Perceptual Influence of Approximate Visibility in Indirect Illumination

Insu Yu[⋆] Andrew Cox[⋆] Min H. Kim[⋆] Tobias Ritschel[†] Thorsten Grosch[†] Carsten Dachsbacher[‡] Jan Kautz[⋆]

University College London[⋆] MPI Informatik[†] Universität Stuttgart[‡]

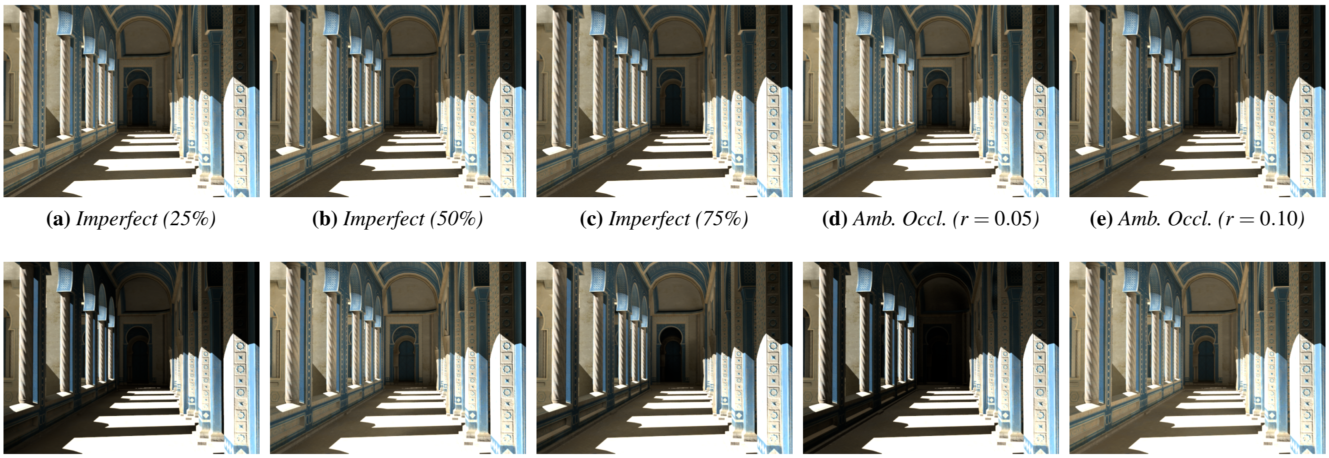

(a) *Imperfect (25%)* **(b)** *Imperfect (50%)* **(c)** *Imperfect (75%)* **(d)** *Amb. Occl.* ($r = 0.05$) **(e)** *Amb. Occl.* ($r = 0.10$)

(f) *Amb. Occl.* ($r = 0.20$) **(g)** *Dir. Amb. Occl.* ($r = 0.05$) **(h)** *Dir. Amb. Occl.* ($r = 0.10$) **(i)** *Dir. Amb. Occl.* ($r = 0.20$) **(j)** *No Visibility*

Figure 1: *Renderings of the arches scene, where the indirect illumination in each image is computed with a different visibility approximation. Our psychophysical study shows that many of these visibility approximations produce images that are perceptually very similar to reference renderings (cf. Fig. 3).*

Abstract

In this paper we evaluate the use of approximate visibility for efficient global illumination. Traditionally, accurate visibility is used in light transport. However, the indirect illumination we perceive on a daily basis is rarely of high frequency nature, as the most significant aspect of light transport in real-world scenes is diffuse, and thus displays a smooth gradation. This raises the question of whether accurate visibility is perceptually necessary in this case. To answer this question, we conduct a psychophysical study on the perceptual influence of approximate visibility on indirect illumination. This study reveals that accurate visibility is not required and that certain approximations may be introduced.

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: visibility, global illumination, perception

1 Introduction

Global illumination effects such as indirect illumination are known to be perceptually important [Stokes et al. 2004] but are often omitted or coarsely approximated due to their high rendering cost—especially in interactive applications. One of the most expensive components when computing indirect illumination is visibility determination, i.e., determining whether two points are mutually visible or not. Commonly, this is done accurately using ray-casting. While methods do exist to speed up accurate visibility queries [Cohen-Or et al. 2003], it is unknown whether accurate visibility for indirect illumination is perceptually important at all. And if it is not, which kinds of approximations are perceptually acceptable?

Nonetheless, previous methods have already used visibility approximations to speed up indirect illumination computation; for instance, Sillion et al. [1995] used volumetric representations, Arikan et al. [2005] ignored occlusion from nearby geometry, Sloan et al. [Sloan et al. 2007] approximated binary visibility with low-order spherical harmonics, and Ritschel et al. [Ritschel et al. 2008b] used imperfect depth.

Our goal is to evaluate the perceptual influence of visibility and visibility approximations on indirect illumination. We conduct a formal study where scenes under global illumination are rendered with different approximations, and are then compared against each other. Approximations in our study include imperfect visibility [Ritschel et al. 2008b], ambient occlusion [Zhukov et al. 1998] and directional ambient occlusion [Sloan et al. 2007; Ritschel et al. 2009]. A psychophysical analysis is performed on the data and we determine which visibility approximations are perceptually acceptable. As it turns out, using visibility approximations to compute indirect illumination can lead to renderings that are considered to be as realistic as accurate solutions. This validates the previous use of visibility approximations in global illumination.

Note that in this paper we are only interested in indirect illumination, as visibility for direct illumination such as soft shadowing has been investigated previously [Annen et al. 2008; Fernando 2005; Akenine-Moeller et al. 2004].

Our main contributions include:

- A formal study on the perceived realism of renderings computed with visibility approximations for indirect illumination.
- A formal study on the similarity of renderings using visibility approximations to accurate reference renderings.

After reviewing previous work in Section 2, we describe the visibility approximations that we have chosen to study the perceptual influence of in Section 3. In Section 4 we present our experimental procedure and discuss an analysis of perceived realism. Finally, we conclude and discuss future work in Section 5.

2 Related Work

Global Illumination Methods Many standard global illumination methods, such as path tracing, photon mapping, and ray-tracing [Dutré et al. 2006] use accurate visibility for indirect as well as direct illumination, usually by intersecting rays with the scene geometry. Even popular techniques that are geared towards efficient computation, such as instant radiosity [Keller 1997] and instant global illumination [Wald et al. 2003], both based on virtual point lights (VPLs), use accurate visibility — either through shadow volumes or ray-casting respectively. Finite element methods, such as radiosity [Cohen and Wallace 1993], are somewhat different, as they require (fractional) visibility between pairs of finite elements. In practice it is evaluated using Monte Carlo sampling or hemicubes [1993].

Methods exist that make a more explicit use of visibility approximations. Sillion and Drettakis [1995] proposed a radiosity method that approximates visibility such that a user-defined feature-size can still be resolved while blurring out smaller features; no distinction was made between direct and indirect illumination. The lightcuts method [Walter et al. 2005] groups a large number of light samples (VPLs representing the indirect illumination) into a hierarchy to speed up rendering. Visibility between a surface point and a node in the hierarchy (a collection of VPLs) is approximated with a single shadow ray. However, the algorithm bounds the resulting error, such that no perceptible artifacts appear. Arikan et al. [2005] propose an approximate global illumination method where irradiance arriving from nearby and distant geometry is decoupled: All incident lighting from nearby geometry is simply integrated without computing visibility at all. Arvo et al. [1994] analyze error in global illumination algorithms, including error from visibility discretization, but do not propose algorithms to make use of this information. Durand et al. [2005] provide a mathematical framework for the analysis of light transport, including visibility. Methods have been proposed to reduce the geometric complexity when computing global illumination solutions [Rushmeier et al. 1993; Christensen et al. 2003; Tabellion and Lamorlette 2004], however, this is different from using approximative visibility, since the shading itself changes when the geometry is approximated.

None of the above methods have considered what the perceptual influence of visibility approximations on indirect illumination really is. In fact, many of the above algorithms do not make a distinction between indirect and direct illumination. In this paper, we will study the influence on indirect illumination and demonstrate that approximations can be perceptually acceptable.

Real-Time Global Illumination Many techniques for real-time global illumination, such as precomputed radiance transfer (PRT) techniques, only support lighting changes and are restricted to static [Sloan et al. 2002] or semi-static scenes [Pan et al. 2007; Iwasaki et al. 2007; Ritschel et al. 2008a], which allow the movement only of rigid objects. These restrictions enable the precomputation of (parameterized) light transport including all visibility information.

Interactive global illumination for dynamic models is possible if visibility for indirect illumination is completely neglected [Dachsbacher and Stamminger 2005; Dachsbacher and Stamminger 2006]. Bunnell [2005] presented GPU-based indirect illumination using a hierarchical link structure, where all visibility between elements was approximated by ambient occlusion [Zhukov et al. 1998]. Recently, Dong et al. [2007] proposed an interactive global illumination method for small scenes, where visibility was implicitly evaluated on the hierarchical finite element structure, i.e., with similar discretization as with omnidirectional shadow mapping. Interactive, dynamic global illumination for moderately complex scenes is possible [Dachsbacher et al. 2007] using a variant of hierarchical radiosity that replaces explicit visibility queries with an iterative process using *antiradiance* (negative light). This also requires using a directionally discretized data structure, akin to omnidirectional shadow mapping. Sloan et al. [2007] demonstrate real-time indirect illumination but only low-frequency incident radiance as well as low-frequency visibility are supported, as these are represented with a small number of spherical harmonics. Ritschel et al. [2008b] have proposed imperfect shadow maps in the context of instant radiosity to speed up the computation of indirect illumination. Their results indicate that imperfect visibility for indirect illumination produces acceptable results.

As outlined, some forms of visibility approximation have been used in the real-time rendering context but their perceptual impact has not been formally studied. We will evaluate this influence and demonstrate a valid basis for interactive global illumination to exploit visibility approximations.

Perceptually-Based Rendering Since computing global illumination solutions is very expensive, perceptually-based rendering methods have been explored. Their goal is to speed up the process by taking the limits of the human visual system into account. Using the visual differences predictor [Daly 1993] to determine whether an approximation is indistinguishable from a reference image is a common approach [Volevich et al. 2000]. Perceptually-based error metrics have also been used to efficiently compute animation sequences [Myszkowski et al. 2001]. Stokes et al. [2004] separate the rendering process into individual illumination components and identify perceptually important ones, which can be used to speed up the computation [Debattista et al. 2005]. The work by Vangorp et al. [2007] predicts how a change in incident illumination affects the appearance of an object, depending on its geometry and material, which can be used to speed up rendering algorithms. Drettakis et al. [2007] use the visual differences predictor in an interactive setting to estimate masking between depth layers of a scene and thereby control the level-of-detail of polygonal meshes. Furthermore, it has been shown that accurate occlusion for glossy reflections is not always necessary [Kozlowski and Kautz 2007]; an observation that has been exploited in rendering glossy reflections from environment maps [Green et al. 2007].

3 Indirect Illumination with Visibility Approximations

We will first introduce the necessary background on global illumination in this section, as well as meaningful visibility approximations, and then study the perceptual influence on scenes with indirect illumination.

3.1 Rendering Equation

The rendering equation [Kajiya 1986] can be written in a simple form to describe global illumination as follows:

$$L(\mathbf{x},\omega) = E(\mathbf{x},\omega) + L_r(\mathbf{x},\omega) \tag{1}$$

By using the operator form, the *reflection* operator $\mathbf{K}$ takes all the incident radiance, performs shading and outputs outgoing radiance L for every point $\mathbf{x}$ and direction ω:

$$L_r(\mathbf{x},\omega) = (\mathbf{K}L_i)(\mathbf{x},\omega) = \int_{\Omega^+} L_i(\mathbf{x},\omega_i) f_r(\mathbf{x},\omega_i \to \omega)(\omega_i \cdot \mathbf{n}_\mathbf{x})\mathrm{d}\omega_i, \tag{2}$$

where $f_r()$ is the BRDF, $\mathbf{x}$ is the current position, $\mathbf{n}_\mathbf{x}$ is the normal at $\mathbf{x}$, ω_i is the light direction, and ω is the outgoing direction. Incident radiance L_i is computed through the *geometry* operator $\mathbf{G}$, which collects radiance from the first hit-point along directions ω:

$$L_i(\mathbf{x},\omega) = (\mathbf{G}L)(\mathbf{x},\omega) = \sum_{\mathbf{y}\in\mathbf{I}} L(\mathbf{y},-\omega)V(\mathbf{x},\mathbf{y}), \tag{3}$$

where $\mathbf{y} \in \mathbf{I}$ are all the intersections along the direction ω starting at $\mathbf{x}$, and $V()$ is the binary visibility between point $\mathbf{x}$ and $\mathbf{y}$. For the purpose of studying visibility approximations, we use a geometry operator, where the amount of light emitted from a point $\mathbf{y}$ toward $\mathbf{x}$ is decoupled from the actual visibility between those two points, i.e., the operator goes through all possible points $\mathbf{y}$ that can contribute to $\mathbf{x}$ along the direction $-\omega$ and explicitly tests visibility between all point pairs. Note this is different from Arvo et al.'s [1994] formulation, who explicitly find the closest point, and allows us to describe visibility approximations.

Outgoing radiance is then computed with the full rendering equation:

$$L(\mathbf{x},\omega) = E(\mathbf{x},\omega) + (\mathbf{KG}L)(\mathbf{x},\omega), \tag{4}$$

with $E()$ being the self-emission at point $\mathbf{x}$. Expanding this equation into a Neumann series yields:

$$L = E + \mathbf{KG}E + (\mathbf{KG})^2E + (\mathbf{KG})^3E + \dots \tag{5}$$

The operator notation also allows us to study visibility approximations that vary with the number of light bounces. For instance, light that is reflected for the first time (geometry operator $\mathbf{G}_1$) might not use an approximation, whereas light that is being reflected the second time (operator $\mathbf{G}_2$) does. For direct lighting (operator $\mathbf{G}_0$), we assume pixel-accurate visibility (which we refer to as $\mathbf{G}_\mathrm{a}$), since we are only interested how indirect illumination is influenced by approximate visibility. This yields the following general form of the Neumann expansion:

$$L = E + \mathbf{K}\dot{\mathbf{G}}_0E + \mathbf{K}\ddot{\mathbf{G}}_1\mathbf{K}\ddot{\mathbf{G}}_0E + \mathbf{K}\dddot{\mathbf{G}}_2\mathbf{K}\dddot{\mathbf{G}}_1\mathbf{K}\dddot{\mathbf{G}}_0E + \dots, \tag{6}$$

where superscript dots refer to path-length (one dot for paths of length $l = 1$, and so on) and subscript numbers refer to the bounce number. This path-length-dependent notation it useful in the context of virtual point lights, where most likely only the last bounce is approximated (e.g., using imperfect shadow maps [Ritschel et al. 2008b]).

Please note that the Neumann series only converges if the spectral radius of the operator $\mathbf{KG} < 1$. When visibility approximations are used, this might be violated, however, we have not encountered this in practice.

Studying visibility approximations for diffuse indirect illumination is a natural start, especially since the dimensionality of the parameter-space is kept rather small. The inclusion of glossy BRDFs would lead to a significant increase, as the BRDF itself introduces more parameters and view-dependency needs to be considered as well. Furthermore, diffuse indirect illumination is perceptually the most important indirect component (according to [Stokes et al. 2004]). Hence, we limit ourselves to studying only diffuse indirect illumination.

3.2 Approximations

In order to study the influence of the visibility function $V(\mathbf{x},\mathbf{y})$, we need to decide on approximations thereof that can lead to efficient rendering techniques. We introduce and examine various possible approximations in the following (see Fig. 2 for an overview).

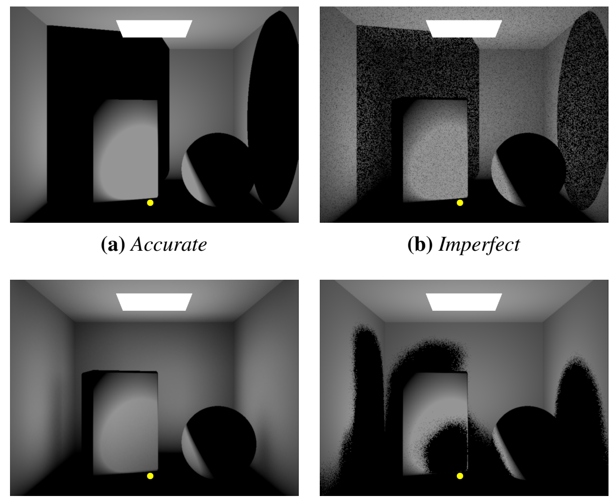

(a) *Accurate* **(b)** *Imperfect* **(c)** *Ambient Occlusion* **(d)** *Dir. Ambient Occlusion*

Figure 2: *Visibility of every point in the scene with a single VPL (in front of the box) is shown. (a) Accurate visibility produces sharp shadow boundaries ($\mathbf{G}_\mathrm{a}$). (b) Imperfect visibility introduces noise ($\mathbf{G}_\mathrm{imp}^{25\%}$). (c) Ambient occlusion produces very smooth visibility ($\mathbf{G}_\mathrm{ao}^{r=0.2}$). (d) Directional ambient occlusion produces partially correct shadows but also adds some extraneous ones ($\mathbf{G}_\mathrm{dao}^{r=0.2}$).*

Imperfect Visibility Recently, *imperfect visibility* [Ritschel et al. 2008b] has been proposed as a means to speed up instant radiosity. We explore this approximation by randomly setting $N\%$ of all visibility queries to either 0 or 1 (uniform distribution). The operator is referred to as $\mathbf{G}_\mathrm{imp}$.

Ambient Occlusion Ambient occlusion was originally proposed as an ambient illumination model [Zhukov et al. 1998] but has also been used as an approximation in indirect illumination [Bunnell 2005]. For a given receiver point $\mathbf{x}$, the percentage of "visible sky" is stored as the ambient occlusion factor

$$a_\mathbf{x} = \frac{1}{\pi}\int_{\Omega+} V_\mathbf{x}^r(\omega_i)(\mathbf{n}_\mathbf{x}\cdot\omega_i)\mathrm{d}\omega_i. \tag{7}$$

Here $V_\mathbf{x}^r$ refers to the visibility from $\mathbf{x}$ in direction ω_i. Note that for enclosed models $V_\mathbf{x}^r$ will always be 0, as any ray will inevitably intersect geometry. In order to support enclosed models, the ray length of ω_i is commonly restricted to a user-defined radius r (blockers beyond r are ignored), which is why we refer to the visibility as $V_\mathbf{x}^r$. The factor $a_\mathbf{x}$ is then used as the visibility for all emitter points $\mathbf{y}$: $V(\mathbf{x},\mathbf{y}) = a_\mathbf{x}$. We will refer to this operator as $\mathbf{G}_\mathrm{ao}$.

Directional Ambient Occlusion Recently, Sloan et al. [2007] proposed to represent binary visibility information using low-order spherical harmonics in the context of indirect illumination. Somewhat related, Ritschel et al. [2009] proposed to use directional screen-space ambient occlusion when gathering irradiance. We will investigate a mixture between the two techniques, *directional ambient occlusion*, which adds a directional component to ambient occlusion. In other words, ambient occlusion is not a scalar anymore but a function that depends on the (hemispherical) direction $\omega_i \in \Omega_+$: $a_{\mathbf{x}}(\omega_i)$. Storing binary visibility for a large number of directions, i.e. $a_{\mathbf{x}}(\omega_i) = V^r_{\mathbf{x}}(\omega_i)$, would be too costly for real-world applications. Instead we will investigate a low-frequency, 5th-order spherical harmonics (SH) approximation to it, which is defined as follows:

$$V(\mathbf{x},\mathbf{y}) = a_{\mathbf{x}}(\omega_i) = \sum_{j=1}^{25} y_j(\omega_i) V^r_{\mathbf{x},j}, \tag{8}$$

where y_j are the spherical harmonics basis functions and $V^r_{\mathbf{x},j}$ are the coefficients of $V^r_{\mathbf{x}}$ when projected into SH. The geometric operator will be referred to as $\mathbf{G}_{\text{dao}}$.

It is educational to consider what this approximations means intuitively: Occlusions between the receiver $\mathbf{x}$ and the emitter VPL $\mathbf{y}$ might be missed if $\mathbf{y}$ is further than distance r away, as blockers that are further than distance r are simply not represented by $a_{\mathbf{x}}(\omega_i)$. Vice versa, emitter points $\mathbf{y}$ that are closer than the distance r might get shadowed by occluders that are actually further away than $\mathbf{y}$, as binary visibility is stored (distance information is lost). In fact, this can be seen in Figure 2d, which shows a phantom shadow on the box.

No Visibility Recent interactive techniques have computed indirect illumination without taking visibility into account at all [Dachsbacher and Stamminger 2005; Dachsbacher and Stamminger 2006]. We will validate whether no visibility can be a useful "approximation" for indirect illumination. The visibility function then reduces to $V(\mathbf{x},\mathbf{y}) = 1$. The geometric operator is called $\mathbf{G}_{\text{no}}$.

There are many other possible approximations. We chose to include these approximations in our study as they show promise for practical algorithms.

3.3 Benefits of Approximate Visibility

The key idea we are exploring in this paper is the observation that accurate visibility might not be perceptually necessary for indirect illumination, and smart approximations can therefore lead to computational savings. Certain global illumination algorithms more easily benefit from approximate visibility. In particular, those algorithms benefit that shoot radiance from one point to all other surfaces, e.g., instant radiosity [Keller 1997] using virtual point lights or lightcuts [Walter et al. 2005]. Of course, a pre-computed visibility field such as (directional) ambient occlusion can be adapted to standard global illumination methods, such as path tracing and ray tracing in order to speed up indirect shadow ray computation.

4 Perceptual Influence of Visibility Approximations

Our goal is to evaluate the influence of visibility approximations presented in Section 3.2 on indirect global illumination. We first measure the perceptual influence of different visibility approximations by carrying out a series of psychophysical experiments. We then analyze this data to evaluate how different approximated solutions affect the perceived realism of renderings under global illumination.

	$l=1$	$l=2$		$l=3$			
	$\mathbf{G}_0^{\cdot}$	$\mathbf{G}_1^{\cdot\cdot}$	$\mathbf{G}_0^{\cdot\cdot}$	$\mathbf{G}_2^{\cdot\cdot\cdot}$	$\mathbf{G}_1^{\cdot\cdot\cdot}$	$\mathbf{G}_0^{\cdot\cdot\cdot}$	...
case 1	$\mathbf{G}_\text{a}$	$\mathbf{G}_\text{imp}^{25\%}$	$\mathbf{G}_\text{a}$	$\mathbf{G}_\text{imp}^{25\%}$	$\mathbf{G}_\text{a}$	$\mathbf{G}_\text{a}$	...
case 2	$\mathbf{G}_\text{a}$	$\mathbf{G}_\text{imp}^{50\%}$	$\mathbf{G}_\text{a}$	$\mathbf{G}_\text{imp}^{50\%}$	$\mathbf{G}_\text{a}$	$\mathbf{G}_\text{a}$	...
case 3	$\mathbf{G}_\text{a}$	$\mathbf{G}_\text{imp}^{75\%}$	$\mathbf{G}_\text{a}$	$\mathbf{G}_\text{imp}^{75\%}$	$\mathbf{G}_\text{a}$	$\mathbf{G}_\text{a}$	...
case 4	$\mathbf{G}_\text{a}$	$\mathbf{G}_\text{ao}^{r=0.05}$	$\mathbf{G}_\text{a}$	$\mathbf{G}_\text{ao}^{r=0.05}$	$\mathbf{G}_\text{a}$	$\mathbf{G}_\text{a}$	...
case 5	$\mathbf{G}_\text{a}$	$\mathbf{G}_\text{ao}^{r=0.1}$	$\mathbf{G}_\text{a}$	$\mathbf{G}_\text{ao}^{r=0.1}$	$\mathbf{G}_\text{a}$	$\mathbf{G}_\text{a}$	...
case 6	$\mathbf{G}_\text{a}$	$\mathbf{G}_\text{ao}^{r=0.2}$	$\mathbf{G}_\text{a}$	$\mathbf{G}_\text{ao}^{r=0.2}$	$\mathbf{G}_\text{a}$	$\mathbf{G}_\text{a}$	...
case 7	$\mathbf{G}_\text{a}$	$\mathbf{G}_\text{dao}^{r=0.05}$	$\mathbf{G}_\text{a}$	$\mathbf{G}_\text{dao}^{r=0.05}$	$\mathbf{G}_\text{a}$	$\mathbf{G}_\text{a}$	...
case 8	$\mathbf{G}_\text{a}$	$\mathbf{G}_\text{dao}^{r=0.1}$	$\mathbf{G}_\text{a}$	$\mathbf{G}_\text{dao}^{r=0.1}$	$\mathbf{G}_\text{a}$	$\mathbf{G}_\text{a}$	...
case 9	$\mathbf{G}_\text{a}$	$\mathbf{G}_\text{dao}^{r=0.2}$	$\mathbf{G}_\text{a}$	$\mathbf{G}_\text{dao}^{r=0.2}$	$\mathbf{G}_\text{a}$	$\mathbf{G}_\text{a}$	...
case 10	$\mathbf{G}_\text{a}$	$\mathbf{G}_\text{no}$	$\mathbf{G}_\text{a}$	$\mathbf{G}_\text{no}$	$\mathbf{G}_\text{a}$	$\mathbf{G}_\text{a}$	...

Table 1: *Visibility approximations investigated for our study (see Fig. 1 for renderings). Visibility approximations are indicated for each path length l, in reverse order of bounces (see Eq. 6). Superscripts indicate the parameter setting for the particular approximation: for imperfect visibility the percentage of corrupted visibility queries, for ambient and directional ambient occlusion the max. ray length r as a fraction of scene diameter.*

4.1 Stimuli

In order to measure the perceptual influence of visibility approximations, we need to define and render a set of meaningful test stimuli. Since visibility approximations are most likely to be used in real-time or interactive rendering applications we decided to opt for pre-rendered 5-second video sequences instead of static images. By using video sequences, we are able to ensure that any temporal artifacts caused by a visibility approximation can be taken into account by the observers. We are interested in "normal" scenes, where direct illumination contributes considerably to the scene, as well as more contrived scenes, which are, for instance, mainly lit by indirect illumination. A full range of scenes will help to determine when approximations are valid and when they are not; our range of scenes is shown in Fig. 3. The use of textures adds realism to a scene but potentially masks shading artifacts [Ferwerda et al. 1997]. Thus we have chosen to limit the use of highly structured textures in all but one scene.

Since there is a large number of possible parameterizations of visibility approximations we had to make a choice of which ones to use. We opted to use parameters that can be used to speed up global illumination algorithms such as lightcuts and instant radiosity. After first experimenting with various $N\%$ for imperfect visibility and distance r for (directional) ambient occlusion algorithms, we computed approximate global illumination solutions (see Table 1 and Fig. 1) for each of our test scenes using:

- Accurate visibility for each bounce yielding the reference solution (Fig. 3),
- Imperfect visibility (25%, 50% and 75% corruption), where only the last indirect bounce is approximated (case 1–3),
- Ambient occlusion with ray length set to 5%, 10%, and 20% of scene diameter (case 4–6),
- Directional ambient occlusion (5th order SH) with the same ray length as above (case 7–9),
- No visibility at all for indirect bounces (case 10).

Note that direct lighting uses accurate visibility in all cases. These 10 approximations to the rendering equation span a considerable variety of parameter settings and visibility approximations. They

Figure 3: *Test scenes. The first scene is a historic building with strong direct lighting that indirectly illuminates the arches (light and camera are moving in video). The tea house scene exhibits strong colored indirect illumination (light is moving in video). The living room scene is mostly lit by direct light with some indirect illumination visible from the orange wall (camera is moving in video). The Sponza scene is lit by strong sunlight, all other illumination is indirect (camera is moving in video).*

are chosen to be directly applicable to existing algorithms, such as instant radiosity, lightcuts and path tracing.

The test scenes were rendered using an instant radiosity (IR) based renderer that has been adapted to account for different visibility approximations. To this end, we compute incident radiance at every point using the formulation of Eq. 6. All our renderings are generated with four indirect bounces of illumination (see Table 1). Further bounces contribute very little energy and were omitted. Since instant radiosity may introduce shading artifacts due to the random walk, we have used a high number of VPLs to avoid such artifaces. All our videos were rendered at an 640×480 resolution and gamma corrected with a global scale factor. The renderings were computed on a cluster of PCs due to the long computation time for the video sequences (100 images per sequence with one to four hours rendering time per image).

4.2 Experimental Procedure

The goal of our psychophysical experiments is not to simply determine whether visibility approximations produce results that are perceptually identical to accurate reference renderings. Instead, we are interested in the following three questions: Do renderings with visibility approximations appear realistic (*perceived realism*)? How *perceptually similar* are these renderings to a reference? Does perceived realism correlate with similarity to reference renderings?

To this end, we conducted two separate psychophysical experiments. The *paired comparison plus category* method [Scheffé 1952] was used to quantify the perceptual similarity of each of these approximate renderings to its reference. The *ordinal rank order* method [Bartleson 1984] was employed to determine the perceived realism of each rendering.

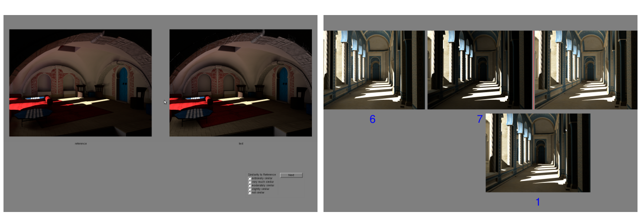

(a) *Similarity to reference* **(b)** *Ranking according to realism*

Figure 4: *Viewing pattern observed by participants. Left: paired comparison plus category experiment. Right: Ordinal rank order experiment.*

In the comparison experiment, a participant was presented with a pair of rendered videos and, asked to answer how *similar to the reference* (left) the test rendering (right) is, using a five-point scoring scale (see Fig. 4a). The technique is thus a combination of a five-point category rating scale and a pair comparison. Participants estimate the difference between a pair and assign a number to this difference. These categories are labeled with the following descriptions: 1. (*not similar*), 2. (*slightly similar*), 3. (*moderately similar*), 4. (*very much similar*), and 5. (*extremely similar*), adapted from Meilgaard et al. [1991]. In the other experiment, a participant was shown a set of rendered videos presented in a row (eleven videos, consisting of the reference rendering and the ten approximations). The participant was then asked to rank the videos in order from highest to lowest by *perceived realism*[1]. The videos are initially loaded in random order after which the participant may freely pan across them and reorder them with a drag-and-drop mouse operation, aided by discrete zoom-in/out and pause functions (see Fig. 4b).

The *paired comparison plus category* experiment and *ordinal rank order* experiment were conducted in two sessions on different days. Fourteen color-normal observers took part in the paired comparison study and eighteen in the ranking experiment. In both experiments, over half of the participants were computer scientists with some imaging background. The participants were given instructions beforehand which contained a brief description of the task (similarity to reference or rank by realism). A short training session was given to familiarize with mouse navigation control and key functions (zoom-in/out and pause). The experiment was conducted in a controlled environment under dim viewing conditions to maximize differences. Participants were asked to adapt 5-10 minutes to the illumination conditions before starting the experiment. A calibrated Dell E248WFP 24" monitor was used with participants seated at a distance of about 60cm. All stimulus videos were presented over a mid-gray background; a short blanking of the screen separated the phases.

In the paired comparison plus category experiment, the participants made 40 estimates (10 approximations compared to a reference for each of 4 scenes), for which they spent about 15 minutes. In the rank order experiment, participants ranked 4 different scenes, sorting 11 videos in each, for which the majority required between 25 and 35 minutes.

The inter-observer coefficient of variation (ratio of the standard deviation to the mean, CV) of the 14 participants of the categorical judgments is 17.40%. The inter-observer CV of the 18 participants of the ranking experiment is 21.99%. Four observers repeated the rank experiment twice in order to judge repeatability. The average CV between the two experiments was 21.93%.

[1]Note that this experiment was conducted first, and hence participant did not know there was a reference rendering.

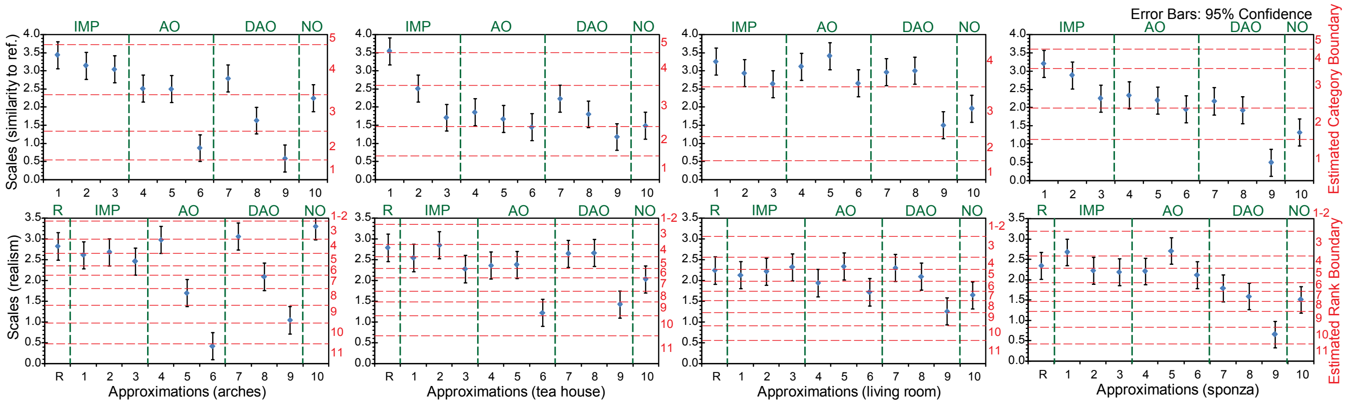

Figure 5: *Analysis of the Similarity-to-Reference and Rank-by-Realism experiment (results are by scene). We show the similarity to the reference video for each approximation (1-10) and the rank-by-realism for each approximation (1-10) including the reference video (R). The estimated category / rank boundaries are indicated as red colored dot-lines; e.g., all the imperfect visibility (IMP) and ambient occlusion (AO) approximations for the* living room *scene were considered* very much similar *(category 4, perceptual scales between 2.67 and 4.05).*

4.3 Results and Analysis

The similarity experiment yielded similarity scores on a five-point scale relating the reference videos to the videos using approximations. We analyzed this data using perceptual scaling. The five-point scores were scaled using the "Law of Categorical Judgment" by Torgerson [1958]. This is an extension to Thurstonian [1927] scaling that allows for several categories. The result of the analysis includes scale values, as well as estimates of the category boundaries. This means the scale values can be related to the original categories (from *not similar* to *extremely similar*). The estimated *scale values* are on a perceptually-uniform scale, which allows one to judge relative differences in similarity of the visibility approximations to the reference. The results are summarized in Fig. 5 (top row).

The imperfect visibility approximations (IMP) are all considered *very much similar* or *moderately similar* to the reference video in all scenes. Hence, this kind of visibility approximations is applicable to a wide range of scenes, especially considering that the realism ranks for videos using this approximation is high as well (see below). Ambient occlusion (AO) of visibility approximation is considered *very much similar* or *moderately similar* to the reference in most scenes, when using a max. radius size of 0.1. However, if the radius is larger than 0.1, it is rated only *moderately similar* in most scenes (statistically significant at 5% level). Directional ambient occlusion (DAO) is also considered *very much similar* or *moderately similar* in most scenes, when using at a max. radius of 0.1 (like ambient occlusion). However, large radii rate relatively worse here and yield only *slightly similar* results in most scenes (statistically significant at 5% level) (see Fig. 7, top left). Surprisingly, the case of no visibility is considered to be *moderately similar* to the reference video, ranking higher than the worst AO and DAO.

Interestingly, the perceived similarity to the reference (when using visibility approximations) seems to be linked to the amount of indirect illumination. The tea house and Sponza scenes have more dominant indirect illumination than the arches and the living room scenes and, as shown in Fig. 6, the similarity to the reference (averaged over all approximations) decreases. However, this decrease is not statistically significant.

Even though raw ordinal rank data has some limitations in terms of quantitative analysis [Guilford 1954], ranking offers an intuitive user interface and can be performed more quickly than a complete pairwise comparison experiment. We analyzed the resulting perceptual

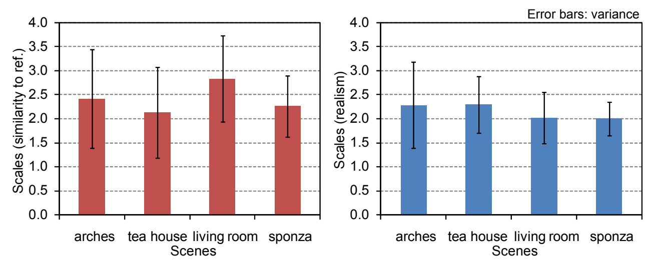

Figure 6: *Scene Dependency. Averaged scales (over all approximations) of similarity-to-reference as well as perceived realism.*

data by adopting Torgerson's [1958] scaling method. The individual results can be found in Figure 5 (bottom row).

The reference video, as well as the visibility approximations using imperfect visibility, ambient occlusion ($r = 0.05$), and direct ambient occlusion ($r = 0.05$) are ranked as equally realistic (statistically significant at 5% level) when averaged over all four scenes (see Fig. 7, bottom left). Based on the results of Fig. 5 and 7, we conclude that using the imperfect visibility approximation of up to 75% does not strongly impact the perceived realism of a scene under global illumination. When using AO and DAO, the chosen radius makes a big difference. Larger radii ($r = 0.2$) were generally perceived as less realistic and this difference in realism is statistically significant for all scenes. In that case, AO and DAO are even perceived as less realistic than the "no visibility" approximation.

The right side of Fig. 7 shows the correlation between the similarity of an approximation to the reference video and the perceived realism (averaged over all scenes). As it turns out, the two properties are correlated with a coefficient of 0.84 when using most approximations (excluding the reference, IMP 20%, and IMP 50%). Imperfect visibility with less than 75% corruption as well as the reference does not make any noticeable difference in perceived realism (see vertical axis). This means that participants considered these videos as being the most realistic as well as *very much similar* to the reference video.

4.4 Discussion

The results show that visibility approximations can be used in global illumination while maintaining an appearance that is perceptually

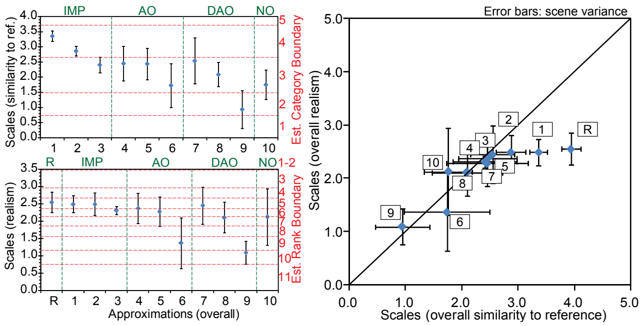

Figure 7: *Analysis of similarity-to-reference and perceived realism (overall). The top left graph shows the similarity-to-reference scales (R refers to reference) averaged over all approximations (1–10, see Table 1). The bottom left graph shows the averaged realism scales of all approximations and the reference video. The right graph depicts the correlation between participants' perceived realism and the similarity-to-reference (from Fig. 5).*

similar to a reference solution. The imperfect visibility approximations generally ranked higher in *perceived realism* than ambient occlusion or directional ambient occlusion. In other words, visibility that is highly corrupted (but at random) is preferred by human eyes over inaccurate visibility such as AO and DAO. The evaluation of the experimental data also shows that most visibility approximations are *very much similar* to the reference when direct illumination is dominant in the view of the scene (e.g., in the living room scene).

Our study therefore validates the use of visibility approximations in previous work, especially the use of imperfect shadow maps in instant radiosity [Ritschel et al. 2008b].

We have chosen to use video sequences as stimuli, which raises the question of whether using static images would have led to different results. Before using video sequences, we conducted a pilot study using static images only (slightly different scenes). The overall result was comparable to the result presented here – with imperfect visibility ranked highest. This is not surprising, as the video sequences do not show any unexpected temporal effects.

5 Conclusions

We have studied the perceptual influence of visibility approximations on indirect illumination. The experiments revealed that using appropriate visibility approximations yields results that are perceptually very similar to reference renderings. Further, many visibility approximations yield renderings that are perceived to be realistic despite perceptible differences to reference renderings.

In the future, we would like to study other approximations that are commonly made in real-time global illumination techniques. For instance, we would like to study whether the use of direct illumination that is only of a low-frequency nature is acceptable. Furthermore, we would like to generalize our user study to include general BRDFs. This is rather challenging as there are many parameters that may influence perception, such as the amount of glossiness, object shape, etc. As a result, any study will only be able to sparsely sample the parameter space.

Acknowledgments

We would like to thank the participants of our study for their effort. We further thank Geomerics Ltd. for the use of their assets. This project was supported by the Technology Strategy Board (Q2047E).

References

AKENINE-MOELLER, T., CHAN, E., HEIDRICH, W., KAUTZ, J., KILGARD, M., AND STAMMINGER, M. 2004. Real-Time Shadowing Techniques. In *ACM SIGGRAPH 2004 Course Notes*.

ANNEN, T., DONG, Z., MERTENS, T., BEKAERT, P., SEIDEL, H., AND KAUTZ, J. 2008. Real-Time, All-Frequency Shadows in Dynamic Scenes. *ACM Transactions on Graphics (Proc. SIGGRAPH) 27*, 3, 34:1–34:8.

ARIKAN, O., FORSYTH, D. A., AND O'BRIEN, J. F. 2005. Fast and Detailed Approximate Global Illumination by Irradiance Decomposition. *ACM Transactions on Graphics (Proc. SIGGRAPH) 24*, 3, 1108–1114.

ARVO, J., TORRANCE, K., AND SMITS, B. 1994. A Framework for the Analysis of Error in Global Illumination Algorithms. In *SIGGRAPH '94*, 75–84.

BARTLESON, C. J. 1984. Chapter 8. In *Optical Radiation Measurements, Vol. 5 Visual Measurements*, C. J. Bartleson and F. Grum, Eds. Academic Press, Orlando, FL.

BUNNELL, M. 2005. Dynamic Ambient Occlusion and Indirect Lighting. In *GPU Gems 2*, M. Pharr, Ed. Addison Wesley, Mar., ch. 2, 223–233.

CHRISTENSEN, P., LAUR, D., FONG, J., WOOTEN, W., AND BATALI, D. 2003. Ray Differentials and Multiresolution Geometry Caching for Distribution Ray Tracing in Complex Scenes. In *Proc. Eurographics Symposium on Rendering*, 543–552.

COHEN, M., AND WALLACE, J. 1993. *Radiosity and Realistic Image Synthesis*. Academic Press Professional, Cambridge, MA.

COHEN-OR, D., CHRYSANTHOU, Y., SILVA, C., AND DURAND, F. 2003. A Survey of Visibility for Walkthrough Applications. *IEEE Transactions on Visualization and Computer Graphics 9*, 3, 412–431.

DACHSBACHER, C., AND STAMMINGER, M. 2005. Reflective Shadow Maps. In *Proc. Symposium on Interactive 3D Graphics and Games*, 203–213.

DACHSBACHER, C., AND STAMMINGER, M. 2006. Splatting Indirect Illumination. In *Proc. Symposium on Interactive 3D Graphics and Games*, 93–100.

DACHSBACHER, C., STAMMINGER, M., DRETTAKIS, G., AND DURAND, F. 2007. Implicit Visibility and Antiradiance for Interactive Global Illumination. *ACM Transactions on Graphics (Proc. SIGGRAPH) 26*, 3, 61:1–61:10.

DALY, S. 1993. The Visible Differences Predictor: An Algorithm for the Assessment of Image Fidelity. In *Digital Images and Human Vision*. MIT Press, 179–206.

DEBATTISTA, K., SUNDSTEDT, V., SANTOS, L. P., AND CHALMERS, A. 2005. Selective Component Based Rendering. In *Graphite 2005*.

DONG, Z., KAUTZ, J., THEOBALT, C., AND SEIDEL, H. 2007. Interactive Global Illumination Using Implicit Visibility. In *Pacific Graphics*, 77–86.

DRETTAKIS, G., BONNEEL, N., DACHSBACHER, C., LEFEBVRE, S., SCHWARZ, M., AND VIAUD-DELMON, I. 2007. An Interactive Perceptual Rendering Pipeline using Contrast and Spatial Masking. In *Proc. Eurographics Symposium on Rendering*, 297–308.

DURAND, F., HOLZSCHUCH, N., SOLER, C., CHAN, E., AND SILLION, F. X. 2005. A Frequency Analysis of Light Transport. *ACM Transactions on Graphics (Proc. SIGGRAPH) 24*, 3, 1115–1126.

DUTRÉ, P., BALA, K., AND BEKAERT, P. 2006. *Advanced Global Illumination*. AK Peters.

FERNANDO, R. 2005. Percentage-Closer Soft Shadows. In *ACM SIGGRAPH 2005 Sketches*, 35.

FERWERDA, J. A., SHIRLEY, P., PATTANAIK, S. N., AND GREENBERG, D. P. 1997. A Model of Visual Masking for Computer Graphics. In *SIGGRAPH '97*, 143–152.

GREEN, P., KAUTZ, J., AND DURAND, F. 2007. Efficient Reflectance and Visibility Approximations for Environment Map Rendering. In *Proc. Eurographics Symposium on Rendering*, 495–502.

GUILFORD, J. 1954. *Psychometric Methods*. McGraw-Hill, New York.

IWASAKI, K., DOBASHI, Y., YOSHIMOTO, F., AND NISHITA, T. 2007. Precomputed Radiance Transfer for Dynamic Scenes Taking into Account Light Interreflection. In *Proc. Eurographics Symposium on Rendering*, 35–44.

KAJIYA, J. 1986. The Rendering Equation. In *Computer Graphics (Proc. SIGGRAPH '86)*, 143–150.

KELLER, A. 1997. Instant Radiosity. In *SIGGRAPH '97*, 49–56.

KOZLOWSKI, O., AND KAUTZ, J. 2007. Is Accurate Occlusion of Glossy Reflections Necessary? In *Symposium on Applied Perception in Graphics and Visualization*, 91–98.

MEILGAARD, M., GAIL, C., AND CARR, B. 1991. *Sensory Evaluation Techniques*, 2nd ed. CRC Press.

MYSZKOWSKI, K., TAWARA, T., AKAMINE, H., AND SEIDEL, H. 2001. Perception-Guided Global Illumination Solution for Animation Rendering. In *SIGGRAPH '01*, 221–230.

PAN, M., WANG, R., LIU, X., PENG, Q., AND BAO, H. 2007. Precomputed Radiance Transfer Field for Rendering Interreflections in Dynamic Scenes. *Computer Graphics Forum (Proc. Eurographics) 26*, 3, 485–493.

RITSCHEL, T., GROSCH, T., KAUTZ, J., AND SEIDEL, H. 2008. Interactive Global Illumination Based on Coherent Surface Shadow Maps. In *Graphics Interface*, 185–192.

RITSCHEL, T., GROSCH, T., KIM, M. H., SEIDEL, H., DACHSBACHER, C., AND KAUTZ, J. 2008. Imperfect Shadow Maps for Efficient Computation of Indirect Il lumination. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia) 27*, 5, 129:1–129:8.

RITSCHEL, T., GROSCH, T., AND SEIDEL, H. 2009. Approximating Dynamic Global Illumination in Image Space. In *Proc. ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*, 75–82.

RUSHMEIER, H., PATTERSON, C., AND VEERASAMY, A. 1993. Geometric Simplification for Indirect Illumination Calculations. In *Graphics Interface*, 227–236.

SCHEFFÉ, H. 1952. An Analysis of Variance for Paired Comparisons. *Jour. Amer. Stat. Assoc. 47*, 381.

SILLION, F., AND DRETTAKIS, G. 1995. Feature-Based Control of Visibility Error: A Multi-Resolution Clustering Algorithm for Global Illumination. In *SIGGRAPH '95*, 145–152.

SILLION, F. 1995. A Unified Hierarchical Algorithm for Global Illumination with Scattering Volumes and Object Clusters. *IEEE Transactions on Visualization and Computer Graphics 1*, 3, 240–254.

SLOAN, P.-P., KAUTZ, J., AND SNYDER, J. 2002. Precomputed Radiance Transfer for Real-Time Rendering in Dynamic, Low-Frequency Lighting Environments. *ACM Transactions on Graphics (Proc. SIGGRAPH) 21*, 3, 527–536.

SLOAN, P.-P., GOVINDARAJU, N., NOWROUZEZAHRAI, D., AND SNYDER, J. 2007. Image-Based Proxy Accumulation for Real-Time Soft Global Illumination. In *Pacific Graphics*, 97–105.

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 Transactions on Graphics (Proc. SIGGRAPH) 23*, 3, 742–749.

TABELLION, E., AND LAMORLETTE, A. 2004. An Approximate Global Illumination System for Computer Generated Films. *ACM Transactions on Graphics (Proc. SIGGRAPH) 23*, 3, 469–476.

THURSTONE, L. L. 1927. A Law of Comparative Judgement. *Psychological Review 34*, 273–286.

TORGERSON, W. 1958. *Theory and Methods of Scaling*. Wiley, New York.

VANGORP, P., LAURIJSSEN, J., AND DUTRÉ, P. 2007. The Influence of Shape on the Perception of Material Reflectance. *ACM Transactions on Graphics (Proc. SIGGRAPH) 26*, 3, 77:1–77:9.

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 Transactions on Graphics 19*, 2, 122–161.

WALD, I., BENTHIN, C., AND SLUSALLEK, P. 2003. Interactive Global Illumination in Complex and Highly Occluded Environments. In *Proc. Eurographics Symposium on Rendering*, 74–81.

WALTER, B., FERNANDEZ, S., ARBREE, A., BALA, K., DONIKIAN, M., AND GREENBERG, D. P. 2005. Lightcuts: A Scalable Approach to Illumination. *ACM Transactions on Graphics (Proc. SIGGRAPH) 24*, 3, 1098–1107.

ZHUKOV, S., IONES, A., AND KRONIN, G. 1998. An Ambient Light Illumination Model. In *Proc. Eurographics Workshop on Rendering*, 45–56.