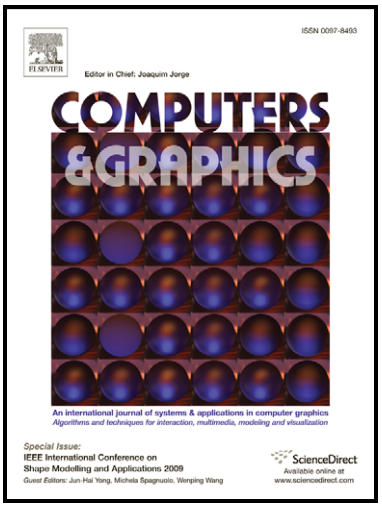

www.elsevier.com/locate/cag

Rendering fur directly into images

Tania Pouli, Martin Pražák, Pavel Zemčík, Diego Gutierrez, Erik Reinhard

PII: S0097-8493(10)00087-7
DOI: doi:10.1016/j.cag.2010.06.004
Reference: CAG 2023

To appear in: *Computers & Graphics*

Received date: 3 June 2009
Revised date: 4 June 2010
Accepted date: 7 June 2010

Cite this article as: Tania Pouli, Martin Pražák, Pavel Zemčík, Diego Gutierrez and Erik Reinhard, Rendering fur directly into images, *Computers & Graphics*, doi:10.1016/j.cag.2010.06.004

Rendering Fur Directly into Images

Tania Pouli[1], Martin Pražák[2], Pavel Zemčík[3], Diego Gutierrez[4], Erik Reinhard[1]

Abstract

We demonstrate the feasibility of rendering fur directly into existing images, without the need to either painstakingly paint over all pixels, or to supply 3D geometry and lighting. We add fur to objects depicted in images by first estimating depth and lighting information and then re-rendering the resulting 2.5D geometry with fur. A brush-based interface is provided, allowing the user to control the positioning and appearance of fur, while all the interaction takes place in a 2D pipeline. The novelty of this approach lies in the fact that a complex, high-level image edit such as the addition of fur can yield perceptually plausible results, even in the presence of imperfect depth or lighting information.

Key words: Picture/Image Generation, Scene Analysis

Email addresses: pouli@cs.bris.ac.uk (Tania Pouli), prazakm@cs.tcd.ie (Martin Pražák), zemcik@fit.vutbr.cz (Pavel Zemčík), diegog@unizar.es (Diego Gutierrez), reinhard@cs.bris.ac.uk (Erik Reinhard)

[1]University of Bristol
[2]Trinity College Dublin
[3]Brno University of Technology
[4]Universidad de Zaragoza, I3A

Figure 1: A variety of realistic fur styles can be achieved through simple user interaction.

1. Introduction

Many applications require the editing of images. Notorious are the touch-ups of photographs in glossy magazines, but also in the film industry effects post-processing is taking ever more advanced forms. A trend in computer graphics is the emergence of algorithms that enable high-level image manipulations [19], including adjustment of the lighting in images [6], re-texturing of objects [7, 29], and more general material replacements [11].

These algorithms typically require some form of depth extraction [30], which on the basis of a single image is an under-constrained problem. Thus,

approximate results are inevitable. The key to successful image manipulations is therefore not to achieve physical accuracy, but aim for visual equivalence [24], i.e. the resulting images may be physically wrong, but should appear perceptually plausible. The human visual system (HVS) helps here, as it is in some cases unable to accurately predict specific features in cluttered environments [21]. At the same time, the HVS makes several assumptions on the nature of the environment, for instance that the geometry is internally consistent [12] or that the viewpoint is chosen in accordance with human physique [9]. Finally, globally convex objects are normally perceived under diffuse lighting according to the *dark-is-deep* paradigm [13], suggesting a relationship between luminance and shape.

After depth extraction, lighting and/or shading can be estimated from the image, for instance by using the background of the object that is being edited [11]. An environment map can be constructed using inpainting to remove the object [3, 5, 26]. Recovery of lighting is a necessarily under-constrained problems as well. With approximate lighting and geometry recovered, new materials can be inserted, and the rendering equation can be re-evaluated, possibly using importance sampling on the recovered environment map to determine a selection of relevant light sources [20]. The result is then an image where the materials of objects have been replaced, but the lighting and geometry are preserved as well as possible.

We are interested in rendering fur directly into images, as shown in Figure 1, for several reasons. Fur is a feature that would be very difficult to draw by hand. Although modern image editing applications make such edits possible, it is a difficult and time consuming process even for the most skillful

artists. Several commercial packages exist to help artists create convincing-looking hair and fur images, based on a traditional 3D representation of a scene (such as *Maya's Paint Effects*, or *Shave and a Haircut*). In contrast to these tools, our method is designed as a post-production process, and uses a single 2D image as input. Our fur editing technique does not require the artist to switch to a 3D environment, and thus it could be easily integrated in any existing 2D pipeline.

Further, the semi-procedural addition of geometry to an image raises the question of how well the new geometry itself would be able to provide masking effects [8], which are required to camouflage the limitations of the depth extraction algorithm. We investigate how different materials affect the estimation of depth and in turn how these effects are counteracted by varying fur properties.

2. Algorithm

Adding fur to an image requires several stages of processing, beginning with image analysis, optional creative user input, fur generation, and rendering. Each of these steps is discussed in the following sections.

2.1. Image Analysis

Before any rendering can take place, ideally a 3D description of the scene would be required. However, the best we can achieve with a single image as input is the computation of an approximate 2.5D depth map, describing the distance between the camera and the nearest surface for every pixel of the input image. This is the classic shape-from-shading problem, which is normally formulated to enable recovery of 3D shape from an image and

Figure 2: Depth maps are rendered from the original viewpoint (top left), which helps to mask artifacts that are more visible from other viewpoints (bottom row). (Note that for illustration purposes, all depth maps in the paper are rendered as a surface with a single light source.)

potentially allow re-rendering of that shape from a different viewpoint. In our case, we will always re-render the image from its original viewpoint, which relaxes the requirement for accuracy somewhat. This is due to the bas-relief ambiguity, the difficulty the HVS appears to have to distinguish surfaces that are related by affine transformations along the line of sight [2]. Thus, oblique illumination causing shearing of the depth map may go unnoticed (Figure 2). These observations were exploited by Khan et al. [11] where luminance values were filtered in order to create a depth map. We follow their approach, which begins by converting the luminance channel of the image into sigmoidal space, followed by the application of a bilateral filter [27]. The sigmoid is then inverted, followed by an optional reshaping step of the luminance signal to obtain the final depth map. This sequence of processing

is sufficient to produce adequate depth maps. The bilateral filtering stage is used exclusively to remove high-contrast high-frequency content, which is typically associated with textured surfaces, and therefore does not carry shape-related information.

2.2. Parametric Fur Generation

We provide two approaches to the generation of fur. In the first case, a user-supplied matte can be used to specify where in the image fur should be inserted. The depth map is then computed for this region, as shown in Section 2.1, and hairs are created and attached to the image algorithmically. To achieve this, each 2×2 block of depth values is converted into two triangles, thereby converting the depth map into a polygonal mesh. To each triangle we attach a number of hairs, determined by the required density, a random variation, and the area of the polygon. With the aid of barycentric coordinates defined on the triangle, the placement of the hairs is randomized as well:

$$\mathbf{h} = \xi_1 \mathbf{p}_0 + \xi_2 \mathbf{p}_1 + (1 - \xi_1 - \xi_2)\mathbf{p}_3 \tag{1}$$

where $\mathbf{p}_0$, $\mathbf{p}_1$ and $\mathbf{p}_2$ correspond to the vertices of the triangle and ξ_1, ξ_2 are two variables randomly selected from a uniform distribution in the range $[0, 1]$ such that $\xi_1 + \xi_2 \leq 1$. From the polygonal model, surface normals are computed for each vertex. These are then linearly interpolated and normalized to compute the surface normal $\mathbf{n}$ for each hair position $\mathbf{h}$. Finally, we compute a desired length for each hair, which can either be specified as a single numeric value, or by supplying a per-pixel map.

A hair is represented by a point $\mathbf{h}$ indicating the position of the root, the surface normal $\mathbf{n}$, a direction randomization vector $\mathbf{r}$ and length l. These parameters, as well as a "gravity" vector $\mathbf{g}$, which bends hair in a user-specified direction, form the input to a particle-based model which then generates the hair by Euler integration [16]. For relatively short hair, a simple Newtonian model suffices:

$$\frac{d^2}{dt^2}\mathbf{p}(t) = \mathbf{g}\ l,\ \text{with}\ \mathbf{p}(t_0) = \mathbf{h},\ \frac{d}{dt}\mathbf{p}(t_0) = \frac{\mathbf{n}+\mathbf{r}}{\|\mathbf{n}+\mathbf{r}\|}l \tag{2}$$

where the position $\mathbf{p}(t)$ represents the resulting function of particle simulation, with discrete values t determined by the integration step (usually 0.02 for normalized vector $\mathbf{n}$ and $\|\mathbf{g}\| < 10$), with the direction randomization vector $\mathbf{r}$ satisfying $\|\mathbf{r}\| < 0.1$.

For longer hair, the interaction between individual hairs as well as the underlying mesh should be considered, and a more advanced particle model involving springs and friction would be more appropriate. However, for fur we ignore these interactions as their effect can be achieved by randomizing the parameters for each hair, which significantly speeds up the simulation.

The result of the particle model for each hair is a set of points representing its shape which is then converted to a cubic Bézier curve. For the purpose of rendering, we represent these curves as flat ribbons that are always facing the camera [4]. This approach is taken to minimize aliasing artifacts. Figure 3 shows a detailed view of the fur.

2.3. Creative User Input

On living creatures, hair and fur shows many irregularities that are difficult to represent procedurally. In particular, the orientation of fur varies

Figure 3: Fur is rendered as flat ribbons facing the camera. Some random luminance and length variation improves the photorealism of the fur.

across the body. Further, the physical values for the density, length or thickness of the hair do not translate directly into parameters that can be inserted into our algorithm. This is due to the fact that by necessity the units our algorithm uses are pixel values. The required hair density or length, on the other hand, depend on the resolution of the photograph and the distance between the object and the camera. Similar arguments can be made for other hair-related parameters such as hair thickness.

In addition, the design of fur may require creative input, determining the average length, direction, and density, for the purpose of telling a story. Therefore, for both technical and creative reasons, it is desirable to have control over the placement and appearance of fur, beyond the simple adjustment of parameters.

To this end, we have implemented a simple-to-use interface where the user can draw and customize fur directly on the image. Brush strokes control the positioning and direction of the fur. Additionally, using a set of sliders the hair length, color, density and flatness can be adjusted for each stroke (Figure 4).

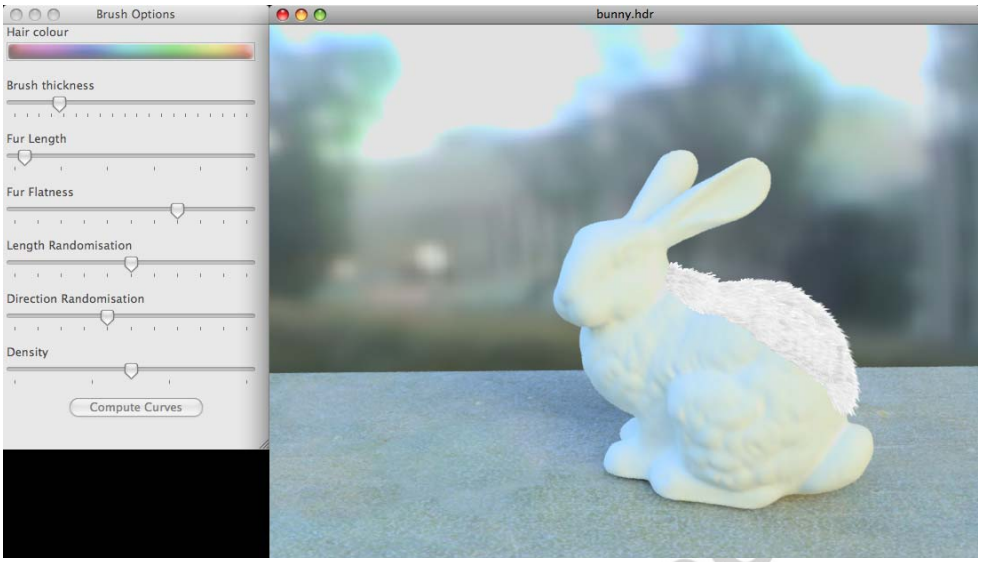

Figure 4: Our tool displays the image on which fur is drawn using brushes (right). The parameters to the brush are also shown (left). Note that as depth is available at the time of drawing, the effect of painting fur on a surface is created.

With the use of the depth map, the brush strokes are converted into a collection of individual hairs. The stroke itself is given by a collection of points $\mathbf{s}_i$, which may optionally be smoothed. From these points we compute normalized stroke direction vectors $\mathbf{d}_i = (\mathbf{s}_{i+1} - \mathbf{s}_i)/\|\mathbf{s}_{i+1} - \mathbf{s}_i\|$. These vectors replace the gravity vector $\mathbf{g}$ in Equation (2). Hair generation is otherwise identical to the algorithm outlined in the preceding section. This approach maps the direction of the brush stroke to the direction of the hair, while taking into account the shape of the object. In essence, the brush stroke can literally be viewed as a hair brush!

2.4. Rendering

The rendering system is supplied with a potentially very large number of hairs which have to be rendered and then composited back into the original image. As we are placing hairs into photographs, we require a state-of-the-art rendering algorithm so that the rendered content will be difficult to differentiate from the photograph. This aspect of our application therefore requires much higher accuracy than was, for instance, required for depth recovery.

One of the earliest methods for fur rendering is due to Kajiya and Kay [10], using three dimensional textures. Van Gelder and Wilhelms [28] manipulate fur in real time but their model can only deal with polylines. The work by Lengyel et al. [15] allows for real-time rendering of fur over arbitrary geometries, at the cost of a double pre-processing step. Lapped texture patches are computed to parameterize the surface [23], followed by the creation of shell textures [14]. Fur appearance at the silhouette of the objects is improved by rendering additional texture cards normal to the surface near such silhouettes. In contrast, the off-line rendering part of our method handles silhouettes naturally and does not require any geometry or texture pre-processing.

Recent work by Zinke et al. [31] allows for real-time rendering of very high-quality hair images by using "aggressive simplifications" of the complex scattering phenomena involved. Their results are on-par with path tracing results, although they also require some pre-computations: as a result, the preprocessing cannot be executed for every frame, precluding our goal for interactive drawing and pre-visualization of hair.

We experimented with two high quality rendering models, namely the

model first proposed by Kajiya and Kay [10], as extended by Banks [1], and the more complex model proposed by Marschner et al. [18]. We have found that the visual quality provided by the latter is closer to the appearance of human hair whilst the first model corresponds better to the appearance of fur.

Of course, this model does not render at interactive rates, which is why we have fitted our user-interface with a fast pre-visualization to give the designer instant feedback. While recently real-time hair rendering algorithms have become available, they still require a few seconds of preprocessing [15, 31]. As a result, the preprocessing cannot be executed for every frame, making interactive drawing and rendering of hair in full quality impossible. The exception is the opacity map based technique augmented with approximate sorting [25], which could be employed in our system.

However, the purpose of our work is to demonstrate the feasibility of augmenting photographs with macro-scale structure. The choice of renderer does not affect our ability to draw conclusions. For simplicity, we therefore resort to a non-photorealistic pre-visualization, which does not require any preprocessing, and is still accurate enough to allow the designer to envisage the final rendered result. This compromise is necessary to create high quality renderings, which are demonstrated in the following section.

The lighting can either be user-specified if high levels of control are needed, or can be derived from the background pixels in the image. In the latter case, the computation of light directions proceeds in a straightforward manner, given that human vision is relatively insensitive to this parameter [21]. As the color of the light source is much more important, we

(a) Lambertian (b) Blinn (c) Phong

Figure 5: The same model was rendered using three different materials, namely a Lambertian surface (a), a specular Blinn surface (b) and a Phong surface (c).

follow the approach of Khan et al. [11], and subject the background pixels to importance sampling to derive a small set of light sources [20]. Finally, the rendered results are composited into the original image using standard compositing techniques [22].

To ensure that satisfactory results can be achieved even with a low fur density, a simple flag is exposed to the user. When enabled, a mostly diffuse surface matching the color of the fur at each location is rendered underneath the fur.

3. Analysis

As outlined in Section 2.1, shape is estimated based on the luminance values of the object of interest in the image. To demonstrate how different materials can affect the accuracy of the recovered depth we rendered the same scene using three different material properties for the horse, approximating Lambertian, Phong and Blinn surfaces. All other scene parameters were the same for all three cases. The rendered images can be seen in Figure 5.

The depth map corresponding to each material was recovered using the same settings across all cases (Figures 6 and 7). Figure 6a shows the ground

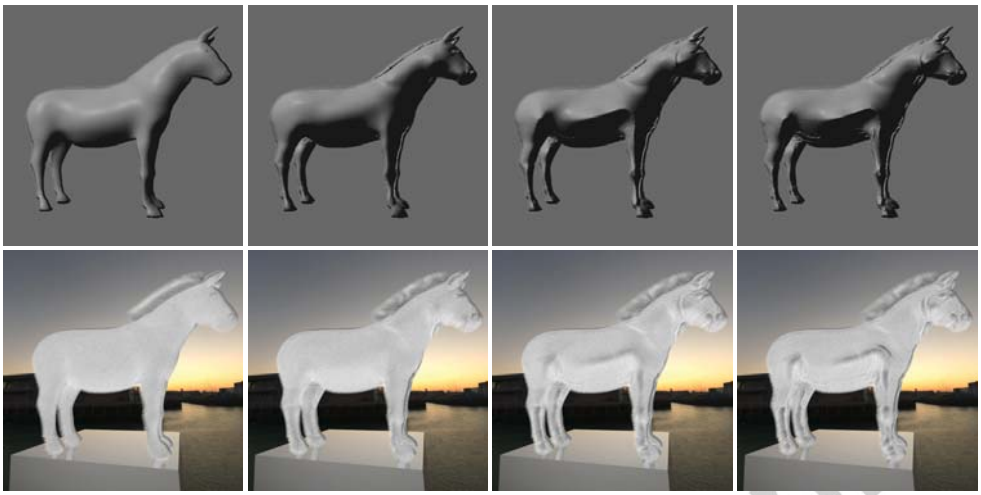

(a) Actual depth (b) Lambertian (c) Blinn (d) Phong

Figure 6: Depths recovered for each of the three materials using the same settings for all cases are shown at the top. The depth recovery parameters were chosen such that any small artifacts would be smoothed out but not highlights due to material specularity. The corresponding fur renderings can be seen at the bottom.

193 truth geometry for the object. Note that all the depths shown in the paper
194 are re-rendered using a specular material for demonstration purposes. As can
195 be seen, the diffuse surface (Figure 6b) allows for the most accurate depth
196 estimation as the properties of the surface are the closest to the assumptions
197 used by our shape recovery algorithm. In the other two cases however, the
198 material properties selected create specular highlights, which when converted
199 to depth, appear as sharp peaks on what should be a smooth surface.
200 Depth was also recovered for each material using the best settings for
201 each case. These settings were chosen manually so that the reconstructed
202 geometry would be as close as possible to the ground truth (Figure 7).
203 To evaluate the masking properties of fur, we rendered fur on each of the

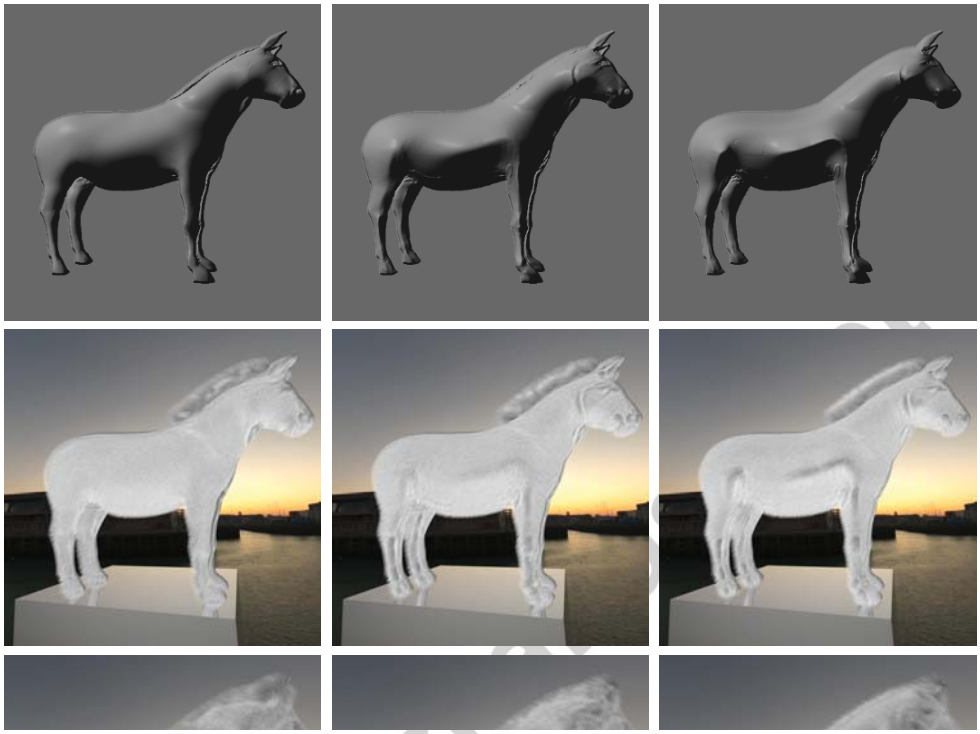

Figure 7: The recovered depths for each of the three materials are shown at the top. The depth recovery parameters were manually chosen for each material to reconstruct the geometry as well as possible. The second and third rows show the corresponding fur renderings using short and longer fur.

Figure 8: Realistic moustaches can also be achieved with our system as they have similar properties to fur. The images on the left show the original photographs.

recovered depth maps for both of the cases described above (Figures 6, 7). In the latter case, fur was rendered using two different lengths that were varied by approximately an order of magnitude. The HDR Visual Differences Predictor (HDR VDP) [17] was employed to compare each image with the respective ground truth. Detailed results and discussion of these comparisons can be found in Appendix A.

4. Results

Parametric fur rendering has the advantage that little user input is required to generate plausible fur. A variety of styles can be achieved through the provided controls as shown throughout the paper (e.g. Figures 1, 8, 9, 10 and 11). Fur of different colors or lengths can be combined in the same image in order to create the desired effect. Additionally, with the aid of user-drawn masks, it is possible to modulate hair parameters to achieve complex effects, as shown in Figure 12.

Figure 9: The material of the left cushion is changed to short fur. The color of the fur is kept the same as the original material (shown on the left).

Figure 10: Fur can be rendered on selected areas of the image. Properties such as color or length can be customized using our interface to achieve a variety of results. The original image is shown at the top left.

Figure 11: Mattes can be used to modulate hair parameters. The matte shown (center) is used to alter hair length (right). The left image shows the unmodulated result.

4.1. Limitations

As discussed previously, our algorithm performs best when the object's material approximates Lambertian surface properties. This should be no surprise as the depth estimation is based on the assumption that luminance relates to depth. This assumption does however break in several occasions.

The simplest case is the presence of specular highlights on the object surface, which can cause spikes in the recovered depth. These can be removed by employing a simple scheme such as the one proposed by Khan et al. [11]. More challenging scenarios arise when complex materials are present. For instance transparent or translucent surfaces would not satisfy the depth recovery assumption, leading to incorrect shapes. Additionally, highly reflective surfaces would require reflection separation in order for the luminance of the surface to be usable.

Other challenging cases arise when the illumination in the scene is highly

Figure 12: Examples of the wide variety of styles that can be achieved using our fur drawing tools. The original image is shown at the top left and the re-rendered depth (shown at a skewed angle) at the bottom left.

directional. If a strong spot light source is lighting the object, self-shadowing can occur, again breaking the original assumption that luminance relates to depth. This creates discontinuities in the depth map where it should be smooth.

We have demonstrated that with limited user input, the semi-procedural addition of macrostructure to an image is possible. Our approach relies on the assumption that the luminance values in the image relate to the depth of the original 3D scene. Although this assumption does not hold when complex, specular surfaces or directional light sources are present, we have shown that it holds for a sufficiently large set of objects.

5. Conclusions

High level image editing, such as material replacements, require a collection of relatively advanced image analysis and rendering algorithms. Some of the algorithmically most troublesome components are shape-from-shading and the detection of appropriate illumination. However, human vision is relatively insensitive to errors in illumination, and the masking effects of the rendered fur allow us to have less than picture-perfect depth estimation.

We have found that objects with diffuse material properties lend themselves best to shape-from-shading algorithms and result in a much more accurate depth reconstruction compared to specular surfaces. In our opinion, (near-) Lambertian surfaces occur often enough for our technique to be practical. As we rerender without changing the viewpoint, our material edits can yield realistic results for both simple and more complicated scenes, including humans, thereby adding to the palette of tools available to 2D artists and designers.

Although we found that specular highlights can cause visible artifacts in the fur rendering, we believe that a possible extention could be to apply a highlight removal pre-process. This could mitigate the effect of shape inaccuracies on the rendered imagery. Further, if and when faster or better fur rendering algorithms become available, these could replace the choice of renderer currently employed without requiring changes to the remainder of our system. Finally, it may be possible to extend this work to include further high-level image edits, perhaps including object deformations, and the addition of more complex geometry.

References

[1] D. C. Banks. Illumination in diverse co-dimensions. In *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques*, pages 327–334, 1994.

[2] P. Belhumeur, D. Kriegman, and A. Yuille. The bas-relief ambiguity. *International Journal of Computer Vision*, 1:33–44, 1999.

[3] M. Bertalmio, G. Sapiro, V. Caselles, and C. Ballester. Image inpainting. In *SIGGRAPH '00: Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques*, pages 417–424, 2000.

[4] A. Daldegan and N. Magnenat-Thalmann. Creating virtual fur and hair styles for synthetic actors. In N. Magnenat-Thalmann and D. Thalmann, editors, *Communicating with Virtual Worlds*, pages 358–370. Springer-Verlag, 1993.

[5] I. Drori, D. Cohen-Or, and H. Yeshurun. Fragment-based image completion. *ACM Transactions on Graphics*, 22(3):303–312, 2003.

[6] E. Eisemann and F. Durand. Flash photography enhancement via intrinsic image relighting. *ACM Transactions on Graphics*, 23(3):673–678, 2004.

[7] H. Fang and J. C. Hart. Textureshop: Texture synthesis as a photographic editing tool. *ACM Transactions on Graphics*, 23(3):354–358, 2004.

[8] J. A. Ferwerda, P. Shirley, S. N. Pattanaik, and D. P. Greenberg. A model of visual masking for computer graphics. In *SIGGRAPH '97: Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques*, pages 143–152, 1997.

[9] W. T. Freeman. The generic viewpoint assumption in a framework for visual perception. *Nature*, 368:542–545, 1994.

[10] J. T. Kayija and T. L. Kay. Rendering fur with three dimensional textures. *Conputer Graphics*, 23(3):271–280, 1989.

[11] E. A. Khan, E. Reinhard, R. W. Fleming, and H. H. Bülthoff. Image-based material editing. *ACM Transactions on Graphics*, 25(3):654–663, 2006.

[12] J. J. Koenderink and A. J. van Doorn. The internal representation of solid shape with respect to vision. *Biological Cybernetics*, 32:211–216, 1979.

[13] M. S. Langer and H. H. Bülthoff. Depth discrimination from shading under diffuse lighting. *Perception*, 29(6):649–660, 2000.

[14] J. Lengyel. Real-time fur. *Eurographics Rendering Workshop*, pages 243–256, 2000.

[15] J. Lengyel, E. Braun, A. Finkelstein, and H. Hoppe. Real-time fur over arbitrary surfaces. In *I3D'01: Proceedings of the 2001 Symposium on Interactive 3D Graphics*, pages 227–232, 2001.

[16] N. Magnenat-Thalmann, S. Hadap, and P. Kalra. State-of-the-art in hair simulation. In *International Workshop on Human Modeling and Animation*, pages 3–9, 2000.

[17] R. Mantiuk, K. Myszkowski, and H. Seidel. Visible difference predicator for high dynamic range images. In *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics*, 2004.

[18] S. R. Marschner, H. W. Jensen, M. Cammarano, S. Worley, and P. Hanrahan. Light scattering from human hair fibers. *ACM Transactions on*

Graphics, 22(3):780–791, 2003.

[19] B. M. Oh, M. Chen, J. Dorsey, and F. Durand. Image-based modeling and photo editing. In *SIGGRAPH '01: Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*, pages 433–442, 2001.

[20] V. Ostromoukhov, C. Donohue, and P. Jodoin. Fast hierarchical importance sampling with blue noise properties. *ACM Transactions on Graphics*, 23(3):488–495, 2004.

[21] Y. Ostrovsky, P. Cavanagh, and P. Sinha. Perceiving illumination inconsistencies in scenes. *Perception*, 34(11):1301–1314, 2005.

[22] T. Porter and T. Duff. Compositing digital images. *Computer Graphics*, 18(3):253–259, 1984.

[23] E. Praun, A. Finkelstein, and H. Hoppe. Lapped textures. In *SIGGRAPH '00: Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques*, pages 465–470, 2000.

[24] G. Ramanarayanan, J. A. Ferwerda, B. Walter, and K. Bala. Visual equivalence: towards a new standard for image fidelity. *ACM Transactions on Graphics*, 26(3):76, 2007.

[25] E. Sintorn and U. Assarsson. Real-time approximate sorting for self shadowing and transparency in hair rendering. In *I3D'08: Proceedings of the 2008 Symposium on Interactive 3D Graphics and Games*, pages 157–162, 2008.

[26] J. Sun, L. Yuan, J. Jia, and H.-Y. Shum. Image completion with structure propagation. *ACM Transactions on Graphics*, 24(3):861–868, 2005.

[27] C. Tomasi and R. Manduchi. Bilateral filtering for gray and color images. In *Proc. IEEE International Conference on Computer Vision*,

pages 836–846, 1998.

[28] A. van Gelder and J. Wilhelms. An interactive fur modeling technique. In *Proceedings of Graphics Interface*, 1997.

[29] S. Zelinka, H. Fang, M. Garland, and J. C. Hart. Interactive material replacement in photographs. In *GI '05: Proceedings of Graphics Interface 2005*, pages 227–232, 2005.

[30] R. Zhang, P. Tsai, J. Cryer, and M. Shah. Shape from shading: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(8):690–706, 1999.

[31] A. Zinke, C. Yuksel, A. Weber, and J. Keyser. Dual scattering approximation for fast multiple scattering in hair. *ACM Transactions on Graphics*, 27(3):32, 2008.

A. The Effect of Different Fur and Material Parameters

As mentioned in Section 3, the same scene was rendered using three different materials in order to evaluate the effect of different fur lengths and surface materials. The depth was recovered for each of these three cases under two different conditions: using the same settings across all cases and using the best settings for each material (derived through manual experimentation). The depth recovery settings for both cases described can be seen in Table 1. The HDR Visual Differences Predictor (HDR VDP) [17] was employed to compare each image with the respective ground truth. Figure 13 shows the results for the renderings that used the same settings across all cases (top row) and the best settings for each material (bottom row) while Figure 14 shows the HDR VDP results for the images with the longer fur. A number of

observations arise from these comparisons. A careful selection of parameters for the depth recovery can improve the results. More specifically, at least 1% fewer pixels were detected by the HDR VDP when using the best settings compared to using the same settings for each of the materials. Further, longer fur increased the percentage of pixels detected as different between each pair. As the roots of the fur are generally darker than the ends, sharp peaks in the geometry can cause the fur direction to change abruptly. This creates the effect of a parting where the darker roots are visible. Longer fur means that these darker regions cover a larger area in the image, resulting in a larger number of pixels appearing different. In particular, we found that 3-5% more pixels were detected by the HDR VDP in this case compared to renderings with shorter fur. Table 2 lists the percentage of pixels with a probability of detection higher than 75% for each of the comparisons. Finally, the results for each of the three materials were compared. Both when the same and the best settings were used, the images rendered using the Lambertian surface were the closest to the ground truth image, with the specular surfaces (Blinn and Phong) resulting in a 2-3% increase in detected pixels in comparison. Diffuse

	Lambertian	**Blinn**	**Phong**
Long	4.99%	6.87%	8.14%
Same	2.09%	4.47%	5.60%
Best	1.85%	3.60%	3.64%

Table 1: Percentage of pixels with a probability of detection $P > 75\%$ when compared with the respective ground truth image for the same settings, best settings and long fur settings cases.

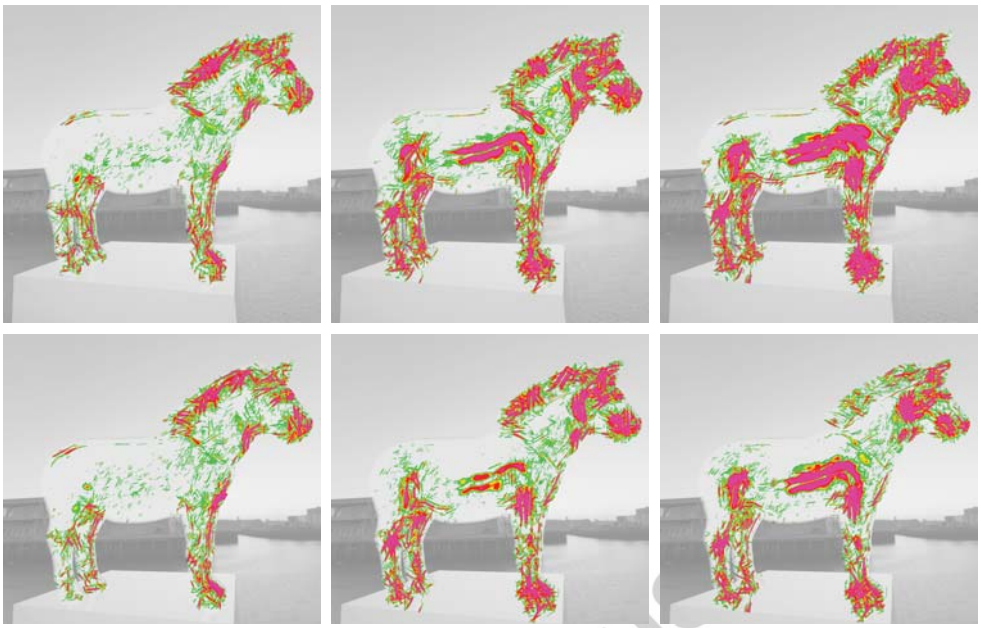

Figure 13: Comparisons of each of the materials with the ground truth using the same depth recovery settings for all cases (top) and the best settings for each material (bottom). The images were compared using the HDR VDP [17].

surfaces, such as the Lambertian example, are much closer to the assumptions employed by our depth recovery algorithm, resulting in a smoother depth map. Highlights in specular surfaces on the other hand can cause sharp peaks on the recovered depth resulting in more visible artifacts when fur is rendered on them. Nevertheless, using shorter fur, most of the results of our algorithm contained under 5% of visibly different pixels when compared to the ground truth, demonstrating that inaccuracies in the depth map can be successfully masked to a large extent by the high frequency properties of the fur.

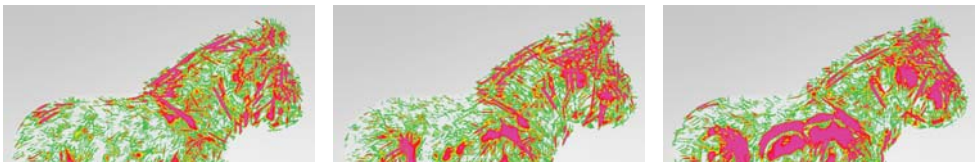

Figure 14: Comparisons of each of the materials with the ground truth for the renderings with the longer fur.

	σ_s	σ_i	**Depth multiplier**
Same	1.2%	16.8%	0.35
Best Lambert	1.3%	12.5%	0.27
Best Blinn	1.4%	65.8%	0.27
Best Phong	2.2%	33.2%	0.21

Table 2: Bilateral filter and depth multiplier parameters used in the two cases described in this section. The parameters for each of the materials are given in the case where the best settings were chosen. Both σ_s and σ_i are given as a percentage of the width of the image and the correspond to the spatial and intensity kernel sizes of the bilateral filter.