

# Spherical Harmonic Exponentials for Efficient Glossy Reflections

A. Silvenoinen<sup>1</sup> P.-P. Sloan<sup>1</sup> M. Iwanicki<sup>1</sup> D. Nowrouzezahrai<sup>2,3,4</sup>

<sup>1</sup>Activision Publishing, Inc. USA

<sup>2</sup>McGill University, Canada

<sup>3</sup>Mila – Quebec AI Institute

<sup>4</sup>CIFAR AI Chair

Our method

Split-sum



**Figure 1:** A production scene from a recently shipped AAA-game (© Activision Publishing, Inc.) Indirect specular illumination for *every material* rendered with our method compared to a mip-chain of prefiltered environment maps using the split-sum method [Kar13]. Our method requires  $200\times$  less memory, enabling alias-free reconstruction over continuous roughness and angular domains using only 33 coefficients per-color channel at roughly equal time (9.95ms vs. 9.85ms). Moreover, it avoids systemic errors caused by the split-sum approximation, e.g., hallucinated and implausible white (orange inset) and green (green inset) reflections, while our method faithfully captures such effects.

## Abstract

We propose a high-performance and compact method for computing glossy specular reflections. Commonly-used prefiltered environment maps have large storage requirements and high error due to constrained treatment of view-dependence. We propose a factorized spherical harmonic exponential representation that exploits new observations of the benefits of log-space reconstruction for reflectance. Our method is compact, properly accounts for view-dependent reflections, and is more accurate than the state-of-the-industry solutions. We achieve higher quality results with an order of magnitude less memory, all with efficient and alias-free reconstruction of glossy reflections from environment lights and continuously-varying material roughness.

## CCS Concepts

• Computing methodologies → Rendering; Reflectance modeling;

## 1. Introduction

Image-based lighting (IBL) is a widely used technique for rendering glossy and specular reflections in interactive applications, such as

video games and architectural visualization. In contrast to interactive ray tracing, which requires high-end hardware support, IBL scales across a wide array of target platforms including mobile devices and virtual reality (VR) headsets. The most commonly used IBL method in industry is the “split-sum” approximation [Kar13; Laz13], which uses a precomputed mip-chain of filtered environment maps to approximate the specular component of light transport at runtime. This approach, however, suffers from two key drawbacks: large memory requirements and approximation error caused primarily by its limited treatment of view dependence.

We present a new, factorized spherical harmonics representation for glossy environment maps that covers 75% of the linear roughness range (with an underlying Cook-Torrance/GGX BRDF). We demonstrate that our method significantly reduces reconstruction error while requiring an order of magnitude less memory, enabling alias-free reconstruction of specular lighting from environment maps at equal performance to “split-sum” methods.

As a representative example, a typical large-scale AAA production scene utilizes  $8 \times 10^3$  precomputed reflection probes: here, even with reduced resolution and GPU-friendly block compression, storing this data would require roughly 700MB of runtime memory, exceeding the available budget on modern target platforms such as PC, Xbox One, Xbox Series X/S, PS4, PS5, and mobile devices; in contrast, our representation requires *only 3MB storage* with an uncompressed 32-bit floating point format. The added memory efficiency of our representation allows us to store all reflection probes in memory, facilitating a scalable level-of-detail mechanism that significantly improves visual quality across a wide range of platforms.

Concretely, our technical contributions are:

- a factorized shading model that accurately reconstructs reflected radiance from environment lights and for continuously-varying roughness, i.e., supporting materials across much of the large perceptual roughness range,
- a spherical harmonics-based compression strategy that leverages band-limiting due to the log transformation,
- a simple, efficient point-wise exponentiation compared to earlier spherical harmonics exponentiation methods [RWS\*06], and
- an efficient and stable log-space fitting strategy that optimizes a proxy of relative MSE in linear space, reducing the sensitivity to outliers that the human visual system is sensitive to.

## 2. Previous Work

Work in interactive, precomputation-based appearance and shading spans several decades. We discuss the most relevant prior works and refer readers to more comprehensive surveys of appearance modeling [DRS08] and precomputed radiance transport [Ram09].

### 2.1. Prefiltered Environment Maps

Kautz et al. [KVHS00] overview prefiltered environment map methods [Gre86; HS99] that treat radially symmetric BRDFs with limited view-dependent effects. In this setting, environment illumination is parametrized by the perfect reflection direction and prefiltering assumes that the viewing and normal directions are aligned. This reduces the dimensionality of the problem, as well as memory cost, but

at the expense of quality and flexibility. At runtime, specular lighting can be computed using a simple lookup. An important extension of this simple approach uses a more elaborate lookup heuristic based on the reflectance profile shape, and averages several lookups to obtain higher quality results with modest added overhead [KM00].

Recent prefiltering-based methods (e.g., [SBN15]) extend into broader classes of BRDFs – allowing for a wider variety of materials to be represented – but often at a cost several orders of magnitude higher in memory than our approach, and with lower performance.

In contrast, our method treats view-dependence explicitly and leverages new and existing observations of log-space fitting to significantly reduce storage requirements whilst retaining high accuracy.

Ramamoorthi and Hanrahan [RH02] drop the radial symmetric constraint with a precomputation-based method for specular reflection under environmental illumination that supported a broader class of BRDFs. Their representation is based on a cubemap of spherical harmonics coefficients that encode view-dependent preintegrated shading. Similarly, we also preintegrate shading, however, our factorized representation uses an order of magnitude less memory as it does not require large high-dimensional cubemaps. Concurrently, Latta and Kolb [LK02] observed that a homomorphic factorization stored in cube maps (with a similar parameterization as in [RH02]) was better suited to the representation. Motivated by these works, we propose a parameterization that – when combined with our factorization and log-space reconstruction – results in both a more compact representation and one that is fast and stable to fit. Notably, to our knowledge, ours is the first method to leverage the fact that log-space lighting exhibits lower frequency content, and thus it is better suited for compact frequency-based representation.

To our knowledge, the most widely adopted *image-based lighting* technique in high-performance graphics (e.g., video games) is the “split sum” [Kar13; Laz13] (Section 3.2). This method is the most relevant prior art, as it currently strikes the most appropriate design trade-off between runtime performance, storage, flexibility, and accuracy. Specifically, Karis’ method [Kar13] – based on the Cook-Torrance BRDF (Section 3.1) – prefilters environment lights with circularly symmetric normal distribution function and approximates reflection with a two-integral product that allows for dynamic runtime materials. As with prefiltering methods, however, this representation also constrains view variability (i.e., due to the circularly symmetric prefiltering kernel), leading to significant error, particularly as roughness increases. In contrast, our representation takes explicit account of view-dependence and yields much lower error while simultaneously reducing the memory cost, all at equal or better performance (Section 5). We additionally support *continuously-varying* roughness settings, avoiding the memory and reconstruction limitations of roughness discretization-based methodologies.

### 2.2. BRDF Models, Approximation, and Representation

**Analytic Models.** The most pervasive reflectance models in interactive rendering rely on analytic representation for (easily-parallelizable) computation, and less so on (high-dimensional) tabulation/memory operations. Blinn’s [Bli77] seminal phenomenological model of specular reflection has been extended to incorporate Fresnel effects [Sch94], as well as leading to the development of

more expressive models, albeit still phenomenological [LFTG97]. Several extensions of this latter work are notably well-suited to our environmental lighting setting, using approximations with either few isotropic cosine [MLH02] or Gaussian [GKD07] lobes, both naturally suited to efficient hardware implementation.

Many works draw on the functional approximation literature, e.g., applying rational function representations to reduce fitting error to measured reflectance [PSS\*12] or in augmenting the parametric space to account for artist-centric control [Bur12; MHH\*12]. Analytic and parametric models of anisotropic reflectance also have a rich history in computer graphics, e.g., [PF90; War92; KSK10], however we focus on isotropic reflectance under complex lighting.

Perhaps the most common parametric reflectance models are *microfacet models*, first introduced to the graphics community by Cook and Torrance [CT82]. These (originally) analytic models incorporate physically inspired shadowing-masking factors based on statistical models of microgeometry orientations, similarly admitting follow-up models exploring the axes of function approximation [Sch94; LS05] and more accurate shadowing-masking terms [Smi67; Bro80; APS00; BBS02; BSH12; LKYU12].

**Data-driven Models.** Goniorelectometric reflectance measurements [MWL\*99] can be used directly as a tabular reflectance model [MPBM03], leading to significant analyses [Bur12; MHH\*12; BSN16] that expose the diversity of reflectance present in even the isotropic context. Ngan et al.'s [NDM05] seminal exploration of the suitability of least-squares parametric fits to the MERL dataset [MPBM03] set a methodological baseline that spurred the development of more sophisticated parametric models [AP07; MHH\*12; BSH12; LKYU12; BSN16].

Hybrid microfacet models can combine both analytic and tabulated factors [NDM05; AP07; WZT\*08; GHP\*08; DHI\*15; BSN16], spanning an interesting gamut of storage/compute/accuracy trade-offs; here, an additional feature axis relies on the *types* of operations required of the reflectance model by a specific rendering engine. For example, evaluation speed, storage, and compatibility with pre-filtered lighting are more important in interactive settings than the ability to perform efficient importance sampling.

**Factorized Representations.** Factorizing reflectance models strikes novel trade-offs in the performance/accuracy/operation design space, as we alluded to above. Kautz and McCool [KM99] decompose tabulated reflectance matrices according to half-vector parameterizations, yielding a compact rank-1 factorization using multi-term sums of outer products. McCool et al. [MAA01] fit a three-product factored model with two unique factors in log-space, implicitly enforcing non-negativity and stabilizing the fitting process. We also fit in log-space for these reasons, but additionally take advantage of the fact that log-space lighting and reflectance have lower-frequency angular variation and thus better suited to compact frequency-based representation; as such, the parameterization we propose in this work relies on a simple exponentiation at runtime, is more memory efficient, stable and easy to fit.

Non-negative matrix factorization is an alternative method to enforce positivity [LS00] and employed in methods focused on efficient importance sampling [LRR04]. Extensions of such factor-

ization approaches to anisotropic reflectance also exist (e.g., [SM02; AS00]) but are outside the scope of our application.

An important distinguishing feature of our method is that – unlike the entirety of the aforementioned factorized representations – it is a *joint representation* of BRDF-integrated incident illumination, not just of reflectance/BRDF. Thus, many symmetries that are utilized in BRDF approximation methods are unavailable in our setting.

### 2.3. Precomputed Radiance Transfer

Many Precomputed Radiance Transfer (PRT) methods [SKS02; LK03] support relighting with glossy BRDFs, typically by computing (and compressing) 4D transfer operators that map incoming radiance to reflected radiance. Here, spherical harmonics-based approaches provide a compact representation for low-frequency, low-dynamic range environment lights; however, the  $L_2$  fit of SH is famously sensitive to outliers, resulting in ringing artifacts even when targeting bandlimited HDR signals.

Our method avoids these artifacts by fitting in log-space, resulting in both quantitative and visual improvements – and, all this, whilst still benefiting from compact (exponentiated) SH representation. To our knowledge, we are the first to exploit the property that the log transform typically *reduces* the angular frequency of a spherical signal, allowing for a more accurate representation with band-limited SH. We also note that our use of logarithm and exponential operations occur “outside” of the SH-projected space, precluding the need for expensive runtime SH exponentiation operations [RWS\*06]: we never explicitly evaluate an SH logarithm or SH exponential but, rather, treat the log-transform implicitly during fitting and evaluate an exponential point-wise at run time after (standard) SH expansion.

All-frequency PRT [NRH03; LSSS04] represent incoming radiance with tabulated compressed basis-spaces, however, with memory requirements that use in high-performance settings.

### 3. Preliminaries

For a fixed view  $\mathbf{v}$  and normal  $\mathbf{n}$ , reflected radiance from an environment light source  $L$  is given by the reflection equation

$$E(\mathbf{v}, \mathbf{n}, \Theta) = \int_{\Omega} L(\mathbf{l})f(\mathbf{v}, \mathbf{n}, \mathbf{l}, \Theta)\langle \mathbf{l}, \mathbf{n} \rangle d\mu(\mathbf{l}), \quad (1)$$

where  $L(\mathbf{l})$  is incoming radiance from emitters,  $f$  is the BRDF with parameters  $\Theta$  (e.g., roughness), and  $\langle \mathbf{l}, \mathbf{n} \rangle d\mu(\mathbf{l})$  is the differential projected solid angle measure. The integral is over the hemisphere  $\Omega$  about the normal  $\mathbf{n}$ .

Although incident radiance is independent of shading point location, it is not uncommon to loosen this constraint by interpolating from several sources, e.g., captured at different locations in space to enable spatially varying specular reflections.

The computational cost of approximating Equation 1 numerically depends most significantly on the underlying BRDF  $f$ . As such, an important first goal (G1) in real-time applications is to seek out representations for reflected radiance  $E$  that are efficient to evaluate. However, in many practical scenarios, memory bandwidth and storage limitations – not just evaluation speed – pose the primary bottleneck. Another important factor in this cost is the degree to

which the angular frequency variation of the environmental source is approximated, typically resulting in supplemental memory/accuracy trade-offs. Therefore, a second goal (G2) is to minimize the memory usage of the targeted representation, e.g., to enable efficient computation of spatially localized glossy reflections. By adopting a compact approximation, we reduce memory usage by an order of magnitude compared to traditional methods (c.f. Section 5), in turn enabling deployment on lower-end hardware. Although this introduces minor approximation error, our evaluation shows that the resulting visual quality remains high, demonstrating that the trade-off is well-justified.

We will derive a factorized representation (Section 4) for the reflected radiance  $E$  that satisfies these two goals. Beforehand, we discuss the state-of-the-industry “split sum” shading approach (Section 3.2) and its most commonly employed realization, based on the Cook-Torrance BRDF (which we also use; Section 3.1).

### 3.1. Cook-Torrance Microfacet Model

Our model will rely on the physically based Cook-Torrance microfacet BRDF model  $f$  [CT82; WMLT07]. We chose this model due to its ubiquity in interactive applications and to compare to the predominant split-sum approximation applied in these applications; note that there is nothing fundamental to our method that precludes the generalization to other BRDF models.

Concretely, the Cook-Torrance BRDF  $f$  is parameterized by a roughness parameter  $\alpha$  and defined as

$$f(\mathbf{v}, \mathbf{n}, \mathbf{l}, \alpha) = F(F_0, \mathbf{v}, \mathbf{h})D(\mathbf{h}, \alpha)V(\mathbf{v}, \mathbf{n}, \mathbf{l}), \quad (2)$$

where  $F$  is a Fresnel function that accounts for increased reflection at grazing angles (parameterized by specular color  $F_0$ ), and  $D$  is a normal distribution function that models the fraction of microfacets that face the viewer. Note that  $F$  and  $D$  are defined in terms of the half-vector (or microfacet normal)  $\mathbf{h}$  instead of the geometric normal  $\mathbf{n}$ . The microfacet normal  $\mathbf{h}$  is the halfway vector between view  $\mathbf{v}$  and light  $\mathbf{l}$  directions, i.e.,  $\mathbf{h} = (\mathbf{v} + \mathbf{l}) / \|\mathbf{v} + \mathbf{l}\|$ . Finally, the normalized visibility term  $V$  is given by

$$V(\mathbf{v}, \mathbf{n}, \mathbf{l}) = G(\mathbf{v}, \mathbf{n}, \mathbf{l}) / (4\langle \mathbf{n}, \mathbf{v} \rangle \langle \mathbf{n}, \mathbf{l} \rangle), \quad (3)$$

where  $G$  is a geometric attenuation term accounting for microfacet shadowing and masking effects [Hei14].

We choose to apply the Schlick Fresnel approximation [Sch94]

$$F(F_0, \mathbf{v}, \mathbf{h}) \approx F_0 + (1 - F_0)(1 - \langle \mathbf{v}, \mathbf{h} \rangle)^5 \quad (4)$$

which decouples the Fresnel term  $F$  from the BRDF  $f$ . Now, the reflection equation 1 simplifies to the sum of base reflectance  $E_0$  and the Fresnel tail reflectance  $E_1$  with

$$E(\mathbf{v}, \mathbf{n}, \alpha) \approx F_0 E_0(\mathbf{v}, \mathbf{n}, \alpha) + (1 - F_0) E_1(\mathbf{v}, \mathbf{n}, \alpha), \quad (5)$$

with base reflectance  $E_0$  and a simplified BRDF  $f_{DV}$  given by

$$E_0(\mathbf{v}, \mathbf{n}, \alpha) = \int_{\Omega} L(\mathbf{l}) f_{DV}(\mathbf{v}, \mathbf{n}, \mathbf{l}, \alpha) \langle \mathbf{l}, \mathbf{n} \rangle d\mu(\mathbf{l}), \quad (6)$$

and Fresnel tail reflectance  $E_1$  given by

$$E_1(\mathbf{v}, \mathbf{n}, \alpha) = \int_{\Omega} L(\mathbf{l}) f_{DV}(\mathbf{v}, \mathbf{n}, \mathbf{l}, \alpha) (1 - \langle \mathbf{v}, \mathbf{h} \rangle)^5 \langle \mathbf{l}, \mathbf{n} \rangle d\mu(\mathbf{l}). \quad (7)$$

Finally the simplified, or “Fresnel-free”, BRDF  $f_{DV}$  is given by

$$f_{DV}(\mathbf{v}, \mathbf{n}, \mathbf{l}, \alpha) = D(\mathbf{h}, \alpha) V(\mathbf{v}, \mathbf{n}, \mathbf{l}). \quad (8)$$

### 3.2. “Split-sum” Approximation

Karis [Kar13] and Lazarov [Laz13] independently proposed a method commonly referred to as the “split sum” approximation. We derive that here, providing clear meaning to each term and approximation step.

Their reformulation of Equation (1) *splits* the original integral into a product of two integrals, as

$$\begin{aligned} E(\mathbf{v}, \mathbf{n}, \Theta) &= \int_{\Omega} L(\mathbf{l}) f(\mathbf{v}, \mathbf{n}, \mathbf{l}, \Theta) \langle \mathbf{l}, \mathbf{n} \rangle d\mu(\mathbf{l}) \\ &= \int_{\Omega} L(\mathbf{l}) \underbrace{\left[ \frac{f(\mathbf{v}, \mathbf{n}, \mathbf{l}, \Theta) \langle \mathbf{l}, \mathbf{n} \rangle}{\int_{\Omega} f(\mathbf{v}, \mathbf{n}, \mathbf{l}, \Theta) \langle \mathbf{l}, \mathbf{n} \rangle d\mu(\mathbf{l})} \right]}_{K(\mathbf{v}, \mathbf{n}, \mathbf{l}, \Theta)} d\mu(\mathbf{l}) \times \int_{\Omega} f(\mathbf{v}, \mathbf{n}, \mathbf{l}, \Theta) \langle \mathbf{l}, \mathbf{n} \rangle d\mu(\mathbf{l}) \end{aligned}$$

where  $K$  in the first integral is interpreted as a (normalized, as per its denominator) weighting function for the incident radiance  $L$ .

By further assuming a normal-aligned view  $\mathbf{v} = \mathbf{n}$  in  $K$ , this term can be approximated as a normalized symmetric function  $\hat{K}(\mathbf{n}, \mathbf{l}, \Theta)$  that admits compact precomputation and efficient lookup at runtime. Specifically, with Cook-Torrance and Schlick’s approximation, and with the  $\Theta$  parameters are  $\alpha$  and  $F_0$ , the  $\mathbf{v} = \mathbf{n}$  assumption allows  $F_0$  to cancel out, leaving  $\hat{K}$  as a function of  $\mathbf{n}$ ,  $\mathbf{l}$  and  $\alpha$ . By fixing  $\alpha$ , the  $\int_{\Omega} L(\mathbf{l}) \hat{K}_{\alpha}(\mathbf{n}, \mathbf{l}) d\mu(\mathbf{l})$  integral can be precomputed for each  $\mathbf{n}$  direction and stored in, e.g., a cubemap indexed by  $\mathbf{n}$ . A common follow-up optimization is to compute and store several such cubemaps – for increasing values of  $\alpha$  – in a mipmap chain; indeed, as  $\alpha$  increases, the angular frequency of  $\hat{K}$  decreases, requiring a lower texture/tabulation resolution.

The second integral in Equation 1 can also be similarly split using Schlick’s approximation into a Fresnel term (as in Equation 5), yielding a two-function factorization. As this term has no dependence on incident radiance, the resulting integral value can be precomputed and evaluated in a canonical coordinate frame with the normal oriented along the  $z$ -axis, and depending only on  $\mathbf{v}$  and  $\alpha$ . Karis [Kar13] chooses to precompute and store this term, whereas Lazarov [Laz13] derives an analytic approximation that is evaluated at runtime.

## 4. Method

We will directly approximate the Fresnel-decomposed reflected radiance  $E \approx F_0 E_0 + (1 - F_0) E_1$  using a separable approximation and a new angular parameterization. Unlike previous work based on prefiltered environment maps using circularly symmetric kernels, we approximate the 4D base reflectance  $E_0$  directly.

We begin by considering the base reflectance  $E_0$ , before treating the Fresnel tail reflectance  $E_1$ .

### 4.1. Base Reflectance

We factorize the 4D base reflectance  $E_0$  as

$$E_0(\mathbf{v}, \mathbf{n}, \alpha) \approx P(\pi_1(\mathbf{v}, \mathbf{n})) Q(\pi_2(\mathbf{v}, \mathbf{n})), \quad (9)$$



**Figure 2: Isolating factors.** We render a sphere ( $\alpha = 0.25$ , LITTLE PARIS EIFFEL TOWER) and isolate the  $P$  (order-4 SH) and  $Q$  (order-2 SH) factors, visualizing their contribution by zeroing-out the  $Q$  or  $P$  coefficients before exponentiation. In general,  $P$  captures higher-frequency detail than  $Q$ , although the effect that increasing the order of  $P$  or  $Q$  has on MSE is scene dependent; see Figure 8.

where  $P$  and  $Q$  are 2D spherical functions parameterized by  $\pi_1$  and  $\pi_2$  respectively. We rely on a novel *reflection-half-reflection* parameterization defined by

$$\pi_1(\mathbf{v}, \mathbf{n}) = \mathbf{r} \quad \text{and} \quad \pi_2(\mathbf{v}, \mathbf{n}) = \mathbf{h}_r, \quad (10)$$

where  $\mathbf{r}$  is the mirror/perfect reflection of  $\mathbf{v}$  with respect to the normal  $\mathbf{n}$ , and  $\mathbf{h}_r$  is the half-vector between  $\mathbf{r}$  and  $\mathbf{n}$ , i.e.,  $\mathbf{h}_r = (\mathbf{n} + \mathbf{r}) / \|\mathbf{n} + \mathbf{r}\|$ . The motivation behind our parameterization is to enable coupling between the separable terms  $P$  and  $Q$  using the half-reflection vector  $\mathbf{h}_r$ , which captures the geometric relationship between the view direction  $\mathbf{v}$ , reflection direction  $\mathbf{r}$ , and surface normal  $\mathbf{n}$ . This shared reference preserves directional correlations important for accurate modeling of glossy reflections.

The main benefit of our separable approximation comes from dimensionality reduction: we approximate a 4D function  $E$  as the product of two 2D functions  $P$  and  $Q$  *without assuming*  $\mathbf{v} = \mathbf{n}$ , avoiding any systemic error due to erroneous circular symmetry assumptions (c.f. Section 5).

#### 4.1.1. Modeling Continuous Roughness

We represent the  $P$  and  $Q$  functions in terms of *low-order spherical harmonic exponentials* and encode roughness  $\alpha$  via an *integrated directional encoding*, i.e., an SH convolution with the von Mises-Fisher distribution  $\mathbf{vMF}(\kappa)$  [VHM\*22]:

$$\begin{aligned} P(\pi_1(\mathbf{v}, \mathbf{n}), \alpha) &= P(\mathbf{r}, \alpha) = \exp\left(\langle \mathbf{vMF}\left(\frac{1}{\alpha}\right) \odot \mathbf{Y}(\mathbf{r}), \mathbf{p} \rangle\right), \text{ and} \\ Q(\pi_2(\mathbf{v}, \mathbf{n}), \alpha) &= Q(\mathbf{h}_r, \alpha) = \exp\left(\langle \mathbf{vMF}\left(\frac{1}{\alpha}\right) \odot \mathbf{Y}(\mathbf{h}_r), \mathbf{q} \rangle\right), \end{aligned} \quad (11)$$

where  $\mathbf{Y}(\omega)$  is a vector of spherical harmonics basis functions evaluated at direction  $\omega$ ,  $\mathbf{p}$  and  $\mathbf{q}$  are vectors of (unknown) spherical harmonic projection coefficients for the (non-exponentiated) functions  $P$  and  $Q$ ,  $\odot$  is a component-wise multiplication, and  $\mathbf{vMF}(\kappa)$



**Figure 3: Fresnel tail reflectance.** Our Fresnel tail approximation (Section 4.2) loses some energy at grazing angles compared to reference but remains visually plausible with MSE of  $8.5 \times 10^{-6}$  (PRETVILLE STREET).

is a spherical harmonic approximation to the von Mises-Fisher distribution, where each band  $l$  is given by

$$\mathbf{vMF}_l(\kappa) = \exp\left(\frac{-l(l+1)}{2\kappa}\right). \quad (12)$$

The motivation for our exponentiation formulation is to linearize the product  $P \times Q$  in log-space, allowing us to solve for the unknown coefficients using simple least-squares while *additionally ensuring* positive values during reconstruction. Unlike prior work [RWS\*06], we never explicitly evaluate SH logs or exponentials; the log is implicit from the solve and the exponential is evaluated point-wise.

Our optimization problem over continuous variables  $\mathbf{p}$ ,  $\mathbf{q}$  and  $\alpha$  is

$$\underset{\mathbf{p}, \mathbf{q}}{\text{minimize}} \left( \log(P(\mathbf{r}, \alpha)Q(\mathbf{h}_r, \alpha)) - \log(E_0(\mathbf{v}, \mathbf{n}, \alpha)) \right)^2. \quad (13)$$

We solve Equation 13 using stochastic linear least-squares by randomly sampling view  $\mathbf{v}$  and normal  $\mathbf{n}$  directions, and roughness values  $\alpha$ . Concretely, for each  $(\mathbf{v}, \mathbf{n}, \alpha)$ -sample, the row constraint using our reflection-half-reflection parametrization is given by

$$\left[ \mathbf{vMF}\left(\frac{1}{\alpha}\right) \odot \mathbf{Y}(\mathbf{r}) \quad \mathbf{vMF}\left(\frac{1}{\alpha}\right) \odot \mathbf{Y}(\mathbf{h}_r) \right] \begin{bmatrix} \mathbf{p} \\ \mathbf{q} \end{bmatrix} = \log(E_0(\mathbf{v}, \mathbf{n}, \alpha)). \quad (14)$$

Finally, we remove one degree of freedom by observing that the constant SH basis function is included twice, leading to an ambiguous sum constraint; to avoid this, we simply omit this DC term from the  $Q$  factor. Runtime reconstruction consists of evaluating Equations 11 using the precomputed vectors  $\mathbf{p}$  and  $\mathbf{q}$ , and at *any* continuous roughness value  $\alpha$ .

## 4.2. Fresnel Tail Reflectance

We have derived a separable approximation of base reflectance  $E_0$  using spherical harmonic exponentials. We now treat the Fresnel tail reflectance  $E_1$ , completing our full factorized model of  $E$ .

We approximate the Fresnel tail reflectance  $E_1$  by assuming a constant environment light  $L$  approximated by  $E_0$ ,

$$E_1(\mathbf{v}, \mathbf{n}, \alpha) \approx E_0(\mathbf{v}, \mathbf{n}, \alpha) \int_{\Omega} f_{DV}(\mathbf{v}, \mathbf{n}, \mathbf{l}, \alpha) (1 - \langle \mathbf{v}, \mathbf{h} \rangle)^5 \langle \mathbf{l}, \mathbf{n} \rangle d\mu(\mathbf{l}). \quad (15)$$

We precompute our Fresnel tail integral approximation for efficient runtime evaluation, as in previous work [Kar13]. Figure 3 compares our Fresnel tail approximation with the ground truth.

## 5. Results

We evaluate our method on eight (8) lighting environments [Pol24] with varying time of day and location. We first compare against split-sum and ground truth (Section 5.1) before analyzing performance (Section 5.2) and performing an ablation study (Section 5.3). We discuss limitations of our method in Section 5.4. In general (i.e., in roughly 95% of the scenarios), we obtain higher quality than the split-sum state-of-the-art, all whilst requiring *significantly* less memory and supporting continuous roughness variation.

### 5.1. Comparisons

All our results use order 4 for  $P$  and order 2 for  $Q$  for consistency although, in practice, these parameters can be adapted based on roughness and environment light bandwidth. We use 64 uniformly sampled directions for both the view and normal directions, and 4 uniformly distributed roughness samples  $\alpha$  in a relaxed range of  $\alpha \in [0.20, 1.0]$  to compute our factorization, for a total of  $64 \times 64 \times 4$  ( $\mathbf{v}, \mathbf{n}, \alpha$ ) samples before solving the linear system in Equation 14. We compare our method to a mip-chain of prefiltered octahedral environment maps with a minimum resolution clamped to  $32 \times 32$  to avoid visible aliasing artifacts under magnification. Both methods use an equal number (1024) of directional samples to estimate their target integrals. Furthermore, we only compute the split-sum mip-chain for equal roughness range for a fair comparison.

We compare to split-sum approximation and ground truth reference in Figures 4 and 5 with roughness values  $\alpha$  distributed according to a perceptual roughness remapping [Bur12]. Our results compare favorable to split-sum approximations, both quantitatively and qualitatively. In fact, our method obtains lower MSE in **95% of our 192 test cases** while using **over  $200\times$  less memory**, enabling alias-free reconstruction of glossy reflections over continuous roughness and angular domains. We compare the MSE error distributions in Figure 7 and visualize MSE error in Figure 9.

### 5.2. Performance

We implement our precomputation on the GPU with CUDA and our runtime in HLSL. We report results on an AMD Ryzen Threadripper 3970X CPU and a NVIDIA RTX A6000 GPU. Our precomputation and runtime performance are similar to the split-sum method: our mean precomputation time is 0.328ms compared to 0.536ms for split-sum, averaged over 1000 runs; the runtime overhead of our reconstruction is roughly 0.1ms, or 1% of total GPU frame time, atop the split-sum cost, at a target resolution of  $1920 \times 1080$ .

### 5.3. Ablation Study

We study the effect that varying the SH degree of  $P$  and  $Q$  has on the reconstruction quality (Figure 8): our method consistently converges as the degree for  $P$  and  $Q$  increases while the relative importance of these factors depends on the environment. Note that the MSE is not

guaranteed to be strictly decreasing as the degree increases, e.g., in BELFAST SUNSET, since log-space optimization is only a proxy for linear space MSE error.

### 5.4. Limitations

As shown in Figure 1, our method supports materials with a wide gamut of continuously varying roughness. The main limitation of our method is the representation capacity of the factorized spherical harmonic exponential model when using low degree spherical harmonics for the factors  $P$  and  $Q$ . In particular, the low roughness regime, i.e.,  $\alpha < 0.25$  is not well suited for a severely bandlimited SH approximation as shown in Figure 6. However, increasing the degrees of  $P$  and  $Q$  helps our model to converge to the ground truth reference. Finally, since we perform fitting over a continuous range of roughness values, the linear least squares optimization tends to prioritize minimizing larger errors that occur in the low-roughness regime. Nevertheless, our method compares favorably to the split-sum approach, which suffers from increasing error at higher roughness due to its circular symmetry assumption.

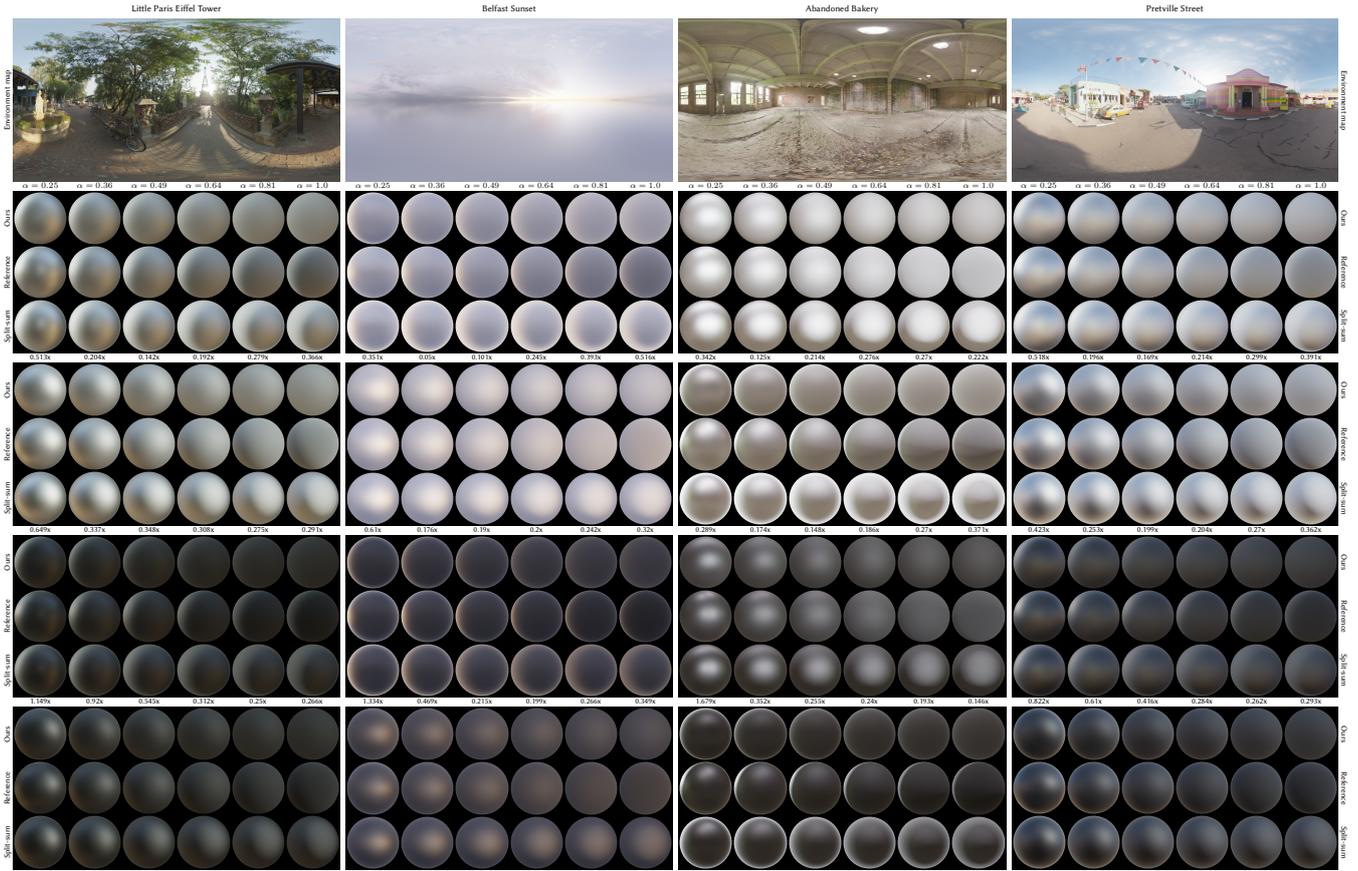
## 6. Conclusion

We introduced a representation for glossy reflections based on factorized spherical harmonic exponentials. Compared to previous work, our method increases reflection quality at significantly reduced memory usage, enabling efficient and alias-free reconstruction over continuous angular and roughness domains.

Our approach extends the widely used linear SH representation of diffuse irradiance to support a wider range of glossy materials [RH02]. This advancement has the potential for immediate practical application in real-time rendering across various hardware platforms, offering scalable, spatially localized glossy reflections that are currently unattainable due to the high memory demands of existing methods. Moreover, because our representation is differentiable, it holds promise for computer vision applications, offering an alternative to linear spherical harmonics expansions in neural radiance fields for modeling view-dependent appearance.

## References

- [AP07] ASHIKHMIN, MICHAEL and PREMOZE, SIMON. *Distribution-based BRDFs*. Tech. rep. Department of Computer Science, University of Utah, Mar. 2007 3.
- [APS00] ASHIKHMIN, MICHAEL, PREMOZE, SIMON, and SHIRLEY, PETER. "A Microfacet-based BRDF Generator". *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH '00. New York, NY, USA: ACM Press/Addison-Wesley Publishing Co., 2000, 65–74. ISBN: 1-58113-208-5. DOI: [10.1145/344779.344814](https://doi.org/10.1145/344779.344814). URL: <http://dx.doi.org/10.1145/344779.344814>.
- [AS00] ASHIKHMIN, MICHAEL and SHIRLEY, PETER. "An Anisotropic Phong BRDF Model". *J. Graph. Tools* 5.2 (Feb. 2000), 25–32. ISSN: 1086-7651. DOI: [10.1080/10867651.2000.10487522](https://doi.org/10.1080/10867651.2000.10487522). URL: <http://dx.doi.org/10.1080/10867651.2000.10487522>.
- [BBS02] BOURLIER, C., BERGIN, G., and SAILLARD, JOSEPH. "One- and two-dimensional shadowing functions for any height and slope stationary uncorrelated surface in the monostatic and bistatic configurations". *Antennas and Propagation, IEEE Transactions on* 50.3 (2002), 312–324. ISSN: 0018-926X. DOI: [10.1109/8.999622](https://doi.org/10.1109/8.999622).



**Figure 4: Comparison to split-sum and ground truth.** We render chrome (top two sub-rows) and glass spheres with  $F_0 = 0.04$  (bottom two sub-rows) under different environments. We compare to split-sum [Kar13] and ray-traced references with roughness  $\alpha \in [0.25, 1.0]$ . We report the MSE ratio at the bottom of every column (ours / split-sum); numbers smaller than one indicate that our method has smaller MSE than split-sum. Our method compares favorably in terms of reconstruction error while using two orders of magnitude less memory and enabling alias-free reconstruction of glossy reflections.

- [Bli77] BLINN, JAMES F. “Models of Light Reflection for Computer Synthesized Pictures”. *SIGGRAPH Comput. Graph.* 11.2 (July 1977), 192–198. ISSN: 0097-8930. DOI: [10.1145/965141.563893](https://doi.org/10.1145/965141.563893). URL: <http://doi.acm.org/10.1145/965141.563893>.
- [Bro80] BROWN, G.S. “Shadowing by non-Gaussian random surfaces”. *Antennas and Propagation, IEEE Transactions on* 28.6 (1980), 788–790. ISSN: 0018-926X. DOI: [10.1109/TAP.1980.1142437](https://doi.org/10.1109/TAP.1980.1142437).
- [BSH12] BAGHER M., MAHDI, SOLER, CYRIL, and HOLZSCHUCH, NICOLAS. “Accurate fitting of measured reflectances using a Shifted Gamma micro-facet distribution”. *Computer Graphics Forum* 31.4 (June 2012), 1509–1518. DOI: [10.1111/j.1467-8659.2012.03147.x](https://doi.org/10.1111/j.1467-8659.2012.03147.x). URL: <http://hal.inria.fr/hal-00702304>.
- [BSN16] BAGHER, MAHDI M., SNYDER, JOHN, and NOWROUZSAHRAI, DEREK. “A Non-Parametric Factor Microfacet Model for Isotropic BRDFs”. 35.5 (July 2016). ISSN: 0730-0301. DOI: [10/f85ktv3](https://doi.org/10/f85ktv3).
- [Bur12] BURLEY, BRENT. *Physically-based Shading at Disney*. Tech. rep. Walt Disney Animation Studios, Aug. 2012, 3, 6.
- [CT82] COOK, ROBERT L. and TORRANCE, KENNETH E. “A Reflectance Model for Computer Graphics”. *ACM Transactions on Graphics* 1.1 (Jan. 1982), 7–24. ISSN: 0730-0301. DOI: [10.1145/357290.357293](https://doi.org/10.1145/357290.357293). URL: <https://doi.org/10.1145/357290.357293>.
- [DHI\*15] DUPUY, JONATHAN, HEITZ, ERIC, IEHL, JEAN-CLAUDE, et al. “Extracting Microfacet-based BRDF Parameters from Arbitrary Materials with Power Iterations”. *Computer Graphics Forum* (2015), 10. URL: <https://hal.inria.fr/hal-01168516>.
- [DRS08] DORSEY, JULIE, RUSHMEIER, HOLLY, and SILLION, FRANOIS. *Digital Modeling of Material Appearance*. San Francisco, CA, USA: Morgan Kaufmann Inc., 2008. ISBN: 0122211812, 9780080556710, 9780122211812.
- [GHP\*08] GHOSH, ABHIJEET, HAWKINS, TIM, PEERS, PIETER, et al. “Practical Modeling and Acquisition of Layered Facial Reflectance”. *ACM Trans. Graph.* 27.5 (Dec. 2008), 139:1–139:10 3.
- [GKD07] GREEN, PAUL, KAUTZ, JAN, and DURAND, FRÉDO. “Efficient Reflectance and Visibility Approximations for Environment Map Rendering”. *Computer Graphics Forum* 26 (2007). URL: <https://api.semanticscholar.org/CorpusID:15451153>.
- [Gre86] GREENE, NED. “Environment Mapping and Other Applications of World Projections”. 6.11 (Nov. 1986). ISSN: 0272-1716. DOI: [10/b7m3gc2](https://doi.org/10/b7m3gc2).
- [Hei14] HEITZ, ERIC. “Understanding the Masking-Shadowing Function in Microfacet-Based BRDFs”. *Journal of Computer Graphics Techniques* 3.2 (June 2014), 32–91 4.



**Figure 5: Comparison to split-sum and ground truth.** Chrome (top two sub-rows) and glass (bottom two sub-rows) spheres with  $F_0 = 0.04$  under various lighting. We compare to split-sum [Kar13] and ground truth with  $\alpha \in [0.25, 1.0]$  and report MSE ratio (bottom of every column; ours / split-sum, as in Figure 4).

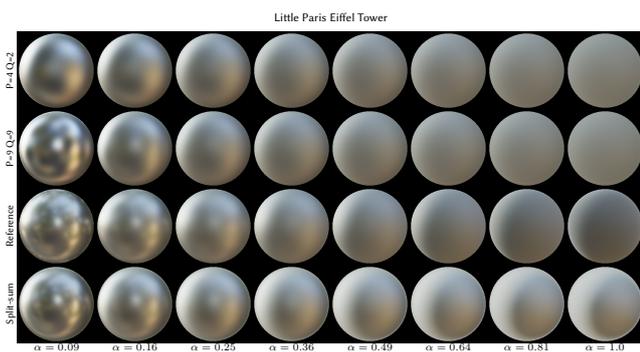
[HS99] HEIDRICH, WOLFGANG and SEIDEL, HANS-PETER. “Realistic, Hardware-Accelerated Shading and Lighting”. Aug. 1999 2.

[Kar13] KARIS, BRIAN. “Real Shading in Unreal Engine 4”. 2013. URL: <https://api.semanticscholar.org/CorpusID:14922512> 1, 2, 4, 6–8.

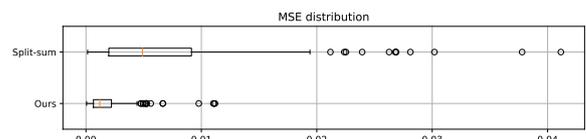
[KM00] KAUTZ, JAN and MCCOOL, MICHAEL. “Approximation of Glossy Reflection with Prefiltered Environment Maps.” Jan. 2000, 119–126 2.

[KM99] KAUTZ, JAN and MCCOOL, MICHAEL D. “Interactive Rendering with Arbitrary BRDFs using Separable Approximations”. *Rendering Techniques*. 1999 3.

[KSK10] KURT, MURAT, SZIRMAJ-KALOS, LÁSZLÓ, and KŘIVÁNEK, JAROSLAV. “An Anisotropic BRDF Model for Fitting and Monte Carlo Rendering”. *SIGGRAPH Comput. Graph.* 44.1 (Feb. 2010), 3:1–3:15. ISSN: 0097-8930 3.

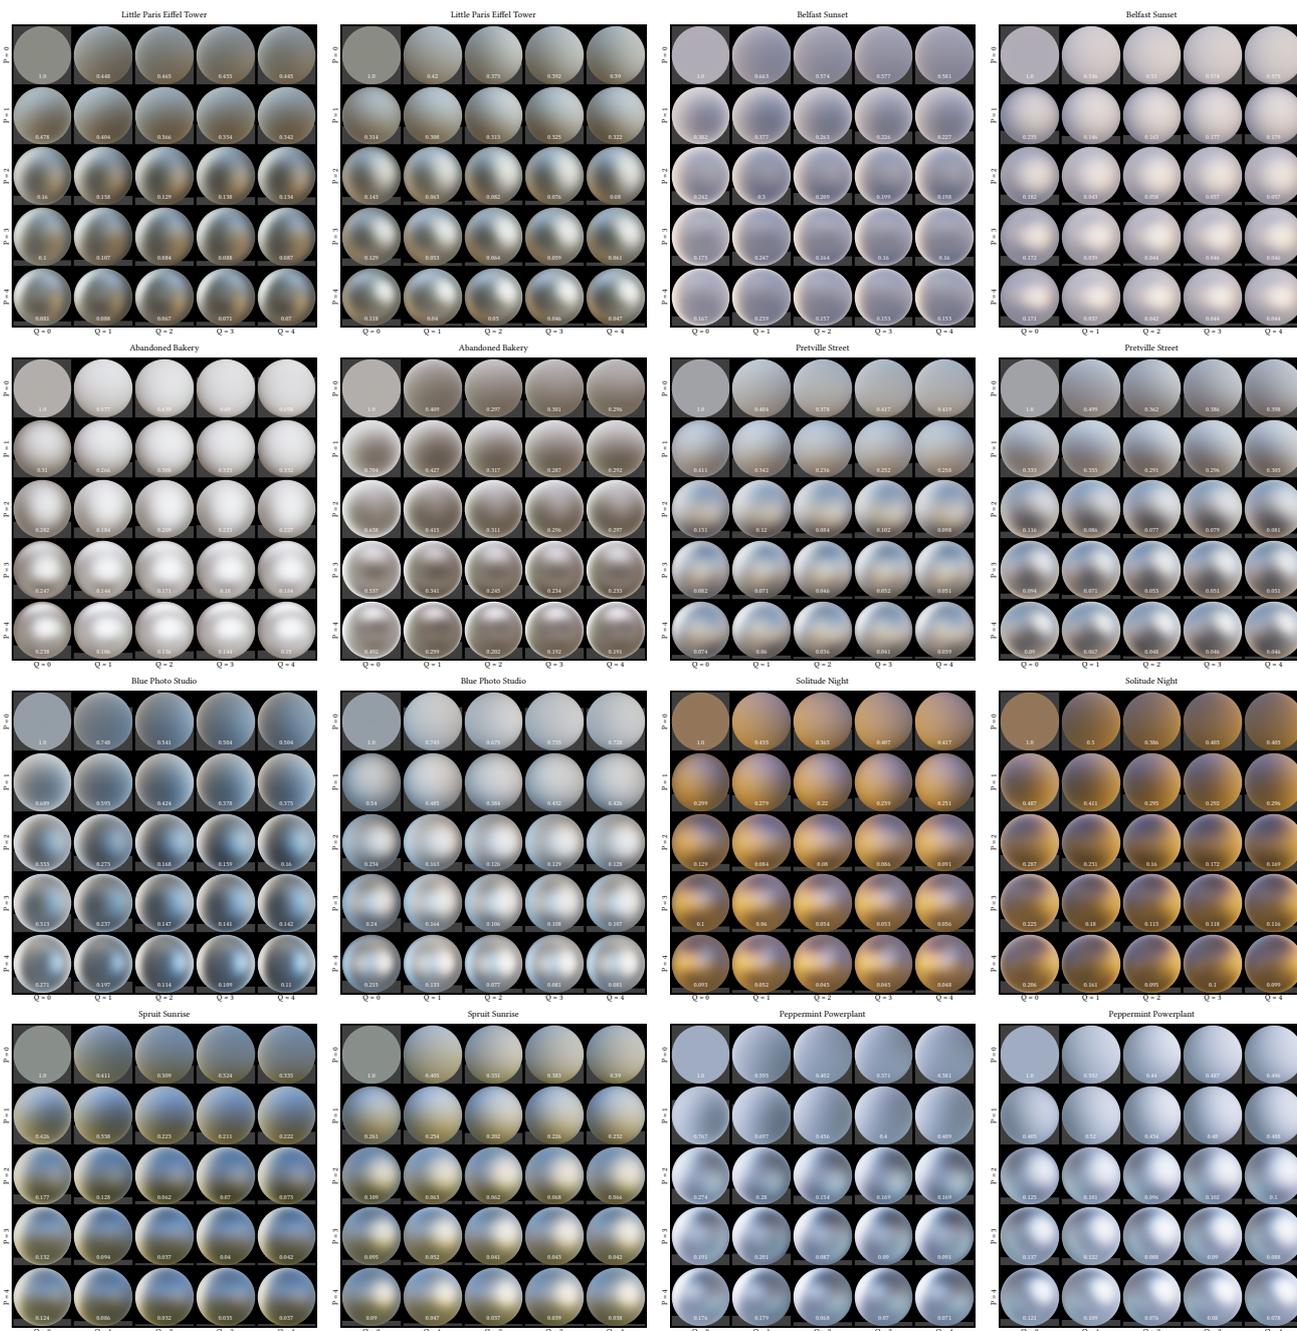


**Figure 6: Limitations.** We solve for our factorized SH representation using a wide range of roughness  $\alpha \in [0.09, 1.0]$  to demonstrate the limitations of our bandlimited model (top row). Increasing the orders of  $P$  and  $Q$  yields results (second row) that are closer to ground truth (third row). We include split-sum results for completeness (bottom).

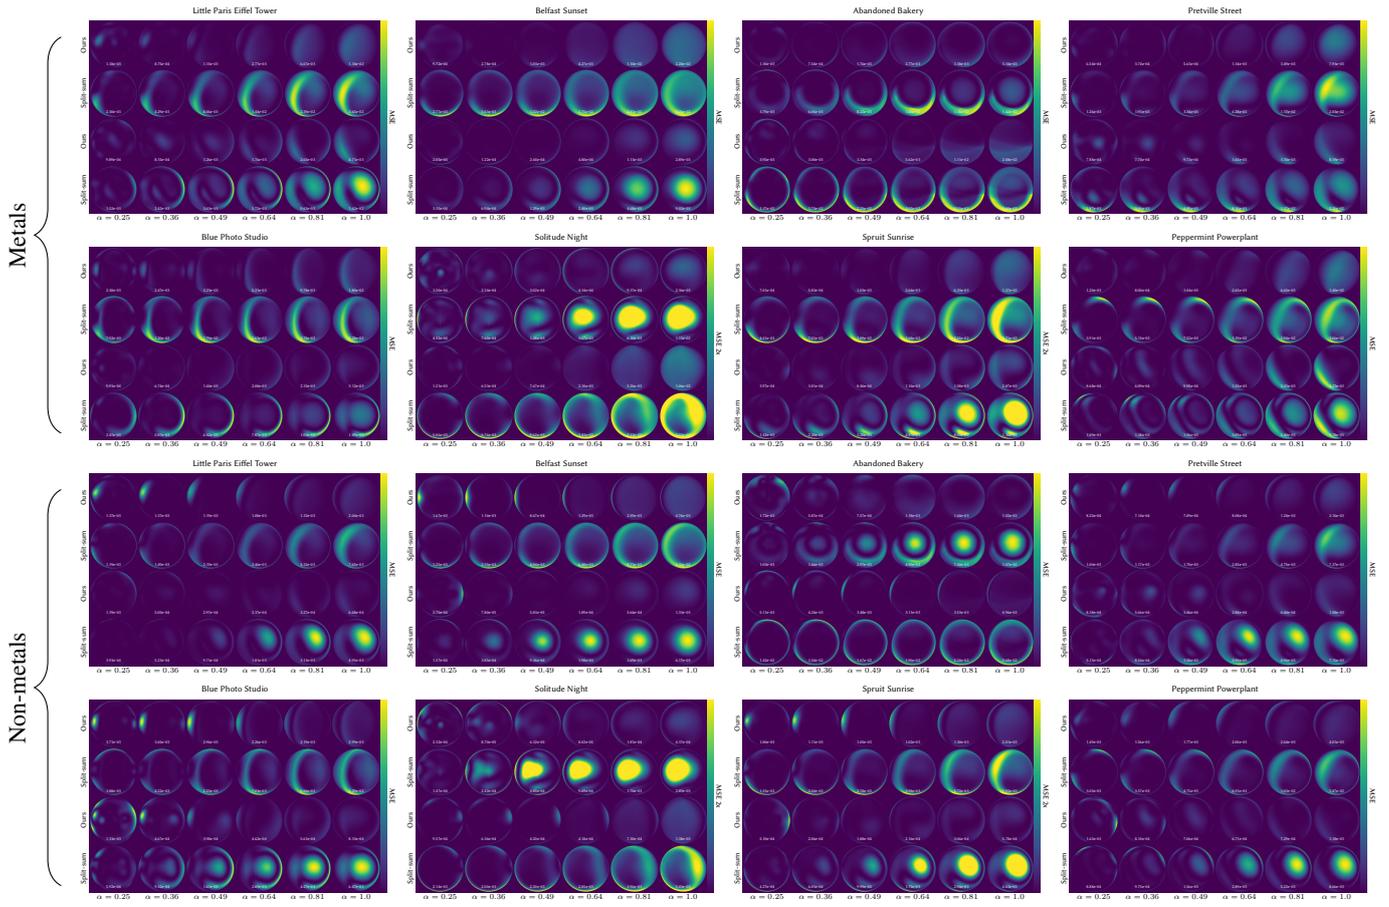


**Figure 7: Error distribution.** MSE error distribution for all 192 results in Figures 4 and 5. Compared to split-sum, our method has lower MSE in 95% of these cases while using  $>200\times$  less memory.

- [KVHS00] KAUTZ, JAN, VÁZQUEZ, PERE-PAU, HEIDRICH, WOLFGANG, and SEIDEL, HANS-PETER. “Unified Approach to Prefiltered Environment Maps”. *Proceedings of the Eurographics Workshop on Rendering Techniques 2000*. Berlin, Heidelberg: Springer-Verlag, 2000, 185–196. ISBN: 3211835350 2.
- [Laz13] LAZAROV, DIMITAR. “Getting More Physical in Call of Duty: Black Ops II”. 2013. URL: <https://blog.selfshadow.com/publications/s2013-shading-course/2,4>.
- [LFTG97] LAFORTUNE, ERIC P. F., FOO, SING-CHOONG, TORRANCE, KENNETH E., and GREENBERG, DONALD P. “Non-linear Approximation of Reflectance Functions”. *Proceedings of SIGGRAPH*. New York, NY, USA: ACM Press/Addison-Wesley Publishing Co., 1997, 117–126 3.
- [LK02] LATTA, LUTZ and KOLB, ANDREAS. “Homomorphic factorization of BRDF-based lighting computation”. *ACM Trans. Graph.* 21.3 (July 2002), 509–516. ISSN: 0730-0301 2.
- [LK03] LEHTINEN, JAAKKO and KAUTZ, JAN. “Matrix radiance transfer”. *Proceedings of the 2003 Symposium on Interactive 3D Graphics*. I3D '03. Monterey, California, 2003, 59–64. ISBN: 1581136455 3.
- [LKYU12] LÖW, JOAKIM, KRONANDER, JOEL, YNNERMAN, ANDERS, and UNGER, JONAS. “BRDF Models for Accurate and Efficient Rendering of Glossy Surfaces”. *ACM Transactions on Graphics (TOG)* 31.1 (2012), 9:1–9:14 3.
- [LRR04] LAWRENCE, JASON, RUSINKIEWICZ, SZYMON, and RAMAMOORTHY, RAVI. “Efficient BRDF Importance Sampling Using a Factored Representation”. *ACM Trans. Graph.* 23.3 (Aug. 2004), 496–505. ISSN: 0730-0301. DOI: [10.1145/1015706.1015751](https://doi.org/10.1145/1015706.1015751). URL: <http://doi.acm.org/10.1145/1015706.1015751> 3.
- [LS00] LEE, DANIEL D. and SEUNG, H. SEBASTIAN. “Algorithms for Non-negative Matrix Factorization”. In *NIPS*. MIT Press, 2000, 556–562 3.
- [LS05] LAZÁNYI, ISTVÁN and SZIRMAY-KALOS, LÁSZLÓ. “Fresnel Term Approximations for Metals.” *WSCG (Short Papers)*. Aug. 1, 2005, 77–80 3.
- [LSS04] LIU, XINGUO, SLOAN, PETER-PIKE, SHUM, HEUNG-YEUNG, and SNYDER, JOHN. “All-Frequency Precomputed Radiance Transfer for Glossy Objects”. June 2004. DOI: [10/gfz7293](https://doi.org/10/gfz7293).
- [MAA01] MCCOOL, MICHAEL D., ANG, JASON, and AHMAD, ANIS. “Homomorphic factorization of BRDFs for high-performance rendering”. *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH '01. New York, NY, USA: Association for Computing Machinery, 2001, 171–178 3.
- [MHH\*12] MCAULEY, STEPHEN, HILL, STEPHEN, HOFFMAN, NATY, et al. “Practical Physically-based Shading in Film and Game Production”. *ACM SIGGRAPH 2012 Courses*. SIGGRAPH '12. Los Angeles, California: ACM, 2012, 10:1–10:7. ISBN: 978-1-4503-1678-1. DOI: [10.1145/2343483.2343493](https://doi.org/10.1145/2343483.2343493) 3.
- [MLH02] MCALLISTER, DAVID K., LASTRA, ANSELMO, and HEIDRICH, WOLFGANG. “Efficient rendering of spatial bi-directional reflectance distribution functions”. *Proceedings of the ACM SIGGRAPH/EUROGRAPHICS Conference on Graphics Hardware*. HWWS '02. Saarbrücken, Germany: Eurographics Association, 2002, 79–88. ISBN: 1581135807 3.
- [MPBM03] MATUSIK, WOJCIECH, PFISTER, HANSPETER, BRAND, MATT, and MCMILLAN, LEONARD. “A Data-driven Reflectance Model”. *ACM SIGGRAPH 2003 Papers*. SIGGRAPH '03. San Diego, California: ACM, 2003, 759–769. ISBN: 1-58113-709-5. DOI: [10.1145/1201775.8823433](https://doi.org/10.1145/1201775.8823433) 3.
- [MWL\*99] MARSCHNER, STEPHEN R., WESTIN, STEPHEN H., LAFORTUNE, ERIC P. F., et al. “Image-based BRDF Measurement Including Human Skin”. *Proceedings of the 10th Eurographics Conference on Rendering*. EGWR'99. Granada, Spain: Eurographics Association, 1999, 131–144. ISBN: 3-211-83382-X. DOI: [10.2312/EGWR/EGWR99/131-144](https://doi.org/10.2312/EGWR/EGWR99/131-144). URL: <http://dx.doi.org/10.2312/EGWR/EGWR99/131-144> 3.
- [NDM05] NGAN, ADDY, DURAND, FRÉDO, and MATUSIK, WOJCIECH. “Experimental Analysis of BRDF Models”. *Proceedings of the Eurographics Symposium on Rendering*. Konstanz, Germany: Eurographics Association, 2005, 117–226 3.
- [NRH03] NG, REN, RAMAMOORTHY, RAVI, and HANRAHAN, PAT. “All-frequency shadows using non-linear wavelet lighting approximation”. *ACM SIGGRAPH 2003 Papers*. 2003, 376–381 3.
- [PF90] POULIN, PIERRE and FOURNIER, ALAIN. “A Model for Anisotropic Reflection”. *SIGGRAPH Comput. Graph.* 24.4 (Sept. 1990), 273–282. ISSN: 0097-8930. DOI: [10.1145/97880.97909](https://doi.org/10.1145/97880.97909). URL: <http://doi.acm.org/10.1145/97880.97909> 3.
- [Pol24] POLY HAVEN. *Environment Maps*. 2024. URL: <https://polyhaven.com/hdris6>.
- [PSS\*12] PACANOWSKI, ROMAIN, SALAZAR-CELIS, OLIVER, SCHLICK, CHRISTOPHE, et al. “Rational BRDF”. English. *IEEE Transactions on Visualization and Computer Graphics* 18.11 (Feb. 2012), 1824–1835. DOI: [10.1109/TVCG.2012.73](https://doi.org/10.1109/TVCG.2012.73). URL: <http://hal.inria.fr/hal-00678885> 3.
- [Ram09] RAMAMOORTHY, RAVI. “Precomputation-Based Rendering”. *Foundations and Trends® in Computer Graphics and Vision* 3.4 (Apr. 2009). ISSN: 1572-2740, 1572-2759. DOI: [10/c8zq2b2](https://doi.org/10/c8zq2b2) 2.
- [RH02] RAMAMOORTHY, RAVI and HANRAHAN, PAT. “Frequency space environment map rendering”. *ACM Trans. Graph.* 21.3 (July 2002), 517–526. ISSN: 0730-0301. DOI: [10.1145/566654.566611](https://doi.org/10.1145/566654.566611). URL: <https://doi.org/10.1145/566654.566611> 2, 6.
- [RWS\*06] REN, ZHONG, WANG, RUI, SNYDER, JOHN, et al. “Real-Time Soft Shadows in Dynamic Scenes Using Spherical Harmonic Exponentiation”. 25.3 (July 2006). DOI: [10/fwjdd592,3,5](https://doi.org/10/fwjdd592,3,5).
- [SBN15] SOLER, CYRIL, BAGHER, MAHDI M., and NOWROUZEZAHRAI, DEREK. “Efficient and Accurate Spherical Kernel Integrals Using Isotropic Decomposition”. *ACM Trans. Graph.* 34.5 (2015), 161:1–161:14. DOI: [10.1145/2797136](https://doi.org/10.1145/2797136). URL: <https://doi.org/10.1145/2797136> 2.
- [Sch94] SCHLICK, CHRISTOPHE. “An Inexpensive BRDF Model for Physically-based Rendering”. *Computer Graphics Forum* 13 (1994), 233–246 2–4.
- [SKS02] SLOAN, PETER-PIKE, KAUTZ, JAN, and SNYDER, JOHN. “Pre-computed Radiance Transfer for Real-Time Rendering in Dynamic, Low-Frequency Lighting Environments”. 21.3 (2002). ISSN: 0730-0301. DOI: [10/fqq3kn3](https://doi.org/10/fqq3kn3) 3.
- [SM02] STEIGLEDER, MAURO and MCCOOL, MICHAEL D. “Factorization of the Ashikhmin BRDF for Real-Time Rendering”. *J. Graphics, GPU, & Game Tools* 7.4 (2002), 61–67 3.
- [Smi67] SMITH, B. “Geometrical shadowing of a random rough surface”. *Antennas and Propagation, IEEE Transactions on* 15.5 (1967), 668–671. ISSN: 0018-926X. DOI: [10.1109/TAP.1967.1138991](https://doi.org/10.1109/TAP.1967.1138991) 3.
- [VHM\*22] VERBIN, DOR, HEDMAN, PETER, MILDENHALL, BEN, et al. “Ref-NeRF: Structured View-Dependent Appearance for Neural Radiance Fields”. *CVPR* (2022) 5.
- [War92] WARD, GREGORY J. “Measuring and Modeling Anisotropic Reflection”. *SIGGRAPH Comput. Graph.* 26.2 (July 1992), 265–272. ISSN: 0097-8930. DOI: [10.1145/142920.134078](https://doi.org/10.1145/142920.134078). URL: <http://doi.acm.org/10.1145/142920.134078> 3.
- [WMLT07] WALTER, BRUCE, MARSCHNER, STEPHEN R., LI, HONGSONG, and TORRANCE, KENNETH E. “Microfacet Models for Refraction through Rough Surfaces”. *Rendering Techniques*. Ed. by KAUTZ, JAN and PATTANAIK, SUMANTA. The Eurographics Association, 2007. ISBN: 978-3-905673-52-4. DOI: [10.2312/EGWR/EGSR07/195-2064](https://doi.org/10.2312/EGWR/EGSR07/195-2064) 4.
- [WZT\*08] WANG, JIAPING, ZHAO, SHUANG, TONG, XIN, et al. “Modeling anisotropic surface reflectance with example-based microfacet synthesis”. Vol. 27. 3. Aug. 2008, 41:1–41:9 3.



**Figure 8: Effect of SH order on the approximation.** We perform an ablation study by varying the SH order of the  $P$  and  $Q$  features at  $\alpha = 0.25$ . The rows show the effect of increasing order for  $P$  and columns for increasing order of for  $Q$ . We report MSE numbers and further normalize these values such that a constant, or degree (0,0) approximation has MSE of 1 for readability. In addition, we visualize the normalized error as grey bars where height is of the bar is directly related to the normalized error values. We observe that the approximation error decreases as the SH degree increases, creating a consistent design space for exploring memory vs. quality trade-offs in different applications.



**Figure 9: Error visualization.** We visualize the error distribution and compare the MSE error between our method and split-sum method by rendering a chrome ball (top two rows) and a glass ball with  $F_0 = 0.04$  (bottom two rows), both with roughness  $\alpha = 0.25$  under different environment lights. The view positions are the same as in Figures 4 and 5. Our method compares favorably against split-sum approximation and avoids the systemic error caused by view assumptions as roughness increases.