

Anisotropic Spherical Gaussians

Kun Xu¹ Wei-Lun Sun¹ Zhao Dong² Dan-Yong Zhao¹ Run-Dong Wu¹ Shi-Min Hu¹
¹TNList, Tsinghua University, Beijing ²Program of Computer Graphics, Cornell University




Figure 1: Comparison of the SG (Spherical Gaussian) based approximation with the ASG (Anisotropic Spherical Gaussian) based approximation in rendering a highly anisotropic metal dish, under an environment light and two local lights. The BRDF of the metal dish is approximated by different number of SGs or SGs in different images. Notice the superior property of ASGs over SGs. The result generated by 1 ASG already matches the path-traced reference well (with a L^2 error of 0.10%), and achieves a high framerate of 125 fps, while, to achieve a similar quality, more than 10 SGs are required, but with much lower framerates (19 fps for 13 SGs or 17 fps for 15 SGs). The L^2 error and the framerates for each configuration are also given in the corresponding subtitle.

Abstract

We present a novel *anisotropic Spherical Gaussian* (ASG) function, built upon the *Bingham distribution* [Bingham 1974], which is much more effective and efficient in representing anisotropic spherical functions than Spherical Gaussians (SGs). In addition to retaining many desired properties of SGs, ASGs are also rotationally invariant and capable of representing all-frequency signals. To further strengthen the properties of ASGs, we have derived approximate closed-form solutions for their *integral*, *product* and *convolution* operators, whose errors are nearly negligible, as validated by quantitative analysis. Supported by all these operators, ASGs can be adapted in existing SG-based applications to enhance their scalability in handling anisotropic effects. To demonstrate the accuracy and efficiency of ASGs in practice, we have applied ASGs in two important SG-based rendering applications and the experimental results clearly reveal the merits of ASGs.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Color, shading, shadowing, and texture

Keywords: Anisotropic Spherical Gaussians, Spherical Gaussians, Anisotropic BRDFs

Links: DL PDF WEB

ACM Reference Format

Xu, K., Sun, W., Dong, Z., Zhao, D., Wu, R., Hu, S. 2013. Anisotropic Spherical Gaussians. ACM Trans. Graph. 32, 6, Article 209 (November 2013), 11 pages. DOI = 10.1145/2508363.2508386 <http://doi.acm.org/10.1145/2508363.2508386>

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
Copyright © ACM 0730-0301/13/11-ART209 \$15.00.
DOI: <http://doi.acm.org/10.1145/2508363.2508386>

1 Introduction

Effective and compact representation of spherical function is beneficial for computer graphics applications, especially rendering. To achieve real-time rendering of complex reflectances (BRDFs) under real-world environmental lighting, recent approaches [Tsai and Shih 2006; Wang et al. 2009; Iwasaki et al. 2012a] adopt Spherical Gaussians (SGs) to effectively represent lighting, BRDFs or visibility for computing light transport. The reason why SGs have been chosen is due to their several desirable properties: First, SGs have varying support sizes (bandwidths) and are rotationally invariant, and hence it is convenient to use SGs to represent all-frequency signals, such as lighting and BRDF, and to rotate them freely; secondly, SGs have closed-form solutions for *integral*, *product*, and *convolution*, which are fundamental operators for evaluating rendering integrals [Kajiya 1986] and many other applications [Han et al. 2007; de Rousiers et al. 2012].

SGs are *isotropic*, or circularly symmetric around its lobe axis. Hence, to faithfully represent most real-world lightings or BRDFs, which are *anisotropic* to some degree, a mixture model of n scattered SGs, or simply an *SG Mixture*, is usually required and applied. Yet, since the SG mixture basis is not orthogonal, a product of two n -term SG mixtures has complexity $O(n^2)$ [Tsai and Shih 2006]. Therefore, using SGs to represent anisotropic functions always has to compromise between accuracy and performance, which is an intrinsic limitation.

To address this limitation, we present a novel anisotropic Spherical Gaussian (ASG) function based on the Bingham distribution [Bingham 1974], which can represent anisotropic spherical functions defined in arbitrary local frame (Section 3). To represent complex anisotropic functions, similar to SGs, a mixture model of scattered ASGs (*ASG Mixture*) needs to be applied. Due to the anisotropic nature of ASGs, a much smaller number of ASGs are usually enough to faithfully represent anisotropic functions, which leads to improvements in both accuracy and performance. Such an example is given in Fig. 1, where a single ASG is able to accurately render

a metal dish with a highly anisotropic BRDF, while 15 SGs are required to achieve a similar quality, but with much lower framerates.

Meanwhile, ASGs still retain those desired properties of SGs (Section 4). It is clear that, by definition, ASGs are still rotationally invariant, and are capable of representing all-frequency signals. Although the *integral*, *product* and *convolution* operators of ASGs have no exact closed-form solutions, we have derived approximate closed-form solutions for all of these operators. Through quantitative validations we show that these approximate solutions have nearly negligible approximation errors, and we expect to find application in fields beyond just computer graphics.

Supported by all these operators, ASGs can be adapted in existing SG-based applications, enhancing their scalability in handling anisotropic effects. To demonstrate the accuracy and efficiency of ASGs in practice, we develop an ASG-based rendering framework (Section 5) to implement two important applications: all-frequency rendering with dynamic BRDFs [Wang et al. 2009] and bi-scale BRDF editing [Iwasaki et al. 2012a] (Section 6), the promising experimental results of which reveal the merits of the ASGs.

2 Related Works

Spherical Gaussians (SGs) in Graphics. Spherical Gaussians (SGs), also known as von Mises-Fisher distribution [Fisher 1953], have been widely adopted in computer graphics to represent spherical functions, such as environment light, light transport functions, BRDFs, etc.

Tsai et al. [2006] used SGs to represent both environment light and light transport functions for all-frequency precomputed radiance transfer (PRT) [Sloan et al. 2002]. Relying on the PRT-based rendering routine, Green et al. [2006] approximated the light transport using sum of SGs, and achieved high-frequency view-dependent effects by non-linear interpolation of per-vertex SG parameters. These two PRT-based techniques mix reflectance and visibility in the precomputation stage, which limit the reflectance of rendered objects to be static (i.e. cannot be changed during rendering). Afterwards, Green et al. [2007] improved their previous method by decoupling visibility from BRDF.

To render dynamic, spatially-varying reflectance, Wang et al. [2009] approximated the microfacet BRDFs using SGs. This is done by first approximating the 2D Normal Distribution Function (NDF) using SGs, and each 2D BRDF slice at a specific view direction can then be approximated by SGs through a newly introduced *spherical warping* operator. However, due to the isotropy of SGs, it requires a large number of SGs to well approximate highly anisotropic BRDFs. Iwasaki et al. [2012a] proposed an SG-based bi-scale BRDF editing method, where large-scale BRDFs are considered as convolutions of small-scale BRDFs and the Bidirectional Visible Normal Distribution (BVNDF) [Wu et al. 2011]. This method represents both small-scale BRDFs and the BVNDF using SGs, and approximates convolved large-scale BRDFs again using SGs through a newly derived closed-form *convolution* operator. To validate the merits of ASGs in practice, we implement all-frequency rendering with dynamic BRDFs [Wang et al. 2009] and bi-scale BRDF editing [Iwasaki et al. 2012a] as two applications of ASGs.

More SG-based graphics applications include: Han et al. [2007] performed accurate normal map filtering by representing normal distribution function using SGs. Olano and Baker [2010] further proposed Linear Efficient Antialiased Normal (LEAN) mapping to achieve real-time filtering of specular highlights in normal maps. Laurijssen et al. [2010] presented a method for efficiently rendering indirect highlights by an SG-based approximation of BRDF lobes and an efficient algorithm to merge multiple SGs into one. Xu et al. [2011] achieved interactive rendering of hairs with dynamic

appearance by approximating hair scattering functions by circular Gaussian, which is a 1D variant of SG. De Rousiers et al. [2012] achieved real-time rendering of rough refractions by approximating the Bidirectional Transmittance Distribution Function (BTDF) using SGs. Yan et al. [2012] introduced a method for accurately rendering homogeneous translucent materials under SG lights, enabled by a derived closed-form integral of BSSRDF with SG lights. Iwasaki et al. [2012b] proposed *integral spherical gaussian*, which enables efficiently integrating SGs over spherical rectangles using summed-area tables.

Other Spherical Functions used in PRT. Besides SGs, a few other functions have been adopted to represent spherical functions in PRT, such as spherical harmonics (SH) [Ramamoorthi and Hanrahan 2001; Sloan et al. 2002], wavelet [Ng et al. 2003], spherical piecewise constant basis functions (SPCBF) [Xu et al. 2008], and pixel basis with biclustering approximation [Sun et al. 2011]. Specifically, SH can smoothly reconstruct low-frequency signals with a small number of coefficients, but they are inefficient in representing high-frequency functions, such as anisotropic lighting or BRDFs. The wavelet basis [Ng et al. 2003] or biclustering approximation [Sun et al. 2011] can well represent all-frequency functions, while temporal flickering or ghosting artifacts might be introduced due to the reconstruction errors of non-linear approximation. SPCBFs approximate the incident lighting based on its piecewise spherical partition, and hence cannot effectively handle high-frequency BRDFs. In comparison, ASGs inherit the nice properties of SGs, and can represent all-frequency functions and rotate them freely. Moreover, the anisotropic nature of ASGs guarantees their capability of representing anisotropic functions efficiently.

Anisotropic Appearance. Anisotropic appearances exhibit change with respect not only to the azimuthal differences between incoming and outgoing directions, but also to different incoming azimuthal directions. In general, most real-world BRDFs are more or less anisotropic, such as brushed metal, satin and hair. The anisotropic BRDF model was first introduced by Kajiya [1985]. A number of anisotropic parametric models have been proposed, including Ward's [1992] and Ashikhmin's [Ashikhmin and Shirley 2000] anisotropic BRDFs. There are also some newly-introduced anisotropic parametric BRDFs [Edwards et al. 2006; Kurt et al. 2010]. Compared with ASGs that support an arbitrary local frame, these parametric BRDF models are limited to a fixed normal-based local frame and hence cannot well represent complex anisotropic effects, such as off-specular highlight. A comprehensive accuracy analysis of parametric BRDF models is detailed in [Ngan et al. 2005].

The microfacet BRDF model [Torrance and Sparrow 1967; Cook and Torrance 1982] assumes that surfaces are composed of many mirror-like microfacets, each of which are purely reflective. Hence, the specular reflection of large-scale BRDFs can be computed based on the normal distribution function (NDF) of the microfacets. As demonstrated in [Ngan et al. 2005], the microfacet model can faithfully convey a large class of real-world reflectances. Recently, Pacanowski et al. [2012] introduced a general representation for BRDFs, called *Rational BRDF*, which can be used to reproduce some anisotropic effects.

Besides surface reflectance, anisotropic appearances are also found to be important in volume scattering [Jakob et al. 2010]. Various techniques have been proposed to model and render the anisotropic appearances of cloth and fabric [Irawan and Marschner 2012; Zhao et al. 2012; Sadeghi et al. 2013].

Directional Statistics. As aforementioned, the definition of ASG is built upon the *Bingham distribution* [Bingham 1974]. During the process of formalizing ASGs, we also experimented with other

widely-used anisotropic distributions in directional statistics, e.g., the *Fisher-Bingham distribution* [Mardia 1975] and the *Kent distribution* [Kent 1982]. However, for the Fisher-Bingham distribution, it is unclear how to efficiently (and analytically) compute its integral. Further, it has 8 parameters and hence requires a non-trivial fitting process and is expensive to compute. The Kent distribution is a simplified version of the Fisher-Bingham distribution, which contains only 5 parameters. However, the product of two Kent distributions is no longer a Kent distribution. In contrast, our Bingham-based ASGs permit approximate closed-form solutions for their integrals, products, and convolutions (see Section 4). More information about all these different distributions are detailed in the directional statistics book [Mardia and Jupp 1999].

3 Definition of an ASG

In this section we define a new family of functions that behave similar to SGs, but are *anisotropic* instead of *isotropic*.

Definition. We define the anisotropic Spherical Gaussian (ASG) to be the function of unit direction \mathbf{v} :

$$G(\mathbf{v}; [\mathbf{x}, \mathbf{y}, \mathbf{z}], [\lambda, \mu], c) = c \cdot S(\mathbf{v}; \mathbf{z}) \cdot e^{-\lambda(\mathbf{v} \cdot \mathbf{x})^2 - \mu(\mathbf{v} \cdot \mathbf{y})^2}, \quad (1)$$

where $\mathbf{z}, \mathbf{x}, \mathbf{y}$ are the *lobe*, *tangent* and *bi-tangent* axes, respectively, and $[\mathbf{x}, \mathbf{y}, \mathbf{z}]$ forms an orthonormal frame; λ and μ are the *bandwidths* for \mathbf{x} - and \mathbf{y} -axes, respectively, satisfying $\lambda, \mu > 0$. c is the *lobe amplitude*; the *smooth term* is defined as $S(\mathbf{v}; \mathbf{z}) = \max(\mathbf{v} \cdot \mathbf{z}, 0)$. The term $e^{-\lambda(\mathbf{v} \cdot \mathbf{x})^2 - \mu(\mathbf{v} \cdot \mathbf{y})^2}$ is denoted as the *exponential term*. Overall, Eq. 1 is referred to as the *geometric form* of an ASG.

This definition of ASG is based on the *Bingham distribution* [Bingham 1974] in directional statistics, which is derived from the intersection of a zero-mean, trivariate Gaussian distribution with the unit sphere. Specifically, the exponential term $e^{-\lambda(\mathbf{v} \cdot \mathbf{x})^2 - \mu(\mathbf{v} \cdot \mathbf{y})^2}$ of ASG is defined to be the same as the exponential term of the Bingham distribution $\mathbb{B}(\mathbf{v})$. Since the Bingham distribution is antipodally symmetric (i.e. $\mathbb{B}(\mathbf{v}) = \mathbb{B}(-\mathbf{v})$ and has two peaks at $\mathbf{v} = \pm \mathbf{z}$, respectively), we incorporate an additional smooth term $S(\mathbf{v}; \mathbf{z})$ to constrain its value in the upper hemisphere while preserving its smoothness. Fig. 2 gives several examples of ASGs with different parameters. Notice that the peak of ASG is at the lobe position $\mathbf{v} = \mathbf{z}$. In the following sections, for simplicity of notation, $G(\mathbf{v}; [\mathbf{x}, \mathbf{y}, \mathbf{z}], [\lambda, \mu], c)$ is denoted as $G(\mathbf{v})$ for short, and the amplitude c is assumed to be 1 if it is omitted.

Algebraic form. The ASG definition can also be equivalently written in algebraic form:

$$G(\mathbf{v}; \mathbf{A}) = S(\mathbf{v}; \mathbf{z}) \cdot e^{-\mathbf{v}^T \mathbf{A} \mathbf{v}}, \quad (2)$$

where \mathbf{A} is a 3×3 symmetric matrix, and \mathbf{z} is its eigen-vector with the smallest eigen-value. Eq. 2 is referred to as the *algebraic form* of an ASG. To prove the equivalence of these two forms (Eq. 1 and Eq. 2), we first apply a singular value decomposition (SVD) to the symmetric matrix \mathbf{A} in the algebraic form (Eq. 2):

$$\mathbf{A} = [\mathbf{x}, \mathbf{y}, \mathbf{z}] \cdot \text{diag}(\lambda_A, \mu_A, \nu_A) \cdot [\mathbf{x}, \mathbf{y}, \mathbf{z}]^T = \lambda_A \mathbf{x} \mathbf{x}^T + \mu_A \mathbf{y} \mathbf{y}^T + \nu_A \mathbf{z} \mathbf{z}^T = (\lambda_A - \nu_A) \mathbf{x} \mathbf{x}^T + (\mu_A - \nu_A) \mathbf{y} \mathbf{y}^T + \nu_A \cdot \mathbf{I},$$

where λ_A, μ_A , and ν_A are the three eigen-values of \mathbf{A} ; \mathbf{x}, \mathbf{y} , and \mathbf{z} are the corresponding eigen-vectors of \mathbf{A} . $\text{diag}(\cdot)$ represents a 3×3 diagonal matrix. The last term in the above equation has used the property that $\mathbf{x} \mathbf{x}^T + \mathbf{y} \mathbf{y}^T + \mathbf{z} \mathbf{z}^T = \mathbf{I}$ (\mathbf{I} is the 3×3 identity matrix). Without loss of generality, we can assume $\lambda_A \geq \mu_A \geq \nu_A$. Hence, Eq. 2 can be written in a geometric form $G(\mathbf{v}; [\mathbf{x}, \mathbf{y}, \mathbf{z}], [\lambda, \mu], c)$ (Eq. 1), where $c = \exp(-\nu_A)$, $[\lambda, \mu] = [\lambda_A - \nu_A, \mu_A - \nu_A]$. Similarly, it is also easy to convert the geometric form to the algebraic form.




Figure 2: Definition of Anisotropic Spherical Gaussians (ASGs). (a) An ASG with parameters $(\lambda = 4, \mu = 1)$. (b) Several ASG examples with parameters $(\lambda = 1, \mu = 1)$, $(\lambda = 10, \mu = 1)$, $(\lambda = 100, \mu = 10)$, and $(\lambda = 300, \mu = 3)$, respectively. Each ASG is illustrated by projecting its upper hemisphere to a disk using Lambert equal-area parameterization.

Relationship to traditional SGs. Our ASG is isotropic when the two bandwidths are the same (i.e., $\lambda = \mu$). However, such an isotropic ASG is not exactly the same as the traditional Spherical Gaussian (i.e. the von Mises-Fisher distribution), which is defined as $G_{iso}(\mathbf{v}; \mathbf{p}, \nu) = \exp(2\nu(\mathbf{v} \cdot \mathbf{p} - 1))$, where \mathbf{p} is the lobe axis (i.e. center), and ν is the bandwidth. However, in practice, they are approximately the same:

$$G_{iso}(\mathbf{v}; \mathbf{p}, \nu) \approx G(\mathbf{v}; [\mathbf{x}, \mathbf{y}, \mathbf{p}], [\nu, \nu]), \quad (3)$$

where \mathbf{x}, \mathbf{y} are two arbitrary orthogonal directions that form a local frame with \mathbf{p} . Detailed evaluation can be found in Section 1 of the supplemental document.

4 Supported Operators

Spherical Gaussians (SGs) possess several desirable properties for graphics applications. First, SGs are rotationally invariant, making them a good choice in representing spherical functions (e.g., lighting) that demand rotation; secondly, the support size (bandwidth) of SGs can be arbitrarily changed and hence they are capable of representing all-frequency signals; more importantly, SGs have closed-form solutions for *integral*, *product*, and *convolution*; These mathematical properties are crucial to make SGs especially suitable for rendering calculations.

Closed-form integrals of spherical distributions are necessary because the rendering process is essentially calculating an integral [Kajiya 1986]. Closed-form products are also important, since rendering often involves multiplications of different functions, such as lighting, BRDF and visibility. Closed-form convolutions provide fundamental support for several graphics applications, such as normal map filtering [Han et al. 2007], bi-scale BRDF calculation [Iwasaki et al. 2012a], and rough refraction [de Rousiers et al. 2012], in which large-scale effective BRDFs (or BTDFs) are obtained by convolving NDF with small-scale BRDFs (or BTDFs). Closed-form convolution is also potentially applicable for rendering indirect highlights [Laurijssen et al. 2010], in which calculations of indirect lighting can be formulated as convolutions of the BRDFs of two bouncing surfaces.

Compared with SGs, ASGs are more general and scalable in representing anisotropic functions. Meanwhile, ASGs still retain those desired properties of SGs. It is easy to see that, by definition, ASGs

are still rotationally invariant, and are capable of representing all-frequency signals. As for the *integral*, *product*, and *convolution* operators, ASGs do not have exact closed-form solutions. However, in the following derivations and validations, we will show and prove that ASGs have approximate closed-form solutions for all these operators, the accuracy of which is enough for graphics applications.

4.1 Integral of an ASG

In this subsection, we will derive an analytic approximation for the integral of an ASG over the unit sphere Ω , written as $\int_{\Omega} G(\mathbf{v}) d\mathbf{v}$. Without loss of generality, we assume $\mathbf{x} = (1, 0, 0)^T$, $\mathbf{y} = (0, 1, 0)^T$, $\mathbf{z} = (0, 0, 1)^T$, and represent \mathbf{v} in spherical coordinates: $\mathbf{v} = (\sin \theta \cos \phi, \sin \theta \sin \phi, \cos \theta)$, where $\theta \in [0, \pi/2]$, $\phi \in [0, 2\pi]$. The integral $\int_{\Omega} G(\mathbf{v}) d\mathbf{v}$ can be rewritten as:

$$\int_{\phi=0}^{2\pi} \left(\int_{\theta=0}^{\pi/2} e^{-\lambda(\sin \theta \cos \phi)^2 - \mu(\sin \theta \sin \phi)^2} \sin \theta \cos \theta d\theta \right) d\phi, \quad (4)$$

where the sine term is the Jacobian from spherical parameterization, and the cosine term is the smooth term of ASG. By denoting $k = \lambda \cos^2 \phi + \mu \sin^2 \phi$, it is easy to find that the inner integral over θ has an analytic solution as (the derivation can be found in Section 2 of the supplemental document):

$$\int_{\theta=0}^{\pi/2} e^{-k \sin^2 \theta} \sin \theta \cos \theta d\theta = \frac{1}{2k} (1 - e^{-k}). \quad (5)$$

By substituting it back, Eq. 4 can be rewritten as:

$$\begin{aligned} \int_{\Omega} G(\mathbf{v}) d\mathbf{v} &= \frac{1}{2} \int_0^{2\pi} \frac{1 - e^{-k}}{k} d\phi = \\ &= \frac{1}{2} \int_0^{2\pi} \frac{1}{\lambda \cos^2 \phi + \mu \sin^2 \phi} d\phi - \frac{1}{2} \int_0^{2\pi} \frac{e^{-\lambda \cos^2 \phi - \mu \sin^2 \phi}}{\lambda \cos^2 \phi + \mu \sin^2 \phi} d\phi. \end{aligned} \quad (6)$$

The following task becomes how to evaluate the two integrals in Eq. 6, respectively. First, it is easy to derive an analytic solution for the first integral:

$$\int_0^{2\pi} \frac{1}{\lambda \cos^2 \phi + \mu \sin^2 \phi} d\phi = \frac{4 \tan^{-1}(\sqrt{\frac{\mu}{\lambda}} \tan \phi)}{\sqrt{\lambda \mu}} \Big|_0^{\pi/2} = \frac{2\pi}{\sqrt{\lambda \mu}}. \quad (7)$$

Regarding the second integral in Eq. 6, unfortunately, it does not have an exact analytic solution. To evaluate it, we approximate the denominator by Gaussians (without loss of generality, here we assume $\lambda \geq \mu$):

$$\frac{1}{\lambda \cos^2 \phi + \mu \sin^2 \phi} \approx \frac{1}{\lambda} + \frac{\lambda - \mu}{\lambda \mu} e^{-\frac{\lambda - \mu}{\mu} \cos^2 \phi}. \quad (8)$$

The detailed derivations of Eq. 8 are presented in Section 3 of the supplemental document. Substituting Eq. 8 into the second integral in Eq. 6 gives:

$$\int_0^{2\pi} \frac{e^{-\lambda \cos^2 \phi - \mu \sin^2 \phi}}{\lambda \cos^2 \phi + \mu \sin^2 \phi} d\phi \approx \frac{e^{-\mu}}{\lambda} \left(F(\nu) + \frac{\nu}{\mu} F\left(\nu + \frac{\nu}{\mu}\right) \right), \quad (9)$$

where $\nu = \lambda - \mu$ is the difference between the two bandwidths, and the 1D function F is defined as $F(a) = \int_0^{2\pi} e^{-a \cos^2 \phi} d\phi$, which in fact is the Bessel function weighted by an exponential term. To efficiently evaluate F in practice, it can be precomputed and stored as




Figure 3: Evaluation of ASG integral approximation. The red curve gives the approximation using Eq. 11, and the blue curve gives the approximation using Eq. 10.

a 1D texture, or approximated by an analytic rational function (the formula can be found in Section 4 of the supplemental document). Finally, substituting Eq. 7 and Eq. 9 back into Eq. 6, the integral of an ASG $\int_{\Omega} G(\mathbf{v}) d\mathbf{v}$ can be approximated as:

$$\int_{\Omega} G(\mathbf{v}) d\mathbf{v} \approx \frac{\pi}{\sqrt{\lambda \mu}} - \frac{e^{-\mu}}{2\lambda} \left(F(\nu) + \frac{\nu}{\mu} F\left(\nu + \frac{\nu}{\mu}\right) \right). \quad (10)$$

In practice, when the two bandwidths λ, μ are not very small (i.e. when $\lambda, \mu > 5$), the integral can also be well approximated as:

$$\int_{\Omega} G(\mathbf{v}) d\mathbf{v} \approx \frac{\pi}{\sqrt{\lambda \mu}}, \quad (11)$$

since the coefficient of the second term (e.g. $e^{-\mu}/(2\lambda)$) becomes rather small.

Validation. To validate the accuracy of approximating the integral of an ASG using either Eq. 10 or Eq. 11, in Fig. 3, we plot the function curves of the original integral and these two approximation formulas. It is clear that the errors are subtle and both approximations are reasonable. We also quantitatively evaluate the bound of the approximation errors: when $\lambda, \mu > 5$, the relative error (using Eq. 10) is smaller than 0.027%, and the relative error (using Eq. 11) is smaller than 0.68%; when $\lambda, \mu > 1$, the relative error (using Eq. 10) is smaller than 5.2%.

Summary. The integral of an ASG can be efficiently evaluated using Eq. 10. When the two bandwidths satisfying $\lambda, \mu > 5$, the integral can be further efficiently approximated by Eq. 11.

4.2 Product of two ASGs

In this subsection, we will show that the product of two ASGs can still be well approximated by another ASG. Given two ASGs expressed by the algebraic form: $G(\mathbf{v}; \mathbf{A}_1) = \max(\mathbf{v} \cdot \mathbf{z}_1, 0) \cdot e^{-\mathbf{v}^T \mathbf{A}_1 \mathbf{v}}$ ($G_1(\mathbf{v})$ for short) and $G(\mathbf{v}; \mathbf{A}_2) = \max(\mathbf{v} \cdot \mathbf{z}_2, 0) \cdot e^{-\mathbf{v}^T \mathbf{A}_2 \mathbf{v}}$ ($G_2(\mathbf{v})$ for short), their product can be written as:

$$G_1(\mathbf{v})G_2(\mathbf{v}) = S(\mathbf{v}; \mathbf{z}_1, \mathbf{z}_2) \cdot e^{-\mathbf{v}^T (\mathbf{A}_1 + \mathbf{A}_2) \mathbf{v}}, \quad (12)$$

where $S(\mathbf{v}; \mathbf{z}_1, \mathbf{z}_2) = S(\mathbf{v}; \mathbf{z}_1) \cdot S(\mathbf{v}; \mathbf{z}_2)$ is a smooth function, and the latter term is essentially the exponential term of an ASG (as shown in Sec. 3). By denoting $\mathbf{A}_3 = \mathbf{A}_1 + \mathbf{A}_2$, and \mathbf{z}_3 as the eigen vector with the smallest eigen value of \mathbf{A}_3 , the above equation can be rewritten as:

$$\begin{aligned} G_1(\mathbf{v})G_2(\mathbf{v}) &= \frac{S(\mathbf{v}; \mathbf{z}_1, \mathbf{z}_2)}{\max(\mathbf{v} \cdot \mathbf{z}_3, 0)} \cdot \left(\max(\mathbf{v} \cdot \mathbf{z}_3, 0) e^{-\mathbf{v}^T \mathbf{A}_3 \mathbf{v}} \right) \\ &= \frac{S(\mathbf{v}; \mathbf{z}_1, \mathbf{z}_2)}{\max(\mathbf{v} \cdot \mathbf{z}_3, 0)} \cdot G(\mathbf{v}; \mathbf{A}_3) \approx S(\mathbf{z}_3; \mathbf{z}_1, \mathbf{z}_2) \cdot G(\mathbf{v}; \mathbf{A}_3). \end{aligned} \quad (13)$$

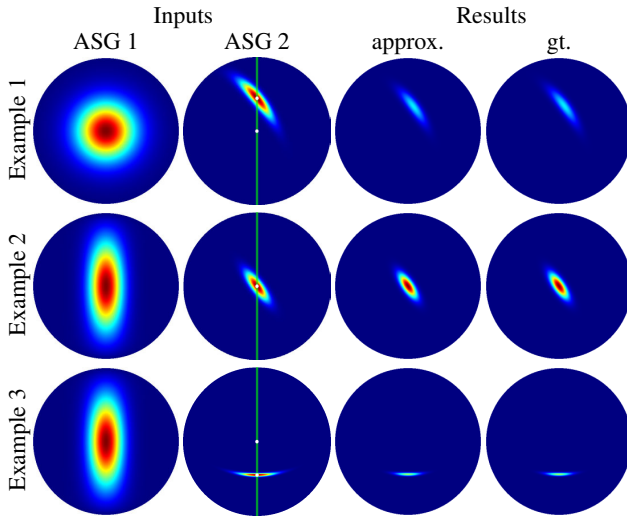


Figure 4: Product of two ASGs. The bandwidth parameters of the first ASG (λ_1 and μ_1) and the second ASG (λ_2 and μ_2) for all examples are listed below. Example 1: $\lambda_1 = \mu_1 = 3$, $\lambda_2 = 100$, $\mu_2 = 5$. Example 2: $\lambda_1 = 10$, $\mu_1 = 1$, $\lambda_2 = 100$, $\mu_2 = 10$. Example 3: $\lambda_1 = 10$, $\mu_1 = 1$, $\lambda_2 = 1000$, $\mu_2 = 10$.




Figure 5: L^2 Error of ASG product and convolution approximations. The error is measured by constraining the two input ASGs to be normalized (i.e., its integral as one).

Here, the smooth term $S(\mathbf{v}; \mathbf{z}_1, \mathbf{z}_2) / \max(\mathbf{v} \cdot \mathbf{z}_3, 0)$ is approximated as a constant of the function value when $\mathbf{v} = \mathbf{z}_3$, i.e., at the peak of the resulted ASG. This approximation is reasonable since the exponential term of ASG often changes much faster than the smooth term, and has been adopted extensively by existing methods [Wang et al. 2009; Xu et al. 2011; Iwasaki et al. 2012a].

Validation. Regarding the error of approximating the product of two ASGs with Eq. 13, the product of the exponential terms in the two ASGs is closed-form, and hence the error is solely due to the approximation of the low-frequency smooth term as a constant. In Fig. 4, we visualize and compare the approximated (approx.) product using a single ASG (Eq. 13) and the ground truth (gt.) product by taking three ASG pairs as examples. It is clear that the approximated results are visually identical to the ground truth. To more accurately identify the error, we further plot the quantitative approximation errors in Fig. 5 (left), of the three examples (Ex.) in Fig. 4. For each example, the error curve is generated by changing the lobe axis \mathbf{z}_2 of the second ASG $G_2(\mathbf{v})$ along the green line, which is shown in the second column of Fig. 4, to ensure the dot product of the two lobe axes \mathbf{z}_1 and \mathbf{z}_2 ranges in $[0, 1]$. Notice that the error is small for all these examples. Please refer to the supplemental document and accompanying video for more validation results.

Summary. The product of two ASGs can again be well approximated by a single ASG, as shown in Eq. 13.

4.3 Convolution of an ASG with an SG

In this subsection, we will derive the formula to compute the convolution of an ASG with an SG. Specifically, Given an ASG $G(\mathbf{v})$, and an SG with center \mathbf{p} and bandwidth ν : $G_{iso}(\mathbf{v}; \mathbf{p}, \nu) = e^{2\nu(\mathbf{v} \cdot \mathbf{p} - 1)} = e^{-\nu(\mathbf{v} - \mathbf{p})^2}$, their convolution $C(\mathbf{p}) = \int_{\Omega} G(\mathbf{v})G_{iso}(\mathbf{v}; \mathbf{p}, \nu) d\mathbf{v}$ is calculated as follows:

$$C(\mathbf{p}) = \int_{\Omega} S(\mathbf{v}; \mathbf{z}) e^{-\lambda(\mathbf{v} \cdot \mathbf{x})^2 - \mu(\mathbf{v} \cdot \mathbf{y})^2 - \nu(\mathbf{v} - \mathbf{p})^2} d\mathbf{v} \approx S(\mathbf{p}; \mathbf{z}) \int_{\Omega} e^{-\lambda(\mathbf{v} \cdot \mathbf{x})^2 - \mu(\mathbf{v} \cdot \mathbf{y})^2 - \nu(\mathbf{v} - \mathbf{p})^2} d\mathbf{v}. \quad (14)$$

Here, the smooth term S is pulled out of the integral as a constant function. We further approximate the center of the SG \mathbf{p} , and direction \mathbf{v} both by a Taylor expansion at $G(\mathbf{v})$'s peak \mathbf{z} : $\mathbf{p} \approx \mathbf{z} + (\mathbf{p} \cdot \mathbf{x})\mathbf{x} + (\mathbf{p} \cdot \mathbf{y})\mathbf{y}$, and $\mathbf{v} \approx \mathbf{z} + (\mathbf{v} \cdot \mathbf{x})\mathbf{x} + (\mathbf{v} \cdot \mathbf{y})\mathbf{y}$, yielding:

$$(\mathbf{v} - \mathbf{p})^2 \approx (\mathbf{v} \cdot \mathbf{x} - \mathbf{p} \cdot \mathbf{x})^2 + (\mathbf{v} \cdot \mathbf{y} - \mathbf{p} \cdot \mathbf{y})^2. \quad (15)$$

Substituting Eq. 15 into Eq. 14, gathering the terms of $\mathbf{v} \cdot \mathbf{x}$ and $\mathbf{v} \cdot \mathbf{y}$, and then pulling terms, which are independent of \mathbf{v} , out of the integral give:

$$C(\mathbf{p}) \approx M(\mathbf{p}) \int_{\Omega} e^{-(\lambda+\nu)(\mathbf{v} \cdot \mathbf{x} - \frac{\nu(\mathbf{p} \cdot \mathbf{x})}{\lambda+\nu})^2 - (\mu+\nu)(\mathbf{v} \cdot \mathbf{y} - \frac{\nu(\mathbf{p} \cdot \mathbf{y})}{\mu+\nu})^2} d\mathbf{v}, \quad (16)$$

where $M(\mathbf{p})$ is computed as:

$$M(\mathbf{p}) = S(\mathbf{p}; \mathbf{z}) \cdot e^{-\frac{\nu\lambda}{\nu+\lambda}(\mathbf{p} \cdot \mathbf{x})^2 - \frac{\nu\mu}{\nu+\mu}(\mathbf{p} \cdot \mathbf{y})^2}, \quad (17)$$

and fortunately, it is in the form of an ASG. The term inside the integral of Eq. 16 can be approximated as a rotated ASG of direction \mathbf{v} , and hence the integral in Eq. 16 can be approximated as:

$$\int_{\Omega} e^{-(\lambda+\nu)(\mathbf{v} \cdot \mathbf{x} - \frac{\nu(\mathbf{p} \cdot \mathbf{x})}{\lambda+\nu})^2 - (\mu+\nu)(\mathbf{v} \cdot \mathbf{y} - \frac{\nu(\mathbf{p} \cdot \mathbf{y})}{\mu+\nu})^2} d\mathbf{v} \approx \int_{\Omega} e^{-(\lambda+\nu)(\mathbf{v} \cdot \mathbf{x}')^2 - (\mu+\nu)(\mathbf{v} \cdot \mathbf{y}')^2} d\mathbf{v} \approx \frac{\pi}{\sqrt{(\lambda+\nu)(\mu+\nu)}},$$

where \mathbf{x}' and \mathbf{y}' are rotated axis directions that are given by $\mathbf{x}' = \mathbf{x} - \frac{\nu}{\lambda+\nu}\mathbf{p}$, $\mathbf{y}' = \mathbf{y} - \frac{\nu}{\mu+\nu}\mathbf{p}$. Substituting the above equations into Eq. 16 gives:

$$C(\mathbf{p}) \approx G(\mathbf{p}, [\mathbf{x}, \mathbf{y}, \mathbf{z}], [\frac{\nu\lambda}{\nu+\lambda}, \frac{\nu\mu}{\nu+\mu}], \frac{\pi}{\sqrt{(\lambda+\nu)(\mu+\nu)}}). \quad (18)$$

Validation. The error of our convolution approximation is mainly caused by the approximation in Eq. 15. This approximation is accurate when the value of $|\mathbf{p} \cdot \mathbf{z} - \mathbf{v} \cdot \mathbf{z}|$ is small, but will probably produce large errors when it is large. However, in practice, when $|\mathbf{p} \cdot \mathbf{z} - \mathbf{v} \cdot \mathbf{z}|$ is large, $|\mathbf{p} - \mathbf{v}|$ is also large, making the value of convolution kernel (i.e., the SG) to be small and contribute less to the whole convolution. Hence, in this case, the approximation error of the whole convolution won't be affected much. In Fig. 6, we visualize and compare the approximated (approx.) convolution using a single ASG (Eq. 18) and the ground truth (gt.) convolution by taking three examples. It is clear that the approximated convolutions are visually indistinguishable from the accurate ones. To more accurately identify the error, again we plot the quantitative approximation errors in Fig. 5 (right), of the three examples (Ex.) in Fig. 6. For each example, the error curve is generated by changing the convolution kernel size (i.e. the bandwidth ν of the SG). This approximation is valid when the bandwidths of the ASG λ and μ are not small. Quantitatively, when $\lambda, \mu > 3$, the L2-error of the approximation is smaller than 0.2% (no matter what the convolution kernel size ν is); when $\lambda, \mu > 1$, the L2-error is bounded in




Figure 6: Convolution of an ASG and an SG. The bandwidth parameters of the ASG (λ and μ) and the SG (ν) in all examples are listed below. Example 1: $\lambda = \mu = 5$, $\nu = 1$. Example 2: $\lambda = 40$, $\mu = 1$, $\nu = 10$. Example 3: $\lambda = 100$, $\mu = 10$, $\nu = 100$.

2.8%. More validations can be found in the supplemental document and accompanying video.

Summary. The convolution of an ASG with an SG can be well approximated by a single ASG using Eq. 18.

5 ASG-based Rendering Framework

In this section, we will show how ASGs can fit into existing SG-based frameworks to benefit rendering applications especially those involving anisotropic materials or lighting. Essentially, we follow the SG-based rendering framework of [Wang et al. 2009], while represent lighting and BRDFs using ASGs instead of SGs, and approximate visibility by extending the spherical signed distance function (SSDF) to accommodate integrations with ASGs. Details are explained in following subsections.

Rendering formulation. Following [Wang et al. 2009], by assuming distant lighting and static scenes, the outgoing radiance $L(\mathbf{o})$ (at a point) under direct illumination (ignoring inter-reflection) is calculated by:

$$L(\mathbf{o}) = \int_{\Omega} L(\mathbf{i})V(\mathbf{i})\rho(\mathbf{i}, \mathbf{o}) \max(\mathbf{i} \cdot \mathbf{n}, 0) d\mathbf{i}, \quad (19)$$

where \mathbf{i} , \mathbf{o} , \mathbf{n} denote lighting, view and normal directions, respectively; L , V denote the spherical lighting and visibility functions, respectively; ρ denotes the 4D BRDF function and Ω denotes the unit sphere. By decomposing the BRDF into the sum of a diffuse component and a specular component: $\rho(\mathbf{i}, \mathbf{o}) = k_d + k_s\rho_s(\mathbf{i}, \mathbf{o})$, the outgoing radiance $L(\mathbf{o})$ can hence be decomposed into a diffuse term L_d and a specular term L_s [Wang et al. 2009]: $L(\mathbf{o}) = k_dL_d + k_sL_s(\mathbf{o})$, where:

$$L_d = \int_{\Omega} L(\mathbf{i})V(\mathbf{i}) \max(\mathbf{i} \cdot \mathbf{n}, 0) d\mathbf{i}, \quad (20)$$

$$L_s(\mathbf{o}) = \int_{\Omega} L(\mathbf{i})V(\mathbf{i})\rho_s(\mathbf{i}, \mathbf{o}) \max(\mathbf{i} \cdot \mathbf{n}, 0) d\mathbf{i}. \quad (21)$$

5.1 Light Approximation

We represent the environment light using ASG mixtures. Firstly, some initial ASGs are determined by preserving the local maximums of the environment light, and then, following [Tsai and Shih




Figure 7: Comparison of ASGs and SGs in fitting environment lights of St. Peter's Basilica and Grace Cathedral.

2006], we use the L-BFGS-B solver [Zhu et al. 1997] to fit the environment light. Local light sources are approximated by equal-bandwidth ASGs following the method from [Wang et al. 2009].

As shown in Fig. 7, we fit two environment maps (EMs) using ASG mixtures and SG mixtures, respectively. Clearly, ASGs are superior to SGs and better capture those anisotropic features in EMs (i.e. some ellipse-like regions).

5.2 BRDF Approximation

Based on the microfacet model [Cook and Torrance 1982], the specular term ρ_s of BRDF can be represented by:

$$\rho_s(\mathbf{i}, \mathbf{o}) = M(\mathbf{i}, \mathbf{o})D(\mathbf{h}), \mathbf{h} = (\mathbf{i} + \mathbf{o})/|\mathbf{i} + \mathbf{o}| \quad (22)$$

where \mathbf{h} is the unit half vector, D is the *normal distribution function* (NDF) and M is a combined function including shadowing and Fresnel terms, which is smooth [Ngan et al. 2005]. Wang et al. [2009] approximate the NDF using SG mixtures, and, in contrast, we represent it with ASG mixtures. In practice, a single ASG is usually enough to approximate lots of complex, highly anisotropic BRDFs.

Parametric BRDFs. For isotropic parametric BRDF, we first fit it as sum of SGs [Wang et al. 2009], and then convert SGs to equal-bandwidth ASGs using Eq. 3. We further approximate two anisotropic parametric BRDF models using ASGs: the Ward model [1992] and the Ashikhmin model [Ashikhmin and Shirley 2000]. The Ward model can be approximated using one ASG:

$$M(\mathbf{i}, \mathbf{o}) = 1/(4\pi\alpha_x\alpha_y\sqrt{(\mathbf{i} \cdot \mathbf{n})(\mathbf{o} \cdot \mathbf{n})}),$$

$$D(\mathbf{h}) = e^{-\frac{2}{1+\mathbf{h} \cdot \mathbf{n}} \left(\frac{(\mathbf{h} \cdot \mathbf{x})^2}{\alpha_x^2} + \frac{(\mathbf{h} \cdot \mathbf{y})^2}{\alpha_y^2} \right)} \approx G(\mathbf{h}; [\mathbf{x}, \mathbf{y}, \mathbf{n}], [\frac{1}{\alpha_x^2}, \frac{1}{\alpha_y^2}]),$$

and the Ashikhmin model can also be approximated by one ASG:

$$M(\mathbf{i}, \mathbf{o}) = \frac{\sqrt{(n_u + 1)(n_v + 1)}F(\mathbf{i} \cdot \mathbf{h})}{8\pi(\mathbf{i} \cdot \mathbf{h}) \cdot \max(\mathbf{n} \cdot \mathbf{i}, \mathbf{n} \cdot \mathbf{o})},$$

$$D(\mathbf{h}) = (\mathbf{h} \cdot \mathbf{n})^{\frac{n_u(\mathbf{h} \cdot \mathbf{u})^2 + n_v(\mathbf{h} \cdot \mathbf{v})^2}{1 - (\mathbf{h} \cdot \mathbf{n})^2}} \approx G(\mathbf{h}; [\mathbf{u}, \mathbf{v}, \mathbf{n}], [\frac{n_u}{2}, \frac{n_v}{2}]),$$

where F is the fresnel function.

As shown in Fig. 8, we approximate the Ashikhmin BRDF using a single ASG (6th col.), and compare the results with the references (7th col., which are generated using standard path tracing.) and the results using different number of SGs (1st to 5th col.). Although 5 SGs are enough in approximating low-anisotropy-ratio BRDFs (e.g. the second row, with anisotropy ratio = 3), the results of using 19




Figure 8: Comparison of ASGs and SGs in fitting Ashikmin BRDFs with different parameters. First row: $n_u = 8, n_v = 800$ (anisotropy ratio: 10); second row: $n_u = 200, n_v = 22$ (anisotropy ratio: 3); third row: $n_u = 20000, n_v = 200$ (anisotropy ratio: 10). The L^2 error for each configuration is also given in the corresponding subtitle.




Figure 10: Spherical Warping.

SGs still exhibit visible differences (see the close-up views) in approximating highly anisotropic BRDFs (e.g. the first or third rows, with anisotropy ratio = 10). In contrast, under all these parameter settings, the results of using only one ASG are visually indistinguishable from the path-traced references, demonstrating the good scalability of ASGs in approximating anisotropic BRDFs.

Measured BRDFs. We represent measured BRDFs using the general microfacet model as in [Wang et al. 2009]. We first fit the NDF $D(\mathbf{h})$ and the smooth function M from the raw measured BRDFs. After that, we fit the NDF using ASG mixtures. For spatially-varying BRDFs, the above procedure is repeated for every point. In Fig. 9, we fit two real-captured anisotropic BRDFs (brushed aluminum and purple satin) using ASGs, and compare the results fitted with ASGs to those fitted with Ashikmin or Ward BRDFs. While these parametric models can approximate the top-row brushed aluminum BRDF well (although the results still exhibit visible differences from the references), they cannot model the bottom-row purple satin BRDF well, since this BRDF exhibits off-specular reflections that cannot be captured in standard normal-based local frame. In contrast, the differences between our results with 3 or 4 ASGs and the references are visually subtle.

Spherical Warping of ASG-based BRDFs. When using microfacet BRDFs, given a view direction \mathbf{o} , we need to obtain the corresponding 2D BRDF slice to integrate with incident lighting and visibility functions. Yet, usually the NDF term of a microfacet BRDF is parameterized using half vector \mathbf{h} instead of incident light direction \mathbf{i} . Therefore, for efficient multiplication of BRDF and incident

lighting, a *spherical warping* strategy as suggested in [Wang et al. 2009] is also required for ASG-based BRDF representation.

Specifically, assuming the NDF $D(\mathbf{h})$ can be approximated by an ASG: $D(\mathbf{h}) \approx G(\mathbf{h}; [\mathbf{x}_h, \mathbf{y}_h, \mathbf{z}_h], [\lambda_h, \mu_h])$ (short for $G_h(\mathbf{h})$ and shown in Fig. 10 (a)), where \mathbf{h} is the half vector of view direction \mathbf{o} and light direction \mathbf{i} : $\mathbf{h}(\mathbf{i}) = \frac{\mathbf{i} + \mathbf{o}}{|\mathbf{i} + \mathbf{o}|}$. We approximate the *warped* NDF expressed in terms of lighting direction \mathbf{i} again using another ASG:

$$G_h(\mathbf{h}(\mathbf{i})) = G_h\left(\frac{\mathbf{i} + \mathbf{o}}{|\mathbf{i} + \mathbf{o}|}\right) \approx G(\mathbf{i}; [\mathbf{x}_i, \mathbf{y}_i, \mathbf{z}_i], [\lambda_i, \mu_i]). \quad (23)$$

To obtain the parameters of the ASG $G_i(\mathbf{i})$ in the above equation, we first preserve the lobe position, and have $\mathbf{z}_i = 2(\mathbf{o} \cdot \mathbf{z}_h)\mathbf{z}_h - \mathbf{o}$. Then, the other parameters can be determined by preserving the local curvature (second-order derivatives) of the exponential term around lobe position $\mathbf{i} = \mathbf{z}_i$. Specifically, denote the exponential order of the warped ASG $G_h(\mathbf{h}(\mathbf{i}))$ as $g(\mathbf{i})$ (such that $G_h(\mathbf{h}(\mathbf{i})) = S(\mathbf{h}(\mathbf{i})) \cdot e^{-g(\mathbf{i})}$), which has the form of:

$$g(\mathbf{i}) = \lambda_h(\mathbf{h}(\mathbf{i}) \cdot \mathbf{x}_h)^2 + \mu_h(\mathbf{h}(\mathbf{i}) \cdot \mathbf{y}_h)^2.$$

To facilitate derivation, at direction $\mathbf{i} = \mathbf{z}_i$, we define a default local frame $[\mathbf{x}'_i, \mathbf{y}'_i, \mathbf{z}_i]$, making the tangent direction \mathbf{x}'_i perpendicular to view direction \mathbf{o} (as shown in Fig. 10 (b)). The exponential order $g(\mathbf{i})$ can hence be approximated as a second order Taylor expansion at $\mathbf{i} = \mathbf{z}_i$:

$$g(\mathbf{i}) \approx (\mathbf{i} \cdot \mathbf{x}'_i, \mathbf{i} \cdot \mathbf{y}'_i) \cdot \mathbf{H}(g) \cdot (\mathbf{i} \cdot \mathbf{x}'_i, \mathbf{i} \cdot \mathbf{y}'_i)^T, \quad (24)$$

where $\mathbf{H}(g)$ is the 2×2 Hessian matrix, which can be further factorized by eigen-decomposition:

$$\mathbf{H}(g) = \begin{pmatrix} \frac{\partial^2 g}{\partial \mathbf{x}'_i{}^2} & \frac{\partial^2 g}{\partial \mathbf{x}'_i \partial \mathbf{y}'_i} \\ \frac{\partial^2 g}{\partial \mathbf{x}'_i \partial \mathbf{y}'_i} & \frac{\partial^2 g}{\partial \mathbf{y}'_i{}^2} \end{pmatrix} = \mathbf{U} \cdot \begin{pmatrix} \lambda_i & 0 \\ 0 & \mu_i \end{pmatrix} \cdot \mathbf{U}^T. \quad (25)$$

The formulas of the second-order derivatives in the Hessian matrix can be found in Section 6 of the supplemental document. Substituting Eq. 25 into Eq. 24 yields:

$$g(\mathbf{i}) \approx \lambda_i(\mathbf{i} \cdot \mathbf{x}_i)^2 + \mu_i(\mathbf{i} \cdot \mathbf{y}_i)^2, \quad (26)$$




Figure 9: Fitting two real captured BRDFs (brushed aluminum and purple satin) from [Ngan et al. 2005]. We also give the fitted results using Ashikmin’s and Ward’s parametric anisotropic models for comparison. Notice that the parametric models cannot fit the bottom data (purple satin) well, since its lobe is off-specular. The L^2 error for each configuration is also given in the corresponding subtitle.

where $[\mathbf{x}_i, \mathbf{y}_i] = \mathbf{U} \cdot [\mathbf{x}'_i, \mathbf{y}'_i]$. Hence the warped ASG $G_h(\mathbf{h}(\mathbf{i}))$ can be approximated by an ASG $G(\mathbf{i}; [\mathbf{x}_i, \mathbf{y}_i, \mathbf{z}_i], [\lambda_i, \mu_i])$ (Fig. 10 (c)). When using an ASG mixture to represent BRDFs, spherical warping is applied to each ASG in the mixture. Detailed validation of spherical warping can be found in Section 7 of the supplemental document.

5.3 Visibility Approximation

We adopt the spherical signed distance function (SSDF), which is presented in [Wang et al. 2009], to approximate visibility. Specifically, given a visibility function $V(\mathbf{i})$, SSDF $\theta_d(\mathbf{i})$ stores the nearest angular distance from the direction (point) \mathbf{i} to the nearest direction (point) on the visibility 0/1 boundary, and its sign is positive (+) when $V(\mathbf{i}) = 1$ or negative (−) when $V(\mathbf{i}) = 0$. By using SSDF, the product integral of an SG with a visibility function can be computed as a 2D function:

$$\int_{\Omega} G_{iso}(\mathbf{i}; \mathbf{p}, \nu) \cdot V(\mathbf{i}) d\mathbf{i} \approx f_h(\theta_d(\mathbf{p}), \nu),$$

where f_h is a sigmoid function composed with a polynomial [Wang et al. 2009]. However, the above equation cannot be directly used for ASGs, since the SSDF representation is essentially isotropic and does not account for azimuthal angle changes. Hence, to compute the product integral of a visibility function with an ASG $G(\mathbf{i}; [\mathbf{x}, \mathbf{y}, \mathbf{p}], [\lambda, \mu])$, we instead define an *effective bandwidth* $\sqrt{\lambda\mu}$, which is motivated by approximating the ellipse-shape support of ASG by an equal-size circle-shape support, to compute the product integral, which is:

$$\int_{\Omega} G(\mathbf{i}; [\mathbf{x}, \mathbf{y}, \mathbf{p}], [\lambda, \mu]) \cdot V(\mathbf{i}) d\mathbf{i} \approx f_h(\theta_d(\mathbf{p}), \sqrt{\lambda\mu}). \quad (27)$$

The product of an ASG and a visibility function can still be approximated as an ASG, only attenuating the amplitude of the original ASG:

$$G(\mathbf{i}; [\mathbf{x}, \mathbf{y}, \mathbf{p}], [\lambda, \mu]) \cdot V(\mathbf{i}) \approx \frac{f_h(\theta_d, \sqrt{\lambda\mu})}{f_h(\frac{\pi}{2}, \sqrt{\lambda\mu})} \cdot G(\mathbf{i}; [\mathbf{x}, \mathbf{y}, \mathbf{p}], [\lambda, \mu]). \quad (28)$$

Finally, we also employ PCA to compress the precomputed SSDF.

The above approximation using effective bandwidth to compute the product integral of an ASG with visibility can incur errors when the anisotropic ratio is large. However, it will only affect the softness of the shadow boundaries (since in other locations, the visibility values are mostly either zero or one). In Fig. 11, we evaluate our visibility approximation. In the test scene, both the teapot and the cylinder have anisotropic BRDFs in both rows, while the plane has

a diffuse/anisotropic BRDF in top/bottom rows, respectively. Notice that in both cases, our approximation without PCA compression (col. (d)) exhibits subtle differences with the references (col. (e)). We also give results generated using SSDFs compressed with 16, 32, and 48 PCA terms, respectively. In practice, we find using 48 PCA terms achieves a good trade-off between accuracy and performance, and hence all the other results in the paper are rendered using 48 PCA terms for visibility approximation.

6 Applications and Results

To demonstrate the merits of ASGs in practice, we implement two dominant SG-based rendering applications, all-frequency rendering with dynamic BRDFs [Wang et al. 2009] and bi-scale BRDF editing [Iwasaki et al. 2012a], using the ASG-based rendering framework. In this subsection, we will describe how to implement these two applications and show the rendering results. Our implementations are done in OpenGL on a PC with Intel Xeon 2.27G CPU and 8 GB memory, and an NVIDIA GeForce GTX 680 graphics card. All the result images are generated with a resolution of 720×480 . The performance of all the generated results in run-time rendering is reported in Table 1. As for the timing of ASG fitting procedure, in our experiments, fitting an environment map (with resolution of 1024×768) costs 23-51 seconds; fitting a measured BRDF (with NDF resolution of 256×256) using 1-3 ASGs takes 0.9-3 seconds.

All-frequency rendering with dynamic BRDFs. To finally achieve all-frequency rendering, we need to evaluate both the diffuse term L_d (Eq. 20) and the specular term $L_s(\mathbf{o})$ (Eq. 21).

By approximating the incident lighting $L(\mathbf{i})$ using ASG mixtures $L(\mathbf{i}) \approx \sum_j l_j G_j(\mathbf{i})$ ($G_j(\mathbf{i})$ is short for $G(\mathbf{i}; [\mathbf{x}_j, \mathbf{y}_j, \mathbf{z}_j], [\lambda_j, \mu_j])$), the diffuse term L_d can be approximated as:

$$\begin{aligned} L_d &\approx \int_{\Omega} \left(\sum_j l_j G_j(\mathbf{i}) \right) V(\mathbf{i}) \max(\mathbf{i} \cdot \mathbf{n}, 0) d\mathbf{i} \\ &\approx \sum_j l_j \max(\mathbf{z}_j \cdot \mathbf{n}, 0) \int_{\Omega} G_j(\mathbf{i}) V(\mathbf{i}) d\mathbf{i}, \end{aligned} \quad (29)$$

Here, since the cosine term $\max(\mathbf{i} \cdot \mathbf{n}, 0)$ is smooth, we approximate it as a constant value and take it out of the integral. The left ASG integral with visibility can be efficiently evaluated using Eq. 27.

For the specular term $L_s(\mathbf{o})$, we first obtain the ASG-based representation of 2D BRDF slice through spherical warping: $\rho_s(\mathbf{i}, \mathbf{o}) \approx M(\mathbf{i}, \mathbf{o}) \sum_k G_k(\mathbf{i})$ (usually $1 \leq k \leq 4$, and $G_k(\mathbf{i})$ is short for $G(\mathbf{i}; [\mathbf{x}_k, \mathbf{y}_k, \mathbf{z}_k], [\lambda_k, \mu_k])$), and then Eq. 21 can be rewritten as:




Figure 11: Evaluation of visibility approximation.

scene	#vert.	BRDF/#ASG.	#E.L./ #P.L.	fps
dish (Fig. 1)	25k	Ashikhmin/1	10/2	125
teapot (Fig. 11)	54k	Ashikhmin/1	10/0	201
dish ball (Fig. 12)	29k	Ward/1	10/2	120
dragon (Fig. 13 (a))	46k	aluminum/4	8/0	83
dish (Fig. 13 (b))	25k	s.v. aluminum/1	14/0	218
pillow (Fig. 13 (c))	6k	s.v. wallpaper/1	8/0	278
dish card(Fig. 13 (d))	17k	s.v. satin/1,s.v. card/1	12/0	245
cloth(Fig. 14)	16k	bi-scale BRDF/2	12/0	182

Table 1: Performance of the results shown in this paper. From left to right, we give the name of the scene, the number of vertices, the BRDF used (and the number of ASGs used to approximate the BRDF), the number of ASGs for environment lights and local point lights, and the framerates.

$$\begin{aligned}
L_s(\mathbf{o}) &\approx \int_{\Omega} L(\mathbf{i}) V(\mathbf{i}) \left(\sum_k M(\mathbf{i}, \mathbf{o}) G_k(\mathbf{i}) \right) \max(\mathbf{i} \cdot \mathbf{n}, 0) d\mathbf{i} \\
&\approx \sum_k M(\mathbf{z}_k, \mathbf{o}) \max(\mathbf{z}_k \cdot \mathbf{n}, 0) \int_{\Omega} L(\mathbf{i}) (V(\mathbf{i}) G_k(\mathbf{i})) d\mathbf{i} \\
&\approx \sum_k M(\mathbf{z}_k, \mathbf{o}) \max(\mathbf{z}_k \cdot \mathbf{n}, 0) \cdot g_k \int_{\Omega} L(\mathbf{i}) G_k(\mathbf{i}) d\mathbf{i}. \quad (30)
\end{aligned}$$

Here, the first equation relies on the fact that both M and the cosine term $\max(\mathbf{i} \cdot \mathbf{n}, 0)$ are smooth, and hence can be approximated as constants and taken out of the integral. The second equation approximates the product of an ASG with visibility again by an ASG (Eq. 28), where $g_k = f_h(\theta_d(\mathbf{z}_k), \sqrt{\lambda_k \mu_k}) / f_h(\frac{\pi}{2}, \sqrt{\lambda_k \mu_k})$.

Next, the specular term is reduced to integrate an ASG with the lighting, in other words, an ASG-filtered lighting. To efficiently evaluate the ASG filtered value using graphics hardware, following [Wang et al. 2009], we first build a pre-filtered mipmap for the environment map (stored as cube map texture). At runtime, instead of using an isotropic texture lookup method (i.e. **textureLod**) to query the filtered lighting intensity, we use OpenGL built-in anisotropic texture lookup method (i.e. **textureGrad**) for that purpose. The gradient directions used for **textureGrad** are calculated at run-time from the elliptical support of the ASG. Note that, current graphics hardware supports a maximal anisotropy of 16 : 1, which is usually enough for most anisotropic BRDFs. Alternatively, we can manually put texture samples to obtain filtered results for extremely anisotropic BRDFs (anisotropy > 16 : 1) or for better quality.

Fig. 12 shows the results of a dish ball scene containing anisotropic Ward BRDFs with different levels of anisotropy and different lo-



Figure 14: Bi-scale BRDF editing of the cloth scene.

cal frames. The BRDF of each object is approximated using only one ASG. Notice the highly anisotropic highlights on the balls and on the dish. In Fig. 13, we show the results with measured anisotropic BRDFs and spatially-varying BRDFs from [Ngan et al. 2005; Lawrence et al. 2006; Wang et al. 2008; Dong et al. 2010]. Fig. 13 (a) shows a dragon scene with the measured brushed aluminum BRDF from [Ngan et al. 2005], which is approximated using 4 ASGs. Fig. 13 (b)-(d) show 3 different scenes with measured spatially-varying BRDFs. Note that the aluminium BRDF in Fig. 13 (b) is highly anisotropic (i.e., anisotropy ratio is about 10), and our method is capable of faithfully reproducing its appearance using only one ASG.

Bi-scale BRDF editing. Let us briefly review the original SG-based bi-scale BRDF editing method [Iwasaki et al. 2012a]. The large-scale effective BRDF $\bar{\rho}$ is the convolution of a small-scale isotropic BRDF ρ_{small} and an effective bi-directional visible normal distribution function (BVNDF) $\bar{\gamma}$:

$$\bar{\rho}(\mathbf{i}, \mathbf{o}) = \int_{\Omega} \rho_{small}(\mathbf{n}, \mathbf{i}, \mathbf{o}) \bar{\gamma}(\mathbf{n}, \mathbf{i}, \mathbf{o}) d\mathbf{n}. \quad (31)$$

The original method approximates both the small-scale isotropic BRDF and the effective BVNDF using SG mixtures, and approximate the convolution of two SGs again by an SG. Hence, the large-scale effective BRDF $\bar{\rho}$ can also be approximated by SG mixtures, and rendering can be achieved using the method of [Wang et al. 2009].

We slightly modify the approximation of BRDFs by representing the effective BVNDF using ASG mixtures, and still represent the small-scale isotropic BRDF using SGs. As shown in Sec. 4.3, the convolution of an ASG with an SG can still be approximated as an ASG. Hence, the large-scale effective BRDF $\bar{\rho}$ can be approximated by ASG mixtures, so that rendering can be achieved using our ASG-based rendering framework. The advantage of using ASGs instead of SGs to represent BVNDF lies in that much less number of ASGs are required for effectively representing highly anisotropic BVNDFs. In Fig. 14, we show the results of ASG-based bi-scale BRDF editing. In this example, we approximate the BVNDF of the




Figure 12: Dish Ball.




Figure 13: Results of measured and spatially varying BRDFs.

“roof shape” small-scale geometry using only 2 ASGs, and in contrast, a large number of SGs are required to reach similar quality. By changing the small-scale geometry, which is shown in the top-left subfigures in Fig. 14, we can modify the large-scale appearance accordingly. Please refer to the accompanying video for interactive editing sequences.

Comparisons with SGs. In the teaser figure (Fig. 1), we further compare the rendering results using a single ASG with those using different number of SGs. The test scene consists of a metal dish exhibiting highly anisotropic appearance, which is modeled by an Ashikhmin BRDF with anisotropy ratio of 10:1. The incident lighting includes an environment map and two local lights. As demonstrated by the rendering results, ASGs are much superior to SGs in dealing with such scenes.

7 Discussions and Conclusion

Scope. In above, we have demonstrated the usage of ASGs in two important rendering applications. Supported by the newly-derived operators, we can easily further adapt ASGs to other existing SG-based applications and improve their scalability in handling anisotropic effects. E.g., ASGs can be applied in normal map filtering [Han et al. 2007] to better represent anisotropic NDFs, or used in indirect highlight rendering [Laurijssen et al. 2010] to deal with anisotropic BRDFs, or employed by real-time rough refraction [de Rousiers et al. 2012] to render anisotropic refraction effects, etc. Besides rendering, since SGs have been widely used for function approximation and nonlinear regression estimation in the field of machine learning, ASGs and its operators with closed-form solutions are also expected to benefit this kind of applications in fields beyond computer graphics.

Conclusion. In summary, we present a novel anisotropic Spherical Gaussian (ASG) function and the approximate closed-form solutions for its *integral*, *product* and *convolution* operators, the accuracy of which is validated by quantitative evaluations. Besides inheriting the nice properties of SGs, ASGs are intrinsically anisotropic and hence can represent anisotropic spherical function much more effectively and efficiently compared with SGs. Based on the ASG-based representations, we have developed an ASG-based rendering framework and implemented two important rendering applications to demonstrate the effectiveness and efficiency

of ASGs in real applications.

There are several future directions for ASGs. First, we currently only derive the convolution operator of an ASG with an SG. To make ASGs more mathematically complete and also benefit future applications, it is valuable to investigate how to compute the convolution of two ASGs. Regarding the convolution of two ASGs to be performing anisotropic blurring using one ASG over another one, it is possible to derive its analytic formula. Secondly, an efficient and stable method to achieve importance sampling from an ASG distribution may also potentially benefit offline rendering research.

Acknowledgements. We thank the reviewers for their valuable comments. This work was supported by National Basic Research Project of China (2012CB316400), Natural Science Foundation of China (61120106007 and 61170153), National High Technology Research and Development Program of China (2012AA011503), PCSIRT and Tsinghua University Initiative Scientific Research Program. Kun Xu is also supported by the CCF-Intel Young Faculty Researcher Program.

References

- ASHIKHMIN, M., AND SHIRLEY, P. 2000. An anisotropic phong brdf model. *Journal of Graphics Tools* 5, 2 (Feb.), 25–32.
- BINGHAM, C. 1974. An antipodally symmetric distribution on the sphere. *Annals of Statistics* 2, 6, 1201–1225.
- COOK, R. L., AND TORRANCE, K. E. 1982. A reflectance model for computer graphics. *ACM Trans. Graph.* 1, 1, 7–24.
- DE ROUSIERS, C., BOUSSEAU, A., SUBR, K., HOLZSCHUCH, N., AND RAMAMOORTHY, R. 2012. Real-time rendering of rough refraction. *IEEE Trans. Vis. Comput. Graph.* 18, 10, 1591–1602.
- DONG, Y., WANG, J., TONG, X., SNYDER, J., LAN, Y., BEN-EZRA, M., AND GUO, B. 2010. Manifold bootstrapping for svbrdf capture. *ACM Trans. Graph.* 29, 4, 98:1–98:10.
- EDWARDS, D., BOULOS, S., JOHNSON, J., SHIRLEY, P., ASHIKHMIN, M., STARK, M., AND WYMAN, C. 2006. The

- halfway vector disk for brdf modeling. *ACM Trans. Graph.* 25, 1 (Jan.), 1–18.
- FISHER, R. 1953. Dispersion on a sphere. *Proc. Roy. Soc. London Ser. A*, 217, 1130, 295–305.
- GREEN, P., KAUTZ, J., MATUSIK, W., AND DURAND, F. 2006. View-dependent precomputed light transport using nonlinear gaussian function approximations. In *Proceedings of I3D*, ACM, 7–14.
- GREEN, P., KAUTZ, J., AND DURAND, F. 2007. Efficient reflectance and visibility approximations for environment map rendering. *Computer Graphics Forum* 26, 3, 495–502.
- HAN, C., SUN, B., RAMAMOORTHY, R., AND GRINSPUN, E. 2007. Frequency domain normal map filtering. *ACM Trans. Graph.* 26, 3.
- IRAWAN, P., AND MARSCHNER, S. 2012. Specular reflection from woven cloth. *ACM Trans. Graph.* 31, 1 (Feb.), 11:1–11:20.
- IWASAKI, K., DOBASHI, Y., AND NISHITA, T. 2012. Interactive bi-scale editing of highly glossy materials. *ACM Trans. Graph.* 31, 6 (Nov.), 144:1–144:7.
- IWASAKI, K., FURUYA, W., DOBASHI, Y., AND NISHITA, T. 2012. Real-time rendering of dynamic scenes under all-frequency lighting using integral spherical gaussian. *Computer Graphics Forum* 31, 727–734.
- JAKOB, W., ARBREE, A., MOON, J. T., BALA, K., AND MARSCHNER, S. 2010. A radiative transfer framework for rendering materials with anisotropic structure. *ACM Trans. Graph.* 29, 4 (July), 53:1–53:13.
- KAJIYA, J. T. 1985. Anisotropic reflection models. *ACM SIGGRAPH Computer Graphics* 19, 3, 15–21.
- KAJIYA, J. T. 1986. The rendering equation. *SIGGRAPH Comput. Graph.* 20, 4 (Aug.), 143–150.
- KENT, J. T. 1982. The fisher-bingham distribution on the sphere. *J. Royal. Stat. Soc.* 44, 1, 71–80.
- KURT, M., SZIRMAY-KALOS, L., AND KŘIVÁNEK, J. 2010. An anisotropic brdf model for fitting and monte carlo rendering. *SIGGRAPH Computer Graphics* 44, 1 (Feb.), 3:1–3:15.
- LAURIJSEN, J., WANG, R., DUTRÉ, P., AND BROWN, B. 2010. Fast estimation and rendering of indirect highlights. *Computer Graphics Forum* 29, 4, 1305–1313.
- LAWRENCE, J., BEN-ARTZI, A., DECORO, C., MATUSIK, W., PFISTER, H., RAMAMOORTHY, R., AND RUSINKIEWICZ, S. 2006. Inverse shade trees for non-parametric material representation and editing. *ACM Trans. Graph.* 25, 3, 735–745.
- MARDIA, K. V., AND JUPP, P. E. 1999. *Directional Statistics*. John Wiley & Sons, Inc.
- MARDIA, K. V. 1975. Statistics of directional data. *J. R. Statist. Soc. B* 37, 3, 349–393.
- NG, R., RAMAMOORTHY, R., AND HANRAHAN, P. 2003. All-frequency shadows using non-linear wavelet lighting approximation. *ACM Trans. Graph.* 22, 3, 376–381.
- NGAN, A., DURAND, F., AND MATUSIK, W. 2005. Experimental analysis of brdf models. In *Proceedings of EGSR*, 117–126.
- OLANO, M., AND BAKER, D. 2010. Lean mapping. In *Proceedings of I3D*, ACM, New York, NY, USA, 181–188.
- PACANOWSKI, R., SALAZAR CELIS, O., SCHLICK, C., GRANIER, X., POULIN, P., AND CUYT, A. 2012. Rational brdf. *IEEE Transactions on Visualization and Computer Graphics* 18, 11, 1824–1835.
- RAMAMOORTHY, R., AND HANRAHAN, P. 2001. An efficient representation for irradiance environment maps. In *Proc. of SIGGRAPH*, ACM, 497–500.
- SADEGHI, I., BISKER, O., DEKEN, J. D., AND JENSEN, H. W. 2013. A practical microcylinder appearance model for cloth rendering. *ACM Trans. Graph.* 32, 2, 14:1–14:12.
- 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, 527–536.
- SUN, X., HOU, Q., REN, Z., ZHOU, K., AND GUO, B. 2011. Radiance transfer biclustering for real-time all-frequency bi-scale rendering. *IEEE Transactions on Visualization and Computer Graphics* 17, 1, 64–73.
- TORRANCE, K. E., AND SPARROW, E. M. 1967. Theory for off-specular reflection from roughened surfaces. *Journal of the Optical Society of America* 57, 9, 1105–1112.
- TSAI, Y.-T., AND SHIH, Z.-C. 2006. All-frequency precomputed radiance transfer using spherical radial basis functions and clustered tensor approximation. *ACM Trans. Graph.* 25, 3, 967–976.
- WANG, J., ZHAO, S., TONG, X., SNYDER, J., AND GUO, B. 2008. Modeling anisotropic surface reflectance with example-based microfacet synthesis. *ACM Trans. Graph.* 27, 3 (Aug.), 41:1–41:9.
- 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.
- WARD, G. J. 1992. Measuring and modeling anisotropic reflection. In *Proceedings of Siggraph*, 265–272.
- WU, H., DORSEY, J., AND RUSHMEIER, H. 2011. Physically-based interactive bi-scale material design. *ACM Trans. Graph.* 30, 6 (Dec.), 145:1–145:10.
- XU, K., JIA, Y.-T., FU, H., HU, S.-M., AND TAI, C.-L. 2008. Spherical piecewise constant basis functions for all-frequency precomputed radiance transfer. *IEEE Transaction on Visualization and Computer Graphics* 14, 2, 454–467.
- XU, K., MA, L.-Q., REN, B., WANG, R., AND HU, S.-M. 2011. Interactive hair rendering and appearance editing under environment lighting. *ACM Trans. Graph.* 30, 6, 173:1–173:10.
- YAN, L.-Q., ZHOU, Y., XU, K., AND WANG, R. 2012. Accurate translucent material rendering under spherical gaussian lights. *Computer Graphics Forum* 31, 7, 2267–2276.
- ZHAO, S., JAKOB, W., MARSCHNER, S., AND BALA, K. 2012. Structure-aware synthesis for predictive woven fabric appearance. *ACM Trans. Graph.* 31, 4 (July), 75:1–75:10.
- ZHU, C., BYRD, R. H., LU, P., AND NOCEDAL, J. 1997. L-bfgs-b: Fortran subroutines for large-scale bound-constrained optimization. *ACM Trans. Math. Software* 23, 4 (Dec.), 550–560.