SIGGRAPH 98
25th International Conference
on Computer Graphics and
Interactive Techniques

COURSE NOTES 5
A BASIC GUIDE TO GLOBAL ILLUMINATION

Sunday, July 19, 1998
Half Day Course

ORGANIZER
Holly Rushmeier
IBM TJ Watson Research Center

LECTURERS

David Banks
NSF Engineering Research Center
Mississippi State University

Holly Rushmeier
IBM TJ Watson Research Center

Peter Shirley
Department of Computer Science
University of Utah

ABSTRACT
Images of the real world are formed by visible light being scattered by surfaces and volumes. The goal of global illumination methods is to simulate the path of light in an environment through the image plane in order to compute realistic images. Not all applications require the accuracy attainable with global illumination methods, and not all global illumination methods are good for all possible lighting effects. In this course the audience will be given a vocabulary and taxonomy for understanding global illumination. Insight into the basic methods will be provided using comparison to physical experiments. The target audience includes: people who are new to graphics who want to be generally informed, people who teach graphics courses but specialize in some other area of graphics, and/or people who think they may need global illumination for their application and want to understand how these methods differ from other rendering techniques.

ABOUT THE SPEAKERS

David C. Banks

Assistant Professor of Computer Science
NSF Engineering Research Center
Mississippi State University, MS 39762
601/325-0528 (voice)
601/325-8997 (fax)
banks@cs.msstate.edu

David Banks received his PhD from the University of North Carolina at Chapel Hill and held a post-doctoral position at the Institute for Computer Applications in Science and Engineering (ICASE) at NASA Langley Research Center. His research interests include applying computer graphics to study large-dimensional problems. He teaches graphics and visualization courses for undergraduate and graduate students.

Holly Rushmeier
Research Staff Member
IBM TJ Watson Research Center
30 Saw Mill River Road
Hawthorne, NY 10532
914/784-7252 (voice)
914/784-7667 (fax)
holly@watson.ibm.com

Holly Rushmeier is a research staff member at the IBM TJ Watson Research Center. She received the BS(1977), MS(1986) and PhD(1988) degrees in Mechanical Engineering from Cornell University. Since receiving the PhD, she has held positions at Georgia Tech, and at the National Institute of Standards and Technology. In 1990, she was selected as a National Science Foundation Presidential Young Investigator. In 1996, she served as the Papers chair for the ACM SIGGRAPH conference, and she is currently Editor-in-Chief of ACM Transactions in Graphics. She has published numerous papers in the areas of data visualization, computer graphics image synthesis and thermal sciences. In the area of global illumination she has worked on the problems of comparing real and synthetic images, imaging participating media, and combining ray tracing and radiosity methods. Most recently she has worked on global illumination methods suitable for image based rendering, accurate tone reproduction for high dynamic range images, and systems for acquiring physical data for realistic rendering.

Peter Shirley
Assistant Professor
3190 Merrill Engineering Building

Department of Computer Science
University of Utah
Salt Lake City, UT 84112
801/581-5290 (voice)
801/581-5843 (fax)
shirley@cs.utah.edu

Peter Shirley is an Assistant Professor at the University of Utah. He has a BA in Physics from Reed College and a Ph.D. in Computer Science from the University of Illinois. He worked at Indiana University and the Cornell Program of Computer Graphics before joining Utah. His research interests include realistic rendering, illustration, and visualization. He has taught several undergraduate and graduate courses on computer graphics in general, and global illumination in particular.

SYLLABUS

1:30 - 2:15 pm
Motivation and Definitions
David Banks and Holly Rushmeier
- presentation notes pp. 2-1 to 2-13

2:15-3:00 pm
Ray Tracing
David Banks
- presentation notes pp. 3-1 to 3-18

(3:00-3:15 pm Break)

3:15 - 4:00 pm
Radiosity
Peter Shirley
- presentation notes pp. 4-1 to 4-6

4:00 - 4:45 pm
Current Trends
Holly Rushmeier
- presentation notes pp. 5-1 to 5-7

4:45 pm Questions and Answers

TABLE OF CONTENTS

A Basic Guide to
Global Illumination

Motivation and Definitions

The purpose of this course is to give an overview of the area of computer graphics that has come to be known as "global illumination". The goal of global illumination is to make images of scenes which are defined numerically (may not physically exist yet) The images are predictive −− they are intended to show how the scene would appear if it were actually built. This is as opposed to artistic images or diagrams which may illustrate an individual's idea of what a scene would look like. Global illumination simulates the physical phenomenon of light transport.

Applications:
– Product appearance design
– Safety design
– Artistic effect acheived by physical means.

There are many reasons to make images. In many cases it is desirable to make "non−photorealistic" images that emulate artistic techniques such as sketching and painting. Global illumination is used when accurate predictions are needed for applications such as: what will the car look like in the showroom? Will the dashboard be visible to a driver at night? Will this theatrical lighting setup achieved the desired dramatic effect for a performance?

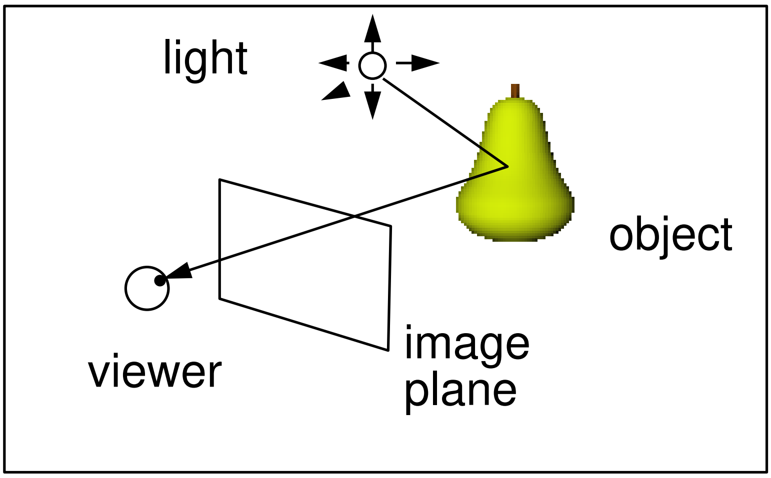

We see things as a result of how they interact with visible light To form an image we select a view point, view direction, image plane and image resolution. We color each pixel in the image according to what object would be visible through that pixel, and what quantity of light would be leaving that object in the direction of the viewer.

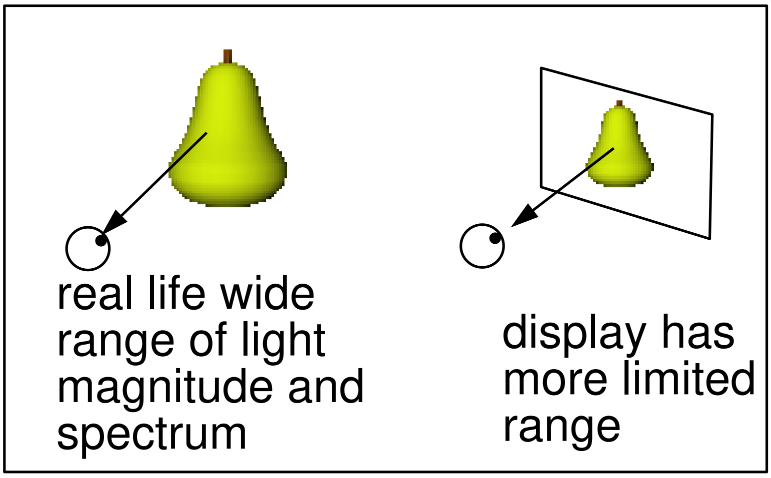

The RGB (red,green,blue) values we ultimately choose will not produce the same quantity of light on our display as we would encounter in real life. Displays have limited color gamuts and dynamic ranges. Mappings are needed to convert the quantity of light we predict to something displayable. These mappings use models of the human visual system. So, unfortunately to completely understand the formation of a realistic image, some knowledge of both the physics of light and the psychophysics of humans is needed.

Isn't lighting built into graphics hardware for making shaded images?

Systems such as Open GL and VRML have "lighting models" that are heuristics that emulate some lighting effects to render objects with shape and texture. However, these systems do not allow the definition of real light sources, physically realizable reflectances, do not include the "inverse square law", often have no shadows (or just sharp ones) and do not account for interreflections

---

Fundamental Effects:

– Direct Illumination

– Shadows

– Interreflections

– Volumes

Let's examine the fundamental effects that are achieved with global illumination that are not achieved by heuristic graphics lighting. Not all of these effects are equally important in every application, and they are certainly not all equally easy to compute. Sometimes they can be approximated by simple methods, but in some cases extensive calculations are required to get an adequate image. It is important to understand the effects critical to an application to determine the most cost effective approach to computing an accurate image.

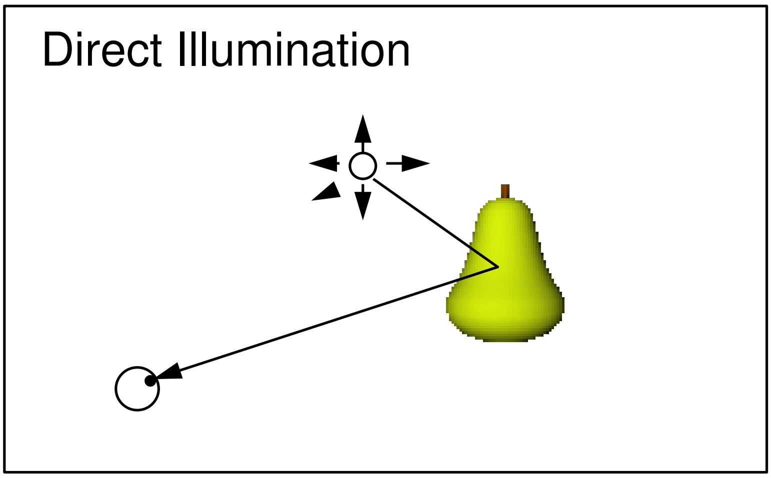

"Direct Illumination" refers to light that arrives at an object directly from the light source and then is reflected to the viewer. To accurately compute direct illumination, appropriate definitions of the geometry, directional, and spectral composition of the light source and the reflectance function of the object are needed. Modeling light sources and reflectances is sometimes referred to as "local illumination." See the section "Global Illumination Input" for more details.

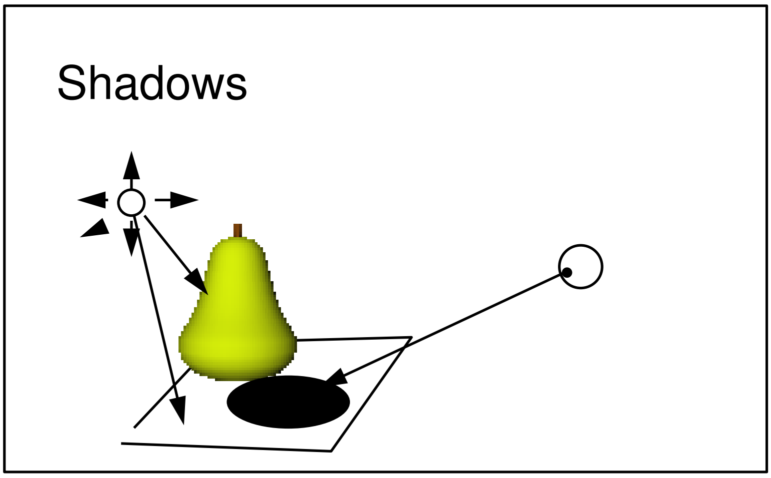

It is also important to find where light from the source does not reach an object. Shadows are an important cue to object locations −− we have the sense that the pear is floating above the plane becase of the location of the black ellipse used to represent its shadow.

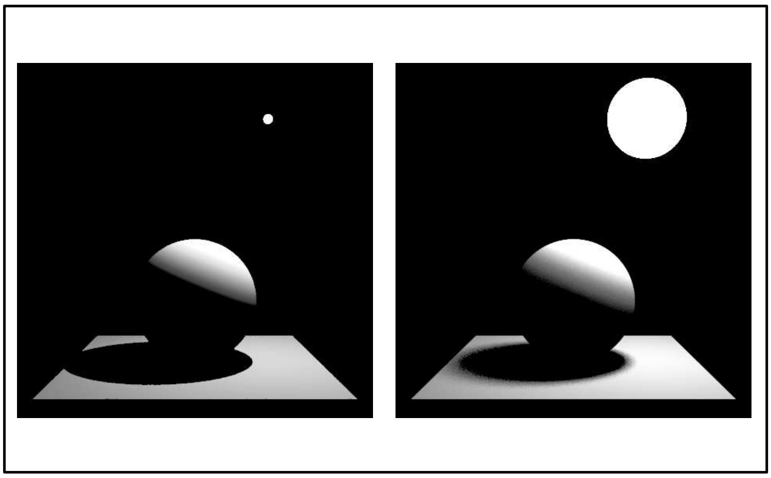

The shadow area where no light reaches is the umbra. There are some points where parts of the light source only are seen. called penumbra, which make the edges of the shadow look fuzzy. Whether the shadow is fuzzy or not depends on the sizes of the source and occlusion relative to the distances to the source and occlusion.

Many algorithms treat just the problem of how to compute shadows.. Classic techniques are Crow's shadow volumes (Crow82) and Williams' shadow maps (Williams78).

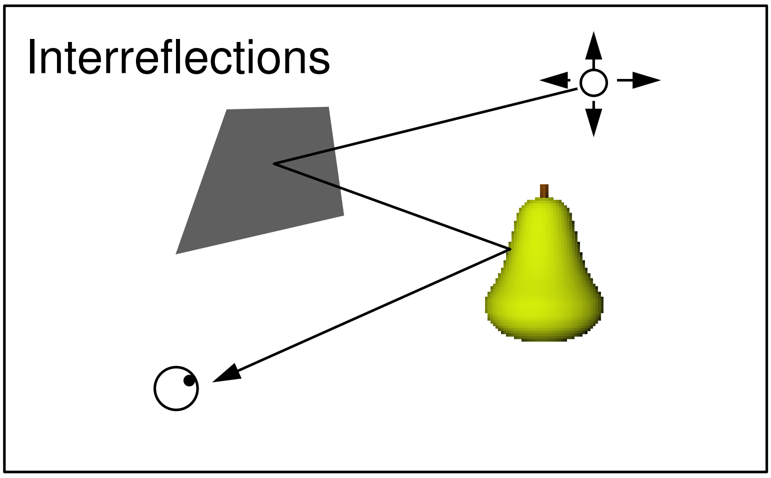

Interreflections are the "global" part of global illumination. The light that ultimately reaches the eye and has an effect on the image often goes through more than one bounce. Interreflections are expensive to compute −− in some scenes where they are not important it may be possible to neglect them or approximate crudely with simple calculations.

Types of Reflectance

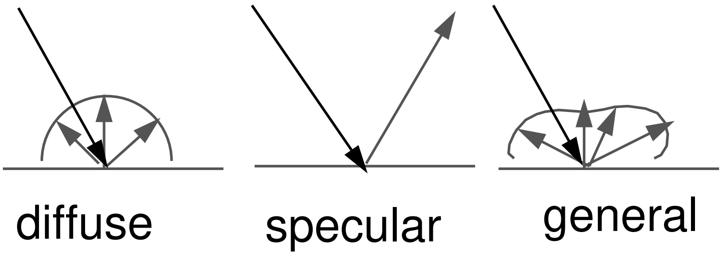

The effect of interreflections depends on the directional properties of the surfaces involved. Diffuse (a.k.a. Lambertian, matte) surfaces are characterised by the fact that you can't make out any objects reflected in the surface. Specular (a.k.a mirror−like) surfaces are characterized by the fact that you can see reflected objects clearly −− i.e. that's why we use them for mirrors ;) Many surfaces are neither of the idealized cases −− and reflections of objects may be seen dimly or fuzzed−out in a general surface

Interreflections: Diffuse Indirect Illumination

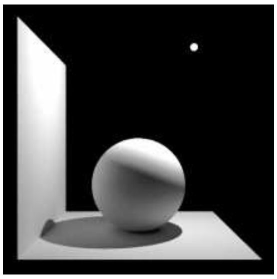

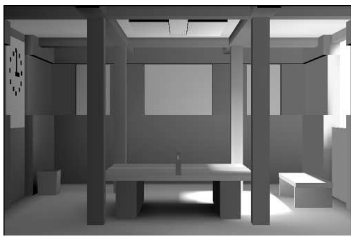

Indirect illumination, generally reflection off diffuse surfaces may cause surfaces which have no direct view of the light source to be illuminated. This is dramatic when there are many surfaces in the scene with no direct view of the source, as in these examples. However these interreflections are expensive to compute, and if most surfaces have a view of a light source, the effect of interreflections might be adequately approximated by a constant value. Good early examples of the effect of indirection illumination are shown in Nishita and Nakamae's 1985 SIGGRAPH paper

Interreflections: Diffuse Color Bleeding

The color of objects depends on their spectral reflectance, and the spectrum of incident light. (This is made a little more comples by the "color constancy" human vision phenomenon −− see "From Solution to Image" in the appendix). If a white wall, for example, is illuminated by indirect illumination from a red wall, the white wall will look somewhat red. This is generally a subtle effect, and was illustrated by the "Cornell Box" (Goral84) shown above in a early incarnation.

Interreflections: Caustics Bright Spots

In graphics, "caustics" refer to bright spots that are the result of a path of reflection or multiple reflections, from the light source, by several specular surfaces, and then finally hitting a diffuse surface. An extreme example is shown here where a spot light on the ceiling at the right is aimed at a mirror which reflects light through a crystal ball which focusses light into a bright spot to the left of the ball on the floor. This is in addition to the bright spot that results from the crystal ball focussing the main ceiling light onto the floor. Combining different paths of interreflection was discussed in Chen et al '91, in which this image appears

Fundamental Effects: Volumes

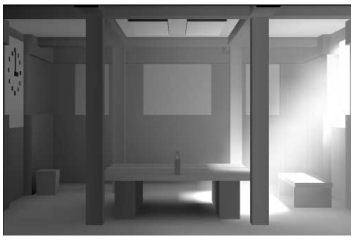

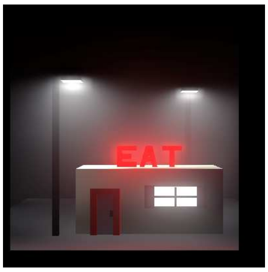

Besides surfaces, volumes of media can interact with light. Most of the time the volumetric medium in our environment (air) does not "participate" in the radiative transfer of visible light. However if there are water droplets (fog or clouds) or particulates (dust or smoke particles) in the air, these volumes of media partipate in the light transfer, and are called "participating media". Details on input and descriptions of participating media can be found in the appendix "Input for Participating Media."

Definitions: RADIANCE

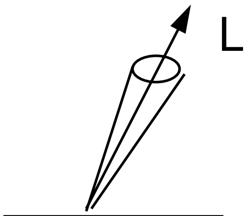

$$\frac{\text{energy}}{\text{time*projected–area*solid–angle}}$$

To form and solve equations for global illumination, we need to get specific on how to define a quantity of light. The basic quantity we want to solve for is radiance L The spectral radiance (i.e. radiance for various wavelengths of light) convolved with spectral functions related to the spectral sensitivities of the human visual system, will ulitimate be what we use to set the value of each pixel in an image.

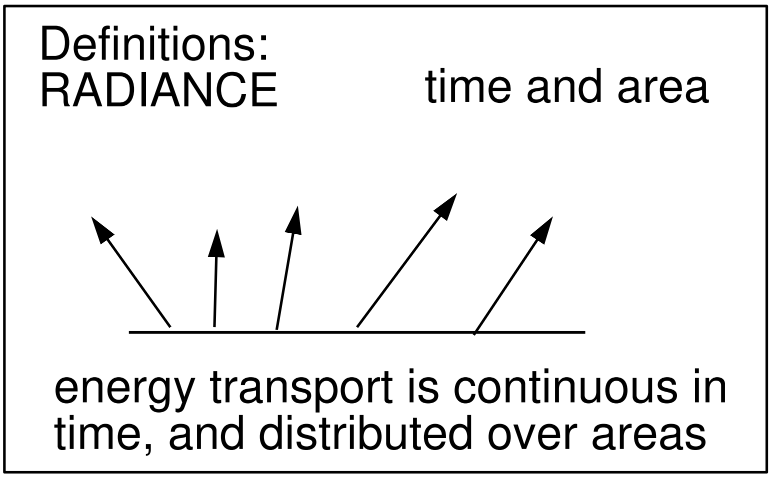

Why this definition? Energy is continually being transferred in our problem. In a still image the rate that visible light per unit time is constant, but it is being transferred continually. A point has no dimension, so strictly speaking there is zero energy leaving a point. We can discuss energy/time at a point though if we express it as energy/(time*area).

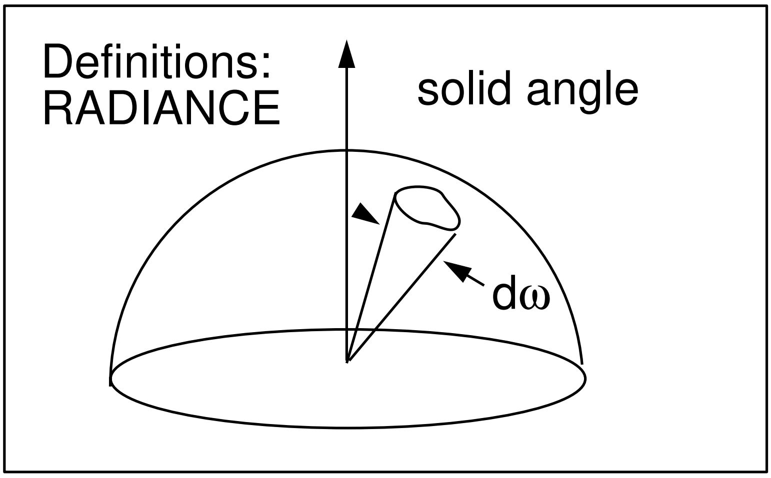

The light leaving a point may be different in each direction. All of the directions around a point are included in a hemisphere over the point. The hemisphere is said to subtend 2π steradians over the surface, analogous to a half circle subtending π radians.

A solid angle is a chunk of that hemisphere of directions. By integrating over all solid angles we can account for either all of the light leaving the surface per unit time and area, or all of the light energy incident.

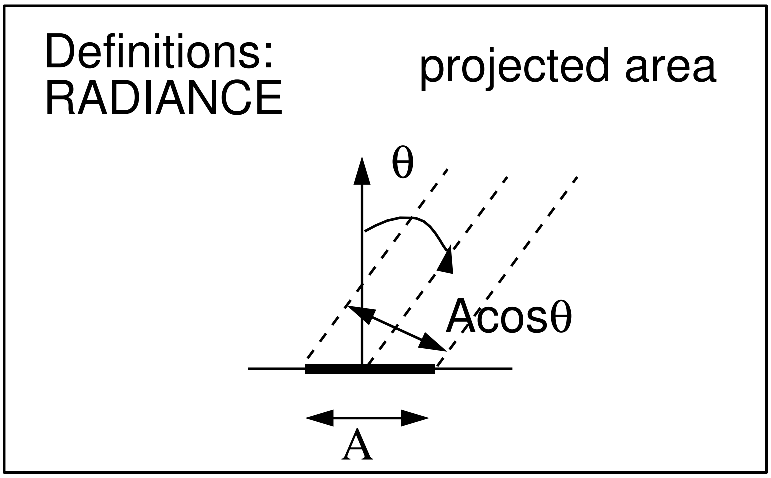

Why projected area and not just area in the radiance definition? When the surface is viewed at an angle, its area is forshortened by the cosine of the angle. By divided by projected area radiance expresses the quantity of light in terms of the effective surface area in the given direction

NOTE: Radiance is defined with respect to a surface, but not necessary a physical solid surface. Radiance is defined for any infinitesmal area specified by a location and surface normal, anywhere in space.

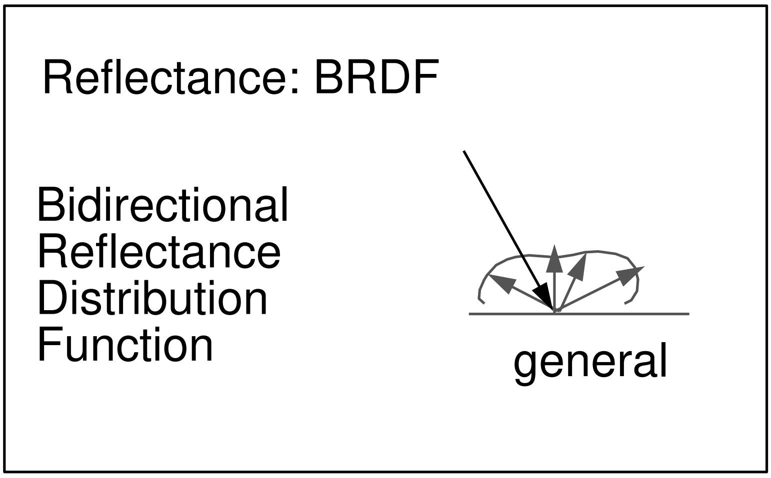

The other key definition is to precisely define the function that describes what happens to light when it is reflected from a surface. A reflectance is the ratio of reflected to incident light energy. A more general function expresses the directionaliuty of reflectance and is NOT a ratio that ranges from 0 to 1, but a distribution function that takes on any non−zero real value.

Reflectance: BRDF

$$f_r(\theta_i,\phi_i,\theta_r,\phi_r) = \frac{L_r(\theta_r,\phi_r)}{L_i(\theta_i,\phi_i)\cos\theta_i d\omega_i}$$

radiance/ energy flux density

The BRDF relates the reflected radiance in a particular direction (indicated here in spherical coordinates −− theta is the polar angle, phi is the azimuthal), to the incident energy flux density. For a general surface fr has a non−zero value for all pairs of incident and reflected directions.

Special BRDF:

$f_r = \rho_d/\pi$	$f_r=\rho_s\,\delta(\theta_m)/\cos\theta$
DIFFUSE	SPECULAR

The BRDF for the idealized surface reflectances have a simple form. A diffuse surface has a BRDF that is the same for all incident and reflected directions. The value ρd is the ratio of reflected to incident light energy. π is in the denominator for the diffuse surface as a result of integrating all directions with a cosθ weighting factor. A specular surface reflects light in only one direction for a given incident direction, so its BRDF is a delta function

Reflected Radiance

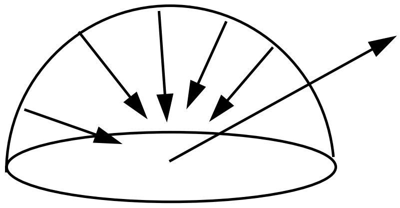

$$L_r(\theta_r,\phi_r) =$$

$$\int f_r\,(\theta_i,\phi_i,\theta_r,\phi_r)\;\; L_i(\theta_i,\phi_i)\cos\theta_i d\omega_i$$

To compute the radiance reflected from a point on a surface, we need to account for the fact that light may be incident from all directions, so we need to integrate over the entire incident hemisphere.

The Rendering Equation

radiance from object

radiance emitted from object

$$L_o(\theta_r,\phi_r) = L_e(\theta_r,\phi_r) \; +$$

$$\int f_r\,(\theta_i,\phi_i,\theta_r,\phi_r)\;\; L_i(\theta_i,\phi_i)\cos\theta_i d\omega_i$$

radiance reflected from object

An object may emit and/or reflect light. The complete rendering equation gives the radiance leaving an object accounting for both effects. This is just the well known equation of radiative transfer as used in heat transfer, illumination engineering, and various area of physics. The seminal paper "The Rendering Equation" by Kajiya in 1986 pointed out that this is the equation we want to solve to generate accurate images, and that in fact all of the approximations that had been made in an attempt to make realistic images were in some way an attempt to solve this equation.