

Estimating Local Beckmann Roughness for Complex BSDFs

Nicolas Holzschuch
Inria

Anton Kaplanyan
Nvidia Research

Johannes Hanika
Karlsruhe Institute of Technology

Carsten Dachsbacher
Karlsruhe Institute of Technology




Figure 1: Indirect lighting (exposure in *b* and *c* increased for printouts) on three test scenes rendered with different materials: (a) multilayer coated plastic material, (b) measured materials on a ring, (c) CTD material on a car. The insets show difference to reference in CIE'76 ΔE . Top: single Gaussian, bottom: our local Gaussian approximation. We can render both analytic (a, c) and measured materials (b) more robustly because the local Gaussian approximation facilitates more even exploration of path space.

Abstract

Many light transport related techniques require an analysis of the blur width of light scattering at a path vertex, for instance a Beckmann roughness. Such use cases are for instance analysis of expected variance (and potential biased countermeasures in production rendering), radiance caching or directionally dependent virtual point light sources, or determination of step sizes in the path space Metropolis light transport framework: recent advanced mutation strategies for Metropolis Light Transport [Veach 1997], such as Manifold Exploration [Jakob 2013] and Half Vector Space Light Transport [Kaplanyan et al. 2014] employ local curvature of the BSDFs (such as an average Beckmann roughness) at all interactions along the path in order to determine an optimal mutation step size. A single average Beckmann roughness, however, can be a bad fit for complex measured materials (such as [Matusik et al. 2003]) and, moreover, such curvature is completely undefined for layered materials as it depends on the active scattering layer. We propose a robust estimation of local curvature for BSDFs of any complexity by using local Beckmann approximations, taking into account additional factors such as both incident and outgoing direction.

Keywords: global illumination, half vector space, light transport

Concepts: •Computing methodologies → Reflectance modeling; Ray tracing;

1 Local Beckmann Roughness Estimation

Our technique is based on the following observation: while for instance Half Vector Space Light Transport uses a Gaussian approximation of the BSDF, this approximation is used only *locally* to compute an average step size for path mutations. For BSDFs that cannot be approximated globally with a Gaussian, we compute a *local* approximation and use this to compute the step size (see Figure 2). We deliberately selected the Beckmann distribution for its

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s). © 2016 Copyright held by the owner/author(s). SIGGRAPH '16, July 24-28, 2016, Anaheim, CA, ISBN: 978-1-4503-4282-7/16/07 DOI: <http://dx.doi.org/10.1145/2897839.2927416>

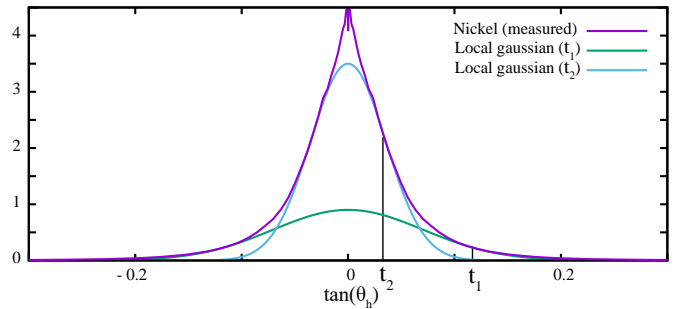


Figure 2: A measured BSDF (MERL nickel) approximated by two local Gaussians at positions $\tan \theta_h = t_1$ and t_2 respectively. These match the local shape much better than a single global Gaussian.

convenient properties, such as closed form multiplication and convolution, easy Fourier transform and computation of bandlimits and confidence intervals, as well as known optimal mutation step size.

Specifically, considering a path vertex \mathbf{v} , with incoming and outgoing directions \mathbf{i} and \mathbf{o} , we are looking for σ_v such that:

$$f_r(\mathbf{i}, \mathbf{o}) \approx k_v e^{-\left(\frac{\tan \theta_h}{\sigma_v}\right)^2}, \quad (1)$$

where θ_h is the zenith angle of the half vector and k_v folds all the other terms of the full BSDF. In a first-order approximation, we find σ_v by computing the derivative of f_r with respect to $\tan \theta_h$:

$$\frac{df_r(\mathbf{i}, \mathbf{o})}{d(\tan \theta_h)} = \frac{-2 \tan \theta_h}{\sigma_v^2} f_r(\mathbf{i}, \mathbf{o}) \quad (2)$$

$$\Rightarrow \sigma_v^2 = (-2 \tan \theta_h) f_r(\mathbf{i}, \mathbf{o}) \left(\frac{df_r(\mathbf{i}, \mathbf{o})}{d(\tan \theta_h)} \right)^{-1} \quad (3)$$

Equation (3) gives the variance for a local approximation of the BSDF by a Gaussian as a function of the reflectance and its derivative. If we have an explicit formulation of the BSDF, we compute the derivative analytically. For measured BSDFs, we compute the derivative through finite differences.

For BSDFs with multiple components, for instance a diffuse lobe plus a glossy lobe, Equation (3) automatically accounts for their relative importance: with a small diffuse lobe and a large glossy lobe, it is almost equal to the standard deviation of the glossy lobe.

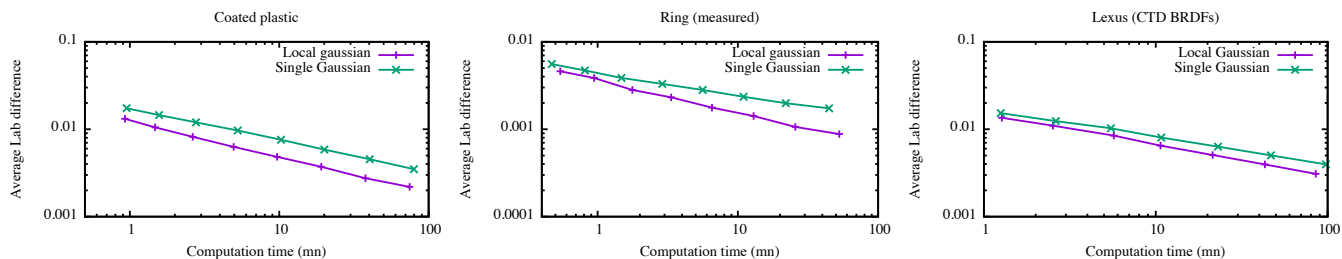


Figure 3: Convergence of the test scenes in Figure 4 as a function of time: average Lab difference per pixel (each point correspond to a different simulation). Using local Beckmann roughness to estimate mutation step size improves performance.

If the diffuse lobe dominates, Equation (3) gives larger values, consistent with the diffuse behaviour.

This local standard deviation can be inserted into the Half Vector Space Light Transport algorithm [Kaplanyan et al. 2014, Sec. 6.1]. For measured and multilobe BSDFs, the standard deviation depends on the distance to the specular reflected direction (see Figure 2): if we are far from the specular lobe, the diffuse component dominates, the estimated roughness is large, and we will explore a large area in path space. In the next iteration step we might be close to the specular lobe at this vertex, the estimated roughness will be small, and we will concentrate our exploration over a small neighborhood.

2 Results

We evaluated our method on different test scenes using measured materials (steel, brass, chrome [Matusik et al. 2003]) and analytic BSDF models with non-Gaussian distribution (see Figure 1 and 4). We used the CTD reflectance model [Holzschuch and Pacanowski 2015], which lets us edit the shininess of the materials while keeping other characteristics. For all test scenes we computed several images with increasing number of mutations per pixels and for each we show the Euclidean distance in Lab space (CIE'76 ΔE) to a reference solution. We compare to the baseline using a global average roughness, i.e. a single Gaussian. Visually, our method provides sharper reflections for highly-glossy materials, especially in combination with a more diffuse substrate. Benefits exist but are less obvious for rougher materials.

Figure 3 shows average Lab difference per pixel as a function of computation time. We used time as a metric to account for the cost of our method. The impact of derivative computation is small; longer runtimes are mainly due to longer Markov chains exploring complicated regions more thoroughly. We get an average speed-up of 50%: half the computation time for the same quality. The gain can be higher, up to 3 \times faster, depending on scene properties.



Figure 4: The test scenes with full global illumination (not only indirect light as Figure 1).

3 Conclusion

We proposed a simple method to extract local Beckmann roughness out of arbitrary (layered or measured) materials. This is a locally good fit to the actual half vector distribution and we showed how to improve Half Vector Space Light Transport by employing this approximation to derive mutation step sizes. The improvements of using a local Beckmann roughness are especially visible for highly glossy materials and long glossy chains, as can be found in the ring scene. We envision that more use cases can profit from this local approximation, such as more precise estimation of light field blur along a path.

References

- HANIKA, J., DROSKE, M., AND FASCIONE, L. 2015. Manifold next event estimation. *Computer Graphics Forum (Proceedings of Eurographics Symposium on Rendering)* 34, 4 (June), 87–97.
- HOLZSCHUCH, N., AND PACANOWSKI, R. 2015. A physically accurate reflectance model combining reflection and diffraction. Research Report RR-8807, INRIA, Nov.
- JAKOB, W. 2013. *Light transport on path-space manifolds*. PhD thesis, Cornell University.
- KAPLANYAN, A., HANIKA, J., AND DACHSBACHER, C. 2014. The natural-constraint representation of the path space for efficient light transport simulation. *Transactions on Graphics (Proceedings of SIGGRAPH)* 33, 4 (August), 1–13.
- MATUSIK, W., PFISTER, H., BRAND, M., AND McMILLAN, L. 2003. A data-driven reflectance model. *ACM Transactions on Graphics* 22, 3 (July), 759–769.
- VEACH, E. 1997. *Robust Monte Carlo Methods for Light Transport Simulation*. PhD thesis, Stanford University.