Linearly Transformed Spherical Harmonics

Jan Allmenröder
contact@jallmenroeder.de
Karlsruhe Institute of Technology
Karlsruhe, Germany

Christoph Peters
christoph.peters@kit.edu
Karlsruhe Institute of Technology
Karlsruhe, Germany

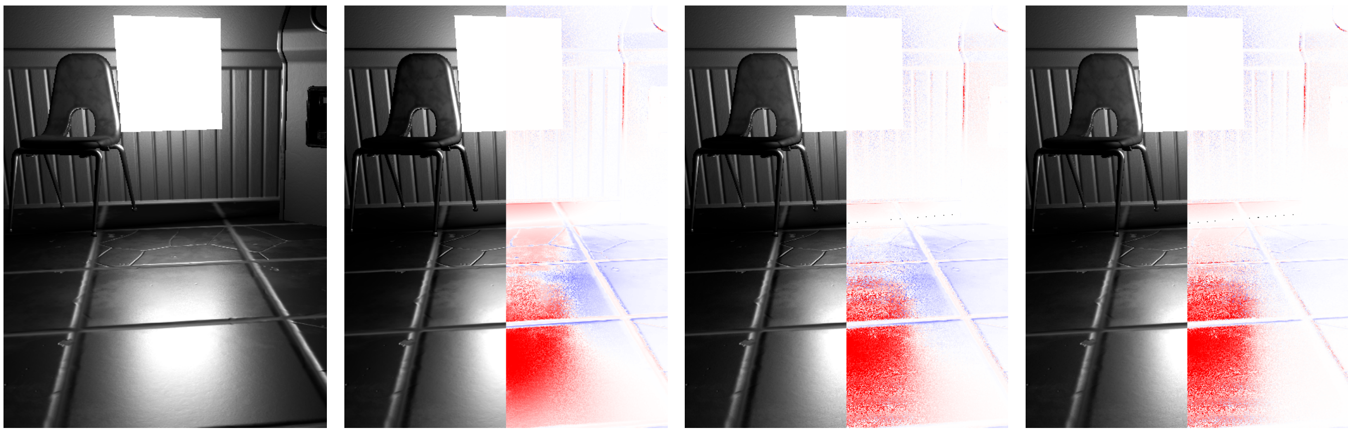

Ground truth | LTC [Heitz et al., 2016] | Our LTSH, $N = 2$ | Our LTSH, $N = 4$

Figure 1: The arcade scene rendered with different specular shading techniques. Red/blue means too bright/dark, respectively.

ABSTRACT

Linearly transformed cosines approximate specular BRDFs in a way that enables efficient shading for Lambertian polygonal area lights. The approximation quality is mostly good but fails to reproduce tear shapes at grazing angle. We swap out the cosine with more general spherical harmonics (SH) expansions. Thus, our linearly transformed SH offers better fits at higher cost. Previous work allows us to integrate these SH expansions over polygons such that shading still works in closed form.

CCS CONCEPTS

• **Computing methodologies → Reflectance modeling**.

KEYWORDS

linearly transformed cosines, LTC, spherical harmonics, LTSH, real-time rendering, polygonal area lights, specular shading

ACM Reference Format:
Jan Allmenröder and Christoph Peters. 2021. Linearly Transformed Spherical Harmonics. In *Proceedings of 25th ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games (i3D 2021 posters).* ACM, New York, NY, USA, 2 pages. https://doi.org/none

i3D 2021 posters, 20-22 April 2021, online

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM…$15.00
https://doi.org/none

1 INTRODUCTION

Realistic real-time rendering requires plausible light sources. Lambertian polygonal area lights are versatile and approximate many real light sources well. However, shading is non-trivial. According to the reflection equation, the radiance reflected by a shading point at $\mathbf{x} \in \mathbb{R}^3$ into direction $\omega_o \in \mathbb{S}^2$ is

$$L(\mathbf{x}, \omega_o) := L_i \int_{\mathbb{P}} f_r(\omega_i, \mathbf{x}, \omega_o) \max(0, \cos\theta_i)\, \mathrm{d}\omega_i, \quad (1)$$

where $\mathbb{P}$ is the solid angle of the polygon, L_i is its emitted radiance, f_r is the bidirectional reflectance distribution function (BRDF) and θ_i is the angle between ω_i and the shading normal.

Linearly transformed spherical distributions (LTSD) [Heitz et al., 2016] approximate the cosine-weighted BRDF in a way that enables fast evaluation of this integral. An LTSD takes the form

$$D(\omega) := D_{\mathrm{orig}}\left(\frac{M^{-1}\omega}{\|M^{-1}\omega\|}\right)\frac{|M^{-1}|}{\|M^{-1}\omega\|^3},$$

where $M \in \mathbb{R}^{3\times 3}$ is a linear transform chosen for an optimal fit and D_{orig} is the density of the original distribution. By construction, integration over an LTSD D is the same as integration over the original distribution D_{orig} in the original domain $\mathbb{P}_{\mathrm{orig}}$:

$$\int_{\mathbb{P}} D(\omega)\, \mathrm{d}\omega = \int_{\mathbb{P}_{\mathrm{orig}}} D_{\mathrm{orig}}(\omega_{\mathrm{orig}})\, \mathrm{d}\omega_{\mathrm{orig}}, \quad (2)$$

$$\mathbb{P}_{\mathrm{orig}} := \left\{ \frac{M^{-1}\omega}{\|M^{-1}\omega\|} \mid \omega \in \mathbb{P} \right\}.$$

Since $\mathbb{P}$ is a spherical polygon, $\mathbb{P}_{\mathrm{orig}}$ is obtained by simply transforming the direction from the shading point $\mathbf{x}$ to each of its vertices with M^{-1}. If the choice for D_{orig} is simple enough, the integral in

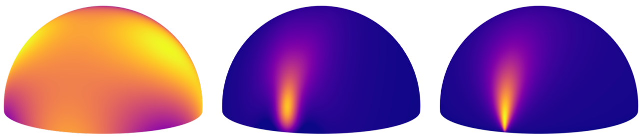

(a) $D_{\text{orig}}(\omega_{\text{orig}})$ **(b)** $D(\omega)$ **(c) Frostbite GGX BRDF**

Figure 2: An SH expansion (a) and its linear transformation (b) approximating a cosine-weighted BRDF (c).

Equation 2 is easy to solve. Linearly transformed cosines (LTC) use a cosine distribution $D_{\text{orig}}(\omega_{\text{orig}}) = \frac{1}{\pi}\max(0, \omega_{\text{orig},z})$ [Heitz et al., 2016]. The authors find that the resulting class of functions provides descent fits for widely used cosine-weighted specular BRDFs. Thus, they precompute a lookup table of transforms M^{-1}, parameterized by the roughness $\alpha \in [0, 1]$ of the BRDF and the inclination of the outgoing light direction ω_o.

Recent work investigates the integration of spherical harmonics (SH) expansions over polygons [Belcour et al., 2018, Wang and Ramamoorthi, 2018]. These works enable us to use SH expansions as original distribution $D_{\text{orig}}(\omega_{\text{orig}})$. Through this larger class of functions, we improve the quality of fits compared to LTCs at an increased cost. We focus on specular BRDFs since LTCs work well for diffuse BRDFs.

2 LINEARLY TRANSFORMED SPHERICAL HARMONICS (LTSH) EXPANSIONS

We use an SH basis consisting of bands $l \in \{0, \ldots, N\}$. Each band holds $2l + 1$ basis functions $Y_{l,m}(\omega_{\text{orig}})$ indexed by $m \in \{-l, \ldots, l\}$ [Wang and Ramamoorthi, 2018]. Thus, our original distribution takes the form

$$D_{\text{orig}}(\omega_{\text{orig}}) = \sum_{l=0}^{N} \sum_{m=-l}^{l} c_{l,m} Y_{l,m}(\omega_{\text{orig}}).$$

While LTCs use a completely fixed distribution, our LTSH have additional degrees of freedom $c_{l,m} \in \mathbb{R}$. Our fit uses them to improve the approximation quality.

For a fixed transformation M, the LTSH density $D(\omega)$ depends linearly on the SH coefficients $c_{l,m}$. Thus, we find the optimal SH coefficients by solving a linear least squares problem. The resulting choice minimizes the L^2-difference between the cosine-weighted BRDF and $D(\omega)$ in the upper hemisphere. Figure 2 shows an approximation of the Frostbite GGX BRDF at grazing angle.

As with LTCs, the search for the transformation M is a non-linear minimization problem and thus harder to solve. In each iteration, our minimizer finds the optimal SH expansion as described above. The residual L^2-error of the least squares solver serves as objective function for the non-linear minimizer. Like the LTC fitting [Heitz et al., 2016], our optimization works in a coordinate frame where the z-axis is aligned with the shading normal and $\omega_{o,y} = 0$. Then the optimal transform should take the form $M = \begin{pmatrix} a & 0 & b \\ 0 & c & 0 \\ d & 0 & 1 \end{pmatrix}$ leaving a four-dimensional search space. We explore it using the `scipy.optimize.curve_fit` optimizer. For LTCs we use Levenberg-Marquardt, for LTSH trust-region reflective. Good initialization of the optimizer is important. We use parameters of neighboring entries in the lookup-table [Heitz et al., 2016]. In the end, the lookup table has the same structure as for LTCs, except that we replace specular albedos by vectors of SH coefficients $c_{l,m}$.

Since our fits are only meaningful in the upper hemisphere, we clip the polygon before shading. Then we transform vertex direction vectors with the appropriate M^{-1} from the table. The approximate shading (without the factor L_i) is given by Equation 2. Thus, we have to integrate the SH expansion with coefficients $c_{l,m}$ from the table over the transformed polygon $\mathbb{P}_{\text{orig}}$. We use the implementation provided by [Wang and Ramamoorthi, 2018]. It is optimized for SH expansions with $N = 8$ and could be further optimized for our setting where $N = 2$ or $N = 4$.

3 RESULTS

Figure 1 shows the performance of LTCs and our LTSH compared to ground truth using the Frostbite GGX BRDF. We achieve small improvements with LTSH of degree 2 and with degree 4 the tails of the specular highlight improve significantly. The continuity of neighboring LTC fits allows for linear interpolation between data in the look-up table. LTSH must use dithering instead, as the fitted parameters are less smooth. This causes more noise compared to LTCs but achieves plausible results in shading. An LTC can be interpreted as LTSH with $N = 1$ and thus, we should always get better fits. Though, the non-linear optimization is challenging and future work may further improve it. Note, that we use a naive fitting implementation that likely performs worse than original LTC.

The improved quality of LTSH comes at the cost of substantially slower shading. The run time of the lighting pass in our deferred shader was 0.44 ms for LTC, 0.95 ms for LTSH with $N = 2$ and 1.78 ms with $N = 4$. The lighting pass without any shading takes 0.14 ms. As with LTCs, the run time cost for LTSH scales linearly with the number of polygon edges but also quadratically in the degree of our SH expansion. The ability to choose SH degrees makes our technique adaptable to different rendering budgets and quality requirements.

Currently, the increased cost is hard to justify through the quality improvement. Besides our approach is prone to slight ringing artifacts. However, improved fitting and more optimized SH integration code, could make LTSH viable as high-quality alternative to LTCs. Code for our fitting procedure and our Falcor renderer is available at http://jallmenroeder.de/LTSH.

ACKNOWLEDGMENTS

The arcade scene was created by Nicholas Hull for NVIDIA. It ships with Falcor, which is the basis of our renderer.

REFERENCES

Laurent Belcour, Guofu Xie, Christophe Hery, Mark Meyer, Wojciech Jarosz, and Derek Nowrouzezahrai. Integrating clipped spherical harmonics expansions. *ACM Trans. Graph.*, 37(2), 2018. ISSN 0730-0301. URL https://doi.org/10.1145/3015459.

Eric Heitz, Jonathan Dupuy, Stephen Hill, and David Neubelt. Real-time polygonal-light shading with linearly transformed cosines. *ACM Trans. Graph.*, 35(4):41:1–41:8, 2016. ISSN 0730-0301. URL http://doi.acm.org/10.1145/2897824.2925895.

Jingwen Wang and Ravi Ramamoorthi. Analytic spherical harmonic coefficients for polygonal area lights. *ACM Trans. Graph.*, 37(4), 2018. ISSN 0730-0301. URL https://doi.org/10.1145/3197517.3201291.