From the ACM SIGGRAPH Asia 2012 conference proceedings.

Practical Hessian-Based Error Control for Irradiance Caching

Jorge Schwarzhaupt* UC San Diego
Henrik Wann Jensen† UC San Diego
Wojciech Jarosz‡ Disney Research Zürich

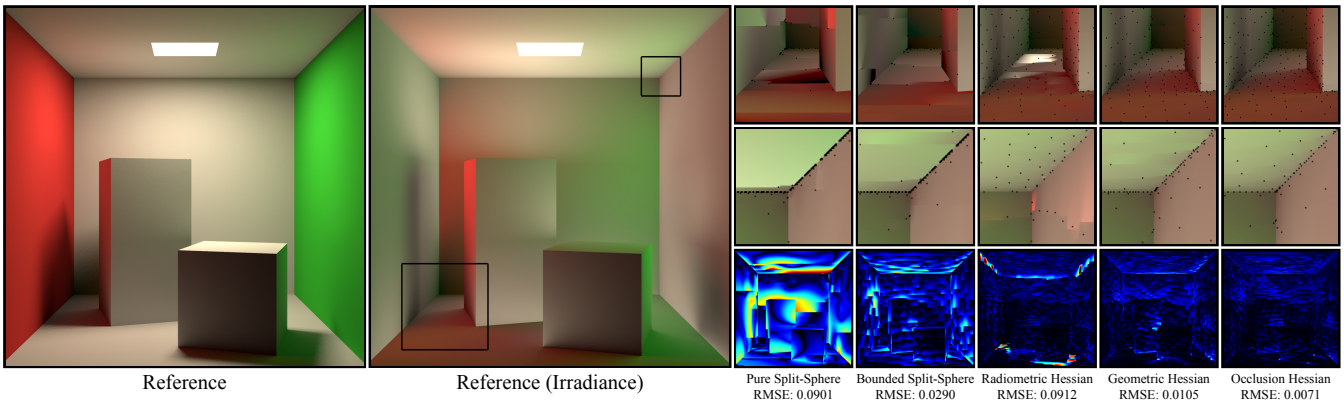

Figure 1: *Our new Occlusion Hessian significantly outperforms both the Pure and the Bounded Split-Sphere (clamped to the gradient and 150px max spacing) for irradiance caching. It also performs significantly better than the recently published occlusion-unaware Hessian error metrics by Jarosz et al. [2012].*

Abstract

This paper introduces a new error metric for irradiance caching that significantly outperforms the classic Split-Sphere heuristic. Our new error metric builds on recent work using second order gradients (Hessians) as a principled error bound for the irradiance. We add occlusion information to the Hessian computation, which greatly improves the accuracy of the Hessian in complex scenes, and this makes it possible for the first time to use a radiometric error metric for irradiance caching. We enhance the metric making it based on the relative error in the irradiance as well as robust in the presence of black occluders. The resulting error metric is efficient to compute, numerically robust, supports elliptical error bounds and arbitrary hemispherical sample distributions, and unlike the Split-Sphere heuristic it is not necessary to arbitrarily clamp the computed error thresholds. Our results demonstrate that the new error metric outperforms existing error metrics based on the Split-Sphere model and occlusion-unaware Hessians.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Raytracing

Keywords: Global illumination, Irradiance caching, Monte Carlo ray tracing

Links: DL PDF WEB

*jschwarz@cs.ucsd.edu
†henrik@cs.ucsd.edu
‡wjarosz@disneyresearch.com

1 Introduction

Computing indirect illumination in complex scenes both accurately and efficiently has been a long-standing challenge in computer graphics. Since its introduction by Ward et al. over twenty years ago, irradiance caching [Ward et al. 1988; Ward and Heckbert 1992] has become a very popular choice for accelerating the computation of diffuse indirect illumination. Irradiance caching exploits the fact that indirect illumination varies slowly across diffuse surfaces by computing it accurately only at a sparse set of locations and interpolating (or extrapolating) between the computed and cached irradiance values whenever possible.

The decision of whether to interpolate or not is central to the quality and efficiency of the algorithm. Ward et al. proposed to determine the cache point radii using the "Split-Sphere" heuristic which approximates an upper-bound on the rate of change of irradiance at the cache points. Unfortunately there are many common failure cases where this basic approach produces unacceptable results. This has led to several modifications and additions to the original metric to more robustly deal with such failure cases, by e.g. clamping to minimum/maximum radii, constraining the radius based on the gradient, enforcing the triangle inequality between cache records, and many more [Křivánek and Gautron 2009]. Though these additions can lead to a more robust approach, they are inherently trying to augment a sub-optimal heuristic and, in the process, introducing many more parameters, which makes it more difficult to control the algorithm.

Instead of modifying the original heuristic, Jarosz et al. [2012] recently proposed an alternative. While analyzing global illumination in 2D, they derived second order derivatives of irradiance for scenes with no occlusions, and showed how this could potentially be applied to obtain an improved error metric for irradiance caching. The core idea was to use a second-order Taylor expansion as a principled error term for the gradient extrapolation used in irradiance caching. By bounding the allowed error, they showed that the cache point radii could be derived from an irradiance Hessian instead of the Split-Sphere heuristic, without relying excessively on clamping and other corrections. Though they demonstrated that this idea shows promise, their preliminary investigation fell short of a full practical

algorithm due to the simplified setting and several practical limitations.

In this paper we propose a new, practical error heuristic for irradiance caching based on irradiance Hessians. Inspired by the preliminary results shown by Jarosz et al., we extend and improve the approach and perform the necessary evaluation to turn this idea into a practical error heuristic for general, complex scenes. The core of our error heuristic is a new formulation of irradiance Hessians which accounts for occlusion changes. This leads to higher-quality images compared to previously proposed techniques (see Figure 1), especially in scenes with complex indirect illumination containing occlussion and indirect penumbra as we will show in our results. Furthermore, our approach naturally supports elliptical cache records by exploiting the anisotropic error information contained within the Hessian. We also show that, compared to modifications of the Split-Sphere heuristic, our approach produces substantially better cache point distributions without relying on radius clamping and excessive parameter tuning. Finally, our formulation is not bound to any particular stratification of the hemisphere, which we show by replacing the traditional polar stratification with the lower-distortion concentric stratification [Shirley and Chiu 1997].

2 Previous Work

Irradiance Caching. Ward et al. [1988] first introduced irradiance caching as well as the Split-Sphere heuristic. During rendering, the cache is queried at every point that needs to be shaded. If valid cache records are found nearby, the cached indirect illumination is interpolated; otherwise, the irradiance is computed using brute-force Monte Carlo ray tracing and a new record is stored in the cache.

To control the placement of cache records, Ward et al. derived a conservative upper bound on the expected variation of irradiance with translation and rotation. They envisioned an imaginary "Split-Sphere" environment that induces the maximum possible change in irradiance for scenes without strong sources of indirect illumination. This error function is used to 1) determine the maximum valid radius for each cache record (enforcing a maximum error tolerance), and 2) derive the weights for each cache record's contribution during irradiance interpolation.

Later, Ward and Heckbert [1992] significantly improved the irradiance reconstruction by computing and storing not only the irradiance, but also its translational and rotation gradients at each cache record. When interpolation is possible, the irradiance at a shading point is approximated using a first-order Taylor extrapolation of the cached irradiance values in the local region. We show how to accurately compute the gradient for arbitrary sampling distributions.

Split-Sphere Extensions. Since its introduction, many modifications have been proposed to either handle more general reflection effects [Křivánek et al. 2005] or to improve the cache point distribution obtained from the Split-Sphere. We only discuss the most relevant prior work here and refer to the recent book by Křivánek and Gautron [2009] for a complete survey.

One of the most common practical modifications is to constrain the radius obtained from the Split-Sphere using 1) minimum and maximum bounds, and 2) the actual gradient computed from the sampled hemispherical environment [Larson and Shakespeare 1998]. Křivánek et al. [2006] also proposed a sophisticated refinement scheme to alleviate some of the shortcomings of the Split-Sphere by iteratively enforcing the triangle inequality and applying neighbor clamping. Herzog et al. [2009] created anisotropic cache records by simply squishing the cache records along the direction of the irradiance gradient by a user-controlled eccentricity parameter.

Illumination Derivatives and Error Control. Derivatives of illumination have also been considered in a wide variety of rendering contexts outside of irradiance caching. Arvo [1994] derived the irradiance Jacobian due to polygonal light sources and used it to compute isolux contours and find extrema of the irradiance field. Likewise, Holzschuch and Sillion [1995] derived radiosity gradients. More recently, Ramamoorthi et al. [2007] performed a first-order analysis of general lighting, deriving a principled net-visibility gradient usable for sparse sampling. Many researchers have investigated error control in the context of radiosity [Shirley 1991; Arvo et al. 1994; Holzschuch and Sillion 1998]. Holzschuch and Sillion [1995; 1998] derived gradients and Hessians of the point-to-area form factor which we utilize in our approach.

Jarosz et al. [2012] presented a 2D theory of light transport, derived occlusion-unaware Hessians of irradiance, and proposed to use these for error control in irradiance caching. They also showed that the irradiance Hessian provides enough information to simultaneously obtain both the orientation and eccentricity of elliptical cache records to optimally adapt to the local anisotropic irradiance curvature. Though their investigation was primarily theoretical, they nonetheless showed promising preliminary results.

Approach & Contributions. In this paper, we follow the same basic recipe as suggested by Jarosz et al. [2012], but propose a number of fundamental improvements that turn it into a practical, well-tested error control for irradiance caching in complex scenes:

- We derive an occlusion-aware Hessian for the error term, accounting for irradiance changes in scenes with occlusions.
- We also automatically obtain an occlusion-aware gradient which (unlike previous formulations) can easily be applied to general hemispherical sample distributions.
- We derive a more perceptually-motivated error criterion based on relative instead of absolute error, further reducing artifacts.
- We show how to make a radiometric (and not geometric) error metric robust, resulting in the first practical error metric for irradiance caching that is radiometrically meaningful.
- In contrast to Jarosz et al. [2012], we address how to take rotations of the surface normals into account, resulting in a complete and practical error-control formulation.
- We validate our proposed metric in complex scenes and compare to previous approaches, including the Split-Sphere with radius and gradient clamping, progressive refinement, and neighbor clamping [Křivánek et al. 2006].

3 Background

Irradiance caching exploits the fact that the reflected radiance L_o on diffuse surfaces can be expressed in terms of the irradiance E:

$$L_o(\mathbf{x}, \vec{\omega}_o) = \frac{\rho(\mathbf{x})}{\pi} E(\mathbf{x}, \vec{\mathbf{n}}), \tag{1}$$

where $\rho \in [0, 1)$ is the diffuse reflectance and $\vec{\mathbf{n}}$ is the surface normal. The irradiance is computed using Monte Carlo ray tracing:

$$E(\mathbf{x}, \vec{\mathbf{n}}) = \int_\Omega L_i(\mathbf{x}, \vec{\omega}_i)(\vec{\mathbf{n}} \cdot \vec{\omega}_i)\, d\vec{\omega}_i, \tag{2}$$

$$\approx \frac{1}{N} \sum_{i=1}^{N} \frac{L(\mathbf{x}, \vec{\omega}_i)(\vec{\mathbf{n}} \cdot \vec{\omega}_i)}{\mathrm{pdf}(\vec{\omega}_i)} = \frac{\pi}{N} \sum_{i=1}^{N} L(\mathbf{x}, \vec{\omega}_i), \tag{3}$$

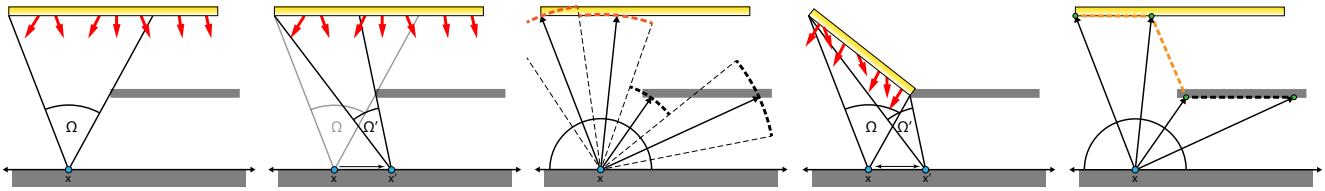

(A) Scene with Occlusions (B) Ward and Heckbert (C) Scene w/o Occlusions (D) Scene Approximation

Figure 2: *The scene depicted in (C) has the same irradiance first- and second-derivative as the one in (A), but there are no occlusions. (B) shows Ward and Heckbert's [1992] interpretation of the stratified sample information used to derive the irradiance gradient, while (D) shows the interpretation we perform in order to approximate (C).*

where the directions $\vec{\omega}_i$ over the hemisphere Ω follow a cosine-weighted distribution ($\text{pdf}(\vec{\omega}_i) = (\vec{\mathbf{n}} \cdot \vec{\omega}_i)/\pi$), canceling out the cosine term in the numerator.

We will follow the approach of Jarosz et al. [2012] and control the placement of cache records by defining the **total error** ϵ_i^t of a cache point i as the integrated difference between the correct irradiance and the extrapolated irradiance over the support of the cache point:

$$\epsilon_i^t = \iint_A |E(\mathbf{x}_i + \Delta\mathbf{x}) - E'(\mathbf{x}_i + \Delta\mathbf{x})| \, d\Delta\mathbf{x}, \tag{4}$$

where A is the area of support of the cache point, $\Delta\mathbf{x}$ is a 2D deviation in the tangent plane, and $E'(\mathbf{x}_i + \Delta\mathbf{x}) = E_i(\mathbf{x}_i) + \nabla_{\mathbf{x}} E_i(\mathbf{x}_i) \cdot \Delta\mathbf{x}$ is the first-order Taylor expansion of the irradiance. To make irradiance caching efficient, the goal is to use the largest possible area A while bounding the error to a certain threshold ϵ^t.

Since this expression depends on the ground truth irradiance E (which we want to avoid computing) we approximate it using a second-order Taylor expansion of the irradiance. The resulting expression relates the cache point error directly to the irradiance Hessian [Jarosz et al. 2012]:

$$\epsilon_i^t \approx \frac{1}{2} \iint_A |\Delta\mathbf{x}^\intercal \, \mathbf{H}_{\mathbf{x}}(E_i) \, \Delta\mathbf{x}| \, d\Delta\mathbf{x} \tag{5}$$

$$= \frac{1}{2} \iint_A \left(|\lambda_1| x^2 + |\lambda_2| y^2 \right) \, dy \, dx, \tag{6}$$

where the irradiance Hessian is a 2×2 matrix defining the second derivatives in the tangent space of the surface. The second line is reformulated in the coordinate system defined by the principle curvatures of irradiance (eigenvectors $\mathbf{v}^{\lambda_1}, \mathbf{v}^{\lambda_2}$ of the Hessian matrix with corresponding eigenvalues/curvatures λ_1 and λ_2) in the vicinity of the cache point.

By transforming to polar coordinates and then carrying out the integration, Equation (6) reduces to an error function that grows with the fourth-power of the deviation along the principle directions $\mathbf{v}^{\lambda_1}$ and $\mathbf{v}^{\lambda_2}$. Inverting this results in an elliptical cache point with radii

$$\left(R_i^{\lambda_1}, \, R_i^{\lambda_2} \right) \approx \sqrt[4]{\frac{4\epsilon^t}{\pi}} \left(\sqrt[4]{\frac{1}{\lambda_1}}, \, \sqrt[4]{\frac{1}{\lambda_2}} \right), \tag{7}$$

where ϵ^t is the error threshold specified by the user.

Jarosz et al. [2012] derived irradiance caching in a 2D setting and proposed the Hessian-based strategy for controlling error in irradiance caching. They derived a Hessian for irradiance assuming no occlusions in the scene. The Hessian based error metric is promising, but it becomes inaccurate in the presence of occlusions and consequently Jarosz et al. presented a Geometric Hessian that ignored the actual radiance values and instead provided a smooth cache distribution throughout the scene. The Geometric Hessian performs well in most scenes, but like the Split-Sphere heuristic it does not adapt in scenes where indirect illumination varies significantly. Consider a floor next to a wall that is either dark or bright. In the case of the dark wall, the sampling density on the floor can be quite low. In the case of the bright wall the floor needs more samples to account for the rapid local changes in the irradiance. The Geometric Hessian (as well as the Split-Sphere heuristic) will use the exact same sampling density in both cases. Having a radiometric error metric would enable the necessary variation in the sampling, and it would also enable a more precise control of the actual error in the extrapolated irradiance, which would lead to a higher quality irradiance estimation throughout the scene.

4 Practical Radiometric Hessian-based Error

4.1 Occlusion-Aware Translation Hessian

The core to our approach is an irradiance gradient and Hessian that accurately considers scenes with occlusions, like the one shown in Figure 2 (A). Ward et al. [1992] previously derived an accurate occlusion-aware irradiance gradient by interpreting the hemispherical samples as a coarse sampling of the surrounding geometry (Figure 2 (B)) and considering the change of strata areas due to occlusions as $\mathbf{x}$ is translated.

A natural next step would be to consider this geometric approximation of the environment, and compute the Hessian of the contribution of each strata. Jarosz et al. [2012] considered this while ignoring occlusions between strata. This makes the computation far simpler, but, unfortunately, ignoring occlusion derivatives is known to produce suboptimal results in irradiance caching [Křivánek et al. 2005].

This problem is in fact the same as computing the occlusion-aware form-factor Hessian of each polygonal stratum, which has been derived previous in the context of radiosity [Holzschuch 1996; Holzschuch and Sillion 1998]. Unfortunately, the discontinuous geometric approximation in Figure 2 (B) corresponds to a pathological case for these methods where the irradiance derivatives are always undefined[1]. Hence, these previous approaches do not produce usable results when applied to Ward et al.'s stratified environment.

Geometric Interpretation of Hemispherical Samples. To address this problem, we take a different approach. Our key insight is that we can convert a scene with occlusions (Figure 2 (A)) into a scene which is radiometrically equivalent at the shade point $\mathbf{x}$, but which contains no occlusions (Figure 2 (C)). Remarkably, this

[1]The derivatives are undefined whenever three or more polygon edges coincide at a single "apparent vertex" from the perspective of the shade point [Arvo 1994].

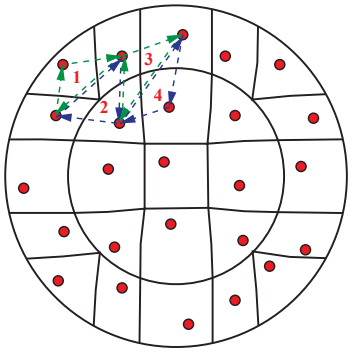

Figure 3: *The implicit connectivity information inherent in the stratified sampling allows efficient triangulation of the environment approximation. Final gather rays are shown as red points, and the implicit triangulation is shown in green.*

modified scene has the same irradiance, irradiance gradient, and irradiance Hessian at $\mathbf{x}$ as the original scene.

Equivalence arises because irradiance at point $\mathbf{x}$ is directly proportional to the solid angle subtended by the uniform diffuse emitter, as seen from $\mathbf{x}$. The gradient and Hessian are, then, proportional to the change in solid angle as $\mathbf{x}$ is translated. In the case of Figure 2 (C), the emitter subtends the same solid angles Ω and Ω', as it does in Figure 2 (A); hence, the irradiance, gradient and Hessian must also be the same. Of course for very large displacements this is not necessarily true, but we are only concerned with this equivalence at a differential scale.

Given this insight about equivalence, our procedure becomes clear. We interpret the hemispherical samples and construct a continuous, piecewise-linear approximation (Figure 2 (D)) of the surrounding scene geometry, which contains no occlusions from the point of view of $\mathbf{x}$. We can then compute an irradiance gradient and Hessian of this triangulated environment which accounts for occlusions occurring in the physical scene, while not explicitly considering occlusions in our calculations.

Assuming we have a triangulated representation of the hemispherical environment, the irradiance at $\mathbf{x}$ defined in Equation (3) can be re-expressed in terms of this approximate geometry as:

$$E(\mathbf{x}, \vec{\mathbf{n}}) \approx \sum_{j=1}^{M} L_{\triangle_j} F_{\triangle_j}(\mathbf{x}), \tag{8}$$

where M is the number of triangles in the tessellated hemisphere, and $L_\triangle$ and $F_\triangle(\mathbf{x})$ are the observed radiance and form-factor of triangle $\triangle$ at $\mathbf{x}$.

The gradient and Hessian of Equation (8) can be readily computed by summing the gradients and Hessians of the triangle form-factors:

$$\nabla_{\mathbf{x}} E \approx \nabla_{\mathbf{x}} \left(\sum_{j=1}^{M} L_{\triangle_j} F_{\triangle_j}(\mathbf{x}) \right) = \sum_{j=1}^{M} L_{\triangle_j} \nabla_{\mathbf{x}} F_{\triangle_j}(\mathbf{x}), \tag{9}$$

$$\mathbf{H}_{\mathbf{x}} E \approx \mathbf{H}_{\mathbf{x}} \left(\sum_{j=1}^{M} L_{\triangle_j} F_{\triangle_j}(\mathbf{x}) \right) = \sum_{j=1}^{M} L_{\triangle_j} \mathbf{H}_{\mathbf{x}} F_{\triangle_j}(\mathbf{x}). \tag{10}$$

Luckily, since our interpretation of the stratified samples induces no occlusion changes with translation, we can compute these form-factor derivatives independently for each triangle while ignoring occlusions. For completeness, we include the formulas for $\nabla_{\mathbf{x}} F_\triangle$ and $\mathbf{H}_{\mathbf{x}} F_\triangle$ as derived by Arvo [1994] and Holzschuch and Sillion [1998] in Appendix A.

The Hessian in Equation (10) is fully 3D, whereas we are interested only in the Hessian across the surface on which it was computed. To project the 3×3 Hessian matrix onto the surface, we define $\mathbf{a} = \mathbf{H}_{\mathbf{x}} E(\mathbf{x}) \mathbf{u}_1$ and $\mathbf{b} = \mathbf{H}_{\mathbf{x}} E(\mathbf{x}) \mathbf{u}_2$ and the tangential Hessian becomes

$$\mathbf{H}_{\mathbf{x}}^{2\times 2} = \begin{bmatrix} \mathbf{u}_1 \cdot \mathbf{a} & \mathbf{u}_1 \cdot \mathbf{b} \\ \mathbf{u}_2 \cdot \mathbf{a} & \mathbf{u}_2 \cdot \mathbf{b} \end{bmatrix} \tag{11}$$

where $\mathbf{u}_1$ and $\mathbf{u}_2$ are any two orthonormal vectors on the tangent space of the surface.

Constructing the Triangulated Environment. To construct our triangulated environment, we consider the hit distance/position for each sample ray we trace. We then go over all the samples and connect neighbors into triangles, defining a triangular mesh that covers the entire hemisphere. Though any triangulation algorithm could be used, we use the connectivity information implicit to the stratification to make this step efficient (see Figure 3). The mesh, along with the incoming radiance information, becomes a 3-dimensional snapshot of the scene, as it is seen from $\mathbf{x}$. Figure 4 illustrates this for a shade point in the Cornell box scene.

The most important features visible in Figure 4 are the large triangles that connect the surface of the occluder geometry to that of the wall. It is these polygons that encode the occlusion information of the scene from the point of view of the sampling location $\mathbf{x}$. While these triangles have a large geometric extent, they subtend a similar solid angle as other triangles, and thus their contribution to irradiance is not disproportionate. As $\mathbf{x}$ is translated, however, the solid angle of these large triangles will grow (or shrink) much faster than for other triangles, leading to a large gradient and Hessian due to occlusion changes.

To evaluate the irradiance gradient and Hessian, we need to define the incoming radiance $L_\triangle$ due to each of the triangles in the mesh. In 2D, as shown in Figure 2 (D), this corresponds to the incoming radiance stored at the vertex that is farthest away from the shading location $\mathbf{x}$. The same idea applies in 3D, though, we choose the farthest of the three vertices of each triangle (instead of two as in the 2D case). Intuitively, vertices that are farthest from $\mathbf{x}$ define the color of objects that become disoccluded during translation of $\mathbf{x}$. Ward and Heckbert [1992] used a similar idea to estimate the differential change in occlusion between strata when computing the gradient. As shown in Figures 5 and 6, this heuristic results in a good approximation of the true first and second derivatives of the irradiance for scenes with significant occlusions. Not surprisingly,

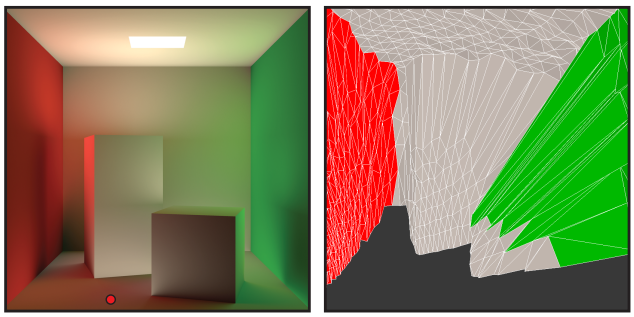

Figure 4: *Example of our mesh-based geometry approximation using the stratified sample data. The left image shows the Cornell Box with a red dot indicating the visualized location in the scene. The right image shows the triangle mesh approximation after applying our method, in this case using a total of 4096 gather rays. Note that rays that hit nothing are not shown.*

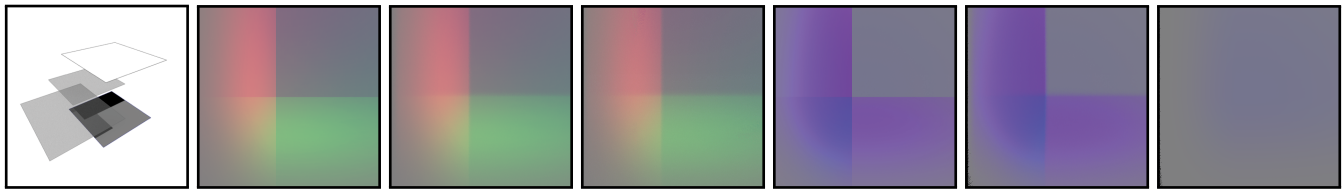

Figure 5: *The left figure shows the scene configuration: a bright area emitter was placed over a large plane, with two occluders located between the light and the plane at different distances. The three middle-left figures show a visualization of the first derivative of the irradiance across the bottom plane, using the mapping $(dE/dx, dE/dy) \rightarrow (r, g)$, while the three right figures show the eigenvalues of the irradiance Hessian using the mapping $(\lambda_1, \lambda_2, \lambda_3) \rightarrow$ (r,g,b). The numerical results use 4K sample rays per pixel.*

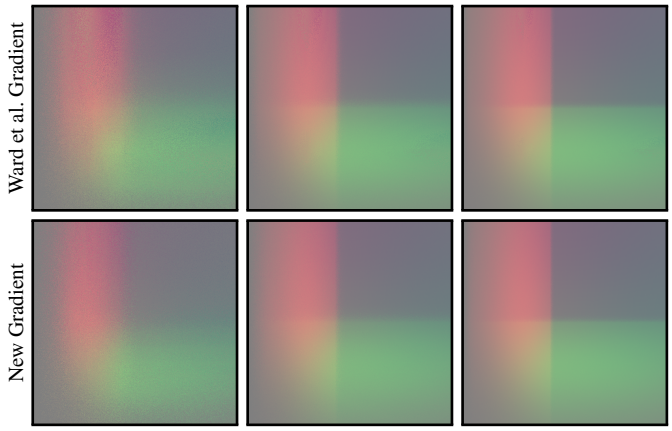

Figure 6: *Quality comparison between the classic Ward et al. irradiance gradient and our new formulation. While the gradient approximations differ slightly – especially at lower gather ray counts the Ward et al. gradient suffers from stronger artifacts – the qualitative results are similar.*

the previous occlusion-less Hessian derived by Jarosz et al. [2012] does not match the ground truth.

Our formulation allows us to implicitly account for occlusion changes in both the gradient and Hessian without resorting to more complex form factor computations which account for occlusions explicitly, and while avoiding the pathological cases that preclude their use in irradiance caching. Furthermore, our formulation has practical benefits over the standard approach proposed by Ward and Heckbert because it allows us to use arbitrary hemispherical distributions without re-deriving the method. Though we tried the standard longitude-latitude stratification, we use Shirley and Chiu's [1997] concentric mapping in our implementation because it produces more regular sampling patterns, leading to better results.

4.2 Cache Record Valid Region

To compute the anisotropic validity region for the cache records, we could now use Equations (6) and (7), while using our improved occlusion-aware Hessian. However, Equation (6) is suboptimal since it defines the error with respect to the absolute variation in irradiance. We would instead like to define a relative error term. This has a number of advantages: 1) it more strongly relates to our visual system's response to contrast instead of absolute changes in intensity and 2) it makes the error independent of the absolute scale of the scene and absolute intensity of the light sources used, making its parameters more compatible across scenes. To accomplish this, we modify the error formula by dividing the Hessian by the indirect irradiance computed at the sample location:

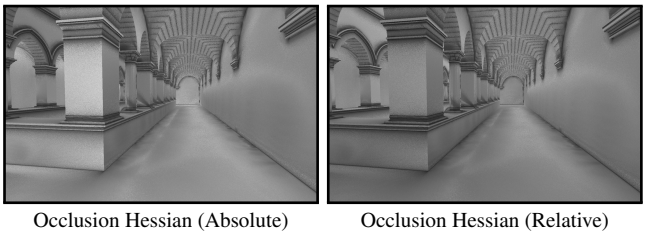

Figure 7: *Visualization of the radii estimated for the Occlusion Hessian using both absolute and relative measures for the total error. A relative measure allows larger radii in bright regions where absolute differences in irradiance induce lower perceived error.*

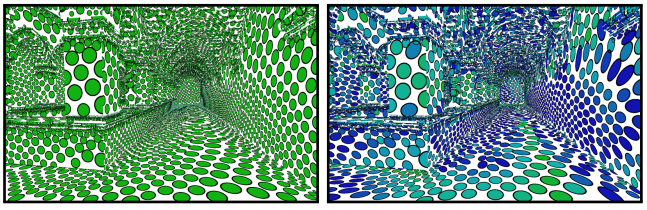

Figure 8: *Comparison of isotropic to anisotropic cache records for the relative Occlusion Hessian metric with the same threshold. The eccentricity of the cache records is visualized as the filled-in color, with green representing isotropic records and dark blue representing maximum anisotropy. Note that we have clamped the major axis at twice the length of the minor axis.*

$$\epsilon_i^r \approx \frac{1}{2} \iint_A \frac{|\Delta\mathbf{x}^\intercal \, \mathbf{H}_\mathbf{x}^{2\times 2}(E_i) \, \Delta\mathbf{x}|}{E_i} \, d\Delta\mathbf{x} \tag{12}$$

$$= \frac{1}{2} \iint_A \left(\frac{|\lambda_1|}{E_i} x^2 + \frac{|\lambda_2|}{E_i} y^2 \right) \, dy \, dx. \tag{13}$$

As in Equation (6), we perform the integration in Equation (13) in polar coordinates and invert the resulting expression. This leads to the following anisotropic cache record radius equation:

$$\left(R_i^{\lambda_1}, \, R_i^{\lambda_2} \right) = \sqrt[4]{\frac{4\epsilon^r E_i}{\pi}} \left(\sqrt[4]{\frac{1}{\lambda_1}}, \, \sqrt[4]{\frac{1}{\lambda_2}} \right), \tag{14}$$

where ϵ^r becomes the primary (relative) error control parameter in our algorithm. The primary differences compared to Equation (7) are that we use our new occlusion-aware Hessian, and that the cache point size is scaled by the fourth-root of the indirect irradiance (more absolute error can be tolerated in bright regions of the scene than in dark regions). In Figure 7 we compare the radius derived using a relative error versus an absolute error in the Sponza scene. The walls in the inner part of the Sponza courtyard have much brighter indirect irradiance, which allows for larger radii when accounting for relative error. For an absolute error, doubling the intensity of the lighting would modify the sample distribution (forcing the user to

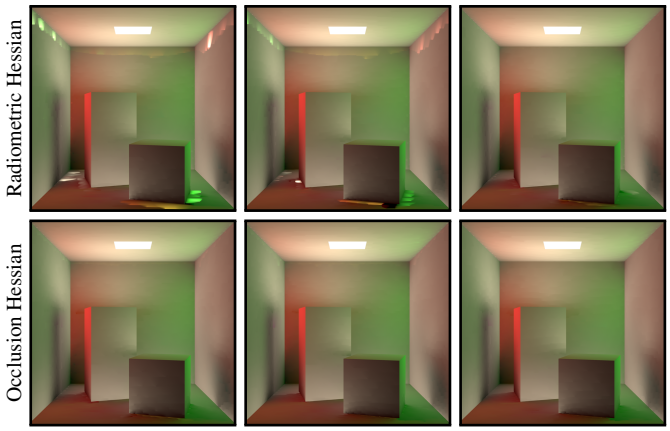

Figure 9: *Comparison of the Radiometric and Occlusion Hessian methods and our improvements for relative error and robustness. The first column shows the result when using an absolute error, while the second column uses a relative measure. The third column adds 1% of the total irradiance E_i to all the triangle radiances $L_{\triangle_j}$ prior to computing the Hessian. While this improves the result for the Radiometric Hessian method, it still retains a few distracting artifacts, that are not present for the Occlusion Hessian.*

manipulate the error threshold to obtain the same image), whereas this would have no effect on our relative metric. In Figure 8 we additionally visualize the effect of allowing anisotropic cache records compared to forcing isotropic records (by using the minimum of the two radii from Equation 14). Anisotropic records adapt their eccentricity to the local irradiance curvature, allowing for fewer cache records (16.5K vs. 20.5K) for the same error threshold. In practice, we clamp the major axis of the elliptical cache records at twice the length of the minor axis, to prevent artifacts that can arise in cases where the irradiance Hessian is locally very small in one direction.

Pathological Cases. One of the fundamental problems with any radiometric approach is that pathological cases (where all or most of the gather rays return black) can result in undefined or infinite radii. When computing full global illumination with many bounces, this is rarely a problem; however, when computing only a single bounce of indirect, this problem occurs more often. Jarosz et al. [2012] noted this issue, which forced them to revert to a completely geometric approach (the Geometric Hessian). To retain the benefits of a radiometric approach, we instead add 1% of the indirect irradiance E_i to all the triangle radiances $L_{\triangle_j}$ before computing the Hessian in Equation (10). In the case where *all* gather rays return black, we set all triangle radiances $L_{\triangle_j} = 1$. This has the effect of switching to an occlusion-aware geometric Hessian only for this special case. Note that we only add 1% during the Hessian computation and do not modify the irradiance stored in the cache.

Evaluation. In Figure 9 we show the effect of incorporating these improvements (moving from an absolute to a relative error, and adding 1% of the irradiance) when applied to both the occlusion-less Radiometric Hessian, and our Occlusion Hessian. These two changes (moving left to right) improve both methods, but the most striking improvement comes from accounting for occlusions in the Hessian itself (bottom row vs. top row). The Radiometric Hessian produces distracting artifacts due to interpolation across steep local gradients. This is because this error-control method is oblivious to occlusion changes, resulting in cache records that are too large for the local irradiance curvature. These artifacts disappear when using our Occlusion Hessian.

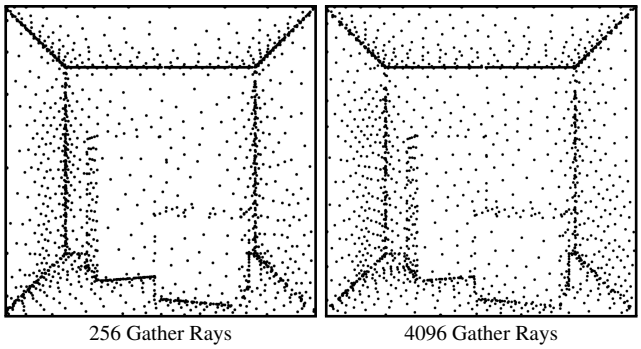

Figure 10: *Our method is robust to low gather ray counts, producing nearly identical cache point distributions with 256 (left) as with 4096 (right) gather rays per cache record.*

To validate that our radius computation is robust to low gather ray counts, in Figure 10 we compare the cache point distribution for the Cornell box scene using 256 gather rays and 4096 gather rays. Even with 16 times fewer gather rays, the cache point distribution remains qualitatively the same.

In Figure 11 we evaluate the Radiometric Hessian and Occlusion Hessian on a simple scene with an indirect occluder. The Occlusion Hessian successfully concentrates samples in regions of rapid irradiance change due to occlusions (the indirect penumbra regions).

4.3 Rotation Error

To account for changes in irradiance due to rotation, we could follow a similar procedure and derive a rotational irradiance Hessian by using rotational form-factor derivatives in Equations (9) and (10). Since the rotational derivatives only account for the change in the cosine foreshortening term as the surface normal is rotated, and no occlusion changes occur, the computation simplifies to the first and second derivatives of a cosine. Unfortunately, this approach does not work well. The problem is that occlusions in fact do occur: for any significant rotation, the change in the cosine factor is insignificant compared to the impact of regions from the lower hemisphere becoming disoccluded and contributing to the irradiance. In effect, rotational derivates in the upper hemisphere ignore all occlusions (the occlusions along the hemispherical boundary), and this has by far the biggest impact on the change in irradiance. Unfortunately, since we can only reliably sample the upper hemisphere, we have little hope of detecting what is under the upper hemisphere to accurately predict the change in irradiance.

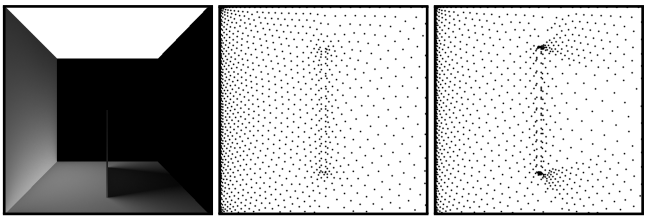

Figure 11: *A simple scene with indirect occlusions (left) where we visualize the emitted radiance on the emitter above and the indirect illumination on the other surfaces. The Radiometric Hessian (middle) ignores occlusions, producing a relatively uniform distribution whereas our Occlusion Hessian method (right) successfully adapts to the irradiance change near the penumbra region.*

In the face of these challenges, we opt for a simple but intuitive strategy as originally proposed by Tabellion and Lamorlette [2004]. We enforce a maximum deviation angle $\Delta\vec{\mathbf{n}}_{\max}$ that is allowed during extrapolation: no cache records are allowed to extrapolate beyond an e.g. 10° deviation in surface normal. We set the default value to $\Delta\vec{\mathbf{n}}_{\max} = 0.2$ radians, but allow the user to modify this if needed.

4.4 Irradiance Storage and Interpolation

With each cache point i we store: the location and normal $(\mathbf{x}_i, \vec{\mathbf{n}}_i)$, the irradiance (E_i), irradiance gradients $(\nabla_{\mathbf{x}} E_i,\ \nabla_{\vec{\mathbf{n}}} E_i)$, the two anisotropic radii $(R_i^{\lambda_1},\ R_i^{\lambda_2})$, and the corresponding eigenvectors $(\mathbf{v}_i^{\lambda_1}, \mathbf{v}_i^{\lambda_2})$.

Note that computing the eigenvalue decomposition for $\mathbf{H}_{\mathbf{x}}^{2\times 2}$ in Equation (13) results in eigenvectors in the tangent-space defined by vectors $\mathbf{u}_1$ and $\mathbf{u}_2$. To make these usable during interpolation, we transform each tangent-space eigenvector $(\mathbf{v}_i')$ into world-space coordinates $(\mathbf{v}_i = [\mathbf{u}_1 \cdot \mathbf{v}_i', \mathbf{u}_2 \cdot \mathbf{v}_i']^\intercal)$ before storing it in the cache.

To interpolate cache records, we use the same first-order extrapolation strategy as proposed by Ward and Heckbert [1992], but with a modified weighting function:

$$E(\mathbf{x}, \vec{\mathbf{n}}) \approx \frac{\sum_{i \in S} \mathrm{w}_i(\mathbf{x})(E_i + \nabla_{\mathbf{x}} E_i \cdot \Delta\mathbf{x}_i + \nabla_{\vec{\mathbf{n}}} E_i \cdot \Delta\vec{\mathbf{n}}_i)}{\sum \mathrm{w}_i(\mathbf{x})} \quad (15)$$

where $\Delta\mathbf{x}_i = \mathbf{x} - \mathbf{x}_i$ and $\Delta\vec{\mathbf{n}}_i = \vec{\mathbf{n}} \times \vec{\mathbf{n}}_i$ are the translational and rotational deviation of the shading location $\mathbf{x}$ with normal $\vec{\mathbf{n}}$ to the cache point.

The set S includes all cache records whose valid regions (Equation (14)) overlap $\mathbf{x}$ and whose normal deviation is less than $\Delta\vec{\mathbf{n}}_{\max}$. To weight the cache records, we propose to use a kernel weighting in both translation and orientation that falls off to zero at the boundaries:

$$\mathrm{w}_i(\mathbf{x}) = k(1 - t_{\mathbf{x}}, 0, 1) \times k(t_{\vec{\mathbf{n}}}, \cos\Delta\vec{\mathbf{n}}_{\max}, 1), \quad (16)$$

where k is a simple tent filter, $k(t, t_{\min}, t_{\max}) = \frac{t - t_{\min}}{t_{\max} - t_{\min}}$, and

$$t_{\mathbf{x}} = \sqrt{\left[\frac{\Delta\mathbf{x}_i \cdot \mathbf{v}_i^{\lambda_1}}{R_i^{\lambda_1}}\right]^2 + \left[\frac{\Delta\mathbf{x}_i \cdot \mathbf{v}_i^{\lambda_2}}{R_i^{\lambda_2}}\right]^2 + \left[\frac{\Delta\mathbf{x}_i \cdot \mathbf{n}_i}{R_i^{\lambda_1}}\right]^2}, \quad (17)$$

$$t_{\vec{\mathbf{n}}} = (\vec{\mathbf{n}} \cdot \vec{\mathbf{n}}_i), \quad (18)$$

are the translational and rotational distances. We use the cache record only if $\mathrm{w}_i(\mathbf{x})$ is greater than zero. Note that k could easily be replaced with a higher-order smooth kernel that falls off to zero, but we did not find this to provide a significant benefit.

Extension to Radiance Caching. Extending our method to radiance caching [Křivánek et al. 2005] would be relatively straightforward. Since radiance caching re-uses irradiance caching's Split-Sphere for error control, we could trivially replace this with our Occlusion Hessian term. For a standard polar stratification, we could project both the radiance and gradient onto spherical harmonics just as in previous work [Křivánek et al. 2005]. If a concentric mapping is used, projecting the radiance would still be trivial, but some work would be needed to derive a projection of the translational radiance gradient onto spherical harmonics.

5 Results

We implemented the new error metric by modifying the native Irradiance Cache implementation in version 2 of PBRT [Pharr and Humphreys 2010]. All results were rendered using 4 samples per pixel on a PC with a 2.66Ghz Intel Core i7-920 CPU. Our implementation of the Split-Sphere heuristic follows Ward's [1988] original formulation and uses the harmonic mean distance. For the Bounded Split Sphere we clamp this radius to a maximum pixel size and to the computed irradiance gradient (increasing the sample density in areas where the first-order gradient magnitude exceeds the Split-Sphere prediction). All methods use 4096 rays per irradiance sample, and we enforce a minimum cache point radius equivalent to the projected size of 1 pixel at the sample location, except where noted. For the Hessian based methods there is no upper bound for the cache point radii. All Hessian based results use the default value $\Delta\vec{\mathbf{n}}_{\max} = 0.2$ radians as the maximum normal deviation angle. We report the render times in Table 1.

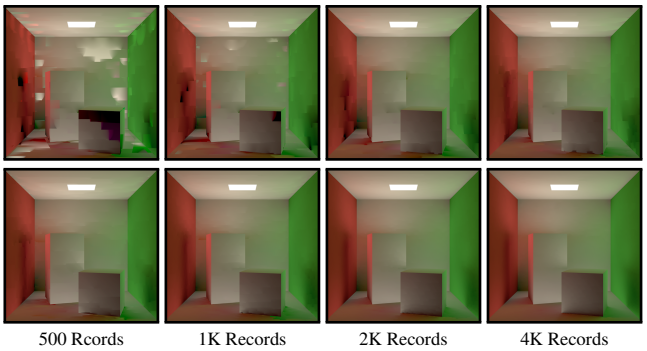

Figure 12: *Comparison of the convergence behavior for the Bounded Split-Sphere (top) and our Occlusion Hessian (bottom).*

Figure 1 shows a comparison of the Occlusion Hessian error control method against the Split-Sphere Heuristic as well as the Geometric and Radiometric Hessian metrics defined by Jarosz et al. [2012]. All images were rendered at a resolution of 1800x1800 pixels with a single bounce of indirect illumination. The error thresholds were adjusted for each method to enable the use of 1700 +/- 2% cache records. With just 1700 cache records the Split-Sphere method shows severe interpolation artifacts. Even with a maximum radius of 150px (a lower threshold is not possible with 1700 cache records) the Bounded Split-Sphere method shows significant errors throughout the rendered image. The Radiometric Hessian works well in lit regions, but has artifacts in shadowed regions. The Geometric Hessian works well in most of the scene, but shows interpolation artifacts on the walls and the ceiling. The Occlusion Hessian produces the best overall result, resolving the detail in the shadow while maintaining a sufficient density of samples on the walls and ceiling to reconstruct the irradiance. For the simple geometry in this scene the Occlusion Hessian added roughly 20% overhead to the computation time. For our more complex scenes this overhead becomes negligible as the time is dominated by ray tracing.

Figure 12 compares the behavior of our Occlusion Hessian method to the Bounded Split-Sphere heuristic with increasing cache point densities. All images used 4096 gather rays, and the Split-Sphere

Table 1: *Rendering times for the scenes and methods presented.*

Method	**Cornell Box** Figure 1	**Sponza** Figure 15	**San Miguel** Figure 17
Split-Sphere	01:02.3	53:08.4	07:15:22.4
Radiometric Hessian	01:06.5	n/a	n/a
Geometric Hessian	01:04.7	n/a	n/a
Occlusion Hessian	01:16.5	54:35.6	07:17:43.7

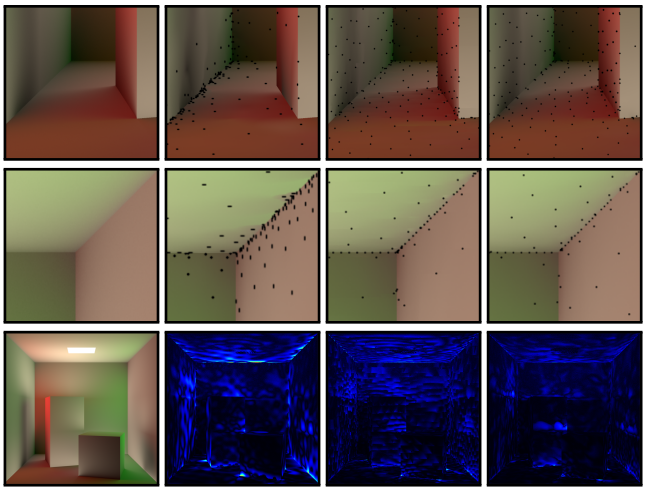

Reference | Enhanced Split-Sphere 123.4s RMSE: 0.0117 | Occlusion Hessian 76.5s RMSE: 0.0071 | Occlusion Hessian (O) 122.3s RMSE: 0.0049

Figure 13: *We show the result of adding progressive refinement, neighbor clamping and an overture pass [Křivánek and Gautron 2009] to the Split-Sphere, and compare to our Occlusion Hessian method without and with an overture pass. While these additions significantly improve the Split-Sphere result, render time is increased, and it still does not match the quality achieved by the Occlusion Hessian without an overture pass.*

used a maximum radius of 150px (a tighter bound was not possible for very limited cache point counts). Note how the Occlusion Hessian produces a high-quality result even with 500 cache records, while the Split-Sphere suffers from large artifacts even with 1000 records. These artifacts occur because the strong constraint on the record count requires setting such a high error threshold that cache records from the ceiling are used on the walls, and vice-versa.

In Figure 13 we extended the Split-Sphere with progressive refinement, neighbor clamping, and an overture pass [Křivánek et al. 2006]. Rendering of the enhanced Split-Sphere image took 123.4s while the Occlusion Hessian rendering, with the overture pass, took 122.6s. The quality obtained by the Occlusion Hessian is visibly better, even though the improvement for the Split-Sphere is significant. Note that the overture pass has less impact on the Occlusion Hessian as the error metric more closely follows the actual error in the irradiance.

The Geometric Hessian is generally robust in most scenes, but since it is a purely geometric metric, like the Split-Sphere it does not adapt to changes in the illumination. In Figure 14 the Cornell box has been changed such that the back wall is completely black. This adds contrast to the indirect illumination, which the Occlusion Hessian is able to detect and consequently it produces a higher quality output with fewer artifacts than the Geometric Hessian.

Figure 15 shows close-ups of Sponza rendered using the Occlusion Hessian and the Split Sphere heuristic (both unbounded and with a maximum radius of 20 pixels). For this scene we used a photon map to add multiple bounces of indirect illumination. The irradiance cache is used to gather the irradiance at the first diffuse surface seen by the eye, and we used 32K +/- 1% cache records for the full images. The resulting images show how the Split-Sphere heuristic has interpolation artifacts in the shadows on the columns, while the Occlusion Hessian is much better at reconstructing the lighting details. Figure 16 shows a direct visualization of the cache record radii (top), and the record support footprints (bottom), for the Sponza

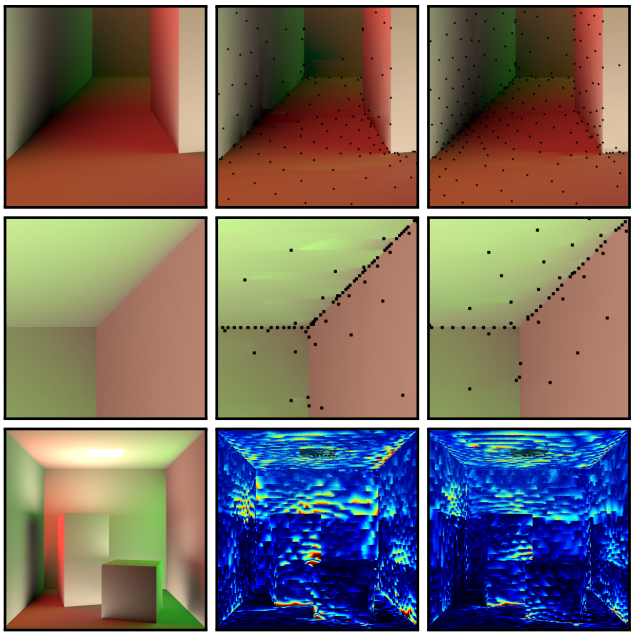

Reference | Geometric Hessian | Occlusion Hessian

Figure 14: *Irradiance rendering of the Cornell Box after changing the back wall so it has an albedo of 0. The Geometric Hessian does not take the radiometry of the scene into account, so its cache point distribution is the same as for the standard Cornell Box, which is sub-optimal. Our new Occlusion Hessian adapts its distribution to match the radiometry in this new configuration, producing a higher quality result. Note that we have scaled the difference image color range to ease comparison.*

scene with the various methods. Note that we visualize the geometric mean of the two radii for the anisotropic methods, and cache record outlines are drawn at 45% of their maximum size to avoid overlap in the visualization. The Split-Sphere images show that additionally bounding the cache point radii by the true first-order gradient leads to smaller radii in areas of high irradiance variance, such as the indirect penumbras visible at the base of the columns. This behavior is, as expected, mirrored by the Occlusion Hessian, which defines the radii in proportion to the rate of change of irradiance. However, the contrast is much higher for the Split-Sphere images – the radii variations are much larger – which explains the worse results obtained from this metric (e.g. sampling the scene outside of the penumbra regions leads to cache points with very large radii and the penumbra regions become severely under-sampled).

Figure 17 shows a view of the PBRT San Miguel scene. The methods were tuned so that 50K cache points (+/- 1%) would be produced. A photon map was used for the secondary light bounces, with 1 million photons stored. Our Occlusion Hessian approach was able to properly define all indirect shadows in the image, and exhibits fewer interpolation artifacts in areas of the scene where the irradiance varies slowly. Rendering times for the Bounded Split-Sphere and the Occlusion Hessian were within 0.5% of each other.

To facilitate detailed comparison, uncropped, full-resolution images for all results are available in the supplemental material.

6 Conclusion & Future Work

In this paper we have described a new method for controlling the approximation error in irradiance caching. The new method is based on computing the second derivative of irradiance - the trans-

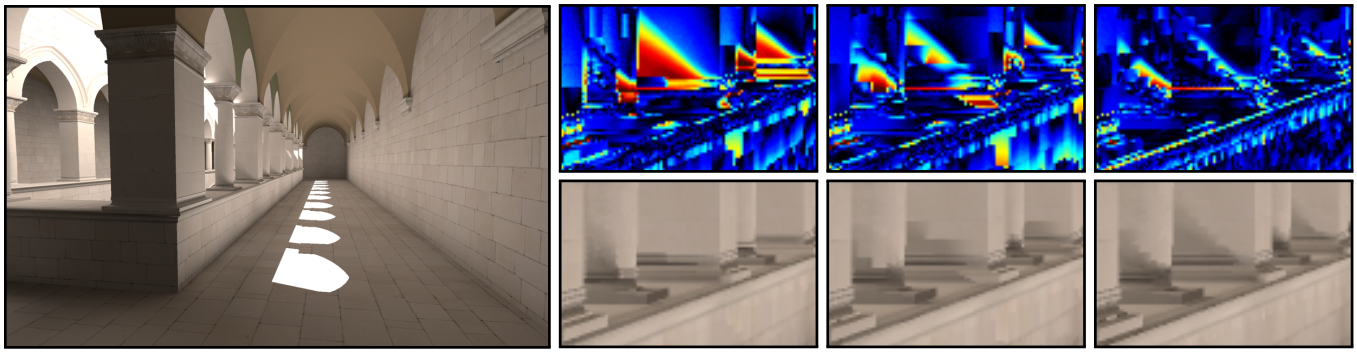

Reference | Pure Split-Sphere | Bounded Split-Sphere | Occlusion Hessian

Figure 15: *In Sponza the Occlusion Hessian is superior to the Split-Sphere in both defining shadow details and eliminating interpolation artifacts across surfaces. The shadows at the base on the column are captured by the Occlusion Hessian, while the Split-Sphere heuristic shows severe interpolation artifacts, even when bounded to 20px and clamped by the gradient.*

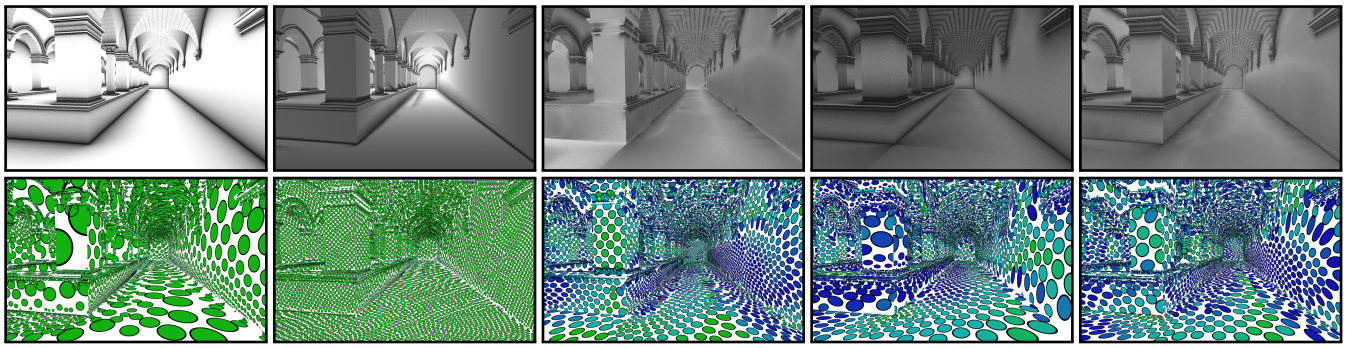

Pure Split-Sphere | Bounded Split-Sphere | Radiometric Hessian | Geometric Hessian | Occlusion Hessian

Figure 16: *Direct visualization of the radii (top) and footprints (bottom) of the cache records for various error metrics. The Pure Split-Sphere, Radiometric and Geometric Hessians are all unable to detect the indirect shadows on the columns, while the Bounded Split-Sphere (clamped to 20px and the gradient) as well as the Occlusion Hessian correctly predict smaller radii in these rapidly changing regions. Unfortunately, the contrast in the Split-Sphere is too high, leading to lost detail due to the greedy nature of the irradiance caching algorithm.*

lation Hessian - while taking object inter-occlusions into account. A second-order Taylor expansion of irradiance is used as an oracle to the true irradiance, in order to estimate the error induced by first-order Taylor expansion during rendering. The Hessian also naturally supports anisotropic cache points. As a result, record density closely matches the rate of change of irradiance in the scene, leading to reduced error in the rendered images when compared to all previous approaches. Furthermore, our method does not rely heavily on user-defined minimum and maximum cache point radii, leading to more intuitive user control. The robustness of our new Occlusion Hessian method also minimizes the need to use other corrective techniques, such as progressive refinement, neighbor clamping, or an overture pass, resulting in a simpler implementation.

Reference | Bounded Split-Sphere | Occlusion Hessian

Figure 17: *An interior of San Miguel with strong indirect illumination. The close-ups show how the Occlusion Hessian is able to resolve the shadows behind the paintings, while the Split-Sphere heuristic (clamped to 20px and the gradient) largely misses these shadows.*

References

ARVO, J., TORRANCE, K., AND SMITS, B. 1994. A framework for the analysis of error in global illumination algorithms. In *Proceedings of SIGGRAPH 94*, Computer Graphics Proceedings, Annual Conference Series, 75–84.

ARVO, J. 1994. The irradiance jacobian for partially occluded polyhedral sources. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, ACM, SIGGRAPH '94, 343–350.

HERZOG, R., MYSZKOWSKI, K., AND SEIDEL, H.-P. 2009. Anisotropic radiance-cache splatting for efficiently computing high-quality global illumination with lightcuts. In *Computer Graphics Forum (Proc. Eurographics)*, Wiley-Blackwell, vol. 28(2), 259–268.

HOLZSCHUCH, N., AND SILLION, F. 1995. Accurate compu-

tation of the radiosity gradient for constant and linear emitters. In *Rendering Techniques 1995 (Proceedings of the Eurographics Workshop on Rendering)*, Springer-Verlag, 186–195.

HOLZSCHUCH, N., AND SILLION, F. X. 1998. An exhaustive error-bounding algorithm for hierarchical radiosity. *Computer Graphics Forum 17*, 4, 197–218.

HOLZSCHUCH, N. 1996. *Le Contrôle de l'Erreur dans la Méthode de Radiosité Hiérarchique*. PhD thesis, Université Joseph Fourier (Grenoble I).

JAROSZ, W., SCHÖNEFELD, V., KOBBELT, L., AND JENSEN, H. W. 2012. Theory, analysis and applications of 2D global illumination. *ACM Transactions on Graphics (Presented at ACM SIGGRAPH 2012) 31*. to appear.

KŘIVÁNEK, J., AND GAUTRON, P. 2009. *Practical Global Illumination with Irradiance Caching*. Synthesis lectures in computer graphics and animation. Morgan & Claypool.

KŘIVÁNEK, J., GAUTRON, P., PATTANAIK, S., AND BOUATOUCH, K. 2005. Radiance caching for efficient global illumination computation. *IEEE Transactions on Visualization and Computer Graphics 11*, 5 (Sept./Oct.), 550–561.

KŘIVÁNEK, J., BOUATOUCH, K., PATTANAIK, S., AND ZÁRA, J. 2006. Making radiance and irradiance caching practical: Adaptive caching and neighbor clamping. In *Rendering Techniques 2006 (Proceedings of the Eurographics Symposium on Rendering)*, 127–138.

KŘIVÁNEK, J., GAUTRON, P., BOUATOUCH, K., AND PATTANAIK, S. 2005. Improved radiance gradient computation. In *SCCG '05: Proceedings of the 21th spring conference on Computer graphics*, ACM Press, 155–159.

LARSON, G. W., AND SHAKESPEARE, R. 1998. *Rendering with Radiance: The Art and Science of Lighting Visualization*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.

PHARR, M., AND HUMPHREYS, G. 2010. *Physically Based Rendering: From Theory to Implementation*, 2nd ed. Morgan Kaufmann.

RAMAMOORTHI, R., MAHAJAN, D., AND BELHUMEUR, P. 2007. A first-order analysis of lighting, shading, and shadows. *ACM Transactions on Graphics 26*, 1 (Jan.), 2:1–2:21.

SHIRLEY, P., AND CHIU, K. 1997. A low distortion map between disk and square. *Journal of Graphics, GPU, and Game tools 2*, 3, 45–52.

SHIRLEY, P. S. 1991. Time complexity of Monte Carlo radiosity. In *Eurographics '91*, 459–465.

TABELLION, E., AND LAMORLETTE, A. 2004. An approximate global illumination system for computer generated films. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2004) 23*, 3 (Aug.), 469–476.

WARD, G. J., AND HECKBERT, P. S., 1992. Irradiance gradients.

WARD, G. J., RUBINSTEIN, F. M., AND CLEAR, R. D. 1988. A ray tracing solution for diffuse interreflection. *SIGGRAPH Comput. Graph. 22* (June), 85–92.

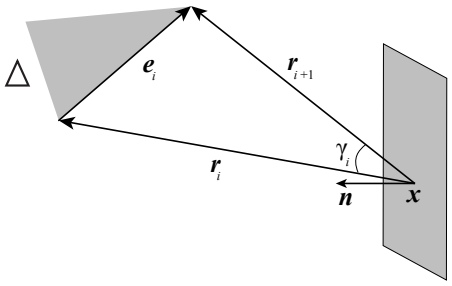

Figure 18: *Illustration of the notation used in the form-factor, its gradient and Hessian.*

A Form Factor, Gradient, and Hessian

The formula for the point-to-triangle form-factor $F_\triangle(\mathbf{x})$ is given by [Holzschuch and Sillion 1998]:

$$F_\triangle(\mathbf{x}) = \frac{1}{2\pi}\mathbf{n} \cdot \sum_i^3 \Gamma_i \tag{19}$$

where i indexes the three vertices, $\mathbf{r}_i$, expressed as vectors from $\mathbf{x}$, and $\Gamma_i = \mathbf{r}_i \times \mathbf{r}_{i+1}$ and has norm γ_i (see Figure 18).

Holzschuch and Sillion derived the following formulas for the gradient and Hessian of the form-factor:

$$\begin{aligned}\nabla_{\mathbf{x}} F_\triangle(\mathbf{x}) = &-\frac{-1}{2\pi}\sum_i \mathbf{n} \times \mathbf{e}_i I_1 \\ &+ 2\mathbf{n} \cdot (\mathbf{r}_i \times \mathbf{r}_{i+1})(\mathbf{r}_i I_2 + \mathbf{e}_i J_2)\end{aligned} \tag{20}$$

and

$$\begin{aligned}\mathbf{H}_{\mathbf{x}} F_\triangle(\mathbf{x}) = &-\frac{1}{\pi}\sum_i Q(\mathbf{n} \times \mathbf{e}_i, \mathbf{r}_i I_2 + \mathbf{e}_i J_2) \\ &- \mathbf{n} \cdot (\mathbf{r}_i \times \mathbf{e}_i) I_2 \mathbf{I}_{3\times 3} \\ &+ 2\mathbf{n} \cdot (\mathbf{r}_i \times \mathbf{e}_i)(Q(\mathbf{r}_i, \mathbf{r}_i) I_3 \\ &+ Q(\mathbf{e}_i, \mathbf{e}_i) K_3 + 2 J_3 Q(\mathbf{r}_i, \mathbf{e}_i)),\end{aligned} \tag{21}$$

where:

$$\begin{aligned}I_1 &= \frac{\gamma_i}{\|\mathbf{e}_i \times \mathbf{r}_i\|}, \\ I_2 &= \frac{1}{2\|\mathbf{e}_i \times \mathbf{r}_i\|^2}\left(\frac{\mathbf{e}_i \cdot \mathbf{r}_{i+1}}{\|\mathbf{r}_{i+1}\|^2} - \frac{\mathbf{e}_i \cdot \mathbf{r}_i}{\|\mathbf{r}_i\|^2} + \|\mathbf{e}_i\|^2 I_1\right), \\ I_3 &= \frac{1}{4}\frac{1}{\|\mathbf{e}_i \times \mathbf{r}_i\|^2}\left(\frac{\mathbf{e}_i \cdot \mathbf{r}_{i+1}}{\|\mathbf{r}_{i+1}\|^4} - \frac{\mathbf{e}_i \cdot \mathbf{r}_i}{\|\mathbf{r}_i\|^4} + 3\|\mathbf{e}_i\|^2 I_2\right), \\ J_2 &= \frac{1}{2\|\mathbf{e}_i\|^2}\left(\frac{1}{\|\mathbf{r}_i\|^2} - \frac{1}{\|\mathbf{r}_{i+1}\|^2}\right) - \frac{\mathbf{e}_i \cdot \mathbf{r}_i}{\|\mathbf{e}_i\|^2} I_2, \\ J_3 &= \frac{1}{4\|\mathbf{e}_i\|^2}\left(\frac{1}{\|\mathbf{r}_i\|^4} - \frac{1}{\|\mathbf{r}_{i+1}\|^4}\right) - \frac{\mathbf{e}_i \cdot \mathbf{r}_i}{\|\mathbf{e}_i\|^2} I_3, \\ K_3 &= \frac{1}{\|\mathbf{e}_i\|^2}\left(I_2 - \|\mathbf{r}_i\|^2 I_3 - 2(\mathbf{r}_i \cdot \mathbf{e}_i) J_3\right), \text{ and} \\ Q(\mathbf{a}, \mathbf{b}) &= \mathbf{a}\mathbf{b}^\intercal + \mathbf{b}\mathbf{a}^\intercal.\end{aligned}$$

Note that $\mathbf{a}$ and $\mathbf{b}$ are column vectors, such that $Q(\mathbf{a}, \mathbf{b})$ results in a 3x3 matrix.

Arvo [1994] presented a formula for the irradiance Jacobian due to a diffusely emitting triangle, which results in an alternate, but equivalent, expression for the form-factor gradient.