

Appendix: Uncertainty for SVBRDF Acquisition via Frequency Analysis

RUBEN WIERSMA, ETH Zurich, Delft University of Technology, Adobe Research, Switzerland

JULIEN PHILIP, Adobe Research, Netflix Eycline Studios, UK

MILOŠ HAŠAN, Adobe Research, USA

KRISHNA MULLIA, Adobe Research, USA

FUJUN LUAN, Adobe Research, USA

ELMAR EISEMANN, Delft University of Technology, The Netherlands

VALENTIN DESCHAINTE, Adobe Research, UK

CCS Concepts: • **Computing methodologies** → **Reflectance modeling**; **3D imaging**; **Appearance and texture representations**.

Additional Key Words and Phrases: Uncertainty, Inverse rendering

ACM Reference Format:

Ruben Wiersma, Julien Philip, Miloš Hašan, Krishna Mullia, Fujun Luan, Elmar Eisemann, and Valentin Deschaintre. 2025. Appendix: Uncertainty for SVBRDF Acquisition via Frequency Analysis. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers (SIGGRAPH Conference Papers '25)*, August 10–14, 2025, Vancouver, BC, Canada. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3721238.3730592>

A Details on- and Extensions to the Convolution Model

In this section, we provide a derivation of the convolution model for reflection. This derivation does not add technical novelty, but provides clarity on certain details that were implied in the original paper by Ramamoorthi and Hanrahan. We also propose a simple extension to the convolution model to include shadowing and masking effects, which were ignored by Ramamoorthi and Hanrahan. Finally, we discuss how to map the parameters of the Torrance-Sparrow BRDF to the Disney Principled BRDF [Burley 2012], which is commonly used in modern rendering pipelines.

A.1 Derivation of the Convolution Model

The outgoing radiance at point p along direction ω_o , $L_o(p, \omega_o)$ is given by

$$B(p, \omega_o) = \int_{H^2(\mathbf{n})} f_r(p, \omega_o, \omega_i) L(p, \omega_i) \cos \theta_i d\omega_i, \quad (1)$$

where f is the BRDF and $L(p, \omega_i)$ is the incident radiance along direction ω_i . For the Torrance-Sparrow BRDF, f is defined as

$$f(p, \omega_o, \omega_i) = K_d + K_s \frac{D(\omega_m) F(\omega_o \cdot \omega_m) G(\omega_i, \omega_o)}{4 \cos \theta_i \cos \theta_o}, \quad (2)$$

Authors' Contact Information: Ruben Wiersma, rubenwiersma@gmail.com, ETH Zurich, and Delft University of Technology, and Adobe Research, Switzerland; Julien Philip, Adobe Research, and Netflix Eycline Studios, UK; Miloš Hašan, Adobe Research, USA; Krishna Mullia, Adobe Research, USA; Fujun Luan, Adobe Research, USA; Elmar Eisemann, Delft University of Technology, The Netherlands; Valentin Deschaintre, deschain@adobe.com, Adobe Research, UK.



This work is licensed under a Creative Commons Attribution-NonCommercial 4.0 International License.

SIGGRAPH Conference Papers '25, August 10–14, 2025, Vancouver, BC, Canada

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1540-2/2025/08

<https://doi.org/10.1145/3721238.3730592>

where ω_m is the half-direction vector $\omega_m = (\omega_i + \omega_o) / \|\omega_i + \omega_o\|$; $D(\omega_m)$ is the normal distribution function; $F(\omega_o \cdot \omega_m)$ is the Fresnel term. Ramamoorthi and Hanrahan simplify this term to $F(\theta_o)$, as the angle θ_o is often close to the angle between ω_o and ω_m ; $G(\omega_i, \omega_o)$ is the shadowing-masking term. Ramamoorthi and Hanrahan ignore G . We assume shadowing and masking are independent statistical events, so that $G(\omega_i, \omega_o) = G(\omega_i)G(\omega_o)$.

There are two important notes about the denominator in Equation 2:

- (1) $1/(4 \cos \theta_o)$ results from the half-direction transform: the distribution of microfacets with a normal ω_m is transformed to the distribution of outgoing directions ω_o that the incoming light ray ω_i reflects toward (see Pharr et al., Equation 9.27).
- (2) $1/(\cos \theta_i)$ cancels out the cosine term applied to the incoming radiance (see Pharr et al., equation 9.30).

We now substitute Equation 2 into Equation 1 and split the equation into diffuse and specular

$$B(p, \omega_o) = K_d \int_{H^2(\mathbf{n})} L(p, \omega_i) \cos \theta_i d\omega_i \quad (3) \\ + K_s \int_{H^2(\mathbf{N})} \frac{D(\omega_m) F(\omega_o \cdot \omega_m) G(\omega_i, \omega_o)}{4 \cos \theta_i \cos \theta_o} L(p, \omega_i) \cos \theta_i d\omega_i$$

This equation is simplified by Ramamoorthi and Hanrahan using the assumptions that F only depends on θ_o and shadowing-masking is ignored. We replace the integral of incoming radiance for diffuse with the symbol for irradiance E .

$$B(p, \omega_o) = K_d E(p) + K_s F(\theta_o) \int_{H^2(\mathbf{N})} \frac{D(\omega_m)}{4 \cos \theta_o} L(p, \omega_i) d\omega_i. \quad (4)$$

Ramamoorthi and Hanrahan rewrite the specular term as a convolution between a filter based on D , and L . Crucially, the domain of D in the Torrance-Sparrow model is the half-angle space. In Ramamoorthi and Hanrahan's derivation, the spherical harmonic representation for this filter, in the paper referred to as S , is derived in incoming-direction space for normal exitance (Ramamoorthi and Hanrahan, Equation 27). This has two consequences:

- (1) We do not have to account for a change of variables and $1/(4 \cos \theta_o)$ can be removed.
- (2) In reality, S depends on the outgoing direction that is observed and thus, the filter changes shape. This variation is ignored with the explanation that "the BRDF filter is essentially symmetric about the reflected direction for small viewing angles, as well as for low frequencies l . Hence, it can be shown by Taylor-series expansions (and verified numerically) that the

corrections to Equation 20 [Equation 13 in our paper] are small under these conditions.”

This means that we can rewrite Equation 4 with a convolution

$$B(p, \omega_o) = K_d E(p) + K_s F(\theta_o) [S * L]_{\omega_o}, \quad (5)$$

which equals Equations 21 and 22 in Ramamoorthi and Hanrahan.

A.2 Shadowing and Masking

The approximate reflection function in Equation 5 does not include the shadowing or masking terms present in microfacet models [Pharr et al. 2023]. These terms model occlusions of incoming light (shadowing) and outgoing light (masking) due to the configuration of the microfacets. Ramamoorthi and Hanrahan argue that these terms can be ignored, because they mostly affect observations made at grazing angles. While this is true for materials with low roughness, we find that ignoring this term can lead to an appearance mismatch for high-roughness materials. We propose a simple way to include these terms in the convolution model to study the effect on uncertainty and BRDF acquisition. Shadowing and masking effects are typically modeled jointly to avoid an over-correction of the radiance (Pharr et al. [2023], Section 9.6.3). However, this joint term breaks the convolution model as presented in Equation 5, as the kernel would depend on both ω_i and ω_o . Therefore, we assume that shadowing and masking are independent and model them as $G_\alpha(\omega_i)G_\alpha(\omega_o)$. We first attenuate the incoming light with the shadowing term, convolve with the BRDF and attenuate the result with the masking term:

$$B(p, \omega_o) \approx K_d E(p) + K_s F(\theta_o) G_\alpha(\theta_o) [S_\alpha * G_\alpha(\theta_i) L(p)]_{\omega_o}, \quad (6)$$

where G_α is the shadowing-masking function in the Trowbridge and Reitz model.

A.3 Torrance-Sparrow to Principled BRDF

Many modern rendering pipelines employ variants of the Disney BRDF [Burley 2012], which is a combination of a diffuse term and a microfacet term with a user-friendly parametrization. The model also contains some additional features beyond the scope of the current work. We can formulate our model using the principled BRDF parameters, rather than the raw parameters of the Torrance-Sparrow model. We parameterize base color, metallicity, and roughness, mapping these to the Torrance-Sparrow model as

$$K_d = R_b, \quad (7)$$

$$K_s = 1, \quad (8)$$

$$R_0 = 0.04 + (R_b - 0.04)m, \quad (9)$$

$$F(\theta_o) = R_0 + (1 - R_0)(1 - \cos \theta_o)^5, \quad (10)$$

$$\alpha = r^2, \quad (11)$$

where R_b is the base color, m is metallicity and r is the roughness. We set K_s to 1, because the specular component is already scaled in the Fresnel term by R_0 . This parameterization is based on Schlick’s approximation [Schlick 1994] and follows the implementation in Mitsuba [Jakob et al. 2022].

B Least-Squares Regularization

We experiment with a number of options for least-squares fitting and show the results in Figure 1. In these experiments, we found that no regularization and constant regularization lead to incorrect estimations of the BRDF, which can be retrieved as the ratio of the outgoing radiance over the incoming radiance power spectrum. Our proposed solution using a non-constant regularizer leads to a smooth interpolation of missing regions. While it also ‘blurs’ the incoming radiance, the relative relation between incoming and outgoing radiance is preserved.

C Sampling Theory

The transformation from the directional domain to spherical harmonics begs the question: do we have enough samples to accurately recover the coefficients of the outgoing radiance? We know from Equation 13 that the BRDF acts as a low-pass filter parameterized by α . We connect this knowledge with sampling theory to derive lower bounds on sampling counts.

The Nyquist-Shannon theorem provides a lower bound on the number of samples required to exactly recover a band-limited signal using a Fourier series. Similar theorems have been developed for spherical harmonics [Driscoll and Healy 1994; McEwen et al. 2011; McEwen and Wiaux 2011]. These state that, to recover a spherical signal with band-limit ℓ^* , the number of samples should be $O(\ell^{*2})$. The sampling rate and related band-limit have direct consequences for BRDF recovery. Assume that the incoming light has been sampled at a high enough rate to be accurately recovered, for example, from projected photographs or a gazing sphere. Then the outgoing light is the weakest link, as it is sampled by moving the camera along N positions around the object. Sampling theory tells us that we can only accurately recover outgoing radiance that is band-limited to $\ell^* < \sqrt{N}$ degrees. Signals with non-zero amplitude in higher degrees will suffer from aliasing.

Fortunately, the BRDF acts as a low-pass filter on the incoming radiance (Equation 13). That means the outgoing radiance can fall into two categories, based on the α parameter of the material ($\alpha = \text{roughness}^2$): α is either too low or α is high enough to recover spherical harmonic coefficients. If α is too low, the low-pass filtering from the BRDF does not band-limit the signal enough to accurately recover with the given sampling rate. The threshold for α can be determined based on Equation 13. Let t be an acceptable attenuation factor for degrees $\ell > \ell^*$. We solve Equation 13 for t to find the lower bound, α' , for accurate recovery

$$\alpha' = \ell^{*2-1} \sqrt{-\ln t}. \quad (12)$$

An acceptable threshold t can be determined empirically, by investigating the reconstruction error for a set of environment maps. To provide some intuition, for $N = 400$ samples and a threshold of $t = 0.5$, $\alpha' \approx 0.07$. Above this threshold, our method can recover α and K_s , provided that the incoming radiance has enough amplitude in the right degrees. This also extends to non-uniform samples, because the Nyquist-Shannon theorem holds for non-uniform samples [Marvasti 2012]. In other words: if a lower bound on α is known, it does not matter where the camera is placed, as long as the average distance to the closest sample is equal to $1/N$. It also means that one

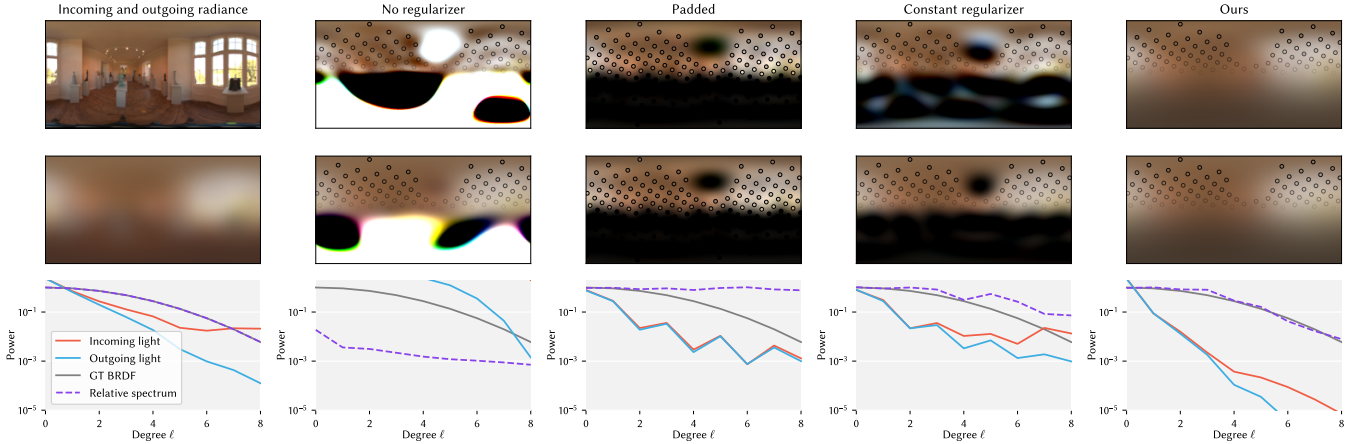


Fig. 1. A comparison of the options for fitting spherical harmonics on sparse samples. In the first column, top: incoming radiance (synthetic) on the sphere mapped to a lat-long grid. Middle: synthesized outgoing radiance as a result of filtering incoming radiance with Equation 13, $\alpha = 0.2$. Bottom: power spectra of incoming and outgoing radiance and ground-truth filter (BRDF) plotted on a log-scale. In the second to fifth column, we show the result of fitting spherical harmonics coefficients to samples and then transforming back to the spatial domain. We use 100 samples from the upper hemisphere of the input radiance (first column) and simulate missing samples due to occlusion or missed camera positions by masking some points, leaving 88 samples. The samples are weighted with $\cos \theta$, which is visualized as the transparency of the samples. We find that our method with a weighted regularizer is able to smoothly interpolate missing values and retrieves the correct BRDF filter in the power spectrum, where other variants overfit, or add dark regions in the upper hemisphere.

can determine the number of required views based on the lowest α that should be recovered: $N \sim \alpha^{-2}$.

It is important to understand what happens if $\alpha < \alpha'$. First, we would be uncertain where α lands between 0 and α' , based on the power spectrum alone. For $0 < \alpha < \alpha'$, Equation 13 is close to 1 for all degrees below ℓ^* . Second, because this situation occurs for low α , the outgoing radiance should be similar to the incoming radiance, up to a scale factor for absorption and transmission. It is unlikely that the spherical harmonics decomposition with significant aliasing will match a filtered version of the incoming radiance. Therefore, we can detect that $\alpha < \alpha'$. In this case, the MSE for any parameter combination ψ is relatively high. Once such a case is detected, we know that our spherical harmonic-based analysis provides no further insights on (un)certainty. There is still a chance for accurate BRDF recovery if $\alpha < \alpha'$. A sample might land on a fortunate spot in the outgoing radiance field. This is the case when there is high local variation in the incoming light around the sample locations, resulting in large changes in radiance for small changes in α .

D Background

Our method builds on prior work in inverse rendering and spherical harmonics. We summarize required background knowledge and refer to related work for further depth.

D.1 Reflection as Convolution

In the inverse rendering framework of Ramamoorthi and Hanrahan [2001], detailed in subsection A.1, estimating the specular BRDF parameters comes down to estimating the convolution kernel S . This can be done efficiently since a convolution in the angular domain can be represented as a multiplication in the spherical harmonics frequency domain. Using this representation, one can find the

convolution kernel through a division of the spherical harmonics coefficients of the outgoing radiance by those of the incoming radiance. This is analogous to kernel estimation for image deblurring in the Fourier domain. Ramamoorthi and Hanrahan derive a number of conclusions from this insight, which we summarize in this section.

One crucial insight concerns the well-posedness of BRDF recovery. Ramamoorthi and Hanrahan state that the recovery of BRDF parameters is ill-posed if the input lighting has no amplitude along certain modes of the filter (BRDF). Those modes cannot be estimated without additional priors on plausible spatial parameter variations. For the microfacet BRDF, this leads to the following conclusion: if the incoming light only contains frequencies $\ell \ll \alpha^{-1}$, multiplying the coefficients of the light with those of the BRDF only results in a small difference, and the inversion of this operation is ill-conditioned. To accurately estimate α , the incoming light used during the capture needs to exhibit sufficiently high frequencies.

This insight is based on a derivation for the coefficients of the microfacet model. The normalized SH-coefficients of the specular component of the BRDF for normal incidence, S , are approximated by

$$\hat{f}_{\ell m} \approx e^{-(\alpha \ell)^2}, \quad (13)$$

which is a Gaussian in the frequency domain with a width determined by α . The kernel is derived from a Beckmann normal distribution function and the α parameter corresponds to the α parameter there. Note that these coefficients do not vary with the Spherical Harmonics order m , since the normal distribution function is isotropic for outgoing rays in the direction of the normal vector ($\theta_o = 0$). An important approximation employed by Ramamoorthi and Hanrahan is that this same kernel can be used for any outgoing direction. While the correct kernel varies with the incoming and outgoing direction, this approximation does not lead to significant error for

inverse rendering [Ramamoorthi and Hanrahan 2001] as they note that “it can be shown by Taylor-series expansions and verified numerically, that the corrections to this filter are small [for low degrees ℓ]”.

In this paper, we expand on the theory established by Ramamoorthi and Hanrahan by improving the reflection as a convolution model’s accuracy and developing the implications for well-posedness into quantifiable metrics on uncertainty without a-priori knowledge on α .

D.2 Spherical Harmonics

Spherical harmonics are a series of orthonormal basis functions on the sphere, indexed by their degree ℓ and order m . We use them to represent incoming and outgoing radiance over incoming and outgoing directions on the unit sphere. Spherical harmonics are analogous to the Fourier series on a flat domain, where the frequency of the Fourier series corresponds to the degree ℓ and order m . We provide a brief overview of properties relevant to our method. For further details, a helpful reference and software package is published by Wieczorek and Meschede [2018].

Any real, square-integrable function on the sphere can be expressed as a spherical-harmonics series:

$$f(\theta, \phi) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} f_{\ell m} Y_{\ell m}(\theta, \phi), \quad (14)$$

where $f_{\ell m}$ is the coefficient for spherical harmonic $Y_{\ell m}(\theta, \phi)$, given as

$$N_{\ell m} = \sqrt{\frac{(2 - \delta_{m0})(2\ell + 1)}{4\pi} \frac{(\ell - m)!}{(\ell + m)!}} \quad (15)$$

$$Y_{\ell m}(\theta, \phi) = \begin{cases} N_{\ell m} P_{\ell}^m(\cos \theta) \cos m\phi & \text{if } m \geq 0, \\ N_{\ell m} P_{\ell}^{|m|}(\cos \theta) \sin |m|\phi & \text{if } m < 0. \end{cases} \quad (16)$$

$N_{\ell m}$ is a normalization factor, δ_{m0} is the Kronecker delta function, which evaluates to 1 when $m = 0$, and P_{ℓ}^m is the associated Legendre function for degree ℓ and order m . The total number of spherical harmonics up to- and including a maximum degree, ℓ^* , equals $(\ell^* + 1)^2$. A useful property of spherical harmonics in our setting is that a rotational convolution on the sphere is equal to multiplication of coefficients in the spherical harmonic domain.

The power spectrum of a spherical function f can be computed from the spherical harmonic coefficients per degree

$$S_f(\ell) = \sum_{m=-\ell}^{\ell} f_{\ell m}^2. \quad (17)$$

The power spectrum is invariant to rotations of the coordinate system. In our context that means the power spectrum is invariant to slight perturbations of the normals at each point.

D.2.1 Computing spherical harmonic coefficients. One can find the SH coefficients for a function f by computing the inner product with the basis functions

$$f_{\ell m} = \int_{S^2} f(\theta, \phi) Y_{\ell m}(\theta, \phi) d\omega. \quad (18)$$

A useful property holds for the coefficient of degree $\ell = 0$, for which the spherical harmonic is constant; $Y_{00}(\theta, \phi) = (4\pi)^{-\frac{1}{2}}$. The corresponding coefficient, f_{00} , is equal to the integral of f times the normalization constant, $(4\pi)^{-\frac{1}{2}}$. The spherical harmonics for higher degrees all integrate to zero¹. This is relevant in the context of rendering, because the total integrated incoming and outgoing radiance can be read from the 0th degree coefficient and that coefficient alone. In the general case, we estimate the coefficient for $f_{\ell m}$ based on samples of f . The sampling method determines how these coefficients are estimated.

Regular sampling. If f is sampled on a grid with equally spaced longitudinal and latitudinal angles, this integral can be accelerated using a fast Fourier transform in the longitudinal direction ϕ and a quadrature rule in the latitudinal direction θ [Driscoll and Healy 1994]. In our setting, this approach can be used for environment maps represented as rectangular textures.

Irregular sampling. During capture the camera is often placed at irregular positions, leading to non-uniform (θ, ϕ) samples. Further, a point on the surface might be observed from only a few positions. We therefore often need to use sparse and irregular samples to fit spherical harmonic coefficients. We do so by fitting the coefficients using least-squares, expressing Equation 14 as a linear system

$$Yc \approx f, \quad (19)$$

where f is a vector of n discrete samples from f , Y is a matrix of size $n \times (\ell^* + 1)^2$ containing the spherical harmonics sampled at the same locations as f , and c is a vector of the $(\ell^* + 1)^2$ coefficients we want to find. We can find c by solving a least-squares system

$$Y^T Yc = Y^T f. \quad (20)$$

To be well posed, this system requires $n > (\ell^* + 1)^2$ independent samples, which can be challenging in the context of sparse sampling, making the system under-constrained. We propose to use a custom regularizer in our work for cases where the number of samples is too low.

E Stanford ORB Reference Results

We include the results table from StanfordORB for reference. These results were obtained under different acquisition condition and cannot be directly compared to our results.

References

- Mark Boss, Raphael Braun, Varun Jampani, Jonathan T Barron, Ce Liu, and Hendrik Lensch. 2021. Nerd: Neural reflectance decomposition from image collections. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 12684–12694.
- Brent Burley. 2012. Physically-based shading at Disney. In *Acm Siggraph*, Vol. 2012. vol. 2012, 1–7.
- J.R. Driscoll and D.M. Healy. 1994. Computing Fourier Transforms and Convolutions on the 2-Sphere. *Advances in Applied Mathematics* 15, 2 (1994), 202–250. <https://doi.org/10.1006/aama.1994.1008>
- Jon Hasselgren, Nikolai Hofmann, and Jacob Munkberg. 2022. Shape, light, and material decomposition from images using Monte Carlo rendering and denoising. *Advances in Neural Information Processing Systems* 35 (2022), 22856–22869.
- Wenzel Jakob, Sébastien Speierer, Nicolas Roussel, Merlin Nimier-David, Delio Vicini, Tizian Zeltner, Baptiste Nicolet, Miguel Crespo, Vincent Leroy, and Ziyi Zhang. 2022. *Mitsuba 3 renderer*. <https://mitsuba-renderer.org>.

¹Because the spherical harmonics are orthonormal, the inner product between any spherical harmonic with $\ell > 0$ and the constant function ($\ell, m = 0$) equals zero.

Table 1. Benchmark Comparison for **Novel Scene Relighting** of Existing Methods from [Kuang et al. 2023]. † denotes models trained with the ground-truth 3D scans and pseudo materials optimized from light-box captures. The rest of results are obtained by optimizing jointly for illumination, geometry and material. *We report these numbers for reference, however they cannot be directly compared to our results.*

| | PSNR-H↑ | PSNR-L↑ | SSIM↑ | LPIPS↓ |
|--|---------|---------|--------------|--------------|
| NVDiffRecMC [Hasselgren et al. 2022] † | 25.08 | 32.28 | 0.974 | 0.027 |
| NVDiffRec [Munkberg et al. 2022] † | 24.93 | 32.42 | 0.975 | 0.027 |
| PhySG [Zhang et al. 2021a] | 21.81 | 28.11 | 0.960 | 0.055 |
| NVDiffRec [Munkberg et al. 2022] | 22.91 | 29.72 | 0.963 | 0.039 |
| NeRD [Boss et al. 2021] | 23.29 | 29.65 | 0.957 | 0.059 |
| NeRFactor [Zhang et al. 2021b] | 23.54 | 30.38 | 0.969 | 0.048 |
| InvRender [Wu et al. 2023] | 23.76 | 30.83 | 0.970 | 0.046 |
| NVDiffRecMC [Hasselgren et al. 2022] | 24.43 | 31.60 | 0.972 | 0.036 |
| Neural-PBIR [Sun et al. 2023] | 26.01 | 33.26 | 0.979 | 0.023 |

Zhengfei Kuang, Yunzhi Zhang, Hong-Xing Yu, Samir Agarwala, Shangzhe Wu, and Jiajun Wu. 2023. Stanford-ORB: A Real-World 3D Object Inverse Rendering Benchmark. arXiv:2310.16044 [cs.CV]

Farokh Marvasti. 2012. *Nonuniform sampling: theory and practice*. Springer Science & Business Media.

Jason D. McEwen, Gilles Puy, Jean-Philippe Thiran, Pierre Vandergheynst, Dimitri Van De Ville, and Yves Wiaux. 2011. Sampling theorems and compressive sensing on the sphere. In *Wavelets and Sparsity XIV*, Manos Papadakis, Dimitri Van De Ville, and Vivek K. Goyal (Eds.), Vol. 8138. International Society for Optics and Photonics, SPIE, 81381F. <https://doi.org/10.1117/12.893481>

Jason D McEwen and Yves Wiaux. 2011. A novel sampling theorem on the sphere. *IEEE Transactions on Signal Processing* 59, 12 (2011), 5876–5887.

Jacob Munkberg, Jon Hasselgren, Tianchang Shen, Jun Gao, Wenzheng Chen, Alex Evans, Thomas Müller, and Sanja Fidler. 2022. Extracting triangular 3d models, materials, and lighting from images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 8280–8290.

Matt Pharr, Wenzel Jakob, and Greg Humphreys. 2023. *Physically based rendering: From theory to implementation*. MIT Press.

Ravi Ramamoorthi and Pat Hanrahan. 2001. A signal-processing framework for inverse rendering. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques (SIGGRAPH '01)*. Association for Computing Machinery, New York, NY, USA, 117–128. <https://doi.org/10.1145/383259.383271>

Christophe Schlick. 1994. An Inexpensive BRDF Model for Physically-Based Rendering. *Comput. Graph. Forum* 13, 3 (1994), 233–246. <https://doi.org/10.1111/1467-8659.1330233>

Cheng Sun, Guangyan Cai, Zhengqin Li, Kai Yan, Cheng Zhang, Carl Marshall, Jia-Bin Huang, Shuang Zhao, and Zhao Dong. 2023. Neural-PBIR Reconstruction of Shape, Material, and Illumination. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 18046–18056.

TS Trowbridge and Karl P Reitz. 1975. Average irregularity representation of a rough surface for ray reflection. *JOSA* 65, 5 (1975), 531–536.

Mark A. Wieczorek and Matthias Meschede. 2018. SHTools: Tools for Working with Spherical Harmonics. *Geochemistry, Geophysics, Geosystems* 19, 8 (2018), 2574–2592. <https://doi.org/10.1029/2018GC007529> arXiv:<https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2018GC007529>

Haoqian Wu, Zhipeng Hu, Lincheng Li, Yongqiang Zhang, Changjie Fan, and Xin Yu. 2023. NeFI: Inverse Rendering for Reflectance Decomposition with Near-Field Indirect Illumination. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 4295–4304.

Kai Zhang, Fujun Luan, Qianqian Wang, Kavita Bala, and Noah Snavely. 2021a. Physg: Inverse rendering with spherical gaussians for physics-based material editing and relighting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 5453–5462.

Xiuming Zhang, Pratul P Srinivasan, Boyang Deng, Paul Debevec, William T Freeman, and Jonathan T Barron. 2021b. Nerfactor: Neural factorization of shape and reflectance under an unknown illumination. *ACM Transactions on Graphics (ToG)* 40, 6 (2021), 1–18.