genBRDF: Synthesizing Novel Analytic BRDFs with Genetic Programming

Adam Brady

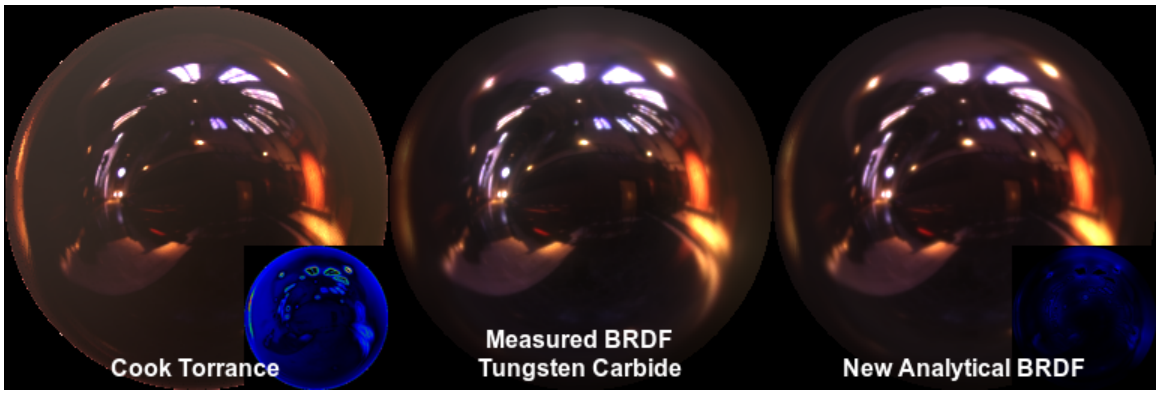

Figure 1: *Comparison of BRDFs modeling the tungsten carbide material from the MERL BRDF database. Each scene consists of a sphere rendered under the grace cathedral lighting environment.* ***Center:*** *Ground truth render produced from measured, tabulated data.* ***Left:*** *Cook Torrance fit to MERL data using parameters learned by [Ngan et al. 2005].* ***Right****: An analytical BRDF produced using our search technique.*

Abstract

We present a framework for automatically synthesizing new analytical BRDFs to model classes of measured materials using a heuristic-based search commonly called Genetic Programming. Our approach can be seen as an extension of traditional approaches to reflectance modeling that rely either on a phenomenological process or a physically-based model. In both cases, our technique augments the human effort involved in deriving suitable mathematical approximations of complex reflectance functions with a large-scale search. We present results that show our technique is able to find novel mathematical expressions that are adjustable, compact, and more accurate than previously available alternatives for representing common classes of materials like metals and dielectrics. We also describe a new fast approximation of a common perceptual BRDF error metric used to make this search tractable.

CR Categories: I.3.3 [Computer Graphics]: Three-Dimensional Graphics and Realism—Display Algorithms I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Radiosity;

Keywords: BRDF, MERL data, genetic programming, meta-heuristic search

Links: DL PDF

1 Introduction

Accurately modeling material appearance plays a critical role in photo-realistic rendering. Despite our understanding of the physics of light propagation, real-world materials include many complex and subtle properties that defy simple mathematical approximations. For this reason, deriving a BRDF that is both accurate enough to meet the demands of modern rendering systems yet general enough to express a range of interesting materials remains a difficult task.

This paper focuses on homogeneous opaque surfaces whose appearance is characterized by the Bidirectional Reflectance Distribution Function (BRDF) [Nicodemus 1965]. Traditionally, BRDFs have been derived either according to some phenomenological process [Phong 1975] hypothesized model of surface micro-geometry and the physics of light propagation [Cook and Torrance 1982b; He et al. 1991]. A more recent trend has been to measure the BRDFs of physical samples and use those measurements either directly, in so called data-driven models [Matusik et al. 2003], or as input to an optimization process that determines the best fitting parameters of an analytical model [Ngan et al. 2005].

Analytical BRDFs are desirable for their compactness and the fact that they often include adjustable parameters that can be used by a designer to author a wide range of materials [Dorsey et al. 2008]. Their main drawback is that they are typically less accurate than data-driven models and often fail to capture subtle aspects of material appearance [Ngan et al. 2005]. Figure 1 illustrates the gap that remains between state-of-the-art analytic BRDFs and measured data in the case of tungsten carbide. This example is representative of the performance of leading analytical BRDFs in fitting measured samples [Ngan et al. 2005]. Indeed, something in the measured data is lost, but it's difficult to say exactly what that is and, more importantly, exactly how to modify the analytical function to achieve a better match.

The goal of this paper is to develop new analytic BRDFs that are more accurate than current alternatives without losing their advantages (compactness and comprising a handful of adjustable parameters suitable for exploration and intuitive material design). Our approach is to perform a large-scale heuristic-based random search,

often called Genetic Programming (GP), that considers hundreds of thousands of symbolic transformations of a few "seed" analytical models with the goal of achieving a more accurate fit to measured data. The result of this search is a series of analytical expressions that trade accuracy for complexity, from which a final model is chosen and further improved by culling unimportant sub-expressions.

This seemingly inefficient "blind" approach to a staggeringly difficult optimization problem like learning new analytical BRDFs has only recently become more feasible due to the increasing availability of inexpensive large-scale computing resources. In other words, just like monkeys and typewriters, directed stochastic search eventually works if you are able to consider enough variants!

We use this framework to synthesize new analytical BRDF models for metals and dielectrics, respectively, relying on the MERL-MIT BRDF database [Matusik et al. 2003] to guide the search. We present results that show these new analytical BRDFs are more accurate than current alternatives and that their mathematical structure offers some interesting insights into real-world reflectance functions. Additionally, we introduce a new fast approximation of a perceptual BRDF error metric [Pereira and Rusinkiewicz 2012] that is necessary to make this search tractable. Finally, we examine the parameters in these new models and demonstrate that they support intuitive material design and conclude with several ideas of how this basic approach to BRDF modeling could be applied to related problems in the field.

2 Related Work

Analytic BRDF Models The large number of analytic BRDF models can be roughly divided into two groups: phenomenological and physically-based. A phenomenological approach aims to reproduce the qualitative aspects of material appearance through mathematical expressions that have no direct connection to the optical properties of the underlying material. This approach led to the seminal Phong BRDF [Phong 1975]. The current paper may be thought of as an automated large-scale phenomenological approach to BRDF modeling. Our GP search is essentially identifying mathematical expressions chosen at random that match measured data better than alternative expressions without consideration of the material's underlying physical structure.

Physically-based analytical BRDF models result from applying some model of light propagation (e.g. geometric, wave-based, etc.) to a hypothesized surface micro-geometry. The Torrance-Sparrow BRDF [Torrance and Sparrow 1992] and its many derivatives [Cook and Torrance 1982a; Ashikhmin and Shirley 2000] are perhaps the most widely used physically-based models. They posit that surfaces are composed of randomly oriented Fresnel mirrors (microfacets) whose aggregate orientations follow some analytical probability distribution function (e.g., the Beckmann distribution [Beckmann and Spizzichino 1963]). Other physically-based models either rely on more general models of light propagation, assume different optical characteristics of the individual microfacets [Oren and Nayar 1994], or consider different categories of microgeometries like oriented grooves for brushed metal [Ashikhmin and Shirley 2000] or fibers [Marschner et al. 2003; Marschner et al. 2005].

The principal advantages of analytical BRDF models are two-fold. They are compact and often contain only a handful of adjustable parameters that enable intuitive material design [Dorsey et al. 2008]. Their main drawback is that they often fail to achieve the fidelity required for modern rendering systems [Ngan et al. 2005]. Our goal is to develop new analytical models that retain these benefits while providing better matches to measured data.

Data-Driven BRDF Models A recent trend in material modeling is to measure the BRDFs of physical samples [Ward 1992; Marschner et al. 1999; Matusik et al. 2003; Ngan et al. 2005]. Using measured data directly for rendering produces realistic images, but tabulated data cannot be easily edited and carries a high storage cost. Therefore, it's more common to fit the parameters of an existing analytic BRDF to measurements [Ngan et al. 2005] or use any number of dimensionality-reduction algorithms to produce more compact and useful representations [Matusik et al. 2003; Lawrence et al. 2004; Peers et al. 2006]. However, these representations still present significant difficulties for material design [Dorsey et al. 2008] and are less compact than analytical models.

We use the MIT-MERL database of isotropic BRDFs to guide our GP search [Matusik et al. 2003]. This database includes 20 metal surfaces (e.g., copper, steel, metallic paints etc.) and 20 common dielectrics like plastics, ceramics, and coated surfaces. We derive a new analytic BRDF for each of these classes that outperform available analytic models [Ngan et al. 2005].

GP Shader Simplification and Code Repair A source of inspiration for this work is the recent work of Sitthi-Amorn et al. [2011], which applied Genetic Programming to the problem of automatic procedural shader simplification. Their results indicate the ability of GP-based searches to identify complex transformations of a symbolic expression that effectively trade speed for accuracy. As in that work, we use a similar internal representation of analytical BRDFs and apply a similar set of expression transformations to enumerate candidate variants during the search.

Finally, recent work in the programming languages and compilers field has also demonstrated the promise of using Genetic Programming for identifying useful high-level code transformations. The pioneering work of Le Goues et al. [2012] uses a GP search to automatically identify and fix bugs in source code. Similar to our technique, their approach represents new candidate programs as an ordered list of edits with respect to some initial expression.

3 GP for Synthesizing New Analytical BRDFs

Our approach to synthesizing new analytical BRDFs is illustrated in Figure 2. Starting from an input analytical BRDF (or multiple BRDFs) we apply a series of expression transformations, or *mutations*, chosen at random to produce multiple BRDF *variants*. For example, a transformation may swap two sub-expressions in a BRDF or remove one of them altogether. Another transformation may swap sub-expressions across two different BRDF variants.

The next step is to *sort* the resulting variants based on their accuracy. This is accomplished by measuring the residual error after performing a standard non-linear optimization that determines the best fitting parameters of each variant with respect to a training set of several measured BRDFs. We use a fast approximation technique for a perceptual BRDF error metric [Pereira and Rusinkiewicz 2012] as the objective function in this non-linear optimization.

Lastly, we select a subset of these variants chosen according to the competing objectives of *exploring* a large portion of this challenging search terrain while *exploiting* favorable variants. We repeat these three steps, mutation–evaluation–selection, over several generations, until the error has converged, and output those BRDFs found to have the lowest residual error across all of the generations.

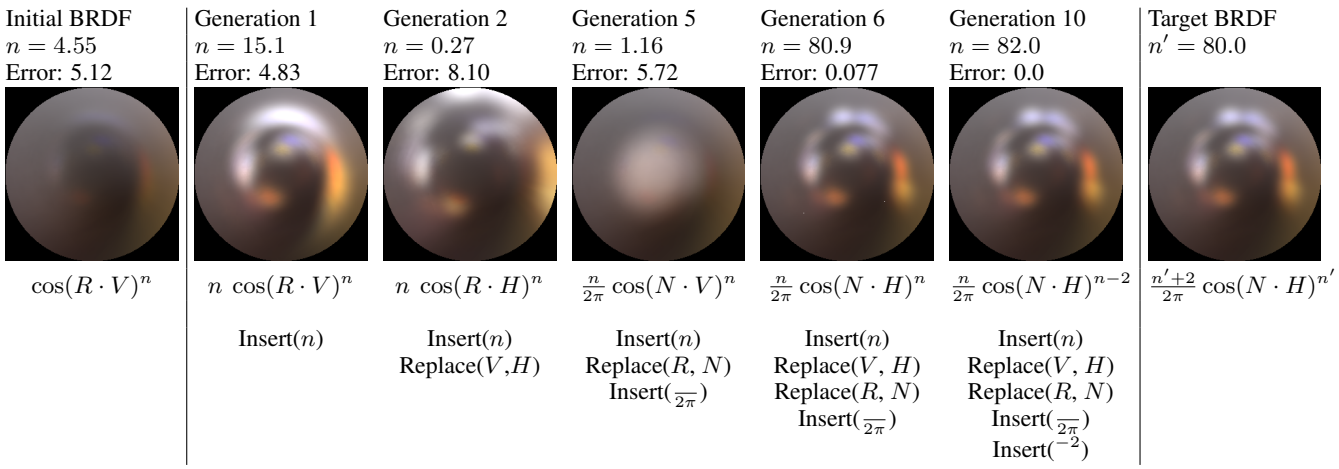

Figure 3: *Selected BRDFs from the example search in Section 3.1. Each BRDF is visualized on a sphere lit by a natural lighting environment. Below is the BRDF expression and the* genome *that produces the expression.*

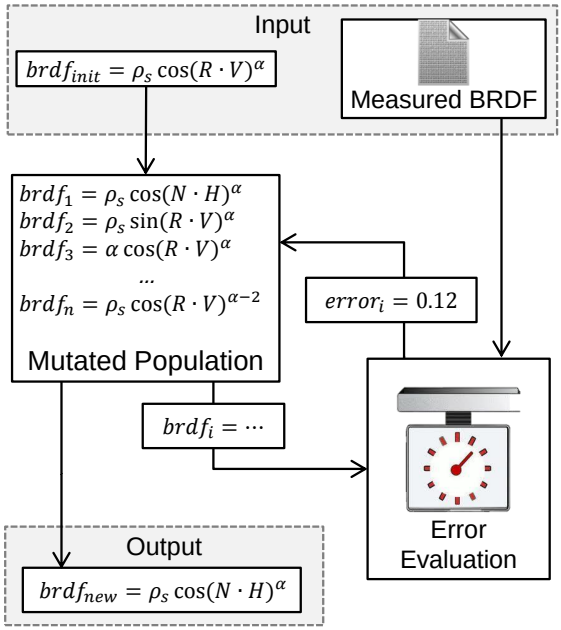

Figure 2: *Architecture diagram for our search.*

3.1 Example

Figure 3 provides a didactic example of our search. In this case, we attempt to synthesize the Blinn-Phong BRDF [Blinn 1977] from the original Phong BRDF [Phong 1975]. Recall that the Blinn-Phong BRDF uses the cosine of the half-angle, $(H \cdot N)$, to express the shape of the specular lobe whereas the original formulation uses the cosine of the angle between the view direction and perfect mirror reflection, $(R \cdot V)$.

We list the best performing variant in each generation below a rendered image that shows it on a sphere in the Grace lighting environment [Debevec 1998]. Below each analytical expression is an ordered list of the edits that were applied to the input BRDF to produce this variant. Above each rendered sphere image is the value of the parameter n found to produce the best match to the target BRDF. (The optimal value of n is determined using the Nelder-Mead simplex non-linear optimization algorithm [Nelder and Mead 1965].) We also list the value of the objective function that guides our GP search. Note that the error is not guaranteed to monotonically decrease from one generation to the next (this property of genetic programming favors exploration over exploitation).

Each generation of variants is produced by making a single edit to the variants in the previous generation. In the example shown, the best performing variant in the first generation contains the parameter n as a multiplicative factor in front of the expression. Note that this provides a better fit to the target BRDF (residual error of 4.83 versus 5.12). Recall that edits may bring in material from any of the other input analytic BRDFs, as in generation 2 when V is replaced by H.

Some variants in the nth generation with lower error are selected to form the $n + 1$th generation. This process repeats until a preset generation limit is reached. Notice how the generation 2 BRDF in Figure 3 includes the best edit from generation 1 as well a new edit that replaces the V vector with H. Note that this change does not lower the error of the BRDF: unlike strict hill-climbing approaches, genetic programming sometimes retains locally suboptimal individuals to allow for a greater exploration of the search space and the hope of eventually reaching a more global optimum. Thus, the ultimate sequence of edits, which reduces error, may include subsequences that increase error when considered alone.

The selected BRDFs from generations 2, 5, and 6 demonstrate why such higher error edits can be acceptable. The BRDF in generation 6 achieved low error by changing the both R and V vectors in the original BRDF to N and H. However, generations 2 and 5 show that changing only *one* of these vectors can increase error. By allowing a diverse population of variants, the search can ultimately discover better solutions by combining edits that are less desirable locally.

By generation 10 in Figure 3 we have learned an expression that perfectly fits the target measured BRDF. Note, however, that the expression learned is semantically but not syntactically identical to Blinn-Phong. The BRDF in generation 10 achieved zero error with $n = 82$, while our data were generation using $n' = 80$. However, since the BRDF learned is mathematically identical when $n' = n + 2$, the generation 10 BRDF perfectly matches the data. This shows that our search can not only come up with accurate solutions, but may do so in novel or unexpected ways.

Having demonstrated that a GP search has the potential to discover compact, low-error analytic BRDFs, we next formalize the details (e.g., mutation, selection, etc.) of our algorithm.

3.2 Algorithm Description

Our algorithm is an iterative, population-based, heuristically-guided search through the space of possible edits to an input BRDF. The goal is to find a sequence of edits that, when applied to the input BRDF, produce a new BRDF that minimizes the observed error with respect to measured reference data. The edits are performed at the abstract syntax tree level and can involve terms from the current individual or from other input BRDFs provided.

Our approach is population-based, in that it maintains a set of diverse variants. In particular, we use an *island model* genetic algorithm [Grosso 1985], a particular type of parallel genetic algorithm that maintains partitions the population into sub-populations that interact only rarely. The biological analogy is an archipeligo of small islands with creatures that normally evolve in isolation. Occasionally (e.g., low tides, strong winds, etc.), however, some individual from one island may *migrate* to another. This sub-population structure has the effect of favoring exploration over exploitation, as each island can diverge more freely without being corralled by the homogenizing effect of selection pressure within one population.

Step 1: Inputs

To begin, our algorithm takes as input (1) an initial target analytical BRDF expression; (2) an optional set of additional analytical BRDF expressions to serve as additional source material; and (3) a set of target measured materials. Measured materials consist of a discrete lookup table that maps incoming (L) and outgoing (V) light vectors to the BRDF response at those vectors $(L, V) \rightarrow brdf(L, V)$. Our algorithm is parametric with respect to a few implementation details (e.g., the size of the population or the number of generations to iterate before terminating, etc.).

Step 2: Creating the Initial Sub-Populations of BRDFs

Each element of the initial population of BRDFs is initialized to a genome containing a single random edit, and those elements are randomly partitioned into equal-sized sub-populations. An analytic BRDF expression can be reified from a genome by sequentially applying each of its edits to a copy of the original. Section Section 3.3 outlines our edit representation in detail.

Step 3: Evaluating the Variants

Next we evaluate how well each new BRDF can fit all target measured materials. We use a fast BRDF error computation detailed in Section 3.4 and a second nonlinear search outlined in Section 3.6 to learn parameter values that minimize fitting error a each material.

We handle multiple materials and color channels by fitting to each color channel in each material separately. We then combine the residual error for each fitting by taking the L^2 norm. This results in a single error value for each new BRDF in the population.

Step 4: Creating a New Population of BRDFs

We advance the search by creating a new generation of BRDFs from the current population. We favor exploration of the space of possible BRDFs (rather than exploitation of existing structures) by ensuring diversity in the new BRDF population. Thus, unlike strict gradient descent, we retain some less fit but diverse individuals into each new iteration.

For each island sub-population of size n, we select $n/2$ BRDFs to act as parents. We use the established tournament selection sampling technique [Miller et al. 1995] with $k = 8$ to select fit parents.

$\langle brdf \rangle$ ::= $\langle node \rangle$

$\langle node \rangle$ ::= $\langle op \rangle$ | $\langle scalar \rangle$

$\langle op \rangle$ ::= $\langle unaryOp \rangle$ ($\langle node \rangle$)
| $\langle binOp \rangle$ ($\langle node \rangle$, $\langle node \rangle$)
| $\langle vecOp \rangle$ ($\langle vector \rangle$ $\langle vector \rangle$)

$\langle unaryOp \rangle$ ::= - | sin | cos | tan | exp | asin | acos | atan | sqrt

$\langle binOp \rangle$ ::= + | - | * | / | max | min | pow

$\langle vecOp \rangle$ ::= dot

$\langle vector \rangle$ ::= $w_i \mid w_o \mid w_h \mid w_r \mid w_n$

$\langle scalar \rangle$ ::= $\langle parameter \rangle$
| $\langle const \rangle$
| $\langle component \rangle$

$\langle parameter \rangle$ ::= $p_0 \mid p_1 \mid p_2 \mid p_3$

$\langle component \rangle$ ::= $w_i.z \mid w_o.z \mid w_h.z \mid w_r.z$

$\langle const \rangle$ ::= 1.0 | 2.0 | π | 4.0

Figure 4: *Grammar describing all BRDFs expressed in our search*

Those parents are then randomly paired and create $n/2$ offspring, in the form of new BRDFs, using one-point crossover operator [Holland 1992] to produce two offspring. *Crossover* is a computational analog of biological operations that form offspring by combining material from parents. Intuitively, parent genomes (edit sequences) $\langle A_1, A_2 \rangle$ and $\langle B_1, B_2 \rangle$ produce offspring $\langle A_1, B_2 \rangle$ and $\langle B_1, A_2 \rangle$. Both the offspring and the parents then append a new random edit to their genome and are added to the new population. This last *mutation step* increases diversity and allows possible improvements to enter the population.

Each island sub-population is handled separately: two parents will always be chosen from the same island (thus favoring exploration).

Step 5: Iteration

Population creation and evaluation repeat for the given number of generations, after which the fittest BRDF found is returned.

Every kth generation, the top m variants from each island x migrate in a ring to island $x + 1$. Previous work in the genetic algorithm community has demonstrated that this simple migration topology performs well on a wide variety of problems [Andalon-Garcia and Chavoya-Pena 2012].

When the desired number of iterations have been performed, the algorithm terminates and reports a Pareto frontier of all variants discovered that represent non-dominated trade-offs between observed error and analytic BRDF size.

3.3 Representing Edits to Analytic BRDFs

We represent new analytical BRDFs using a series of edits to the abstract syntax tree (AST) [Aho et al. 1986] of an existing BRDF. The full grammar of expressible BRDFs is shown in Figure 4. This section outlines the types of edits that can be applied.

$Swap(\langle node \rangle_1, \langle node \rangle_2)$ – Given two nodes in the tree, replace $node_1$ with $node_2$ and $node_2$ with $node_1$. To avoid creating cy-

cles, no $Swap()$ edits are created where either node is a subtree of the other.

$Insert(\langle node\rangle_{new}, \langle node\rangle_{old})$ – Insert $\langle node\rangle_{new}$ at the location once occupied by $\langle node\rangle_{old}$. Rather than synthesize a node to insert, all inserted nodes are drawn from the *codebank*. The codebank is populated with subexpressions drawn from existing analytical BRDFs.

$Replace(\langle node\rangle_{old}, \langle node\rangle_{similar})$ – Unlike $Insert$, this edit replaces $\langle node\rangle_{old}$ with a similar node, leaving any children untouched. Specifically, a $\langle binOp\rangle$ is replaced by another $\langle binop\rangle$ production, a $\langle unaryOp\rangle$ with another $\langle unaryOp\rangle$ production, and a $\langle scalar\rangle$ is replaced with a different $\langle scalar\rangle$ production. For example, this edit could replace sin with cos, $+$ with $*$, or 2.0 with p_0.

$Delete(\langle node\rangle)$ – This operation replaces $\langle node\rangle$ with the constant 1.0.

3.4 BRDF Similarity Metric

This section outlines the BRDF similarity metric used by our search. The next section outlines how this new metric admits a large precomputation step, allowing for rapid evaluation of new BRDFs. Note that for an analytical BRDF, this metric assumes values have been chosen for each parameter. In Section 3.6 we detail how we choose these parameter values using this similarity metric.

Our new metric is approximates an image-based metric used by Pereira and Rusinkiweicz [Pereira and Rusinkiewicz 2012], where a BRDF is applied to a rendered sphere lit by an artificial lighting environment. While these renders are expensive, Pereira and Rusinkiweicz offer a fast analytical expression for their metric that admits a precomputation step. However, this fast precomputation requires all BRDFs be expressed as θ_h curves. Since we do not observe such a constraint, we have instead created an approximation to their metric that also admits a large precomputation step. This approximation enables our search explore hundreds of thousands of possible BRDFs within our computational budget.

In the non-analytic form of the Pereira and Rusinkweicz metric, each element in the L^2 norm corresponds to a pixel. When producing this value using Monte Carlo integration with importance sampling, each pixel is computed using a finite sum of light contributions. The contribution of this sample to the final pixel (L_o) can be expressed using the incident light, $L_I(w_i)cos(w_i)$, the BRDF of the material ($f(w_i, w_o)$), and the probability of choosing that sample $p(w_i)$ as shown in equation (1).

$$L_o(f_r, w_i, w_o) = \frac{f_r(w_i, w_o)L(w_i)cos\theta_i}{p(w_i)} \tag{1}$$

Equation (2) shows the non-analytic form of the Pereira nd Rusinkweicz metric whem comparing two BRDFs (f_1 and f_2). Since each pixel corresponds to a single value of ω_o, the metric can be expressed as a sum of squared pixel value differences , where each pixel is a sum of contributions for all ω_i.

$$S(f_1, f_2) = \sqrt{\sum_{\omega_o}\left(\sum_{\omega_i} L_o(f_1, w_i, w_o) - \sum_{\omega_i} L_o(f_2, w_i, w_o)\right)^2} \tag{2}$$

$$S'(f_1, f_2) = \sqrt{\sum_{\omega_o}\sum_{\omega_i}\left(L_o(f_1, w_i, w_o) - L_o(f_2, w_i, w_o)\right)^2} \tag{3}$$

Our new metric instead takes the L^2 norm of each incident light sample contribution as shown in (3). Rather than measure similarity by comparing aggregate reflectance at many ω_o (i.e., pixels), we compare individual reflectance events as modulated by the BRDF. Note that if two BRDFs observe high similarity using our metric, they should necessarily observe high similarity using image-based L^2.

3.5 Pre-computation Optimization for BRDF Similarity

In the last section, we presented our new image-based similarity metric and how it compares indidividual light samples contributions rather than converged pixel values. A benefit of this change is the ability to precompute large portions of this similarity metric and remove redundant evaluations of the BRDF.

The first step is to discretize the BRDF using half-angle coordinates [Rusinkiewicz 1998]. For isotropic materials, this reduces the BRDF dimensionality in spherical coordinate from four (θ_i, ϕ_i, θ_o, ϕ_i) to three (θ_h, 0, θ_d, ϕ_d). Next, we discretize θ_h, θ_d, ϕ_d using 1 degree divisions along θ_d and ϕ_d, with a nonlinear mapping for θ_h [Matusik et al. 2003]. We chose this binning resolution as publically available measured BRDF databases are tabulated using the same format.

We perform this precomputation by performing a single high-resolution render of the target scene with sufficient samples to ensure convergence of the image. For each light sample, we compute the equation as shown in Equation (1), but replate the BRDF with 1. Rather than compute the pixel value, this light contribution without the BRDF is recorded in a table indexed by its half angle representation.

The result is a table that maps a small, discrete region of the BRDF to the aggregate amount of light it would have contributed to the final image (4). Similarity is then computed by comparing each BRDF evaluated a (θ_h, 0, θ_d, ϕ_d) with the cached light value from the table ($T[\theta_h, \theta_d, \phi_d]$).

$$S'(f_1, f_2) = \sum_{(\theta_h, \theta_d, \phi_d)} \left(T[\theta_h, \theta_d, \phi_d]f_1 - T[\theta_h, \theta_d, \phi_d]f_2\right)^2 \tag{4}$$

This precomputation compresses an arbitrarily high resolution render with arbitrary many light samples to a fixed-length table. We remove any redundant BRDF calculations and still allow fast parallel execution on CPU or GPU hardware. While we are aggregating multiple light samples that index into the same discrete region of the BRDF, traditional image-based metrics perform a similar aggregation. The next section will show how this fast evaluation of BRDF similarity can be used to fit parameter values to arbitrary analytical BRDFs.

3.6 Fitting BRDFs to Measured Data

For each analytical BRDF produced by our search, we must fit a unique set of parameter values to each target measured BRDF. The residual errors of these fittings guide our GP search, making this a critical step in our process.

We fit parameters using a Nelder-Meade simplex search[Nelder and Mead 1965] to guided by our BRDF error metric. For a given material and RGB color channel, we randomly initialize each parameter then invoke the simplex search. Because we allow our search significant freedom in generating expressions, after the simplex search converges we repeat the random initialization and search multiple times before reporting the lowest error found. Once all color channels and materials have been fit, we take the L^2 norm of all color

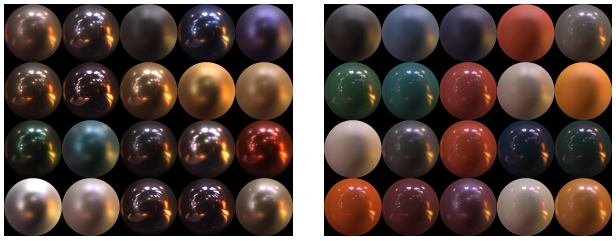

Metals Dielectrics

Figure 5: *Selected 40 metal and dielectric materials from the MERL BRDF database.*

channels and materials. The result is a single error value representing how well a BRDF fits the set of target measured BRDFs.

4 Experimental Setup

This section outlines how we will produce new BRDFs and evaluate them for accuracy, generality, and tunability. We will produce a new analytical BRDF for two material classes, metals and *dielectrics* (e.g., plastics, acrylics, and ceramics). For each material class we chose 20 isotropic, measured BRDFs from the 100 materials in the MERL MIT BRDF database (Figure 5). The materials chosen for each class were manually selected to represent a broad range of material properties.

$$brdf_{init} = k_d + k_s \frac{F}{\pi} \frac{DG}{(N \cdot L)(N \cdot V)} \tag{5a}$$

$$D = \frac{1}{m^2 cos^4 \delta} e^{-[tan(\delta)/m]^2} \tag{5b}$$

$$G = min\{1, \frac{2(N \cdot H)(N \cdot V)}{(V \cdot H)}, \frac{2(N \cdot H)(N \cdot L)}{(V \cdot H)}\} \tag{5c}$$

$$F = R_0 + (1 - R_0)(1 - (N \cdot H))^5 \tag{5d}$$

Each BRDF will be learned by modifying the Cook Torrance[Cook and Torrance 1982a] BRDF shown in (5). We will learn each new BRDF using 4 of the 20 materials in each material class. We chose Cook Torrance as it models this dataset better than other existing analytical BRDFs[Ngan et al. 2005]. We will test our new BRDF against Cook Torrance using the remaining 16 materials in each class. Each search will consists of 100 generations with a population size of 4096. We split each population into 4 islands of size 1024. Each island sends its best variant to one neighboring island every 5 generations. We allow each new BRDF to include at most four free parameters for each color channel.

We will demonstrate accuracy by comparing each new BRDF to Cook Torrance using three error metrics. The first metric will be the fast precomputed error metric used by our search. The remaining two metrics compare images of representative scenes rendered using either a measured or analytical BRDF. We will compare ground-truth renders produced from measured BRDF data to renders using either our new BRDF or Cook Torrance. We will compare images using both L^2 error in RGB space as well as the perceptually-based structural similarity metric (SSIM) [Wang et al. 2004].

We evaluate the generality by comparing how well each new BRDF fits the 16 materials not seen during the search. If our new BRDFs are general, we should observe acceptable levels of accuracy in fitting the remaining 16 materials.

Existing analytical BRDFs like Cook Torrance can also be manually adjusted to fine-tune a material's appearance. A tunable BRDF should produce a valid BRDF over a continuous range of values for each parameter. Each parameter should also offer some intuitive effect on the BRDF's appearance. We evaluate tunability by sweeping each parameter over an empirically determined range while keeping all other parameters fixed. We will show how valid BRDF is produced across this range, and offer a high-level description of each parameter's effect.

Each experiment was run on the Longhorn visualization cluster at the Texas Advanced Computing Center. We chose to use GPU cluster resources due to the parallel nature of our error evaluation and population evaluation. Each node on Longhorn contains two NVIDIA Quadro FX 5800 GPUs, 8 2.5GHz Intel Nehalem CPU cores, and 48GB of RAM. We evaluate 8 BRDFs simultaneously on a single node. Each experiment used approximately 1000 CPU-hours, where CPU-hours are measured per-core.

5 New Metal BRDF

The minimized form of our new Metal BRDF is shown in equation (6). Similar to Cook Torrance, our new BRDF contains four free parameters: k'_d, k'_s, α, and β. δ represents the angle between N and H. The expression consists of a diffuse color parameter (k'_d) and a specular term. The specular term consists of a specular color parameter (k'_s) multiplied by three subexpressions, D', G', and F'. Each of these subexpressions show several novel modifications to the Cook Torrance microfacet distribution (D), geometry shadowing (G), and fresnel (F) terms. Note that these modifications result in a more compact expression. Our new BRDF contains only 41 AST nodes compared to 84 for Cook Torrance.

While each subexpression roughly parallels Cook Torrance, each models the BRDF very differently. Unlike G in Cook Torrance, G' includes three tunable parameters. This allows the BRDF to control transitions between the left or right term in the min expression by adjusting the values of β and α. Such tunable selection of subexpressions does not exist in Cook Torrance. The D' term is also much simpler than the microfacet distribution term (D) in Cook Torrance, removing the normalization and exponent square terms. Finally, the F' contains two edits to Schlick's approximation for the Fresnel equations. Both the $(1 - R_0)$ subexpression is removed and the view vector V replaces the surface normal N.

$$brdf_m = \frac{k'_d}{\pi} + k'_s D' G' F' \tag{6a}$$

$$D' = e^{-\delta/\beta} \tag{6b}$$

$$G' = min\{\frac{\alpha}{k'_d}(V \cdot H), \frac{4\beta}{tan(\delta)(N \cdot V)(N \cdot L)}\} \tag{6c}$$

$$F' = \alpha + (1 - (V \cdot H))^5 \tag{6d}$$

5.1 Metal BRDF Accuracy

We learned this new metal BRDF by training on the Steel, Two Layer Gold, Green Metallic Paint, and Red Ornament materials from the MERL database. Our fits for these materials are shown in shown in Figure 7 with error reductions shown in Figure 6. In Figure 7, the center image of each row is a render of our target measured BRDF. We chose these materials as they represent a wide range of metal appearances.

The first row of spheres in Figure 7 compares our new BRDF to Cook Torrance in modeling the Steel BRDF. Note from the table

Metal BRDF Training Set

	L^2 Error			1-SSIM Error		
MERL Material	CT	New	%Impr	CT	New	%Impr
Steel	.083	.052	36.9%	.125	.036	71.3%
Two Layer Gold	.087	.028	67.9%	.077	.032	58.5%
Green Metallic Paint	.079	.040	49.4%	.078	.035	55.5%
Red Ornament	.096	.084	11.9%	.160	.072	54.8%
Average	.086	.051	41.5%	.110	.044	60.0%

Metal BRDF Testing Set

	L^2 Error			1-SSIM Error		
MERL Material	CT	New	%Impr	CT	New	%Impr
Tungsten Carbide	.082	.025	69.3%	.160	.016	90.0%
Blue Metallic Paint	.054	.024	55.8%	.043	.006	86.6%
Gold Metallic Paint	.098	.028	70.8%	.101	.015	85.1%
Silver Metallic Paint	.143	.038	73.6%	.117	.027	76.7%
Silver Paint	.098	.032	66.9%	.032	.008	75.2%
Aluminium	.092	.053	42.4%	.139	.037	73.3%
Nickel	.075	.039	48.3%	.094	.026	72.8%
Chrome	.071	.036	48.8%	.102	.032	68.5%
Gold-paint	.086	.027	68.5%	.032	.011	64.3%
Gold Ornament	.118	.086	27.2%	.149	.054	64.0%
Black Oxidized Steel	.061	.036	41.0%	.054	.029	47.1%
Hematite	.093	.087	6.0%	.192	.123	35.9%
Green Ornament	.064	.059	8.3%	.084	.060	29.3%
Blue Ornament	.077	.073	5.5%	.091	.066	27.0%
Brass	.081	.111	-37.3%	.134	.099	25.8%
Aluminum Bronze	.100	.066	34.4%	.107	.089	16.8%
Average	.087	.051	39.3%	.102	.044	58.6%

Figure 6: *The 4 materials used during training for our new metal BRDF.*

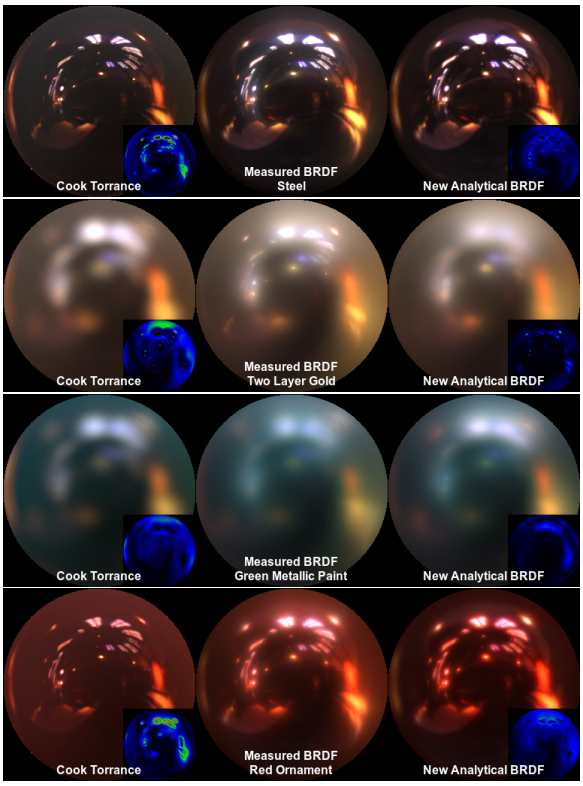

Figure 7: *Renders metal materials used during for learning our new metal BRDF.*

that our new BRDF reduces L^2 and $1-SSIM$ error by 36.9% and 71.3% respectively. We chose this material as it represents a common, simple, highly specular metal. Note that the measured BRDF includes a small amount of anisotropic noise around its specular highlights. This is due to inter-lens reflection during capture and should not be modeled. Qualitatively, notice how reflections in the Cook Torrance render appear darker and flatter. This is shown in the inset L^2 error plots, where Cook Torrance shows higher amounts of error (green and red areas) around bright reflections. Our new metal BRDF more accurately models the falloff of these reflections, resulting in lower error in the inset L^2 error plots (blue areas).

Two Layer Gold is a more complex BRDF, with both a highly specular and a diffuse reflection present. Materials with a thin translucent film layered over a highly specular surface can exhibit this behavior. While our new BRDF halves the image error in both metrics compared to Cook Torrance, the second highly specular reflection is not modeled by either BRDF. This is somewhat expected, as multi-layer materials are typically modeled using multiple specular lobes. We leave fitting multiple specular lobes of our new BRDF as future work.

The Green Metallic Paint BRDF represents a simpler diffuse material. However, the large specular lobe of this material falls off more gradually than the Cook Torrance BRDF. This is demonstrated by the window reflections in the upper quarter of the rendered spheres. In both the new and measured BRDFs, notice how the bright window reflections smoothly transition upward to light green. Cook Torrance, however, makes this transition much more sharply and to a darker shade of green. While our new BRDF has slightly brighter highlights, again it reduces both error metrics by about half.

The final Red Ornament material is highly specular, but falls off more gradually than the raw Steel BRDF. This manifests as an almost glowing appearance to the measured material. Note that this glow around the windows is absent in Cook Torrance, but captured by our new metal BRDF.

5.2 Metal BRDF Generality

A useful analytical BRDF should be generally applicative to a wide variety of materials. Figure 6 shows improvements over Cook Torrance when fitting 16 metal materials not seen during our search. All 16 testing materials show an improvement in 1-SSIM error compared to Cook Torrance, with 15 materials showing an improvement in L^2 error. We average improvements in 1-SSIM of 60% for the 4 training materials and 59% for testing materials with similar improvements of 41% and 39% for L^2 train and test. Our best improvement of 90% was observed in our testing set. These results offer strong evidence that the metal BRDF is generally applicable to a wide range of metal materials.

5.3 Tunability

Figure 8 shows the affect of adjusting each parameter in our new metal BRDF. Each sphere is rendered using the grace cathedral lighting environment. The middle sphere in each row is rendered using the median values for each parameter found when fitting the 20 dielectric materials from the MERL dataset ($k'_d = 0.01$, $k'_s = 4.0$, $\alpha = 0.05$, and $\beta = 0.05$). Each row sweeps a single parameter from $\frac{1}{4}$, $\frac{1}{2}$, 1, 2, and 4 times it's median value. Figure 8 shows how each parameter produces a valid, smooth transition in BRDF appearance.

We will now offer a qualitative description of each parameter. Both k'_d and k'_s bear a strong resemblence to the diffuse intensity (k_d) and specular intensity (k_s) terms in Cook Torrance. Interestingly, notice how reducing k'_d in the leftmost two images of the first row

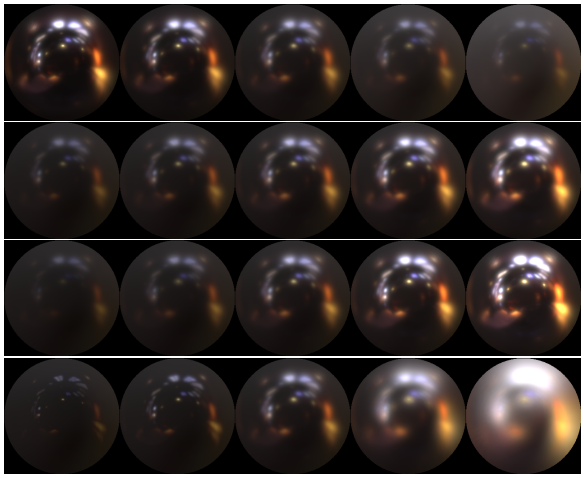

Figure 8: *A demonstration of manual parameter adjustment for the new metal BRDF. Each row of spheres shows the change in the BRDF appearance when manipulating a single parameter while keeping all others fixed.*

in Figure 8 causes an increase in specular intensity. This is due to the G' term in equation (6) including k'_d. Modifying α shows a similar effect as k'_s, but serves a dual-role in helping k'_d control the left portion of G'. Finally, β plays a strong role in determining roughness in the BRDF, but also modulates intensity. We speculate that these nuanced relationships between parameters play a key role in achieving a high degree of fitting flexibility and accuracy.

6 New Dielectric BRDF

Our new BRDF for modeling dielectrics is shown in (7). It also includes four free parameters, k'_d, k'_s, α, and β with δ representing the angle between N and H. Like Cook Torrance and our new Metal BRDF, at a high level this new expression consists of a diffuse term and specular term modulated by a color parameters k'_d and k'_s. However, compared to our new metal BRDF, this BRDF shows significantly more novel modifications while also including unmodified Cook Torrance terms. The diffuse term now includes the unique subexpression R_d, and the specular contribution can be broken into two lobes, S_1 and S_2. Like our metal BRDF, we observe tunable parameters inside a min. Unlike our metal BRDF, right min subexpression is much more complex, including the unmodified Cook Torrance F and D terms parameterized by α and β.

$$brdf_d = k'_d R_d + S_1 S_2 \tag{7a}$$

$$R_d = (1 + cos(\delta)^{1-\alpha}) \tag{7b}$$

$$S_1 = min\{\frac{1}{\sqrt{k'_d}}, \frac{F(\alpha, N \cdot H) D(\beta)}{\pi (V \cdot H)(N \cdot L)(N \cdot V)}\} \tag{7c}$$

$$S_2 = \frac{e^{-tan(\delta)/k'_s}}{\alpha^2(\alpha+2)\beta} \tag{7d}$$

It is interesting to note that each novel term in this expression is essential for fitting success. Empirical attempts to remove or simplify subexpressions such as the $\alpha^2(\alpha+2)\beta$ normalization resulted in significantly reduced fitting performance. While a mathematician may not choose such a normalization, our GA has found a solution that does not require additional parameters and (we will show) reduces fitting error.

MERL Material	L^2 Error CT	L^2 Error New	L^2 Error %Impr	1-SSIM Error CT	1-SSIM Error New	1-SSIM Error %Impr
	Dielectric BRDF Training Set					
Maroon Plastic	.074	.019	74.8%	.044	.006	85.6%
PVC	.074	.026	65.3%	.063	.010	83.4%
Green Acrylic	.074	.021	71.8%	.098	.019	81.1%
Dark Blue Paint	.051	.017	66.5%	.036	.013	62.6%
Training Set Average	.068	.021	69.6%	.060	.012	78.2%
	Dielectric BRDF Testing Set					
Green Phenolic	.055	.013	75.9%	.090	.007	92.1%
Blue Phenolic	.053	.014	73.4%	.098	.008	91.5%
Violet Phenolic	.059	.016	72.2%	.058	.006	89.4%
Green Plastic	.064	.025	60.7%	.061	.007	88.2%
Gray Plastic	.077	.025	67.7%	.048	.009	80.8%
Orange Phenolic	.080	.032	59.7%	.046	.009	79.7%
Dark Red Paint	.083	.023	71.8%	.025	.006	76.7%
Pure Rubber	.121	.030	74.8%	.026	.007	74.4%
Orange Paint	.097	.043	55.9%	.023	.007	70.5%
Black Soft Plastic	.075	.025	66.7%	.066	.020	70.3%
Neoprene Rubber	.101	.033	66.8%	.026	.008	67.6%
Blue Rubber	.073	.015	79.1%	.026	.010	63.0%
Violet Acrylic	.085	.061	27.9%	.078	.031	60.6%
Yellow Matte Plastic	.088	.036	58.7%	.038	.018	52.0%
White Acrylic	.105	.038	64.1%	.046	.023	50.1%
Red Specular Plastic	.071	.035	51.1%	.047	.030	36.0%
Test Set Average	.080	.029	64.1%	.050	.013	71.4%

Figure 9: *The 4 materials used during training for our new dielectric BRDF.*

6.1 Accuracy

We learned this new BRDF on the Dark Blue Paint, Green Acrylic, Maroon Plastic, and PVC materials from the MERL BRDF database. Figure 10 shows renders of each of these materials using Cook Torrance, measured MERL data, and our new dielectric BRDF. Improvements in L^2 and $1 - SSIM$ error for these fits are shown in Figure 9. For the four training materials we reduce L^2 and $1 - SSIM$ error by 70% and 78% respectively.

The first row of spheres in Figure 10 compares our BRDF to Cook Torrance in fitting diffuse dielectric materials. Notice the left edge the Cook Torrance sphere includes an orange glow from fresnel reflectance that is absent in both our new BRDF and the measured data. Also note how the falloff of the specular highlights in the upper edge of the sphere are better represented in our BRDF.

The Green Acrylic and Maroon Plastic in rows 2 and 3 of Figure 10 represent two subclasses of highly specular dielectrics. In both materials, notice how the brightness of the specular window reflections are better matched using our new BRDF. Note again how the Cook Torrance fit misrepresents fresnel reflectance at grazing angles.

Of the 20 dielectric materials chosen from the MERL dataset, most are either highly specular or highly diffuse. We chose the PVC material (4th row in Figure 10) as it represents a dielectric in the middle of these extremes.

6.2 Generality

The results for our generality analysis are shown in Figure 9. For all 20 dielectric materials our new dielectric BRDF improves over Cook-Torrance, with some test set materials showing error values one tenth that of Cook Torrance. Using SSIM, we improve the materials seen during our search by 78.2%, with held out materials

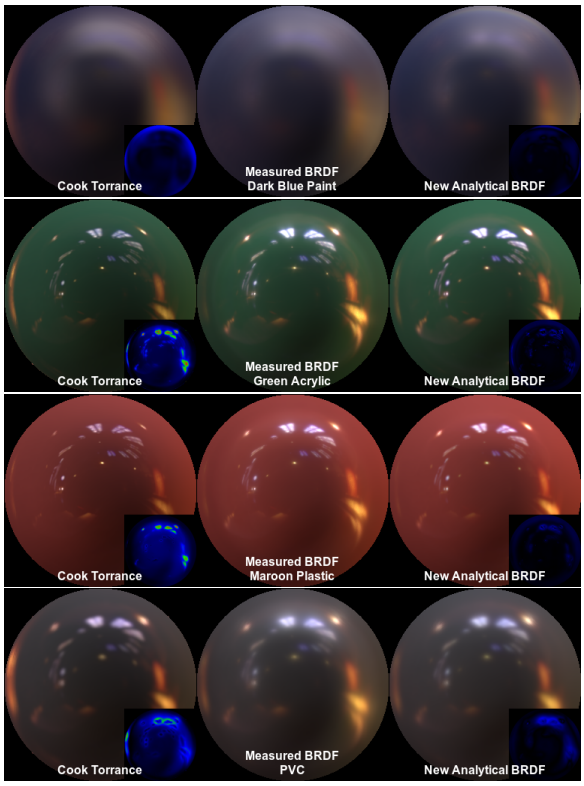

Figure 10: *Renders of dielectric materials used for learning our new dielectric BRDF.*

showing an average improvement of 71.4%. Similar average improvements of 69.6% and 64.1% were seen using L^2 error. These results offer strong support that our dielectric BRDF does not overfit to the materials seen during testing.

6.3 Tunablity

Our parameter sweep for the new dielectric BRDF is shown in Figure 11. Each sphere is rendered using the grace cathedral lighting environment. The middle sphere in each row is rendered using the median values for each parameter found when fitting the 20 dielectric materials from the MERL dataset ($k_d' = 0.007$, $k_s' = 0.016$, $\alpha = 0.16$, and $\beta = 1.03$). Each row sweeps a single parameter from $\frac{1}{4}$ to 4 times its median value while holding the other 3 parameters fixed. Note that a valid BRDF is produced across all sweeps.

Since our dielectric BRDF contains many novel modifications to Cook Torrance, the effects of adjusting each parameter show similar novelty. k_d' is most similar to Cook Torrance, modulating the diffuse color contribution of the sphere withing affecting specular intensity. Like Cook Torrance, the remaining parameters each adjust the specular lobe intensity. Increasing k_s' increases specular highlight intensity while also increasing roughness. Increasing the remaining parameters α and β reduce specular intensity, but show a less pronounced effect on roughness.

7 Limitations and Future Work

Our technique relies on the existance of measured BRDF data to produce new BRDFs. While we have access to 100 materials from the public MERL MIT BRDF database, capturing additional materials can require significant equipment and time commitments. However, we believe that current databases already contain a broad range of materials suitable learning BRDFs to fit broad classes of materials. Dielectrics and metals are both well-represented, which encompass many commong materials. While we did not attempt to fit materials such as fabrics or woods, there is nothing inherit to our fitting process that would limit our ability to produce BRDFs for these material classes.

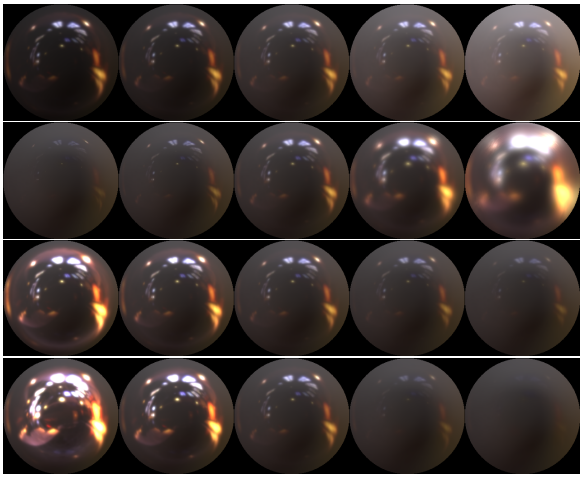

Figure 11: *A demonstration of manual parameter adjustment for the new metal BRDF. Each row of spheres shows the change in the BRDF appearance when manipulating a single parameter while keeping all others fixed.*

Our approach also only considers isotropic, opaque materials. We chose to only focus on these types of materials due to the wealth of available data, but also our ability to quickly precompute an error evaluation for isotropic materials. This simplifying assumption allowed us to test our technique by evaluating more candidate BRDFs. Expanding our technique to synthesize anisotropic BRDFs is left as future work. As long as we can reliably fit parameters to the candidate expression, our search can produce new BRDFs of arbitrary dimensionality.

For future work, we observe that our grammar for constructing new BRDFs could be expanded to include the creation of arbitrary vectors. Our grammar could be trivially changed to allow the creation, addition, subtraction, and cross product of arbitrary vectors. While we include all vectors used by standard analytical BRDFs, allowing our search to create new vectors may lead to additional discoveries. For example, if we did not include the half-angle vectors in our search, we could not produce the Blinn Phong BRDF shown in our example.

8 Conclusion

We have shown that our technique can automatically produce new analytical BRDFs that match measured data better than any previous analytical BRDF. We have presented two new analytical BRDFs for metals and dielectrics that are compact, adjustable, and generalizable to a broad set of materials. We have shown that these BRDFs obey standard BRDF properties such as reciprocity and non-negativity. A key insight into this success is that existing analytical BRDFs can come close to fitting many materials, but require edits at the expression level to further improve fitness. Using GP search we are able to leverage these partial solutions and search a diverse population of candidate BRDFs to find a solution that best

matches our visual perception. We evaluate these candidates using a fast BRDF similarity metric for isotropic materials that leverages a precompution optimization to allow our search to evaluate hundreds of thousands of materials.

References

AHO, A. V., SETHI, R., AND ULLMAN, J. D. 1986. *Compilers principles, techniques, and tools*. Addison-Wesley, Reading, MA.

ANDALON-GARCIA, I. R., AND CHAVOYA-PENA, A. 2012. Performance comparison of three topologies of the island model of a parallel genetic algorithm implementation on a cluster platform. In *CONIELECOMP*, 1–6.

ASHIKHMIN, M., AND SHIRLEY, P. 2000. An anisotropic phong brdf model. *J. Graph. Tools 5*, 2 (Feb.), 25–32.

BECKMANN, P., AND SPIZZICHINO, A. 1963. *The scattering of electromagnetic waves from rough surfaces*. International series of monographs on electromagnetic waves. Pergamon Press; [distributed in the Western Hemisphere by Macmillan, New York].

BLINN, J. F. 1977. Models of light reflection for computer synthesized pictures. *SIGGRAPH Comput. Graph. 11*, 2 (July), 192–198.

COOK, R. L., AND TORRANCE, K. E. 1982. A reflectance model for computer graphics. *ACM Trans. Graph. 1*, 1 (Jan.), 7–24.

COOK, R., AND TORRANCE, K. 1982. A reflectance model for computer graphics. *ACM Transactions on Graphics (TOG) 1*, 1, 7–24.

DEBEVEC, P. 1998. Rendering synthetic objects into real scenes: bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques*, ACM, New York, NY, USA, SIGGRAPH '98, 189–198.

DORSEY, J., RUSHMEIER, H., AND SILLION, F. 2008. *Digital Modeling of Material Appearance*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.

GROSSO, P. B. 1985. *Computer simulations of genetic adaptation: parallel subcomponent interaction in a multilocus model*. PhD thesis, Ann Arbor, MI, USA. AAI8520908.

HE, X. D., TORRANCE, K. E., SILLION, F. X., AND GREENBERG, D. P. 1991. A comprehensive physical model for light reflection. *SIGGRAPH Comput. Graph. 25*, 4 (July), 175–186.

HOLLAND, J. H. 1992. *Adaptation in natural and artificial systems*. MIT Press, Cambridge, MA, USA.

LAWRENCE, J., RUSINKIEWICZ, S., AND RAMAMOORTHI, R. 2004. Efficient brdf importance sampling using a factored representation. *ACM Trans. Graph. 23*, 3 (Aug.), 496–505.

LE GOUES, C., DEWEY-VOGT, M., FORREST, S., AND WEIMER, W. 2012. A systematic study of automated program repair: fixing 55 out of 105 bugs for $8 each. In *Proceedings of the 2012 International Conference on Software Engineering*, IEEE Press, Piscataway, NJ, USA, ICSE 2012, ACM, 3–13.

MARSCHNER, S. R., WESTIN, S. H., LAFORTUNE, E. P. F., TORRANCE, K. E., AND GREENBERG, D. P. 1999. Image-based brdf measurement including human skin. In *Proceedings of the 10th Eurographics conference on Rendering*, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, EGWR'99, 131–144.

MARSCHNER, S. R., JENSEN, H. W., CAMMARANO, M., WORLEY, S., AND HANRAHAN, P. 2003. Light scattering from human hair fibers. In *ACM SIGGRAPH 2003 Papers*, ACM, New York, NY, USA, SIGGRAPH '03, 780–791.

MARSCHNER, S. R., WESTIN, S. H., ARBREE, A., AND MOON, J. T. 2005. Measuring and modeling the appearance of finished wood. *ACM Trans. Graph. 24*, 3 (July), 727–734.

MATUSIK, W., PFISTER, H., BRAND, M., AND MCMILLAN, L. 2003. A data-driven reflectance model. In *ACM SIGGRAPH 2003 Papers*, ACM, New York, NY, USA, SIGGRAPH '03, 759–769.

MILLER, B. L., MILLER, B. L., GOLDBERG, D. E., AND GOLDBERG, D. E. 1995. Genetic algorithms, tournament selection, and the effects of noise. *Complex Systems 9*, 193–212.

NELDER, J. A., AND MEAD, R. 1965. A simplex method for function minimization. *The computer journal 7*, 4, 308–313.

NGAN, A., DURAND, F., AND MATUSIK, W. 2005. Experimental analysis of brdf models. In *Proceedings of the Eurographics Symposium on Rendering*, Eurographics CHANGEME, 117–226.

NICODEMUS, F. E. 1965. Directional reflectance and emissivity of an opaque surface. *Appl. Opt. 4*, 7 (Jul), 767–773.

OREN, M., AND NAYAR, S. K. 1994. Generalization of lambert's reflectance model. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, ACM, New York, NY, USA, SIGGRAPH '94, 239–246.

PEERS, P., VOM BERGE, K., MATUSIK, W., RAMAMOORTHI, R., LAWRENCE, J., RUSINKIEWICZ, S., AND DUTRÉ, P. 2006. A compact factored representation of heterogeneous subsurface scattering. In *ACM SIGGRAPH 2006 Papers*, ACM, New York, NY, USA, SIGGRAPH '06, 746–753.

PEREIRA, T., AND RUSINKIEWICZ, S. 2012. Gamut mapping spatially varying reflectance with an improved brdf similarity metric. *Comp. Graph. Forum 31*, 4 (June), 1557–1566.

PHONG, B. T. 1975. Illumination for computer generated pictures. *Commun. ACM 18*, 6 (June), 311–317.

RUSINKIEWICZ, S. M. 1998. A new change of variables for efficient brdf representation. In *In Eurographics Workshop on Rendering*, 11–22.

SITTHI-AMORN, P., MODLY, N., WEIMER, W., AND LAWRENCE, J. 2011. Genetic programming for shader simplification. *ACM Trans. Graph. 30*, 6 (Dec.), 152:1–152:12.

TORRANCE, K. E., AND SPARROW, E. M. 1992. Radiometry. Jones and Bartlett Publishers, Inc., USA, ch. Theory for off-specular reflection from roughened surfaces, 32–41.

WANG, Z., BOVIK, A. C., SHEIKH, H. R., AND SIMONCELLI, E. P. 2004. Image quality assessment: From error visibility to structural similarity. *IEEE TRANSACTIONS ON IMAGE PROCESSING 13*, 4, 600–612.

WARD, G. 1992. Measuring and modeling anisotropic reflection. In *ACM SIGGRAPH Computer Graphics*, vol. 26, ACM, 265–272.