

# Specular Lobe-Aware Filtering and Upsampling for Interactive Indirect Illumination

Y. Tokuyoshi

Square Enix Co., Ltd., Japan

---

## Abstract

Although geometry-aware filtering and upsampling have often been used for interactive or real-time rendering, they are unsuitable for glossy surfaces because shading results strongly depend on the BRDFs. This paper proposes a novel weighting function of cross bilateral filtering and upsampling to measure the similarity of specular lobes. The difficulty is that a specular lobe is represented with a distribution function in directional space, whereas conventional cross bilateral filtering evaluates similarities using the distance between two points in a Euclidean space. Therefore, this paper first generalizes cross bilateral filtering for the similarity of distribution functions in a non-Euclidean space. Then, the weighting function is specialized for specular lobes. Our key insight is that the weighting function of bilateral filtering can be represented with the product integral of two distribution functions corresponding to two pixels. In addition, we propose spherical Gaussian based approximations to calculate this weighting function analytically. Our weighting function detects the edges of glossiness, and adapts to all-frequency materials using only a camera position and G-buffer. These features are not only suitable for path tracing, but also deferred shading and non-ray tracing based methods such as voxel cone tracing.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation—

---

## 1. Introduction

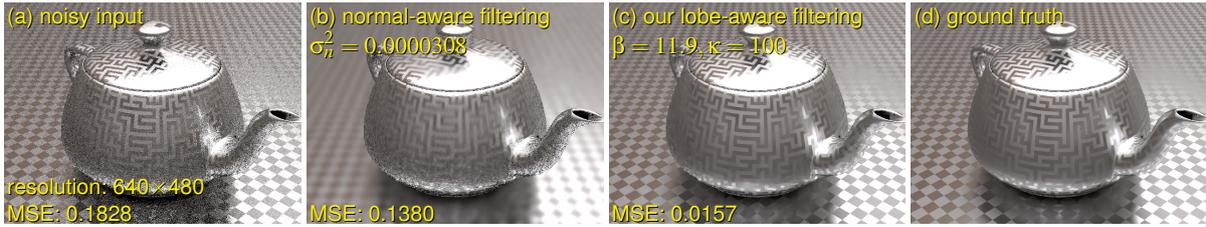
Computing indirect illumination in real-time frame rates is a major challenge with high industrial impact. However, it can be difficult for high-resolution frame buffers. In order to avoid this problem, denoising filters for Monte-Carlo (MC) rendering or upsampling techniques have often been used for post-processing. While sophisticated filtering methods have been developed, simple cross bilateral filtering [PSA\*04, ED04] and upsampling [KCLU07] are often employed for interactive or real-time applications due to their faster speed and lower memory usage. In deferred shading pipelines, a G-buffer can be inexpensively utilized for such filtering. These techniques are often called geometry-aware filtering and upsampling.

However, geometry-aware filtering can produce overblurring and underblurring errors for scenes with several glossy materials (Fig. 1). This is because they evaluate only the similarity of the surfaces by using a normal-aware weighting function and depth-aware weighting function, whereas actual illumination appearance is strongly affected by the specular lobe which is represented with a bidirectional re-

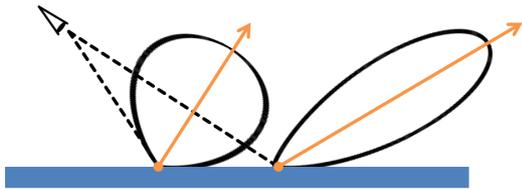
flectance distribution function (BRDF). For glossy indirect illumination, the similarity of the specular lobes should be evaluated as shown in Fig. 2.

Bilateral filtering evaluates similarities using the distance between two points in a high-dimensional Euclidean space, whereas specular lobes are distribution functions in directional space. Therefore, we represent the weighting function as a product integral of two distribution functions to generalize cross bilateral filtering. Using this generalization, we introduce a novel weighting function to measure the specular lobe similarity. Our weighting function inexpensively handles all-frequency materials similar to previous BRDF based bandwidth controlling methods such as [MWRD13]. These methods often use only the sharpness of a normal distribution function (NDF) for their BRDF based bandlimits, but the appropriate kernel bandwidth should be according to the reflection lobe sharpness which depends on a view direction and surface normal. Using reflection lobes, this paper takes into account such view dependent properties. In addition, our method detects the edges of lobe sharpness.

In order to calculate this weighting function analytically,



**Figure 1:** Quality comparison for glossy surfaces. The quality is evaluated using mean squared error (MSE) metric. The teapot and floor have different BRDF models (the Kelemen model [KSK01] with the Phong NDF and Cook-Torrance model [CT82] with the Beckmann NDF). Both surfaces have a glossiness map and roughness map, respectively. Conventional geometry-aware filtering using surface normals (b) produces noticeable overblurring and underblurring errors, while our specular lobe-aware filtering (c) does not. Parameters of weighting functions  $\sigma_n^2$  and  $\beta$  (described in Sect. 2.2 and 3) are optimal for this image.



**Figure 2:** Difference of specular lobes. Normal-aware weighting functions neglect this difference.

this paper additionally proposes an approximation using spherical Gaussians (SGs). Furthermore, an approximation using anisotropic spherical Gaussians (ASGs) is proposed as an option for highly anisotropic BRDFs. These approaches avoid precomputations and additional memory footprints for parametric BRDFs thanks to Wang’s on-the-fly approximation [WRG\*09]. Since our lobe-aware filtering requires only a camera position and G-buffer (i.e. depth, normal, reflectance, and BRDF parameter buffer), it is not only suitable for path tracing [Kaj86], but also deferred shading and non-ray tracing based methods such as voxel cone tracing [CNS\*11].

The contributions of this paper are as follows.

- Cross bilateral filtering is generalized in order to evaluate the similarity of distribution functions in a non-Euclidean space.
- Specular lobe-aware filtering and upsampling are introduced based on the above generalization for glossy indirect illumination. This filtering takes into account all-frequency materials and edges of lobe sharpness with less material-dependent parameter tuning than normal-aware filtering.
- Analytic approximations using SGs and ASGs for the specular lobe-aware weighting function are proposed. They are completely dynamic and has no additional storage cost for deferred shading pipelines with parametric BRDFs.

## 2. Background

### 2.1. Related work

**Bilateral filters.** The bilateral filter [SB97, TM98] is a single-pass edge-preserving filter often used for denoising purposes. The weighting function of this filter measures the photometric similarity between the target pixel and sample pixel. Since the photometric similarity also has noise, cross bilateral filtering [PSA\*04, ED04] uses additional noiseless images as guidance to measure the similarity. The same weighting functions as used in cross bilateral filtering are also used in joint bilateral upsampling [KCLU07]. The non-local (NL) means filter proposed by Buades et al. [BCM05b, BCM05a] is a generalized bilateral filter. For this filter, the weight for each pair of pixels is determined by the similarity of small image patches centered at the two pixels. This screen-space generalization is more robust, but slower than the conventional bilateral filtering. On the other hand, this paper generalizes cross bilateral filtering to evaluate the similarity of two parametric distribution functions in a high-dimensional non-Euclidean space whose parameters are given by guidance images. This is inexpensive and suitable for global illumination rendering.

**Global illumination with filtering.** Geometry-aware filtering is generally represented by cross bilateral filtering with a depth buffer and normal buffer. Such filtering has been often used for post-processing of interactive global illumination. Wald et al. [WKB\*02] used a discontinuity buffer for instant radiosity [Kel97] with interleaved sampling. To accelerate interleaved sampling and filtering, G-buffer splitting and hardware-supported filtering [SIMP06] were employed for deferred shading based instant radiosity [LSK\*07, REH\*11] and voxel cone tracing [CNS\*11, Mit12]. Knecht [Kne09] extended screen-space filtering to the temporal axis using reverse reprojection caching [NSL\*07]. Dammertz et al. [DSHL10] proposed the À-trous filter that incorporates a wavelet formulation into the bilateral filter. More recently, geometry-aware filtering based on guided image filtering [HST10] was also proposed for diffuse indirect illumina-

tion [BEM11]. Geometry-aware upsampling was employed for a frame buffer level-of-detail technique proposed by Yang et al. [YSL08]. Herzog et al. [HEMS10] introduced spatio-temporal upsampling using temporal coherence. Upsampling is also effective for voxel cone tracing on specular surfaces [Mit12], because rendering high-frequency materials is known to be slower than low-frequency materials.

However, geometry-aware filtering can produce an overblurring error for glossy surfaces. Such problems have been avoided by using more robust filtering methods with additional computation and memory footprints. Egan et al. [EHDR11] discussed screen-space bandlimits for glossy BRDFs in their frequency analysis, since a BRDF acts as a low-pass filter at each bounce. Based on this analysis, axis-aligned filtering [MWR12, MWRD13] determines screen-space bandlimits using the depth of secondary-rays and the BRDF at each pixel, but it does not focus on high-frequency BRDFs and high-frequency changes of BRDF parameters such as glossiness mapping. In addition, they use only the sharpness of an NDF, and neglect a view direction and surface normal dependent properties for their Blinn-Phong BRDF based bandlimit. For glossy surfaces lit from environment maps, Bagher et al. [BSS\*12] proposed an inexpensive upsampling technique by predicting bandwidths of BRDFs and environment maps with importance sampling. Segovia and Wald [SW10] blur incident light with a spherical harmonic approximation to preserve the edges of moderately glossy materials. Other than glossy surfaces and changes of BRDF parameters, there are difficult problems such as depth-of-field, motion blur, anti-aliasing, participating media, and high-frequency incident light. To reduce their noise, many methods have been developed especially for off-line rendering, such as error estimation for Gaussian or bilateral filtering [RKZ11, SD12, LWC12], NL means filtering based approaches [RKZ12, RMZ13, MJL\*13], wavelet based approaches [ODR09, KS13], and reconstruction using a light field [LALD12]. The above bilateral filtering based methods control the screen-space kernel bandwidth using the pixel information (e.g. BRDF and estimated error).

This paper focuses only on the problem of glossy surfaces for interactive or real-time applications. Unlike the above screen-space kernel bandwidth controlling methods, our method determines the directional kernel bandwidth utilizing BRDFs of the both target pixel and sample pixel. Therefore, our weighting function does not only take into account all-frequency materials, but also high-frequency changes of specular lobe sharpness. Although this simple approach without error estimation does not always produce the optimal weight unlike the above off-line error estimation based methods, it inexpensively avoids the risk of increasing overblurring errors in real-time frame rates.

**Spherical Gaussians.** SGs are often used for approximating environmental lighting in real-time applications [TS06, WRG\*09, IDN12], because SGs have closed-form solutions

for the integral, product, and product integral, which are fundamental operators to evaluate rendering integrals. Xu et al. [XSD\*13] proposed ASGs to represent anisotropic spherical functions. These ASGs have the approximate integral, product of two ASGs, and product integral of an ASG and SG. This paper employs SGs and ASGs to evaluate the similarity of two specular lobes.

## 2.2. Conventional geometry-aware filtering

In cross bilateral filtering (or upsampling), the filtered value  $\hat{I}(i)$  of target pixel  $i$  is estimated with a weighted averaging operation as follows:

$$\hat{I}(i) = \frac{\sum_j W(i, j) I(j)}{\sum_j W(i, j)}, \quad (1)$$

where  $I(j)$  is the value of the input image at sample pixel  $j$ , and weight  $W(i, j)$  is given as

$$W(i, j) = f(\mathbf{s}_i, \mathbf{s}_j) \prod_k w_k(i, j),$$

where  $\mathbf{s}_i$  is the screen-space position of pixel  $i$ ,  $f(\mathbf{s}_i, \mathbf{s}_j)$  is a screen-space filtering kernel, and  $k$ th weighting function  $w_k(i, j)$  is given by

$$w_k(i, j) = g\left(J_k(i) - J_k(j), \sigma_k^2\right), \quad (2)$$

where  $g$  is the Gaussian function,  $\sigma_k^2$  is the user-specified variance parameter, and  $J_k(i)$  is a noiseless guidance image. This weighting function measures a similarity between target pixel  $i$  and sample pixel  $j$ .

For global illumination, a depth buffer and normal buffer are often employed for the guidance images  $J_k(i)$ . This is because the shading similarity can be measured with these parameters when the BRDF is the Lambertian model with a constant reflectance. The weighting function can be defined as

$$W(i, j) = f(\mathbf{s}_i, \mathbf{s}_j) w_z(i, j) w_n(i, j),$$

$$w_z(i, j) = g\left(z_i - z_j, \sigma_z^2\right), \quad (3)$$

$$w_n(i, j) = g\left(\mathbf{n}_i - \mathbf{n}_j, \sigma_n^2\right), \quad (4)$$

where  $z_i$  and  $\mathbf{n}_i$  are the depth value and surface normal at the pixel  $i$  respectively. To avoid the influence of the spatially-varying reflectance on filtering, the reflectance and indirect illumination are separated and stored in different buffers, and blended after the filtering [SIMP06]. This geometry-aware filtering (or upsampling) is well established for low-frequency BRDFs. However, for specular surfaces, the assumption of geometry-aware filtering is violated. Therefore, a specular lobe similarity has to be taken into account, because the shading result is very sensitive to the lobe. Since specular lobes are distribution functions in directional space, cross bilateral filtering has to be generalized for such non-Euclidean distributions.

### 3. Distribution-Aware Filtering

In this section, cross bilateral filtering is generalized for the similarity of distributions. This paper represents the pixel value  $I(i)$  with the product integral of a pixel-independent function  $Q(\mathbf{x})$  and pixel-dependent distribution function  $a_i(\mathbf{x})$  as follows:

$$I(i) = o_i + \int_{\Omega} Q(\mathbf{x}) a_i(\mathbf{x}) d\mathbf{x}, \quad (5)$$

where  $\mathbf{x}$  is the position in a high-dimensional domain  $\Omega$  and  $o_i$  is a noisy offset. In this representation, since  $a_i(\mathbf{x})$  can be considered as a feature of the pixel, the similarity of pixels  $i$  and  $j$  is represented with the similarity of  $a_i(\mathbf{x})$  and  $a_j(\mathbf{x})$ . For bilateral filtering, let a screen-space position and luminance be  $\mathbf{s} \in \mathbb{R}^2$  and  $Y \in \mathbb{R}$  respectively, then  $\mathbf{x} = \{\mathbf{s}, Y\}$ ,  $\Omega = \mathbb{R}^3$ ,  $Q(\mathbf{x}) = Y$ , and  $a_i(\mathbf{x}) \propto g(Y - I(i), \tau^2) \delta(\mathbf{s} - \mathbf{s}_i)$ , where  $\tau^2$  is a user-specified constant value in conventional bilateral filtering, and  $\delta(\mathbf{s} - \mathbf{s}_i)$  is Dirac's delta function. For cross bilateral filtering, the parameters of  $a_i(\mathbf{x})$  are given by guidance images as:  $a_i(\mathbf{x}) \propto g(Y - J(i), \tau^2) \delta(\mathbf{s} - \mathbf{s}_i)$ .

The weighting function is represented based on the similarity between the two distributions  $a_i(\mathbf{x})$  and  $a_j(\mathbf{x})$ . This similarity is obtained by computing the product integral of two distributions. Since  $a_i(\mathbf{x})$  can have a delta function, this paper computes the product integral of smoothed distributions of  $a_i(\mathbf{x})$  and  $a_j(\mathbf{x})$  to control the kernel bandwidth using a user-specified smoothing kernel  $b(\mathbf{x}', \mathbf{x})$ . The smoothed distribution  $c_i(\mathbf{x})$  is given by the following convolution:

$$c_i(\mathbf{x}) = \int_{\Omega} a_i(\mathbf{x}') b(\mathbf{x}', \mathbf{x}) d\mathbf{x}', \quad (6)$$

and its normalized distribution is given as

$$p_i(\mathbf{x}) = \frac{c_i(\mathbf{x})}{\sqrt{\int_{\Omega} (c_i(\mathbf{x}'))^2 d\mathbf{x}'}}. \quad (7)$$

The similarity of two pixels can be represented by the product integral of  $p_i(\mathbf{x})$  and  $p_j(\mathbf{x})$  as follows:

$$q_{i,j} = \int_{\Omega} p_i(\mathbf{x}) p_j(\mathbf{x}) d\mathbf{x}, \quad (8)$$

Since  $p_i(\mathbf{x})$  and  $p_j(\mathbf{x})$  are normalized, the range of  $q_{i,j}$  is  $[0, 1]$ . When  $p_i(\mathbf{x})$  and  $p_j(\mathbf{x})$  are the same,  $q_{i,j} = 1$ . Finally, this paper additionally introduces a user-specified power parameter  $\beta$  to control the influence of the frequency of  $a_i(\mathbf{x})$  and  $a_j(\mathbf{x})$  for the weighting function as follows:

$$w(i, j) = q_{i,j}^{\beta}.$$

Cross bilateral filtering can be represented by this formulation, since this is a superset of the Gaussian based weighting function (Eq. (2)) as described in the supplemental material. In addition, this reformulation can use an individual variance parameter  $\tau_i^2$  given by a guidance image for each  $a_i(\mathbf{x})$ . When noise distribution is given by  $\tau_i^2$  for each pixel



**Figure 3:** Distribution-aware filtering with Gaussian functions represented with guidance images (b) and (c). The noise of the input image (a) is proportional to the variance (c). Our filtering (e) preserves the edges of the eye, while cross bilateral filtering (d) does not. Both  $\sigma_k^2$  and  $\beta$  are optimal for this image.

$i$ , this generalized filtering produces more accurate results than conventional cross bilateral filtering as shown in Fig. 3. Although similar results can be produced by controlling screen-space kernel bandwidth according to each variance parameter  $\tau_i^2$ , our weighting function detects the edges between two parameters  $\tau_i^2$  and  $\tau_j^2$  unlike the bandwidth controlling technique.

For conventional cross bilateral filtering, the luminance space is used to measure the photometric similarity, and each  $a_i(\mathbf{x})$  is parameterized by a constant variance in a Euclidean space. However, this reformulation supports an arbitrary distribution in a non-Euclidean space for each pixel. This feature is suitable for global illumination given by the integral in non-Euclidean space.

### 4. Specular Lobe-Aware Filtering and Upsampling

Indirect illumination is represented by the rendering equation [Kaj86] as follows:

$$L_r(\mathbf{y}, \Psi) = \int_{S^2} L(\mathbf{y}, \omega) \rho(\mathbf{y}, \Psi, \omega) \max(\mathbf{n} \cdot \omega, 0) d\omega,$$

where  $L_r(\mathbf{y}, \Psi)$  is the reflected radiance of indirect illumination at the position  $\mathbf{y}$  and direction  $\Psi$ ,  $L(\mathbf{y}, \omega)$  is the incident radiance of indirect illumination,  $\rho(\mathbf{y}, \Psi, \omega)$  is the BRDF, and  $\mathbf{n}$  is the surface normal. Therefore, the pixel value of the indirect illumination buffer is represented as

$$\begin{aligned} I(i) &= \varepsilon_i + \frac{L_r(\mathbf{y}_i, \Psi_i)}{R_i} \\ &= \varepsilon_i + \int_{\mathbb{R}^3} \int_{S^2} L(\mathbf{y}, \omega) \rho_i(\omega) \max(\mathbf{n}_i \cdot \omega, 0) \delta(\mathbf{y} - \mathbf{y}_i) d\omega dy, \end{aligned}$$

where  $\varepsilon_i$  is the estimation error,  $\mathbf{y}_i$  and  $\Psi_i$  are the position and direction at the first bounce respectively,  $R_i$  is the reflectance

of the BRDF. For simplicity, this paper represents the normalized BRDF at the first bounce as:  $\rho_i(\omega) = \frac{\rho(\mathbf{y}_i, \mathbf{y}_i, \omega)}{R_i}$ . For our filtering,  $\mathbf{o}_i = \mathbf{e}_i$ ,  $\mathbf{x} = \{\mathbf{y}, \omega\}$ ,  $\Omega = \{\mathbb{R}^3, S^2\}$ ,  $Q(\mathbf{y}, \omega) = L(\mathbf{y}, \omega)$  and  $a_i(\mathbf{y}, \omega) = \rho_i(\omega) \max(\mathbf{n}_i \cdot \omega, 0) \delta(\mathbf{y} - \mathbf{y}_i)$  in Eq. (5). This section introduces a directional weighting function  $w_\omega(i, j)$  instead of the normal-aware weighting function  $w_n(i, j)$ . The positional weighting function  $f(\mathbf{s}_i, \mathbf{s}_j) w_z(i, j)$  is the same as the conventional Gaussian based weighting function given by Eq. (3), because the positional distribution is represented by  $\delta(\mathbf{y} - \mathbf{y}_i) = \delta(\mathbf{s} - \mathbf{s}_i) \delta(z - z_i)$  which derives Eq. (3). In the following, for simplicity, we discuss only the directional distribution function which is represented by

$$a_i(\omega) = \rho_i(\omega) \max(\mathbf{n}_i \cdot \omega, 0).$$

However, computing the exact similarity can be expensive for this case. Since the shape of the reflection lobe depends on the BRDF model, many combinations of BRDFs have to be taken into account for the product integral of two lobes. Furthermore, there is not always an analytical solution of the product integral. Although the solution can be obtained by numerical precomputation based approaches, they require an additional memory footprint and access cost. In addition, long precomputation time can be necessary for a huge variety of materials. Therefore, this paper employs SG approximations for parametric BRDFs as described in the following subsection. Furthermore, we explain our weighting function is an extended formulation of the conventional normal-aware weighting function.

#### 4.1. Specular lobe similarity using SGs

In this subsection, we first derive a formulation of the specular lobe-aware weighting function using SG mixtures. Then, a more approximated weighting function is introduced for simplicity. This approximated formulation is a simple extension of the normal-aware weighting function and suitable for real-time applications. Although the SG representation induces additional approximation errors, this is practical for dynamic scenes with several materials, especially for filtering one-bounce global illumination.

**Weighting function.** An SG is a type of spherical function and is represented by the following equation:

$$G(\omega, \xi, \lambda) = g\left(\omega - \xi, \frac{1}{\lambda}\right) = \exp(\lambda((\omega \cdot \xi) - 1)),$$

where  $\xi$  is the lobe axis and  $\lambda$  is the lobe sharpness. Specular lobe  $a_i(\omega)$  is approximated with an SG mixture as

$$a_i(\omega) \approx \sum_l^{N_i} \mu_{i,l} G(\omega, \xi_{i,l}, \lambda_{i,l}).$$

where  $N_i$  is the number of lobes and  $\mu_{i,l}$  is the lobe amplitude. Each SG parameter is analytically obtained on-the-fly for parametric BRDFs such as the Cook-Torrance model,

Ward model and Blinn-Phong model [WRG\*09] (please refer to the supplemental material). Although our main target is such parametric BRDFs of dynamic scenes (i.e. BRDF parameters can change dynamically), measured BRDFs can also be used as an option for static materials. For measured BRDFs, we represent a BRDF using a microfacet model, and fit its NDF using an SG mixture in preprocessing [WRG\*09]. Since these lobe axes and sharpness are determined by using the surface normal and view direction, such view-dependent effects are taken into account for the weighting function. An SG is also employed for the smoothing kernel for specular lobes as

$$b(\omega', \omega) = G(\omega', \omega, \kappa), \quad (9)$$

where  $\kappa$  is the user-specified lobe sharpness. For global illumination using BRDF importance sampling, the noise produced at the first bounce is controlled by parameter  $\beta$ , while other noises are controlled by this  $\kappa$ . For filtering one-bounce indirect illumination,  $\kappa = \infty$  can be assumed. By using Iwasaki's approximate product integral of SGs [IDN12], the smoothed distribution  $c_i(\omega)$  is approximated with an SG mixture as follows:

$$c_i(\omega) \approx \sum_l^{N_i} \bar{\mu}_{i,l} G(\omega, \xi_{i,l}, \bar{\lambda}_{i,l}). \quad (10)$$

where  $\bar{\mu}_{i,l} = \frac{2\pi\mu_{i,l}}{\lambda_{i,l} + \kappa}$  and  $\bar{\lambda}_{i,l} = \frac{\lambda_{i,l}\kappa}{\lambda_{i,l} + \kappa}$ . By using the exact product integral of SGs derived in [TS06], the normalization factor of this distribution for Eq. (7) and the product integral Eq. (8) are computed. Thus, the approximated specular lobe-aware weighting function is obtained analytically as follows:

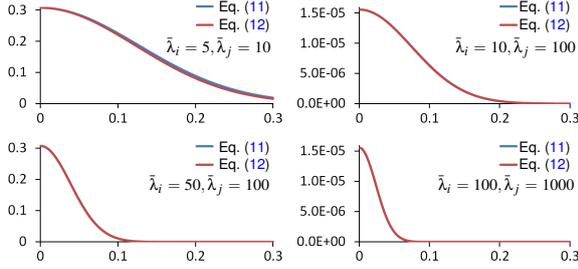
$$w_\omega(i, j) \approx \left( \frac{\alpha_{i,j}}{\sqrt{\alpha_{i,i}\alpha_{j,j}}} \right)^\beta, \quad (11)$$

where  $\alpha_{i,j} = \sum_l^{N_i} \sum_m^{N_j} \frac{\bar{\mu}_{i,l} \bar{\mu}_{j,m} \sinh(\|\bar{\lambda}_{i,l} \xi_{i,l} + \bar{\lambda}_{j,m} \xi_{j,m}\|)}{\exp(\bar{\lambda}_{i,l} + \bar{\lambda}_{j,m}) \|\bar{\lambda}_{i,l} \xi_{i,l} + \bar{\lambda}_{j,m} \xi_{j,m}\|}$ . Since this multi-lobe formulation has complexity  $O(N_i N_j)$ , small numbers for  $N_i$  and  $N_j$  are recommended.

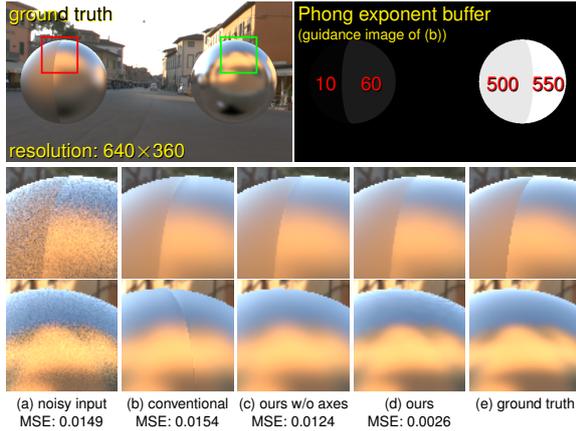
**More approximated weighting function.** For more time-sensitive applications, this paper restricts specular lobe  $a_i(\omega)$  to a single SG lobe (i.e.  $N_i = 1$ ). This can be done by merging distribution functions on-the-fly [Tok05, LWDB10] for multi-lobe BRDFs. For the single lobe representation, since the amplitudes  $\bar{\mu}_{i,l}$  and  $\bar{\mu}_{j,m}$  will be eliminated by normalization (Eq. (7)), the calculation of these amplitudes can be omitted. In addition, Iwasaki's approximate product integral is used for Eq. (7) and (8). Thus, the approximated weighting function is derived as

$$w_\omega(i, j) \approx \left( \frac{2\sqrt{\bar{\lambda}_i \bar{\lambda}_j}}{\bar{\lambda}_i + \bar{\lambda}_j} \right)^\beta G\left(\xi_i, \xi_j, \frac{\beta \bar{\lambda}_i \bar{\lambda}_j}{\bar{\lambda}_i + \bar{\lambda}_j}\right), \quad (12)$$

where the subscript  $l$  is omitted because  $N_i = N_j = 1$ . If lobe sharpness  $\bar{\lambda}_i$  and  $\bar{\lambda}_j$  are not small, Iwasaki's approxi-

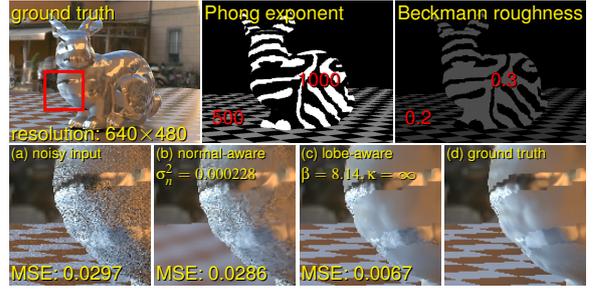


**Figure 4:** Plots of Eq. (11) and (12) for various lobe sharpness ( $N_i = 1$ ,  $\beta = 20$ ). The horizontal axis is  $\arccos(\xi_i \cdot \xi_j)$ .



**Figure 5:** Comparison for different glossiness. Each sphere in the target image (upper left) has two materials whose differences of Phong exponents are 50 (upper right). Although these differences are the same, the right sphere does not have noticeable edges unlike the left sphere. Conventional cross bilateral filtering (b) using the Phong exponent buffer overestimates the edges of the right sphere, and produces artefacts. Filtering (c) uses  $w_\omega(i, j) = \left( \frac{2\sqrt{\lambda_i \lambda_j}}{\lambda_i + \lambda_j} \right)^\beta$  which represents the similarity of lobe sharpness. This can avoid such overestimation. Moreover, taking lobe axes into account (d), errors are reduced without additional parameters.  $\sigma_k^2 = 763$  and  $\beta = 5.05$  (the optimal parameters for this image).

mation error is negligible as shown in Fig. 4. This approximation is more inexpensive than Eq. (11), and more suitable for real-time applications. This formulation is an extension of the normal-aware weighting function  $w_n(i, j) = g(\mathbf{n}_i - \mathbf{n}_j, \sigma_n^2) = G\left(\mathbf{n}_i, \mathbf{n}_j, \frac{1}{\sigma_n^2}\right)$ . Our lobe axes are not limited to normal vectors. The term  $\frac{\beta \bar{\lambda}_i \bar{\lambda}_j}{\lambda_i + \lambda_j}$  acts like previous BRDF based bandwidth controlling methods [MWRD13] in directional space, while  $\frac{2\sqrt{\lambda_i \lambda_j}}{\lambda_i + \lambda_j}$  represents the similarity of the specular lobe sharpness between two pixels. The normal-aware weighting function is the special case of our lobe-



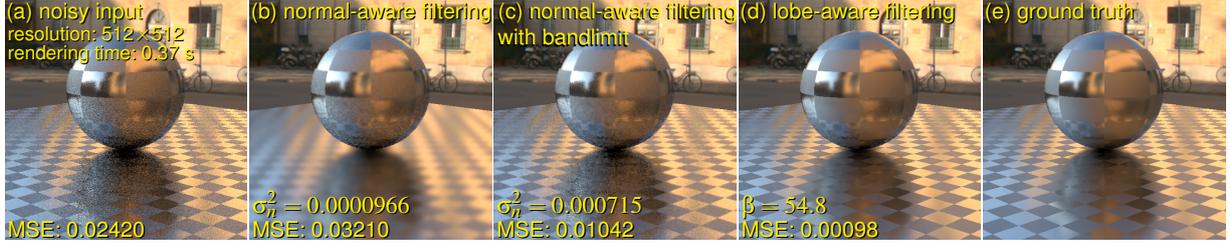
**Figure 6:** The scene has Kelemen BRDFs with the Phong NDF and Cook-Torrance BRDFs with the Beckmann NDF. For this scene, the distance of different BRDF parameters (i.e. Phong exponent and Beckmann roughness) cannot be measured directly. Unlike normal-aware filtering (b), this paper measures the similarity of such different BRDF models by using the SG representation (c). Both  $\sigma_n^2$  and  $\beta$  are optimal for this image.

aware weighting function for diffuse surfaces, because lobe axes are normal vectors and the sharpness is constant for diffuse lobes (please refer to the supplemental material).

**Comparison with conventional filtering.** When a scene has only a single BRDF model, although BRDF parameters (e.g. Phong exponent) can be directly used as guidance for conventional cross bilateral filtering, this can overestimate edges of lobe sharpness as shown in Fig. 5. In this figure, the difference of Phong exponents cannot represent actual visual differences. On the other hand, our metric is more suitable to represent such visual differences. Moreover, for one-bounce global illumination, our weighting function can be controlled by only the single parameter  $\beta$ , while the direct use of BRDF parameters increases the user-specified parameters in addition to  $\sigma_n^2$ . The number of parameters is different depending on the BRDF model. Furthermore, when a surface has several different BRDF models, the distance of these BRDF parameters cannot be measured directly (see Fig. 6). Multi-lobe BRDFs (e.g. layered materials) also have the same problem. Therefore, instead of the direct use of BRDF parameters as guidance images, this paper translates specular lobes into SGs, and measures their similarity.

## 4.2. Specular lobe similarity using ASGs

Instead of SGs, ASGs can be employed for the specular lobe similarity of highly anisotropic BRDFs. This approach has additional approximation errors especially for low-frequency lobes, and requires more expensive computation for each lobe. Therefore, it is proposed as an option for non-low-frequency BRDFs and non-time-sensitive applications.



**Figure 7:** Quality comparison for a scene with several sharpness materials. The sphere has a glossiness map. The floor has the Kelemen BRDF with the Phong NDF and Cook-Torrance BRDF with the Beckmann NDF for highly and roughly glossy tiles, respectively. Normal-aware filtering (b) induces overblurring and underblurring errors. Normal-aware filtering with screen-space bandlimit based on the BRDF at each target pixel (c) still has the noticeable error. Our filtering (d) produces closer images to the ground truth preserving edges, though scene has different BRDF models. Both  $\sigma_n^2$  and  $\beta$  are optimal for this image. For (c), the optimal screen-space bandlimit parameter is also chosen.

Xu's ASG is defined as the following equation:

$$\begin{aligned} \hat{G}(\omega, \xi_x, \xi_y, \xi_z, \lambda_x, \lambda_y) \\ = \max(\omega \cdot \xi_z, 0) \exp\left(-\lambda_x(\omega \cdot \xi_x)^2 - \lambda_y(\omega \cdot \xi_y)^2\right), \end{aligned}$$

where  $\xi_x, \xi_y, \xi_z$  are orthonormal vectors, and  $\lambda_x$  and  $\lambda_y$  are the bandwidth parameters. The approximate product integral of an ASG and SG is closed in ASG basis, and the approximate product integral of two non-low-frequency ASGs is also available (please refer to the supplemental material). Therefore, the weighting function using ASGs is derived in the same manner as the SG based weighting function. The specular lobe is approximated with an ASG mixture as

$$a_i(\omega) \approx \sum_l^{N_i} \mu_{i,l} \hat{G}(\omega, \xi_{x,i,l}, \xi_{y,i,l}, \xi_{z,i,l}, \lambda_{x,i,l}, \lambda_{y,i,l}).$$

An SG is also employed for the smoothing kernel as Eq. (9). Using Xu's approximate product integral of an ASG and SG, the smoothed distribution  $c_i(\omega)$  defined as Eq. (6) is closed in ASG basis as follows:

$$c_i(\omega) \approx \sum_l^{N_i} \bar{\mu}_{i,l} \hat{G}(\omega, \xi_{x,i,l}, \xi_{y,i,l}, \xi_{z,i,l}, \bar{\lambda}_{x,i,l}, \bar{\lambda}_{y,i,l})$$

where  $\bar{\mu}_{i,l} = \frac{2\pi\mu_{i,l}}{\sqrt{(2\lambda_{x,i,l}+\kappa)(2\lambda_{y,i,l}+\kappa)}}$ ,  $\bar{\lambda}_{x,i,l} = \frac{\lambda_{x,i,l}\kappa}{2\lambda_{x,i,l}+\kappa}$ , and  $\bar{\lambda}_{y,i,l} = \frac{\lambda_{y,i,l}\kappa}{2\lambda_{y,i,l}+\kappa}$ . By using the approximate product integral of two ASGs, the normalization of  $c_i(\omega)$  (Eq. (7)) and the lobe similarity (Eq. (8)) are analytically obtained. Thus, the weighting function is given as follows:

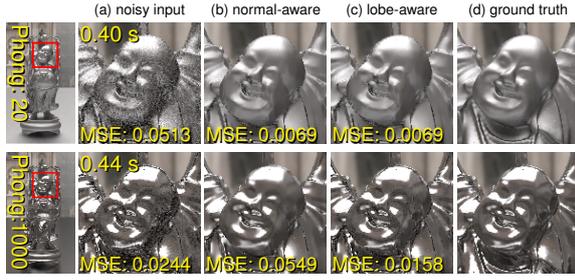
$$w_{\omega}(i, j) \approx \left( \frac{\gamma_{i,j}}{\sqrt{\gamma_{i,i}\gamma_{j,j}}} \right)^\beta, \quad (13)$$

where  $\gamma_{i,j} = \sum_l^{N_i} \sum_m^{N_j} \bar{\mu}_{i,l} \bar{\mu}_{j,m} \int_{S^2} \hat{G}_{i,l}(\omega) \hat{G}_{j,m}(\omega) d\omega$ , and  $\hat{G}_{i,l}(\omega) = \hat{G}(\omega, \xi_{x,i,l}, \xi_{y,i,l}, \xi_{z,i,l}, \bar{\lambda}_{x,i,l}, \bar{\lambda}_{y,i,l})$ .

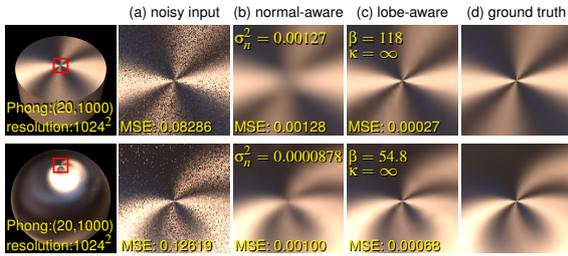
## 5. Experimental Results

In the following, we show the rendering results with our specular lobe-aware filtering and upsampling on an Intel<sup>®</sup> Xeon<sup>®</sup> W5590 and AMD Radeon<sup>™</sup> HD 6990. Similar to [Mit12], diffuse and specular indirect illumination are stored in different buffers for filtering and upsampling. In this section, all images are specular indirect illumination buffers, and reflectances and diffuse terms are separated from the images. To evaluate filtering accuracy, this section does not employ hardware-supported filtering [SIMP06] which produces an approximation error. Kelemen's microfacet BRDFs with the Phong NDF and Ashikhmin's NDF [AS00] are used for isotropic and anisotropic BRDFs, respectively. For Fig. 7, the Cook-Torrance BRDF with the Beckmann NDF is additionally used.

**Denosing for path tracing.** Figs. 7 and 8 show the quality comparison between the conventional normal-aware filtering (Eq. (4)) and our specular lobe-aware filtering (Eq. (12)) for path traced images which are generated by our unoptimized CPU implementation. In these experiments,  $\kappa = 100$  is used to denoise indirect illumination after the second bounce. The scene of Fig. 7 has a glossiness map (i.e. texture of Phong exponents) and two different BRDF models. Normal-aware filtering (b) cannot detect the edges of the glossiness and induces overblurring errors on these edges especially for the floor. Nonetheless, noises are still noticeable on the sphere, since small  $\sigma_n^2$  is used to narrow down the filtering bandwidth for the edges of the glossiness. Even with the use of screen-space bandlimit based on the BRDF at each target pixel similar to [MWRD13], the overblurring and underblurring error is still noticeable (c). On the other hand, our filtering (d) avoids these errors by evaluating the specular lobe similarity between the target pixel and sample pixel. Fig. 8 shows different glossy materials. In order to reduce overblurring error, conventional normal-aware filtering needs careful parameter tuning depending on the sharpness of materials. On the other hand, our filtering avoids overblurring errors



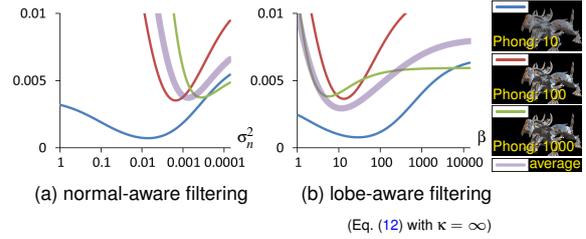
**Figure 8:** Quality comparison for glossy models (resolution:  $360 \times 640$ ).  $\sigma_n^2 = 0.00437$  and  $\beta = 13.1$  are used for both rows, which are optimal for the upper row. Normal-aware filtering (b) reduces the error for rough glossy surfaces (upper row), but for highly glossy surfaces (lower row) increases the error due to overblurring using the same parameter. On the other hand, our lobe-aware filtering (c) reduces the error for both images with the same parameter.



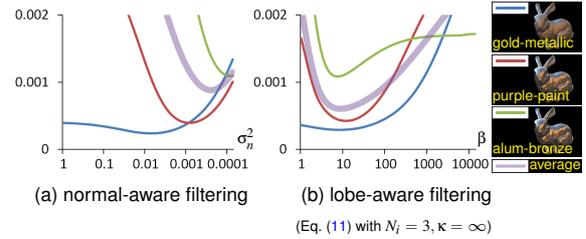
**Figure 9:** Filtering for the anisotropic BRDF. Our filtering (c) preserves edges of anisotropy and reduces the error, while conventional normal-aware filtering (b) does not. Both  $\sigma_n^2$  and  $\beta$  are optimal for these images.

for both scenes using the same parameters, since a sharper weighting function is used for sharper materials similar to the previous screen-space bandwidth controlling approach. This is suitable for path tracing with BRDF importance sampling, because it reduces noise for high-frequency materials. In addition, unlike the screen-space bandwidth controlling approach, our weighting function does not require additional material-dependent parameters, and is usable for more complex BRDFs (e.g. anisotropic BRDFs shown in the next paragraph).

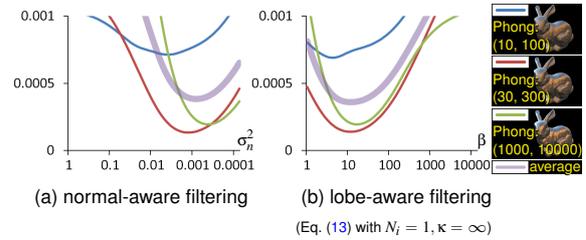
**Anisotropic BRDFs.** Fig. 9 shows filtering results of a highly anisotropic BRDF with MC environmental lighting on the GPU. Since normal-aware filtering (b) cannot detect the edges of anisotropy, overblurring errors are produced on the edges. Our specular lobe-aware filtering (c) using ASGs (Eq. (13) with  $N_i = 1$ ) preserves these visually important edges without increasing user-specified parameters. The upper and lower images have the same material and lighting condition. Nonetheless, the optimal  $\sigma_n^2$  of the lower image is



**Figure 10:** MSEs for various  $\sigma_n^2$  and  $\beta$  for single lobe BRDFs.



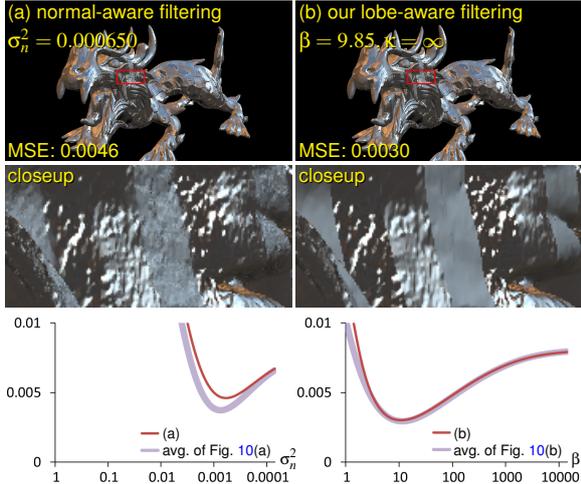
**Figure 11:** MSEs for various  $\sigma_n^2$  and  $\beta$  for multi-lobe BRDFs (3 SGs fitted from [MPBM03]).



**Figure 12:** MSEs for various  $\sigma_n^2$  and  $\beta$  for anisotropic BRDFs.

about 14 times smaller than the upper image, while the optimal  $\beta$  of the lower image is only about two times smaller.

**Parameter tuning.** We here employ MC environmental lighting on the GPU to evaluate the filtering accuracy with various parameters. Fig. 10 shows MSEs of normal-aware filtering (a) and our specular lobe-aware filtering (b) for various Phong exponents. Each image has only a single material. For these single material cases, each MSE of our filtering at the optimal parameter is almost the same or slightly larger than normal-aware filtering. They are thought to be caused by the SG approximation error. However, optimal parameters of lobe-aware filtering are more densely distributed than normal-aware filtering. Thus, the average of these errors (purple line) is reduced at the optimal parameter by using our filtering. In other words, our filtering can have lower errors than conventional normal-aware filtering for scenes with several materials. Fig. 11 shows MSEs using SG mixtures for multi-lobe BRDFs (3 SGs) fitted from measured data [MPBM03]. Fig. 12 shows MSEs for anisotropic



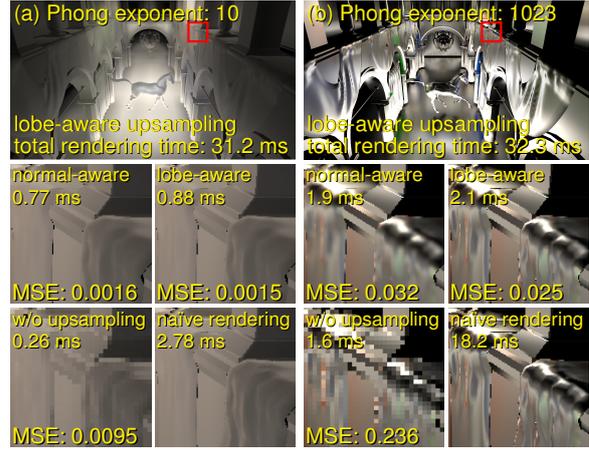
**Figure 13:** MSEs for various  $\sigma_n^2$  and  $\beta$  for a scene with several materials. All the materials of Fig. 10 (i.e. Phong exponent 10, 100, and 1000) are used.

**Table 1:** Filtering time for 1920×1080 resolution (ms).

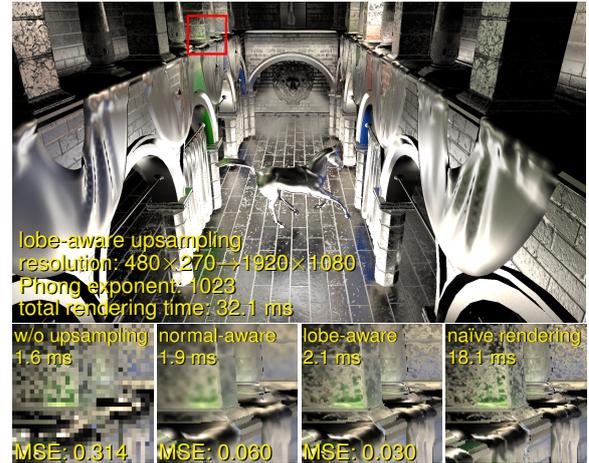
kernel radius	Eq. (4)	Eq. (12)	Eq. (11)	Eq. (13)
			( $N_i = 3$ )	( $N_i = 1$ )
8 pixels	2.43	3.29	15.2	17.7
16 pixels	4.56	6.03	27.6	30.8

BRDFs. In these cases, optimal parameters of lobe-aware filtering are also more densely distributed, and also reduce the averaged error at the optimal parameter. Fig. 13 shows MSEs for an image which has all the materials of Fig. 10. While normal-aware filtering (a) has a larger error than the average of Fig. 10(a), our lobe-aware filtering (b) produces a similar result to the average of Fig. 10(b). As described in earlier paragraphs, our filtering adapts all-frequency BRDFs and preserves edges of materials without increasing user-specified parameters for one-bounce global illumination. In addition, these densely distributed optimal parameters can reduce material-dependent parameter tuning.

**Performance.** Table 1 shows the filtering time on the GPU. Although our specular lobe-aware filtering using Eq. (12) is more expensive than normal-aware filtering (Eq. 4), it performs in real-time frame rates. The SG mixture based formulation (Eq. (11)) and ASG based formulation (Eq. (13)) are applicable for interactive frame rate applications. For filtering, the main bottleneck is the evaluation of weighing functions for many samples. On the other hand, for upsampling, since a smaller number of samples are evaluated, this cost cannot be a considerable problem compared to a denoising filter as shown in the next paragraph.



**Figure 14:** Glossy indirect illumination using [CNS\*11, HEMS10] with our weighting function ( $\sigma_n^2 = 0.01$ ,  $\beta = 20$ ,  $\kappa = 100$ , resolution: 480×270→1920×1080). The material changes dynamically, and images (a) and (b) are different frames in the same scene. Each computation time is the total of specular cone tracing and spatial upsampling.



**Figure 15:** Scene with normal maps and normal mipmap filtering [Tok05]. The quality difference between normal-aware upsampling and our lobe-aware upsampling is larger than Fig. 14(b).



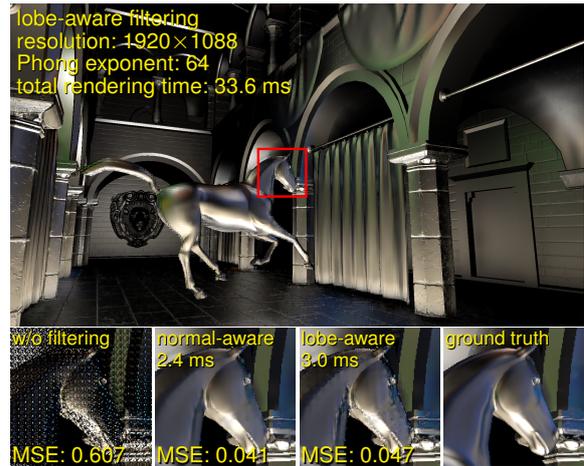
**Figure 16:** Adaptive recomputation for the spatio-temporal upsampled image (Fig. 15) evaluating low-weight pixels. The horse is a dynamic object. A threshold 0.25 is used for resampling in this experiment.

**Upsampling for voxel cone tracing.** Fig. 14 shows rendered images using voxel cone tracing and spatio-temporal upsampling with our weighting function. Since our method is not specific for ray tracing, it is applicable for such non-ray tracing based methods. Furthermore, a specular lobe is well represented with a cone, and only a few samples (e.g. 4 samples) are evaluated for spatial upsampling. For upsampling, since samples are sparse unlike filtering,  $\kappa = 100$  is used to smoothly interpolate between such sparse samples. In this experiment, we choose parameters  $\sigma_n^2$  and  $\beta$  heuristically, since this scene is completely dynamic (i.e. camera, light and objects can be moved). In addition, the material also changes dynamically. When the BRDF is low frequency (a), normal-aware upsampling with  $\sigma_n^2 = 0.01$  produces accurate results. However, when the material is changed to the higher-frequency BRDF (b), noticeable errors are produced with the same  $\sigma_n^2$ . On the other hand, our weighting function adapts all-frequency materials with the same  $\beta$ . Fig. 15 shows the same scene as Fig. 14 with normal maps. Since the G-buffer is rendered with normal mipmap filtering [Tok05], it does not only have high-frequency changes of normals, but also high-frequency changes of sharpness of NDFs. For this more difficult scene, the quality difference between our upsampling and normal-aware upsampling becomes larger. Our weighting function is more suitable than the normal-aware weighting function for such practical scenes.

**Adaptive sampling.** For bilateral filtering based methods, undersampling artefacts (e.g. noise and aliasing) can remain on high-frequency geometry by reducing the weights. These errors can be avoided by recomputing the pixel values adaptively as shown in Fig. 16, since these errors are simply detectable by evaluating the total weights (denominator of Eq. (1)) unlike overblurring errors. Our specular lobe-aware filtering and upsampling are more suitable than normal-aware methods for such simple adaptive sampling schemes. Normal-aware methods can produce overblurring errors which are more difficult to detect than undersampling errors in real-time frame rates.

## 6. Limitations and Future Work

**Discontinuous regions.** Bilateral filtering based methods can produce undersampling for discontinuous regions. Our method inherits this limitation from conventional geometry-aware filtering. For a silhouette with some curvature which can have high-frequency illumination appearance, our weighting function tends to produce undersampling artefacts. This is because specular lobe axes and sharpness are more different on such silhouettes. However, spatio-temporal filtering [Kne09] and upsampling [HEMS10] reduces undersampling for static objects. For dynamic objects, adaptive sampling is effective, because only rapidly moving objects are recomputed (Fig. 16). Even if adaptive sampling is not used, errors on such moving objects are less noticeable than static objects. Using such spatio-temporal filtering and



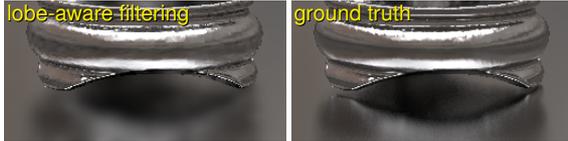
**Figure 17:** Instant radiosity using [REH\*11, Kne09] with our weighting function for a dynamic scene. Each computation time is the total of spatial filtering for specular and diffuse terms with hardware-supported filtering. For instant radiosity, our filtering accuracy can be lower than normal-aware filtering because BRDF importance sampling is unused and specular lobes are clamped.

adaptive sampling frameworks, the limitation of bilateral filtering is reducible.

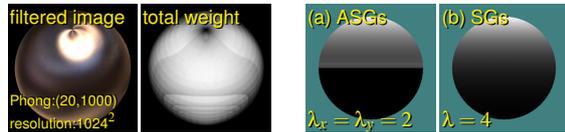
**Importance sampling.** Since the directional bandwidth of our filtering depends on the specular lobe sharpness, input images should be rendered using BRDF importance sampling similar to [BSS\*12]. Fig. 17 shows real-time instant radiosity with spatio-temporal filtering using our weighting function on the GPU. Unlike path tracing, since instant radiosity cannot use BRDF importance sampling, specular lobe-aware filtering can produce lower-quality images. In addition, clamping specular lobes in instant radiosity violates our BRDF based lobe representation. Our method inherits these limitations from previous BRDF based bandwidth controlling such as [MWRD13] whose screen-space bandlimit is determined by using the BRDF of the target pixel. However, unlike them, our filtering takes into account high-frequency changes of BRDF parameters such as glossiness mapping. In the future, we would like to develop interactive path tracing with our filtering on GPUs.

**Incident radiance.** Similar to normal-aware filtering, our weighting function neglects the high-frequency changes of incident radiance (Fig. 18). This limitation can be avoided using secondary-ray information [MWRD13] with an additional memory footprint. The combination of our weighting function and such methods is thought to be more robust.

**Optimal parameters.** Generally, since finding optimal parameters of filtering (i.e. balancing bias and variance) is a



**Figure 18:** Same scene as Fig. 8 bottom. Similar to normal-aware filtering, our method can produce an overblurring error neglecting the change of incident radiance.



**Figure 19:** The weight image using ASGs has banding artefacts due to the precision error.

**Figure 20:** Approximate product integral of two ASGs (a) and exact product integral of two SGs (b).

difficult problem, this has been often investigated for off-line methods such as error estimation. Since this paper focuses on real-time applications, it does not provide methods to find optimal parameters for arbitrary scenes. Our method can reduce only material-dependent parameter tuning. For other factors (e.g. noise amount and discontinuity), we would like to investigate the combination of error estimation and our method.

**Errors of ASGs.** The product of two ASGs, which is calculated to obtain the product integral, can have a large precision error for highly anisotropic lobes as shown in Fig. 19. To calculate this product, eigenvalues of a  $3 \times 3$  symmetric matrix are computed. In our implementation, these eigenvalues are analytically solved by using Cardano's formula, but single precision floating points can be insufficient for this formula. To reduce this precision error, double precision can be used at the expense of performance for current commodity GPUs. Another error of the product is caused by the lobe axis  $\xi_z$ . For example, when  $\xi_{z,i} \cdot \xi_{z,j} < 0$  and  $\lambda_x = \lambda_y$ , the ideal lobe axis of the product is not obtained by Xu's analytical approximation. Since this error is noticeable for low-frequency lobes unlike SGs (Fig. 20), ASGs are recommended for higher-frequency lobes. These limitations are not problematic only for our filtering, but also other applications such as lighting via analytical product integrals. To improve the practicality, we have to investigate more precise and efficient implementations of ASGs.

## 7. Conclusions

This paper has presented a generalization of cross bilateral filtering for specular lobe-aware filtering and upsampling to accelerate glossy indirect illumination rendering. This generalization evaluates the similarity of distribution functions

in a non-Euclidean space. When a scene has several materials, our specular lobe-aware weighting function reduces overblurring and underblurring errors by detecting the edges of BRDFs and adapting to all-frequency materials with less material-dependent parameter tuning than the normal-aware weighting function. This is effective for MC rendering with BRDF importance sampling. In addition, this paper has presented approximations using SGs and ASGs to calculate this weighting function analytically. SG based filtering performs in a few milliseconds without any precomputation or additional memory footprints for dynamic scenes with parametric BRDFs. Since our method requires only a camera position and G-buffer, it is also suitable for deferred shading and non-ray tracing based methods such as voxel cone tracing.

Our weighting function can be easily integrated with geometry-aware filtering based frameworks, since the normal-aware weighting function can be replaced with our weighting function for specular surfaces. This paper demonstrated spatio-temporal filtering, spatio-temporal upsampling, and adaptive sampling with our weighting function to improve real-time global illumination. It is thought to be also applicable for other filtering or upsampling methods such as  $\tilde{\Lambda}$ -trous filtering, frequency analysis based methods, and error estimation based methods.

## Acknowledgements

The light probe images and polygon models are courtesy of P. Debevec, The Stanford 3D Scanning Repository, M. Dabrovic, F. Meinel, R. W. Sumner, and J. Popovic. The authors would like to thank the anonymous reviewers for valuable comments and helpful suggestions.

## References

- [AS00] ASHIKHMIN M., SHIRLEY P.: An anisotropic phong brdf model. *J. Graph. Tools* 5, 2 (2000), 25–32. 7
- [BCM05a] BUADES A., COLL B., MOREL J.-M.: A non-local algorithm for image denoising. In *Proc. CVPR'05* (2005), pp. 60–65. 2
- [BCM05b] BUADES A., COLL B., MOREL J. M.: A review of image denoising algorithms, with a new one. *Multiscale Model. Simul.* 4 (2005), 490–530. 2
- [BEM11] BAUSZAT P., EISEMANN M., MAGNOR M.: Guided image filtering for interactive high-quality global illumination. *Comput. Graph. Forum* 30, 4 (2011), 1361–1368. 3
- [BSS\*12] BAGHER M. M., SOLER C., SUBR K., BELCOUR L., HOLZSCHUCH N.: Interactive rendering of acquired materials on dynamic geometry using bandwidth prediction. In *Proc. I3D'12* (2012), pp. 127–134. 3, 10
- [CNS\*11] CRASSIN C., NEYRET F., SAINZ M., GREEN S., EISEMANN E.: Interactive indirect illumination using voxel cone tracing. *Comput. Graph. Forum* 30, 7 (2011), 1921–1930. 2, 9
- [CT82] COOK R. L., TORRANCE K. E.: A reflectance model for computer graphics. *ACM Trans. Graph.* 1, 1 (1982), 7–24. 2
- [DSHL10] DAMMERTZ H., SEWITZ D., HANIKA J., LENSCH H. P. A.: Edge-avoiding  $\tilde{\Lambda}$ -trous wavelet transform for fast global illumination filtering. In *Proc. HPG'10* (2010), pp. 67–75. 2

- [ED04] EISEMANN E., DURAND F.: Flash photography enhancement via intrinsic relighting. *ACM Trans. Graph.* 23, 3 (2004), 673–678. 1, 2
- [EHDR11] EGAN K., HECHT F., DURAND F., RAMAMOORTHI R.: Frequency analysis and sheared filtering for shadow light fields of complex occluders. *ACM Trans. Graph.* 30, 2 (2011), 9:1–9:13. 3
- [HEMS10] HERZOG R., EISEMANN E., MYSZKOWSKI K., SEIDEL H.-P.: Spatio-temporal upsampling on the gpu. In *Proc. I3D'10* (2010), pp. 91–98. 3, 9, 10
- [HST10] HE K., SUN J., TANG X.: Guided image filtering. In *Proc. ECCV'10* (2010), pp. 1–14. 2
- [IDN12] IWASAKI K., DOBASHI Y., NISHITA T.: Interactive bi-scale editing of highly glossy materials. *ACM Trans. Graph.* 31, 6 (2012), 144:1–144:7. 3, 5
- [Kaj86] KAJIYA J. T.: The rendering equation. *SIGGRAPH Comput. Graph.* 20, 4 (1986), 143–150. 2, 4
- [KCLU07] KOPF J., COHEN M. F., LISCHINSKI D., UYTENDAELE M.: Joint bilateral upsampling. *ACM Trans. Graph.* 26, 3 (2007). 1, 2
- [Kel97] KELLER A.: Instant radiosity. In *Proc. SIGGRAPH'97* (1997), pp. 49–56. 2
- [Kne09] KNECHT M.: *Real-Time Global Illumination Using Temporal Coherence*. Master's thesis, Institute of Comput. Graph. Algorithms, Vienna U. Technol., 2009. 2, 10
- [KS13] KALANTARI N. K., SEN P.: Removing the noise in monte carlo rendering with general image denoising algorithms. *Comput. Graph. Forum* 32, 2 (2013), 93–102. 3
- [KSK01] KELEMEN C., SZIRMAY-KALOS L.: A microfacet based coupled specular-matte brdf model with importance sampling. In *EG'01 Short Presentations* (2001), pp. 25–34. 2
- [LALD12] LEHTINEN J., AILA T., LAINE S., DURAND F.: Reconstructing the indirect light field for global illumination. *ACM Trans. Graph.* 31, 4 (2012), 51:1–51:10. 3
- [LSK\*07] LAINE S., SARANSAARI H., KONTKANEN J., LEHTINEN J., AILA T.: Incremental instant radiosity for real-time indirect illumination. In *Proc. EGSR'07* (2007), pp. 277–286. 2
- [LWC12] LI T.-M., WU Y.-T., CHUANG Y.-Y.: SURE-based optimization for adaptive sampling and reconstruction. *ACM Trans. Graph.* 31, 6 (2012), 194:1–194:9. 3
- [LWDB10] LAURIJSSSEN J., WANG R., DUTRÉ P., BROWN B. J.: Fast estimation and rendering of indirect highlights. *Comput. Graph. Forum* 29, 4 (2010), 1305–1313. 5
- [Mit12] MITTRING M.: The technology behind the unreal engine 4 elemental demo. In *SIGGRAPH Courses* (2012). 2, 3, 7
- [MJL\*13] MOON B., JUN J. Y., LEE J., KIM K., HACHISUKA T., YOON S.-E.: Robust image denoising using a virtual flash image for monte carlo ray tracing. *Comput. Graph. Forum* 32, 1 (2013), 139–151. 3
- [MPBM03] MATUSIK W., PFISTER H., BRAND M., MCMILLAN L.: A data-driven reflectance model. *ACM Trans. Graph.* 22, 3 (2003), 759–769. 8
- [MWR12] MEHTA S. U., WANG B., RAMAMOORTHI R.: Axis-aligned filtering for interactive sampled soft shadows. *ACM Trans. Graph.* 31, 6 (2012), 163:1–163:10. 3
- [MWRD13] MEHTA S. U., WANG B., RAMAMOORTHI R., DURAND F.: Axis-aligned filtering for interactive physically-based diffuse indirect lighting. *ACM Trans. Graph.* 32, 4 (2013), 96:1–96:12. 1, 3, 6, 7, 10
- [NSL\*07] NEHAB D., SANDER P. V., LAWRENCE J., TATARCHUK N., ISIDORO J. R.: Accelerating real-time shading with reverse reprojection caching. In *Proc. Graph. Hardw. '07* (2007), pp. 25–35. 2
- [ODR09] OVERBECK R. S., DONNER C., RAMAMOORTHI R.: Adaptive wavelet rendering. *ACM Trans. Graph.* 28, 5 (2009), 140:1–140:12. 3
- [PSA\*04] PETSCHNIG G., SZELISKI R., AGRAWALA M., COHEN M., HOPPE H., TOYAMA K.: Digital photography with flash and no-flash image pairs. *ACM Trans. Graph.* 23, 3 (2004), 664–672. 1, 2
- [REH\*11] RITSCHEL T., EISEMANN E., HA I., KIM J. D., SEIDEL H.-P.: Making imperfect shadow maps view-adaptive: High-quality global illumination in large dynamic scenes. *Comput. Graph. Forum* 30, 8 (2011), 2258–2269. 2, 10
- [RKZ11] ROUSSELLE F., KNAUS C., ZWICKER M.: Adaptive sampling and reconstruction using greedy error minimization. *ACM Trans. Graph.* 30, 6 (2011), 159:1–159:12. 3
- [RKZ12] ROUSSELLE F., KNAUS C., ZWICKER M.: Adaptive rendering with non-local means filtering. *ACM Trans. Graph.* 31, 6 (2012), 195:1–195:11. 3
- [RMZ13] ROUSSELLE F., MANZI M., ZWICKER M.: Robust denoising using feature and color information. *Comput. Graph. Forum* 32, 7 (2013), 121–130. 3
- [SB97] SMITH S. M., BRADY J. M.: SUSAN - a new approach to low level image processing. *Int. J. Comput. Vis.* 23, 1 (1997), 45–78. 2
- [SD12] SEN P., DARABI S.: On filtering the noise from the random parameters in monte carlo rendering. *ACM Trans. Graph.* 31, 3 (2012), 18:1–18:15. 3
- [SIMP06] SEGOVIA B., IEHL J.-C., MITANCHEY R., PÉROCHE B.: Non-interleaved deferred shading of interleaved sample patterns. In *Proc. Graph. Hardw. '06* (2006), pp. 53–60. 2, 3, 7
- [SW10] SEGOVIA B., WALD I.: *Screen Space Spherical Harmonics Filters for Instant Global Illumination*. Tech. rep., Intel Corp., 2010. 3
- [TM98] TOMASI C., MANDUCHI R.: Bilateral filtering for gray and color images. In *Proc. ICCV'98* (1998), pp. 839–846. 2
- [Tok05] TOKSVIG M.: Mipmapping normal maps. *J. Graph. Tools* 10, 3 (2005), 65–71. 5, 9, 10
- [TS06] TSAI Y.-T., SHIH Z.-C.: All-frequency precomputed radiance transfer using spherical radial basis functions and clustered tensor approximation. *ACM Trans. Graph.* 25, 3 (2006), 967–976. 3, 5
- [WKB\*02] WALD I., KOLLIG T., BENTHIN C., KELLER A., SLUSALLEK P.: Interactive global illumination using fast ray tracing. In *Proc. EGWR'02* (2002), pp. 15–24. 2
- [WRG\*09] WANG J., REN P., GONG M., SNYDER J., GUO B.: All-frequency rendering of dynamic, spatially-varying reflectance. *ACM Trans. Graph.* 28, 5 (2009), 133:1–133:10. 2, 3, 5
- [XSD\*13] XU K., SUN W.-L., DONG Z., ZHAO D.-Y., WU R.-D., HU S.-M.: Anisotropic spherical gaussians. *ACM Trans. Graph.* 32, 6 (2013), 209:1–209:11. 3
- [YSL08] YANG L., SANDER P. V., LAWRENCE J.: Geometry-aware framebuffer level of detail. *Comput. Graph. Forum* 27, 4 (2008), 1183–1188. 3