

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

Victoria L. Cooper James C. Bieron Pieter Peers
College of William & Mary

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. Next, we compute a linear solution per material class. Finally, we select the best candidate as the final estimate. We demonstrate the efficacy and robustness of our method using the MERL BRDF database under a variety of natural lighting conditions.

1. Introduction

Data-driven appearance models [MPBM03a] express the Bidirectional Reflectance Distribution Function (BRDF) of a homogeneous 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. Recent advances have shown great promise in reconstructing a data-driven BRDF from very few measurements [NJR15, XNY*16]. However, these methods rely on controlled directional or point lighting.

In this paper we aim to narrow the gap between inverse rendering with data-driven BRDF models and analytical BRDF models under natural lighting while retaining the robustness and simplicity of linear parameter estimation for data-driven models. 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 [RRF*16].

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. We therefore, first approximate the space of homogeneous BRDFs with a Gaussian mixture model. Each normal distribution in the Gaus-

sian 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 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 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 an efficient linear least squares approximation per material class. Finally, we select the most likely provisional least squares solution based on the maximum a-posteriori error.

2. Related Work

We focus this discussion of prior work on: reflectance modeling under natural lighting, and appearance modeling with a data-driven reflectance model. We refer to the surveys of Dorsey et al. [DRS08], and Weinmann and Klein [WK15] 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 [ON16, PCDS12, DCP*14, ZCD*16, XDPT16]. These methods 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, albeit limited to a homogeneous material and under *known* natural lighting.

A second subset models surface reflectance from just a *single* photograph under natural lighting, using deep learning [LDPT17, YLD*18, LSC18, LXR*18], or without deep learning [RH01, RVZ08, RZ10, LN16, BM15]. Our method espouses the same goal as this second subset. However, we explicitly desire to recover a data-driven model [MPBM03a].

Data-driven Reflectance Model In seminal work, Matusik et al. [MPBM03a] 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. Nielsen et al. [NJR15] and Xu et al. [XNY*16] show that with appropriate regularization a good data-driven BRDF can be reconstructed from very few observations. 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.

3. Overview

Data-driven BRDF In this paper, we follow the data-driven BRDF model of Matusik et al. [MPBM03b] that characterizes the BRDF ρ as a linear combination of n measured materials $b_i, i \in [1, n]$: $\rho = Bw$, 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 parameterization of the MERL BRDF database [MPBM03a], and $p = 90 \times 90 \times 180$. We do not apply any logarithmic compression as in prior work. Furthermore, we consider each color channel of the 100 MERL BRDFs as a basis BRDF (i.e., $n = 300$).

Natural Lighting 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:

$$y(\omega_o) = \int_{\Omega} \rho(\omega_i, \omega_o) \cos(\theta_i) L(\omega_i) d\omega_i, \quad (1)$$

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 1 in terms of corresponding basis observations $y = Yw$, where the weights w are the same as before, and thus can be used to reconstruct ρ . The basis images $Y = [y_0, \dots, y_n]$ are the observations of the measured basis BRDFs b_i under the same conditions.

Problem Statement Prior work relied on a dynamic range compression function to obtain good data-driven BRDF reconstructions. However, this compression function cannot be used when linearly estimating the weights w from observations under natural lighting. 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 and/or a non-linear optimization procedure for estimating the weights w . We will assume that the observations are in the form of a high dynamic range reflectance map provided as a visualization of a sphere under the target illumination.

Maximum a-posteriori Optimization Our goal is to find the most likely weights w , relying on a linear estimation process. Using

Bayes' theorem, we can formulate the log-likelihood maximum a-posteriori estimation of w as:

$$\operatorname{argmin}_w (\log P(y|\rho) + \log P(\rho)). \quad (2)$$

In order to solve this minimization, we need a model of the likelihood of the BRDF estimation ρ , a model for the conditional probability of the observation y given the estimated BRDF ρ .

4. BRDF Likelihood Modeling

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

$$P(\rho) = \sum_{j=1}^k \pi_j \mathcal{N}(\rho | \mu_j, \Sigma_j), \quad (3)$$

where π_j are the mixing coefficients of the j -th normal distribution \mathcal{N} with mean μ_j and covariance matrix Σ_j .

Due to the limited number of basis BRDFs (i.e., 300), we cannot directly perform Expectation-Maximization (EM) to compute the mixture model. To resolve this issue, we perform EM in a reduced space, and only keep the N largest singular values.

We found that $N = 4$ offers a good balance between accuracy and numerical stability. Furthermore, we set the number of Gaussian mixtures to $K = 4$, which 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).

5. Data-driven Model Estimation

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

$$P(y|\rho) = \mathcal{N}(Yw - y | \mu, \sigma), \quad (4)$$

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 camera noise.

Directly solving the MAP estimation poses two problems: (1) the resulting equation (Equation 2) is highly non-linear and difficult to solve, and (2) $P(\rho)$ only linearly constrains 4 coefficients of w due to the dimension reduction.

Linear MAP Approximation To alleviate the above two practical issues, we exploit the observation that the overlap between the Gaussians in the Gaussian mixture model is limited. We therefore propose to compute a candidate BRDF for each material class $j \in [1, k]$:

$$\operatorname{argmin}_{w^{(j)}} \left(\|Y^{(j)} w^{(j)} - y\|^2 + \lambda_j \frac{\|w^{(j)} - \mu_j'\|^2}{\Sigma_j^2} \right), \quad (5)$$

with μ_j' and Σ_j the mean and standard deviations of the j -th material class. Given the set of candidate solutions $w' = \{w^{(1)}, \dots, w^{(k)}\}$, we then rely on Equation 4 to pick the best candidate from w' .

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 4. 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 (Equation 1). 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 5 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 solution from w' that minimizes Equation 4.

6. 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 [Deb98] using Mitsuba [Jak10]; as noted in section 3, we will directly use this rendered image as a representation of the reflectance map. We use the BRDFs in the MERL database [MPBM03a] 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 BRDF corresponding to any of the three color channels of the BRDF), and only use these 297 MERL BRDFs for reconstruction.

Reconstruction Results Figure 1 shows reconstructions of 4 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 natural lighting (i.e., *Uffizi Gallery*; different than the lighting condition under which the BRDF was reconstructed) and a directional light. 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 extended report [CBP19] for a greater in depth analysis and additional reconstructions on real-world captured reflectance maps.

7. 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.

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

References

- [BM15] BARRON J. T., MALIK J.: Shape, illumination, and reflectance from shading. *PAMI* (2015). 23
- [CBP19] COOPER V. L., BIERON J. C., PEERS P.: Estimating homogeneous data-driven brdf parameters from a reflectance map under known natural lighting. *CoRR abs/1906.04777* (2019). 25
- [DCP*14] DONG Y., CHEN G., PEERS P., ZHANG J., TONG X.: Appearance-from-motion: Recovering spatially varying surface reflectance under unknown lighting. *ACM Trans. Graph.* 33, 6 (2014), 193:1–193:12. 23
- [Deb98] DEBEVEC P.: Light probe gallery. <http://www.pauldebevec.com/Probes/>, 1998. 25
- [DRS08] DORSEY J., RUSHMEIER H., SILLION F.: *Digital Modeling of Material Appearance*. Morgan Kaufmann Publishers Inc., 2008. 23
- [Jak10] JAKOB W.: Mitsuba: Physically based renderer. <https://www.mitsuba-renderer.org>, 2010. 25
- [LDPT17] LI X., DONG Y., PEERS P., TONG X.: Modeling surface appearance from a single photograph using self-augmented convolutional neural networks. *ACM Trans. Graph.* 36, 4 (July 2017), 45:1–45:11. 23
- [LN16] LOMBARDI S., NISHINO K.: Reflectance and illumination recovery in the wild. *PAMI* 38, 1 (2016), 129–141. 23
- [LSC18] LI Z., SUNKAVALLI K., CHANDRAKER M. K.: Materials for masses: Svbrdf acquisition with a single mobile phone image. In *ECCV* (2018). 23
- [LXR*18] LI Z., XU Z., RAMAMOORTHY R., SUNKAVALLI K., CHANDRAKER M.: Learning to reconstruct shape and spatially-varying reflectance from a single image. *ACM Trans. Graph.* 37, 6 (Dec. 2018). 23
- [MPBM03a] MATUSIK W., PFISTER H., BRAND M., McMILLAN L.: A data-driven reflectance model. *ACM Trans. Graph.* 22, 3 (July 2003), 759–769. 23, 24, 25
- [MPBM03b] MATUSIK W., PFISTER H., BRAND M., McMILLAN L.: Efficient isotropic BRDF measurement. In *Rendering Techniques* (2003), pp. 241–248. 24
- [NJR15] NIELSEN J. B., JENSEN H. W., RAMAMOORTHY R.: On optimal, minimal brdf sampling for reflectance acquisition. *ACM Trans. Graph.* 34, 6 (Oct. 2015). 23, 24
- [ON16] OXHOLM G., NISHINO K.: Shape and reflectance estimation in the wild. *PAMI* 38, 2 (Feb. 2016), 376–389. 23
- [PCDS12] PALMA G., CALLIERI M., DELLEPIANE M., SCOPIGNO R.: A statistical method for svbrdf approximation from video sequences in general lighting conditions. *Comput. Graph. Forum* 31, 4 (2012), 1491–1500. 23
- [RH01] RAMAMOORTHY R., HANRAHAN P.: A signal-processing framework for inverse rendering. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques* (2001), SIGGRAPH '01, pp. 117–128. 23
- [RRF*16] REMATAS K., RITSCHEL T., FRITZ M., GAVVES E., TUYTELAARS T.: Deep reflectance maps. In *CVPR* (2016). 23
- [RVZ08] ROMEIRO F., VASILYEV Y., ZICKLER T.: Passive reflectometry. In *ECCV* (2008), pp. 859–872. 23
- [RZ10] ROMEIRO F., ZICKLER T.: Blind reflectometry. In *ECCV* (2010), pp. 45–58. 23
- [WK15] WEINMANN M., KLEIN R.: Advances in geometry and reflectance acquisition. In *ACM SIGGRAPH Asia, Course Notes* (2015). 23
- [XDPT16] XIA R., DONG Y., PEERS P., TONG X.: Recovering shape and spatially-varying surface reflectance under unknown illumination. *ACM Trans. Graph.* 35, 6 (December 2016). 23
- [XNY*16] XU Z., NIELSEN J. B., YU J., JENSEN H. W., RAMAMOORTHY R.: Minimal brdf sampling for two-shot near-field reflectance acquisition. *ACM Trans. Graph.* 35, 6 (Nov. 2016), 188:1–188:12. 23, 24

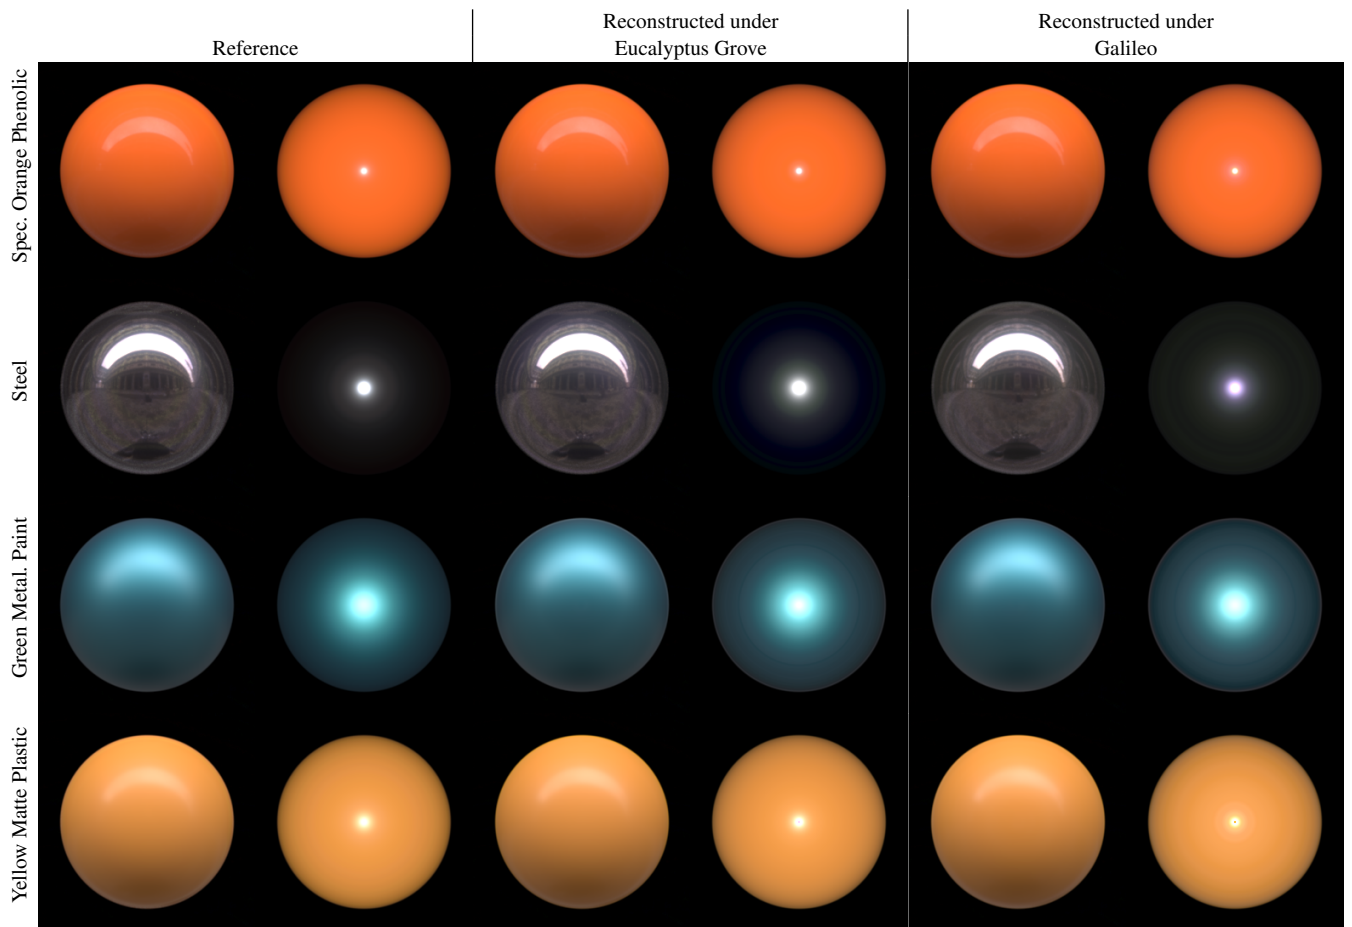


Figure 1: Data-driven BRDF reconstructions from a reflectance map under the Eucalyptus Grove and the Galileo’s Tomb light probe. We visualize the reference and reconstructed BRDFs under the Uffizi Gallery light probe and a directional light.

[YLD*18] YE W., LI X., DONG Y., PEERS P., TONG X.: Single photograph surface appearance modeling with self-augmented CNNs and inexact supervision. *Computer Graphics Forum* 37, 7 (Oct 2018). 23

[ZCD*16] ZHOU Z., CHEN G., DONG Y., WIPF D., YU Y., SNYDER J., TONG X.: Sparse-as-possible svbrdf acquisition. *ACM Trans. Graph.* 35 (November 2016). 23