



PHOTOREALISTIC RENDERING:

MAKING THE **VIRTUAL**
INTO REALITY

Patricia Daukantas

Studies of the optical properties of translucent materials help computer graphics wizards dream up ever more realistic scenes of fantastic movie characters and industrial design prototypes.



Awizened old man leaps onto a rock. Furrows run across his forehead to his nearly bald pate. His translucent skin hints at the network of blood vessels underneath.

The translucence of his skin—the scattering and diffusion of light in amounts varying with depth—is an important component of what makes this character look so real. Except he isn’t real. He is Gollum, the virtual character brought to “life” in the *Lord of the Rings* film trilogy.

The topic of subsurface scattering has taken hold in the computer graphics community, especially over the past 10 years. Movie studios have driven graphics experts to reach beyond the surface topology of objects and people to illuminate the deeper layers, because a big investment in jaw-dropping special effects can translate to blockbuster success at the box office. But these modeling techniques also have great usefulness in industrial design, especially for producers of big-ticket items that are expensive to prototype out of real materials.

Photorealism is the craft of making a synthetically generated image look like a photograph of the real thing. When designing a scene, the digital artist must specify the geometry of the objects that will appear, assign materials to those objects, and place the light sources. The computer then generates a picture of the scene by calculating how light is reflected from the objects and how it ultimately reaches the eye or the virtual camera.



A face rendered by two different computer models: (left) BRDF, or bidirectional reflection distribution function, and (right) BSSRDF, or bidirectional surface scattering reflection distribution function. The BRDF face looks more like hard plastic than human skin, while the BSSRDF face has a softer appearance. These images appeared in Jensen et al.'s SIGGRAPH 2001 paper, which won a 2003 Academy Award for technical achievement.

Aided by increasingly fast computers, academic computer scientists and film-industry graphics specialists have, over the past decade, developed powerful models for rendering photorealistic images. Studios have built on the theoretical models developed by researchers such as Henrik Wann Jensen, associate professor of computer science at the University of California at San Diego (UCSD). A software engineer at one of the world's top digital visual effects companies called Jensen "one of the pioneers in simulating wider-scale interactions of light and matter." In his research, Jensen seeks to understand what happens when light interacts with translucent materials, including human skin, leaves and other organic matter.

Many times, of course, artists, game designers or movie directors want to depict completely virtual or non-existent things—futuristic vehicles or a humanoid monster with a beard full of octopus tentacles—that are composed of real materials. Yet they want the objects to look realistic enough so that, in the context of their surroundings, viewers believe that they are seeing reality.

The state-of-the-art in realistic computer graphics goes well beyond the images generated by popular computer games. Gaming systems use slightly different techniques to generate pictures at a high frame rate, so their pictures tend to

The audience sees a single frame for only $1/24$ of a second, but huge computing "farms" can take several hours to render each frame. A high-definition still image, which creates an even higher demand for realism, can take much longer.

be somewhat diffuse and textured. Plus, there's the issue of image resolution.

Although many computer games are based on popular Hollywood movies, they fall behind professional films in image quality. The Nintendo Wii handles only 480 pixels of vertical resolution; Microsoft's Xbox 360 and Sony's PlayStation 3 use 1,080 pixels vertically. By contrast, most Hollywood special effects are created with a minimum of 4,000 pixels of vertical resolution. Home game boxes don't have the memory or processing power to generate that kind of resolution. Thus, although the realism

of computer games has dramatically improved over the past two decades, no one would mistake a game's screenshot for a photograph.

Rendering, as done by computer graphics artists and engineers in the movie industry, is the act of compiling all the necessary information about a digitally imaged scene into the final frame of film. The audience sees a single frame for only $1/24$ of a second, but huge computing "farms" can take several hours to render each frame. A high-definition still image, which creates an even higher demand for realism, can take much longer.

Early rendering algorithms for synthetic humans were created with the assumption that human skin was a homogenous material. However, such "surface modeling" describes only the three-dimensional boundaries of objects, as if light bounced off these boundaries or surfaces without interacting with the matter, said Luca Fascione, senior R&D software engineer for Weta Digital in Wellington, New Zealand. Computer models based solely on surface modeling made people's skin look a little like wax (see above left).

Natural materials such as skin or hair have many optically translucent layers, so the interaction of light with the strata, not just the surface, is critical to creating realism. In subsurface scattering, or the diffusion of light through turbid media,

the impinging light travels through the illuminated material and then exits through a different point from which it entered, Fascione said. Enhancing rendering technology with these subsurface scattering effects adds to the realism of virtual skin, leaves, fingernails, teeth and other biological materials.

Skin is obviously a net absorber of visible light; after all, when we look at ourselves, we don't see the underlying bone and muscle. But if we close our eyes when we are in sunshine or a brightly lit room, we still see a bit of red glow, not total darkness. A small percentage of the photons impinging on the eyelid are getting through to the eye.

At first glance, one would think that artists could almost ignore the layered structure of skin and just worry about depicting its correct tone and surface texture. Neurologically, though, humans are wired to be so sensitive to people's appearances that, if anything appears even slightly "wrong," they will notice it right away.

Optical scientists have conducted numerous studies, many with a medical orientation, on the interaction of laser light with skin. Some of the approximations originally derived for photodynamic therapy and laser surgery have led to computational subsurface-scattering models that moviemakers now use to simulate the appearance of skin and hair in *King Kong*, the *Pirates of the Caribbean* trilogy and other films.

Graphic artists could render any type of milk by specifying the relative amounts of riboflavin, protein and fat it contains.

Lorenz-Mie scattering theory

A technique that's widely used in computer graphics is based on Mie theory, also called Lorenz-Mie theory. In 1908, German physicist Gustav Mie published his analytical solution for the scattering of electromagnetic radiation off of homogenous spheres—a theory that has found wide use in atmospheric physics, astrophysics and other situations involving non-absorbing media. (The theory actually predates Mie: In 1890, Danish scientist Ludvig Lorenz wrote a technical report in which he independently developed the same theory; since he wrote his report in Danish, it was forgotten until Mie rediscovered it.)

Studying Lorenz-Mie theory inspired Jensen and colleagues to generalize the theory so that it can address non-spherical particles as well as particles embedded in absorbing media. For example, a glass of milk contains not just water but also vitamin B2 (riboflavin) molecules

that absorb light, and proteins that scatter light. Proteins, with a maximum radius of 150 nm, and fat globules in milk, with diverse radii of up to 10 μm , scatter the shorter wavelengths preferentially, in the effect known as Rayleigh scattering, a subset of the Lorenz-Mie theory. Graphic artists could render any type of milk by specifying the relative amounts of riboflavin, protein and fat it contains (see below).

According to Jensen, white ice is a substance where Mie theory fails completely because it is an absorbing medium that contains air pockets of non-uniform size that scatter the light. (Very clear ice appears faintly blue because it absorbs some red and green light.) An air pocket acts like a particle embedded in this absorbing material.

Learning by rendering milk


Liquid solutions of common substances make an excellent test case for deriving an appearance model based on Lorenz-Mie theory. Researchers can vary concentrations and types of solutes, and they can easily compare their measurements of real-life samples to the results of their computational models. Jensen and two Danish colleagues studied the scattering of light through milk and its solid components: proteins, fat and vitamins.

To start with, the researchers took a laser pointer to a supermarket and shined the pointer on various types of milk and

Henrik Wann Jensen



Photorendered milk and its components. Jensen and colleagues computed the optical properties of the solutions using Lorenz-Mie theory. From left to right, the glasses contain water, water and riboflavin, water and protein, water and fat, skim milk, reduced-fat milk and whole milk.



UCSD student Iman Sadeghi won the grand prize in Jensen's 2007 rendering class with this image. Sadeghi rendered this 1,024 x 768 image with photon mapping and a variety of other software tools.

other beverages. They purchased the ones that seemed to be scattering light in an interesting way, brought them back to the laboratory and measured the optical properties using well-known optical techniques that were new in the computer graphics community.

A glass of skim (or fat-free) milk that is illuminated from one side with a white light in a dark room will look slightly bluish when the light is viewed through the milk in transmission and reddish when it is reflected off the milk.

"What you're seeing is like a little sunset in the glass," Jensen said. The proteins in the milk Rayleigh-scatter the shorter wavelengths of light, while longer, redder wavelengths are able to penetrate to the opposite side of the glass. The effect is most evident in a dark room—you can try this yourself at home with a bright white-light source.

"It is kind of surprising, because we usually think of milk as a white material," Jensen said. "But that's because we usually see it illuminated from all sides."

Replacing skim milk with reduced-fat (2 percent fat) or whole milk (3.25 to 3.5 percent fat) produces a different result.

The broadband scattering effects of the fat dominate the short-wavelength scattering effect of the proteins, so the milk looks whiter again.

This model can be used for ocean water and other translucent materials with heterogeneous particles mixed together. Jensen and his colleagues have also researched the scattering properties of 50 other materials that could be diluted in water.

Teaching rendering algorithms

Graphics rendering involves a lot of computer modeling and coding, and that's why Jensen tells would-be computer graphics artists that a computer-science background helps. After all, professionals are not afraid of making programs with thousands or even hundreds of thousands of lines of code. Fascione, of Weta Digital (*King Kong*, *Lord of the Rings*), has a master's degree in pure mathematics.

At UCSD, Jensen teaches a class called Rendering Algorithms. His students start out by building a system that can generate pictures with ray tracing.

They keep refining their models, and then at the end of the 10-week course Jensen holds a competition to determine who can make the most realistic image. Outside experts are called upon to evaluate the students' work.

Most of the UCSD students were unfamiliar with rendering before they took the course. "The ones who have the nicest pictures are the ones who look very sleepy when they present the results, because it does involve a lot of work to get very realistic results," Jensen said.

Movies and photorealism

Once we know how to simulate light that reflects around a scene, how do we make the individual objects—and the materials of which they are made—look realistic? Of course, various classes of materials—metals, plastics, dielectrics, glasses and so forth—all have different properties, which must be simulated accurately to create a realistic scene. "It's not enough to know that an object is red or blue," Jensen said. "You have to know what kind of material it's made of."

Motion pictures are not necessarily photorealistic; they depict a world the way the director wants it to be, and the scenes are lit a certain manner to convey a mood. Of course, non-human characters such as those in the *Pirates of the Caribbean* series must look reasonably natural in their setting.

The creators of movies set in the distant past, such as *Jurassic Park*, have a certain freedom, since people are not accustomed to looking at long-extinct beasts. However, characters that are closer to humanlike, such as Davy Jones (the ghost-ship captain from the *Pirates of the Caribbean* series) or Gollum (the troll-like character from the *Lord of the Rings* trilogy), creates expectations in the viewer's mind. If Gollum had looked like he had been made of plastic, the audience would have thought, "Hey, this doesn't look real, and something's wrong."

Special-effects artists used the technique from Jensen's 2001 paper to make Davy Jones look highly translucent, almost like an octopus. That

translucence, and the more human-like aspects of his appearance, combined to produce the realism. The entertainment industry widely praised the computer graphics used to create Davy Jones, and the movie *Dead Men's Chest* (the *Pirates* sequel) earned a 2006 Academy Award for best visual effects.

The movie industry has adopted new technologies almost as fast as they come out. Several times over the years, Jensen has been contacted by people from movie companies right after having had a paper accepted by a computer graphics conference. These individuals will reach out to him and say, "We want to use this for our upcoming movie! Can you tell us more about it?" The studios push the technology envelope because it can help them make their rendering pipelines more efficient.

"We have to be able to render relatively fast—not as fast as a video game, but still fast—because we have to render a lot of pictures, and that takes a lot of time," Fascione said. "And, of course, computer time is extremely expensive." Part of the art and science of computer graphics is to find physical approximations that are "visually plausible" but compute efficiently.

According to Jensen and colleagues, solving the full radiative transfer equation can simulate the subsurface transport of light accurately, but at an enormous computational cost. By starting with a dipole diffusion approximation and later a multipole diffusion approximation, they developed a model that computers can use to crank out images in minutes rather than hours.

Pixar Animation Studios developed a popular software platform called PhotoRealistic RenderMan and now sells it to the computer graphics industry. The film studios that use RenderMan have enormous "render farms," with many thousands of CPUs that work constantly to produce scenes for movies.

RenderMan has been one of the most widely used software packages, especially for films that have gone on to win Academy Awards for visual effects (*King Kong*, all three *Lord of the Rings*

RenderMan has been one of the most widely used software packages, especially for films that have gone on to win Academy Awards for visual effects (*King Kong*, all three *Lord of the Rings* installments and the second *Pirates* film).

installments and the second *Pirates* film). A second company, Image Metrics Ltd., also makes facial animation software for movies and games.

Some studios that make special-effects-heavy films have developed their own custom tools. One example is Weta Digital in New Zealand, which developed a software application called Massive that animated enormous numbers of creatures running across fields and battlefields without the need to render each creature individually. According to Jensen, Massive uses artificial intelligence techniques to give the moving crowds verisimilitude.

From the rendering point of view, it's almost easier to render characters or objects that are moving, because the viewers won't have time to study all the details of bodies in motion. "If somebody's running across the screen, you don't have to render every pore in their skin, because nobody's going to see that anyway," Jensen said.

In order to make realistic images, programmer-artists still have to tell RenderMan how the individual objects reflect light, and that's done using "shaders," Jensen said. A shader is a customized extension piece of computer code that tells RenderMan how an object in a scene reflects light, so that RenderMan

can render the individual film frames of that object. Thus, the movie studios need software developers to make the communications bridge between RenderMan and the artists who dreamed up the scenes in the first place.

"I know it sounds tricky, but the software pipeline is actually fairly significant and fairly large at the studios," Jensen said. "Even though they buy RenderMan as their software package, they can have thousands or even tens of thousands or even hundreds of thousands of lines of code that they use on top of that to make effects specific for a given movie."

Looking forward

Thanks to physical measurements, advances in computer modeling and a dash of imagination, subsurface scattering comes closer than ever to creating realistic humanoid characters, plants and biological materials. As computers grow ever faster and researchers find even better virtual models of optical phenomena, the line between "photography" and "virtual objects" will likely become even more blurred. ▲

The author thanks OPN contributing editor Brian Monacelli for his detailed comments.

Patricia Daukantas (pdauka@osa.org) is the senior writer/editor of *Optics & Photonics News*.

[References and Resources]

- >> H.W. Jensen et al. Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH'2001), p. 511–518 (Los Angeles, August 2001).
- >> S.R. Marschner et al. ACM (Assoc. Comput. Mach.) Trans. Graphics **22**(3), 780 (Proceedings of SIGGRAPH'03, San Diego, July 2003).
- >> C. Donner and H.W. Jensen. ACM (Assoc. Comput. Mach.) Trans. Graphics **24**(3), 1032 (Proceedings of SIGGRAPH 2005, Los Angeles, July 2005).
- >> C. Donner and H.W. Jensen. J. Opt. Soc. Am. A **23**, 1382 (2006).
- >> N. Joshi et al. Opt. Lett. **31**, 936 (2006).
- >> J.R. Frisvad et al. ACM (Assoc. Comput. Mach.) Trans. Graphics **26**(3), Article No. 60 (Proceedings of SIGGRAPH 2007, San Diego, August 2007).
- >> Check out the online version of this article at www.osa-opn.org for a list of the OSA journal articles that Jensen has cited in SIGGRAPH conference proceedings.