

An Artist Friendly Hair Shading System

Iman Sadeghi^{1,2}

Heather Pritchett¹

Henrik Wann Jensen²

Rasmus Tamstorf¹

¹Walt Disney Animation Studios

²University of California, San Diego



Figure 1: Rendering results using our artist friendly hair shading system. For film production it is important to see that the overall appearance holds up not just for a still image, but also when the hair is animated as shown here.

Abstract

Rendering hair in motion pictures is an important and challenging task. Despite much research on physically based hair rendering, it is currently difficult to benefit from this work because physically based shading models do not offer artist friendly controls. As a consequence much production work so far has used ad hoc shaders that are easier to control, but often lack the richness seen in real hair. We show that physically based shading models fail to provide intuitive artist controls and we introduce a novel approach for creating an art-directable hair shading model from existing physically based models. Through an informal user study we show that this system is easier to use compared to existing systems. Our shader has been integrated into the production pipeline at the Walt Disney Animation Studios and is being used in the production of the upcoming animated feature film *Tangled*.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and realism—Color, shading, shadowing, and texture

Keywords: Hair shading, artist control, single scattering, multiple scattering

1 Introduction

Almost all characters in movies, games and other digitally created content have some kind of hair or fur on their bodies. Our eyes are very sensitive to the appearance of hair and we can observe subtle inaccuracies in its appearance. In fact, hair appearance has been shown to be one of the most important features of avatar personalization [Ducheneaut et al. 2009]. As such it is critical to be able to render good looking hair and fur. In this context it is important to recognize that “good looking” does not necessarily mean scientifically accurate. The appearance is the result of a creative process and the most important criterion is that it is aesthetically pleasing and

that it fits within the universe of the character. This definition subsumes “photo-realistic” as a special case, but in general it is much broader.

Hair rendering is challenging because it requires capturing the complex behavior of light scattering events inside the hair volume which is computationally very expensive. There has been much research on physically based hair rendering, but it remains difficult for artists to really benefit from these results. The main drawback of using physically based shaders in creative endeavors is the lack of suitable controls. Because of this, most production work so far has used ad hoc shaders which are typically more art-directable. However, ad hoc shaders fail to capture the details in light scattering inside the hair volume and often produce inconsistent results under different lighting conditions. Overall this leads to a less rich appearance which then limits the universe in which the characters can live.

In the following we illustrate how physically based shaders fail to satisfy the controllability required in a creative environment. Based on this, we then present a new approach that can produce an art-directable hair shading model using the physical properties of hair fibers. It also handles different lighting conditions while giving full control over all visually important aspects of the hair appearance to the artists. In practice it has been shown to be versatile enough to handle all types of hair as well as fur. Our evaluation shows that using our new shading model, artists tend to achieve the desired appearance more easily compared to both a physically based hair shader and an ad hoc hair shader which has previously been used in feature film production.

The outline of the remainder of the paper is as follows. In the next section we briefly consider related work, while Section 3 defines the controls needed by artists. It also explains why physically based shading models fail to satisfy these needs. In Section 4 we present our new approach for producing an art-directable hair shading model based on existing physically based models. In Section 5 we show how we have applied our approach to both the single scattering and multiple scattering components of hair. We proceed by presenting some rendering results in Section 6. The results of our evaluation are summarized in Section 7, and we end with the conclusion and future work in Section 8.

2 Related Work

Extensive work has been done in the research community to capture the exact behavior of light scattering by human hair fibers. The first prominent work in the field of hair rendering is the classical model of Kajiyama and Kay [1989]. Since then, researchers have investi-

ACM Reference Format

Sadeghi, I., Pritchett, H., Jensen, H., Tamstorf, R. 2010. An Artist Friendly Hair Shading System. *ACM Trans. Graph.* 29, 4, Article 56 (July 2010), 10 pages. DOI = 10.1145/1778765.1778793 <http://doi.acm.org/10.1145/1778765.1778793>.

Copyright Notice

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701, fax +1 (212) 869-0481, or permissions@acm.org.
© 2010 ACM 0730-0301/2010/07-ART56 \$10.00 DOI 10.1145/1778765.1778793 <http://doi.acm.org/10.1145/1778765.1778793>

gated the scattering of light by hair fibers both in the case of single scattering [Kim 2002; Marschner et al. 2003] and multiple scattering [Moon and Marschner 2006; Zinke and Weber 2006; Moon et al. 2008; Zinke et al. 2008]. For a survey on hair rendering please refer to [Ward et al. 2007]. These methods are all concerned with physical correctness of the results and do not consider the controllability of the hair shading model.

Unfortunately, matching a desired appearance by tweaking the physically based parameters (e.g. absorption coefficient, and index of refraction) is a time-consuming and tedious task [Zinke et al. 2009; Bonneel et al. 2009]. There has been some effort to estimate the values of those physically based parameters by analyzing a single photograph [Zinke et al. 2009; Bonneel et al. 2009]. Paris et al. [2008] also presented an image-based rendering method that is able to capture the hair style and hair color of a person by using a complex capturing setup.

All of these methods enable artists to render hair with an appearance similar to a photographic reference. However, they do not provide artists any controls for further adjustments. Also, they can only produce results for which there is already a photographic reference. In practice, art direction often goes beyond what can be captured in a photo from the real world, so it is important for a hair shader to be able to extrapolate beyond the physical range. Furthermore, it is important to realize that a physically accurate model requires equally accurate input to generate realistic results. When the goal is to create aesthetically pleasing imagery rather than doing scientific simulation of light interaction with hair, it may therefore be necessary to allow for non-physical behavior in order to compensate for imperfect input-geometry, animation, or lighting.

Beside physically based shading models, there has been a lot of work on implementing ad hoc shaders in production environments. Goldman [1997] introduced his *fakefur* rendering method by modifying the model of Kajiya and Kay. He provided simple parameters for controlling the amount of light scattering in frontlit and backlit scenes. Other hair rendering methods have been introduced that are based on simple shading models like a Lambertian [Apodaca et al. 2000; Apodaca et al. 2002] or the Kajiya and Kay model [Apodaca et al. 2001; Neulander 2004]. These methods have introduced numerous tricks to handle the backlighting situation better as well as finding a hair tangent based on the surface normal and the hair geometry to get more pleasing results. Petrovic et al. [2005] presented a volumetric representation that was easier for the artist to light. They used the model of Kajiya and Kay for shading. These shaders have more intuitive and easy-to-understand controls and they can produce art-directed appearances. However, as previously mentioned, the results often do not hold up under different lighting conditions and they tend to miss some of the visual complexity which characterizes real hair.

Our work is a form of Participatory Design where all end users of a system are involved in the design process to ensure the usability of the final product [Schuler and Namioka 1993]. Participatory design is a model of User-Centered Design [Norman and Draper 1986; Norman 2002] and is known to be a challenging process in large development environments [Grudin 1993]. To the best of our knowledge, we are the first to present a user-centered design approach for a physically-based rendering component. However, unlike most work in this area, we focus on reformulating the underlying computations rather than modifying the user interface to enhance the usability of the system.

3 Artist Friendly Control Parameters

Making it easy to control the behavior of shading modules is critical in most creative applications. Art directors usually have specific

comments about the appearance of characters and how they want to modify the appearance, and it is important that artists have tools with the right controls to achieve this creative vision.

Having said that, it is important to note that there is no universal “artist friendly” system. Different artists often have different needs and concerns regarding the final appearance. These concerns vary over time, and they are different between production departments as well as between individuals. In particular, we have found the needs early on in the creative process to be different compared to those late in the process. Early on, when the design space is being explored, a few controls tend to be favored, while a lot of controls are desired late in the process when specific details are being tweaked.

In general, it is therefore impossible to satisfy the needs of all artists in one shading model. However, there are some simple criteria which are common among most users, and usually physically based shaders fail to satisfy these criteria as described below.

3.1 Intuitive behavior

The first requirement is that the control parameters should correspond to visually distinct features and behave predictably.

In the physically based world, the appearance of materials is being determined by intrinsic properties (e.g., index of refraction and absorption coefficients). These physically based properties have complex and unintuitive effects on the final hair appearance [Zinke et al. 2009; Bonneel et al. 2009]. This makes it very hard even for trained artists, to guess the shader parameter values in order to get a desirable appearance [Mihashi et al. 2005].




Figure 2: Final hair colors and their corresponding RGB absorption coefficients. Coming up with appropriate values for the absorption coefficients to get a desired hair color is not intuitive. Image is reproduced from [Zinke et al. 2009] © 2009 Association for Computing Machinery, Inc. Reprinted by permission.

One such example is hair color which is determined by absorption coefficients. These measure the wavelength based attenuation of light as it passes through the hair medium, but there is no obvious relation between the absorption coefficients and the hair color. Adding to the complexity is the fact that the final hair color (especially for light-colored hair) also depends on the number of light scattering events inside the hair volume. Figure 2 shows some hair rendering results and their corresponding absorption coefficient values.

A more intuitive control parameter for changing the color of hair would be a simple color variable that has direct impact on the color of hair. However, for a physically based shader this may not be easy to implement as the relationship is often highly non-linear especially when multiple scattering is involved.

3.2 Decoupling

From an artist’s point of view, changes to one visually distinct feature (e.g. brightness of primary highlight, color of the secondary highlight) should not affect other visually distinct features. This is critical because in many cases art directors ask for a change in one feature and expect the rest of the appearance to remain unchanged.

Physical material properties are unfortunately inherently interconnected: Changing one physical property will in most cases affect all the visually distinct features of the final appearance. As an example, changing the index of refraction of hair will affect the color, intensity, and position of different highlights. Even if only one of these needed to change. See Figure 3 (top row) for a visualization of these effects.

Also, energy conservation forces any physically based scattering function to integrate to a value less than (or equal to) one. Therefore, if an artist makes one of the subcomponents of the scattering function very large, other subcomponents have to become smaller. As an example, if we increase the intensity of one of the highlights, the intensity of the other two highlights must be reduced to conserve energy. Similarly, if we consider a single highlight, increasing its width will reduce the observed intensity since the energy gets distributed over a larger area. See Figure 3 (middle row) for a visualization of this effect on the primary highlight. These coupled behaviors reduce art-directability and are undesirable from a usability point of view. Instead, in this example, we would like to be able to change the width of the highlight while keeping the color, intensity, and position of the highlight intact. See Figure 3 (bottom row).




Figure 3: Comparison between coupled and decoupled control parameters. Top row : Changing the index of refraction in a physically based shader affects all visual components of the appearance at the same time and is an example of a coupled (and unintuitive) control parameter. Middle row : Increasing the width of a highlight in a physically based shader must reduce the intensity to preserve energy. This coupled behavior is undesirable from an artist's point of view. Bottom row : An example of decoupled control parameter for changing the width of a highlight which will not affect any other aspects of the highlight.

3.3 Going beyond reality

Physically based models follow the rules of physics and can therefore not produce non-physical results. However, in the creative world the only limitation is the imagination, and art directors are often interested in appearances which are not feasible in the real world. To complicate matters, with a physically based shader if one attempts to modify the hair scattering functions to accommodate such non-physical requests, undesirable side effects may occur. For instance, if the modified scattering function is energy absorbent, the multiple scattering component (which is very important for the overall hair color) might disappear. On the other hand, if the scattering function is energy producing, then the multiple scattering component will blow up (Figure 4). This happens because the overall energy gets increased with every bounce. These results, which are consistent with the rules of physics, are just as undesirable from an artist's point of view as the coupling and unintuitive behavior

mentioned above. Our goal is therefore to create controls which cover the physically correct domain, but which extend seamlessly into the non-physical domain. For discussion on how we address this problem refer to Section 5.2.3.



Figure 4: Unexpected behavior of a physically based shading model after some non-physical modifications to the underlying hair scattering functions. Left : A physically based setting for the shader will result in a reasonable multiple scattering component. Center : When the modified single scattering functions absorb energy the multiple scattering component might disappear. Right : When the single scattering functions produce energy, the multiple scattering component might blow up.

4 Our Approach

As explained above, the reason that physically based shading models have limited artist controls is that physically based scattering functions (f_s) are defined over the domain of material properties. We refer to these material properties as *Physically Based Controls (PBC)* which include parameters like index of refraction η , absorption coefficient σ_a , etc.

$$f_s = f(\omega_i, \omega_r, \sigma_a, \eta, \dots) = f(\omega_i, \omega_r, \{PBC\}) \quad (1)$$

where ω_i and ω_r are lighting and viewing directions.

Our goal is to produce a pseudo scattering function f'_s that approximates f_s but is defined on a different domain of parameters which have intuitive visual meanings to the artists and are separate for all visually meaningful components. We refer to these intuitive, decoupled, and meaningful parameters as *Artist Friendly Controls (AFC)*.

$$f'_s = f(\omega_i, \omega_r, \{AFC\}) \approx f_s \quad (2)$$

The pseudo scattering function f'_s is not limited to the rules of physics and can produce a larger range of appearances than f_s . Therefore, f'_s can produce super-natural appearances as well as physically based results.

To produce f'_s from f_s , we propose the following steps:

1. Examination: Examine the exact behavior of a physically based scattering function f_s over the domain of some material properties $\{PBC\}$.
2. Decomposition: Decompose the behavior of f_s into visually separate and meaningful scattering sub-functions f_{s_i} . Defining meaningful subcomponents is a subjective task and should be done with the help of the artists who will be the end users of the system [Norman and Draper 1986].
3. Defining AFCs: For each sub-component f_{s_i} , define artist friendly control parameters AFC_{ij} which are intuitive, decoupled, and visually meaningful for the artists. These AFCs define qualities like color, intensity, size, shape, and position of the visual features. This step is also subjective and should be done with the help of artists that will be using the system.
4. Reproduction: Create pseudo scattering functions f'_{s_i} that approximate the qualitative behavior of decomposed scattering functions f_{s_i} over the domain of $\{AFC_{ij}\}$.




Figure 5: A schematic visualization of our approach. The first column shows the physically based hair scattering function f_s which is defined over a set of physically based parameters $\{PBC\}$. The second column shows decomposed scattering functions f_{s_i} . In the third column, meaningful artist friendly controls AFC_{ij} are being defined for each subcomponent. In the fourth column, new pseudo scattering functions f'_{s_i} are introduced to approximate each f_{s_i} function. These functions are defined over the domain of AFC_{ij} . In the last column, all of these pseudo scattering functions are combined to give us the final pseudo scattering function f'_s that approximates f_s .

5. Recombination: Combine the approximated pseudo scattering functions f'_{s_i} to get one pseudo scattering function f'_s . The final pseudo scattering function f'_s approximates f_s and is defined over the domain of artist friendly control parameters $\{AFC_{ij}\}$.

See Figure 5 for a schematic visualization of this approach.

5 Applying Our Approach

In this section we explain how we have applied our approach to the single and multiple scattering components.

5.1 Single Scattering

5.1.1 Examination

In computer graphics, the prominent work on single scattering properties of hair fibers is by Marschner et al. [2003]. According to their measurements, single scattering has three main subcomponents: 1) The light that reflects off the surface of hair (aka primary highlight), 2) light that has transmitted through the hair medium (aka transmission highlight), and 3) light that has been internally reflected off the inner surface of the hair (aka secondary highlight). We will refer to these components as R, TT, and TRT, respectively (Figure 6 left).

Due to the presence of tilted cuticles, these three components will be reflected in three different angles around the hair fiber, forming 3 different cones. The R component has the color of the light source and usually appears as a bright highlight. The TT component appears in back lighting situations and is the bright halo around the hair. The TRT component appears above the primary highlight and has the color of the hair. This component contains some randomized looking sharp peaks that are basically caustics formed as the light passes through the hair fibers. Their randomized appearance is due to the fact that hair fibers have elliptical cross sections and are oriented randomly.

Marschner et al. [2003] showed that one can decompose the scattering function of a hair fiber into three longitudinal functions $M(\theta)$ and three azimuthal functions $N(\phi)$. See Figure 6 for a qualitative visualization of these six functions. Marschner et al. [2003] defined the final hair scattering function f_s as:

$$f_s(\theta, \phi) = \sum_X M_X(\theta) N_X(\phi) / \cos^2 \theta \quad (3)$$

where subscript $X \in \{R, TT, TRT\}$ represents one of the three subcomponents.

The longitudinal scattering functions $M_X(\theta)$ have been modeled as unit-integral, zero-mean Gaussian functions. The variance of these Gaussian functions represents the longitudinal width of each highlight:

$$M_X(\theta) = g(\beta_X^2, \theta_h - \alpha_X) \quad (4)$$

Here g is a unit-integral zero-mean Gaussian function, β_X^2 represents the variance of the lobe, α_X represents its longitudinal shift, and θ_h is the longitudinal half angle between incoming and outgoing light directions.

Marschner et al. [2003] proceed to compute the azimuthal scattering functions assuming that the hair fibers have circular cross sections. The important observation in this context is that the final shape of these scattering functions is relatively easy to characterize, and that it is qualitatively similar for different types of hair (Figure 6 right). The exception is the behavior of the glints which is very complex. However, for our purposes it is sufficient to use a simplified model as described below.




Figure 6: Single scattering subcomponents: R, TT, and TRT. (left) Three different paths that light can take after intersecting a hair fiber. (middle) Longitudinal scattering functions $M_x(\theta)$ which are three cones with different apex angles. (right) Azimuthal scattering functions $N_x(\phi)$.

5.1.2 Decomposition

For the decomposition step, we asked a team of artists to identify appearance properties that they want to control. They came up with four components: primary highlight, secondary highlight, glints, and the rim light when the hair is backlit.

These four components align with the underlying physically based calculations very nicely. They are basically the R, TT, and TRT components where the TRT component is being decomposed into two subcomponents: glints and TRT excluding the glints.

5.1.3 Defining AFCs

Defining decoupled and meaningful artist friendly controls for each component means defining qualities like color, intensity, size,

shape, and position of each decomposed component. The artists came up with the following AFCs:

R: Color, intensity, longitudinal position, and longitudinal width.

TT: Color, intensity, longitudinal position, longitudinal width, and azimuthal width.

TRT minus glints: Color, intensity, longitudinal position, and longitudinal width.

Glints: Color, intensity, and frequency of appearance.

Please note that the decomposition and the choice of control parameters are subjective choices and could be done differently depending on the needs of the artists.

5.1.4 Reproduction

Our task is to approximately reproduce all of the decomposed sub-components based on their defined AFCs.

We do this by simulating the longitudinal and azimuthal scattering functions separately. To simulate the longitudinal function, we use Gaussian functions similar to the original paper with the only difference being that the new functions are unit height instead of unit-area. This way, changing the width will not affect the brightness of the highlight. We have thus defined the pseudo longitudinal scattering functions as follows:

$$M'_X(\theta) = g'(\beta_X^2, \theta_h - \alpha_X) \quad (5)$$

where $X \in \{R, TT, TRT\}$, g' is a unit-height zero-mean Gaussian function, β_X^2 represents the longitudinal width of component X , and α_X is its longitudinal shift.

The azimuthal scattering functions are more complex and we have to simulate each one separately. We have to keep in mind that we want to keep the peak of these functions constant so that they will not affect the brightness of each component.

The azimuthal scattering function for the primary highlight $N_R(\phi)$ is like an upside down heart shape (Figure 7). We reproduce a simple approximation of this component according to the work by Kim [2002] as the equation below :

$$N'_R(\phi) = \cos(\phi/2) \quad 0 < \phi < \pi \quad (6)$$

With this approximation we are ignoring the Fresnel term for simplicity but it can be added if more accurate results are desired.




Figure 7: Visualizing the primary highlight's control parameters on its corresponding azimuthal scattering function (left) and a frontlit rendering (right). (a) The intensity I_R , (b) the longitudinal shift α_R , and (c) the longitudinal width β_R^2 of the primary highlight.

The azimuthal scattering function of the transmission component N_{TT} is a sharp forward directed lobe (Figure 8). We simply reproduce it as a Gaussian with unit height and controllable azimuthal width:

$$N'_{TT} = g'(\gamma_{TT}^2, \pi - \phi), \quad (7)$$




Figure 8: Visualizing the transmission component's control parameters on the corresponding azimuthal scattering function (left) and a backlit rendering (right). (a) The intensity I_{TT} , (b) the azimuthal width γ_{TT}^2 , (c) the longitudinal shift α_{TT} , and (d) the longitudinal width β_{TT}^2 of the transmission highlight.

where γ_{TT}^2 is the azimuthal width.

For the secondary highlight we have more control parameters because of glints. Due to the eccentricity of the human hair fibers, the number, intensity, and the azimuthal direction of the glints varies based on the orientation of each hair. However, since we are only concerned with the final visual impact of the glints, we assume that glints are two sharp peaks with the same intensity that are always coming back toward the incoming light direction. We add a random shift to their azimuthal direction to get the randomized appearance. This very simplified model for glints produces visually acceptable results for our purpose. We give the artist controls over the relative brightness of the glints and the frequency of their appearance.

$$N'_{TRT-G} = \cos(\phi/2) \quad (8)$$

$$N'_G = I_g g'(\gamma_g^2, G_{angle} - \phi) \quad (9)$$

$$N'_{TRT} = N'_{TRT-G} + N'_G \quad (10)$$

Here, I_g is the relative intensity of glints over the intensity of secondary highlight, and γ_g^2 is the azimuthal width of the glints. Increasing the azimuthal width of the glints makes them appear more often and decreasing their width will reduce their presence. G_{angle} is the half angle between two glints. To give a randomized appearance to the glints, G_{angle} is different for each hair strand and has a randomized value between 30° and 45° . See Figure 9.




Figure 9: Visualizing the secondary highlight's control parameters on the corresponding azimuthal scattering function (left) and a frontlit rendering (right). (a) The intensity of the secondary highlight I_{TRT} , (b) the relative intensity of glints I_g , (c) the azimuthal width of the glints γ_g^2 , (d) the half angle between the glints G_{angle} , (e) the longitudinal shift α_{TRT} , and (f) the longitudinal width β_{TRT}^2 of the secondary highlight.

To embed the control for color and brightness of each component we simply multiply each one by a scalar variable and a color variable :

$$f'_X = C_X I_X M'_X(\theta) N'_X(\phi) \quad (11)$$

where $X \in \{R, TT, TRT\}$ and C_X and I_X are the color and intensity of component X , respectively. In practice these values can be controlled manually, procedurally or through painted maps.

5.1.5 Recombination

To combine the results we have to add all the components together and divide by \cos^2 to account for the projected solid angle of the specular cone [Marschner et al. 2003].

$$f'_s = \sum_X f'_X / \cos^2(\theta) \quad (12)$$

Note that f'_s is not energy preserving, which we will discuss further in Section 5.2.3.

5.2 Multiple Scattering

Considering multiple scattered light is critical for correct reproduction of hair color, especially for light colored hair. To capture the exact behavior of multiple scattered light one needs elaborate methods like brute force path tracing [Zinke and Weber 2007], photon mapping [Moon and Marschner 2006; Zinke and Weber 2006], or other grid based approaches [Moon et al. 2008]. Path tracing approaches are computationally expensive and their results converge very slowly. Photon mapping and grid based approaches are faster than path tracing methods but are still relatively expensive. All of these methods require ray tracing capabilities and are very costly in a production environment.

Another class of methods try to approximate the multiple scattering component by considering the physical properties of human hair fibers. The most prominent work in this category is the Dual Scattering model [Zinke et al. 2008]. This method is fast and relatively accurate, and with some considerations it can be used efficiently in production without the use of any extra data structures or any ray tracing steps [Sadeghi and Tamstorf 2010].

We chose the Dual Scattering model as our physically based scattering function. Here we explain how we have applied our approach to this model to produce the pseudo scattering function for the multiple scattering component.

5.2.1 Examination

The Dual Scattering method approximates the multiple scattering function as a combination of two components: global multiple scattering and local multiple scattering.

Global multiple scattering accounts for the light that reaches the neighborhood of the shading point. It is dependent on the orientations of all the hairs between the light source and that point. It requires calculating the forward scattering transmittance and the spread of the light that reaches the shading point from all light sources. Global multiple scattering will be computed for different points separately.

Local multiple scattering approximates the scattering events within the local neighborhood of the shading point. It is only dependent on the longitudinal inclination of the hair strand at the shading point, and assumes that all surrounding hairs in the shading region have the same orientation and that there is an infinite number of them. For more details about the Dual Scattering method refer to the original paper [Zinke et al. 2008]. Decomposing the dual scattering component into meaningful components is not as straight-forward as it is for the single scattering component. To find meaningful components, we visualized all the terms involved in the computation of the final result of the model (Figure 10) and asked our artists to choose the ones that have intuitive meanings for them. The artists came up with two groups of components:

Forward Scattering component (F.S.) This includes the $f_s^{scatter}$ term from the dual scattering model, which computes the light that scatters forward and maintains its forward directionality inside the hair volume. This component is very important in backlit situations.




Figure 10: Different quantities involved in the computation of the (a) final results of the Dual Scattering method: (b) Single scattering components f_s , (c) average backscattering attenuation A_b , (d) multiple backscattering distribution function for direct lighting f_{back}^{direct} , (e) multiple backscattering distribution function for indirect lighting $f_{back}^{scatter}$, (f) average backscattering spread S_b , (g) single scattering for indirect lighting $f_s^{scatter}$, (h) $F^{scatter}$ term, and (i) F^{direct} term.

Backscattering component (B.S.) This includes the terms f_{back}^{direct} and $f_{back}^{scatter}$ from the dual scattering model. These are multiple backscattering distribution functions that represent the light that goes into the hair volume and comes back to the surface. f_{back}^{direct} computes this quantity for direct illumination and $f_{back}^{scatter}$ computes this for indirect illumination. Both of these components are smooth Gaussian functions in the Dual Scattering model.

Note that $f_{back}^{scatter}$ and f_{back}^{direct} are very similar quantities. f_{back}^{direct} is being used in the computation of F^{direct} while $f_{back}^{scatter}$ is being used in the computation of $F^{scatter}$ and accounts for the variance of forward scattering in the longitudinal directions σ_f^2 (Figure 11 d and e). In the original paper [Zinke et al. 2008], the term f_{back} is being used for both of these quantities. For further discussion of this refer to [Sadeghi and Tamstorf 2010].




Figure 11: Visualizing the decomposed components of our hair shading model. (a) Final rendering result. (b) Primary highlight component R' , (c) secondary highlight component TRT' , (d) transmission component TT' , (e) backscattering component $B.S.'$, and (f) forward scattering component $F.S.'$.

5.2.2 Defining AFCs

At this point we need to define artist friendly controls like the color, intensity, size, shape, and position for all of the decomposed components. However, all of these components are already indirectly defined by the values chosen for the single-scattering components. By overriding these values with artist defined values, we will lose the richness and details in the appearance of those components. Therefore, instead of overriding the color variable we provide adjustment control parameters for modifying these components. The artists came up with the following AFCs:

F.S. Color Adjust and Intensity Adjust for modifying the computed color and intensity values.

B.S. Color Adjust, Intensity Adjust, Longitudinal Shift Adjust, and Longitudinal Width Adjust.

Setting all these control parameters to their default values gives the original results of the dual scattering model.

5.2.3 Reproduction and Recombination

To reproduce the dual scattering results we use the original algorithm and replace the single scattering component f_s^{direct} with our pseudo scattering function f'_s and embed the defined artist controls into the f_{back}^{direct} , $f_{back}^{scatter}$ and $f_s^{scatter}$ components.

However, as mentioned in Section 3.3, replacing the physically based scattering function of f_s^{direct} with the non-physically based model f'_s can cause problems since f'_s is not necessarily energy conserving. For the purpose of computing the multiple scattering component we solve this problem by normalizing the pseudo scattering function f'_s to make it energy conserving :

$$f_s^{norm}(\theta, \phi) = \frac{f'_s(\theta, \phi)}{\int_{\Omega} f'_s(\theta', \phi') d\theta' d\phi'} \quad (13)$$

Here Ω is the full sphere around the shading point. We use this normalized version of the single scattering component for all the computations in the dual scattering model. This will prevent the multiple scattering component from disappearing or blowing up (Figure 4). For single scattering we still use the un-normalized function f'_s which means that the final result of the hair shader is not guaranteed to be energy conserving. If this is a problem, another normalization step similar to the one described above can be applied to the final result of the shader.

Combining all of the above, we get the pseudo code shown below for reproducing the results of dual scattering with embedded artist controls. The modifications to the physically based version given in [Zinke et al. 2008] are highlighted in blue.

```
// Precompute  $\bar{A}_b(\theta)$ ,  $\bar{\Delta}_b(\theta)$ , and  $\bar{\sigma}_b^2(\theta)$  from  $f_s^{norm}$  for  $0 < \theta < \pi$ 
F( $T_f, \bar{\sigma}_f^2$ , directFraction)
// Backscattering for direct and indirect lighting
 $f_{back}^{direct} \leftarrow 2\bar{A}_b(\theta)g(\theta_h - \bar{\Delta}_b(\theta) + \alpha_{Back}, \bar{\sigma}_b^2(\theta) + \beta_{Back}) / (\pi \cos^2 \theta)$ 
 $f_{back}^{scatter} \leftarrow 2\bar{A}_b(\theta)g(\theta_h - \bar{\Delta}_b(\theta) + \alpha_{Back}, \bar{\sigma}_b^2(\theta) + \bar{\sigma}_f^2(\theta) + \beta_{Back}) / (\pi \cos^2 \theta)$ 
 $f_{back}^{direct} \leftarrow C_{Back} I_{Back} f_{back}^{direct}$ 
 $f_{back}^{scatter} \leftarrow C_{Back} I_{Back} f_{back}^{scatter}$ 
// Longitudinal functions for direct and indirect lighting
 $M'_X \leftarrow g'(\theta_h - \alpha_X, \beta_X^2)$ 
 $M_X^G \leftarrow g'(\theta_h - \alpha_X, \beta_X^2 + \bar{\sigma}_f^2)$ 
// Azimuthal functions for indirect lighting
 $N_X^G(\theta, \phi) \leftarrow \frac{2}{\pi} \int_{\pi/2}^{\pi} N'_X(\theta, \phi') d\phi'$ 
// Single scattering for direct and indirect lighting
 $f_s^{direct} \leftarrow \sum M'_X N'_X(\theta, \phi)$ 
 $f_s^{scatter} \leftarrow \sum M_X^G N_X^G(\theta, \phi)$ 
 $f_s^{scatter} \leftarrow C_{Forward} I_{Forward} f_s^{scatter}$ 
 $F^{direct} \leftarrow \text{directFraction}(f_s^{direct} + d_b f_{back}^{direct})$ 
 $F^{scatter} \leftarrow (T_f - \text{directFraction}) d_f (f_s^{scatter} + \pi d_b f_{back}^{scatter})$ 
// Combine the direct and indirect scattering components
return ( $F^{direct} + F^{scatter}$ ) cos  $\theta_i$ 
```

The symbols used here are as follows : \bar{A}_b for the average backscattering attenuation, $\bar{\Delta}_b$ for the average longitudinal shift, $\bar{\sigma}_b^2$ for the average backscattering variance, T_f for the front scattering transmittance, and $\bar{\sigma}_f^2$ for the front scattering variance. M_X^G and N_X^G are the averaged forward scattered longitudinal and azimuthal scattering functions, respectively. All of these terms are taken directly from the Dual Scattering paper [Zinke et al. 2008].

The terms M'_X and N'_X are the pseudo scattering functions defined in Section 5.1.4. $I_{Forward}$ and I_{Back} are control parameters for adjusting the intensity values, while $C_{Forward}$ and C_{Back} are control parameters for adjusting the color values of the forward scattering and backscattering components respectively. Finally, α_{Back} and β_{Back} are control parameters for adjusting the longitudinal shift and the longitudinal width of the back scattering components.

5.3 Decoupling Single and Multiple Scattering

Multiple scattering computations are based on the single scattering functions, and there is an inherent relationship between these two components since multiple scattering is basically the effect of many single scattering events.

However, this relationship can be problematic for art direction. An art director might request a change of the appearance of the single scattering (e.g. color of the primary highlight) and yet want to keep everything else untouched. Unfortunately, if the artist changes the single scattering components, it will also affect the multiple scattering component.

To address this problem, we have provided the ability to break the link between single and multiple scattering by having two sets of parameters for the single scattering components. One of these sets will feed into the computations of multiple scattering and one will be used as the parameters of the single scattering. These two sets are linked together by default, but the artist has the ability to break this link at any point through the shader GUI.

6 Rendering Results

We have implemented our hair shading system in RenderMan and it has been fully integrated into the production pipeline at the Walt Disney Animation Studios. It is currently being used in the production of the upcoming animated feature film Tangled and has proven to be very versatile in handling different types of hair and fur successfully. Figure 12 shows some frames from the film.

Figure 1 shows frames from a short animation which is also using our hair shader, and which is included in the accompanying video. This shows that the overall appearance holds up nicely when animated. This is important in film production since many shading models can produce convincing still images but perform poorly with animation.

A common concern when introducing a physically inspired shading system in a production environment is that it might not have the same amount of flexibility as an existing ad hoc system. To address this issue, we have recreated the look for a character which was originally created using an ad hoc shader. Figure 13 shows the look of this character using both the original shader and our new shader. Using our new shader, we were able to replicate the art direction and the overall appearance of the character. Note that the groom for this character was created specifically for the ad hoc shader. Since the new shader tends to reveal many more details in the groom than the ad hoc shader, a more exact match between the two results would require changes to the underlying groom. We did not do that for this test as we wanted to focus on the effect of the shader itself. However, the underlying grooms were adapted to the new shader in order to produce the results in Figure 1.



Figure 12: The hair shading system has been used successfully for rendering the long blond hair of the character Rapunzel, the short brown hair of the character Flynn, and the white hair/fur on the horse. The three frames are from the upcoming animated feature film Tangled.




Figure 13: Reproducing the look for the character Penny using the original hair shader (top row), and our new hair shader (bottom row).

Figure 14 shows an art reference and a matching rendering using our new shader. It is important to note that in this case, the art directors were not looking for an exact match between the rendered result and the painted reference image. The painted reference only served as the initial guide for the overall look and feel of the rendered image.




Figure 14: Painted art reference (left) and a corresponding rendering (right) using our new hair shader. Faces have been taken out intentionally.

7 Evaluation

To evaluate the usability of our shader we organized an informal user study at our studio. Our goal was to compare our shader to the current state of the art, by comparing to the existing ad hoc production shader used in the feature film BOLT and a physically based shader that has been used in production. The former is based on Kajiya-Kay’s model, but has many layers of tweaks and controls built on top of it, and is based on years of accumulated experience from numerous feature films. The latter is an implementation of Marschner’s model and Zinke’s BCSDf model for single scattering.

It uses a version of what later became the dual-scattering model for the multiple scattering. We will refer to the three shaders as the New shader, the Production shader and the Research shader.

7.1 Setup

We gathered three groups of artists from our look development and lighting departments with varying levels of experience. We assigned one of the shaders to each group of artists and trained each group on using their assigned shader. Among those who finished the user study, 6 artists were assigned to the New shader, 4 artists were assigned to the Research shader, and 3 artists were assigned to the Production shader.

As reference material for the test we captured photos of a natural blond hair wig illuminated by a single directional light. Blond hair provides the most challenging test for most hair shaders, since its appearance depends on both single and multiple scattering. We provided 4 photos to the artists corresponding to different lighting directions, but captured 8 photos to be used for the evaluations. The goal was to evaluate the behavior of the shaders under both known and unknown lighting directions. In particular we wanted to be able to identify situations in which a shader could be fine tuned to match any given reference, but would require additional tweaks for every new lighting situation.

See Figure 15. The reason we chose photographic reference was to ensure a fair comparison for the physically based shader and to give us an unbiased ground truth.

We groomed a hair model similar to the wig in our photo shoot setup, and placed it in a scene which contained 8 directional lights with the same position, orientation and intensity as the ones used in our photo shoot. Given this we asked the artists to tweak the parameters of their assigned shader in order to come up with a fixed set of parameter values to match the appearance of all 4 given photo references. We limited the amount of time that each artist could spend on this task to 4 hours which is comparable or a little less than what they would spend on a similar task in production.

7.2 Ranking

Upon completion of the assignment we used the submitted shader parameters of each artist to render all 8 lighting directions. We then anonymized and randomized the order of the rendering results and encouraged everyone at our studio to rank the results from best to worst based on the photographic references. We got 36 responses from a mix of expert and non-expert volunteers. Figure 16 visualizes the distribution of rankings for each shader. All rendering results and the details of their evaluation ranks are available in the supplemental materials.




Figure 15: Reference photographs used for the user study. The labels indicate the location of the light. Only images with black labels were given to the artists at the time of the user study. All images, however, were used to evaluate the results.




Figure 16: Normalized distribution of the rankings for all the shading models (1 is best, 13 is worst). The New shader (red) is generally ranked better than the Research shader (blue) while the Production shader (green) is somewhat inconsistent with really good results and some not so good. The numbers have been normalized based on the number of samples in each group.

7.3 Discussion

There are many confounding factors in the evaluation described above which prevent us from drawing any statistically significant conclusions. One is the small sample size, which is difficult to address since the population of people with the right skillset is fairly small. Another is the fact that every artist did not repeat the experiment with all three shaders. This means that an exceptionally good artist can skew the results toward one shader.

With these caveats in mind, there are still a number of trends apparent in the results which are consistent with the experiences we have had with the use of these shaders in production. The Research shader generally performs poorly. While the artists are able to get the desired appearance for some lighting directions, they fail to get a consistent appearance under all lighting directions. See Figure 17. In practice they were also a lot more frustrated with the overall experience using this shader.




Figure 17: Comparing one of the results of the New shader (middle row) with one of the results of the Research shader (bottom row). The artists have successfully matched the frontlit image (middle column) with the photo reference. However, due to lack of control when using the Research shader, the artist was not able to match the hair appearance in the other two lighting setups (left and right columns).

The Production shader can do surprisingly well with an experienced user, and the best results are clearly better than the best result produced by the Research shader. However, it often fails to recreate the appearance of real hair. In particular, it misses the secondary highlight, and has problems with the bright transmission component. It also tends to miss the variation in lighting inside the hair volume and it often gives a flat appearance to the hair. This is shown in Figure 18, but has also been observed during production use.




Figure 18: Comparing one of the results of the New shader (middle row) with one of the results of the Production shader (bottom row). Unlike the New shader, the Production shader fails to correctly produce the secondary highlight (left column), to capture the details of the light scattering inside the hair volume (middle column), and to produce the bright halo around the hair in the backlit situation (right column).

In contrast to the above, the new shader generally performs very well and is able to produce physically plausible results for all the lighting directions. While it is not flawless, its intuitive controls let most artists produce good results with a minimal amount of training. In our video we show that this is also true for different lighting conditions (indoor / outdoor, direct / indirect, etc.), and in practice we have received almost exclusively positive feedback from its use in production.

7.4 Performance

Performance was not the main focus of our work but our shading model has turned out to be more efficient than the other two shaders. We compared the performance of the different shaders by rendering the frontlit image of the user study in 1024×1024 resolution on a 2.66 GHz Intel Xeon 5150 machine. The hair model consisted of 140 guide hairs which were used to generate more than 100000 hairs procedurally at render time. We used the instances of each shader with the highest rankings for these measurements. The results indicate that our new shader is around 3.3 times faster than the Production shader and 1.5 times faster than the Research shader. This is a significant improvement, especially in movie production where even small improvements add up quickly. The Research shader consumes 1.6 times more memory than our shader and the Production shader uses 1.3 times more memory.

8 Summary and Future Work

We have addressed the problem of art-directability of physically-based hair shading models. In particular we have defined the basic requirements for artist friendly systems and shown that physically based models fail to satisfy those requirements. Based on this we have introduced a new approach for creating a physically inspired but art-directable hair shading model. In our experience the new shading model is more intuitive to use and produces better results than the previous shading models used in production.

One interesting area that needs more investigation is the interplay between the underlying groom and the hair shader. Another avenue for future work is to investigate the applicability of our approach to materials other than hair. We speculate that our approach is applicable to a much broader range of materials in appearance modeling.

9 Acknowledgments

We would like to thank everyone at the Walt Disney Animation Studios who helped us with this project as well as Maryann Simmons, Will Chang and Krystle de Mesa for proof reading a draft of this paper. We would also like to thank the anonymous reviewers for their helpful comments. All images © Disney Enterprises, Inc. unless otherwise noted.

References

- APODACA, T., GRITZ, L., LANCASTER, T., PRATER, M., AND BREDOW, R. 2000. Advanced RenderMan 2: To RLINFINITY and Beyond. SIGGRAPH 2000 Course 40.
- APODACA, T., GRITZ, L., PHARR, M., HERY, C., BJORKE, K., AND TREWEEK, L. 2001. Advanced RenderMan 3: Render Harder. SIGGRAPH 2001 Course 48.
- APODACA, T., GRITZ, L., PHARR, M., GOLDMAN, D., LANDIS, H., QUARONI, G., AND BREDOW, R. 2002. RenderMan in Production. SIGGRAPH 2002 Course 16.
- BONNEEL, N., PARIS, S., VAN DE PANNE, M., DURAND, F., AND DRETTAKIS, G. 2009. Single Photo Estimation of Hair Appearance. *Computer Graphics Forum (Proceedings of the Eurographics Symposium on Rendering)* 28, 4, 1171–1180.
- DUCHENEAUT, N., WEN, M.-H., YEE, N., AND WADLEY, G. 2009. Body and Mind: A Study of Avatar Personalization in Three Virtual Worlds. In *CHI '09: Proceedings of the 27th international conference on Human factors in computing systems*, ACM, New York, NY, USA, 1151–1160.
- GOLDMAN, D. B. 1997. Fake Fur Rendering. In *Proceedings of SIGGRAPH 97*, ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 127–134.
- GRUDIN, J. 1993. Obstacles to Participatory Design in Large Product Development Organizations. In *Participatory design: principles and practices*. L. Erlbaum Associates, Hillsdale, NJ, USA, 99–119.
- KAJIYA, J. T., AND KAY, T. L. 1989. Rendering Fur With Three Dimensional Textures. In *Computer Graphics (Proceedings of SIGGRAPH 89)*, ACM, New York, NY, USA, 271–280.
- KIM, T.-Y. 2002. *Modeling, Rendering and Animating Human Hair*. PhD thesis, Los Angeles, CA, USA.
- MARSCHNER, S. R., JENSEN, H. W., CAMMARANO, M., WORLEY, S., AND HANRAHAN, P. 2003. Light Scattering from Human Hair Fibers. *ACM Transactions on Graphics* 22, 3, 780–791.
- MIHASHI, T., TEMPELAAR-LIETZ, C., AND BORSHUKOV, G. 2005. Generating Realistic Human Hair for “The Matrix Reloaded”. In *SIGGRAPH '05: ACM SIGGRAPH 2005 Courses*, ACM, New York, NY, USA, 17.
- MOON, J. T., AND MARSCHNER, S. R. 2006. Simulating Multiple Scattering in Hair Using a Photon Mapping Approach. *ACM Transactions on Graphics* 25, 3, 1067–1074.
- MOON, J. T., WALTER, B., AND MARSCHNER, S. 2008. Efficient Multiple Scattering in Hair Using Spherical Harmonics. *ACM Transactions on Graphics* 27, 3, 31:1–31:7.
- NEULANDER, I. 2004. Quick Image-Based Lighting of Hair. In *SIGGRAPH '04: ACM SIGGRAPH 2004 Sketches*, ACM, New York, NY, USA, 43.
- NORMAN, D. A., AND DRAPER, S. W., Eds. 1986. *User Centered System Design; New Perspectives on Human-Computer Interaction*. L. Erlbaum Associates, Hillsdale, NJ, USA.
- NORMAN, D. A. 2002. *The Design of Everyday Things*. Basic Books, New York, NY, USA.
- PARIS, S., CHANG, W., KOZHUSHNYAN, O. I., JAROSZ, W., MATUSIK, W., ZWICKER, M., AND DURAND, F. 2008. Hair Photobooth: Geometric and Photometric Acquisition of Real Hairstyles. *ACM Transactions on Graphics* 27, 3, 30:1–30:9.
- PETROVIC, L., HENNE, M., AND ANDERSON, J. 2005. Volumetric Methods for Simulation and Rendering of Hair. Tech. Rep. 06-08, Pixar Animation Studios.
- SADEGHI, I., AND TAMSTORF, R. 2010. Efficient Implementation of the Dual Scattering Model in RenderMan. Tech. rep., Walt Disney Animation Studios.
- SCHULER, D., AND NAMIOKA, A., Eds. 1993. *Participatory Design: Principles and Practices*. L. Erlbaum Associates, Hillsdale, NJ, USA.
- WARD, K., BERTAILS, F., KIM, T.-Y., MARSCHNER, S. R., AND CANI, M.-P. 2007. A Survey on Hair Modeling: Styling, Simulation, and Rendering. *IEEE Transactions on Visualization and Computer Graphics* 13, 2, 213–234.
- ZINKE, A., AND WEBER, A. 2006. Global Illumination for Fiber Based Geometries. In *Electronic proceedings of the Ibero American Symposium on Computer Graphics (SIACG)*.
- ZINKE, A., AND WEBER, A. 2007. Light Scattering from Filaments. *IEEE Transactions on Visualization and Computer Graphics* 13, 2, 342–356.
- ZINKE, A., YUKSEL, C., WEBER, A., AND KEYSER, J. 2008. Dual Scattering Approximation for Fast Multiple Scattering in Hair. *ACM Transactions on Graphics* 27, 3, 32:1–32:10.
- ZINKE, A., RUMP, M., LAY, T., WEBER, A., ANDRIYENKO, A., AND KLEIN, R. 2009. A Practical Approach for Photometric Acquisition of Hair Color. *ACM Transactions on Graphics* 28, 5, 165:1–165:9.