Modeling Vellus Facial Hair from Asperity Scattering Silhouettes

Chloe LeGendre
USC Institute for Creative Technologies

Loc Hyunh
USC Institute for Creative Technologies

Shanhe Wang
USC Institute for Creative Technologies

Paul Debevec
USC Institute for Creative Technologies

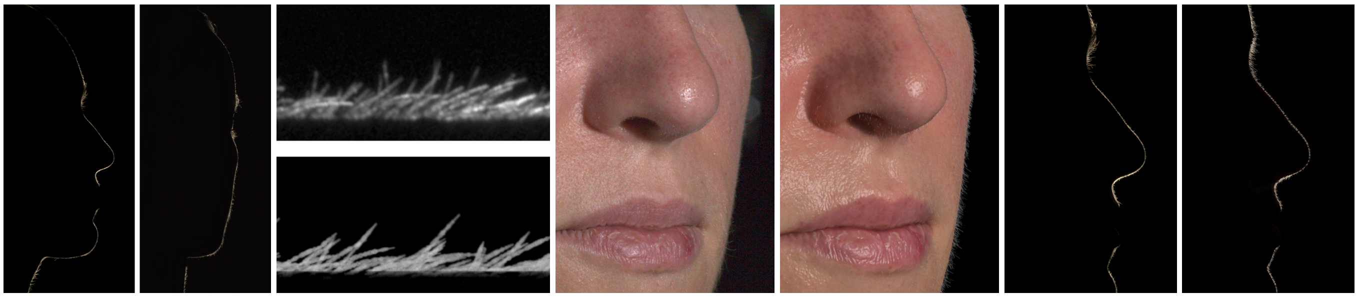

(a) silhouette photos (b) real/rendered patches (c) photo (d) rendering (e) photo (f) rendering

Figure 1: (a) Photographed backlit facial silhouettes. (b) Top: Vellus hair segment sampled from the straightened silhouette; Bottom: A rendering from a library of synthetic backlit vellus hair patches which best matches the top's image statistics. (c,e) Face reference photographs. (d,f) 3D face model rendered with vellus hair generated using our method.

ABSTRACT

We present a technique for modeling the vellus hair over the face based on observations of asperity scattering along a subject's silhouette. We photograph the backlit subject in profile and three-quarters views with a high-resolution DSLR camera to observe the vellus hair on the side and front of the face and separately acquire a 3D scan of the face geometry and texture. We render a library of backlit vellus hair patch samples with different geometric parameters such as density, orientation, and curvature, and we compute image statistics for each set of parameters. We trace the silhouette contour in each face image and straighten the backlit hair silhouettes using image resampling. We compute image statistics for each section of the facial silhouette and determine which set of hair modeling parameters best matches the statistics. We then generate a complete set of vellus hairs for the face by interpolating and extrapolating the matched parameters over the skin. We add the modeled vellus hairs to the 3D facial scan and generate renderings under novel lighting conditions, generally matching the appearance of real photographs.

CCS CONCEPTS

•Computing methodologies →Image-based rendering;

KEYWORDS

Vellus hair, Asperity scattering, Digital humans

DOI: http://dx.doi.org/10.1145/3084363.3085057

ACM Reference format:
Chloe LeGendre, Loc Hyunh, Shanhe Wang, and Paul Debevec. 2017. Modeling Vellus Facial Hair from Asperity Scattering Silhouettes. In *Proceedings of SIGGRAPH '17 Talks, Los Angeles, CA, USA, July 30 - Aug 03, 2017,* 2 pages.
DOI: http://dx.doi.org/10.1145/3084363.3085057

1 INTRODUCTION

Vellus facial hair, or "peach fuzz", contributes to skin's signature "velvety" appearance [Koenderink and Pont 2003]. These nearly transparent hairs, less than 0.03mm in diameter and up to 2mm in length, scatter light along the surface of the face through *asperity scattering*. Experienced facial modeling artists will manually add vellus hair to their facial models to increase the realism of a digital character. While techniques have been presented to derive individualized models of the hair on the head from photographs (e.g. [Beeler et al. 2012; Paris et al. 2004]); deriving a vellus hair model particular to a given individual has remained elusive since these tiny hairs are hard to individually observe in photographs.

2 SILHOUETTE ACQUISITION

Vellus hairs are easiest to see along the silhouette of the face when lit from behind, due to *asperity scattering* [2003]. To image a subject's vellus hairs, we photograph the face with a Canon 1DX Mark II DSLR camera at a distance of 1.5m with a 180mm macro lens and a bright incandescent spotlight 2m behind their head. We capture both profile and three-quarters views as shown in Fig. 1a.

3 VELLUS HAIR MODEL

We render a library of backlit vellus hair patch samples with different values for various geometric parameters, including hair density, hair length, standard deviation of the hair length, principal orientation, and standard deviation of the orientation, creating a library of 9,216 total images. Facial curvature impacts the observed density of the hair fibers along the silhouette contour, so we render vellus hair

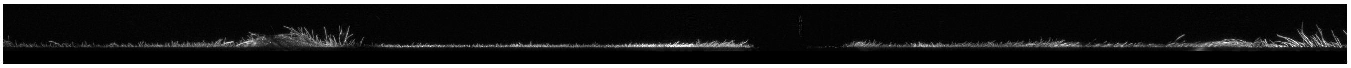

Figure 2: Straightened backlit facial silhouette made by resampling the original photo along lines orthogonal to the contour.

patches on cylinders of various radii. Each hair fiber is modeled as a 2^{nd} degree NURBS curve with three control points. We use the Arnold® ray tracer and the open source curve shader alHair[1] based on Marschner et al. [2003]. Our supplemental material shows the full list of geometric parameters and their value ranges, along with shader parameter values and example library images.

4 IMAGE STATISTICS

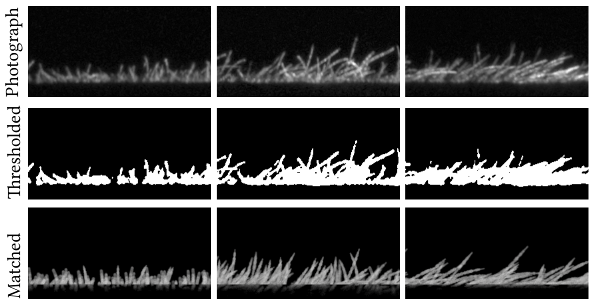

Figure 3: Top: Photographed backlit vellus hair. Middle: Binary threshold of top. Bottom: Matched library renderings.

We compute a set of image statistics for each rendering in the library: **Orientation**: We convolve each image with a set of oriented second-derivative-of-Gaussian filters (every 10°), scaled to expected hair width ($\sigma = 1.5$ pixels). We compute an orientation histogram for the image by computing for each pixel the orientation of the maximum filter response and then adding the response magnitude to the appropriate orientation bin. We compute the (1) mean and (2) standard deviation of this orientation histogram. **Length and area**: We threshold the images to binary and count the (3) total number of white pixels. We also sum across the binary image rows to produce a single column image, computing the (4) mean and (5) standard deviation of this energy histogram. We normalize the image statistics, dividing by the range for each parameter.

5 PARAMETER FITTING

We resample each silhouette image along the facial contour to create a straightened vellus hair profile as in Fig. 2. We then break the straightened silhouette into 1cm linear segments as in Fig. 3. We determine the out-of-plane radius of curvature for each segment by intersecting a plane with the 3D model of the subject at the center of each segment. We then threshold each segment into "hair" and "non-hair" pixels based on the average pixel value from the region below the localized silhouette as shown in Fig. 3; this allows us to factor out the strength of the scattered light. We then compute our image statistics on each segment. We find the closest matching hair patch in the rendered library with the same facial curvature and nearest image statistics, using a Euclidean distance metric in the 5-dimensional parameter space. Example library matches are shown in the bottom row of Fig. 3, with others shown in supplemental material. We retrieve the geometric vellus hair parameters for each match, sparsely populating a map of the face with these values and mirroring the three-quarters view values across the profile. We then interpolate/extrapolate the sparse parameter values across the entire facial surface. Finally, we procedurally generate vellus hair on a low-res 3D model of the subject, using stratified sampling to position individual hair fibers on each model facet, assigning hair length, standard deviation of length, fiber density, orientation and curvature according to the geometric parameter maps.

[1]http://www.anderslanglands.com/alshaders/alHair.html

6 RESULTS

Using this technique we generated 104,752 vellus hairs over the scan's face and neck and made several renderings in Arnold. We tested the model's ability to match a photographed backlit lighting condition by comparing Figs. 1e and 1f. While we generate a different population of hairs than the photograph, the overall lengths, orientations, and local variability match reasonably well across the silhouette. We also explored the case of matching frontal lighting in Figs. 1c and 1d; Fig. 4 shows a detail of the real and rendered vellus hairs in a shadow terminator. While some differences in skin shading are noticeable, the placement and appearance of vellus hairs are similar. Additional results are presented in supplemental material. In future work, we would like to use additional silhouette images for a more complete model of the hairs over the face.

Figure 4: Photo (left) and rendering with vellus hair (right).

Acknowledgements

The authors wish to thank Randy Hill, Kathleen Haase, Christina Trejo, Marcel Ramos, and Dai Liu for their support of this work. This work was sponsored by the U.S. Army Research Laboratory (ARL) under contract W911NF-14-D-0005. The content of the information does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred.

REFERENCES

Thabo Beeler, Bernd Bickel, Gioacchino Noris, Paul Beardsley, Steve Marschner, Robert W Sumner, and Markus Gross. 2012. Coupled 3D reconstruction of sparse facial hair and skin. *ACM Transactions on Graphics (TOG)* 31, 4 (2012), 117.

Jan Koenderink and Sylvia Pont. 2003. The Secret of Velvety Skin. *Mach. Vision Appl.* 14, 4 (Sept. 2003), 260–268.

Stephen R Marschner, Henrik Wann Jensen, Mike Cammarano, Steve Worley, and Pat Hanrahan. 2003. Light scattering from human hair fibers. In *ACM Transactions on Graphics (TOG)*, Vol. 22. ACM, 780–791.

Sylvain Paris, Hector M Briceño, and François X Sillion. 2004. Capture of hair geometry from multiple images. In *ACM Transactions on Graphics (TOG)*, Vol. 23. ACM, 712–719.