Environment Map Based Lighting for Reflectance Transformation Images

Michael Ludwig*, Gary Meyer†
University of Minnesota, Department of Computer Science
*mludwig@cs.umn.edu, †meyer@cs.umn.edu

***Abstract*—Environment map based lighting has proven to be an effective technique in simulations and real time rendering for adding realism and conveying complex illumination to the viewer. We present a novel algorithm for using environment maps to light and render polynomial texture maps. The technique is demonstrated with reflectance transformation images stored using both polynomial texture maps and hemispherical harmonics. The limitations of employing hemispherical harmonics in this context are discussed.**

***Index Terms*—environment mapping, rti, ptm**

I. Introduction

Reflectance transformation imaging (RTI), first introduced as polynomial texture maps [1], has been applied in the field of cultural heritage to capture the reflectance properties of the complex materials used in paintings, sculptures, and other artifacts, and to demonstrate how they behave under different lighting conditions. Computed from photographs, RTI data sets are highly faithful to the spatially-varying details of the object and essentially interpolate between the photographs to simulate appearance changes under new lighting directions. Existing research has focused on either improving the model constructed from the photographs–expanding the classes of materials that can be represented–and on manipulating the presentation of the RTI data for informational and visualization purposes. The types of visualizations developed have been driven by cultural heritage researchers and their needs to better analyze and understand the artifacts that were captured. In this paper we present an alternative approach to rendering RTIs: relighting the data set with entirely new environments. This can be used by patrons, visitors, and researchers to view how an object would appear in a variety of conditions such as time-of-day, location, or the environment of the original artist. It can also be a useful tool for lighting design when determining how best to present an artifact in a gallery.

There are a number of RTI techniques that allow for viewing an object lit from a novel direction. However, they achieve this in methods incompatible with one another, necessitating unique environment lighting solutions for each RTI method. We present environment lighting of RTI data sets computed using polynomial texture maps [1] and hemispheric harmonics [2], perhaps the two most widely used methods. The rendering algorithms are described for each technique and then the results are analyzed. In particular, hemispheric harmonics, which tend to be heralded as the ideal solution, perform poorly in this application.

Section II provides an overview of both RTI and environment lighting. Next, Section III describes how to use environment maps to render RTI modeled as hemispheric harmonics and polynomial texture maps. Lastly, Section IV demonstrates the two environment lighting methods on a number of RTI data sets and concludes with a discussion and comparison between the RTI models.

II. Background

Numerous methods have been developed to model the reflectance of objects. Bidirectional reflectance distribution functions (BRDF) [3] capture the microscopic behavior of a material. Expensive camera and lighting gantries have been constructed and carefully calibrated, allowing both geometry and BRDF acquisition from real-world objects [4], [5]. Although expensive, these approaches produce models and materials that can be easily rendered with conventional algorithms for environment re-lighting. Bidirectional texture functions [6]–applied to cultural heritage by Schwartz et al. [7]–also afford re-lighting with new environments. Unfortunately, in addition to requiring hundreds or thousands of photographs and expensive acquisition equipment, these approaches do not map well to modern graphics hardware.

Reflectance transformation imaging takes a simpler approach, that makes it significantly more affordable for cultural institutions. A single camera is required and it is used to photograph an object from a fixed location. Multiple captures with a single small area light are made to record how the artifact behaves from different lighting directions. Polynomial texture maps (PTM) [1] popularized this capture methodology and fit the resulting photographs with per-pixel biquadratic functions parameterized by the light direction. Hemispheric harmonics [2] can be used to fit the photographic data to modified spherical harmonic basis functions. Both of these approaches are effective at modeling the diffuse properties of an object, although hemispheric harmonics are not as biased when some photographs include specular highlights. Newer techniques have tackled issues such as identifying shadows [8], improved surface normal reconstruction [9], or including specular surface properties [10] in the model. Because these last approaches have only recently appeared, we have not developed environment lighting methods for their models.

Rendering an RTI data set is essentially interpolating the captured photographs to reconstruct the object as if it were lit from the user-selected direction, although it is limited to the type and characteristics of the light originally used. Cultural heritage research has benefited from alternative rendering methods for RTIs. For example, Palma et al. [11] presents a number of non-photorealistic algorithms for enhancing the surface appearance of an object. Environment mapping the RTI data set can be considered as another visualization method: rendering the artifact as if it were in a novel environment, as done by Debevec [12]. An environment map captures all incoming light radiance from every direction to a given point in space. Thus, it is an efficient means of representing the complex light sources present in the real world. Environment maps are often represented as spherical harmonics [13], [14] making them convenient to apply to RTIs modeled as hemispheric harmonics. The alternative to the spherical harmonics representation of an environment is to sample it into discrete light sources that capture the most predominant light sources [15]. Sampling is a necessary part of our algorithm for PTM environment rendering.

III. Re-lighting RTIs

Existing literature on environment re-lighting assumes the object is either a perfect mirror or that its material is represented as a BRDF. A BRDF is a four dimensional function that returns the fraction of incoming light that's reflected in an outgoing direction. To determine the appearance of a BRDF for a given outgoing direction, the rendering equation [16] must be evaluated:

$$L(\omega_o) = \int_{\Omega} f(\omega_o, \omega_i) L(\omega_i) cos\theta d\omega_i \tag{1}$$

In Equation 1, ω represents a solid angle, f is the BRDF, L is the radiance along a particular solid angle–either incoming or outgoing, and θ is the angle of ω_i with respect to the surface normal.

A naive interpretation of an RTI is that it's a BRDF with a fixed outgoing direction. However, this is incorrect as the RTI models the final appearance or rendering of a material (i.e. $L(\omega_o)$), including any shadows present on the object, the exposure and other camera settings from the capture, and most importantly the lighting used. Physical light sources have a non-zero solid angle, so treating the RTI as a BRDF effectively widens the BRDF lobes. This can be likened to "faxing a fax", but in this case the errors are manifested as an overall increase in brightness and flattening of appearance as if the light used in the scene was made larger. To produce compelling renderings, this must be taken into account by the algorithm. One such approach is described for PTMs in Section III-B.

A. Hemispheric harmonic RTIs

Gautron et al. [2] defines the hemispheric harmonic (HSH) basis functions, $H_l^m(\theta, \phi)$, which are based on the real spherical harmonic (SH) functions , $S_l^m(\theta, \phi)$. In their initial work, HSH was used to represent BRDFs and they outlined an approach to render reflectance functions within an environment using their basis:

1) Compute the SH coefficients for the environment.
2) Rotate the SH coefficients into the coordinate frame of the object being rendered.
3) Convert the SH coefficients into HSH coefficients using a change-of-basis matrix defined in [2].
4) Take the dot product between the HSH coefficients of the material and those of the environment. This is the final result.

Hemispheric harmonics can also represent RTI data. A least squares fitting of HSH coefficients can be found for each pixel in the RTI with respect to the source pixel's values that were photographed. Each pixel within the RTI has its own coefficients, so to reduce memory pressure have been limited to order-3 hemispheric harmonics (9 coefficients), which is sufficient to represent the diffuse characteristics of an object.

It is straightforward to apply the environment rendering algorithm highlighted above to HSH coefficients representing RTI data The transformation from the environment's coordinate frame to that of the RTI is the same for every pixel of the RTI. The transformed environment coefficients, now in the hemispheric basis, can be sent to the GPU and combined per-pixel with the coefficients in the RTI for real-time rendering. Conveniently, the rotation matrices and change-of-basis matrices for order-3 harmonics are provided directly in [14] and [2] respectively.

There are two critical problems with rendering RTIs described by hemispheric harmonics. The first is caused by the real size of the lights used during RTI capture, as described in the beginning of Section III. The dot product between two sets of HSH coefficients is an approximation of the integral of the product of the reflectance function and environment. Thus, the issue manifests itself when HSH RTI's are rendered with a new environment map, which is demonstrated in Section IV. The second is from the limited order of HSH basis functions used; although adequate for describing a material, higher orders are required to properly describe many environments. This situation is illustrated in the bottom of Fig. 1: sharp boundaries between lights are expanded and softened. When they are present, high-intensity lights such as the sun can have significant impacts on the overall approximation. The following section describes a new approach to environment rendering RTIs described by PTMs that avoids these problems.

B. Polynomial texture maps

Polynomial texture maps (PTMs) are described by biquadratic functions, parameterized by the projected coordinates of the incoming light direction:

$$L(l_u, l_v) = a_0 l_u^2 + a_1 l_v^2 + a_2 l_u l_v + a_3 l_u + a_4 l_v + a_5 \tag{2}$$

Every pixel in the PTM has its own set of coefficients $\{a_i\}$ for $i = 0, \ldots, 5$. Evaluating this polynomial for a particular lighting direction (l_u, l_v) approximates what the material would reflect if the light used during photography came from there.

To efficiently evaluate environment lighting for every single pixel within a PTM, the polynomial must be restructured. Consider the environment lighting equation for a particular

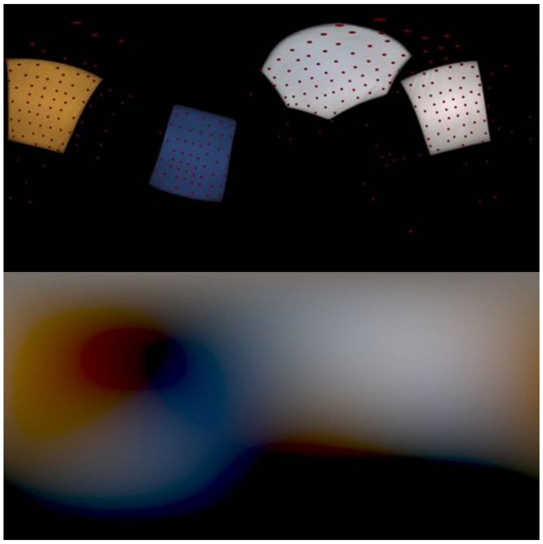

Fig. 1. Comparison between a studio environment with light boxes (top), and its reconstruction from order-3 spherical harmonics (bottom). Red dots mark the sampled light positions described in Section III-B.

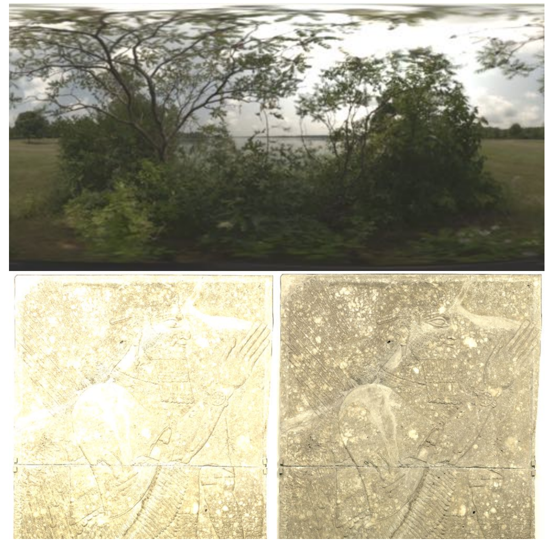

Fig. 2. Environment (top), HSH (left), PTM (right): Winged Genius Assyrian bas relief in an overcast, forested lakeside. The HSH is clearly overexposed, as if it were under brighter sun than the environment suggests. Section III-A predicted such a brightening.

pixel, where I is the final intensity for the pixel, and $E(l_u, l_v)$ is the radiance of the environment along that direction:

$$\begin{aligned} I &= \sum_{l_u,l_v} E(l_u,l_v)L(l_u,l_v) \\ &= \sum_{l_u,l_v} a_0E(l_u,l_v)l_u^2 + a_1E(l_u,l_v)l_v^2 + a_2E(l_u,l_v)l_ul_v + \\ &\quad a_3E(l_u,l_v)l_u + a_4E(l_u,l_v)l_v + a_5E(l_u,l_v) \\ &= a_0\sum_{l_u,l_v} E(l_u,l_v)l_u^2 + a_1\sum_{l_u,l_v} E(l_u,l_v)l_v^2 + \\ &\quad a_2\sum_{l_u,l_v} E(l_u,l_v)l_ul_v + a_3\sum_{l_u,l_v} E(l_u,l_v)l_u + \\ &\quad a_4\sum_{l_u,l_v} E(l_u,l_v)l_v + a_5\sum_{l_u,l_v} E(l_u,l_v) \end{aligned} \tag{3}$$

The six summations in Equation 3, before being multiplied by the a_i coefficients, are independent of the coefficients. This means they can be evaluated once for an entire PTM and environment and applied to each pixel's coefficients. To support real-time rotations of the environment, it is useful to decompose an environment map into discrete lights. In practice we use structured importance sampling [15] or a simpler binary partitioning scheme described by Reinhard et al. [17] to describe an environment as several hundred to thousand directional lights (see Fig. 1).

After the environment map is decomposed into a set of point lights, the lights are filtered to maintain a minimum angle between them. A greedy approach of repeatedly collapsing lights that failed to meet the threshold has worked well. The minimum angle can be selected by photographing the light source as seen by the artifact and measuring its angular width, or by estimating its width.

IV. Results

In this section, we demonstrate environment lighting applied to RTIs modeled as both HSHs and PTMs. Fig. 2, 3a, and 3b demonstrate the HSH technique previously described by Gautner and our novel PTM environment map lighting technique. Each figure displays the environment along the top and then a comparison between the HSH (left) and PTM (right) images. A case-by-case analysis of the two techniques is given beneath each figure.

The poor performance of hemispheric harmonic RTIs when using environment maps containing high frequencies is a heretofore unrealized limitation of that model for use in RTI. In general, it's been considered the work horse for robustly capturing surface normals and diffuse coloration of many artifacts, with less error than PTMs. For non-scientific purposes, where environment lighting can provide an interesting online or interactive exhibit, PTMs may now be a compelling alternative. The newer approaches described by Drew [8] and Zhang [9] rely on higher order polynomials than the biquadratic used in PTMs. However, the equation restructuring shown in Equation 3 applies to any polynomial of similar form. The methods described here should extend to this work as well.

In this paper we have described in detail how to re-light RTI data sets in environment maps, where the RTI data is represented as either hemispheric harmonics or polynomial texture maps. We have shown where HSH models break down in this new application and provide evidence for considering PTMs again. It is our intention to make this rendering software available for cultural heritage purposes so that museums and curators may take advantage of these rendering techniques.

(a) Environment (top), HSH (left), PTM (right): The whitening of the gold leaf in the HSH image is incorrect: it should still be golden. Also the HSH appears to be more significantly lit from overhead, indicated by the dark shadow along the top of the frame, which is consistent with the bright track lights dominating the SH environment representation.

(b) Environment (top), HSH (left), PTM (right): The PTM image shows both blue and orange lights contributing to the shading of the artwork. The HSH has much less blue and appears dominated by the orange light, and is significantly brighter overall. This environment was shown in spherical harmonics in Fig. 1.

Fig. 3. Tempera on panel with gold leaf of Saint Sirus by Vincenzo Foppa. The left environment was captured from a museum gallery with skylight. The right environment is an extreme studio environment with colored light boxes.

ACKNOWLEDGMENT

The authors would like to thank the Minneapolis Institute of Art and Charles Walbridge for providing the photographic data sets. Additionally, we are thank of Seth Berrier for capturing the environment maps used in the examples. Funded by the UM OVPR Grant in Aid Program.

REFERENCES

[1] T. Malzbender, D. Gelb, and H. Wolters, "Polynomial texture maps," in *Proceedings of SIGGRAPH '01*. New York, NY, USA: ACM, 2001, pp. 519–528.

[2] P. Gautron, J. Krivanek, S. Pattanaik, and K. Bouatouch, "A novel hemispherical basis for accurate and efficient rendering," in *Proceedings of the Fifteenth Eurographics Conference on Rendering Techniques*, 2004.

[3] G. Ward, "Measuring and modeling anisotropic reflection," *ACM SIGGRAPH Computer Graphics*, vol. 26, no. 2, pp. 265–272, 1992.

[4] A. Gardner, C. Tchou, T. Hawkins, and P. Debevec, "Linear light source reflectometry," *ACM Transactions on Graphics*, vol. 22, no. 3, pp. 749–758, 2003.

[5] B. Tunwattanapong and G. Fyffe, "Acquiring reflectance and shape from continuous spherical harmonic illumination," *ACM Transactions on Graphics*, vol. 32, no. 4, pp. 109:1–109:12, 2013.

[6] K. J. Dana, B. van Ginneken, S. K. Nayar, and J. J. Koenderink, "Reflectance and Texture of Real-World Surfaces," *ACM Trans. Graph.*, vol. 18, no. 1, pp. 1–34, Jan. 1999.

[7] 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*, 2011.

[8] M. S. Drew, Y. Hel-Or, T. Malzbender, and N. Hajari, "Robust estimation of surface properties and interpolation of shadow/specularity components," *Image Vision Comput.*, vol. 30, no. 4-5, pp. 317–331, 2012.

[9] M. Zhang, "Robust surface normal estimation via greedy sparse regression," Master's thesis, Simon Fraser University, 2014.

[10] M. Zhang, M. S. Drew, and M. S Drew, "Efficient robust image interpolation and surface properties using polynomial texture mapping," *EURASIP Journal on Image and Video Processing*, vol. 2014, p. 25, 2014.

[11] G. Palma, M. Corsini, P. Cignoni, R. Scopigno, and M. Mudge, "Dynamic shading enhancement for reflectance transformation imaging," *Journal on Computing and Cultural Heritage*, vol. 3, no. 2, pp. 1–20, Sep. 2010.

[12] P. Debevec, "Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography," in *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques*, 1998, pp. 189–198.

[13] R. Ramamoorthi and P. Hanrahan, "An efficient representation for irradiance environment maps," *Siggraph*, vol. 64, pp. 497–500, 2001.

[14] R. Green, "Spherical harmonic lighting: The gritty details," *Archives of the Game Developers Conference*, 2003.

[15] S. Agarwal and R. Ramamoorthi, "Structured importance sampling of environment maps," *ACM Transactions on Graphics*, vol. 22, no. 3, pp. 605–612, 2003.

[16] J. T. Kajiya, "The rendering equation," *SIGGRAPH Comput. Graph.*, vol. 20, no. 4, pp. 143–150, Aug. 1986.

[17] E. Reinhard, G. Ward, S. Pattanaik, P. Debevec, W. Heidrich, and K. Myszkowski, *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting*, 2nd ed. Burlington, MA: Morgan Kaufmann, 2010.