

# Fake Fur Rendering

Dan B Goldman\*

Industrial Light and Magic



## Abstract

A probabilistic lighting model is presented for thin coats of fur over skin. Previous methods for rendering furry objects and creatures have addressed the case where individual strands or tufts of hair may be resolvable at the pixel level. These methods are often computationally intensive. However, a large class of real-world cases where individual hairs are much smaller than the size of a pixel can be addressed using a probabilistic model for the expected value of reflected light within a small surface area. Under the assumption that hair parameters are slowly varying across the skin, lighting calculations are performed on a reference hair with prefiltered parameters. The reflected light from individual hairs and from the skin below is blended using the expectation of a ray striking a hair in that area as the opacity of the fur coating. Approximations for hair-to-hair shadowing and hair-to-skin shadowing can be made using the same hit-expectation model. Our system can be implemented in existing commercial surface-rendering software at a much lower computational cost than typical resolvable-hair methods.

**CR Categories and Subject Descriptors:** I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism - Color, shading, shadowing and texture.

**Additional Keywords:** natural phenomena, animals, fur, anisotropic shading

## 1. INTRODUCTION

At first glance, fur appears to be one natural phenomenon which doesn't "cheat" easily. The appearance of a furry object is so distinctive that many approaches have utilized hair-by-hair methods to achieve an acceptable level of realism. Indeed, some of the most striking images of creatures with fur have been rendered using techniques in which individual strands of hair are visible [11][19]. But since many real-world creatures have millions of

hairs, object-based rendering techniques typically have running times which do not decrease as the image size decreases. Volumetric rendering algorithms render smaller images faster, but tend to be memory and computation-intensive.

Our solution does not attempt to address the 'closeup' situations which are well-handled by the existing models. Instead we address the common case where hair geometry is not visible at the final image resolution, but the visual characteristics of fur, such as glossy sheen and soft illumination, are still observed. Our model falls into a class of secondary approximations, in which hairs are not rendered directly, but are used as the underlying model for the furry surface's lighting properties.

## 2. RELATED WORK

Most attempts to render fur have used brute force methods, representing hairs with large numbers of polygons or particles [7],[15],[20],[13],[2]. The primary drawbacks of these types of methods are severe aliasing and/or computational costs which, in some algorithms, actually increase as the subject decreases in screen size.

Kajiya [12] has addressed the illumination and rendering of hairs using a volume technique, by precomputing a volume 'texel' which is tiled across a furry surface. Hair geometry is rasterized into this texel, and final rendering is accomplished using volume rendering. This technique lends itself well to uniformly furry surfaces which can be tiled using a small number of such texels. Others eschew texels in favor of procedural hair generation [18]. Unfortunately, these are some of the more computationally intensive methods available.

Many lighting models for complex surfaces take a probabilistic approach to microstructure [4],[21]. Recent work [22] has extended this paradigm to more complex surfaces by describing a general method for estimating the bidirectional reflectance distribution function (BRDF) via Monte Carlo sampling and parametrizing it using spherical harmonics. This method is well suited to complex but uniformly patterned surfaces.

A number of proprietary fur renderers have been developed, but published details are rare [8][17]. The true ancestors of this work are the proprietary renderers used at our facility for high-detail fur rendering [19] ([11]). These renderers were useful not only for creating reference images for quality comparisons, but also as working models forming a basis for comparisons of tradeoffs and limitations of alternative rendering methods.

## 3. FAKE FUR RENDERING

We call our probabilistic fur rendering algorithm method the 'fakefur' algorithm, to distinguish it from our proprietary high-detail method not covered here, which by comparison became known as the 'realfur' algorithm. Despite the nomenclature, both methods

---

\*P.O. Box 2459, San Rafael, CA 94912.  
email: dgoldman@cs.stanford.edu

are approximations of varying degrees to the appearance of real mammalian fur.

It should be noted that the goal of this method is to render creatures to be composited into live-action feature films. Although the model can be used for other purposes, this goal has motivated its parametrization and allowed us to ignore certain lighting behaviours uncommon in such a setting:

The fakefur rendering method hinges on a probabilistic method for computing fur visibility, which we call the fakefur opacity function.

### 3.1. Outline

A sketch of the fakefur illumination process for a given area is as follows:

- I. Compute the mean hair geometry within the sample region. This is the 'reference hair'.
- II. For each light:
  1. Using the fakefur opacity function, compute the hair-over-hair shadow attenuation.
  2. Compute the reflected luminance of the average hair in the sample region.
  3. Using the fakefur opacity function, compute the hair-over-skin shadow factor.
  4. Compute the reflected luminance of the underlying skin.
  5. Using the fakefur opacity function, compute the hair/skin visibility ratio.
  6. Blend the reflected luminance of the skin and hair using the visibility ratio to obtain the final reflected luminance of the sample region.
- III. Sum the reflected luminances for each light to obtain the total reflected luminance for the sample region.

### 3.2. Parametrization

We parameterize hair geometry on a surface by hair length, hair radius, density of hairs, and hair tangents at the root and tip. The reflectivity of individual hairs is parameterized by diffuse reflectivity, specular reflectivity, specular exponent, and several directionality factors for reflectivity/transmissivity control and Lambertian macro-behaviour. Of the aforementioned parameters, only the diffuse reflectivity is wavelength dependent.

The parameters for the hairs in a particular region may vary over the surface, either in a procedural manner or defined via texture maps.

### 3.3. The Fakefur Illumination Function

To describe the reflected luminance of a single hair, we use a modified version of the hair reflectance model described in [12]. The equations from [12] (with some notation modified for consistency) are:

$$\Psi_{\text{diffuse}} = K_d \sin(\bar{T}, \bar{L}) \quad (1)$$

$$\Psi_{\text{specular}} = K_s [(\bar{T} \cdot \bar{L})(\bar{T} \cdot \bar{E}) + \sin(\bar{T}, \bar{L})\sin(\bar{T}, \bar{E})]^p \quad (2)$$

$$\Psi_{\text{hair}} = \Psi_{\text{diffuse}} + \Psi_{\text{specular}} \quad (3)$$

The vectors  $\bar{T}$ ,  $\bar{L}$ , and  $\bar{E}$  represent the normalized hair tangent vector, the normalized light direction vector, and the normalized eye direction vector, respectively.  $\Psi$  and its subscripts are the reflectivity components.

One limitation of this model is its lack of directionality: hairs are fully lit even if  $\bar{L}$  is opposite  $\bar{V}$ . We are interested in both reflection and transmission. To increase directionality, we utilize two new attenuation factors, introduced by [24], which may be used to tune the relative transmissivity and reflectivity of a hair.

We first characterize the relative directionality of a given incident light ray, eye ray, and hair tangent using the cosine of the dihedral angle between the planes containing each pair.

$$\kappa = \cos(\bar{T} \times \bar{L}, \bar{T} \times \bar{E}) = \frac{(\bar{T} \times \bar{L}) \cdot (\bar{T} \times \bar{E})}{|\bar{T} \times \bar{L}| |\bar{T} \times \bar{E}|} \quad (4)$$

Note that when  $\bar{L}$  and  $\bar{E}$  strike the same side of the hair (frontlighting),  $\kappa > 0$ , and when  $\bar{L}$  and  $\bar{E}$  lie on opposite sides of the hair (backlighting),  $\kappa < 0$ .

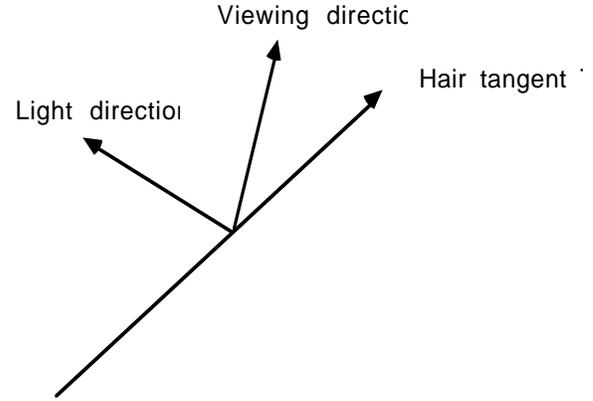


Figure 1. Frontlighting

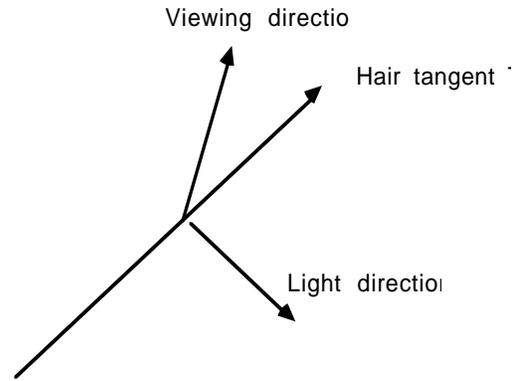


Figure 2. Backlighting

We represent the amounts of forward and backward scattering of the hair by the parameters  $\rho_{\text{transmit}}$  and  $\rho_{\text{reflect}}$ , which vary in the range  $[0,1]$ . Then our directional attenuation factor  $f_{\text{dir}}$  is computed as follows:

$$f_{\text{dir}} = \frac{1+\kappa}{2} \rho_{\text{reflect}} + \frac{1-\kappa}{2} \rho_{\text{transmit}} \quad (5)$$

White or gray hairs are well-represented by nearly equal reflectance and transmission coefficients. Hairs with more pigment will have much higher reflectance than transmission. When both  $\rho_{\text{reflect}}$  and  $\rho_{\text{transmit}}$  are 1, no attenuation occurs.

We also include a surface normal factor as a quick and dirty way to adjust shadowing. Since the layer of fur is approximated geometrically by a flat surface, if we were to use only shadow mapping or shadow tracing to determine the shadowed areas, a hard termination line would result. Instead we specify a smooth gradation from full illumination to full shadow:

$$f_{\text{surface}} = 1 + \rho_{\text{surface}} (\text{smoothstep}(\bar{N} \cdot \bar{L}, \omega_{\min}, \omega_{\max}) - 1) \quad (6)$$

where  $\bar{N}$  is the normalized surface normal, and smoothstep is the smooth Hermite interpolation between  $\omega_{\min}$  and  $\omega_{\max}$ :

$$\text{smoothstep}(x, a, b) = \begin{cases} 0, & \forall x < a \\ 1, & \forall x > b \\ -2\left(\frac{x-a}{b-a}\right)^3 + 3\left(\frac{x-a}{b-a}\right)^2, & \forall a < x < b \end{cases} \quad (7)$$

When  $\rho_{\text{surface}}$  is 0, no attenuation occurs.  $\omega_{\min}$  and  $\omega_{\max}$  are the cosines of the starting and ending shadow termination angles, and can be easily adjusted to match reference images.

In our model, both  $f_{\text{dir}}$  and  $f_{\text{surface}}$  are multiplied into the r.h.s. of equation [III.3] above:

$$\Psi_{\text{hair}} = f_{\text{dir}} f_{\text{surface}} (\Psi_{\text{diffuse}} + \Psi_{\text{specular}}) \quad (8)$$

If hairs are bent along their trajectory, or if the hair parameters are otherwise varying from root to tip, multiple samples along the reference hair can be computed for  $\Psi_{\text{hair}}$  and averaged. In practice, we found using a very small number of samples (3 or 4) is sufficient.

Because the hairs in question are relatively short and slowly varying, we disregard the possibility of hairs with widely disparate parameters reflecting light in the same sample region.

Like [12], this model is a first-order approximation, which is most accurate when the hair albedo is low. No secondary scattering of light off of hairs onto other hairs or onto skin is considered.

### 3.4. The Fakefur Opacity Function

The fur opacity function, denoted  $\alpha_f$ , computes the mean opacity of a patch of fur as viewed from a given angle.  $\alpha_f$  is a function of the hair geometry, the distribution of hairs, and the viewing angle. In general, both the hair geometry and the distribution of hairs can be quite complex, and we must make some simplifying assumptions in order to generate an easily computable form for  $\alpha_f$ .

We make the following assumptions concerning hair geometry:

- Hairs are truncated cones of radius  $r_b$  at their base,  $r_t$  at their tip and length  $l$ .

- $l \gg r_b$  (9)

- $r_b \geq r_t$  (10)

In general, the projection of a truncated cone into a viewing plane is the union of two ellipses (the projection of the base and

top) and a trapezoid (the projection of the sides). The area of the projection of the base and top are proportional to  $r^2$ , while the area of the projection of the sides is proportional to  $l(r_b + r_t)$ . So the constraint  $l \gg r_b, r_t$  implies that  $r^2$  is vanishingly small for most viewing directions. Therefore, we will consider only the projection of the sides onto the viewing plane.

Under these assumptions, we compute the area of the projection of a hair onto the viewing plane as the projection of its trapezoidal profile,

$$A_h = l(r_b + r_t) / 2 \quad (11)$$

$$A'_h = A_h \sin(\bar{E}, \bar{T}) \quad (12)$$

We make the following assumptions concerning distribution of hairs:

- All hairs in the sampled region share identical geometry and orientation.
- The distribution of hairs in a small region has Poisson characteristics: Within a zone of uniform density, a sample of half the size will contain half the hairs and hairs are placed independently of each other.

It may be noted that the distribution of hairs on mammal fur seems to follow a Poisson-disk pattern, not the Poisson pattern described by our model [12].<sup>1</sup> Nevertheless, the assumption of a Poisson pattern vastly simplifies the computation of the fur opacity function, and as we will see in the following section, does not significantly alter the results.

Under the above assumptions, we compute the average area on the skin covered by  $n_i$  hairs:

$$A_s = \frac{n_i}{D} \quad (13)$$

where  $n_i$  is a constant denoting the number of hairs in a sample region and  $D$  is the local density of hairs. The projection of that area is

$$A'_s = A_s (\bar{E} \cdot \bar{N}) \quad (14)$$

Thus, the coverage of a single hair in this area, and the probability of a random ray striking the single hair from direction  $\bar{E}$  is

$$\alpha_h = \frac{A'_h}{A'_s} = \frac{A_h \sin(\bar{E}, \bar{T})}{A_h \bar{E} \cdot \bar{N}} = \frac{A_h}{A_s} g(\bar{E}, \bar{T}, \bar{N}) \quad (15)$$

We isolate the projection-dependent part of  $\alpha_h$  above as the fakefur projection function

$$g(\bar{E}, \bar{T}, \bar{N}) = \frac{\sin(\bar{E}, \bar{T})}{\bar{E} \cdot \bar{N}}. \quad (16)$$

The coverage of the entire distribution of hairs, assuming their independence, is computed as:

$$\alpha_f = 1 - (1 - \alpha_h)^{n_i} = 1 - \left(1 - \frac{DA_h g(\bar{E}, \bar{T}, \bar{N})}{n_i}\right)^{n_i} \quad (17)$$

As the number of hairs in the sample region increases, this

<sup>1</sup>There is some ambiguity in the literature concerning the difference between a Poisson pattern and a Poisson-disk pattern. In this paper, a Poisson pattern refers to independently distributed samples, while a Poisson-disk pattern is defined, as in [9], as one in which "no two samples are closer together than some distance  $r_p$ " defining a non-overlapping radius surrounding each sample. [9] notes that "we also usually want the samples to be as close together as the disks allow."

simplifies to:

$$\lim_{n_i \rightarrow \infty} \alpha_f = 1 - \frac{1}{e^{DA_h g(\bar{E}, \bar{T}, \bar{N})}} \quad (18)$$

This is the fakefur opacity function.

### 3.5. Using the Fakefur Opacity Function

In the illumination process, the fakefur opacity function is used for three separate computations: hair-over-skin shadows, hair-over-hair shadows, and hair-over-skin visibility.

Hair-over-skin shadows are handled by computing the opacity of the fur as seen from the light direction, and attenuating the light intensity by this opacity before illuminating the skin:

$$\lambda_{\text{skin}} = I[1 - \alpha_f(\bar{L})]\Psi_{\text{skin}} \quad (19)$$

where  $I$  is the illuminance, and  $\lambda$  and its subscripts are the reflected luminance and its summed skin and hair components, respectively.

Hair-over-hair shadows are simulated by using some fraction of the hair-over-skin shadows to similarly attenuate the hair illuminance:

$$\lambda_{\text{hair}} = I[1 - s\alpha_f(\bar{L})]\Psi_{\text{hair}} \quad (20)$$

A physical model should include an integral summing the shadowed regions along a hair. The tips of the hairs will be unshadowed by other hairs, while the roots of the hairs will be completely shadowed by other hairs. We approximate this integral with the constant  $s$ . This constant can be adjusted to increase or decrease the density of the hair-over-hair shadows, but a value of  $s = 0.5$  seems to work well for essentially straight cylindrical hairs. This corresponds to a coat of fur in which, on average, half of each hair is in shadow and half is not in shadow.

The hair-over-skin visibility computation is the simplest. The opacity of the fur as seen from the camera viewing direction is computed, and this value used to blend the skin luminance with the hair luminance:

$$\lambda = \alpha_f(\bar{E})\lambda_{\text{hair}} + [1 - \alpha_f(\bar{E})]\lambda_{\text{skin}} \quad (21)$$

### 3.6 Large-scale geometry

Certain other steps involving large-scale geometry are not included in the above outline. The three most notable omissions are shadows cast by skin onto hairs (skin-on-hair shadows), shadows cast by skin onto other skin surfaces (skin-on-skin shadows), and the skin illumination model. These are not central to the algorithm, and are well-handled by existing methods, so they will be covered fairly briefly:

An implementation emphasizing physical accuracy might opt for ray-traced shadows [5] and a skin substructure illumination model such as [10]. However, in keeping with the high priority of efficiency, our implementation applies the most expedient methods available: For the scales at which we wish to render furry things, shadow maps [23] are adequate mechanisms for skin-on-hair and skin-on-skin shadows. And since the underlying skin is visible only in a few areas where fur is thin or sparse, we use a variant of the Torrance-Sparrow illumination model [21] for computing the

reflected luminance of the skin.

## 4. DISCUSSION

### 4.1. Validating the Opacity Approximation

Since the fakefur opacity function is so essential to the illumination equation, it's important to establish that it's a valid approximation.

How well does  $\alpha_f$  approximate the density of real fur? The distribution of hairs on mammal fur has been observed to be distributed in a poisson-disk pattern [12]. Our model, on the other hand, assumes that within a local region hairs are distributed independently. We might expect that our model underestimates opacity, because hairs in a poisson-disk distribution will overlap less frequently than in a poisson distribution.

However, the approximation can be justified as follows: Very short hairs can only overlap other hairs whose roots lie immediately adjacent. In a poisson-disk distribution, there is high correlation between adjacent hairs, while in an independent scatter, there is no correlation. Therefore, for short hairs, the fur opacity function is indeed a poor approximation of a poisson-disk distribution's opacity. However, as hairs grow longer relative to the distance between their roots, they may overlap hairs whose roots lie far from their own. In a poisson-disk distribution such that the disk radius is much larger than the hair radius, there are no direct placement constraints on hairs which lie further away, so the odds of overlapping hairs are well approximated by the poisson distribution. (This is not true if the disk radius is not much larger than the hair radius, but in this case the density is probably very close to 1, by the assumption of equation 9, so the error is very small.)

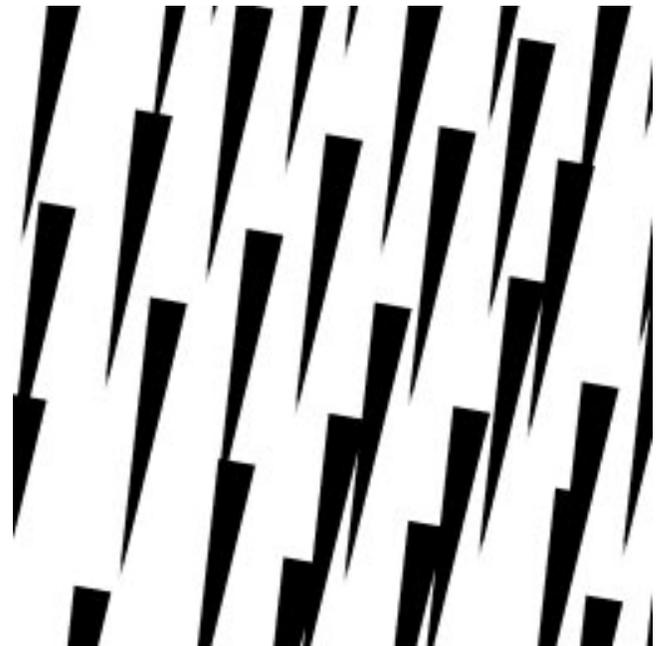


Figure 3. Poisson-disk pattern of triangles scan-converted into a 512x512 grid.  $r = .1, l = 1.5$

We have verified empirically that the poisson distribution

density model approximates a poisson-disk distribution's density quite closely when hairs are long relative to the distance between them. Figure 3 shows a poisson-disk distribution of identically oriented hairs with  $l = 1.5$ ,  $r_b = .2$ , and  $r_t = 0$  scan-converted into a 512x512 buffer of size 3.5. Poisson scatters of varying densities were similarly scan-converted and the resulting coverage  $\alpha_f$  plotted against  $D$ . The results are shown in figure 4. The fur opacity function approaches 1 slightly more slowly than the empirical data, as expected, but the discrepancy is small, and narrows even further for longer hairs.

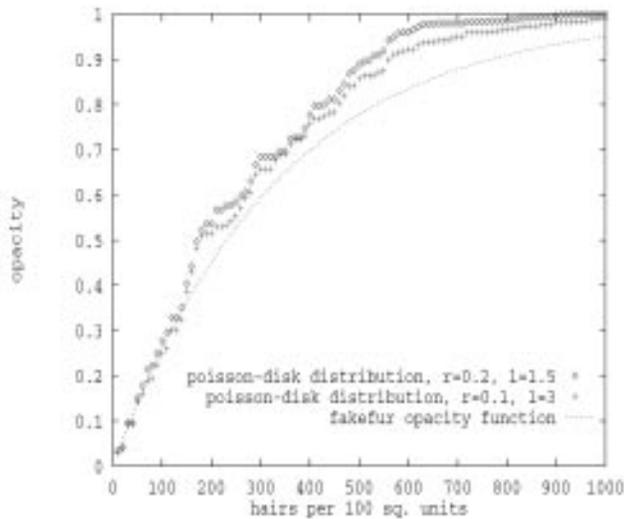


Figure 4. Fakefur opacity function compared to poisson-disk distribution densities.

## 4.2. Special Cases

There are two special cases which the above model does not consider.

The first special case is the 'hot spot': the viewing vector  $\bar{E}$  is close to the illumination vector  $\bar{L}$ . Because the shadows of hairs are almost entirely occluded from view by the hairs themselves, the overall brightness is greater than predicted by this model. This effect can be observed by looking at one's own shadow in a patch of grass on a sunny day. The area immediately surrounding the shadow, where the viewing and illumination vectors are very close, appears brighter than the rest of the grass.

The second special case is the 'halo': the viewing and illumination vectors are nearly opposite. In this case, the approximation of the fur layer as a flat surface and the use of shadow maps underestimates the strong transmission of light through the fur around the silhouette edges of the creature.

Although these are easily observed in the real world, such illumination environments are generally avoided by cinematographers: When frontlighting, it is common to offset the frontmost lights by several degrees from the camera axis, and when backlighting, lights are often placed above and out of frame to separate foreground and background without appearing unnatural. Therefore, we have not found these limitations to be of great practical concern for our usage.

## 5. RESULTS

The fakefur algorithm was used with resounding success to simulate the appearance of dog fur in live-action feature films [1][14]. The algorithm was implemented using the RenderMan shading language and PhotoRealistic RenderMan rendering software.

The color images below illustrates the degree of realism attainable with this technique. In figures 5 and 6., only the two adult dogs were filmed on location. All of the puppies in these two images are computer-generated models illuminated using only the fakefur method.<sup>2</sup>

The fakefur model was eventually used to render most of the computer-generated dalmatians featured in this film. Where additional detail and nonlocal effects were required in closeups, the 'realfur' hair renderer was employed. Figures 7 and 8 illustrate the close match that was achieved between the 'realfur' model and the probabilistic 'fakefur' model. Images were generated at a variety of scales to verify the similarity of appearance. Note in particular the specular sheen on the ears, the effect of high opacity along the silhouette edges, and the pinkness of the underlying skin showing through the thin coat of fur, which is itself colored a neutral off-white.

Although actual rendering times are highly dependent on hardware and software, the 'fakefur' images in figures 7 and 8 rendered about 6 times faster than the 'realfur' images on the same machines.

As shown in this plate, the two rendering techniques generate almost indistinguishable images at a sufficiently small scale. This allowed us to utilize both techniques in the same shots. In some cases, individual dalmatians were rendered using the fakefur method when distant from the camera, and using the realfur method as they approached (or vice versa). The dissolve between the two is essentially invisible.

## 6. FUTURE WORK

Some of the limitations of this work are fundamental, such as the lack of high-frequency detail. However, others are merely simplifications and could be improved with some additional work.

The constraint that hairs must be short is imposed in order to enforce locality of texture influence and to avoid geometry displacement away from the underlying surface. This constraint could be relaxed by convolving the hair parameters with a variable length and direction linear kernel [3] before applying them, and by including a displacement computation for the underlying surface.

The current hair-to-hair shadow formula could be made more accurate by extending the fur opacity function into three-dimensions, taking into account the change of hair radius and the bend of the hairs along their length.

The model has an unwieldy number of parameters. Although some may be determined by direct measurement, many must be assigned by trial and error. This allows considerable freedom for aesthetic considerations, but makes achieving a specific appearance somewhat cumbersome. In the future we hope to find ways of reducing the number of free parameters.

Although this model itself is applied here to a single breed of

<sup>2</sup>Color plates have been color-corrected to match photographic film response.

dog, it can be applied to any animal with reasonably short fur. The concepts underlying the fakefur model also show promise in creating effective lighting models for certain classes of fabrics.

## 7. ACKNOWLEDGEMENTS

The Dalmatian model shown here was created by Geoff Campbell, Kyle Odermatt, and Wayne Kennedy. Textures and hairstyles by Carol Hayden. The shots pictured in Color Plate I were animated by Mike Eames, John Campanaro, and Daniel Jeannette. Color Plate I.a. was lit and composited by Samir Hoon.

Thanks to Barry Armour, Gail Currey, Christian Rouet, Doug Smythe, Chrissie England, Jeff Yost, and Florian Kainz. Thanks also to the SIGGRAPH paper reviewers, whose insightful comments helped me close a number of holes.

Very special thanks to Jocelyn Lamm and Ellen Pasternack at ILM, and Mary Lippold at Disney Rights Administration, and to Disney for allowing some of the above images to appear in this paper.

## 8. REFERENCES

- [1] *101 Dalmatians*, Walt Disney Studios 1996.
- [2] Anjyo, Ken-Ichi, Yoshiaki Usami, and Tsuneya Kurihara, "A Simple Method for Extracting the Natural Beauty of Hair." In Edwin C. Catmull, editor, *Computer Graphics (SIGGRAPH 92 Conference Proceedings)*, volume 26, pages 111-120. Addison Wesley, July 1992. ISBN 0-89791-479-1.
- [3] Cabral, Brian, and Leith Leedom, "Imaging Vector Fields Using Line Integral Convolution." In James T. Kajiya, editor, *SIGGRAPH 93 Conference Proceedings*, Annual Conference Series, pages 263-270. ACM SIGGRAPH, Addison Wesley, 1993. ISBN 0-89791-601-8.
- [4] Cook, Robert L. and Kenneth Torrance, "A Reflectance Model for Computer Graphics," In *Computer Graphics (SIGGRAPH 81 Conference Proceedings)*, volume 15(3), pages 307-316. ACM SIGGRAPH, Addison Wesley, August 1981.
- [5] Cook, Robert L., "Shade Trees." In *Computer Graphics (SIGGRAPH 84 Conference Proceedings)*, volume 18(3), pages 223-231. ACM SIGGRAPH, Addison Wesley, July 1984.
- [6] Cook, Robert L., Tom Porter, and Loren Carpenter, "Distributed Ray Tracing." In *Computer Graphics (SIGGRAPH 84 Conference Proceedings)*, volume 18(3), pages 137-145. ACM SIGGRAPH, Addison Wesley, July 1984.
- [7] Csuri, C., et al., "Towards an interactive high visual complexity animation system." In *Computer Graphics (SIGGRAPH 79 Conference Proceedings)*, volume 13(2), pages 289-299. ACM SIGGRAPH, Addison Wesley, August 1979.
- [8] Duncan, Jodi. "The Island of Dr. Moreau: Moreau's Menagerie," *Cinefex* 68, pages 59-65, 123-124, and 142, December 1996.
- [9] Glassner, Andrew S. *Principles of Digital Image Synthesis*, Morgan Kaufmann Publishers, 1995. ISBN 1-55860-276-3
- [10] Hanrahan, Pat, and Wolfgang Krueger. "Reflection from Layered Surfaces due to Subsurface Scattering." In James T. Kajiya, editor, *SIGGRAPH 93 Conference Proceedings*, Annual Conference Series, pages 165-174. ACM SIGGRAPH, Addison Wesley, August 1993. ISBN 0-89791-601-8.
- [11] *Jumanji*, Tri-Star Pictures, 1996.
- [12] Kajiya, James T., and Timothy L. Kay, "Rendering Fur with Three Dimensional Textures." In *Computer Graphics (SIGGRAPH 89 Conference Proceedings)*, volume 23(3), pages 271-277. ACM SIGGRAPH, Addison Wesley, July 1989. ISBN 0-89791-312-4.
- [13] LeBlanc, André M., Russell Turner, and Daniel Thalmann, "Rendering Hair using Pixel Blending and Shadow Buffers," In *The Journal of Visualization and Computer Animation*, Vol 2, pages 92-96, 1991. ISSN 1049-8907
- [14] *Mars Attacks*, Warner Brothers, 1996.
- [15] Miller, Gavin S.P., "From Wire-Frames to Furry Animals.", In *Graphics Interface '88 Proceedings*, pages 138-145, 1988.
- [16] Neyret, Fabrice. "A General And Multiscale Method For Volumetric Textures." In *Graphics Interface '95 Proceedings*, pages 83-91, May 1995.
- [17] Peishel, Bob, "Feline Fabrication." *Cinefex* 56, November 1993, 17-18.
- [18] Perlin, Ken, and Eric M. Hoffert, "Hypertexture." In *Computer Graphics (SIGGRAPH 89 Conference Proceedings)*, volume 23(3), pages 253-262. Addison Wesley, July 1989.
- [19] Pourroy, Janine, "The Game Board Jungle." *Cinefex* 64, pages 54-71, December 1995.
- [20] Rosenblum, Robert E., Wayne E. Carlson, and Edwin Tripp III, "Simulating the Structure and Dynamics of Human Hair: Modeling, Rendering, and Animation." In *The Journal of Visualization and Computer Animation*, volume 2, pages 141-148, 1991. ISSN 1049-8907
- [21] Torrance, Kenneth, and E.M. Sparrow, "Theory for Off-Specular Reflection from Roughened Surfaces." In *Journal of the Optical Society of America*, volume 57(9), pages 1105-1114, September 1967.
- [22] Westin, Stephen .H., James R. Arvo, and Kenneth E. Torrance, "Predicting Reflectance Functions from Complex Surfaces" In Edwin C. Catmull, editor, *Computer Graphics (SIGGRAPH 92 Conference Proceedings)*, volume 26(2), pages 255-264. Addison Wesley, July 1992.
- [23] Williams, Lance, "Casting Curved Shadows on Curved Surfaces." In *Computer Graphics (SIGGRAPH 78 Conference Proceedings)*, volume 12(3), pages 270-274. Addison Wesley, August 1978.
- [24] Yost, Jeffrey, "Fur Lighting Parameters", ILM internal memo, 1995.



Figure 5. A frame from the film *101 Dalmatians*. © Disney 1996, All Rights Reserved.



Figure 6. A frame from the film *101 Dalmatians*. © Disney 1996, All Rights Reserved.

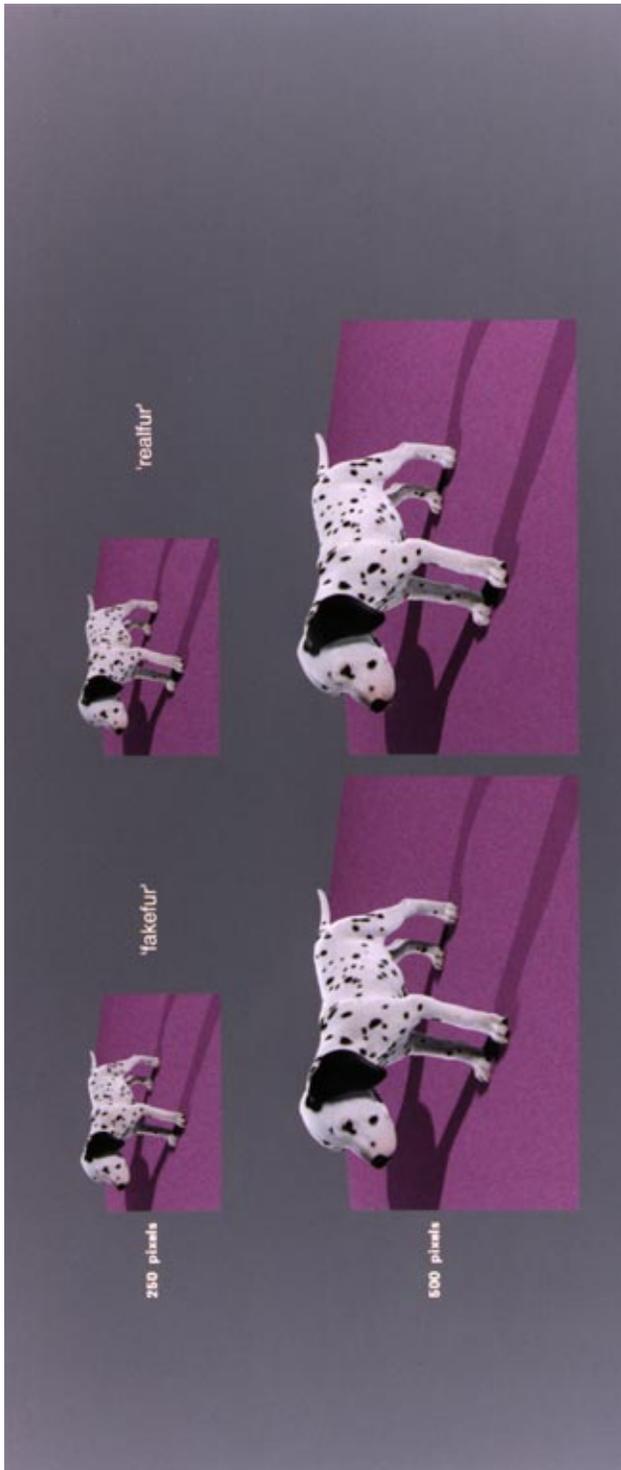


Figure 7. Comparison of 'realfur' and 'fakefur' methods at different scales.

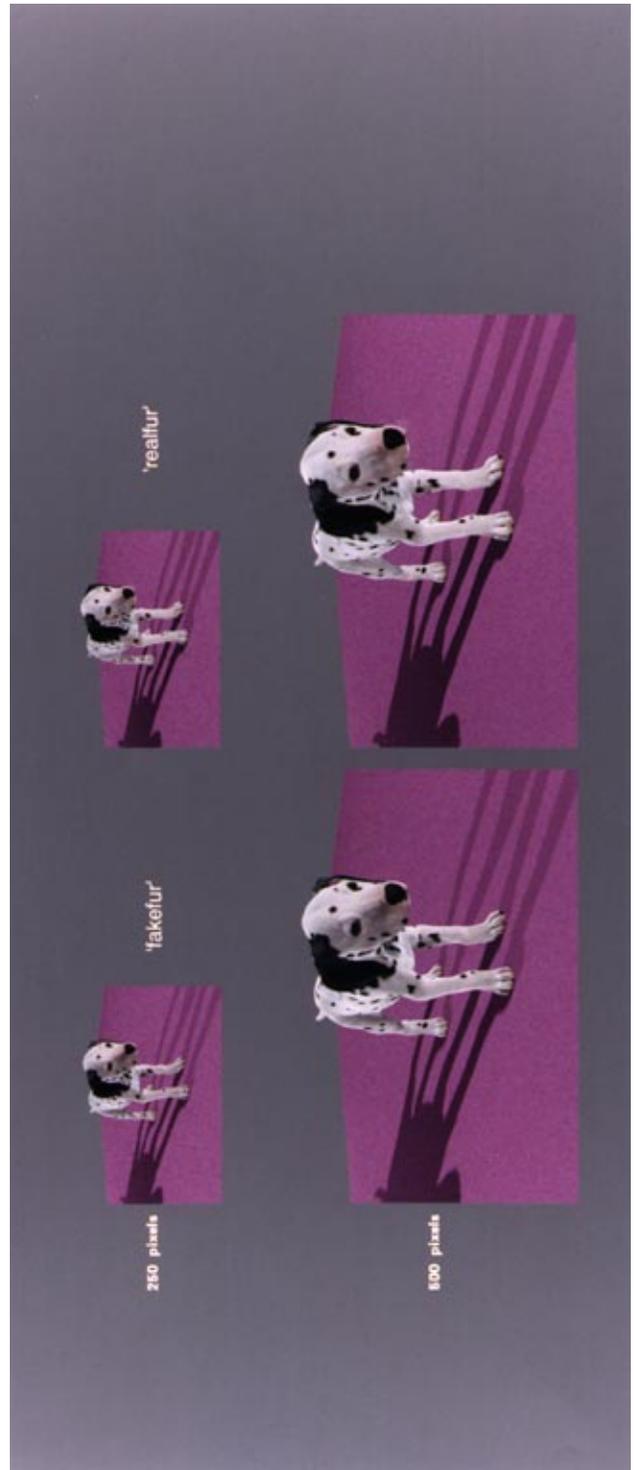


Figure 8. Comparison of 'realfur' and 'fakefur' methods at different scales.