Approximating Per Vertex BRDFs Using Multiple Light Directions

Daniel Coutinho, Ricardo Marroquim
Universidade Federal do Rio de Janeiro
UFRJ
Rio de Janeiro, Brazil
www.lcg.ufrj.br

Matteo Dellepiane
Instituto di Scienza e Tecnologie dell'Informazione
CNR
Pisa, Italy
http://vcg.isti.cnr.it/

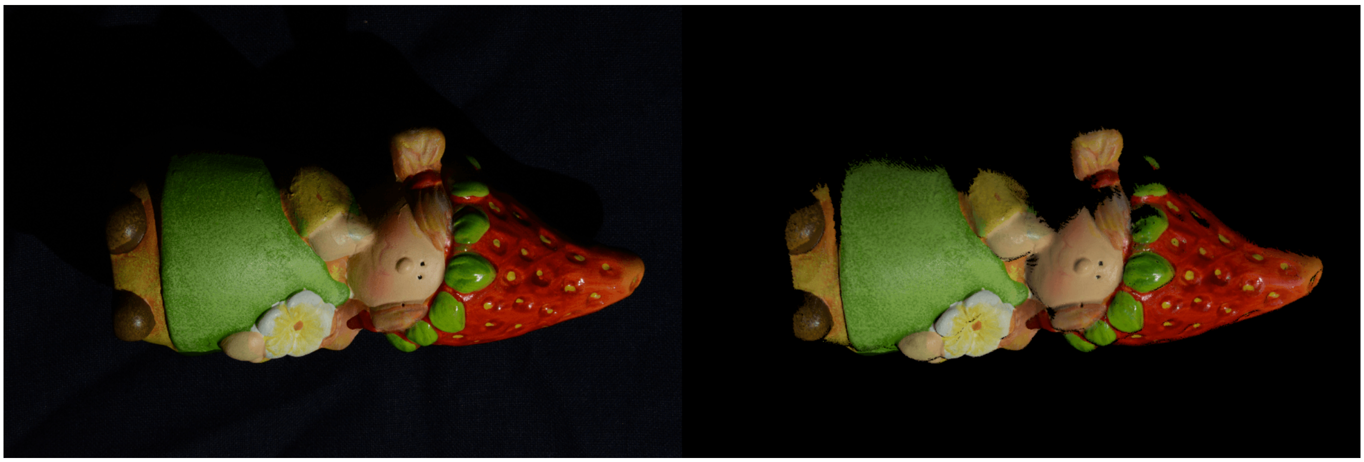

Fig. 1. Illustration of our method, left we can see a photograph of the real object and on the right our result. (right).

***Abstract*—This work presents an acquisition system that approximates a simple BRDF model per vertex of a real world object. Given a digital mesh from the physical object and photos taken with different light directions, our system faithfully represents the material that composes the object. The main advantage of our work is the ability to generate online results and provide immediate feedback to the user. We present an efficient system that doesn't require complex devices or an over controlled acquisition environment.**

***Keywords*-BRDF Approximation; Phong Model; Interactive Rendering;**

I. Introduction

Photorealistic rendering has been a goal in Computer Graphics almost since its beginning. A lot of focus has been placed on the development of global illumination algorithms[1], however, even the most advanced techniques rely on the description of the object's material properties. BRDFs are models that mimic the material behavior but for their parameters to be settled, some measurement of the specific material must be carried on. BRDF Approximation has been a topic of research in the last decades [2]. Our work consists in approximating a simple BRDF model using only photos with an associated light direction.

Even though there are works that almost generally solve the BRDF acquisition problem, they usually have some limitations regarding the object's size and position or special acquisition environments. Our work presents no restrictions about the object's dimensions and its position. Furthermore, some acquisition environment conditions are desirable, but not mandatory. Besides the desire of rendering photorealistic scenes, our work supplies the demand for other areas of research that focus on digital techniques applied to Cultural Heritage. A field focused on preservation, restoration and dissemination of historical artifacts. Although we can faithfully create a geometric copy of a real world object and extract information about its color, there's still room for improvements regarding the extraction of its reflectance properties.

A common pipeline for digitizing a historical artifact is to acquire its geometry and photos in situ, followed by offline post-processing. It is not uncommon during this step to notice that there is not enough data to fully represent the object. In certain campaigns this may be a major problem, since another trip to the site may not be viable.

Contributions: The main contribution of our work is a fast algorithm that can approximate a simple BRDF model of a real world object. Our system can provide the user with feedback about all the information required to faithfully represent the object and overall quality, and also suggest good light directions in order to have a better result with less data.

We are not aware of other tools used in the field that can provide such online response.

II. Related work

There are some approaches used in estimating BRDFs of real objects: controlled lighting solutions which use various images with fixed viewpoint but varying light directions, techniques that try to perform the acquisition in a general lighting environment, and some of them design specific devices for this task.

Lensch et. al [3] presents a fitting process using only a professional digital camera, a reflecting sphere and a dark room. They use the Lafortune BRDF model in their approach, but instead of having a BRDF model per vertex, they divide the mesh in clusters of BRDFs. Due to the clusterization it may happen that some areas may not be represented. The authors extended the previous work in [4] by changing the calibration of the light source position and estimating normal maps in order to refine the mesh's geometric details. The work of [5] fits a Polynomial Texture Map by solving a linear system for given N images using singular value decomposition. They also show that it is possible to apply filter behaviors on the PTM and some lighting models such as anisotropic surfaces and fresnel effects can be modeled. The work of [6] tries to recover at the same time the shape and the spatially-varying BRDFs of objects. The shading model is the isotropic Ward. Following the same approach the work of [7] recovers both shape and BRDF simultaneously, but instead of using a parametric reflectance model, it uses a bi-variate approximation of measured isotropic BRDFs, which they argue that it can represent a broader number of materials.

In [8] a statistical method for the estimation of Spatially Varying BRDFs is provided. Their approach is based in video sequences with fixed but general lighting conditions. A user assisted clustering process is also performed, since in the video some object areas may have not been appropriately specularly sampled. Some limitations are presented in this work due to the input data and the Phong model, it may also present blur effects. In addition, the clustering step may sometimes require too much manual intervention. The work presented by [9] also tackles the unknown lighting conditions using a video, but in this case the object is rotated around its axis. As BRDF model they chose the isotropic microfacet model.

[10] designed a complex coaxial optical scanner capable of synchronously acquiring shape and spatially varying reflectance. Their device consists of a pair of assemblies each containing a coaxial camera and a light source. The model used is Cook-Torrance. [11] created a minidome with the goal of digitizing cuneiform tablets. Their dome with a radius of 50cm consists of a single camera placed on top and 256 white power LEDs positioned on the knots and edge centers. The entire digitization process is programmable and fully automatic. One special work with the usage of a dome is from [12], currently producing state-of-art results. Their dome consists of 151 DSLR cameras taking HDR sequences and one LED-Projector mounted on a tripod placed at five to eight different positions, projecting 38 different patterns. Instead of calculating BRDF functions, they create a *Bidirectional Texture Function* for the mesh. All of the related works with devices present as limitation the size and position of objects that can be used and its practicality in field application.

III. Technique overview

The goal of our algorithm is to approximate the BRDF of an object given its mesh and photos from a fixed viewpoint but with different light positions. In order for our method to work we expect the mesh to be aligned with the real world object in the photo. We project the mesh's vertices onto each image and read the pixel color information. All calculations are done independently per vertex in the fragment shader, where we fit a Phong model for every vertex:

$$f(L,E) = k_d(\vec{N}\cdot\vec{L})C_{rgb} + k_s(\vec{R}\cdot\vec{E})^{\alpha}S_{rgb}$$

The unknown parameters to be computed for each vertex are: the diffuse coefficient k_d, the specular coefficient k_s, the shininess α and the diffuse color C_{rgb}.

We may acquire all the photos interactively in our acquisition procedure or work with a prepared set. If we acquire the photos interactively, we start our algorithm with the calibration procedure. This includes locating the sphere, retrieving the color chart's values, and computing the mesh-image alignment. After the calibration, we start the BRDF approximation *per se*. We compute an initial diffuse color for each vertex. Following, there's the computation of an approximate specular coefficient. Finally, we perform an optimization on the BRDF's parameters.

We use the Hough transform to automatically detect the sphere position, and the highlight center projected in the sphere to calculate the light direction. In order to simplify the solution we assume that the distances between both camera and light are large with respect to the object, then we can approximate our view direction $\vec{V}$ as the z axis in the world space. To calculate the light direction $\vec{L}$, we need only to reflect the view vector with respect to the normal at the highlight position.

The Color Checker is a chart consisting of color samples designed to obtain a measure of how your photos deviate from the true colors. The model used in our system has 24 colors. The color chart detection consists merely in detecting the squares with respective colors in the image. Then, we apply a linear regression for each color channel.

The effectiveness of our algorithm relies heavily on the alignment between the mesh and the photo. We perform the Mutual Correspondences algorithm [13]. The goal is to find the position, orientation and focal length of the camera. The algorithm tries to optimize these variables by simultaneously maximizing the mutual information and minimizing the error of manually-set correspondences.

An important aspect is to use as few data as possible. Our system is able to provide a new light direction as suggestion to cover the remaining non covered vertices. We generate a new light direction for each vertex. Instead of testing all new light directions, we employ a spatial bin division scheme, where we allocate each light into the correct bin and select the best one. This strategy is also used to generate photos for vertices which do not have a good specularity coverage.

IV. Diffuse color

As a first approximation of the diffuse color we assume our object is lambertian, but only reflects half of the received radiance. It is worth mentioning that for each photo we render our mesh as seen by the light so we can mark which vertices are illuminated. We also do not considerate vertices that are in highlight or in grazing angles.
The color of a vertex v is the weighted average of all corresponding pixel colors C_i, where the weights are simply the product $\vec{N} \cdot \vec{L_i}$, and $\vec{L_i}$ is the light direction of photo i . However, the pixel color does not represent the true object's reflectance value, it is affected by the light interaction, thus, we need to take this factor in consideration:

$$C_{rgb} = \frac{f(L,E)}{k_d(\vec{N}\cdot\vec{L})}$$

V. Specular coefficient

After capturing all photos and calculating a basic diffuse color, we go through the whole dataset estimating specular coefficients for all the vertices that have been *specularly covered*. For every photo we subtract the diffuse color from the pixel P_i and take out the specular value:

$$k_s^i(\vec{R_i}\cdot\vec{E})^\alpha S_{rgb} = P_i - k_d(\vec{N}\cdot\vec{L})C_{rgb}$$

The superscript i in k_s^i means that this is the specular coefficient k_s from the photo i. Let $(\vec{R_i}\cdot\vec{E})^\alpha S_{rgb}$ be $\vec{S_i}$, and S_i^+ its Moore-Penrose pseudoinverse. We can now calculate k_s:

$$k_s^i = (P_i - k_d(\vec{N}\cdot\vec{L})C_{rgb})S_i^+$$

The final k_s is simply the average of all k_s^i obtained from each photo.

VI. Optimization cycle

The next phase of our system is to optimize the parameters, but instead of improving all simultaneously, we perform three optimization steps: first the coefficients, followed by the shininess parameter, and finally the base color. As we improve our color estimation, the other parameters need to be reestimated, hence we begin a cycle repeating these steps. Our optimization scheme seeks to reduce the squared error between the pixel and the resulting color of the vertices. Note that the optimization is global only on a per vertex basis, we do not take into account neighboring vertices.
In order to simplify our future notation, we will designate the resultant color from the Phong model of the vertex i as C_i, $d_i = (\vec{N}\cdot\vec{L})$ and $s_i = (\vec{R}\cdot\vec{E})$.
What remains is to calculate the Jacobian for every set of parameters. We show how to calculate the derivative in each following subsections.

A. Diffuse and Specular Coefficients

We need the derivatives $\frac{\partial f}{\partial k_d}$ and $\frac{\partial f}{\partial k_s}$ which are easily calculated:

$$\frac{\partial f}{\partial k_d} = -2(P_i - C_i)(d_i C_{rgb})$$
$$\frac{\partial f}{\partial k_s} = -2(P_i - C_i)(s_i^\alpha S_{rgb})$$

B. Shininess

Up to now we just used an initial empirical value of $\alpha = 10$. Although not optimal, the chosen value provided satisfactory results as a starting point. The minimum point of f with respect to α is the critical point when $\frac{df}{d\alpha} = 0$, which is: $\frac{df}{d\alpha} = -2(P_i - k_d d_i C_{rgb})(k_s S_{rgb} ln(s_i) s_i^\alpha) + 2((k_s S_{rgb})^2 (s_i^\alpha)^2 ln(s_i)$
Finding the root of this equation is not a simple task since we have an exponential equation of α. To overcome this we use the classic Newton's method to find the zero of an equation. Thus we need the second derivative $\frac{d^2 f}{d\alpha^2}$:
$\frac{d^2 f}{d\alpha^2} = -2(P_i - k_d d_i C_{rgb})(k_s S_{rgb} s_i^\alpha ln^2(s_i)) + 4(k_s S_{rgb} s_i^\alpha ln(s_i))^2$
As this is a global optimization, our optimization actually finds the root of the function which is the sum of the derivatives for each photo.

C. Color

The final optimization step is to improve our basic diffuse color. For this matter, we need the partial derivatives $\frac{\partial f}{\partial r}$, $\frac{\partial f}{\partial g}$ and $\frac{\partial f}{\partial b}$.

$$\frac{\partial f}{\partial r} = -2(Pr_i - Cr_i)(k_d d_i)$$
$$\frac{\partial f}{\partial g} = -2(Pg_i - Cg_i)(k_d d_i)$$
$$\frac{\partial f}{\partial b} = -2(Pb_i - Cb_i)(k_d d_i)$$

VII. Results and Discussions

We performed a number of tests to analyze the results of our BRDF approximation algorithm. We test it with different materials and examine its behavior.

Nana: The first object analyzed is the Nana doll, composed of a very specular head and a diffuse body. Besides, groups of similar colors, for example its green belly, tend to present a high color variation, *i.e.*, many close points present different green tonalities, which calls for a good SVBRDF approximation. Results show that our algorithm is able to faithfully capture highlight areas and perform an efficient diffuse color extraction (Figure 2).


(a) Photograph of Nana


(b) Rendering of Nana

Fig. 2. Results for the Nana

However, there are some drawbacks mainly regarding the diffuse and specular coefficients (Figure 3). The error of specular photos are guiding the optimization process to set the specular coefficient too high while setting the diffuse coefficient too low.

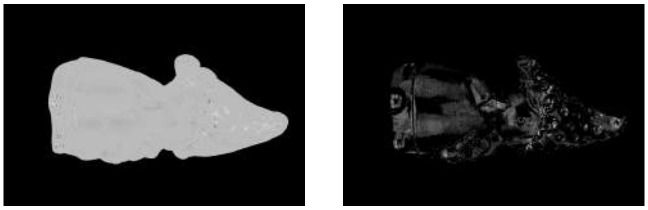

(a) Diffuse coefficient (b) Specular coefficient

Fig. 3. Nana's coefficients. Observe the dark dots present on Nana's head in the diffuse coefficients photo and how these match the white dots in the specular coefficients photo.

Buddha: The Buddha is a small statue composed of a highly specular golden part representing the Buddha's body and a less shiny, but still moderately specular surface composing its robe. The robe is painted in a dark red color. There is also the Budda's hair composed of a mostly dark diffuse surface with some golden spots in the middle. Due to the golden material, we already expected to face some problems with the Buddha. We show results in Figure 4.

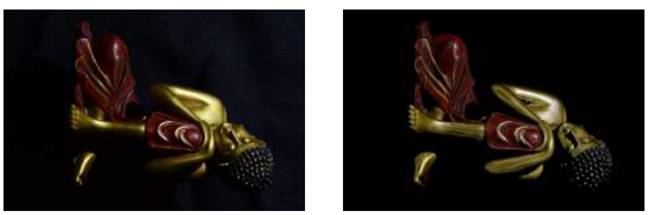

(a) Photograph of Buddha (b) Rendering of Buddha

Fig. 4. Results for the Buddha. While we can adequately approximate the robe's BRDF, and the highlight shape to some extent, the problem experienced with the diffuse holes prevents a better reflectance approximation.

For our datasets, the total time of execution was: Nana 42.7443s and Buddha 44.9737s. These results include all the steps of the BRDF approximation without the calibration procedure.

VIII. Conclusion and Future Works

In this work we have presented a system for simple, fast and faithful BRDF approximation using only photos with varying light directions. There are some drawbacks regarding shiny objects, but we believe a few corrections in our optimization algorithm could solve the problem. We also provide a tool for CH professionals who need to perform digitization of historical artifacts in possible distant locations. To the best of our knowledge, there is no tool in the research community that can provide the same immediate feedback.

Nevertheless, there is always room for improvement and experimentation. The main drawback in performing the tests is the maximum number of photos used in our GPU. Also, a recurrent artifact in our results was some very low diffuse coefficients after the optimization approach biased towards the specular photos. Although Phong's model provides convincing results, there may be a need to use another reflectance model for objects with complicated materials. One idea would be to make sure in situ we capture all the relevant data for a good BRDF approximation, and perform a more time-consuming optimization approach offline.

References

[1] T. Ritschel, C. Dachsbacher, T. Grosch, and J. Kautz, "The state of the art in interactive global illumination," *Comput. Graph. Forum*, vol. 31, no. 1, pp. 160–188, Feb. 2012. [Online]. Available: http://dx.doi.org/10.1111/j.1467-8659.2012.02093.x

[2] T. Weyrich, J. Lawrence, H. P. Lensch, S. Rusinkiewicz, and T. Zickler, "Principles of appearance acquisition and representation," *Foundations and Trends in Computer Graphics and Vision*, vol. 4, no. 2, pp. 75–191, Oct. 2009.

[3] H. Lensch, J. Kautz, M. Goesele, W. Heidrich, and H.-P. Seidel, "Image-based reconstruction of spatially varying materials," in *Rendering Techniques 2001*, ser. Eurographics, S. Gortler and K. Myszkowski, Eds. Springer Vienna, 2001, pp. 103–114. [Online]. Available: http://dx.doi.org/10.1007/978-3-7091-6242-2_10

[4] H. P. A. Lensch, J. Kautz, M. Goesele, W. Heidrich, and H.-P. Seidel, "Image-based reconstruction of spatial appearance and geometric detail," *ACM Trans. Graph.*, vol. 22, no. 2, pp. 234–257, Apr. 2003. [Online]. Available: http://doi.acm.org/10.1145/636886.636891

[5] T. Malzbender, D. Gelb, and H. Wolters, "Polynomial texture maps," in *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*, ser. SIGGRAPH '01. New York, NY, USA: ACM, 2001, pp. 519–528. [Online]. Available: http://doi.acm.org/10.1145/383259.383320

[6] D. Goldman, B. Curless, A. Hertzmann, and S. Seitz, "Shape and spatially-varying brdfs from photometric stereo," in *Computer Vision, 2005. ICCV 2005. Tenth IEEE International Conference on*, vol. 1, Oct 2005, pp. 341–348 Vol. 1.

[7] N. G. Alldrin, T. Zickler, and D. J. Kriegman, "Photometric stereo with non-parametric and spatially-varying reflectance." in *CVPR*. IEEE Computer Society, 2008. [Online]. Available: http://dblp.uni-trier.de/db/conf/cvpr/cvpr2008.html#AlldrinZK08

[8] G. Palma, M. Callieri, M. Dellepiane, and R. Scopigno, "A statistical method for svbrdf approximation from video sequences in general lighting conditions," *Comp. Graph. Forum*, vol. 31, no. 4, pp. 1491–1500, Jun. 2012. [Online]. Available: http://dx.doi.org/10.1111/j.1467-8659.2012.03145.x

[9] Y. Dong, G. Chen, P. Peers, J. Zhang, and X. Tong, "Appearance-from-motion: Recovering spatially varying surface reflectance under unknown lighting," *ACM Trans. Graph.*, vol. 33, no. 6, pp. 193:1–193:12, Nov. 2014. [Online]. Available: http://doi.acm.org/10.1145/2661229.2661283

[10] M. Holroyd, J. Lawrence, and T. Zickler, "A coaxial optical scanner for synchronous acquisition of 3d geometry and surface reflectance," in *ACM SIGGRAPH 2010 Papers*, ser. SIGGRAPH '10. New York, NY, USA: ACM, 2010, pp. 99:1–99:12. [Online]. Available: http://doi.acm.org/10.1145/1833349.1778836

[11] G. Willems, F. Verbiest, W. Moreau, H. Hameeuw, K. Van Lerberghe, and L. Van Gool, "Easy and cost-effective cuneiform digitizing," in *The 6th International Symposium on Virtual Reality, Archaeology and Cultural Heritage (VAST 2005)*, 2005, pp. 73–80.

[12] C. Schwartz, M. Weinmann, R. Ruiters, and R. Klein, "Integrated high-quality acquisition of geometry and appearance for cultural heritage," in *The 12th International Symposium on Virtual Reality, Archeology and Cultural Heritage VAST 2011*, Eurographics Association. Eurographics Association, Oct. 2011, pp. 25–32. [Online]. Available: http://diglib.eg.org/EG/DL/WS/VAST/VAST11/025-032.pdf

[13] M. Sottile, M. Dellepiane, P. Cignoni, and R. Scopigno, "Mutual correspondences: An hybrid method for image-to-geometry registration." in *Eurographics Italian Chapter Conference*, E. Puppo, A. Brogni, and L. D. Floriani, Eds. Eurographics, 2010, pp. 81–88. [Online]. Available: http://dblp.uni-trier.de/db/conf/egItaly/egItaly2010.html#SottileDCS10