

Measuring and Modeling the Appearance of Finished Wood

Stephen R. Marschner

Stephen H. Westin

Adam Arbree

Jonathan T. Moon

Cornell University

Abstract

Wood coated with transparent finish has a beautiful and distinctive appearance that is familiar to everyone. Woods with unusual grain patterns, such as tiger, burl, and birdseye figures, have a strikingly unusual directional reflectance that is prized for decorative applications. With new, high resolution measurements of spatially varying BRDFs, we show that this distinctive appearance is due to light scattering that does not conform to the usual notion of anisotropic surface reflection. The behavior can be explained by scattering from the matrix of wood fibers below the surface, resulting in a subsurface highlight that occurs on a cone with an out-of-plane axis. We propose a new shading model component to handle reflection from subsurface fibers, which is combined with the standard diffuse and specular components to make a complete shading model. Rendered results from fits of our model to the measurement data demonstrate that this new model captures the distinctive appearance of wood.

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

Keywords: wood, fibers, anisotropic reflection, optical scattering, rendering

1 Introduction

The distinctive appearance of finished wood is commercially important. Wood is used for fine furnishings and decorative accents in homes, business places, and luxury automobiles because of its rich, attractive appearance. Hitherto, wood has been represented in computer renderings using color texture maps, sometimes combined with spatially varying surface gloss. Such models duplicate the appearance of simulated wood grain such as that commonly used for plastic laminate table tops, but they fail to reproduce the particular depth and richness of finely finished wood.

The distinctive optical behavior of wood is most dramatic in woods with unusual grain patterns, or *figures*, such as tiger, burl, and birdseye. In boards of figured wood, irregularities in the growth of the tree result in wood fibers whose angle with the surface varies substantially across the board. This causes spatial variation, not just in color, but also in the directional scattering characteristics. In straight-grained wood, the same kind of reflection occurs, but with subtler spatial variation. This paper examines the scattering behavior that underlies the appearance of wood: we predict it from the structure, verify it experimentally, and propose what we believe to be the first shading model that reproduces wood’s distinctive appearance, as we demonstrate in side-by-side comparisons.

Previous models and simulations of anisotropic surfaces have dealt with materials, such as brushed metal, whose appearance is

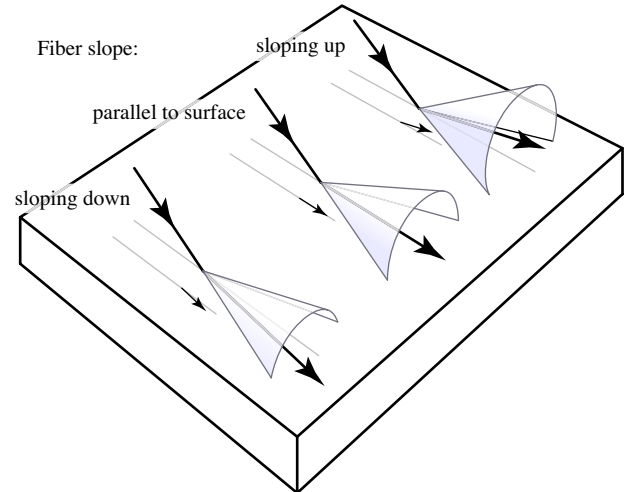


Figure 1: The basic shape of the *subsurface* reflection from fibers inclined downward, parallel to the surface, and inclined upward. Models for anisotropic *surface* reflection produce results similar to the center drawing.

produced by surface reflection. Wood is different from these materials because its appearance is mainly caused by subsurface reflection. The characteristic highlight in wood, like an anisotropic surface highlight, emerges on a cone. But the basic result of this paper is that the axis of the cone for this *subsurface highlight* is not constrained to lie in the surface, as it generally is for surface reflection (Figure 1). Because the axis is determined by the wood’s three-dimensional grain direction, it can and does point into or out of the surface and it varies from one part of the surface to another. It is the separation of this colored subsurface highlight from the white surface highlight that distinguishes wood from other anisotropic materials. It is the variation in fiber inclination across the surface that gives wood—especially figured wood—its distinctive appearance.

In Section 3 we present new light reflection measurements that demonstrate how well this conceptual framework fits the actual reflectance of wood, and in Section 4 we propose a new shading model that captures the essence of wood’s appearance for rendering. Finally, in Section 5 we demonstrate that this model can be fit to the measurements, producing a four-parameter BRDF texture that can be used to make realistic renderings of the measured samples.

1.1 Prior work

In computer graphics, most work on wood has focused on how to automatically generate a color texture map to control the Lambertian component of a dichromatic reflectance model. Wood grain is a standard application for 3D texturing [Peachey 1985; Lewis 1989; Lefebvre and Poulin 2000], which is quite successful in generating typical patterns due to growth rings. The same texture functions may also be used to adjust the parameters of the specular reflection

on the surface, or to create a bump map that simulates dips in the surface due to finish absorbing into more porous parts of the wood. However, to our knowledge no one has modeled the light scattering properties of the material itself.

The literature on wood contains a great deal of information about the structure of wood from the macroscopic to the molecular level [Lewin and Goldstein 1991; Wilson and White 1986]. Spiral grain and other anomalous growth patterns are well documented [Harris 1989], and it is widely recognized that nonuniform fiber directions give rise to decorative effects. These researchers have focused on the biology of wood growth and the effects of grain structure on the mechanical properties of wood. Light diffusion in wood is used to automatically identify grain direction and detect defects in wood [Nyström 2003], and the optics of the separate wood fibers is of interest as it pertains to paper manufacturing [Saarinen and Muinonen 2001]. However, all this research has been motivated by the needs of the wood and paper industries, and to our knowledge no one has examined angular scattering in a way that is relevant to appearance.

We report measurements of the bidirectional reflectance distribution function (BRDF) of wood, including its spatial variation across the surface. Our measurement approach is similar to prior work on measuring spatially varying BRDFs [Dana et al. 1999; McAllister 2002; Gardner et al. 2003]. Like the measurements of McAllister and Gardner et al., our data is dense enough that we can independently fit a model at every point. The difference in our work is that we are examining a different kind of material and observing a previously unstudied mode of reflection.

This paper examines scattering that occurs below the surface, and prior work both on BRDFs [Hanrahan and Krueger 1993; Wolff et al. 1998] and on spatial subsurface diffusion [Stam 1995; Jensen et al. 2001] has also considered subsurface scattering. That work focuses mainly on diffuse reflection or on isotropic, rather smooth, BRDFs, and was based on standard volume scattering. In contrast, in this paper we model a strongly directional, anisotropic, subsurface effect based on a completely different mechanism.

The primary feature of subsurface reflection in wood is strong anisotropy. Many anisotropic reflection models have been proposed, ranging from physically based models for particular types of microgeometry [Poulin and Fournier 1990; He et al. 1991] to microfacet models [Ashikhmin and Shirley 2000] to more empirical or qualitative models [Kajiya 1985; Ward 1992; Lu et al. 2000]. In all cases, the reflection is assumed to be caused by long, narrow features on the surface. Barring asymmetric microstructure, long parallel features present local surface normals that are predominantly perpendicular to the axis of elongation. Because of this, light will reflect into a cone at the specular angle, and this is the basic prediction of all the aforementioned models. Note that Ashikhmin and Shirley's [2000] model does allow asymmetric microstructure, and that paper discusses a model for surface reflection from velvet that actually behaves a bit like subsurface reflection from wood. Also, Pont and Koenderink [2003] observed a double specular reflection from woven materials that, like the phenomena modeled in this paper, is caused by reflection from fibers not parallel to the surface.

Our model is related to recent work on reflection from hair [Marschner et al. 2003], which established the basic geometry and radiometry of reflection and refraction in transparent fibers. Wood, as we will see, is very much like hair turned inside out: instead of a bundle of near-parallel dielectric fibers in air, we have a solid block of dielectric filled with near-parallel air tubes. For this reason some of the results of that work can also be applied to wood.

2 Wood anatomy

Wood has a complex structure that has been studied extensively with regards to its biological function in trees and its mechanical properties as timber. Starting from an understanding of wood's




Figure 2: An electron micrograph of hardwood, showing narrow vertical fibers, large vessels, and horizontal rays. [NC Brown Center for Ultrastructure Studies, SUNY College of Environmental Science and Forestry, Syracuse, NY]

structure, we can make some predictions about how it will reflect light.

2.1 Structure and function

In living trees, wood serves to conduct water from the roots to the stems and leaves, to store starch, and to support the tree. The anatomy of wood follows from these functions: it is composed mainly of highly elongated cells that perform the conduction and support roles, resulting in a highly anisotropic structure. The walls of the cells are thick (for support) and composed mainly of cellulose, hemicellulose, and lignin. [Lewin and Goldstein 1991]

The three-dimensional structure of wood (Figure 2) includes both longitudinal cells, running along the trunk, and transverse cells, which form *rays* that are oriented radially. The longitudinal cells have a yearly variation in density and color that leads to concentric *growth rings*. The rays run perpendicular to the growth rings, typically account for 10% to 20% of the volume, and are generally more prominent in hardwoods than softwoods. Hardwoods also contain large longitudinal *vessels* that can be as wide as 0.3 mm and are often quite visible. [Lewin and Goldstein 1991]

It is widely recognized that the color texture seen on a cut face of wood represents the intersection of a plane with the three-dimensional structure of growth rings and rays. But at the same time, the angle at which the fibers meet the slicing surface creates an orientation texture, with a three-dimensional fiber direction that varies across the surface. These two types of textures are used in our shading model for wood.

2.2 Optical scattering predictions

When wood is dried, the interiors of the cells become filled with air. Following previous work on scattering from individual wood fibers that make up paper [Saarinen and Muinonen 2001], we assume that the cell walls, despite their complex structure, can be approximated as a homogeneous dielectric material with refractive index 1.55. For convenience we will refer to this material as cellulose, even though it has other constituents. This means that the interior of wood, optically, is a block of cellulose (which behaves much like glass) with parallel air tubes running through it.

As light travels through the wood, it encounters air-cellulose interfaces that are parallel to the fiber direction. As with reflections from fibers [Kajiya and Kay 1989], the light that reflects from these




Figure 3: Light reflecting from a matrix of subsurface fibers. The reflection from the fibers has equal incident and reflected angles, but because the fibers are tilted the angles are different at the surface.

interfaces is confined to a cone that is centered on the fiber direction and contains the opposite of the incident direction—that is, all the reflected rays retain the same inclination to the fiber axis (Figure 7). As is the case with refracted paths through the solid dielectric cylinders that make up hair [Marschner et al. 2003], the geometry of Snell’s law guarantees that rays that refract from cellulose to air and back to cellulose will also maintain the same inclination. So once a ray has entered the wood, it can change its direction within the specular cone, but it cannot change its inclination—the angle between the ray direction and the fiber direction remains constant.

Because these interfaces are highly reflective from the dielectric side, and because most light will interact with several fibers before it escapes the surface, we assume that by the time the light emerges it is spread uniformly around the specular cone. Because the fibers are not exactly parallel and the interfaces are not exactly smooth, the reflected light will be spread out somewhat from the ideal cone.

To summarize, the key prediction is that the power in a beam of light that enters the subsurface of wood will emerge spread out on a cone that has the same inclination to the fiber direction as the incident beam. That is, we expect to see a bright reflection approximately when

$$\mathbf{v}_r \cdot \mathbf{u} = -\mathbf{v}_i \cdot \mathbf{u}$$

where \mathbf{v}_i and \mathbf{v}_r are the incident and reflected directions, and \mathbf{u} is the fiber direction.

When wood is planed and sanded to produce a flat surface, the result is a layer of partial (and damaged) cells on the surface, with the intact structure below the surface. The first interface presented to incident light is therefore irregular, and much of the light is reflected diffusely. This white surface scattering is the reason that the color of unfinished wood is light and unsaturated. When varnish or other clear finish is applied to the wood, this first surface is effectively eliminated because the refractive index of the finish is close to that of the wood. For this reason, the distinctive reflection behaviors of wood, which arise from the more organized structure of the subsurface, are much more prominent in finished wood: one of the reasons, aside from protecting the wood from moisture, for applying clear finishes. So for finished wood we expect to see the subsurface highlight more clearly than in unfinished wood.

One final modification is required for surfaces that are coated with a finish. Light is refracted as it enters the surface and again as it leaves. This means that light scattered from global direction \mathbf{v}_i into global direction \mathbf{v}_r actually interacts with the fibers at the modified directions $s(\mathbf{v}_i)$ and $s(\mathbf{v}_r)$, where $s(\mathbf{v})$, given by Snell’s Law, is the refracted direction inside the material corresponding to the direction \mathbf{v} in air. So when \mathbf{v}_i and \mathbf{v}_r are outside a layer of finish the previous equation needs to be modified to:

$$s(\mathbf{v}_r) \cdot \mathbf{u} = -s(\mathbf{v}_i) \cdot \mathbf{u}. \quad (1)$$

The geometry of this situation for a path that remains in the vertical plane is shown in Figure 3.

3 Measurements

To learn about the appearance of wood surfaces, we made detailed spatially varying BRDF measurements on samples of finished hardwoods. The four samples we use to illustrate this paper include American walnut, maple with a curly or “fiddleback” figure, and African padauk with a strong interlocked grain figure. The padauk and one sample of walnut were finished with a high-gloss varnish; the other sample of walnut was finished with a satin lacquer; and the maple was finished with a high-gloss lacquer. The three woods provide examples of relatively straight grain (walnut), curly grain (maple) which takes the form of sinusoidal variations in direction *along* the grain, and interlocked grain (padauk) in which the direction varies radially, *across* the grain.

We carried out the measurements using a Spherical Gantry, a four-axis motion device that can position a camera and a light source independently anywhere on the hemisphere under computer control. Using this apparatus we can access any combination of light source and viewing directions, except for when the viewing direction is within about 12° of the source direction, when the camera will shadow part of the sample. Further details of the apparatus are given in the appendix.

We captured two different types of datasets. For the first type, we kept the camera fixed looking perpendicular to the surface and moved the light source to positions with incidence angle up to 75° , on a 31 by 31 grid spaced uniformly with respect to solid angle using Shirley and Chiu’s [Shirley and Chiu 1997] concentric mapping. We made this kind of measurement for the varnished walnut, lacquered curly maple, and varnished padauk samples. These datasets have the same sampling density in all frames, and they are useful for studying the variation in properties across the surface and for identifying parameters of the model. Figure 4 shows the images for four light positions for each of the three samples. It is convenient to use the reciprocal interpretation of the reflection geometry, interchanging the source and detector to think of these measurements as capturing the scattering diagram from every surface point for illumination in the normal direction. Figure 5 shows the intensity over the hemisphere for each of four surface points on each sample.

The second type of dataset we captured consists of a lower resolution grid for each of several camera positions. For space reasons we show only the lacquered walnut, which consists of 11 by 11 grids for each of four camera positions looking along the grain direction (indicated in Figure 6). These datasets provide more information about the four-dimensional BRDF and let us see changes in reflection with incidence angle. Figure 6 shows images for one source position and all four camera positions, and hemispheres for one surface point and all four camera positions.

The results of the measurements are consistent in overall form across the samples we measured, indicating that the basic reflection mechanism is the same in these different types of wood. There is, as expected, a white surface highlight (most visible in the moving-camera data) on the finished surfaces in the specular direction, and its magnitude increases with incidence angle. There is a colored diffuse component that contributes approximately uniformly to the entire hemisphere. But the remarkable feature of these reflection functions is a very distinct subsurface highlight that is dramatically anisotropic. However, this subsurface highlight does not lie on a cone with the same inclination to the surface tangent as the incident direction, as the highlight from an anisotropic surface reflection must do. Instead, it lies on a different cone, and the cone varies across the surface even for a single illumination direction.

For instance, in data taken with a fixed, overhead camera, the reflection from a grooved surface would occur in a vertical plane




Figure 4: Selected images with fixed overhead view and varying light positions. The light position for each image is marked on the hemisphere to its left. The three samples are varnished walnut, lacquered curly maple, and varnished padauk. The maple has undulations in the grain that lead to shifting horizontal highlights, while the padauk has interlocked grain, which leads to shifting vertical highlights.




Figure 5: Plots of reflected intensity over the incident hemisphere (projected to the equatorial disc) for selected surface points. The point for each hemisphere is marked on the image to its left. We can conclude that the grain in the walnut sample tilts slowly into the paper as we move from bottom to top, while the maple and padauk exhibit periodic changes along the vertical and horizontal directions, respectively.




Figure 6: Non-perpendicular views of the second walnut sample. The top row shows rectified images for the four camera positions under a single illumination direction. Note the reversal of contrast across the region near the knot. The second row shows reflected intensity over the incident hemisphere for each camera direction at a single surface point.

passing exactly through the surface normal, but on the padauk sample the anisotropic reflection component can be seen to swing back and forth by as much as 30 to 40 degrees to either side of the normal as we look at different parts of the grain pattern. In less exotic woods, the angle changes slowly across most of the surface, as in the walnut samples, but it can change quickly near irregularities, such as the knot visible in the lacquered sample (Figure 6).

The color and intensity of the subsurface reflection change rapidly across the surface, with the largest changes associated with moving from one part of the anatomy to another. In all these samples the slicing planes are approximately tangent to the growth rings (they are plainsawn, rather than quartersawn, boards), so the rays are at a steep angle to the surface and tend to appear dark (Figure 5(b)4). Vessels in the wood also tend to cause elongated pits in the surface, which appear as dark areas with weaker anisotropy. Changes in pigmentation between the earlywood and latewood parts of the growth rings also cause changes in intensity: Figure 5(a)2 shows a point in the denser, darker latewood that has a weak subsurface highlight.

4 A shading model for wood

In Section 2 we discussed the basic phenomenon of reflection from subsurface fibers, with the prediction that the fiber highlight is observed when the refracted incident and exitant directions have the same inclination to the fiber direction. In this section we extend this to a physically motivated phenomenological reflectance model. Our model is motivated by analyzing the physical structure, but in the interest of simplicity it does not attempt to account exactly for the radiometry.

The overall form of the model for the BRDF, f_r , is as follows:

$$f_r(\mathbf{v}_i, \mathbf{v}_r) = f_s(\mathbf{v}_i, \mathbf{v}_r) + T_i T_r (\rho_d + f_f(\mathbf{u}, \mathbf{v}_i, \mathbf{v}_r))$$

The two terms represent surface and subsurface reflection, respectively, and the two parts of the subsurface term represent uniform diffuse reflection and fiber reflection.

The surface BRDF, f_s , is a standard dielectric interface reflection, broadened due to surface roughness if appropriate. We have not included the surface term in the direct comparisons with data, and in the renderings we have used a simple Phong model.

The two subsurface terms are ρ_d , the uniform diffuse reflectance, and f_f , the fiber reflection BRDF. Both are attenuated by transmission into and out of the surface of the finish:

$$T_i T_r = (1 - F(\mathbf{v}_i))(1 - F(\mathbf{v}_r))$$

where $F(\mathbf{v})$ is the Fresnel reflection factor for a direction \mathbf{v} on the air side of the interface. This makes the diffuse term into a Fresnel-corrected diffuse term like those used by Ashikhmin and Shirley [2000] and Wolff [1998], and it causes both subsurface terms to drop to zero at grazing incidence or grazing exitance.

The subsurface fiber reflection term assumes that the incident light that passes through the surface is distributed uniformly around the specular cone for the local fiber direction. Because of imperfections in the structure, the reflection will be spread out from the cone, and we model that spread as a Gaussian:

$$f_f(\mathbf{u}, \mathbf{v}_i, \mathbf{v}_r) = k_f \frac{g(\beta, \psi_h)}{\cos^2(\psi_d/2)}$$

where

$$\begin{aligned} \psi_i &= \sin^{-1}(s(\mathbf{v}_i) \cdot \mathbf{u}) & \psi_d &= \psi_r - \psi_i \\ \psi_r &= \sin^{-1}(s(\mathbf{v}_r) \cdot \mathbf{u}) & \psi_h &= \psi_r + \psi_i \end{aligned}$$

and $g(\sigma, x)$ is a normalized Gaussian with zero mean and standard deviation σ . The fiber reflection color k_f and the highlight width β are parameters of the model. The angles ψ_i and ψ_r are the inclinations of the incident and reflected directions to the fiber direction (they are signed angles and are zero when the direction is perpendicular to the fiber). The angle ψ_h is the *half angle*, the inclination of the bisector of $s(\mathbf{v}_i)$ and $s(\mathbf{v}_r)$, and when it is zero $s(\mathbf{v}_r)$ is in the specular cone of $s(\mathbf{v}_i)$. The angle ψ_d is the *difference angle*, the total inclination difference between $s(\mathbf{v}_i)$ and $s(\mathbf{v}_r)$.

The factor of $\cos^{-2}(\psi_d/2)$ is the geometry factor for specular reflection into a cone: the projected solid angle of the differential cone into which the incident light is channeled is proportional to $\cos^2 \psi$ where ψ is the reflection angle. We use $\psi_d/2$ to extend this notion to non-ideal specular reflection while maintaining reciprocity.

4.1 Fitting the model to data

We fit our reflection model independently at each pixel, using all the illumination directions for a single camera direction. The large




Figure 8: Fitting results for padauk. This figure shows the fitted spatially varying BRDF model evaluated for conditions that match the measured images shown in Figures 4 and 5.

number of fitting operations necessitates a simple fitting procedure. We use a two-step process: first we estimate the fiber axis, then we estimate the diffuse and fiber colors (ρ_d and k_f). We do not compute a separate highlight width per pixel because we have found a fixed highlight width for each type of wood is sufficient for an appearance match.

After excluding a region around the specular direction to avoid the surface highlight, we assume the brightest n samples are near the subsurface specular cone. Each of these samples has an incident direction and an exitant direction, and our model relates these vectors to the fiber direction: the subsurface incident direction $s(\mathbf{v}_i)$ and the subsurface reflection vector $s(\mathbf{v}_r)$ should have nearly equal and opposite dot products with the fiber direction \mathbf{u} . This leads to a simple linear system:

$$\begin{aligned} \mathbf{u}^* &= \underset{\mathbf{u}}{\operatorname{argmin}} \sum_k (s(\mathbf{v}_{i,k}) \cdot \mathbf{u} + s(\mathbf{v}_{r,k}) \cdot \mathbf{u})^2 \\ &= \underset{\mathbf{u}}{\operatorname{argmin}} \|A\mathbf{u}\|^2 \end{aligned}$$




Figure 7: Geometry for fiber reflection model. Specular reflection from the subsurface fibers occurs when the refracted reflection direction, $s(\mathbf{v}_r)$, is in the cone centered on the fiber axis \mathbf{u} and containing the opposite of the refracted incident direction, $-s(\mathbf{v}_i)$.

where row k of the n -by-3 matrix A is $(s(\mathbf{v}_{i,k}) + s(\mathbf{v}_{r,k}))^T$. The solution \mathbf{u}^* to this problem is the third singular vector of A , corresponding to its smallest singular value. The results in this paper were computed using 20% of the total number of samples for n and weighting the rows of the system by the fourth power of the intensity.

Once we have computed the fiber axis, we use simple heuristics to estimate the two colors: the diffuse color is the 35th percentile of the valid data points, and the fiber color is the 97th percentile.

In the fitting process we avoid shadowed samples and samples contaminated by surface reflection by ignoring all samples that are within an “exclusion angle” of retroreflective or specular configurations. It is particularly important to exclude specular samples because they will bias the axis by introducing bright points that are not on the correct cone.

5 Results

We applied the fitting process described in the previous section to derive maps of \mathbf{u} , ρ_d , and k_f from the overhead datasets; a constant β was used for each sample.

We used only the overhead data because it provides enough information to determine all the parameters and it has uniformly high resolution, unlike the oblique views. The fits were performed on subsampled versions of the full data, with 1/3 megapixel images and 256 light directions, and the computation took about 30 minutes using a Java program on a current PC.

To illustrate the fit of the model to the data, we evaluated the model using the appropriate illumination and viewing directions to reconstruct the original frames. Figure 8 shows representative images and hemispheres computed from the model of the padauk sample; compare to the original data in Figures 4 and 5. Although there are detectable differences, the visual match is very good and, just as important, the reconstructed images are plausible. More views for all samples are available in the accompanying video.

A more challenging test of the model is whether it can predict the images from camera positions other than the overhead view it was fit to. Figure 9 shows representative images and incident hemispheres for non-orthogonal views of the walnut sample that were computed from the model parameters determined from the overhead view. These results can be compared with the measurements

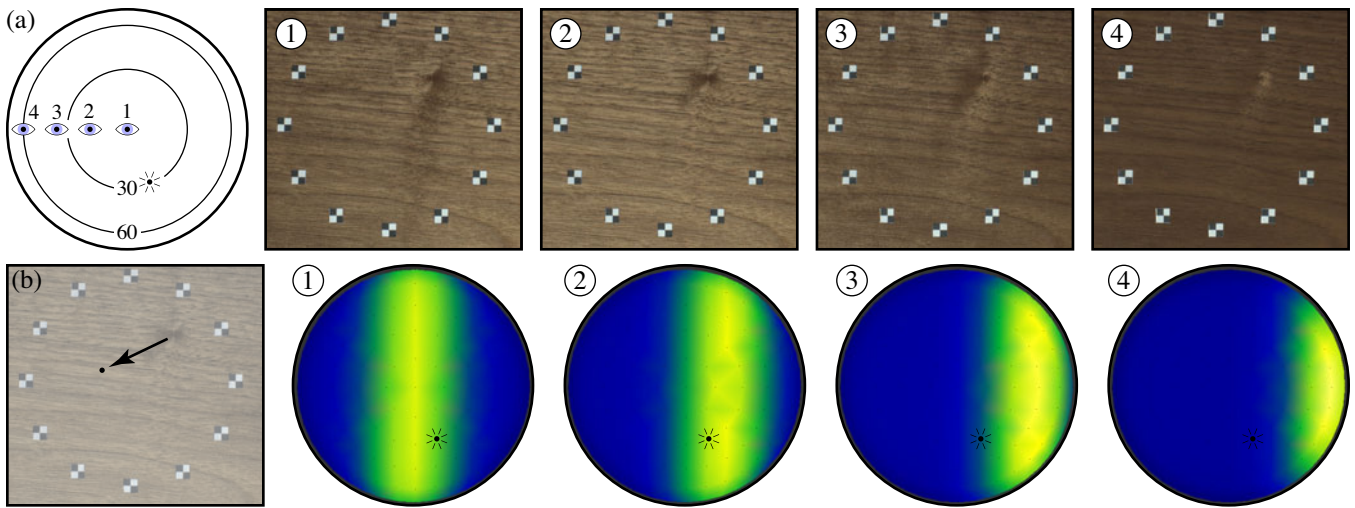


Figure 9: Extrapolating from overhead views to oblique views. The model was fit to only the data with the camera directly overhead, and it was evaluated to reconstruct images that match the oblique-viewing measurement images shown in Figure 6. The contrast reversal near the knot is correctly predicted. Note that the plots in Figure 6 include surface reflection, which is not included here.

for the same sample in Figure 6. The reconstructed images continue to be a good qualitative match.

The video shows another extrapolation: the model (still fit to overhead views) is used to render an animation of each sample rotating on a turntable, and the result is compared with actual video of the sample under the same conditions. In these animations the camera is neither overhead nor aligned with the fiber direction, further demonstrating the model’s ability to generalize.

The video also compares our model to Ward’s surface-oriented model [1992]. Ward’s model is supplied with the parameters derived from fitting the new model; the fiber color (adjusted by a constant factor) becomes Ward’s specular coefficient, and the fiber direction is projected into the plane of the surface to determine the tangent frame. Ward’s roughness is chosen to match the width of the observed highlight. Note that we are applying Ward’s model to the subsurface fiber reflection, whereas he used it to model the highlight from an anisotropic varnish surface. The results show effects of variation in the projected axis direction, but they do not match the behavior of the real wood, demonstrating the importance of 3D anisotropy directions.

The renderings in Figure 10, made with RenderMan, illustrate the same comparison. Ward’s model captures only the changes in appearance due to in-plane variation in the fiber direction. For example, the padauk and maple appear much more uniform in the top row, and the contrast reversal around the knot in the walnut sample is completely missed.

6 Conclusion

Wood, which is valued as a decorative material and appears commonly in rendered scenes, has hardly been studied for directional optical scattering. Our new, high resolution measurements show that wood behaves differently from other anisotropic materials, confirming the intuition that wood has a distinctively different appearance from other materials. We have proposed a new reflection model that describes an entirely new kind of scattering: strongly directional, anisotropic subsurface scattering. The simple model of a classical fiber reflection from a tilted axis viewed through a refracting surface matches the data well and correctly predicts how the highlight moves with varying incidence angle. While further

refinements remain to be made, we can conclude that we have described the basic cause of the distinctive appearance of wood.

7 Acknowledgements

The authors would like to thank Jim Ferwerda for providing wood samples, Piti Irawan for finding many valuable references, and Ken Torrance for providing guidance and laboratory space for sample preparation. This work was supported by an NSF CAREER grant (#0347303), and the second author was supported by an NSF ITR grant (#0113851). Pixar Animation Studios donated the RenderMan software. The Cornell Spherical Gantry is based on a design by Marc Levoy, who generously provided the plans, and was constructed using startup funds from Cornell University.

References

ASHIKHMIN, M., AND SHIRLEY, P. S. 2000. An anisotropic phong brdf model. *Journal of Graphics Tools* 5, 2, 25–32.

DANA, K. J., VAN GINNEKEN, B., NAYAR, S. K., AND KOENDERINK, J. J. 1999. Reflectance and texture of real-world surfaces. *ACM Transactions on Graphics* 18, 1, 1–34.

GARDNER, A., TCHOU, C., HAWKINS, T., AND DEBEVEC, P. 2003. Linear light source reflectometry. *ACM Transactions on Graphics* 22, 3 (July), 749–758.

HANRAHAN, P., AND KRUEGER, W. 1993. Reflection from layered surfaces due to subsurface scattering. In *Proc. SIGGRAPH 93*, 165–174.

HARRIS, J. M. 1989. *Spiral Grain and Wave Phenomena in Wood Formation*. Springer-Verlag, Berlin.

HE, X. D., TORRANCE, K. E., SILLION, F. X., AND GREENBERG, D. P. 1991. A comprehensive physical model for light reflection. *Computer Graphics (Proc. SIGGRAPH)* 25, 4, 175–186.

JENSEN, H. W., MARSCHNER, S. R., LEVOY, M., AND HANRAHAN, P. 2001. A practical model for subsurface light transport. In *Proc. SIGGRAPH 2001*, 511–518.



Figure 10: Top row: three measured samples rendered using Ward’s model for the subsurface reflection, with parameters derived from our fitting results. Bottom row: samples rendered with the new model. Each is shown in two different orientations. Left to right: curly maple, padauk, walnut. See the accompanying video for a more complete comparison.

KAJIYA, J. T., AND KAY, T. L. 1989. Rendering fur with 3D textures. In *Proceedings of SIGGRAPH 1989*, 271–280.

KAJIYA, J. T. 1985. Anisotropic reflection models. *Computer Graphics (Proc. SIGGRAPH) 19*, 3, 15–21.

LEFEBVRE, L., AND POULIN, P. 2000. Analysis and synthesis of structural textures. In *Graphics Interface*, 77–86.

LEVOY, M., 2004. The Stanford Spherical Gantry. <http://graphics.stanford.edu/projects/gantry>.

LEWIN, M., AND GOLDSTEIN, I. S., Eds. 1991. *Wood Structure and Composition*. Dekker, New York.

LEWIS, J.-P. 1989. Algorithms for solid noise synthesis. In *Proc. SIGGRAPH 89*, 263–270.

LU, R., KOENDERINK, J. J., AND KAPPERS, A. M. L. 2000. Specularities on surfaces with tangential hairs or grooves. *Computer Vision and Image Understanding 78*, 320–335.

MARSCHNER, S. R., JENSEN, H. W., CAMMARANO, M., WORLEY, S., AND HANRAHAN, P. 2003. Light scattering from human hair fibers. *ACM Transactions on Graphics (Proc. SIGGRAPH) 22*, 3, 780–791.

MCALLISTER, D. K. 2002. *A Generalized Surface Appearance Representation for Computer Graphics*. PhD thesis, University of North Carolina at Chapel Hill.

NYSTRÖM, J. 2003. Automatic measurement of fiber orientation in softwoods by using the tracheid effect. *Computers and Electronics in Agriculture 41*, 1–3 (Dec.), 91–99.

PEACHEY, D. R. 1985. Solid texturing of complex surfaces. In *Proc. SIGGRAPH 85*, 279–286.

PONT, S. C., AND KOENDERINK, J. J. 2003. Split off-specular reflection and surface scattering from woven materials. *Applied Optics 42*, 1526–1533.

POULIN, P., AND FOURNIER, A. 1990. A model for anisotropic reflection. *Computer Graphics (Proc. SIGGRAPH) 24*, 4, 273–282.

SAARINEN, K., AND MUINONEN, K. 2001. Light scattering by wood fibers. *Applied Optics 40*, 28 (Oct.), 5064–5077.

SHIRLEY, P., AND CHIU, K. 1997. A low distortion map between disk and square. *Journal of Graphics Tools 2*, 3, 45–52.

STAM, J. 1995. Multiple scattering as a diffusion process. In *Eurographics Rendering Workshop 1995*, 41–50.

WARD, G. J. 1992. Measuring and modeling anisotropic reflection. *Computer Graphics (Proc. SIGGRAPH) 26*, 2, 265–272.

WILSON, K., AND WHITE, D. J. B. 1986. *The Anatomy of Wood: Its diversity and variability*. Stobart and Son, London.

WOLFF, L. B., NAYAR, S. K., AND OREN, M. 1998. Improved diffuse reflection models for computer vision. *International Journal of Computer Vision 30*, 1, 55–71.

A Measurement details

The measurement apparatus was built on the Cornell Spherical Gantry, a general-purpose four-axis motion system designed for scattering measurement. The gantry is similar in design to the Stanford Spherical Gantry [Levoy 2004], allowing precise placement of an illuminator and a camera, mounted on two articulated arms, over the sphere of directions looking inward at a sample.

The illuminator was a tungsten source directed through a fiber optic light guide, and we used a cooled 1.3-megapixel monochrome digital scientific camera, with a liquid crystal RGB filter between the lens and camera. Multiple exposures were used to extend the effective dynamic range, though there are still some saturated pixels in the surface specular reflection, which appear black in the data. The bare end of the light guide creates a circular source subtending a solid angle of 2×10^{-5} sr. For the overhead datasets the camera’s aperture, at $f/5.6$, subtended 8×10^{-5} sr; for the multi-view datasets we used $f/11$ to obtain the required depth of field, resulting in a solid angle of 2×10^{-5} sr. The distance to the source was about 1 m, and the distance to the camera was about 0.5 m. The width of the measured area of the sample was about 100 mm, so the direction to the camera varies up to 7° from the nominal direction at the corners of the sample, and the direction to the light source varies up to 4° from nominal. In our analysis of the data we use the true camera and source directions, resulting in slightly different angles for each point on the surface. The data were not corrected for the camera’s flat-field response or for directional nonuniformity in the light source intensity; calibration tests suggest that these factors together contribute about a $\pm 10\%$ variation across the sample.

In order to interpret the images from different camera positions as spatially varying BRDF measurements, it is necessary to warp them into a common coordinate system, and we attached fiducial markers to the surface to aid in this process. Once the markers were found and correspondence established using an approximate model for the camera motion, the data were resampled into a single overhead view. The same mechanism was used even when the camera was fixed overhead, because in order to prevent shadowing by the arm holding the camera it was necessary to rotate the camera about its axis for some of the images.