



There is a well known property of rough surface reflectance: at grazing angles a rough surface becomes mirror-like. These three images have identical material properties, but near grazing angles, the specular is much more apparent and tends to a more mirror-like reflection. Earlier and simpler BRDF models like the Phong and Blinn-Phong models lack this behaviour: the specular lobe in these models does not tighten and brighten at grazing angles.



It is natural to then ask: do models commonly used for fiber scattering have a similar behaviour?



In this talk we show how all current fiber reflectance models in fact lack a visually-important missing behaviour that is analogous to this grazing-angle/mirror-like rough surface behaviour. We track the cause of this to how these models are written as a sum of separable lobe functions. We then present a new BCSDF derivation using non-separable lobes which we find to be practical and analytic. We present results of the first ground truth MC validation of far-field reflectance from rough dielectric fibers and use this to validate our new derivations. Finally, we report on our findings of applying this new model, and show some results.



The scope of this talk then, is that we perform a new accuracy evaluation of an entire class of complexity compression tools for fiber rendering...



Statistical models are used all the time to compress complexity in rendering. It is essential to avoid unnecessary complexity in order to make rendering a scene possible. For example, we would need to model the water droplets on the left to synthesis that photograph, but on the right, we might specify only a volumetric density of the waterfall, and use a volume scattering model in place of millions of water droplets.



Similarly, for a rough 2D correlated random medium an explicit surface representation is unnecessary at a distance, and we typically use a BRDF. (Exciting work this year at SIGGRAPH has begun to explore the regime between explicit structure and BRDF reflectance for sparkly surfaces).



There is a 3rd class of light transport atomics for mediums with strong 1D correlations – its representation is called the BCSDF and it is applied to curve primitives.



Thus, the focus of this talk is the rendering of thin fibers using a far-field approximation. The reflectance is thus described by a BCSDF, one of several common light transport path 'atoms'.

Two steps of Appearance Modeling



We note that previous comparisons of hair appearance modelling to measured photographs of human hair implicitly include two important elements in one step: first, the proposal of a physical/optical model for a human hair and, in addition, derivation of a BCSDF that approximates the transport of light within the physical model. Monte Carlo lets us easily analyse the accuracy of the second step in isolation by providing ground truth for the physical model (not relying on measurements of real hair).



Far-field models consider a uniform illumination of parallel rays, and predict the total resulting reflectance distributions, as if the fiber were infinitely thin.



Marschner et al. proposed a powerful method to compactly express the complicated behaviour or light paths as they traverse through a fiber.



The key idea is to project the behaviour into the normal plane and look at a 2D problem in flatland.



On the left we follow a light path interacting the a hair fiber. You see that it can support tilted scales that affect the Fresnel refractions and reflections (a property of human hair).



The first interaction that we consider is the first lobe R – for reflection. The right image shows the projection of this path onto the normal plane, where a simple reflection law holds.



the second lobe corresponds to light that transmits twice – and this is called TT.



We carry this forward – created TRT – a colored highlight that sits next to the white R highlight on the fiber, but is offset from it because of the tilted scales.



The Marschner et al. style of expressing fiber reflectance is to split these two longitudinal and azimuthal behaviours apart, work them out in detail, and then stick them back together.



The component distributions are typically denoted M and N.

The Factored Lobe BCSDF family

• Marschner et al. 2003:

0

• Express BCSDF as sum of separable lobes

$$S(\theta_i, \theta_r, \phi) = \sum_{p=0}^{\infty} S_p(\theta_i, \theta_r, \phi)$$

$$S_p(\theta_i, \theta_r, \phi) = M_p(\theta_i, \theta_r) N_p(\theta_i, \theta_r, \phi)$$

To achieve more rich behaviours, most recent BCSDFs have followed Marschner et al. from 2003 and split the total reflectance of the fiber into a sum of different lobes, Sp. Each lobe is then a product of two functions M and N. Let's quickly have a look at what this describes in more detail.

weta



Starting first with the N functions, describing the variation in exitance around each exitant cone.



previous models such as Zinke and weber 2007 produce a very wide variety of shapes, depending on inclination to the fiber, ior, and micro roughness, and our MC validation showed these models to be very accurate. Additional details are in the supplemental material.



However, let's look closely at what the M function describes. Because the surface is rough, interactions with the fiber leave in a variety of cones.



However, previous M functions (typically just a Gaussian or Cauchy distribution) are independent of the azimuthal direction (phi). This means that the longitudinal width of the outgoing distribution of light is fixed as you consider various exitant phi angles.



To test the accuracy of these phi-invariant M functions, we performed a number of Monte Carlo reflectance simulations for rough fibers. These provide ground truth R, TT, TRT, and higher lobes for rough dielectric cylinders with a variety of properties and for a variety of incident angles.



As an example, one such simulation for black hair (R lobe only) is shown here. The horizontal axis of the photon splat image corresponds to changing the phi angle. Here in the simulated data we see the grazing->mirror-like property that we know well for physically-based BRDFs.



To see why this happens consider the various paths that create different portions of the R lobe. In the middle of our plot, phi = 0, this corresponds to striking the center of the hair and reflecting backward.



Here, a small variation in the surface normal can yield a large deviation in exiting angle theta_o, so the lobe is wide (showed here having a large vertical extent).



However, if striking the side of the fiber and grazing forward in the case of back lighting...



deviations in the surface micronormals due to roughness yield very little change in longitudinal angle theta_o, and so the highlight does not widen due to roughness and tightens.



Now considering previous BCSDF models: we see a distinctly different result for MC vs previous models (d'Eon et al. 2011). Here we show three results, the MC, our new model, which I will describe soon, and a previous model. We see that the previous model does not exhibit a tightening/brightening effect for |phi| near pi – the previous lobe is about the same vertical width for any value of phi.



For lower roughnesses, we see another behaviour lacking from previous BCSDFs: the exiting inclination is not constant when the tilted scales perturb the reflection direction (for a similar reason that the lobe is differing widths for various phi values).



The key idea in our new lobe derivations is analytic 'tracing' of the **specular cone angle**: the exiting energy for each mode assuming only the tilted scales, but no roughness.



We form new M functions for rough fibers using the mean and variance of this specular cone function, theta_cone. The derivation more accurately accounts for the micro normal distribution of the rough cylinder surface. The lobe is centered at the mean value of this function, and the lobe width is proportional to this function's first derivative with respect to deviating the micro normal in the longitudinal direction. See the supplementary material for further details.



Where in previous models the scale shift of M was azimuthally invariant (a constant), we arrive at the following expression that accurately matches MC.



The new R lobe width contains a cos(phi/2) term that produces the lobe width variation seen in our MC simulation, that previous models lack, and is added to hair rendering systems efficiently and easily.



Here we show R only – black shiny fibers. The old model and new model look nearly identical for front lit, since we tuned the roughness to match for the front lit case.



Moving the light source behind the hair now, relative to the viewpoint, we see that the old separable-lobe model maintains a similar lobe breadth, whilst the new non-separable model exhibits the lobe tightening/mirror-effect that was see in the ground-truth simulation. We found a MC render of Cook-torrance cylinder identical to our result.



Other lobes are generalised in a similar fashion, but more complicated. It is easy to derive any order using Nusselt sphere coordinates and auto-differentiation, but the results get trig-heavy quite quickly. Also, the visual difference of using more accurate higher-order terms was much more minor in practice. Complete details in the supp pdf.



Light source above, blond hair, all lobes, fairly subtle difference.



Light source above, blond hair, all lobes, fairly subtle difference.



Please see supplementary details for...



Monte carlo lets us analyse the accuracy of BCSDF derivation with great scrutiny. With more accurate BCSDF derivations derived in this work which we found to very closely match ground truth, it is then appropriate to return to comparisons of light transport simulations using these BCSDFs with photographs of human hair. Our initial findings along these lines found significant discrepancies between the simulation and the photographs that were only lessened by widening the TT and TRT lobe widths much much wider that those predicted by our model. Thus, we conclude that the physical model of a rough dielectric cylinder with a homogenous non-scattering interior is not sophisticated enough for human hair.



In summary, we have presented the findings of the first ground-truth study of the accuracy of analytic reflectance models for far-field scattering from rough dielectric cylinders with uniformly absorbing interiors.



Sincere thanks to Weta Digital, especially Joe Letteri and Luca Fascione for supporting and encouraging this research.

