A Framework for Realistic Image Synthesis

Donald P. Greenberg* *Kenneth E. Torrance** *Peter Shirley†*
James Arvo‡ *James A.Ferwerda** *Sumanta Pattanaik**
Eric Lafortune* *Bruce Walter** *Sing-Choong Foo***
Ben Trumbore*

Program of Computer Graphics*
Cornell University

Abstract

Our goal is to develop physically based lighting models and perceptually based rendering procedures for computer graphics that will produce synthetic images that are visually and measurably indistinguishable from real-world images. Fidelity of the physical simulation is of primary concern. Our research framework is subdivided into three sub-sections: the local light reflection model, the energy transport simulation, and the visual display algorithms. The first two subsections are physically based, and the last is perceptually based.

We emphasize the comparisons between simulations and actual measurements, the difficulties encountered, and the need to utilize the vast amount of psychophysical research already conducted. Future research directions are enumerated. We hope that results of this research will help establish a more fundamental, scientific approach for future rendering algorithms. This presentation describes a chronology of past research in global illumination and how parts of our new system are currently being developed.

CR Categories and Subject Descriptors: I.3.0 [Computer Graphics]: General; I.3.6 [Computer Graphics]: Methodology and Techniques.

Additional Key Words and Phrases: Realistic Image Synthesis, Light Reflection, Perception.

1 INTRODUCTION

From its infancy in the late 1960's, the quality of computer graphics images have improved at fantastic rates. The initial renderings of simple environments with only direct lighting have been transformed into pictures of complex scenes with shadows, shading, and global interreflections. For several decades now, high quality simulations have been used for a large number of tasks such as pilot training, automotive design, and architectural walkthroughs [GREE91]. The entertainment industry has developed techniques for creating startling special effects and realistic simulations with dramatic results. Even virtual reality games use convincing imagery with great success.

But are these images correct? Would they accurately represent the scene if the environment actually existed? In general, the answer is no; yet the results are appealing because the resulting images are believable.

If we could generate simulations that were guaranteed to be correct, where the algorithms and resulting pictures were accurate representations, then the simulations could be used in a predictive manner. This would be a major paradigm shift for the computer graphics industry, but would have much broader applicability than just picture making.

A look at how accurate simulations are used in other areas might clarify this point. The entire electronics industry is now based on simulations for chip design; these simulations are standardly used for testing and modifications prior to fabrication. In color science, we utilize the response matching functions for color transformations without recreating the matching experiments. Why can't we use computer graphics algorithms for the testing and development of printing technologies, photographic image capture or display devices? Why can't these accurate but artificial scenes be used for algorithmic development in image processing, robotics and machine vision? If we knew that the simulated images were correct, we could easily control and isolate the variables, obtain any precision or resolution desired, and avoid the difficulties and constraints of experimental measurements.

However, in order to be *predictive*, we must prove that the simulations are correct. *Fidelity* is the key. This difficult task requires a major multi-disciplinary effort between physicists, computer scientists, and perception psychologists and also requires experimental measurements and comparisons. Unfortunately to date there has been relatively little work done in correlating the results of computer graphics simulations with real scenes. However, with more accurate image acquisition and measurement devices, these compaisons will be achievable if we can generate adequate computer simulations.

From early computer generated images such as the Phong goblet (1975) to the synthesized pictures of today, there has been an exponential growth in the complexity of environments (Figure 1). This increased complexity has also led to an exponential growth in computational costs for realistic rendering (Figure 2). But the available processing power has also increased exponentially.
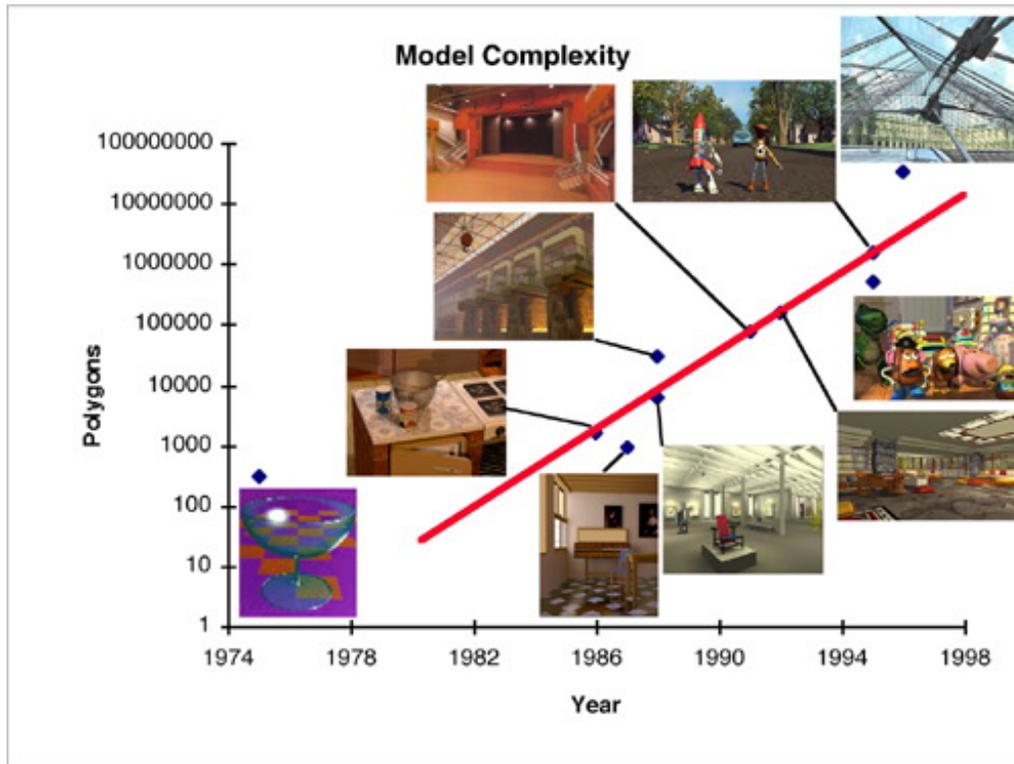
---

*580 Frank H.T. Rhodes Hall, Ithaca, NY 14853
† Department of Computer Science, University of Utah
‡ Department of Computer Science, California Institute of Technology
** Currently with Blue Sky Studios, Harrison, NY

Figure 1: *The polygonal complexity of several familiar models is plotted against the year in which the images were generated. Each scene may be considered "complex" for the time when it was rendered. The number of polygons in complex scenes is increasing exponentially as the years go by. Some polygon counts are estimated from available information [DISN97], and are for illustrative purposes only.*
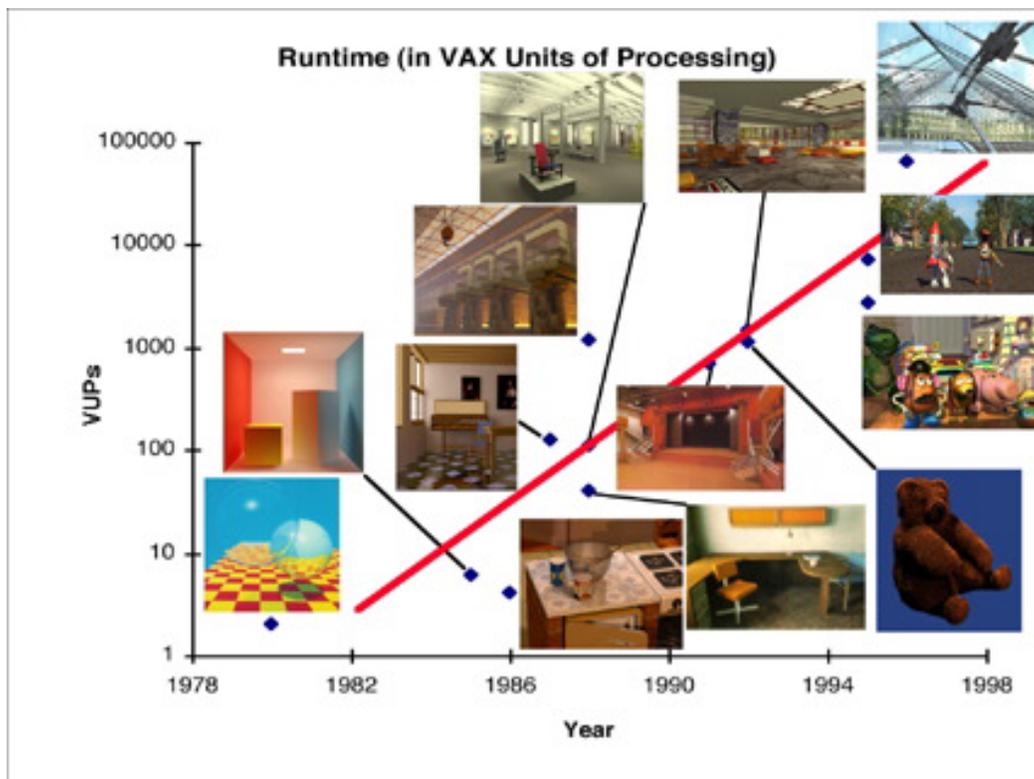

Figure 2: *The computation time for images of "complex" scenes is plotted against the year in which the images were generated. Computation time is presented in VAX Units of Processing (VUPs) [HENN97], which approximate the number of hours a DEC VAX 11/780 would require to perform the same computation. The processing power used to generate complex images is increasing exponentially with passing years. Some computation times are estimated from available information [DISN97], and are for illustrative purposes only.*

According to Moore's Law, with a doubling of chip density every 18 months, we now have approximately 4000 times the processing power that was available when the first graphics algorithms were developed (Figure 3). There has also been a concomitant increase in memory capacity, which offsets the constraint on environment complexity, as well as a vast reduction in cost per compute cycle. As we look toward the future a combination of increasing computation power and algorithmic improvements will certainly allow us to compute images that are physically and perceptually correct.

The purpose of this monograph is to describe in general terms our long-term development efforts attempting to achieve these tasks, to describe the difficulties encountered, and to encourage the computer graphics community to develop physically based algorithms of great realism and fidelity. Although there are many frontiers for future research in computer graphics, for physically based realistic image synthesis, three areas are especially critical: local light reflection models, light transport simulation, and perceptually based issues.
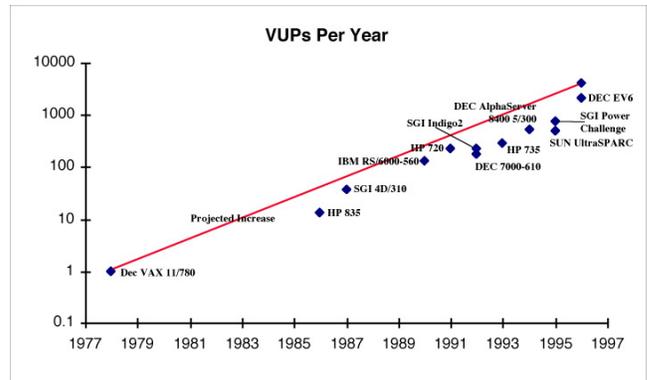
Figure 3: *The computational power of common computers is plotted against the year they were introduced. Computational power is presented in VAX Units of Processing (VUPs) [HENN97], which approximate how many times more powerful a machine is relative to a DEC VAX 11/780. The line represents a doubling of processing power every 18 months. VUPs were calculated using the LINPACK benchmark [DONG97]. Some machine introduction dates are approximate.*
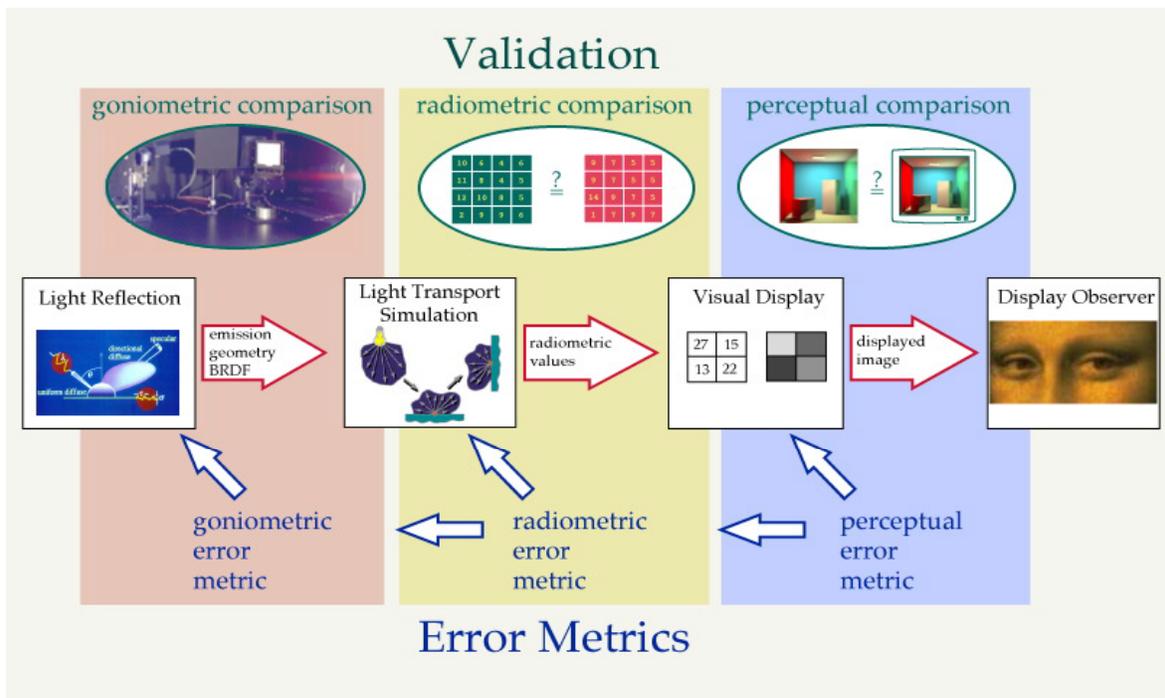
Figure 4: *The entire system is subdivided into three separate stages, the local light reflection model, the global light transport simulations, and the image display procedures. The first two stages simulate and compare physical processes, whereas the last stage occurs in the perceptual domain. Note that the squares represent processes and simulations, the ovals represent comparison measurements, and the bold arrows contain only data. As shown in the smaller arrows (bottom) each stage is related to itself and its predecessor stage by providing a "feedback loop" to reduce computations, resolution, or representations when below a given error threshold.*

Our specific long term goals are:

Light Reflection Models
- The development of a general purpose wavelength-dependent model or models for arbitrary reflectance functions including the effects of subsurface scattering and re-emission, texturing and surface anisotropy
- Validation of the local light reflection model through comparisons with measured physical experiments

- A means for representing this model in a compact, data-efficient form useful for progressive rendering algorithms
- Establishment and distribution of the reflectance characteristics of materials whose accuracy has been verified by measurements

Light Transport Simulation
- Creation of global illumination procedures capable of accurately simulating the light energy transport within complex geometric environments composed of surfaces with arbitrary reflection functions

- Validation of the global energy transport algorithms through comparisons to measured physical environments
- Creation of automatic and adaptive techniques for progressive energy transport algorithms
- Development of physical error estimates for radiometric computations

Perceptual Issues
- Creation of photorealistic synthetic images which are perceptually indistinguishable from real scenes
- Use of perceptual metrics to establish realistic visual display methods for a wide variety of display devices and viewing conditions
- Improving the efficiency of global illumination algorithms through the utilization of perceptual error bounds

For the past decade or more we have been developing a system to test, validate, and improve the fidelity and efficiency of computer graphics algorithms. An overview of this system is shown in Figure 4. This system is structured into three separate sub-sections dealing with the local light reflection model, the global light transport simulation, and the image display. Each of these stages is discussed in greater detail in following sections. What is of paramount importance is that at each stage, simulations are compared with measured experiments.

For the first stage, our ultimate goal is to derive an accurate, physically based local light reflection model for arbitrary reflectance functions. A measurement laboratory has been assembled to goniometrically measure and compare the local reflection model with a large number of samples. If the simulation model is correct, it becomes possible to send accurate data in terms of geometry, emission, and reflectance functions to the next stage.

With this information, it is then necessary to accurately simulate the physical propagation of light energy throughout the environment. For arbitrary reflectance functions and complex geometries, current simulation procedures are computationally excessive to say the least. As will be described later, most global illumination algorithms use simplifying assumptions, and although images of spectacular quality have been achieved, none have really guaranteed physical accuracy. If, however, it were feasible to simulate these reflection and transport processes, then once again we could measure and compare the resulting radiometric scene values.

Two factors are worth emphasizing. One is that the first two stages deal with physically based simulations only. The second is that we have not yet created a "picture", but are only comparing measured and simulated radiant energy on an image plane with full dynamic range and infinite resolution.

If the results of the first two physical stages are accurate, we then can proceed to the third phase of creating and perceptually comparing images. Since any comparison must utilize the human vision system, this stage occurs entirely in the perceptual domain. The computational processes must account for the limited dynamic range, limited spatial resolution, and limited color gamut of the display or printing devices. But the "mappings" should also account for the viewer's position and focus, state of adaptation, and the vast, complex, and mostly unknown relationships between the spatial, temporal and chromatic attributes of the scene.

One major benefit of this research will be to reduce the computational expense of the global illumination algorithms. An inherent cause of the slowness of these algorithms is that an excessive amount of time is spent computing scene features that are measurably unimportant and perceptually below the visible thresholds of the average human observer. Algorithms could be substantially accelerated if we can develop error metrics that correctly predict the perceptual thresholds of scene features. The establishment of these techniques will not only allow realistic visual display, but will also provide a feedback loop for reducing the magnitude of the physical computations. This will improve the *efficiency* of the global illumination algorithms.

In the following three sections, each of these stages, and the inherent difficulties in the measurement, testing and simulations are described. The light reflection stage is presented in the most detail since the research is relatively mature, though not yet "mainstream" in the computer graphics community. The light transport stage is described in less detail since the field is now mature, and the algorithms are generally well-known to the SIGGRAPH audience. The third stage dealing with perceptual investigations is an area of research just beginning, and results have not yet been achieved, as they depend on the accuracy of the prior two stages.

2 LIGHT REFLECTION

2.1 Light Reflectance Models

Light reflectance models have always been of great interest to the computer graphics community. The most commonly used model was derived approximately twenty-five years ago at the University of Utah [PHON75]. The Phong direct lighting model is a clever scheme using a simple representation, but it is neither accurate in the sense that it represents the true reflection behavior of surfaces, nor is it energy consistent. The arbitrary nature in which we assign the specular and diffuse coefficients and the associated energy of these components is not physically correct. Yet the entire industry is based on these early investigations, and all major manufacturers today use the same computationally efficient but overly simplified shading model.

Blinn [BLIN77] introduced a more accurate model, based on earlier work by Torrance and Sparrow [TORR67]. The model was based on geometric optics, thus accounting for self-shadowing effects from surface facets, and for the off-specular peaks that occur at grazing angles. By 1981, Cook [COOK81] further improved the reflection model by relating the solid angle of the illuminating light source and incorporating the Fresnel color shift effects of the specular at grazing angles. The work was particularly successful in depicting metallic objects.

None of these models, however, was sufficiently general. A comprehensive model of how light reflects or transmits when it hits a surface, including its subsurface interactions, needs to be developed. The resulting bidirectional reflectance distribution function (BRDF) is a function of the wavelength, surface roughness properties, and the incoming and outgoing directions. The BRDF should correctly predict the diffuse, directional diffuse, and specular components of the reflected light.

In 1991, He [HE91] presented a sophisticated model based on physical optics and incorporating the specular, directional diffuse, and

uniform diffuse reflections by a surface. The reflected light pattern depends on wavelength, incidence angle, two surface roughness parameters, and a surface refractive index. The formulation is self-consistent in terms of polarization, surface roughness, masking/shadowing, and energy. However, the model is extremely cumbersome to compute, since it contains an infinite summation that converges very slowly. This infinite summation term can be closely approximated by a spline surface and stored as a small look-up table of control points, vastly accelerating the computation of the full BRDF [HE92].

Poulin and Fournier [POUL90] constructed a similar model assuming a surface of cylindrical facets. Oren and Nayar [OREN94] also derived a model for more non-Lambertian diffuse reflection, based on diffuse micro-facets.

These models are applicable to a wide range of materials and surface finishes. For more complex surfaces, such as layered surfaces or thin films, analytical derivations are often too complicated. In those cases, Monte Carlo methods have been applied for simulating the local reflectance properties on a micro-scale [KAJI85] [CABR87] [HANR93] [GOND94]. For these models, and for reflectance models and measurements in general, a good means of representation is important.

2.2   Light Reflectance Measurement

In attempting to simulate the reflection and transport processes of light propagation, it is first necessary to measure and compare the material light reflection models (BRDF's) as well as measure the goniometric diagrams and spectral distribution of the illumination sources. For this reason it was necessary to set up a sophisticated light measurement laboratory (Figure 5).
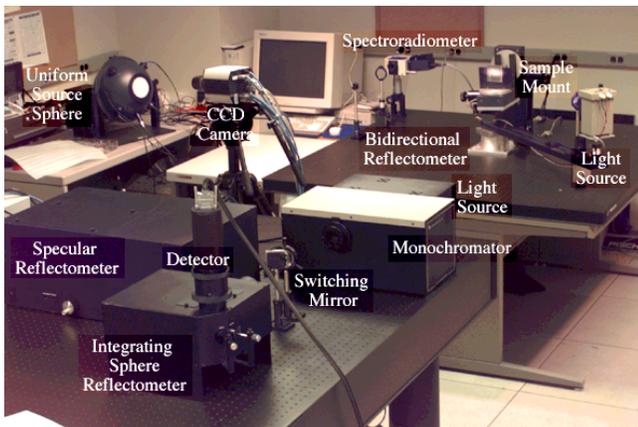

Figure 5: *Light Measurement Laboratory (1995)*

Difficulties for the measurement experiments are numerous and basically fall into three categories:

- Information on the calibration and precision of measurement equipment is scarce
- Physical constraints sometimes make the measurement a complex task, particularly for systems involving dense data sets
- Certain phenomena play a large role in reflection models, and yet are extremely difficult to measure

We would like to address each of these issues, at least by example, to show the problems.

To show the difficulties with the calibration and precision of the measurement equipment, let us just examine the tasks that had to be completed to use our high-resolution liquid-cooled camera, with a 1280 x 1024 CCD array. If you have either a guaranteed flat field illumination source, or a guaranteed flat field reflection from a surface, every pixel in the liquid cooled camera should yield an identical numerical intensity value. Using an integrating sphere uniform light source, we could produce a flat field, but there was considerable variation in the camera on a pixel by pixel basis. This could be caused by a variation in the sensors across the image plane or possibly by unequal thermal distribution.

Discrepancies could also be caused by lens effects, such as aberration or defects in the lens geometry, focusing and depth of field effects, selective absorption at different wavelengths, scattering within the lens and, most evident, the electronic and mechanical shutter effects. In either case, it was necessary to determine the behavior and response of the entire camera system over its full dynamic range before using the measurements. We did this with painstaking care so that we ultimately obtained different calibration factors for each pixel for each wavelength band.

Furthermore, if we wanted to deal with recording the intensities across the full spectral range of visible light we had to use a significant number of filters. We found that using eight band pass filters presented us with satisfactory data (Figure 6). Although this is not sufficient to reconstruct a continuous spectral distribution, we are sure of the accuracy of the individual basis functions. However, since we can compare the simulation of radiant energy at each one of these individual wavelength bands, we can make experimental comparisons without having to reconstruct the full spectrum image.
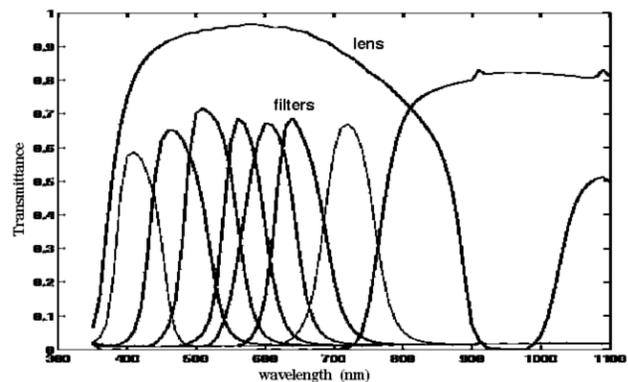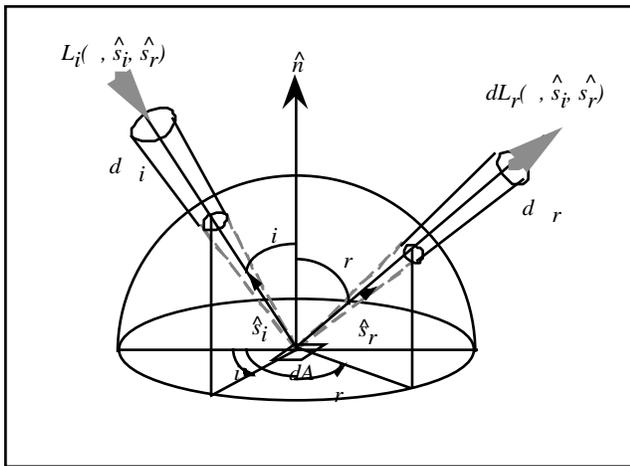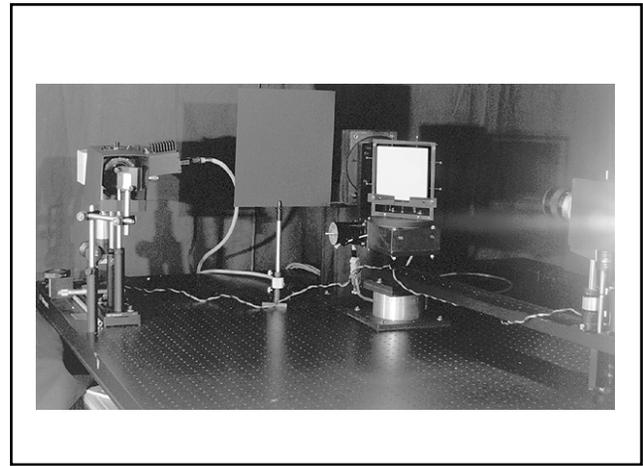

Figure 6: *Spectral transmittance of filters used in spectral image reconstruction*
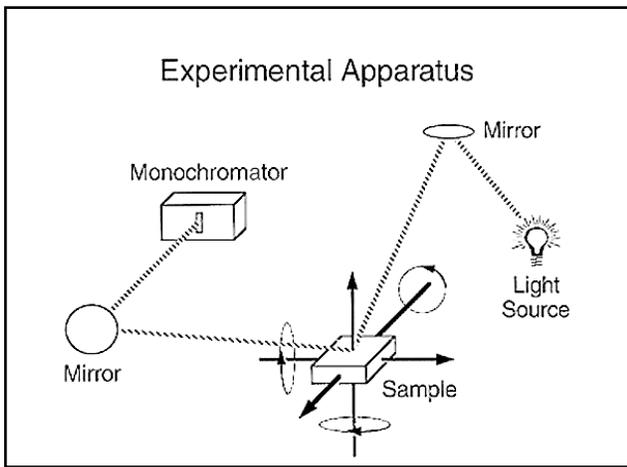
Let us now look at some of the physical constraints which make reflectance measurement a complex task. A gonioreflectometer is used to measure the full bidirectional reflectance distribution function (BRDF) for every wavelength and for every incidence or reflectance direction on the material surface. To do this the gonioreflectometer must have three degrees of freedom for an isotropic sample. Two of these are obtained by allowing the plane of the sample to be rotated and tilted. The third pivots the direction of the light source. For each position the sample must be illuminated by a constant light source with a uniform illuminating field (Figure 7).
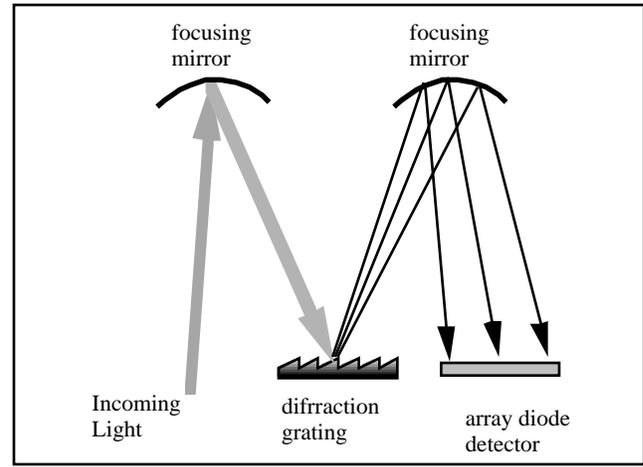
$L_i(, \hat{s_i}, \hat{s_r})$

$\hat{n}$

$dL_r(, \hat{s_i}, \hat{s_r})$

$d \quad _i$

$d \quad _r$

$_i$

$_r$

$\hat{s_i}$

$\hat{s_r}$

dA

(a)

(b)

Experimental Apparatus

Monochromator

Mirror

Mirror

Light
Source

Sample

(c)

focusing
mirror

focusing
mirror

Incoming
Light

diffraction
grating

array diode
detector

(d)

Figure 7:
a) Diagram of incidence and reflection geometry
b) Photograph of gonioreflectometer
c) Schematic of gonioreflectometer
d) Schematic of array diode spectrometer, the detector for the gonioreflectometer

Difficulties arise because of mechanical and optical constraints; for example, holding the sample in place can physically obscure the light source under particular conditions. We had to modify our equipment several times to obtain almost all of the hemispherical values. High grazing angles are also difficult since a small detector solid angle is required to measure the reflected energy. Unfortunately, this grazing region is where some of the most interesting characteristics occur, such as the off-specular peak (Figure 8).

To complicate the situation, the variation of the BRDF is relatively gradual in the diffuse and directional diffuse regions, but we need to take very closely spaced samples in the specular direction. Thus, the patterns used to measure the samples should not be uniform, but concentrate in the areas of greatest change (Figure 9).

Of even greater concern, perhaps, is the large role that phenomena which are difficult to model play in the BRDF's. Polarization has a very large effect and we have had to measure our isotropic samples

with two measurements after depolarizing the source and correcting for the bias of the detector. Measuring the anisotropy of the surface is difficult and we do not have the equipment within our laboratory to actually measure the microscopic deviations in the surface roughness, although we can make some statistical straight-line stylus probe measurements.

Since the reflection models are statistically based, we need to at least have the parameters describing the surface geometry. Fluorescence and phosphorescence also have substantial effects and these phenomena are material dependent. Even worse with fluorescence, light which may enter at a given wavelength may be re-emitted at many other wavelengths as there is crosstalk between color channels. Subsurface scattering depends on how far the light penetrates underneath the surface, and is a material property usually defined by the coefficient of extinction, but we have no mechanism for measuring this. Thus many problems still exist.
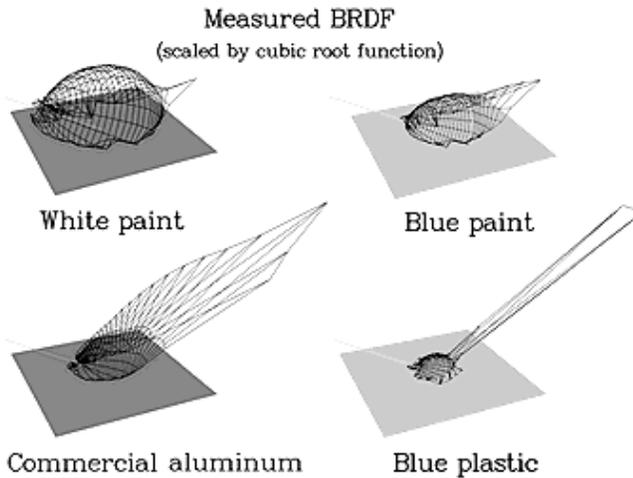
Measured BRDF
(scaled by cubic root function)

White paint

Blue paint

Commercial aluminum

Blue plastic

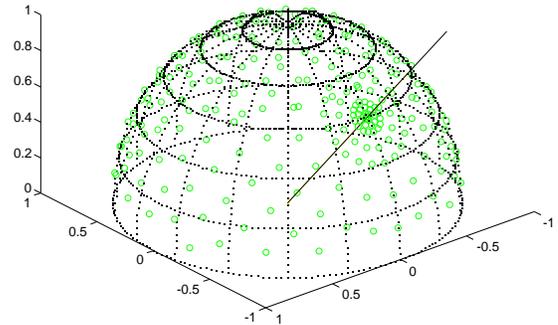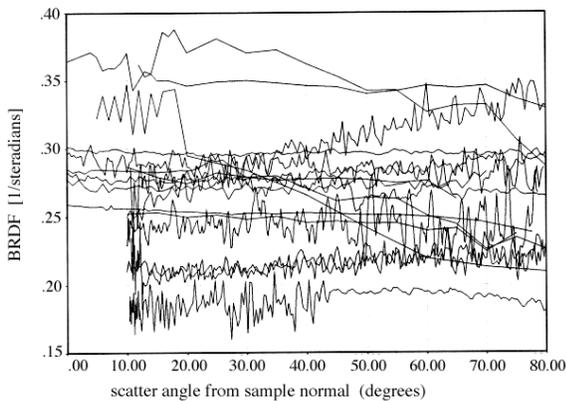**Figure 8:** *Measured BRDF for four isotropic materials.*

Figure 9: *Sampling positions of BRDF measurements on the hemisphere above the sample surface. The solid line indicates the specular reflection direction.*

(a)

(b)

Figure 10: *(from Leonard, 1988)*
(a) Measured BRDF curves for the diffuse white sample by various facilities participating in the BRDF round robin. The specular beam would be located at 10°.
(b) Measured BRDF curves of the diffuse black sample
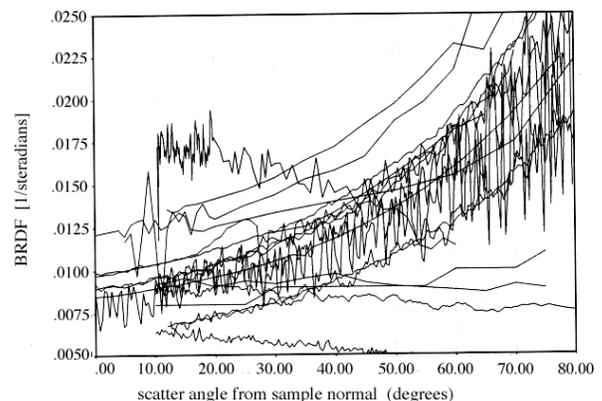
Once we have set up a BRDF experiment, the reflected light is measured in many directions on the reflection hemisphere. At each sample point the light is reflected through a diffraction grating so that we can measure the response at a thousand different wavelengths simultaneously. This creates a data explosion and thus both data storage and time for measurement become large[1]. However, despite all of these difficulties, we are now attempting to correlate our simulated light reflection model with our measured experimental results. Initial tests look very promising, at least for a small number of samples.

As stated previously, we are not the only ones to have difficulties with these measurements [LEON88] [LEON89]. In 1988, eighteen "scatter" measurement facilities participated in a round robin mea-

suring four two-inch diameter samples of a white diffuse surface, a black diffuse surface, an industrial grade molybdenum mirror and an aluminum mirror, both very smooth. The results showed an enormous range of deviation, confirming the difficulty of the task (Figure 10).

2.3   Light Reflectance Representation

Ultimately, what is necessary is a compact representational scheme which can accurately describe the dominant behavior of a BRDF. The functions must capture the diffuse, directional diffuse and specular characteristics, including the off-specular peaks, but must also be energy consistent and obey the laws of reciprocity. Furthermore, the representation method should be suitable for progressive algorithms, monotonically converging to a correct solution. Several researchers have used spherical harmonics with some success, but this representation has problems with ringing and negative values in the approximation [CABR87] [SILL91] [WEST92]. Work based on wavelets has been presented by Schröder [SCHR95]. This representation seems ideally suited for progressive algorithms and shows

---

[1]A typical set of data for a single isotropic material at 10 degree intervals for the incoming directions, with 800 outgoing directions for each incoming direction, at 8 wavelengths yields approximately 230 Kbytes. The potential size of accumulated material data indicates the utility of an accurate reflectance model with a compact means of representation.

great promise. Fournier [FOUR95] proposed the use of separable bicubic polynomials, and recently Koenderink etal. [KOEN96] presented a compact scheme based on Zernike polynomials. Ward [WARD92] and Lafortune [LAFO97] presented models with an emphasis on computational efficiency, taking into account that reflectance functions are generally only known to a limited accuracy that is sufficient for global illumination computations.

3   LIGHT TRANSPORT

3.1   Light Transport Theory and History

Once the emission, geometry, and reflection functions (BRDF's) are known, we can then simulate the light transport. The general equations have been well known [KAJI86], but until recently neither the processing power nor the physically based reflection models were available to perform accurate simulations.

The following equation is a modified form of Kajiya's rendering equation [IMME86] which expresses the outgoing radiance at a surface point x in terms of the emission from x and incoming radiance at x.

$$L\_o(\ ) = E(\ ) + \ \ (\ ,\ \ ')\ L\_i(\ ')d\ \ '\quad(1)$$

In the equation, ', are respectively the hemispherical incoming and outgoing directions, L_o( ), E( ) are respectively the total radiance and emitted radiance along w, L_i(') is the incoming radiance along ', and (, ') is the bidirectional reflectance distribution function (BRDF). Incoming radiance at any surface point x along ' in a scene is due to the outgoing radiance at another surface point x' visible to x along the direction opposite to ' in the same scene. In a complex scene the computation of the BRDF function and the visibility along the hemispherical directions itself is computationally expensive. To compute the solution of Equation (1), with complex BRDF and accurate visibility along all incoming and outgoing directions, for the outgoing radiances at all points in the scene is a formidable problem. Most algorithms make simplifying assumptions, for the BRDF function, for the visibility computation, and for the solution of the integral equation, yet still produce images of startling quality and realism. The two most common methods used are ray-tracing, introduced to the graphics community in 1979 [WHIT80], and radiosity, first presented five years later [GORA84]. Although during the past fifteen years many improvements have been made, neither of these commonly used algorithms are exact, each neglecting various and significant mechanisms of light transport.

3.1.1   Ray Tracing

View-dependent ray tracing methods originally computed only some of the transport paths, but accounted for specular transport in a visually compelling manner [WHIT80]. In essence, ray tracing reduces the BRDF expression to only include the path in specular direction [KAJI86], thus simplifying the computations but ignoring diffuse-diffuse and specular-diffuse interactions [WALL87]. Cook added a probabilistic framework to ray tracing to account for more physical effects [COOK84]. Kajiya extended Cook's framework to include all transport paths for a probabilistic view-dependent solution [KAJI86]. The main limitation of Kajiya's method in practice is that it takes a great deal of computation time. Many researchers have attempted to improve the methods of gathering statistics in Kajiya's methods, and this

is still an active and fruitful area of work [LAFO96] [VEAC95]. The main theoretical limitation of Kajiya's method is that although it is an unbiased method, meaning it will converge to the correct answer, its variance is unknown. The variance can itself be statistically estimated, but important statistical outliers can cause these error estimates to be misleading [LEE85].

Another important extension to Cook's framework was brought out by Ward [WARD88][WARD94a] from the observation that illumination due to diffuse-diffuse interactions changes gradually over surfaces and it is possible to reuse the computation at any surface point for points in close proximity. Though Ward's method ignores other indirect interactions such as specular-diffuse interactions, it provides a practical method for near accurate view-dependent solution of Equation (1) for many complex scenes.

3.1.2 Boundary Element Methods

Traditionally, view-independent radiosity-type solutions have been computed by boundary element methods. These methods work by interleaving the computation of the global light transport and the local lighting representation. In essence, these approaches model the transport processes by determining the "form-factor", the percentage of illumination leaving one surface element and reaching another. With the typical assumption of diffuse (Lambertian) reflection only, computations are based on geometric relationships only (shape, size, orientation, distance, and occlusion).

Although computationally expensive due to the complex visibility problem in real scenes, and matrices which are large due to the number of elements, equations can be solved (by Gaussian elimination) due to diagonal dominance [COHE85]. For complex environments, it is not practical to explicitly solve the full set of energy equations. Thus most solutions iteratively compute a portion of the global transport and update local representations until some convergence criteria are reached. To create high quality images, the requirement of very fine local representations, particularly in areas of high illumination gradients, e.g. shadow boundaries, gives rise to an exponential increase in elements. This combination of operations involving high global and high local complexity causes an explosion in resource consumption in terms of both memory and time.

Much research has gone into improving the basic finite element or boundary element method. Several techniques such as hierarchical radiosity [HANR91] and clustering [SMIT94] greatly reduce the time required at the expense of additional data structures and greater memory usage. Consequently, it is usually memory that limits the maximum input size and solution quality. To overcome this problem, researchers have tried various ways to reduce global or local complexity. Discontinuity meshing [CAMP90] [LISC93] attempts to precompute the potential locations of shadows to allow for a more compact local representation. This can produce dramatically better shadows, but it does not handle other lighting features, such as caustics and shadows from secondary sources, and does not scale well to large environments. Teller et al. [TELL94] try to partition the environment into small, weakly interacting subsets to lower the effective global complexity. If such a partitioning can be found then the computation can be ordered to use virtual memory efficiently; however such a partitioning may not exist (e.g., a hotel atrium). Smits et al. [SMIT94] reduce the local complexity by abandoning the idea of displaying the solution directly. Instead, the solution is computed at a low resolution and a computationally expensive local-gather pass is required to display the solution.

However, despite these impressive advances in reducing computational tasks, most of these schemes to date have been restricted to diffuse environments and static scenes. For more exact solutions, what is necessary is a physically based approach that can handle complex geometric environments with arbitrary reflectance functions, resulting in accurate solutions with known error characteristics.
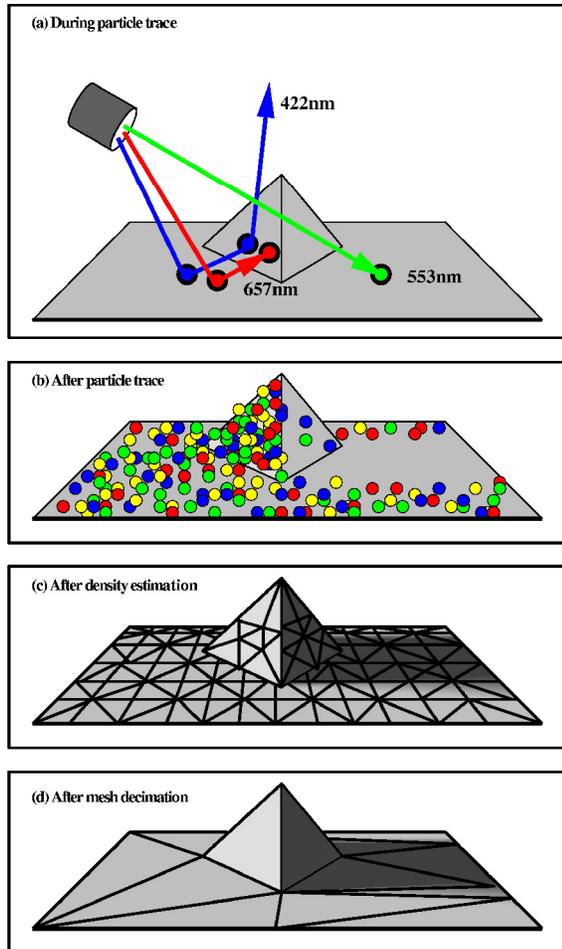
Figure 11: *(a) and (b) Particle Tracing: Power-carrying particles, each with a specific wavelength, are emitted from each luminaire using an appropriate spectral radiant intensity distribution and tracked as they travel through the environment until they are absorbed. Each time a particle hits a surface it is probabilistically absorbed or scattered in a new direction according to the BRDF of the surface. A list of all particle "hit points" is generated and saved.*
(c) Density Estimation: The stored hit points are used to construct approximate lighting functions on each surface. The illumination on a surface is proportional to the density of the hit points. To reconstruct luminous and chromatic exitance, the hits are weighted by the surface reflectance and the CIE XYZ response functions and local linear density estimation is applied. The result is a Gouraud-shaded mesh of triangles with three color channels suitable for direct display.
(d) Mesh Decimation: The initial mesh is generated conservatively and can be decimated by progressively removing vertices as long as the resulting change is below a perceptually-based threshold. Although not part of the simulation/measurement paradigm, it is useful to have the capability to rapidly display the rendered scene.

3.2   Density Estimation Framework

For our predictive simulations, we have chosen a density estimation framework which avoids the combination of high local and global complexity by splitting light transport and lighting representation into separate computational stages [WALT97a]. In the transport stage we compute the flow of light between surfaces without ever explicitly reconstructing the lighting on surfaces. Particle tracing is a natural and robust way to simulate this light flow. The representation stage then uses information from the transport stage to explicitly reconstruct the lighting on each surface. Since the intensity of the lighting is proportional to the density of light particles, the reconstruction is a density estimation problem [SILV86]. It should be emphasized that this is not a practical way to generate synthetic images, but does allow us to achieve the accuracy desired for complex environments with arbitrary reflectance functions.

Particle tracing has been used by many other researchers [APPE68] [ARVO86] [PATT93] to compute illumination, and Heckbert [HECK90] first noted that reconstructing lighting from particles is a density estimation problem. Since then a variety of different density estimation techniques have been applied, including histograms [HECK90], kernel methods [CHEN91] [COLL94], and splines [REDN95]. However, the fundamental difference between our framework and previous work is the separation of the transport and reconstruction stages.

The particle tracing stage computes a statistical simulation of global light transport. Since light particles do not interact there is no need to explicitly reconstruct the lighting function, and instead we simply record some information about the particle histories. Thus the particle tracing can work directly with the raw input geometry and has high global complexity but minimal local complexity.

The lighting reconstruction stage uses the recorded particle histories to estimate the lighting function on each surface. Because all of the global transport was handled in the previous phase, we can reconstruct the lighting on each surface independently. Reconstruction on a surface has high local complexity but no global complexity. Because each stage has only high global or high local complex-
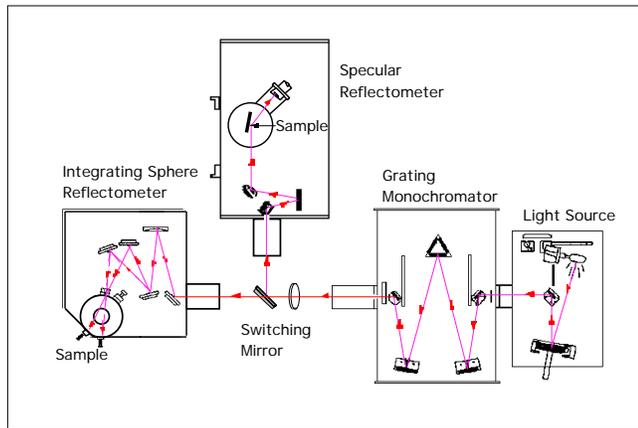

(e) An example scene rendered using Particle Tracing and Density estimation framework

ity but not both, the individual stages require fewer resources than finite element methods, especially in terms of memory. But we have not completely escaped the complexity problem. The combined high global and high local complexity is contained in the voluminous particle history data.

Our specific implementation is composed of three phases as shown in Figure 11. The system has been implemented because it is the only physically-based method we know which can handle complex geometries, general reflection functions, and whose error can be characterized. Although computational requirements are enormous, a major benefit of this approach is that particle tracing can easily exploit coarse-grain parallelism, thus reducing computation time.

The current implementation reconstructs irradiance function over the surfaces, which amounts to the reconstruction of radiance function for Lambertian surfaces only. However, the density estimation framework could be extended to reconstruct the radiance function for non-diffuse surfaces as well. The particle tracing phase already handles non-diffuse light transport. Depending on the degree of non-

diffuseness, more particles would be required to get an acceptable solution. The major unresolved issue is to find an efficient mechanism for storing the directional information. Solutions to this problem are currently being investigated. [WALT97b].

## 3.3 Light Transport Measurement and Image Validation

In keeping with our goal of comparing photorealistic image simulations against real world scenes, several steps are required. These include accurately measuring the input physical data, as well as acquiring a CCD camera image of the scene. The input data include scene geometry, light sources, and material reflectance properties, measured with appropriate equipment. The CCD measurement is acquired with our calibrated camera previously described. Some of the equipment for these measurements has been shown in Figure 5. The objective is to directly compare a simulation with a calibrated physical measurement.

Accurate simulations require accurate input data. We use the gonioreflectometer (Figure 7) to obtain the reflectance of surfaces with complex, strongly directional behavior. For surfaces that are specular or nearly ideal diffuse, we can use either a specular reflectometer or an integrating sphere reflectometer. Schematics are shown in Figure 12. These provide fast, accurate spectral measurements of the reflectance. The two reflectometers can also be used for transmission measurements, and the monochromator and detectors can be arranged for light source measurements. For measurements of scene geometry, we have relied on direct mensuration, or for small objects, a Cyberware scanner.

The CCD camera is used to acquire a 2D image of a scene. After calibration, twelve-bit images at any wavelength band can be obtained. An example image obtained through a 550nm filter is shown in Figure 13.

Figure 12: *Schematics of equipment for measuring the diffuse and specular reflectances.*

(a)
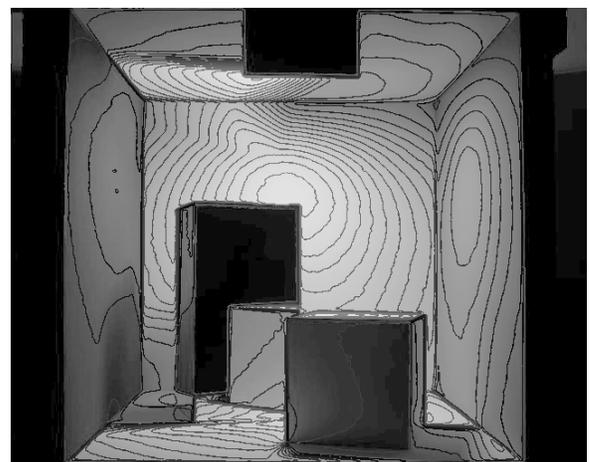
(b)

Figure 13:
(a) Calibrated CCD camera image via 550nm ($\Delta\lambda = 80nm$) filter
(b) The same camera image showing radiance contours

3.4 Future Directions

Before describing the perceptual issues, it is worthwhile examining our future directions in modeling the physical behavior of light reflection and light transport. Though significant progress has been made in modeling the surface BRDF, the model is far from complete. Properties such as polarization and anisotropy need to be well accounted for. Subsurface scattering which contributes towards the diffuse component of the BRDF is not well understood and is being handled empirically. Surface properties other than the BRDF which affect light interaction such as transmission, fluorescence, phosphorescence are either completely ignored or are being modeled empirically. These need to be correctly accounted for.

In the context of the simulation of light transport, we have introduced particle tracing and the density estimation framework. This framework can in principle handle all types of surface properties accurately. However, as mentioned earlier, further research is necessary to efficiently store and estimate the radiance function resulting from such accurate simulations.

4 PERCEPTION

A major goal of realistic image synthesis is to create an image that is perceptually indistinguishable from an actual scene. This is illustrated in Figure 14. This is not a trick photograph. The person in the figure is holding a real physical image generated by the rules of photographic tone reproduction.

Generating a visual image is the last stage of realistic image synthesis. At the end of the light transport process outlined in Figure 4 we have a global illumination solution that represents the radiometric values at every point in a three-dimensional scene. The final stage in image synthesis involves mapping these simulated scene radiances to display radiances to produce a visual image. This is an underappreciated and very important part of the image synthesis process that must take into account the physical characteristics of the display device, the perceptual characteristics of the observer, and the conditions under which the image will be viewed.

While the physically-based rendering methods described in the previous sections make it possible to accurately simulate the radiometric properties of scenes, this physical accuracy does not guarantee that the images displayed at the end of the process will have a realistic visual appearance. There are two reasons for this. First, current display devices are limited in a number of ways including spatial resolution, temporal resolution, absolute and dynamic luminance range, and color gamuts. Second, the scene observer and the display observer may be in very different visual states and this can affect how they perceive the visual information before them.

Current display technologies place fundamental limits on the fidelity of the display process. In the spatial domain, displays have fixed addressability and resolution, and are bounded in extent. In the temporal domain, they have fixed refresh rates and discrete update intervals. In luminance, both the absolute and dynamic ranges producible on displays are small relative to the ranges that can be measured in real scenes. Finally, in color the displays are trichromatic and have limited gamuts. The fact that display devices work as well as they do in creating acceptable visual representations of scenes is due to the fact that the human visual system is limited as well.

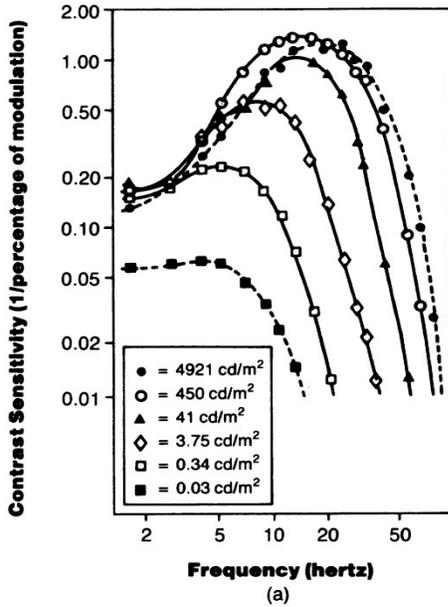## 4.1 Psychophysical Models Of Vision

For the last 150 years psychophysicists have measured the characteristics of human visual function. The contrast sensitivity function (CSF) shown in Figure 15 plots the spatial transfer properties of vision. The high frequency cutoff of the CSF indicates the spatial resolution of the visual system which is on the order of 60 cycles/degree or 1 minute of visual angle. The temporal contrast sensitivity function shown in Figure 16 plots the temporal response properties of the visual system.

In this case the high frequency cutoff indicates the limit of flicker sensitivity which at high levels of illumination is approximately 75-80 Hz. In the luminance domain, the threshold-versus-intensity functions show the relationship between just noticeable differences in intensity (JND) and the background illumination level. Over a wide range of illumination levels the visual system obeys Weber's law which states that the size of the JND is a constant proportion of the background
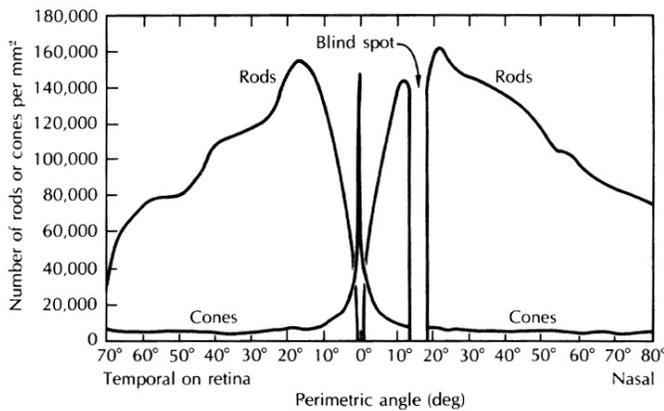

Figure 14: *The goal of realistic image synthesis: an example from photography [STRO86].*
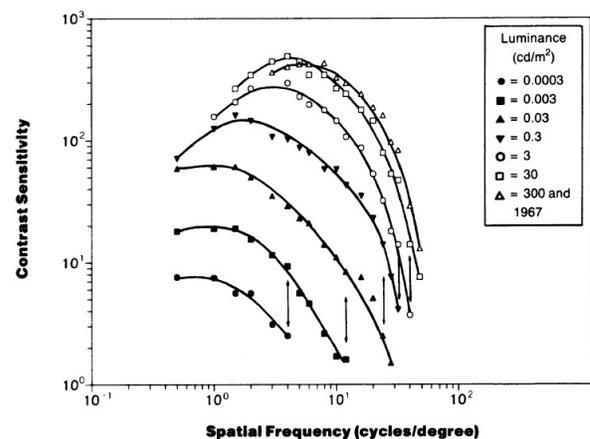
Figure 15: *The spatial contrast sensitivity function. The curves plot the relationship between contrast sensitivity and spatial frequency for detection of sine-wave gratings. The different curves reflect contrast sensitivity at the different levels of illumination shown in the inset, after [EVANN67], [BOFF86].*
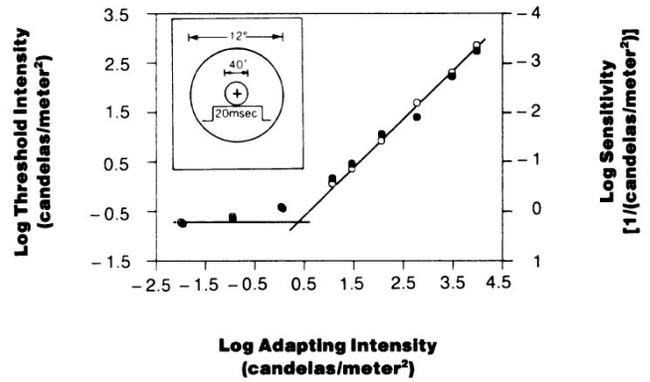
level (Figure 17). In color vision the shapes and sizes of the MacAdam ellipses on the CIE chromaticity diagram indicate that color discrimination is not uniform within the spectral gamut, but varies with chromaticity and with the direction of the chromatic difference (Figure 18). Lastly, our visual acuity decreases dramatically with the distance from the central fovea. Although acuity can be as good as 20/10 in the central field, at only 10 degrees in the periphery (roughly a palm's width at arm's length) acuity is reduced to 20/80 (Figure 19). This is an important phenomenon, especially when considering the necessary resolution for immersive displays.
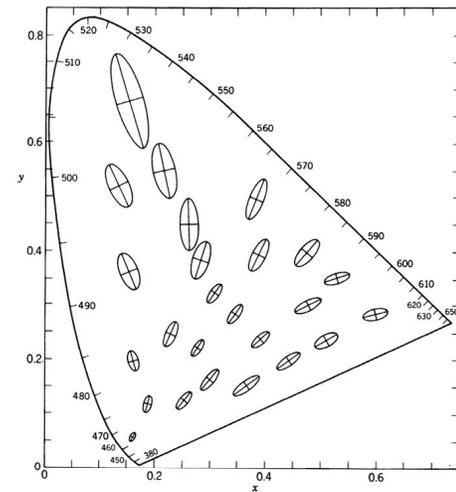

Figure 17: *Contrast threshold-versus-intensity function for the disk target and adapting annulus background shown in the inset. The abscissa shows the luminance of the background. The ordinate shows the threshold luminance necessary to detect the target. At low adapting luminances the background has no effect on the target. At higher luminances the threshold increases in constant proportion to the background luminance following Weber's law: ^{l}threshold= ^{kl}background. After [MUEL50], from [BOFF86].*

Figure 16: *The temporal contrast sensitivity function. The curves plot the relationship between contrast sensitivity and temporal frequency for a large white disk modulated sinusoidally . The curves reflect contrast sensitivity at the different levels of illumination shown in the inset. After [KELL61], from[BOFF86].*

Figure 18: *The MacAdam's ellipses for color discrimination plotted on the CIE chromaticity diagram. Ellipses show measured color difference thresholds at different chromaticities. The axes of the plotted ellipses are ten times their actual length [WYSZ82].*
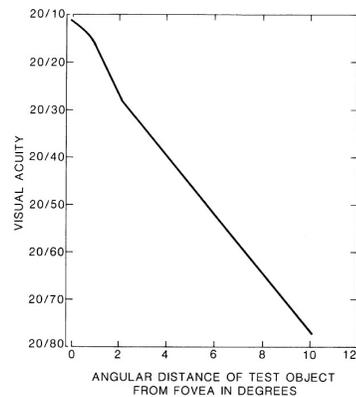

Figure 19: *Densities of rods and cones across the retina. Cone density is highest in the central fovea and falls off rapidly with eccentricity. Rod density is minimal in the fovea, increasing rapidly out to approximately 20 degrees peripherally and declining thereafter. The blind spot is a region devoid of photoreceptors located at about 17 degrees on the nasal retina. From Pirenne (1967). The graph on the right shows how acuity varies with retinal eccentricity. Acuity is approximately 20/10 in the fovea, this has dropped to about 20/80 ten degrees peripherally[Boff86].*

4.2 Use of Psychophysical Models In Imaging Systems

Reliance on psychophysical measurements is not new. Imaging system designers have used visual models for decades to improve the quality and reduce the bandwidth and computational load of imaging systems. In photography, subjective tone reproduction and preferred gamma curves incorporate Weber's law and simultaneous contrast effects. In color printing, knowledge of the trichromatic nature of vision allows full color reproduction from a small number of inks, while an awareness of spatial integration in vision has led to halftoning and color dithering techniques. Designers of simulation systems have taken advantage of differences in resolution across the visual field to reduce the level of detail for objects outside of the focal region. Finally, designers of image coding and compression systems such as NTSC, JPEG, and MPEG have used the spatial, temporal, and chromatic limits of vision to determine bandwidths and quantization levels for visual features of different scales, choose refresh rates and motion prediction methods for image sequences, and guide the choice of color coding schemes.

4.3 Using Psychophysical Models In Realistic Image Synthesis

In computer graphics we are only beginning to take similar advantage of visual perception in realistic image synthesis. Meyer [MEYE80] [MEYE86] has used results from color vision to develop color reproduction and spectral sampling strategies. Mitchell [MITC87] used a visual model to generate antialiased images at low sampling densities. Tumblin [TUMB93], Ward [WARD94b] and Ferwerda [FERW96] have developed tone reproduction schemes based on models of brightness and contrast perception and on a model of visual adaptation to realistically map the results of global illumination solutions to display devices. Spencer [SPEN95] developed a model of visual glare to both simulate the disabling effects of glare on visibility and to increase the apparent dynamic range of displayed images. Bolin [BOLI95] developed a frequency-based ray tracer that decides where to send rays based on the a model of the visibility of sampling artifacts. Recently Ferwerda [FERW97] has introduced a model of visual masking to predict when textures will hide visible artifacts in rendered scenes.

To improve the visual realism of synthetic images, it will be necessary to continue to "mine" the psychophysics literature for visual models that can be applied in computer graphics. A better understanding of the spatial, temporal, chromatic, and three-dimensional properties of vision can lead to more realistic and more efficient graphics algorithms.

4.4 Research Directions In Perceptually-Based Realistic Image Synthesis

To produce realistic images we need to model not only the physical behavior of light, but also the parameters of perceptual response. By modeling the transformations that occur during visual processing we can develop mappings from simulated scene radiances to display radiances to produce images that are as realistic as possible. Our goal is to show that these images can be predictive of what a observer standing in the physical scene would see. Validation of the predictive aspect of the images is a key component of the framework. Models of visual processing will also allow us to create perceptually-based error metrics for our rendering algorithms that reduce the computational demands of rendering while preserving the visual fidelity of the rendered images.

For our research framework, the approach used is based on the idea of a tone *reproduction operator* introduced by Tumblin [TUMB93] (Figure 20). The oval on the left represents the scene radiances simulated by our light transport algorithms. A hypothetical scene observer receiving these radiances will have a particular visual experience. On the right is a display observer looking at a display device driven by a graphics frame buffer. Our goal in realistic image synthesis is to have the display observer have the same visual experience as the scene observer. The tone reproduction operator maps the simulated scene radiances to the display radiances with the goal of producing a perceptual match between the display and the scene. There are two major components; the first is a model of the physical transfer properties of the display device including information about the absolute and dynamic range limits of the display, gamma correction factors, monitor white point, and color gamut. The second component of the tone reproduction operator is a visual model of the scene and display observers.
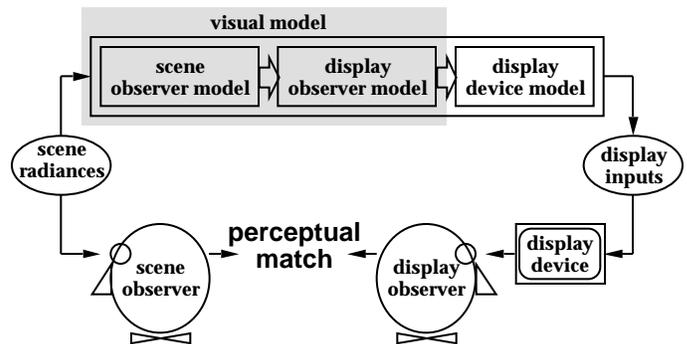
Figure 20: *The tone reproduction operator. A scene observer shown on the left receives the scene radiances and has a particular visual experience. We wish to have the display observer shown on the right have the same visual experience so the displayed image is a perceptual match to the scene. The tone reproduction operator maps the (simulated) scene radiances to display radiances taking into account the limits of the display device and the visual states of the scene and display observers. After [TUMB93].*

An accurate visual model is the essential component of a tone reproduction operator. The model allows us to characterize the visual states of the scene and display observers and allows us to relate them to determine the mapping from simulated scene radiances to display radiances.

Because the tone reproduction operator produces a perceptual match between the image and the scene this allows the images to be used predictively. Images produced by this method can be used quantitatively in areas such as illumination engineering, transportation and safety design, and visual ergonomics.

To claim that images generated by our visually-based tone reproduction operators are predictive, we have to validate this in com-

parison experiments. The results of such experiments will allow us to tune the visual models so the images we create are truly predictive. Furthermore, an experimentally validated visual model will also allow us to use the model in lieu of actual comparison experiments for the development of perceptually-based error metrics. These perceptually-based error metrics along with the previously determined physically-based error metrics will allow us to create more realistic and efficient image synthesis algorithms. If the end product of a simulation is a visual image then an efficient "steepest ascent" path can be derived to obtain a high fidelity visual solution with fewer computational demands.

We are just beginning this work but we believe that predictive visual models of phenomena like these are at the heart of future advances in computer graphics. To develop these models we will need to work together with experimental psychologists. But part of the difficulty in using the results of psychophysical experiments in graphics is that the experiments are typically carried out under extremely reductionistic conditions that only probe one dimension of visual function at a time. This makes it very difficult to generalize from these results to determine their effects in complex computer generated images, where there are significant interactions between the different dimensions of visual experience, and "crosstalk" between visual mechanisms is the rule rather than the exception. Two examples will make this clear.

The first example is Benham's top [HURV81]. Here a toy top is painted with the black and white pattern shown on the left in Figure 21. When the top is spun an observer sees red, yellow, and green rings even though there is no color in the pattern. The effect is due to interactions between the spatial, temporal, and chromatic mechanisms in vision. Current psychophysical models are only beginning to be able to explain these kinds of interactions. The second example shows that these interactions occur at even higher levels in the visual system. In Figure 22, adapted from the work of Adelson [ADEL93], the elements of the checkerboard block on the left and the flat pattern on the right have the same reflectances but the three-dimensional organization of the block makes us perceive them very differently.
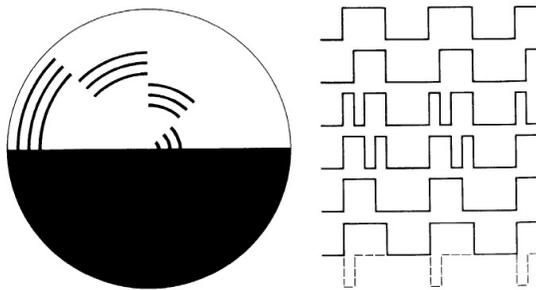
Figure 21: *Benham's top. When the black and white patterned top shown on the left is rotated at 5-10 rev./sec. colored rings are seen. The light intensity distribution of the rotating pattern as a function of time is shown on the right. Spatiotemporal interactions between antagonistic, spectrally-opponent color mechanisms account for the phenomenon. Thus while it is a convenient fiction for vision scientists to measure and characterize the visual system in terms of its simple responses to pattern, motion, and color, in fact the system is unified and the same visual mechanisms are at work in all cases [HURV81].*
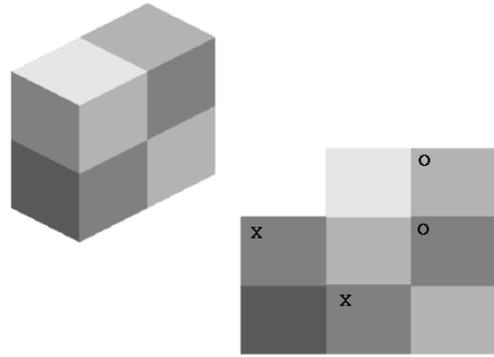
Figure 22: *Interactions between apparent reflectance, apparent illumination, and 3d spatial organization. The two patterns have the same reflectances. The 3d spatial organization of the block produces differing interpretations of the shaded regions in terms of lightness and brightness The two regions marked with X's have the same reflectance, but in the block they are perceived as different colors under different levels of illumination. The regions marked O's show another effect. Here the two regions have different reflectances but appear to be the same color, again under different illumination levels. After [ADEL93].*

Thus there are complex interactions between apparent reflectance, apparent illumination, and three-dimensional spatial organization that dramatically affect our perceptions of the identical visual stimulus. These interactions have implications for object recognition, color constancy, and other higher order visual phenomena. The quantitative aspects of these relationships are still not well understood.

In order to better understand normal vision outside the laboratory, vision researchers need to work with natural, complex, dynamic, three-dimensional visual displays. In the past, the problem has been that in these environments there has been little control to change the features of the environment, which has made careful experimentation difficult. By using the physically-based global illumination image synthesis methods, we should be able to provide rich visual environments for perceptual experiments that also allow precise control over the environmental features. Through this collaboration we hope to develop a better understanding of the complex interacting processes that underlie visual perception, which will allow us to develop more comprehensive visual models for computer graphics that will improve the realism and efficiency of our algorithms.

5 CONCLUSION

Our quest has been to develop physically based lighting models and perceptually based rendering procedures to produce synthetic images visually indistinguishable from real world scenes. To obtain *fidelity* we have subdivided our research into three parts: the local light reflection model, the energy transport phase, and the display procedures. The first two portions occur entirely in the physical domain, whereas the last is based entirely in the perceptual domain. All sections involve the testing and measurement necessary to guarantee *fidelity* and thus *predictive* algorithms. Since each stage is related to the next, feedback loops enable errors to be evaluated, a subsequent reduction in computational expense, and ultimately, we hope, more *efficient* algorithms.

The approaches and algorithms utilized are not currently practical, and require excessive computational resources. They do, however, yield important scientific insights into the physical processes of light reflection and light transport and clarify the computational bottlenecks. And because the algorithms are physically correct, they have already been used for simulating radiant heat exchange in turbine design, canopy detection in aerial reconnaissance, theater lighting and architectural and automotive design. With computing power increasing exponentially, global illumination algorithms will eventually become the norm. We hope that this research will ultimately help provide a better scientific foundation for future rendering algorithms.

Although the work presented has primarily been derived from research at one university, the effort is expanding to other universities, and there have been significant contributions by private and national laboratories. The research efforts should now be enlarged to a greater portion of the SIGGRAPH community. Only with a greater concentrated focus on these issues can we hope to improve the fidelity of our rendering algorithms.

Acknowledgments

References

[ADEL93] E.H. Adelson. Perceptual Organization and the Judgment of Brightness. In *Science*, vol. 262, pp. 2042-2044, 1993.

[APPE68] A. Appel. Some Techniques for Shading Machine Renderings of Solids. In *AFIPS 1968 Spring Joint Computing Conference*, pp. 37-49, 1968.

[ARVO86] James Arvo. Backward Ray Tracing. Developments in Ray Tracing. In *Computer Graphics* Course Notes, Annual Conference Series, 1986, ACM SIGGRAPH, pp. 259-263, 1986.

[BLIN77] James F. Blinn. Models of Light Reflection for Computer Synthesized Pictures. In *Computer Graphics,* Proceedings, Annual Conference Series, 1977, ACM SIGGRAPH, pp. 192-198.

[BOFF86] Kenneth R. Boff, Lloyd Kaufman, James P. Thomas. *Handbook of Perception and Human Performance.* John Wiley and Sons, 1986.

[BOLI95] Mark R. Bolin and Gary W. Meyer. A Frequency Based Ray Tracer. In *Computer Graphics,* Proceedings, Annual Conference Series, 1995, ACM SIGGRAPH, pp. 409-418.

[CABR87] Brian Cabral, Nelson Max and Rebecca Springmeyer. Bidirectional Reflectance Functions from Surface Bump Maps. In *Computer Graphics* , 21(4), Proceedings, Annual Conference Series, 1987, ACM SIGGRAPH, pp. 273-282.

[CAMP90] A. T. Campbell III, and Donald Fussel. Adaptive Mesh Generation for Global Diffuse Illumination. In *Computer Graphics*, 24(4), Proceedings, Annual Conference Series, 1990, ACM SIGGRAPH, pp. 155-164.

[CHEN91] Shenchang Eric Chen, Holly Rushmeier, Gavin Miller, and Douglass Turner. A Progressive Multi-Pass Method for Global Illumination. In *Computer Graphics*, 25(4), Proceedings, Annual Conference Series, 1991, ACM SIGGRAPH, pp. 165-174.

[COHE85] Michael Cohen and Donald Greenberg. The HEMI-CUBE, A Radiosity Solution for Complex Environments. In *Computer Graphics*, 19(3), Proceedings, Annual Conference Series, 1985, ACM SIGGRAPH, pp. 31-40.

[COLL94] Steven Collins. Adaptive Splatting for Specular to Diffuse Light Transport. In *Proceedings of the Fifth Eurographics Workshop on Rendering*, pp. 119-135, June 1994.

[COOK81] Robert L. Cook and Kennneth E. Torrance. A Reflectance Model for Computer Graphics. In *Computer Graphics* , 15(3), Proceedings, Annual Conference Series, 1981, ACM SIGGRAPH, pp. 307-316.

[COOK84] Robert L. Cook, Thomas Porter, and Loren Carpenter. Distributed Ray Tracing. In *Computer Graphics,* Proceedings, Annual Conference Series, 1984, ACM SIGGRAPH, pp. 137-145.

[DISN97] Anonymous, "About The Film," Walt Disney Home Video, 1996. http://www.disney.com/DisneyVideos/ToyStory/about/abfilm.htm.

[DONG97] Dongarra, Jack, "Linpack Benchmark," The Performance Database Server, May 5, 1997. http://performance.netlib.org/performance/html/linpack.data.col0.html.

[FERW96] J.A. Ferwerda, S. Pattanaik, P. Shirley, and D.P. Greenberg. A Model of Visual Adaptation for Realistic Image Synthesis. In *Computer Graphics,* Proceedings, Annual Conference Series, 1996, ACM SIGGRAPH, pp. 249-258.

[FERW97] J.A. Ferwerda, S. Pattanaik, P. Shirley, and D.P. Greenberg. A Model of Visual Masking for Computer Graphics. In *Computer Graphics,* Proceedings, Annual Conference Series, 1997, ACM SIGGRAPH, in press.

[FOUR95] A. Fournier. Separating reflection functions for linear radiosity. In Proceedings of the *Sixth Eurographics Workshop on Rendering*, pp. 383-392, Dublin, Ireland, June 1995.

[GOND94] J.S. Gondek, G.W. Meyer, and J.G. Newman. Wavelength dependent reflectance functions. In *Computer Graphics,* Proceedings, Annual Conference, 1994, ACM SIGGRAPH, pp. 213-220, Orlando, Florida.

[GORA84] Cindy M. Goral, Kenneth E. Torrance, and Donald P. Greenberg. Modeling the Interaction of Light between Diffuse Surfaces. In *Computer Graphics*, 18(4), Proceedings, Annual Conference Series, 1984, ACM SIGGRAPH, pp. 213-222.

[GREE91] Donald P. Greenberg. Computers in Architecture. In *Scientific American*, pp. 104-109, February 1991.

[HANR91] Pat Hanrahan, David Salzman, and Larry Aupperle. A Rapid Hierarchical Radiosity Algorithm. In *Computer Graphics,* 25(4), Proceedings, Annual Conference Series, 1991, ACM SIGGRAPH, pp. 197-206.

[HANR93] P. Hanrahan and W. Krueger. Reflection from layered surfaces due to subsurface scattering. In *Computer Graphics,* Proceedings, Annual Conference Series, 1993, ACM SIGGRAPH, pp. 165-174, Anaheim, California.

[HE91] Xiao D. He, Kenneth E. Torrance, Francois X. Sillion, and Donald P. Greenberg. A Comprehensive Physical Model for Light Reflection. In *Computer Graphics*, 25(4), Proceedings, Annual Conference Series, 1991, ACM SIGGRAPH, pp. 175-186.

[HE92] Xiao D. He, Patrick O. Heynen, Richard L. Phillips, Kenneth E. Torrance, David H. Salesin, and Donald P. Greenberg. A Fast and Accurate Light Reflection Model. In *Computer Graphics*, 26(2), Proceedings, Annual Conference Series, 1992, ACM SIGGRAPH, pp. 253-254.

[HECK90] Paul S. Heckbert. Adaptive Radiosity Textures for Bidirectional Ray Tracing. In *Computer Graphics*, 24(3), Proceedings, Annual Conference Series, 1990, ACM SIGGRAPH, pp. 145-154.

[HENN97] Henning, John L., "How Many VUPS Is That Alpha In The Window?" Digital Equipment Corporation, Feb. 1997. http://www.europe.digital.com/info/alphaserver/ performance/vups_297.html.

[HURV81] L.M. Hurvich. In *Color Vision*. Sunderland MA: Sinauer Assoc. Inc., 1981.

[IMME86] David S. Immel, Michael F. Cohen, and Donald P. Greenberg. A Radiosity Method for Non-diffuse Environment. In *Computer Graphics*, 20(4), Proceedings, Annual Conference Series, 1986, ACM SIGGRAPH, pp. 133-142.

[KAJI85] J. Kajiya. Anisotropic reflectance models. In *Computer Graphics,* 19(4), Proceedings, Annual Conference, 1985, ACM SIGGRAPH , pp. 15-21.

[KAJI86] James T. Kajiya. The Rendering Equation. In *Computer Graphics*, 20(4), Proceedings, Annual Conference Series, 1986, ACM SIGGRAPH, pp. 143-150.

[KELL61] D.H. Kelly. Visual Response to Time-Dependent Stimuli. I. Amplitude Sensitivity Measurements. In *J. Opt. Soc. Am.* 51, pp. 422-429, 1961.

[KOEN96] J.J. Koenderink, A.J. van Doorn, and M. Stavridi. Bidirectional reflection distribution function expressed in terms of surface scattering modes. In *European Conference on Computer Vision*, pp. 28-39, 1996.

[LAFO96] Eric Lafortune. Mathematical Methods and Monte Carlo Algorithms for Physically Based Rendering. In Ph.D. Thesis, Katholieke Universiteit Leuven, Belgium, February 1996.

[LAFO97] E. Lafortune, S.C. Foo, K.E. Torrance, D.P. Greenberg. Non-Linear Approximation of Reflectance Functions. In *Computer Graphics,* Proceedings, Annual Conference Series, 1997, ACM SIGGRAPH, Los Angeles, California.

[LEE85] M.E. Lee, R. A. Redner, and S.P. Uselton. Statistically Optimized Sampling for Distributed Ray Tracing. In *Computer Graphics,* 19(3), Proceedings, Annual Conference Series, 1985, ACM SIGGRAPH, pp. 61-67.

[LEON88] Thomas A. Leonard and Michael Pantoliano. BRDF Round Robin, Stray Light and Contamination in Optical Systems. In S.P.I.E. Vol. 967, Society of Photo-Optical Instrumentation Engineers, 1988.

[LEON89] Thomas A. Leonard, Michael Pantoliano, and James Reilly. Results of a CO2 BRDF Round Robin, Scatter from Optical Components. In S.P.I.E. Vol 1165, Society of Photo-Optical Instrumentation Engineers, 1989.

[LISC93] Dani Lischinski, Filippo Tampieri, and Donald P. Greenberg. Combining Hierarchical Radiosity and Discontinuity Meshing. In *Computer Graphics,* Proceedings, Annual Conference Series, 1993, ACM SIGGRAPH, pp. 199-208.

[MEYE80] G.W. Meyer and D.P. Greenberg. Perceptual Color Spaces for Computer Graphics. In *Computer Graphics* Proceedings, Annual Conference Series, 1980, ACM SIGGRAPH, pp. 254-261.

[MEYE86] G.W. Meyer. Color Calculation for and Perceptual Assessment of Computer Graphic Images. In Unpublished Ph.D. thesis, Cornell University, 1986.

[MITC87] D.P. Mitchell. Generating Antialiased Images at Low Sampling Densities. In *Computer Graphics,* 21(4), Proceedings, Annual Conference Series, 1987, ACM SIGGRAPH, pp. 463-474.

[MUEL50] C.G. Mueller. Frequency of Seeing Functions for Intensity Discriminations at Various Levels of Adapting Intensity. In *J. Gen. Psych.*, 1950.

[OREN94] M. Oren and S.K. Nayar. Generalization of Lambert's reflectance model. In *Computer Graphics,* Proceedings, Annual Conference Series, 1994, ACM SIGGRAPH, pp. 239-246, Orlando, Florida, July 1994.

[PATT93] S. N. Pattanaik. Computational Methods for Global Illumination and Visualization of Complex 3D Environments. In Ph.D. thesis, Birla Institute of Technology & Science, Computer Science Department, Pilani, India, February 1993.

[PHON75] Bui-Tuong Phong. Illumination for Computer Generated Images. In *Communications of the ACM*, 18(6):311-317, June 1975.

[PIRE67] M.H. Pirenne. In *Vision and the Eye*, 2nd edition. London: Associated Book Publishers, 1967.

[POUL90] P. Poulin and A. Fournier. A Model for Anisotropic Reflection. In *Computer Graphics,* Proceedings, 24(4), Annual Conference Series, 1990, ACM SIGGRAPH, pp. 273-282, August 1990.

[REDN95] R.A. Redner, M.E. Lee, and S.P. Uselton. Smooth B-Spline Illumination Maps for Bidirectional Ray Tracing. In *ACM Transactions on Graphics*, 14(4), October 1995.

[SCHR95] Peter Schröder, Win Sweldens. Spherical Wavelets: Efficiently Representing Functions on the Sphere. In *Computer Graphics*, Proceedings, Annual Conference Series, 1995, ACM SIGGRAPH, pp. 161-171.

[SILL91] Francois X. Sillion, James Arvo, Stephen Westin, and Donald Greenberg. A Global Illumination Algorithm for General Reflection Distributions. In *Computer Graphics*, 25(4), Proceedings, Annual Conference Series, 1991, ACM SIGGRAPH, pp. 187-196.

[SILV86] B. W. Silverman. In *Density Estimation for Statistics and Data Analysis*. Chapman and Hall, London, 1986.

[SMIT94] Brian E. Smits, James R. Arvo, and Donald P. Greenberg. A Clustering Algorithm for Radiosity in Complex Environments. In *Computer Graphics*, 28(3), Proceedings, Annual Conference Series, 1994, ACM SIGGRAPH, pp. 435-442.

[SPEN95] G. Spencer, P. Shirley, K. Zimmerman, and D.P. Greenberg. Physically-Based Glare Effects for Computer Generated Images. In *Computer Graphics*, Proceedings, Annual Conference Series, 1995, ACM SIGGRAPH, pp. 325-334.

[BREA88] Robert P. Breault. Stray Light and Contamination in Optical Systems. In S.P.I.E., 1988, p.234.

[STRO86] L. Stroebel, J. Compton, I. Current, and R. Zakia. In *Photographic Materials and Processes*. Boston: Focal Press, 1986.

[TELL94] Seth Teller, Celeste Fowler, Thomas Funkhouser, and Pat Hanrahan. Partitioning and Ordering Large Radiosity Calculations. In *Computer Graphics*, 28(3), Proceedings, Annual Conference Series, 1994, ACM SIGGRAPH, pp. 443-450.

[TORR67] K.E. Torrance and E.M. Sparrow. Theory for Off-Specular Reflection from Roughened Surfaces. In *Journal of the Optical Society of America* 57(9), September 1967.

[TUMB93] J. Tumblin and H. Rushmeier. Tone Reproduction for Realistic Images. In *IEEE Computer Graphics and Applications*, 13(6), pp. 42-48, 1993.

[VANN67] F.L. Van Nes and M.A. Bouman. Spatial Modulation Transfer in the Human Eye. In *J. Opt. Soc. Am.* 57, pp. 401-406, 1967.

[VEAC95] Eric Veach and Leonidus J. Guibas. Optimally Combining Sampling Techniques for Monte Carlo Rendering. In *Computer Graphics*, Proceedings, Annual Conference Series, 1995, ACM SIGGRAPH, pp. 419-428.

[WALL87] John Wallace, Michael Cohen, and Donald Greenberg. A Two-Pass Solution to the Rendering Problem. In *Computer Graphics*, 21(4), Proceedings, Annual Conference Series, 1987, ACM SIGGRAPH, pp 311-320.

[WALT97a] Bruce Walter, Philip M. Hubbard, Peter Shirley, and Donald P. Greenberg. Global Illumination Using Local Linear Density Estimation. In *ACM Transactions on Graphics*, 1997.

[WALT97b] Bruce Walter, Gun Alppay, Eric Lafortune, Sebastian Fernandez, and Donald P. Greenberg. Fitting Virtual Light for Non-Diffuse Walkthroughs. In *Computer Graphics*, Proceedings, Annual Conference Series, 1997, ACM SIGGRAPH.

[WARD88] Gregory Ward, Francis Rubinstein and Robert Clear, A Ray Tracing Solution for Diffuse Interreflection. In *Computer Graphics*, Proceedings, Annual Conference Series, 1988, 22(4), ACM SIGGRAPH.

[WARD92] G.J. Ward. Measuring and Modeling Mnisotropic Reflection. In *Computer Graphics,* 26(2), Proceedings, Annual Conference Series, 1992, ACM SIGGRAPH, pp. 265-272.

[WARD94a] Gregory Ward. The RADIANCE Lighting Simulation and Rendering System. In *Computer Graphics*, Proceedings, Annual Conference Series, 1994, 28(4), ACM SIGGRAPH, pp 459-472.

[WARD94b] G. Ward. A Contrast-Based Scalefactor for Luminance Display. In P.S. Heckbert (Ed.). In *Graphics Gems IV*, Boston: Academic Press Professional, 1994.

[WEST92]   S.H. Westin, J.R. Arvo, and K.E. Torrance. Predicting
Reflectance Functions from Complex Surfaces. In *Computer Graphics,* 26(2), Proceedings, Annual Conference
Series, 1992, ACM SIGGRAPH, pp. 255-264, July
1992.

[WHIT80]   Turner Whitted. An Improved Illumination Model for
Shaded Display. In *Communications of the ACM*, 23(6),
pp. 343-349, June 1980.

[WYSZ82]   G. Wyszecki and W.S. Stiles. In *Color Science: Concepts and Methods, Quantitative Data and Formulae
(2nd edition)*. New York: Wiley, 1982.