

Processing and Editing of Faces using a Measurement-Based Skin Reflectance Model

Bernd Bickel * Tim Weyrich * Wojciech Matusik † Hanspeter Pfister † Craig Donner ‡ Chien Tu †
Janet McAndless † Jinho Lee † Addy Ngan § Henrik Wann Jensen ‡ Markus Gross *

Introduction

In extension to our paper *Analysis of Human Faces using a Measurement-Based Skin Reflectance Model* we describe the almost fully automated face processing pipeline including the custom-built acquisition system, calibration procedures, and rendering. The underlying model is intuitive, physically plausible, and allows easy face editing. Generated image quality approaches that of real photographs.

Data Acquisition

We use a custom built measurement-device consisting of a geodesic dome with 150 LED light sources and 16 cameras. Additionally, a commercial face-scanning system from 3QTech (www.3dmd.com) captures the complete 3D face geometry. Both systems have to be calibrated carefully. This involves vignetting, color calibration, and light source calibration, including the individual intensity profile of each light source. External camera calibration is performed by sweeping a red LED through the volume of interest, and a non-linear bundle adjustment is used to derive camera parameters that are consistent with a single 3D trajectory of the LED. The 3QTech scanner has its own proprietary calibration procedure. A correspondence between 3QTech and camera coordinate system is established by scanning a calibration target containing markers with both systems. We apply the Procrustes method to find a rigid transformation and scale between them.

The system sequentially turns on each light while simultaneously capturing images with all 16 cameras. A complete scanning sequence takes about 25 seconds, acquiring an HDR reflectance field of the human face. We clean the output mesh by manually cropping non-facial areas and fixing non-manifold issues. Loop-subdivision is applied to obtain a mesh with about 800,000 vertices. For further processing, the reflectance field is re-parametrized at the vertices to form lumitexels. On average, a lumitexel contains about 900 BRDF samples per color channel, with many up to 1,200 samples.

Geometry Processing

The acquired 3D mesh does not accurately capture fine facial details such as wrinkles. We estimate normals at each lumitexel from the reflectance data using photometric stereo, for each camera view a normal map is computed. Using the method by Nehab et al. [2005] we correct low-frequency bias of the normals. For each camera view, the cosines between viewing rays and surface normals are stored as a blending weights texture. After smoothing discontinuities in the weight textures we compute the final *diffuse normal map* by averaging all normals at each vertex weighted with their blending weights. Finally, we compute a high-resolution geometry by

applying the method of Nehab et al. As specular highlights are influenced by micro-geometry such as pores or very fine wrinkles, that is not accurately captured in the face geometry or diffuse normal map, we use *micro-normals* to render the specular component. For each camera viewpoint the sample with maximum value is chosen at each lumitexel. The halfway vector between the view direction and the sample's light source direction is the micro-normal for this view. We obtained the best results storing these normals in per-camera micro-normal maps and using angular interpolating for between-camera views.

Measurement-Based Skin Reflectance Model

Our model decomposes the measured skin data into a spatially-varying analytic BRDF, a diffuse albedo map, and diffuse subsurface scattering. We assume homogeneous subsurface scattering and approximate it using the dipole diffusion approximation by Jensen et al. [2001] and a modulation texture. Intuitively, the modulation texture represents an absorptive film of zero thickness on top of a homogeneously scattering material that only influences light on its way out. The specular component is represented by the Torrance-Sparrow BRDF model. All parameters can be estimated robustly from the acquired data. The skin reflectance model is demonstrated using Henrik Wann Jensen's high-quality ray tracer *Dali*.

Face Editing

The skin reflectance model allows fast, easy, and intuitive editing. For example, we can change the skin type of a model by transferring statistics of albedo and reflectance parameters, or use texture synthesis/analysis to transfer local skin features, such as moles and freckles. The combination of these approaches allows us to convincingly age a face or transfer facial appearance properties between individuals.



Figure 1: Left: Side-by-side comparison of a real photograph and a rendered image. Right: Face editing using statistics for transferring albedo and BRDF from a sample population.

References

- JENSEN, H. W., MARSCHNER, S. R., LEVOY, M., AND HANRAHAN, P. 2001. A practical model for subsurface light transport. In *Proc. SIGGRAPH '01*, ACM Press, 511–518.
- NEHAB, D., RUSINKIEWICZ, S., DAVIS, J., AND RAMAMOORTHY, R. 2005. Efficiently combining positions and normals for precise 3d geometry. *ACM Trans. Graph.* 24, 3, 536–543.

*ETH Zürich, Email: [bickel,weyrich,grossm]@inf.ethz.ch

†MERL, Email: [matusik,pfister,ctu,leejh,mcandless]@merl.com

‡UCSD, Email: [cdonner@graphics.henrik@cs].ucsd.edu

§MIT, Email: addy@graphics.csail.mit.edu