painting was finished, Brunelleschi drilled a hole in it so that peering from the back one could see the painting reflected in a mirror. The painting positioned in this manner could be directly compared to its subject. When the mirror was removed, the Baptistry took the place of the reflected painting.

Such was the faithfulness of the painting to the real scene, that it is said viewers had difficulty perceiving any difference as the mirror was removed and replaced. Although historians do not know the precise nature of Brunelleschi's technique, it was recognized at the time as a means to reproduce, in a painting, the manner in which objects change their apparent size with increasing distance from the viewer, the fundamental idea of linear perspective.

Brunelleschi's experiments inspired a proponent in the person of Leon Battista Alberti, a scholar, painter, architect, and man of letters. Impressed by the demonstration, Alberti sought to codify Brunelleschi's findings and to make them the basis of a new mode of representation. In his seminal *On Painting,* Alberti describes a method for drawing in perspective that is rigorous and systematic. Following this technique, the space depicted in the painting is continuous with that of the observer. It is a transitional moment in the history of art: a painting becomes, in essence, a window open onto the world. *On Painting* stands today as one of the most influential works in the history of painting and a defining work of the Renaissance.

Alberti devised a practical method for perspectival representation that has special significance for computer graphics. The German artist Albrecht Dürer (1471–1528) illustrated the basic technique in a series of woodcuts, one of which is shown at right. ~~Alberti~~ placed a fine veil with black threads woven in a grid pattern between the painter and the scene. In front of the painter, he placed a second grid. He then fixed the point of view, directing the artist's gaze through the veil. The artist then painted on the grid the colors seen through each point using the black threads and grid lines as a guide. In this manner a geometrically correct image could be created point by point.

Alberti's technique is a method for projecting the three-dimensional scene in perfect proportion onto a two-dimensional plane. A tree in the background, for instance, occupies a smaller number of grid squares than does a tree of the same size in the foreground. Thus, scale is properly determined. Alberti's methodology is based upon the notion of the vanishing point: a road extending toward the horizon will diminish until the two sides occupy the same point. Indeed, any object moved farther and farther away, will subtend a smaller and smaller region of the grid until it occupies just a single grid point.

Alberti's method, although designed to enforce geometric perspective, coincidentally determines the color for each point in the image when transcribing a real scene. The color for any point in the image plane can be determined by defining a straight line that starts at the viewpoint, passes through the image plane, and encounters a surface in the scene. Of course, as noted in the previous chapter, the color of the point on the encountered surface is not only a property of the surface itself but also a function of the illumination of the scene.

Alberti identified four elements that are reiterated in computer graphics: the scene, sources of light, a visual pyramid that is transected by a picture plane, and a station point. In computer graphics, the term *image plane* is generally preferred to picture plane, and *viewpoint* is used rather than station point. In Brunelleschi's experiment, the viewpoint is the position of the artist's eye. In Dürer's woodcut, the artist places his eye at the tip of a small obelisk on a table, fixing the viewpoint as he works. The visual pyramid is formed by extending lines from the viewpoint through the four corners of the veil, with the veil itself transecting the pyramid and defining an image plane. The scene is simply the model reclining on two

PERSPECTIVE ILLUSTRATED In this woodcut, Albrecht Dürer depicts Alberti's schema for translating points in a scene to points in a drawing. A grid, shown here in a frame, defines an image plane from a particular viewpoint, indicated by the obelisk. The artist translates details from this grid to a corresponding grid on his drawing. An image constructed in this manner faithfully reproduces the spatial relationships between objects in the scene.

pillows, and the illumination is the daylight streaming in the back windows. The light and surfaces together determine the shading and shadow that the artist seeks to replicate with paint.

In the computer, of course, objects do not exist except as mathematically defined surfaces at different positions in a simulated three-dimensional space. Likewise, surface reflectances, as well as the positions, shapes, and spectral composition of light sources, and the viewing pyramid are all described mathematically. The guideline squares in Alberti's veil are large so that whole objects in the scene may fit within a single square. The painter uses the guidelines to maintain control over relative distances and sizes of the various elements in the scene. In computer graphics, however, Alberti's large grid squares have been replaced by pixels, rectangular elements of the computer screen that are, ideally, small enough to be imperceptible. Each pixel, furthermore, has only one color. Pixel color is determined by computing the light scattered from the scene through the pixel to the viewpoint. Since this light depends upon both the objects and the illumination, a great challenge of computer graphics is to simulate with mathematics all of the nuances of the play of light on surfaces within the synthetic scene.

From Simple to Complex: Defining Object Geometry

Before the interaction of light and surfaces can be simulated, all of the objects in the scene must be mathematically defined. Objects are defined as sets of points in three dimensions satisfying geometrical equations. Usually, the three-dimensional geometry is specified relative to a Cartesian coordinate system with three axes, x, y, and z. For example, a sphere can be described by the three-variable equation $(x - x_0)^2 + (y - y_0)^2 + (z - z_0)^2 = r^2$, where the sphere has its origin at the point (x_0, y_0, z_0) and a radius equal to r. Points whose x, y, and z values satisfy the equation form the surface of the sphere.

Typically, a computer graphics program is able to process only a relatively small number of geometric shapes. These objects, known as primitives, might include spheres, cylinders, and other forms that are easily characterized with simple mathematical formulae. Additionally, a number of points in three dimensions may be connected with lines to form the

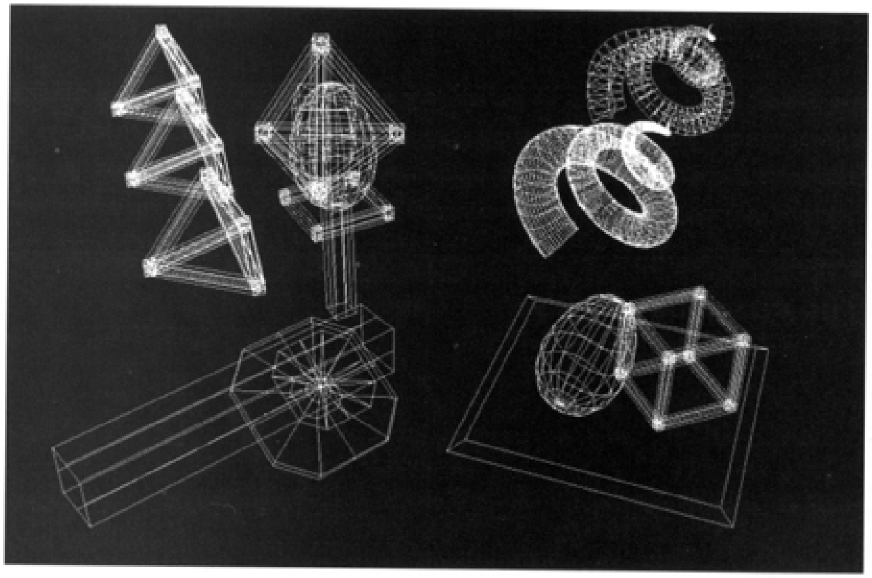

WIREFRAME RENDERING In three-dimensional computer graphics, objects are first defined mathematically as surfaces in the computer and only later acquire more familiar visible characteristics. Even the most complex objects must be assembled, in some manner, from geometric primitives such as spheres, cylinders, polygons, and curved surface patches. Each spiral in this illustration has been made by sweeping a two-dimensional cross-section along a three-dimensional curve. The eggs and the table support have been generated by rotating a two-dimensional curve around an axis. The tetrahedra and cube are constructed from cylinders, spheres and polygons. The teapot lid is comprised of curved patches. In this simple rendering, called a wireframe, all of these surfaces have been converted to polygons and only their edges are drawn. In subsequent images in this chapter (rendered by Matt Pharr of Stanford University), these objects will be illuminated with synthetic light using various computational models.

boundaries of a *polygon,* another kind of primitive. From this toolkit of simple shapes, more complex objects are created. Even visually intricate scenes can be built up of combinations of appropriately scaled and positioned geometric primitives.

For shapes that do not lend themselves easily to geometric descriptions a *polygon mesh* may be used. A polygon mesh is a set of connected, flat polygons that are chosen to approximate the irregular geometry of a surface. The polygon mesh method works well for objects composed of flat surfaces, but it often is a poor approximation for curved surfaces. To overcome this problem, the number of polygons can be increased, or alternatively, a *parametric mesh* can be utilized. As with a polygon mesh, a parametric mesh approximates a surface by stitching together a number of smaller patches. The difference is that these patches, unlike polygons, can be curved.

While it is certainly possible to define objects by writing geometric equations for their constituent primitives in computer code, it is much more convenient to use specialized applications that provide a *graphical interface* so that objects can be generated and modified by a direct, visual interaction with the computer. These applications, known as *modelers,* make it possible to automatically generate and then visually manipulate primitives to create three-dimensional forms. A primitive can be rotated, translated, and scaled before being positioned. Custom shapes are made by modifying a primitive: a silhouette edge such as a two-dimensional curve can be rotated around an axis to create a body of rotation, a process that can be compared to lathing in woodworking. A two-dimensional cookie-cutter shape may be dragged through a path in three dimensions to create a variation of a cylinder, a process that has kinship with extrusion.

Modelers make computer graphics much more practical and make it possible to focus on design rather than on mathematics. It would be a formidable task for a designer to specify intricate objects point by point. In these methods of lathing or extrusion, for example, the computer itself is used to create surface geometry. This approach is generally known as *procedural modeling* and is essential for creating complex objects such as clouds, mountains, and trees in the computer, as we shall discover in Chapter 5, "A Sorcerer's Apprentice."

Visible Surfaces

Up to this point, we have defined the surfaces of objects only as points satisfying equations: we still have no image of the scene to view. To synthesize an image, we need to model the light falling on the scene that is reflected to the image plane. This process, known as *rendering*, involves the other three elements defined by Alberti: the viewpoint, visual pyramid, and sources of light.

The viewpoint and viewing direction are specified in the computer in the same three-dimensional mathematical space that contains the scene. The angular extent of the image measured from the viewpoint, the field of view, uniquely determines a visual pyramid. The image is a slice of pixels through this pyramid, equivalent to Alberti's veil.

Unless an object is transparent, we should see only the portion of its surface that is unobstructed from our viewpoint. The next step in rendering is, therefore, to find the surface closest to the viewpoint for each pixel in the final image. By starting at the viewpoint, we can project a ray through each pixel into the three-dimensional scene. Since both the rays and the objects are described mathematically, the points that a ray intersects can be precisely computed.

A single ray through one pixel might intersect the front and back surfaces of several objects. It might, for instance, first intersect a flat polygon and then a sphere at two points, once on the front surface and once on

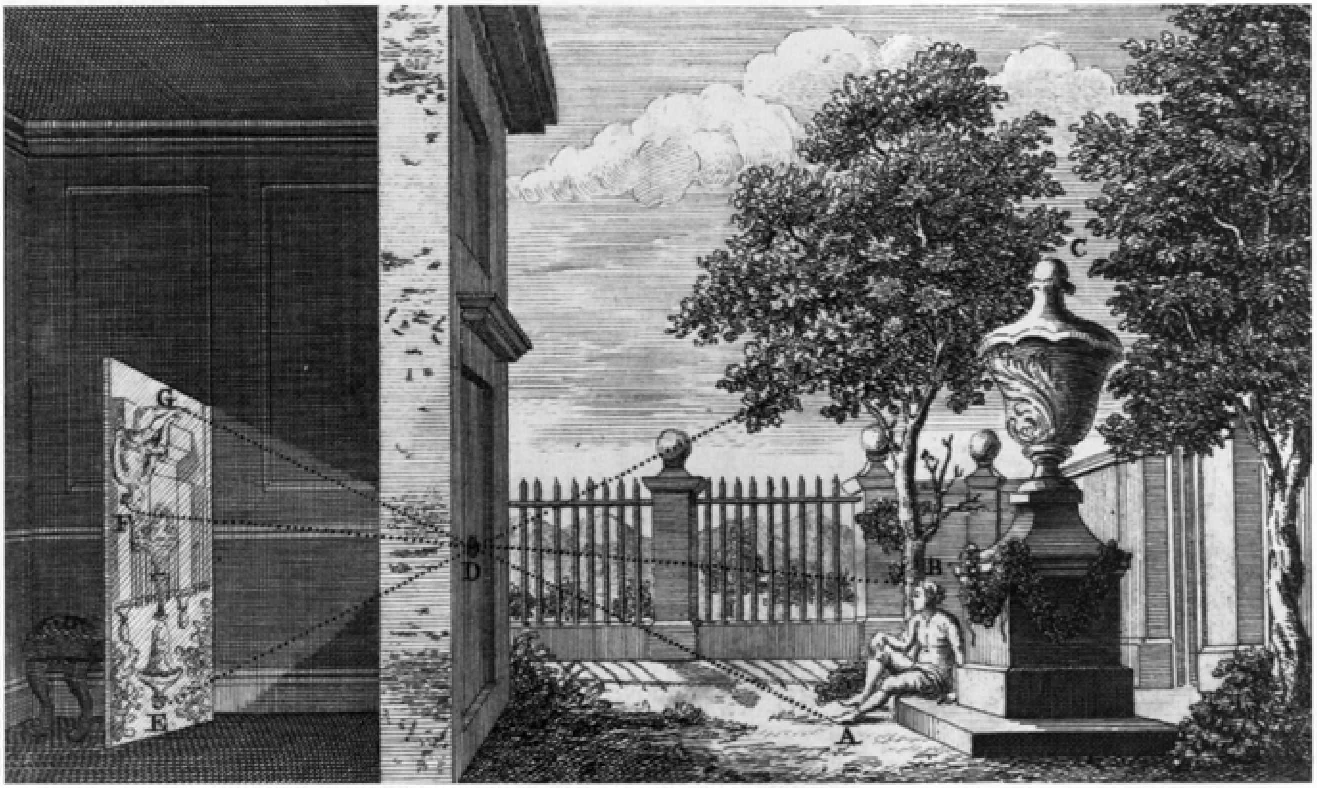

LIGHT TRAVELS IN STRAIGHT LINES Four centuries before Brunelleschi and Alberti, the Arabian mathematician and physicist Alhazen (965–1039) rediscovered classical theories of optics. He challenged the orthodoxy of his day that vision is mediated by rays emanating outward from the eyes and proposed the correct view that light is reflected to the eyes from the external world. Alhazen studied the manner in which light can project an image in a *camera obscura,* a dark room with a hole in one wall that admits light. As shown in this eighteenth-century etching, the light entering the room projects an inverted image. Through his experiments with the camera obscura, Alhazen concluded that light travels in straight lines and extended this insight to the workings of the eye and a theory of vision. His ideas influenced Leonardo da Vinci and other Renaissance painters and formed the conceptual foundations for perspective painting.

the back surface. The first object intersected by the ray usually is not known beforehand, however, so the intersections with all objects may need to be computed. The visible surface is simply the intersection with the shortest distance to the viewpoint. This procedure, determining which surfaces need to be rendered, is called *hidden surface removal.* When objects are intended to be transparent, however, there are no hidden

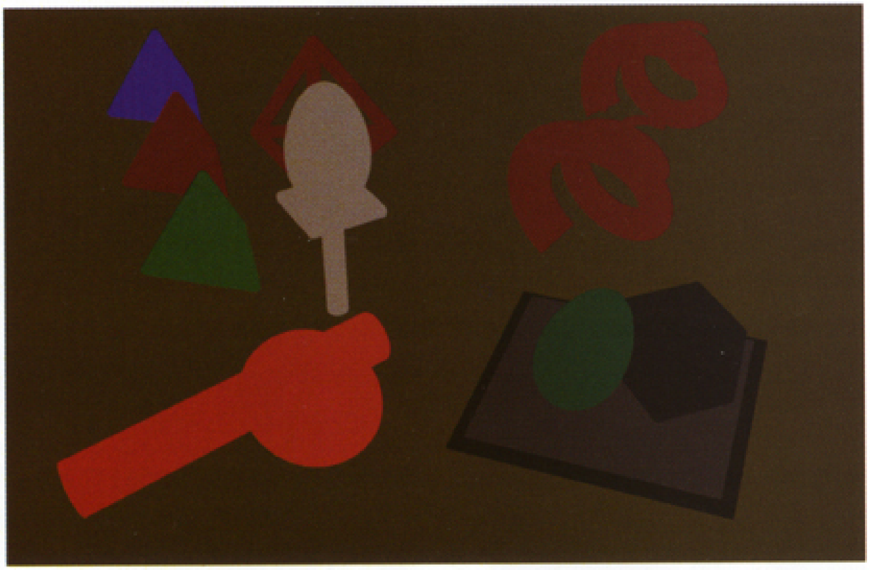

Hidden Surface Removal The computer must be instructed to render only surfaces visible from a particular perspective. In this image, occluded surfaces have been removed by identifying the surfaces that are closest to the viewpoint. Each surface is rendered with a single color and there is no illumination.

surfaces to be removed. Hidden surface removal is bypassed, and the rendered image depicts front and back surfaces.

The final step in image synthesis is the determination of the color of the visible surface represented by each pixel. This requires that the computer model the subtle interaction of all of the illuminants with surfaces.

Simulating Light and Surface

Our brief survey of some of the basic competencies of human vision suggests a certain reciprocity between computing and perceiving. Visual computing is used to synthesize an image in which light and surface are combined so that they can be later separated by vision. When all goes well, the image appears realistic.

The Envelope of Light

The development of algorithms to simulate light in the computer inevitably requires a rich education in the properties of the visual world and particularly the intricate nature of the illumination environment. In the natural world, light is composed of many different wavelengths, familiar as the visible spectrum from violet to red. A given light source will radiate

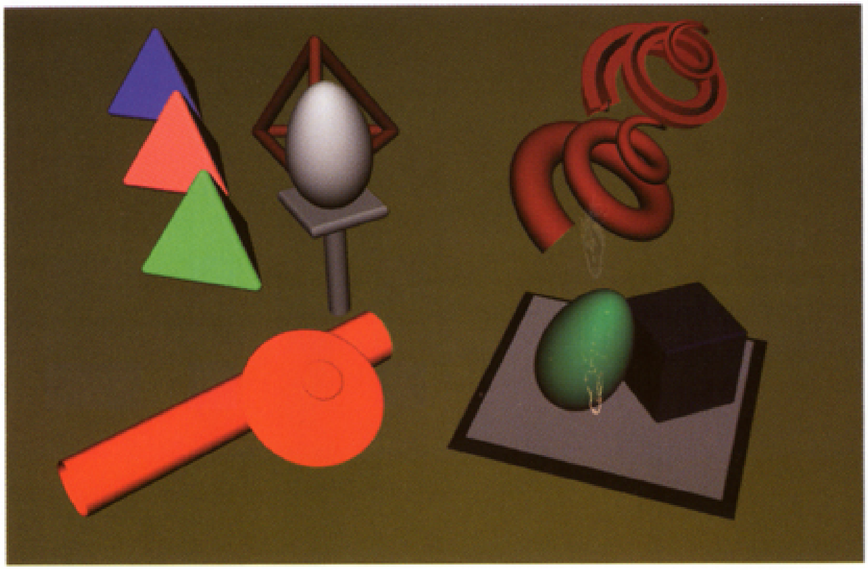

Diffuse Reflection A source of illumination has been introduced into the scene. The computer models the way in which this light reflects from each of the surfaces. In this image, the surfaces are modeled as purely diffuse reflectors. The reflected light does not depend upon the location of the viewer but only on the orientation of the surfaces with respect to the light source.

different proportions of energy at different wavelengths, the light's *spectral power distribution*. Light in a scene might emanate from multiple sources, each with its own spectral power distribution, shape, and position. For example, a small, tungsten filament emits light that consists predominantly of red wavelengths with very little energy in the blue portion of the spectrum. The small size of the filament also creates sharp shadows. Fluorescent light, on the other hand, principally comprises a small number of wavelengths of light produced by the energized gas, but it is softened by a coating that fluoresces broadly across the visible spectrum. Fluorescent lights, also in contrast to tungsten light bulbs, create very soft shadows, or *penumbrae*.

Light reflected from a surface is a product of the spectral power distribution of the illuminant and the surface reflectance.

Typically, the surface reflectance of an object changes as the orientation of the surface is altered with respect to both the illumination direction and the viewing direction. A simple experiment with a flashlight beam reflected off a flat matte board illustrates the effect of surface orientation. As the board is rotated to face your eyes or to face the flashlight, the amount of reflected light will vary, but only gradually. This type of reflection, called *diffuse reflection* (or *Lambertian reflection,* after the French mathematician who studied it, Johann Lambert, 1728–1777) is a function will cause certain wavelengths to be absorbed and converted to heat energy rather than reradiated, so that the spectrum is altered in reflection.

Nonmetals such as glass, transparent plastic, and water, on the other hand, have low conductivity. Having no conducting electrons and typically no other mechanism to absorb visible light, uniform nonmetals are generally transparent. When light falls on the surface of a nonmetal, a portion of the light is reflected and the remainder passes entirely through the material. For this reason, when we look at a pane of glass, we can see a reflection of ourselves as well as light through the glass.

As light passes from one transparent or translucent material to another, its speed, and consequently its direction, is changed. This property of glass is exploited in optical lenses. The best explanation for this bending, or *refraction,* can be found in the famous equations of light and magnetism formulated by the nineteenth-century Scottish physicist James Clerk Maxwell. Every transparent material has an *index of refraction,* derived from Maxwell's equations, that determines, for different angles of incidence, the angle through which light will bend when it passes from a vacuum into the material. Air, for example, has a very low index of refraction, meaning that light scarcely is affected by passing through it, while water, glass, and many familiar transparent solids and liquids have larger indices of refraction, which lead to more pronounced bending of light.

The term *transmittance* describes the tendency of a transparent or translucent object to allow different wavelengths of light to pass through it. Light passes not only through some solids and liquids but also through gaseous phenomena such as fog, mist, and smoke. The contribution of light passing through an object or volume to a viewpoint is a product of its transmittance and the spectral power distribution of the incident illuminant.

Light sources, surface reflectance, and transmittance, therefore, all contribute to the amount of light passing from a point in a scene to a viewpoint. The environment of light, however, has one further subtlety. In the natural world, light will scatter from object to object and its spectral power distribution will change with each colored surface it encounters. The light scattered from a white wall near a blue rug, for example, has a bluish tint because some light scatters from the rug to the wall to the viewpoint. The blue light reflected from the rug effectively acts as another source of illumination.

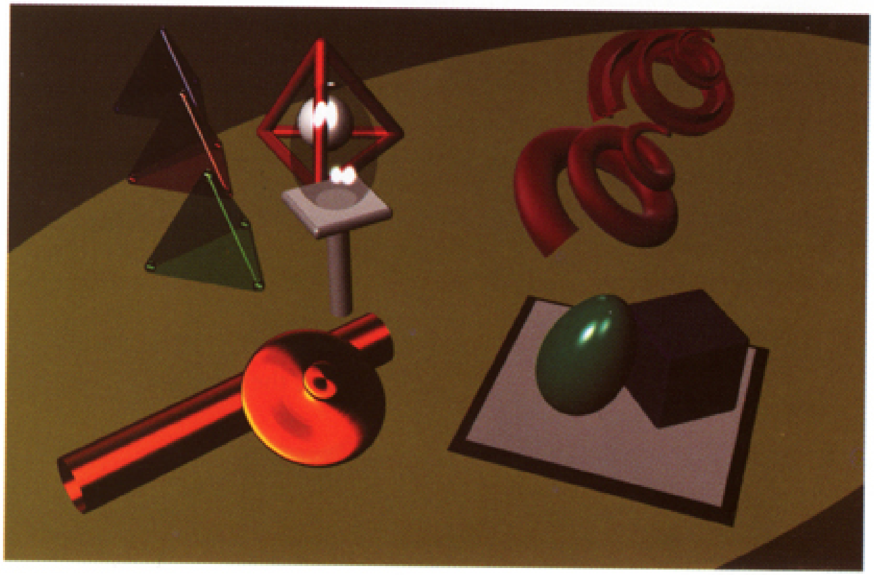

Spotlight In this image, the computer is modeling a second light source which is evident in added highlights. The computer simulates the effect of a spotlight by limiting the influence of the illumination to a certain angular extent. Part of the image, particularly the topmost tetrahedron, is outside the cone of the spotlight illumination.

One has only to glance around a room illuminated by a single lamp to observe that indirect reflections are quite significant. It is unlikely that any surface, even those hidden from the lamp, will be totally dark because the ceilings and walls reflect light and scatter illumination. Light arising from such interreflections is commonly referred to as *ambient* illumination and it constitutes an important component of realistic, perceptually satisfying images. In daylight, scattering from the atmosphere provides a high level of ambient illumination.

Illumination Models

If computer graphics images are to appear verisimilar, the computer must be formally instructed to calculate each of these qualities of light. The first step, therefore, is to propose a mathematical model that describes how light behaves, how it is transported to and reflected by different surfaces. In computer graphics, the terms *illumination model, shading model,* and *reflection model* are used somewhat interchangeably to describe this mathematical model.

An illumination model describes two properties of light: the manner in which light is radiated from light sources and its behavior when it encounters surfaces, including, in some models, interreflections. The computer