

Document Version

Final published version

Licence

CC BY-NC

Citation (APA)

Wiersma, R., Philip, J., Hašan, M., Mullia, K., Luan, F., Eisemann, E., & Deschaintre, V. (2025). Uncertainty for SVBRDF Acquisition using Frequency Analysis. In S. N. Spencer (Ed.), *Proceedings - SIGGRAPH 2025 Conference Papers* Article 169 (Proceedings - SIGGRAPH 2025 Conference Papers). Association for Computing Machinery (ACM). <https://doi.org/10.1145/3721238.3730592>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership. Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.



Uncertainty for SVBRDF Acquisition using Frequency Analysis

RUBEN WIERSMA, ETH Zurich, Switzerland, Delft University of Technology, Netherlands, and Adobe Research, USA
 JULIEN PHILIP, Adobe Research, United Kingdom and Netflix Eycline Studios, United Kingdom
 MILOŠ HAŠAN, Adobe Research, USA
 KRISHNA MULLIA, Adobe Research, USA
 FUJUN LUAN, Adobe Research, USA
 ELMAR EISEMANN, Delft University of Technology, Netherlands
 VALENTIN DESCHAINTE, Adobe Research, United Kingdom

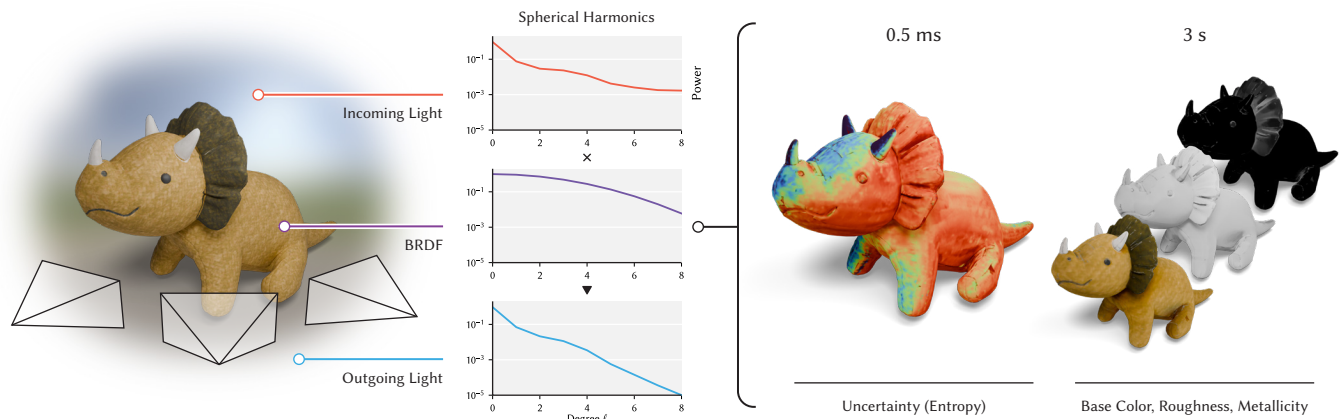


Fig. 1. We quantify uncertainty for SVBRDF acquisition from multi-view captures (left) using entropy (right). We significantly accelerate the, otherwise heavy, computation in the frequency domain (spherical harmonics power spectrum, center), leading to a practical, efficient method.

This paper aims to quantify uncertainty for SVBRDF acquisition in multi-view captures. Under uncontrolled illumination and unstructured viewpoints, there is no guarantee that the observations contain enough information to reconstruct the appearance properties of a captured object. We study this ambiguity, or uncertainty, using entropy and accelerate the analysis by using the frequency domain, rather than the domain of incoming and outgoing viewing angles. The result is a method that computes a map of uncertainty over an entire object within a millisecond. We find that the frequency model allows us to recover SVBRDF parameters with competitive performance, that the accelerated entropy computation matches results with a physically-based path tracer, and that there is a positive correlation between error and uncertainty. We then show that the uncertainty map can be applied to

improve SVBRDF acquisition using capture guidance, sharing information on the surface, and using a diffusion model to inpaint uncertain regions. Our code is available at https://github.com/rubenwiersma/svbrdf_uncertainty.

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. Uncertainty for SVBRDF Acquisition using 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, 12 pages. <https://doi.org/10.1145/3721238.3730592>

1 Introduction

Recovering realistic appearance properties for digital objects is inherently challenging. We propose to estimate the uncertainty of the process, given a set of views of an object and information about the incident lighting. This enables applications such as viewpoint planning and information sharing for improved acquisition.

Appearance is typically modeled as a six-dimensional function, depending on the incoming and outgoing light directions and locations on the surface. Accurately reconstructing this high-dimensional function requires many samples and, ideally, control of both lighting and camera positions. Such capture requires complex and expensive



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

SIGGRAPH Conference Papers '25, Vancouver, BC, Canada

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

ACM ISBN 979-8-4007-1540-2/25/08

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

hardware, such as light stages, and is particularly difficult for on-site capture. It is more convenient to *passively* photograph the object from a set of viewpoints, without control of the lighting.

Passive captures often lead to ambiguity about the material of an object, as shown in Figure 2. Depending on the lighting and viewpoint, it may be hard to distinguish matte from glossy materials. On complex, curved shapes, such ambiguities can arise, e.g., from occlusion, shadowing, and missing viewpoints. In this work, our main goal is to estimate the ambiguity, or *uncertainty*, when recovering SVBRDF parameters. The result is a map on the surface, quantifying how certain we are about the parameter estimates (Figure 1). We demonstrate that this information can improve SVBRDF recovery through capture guidance, information sharing, and inpainting of uncertain areas. We also show that our contributions can be used for fast recovery of SVBRDF parameters.

To quantify uncertainty, we use entropy. Entropy measures how ‘spread out’ a probability distribution is: in the overcast setting (left of Figure 2), the planes being glossy or matte is equally likely. The probability distribution is uniform, leading to high entropy. In the setting on the right, it is very unlikely that the plane on the right is matte. The probability distribution is concentrated around the glossy material, resulting in low entropy. To compute entropy, we require knowledge of the posterior probability distribution over the entire space of possible parameters. This requires rendering the scene from every input viewpoint for all possible parameter combinations. Prior approaches for estimating uncertainty handle this by sparsely sampling the parameter space. For example, Lensch et al. [2003b] and Goli et al. [2024] use the inverse of the Hessian around the optimized parameters as a proxy for uncertainty. This approach, however, is local in parameter space and requires running an optimization first. Zhou et al. [2024] sample the full parameter space using stochastic particle-optimization sampling (SPOS) [Zhang et al. 2020]. While this approach considers the full parameter space, it is slow to compute.

Rather than focusing on reducing the number of samples, we accelerate the computation of the probability for each parameter combination using frequency-domain analysis [Ramamoorthi and Hanrahan 2001]. The resulting approach is lightweight and highly parallelizable, making entropy as a measure of uncertainty practical for SVBRDF recovery. The frequency-domain analysis considers reflection functions as convolutions of the BRDF over the incoming light, easily computed in the frequency domain. We show that the entire reflection function can be analyzed within the frequency domain, given careful approximations. This eliminates repeated transforms to the angular domain. Second, we show that the analysis of uncertainty within the frequency domain can be accelerated by only considering the power spectra of spherical harmonics. We also find that uncertainty is determined by two parameters, the specular coefficient and roughness, further reducing the complexity of the problem. To be practical, these contributions require a spherical-harmonics transform for sparse and irregular samples of the incoming and outgoing radiance, as these are sampled through unstructured photographs. We therefore propose a robust spherical-harmonics transform that uses least-squares fitting on sparse and irregular points, with regularization specialized for natural lighting.




Fig. 2. Passive capture can lead to ambiguity about the material of an object. Which plane is glossy and which is matte? While it is hard to see in the left scene, the materials are the same in both scenes (left matte, right glossy).

We validate our method by comparing our entropy estimation to entropy computed with a state-of-the-art path tracer [Jakob et al. 2022]. We find that our approach correlates strongly ($\rho = 0.9$) and runs in a fraction of the time (1/700 000). We also evaluate the relationship between entropy and reconstruction error and show three applications of our uncertainty evaluation: capture guidance, information sharing between certain and uncertain areas on the surface, and inpainting regions with high uncertainty using a diffusion model. Finally, we demonstrate that our framework can be used to optimize SVBRDF parameters with state-of-the-art relighting results and a 10x speedup and use these results to initialize other approaches.

In summary, we propose to compute uncertainty for the common passive capture setup, through the following contributions:

- Uncertainty for SVBRDF acquisition through entropy.
- A practical and fast entropy computation, following from an analysis in the frequency domain and power spectrum simplification.
- Demonstration of applications for uncertainty.

2 Related Work

We provide a detailed overview of related work in uncertainty estimation for (SV)BRDF recovery and a brief analysis of how approaches for (SV)BRDF recovery deal with ambiguity.

2.1 Uncertainty Estimation

Similar to us, prior works study uncertainty from a Bayesian perspective, analyzing the posterior probability distribution over BRDF parameters. Most works sample parameters sparsely, to avoid costly rendering operations. We distinguish approaches that sample the parameter space locally (Laplace’s approximation) or globally. Laplace’s approximation assumes the posterior follows a Gaussian distribution centered around the optimal parameters and uses the inverse of the Hessian as a covariance matrix, representing uncertainty. Lensch et al. [2003b] use this approach to select viewpoints that minimize uncertainty. In another work, Lensch et al. [2003a] use the covariance matrix to optimally split clusters of materials. Recently, Goli et al. [2024] use Laplace’s approximation to quantify uncertainty for NeRFs [Mildenhall et al. 2020]. They also show a correlation between uncertainty and absolute error. This information is used to remove spurious geometry. In our work, we approximate the rendering equation, which allows us to sample the posterior probability distribution globally along a dense grid. We directly use entropy to measure uncertainty, rather than the covariance matrix.

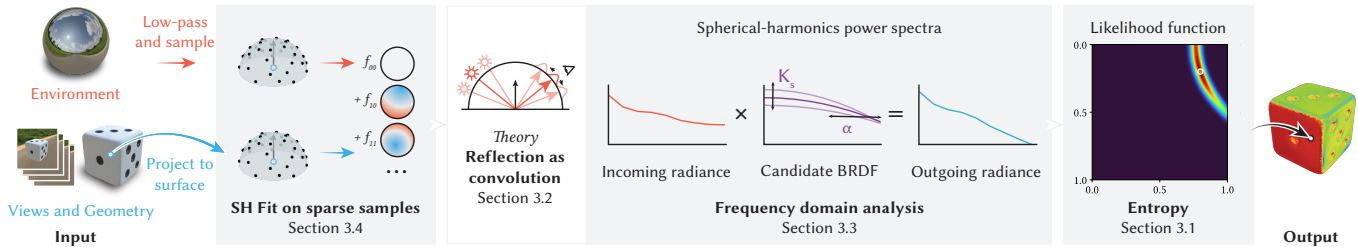


Fig. 3. An overview of our pipeline showing input, output, and the proposed algorithm. The input to our approach is a set of photographs of an object from multiple viewpoints and the associated camera extrinsics and intrinsics. We also provide the input lighting as an environment map and object geometry. We first estimate spherical harmonic coefficients on both the incoming and outgoing radiance and then estimate the effect of different BRDF filters within the power spectrum. This is used to compute a measure of uncertainty for the predicted parameters of the acquisition.

Zhou et al. [2024] sample the full parameter space using stochastic particle-optimization sampling (SPOS) [Zhang et al. 2020]. SPOS optimizes samples through particle-optimization to be distributed according to a target function that resembles the posterior. Then, the variance of these samples is used as a proxy for uncertainty. These contributions are orthogonal to our work. Because we evaluate the rendering equation in the frequency domain, our approach takes roughly one millisecond to compute uncertainty, while Zhou et al. [2024] report “a couple to 20 minutes” on comparable hardware.

Rodriguez-Pardo et al. [2023] used Monte-Carlo dropout to estimate uncertainty in the context of single-view material estimation, taking inspiration from Bayesian methods ([Gal and Ghahramani 2016]). Our method is designed for multi-view captures and measures uncertainty in the input captures, rather than the uncertainty of a predictive model.

2.2 Frequency-Based Light Transport

Our work builds on the idea that BRDFs can be expressed as low-pass filters in the frequency domain [Durand et al. 2005; Ramamoorthi and Hanrahan 2001]. This concept is widely adopted in interactive rendering, e.g., [Bagher et al. 2013]. We apply it in the context of material acquisition, where it has been used for controlled illumination [Aittala et al. 2013; Ghosh et al. 2007], controlling the frequency of light patterns to estimate the BRDF filter parameter through deconvolution of the reflected light. We do not assume control of the light and operate in the spherical-harmonics frequency domain rather than the Fourier domain [Aittala et al. 2013] or custom basis functions [Ghosh et al. 2007]. Our analysis therefore works with arbitrary natural lighting environments.

Our work is closest in spirit to that of Ramamoorthi and Hanrahan [2001]. Ramamoorthi and Hanrahan’s analysis allows one to conclude whether the BRDF recovery is well-posed or not, given prior knowledge of the width of the normal distribution function. Our uncertainty estimation is easier to apply as it gives a more precise and continuous answer to the question of uncertainty and does not assume prior knowledge of the material. Second, our approach works directly in the frequency domain and supports arbitrary light setups by using a robust spherical-harmonics transform. Finally, we improve the BRDF model to include shadowing and masking

and demonstrate that our improvements to the method result in competitive performance in (SV)BRDF capture.

2.3 (SV)BRDF Recovery

Our method operates in the context of (SV)BRDF recovery using optimization-based methods [Loubet et al. 2019; Munkberg et al. 2022; Nimier-David et al. 2021; Vicini et al. 2022]. These methods optimize BRDF parameters and, optionally, geometry and lighting, to reproduce the appearance of a set of input views. Existing approaches try to reduce the ambiguity in their capture, which can lead to suboptimal convergence. A direct way to reduce ambiguity is to use many photographs and controlled lighting [Aittala et al. 2013; Dupuy and Jakob 2018; Nam et al. 2018] or object orientations [Dong et al. 2014]. Some rely on specialized hardware to capture polarimetric information [Hwang et al. 2022]. Others use prior knowledge to compensate for limited information, such as stationarity of the captured materials [Aittala et al. 2016, 2015; Henzler et al. 2021; Xu et al. 2016].

Another way to reduce ambiguity is to use data-based priors. This allows the capture of (SV)BRDFs from as little as one image [Deschaintre et al. 2018, 2019, 2020; Guo et al. 2021; Li et al. 2017; Luo et al. 2024; Martin et al. 2022; Shah et al. 2023; Shi et al. 2020; Vecchio et al. 2023; Zhou et al. 2023; Zhou and Kalantari 2021; Zhou and Khademi Kalantari 2022]. For material acquisition on curved surfaces, approaches have been proposed to extract material parameters from single, multiple flash, or multi-focal photographs [Boss et al. 2020; Deschaintre et al. 2021; Fan et al. 2023; Li et al. 2018]. More recently, diffusion priors have been used to complete the missing lighting information in the context of acquisition [Lyu et al. 2023].

Our method is intended to work in conjunction with general material acquisition approaches. In our applications, we show several ways how our uncertainty analysis can be used to improve acquisition, both during optimization (e.g., by sharing information) and by using data-based priors (e.g., through diffusion model inpainting).

3 Method

Our goal is to evaluate uncertainty in material reconstruction during a ‘passive’ capture. The result is a mapping on the surface that denotes where the signal is insufficient to guarantee an accurate

reconstruction. The input to our method is a set of views from multiple viewpoints and we assume that the camera extrinsics, intrinsics, object geometry and HDR environmental lighting are known but not controlled.

Figure 3 provides an overview of our method. In this section, we work our way back from the desired output, entropy, to a practical and efficient implementation. Entropy requires evaluating the reconstruction error of many material parameters for the entire surface, which we accelerate by building on the signal processing framework proposed by Ramamoorthi and Hanrahan [2001] (frequency domain analysis). We analyze the uncertainty problem in the frequency domain and find that, given some approximations, we can evaluate entropy entirely using the power spectrum of the spherical harmonics representation of the incoming and outgoing radiance. Finally, we show how to robustly transform sparse and irregular radiance samples in the angular domain to spherical harmonics.

3.1 Entropy and Uncertainty

Entropy, in the context of information theory, measures the uncertainty about an unknown variable x over the possible states X [Shannon 1948]. Given a probability distribution $p(x)$, entropy is defined in the discrete setting as

$$H = - \sum_{x \in X} p(x) \log p(x). \quad (1)$$

Intuitively, entropy measures the ‘spread’ of the probability distribution $p(x)$. For example, a uniform distribution has maximum entropy and a narrow, concentrated distribution has low entropy. We are interested in measuring the uncertainty over the continuous parameter combinations ψ for a BRDF at each point on the surface (e.g., roughness, metallicity, base color), given the observed outgoing radiance B and incoming radiance L . For simplicity of exposition, we will consider n evenly spaced discrete parameter choices ψ . This discretization can be adapted to incorporate perceptual non-linearity¹. Computing continuous entropy also requires discretization, because an analytical solution is not available for the integral, and results in the same formulation for entropy we derive here.

First, we need to know the *likelihood* that the parameters ψ explain the observations B , given L , $p(\psi|B, L)$. This likelihood can be derived using Bayes’ rule from $p(B|\psi, L)$, the *probability* that we observe radiance B , given the parameters ψ and incoming light L

$$p(\psi|B, L) \propto p(B|\psi, L)p(\psi). \quad (2)$$

In our experiments, we assume that $p(\psi)$ is uniform, but one could include priors on materials based on real-world data without affecting the remainder of this discussion. If our measurements and models were perfect, $p(B|\psi, L)$, would be 1 for the radiance observation predicted by our BRDF model and 0 otherwise. However, we encounter measurement noise and our renders are approximate. If we assume this error follows a Gaussian distribution with variance σ^2 , we can model the associated probability distribution as

$$p(B|\psi, L) = \frac{1}{\sigma\sqrt{2\pi}} \exp \left[-\frac{d(B, r(\psi, L))}{2\sigma^2} \right], \quad (3)$$

¹Because we use the roughness parameterization from Burley [2012] (Eq. 11 in the Supplement), perceptual nonlinearity is already partly accounted for (p.15 in Burley [2012]).

where d is the distance function between the renders $r(\psi, L)$ and the observed radiance B . Typically, d is the $L1$ - or $L2$ norm over the difference between the set of views and their corresponding renders for every viewpoint. To compute entropy, we need to normalize $p(\psi|B, L)$ to be a proper probability distribution, since, as a likelihood, it does not integrate to one:

$$\bar{p}(\psi|B, L) = \frac{p(\psi|B, L)}{\sum_{\psi \in \Psi} p(\psi|B, L)}. \quad (4)$$

Finally, we would like the maximum entropy to be one for interpretability. Therefore, we divide entropy by $\log(n)$, the maximum entropy corresponding to a uniform distribution

$$H = -\frac{1}{\log(n)} \sum_{\psi \in \Psi} \bar{p}(\psi|B, L) \log \bar{p}(\psi|B, L). \quad (5)$$

Figure 4 shows three examples of likelihood over the roughness, α , and the specular coefficient K_s , given different lighting conditions and ground-truth parameters (marked by a red dot). The heading H denotes the corresponding entropy. An important observation is that these distributions do not follow a Gaussian distribution, which is the assumed distribution in Laplace’s approximation. This demonstrates the importance of sampling the full parameter space, rather than a local approximation.

3.2 Reflection as a Convolution

To compute entropy, we must evaluate the rendering function r for the entire parameter space and for all viewpoints. For example, evaluating Equation 5 at 16 intervals per parameter ($n = 4096$) for 60 viewpoints requires over 45 minutes with Mitsuba, a highly optimized path tracer. To make our approach practical, we approximate the rendering equation using spherical harmonics, based on the signal processing framework for inverse rendering by Ramamoorthi and Hanrahan [2001]. Ramamoorthi and Hanrahan show that the reflection function can be approximated with a spherical convolution of the BRDF over the incoming radiance. Assuming an isotropic microfacet Torrance and Sparrow [1967] BRDF, combined with a Lambertian term, they derive the following approximation for outgoing radiance B at point p in the view direction ω_o

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

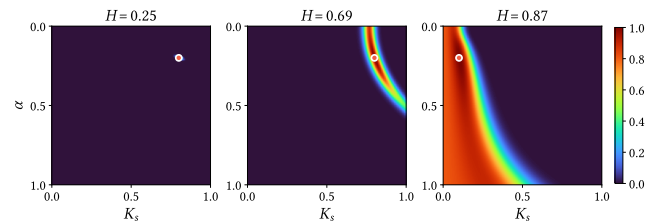


Fig. 4. Examples of the likelihood and entropy H for several parameter/lighting combinations. Left: An ideal situation, where the lighting is a dirac delta. Center: A situation where the lighting is too low-frequency to recover a good α value. Right: A situation where the incoming radiance is ideal, but the specular component is too low to get a proper recovery for α .

where K_d and K_s are diffuse and specular terms; $E(p)$ is the irradiance integrated over the hemisphere; $F(\theta_o)$ is a simplified Fresnel term, which only depends on the outgoing direction; $L(p)$ is the incoming radiance; and S_α is a filter parametrized by the normal distribution width, α . This filter is derived from the normal distribution function of the surface. The $*$ operator represents convolution. We provide a derivation of this approximation in the Supplement and detail there how to integrate the Shadowing and Masking terms [Pharr et al. 2023] into the convolution perspective, which improves accuracy for high-roughness materials. The microfacet model that underlies this reflection model is still used in state-of-the-art inverse rendering research [Hasselgren et al. 2022; Munkberg et al. 2022; Sun et al. 2023] using the parameters from the Disney Principled BRDF [Burley 2012]. In the Supplement, we show how to map the Principled BRDF parameters to our model and in our experiments, we validate that the model performs on par with a state-of-the-art path tracer for inverse rendering.

3.3 Frequency Domain Analysis

Spherical harmonics-based rendering already significantly accelerates the evaluation of our uncertainty metric. We propose to go one step further for uncertainty estimation. The implementation of the rendering equation described in Equation 6 computes convolution using spherical harmonics, but returns to the angular domain to evaluate radiance samples of B . Instead, we propose to evaluate the error function d within the frequency domain, by first transforming the outgoing radiance B to the frequency domain. This yields a significant speedup, as we only need to transform the outgoing radiance to the frequency domain once, rather than computing the inverse transform for every parameter combination. Moreover, we can evaluate the loss in the low-dimensional power spectrum, which allows us to evaluate all parameter combinations in parallel. The resulting speedup is of multiple orders of magnitude. This comes at the cost of some approximations, which we detail in this section. We show how to transform the outgoing radiance to the frequency domain in the next section.

The first required approximation is to omit the Fresnel, shadowing, and masking terms (detailed in the Supplement). These terms are applied as a pointwise multiplication in the angular domain, which is difficult to compute in the frequency domain. In Section 4.2, we investigate the impact of this approximation and find that it is acceptable for uncertainty estimation. We can now rewrite Equation 6 fully in the frequency domain, where $L_{\ell m}$ denotes the spherical harmonic coefficients for the incoming radiance and $B_{\ell m}^\psi$ the coefficients for the predicted outgoing radiance, given parameters $\psi = (K_d, K_s, \alpha)$

$$B_{\ell m}^\psi = \begin{cases} \pi^{-\frac{1}{2}} K_d E(p) + K_s L_{00}, & \text{if } \ell = 0, \\ K_s e^{-(\alpha\ell)^2} L_{\ell m}, & \text{otherwise.} \end{cases} \quad (7)$$

From this expression, we make two observations: (1) We cannot recover the ratio between K_d and K_s from degree $\ell = 0$ alone, as we have two unknowns and only one coefficient. (2) α has no effect on degree $\ell = 0$. It follows that we can only recover K_s and α from degrees $\ell > 0$. Once K_s is known, we can estimate K_d from degree $\ell = 0$. Put simply, once we know the contribution of the specular

component, we can recover the diffuse component exactly, up to measurement error. This leads to an important insight: **we can reduce uncertainty evaluation to the specular component parameters K_s and α** . This simplification comes naturally in the frequency domain, because we can limit our analysis to degrees $\ell > 0$, but would be challenging in the angular domain.

3.3.1 Power spectrum. We propose to further simplify our analysis using the power spectrum of the spherical harmonics coefficients of the incoming and outgoing radiance. The power spectrum of the incoming radiance L can be computed per degree from the spherical harmonic coefficients

$$S_L(\ell) = \sum_{m=-\ell}^{\ell} L_{\ell m}^2 \quad (8)$$

The power spectrum is invariant to rotations of the coordinate system, which, in our context, makes it invariant to slight perturbations of the normals at each point. Using Equation 7, we can compute the power spectrum for a hypothetical outgoing radiance for degrees $\ell > 0$, given a set of parameters ψ as

$$\begin{aligned} S_{B^\psi}(\ell) &= \sum_{m=-\ell}^{\ell} (K_s e^{-(\alpha\ell)^2} L_{\ell m})^2 \\ &= K_s^2 e^{-2(\alpha\ell)^2} \sum_{m=-\ell}^{\ell} L_{\ell m}^2 = K_s^2 e^{-2(\alpha\ell)^2} S_L(\ell) \end{aligned} \quad (9)$$

We can now express the error function d , with $\psi = (K_s, \alpha)$, based on the error between the observed power spectrum and predicted power spectrum

$$d(B, r(\psi, L)) = \sum_{\ell=1}^{\ell^*} (S_B(\ell) - K_s^2 e^{-2(\alpha\ell)^2} S_L(\ell))^2 \quad (11)$$

$\ell = 0$ is left out, because it does not contribute to uncertainty. Using the power spectrum reduces the number of summands in this equation from $(\ell^* + 1)^2 - 1$ to ℓ^* , because we do not need to evaluate Equation 7 for all coefficients. This lets us parallelize the evaluation of the objective for many parameter combinations and locations on the surface, making evaluation near instantaneous (see Section 4).

The combination of entropy with frequency analysis allows us to derive similar observations to Ramamoorthi and Hanrahan [2001], as summarized in the Supplement, Section D.1. For example, uncertainty is high when the lighting contains only low frequencies, because many roughness values would give similar results (Figure 4, center). Low albedo also causes high uncertainty, because less light is reflected, amplifying the relative effect of measurement noise (Figure 4, right).

3.4 Fitting Spherical Harmonic Coefficients

To analyze the reflection function in the frequency domain, we need a spherical harmonics transform for the incoming and outgoing radiance samples. However, the samples of outgoing radiance are captured in a multi-view capture setup, meaning they are sparse and irregularly distributed on the upper hemisphere. The typical approach would be to fit spherical harmonic coefficients using least-squares (see Supplement). However, we find that naive least-squares

fitting is insufficient to recover BRDF parameters in the frequency domain (detailed in the Supplement). We address this by adding a weighted $L2$ norm on the spherical harmonic coefficients in the least-squares formulation

$$(\mathbf{Y}^T \mathbf{Y} + \lambda \mathbf{W}) \mathbf{c} = \mathbf{Y}^T \mathbf{f}, \quad (12)$$

where \mathbf{f} is a vector of n discrete samples of the signal, \mathbf{Y} is a matrix of size $n \times (\ell^* + 1)^2$ containing the spherical harmonics evaluated at the sampling points, and \mathbf{c} is a vector of the $(\ell^* + 1)^2$ coefficients we want to find. \mathbf{W} is a diagonal weight matrix. If constant, this weight matrix leads to poor fitting as the regularization is too strong on the lower spherical-harmonics degrees. We set the weight equal to e^ℓ , increasing the strength of regularization for higher spherical harmonics degrees. This is supported by the observation that many natural images have a power spectrum with exponential decay [Fleming et al. 2003]. Intuitively, this regularizer encourages filling unknown regions with low-frequency information. We weigh the samples based on their elevation angle θ with $|\cos(\theta)|$, since observations at grazing angles are more likely to present measurement error.

Aliasing and ringing. The incoming signal should be bandlimited to roughly $v^{\frac{1}{2}}$ to avoid aliasing, where v is the number of input views. We ensure this is the case for the incoming radiance by filtering the input environment map with the BRDF function for $\alpha' = v^{-\frac{1}{2}}$. Our uncertainty and parameter estimation is therefore approximate for $\alpha < \alpha'$. A more in-depth analysis is provided in the Supplement. In practice, we find that the loss in accuracy is acceptable in comparison to the gains in efficiency (subsection 4.2). Ringing, which is often an issue for spherical-harmonics based approaches [Sloan 2017], does not impact our method, because we do not represent the BRDF directly as spherical harmonics, but fit a parametric BRDF model to the first ℓ^* degrees in the power spectrum. The pre-filtering and regularization of spherical-harmonics fitting ensure that ringing does not impact how we handle incoming and outgoing radiance.

4 Experiments

Implementation. We implement our method in PyTorch. Mitsuba 3 [Jakob et al. 2022] is used to render results and as a reference method. Each timing result is computed on an NVIDIA RTX4090 GPU and 32-core Intel Core i9-14900KF CPU. Code for our experiments is available at https://github.com/rubenwiersma/svbrdf_uncertainty.

Datasets and tasks. The experiments are performed on Stanford ORB [Kuang et al. 2023] and on synthetic scenes. Stanford ORB is an in-the-wild benchmark that contains fourteen objects, each captured three times in different scenes (lighting environments), selected from a total of seven scenes. The lighting environment is captured through a chrome ball and stored as a lat-long environment map. Each object is also scanned in a separate stage, providing high-quality geometry. For our evaluation, we focus on the (SV)BRDF recovery step and use the provided geometry.

For the synthetic benchmark, we gathered a selection of fifteen objects with spatially varying BRDF textures for base color, roughness and metallicity. The benchmark is rendered with four environment

Table 1. Benchmark Comparison: Stanford ORB [Kuang et al. 2023] Relighting. Each method has access to the ground-truth lighting and geometry. See supplement for additional methods and different acquisition conditions.

	PSNR-H \uparrow	PSNR-L \uparrow	SSIM \uparrow	LPIPS \downarrow	Time
NVDiffRec [Munkberg et al. 2022]	24.319	31.492	0.969	0.036	142.14s
Mitsuba [Jakob et al. 2022]	26.601	34.195	0.977	0.032	52.49s
SH Model	26.525	33.796	0.977	0.029	5.07s
SH Model - Power Spectrum	24.494	31.215	0.971	0.035	1.78s
SH Model init + Mitsuba (1 epoch)	26.918	34.386	0.978	0.031	16.07s

Table 2. Synthetic benchmark: Mitsuba vs SH Model - PSNR to ground-truth and correlation between SH-model entropy and error for Mitsuba.

	Photo Studio	Museum	Overcast	Rural Road	Avg.	Time
Mitsuba	17.07	17.69	16.35	17.32	17.11	74.54s
Mitsuba global sharing	18.07	18.38	17.96	17.27	17.92	77.81s
SH Model	15.23	15.63	15.60	16.55	15.75	3.09s
SH Model global sharing	17.34	18.10	18.11	18.95	18.13	3.15s
Error-Entropy correlation	0.27	0.20	0.19	0.24	0.22	

maps representing varying challenges: two indoor scenes, one outdoor scene with a clear sky and sun, and one overcast outdoor scene. Because ground-truth material textures are available for the synthetic scenes, we can quantitatively evaluate our optimization results and validate uncertainty directly on the optimized parameters. The objects are rendered in Mitsuba at 512×512 resolution and 256 samples per pixel. Renders and ground-truth views for both Stanford ORB and the synthetic benchmark are included in the Supplement.

4.1 Quality of Inverse Rendering

As a first validation, we are interested in the performance of the spherical-harmonics model (SH model) evaluated in the angular domain, including our proposed additions for shadowing and masking, on the task of inverse rendering. This validates that the SH model is a relevant proxy for more recent and exact models in inverse rendering and lets us evaluate the model without the approximations required for the power spectrum simplification. We will also use the SH model in our applications of entropy, where this model acts as a baseline to compare the applications with. We optimize PBR material textures to match renderings from multiple views (512×512 textures of base color, roughness, metallicity), UV-mapped to the surface. We use a simple mapping between the PBR materials and the parameters used in the SH model, which is described in the Supplement. The parameters are optimized using gradient descent with an Adam optimizer that minimizes an L1 reconstruction loss and a TV-norm regularizer in texture space weighted by 1×10^{-2} .

Our main comparison target is Mitsuba 3 [Jakob et al. 2022], a complete differentiable path tracer. When optimizing with Mitsuba, we sample one ray per pixel in each view and randomly select a batch of rays per optimization step for better convergence. Each primary ray is sampled 32 times to sample reflection directions well. The relighting results on Stanford ORB are shown in Table 1 and visual comparisons in Figure 5. We find that our SH model performs on par with Mitsuba but requires only 1/10 of the time despite little




Fig. 5. Material recovery results on Stanford ORB for Mitsuba (52.49s on average), the SH Model and our proposed global sharing application (5.07s on average).

optimization on our side, as we use pure PyTorch. Our power spectrum variant, which ignores Fresnel, shadowing, and masking, still performs on par with NVDiffRec, with a 30-fold speedup compared to Mitsuba. On the synthetic benchmark, we compare the optimized textures to the ground-truth textures used to render the dataset. The results are presented in Table 2 and visual comparisons in Figure 6 and the Supplement. The SH Model achieves 1.36dB lower PSNR than Mitsuba on average. The results for Mitsuba are the best-case scenario, given that the dataset was rendered with Mitsuba, using the same BRDF model used for material estimation. Moreover, the dataset contains particularly challenging objects for our method (e.g., with significant interreflections). Split out over base color (Mitsuba: 20.54dB, SH: 18.77dB), roughness (Mitsuba: 17.84dB, SH: 14.71dB) and metallicity (Mitsuba: 12.94dB, SH: 13.78dB), we find that roughness is particularly challenging for the SH model. This is expected, because the effect of roughness is approximated by the SH model. Given that our goal is to estimate uncertainty, we find the performance of the SH-model on this stress-test acceptable.




Fig. 6. Material recovery results on the synthetic dataset for Mitsuba, the SH Model, and the SH Model with global sharing enabled by entropy. We see that global sharing helps smooth and improve the results obtained with the SH Model without blurring the maps.

In subsection 5.2, we discuss how to improve these results using information sharing (‘global sharing’ rows in Table 2).

4.2 Quality of Approximate Entropy

Our method simplifies the reflection equation, so it can be evaluated in the power spectrum of spherical harmonics. This yields significant performance gains, but involves approximations. To study the validity of these approximations, we compare the entropy computed with our method to the entropy computed using Mitsuba and a variant of our method that evaluates the reconstruction error in the angular instead of the frequency domain. This allows us to evaluate the signal-processing framework for inverse rendering without additional approximations. While using a path-tracing renderer like Mitsuba to compute entropy is impractical for most use cases, it provides a close proxy for ‘ground-truth’ uncertainty.

We compute entropy for every scene in Stanford ORB with each approach and use the Pearson correlation coefficient ρ between the entropy maps as a quality measure. We do not compare exact




Fig. 7. Entropy estimation results on Stanford ORB for Mitsuba (5m50s on average) and our method (0.0005s on average). This ensures our approximation still produces entropy close to that computed, slower, with Mitsuba

entropy values, as both methods compute the loss in a different space, making it difficult to exactly compare the resulting probability distributions. For Mitsuba, we discretize the parameter space into $8 \times 8 \times 8$ parameter combinations (roughness, metallicity, base color) and 8×8 for our method (specularity and roughness, see explanation below Equation 7). We find a high correlation between entropy computed with Mitsuba and both our method evaluated in the angular domain $\rho = 0.94$ and in the power spectrum $\rho = 0.88$ as shown in Table 3. This validates that the entropy we compute with our power spectrum based model is representative of one computed with a physically-based renderer, which confirms that our approximations are permissible for uncertainty. Visual comparisons of entropy in Figure 7 confirm these findings. Further, as our power spectrum approximation is parallelizable, we obtain extreme speedups over both the angular-domain approach (3,320 \times) and Mitsuba (700,000 \times) making uncertainty estimation practical. For a higher (and more accurate) sampling density of $16 \times 16 \times 16$, Mitsuba takes roughly 47 minutes, the angular-domain approach 13.2s, and our power-spectrum method only 0.0006s.

4.3 Entropy and Error

An important part of our motivating argument is that high certainty is predictive of low error. This makes uncertainty useful for various applications, e.g., information sharing across areas, or inpainting of uncertain areas using a diffusion model. We argue that regions with high uncertainty are more likely to converge to an incorrect parameter, simply because an incorrect parameter yields a similar

Table 3. Correlation of entropy computed with Mitsuba and our method.

	Entropy Mitsuba \uparrow	Time
Entropy Mitsuba	1.00 ± 0.00	5m50s
Entropy angular	0.94 ± 0.04	1.66s
Entropy power spectrum	0.88 ± 0.05	0.0005s

error to the correct parameter. However, an uncertain region does not necessarily converge to an incorrect result: the correct parameter choice could still have the maximum likelihood, even by a small margin, or the optimization could converge on the correct choice by chance. We therefore do not expect a perfect correlation between error and entropy. At most, one could expect a correlation of 0.5.

We study the relationship between entropy computed with our power-spectrum model and error in the texture maps estimated with Mitsuba. We intentionally test the relationship between different models to show the generalizability of our approach. In the last row of Table 2 we show the correlation scores, which are 0.22 on average and 0.27 for the photostudio scene. We also show qualitative results in Figure 8. These results are optimized on images rendered under the “rural asphalt road” environment. We observe a match between high entropy- and high error regions. To illustrate the usefulness of entropy, we show a mask of all regions with above-average error and a mask of regions with above-average entropy. Such masks could be used to remove erroneous regions. We show that the entropy-based mask covers nearly all regions of high error, while still being sparse. This demonstrates that entropy can serve as an a-priori guide to reduce error.

5 Applications

To demonstrate use of uncertainty and practicality of our method, we propose three applications of entropy in the context of SVBRDF recovery: guiding capture, sharing information, and inpainting uncertain regions with a diffusion model.

5.1 Guiding Capture

Entropy can be applied during capture as a measure to define the most informative viewpoints. This is challenging, because we do not know what information can be captured from a viewpoint before capturing it. Therefore, we propose an approach to compute the *expected* entropy (i.e., information gain) from adding a new viewpoint: we compute the entropy for hypothetical observations B'_ψ for a parameter combination ψ , given incoming radiance observed from a hypothetical set of viewpoints, L' , sampled from the known lighting. The entropy for a hypothetical observation is weighed by the probability of the parameters responsible for that observation, $\bar{p}(\psi|B, L')$ from Equation 4. The result is expected entropy

$$E[H] = \sum_{\psi \in \Psi} \bar{p}(\psi|B, L') H(B'_\psi, L'). \quad (13)$$

We use our power spectrum approximation to accelerate this computation: we first estimate the spherical harmonics coefficients for L' and compute the hypothetical radiance B' only in the power spectrum, based on $S_{L'}$. A set of viewpoints is then selected with a greedy strategy, which we call *best-view selection*. Starting with




Fig. 8. Comparison of MSE for Mitsuba and entropy computed with the SH Model. We observe that where there is error, entropy tends to be higher. We also show a mask for above-average MSE and entropy. The *Mask Coverage* column shows which regions of the MSE mask are not covered by the entropy mask and should, ideally, be completely black.

v' evenly spaced viewpoints, we compute the expected entropy for V candidate viewpoints and select the viewpoint with the lowest expected entropy. This is repeated until v desired viewpoints are selected. We compare this algorithm with farthest point sampling (FPS) of the V candidate viewpoints as a baseline. Our hypothesis is that the best-view selection yields better performance, given the same number of viewpoints. In our experiment, we use $V = 100$ candidate points and $v' = 10$ starting points. The same starting points are used for the best-view selection and FPS baseline. We evaluate both approaches on the synthetic dataset and compare the PSNR gain and the average entropy over all shapes using the additional views (see Figure 9). We observe a considerable improvement in both the PSNR gain and entropy: the number of views required to reach +1 PSNR gain with the FPS baseline is 60% higher than the number of views the best-view selection requires. We also compare the results with a paired t-test (pairing the same scenes) and see a statistically significant improvement between the best-view selection and the FPS baseline ($p = 1 \times 10^{-8}$).

This application shows the importance of our efficient power spectrum approximation, as it involves a nested entropy computation. Using our method, capture guidance requires roughly one second *per added viewpoint*, while this would take roughly 45 minutes in the angular domain and more than nine days using Mitsuba.

5.2 Sharing Information

Uncertain regions on the surface could be improved with information from more certain regions. Prior works typically use a regularizer, enforcing smoothness of the parameters over the surface (e.g., by adding a TV-norm term [Hasselgren et al. 2022; Munkberg et al. 2022; Sun et al. 2023]). Information is consequently shared locally. However, a smoothness term does not distinguish between certain and uncertain points, blurring regions that should actually be preserved. Given our measure of uncertainty, we can develop a more informed strategy: we would like to share information from

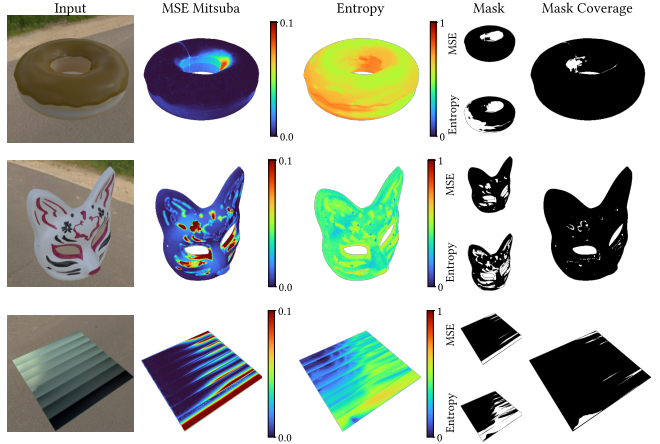


Fig. 9. Plots of PSNR gain and entropy for varying input view counts on the synthetic benchmark with the Museum environment. We observe a significant improvement of our best-view selection (Best) over farthest point sampling (FPS) on a paired t-test ($p = 1 \times 10^{-8}$). Using farthest point sampling requires 60% more views compared to using our entropy-based capture guidance.

the most certain points to less certain points with a similar appearance and allow for sharing to happen with non-neighboring points. We choose a set of certain points, C (local minima in the uncertainty map H). Next, we add a term to the regularizer that enforces other points' parameters to be similar to the selected points:

$$\mathcal{L}_S = \frac{1}{T} \sum_p \sum_{q \in C} w_{pq} \|\psi_p - \psi_q\|^2, \quad (14)$$

where T is the number of discrete points on the surface (in our experiments, the number of texels), ψ_p is the parameter at point p , and w_{pq} is a weight based on the similarity of the materials at points p and q . Since we do not know a priori whether two regions consist of the same material, we use the KL-divergence between the probability distributions P_p and P_q over the possible material parameters as a similarity metric

$$w_{pq} = \exp \left[-\frac{D_{KL}(P_p || P_q)^2}{2\tau^2} \right]. \quad (15)$$



Fig. 10. Inpainting low-certainty results with the RGB \leftrightarrow x diffusion model [Zeng et al. 2024]. We show the result of the optimization, results for inpainting and a full diffusion model prediction on base color. Results are shown with 10/20 steps of complete diffusion after 40/30 steps of inpainting, striking a balance between preserving the information from the optimization and using the diffusion model to inpaint uncertain areas.

The parameter τ determines the selectiveness in finding similar regions. We set τ as a fraction of the average KL-divergence between C and all other points (0.4 in our experiments). The results of this approach are compared with the baseline models using a TV-norm in Table 2. We observe a +1.5 dB PSNR improvement for the SH model and +0.5 dB for Mitsuba. Visual comparisons reflect these results (Figure 6 and Figure 5, global sharing). We only show results for the SH model as the results for Mitsuba with global sharing are visually similar, albeit slightly better.

5.3 Diffusion Model Guidance

We can use the material map estimated with an optimization method, together with our uncertainty map to extract material maps with a diffusion model [Ho et al. 2020]. As a proof of concept, we inpaint regions in the estimated albedo map with high uncertainty using blended diffusion [Avrahami et al. 2022]. For the diffusion model, we use RGB \leftrightarrow x [Zeng et al. 2024], which estimates material maps in image space for shaded images. Following the blended diffusion process, we perform 50 diffusion steps and replace the diffusion result

with the optimized material map for the first few steps, followed by n diffusion steps on the full material map. In Figure 10, we show results on Stanford ORB and the base color parameter for four different variants: no inpainting, inpainting followed by 10 and 20 full diffusion steps, and full diffusion. We observe that the results with 10 full diffusion steps (i.e., more influence of optimization results) strike a good balance between the correct parameter values and the diffusion model result.

5.4 Optimization: Initialization

We show that our fast estimate of BRDF parameters can be used as a high-quality initialization for other inverse rendering approaches. In this experiment, we initialize an optimization with Mitsuba using the textures from the SH Model. This is optimized for one epoch, using only 16 samples per pixel (half the samples we used for the other experiments) and the results are included in Table 1 (last row). This single epoch only takes 11s, resulting in a total of 16s, combined with the initialization. This is over three times faster than the 52s required with random initialization and achieves better performance (+0.3dB PSNR-H).

6 Challenges and conclusion

In summary, we present an uncertainty estimation method for multi-view capture of materials on objects using frequency domain analysis. We use least-squares fitting and a power spectrum approximation to efficiently compute entropy as uncertainty. We validate our uncertainty estimation and demonstrate its benefits through multiple applications, improving the quality of captured assets.

A future challenge for practical use is to enable spherical harmonics decompositions on LDR images. Because LDR signals are typically clamped, they introduce frequencies that are not present in the original signal, conflicting with the frequency analysis. Another direction is extending the framework for situations with unknown lighting, which requires computing entropy over hypothetical lighting setups. We believe frequency analysis for uncertainty estimation provides a powerful and practical tool for understanding and improving object capture approaches and look forward to its applications in future work.

Acknowledgments

We would like to thank Yash Belhe and Georgios Kopanas for insightful early discussions. This work was partially supported by the convergence theme AI, Data and Digitalisation of Erasmus University, the Erasmus medical centre and the Delft University of Technology via the Immersive Tech AI lab.

References

- Miika Aittala, Timo Aila, and Jaakko Lehtinen. 2016. Reflectance modeling by neural texture synthesis. *ACM Trans. Graph.* 35, 4, Article 65 (jul 2016), 13 pages. <https://doi.org/10.1145/2897824.2925917>
- Miika Aittala, Tim Weyrich, and Jaakko Lehtinen. 2013. Practical SVBRDF capture in the frequency domain. *ACM Trans. Graph.* 32, 4 (2013), 110–1.
- Miika Aittala, Tim Weyrich, Jaakko Lehtinen, et al. 2015. Two-shot SVBRDF capture for stationary materials. *ACM Trans. Graph.* 34, 4 (2015), 110–1.
- Omri Avrahami, Dani Lischinski, and Ohad Fried. 2022. Blended Diffusion for Text-Driven Editing of Natural Images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 18208–18218.

- Mahdi M. Bagher, Cyril Soler, Kartic Subr, Laurent Belcour, and Nicolas Holzschuch. 2013. Interactive Rendering of Acquired Materials on Dynamic Geometry Using Frequency Analysis. *IEEE Trans. Vis. Comput. Graph.* 19, 5 (2013), 749–761. <https://doi.org/10.1109/TVCG.2012.298>
- Mark Boss, Varun Jampani, Kihwan Kim, Hendrik Lensch, and Jan Kautz. 2020. Two-shot spatially-varying brdf and shape estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 3982–3991.
- Brent Burley. 2012. Physically-based shading at Disney. In *Acm Siggraph*, Vol. 2012. vol. 2012, 1–7.
- Valentin Deschaintre, Miika Aittala, Fredo Durand, George Drettakis, and Adrien Bousseau. 2018. Single-image svbrdf capture with a rendering-aware deep network. *ACM Transactions on Graphics (ToG)* 37, 4 (2018), 1–15.
- Valentin Deschaintre, Miika Aittala, Frédo Durand, George Drettakis, and Adrien Bousseau. 2019. Flexible svbrdf capture with a multi-image deep network. In *Computer graphics forum*, Vol. 38. Wiley Online Library, 1–13.
- Valentin Deschaintre, George Drettakis, and Adrien Bousseau. 2020. Guided Fine-Tuning for Large-Scale Material Transfer. *Computer Graphics Forum (Proceedings of the Eurographics Symposium on Rendering)* 39, 4 (2020). <http://www.sop.inria.fr/revues/Basilic/2020/DBB20>
- Valentin Deschaintre, Yiming Lin, and Abhijeet Ghosh. 2021. Deep polarization imaging for 3D shape and SVBRDF acquisition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 15567–15576.
- Yue Dong, Guojun Chen, Pieter Peers, Jiawan Zhang, and Xin Tong. 2014. Appearance-from-motion: Recovering spatially varying surface reflectance under unknown lighting. *ACM Transactions on Graphics (TOG)* 33, 6 (2014), 1–12.
- Jonathan Dupuy and Wenzel Jakob. 2018. An Adaptive Parameterization for Efficient Material Acquisition and Rendering. *Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 37, 6 (Nov. 2018), 274:1–274:18. <https://doi.org/10.1145/3272127.3275059>
- Frédo Durand, Nicolas Holzschuch, Cyril Soler, Eric Chan, and François X. Sillion. 2005. A frequency analysis of light transport. *ACM Trans. Graph.* 24, 3 (jul 2005), 1115–1126. <https://doi.org/10.1145/1073204.1073320>
- Chongrui Fan, Yiming Lin, and Abhijeet Ghosh. 2023. Deep Shape and SVBRDF Estimation using Smartphone Multi-lens Imaging. *Computer Graphics Forum* (2023). <https://doi.org/10.1111/cgf.14972>
- Roland W. Fleming, Ron O. Dror, and Edward H. Adelson. 2003. Real-world illumination and the perception of surface reflectance properties. *Journal of Vision* 3, 5 (07 2003), 3–3. <https://doi.org/10.1167/3.5.3> arXiv:<https://arxiv.org/abs/10.1167/3.5.3>
- Yarin Gal and Zoubin Ghahramani. 2016. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*. PMLR, 1050–1059.
- Abhijeet Ghosh, Shruthi Achutha, Wolfgang Heidrich, and Matthew O’Toole. 2007. BRDF Acquisition with Basis Illumination. In *2007 IEEE 11th International Conference on Computer Vision*. IEEE, Rio de Janeiro, Brazil, 1–8. <https://doi.org/10.1109/ICCV.2007.4408935>
- Lily Goli, Cody Reading, Silvia Sellán, Alec Jacobson, and Andrea Tagliasacchi. 2024. Bayes’ Rays: Uncertainty Quantification in Neural Radiance Fields. *CVPR* (2024).
- Jie Guo, Shuichang Lai, Chengzhi Tao, Yuelong Cai, Lei Wang, Yanwen Guo, and Ling-Qi Yan. 2021. Highlight-aware two-stream network for single-image SVBRDF acquisition. *ACM Trans. Graph.* 40, 4, Article 123 (jul 2021), 14 pages. <https://doi.org/10.1145/3450626.3459854>
- 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.
- Philipp Henzler, Valentin Deschaintre, Niloy J Mitra, and Tobias Ritschel. 2021. Generative Modelling of BRDF Textures from Flash Images. *ACM Trans Graph (Proc. SIGGRAPH Asia)* 40, 6 (2021).
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising Diffusion Probabilistic Models. *arXiv preprint arxiv:2006.11239* (2020).
- Inseung Hwang, Daniel S. Jeon, Adolfo Muñoz, Diego Gutierrez, Xin Tong, and Min H. Kim. 2022. Sparse Ellipsometry: Portable Acquisition of Polarimetric SVBRDF and Shape with Unstructured Flash Photography. *ACM Transactions on Graphics (Proc. SIGGRAPH 2022)* 41, 4 (2022).
- 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>.
- Zhengfei Kuang, Yunzhi Zhang, Hong-Xing Yu, Samir Agarwala, Shangzhe Wu, and Jianjun Wu. 2023. Stanford-ORB: A Real-World 3D Object Inverse Rendering Benchmark. arXiv:2310.16044 [cs.CV]
- Hendrik Lensch, Jan Kautz, Michael Goesele, Wolfgang Heidrich, and Hans-Peter Seidel. 2003a. Image-based Reconstruction of Spatial Appearance and Geometric Detail. *ACM Transactions on Graphics (TOG)* 22 (04 2003), 234–257. <https://doi.org/10.1145/636886.636891>
- Hendrik P.A. Lensch, Jochen Lang, Asla M. Sá, and Hans-Peter Seidel. 2003b. Planned Sampling of Spatially Varying BRDFs. *Computer Graphics Forum* 22, 3 (2003), 473–482. <https://doi.org/10.1111/1467-8659.00695> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/1467-8659.00695>
- Xiao Li, Yue Dong, Pieter Peers, and Xin Tong. 2017. Modeling surface appearance from a single photograph using self-augmented convolutional neural networks. *ACM Transactions on Graphics (ToG)* 36, 4 (2017), 1–11.
- Zhengqin Li, Zexiang Xu, Ravi Ramamoorthi, Kalyan Sunkavalli, and Manmohan Chandraker. 2018. Learning to reconstruct shape and spatially-varying reflectance from a single image. *ACM Transactions on Graphics (TOG)* 37, 6 (2018), 1–11.
- Guillaume Loubet, Nicolas Holzschuch, and Wenzel Jakob. 2019. Reparameterizing Discontinuous Integrands for Differentiable Rendering. *Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 38, 6 (Dec. 2019). <https://doi.org/10.1145/3355089.3356510>
- Xuejiao Luo, Leonardo Scandolo, Adrien Bousseau, and Elmar Eisemann. 2024. Single-Image SVBRDF Estimation with Learned Gradient Descent. *Computer Graphics Forum* 43 (2024), i–iii. <https://doi.org/10.1111/cgf.15018>
- Linjie Lyu, Ayush Tewari, Marc Habermann, Shunsuke Saito, Michael Zollhöfer, Thomas Leimkühler, and Christian Theobalt. 2023. Diffusion Posterior Illumination for Ambiguity-aware Inverse Rendering. *ACM Transactions on Graphics* 42, 6 (2023).
- Rosalie Martin, Arthur Roullier, Romain Rouffet, Adrien Kaiser, and Tamy Boubekeur. 2022. MaterIA: Single Image High-Resolution Material Capture in the Wild. *Computer Graphics Forum (Proc. EUROGRAPHICS 2022)* to appear, to appear (2022), to appear.
- Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. 2020. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In *ECCV*.
- 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.
- Giljoo Nam, Joo Ho Lee, Diego Gutierrez, and Min H. Kim. 2018. Practical SVBRDF Acquisition of 3D Objects with Unstructured Flash Photography. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia 2018)* 37, 6 (2018), 267:1–12. <https://doi.org/10.1145/3272127.3275017>
- Merlin Nimier-David, Zhao Dong, Wenzel Jakob, and Anton Kaplanyan. 2021. Material and Lighting Reconstruction for Complex Indoor Scenes with Texture-space Differentiable Rendering. In *Eurographics Symposium on Rendering - DL-only Track*, Adrien Bousseau and Morgan McGuire (Eds.). The Eurographics Association. <https://doi.org/10.2312/sr.20211292>
- 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>
- C. Rodriguez-Pardo, H. Dominguez-Elvira, D. Pascual-Hernandez, and E. Garces. 2023. UMat: Uncertainty-Aware Single Image High Resolution Material Capture. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, Los Alamitos, CA, USA, 5764–5774. <https://doi.org/10.1109/CVPR52729.2023.00558>
- Viraj Shah, Svetlana Lazebnik, and Julien Philip. 2023. JoIN: Joint GANs Inversion for Intrinsic Image Decomposition. *arXiv* (2023).
- Claude Elwood Shannon. 1948. A mathematical theory of communication. *The Bell system technical journal* 27, 3 (1948), 379–423.
- Liang Shi, Beichen Li, Miloš Hašan, Kalyan Sunkavalli, Tamy Boubekeur, Radomir Mech, and Wojciech Matusik. 2020. MATch: Differentiable Material Graphs for Procedural Material Capture. *ACM Trans. Graph.* 39, 6 (Dec. 2020), 1–15.
- Peter-Pike Sloan. 2017. Deringing spherical harmonics. In *SIGGRAPH Asia 2017 Technical Briefs, Bangkok, Thailand, November 27 - 30, 2017*, Diego Gutierrez and Hui Huang (Eds.). ACM, 11:1–11:4. <https://doi.org/10.1145/3145749.3149438>
- Cheng Sun, Guangyan Cai, Zhengqin Li, Kai Yan, Cheng Zhang, Carl Marshall, Jia-Bin Huang, Shuang Zhao, and Zhao Dong. 2023. Neural-PBR Reconstruction of Shape, Material, and Illumination. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 18046–18056.
- K. E. Torrance and E. M. Sparrow. 1967. Theory for Off-Specular Reflection From Roughened Surfaces*. *J. Opt. Soc. Am.* 57, 9 (Sep 1967), 1105–1114. <https://doi.org/10.1364/JOSA.57.001105>
- Giuseppe Vecchio, Rosalie Martin, Arthur Roullier, Adrien Kaiser, Romain Rouffet, Valentin Deschaintre, and Tamy Boubekeur. 2023. ControlMat: Controlled Generative Approach to Material Capture. *arXiv preprint arXiv:2309.01700* (2023).
- Delio Vicini, Sébastien Speierer, and Wenzel Jakob. 2022. Differentiable Signed Distance Function Rendering. *Transactions on Graphics (Proceedings of SIGGRAPH)* 41, 4 (July 2022), 125:1–125:18. <https://doi.org/10.1145/3528223.3530139>
- Zexiang Xu, Jannik Boll Nielsen, Jiyang Yu, Henrik Wann Jensen, and Ravi Ramamoorthi. 2016. Minimal BRDF sampling for two-shot near-field reflectance acquisition. *ACM*

- Transactions on Graphics (TOG)* 35, 6 (2016), 1–12.
- Zheng Zeng, Valentin Deschaintre, Iliyan Georgiev, Yannick Hold-Geoffroy, Yiwei Hu, Fujun Luan, Ling-Qi Yan, and Miloš Hašan. 2024. RGB \leftrightarrow X: Image decomposition and synthesis using material- and lighting-aware diffusion models. In *ACM SIGGRAPH 2024 Conference Papers* (Denver, CO, USA) (*SIGGRAPH '24*). Association for Computing Machinery, New York, NY, USA, Article 75, 11 pages. <https://doi.org/10.1145/3641519.3657445>
- Jianyi Zhang, Ruiyi Zhang, Lawrence Carin, and Changyou Chen. 2020. Stochastic Particle-Optimization Sampling and the Non-Asymptotic Convergence Theory. In *The 23rd International Conference on Artificial Intelligence and Statistics, AISTATS 2020, 26-28 August 2020, Online [Palermo, Sicily, Italy]* (*Proceedings of Machine Learning Research, Vol. 108*), Silvia Chiappa and Roberto Calandra (Eds.). PMLR, 1877–1887. <http://proceedings.mlr.press/v108/zhang20d.html>
- Xilong Zhou, Miloš Hašan, Valentin Deschaintre, Paul Guerrero, Kalyan Sunkavalli, and Nima Khademi Kalantari. 2023. A Semi-Procedural Convolutional Material Prior. In *Computer Graphics Forum*. Wiley Online Library.
- Xilong Zhou and Nima Khademi Kalantari. 2021. Adversarial Single-Image SVBRDF Estimation with Hybrid Training. *Computer Graphics Forum* (2021).
- Xilong Zhou and Nima Khademi Kalantari. 2022. Look-Ahead Training with Learned Reflectance Loss for Single-Image SVBRDF Estimation. *ACM Transactions on Graphics* 41, 6 (12 2022). <https://doi.org/10.1145/3550454.3555495>
- Zhiqian Zhou, Cheng Zhang, Zhao Dong, Carl Marshall, and Shuang Zhao. 2024. Estimating Uncertainty in Appearance Acquisition. (2024).