

Visual Equivalence: Towards a New Standard for Image Fidelity

Ganesh Ramanarayanan*

James Ferwerda†

Bruce Walter†

Kavita Bala*

Program of Computer Graphics and Computer Science Department
Cornell University




Figure 1: Demonstration of visual equivalence. (a) shows an illumination map, and (b) shows a warped version of the same map. The scenes shown in (c) and (d) were rendered with the original and warped maps respectively. Notice that in both images the dragon has the same appearance; i.e. the images are *visually equivalent*. (e) shows the output of a visible differences predictor (VDP), indicating the reflection patterns in the two images are visibly different.

Abstract

Efficient, realistic rendering of complex scenes is one of the grand challenges in computer graphics. Perceptually based rendering addresses this challenge by taking advantage of the limits of human vision. However, existing methods, based on predicting visible image differences, are too conservative because some kinds of image differences do not matter to human observers. In this paper, we introduce the concept of *visual equivalence*, a new standard for image fidelity in graphics. Images are visually equivalent if they convey the same impressions of scene appearance, even if they are visibly different. To understand this phenomenon, we conduct a series of experiments that explore how object geometry, material, and illumination interact to provide information about appearance, and we characterize how two kinds of transformations on illumination maps (blurring and warping) affect these appearance attributes. We then derive *visual equivalence predictors* (VEPs): metrics for predicting when images rendered with transformed illumination maps will be visually equivalent to images rendered with reference maps. We also run a confirmatory study to validate the effectiveness of these VEPs for general scenes. Finally, we show how VEPs can be used to improve the efficiency of two rendering algorithms: Lightcuts and precomputed radiance transfer. This work represents some promising first steps towards developing perceptual metrics based on higher order aspects of visual coding.

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

Keywords: perception, human visual system, appearance

*{graman,kb}@cs.cornell.edu

†{jaf,bjw}@graphics.cornell.edu

1 Introduction

Realistic and efficient image synthesis, particularly the rendering of complex scenes, is one of the grand challenges in computer graphics. To address this challenge, researchers have developed perceptually based methods that take advantage of the limits of human vision, performing less computation while assuring visual fidelity. Existing methods typically incorporate visible differences predictors (VDPs) [Daly 1993] to determine when an approximate solution will be visually indistinguishable from a reference solution. While this approach is well-founded, it is too conservative because some kinds of visible image differences simply do not matter to human observers.

In this paper, we introduce the concept of *visual equivalence*, a new standard for image fidelity in graphics. Two images are visually equivalent if they convey the same impressions of scene appearance, even if they are visibly different in a side-by-side comparison. Figure 1 illustrates the phenomenon.

As a first step in understanding visual equivalence, we focus on the problem of rendering objects with illumination maps. We conduct a series of psychophysical experiments that explore how object geometries, materials, and illumination interact to provide information about appearance, and we quantify how two sets of transformations on illumination maps, blurring and warping, affect these appearance attributes. Based on these experiments, we develop *visual equivalence predictors* (VEPs), metrics predicting when images rendered with the transformed maps will convey the same appearance as images rendered with reference maps. We then run a confirmatory study to validate that these VEPs generalize, accurately reflecting observers' judgments about visual equivalence for new geometries, materials, and illuminations. Finally, we demonstrate how VEPs can be used to improve the efficiency of two rendering algorithms: Lightcuts [Walter et al. 2005] and wavelet compression for precomputed radiance transfer [Ng et al. 2004].

The research described in this paper represents some promising first steps toward developing new perceptual metrics for realistic image synthesis based on higher order aspects of human visual coding. We believe that this work can serve as the foundation of advanced perceptually based image synthesis algorithms that will dramatically lower computational costs while assuring the fidelity of the resulting images.

2 Prior Work

2.1 Psychophysics of Appearance

Why do things look as they do? This question posed in 1935 by psychologist Kurt Koffka [Koffka 1935] has been a central theme in the field of perception psychology. The way human vision untangles the patterns of light in the retinal images to perceive the shapes, materials, and illumination of objects is the subject of a vast research literature (surveyed in [Palmer 1999]).

Shape perception: The central problem in shape perception is how the visual system recovers the three-dimensional shapes of objects from two-dimensional retinal images. Many potential sources of information have been identified, including shading, shadows, perspective, reflections, texture, motion, and occlusion [Gibson 1979; Rock 1983]. Modern work has focused on trying to measure the efficacy of different sources [Todd and Mingolla 1983; Cavanagh and Leclerc 1989; Cutting and Millard 1984; Fleming et al. 2004], and build computational models of how vision might combine different sources [Knill and Richards 1996] to recover shape. Researchers have also discovered how the limitations of fixed viewpoints make it mathematically impossible to recover the exact shape of objects under certain conditions [Belhumeur et al. 1997].

Material Perception: Historically, there has been relatively little research on material perception [Beck 1972]. However, there is now a growing literature that spans computer graphics and vision research. Active research areas include studies of surface lightness and color [Gilchrist et al. 1999; Adelson 2000; Brainard and Maloney 2004], gloss and specularities [Blake and Bülthoff 1990; Ferwerda et al. 2001; Westlund and Meyer 2001; Fleming et al. 2003; Fleming et al. 2004; Hartung and Kersten 2002; Todd et al. 2004; Xiao and Brainard 2006], translucency and subsurface scattering [Robilotto et al. 2002; Fleming and Bülthoff 2005], and surface texture [Dana et al. 1999; te Pas and Pont 2005; Ho et al. 2006].

Illumination Perception: While accurate rendering of illumination is a central issue in graphics, in vision research it is often regarded as something to be discounted to achieve constancy [Gilchrist et al. 1999; Adelson 2000; Brainard and Maloney 2004; Wandell 1993]. Recently, greater focus has been placed on understanding perception of illumination per se. Many researchers [Todd and Mingolla 1983; Koenderink et al. 2003; te Pas and Pont 2005; Khang et al. 2006] have looked at the ability to estimate the directionality and complexity of illumination fields. Dror et al. [2004] have characterized natural illumination statistics, Pont and Koenderink [2004] have studied surface illuminance flow, and Ostrovsky et al. [2005] have studied tolerance for illumination inconsistencies.

2.2 Perceptually based rendering

Due to the computational expense of realistic image synthesis, perceptually based rendering has been an active research area over the past fifteen years. The central goal of perceptually based rendering is to improve the efficiency of realistic image synthesis by taking the limits of human vision into account. A popular approach has been to use visible differences predictors (VDPs) [Daly 1993], which can determine whether images rendered using an approximate light transport simulation will be indistinguishable from a reference solution. Myszkowski [2002] provides a recent review.

While VDPs offer a principled approach to perceptually based rendering, the overhead of evaluating VDPs and the conservative threshold metrics they incorporate has limited the performance gains that have been achieved. Recently a number of researchers have been looking for ways to improve the efficiency of perceptually based rendering by taking advantage of other aspects of visual




Figure 2: Material and Geometry. An object’s material and geometry affect how it reflects the incoming illumination field. (a) A glossy material blurs the illumination map, producing a blurry reflection (see bottom array) of the checkerboard environment. (b) A shiny, bumpy object disorders the illumination map, as shown in the bottom array.

processing, including perceptual salience [Yee et al. 2001; Anson et al. 2006], attention and change blindness [Cater et al. 2002; Cater et al. 2003; Chalmers et al. 2003], natural image statistics [Reinhard et al. 2004], perception of illumination components [Stokes et al. 2004; Debattista et al. 2005], and task importance [Ferwerda and Pellacini 2003; Sundstedt et al. 2004].

Researchers have also started to develop perceptual metrics to optimize and control the appearance of the geometric and material models used in rendering. There is a large body of work using VDPs, saliency, attention, and related metrics for geometric level-of-detail schemes [Luebke and Hallen 2001; Williams et al. 2003; Reddy 2001; Brown et al. 2003; Howlett et al. 2004; Parkhurst and Niebur 2004; Lee et al. 2005; Qu and Meyer 2006]. Cleju and Saupé [2006] and Watson et al. [2001] have explored higher level metrics for evaluating model fidelity. Rogowitz and Rushmeier [2001] have characterized the consequences of substituting texture for geometry. Others have developed psychophysically based models for describing and controlling material appearance [Pellacini et al. 2000; Westlund and Meyer 2001; Khan et al. 2006; Meseth et al. 2006].

2.3 The Need for a New Approach

While there are large literatures on the psychophysics of appearance and perceptually based rendering, efficient rendering of complex scenes remains a difficult, unsolved problem. One reason is that perceptual metrics currently used in graphics (such as VDPs) only consider the earliest levels of visual coding, and are therefore too conservative with respect to the kinds of approximations that can be applied in the rendering process. We believe that by understanding the visual coding of higher level scene properties, such as shapes, materials, and lighting, it should be possible to develop powerful new classes of *appearance-preserving* rendering algorithms that will realistically represent the way scenes look, while aggressively reducing computational costs. In the following sections we investigate the concept of visual equivalence, which will serve as the foundation of our new approach.

3 Visual Equivalence

Scene appearance is determined by the complex interactions of a scene’s geometry and materials with the incoming illumination field. Our goal is to understand when these interactions result in images that are visually equivalent.

Because the space of interactions between geometry, material, and illumination is so large, in this work we focus on the effects of illumination. To do this, we need to (a) pick an encoding of the incoming illumination field and (b) define classes of transformations we can apply to the illumination field so that we can measure when visual equivalence occurs. We chose illumination maps as representations of illumination for two reasons. First, HDR environment

maps of natural illumination are widely used and readily available. Second, recent research has demonstrated the importance of natural illumination (which can be encoded in such maps) for tasks such as material appearance perception [Fleming et al. 2003]. Now, we must define classes of transformations we can apply to these maps.

3.1 Illumination Transformations

The space of transformations on illumination maps is potentially infinite - how do we decide which ones are worth studying? To pick a set of transformations, we drew inspiration from how object material and geometry interact with the illumination field. All opaque objects can be thought of as mirrors in the sense that they reflect the light they receive from their surroundings. Some objects, such as flat metal surfaces, are good mirrors in that they produce regular, high contrast, coherent reflections of the environment. Other objects are bad mirrors, though bad mirrors can be bad in different ways.

Material. A glossy material reflects a *blurred* version of the incoming illumination field; the degree of the blur depends on the width of the material’s specular lobe. Figure 2-(a) shows how material blurs a 1D checkerboard environment map. This effect is exploited in interactive rendering using prefiltered environment maps [Cabral et al. 1999; Kautz et al. 2000; Ramamoorthi and Hanrahan 2002].

Geometry. The geometry of an object can also affect how it reflects the incoming illumination field. Figure 2-(b) shows how bumpy geometry disorders access into the 1D checkerboard environment map. Characterizing this access pattern for arbitrary geometry is complex. However, when we look at a small surface patch, we can think of how it accesses the incoming illumination as a *warp* on the illumination map. The extent of the warp depends on the local curvature of the surface.

Choice of Transformations. Based on these insights, we decided to study the following two transformations on illumination maps: *blurs and warps*.

Object Complexity. The above discussion, in addition to inspiring particular illumination transformations, also illustrates the importance of characterizing the effects of illumination on a variety of geometries and materials. One would expect that more illumination transformations are permissible on complex objects than on simple ones. For this reason, we decided to study a high-to-low-gloss range of materials, and a smooth-to-bumpy range of geometries. Our exact set of objects is shown in Figure 3 and discussed in Section 4.1.

3.2 Criteria for Visual Equivalence

We want to find illumination transformations that preserve scene appearance. When do two images convey the “same” scene appearance? VDPs only predict when observers detect contrast differences between a pair of images. For visual equivalence, we need to use higher-level criteria. Given an object, and its associated geometry and material properties, the images produced by the reference and transformed illumination maps are *visually equivalent* if:

- The object’s shape and material are judged to be the same in both images, and
- In a side-by-side comparison, a person is unable to correctly identify which object has been rendered with the reference map.

Note that this definition specifically does not require that the images be indistinguishable in the VDP sense.




Figure 3: The geometries and materials used in the experiments. Specific parameters are given in the text.

4 Experiments

To understand the effects of illumination transformations on visual equivalence, we conducted a series of psychophysical experiments measuring how illumination changes interact with object geometry and material properties to affect appearance. Specifically, we designed studies to answer the following questions: 1) when do changes in illumination produce visible differences in images; and 2) when do changes in illumination produce renderings that are visually equivalent despite these differences? In the following sections, we describe our stimuli and experimental procedure.

4.1 Stimuli

First, we created a set of images that would allow us to systematically explore interactions between geometry, materials, and illumination. Figure 3 shows the test objects in our stimulus set. The scene model consists of a bumpy ball-like test object on a brick patio flanked by two pairs of children’s blocks (see Figure 5). The following paragraphs describe the object, scene, and rendering parameters we used to generate the images.

Geometry: We created four object geometries (G0-G3), shown in the rows of Figure 3. G0 is a sphere (approximated by a triangle mesh with 164k vertices and normal interpolation) while G1-G3 are modified spheres with bumps of uniform amplitude and increasing spatial frequency. G1-G3 were created by applying a scaled cube of Perlin noise [Perlin 2002] to the sphere mesh with cube sizes of {8, 4, 2}, which was judged to produce objects of roughly uniformly increasing “bumpiness”. We chose these geometries because: 1) their functional definitions should facilitate formal analysis of the effects of geometry on appearance; 2) there is a precedent in the shape perception literature for similar geometries [Todd and Mingolla 1983; Todd et al. 2004]; and 3) there are recent studies that point to the importance of mesoscale surface variations in the perception of material and illumination properties [te Pas and Pont 2005].

Materials: The columns of Figure 3 show the materials used in the experiments, which represent brushed aluminums with different de-



Figure 4: The illumination transformations used in the experiments, and the resulting renderings on object G1/M0.

degrees of microscale surface roughness. Materials were defined using an isotropic version of the Ward light reflection model [Ward 1992]. The Ward model uses three parameters to describe surface reflectance properties: ρ_d (diffuse reflectance), ρ_s (specular reflectance), and α (specular lobe width). For all materials, $\rho_d = 0.15$ and $\rho_s = 0.19$. α values for M0 through M3 were set to $\{0.01, 0.06, 0.11, 0.16\}$ respectively. We chose these parameters to: 1) span a significant range of high-to-low-gloss reflectance (producing visually salient reflected images); and 2) represent perceptually equal changes in gloss appearance [Pellacini et al. 2000].

Illumination: Recent studies have demonstrated the importance of real-world illumination for the accurate perception of shape and material properties [Fleming et al. 2004; Fleming and Bühlhoff 2005]. We used Debevec’s GROVE (UC Berkeley Eucalyptus Grove) HDR environment map [Debevec and Malik 1997]. We chose this map in particular because Fleming et al. [2003] found that it allowed subjects to most accurately discriminate material properties. Starting with the original GROVE map, we first generated a *reference map* that incorporated the other components of our scene (i.e., the brick patio and the children’s blocks, which provided high contrast, colored features one would expect to see in object reflections). We then generated two sets of transformed maps, via blurring and warping:

- **Blurs:** We convolved the reference map with progressively larger Gaussian blurring kernels. The sizes of these kernels roughly corresponded to Ward α values of $\{0.01, 0.035, 0.06, 0.085, 0.11\}$. The top row of Figure 4 shows a section of each blurred illumination map, and below it the result of applying the blurred illumination to an object with geometry G1 and material M0.
- **Warps:** We wanted to create illumination distortions similar to those seen in the reflection on a bumpy surface. We first created bumpy spheres by applying scaled Perlin noise to a sphere, as described by the geometry, using scaling factors of $\{\sqrt{2}, 1, \sqrt{1/2}, 1/2, \sqrt{1/8}\}$ respectively. Then for each direction in a warped map, we found the corresponding surface point as seen from the center of the bumpy object and used its surface normal to lookup in the reference map. Lastly, we renormalized the warped maps to have the same overall energy as the reference. The bottom row of Figure 4 shows a section of each warped illumination map, and below it and the result of apply-




Figure 5: Browser-based interface used in the experiment. Shown: Illumination task, G1/M0 object, rendered with warp5 illumination map on left, and reference map on right.

ing the warped illumination to an object with geometry G1 and material M0.

Rendering and display: The images were rendered at 484x342 using a Monte Carlo path tracer. The test object was illuminated by either the reference map, a blurred map, or a warped map, while the rest of the scene was always illuminated by the original GROVE map. Overall, 176 images were rendered for the stimulus set (4 geometries x 4 materials x 11 illuminations (reference, 5 blurs, 5 warps)). For display, the images were tone mapped using a global sigmoid [Tumblin and Rushmeier 1993] that was tuned to the characteristics of the LCD display (Dell 2000FP, 20” diagonal, 1600x1200 resolution, sRGB color space, max luminance 200 cd/m^2 , 60:1 dynamic range, gamma 2.2). The images were viewed under dim office lighting conditions. At a nominal 24” viewing distance each image subtended 11.6 degrees of visual angle and each test object subtended 7 degrees.

4.2 Procedure

The images in the stimulus set were presented to subjects in pairs using the browser-based interface shown in Figure 5. The pairs always showed objects with the same geometry and material properties. The only difference was that one was lit by the reference map and the other by a transformed (blurred or warped) map. In some cases (Tasks 1 and 4) a third reference image was displayed above the test pair.

Separate experiments were conducted for the “blurred” and “warped” map image sets. An experiment consisted of four related tasks:

1. **Image differences task:** In this task, subjects were shown a reference image, and a pair of test images, all of the same object. The reference image and one of the test images were rendered with the reference map; these images were identical. The other test image was rendered with one of the transformed maps. Subjects were asked: “Which test image is the same as the reference image?”. The purpose of this task was to de-

termine when the illumination transformations produce images that are visibly different (in the VDP sense) from the reference.

2. **Shape task:** In this task, subjects were shown two images of the same object. One object was rendered with the reference map the other was rendered with a transformed map. Subjects were asked: “**Are the left and right test objects the same shape?**”. The purpose of this task was to determine if the transformed maps produce illusory changes in the apparent shapes of the objects.
3. **Material task:** In this task, subjects were shown two images of the same object. One object was rendered with the reference map the other was lit with a transformed map. Subjects were asked: “**Are the left and right test objects made of the same material?**”. The purpose of this task was to determine if the transformed maps produce illusory changes in the apparent material properties of the test objects.
4. **Illumination task:** In this task, subjects were shown a reference image, and a pair of test images. The reference image showed an ideal mirror sphere rendered with the reference map. The test images showed identical objects, one rendered with the reference map and one rendered with a transformed map. Subjects were asked: “**Which test object is lit the same as the reference object?**”. The purpose of this task was to determine if subjects can use surface reflection patterns to detect differences in scene illumination.

Each subject performed the image differences task first. To reduce the experiment size, if the subject incorrectly identified which image was the same as the reference, then they could not distinguish the reference from the test image, and that pair was not shown in later tasks. The shape, material, and illumination tasks were then delivered in random order. Within each task both the overall order of presentation and left/right positions of the images were randomized across trials.

Subjects entered their responses with a keyboard and mouse using the buttons shown in the interface. Response times were recorded, but the trials themselves were open-ended and subjects could take as much time as they needed. On average, subjects took approximately 45 minutes to complete all four tasks.

Overall, 30 subjects participated in the experiments (15 each in the blurred map and warped map conditions). The subjects were university students, staff, and faculty (ages 20 to 50). Many had technical backgrounds, but none in computer graphics. All were naïve to the design and purpose of the experiments, and all had normal or corrected-to-normal vision.

5 Results of the Experiments

Our results are summarized in Figure 6. Recall that we are interested in answering two questions: 1) when do changes in illumination produce images that are visibly different; and 2) when do different illuminations produce images that are visually equivalent (i.e. same scene appearance) despite these differences? Task 1 addressed the image differences question and Tasks 2, 3, and 4 addressed the appearance question.

5.1 Task 1 - Image Differences

Task 1 tested if subjects could see any differences between images rendered with the transformed maps and the reference map. Performance on the task was measured using a 75% threshold (2AFC) [Gescheider 1997]. In Figure 6, the cases where images were indistinguishable are indicated with solid green circles; in all

other cases, visible differences were detected. The top row summarizes the results for the blurred maps, with the blur increasing across the panels from left to right. Overall, as the blur increases the illumination transform becomes more detectable, as expected. Within this broad trend, visibility of differences is affected by the object geometry and material. In general, more blur is tolerated as the material gloss decreases (i.e., moving up from M0 to M3 within each panel). There is also a small effect of geometry; the blur is less detectable as surface bumpiness increases (moving right from G0 to G3 within each panel). Although, as with any psychophysical experiment, there is some noise in the data, these results are well understood within characterizations of object-illumination interactions as a filtering operation [Ramamoorthi and Hanrahan 2002].

The results for the warped maps are shown in the bottom row of Figure 6. Again, overall detectability of differences between the test and reference images increases with the warp level. However, unlike the blur results, which showed a gradual increase in differences across the range, here the effect is dramatic. While some images rendered with warp1 are indistinguishable from the reference, virtually all images rendered with warp2 and above were judged to be visibly different (very few solid green circles). This reflects the reports of subjects who said that it was relatively easy to see this kind of image difference because it produces distinctly different reflection patterns in the surfaces. The detectability of this kind of transformation on the illumination map is relatively unaffected by the object geometry, and only modestly affected by material properties.

We also ran an open source VDP¹ [Mantiuk et al. 2005] on our test pairs to predict the visibility of differences and it generally agreed with the judgments of our subjects.

The Task 1 results show that: 1) the transformations we have made on the illumination maps often produce images that are visibly different from reference images, and 2) the blur and warp transformations appear to be different in kind, where the detectability of blur increases gradually and is affected by object properties, but warping is reliably detected at all but the lowest warp level, and detection is largely independent of object properties.

5.2 Tasks 2, 3, and 4 - Changes in Appearance

The focus of Task 1 was on simple detectability of pixel or region-wise image differences without regard for image content (similar to measures produced by VDPs). The focus of Tasks 2, 3, and 4 was on how the blur and warp transformations affect the appearance attributes of the test objects (shape, material, and lighting respectively). Since distortion of any appearance attribute indicates a failure of visual equivalence, in reporting the results we have combined the data from the three tasks. Specifically, if for Task 2 or Task 3 subjects consistently reported that one of the object’s attributes looked different (50% threshold for yes/no design), or if for Task 4 the correctly rendered object was picked at least 75% of the time (2AFC threshold), then the transformed illumination map did not preserve appearance for that object.

We refer again to the top row of Figure 6 for the results of Tasks 2, 3, and 4 for the blurred maps. Red squares indicate where subjects saw changes in appearance for that particular object/illumination combination; green circles (solid or with a dot) indicate where appearance was preserved. As with Task 1, there appears to be an interaction between illumination blur and object material properties. At the lowest blur level (blur1), with one exception, subjects saw no differences in appearance between objects rendered with

¹<http://www.mpi-inf.mpg.de/resources/hdr/vdp>

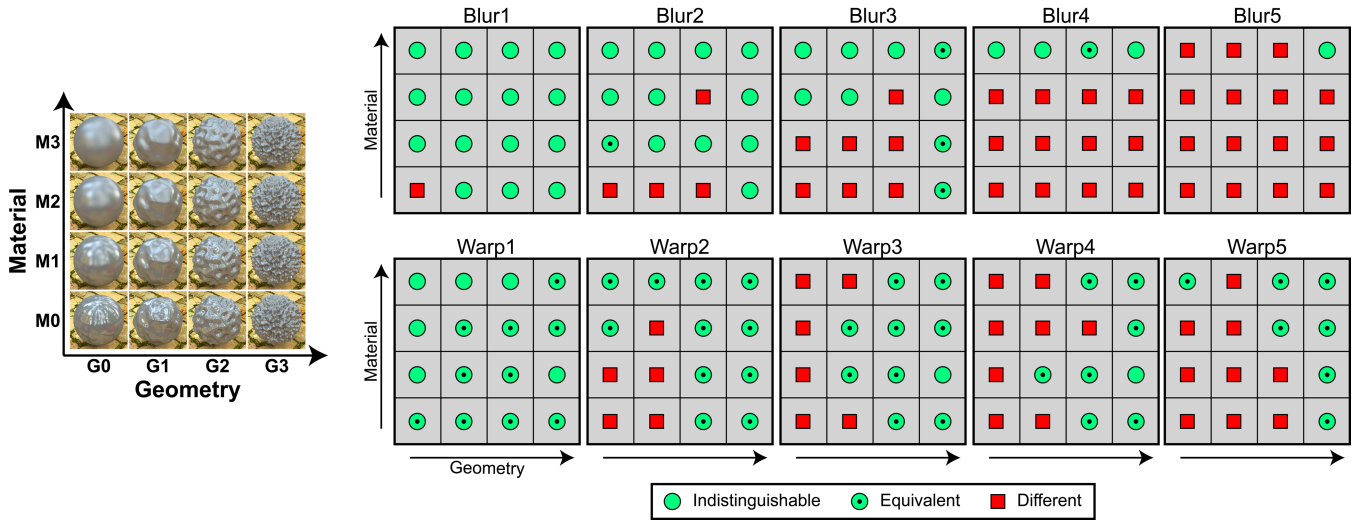


Figure 6: Summary of results for the experiments. The top and bottom rows show results for the blur and warp studies, respectively. Green circles mark cases where objects rendered with transformed illumination maps had the same appearance as objects rendered with reference maps. Red squares show cases where object appearance was different from from the reference. Among the green circles (same appearance), there is a further distinction between the cases where images were indistinguishable (solid circles) and cases where the images were different but conveyed the same appearance (circles with a dot). Visual equivalence is represented by this latter set.

the blurred and reference maps, but as blur increases (blur2-blur5) more objects were affected.

The results of Tasks 2, 3, and 4 for the warped maps are shown in the bottom row of Figure 6. Similar to the blur results, smaller warps have less effect on appearance than greater warps; however, even at the higher warp levels (warp3-warp5) there are significant regions where subjects said the warp and reference rendered objects appeared the same.

5.3 Significance of the Results

So what have we learned from the experiments?

1. First, for the range of objects we tested, it is often possible to significantly transform the illumination maps used in rendering without affecting the object’s appearance (green circles);
2. Second, it is not just that the images produced by the reference and transformed maps are visually indistinguishable from each other, because in many cases the images were clearly different; rather, it is that images could be *visually equivalent* (i.e. same scene appearance) despite these differences (green circles vs. green circles with dots);
3. Third, while the effects of the blur transformation are largely predicted by visible differences (Task 1), the effects of the warp transformation depend much more strongly on properties of appearance (Tasks 2, 3, and 4), which are frequently preserved even for large warps (many more green circles with dots for the warp transformation).

6 Visual Equivalence Predictors

To use the findings of the experiments, we would like metrics that can predict when images rendered with transformed illumination maps will be visually equivalent with reference renderings. In this section, we derive these visual equivalence predictors (VEPs) based on our experiment data. We then describe geometry and material measures to apply our VEP metrics to objects other than those tested in the experiment.

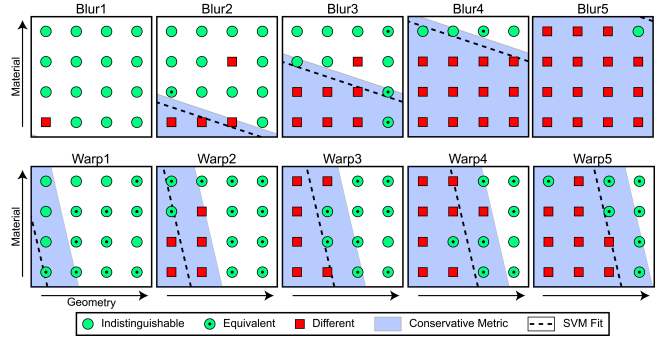


Figure 7: Visual equivalence predictors (VEPs) for illumination transformations. The dashed lines represent the original SVM plane fit, and the blue regions represent a more conservative partitioning. For blur, the planes have a shallow slope, indicating the main interaction is with material properties. For warp, the planes have a steep slope, indicating the main interaction is with geometry.

6.1 Defining the Metrics

For each geometry, material, and illumination transformation included in the experiment, we know whether the resulting rendering was visually equivalent to a reference rendering or not (Figure 6). To define metrics, we need classifiers that accurately separate the results into “good” (green) and “bad” (red) for the blur and warp transformations.

Support Vector Machines (SVM) [Vapnik 1995] are a well-known technique used in machine learning for classification and regression. A support vector machine finds an optimal separating hyper-plane between members and non-members of a given class in an abstract space. Depending on the kernel functions used, SVMs can function as linear or non-linear classifiers. SVMs are good at simultaneously minimizing classification error and maximizing the geometric margin.

We used the popular SVM software *SVMlight*² [Joachims 1999] to

²<http://svmlight.joachims.org>

find a linear classification of our data. Figure 7 shows the resulting plane fits, which classify the results with 90% accuracy. Given a point (g, m, i) in the geometry, material and illumination transformation axes, the equations of the original SVM separating planes for blur B and warp W are:

$$B: 0.181g + 0.546m - 0.728i + 1.177 = 0 \quad (1)$$

$$W: 0.772g + 0.178m - 0.456i + 0.500 = 0 \quad (2)$$

These planes are indicated by the dashed lines in Figure 7. Notice how most reds (different appearance) lie on one side of these lines, and most greens (same appearance) lie on the other side. While this is the best fit, we may want to use more conservative metrics since the misclassification penalty for our data is asymmetric (it is worse to misclassify a configuration as visually equivalent when it actually is not). Conservative plane fits are given by:

$$BC: 0.181g + 0.546m - 0.728i + 1.027 = 0 \quad (3)$$

$$WC: 0.772g + 0.128m - 0.456i - 0.299 = 0 \quad (4)$$

which classify the blue regions in Figure 7 as having a different appearance. Our conservative metrics have the property that no instance of different appearance is misclassified, except for three cases in the blur experiment where the difference is so subtle it cannot be detected by a VDP. We attribute this to minute changes in the brightness of certain highlights that users learned to spot during the course of the experiment.

We considered the possibility of obtaining a non-linear fit for our data. However, computing higher dimensional fits for datasets of our size (80 points) is not encouraged with SVMs, so we leave this as future work.

6.2 Applying the Metrics to Novel Scenes

To apply our metrics to scenes with arbitrary geometries and materials, we need to be able to map these properties into their corresponding positions in the configuration space defined by the test objects in the experiment. In particular, given a new object, we would like to derive the closest corresponding g and m values that describe it. The metrics will then predict levels and types of illumination transformations that result in visual equivalence for that object.

Geometry. Given a general geometric model, what g value should be associated with it? Many characterizations of geometry are possible [Funkhouser and Shilane 2006; DeCarlo et al. 2003]. We use a relatively simple measure that proved effective: we compute the local standard deviation of the surface normal per degree of visual angle squared (42x42 pixels), and associate the center pixel of that region with the test object (G0-G3) whose average local standard deviation is most similar. Note this measure takes the viewpoint into consideration, so the same object may correspond to different values depending on viewing distance and angle. The average local standard deviations were $\{0.157, 0.245, 0.408, 0.592\}$ for G0 to G3, respectively, given the camera view used to render the test stimuli. Figure 8 shows a false colored image of the bunny and dragon with automatically computed g values. This can be treated as a pixel-wise measure, or it can be averaged to associate a single g value with a new object.

Material. Any Ward material can be linearly mapped into our space based on its c and d values using the equations from [Pellacini et al. 2000], provided that $c \leq 0.221$, the contrast gloss for our test material, which is on the upper end of the scale of glossiness. In fact, the less you can see reflected in the object, the more aggressively our metrics can be applied. Mappings for other material models are left as future work.




Figure 8: Geometry measure for bunny and dragon.

#	Scene	Predicted/ExpResult	#	Scene	Predicted/ExpResult
0	Bunny M1 Blur2 Campus	● / ●	7	G1 M0 Warp4 Galileo	■ / ■
1	Bunny M1 Blur4 Campus	■ / ■	8	Dragon M2 Warp3 Grove	● / ●
2	Dragon M1 Blur2 Grove	● / ●	9	Dragon M2 Warp5 Grove	● / ●
3	Dragon M1 Blur4 Grove	■ / ■	10	Bunny M1 Warp2 StPeters	● / ●
4	Dragon M2 Blur2 Galileo	● / ●	11	Bunny M1 Warp5 StPeters	■ / ■
5	Dragon M2 Blur4 Galileo	■ / ■	12	Dragon M1 Warp2 Galileo	● / ●
6	G1 M0 Warp1 Galileo	● / ●	13	Dragon M1 Warp5 Galileo	■ / ■

Figure 10: Results of our follow-up experiment showing the predictive power of our metrics. We achieve 93% accuracy, with our only mismatch being a conservative classification. The visually equivalent image pairs in Figures 1 and 9 were all validated in this experiment.

Illumination. Once we compute the g and m value for an arbitrary model, we can use our metrics to check which illumination transformations lie on the allowed (visually equivalent) side of the separating plane. We can then pick the most advantageous one for the particular application. Blur values directly correspond to a filter width or downsampling that can be used on an illumination map. For warps, we computed the average angular displacement of pixels in the warped maps ($\{0.074, 0.123, 0.154, 0.241, 0.310\}$ radians for warp1 to warp5, respectively) and use this to characterize arbitrary warps.

6.3 Metric Generalization and Validation

To test how well our metrics work on objects and illuminations not studied in the experiment, we created new scenes using more of Debevec’s HDR environment maps (GALILEO, STPETERS, and CAMPUS). We also used the Stanford bunny and dragon geometries, which have g values of roughly 1.5 and 2.5 respectively, based on pixel averaging of the geometry measure described earlier. Figure 9 shows some examples, using the conservative metrics of Equations 3 and 4. In Figure 9-(a), we generalize across illumination only, replacing GROVE with GALILEO but still using our test geometries. In Figure 9-(b), we generalize across geometry only, replacing our test geometries with the dragon model. In Figure 9-(c), we generalize across both geometry and illumination with the bunny in STPETERS. Figure 1 shows full generalization across geometry, illumination, and material for warp2.

To confirm the predictions of our metrics with the new scenes, we performed a follow-up experiment with 10 additional subjects, identical to the experiment described in Section 4.2. The results for these scenes are summarized in Figure 10. Of 14 test stimuli, 7 predicted different and 7 predicted equivalent by our metrics, we achieved 93% accuracy relative to the subjects’ judgments, with our only mismatch being a conservative classification.




Figure 9: Predicting visual equivalence for novel scenes. (a) generalization across illumination; (b) generalization across geometry; (c) generalization across both geometry and illumination.



Figure 11: Application of visual equivalence predictors to the Lightcuts rendering algorithm. Using a warp3 based error threshold reduced the rendering time in half, while the warp5 based threshold was correctly predicted as not acceptable.

7 Using Visual Equivalence Predictors

Visual equivalence has a wide variety of potential applications which we have only started to explore. Using new metrics of appearance, it should be possible to create algorithms that reduce image generation costs and enable greater data compression. We have applied our illumination metrics to two such applications with promising early results: high quality rendering using Lightcuts [Walter et al. 2005], and wavelet compression of illumination for precomputed radiance transfer (PRT) style applications [Ng et al. 2004].

Lightcuts is a scalable algorithm for computing the illumination from complex sources including area lights, HDR environment

maps, and indirect illumination. It converts these sources into many point lights, clusters them into a light hierarchy, and adaptively chooses an appropriate clustering for each point to be illuminated. Ordinarily it refines/splits light clusters until a 2% error criteria is met. We modified Lightcuts to not refine light clusters whose solid angle falls below a fixed threshold based on the angular deviations derived for the warp metric. The results of this modification of the illumination approximation resemble the effects seen in the warp experiment. For the dragon model with the M2 material lit by the GROVE HDR map, our metrics correctly predict that thresholds based on warp3 produce a visibly different but equivalent image, but thresholds based on warp5 produce objectionable appearance changes (see Figure 11). Rendering using the warp3 based thresholds reduced the rendering time from 143 to 77 seconds.

Precomputed Radiance Transfer techniques use specialized and highly compressed representations of material and illumination properties to quickly compute images of complex objects under complex illumination, often at interactive rates. For example, [Ng et al. 2004] use Haar wavelets and lossy compression where only a small number of the strongest wavelets coefficients are kept. Using fewer coefficients reduces memory and computational cost, but also introduces blurring and reduces the contrast in the illumination, which can affect material perception. We have started to experiment with pre-warping of the illumination before the lossy wavelet compression such that the compressed image retains more of the con-

trast of the original image. Currently we swap nearby pixels in the environment map to better align features with the block boundaries of the Haar basis while ensuring that pixels remain within some maximum deviation of their original location, based on the angular deviations from the warp metric. For the dragon model with material M1 and the GROVE map, we have been able to produce PRT-style rendering using an environment map encoded with 2× fewer coefficients (50 instead of 100) while achieving equivalent fidelity.

Discussion: So far, using the warp metric, we have demonstrated gains for fixed-camera renderings of objects under distant lighting (the applications of the blur metric are more straightforward). These proof-of-concept examples, while promising, are only a first step to leveraging the concept of visual equivalence. We believe much larger gains are possible by incorporating visual equivalence into the basic design of new algorithms. Shifting from restrictive early vision models to higher level vision metrics will allow much greater freedom to find more efficient algorithms and solutions.

8 Conclusions and Future Work

In this paper, we have introduced the concept of visual equivalence as a new approach for perceptually based rendering. In a series of psychophysical experiments, we have characterized conditions under which two classes of transformations on illumination maps (blurring and warping) yield images that are visually equivalent to reference solutions. On the basis of these experiments, we have derived metrics for predicting visual equivalence for these transformations, and in a follow-up experiment we have validated that the predictive power of the metrics generalizes across different geometries, materials, and illumination maps. Finally, we have shown how these metrics can be applied to two existing rendering algorithms to increase efficiency while maintaining image fidelity.

While these initial findings are encouraging, there is still much work to be done. In the short term, we would like to do further testing of our metrics and analyses of our data using models of psychometric functions. We would also like to identify other classes of transformations on illumination maps that yield visual equivalence, such as rotation and blocking. One promising possibility might be to do energy-based segmentation and then apply different transformations to the different segments. Another would be to investigate how different kinds of disordering interact with the statistical properties of natural illumination.

In the long term, we would like to use our visual equivalence predictors to develop advanced perceptually based rendering algorithms. This will require an understanding of how exactly our metrics apply in the context of global illumination and multiple objects. We could also investigate the effect of moving cameras, which motivates a study of visual equivalence for animation and video. Eventually, we would like to extend the concept of visual equivalence to include transformations on geometry and material properties and develop comprehensive metrics that could serve as the foundation for dramatic advances in realistic and efficient perceptually based image synthesis.

Acknowledgments

We thank Thorsten Joachims and Filip Radlinski for useful discussions regarding SVM, the members of the Cornell Program of Computer Graphics for valuable suggestions and feedback, and the experiment subjects for their time and patience. Finally, we thank the reviewers for their constructive and helpful comments. This work was supported by NSF CAREER 0644175, NSF 0539996, NSF ITR 0113310, NSF ITR 0205438, NIH NEI/SBIR EY017238-

01, the Intel Corporation, the Eastman-Kodak Corporation, and an Affinito-Stewart Award.

References

- ADELSON, E. H. 2000. *Lightness perception and lightness illusions*, 2nd ed. The MIT Press, Cambridge, MA, 339–351.
- ANSON, O., SUNDESTEDT, V., GUTIERREZ, D., AND CHALMERS, A. 2006. Efficient selective rendering of participating media. In *APGV '06: Proceedings of the 3rd symposium on Applied perception in graphics and visualization*, 135–142.
- BECK, J. 1972. *Surface Color Perception*. Cornell University Press, Ithaca, NY.
- BELHUMEUR, P. N., KRIEGMAN, D. J., AND YUILLE, A. L. 1997. The bas-relief ambiguity. *CVPR '97: IEEE conference on Computer vision and pattern recognition 00*, 1060–1066.
- BLAKE, A., AND BÜLTHOFF, H. H. 1990. Does the brain know the physics of specular reflection? *Nature* 394, 165–168.
- BRAINARD, D. H., AND MALONEY, L. T. 2004. Perception of color and material properties in complex scenes. *Journal of Vision* 4, 9, 2–4.
- BROWN, R., COOPER, L., AND PHAM, B. 2003. Visual attention-based polygon level of detail management. In *GRAPHITE '03: Proceedings of the 1st international conference on Computer graphics and interactive techniques in Australasia and South East Asia*, 55–62.
- CABRAL, B., OLANO, M., AND NEMEC, P. 1999. Reflection space image based rendering. In *SIGGRAPH '99: Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 165–170.
- CATER, K., CHALMERS, A., AND LEDDA, P. 2002. Selective quality rendering by exploiting human inattentive blindness: looking but not seeing. In *VRST '02: Proceedings of the ACM symposium on Virtual reality software and technology*, 17–24.
- CATER, K., CHALMERS, A., AND DALTON, C. 2003. Varying rendering fidelity by exploiting human change blindness. In *GRAPHITE '03: Proceedings of the 1st international conference on Computer graphics and interactive techniques in Australasia and South East Asia*, 39–46.
- CAVANAGH, P., AND LECLERC, Y. G. 1989. Shape from shadows. *Journal of Experimental Psychology: Human Perception and Performance* 15, 1, 3–27.
- CHALMERS, A., CATER, K., AND MAFLIOLI, D. 2003. Visual attention models for producing high fidelity graphics efficiently. In *SCCG '03: Proceedings of the 19th spring conference on Computer graphics*, ACM Press, New York, NY, USA, 39–45.
- CLEJU, I., AND SAUPE, D. 2006. Evaluation of supra-threshold perceptual metrics for 3D models. In *APGV '06: Proceedings of the 3rd symposium on Applied perception in graphics and visualization*, 41–44.
- CUTTING, J. E., AND MILLARD, R. T. 1984. Three gradients and the perception of flat and curved surfaces. *Journal of Experimental Psychology: General* 113, 2, 198–216.
- DALY, S. 1993. The visible differences predictor: an algorithm for the assessment of image fidelity. In *Digital Images and Human Vision*, A. B. Watson, Ed. MIT Press, 179–206.

- DANA, K. J., VAN GINNEKEN, B., NAYAR, S. K., AND KOENDERINK, J. J. 1999. Reflectance and texture of real world surfaces. *ACM Transactions on Graphics* 18, 1, 1–34.
- DEBATTISTA, K., SUNDSTEDT, V., SANTOS, L. P., AND CHALMERS, A. 2005. Selective component-based rendering. In *GRAPHITE '05: Proceedings of the 3rd international conference on Computer graphics and interactive techniques in Australasia and South East Asia*, 13–22.
- DEBEVEC, P. E., AND MALIK, J. 1997. Recovering high dynamic range radiance maps from photographs. In *SIGGRAPH '97*, 369–378.
- DECARLO, D., FINKELSTEIN, A., RUSINKIEWICZ, S., AND SANTELLA, A. 2003. Suggestive contours for conveying shape. In *ACM Transactions on Graphics*.
- DROR, R. O., WILLSKY, A. S., AND ADELSON, E. H. 2004. Statistical characterization of real-world illumination. *Journal of Vision* 4, 9, 821–837.
- FERWERDA, J., AND PELLACINI, F. 2003. Functional difference predictors (FDPs): Measuring meaningful image differences. *Asilomar Conference on Signals, Systems, and Computers*, 1388–1392.
- FERWERDA, J. A., PELLACINI, F., AND GREENBERG, D. P. 2001. A psychophysically-based model of surface gloss perception. In *Proceedings of the SPIE: Human Vision and Electronic Imaging VI*, vol. 4299, 291–301.
- FLEMING, R. W., AND BÜLTHOFF, H. H. 2005. Low-level image cues in the perception of translucent materials. *ACM Transactions on Applied Perception* 2, 3, 346–382.
- FLEMING, R. W., DROR, R. O., AND ADELSON, E. H. 2003. Real-world illumination and the perception of surface reflectance properties. *Journal of Vision* 3, 5, 347–368.
- FLEMING, R. W., TORRALBA, A., AND ADELSON, E. H. 2004. Specular reflections and the perception of shape. *Journal of Vision* 4, 9, 798–820.
- FUNKHOUSER, T., AND SHILANE, P. 2006. Partial matching of 3D shapes with priority-driven search. In *Fourth Eurographics Symposium on Geometry Processing*, 131–142.
- GESCHEIDER, G. 1997. *Psychophysics: The Fundamentals*. Lawrence Erlbaum Associates, Mahwah, NJ.
- GIBSON, J. 1979. *The Ecological Approach to Visual Perception*. Houghton Mifflin, Boston, MA.
- GILCHRIST, A., KOSSYFIDIS, C., BONATO, F., AGOSTINI, T., CATALIOTTI, J., LI, X., SPEHAR, B., ANNAN, V., AND ECONOMOU, E. 1999. An anchoring theory of lightness perception. *Psychological Review* 106, 4, 795–834.
- HARTUNG, B., AND KERSTEN, D. 2002. Distinguishing shiny from matte. *Journal of Vision* 2, 7, 551–551.
- HO, Y.-X., LANDY, M. S., AND MALONEY, L. T. 2006. How direction of illumination affects visually perceived surface roughness. *Journal of Vision* 6, 5, 634–648.
- HOWLETT, S., HAMIL, J., AND O’SULLIVAN, C. 2004. An experimental approach to predicting saliency for simplified polygonal models. In *APGV '04: Proceedings of the 1st Symposium on Applied perception in graphics and visualization*, 57–64.
- JOACHIMS, T. 1999. Making large-scale support vector machine learning practical. In *Advances in kernel methods: support vector learning*. MIT Press, Cambridge, MA, 169–184.
- KAUTZ, J., VÁZQUEZ, P.-P., HEIDRICH, W., AND SEIDEL, H.-P. 2000. A unified approach to prefiltered environment maps. In *Proceedings of the Eurographics Workshop on Rendering Techniques 2000*, Springer-Verlag, London, UK, 185–196.
- KHAN, E. A., REINHARD, E., FLEMING, R. W., AND BÜLTHOFF, H. H. 2006. Image-based material editing. *ACM Transactions on Graphics* 25, 3, 654–663.
- KHANG, B.-G., KOENDERINK, J. J., AND KAPPERS, A. M. L. 2006. Perception of illumination direction in images of 3-D convex objects: Influence of surface materials and light fields. *Perception* 35, 5, 625–645.
- KNILL, D. C., AND RICHARDS, W. 1996. *Perception as Bayesian Inference*. Cambridge University Press, New York, NY.
- KOENDERINK, J. J., VAN DOORN, A. J., KAPPERS, A. M. L., TE PAS, S. F., AND PONT, S. C. 2003. Illumination direction from texture shading. *Journal of the Optical Society of America A* 20, 6, 987–995.
- KOFFKA, K. 1935. *Principles of Gestalt Psychology*. Harcourt, Brace and World, New York, NY.
- LEE, C. H., VARSHNEY, A., AND JACOBS, D. W. 2005. Mesh saliency. *ACM Trans. Graph.* 24, 3, 659–666.
- LUEBKE, D. P., AND HALLEN, B. 2001. Perceptually driven simplification for interactive rendering. In *Proceedings of the 12th Eurographics Workshop on Rendering Techniques*, 223–234.
- MANTIUK, R., DALY, S., MYSZKOWSKI, K., AND SEIDEL, H.-P. 2005. Predicting visible differences in high dynamic range images - model and its calibration. In *Proceedings of the SPIE: Human Vision and Electronic Imaging X*, vol. 5666, 204–214.
- MESETH, J., MÜLLER, G., KLEIN, R., RÖDER, F., AND ARNOLD, M. 2006. Verification of rendering quality from measured BTFs. In *APGV '06: Proceedings of the 3rd symposium on Applied perception in graphics and visualization*, 127–134.
- MYSZKOWSKI, K. 2002. Perception-based global illumination, rendering, and animation techniques. In *SCCG '02: Proceedings of the 18th spring conference on Computer graphics*, 13–24.
- NG, R., RAMAMOORTHY, R., AND HANRAHAN, P. 2004. Triple product wavelet integrals for all-frequency relighting. *ACM Transactions on Graphics*, 477–487.
- OSTROVSKY, Y., CAVANAGH, P., AND SINHA, P. 2005. Perceiving illumination inconsistencies in scenes. *Perception* 34, 11, 1301–1314.
- PALMER, S. E. 1999. *Vision science: From Photons to Phenomenology*. Bradford Books/MIT Press, Cambridge, MA.
- PARKHURST, D., AND NIEBUR, E. 2004. A feasibility test for perceptually adaptive level of detail rendering on desktop systems. In *APGV '04: Proceedings of the 1st Symposium on Applied perception in graphics and visualization*, 49–56.
- PELLACINI, F., FERWERDA, J. A., AND GREENBERG, D. P. 2000. Toward a psychophysically-based light reflection model for image synthesis. In *SIGGRAPH '00: Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, 55–64.
- PERLIN, K. 2002. Improving noise. *ACM Transactions on Graphics* 21, 3 (July), 681–682.
- PONT, S. C., AND KOENDERINK, J. J. 2004. Surface illuminance flow. In *3DPVT '04: Proceedings of the 2nd international sym-*

- posium on 3D data processing, visualization and transmission, 2–9.
- QU, L., AND MEYER, G. W. 2006. Perceptually driven interactive geometry remeshing. In *SI3D '06: Proceedings of the 2006 symposium on Interactive 3D graphics and games*, 199–206.
- RAMAMOORTHI, R., AND HANRAHAN, P. 2002. Frequency space environment map rendering. In *ACM Transactions on Graphics*, 517–526.
- REDDY, M. 2001. Perceptually optimized 3D graphics. *IEEE Computer Graphics and Applications* 21, 5, 68–75.
- REINHARD, E., SHIRLEY, P., ASHIKHMIN, M., AND TROSCIANKO, T. 2004. Second order image statistics in computer graphics. In *APGV '04: Proceedings of the 1st Symposium on Applied perception in graphics and visualization*, 99–106.
- ROBILOTTO, R., KHANG, B.-G., AND ZAIDI, Q. 2002. Sensory and physical determinants of perceived achromatic transparency. *Journal of Vision* 2, 5, 388–403.
- ROCK, I. 1983. *The Logic of Perception*. MIT Press, Cambridge, MA.
- ROGOWITZ, B. E., AND RUSHMEIER, H. E. 2001. Are image quality metrics adequate to evaluate the quality of geometric objects? In *Proceedings of the SPIE: Human Vision and Electronic Imaging VI*, vol. 4299, 340–348.
- 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* 23, 3 (Aug.), 742–749.
- SUNDSTEDT, V., DEBATTISTA, K., AND CHALMERS, A. 2004. Selective rendering using task-importance maps. In *APGV '04: Proceedings of the 1st symposium on Applied perception in graphics and visualization*, 175–175.
- TE PAS, S. F., AND PONT, S. C. 2005. A comparison of material and illumination discrimination performance for real rough, real smooth and computer generated smooth spheres. In *APGV '05: Proceedings of the 2nd symposium on Applied perception in graphics and visualization*, 75–81.
- TODD, J. T., AND MINGOLLA, E. 1983. Perception of surface curvature and direction of illumination from patterns of shading. *Journal of Experimental Psychology: Human Perception and Performance* 9, 4, 583–595.
- TODD, J. T., NORMAN, J. F., AND MINGOLLA, E. 2004. Lightness constancy in the presence of specular highlights. *Psychological Science* 15, 1, 33–39.
- TUMBLIN, J., AND RUSHMEIER, H. 1993. Tone reproduction for realistic images. *IEEE Computer Graphics and Applications* 13, 6, 42–48.
- VAPNIK, V. 1995. *The Nature of Statistical Learning Theory*. Springer-Verlag, Berlin.
- 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* 24, 3, 1098–1107.
- WANDELL, B. A. 1993. Color appearance: The effects of illumination and spatial pattern. In *Proceedings of the National Academy of Sciences*, vol. 90, 9778–9784.
- WARD, G. J. 1992. Measuring and modeling anisotropic reflection. In *SIGGRAPH '92: Proceedings of the 19th annual conference on Computer graphics and interactive techniques*, vol. 26, 265–272.
- WATSON, B., FRIEDMAN, A., AND MCGAFFEY, A. 2001. Measuring and predicting visual fidelity. In *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, 213–220.
- WESTLUND, H. B., AND MEYER, G. W. 2001. Applying appearance standards to light reflection models. In *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, 501–51.
- WILLIAMS, N., LUEBKE, D., COHEN, J. D., KELLEY, M., AND SCHUBERT, B. 2003. Perceptually guided simplification of lit, textured meshes. In *SI3D '03: Proceedings of the 2003 symposium on Interactive 3D graphics*, 113–121.
- XIAO, B., AND BRAINARD, D. H. 2006. Color perception of 3D objects: constancy with respect to variation of surface gloss. In *APGV '06: Proceedings of the 3rd symposium on Applied perception in graphics and visualization*, 63–68.
- 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, 39–65.