

Rational BRDF: Supplementary Results

Romain Pacanowski, Oliver Salazar Celis, Christophe Schlick, Xavier Granier, Pierre Poulin and Annie Cuyt

1 MORE ON BRDF APPROXIMATION

Parametrization: $(\theta_h, \theta_d, \phi_d)$ vs. (θ_h, θ_d)

The most comprehensive and accurate database of measured isotropic BRDF, known as the *MERL-MIT BRDF database*, combines over 1 billion individual BRDF measurements generated by Matusik et al. [?]. The measured BRDF data are available with a $90 \times 90 \times 180$ angular sampling in $(\theta_h, \theta_d, \phi_d)$, which represents a storage amount of 33 MB per material.

As noted by Romeiro et al. [?], projecting the isotropic materials of the MERL-MIT database on the (θ_h, θ_d) subspace generates relatively few visual artefacts compared to the ground-truth image obtained by using the whole 3D $(\theta_h, \theta_d, \phi_d)$ parametrization. This projection naturally generates 90×90 clusters, each containing theoretically 180 reflectance values. In practice, many clusters contain less than 180 values, as many $(\theta_h, \theta_d, \phi_d)$ triplets generate viewing and/or lighting directions that are located beneath the tangent plane.

We have analyzed the result of this clustering process by computing mean and variance values of each individual cluster, presented in Figure 1 for our four selected materials: beige-fabric (almost perfect lambertian reflection), blue-metallic-paint (glossy reflection with strong chromatic behavior), nickel (specular reflection), chrome (almost perfect mirror reflection).

For diffuse and glossy materials, the variance in each cluster is always very low compared to the absolute magnitude of the corresponding reflectance, whereas for specular and almost-mirror materials, the variance increases. This can be explained by mainly two reasons: first, the angular resolution may become insufficient to accurately sample the reflectance behavior of high-frequency materials, and second, these materials are much more prone to measurement noise. Note that even for low-frequency materials, acquisition noise can be observed at grazing angles, which naturally increases the variance of the corresponding clusters.

- R. Pacanowski is with CNRS - LP2N. IOGS. Bordeaux University.
- O. Salazar Celis and A. Cuyt are with the Department of Mathematics and Computer Science, University of Antwerpen.
- C. Schlick and X. Granier are with the LaBRI - INRIA Bordeaux University. E-mail: [schlick|granier]@labri.fr
- P. Poulin is with the Department I.R.O., Université de Montréal. E-mail: poulin@iro.umontreal.ca

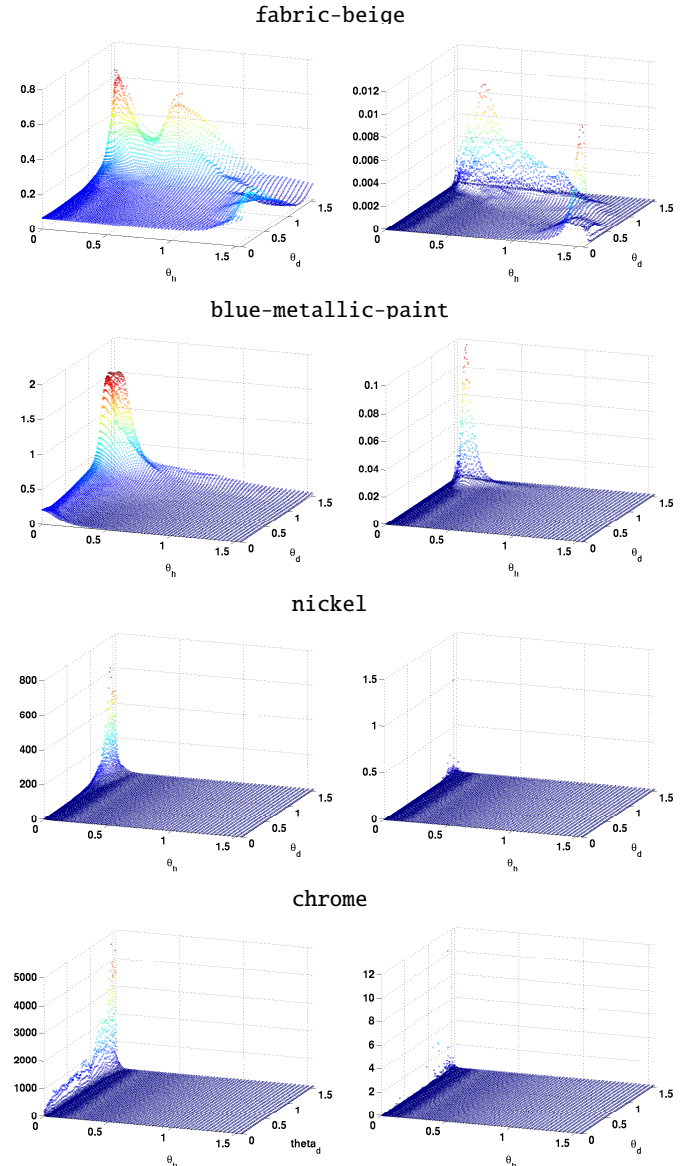


Fig. 1. (left) Mean and (right) normalized variance values obtained when projecting BRDF data from the original 3D $(\theta_h, \theta_d, \phi_d)$ space into the 2D (θ_h, θ_d) subspace.

In such cases, filtering may be employed to reduce variance before starting the fitting process. However, to avoid any statistical bias, we have only performed data filtering *within* each individual cluster, and not among neighboring clusters. After having tested several averaging procedures to find a representative sample for each cluster (arithmetic, geometric

or harmonic mean, root mean square, etc.), we have found that taking the median reflectance value appears to be the most robust choice, as it naturally removes outliers within the cluster.

For glossy and specular materials, most interesting phenomena are located on the isolines $\theta_n = 0$ and $\theta_d = 0$ as shown in Figure 1. These two curves are shown in Figure 2, where the data coming from all three color channels are presented.

As noted in the main paper, we have also observed that the filtering operator involved by the 2D projection (using either the mean or the median values of the cluster) sometimes even improves the visual aspect of the material, by removing some acquisition artefacts. As can be seen in Figure 3, when rendering some materials from the MERL-MIT database with the original 3D data, lens-flare-like visual artefacts can be observed in the resulting images. This phenomenon remains present, even when removing the noisier data that are located at grazing angles ($> 80^\circ$), as recommended by several researchers whom have worked on the MERL-MIT database in the past. When rendering the same scene with the projected 2D data, these artefacts are totally removed.

Additional Visual Comparison

Figures 4 and 5 present additional comparisons with our four selected MERL-MIT materials, by showing the visual difference (presented as a difference image in the *Lab* color space) between similar renderings obtained with either the original 3D data, the projected 2D data, or the rational BRDF fit. Figure 4 only uses one directional light source colinear to the view direction, while Figure 5 uses 1024 directional light sources sampled by importance from the surrounding environment map.

2 MORE ON CDF APPROXIMATION

A Brief Primer on Monte-Carlo Rendering

Over the years, Monte-Carlo-based techniques have become the standard approach to generate realistic images of arbitrary 3D scenes in the context of global illumination. As the convergence speed of such techniques is typically proportional to the square root of the number of samples, it is critical to exploit efficient variance reduction techniques, the most common one being *Importance Sampling* (IS). The principle of IS used in Monte-Carlo rendering is to perform a stochastic evaluation of the rendering equation:

$$L(\mathbf{v}) = \int \rho(\mathbf{v}, \mathbf{l}) (\mathbf{n} \cdot \mathbf{l}) L(\mathbf{l}) d\mathbf{l} \quad (1)$$

by defining an *estimator* that averages the contribution of K random light directions \mathbf{l}_k selected according to a conditional *Probability Density Function* (PDF) $\text{PDF}(\mathbf{l}_k | \mathbf{v}) = \text{PDF}_{\mathbf{v}}(\mathbf{l}_k)$:

$$L(\mathbf{v}) \approx \frac{1}{K} \sum_{k=1}^K \frac{1}{\text{PDF}_{\mathbf{v}}(\mathbf{l}_k)} \rho(\mathbf{v}, \mathbf{l}_k) (\mathbf{n} \cdot \mathbf{l}_k) L(\mathbf{l}_k). \quad (2)$$

When there is no a priori knowledge about the incident lighting, the optimal choice for the PDF is to be proportional to the BRDF multiplied by $(\mathbf{n} \cdot \mathbf{l})$ (i.e. the cosine of θ_l):

$$\text{PDF}_{\mathbf{v}}(\mathbf{l}) \propto \rho(\mathbf{v}, \mathbf{l}) (\mathbf{n} \cdot \mathbf{l}).$$

As \mathbf{l} is defined by its spherical coordinates, the PDF $\text{PDF}_{\mathbf{v}}(\mathbf{l})$ is actually a bivariate function $\text{PDF}_{\mathbf{v}}(\theta_l, \phi_l)$ which can be decomposed as a product of a 1D PDF $\text{PDF}_{\mathbf{v}}(\theta_l)$ with a 1D marginal PDF $\text{PDF}_{\mathbf{v}}(\phi_l | \theta_l)$.

Each individual sample of the IS process for Monte-Carlo rendering consists in generating a pair of uniform random numbers $(\mu, \tau) \in [0, 1)$ and computing the corresponding vector $\mathbf{l}(\theta_l, \phi_l)$ by successively inverting the two *Cumulative Distribution Functions* (CDF) corresponding to the pair of previous PDFs:

$$\theta_l = \text{CDF}_{\mathbf{v}}^{-1}(\mu) \quad \text{and} \quad \phi_l = \text{CDF}_{\mathbf{v}}^{-1}(\tau | \theta_l).$$

Direct Use of CDF⁻¹

As recalled in the main paper, two different approaches have been proposed in the past to obtain efficient evaluation of these inverse CDFs: either use a specific family of CDFs which allows a closed-form expression for its inverse function, or perform an on-the-fly numerical inversion from a tabulated version of the CDF. We propose here an innovative technique intended to combine the strengths of both approaches. The idea is to *directly* define a closed-form expression for the inverse CDF, without preliminary analytical formulation of neither the CDF nor the PDF. This offers several interesting features:

- With a well-chosen representation for the inverse CDF, the generation of each sample light direction required by the Monte-Carlo estimator can be done very efficiently, whatever the complexity of the corresponding BRDF.
- Both the CDF and the inverse CDF are monotonic functions, whereas the PDF may present many oscillations. Thus the inverse CDF is usually much more accurately approximated by a numerical process, compared to the corresponding PDF.
- When the inverse CDF is derivable, there is no need neither to store nor to compute the PDF, as one can derive a Monte-Carlo estimator entirely based on the inverse CDF (see below).

About Feature (a): Starting from a measured BRDF scaled by $\mathbf{n} \cdot \mathbf{l}$, we create a tabulated version for each of the inverse CDFs: a 1D table for $\text{CDF}_{\mathbf{v}}^{-1}(\mu)$ and a 2D table for $\text{CDF}_{\mathbf{v}}^{-1}(\tau | \theta_l)$. Then, as we did for BRDFs, we use the fitting technique, described in the main paper, to get a bivariate RF $r_{m,n}(\mu)$ that approximates $\text{CDF}_{\mathbf{v}}^{-1}(\mu)$ and a trivariate RF $r_{m,n}(\tau | \theta_l)$ that approximates $\text{CDF}_{\mathbf{v}}^{-1}(\tau | \theta_l)$.

As a result, generating each random vector \mathbf{l}_k required by the Monte-Carlo estimator can be done very efficiently as it involves only the computation of a pair of low-degree rational functions. Moreover, this is always done in constant time,

which is not the case with approaches based on numerical inversion of tabulated CDFs.

About Feature (b): In the main paper, we have presented two sample plots, respectively of $\text{PDF}_{\mathbf{v}}(\phi_l | \theta_l)$ and $\text{CDF}_{\mathbf{v}}^{-1}(\tau | \theta_l)$, for our four selected MERL-MIT materials at $\theta_v = \pi/4$. Figures 6 to 9 present a more complete version of these plots, showing the variation when θ_v goes from 0 to $\pi/2$. As can be observed, the shapes of the inverse CDF are *always much simpler* than of their corresponding PDF. This strongly supports our choice to fit inverse CDFs instead of PDFs. As already mentioned, the steep variations present in most plots can be accurately captured by RFs.

About Feature (c): Starting from the inverse CDF, the corresponding PDF can be easily obtained using the following equations (as demonstrated in the main paper):

$$\frac{\partial \text{CDF}_{\mathbf{v}}^{-1}}{\partial \mu}(\mu) = \frac{1}{\text{PDF}_{\mathbf{v}}(\theta_l)} ; \frac{\partial \text{CDF}_{\mathbf{v}}^{-1}}{\partial \tau}(\tau | \theta_l) = \frac{1}{\text{PDF}_{\mathbf{v}}(\phi_l | \theta_l) \sin \theta_l} \quad (3)$$

A nice side-effect of this latter equation is that we can also replace the PDF in the standard Monte-Carlo estimator (see Equation 2) of the rendering equation, to get a new estimator that is directly based on the inverse CDF:

$$L(\mathbf{v}) \approx \frac{1}{K} \sum_{k=1}^K \alpha_{\mathbf{v}}(\mu_k, \tau_k) \rho(\mathbf{v}, \mathbf{h}_k) (\mathbf{n} \cdot \mathbf{l}_k) L(\mathbf{l}_k) \quad (4)$$

with

$$\alpha_{\mathbf{v}}(\mu, \tau) = \partial_{\mu} \text{CDF}_{\mathbf{v}}^{-1}(\mu) \partial_{\tau} \text{CDF}_{\mathbf{v}}^{-1}(\tau | \theta_l) \sin \theta_l.$$

Moreover as inverse CDFs are defined as rational functions $R(x) = P(x)/Q(x)$, their partial derivatives can be trivially obtained by:

$$\frac{\partial R}{\partial x}(x) = \frac{\frac{\partial P}{\partial x}(x) - R(x) \frac{\partial Q}{\partial x}(x)}{Q(x)}.$$

This offers very compact storage, as we can analytically define the partial derivatives of polynomials $\frac{\partial P}{\partial x}(x)$ and $\frac{\partial Q}{\partial x}(x)$ on-the-fly. However, for faster processing, the coefficients of these new polynomials may also be precomputed and stored with an offline process.

ACKNOWLEDGMENTS




Fig. 2. Wavelength-dependent BRDFs of our four selected MERL-MIT materials, plotted for $\theta_h = 0$ (top) or $\theta_d = 0$ (bottom). Noise is typically located near grazing angles ($\theta_d \geq 80^\circ = 1.39$ radians) for all materials, but its amplitude depends on the specularity. Note the specific chromatic behavior of the blue-metallic-paint material. As the decrease is very steep for nickel and chrome, the plots on the bottom row have been clamped for better illustration.




Fig. 3. Lens-flare-like acquisition artefacts are particularly visible on different MERL-MIT materials when rendering with a small set of directional light sources.

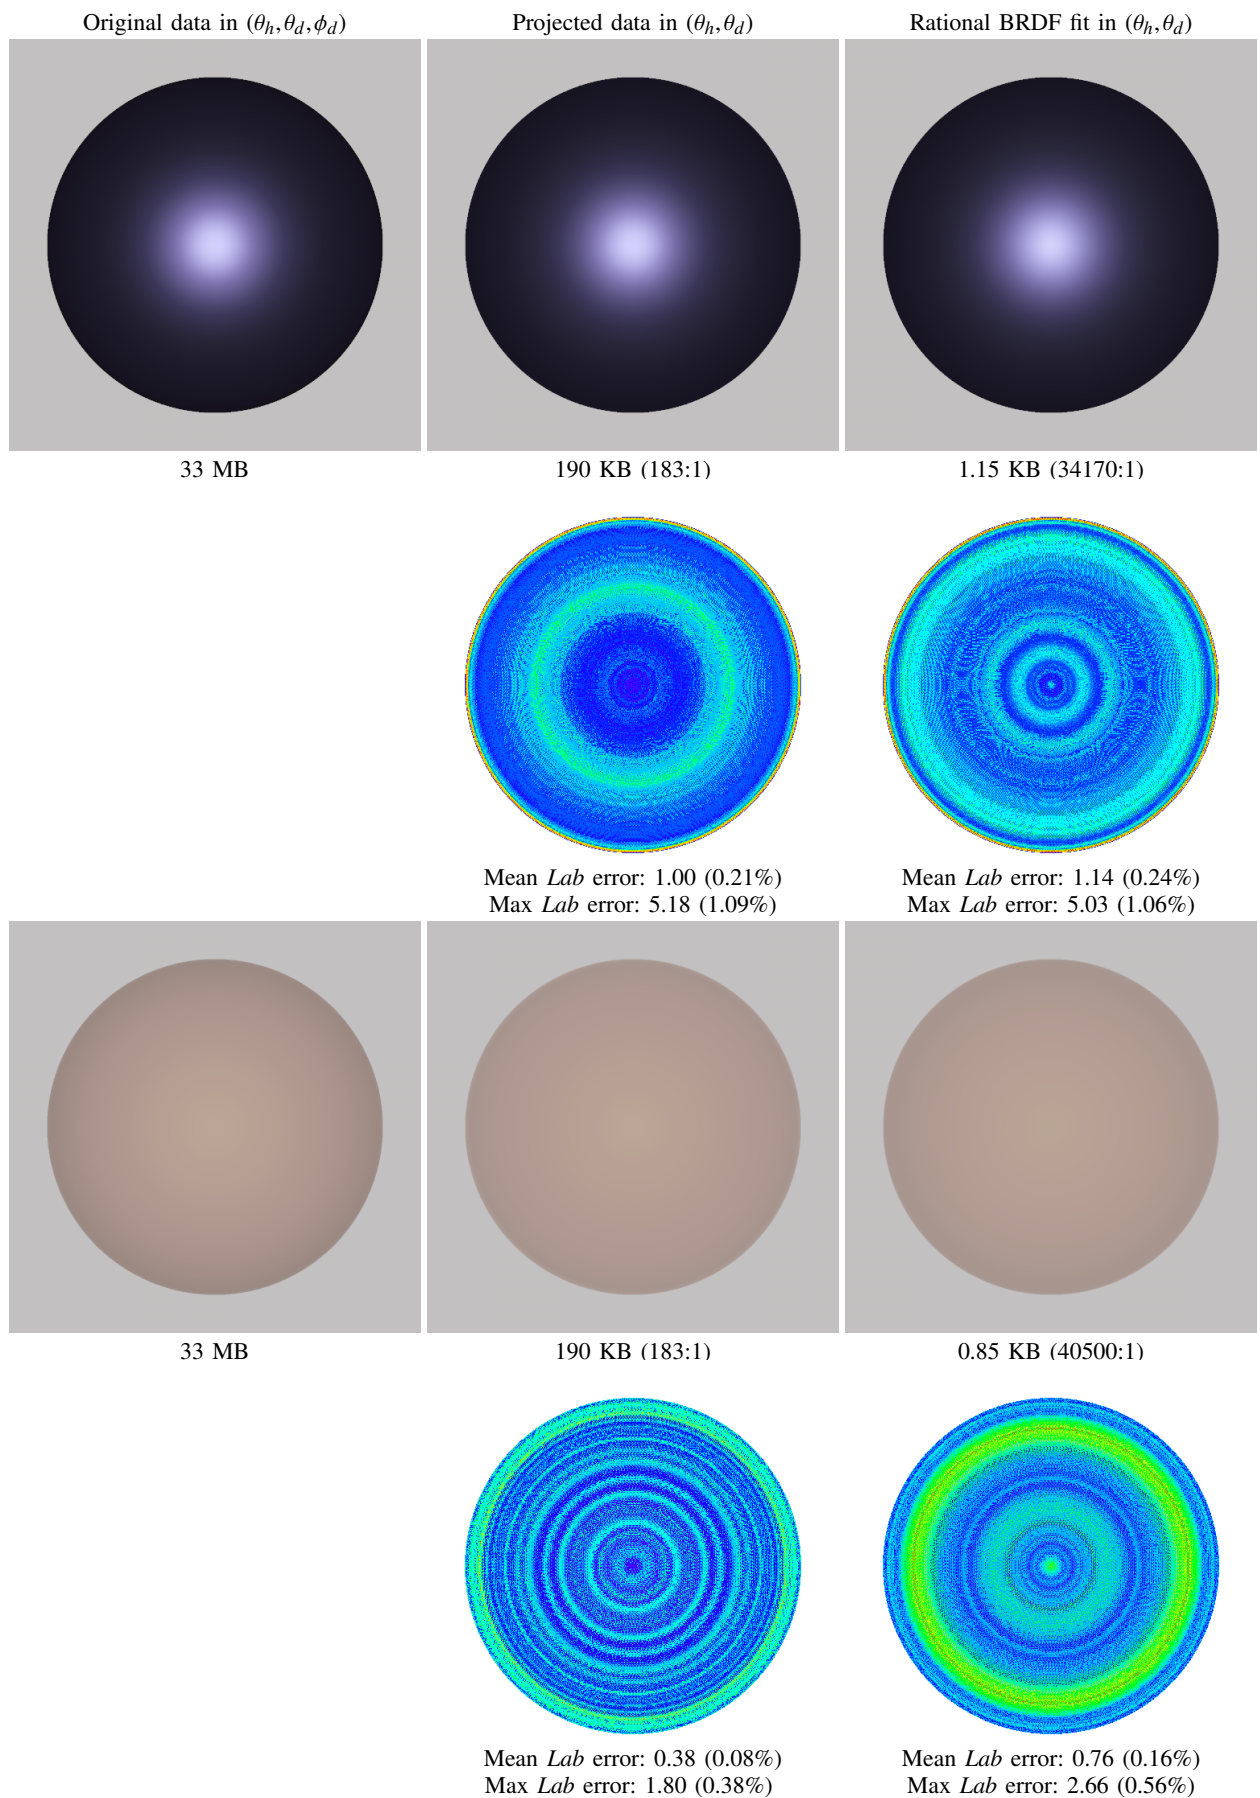


Fig. 4. blue-metallic-paint (top) and fabric-beige (bottom) rendered with one directional light colinear with the view direction.

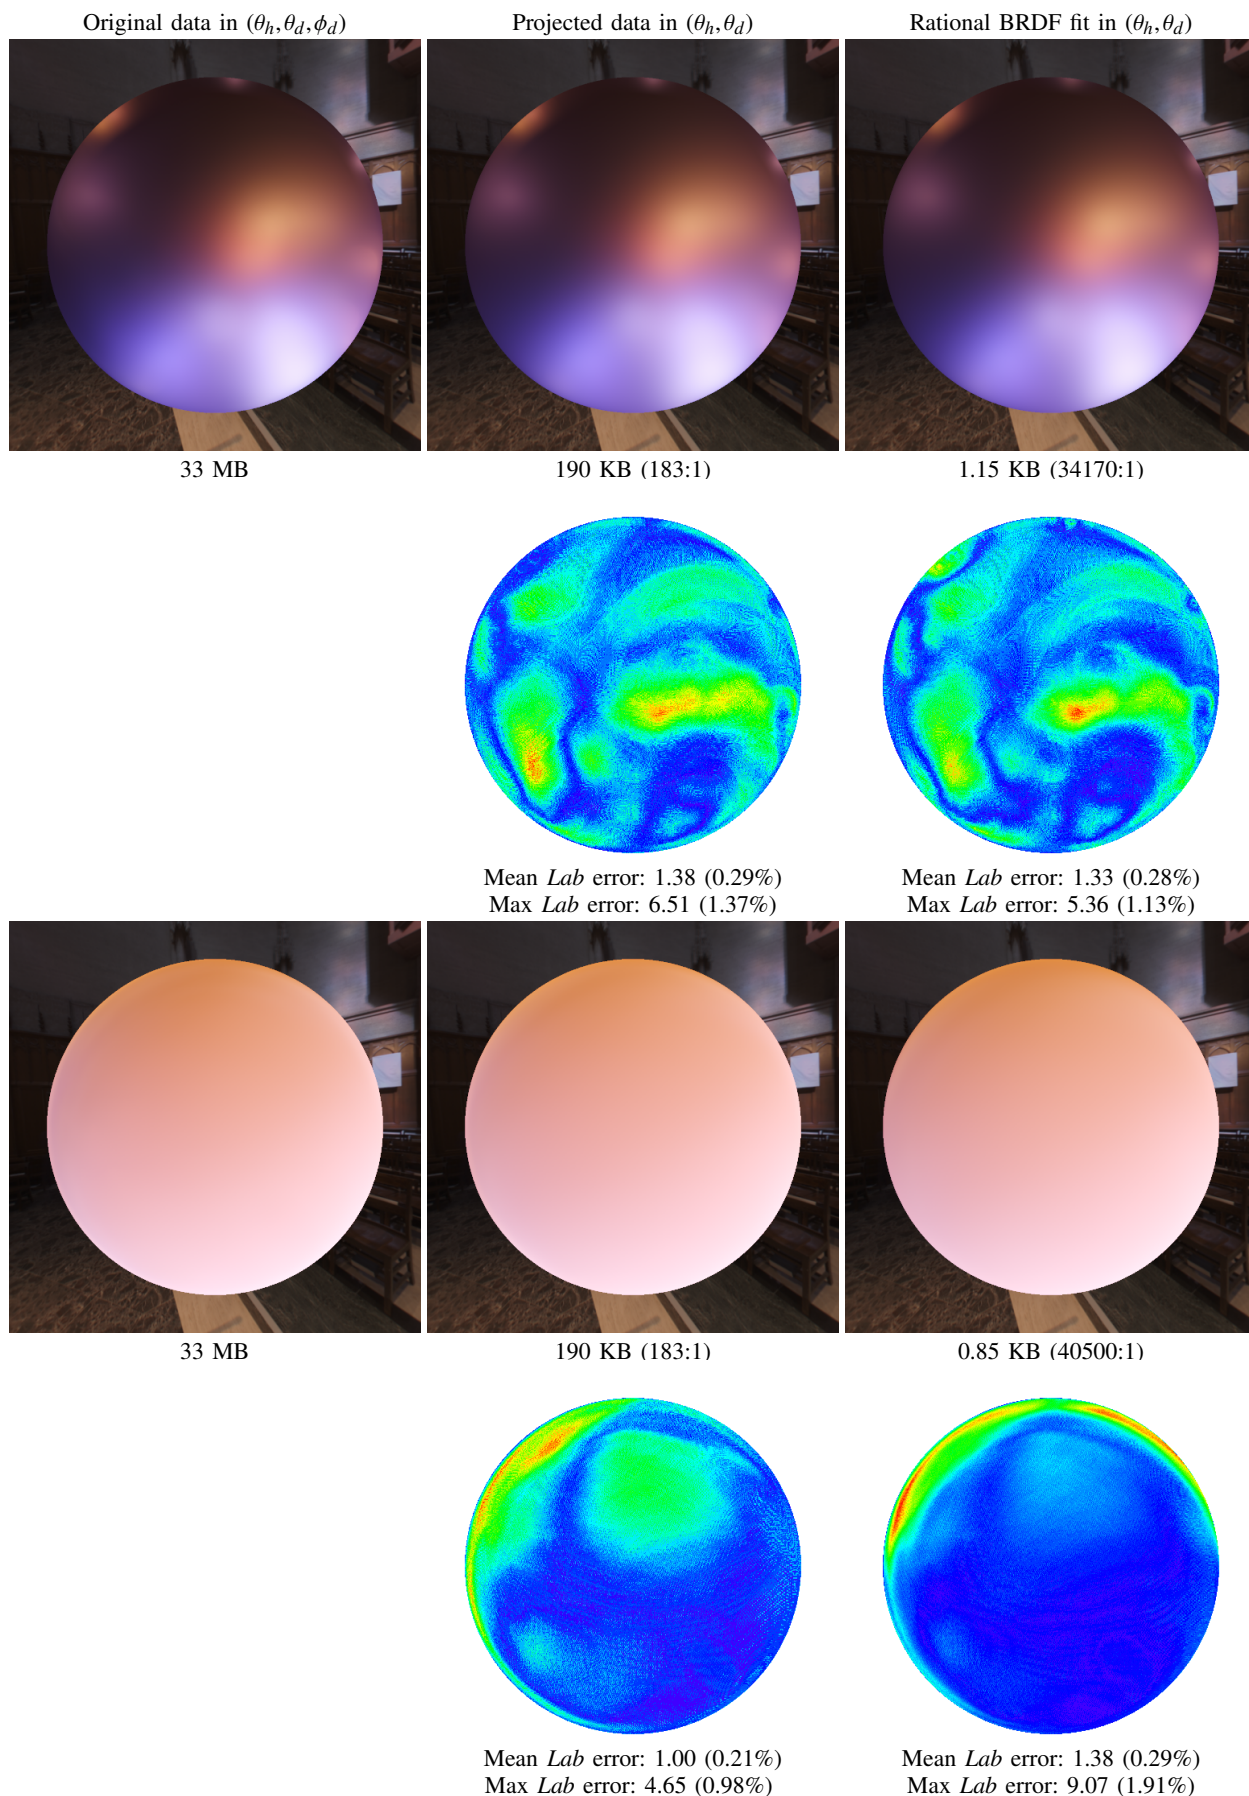


Fig. 5. blue-metallic-paint (top) and fabric-beige (bottom) rendered with 1024 lights selected from the surrounding environment map.

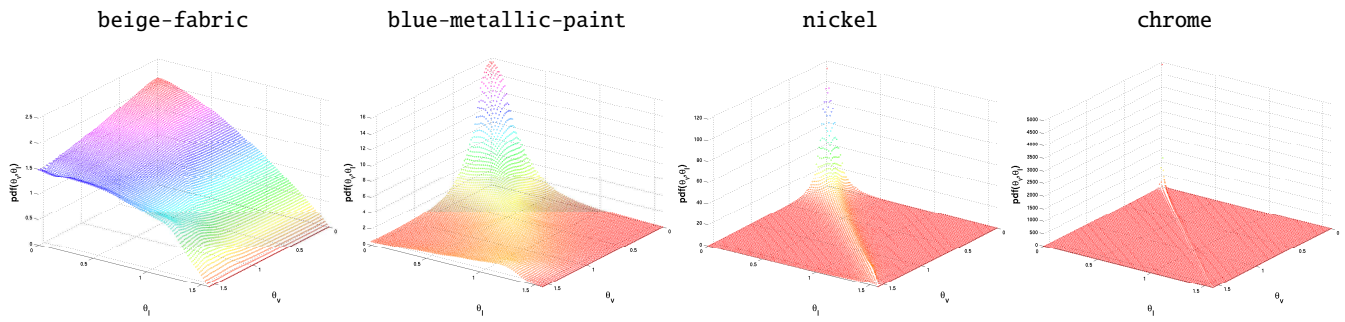


Fig. 6. Visualization of 2D $\text{PDF}_v(\theta_l)$ for different BRDFs.




Fig. 7. Visualization of the slice $\theta_v = 40^\circ$ for the 3D $\text{PDF}_v(\phi_l | \theta_l)$ of different BRDFs.

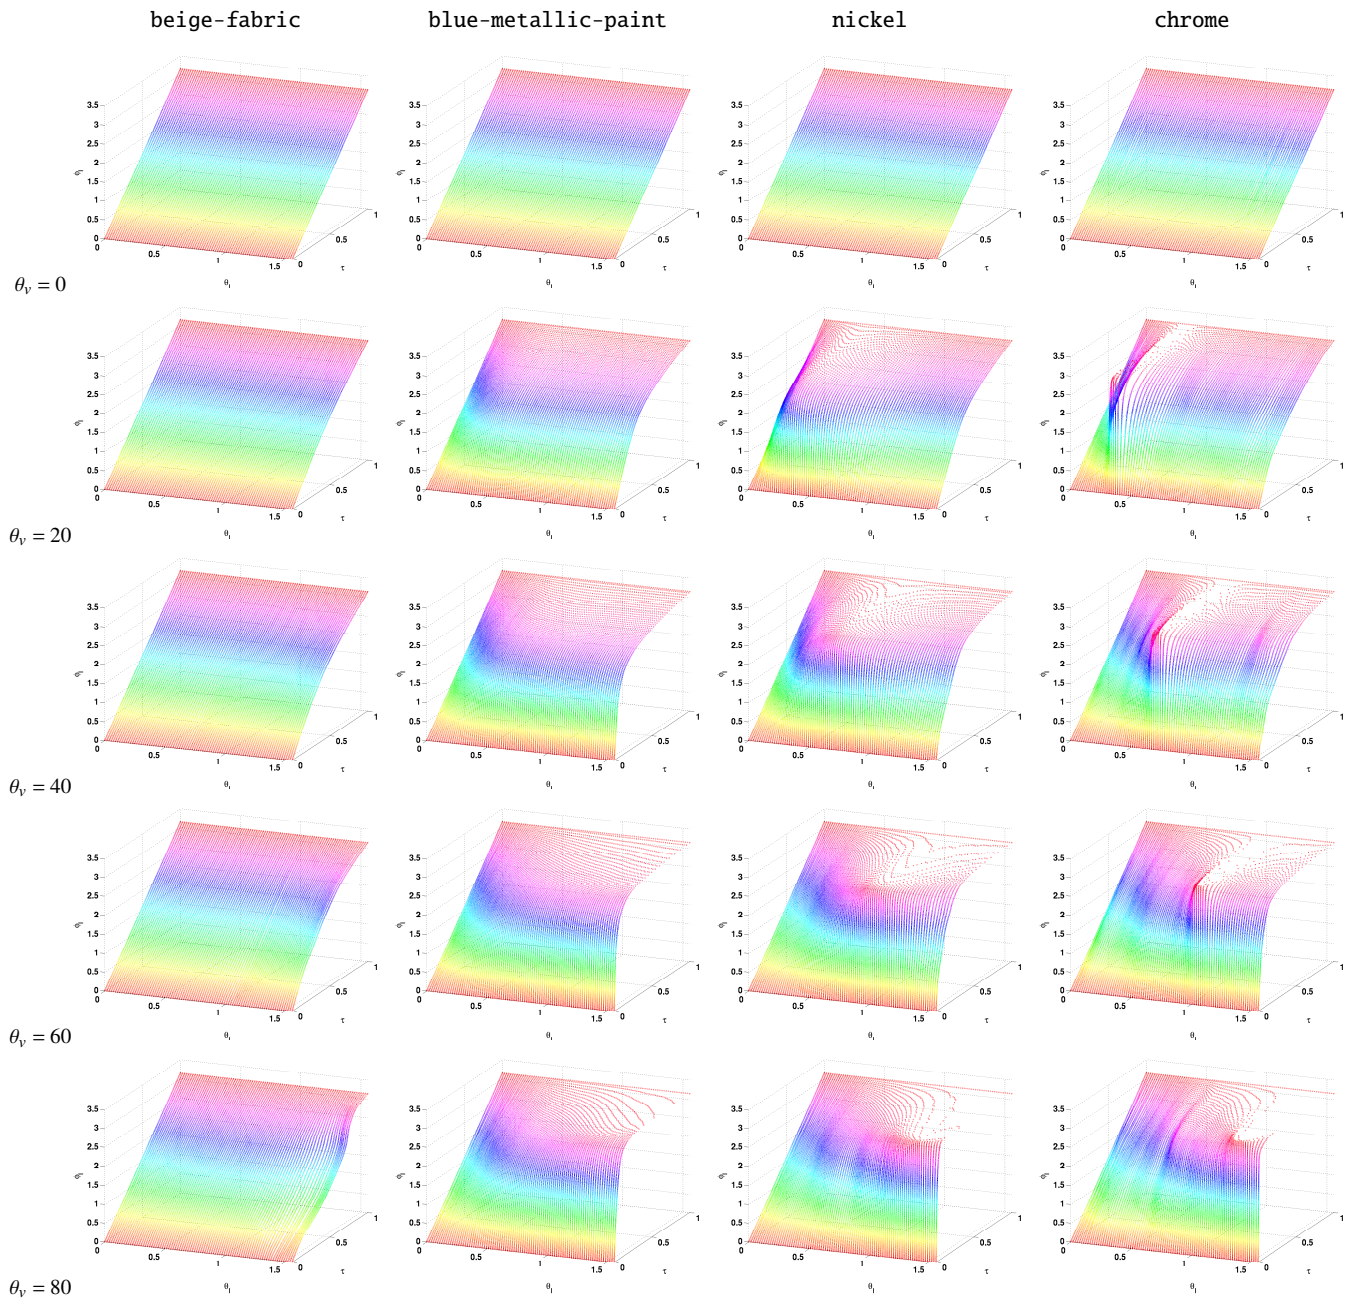


Fig. 8. Visualization of slices for different values of θ_v , of the 3D $CDF_v^{-1}(\phi_l | \theta_l, \tau)$ for different BRDFs.

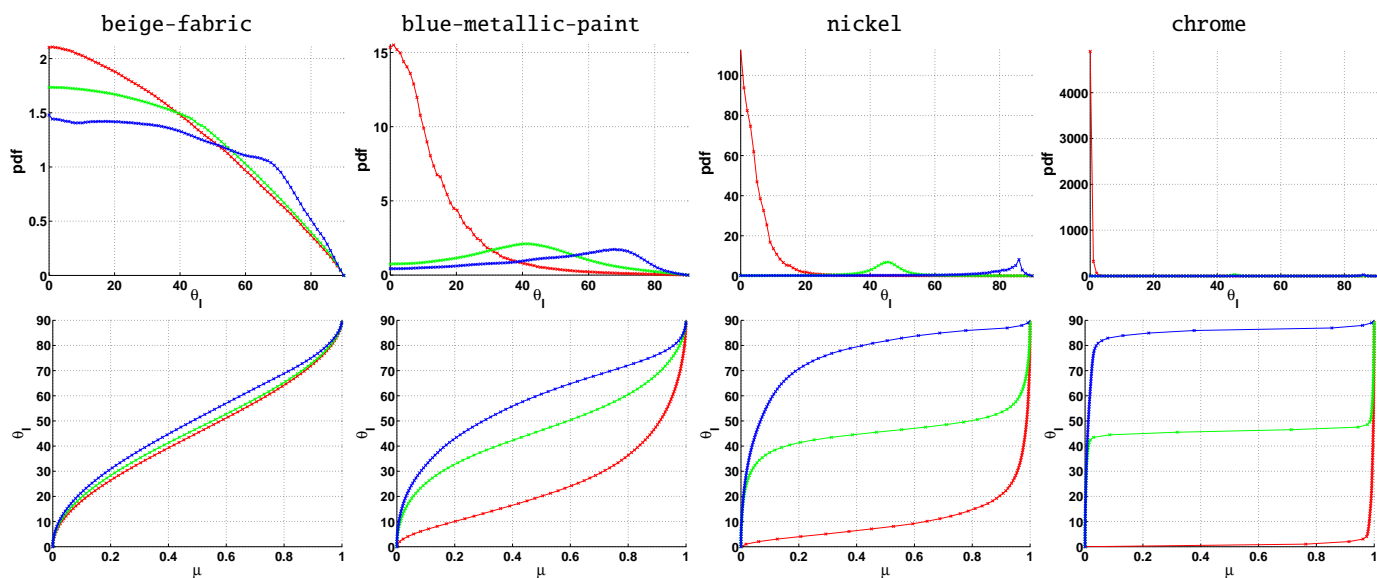


Fig. 9. (top row) $\text{PDF}_v(\theta_l)$ and (bottom row) $\text{CDF}_v^{-1}(\mu)$ for our four selected MERL-MIT materials, measured for (red curve) $\theta_v = 0^\circ$, (green curve) $\theta_v = 45^\circ$ and (blue curve) $\theta_v = 85^\circ$.