

# Is Accurate Occlusion of Glossy Reflections Necessary?

Oscar Kozlowski\*  
University College London

Jan Kautz†  
University College London

## Abstract

Much research in recent times has been conducted towards real-time rendering of accurate glossy reflections under direct, natural illumination including correct occlusions. The view dependent nature of these reflections will always cause this computation to be expensive unless heavily approximated. There also remains a question as to whether humans are even capable of noticing the difference in accuracy or whether our perception of the realism of the scene remains unchanged and thus the extra effort expended in rendering accurate reflections is effectively wasted. We conduct a user study to analyse any decline in perceived realism of glossy scenes rendered with a variety of specular occlusion approximations under a multitude of BRDFs, lighting environments and camera orientations. We demonstrate that although no one approximation is always suitable, it is rare to have a scene whose computational complexity cannot be decreased to some degree.

**CR Categories:** I.3.3 [Computer Graphics]: Picture/Image Generation—Bitmap and frame buffer operations; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Color, Shading, Shadowing and Texture

**Keywords:** glossy reflections, shadowing, perception

## 1 Introduction

The human perception of visual stimuli has been studied to some degree for centuries. Early art focused on recognisable shapes representing the subject matter. Renaissance artists were among the first to formalise the rules of realistic shading of diffuse objects as well as more general visual cues such as perspective. In recent times, increasing amounts of effort have been made to try to understand exactly how the visual system works and how information from complicated lighting effects such as shadowing, texturing and specularities is perceived and interpreted [Todd et al. 1997; Fleming et al. 2004]. The invention of the computer has brought the same millenia-old desire of conveying realism through images into a digital age, the ultimate goal being the indistinguishable replication of reality.

Even with exponentially increasing computational power, our drive for constantly more complicated and detailed scenes means that we are not much closer to achieving our dream goal. Indeed, it is likely that we will never be able to fully simulate reality, so instead, research has turned back on focusing on determining the extents of our visual perception system, and if it would be possible to fool it sufficiently without replicating reality in full.

Much research has been undertaken to obtain knowledge about the visual system and its perception of materials [Hunter and Harold

1987; Adelson and Pentland 1996; Fleming et al. 2001; Fleming et al. 2003] and shapes [Blake and Bülthoff 1990; Todd et al. 1997; Fleming et al. 2004]. In turn this knowledge can be used to reduce computation, and has been exploited in computer graphics for efficient image synthesis [Ramasubramanian et al. 1999; Prikryl and Purgathofer 1999; Volevich et al. 2000; Myszkowski et al. 2001; Stokes et al. 2004].

Research on specular highlights and reflections [Fleming et al. 2004] has demonstrated that even in the absence of all other visual cues, a strong sense of shape is achievable from solely the perfect reflection of the environment. They argue that the visual system treats the reflection as a texture that has been warped by projecting it onto the second derivative of the the surface of the object, and is able to utilise this information to determine the likely 3D shape of the mesh. This result implies that inaccurate specular reflections can possibly contradict the other visual cues and undermine the perceived realism of a scene.

Due to their view-dependent nature, specular and glossy reflections from direct, natural illumination (e.g., from a high-dynamic range environment map) are computationally extremely expensive. Fully accurate glossy occlusion can take hours to compute each frame and must be significantly approximated using a variety of techniques to achieve interactive or real-time frame rates. A trade-off must therefore be made between a sufficient level of approximation for reducing the computational complexity and ensuring that our perception of realism does not needlessly suffer. This balancing act also varies between Bi-directional Reflectance Distribution Functions (BRDFs) of varying degrees of specularity, different lighting environments and even the orientation of the scene with respect to the camera, making the balance very difficult to achieve. This balance has never been studied so our contribution is the analysis of the decline in perceived realism when rendered with a variety of approximating techniques. Our experiment is conducted on images rendered under direct illumination only. This implies that reflections only consist of the lighting environment; objects occlude reflections of the lighting environment only. This is a valid approach because all of the approximating techniques that are commonly used (and are examined here) are geared toward the real-time problem domain where inter-reflections are not considered.

## 2 Related Work

**Precomputed Radiance Transfer.** Recently, Precomputed Radiance Transfer (PRT) [Sloan et al. 2002] has created significant interest in rendering scenes under natural illumination at interactive or real-time rates. The original form of PRT allows us to render diffuse and glossy objects with direct and indirect illumination. When using spherical harmonics as basis functions, the natural illumination is limited to low frequencies [Sloan et al. 2002; Kautz et al. 2002; Sloan et al. 2005]. This can be alleviated using Wavelets [Ng et al. 2003; Ng et al. 2004] enabling higher-frequency reflections.

Interestingly, almost all improved PRT techniques that incorporate high-frequency reflections or allow dynamic scenes of rigid objects are limited to *direct* illumination only [Wang et al. 2004; Ng et al. 2004; Kautz et al. 2004; Kontkanen and Laine 2005; Zhou et al. 2005; Green et al. 2006; Sun and Mukherjee 2006; Tsai and Shih 2006]. This limitation stems from various technical reasons, e.g., prohibitively long precomputation times [Green et al. 2006; Tsai and Shih 2006], expensive run-time [Zhou et al. 2005; Green et al.

\*e-mail: o.kozlowski@cs.ucl.ac.uk

†e-mail: j.kautz@cs.ucl.ac.uk

2006], or techniques being designed only for direct lighting [Wang et al. 2004; Ng et al. 2004; Kautz et al. 2004; Zhou et al. 2005; Sun and Mukherjee 2006]. For diffuse materials, shadowing is very obvious, and enabling accurate high-frequency direct illumination is visually very important. However, until now it has been unclear how accurate self-occlusion needs to be when rendering glossy reflections. If accurate self-occlusion for glossy reflections is visually not that important, then this can potentially save much computation, especially in real-time applications.

**Efficient Global Illumination using Perception.** A few techniques have been proposed to make global illumination more efficient based on human perception [Ramasubramanian et al. 1999; Prikryl and Purgathofer 1999; Volevich et al. 2000; Myszowski et al. 2001; Stokes et al. 2004]. The paper by Stokes et al. [2004] is most related to our research. They propose to render direct, indirect diffuse, indirect glossy, and indirect specular as separate components, which are then summed up to yield a final image. Depending on the scene and view-point not all components are needed. In contrast, we wish to study the visual realism of the direct glossy component, depending on its accuracy.

**Perception of Shape and Materials.** Much work has been done on the perception of shape and materials. Glossy and specular reflections give important visual cues on the shape of an object [Blake and Bülthoff 1990; Todd et al. 1997; Fleming et al. 2004], which humans are able to exploit. The perception of materials itself has also gathered significant attention [Hunter and Harold 1987; Adelson and Pentland 1996; Fleming et al. 2001; Fleming et al. 2003]. However, to our best knowledge the perception of occluded glossy reflections has not been studied yet and is addressed in this paper.

### 3 Rendering Background

Let us first introduce the rendering equation [Immel et al. 1986], which is used to compute the amount of reflected radiance  $L_{out}$  at a point  $\mathbf{p}$  on the surface along the viewing direction  $\mathbf{v}$ :

$$L_{out}(\mathbf{p}, \mathbf{v}) = \int L_{in}(\mathbf{p}, s) f_r(\mathbf{p}, s, \mathbf{v}) (s \cdot \mathbf{n}) \partial s, \quad (1)$$

where  $\mathbf{n}$  is the surface normal,  $L_{in}$  is the incoming radiance,  $f_r$ <sup>1</sup> is the BRDF. This equation accounts for direct as well as indirect illumination, as the incoming radiance  $L_{in}$  includes indirect radiance  $L_{indirect}$  reflected from other parts of the scene, as well as radiance  $L_s$  coming from light sources:

$$L_{in}(\mathbf{p}, s) = (1 - V(\mathbf{p}, s)) L_{indirect}(\mathbf{p}, s) + V(\mathbf{p}, s) L_s(\mathbf{p}, s) \quad (2)$$

where  $V$  is a binary visibility classifier that indicates whether a light source or another point in the scene is visible. Note that the indirect radiance  $L_{indirect}$  can be written as  $L_{indirect}(p, l) = L_{out}(p', -l)$  with  $p'$  being the first visible point from  $\mathbf{p}$  along  $s$ .

As mentioned in the previous section, most all-frequency PRT techniques only support direct lighting. Furthermore, they make the assumption that the light source  $L_s$  is given as a high-dynamic range environment map, representing distant natural illumination:

$$L_{out}(\mathbf{p}, \mathbf{v}) = \int L_s(s) f_r(\mathbf{p}, s, \mathbf{v}) V(\mathbf{p}, s) (s \cdot \mathbf{n}) \partial s. \quad (3)$$

Note that we have dropped  $L_s$ 's dependency on  $\mathbf{p}$ , and that we have removed the indirect lighting contribution.

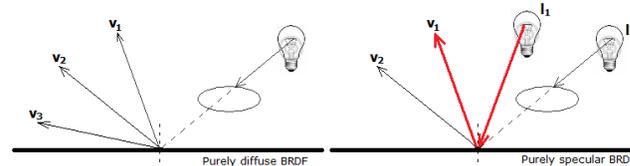
<sup>1</sup>Note that the BRDF  $f_r$  can usually only be defined in the local tangent frame at  $\mathbf{p}$ . Throughout the paper we assume an implicit conversion to that local coordinate system.

This integral itself and especially the visibility term  $V$  within is computationally expensive, so numerous techniques have been developed to approximate the result by adding some constraining assumptions. E.g., environment mapping techniques simplify this even further by dropping the visibility term [Greene 1986], whereas PRT techniques try to accurately evaluate the equation but commonly only allow the illumination to vary (static scene or at least rigid objects).

In the following sections, we examine renderings done using Equation 3 (i.e., direct lighting under natural illumination) with varying approximations of the visibility and draw conclusions as to which ones are perceptually valid and can also result in computational gains.

#### 3.1 Direct vs. Specular Occlusion

Objects illuminated within a scene effectively partition three dimensional space into regions containing varying amounts of scattered radiance. Illumination on the surfaces of objects can be thought of as the projection of this radiance along their direction of travel. What are commonly referred to as shadows thus arise from a lack of scattered radiance travelling along the direction of projection. It is worth noting that the dependency of the direction of this virtual projection differs between directly cast and specular occlusions.



**Figure 1:** Left: A purely diffuse BRDF scatters radiance independently of viewing direction. All possible viewing directions see a shadow caused by the occlusion of the light. Right: A purely specular BRDF reflects radiance only in the perfect reflection angle. Therefore viewing direction  $\mathbf{v}_1$  see a reflection of light  $l_1$  but view  $\mathbf{v}_2$  sees a shadow (light  $l_2$  is occluded along the perfect reflectance direction).

As depicted in Figure 1 left, a purely diffuse BRDF  $f_r = k_d/\pi$  results purely in the ubiquitous shadows we observe in our daily lives. These shadows arise from *direct occlusions*; light arriving from the light source is directly blocked, casting a shadow which is independent of the viewing direction. The projection effectively occurs along the direction  $s$ .

Now imagine a scene with only a perfectly specular BRDF (no diffuse contribution) as in Figure 1 right. In this case, all radiance reaching the viewer has not interacted with the scene at all or undergone perfect reflections only. A reflection of the light source at point  $\mathbf{p}$  is only visible if the viewing direction  $\mathbf{v}$  is a perfect reflection of the incident light direction  $l$  around the normal at that point. Conversely *specular occlusions* are an absence of this reflected radiance. Specular occlusions behave very differently due to their dependency on the viewing direction.

Glossy (i.e. not perfectly specular) BRDFs and dull glossy<sup>2</sup> BRDFs consist of diffuse and specular components which scatter light independent and dependent on the viewing direction respectively. Thus scenes with these BRDFs will exhibit both directly cast occlusions and specular occlusions.

Accurate direct occlusions, i.e. shadows, from natural illumination are handled quite well and efficiently by many real-time shading

<sup>2</sup>We define the term 'dull glossy' as a glossy BRDF with a broad lobe such that reflections appear blurred.

methods [Ng et al. 2003; Sloan et al. 2003; Ng et al. 2004; Sun and Mukherjee 2006; Tsai and Shih 2006]. Correct specular occlusion can be achieved interactively as well, however, they are much more expensive to evaluate (due to view-dependency) and require more memory [Wang et al. 2004; Ng et al. 2004; Green et al. 2006; Sun and Mukherjee 2006; Tsai and Shih 2006].

## 4 Shadowing of Glossy Reflections

We have conducted an experiment using 5 differing techniques, which approximate the specular visibility term to analyse if fully accurate occlusion of specular highlights and reflections is perceptually really necessary. For all images, the diffuse contribution and its occlusion has not been approximated. Although each of the techniques could be extended to account for inter-reflections, they are all geared towards deployment in real-time rendering situation (see above), and so, we have not rendered them with inter-reflections to more accurately portray the potential of the approximations in context.

Note that in the following derivations, we only look at the glossy component of the BRDF. The diffuse shading is always computed accurately with path tracing, evaluating the following integral:

$$L_{out}^d(\mathbf{p}, \mathbf{v}) = \int L_{in}(\mathbf{p}, s) f_r^d(\mathbf{p}, s, \mathbf{v})(s \cdot \mathbf{n}) \partial s, \quad (4)$$

with  $f_r^d$  being the diffuse component of the BRDF.  $L_{out}^d$  is then added to the glossy result, which in term is computed from the glossy component of the BRDF  $f_r^g$ ; see the following.

### 4.1 Discretisation of Visibility

The standard approach to speeding up the reflection equation is by assuming various constraints, which allow us to re-arrange the equation such that a portion of it can be pre-computed offline. The simplest of these techniques is the discretisation of the specular visibility across the continuous hemisphere into a discrete hemicube [Cohen and Wallace 1993]. The reflection equation now changes to Equation 5 where  $V_d$  is the discretised visibility and  $d$  is the ray direction  $s$  snapped to the nearest texel of the discretised hemicube.

$$L_{out}^g(\mathbf{p}, \mathbf{v}) = \int L_s(s) f_r^g(\mathbf{p}, s, \mathbf{v}) V_d(\mathbf{p}, d)(s \cdot \mathbf{n}) \partial s. \quad (5)$$

The discretisation of the visibility is a common starting point for accelerating path tracing and some real-time applications as the visibility can be pre-calculated at select points across a mesh and interpolated at run-time. Low resolution hemicubes suffer from artifacts at occlusion boundaries whereas large resolution hemicubes become prohibitively storage cost expensive as the mesh complexity increases. A trade-off must be made and we have chosen to use a 32x32 cubemap to discretise and store the visibility. Discretisation of visibility has been used before [Heidrich et al. 2000], and is implicitly done in Wavelet PRT as well, where 64x64 cubemaps are commonly used [Ng et al. 2003; Ng et al. 2004; Sun and Mukherjee 2006].

### 4.2 Visibility using Spherical Harmonics

Spherical harmonics define an orthonormal basis over a sphere, analogous to the Fourier Transform over the 1 dimensions circle. Signals projected into a spherical harmonics basis have the advantage of a compact representation, are rotationally invariant and low frequencies can be accurately represented using only a few bands (we utilise only 100 coefficients). Higher frequencies are unfortunately band-limited (i.e. smoothed without aliasing) with a low-order projection. Spherical harmonics have been commonly utilised since their use in real-time computer graphics by Sloan et al. [2002]

in their PRT technique especially for low-frequency diffuse radiance contributions. In their work, real-time applications project the lighting environment into a spherical harmonics basis and thus reduce the integral into a simple dot product. We make a slightly different approximation. We only project the visibility into spherical harmonics, but keep the full high-frequency incident lighting, since we want to study the effect of visibility on glossy reflections and not the effect of low-frequency lighting. Thus, we query the spherically projected visibility  $V_{SH}$  for a given direction and use the result (clamped between 0 and 1) as an approximation for the visibility:

$$L_{out}^g(\mathbf{p}, \mathbf{v}) = \int L_s(s) f_r^g(\mathbf{p}, s, \mathbf{v}) V_{SH}(\mathbf{p}, s)(s \cdot \mathbf{n}) \partial s. \quad (6)$$

### 4.3 Directional Ambient Occlusion using Spherical Harmonics

Glossy BRDFs' specular lobes are usually centered on the perfect reflection direction, and thus, the perfectly reflected ray carries the most potential radiance. Projecting the visibility into a spherical harmonics basis (especially a low-order basis) has the effect of smoothing the previously binary visibility across the sphere. Every direction within the spherical harmonics distribution now approximates not only its own visibility, but also the visibilities of angularly nearby rays. Decreasing the number of spherical harmonics coefficients used, increases the smoothing across the sphere, and thus, each ray approximates the visibility of an increasing angle of rays (to a decreasing level of precision). Querying the spherical harmonics distribution for the perfect reflection direction  $r$ , thus results in an approximation for the visibility of an angular segment of the specular lobe

$$L_{out}^g(\mathbf{p}, \mathbf{v}) = V_{SH}(\mathbf{p}, r) \int L_s(s) f_r^g(\mathbf{p}, s, \mathbf{v})(s \cdot \mathbf{n}) \partial s \quad (7)$$

This approximation considerably reduces the computational complexity (the visibility for potentially hundreds of specular rays is approximated by only one) and is extremely suitable for real-time applications. The danger is that the visibility of the lobe of some BRDFs will not be fully captured or will be smoothed too much resulting in artifacts. We are experimenting on a range of glossy BRDFs utilising 100 spherical harmonic coefficients for the visibility. This approximation was used by [Green et al. 2007] but was not compared to ground truth images, making it difficult to judge the approximation quality. It can also easily be combined with other filtered environment map rendering techniques [Ramamoorthi and Hanrahan 2002; Kautz and McCool 2000; Kautz et al. 2000].

### 4.4 Ambient Occlusion

By disregarding all directional information about the visibility across the hemisphere, visibility is simplified back to a constant. This constant is known as ambient occlusion [Zhukov et al. 1998] and is defined in its simplest terms as the ratio between the number of unoccluded rays from a point on the surface to the lighting environment and the total number of queried rays. This simplifies the rendering equation to:

$$L_{out}^g(\mathbf{p}, \mathbf{v}) = \frac{\int_{\Omega^+} V(\mathbf{p}, s) \partial s}{2\pi} \cdot \int L_s(s) f_r^g(\mathbf{p}, s, \mathbf{v})(s \cdot \mathbf{n}) \partial s, \quad (8)$$

where the first term is the ambient occlusion term, which can be computed independently of the lighting. This model is commonly used [Méndez et al. 2003; Kontkanen and Laine 2005; Bunnell 2005] in the case of self-occlusion on diffuse BRDFs. In this work, we evaluate its use applied to self-occlusion on glossy BRDFs.

## 4.5 No Occlusion

The final approximating approach we utilise is to completely disregard visibility and assume that all specular rays are always visible:

$$L_{out}^g(\mathbf{p}, \mathbf{v}) = \int L_s(s) f_r^g(\mathbf{p}, s, \mathbf{v})(s \cdot \mathbf{n}) \partial s. \quad (9)$$

This corresponds to traditional filtered environment mapping [Ramamoorthi and Hanrahan 2002; Kautz and McCool 2000; Kautz et al. 2000].

## 4.6 Ground Truth

For the purposes of a thorough analysis, we produced ground truth images containing fully accurate specular self-occlusions and inter-reflections by numerically integrating the full rendering equation (Eq. 1) using a Monte Carlo path tracer.

## 5 Experimental Setup

The experiment was conducted on 24 people of mixed age and split approximately equally between backgrounds in computer science and not. The only requirement imposed on test subjects was their having perfect or corrected to perfect vision.

There were two sections in the experiment, the first consisted of a random order of 60 images (2 repeats of 5 scenes of 6 images) being presented one after the other to the participant. The order was constrained such that no two images of the same scene could appear consecutively and a blank image was displayed for 0.5 seconds between each picture. These measures minimised the probability of opportunity to directly compare images of a particular scene. The participants were asked to rate each picture from 1–5 (low to high) on two scales, the first being defined as realism (in terms of lighting, shadowing and reflections) and the second scale being defined as how pleasing the overall image was to the participant.

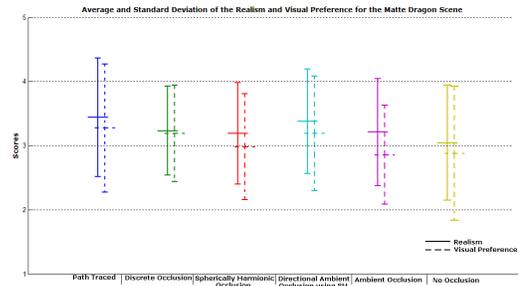
The second section of the experiment allowed the user to directly compare all 6 images from each scene and was then asked to rank them from most to least realistic image. Both tests are insightful as although most computer graphic applications would not provide the user with a reference image for direct comparison, there would be some circumstances where the user may have prior knowledge (potentially learned elsewhere) of the visual stimuli.

## 6 Results

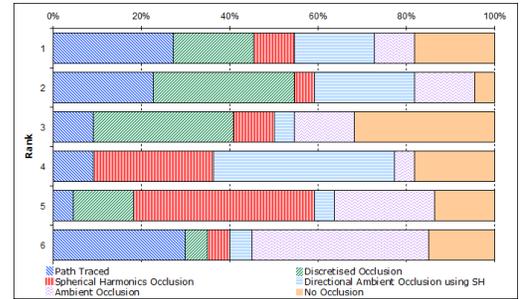
Each participant has provided two 'Realism' and two 'Visual Preference' scores for each image (two repeats); these repeats are averaged together to form the score that the participant has given that image. These results from each person are averaged together to calculate the average scores and standard deviation of the entire test group for each image, which are then plotted using box plots below. Additionally the rankings for the images within each scene are tabulated and also presented below in percent bar charts; these charts show the frequency each image was graded in a particular rank.

### 6.1 Dull Glossy Dragon Scene

The dull glossy dragon illuminated in Grace Cathedral (Figure 13) demonstrates that even with a myriad of shadows cast by a multitude of light sources, the differences in the resultant images are too subtle to cause significant variation in the opinion of their realism (Figure 2). Looking more closely at the images, only the path traced and discretised occlusion images replicated the large shadow directly below the dragon with any detail, but of these, only the path



**Figure 2:** A box plot plotting the average and standard deviation of the 'Realism' and 'Visual Preference' scales for the dull glossy dragon illuminated in Grace Cathedral scene.



**Figure 3:** A percent bar chart plotting the rankings for the dull glossy dragon illuminated in Grace Cathedral.

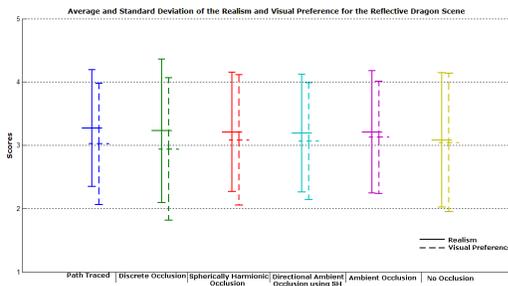
traced image was well received. This is likely to be due to the extra interreflections adding a fine, but noticeable layer of additional detail over this area causing the image to appear less computer-generated. Both the spherical harmonics occlusion and ambient occlusion techniques produced significantly smoothed shadow boundaries along the dragon's torso, resulting in an increased diffuse-like appearance reminiscent of early computer generated imagery. In contrast, the directional ambient occlusion using spherical harmonics technique yielded an appearance evocative of sharpened shadow boundaries on dull glossy BRDF surfaces such as the dragon's torso. This did however retain a much larger proportion of the specular highlights, resulting in an image that was regarded as highly as the ground truth path traced image.

Given the opportunity to directly compare the images, the rankings in Figure 3 depict that with the exception of ambient occlusion and spherical harmonics occlusion, each of the images are similar in realistic plausibility. Somewhat surprisingly, the ground truth path traced image ranks largely at the bottom as well as the top of the rankings. This may be possibly accounted for by the amount of residual noise in the path-traced images; some people may have thought that this noise added to the realism and some that it detracted from it.

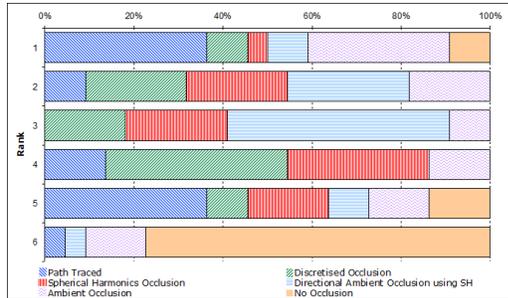
### 6.2 Highly Reflective Dragon Scene

The high complexity of the dragon mesh (and thus an unstable second derivative of the surface) causes the specular reflections to be warped such that our visual system is probably unable to determine if there are any visual contradictions (Figure 12). Thus, all the images score almost identically in Figure 4.

However, Figure 5 shows how direct comparison of the same scene causes the ranking to become more polarised. No occlusion is universally regarded as the least realistic image due to its obvious over-bright highlights in comparison to the other images. Ambient occlusion is rated as one of the most realistic, but this is probably due to the fact that the bright incoherent highlights are made more pleas-



**Figure 4:** A box plot plotting the average and standard deviation of the 'Realism' and 'Visual Preference' scales for the highly reflective dragon illuminated in Grace Cathedral scene.



**Figure 5:** A percent bar chart plotting the rankings for the highly reflective dragon illuminated in Grace Cathedral.

ing to the eye by being effectively darkened. Directional ambient occlusion using spherical harmonics produces an identical image to spherical harmonics occlusion with the exception of a very soft shadow directly underneath the dragon. This shadow is actually an artifact as although the perfect reflection direction may be occluded at that point on the surface of the groundplane, the vast majority of the remainder of the hemisphere is not occluded. However, this shadow appears to be visually pleasing and has thus, caused this image to be ranked higher in Figure 5 (but not statistically significant).

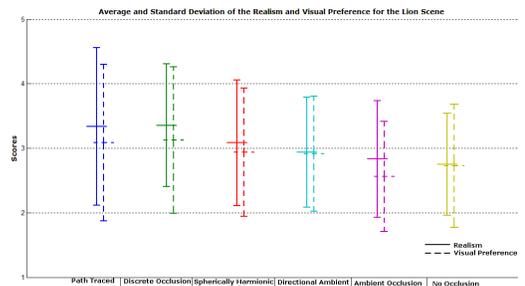
### 6.3 Dull Glossy Chinese Lion Scene

This scene (Figure 14) contains multiple small lightsources which are reflected in the groundplane and the various approximating techniques have difficulty in shading some of these sufficiently. This results in Figure 6 demonstrating a visually clear (but statistically insignificant) preference for more accurate occlusion classification. Both the no occlusion and ambient occlusion images contain an obviously incorrect bright highlight on the groundplane next to the lion's left front paw which in all likelihood caused the test participants to immediately question the realism of the scene. The spherical harmonics occlusion and directional ambient occlusion using spherical harmonics fail to reproduce in sufficient detail the large reflection of the lion in the groundplane, and are rated down accordingly.

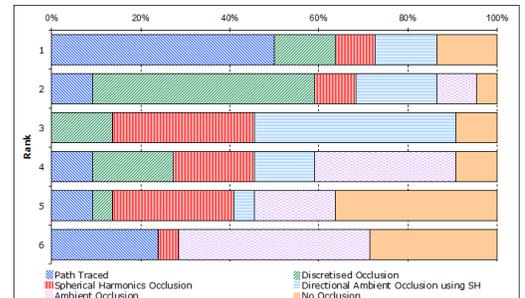
Figure 7 supports the clear tendency to prefer the more accurate occlusion techniques. The reflections of the small bright light sources provide such a strong visual cue to the orientation of the scene and specularity of the groundplane, that the lack of clarity in the reflection of the lion is made all the more apparent.

### 6.4 Dull Glossy Tweety Bird Scene

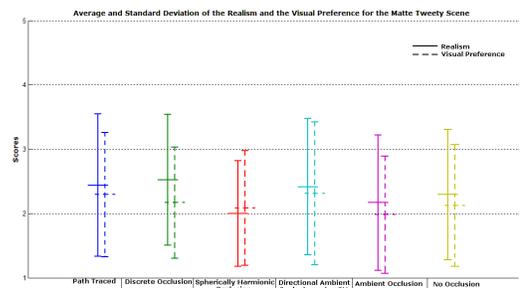
In direct contrast to the last scene, the Eucalyptus Grove lighting environment in this scene consists of only 1 very large light-source (Figure 12) and thus, a mapping of exactly which texel of



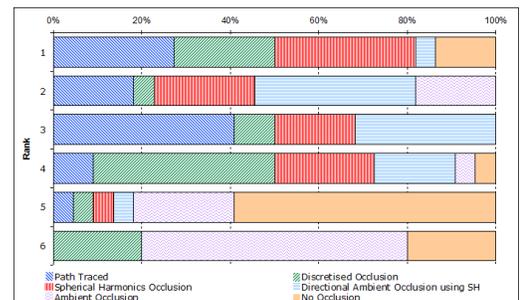
**Figure 6:** A box plot plotting the average and standard deviation of the 'Realism' and 'Visual Preference' scales for the dull glossy chinese lion illuminated in St. Peter's Basilica scene.



**Figure 7:** A percent bar chart plotting the rankings for the dull glossy chinese dragon illuminated in St. Peter's Basilica.



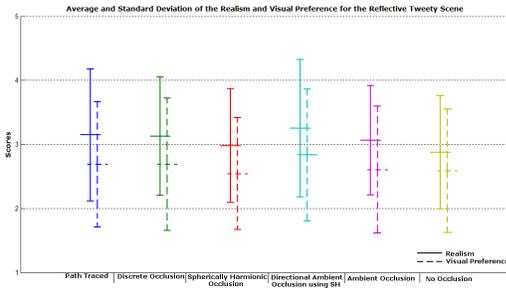
**Figure 8:** A box plot plotting the average and standard deviation of the 'Realism' and 'Visual Preference' scales for the dull glossy tweety bird illuminated in the Eucalyptus Grove.



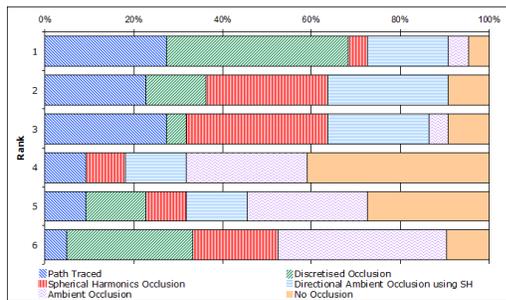
**Figure 9:** A percent bar chart plotting the rankings for the dull glossy tweety bird illuminated in the Eucalyptus Grove.

the environment map is being reflected at a particular point on the groundplane is difficult to establish. Even under highly approximated occlusion, the cues generated are too weak to contradict the other cues, and so, our visual perception system deems them all as equally plausible (Figure 8). Once again however, given an opportunity for direct comparison, a clear ranking is established with a preference towards higher accuracy occlusion classifiers.

## 6.5 Highly Reflective Tweety Bird Scene



**Figure 10:** A box plot plotting the average and standard deviation of the ‘Realism’ and ‘Visual Preference’ scales for the highly reflective tweety bird illuminated in the Eucalyptus Grove.



**Figure 11:** A percent bar chart plotting the rankings for the highly reflective tweety bird illuminated in the Eucalyptus Grove.

This scene consists of an extremely dull glossy groundplane with a reflective tweety bird mesh resting on top (Figure 12); very little difference can be seen between the images and this is reflected in the scores in Figure 10. Figure 11 supports this as even with direct comparison, there is little significant preference for any image. Only ambient occlusion and no occlusion suffer in the rankings and this is because these images stand out against the others in terms of brightness (the images are darker and brighter than the others respectively) and are thus ranked low by the process of elimination.

## 7 Conclusions

In this paper, we have asked the question, how much, if any, and under what circumstances does approximating the specular occlusion not impact significantly on our perception of scene realism. We have concluded that the vast majority of glossy scenes can be approximated to some degree and that the most effective (in terms of computational complexity reduction versus decrease in perceived realism) approximation is directional ambient occlusion using spherical harmonics.

Highly complicated meshes are barely affected by any specular occlusion approximations as the second derivative of the surface is sufficiently complicated (and thus, the reflections are abundantly warped) to confuse our visual perception system in its plausibility decision task. Equally, glossy BRDFs with a broad lobe are minimally affected because even under no approximation, the visual appearance is smoothed due to the greater scattering of light. Finally, scenes where the environments consist of large light sources such as the Eucalyptus Grove are also suitable for approximations. This is because even for highly specular BRDFs, at any point on such a surface, the ratio between unoccluded and occluded rays is likely to be similar to other spatially nearby points. This creates very few harsh discontinuities in exit radiance which makes the scene more visually plausible. Only scenes consisting of highly specular, simple

meshes (such as the ground plane or other simple geometry) in the presence of small, bright light sources suffer significantly under all form of approximation. This type of scenes produce sharp discontinuities in the exit radiance at neighbouring points on the surface which our visual system is drawn towards. Because of extra focus at these harsh boundaries, any incorrect occlusion is registered and the realism of the scene suffers overall.

However, given an opportunity for direct comparison (either through a reference image or prior knowledge), approximations become noticeable in every type of glossy scene. Although our visual perception system can be fooled into accepting plausible scenes, given a strict choice between images, we become substantially more demanding in our perception of realism.

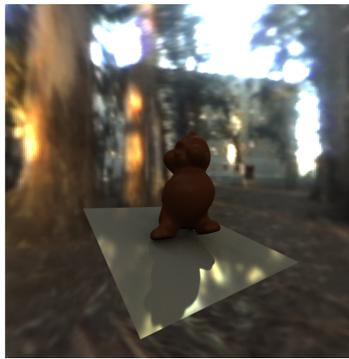
In summary, the various approximations considerably reduce the computational complexity (from the realm of hours per frame to near real-time rates) with little loss of perceived realism except in certain circumstances. With care being taken with regard to the setup of the scene, techniques such as directional ambient occlusion using spherical harmonics have considerable use in the domain of approximating specular occlusion.

## References

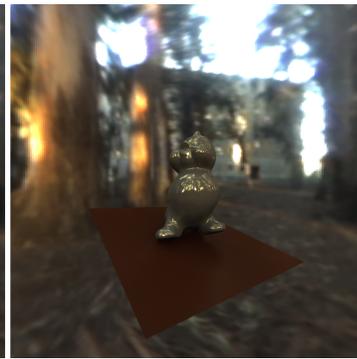
- ADELSON, E., AND PENTLAND, A. 1996. *Perception as Bayesian Inference*. Cambridge University Press, ch. The perception of shading and reflectance.
- BLAKE, A., AND BÜLTHOFF, H. 1990. Does the brain know the physics of specular reflection? *Nature* 343, 165–168.
- BUNNELL, M. 2005. Dynamic Ambient Occlusion and Indirect Lighting. In *GPU Gems 2*, M. Pharr, Ed. Addison Wesley, Mar., ch. 2, 223–233.
- COHEN, M., AND WALLACE, J. 1993. *Radiosity and Realistic Image Synthesis*. Academic Press Professional, Cambridge, MA.
- FLEMING, R., DROR, R., AND ADELSON, E. 2001. How do humans determine reflectance properties under unknown illumination? In *Proceedings of CVPR Workshop on Identifying Objects Across Variations in Lighting: Psychophysics and Computation*, 3–10.
- FLEMING, R., DROR, R., AND ADELSON, E. 2003. Real-world illumination and the perception of surface reflectance properties. *Journal of Vision* 3, 5 (jul), 347–368.
- FLEMING, R., TORRALBA, A., AND ADELSON, E. 2004. Specular reflections and the perception of shape. *Journal of Vision* 9, 798–820.
- GREEN, P., KAUTZ, J., MATUSIK, W., AND DURAND, F. 2006. View-dependent precomputed light transport using nonlinear gaussian function approximations. In *Proceedings of ACM Symposium in Interactive 3D Graphics and Games*, 7–14.
- GREEN, P., KAUTZ, J., AND DURAND, F. 2007. Efficient reflectance and visibility approximations for environment map rendering. *To Appear in Proceedings of Eurographics 2007*.
- GREENE, N. 1986. Environment Mapping and Other Applications of World Projections. *IEEE Computer Graphics & Applications* 6, 11 (Nov.), 21–29.
- HEIDRICH, W., DAUBERT, K., KAUTZ, J., AND SEIDEL, H.-P. 2000. Illuminating Micro Geometry Based on Precomputed Visibility. In *Proceedings of ACM SIGGRAPH*, 455–464.
- HUNTER, R., AND HAROLD, R. 1987. *The Measurement of Appearance (2nd edition)*. Wiley, New York.
- IMMEL, D., COHEN, M., AND GREENBERG, D. 1986. A Radiosity Method for Non-Diffuse Environments. In *Proceedings of ACM SIGGRAPH*, 133–142.



Reflective Dragon illuminated in Grace Cathedral



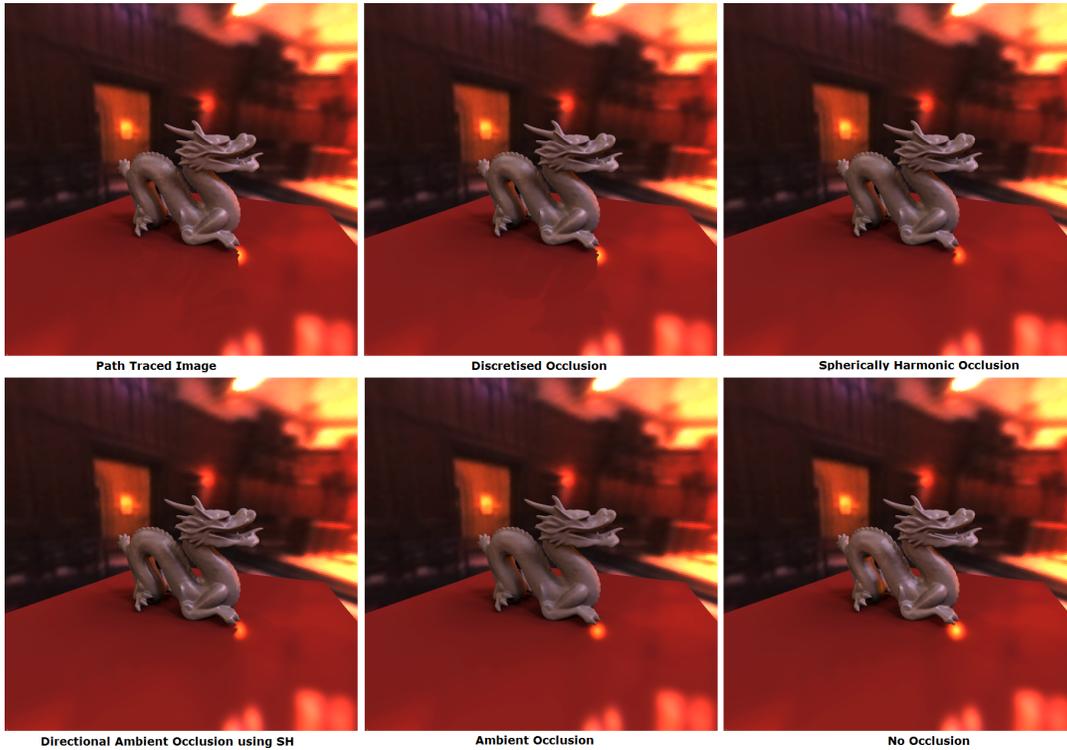
Dull Glossy Tweety illuminated in the Eucalyptus Grove



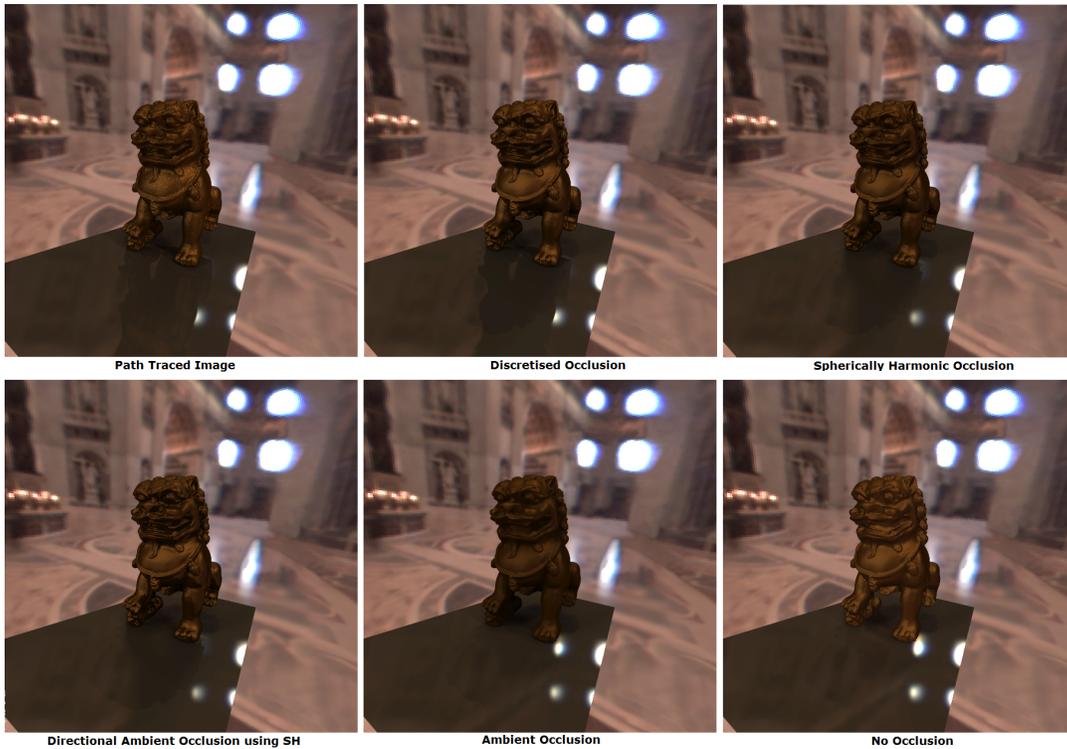
Reflective Tweety illuminated in the Eucalyptus Grove

**Figure 12:** Three out of the five scenes we have used for our user-study (path traced images are shown). The remaining two scenes can be seen in Figure 13 and 14.

- KAUTZ, J., AND MCCOOL, M. 2000. Approximation of Glossy Reflection with Prefiltered Environment Maps. In *Graphics Interface*, 119–126.
- KAUTZ, J., VÁZQUEZ, P.-P., HEIDRICH, W., AND SEIDEL, H.-P. 2000. A Unified Approach to Prefiltered Environment Maps. In *Eleventh Eurographics Workshop on Rendering*, 185–196.
- KAUTZ, J., SLOAN, P.-P., AND SNYDER, J. 2002. Fast, Arbitrary BRDF Shading for Low-Frequency Lighting Using Spherical Harmonics. In *Thirteenth Eurographics Workshop on Rendering*, 301–308.
- KAUTZ, J., LEHTINEN, J., AND AILA, T. 2004. Hemispherical Rasterization for Self-Shadowing of Dynamic Objects. In *Proceedings of the Eurographics Symposium on Rendering*, 179–184.
- KONTKANEN, J., AND LAINE, S. 2005. Ambient Occlusion Fields. In *Proceedings of ACM SIGGRAPH 2005 Symposium on Interactive 3D Graphics and Games*, 41–48.
- MÉNDEZ, A., SBERT, M., AND CATÁ, J. 2003. Real-time obscurances with color bleeding. In *SCCG '03: Proceedings of the 19th Spring Conference on Computer Graphics*, 171–176.
- MYSZKOWSKI, K., TAWARA, T., AKAMINE, H., AND SEIDEL, H.-P. 2001. Perception-guided global illumination solution for animation rendering. In *Proceedings SIGGRAPH 2001*, 221–230.
- NG, R., RAMAMOORTHI, R., AND HANRAHAN, P. 2003. All-Frequency Shadows Using Non-linear Wavelet Lighting Approximation. *ACM Transactions on Graphics* 22, 3 (July), 376–381.
- NG, R., RAMAMOORTHI, R., AND HANRAHAN, P. 2004. Triple Product Wavelet Integrals for All-Frequency Relighting. *ACM Trans. Graph.* 23, 3 (Aug.), 477–487.
- PRIKRYL, J., AND PURGATHOFER, W. 1999. Perceptually driven termination for stochastic radiosity. In *Proceedings of WSCG99*, 418–425.
- RAMAMOORTHI, R., AND HANRAHAN, P. 2002. Frequency Space Environment Map Rendering. In *Proceedings of ACM SIGGRAPH*, 517–526.
- RAMASUBRAMANIAN, M., PATTANAIK, S. N., AND GREENBERG, D. P. 1999. A perceptually based physical error metric for realistic image synthesis. In *Proceedings SIGGRAPH 1999*, 73–82.
- SLOAN, P.-P., KAUTZ, J., AND SNYDER, J. 2002. Precomputed Radiance Transfer for Real-Time Rendering in Dynamic, Low-Frequency Lighting Environments. In *Proceedings of ACM SIGGRAPH*, 527–536.
- SLOAN, P.-P., HALL, J., HART, J., AND SNYDER, J. 2003. Clustered Principal Components for Precomputed Radiance Transfer. *ACM Transactions on Graphics* 22, 3 (July), 382–391.
- SLOAN, P.-P., LUNA, B., AND SNYDER, J. 2005. Local, Deformable Precomputed Radiance Transfer. *ACM Transactions Graphics* 24, 3 (Aug.).
- STOKES, W. A., FERWERDA, J. A., WALTER, B., AND GREENBERG, D. P. 2004. Perceptual illumination components: a new approach to efficient, high quality global illumination rendering. *ACM Trans. Graph. (SIGGRAPH 2004)* 23, 3, 742–749.
- SUN, W., AND MUKHERJEE, A. 2006. Generalized Wavelet Product Integral for Rendering Dynamic Glossy Objects. *ACM Trans. Graph.* 25, 3, 955–966.
- TODD, J., NORMAN, J., KOENDERINK, J., AND KAPPERS, A. 1997. Effects of texture, illumination, and surface reflectance on stereoscopic shape perception. *Perception* 26, 807–822.
- TSAI, Y.-T., AND SHIH, Z.-C. 2006. All-frequency precomputed radiance transfer using spherical radial basis functions and clustered tensor approximation. *ACM Trans. Graph.* 25, 3, 967–976.
- VOLEVICH, V., MYSZKOWSKI, K., KHODULEV, A., AND KOPYLOV, E. A. 2000. Using the visual differences predictor to improve performance of progressive global illumination computation. *ACM Trans. Graph.* 19, 2, 122–161.
- WANG, R., TRAN, J., AND LUEBKE, D. 2004. All-frequency relighting of non-diffuse objects using separable brdf approximation. In *Eurographics Symposium on Rendering*, 345–354.
- ZHOU, K., HU, Y., LIN, S., GUO, B., AND SHUM, H.-Y. 2005. Precomputed Shadow Fields for Dynamic Scenes. *ACM Transactions on Graphics* 24, 3, 1196–1201.
- ZHUKOV, S., IONES, A., AND KRONIN, G. 1998. An Ambient Light Illumination Model. In *Eurographics Rendering Workshop 1998*, 45–56.



**Figure 13:** Grace Cathedral casts multiple shadows onto a dull glossy dragon statue. We compare the images produced utilising a variety of specular occlusion approximations. Only a few, subtle differences can be seen between the images explaining why the scores were so similar when no direct comparison was allowed.



**Figure 14:** St Peter's Basilica casts sharp shadows and reflections onto a dull glossy chinese lion statue. We compare the images produced utilising a variety of specular occlusion approximations. It is easy to see the preference for more accurate occlusion in this scene. Ambient occlusion and no occlusion images contain an obviously wrong highlight next to the lion's left front paw. The spherical harmonics occlusion and directional ambient occlusion using SH images fail to replicate the reflection of the lion with enough detail.