

genBRDF: Discovering New Analytic BRDFs with Genetic Programming

Adam Brady¹ Jason Lawrence¹ Pieter Peers² Westley Weimer¹

¹University of Virginia ²The College of William & Mary



Figure 1: A comparison between Cook-Torrance [1982], L w et al.’s smooth surface model [2012], and a new analytic BRDF model found by genBRDF, all fit with L w et al.’s E_2 metric to three materials from the MERL-MIT BRDF database [Matusik et al. 2013].

Abstract

We present a framework for learning new analytic BRDF models through Genetic Programming that we call genBRDF. This approach to reflectance modeling can be seen as an extension of traditional methods that rely either on a phenomenological or empirical process. Our technique augments the human effort involved in deriving mathematical expressions that accurately characterize complex high-dimensional reflectance functions through a large-scale optimization. We present a number of analysis tools and data visualization techniques that are crucial to sifting through the large result sets produced by genBRDF in order to identify fruitful expressions. Additionally, we highlight several new models found by genBRDF that have not previously appeared in the BRDF literature. These new BRDF models are compact, and more accurate than current state-of-the-art alternatives.

Keywords: BRDF, isotropic, analytic, genetic programming

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 common simplifying assumptions. For this reason, deriving BRDF models that are accurate and compact enough to meet the demands of modern rendering systems, while remaining general enough to express a range of interesting materials, remains an open problem.

This paper focuses on homogeneous opaque surfaces whose appearance is characterized by the Bidirectional Reflectance Distribution Function (BRDF) [Nicolodemos et al. 1977]. Traditionally, BRDF models have been derived manually either according to some phenomenological process [Phong 1975] or based on the physics of light propagation [Cook and Torrance 1982; He et al. 1991]. A more recent trend has been to measure the BRDFs of physical samples and use those measurements either directly [Matusik et al. 2003], or as input to an optimization process that determines the best fitting parameters of an analytic model [Ngan et al. 2005].

Analytic BRDFs are desirable for their compactness and because they often have adjustable parameters that designers can use to specify 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 analytic BRDF models and measured data in the case of three material samples from the MERL-MIT database [Matusik et al. 2003]. Note how the Cook-Torrance model overblurs the high-

lights, while Löw et al.’s model oversharpens them. Deriving new analytic BRDF models or even modifying existing ones in a principled way to achieve a better match is a difficult task as it is hard to pinpoint which component of the BRDF model needs improvement.

This paper presents a framework for discovering new analytic BRDF models. This approach can be seen as an extension of traditional methods that rely either on a phenomenological or empirical process. Our technique augments the human effort involved in deriving mathematical expressions that accurately characterize complex high-dimensional reflectance functions through Genetic Programming (GP). We perform a large-scale heuristic-based random search that considers hundreds of thousands of symbolic transformations of a few “seed” BRDF expressions with the goal of adapting them to achieve a more accurate fit to measured data. The result of this search is a large set (i.e., thousands) of analytic expressions that trade accuracy for complexity. We introduce analysis techniques and visualization tools to help identify fruitful BRDF expressions within these large result sets.

This seemingly inefficient “blind” approach to a staggeringly difficult optimization problem like learning new analytic BRDF models has only recently become feasible due to the availability of inexpensive large-scale computing resources. Our approach to enumerating potential analytic BRDF expressions also incorporates a number of domain-specific constraints that limit the search space. Finally, we present several experiments that learn different-sized subterms of the BRDF (e.g., the entire specular term vs. only the microfacet normal distribution function) in order to study the trade-off between the size of the search space and the computational effort required to identify useful reflectance models.

The main contribution of this paper is not just another BRDF model (or set of new BRDF models), but rather a framework capable of learning superior analytic BRDF models. We demonstrate the versatility of genBRDF by synthesizing new analytic BRDF models for isotropic materials, relying on the MERL-MIT BRDF database [Matusik et al. 2003] to guide the search. We evaluate a selected subset of these models and show that they outperform the current state-of-the-art. Additionally, we analyze one of these models in detail to illustrate how our genetic search works and the types of expressions it is able to discover. Finally, we apply genBRDF to the more focused task of synthesizing a new BRDF model for metals and, lastly, for one specific metal from the MERL-MIT database. We believe genBRDF will serve as a valuable tool in the development of future analytic models for material reflectance and, more broadly, the many other complex high-dimensional functions that arise in graphics.

2 Related Work

Modeling Surface Reflectance The manner in which light is scattered at a point on an opaque surface is fully described by the Bidirectional Reflectance Distribution Function (BRDF) [Nicodemus et al. 1977], a function of 4 variables that relates incident irradiance to exitant radiance. Many analytic BRDF models have been proposed which range from phenomenological to physically-based. On one end of the spectrum are pure phenomenological models, such as the well known Phong BRDF model [1975], that aim to reproduce the qualitative aspects of material appearance without any regard to the underlying physical process of light-matter interaction. On the other end of the spectrum are physically-based models, such as the models by He et al. [1991] and Torrance and Sparrow [1992]. These are based on a theory of light propagation (e.g., geometric optics, wave-based, etc.) and an idealized model of the material surface’s microgeometry. Furthermore, physically-

based models are often supplemented with phenomenological components in order to express certain underlying physical processes that are difficult to model exactly [Ashikhmin et al. 2000]. The current paper can be seen as an automated large-scale phenomenological approach to BRDF modeling. Our genetic search identifies mathematical expressions chosen at random that match measured data better than alternative expressions without consideration of the physics of light propagation or the material’s underlying physical structure.

The availability of large measured datasets of surface reflectance [Matusik et al. 2003] has enabled experimental validation of existing analytic BRDF models [Ngan et al. 2005] and fueled the development of new BRDF models (e.g., [Löw et al. 2012; Bagher et al. 2012; Walter et al. 2007]). Our method uses the MERL-MIT database of measured isotropic reflectance functions [Matusik et al. 2003] to guide a large-scale genetic search for better analytic BRDF expressions. Unlike prior work, we do not propose a single specific model, but rather a new approach for learning new BRDF models through Genetic Programming.

Genetic Programming Genetic Programming (GP) is a general machine learning strategy that uses an evolutionary methodology to search for a set of instructions (i.e., model) that optimize some fitness criterion [Koza 1992]. From an initial population of candidate models, a new evolved generation of candidates is produced through a series of pre-defined mutation and combination operations (i.e., cross-over). After each generation, a new population is chosen that favors the fittest candidates and the process repeats.

Although GP has been applied to many problems in computer science, it has not yet been used extensively in computer graphics. A notable exception is recent work by Sithi-Amorn et al. [2011] who described a GP approach to the problem of automatic procedural shader simplification. We use a similar internal representation of analytic BRDF expressions, but instead of simplifying mathematical expressions, we attempt to grow new ones. Additionally, we use an *island model* genetic algorithm [Grosso 1985] which is designed to favor exploring the space of possible BRDF variants over narrowly searching within profitable regions. We found this approach to be more suitable for this more challenging synthesis problem.

3 Overview

We apply Genetic Programming (GP) to the task of learning new analytic expressions for real-world surface reflectance functions. At a high-level our algorithm iterates between the following two steps:

- *Expansion:* Given a population of analytic BRDF expressions, compute a new population by randomly mutating individual expressions and randomly combining parts of multiple expressions.
- *Selection:* Sort the resulting population based on a fitness score and choose a subset that favors the fittest to survive into the next generation. Fitness is equal to the residual error after fitting the free parameters of each BRDF variant to a training set of measured materials.

The initial population is constructed from one or more existing analytic BRDF models and the search is terminated after a fixed number of generations.

Figure 2 provides a didactic example of how our genetic search works when the starting point is the original (“R dot V”) Phong BRDF model [Phong 1975] and the target (training set) is the (“H dot N”) Blinn-Phong BRDF model [Blinn 1977]. This sequence highlights a single variant within each generation, visualized as a

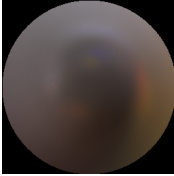
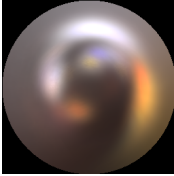
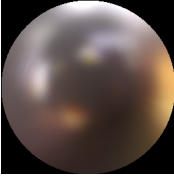

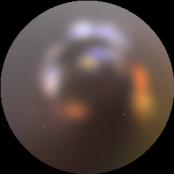
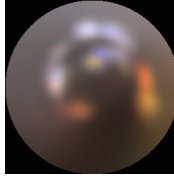
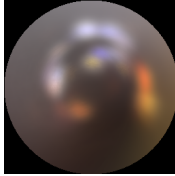
| Initial BRDF | Generation 1 | Generation 2 | Generation 5 | Generation 6 | Generation 10 | Target BRDF |
|---|---|---|---|---|---|---|
| $n = 4.55, \text{Err.: } 5.12$ | $n = 15.1, \text{Err.: } 4.83$ | $n = 0.27, \text{Err.: } 8.10$ | $n = 1.16, \text{Err.: } 5.72$ | $n = 80.9, \text{Err.: } 0.077$ | $n = 82.0, \text{Err.: } 0.0$ | $n' = 80.0$ |
|  |  |  |  |  |  |  |
| $\cos(R \cdot V)^n$ | $n \cos(R \cdot V)^n$ | $n \cos(R \cdot H)^n$ | $\frac{n}{2\pi} \cos(N \cdot V)^n$ | $\frac{n}{2\pi} \cos(N \cdot H)^n$ | $\frac{n}{2\pi} \cos(N \cdot H)^{n-2}$ | $\frac{n'+2}{2\pi} \cos(N \cdot H)^{n'}$ |
| | Insert(n) | Insert(n) Replace(V, H) | Insert(n) Replace(R, N) Insert($\frac{2\pi}{n}$) | Insert(n) Replace(V, H) Replace(R, N) Insert($\frac{2\pi}{n}$) | Insert(n) Replace(V, H) Replace(R, N) Insert($\frac{2\pi}{n}$) Insert(-2) | |

Figure 2: A didactic example of *genBRDF* shows how the original (“ $R \cdot V$ ”) Phong BRDF model [Phong 1975] is evolved into the (“ $H \cdot N$ ”) Blinn-Phong BRDF model [Blinn 1977]. For selected variants and generations, a visualization of the resulting BRDF expression is shown (top) as well as the sequence of edits (with respect to the initial BRDF expression) that produced the variant (bottom).

sphere rendered in a natural lighting environment above the corresponding analytic expression (we omit the diffuse and specular parameters for brevity). The value of the parameter n found to give the best fit to the training data along with the residual error of that best fit are reported above the sphere images.

Note that the error does not monotonically decrease over consecutive generations (e.g., the second generation). Unlike strict hill climbing approaches, GP intentionally retains a fraction of suboptimal variants to allow for greater *exploration* of the search space in the hopes of eventually reaching a better local optimum. Thus, the final sequence of edits, which reduce the error to zero, often include subsequences that increase error when considered alone. (Note that the solution found in generation 10 is identical to the target with the substitution $n = n' + 2$.)

4 Genetic Algorithm

Our algorithm is an iterative, population-based, heuristically-guided search through the space of BRDF expressions. The goal is to find new BRDF variants that minimize an error function with respect to a training set of measured reflectance functions.

Analytic BRDF Representation We represent an analytic BRDF by its Abstract Syntax Tree (AST) [Aho et al. 1986] using the grammar shown in Figure 3. This grammar was designed to capture constants and mathematical operators commonly found in existing BRDFs. The unary operators ‘pos’ and ‘clamp’ are defined as:

$$\text{pos}(x) \equiv \min(x, 0)/x \quad (1)$$

$$\text{clamp}(x) \equiv \min(\max(x, -1), +1). \quad (2)$$

Of special note is that our grammar also includes a set of n free BRDF parameters ($p_0 \dots p_{n-1}$) akin in nature to the sharpness parameter in Blinn-Phong [1977], roughness in Cook-Torrance [Cook and Torrance 1982], or R_0 in Schlick’s approximation of Fresnel reflectance [Schlick 1994].

Edit Operations Edits (or mutations) are performed at the AST level and can involve multiple terms from a single BRDF expression or several terms from multiple BRDF expressions. Instead of explicitly tracking mutated ASTs, we record a sequences of edits applied to a starting BRDF expression. For instance, in the example in Figure 2, the list of edits with respect to the initial Phong BRDF model (shown at the bottom) would be stored for each variant.

| | | |
|------------------------------------|-----|--|
| $\langle \text{brdf} \rangle$ | ::= | $\langle \text{node} \rangle$ |
| $\langle \text{node} \rangle$ | ::= | $\langle \text{op} \rangle \mid \langle \text{scalar} \rangle$ |
| $\langle \text{op} \rangle$ | ::= | $\langle \text{unaryOp} \rangle (\langle \text{node} \rangle)$ $\mid \langle \text{binOp} \rangle (\langle \text{node} \rangle, \langle \text{node} \rangle)$ $\mid \langle \text{vecOp} \rangle (\langle \text{vector} \rangle, \langle \text{vector} \rangle)$ |
| $\langle \text{unaryOp} \rangle$ | ::= | $- \mid \sin \mid \cos \mid \tan \mid \exp \mid \text{asin} \mid \text{acos} \mid \text{atan} \mid \text{sqrt} \mid \text{pos} \mid \text{clamp}$ |
| $\langle \text{binOp} \rangle$ | ::= | $+ \mid - \mid * \mid / \mid \max \mid \min \mid \text{pow}$ |
| $\langle \text{vecOp} \rangle$ | ::= | dot |
| $\langle \text{vector} \rangle$ | ::= | $L \mid V \mid H \mid R \mid N$ |
| $\langle \text{scalar} \rangle$ | ::= | $\langle \text{parameter} \rangle$ $\mid \langle \text{const} \rangle$ $\mid \langle \text{component} \rangle$ |
| $\langle \text{parameter} \rangle$ | ::= | $p_0 \mid p_1 \mid \dots \mid p_{n-1}$ |
| $\langle \text{component} \rangle$ | ::= | $L_z \mid V_z \mid H_z \mid R_z$ |
| $\langle \text{const} \rangle$ | ::= | $1.0 \mid 2.0 \mid \pi \mid 4.0$ |

Figure 3: The grammar used to generate analytic BRDF expressions in our genetic search.

We allow the following edits:

- **Swap($\langle \text{node} \rangle_1, \langle \text{node} \rangle_2$):** Given two nodes in the AST, replace node_1 with node_2 and node_2 with node_1 . To avoid creating cycles, no *Swap()* edits are created when either node is a descendant of the other.
- **Insert($\langle \text{node} \rangle_{\text{new}}, \langle \text{node} \rangle_{\text{old}}$):** Insert $\langle \text{node} \rangle_{\text{new}}$ at the location once occupied by $\langle \text{node} \rangle_{\text{old}}$. Inserted nodes are drawn from a codebook of sub-expressions created from existing analytic BRDF models. In this way, our search leverages expressions previously found to be useful in modeling BRDFs while also exploring the space of possible formulations.
- **Replace($\langle \text{node} \rangle_{\text{old}}, \langle \text{node} \rangle_{\text{similar}}$):** Unlike *Insert*, this edit replaces $\langle \text{node} \rangle_{\text{old}}$ with a similar node, leaving any chil-

dren untouched. Specifically, a $\langle binOp \rangle$ is replaced by another $\langle binOp \rangle$, a $\langle unaryOp \rangle$ with another $\langle unaryOp \rangle$, and a $\langle scalar \rangle$ is replaced with a different $\langle scalar \rangle$. For example, this edit could replace \sin with \cos , $+$ with $*$, or the constant 2.0 with parameter p_0 .

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

Algorithm To favor exploration (i.e., considering a diverse set of BRDF variants) over exploitation (i.e., narrowly focusing on many different BRDF variants with subtle changes), we employ an *island model* genetic algorithm [Grosso 1985]. This involves partitioning the initial population into sub-populations that only rarely interact. The biological analogy is an archipelago of small islands with creatures that normally evolve in isolation. Occasionally, however, an organism on one island may *migrate* to another. Such a strategy encourages exploration of the space of BRDFs because each island is freed from the homogenizing effect of selection pressure within one population. We follow a simple, yet effective migration strategy [Andalon-Garcia and Chavoya-Pena 2012] between 4 islands: in every other 5th generation, the best variant from each island i migrates to island $((i + 1) \bmod 4)$.

At each iteration, a new population evolves from the current population. For each island’s sub-population of size N , we select $N/2$ BRDF variants based on their fitness to act as parents of the next generation of BRDF variants. Rather than selecting the $N/2$ best variants, we use tournament selection [Miller et al. 1995], i.e., the best variant from $k = 8$ randomly selected variants is retained. This further encourages exploration. These parent BRDF variants are randomly paired to create $N/2$ new offspring using a one-point crossover operator [Holland 1992]. That is, two edit sequences $\langle A_1, A_2 \rangle$ and $\langle B_1, B_2 \rangle$ produce offspring $\langle A_1, B_2 \rangle$ and $\langle B_1, A_2 \rangle$. The $N/2$ offspring and the $N/2$ parents then form the new generation of size N . To further increase diversity, one additional random mutation is applied to each variant.

Fitness The fitness of each variant is equal to the residual error after fitting its free parameters $(p_0 \dots p_{n-1})$ to a selected set of training materials. We enforce non-negativity of the free parameters by optimizing in the log domain, and fit each color channel separately. We estimate the BRDF parameters using the Nelder-Mead simplex search [1965] with 3 restarts guided by the E_2 BRDF error metric proposed by Löw et al. [2012]:

$$E_2(p, q) = \iiint (g_2(L, V; p) - g_2(L, V; q))^2 \sin \theta_o d\phi_o d\theta_o d\theta_i, \quad (3)$$

where $g_2(L, V; p) = \ln(1 + \cos \theta_i \rho(L, V; p))$. We sample V at half the resolution used by Löw et al. to reduce running times. In particular, we sample θ_i in 10 degree steps, and θ_o and ϕ_o in 2 degree steps. Löw et al.’s [2012] E_2 metric is originally intended for fitting BRDF models, and assumes that the BRDF model obeys reciprocity. However, the grammar listed in Figure 3 does not guarantee that BRDF variants are reciprocal. To encourage reciprocity, we augment the E_2 metric to include reciprocal sample directions, doubling the total number of samples. Lastly, unlike Löw et al., we sample the full hemisphere, including directions near grazing angles.

Input and Initialization Our algorithm takes as input:

1. the maximum number of free BRDF parameters n allowed in the BRDF variants;
2. an existing analytic BRDF model, represented as an AST, which serves as the starting point for the search, such as the

Phong BRDF model in the example shown in Figure 2;

3. an optional set of additional analytic BRDF models from which a codebook of sub-expressions is formed;
4. a training set of measured materials (e.g., a subset of the MERL-MIT database [Matusik et al. 2003]); and
5. the per-island population size N and the maximum number of iterations

The codebook of BRDF sub-expressions is formed by taking the union of all the subtrees within the AST representations of a set of BRDF expressions. For example, if the set of BRDF expressions was the singleton $\{\cos(R \cdot V)^n\}$, then the codebook would be the set of sub-expressions $\{\cos(R \cdot V)^n, \cos(R \cdot V), n, R \cdot V, R, V\}$. In addition, we require special handling of BRDF free parameters, which may vary in meaning between different BRDF models. For example, suppose that p_0 encodes the sharpness parameter in the Blinn-Phong model but p_0 encodes albedo in another BRDF model. Naïvely constructing the codebook from the subexpressions can result in an unintentional linking of BRDF parameters that encode different phenomena. We desire to combine subexpressions involving those parameters (i.e., exploring the search space by combining the best of both worlds) but the arbitrary labeling is a complication. Thus, if we are considering at most n free BRDF parameters and we encounter an expression p_i , we replace it with $p_0 \dots p_{n-1}$. For example, if an input BRDF expression contains p_1/π , we add $\{p_0/\pi, p_1/\pi, \dots, p_{n-1}/\pi\}$ to the codebook.

From the starting BRDF expression and codebook, we create an initial population by applying a random edit to the starting BRDF expression, and uniformly distribute these variants over the islands.

5 Experiments

We performed a series of experiments using the genBRDF framework described in the previous sections. The purpose of these experiments was three-fold. First, we wanted to identify analytic expressions not previously used for BRDF modeling that outperform the current state-of-the-art in terms of accuracy, compactness, or both.

Second, we wanted to study the feasibility of, and develop best practices for, applying a large-scale genetic search to the task of learning new analytic BRDF models. To this end, we observe that although our GP search can generate arbitrary expressions within the chosen grammar (Figure 3), BRDFs are often the product of several terms, each having a specific physical interpretation. For example, the *geometry* term in the Cook-Torrance BRDF model is meant to capture the effects of microfacets shadowing and masking one another [Torrance and Sparrow 1992]. Consequently, we hold fixed a different-sized portion of the starting BRDF expression in each experiment, thereby allowing the genetic search to focus on a narrower part of the problem. The goal was to identify the “sweet spot” that strikes the right balance between the size of the search space and the computational effort required to identify useful expressions.

Finally, our third experimental goal was to develop analysis techniques and visualization tools that help quickly identify useful BRDF expressions in the set of variants returned by our search.

We performed the following experiments:

FULL Experiment We allowed the GP search to manipulate all of the terms in the BRDF $f_r(L, V)$ and used the Cook-Torrance

BRDF model as the starting point:

$$f_r(L, V) = \frac{k_d}{\pi} + k_s \frac{F}{\pi} \frac{DG}{(N \cdot L)(N \cdot V)} \quad (4a)$$

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

$$G = \min \left\{ 1, \frac{2(N \cdot H)(N \cdot V)}{(V \cdot H)}, \frac{2(N \cdot H)(N \cdot L)}{(V \cdot H)} \right\} \quad (4c)$$

$$F = R_0 + (1 - R_0)(1 - (V \cdot H))^5, \quad (4d)$$

where δ is the angle between the half-direction (H) and the normal (N). The four free parameters (k_d , k_s , m , R_0) correspond to the amplitude of the diffuse lobe and specular lobe, the root mean square slope of the Beckmann microfacet distribution function [Beckmann and Spizzichino 1963], and the reflectance at normal incidence, respectively. This particular approximation of the Fresnel term is due to Schlick [1994].

Our codebook consisted of the union of all valid sub-expressions in the Ward [1992], Blinn-Phong [1977], and Cook-Torrance [1982] models. We allowed variants to incorporate up to four free parameters ($p_0 \dots p_3$).

In all of our experiments, BRDF variants produced during the search were forced to be non-negative by enclosing the generated code inside the max function: $\max(f_r, 0)$. We also return zero whenever L or V are found to lie below the horizon.

The inner loop of our search involves computing the best fitting values for the four BRDF parameters in each of the three color channels independently (Section 4). The final fitness value for each variant is the sum of these residual errors. In all of our experiments, we let the GP search run for 100 generations using 4 islands, each with a population of 1,024, yielding a total of 409,600 BRDF variants.

SPEC Experiment In this experiment we restrict the search to the specular term, $s_r(L, V)$, instead of attempting to estimate the entire BRDF:

$$f_r(L, V) = \frac{k_d}{\pi} + k_s s_r(L, V). \quad (5)$$

We use the corresponding specular term of the Cook-Torrance model as a starting point, and thus $s_r(L, V)$ subsumes the Fresnel (F), geometry (G), and microfacet distribution (D) terms. The codebook consisted of the union of all valid sub-expressions in the specular components from the Ward [1992], Blinn-Phong [1977], and Cook-Torrance [1982] models, and we allow for three free parameters ($p_0, \dots p_2$) in the GP search. Note that GP expressions generated during our search may not incorporate the parameters k_d and k_s because these appear in the fixed parts of the initial BRDF. However, k_d and k_s are included when computing the best fit (over a total of 5 parameters).

On the one hand, formulating the problem as estimating $s_r(L, V)$ embedded within a classical “diffuse plus specular” BRDF formula reduces the difficulty of finding useful expressions. On the other, it biases the search towards a particular structure and makes it more likely to overlook expressions that may produce superior fits [Oren and Nayar 1994; Kelemen et al. 2001]. These experiments were designed to analyze this important trade-off.

NDF Experiment As a final experiment, we further restricted the search to consider only the normal distribution and geometry terms of a microfacet BRDF [Walter et al. 2007], $n_r(L, V)$. The full BRDF expression is given as

$$f_r(L, V) = \frac{k_d}{\pi} + k_s \frac{n_r(L, V)F}{4(V \cdot N)(L \cdot N)}, \quad (6)$$




Figure 4: The set of 40 representative materials from the MERL-MIT BRDF database used for evaluation. Metals are shown on the left, dielectrics are shown on the right. Our training set (8 materials) are indicated with green borders.

where F is given in Equation (4d). We allow the term $n_r(L, V)$ estimated by our search to have two parameters and so, as with the SPEC experiment, the final BRDF has a total of five parameters.

For this experiment, our codebook was composed of the union of subexpressions from well-known NDF and geometry terms: the Beckmann [1963] distribution, Walter et al.’s [2007] GGX and Phong distribution, and Löw et al.’s [2012] normal distribution.

5.1 Training Data

To compute the fitness of each BRDF variant, we use 8 materials from the MERL-MIT database [Matusik et al. 2003] as our training set: *gold-metallic-paint2*, *specular-blue-phenolic*, *red-metallic-paint*, *green-plastic*, *pvc*, *tungsten-carbide*, *dark-red-paint*, and *blue-metallic-paint* (highlighted in Figure 4). These materials were chosen to achieve a variety of different dielectrics and metals.

5.2 Parallel Evaluation

In the course of each experiment, we evaluate the fitness of a very large number of BRDF expressions: 409,800 BRDFs \times 8 materials \times 3 color channels = 9,830,400 BRDF fits. An important benefit of the genBRDF framework is that it permits parallel execution. In our implementation, a single server runs the GP search, which generates and distributes BRDF variants to individual clients working in parallel, where they are compiled, their parameters are fit to the training data, and the fitness value is returned to the server. Specifically, we implemented the BRDF fitting procedure in CUDA and used a 24 node NVidia Tesla M2075 GPU cluster. A single experiment requires between 30 to 60 hours to complete, which would correspond to between 30 to 60 days on a single GPU.

5.3 Pareto Frontier

The genetic search produces a large number of analytic BRDF variants. Simply selecting the most fit variant runs the risks of overlooking perceptual artifacts that are not captured by our fitness measure. In addition to visual fidelity, other criteria such as expression length (defined as the string length of the BRDF expression) are important: we seek the most accurate BRDF variant that is no more complex than absolutely necessary.

To consider the incompatible dimensions of fitness and length simultaneously, we construct a Pareto frontier of BRDF variants. A




Figure 5: Pareto frontiers for the experiments described in Section 5.

Pareto frontier retains only non-dominated solutions, ruling out inferior BRDF expressions that are larger without being more accurate as well as BRDF expressions that are less accurate without being smaller. This step is critical, because it removes from consideration a significant number of dominated, inferior BRDF variants. The use of Pareto frontiers is a standard practice for enumerating the optimal set of choices among multiple dimensions when those dimensions cannot be directly compared [Sitthi-Amorn et al. 2011].

Figure 5 shows the Pareto frontiers of our three experiments, each of which contains between 30 and 45 BRDF variants. An ideal BRDF expression would minimize both fitness and length (lower left corner). Thus, an axis-hugging curve that approaches the origin is preferable. The points on a Pareto frontier are formally incomparable and must thus be further inspected or filtered manually. Any knee in a Pareto frontier, such as at point “E” in Figure 5, represents a natural choice for further inspection.

5.4 Image-Driven BRDF Evaluation

Despite the significant reduction in problem size by considering only the Pareto frontier, it is still a difficult and time-consuming task to analyze 30 to 45 machine-generated analytic BRDF expressions by hand. Simply selecting the variant with the best fitness value is inadequate, as this risks overlooking perceptual artifacts that are not captured by our fitness measure and missing more compact expressions that provide high fidelity fits to the training set.

Inspired by prior work [Ngan et al. 2005], we guide the process of analyzing the BRDF variants on the Pareto frontier by comparing images of spheres rendered under natural lighting. Specifically, we use the *grace cathedral* light probe¹ and compute both the L_2 error on absolute radiance values along with the Structural Similarity Image Metric (SSIM) [Wang et al. 2004] on tone-mapped pixel values — we found that a simple gamma 2.2 mapping works well in practice.

Specifically, we fit each BRDF variant on the Pareto frontier to a representative selection of 40 materials from the MERL-MIT database (Figure 4). This set includes an even number of metals and dielectrics to provide a balanced evaluation. (Please refer to the supplemental materials for an exact listing.) We compare these fits to the best reported fits for what we consider to be the state-of-the-art in analytic isotropic BRDF models: Cook-Torrance [1982], Löw et al.’s [2012] smooth surface (or Rayleigh-Rice) model and Löw et al.’s microfacet (MF) model, and the model proposed by

Bagher et al. [2012]. We use the parameter fits reported by Ngan et al. [2005] for the Cook-Torrance BRDF, and those provided by the authors of the Löw et al. and Bagher et al. models, respectively. Additionally, we also fit the Cook-Torrance BRDF ourselves using the E_2 metric [Löw et al. 2012] and the Nelder-Mead simplex algorithm with 100 restarts.

5.5 Reciprocity

Although we encourage reciprocity in the way our fitness measure is constructed (Section 4) this is not a hard constraint. Instead, we achieve reciprocity by explicitly culling variants that exceed a numerically computed reciprocity threshold. Specifically, we compute:

$$\sqrt{\frac{1}{\#T} \sum_{p \in T} \int \int (f_r(L, V; p) - f_r(V, L; p))^2 dLdV}, \quad (7)$$

where T is the set of best-fit BRDF parameters for our training set. We evaluate Equation 7 using Monte Carlo integration with 10,000 uniformly generated (L, V) pairs. We conclude that a BRDF is most likely reciprocal if this measure is below 10^{-5} to allow for floating-point rounding errors. We only consider BRDF variants that pass this test when constructing the Pareto frontier.

We also investigated modifying the BRDF grammar in order to enforce reciprocity during the genetic search. In particular, we performed the *FULL Experiment* described above, but returned the following expression: $0.5(f_r(L, V) + f_r(V, L))$, which is always reciprocal. However, we found that this approach restricts the search space too much and we were not able to produce useful BRDF expressions. Future work could investigate alternative grammars that are more flexible, but still designed to guarantee reciprocity in the generated BRDF models.

6 Results

The three experiments described in the previous section produced a total of ~ 120 BRDF variants across their respective Pareto frontiers. These BRDF variants reveal a variety of different analytic expressions. They also represent a smooth spectrum of options that trade accuracy for length, as can be seen by their corresponding Pareto frontiers (Figure 5). Table 1 lists a selection of 5 BRDF variants from this total. The corresponding location on the Pareto frontiers for each of these BRDF variants is marked in Figure 5. A visual comparison to ground truth, Cook-Torrance [2005], and Löw et al.’s [2012] smooth surface model is shown in Figure 10. As can be seen, all models produce slightly different, but accurate matches.

Each entry in Table 1 summarizes a single BRDF variant. BRDF models A and D correspond to variants 015-000599 and 036-003067 in the supplementary document “NDF Experiment Result” respectively; BRDF models B and E correspond to variants 025-000566 and 026-000754 in the supplementary document “SPEC Experiment Result”; and BRDF model C corresponds to the variant 025-001330 in the supplementary document “FULL Experiment Result.” For each entry we list the fitness ($\times 10^{-3}$) and expression length of the variant (as returned by genBRDF, without algebraic simplifications). The expressions shown in this table have been algebraically simplified for the sake of clarity (e.g., $1.0 \cdot X$ was simplified to X , etc.). The last column list how many times the BRDF variant outperformed *all* the competing models according to the L_2 and SSIM metric on the test set of 40 materials (Section 5.4).

At a high level, BRDF models A and E have similar expressions after taking into account the fixed components in the NDF and SPEC experiments (i.e., $NDF = SPEC \cdot F / (4(L \cdot N)(V \cdot N))$),

¹<http://www.pauldebevec.com/Probes/>

| BRDF | Exp. | ID | Fit (10^{-3}) | Len. | (Simplified) Expression | L_2 /SSIM |
|------|------|------------|-------------------|------|--|-------------|
| A | NDF | 015-000599 | 0.10661 | 120 | $n_r = \frac{(V \cdot N)(L \cdot N)e^{-(p_0 \tan \delta)^{p_1}}}{(L \cdot H)}$ | 20/20 |
| B | SPEC | 025-000566 | 0.10112 | 308 | $s_r = e^{-\left(\frac{\tan \delta}{p_1}\right)^{p_2}} \min \left\{ 1, \frac{p_1(p_0+2)}{\tan \delta} \right\} \frac{F(p_0, (L \cdot H))}{(p_0) \cos \delta}$ | 28/24 |
| C | FULL | 025-001330 | 0.10922 | 314 | $f_r = \frac{p_0}{\pi} + \frac{p_1}{\pi} e^{-\frac{p_3}{p_3^2(L \cdot H)}} \min \left\{ \frac{p_0 p_3 (p_1+2)}{\tan \delta}, (L \cdot H), 1 \right\} F(p_2, (L \cdot H))$ | 23/22 |
| D | NDF | 036-003067 | 0.09941 | 272 | $n_r = \min \left\{ 1 + p_0, \frac{2(V \cdot N)(L \cdot N)}{(L \cdot H)} \right\} \min \left\{ (L \cdot H), \frac{p_0}{p_1 \tan \delta}, \frac{2(H \cdot N)}{p_0} \right\} e^{-\frac{\tan \delta}{p_0}}$ | 29/26 |
| E | SPEC | 026-000754 | 0.10673 | 170 | $s_r = e^{-\left(\frac{\tan \delta}{p_1}\right)^{p_2}} \frac{p_0 + (1 - (L \cdot H))^5}{(L \cdot H)}$ | 19/21 |

Table 1: A representative selection of BRDF models discovered in the course of our experiments with genBRDF.

and a reordering of free parameters ($A \rightarrow E : \{p_0, p_1, R_0\} \rightarrow \{1/p_1, p_2, p_0\}$). Algebraically, the only difference is the factor $(1 - p_0)$ that appears in the Fresnel component. Bear in mind that these variants were independently discovered in separate experiments with different codebooks. BRDF model A appeared in generation 15 of the more constrained NDF experiment, whereas BRDF model E appeared 11 generations later (26) in the less constrained SPEC experiment.

BRDF models C and D illustrate how the genetic search is capable of producing much longer expressions. Originally of similar length as BRDF models C and D, the expression of BRDF model B after algebraic simplification is more compact. Genetic programming is essentially a “blind” search through the space of analytic BRDF expressions, and it does not take in account the algebraic or physical implications of mutations on the expressions. Hence, manual simplification and interpretation is still necessary after a good variant has been found.

Although BRDF model C is slightly less accurate than the other variants, it is worth noting that it has only 4 parameters, whereas the others have 5. Furthermore, note that it outperformed the existing 5-parameter models from prior work in 23 out of the 40 test materials.

We consistently observe a new two-parameter microfacet normal distribution function (NDF) within these expressions. This NDF can express a wider range of specular lobe shapes and achieves more accurate fits to measured data than current models. To gain further insight into this new expression, we select BRDF model A for a more in-depth analysis.

6.1 BRDF Model A²

The following BRDF expression was produced by the *NDF Experiment*:

$$n_R(L, V) = \frac{(V \cdot N)(L \cdot N)}{(L \cdot H)} D' \quad (8a)$$

$$D' = e^{-\left(\frac{\tan \delta}{\beta^2}\right)^\alpha} \quad (8b)$$

²This BRDF corresponds to variant 015-000599 in the supplemental document “NDF Experiment Results”. Note that for the sake of clarity, we performed several changes of variables and algebraic simplifications. In particular, we substituted the parameter named p_0 in the supplemental document with $1/(\beta^2)$. This was done to give the final BRDF parameters more intuitive ranges and to make them similar to their closest counterparts in the existing BRDF literature.




Figure 6: Example of a false color plot of the relative difference in L_2 error on the forty validation materials (vertical axis) of Löw et al.’s smooth surface model and the variants on the Pareto frontier from the NDF experiment (horizontal axis). Blue indicates the variant has lower error, while red indicates the opposite. Highlighted is BRDF variant A. Refer to the supplemental document “NDF Experiment Results” for comparisons to other BRDF models.

After a few substitutions and algebraic simplifications, this corresponds to the BRDF

$$f_r(L, V) = \frac{k_d}{\pi} + k_s \frac{D' F}{4(L \cdot H)} \quad (9)$$

The parameters α and β govern the shape of the specular lobe and are discussed below. Figures 1 and 10 compare this BRDF model to what we consider to be the state-of-the-art in analytic isotropic BRDF models (Section 5). The latter also shows the best fits for all of these BRDF variants in Table 1 for two materials in our training set (*gold-metallic-paint2* and *red-metallic-paint*) plus four materials that were not part of the training set. (Please see the supplemental document *NDF Experiment Results* for the complete comparisons to all forty materials.)

Of the selected forty materials in the MERL-MIT database that we tested, and under the testing conditions (e.g. light probe, geometry, etc.), BRDF model A produced the most accurate fits 20 times when the L_2 distance metric was used to judge image fidelity and 20 times when the more perceptually meaningful SSIM measure was used. The second-best performing BRDF model was Löw et al.’s smooth surface model, which produced the most accurate fits only 10 times and 5 times for the L_2 and SSIM quality measures, respectively. The more complex model of Bagher et al. [2012] has 7 free parameters, requiring a complex fitting procedure, but did not perform as well. Furthermore, in the cases where BRDF model A did not produce the best fits, it was very close. This can be seen in the false color map shown in Figure 6, which visualizes, for all

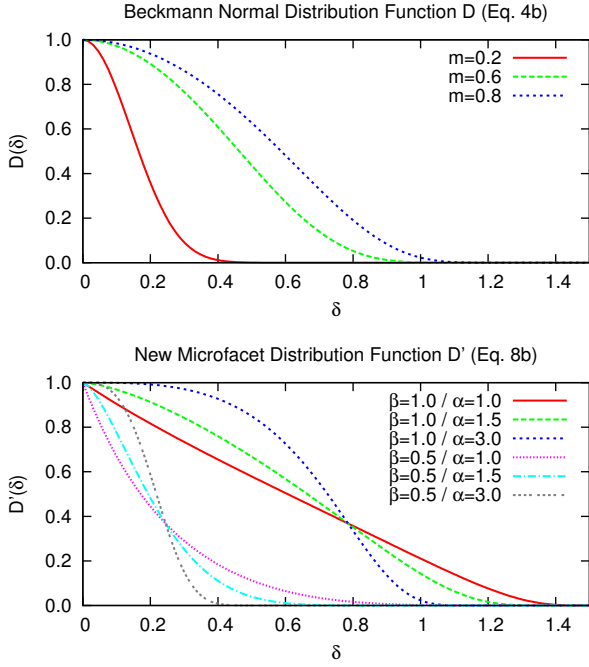


Figure 7: Comparison of (top) the Beckmann microfacet normal distribution function [Beckmann and Spizzichino 1963] with (bottom) a new two-parameter NDF found by our genetic search.

variants on the Pareto frontier for the NDF experiment, the relative improvement in L_2 error over Löw et al.’s smooth surface model when BRDF model A gives the lowest error (blue) or the relative difference when BRDF model A does not yield the lowest error (red). We relied heavily on these types of visualizations to help identify interesting parts of the result set.

Discussion We chose this particular BRDF variant because it contains understandable pieces, captures some interesting trends that we observed across the entire result set, and occurs near a clear “knee” in the Pareto frontier (Figure 5). Moreover, the accuracy of such a compact analytic BRDF expression is noteworthy.

We believe that the two-parameter normal distribution function D' (Equation 8b) has not been previously studied in the BRDF literature. This expression appears repeatedly in our search results, leading us to conclude that it offers a stable and reliable way of reproducing the shapes of many common specular lobes. Figure 7 illustrates the effect of these two parameters (α and β) on the shape of the specular lobe and offers a comparison to the often used Beckmann distribution [Beckmann and Spizzichino 1963]. Note that D' is able to create spikier highlights, which are consistent with the shape of many of the lobes in the BRDF database, and helps explain its superior performance. Note that the smooth surface model proposed by Löw et al. [2012] also includes a two-parameter NDF that can produce similar (but slightly different) shapes.

Figure 8 shows the effect of adjusting either α or β while holding the other fixed. As our model is not normalized, we adjust k_s accordingly to provide equal overall brightness. k_d was kept to zero. We conclude that the parameters α and β determine the *strength* and *clarity* of the specular highlight, respectively. In other words, larger values of β^2 produce highlights that are softer and appear to glow near the edges of strong light sources, whereas smaller values produce sharper and more clearly defined highlights.




Figure 8: A visualization of the effect that α and β have in BRDF model A on the material appearance. Each row of spheres shows the effect of manipulating a single parameter. k_s was scaled appropriately to provide equal overall brightness.

BRDF model A can also be regarded as having a significantly simpler geometry term, $(1/(4(L \cdot H)))$, than current models. Compared to the Cook-Torrance BRDF model, for example, BRDF model A has dropped the $(L \cdot N)$, $(V \cdot N)$ and $(H \cdot N)$ terms that appear in the numerator and denominator of the geometry term, respectively, depending on the value of the ternary *min* function. While we cannot offer an explanation of this simpler term based on first principles, it does achieve the qualitative goal of increasing the magnitude of the BRDF near “grazing” configurations ($(L \cdot H) \rightarrow 0$) and it does so with a far simpler expression. Finally, note that BRDF model A is reciprocal by inspection since $(L \cdot H) = (V \cdot H)$.

6.2 METAL Experiment

Although the selected models from the three previous experiments produce good fits to our test set of forty materials, we did notice that they tend to fit dielectrics better than metals. We used genBRDF to perform two more targeted experiments to produce better models for metals. Specifically, we repeated the NDF experiment with a different training set of 8 metals: *gold-metallic-paint*, *steel*, *aluminium*, *nickel*, *blue-metallic-paint*, *blue-metallic-paint2*, *red-metallic-paint*, and *brass*. From this experiment we briefly highlight one variant 043-001155; please refer to “METAL Experiment Result” in the supplementary documents for a detailed report.

This variant has a fitness of 0.0963×10^{-3} and an expression length of 353. The full BRDF after simplification and parameter substitution is:

$$n_R(L, V) = \frac{D'G'}{\alpha^{-2}\beta^{-2}\tan^2\delta + c}$$

$$G' = \min \left\{ 1 + \beta^2(1 + \cos\delta), 2\frac{(L \cdot N)(V \cdot N)}{(L \cdot H)}, \frac{1}{\beta^2 \tan\delta} \right\}$$

where $c = \cos^4(1) = 0.9994$ is an example of a vestigial remnant that can most likely be replaced by the constant “1”. This model again exhibits the same two-parameter NDF D' as seen before, in addition to a somewhat more complex geometry term. It performs best in terms of the L_2 (resp. SSIM) error for 15 (reps. 13) times over the 20 validation metal materials (left side in Figure 4). For comparison, Löw et al.’s smooth surface model was the second best (2/5). The selected BRDF variants in Table 1 performed better than the alternative BRDF models from prior work 3/7, 10/9, 6/6, 9/10, and 3/7 times respectively. Figure 9 shows a visual comparison for two selected materials.

Lastly, we observe that, even after learning new metal BRDF models, for some materials only few of the found expressions excel.




Figure 9: A selection of fitted materials using a selected *genBRDF* from the METAL experiment compared to Löw et al.’s smooth surface model.

An example of such a material is *chrome*. We therefore repeated the SPEC experiment using only the *chrome* material sample in the MERL-MIT database as the training set. As expected, this produced several superior analytic expressions for this specific material when compared to any of the other BRDF variants discovered by our framework in the earlier experiments. For example, the L_2 (SSIM) error on BRDF model D is 0.021 (0.969) compared to 0.012 (0.976) for variant 049-00770. Please refer to the supplemental document titled “CHROME Experiment Results” for a complete analysis.

7 Discussion

Our primary contribution, the proposed genetic search framework, produces BRDF expressions that are only strictly valid for the materials in the training set. In other words, they are not guaranteed to produce valid (or even plausible) BRDFs for any values of their parameters — a trait shared by other *data-driven* BRDF models [Matusik et al. 2003]. In fact, one could interpret the models found by our search as an analytical encoding/compression of the training data. However, our analysis shows that when the resulting models are fit to other measured materials they still produce good results.

Note that unless they are explicitly enforced by either the grammar or fitness function, common properties of BRDFs are not guaranteed to be respected by the computed models. Besides reciprocity, energy conservation is another example of such property. Note that energy conservation issues are largely mitigated by fitting these models to measured materials that presumably conserve energy themselves. We validated energy conservation numerically for the BRDF variants in Table 1 and found that they all conserve energy for the fitted parameters. Also, note that our models are not normalized in the sense that the so-called diffuse and specular “albedos” can exceed 1. The effects of other parameters are typically non-intuitive, and in general, care should be taken when forming physical interpretations of the terms and parameters in models found by *genBRDF*.

More generally, there are a number of factors that impact the results of our search:

- **Training Materials** As noted above, the generated expressions and corresponding parameters are only strictly valid for the training material set. As such the choice of this training set is critical in determining how the computed models will extrapolate to other materials. For example, if the BRDF vari-

ants were trained on only plastics, then it is unlikely that the model will retain similar accuracy for metals

- **Initial BRDF and Codebook** As with many optimization algorithms, the starting point plays a major role in which part of the search space is explored. Similarly, the codebook (i.e., the “primordial soup” of possible components) can be seen as operators that allow us to “jump” over ridges and disconnected areas in the fitness landscape, enabling exploration of a much larger portion of the search space. In other words, the the more varied the codebook, the more “intelligent” the genetic search can be.
- **Fitness Metric** The fitness metric plays a crucial role in the search process. Care has to be taken when applying existing BRDF fitting functions or BRDF similarity metrics. Typically, these metrics are either too computationally expensive, or they rely on specific properties of BRDF models and physical materials (e.g., reciprocity, continuity, symmetries, etc.). However, the outcome of these metrics might not be meaningful if the BRDF variant does not exhibit one of these properties.
- **Non-linear Optimization Algorithm** The optimization algorithm used for fitting the BRDF parameters also implicitly restricts the search space to functions that can be reliably optimized. A positive consequence of this is that the found BRDF expressions are likely well-behaved when using the Nelder-Mead simplex algorithm to fit their parameters.

Our analysis relies on the L_2 and SSIM metrics on a specific scene and lighting to find good expressions. Different metrics such as the one proposed by Pereira et al. [2012] or CIEDE2000 [Sharma et al. 2005] could potentially produce different results. However, unlike the fitness metric this choice does not impact the genetic search, only the analysis.

Finally, note that the variants produced by our search cannot immediately be used for importance sampling without further analysis. To generate the images in this paper, we followed the common convention of using a proxy-BRDF to guide importance sampling. We implemented both a GPU-based BRDF renderer (used for generating the overview reports) as well as a high-quality Mitsuba [Jakob 2010] plug-in. For the former, we use a Blinn-Phong proxy to generate BRDF samples, while for the latter we use Löw et al.’s smooth surface model as a proxy. As a side note, we found that of the alternative BRDF models we considered, Löw et al.’s smooth surface model outperformed the competition in our experiments.

8 Conclusion

We introduced *genBRDF*; a framework for discovering new analytic BRDFs. We adapted genetic programming for this purpose and introduced a suitable grammar and fitness function. Our technique is not meant to replace, but rather augment, the human effort involved in deriving expressive analytic models. In addition, we described a strategy and set of tools for rapidly analyzing the large result sets returned by our genetic search. We demonstrated the effectiveness of our method by describing four experiments that demonstrate new BRDF expressions that in many cases exceed the best currently available alternatives.

Our search strategy intentionally favors exploration over exploitation. For future research we would like to consider alternative strategies that perhaps favor exploration in the beginning, but gradually place a heavier emphasis on exploitation as the search progresses. Other avenues for research include: learning specialized BRDF expressions of each material in the MERL-MIT database independently in order to establish a baseline model for future BRDF development; and finding a fitness metric that better pre-

dicts perceptual quality of the BRDF variants. Finally, we believe the genBRDF framework could be adapted to many other modeling problems in computer graphics that would benefit from having compact analytic models of complex high-dimensional functions.

Acknowledgements

We wish to thank the reviewers for their helpful and constructive feedback. We also acknowledge valuable discussions with David Luebke and Tomas Akenine-Möller. This work was partially funded by NSF grants IIS-1011444, IIS-1217765, IIS-1350323, CCF-0747220, CCF-0954024, and CCF-0905373, DARPA grant P-1070-113237 and gifts from NVIDIA and Google. This work was performed in part using the computational facilities at the College of William & Mary which were provided with the assistance of the National Science Foundation, the Virginia Port Authority, and Virginia's Commonwealth Technology Research Fund.

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., PREMOZE, S., AND SHIRLEY, P. 2000. A microfacet-based BRDF generator. In *Siggraph 2000, Computer Graphics Proceedings*, 65–74.
- BAGHER, MAHDI, M., SOLER, C., AND HOLZSCHUCH, N. 2012. Accurate fitting of measured reflectances using a Shifted Gamma micro-facet distribution. *Computer Graphics Forum* 31, 4 (June), 1509–1518.
- BECKMANN, P., AND SPIZZICHINO, A. 1963. *The scattering of electromagnetic waves from rough surfaces*. International series of monographs on electromagnetic waves. Pergamon Press.
- 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, 7–24.
- DORSEY, J., RUSHMEIER, H., AND SILLION, F. 2008. *Digital Modeling of Material Appearance*. Morgan Kaufmann Publishers Inc.
- GROSSO, P. B. 1985. *Computer simulations of genetic adaptation: parallel subcomponent interaction in a multilocus model*. PhD thesis, Ann Arbor, MI, USA.
- HE, X. D., TORRANCE, K. E., SILLION, F., AND GREENBERG, D. P. 1991. A comprehensive physical model for light reflection. In *SIGGRAPH'91 conference proceedings*.
- HOLLAND, J. H. 1992. *Adaptation in natural and artificial systems*. MIT Press, Cambridge, MA, USA.
- JAKOB, W., 2010. Mitsuba renderer. <http://www.mitsuba-renderer.org>.
- KELEMEN, C., SZIRMAY-KALOS, L., AND SZIRMAY-KALOS, L. 2001. A microfacet based coupled specular-matte brdf model with importance sampling. In *Eurographics (short presentations)*.
- KOZA, J. R. 1992. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. MIT Press.
- LÖW, J., KRONANDER, J., YNNERMAN, A., AND UNGER, J. 2012. Brdf models for accurate and efficient rendering of glossy surfaces. *ACM Transactions on Graphics (TOG)* 31, 1 (January).
- MATUSIK, W., PFISTER, H., BRAND, M., AND MCMILLAN, L. 2003. A data-driven reflectance model. *ACM Trans. Graph.* 22, 3, 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. *Eurographics Symposium on Rendering 2005*, 117–226.
- NICODEMUS, F. E., RICHMOND, J. C., HSIA, J. J., GINSBERG, I. W., AND LIMPERIS, T. 1977. Geometric considerations and nomenclature for reflectance. *Monograph 161, National Bureau of Standards (US)*.
- 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*, 239–246.
- PEREIRA, T., AND RUSINKIEWICZ, S. 2012. Gamut mapping spatially varying reflectance with an improved brdf similarity metric. *Comput. Graph. Forum* 31, 4, 1557–1566.
- PHONG, B. T. 1975. Illumination for computer generated pictures. *Commun. ACM* 18, 6, 311–317.
- SCHLICK, C. 1994. An inexpensive brdf model for physically-based rendering. *Computer Graphics Forum* 13, 3, 233–246.
- SHARMA, G., WU, W., AND DALAL, E. N. 2005. The ciede2000 color-difference formula: Implementation notes, supplementary test data, and mathematical observations. *Color Research & Application* 30, 1, 21–30.
- SITTHI-AMORN, P., MODLY, N., WEIMER, W., AND LAWRENCE, J. 2011. Genetic programming for shader simplification. *ACM Trans. Graph.* 30, 6 (Dec.).
- TORRANCE, K. E., AND SPARROW, E. M. 1992. Radiometry. ch. Theory for off-specular reflection from roughened surfaces, 32–41.
- WALTER, B., MARSCHNER, S. R., LI, H., AND TORRANCE, K. E. 2007. Microfacet models for refraction through rough surfaces. In *Rendering Techniques*, 195–206.
- WANG, Z., BOVIK, A. C., SHEIKH, H. R., AND SIMONCELLI, E. P. 2004. Image quality assessment: From error visibility to structural similarity. *IEEE Trans. Image Proc.* 13, 4, 600–612.
- WARD, G. 1992. Measuring and modeling anisotropic reflection. In *ACM SIGGRAPH Computer Graphics*, vol. 26, 265–272.




Figure 10: A selection of five materials discovered using our genBRDF framework compared to Cook-Torrance and Löw et al.'s smooth surface model.