EUROGRAPHICS 2018 / D. Gutierrez and A. Sheffer
(Guest Editors)

Volume 37 (2018), Number 2

A Versatile Parameterization for Measured Material Manifolds

C. Soler[1], K. Subr[2] and D.Nowrouzezahrai[3]

[1]INRIA - Grenoble University, France, [2] Univ. of Edinburgh, UK, [3] McGill University, Canada

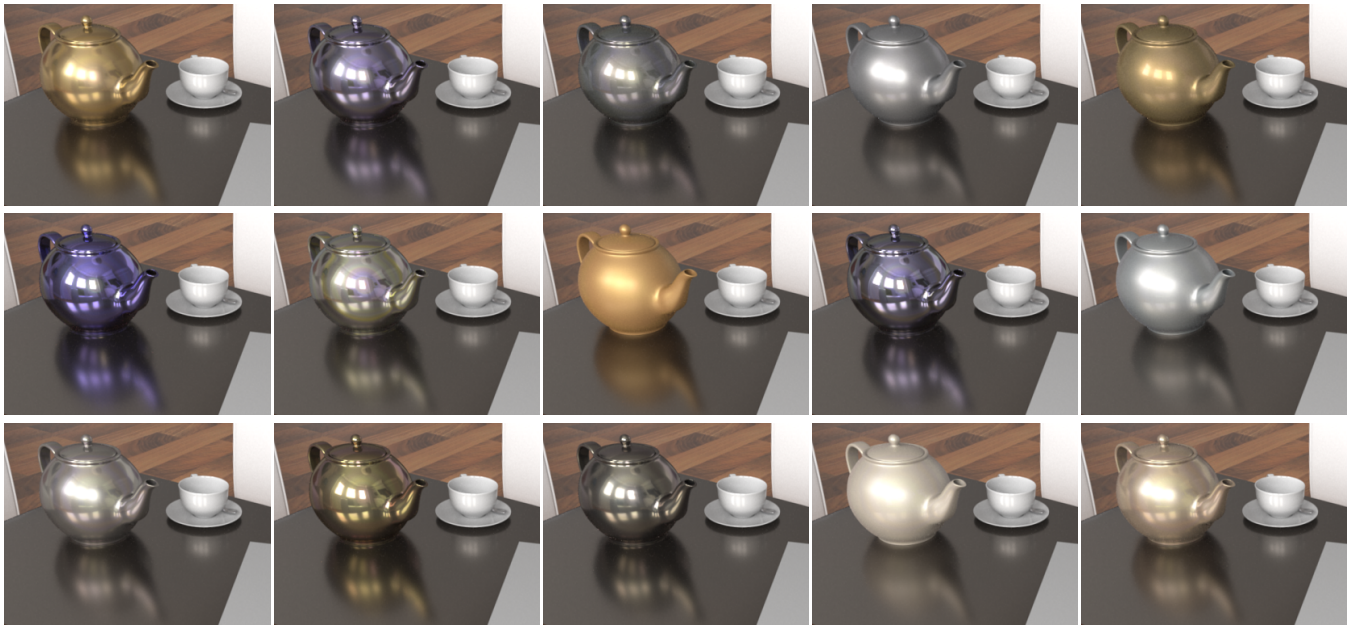

Figure 1: *Four of the images above are rendered with measured BRDFs from the MERL dataset [MPBM03b]. The remaining* 11 *are rendered with interpolated BRDFs generated from our parameterization of the non-linear manifold containing MERL materials. We explore this manifold to produce high-quality BRDFs that retain the physical properties and perceptual aspect of the original measured materials.*

Abstract

A popular approach for computing photorealistic images of virtual objects requires applying reflectance profiles measured from real surfaces, introducing several challenges: the memory needed to faithfully capture realistic material reflectance is large, the choice of materials is limited to the set of measurements, and image synthesis using the measured data is costly. Typically, this data is either compressed by projecting it onto a subset of its linear principal components or by applying non-linear methods. The former requires prohibitively large numbers of components to faithfully represent the input reflectance, whereas the latter necessitates costly algorithms to extrapolate reflectance data. We learn an underlying, low-dimensional non-linear reflectance manifold amenable to the rapid exploration and rendering of real-world materials. We show that interpolated materials can be expressed as linear combinations of the measured data, despite lying on an inherently non-linear manifold. This allows us to efficiently interpolate and extrapolate measured BRDFs, and to render directly from the manifold representation. To do so, we rely on a Gaussian process latent variable model of reflectance. We demonstrate the utility of our representation in the context of both high-performance and offline rendering with materials interpolated from real-world captured BRDFs [MPBM03b].

1 Introduction

Realistic image synthesis relies on having accurate representations of the reflective characteristics of real-world surfaces. A powerful method for obtaining this information is to *acquire* the bidirectional reflectance distribution function (BRDF) profiles directly from real-world materials. To do so typically requires capturing raw data and 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

(*e.g.* four million for the MERL dataset), it has been shown that the subspace of real-world BRDFs is of substantially lower dimensionality [MPBM03b]. Direct exploration of this subspace, which we call the manifold of BRDFs, would lead to many exciting applications: real-world BRDF interpolation for material design, BRDF inference, and data completion for partially-observed/sampled BRDFs, to name a few.

Previous approaches apply dimensionality reduction techniques to BRDF data. Linear dimensionality reduction (*i.e.* variants of PCA) require high-dimensional linear spaces (e.g., about 45D for the MERL dataset) in order to faithfully reproduce measured materials. The same applies to dictionary-based BRDFs representations [XNY*16, **?**]. When the dimension is low (namely 2 to 5 dimensions) neither kernel approximations [MPBM03b] nor projections on linear basis functions [SBN15] are sufficiently accurate for rendering applications. This is unfortunate since linear approximations often yield fast rendering techniques. 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 behave poorly at low sampling rates (*i.e.* few measured BRDFs) and/or noisy measurements – two properties of modern BRDF datasets. Charting methods have been used to learn tighter (*i.e.* 10D) non-linear BRDF manifolds [MPBM03b], but these approaches require complex algorithms for obtaining extrapolated BRDFs. Euclidean embeddings can provide a latent space that is useful for studying relationships between BRDFs, such as perceptual distances [WAKB09], but they do not permit interpolation and exploration within the embedded space.

One common approach to addressing these shortcomings is to fit non-linear low-dimensional analytic models to the acquired reflectance data and to interpolate directly in the parametric space [NDM05, WMLT07, BSH12, HP17]. Although identifying such models is non-trivial, their parameterizations naturally permit interpolation and also lead to low-dimensional manifolds: dimensionality is pre-determined by the model's degrees of freedom. Fitting parametric models can however prove numerically unstable, especially in the presence of multimodal reflectance profiles common to real-world materials. Another important drawback is that these fits provide no smoothness guarantees in the mapping from the parametric space to the measured BRDFs. This causes interpolated BRDFs to suffer from abrupt appearance transitions, even when interpolating between two similar BRDFs. Finally, realistic rendering of images with parametric BRDFs still requires costly numerical integration techniques for estimating the rendering (or reflection) equation; typically, these methods propose and rely on importance sampling schemes to accelerate this numerical integration, but the integration task remains necessary. This is summarized in Figure 2.

We devise a manifold that spans measured reflectance using a Gaussian Process Latent Variable Model (GPLVM). Our parameterization naturally yields a non-linear manifold that *exactly* interpolates the input (measured) data. Moreover, as the GPLVM maintains a linear relationship between training and interpolated samples, we can quickly extrapolate BRDFs over the non-linear manifold. We identify and leverage ties in the areas of dimensionality reduction, material appearance and light transport, to allow us to:

1. Adapt GPLVM, a well-known probabilistic dimensionality reduction tool, to the problem of optimizing for a low-dimensional latent space that maps to the space to measured BRDFs,
2. Use our mapping to efficiently traverse points on the *non-linear* manifold of acquired BRDFs using a simple *linear* interpolation of the original measured BRDFs data, and
3. Analyze the stability and physical correctness of interpolated BRDFs, and assess the plausibility of the interpolated/extrapolated BRDFs.

We demonstrate the utility of our BRDF manifold on interactive rendering and approximate global illumination applications.

Method	*NL*	*DR*	*Stab.*	*Map*	*Intp.*	*Eff.*
PCA	no	poor	good	$H \to L$	no	good
Dict.	no	none	good	$L \to H$	yes	good
Kern. PCA	yes	good	good	$H \to L$	no	OK
MDS	yes	good	good	$H \to L$	N/A	OK
Isomap	yes	good	good	$H \to L$	N/A	poor
Param.	yes	good	poor	$L \to H$	no	good
Ours	yes	good	yes	$L \to H$	yes	good

Figure 2: *Summary of dimensionality reduction approaches for measured BRDFs.* LEGEND*: Parametric (Param.), ability to capture non-linearity (NL), effectiveness of dimensionality reduction (DR), stability (Stab.), mapping from low-to-high dimensions ($L \to H$), interpolation capability (Intp.) and measurement efficiency (Eff.).*

2 Related Work

BRDFs maps pairs of directions to real values: $\rho : \omega_i \times \omega_r \to \Re$, for each color channel. The BRDF is the ratio of the reflected radiance towards ω_r to the differential radiance from light incident from ω_i, at a shading point $\mathbf{s}$ on a surface. The area of material appearance comprises over four decades of prior art, and we refer readers to comprehensive surveys on the measurement, modelling, analysis and rendering of materials [DRS08, WdBKK15, GGG*16]. Here, we instead focus on works that are most relevant to our problem.

Measurement and modelling acquired BRDFs: Marschner and colleagues [MWL*99] measure real-world materials at many incident and reflected directions, sampled over the 4D domain. Matusik et al.'s dataset includes 100 materials, each sampled with four million direction pairs [MPBM03b]. The size of these datasets has motivated more compact representations that trade accuracy for size, such as factorized tabulations [SM02], matrix decompositions [KM99], non-negative matrix factorisations [LRR04], inverse shade trees [LBAD*06b] and Tucker decompositions [BOK11]. Recent methods capture material properties with impressively few measurements [GVPG15, AAL16, NLW*16, XNY*16], however no large datasets have been created using them.

Analysis: BRDF analysis has been approached from roughly two directions: basis function approximations for individual BRDFs, and the larger study of the entire space of BRDFs. In the former, bases used during analysis include the spherical harmonics (SH) [WAT92], spherical wavelets [SS95], clustering-based empirical bases [LKG*03], constrained basis decompositions [LBAD*06b] and rotated zonal harmonics [SBN15]. When only partial/sparse observations of a single BRDF are available, Gaussian Process (GP)

regression has proven effective for BRDF completion [HLW15] by interpolating a single BRDF in angular space (whereas we interpolate accross BRDFs in a *ad-hoc* latent space). Radiometric studies of the space of BRDFs apply dimensionality reduction tools directly on the measured data. Linear approaches are unable to identify sufficiently small subspaces [MPBM03b] to facilitate practical exploration, whereas many non-linear dimensionality reduction methods (e.g., MDS, ISoMap, LLE) yield compact embeddings without explicitly providing mappings between the measured space and the manifold. Alternatively, one can fit parametric models [NDM05, AP07, BSH12, LKYU12] to the captured data to model variation across measured BRDFs, however fitting can become numerically unstable, especially for materials with multimodal reflectance distributions.

Perceptual space of BRDFs: Several works aim to understand the perceptual properties of BRDFs, often driven by user studies. This has led to reparameterizations of specific BRDFs with respect to perceptual metrics, as well as identifying semantically meaningful axes of variation for the BRDF manifold (e.g., color and gloss) [PFG00]. Furthermore, correspondences between parametric and perceptual spaces have also been formed [WM01]. Wills et al. [WAKB09] show that linear interpolation in BRDF space does not result in a linear blend of materials in the perceptual space. They obtain an embedding of BRDFs using MDS, rendering 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 [SGM*16], allowing for impressive exploration of the manifold of acquired materials from a perceptually semantic perspective (see Section 6).

Rendering: Although we focus on parameterizing the space of measured BRDFs, we briefly discuss the state of the art in rendering acquired materials since our model also enables rapid rendering. Many methods can directly render acquired materials, from accurate (but slow) physically-based methods to coarser (but faster) approximations. For interactive rendering, the compression and representation of an acquired BRDF is usually tightly coupled to a specialized rendering algorithm. View-light factorization using SVD [KM99] offers a simple rendering algorithm for a specific BRDF, and trilinear tensor factorization extensions [SZC*07] improve compression and allow for interactive exploration BRDFs. By projecting BRDFs onto SH, shading can be efficiently reformulated as a scalar product in frequency space for arbitrary 4D BRDFs [SKS02, KSS02]. More recently, an efficient spherical filtering approach using isotropic spherical decompositions [SBN15] has enabled more accurate interactive rendering solutions for arbitrary BRDFs. We demonstrate the versatility of our BRDF manifold by building atop this ISD to enable efficient real-time rendering of interpolated materials. Some of the aforementioned techniques treat visibility under certain constraints, such as static view, geometry, or with costly precomputed representations. Sun et al [SZC*07] additionally treat global illumination using precomputed transfer tensors, and a large body of work on precomputed rendering [Ram09] demonstrate methods that trade speed for accuracy in this domain. Xu et al [XCM*14] render 1-bounce interreflections, however it is unclear how their spherical-Gaussian representation can be used to realistically 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 $f : X \to \Re$, 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 [RW06].

Regression: Our goal is to predict the value z_* at arbitrary locations $\mathbf{x}_* \in X$ given pairs of observed values $(\mathbf{x}_i, z_i)$ at training locations $\mathbf{x}_i \in X$, where $i = 0, 1, ..N-1$. By definition, if we denote the vector of values $\mathbf{z} = [z_0, z_1, ..., z_{N-1}]^\intercal$ of the Gaussian process, then it follows that $\mathbf{z} \sim \mathcal{N}(\mu_z, \mathbf{K})$ where μ_z and $\mathbf{K}$ are the mean and covariance of the Gaussian distribution. The elements of the covariance matrix are $\mathbf{K}_{ij} = c(\mathbf{x}_i, \mathbf{x}_j)$ where $c : X \times X \to \Re$ 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 1D subspace of coordinates results in a 1D Gaussian distribution. So, $z_* \sim \mathcal{N}(\mu_{z*}, \sigma^2_{z*})$ with mean and variance that can be shown to satisfy [RW06]:

$$\mu_{z*} = \mathbf{k}_*^\intercal \mathbf{K}^{-1} \mathbf{z}, \tag{1}$$

$$\sigma^2_{z*} = c(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}_*^\intercal \mathbf{K}^{-1} \mathbf{k}_* \quad \text{and} \tag{2}$$

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

Interpolation is formulated as a prediction of output values μ_{z*} which requires: evaluating $\mathbf{k}_*$, the input covariance function between each training (observed) and test location, solving for $\mathbf{K}^{-1}\mathbf{z}$ (a linear system of dimension N), and computing an inner product of two vectors. This method models non-linearities by virtue of the non-linearity of the covariance function. The above may be extended from the case of a single output variable z (resp. z_i) to a d-dimensional output, where the function being learned is $f : X \to \Re^d$, by simply replacing the original observation vector $\mathbf{z}$ by an observation matrix $\mathbf{Z}$ in Eq. 1, each column of which is independently extrapolated (step 3 in fig. 3).

Latent Variable Model (LVM): In our situation, $\mathbf{z}_i \in \Re^d$ 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. Although $\mathbf{x}_i \in \Re^q$ may be chosen arbitrarily and associated with corresponding $\mathbf{z}_i$, the interpolation at some $\mathbf{x}_* \in \Re^q$ is not expected to be useful since $\mathbf{K}$ (the covariance function evaluated at all-pairs of the chosen $\mathbf{x}_i$) is implicitly dependent on the arbitrarily chosen $\mathbf{x}_i$. Although the "kernel trick" may be used in conjunction with PCA for learning non-linear mappings from observed to latent variables, they are not easily invertible. For our application of exploring the manifold of BRDFs, we seek a dual of kernel PCA, which learns the mapping from the latent space to the space of measurements (rather than the other way around). GPLVM [Law05] achieves this by optimizing the likelihood of the latent variables $\mathbf{x}_i$ given the measured data, under a GP prior on the mapping (a given covariance function). The non-linearity in the mapping is governed by the covariance function. A key property that we exploit in our work is that the interpolations remain linear with respect to the observations.

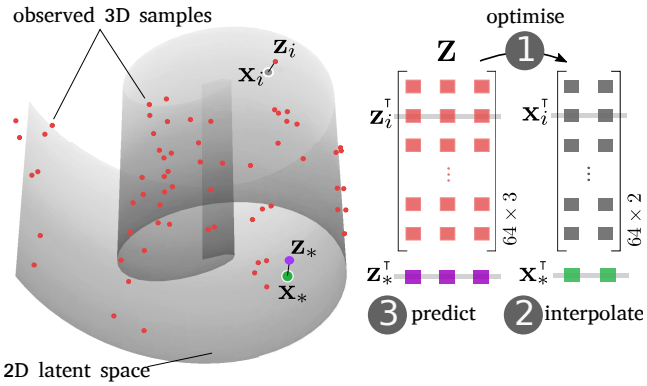

Figure 3: *Illustration of a 2D latent space (grey surface) of 64 3D measurements (red points); so $N = 64$, $d = 3$ and $q = 2$. Also shown are the latent variables associated with the measurements (grey dot), the interpolated latent variable (green dot) and the interpolated measurement (purple dot).*

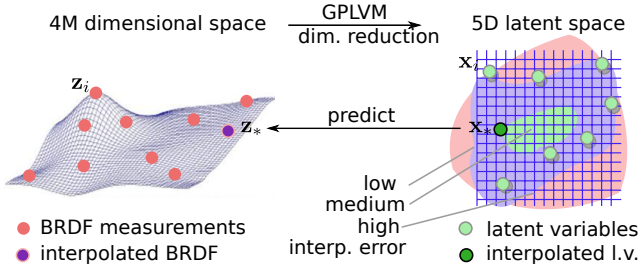

Figure 4: *We use GPLVM to identify a low-dimensional non-linear manifold on which latent variables of the measured BRDF values lie. We interpolate the latent variables and map the interpolated vector to the data space to obtain interpolated BRDF $\mathbf{z}_*$ which is a linear combination of the observations $\mathbf{z}_i$.*

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 $\mathbf{z}_i^\mathsf{T}$ 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)

$$\mathbf{z}_*^\mathsf{T} = \mathbf{b}_{x*}^\mathsf{T}\,\mathbf{Z} \tag{4}$$

where $\mathbf{b}_{x*}^\mathsf{T} = \mathbf{k}_*^\mathsf{T}\,\mathbf{K}^{-1}$. 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 $\mathbf{Z}$.

Properties: To summarize, we choose the latent space generated by GPLVM because it offers the following key properties: non-linear mapping from an optimized, low-dimensional latent space to high-dimensional observations; linearity with respect to observations; guaranteed interpolation of observed data regardless of the choice of latent variables $\mathbf{x}_i$ (if $\mathbf{x}_* = \mathbf{x}_i$, then Eq.4 yields $\mathbf{z}_* = \mathbf{z}_i$); continuity in the interpolated observations as long as the covariance function is continuous; flexibility in the choice of the dimension of the latent variable space.

3 Parameterizing the BRDF Manifold

We learn the manifold of acquired BRDFs using GPLVM (reviewed in sec. 2.1). The MERL dataset [MPBM03b] 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 $100 \times 4M$. The whole training set is first centered at 0, so that extrapolated values far from the input data converge to the average of all input BRDFs (instead of 0). We calculate optimized latent variables, obtaining a manifold where the spacing between latent variables is in accordance to the variation of the data. The output of this step is a matrix $\mathbf{X}$ of size $100 \times q$ whose rows are the latent variables $\mathbf{x}_i^\mathsf{T}$. Best results are obtained with $q = 5$ for the full MERL database although $q = 2$ used in our video, for easy illustration, yields excellent results. Smaller sets of similar materials are very well approximated with $q = 2$.

Choice of covariance function: We use a shifted squared-exponential function (widely used in the GP literature),

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

where ℓ and μ are hyperparameters that correspond to the *characteristic length scale* and *noise-filtering parameter* respectively. Since the expected number of level-zero upcrossings for a 1D stationnary GP with this covariance function is $(2\pi\ell)^{-1}$, a high value for ℓ leads to a smoother function (See e.g. [RW06] chapter 4). A small value for μ (10^{-5}) significantly improves numerical stability when inverting $\mathbf{K}$ at the cost of introducing a negligible discontinuity in the interpolant. We choose this covariance function because of its smoothness and local support, which translates into both smooth transitions across observed BRDFs (see sec. 4.1 for detailed discussion) and limited and controllable local influence of neighbor BRDFs.

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

$$L = -\frac{d}{2}\log|\mathbf{K}| - \frac{1}{2}\mathrm{tr}\left(\mathbf{K}^{-1}\mathbf{Z}\mathbf{Z}^T\right) \tag{6}$$

We perform this optimisation using direct local search [HJ61] which offers an efficient calculation scheme in our case since this only requires evaluating the cost function L once per variable $\mathbf{x}_i$. We maintain both the inverse and the determinant while changing each row and column of $\mathbf{K}$ using the Sherman-Morisson [PTVF07] and matrix determinant [Har97] formulas twice each. Figure 5 shows the evolution of the log-likelihood, over successive iterations, when fitting the full MERL database with $q = 2$. The corresponding map is displayed in Figure 10. We do not perform complicated optimisation methods such as scaled conjugate gradients (gradients of the log-likelihood may be calculated using the chain rule), even though the latent variables can be jointly optimised with the hyperparameter, because the computational complexity of gradient-based methods scales cubically with N. We initialise the latent variables using truncated linear PCA. See sec. 6 for a discussion of these choices.

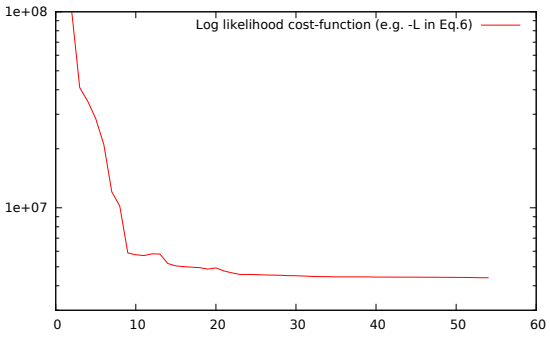

Figure 5: *Convergence of the log-likelyhood when fitting the full MERL database using a latent space of dimension 2*

Note that in this setup, ℓ can be chosen arbitrarily, since simultaneously scaling ℓ and the latent points $\mathbf{x}_i$ will keep the covariance and the log-likelyhood unchanged. Therefore we set $\ell = 1$. The second hyperparameter μ is only needed to maintain $\mathbf{K}$ well conditioned and is set to $\mu = 10^{-4}$ in all our experiments.

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) using eq. 4. The choice of $\mathbf{x}_*$ depends on the application for which the BRDF manifold needs to be traversed. For interactive exploration of the space of acquired BRDFs, we either limit ourselves to 2D latent spaces or display 2D slices of 5D latent spaces along with their corresponding projections of $\mathbf{x}_i$. The user then manually selects and drags a point within this subspace as $\mathbf{x}_*$ with real-time interpolation of the BRDFs. Further, we also tested interpolations between two chosen BRDFs along the 1D trajectory between them in the latent space.

4 Evaluation

4.1 Behavior of the interpolant

In this section, we bound the variations of interpolated data $\mathbf{z}_*$ when our mapping is applied to $\mathbf{x}_*$ in latent space. Since we have already established that $\mathbf{z}_*$ interpolates measured data $\mathbf{z}_i$ when $\mathbf{x}_* = \mathbf{x}_i$, we study the interpolant at other regions in the latent space. Equation 4 suggests that the quality of the interpolant depends on the covariance function and the inverted matrix $\mathbf{K}^{-1}$. More precisely, the behavior of the interpolant is dictated by two factors: the chosen hyperparameters (ℓ in our case) as well as the particular choice of $\mathbf{x}_i$.

We derived, via a first order approximation of the Gaussian covariance functions, that for every choice of $\mathbf{x}_*$ in the latent space the mapping $\mathbf{z}_*$ can be connected to at least one of the training data points using (See proof in supplementary material):

$$\|\mathbf{z}_* - \mathbf{z}_j\| \quad \leq \quad \frac{e^{-1/2}\sqrt{2N}}{\ell} \; \kappa(\mathbf{K}) \; \|\mathbf{Z}\| \; \|\mathbf{x}_* - \mathbf{x}_j\|. \qquad (7)$$

Here $\mathbf{z}_j$ (where $j = \beta(\mathbf{x}_*)$ depends on $\mathbf{x}_*$) is the latent variable associated with one of the input measurements (most of the time corresponding to the closest point in latent space), $\|.\|$ the L_2 norm on vectors (Frobenius norm on matrices) and $\kappa(\mathbf{K})$ the condition number of $\mathbf{K}$. Although we also provide a tightened bound using a second order approximation, eq. 7 already shows that it is desirable to optimize the ratio between the density of latent variables and hyperparameter ℓ. We do this by keeping ℓ fixed and optimizing the log-likelihood of the latent variables given the measured data under the GP prior. The reason for this choice is that rhs of Equation 6 reduces the condition number since the eigenvalues of $\mathbf{K}$ are less than 1.

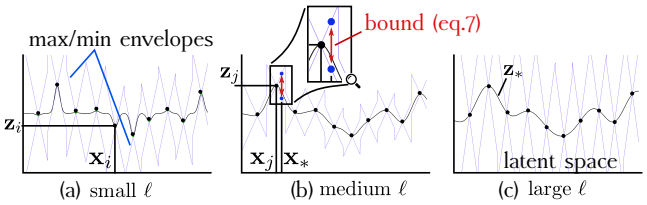

Figure 6: *Illustration of the bound on the behavior of the interpolant derived using a first order approximation.*

The derived bound can be verified to be zero when $\mathbf{x}_*$ is one of $\mathbf{x}_i$ since $\beta(\mathbf{x}_i) = i$. Another way to understand the role of $j = \beta(\mathbf{x}*)$ in Equation 7 is that the interpolant stays close to the value at training data points in proportion $\|\mathbf{Z}|\kappa(\mathbf{K})/\ell$. Equation 7 also suggests that choosing ℓ to be very small or very large causes the maximum deviation to worsen. The bound can be tightened for the special case that ℓ approaches zero. In this case, which is of purely academic interest, the interpolated BRDF is the average of all measurements everywhere in the domain except very close to the $\mathbf{x}_i$ (See Fig.6). We also show in the supplemental document that a much tighter bound of the same nature can be derived if considering the next derivatives of the covariance function.

4.2 Physical correctness

Since the interpolated data is linear in the measurement matrix $\mathbf{Z}$ (eq. 4), any properties defined using linear operators of the measured data are preserved after mapping from $\mathbf{x}_*$. Therefore, the interpolated BRDF (1) obeys Helmholtz reciprocity; (2) implicitly interpolates albedo and (3) applies to reflectivity measurements along a fixed direction of incidence.

The Gaussian process does not provide a theoretical guarrantee that the albedo of interpolated BRDFs will be limited to $[0,1]$. However, the optimization of the log-likelyhood inherently limits the oscillations of the interpolant in between observed data because it tends to keep $\kappa(\mathbf{K})$ as small as possible, tightenning the bound in Eq. 7. To illustrate this, we visualize the albedo of the interpolated BRDFs in the extreme case where we learned a 2D manifold of the entire MERL database. Figure 7 shows a portion of the manifold (for clarity) along with iso-lines of albedo. We define albedo as the ratio between the total amount of radiance reflected by a material under a uniform incident distribution and the irradiance of that distribution:

$$a(\mathbf{x}) = \frac{1}{\pi} \int_{\Omega} \int_{\Omega} f_{\mathbf{x}}(\omega_i, \omega_o)(\omega_i.\mathbf{n}) \, d\omega_i \, d\omega_o,$$

where $f_{\mathbf{x}}$ is the interpolated brdf obtained by mapping $\mathbf{x}$ into the observation space, $\mathbf{n}$ is the vertical axis, and Ω the upper hemisphere of incident directions. Every pixel in figure 7 corresponds to the albedo of the interpolated BRDF at the corresponding 2D point of

the BRDF manifold. The albedo can be observed to remain stable in the entire domain. For latent points $\mathbf{x}_*$ that are far away from the training data, the interpolated BRDF $\mathbf{z}_*$ approaches the average of all training data thus preventing unexpected oscillations.

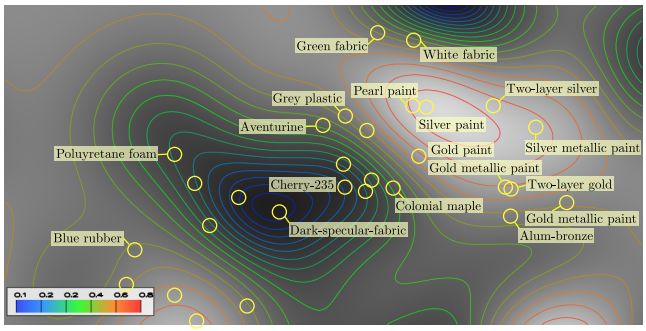

Figure 7: *Zoomed region of a 2D manifold of MERL BRDFs juxtaposed on a map of the of the interpolated BRDFs and its iso-curves. The albedo of our interpolated model is observed to be limited to* $[0, 1]$ *and it smoothly interpolates the albedo of the input materials.*

Unlike kernel PCA (or any linear decomposition over a set of basis functions), our method linearly interpolates training data. Regardless of the dimensionality of the latent space, our method will always faithfully reproduce the measured (training) BRDF data. Thus, even with a manifold of dimension 2, we can perfectly interpolate the whole MERL database. A PCA approximation of the same data with subspaces up to dimension 40 are unable to accurately reproduce measured (input) BRDFs; a minimum of 50 to 60 eigenfunctions [MPBM03b] are required for most materials in this database.

PCA, dim. 20 PCA, dim. 40 Ours, dim. 2 MERL data

Figure 8: *Comparison between our technique and the kernel trick method when it comes to reproducing the input data for a material in the MERL database (*`color-changing-paint3`*). Our method always reproduces the training materials to the last bit whatever the dimension of the parameter space, which is not the case for kernel approximations (here done with PCA at dimensions* 20 *and* 40*).*

4.3 Plausibility of interpolated BRDFs

Our choice of the manifold is not perceptually motivated. Although our main contribution is to present objective and empirical benefits of the parameterization, we also assessed our interpolated BRDFs using simple qualitative as well as quantitative experiments. The accompanying video shows several objects rendered in real-time with interpolated BRDFs corresponding to various points selected interactively in our manifold. These experiments are consistent with the behavior that we observed, that the resulting materials remain plausible everywhere in the latent space. Figure 1 shows images rendered using 11 materials chosen from randomly sampled points in the latent space and 4 images corresponding to actual BRDFs that were present in the MERL database. It is not obvious that the 4 measured materials are: row 1, column 4 (`two-layer silver`); row 2, column 1 (`chrome-steel`); row 2, column 3 (`gold-paint`); and row 3, column 2 (`brass`). Although a proper evaluation of perceptual correctness would require a thorough user study, which is beyond the scope of this work, this experiment serves as another example that suggests that our interpolated materials are plausible.

To study the behaviour of the interpolated BRDFs quantitatively, we chose interpolating paths in our latent space between arbitrary pairs of BRDFs and divided them into 20 equal intervals. At each of these 20 keypoints, we computed the perceptual distance [PR12] of the interpolated BRDF from the source and destination BRDFs. Figure 9 shows three examples: (a) blue-metallic-paint to dark-blue-paint; (b) teflon to silver-paint; and (c) black-oxidized-steel to color-changing-paint3. The blue curves and red curves are distances of the interpolated BRDFs to the source (BRDF A) and destination (BRDF B), and ideal (perceptually uniform) behavior is shown using dashed lines. Figure 9(c) is affected by the spectral behavior of color-changing-paint3.

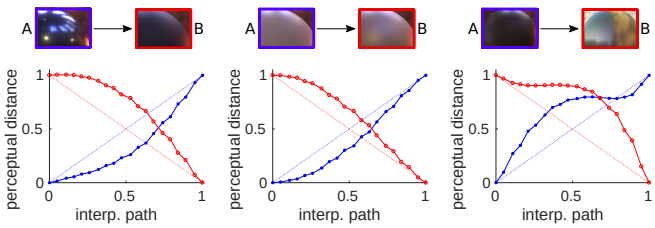

(a) (b) (c)

Figure 9: *Perceptual distances [PR12] (Y-axis) between interpolated materials (data points) at regular distances in our latent space (dimension* 5*), between BRDF A (blue curve) and BRDF B (red curve). The X- and Y- axes are normalized by the distance between A and B and the perceptual distance between BRDF A and BRDF B respectively.*

4.4 Computation time

The computational complexity of a single step of fitting the latent variables is $O(N^2)$. For the full MERL database, the total fitting time on a single core of an Intel i7 processor is 122 seconds. For interpolation, we precompute and store $\mathbf{K}^{-1}\mathbf{Z}$ for the training data. Therefore computing an interpolated BRDF amounts to calculating the correlation vector $\mathbf{k}_*$ (N evaluations of the covariance function) followed by a matrix-vector multiplication $\mathbf{K}^{-1}\mathbf{Z}$. Together this takes less than 1 second on a single CPU. For rendering, we instead interpolate data of much smaller size (such as zonal harmonic coefficients of the BRDF as described in section 5.1) for which the computation time is negligible.

5 Use cases

In this section we demonstrate the utility of our BRDF manifold parameterisation in different scenarios. We show that our extrapolation method integrates well with existing rendering techniques.

5.1 Interactive BRDF manifold exploration

Estimation of the reflectance integral, which describes the radiance arriving from a point $\mathbf{r}$ along a direction ω_o towards the centre of projection through pixel $\mathbf{p}$, forms the crux of the rendering:

$$I(\mathbf{p},\omega_o) = \int_{\mathcal{S}^2} L(\mathbf{r},\omega)\ \rho(\mathbf{r},\omega_o,\omega)\ v(\mathbf{r},\omega)\ \max(0,\omega.\mathbf{n})\ d\omega\,, \quad (8)$$

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 $\mathbf{r}$.

Fixed view: For a given ω_o, rendering a pixel $\mathbf{p}$ with a BRDF from the manifold can therefore be computed as

$$I_*(\mathbf{p}) = \sum_{i=1}^{N} \mathbf{b}^i_{x*}\ I_i(\mathbf{p}) \quad (9)$$

due to the linearity of Equation 8 wrt ρ and because any material generated by our manifold parameterization is a linear combination of the N measured BRDFs (eq. 4): $\rho_* = \sum_{i=1}^{N} \mathbf{b}^i_{x*}\ \rho_i$, where I_i is the image rendered with training material ρ_i. For some applications, such as material design, pre-rendered images (one image for BRDF in the training set) may be used to explore the interpolated appearances on the BRDF manifold *without recalculating the reflectance integral* for a fixed view. Images using interpolated materials can be obtained by linear interpolation of the pre-rendered images using $\mathbf{b}_{x*}$ as coefficients.

Dynamic view/geometry/lighting: Using the above argument, any algorithm that expresses the BRDF (and hence the image itself) by projection onto a linear basis can be adapted to use our interpolated BRDF with minimal implementational changes. We demonstrate this using the example of a recent algorithm [SBN15] which expresses the BRDF as a sum of rotated zonal harmonics (RZH) – 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 $(l_{max}+1)^2$ ZH along $2l_{max}+1$ fixed axes $\mathbf{a}_m$, where l_{max} is the degree, together form a basis that exactly spans the space of SH. Using this, they compute the shading equation in real time for large values of l_{max}. For directional (distant) lighting, where $L(.,\omega) = E(\omega)$ (temporarily ignoring the visibility term for simplicity), they derived the reflectance equation

$$I(\mathbf{p},\omega) = \sum_{l=0}^{l_{max}} \sum_{m=-l}^{l} (E \otimes \mathcal{Y}^0_l)(\mathbf{R}_n^{-1}\mathbf{a}_m)\ \lambda^m_l(\mathbf{R}_n^{-1}\omega_o). \quad (10)$$

$\mathbf{R}_n$ is a rotation that maps global into local directional coordinates so that the up direction is aligned with the shading normal $\mathbf{n}$, $E \otimes \mathcal{Y}^0_l$ denotes spherical convolution of the illumination and zonal harmonic $\mathcal{Y}^0_l$, and λ^m_l are coefficients of the BRDF projected onto rotated ZH. Because λ^m_l linearly depends on the BRDF, there exists a constant matrix $\mathbf{P_a}$ (that depends only on directions $\{\mathbf{a}_m\}$), so that the vector Λ_i^T of the $(l_{max}+1)^2$ zonal harmonic coefficients associated with reflectance ρ_i is $\Lambda_i^\mathsf{T} = \mathbf{z}_i^\mathsf{T}\mathbf{P_a}$. The interpolated ZH coefficients corresponding to ρ_* are consequently

$$\Lambda_*^\mathsf{T} \;=\; \mathbf{z}_*^\mathsf{T}\mathbf{P_a} \;=\; \mathbf{b}_{x*}^\mathsf{T}\,\mathbf{Z}\,\mathbf{P_a} \;=\; \mathbf{b}_{x*}^\mathsf{T}\,\Lambda \quad (11)$$

where the matrix Λ is formed by stacking the Λ_i^T as its rows. So Λ_* can be computed without the need for explicitly determining $\mathbf{P_a}$. Λ, which is used as $\mathbf{Z}$ in Eq.4, is $N \times (l_{max}+1)^2$. Given that Λ_* is the set of ZH coefficients for the interpolated material, we *use the shader of Soler* et al. *[SBN15] 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. This technique was used to display the dragons in Figure 10 at 25 fps (See also our accompanying video for an interactive example) using a 2D parameterization of the manifold from the full MERL data set [MPBM03b].

5.2 Real-time material editing with 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 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 or twice or more (L_{2+}) as depicted in figure 11. Since L_{1-} contains paths with at most one interaction involving ρ, its contribution to $L(\mathbf{r},\omega_o)$ is linear (affine, to be precise) in ρ. Substituting this in eq. 8 results in a separation of the image I, where the material to be modified is ρ, into $I = I^{1-} + I^{2+}$ where I^{1-} is an image that is entirely affine in ρ and I^{2+} contains the remainder of the energy. Due to this linearity, by construction,

$$I_i^{1-}(\mathbf{p},\omega) \;=\; T\,\mathbf{z}_i \quad (12)$$

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. 12 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 changed material. Here again, leveraging linearity, the image I_*^{1-} where ρ is replaced with ρ_* is obtained by interpolating pre-rendered images:

$$I_*^{1-}(\mathbf{p},\omega) = \sum_{i=1}^{N} \mathbf{b}^i_{x*}\ I_i^{1-}(\mathbf{p},\omega). \quad (13)$$

In practice we expect $L_{2+} \ll L_{1-}$, since the measure of multiple-bounce paths for which more than one of the bounces on ρ is expected to be small. We observe that applying this interpolation to calculate I_* directly rather than I_*^0 produces plausible results. A measurement of the approximation error is shown in Figure 13.

Further, figure 12 visualizes screenshots from a live session where the user interpolates in the manifold between two BRDFs from the training set. The images using interpolated materials (shown alongside) are obtained in real time 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 excellent approximations of images obtained by solving for global illumination with the interpolated material.

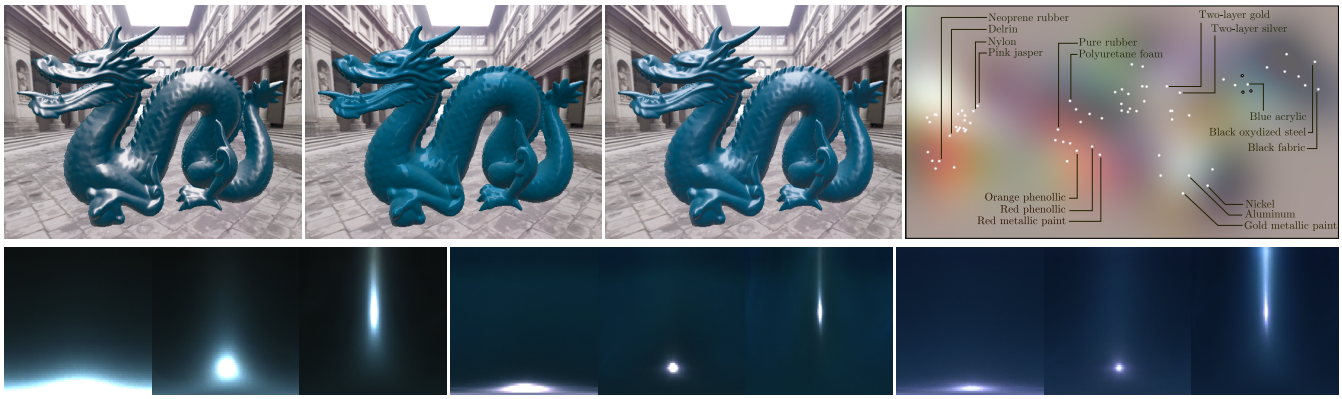

Figure 10: *Real-time exploration of our BRDF manifold. 3 materials (circle points) are chosen manually, in the vicinity of* `blue-acrylic`*, by clicking in the 2D latent space (top right). The latent variables were optimised using all* 100 *materials of the MERL database. Our proposed construction of 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. 5.1) using zonal harmonics up to* $L = 40$ *[SBN15]. The slices of the interpolated BRDFs are also visualised. Please see the accompanying video for a live demonstration.*

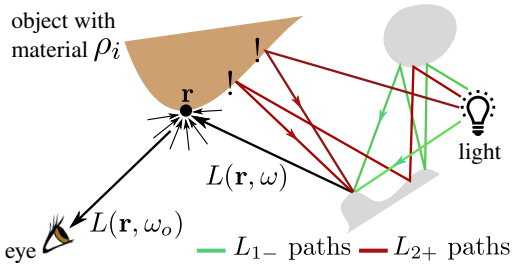

Figure 11: *An illustration of multi-bounce paths contributing to the reflectance integral at a point* $\mathbf{r}$ *on a material* ρ_i *which we wish to modify. We seggregate contributions to* $L(\mathbf{r}, \omega)$*, the radiance along one incident direction* ω *at* $\mathbf{r}$*, into* L_0 *(green) and* L_{1+} *(red) based on whether the paths contain a point on object with* ρ_i *before arrival (!) at* $\mathbf{r}$ *or not respectively:* $L(\mathbf{r}, \omega) = L_0(\mathbf{r}, \omega) + L_{1+}(\mathbf{r}, \omega)$.

5.3 Importance sampling

Although the above use cases do not require numerical integration, the ability to draw samples distributed according to 2D slices of BRDF may be useful for general rendering approaches that wish to use our representation of the BRDF manifold. There are multiple ways of importance sampling from our interpolated BRDFs. 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 we would then 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.

6 Discussion

Comparison to BRDF interpolation methods: Existing point-to-point BRDF interpolation methods [BvdPPH11] are excellent at interpolating between pairs of materials. However they do not provide a full parameterization of the BRDF manifold, which makes it impossible to explore the dataset outside the "optimal path" computed by the interpolation method.

Other 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 [HLW15] for a single BRDF. Georgoulis used GPs to overcome the problem of ill-posedness while performing BRDF inference [GVPG15], by working in the (much smaller) latent space. In this paper, we exploit the linearity of the interpolated variables with respect to the observed data for computer graphics applications.

Editing multiple materials simultaneously: Our discussion through the paper has been focused 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 13 to compute intermediate points.

Anisotropic materials: Assuming that the measurement spaces are consistent over the different measured data, there is no fundamental difference in using anisotropic BRDFs. A larger set of measurements would simply increase the size of vectors in Eq.4 and our approach would consume more larger computational resources. Both use-cases, the real-time applications as well as GI interpolations, naturally extend to anisotropic BRDFs. The consequences of the larger memory footprint of measurements is not a problem for the former since the training data is never stored on the GPU; only the ZH coefficients of the interpolated BRDF are stored.

Colored materials: There is no trivial way to separate spectral behavior of measured BRDFs. Materials such as `color-changing-paint3`, for instance, do not have a mean-

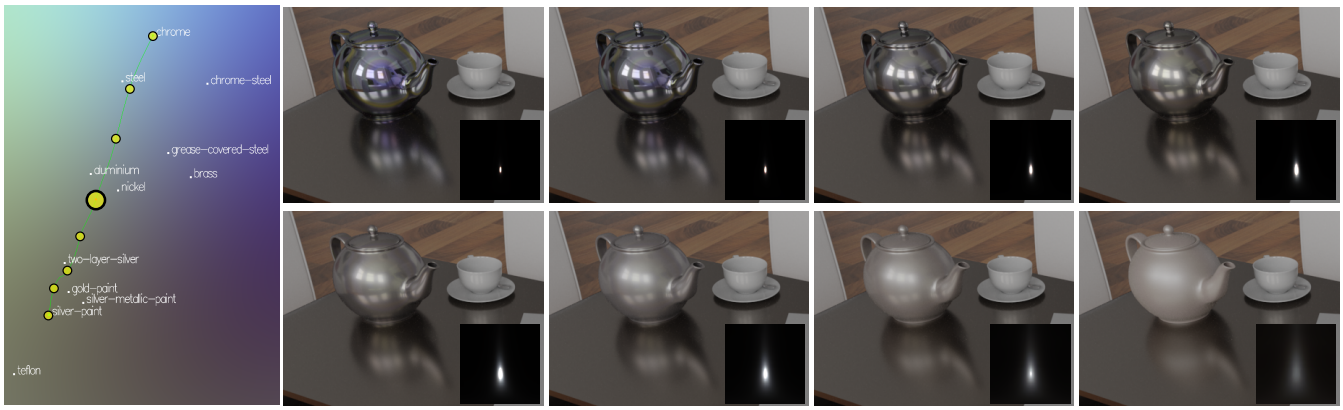

Figure 12: *Snapshots captured during a live material editing session where the material on the teapot was continuously interpolated between MERL* `chrome` *and* `silver-paint`*, in a 5D latent space. Slices of the corresponding BRDFs are displayed as insets.* ***Only the first (top left) and last (bottom right) images were pre-rendered*** *with full global illumination using Mitsuba [Jak10].* ***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 13. Note: the albedo displayed is for the 2D slice and therefore may not correspond to the materials which white dots are obtained after projection in 2D. Please see the accopanying video for more.* ***Note to reviewers****: In this figure, and Fig.13 show white squares on MacOS, and need Acrobat to display correctly.*

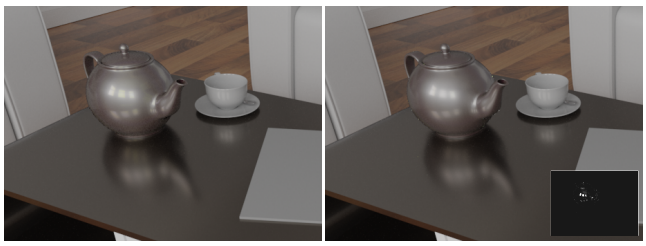

Figure 13: *Comparison of our blended image (left), against a reference (right) computed with global illumination (using the corresponding* interpolated reflectance *data) for the midpoint along the path shown in Fig.12. The differences, due to the absence of L_{2+} paths in our solution, are barely perceivable (and well below the numerical convergence error between the two images in the scaled difference inset).*

ingful average color. We concatenate all three channels into the observation (data) matrix **Z**, thus tripling the dimensionality of measurements. We experimented with both RGB and CIE Lab without significant differences and therefore used concatenated RGB measurements for all experiments in this paper.

Use in material design: Our method inherently interpolates BRDFs from the training database while retaining their physical properties and perceptual aspects, producing very realistic results that we believe could easily fit material design. It is of course clear that the interpolant will not "guess" accurate information about a missing measured material that could not be expressed as a linear combination of the materials in the training data. Similarly, extrapolating far from the training data points in an arbitrary direction will eventually reproduce the average of the training data.

Since our goal is to identify the non-linear manifold of measured data, we chose key-points for interpolated BRDFs arbitrarily without addressing perceptual attributes. Recent work on exploring the intuitive space of materials [SGM*16] 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, their method does not lend itself to efficient rendering. We hope that our manifold will trigger future research towards perceptually uniform traversal of measured BRDF manifolds.

Choice of dimensionality: While the training data is always reproduced exactly whatever the dimension of the latent space, consistent interpolation of materials depend on whether the path between two materials in latent space has other materials "nearby". While choosing a low-dimensional space (e.g. 2D) is advantageous for interactive editing, it also increases the probability of finding training points near any chosen path. Latent point optimization will naturally try to cluster materials that belong to a coherent class (e.g. metals), with the benefit that interpolated materials will share some characteristics of the materials at those nearby points. For the case where the training data contains many different features (color, material class, shininess,...), it is better to give enough dimensionality to the optimizer to consistently cluster the input data. We advise to use $d = 5$ for the entire MERL database.

Limitations: In some situations where the training data is noisy, it might be desirable to use the low dimensionnality representation to also “smooth” the input data. By construction, our manifold approximation allows this, if using a non negligible parameter μ in the covariance function of Equation 5. In this case however, our derivations for stability (performed without μ) become very conservative.

7 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 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=4M) 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 combine this with existing techniques to enable real-time rendering and fast blending of precomputed images with global illumination, for objects with interpolated materials.

Although there is no guarranty that every possibly existing material can be found in our manifold parameterisation, it is tempting to use it to fit pictures of real word materials in the hope to "capture" the BRDF from a single image, or completing missing BRDF measurements. A good image fit indeed is likely to produce a very sensible candidate for the BRDF being displayed in the matched picture.

References

[AAL16] AITTALA M., AILA T., LEHTINEN J.: Reflectance modeling by neural texture synthesis. *ACM Trans. Graph. 35*, 4 (July 2016), 65:1–65:13. 2

[AP07] ASHIKHMIN M., PREMOZE S.: *Distribution-based BRDFs*. Tech. rep., Dep. of Computer Science, University of Utah, March 2007. 3

[BOK11] BILGILI A., ÖZTÜRK A., KURT M.: A general BRDF representation based on tensor decomposition. *Comput. Graph. Forum 30*, 8 (2011). 2

[BSH12] BAGHER M. M., SOLER C., HOLZSCHUCH N.: Accurate fitting of measured reflectances using a Shifted Gamma micro-facet distribution. *Computer Graphics Forum 31*, 4 (June 2012), 1509–1518. 2, 3

[BvdPPH11] BONNEEL N., VAN DE PANNE M., PARIS S., HEIDRICH W.: Displacement interpolation using lagrangian mass transport. *ACM Trans. Graph. 30*, 6 (Dec. 2011), 158:1–158:12. 8

[DRS08] DORSEY J., RUSHMEIER H., SILLION F.: *Digital Modeling of Material Appearance*. Morgan Kaufmann Inc., San Francisco, CA, USA, 2008. 2

[GGG*16] GUARNERA D., GUARNERA G., GHOSH A., DENK C., GLENCROSS M.: Brdf representation and acquisition. *Comput. Graph. Forum 35*, 2 (May 2016), 625–650. 2

[GVPG15] GEORGOULIS S., VANWEDDINGEN V., PROESMANS M., GOOL L. V.: A gaussian process latent variable model for brdf inference. In *ICCV* (2015). 2, 8

[Har97] HARVILLE D. A.: *Matrix Algebra From a Statistician£s Perspective*. 1997. 4

[HJ61] HOOKE R., JEEVES T.: "direct search" solution of numerical and statistical problems. *Journal of the Association for Computing Machinery (ACM) 8*, 2 (1961), 212£229. 4

[HLW15] HAO J., LIU Y., WENG D.: *A BRDF Representing Method Based on Gaussian Process*. Springer International Publishing, Cham, 2015, pp. 542–553. 3, 8

[HP17] HOLZSCHUCH N., PACANOWSKI R.: A Two-Scale Microfacet Reflectance Model Combining Reflection and Diffraction. *ACM Transactions on Graphics 36*, 4 (July 2017), 12. Article 66. 2

[Jak10] JAKOB W.: Mitsuba renderer, 2010. http://www.mitsuba-renderer.org. 9

[KM99] KAUTZ J., MCCOOL M. D.: Interactive rendering with arbitrary BRDFs using separable approximations. In *Rendering Techniques* (1999). 2, 3

[KSS02] KAUTZ J., SLOAN P.-P., SNYDER J.: Fast, arbitrary BRDF shading for low-frequency lighting using spherical harmonics. In *Proceedings of the 13th Eurographics Workshop on Rendering* (Aire-la-Ville, Switzerland, Switzerland, 2002), EGRW '02, Eurographics Association, pp. 291–296. 3

[Law05] LAWRENCE N.: Probabilistic non-linear principal component analysis with gaussian process latent variable models. *J. Mach. Learn. Res. 6* (Dec. 2005), 1783–1816. 3

[LBAD*06a] LAWRENCE J., BEN-ARTZI A., DECORO C., MATUSIK W., PFISTER H., RAMAMOORTHI R., RUSINKIEWICZ S.: Inverse shade trees for non-parametric material representation and editing. *ACM Transactions on Graphics (Proc. SIGGRAPH) 25*, 3 (July 2006).

[LBAD*06b] LAWRENCE J., BEN-ARTZI A., DECORO C., MATUSIK W., PFISTER H., RAMAMOORTHI R., RUSINKIEWICZ S.: Inverse shade trees for non-parametric material representation and editing. In *ACM SIGGRAPH 2006 Papers* (New York, NY, USA, 2006), SIGGRAPH '06, ACM, pp. 735–745. 2

[LKG*03] LENSCH H. P. A., KAUTZ J., GOESELE M., HEIDRICH W., SEIDEL H.-P.: Image-based reconstruction of spatial appearance and geometric detail. *ACM Trans. Graph. 22*, 2 (Apr. 2003), 234–257. 2

[LKYU12] LÖW J., KRONANDER J., YNNERMAN A., UNGER J.: BRDF models for accurate and efficient rendering of glossy surfaces. *ACM Transactions on Graphics (TOG) 31*, 1 (January 2012), 9:1–9:14. 3

[LRR04] LAWRENCE J., RUSINKIEWICZ S., RAMAMOORTHI R.: Efficient BRDF importance sampling using a factored representation. *ACM Trans. Graph. 23*, 3 (Aug. 2004), 496–505. 2

[MPBM03a] MATUSIK W., PFISTER H., BRAND M., MCMILLAN L.: A data-driven reflectance model. In *ACM SIGGRAPH 2003 Papers* (New York, NY, USA, 2003), SIGGRAPH '03, ACM, pp. 759–769.

[MPBM03b] MATUSIK W., PFISTER H., BRAND M., MCMILLAN L.: A data-driven reflectance model. In *ACM SIGGRAPH 2003 Papers* (New York, NY, USA, 2003), SIGGRAPH '03, ACM, pp. 759–769. 1, 2, 3, 4, 6, 7

[MWL*99] MARSCHNER S. R., WESTIN S. H., LAFORTUNE E. P. F., TORRANCE K. E., GREENBERG D. P.: Image-based BRDF measurement including human skin. In *Proceedings of the 10th Eurographics Conference on Rendering* (Aire-la-Ville, Switzerland, Switzerland, 1999), EGWR'99, Eurographics Association, pp. 131–144. 2

[NDM05] NGAN A., DURAND F., MATUSIK W.: Experimental analysis of BRDF models. In *Proceedings of the Eurographics Symposium on Rendering* (2005), Eurographics Association, pp. 117–226. 2, 3

[NLW*16] NAM G., LEE J. H., WU H., GUTIERREZ D., KIM M. H.: Simultaneous acquisition of microscale reflectance and normals. *ACM Trans. Graph. 35*, 6 (Nov. 2016), 185:1–185:11. 2

[PFG00] PELLACINI F., FERWERDA J. A., GREENBERG D. P.: Toward a psychophysically-based light reflection model for image synthesis. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 2000), SIGGRAPH '00, ACM Press/Addison-Wesley Publishing Co., pp. 55–64. 3

[PR12] PEREIRA T., RUSINKIEWICZ S.: Gamut mapping spatially varying reflectance with an improved BRDF similarity metric. *Computer Graphics Forum (Proc. Eurographics Symposium on Rendering) 31*, 4 (June 2012). 6

[PTVF07] PRESS W. H., TEUKOLSKY S. A., VETTERLING W. T., FLANNERY B. P.: *Numerical Recipes 3rd Edition: The Art of Scientific Computing*, 3 ed. Cambridge Univ. Press, New York, NY, USA, 2007. 4

[Ram09] RAMAMOORTHI R.: Precomputation-based rendering. *Found. Trends. Comput. Graph. Vis. 3*, 4 (Apr. 2009), 281–369. 3

[RW06] RASMUSSEN C. E., WILLIAMS C.: *Gaussian Processes for Machine Learning*. 2006. 3, 4

[SBN15] SOLER C., BAGHER M., NOWROUZEZAHRAI D.: Efficient and Accurate Spherical Kernel Integrals using Isotropic Decomposition. *ACM Transactions on Graphics 34*, 5 (Nov. 2015), 14. 2, 3, 7, 8

[SGM*16] SERRANO A., GUTIERREZ D., MYSZKOWSKI K., SEIDEL H.-P., MASIA B.: An intuitive control space for material appearance. *ACM Trans. Graph. 35*, 6 (Nov. 2016), 186:1–186:12. 3, 9

[SKS02] SLOAN P.-P., KAUTZ J., SNYDER J.: Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. *ACM Trans. Graph. 21*, 3 (July 2002), 527–536. 3

[SM02] STEIGLEDER M., MCCOOL M. D.: Factorization of the Ashikhmin BRDF for real-time rendering. *J. Graphics, GPU, & Game Tools 7*, 4 (2002), 61–67. 2

[SS95] SCHRÖDER P., SWELDENS W.: Spherical wavelets: Efficiently representing functions on the sphere. In *Proceedings of the 22Nd Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 1995), SIGGRAPH '95, ACM, pp. 161–172. 2

[SZC*07] SUN X., ZHOU K., CHEN Y., LIN S., SHI J., GUO B.: Interactive relighting with dynamic brdfs. *ACM Trans. Graph. 26*, 3 (July 2007). 3

[WAKB09] WILLS J., AGARWAL S., KRIEGMAN D., BELONGIE S.: Toward a perceptual space for gloss. *ACM Trans. Graph. 28*, 4 (Sept. 2009), 103:1–103:15. 2, 3

[WAT92] WESTIN S. H., ARVO J. R., TORRANCE K. E.: Predicting reflectance functions from complex surfaces. In *Proceedings of the 19th Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 1992), SIGGRAPH '92, ACM, pp. 255–264. 2

[WdBKK15] WEINMANN M., DEN BROK D., KRUMPEN S., KLEIN R.: Appearance capture and modeling. In *SIGGRAPH Asia 2015 Courses* (2015), ACM. 2

[WM01] WESTLUND H. B., MEYER G. W.: Applying appearance standards to light reflection models. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 2001), SIGGRAPH '01, ACM, pp. 501–51. 3

[WMLT07] WALTER B., MARSCHNER S. R., LI H., TORRANCE K. E.: Microfacet models for refraction through rough surfaces. In *Rendering Techniques* (2007), Eurographics Association, pp. 195–206. 2

[XCM*14] XU K., CAO Y.-P., MA L.-Q., DONG Z., WANG R., HU S.-M.: A practical algorithm for rendering interreflections with all-frequency brdfs. *ACM Trans. Graph. 33*, 1 (Feb. 2014), 10:1–10:16. 3

[XNY*16] XU Z., NIELSEN J. B., YU J., JENSEN H. W., RAMAMOORTHI R.: Minimal brdf sampling for two-shot near-field reflectance acquisition. *ACM Trans. Graph. 35*, 6 (Nov. 2016), 188:1–188:12. 2