

Estimating Homogeneous Data-driven BRDF Parameters from a Reflectance Map under Known Natural Lighting

Victoria L. Cooper, James C. Bieron, Pieter Peers

Abstract—In this paper we demonstrate robust estimation of the model parameters of a fully-linear data-driven BRDF model from a reflectance map under known natural lighting. To regularize the estimation of the model parameters, we leverage the reflectance similarities within a material class. We approximate the space of homogeneous BRDFs using a Gaussian mixture model, and assign a material class to each Gaussian in the mixture model. We formulate the estimation of the model parameters as a non-linear maximum a-posteriori optimization, and introduce a linear approximation that estimates a solution per material class from which the best solution is selected. We demonstrate the efficacy and robustness of our method using the MERL BRDF database under a variety of natural lighting conditions, and we provide a proof-of-concept real-world experiment.

Index Terms—Homogeneous BRDF, Data-driven, Natural Lighting.

I. INTRODUCTION

DATA-DRIVEN appearance models [1], [2] express the Bidirectional Reflectance Distribution Function (BRDF) of a homogeneous material (i.e., a spatially invariant material) as a linear combination of a large set of measured “basis” BRDFs. The key assumption is that this large set of basis BRDFs covers the full space of BRDFs, and any BRDF in this space can be represented as convex combination of these basis BRDFs, thereby inheriting all the intricate reflectance details present in the measured basis BRDFs that can be difficult to model with analytical BRDF models. Compared to analytical BRDF models that require an expensive and fragile non-linear optimization to estimate the model parameters from reflectance measurements, data-driven BRDF models, by virtue of its linear nature, only require a linear least squares to estimate the model parameters. Recent advances have shown great promise in reconstructing a data-driven BRDF from very few measurements [3], [4]. However, these methods rely on controlled directional or point lighting. A key problem in generalizing prior methods to natural lighting is that these prior methods require a non-linear encoding (e.g., logarithmic) to compress the dynamic range of the basis BRDFs in order to regularize the estimation of the model parameters. Such non-linear encoding can only be undone after linear parameter estimation if the observations consist of direct BRDF observations (i.e., a single view and a single light direction per observation). In contrast, observations under natural lighting are the result of an integration of the BRDF times lighting over all directions,

and only linear transformations of the BRDF are transparent to this integration.

In this paper we aim to narrow the gap between inverse rendering with data-driven isotropic BRDF models and analytical BRDF models under natural lighting while retaining the robustness and simplicity of linear parameter estimation for data-driven models. We consider our work a first exploration in this direction that demonstrates that robust linear data-driven isotropic BRDF model parameter estimation under natural lighting is feasible, rather than introducing a practical and/or competitive method to current advanced inverse rendering methods that use analytical BRDF models as a basis. To focus our exploration, we will a-priori assume that the natural lighting is known and that we have a full characterization of the material reflectance under this lighting condition in the form of a reflectance map [5] (i.e., a photograph of a spherical sample), and investigate extensions to other (a-priori known) shapes in section VIII.

We desire to retain the advantages of a linear parameter estimation process, and therefore avoid non-linear encoded basis BRDFs, and directly estimate the data-driven BRDF model parameters from unmodified basis BRDFs. To regularize the estimation of the model parameters from a reflectance map under natural lighting, we leverage the reflectance similarities between BRDFs in a material class. Intuitively, we expect that it is easier to express the BRDF as a combination of a small set of similar materials than from a large set of BRDFs that span a larger spectrum of more varied materials. In other words, we aim to more accurately reconstruct a data-driven BRDF by leveraging material class (e.g., metals, paints, diffuse, and phenolics) specific knowledge. To define the material classes, we first approximate the space of homogeneous BRDFs with a Gaussian mixture model. Each normal distribution in the Gaussian mixture model represents a material class, and we assign each basis material to the class with the highest likelihood. We formulate the estimation of the data-driven BRDF model parameters as a maximum a-posteriori optimization that maximizes the likelihood that the model parameters explain the observations, as well as the likelihood that the reconstructed model belongs to the material class. However, this formulation is highly non-linear and difficult to minimize. We therefore exploit the additional observation that in high dimensional spaces everything is distant, and approximate the maximum a-posteriori optimization by computing a set of provisional data-driven BRDF solutions for each material class using an efficient linear least squares approximation. Finally, we select

V. Cooper is with Randolph-Macon College, Ashland, VA, 23005.

V. Cooper, J. Bieron, and P. Peers are with the College of William & Mary, Williamsburg, VA, 23185.

the most likely provisional least squares solution based on the maximum a-posteriori error.

We demonstrate the efficacy of our solution using the isotropic MERL BRDF database under a variety of natural lighting conditions. Furthermore, we provide a proof-of-concept real-world experiment to demonstrate that our results generalize beyond the ideal simulated experiments on the MERL BRDF database. An initial version of this work was presented at the 7th Annual Eurographics Workshop on Material Appearance Modeling [6].

II. RELATED WORK

We focus this discussion of prior work on the two key properties of our method: reflectance modeling under natural lighting, and appearance modeling with a data-driven reflectance model. We refer to the surveys of Dorsey et al. [7], and Weinmann and Klein [8] for an in-depth general overview of appearance modeling.

Reflectance Modeling under Natural Lighting A first subset of methods models surface reflectance from multiple photographs under natural lighting. Oxholm and Nishino [9] model shape and homogeneous reflectance from multiple photographs under known natural lighting. Palma et al. [10], Dong et al. [11], and Zhou et al. [12] recover spatially-varying surface reflectance under *unknown* natural lighting from a dense sampling of multiple views or multiple rotations of a subject with known shape. Xia et al. [13] extend the method of Dong et al. [11] to model spatially-varying reflectance under unknown natural lighting *and* unknown shape. These model all rely on non-linear reflectance models and estimation processes. In contrast, we employ a linear data-driven BRDF model and rely on a linear estimation process.

A second subset of methods models surface reflectance from just a *single* photograph of an object under natural lighting. In seminal work, Ramamoorthi and Hanrahan [14] lay out a spherical harmonics framework for estimating general homogeneous reflectance functions modeled by a spherical harmonics expansion. Romeiro et al. [15], [16] model the homogeneous surface reflectance using a bivariate data-driven model from an object with known shape under known and unknown natural lighting respectively. Similarly, Lombardi et al. [17] also estimate natural lighting and homogeneous surface reflectance modeled by the DSBDRF reflectance model [18]. Lin et al. [19] exploit color statistics to estimate the homogeneous reflectance from a single photograph of an object with unknown shape and known lighting. Finally, Barron and Malik [20] recover shape, lighting and spatially-varying albedo from a single photograph under unknown natural lighting. However, Barron and Malik only consider diffuse reflectance. Our method espouses the same overall goal as this second subset of methods. A reflectance map can potentially be obtained from a single observations of a convex object of known shape (e.g., sphere) or using the deep learning method of Rematas et al. [5]. However, we explicitly desire to recover a data-driven model [1] based on real-world measured reflectance.

A third subclass of methods relies on deep learning to infer reflectance properties under unknown natural lighting from a

single image. Li et al. [21] and Ye et al. [22] estimate the parameters of an analytical BRDF model [23] for a spatially-varying material. Both Li et al. and Ye et al. focus on augmenting the training data with unlabeled photographs in order to reduce the number of required labeled training data (i.e., measured SVBRDFs). Li et al. [24] present a network structure and a novel post-processing step based on conditional random fields to estimate spatially-varying reflectance parameters for an analytical micro-facet BRDF model [25]. Finally, Li et al. [26] propose a cascading network structure to iteratively estimate and refine the shape and spatially-varying surface reflectance. All of the above methods express the surface reflectance using an analytical BRDF model. In contrast, we express the surface reflectance using a more expressive data-driven model, albeit limited to a homogeneous material and under *known* natural lighting.

Data-driven Reflectance Model In seminal work, Matusik et al. [1] presented a data-driven BRDF model that expresses the surface reflectance as a weighted combination of a large set of measured BRDFs. To handle the large dynamic range between the specular peaks and the diffuse reflectance, a log-encoding is first applied to the measured basis BRDFs. Matusik et al. propose two models: a PCA based 45D linear model, and non-linear, charting based, 15D model. In follow up work, Matusik et al. [2] use the linear PCA model and show that 800 well selected and controlled view-light direction pairs are sufficient for estimating the BRDF. Nielsen et al. [3] show that by adding a Tikhonov regularization to the estimation of a log-relative encoded linear data-driven model, a good BRDF estimate can be obtained from less than 20 optimized and controlled view-light direction pairs, and for 5 photographs of a sphere lit by optimized directional light sources. Xu et al. [4] build on the method of Nielsen et al., and show that with an improved error metric, a log-relative encoded linear data-driven model can be recovered from just 2 near-field observations (photographs) under controlled directional lighting. Várva and Filip [27] extend the method of Nielsen et al. and show that 5 bidirectional samples in combination with 8 rotations is sufficient to faithfully reconstruct *anisotropic* materials; an identical log-relative encoding is used to regularize the estimation. All of the above methods estimate a data-driven BRDF from observations under directional lighting, and regularize the estimation using a non-linear encoding of the measured BRDFs. In contrast, our method uses a fully linear model and reconstructs the data-driven BRDF model from a reflectance map under uncontrolled known natural lighting.

In addition to Matusik et al.'s data-driven model, other encodings of the manifold of BRDFs have been proposed. Serrano et al. [28] design an intuitive control space for predictable editing of captured BRDFs. Key to their method is a learned mapping from perceptual attributes of the BRDFs to the underlying PCA coefficients of a log-relative encoding of the BRDF space. Serrano et al.'s control space is specially designed for intuitive editing, and is not suited for BRDF estimation. Soler et al. [29] model the the reflectance manifold with Gaussian processes, allowing them to interpolate and extrapolate measured BRDFs, and to directly render from the latent space. While this model allows Soler et al. to quickly

map from latent space to BRDF space, the reverse is difficult, making the mapping unsuitable for reconstructing BRDFs from measurements.

III. OVERVIEW

Data-driven BRDF The reflectance behavior of a homogeneous material is described by the bidirectional reflectance distribution function (BRDF) $\rho(\omega_i, \omega_o)$: a $4D$ function defined as the ratio of incident irradiance for an incident direction ω_i over the outgoing radiance for an outgoing direction ω_o .

In this paper, we follow the data-driven BRDF model of Matusik et al. [2] that characterizes the BRDF ρ as a linear combination of a large set of n measured materials $b_i, i \in [1, n]$. The underlying idea is that the set of measured BRDFs spans the space of BRDFs, and any material's BRDF should lie in this space:

$$\rho = Bw, \quad (1)$$

where we stack the BRDF ρ and basis BRDFs b_i in a vector of length p , and form the matrix B by stacking each basis vector in a column: $B = [b_1, \dots, b_n]$. The model parameters are stacked in a vector w of n scalar weights. We directly use the BRDF parameterization of the MERL BRDF database [1], and $p = 90 \times 90 \times 180$. Furthermore, similar as in Nielsen et al. [3], we consider each color channel of the 100 MERL BRDFs as a basis BRDF, and thus $n = 300$.

Due to the large dynamic range between specular peaks versus diffuse reflectance, prior work [1], [3], [4] has applied a non-linear compression function ζ to make the estimation of w less sensitive to errors on the (large) specular peaks:

$$\rho' = B'w', \quad (2)$$

where $B' = [\zeta(b_1), \dots, \zeta(b_n)]$. An expansion ζ^{-1} is applied to the compressed BRDF ρ' after computation of the weights. A common compression function is the logarithmic function, in which case Equation 2 becomes a homomorphic factorization.

Natural Lighting Prior work relied on point sample measurements of ρ for a set of incoming-outgoing direction pairs to estimate the weights w . In contrast, in this paper we aim to estimate the weights w from an observation under natural lighting. Assuming the lighting L is distant (i.e., it only depends on the incident direction $\omega_i = (\phi_i, \theta_i)$), and ignoring interreflections, we can formulate the observed radiance y as:

$$y(\omega_o) = \int_{\Omega} \rho(\omega_i, \omega_o) \cos(\theta_i) L(\omega_i) d\omega_i, \quad (3)$$

where $\cos(\theta_i)$ is the foreshortening, and Ω is the upper hemisphere of incident directions. Due to linearity of light transport, we can express Equation 3 in terms of corresponding basis observations y :

$$y = Yw, \quad (4)$$

where the weights w are the same as in Equation 1, and thus can be used to reconstruct ρ . The basis images $Y = [y_1, \dots, y_n]$ are the observations of the measured basis BRDFs b_i under the same conditions:

$$y_i = \int_{\Omega} b_i(\omega_i, \omega_o) \cos(\theta_i) L(\omega_i) d\omega_i. \quad (5)$$

Problem Statement As noted before, the dynamic range compression function ζ is essential in obtaining good data-driven BRDF reconstructions, even in the case of a very dense point sampling of light and view directions [30]. However, this compression function cannot be used when linearly estimating the weights w from observations under natural lighting. This can be seen by inserting Equation 2 in Equation 5:

$$\zeta(y_i) \neq y_i^{\zeta} = \int_{\Omega} \zeta(b_i(\omega_i, \omega_o)) \cos(\theta_i) L(\omega_i) d\omega_i. \quad (6)$$

In other words, the non-linear compression of the observation is not equivalent to the observation under natural lighting of the non-linearly compressed BRDFs. While not a problem for the basis BRDFs b_i , since we can generate the corresponding images y_i^{ζ} with any rendering system directly from the non-linear encoded basis BRDFs $\zeta(b_i)$, it is a problem for ρ , because we can only observe y the resulting radiance of ρ under natural lighting, not the reflected radiance of its non-linear compressed form $\zeta(\rho)$, and hence we do not have access to y^{ζ} . Consequently, the key problem we aim to address in this paper is to find the data-driven weights w from the observation y without relying on a non-linear compression function ζ and/or a non-linear optimization procedure for estimating the weights w .

Maximum a-posteriori Optimization Formally, our goal is to find the most likely weights w , relying on a linear estimation process, such that the conditional probability of the reconstructed data-driven homogeneous BRDF ρ is maximized given a reflectance map y under known natural lighting L :

$$\underset{w}{\operatorname{argmax}} P(\rho|y). \quad (7)$$

We will assume that the observations are in the form of a high dynamic range reflectance map (i.e., a full characterization of the reflectance radiance of a homogeneous BRDF for a fixed lighting condition). In the remainder of this paper, we will assume that the reflectance map is provided in the form of a visualization of a sphere under the target illumination (refer to Figure 4, first column of examples of measured reflectance maps).

Using Bayes' theorem, we can formulate the maximum a-posteriori (MAP) estimation of w as:

$$\underset{w}{\operatorname{argmax}} \frac{P(y|\rho)P(\rho)}{P(y)}. \quad (8)$$

Rewriting in terms of the log-likelihood, and noting that $P(y)$ is constant (i.e., the observation is given), we obtain:

$$\underset{w}{\operatorname{argmin}} (\log P(y|\rho) + \log P(\rho)). \quad (9)$$

Intuitively, this minimization searches for the data-driven BRDF parameters (w) such that it best explains the observed reflectance map y (first term) and that produces a data-driven BRDF ρ that is plausible (second term). The latter term can be interpreted as expressing the likelihood that there exists a real-world material that matches the data-driven BRDF reconstruction.

Solution Preview In order to solve the minimization problem in Equation 9, we need (for the second term) a model of

the likelihood of the BRDF estimation ρ (section IV), and (for the first term) a model for the conditional probability of the observation y given the estimated BRDF ρ (section V). However solving Equation 9 is not trivial as we will show in section V. We will therefore derive an efficient linear strategy for approximating this minimization.

In brief, our final solution (which we formally derive in the next two sections), is to compute the best BRDF reconstruction assuming the material belong to a material class (which are determined by a Gaussian Mixture Model detailed in section IV). We will compute such a provisional solution for each class, and finally select from these provisional solutions, the data-driven BRDF that produces the most accurate rerendering of the observation. A formal algorithmic summary can be found in section VI.

IV. BRDF LIKELIHOOD MODELING

Gaussian Mixture Model We propose to model the likelihood of BRDFs by a Gaussian mixture model (GMM) with K Gaussians:

$$P(\rho) = \sum_{j=1}^K \pi_j \mathcal{N}(\rho | \mu_j, \Sigma_j), \quad (10)$$

where π_j are the mixing coefficients of the j -th normal distribution \mathcal{N} with mean μ_j and covariance matrix Σ_j . Intuitively, we write the likelihood of the BRDF as the sum of how likely the BRDF is to belong to any of the K material classes; each normal distribution models the variations of BRDFs in a material class. However, we will not assign these classes manually, but compute the optimal classification based on a set of representative BRDFs (i.e., MERL BRDF database).

Expectation-Maximization An effective method for computing the parameters $\Theta = (\pi, \mu, \Sigma)$ is the Expectation Maximization algorithm using the MERL BRDFs b_i as observations. For this we define a latent variable $\gamma_j(b_i)$ that indicates the likelihood of the j -th Gaussian given a MERL BRDF b_i :

$$\gamma_j(b_i) = P(j|b_i), \quad (11)$$

$$= \frac{P(j)P(b_i|j)}{P(b_i)}, \quad (12)$$

$$= \frac{\pi_j \mathcal{N}(b_i | \mu_j, \Sigma_j)}{\sum_{j=1}^K \pi_j \mathcal{N}(b_i | \mu_j, \Sigma_j)}. \quad (13)$$

Expectation minimization iterates between estimating the latent variable $\gamma_j(b_i)$ (E-step, Equation 13), and the model parameters (M-step):

$$\pi_j = \frac{1}{n} \sum_i \gamma_j(b_i), \quad (14)$$

$$\mu_j = \frac{\sum_i \gamma_j(b_i) b_i}{\pi_j}, \quad (15)$$

$$\Sigma_j = \frac{\sum_i \gamma_j(b_i) (b_i - \mu_j)(b_i - \mu_j)^T}{\pi_j}. \quad (16)$$

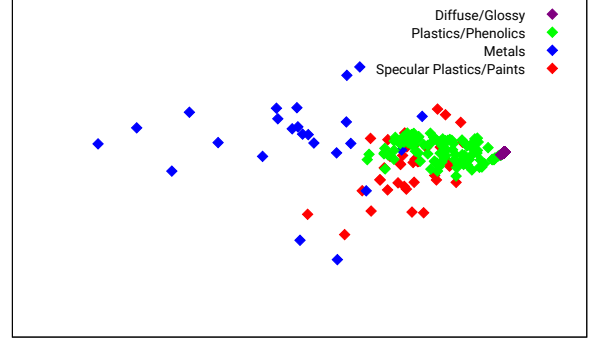


Fig. 1. 2D multi-dimensional scaling of the projected MERL BRDFs $\hat{U}^T B$ and a color-coding of the respective material classes derived from the 4D approximation of the BRDF likelihood modeled by a Gaussian mixture model.

We iterate until the log-likelihood over the MERL BRDFs converges:

$$\log P(B|\Theta) = \sum_i \log \sum_j \pi_j \mathcal{N}(\rho | \mu_j, \Sigma_j). \quad (17)$$

To bootstrap the EM algorithm, we perform a standard k-means clustering, and initialize π_j as the ratio of assigned BRDFs to the j -th cluster over the total number of MERL BRDFs (i.e., n).

Curse of Dimensionality A practical problem is that the number of observations n is significantly lower than the dimensionality of the space (i.e., p). In other words, each BRDF contains more samples than we have measured BRDFs in the MERL dataset. Consequently, the above procedure is not stable. We therefore apply a singular value decomposition (SVD) to express the observations in an n dimensional space U :

$$B = USV^T. \quad (18)$$

However, this is still a 300 dimensional space. A key issue is that even for a moderate number of dimensions any distance is very large, and thus the distance to the means μ_j are large too. Consequently, the likelihood of each Gaussian mixture (Equation 13) will always be very low and it can potentially cause numerical instabilities. To resolve this issue, we perform expectation maximization in a reduced space, and only keep the coefficients belonging to the N largest singular values. In other words, we perform expectation maximization (i.e., soft clustering) on a projection to an N dimensional subspace, and approximate the likelihood: $P(\rho) \approx P(\tilde{U}^T \rho)$, where \tilde{U} is the N dimensional basis (i.e., the first N vectors in U ; the most optimal N basis vectors that best describe the space of BRDFs).

Discussion We found that $N = 4$ offers a good balance between accuracy and numerical stability. A second parameter that needs to be set is the number of Gaussian mixtures K (or material classes). If the number of Gaussians is too low, then $P(\tilde{U}^T \rho)$ only offers a coarse approximation. However, we also found that for increasing number of K , the algorithm tends to subdivide the same Gaussian distribution, essentially over-fitting to 'special case' BRDFs (such as *Steel* which exhibits

acquisition artifacts). In practice we found that $K = 4$ offers a good approximation that nicely categorizes the materials in four recognizable distinct material classes: “diffuse and glossy” materials (137 materials), “plastics/phenolics” (99 materials), “metals” (24), and “specular plastics/paints” (40 materials); we determine membership to a material class by assigning the material to the material class with the maximum $\gamma_j(b_i)$ likelihood. In section VIII we further analyze our choice of N and K . Figure 1 shows a plot of a 2D multi-dimensional scaling of the 4D projected coordinates of the MERL BRDFs, as well as a color-coding to indicate for which material class the material has the highest affinity. Note that even though the diffuse-like material class contains 137 materials, the multi-dimensional scaling places them all close together. Please refer to the supplemental material for an exhaustive list of which material belongs to which material class.

V. DATA-DRIVEN MODEL ESTIMATION

MAP Estimation We express the likelihood of the observation given an estimate of the BRDF as:

$$P(y|\rho) = \mathcal{N}(Yw - y|\mu, \sigma), \quad (19)$$

where μ and Σ is the expected mean error and standard deviation on the reconstructions, and Yw is the rendering of the estimated BRDF under the target natural lighting. We assume that the mean error is close to zero ($\mu = 0$), and σ is proportional to the expected measurement error (e.g., camera noise).

Given the likelihood $P(\hat{U}^T \rho)$ expressed by the Gaussian mixture model (Equation 10), we can then formulate the MAP estimation (Equation 9) as:

$$\operatorname{argmin}_w \left(\frac{\|Yw - y\|^2}{\sigma^2} + \log \sum_j \pi_j \mathcal{N}(\hat{U}^T Bw | \mu_j, \Sigma_j) \right). \quad (20)$$

The first term is the data term that indicates how well (a visualization of) the BRDF $\rho = Bw$ can explain the observation y , and the second term indicates how plausible the reconstructed BRDF ρ (projected in the 4 dimensional space \hat{U}) is.

However, directly solving for the BRDF weights w using Equation 20 is not practical because of two key issues:

- 1) *Non-linear*: Equation 20 is highly non-linear and difficult to optimize due to the log of the Gaussian mixture model in the second term.
- 2) *Gaussian Mixture Model Accuracy for $P(\rho) \approx P(\hat{U}^T \rho)$* : We approximated the likelihood of the BRDF by a 4 dimensional Gaussian mixture model. This reduction in dimensionality was necessary due to the curse of dimensionality. However, it also implicitly assumes that the BRDF lies not too far from the space of plausible BRDFs. Since the likelihood is only determined based on 4 dimensions (and thus only regularizes these four), the other 296 dimensions can be set to any value (including unreasonable values that result in an implausible BRDF).

Linear MAP Approximation To alleviate the above two practical issues, we exploit the observation that the likelihood

of a basis BRDF b_i belonging to a material class m is for most basis BRDFs equivalent to an indicator function:

$$\gamma_j(b_i) \approx \delta_{j,m}. \quad (21)$$

This implies that the overlap between the Gaussians in the Gaussian mixture model is limited. In other words, there is little overlap between the different material classes, and the classification of the MERL BRDFs to a material class is unambiguous. Armed with this observation, we therefore propose to compute a provisional candidate BRDF solution for each material class $j \in [1, K]$:

$$\operatorname{argmin}_{w^{(j)}} (\log P(y|\rho, j) + \log P(\rho|j)). \quad (22)$$

Given the set of candidate solutions $w' = \{w^{(1)}, \dots, w^{(K)}\}$, we then pick the candidate that best reconstructs the target BRDF. **Per-Material Class Linear Data Term** We define the data-term similarly as in the general non-linear case, except that we only use the basis BRDFs that belong to the same material class:

$$\log P(y|\rho, j) = \|Y^{(j)}w^{(j)} - y\|^2, \quad (23)$$

where $Y^{(j)}$ is the set of observations that correspond to the basis BRDFs assigned to the j -th material class (i.e., the materials b_i for which $\gamma_j(b_i)$ is maximal).

Per-Material Class Linear Likelihood Term We express the per-material class likelihood by a single Gaussian model. We directly compute this probability on the BRDF weights $w^{(j)}$:

$$P(\rho|j) = \mathcal{N}(w^{(j)}, \mu'_j, \Sigma'_j), \quad (24)$$

where: $\mu'_j = (\frac{1}{c_j}, \dots, \frac{1}{c_j})^T$, and c_j is the number of basis BRDFs in the j -th material class. Note that $Y^{(j)}\mu'_j$ is equivalent to the mean BRDF of the material class, and μ'_j the corresponding coordinate in the j -th BRDF subspace. Σ'_j is the covariance matrix computed from all the coordinates of the MERL BRDFs that belong to the j -th material class.

Linear Least Squares Estimation Both Equation 23 and (the log likelihood of) Equation 24 are quadratic terms that define a linear system in terms of w that can be solved using a regular linear least squares. However, both terms can have a vastly different magnitude. The magnitude of the data-term depends on the error on the rendered image of the estimated BRDF. This image error depends on the resolution, the overall intensity of the lighting, and the reflectivity of the material. Similarly, the magnitude of the likelihood term depends on the number of basis BRDFs per material class. We therefore add a balancing term:

$$\lambda_j = \frac{\lambda \|y\|^2}{c_j}, \quad (25)$$

where $\|y\|^2$ is the total squared pixel intensities in the observation. We expect that the overall intensity of the observation is directly proportional to the lighting intensity and reflectivity of the BRDF, and hence the overall scale of the image error. λ is a user set constant that depends on the qualities of the lighting. An ill-conditioned lighting condition requires a larger λ value (e.g., a low frequency lighting environment is ill-conditioned for estimating specular properties [14]). In practice we found

that $\lambda = 0.5$ works well for many lighting environments, and forms a good starting point for fine-tuning λ .

The final linear least squares is:

$$\operatorname{argmin}_{w^{(j)}} \left(\|Y^{(j)}w^{(j)} - y\|^2 + \lambda_j \frac{\|w^{(j)} - \mu'_j\|^2}{\Sigma_j'^2} \right) \quad (26)$$

Selection Ideally, we would like to select the best candidate solution from w' by evaluating Equation 20. However, by a-priori assuming that a BRDF belongs to a material class j , it is possible that there is a significant mismatch between the target material and the material class. For example, attempting to model a mirror-like specular material using the diffuse material class is unlikely to produce a satisfactory result. Consequently, we cannot simply rely on the likelihood $P(\hat{U}^T \rho)$ based on the 4 dimensional Gaussian mixture model to select the best solution from w' (i.e., the other 296 dimensions can be arbitrarily wrong). We will therefore further exploit the observation of the limited overlap of the Gaussians in the mixture model, and approximate the solution per material class by enforcing that it lies in the convex hull of the subspace spanned by the BRDFs assigned to the material class, and only rely on Equation 19 to pick the best candidate from w' . We ignore the standard deviation (i.e., $\sigma = 1$) in Equation 19 as it only acts as a scale (in the log-likelihood) that does not affect the selection of the best reconstruction (i.e., minimum log-likelihood).

Color Our discussion until now only considered monochrome BRDFs; we used all color channels from the MERL BRDFs as separate basis BRDFs. A straightforward strategy for estimating a non-monochrome BRDF with three color channels, would be to execute the estimation separately for each color channel, and combine the three reconstructed monochrome BRDF into a single RGB BRDF. However, it is possible that a solution from a different material classes j is selected for each of the three color channels. Because the set of basis BRDFs for each material class are disjunct, there can be slight differences in the constructed BRDF shape for each color channel, which in turn can result in color artifacts in the combined BRDF. We circumvent this potential problem by combining the three color channels after obtaining the candidate BRDFs, and performing the selection on the RGB BRDF instead of each color channel separately. Hence, each color channel will be reconstructed with the same set of basis BRDFs.

VI. ALGORITHM SUMMARY

In summary, given a reflectance map y under known natural lighting L , and given a user provided balance parameter λ , we compute the data-driven BRDF $\rho = Bw$ as:

- 1) We precompute the Gaussian mixture model using the EM algorithm detailed in section IV. Note, this precomputation only needs to happen once for the MERL BRDF database, and is independent of the lighting.
- 2) We precompute Y by rendering a sphere with each basis BRDF b_i under the natural lighting L (Equation 5). This precomputation needs to happen for every lighting condition.

- 3) We compute the candidate solutions $w'_{\{r,g,b\}}$ for each material class by solving the linear least squares in Equation 26 per color channel.
- 4) We combine the monochrome BRDFs to a 3-channel BRDF: $w' = \{(w'_{r,1}, w'_{g,1}, w'_{b,1}), \dots, (w'_{r,K}, w'_{g,K}, w'_{b,K})\}$.
- 5) Finally, we select the candidate from w' that minimizes Equation 19 (i.e., the solution that produces the most similar rendering to the given reflectance map y).

VII. RESULTS

Experiment Setup We demonstrate our method on simulated reflectance maps in order to fully control all parameters. We generate the reflectance maps under natural lighting, by rendering a sphere lit by a light probe [31] using Mitsuba [32]; as noted in section III, we will directly use this rendered image as a representation of the reflectance map. All generated images are radiometrically linear, and we only tone map them for display. All results shown in this paper were tone mapped by a simple gamma 2.2 correction and a virtual exposure (i.e., scale factor) of 1.0; all pixel values above 1.0 and below 0.0 are clipped to the respective clipping values. We use the BRDFs in the MERL database [1] for generating reflectance maps. For each MERL BRDF, we compute a novel Gaussian mixture model on the 297 remaining MERL BRDFs (i.e., we exclude the basis BRDFs corresponding to any of the three color channels of the BRDF), and only use these 297 MERL BRDFs for reconstruction. Consequently, any reconstruction of a BRDF from the MERL BRDF database is computed using a different set of basis BRDFs. As noted in the prior sections, we compute the Gaussian mixture model on a $N = 4$ dimensional reduced space, and use $K = 4$ Gaussians in the mixture model. All reconstructions are generated with a fixed balancing factor $\lambda = 0.5$. All reconstructed BRDFs preserve, by virtue of the data-driven BRDF model, reciprocity. We do not enforce energy conservation, although this could be added by using a constrained linear least squares procedure.

Reconstruction Results Figure 2 shows reconstructions of 7 selected materials under two different light probes (i.e., *Eucalyptus Grove* and *Galileo's Tomb*). For each reconstruction (and the reference), we show a visualization of the reference/reconstructed BRDF under a *new* natural lighting condition (i.e., *Uffizi Gallery*; different than the lighting condition under which the BRDF was reconstructed) and two different light directions (i.e., a slice of the BRDF for a single incident direction for all outgoing directions). These results show that our method is able to reconstruct plausible BRDFs for a wide range of materials from a reflectance map under natural lighting. We refer to the supplemental material for the reconstructions under different natural lighting conditions.

Per-Material Class Reconstruction Figure 3 illustrates, for a selection of 4 materials, reconstructed under the *Uffizi Gallery* light probe, that the reconstructions per material class are different, and that depending on the material a different class' reconstruction is selected. We show a visualization of the reference BRDF and the reconstructions per cluster under a natural lighting condition (i.e., *Eucalyptus Grove*)




Fig. 2. Data-driven BRDF reconstructions from a reflectance map under the *Eucalyptus Grove* and the *Galileo's Tomb* light probe. We visualize the reference and reconstructed BRDFs under the *Uffizi Gallery* light probe and two different lighting directions (0° and 136°).

and a directional light. We also list the log-likelihood of the observation given the BRDF (Equation 19) below each cluster, and mark the final selected solution (i.e., minimum). While not all clusters' reconstructions are plausible, the selection process tends to pick the most plausible reconstruction. The *Metal* material class is smallest and exhibits a large diversity. As a result, we occasionally observe minor ringing artifacts in the reconstructions. We posit that a denser sampling of the *Metal* material class would reduce these ringing artifacts. For reference, we also show the linear least squares solution: $\arg\min_w ||Yw - y||^2$. As expected this yields the lowest reconstruction error (since it explicitly optimizes for this). However, the linear least squares solution does not always yield a plausible result when visualized under a different lighting condition. This is not only clearly visible under the directional light source, but also under other natural lighting conditions other than the original observed lighting (e.g., the black spot in the center of the visualizations under the *Eucalyptus Grove* light probe for *Steel* and *Red Metallic Paint*).

Captured Reflectance Map Validation To validate our method on other materials than the MERL BRDF database, we performed the following proof-of-concept experiment. We acquired three spheres with different materials (i.e., *Dense Orange Foam*, *Blue Plastic*, and *Dark Bronze*) under two different natural lighting environments shown in the insets. Next, we estimate data-driven BRDF parameters under the indoor *Chapel* lighting, and render the sphere under a novel outdoor lighting condition. We mask out any measured reflectance values that deviate from the expected measurement conditions (e.g., the dimple on the *Blue Plastic*, and the near field reflection from the stand). As can be seen in Figure 4, the rerendered reflectance maps closely resemble the acquired reference maps under the novel lighting condition (2nd row); for reference we also include the visualizations of the recovered BRDF under the capture lighting (1st row) and directional lighting (3rd row). Note, the *Dark Bronze* material exhibits anisotropic reflectance which adversely impacts the reconstruction. Nevertheless, the reconstructed BRDF remains




Fig. 3. Reconstructions for each material class for 4 selected materials observed under the *Uffizi Gallery* light probe, and revisualized under the *Eucalyptus Grove* light probe and directional lighting. We list the log-likelihood error on the observations, and mark the best solution. In addition we provide a comparison against a naive linear least squares reconstruction with the full MERL BRDFs.




Fig. 4. Reconstructions for each material class for 3 captured materials observed under indoor natural lighting (first row), and revisualized under the capture lighting condition (1st row), a novel outdoor natural lighting (2nd row) and directional lighting (3rd row). For each natural lighting condition we also provide a reference photograph. We list the log-likelihood error, and mark the best solution. In addition we provide a comparison against a naive linear least squares reconstruction.

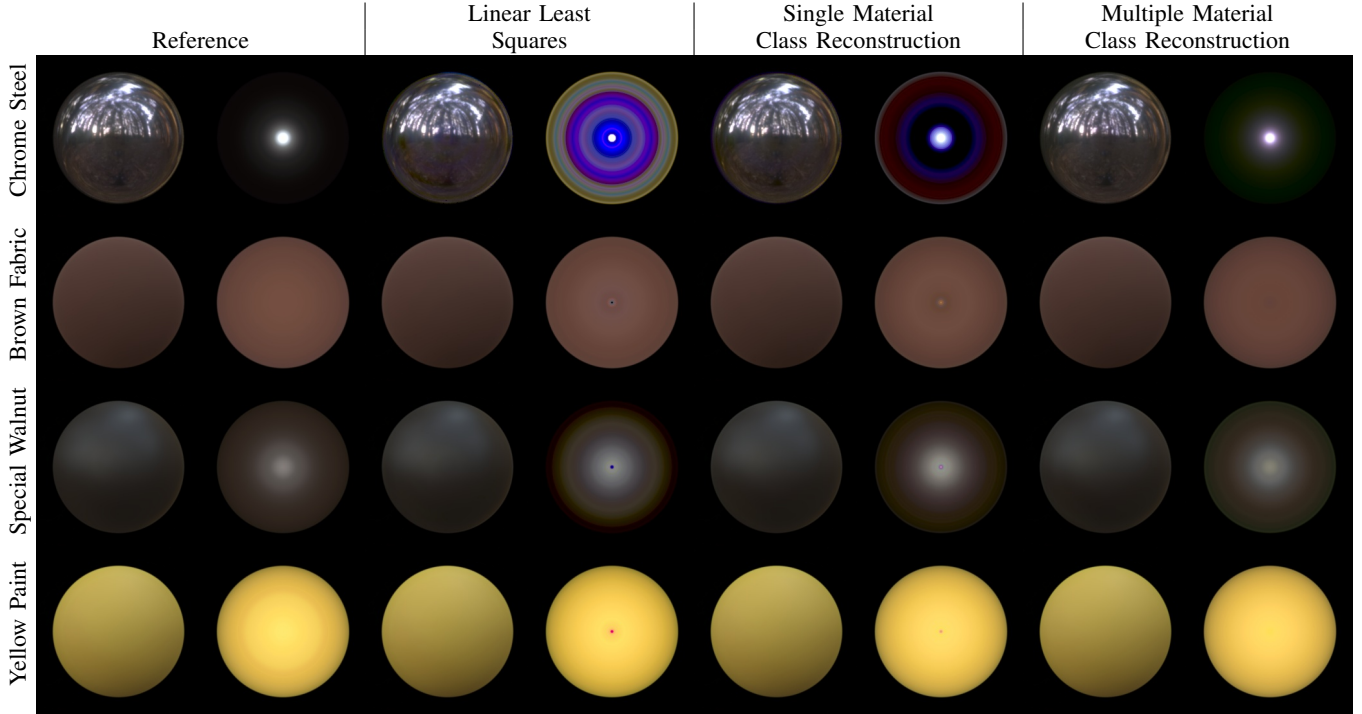


Fig. 5. Comparison between naive linear least squares data-driven BRDF reconstruction, single material class reconstruction, and our multiple material class reconstruction under *Grace Cathedral* and revisualized under the *Eucalyptus Grove* light probe and a directional light. For diffuse-like materials, both the linear least squares and single material class reconstructions exhibit a central “spike” visible under the directional lighting. For metals, strong ringing artifacts can be observed for both the linear least squares and single material class solutions.

plausible. For reference, we also include a least squares data-driven BRDF reconstruction. In addition, we also show visualizations of the reconstructed BRDFs lit by a directional light source to better demonstrate the plausibility of the reconstructions.

VIII. DISCUSSION

Comparison: Single Material Class Reconstruction To gain insight in the importance of reconstructing the BRDF per material class, we compare the reconstruction quality of the BRDF from a single material class to our multi-material class solution (Figure 5). Our results demonstrate that using a single material class improves on a naive linear least squares. However, our solution with multiple material classes outperforms the single material class case. Note that we optimized λ for the single cluster case to produce an as optimal result given the lighting conditions. In this case we reconstructed the BRDF under *Grace Cathedral* lighting using a $\lambda = 0.5$ for the single cluster case. Note, that the single material class reconstruction (Equation 26) is similar to Nielsen et al.’s [3] method, without applying a non-linear encoding of the BRDF. Additional minor differences are that Nielsen et al. subtract the median instead of the mean before computing the linear least squares and assume a unit standard deviation. Furthermore, the single class reconstruction is also similar to Romeiro et al.’s [15] method, using a linear data-driven BRDF model instead of the bivariate model. Since we a-priori assume a linear BRDF model, we want to explore the differences between the reconstruction methods, not the BRDF model representations.

In general, we found that overall our method outperforms a single material class reconstruction. The single material reconstruction tends to work equally well on phenolic and plastic materials as these are similar in BRDF shape to the mean material. However, the single material class reconstruction fails for diffuse and metal-like materials. While less strong than for the naive least squares, for diffuse materials we can observe a central “spike” visible under the directional lighting. For the metal-like materials, we typically observe strong ringing artifacts.

Impact of Lighting Accuracy Our method relies on prior knowledge of the lighting in the reflectance map. Practically, this requires capturing a light probe of the environment before or after capturing the reflectance map of the material at exactly the same location. However, in practice the light probe and the reflectance probe do not have the same size, and perfect alignment is difficult. To better understand the impact of lighting accuracy, we perform the following experiment. Given the reflectance maps of the MERL BRDFs captured under a light probe, we use our method to reconstruct the data-driven BRDF using a rotated light probe. Hence, the light probe during capture and during reconstruction contain the same lighting features, but they are not identical (i.e., rotated). Table I summarizes the relative log-likelihood errors for different light probes and material classes; visualizations for each material can be found in the supplemental material. From this experiment, we conclude that:

- 1) Diffuse materials are robust in misaligned lighting for up to 2 degrees, after which ringing artifacts become

TABLE I

RELATIVE LOG-LIKELIHOOD ERROR FOR MISALIGNED RECONSTRUCTION
LIGHT PROBES FOR ALL MERL BRDFs UNDER FOUR DIFFERENT
LIGHTING ENVIRONMENTS.

Light Probe	Rot.	Diffuse	Plastics & Phenolics	Metals	Spec. Paints & Plastics
Eucalyptus Grove	1°	0.42	1.93	3.33	0.42
	2°	1.67	9.43	10.69	8.61
	5°	10.76	82.27	73.18	116.92
	10°	46.55	180.75	121.30	250.17
Grace Cathedral	1°	2.11	3.17	2.18	0.99
	2°	8.44	13.48	9.93	6.33
	5°	51.73	78.53	47.57	58.38
	10°	203.12	200.82	255.99	226.35
Galileo's Tomb	1°	17.84	12.28	7.75	8.78
	2°	72.90	47.85	51.76	71.53
	5°	504.72	335.09	419.84	899.47
	10°	2358.62	1012.95	508.79	1678.51
Uffizi Gallery	1°	0.25	0.73	2.39	0.29
	2°	1.02	1.91	4.16	1.00
	5°	6.58	8.88	5.12	10.08
	10°	28.70	39.67	15.35	67.98

visible.

- 2) Metals are extremely sensitive to rotations (i.e., a larger relative error compared to the other material classes). Even at 1 degree they exhibit artifacts. This is likely due to three key reasons. First, the material class of metals already contains very few exemplars. Second, the variation in the metal material class is very large. Third, metals are typically very specular, and thus, rotation misalignments are more noticeable. These three reasons combined make metals very sensitive to noise and/or calibration errors.
- 3) Some light probes are more sensitive to lighting calibration artifacts. *Galileo's Tomb* and *Grace Cathedral* are more sensitive to calibration errors, while *Eucalyptus Grove* and *Uffizi Gallery* are more robust. We posit that small point lighting in the environment makes it more prone to calibration errors, while area light sources act as a low-pass filter, and therefore are less sensitive to minor rotations.

Robustness with Respect to Lighting The accuracy of estimating surface reflectance under natural lighting depends on the qualities of the lighting. Clearly, a fully uniform lighting environment will not provide sufficient cues to reconstruct the material properties. Based on prior work, general guidelines for good lighting environments are those which contain the right variety in light frequencies [14], especially those with strong edges in the lighting [11]. From experience, light probes like *Eucalyptus Grove* and *Uffizi Gallery* are well suited for BRDF estimation (i.e., uniformly distributed sharp edged area lights). *Grace Cathedral* suffers from issues with color fidelity (i.e., non-color neutral probes). *Galileo's Tomb* has trouble for difficult materials such as metals and is more sensitive to rotations (likely due to the non-uniform distribution of light sources).

Even if the lighting environment is suitable, there is still the open question how robust the method is. To validate this, we investigate the stability of the reconstruction when rotating the same lighting environment. Table II lists the stability over

TABLE II

STABILITY OF BRDF RECONSTRUCTION UNDER ROTATED LIGHTING
(AVERAGED OVER FOUR DIFFERENT LIGHTING ENVIRONMENTS).

Angle	Diffuse	Plastics & Phenolics	Metals	Spec. Paints & Plastics	Change Class
45°	30.00	4.41	1.37	2.57	4%
90°	23.09	6.65	1.47	11.08	7.75%
135°	47.33	14.42	1.36	15.61	13%
180°	49.04	12.64	2.24	6.92	10.25%
225°	13.27	4.90	1.31	2.22	7.75%
270°	8.17	4.61	2.32	1.17	9.5%
315°	5.71	3.23	1.60	1.37	4.24%

four different light probes for the four material classes for every 45° of rotation. Besides the relative log-likelihood error with respect to the zero degree rotation reconstruction, we also list how many of the reconstructions changed material class compared to the reconstruction at zero degrees. From this we observe that the method is robust to rotation and that the largest difference is when the light probe is between 135° and 180° rotated. Furthermore, we observe that specular materials are less affected than materials with a strong diffuse component. Both are related to the fact that the diffuse component is determined by the irradiance from the frontal hemisphere, whereas specular reflects the full environment lighting. Visually, we observe limited differences over all rotations; please refer to the supplemental material for exhaustive visualizations for each material.

Gaussian Mixture Model The two key parameters in our Gaussian mixture model are: K , the number of clusters or material classes, and N , the dimensionality of the subspace in which we cluster. Both parameters are independent and serve a different purpose in our algorithm. N , the dimensionality of the subspace, serves to improve numerical accuracy. First, the dimensionality of the data is limited by the number of measurements due to the Expectation-Minimization algorithm. Second, the higher the dimensionality of the space, the larger the distances, which can quickly reach the numerical limits of the floating point machine representations. The number of clusters K , on the other hand, depends on the diversity of the material types in the BRDF exemplars. For example, if the BRDF exemplars only capture diffuse reflectance then less clusters are needed than for example a BRDF dataset that also includes specular materials.

Our choice of $N = 4$ is directed by the observation that when increasing the number of dimensions, the clustering remains stable. However, using less dimensions yields a different clustering for a fixed number of clusters. When $N > 8$ we observe numerical issues that negatively impact clustering.

We decided on $K = 4$ material classes as this yields easily identifiable classes, and because we found that more clusters did not benefit the reconstruction much. Table III lists, for each cluster, for how many reconstructions under natural lighting the material class featured the best reconstruction. We averaged the result over four different lighting environments (i.e., *Uffizi Gallery*, *Grace Cathedral*, *Galileo's Tomb*, and *Eucalyptus Grove*). Our experiments indicate that independent of the number of clusters, only at most 4 of the material classes are actively used. In general, we found that for more

TABLE III
MATERIAL CLASS DISTRIBUTION OF MERL BRDF RECONSTRUCTIONS
(AVERAGED OVER FOUR LIGHTING ENVIRONMENTS) FOR $K = [2, 7]$
MATERIAL CLASSES.

K	Material Class (Cluster)						
	1	2	3	4	5	6	7
2	55.50	44.50					
3	46.75	45.25	8.00				
4	48.25	44.25	5.50	2.00			
5	48.75	37.25	9.25	4.75	0.00		
6	50.50	42.00	3.75	3.50	0.25	0.00	
7	49.75	45.25	4.00	0.50	0.50	0.00	0.00

than 4 material classes, the already undersampled *Metal* class was further subdivided, which eventually results in too small clusters for computing the necessary statistics for a Gaussian mixture model. We argue that, unless more measured basis BRDFs are used, there is no benefit in using more than 4 material classes.

Selection Criterion While our selection criterion does in the majority of cases select the best reconstruction from the different material classes, we found that in a few cases it does not select the best reconstruction, and a better reconstruction can be observed in another material class. Ideally, the selection criterion should not only include the data term (Equation 19), but evaluate the full non-linear MAP estimation loss (Equation 20). We observe that for cases where our current selection criterion prefers a suboptimal solution, that the accompanying likelihood term $P(\hat{U}^T \rho)$ is relatively large. However, a challenge is that the range of the data-term and the likelihood cover a different range due to: (a) the omission of a standard deviation scale in the data-term, and (b) the dimension reduction in the likelihood term. Finding a good balancing term is non-trivial and an interesting avenue for future research.

Balancing Factor λ We currently used a λ balancing factor of 0.5 for all our reconstructions. This λ is a compromise to produce the best result over all materials. Despite the material class and scene dependent scale factor (Equation 25), we observe that this lambda terms tends to affect the “diffuse” and “plastics / phenolics” stronger, and the “metals” and “specular plastics / paints” less. These latter two material classes exhibit not only a lower number of materials (for which we compensate), but we can also observe in Figure 1 that they are also spread out further. Consequently, the density of these material classes is significantly lower. This lower density implies that the material class is very diverse in BRDF types and that the MERL BRDF database does not densely sample these material types. Taking in account this density difference is another interesting avenue for future work. In general, we find that reconstructions from these material classes are less often selected.

Extension Beyond Reflectance Maps In all of the above results and evaluations, we have assumed the availability of a reflectance map in the form of a visualization of the reflectance on a spherical geometry. In Figure 6 we explore how well our method extends beyond reflectance maps, and to visualizations with other known geometries. Instead of precomputing Y by rendering a sphere with each basis material b_i under

natural lighting, we now render the respective geometry under the target lighting. All other steps are identical. As shown in Figure 6, we obtain qualitatively similar results for the *Buddha*, *Dragon*, *Stanford Bunny*, and *Utah Teapot* shapes. As expected, our method will fail for shapes that exhibit little normal variation such as planes and cubes. However, for such shapes, inverse rendering methods using analytical models will also fail.

Relation to Prior Work Matusik et al. [1] showed that the log-encoded BRDF space can be accurately modeled by a 45D linear subspace and a 15D non-linear manifold. While a linear model is computationally more convenient, a non-linear model offers a tighter fit to the space of BRDFs, and consequently, it contains less implausible BRDFs. Our Gaussian mixture based model can be seen as a piecewise linear approximation of the non-linear manifold of BRDFs. In contrast to Matusik et al., we work directly on the space spanned by the basis BRDF (i.e., without log-encoding). However, as shown in Figure 1 this manifold is highly non-linear too. While less tight than a full 15D non-linear model, our Gaussian mixture models strikes a balance between tightness and the ability to robustly identify the piecewise linear subspace to which the observations under natural lighting belong.

An implausible BRDF lies inside the linear subspace spanned by the linear model, but outside the non-linear BRDF manifold. Ideally, we would like to bias these implausible solutions towards the non-linear manifold to obtain a more plausible solution. Tikhonov regularization biases the reconstruction towards a mean BRDF, assuming that the solution is more plausible when closer to the mean. However, such a regularization is only efficient if the modeled space resembles a hypersphere. Nielsen et al. [3] model the BRDF space as a hypersphere by scaling the PCA basis BRDFs by the singular values. However, as noted before, the BRDF space is highly non-linear and such a hypersphere is not a tight model. Intuitively, biasing the reconstruction of a diffuse material towards the mean or median BRDF is suboptimal; the mean or median BRDF contains a rough specular lobe. Consequently, it is possible that biasing pushes the solution towards a point away from the non-linear manifold. Our solution represents the space of BRDFs as a sum of (rescaled) hyperspheres: the per-material class linear likelihood term (Equation 24) biases the solution to the mean (μ'_j) of the local hypersphere (rescaled by Σ'_j). Since each Gaussian subspace is more tight, biasing towards the mean has a lower likelihood of ending away from the non-linear manifold.

Practical Outlook As stated in the introduction, our aim is to narrow the gap between inverse rendering with data-driven isotropic BRDF models compared to analytical models. We consider our work a first exploration in this direction to demonstrate that data-driven BRDF model parameter estimation under uncontrolled natural lighting is possible, rather than introducing a competitive and/or practical method that displaces current state-of-the-art inverse rendering methods.

Our choice for furthering data-driven BRDF models is driven by a number of developments. First, recent work has shown that the classical fitting of analytical BRDF models is far from optimal [33], and preliminary research indicates




Fig. 6. Data-driven reconstructions for each of the five geometries (*Sphere*, *Buddha*, *Bunny*, *Dragon*, and *Teapot*) observed under *Eucalyptus Grove* and revisualized as a sphere under directional lighting.

similar issues for inverse rendering [34]. Second, data-driven models are by design better suited for robust linear estimation of their parameters compared to fragile non-linear optimization required for analytical models. Linear parameter estimation has the advantage that the global minimum is always reached, and that the method does not require a suitable starting point. This makes linear methods better suited for time-critical applications where the estimation process cannot fail (e.g., VR/AR applications). Furthermore, the linear nature of the estimation process also makes it potentially more amendable to a GPU implementation.

Our method is not yet suited for real-time applications. Our current bottle-neck is the rendering of B ; the basis BRDFs under the target lighting. In our current unoptimized implementation, clustering takes less than 3 seconds, and solving the linear systems takes less than a second. However, for rendering we rely on an offline rendering system (Mitsuba [32]) without using importance sampling for the BRDFs which takes about a day for the whole set B .

Currently, direct practical applications of our method are hampered by the requirement of known shape and lighting. This is in line with prior methods that rely on the data-driven BRDF [3], [4], [27], and many methods that rely on analytical models. Recently, a number of methods have been proposed that jointly estimate shape and/or lighting. However, due to

the larger number of free BRDF variables, directly adapting such methods to data-driven BRDF models is challenging. We hope our method will encourage exploration in this direction for data-driven BRDFs too.

We believe our method could be useful as a preconditioner for inverse rendering methods that reconstruct an analytical BRDF model under natural lighting too. Similar to Kang et al. [35] who first reconstruct dense measurements before fitting an analytical model, our data-driven reconstruction can serve as reconstructed measurements to fit an analytical BRDF model, allowing direct use of advanced BRDF fitting metrics [33] in inverse rendering. Additionally, more advanced learning methods than MAP likelihood estimation could also benefit from our observation that clustering materials (i.e., recognition) can aid in BRDF reconstruction.

IX. CONCLUSION

In this paper we presented a novel method for estimating the parameters of a fully linear data-driven BRDF model from a reflectance map under uncontrolled, but known, natural lighting. Our estimation method does not require any non-linear optimization, and only requires solving 4 linear least squares problems. Our method requires modest precomputations: a Gaussian mixture model clustering for the basis BRDFs, and for each natural lighting conditions, renderings of each basis

material. We demonstrated the accuracy and robustness of our method on the MERL BRDF database, and validated our method on real-world measurements.

For future work we would like to explore better selection criteria and a per-material class λ_j density correction factor.

ACKNOWLEDGMENT

This work was supported in part by NSF grants IIS-1350323 and IIS-1909028, and gifts from Google, Activision, and Nvidia.

REFERENCES

- [1] W. Matusik, H. Pfister, M. Brand, and L. McMillan, “A data-driven reflectance model,” *ACM Trans. Graph.*, vol. 22, no. 3, pp. 759–769, Jul. 2003.
- [2] —, “Efficient isotropic BRDF measurement,” in *Rendering Techniques*, 2003, pp. 241–248.
- [3] J. B. Nielsen, H. W. Jensen, and R. Ramamoorthi, “On optimal, minimal brdf sampling for reflectance acquisition,” *ACM Trans. Graph.*, vol. 34, no. 6, Oct. 2015.
- [4] Z. Xu, J. B. Nielsen, J. Yu, H. W. Jensen, and R. Ramamoorthi, “Minimal brdf sampling for two-shot near-field reflectance acquisition,” *ACM Trans. Graph.*, vol. 35, no. 6, pp. 188:1–188:12, Nov. 2016.
- [5] K. Rematas, T. Ritschel, M. Fritz, E. Gavves, and T. Tuytelaars, “Deep reflectance maps,” in *CVPR*, 2016.
- [6] V. L. Cooper, J. C. Bieron, and P. Peers, “Estimating homogeneous data-driven BRDF parameters from a reflectance map under known natural lighting,” in *Eurographics Workshop on Material Appearance Modeling*, June 2019.
- [7] J. Dorsey, H. Rushmeier, and F. Sillion, *Digital Modeling of Material Appearance*. Morgan Kaufmann Publishers Inc., 2008.
- [8] M. Weinmann and R. Klein, “Advances in geometry and reflectance acquisition,” in *ACM SIGGRAPH Asia, Course Notes*, 2015.
- [9] G. Oxholm and K. Nishino, “Shape and reflectance estimation in the wild,” *IEEE PAMI*, vol. 38, no. 2, pp. 376–389, Feb. 2016.
- [10] G. Palma, M. Callieri, M. Dellepiane, and R. Scopigno, “A statistical method for svbrdf approximation from video sequences in general lighting conditions,” *Comput. Graph. Forum*, vol. 31, no. 4, pp. 1491–1500, 2012.
- [11] Y. Dong, G. Chen, P. Peers, J. Zhang, and X. Tong, “Appearance-from-motion: Recovering spatially varying surface reflectance under unknown lighting,” *ACM Trans. Graph.*, vol. 33, no. 6, pp. 193:1–193:12, 2014.
- [12] Z. Zhou, G. Chen, Y. Dong, D. Wipf, Y. Yu, J. Snyder, and X. Tong, “Sparse-as-possible svbrdf acquisition,” *ACM Trans. Graph.*, vol. 35, November 2016.
- [13] R. Xia, Y. Dong, P. Peers, and X. Tong, “Recovering shape and spatially-varying surface reflectance under unknown illumination,” *ACM Trans. Graph.*, vol. 35, no. 6, December 2016.
- [14] R. Ramamoorthi and P. Hanrahan, “A signal-processing framework for inverse rendering,” in *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*, ser. SIGGRAPH ’01, 2001, pp. 117–128.
- [15] F. Romeiro, Y. Vasilyev, and T. Zickler, “Passive reflectometry,” in *ECCV*, 2008, pp. 859–872.
- [16] F. Romeiro and T. Zickler, “Blind reflectometry,” in *ECCV*, 2010, pp. 45–58.
- [17] S. Lombardi and K. Nishino, “Reflectance and illumination recovery in the wild,” *IEEE PAMI*, vol. 38, no. 1, pp. 129–141, 2016.
- [18] K. Nishino and S. Lombardi, “Directional statistics-based reflectance model for isotropic bidirectional reflectance distribution functions,” *J. Opt. Soc. Am. A*, vol. 28, no. 1, pp. 8–18, Jan 2011.
- [19] Y. Lin, P. Peers, and A. Ghosh, “On-site example-based material appearance acquisition,” *Computer Graphics Forum*, vol. 38, no. 4, July 2019.
- [20] J. T. Barron and J. Malik, “Shape, illumination, and reflectance from shading,” *IEEE PAMI*, 2015.
- [21] X. Li, Y. Dong, P. Peers, and X. Tong, “Modeling surface appearance from a single photograph using self-augmented convolutional neural networks,” *ACM Trans. Graph.*, vol. 36, no. 4, pp. 45:1–45:11, Jul. 2017.
- [22] W. Ye, X. Li, Y. Dong, P. Peers, and X. Tong, “Single photograph surface appearance modeling with self-augmented CNNs and inexact supervision,” *Computer Graphics Forum*, vol. 37, no. 7, Oct 2018.
- [23] G. J. Ward, “Measuring and modeling anisotropic reflection,” *SIGGRAPH Comput. Graph.*, vol. 26, no. 2, pp. 265–272, 1992.
- [24] Z. Li, K. Sunkavalli, and M. K. Chandraker, “Materials for masses: Svbrdf acquisition with a single mobile phone image,” in *ECCV*, 2018.
- [25] B. Walter, S. R. Marschner, H. Li, and K. E. Torrance, “Microfacet models for refraction through rough surfaces,” in *Rendering Techniques*, 2007, pp. 195–206.
- [26] Z. Li, Z. Xu, R. Ramamoorthi, K. Sunkavalli, and M. Chandraker, “Learning to reconstruct shape and spatially-varying reflectance from a single image,” *ACM Trans. Graph.*, vol. 37, no. 6, Dec. 2018.
- [27] R. Vávra and J. Filip, “Minimal sampling for effective acquisition of anisotropic brdfs,” *Computer Graphics Forum*, vol. 35, no. 7, pp. 299–309, 2016.
- [28] A. Serrano, D. Gutierrez, K. Myszkowski, H.-P. Seidel, and B. Masia, “An intuitive control space for material appearance,” *ACM Trans. Graph.*, vol. 35, no. 6, 2016.
- [29] C. Soler, K. Subr, and D. Nowrouzezahrai, “A versatile parameterization for measured material manifolds,” *Computer Graphics Forum*, vol. 37, no. 2, pp. 135–144, 2018.
- [30] M. M. Bagher, J. Snyder, and D. Nowrouzezahrai, “A non-parametric factor microfacet model for isotropic brdfs,” *ACM Trans. Graph.*, vol. 35, no. 5, pp. 159:1–159:16, 2016.
- [31] P. Debevec, “Light probe gallery,” <http://www.pauldebevec.com/Probes/>, 1998.
- [32] W. Jakob, “Mitsuba: Physically based renderer,” <https://www.mitsuba-renderer.org>, 2010.
- [33] J. C. Bieron and P. Peers, “An adaptive brdf fitting metric,” *Computer Graphics Forum*, vol. 39, no. 4, July 2020.
- [34] —, “An adaptive metric for brdf appearance matching,” in *Eurographics Workshop on Material Appearance Modeling*, June 2020.
- [35] K. Kang, Z. Chen, J. Wang, K. Zhou, and H. Wu, “Efficient reflectance capture using an autoencoder,” *ACM Trans. Graph.*, vol. 37, no. 4, 2018.



Victoria L. Cooper is an assistant professor at Randolph-Macon College since 2019. Cooper received her Masters degree and her Ph.D. from the College of William & Mary in 2014 and 2021 respectively. Her research interests are in data-driven appearance modeling under uncontrolled natural lighting.



James C. Bieron is a Ph.D. candidate at the College of William & Mary supervised by Pieter Peers. He received his Masters degree in Computer Science at the College of William & Mary in 2016. His primary research interests are in appearance modeling, inverse rendering, and BRDF fitting.



Pieter Peers is an associate professor at the College of William & Mary. Peers received his Masters Degree and his Ph.D. from the Katholieke Universiteit Leuven, Belgium. His primary research interests are in appearance modeling, computational lighting, and inverse rendering.