USC-ICT Technical Report ICT-TR-01-2012: Measurement-Based Synthesis of Facial Microgeometry

Paul Graham Borom Tunwattanapong Jay Busch Xueming Yu Andrew Jones
Paul Debevec Abhijeet Ghosh

(a) Rendering from original facial scan (b) Rendering using synthesized microstructure (c) Comparison photograph

Figure 1: *(a) Rendering with scanned mesostructure (4K displacement map). (b) Rendering with synthesized microstructure (16K displacement map). (c) Photograph under flash illumination. The digitized skin patches used for microstructure synthesis can be seen in Fig. 3, Subject 2.*

Abstract

We present a technique for generating microstructure-level facial geometry by augmenting a mesostructure-level facial scan with detail synthesized from a set of exemplar skin patches scanned at much higher resolution. We use constrained texture synthesis based on image analogies to increase the resolution of the facial scan in a way that is consistent with the scanned mesostructure. We digitize the exemplar patches with a polarization-based computational illumination technique which considers specular reflection and single scattering. The recorded microstructure patches can be used to synthesize full-facial microstructure detail for either the same subject or to a different subject. We show that the technique allows for greater realism in facial renderings including more accurate reproduction of skin's specular roughness and anisotropic reflection effects.

Keywords: microgeometry, texture synthesis, computational illumination

1 Introduction

The way our skin reflects light is influenced by age, genetics, health, emotion, sun exposure, substance use, and skin treatments. The importance of skin reflectance is underscored by the worldwide cosmetics industry which sells a myriad of products to achieve specific skin appearance results. As virtual human characters become increasingly prevalent in linear and interactive storytelling, the need for measuring, modeling, and rendering the subtleties of light reflection from skin also becomes increasingly important.

While great strides have been made in simulating the scattering of light beneath skin [Hanrahan and Krueger 1993; Jensen et al. 2001; Donner and Jensen 2005; D'Eon and Irving 2011], somewhat less recent attention has been given to surface reflection. While the shine of our skin — the specular reflection from the epidermal cells of the *stratum corneum* — is a small fraction of the incident light, its lack of scattering provides a clear indication of skin's surface shape and condition. And under concentrated light sources, the narrower specular lobe produces highlights which can dominate appearance, especially for smoother, oilier, or darker skin.

Current face scanning techniques [Weyrich et al. 2006; Ma et al. 2007; Beeler et al. 2010] as well as high-resolution scanning of facial casts (e.g. [XYZRGB], [Acevedo et al. 2010]) provide submillimeter precision, recording facial *mesostructure* at the level of pores, wrinkles, and creases. Nonetheless, the effect of surface roughness continues to shape specular reflection at the level of *microstructure* [Torrance and Sparrow 1967] — surface texture at the scale of microns. The current absence of such microstructure may be a reason why digital humans can still appear uncanny in close-ups, which are the shots most responsible for conveying the thought and emotion of a character.

In this work, we present a texture synthesis technique to increase the resolution of mesostructure-level facial scans using surface microstructure digitized from skin samples about the face. We digitize the skin patches using macro photography and polarized gradient illumination [Ma et al. 2007] at approximately 10 micron precision, and we make point-source reflectance measurements to characterize the specular reflectance lobe at this smaller scale. We then employ a constrained texture synthesis algorithm based on Image Analogies [Hertzmann et al. 2001] to synthesize appropriate surface microstructure per-region, blending the regions to cover the whole entire face. We show that renderings made with microstructure-level facial models preserve the original scanned mesostructure and exhibit surface reflection which is significantly more consistent with real photographs of the face.

2 Related Work

Our work builds on a variety of results in skin reflectance modeling, surface detail acquisition, and texture synthesis.

Skin Reflectance Modeling Physically-based models of light reflection from rough surfaces are typically based on a microfacet distribution model [Torrance and Sparrow 1967; Cook and Torrance 1982]. [Marschner et al. 1999] photographed foreheads from a sampling of lighting and viewing directions to tabulate skin BRDFs for rendering. [Cula et al. 2005] used a moving light source and camera to record a bidirectional database of high-resolution skin samples in a clinical setting, building a texton histogram model to detect skin abnormalities. Using high-resolution surface orientation measurements, [Ghosh et al. 2008] fit the microfacet BRDF model of [Ashikhmin et al. 2000] to specular reflectance data across various patches of human faces.

We acquire detailed surface shape and reflectance of example skin patches to increase the resolution of a larger-scale model. This approach echoes other work which extrapolates detailed samples of reflectance over complete models, including [Weyrich et al. 2006] which samples subsurface scattering parameters with a special probe, and [Dong et al. 2010] which uses a BRDF probe to add detailed BRDFs to entire objects observed in relatively few lighting conditions. Our work also relates to the basis BRDF modeling approach of [Lensch et al. 2003] and the reflectance sharing approach of [Zickler et al. 2006].

Surface Detail Measurement Key to our work is measuring surface detail from images taken with a fixed viewpoint and varying lighting. Classic photometric stereo [Woodham 1978] derives surface orientations of Lambertian surfaces from three point light directions. [Rushmeier et al. 1997] used photometric stereo to produce bump maps for increasing the detail on rendered surfaces. For semi-translucent materials such as skin, subsurface scattering blurs the surface detail recoverable from traditional photometric stereo significantly [Ramella-Roman 2008]. [Debevec et al. 2000], [Ramella-Roman 2008], [Weyrich et al. 2006], and [Chen et al. 2006] analyze *specular* surface reflection to obtain more precise surface orientation measurements of translucent materials. [Ma et al. 2007] uses polarization difference imaging to isolate the specular reflection under gradient lighting conditions, allowing specular surface detail to be recorded in a small number of images. The GelSight system [Johnson et al. 2011] pushes silver-coated gel against a sample and uses photometric stereo to record surface microgeometry at the level of a few microns, which achieves the resolution required for our technique. In this work, we measure microgeometry using a polarized gradient illumination technique based on [Ghosh et al. 2011] since it requires no contact with the sample and permits the material's actual reflectance to be observed during measurement for BRDF fitting.

Texture Synthesis A recent review of example-based texture synthesis may be found in [Wei et al. 2009]. To summarize a few results, 2D texture synthesis algorithms (e.g. [Heeger and Bergen 1995], [Efros and Leung 1999]) have been extended to arbitrary manifold surfaces [Wei and Levoy 2001; Ying et al. 2001; Lefebvre and Hoppe 2006] and can also synthesize displacement maps [Ying et al. 2001], measured reflectance properties [Tong et al. 2002] and skin color [Tsumura et al. 2003]. Some techniques [Efros and Freeman 2001; Hertzmann et al. 2001] permit *texture transfer*, transforming a complete image so that is has textural detail of a given sample. Techniques such as [Hertzmann et al. 2001], [Freeman et al. 2002], and [Lefebvre and Hoppe 2005] also permit *super-resolution*, adding detail to a low-resolution image based on examples. Especially relevant to our work are *constrained texture synthesis* techniques [Wang and Mueller 2004; Ramanarayanan and Bala 2007] which add plausible detail to a low-resolution image without changing its low frequency content. In our work, we adapt the Image Analogies [Hertzmann et al. 2001] framework to perform constrained texture synthesis of skin microstructure onto facial scans while preserving scanned facial mesostructure.

Some techniques have been proposed to increase the detail present in facial images and models, although at significantly coarser resolution than we address in our work. The face hallucination technique of [Liu et al. 2001] creates recognizable facial images by performing constrained facial texture synthesis onto a parametric face model. [Liu et al. 2004] uses image processing to apply facial detail from one subject to another, allowing image-based aging effects. Most closely related to our work is [Golovinskiy et al. 2006], which synthesizes skin mesostructure (but not microstructure) learned from higher-quality facial scans onto otherwise smooth facial shapes. In addition to working at a very different scale (we assume that wrinkles, creases, and larger pores are already apparent in the source scan) than our work, the synthesis method is based on matching per-region frequency statistics rather than patch-based texture synthesis, and is not designed to match to existing mesostructure.

Our process of sampling skin microgeometry follows recent work which models or measures materials at the micro scale in order to better predict their appearance at normally observable larger scales. [Marschner et al. 2005] observes an electron micrograph of a cross-section of wood to motivate a reflection model which better predicted anisotropic reflection effects. [Zhao et al. 2011] makes direct use of Micro CT imaging of fabric samples to model the volumetric scattering of complete pieces of cloth.

3 Recording Skin Microstructure

Acquisition We record the microstructure of skin patches using either of two systems to create polarized gradient illumination. For both, we stabilize the skin patch relative to the camera by having the subject place their skin against a 24mm by 16mm aperture in a thin metal plate. This plate is firmly secured 30cm in front of the sensor of a Canon 1D Mark III camera with a Canon 100mm macro lens stopped down to f/16, just enough depth of field to image the sample when properly focused. The lens achieves 1:1 macro magnification, so each pixel of the 28mm $\times$ 19mm sensor images about seven microns (7μm) of skin.

Our small capture system (Fig. 2(a)) is a 12-light dome (half of a deltoidal icositetrahedron) similar to those used for acquiring Polynomial Texture Maps [Malzbender et al. 2001], with the addition that each light can produce either of two linear polarization conditions. The polarization pattern is similar to that of [Ghosh et al. 2011] except that the polarization orientation of each light is specifically optimized for a single camera viewpoint. The difference between images acquired under parallel- and cross polarized states records the polarization-preserving reflectance of the sample, attenuating subsurface reflectance. In approximately two seconds, we acquire polarized gradient illumination conditions to record surface normals. We compensate for any subject motion using the joint photometric alignment technique of [Wilson et al. 2010]. For BRDF fitting, we additionally capture a single-light image in both cross- and parallel-polarized lighting conditions.

For especially smooth or oily skin patches, the twelve light positions can produce separated specular highlights, which can bias surface normal measurement. To address this, for some subjects we placed the macro photography camera and metal aperture frame inside the same 2.5m-diameter polarized LED sphere used for facial scanning (Fig. 2(b)). While the camera occludes some light directions from reaching the sample, the hemispherical coverage of the incident light directions is denser and more complete than the 12-light setup, allowing the gradient illumination to yield well-conditioned surface normal measurements for all patches we tested.

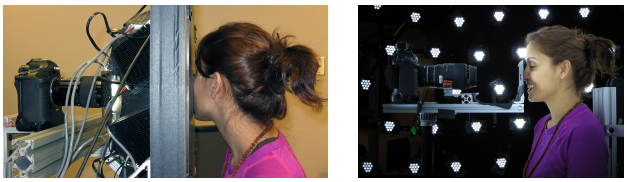

(a) (b)

Figure 2: *Acquisition setups for skin microgeometry (a) Twelve-light hemispherical dome (at end of camera lens) capturing a cheek patch. (b) LED Sphere capturing the tip of the nose, with the camera inside the sphere.*

Since a single LED light source was not bright enough to illuminate the sample for specular BRDF observation, a pair of horizontally and vertically polarized camera flashes (Canon Speedlite 580EX II) were fired to record the point-light condition from close to the normal direction. The camera mounts in both setups were re-enforced to mechanically eliminate vibrations and flexing which would blur the captured imagery.

Surface Normal and BRDF Estimation We compute a per-pixel surface orientation map from the polarized gradients, as well as specular and subsurface albedo maps, as in [Ghosh et al. 2011]. Figure 3 shows the geometry of five skin patches digitized for two subjects, including regions of the forehead, temple, cheek, nose, and chin. Due to the flat nature of the skin patches, we only visualize the x and y components of the surface normals with yellow and cyan colors respectively. Note that the skin microstructure is highly differentiated across both individuals and facial regions.

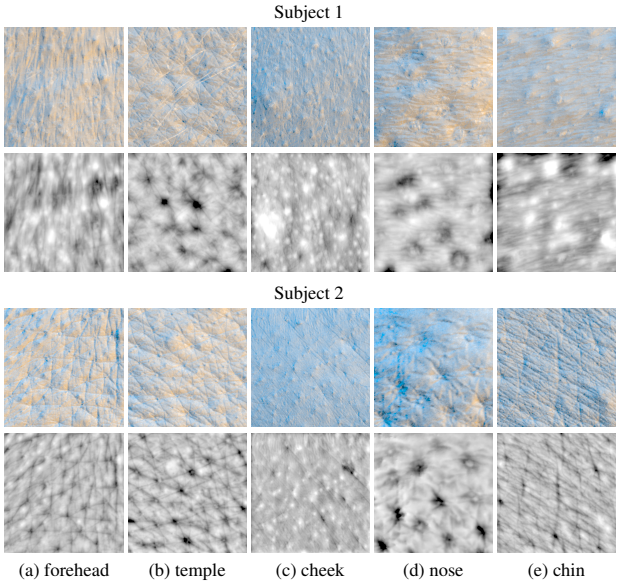

(a) forehead (b) temple (c) cheek (d) nose (e) chin

Figure 3: *Measured skin patches from different facial regions of two different subjects. (Top two rows) Caucasian male subject. (Bottom two rows) Asian female subject. (Rows one and three) Surface normals. (Rows two and four) Displacements.*

Using the polarization difference point-lit image, we also tabulate a specular lobe shape and single scattering model parameters as in [Ghosh et al. 2008]. With light pressure, the skin protrudes slightly through the metal aperture frame, providing a slightly convex surface which exhibits a useful range of surface normals for BRDF estimation. Using the specular BRDF model of [Torrance and Sparrow 1967], we found that *two* lobes of a Beckmann distribution as in [Cook and Torrance 1982] fit the data well. In order to better fit

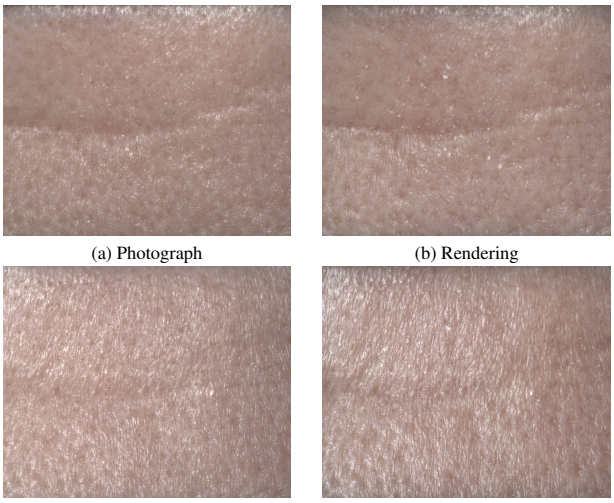

(a) Photograph (b) Rendering

(c) Photograph (d) Rendering

Figure 4: *(a) Parallel-polarized photograph of a forehead skin patch of a male subject, lit slightly from above. (b) Validation rendering of the patch under similar lighting using the surface normals, specular albedo, diffuse albedo, and single scattering maps estimated from the 12-light dome, showing visual similarity. (c,d) Similar images from a different light source direction.*

the specular and single scattering model parameters, we factor the observed polarization preserving reflection under constant full-on illumination into two separate specular and single scattering albedo maps. Here, we estimate the single scattering albedo as the difference between observed polarization preserving reflectance and average hemispherical specular reflectance of a dielectric surface with index of refraction $\eta = 1.33$, which is about 0.063.

Figures 4 and 5 show skin patch samples and validation renderings made using the estimated subsurface albedo, specular albedo, specular normals, and specular BRDF, showing close visual matches of the model to the photographs. At this zoomed-in scale, where much of the surface roughness variation is evident geometrically, we found that a single two-lobe specular BRDF estimate to be sufficient over each sample, and that variation in the reflectance parameter fits were quite modest compared to the differences observed at the mesostructure scale as in [Ghosh et al. 2008].

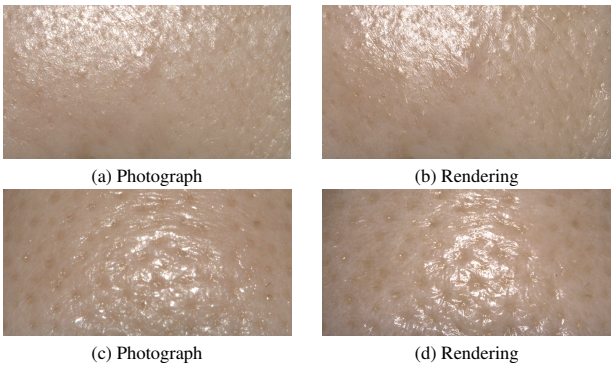

(a) Photograph (b) Rendering

(c) Photograph (d) Rendering

Figure 5: *(a) Parallel-polarized photograph of a cheek patch for a female subject, lit slightly from above. (b) Validation rendering of the cheek patch lit from a similar direction using reflectance maps estimated from the LED sphere, showing a close match. (c,d) Corresponding images from subject's nose patch, also showing a close match.*

4 Facial Microstructure Synthesis

From the skin microstructure samples, we employ constrained texture synthesis to generate skin microstructure for an entire face. To do this, we use the surface mesostructure evident in a full facial scan to guide the texture synthesis process for each facial region, and we then merge the synthesized facial regions into a full map of the microstructure.

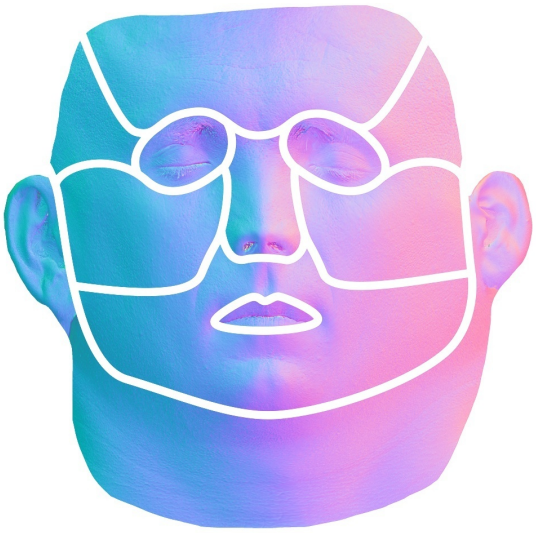

Figure 6: *A world-space normal map from the Subject 1 facial scan, with the segmented regions for texture synthesis indicated.*

We begin with full facial scans recorded using the multiview polarized gradient illumination technique of [Ghosh et al. 2011], which produces an ear-to-ear polygon mesh of approximately 5 million polygons, 4K (4096 $\times$ 4096 pixel) diffuse and specular albedo maps, and a world-space normal map. We believe our technique could also work with other high-resolution facial capture techniques such as [Beeler et al. 2010] and [XYZRGB].

We create the texture coordinate space for the facial scan using the commercial product *Unfold3D* in a way which best preserves surface area and orientation with respect to the original scan. This allows us to assume that the relative scale and orientation of the patches is constant with respect to the texture space; if this were not the case, then an anisometric texture synthesis technique such as in [Lefebvre and Hoppe 2006] could be employed.

We transform the normal map to tangent space, and use multiresolution normal integration to construct the 4K displacement map which best agrees with the normal map. An artist then segments this map into five regions: forehead, temples, nose, cheeks, and chin (Fig. 6(a)).

To synthesize appropriate skin microstructure over the mesostructure present in our facial scans, we employ constrained texture synthesis in the framework of Image Analogies [Hertzmann et al. 2001], which synthesizes an image B' from an image B following the relationship of a pair of example images A and A'. In our case, B is the mesostructure-level displacement map for a region of the face, such as the forehead or cheek (see Fig. 7). Our goal is to synthesize B', a higher-resolution of this region exhibiting appropriate microstructure. We form the A to A' analogy by taking an exemplar displacement map of the microstructure of a related skin patch to be A'. The exemplar patch A' typically covers less than a square centimeter of skin surface, but at about ten times the resolution of

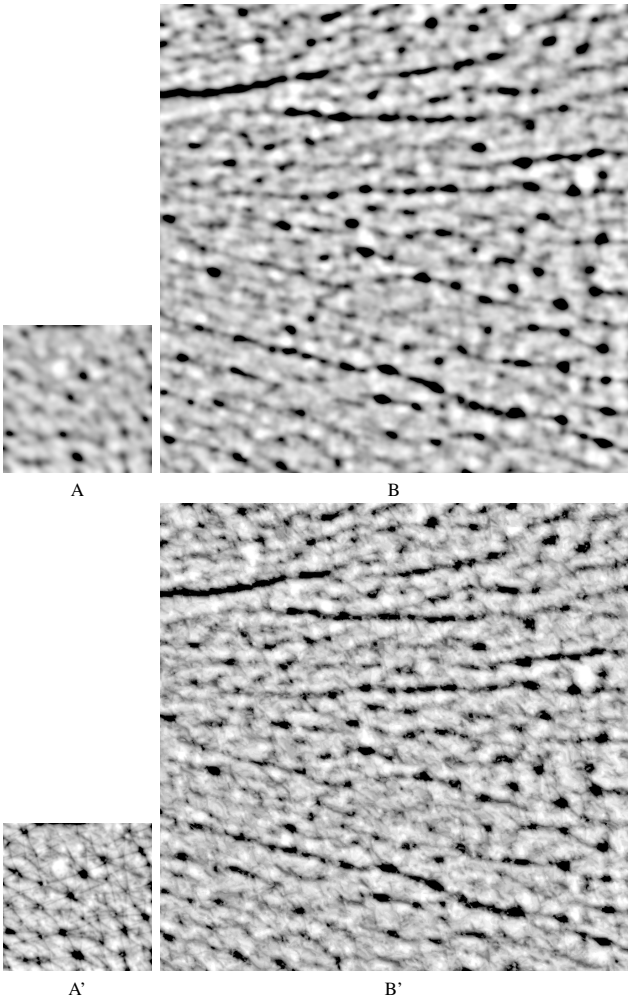

Figure 7: Microstructure Synthesis *We add microstructural detail to a scanned facial region B using the analogous relationship between an exemplar microstructure map A′ and a blurred version of it A which matches the mesostructural detail of B. The synthesized result with both meso- and microstructure is B′. These displacement maps are from the* temple *region of a female subject.*

the mesostructure detail. To form A, we blur the exemplar A' with a Gaussian filter to suppress the high frequencies not present in the input mesostructure map B. A thus represents the mesotructure of the exemplar. The synthesis process then proceeds in the image analogies framework to produce the output surface shape B', with both mesostructure and microstructure, given the input mesostructure B, and the pair of exemplars relating mesostructure A to its microstructure A'.

As values from A' replace values in B, it is possible for the low-to-medium frequency mesostructure of B to be unintentionally replaced with mesostructure from A'. To ensure preservation of mesostructure details in the scan data, we run a high-pass filter on both the input and the exemplar displacement maps. This technique separates the low-to-medium frequency features from the high-frequency features, allowing the medium frequency features present in the high pass of B to guide the synthesis of the high-frequency features from A'.

Our synthesis process employs a similar *BestApproximateMatch* based on an approximate-nearest-neighbor (ANN) search and *BestCoherenceMatch* based on Ashikhmin [Ashikhmin 2001] as discussed in [Hertzmann et al. 2001]. However, we introduce a weight-

ing parameter α to control the relative importance of a pixel neighborhood in A' matching B' compared to the importance $(1-\alpha)$ of the match between the corresponding pixel neighborhood in B and A. We found this to be a useful control parameter because certain visually important details that exist in A' can be entirely absent in the lower frequency mesostructures maps A and B. We found that an α value around 0.5 produced good results.

We carry out the texture synthesis in a multiresolution fashion, increasing the window size for the neighborhood matching from a 5×5 pixel window at the lowest level (4K resolution) to a 13×13 pixel window at the highest level (16K resolution) to match the increase in size of features at each level of the synthesis. We also apply principle component analysis (PCA) in order to speedup the synthesis. We use PCA to reduce the dimensionality of the search space to n, where $n \times n$ is the original pixel window. PCA reduced the synthesis time by a factor of three without any qualitative decreases in the quality of the result. We employ equal weighting to *BestApproximateMatch* and *BestCoherenceMatch* by setting the coherence parameter $\kappa = 1$ in the synthesis process. Since the specular albedo map is highly correlated to the surface meso- and microstructure, we also synthesize a 16K specular albedo map as a by-product of the microstructure displacement synthesis process by borrowing the corresponding pixels from the specular albedo exemplar.

5 Results

Creating Renderings Once we have synthesized microstructure-level displacement maps for a face, we can create renderings using the subsurface albedo, specular albedo, and single scattering coefficients using any number of standard rendering techniques. To generate the renderings in this section, we use a local specular reflection model with two lobes per skin region estimated as described in Section 3. For efficiency, the subsurface reflection is simulated using the hybrid normals rendering technique of [Ma et al. 2007] from the gradient illumination data of the full facial scan, though in practice a true scattering simulation would be preferable. Single scattering, estimated from the exemplars, is rendered with the process used in [Ghosh et al. 2008]. We upscaled the original scanned data to fill in the regions where we did not synthesize microgeometry (lips, neck, eye brows, etc). For the upper eyelids, we synthesised microstructure using the measured forehead microstructure exemplar.

Image Sizes and Resampling In the rendering process, the subsurface albedo and subsurface normal maps remain at the original 4K resolution of the facial scan, as does the polygon geometry. The synthesized 16K microstructure displacement map is converted to a normal map for rendering and used in conjunction with the 16K synthesized specular albedo map. To avoid artifacts from normal map resampling or aliasing, full-face renderings are hardware rendered using OpenGL to a large 16K (16384×16384 pixels) *half float* frame buffer, and then resized to 4K using radiometrically linear pixel blending, requiring approximately 1GB of GPU memory.

Fig. 1 shows a high-resolution point-light rendering of a female subject using a synthesized 16K microstructure displacement map (b) compared to using just a 4K mesostructure displacement map from the original scan (a) and a reference photograph under flash illumiation (c). The 16K rendering includes more high-frequency specular detail, and better exhibits skin's characteristic effect of isolated "glints" far away from the center of the specular highlight. A similar result is shown in Fig. 8 where a point-light rendering of a male subject's forehead using synthesized microstructure is a better match to a validation photograph compared to the rendering of the original scan with mesostructure detail.

Description	Subject 1	Subject 2
forehead	m1=0.150, m2=0.050, w=0.88	m1=0.150, m2=0.050, w=0.60
temple	m1=0.150, m2=0.075, w=0.55	m1=0.175, m2=0.050, w=0.80
cheek	m1=0.150, m2=0.125, w=0.60	m1=0.100, m2=0.075, w=0.50
nose	m1=0.100, m2=0.075, w=0.80	m1=0.100, m2=0.050, w=0.50
chin	m1=0.125, m2=0.100, w=0.90	m1=0.150, m2=0.050, w=0.75

Table 1: *Microscale two-lobe Beckmann distribution parameters obtained for the different skin patches across two subjects of Fig. 3.*

Description	Subject 1	Subject 2
forehead	m1=0.250, m2=0.125, w=0.85	m1=0.250, m2=0.125, w=0.80
temple	m1=0.225, m2=0.125, w=0.80	m1=0.225, m2=0.150, w=0.70
cheek	m1=0.275, m2=0.200, w=0.60	m1=0.225, m2=0.150, w=0.50
nose	m1=0.175, m2=0.100, w=0.65	m1=0.150, m2=0.075, w=0.80
chin	m1=0.250, m2=0.150, w=0.35	m1=0.300, m2=0.225, w=0.15

Table 2: *Mesoscale two-lobe Beckmann distribution parameters obtained for different facial regions across two subjects.*

Fig. 9 shows displacement maps (top row), normal maps (x and y components only, middle row), and point-light renderings (bottom row) of a male forehead region generated with different synthesis processes. Fig. 9(a) shows a region from an original mesostructure-only scan, with no synthesis to add microstructure detail. The specular reflection, rendered with corresponding mesoscale BRDF fit, is quite smooth as a result, and the skin reflection is not especially realistic. Fig. 9(b) shows the result of our microstructure synthesis process using an exemplar skin patch measurement from the same subject (the "forehead" of Subject 1 in Fig. 3(a)). The specular reflection, rendered with a microscale BRDF fit, is broken up and shows greater surface detail, while the mesostructure of the forehead crease is preserved. Fig. 9(c) shows the result of using a forehead patch from a different male subject as the exemplar for adding microstructure. Although the fine skin texture is different, the synthesized geometry and rendering and is still very plausible, suggesting cross-subject microstructure synthesis to be a viable option.

The last column of Fig. 9 tests the importance of the mesostructure constraints during texture synthesis. Fig. 9(d) was generated by setting the α parameter to 1.0, ignoring mesostructure matching constraints in the matching process, and then blindly embossing the synthesized detail onto the mesostructure of the original scan. Fig. 9(b), however, synthesizes detail in a way which tries to follow the mesostructure, so pores and creases in the scan will tend to draw upon similar areas in the microstructure exemplar for their detail. As a result, the constrained synthesis column (b) produces a more plausible result which better reinforces the scanned mesostructure than (d).

Table 1 shows specular BRDF lobe fits for different skin patches across two subjects measured using our skin patch measurement setups. Table 2 presents comparison Beckmann distribution fits for similar facial regions obtained at the mesostructure scale from a face scan. The parameter w is the weight of the convex combination of the two lobes $m1$ and $m2$. As can be seen, the BRDF lobes estimates at the microstructure scale exhibit reduced specular roughness compared to the mesostructure scale BRDF estimate as well as significantly less variation across skin patches. This agrees with the theory that at sufficiently high resolution, the surface microgeometry variation is responsible for the appearance of specular roughness.

Fig. 10 shows comparison renderings of a small patch of forehead shown in Fig. 9, at different scales of modeling. The original scanned data with mesostructure detail and mesoscale BRDF fit results in a broad specular reflection that misses the sharp "glints" (a). Rendering the scanned mesoscale surface detail with a microscale BRDF fit results in a qualitiative improvement in the result at this

(a) Rendering using original scan (b) Rendering using synthesized microstructure (c) Comparison photograph

Figure 8: *Renderings of original facial scan with mesostructure detail (a), and with synthesized microgeometry (b) compared to to a photograph under flash illumination (c). The digitized skin patches used for microstructure synthesis can be seen in Fig. 3, Subject 1.*

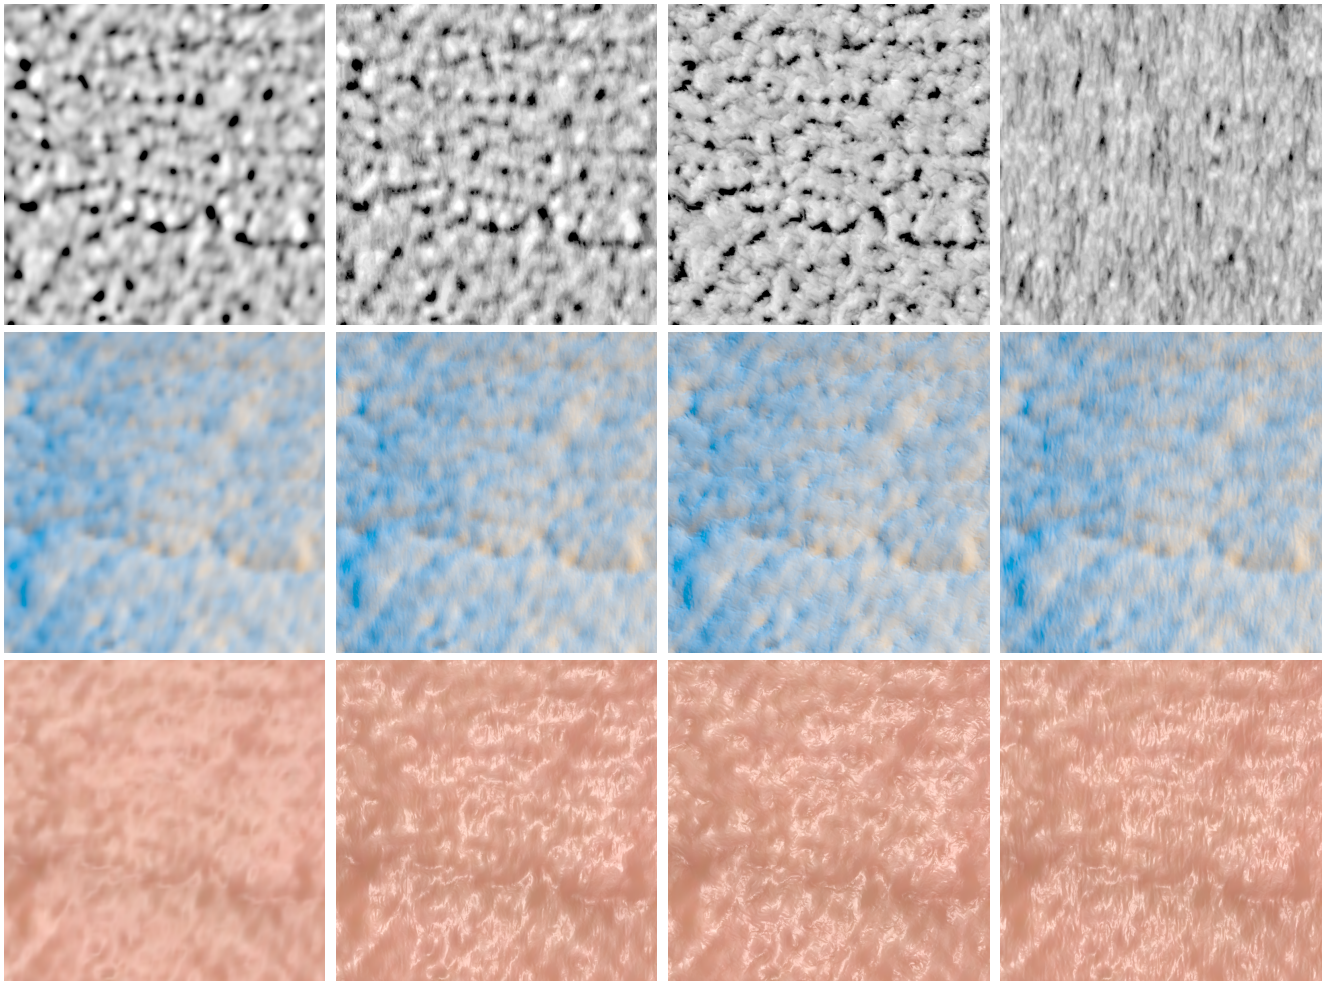

(a) Original mesostructure (b) Same-subject synthesis (c) Cross-subject synthesis (d) Unconstrained synthesis

Figure 9: *Microstructure synthesis with different exemplars and constraints. The top row shows displacement maps, the middle row shows normal maps and the bottom row shows point-light renderings using the maps. (a) Original mesostructure. (b) With microstructure synthesized from the same subject. (c) With microstructure synthesized from a different subject. (d) Without constraining the synthesis to match the underlying mesostructure.*

scale of visualization (b). However, the specular reflection still appears a bit too smooth to be realisitc. Rendering with synthesized microstructure in conjunction with the microscale BRDF fit achieves the best qualitiative rendering result (c).

Fig. 12 shows a full-facial rendering with a 16K microstructure displacement map created by merging synthesized skin regions from the skin patches seen for "Subject 2" in Fig. 3. The rendering was created at 16K and filtered down to 4K for presentation. Here, we simulate a "parallel polarized" point lighting condition in order to accentuate the observed specular highlights. Figure 11 provides additional such renderings of faces with synthesized microstructure at 16K resolution.

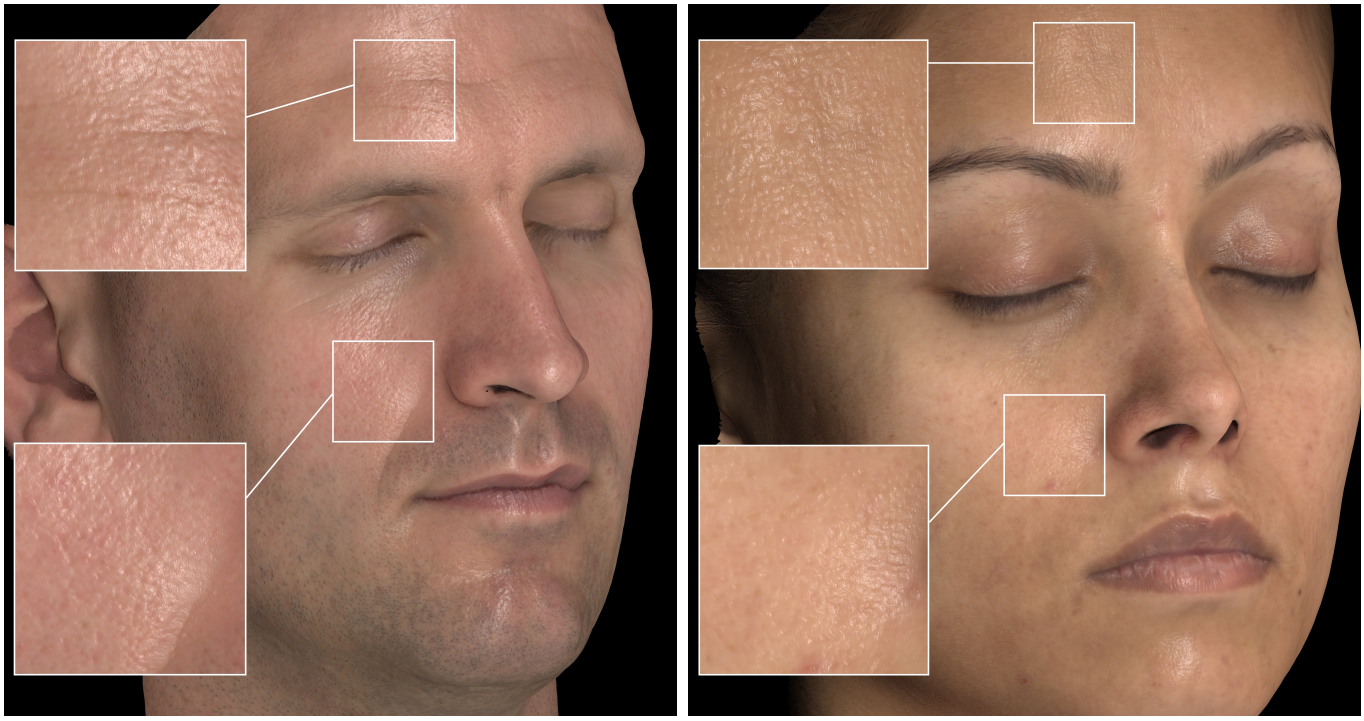

Figure 11: *Renderings from 16K displacement maps with synthesized microstructure.*

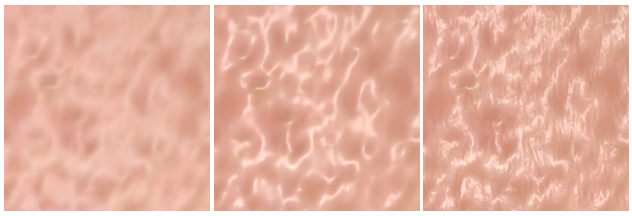

(a) Original scan (b) Microscale BRDF (c) Our result

Figure 10: *Rendering original scan data (geometry + BRDF fit) (a), compared to rendering scanned geometry with microscale BRDF fit (b), and rendering with synthesized microstructure + microscale BRDF fit (c).*

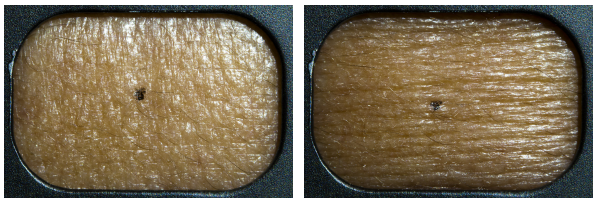

(a) (b)

Figure 13: *(a) 15mm × 10mm forehead patch from a neutral expression, with marked reference point. (b) The same forehead patch during a raised-eyebrows expression, exhibiting strongly anisotropic microstructure, with submillimeter furrows.*

6 Future Work

Our results suggest several avenues for future work. The efficiency of our texture synthesis and rendering algorithms could clearly be increased using GPU acceleration with more advanced texture synthesis algorithms and acceleration transformations such as appearance-space texture synthesis [Lefebvre and Hoppe 2006] and PatchMatch [Barnes et al. 2009]. We do not yet address real-time rendering, for which GPU-accelerated normal distribution functions [Tan et al. 2005; Han et al. 2007] will be required. The complexity of skin microstructure may benefit from techniques developed for rendering especially high-frequency normal distributions such as car paint [Rump et al. 2009].

We currently ignore the face's velvety fine hair and the *asperity scattering* [Koenderink and Pont 2003] it contributes to skin reflectance. We believe that side lighting conditions during skin measurement could allow the hairs to be isolated, and models of the hairs could be added onto the synthesized textures, further increasing realism. It would also be of interest to record the effects of cosmetics on skin reflectance at the microstructure scale, potentially enabling accurate simulations of cosmetic applications.

Since actors are often scanned in a variety of expressions, it would be of interest to synthesize consistent microstructure across all of the expressions. This may not be straightforward, as we can observe that microstructure changes dramatically as skin stretches and contracts (Fig. 13) in the course of expression formation. We believe that recording skin microstructure under calibrated amounts of stress and shear will provide useful data for microstructure dynamics simulation, increasing the realism of animated digital characters.

7 Conclusion

In this work, we have presented a practical technique for adding microstructure-level detail to high-resolution facial scans, as well as microscale skin BRDF analysis, allowing renderings to exhibit considerably more realistic patterns of specular reflection. We believe this can produce a significant improvement in the realism of computer-generated digital characters. Our initial experiments with cross-subject microstructure synthesis also suggests the applicability of this technique to a wide variety of facial scans using a small database of measured microgeometry exemplars.

Figure 12: *A facial rendering made using a 16K high-resolution displacement map synthesized from the "Subject 2" skin microgeometry exemplars in Figure 3.*

References

ACEVEDO, G., NEVSHUPOV, S., COWELY, J., AND NORRIS, K. 2010. An accurate method for acquiring high resolution skin displacement maps. In *ACM SIGGRAPH 2010 Talks*, ACM, New York, NY, USA, SIGGRAPH '10, 4:1–4:1.

ASHIKHMIN, M., PREMOZE, S., AND SHIRLEY, P. S. 2000. A microfacet-based BRDF generator. In *Proceedings of ACM SIGGRAPH 2000*, 65–74.

ASHIKHMIN, M. 2001. Synthesizing natural textures. In *Proceedings of the 2001 symposium on Interactive 3D graphics*, ACM, New York, NY, USA, I3D '01, 217–226.

BARNES, C., SHECHTMAN, E., FINKELSTEIN, A., AND GOLDMAN, D. B. 2009. Patchmatch: a randomized correspondence algorithm for structural image editing. *ACM Trans. Graph. 28* (July), 24:1–24:11.

BEELER, T., BICKEL, B., BEARDSLEY, P., SUMNER, B., AND GROSS, M. 2010. High-quality single-shot capture of facial geometry. *ACM Trans. Graph. 29* (July), 40:1–40:9.

CHEN, T., GOESELE, M., AND SEIDEL, H. P. 2006. Mesostructure from specularities. In *CVPR*, 1825–1832.

COOK, R. L., AND TORRANCE, K. E. 1982. A reflectance model for computer graphics. *ACM TOG 1*, 1, 7–24.

CULA, O. G., DANA, K. J., MURPHY, F. P., AND RAO, B. K. 2005. Skin texture modeling. *International Journal of Computer Vision 62*, 97–119. 10.1023/B:VISI.0000046591.79973.6f.

DEBEVEC, P., HAWKINS, T., TCHOU, C., DUIKER, H.-P., SAROKIN, W., AND SAGAR, M. 2000. Acquiring the reflectance field of a human face. In *ACM SIGGRAPH*.

D'EON, E., AND IRVING, G. 2011. A quantized-diffusion model for rendering translucent materials. In *ACM SIGGRAPH 2011 papers*, ACM, New York, NY, USA, SIGGRAPH '11, 56:1–56:14.

DONG, Y., WANG, J., TONG, X., SNYDER, J., LAN, Y., BEN-EZRA, M., AND GUO, B. 2010. Manifold bootstrapping for svbrdf capture. *ACM Transactions on Graphics 29*, 4 (July), 98:1–98:10.

DONNER, C., AND JENSEN, H. W. 2005. Light diffusion in multilayered translucent materials. *ACM TOG 24*, 3, 1032–1039.

EFROS, A. A., AND FREEMAN, W. T. 2001. Image quilting for texture synthesis and transfer. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, ACM, New York, NY, USA, SIGGRAPH '01, 341–346.

EFROS, A. A., AND LEUNG, T. K. 1999. Texture synthesis by non-parametric sampling. In *Proceedings of the International Conference on Computer Vision-Volume 2 - Volume 2*, IEEE Computer Society, Washington, DC, USA, ICCV '99, 1033–.

FREEMAN, W., JONES, T., AND PASZTOR, E. 2002. Example-based super-resolution. *Computer Graphics and Applications, IEEE 22*, 2 (mar/apr), 56 –65.

GHOSH, A., HAWKINS, T., PEERS, P., FREDERIKSEN, S., AND DEBEVEC, P. 2008. Practical modeling and acquisition of layered facial reflectance. *ACM Transactions on Graphics 27*, 5 (Dec.), 139:1–139:10.

GHOSH, A., FYFFE, G., TUNWATTANAPONG, B., BUSCH, J., YU, X., AND DEBEVEC, P. 2011. Multiview face capture using polarized spherical gradient illumination. In *Proceedings of the 2011 SIGGRAPH Asia Conference*, ACM, New York, NY, USA, SA '11, 129:1–129:10.

GOLOVINSKIY, A., MATUSIK, W., PFISTER, H., RUSINKIEWICZ, S., AND FUNKHOUSER, T. 2006. A statistical model for synthesis of detailed facial geometry. *ACM Transactions on Graphics 25*, 3 (July), 1025–1034.

HAN, C., SUN, B., RAMAMOORTHI, R., AND GRINSPUN, E. 2007. Frequency domain normal map filtering. *ACM Trans. Graph. 26* (July).

HANRAHAN, P., AND KRUEGER, W. 1993. Reflection from layered surfaces due to subsurface scattering. In *Proceedings of SIGGRAPH 93*, 165–174.

HEEGER, D. J., AND BERGEN, J. R. 1995. Pyramid-based texture analysis/synthesis. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, ACM, New York, NY, USA, SIGGRAPH '95, 229–238.

HERTZMANN, A., JACOBS, C. E., OLIVER, N., CURLESS, B., AND SALESIN, D. H. 2001. Image analogies. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, ACM, New York, NY, USA, SIGGRAPH '01, 327–340.

JENSEN, H. W., MARSCHNER, S. R., LEVOY, M., AND HANRAHAN, P. 2001. A practical model for subsurface light transport. In *Proceedings of ACM SIGGRAPH 2001*, 511–518.

JOHNSON, M. K., COLE, F., RAJ, A., AND ADELSON, E. H. 2011. Microgeometry capture using an elastomeric sensor. *ACM Trans. Graph. 30* (Aug.), 46:1–46:8.

KOENDERINK, J., AND PONT, S. 2003. The secret of velvety skin. *Mach. Vision Appl. 14* (September), 260–268.

LEFEBVRE, S., AND HOPPE, H. 2005. Parallel controllable texture synthesis. In *ACM SIGGRAPH 2005 Papers*, ACM, New York, NY, USA, SIGGRAPH '05, 777–786.

LEFEBVRE, S., AND HOPPE, H. 2006. Appearance-space texture synthesis. *ACM Trans. Graph. 25* (July), 541–548.

LENSCH, H. P. A., KAUTZ, J., GOESELE, M., HEIDRICH, W., AND SEIDEL, H.-P. 2003. Image-based reconstruction of spatial appearance and geometric detail. *ACM TOG 22*, 2, 234–257.

LIU, C., SHUM, H.-Y., AND ZHANG, C.-S. 2001. A two-step approach to hallucinating faces: global parametric model and local nonparametric model. In *Computer Vision and Pattern Recognition, 2001. CVPR 2001. Proceedings of the 2001 IEEE Computer Society Conference on*, vol. 1, I–192 – I–198 vol.1.

LIU, Z., ZHANG, Z., AND SHAN, Y. 2004. Image-based surface detail transfer. *Computer Graphics and Applications, IEEE 24*, 3 (may-june), 30 –35.

MA, W.-C., HAWKINS, T., PEERS, P., CHABERT, C.-F., WEISS, M., AND DEBEVEC, P. 2007. Rapid acquisition of specular and diffuse normal maps from polarized spherical gradient illumination. In *Rendering Techniques*, 183–194.

MALZBENDER, T., GELB, D., AND WOLTERS, H. 2001. Polynomial texture maps. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, ACM, New York, NY, USA, SIGGRAPH '01, 519–528.

MARSCHNER, S. R., WESTIN, S. H., LAFORTUNE, E. P. F., TORRANCE, K. E., AND GREENBERG, D. P. 1999. Image-based BRDF measurement including human skin. In *Rendering Techniques*.

MARSCHNER, S. R., WESTIN, S. H., ARBREE, A., AND MOON, J. T. 2005. Measuring and modeling the appearance of finished wood. *ACM Trans. Graph. 24* (July), 727–734.

RAMANARAYANAN, G., AND BALA, K. 2007. Constrained texture synthesis via energy minimization. *IEEE Transactions on Visualization and Computer Graphics 13* (January), 167–178.

RAMELLA-ROMAN, J. C. 2008. Out of plane polarimetric imaging of skin: Surface and subsurface effect. In *Optical Waveguide Sensing and Imaging*, W. J. Bock, I. Gannot, and S. Tanev, Eds., NATO Science for Peace and Security Series B: Physics and Biophysics. Springer Netherlands, 259–269. "10.1007/978-1-4020-6952-9_12".

RUMP, M., SARLETTE, R., AND KLEIN, R. 2009. Efficient resampling, compression and rendering of metallic and pearlescent paint. In *Vision, Modeling, and Visualization*, M. Magnor, B. Rosenhahn, and H. Theisel, Eds., 11–18.

RUSHMEIER, H., TAUBIN, G., AND GUÉZIEC, A. 1997. Applying shape from lighting variation to bump map capture. In *Rendering Techniques*, 35–44.

TAN, P., LIN, S., QUAN, L., GUO, B., AND SHUM, H.-Y. 2005. Multiresolution reflectance filtering. In *Rendering Techniques 2005: 16th Eurographics Workshop on Rendering*, 111–116.

TONG, X., ZHANG, J., LIU, L., WANG, X., GUO, B., AND SHUM, H.-Y. 2002. Synthesis of bidirectional texture functions on arbitrary surfaces. *ACM Trans. Graph. 21* (July), 665–672.

TORRANCE, K. E., AND SPARROW, E. M. 1967. Theory of off-specular reflection from roughened surfaces. *J. Opt. Soc. Am. 57*, 1104–1114.

TSUMURA, N., OJIMA, N., SATO, K., SHIRAISHI, M., SHIMIZU, H., NABESHIMA, H., AKAZAKI, S., HORI, K., AND MIYAKE, Y. 2003. Image-based skin color and texture analysis/synthesis by extracting hemoglobin and melanin information in the skin. *ACM TOG 22*, 3, 770–779.

WANG, L., AND MUELLER, K. 2004. Generating sub-resolution detail in images and volumes using constrained texture synthesis. In *Proceedings of the conference on Visualization '04*, IEEE Computer Society, Washington, DC, USA, VIS '04, 75–82.

WEI, L.-Y., AND LEVOY, M. 2001. Texture synthesis over arbitrary manifold surfaces. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, ACM, New York, NY, USA, SIGGRAPH '01, 355–360.

WEI, L.-Y., LEFEBVRE, S., KWATRA, V., AND TURK, G. 2009. State of the art in example-based texture synthesis. In *Eurographics 2009, State of the Art Report, EG-STAR*, Eurographics Association.

WEYRICH, T., MATUSIK, W., PFISTER, H., BICKEL, B., DONNER, C., TU, C., MCANDLESS, J., LEE, J., NGAN, A., JENSEN, H. W., AND GROSS, M. 2006. Analysis of human faces using a measurement-based skin reflectance model. *ACM Transactions on Graphics 25*, 3 (July), 1013–1024.

WILSON, C. A., GHOSH, A., PEERS, P., CHIANG, J.-Y., BUSCH, J., AND DEBEVEC, P. 2010. Temporal upsampling of performance geometry using photometric alignment. *ACM Trans. Graph. 29* (April), 17:1–17:11.

WOODHAM, R. J. 1978. Photometric stereo: A reflectance map technique for determining surface orientation from image intensity. In *Proc. SPIE's 22nd Annual Technical Symposium*, vol. 155.

XYZRGB. 3D laser scanning - XYZ RGB Inc. http://www.xyzrgb.com/.

YING, L., HERTZMANN, A., BIERMANN, H., AND ZORIN, D. 2001. Texture and shape synthesis on surfaces. In *Proceedings of the 12th Eurographics Workshop on Rendering Techniques*, Springer-Verlag, London, UK, 301–312.

ZHAO, S., JAKOB, W., MARSCHNER, S., AND BALA, K. 2011. Building volumetric appearance models of fabric using micro ct imaging. *ACM Trans. Graph. 30* (Aug.), 44:1–44:10.

ZICKLER, T., RAMAMOORTHI, R., ENRIQUE, S., AND BELHUMEUR, P. N. 2006. Reflectance sharing: Predicting appearance from a sparse set of images of a known shape. *PAMI 28*, 8, 1287–1302.