Pacific Graphics 2025
M. Christie, N. Pietroni, and Y.-S. Wang
(Guest Editors)

Volume 44 (2025), Number 7

Supplementary: LTC-IR: Multiview Edge-Aware Inverse Rendering with Linearly Transformed Cosines

Dabeen Park[1†], Junsuh Park[1†], Jooeun Son[2], Seungyoung Lee[2] and Joo-Ho Lee[1]

[1] Sogang University, Seoul, South Korea
{ysbeen00, pjun, jhleecs}@sogang.ac.kr
[2] POSTECH, Pohang, South Korea
{jeson, leesy}@postech.ac.kr

1. Double Axial Moments

Assuming $\mathbf{A}$ is the spherical projection of a polygonal light source $\mathbf{P}$, previous work [Arv95] has provided a method to compute the integral of products of dot products over a subset $\mathbf{A}$ of the spherical domain Ω with respect to two direction vectors $\mathbf{a}$ and $\mathbf{b}$ in closed form. This integral is defined as

$$\tau^{m,n}(\mathbf{A},\mathbf{a},\mathbf{b}) = \int_{\mathbf{A}\subset\Omega} (\mathbf{a}\cdot\mathbf{i})^m (\mathbf{b}\cdot\mathbf{i})^n \mathrm{d}\mathbf{i}.$$

For our purposes, we are interested in the case where $m = 1$ and $n = 1$. Let us assume region $\mathbf{A}$ is a spherical triangle defined by vertices $(\mathbf{v}_0, \mathbf{v}_1, \mathbf{v}_2)$, where each vertex lies on the unit sphere, $\|\mathbf{v}_i\| = 1$. The specified integral we need is:

$$\begin{aligned}\tau^{1,1}(\mathbf{A},\mathbf{a},\mathbf{b}) &= \int_{\mathbf{A}\subset\Omega} (\mathbf{a}\cdot\mathbf{i})^1 (\mathbf{b}\cdot\mathbf{i})^1 \mathrm{d}\mathbf{i} \\ &= \frac{1}{3}((\mathbf{a}\cdot\mathbf{b})\sigma(\mathbf{A}) - (\mathbf{b}\cdot\mathbf{n}') \int_{\partial\mathbf{A}} (\mathbf{a}\cdot\mathbf{i})\mathrm{d}\mathbf{i}),\end{aligned}$$

where $\sigma(\mathbf{A})$ denotes the solid angle of the spherical projection $\mathbf{A}$, $\mathbf{n}'$ is the outward-pointing normal of the base of the spherical triangle $\mathbf{A}$, and $\int_{\partial\mathbf{A}}(\mathbf{a}\cdot\mathbf{i})\mathrm{d}\mathbf{i}$ represents the line integral term over the boundary $\partial\mathbf{A}$ of region $\mathbf{A}$.

The solid angle $\sigma(\mathbf{A})$ of the spherical triangle is given by Girard's theorem:

$$\sigma(\mathbf{A}) = \sum_{i=0}^{2} \left(\frac{(\mathbf{v}_{i+1}\times\mathbf{v}_i)\cdot(\mathbf{v}_{i+2}\times\mathbf{v}_i)}{\|\mathbf{v}_{i+1}\times\mathbf{v}_i\|\|\mathbf{v}_{i+2}\times\mathbf{v}_i\|} \right) - \pi,$$

where vertex indices are taken modulo 3, i.e., $\mathbf{v}_{i+3} = \mathbf{v}_i$.

Next, the line integral term $\int_{\partial\mathbf{A}}(\mathbf{a}\cdot\mathbf{i})\mathrm{d}\mathbf{i}$ is computed by summing contributions from each edge of the spherical triangle. For an edge from $\mathbf{v}_i$ to $\mathbf{v}_{i+1}$, the normal to the plane containing the great circle arc connecting $\mathbf{v}_i$ and $\mathbf{v}_{i+1}$ is $\mathbf{u} = \frac{\mathbf{v}_i\times\mathbf{v}_{i+1}}{\|\mathbf{v}_i\times\mathbf{v}_{i+1}\|}$. Define an orthonormal basis for this edge arc:

$$\mathbf{s} = \mathbf{v}_{i+1}, \quad \mathbf{t} = \frac{\mathbf{u}\times\mathbf{s}}{\|\mathbf{u}\times\mathbf{s}\|}.$$

Then project the direction vector $\mathbf{a}$ onto this basis:

$$\mathbf{m} = \mathbf{a}\cdot\mathbf{s}, \quad \mathbf{n} = \mathbf{a}\cdot\mathbf{t}, \quad \mathbf{l} = \sqrt{\mathbf{m}^2 + \mathbf{n}^2}.$$

The angle subtended by the arc and the phase angle are:

$$\theta = cos^{-1}(\mathbf{v}_{i+1}\cdot\mathbf{v}_i), \quad \phi = sgn(\mathbf{n})cos^{-1}\left(\frac{\mathbf{m}}{\mathbf{l}}\right),$$

where $sgn(x)$ is the sign function, returning -1 or 1. The total line integral is the sum over the three edges:

$$\int_{\partial\mathbf{A}} (\mathbf{a}\cdot\mathbf{i})\mathrm{d}\mathbf{i} = \sum_{i=0}^{2} l_i \left[sin(\theta_i - \phi_i) - sin(-\phi_i)\right],$$

where l_i, θ_i, ϕ_i are computed for each edge i.

2. LTC Mapper

Our LTC mapper network is a fully-connected neural network (multi-layer perceptron network) with five layers. It learns a function that estimates the parameters of an LTC lobe approximating a GT specular lobe for a given reflection state. The state of surface reflection includes reflection geometry and material properties. Since we assume that materials have an isotropic BRDF and the lobe shape is determined by surface roughness, the incident angle and roughness are given to the mapper as input. We train our network by sampling ray directions of arbitrary reflection states and penalizing the differences on samples.

The network takes the incident angle θ and roughness σ as inputs. It outputs five parameters $[\theta', \alpha, m_{11}, m_{22}, m_{31}]$. θ' is a residual angle to modulate the reflection axis. α estimates an amplitude of the lobe and $[m_{11}, m_{22}, m_{31}]$ construct an shearing and scaling matrix $\mathbf{S}$:

$$\theta_t = \theta + \theta',$$

$$\mathbf{S} = \alpha \begin{bmatrix} m_{11} & 0 & 0 \\ 0 & m_{22} & 0 \\ m_{31} & 0 & 1 \end{bmatrix},$$

$$\mathbf{R}(\mathbf{p}) = \begin{bmatrix} cos(\theta_t) & 0 & sin(\theta_t) \\ 0 & 1 & 0 \\ -sin(\theta_t) & 0 & cos(\theta_t) \end{bmatrix}.$$

We present the accuracy of the LTC mapper across various

Roughness	0.05	0.1	0.2	0.3	0.4
RMSE (%)	0.62	0.2	0.3	1.02	2.44

Table 1: *The relative errors between GGX BRDF lobes and their predicted LTC lobes.*

Figure 1: *Qualitative evaluation for a diffuse object.*

roughness values in Tab. 1. As the roughness decreases, the approximated cosine lobes align more closely with the original LTC lobes. The relative error remains consistently small.

3. Relighting Comparison with Baselines

We conduct relighting comparison with baselines that support relighting in both qualitative and quantitative ways shown in Fig. 2 and Tab. 2. We include NMF [MVKFK23], which accounts for indirect illumination, as an additional baseline. NMF struggles to correctly estimate material parameters, as its neural BRDF modeling compensates for the prediction error of classic BRDF parameters. EnvIDR [LCL*23] fails to synthesize the consistent material appearance because it does not explicitly model and estimate the physical roughness for surface reflection. InvRender [ZSH*22] cannot reconstruct high-frequency details due to the frequency limit of its spherical Gaussian lighting representation, resulting in blurred relighting results. We exclude NeRO [LWL*23] from evaluation because it does not use its light representation for relighting. Our method achieves the best performance among the STAR methods, thanks to our physically plausible rendering framework.

4. Evaluation of Diffuse Object

Fig. 1 shows that our method can accurately estimate materials for a given object as our method accurately reconstructs lighting by utilizing background rays.

5. Gradation-Specular Lighting Approximation

We approximate the integral by multiplying it by $\frac{\|\mathbf{Mi}'\|}{\|\mathbf{Mm}'\|}$, which changes the shapes of the LTC lobes. Fig. 3 illustrates the shapes of the original LTC lobes and their approximations. Although this introduces a numerical bias, the approximated lobes closely match the shapes of the original lobes.

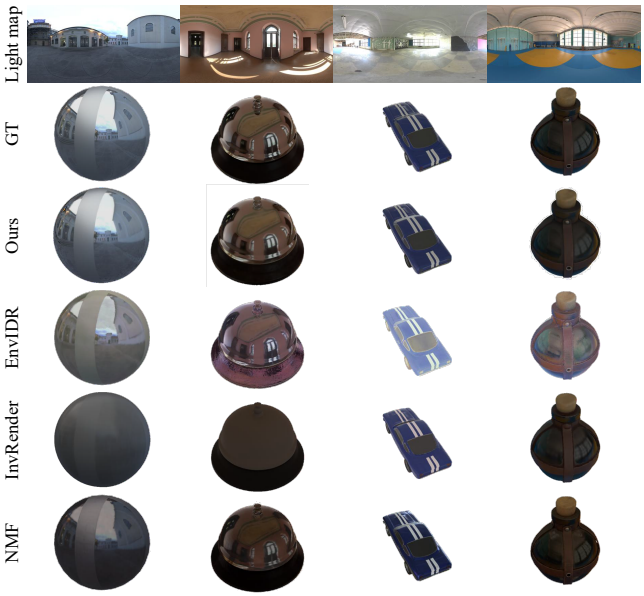

Figure 2: *Relighting comparison with baselines.*

	Ours		EnvIDR		InvRender		NMF	
	PSNR	LPIPS	PSNR	LPIPS	PSNR	LPIPS	PSNR	LPIPS
Potion	31.312	0.032	27.487	0.018	12.353	0.056	30.762	0.041
Musclecar	24.700	0.053	20.047	0.040	8.118	0.070	21.850	0.066
Hetero_ball	30.586	0.033	13.554	0.040	21.359	0.029	12.769	0.079
Tbell	24.857	0.055	19.918	0.046	13.036	0.098	19.451	0.075

Table 2: *Quantitative comparison for relighting given in Fig. 2. We mark the first, second, and third place with different colors.*

References

[Arv95] ARVO J.: Applications of irradiance tensors to the simulation of non-lambertian phenomena. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques* (1995), pp. 335–342. 1

[LCL*23] LIANG R., CHEN H., LI C., CHEN F., PANNEER S., VIJAYKUMAR N.: Envidr: Implicit differentiable renderer with neural environment lighting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (October 2023), pp. 79–89. 2

[LWL*23] LIU Y., WANG P., LIN C., LONG X., WANG J., LIU L., KOMURA T., WANG W.: Nero: Neural geometry and brdf reconstruction of reflective objects from multiview images. *ACM Transactions on Graphics (ToG) 42*, 4 (2023), 1–22. 2

[MVKFK23] MAI A., VERBIN D., KUESTER F., FRIDOVICH-KEIL S.: Neural microfacet fields for inverse rendering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2023), pp. 408–418. 2

[ZSH*22] ZHANG Y., SUN J., HE X., FU H., JIA R., ZHOU X.: Modeling indirect illumination for inverse rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022), pp. 18643–18652. 2

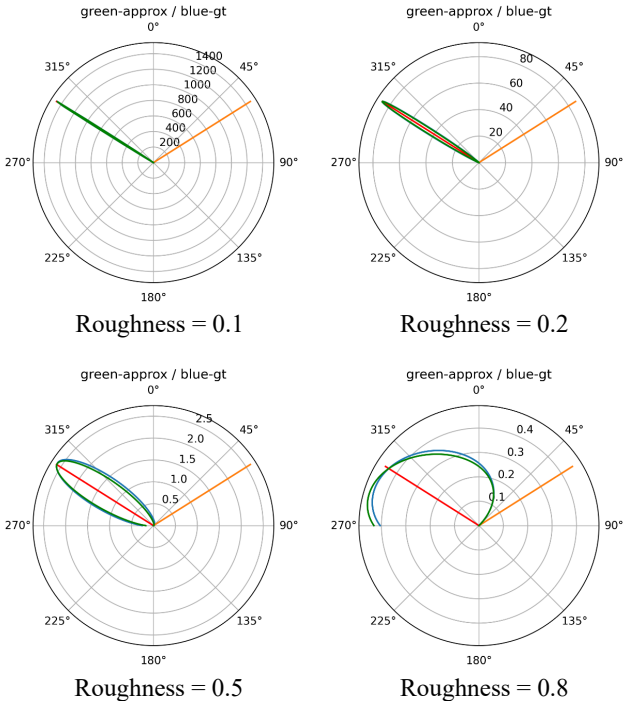

Figure 3: *LTC lobes and their approximations. We compare original LTC lobes (green)* $\left\langle (0,0,\frac{1}{\pi}),\boldsymbol{i}'\right\rangle$ *with their approximations (blue),* $\frac{\|\boldsymbol{Mi}'\|}{\|\boldsymbol{Mm}'\|}\left\langle (0,0,\frac{1}{\pi}),\boldsymbol{i}'\right\rangle$ *for various roughness (0.1, 0.2, 0.5, 0.8) at a fixed incident angle of 51.58 degrees. Plots are in a polar coordinate system with respect to the surface normal. Note that orange lines indicate the viewing directions, and red lines indicate the reflection directions.*