



HAL
open science

A Practical Non-Linear Parameterization of the BRDF Manifold

Cyril Soler, Kartic Subr, Derek Nowrouzezahrai

► **To cite this version:**

Cyril Soler, Kartic Subr, Derek Nowrouzezahrai. A Practical Non-Linear Parameterization of the BRDF Manifold. [Research Report] RR-9069, INRIA. 2017, pp.20. hal-01523874

HAL Id: hal-01523874

<https://inria.hal.science/hal-01523874>

Submitted on 17 May 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



A Practical Non-Linear Parameterization of the BRDF Manifold

Cyril Soler, Kartic Subr, Derek Nowouzezahrai

**RESEARCH
REPORT**

N° 9069

May 2017

Project-Team Maverick



A Practical Non-Linear Parameterization of the BRDF Manifold

Cyril Soler, Kartic Subr, Derek Nowrouzezahrai

Project-Team Maverick

Research Report n° 9069 — May 2017 — 20 pages

Abstract: Real-world reflectance data can be used to improve the realism of synthesized images, albeit with many challenges: memory footprints can be large, profiles are limited to a finite (usually small) set of materials and rendering with measured data can be costly. Since the observation space (number of reflectance measurements) is usually much larger than the underlying space of real-world reflectance profiles, a typical optimisation strategy identifies principal components in the data to directly render from compressed representations of the measurements. We directly learn an underlying low-dimensional non-linear reflectance manifold amenable to rapid exploration and rendering of the space of real-world materials. We show that interpolated materials can be expressed as linear combinations of the measured data, despite lying on a non-linear manifold. This allows us to efficiently interpolate, extrapolate and render directly from the manifold. We apply a Gaussian process latent variable model to represent the reflectance manifold, demonstrating its utility in the context of high-performance and realistic rendering with materials that are interpolations of acquired BRDFs (from the popular MERL dataset [Matusik et al. 2003a]).

Key-words: BRDF interpolation

**RESEARCH CENTRE
GRENOBLE – RHÔNE-ALPES**

Inovallée
655 avenue de l'Europe Montbonnot
38334 Saint Ismier Cedex

Une paramétrisation non linéaire mais versatile du manifold des BRDFs

Résumé : Les données de réflectance issues de mesures goniophotométriques ajoutent un fort niveau de réalisme aux images de synthèse, mais au prix d'une consommation mémoire et d'un coup de calcul conséquents, le tout pour un nombre limité de matériaux possibles. Étant donné que la dimensionnalité de l'espace dans lequel les fonctions de réflectance évoluent est relativement faible par rapport à l'espace d'échantillonnage, une stratégie classique est d'avoir recours à l'analyse en composantes principales afin de pouvoir directement exprimer le rendu dans un espace de dimension faible. Dans ce papier, nous adoptons une approche non linéaire de réduction de dimensionnalité, en utilisant une technique qui préserve la linéarité par rapport aux données d'entrée du problème. Cela nous permet d'exprimer efficacement tout calcul dépendant linéairement des fonctions de réflectance, et d'extrapoler les résultats de ce calcul sur tout le manifold. Nous utilisons pour cela des processus Gaussiens sur la base de MERL [Matusik et al. 2003a], dont nous démontrons l'utilité pour des tâches de rendu temps réel et de simulation de l'éclairage.

Mots-clés : Fonctions de réflectance, Interpolation de matériaux

1 Introduction

Incorporating accurate representations of the reflective characteristics of real-world surfaces is crucial to the photorealism of synthesised images, where the appearance of opaque surfaces is modelled using Bidirectional Reflectance Distribution Functions (BRDFs). Although BRDFs may be specified analytically, many methods capable of *acquiring* BRDFs from real-world materials have been proposed. These methods typically capture raw data by exhaustively tabulating reflectance from many sampled incident and reflected directions. While the dimensionality of the space of BRDFs spanned by these measurements can be arbitrarily large (i.e., four million for the MERL dataset), it is known that the subspace of real-world BRDFs is of comparably low dimensionality [Matusik et al. 2003b]. Directly exploring this subspace, which we call the manifold of BRDFs, opens up many exciting applications: real-world BRDF interpolation for material design, BRDF inference, data completion for partially-observed BRDFs, to name a few.

Linear interpolation of captured BRDF data leads to rendering artifacts, such as “ghosting” when highlights for different BRDFs are not co-incident. Non-linear BRDF interpolation methods [Bonneel et al. 2011, Bonneel et al. 2016] work well, but they do not provide a low dimensional parameterisation of the manifold. This is an important limitation when interpolating materials with high-dimensional measurements. Another drawback of such direct interpolation methods is that their interpolation weights may be directionally-dependent, leading to under-constrained systems that do not respect the physically-based properties of reflection, such as reciprocity.

Several works apply dimensionality reduction techniques to BRDF data. Linear dimensionality reduction (i.e., variants of PCA) results in high-dimensional manifolds (e.g., about 45D for the MERL dataset). Here, the calculation of orthogonal 4D basis functions is challenging due to the aforementioned physical constraints that must be satisfied by real-world BRDFs, such as non-negativity and reciprocity. Standard methods for non-linear dimensionality reduction are not very useful for reasoning about the BRDF manifold: these methods exploit local relations, and so they tend to cope poorly with low sampling rates (i.e., few measured BRDFs) and noisy measurements, two artifacts present in modern BRDF datasets. Charting methods have been used to learn tighter (i.e., 10D) non-linear BRDF manifolds [Matusik et al. 2003b], but these approaches cannot guarantee that arbitrary points on the manifold correspond to valid BRDF approximations. A Euclidean embedding can provide a latent space that is useful for studying relationships between BRDFs, including perceptual distances [Wills et al. 2009], but it does not allow for interpolative exploration of the embedded space.

One approach for resolving these issues is to fit low-dimensional parametric (analytical) models to the acquired reflectance data and to perform interpolation directly in the parametric space [Ngan et al. 2005, Walter et al. 2007, Bagher et al. 2012]. Although identifying such models is non-trivial, their parameterisations naturally allow for interpolation and also guarantee low-dimensional manifolds (as the dimensionality is pre-determined based on the analytical model’s degrees of freedom). Unfortunately, fitting parametric models is often numerically unstable, especially in the presence of the multiple reflectance lobes common to real-world materials. Another important drawback is that these fits provide no smoothness guarantees between the mapping from the parametric space to the measured BRDFs, causing interpolated BRDFs to suffer from abrupt appearance transitions, even when interpolating between two similar BRDFs. Finally, rendering images with BRDFs interpolated in the parametric space still requires costly numerical integration for estimating the rendering (or reflection) equation; typically, these methods propose importance sampling schemes to accelerate the numerical integration, but the integration task remains necessary.

We instead develop an approach, that does not require any run-time numerical integration,

to efficiently render images with materials interpolated from a library of acquired BRDFs. Our method is based on the observation of two important properties of real-world BRDF manifolds, both of which facilitate interpolative exploration: first, we observe that the manifold (or latent space) should be formed using non-linear dimensionality reduction techniques capable of coping with noisy measurements; secondly, even though the interpolation coefficients of different input BRDFs may be non-linearly related, it is desirable that the values of interpolated BRDFs remain linear with respect to the input measurements (see 2.1). We use a Gaussian Process Latent Variable Model (GPLVM) to meet these design properties, allowing us to perform accurate and efficient rendering of interpolated materials, both in the context of direct illumination from distant environmental lighting and approximate global illumination.

2 Related Work

The BRDF is a function that maps pairs of directions to real values: $\rho : \omega_i \times \omega_r \rightarrow \mathfrak{R}$ for each color channel. The value of the BRDF is the ratio of the reflected radiance along ω_r , at a shading point \mathbf{s} in a geometric scene, to the differential radiance due to light incident along ω_i at \mathbf{s} . There is a large body of work pertaining to the study of material appearance over the past four decades, and we point readers to comprehensive surveys on the measurement, modelling, analysis and rendering of materials [Dorsey et al. 2008, Weinmann et al. 2015, Guarnera et al. 2016]. Here, we instead focus on works that are most relevant to our goals, dealing with spatially-invariant materials.

Measurement and modelling of acquired BRDFs

Marschner and colleagues [Marschner et al. 1999] measure BRDF reflectivities at multiple pairs of incident and reflected angles, sampled over the 4D domain, and Matusik et al. similarly publish a dataset [Matusik et al. 2003a] with four million such sampled measurements, for 100 different materials. Due to the amount of data, a common approach for building practical representations has been to combine some form of BRDF parameterisation with numerical approximation, such as tabulation [Steigleder and McCool 2002], matrix decomposition [Kautz and McCool 1999], non-negative matrix factorisation [Lawrence et al. 2004], inverse shade trees [Lawrence et al. 2006a] or Tucker tensor decompositions [Bilgili et al. 2011]. These methods vary in the accuracy-storage trade-offs they make. Recent methods for capturing material properties require impressively few measurements [Georgoulis et al. 2015, Aittala et al. 2016, Nam et al. 2016, Xu et al. 2016], however large datasets created with these methods are not available.

Analysis

BRDF analysis has been approached from roughly two directions: basis function approximations of individual BRDFs, and the study of the entire space of BRDFs. For the former, bases used for analysis include the spherical harmonics (SH) [Westin et al. 1992], spherical wavelets [Schröder and Sweldens 1995] empirical bases using clustering algorithms [Lensch et al. 2003] constrained basis decompositions [Lawrence et al. 2006b] and rotated zonal harmonics [Soler et al. 2015]. When furnished with only partial observations of a single BRDF, Gaussian Process (GP) regression has proven effective for BRDF completion [Hao et al. 2015]. Radiometric studies of the space of BRDFs apply tools for dimensionality reduction directly on the measured data. Linear approaches are unable to identify sufficiently small subspaces [Matusik et al. 2003a] to facilitate practical exploration, whereas many non-linear dimensionality reduction tools (e.g., MDS, ISoMap, LLE) yield compact embeddings without explicitly providing mappings between the measured space and the manifold.

An alternative approach fits parametric models [Ngan et al. 2005, Ashikhmin and Premoze 2007, Bagher et al. 2012, Löw et al. 2012] to the captured data in order to model the variation across different measured BRDFs. Here, the fitting process can be numerically unstable, especially for materials with multiple reflectance lobes.

Perceptual space of BRDFs

Several works aim to understand the perceptual properties of BRDFs, often driven by user studies. This has led to reparameterisations of a specific BRDF with respect to perceptual metrics, as well as for identifying semantically meaningful axes of variation for the BRDF manifold (e.g., color and gloss) [Pellacini et al. 2000]. Further work led to the development of correspondences across these two parametric and perceptual spaces [Westlund and Meyer 2001]. Wills et al [Wills et al. 2009] show that linear interpolation in BRDF space does not result in a linear blend of materials in the perceptual space. They used MDS to obtain an embedding of BRDFs and rendered images by traversing their embedding. Since MDS does not provide a mapping between the two spaces, their method is unable to interpolate materials within the embedded space. A recent technique identified an intuitive control space for materials [Serrano et al. 2016], allowing for impressive exploration of the manifold of acquired materials from a perceptually semantic perspective (see Section 5).

Rendering

Many methods can directly render acquired materials, spanning accurate (but slow) physically-based methods to coarser (but faster) approximations. For real-time rendering, the compression and representation of the acquired BRDFs necessarily remain tightly coupled to the direct rendering algorithm materials. While view-light factorisation using SVD [Kautz and McCool 1999] offers a simple rendering algorithm for a specific BRDF, the use of trilinear tensor factorisation [Sun et al. 2007] improves compression thereby allowing exploration of the space of BRDFs. The projection of BRDFs onto the SH basis leverages the simplicity of convolution in this space for rendering, and this has been exploited for rendering isotropic [Sloan et al. 2002] and arbitrary 4D BRDFs [Kautz et al. 2002]. Most recently, a framework for accelerating spherical filtering with isotropic spherical decompositions (ISD) [Soler et al. 2015] has led to a real-time frequency-domain BRDF rendering solution. One of our applications builds atop the ISD to demonstrate how our representation enables efficient real-time rendering of interpolated materials. Some of the above methods can handle visibility (shadows) but with restrictions, such as static views, static geometry or the use of heavy precomputation. Sun et al [Sun et al. 2007] handles global interreflections using precomputed transfer tensors, and a large body of work on precomputed rendering [Ramamoorthi 2009] demonstrate methods that trade speed for accuracy. Xu et al [Xu et al. 2014] render 1-bounce interreflections, however it is unclear how their spherical-Gaussian representation can be used to render measured BRDFs.

2.1 A Overview of Gaussian Processes

A Gaussian process (GP) is a collection of random variables, any finite number of which stem from a joint Gaussian distribution. If the random values represent the evaluations of some function $f : X \rightarrow \mathfrak{R}$, their associated GP implicitly models distributions over the space of functions. Here, we overview how GPs can be used to perform regression (interpolation) and to optimise low-dimensional latent variables. We limit our review of GPs to the extent that is necessary for understanding our problem, and we refer interested readers to a comprehensive reference on this topic [Rasmussen and Williams 2006].

Regression


Our goal is to predict the value y_* at arbitrary locations $\mathbf{x}_* \in X$ given pairs of observed values (\mathbf{x}_i, y_i) at training locations $\mathbf{x}_i \in X$, where $i = 0, 1, \dots, N-1$. By definition, if we denote the vector of values $\mathbf{y} = [y_0, y_1, \dots, y_{N-1}]^\top$ of the Gaussian process, then it follows that $\mathbf{y} \sim \mathcal{N}(\mu_y, \mathbf{K})$ where μ_y and \mathbf{K} are the mean and covariance of the Gaussian. The elements of the covariance matrix are $\mathbf{K}_{ij} = c(\mathbf{x}_i, \mathbf{x}_j)$ where $c : X \times X \rightarrow \mathbb{R}$ is a covariance function of the users specification. Here, c can be thought of as a *kernel*, and it is key to modeling the non-linearity of the underlying function. Due to the consistency (or marginalisation) property of the GPs, "slicing" a GP along any subset of coordinates results in a 1D Gaussian distribution. So, $y_* \sim \mathcal{N}(\mu_{y_*}, \sigma_{y_*}^2)$ with mean and variance that can be shown to satisfy [Rasmussen and Williams 2006]:

$$\mu_{y_*} = \mathbf{k}_*^\top \mathbf{K}^{-1} \mathbf{y}, \quad (1)$$

$$\sigma_{y_*}^2 = c(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}_*^\top \mathbf{K}^{-1} \mathbf{k}_* \quad \text{and} \quad (2)$$

$$\mathbf{k}_* = [c(\mathbf{x}_0, \mathbf{x}_*), c(\mathbf{x}_1, \mathbf{x}_*), \dots, c(\mathbf{x}_{N-1}, \mathbf{x}_*)]^\top. \quad (3)$$

Interpolating observed values is equivalent to determining a μ_{y_*} which requires: evaluating \mathbf{k}_* , the input covariance function between each training (observed) and test location, computing the inverse of \mathbf{K} (an $N \times N$ matrix), computing a matrix-vector product $\mathbf{K}^{-1} \mathbf{y}$, and computing an inner product of two vectors. This method models non-linearities by virtue of the non-linearity of the covariance function. However, the prediction **remains linear** in \mathbf{y} (Eq. 1). The uncertainty (variance) in the interpolation is given by $\sigma_{y_*}^2$. The above may be extended from the case of a single output variable y to a d -dimensional output, where the function being learned is $f : X \rightarrow \mathbb{R}^d$, by simply replacing the observation vector \mathbf{y} by an observation matrix \mathbf{Z} in Eq. 1, where **each matrix column** is independently extrapolated using Eq. 1 (see Figure below, step 3).



Latent Variable Model (LVM)

In some situations, $\mathbf{z}_i \in \mathbb{R}^d$ are observed but the corresponding \mathbf{x}_i are unknown: e.g., each \mathbf{z}_i could have $d = 4$ million measurements of a single acquired BRDF. One possible solution would be to associate arbitrary valuations of $\mathbf{x}_i \in \mathbb{R}^q$ to the corresponding \mathbf{z}_i and to perform interpolation at some $\mathbf{x}_* \in \mathbb{R}^q$. This is suboptimal as the interpolation results depend heavily on \mathbf{K} (the covariance function evaluated at all-pairs of the chosen \mathbf{x}_i). If $d \gg q$ and the mapping is non-linear, then unsupervised learning of the mapping from observed-to-latent variables corresponds to the classical non-linear dimensionality reduction problem. Alternatively, the GPLVM [Lawrence 2005] approach places a GP prior on the mapping (choosing a covariance function) before optimising the latent variables \mathbf{x}_i of this mapping. The benefits of learning non-linear mappings relies on making a suitable choice for the covariance function, which we will discuss, and a key property that we exploit in our work is that the interpolations remain linear with respect to the observations.

Linearity of interpolation using GPLVM

Let $\mathbf{Z}_{N \times d}$ be the matrix of N observations stacked so that the i^{th} row is \mathbf{z}_i^T and the j^{th} column is a vector composed of the j^{th} dimensional components of all N observations. The output of GPLVM is N optimised q -dimensional latent variables \mathbf{x}_i . Then, the problem of traversing the manifold (latent space) is identical to regression. Given some traversal location \mathbf{x}_* , the goal is to predict the corresponding extrapolated observation \mathbf{z}_* (akin to Eq. 1)

$$\mathbf{z}_*^T = \mathbf{b}_{x_*}^T \mathbf{Z} \quad (4)$$

where $\mathbf{b}_{x_*}^T = \mathbf{k}_*^T \mathbf{K}^{-1}$. Although \mathbf{b}_{x_*} is non-linear with respect to the latent variables \mathbf{x}_i , the extrapolated data is still linear with respect to the observed data.

Properties

We choose the latent space generated by GPLVM because it offers the following key properties:

1. linearity with respect to observations;
2. guaranteed interpolation of observed data regardless of the choice of latent variables \mathbf{x}_i . Replacing \mathbf{x}_* by one of the \mathbf{x}_i turns \mathbf{k}_*^T into line i of \mathbf{K} , which through Eq.4 leads $\mathbf{z}_* = \mathbf{z}_i$;
3. guaranteed continuity in the interpolated observations as long as the covariance function is continuous;
4. knowledge of uncertainty in prediction which may be used as a measure of confidence in the interpolant. Prediction indeed is the mean of a Gaussian random variable which variance can naturally be interpreted as a confidence value.

2.2 Contributions

Our work identifies and leverages key connections in the areas of dimensionality reduction, material appearance and local light transport, to make the following contributions:

1. we show that a non-linear manifold of acquired BRDFs can be directly traversed to produce interpolated BRDFs that are linear combinations of the acquired BRDFs;

2. we apply well-known numerical tool (Gaussian Processes) to a new problem, computing an optimised mapping from observed data to the non-linear manifold of BRDFs, all while preserving important properties of physically-based real-world BRDFs;
3. we show that the uncertainty in GP prediction can be used to drive “sensible” exploration of the underlying manifold; and,
4. we present two new, efficient rendering algorithms that leverage our BRDF interpolation technique – one for real-time rendering of direct illumination and another that includes transport along multi-bounce paths using pre-rendered images.

3 Rendering with interpolated materials

In this section we describe the three stages of our algorithm.

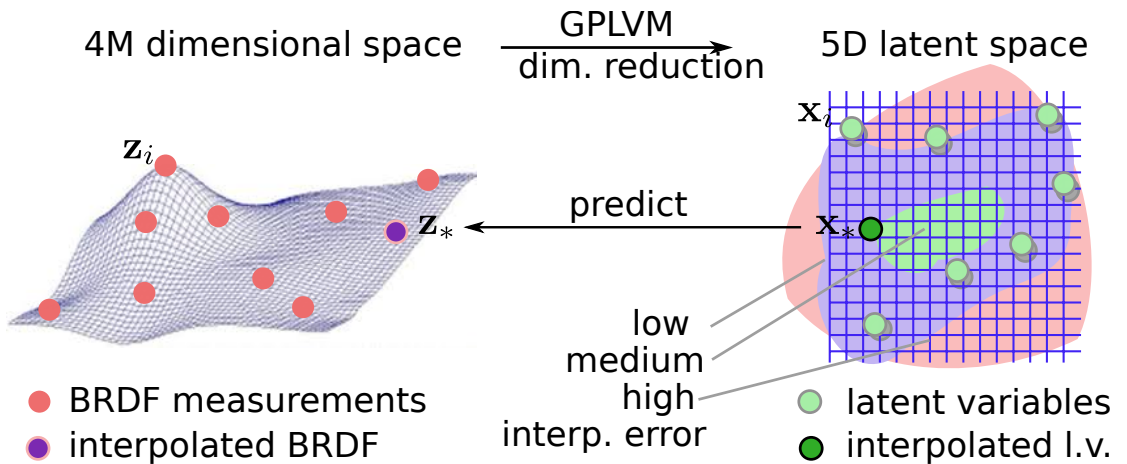


Figure 1: We use GPLVM to identify a low-dimensional non-linear manifold on which latent variables of the measured BRDF values lie. We interpolate the latent variables and map the interpolated vector to the data space to obtain interpolated BRDF \mathbf{z}_* which is a linear combination of the observations \mathbf{z}_i .

3.1 Learning the BRDF manifold

We learn the manifold of acquired BRDFs using GPLVM (see sec. 2.1). The MERL dataset [Matusik et al. 2003b] contains $N = 100$ materials, each with $d = 4M$ (four million) measurements. Each measurement records a scalar measurement of the reflectance for a specific pair of incident and reflected directions. Thus, the size of our observation matrix \mathbf{Z} is $100 \times 4M$. Although latent variables may be chosen arbitrarily, we perform an optimisation to calculate them. This results in a manifold where the spacing between latent variables is in accordance to the L_2 distance in the data. The output of this step is a matrix \mathbf{X} of size $100 \times q$ whose rows are the latent variables \mathbf{x}_i^\top . Best results are obtained with $q = 5$ for the full MERL database (although $q = 2$, used in our video

for clarity gives excellent results), whereas smaller and more consistent sets of similar materials can be very well approximated with $q = 2$.

Choice of covariance function

We use the shifted squared-exponential function (the most widely used kernel in the GP literature), which is

$$c(\mathbf{x}, \mathbf{x}') = \mu\delta(\mathbf{x}, \mathbf{x}') + e^{-\|\mathbf{x}-\mathbf{x}'\|^2/2\ell^2}, \quad (5)$$

where ℓ and μ are hyperparameters. ℓ is often referred to as the *characteristic length scale* since the mean number of level-zero upcrossings for a 1D GP with this covariance function is $(2\pi\ell)^{-1}$. A high value for ℓ leads to a smoother function. μ is a noise-filtering parameter. A very small value (10^{-5}) improves numerical stability a lot (in inverting \mathbf{K}) at the cost of an invisible discontinuity in the interpolant. We choose this covariance function because of its smoothness (it has mean square derivatives of all orders), which will translate into smooth transitions across observed BRDFs (see sec. 5 for detailed discussion).

Optimisation

We obtain optimised latent variables \mathbf{x}_* , by maximising the log-likelihood of the GP for a fixed choice of ℓ and μ :

$$L = -\frac{d}{2} \log |\mathbf{K}| - \frac{1}{2} \text{tr} (\mathbf{K}^{-1} \mathbf{Z}\mathbf{Z}^T) \quad (6)$$

We perform this optimisation using direct local search [Hooke and Jeeves 1961] which offers a very efficient calculation scheme in our case since it only requires evaluating the cost function above while changing a single variable \mathbf{x}_i at once. We therefore maintain both the inverse and the determinant while changing a single line and column of \mathbf{K} using twice the Sherman-Morrisson [Press et al. 2007] and matrix determinant [Harville 1997] formulas. Figure 2 shows the evolution of the log-likelihood when fitting the full MERL database in dimension 2. The corresponding map is displayed in Figure 3. More complicated optimisation methods exist, such as scaled conjugate gradients (the gradients of the log-likelihood are calculated using the chain rule), where the latent variables can be jointly optimised with the hyperparameter. We chose the derivate-free method instead because gradient computations here scale cubically with N . We initialise the latent variables using truncated linear PCA. See sec. 5 for a discussion of these choices.

3.2 Interpolating materials

Given the latent variables \mathbf{x}_i from the previous step and a new location \mathbf{x}_* , we calculate the interpolated BRDF (observed) exactly as in eq. 4. The main questions are then how \mathbf{x}_* can be chosen and what the properties of the interpolated BRDF are.

Choice of \mathbf{x}_*

The choice of \mathbf{x}_* depends on the application for which the BRDF manifold needs to be traversed. If the goal is interactive exploration of the space of acquired BRDFs then even the reduced q -dimensional latent space (e.g. $q = 5$) is uneasy. In our examples of interactive exploration, we display 2D slices of the latent space, the corresponding projections of \mathbf{x}_i and the uncertainty of predictions across this slice. The user then manually selects a point within this subspace as \mathbf{x}_* . We also allow the user to explore successive interpolations along a 1D trajectory in the latent




Figure 2: Convergence of the log-likelihood when fitting the full MERL database using a latent space of dimension 2

space by specifying the endpoints of the trajectory. We then generate the path between the endpoints using an optimisation over the maximum of the uncertainty (variance of prediction) along the trajectory and the length of this trajectory (see sec. 4).

Properties of the interpolated data

Given the value of \mathbf{x}_* , we calculate the interpolated brdf data as $\mathbf{z}_*^T = \mathbf{b}_{x_*}^T \mathbf{Z}$ (eq. 4). Here \mathbf{b}_{x_*} contains non-linearities as a function of \mathbf{X} and \mathbf{x}_* , but the interpolated data is linear in the measurement matrix \mathbf{Z} . Thus, any properties defined using linear operators of the measured data are retained. This is a key property for BRDFs since it guarantees that the interpolated BRDF (1) obeys Helmholtz reciprocity; (2) implicitly interpolates albedo and (3) applies to reflectivity measurements along a fixed direction of incidence. We exploit these properties to develop fast rendering algorithms using the interpolated BRDFs.

3.3 Rendering

The output of the previous step is \mathbf{z}_* , an interpolated BRDF whose values are densities corresponding to 4D points in the same order that they were listed in the library of acquired BRDF measurements. Here we describe how images may be rendered using our interpolated materials \mathbf{z}_* in different scenarios. The central equation of interest for this is the reflectance integral, which describes the radiance arriving from a point \mathbf{r} in space along a direction ω_o towards the centre of projection through pixel \mathbf{p} :

$$I(\mathbf{p}, \omega_o) = \int_{S^2} L(\mathbf{r}, \omega) \rho(\mathbf{r}, \omega_o, \omega) v(\mathbf{r}, \omega) \max(0, \omega \cdot \mathbf{n}) d\omega, \quad (7)$$

where $L(\mathbf{r}, \omega)$ is the incident radiance at \mathbf{r} along ω , ρ is the BRDF at \mathbf{r} , $v(\mathbf{r}, \omega)$ is the visibility of the source of L at \mathbf{r} along direction ω and \mathbf{n} is the normal at \mathbf{r} . We refer to the situation

when L describes radiance that is directly arriving from a light source to \mathbf{r} as *direct reflectance*. For the general case, where light undergoes multiple bounces before the radiance is incident at some point \mathbf{p} on the image plane, L refers to radiance that is arriving at \mathbf{r} from multi-bounce lightpaths.

Direct reflectance + fixed view

Thanks to the linearity of Equation 7 with respect to the reflectance ρ , rendering a pixel with a BRDF at \mathbf{r} that is a linear combination of N measured BRDFs, $\rho_* = \sum_{i=1}^N \mathbf{b}_{x_*}^i \rho_i$, can be expressed as

$$I_*(\mathbf{p}, \omega) = \sum_{i=1}^N \mathbf{b}_{x_*}^i I_i(\mathbf{p}, \omega) \quad (8)$$

where I_i is the image rendered with material ρ_i . The image rendered using the interpolated material is therefore a linear combination of the images rendered with each material. For some applications, such as material design, pre-rendered images may be used to explore the interpolated appearances on the BRDF manifold *without recalculating the reflectance integral* for a fixed view. The images pre-rendered for the different materials may directly be linearly interpolated using elements of the vector \mathbf{b}_{x_*} as coefficients.

Direct reflectance + dynamic view/geometry/lighting

Any algorithm that expresses the BRDF linearly in terms of a basis may be extended to accommodate our interpolated BRDF with minimal implementational changes. We demonstrate this using the example of a recent algorithm [Soler et al. 2015] which expresses the BRDF as a sum of rotated zonal harmonics (ZH) – special spherical harmonics (SH) that are invariant to rotations through a particular fixed axis. Their work exploits the property that a statically chosen set of $(L + 1)^2$ ZH along $2L + 1$ fixed axes \mathbf{a}_m , where L is the degree, together form a basis that exactly spans the space of SH while allowing to compute the shading equation in real time for large values of L (typically up to $L = 40$ in our video). For directional (distant) lighting, where $L(\cdot, \omega) = E(\omega)$ (temporarily ignoring the visibility term for simplicity), they derived the reflectance equation

$$I(\mathbf{p}, \omega) = \sum_{l=0}^L \sum_{m=-l}^l (E \otimes \mathcal{Y}_l^0)(\mathbf{R}_n^{-1} \mathbf{a}_m) \lambda_l^m(\mathbf{R}_n^{-1} \omega_o). \quad (9)$$

\mathbf{R}_n is a rotation that maps global into local directional coordinates so that the up direction is aligned with the shading normal \mathbf{n} , $E \otimes \mathcal{Y}_l^0$ denotes spherical convolution of the illumination and zonal harmonic \mathcal{Y}_l^0 , and λ_l^m are coefficients of the BRDF projected onto rotated ZH. Because λ_l^m linearly depends on the BRDF, there exists a constant matrix $\mathbf{P}_\mathbf{a}$ (that depends only on directions $\{\mathbf{a}_m\}$), so that the vector Λ_i^\top of the $(L + 1)^2$ zonal harmonic coefficients associated with reflectance ρ_i is

$$\Lambda_i^\top = \mathbf{z}_i^\top \mathbf{P}_\mathbf{a}.$$

The interpolated ZH coefficients corresponding to ρ_* are consequently


$$\Lambda_*^\top = \mathbf{z}_*^\top \mathbf{P}_\mathbf{a} = \mathbf{b}_{x_*}^\top \mathbf{Z} \mathbf{P}_\mathbf{a} = \mathbf{b}_{x_*}^\top \Lambda \quad (10)$$

where the matrix Λ is formed by stacking the Λ_i^\top as its lines. So Λ_* can be computed without the need for explicitly determining $\mathbf{P}_\mathbf{a}$. Λ has a size $N \times (L + 1)^2$ where N is the number of measured BRDFs in the input library and L is the chosen SH degree, and is used as \mathbf{Z} in Eq.4.

Given that Λ_* is the set of ZH coefficients for the interpolated material, we use the shader of Soler et al. [Soler et al. 2015] without any implementational changes by simply providing it with Λ_* for real-time rendering of the interpolated material. Due to this simplicity, our interpolation can be used with either variant of their real-time shader: static geometry with the visibility term or dynamic geometry but without visibility.

Global illumination

Let \mathbf{r} be the point of last bounce to the eye, located anywhere in a scene which contains an object whose material ρ_i we wish to modify. We separate the paths of light arriving at \mathbf{r} into two classes $L(\mathbf{r}, \omega) = L_{1-}(\mathbf{r}, \omega) + L_{2+}(\mathbf{r}, \omega)$ based on whether the paths bounce at most once (L_{1-}) on ρ_i or twice or more (L_{2+}) as depicted in the figure below:



Since L_{1-} contains paths with at most one interaction with ρ_i , its contribution to $L(\mathbf{r}, \omega_o)$ is linear (affine, to be accurate) in ρ . L_{2+} includes radiance along all other paths. Substituting this in eq. 7 results in a separation of the image I_i , where the material to be modified is ρ_i , into $I_i = I_i^{1-} + I_i^{2+}$ where I_i^{1-} is an image that is entirely affine in ρ_i and I_i^{2+} contains the remainder of the energy. Due to this linearity, by construction,

$$I_i^{1-}(\mathbf{p}, \omega) = T \mathbf{z}_i \quad (11)$$

where T corresponds to a non-conventional form of the transport matrix. Rather than expressing the radiance at the image plane through linear transport from the light source, eq. 11 represents the image as a linear combination of the measured 4D reflectance data for light bouncing at most once on ρ_i . T includes information about the geometry and lighting in the scene. Note that this is different from direct reflection because T includes multibounce paths to the exception of paths that contain more than 1 reflection off the surface with the material ρ_i . The image I_*^{1-} where ρ_i is replaced with ρ_* is then obtained by interpolating pre-rendered images linearly (as in eq. 8):

$$I_*^{1-}(\mathbf{p}, \omega) = \sum_{i=1}^N \mathbf{b}_{x*}^i I_i^{1-}(\mathbf{p}, \omega). \quad (12)$$

In practice, we expect $L_{2+} \ll L_{1-}$ in general, since the measure of multiple-bounce paths for which more than one of the bounces on ρ is expected to be small, and the energy along these paths is expected to be small as well (compared to the total energy arriving at \mathbf{r} along ω). We observed that applying this interpolation to calculate I_* directly rather than I_*^0 produces very plausible results.

4 Results

Real time exploration of the BRDF manifold

Figure 3 illustrates exploration of the BRDF manifold with images rendered at 25 fps, a primary application for which would be material design. Three different \mathbf{x}_* (red points) are chosen by clicking and dragging the mouse within the 2D latent space (image on far right). All 101 materials from the MERL data set were used to optimise the latent variables. The corresponding rendered images and BRDF slices (below) are shown. The only precomputation necessary is the optimisation of the latent variables associated with each material from the acquired data (which takes less than a minute on a modern CPU). For a fixed set of measured BRDFs, our method allows exploration of the BRDF space while rendering *all combinations of dynamic geometry, view points and lighting at real-time*.



Figure 3: Real-time exploration of our BRDF manifold. 3 materials (red points) are chosen manually, in the vicinity of `blue-acrylic`, by clicking in the 2D latent space (top right). The latent variables were optimised using all 101 materials of the MERL database. Our proposed construction of the BRDF manifold lends itself to **real-time rendering (25 fps)** using interpolated materials along with **all combinations of dynamic geometry, view points and lighting** (see sec. 3.3) using zonal harmonics up to $L = 40$ [Soler et al. 2015]. The slices of the interpolated BRDFs are also visualised. Please see the accompanying video for a live demonstration.

Interactive manipulation of materials in images with global illumination

Figure 4 visualises screenshots from a live session where the user navigates in the latent BRDF space. The images produced by our method using interpolated materials (shown alongside) are obtained by blending images that were pre-rendered using training data (measured materials) with global illumination. Interpolation coefficients are computed in real time and applied to the precomputed images, resulting in blended images that are good approximations of images obtained by solving for global illumination with the interpolated material (see figure 6).

Best path interpolation of BRDFs

The variance of the posterior probability ($\sigma_{z_*}^2$ at a given \mathbf{z}_* similar to $\sigma_{y_*}^2$ in eq. 1) of GPs corresponds to error or uncertainty in prediction. We exploit knowledge of this to explore the manifold. Given two materials, we find the shortest path between them, where cost is defined using a combination of the length of the path and the maximum variance along that path. The computation of optimal paths (motivated in sec. 3.2) in 5D (and greater) is costly. We find




Figure 4: Snapshots captured during a live material editing session where the material on the teapot was continuously interpolated between MERL **chrome** and **silver-paint**. Slices of the corresponding BRDFs are displayed as insets. **Only the first (top left) and last (bottom right) images were pre-rendered** with full global illumination using Mitsuba [Jakob 2010]. **Other images were computed in real-time** using linear interpolation of pre-rendered images using our technique to compute blending weights. Note the consistent change in the reflection of the teapot on the table. The bigger dot corresponds to the material used for the error test in Figure 6. Please see the accompanying video for more. **Note to reviewers:** this figure, and Fig.6 show white squares on MacOS, and need Acrobat to display correctly.

optimised paths to explore the BRDF manifold using Dijkstra’s algorithm for shortest paths, on a coarse discretisation of the latent space. As an illustration of this (see fig. 4), we visualise the shortest path computed in a 5D latent space computed using 10 metals (small white dots) from the MERL database. Even for this selective subset of materials, the non-linearity of the space is obvious (see other 2D slices of the space shown in fig. 5). The blue areas in the figure represent values of \mathbf{x}_* for which the error in the corresponding \mathbf{z}_* is low. Green and red areas indicate medium and high errors. To the right of the latent space, in fig. 4, we show images rendered with materials corresponding to the 8 points – two endpoints (MERL chrome and silver paint) and six arbitrarily chosen \mathbf{x}_* (yellow circles) that lie on the optimised path. The images with materials corresponding to intermediate points were computed using linear interpolation of pre-rendered (with global illumination) images of metals in the database (as in sec. 3.3). Note that best-path interpolation merely produces the relevant interpolated points. Although this example uses the interpolated materials to render using pre-rendered images, the same interpolated materials can indeed be used for real-time exploration (see accompanying video).

5 Discussion

Interpolation step-size

Determining the rate at which the manifold must be traversed, for a perceptually uniform transition of the interpolated BRDF, is non-trivial [Wills et al. 2009]. For example, in figure 4, it is unclear how to choose the key-points of interpolation so that there is a uniform transition from MERL-chrome to MERL-silver-paint. Our goal here is orthogonal – to develop tools to identify the non-linear manifold as well as for its fast exploration. The ability to cope with the non-linearity establishes an essential foundation for elegant connections to perceptually-motivated

studies.

Fast interpolation

We store precomputed products $\mathbf{K}^{-1}\mathbf{Z}$ for the training data. The cost of obtaining an interpolated BRDF requires computing the correlation vector \mathbf{k}_* which amounts to N calculations of the covariance function (between \mathbf{x}_* and each of the \mathbf{x}_i).

Choice of hyperparameters

Figure 7 illustrates the influence of the two parameters ℓ and μ using the Gaussian Process for 1D regression over a fixed set of points. If ℓ is too small, the influence of the training samples is very local (fig. 7a), and the uncertainty prediction (red) is very high between training samples. If ℓ is approximately equal to the mean distance between samples (fig. 7b), the interpolant has a low variance between data points. When ℓ is large (fig. 7c), the interpolant 'buckles' and bends overly between points that are close to each other. For very large values of ℓ , numerical instability causes the interpolation to fail (fig. 7d). The use of a diagonal term μ with the covariance (See Eq.5) introduces the ability to cope with training points that are too close to each other, without bending the interpolant by sacrificing continuity at the training points. Note that we do not need to optimize for ℓ since it acts as a global scale parameter that is compensated by the optimization for the latent variables \mathbf{x}_i , we therefore arbitrarily set $\ell = 1$. μ is chosen to be as small as possible ($\mu = 10^{-6}$ in our examples) while ensuring numerical stability.

Importance sampling

None of the rendering methods proposed in this paper rely on integration via Monte Carlo sampling and so importance sampling is not relevant. Nonetheless, the ability to draw samples distributed according to 2D slices of BRDF may be useful for other rendering approaches that wish to use our representation of the BRDF manifold. There are multiple ways of importance sampling from our interpolated BRDF. The straightforward way would be to exploit linearity and interpolate precomputed cumulative distribution functions (CDFs) associated with each of the materials. The CDF of the interpolated BRDF slice is easily computed on-the-fly. This method, although straightforward to implement, would introduce the cost of numerical inversion of the CDF while generating samples. Some renderers generate importance samples by first fitting parametric models (with prescribed importance sampling algorithms) to the acquired BRDFs. In that case, the parameters for each \mathbf{z}_i could be set as the latent variables \mathbf{x}_i . Instead of generating optimised latent variables our optimisation would then be used to optimise the hyperparameters. The resulting \mathbf{x}_* would correspond to the parameters for the interpolated BRDF and importance sampling could be performed as prescribed by the chosen parametric model.

Connections to work using GPs

GPs are popular tools that have been widely used. As explained in sec. 2 GPs have been explored for regression to complete missing BRDF data [Hao et al. 2015] for a single BRDF. Georgoulis used GPs to overcome the problem of ill-posedness while performing BRDF inference [Georgoulis et al. 2015], by working in the (much smaller) latent space. In this paper, we exploit the linearity of the regressed variables with respect to the observed data for fast rendering.




Figure 5: 3 slices of the 5D latent space where the path between chrome and silver-paint was optimised (fig. 4).

Choosing key-points for interpolation

We have chosen arbitrary key-points in the latent space as points \mathbf{x}_* for the choice of interpolated materials. Recent work on exploring the intuitive space of materials [Serrano et al. 2016] uses data from user-studies to learn non-linear mappings from the top 5 principal components to perceptually-meaningful attributes. They demonstrate impressive applications such as artistic exploration of the space of plausible materials. However, since their mapping to the perceptual attributes is non-linear and their interpolated BRDFs are non-linear in the measurements, the method does not lend itself to efficient rendering. Further, the domain of the mapping is limited to a linear subspace of the measured data. Our method, on the other hand, produces a non-linear manifold and yet retains the desirable property that interpolated BRDFs are linear in the observations. Interesting areas of future research would be to either modify our covariance function to account for perceptual attributes or to define a distance metric on the manifold using which perceptually-optimal interpolation key-points may be obtained.

Editing multiple materials simultaneously

Our discussion through the paper has been focussed on modifying one of the BRDFs in the scene. This trivially generalises to real-time rendering of multiple materials using ZH, for materials either on the same manifold or on different manifolds. For interpolation of pre-rendered GI images, editing p materials requires a multi-linear interpolation of dimension p , applying Equation 12 to compute intermediate points.

Anisotropic materials

There is no fundamental difference in using anisotropic BRDFs (provided that the parameterisation of these functions is consistent over the training data), which would simply appear as larger vectors in Eq.4 with a significantly larger computation time. Both our GI and real-time applications naturally extend to anisotropic BRDFs, the later being facilitated by the fact that the training data is never stored on the GPU: only the ZH coefficients of the interpolated BRDF currently displayed is.




Figure 6: Comparison of our blended image (top left), against a reference computed with global illumination (top right) for the midpoint of the path in Fig.4 (top-left). The reference image was generated using the corresponding *interpolated reflectance* data. The differences (right) are due to the absence of L_{2+} paths in our solution.




Figure 7: Overview of the effects of meta-parameters on Gaussian process interpolation with gaussian covariance (Eq. 5, see section text): (a) too small a ℓ create a local interpolant a high variance. (b) correct ℓ : low variance in between data points. (c) when ℓ is too large, the interpolant appears overly "bent". (d) even larger ℓ causes numerical instability. (e) adding a diagonal regulation term on the covariance allows to cope for close data points without bending the interpolant.

Limitations

One of the drawbacks of our interpolation scheme is that it does not guarantee conservation of energy and positivity of the resulting BRDF. The former requires that for any fixed incoming (resp. outgoing) angle, the integral of BRDF density over all outgoing (resp. incoming) angles is less than unity. As for positivity, one possibility would be to interpolate $\log(\mathbf{Z})$ instead of \mathbf{Z} at the cost of losing linearity and all the associated benefits. In practice, we observed that these two problems are insignificant as long as interpolation is restricted to regions with low prediction uncertainty (low variance).

6 Conclusion

We have presented a method for learning and traversing a non-linear manifold of measured BRDFs. The input to our method is a set of reflectivity measurements made at locations in t

the 4D domain of BRDFs. The locations are obtained by densely sampling the space composed of incident and exitant angles. First we obtain the mapping from the measurement space ($d=4$ M) to a much smaller latent space ($q=2$). For novel points in this latent space, obtained by interpolating the latent variables associated with the measured BRDFs, we use the mapping to calculate the corresponding high-dimensional point. The computed high-dimensional point corresponds to the virtual measurements associated with the interpolated latent variable. The key property of our method is that these virtual measurements can be calculated as linear combinations of the measured data. We exploit this to obtain real-time rendering and fast blending of precomputed images with global illumination, using interpolated materials.

References

- [Aittala et al. 2016] AITITALA, M., AILA, T., AND LEHTINEN, J. 2016. Reflectance modeling by neural texture synthesis. *ACM Trans. Graph.* 35, 4 (July), 65:1–65:13.
- [Ashikhmin and Premoze 2007] ASHIKHMIN, M., AND PREMOZE, S. 2007. Distribution-based BRDFs. Tech. rep., Department of Computer Science, University of Utah, March.
- [Bagher et al. 2012] BAGHER, M., SOLER, C., AND HOLZSCHUCH, N. 2012. Accurate fitting of measured reflectances using a Shifted Gamma micro-facet distribution. *Computer Graphics Forum* 31, 4 (June), 1509–1518.
- [Bilgili et al. 2011] BILGILI, A., ÖZTÜRK, A., AND KURT, M. 2011. A general BRDF representation based on tensor decomposition. *Comput. Graph. Forum* 30, 8.
- [Bonneel et al. 2011] BONNEEL, N., VAN DE PANNE, M., PARIS, S., AND HEIDRICH, W. 2011. Displacement Interpolation Using Lagrangian Mass Transport. *ACM Transactions on Graphics (SIGGRAPH ASIA 2011)* 30, 6.
- [Bonneel et al. 2016] BONNEEL, N., PEYRÉ, G., AND CUTURI, M. 2016. Wasserstein Barycentric Coordinates: Histogram Regression Using Optimal Transport. *ACM Transactions on Graphics (SIGGRAPH 2016)* 35, 4.
- [Dorsey et al. 2008] DORSEY, J., RUSHMEIER, H., AND SILLION, F. 2008. *Digital Modeling of Material Appearance*. Morgan Kaufmann Inc., San Francisco, CA, USA.
- [Georgoulis et al. 2015] GEORGIOULIS, S., VANWEDDINGEN, V., PROESMANS, M., AND GOOL, L. V. 2015. A gaussian process latent variable model for brdf inference. In *ICCV*.
- [Guarnera et al. 2016] GUARNERA, D., GUARNERA, G., GHOSH, A., DENK, C., AND GLENCROSS, M. 2016. Brdf representation and acquisition. *Comput. Graph. Forum* 35, 2 (May), 625–650.
- [Hao et al. 2015] HAO, J., LIU, Y., AND WENG, D. 2015. *A BRDF Representing Method Based on Gaussian Process*. Springer International Publishing, Cham, 542–553.
- [Harville 1997] HARVILLE, D. A. 1997. *Matrix Algebra From a Statistician's Perspective*.
- [Hooke and Jeeves 1961] HOOKE, R., AND JEEVES, T. 1961. "direct search" solution of numerical and statistical problems. *Journal of the Association for Computing Machinery (ACM)* 8, 2, 212–229.
- [Jakob 2010] JAKOB, W., 2010. Mitsuba renderer. <http://www.mitsuba-renderer.org>.
- [Kautz and McCool 1999] KAUTZ, J., AND MCCOOL, M. D. 1999. Interactive rendering with arbitrary BRDFs using separable approximations. In *Rendering Techniques*.
- [Kautz et al. 2002] KAUTZ, J., SLOAN, P.-P., AND SNYDER, J. 2002. Fast, arbitrary BRDF shading for low-frequency lighting using spherical harmonics. In *Proceedings of the 13th Eurographics Workshop on Rendering*, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, EGRW '02, 291–296.
- [Lawrence et al. 2004] LAWRENCE, J., RUSINKIEWICZ, S., AND RAMAMOORTHY, R. 2004. Efficient BRDF importance sampling using a factored representation. *ACM Trans. Graph.* 23, 3 (Aug.), 496–505.
- [Lawrence et al. 2006a] LAWRENCE, J., BEN-ARTZI, A., DECORO, C., MATUSIK, W., PFISTER, H., RAMAMOORTHY, R., AND RUSINKIEWICZ, S. 2006. Inverse shade trees for non-parametric material representation and editing. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 25, 3 (July).

- [Lawrence et al. 2006b] LAWRENCE, J., BEN-ARTZI, A., DECORO, C., MATUSIK, W., PFISTER, H., RAMAMOORTHY, R., AND RUSINKIEWICZ, S. 2006. Inverse shade trees for non-parametric material representation and editing. In *ACM SIGGRAPH 2006 Papers*, ACM, New York, NY, USA, SIGGRAPH '06, 735–745.
- [Lawrence 2005] LAWRENCE, N. 2005. Probabilistic non-linear principal component analysis with gaussian process latent variable models. *J. Mach. Learn. Res.* 6 (Dec.), 1783–1816.
- [Lensch et al. 2003] LENSCH, H. P. A., KAUTZ, J., GOESELE, M., HEIDRICH, W., AND SEIDEL, H.-P. 2003. Image-based reconstruction of spatial appearance and geometric detail. *ACM Trans. Graph.* 22, 2 (Apr.), 234–257.
- [Löw et al. 2012] LÖW, J., KRONANDER, J., YNNERMAN, A., AND UNGER, J. 2012. BRDF models for accurate and efficient rendering of glossy surfaces. *ACM Transactions on Graphics (TOG)* 31, 1 (January), 9:1–9:14.
- [Marschner et al. 1999] MARSCHNER, S. R., WESTIN, S. H., LAFORTUNE, E. P. F., TORRANCE, K. E., AND GREENBERG, D. P. 1999. Image-based BRDF measurement including human skin. In *Proceedings of the 10th Eurographics Conference on Rendering*, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, EGWR'99, 131–144.
- [Matusik et al. 2003a] MATUSIK, W., PFISTER, H., BRAND, M., AND McMILLAN, L. 2003. A data-driven reflectance model. In *ACM SIGGRAPH 2003 Papers*, ACM, New York, NY, USA, SIGGRAPH '03, 759–769.
- [Matusik et al. 2003b] MATUSIK, W., PFISTER, H., BRAND, M., AND McMILLAN, L. 2003. A data-driven reflectance model. In *ACM SIGGRAPH 2003 Papers*, ACM, New York, NY, USA, SIGGRAPH '03, 759–769.
- [Nam et al. 2016] NAM, G., LEE, J. H., WU, H., GUTIERREZ, D., AND KIM, M. H. 2016. Simultaneous acquisition of microscale reflectance and normals. *ACM Trans. Graph.* 35, 6 (Nov.), 185:1–185:11.
- [Ngan et al. 2005] NGAN, A., DURAND, F., AND MATUSIK, W. 2005. Experimental analysis of BRDF models. In *Proceedings of the Eurographics Symposium on Rendering*, Eurographics Association, 117–226.
- [Pellacini et al. 2000] PELLACINI, F., FERWERDA, J. A., AND GREENBERG, D. P. 2000. Toward a psychophysically-based light reflection model for image synthesis. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques*, ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, SIGGRAPH '00, 55–64.
- [Press et al. 2007] PRESS, W. H., TEUKOLSKY, S. A., VETTERLING, W. T., AND FLANNERY, B. P. 2007. *Numerical Recipes 3rd Edition: The Art of Scientific Computing*, 3 ed. Cambridge University Press, New York, NY, USA.
- [Ramamoorthi 2009] RAMAMOORTHY, R. 2009. Precomputation-based rendering. *Found. Trends. Comput. Graph. Vis.* 3, 4 (Apr.), 281–369.
- [Rasmussen and Williams 2006] RASMUSSEN, C. E., AND WILLIAMS, C. 2006. *Gaussian Processes for Machine Learning*.
- [Schröder and Sweldens 1995] SCHRÖDER, P., AND SWELDENS, W. 1995. Spherical wavelets: Efficiently representing functions on the sphere. In *Proceedings of the 22Nd Annual Conference on Computer Graphics and Interactive Techniques*, ACM, New York, NY, USA, SIGGRAPH '95, 161–172.
- [Serrano et al. 2016] SERRANO, A., GUTIERREZ, D., MYSKOWSKI, K., SEIDEL, H.-P., AND MASIA, B. 2016. An intuitive control space for material appearance. *ACM Trans. Graph.* 35, 6 (Nov.), 186:1–186:12.
- [Sloan et al. 2002] SLOAN, P.-P., KAUTZ, J., AND SNYDER, J. 2002. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. *ACM Trans. Graph.* 21, 3 (July), 527–536.
- [Soler et al. 2015] SOLER, C., BAGHER, M., AND NOWROUZSAHRAI, D. 2015. Efficient and Accurate Spherical Kernel Integrals using Isotropic Decomposition. *ACM Transactions on Graphics* 34, 5 (Nov.), 14.
- [Steigleder and McCool 2002] STEIGLEDER, M., AND MCCOOL, M. D. 2002. Factorization of the Ashikhmin BRDF for real-time rendering. *J. Graphics, GPU, & Game Tools* 7, 4, 61–67.

- [Sun et al. 2007] SUN, X., ZHOU, K., CHEN, Y., LIN, S., SHI, J., AND GUO, B. 2007. Interactive relighting with dynamic brdfs. *ACM Trans. Graph.* 26, 3 (July).
- [Walter et al. 2007] WALTER, B., MARSCHNER, S. R., LI, H., AND TORRANCE, K. E. 2007. Microfacet models for refraction through rough surfaces. In *Rendering Techniques*, Eurographics Association, 195–206.
- [Weinmann et al. 2015] WEINMANN, M., DEN BROK, D., KRUMPEN, S., AND KLEIN, R. 2015. Appearance capture and modeling. In *SIGGRAPH Asia 2015 Courses*, ACM.
- [Westin et al. 1992] WESTIN, S. H., ARVO, J. R., AND TORRANCE, K. E. 1992. Predicting reflectance functions from complex surfaces. In *Proceedings of the 19th Annual Conference on Computer Graphics and Interactive Techniques*, ACM, New York, NY, USA, SIGGRAPH '92, 255–264.
- [Westlund and Meyer 2001] WESTLUND, H. B., AND MEYER, G. W. 2001. Applying appearance standards to light reflection models. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*, ACM, New York, NY, USA, SIGGRAPH '01, 501–51.
- [Wills et al. 2009] WILLS, J., AGARWAL, S., KRIEGMAN, D., AND BELONGIE, S. 2009. Toward a perceptual space for gloss. *ACM Trans. Graph.* 28, 4 (Sept.), 103:1–103:15.
- [Xu et al. 2014] XU, K., CAO, Y.-P., MA, L.-Q., DONG, Z., WANG, R., AND HU, S.-M. 2014. A practical algorithm for rendering interreflections with all-frequency brdfs. *ACM Trans. Graph.* 33, 1 (Feb.), 10:1–10:16.
- [Xu et al. 2016] XU, Z., NIELSEN, J. B., YU, J., JENSEN, H. W., AND RAMAMOORTHY, R. 2016. Minimal brdf sampling for two-shot near-field reflectance acquisition. *ACM Trans. Graph.* 35, 6 (Nov.), 188:1–188:12.



**RESEARCH CENTRE
GRENOBLE – RHÔNE-ALPES**

Inovallée
655 avenue de l'Europe Montbonnot
38334 Saint Ismier Cedex

Publisher
Inria
Domaine de Voluceau - Rocquencourt
BP 105 - 78153 Le Chesnay Cedex
inria.fr

ISSN 0249-6399