Rational BRDF

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

Abstract—Over the last two decades, much effort has been devoted to accurately measuring *Bidirectional Reflectance Distribution Functions* (BRDFs) of real-world materials and to use efficiently the resulting data for rendering. Because of their large size, it is difficult to use directly measured BRDFs for real-time applications, and fitting the most sophisticated analytical BRDF models is still a complex task. In this paper, we introduce **Rational BRDF**, a general-purpose and efficient representation for arbitrary BRDFs, based on *Rational Functions* (RFs). Using an adapted parametrization, we demonstrate how Rational BRDFs offer 1) a more compact and efficient representation using low-degree RFs, 2) an accurate fitting of measured materials with guaranteed control of the residual error, and 3) efficient importance sampling by applying the same fitting process to determine the inverse of the *Cumulative Distribution Function* (CDF) generated from the BRDF for use in Monte-Carlo rendering.

Index Terms—BRDF, fitting, importance sampling, Monte-Carlo rendering

✦

1 MOTIVATION AND PREVIOUS WORK

T HE BRDF is the keystone of the rendering equation [1], and thus of any lighting simulation process, since it describes the light reflection at a surface location. Consequently, building a comprehensive BRDF representation suitable to explain, describe, and/or simulate all complex phenomena involved in this process remains an active topic both in physics and computer graphics. The large number of scientific publications dealing with BRDF can mainly be divided into two families, depending on whether they propose an analytical formulation, or rather a numerical process to approximate measured BDRF data.

1.1 Closed-Form versus Numerical Representations

Papers in the first family present various models intended to represent some observed phenomena using either an empirical or a theoretical framework. These papers usually propose a closed-form formulation for the BRDF, parametrized in a manner that allows for the approximation of a limited set of real-world reflectance behaviors. These models range from efficient ad hoc formulations such as Phong's model [2], to sophisticated ones that include the

complex effects of wave optics [3], [4], [5]. The most common group of models [6], [7], [8], [9], [10], [11] exploit the *microfacet theory* [12].

Papers in the second family seek instead an efficient approach for representing the measured data using a set of basis functions, by means of standard linear decomposition techniques. Of these, some focus on (hemi)spherical basis functions, such as spherical harmonics (SHs) [13], [14], Zernike polynomials [15], spherical wavelets [16], or spherical radial basis functions [17]. Others apply dimension-reduction techniques such as homomorphic factorization [18], singular value decomposition [19], [20], or nonnegative factorization [21], [22]. Their main limitation resides in the fact that the number of coefficients required to get an accurate result grows quickly with the directional frequency of the BRDF (*quadratically*, as shown by Mahajan et al. in [23]), and thus becomes intractable for highly specular BRDFs.

Measured data may also be projected onto the parametric space of analytical BRDF models. In fact, some models are specifically designed for numerical fitting [7], [24], [25]. Compared to linear decomposition, fitting can easily handle materials with sharp specularity. However, this typically involves nonlinear minimization techniques (e.g., the Levenberg-Marquardt algorithm), which are quite slow, and reaching a global minimum of the function to be optimized is in general not guaranteed. Moreover, as analytical models have usually been designed for a specific class of materials, they cannot accurately fit arbitrary BRDFs [26]. Using multiple lobes or combining different models may improve accuracy, but due to the many local minima, the optimization quickly becomes impractical [21], [26].

1.2 BRDF Importance Sampling (IS)

Accurate fitting of measured data only solves half of the problem. Plugging the fitted data into a rendering engine, while preserving accurate and efficient computation, is far from straightforward. Over the years, Monte-Carlo-based techniques (see [27] for a comprehensive survey) have become the standard approach for generating realistic

- *R. Pacanowski is at Université de Bordeaux with LP2N (CNRS: UMR5298), and Institut d'Optique Graduate School—F-33400 Talence, France. E-mail: romain.pacanowski@institutoptique.fr.*
- *O. Salazar Celis and A. Cuyt are with the Department of Mathematics and Computer Science, University of Antwerpen, Campus Middelheim, Building G, Middelheimlaan 1, B-2020 Antwerpen, Belgium. E-mail: {oliver.salazarcelis, annie.cuyt}@ua.ac.be.*
- *C. Schlick is at Université Bordeaux Segalen with LaBRI (CNRS: UMR5800), F-33400 Talence, France. E-mail: schlick@labri.fr.*
- *X. Granier is at INRIA Bordeaux Sud-Ouest, with LaBRI (CNRS: UMR5800—Université de Bordeaux) and LP2N (Institut d'Optique Graduate School: CNRS: UMR5298—Université de Bordeaux), F-33400 Talence, France. E-mail: xavier.granier@inria.fr.*
- *P. Poulin is with the Département I.R.O., Université de Montréal, C.P. 6128, succ. Centre-ville, Montréal, Québec H3C 3J7, Canada. E-mail: poulin@iro.umontreal.ca.*

images. As their convergence speed is proportional to the inverse square root of the number of samples and to the initial variance, it is critical to exploit efficient variance reduction techniques, the most common one being *Importance Sampling* according to the BRDF.

As detailed in Section 4.1, performing IS requires the inverse of the BRDF's *Cumulative Distribution Function* (CDF). There are basically two approaches to compute the inverse of the CDF. The first is to fit the measured data with an analytical BRDF model [7], [8], [9], [10], [24], [28] that offers a readily invertible CDF. In addition to the previously mentioned weaknesses of nonlinear fitting, all these approaches (except for [28]) ignore the cosine factor that scales the BRDF according to the incident light direction, reducing the efficiency of IS for grazing angles of light. The second approach consists of tabulating the CDF into a sorted data structure (e.g., binary search tree) and computing the inverse function on the fly in this structure [21], [29], [30]. A major benefit of this approach is that the cosine scaling factor can be trivially included, greatly improving the efficiency of IS. Unfortunately, the storage cost is several orders of magnitude higher than with the first approach, and the iterative data retrieval process has a nonconstant computation cost.

In this paper, we introduce the following contributions:

- a general framework based on **Rational Functions (RFs)**, which efficiently represents BRDFs and CDFs without having to separate diffuse and specular components.
- an associated fitting technique that scales linearly with the desired accuracy and memory footprint. The involved optimization is that of a strictly convex function for which the global minimum is guaranteed to be reached, provided that a feasible solution exists.
- a new Monte-Carlo estimator for importance sampling rendering, which does not require storage of the Probability Density Function (PDF) when combined with our representation.

2 RATIONAL FUNCTIONS FRAMEWORK (RFF)

In approximation theory, Rational Functions are recognized for their greater expressivity compared to polynomials. They are preferred in several numerical approximation problems in scientific computing [31]. A Rational Function of a finite dimensional vector $\boldsymbol{x}$ of real variables x_i is

$$r_{n,m}(\boldsymbol{x}) = \frac{p_{n,m}(\boldsymbol{x})}{q_{n,m}(\boldsymbol{x})} = \frac{\sum_{j=0}^{n} p_j \, b_j(\boldsymbol{x})}{\sum_{k=0}^{m} q_k \, b_k(\boldsymbol{x})}, \quad (1)$$

where the $n + 1$ (resp. $m + 1$) coefficients of the numerator (resp. denominator) are represented by the real numbers p_j (resp. q_k), and where $b_j(\boldsymbol{x})$ and $b_k(\boldsymbol{x})$ are multivariate basis functions. We use multinomials in this paper, because they can be evaluated efficiently. We order them by increasing total degree, for example, in the bivariate case: $b_0 = 1, b_1 = x_1, b_2 = x_2, b_3 = x_1^2, b_4 = x_2^2, b_5 = x_1 \, x_2, b_6 = x_1^3, \ldots$ Therefore, for a given degree we favor adding first smoother basis functions (e.g., x_1^2) rather than more oscillating ones (e.g., $x_1 \, x_2$). Furthermore, both $p_{n,m}(\boldsymbol{x})/$

$q_{n,m}(\boldsymbol{x})$ and $\alpha p_{n,m}(\boldsymbol{x})/\alpha q_{n,m}(\boldsymbol{x})$ take the same function values for finite nonzero α, and the coefficients p_j and q_k need only be determined up to a multiplicative constant that can be used to normalize the representation of $r_{n,m}(\boldsymbol{x})$. Therefore, $r_{n,m}(\boldsymbol{x})$ has no more than $n + m + 1$ free coefficients.

RFs are ideal for approximating data that exhibit abrupt changes which are characteristic for specular lobes. An illustration of approximation of lobe-like functions using RFs is given in Fig. 2, where it can be observed that a low degree RF can easily represent abrupt variations followed by regions of almost constant values, whereas a polynomial with the same number of coefficients cannot. Such combinations of steep changes with flat regions are quite common in measured BRDF data and their corresponding CDF. However, in computer graphics, RFs have seldom been employed (except for the ad hoc BRDF model proposed by Schlick [8]).

Algorithm 1 presents an overview of our fitting procedure based on the work of Salazar Celis et al. [32]. A preprocessing step is required before the fitting itself, for data reparametrization and noise removal. Contrary to classical nonlinear methods, our Rational Functions Framework allows full control of the residual error by setting intervals F_i around each measured value, and the involved optimization is guaranteed to converge to a global minimum (cf. Section 2.1). Besides the measurements, the memory budget (i.e., the maximum number of coefficients), and the interval widths, the key part of the algorithm requires solving a quadratic programming problem $\mathcal{P}$ which is detailed in the following subsections.

Algorithm 1. Overview of the fitting procedure.

    $pdata = $ PREPROCESSED $data$
    $C_{max} = $ RF maximum number of coefficients
    $\{F_i\} = $ intervals for measured value $\boldsymbol{x}_i$ of $pdata$
    $b_0, \ldots, b_{C_{max}-1}$ are ordered basis functions
1: FITDATA ($pdata$, C_{max}, $\{F_i\}$)

2: **function** FITDATA ($data$, C_{max}, $\{F_i\}$)
3:     FITRATFUNC ($data$, C_{max}, $\{F_i\}$)
4:     **while** no solution **do**
5:         increase the width of all intervals $\{F_i\}$
6:         solution = FITRATFUNC ($data$, C_{max}, $\{F_i\}$)
7:     **end while**
8:     **return** solution
9: **end function**

10: **function** FITRATFUNC ($data$, C_{max}, $\{F_i\}$)
11:     $S = \emptyset$ $S_{best} = \emptyset$ $ccn = \infty$
12:     **for** $n = 0, \ldots, C_{max} - 1$ **do**
13:         **for** $l = n, \ldots, 0$ **do**
14:             $cn = $ CONDITIONNUMBER ($\mathbf{A}(l, n - l)$)
15:             $S = $ SOLVE ($\mathcal{P}(l, n - l)$, $data$, $\{F_i\}$)
16:             **if** $S \neq \emptyset$ and $cn < ccn$ **then**
17:                 $ccn = cn$
18:                 $S_{best} = S$
19:             **end if**
20:         **end for**
21:     **end for**

22:     **return** S_{best}
23: **end function**

2.1 Problem Statement for Fitting Data with RFs

Assume $s + 1$ measured values f_i, each of them located at a vector x_i. Let $F_i = [\underline{f_i}, \overline{f_i}]$ represent a real-valued interval placed around each f_i. We would like to find a Rational Function $r_{n,m}(x)$ that interpolates these intervals:

$$\forall i = 0, \ldots, s \qquad \underline{f_i} \le r_{n,m}(x_i) \le \overline{f_i}, \tag{2}$$

and this, for the smallest possible value of $n + m$, with $n + m \le s$ (usually $n + m \ll s$).

It can be shown [32] that a robust solution for $r_{n,m}(x_i)$ can be obtained by solving the following quadratic programming problem $\mathcal{P}(n,m)$:

$$\arg \min_{c \in \mathbb{R}^{n+m+2}} |c|_2$$
subject to
$$\mathbf{A}_{n,m}^{(j)} c - \delta |\mathbf{A}_{n,m}^{(j)}|_2 \ge 0, \qquad j = 1, \ldots, 2s+2$$
with $c = (p_0, \ldots, p_n, q_0, \ldots, q_m)^t$

where $\mathbf{A}_{n,m}^{(j)}$ denotes the jth row of matrix

$\mathbf{A}_{n,m} =$
$$\begin{pmatrix} b_0(x_0) & \ldots & b_n(x_0) & -\underline{f_0}\,b_0(x_0) & \ldots & -\underline{f_0}\,b_m(x_0) \\ \vdots & & \vdots & \vdots & & \vdots \\ b_0(x_s) & \ldots & b_n(x_s) & -\underline{f_s}\,b_0(x_s) & \ldots & -\underline{f_s}\,b_m(x_s) \\ -b_0(x_0) & \ldots & -b_n(x_0) & \overline{f_0}\,b_0(x_0) & \ldots & \overline{f_0}\,b_m(x_0) \\ \vdots & & \vdots & \vdots & & \vdots \\ -b_0(x_s) & \ldots & -b_n(x_s) & \overline{f_s}\,b_0(x_s) & \ldots & \overline{f_s}\,b_m(x_s) \end{pmatrix}$$

and where $|\cdot|_2$ denotes the euclidean norm. How to choose the real value δ is discussed in [32]. Furthermore, if $\mathcal{P}$ has a solution, it is **unique** and the returned function $r_{n,m}$ is **pole-free** at each measured value, i.e., $\forall i = 0, \ldots, s,\ q_{n,m}(x_i) > 0$. Whether a solution exists for a specific n and m depends on the relation between the width of the intervals F_i and the values n, m, and s. Algorithm 1 may find multiple solutions, but it always selects the most stable solution by choosing the one associated with the lowest condition number of matrix $\mathbf{A}_{n,m}$. Although theoretically possible, multiple solutions are rare for BRDF/CDF fitting. Finally, solving $\mathcal{P}$ can be done with any quadratic programming solver; in this paper, all RFs are obtained using the *qpas* routine of the freely available *QPC Matlab interface* [33].

In the next section, we introduce linear constraints that are easily added to $\mathcal{P}$ when fitting BRDFs or inverse CDFs.

2.2 Additional Constraints

2.2.1 Value Constraints
Since BRDFs are nonnegative functions, we build value constraints on the function as a whole into the intervals right from the start:

$$\forall i = 0, \ldots, s \quad \underline{f_i} \ge 0.$$

For 2D inverse CDFs (cf. Section 4), we make sure that

$$\forall i = 0, \ldots, s \quad 0 \le \underline{f_i} \le \overline{f_i} \le \pi/2$$

and similarly for 3D inverse CDFs

$$\forall i = 0, \ldots, s \quad 0 \le \underline{f_i} \le \overline{f_i} \le \pi.$$

Even though these types of value constraints are only valid at discrete positions, we found that this approach works satisfactorily.

2.2.2 Symmetry Constraints
Helmholtz reciprocity (or symmetry) is one of the main properties of physically based BRDFs and, by construction, inverse CDFs inherit it. Therefore, it is important to guarantee that the fitted RFs approximating BRDFs or inverse CDFs preserve this property. Thanks to the parametrization that we use for BRDFs (cf. Section 3.1), we do not need additional constraints since the half-angle parametrization already includes the symmetry property.

However, for inverse CDFs, we use the classical light-view parametrization (cf. Section 4.1), and therefore we enforce symmetry by grouping basis functions that should have the same coefficients. For example, assume we would like to fit data with a trivariate RF

$$r_{n,m}(x_1, x_2, x_3) = c_0 x_1^3 x_2^4 x_3 + c_1 x_2^3 x_1^4 x_3 \quad (c_0, c_1) \in \mathbb{R}^2$$

and enforce its symmetry for the variables x_1 and x_2, i.e., $r_{n,m}(x_1, x_2, x_3) = r_{n,m}(x_2, x_1, x_3)$. Instead of adding equality constraints between c_0 and c_1 in $\mathcal{P}$, we group the two basis functions into one:

$$r_{n,m}(x_1, x_2, x_3) = c_0 \left((x_1^3 x_2^4 + x_2^3 x_1^4)\, x_3 \right).$$

This reduces the size of matrix $\mathbf{A}_{n,m}$ and hence the size of the quadratic programming problem.

2.2.3 Monotonicity Constraints
Since inverse CDFs are positive monotonic functions, it is crucial to enforce the condition that the derivative of the fitted function is nonnegative. In other words, we want to ensure the monotonicity of the rational function $r_{n,m}(x) = p_{n,m}(x)/q_{n,m}(x)$ in each of the variables x_j:

$$\frac{\partial r_{n,m}}{\partial x_j}(x) \ge 0. \tag{3}$$

Assuming that $q_{n,m}(x) > 0$, Inequality (3) becomes

$$\frac{\partial p_{n,m}}{\partial x_j}(x) \ge \frac{p_{n,m}(x)}{q_{n,m}(x)} \frac{\partial q_{n,m}}{\partial x_j}(x).$$

By using (2), the previous inequality is achieved for all sample positions x_i if

$$\frac{\partial p_{n,m}}{\partial x_j}(x_i) \ge \frac{\partial q_{n,m}}{\partial x_j}(x_i)\,\underline{f_i} \quad \text{and} \quad \frac{\partial p_{n,m}}{\partial x_j}(x_i) \ge \overline{f_i}\,\frac{\partial q_{n,m}}{\partial x_j}(x_i). \tag{4}$$

This last constraint consists of two additional linear inequalities for the coefficients $p_0, \ldots, p_n$ and $q_0, \ldots, q_m$ per data point x_i. The effect of adding these constraints can be seen in Fig. 3, where we show two fitted rational functions, Fig. 3(left) without and Fig. 3(right) with the added discrete monotonicity constraints.

2.3 Solving $\mathcal{P}$ Efficiently
The speed of the resolution of $\mathcal{P}$ greatly depends on the number $s + 1$ of intervals to be interpolated. Furthermore,

the faster $\mathcal{P}$ is solved, the sooner the algorithm is going to find the optimal solution and return. To quickly find the couple (n, m), we propose an adaptive procedure that operates on a selected small subset of the whole data. The key observation here is that if the current tested $r_{n,m}(\boldsymbol{x})$ RF is not a solution for the selected subset, it cannot be for the whole data set.

Of the given $s + 1$ data intervals, a small number s_0 of intervals are uniformly selected and $r_{n_0,m_0}^{(0)}(\boldsymbol{x})$ is computed such that it satisfies (2) for these s_0 data; these s_0 intervals are called the *training data*. Then, it is checked how many of the original $s + 1$ interval interpolation conditions are automatically satisfied by $r_{n_0,m_0}^{(0)}(\boldsymbol{x})$ in addition to the s_0 imposed conditions. Usually, this is quite a large number. The $s + 1 - s_0$ intervals are called the *verification data*. Among the violated interval interpolation conditions, we select $s_1 - s_0$ ($s_1 - s_0 = 1$ in our implementation) additional data points to compute $r_{n_1,m_1}^{(1)}(\boldsymbol{x})$ that satisfies (2) for these s_1 data points. In other words, we update the set of training data. These $s_1 - s_0$ additional training data are selected according to where $r_{n_0,m_0}^{(0)}(\boldsymbol{x}_i)$ deviates most from the given intervals F_i. With $r_{n_1,m_1}^{(1)}(\boldsymbol{x})$, we then check the new $s + 1 - s_1$ verification data again. The procedure is repeated until the obtained rational model satisfies all verification data.

3 RATIONAL FUNCTIONS FOR BRDF

3.1 Parametrization

A BRDF ρ is a 4D nonnegative function defined on the euclidean product of two hemispheres and depending both on the lighting $\boldsymbol{l}$ and viewing $\boldsymbol{v}$ directions. The most common way to express $\boldsymbol{l}$ and $\boldsymbol{v}$ is to use the standard zenith and azimuthal angles (θ_l, ϕ_l) and (θ_v, ϕ_v): $\rho(\theta_l, \phi_l, \theta_v, \phi_v)$, often denoted by $\rho(\boldsymbol{l}, \boldsymbol{v})$.

Although most acquisition devices (e.g., gonioreflectometers) cover the acquisition space by uniformly sampling these coordinates, the computer graphics community prefers to use the BRDF parametrization introduced by Rusinkiewicz [34]: $\rho(\boldsymbol{l}, \boldsymbol{v}) = \rho(\boldsymbol{h}, \boldsymbol{d}) = \rho(\theta_h, \phi_h, \theta_d, \phi_d)$ where $\boldsymbol{h}$ is the half-vector used in the microfacet theory, and (θ_d, ϕ_d) are the spherical coordinates of vector $\boldsymbol{l}$ expressed in a rotated frame where $\boldsymbol{h}$ defines the north pole direction. This parametrization offers several advantages:

1. in the case of isotropic materials (which are invariant when rotating the surface around its normal vector), one can set $\phi_h = 0$ without any loss of generality,
2. the reciprocity of light propagation ensures that ϕ_d can be limited to the range $[0, \pi]$, again without any loss of generality,
3. as noted by Romeiro et al. [35], for a very common class of materials called *bilateral symmetric*, the domain of ϕ_d can even be reduced to the range $[0, \pi/2]$,
4. as several phenomena involved in light reflection are mostly decorrelated along the different parameter axes, it is relatively safe to factor the 4D function into a product (or a sum of products) of two 2D functions, as done in work on factorization of BRDFs and SVBRDFs [21], [36].

As the BRDF is intrinsically 3D/4D for isotropic/anisotropic materials, further dimension reduction cannot be achieved without generating some undesired folding of the domain. This can be seen as a standard aliasing process, where an infinite number of configurations of the 4D space are projected onto a single configuration in the lower dimensional space. However, some recent work has shown that for most isotropic materials, the BRDF can be projected onto a well-chosen 2D parametric space without severe visual degradation [35], [37]. Romeiro et al. [35] have observed that almost all materials of the MERL-MIT database [20] are visually well approximated by a projection onto (θ_h, θ_d) (see Fig. 7, top and middle rows). Incidentally, we have observed that the filtering operator involved when projecting measured data on (θ_h, θ_d) sometimes also results in more visually pleasing rendering results (see Fig. 4).

In this paper, we propose to go a step further and show that we can accurately approximate the whole 2D projection of the BRDF $\rho(\theta_h, \theta_d)$ by a single RF $r_{m,n}(\theta_h, \theta_d)$, defined as in (1), that we call a **Rational BRDF**

$$\rho(\boldsymbol{l}, \boldsymbol{v}) \approx \rho(\theta_h, \theta_d) \approx r_{n,m}(\theta_h, \theta_d). \qquad (5)$$

As illustrated in the next sections, a single low-complexity RF is able to approximate isotropic BRDF data well, without multiple lobe fittings like those required by the current standard techniques, even when several phenomena (diffuse reflection, forward or backward glossy or specular reflections, Fresnel effects, etc.) are simultaneously observed on a given material.

Although isotropic BRDF measurements have been intensively studied, much less work has been done on fitting measured anisotropic materials, basically because much less anisotropic data are publicly available. In Section 5, we present our adaptation of Rational BRDF to an anisotropic data model, and discuss the need for high quality measured 4D data sets.

3.2 Rational Approximation of BRDF

The most comprehensive and accurate database of measured isotropic BRDFs, known as the *MERL-MIT BRDF database*, combines over 1 billion individual BRDF measurements generated by Matusik et al. [20]. 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 (the database was last updated in 2006). We have tested our RF approximation technique on many different materials from the database, but for clarity of presentation, we focus here on four representatives of common BRDF families: `beige-fabric` (almost perfect Lambertian reflection), `blue-metallic-paint` (glossy reflection with strong chromatic behavior), `nickel` (specular reflection), and `chrome` (almost perfect mirror reflection), which are presented in Fig. 7(top row).

For each of these materials, we fit, as described in Section 2, a bivariate rational function $r_{n,m}(\theta_h, \theta_d)$ that satisfies (2). Of course, the behavior of the approximation is determined by the choice of the intervals $F_i = [\underline{f_i}, \overline{f_i}]$ containing $\rho(\theta_h^i, \theta_d^i)$, $(i = 0, \ldots, 90 \times 90 - 1)$, for each of the BRDF values. Ideally, the widths of these intervals are calibrated such that they respect the accuracy of each individual measurement. Since

Original data size: BRDF = 99 MB, Tabulated CDF+PDF ≃30 MB   *Our approach: BRDF = 1.67 KB, Inverse CDF = 0.600 KB*

Fig. 1. Monte-Carlo rendering with 2,048 samples/pixel for a scene with three measured BRDFs from the MERL-MIT database (`blue-metallic` on the dragon, `beige-fabric` on the floor, `nickel` on the sphere). Our approximation of the BRDFs and the inverse CDFs, based on Rational Functions, provides efficient importance sampling with a negligible memory footprint: with less than 1 KB of storage, our IS technique (right) offers equivalent quality (mean *Lab* difference is 0.77 and max 7.03 on low-dynamic range images) compared to the reference solution (left) obtained from tabulated data of ≃30 MB. These tabulated CDF+PDF data have been generated by resampling the BRDF in $(\theta_v, \theta_l, \phi_l)$ at $90 \times 90 \times 180$. Furthermore, the rendering time of our approach is 10 percent faster.

measurement error bounds are not explicitly available here, the interval widths are taken such that the renderings are visually satisfactory, while leaving a reasonable number $m + n$ of coefficients (cf. Fig. 6) in the approximation. As a general guideline, we control the interval widths with the following formula:

$$F_i = [\underline{f_i}, \overline{f_i}] = [\rho(\theta_h^i, \theta_d^i)(1 - \epsilon_i), \; \rho(\theta_h^i, \theta_d^i)\,(1 + \epsilon_i)],$$

where ϵ_i specifies the desired relative error. Fig. 8 presents the maximum relative errors obtained for BRDFs. However, the interval widths (and therefore the ϵ_i) are usually not set uniformly. Instead smaller interval widths near the hemisphere pole and more relaxed widths are chosen near grazing angles.

For all materials, we noted increased acquisition noise when approaching grazing angles ($\theta_h > 60°$ and/or $\theta_d > 60°$). In the case of highly specular materials, we also noted that the acquired BRDF values around the center of the lobe ($\theta_h = 0$) are particularly large and subject to noise as illustrated in Fig. 5 on `chrome` and `nickel` BRDFs. Small interval widths are thus chosen on data with low noise, leading to a very accurate fit in such a configuration (as also shown in Fig. 5). However, as confirmed by Figs. 1 and 7, even in the case of materials including very noisy data and therefore larger interval widths, the rendering obtained with the fitted function has a very low visual impact when compared to the rendering based on the original data.

Fig. 7(bottom row) presents the rendering obtained with the Rational BRDF approximation of our four selected materials. The environment map is approximated by using 1,024 directional light sources selected according to their power in the surrounding environment map. As can be observed, the visual error (measured in *Lab* color space) is approximately of the same magnitude as with the 2D projected data, but our Rational BRDFs require between 165 and 1,000 times less storage memory (in our current implementation, each coefficient p_j and q_k of the RF is stored as an 8-byte double-precision float). When compared to the original 3D data, the compression rate ranges from 28,700:1 to 174,000:1. Note that using 4-byte standard floats is likely to be sufficient in many cases. Finally, for configurations at grazing

angles and for rendering purposes only, we extend our RF by clamping light and view directions that exceed 75 degrees. As shown in Fig. 1, this does not introduce visual artifacts. For all our figures, we used the same clamping for the data and the RFs in order to illustrate the accuracy of the fitting.

The final approximation errors are summarized in the table of Fig. 8: we have computed for each material (column 1) the maximum relative error of the 2D data using the (θ_h, θ_d) parametrization (column 2), our 2D RF fitting against the 2D data (column 3) and against the original 3D data (column 4), and using a nonlinear fitting procedure with the modified Ashikhmin-Shirley BRDF model as described by Ngan et al. [26] (column 5). As expected, for every material and regardless of the representations, the fitting errors grow with the specularity of the material. As can be observed, the overall errors of the rational approximations and the Ngan fitting are quite similar. However, the interval interpolation process has the advantage of offering full control both on the location and the magnitude of the errors. We chose to compare with the Ashikhmin-Shirley model because, as stated by Ngan et al. [26]: "for single specular lobe, Cook-Torrance, Ashikhmin-Shirley, and the He models perform well for most of the 100 isotropic materials." Furthermore, the effect of adding one
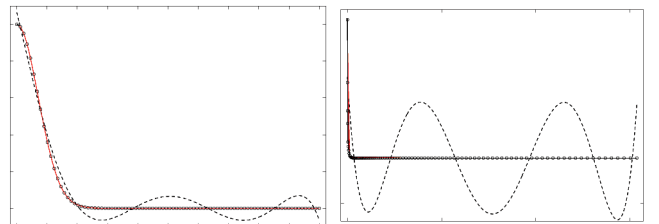

Fig. 2. Two low-degree univariate approximations with seven unknowns: a polynomial of degree 6 (dashed line) and a rational function (full red line) consisting of a polynomial of degree 1 and 5 in the numerator and denominator, respectively. The data are shown as dots. The RFs clearly follow the data much better and do not suffer from oscillations as polynomials do. Left. The rapidly decreasing function e^{-x^2}, sampled at 90 equidistant points. The maximum absolute error of the polynomial is 0.0689 while it is 0.00117 for the rational function. Right. The steeply decreasing lobe corresponding to the `chrome-steel` material from the MERL-MIT database. The maximum absolute error of the polynomial is 138 while it is only 2.61 for the rational function.
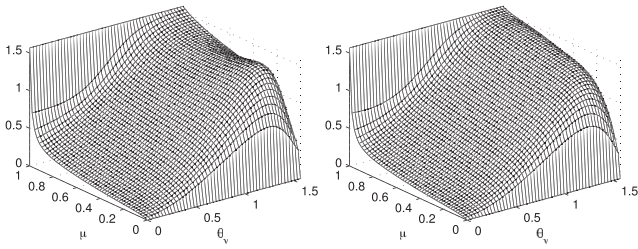
Fig. 3. Two low-degree bivariate rational approximations to $\mathrm{CDF}_v^{-1}(\mu)$ for the `nickel` MERL-MIT material. The approximation should be monotonically increasing in μ. Left. Rational interval interpolant obtained without Constraint (4), clearly not monotonically increasing in μ. Right. Rational interval interpolant obtained with Constraint (4), now monotonically increasing in μ.

lobe (hence fitting with a two-lobe BRDF model), as shown by Ngan et al. [26] (in their Section 4.2), indeed reduces the fitting error (by approximately 25 percent) for 26 of the 31 materials. However, according to these authors, the fitting process quickly becomes unstable with more than three lobes and the benefit of adding a lobe is only marginal. In conclusion, the errors presented in the table of Fig. 8 for the Ashikhmin-Shirley model were improved by at most 25 percent when adding a second lobe. Therefore, a model from a two-lobe fitting would still not be competitive (in terms of the error and the numerical stability of the technique) against our RF fitting technique.

Finally, the rightmost column (column 6) of the table of Fig. 8 shows the maximum relative error against the 3D data when using least squares bivariate polynomials with the **same** number of coefficients as the RFs. Except for the low frequency `beige-fabric` material, the errors are several orders of magnitude higher than the ones obtained with RFs. Furthermore, these errors show that approximating high-frequency signals with RFs is more efficient than using polynomials. As shown by Mahajan et al. [23], the required number of coefficients for polynomial functions (e.g., spherical harmonics) grows quadratically with the directional frequency of the BRDF. Note that it is not straightforward to fit directly spherical harmonics using the (θ_h, θ_d) because these angles do not span a spherical domain. Other approaches such as those of Sillion et al. [38] or Westin et al. [14] have used classical spherical parametrization but require between 80 and 133 coefficients for diffuse or glossy materials, and become impractical for high-frequency BRDF.
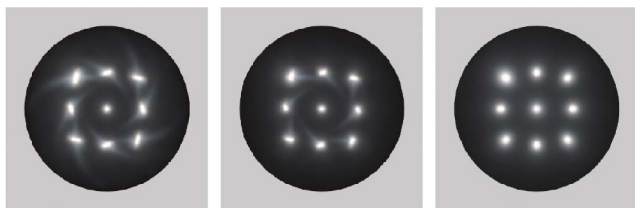

Fig. 4. Lens-flare-like acquisition artifacts of the `grease-covered-steel` MERL-MIT material illuminated by nine directional sources. Left. Original 3D data. Middle. 3D data without noisy grazing angles (i.e., θ_h or $\theta_d > 80°$). Right. 2D projection on (θ_h, θ_d) mostly removes all artifacts.
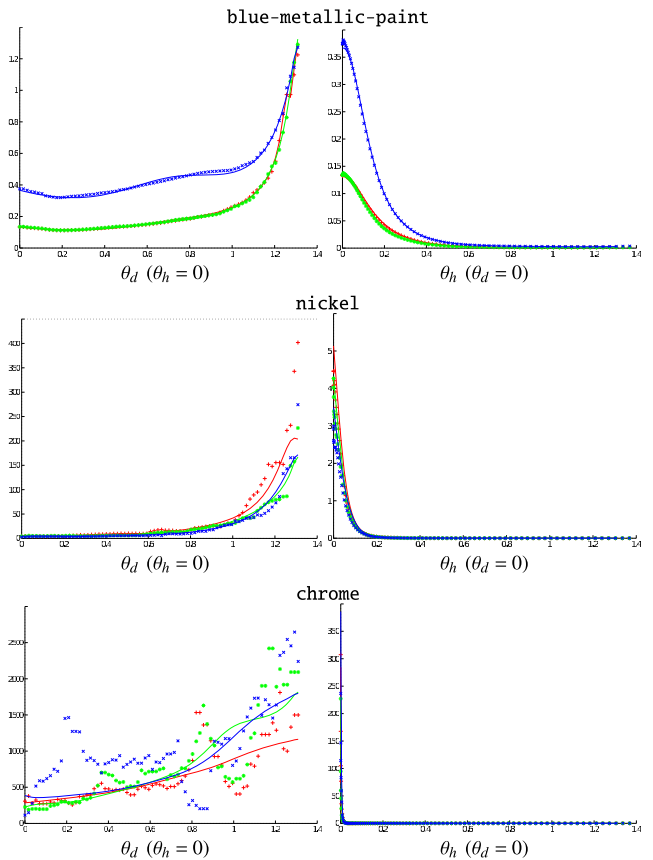

Fig. 5. Specular materials' BRDF data (dots) without grazing angles (i.e., θ_h or $\theta_h > 75°$.) and corresponding RF approximations (full line) per color channel (RGB).

4 RATIONAL FUNCTIONS FOR CDF

4.1 Parametrization

As mentioned in Section 1.2, efficient Importance Sampling strategies significantly increase the rate of convergence of Monte-Carlo rendering techniques. In this section, we present quasioptimal IS of arbitrary BRDFs by means of RF approximation. For the sake of brevity, we only detail our procedure for the standard (l, v) parametrization. Note that this parametrization has been proven to be well suited for glossy surfaces [21].

The principle of IS rendering is to define a *stochastic estimator* for the reflected radiance $L(v)$ in direction v by averaging the contribution from K random light directions l_k selected according to a conditional *Probability Density Function* $\mathrm{PDF}(l_k \mid v) = \mathrm{PDF}_v(l_k)$:

Material	$(n+1, m+1)$			Maximum degree		
	R	G	B	R	G	B
`beige-fabric`	(18,12)	(21,12)	(18,12)	(5,4)	(5,4)	(5,4)
`blue-metallic`	(20,28)	(22,24)	(16,16)	(5,6)	(6,6)	(5,5)
`nickel`	(27,16)	(40,9)	(36,12)	(6,5)	(8,3)	(7,4)
`chrome`	(84,17)	(47,14)	(50,26)	(12,5)	(9,4)	(10,6)

Fig. 6. The data for the fitted RF representation for each material 2D BRDF are shown for the denominator and the numerator, for each (R,G,B) color channel: (center column) number of coefficients, (right column) maximum degree.
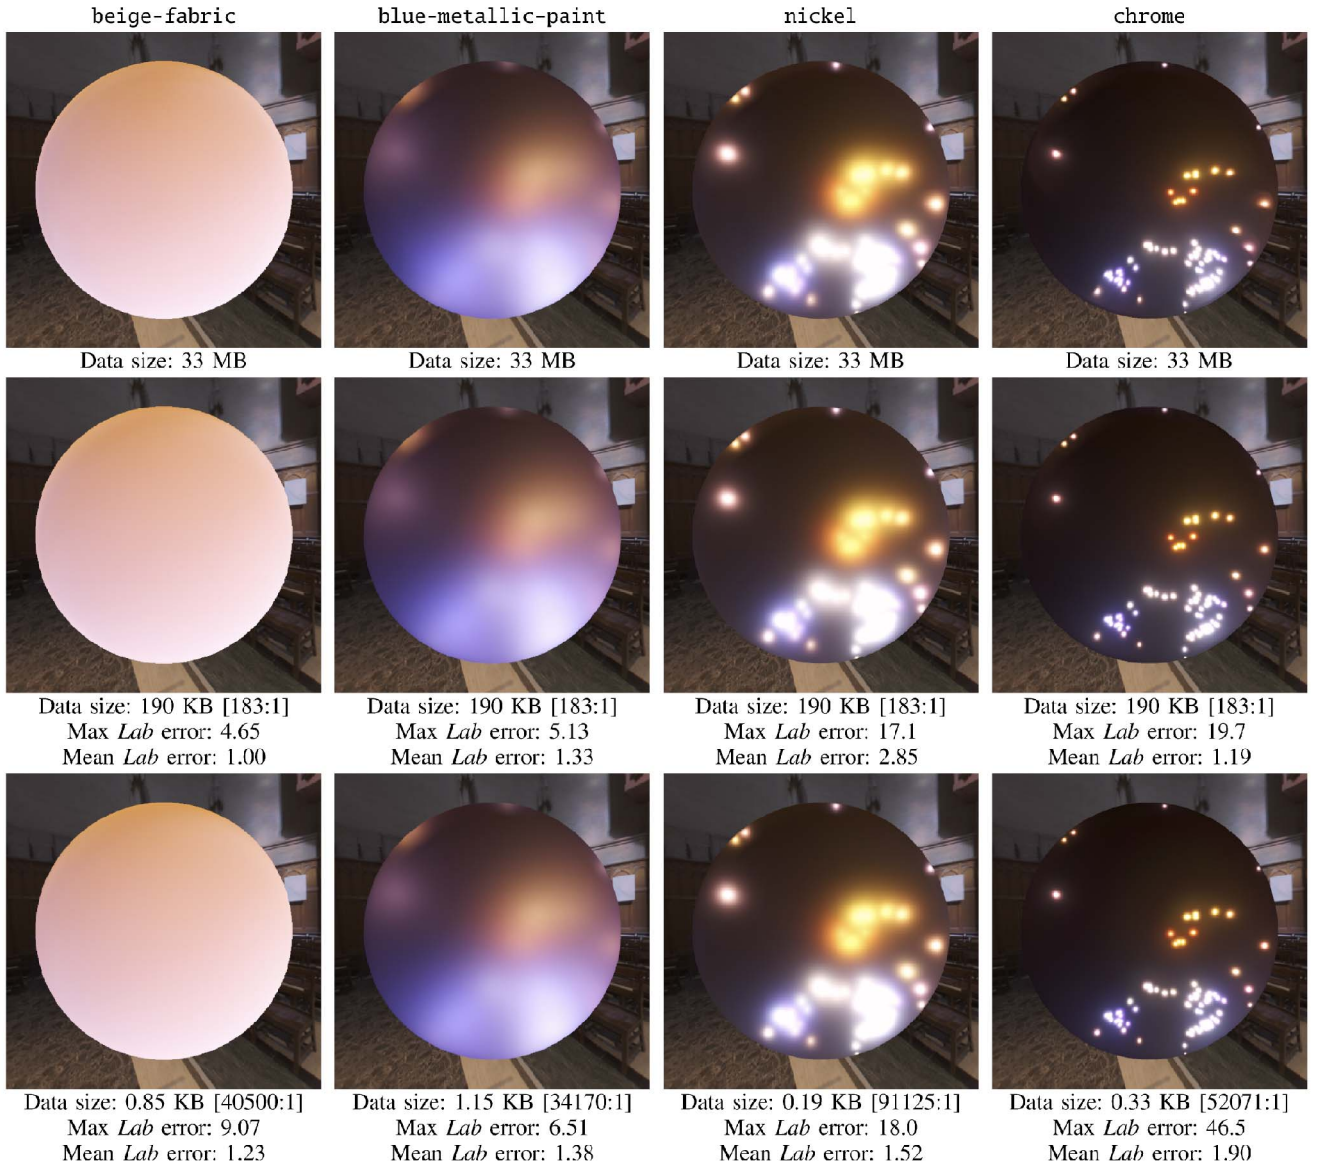
Fig. 7. From top to bottom: original measured data in $(\theta_h, \theta_d, \phi_d)$, projected data in (θ_h, θ_d), rational BRDF fit in (θ_h, θ_d). All scenes are rendered using 1,024 directional light sources sampled from the surrounding environment map. Each Rational BRDF generates a similar magnitude of Lab visual errors as the ones produced by the 2D projected data, but requires between 165 and 1,000 times less storage.

$$L(\boldsymbol{v}) \approx \frac{1}{K} \sum_{k=1}^{K} \frac{\boldsymbol{n} \cdot \boldsymbol{l}_k}{\text{PDF}_{\boldsymbol{v}}(\boldsymbol{l}_k)} \rho(\boldsymbol{l}_k, \boldsymbol{v}) L(\boldsymbol{l}_k). \qquad (6)$$

When there is no prior knowledge about the incident lighting, the optimal choice for the estimator is to use a $\text{PDF}_{\boldsymbol{v}}(\boldsymbol{l}_k)$ that is proportional to the BRDF scaled by $\boldsymbol{n} \cdot \boldsymbol{l}_k$. Each random light direction $\boldsymbol{l} = (\theta_l, \phi_l)$ required by the estimator can then be obtained by generating a pair of uniform random numbers (μ, τ) and inverting a pair of *Cumulative Distribution Functions*, obtained from the integration of the selected PDF:

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

We propose an innovative technique intended to combine the strengths of previous approaches. The idea is to *directly* define a closed-form expression for the inverse CDF, without preliminary analytical formulation of either the CDF or the PDF.

Starting from a measured BRDF scaled by $\boldsymbol{n} \cdot \boldsymbol{l}$, in other words $\rho(\boldsymbol{l}, \boldsymbol{v}) \cos \theta_l$, we first compute a tabulated version of the inverse CDFs

- a 2D table for $\text{CDF}^{-1}(\theta_v, \mu)$ where $\theta_v \in [0, \pi/2]$ indexes the view direction and $\mu \in [0, 1]$ is a random sample,
- a 3D table for $\text{CDF}^{-1}(\theta_v, \tau | \theta_l)$ where $\theta_l \in [0, \pi/2]$ and $\tau \in [0, 1]$ is a second random sample.

Since we are using isotropic materials, there is no azimuthal dependency on the view direction $\boldsymbol{v} = (\theta_v, \phi_v)$, i.e., $\phi_v = 0$. Then, we approximate these tabulated CDF data with a bivariate and trivariate RF, respectively:

$$\begin{cases} \theta_l = \text{CDF}^{-1}(\theta_v, \mu) \approx r_{n_\theta, m_\theta}(\theta_v, \mu) \\ \phi_l = \text{CDF}^{-1}(\theta_v, \tau | \theta_l) \approx r_{n_\phi, m_\phi}(\theta_v, \theta_l, \tau). \end{cases}$$

Note that since the CDF is the integral of its corresponding PDF, the shapes of the inverse CDFs are *always much*

	2D proj. Romeiro 33000 KB	Rational Functions 0.19 KB – 1.15 KB		Non-linear A&S 0.0625 KB	Polynomial Functions 0.19 KB – 1.15 KB
	Error on 3D data	Error on 2D data	Error on 3D data	Error on 3D data	Error on 3D data
`beige-fabric`	(0.71, 0.73, 0.69)	(0.71, 0.60, 0.55)	(0.71, 0.71, 0.69)	(1.66, 1.75, 1.68)	(0.71, 0.72, 0.71)
`blue-metallic`	(0.80, 0.80, 1.0)	(0.25, 0.42, 0.29)	(0.74, 0.71, 1.01)	(28.49, 34.80, 17.50)	(43.47, 77.77, 69.02)
`nickel`	(1.12, 2.44, 2.95)	(0.77, 0.78, 0.86)	(0.77, 2.45, 1.95)	(4.99, 4.11, 4.64)	(47697, 49399, 152752)
`chrome`	(5.67, 4.01, 3.99)	(1.72, 1.97, 3.36)	(3.29, 2.38, 3.38)	(65.04, 133.63, 148.42)	(26941, 1323656, 1125357)

Fig. 8. Maximum relative error of the approximation, per (R,G,B) color channel. In the usual order: 2D projected data compared to 3D original data; rational BRDF compared to the 2D projected data, and compared to the 3D original data; nonlinear fitting of the Ashikhmin-Shirley analytic model compared to the original data, and least squares bivariate polynomials compared to the original 3D data. All data at grazing angles ($>80°$) have been discarded.

simpler than the PDFs, as shown in Fig. 9, and RFs can easily reproduce the abrupt variations that are present in most data sets.

The advantage of this approach is twofold. First, due to the compact representation of RFs, generating each random vector l_k is very efficient as it only involves the evaluation of low-degree multivariate RFs. This is actually done in constant time, contrary to approaches based on the numerical inversion of tabulated CDFs. Second, it allows for a more effective IS strategy by using our new Monte-Carlo estimator (the derivation is available in the dedicated companion document, Section 2), which is directly based on the inverse CDF instead of the PDF:

$$L(v) \approx \frac{1}{K} \sum_{k=1}^{K} \alpha_v(\mu_k, \tau_k) \rho(v, l_k)(n \cdot l_k) L(l_k)$$

$$\text{with } \alpha_v(\mu, \tau) = \frac{\partial \text{CDF}_v^{-1}}{\partial \mu}(\mu) \frac{\partial \text{CDF}_v^{-1}}{\partial \tau}(\tau \mid \theta_l) \sin \theta_l \quad (7)$$

$$\text{and } \theta_l = \text{CDF}_v^{-1}(\mu).$$

Since both inverse CDFs are RFs, their partial derivatives can easily be evaluated on the fly.

Unlike previous approaches (e.g., [21]), which use different parametrizations depending on the type of material, we have decided to use the (θ_l, ϕ_l) parametrization for all materials. This provides a consistent choice and we leave for future work the question of the most suitable parametrization for importance sampling independent of the material type. Nevertheless, as illustrated in Section 4.2, the whole IS strategy proposed here results in a very efficient computation, whatever the complexity of the underlying BRDF, with extremely compact storage compared to prior work.

Fig. 9. Comparisons of the PDF_v and the inverse CDF_v^{-1} for `nickel` material. The measured inverse CDF curves traced in red for $\theta_v = 0°$, in green for $\theta_v = 45°$, and in blue for $\theta_v = 85°$, are always simpler than the ones for the PDF.

4.2 Rational Approximation of CDFs

We proceed as follows: First, we compute two tabulated versions of the inverse CDFs for each of the four selected materials. The values of $\text{CDF}_v^{-1}(\mu) \in [0, \pi/2]$ are computed on a 91×91 grid. Similarly, the values of $\text{CDF}_v^{-1}(\tau|\theta_l) \in [0, \pi]$ are computed on a $91 \times 91 \times 128$ grid, which size is then reduced by a factor of two using the fact that the inverse CDF is symmetric. Second, to accelerate the resolution of $\mathcal{P}$, we use 300 measured values as training data (cf. Section 2.2).

As explained in Section 2.2, we guarantee that the fitted RF for the 3D inverse CDF is symmetric with respect to θ_v and θ_l, i.e., $r_{n_\phi, m_\phi}(\theta_v, \theta_l, \tau) = r_{n_\phi, m_\phi}(\theta_l, \theta_v, \tau)$. Moreover, we guarantee that both RFs $r_{n_\theta, m_\theta}(\theta_v, \mu)$ and $r_{n_\phi, m_\phi}(\theta_v, \theta_l, \tau)$ are monotonically increasing with respect to μ and τ, respectively. Besides the monotonicity and symmetry properties, we also enforce boundary conditions on the CDFs:

$$\text{CDF}^{-1}(\theta_v, \ \mu \ = \ 0) = 0 \qquad \text{CDF}^{-1}(\theta_v \ = \ 0, \ \theta_l, \tau) = \pi\tau$$

$$\text{CDF}^{-1}(\theta_v, \ \mu \ = \ 1) = \frac{\pi}{2} \qquad \text{CDF}^{-1}(\theta_v, \ \theta_l \ = \ 0, \tau) = \pi\tau$$

by imposing the following rational forms:

$$r_{n_\theta, m_\theta}(\theta_v, \mu) = \frac{\pi}{2} \mu + \mu(1 - \mu) \frac{p_{n_\theta, m_\theta}(\theta_v, \mu)}{q_{n_\theta, m_\theta}(\theta_v, \mu)} \quad (8)$$

$$r_{n_\phi, m_\phi}(\theta_v, \theta_l, \tau) = \pi \tau + \tau(1 - \tau) \theta_v \theta_l \frac{p_{n_\phi, m_\phi}(\theta_v, \theta_l, \tau)}{q_{n_\phi, m_\phi}(\theta_v, \theta_l, \tau)}. \quad (9)$$

For each of the four selected materials, we apply the fitting algorithm described in Section 2, complemented with the corresponding modifications and new constraints mentioned in Section 2.2. The choice of the intervals $F_i = [\underline{f_i}, \overline{f_i}]$, now containing either θ_l^i or ϕ_l^i, for each of the inverse CDF values, completely determines the behavior of the approximations. Since these inverse CDFs guide the IS process, the interval widths are taken such that the resulting variance is acceptable, while leaving the number $m + n$ of coefficients in the approximation manageable. As illustrated below, this implies a fairly accurate $\text{CDF}_v^{-1}(\mu)$ approximation, and capturing the main trend of $\text{CDF}_v^{-1}(\tau|\theta_l)$. For the 2D $\text{CDF}_v^{-1}(\mu)$, we define the intervals as:

$$\begin{cases} \underline{f_i} &= \theta_l^i - \epsilon(1 + \theta_l^i) \\ \overline{f_i} &= \theta_l^i + \epsilon(1 + \theta_l^i) \end{cases} \qquad i = 0, \dots, 91 \times 91 - 1$$
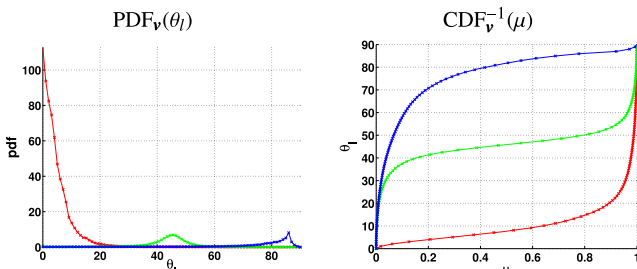
	beige-fabric	blue-metallic	nickel	chrome
Factored	na	$\simeq 200$	$\simeq 200$	na
Cascade PDF	374.16	1685.64	2057.51	na
A. Disc	272.62	715.91	846.15	na
RFs	0.117	0.211	0.230	0.172

Fig. 10. Comparison of the CDF and PDF data size in KB for different techniques: Factored representation [21], Cascade PDF [29], A. Disc [30], and Rational Functions. Our Rational Function approach consumes always less memory than the previous approaches.

with $\epsilon = 0.015$ for beige-fabric, $\epsilon = 0.035$ for blue-metallic, $\epsilon = 0.05$ for nickel, and $\epsilon = 0.1$ for the highly specular chrome. In a similar way, we define the intervals for 3D $\mathrm{CDF}_v^{-1}(\tau|\theta_l)$ according to ϕ_l with $\epsilon = 0.015$ for beige-fabric and blue-metallic, $\epsilon = 0.72$ for nickel, and $\epsilon = 0.8$ for chrome. For the 3D $\mathrm{CDF}_v^{-1}(\tau|\theta_l)$, we relax the intervals near $\theta_v = \theta_l$ and near grazing angles $\theta_v, \theta_l > 80°$. In order to keep the number of coefficients manageable in $r_{n_\phi, m_\phi}(\theta_v, \theta_l, \tau)$, we restrict the partial degree in τ to at most 2, which is sufficient for capturing the steep increase in τ.

Fig. 11 shows the 2D approximations of $\mathrm{CDF}_v^{-1}(\mu)$ whereas Fig. 12 shows the 3D approximations of $\mathrm{CDF}_v^{-1}(\tau|\theta_l)$. The relative errors shown are

$$\frac{|\mathrm{CDF}^{-1} - r_{n,m}|}{1 + |\mathrm{CDF}^{-1}|},$$

where $r_{n,m}$ is the corresponding rational approximation. Note from these figures that the number of coefficients needed is always very moderate. This indicates that the proposed forms in (8) and (9) are adequate and capture most of the information. The table of Fig. 14 presents the fitting time for each of the 2D $\mathrm{CDF}_v^{-1}(\mu)$. These timings are directly proportional to the number of coefficients tested.

Fig. 13 further compares, for each material, the rendering quality obtained with tabulated CDFs and PDFs (first row) against approximated CDFs (second row). We also compute the variance of the sphere pixels in each image to compare the performance of the IS strategy based on our RF CDF with that of the tabulated data. To provide a fair comparison of the sampling efficiency, our renderer uses the BRDF combined either with the tabulated CDF or its RF approximation. For the blue-metallic, beige-fabric, and nickel materials, the resulting variance is as low as the reference solution obtained with tabulated data. For the chrome material, the variance is ten times larger than the reference solution, but the *Lab* difference between the images shows that our RF approximation still provides a visually satisfactory result. Chrome remains the most challenging material. However, to our knowledge, previous approaches could not even represent the chrome data.

Note that our technique **guarantees** the existence of a representable inverse CDF **without** introducing a bias in the sampling process because the RFs preserve monotonicity. However, the speed of convergence of the IS is directly dependent on the overall quality of the fitted function.

Regarding the size of our functions, Fig. 10 shows that with RF approximations, we achieve high compression rates. For example, the nickel material in our approach is $7,600\times$ (resp. $3,100\times$) more compact than the one proposed by Lawrence et al. [21] (resp. [30]). The numbers from the first row of the table have been directly reported from the original paper whereas for the other rows, we have reported the lowest available sizes, which are coming from Montes' PhD thesis [39]. Notice that previous approaches do not provide any results for specular chrome whereas our technique may still be used to approximate an importance function. This high compression ratio directly results from the properties of RFs, which can efficiently represent steep changes with low-degree approximations, as well as from
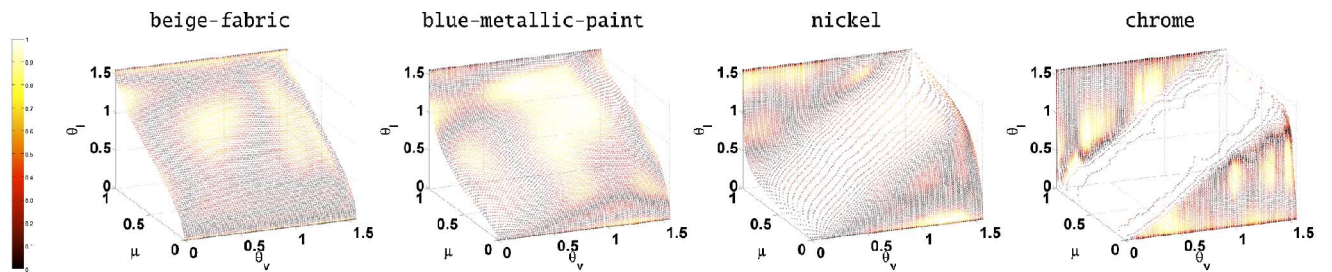

Fig. 11. Rational Function approximation of the bivariate $\mathrm{CDF}_v^{-1}(\mu)$ for our four selected MERL-MIT materials. For each material, the colormap represents the fitting relative error and ranges from the minimum relative error (black) to the maximum relative error (yellow). The number of coefficients is expressed as a pair (numerator, denominator).
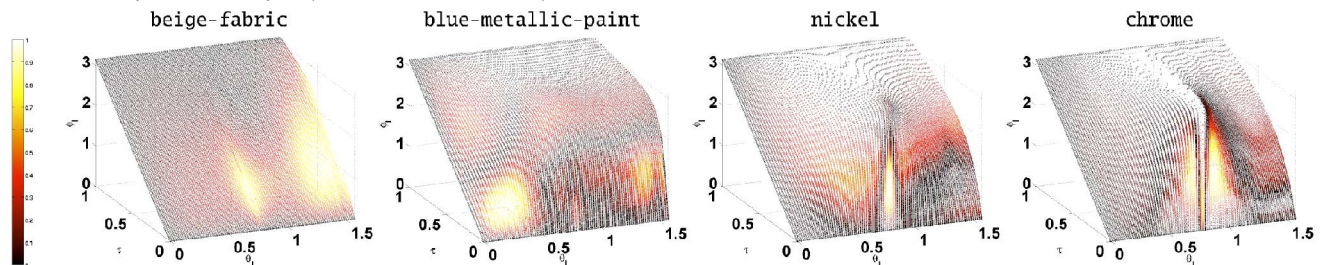

Fig. 12. Slices ($\theta_v = 45°$) of the Rational Function approximation of the trivariate $\mathrm{CDF}_v^{-1}(\tau|\theta_l)$ for our four selected MERL-MIT materials. For each material, the colormap represents the fitting relative error and ranges from the minimum relative error (black) to the maximum relative error (yellow). The mean and max relative errors are given for the whole 3D data. The number of coefficients is expressed as a pair (numerator, denominator).

| beige-fabric | blue-metallic-paint | nickel | chrome |

Tabulated CDFs and PDFs

Variance (0.0153, 0.0084, 0.0088)   Variance (0.0006, 0.0004, 0.0033)   Variance (0.0025, 0.0020, 0.0022)   Variance (0.0077, 0.0062, 0.0095)

Rational Function Approximation of the inverse CDFs

Variance (0.0152, 0.0084, 0.0088)   Variance (0.0009, 0.0007, 0.0048)   Variance (0.0017, 0.0141, 0.0156)   Variance (0.0701, 0.0631, 0.0937)
Max *Lab* error: 2.09   Max *Lab* error: 10.2   Max *Lab* error: 2.90   Max *Lab* error: 24.25
Mean *Lab* error: 0.43   Mean *Lab* error: 1.19   Mean *Lab* error: 0.44   Mean *Lab* error: 1.13
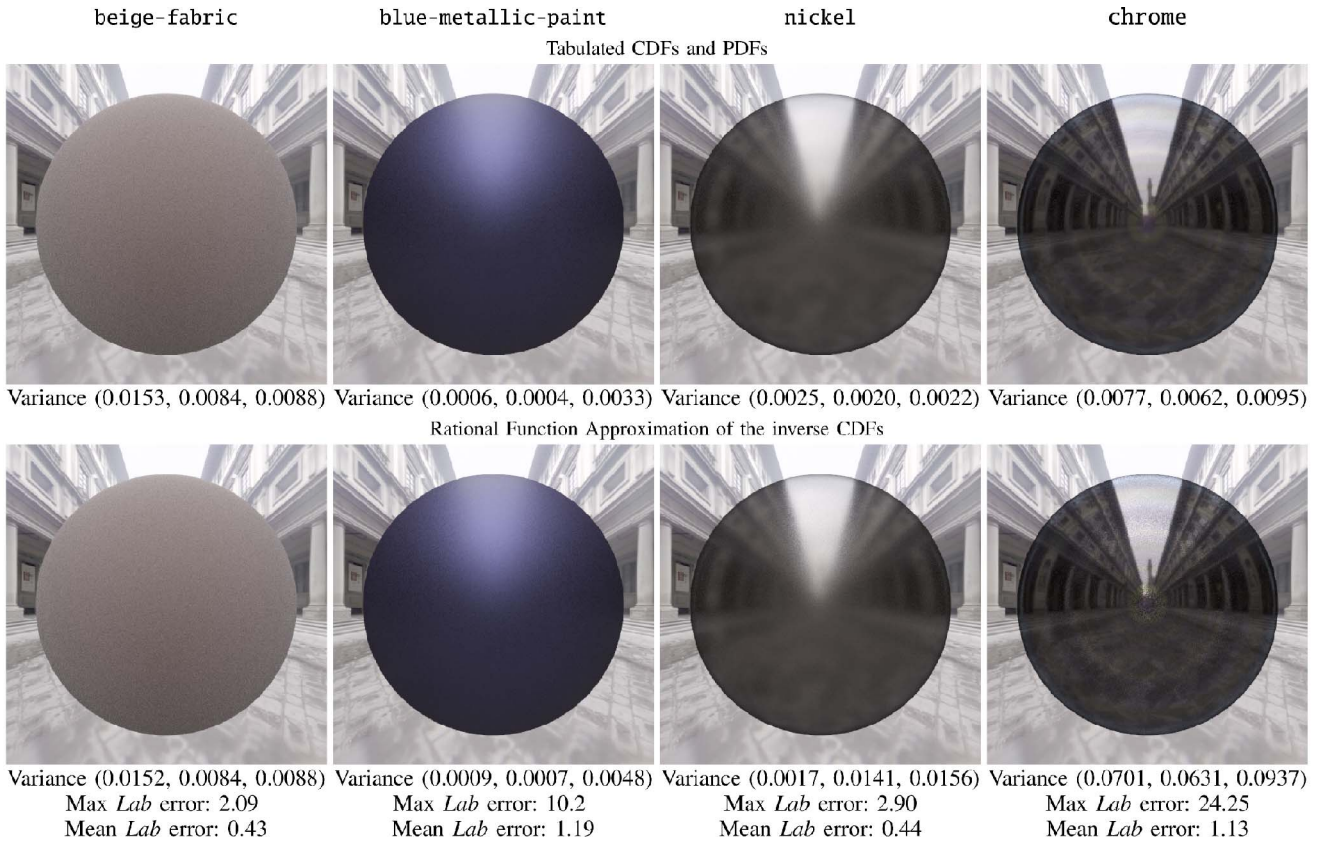
Fig. 13. Comparison of Importance Sampling with the Rational Function approximation of inverse CDFs and the original tabulated data. Four hundred samples per pixel were used. The numbers under each image indicate the mean variance per color channel. Compared to the original tabulated data where the different CDFs and PDFs require 10 MB for each material, our RF approach requires only between 0.117 and 0.230 KB.

the fact that we do not store any PDFs. As explained earlier, the PDF is computed on the fly from the inverse CDF during the rendering step and thus does not require any supplemental storage.

Finally, the efficiency of the IS process can be observed in Fig. 1 which presents a global illumination rendering of a more complex scene. Compared to the reference solution, resulting in a small *Lab* difference, our solution achieves the same overall rendering quality but with a drastically lower memory footprint (90 MB + 30 MB versus 1.67 KB + 0.6 KB) when using the 3D measured data, and (570 KB + 30 MB versus 1.67 KB + 0.6 KB) when using the 2D projected data.

5 ANISOTROPIC BRDF MODEL USING RFS

Ward [7] observed that for many anisotropic materials, the variation of the reflected intensity when rotating the surface around its normal vector often consists of a simple scaling factor applied to an average isotropic lobe. Based on that observation, we propose the following model for aniso-tropic BRDFs as a product of two rational functions:

$$\rho(\boldsymbol{l}, \boldsymbol{v}) \approx r_{m',n'}^{a}(\phi_h) r_{m,n}^{i}(\theta_h, \theta_d),$$

where $r_{m,n}^{i}(\theta_h, \theta_d)$ represents a standard isotropic Rational BRDF and $r_{m',n'}^{a}(\phi_h)$ is a scaling factor defining the anisotropic variation. This scaling factor agrees with classical brushed-metal behavior [7] and with several other anisotropic models [8], [10], [28].

To validate this model, we have chosen to generate "hybrid" anisotropic BRDFs by postprocessing measured isotropic materials. We may view the scaling function as an approximation of the directional variation generated by weaving mesostructures of textiles, or regular geometric structures of crystals. Actually, the whole process can be seen as a combination of the hybridization technique, based on a linear combination of materials, proposed by Matusik et al. [20], and the mesostructure simulation technique used by Westin et al. [14] to predict reflectance from complex materials. Fig. 15 presents results obtained with different anisotropic transformations of the `blue-metallic-paint` material.

In theory, there is nothing preventing our RF method from approximating a measured full 4D BRDF. The results obtained so far suggest that the 3D subset $(\theta_h, \theta_d, \phi_h)$ of Rusinkiewicz's parametrization is sufficient for some anisotropic BRDFs, similar to the use of (θ_h, θ_d) for isotropic BRDFs [35]. Reducing the dimension to 3D leads to fewer constraints for the interval interpolation technique [32], and results in a large speedup of the fitting process. Unfortu-nately, far less data are publicly available. We experimented
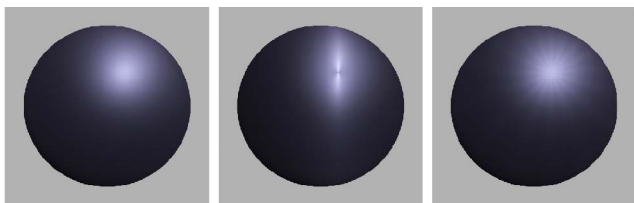
beige-fabric	blue-metallic	nickel	chrome
2.82	5.74	19.19	6.44

Fig. 14. Fitting timings in seconds computed with our *Matlab* program on a Intel Xeon L5420@2.5 GHz for the 2D inverse CDF. These timings include all iterations of the fitting algorithm detailed in Section 2.

Fig. 15. A `blue-metallic-paint` sphere illuminated by a single directional light source, with different anisotropic transformations using a univariate function $r^a_{m',n'}(\phi_h)$. The original isotropic material is shown on the left. The middle image corresponds to an elliptical scaling and presents the typical butterfly pattern of brushed metals. The image on the right corresponds to a star-shaped scaling function.

with another MIT BRDF database [26] containing four measured anisotropic materials. However, as noted in previous work [10], the embedded data contain significant acquisition noise, and present large areas in its parametric domain where the measured BRDF values are quite sparse. These two issues lead to artifacts in the fitted data. Smoothing and good quality data completion [11] should improve our results, but this remains one of our future research directions.

6 CONCLUSION AND FUTURE WORK

In this paper, we have introduced *Rational BRDFs*, a generic and compact representation of arbitrary BRDFs, based on Rational Functions. By employing a subset of the BRDF parametrization introduced by Rusinkiewicz [34], isotropic BRDFs can be represented as bivariate RFs. Very compact approximations of measured BRDFs are obtained with a memory footprint that is usually less than one kilobyte for arbitrary isotropic materials. Moreover, the same approximation process applies to the inverse CDF generated from the corresponding BRDF multiplied by the cosine factor. Combined with our new Monte-Carlo estimator which exploits the RF formulation, our approach offers a quasioptimal IS scheme, with compact storage compared to prior work. We have shown that the Rational BRDF is suitable to reproduce some anisotropic effects on a 3D subspace of Rusinkiewicz's parametrization [34] by the introduction of a univariate rational function that scales an isotropic Rational BRDF according to its orientation.

All of our results are the foundations for future research. First, we have started investigating acquisition techniques to obtain dense BRDF measurements of materials, focusing on families of BRDFs that are currently lacking in public databases such as materials exhibiting retroreflective behavior, and various types of anisotropic reflections. Based on this data set, we may extend the previous studies on parametrization [35], [37] to more generic materials. In investigating extensions to the work of Lawrence et al. [21], we also hope to find a set of best suited parametrizations for IS.

Second, as shown by the anisotropic transformation of `blue-metallic-paint` presented above, generating hybrid materials by combining analytical and measured BRDFs appears to be a promising direction of research. We are currently formalizing this process to define a flexible framework for intuitive BRDF editing.

Finally, for applications where lighting is exclusively defined by environment maps, we intend to generalize our IS scheme to directly process the product of the BRDF with the environment map, following ideas presented by Jarosz et al. [40].

REFERENCES

[1] J.T. Kajiya, "The Rendering Equation," *Proc. ACM SIGGRAPH '86*, pp. 143-150, 1986.
[2] B. Phong, "Illumination for Computer Generated Pictures," *Comm. ACM*, vol. 18, no. 6, pp. 311-317, 1975.
[3] X. He, K. Torrance, F. Sillion, and D. Greenberg, "A Comprehensive Physical Model for Light Reflection," *Proc. ACM SIGGRAPH '91*, pp. 175-186, 1991.
[4] J. Stam, "Diffraction Shader," *Proc. ACM SIGGRAPH '99*, pp. 75-84, 1999.
[5] I. Icart and D. Arquès, "A Physically-Based BRDF Model for Multilayer Systems with Uncorrelated Rough Boundaries," *Proc. Eurographics Workshop Rendering Techniques (EGWR '00)*, pp. 353-364, 2000.
[6] R. Cook and K. Torrance, "A Reflectance Model for Computer Graphics," *ACM Trans. Graphics*, vol. 1, no. 1, pp. 7-24, 1982.
[7] G. Ward, "Measuring and Modeling Anisotropic Reflection," *Proc. ACM SIGGRAPH '92*, pp. 265-272, 1992.
[8] C. Schlick, "An Inexpensive BRDF Model for Physically-Based Rendering," *Proc. Conf. Eurographics '94*, pp. 233-246, 1994.
[9] D. Edwards, S. Boulos, P. Shirley, M. Ashikhmin, M. Stark, and C. Wyman, "The Halfway Vector Disk for BRDF Modeling," *ACM Trans. Graphics*, vol. 25, no. 1, pp. 1-18, 2006.
[10] M. Ashikhmin and S. Premoze, "Distribution-Based BRDFs," technical report, Univ. of Utah, unpublished, 2007.
[11] J. Wang, S. Zhao, X. Tong, J. Snyder, and B. Guo, "Modeling Anisotropic Surface Reflectance with Example-Based Microfacet Synthesis," *Proc. ACM SIGGRAPH '08*, article 41, 2008.
[12] K. Torrance and E. Sparrow, "Polarization, Directional Distribution, and Off-Specular Peak in Light Reflected from Roughened Surfaces," *J. Optical Soc. Am.*, vol. 57, no. 9, pp. 1105-1114, 1967.
[13] B. Cabral, N. Max, and R. Springmeyer, "Bidirectional Reflection Functions from Surface Bump Maps," *Proc. ACM SIGGRAPH '87*, pp. 273-281, 1987.
[14] S. Westin, J. Arvo, and K. Torrance, "Predicting Reflectance Functions from Complex Surfaces," *Proc. ACM SIGGRAPH '92*, pp. 255-264, 1992.
[15] J. Koenderink, A. Doorn, and M. Stavridi, "BRDF Expressed in Terms of Surface Scattering Modes," *Proc. European Conf. Computer Vision*, pp. 28-39, 1996.
[16] P. Schröder and W. Sweldens, "Spherical Wavelets: Efficiently Representing Functions on the Sphere," *Proc. ACM SIGGRAPH '95*, pp. 161-172, 1995.
[17] T. Zickler, S. Enrique, R. Ramamoorthi, and P. Belhumeur, "Reflectance Sharing: Image-Based Rendering from a Sparse Set of Images," *Proc. Eurographics Symp. Rendering (EGSR '05)*, pp. 253-265, 2005.
[18] M. McCool, J. Ang, and A. Ahmad, "Homomorphic Factorization of BRDFs for High-Performance Rendering," *Proc. ACM SIGGRAPH '01*, pp. 171-178, 2001.
[19] J. Kautz and M. McCool, "Interactive Rendering with Arbitrary BRDFs Using Separable Approximations," *Proc. Eurographics Workshop Rendering (EGWR '99)*, pp. 247-260, 1999.
[20] W. Matusik, H. Pfister, M. Brand, and L. McMillan, "A Data-Driven Reflectance Model," *Proc. ACM SIGGRAPH '03*, pp. 759-769, 2003.
[21] J. Lawrence, S. Rusinkiewicz, and R. Ramamoorthi, "Efficient BRDF Importance Sampling Using a Factored Representation," *Proc. ACM SIGGRAPH '04*, pp. 496-505, 2004.
[22] P. Weistroffer, K. Walcott, G. Humphreys, and J. Lawrence, "Efficient Basis Decomposition for Scattered Reflectance Data," *Proc. Eurographics Symp. Rendering (EGSR '07)*, pp. 207-218, 2007.
[23] D. Mahajan, Y.-T. Tseng, and R. Ramamoorthi, "An Analysis of the In-Out BRDF Factorization for View-Dependent Relighting," *Proc. Eurographics Symp. Rendering (EGSR '08)*, pp. 1137-1145, 2008.
[24] E. Lafortune, S.-C. Foo, K. Torrance, and D. Greenberg, "Non-Linear Approximation of Reflectance Functions," *Proc. ACM SIGGRAPH '97*, pp. 117-126, 1997.

[25] J. Wang, P. Ren, M. Gong, J. Snyder, and B. Guo, "All-Frequency Rendering of Dynamic, Spatially-Varying Reflectance," *Proc. ACM SIGGRAPH Asia '09,* article 133, 2009.

[26] A. Ngan, F. Durand, and W. Matusik, "Experimental Analysis of BRDF Models," *Proc. Eurographics Symp. Rendering (EGSR '05),* pp. 117-226, 2005.

[27] P. Dutré, K. Bala, and P. Bekaert, *Advanced Global Illumination.* A.K. Peters, 2006.

[28] M. Kurt, L. Szirmay-Kalos, and J. Křivánek, "An Anisotropic BRDF Model for Fitting and Monte Carlo Rendering," *ACM SIGGRAPH Computer Graphics,* vol. 44, no. 1, article 3, 2010.

[29] J. Lawrence, S. Rusinkiewicz, and R. Ramamoorthi, "Adaptive Numerical CDF for Efficient Importance Sampling," *Proc. Eurographics Symp. Rendering (EGSR '05),* pp. 11-20, 2005.

[30] R. Montes, C. Urena, R. Garcia, and M. Lastra, "Generic BRDF Sampling: A Sampling Method for Global Illumination," *Proc. Third Int'l Conf. Computer Graphics Theory and Applications (GRAPP '08),* pp. 191-198, 2008.

[31] A. Cuyt, "Recent Applications of Rational Approximation Theory: A Guided Tour," *Proc. Conf. Numerical Analysis and Computational Math.,* G. Psihoyios, ed., pp. 50-52, 2003.

[32] O. Salazar Celis, A. Cuyt, and B. Verdonk, "Rational Approximation of Vertical Segments," *Numerical Algorithms,* vol. 45, pp. 375-388, 2007.

[33] A. Wills, *QPC-Quadratic Programming in C,* School of Electrical Eng. and Computer Science, Univ. of Newcastle, Callaghan, Australia, http://sigpromu.org/quadprog/. 2010.

[34] S. Rusinkiewicz, "A New Change of Variables for Efficient BRDF Representation," *Proc. Eurographics Workshop Rendering (EGWR '98),* pp. 11-22, 1998.

[35] F. Romeiro, Y. Vasilyev, and T. Zickler, "Passive Reflectometry," *Proc. European Conf. Computer Vision,* pp. 859-872, 2008.

[36] A. Ben-Artzi, R. Overbeck, and R. Ramamoorthi, "Real-Time BRDF Editing in Complex Lighting," *Proc. ACM SIGGRAPH '06,* pp. 945-954, 2006.

[37] M. Stark, J. Arvo, and B. Smits, "Barycentric Parameterizations for Isotropic BRDFs," *IEEE Trans. Visualization and Computer Graphics,* vol. 11, no. 2, pp. 126-138, Mar. 2005.

[38] F.X. Sillion, J.R. Arvo, S.H. Westin, and D.P. Greenberg, "A Global Illumination Solution for General Reflectance Distributions," *Proc. ACM SIGGRAPH '91,* pp. 187-196, 1991.

[39] R. Montes, "An Importance Sampling Method for Arbitrary BRDFs Used in Global Illumination Applications," PhD dissertation, Univ. of Granada, 2008.

[40] W. Jarosz, N. Carr, and H. Jensen, "Importance Sampling Spherical Harmonics," *Proc. Conf. Eurographics '09,* pp. 577-586, 2009.

Romain Pacanowski received the engineer degree from EFREI and the MS degree from the University of Bordeaux. He received the PhD degree from the University of Bordeaux, France, and Université de Montréal, Canada. He was a postdoctoral fellow at the CEA in 2010 and for Disney Interactive Media Group in 2011. Since December 2011, he is working at the CNRS as a research-engineer. His research interests include realistic rendering as well as appearance modeling and acquisition.

Oliver Salazar Celis received the MSc degree in computer science in 2004 and the Doctor Scientiae degree in 2008, both from the University of Antwerp, Belgium. From October 2008 until December 2011, he was a postdoctoral researcher in the Department of Mathematics and Computer Science at the same university. His research interests include rational approximations, numerical analysis, optimization and their use in applications ranging from telecommunication networks to signal processing. Currently, his research focuses on modeling and numerical aspects in computational finance.

Christophe Schlick received the PhD degree in 1992 for his work on BRDF models and Monte Carlo techniques. He is a professor in computer science at the University of Bordeaux 2 (France) where he has recently headed the Applied Mathematics and Computer Science Department. His research interests include many aspects of computer graphics, including global illumination, procedural texture and geometric synthesis, curves and surfaces, point-based modeling and rendering.

Xavier Granier received the engineer and MS degrees from the Grenoble Institute of Technology and the PhD Diploma for University Joseph Fourier in Grenoble, France. He is currently at INRIA Bordeaux Sud-Ouest, France, as a research scientist. His research interests include realistic lighting and expressive rendering, appearance modeling and acquisition. He is a member for ACM and EUROGRAPHICS.

Pierre Poulin received the MSc degree from the University of Toronto and the PhD degree from the University of British Columbia, both in computer science. He is a full professor in the Computer Science and Operations Research Department of the Université de Montréal. He has served on program committees of more than 45 international conferences. His research interests cover a wide range of topics, including image synthesis, image-based modeling, procedural modeling, natural phenomena, scientific visualization, and computer animation.

Annie Cuyt received the Doctor Scientiae degree in 1982 from the University of Antwerp, summa cum laude and with the felicitations of the jury. She is a full professor at the Department of Mathematics and Computer Science of the University of Antwerp. Subsequently, she was a research fellow with the Alexander von Humboldt Foundation, Germany, and she was honored with a Masuda Research Grant, Japan. She is the author of more than 100 publications in international journals and conference proceedings, the author or editor of several books and the organizer of a number of international events. Her current interests are in rational approximation theory and its applications in scientific computing.

▷ **For more information on this or any other computing topic, please visit our Digital Library at** www.computer.org/publications/dlib.