



HAL
open science

Efficient and Accurate Spherical Kernel Integrals using Isotropic Decomposition

Cyril Soler, Mahdi Bagher, Derek Nowrouzezahrai

► **To cite this version:**

Cyril Soler, Mahdi Bagher, Derek Nowrouzezahrai. Efficient and Accurate Spherical Kernel Integrals using Isotropic Decomposition. ACM Transactions on Graphics, 2015, 34 (5), pp.14. 10.1145/2797136 . hal-01187865

HAL Id: hal-01187865

<https://inria.hal.science/hal-01187865v1>

Submitted on 17 May 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NoDerivatives 4.0 International License

Efficient and Accurate Spherical Kernel Integrals using Isotropic Decomposition

CYRIL SOLER

INRIA - Grenoble University

and

MAHDI M. BAGHER

University of Montreal, Microsoft corporation

and

DEREK NOWROUZEZAHRAI

University of Montreal

Spherical filtering is fundamental to many problems in image synthesis, such as computing the reflected light over a surface or anti-aliasing mirror reflections over a pixel. This operation is challenging since the profile of spherical filters (e.g., the view-evaluated BRDF or the geometry-warped pixel footprint, above) typically exhibits both spatial- and rotational-variation at each pixel, precluding precomputed solutions. We accelerate complex spherical filtering tasks using *isotropic spherical decomposition* (ISD), decomposing spherical filters into a linear combination of simpler *isotropic* kernels. Our general ISD is flexible to the choice of the isotropic kernels, and we demonstrate practical realizations of ISD on several problems in rendering: shading and prefiltering with spatially-varying BRDFs, anti-aliasing environment mapped mirror reflections, and filtering of noisy reflectance data. Compared to previous basis-space rendering solutions, our shading solution generates ground truth-quality results at interactive rates, avoiding costly reconstruction and large approximation errors.

Categories and Subject Descriptors: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—*Image Filtering*

Additional Key Words and Phrases: Spherical filtering, Zonal harmonics, Measured BRDFs, Interactive rendering

1. INTRODUCTION

Signals defined on the sphere appear in the computer graphics, astrophysics, medical imaging, and neutron transport communities. Filtering these signals is fundamental to several important problems in image synthesis. A fundamental challenge is that the con-

volution of two arbitrary spherical functions does not, in general, result in another spherical signal; instead, this convolution yields a function on the 3D rotation group $\mathcal{SO}(3)$. The dimensionality of $\mathcal{SO}(3)$ precludes typical precomputed solutions, where the filtered result is tabulated at discrete “locations” (i.e., an orthonormal 3D coordinate frame) in the domain. Using such a discretization, particularly at the resolution necessary to generate accurate results for spherical filtering problems in computer graphics, is difficult and prohibitively costly.

In the special case where one of the two spherical functions in the convolution exhibits *circular symmetry*, the result of the convolution simplifies to a spherical function, allowing for precomputed solutions (albeit still at restrictive, low discretization resolutions): indeed, this approach has been used in graphics, e.g., when representing diffuse and glossy reflection maps [Greger et al. 1998; Heidrich and Seidel 1999]. A more efficient solution to this special case scenario leverages frequency-space decompositions with spherical harmonics (SH), where the Funke-Hecke theorem applies directly (e.g., [Ramamoorthi and Hanrahan 2001]). We use the terms “convolution” and “filtering” interchangeably, with the convention that the filter’s shape and orientation may both vary spatially.

We propose an *Isotropic Spherical Decomposition* (ISD) to represent arbitrary spherical kernels as weighted sums of shifted *circularly symmetric* kernels, reducing the problem of general spherical convolution to that of the aforementioned simpler special-case. Our decomposition allows filtering with time complexity *linear* in the decomposition size (e.g., quadratic in L_{max}), and space complexity *linear* in the kernel’s maximum frequency bandwidth L_{max} . In comparison, the most efficient spherical frequency-space convolution techniques have time complexity that is *cubic* in L_{max} (due to the rotation necessary to align the filter to the signal) and storage cost that is *quadratic* in L_{max} .

Alternate solutions for decomposing spherical filters based on fitting a set of function lobes, both isotropic [Wang et al. 2009] or anisotropic [Xu et al. 2013], do exist. These approaches can fit both the lobe orientations and scales, providing reasonable filter approximations with very few basis functions in many cases; however, the non-linear nature of the lobe fitting procedures preclude any predictable approximation error in the filter representation, particularly since fitting more functions is prone to reduced numerical stability. We instead rely on a different solution: our framework *fixes* lobe orientations (i.e., directions), and we construct bases that span (in an L_2 sense) band limited function spaces in the frequency domain. As a consequence, our basis construction completely avoids

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2015 ACM 0730-0301/2015/15-ARTXXX \$10.00

DOI 10.1145/XXXXXXX.YYYYYYY

<http://doi.acm.org/10.1145/XXXXXXX.YYYYYYY>

any non-linear fitting, requiring only the precomputed solution of a simple linear system.

A side effect of L_2 convergence is that the number of basis functions required to represent signals with increasing frequency content has to grow quadratically. Our contribution is to construct complete L_2 band limited bases using rotated copies of only a few mother functions, leading to pre-filtering storage costs that scale linearly with the maximum bandwidth L_{max} . In other words, a single pre-filtered map can be used to reconstruct any (and all) of the $\mathcal{O}(L_{max})$ band- L_{max} basis functions. This is one key to the efficiency of our method. Another important advantage of our representation is that it supports signal rotation with zero additional computation costs: filtering with a rotated input spherical function reduces to lookups with rotated spherical coordinates, completely avoiding the costly rotation operations of, e.g., spherical harmonic-based filtering techniques. This affords us the flexibility of using many more basis functions than traditional SH rendering methods, while still maintaining interactivity.

Our theory is flexible in the choice of the underlying mother basis functions. For example, enforcing non-negativity in the mother functions can lead to the application of hardware prefiltering when precomputing pre-filtered data. Since we base our analysis on approximations of increasing frequency bandwidth, a natural common space for us to frame our formulation is the space of harmonic functions over the sphere: harmonic functions have an extremely compact support in the frequency domain and are steerable, and offer L_2 convergence, which makes them perfect candidates. We show indeed that our theory leads to computations with very compact forms when employing harmonic functions as a common denominator, especially when compared to other potential base representations. The use of harmonic functions also eliminates the need to explicitly store pre-filtered data, as this data can be directly represented as a spherical harmonic expansion.

The idea of isotropic decomposition has also been used in the planar domain for texture filtering [McCormack et al. 1999]. The spherical domain however brings additional properties which offer the possibility to exactly represent any bandlimited function with a finite set of rotationally symmetric kernels. To our knowledge, no such translation-invariant basis exists in the plane.

We present a general theory of ISD (Section 4) before exploring realizations of ISD using different isotropic basis functions (Section 5). We then devise several efficient ISD algorithms, reducing spherical convolution to simple (and sometimes sparse) linear algebra operations, and we demonstrate its practical benefits on three challenging rendering applications (Section 6):

- real time rendering of anisotropic and isotropic measured BRDFs under environmental lighting, including convergence to ground truth; we additionally extend this use case to handle visibility without the need for any spherical harmonic rotations,
- anti-aliasing high-frequency environmental mirror reflections over a pixel footprints, and
- an efficient solution for filtering extremely sparse, noisy anisotropic BRDF datasets captured from the real-world.

None of our applications require any precomputation beyond a negligible zonal harmonic filter decomposition, and we can therefore support geometry, view, and lighting manipulation without incurring any additional cost at run-time. We additionally demonstrate that our technique is compatible with existing self-shadowing solutions, with the usual added cost of visibility pre-computation and storage, which precludes dynamic geometry.

2. PREVIOUS WORK

The development of efficient and accurate filtering approaches is a long-standing problem in many fields with several decades of prior work, however fewer approaches explicitly address the problem of *spherical* signal filtering. If we first consider the simpler case of (2D) *image* filtering, the naïve application of an $n \times n$ -sized filter to an image has time complexity $\mathcal{O}(n^4)$, and the *fast Fourier transform* (FFT) permits an optimization to $\mathcal{O}(n^2 \log n^2)$ for periodic domains. Here, we assume that the filter size is chosen according to the dimensions of the image or the size of image features, making it proportional to the image’s effective resolution. In the spherical domain however, due to an additional degree of (rotational) freedom, the brute-force filtering approach has time complexity $\mathcal{O}(n^5)$ and no existing frequency-space approach can handle the *general* spherical image processing problem (see Section 3). Given this discrepancy between “standard” image filtering and spherical image filtering, we focus on spherical signal processing techniques that most closely motivate our approach. Furthermore, in Section 5 we illustrate several realizations of our ISD, one of which corresponds to the *first* frequency-space spherical filtering approach to support arbitrary spherical kernel and signals.

Spherical Function Representation. Several discrete spherical signal representations are based on applying subdivision schemes to platonic solids, including cube and octahedron [Miller and Hoffman 1984; Grski et al. 2005; Wan et al. 2007], dodecahedra [Schröder and Sweldens 1995], and icosahedra [Fekete 1990]. These representations can attain arbitrary accuracy at the cost of additional storage, however using them to perform spherical filtering still reduces to a costly brute-force numerical integration over their domain.

Basis representations are an alternative to these tabulated forms. Spherical harmonics [Hobson 1931] are a frequency-space representation with several attractive properties including orthonormality, trivial L_2 projection and reconstruction, and a convolution formulation for *circularly symmetric* kernels (which we leverage in Section 5.1). SH representations grow quadratically with a signal’s frequency content. Spherical wavelets [Schröder and Sweldens 1995; Mallat 2008] are a multi-resolution basis that permit efficient representation of localized signals, however they do not admit any accelerated filtering operation: the time complexity of filtering in a wavelet domain is identical to that of filtering naïve tabulated representations, albeit with a lower constant of computation due to sparsity and dynamic programming solutions. Data-driven bases have been used in BRDF and shading compression [McCool et al. 2001] yet, like wavelets, they can be interpreted as a signal discretization that still mandates brute-force integration; in this case, the constant of computation is reduced according to the compactness of the data-driven basis instead of its sparsity. Spherical piecewise constant bases have also been used in the context of all frequency lighting [Xu et al. 2008], with similar tradeoffs as spherical wavelet representations.

Spherical Gaussians (SGs) [Green et al. 2007; Xu et al. 2013] admit compact representations and are the only basis capable of efficiently filtering spherical signals. SGs approximations are typically fit using an expensive and unstable non-linear operation that precludes using many basis function terms in an expansion. In practice, usually only a handful of SGs are employed, resulting in a large approximation error. Our fitting is stable, requires the *a priori* inversion of a simple linear system once for a given basis, grows linearly with the signal’s frequency content, and is capable of *perfect* reconstruction.

Spherical Image Filtering. Spherical image processing approaches have recently been applied to (spherical) image stitching problems in panorama composition and astrological image processing. Kazhdan and colleagues [Kazhdan et al. 2010; Kazhdan and Hoppe 2010] adapt multi-grid Poisson equation solvers to the spherical domain, investigating the suitability of many of the aforementioned spherical representations for these tasks.

Another common application of spherical filtering is for realistic image synthesis: BRDFs are 4D functions that model the reflectance of a material, where every view-dependent 2D “slice” corresponds to a spherical kernel; the shading at a point corresponds to an integration of a BRDF slice against (spherical) incident illumination. Many of the basis representations outlined above have been applied to this problem, each with their own advantages and disadvantages. SH representations have been used in the context of *pre-computed radiance transfer* (PRT) techniques [Sloan et al. 2002; Kautz et al. 2002] and, in the special case of a single circularly symmetric filter kernel, they can leverage the Funke-Hecke theorem to accelerate the convolution between the BRDF slice and incident illumination [Ramamoorthi and Hanrahan 2001; Ramamoorthi and Hanrahan 2002; Kautz et al. 2004]. Spherical and warped planar wavelets have also been applied to PRT but, in order to resolve the general filtering problem, they require a brute-force tabulation of re-oriented filter kernels [Ng et al. 2003; Ng et al. 2004; Ma et al. 2006]. SGs have also been applied to related problems in interactive rendering [Green et al. 2007; Han et al. 2007; Wang et al. 2009; Xu et al. 2013], however they admit costly and unstable fitting solutions and are generally unable to accurately represent arbitrary kernels.

Křivánek and Colbert propose an approximate, but efficient solution to the spherical filtering problem for unshadowed shading from environment maps. Their approach uses a small number of lookups into pre-filtered distant illumination maps, where sampling directions are driven by BRDF importance sampling [Křivánek and Colbert 2008], and can be interpreted as a biased Monte Carlo estimator of the reflection equation. While the overarching premise of sampling into pre-filtered data overlaps our ISD rendering algorithms, we instead formulate a theory of isotropic filter decomposition that allows us to construct sampling patterns and pre-filtered data in order to perfectly filter band limited spherical functions. Finally, our computation method differs as it uses an explicit calculation instead of a Monte-Carlo estimator.

Filter Decomposition. Our approach is conceptually similar to steerable filters (SFs) [Freeman and Adelson 1991]. SFs decompose (potentially anisotropic) filters as a sum of oriented *anisotropic* filters, whereas we decompose spherical kernels using *isotropic* filters. Among other benefits, this allows us to prefilter our signal using only the (canonically oriented) isotropic kernels in order to apply arbitrary spatially-varying filters using simple sampling operations. We can avoid prefiltering altogether, exchanging performance and memory for accuracy, all while maintaining the same L_2 convergence behavior of our method.

Gastal and Oliveira [2012] evaluate filters at carefully chosen sample positions and interpolate these values in order to perform approximate high-dimensional filtering with time complexity $\mathcal{O}(dn^2)$ for d -dimensional signals, however they do not consider signals over the sphere. Unlike [Gastal and Oliveira 2012], any well-distributed set of alignment vectors will yield a suitable ISD decomposition. Isotropic filter decompositions have been applied to anisotropic texture filtering [McCormack et al. 1999] in the plane however, in the spherical domain, we leverage properties of zonal harmonics which will allow us to provide a similarly efficient




Fig. 1. Illustrating the ISD: we decompose an input filter (in purple) into a sum of displaced copies of a single isotropic mother kernel ϕ_i at each frequency band l . This allows us to efficiently filter the input signal with any rotated version of the anisotropic kernel, all with only a few lookups into prefiltered data.

filtering solution but with controllable error and without the need for any offline non-linear optimization.

One of our realizations of ISD in Section 5.1 generalizes recent work on zonal harmonic factorization [Nowrouzezahrai et al. 2012] that presented a similar zonal harmonics decomposition, but instead focused on transfer matrix sparsity and the exclusive application to SH rotation. Our contribution atop this existing work include a formalization of the spherical filtering problem with *arbitrary* rotationally symmetric functions and its generalization to any basis of functions, with ZH being simply an instance of the theory used as a reference basis. We additionally derive efficient transfer matrix computation and filtering for the general case problem, study the numerical stability across different bases, and present three novel applications: real-time shading of measured — both isotropic and anisotropic — materials, efficient filtering of sparse and noisy 4D BRDF data, and antialiasing of environment map mirror reflections.

3. OVERVIEW AND NOTATION

A spherical function f maps directions ω on the unit sphere Ω to real numbers (or any vector space, e.g. RGB colors). Applying a *general* spherical filter $k : \Omega \rightarrow \mathbb{R}$ to an input spherical signal f results in a function \hat{f} that maps frames from the 3D rotation group $\mathcal{SO}(3)$ to real numbers, $\hat{f} : \mathcal{SO}(3) \rightarrow \mathbb{R}$. Therefore, for a rotation \mathbf{R} , $\hat{f}(\mathbf{R})$ is the value of the filtered signal:

$$\hat{f}(\mathbf{R}) = \int_{\Omega} k(\mathbf{R}^{-1}\omega)f(\omega)d\omega \quad (1)$$

Estimating this integral numerically in real time is challenging: in the general case, prohibitively large (spherical) sampling rates are required for numerical integration techniques, and precomputing the filtered signal for all possible rotations \mathbf{R} necessitates an impractical amount of storage (and preprocessing time).

In the special case of a circularly symmetric filter about axis ω_a , namely, $k(\omega) = \kappa(\omega \cdot \omega_a)$ for all ω , we have

$$\hat{f}(\mathbf{R}) = \int_{\Omega} \kappa(\mathbf{R}^{-1}\omega \cdot \omega_a)f(\omega)d\omega = \int_{\Omega} \kappa(\omega \cdot \mathbf{R}\omega_a)f(\omega)d\omega, \quad (2)$$

and therefore $\hat{f}(\mathbf{R})$ is constant for all rotations that leave ω_a in place, leaving only two dimensions where $\hat{f}(\mathbf{R})$ actually varies. As such, \hat{f} is commonly parameterized by ω_a , and written as $\hat{f}(\omega_a)$. Here, the precomputation dimensionality is reduced and, in cases where the signal f can be expressed in a basis that allows analytical computation of the integral, real time computation is plausible.

Our main contribution is to propose a generic method for interactively computing the solution to Equation 1, leveraging a novel

$\mathcal{SO}(3)$	3D rotation space
\mathbb{R}	Real numbers
\mathbf{R}	Rotation matrix
$\langle f, g \rangle$	Spherical inner product
Ω	Unit sphere of directions
y_l^m	Real band- l SH basis function
N_l^m	SH basis function normalization term
π_l	notation for $\sqrt{4\pi/2l+1}$
P_l	Band- l Legendre polynomial function
P_l^m	Associated Legendre polynomial
M, M_l	Transfer matrices between zonal harmonics and SH
ω	Spherical direction in Ω : $(\sin \theta \cos \phi, \sin \theta \sin \phi, \cos \theta)$
k	Spherical kernel, mapping Ω to \mathbb{R}
\mathbf{z}	3rd vector of the canonical basis of \mathbb{R}^3
f, \hat{f}	Spherical signal, filtered by kernel k
$f \circ \mathbf{R}$	Rotated signal, expressed by composition of functions

Fig. 2. Summary of notation used in our paper.

isotropic decomposition of k and using isotropic basis functions that permit very efficient evaluation of Equation 2.

4. ISOTROPIC SPHERICAL DECOMPOSITION

Consider decomposing k into a weighted sum of basis functions ϕ^m with coefficients λ_m so as to express $\hat{f}(\mathbf{R})$ as a linear combination of spherical integrals of the basis functions with the signal:

$$\hat{f}(\mathbf{R}) = \sum_m \lambda_m \int_{\Omega} \phi^m(\mathbf{R}^{-1}\omega) f(\omega) d\omega. \quad (3)$$

In and of itself, this expression does not permit a more efficient computation of $\hat{f}(\mathbf{R})$, since precomputing (over the 3D space of rotations) the spherical convolution of the signal f with all basis functions ϕ^m is generally too costly. However, suppose that we can enforce two constraints on the form of the basis functions ϕ^m :

- (1) each ϕ^m is a rotated copy of a unique mother function ϕ : $\phi^m(\omega) = \phi(\mathbf{R}_m^{-1}\omega)$ where \mathbf{R}_m is a rotation, and
- (2) ϕ is rotationally symmetric about axis \mathbf{z} and so there exists a function g such that $\phi(\omega) = g(\omega \cdot \mathbf{z})$.

Under these assumptions, Equation 3 simplifies to:

$$\begin{aligned} \hat{f}(\mathbf{R}) &= \sum_m \lambda_m \int_{\Omega} \phi(\mathbf{R}_m^{-1}\mathbf{R}^{-1}\omega) f(\omega) d\omega \\ &= \sum_m \lambda_m \int_{\Omega} g(\mathbf{z} \cdot \mathbf{R}_m^{-1}\mathbf{R}^{-1}\omega) f(\omega) d\omega \\ &= \sum_m \lambda_m \int_{\Omega} g(\mathbf{R}\mathbf{R}_m\mathbf{z} \cdot \omega) f(\omega) d\omega \\ &= \sum_m \lambda_m (\phi \otimes f)(\mathbf{R}\mathbf{R}_m\mathbf{z}). \end{aligned} \quad (4)$$

Equation 4 now expresses the convolution as a linear combination of values sampled from the spherical convolution of ϕ and f and, since ϕ is circularly symmetric, this convolution results in simpler spherical function instead of a function on the rotation space.

Generalization to Many Frequency Bands. It is unlikely that ϕ_m functions obtained from a single mother function ϕ will span a space dense enough to correctly approximate all possible

kernels k : for instance, ϕ may be band-limited, and therefore rotated copies of ϕ will only be able to represent signals within the band-limit.

In order to generalize, we decompose filters k across individual frequency bands l where, for each band, we apply the preceding theory with several (rotated) isotropic kernels ϕ_l^m , each obtained from a common band- l band-limited “mother function” ϕ_l : $\phi_l^m(\omega) = \phi_l((\mathbf{R}_l^m)^{-1}\omega)$, where each \mathbf{R}_l^m is a rotation. The full *isotropic spherical decomposition* of k is thus (see Figure 1):

$$k(\omega) \approx \sum_{l=1}^L \sum_m \lambda_l^m \phi_l^m(\omega). \quad (5)$$

We introduce a band limiting approximation due to the projection of the kernel onto a finite set of harmonic basis functions. If the kernel has a natural band limit of L_{max} , the approximation error will reach zero at $L = L_{max}$, however for arbitrarily non-bandlimited kernels, the projection will converge to the correct $k(\omega)$ as L approaches infinity. We will propose specific isotropic mother functions ϕ_l^m and discuss their benefits and properties in Section 5.

Using the ISD. Following the derivations of Equation 4 with ϕ_l and ϕ_l^m , instead of ϕ and ϕ^m , leads to:

$$\hat{f}(\mathbf{R}) \approx \sum_{l,m} \lambda_l^m (f \otimes \phi_l)(\mathbf{R}\mathbf{R}_l^m\mathbf{z}) = \sum_{l,m} \lambda_l^m \hat{f}_l(\mathbf{R}\mathbf{R}_l^m\mathbf{z}) \quad (6)$$

where $\hat{f}_l = f \otimes \phi_l$ is the spherical convolution of the signal with each (per-band) isotropic *mother kernel* ϕ_l .

We note from Equation 6 that, once the \hat{f}_l functions are (pre)computed, evaluating the convolved signal \hat{f} for any rotation \mathbf{R} simply reduces to the (weighted) evaluation of \hat{f}_l at the rotated directions $\mathbf{R}\mathbf{R}_l^m\mathbf{z}$.

Choosing mother kernels ϕ_l that satisfy the ISD rotation and isotropy conditions might sound challenging, however we are afforded certain uncommon flexibilities: for example, while orthogonality of the ϕ_l^m would simplify the computation of the weights λ_l^m , it is actually not required when computing the integral. The next section we show that rotated zonal harmonics meet our ISD requirements, in addition to providing a very efficient decomposition scheme suitable for any other potential function bases.

5. ISOTROPIC SPHERICAL BASES

Of the common representations for spherical signals reviewed in Section 2, including spherical wavelets [Schröder and Sweldens 1995], frequency-space representations (e.g. SH), and SGs [Mardia and Jupp 2009], only SGs and the zonal harmonic (ZH) subset of SH fulfil our isotropy requirement; of these two, only the latter affords optimal L_2 signal representation *and* rotational-invariance, motivating our initial investigation below (Section 5.1).

5.1 Rotated Zonal Harmonics (RZH)

We build upon the ZH subset of SH basis functions, motivated by recent work on *zonal harmonic factorization* [Nowrouzezahrai et al. 2012], as they afford both frequency-localization and isotropy (i.e., circular/rotational symmetry). We begin with a brief overview of SH, before deriving a non-orthogonal basis that satisfies the ISD requirements presented in Section 4.

Spherical Harmonics. Real SH basis functions are defined as:

$$y_l^m(\omega) = \begin{cases} \sqrt{2}N_l^m P_l^m(\cos\theta) \cos(m\phi) & \text{if } m > 0 \\ N_l^m P_l^0(\cos\theta) & \text{if } m = 0 \\ \sqrt{2}N_l^{-m} P_l^{-m}(\cos\theta) \sin(-m\phi) & \text{if } m < 0 \end{cases} \quad (7)$$

where basis functions are indexed by frequency band l and function index m , N_l^m is a normalization term, and P_l^m are the Associated Legendre Polynomials. A full degree- L SH expansion of a function includes all bands $l \leq L$, and the $2l+1$ basis functions (with $-l \leq m \leq l$) in each band, for a total of $s = (L+1)^2$ terms.

Each band l contains functions of fixed bandwidth in the (spherical) frequency domain, and the $2l+1$ band- l basis functions span the space of l -band-limited spherical functions. Unfortunately, we cannot directly utilize band- l basis functions to satisfy ISD since they are not all isotropic. However, Nowrouzezahrai et al. [2012] show that the same band- l space is spanned by rotated copies of the band- l zonal harmonic y_l^0 , which happens to be rotationally symmetric w.r.t. the vertical axis \mathbf{z} . As such, we choose the band- l ZH as our mother kernels: $\phi_l = y_l^0$.

ISD with Rotated Zonal Harmonics. Given this choice, the two remaining steps to define our ISD are: determining the rotations \mathbf{R}_l^m for our isotropic kernels $\phi_l^m = y_l^0((\mathbf{R}_l^m)^{-1}\omega)$, and computing the decomposition weight λ_l^m . We will see shortly that our choice of ZHs for the mother kernel will permit an analytic solution for the decomposition weights.

To simplify our exposition, we first consider a single band l and later extend our formulation (trivially) to degree- L reconstructions. Let $\mathbf{Z} = \{\mathbf{z}_l^{-l}, \dots, \mathbf{z}_l^l\}$ be a set of $2l+1$ unit vectors in Ω . Let \mathbf{R}_l^m be a rotation that aligns the \mathbf{z} axis onto \mathbf{z}_l^m , we have¹

$$\phi_l^m(\omega) = y_l^0((\mathbf{R}_l^m)^{-1}\omega). \quad (8)$$

Any non-degenerate set of directions \mathbf{Z} yields an isotropic kernel basis (with elements as in Equation 8) that spans the same space as band- l SH functions [Nowrouzezahrai et al. 2012], while satisfying ISD's isotropy and displacement properties.

Given the isotropic kernels ϕ_l^m (with their corresponding symmetry axis \mathbf{z}_l^m), we need only determine weights λ_l^m to be able to decompose any spherical kernel k according to Equation 5. We first express the l -band-limited component k_l of k with its expansion in the band- l SH basis as $k_l(\omega) = \sum_{m=-l}^l c_l^m y_l^m(\omega)$. Here, the band- l projection coefficients $c_l^m = \langle k_l, y_l^m \rangle$ are computed using analytic or numerical integration, depending on the form of k (i.e., we support both analytic and tabulated filters; see Section 6.1).

Once we compute the filter's (band- l) SH projection coefficients, we devise a linear mapping between the c_l^m and ISD weights λ_l^m by leveraging the SH addition theorem [Hobson 1931] that expresses rotated (band- l) ZHs as a weighted sum of (band- l) SH functions:

$$y_l^0((\mathbf{R}_l^m)^{-1}\omega) = \pi_l \sum_{m'=-l}^l y_l^{m'}(\mathbf{z}_l^m) y_l^{m'}(\omega), \quad (9)$$

where $\pi_l = \sqrt{4\pi/2l+1}$ is a convolution normalization coefficient. Combining Equations 8 and 9, and substituting into Equation 5, we obtain a matrix equation across all m functions in band- l , relating λ_l^m to c_l^m . Let $\mathbf{c}_l = [c_l^{-l}, \dots, c_l^l]$ denote the vector of k_l 's band- l SH projection coefficients, and $\mathbf{\Lambda}_l = [\lambda_l^{-l}, \dots, \lambda_l^l]$ denote the

¹An infinite number of such rotations exist and all fit the theory. In our implementation we normalize vectors $\mathbf{z} \times \mathbf{z} \times \mathbf{z}_l^m$, $\mathbf{z} \times \mathbf{z}_l^m$ and \mathbf{z}_l^m to form the rotation.




Fig. 3. Our RZH basis spans the same space than SH: although their respective projection coefficients are quite different, reconstructed BRDF lobes (gold-metallic-pain BRDF from the MERL dataset) are identical in both bases; however, the rotation operation into the local shading frame at each pixel is trivial to compute with our RZH representation, whereas SH requires a complex and costly rotation. Indeed, rendering with RZH is – in this example – **30 times faster** than with SH using the most efficient rotation method implemented on the GPU [Kautz et al. 2002].

vector of (unknown) ISD weights:

$$\mathbf{c}_l = \mathbf{M}_l \mathbf{\Lambda}_l, \text{ with } (\mathbf{M}_l)_{ij} = \pi_l y_l^i(\mathbf{z}_l^j), \quad (10)$$

In practice, we begin by distributing $2l+1$ directions \mathbf{z}_l^m (e.g. using low-discrepancy patterns on the sphere), and then precompute and store matrix \mathbf{M}_l and its inverse \mathbf{M}_l^{-1} . Given any function with known spherical harmonic coefficients c_l^m , we obtain its ISD at the cost of one matrix-vector product using:

$$k_l(\omega) = \sum_{m=-l}^l \lambda_l^m y_l^0(\mathbf{R}_l^m \omega) \quad \text{where } \mathbf{\Lambda}_l = \mathbf{M}_l^{-1} \mathbf{c}_l. \quad (11)$$

Note that, unlike the sampling locations of Gastal and Oliveira [2012], we can choose any distribution of the \mathbf{z}_l^m (that lead to a non-singular \mathbf{M}_l [Lessig et al. 2012]) and, with the resulting \mathbf{M}_l , the space spanned by our rotated ZH kernels is the same space spanned by the $2l+1$ band- l SH functions, by construction. Fortunately, only very few degenerate direction sets lead to a non-invertible \mathbf{M}_l , posing no difficulty in practice.

Extension to Order- L . The full $s = (L+1)^2$ -term spherical ISD is

$$k(\omega) \approx \sum_{l=0}^L \sum_{m=-l}^l \lambda_l^m \phi_l^m(\omega) = \sum_l \sum_m \lambda_l^m y_l^0(\mathbf{R}_l^m \omega), \quad (12)$$

and can be easily formed as a simple combination across the band- l spherical ISDs (Equation 11) for every band $l \leq L$. We are able to perform this concatenation since each band- l isotropic kernel set spans an isolated frequency-band-limited space and linear reconstruction across bands is L_2 -optimal from the properties of SH.

We similarly rewrite the system of equations in Equation 10 across all bands $l \leq L$, $\mathbf{c} = \mathbf{M} \mathbf{\Lambda}$, where the three terms are degree- L generalizations of their band- l counterparts. As such, \mathbf{M} (and $\mathbf{\Lambda}$) have block-diagonal structure with $(2l+1) \times (2l+1)$ matrix sub-blocks along the diagonal.

Note that it is possible to share the displacement directions \mathbf{z}_l^m across bands using the “lobe sharing” scheme presented by Nowrouzezahrai et al. [2012], which means that for all l and m , $\mathbf{z}_{l+1}^m = \mathbf{z}_l^m$, noted \mathbf{z}^m . We discuss the influence of the displacement direction choice on the stability of numerical calculations in Section 7.

Efficient Spherical Filtering with ISD. Recall from Equation 6 that the input signal must be convolved against the ISD mother kernels ϕ_l in order to compute the spherical kernel integrals. We (pre)convolve the band- l component of the input signal f with that band’s mother function $\phi_l = y_l^0$ in order to filter using Equation 6.

Our choice of mother kernels affords an interesting accelerated $\mathcal{O}(L)$ time complexity algorithm for computing each of the L (pre)convolutions: given the SH coefficients f_l^m of the signal, we compute the SH coefficients of the pre-filtered function, using the SH convolution theorem [Ramamoorthi and Hanrahan 2004], as

$$\hat{f}_l^m = \pi_l f_l^m. \quad (13)$$

Interestingly, these pre-convolved coefficients simply correspond to a per-band scaling of its frequency spectrum (i.e., its SH coefficients).

Depending on the application, the pre-convolved functions can be evaluated directly (as necessary for filtering in Equation 6) using either a pre-tabulation e.g. in a cubemap or simply computing its SH expansion on the fly as

$$\hat{f}_l(\omega) = \sum_m \hat{f}_l^m y_l^m(\omega) = \pi_l \sum_m f_l^m y_l^m(\omega). \quad (14)$$

Drawing parallels to Steerable Filters (Section 2, [Freeman and Adelson 1991]), spherical extensions of SFs would amount to mapping SH basis functions onto (potentially rotated) SH basis functions, requiring $(L+1)^2$ (instead of our L) pre-convolved lookup functions.

Figure 3 illustrates a view-evaluated BRDF lobe from the MERL database that is decomposed onto SH and our RZH and, although the coefficient values are different, the two ($L=20$) approximations are identical, as predicted by our theory, and both are very close to the original data. We apply spherical ISDs to realistic interactive rendering in Section 6.1, where the environment lighting acts as a signal that is filtered by spatially- and orientation-varying BRDFs.

5.2 ISD with Arbitrary Mother Functions

In the general case, we devise a straightforward approach to realize ISD using any (non-ZH) mother kernel functions ϕ_l , such as Gaussian kernels. This could be interesting or necessary if, for example, to maintain positiveness in the pre-convolved signals or to leverage hardware-accelerated convolution filters. Fortunately, the general case is not more complicated than the ZH realization. While M is not likely to exhibit the (band-diagonal) structure we observed in the ZH case, its rows still rely only on simple evaluations of rotated SH coefficients, and so it remains easy to compute.

In this general case, M_{ij} are the coupling coefficients of circularly symmetric functions ϕ_l^m over the harmonic function $y_l^{m'}$, and

we adopt the single-indices $i = (l', m')$ and $j = (l, m)$, yielding:

$$\begin{aligned} M_{ij} &= \langle y_{l'}^{m'}, \phi_l^m \rangle \\ &= \langle y_{l'}^{m'}, \phi_l \circ (\mathbf{R}_l^m)^{-1} \rangle \\ &= \langle y_{l'}^{m'} \circ \mathbf{R}_l^m, \phi_l \rangle \\ &= \sum_{-l' \leq k \leq l'} W_{l'}^{m',k} ((\mathbf{R}_l^m)^{-1}) \langle y_{l'}^k, \phi_l \rangle \\ &= W_{l'}^{m',0} ((\mathbf{R}_l^m)^{-1}) \langle y_{l'}^0, \phi_l \rangle, \end{aligned}$$

where $W_{l'}^{m',k}((\mathbf{R}_l^m)^{-1})$ are the Wigner coefficient that express the projection of $y_{l'}^{m'}$, rotated by $(\mathbf{R}_l^m)^{-1}$, onto $y_{l'}^k$. To simplify the sum, we exploit the circular symmetry of ϕ_l that leads to only 1 non-zero projection coefficient at $k=0$.

In order to compute M_{ij} , we exploit the simplified form of the Wigner coefficients in this case:

$$W_{l'}^{m',0}((\mathbf{R}_l^m)^{-1}) = W_{l'}^{0,m'}(\mathbf{R}_m) = \pi_{l'} y_{l'}^{m'}(\mathbf{z}_l^m), \quad (15)$$

and we define the projection of the mother functions ϕ_l onto ZH as

$$\beta_{l'l} = \langle \phi_l, y_{l'}^0 \rangle. \quad (16)$$

Now we can write the expression for M_{ij} directly as:

$$M_{ij} = \beta_{l'l} \pi_{l'} y_{l'}^{m'}(\mathbf{z}_l^m). \quad (17)$$

This means that, in practice, we need only compute the coefficients $\beta_{l'l}$ of the ZH decomposition of each ϕ_l to obtain the elements of M : for example, we see that with $\phi_l = y_l^0$ we obtain $\beta_{l'l} = \delta(l-l')$, and M simplifies to the block-diagonal form

$$M_{ij} = \pi_{l'} y_{l'}^{m'}(\mathbf{z}_l^m) \delta(l-l') \quad (18)$$

as derived earlier in Equation 10.

We note an important conclusion: Equation 17 implies that any ISD realization (with any circularly symmetric mother kernels) can be specified completely by its set of $\beta_{l'l}$ coefficients and the directions \mathbf{z}^m at each level l . Moreover, the approximation error of the realization will be driven by how accurately the mother kernel functions are represented in SH. We therefore recommend choosing mother kernels ϕ_l at each level with increasing bandwidth coverage so that ϕ_l can be reasonably well represented by functions $y_{l'}^0$ with $l' \leq l$. In practice, any function set works provided that M can be inverted, but function sets that exhibit good discrepancy, i.e., uniform coverage in the frequency domain produce transfer matrices with better conditioning. Spherical Gaussians, for instance, appear to be a good practical choice if setting:

$$\phi_l(\omega) = e^{-(1-\omega \cdot \mathbf{z})l^2}. \quad (19)$$

Another metric to consider when designing ISDs is the resulting conditioning of M , which would affect the numerical stability of the decomposition as well as complicating the numerical pre-computation of M^{-1} (see Section 7.2 for more details).

6. APPLICATIONS AND RESULTS

We apply ISD to three common problems in realistic rendering: unshadowed and shadowed shading with complex BRDFs under environmental lighting captured from the real world, and filtering sharp mirror reflections over a pixel’s footprint; as a side-effect, we also develop the first strategy to reliably filter noisy and sparse measured anisotropic BRDFs from the MERL dataset. In each case we compare favorably to state-of-the-art approaches.




Fig. 4. Convergence of spherical ISD shading for increasing L . Top to bottom: isotropic alum-bronze (with `pisa` illumination), isotropic gold-paint and anisotropic yellow-satin (both using `grace cathedral` illumination). Reference images were ray-traced using 300K samples per pixel. False color error images in the bottom row visually illustrate the convergence of our RZH approximation to the reference rendering. Note that the dark region in the center of the spheres on the last row of renderings is indeed part of the underlying reflectance input data.

6.1 Accurate, Interactive Shading with Arbitrary BRDFs

The reflection of distant spherical incident lighting E at a point x , with BRDF $\rho(x, \omega_o, \omega)$, towards a viewer ω_o is

$$I(x, \omega_o) = \int_{\Omega} E(\omega) \rho(x, \mathbf{R}_n \omega, \mathbf{R}_n \omega_o) \max(\cos \theta, 0) d\omega, \quad (20)$$

where the incident and outgoing directions ω and ω_o are expressed in global (world) coordinates, the BRDF ρ is parameterized in the local (surface) coordinate frame at x , and the rotation $\mathbf{R}_n \in \mathcal{SO}(3)$ transforms from global- to local-coordinates at x (where \mathbf{n} is the normal at x and θ is the angle between \mathbf{n} and ω).

We treat each spatially-varying view-slice of the cosine-weighted BRDF as a filter, $k(\omega) = \rho(x, \omega, \omega_o) \max(\cos \theta, 0)$, and the environment map as our signal $f(\omega) = E(\omega)$. Following the process outlined in Section 5.1, we choose $2L+1$ well-distributed directions \mathbf{z}^m — for this application, we share the directions across bands (see Section 7) — and compute ISD weights for an ensemble of filters (one for each BRDF slice):

$$\rho(x, \omega, \omega') \max(\cos \theta, 0) = \sum_l \sum_m \lambda_l^m(x, \omega') \phi_l^m(\omega). \quad (21)$$

Substituting Equation 21 into Equation 20, or applying Equation 6, yields the ISD equation for shading (we drop x for brevity, and we drop the l subscript on ISD displacement directions $\mathbf{z}_l^m \equiv \mathbf{z}^m$ to

make explicit that “lobe sharing” permits us to store a single set of directions *across* frequency bands l):

$$\begin{aligned} I(\omega_o) &\approx \sum_{l=0}^L \sum_{m=-l}^l \lambda_l^m(\mathbf{R}_n \omega_o) \int_{\Omega} E(\mathbf{R}_n^{-1} \omega) \phi_l^m(\omega) d\omega \\ &\approx \sum_{l,m} \lambda_l^m(\mathbf{R}_n \omega_o) \underbrace{(E \otimes y_l^0)}_{\hat{f}_l(\omega) = \hat{E}_l(\omega)}(\mathbf{R}_n^{-1} \mathbf{z}^m) \\ &\approx \sum_{l,m} \pi_l \lambda_l^m(\mathbf{R}_n \omega_o) \sum_{k=-l}^l e_l^k y_l^k(\mathbf{R}_n^{-1} \mathbf{z}^m), \end{aligned} \quad (22)$$

where e_l^m are the SH projection coefficients of the lighting E . Here, again, any approximation error is due only to the band limiting of the finite sum. In practice, the approximation converges to the exact value of $I(\omega_o)$ as L tends to the first band limit of either the BRDF or the lighting.

The (spatially- and view-varying) filter weights $\lambda_l^m(x, \omega')$ can be pre-tabulated for data-driven BRDFs (e.g. the MERL BRDF dataset [Matusik et al. 2003]), or computed analytically from their analytic SH coefficients for phenomenological BRDFs e.g., Lambertian and Phong BRDFs. We illustrate rendering results for several environment maps with spatially-varying and anisotropic BRDFs (Figures 9 and 5) of varying glossiness, as well as demonstrating convergence with ground truth renderings computed using importance-sampling Monte Carlo integration (Figure 4). Com-

pared to recent *anisotropic spherical Gaussians*, our approach generates more accurate renderings, with high performance, even with noisy input data (Figure 11).



Fig. 5. We illustrate a variety of isotropic materials for various scenes and lighting environments. All renderings were captured in real-time with performance ranging from 30 to 70 Hz.

Radial Reflection Maps for General BRDFs.

In the special case of distant illumination and BRDF lobes each comprised *exclusively* of radially symmetric light sources, Stone has shown that low-dimensional precomputed reflection maps can be leveraged for high-performance shading [Stone 2009]. By first

reformulating this approach in the context of ISD, we will show how to further reduce its storage requirements as well as generalizing it to arbitrary BRDF models at effectively no extra cost.

We decompose radially symmetric lighting E along axis ζ onto ZH,

$$E(\omega) = \sum_l e_l^0 y_l^0(\mathbf{R}_\zeta^{-1}\omega), \quad (23)$$

where \mathbf{R}_ζ aligns ζ along \mathbf{z} . Stone precomputes the convolution of $\rho(\omega) \cos \theta$ with E , for BRDFs that are sums of radially symmetric lobes, and demonstrates his approach for simple Phong lobes. This bears close similarity to ISD; Equation 22 boils down to

$$I(\omega_o) = \sum_{l,m} \pi_l \lambda_l^m(\mathbf{R}_n \omega_o) e_l^0 y_l^0(\mathbf{R}_\zeta^{-1} \mathbf{R}_n^{-1} \mathbf{z}^m), \quad (24)$$

in this case and, instead of precomputing and storing the convolutions, we can evaluate it on the fly using the SH coefficients e_l^0 , hence generalizing Stone’s work to any form of BRDF.

All Frequency Shading with Visibility

We extend Equation 22 to include the effects of self-occlusion by incorporating the visibility into the shading integral (in world coordinates, and dropping x for brevity):

$$I(\omega_o) = \int_{\Omega} E(\omega) v(\mathbf{R}_n \omega) \rho(\mathbf{R}_n \omega, \mathbf{R}_n \omega_o) \max(\cos \theta, 0) d\omega, \quad (25)$$

where v is the visibility function expressed in the local coordinate system at x .

By construction, our framework is compatible with rotationally symmetric spherical visibility representations such as the Spherical Signed Distance Function (SSDF) proposed by Wang et al. [2009], provided that ϕ_l has a local support. For this, the most appropriate choice for ϕ_l is a spherical gaussian, since one can then directly reuse [Wang et al. 2009, Eq. 31]. However, we present a new formulation based on zonal harmonics that generalizes it to visibility representations without local support and with provable convergence behavior.

We re-use the development of Equation 22, but now substituting the distant illumination for visibility, and the BRDF with the product of the distant illumination and BRDF, all expressed in the global coordinate system of the rendered object. For this, we impose two additional requirements:

- (1) the visibility v must be projected into SH at each vertex of the mesh, in the global coordinate system of the object, and
- (2) we must be able to express the RZH decomposition of $E \times \rho$, also in global (object) coordinates.

We can obtain these two quantities without the need for any costly SH rotation. We show that (see Appendix A):

$$\rho(\mathbf{R}_n \omega_o, \mathbf{R}_n \omega) E(\omega) = \sum_{l,m} y_l^0(\mathbf{R}_m^{-1} \omega) s_l^m,$$

where

$$s_l^m = \sum_{l',m'} \gamma_{l',m'}^{l,m} \rho_{l'}^{m'} \text{ and } \rho_{l'}^{m'} = \pi_{l'} \sum_m \lambda_{l'}^m(\mathbf{R}_n \omega_o) y_{l'}^m(\mathbf{R}_m^{-1} \mathbf{R}_n \mathbf{z}).$$

The coefficients $\gamma_{l',m'}^{l,m}$ represent the pre-computed projection of $y_l^m \times E$ onto the rotated zonal harmonics $y_{l'}^0(\mathbf{R}_m^{-1} \omega)$. We also pre-




Fig. 6. Left to right (Zoom-ins in reading order): no antialiasing, FXAA antialiasing, antialiasing with ISD and isotropic Gaussian kernels, and 16×16 antialiasing. Although only approximate (see areas of high-frequency geometric variations), our results are visually pleasing and more closely match ground truth at a fraction of its cost.

compute and store the SH coefficients of visibility v_l^m in the object's global coordinate frame such that:

$$v(\omega) = \sum_{l,m} v_l^m y_l^m(\omega).$$

When computing $I(\omega_o)$ at each pixel, we perform the following steps, none of which requires any explicit SH rotation operation (as a direct result of our RZH formulation):

- (1) compute the vector of $(L+1)^2$ values ρ_l^m ,
- (2) multiply that vector by the precomputed matrix containing the $(L+1)^2 \times (L+1)^2$ values $\gamma_{l',m'}^{l,m}$, to obtain the s_l^m , and
- (3) apply Equation 22, replacing e_l^k by v_l^k and $\lambda_l^m(\mathbf{R}_n \omega_o)$ by s_l^m .

In these steps, the only pre-computation needed is the static coefficients $\gamma_{l',m'}^{l,m}$ and the SH decomposition of visibility. In practice, we store SH coefficients for v in a texture atlas, and spatially interpolate between these values since they are all expressed in the same (object-space) global coordinate system. Figure 7 illustrates the results and performance of our GPU implementation of the aforementioned shading process.

The shader to perform these steps effectively involves evaluating the shader for Equation 22 twice (i.e., once for step 1 without addition and once for step 3 with addition), separated by a matrix-vector product for step 2.

6.2 Antialiasing Environment Map Reflections

Anti-aliasing mirror reflections of an environment mapping is typically accomplished using simple multi-sampling schemes in screen space followed by screen-space reconstruction filters e.g., a Gaussian filter g . These solutions scale linearly with the super-sampling rate and e.g. a 16×16 anti-aliased result would require as much as 256 times the cost of calculating the original (aliased) image, especially on the GPU where non coherent texture lookups tend to incur heavy performance penalties due to reduced cache coherence.

We instead *directly* formulate the anti-aliased reflection at a pixel p as an integral of the distant illumination and the set of mirror reflection directions over the surface projected onto the pixel's footprint, weighted by a Gaussian filtering/reconstruction weight g over p :

$$I(p) = \int_{|\Delta p| < 1} g(\Delta p) E(r(\mathbf{n}(p + \Delta p), \omega_o)) d\Delta p, \quad (26)$$

where ω_o is the view direction through the center of pixel p (supposed constant across pixel p), $\mathbf{n}(p')$ is the normal of the surface at

point p' inside the pixel's footprint Δp , and $r(\mathbf{n}, \omega) = \omega - 2(\mathbf{n} \cdot \omega)\mathbf{n}$ is the mirror reflection operator.

We use a second-order approximation of r based on the screen-space shape operator $\nabla \mathbf{n}$ that denotes the orthogonal displacement of the reflected direction w.r.t. ω_i as a function displacements within the pixel, where H is a matrix that converts this 2D displacement into a 3D vector in world space. Applying this approximation to Equation 26 yields:

$$I(p) \approx \int_{|\Delta p| < 1} g(\Delta p) E(\omega_i + H \nabla \mathbf{n}(\Delta p)) d\Delta p, \quad (27)$$

and changing integration variables to spherical solid angles gives:

$$I(p) \approx \int_{\Omega} \underbrace{g((\nabla \mathbf{n})^{-1} H^+ (\omega - \omega_i))}_{g_p(\omega)} E(\omega) |\det(\nabla \mathbf{n})| d\omega \quad (28)$$

Equation 28 states that the reflected, anti-aliased pixel color is the result of spherically filtering the incident illumination with an anisotropic Gaussian filter g_p with a pixel-dependent profile; for nearly flat surfaces, $\nabla \mathbf{n}$ has very small eigenvalues and g_p approximates a Dirac, as expected.

With ISDs, we can compute $I(p)$ (Equation 28) in several ways:

- (1) decomposing g_p into RZH and apply Equation 22, precomputing λ_l^m coefficients for various anisotropic Gaussian profiles;
- (2) decompose g_p into any other isotropic basis using Equation 17 e.g., Von-Mises functions, and shading using lookups into pre-filtered environment maps;
- (3) approximating g_p with a single isotropic Gaussian ϕ_l aligned along the reflected direction ω_i , which we then decompose onto zonal harmonics up to degree L using Equation 16 as

$$I(p) = \sum_{l'=0}^L \beta_{l'} f \otimes y_l^0(\omega_i) = \sum_{l'=0}^L \beta_{l'} \pi_{l'} \sum_{m=-l'}^{l'} e_{l'}^m y_{l'}^m(\omega_i). \quad (29)$$

We found experimentally that the orientation and flatness of the Gaussian filter are likely to vary significantly from pixel to pixel, and so option #3 appears to be the good choice, even if it introduces an additional approximation. Its computation cost is equal to performing a simple SH reconstruction, and it requires no storage other than the precomputed $\beta_{l'}$ and $e_{l'}^m$ coefficients. Rendering is rather fast since Eq.29 has one less nested loop than Eq.22, and avoids memory lookups in the illumination map. We show the results of this approach in Figure 6 and supply shader code for computing matrix $\nabla \mathbf{n}$ in our supplemental material.




Fig. 7. Shading with visibility, and comparisons to shading without visibility and to a ray-traced reference. Note the replacement of unshadowed highlights on the back of the dragon by shadows, and the darkening in the interior of the mouth. Illumination: Grace cathedral. Material: gold-paint.

6.3 Filtering Sparse/Noisy Anisotropic BRDFs

Anisotropic BRDFs often come as noisy and sparse data (such as the four anisotropic BRDFs of the MERL database), making them difficult to use when shading with high-frequency distant illumination. We use our framework to prefilter the data in the $\Omega \times \Omega$ domain with a separable 4D von Mises Fisher (vMF) kernel, obtaining very promising results:

$$\tilde{\rho}(\omega_i, \omega_o) = (\rho \otimes m)(\omega_i, \omega_o) \text{ with } m(\omega_i, \omega_o) = K e^{\kappa(\omega_i \cdot \mathbf{z} + \omega_o \cdot \mathbf{z})},$$

where K normalizes the filter integral over $\Omega \times \Omega$. This operation is analogous to convolving a 2D image with a separable Gaussian kernel to remove noise, however in the spherical domain, and we project the 4D BRDF onto SH with respect to both ω_i and ω_o (first project the ω_i dimension, then each ω_o “slice”),

$$\rho(\omega_i, \omega_o) = \sum_{l_i, m_i} \sum_{l_o, m_o} C_{l_i, l_o}^{m_i, m_o} y_{l_i}^{m_i}(\omega_i) y_{l_o}^{m_o}(\omega_o). \quad (30)$$

We can now simply apply Equation 13 to convolve against the vMF kernel, multiplying the coefficients by the ZH coefficients $\{k_l\}_{l \geq 0}$ of the vMF given by:

$$k_l = \int_{\Omega} (\sqrt{K})^{-1} e^{\kappa \omega \cdot \mathbf{z}} y_l^0(\omega) d\omega. \quad (31)$$

The final, filtered BRDF we use for shading is:

$$\tilde{\rho}(\omega_i, \omega_o) = \sum_{l_i, m_i} \sum_{l_o, m_o} \frac{k_{l_i} k_{l_o}}{\pi_{l_i} \pi_{l_o}} C_{l_i, l_o}^{m_i, m_o} y_{l_i}^{m_i}(\omega_i) y_{l_o}^{m_o}(\omega_o). \quad (32)$$

By construction, this method produces an L_2 -optimal approximation of the BRDF data with band limited functions in the 4D domain. It is not surprising that this eliminates the noise in the dataset while retaining the low and medium frequency content of the BRDF (see Figure 11, for an example, shaded with a high-frequency environment map, for a typical anisotropic material of the MERL database). Furthermore, it preserves (again by construction) both the energy-conservation integration property of the BRDF, as well as its reciprocity.

7. IMPLEMENTATION, PERFORMANCE AND EXTENSIONS

We provide all implementation-specific details, outside of directly implementing Equations 22 and 29 in Section 6, necessary to reproduce our results, as well as discussing performance/memory requirements. We provide the source code of all our shaders used in the Applications Section in a supplemental document.

7.1 Implementation Details

For environment map shading with BRDF filters we implement Equation 22 in a simple GLSL shader (see supplemental material): for data-driven BRDF filters from the MERL BRDF dataset [Matusik et al. 2003], we pre-tabulate the ISD weights per (cosine-weighted) BRDF view-slice $\lambda_l^m(\omega_o)$, for each discrete BRDF, in a cubemap indexed by the view direction ω_o (in local surface coordinates); for analytic BRDFs (e.g., Lambertian, Phong) we hard-code analytic ISD weights. Equation 7 is not a numerically stable method for evaluating high-degree SH basis functions, and we implemented stable recurrence formulae for these evaluations based on the principles suggested by Sloan [2008] (also provided in our supplemental material).

A subtle, albeit important, shader implementation trade-off exists when implementing ISDs: when evaluating the preconvolved signals $\hat{f}_l(\omega)$ (Equation 14) in Equation 22, we can either pre-evaluate and store the functions in cubemaps (Equation 22, middle line) or explicitly compute the band- l SH expansion (Equation 22, last line). This amounts to a texture-lookup vs. ALU opcode trade-off in the shader design, and we found that the pre-convolved cubemaps were slower for our use-cases, probably due to cache misses in the texture lookups. They also incur a higher memory cost since, as we increase the L_{max} , the cubemap resolutions have to be increased to avoid numerical imprecisions. All our screen-captures in the paper (and accompanying videos) instead use the on-the-fly reconstruction solution.

7.2 Performance analysis

Speed and memory. All our results were captured on an Intel Xeon 1.2GHz with 24 GB of RAM and an NVIDIA GTX 780TI with 3 GB of VRAM. We provide performance statistics both in Hz and *filtering operations per second* (FOPS), and our video results were all captured interactively and given in Table I. The memory footprint takes into account the storage of the $2(L+1)^2$ coefficients λ_l^m and e_l^m in Equation 22, tabulated by the view direction in 1D (resp. 2D) for isotropic (resp. anisotropic) BRDFs.

Table I. Performance, memory statistics, for interactive environment shading with BRDFs (Section 6.1) at resolution 1024×768 . The scene used is the chair of Figure 5

$s = L^2 =$	8^2	12^2	16^2	31^2
Speed [Hz]	59	46	29	11
Speed [FOPS]	46.4M	36.2M	22.8M	8M
Mem.(isotropic) [MB]	0.1	0.2	0.4	1.5
Mem.(anisotropic) [MB]	12	24	48	192




Fig. 8. Qualitative comparison of our technique to shading with Spherical Gaussians, using three lobes [Wang et al. 2009] for the two materials `gold-paint` and `color-changing-paint3` from the MERL database. The lack of stability in the SG fitting process limits the number of basis functions that can be fit, hence reducing the accuracy of the result. At equal speed, our renderings are much closer to ground truth. All SG renderings used in our comparisons have been provided by the authors of [Wang et al. 2009] to match our supplied viewing parameters.




Fig. 9. Convergence of ISD rendering with zonal harmonics applied to the `brushed-aluminum` material from the MERL database [Matusik et al. 2003]. Our method offers more accurate reproduction of the material than anisotropic spherical Gaussians [Xu et al. 2013] which happen to share fitting parameters across viewing directions. Image resolution: 1024×768 . ASG parameters graciously provided by Xu et al. upon our request.

For environmental shading with complex BRDFs, we note that performance scales (sub)linearly with s . We illustrate consistent convergence to ground truth results in Figure 4, and rendering results with spatially-varying anisotropic BRDFs and varying BRDF glossiness (corresponding to different BRDF filter sizes) in Figures 5 and 9.

Comparison to Spherical Gaussians. In Figure 8 we compare our ISD shading to the SG technique proposed by Wang et al. [2009]. Since fitting a mixture of SGs is a numerically difficult and unstable, Wang et al. is limited to high-performance rendering of fits with low-accuracy and only a few SGs; in contrast, we are able to trivially increase the number of lobes in our representation, at a reasonable cost in performance, resulting in nearly pixel perfect renderings compared to ground truth computed using importance-sampled Monte Carlo integration. Indeed, our approach guarantees L_2 optimal reconstruction, whereas all existing SG approaches provide no such convergence guarantees.




Fig. 10. Two local spherical light sources light the dragon mesh with a `white-fabric2` BRDF in the Grace Cathedral (left), and anisotropic `brushed-aluminum` BRDF in St. Peter’s Basilica (right).

Table II. Comparison between our method and methods based on Gaussian approximations

	[Wang2009]	[Xu2013]	Our method
Visibility	SSDF	2D SSDF	ZH
Rotation	None	Warping	None
Fitting	Hard	Hard	Trivial
Accuracy	Low	Medium	High
Speed	Real time	Real time	Real time
Lobe Directions	Fitted	Fitted	Constant

We similarly compare ISD shading with anisotropic BRDFs to the recent *anisotropic SGs* work of Xu et al. [2013]. While, in some cases, 1 to 3 ASGs provide a sufficiently good approximation, this is not generally true for e.g. measured BRDFs, such as the `red-velvet` material illustrated in Figure 11. Here, ISDs render with significantly better visual fidelity compared to the results generated using the method of Xu et al. (with fitting data courtesy of the authors). Since ISDs do not rely on costly SH rotation nor (A)SG warping, we can afford to use a large value for L while maintaining fast rendering performance. Table II summarizes the key capabilities of ISD, SGs and ASGs.

Limitations. Rendering with ISDs and RZH inherits some of the limitations of SH rendering, as both families of functions span the same space. Extremely specular BRDFs usually require high L to display without Gibbs ringing artifacts, as illustrated in Figure 12. However, since ISDs completely avoid the local-to-global frame rotation during rendering, we are able to increase the maximum reconstruction degree L much higher than typical SH-based techniques, all while maintaining interactive performance. ISDs are particularly well-suited to complex, spatially-varying anisotropic




Fig. 11. ASGs cannot accurately represent the anisotropic measured data from the MERL dataset: ASG lobe fitting is costly and unstable when fitting many lobes. Furthermore, the ASG representation uses the same lobe profiles across viewing directions (according to the parameters provided to us courtesy of [Xu et al. 2013]). Although slower, ISD using RZH offers a much better accuracy and a flexible trade-off between accuracy and speed (illumination: Grace Cathedral; Material: purple-satin.)

filters with *large* extents: these types of filters were traditionally hard to treat with previous techniques.




Fig. 12. Zoom-in of our second result in Figure 8 for different values of L . Image contrast and luminance have been highly exaggerated in order to better illustrate Gibbs oscillations due to high frequencies in the BRDF; note, these artifacts are hardly visible beyond $L = 30$ and completely disappear at $L \approx 50$. Timings correspond to the full image rendering at 1024×768 . Material: color changing paint 3.

Numerical Stability. The conditioning of M gives important clues as to the stability of numerical evaluations of Equations 22 and 29. We compute the condition number of M for various mother function profiles ϕ_l as a function of increasing L , each time with and without shared directions between bandwidth levels (Figure 13): we used a Phong profile ($\phi_l(\omega) = (\omega \cdot \mathbf{z})^{4l}$), a Gaussian profile ($\phi_l(\omega) = e^{-(1-\omega \cdot \mathbf{z})^2}$), a Legendre profile corresponding to RZH ($\phi_l(\omega) = P_l(\omega \cdot \mathbf{z})$), and an hybrid profile (Gauss-Legendre) that is the product of the last two profiles in an attempt of forcing ZHs to have a more localized support. From this experiment we conclude that sharing directions across levels has little impact on the conditioning of ISD with RZH profiles, but has a high impact on numerical stability for Phong and Gaussian-based profiles. In these cases, we recommend not sharing directions across bandwidth levels in order to maintain numerically stable reconstruction. This behavior is simple to explain: for functions with local support, mother functions of successive levels will share similar profile shapes (i.e., the rows of M will be almost identical) if lobes are shared between levels. The global support of zonal harmonics (i.e., the Legendre profile) do not suffer from this behavior.

A similar issue exists in the frequency domain: when constructing an ISD basis from arbitrary functions (e.g. Gaussians), the conditioning of M is improved when these functions span disjoint regions of the frequency spectrum. Approximation errors will arise, in practice, when a region of the spectrum beyond is not covered by maximum frequency of the mother functions.

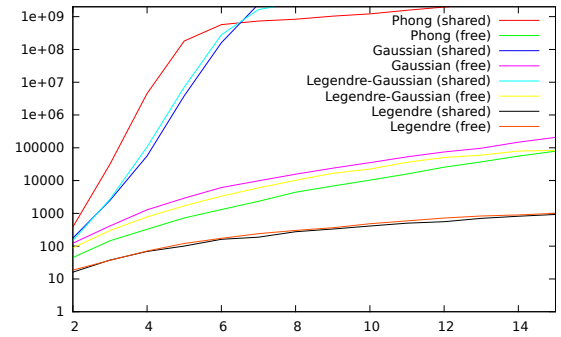


Fig. 13. Condition number of matrix M (with size $(L+1)^2 \times (L+1)^2$) for increasing L .

7.3 Extensions

Spherical Lights. Shading from point sources are best left to specialized techniques, however local spherical light sources fit nicely into our framework: we decompose these sources as a sum of zonal harmonics after aligning the source direction to \mathbf{z} with a proper rotation $\mathbf{R}_i(x)$ at each shading point x :

$$E(x, \omega) = \sum_l e_l y_l^0(\mathbf{R}_i^{-1} \omega), \quad (33)$$

where the e_l can be computed analytically as a function of (the ratio between) the source's size and distance to the x , and Equation 22 simplifies considerably since $e_l^k = 0$ for $k \neq 0$:

$$I(\omega_0) = \sum_{l,m} \sqrt{4\pi/2l+1} \lambda_l^m(\mathbf{R}_n \omega_0) e_l y_l^0(\mathbf{R}_i \mathbf{R}_n^{-1} \mathbf{z}^m). \quad (34)$$

Figure 10 illustrates spherical sources and anisotropic BRDFs shaded using ISD.

Parametric BRDFs. Since our method handles measured BRDFs, parametric models can be treated the same way by tabulating coefficients $\{\lambda_l^m(\mathbf{v})\}_{l,m}$ where \mathbf{v} is the vector of parameters for the BRDF. Naturally, simpler analytical BRDF models can be handled more efficiently using previous work, e.g. [Ramamoorthi and Hanrahan 2001]; this does not preclude their integration using ISD.

For analytic BRDFs, we can analytically determine the maximum required bandwidth level, in some cases. For a Phong lobe with exponent s , existing work [Ramamoorthi and Hanrahan 2004] reports that the maximum bandwidth is $O(\sqrt{s})$ which suggests that we should use $L_{max} = \sqrt{s}$. For example, renderings with $L_{max} = 30$ support glossy reflections on the order of $s = 900$.

Extension to Other Filtering Domains. We believe that ISDs can be extended beyond the spherical domain and we briefly discuss two other domains of interest, providing example applications for future work:

Lightfield Filtering. Instead of using distant illumination to shade an object with arbitrary BRDFs, one can imagine using a full 4D lightfield. Pre-filtering 4D lightfield data using tailored basis spaces is not currently possible, since each point on the object filters a different 2D slice of the 4D illumination. We believe that it could be possible to extend ISDs to use a set of 4D separable filters (e.g., high-dimensional Gaussians) to filter individual 2D slices of the lightfield with rotationally symmetric filters aligned along each shading point on the object. A variant of Equation 22 can be adopted to render objects lit entirely by the lightfield, using only look-ups into mipmapped 4D datasets.

Image Filtering. We have leveraged the properties of ZH to form a theory of efficient filtering using isotropic kernel decompositions on the sphere. We have additionally explored ZH-like extensions of ISDs to image filtering on the 2D plane, for image processing applications. Here, a naive solution would be to transform each rotated ZH using a sphere-to-plane isomorphism; this certainly generates a useable basis, but we found no such transform that produces isotropic functions that would remain steerable on the plane. However, we have experimented with a filter decomposition using a set of pre-shifted isotropic functions, such as displaced Gaussians, that was able to maintain all of the properties of our isotropic decomposition and filtering theory. Moreover, working in the plane afforded some additional flexibility in the lobe direction fitting, and we found it possible to optimize for sparsity (i.e., the number of non-zero coefficients) on such planar ISD decompositions, using compressive sensing methods [Wright et al. 2009]. Given a planar-ISD, we can efficiently apply “convolutions” with non-constant and anisotropic filters in the image domain.

8. CONCLUSION AND FUTURE WORK

We presented a general theory of Isotropic Spherical Decomposition (ISD) based on multi-frequency displaced isotropic kernels. Unlike many other representations, ours combines several important properties for efficient filtering: it is compact, affords a trivial and efficient rotation formulation, has controllable error vs. performance behavior, requires only a modest memory footprint, and is very efficient and trivially parallelizable.

We have fully developed a realization of ISD using ZH mother kernels, as well as outlining “recipes” for general ISD realizations. We have applied ISD with ZH and Gaussian mother kernel profiles to important problems in rendering: interactive shading with arbitrary BRDFs (including anisotropic BRDFs from the MERL dataset), and efficient anti-aliasing of mirror reflections. We compare favorably to state-of-the-art SG and ASG representations, as well as popular interactive anti-aliasing approximations.

As future work, we believe ISD-like decompositions could benefit filtering in other domains, and we briefly discussed two such applications: real time shading of objects under 4D mipmapped light-

field illumination, and image processing with anisotropic filters decomposed into e.g. Gaussian mixtures using compressive sensing.

We provide basic GPU implementations of our approach in supplemental material and we discuss all the relevant implementation details necessary to reproduce our work.

Acknowledgments

This work was supported in part by a grant “ALTA” ANR-11-BS02-006, an NSERC Discovery grant, and a MITACS Engage Cluster grant. The Bunny and Dragon models are provided by the Stanford Computer Graphics Lab data repository.

REFERENCES

- FEKETE, G. 1990. Rendering and managing spherical data with sphere quadtrees. In *Proceedings of the 1st Conference on Visualization '90. VIS '90*. IEEE Computer Society Press, Los Alamitos, CA, USA, 176–186.
- FREEMAN, W. T. AND ADELSON, E. H. 1991. The design and use of steerable filters. *IEEE Trans. Pattern Anal. Mach. Intell.* 13, 9 (Sept.), 891–906.
- GASTAL, E. S. L. AND OLIVEIRA, M. M. 2012. Adaptive manifolds for real-time high-dimensional filtering. *ACM TOG* 31, 4, 33:1–33:13. Proceedings of SIGGRAPH 2012.
- GREEN, P., KAUTZ, J., AND DURAND, F. 2007. Efficient reflectance and visibility approximations for environment map rendering. *Computer Graphics Forum* 26, 495–502.
- GREGER, G., SHIRLEY, P., HUBBARD, P. M., AND GREENBERG, D. P. 1998. The irradiance volume. *IEEE Comput. Graph. Appl.* 18, 2 (Mar.), 32–43.
- GRSKI, K. M., HIVON, E., BANDAY, A. J., WANDEL, B. D., HANSEN, F. K., REINECKE, M., AND BARTELMANN, M. 2005. Healpix: A framework for high-resolution discretization and fast analysis of data distributed on the sphere. *The Astrophysical Journal* 622, 2, 759.
- HAN, C., SUN, B., RAMAMOORTHI, R., AND GRINSPUN, E. 2007. Frequency domain normal map filtering. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2007)* 26, 3, 28:1–28:12.
- HEIDRICH, W. AND SEIDEL, H.-P. 1999. Realistic, hardware-accelerated shading and lighting. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques. SIGGRAPH '99*. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 171–178.
- HOBSON, E. 1931. *The theory of spherical and ellipsoidal harmonics*. The University press, Cambridge.
- KAUTZ, J., DAUBERT, K., AND SEIDEL, H.-P. 2004. Advanced environment mapping in vr applications. *Computers and Graphics* 28, 1, 99–104.
- KAUTZ, J., SLOAN, P.-P., AND SNYDER, J. 2002. Fast, arbitrary brdf shading for low-frequency lighting using spherical harmonics. In *Proceedings of the 13th Eurographics Workshop on Rendering. EGRW '02*. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 291–296.
- KAZHDAN, M. AND HOPPE, H. 2010. Metric-aware processing of spherical imagery. In *ACM SIGGRAPH Asia 2010 Papers. SIGGRAPH ASIA '10*. ACM, New York, NY, USA, 149:1–149:10.
- KAZHDAN, M., SURENDRAN, D., AND HOPPE, H. 2010. Distributed gradient-domain processing of planar and spherical images. *ACM Trans. Graph.* 29, 2 (Apr.), 14:1–14:11.
- KŘIVÁNEK, J. AND COLBERT, M. 2008. Real-time shading with filtered importance sampling. *Computer Graphics Forum - CGF* 27, 4, 1147–1154.

- LESSIG, C., DE WITT, T., AND FIUME, E. 2012. Efficient and accurate rotation of finite spherical harmonics expansions. *Journal of Computational Physics* 231, 2, 243–250.
- MA, W.-C., HSIAO, C.-T., LEE, K.-Y., CHUANG, Y.-Y., AND CHEN, B.-Y. 2006. Real-time triple product relighting using spherical local-frame parameterization. *The Visual Computer* 22, 9-11, 682–692.
- MALLAT, S. 2008. *A Wavelet Tour of Signal Processing, Third Edition: The Sparse Way*, 3rd ed. Academic Press.
- MARDIA, K. AND JUPP, P. 2009. *Directional Statistics*. Wiley Series in Probability and Statistics. Wiley.
- MATUSIK, W., PFISTER, H., BRAND, M., AND MCMILLAN, L. 2003. A data-driven reflectance model. *ACM Transactions on Graphics* 22, 3 (July), 759–769.
- MCCOOL, M. D., ANG, J., AND AHMAD, A. 2001. Homomorphic factorization of brdfs for high-performance rendering. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH '01. ACM, New York, NY, USA, 171–178.
- MCCORMACK, J., FARKAS, K. I., PERRY, R., AND JOUPEI, N. P. 1999. Simple and table feline: Fast elliptical lines for anisotropic texture mapping. Research report, HP labs.
- MILLER, G. S. AND HOFFMAN, C. R. 1984. Illumination and reflection maps: Simulated objects in real and simulated environment.
- NG, R., RAMAMOORTHY, R., AND HANRAHAN, P. 2003. All-frequency shadows using non-linear wavelet lighting approximation. In *ACM SIGGRAPH 2003 Papers*. SIGGRAPH '03. ACM, New York, NY, USA, 376–381.
- NG, R., RAMAMOORTHY, R., AND HANRAHAN, P. 2004. Triple product wavelet integrals for all-frequency relighting. In *ACM SIGGRAPH 2004 Papers*. SIGGRAPH '04. ACM, New York, NY, USA, 477–487.
- NOWROUZEZAHRAI, D., SIMARI, P., AND FIUME, E. 2012. Sparse zonal harmonic factorization for efficient sh rotation. *ACM Transactions on Graphics* 31, 3, 1–9.
- RAMAMOORTHY, R. AND HANRAHAN, P. 2001. An efficient representation for irradiance environment maps. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*. SIGGRAPH '01. ACM, New York, NY, USA, 497–500.
- RAMAMOORTHY, R. AND HANRAHAN, P. 2002. Frequency space environment map rendering. *ACM Trans. Graph.* 21, 3 (July), 517–526.
- RAMAMOORTHY, R. AND HANRAHAN, P. 2004. A signal-processing framework for reflection. *ACM Trans. Graph.* 23, 4 (Oct.), 1004–1042.
- SCHRÖDER, P. AND SWELDENS, W. 1995. Spherical wavelets: efficiently representing functions on the sphere. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*. SIGGRAPH '95. ACM, New York, NY, USA, 161–172.
- SLOAN, P.-P. 2008. Stupid spherical harmonics (sh) tricks.
- SLOAN, P.-P., KAUTZ, J., AND SNYDER, J. 2002. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. *ACM Trans. Graph.* 21, 3 (July), 527–536.
- STONE, J. 2009. Radially-symmetric reflection maps. In *SIGGRAPH 2009: Talks*. SIGGRAPH '09. ACM, New York, NY, USA, 24:1–24:1.
- WAN, L., WONG, T.-T., AND LEUNG, C.-S. 2007. Isocube: Exploiting the cubemap hardware. *IEEE Trans. Vis. Comput. Graph.* 13, 4, 720–731.
- WANG, J., REN, P., GONG, M., SNYDER, J., AND GUO, B. 2009. All-frequency rendering of dynamic, spatially-varying reflectance. *ACM Trans. Graph.* 28, 5, 133:1–133:10.
- WRIGHT, S. J., NOWAK, R. D., AND FIGUEIREDO, M. A. T. 2009. Sparse reconstruction by separable approximation. *Trans. Sig. Proc.* 57, 7 (July), 2479–2493.
- XU, K., JIA, Y.-T., FU, H., HU, S.-M., AND TAI, C.-L. 2008. Spherical piecewise constant basis functions for all-frequency precomputed radi-

ance transfer. *IEEE Transactions on Visualization and Computer Graphics* 14, 2, 454–467.

- XU, K., SUN, W.-L., DONG, Z., ZHAO, D.-Y., WU, R.-D., AND HU, S.-M. 2013. Anisotropic spherical gaussians. *ACM Trans. Graph.* 32, 6 (Nov.), 209:1–209:11.

APPENDIX

A. SHADING WITH VISIBILITY

We derive the decomposition of the BRDF-weighted distant illumination $\rho(\mathbf{R}_n\omega_o, \mathbf{R}_n\omega)E(\omega)$ onto rotated zonal harmonics. From the definition of λ_l^m in Equation 12 (and by sharing lobe directions), we have:

$$\begin{aligned} \rho(\mathbf{R}_n\omega_o, \mathbf{R}_n\omega) &= \sum_{l,m} \lambda_l^m(\mathbf{R}_n\omega_o) y_l^0(\mathbf{R}_m^{-1}\mathbf{R}_n\omega) \\ &= \sum_{l,m} \lambda_l^m(\mathbf{R}_n\omega_o) \sum_{m'} \pi_l y_l^{m'}(\mathbf{R}_m^{-1}\mathbf{R}_n\mathbf{z}) y_l^{m'}(\omega) \\ &= \sum_{l,m'} y_l^{m'}(\omega) \pi_l \underbrace{\sum_{m'} \lambda_l^m(\mathbf{R}_n\omega_o) y_l^{m'}(\mathbf{R}_m^{-1}\mathbf{R}_n\mathbf{z})}_{\text{Defined as } \rho_l^{m'}}. \end{aligned}$$

We define $\gamma_{l,m}^{l',m'}$ the RZH projection coefficient of $y_l^m E(\omega)$ onto the rotated zonal harmonic $y_{l'}^0(\mathbf{R}_{m'}^{-1}\omega)$ such that:

$$y_l^m(\omega)E(\omega) = \sum_{l',m'} \gamma_{l,m}^{l',m'} y_{l'}^0(\mathbf{R}_{m'}^{-1}\omega).$$

Then we have

$$\begin{aligned} \rho(\mathbf{R}_n\omega_o, \mathbf{R}_n\omega)E(\omega) &= \sum_{l,m} y_l^m(\omega) \rho_l^m E(\omega) \\ &= \sum_{l,m} \rho_l^m \sum_{l',m'} \gamma_{l,m}^{l',m'} y_{l'}^0(\mathbf{R}_{m'}^{-1}\omega) \\ &= \sum_{l',m'} y_{l'}^0(\mathbf{R}_{m'}^{-1}\omega) \underbrace{\sum_{l,m} \rho_l^m \gamma_{l,m}^{l',m'}}_{\text{Defined as } s_{l'}^{m'}}. \end{aligned}$$

In this form, it is very easy to express $\rho \times E$ in a different coordinate system with negligible computation cost, by simply pre-multiplying each $\mathbf{R}_{m'}$ by any additional rotation.