VR, architecture, and interior design are highlighted. The course is partly based on our previous SIGGRAPH Asia 2016 and SIGGRAPH 2017 Courses. We have revised the material taking into account the feedback from beginners, students, artists and researchers to increase the educational value. Non-experts in the field can easily follow the course and benefit from our comprehensive review. After our introduction to the fundamental concepts, the course gradually moves to most advanced and up to date concepts on both the representation and acquisition side, such as the use of deep learning based techniques to estimate BRDFs/SVBRDFs, as well as data-driven techniques for physically-based rendering.

2 SYLLABUS

The outline of the half-day course is reported below.

- (1) Introduction and Course Overview – Mashhuda Glencross, 5 minutes;
- (2) Materials for Virtual Reality: A Case Study – Ian Hall, 25 minutes
 - (a) Barriers to use in VR.
 - (b) Interior Design VREs.
- (3) Material Appearance representation (Sec 4) – Giuseppe Claudio Guarnera, 45 minutes;
 - (a) Taxonomy of the reflectance functions (BRDF, BTDF, BSDF, BTF, BSSRDF).
 - (b) BRDF parameterization and properties.
 - (c) Physically based models.
 - (d) Appearance models for fabric.
- (4) Material appearance constancy (Sec 5) – Dar'ya Guarnera, 20 minutes;
 - (a) BRDF parameter remapping
- (5) Material capture Methods (Sec 6) – Mashhuda Glencross, 20 minutes;
 - (a) Acquisition Setups and trade-offs.
 - (b) Gonioreflectometers.
 - (c) Image Based acquisition.
 - (d) Catadioptric setups.
 - (e) Linear Light Source setups.
 - (f) LCD Light Source setup.
 - (g) Flash Light setups.
 - (h) Spherical Gantries setups.
- (6) Break, 15 minutes.
- (7) Applying Deep Learning (Sec. 6.7) – Giuseppe Claudio Guarnera, 20 minutes.
 - (a) Deep learning for material appearance estimation.
- (8) General data-driven approach to BSDFs based on detailed measurements (Sec 4.6) – Gregory J. Ward, 70 minutes
 - (a) PAB-opto goniophotometer
 - (b) Interpolating BSDFs using radial basis functions and advection between incident directions
 - (c) BSDF representations appropriate for rendering
 - (d) How to best make use of free, available software and APIs
- (9) Questions, 5 minutes

ORGANIZERS AND PRESENTERS BIOGRAPHIES

Giuseppe Claudio Guarnera is an Associate Professor at NTNU (Norway) and a Lecturer at the University of York (UK). He received his PhD in Computer Science from the University of Catania (Italy). During his PhD, he also worked at the USC Institute for Creative Technologies (US). His research interests include applications of Computer Vision to Computer Graphics and human perception of materials.

Dar'ya Guarnera is a Research Associate at NTNU (Norway). She obtained her PhD in computer science from Loughborough University. Before this, she gained her Postgraduate Certificate in Education and first-class degree in Computer Science from Liverpool Hope University. She has a BSc in Architecture from Odessa State Academy of Civil Engineering and Architecture. Her interests include 3D architectural visualisation, virtual materials, art and sculpture.

Gregory J. Ward has been a computer graphics researcher since 1985, and is the principal author of the Radiance rendering system. His research includes global illumination, reflectance models, high dynamic range image capture and display, general image processing, and human perception. He is the recipient of the 2007 ACM SIGGRAPH Computer Graphics Achievement Award, the 2012 International Building Performance Association Achievement Award, and the 2018 Daylight Award for Daylight Research. He is currently employed by Dolby Laboratories, and consulting for Irystec and the Lawrence Berkeley National Laboratory.

Mashhuda Glencross is a Senior Lecturer at The University of Queensland. Prior to this she was a faculty member at Loughborough University (UK). She obtained her PhD in Computer Science from the University of Manchester. She has also worked at ARM in Cambridge and as a postdoc at the University of Manchester. Her research interests include virtual reality, 3D reconstruction/relighting, material appearance acquisition, novel user interfaces and human visual perception. She has held a number of SIGGRAPH conference committee roles and is currently a Director at Large on the Executive Committee.

Ian Hall is the co-founder and managing director of Pixel Tours Inc, a design consultancy specializing in human factors engineering and development. Together with KiSP Inc, a pioneer in commercial 3D visualization, Ian has been leading the development of Yulio.com, a virtual reality platform for architects and interior designers. After earning his degree in electrical engineering from the University of Waterloo, he served as the product strategist for a series of startups focused on real time control and visualization. In 1999, he became the director of product management for True-spectra Inc., later acquired by Adobe Inc.

3 MATERIAL APPEARANCE: REFLECTANCE FUNCTIONS AND PROPERTIES

The BRDF is a radiometric function, currently used to varying levels of accuracy in photorealistic rendering systems. It describes, in the general case, how incident energy is redirected in all directions across a hemisphere above the surface.

Historically, the BRDF was defined and suggested over the more generalised BSSRDF (Bidirectional Scattering- Surface Reflectance-Distribution function) [Jensen et al. 2001] by Nicodemus [Nicodemus et al. 1977], as a simplified reflectance representation for opaque surfaces: the BRDF assumes that light entering a material leaves the material at the same position, whereas the BSSRDF can describe light transport between any two incident rays on a surface. Many common translucent materials like milk, skin and alabaster cannot be represented by a BRDF since they are characterised by their subsurface scattering behavior that smooths the surface details, with the light shining through them [Goesele et al. 2004]. These materials are expensive to render and many techniques have been proposed [Jensen et al. 2001], [Dachsbacher and Stamminger 2003], [Hao et al. 2003], [Donner et al. 2008], [D'Eon and Irving 2011], [Klehm et al. 2015].

The BRDF is defined as the ratio of the outgoing radiance to the incoming irradiance:

$$f_r(\mathbf{v}_i, \mathbf{v}_r) = \frac{dL_r(\mathbf{v}_r)}{dE_i(\mathbf{v}_i)} = \frac{dL_r(\mathbf{v}_r)}{L_i(\mathbf{v}_i) \cos \theta_i d\omega_i} \quad (1)$$

where $\mathbf{v}_i$ and $\mathbf{v}_r$ are vectors describing the incident (i) and exitant (r) directions, E_i is incident irradiance (*i.e.* the incident flux per unit area of the surface), L_i is incident radiance (*i.e.* the reflected flux per unit area per unit solid angle) and L_r is the reflected radiance; the units of a BRDF are inverse steradian [$1/sr$]. Slide 11 shows the geometry of the BRDF and the vectors used for parameterisation:

- $\mathbf{n}$ is the normal at a specific point p on the surface
- $\mathbf{t}$ is the tangent vector. It is perpendicular to the normal $\mathbf{n}$ and hence it is tangent to the surface at p .

- $\mathbf{b}$ is the bi-tangent vector, defined as $\mathbf{b} = \mathbf{n} \times \mathbf{t}$. In literature it is also named as binormal vector.
- $\mathbf{h}$ is the halfway vector [Rusinkiewicz 1998], defined as: $\mathbf{h} = \frac{(\mathbf{v}_i + \mathbf{v}_r)}{\|\mathbf{v}_i + \mathbf{v}_r\|}$.

The vectors $\mathbf{t}$, $\mathbf{b}$ and $\mathbf{n}$ define a local reference frame; another possibility is to have a local reference frame in which one of the axis is aligned with $\mathbf{h}$ and the other two are given by $\mathbf{b}' = \frac{(\mathbf{n} \times \mathbf{h})}{\|\mathbf{n} \times \mathbf{h}\|}$ and $\mathbf{t}' = \mathbf{b}' \times \mathbf{h}$.

There exist other coordinate system and parameterisation, especially suited for dimensionality reduction of some isotropic BRDF models, for instance the barycentric coordinate system with respect to a triangular support proposed by Stark *et al.* [Stark et al. 2005], or the hybrid model described by Barla *et al.* which could lead to a better repartition of samples to cover most of the effects of materials [Barla et al. 2015]. Edwards *et al.* proposed a framework for transforming the halfway vector $\mathbf{h}$ into different domains to enforce some property of the BRDFs but compromising others.

The BRDF aims to represent the reflectance characteristics of homogeneous materials. An extension of the BRDF concept to non-homogeneous materials is given by the Spatially Varying BRDF (SVBRDF), which can be viewed as a spatial collection of BRDFs distributed over the surface, to simulate the appearance of smooth materials [Haindl and Filip 2013]. The SVBRDF parameterisation includes extra parameters with respect to the BRDF, to take into account the location over the surface: $f_r(x, y, \mathbf{v}_i, \mathbf{v}_r)$. Capturing a SVBRDF generally requires long measurements and processing and can require large, specialised and sometimes expensive hardware rigs to capture the reflectance data. SVBRDF models represent surfaces that are nearly flat and opaque, since the model is restricted by the BRDF reciprocity and energy conservation properties and can not capture subsurface scattering.

Many real world surfaces exhibit local variations not only in reflectance but also in small-scale geometry, causing mesoscopic effects like inter-reflections, self-occlusions and self-masking, not captured by a SVBRDF representation [Haindl and Filip 2013]. Dana *et al.* [Dana et al. 1999] suggested that the BRDF is suitable to characterise surface variation at a coarse scale and introduced the term BTF (Bidirectional Texture Function), an image-based representation that includes small-scale geometry and can describe a fine-scale appearance of a rough surface. The aforementioned mesoscopic effects are difficult to quantify and separate from the measured data, hence BTF acquisition generally needs a large number of samples of the surface as well as high-end hardware support, due to lengthy acquisition times and storage demands [Haindl and Filip 2013]. Nevertheless, there exist low cost acquisition setups, like the kaleidoscopic device by Han and Perlin [Han and Perlin 2003]. BTFs generally result in very realistic material appearance, since BTFs capture spatial variations of the whole surface, unlike BRDFs. Despite of this, many BTF datasets can be approximated as a sparse linear combination of rotated analytical BRDFs [Wu et al. 2011]. The first BTF database, described in [Dana et al. 1999], contains 61 real-world surfaces, each observed under 205 different combinations of lighting and viewing illuminations (plus 205 additional measurements for anisotropic surfaces), consists of over 14,000 images.

For a homogeneous material, an intermediate representation between the BSSRDF and the BRDF, comprising scattering effects for both reflection and transmission, is given by the Bidirectional Scattering Distribution Function (BSDF). The concept of BSDF is generally understood as a sum of a BRDF and a Bidirectional Transmittance Distribution Function (BTDF), the latter modelling how the light passes through a (semi)transparent surface [Walter et al. 2007], [Haindl and Filip 2013].

The taxonomy of the reflectance functions introduced in this Section, showing their relationships with the BRDF along with their parameterisation, is reported in Slide 8: each of them can be considered as a special case of the BRDF [Ashikhmin et al. 2001]. A BRDF should respect some basic physical properties, namely non-negativity, reciprocity and energy conservation:

- non-negativity: the BRDF is a non-negative function, hence for any pair of incident and outgoing direction $f_r(\mathbf{v}_r, \mathbf{v}_i) \geq 0$;
- the Helmholtz reciprocity principle states that the light path is reversible, for any pair of incident and outgoing direction: $f_r(\mathbf{v}_i, \mathbf{v}_r) = f_r(\mathbf{v}_r, \mathbf{v}_i)$. This principle holds only for corresponding states of polarisation for incident and emerging fluxes, whereas large discrepancies might occur for non-corresponding states of polarisation [Clarke and Parry 1985]. In designing a rendering system possible non-reciprocity should be taken into account [Veach 1997].
- Energy conservation assumes that the energy reflected cannot exceed incident energy [Dorsey et al. 2007]: $L_r \leq E_i$ hence over the unit hemisphere Ω_+ above the surface

$$\forall \mathbf{v}_i, \int_{\Omega_+} f_r(\mathbf{v}_i, \mathbf{v}_r)(\mathbf{v}_r \cdot \mathbf{n}) d\omega_r \leq 1 \quad (2)$$

3.1 Lambertian, specular and glossy BRDFs

There are many reflectance models that are simplified subsets of the BRDF function. One of the simplest reflectance models is the Lambertian model, which represents the perfect diffuse reflectance and is often used in many interactive applications, since it requires no recalculation with the change of viewing direction. The model simply assumes that the surface reflects light uniformly in all directions with the same radiance (see Slide 19, in light blue) constant with v_r unlike other BRDF models: $f_r(\mathbf{v}_i, \mathbf{v}_r) = \rho_d/\pi$, where ρ_d is the diffuse albedo.

In the case of a pure specular BRDF all the light is reflected in a single direction, for a given incident direction (see Slide 19, in blue). In fact, light that is incident within a differential solid angle $d\omega_i$ from direction (θ_i, ϕ_i) is reflected in a differential solid angle $d\omega_r$ in direction $(\theta_i, \phi_i + \pi)$, hence the pure specular BRDF can be formalised with a double Dirac delta function: $f_r(\mathbf{v}_i, \mathbf{v}_r) = \rho_s \delta(\theta_i - \theta_r) \delta(\phi_i + \pi - \phi_r)$, where $\rho_s = L_r/L_i$ is the specular albedo. Perfect specularity is valid only for highly polished mirrors and metals.

Surfaces not perfectly smooth, which have some roughness at the micro-geometry level, have a glossy appearance and show broader highlights, other than specular reflections (see Slide 19, in purple).

Some materials, like the surface of the moon or some biological tissues, show a phenomenon called retro-reflection in which light is scattered not only in the forward direction but also in the direction of the illuminant.

BRDFs can be classified by taking into account the characteristics of the material to represent:

- *Isotropic* BRDFs are able to represent materials whose reflection does not depend on the orientation of the surface, since the reflectance properties are invariant to rotations of the surface around $\mathbf{n}$.
- *Anisotropic* BRDFs can describe materials whose reflection change with respect to rotation of the surface around $\mathbf{n}$; this class includes materials like brushed metal, satin, velvet and hair.

The Fresnel effect predicts the fraction of power which is reflected and transmitted and has a great impact on the appearance. Many basic BRDF models have lost importance in the context of physically based modelling because they do not account for a Fresnel term. For conductive materials, like metals, the fraction of light reflected by pure specular reflection is roughly constant for all angles of incidence, whereas for non-conductive materials (dielectrics), the amount of light reflected increases at grazing angles. For a comparative example of the behaviour of metals and dielectrics. The fraction of light reflected is called Fresnel reflectance, which can be obtained from the solution of Maxwell's equations and depends also on the polarisation state of the incident light. For unpolarised light, the Fresnel reflectance $\mathfrak{F}$ at the interface between the surface and the air is given by

$$\mathfrak{F}(\eta, \theta_i, \theta_t) = \frac{1}{2} \left[\left(\frac{\eta \cos \theta_i - \cos \theta_t}{\eta \cos \theta_i + \cos \theta_t} \right)^2 + \left(\frac{\cos \theta_i - \eta \cos \theta_t}{\cos \theta_i + \cos \eta \theta_t} \right)^2 \right], \quad (3)$$

where η is the index of refraction of the surface and θ_t is the angle of transmission. In Computer Graphics, it is very common to use Schlick's approximation of the Fresnel reflectance [Schlick 1994]: $\mathfrak{F}(\theta) = \mathfrak{F}(0) + (1 - \mathfrak{F}(0))(1 - \cos(\theta))^5$; in Section 4 we will generally use the symbol $\mathfrak{F}$ to refer either to the exact Fresnel reflectance or one of its approximations.

4 TAXONOMY OF THE BRDF MODELS

Phenomenological models are entirely based on reflectance data, which is fitted to analytical formulas, thus approximating the reflectance and reproducing characteristics of real world materials. Some of the most important phenomenological models are described in Sections 4.1 and 4.2.

Physically-based models, reported in Sections 4.3 and 4.4, are based on Physics and Optics with the assumption that the surface is rough at a fine scale, therefore described by a collection of micro facets with some distribution D of size and direction. The most common mathematical model has the form:

$$f_r(\mathbf{v}_i, \mathbf{v}_r) = \frac{D \cdot G \cdot \mathfrak{F}}{4 \cos \theta_i \theta_r} \quad (4)$$

which also takes into account the Fresnel term $\mathfrak{F}$. Effect like masking and self-shadowing [Akenine-Möller et al. 2008] depend on the projected area of the microfacets and hence on the distribution D , generally described by the geometrical attenuation term G [Heitz 2014].

A measured BRDF can be fitted to analytic models and employed to reconstruct the BRDF, thus significantly reducing storage size. The down side of this strategy is related to the inflexibility of many

models, hard to edit and able to represent only limited classes of materials. A different solution is to approximate measured BRDFs with a suitable function space, *e.g.* spherical harmonics or wavelets, weighted sum of separable functions or product of functions. We refer to this class of models as Data-Driven models (Sec. 4.6).

Many medium, such as hair, fur, cloth and knitwear are difficult to describe by a surface model. These materials, and objects with highly complex boundary, are better described by volumetric appearance models [Kajiya and Kay 1989; Perlin and Hoffert 1989; Xu et al. 2001], in particular for closer viewing distance, whereas BRDFs can be used from farther away. Appearance models for fabric are reported in Section 4.5.

An important aspect is the practicality of a model in a rendering system, which requires a suitable technique for importance sampling. When calculating the radiance direction of a surface in a scene, accounting for the contribution of light from all possible directions is expensive to compute, therefore Monte Carlo techniques are used to estimate the values with fewer samples [Haines 1991], based on a stochastic process. However, the number of samples should be sufficient to produce consistent estimations, otherwise the results will vary significantly.

A Table which summarizes the characteristics of the reflectance models described in this Course is reported in [Guarnera et al. 2016].

4.1 Phenomenological models for isotropic materials

One of the earliest models for non-Lambertian surfaces is the Phong [Phong 1975] model, based on the cosine law and useful for isotropic materials with a slightly rough surface: $f_r(\mathbf{v}_i, \mathbf{v}_r) = k_s(\mathbf{v}_r \cdot \mathbf{r}_{v_i})^n$, where k_s is a specular constant in the range $[0, \infty]$, $\mathbf{r}_{v_i}$ is the direction of $\mathbf{v}_i$ after being perfectly reflected and n controls the shape of the specular highlight. This model does not take into account energy conservation nor reciprocity. Moreover, it does not capture the reflection behaviour of real surfaces at grazing angles. It is not normalised, however some normalisation factors for cosine lobes have been proposed, either based on double-axis moments [Arvo 1995] or with the simpler option of a power series in $(\mathbf{n} \cdot \mathbf{h})$ with a suitable sequence of exponents [Lewis 1994].

The Blinn-Phong reflection model [Blinn 1977], is derived from the Phong model, but makes use of the halfway vector $\mathbf{h}$ and the normal $\mathbf{n}$ instead of the reflection vector $\mathbf{r}_{v_i}$, thus reducing the computational cost associated with the need to constantly calculate the latter vector:

$$f_r(\mathbf{v}_i, \mathbf{v}_r) = k_s(\mathbf{n} \cdot \mathbf{h})^n \quad (5)$$

Although it has been used as the default shading model for OpenGL and Direct3D until recent times, it shares the same limitations of the Phong model, being physically not plausible and not able to capture metallic and mirrored appearance. Since it follows the cosine function, if n goes to infinity the reflected radiance and the albedo converges to zero towards grazing angles.

Brady *et al.* [Brady et al. 2014] used Genetic Programming to learn new analytic BRDF models. A few basic BRDF models are used as a starting point (seeds), on which symbolic transformations are applied. The fitness function calculates the residual error of each variant after fitting the free parameters to the training set of isotropic materials, a subset of the MERL-MIT database [Matusik

et al. 2003]. The random search is heuristic-based, trying to adapt the starting models to the measured ground truth data and tends to produce a large set of candidates expressions. To allow a better exploration of the search space some suboptimal variations that increase the error are allowed and for the same purpose an island model genetic algorithm is used, allowing only sporadic interactions between sub-populations. The grammar does not guarantee that the resulting models respect energy conservation and reciprocity, hence these properties need to be taken into account by the fitness function; in [Brady et al. 2014] is reported a table with some variants for which the properties have been numerically verified.

4.2 Phenomenological models for anisotropic materials

The Ward reflectance model [Ward 1992] is able to represent both isotropic and anisotropic reflection; it combines specular and diffuse components of reflectance, representing specular peaks through Gaussian distributions. The model specifies an efficient for Monte Carlo sampling. The Ward model has four parameters, which can be set independently, therefore it can be fitted to a large class of measured data. The anisotropic model makes use of the two parameters α_x and α_y to control the width of the gaussian lobe in the two principal directions of anisotropy:

$$f_r(\mathbf{v}_i, \mathbf{v}_r) = \frac{\rho_d}{\pi} + \frac{\rho_s}{\sqrt{\cos(\theta_i) \cos(\theta_r)}} \cdot \frac{e^{-\tan^2(\theta_h) \left(\frac{\cos^2 \theta_h}{\alpha_x^2} + \frac{\sin^2 \theta_h}{\alpha_y^2} \right)}}{4\pi \alpha_x \alpha_y} \quad (6)$$

where ρ_s controls the magnitude of the lobe and $4\pi\alpha^2$ is a normalisation factor. The isotropic Ward model is obtained by setting $\alpha_x = \alpha_y$.

The model does not obey the principle of energy conservation at grazing angles, which has been investigated in [Neumann et al. 1999], [Dür 2006], [Geisler-Moroder and Dür 2010]. A different normalisation factor has been proposed in [Dür 2006] to prevent numerical instabilities and to correct the loss of energy at flat angles, specifically $(4 \cos(\theta_i) \cos(\theta_r))$ instead of $(4 \sqrt{\cos(\theta_i) \cos(\theta_r)})$, however it shares the problem of diverging to infinity with the original Ward model. A new physically plausible version of the model has been proposed in [Geisler-Moroder and Dür 2010], which meets the energy conservation principle even at grazing angles by using the following normalisation factor:

$$\frac{2(1 + \cos \theta_i \cos \theta_r + \sin \theta_i \sin \theta_r \cos \phi_r - \phi_r)}{(\cos \theta_i \cos \theta_r)^4} \quad (7)$$

Neumann *et al.* [Neumann et al. 1999] proposed some modifications and correction factors for the reciprocal [Phong 1975], [Lafortune and Willems 1994], [Blinn 1977] and [Ward 1992] models. The correction factors can be seen as shadowing and masking terms to make the models physically plausible. Moreover the modified models can be used to render metals and other specular objects and for each of them an importance sampling procedure is described.

The Lafortune [Lafortune et al. 1997] model is a flexible, empirical model designed to fit measurements from real surfaces and compactly represent them [Westin et al. 2004]. The model is a generalisation of the cosine lobe model with multiple steerable lobes, based on the Phong model. The primitive functions obey the

Energy Conservation and Reciprocity principles. This model allows lobe specification on the surface in terms of shape and direction, by simply setting up to 3 parameters and an exponent :

$$f_r(\mathbf{v}_i, \mathbf{v}_r) = \frac{\rho_d}{\pi} + \sum_{l=1}^N \left(C_{x,l} v_{ix} v_{rx} + C_{y,l} v_{iy} v_{ry} + C_{z,l} v_{iz} v_{rz} \right)^{n_l} \quad (8)$$

where N is the number of lobes, C_x, C_y, C_z are parameters which absorb the specular albedo and control retro-reflections (by setting C_x, C_y and C_z to positive values), anisotropy (with $C_x \neq C_y$) and off-specular peaks (if C_z is smaller than $-C_x = -C_y$). Lafortune's reflection model can represent generalised diffuse reflectance as the model is able to reflect radiance evenly in all directions, by setting $C_x = C_y = 0$; the Lambertian model can be obtained by setting $N = 0$. A comparative study shows that the Lafortune model performs better than the Phong, Ward and He *et al.* models in representing measured BRDFs like white paper, rough plastic, rough aluminium and metal, since it was designed to fit almost any BRDF data [Westin et al. 2004].

4.3 Physically based models for single-layered material

The Cook-Torrance model [Cook and Torrance 1982] takes into account both specular and diffuse reflections, the latter modeled as Lambertian reflections. As for the specular component, the model assumes that only the fraction of the facets oriented in the direction of $\mathbf{h}$ contributes to the final reflection, moreover it accounts for how many facets are visible from different view angles and how they reflect light [Weyrich et al. 2009]. These factors are modeled respectively through the functions D , G and $\mathfrak{F}$:

$$f_{r,s}(\mathbf{v}_i, \mathbf{v}_r) = \frac{\mathfrak{F}(\theta_r) D(\mathbf{h}) G(\mathbf{v}_i, \mathbf{v}_r)}{\pi \cos(\theta_r) \cos(\theta_i)} \quad (9)$$

The expression of the distribution $D(\mathbf{h})$ is generally a Gaussian: $D(\mathbf{h}) = \cos(\theta_r) \exp(-\frac{\alpha}{m})^2$, where α is the angle between $\mathbf{v}_i$ and the reflected $\mathbf{v}_r$ and m is a roughness parameter. The attenuation term G includes both the shadowing and masking effects:

$$G(\mathbf{v}_i, \mathbf{v}_r) = \min \left(1, \frac{2(\mathbf{n} \cdot \mathbf{h})(\mathbf{n} \cdot \mathbf{v}_r)}{\mathbf{v}_r \cdot \mathbf{h}}, \frac{2(\mathbf{n} \cdot \mathbf{h})(\mathbf{n} \cdot \mathbf{v}_i)}{\mathbf{v}_r \cdot \mathbf{h}} \right) \quad (10)$$

One of the important contributions of this work is the formulation of the Fresnel term $\mathfrak{F}$, which represents the reflection of polished microfacets, approximated with the following expression:

$$\mathfrak{F}(\theta) = \frac{(g - c)^2}{2(g + c)^2} \left(1 + \frac{(c(g + c) - 1)^2}{(c(g - c) + 1)^2} \right) \quad (11)$$

where $c = \mathbf{v}_r \cdot \mathbf{h}$ and $g = \eta^2 + c^2 - 1$, being η the index of refraction. The Cook-Torrance model can properly model metals, plastic with varying roughness and view-dependent changes in colour, although it does not follow the energy conservation principle in the entire hemisphere; additional drawbacks are the not intuitive parameters.

Oren-Nayar [Oren and Nayar 1994] enhanced the Lambertian model for rough diffuse surfaces, to describe in a more realistic way the behaviour of real-world materials like concrete, sand and cloth, which show increasing brightness as the viewing direction approaches the light source direction, rather than being independent of the viewing direction. A rough diffuse surface is modelled as

a collection of long symmetric V-cavities, each of which consists of two microfacets with a Lambertian reflectance; microfacets orientated toward the light source diffusely reflect some light back to the light source (backscatter). The model takes into account masking, shadowing and inter-reflections. The expression is given by:

$$f_r(\mathbf{v}_i, \mathbf{v}_r) = \frac{\rho_d}{\pi} (A + B \max(0, \cos(\phi_i - \phi_r)) \sin(\alpha) \tan(\beta)) \quad (12)$$

where $\alpha = \max(\theta_r, \theta_i)$; $\beta = \min(\theta_r, \theta_i)$; given the surface roughness σ , the expressions for A and B are:

$A = 1 - [(0.5 \cdot \sigma^2)/(\sigma^2 + 0.33)]$; $B = (0.45 \cdot \sigma^2)/(\sigma^2 + 0.09)$. This model, widely used in computer graphics, obeys the reciprocity principle and reduces to the Lambertian model when $\sigma = 0$.

Walter *et al.* [Walter et al. 2007] extend the microfacets theory introduced by [Cook and Torrance 1982] to simulate transmission through etched glass and other rough surfaces, thus taking into account the BSDF. The work by Smith [Smith 1967], which investigated the geometrical self-shadowing of a surface described by Gaussian statistics, is also extended by deriving a shadowing function from any microfacet distribution D ; the BRDF component follows 4.

The distribution D is different from previous models and has been developed to better fit measured data; it is named GGX and has the following expression:

$$D(\mathbf{h}) = \frac{\alpha_g^2 \chi^+(\mathbf{h} \cdot \mathbf{n})}{\pi \cos^4 \theta_h (\alpha_g^2 + \tan^2 \theta_h)^2} \quad (13)$$

where α_g^2 is a width parameter and $\chi^+(x)$ is equal to one if $x > 0$ and zero if $x \leq 0$. The GGX distribution has a stronger tail than commonly used distributions, such as Beckmann and Phong, and thus tends to have more shadowing; in [Bagher et al. 2012] it has been observed that the GGX distribution is identical to the Trowbridge-Reitz distribution [Trowbridge and Reitz 1975]. From D it is possible to derive a simple sampling equation and the expression of G , which is given by:

$$G(\mathbf{v}_i, \mathbf{v}_r, \mathbf{h}) \approx G_1(\mathbf{v}_i, \mathbf{h}) G_1(\mathbf{v}_r, \mathbf{h}) \quad (14)$$

$$G_1(\mathbf{v}_x, \mathbf{h}) = \chi^+ \left(\frac{\mathbf{v}_x \cdot \mathbf{h}}{\mathbf{v}_x \cdot \mathbf{n}} \right) \frac{2}{1 + \sqrt{1 + \alpha_g^2 \tan^2 \theta_x}} \quad (15)$$

The GGX distribution fails to properly capture the glowy highlights of highly polished surfaces like the chrome sample in the MERL database [Matusik et al. 2003], with a narrow specular peak and a much wider specular tail [McAuley et al. 2012]. An anisotropic extension of the distribution, named Generalised-Trowbridge-Reitz, has been proposed by Burley [McAuley et al. 2012]; a symmetric extension of the GGX to the entire ellipsoid domain, suitable for volumetric anisotropic materials, is described by Heitz *et al.* [Heitz et al. 2015].

Bagher *et al.* suggested a function of $\tan^2 \theta_h^{-p}$ for the distribution D , where p depends on the model [Bagher et al. 2012], in order to enhance data fitting for single-layered materials like metals, metallic paints and shiny plastics, otherwise very difficult to fit with commonly used distributions and generally requiring several lobes, due to the shape of the decrease in their BRDFs, close to exponential at large angles but sharper at small angles. The model presented is the Cook-Torrance [Cook and Torrance 1982], in which the microfacets distribution is designed to efficiently and accurately

approximate measured data. The distribution resulting from the suggested slope is called SGD (Shifted Gamma Distribution):

$$D(\theta_h) = \frac{\chi_{[0, \pi/2]}(\theta_h) \alpha^{p-1} e^{-\frac{\alpha^2 + \tan^2 \theta_h}{\alpha}}}{\pi \cos^4 \theta_h \Gamma(1-p, \alpha) (\alpha^2 + \tan^2 \theta_h)^p} \quad (16)$$

where α is a fitting parameter, $\chi_{[0, \pi/2]}(\theta_h)$ is equal to 1 if $\theta_h < \pi/2$ and 0 otherwise, Γ is the incomplete Gamma function:

$$\Gamma(1-p, \alpha) = \int_{\alpha}^{\infty} t^{-p} e^{-t} dt. \quad (17)$$

From the SGD it is possible to derive the shadowing function G and a sampling method.

Low *et al.* [Löw et al. 2012] proposed two isotropic models for glossy surfaces, based either on the Rayleigh-Rice light scattering theory (smooth surface BRDF) or on the microfacet theory (microfacet BRDF). Both models make use of a modified version of the ABC model [Church et al. 1990], [Church and Takacs 1991], which was originally formulated to fit the Power Spectral Density of some measured smooth surfaces. The PSD describes the surface statistics in terms of the spacial frequencies f_x and f_y , which depend on the wavelength λ of the incident light:

$$f_x(\mathbf{v}_i, \mathbf{v}_r) = (\sin \theta_r \cos \phi_r - \sin \theta_i) / \lambda; \quad (18)$$

$$f_y(\mathbf{v}_i, \mathbf{v}_r) = (\sin \theta_r \sin \phi_r) / \lambda. \quad (19)$$

The ABC model [Church et al. 1990], [Church and Takacs 1991] is able to model the inverse power law shape PSD of polished data, and it is given by:

$$PSD(f) = A' / \left(1 + B^2 f^2\right)^{\frac{C+1}{2}} \quad (20)$$

where A is determined by low-frequency spectral density, $B = 2\pi l_0$, l_0 is the autocorrelation length, $C > 0$, $f = \sqrt{f_x^2 + f_y^2}$, $A' = \Gamma((c+1)/2)AB/[2\Gamma(c/2)\sqrt{\pi}]$ and Γ is the gamma function. In [Löw et al. 2012] the ABC model is simplified to $S(f) = a/(1 + bf^2)^c$, where the mapping of the new parameters to the original ABC is: $a = A'$, $b = B^2$ and $c = (C+1)/2$; in practice narrower specular peaks are obtained by increasing b , whereas c controls the fall-off rate of wide-angle scattering. The smooth surface BRDF has the following expression:

$$f_r(\mathbf{v}_i, \mathbf{v}_r) = (k_d/\pi) + O \mathfrak{I}(\theta_d) S(\|\mathbf{D}_p\|) \quad (21)$$

where k_d is a scaling factor for the Lambertian term, O is a modified obliquity factor, $\mathfrak{I}(\theta_d)$ is the Fresnel term in Equation 9 with extinction coefficient set to zero and aimed to approximate the reflectivity polarisation factor, which depends on the surface material properties. $\mathbf{D}_p$ is the projected deviation vector, defined as $\mathbf{D}_p = \mathbf{v}_r - \mathbf{r}_{v_i, p}$, where $\mathbf{r}_{v_i, p}$ is the projection of $\mathbf{v}_r$ on the surface tangent plane and $\mathbf{r}_{v_i, p}$ is the projection of the mirror direction of $\mathbf{v}_i$ on the surface tangent plane. To deal with unreliable data near grazing angles, the value suggested for the obliquity factor O is 1 instead of the typical definition of $O = \cos \theta_i \cos \theta_r$. The microfacet model is based on Cook-Torrance [Cook and Torrance 1982] and makes use of the modified ABC distribution:

$$f_r(\mathbf{v}_i, \mathbf{v}_r) = \frac{k_d}{\pi} + \frac{\mathfrak{I}(\theta_h) S\left(\sqrt{1 - \mathbf{h} \cdot \mathbf{n}}\right) G(\mathbf{v}_i, \mathbf{v}_r)}{\mathbf{v}_i \cdot \mathbf{n} \mathbf{v}_r \cdot \mathbf{n}} \quad (22)$$

where $\mathfrak{I}$ and G are the same as in Equation 9, S is the modified ABC distribution and k_d is again a scaling factor for the diffuse component; the parameter a of S is used as a scaling factor for the specular term, hence the distribution is not normalised. The model is reciprocal but does not obey energy conservation. Both models provide accurate fits to measured data, with the microfacet model showing lower errors, and accurately represent scattering from glossy surfaces with sharp specular peaks and non-Lambertian wide angle scattering. For both models an efficient importance sampling strategy is suggested.

The discrete stochastic model by Jakob *et al.* [Jakob et al. 2014] extends the microfacet theory by replacing the continuous distribution of microfacets in the Cook-Torrance model [Cook and Torrance 1982] with a discrete one, thus assuming that a surface consists of a high but finite number of scattering particles. This assumption facilitates modelling a controllable, non-smooth spatially varying BRDF appearance of a glittery surface, like mica flakes, ice crystals, metallic car paint and craft glitter for decorations. The notion of multiscale BRDF is introduced, which takes into account finite areas and solid angles rather than single points and directions:

$$f_r(A, \mathbf{v}_i, \omega_r) = \frac{(\mathbf{v}_i \cdot \mathbf{h}) \mathfrak{I}(\mathbf{v}_i \cdot \mathbf{h}) D(A, \omega_h) G(\mathbf{v}_i, \mathbf{v}_r, \mathbf{n})}{a(A) \sigma(\omega_r) (\mathbf{v}_i \cdot \mathbf{n}) (\mathbf{v}_r \cdot \mathbf{n})} \quad (23)$$

where A is the area around the point p into account, $a(A)$ its surface area, $\omega_h := \{(\mathbf{v}_i + \mathbf{v}_r / \|\mathbf{v}_i + \mathbf{v}_r\|), \mathbf{v}_r \in \omega_r\}$ is the set of microfacet normals that reflect from $\mathbf{v}_i$ into the finite solid angle ω_r around $\mathbf{v}_r$, $\sigma(\omega_r)$ is the area of ω_r on the unit sphere, $\mathfrak{I}$ is the fresnel term, G models shadowing and masking. The discrete multiscale microfacets distribution D is defined as:

$$D(A, \omega_h) = \frac{1}{N} \sum_{k=1}^N 1_{\omega_h}(\mathbf{v}_h^k) 1_A(p^k) \quad (24)$$

where p^k and $\mathbf{v}_h^k$ are the position and normal of the k^{th} microfacet of a list of N microfacets, 1_A and 1_{ω_h} are the indicator functions of the sets A and ω_h respectively. The indicator functions control the appearance of the surface, since they determine which microfacets in A reflect light into the solid angle ω_r around $\mathbf{v}_r$: a high number of participating facets gives a smoother appearance than a low number, which gives instead a strongly glittery appearance. An efficient implementation of the model is discussed, together with an importance sampling strategy for Monte Carlo renderers.

4.4 Physically based models for multi-layered material

The multilayered model by Ershov *et al.* [Ershov et al. 2001] represents car paint and consists of binder pigment particles, flakes and flake coatings. The model approximates the BRDF of each sub-layer and then merges sub-layers together and it is able to produce realistic appearance for car paints and models their components (binder, pigment particles, flakes). However, due to the complexity of the layered model, the computational time is significantly high. An later model is simplified to bi-layered materials and presents a substrate layer as a solid paint film where the reflectance is Lambertian and a transparent binder layer with embedded flakes. Flakes are considered as partially transparent coloured mirrors, with the assumption

that the reflectance of flakes does not depend on the incident direction and inter-reflections between flakes, so their interaction with light is modelled using constant reflectance values [Ershov et al. 2004]. A large number of parameters are required and not all of them can be directly measured. This model is suitable for interactive design of automotive paints, by solving through optimisation the problem of finding pigment composition of a paint from its bidirectional reflectance distribution function.

To simulate both smooth and rough multi-layered materials, Weidlich and Wilkie [Weidlich and Wilkie 2007] proposed to combine several microfacet based layers into a single physically plausible BRDF model. Their model assumes that any microfacet is large in relation to the layer thickness, models the absorption of part of light when it travels inside a transparent material and include a total reflection term, when light propagates at an angle of incidence greater than the critical angle; the simplicity of the model does not allow reproducing effects like iridescence.

Rump *et al.* [Rump et al. 2008] proposed the first hybrid analytical BRDF and image-based BTF representation to describe the complex reflectance behaviour of car paint. The appearance of metallic car paint is separated into the homogeneous BRDF part, which describes the reflection behaviour of the base and the top layer of the paint, and the spatially varying BTF part, which is caused by the aluminium flakes. The homogeneous part is represented by a multi lobe version of the Cook Torrance model [Cook and Torrance 1982]. In order to account for the characteristics of pearlescent paint, which show view-dependent off-specular colour changes, the model includes a spectral view and light dependent part. The BRDF parameters are derived from the BTF measurements by means of a fitting procedure; the BRDF is calculated for every pixel and subtracted in the RGB space from the captured images. The resulting images contain only flakes data and they are used for a copy and paste synthesis approach.

Kurt *et al.* [Kurt et al. 2010] proposed a physically plausible BRDF model based on the halfway vector representation and Beckmann distribution. The basic BRDF model they propose is the sum of a pure Lambertian term and a single specular lobe, which can be readily extended to multiple specular lobes representation, to model mixture materials like a car paint:

$$f_r(\mathbf{v}_i, \mathbf{v}_r) = \frac{k_d}{\pi} + \sum_{l=1}^N \frac{k_{sl} \mathfrak{F}_l(\mathbf{v}_r, \mathbf{h}) D_l(\mathbf{h})}{4(\mathbf{v}_r \cdot \mathbf{h})((\mathbf{v}_i \cdot \mathbf{n})(\mathbf{v}_r \cdot \mathbf{n}))^{\alpha_l}} \quad (25)$$

where N is the number of lobes, k_d is the diffuse albedo, k_{sl} is the specular reflectivity per-lobe, $\mathfrak{F}_l$ is a per-lobe Fresnel term, D_l a per-lobe normalised microfacet distribution, α_l is a set of parameters which needs to be chosen carefully to enforce energy conservation.

4.5 Appearance Models for Fabric

Realistic representation of textiles is essential in Computer Graphics, since cloth is an important component of many virtual scenes, from interior design to fashion.

A woven textile can be considered at the fiber scale, yarn scale and weave pattern scale. At the micro-scale, the geometric (*e.g.* cross section) and optical properties (*e.g.* index of refraction) of the fibers, which are grouped together to form yarns, influence the

light transport in a scene and hence can have a significant impact on the overall appearance of a fabric swatch.

A huge number of short staple fibers (*e.g.* cotton) can be twisted together to make a single thread, called spun yarn. Strands of yarn (plies) are twisted together to make a thicker yarn and depending on the direction of the twisting the yarn will have either *s*-twist or *z*-twist. Small fibers can protrude from a yarn, causing hairiness. For materials like silk, long and continuous fibers are grouped together to form filament yarns.

Once yarns have been colored or brightened they can be used for weaving, the process of interlacing two perpendicular sets of yarns (warps and wefts) by means of a loom. Warps are fixed horizontally to the loom while the wefts are inserted crossways while some of the warps are lowered or raised; the interlaced yarns hold together by friction. The process of interlacing the yarn is controlled by the weave pattern, a binary matrix which defines the crossings: typically 0 indicates the warp overlaps the weft at the corresponding yarn crossing, whereas 1 means that it lies below. The weave pattern determines many mechanical properties of fabric and also greatly influence its appearance.

To capture the view-dependent appearance of cloth BTFs have been often employed. In [Sattler et al. 2003] the fabric BTF is acquired from a rectangular probe and PCA of the acquired data is used to generate a view-dependent texture-map, mapped onto the simulated geometry. One of the main issues with the use BTF is the difficult handling of silhouette effects, where the quality is limited as well as under grazing illumination.

Distant viewing of cloth has been simulated in [Sadeghi et al. 2013], where a shading model based on light scattering measurements at both weave and yarn level is reported. The surface-based BSDF model takes into account shadowing and masking between neighboring threads.

Jakob *et al.* [Jakob et al. 2010] introduced a generalisation to anisotropic scattering structures, exploited also for volumes acquired by CT scans [Zhao et al. 2011]. More recently, collections of individual fibers have been used for fabric representation [Khurgurn et al. 2015].

Computed Tomography imaging (CT) has been used by Zhao *et al.* [Zhao et al. 2011] to derive yarn and fiber geometry information. From these, they build volumetric models which are fitted to the optical information, *i.e.* color and texture, derived from photographs. The resulting model allows to synthesize large regions of fabric without tiling artifacts, however it can only reproduce textiles that have been previously scanned.

In a later work, Zhao *et al.* [Zhao et al. 2012] used CT scanning to acquire a database of volumetric 3D models of complex woven textiles. Realistic close up renderings are achieved from two types input, the description of the material weave pattern, and a few examples of simpler fabrics. Given a weave pattern in input which specifies the structure of the output, a new sample can be generated from similar or simpler exemplars of a textile, by copying data from the reference exemplars at each yarn crossing to match the requested structure. Their method allows quick creation of new exemplars of fabrics due to a fast synthesis algorithm and is useful for generating volumetric models of complex and spatially varying woven textiles.

In Schröder *et al.* [Schröder et al. 2015] a single image as input is used to derive the weave pattern, yarn colors and local irregularities such as thinning and thickening of yarns, by means of an iterative process. The optical properties of the fibers, assumed to be dielectrics with elliptical cross-section, are described by a bidirectional curve scattering distribution function (BCSDF), where the parameters are manually set. Khungurn *et al.* [Khungurn et al. 2015] proposed a framework to optimize for the parameters of a given appearance model (either volumetric or at fiber level) that best match a set of photographs of a fabric sample under many different lighting conditions. The described framework can handle only fabrics with a single yarn type and color.

In [Aliaga et al. 2017] low-level optical and geometric properties of fibers are used to derive a physically-based scattering model, allowing accurate definition of the reflectance field of cloth fibers.

In the context of the contract furniture business, woven cloth model acquisition must be rapid, since the volume of cloth swatches that need to be routinely processed is huge. Nevertheless, results must be renderable in high quality even on a mobile device, for VR applications. Guarnera *et al.* [Guarnera et al. 2017] developed an image-based technique for reverse engineering woven fabrics at a yarn level. It determines, from a single digital image, captured with a DSLR camera under controlled uniform lighting, the woven cloth structure and reflectance properties.

4.6 Data-driven models

The general idea behind this class of models use the measured data to derive some ad-hoc model, generally based on the principle that a continuous function can be represented by a linear combination of basis functions and a mixture of basis functions can be used for interpolation.

Tensors are a generalisation of scalars and vectors to higher orders and their rank is defined by the number of directions, *e.g.* a scalar is a zero-order tensor and a vector a first-order tensor. Tensor representation has been used for interactive modification of the material properties and relighting by Sun *et al.* [Sun et al. 2007].

Bilgili *et al.* [Bilgili et al. 2011] proposed to represent 4-D measured BRDFs data as a function of tensor products, factorised using Tucker decomposition [Tucker 1966], a generalisation of higher order principal component analysis which decomposes a tensor into a set of matrices and one small core tensor.

More recently, tensor representation has been used by Ward *et al.* [Ward et al. 2012, 2014] to represent anisotropic materials with no assumption on the reflectance and scattering behaviour, particularly useful in presence of unusual scattering properties. The measured data is fitted to a series of radial basis functions in order to derive a continuous representation from the sparse input 4-D measurements. The incident and reflected hemispheres are projected onto disks and mapped over the unit square; the four dimensions given by the two squares define a rank-4 tensor, subdivided into a tensor tree for fast Monte Carlo sample generation. The tensor tree representation adaptively subdivides sharp peaks of the BRDF in different regions of the distribution, with an additional averaging step between incident and reflected direction to account for Helmholtz reciprocity. For more details, please see the included papers [Ward et al. 2012, 2014].

5 BRDF PARAMETER REMAPPING TECHNIQUES

The need to manually match the appearance of a material in two or more different rendering tools is common in digital 3D product design, due to the wide range of tools and material models commonly used, and a lack of standards to exchange materials data. In fact, the appearance of a virtual material depends on the underlying BRDF model implementation; the same material model can be implemented differently in different rendering tools [Sztrajman et al. 2017], and even within the same tool there is no guarantee that different versions render a material identically.

Ngan *et al.* [Ngan et al. 2006] proposed an approach to navigate the appearance space spanned by analytical BRDF models, accounting also for variations within a model; reflectance neighbours in other models are found by using precomputed conversions and multi-linear interpolation. To measure the distance between BRDFs, an image-based L^2 metric is used.

Sztrajman *et al.* [Sztrajman et al. 2017] suggested two image-based strategies for matching the appearance of a BRDF model (called source) to another one (target). The framework is based on a nonlinear optimization performed by means of the Trusted Region Reflective method, which measures the difference between source and target renderings in image space, using a L^2 metric; starting from an user-provided initial guess of the parameters for the target model, the optimization tries to find a set of parameters which minimizes the L^2 metric, generating new rendering with the target model at each step. The most reliable results are obtained by remapping specular and diffuse terms independently (two-stage method). In order to address layered materials, for which the assumption of independence between diffuse and specular terms might not hold, a third stage takes in input the results of the two-stage method, performing an additional optimization in which specular and diffuse are coupled together. However, coupling diffuse and specular seems to causes instabilities [Sztrajman et al. 2017].

Guarnera *et al.* proposed a genetic algorithm-based approach to derive a sensible mapping for parametric BRDF models [Guarnera et al. 2019; Guarnera et al. 2018], without the need for an user-provided initial guess. Their pipeline, aimed at providing a perceptually accurate remapping also for anisotropic BRDFs and SVBRDFs, relies on a metric in image space which exploits features to which the human visual system is very sensitive, such as color and gradient differences. The framework can incorporate knowledge about the semantic and functional relationship among source and target parameters (constrained remapping), allowing a significant reduction of the computational time with respect to the alternative black box-like approach (unconstrained remapping).

6 REFLECTANCE ACQUISITION SETUPS

Measuring how a surface interacts with light can be a time consuming and expensive procedure, potentially generating a vast amount of data.

The setup of a typical measurement device includes a light source to uniformly illuminate a large area of a surface and a detector to measure a small area within the illuminated region [Ashikhmin et al. 2001]. Various systems with different degrees of accuracy and costs have been constructed to measure reflectance functions, ranging

from gonireflectometers to image based measurement systems; low cost setups have also been investigated [Filip et al. 2015; Han and Perlin 2003; Rushmeier et al. 2015]. By dropping the assumption that a material is homogeneous and opaque, many techniques for BRDF measurement can be adapted for more complex reflectance functions (SVBRDFs, BTFs, BSSRDFs).

Under certain assumptions, also setups used to acquire objects geometry through the classical photometric stereo technique [Woodham 1980], where the point of view is kept constant between successive images while the direction of incident illumination varies, have been successfully used to recover BRDF and SVBRDF of non-lambertian surfaces [Alldrin et al. 2008; Chen et al. 2006; Georgiades 2003; Goldman et al. 2005; Hertzmann and Seitz 2005; Holroyd et al. 2008; Zickler et al. 2006].

In general, disentangling the interactions between geometry, reflectance and light, is an ill-posed task, since different combinations of these factors can lead to the same appearance. As a result, most setups need more than one measurement to overcome ambiguities.

Numerical simulation [Ashikmin et al. 2000; Cabral et al. 1987; Dong et al. 2015; Hanrahan and Krueger 1993; Westin et al. 1992] represents, for some complex materials, a possible alternative to a measurement device. The material appearance is described by the result of the simulation of the light interaction with the surface (and sub-surface) structure. Given a geometry that can be ray-traced, Westin *et al.* in their seminal work [Westin et al. 1992] describe a method to simulate scattering hierarchically, by using the result of the simulation at a scale to generate the BRDF for a larger scale.

6.1 Gonireflectometers

The gonireflectometer measures the spectral reflectance of surfaces, it covers specular and diffuse reflectance depending on the settings of the device. The construction of the device was described by Nicodemus and used in the experimental development of a number of BRDF models, including Torrance and Sparrow [Torrance and Sparrow 1967], Blinn [Blinn 1977], He *et al.* [He et al. 1991] and many others. The gonireflectometer design can allow different Degrees of Freedom (DoF), thus allowing different kind of measurements.

Foo [Foo 1997] designed a three axis automated gonireflectometer with two degrees of freedom. The measuring system consists of a light source moving around a sample, a stationary detector and a folding mirror. The system can measure the reflection at high grazing angles (up to 86 degrees) and allows high dynamic range measurements, making it considerably precise. Unfortunately, this setup can only measure isotropic BRDFs; a similar setup is described by Li *et al.* [Li et al. 2006].

Riviere *et al.* [Riviere et al. 2012] used an in-plane multispectral polarised reflectometer. The measurement setup consists of a lighting system with three linearly polarised laser sources; the polarised detection system is based on the Fresnel equation to identify polariser's axes. It allows sampling at zero lighting angles and it is fully calibrated for polarised and multispectral in-plane BRDF measurements. Polarised measurements are used to distinguish the different scattering processes in BRDF directional components. This measurement system is suggested for analysis of physical measurements of the optical surface and for laser-imaging applications.

6.2 Image based measurement

Image-based BRDF measurement makes use of photographs of an object and requires only general-purpose equipment, thus lowering the cost of the process. The data can be measured quickly and completely through a series of photographs taken of a surface. These photographs capture light reflected from various surface orientations. However, to measure the wavelength spectrum of the BRDF requires more time per measurement [Marschner 1998].

Marschner *et al.* [Marschner et al. 1999] presented a rapid, complete and accurate isotropic BRDF measurement setup for a broad range of homogeneous materials, including human skin. It can achieve high resolution and accuracy over a large range of illumination and reflection directions. This setup consists of a hand-held digital camera, equipped with a standard CCD sensor with RGB colour filter array, and an industrial electronic flash light source, which suffice to measure surfaces with simple shapes, *e.g.* spherical and cylindrical which can be defined analytically; for more complex irregular shapes a 3D scanner is required in addition. The camera, characterised in terms of Optoelectronic Conversion Function (OECF) in order to know the radiance reflected to the camera and the irradiance due to the source, moves from near the light source, to measure near retro-reflections, to opposite the light source, in order to measure grazing-angle reflection. Some additional photographs are taken to measure the location and intensity of the light source, the camera pose and the sample pose. About 30 images from different positions are required to cover the three-dimensional BRDF domain. Each pixel in the images is used to derive one sample in the domain of the BRDF, thanks to the estimated relationship between the geometry of the sample and the position of the camera, light source and sample, through bundle adjustment. A typical measurement session takes up to half an hour.

A more recent development by Matusik *et al.* [Matusik et al. 2003], similar to Marschner *et al.* [Marschner et al. 1999], has been used to measure 100 isotropic materials. Matusik's data-driven method is described in Section 4.6.

Ngan *et al.* [Ngan et al. 2005] presented an anisotropic BRDF acquisition setup for flat and flexible samples. To deal with the anisotropy, strips of the material at different orientations obtained from flat samples are wrapped around a cylinder, which can be tilted by means of a precision motor in order to account for the missing degree of freedom with respect to a sphere. A light source rotates around the cylinder while the target is captured by a fixed camera, enabling the capture of the full 4d BRDF. For each light and target position a set of 8 pictures with different exposures is taken, to form an HDR image. The sampling density of the light and the cylinder tilting can be adjusted to increase the resolution of the measured BRDF, whereas the main limitation in the resolution is due to the limited number of material strips which can be wrapped around the cylinder.

The reflectance acquisition setup proposed by Naik *et al.* [Naik et al. 2011] exploits space-time images captured by a time-of-flight camera. Two different setups are described, both based on indirect viewing with 3-bounce scattering and making use of two known Lambertian materials, respectively the source S and the receiver R, while P is the patch to measure. In the first setup, the laser illuminates S, and the camera views R, thus measuring P indirectly.

As for the second configuration, it is based on an around the corner viewing in which P is not directly visible to the camera, whereas S and R are the same surface. The light is multiplexed along different transport paths and some of them might have the same length, hence the light can arrive along multiple paths at the same point at the same time. For this reason the measurements of the material need to be decoded, by solving a sparse underdetermined system; the system is solved by recovering the parameters the Ashikhmin-Premoze model [Ashikhmin and Premoze 2007] (see Section 4.2), using the halfway vector parameterisation. When the multiplexing does not cause ambiguities, in order to measure the parameters of a material it is enough to analyse the streak images to find the specular peak. This setup enables to take many BRDF measurements simultaneously, but it requires an ultra-fast camera; moreover it suffers from a low signal to noise ratio due to the multiple bounces, the size of patches and the maximum sharpness of the reflectance function are limited by the hardware and the range of measurable incoming and outgoing directions is limited by the geometry of the setup.

6.3 Catadioptric measurement setups

Catadioptric optical systems makes use of both reflected and refracted light, in order to reduce aberrations. The resulting imaging setups are generally efficient image based BRDF acquisition devices, usually without any moving parts.

The imaging gonioreflectometer described by Ward [Ward 1992] measures anisotropic surfaces by repeating measurement process under various orientations. It captures the entire hemisphere of reflected and refracted directions at the same time. This device cannot measure sharp specular peaks nor take measurements at high grazing angles.

Dana *et al.*'s [Dana 2001] measuring device makes use of a parabolic mirror that densely covers a relatively small solid angle. Planar translations of the light source are required to cover various incident directions; similarly translations of the sample are necessary in order to scan the surface for spatial variations in reflectance.

The measurement device presented in [Han and Perlin 2003] is based on the principle of the kaleidoscope and consists of a tapered tube whose inner walls are lined with front-surface mirrors. A single camera captures the kaleidoscopic image, in which the subimages represent the same sample seen simultaneously from many different viewpoints. The sample is illuminated by a DLP projector, which shares the optical path with the camera by means of a 45° beam splitter. The properties of the sample are measured through a sequence of pictures with different illumination images, which illuminate the sample from a known range of incoming directions due to the unique sequence of reflections from the kaleidoscopic walls. The advantages of this setting, suitable for BTfFs and BSSRDFs, are the absence of moving parts which enables quick measurements and guarantee perfect registration of the measurements and the low cost; radiometric and geometric calibration need to be performed only once.

Mukaigawa *et al.* [Mukaigawa *et al.* 2007] built two measurement systems for anisotropic BRDFs which use a projector as the light source, placed at the focal point of an ellipsoidal mirror, a camera and a beam splitter, since the camera and the projector cannot be

located at the same position; a slightly different setup makes use of a half mirror. The number of acquired images depends on the sampling of the lighting direction and viewing direction, which needs to be estimated based on the accuracy required. The acquired data are then fitted to the Ward anisotropic reflection model.

Ghosh *et al.* [Ghosh *et al.* 2007], [Ghosh *et al.* 2010b] describe a measurement device that consists of a camera focusing on a zone of reflected directions, a light source with a beam splitter, a mirrored dome and mirrored parabola. The focus of the illumination beam is on the mirrored components that the beam reflects back to its origin. This setup allows BRDF measurement over a continuous region, from 9° to 57° off normal, corresponding to about 51% of the hemisphere, with a specially designed orthonormal zonal basis function illumination, which results in a very rapid BRDF acquisition and in a better signal to noise ratio compared to point-sampling the incident directions [Mukaigawa *et al.* 2007]. The measurements are then projected into a spherical harmonics basis or fitted to an analytical reflection model. The downside is due to the oscillations which affect the zonal basis in the proximity of discontinuities or strong gradients, thus limiting the acquisition of specular materials.

6.4 Spherical and Hemispherical Gantry

Malzbender *et al.* [Malzbender *et al.* 2001] built a hemispherical device with 50 strobe light source, where the camera is placed at the apex. The samples, almost flat, are placed on the floor and illuminated by a single light source at a time. The acquired data are represented by Polynomial Texture Maps (PTM), in which for each fitted texel the coefficients of the following polynomial are fitted to the data and stored as a map:

$$L(u, v; l_u, l_v) = a_0(u, v) l_u^2 + a_1(u, v) l_v^2 + a_2(u, v) l_u l_v + a_3(u, v) l_u + a_4(u, v) l_v + a_5(u, v) \quad (26)$$

where L is the surface luminance at (u, v) , the local coordinates of the texture and (l_u, l_v) are the projection of the normalised light vector at that coordinate. PTMs can lead to good quality renderings, in particular for diffuse samples.

A hemispherical device for anisotropic BRDF measurement was presented by Ben-Ezra *et al.* [Ben-Ezra *et al.* 2008]. The rationale of this setup is that after an accurate radiometric and geometric calibration, LEDs can be used as light sources and as detectors, without needing any moving parts nor cameras and thus allowing fast acquisition times. In their implementation 84 LEDs pointing toward the centre of the hemisphere are used. During the acquisition, each LED is switched on, in turn acting as an emitter, while all others measure the reflected light from the sample. The SNR of the measurements can be increased by multiplexed illumination and the use of different colours for the LEDs allows capture of multispectral data. Since a LED cannot be used at the same time as an emitter and detector this setup cannot be used to measure retro-reflection and offers a lower resolution compared to camera-based setups.

The measurement device presented by Rump *et al.* [Rump *et al.* 2008] consists of a hemispherical gantry with 151 cameras uniformly distributed; the cameras flashes are used as light sources and for each flash all the cameras take a picture of the subject, giving a total of $151 \times 151 = 22,801$ pictures, which can be increased

by taking HDR sequences. The gantry is capable of supporting projectors in order to illuminate the subject with structured light.

Ghosh *et al.* [Ghosh et al. 2009] proposed a setup suitable for roughly specular objects of any shape, based on a LED sphere with 150 controllable lights linearly polarised, with the subject placed at the centre of the sphere. It can be used to estimate spatially varying BRDFs for both isotropic and anisotropic materials, using up to 9 polarised second order spherical gradient illumination patterns. For specular reflections, specular albedo, reflection vector and specular roughness can be directly estimated from the 0^{th} , 1^{st} [Ma et al. 2007] and 2^{nd} order [Ghosh et al. 2009] statistics respectively. In the same work two additional setups are described. The second setup is suitable for flat objects and uses as the light source a LCD monitor, placed very close to the subject, which clearly offers a smaller coverage of incident direction but with a higher resolution than the LED sphere. The third setup makes use of a roughly specular hemisphere which reflects the light emitted by a projector on the subject placed at the centre of the hemisphere, thus allowing a dense sampling; the camera observes the subject from the apex of the hemisphere.

The analysis of the Stokes reflectance field of circularly polarised spherical illumination has been exploited by Ghosh *et al.* [Ghosh et al. 2010a] to estimate the specular and diffuse albedo, index of refraction and specular roughness for isotropic SVBRDFs, assuming known surface orientation. Three different setups are used to demonstrate the technique, similar to the ones described in [Ghosh et al. 2009] but with the light sources covered with right circular polarisers. Four pictures of the subject are required to measure the Stokes field, three of them with differently oriented linear polarisers in front of the camera and one with a circular polariser.

The same framework based on the analysis of the Stokes reflectance field has been further exploited by Guarnera *et al.* [Guarnera et al. 2012] and it is extended to cover also unpolarised illumination, to obtain a per-pixel estimate of the surface normal from the same input data as in [Ghosh et al. 2010a]. The proposed setup makes use of a LED sphere with 346 controllable lights unpolarised/circularly polarised; the surface normals estimation is demonstrated also with uncontrolled outdoors measurement under overcast and hence unpolarised sky, by capturing a reference dielectric sphere in the same environment.

Tunwattanapong *et al.* [Tunwattanapong et al. 2013] proposed a spinning spherical reflectance acquisition apparatus. A 1m semi-circular arc with 105 LED focused toward the centre rotates about the vertical axis at 1rpm, sweeping out continuous spherical harmonic illumination conditions. They demonstrated that 44 pictures are enough to estimate anisotropic SVBRDFs and the 3D geometry of very specular or diffuse objects. This technique further generalises the approach by Ghosh *et al.* [Ghosh et al. 2009], since it can be applied to higher-order spherical harmonic illumination (up to 5^{th} order), which allows obtaining diffuse/specular separation without relying on polarisation.

Gardner *et al.* [Gardner et al. 2003] built a low cost linear light source apparatus to capture flat samples making use of a fixed camera for imaging and a structured light diode. The light source is a 50 cm long neon tube, which is translated horizontally over the surface of the subject and moved in sync with the camera

acquisitions. The reflectance model used to fit the measured data is the isotropic model by Ward [Ward 1992], given the camera and light source positions at each frame. The laser projects a laser stripe, which is deformed by surface variations and used in order to recover the geometry, together with two scans of the light source, in a diagonal direction. A cabin light box, with two diffused cathode tubes are used as a sample holder and to project a even diffuse white light on the surface and allows measurement of the transmitted and reflected light. Overall, the system allows recovery of the diffuse and specular colours, specular roughness, surface normals and per pixel translucency for isotropic samples.

In Ren *et al.* [Ren et al. 2011] a hand-held linear light source device, together with a BRDF chart is employed to obtain spatially varying isotropic BRDFs from a video taken with a mobile phone in LDR. The BRDF chart consists of 24 square flat tiles, with known BRDFs. The tiles are made of specular materials, except one which is a diffuse standard for camera calibration (exposure and white balance). The light source is a 40cm florescent tube, slowly moved by hand over the surface and the chart, which needs to be placed alongside. This approach requires solving a number of issues, since the camera and the light source need to be placed close to the sample and the light is moved manually. Consequently, the camera and light position are unknown, as well as the SVBRDF of the sample. Saturated values from LDR acquisition are repaired using the values in the neighbourhood and the reflectance responses are normalised and hence aligned by a dynamic time warping algorithm. Aligned samples are then used for BRDF reconstruction.

Chen *et al.* [Chen et al. 2014] present a similar setup to Gardner *et al.* [Gardner et al. 2003], scanning a linear light source over a flat sample but with the significant advantage of capturing anisotropic surface reflectance. The basic assumption is that a microfacet model can be used to model the anisotropic surface reflectance. To observe the specular reflection they modulate the illumination along the light source, by means of a transparent mask. They propose two different setups which differ in form factor and employ the same 35cm CCFL lamp and DSLR camera. The desktop form factor scanner scans a linear light source over the sample, observing the SVBRDF by means of the camera; as for the hand-held form factor scanner, the sample moves with respect to the camera and the linear light source, which instead have a fixed relative position. Finally a cylindrical lens is employed to capture in a single picture a scanline of the sample. One constant lighting pattern, together with two phase shifted sinusoidal patterns suffices to reconstruct the surface reflectance.

6.5 LCD Light Source setups

Francken *et al.* [Francken et al. 2008] make use of commodity hardware such as a LCD display and a SLR camera to recover detailed normal maps of specular objects, based on the observation that the normal of a specular pixel is the halfway vector between the light direction and the view direction. To identify the light direction among n different light sources they make use of a gray code lighting patterns, by taking $O(\log_2 n)$ pictures. The accuracy of the estimated normal map depends on the number of sampled light sources.

In Aittala *et al.* [Aittala et al. 2013] a low cost capture setup for SVBRDFs is presented. Their work relies on the design of the image formation model and uses a Fourier basis for the measurements. Isotropic BRDFs are reconstructed through Bayesian inference, since the model is analytically integrable.

The capture set up by Wang *et al.* [Wang et al. 2011], consists of a vision camera and a regular LCD, used as an area light source. It allows rapid measurement of a stationary, isotropic, glossy and bumpy surface, describing its appearance with a dual-level model, which consists of the specular and diffuse relative albedos, two surface roughness parameters and a 1D power spectrum over frequencies for visible surface bumps. Two images are required for calibration, since the LCD radiance is dependent on the viewing angle. To establish the pose of the surface with respect to the camera a target is placed on the surface. At the micro-scale the reflectance is characterized with the Cook-Torrance model and the distribution D is assumed to be Gaussian, where the standard deviation represents the roughness; similarly at the mesoscale level roughness is approximated in terms of the standard deviation. The effect of the roughness at the microscale is assumed to be a blurring of perfect mirror reflections, whereas at the mesoscale it determines a permutation of the pixels. The surface is illuminated with a half-black, half-white image with a vertical edge, and the overall roughness is estimated by fitting a Gaussian filter that blurs the step-edge image to produce the observed one. To separate the roughness for the two different scales, all pixels are sorted by intensity and reshaped back in column-major order, thus removing the permutation induced by the mesoscale roughness; the slope of the segment obtained by averaging over the rows of the sorted image is used to estimate the microscale roughness. This approach can produce visually plausible results for highly glossy man-made indoor surfaces, including some paints, metals and plastics.

Riviere *et al.* [Riviere et al. 2015] propose a mobile reflectometry solution based on a mobile device's LCD panel as extended illumination source, statically mounted at a distance of 45cm above a isotropic planar material sample, at normal incidence, in a dimly lit room. The linear polarisation of the LCD panel is exploited for diffuse/specular separation, by taking two pictures of the sample with a differently orientated plastic sheet linear polariser in front of the device camera. Albedo, surface normals and specular roughness are estimated by illuminating the sample with the same lighting patterns described in [Ghosh et al. 2009]. Due to the limited size of the LCD panel and the position of the front camera, this setup can only acquire $5\text{cm} \times 5\text{cm}$ area of the sample; for larger samples an appearance transfer approach, that relies on additional measurements under natural illumination, is used.

6.6 Flash Illumination and other Capture Setups

Backscattering data can be used to extract an appropriate distribution for microfacets BRDF models [Ashikhmin and Premoze 2007; Dupuy and Jakob 2018]. Based on this observation, mobile devices equipped with a flash light, typically positioned near the back camera, represent near-coaxial setups particularly useful to capture the backscatter surface reflectance to be fitted in a microfacets BRDF model [Riviere et al. 2015], since they allow to extract an

appropriate distribution of microfacets [Ashikhmin and Premoze 2007], by observing high frequencies for specular components in the highlights areas and the diffuse components in other areas [Li et al. 2018].

Riviere *et al.* [Riviere et al. 2015] mobile flash-based acquisition setup estimates the diffuse and specular albedo, surface normal and specular roughness of a planar material sample, with spatially varying isotropic surface reflectance. The back camera and flash light of a mobile device are used for a hand-held acquisition of a video in a dimly lit room, capturing data of the sample from several directions over the upper hemisphere. For reflectance calibration the diffuse grey squares of an X-Rite ColorChecker are used. The top view of the sample at normal incidence is used as a reference to register the other frames. To estimate the lighting and view directions the magnetometer/accelerometer sensors or 3D tracking can be used. The surface normal of each point is computed as the weighted average of the brightest reflection direction, the diffuse albedo is estimated as the trimmed median of the measured intensities, whereas the specular albedo is estimated from the hemispherical integral of the diffuse subtracted measurements. The specular roughness is obtained by fitting the observed backscattering profile to the [Walter et al. 2007] model (see Section 4.3). Some blurring in the reflectance maps can be introduced by misalignments and motion blur. The limited number of lighting directions suggests the use only for rough specular materials.

Aittala *et al.* [Aittala et al. 2015] mobile measurement setup for stationary materials consists of a single mobile device with on-board flash light. Given a flash-no-flash image pair of a textured material of known characteristic size, a multi-stage reconstruction pipeline allows to capture the full anisotropic SVBRDF. The input images are registered through a homography, computed from manually specified points of correspondence. The flash image provides an approximate retro-reflective measurement for each pixel, that combines the effect of surface normal and BRDF, whereas the other image is used as a guide to identify points on the surface with similar local reflectance. Since there is only one observation per pixel, it is assumed that multiple points on the surface share the same reflectance properties and that can be identified under ambient lighting to be combined together. The input is organised into regular tiles approximately of the same size of the repeating texture pattern, assumed to contain a random rearrangement of the same BRDF values. A master-tile is selected for relighting and lumitexels, (*i.e.* data structures to store the geometric and photometric data of one point [Lensch et al. 2003]), are obtained for it. The lumitexels are regularised using a preliminary SVBRDF fit and augmented by transferring high-frequency detail from similarly lit tiles to reduce blurring. The augmented lumitexels are used in a non-linear optimizer to fit an analytic SVBRDF model and the solution is finally reverse-propagated to the full image. This setup limits the input to the retro-reflective slice of the BRDF, hence the Fresnel effect, shadowing and masking are assumed to have typical behaviour and modelled with the BRDF model A [Brady et al. 2014] (see Section 4.1). The camera field of view represents an upper limit on the width of the specular lobes which can be observed.

6.7 Deep learning for material appearance estimation

As seen in the previous sections, the ambiguities arising by the complex interactions between light, reflectance and shape, are typically mitigated by taking more than one measurement of a material. Alternatively, inverse rendering can be solved by introducing a priori assumptions, thus constraining the space of plausible solutions. For example, a typical approach to intrinsic images [Barrow and Tenenbaum 1978] is to assume one or more properties to be known and rather simple, such as a point lights and lambertian reflectance, while estimating the other properties. More complex priors, such as a statistical model of spherical harmonic lighting in natural scenes, help in addressing inverse rendering of arbitrary shapes, reflectances and lighting in the wild [Yu and Smith 2019].

In recent years, automatically learned priors from large datasets have become popular, thanks to the use of Convolutional Neural Networks (CNNs). Rematas *et al.* [Rematas et al. 2016] proposed the use of deep learning to derive a matcap (*i.e.* the image of a sphere in orthographic projection, in which lighting and material properties are conveyed together) from a single RGB image depicting an object with complex shape and specular, homogeneous material. They describe two different variants of their approach, with the first one directly estimating a mapping between the object image and its matcap by means of a CNN. Instead, the second variant starts with CNN-based estimation of a per-pixel surface normal map, used to map the observed RGB values on a sphere, thus producing a sparse reflectance map; finally a dense reflectance map is estimated from the sparse one, by means of a second CNN.

Single-image SVBRDFs reconstruction has been addressed in [Deschaintre et al. 2018], assuming in input a flash photograph, since they contain backscattering data, useful to estimate material properties (see Section 6.6). Since different areas of the flash image convey either information about the specular or the diffuse lobes, to derive a consistent solution for the image it is important to exchange information across distant image region. For this purpose, the typical encoder-decoder architecture is augmented with a secondary network, aimed at extracting global feature at each stage of the main network, the latter focusing on local features. The training stage relies on a similarity metric which compares the appearance of estimated and ground truth material, rendered under several randomized combinations of viewing and incoming directions lighting and viewing directions.

While [Deschaintre et al. 2018] focuses on almost flat geometries, in [Li et al. 2018] SVBRDF and geometry estimation is extended to arbitrary shapes, by means of network with a cascade design. In such a design, each stage estimates shape and SVBRDF parameters, providing the results in input to subsequent stages, along with the error. Input to their system is a single RGB image, acquired under a combination of uncontrolled environment lighting and flash lighting, with the latter assumed to be dominant.

7 CONCLUSION

In this course we introduced the problem of material representation for Computer Graphics, with a particular focus on Virtual Reality applications. We described the state of the art on material models, their representations and acquisition setups. Most models are

limited to a specific material group and even generalised models might not be able cover a broad range of materials nor variations of a material within one group. Currently, material modelling might involve a great deal of manual effort from artists, and the broad range of material models and complexity of the parameters requires from an artist an understanding of the underlying representation and material's micro/ macrostructure.

Due to the wide range of tools and material models commonly used, and the lack of standards to exchange materials data, it is common for a digital artist to manually match the appearance of a material in two or more different rendering tools. Recent BRDF parameter remapping techniques promise to reduce the burden on the artist, automatically providing the user with appearance preserving mapping between a pair of different material models, even across different rendering platforms.

On the material appearance acquisition side, while many setups require samples of a specific size and shape, in recent years light-weight appearance acquisition setups, for instance based on flash photography, and deep-learning based approaches, have become suitable to acquire SVBRDFs from a single input photograph, even in presence of arbitrary shapes.

ACKNOWLEDGMENTS

This work is partly supported by the Research Council of Norway under Grant "MUVApp" N-250293 and Grant "Spectraskin" N-288670.

REFERENCES

- Miika Aittala, Tim Weyrich, and Jaakko Lehtinen. 2013. Practical SVBRDF Capture in the Frequency Domain. *ACM Trans. Graph.* 32, 4, Article 110 (July 2013), 12 pages. <https://doi.acm.org/10.1145/2461912.2461978>
- Miika Aittala, Tim Weyrich, and Jaakko Lehtinen. 2015. Two-shot SVBRDF Capture for Stationary Materials. *ACM Trans. Graph.* 34, 4, Article 110 (July 2015), 13 pages. <https://doi.acm.org/10.1145/2766967>
- Tomas Akenine-Möller, Eric Haines, and Naty Hoffman. 2008. *Real-Time Rendering 3rd Edition*. A. K. Peters, Ltd., Natick, MA, USA, 1045 pages.
- Carlos Aliaga, Carlos Castillo, Diego Gutierrez, Miguel A Otaduy, Jorge Lopez-Moreno, and Adrian Jarabo. 2017. An appearance model for textile fibers. In *Computer Graphics Forum*, Vol. 36. Wiley Online Library, 35–45.
- Neil Alldrin, Todd Zickler, and David Kriegman. 2008. Photometric stereo with non-parametric and spatially-varying reflectance. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*. IEEE, 1–8.
- James Arvo. 1995. Applications of Irradiance Tensors to the Simulation of non-Lambertian Phenomena. In *Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '95)*. ACM, New York, NY, USA, 335–342.
- Michael Ashikhmin and Simon Premoze. 2007. *Distribution-based BRDFs*. Technical Report. 10 pages.
- M. Ashikhmin, P. Shirley, S. Marschner, and J. Stam. 2001. State of the Art in Modeling and Measuring of Surface Reflection. In *ACM SIGGRAPH 2001 Courses*. ACM, 1.
- Michael Ashikhmin, Simon Premoze, and Peter Shirley. 2000. A Microfacet-based BRDF Generator. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '00)*. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 65–74.
- Mahdi M Bagher, Cyril Soler, and Nicolas Holzschuch. 2012. Accurate fitting of measured reflectances using a Shifted Gamma micro-facet distribution. In *Computer Graphics Forum*, Vol. 31. Wiley Online Library, 1509–1518.
- Pascal Barla, Laurent Belcour, and Romain Pacanowski. 2015. In Praise of an Alternative BRDF Parametrization. In *Workshop on Material Appearance Modeling (Proceedings of the Workshop on Material Appearance Modeling 2015)*. Darmstadt, Germany.
- HG Barrow and JM Tenenbaum. 1978. Recovering Intrinsic Scene Characteristics from Images. *Computer Vision Systems* (1978), 3–26.
- M. Ben-Ezra, Jiaping Wang, B. Wilburn, Xiaoyang Li, and Le Ma. 2008. An LED-only BRDF measurement device. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*. 1–8. <https://doi.org/10.1109/CVPR.2008.4587766>

- Ahmet Bilgili, Aydın Öztürk, and Murat Kurt. 2011. A General BRDF Representation Based on Tensor Decomposition. *Computer Graphics Forum* 30, 8 (2011), 2427–2439.
- James F. Blinn. 1977. Models of Light Reflection for Computer Synthesized Pictures. *SIGGRAPH Comput. Graph.* 11, 2 (July 1977), 192–198.
- Adam Brady, Jason Lawrence, Pieter Peers, and Westley Weimer. 2014. genBRDF: Discovering New Analytic BRDFs with Genetic Programming. *ACM Trans. Graph.* 33, 4, Article 114 (July 2014), 11 pages.
- Brian Cabral, Nelson Max, and Rebecca Springmeyer. 1987. Bidirectional Reflection Functions from Surface Bump Maps. *SIGGRAPH Comput. Graph.* 21, 4 (Aug. 1987), 273–281.
- Guojun Chen, Yue Dong, Pieter Peers, Jiawan Zhang, and Xin Tong. 2014. Reflectance Scanning: Estimating Shading Frame and BRDF with Generalized Linear Light Sources. *ACM Trans. Graph.* 33, 4, Article 117 (July 2014), 11 pages. <http://doi.acm.org/10.1145/2601097.2601180>
- Tongbo Chen, Michael Goesele, and Hans-Peter Seidel. 2006. Mesostructure from specular. In *Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on*, Vol. 2. IEEE, 1825–1832.
- Eugene L. Church and Peter Z. Takacs. 1991. Optimal estimation of finish parameters. In *Proc. SPIE*, Vol. 1530. 71–85.
- E. L. Church, P. Z. Takacs, and T. A. Leonard. 1990. The Prediction Of BRDFs from Surface Profile Measurements. In *Proc. SPIE*, Vol. 1165. 136–150.
- F.J.J. Clarke and D.J. Parry. 1985. Helmholtz Reciprocity: its validity and application to reflectometry. *Lighting Research and Technology* 17, 1 (1985), 1–11.
- Robert I. Cook and Kenneth E. Torrance. 1982. A reflectance model for computer graphics. *ACM Transactions on Graphics (TOG)* 1, 1 (1982), 7–24.
- Carsten Dachsbacher and Marc Stamminger. 2003. Translucent Shadow Maps. In *Proceedings of the 14th Eurographics Workshop on Rendering (EGRW '03)*. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 197–201.
- K. J. Dana. 2001. BRDF/BTF measurement device. In *Computer Vision, 2001. ICCV 2001. Proceedings. Eighth IEEE International Conference on*, Vol. 2. 460–466 vol.2. <https://doi.org/10.1109/ICCV.2001.937661>
- Kristin J Dana, Bram Van Ginneken, Shree K Nayar, and Jan J Koenderink. 1999. Reflectance and texture of real-world surfaces. *ACM Transactions on Graphics (TOG)* 18, 1 (1999), 1–34.
- Eugene D'Eon and Geoffrey Irving. 2011. A Quantized-diffusion Model for Rendering Translucent Materials. *ACM Trans. Graph.* 30, 4, Article 56 (July 2011), 14 pages.
- Valentin Deschaintre, Miika Aittala, Fredo Durand, George Drettakis, and Adrien Bousseau. 2018. Single-image SVBRDF Capture with a Rendering-aware Deep Network. *ACM Trans. Graph.* 37, 4, Article 128 (July 2018), 15 pages. <https://doi.org/10.1145/3197517.3201378>
- Zhao Dong, Bruce Walter, Steve Marschner, and Donald P. Greenberg. 2015. Predicting Appearance from Measured Microgeometry of Metal Surfaces. *ACM Trans. Graph.* 35, 1, Article 9 (Dec. 2015), 13 pages.
- Craig Donner, Tim Weyrich, Eugene d'Eon, Ravi Ramamoorthi, and Szymon Rusinkiewicz. 2008. A Layered, Heterogeneous Reflectance Model for Acquiring and Rendering Human Skin. *ACM Trans. Graph.* 27, 5, Article 140 (Dec. 2008), 12 pages.
- J. Dorsey, H. Rushmeier, and F. Sillion. 2007. *Digital Modeling of Material Appearance*. Morgan Kaufmann.
- Jonathan Dupuy and Wenzel Jakob. 2018. An Adaptive Parameterization for Efficient Material Acquisition and Rendering. *ACM Trans. Graph.* 37, 6, Article 274 (Dec. 2018), 14 pages.
- Arne Dür. 2006. An Improved Normalization for the Ward Reflectance Model. *Journal of Graphics, GPU, and Game Tools* 11, 1 (2006), 51–59.
- Sergey Ershov, Roman Durikovic, Konstantin Kolchin, and Karol Myszkowski. 2004. Reverse engineering approach to appearance-based design of metallic and pearlescent paints. *The Visual Computer* 20, 8–9 (2004), 586–600.
- Sergey Ershov, Konstantin Kolchin, and Karol Myszkowski. 2001. Rendering Pearlescent Appearance Based On Paint-Composition Modelling. In *Computer Graphics Forum*, Vol. 20. Wiley Online Library, 227–238.
- Jiří Filip, Michal Havlíček, and Radomír Vávra. 2015. Adaptive highlights stencils for modeling of multi-axial BRDF anisotropy. *The Visual Computer* (2015), 1–11.
- Sing Choong Foo. 1997. *A gonioreflectometer for measuring the bidirectional reflectance of material for use in illumination computation*. Ph.D. Dissertation. Cornell University.
- Yannick Francken, Tom Cuypers, Tom Mertens, Jo Gielis, and Philippe Bekaert. 2008. High quality mesostructure acquisition using specularities. In *CVPR*.
- Andrew Gardner, Chris Tchou, Tim Hawkins, and Paul Debevec. 2003. Linear Light Source Reflectometry. *ACM Trans. Graph.* 22, 3 (July 2003), 749–758. <http://doi.acm.org/10.1145/882262.882342>
- David Geisler-Moroder and Arne Dür. 2010. A New Ward BRDF Model with Bounded Albedo. *Computer Graphics Forum* 29, 4 (2010), 1391–1398.
- Athinodoros S. Georgiades. 2003. Recovering 3-D Shape and Reflectance from a Small Number of Photographs. In *Proceedings of the 14th Eurographics Workshop on Rendering (EGRW '03)*. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 230–240.
- A Ghosh, S. Achutha, W. Heidrich, and M. O'Toole. 2007. BRDF Acquisition with Basis Illumination. In *Computer Vision, 2007. ICCV 2007. IEEE 11th International Conference on*. 1–8. <https://doi.org/10.1109/ICCV.2007.4408935>
- Abhijeet Ghosh, Tongbo Chen, Pieter Peers, Cyrus A. Wilson, and Paul Debevec. 2009. Estimating Specular Roughness and Anisotropy from Second Order Spherical Gradient Illumination. In *Proceedings of the Twentieth Eurographics Conference on Rendering (EGSR '09)*. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 1161–1170.
- Abhijeet Ghosh, Tongbo Chen, Pieter Peers, Cyrus A. Wilson, and Paul Debevec. 2010a. Circularly Polarized Spherical Illumination Reflectometry. *ACM Trans. Graph.* 29, 6, Article 162 (Dec. 2010), 12 pages.
- Abhijeet Ghosh, Wolfgang Heidrich, Shruthi Achutha, and Matthew O'Toole. 2010b. A basis illumination approach to BRDF measurement. *International journal of computer vision* 90, 2 (2010), 183–197.
- Michael Goesele, Hendrik P. A. Lensch, Jochen Lang, Christian Fuchs, and Hans-Peter Seidel. 2004. DISCO: Acquisition of Translucent Objects. *ACM Trans. Graph.* 23, 3 (Aug. 2004), 835–844.
- Dan B. Goldman, Brian Curless, Aaron Hertzmann, and Steven M. Seitz. 2005. Shape and Spatially-Varying BRDFs from Photometric Stereo. In *Proceedings of the Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1 - Volume 01 (ICCV '05)*. IEEE Computer Society, Washington, DC, USA, 341–348.
- D. Guarniera, G.C. Guarniera, A. Ghosh, C. Denk, and M. Glencross. 2016. BRDF Representation and Acquisition. *Computer Graphics Forum* 35, 2 (2016), 625–650. <https://doi.org/10.1111/cgf.12867>
- D. Guarniera, G. C. Guarniera, M. Toscani, M. Glencross, B. Li, J. Y. Hardeberg, and K. Gegenfurtner. 2019. Perceptually Validated Cross-Renderer Analytical BRDF Parameter Remapping. *IEEE Transactions on Visualization and Computer Graphics* (2019), 1–1. <https://doi.org/10.1109/TVCG.2018.2886877>
- Dar'ya Guarniera, Giuseppe Claudio Guarniera, Matteo Toscani, Mashhuda Glencross, Baihua Li, Jon Yngve Hardeberg, and Karl R. Gegenfurtner. 2018. Perceptually Validated Analytical BRDFs Parameters Remapping. In *ACM SIGGRAPH 2018 Talks (SIGGRAPH '18)*. ACM, New York, NY, USA, Article 17, 2 pages. <https://doi.org/10.1145/3214745.3214807>
- Giuseppe Claudio Guarniera, Pieter Peers, Paul Debevec, and Abhijeet Ghosh. 2012. Estimating Surface Normals from Spherical Stokes Reflectance Fields. In *Computer Vision - ECCV 2012. Workshops and Demonstrations*, Andrea Fusiello, Vittorio Murino, and Rita Cucchiara (Eds.). Lecture Notes in Computer Science, Vol. 7584. Springer Berlin Heidelberg, 340–349.
- Giuseppe Claudio Guarniera, Peter Hall, Alain Chesnais, and Mashhuda Glencross. 2017. Woven Fabric Model Creation from a Single Image. *ACM Trans. Graph.* 36, 5, Article 165 (Oct. 2017), 13 pages. <https://doi.org/10.1145/3132187>
- Michal Haindl and Jiří Filip. 2013. Spatially Varying Bidirectional Reflectance Distribution Functions. In *Visual Texture*. Springer London, 119–145.
- Eric A Haines. 1991. Beams O' light: Confessions of a hacker. *SIGGRAPH '91 Course Notes - Frontiers in Rendering* (July 1991).
- Jefferson Y. Han and Ken Perlin. 2003. Measuring Bidirectional Texture Reflectance with a Kaleidoscope. *ACM Trans. Graph.* 22, 3 (July 2003), 741–748.
- Pat Hanrahan and Wolfgang Krueger. 1993. Reflection from Layered Surfaces Due to Subsurface Scattering. In *Proceedings of the 20th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '93)*. ACM, New York, NY, USA, 165–174.
- Xuejun Hao, Thomas Baby, and Amitabh Varshney. 2003. Interactive Subsurface Scattering for Translucent Meshes. In *Proceedings of the 2003 Symposium on Interactive 3D Graphics (I3D '03)*. ACM, New York, NY, USA, 75–82.
- Xiao D. He, Kenneth E. Torrance, François X. Sillion, and Donald P. Greenberg. 1991. A Comprehensive Physical Model for Light Reflection. *SIGGRAPH Comput. Graph.* 25, 4 (July 1991), 175–186.
- Eric Heitz. 2014. Understanding the Masking-Shadowing Function in Microfacet-Based BRDFs. *Journal of Computer Graphics Techniques (JCGT)* 3, 2 (30 June 2014), 48–107.
- Eric Heitz, Jonathan Dupuy, Cyril Crassin, and Carsten Dachsbacher. 2015. The SGGX Microflake Distribution. *ACM Trans. Graph.* 34, 4, Article 48 (July 2015), 11 pages.
- Aaron Hertzmann and Steven M Seitz. 2005. Example-based photometric stereo: Shape reconstruction with general, varying brdfs. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 27, 8 (2005), 1254–1264.
- Michael Holroyd, Jason Lawrence, Greg Humphreys, and Todd Zickler. 2008. A Photometric Approach for Estimating Normals and Tangents. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia 2008)* 27, 5 (2008).
- Wenzel Jakob, Adam Arbree, Jonathan T. Moon, Kavita Bala, and Steve Marschner. 2010. A Radiative Transfer Framework for Rendering Materials with Anisotropic Structure. *ACM Trans. Graph.* 29, 4, Article 53 (July 2010), 13 pages.
- Wenzel Jakob, Miloš Hašan, Ling-Qi Yan, Jason Lawrence, Ravi Ramamoorthi, and Steve Marschner. 2014. Discrete stochastic microfacet models. *ACM Trans. Graph.* 33, 4 (July 2014), 115:1–115:10. <https://doi.org/10.1145/2601097.2601186>
- Henrik Wann Jensen, Stephen R. Marschner, Marc Levoy, and Pat Hanrahan. 2001. A Practical Model for Subsurface Light Transport. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '01)*. ACM, New York, NY, USA, 511–518.
- J. T. Kajiya and T. L. Kay. 1989. Rendering Fur with Three Dimensional Textures. *SIGGRAPH Comput. Graph.* 23, 3 (July 1989), 271–280.

- Pramook Khungurn, Daniel Schroeder, Shuang Zhao, Kavita Bala, and Steve Marschner. 2015. Matching Real Fabrics with Micro-Appearance Models. *ACM Trans. Graph.* 35, 1, Article 1 (Dec. 2015), 26 pages. <https://doi.org/10.1145/2818648>
- Oliver Klemh, Fabrice Rousselle, Marios Papas, Derek Bradley, Christophe Hery, Bernd Bickel, Wojciech Jarosz, and Thabo Beeler. 2015. Recent Advances in Facial Appearance Capture. *Computer Graphics Forum (Proceedings of Eurographics)* 34, 2 (May 2015), 709–733.
- Murat Kurt, László Szirmay-Kalos, and Jaroslav Krivánek. 2010. An Anisotropic BRDF Model for Fitting and Monte Carlo Rendering. *SIGGRAPH Comput. Graph.* 44, 1, Article 3 (Feb. 2010), 15 pages.
- Eric P. Lafortune and Yves D. Willems. 1994. *Using the Modified Phong Reflectance Model for Physically Based Rendering*. Technical Report.
- Eric P. F. Lafortune, Sing-Choong Foo, Kenneth E. Torrance, and Donald P. Greenberg. 1997. Non-linear Approximation of Reflectance Functions. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '97)*. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 117–126.
- Hendrik P. A. Lensch, Jan Kautz, Michael Goesele, Wolfgang Heidrich, and Hans-Peter Seidel. 2003. Image-based Reconstruction of Spatial Appearance and Geometric Detail. *ACM Trans. Graph.* 22, 2 (April 2003), 234–257.
- Robert R. Lewis. 1994. Making Shaders More Physically Plausible. *Computer Graphics Forum* 13, 2 (1994), 109–120.
- Hongsong Li, Sing-Choong Foo, Kenneth E Torrance, and Stephen H Westin. 2006. Automated three-axis gonio-reflectometer for computer graphics applications. *Optical Engineering* 45, 4 (2006), 043605–043605.
- Zhengqin Li, Zexiang Xu, Ravi Ramamoorthi, Kalyan Sunkavalli, and Manmohan Chandraker. 2018. Learning to Reconstruct Shape and Spatially-varying Reflectance from a Single Image. *ACM Trans. Graph.* 37, 6, Article 269 (Dec. 2018), 11 pages. <https://doi.org/10.1145/3272127.3275055>
- J. Löw, J. Kronander, A. Ynnerman, and J. Unger. 2012. BRDF models for accurate and efficient rendering of glossy surfaces. *ACM Transactions on Graphics (TOG)* 31, 1 (2012), 9.
- Wan-Chun Ma, Tim Hawkins, Pieter Peers, Charles-Felix Chabert, Malte Weiss, and Paul Debevec. 2007. Rapid Acquisition of Specular and Diffuse Normal Maps from Polarized Spherical Gradient Illumination. In *Proceedings of the 18th Eurographics Conference on Rendering Techniques (EGSR '07)*. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 183–194.
- Tom Malzbender, Dan Gelb, and Hans Wolters. 2001. Polynomial Texture Maps. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '01)*. ACM, New York, NY, USA, 519–528.
- Steve Marschner. 1998. *Inverse Rendering for Computer Graphics*. Technical Report. Cornell University.
- Stephen R Marschner, Stephen H Westin, Eric PF Lafortune, Kenneth E Torrance, and Donald P Greenberg. 1999. Image-based BRDF measurement including human skin. In *Rendering Techniques 99*. Springer, 131–144.
- Wojciech Matusik, Hanspeter Pfister, Matt Brand, and Leonard McMillan. 2003. A Data-driven Reflectance Model. *ACM Trans. Graph.* 22, 3 (July 2003), 759–769.
- Stephen McAuley, Stephen Hill, Naty Hoffman, Yoshiharu Gotanda, Brian Smits, Brent Burley, and Adam Martinez. 2012. Practical Physically-based Shading in Film and Game Production. In *ACM SIGGRAPH 2012 Courses (SIGGRAPH '12)*. ACM, New York, NY, USA, Article 10, 7 pages.
- Yasuhiro Mukaigawa, Kohei Sumino, and Yasushi Yagi. 2007. Multiplexed illumination for measuring BRDF using an ellipsoidal mirror and a projector. In *Computer Vision-ACCV 2007*. Springer, 246–257.
- Nikhil Naik, Shuang Zhao, Andreas Velten, Ramesh Raskar, and Kavita Bala. 2011. Single View Reflectance Capture Using Multiplexed Scattering and Time-of-flight Imaging. *ACM Trans. Graph.* 30, 6, Article 171 (Dec. 2011), 10 pages. <http://doi.acm.org/10.1145/2070781.2024205>
- László Neumann, Attila Neumann, and László Szirmay-Kalos. 1999. Compact Metallic Reflectance Models. *Computer Graphics Forum* 18, 3 (1999), 161–172.
- Addy Ngan, Frédo Durand, and Wojciech Matusik. 2005. Experimental Analysis of BRDF Models. In *Proceedings of the Sixteenth Eurographics Conference on Rendering Techniques (EGSR '05)*. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 117–126. <https://doi.org/10.2312/EGWR/EGSR05/117-126>
- Addy Ngan, Frédo Durand, and Wojciech Matusik. 2006. Image-driven Navigation of Analytical BRDF Models. In *Proceedings of the 17th Eurographics Conference on Rendering Techniques (EGSR '06)*. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 399–407. <https://doi.org/10.2312/EGWR/EGSR06/399-407>
- FE Nicodemus, JC Richmond, JJ Hsia, IW Ginsberg, and T Limperis. 1977. Geometrical considerations and nomenclature for reflectance, Natl. Bur. Stand. Rep., NBS MN-160 (1977).
- Michael Oren and Shree K Nayar. 1994. Generalization of Lambert's reflectance model. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*. ACM, 239–246.
- K. Perlin and E. M. Hoffert. 1989. Hypertexture. *SIGGRAPH Comput. Graph.* 23, 3 (July 1989), 253–262.
- Bui T Phong. 1975. Illumination for Computer Generated Pictures. *Commun. ACM* 18, 6 (June 1975), 311–317.
- Konstantinos Rematas, Tobias Ritschel, Mario Fritz, Efstratios Gavves, and Tinne Tuytelaars. 2016. Deep Reflectance Maps. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Peiran Ren, Jiaping Wang, John Snyder, Xin Tong, and Baining Guo. 2011. Pocket Reflectometry. *ACM Trans. Graph.* 30, 4, Article 45 (July 2011), 10 pages.
- J. Riviere, P. Peers, and A. Ghosh. 2015. Mobile Surface Reflectometry. *Computer Graphics Forum* (2015), n/a–n/a. <http://dx.doi.org/10.1111/cgf.12719>
- Nicolas Riviere, Romain Ceolato, and Laurent Hespel. 2012. Multispectral polarized BRDF: design of a highly resolved reflectometer and development of a data inversion method. *Opt. Appl* 42 (2012).
- Martin Rump, Gero Müller, Ralf Sarlette, Dirk Koch, and Reinhard Klein. 2008. Photo-realistic Rendering of Metallic Car Paint from Image-Based Measurements. *Computer Graphics Forum* 27, 2 (April 2008), 527–536.
- Holly Rushmeier, Yitzhak Lockerman, Luke Cartwright, and David Pitera. 2015. Experiments with a low-cost system for computer graphics material model acquisition. In *Measuring, Modeling, and Reproducing Material Appearance 2015*, Vol. 9398. International Society for Optics and Photonics, 939806.
- Szymon Rusinkiewicz. 1998. A New Change of Variables for Efficient BRDF Representation. In *Rendering Techniques 98*, G. Drettakis and N. Max (Eds.). Springer-Verlag, Wien, Austria. Proceedings of the Workshop held in Vienna, Austria, between Jyue 29th and July 1st 1998.
- Iman Sadeghi, Oleg Bisker, Joachim De Deken, and Henrik Wann Jensen. 2013. A Practical Microcylinder Appearance Model for Cloth Rendering. *ACM Trans. Graph.* 32, 2, Article 14 (April 2013), 12 pages. <https://doi.org/10.1145/2451236.2451240>
- Mirko Sattler, Ralf Sarlette, and Reinhard Klein. 2003. Efficient and Realistic Visualization of Cloth. In *Proceedings of the 14th Eurographics Workshop on Rendering (EGRW '03)*. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 167–177. <http://dl.acm.org/citation.cfm?id=882404.882429>
- Christophe Schlick. 1994. An Inexpensive BRDF Model for Physically-based Rendering. *Computer Graphics Forum* 13, 3 (1994), 233–246.
- K. Schröder, A. Zinke, and R. Klein. 2015. Image-Based Reverse Engineering and Visual Prototyping of Woven Cloth. *Visualization and Computer Graphics, IEEE Transactions on* 21, 2 (Feb 2015), 188–200. <https://doi.org/10.1109/TVCG.2014.2339831>
- B Smith. 1967. Geometrical shadowing of a random rough surface. *Antennas and Propagation, IEEE Transactions on* 15, 5 (1967), 668–671.
- M.M. Stark, J. Arvo, and B. Smits. 2005. Barycentric parameterizations for isotropic BRDFs. *Visualization and Computer Graphics, IEEE Transactions on* 11, 2 (March 2005), 126–138.
- Xin Sun, Kun Zhou, Yanyun Chen, Stephen Lin, Jiaoying Shi, and Baining Guo. 2007. Interactive Relighting with Dynamic BRDFs. *ACM Trans. Graph.* 26, 3, Article 27 (July 2007).
- Alejandro Strajman, Jaroslav Krivánek, Alexander Wilkie, and Tim Weyrich. 2017. Image-based Remapping of Material Appearance. In *Proc. 5th Workshop on Material Appearance Modeling (MAM '17)*, Reinhard Klein and Holly Rushmeier (Eds.). The Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 5–8. <https://doi.org/10.2312/mam.20171323>
- Kenneth E Torrance and Ephraim M Sparrow. 1967. Theory for off-specular reflection from roughened surfaces. *JOSA* 57, 9 (1967), 1105–1112.
- T. S. Trowbridge and K. P. Reitz. 1975. Average irregularity representation of a rough surface for ray reflection. *J. Opt. Soc. Am.* 65, 5 (May 1975), 531–536.
- Ledyard R Tucker. 1966. Some mathematical notes on three-mode factor analysis. *Psychometrika* 31, 3 (1966), 279–311.
- Borom Tunwattanaapong, Graham Fyfe, Paul Graham, Jay Busch, Xueming Yu, Abhijeet Ghosh, and Paul Debevec. 2013. Acquiring Reflectance and Shape from Continuous Spherical Harmonic Illumination. *ACM Trans. Graph.* 32, 4, Article 109 (July 2013), 12 pages. <http://doi.acm.org/10.1145/2461912.2461944>
- Eric Veach. 1997. *Robust Monte Carlo Methods for Light Transport Simulation*. Ph.D. Dissertation. Stanford University.
- Bruce Walter, Stephen R Marschner, Hongsong Li, and Kenneth E Torrance. 2007. Microfacet models for refraction through rough surfaces. In *Proceedings of the 18th Eurographics conference on Rendering Techniques*. Eurographics Association, 195–206.
- Chun-Po Wang, Noah Snavely, and Steve Marschner. 2011. Estimating Dual-scale Properties of Glossy Surfaces from Step-edge Lighting. In *Proceedings of the 2011 SIGGRAPH Asia Conference (SA '11)*. ACM, New York, NY, USA, Article 172, 12 pages. <http://doi.acm.org/10.1145/2024156.2024206>
- Greg Ward, Murat Kurt, and Nicolas Bonneel. 2012. *Framework for Sharing and Rendering Real-World Bidirectional Scattering Distribution Functions*. Technical Report LBNL-5954E. Lawrence Berkeley National Laboratory. 24 pages.
- Greg Ward, Murat Kurt, and Nicolas Bonneel. 2014. Reducing Anisotropic BSDF Measurement to Common Practice. In *Proceedings of the Eurographics 2014 Workshop on Material Appearance Modeling: Issues and Acquisition (MAM '14)*. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 5–8.
- Gregory J. Ward. 1992. Measuring and Modeling Anisotropic Reflection. *SIGGRAPH Comput. Graph.* 26, 2 (July 1992), 265–272.
- Andrea Weidlich and Alexander Wilkie. 2007. Arbitrarily Layered Micro-facet Surfaces. In *Proceedings of the 5th International Conference on Computer Graphics and Interactive Techniques in Australia and Southeast Asia (GRAPHITE '07)*. ACM, New

- York, NY, USA, 171–178.
- Stephen H. Westin, James R. Arvo, and Kenneth E. Torrance. 1992. Predicting Reflectance Functions from Complex Surfaces. *SIGGRAPH Comput. Graph.* 26, 2 (July 1992), 255–264.
- Stephen H. Westin, Hongsong Li, and Kenneth E. Torrance. 2004. A comparison of four BRDF models. In *Eurographics Symposium on Rendering*. *pages.* 1–10.
- T. Weyrich, J. Lawrence, H. Lensch, S. Rusinkiewicz, and T. Zickler. 2009. Principles of appearance acquisition and representation. *Foundations and Trends® in Computer Graphics and Vision* 4, 2 (2009), 75–191.
- Robert J. Woodham. 1980. Photometric Method For Determining Surface Orientation From Multiple Images. *Optical Engineering* 19, 1 (1980), 191139–191139–.
- Hongzhi Wu, Julie Dorsey, and Holly Rushmeier. 2011. A sparse parametric mixture model for BTF compression, editing and rendering. In *Computer Graphics Forum*, Vol. 30. Wiley Online Library, 465–473.
- Ying-Qing Xu, Yanyun Chen, Stephen Lin, Hua Zhong, Enhua Wu, Baining Guo, and Heung-Yeung Shum. 2001. Photorealistic Rendering of Knitwear Using the Lumislice. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '01)*. ACM, New York, NY, USA, 391–398.
- Ye Yu and William A. P. Smith. 2019. InverseRenderNet: Learning Single Image Inverse Rendering. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Shuang Zhao, Wenzel Jakob, Steve Marschner, and Kavita Bala. 2011. Building Volumetric Appearance Models of Fabric Using Micro CT Imaging. *ACM Trans. Graph.* 30, 4, Article 44 (July 2011), 10 pages.
- Shuang Zhao, Wenzel Jakob, Steve Marschner, and Kavita Bala. 2012. Structure-aware Synthesis for Predictive Woven Fabric Appearance. *ACM Trans. Graph.* 31, 4, Article 75 (July 2012), 10 pages. <https://doi.org/10.1145/2185520.2185571>
- Todd Zickler, Ravi Ramamoorthi, Sebastian Enrique, and Peter N. Belhumeur. 2006. Reflectance Sharing: Predicting Appearance from a Sparse Set of Images of a Known Shape. *IEEE Trans. Pattern Anal. Mach. Intell.* 28, 8 (Aug. 2006), 1287–1302.

Reducing Anisotropic BSDF Measurement to Common Practice

Greg Ward¹, Murat Kurt² and Nicolas Bonneel³

¹Lawrence Berkeley National Laboratory

²International Computer Institute, Ege University

³Harvard University and LIRIS/CNRS

Abstract

We address the problem of measuring and representing reflection and transmission for anisotropic materials without relying on mathematical models or a large sample database. By eliminating assumptions of material behavior, we arrive at a general method that works for any surface class, from metals to fabrics, fritted glazing, and prismatic films. To make data gathering practical, we introduce a robust analysis method that interpolates a sparse set of incident angle measurements to obtain a continuous function over the full 4-D domain. We then convert this interpolant to a standard representation tailored for efficient rendering and supported by a common library that facilitates data sharing. We conclude with some remaining challenges to making anisotropic BSDF measurements truly practical for rendering.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Color, shading, shadowing, and texture

1. The Challenge

The principal difficulty with measured anisotropic Bidirectional Scattering Distribution Functions (BSDFs) is in gathering enough incident and reflected directions to completely characterize a material. The usual approach is to take a small number of accurate measurements or, equivalently, a large number of noisy measurements, and fit them to an appropriate Bidirectional Reflectance Distribution Function (BRDF) or Bidirectional Transmittance Distribution Function (BTDF) model [WMLT07]. This fitting process smooths out noise and reduces the number of dimensions by orders of magnitude, but requires that the data fit a known model a priori. In many applications, this undercuts the purpose of taking measurements, which is to *discover* and *characterize* a material's behavior. Particularly in cases where a material has been custom-designed to have unusual scattering properties, measurements are needed for the very reason that there are no models to fit the data. Data-driven methods have been proposed to solve this problem for dense isotropic BRDF data [MPBM03, PCS*12], but achieving good results with sparse anisotropic data is an open challenge.

Complete anisotropic BSDF measurements are of course possible, but they require days if not weeks of measurement time for a single material in order to gather hundreds of mil-

lions of data points for a simple interpolation scheme. For this reason, most measurements and most research to date have been restricted to isotropic measurements.

A secondary factor curtailing the adoption of measured BSDFs for rendering is the absence of any widely supported representation. There are a number of informal standards that typically list measured values at specific angles on an angular coordinate grid or individual point samples, but there is no obvious way to apply such data. If the data is dense enough, linear interpolation may be used to estimate unmeasured positions, but this threshold is almost never reached for anisotropic distributions. In addition to the impractical measurement times, there is the problem of missing angles due to instrument self-interference. A robust interpolation method is required to arrive at a complete BSDF representation, which must be stored in a standard format supported by a majority of physically-based rendering software packages. Only then will we have a common practice for measured materials.

2. The Vision

Our solution to anisotropic BSDF data measurement is outlined in five parts. First, we describe a commercial goniophotometer that captures dense enough scattering angles to char-




Figure 1: An overview of pgII goniophotometer [AB14].

acterize peaks in the BSDF. Second, we show how to smooth and interpolate incomplete measurements captured by this device. Third, we convert the intermediate representation to a standard format suitable for rendering. Fourth, we describe a BSDF library that supports this format and provides functions for interrogating and generating stratified Monte Carlo samples in a space- and time-efficient manner. Fifth, we suggest a means for sharing measured data between software packages and users, leveraging these techniques.

2.1. Anisotropic BSDF Measurements

We employ the pgII goniophotometer designed and built by Peter Apian-Bennwitz of Freiburg, Germany to measure materials for lighting and energy simulation. This apparatus brings a high degree of automation, programmability, and precision to the characterization of flat material samples [AB14] (See Fig. 1). A sample rotator allows us to measure almost any desired incident and reflected angles up to about 82° from normal, excepting some unavoidable source-detector interference near retroreflection. Also, the total measurement time goes up linearly with the number of incident directions, moving in the direction of "days" as we pass a few dozen angles.

Scattered directions are captured at a high sampling rate during sweeps of the detector arm, which follow longitudinal paths, but may be programmed to capture higher density in areas of interest using additional spiral or sine-wave patterns. This yields many low-noise BSDF measurements along overlapping longitudes with supplementary samples like those shown in Fig. 2(a). The smooth and continuous motion of the detector arm is key to the outstanding accuracy and repeatability of this device. Typically, incident directions are spaced apart 10° or 15° in altitude and 15° or greater in azimuth as shown in Fig. 2(b). This results in between 50 and 100 incident directions for each hemisphere – two hemispheres are captured for materials that transmit as well as reflect light. The number of incident angles required may often be reduced by a factor of two or four if the material's microstructure has bilateral or quadrilateral symmetry, respectively. The number of scattered directions usually runs into tens or hundreds of thousands of sample points per incident direction that do not fit a grid pattern. The only assumption we make is that there are enough scattered di-




Figure 2: (a) Example measured pgII scattering directions for a single incidence. (b) Example set of incident directions for a material with bilateral symmetry. (c) Incident directions from (b) organized into triangle mesh for interpolation.

rections to capture the important peaks in the data, although these peaks may be quite distinct between different incident measurements, as shown in Fig. 3(a,b).

2.2. BSDF Interpolation

One reason we do not see direct renderings of measured materials is that there are invariably holes and noise in the data. There must always be some BSDF model, or another intermediate representation that fills in and smooths out the raw measurements to suit them for rendering. Anisotropic materials are particularly challenging in this regard, and most data-driven methods are isotropic, leaving us only with a handful of mathematical reflectance models, and nothing at all for transmission.

To resolve this problem, we have created a method for interpolating reflectance and transmittance data that makes no assumptions about material behavior and requires no database of measurements other than the particular BSDF of interest. We extend the mass-transport solution of Bonneel et al. [BvdPPH11] to drive a set of radial basis functions fitted to each measured distribution. This allows us to interpolate between sparse incident directions, arriving at a continuous, smooth description of the BSDF over the entire four-dimensional (4-D) domain. Incident direction measurements are organized into a spherical Delaunay mesh like the one shown in Fig. 2(c).

An example of our interpolation technique is shown in Fig. 3, where we compare a naive linear interpolation method to our mass-transport solution. While changing the linear interpolation coordinates to a Rusinkiewicz [Rus98] basis would certainly improve results for isotropic materials, it relies on assumptions about the surface microstructure that do not apply to anisotropic materials in any general way. We could in principle design a custom coordinate system that worked for a particular material, but only if we knew how it behaved, which is presumably why we need measurements in the first place. We opt for the much simpler and more direct solution of employing an interpolation method that fits whatever behavior we measure. Our interpolant is a spherical mesh of radial basis function sets (or systems), one set per incident angle vertex, which describes the scattering distribution. The vertices are interpolated in two dimensions




Figure 3: (a) Reflectance distribution for one incident measurement direction. (b) Distribution at another incidence. (c) Linear interpolation of three distributions. (d) Lagrangian mass transport.

using *transport plans* computed along each edge with a mass transport solver. Each plan is represented as a sparse matrix, and each radial basis system is a list of Gaussian lobe positions, widths, and maximum values. This reduced data is then written to an intermediate file, which is used to generate a BSDF representation for rendering as described below.

2.3. BSDF Representation

While we could render directly from our interpolant, it is more efficient to convert it to a form that is tailored for fast queries and Monte Carlo sample generation. To this end, we have created what we call the BSDF *tensor tree* representation [WKB12]. The incident and reflected hemispheres are projected onto disks, then mapped to Cartesian coordinates over the unit square using the Shirley-Chiu formula [SC97]. These two squares represent the four dimensions of a rank-4 tensor that gets subdivided as a hextree (16 children per node). In this manner, high-resolution peaks anywhere in a distribution may be captured without requiring equivalent high data density everywhere. In the case of anisotropic BRDF measurements, we can add an averaging step between complimentary incident and reflected directions according to Helmholtz reciprocity. This guarantees physical behavior while minimizing unwanted variations in the reconstruction.

For stratified importance sampling, we need to sort our leaf nodes at a given incident direction into a one-dimensional (1-D) sequence that preserves locality. We employ the Hilbert traversal shown in Fig. 4, which maximizes locality while keeping a direct relationship to the output quadtree branching [GL96]. Working in a square slice corresponding to the exiting hemisphere for the queried incident vector, we order a cumulative table along the Hilbert path. This path traverses the $[0, 1]^2$ square with a $[0, 1]$ line segment such that any given fraction of the line segment covers the same fraction of the area. This is the key to converting our two-dimensional (2-D) sampling domain into a 1-D cumulative table based on projected area. Larger leaf nodes will skip further along the path, but we never have the problem of re-entering a node since the Hilbert curve respects the quadtree’s boundaries at each level of detail.

While we have found the tensor tree to be fast and efficient for rendering, other applications may prefer a different representation. For example, the three-phase method




Figure 4: A Hilbert 2-D space-filling curve drawn to the 4th level. Different regions illustrate a hypothetical tensor tree’s subdivision of exiting directions where samples might go.

for annual daylight simulation [MJA*13] works best when a BSDF is stored as a fixed-size matrix corresponding to incoming and outgoing directions using a particular subdivision of the hemisphere. Starting with our mass-transport interpolant, we can compute such a matrix representation.

Fig. 5 shows a rendering comparison between these two representations, demonstrating the superiority of the tensor tree in this case. However, when we want to simulate the appearance of a space as daylight changes and the shades on the window adapt dynamically, a matrix representation reduces the calculation time to a tiny fraction of what it would be otherwise, and permits one to try out different shading solutions and controls at little added cost.

A standard XML file format for the matrix and tensor tree representations has already been defined, and we expect to refine and extend this over time. Supporting this standard is a C library that interprets the XML files in a backwards-compatible fashion, so updating the library is all that is needed to support extensions to the XML representation.

2.4. BSDF Software Library

We have implemented the following queries for matrix and tensor tree sampling in an ANSI-C library with a fixed API:

- Get a BSDF value for a pair of directions.
- Get the directional hemispherical scattering for a given incident direction.
- Get the projected solid angle sample size for one or two directions.
- Get a probability-based importance direction and weight for the given input direction.

The caller may in practice treat the BSDF as a single entity, or may access it as components. These components include transmission, front reflection and back reflection, each further divided into Lambertian and non-Lambertian components, any of which may be zero. The library further supports multiple non-Lambertian components, which may provide for more efficient representations in the future.

In the simplest use of our library, the application makes a single call at each surface evaluation to generate a sample ray direction and weight using method (d). The weight would always be the same, although the spectrum might




Figure 5: (a) Matrix-based BRDF rendering of anisotropic Ward model. (b) Reference image. (c) Tensor tree BRDF rendering of the same model. Insets show color-coded maps which represent a probability that an average observer will notice a difference between the reference image and rendered image [MKRH11].

change for different importance samples. Multiple rays can be generated to reduce variance, and the tensor tree representation maintains any stratification present in the passed random variable. New representations as well as extensions to existing types may be added to our library without altering the interface, which is more general than the existing XML specification. The library also supports loading and caching BSDF data and vector operations for surface reorientation.

While we have tested this API for our own purposes, our goal is to provide it to the rendering and simulation communities as a C library to facilitate BSDF data sharing and re-use based on a standardized XML format.

3. Data Sharing and Remaining Challenges

It is unclear how BSDF data sharing will eventually work, but once we have the capability to characterize materials, a standard means to represent them, and widespread software to make use of these data, manufacturers will most likely want their products included. Similar to the IES and EULUMDAT standards for luminaire data, we expect to see companies offering measured BSDFs they either create themselves or hire an independent laboratory to create for them. We plan to seed this process by measuring a number of materials and providing them with our BSDF library.

In order to create sufficiently compelling examples of this new capability, we have a few issues we must still address. First, we need to find a satisfactory method for extrapolating data to grazing angles. Backlit appearance is important to our judgement of materials, and getting this right is challenging, even for smooth surfaces. The limited ability of most goniophotometers to measure low angles due to edge effects and illumination/acceptance apertures forces us to rely on grazing models to fill in these important regions. Unfortunately, this is also where most mathematical models are weak and do not represent real-world materials accurately. Further research is needed to resolve this problem. Second, there is the question of what to do when there are very narrow peaks in the distribution. While the tensor tree can resolve peaks down to a fraction of a degree, mirror-like re-

flections and glass-like transmission beg for a different representation altogether. Is it possible to represent directional-diffuse with such specular peaks together in a common format? We have not found a good solution to this issue, and we do not know how critical it will be in practice. Third, we need to experiment with different spectral representations to find the best compromise between accuracy and computation requirements. There probably is no "one-size-fits-all" solution, but having a small set of common approaches would be very helpful.

Acknowledgements

This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Building Technologies Program, of the U.S. Department of Energy, under Contract No. DE-AC02-05CH11231 and by the California Energy Commission through its Public Interest Energy Research (PIER) Program on behalf of the citizens of California.

References

- [AB14] APIAN-BENNEWITZ P.: Building material examples (BME) BRDF and BSDF database., 2014. URL: <http://www.pab.eu/gonio-photometer/demodata/bme/>. 2
- [BvdPPH11] BONNEEL N., VAN DE PANNE M., PARIS S., HEIDRICH W.: Displacement interpolation using lagrangian mass transport. *ACM TOG* 30, 6 (2011), 158:1–158:12. (Proc. SIGGRAPH Asia '11). 2
- [GL96] GOTSMAN C., LINDENBAUM M.: On the metric properties of discrete space-filling curves. *IEEE Transactions on Images Processing* 5, 5 (1996), 794–797. 3
- [MJA*13] MCNEIL A., JONSSON C., APPLEFELD E., WARD G., LEE E.: A validation of a ray-tracing tool used to generate bi-directional scattering distribution functions for complex fenestration systems. *Solar Energy* 98, Part C (2013), 404–414. 3
- [MKRH11] MANTIUK R., KIM K. J., REMPEL A. G., HEIDRICH W.: HDR-VDP-2: A calibrated visual metric for visibility and quality predictions in all luminance conditions. *ACM TOG* 30, 4 (2011), 40:1–40:14. (Proc. SIGGRAPH '11). 4
- [MPBM03] MATUSIK W., PFISTER H., BRAND M., McMILLAN L.: A data-driven reflectance model. *ACM TOG* 22, 3 (2003), 759–769. (Proc. SIGGRAPH '03). 1
- [PCS*12] PACANOWSKI R., CELIS O. S., SCHLICK C., GRANIER X., POULIN P., CUYT A.: Rational BRDF. *IEEE TVCG* 18, 11 (2012), 1824–1835. 1
- [Rus98] RUSINKIEWICZ S. M.: A new change of variables for efficient brdf representation. In *Proc. of Eurographics Workshop on Rendering* (1998), Springer, pp. 11–22. 2
- [SC97] SHIRLEY P., CHIU K.: A low distortion map between disk and square. *Journal of Graphics Tools* 2, 3 (1997), 45–52. 3
- [WKB12] WARD G., KURT M., BONNEEL N.: *A Practical Framework for Sharing and Rendering Real-World Bidirectional Scattering Distribution Functions*. Tech. Rep. LBNL-5954E, Lawrence Berkeley National Laboratory, October 2012. 3
- [WMLT07] WALTER B., MARSHNER S. R., LI H., TORRANCE K. E.: Microfacet models for refraction through rough surfaces. In *Proceedings of the 18th Eurographics Conference on Rendering Techniques* (2007), EGSR'07, pp. 195–206. 1