

Image-driven Navigation of Analytical BRDF Models

Addy Ngan¹

Frédo Durand¹

Wojciech Matusik²

¹Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology

²Mitsubishi Electric Research Laboratories (MERL)

Abstract

Specifying parameters of analytic BRDF models is a difficult task as these parameters are often not intuitive for artists and their effect on appearance can be non-uniform. Ideally, a given step in the parameter space should produce a predictable and perceptually-uniform change in the rendered image. Systems that employ psychophysics have produced important advances in this direction; however, the requirement of user studies limits scalability of these approaches. In this work, we propose a new and intuitive method for designing material appearance. First, we define a computational metric between BRDFs that is based on rendered images of a scene under natural illumination. We show that our metric produces results that agree with previous perceptual studies. Next, we propose a user interface that allows for navigation in the remapped parameter space of a given BRDF model. For the current settings of the BRDF parameters, we display a choice of variations corresponding to uniform steps according to our metric, in the various parameter directions. In addition to the parametric navigation for a single model, we also support neighborhood navigation in the space of all models. By clustering a large number of neighbors and removing neighbors that are close to the current model, the user can easily visualize the alternate effects that can only be expressed with other models. We show that our interface is simple and intuitive. Furthermore, visual navigation in the BRDF space both in the local model and the union space is an effective way for reflectance design.

Categories and Subject Descriptors (according to ACM CCS): I.3.6 [Computer Graphics]: Interaction techniques, I.3.7 [Computer Graphics]: Color, shading, shadowing, and texture

1. Introduction

Appearance modeling is a crucial aspect of image synthesis. One common way to represent appearance in computer graphics is to use the Bidirectional Reflectance Distribution Function (BRDF), which captures the interaction of light and matter. BRDFs can be represented using different analytic models. However, the specification of the model parameters is often difficult because these parameters can be non-intuitive for artists and can have non-uniform effects on the rendered image. For example, it is well known that changing the Phong exponent from 5 to 10 has a much more noticeable effect than changing it from 95 to 100. This non-uniformity makes the design of desired BRDFs difficult for users. Ideally, a given step in the parameter space should produce a predictable and perceptually-uniform change of the rendered image. To alleviate this issue, software developers often include ad-hoc remapping of the parameters, for example displaying the log of the Phong exponent. However, a more systematic solution is desirable in order to make

material modeling more intuitive. Systems that employ psychophysics have produced important advances in this direction [PFG00]. The main drawback of these systems is required extensive user studies, which limits their scalability. In this work, we take an approach that is not strictly perceptually-based but that facilitates user navigation in the space of analytical BRDFs.

Another difficulty often faced by users is the choice of analytic reflectance model. BRDF models differ in the class of materials they are capable of representing and it takes considerable expertise to pick the right model for a desired material. Various models differ visually in subtle ways that can be important to the user but are hard to guess. In our work, we seek to blur the boundary between models by allowing users to navigate freely across a variety of models and access the nuances afforded by different analytical BRDFs.

Our visual navigation interface presents the user with a set of images sampling the reflectance neighborhood of the current material selection. The user simply clicks on the image

he likes best, and the reflectance corresponding to this image becomes the new selection. This way of navigation simplifies the user's task by removing the need for a mental model of the parameter mapping. Indeed, the user directly *sees* the potential effect of a parameter change before effecting it. This interface requires that the spacing of the BRDF variations shown to the user be as perceptually uniform as possible. In this work, we propose a metric for BRDFs based on rendered images. We show that this metric is uniform with respect to perceptual measures. Using this metric we can find a set of equidistant neighbors in the current model by independently changing each of the parameters. In addition, our metric enables us to define the notion of a neighborhood in the space of all models. In this context, our interface permits easy conversion between BRDF models and it reveals local differences between different models. In particular, it emphasizes neighbors that do not have a similar counterpart in the current model.

1.1. Related Work

In order to leverage the industry's expertise in material specification, Westlund and Meyer [WM01] applied appearance standards to establish correspondences between measurement scales and parameters of analytic BRDF models. They *measure* instances of the analytical models by rendering BRDFs at several angles specified by the industry appearance standards. This produces a one-to-one mapping between the industry gloss values and the gloss parameters of a few analytical models.

Pellacini et al. [PFG00], inspired by the work on perceptually uniform color spaces (e.g., CIELAB), proposed a perceptual adaptation of Ward's BRDF model [War92]. They reparameterize the model based on psychophysical experiments, in which subjects are asked to assign numbers to describe the apparent differences between rendered images of different BRDFs. Next, they apply multidimensional scaling (MDS) techniques to recover the perceptual axes and scaling of the gloss space according to the reported differences. They demonstrate that the reparameterized model is easier to use because it is perceptually uniform. However, it is difficult to extend the same method to more general BRDF models since the higher dimensionality and increase in the number of samples quickly render human experiments impractical.

The work of Pellacini et al. is the main inspiration behind our approach. We replace their psychophysical experiments with a computational surrogate that is not strictly perceptual but provides better scalability. In addition, we present a new visual interface for BRDF parameter specification.

A number of computational metrics have been proposed to capture human visual performances, e.g., [Dal93, RWP*95]. These metrics have had successful applications in realistic image synthesis, e.g., [BM98, Mys98, RPG99].

Our work deals with a stimulus space that is simpler than most related work because our images are perfectly aligned and the parameter space of analytical BRDF models has a smooth effect on the rendered images.

We also build on recent results from Dror et al. [DAW01] that indicate that a single image of a material under unknown illumination is often enough for reflectance recognition by a human. Furthermore, Fleming et al. [FDA01] show that humans can recognize surface reflectance under natural illumination with high accuracy from a single image. Leveraging these results, we assume that reflectance can be compared meaningfully using a single rendered image for each BRDF.

The most popular line of research for intuitive specification in realistic rendering is inverse or goal-based rendering, e.g., [SDS*93, KPC93]. This approach simplifies the task of the user and removes the need for mentally intricate reverse-engineering [Dur02]. However, it requires that users be able to assign absolute goals and have a precise idea of what they desire to achieve.

The recently-published BRDFShop system [CPK06] also seeks to facilitate the specification of material appearance. It builds on a painting interface and the development of an extended Ward model while we focus on the specification of parameters for existing BRDF models.

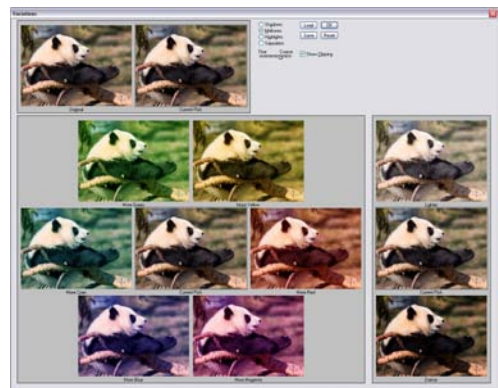


Figure 1: Photoshop's Variations interface.

Our work also builds upon the research in user interface design. Approaches such as design galleries [MAB*97] and the Adobe Photoshop *Variations* interface [Bak01] present the user with various visual options. Design galleries deal with complex multidimensional and discontinuous parameter spaces and use a computational metric to maximize the dispersion of the output images. The Adobe Photoshop "Variations" interface, currently under the image/adjustments menu, facilitates the alteration of chromaticity, brightness and saturation of an image (Fig. 1). The interface displays the current modified image as well as several variations along the various chromaticity, saturation, or brightness axes. It makes color correction intuitive because

the user does not need to form a mental model of the effect of the various parameters. The options are presented in a visual form that can be directly judged by the user based on relevant subjective criteria [Bak01]. The success of this interface has strongly influenced our work on BRDF specification.

1.2. Overview

The technical contributions that make our interface possible are as follows. First, we define a computational metric based on rendered images to measure distance between BRDFs. We show that our metric produces similar results to previous perceptual studies (Section 2). Next, we create an interface where the user navigates within a BRDF model with parameters remapped using our metric. For the current parameter settings, we display a choice of *variations* corresponding to uniform steps in the various parameter directions (Section 3.3). The purely-numerical control for the user is reduced to one slider that sets the desired distance from the current BRDF to the neighbors. In addition to local navigation within a model, we also allow the user to visualize the neighbors in other models, with an emphasis on maximizing the palette of possible appearances. In order to maximize the diversity of materials offered to the user, we discard neighbors that can be well-approximated by the current model, and we choose the ones that are most different from each other (Section 4). In Section 3.2 we discuss the embedding of the BRDF models into a unified low-dimensional space, which is essential for constructing the neighborhoods at interactive rates.

2. BRDF Metric

Comparison of BRDFs has typically been used for fitting measured data to an analytic model. One popular choice is the squared difference of the two BRDFs multiplied by the cosine of one or both incident and outgoing angles, integrated over the hemisphere [LFTG97]. However, the criteria of a good fitting (near-threshold) metric do not necessarily make it appropriate for our task, which focuses on suprathreshold perceptual uniformity. In particular, BRDF-space metrics like the one mentioned tend to overfit near the mirror direction, as the values are often orders of magnitude higher than the average of the whole BRDF. In the first row of Figure 8, we show the Ward model varying along the roughness axis (α) uniformly spaced according to the BRDF-space L^2 metric described above. We observe that the distance is much larger when the BRDF is sharp (mirror-like), and as a result the samples are concentrated near the sharper range. This motivates the development of a metric that better captures the visual effect of BRDFs.

2.1. Image-driven Metric

Our metric follows from two decisions. The most important decision is to define the distance between two BRDFs as

the difference between the rendered images with the given BRDFs under a natural environment map. The second decision concerns the precise image difference, and our choice is a compromise between perceptual motivation, simplicity, and computational efficiency.

As BRDFs are used to define object appearance in rendered images in most applications, we choose to define our metric in the image domain. While it is clear that a single image does not have enough information to uniquely define a BRDF, we leverage results from Dror et al. [DAW01] and Fleming et al. [FDA01] that show that a single image can capture a large part of the material characteristics if the illumination is natural. In this work, we choose to use the image of a sphere rendered with a given natural environment map, in practice the Grace Cathedral, courtesy of Paul Debevec. For the rest of the discussion, we will use the term *BRDF image* to represent this particular scene rendered with the corresponding BRDF. Note that, ideally, the environment map of the intended final rendering could be used at the cost of increased computation.

The image difference that we choose is the L^2 difference between the cubic roots of the RGB channels of two BRDF images. The images are represented in floats with high-dynamic range, without any tone mapping. The choice of the cubic root is inspired by the luminance mapping of the perceptually uniform CIELAB color space [Fai98] and related to early tone mapping operators [TR93]. In our application, we do not use more perceptually-uniform CIELAB for computational reasons: the cubic root of the RGB channels is a slightly simpler formula that allows us to precompute BRDF images and enable approximation based on principal components analysis (Section 3.2). However, we advocate the use of CIELAB if precomputation is not an issue. More comprehensive visual difference predictors could also be used, but would be even more computationally expensive. In practice, we have found that our simple metric yields uniform spacing.

2.2. Metric Evaluation

To validate our metric with perceptual measures, we compare the reported distance of our metric to the psychophysical experiments by Pellacini et al. [PFG00]. In their work, 11 renderings of the Ward model with different roughness are shown to human subjects, and the subjective *gloss ratings* are reported. The reported values together with the linear fit proposed are shown in Figure 2. In order to provide a comparison with their results, we sample the same parameter range in the Ward model, and the distance of each sample to its next neighbor is computed according to our metric. The cumulative distance starting from the first BRDF ($d = 0.8$) is plotted in blue in the figure. As there is an unknown calibration scale between our metric and the reported ratings, we choose the scale to give a best fit to the data (which explains why the curves do not meet at $d = 0.8$.) The same cal-




Figure 2: Comparing the gloss ratings reported by human subjects and the 10-steps cumulative distances reported by our image-driven metric (Blue), and the direct L^2 metric with LAB (Green) and RGB (Red). The linear fit proposed by Pellacini et al. is also shown (Black). d is the distinctness-of-image parameter defined as $d = 1 - \alpha$, where α is the roughness parameter in the original Ward model. Adapted from [PFG00] with author’s permission.

ibration is applied to two alternative metrics on the BRDF images: RGB (without cubic root) and LAB L^2 difference. Although the shape of the function for our metric is slightly different from Pellacini et al.’s linear regression, our metric is consistent with the ratings reported. The LAB L^2 result is marginally better than our metric, while the RGB L^2 metric, in contrast, deviates significantly from both estimates and overly emphasize differences close to the mirror-like range. To summarize, the form of this image-difference metric allows us to reduce the computation cost while offering similar performance to the LAB metric.

More importantly, it is the use of rendered images as opposed to BRDF-space differences that makes our metric unique. In contrast, experiments with metric directly computed in the BRDF angular domain are unsuccessful. Even with the cubic root remapping applied directly to the BRDF data, the results are as poor as the RGB L^2 image metric. We hypothesize that the use of rendered images imposes a convolution to the BRDF [RH01, DHS*05] and emphasizes different features than the original BRDF data, which better captures our perception of material.

In Figure 8 we compare our image-driven metric to the BRDF-space metric defined in the angular domain. In each of the two rows the Ward model is varied over the same parameter range, and samples are chosen uniformly according to the BRDF-space metric (first row) and our image metric (second row). Another example is shown in Figure 9 for the Lafortune model, where we compare the uniform parameter spacing (first row) and our image metric (second row), when

the parameter c_z is varied with a fixed exponent n . In both cases, the spacing with our metric is much more uniform in terms of appearance, and as such would be easier for the user to control. We have found this observation consistent across the different parameters of the BRDF models we use.

For the Lafortune model, the c_z parameter is particularly difficult to navigate without our remapping. This is not a critical issue for its initial motivation, data fitting, but the model is more and more used in other applications because it is flexible and efficient to evaluate. When $c_z = -c_x$, the Lafortune lobe is equivalent to the Phong lobe. However, the model is very sensitive with respect to c_z when the exponent is high: even small deviation of c_z can lead to huge bias of the lobe towards normal or grazing angle. In the first row of Figure 9 an important region of the parameter space is skipped when c_z is varied linearly, as most of the visually interesting behavior of the model is highly concentrated near the value when $c_z = -c_x$. Our metric is able to offer a much more uniform spacing, and would allow a user to interactively explore the expressiveness of the model more easily.

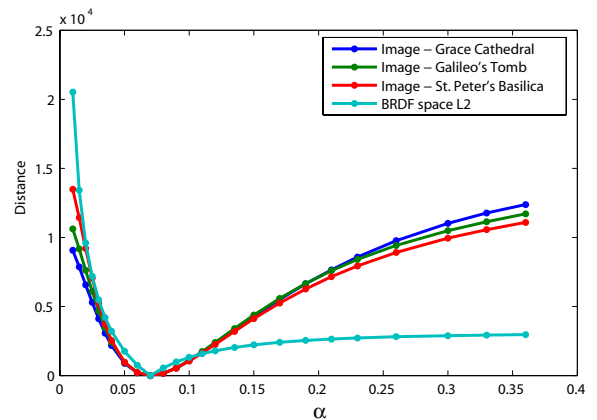


Figure 3: Plot of distances from the Ward BRDF at $\alpha = 0.07$ to 25 samples ranging from $\alpha = 0.01$ to 0.36. We compute our image-based metric with renderings using three different environment maps and compare to the BRDF-space L^2 metric. The distances from the different metrics are brought to the same scale by minimizing the least-square errors over all (25×24) pairwise distances.

Scene Dependence Next, we evaluate the sensitivity of our metric to the choice of the particular environment map. We choose a set of 25 Ward BRDFs and render the corresponding images with two additional environment maps. We compute all pairwise distances (25×24) using the image metric on each set of images, and also the BRDF-space L^2 metric for comparison. We find the scale factor between the metrics by least square fitting. We have found that over 85% of the distances for the two additional environment maps are within 20% of our reference distances. In Figure 3 we plot the distances of the 25 samples from a particular sample at $\alpha =$

0.07. We conclude that our metric is consistent regardless of the used environment map. We emphasize, however, that it is critical to use a *natural* environment map [DLAW01] that exhibits enough complex features. In particular, trivial environment maps such as a constant grey sphere would ignore highlight shapes. Fleming et al. [FDA01] have shown that complex natural illumination greatly facilitates the recognition of materials. Furthermore, Dror et al. [DAW01] have demonstrated that a number of characteristics of rendered BRDFs are invariant to a large class of real-world environment maps.

We have not evaluated our metric’s dependence on scene geometry. As each pixel is compared independently, our metric only depends on the distribution of normals instead of the actual shape. For a rendered sphere, normals that are close to grazing (with respect to the view direction) are sampled more sparsely. If a different bias is desired, we can alter our metric to weigh the pixels in a non-uniform way. However, the effects of using non-convex objects would require further study, as self-shadowing and other global illumination effects come into play.

3. Fast Distance Computation and Navigation Interface

Our navigation interface presents the user with choices that are distributed uniformly according to our metric. This requires on-the-fly image generation and distance computation. We can achieve fast image generation using pre-computation and PCA. However, we still need to compute many image differences and the need for image reconstruction and cubic-root non-linearity makes this computation costly. This is why we introduce a new embedding where our metric corresponds to the Euclidean distance, which affords dramatic speed up. Note that this embedding space is different from the one used for image generation. We then present our interface.

3.1. Image Pre-rendering

As images of arbitrary BRDFs under an environment map cannot be rendered interactively, we pre-render a set of images for each model by sampling its non-linear parameters. Linear parameters are applied on demand. The images are rendered at a resolution of 320 by 320 and principal component analysis (PCA) reduces the data size. The sampling density of each model is shown in Table 1. For simplicity, the isotropic single-lobe version of the Lafortune model is employed. With this assumption, we can reduce redundant degree of freedom by setting $c_x^2 + c_y^2 + c_z^2 = 1$. As $c_x = c_y$ for an isotropic Lafortune lobe, the parameters c_z and n are enough to fully specify the lobe up to a scaling factor. The scaling factor can be ignored as we normalize the lobe approximately such that the BRDF have the same brightness when c_z or n is varied. In the current implementation, we only sample c_x in the negative range which precludes retro-

Model	Sampling Grid	Parameters
Ward	$(9) \times 15 \times 15$	α_x, α_y
Blinn-Phong	$(9) \times 15$	n
Cook-Torrance	$(9) \times 11 \times 13$	F_0, m
He et al.	$(9) \times 8 \times 9 \times 5$	σ, τ, n
Lafortune et al.	$(9) \times 14 \times 13$	c_z, n

Table 1: Sampling density of various models. The first dimension in parentheses corresponds to the (ρ_d/ρ_s) dimension, which is only applicable to the embedding calculation, but not the linear prerendering.

reflection-like behavior. This limitation can be easily lifted by expanding our pre-rendering domain.

3.2. Embedding in a Unified Euclidean Space

In order to facilitate uniform navigation according to the metric, distances in the local neighborhood are required within a model and between models. In theory, one could compute the metric from the images on the fly, but the computational cost is too high when a number of neighbors needs to be considered. This is why we introduce a fast approximation of the metric based on the embedding in a Euclidean space with PCA. This is different from typical usage of PCA for compression purposes: we only want to compute distances and do not need to reconstruct the data points.

Observing that our metric is defined as the Euclidean distance between the cubic root of the two images, we can approximate the metric by embedding the set of post cubic root images in a low-dimensional vector space. However, in this case we can no longer ignore the linear parameters of the BRDF due to the cubic root. We can write the post cubic root image I as

$$I(\rho_d, \rho_s, \mathbf{p}) = (\rho_d I_{diffuse} + \rho_s I_{specular}(\mathbf{p}))^{\frac{1}{3}} \\ = \rho_s^{\frac{1}{3}} \left(\frac{\rho_d}{\rho_s} I_{diffuse} + I_{specular}(\mathbf{p}) \right)^{\frac{1}{3}} \quad (1)$$

where $\rho_s^{\frac{1}{3}}$ can be seen as a global scale. As a result, in addition to the nonlinear parameters, we also need to sample along the $\frac{\rho_d}{\rho_s}$ parameter. We sample each of the five models with regular grids (Table 1). Next we seek to embed all five models into a single embedding space. Conventional PCA, which requires the data to be mean-centered, is unsuitable for our purpose, as linear scaling of the image does not correspond to a simple scaling of the embedding coordinates. We enforce the center to be zero to allow linear scaling and use uncentered PCA [Jol02] to compute the embedding coordinates. While uncentered PCA does not have the optimality properties of standard PCA, the approximation errors still decay very quickly. Distance between the BRDFs according to our image metric can be approximated efficiently in this

space using the Euclidean distance. In practice, we conservatively use the first 200 coefficients, which means that our image metric becomes the Euclidean distance for 200 dimensions, as opposed to the difference between cubic roots of 320×320 pixels, thereby achieving dramatic speed-up. This allows us to employ a more sophisticated algorithm for the neighborhood construction when navigating from one model to another (Section 4).

3.3. Interface Overview

Our navigation interface is based on visual variations (Figure 4). The user is presented with a number of neighboring BRDF images surrounding the current one. The user can choose the desired image by clicking on it. There is only one real-valued slider in the interface – the radius of the neighborhood. A typical user navigation begins with a large radius, which will then be reduced gradually as the user converges to the desired BRDF.




Figure 4: Screenshot of the navigation interface. The current model is the Cook-Torrance model, and the user is at the roughness/grazing tab. The center image is the current BRDF, and the surrounding ones are the four equidistant neighbors.

Most BRDF models have more than a few parameters, which is why we divide the neighbors into three different tabs. In the first two tabs, we show the variations in the color and intensity of the diffuse and the specular scales of the BRDF. In the third tab, we show the variations due to the other parameters depending on the chosen model (roughness, Fresnel factor, etc.), which typically affect the shape of the specular lobe. In this tab, the neighbors are chosen according to our image-driven metric described in Section 2.1.

The current implementation of the interface allows

the user to navigate within the space of 5 different BRDF models. These include Blinn-Phong [Bli77], Ward [War92], Cook-Torrance [CT81], Lafortune [LFTG97] and He [HTSG91]. The anisotropic version of the Ward model is employed while the other models are enforced to be isotropic. We use the implementation of the He et al. model by Rusinkiewicz where polarization is ignored [Rus]. Adding more models to our interface would be straightforward.




Figure 5: Equidistant neighbors (small red circles) are found by walking along the isoparameter lines (dashed lines). The neighbor at the desired distance is found when the segment intersect the circle (sphere). The orange lines highlight the grid samples that are queried during the search. The black lines indicate other grid samples.

Let the user-specified distance be ϵ . In the diffuse/specular color tabs, we define the neighborhood to be the sphere of radius ϵ in the LAB space, centered around the current color. We display 2 neighbors in the L (luminance) direction, and 8 neighbors in the a-b chromaticity plane. We could interchangeably have used the cubic root of R, G, and B for these parameters as well, but since we have no computational limitation, we have chosen to use standard CIELAB.

The specular lobe tab (third tab) depends on the parameters used in the chosen analytical model. For these parameters we need to find equidistant neighbors according to our image-driven metric. For illustration, let us assume that the BRDF has only two parameters, and thus the sample grid for the distance precomputation is a 2D lattice (Figure 5). First, we find the embedding coordinate of the current BRDF using a multi-linear interpolation. Next we seek neighbors in the 4 different parametric directions. Starting from the current point (red point in the center), we walk along the chosen parameter keeping the other parameters unchanged (dashed line). As we advance along the path, the current segment is intersected with the ϵ circle. The distance computation for this intersection exploits the post-cubic-root PCA coefficients described above. This algorithm is generalized to higher-dimensions in our application.

4. Navigating Across Different Models

An important choice in material reflectance design is that of the analytical model itself. BRDF models differ in their expressiveness and the class of materials that they can represent. Our interface does not require BRDF expertise or trial-and-error of different models. In order to achieve this, we show images of material appearance that are at a certain distance from the current pick but cannot be obtained with the current analytical model. This requires two different features: we need to convert between two different analytical models, and we need to evaluate which material appearance afforded by other BRDF models cannot be achieved using the current analytical model.

We compute conversions between the various analytical models by fitting a discrete set of samples of each model to every other model. In practice, the fit is computed using the L^2 BRDF-space metric because it is more computationally efficient than the image-driven metric. In addition, we are seeking near-threshold matches, which is different from spacing BRDFs uniformly at a suprathreshold distance. Our experiments have indicated that fits obtained using the two metrics are similar, which contrasts the different suprathreshold behavior observed in Section 2.1. Fitting a single BRDF to a target model takes about 10 minutes on average on a single PC, but the computation needs to be performed only once. We sample the BRDF models with the same grid as for the embedding space (Table 1) and store all the pairwise conversions.




Figure 6: Illustration of the manifolds spanned by two analytical BRDF models in an abstract unified BRDF space. Point on the black curves represent instance of two BRDF models. Given a current BRDF, we want to find BRDFs on an alternative model that are close but cannot be represented by the current analytical model. We wish to propose materials such as B' because its distance $d_{B'}$ to the current model is large. See text for detail.

To construct the desired neighborhood, we first want to find neighbors in all models which are at distance ϵ from the current BRDF. Using the precomputed conversion and multi-linear interpolation, we can find the BRDF A' closest to the current pick A on the manifold spanned by a different analytical model (Fig. 6). Starting from point A' in the new model, we can search for neighbors in the same way

as described in the previous section. The only difference is that the ϵ sphere is now centered around A instead of A' . We repeat the same search for each model pair.

For each neighbor (B', C') we look up the best-fit BRDF in the current model and compute the distance ($d_{B'}, d_{C'}$). We remove neighbors that are less than a threshold distance δ from the current model. We empirically choose $\delta = 0.3\epsilon$.

As a final step, we further trim the number of neighbors by clustering. This is necessary for cases when the number of neighbors is too large to display in the interface. In this case, we would like to show neighbors that are most different from each other. A standard solution is to use k -means clustering on the set of neighbors represented by their embedding vectors. A single neighbor is then chosen randomly from each cluster to form the neighborhood. While the k -means algorithm, in general, does not give the global optimum, the method works well in our application. Note that a large number of metric computations happen in this phase and our low-dimensional embedding is crucial for reducing the computation time. In fact, the update time to move to a new selection is less than 0.7 seconds in the worst case.

This solution can be seen as a simplified way of navigating in the manifold spanned by all BRDF models in the spirit of Matusik et al.'s work on data-driven BRDFs [MPBM03]. Our manifold-hopping solution allows us to implicitly navigate such a global manifold without explicitly building it. Topological and metric information about this manifold is encoded by the conversion and distance information.




Figure 7: The conversion tab showing the neighbors in the union space of all models. Note that all the neighbors show some effects which are not expressible with the current model (Blinn-Phong).

5. Discussion

In this paper we have proposed a new distance metric for BRDFs. Each BRDF is represented by a corresponding rendered image under a natural illumination environment. We show that while this metric is not directly derived from psychophysics measurements, it reasonably reflects the visual differences between materials in a uniform way. Given this metric, we have built a user interface for navigating in the BRDF space. Our interface is intuitive and simple: the only non-visual parameter is the neighborhood radius. Neighbors in different parameter directions at the specified distance from the current BRDF are shown and can be selected by a mouse click. In addition, conversions between models are precomputed and the user can readily see the neighbors in other models and switch to them if desired. With our interface, the user is not required to understand the intricate complexities and differences between the different models. Instead, the user can always see the neighboring BRDFs across different models and freely jump between them. The accompanying video demonstrates our method and user interface.

In our experience, the ability to *previsualize* the result of the next navigation step is key to the effectiveness of the interface. In contrast to slider-based interfaces, our approach allows the user to directly see possible options and does not require creation of a mental model for the effects of various parameters.

The dependence of our metric on an environment map deserves further study. On one hand, we believe that the environment map provides images that better represent real-world usage of BRDFs; and it has been shown that the complexity of real-world environment maps greatly facilitate the recognition of materials by humans. On the other hand, there is something arbitrary in choosing a given environment map, even though our experiments show that the metric is robust to this choice. We believe that the most exciting question is to further characterize what is special about natural environment maps using tools such as Fourier analysis, wavelets, and derivatives [DLAW01]. Once natural environment maps are better understood, one can hope to directly define a metric in the BRDF domain. Recent findings on the signal processing interaction between illumination and BRDFs will likely be important [RH01, DHS*05].

A second limitation of our current approach is dictated by the original parameterization. The different parameters of a BRDF model are not necessarily “perceptually orthogonal” and while the different choices we present around the current selection are on a perceptual circle, they might not be uniformly distributed on this circle. Pellacini et al. [PFG00] defined new axes for the Ward model based on multi-dimensional scaling of the reported distances. It would be a bigger challenge to define new axes that are meaningful and uniform in the unified space of multiple BRDF models.

There are other avenues for future work. First, we currently only support a single specular lobe in our interface.

Additional lobes increase the dimensionality of the space, and organizing them in a meaningful way is challenging. Next, our interface is limited to a predefined scene and environment. To allow for arbitrary scenes/environments, pre-computed rendering techniques can be employed [RH02, SKS02]. The choice of the initial BRDF at the beginning of navigation is another important issue. An interface similar to design galleries [MAB*97] would greatly facilitate this first step and refine our *k*-means approach. While we have leveraged data from Pellacini’s perceptual study to validate our image-driven metric, a logical next step is to perform a formal user study in order to better evaluate the effectiveness of our interface. Finally, we believe that visual interfaces such as design galleries and variations-based interfaces like the one we have presented have important applications in all areas of computer graphics, and can significantly enhance user’s experience when dealing with complex parameter spaces.

Acknowledgement We thank Eric Chan, Jan Kautz, Barb Cutler, Kari Pulli, Tom Mertens, Soonmin Bae, Sara Su, Tilke Judd and the anonymous reviewers for valuable feedback on this paper, and Tom Buehler for producing the accompanying video. This work was supported by a National Science Foundation CAREER award 0447561 “Transient Signal Processing for Realistic Imagery.” Frédo Durand acknowledges a Microsoft Research New Faculty Fellowship.

References

- [Bak01] BAKER A.: Learning from Photoshop’s variations tool. <http://www.merges.net/theory/20010308.html>, Mar. 2001.
- [Bli77] BLINN J.: Models of light reflection for computer synthesized pictures. In *Proceedings of SIGGRAPH 1977* (1977), pp. 192–198.
- [BM98] BOLIN M. R., MEYER G. W.: A perceptually based adaptive sampling algorithm. In *Proceedings of SIGGRAPH 1998* (1998), pp. 299–310.
- [CPK06] COLBERT M., PATTANAIK S., KRIVANEK J.: BRDF-Shop: Creating physically correct bidirectional reflectance distribution functions. *IEEE Comput. Graph. Appl.* 26, 1 (2006), 30–36.
- [CT81] COOK R. L., TORRANCE K. E.: A reflectance model for computer graphics. In *Proceedings of SIGGRAPH 1981* (Aug. 1981), pp. 307–316.
- [Dal93] DALY S.: *The Visible Differences Predictor: An algorithm for the assessment of image fidelity*. Digital Image and Human Vision. MIT Press, 1993, pp. 179–206.
- [DAW01] DROR R., ADELSON E., WILLSKY A.: Recognition of surface reflectance properties from a single image under unknown real-world illumination. In *Proceedings of the Workshop on Identifying Objects Across Variations in Lighting at CVPR* (Hawaii, Dec. 2001).

- [DHS*05] DURAND F., HOLZSCHUCH N., SOLER C., CHAN E., SILLION F. X.: A frequency analysis of light transport. *ACM Transactions on Graphics* 24, 3 (Aug. 2005), 1115–1126.
- [DLAW01] DROR R., LEUNG T., ADELSON E., WILL-SKY A.: Statistics of real-world illumination. In *Proceedings of CVPR* (Hawaii, December 2001).
- [Dur02] DURAND: An invitation to discuss computer depiction. In *Proc. NPAR* (2002).
- [Fai98] FAIRCHILD: *Color Appearance Models*. Addison-Wesley, 1998.
- [FDA01] FLEMING R., DROR R., ADELSON E.: How do humans determine reflectance properties under unknown illumination? In *Proceedings of the Workshop on Identifying Objects Across Variations in Lighting at CVPR* (Hawaii, Dec. 2001).
- [HTSG91] HE X. D., TORRANCE K. E., SILLION F. X., GREENBERG D. P.: A comprehensive physical model for light reflection. In *Proceedings of SIGGRAPH 1991* (1991), pp. 175–186.
- [Jol02] JOLLIFFE I.: *Principal Component Analysis*. Springer Verlag, New York, 2002.
- [KPC93] KAWAI J. K., PAINTER J. S., COHEN M. F.: Radioptimization - goal based rendering. In *Proceedings of SIGGRAPH 1993* (Aug. 1993), pp. 147–154.
- [LFTG97] LAFORTUNE E. P. F., FOO S.-C., TORRANCE K. E., GREENBERG D. P.: Non-linear approximation of reflectance functions. In *Proceedings of SIGGRAPH 1997* (1997), pp. 117–126.
- [MAB*97] MARKS J., ANDALMAN B., BEARDSLEY P. A., FREEMAN W., GIBSON S., HODGINS J. K., KANG T., MIRTICH B., PFISTER H., RUMI W., RYALL K., SEIMS J., SHIEBER S.: Design galleries: A general approach to setting parameters for computer graphics and animation. In *Proceedings of SIGGRAPH 1997* (1997), pp. 389–400.
- [MPBM03] MATUSIK W., PFISTER H., BRAND M., MCMILLAN L.: A data-driven reflectance model. *ACM Transactions on Graphics* 22, 3 (July 2003), 759–769.
- [Mys98] MYSZKOWSKI K.: The visible differences predictor: Applications to global illumination problems. In *Rendering Techniques '98* (1998), pp. 233–236.
- [PFG00] PELLACINI F., FERWERDA J. A., GREENBERG D. P.: Toward a psychophysically-based light reflection model for image synthesis. In *Proceedings of ACM SIGGRAPH 2000* (2000), pp. 55–64.
- [RH01] RAMAMOORTHY R., HANRAHAN P.: A signal-processing framework for inverse rendering. In *Proceedings of SIGGRAPH 2001* (2001), pp. 117–128.
- [RH02] RAMAMOORTHY R., HANRAHAN P.: Frequency space environment map rendering. In *Proceedings of SIGGRAPH 2002* (2002), pp. 517–526.
- [RPG99] RAMASUBRAMANIAN M., PATTANAIK S. N., GREENBERG D. P.: A perceptually based physical error metric for realistic image synthesis. In *Proceedings of SIGGRAPH 1999* (Aug. 1999), pp. 73–82.
- [Rus] RUSINKIEWICZ S.: bv - a BRDF browser. <http://graphics.stanford.edu/~smr/brdf/bv/>.
- [RWP*95] RUSHMEIER H., WARD G., PIATKO C., SANDERS P., RUST B.: Comparing real and synthetic images: Some ideas about metrics. In *Rendering Techniques '95* (1995), Springer-Verlag Wien New York, pp. 82–91.
- [SDS*93] SCHOENEMAN C., DORSEY J., SMITS B., ARVO J., GREENBERG D.: Painting with light. In *Proceedings of SIGGRAPH 1993* (1993), pp. 143–146.
- [SKS02] SLOAN P.-P., KAUTZ J., SNYDER J.: Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. *ACM Transactions on Graphics* 21, 3 (July 2002), 527–536.
- [TR93] TUMBLIN J., RUSHMEIER H. E.: Tone reproduction for realistic images. *IEEE Computer Graphics & Applications* 13, 6 (Nov. 1993), 42–48.
- [War92] WARD G.: Measuring and modeling anisotropic reflection. In *Proceedings of SIGGRAPH 1992* (1992), pp. 265–272.
- [WM01] WESTLUND H. B., MEYER G. W.: Applying appearance standards to light reflection models. In *Proceedings of SIGGRAPH 2001* (2001), pp. 501–510.