# Perceptual Illumination Components: A New Approach to Efficient, High Quality Global Illumination Rendering

William A. Stokes* James A. Ferwerda Bruce Walter Donald P. Greenberg
Program of Computer Graphics, Cornell University†

## Abstract

In this paper we introduce a new perceptual metric for efficient, high quality, global illumination rendering. The metric is based on a rendering-by-components framework in which the direct, and indirect diffuse, glossy, and specular light transport paths are separately computed and then composited to produce an image. The metric predicts the perceptual importances of the computationally expensive indirect illumination components with respect to image quality. To develop the metric we conducted a series of psychophysical experiments in which we measured and modeled the perceptual importances of the components. An important property of this new metric is that it predicts component importances from inexpensive estimates of the reflectance properties of a scene, and therefore adds negligible overhead to the rendering process. This perceptual metric should enable the development of an important new class of efficient global-illumination rendering systems that can intelligently allocate limited computational resources, to provide high quality images at interactive rates.

**CR Categories:** I.3.7 [Three-Dimensional Graphics and Realism];

**Keywords:** perception, global illumination, reflection components

## 1 Introduction

Global illumination effects, while necessary for improved realism, are often omitted because of their high cost. Traditionally, global illumination simulations have only been available from computationally intensive off-line systems or low quality interactive alternatives. The goal of our research is to enable the production of high quality global illumination renderings at interactive rates by approaching the rendering problem from a perceptual standpoint.

We propose that an efficient approach to providing global illumination effects is to split up the global illumination simulation into direct and indirect components. By quantifying the perceptual importances of these components, we can construct a perceptual quality metric that can be used to efficiently allocate computational resources and maximize image quality within system constraints.

To develop this new perceptual metric, we first conduct a series of experiments in which subjects rank the quality of global illumination images rendered using different combinations of direct and indirect components. Using the data provided by these experiments, we compute the perceptual importance of each of the components. We then construct a model that can predict these importance values by measuring the physical reflectance properties of the objects visible in a particular image. Finally we formulate the perceptual metric that can be used to predict the visual quality of different global illumination renderings from knowledge of the physical scene characteristics and the illumination components used in the rendering.

*e-mail:(wuz,jaf,bjw,dpg)@graphics.cornell.edu
†580 Rhodes Hall, Ithaca NY, 14853

Figure 1: Office and kitchen scenes

We believe this new metric will enable the development of an important new class of efficient global-illumination rendering systems that can intelligently allocate computational resources between the components of the global illumination simulation to provide high quality renderings at interactive rates.

## 2 Related Work

Perception has become an important research topic in computer graphics, especially in the area of rendering. Much of the work on perceptually-based rendering has focused on two goals: 1) developing perceptual metrics that can be used to establish stopping criteria for high quality rendering systems [Mitchell 1987; Gaddipatti et al. 1997; Gibson and Hubbold 1997; Prikryl and Purgathofer 1999; Walter et al. 1997; Neumann et al. 1998; Hedley et al. 1997; Tamstorf and Jensen 1997; Ramasubramanian et al. 1999; Walter et al. 2002] and 2) using perceptual metrics to optimally manage resource allocation for efficient rendering algorithms [Meyer and Liu 1992; Bolin and Meyer 1995; Bolin and Meyer 1998; Myszkowski 1998; Myszkowski et al. 2000; Myszkowski et al. 2001; Gibson and Hubbold 2000; Haber et al. 2001; Volevich et al. 1999; Yee et al. 2001; Tole et al. 2002; Dumont et al. 2003].

We share this community's interest in taking advantage of perception to improve the quality and efficiency of rendering algorithms, but our perspective is unique in that we will start by taking a rendering-by-components approach to the global illumination problem. The rendering-by-components strategy itself is not new [Shirley 1990], but to our knowledge there has been little work to develop perceptual metrics for component-based rendering.

## 3 Perceptual illumination components

In global illumination rendering, the radiance of a pixel in the rendered image can be determined by finding the intersection between the pixel's view ray and a surface, and calculating the sum of the energy emitted by the surface and the energy reflected in the direction of the view ray. This can be expressed as:

$$L^{out}(x) = L^{emission} + \int^{incoming}_{directions} L^{in} \cdot Brdf() \qquad (1)$$


Figure 2: Viewpoints tested in the office and kitchen scenes

### 3.1 Rendering by components

The starting point of our work is the insight that the global illumination simulation process can be successfully modeled by calculating light transport along direct and indirect illumination paths and combining the results. This allows the integral term in Equation 1 to be split up into direct and indirect components:

$$\int^{\substack{incoming\\directions}} L^{in} \cdot Brdf() = \begin{cases} \int L^{direct} \cdot Brdf() & + \\ \int L^{indirect} \cdot Brdf() \end{cases} \tag{2}$$

The indirect component can be further defined as the sum of the contributions from three pure indirect transport paths (indirect diffuse, indirect glossy, and indirect specular) and a fourth set of hybrid paths that account for the interactions between the pure paths. If the hybrid path interactions are negligible, the overall expression for the indirect component can be written as:

$$\int L^{indirect} \cdot Brdf() \approx \begin{cases} \int^{indirect} L_{diffuse} & + \\ \int^{indirect} L_{glossy} & + \\ \int^{indirect} L_{specular} \end{cases} \tag{3}$$

### 3.2 Perceptual illumination components

The key idea we are going to explore in this paper is the observation that the different direct and indirect illumination components are probably not of equal importance with respect to their contributions to the visual quality of global illumination renderings. Speculation is often made about the relative importance of the different components, but to our knowledge, there have not been any attempts to formalize these assertions, or to quantify the perceptual importances of the direct and indirect components in global illumination rendering. In the following section of the paper we describe a series of psychophysical experiments we conducted to measure the effects of the different illumination components on the perceived quality of a set of global illumination images.

## 4 Experiments

Our goal is to develop a new perceptual metric for efficient, high-quality, global illumination rendering that can predict the consequences for visual image quality of approximations to full global illumination simulations. Having adopted the rendering by components framework described in the previous section, our task is first to measure, and then to model how the different illumination components affect judgments of image quality. To accomplish this we have conducted a series of psychophysical experiments.

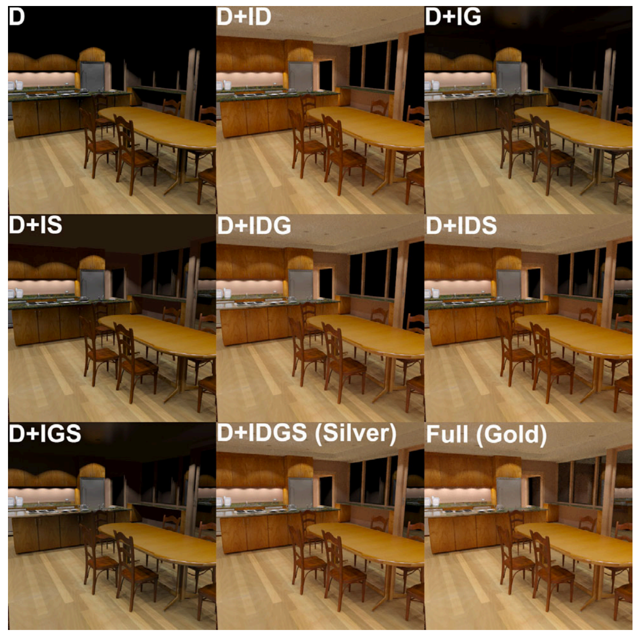


Figure 3: A complete set of composites plus the gold standard image for a single viewpoint

### 4.1 Stimuli

To measure how different illumination components affect perceived image quality, we needed to define and render a set of images to be used in the experiments. Ideally we would like these images to be representative of typical global illumination renderings so our results can be applied beyond the specific conditions of our experiments. For this reason we rejected the idea of rendering a simple “blocks world” environment and instead generated images from detailed models of two real scenes. Example images are shown in Figure 1. Together the office and kitchen scenes span a significant portion of the range of typical interior environments we encounter in the world, with the office scene having relatively uniform illumination and mostly matte surfaces, and the kitchen scene having more dramatic variations in illumination and a larger proportion of glossy and specular surfaces.

To construct test image sets for our experiments, we rendered six views of each scene using a custom component-based Monte Carlo path tracer. Figure 2 shows these views, which were chosen to be representative of viewpoints that a real observer might occupy, and to show a range of different objects and materials (i.e. avoiding views of blank walls, ceilings, desktops, etc.).

For each view, we generated separate renderings of the direct and indirect diffuse, glossy, and specular illumination components. We created eight composite images of each view by combining all permutations of the three indirect components with the baseline direct component. Figure 3 shows the set of composite images for one view of the kitchen scene. In addition, for each view we generated a ”gold standard” full global illumination rendering, so in total there were nine images in each of 12 test sets (6 viewpoints per scene x 2 scenes).

Each image was rendered at 512×512 pixels and printed at 4”×4” using a Kodak XLS 8600 dye sublimation printer. Each image was then mounted on 1/4” foamboard to allow easy manipulation.

Since the rendering-by-components framework is based on the idea that good approximations to full global illumination renderings can be achieved by combining separate simulations of the direct and pure indirect illumination components, one issue we were interested in exploring was whether the hybrid indirect paths neglected by the framework, have a significant effect on visual image quality. This was our rationale for including the gold standard in each test set. By comparing the quality ratings given to this full global illumination rendering and the pure component “silver standard” (direct plus pure indirect diffuse, glossy, and specular), we can determine the perceptual importance (or lack thereof) of the hybrid paths. From informal observations we expect that the contribution of the hybrid paths to visual image quality will be small since as is shown in Figure 3, there is very little visual difference between the gold and silver standard images.

In computing the indirect components, we needed to set a cutoff for path depth. Too severe a cutoff could result in large errors in radiance estimates, while too lax a cutoff would be inefficient. We chose to set the cutoff at four bounces. We found that images with greater path depths were visually indistinguishable from the four-bounce images.

Another issue we had to address was that the inclusion and omission of different components produced variations in contrast and hue in the composite images. To minimize these differences we equalized the images in each set using the techniques described in Appendix A.

## 4.2 Procedure

Figure 4: Ranking a set of images

To measure the relationships between the presence of different illumination components and visual image quality we ran a series of ranking experiments [Guilford 1954]. The procedure is illustrated in Figure 4. On each trial, a subject was given the set of composite images for one viewpoint. The subject was then asked to place the images in order from lowest to highest by perceived quality. We randomized the order in which each subject received the sets. The ”office” and ”kitchen” ranking experiments were conducted in two sessions on different days. Ten subjects participated in the experiments. Both expert (computer graphics graduate students) and non-expert (university graduates and undergraduates) participated. All were naive to the design and goals of the experiments, and all had normal or corrected to normal vision.

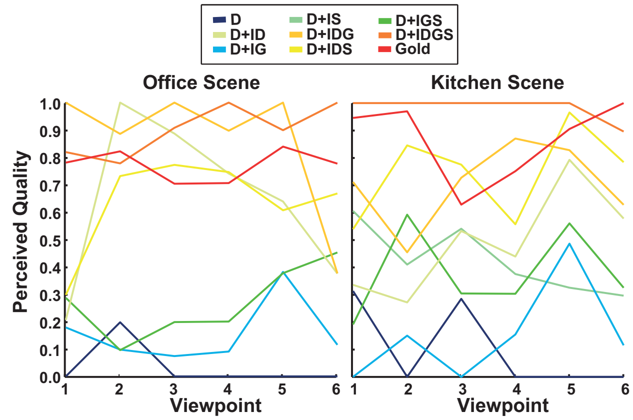


Figure 5: Experimental results: quality scores for the different composites

## 4.3 Results and Preliminary Analysis

The data generated by the experiments consisted of the rank orderings produced by each subject for each image set. Rank data is strictly ordinal by nature which presents significant limitations in terms of quantitative analysis and modeling [Guilford 1954]. For example, while average rankings may indicate that one image is consistently perceived to be of higher quality than another, rankings alone cannot specify *how much* higher in quality one image is than another. To place the images on an interval scale that allows us to quantify the contributions to quality of the different illumination components, we applied Thurstonian scaling techniques from visual psychophysics [Torgerson 1958].

In Thurstonian scaling, the variance in the rank position given to an image by different subjects is used to derive rank distributions for each image. The overlap in these distributions is taken as a measure of the psychological distance between the images with respect to visual quality. By determining these parameters for all the images in a set, a perceptually-linear interval scale that indicates the relative differences in quality between the images can be derived.

Using the technique described above, we derived perceived quality scales for each of the 12 image sets (6 viewpoints per scene × 2 scenes). The scales are summarized in Figure 5. There are several preliminary observations that we can make about this data.

- First, across all the viewpoints there is a fairly consistent ordering of the different composite images, with direct-only images judged to be lowest in quality, and the silver and gold standard images judged to be highest.

- Within the broad middle range, the presence of the indirect diffuse component appears to be an important factor, with images that include indirect diffuse generally being ranked higher than those that exclude it. Additionally, the indirect glossy and specular components appear to have smaller modulating effects within these larger trends.

- Finally, there are also clearly significant variations in the perceptual importances of the different components across viewpoints and scenes that will need to be accounted for by our metric.

## 4.4 Determining the perceptual importances of the illumination components

The scaling procedure we applied in the previous section allowed us to calculate perceived quality scores for the composite images in each set. The next step toward our goal of developing a metric that

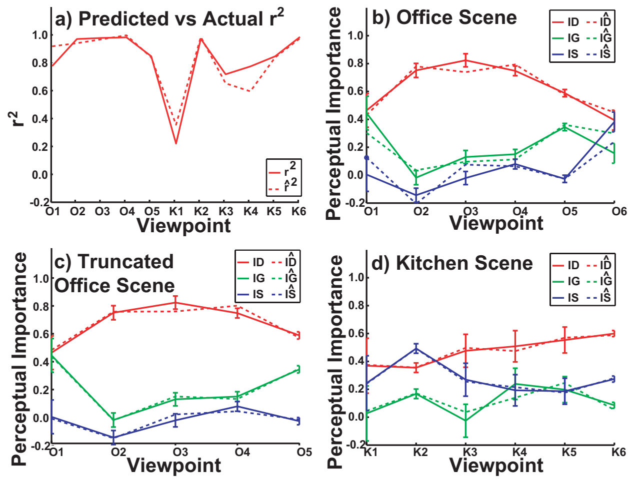

Figure 6: Computed and predicted $r^2$ values and perceptual importances for the office and kitchen scenes.

can predict these scores, is to relate them to the presence or absence of the different indirect illumination components, and derive measures of the perceptual importance of each component with respect to perceived quality.

We do this by performing a linear regression on the quality scores in which the score (Q) is modeled as the linear combination of the components. This model takes the form:

$$Q_i = a_i + a_{i,d} \cdot ID + a_{i,g} \cdot IG + a_{i,s} \cdot IS \tag{4}$$

where $ID$, $IG$, and $IS$ are binary variables that indicate the inclusion or exclusion of the indirect components, and $a_{i,d}$, $a_{i,g}$, and $a_{i,s}$ are weights returned by the regression that quantify the perceptual importance of each component in viewpoint $i$.

We performed separate regressions for each viewpoint in each scene. Figure 6 summarizes the perceptual importance factors returned by the regressions. We can observe several trends in this data.

- First, in all but one viewpoint, the indirect diffuse component ($ID$) has the highest perceptual importance. This squares with informal observations and comments made by the subjects that they did not like images containing large black regions (which would be the case for shadowed areas).
- Next, the regression showed that the indirect glossy ($IG$) and indirect specular ($IS$) components were of lower and roughly equal perceptual importance, with some variation across viewpoints. This result is also in line with observations and subject reports that the absence of reflections on surfaces that should have them (e.g. windows, monitors, etc.), is disturbing and reduces the visual quality of the image, but at the same time, surface reflections that mask underlying detail (such as the reflections on the marble countertop in the kitchen scene), are also sometimes judged to reduce image quality.
- Finally, an analysis showed that addition of the hybrid paths did not have a significant effect on perceived image quality (t-test: $t_{22} = 0.32, p = 0.79$), so we can conclude that under the conditions studied their perceptual importance is low. This result provides further support for the rendering-by-components framework.

At this point we have taken several important steps toward our goal of developing a perceptual metric for global illumination rendering-by-components. Through the experiments we have measured the relationships between perceived image quality and the presence or absence of the different illumination components. From our subsequent analysis we have derived values for the perceptual importance of each of the components with respect to the measured quality scores. What remains to be done is: 1) to find a method by which we can predict the perceptual importances of the components from some indicator of the physical properties of a scene, and 2) to formulate the metric so we can predict and/or specify the visual quality of a rendering by knowing the scene characteristics and the illumination components used. This work will be described in the following section.

## 5 Formulating the metric

From our experiments, we have derived measures of the perceptual importances of the different illumination components. To proceed with formulating a perceptual metric that can be used to guide the rendering process, we need indicators that specify how the physical characteristics of the scene are related to the magnitudes of the different illumination components. While these component magnitudes are ultimately a complex function of scene geometry, lighting, and material properties, we believe that useful indicators can be based on material properties alone. Our reasoning is as follows. In a scene without specularly reflecting materials, the illumination contribution of the indirect specular path would be null, and therefore the perceptual importance of that component should be zero. Similarly, in a scene with more matte than glossy materials, the indirect diffuse component will on average have greater influence on the scene radiances than the glossy component, and therefore its influence on the final appearance of the image, and its perceptual importance, are likely to be greater. Similar arguments in support of material-based indicator variables can be constructed for the other illumination components. Undoubtedly one can create scenes where these principles will break down, but we believe that this approach will be useful for a wide range of scenes.

One distinct advantage of employing material-based indicators, is that indicator values can be calculated online during computation of the direct illumination component. Since any global illumination rendering system will have to compute the direct component anyway, this approach adds negligible overhead to the rendering process, which should be of great advantage for implementing interactive systems.

### 5.1 Defining the indicators

We explored several formulations of the indicator variables.

- lobe counts: percentage of pixels with a diffuse/glossy/specular lobe
- image reflectivities: percentage of total image reflectance contributed by each component

$$X_{\{d,g,s\}} = \frac{\Sigma_i^{\#pixels} \rho_{i,\{d,g,s\}}}{\Sigma_i^{\#pixels} \rho_{i,d} + \rho_{i,g} + \rho_{i,s}} \tag{5}$$

- pixel reflectivities: average percentage of pixel reflectance contributed by each component

$$X_{\{d,g,s\}} = \frac{1}{\#pixels} \cdot \Sigma_i^{\#pixels} \frac{\rho_{i,\{d,g,s\}}}{\rho_{i,d} + \rho_{i,g} + \rho_{i,s}} \tag{6}$$

### 5.2 Indicator response functions

One final step we need to perform before we can construct and test different metrics, is to define response functions for the indicators that specify how sensitive the metric is to any particular indicator. We experimented with three types of response functions. The simplest was a linear function of the form,

$$f_1(x) = b_1 + b_2 \cdot x \tag{7}$$

Here sensitivity to the indicator is controlled by the constant $b_2$ which is fixed for all values of an indicator. The second response function we tested,

$$f_2(x) = b_1 - b_2^{\frac{1}{x+1}} \tag{8}$$

is also linear for low and moderate indicator values, but asymptotes to a ceiling to limit the influence of large values. Finally we also experimented with an s-shaped response function of the form,

$$f_3(x) = b^{\frac{-1}{x}} \text{ where we define } f_3(0) = 0 \tag{9}$$

that in addition to limiting the response to large values, also limits the influence of small values that might be contaminated by sampling noise.

### 5.3 Modeling perceptual importance

We are now ready to construct and test different metrics for predicting the visual quality of component-rendered global illumination images. Since our experiments have shown that the perceptual importance of the components varies across viewpoint and scene, we need to include these factors in our metrics

Using the indicator variables and response functions defined above, we can model the perceptual importance $a_{i,j}$ of the $j$th indirect illumination component for viewpoint $i$ as:

$$\hat{a}_{i,j} = w_{1,j} + w_{d,j} \cdot f_i(x_{i,d}) + w_{g,j} \cdot f_i(x_{i,g}) + w_{s,j} \cdot f_i(x_{i,s}) \tag{10}$$

In this model $x_{i,d}$, $x_{i,g}$, and $x_{i,s}$ are the diffuse, glossy, and specular indicators for viewpoint $i$, and $w_{1,j}$, $w_{d,j}$, $w_{g,j}$, and $w_{s,j}$ are weights determined by regressing on the perceptual importances.

### 5.4 Selecting the best model

Using this general model, we ran a series of regressions to try to find the combination of indicators and response functions that best predicted the perceptual importance factors we measured in the experiments. Of the three indicators we tested, the regressions showed that the second indicator: image reflectivities, provided the best prediction of the experimental results. The lobe counts indicator worked reasonably well for the specular component, but failed to predict the importances of the diffuse and glossy components. The performance of the pixel reflectivity indicator was similar to, but slightly worse than the image reflectivity indicator. Of the three response functions we tested, the regressions also indicated that the simple linear function yielded the best-fitting model.

While overall, this model was reasonably good at predicting the perceptual importances of the different components, its ability to predict the importance of the indirect glossy component was poor relative to the others. We believe this is because glossy materials have a wide range of appearances, from almost matte to almost mirror-like. We found that by defining separate high gloss and low gloss indicators we were able to produce a model that was better at predicting the perceptual importance of the indirect glossy component. We experimented with various split points and found that the best fit occurred when Ward materials with lobe widths $> 0.05$, and Phong materials with exponents $< 320$ were used to calculate the low gloss indicator ($x_{lg}$). Materials with lobe widths $\leq 0.05$ and exponents $\geq 320$ were used to calculate the high gloss indicator ($x_{hg}$) [Ward 1992; Blinn 1977].

The expanded form of the model is shown below. Since we have decided to use linear response functions for the indicators, the $f_i()$ terms in Equation 10 can be dropped, because the constants contained in these terms can be folded into the indicator weights $w_{i,j}$. Also it is not necessary to explicitly include the specular indicator variable in this formulation because $x_{i,d} + x_{i,g} + x_{i,s} = 1$ and therefore $x_{i,s}$ can be derived from the other two variables. This produces the simplified expression:

$$\hat{a}_{i,j} = w_{1,j} + w_{d,j} \cdot x_{i,d} + w_{hg,j} \cdot x_{i,hg} + w_{lg,j} \cdot x_{i,lg} \tag{11}$$

Finally, we found that we could improve the model's ability to predict the perceptual importances of the different components by adding a factor $r^2$ that is a statistical measure of the reliability of the data we are attempting to model. Since we are modeling perceptual importances that are themselves derived from regressions on the quality scores measured in the experiments, there will be higher $r^2$ values associated with viewpoints where these regressions provided a better fit to the quality scores, and higher $r^2$ values indicate that the reliability of the perceptual importance estimates are better for those views. Incorporating this factor into the model produces the expression:

$$\hat{a}_{i,j} = w_{1,j} + (1 - r^2) \cdot w_{2,j} + r^2 \cdot \left[ w_{d,j} \cdot x_{i,d} + w_{hg,j} \cdot x_{i,hg} + w_{lg,j} \cdot x_{i,lg} \right] \tag{12}$$

With this model we were able to predict more than 70% of the variance in the perceptual importances of the components we measured in the experiments.

Unfortunately, these $r^2$ values would not normally be available to a rendering algorithm, since they are the product of the analysis of the data from the ranking experiment. If we want to use this model in our final metric we need to be able to estimate $r^2$ from other, more accessible data. We found that the following expression:

$$\hat{r}_i^2 = c_1 + c_d \cdot x_{i,d} + c_{hg} \cdot x_{i,hg} + c_{lg} \cdot x_{i,lg} \tag{13}$$

is a good estimator of the $r^2$ values we obtained from our analyses of the experimental data.

By estimating these $r^2$ values and calculating the indicator values for each viewpoint, we were able to run regressions to determine the weights ($w_{i,j}$) in Equation 12 that best model the perceptual importances of the indirect illumination components for each scene. These parameters are tabulated in Appendix B.

We found that the model's predictive power improved slightly when we excluded viewpoint 6 of the office scene. In this particular viewpoint, the contrast adjustment procedure (described in Section 4.1 and Appendix A) added a significantly larger ambient term to the images than in the other cases, which we suspect distorted the experimental measures of the perceptual importances of the illumination components for this view, and made it an outlier in subsequent analysis and models.

Figure 6 illustrates the predictive abilities of the models. Figure 6a compares observed $r^2$ values(solid line) and those predicted (dashed line) by Equation 13. Figures 6b and 6c show the correspondence between the experimentally measured importances (solid lines) and the importances predicted by Equation 12 (dashed lines) for the office scene. In Figure 6c viewpoint 6 has been removed. Finally, Figure 6d shows the correspondence between the measured and predicted importances for the kitchen scene. It is clear from the close correspondence between the solid and dashed lines in these graphs, that the model is very good at predicting the perceptual importances we measured in our experiments, and therefore should perform well as the foundation of our perceptual quality metric.

### 5.5 Formulating the perceptual metric

Given the model for the perceptual importances of the different illumination components defined in Equation 12, we can now finally

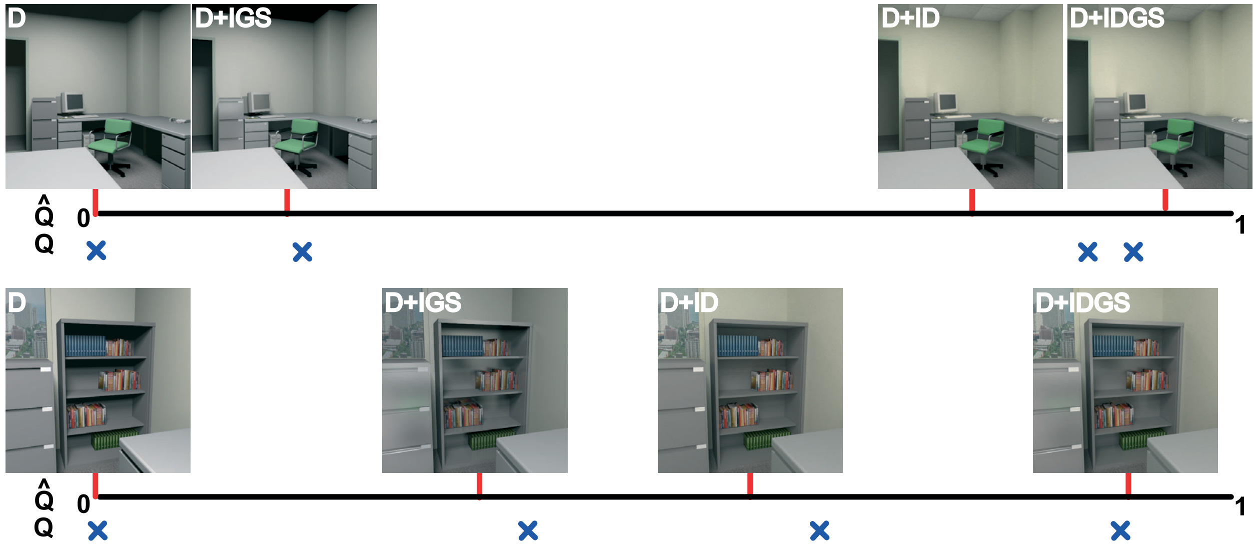

Figure 7: Applying the metric to predict the quality values of composites for two viewpoints in the office scene.

formulate our full perceptual quality metric as:

$$\hat{Q}_i = a_i + \hat{a}_{i,d} \cdot ID + \hat{a}_{i,g} \cdot IG + \hat{a}_{i,s} \cdot IS \tag{14}$$

where $\hat{Q}_i$ is the quality score predicted by the metric for an image of viewpoint $i$, (defined by some composite of illumination components), and the $\hat{a}_{i,j}$s are the perceptual importances of the different components.

# 6 Using the metric

In the previous section we formulated a new perceptual metric for predicting the visual quality of component-rendered global illumination images. In this section we will demonstrate the metric's predictive abilities and illustrate how the metric could be used in an interactive rendering system.

Figure 7 shows two sets of component-rendered images of the office scene. Each of the images is positioned with respect to the visual quality scale $\hat{Q}$ predicted by the new perceptual metric. There are several things to observe in this figure.

- First, is the correspondence between the quality values ($\hat{Q}$) predicted by the metric for these images, and the values measured in our experiment ($Q$), indicated by the X's below each scale. The similarity of the measured and predicted values is a confirmation that the metric is indeed capable of modeling the perceptual importances of the different illumination components and their contributions to perceived quality.

- Next, the metric's ability to quantify the quality of these images can also be confirmed by noticing the correspondence between the locations of the images along the quality scale and the similarities and differences in their visual appearances. Notice that on the upper scale the composites $D+ID$ and $D+IDGS$ placed near the high end of the scale are both similar to each other in visual quality, and are of distinctly higher quality than the composites $D$ and $D+IGS$ that the metric placed near the lower end of the scale. Similarly, the relatively equidistant positions of the images on the lower scale, accurately reflect the moderate increments in visual quality that can be observed by comparing adjacent images on the scale.

- Finally, the power of the metric to predict view-specific differences in the perceptual importances of the illumination components can be seen by comparing the central pair of images on each scale. On the upper scale, because of the material characteristics of the objects in this view, the metric correctly predicts that the indirect diffuse component makes a much greater contribution to quality than the glossy or specular components, so composite $D+ID$ is of substantially higher quality than composite $D+IGS$. However, for the view shown on the lower scale, where the material characteristics are substantially different, adding the the indirect diffuse component $D+ID$ does not produce as great an improvement in image quality, and conversely, adding the indirect glossy and specular components $D+IGS$ produces a relatively greater improvement than it did for the other view.

The three capabilities described above confirm that our new metric can accurately predict both the perceptual importances of different illumination components, and the resulting perceived quality values of component-rendered images.

## 6.1 A perceptually-based component renderer.

Figure 8: A hypothetical perceptual component rendering system.

Figure 8 illustrates how our new metric could be used in an interactive global illumination rendering system. Assuming there are not sufficient resources to do full global illumination rendering in real time, a user could specify whether to optimize system performance to a constant quality level, or a constant frame rate. Given these user preferences, for each viewpoint the system would: gather information about the materials visible in that view as part of an initial direct illumination pass; calculate indicator values and the perceptual importances of the indirect illumination components; dynamically allocate system resources to the computation of the different components (as shown by the pie chart on the lower right of Figure 8); and composite the components for display. Taking advantage of frame-to-frame coherence in estimating the importances and allocating resources would be likely to lead to even further improvements in performance.

## 7 Conclusions/Future Work

In this paper we have developed a new perceptual metric for efficient, high quality, global illumination rendering. Using the rendering-by-components framework, the metric can predict the perceptual importances of the indirect illumination components, and their contributions to the visual quality of the resulting image. We have demonstrated the predictive accuracy of the metric and have shown how it could be used in a global illumination rendering system.

An important aspect of this perceptual metric compared to others that have been developed, is that because it is based on simple measures of scene reflectances that can be gathered during calculation of the direct illumination component, it adds negligible overhead to the global illumination rendering process. This should make it attractive for use in interactive rendering systems.

Although we feel the current metric makes an important contribution, there is of course always much work to be done. First, we would like to explore improvements to the material-based indicators used to estimate the importance of the indirect glossy component, and find more meaningful ways to divide the wide range of surface reflectance properties that are currently defined as glossy. Work by [Pellacini et al. 2000] might provide a useful starting point. Second, the parameters used in the current metric are viewpoint independent, but must be tuned for a particular scene. As a next step, we would like generalize the metric so it can automatically adapt to different scene characteristics. Third, we would like to implement a rendering-by-components system that uses our metric. This will entail solving a number of important problems posed by the component approach including dynamic resource allocation, methods for gracefully approximating the components, and understanding the perceptual consequences of different approximation methods. Finally, there is always much more work to be done to increase our understanding of human perception to develop more sophisticated and effective perceptual metrics for global illumination rendering and other aspects of computer graphics.

## Acknowledgements

The authors would like to thank all our subjects for participating in the experiments, Jeremiah Fairbank for building the kitchen model, and Houston H. Stokes for statistical advice. We also acknowledge the support of NSF grant ACI-0205438 and Intel Corporation for generously providing funding and equipment.

## References

BLINN, J. F. 1977. Models of light reflection for computer synthesized pictures. vol. 11, 192–198.

BOLIN, M. R., AND MEYER, G. W. 1995. A frequency based ray tracer. In *Proceedings of SIGGRAPH 95*, Computer Graphics Proceedings, Annual Conference Series, 409–418.

BOLIN, M. R., AND MEYER, G. W. 1998. A perceptually based adaptive sampling algorithm. In *Proceedings of SIGGRAPH 98*, Computer Graphics Proceedings, Annual Conference Series, 299–310.

DUMONT, R., PELLACINI, F., AND FERWERDA, J. A. 2003. Perceptually-driven decision theory for interactive realistic rendering. *ACM Transactions on Graphics 22*, 2 (Apr.), 152–181.

GADDIPATTI, A., MACHIRAJU, R., AND YAGEL, R. 1997. Steering image generation with wavelet based perceptual metric. *Computer Graphics Forum 16*, 3 (Aug.), 241–252.

GIBSON, S., AND HUBBOLD, R. J. 1997. Perceptually-driven radiosity. *Computer Graphics Forum 16*, 2, 129–141.

GIBSON, S., AND HUBBOLD, R. J. 2000. A perceptually-driven parallel algorithm for efficient radiosity simulation. *IEEE Transactions on Visualization and Computer Graphics 6*, 3 (July), 220–235.

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

HABER, J., MYSZKOWSKI, K., YAMAUCHI, H., AND SEIDEL, H.-P. 2001. Perceptually guided corrective splatting. *Computer Graphics Forum 20*, 3, 142–152.

HEDLEY, D., WORRALL, A., AND PADDON, D. 1997. Selective culling of discontinuity lines. In *Eurographics Rendering Workshop 1997*, 69–80.

MEYER, G. W., AND LIU, A. 1992. Color spatial acuity control of a screen subdivision image synthesis algorithm. In *Proceedings of SPIE 92*, Human Vision, Visual Processing, and Digital Display, 387–399.

MITCHELL, D. P. 1987. Generating antialiased images at low sampling densities. In *Computer Graphics (Proceedings of SIGGRAPH 87)*, vol. 21, 65–72.

MYSZKOWSKI, K., ROKITA, P., AND TAWARA, T. 2000. Perception-based fast rendering and antialiasing of walkthrough sequences. *IEEE Transactions on Visualization and Computer Graphics 6*, 4, 360–379.

MYSZKOWSKI, K., TAWARA, T., AKAMINE, H., AND SEIDEL, H.-P. 2001. Perception-guided global illumination solution for animation rendering. In *Proceedings of ACM SIGGRAPH 2001*, Computer Graphics Proceedings, Annual Conference Series, 221–230.

MYSZKOWSKI, K. 1998. The visible differences predictor: Applications to global illumination problems. In *Eurographics Rendering Workshop 1998*, 223–236.

NEUMANN, L. O., MATKOVIC, K., AND PURGATHOFER, W. 1998. Perception based color image difference. *Computer Graphics Forum 17*, 3, 233–242.

PELI, E. 1991. Contrast in complex images. *J. Opt. Soc. Am. A 7*, 10, 2032–2040.

PELLACINI, F., FERWERDA, J. A., AND GREENBERG, D. P. 2000. Toward a psychophysically-based light reflection model for image synthesis. In *Proceedings of ACM SIGGRAPH 2000*, Computer Graphics Proceedings, Annual Conference Series, 55–64.

POYNTON, C. A. 1996. *A Technical Introduction to Digital Video*. Wiley, New York.

PRIKRYL, J., AND PURGATHOFER, W. 1999. Perceptually-driven termination for stochastic radiosity. In *Seventh International Conference in Central Europe on Computer Graphics and Visualization (Winter School on Computer Graphics)*.

RAMASUBRAMANIAN, M., PATTANAIK, S. N., AND GREENBERG, D. P. 1999. A perceptually based physical error metric for realistic image synthesis. In *Proceedings of SIGGRAPH 99*, Computer Graphics Proceedings, Annual Conference Series, 73–82.

SHIRLEY, P. 1990. A ray tracing method for illumination calculation in diffuse-specular scenes. In *Graphics Interface '90*, 205–212.

TAMSTORF, R., AND JENSEN, H. W. 1997. Adaptive sampling and bias estimation in path tracing. In *Eurographics Rendering Workshop 1997*, 285–296.

TOLE, P., PELLACINI, F., WALTER, B., AND GREENBERG, D. P. 2002. Interactive global illumination in dynamic scenes. *ACM Transactions on Graphics 21*, 3 (July), 537–546.

TORGERSON, W. S. 1958. *Theory and methods of scaling*. Wiley, New York.

VOLEVICH, V., MYSZKOWSKI, K., KHODULEV, A., AND KOPYLOV, E. A. 1999. Perceptually-informed progressive global illumination solution. Tech. rep., University of Aizu, Aizu Wakamatsu, Japan, Feb.

WALTER, B., HUBBARD, P. M., SHIRLEY, P., AND GREENBERG, D. F. 1997. Global illumination using local linear density estimation. *ACM Transactions on Graphics 16*, 3 (July), 217–259.

WALTER, B., PATTANAIK, S. N., AND GREENBERG, D. P. 2002. Using perceptual texture masking for efficient image synthesis. *Computer Graphics Forum 21*, 3, 393–399.

WARD, G. J. 1992. Measuring and modeling anisotropic reflection. In *Computer Graphics (Proceedings of SIGGRAPH 92)*, vol. 26, 265–272.

YEE, H., PATTANAIK, S., AND GREENBERG, D. P. 2001. Spatiotemporal sensitivity and visual attention for efficient rendering of dynamic environments. *ACM Transactions on Graphics 20*, 1 (Jan.), 39–65.

## Appendix A: Equalizing contrast and hue

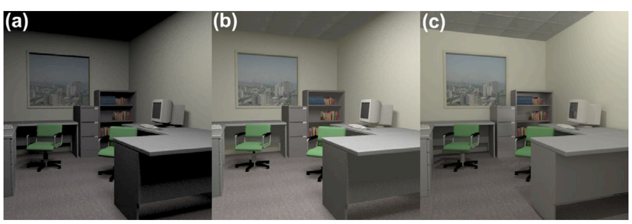


Figure 9: (a) Direct illumination image. (b) Contrast equalized version of (a). (c) Full global illumination solution. Notice the contrast differences between (a) and (c) and the similarity of (b) and (c).

One consequence of the rendering-by-components and compositing approach, is that including or excluding components from a composite can introduce contrast differences with respect to the other composites in a set. Figure 9 illustrates this problem. Image (a) only contains direct illumination. Notice that image contrast is greater than in (c), the full global illumination solution. To minimize this artifact we had to develop a contrast equalization technique.

In any set of composites, the full global illumination solution (gold standard) has the lowest contrast. By adding an "ambient" image to the other composites in the set, it is possible to minimize contrast differences. To accomplish this we first calculated an albedo image that represented the diffuse colors of the surfaces visible from that viewpoint. For each composite image + albedo image combination, we used a binary search method to find a scaling factor for the albedo image that minimized contrast differences with respect to the full global illumination solution.

Image contrasts were calculated using a scale-dependent local contrast measure that [Peli 1991] has shown is well correlated with perceived contrast. First, we derived a luminance image from RGB values by the technique described in [Poynton 1996]. Next we constructed an image pyramid and calculated pixel contrast $C_{i,j}$ at different spatial scales using the equation:

$$C_{i,j} = \frac{L'_{i,j} - L_{i,j}}{L'_{i,j}} \tag{15}$$

where $L_{i,j}$ is the luminance of pixel $i$ at pyramid level $j$ and $L'_{i,j}$ is the luminance of the pixel in a bi-linearly upsampled version of the next highest pyramid level. To combine the contrast measures at different spatial scales, it is necessary to weight the contrasts at each level by the factor $W_j = \frac{B_w \cdot B_h}{C_w \cdot C_h}$ where $B_w$ and $B_h$ are the image dimensions at the base level and $C_w$ and $C_h$ are the dimensions at the current mip-map level $j$. Thus a summary contrast measure of the image can be defined as:

$$C = \sum_{i=0}^{m} \sum_{j=0}^{n} W_j \cdot C_{i,j} \tag{16}$$

The effectiveness of this contrast equalization technique can be seen in Figure 9 by comparing image (b), the contrast equalized version of (a) with image (c), the full global illumination solution.


Figure 10: (a) Direct illumination, (b) contrast equalized version, (c) contrast equalized using hue shifted albedo image, and (d), full global illumination solution. Notice the hue differences between (a,b) and (d), and the similarity of (c) and (d).

Another side effect of the compositing process was a hue shift between images containing different combinations of indirect illumination components. To minimize this artifact we hue shifted the albedo image so its average hue was the same as the average hue of the indirect components. This was done before contrast equalization so any luminance changes in the albedo image due to the hue shifting would be compensated for during calculation of the albedo scale factors. Figure 10 shows an example of this hue shift artifact and how our procedure minimizes this artifact while equalizing contrast with respect to the full global illumination solution.

## Appendix B: Parameters used in the metric

### Predicted $r^2$ values:

| $c_1$ | $c_d$ | $c_{hg}$ | $c_{lg}$ |
|---|---|---|---|
| 6.07 | -5.26 | -8.52 | -4.74 |

### Predicted importance factors: office scene

| | $w_1$ | $w_2$ | $w_d$ | $w_{hg}$ | $w_{lg}$ |
|---|---|---|---|---|---|
| Indirect Diffuse: | -3097.21 | 2566.22 | 3206.90 | 0 | 2911.57 |
| Indirect Glossy: | 5344.02 | -4439.80 | -5529.39 | 0 | -5026.76 |
| Indirect Specular: | 1730.64 | -1446.58 | -1789.08 | 0 | -1630.97 |

### Predicted importance factors: kitchen scene

| | $w_1$ | $w_2$ | $w_d$ | $w_{hg}$ | $w_{lg}$ |
|---|---|---|---|---|---|
| Indirect Diffuse: | -0.509 | 0.80 | 1.08 | 0.49 | 1.70 |
| Indirect Glossy: | -12.12 | 11.20 | 13.28 | 21.64 | 6.37 |
| Indirect Specular: | 5.45 | -4.97 | -5.54 | -7.66 | -3.65 |