

Image-Based BRDF Design

Ezra Davis¹, Weiqi Shi¹, Holly Rushmeier¹, Julie Dorsey¹, Hongzhi Wu²

¹Department of Computer Science, Yale University, New Haven, CT USA

²State Key Lab of CAD&CG, Zhejiang University, Hangzhou, China

Abstract

Material appearance design usually requires an unintuitive selection of parameters for an analytical BRDF (bidirectional reflectance distribution functions) formula or time consuming acquisition of a BRDF table from a physical material sample. We propose a material design system that can take visual input in the form of images of an object with known geometry and lighting and produce a plausible BRDF table within a class of BRDFs. We use the principal components of a large dataset of BRDFs to reconstruct a full BRDF table from a sparse input of BRDF values at pairs of incoming and outgoing light directions.

To get visual user input, we allow the user to provide their own object and then generate guide images with selected lighting directions. Once the user shades the guide images, we construct the reflectance table and allow the user to iteratively refine the material appearance.

We present preliminary results for image-based design, and discuss the open issues remaining to make the approach practical for mainstream use.

Introduction

Material appearance design has a wide variety of applications, both virtual and physical. The design of virtual material appearance is needed in films and games to assist in setting a mood, conveying a sense of realism and telling a story. Virtual material appearances must be designed to create effective and reliable educational and training materials for either video or virtual/artificial reality in applications such as industrial training and cultural heritage education. The design of physical material appearance is critical in the design of essentially all consumer products – including clothing, furniture, and automobiles.

Over the last thirty years in the field of computer graphics a variety of computer systems have been developed for the design of virtual and physical objects. Extensive work has been done to allow users to design the geometry of objects, but much less to allow them to design materials. In many systems material appearance has been limited to specifying the color of diffuse reflectance, the size of a specular lobe, and spatial variations of these simple characteristics. Much recent research in computer graphics has focused on mathematical representations of material appearance, and how those representations can be used efficiently in generating an image. Research in the specification of the representations beyond simple diffuse plus specular has lagged. As a result, current systems are strongly tied to the specifics of mathematical representations. The systems are difficult to use, and limited in their expressive power.

In this work we explore a new approach for the design of material appearance (illustrated in Figure 1), and in particular for designing BRDFs (bidirectional reflectance distribution functions.)

We seek to address two major problems – the unintuitive nature of current systems, and the limited gamut of appearance they allow. We pursue an image-based approach, so that the designer can focus solely on the appearance, rather than on the mathematical form used to define a BRDF. An image-based approach is not restricted to specifying BRDFs within the gamut of analytical functions.

Designing with images however raises new issues. What is the best numerical representation for the BRDF in this context? How can the numerical representation be extracted from the user prepared images? We address these questions by using BRDF tables as our base representation and by using user prepared images with known geometry and lighting. Can users prepare the images we propose as input? Do they find creating such images more intuitive than current systems? We demonstrate some image editing/painting techniques for creating the images, and report feedback from some early users of a prototype system.

Previous Work

Our work builds on previous efforts in designing interfaces for BRDF specification and in populating dense BRDF tables from sparse measurements.

BRDF creation systems

There is a long history of systems for creating BRDF models to render 3D scenes [5, 11]. Most systems let artists control the inputs to simple, well known, analytic BRDF models. The left image in Figure 2 shows a typical example of a simple BRDF editor in the Blender software system [2]. The simple editor allows the artist to switch between a few common analytic BRDF equations and modify their parameters, typically roughness, diffuse color, and specular color. The right image in Figure 2 shows a more complex visual programming interface (also in Blender) inspired by Abram and Whitted’s “Building Block Shaders” [1] and subsequent work by McGuire et al. [8]. Visual programming interfaces are now common and allow artists to have much more control over the final result. Such visual programming interfaces however require a large number of parameters. Zsolnai-Fehér et al. [14] streamline the process of setting the 19 parameters for a BRDF specification by presenting the user with a gallery of models and letting the user rate successive galleries of algorithmically selected BRDFs until the artist is happy with the results.

Research projects have introduced painting-based material design approaches that avoid tuning parameters. Sloan et al. [12] presented the idea of painting on a spherical canvas to capture custom artistic shading models for nonphotorealistic rendering. However, the gamut of materials their system can create is not physically plausible. Based on their idea, Colbert et al. [3] proposed an intuitive painting mechanism, BRDF-Shop, which takes




Figure 1. Creation of a BRDF table via PCA (principal component analysis) from image(s) created with known geometry and known lighting directions.

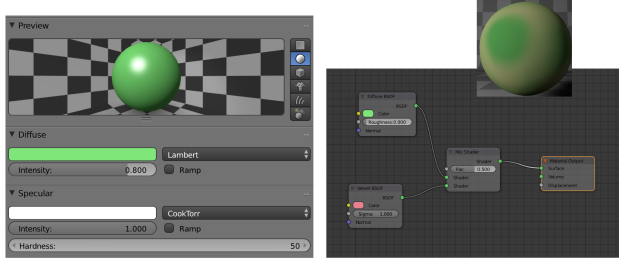


Figure 2. Common material appearance design interfaces: simple parameter adjustment (left), visual programming interface (right).

the painted highlights as input and approximates them with analytical BRDF models. Our method, in contrast, allows users to paint the full material without being limited to shaping highlights.

Reconstruction of BRDF tables

An alternative to specifying analytical BRDF models is to use measured values, such as those in the MERL database of isotropic measured BRDFs [7]. Acquiring measured BRDFs requires a material sample, a goniometer or similar measuring device, and a time-consuming measurement process. Large numbers of measurements are required to populate a BRDF densely enough so it can be used for rendering. Nielsen et al. [9] proposed a technique to speed up material acquisition by taking a few samples of the material and filling in unknown values by using the principle components of a database of previously measured BRDFs. Romeiro and Zickler [10] use a similar method (non-negative matrix factorization) to extract BRDFs and environment maps from a single image. We use a similar method to Nielsen et al. to reconstruct dense BRDF tables from sparse user input.

Proposed New System

Overview

We created a new prototype system that allows users to visually specify material appearances via a user interface and reconstructs BRDFs using a data-driven solution. Our interface provides the guidance to help users visually design materials using any image-editing software. Given an appropriate image with the desired material, our system first extracts BRDF values to create a sparse BRDF table and then conducts PCA (principle component analysis) with a simulated BRDF dataset to recreate the full BRDF for rendering. We also created a user interface to guide users through the image-based material design process.




Figure 3. Screenshot of the proposed BRDF creation interface early in the BRDF creation process.

Pipeline

Image-Based Material Editing. It is known that humans can be insensitive to changes in illumination, and describing how lighting interacts with material appearance is difficult. To help users visually design materials, our system provides guide image(s) with user specified geometry and lighting conditions. In a guide image, the object is lit by a single directional light, which simplifies users' visualization of the material and aids in extracting datapoints for the BRDF. The guide image is shaded with a simple diffuse shader as a starting point. Notice that the lighting used for the guide image is white balanced and normalized, which allows us to extract the BRDF values later using image-based BRDF techniques. Given an exported guide image, users can augment the material and paint the image by modifying the existing shading, adjusting the contrast and color balance, painting over the image using brushes, or applying other effects using any image editor. The modified image is then imported into our software.

BRDF Table Format. Since the image with designed materials has fixed viewing directions and known lighting and normals, each pixel in the modified image can be treated as a BRDF sample. After users finish their design, we extract the BRDF values from the given image and store them in a sparse BRDF table. Our BRDF format is similar to UTIA's material format [6] but we use a higher resolution to avoid mapping too many pixels to the same location in the BRDF table. Our BRDF table has 1620 lighting directions and 1620 viewing directions, where lighting and viewing directions are the combination of 45 and 36 uniformly sampled




Figure 4. The BRDF table format used in our system.

points from the polar angle θ and the azimuthal angle ϕ respectively in spherical coordinates (The UTIA format is 288×288 , where $288 = 6$ polar angles $\times 48$ azimuthal angles.) We picked our BRDF table structure so that it minimizes the discontinuities between patches. Our BRDF table (Figure 4) contains patches of constant θ with varying ϕ . There are no sharp discontinuities between the patches because the ϕ value goes from 0 to 2π in each patch. While there may be more meaningful spatial orderings, this ordering, when used for our BRDF reconstruction, provides results that are accurate enough.

Acquiring a Sparse BRDF. To acquire a sparse BRDF table from an artist’s rendition, we process each pixel in the painted image with the following steps:

1. Construct a frame coordinate system tangent to the surface. This serves as the BRDF’s coordinate system. We use the pixel’s normal and a tangent direction to indicate where $\phi = 0$. To select the first tangent direction, we project a previously chosen vector onto the surface. In the future, we could potentially use the object’s shape to control the direction of $\phi = 0$.
2. We next convert the view and lighting direction into our new coordinate system, using spherical coordinates.
3. We write the color of that pixel into the location indicated by the transformed values of the view and lighting directions.

Since users are asked to create images of objects that may contain multiple locations with the similar normal and lighting combinations, it is possible that the same points on the sparse BRDF table are set multiple times with different values from the input image. However, in our experiments, the number of the inconsistent points are less than 1% of the total number of data points, which applies to all the objects we tested. We examined different strategies to handle those inconsistent points: by averaging all the possible values, or selecting the minimum or maximum value for each single inconsistent point. We found there is no significant difference in the reconstructed BRDF results (Figure 5). Therefore, we believe the inconsistent points do not significantly influence the final results.

Dense BRDF Reconstruction. Inspired by the work from Nielsen et al. [9], we conduct PCA to reconstruct a full BRDF from the sparse BRDF table. There are two differences between our method and the method described by Nielsen et al.: We start with a more, perhaps less optimal, samples in our sparse BRDF, and use a different BRDF database.

Unlike Nielsen’s work which uses the MERL dataset, we created a new database which contains anisotropic BRDFs. We generated these BRDFs using a bi-scale material model based on




Figure 5. Results from different strategies to handle inconsistent points: one of several painted input images (left), the reconstruction using the average (center left), reconstruction using the minimum (center right) and reconstruction using the maximum (right).

the method by Wu et al. [13] (see the paper for simulation details). From this large database, we apply PCA to extract the first 20 principle components for reconstruction. Similar to the work of Nielsen et al., we apply the ridge regression to reconstruct the full table from the sparse one. We experimented with up to a hundred principle components, but did not notice any significant gains in quality of the reconstruction compared to Nielsen et al.’s 20 components; additionally, 20 components allows us to get significantly faster feedback, only taking approximately 4.5 seconds to reconstruct a BRDF on a 2013 Macbook Pro, as opposed to many minutes for 100 components.

User Interface. We created a user interface to guide users through the image-based materials design process. The user interface is shown in Figure 3. Our user interface has six steps, from setting up the guide image to downloading the final BRDF. To create a material, the user first selects an object to paint over (step 1) and the lighting conditions (step 2) for that object. Organic or curved objects are better than angular flat objects for reconstruction because they provide more unique normal directions. The user next downloads the guide image (step 3) and edits it in an external image editor (step 4). The next step is to import the painted image into our application (step 5) and it will take that image and reconstruct a full BRDF to render a preview (step 6). If the user is unsatisfied with the result, they can create a second or more images to provide additional information for creating the full BRDF by repeating the process, starting with designing a second guide image.

There are a few additional visualization tools available to aid in designing a BRDF:

- Background: There is the option to add a background cube map to help gain additional context in the guide image.
- Shadows: Adds shadows cast from the guide object, which is very useful in determining the light direction.
- Full Reconstruction: Turns on or off the visualization of the BRDF, if available.
- Sparse Reconstruction: Shows a reconstruction of the object using only information in the BRDF directly contained by the previous painted image(s).

Preliminary Results

Preliminary versions of the interface were provided to members of our research group, an artist, and a computer science undergraduate to explore how difficult the interface is to use. This did not constitute a formal user study because we asked different questions to different people during the development process. The testers had varying levels of familiarity with alternate material creation methods and digital painting.

All of the students we had test the interface would be willing




Figure 6. An example of a user's painted image (left) and a preview of the resulting BRDF (right).




Figure 7. Another example of a user's painted image (left) and a preview of the resulting BRDF (right). Note that we can get a good result even from a simple painted image. Also note the small amount of blue in the resulting BRDF that doesn't appear in the original painting.

to use a system like this to create novel BRDFs, though one of the people we interviewed would rather use a traditional method when possible. Overall, the user experience was intuitive except for the light direction controls; both student 1 and 2 said that would rather have the ϕ and θ values control elevation and azimuth of the light source instead of angle relative to the camera. Student 3 also had some issues with initially determining the light controls. She used the background image to help figure out how the controls worked, but the other students didn't consider it useful, and even student 3 would consider removing the background from the interface.

All of the students considered the preview of the object with the reconstructed BRDF far more useful than the sparse BRDF preview, and both student 1 and 2 did not consider the sparse BRDF to be useful in either visualizing what they already drew or in making a consistent second image.

All of the students considered the material created after using a single edited or painted image to be close enough to what they envisioned that they did not need to go back and draw from another light direction despite some visual artifacts (mostly colors showing up in the BRDF that did not appear in the original painted image). In fact, student 2 had some issues drawing a specular highlight when attempting to mimic a satin material, but due to how the reconstruction works, felt that the reconstructed material was actually better than what he'd drawn (Figure 6). Student 3 said that if she was attempting to precisely match a material she would definitely have drawn a second image to correct some blue highlights that did not appear in her drawn image, visible in Figure 7. If we reconstructed the BRDF with a database of BRDF tables with more variety, there would have been fewer artifacts.

All of the students considered the shadow very useful in determining the light direction.




Figure 8. Our system does not handle textured images, simply assuming that they are intended to represent angular variations on a single point. Left is the painted image, right is a preview of the reconstructed BRDF.

Limitations and Future Work

While it demonstrates a “proof-of-concept”, our BRDF creation system has several limitations. Our future work will be focused on the following challenges:

- In our system, artists create their images in traditional low dynamic range image formats, and we produce a low dynamic range BRDF. Most measured materials require a high dynamic range BRDF table to fully capture the specular values. Our future research will investigate ways of converting a low dynamic range artistic rendering into a high dynamic range BRDF.
- Due to how we create the sparse BRDF, we cannot handle textured objects and cannot produce spatially varying BRDFs (SVBRDFs). There are systems for extracting simple analytic SVBRDFs from a single image [4]. We will pursue research in extracting more complex SVBRDFs.
- We use a single method for reconstructing full BRDF tables from sparse BRDFs. We will investigate alternate reconstruction methods. Additionally, we don't enforce reciprocity or any form of smoothness in our reconstructions.
- We used a relatively limited database of previously rendered BRDFs; expanding this will allow us to create a larger variety of BRDF tables.
- Alternate methods of presenting lighting direction and context of the guide image could potentially improve the user experience and result in BRDF tables that reflect the artists conception more accurately.
- We provide two methods of receiving feedback on the current state of the BRDF table, sparse and full reconstruction of the guide object. We will explore alternate methods of providing feedback and asking for refinement of the table.

References

- [1] Gregory D. Abram and Turner Whitted, Building block shaders, Proceedings of the 17th annual conference on Computer graphics and interactive techniques (SIGGRAPH '90) pp. 283-288 (1990).
- [2] Blender Foundation, Blender, blender.org (accessed January 2019).
- [3] Mark Colbert, Sumanta Pattanaik, and Jaroslav Krivanek, BRDF-Shop: Creating physically correct bidirectional reflectance distribution functions. IEEE Computer Graphics and Applications, 26, pp. 30-36 (2006).
- [4] Valentin Deschaintre, Miika Aittala, Fredo Durand, George Drettakis, and Adrien Bousseau, Single-image SVBRDF capture with a rendering-aware deep network, ACM Trans. Graph. 37, 4, Article 128 pp. 1-15 (2018).

- [5] Julie Dorsey, Holly Rushmeier, and François Sillion. Digital modeling of material appearance. Elsevier, 2008, pp. 227-242.
- [6] Michal Haindel, Jiří Filip and Radomír Vávar, Digital Material Appearance: the Curse of Tera-Bytes, ERCIM News 90, pp.49-50 (2012).
- [7] Wojciech Matusik, Hanspeter Pfister, Matt Brand, and Leonard McMillan, A data-driven reflectance model, ACM Trans. Graph. 22, 3, pp. 759-769 (2003).
- [8] Morgan McGuire, George Stathis, Hanspeter Pfister, and Shriram Krishnamurthi, Abstract shade trees, Proc. Symposium on Interactive 3D graphics and games, pp. 79-86 (2006).
- [9] Jannik Boll Nielsen, Henrik Wann Jensen, and Ravi Ramamoorthi, On optimal, minimal BRDF sampling for reflectance acquisition. ACM Trans. Graph. 34, 6, Article 186, pp. 1-11 (2015).
- [10] Fabiano Romeiro and Todd Zickler, Blind reflectometry, Proceedings of the 11th European Conference on Computer Vision: Part I (ECCV'10), pp. 45-58 (2010).
- [11] Thorsten-Walter Schmidt, Fabio Pellacini, Derek Nowrouzezahrai, Wojciech Jarosz, and Carsten Dachsbacher, State of the art in artistic editing of appearance, lighting and material, Computer Graphics Forum, 35, pp. 216-233 (2016).
- [12] Peter-Pike Sloan, William Martin, Amy Gooch, and Bruce Gooch. The lit sphere: A model for capturing NPR shading from art, Proc. Graphics interface, pp. 143-150 (2001).
- [13] Hongzhi Wu, Julie Dorsey, and Holly Rushmeier, Physically-based interactive bi-scale material design, ACM Trans. Graph. 30, 6, Article 145 pp. 1-11 (2011).
- [14] Károly Zsolnai-Fehér, Peter Wonka, and Michael Wimmer, 2018. Gaussian material synthesis. ACM Trans. Graph. 37, 4, Article 76, pp. 1-14 (2018).

Author Biography

Ezra Davis is a graduate student in computer science at Yale University. He received his BS degree in computer science with minors in electrical and computer engineering and interactive media and game design from Worcester Polytechnic Institute (WPI).