Measurement-Based Synthesis of Facial Microgeometry

Paul Graham, Borom Tunwattanapong, Jay Busch, Xueming Yu, Andrew Jones, Paul Debevec, and Abhijeet Ghosh [†]

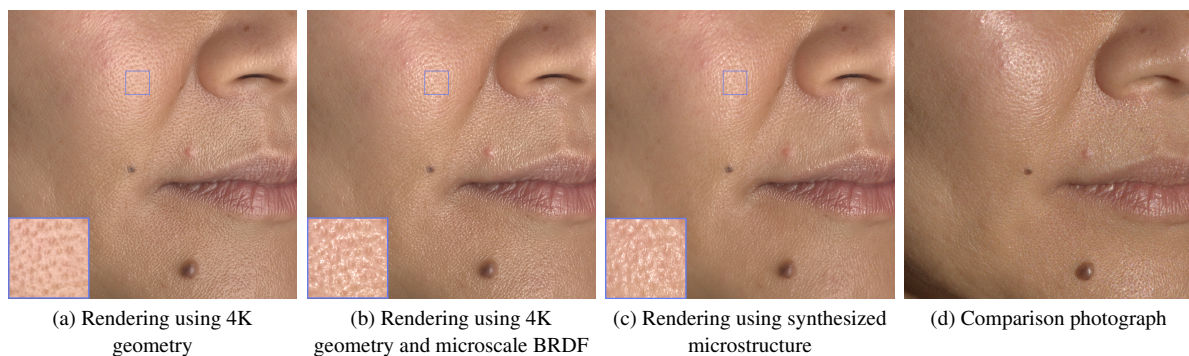USC Institute for Creative Technologies, Los Angeles, CA, USA

| (a) Rendering using 4K geometry | (b) Rendering using 4K geometry and microscale BRDF | (c) Rendering using synthesized microstructure | (d) Comparison photograph |

Figure 1: *(a) Rendering with scanned mesostructure (4K displacement map). (b) Rendering with scanned mesostructure (4K displacement map) and microstructure measured BRDF. (c) Rendering with synthesized microstructure (16K displacement map). (d) Photograph under flash illumination. See Fig. 3, Subject 2 for the skin patches used in microstructure synthesis.*

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. Additionally, we make point-source reflectance measurements of the skin patches to characterize the specular reflectance lobes at this smaller scale and analyze facial reflectance variation at both the mesostructure and microstructure scales. 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 reflection effects.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [COMPUTER GRAPHICS]: Three-Dimensional Graphics and Realism

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 [HK93, JMLH01, DJ05, DI11], somewhat less attention has been given to surface reflection.

[†] Currently at Imperial College London

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 [WMP*06, MHP*07, BBB*10] as well as high-resolution scanning of facial casts (e.g. [XYZ], [ANCN10]) 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* [TS67] — 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 synthesis approach for increasing 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 [MHP*07] at approximately 10 micron precision. Additionally, we make point-source reflectance measurements to characterize the specular reflectance lobes at this smaller scale and analyze facial reflectance variation at both the mesostructure and microstructure scales. We then employ a constrained texture synthesis algorithm based on Image Analogies [HJO*01] to synthesize appropriate surface microstructure per-region, blending the regions to cover the whole entire face. We show that renderings made with microstructure-level geometry and reflectance 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.

2.1. Skin Reflectance Modeling

Physically-based models of light reflection from rough surfaces are typically based on a microfacet distribution model [TS67, CT82]. BRDFs for rendering have been tabulated from photographed foreheads under a sampling of lighting conditions and viewing directions [MWL*99]. Using a moving light source and camera to record a bidirectional database of skin samples, a texton histogram model has been built to detect skin abnormalities [CDMR05]. Ghosh et al. [GHP*08] fit a microfacet BRDF model [APS00] to specular reflectance data across various patches of human faces from high-resolution surface orientation measurements.

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 sampling subsurface scattering parameters with a special probe [WMP*06], and using a BRDF probe to add detailed BRDFs to entire objects observed in relatively few lighting conditions [DWT*10]. Our work also relates to basis BRDF modeling [LKG*03] and reflectance sharing [ZREB06].

2.2. Surface Detail Measurement

Key to our work is measuring surface detail from images taken with a fixed viewpoint and varying lighting. Classic photometric stereo [Woo78] derives surface orientations of Lambertian surfaces from three point-light directions. Bump maps produced from photometric stereo have been used to increase the detail on rendered surfaces [RTG97]. For semi-translucent materials such as skin, subsurface scattering blurs the surface detail recoverable from traditional photometric stereo significantly [RR08]. *Specular* surface reflection analysis has been used to obtain more precise surface orientation measurements of translucent materials [DHT*00, RR08, WMP*06, CGS06]. Ma et al. [MHP*07] 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 [JCRA11] pushes silver-coated gel against a sample and uses photometric stereo to record surface microgeometry at the level of a few microns. While this setup achieves the resolution required for our technique, it does not enable reflectance measurements. In our work, we measure microgeometry using a polarized gradient illumination technique [GFT*11] since it requires no contact with the sample and permits the material's actual reflectance to be observed during measurement for BRDF fitting.

2.3. Texture Synthesis

Wei et al. [WLKT09] compiled a recent review of example-based texture synthesis. To summarize a few results, 2D texture synthesis algorithms [HB95, EL99] have been extended to arbitrary manifold surfaces [WL01, YHBZ01, LH06] and can also synthesize displacement maps [YHBZ01], measured reflectance properties [TZL*02] and skin color [TOS*03]. Some techniques [EF01, HJO*01] permit *texture transfer*, transforming a complete image so that it has textural detail of a given sample. *Super-resolution* techniques add detail to a low-resolution image based on examples [HJO*01, FJP02, LH05]. Especially relevant to our work are *constrained texture synthesis* techniques [WM04, RB07] which add plausible detail to a low-resolution image without changing its low frequency content. In our work, we adapt the Image Analogies [HJO*01] framework to perform con-

strained 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. Face hallucination [LSZ01] creates recognizable facial images by performing constrained facial texture synthesis onto a parametric face model. Image processing applies facial detail from one subject to another, allowing image-based aging effects [LZS04]. Most closely related to our work is Golovinskiy et al. [GMP*06], which synthesizes skin mesostructure (but not microstructure) from higher-quality facial scans onto otherwise smooth facial shapes. In addition to working at a very different scale than our work (we assume that wrinkles, creases, and larger pores are already apparent in the source scan), 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. Observations of a cross-section of wood from an electron micrograph have motivated reflection models which better predicted anisotropic reflection effects [MWAM05]. Direct use of Micro CT imaging of fabric samples have been used to model the volumetric scattering of complete pieces of cloth [ZJMB11].

3. Recording Skin Microstructure

3.1. Acquisition

We record the microstructure of skin patches using one 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 $\times$ 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 square microns ($7\mu\text{m}^2$) 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 [MGW01], with the addition that each light can produce either of two linear polarization conditions. We modified the polarization orientations of a multiview polarization pattern [GFT*11] such that each light is specifically optimized for a single 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 joint photometric alignment [WGP*10]. For BRDF fitting, we additionally capture a single-light image in both polarized lighting conditions.


Figure 2: *Acquisition setups for skin microgeometry (a) 12-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.*

For especially smooth or oily skin patches, the 12 light positions can produce separated specular highlights, which can bias surface normal measurement. To address this, 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. 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 reinforced to remove mechanically vibrations and flexing which would blur the captured imagery.

3.2. 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. 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.

Using the polarization difference point-lit image, we also tabulate a specular lobe shape and single scattering model parameters [GHP*08]. 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 a specular BRDF
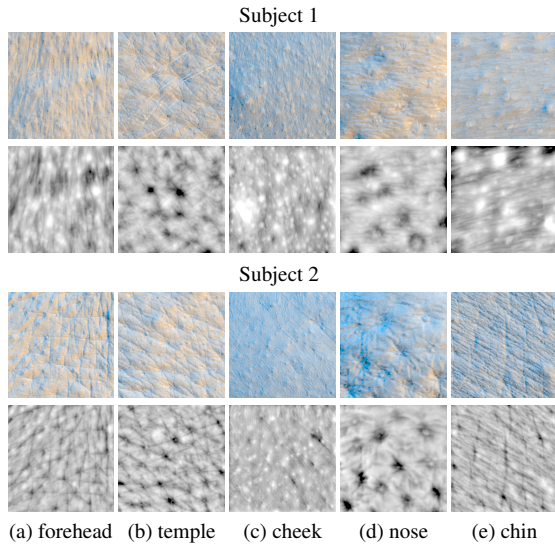
Subject 1

Subject 2

(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.*

model [TS67], we found that *two* lobes of a Beckmann distribution [CT82] fit the data well. In order to better fit 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 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 (see Table 1) compared to the differences observed at the mesostructure scale (see Table 2).

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 then merge the synthesized facial regions into a full map of the microstructure.

We begin with full facial scans recorded using a multiview

(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.*


(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.*

polarized gradient illumination technique [GFT*11], which produces an ear-to-ear polygon mesh of approximately five million polygons, 4K (4096×4096 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 [BBB*10, XYZ].

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

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

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 [LH06] 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 segments this map into forehead, temples, nose, cheeks, and chin regions (Fig. 6). Regions overlap enough so they can be blended together using linear-interpolation.

To synthesize appropriate skin microstructure over the mesostructure present in our facial scans, we employ constrained texture synthesis in the framework of Image Analogies [HJO*01], 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 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 exem-


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 mesostructure and microstructure is B′. These displacement maps are from the* temple *region of a female subject.*

plar 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'. We then recombine the frequencies to obtain a final result.

Our synthesis process employs a similar *BestApproximateMatch* based on an approximate-nearest-neighbor (ANN) search and *BestCoherenceMatch* based on Ashikhmin [Ash01]. However, we introduce a weighting 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 speed up 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 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 mesostructure 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

5.1. 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 standard rendering technique. 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 a hybrid normals rendering technique [MHP*07] 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 also rendered [GHP*08]. 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.

5.2. 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 created using an OpenGL GPU shader 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 (c) compared to using just the 4K mesostructure displacement map from the original scan (a) and a 4k rendering using the microscale measured BRDF (b) as well as a reference photograph under flash illumiation (d). 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.

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 is still very plausible, suggesting cross-subject microstructure synthesis to be a viable option.

Fig. 9(d) tests the importance of the mesostructure constraints during texture synthesis. This column 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 (Fig. 9(b)), produces a more plausible result which better reinforces the scanned mesostructure than the texture synthesis column (Fig. 9(d)).

Table 1 shows specular BRDF lobe fits for different skin patches across two subjects measured using our microstructure skin patch measurement setups. Table 2 presents comparison Beckmann distribution fits for similar facial regions obtained at the mesostructure scale from a face scan. Table 3 shows the results of a cross-verification done by low-pass filtering the skin patches and measuring the BRDF fits at a "mesostructure scale." The parameter w is the weight of

| (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). See Fig. 3, Subject 1 for the skin patches used in microstructure synthesis.*
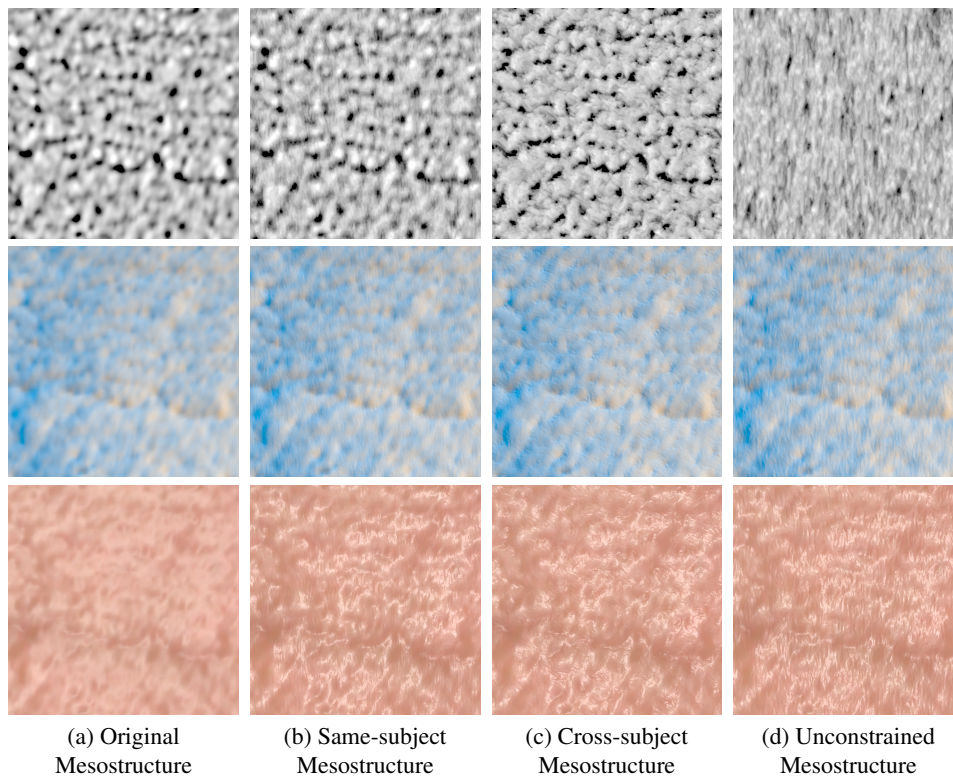

| (a) Original Mesostructure | (b) Same-subject Mesostructure | (c) Cross-subject Mesostructure | (d) Unconstrained Mesostructure |

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.*

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. Table 3 confirms that low-pass filtering the

microstructure also results in BRDF fits with wider roughness, similar to the mesostructure scale BRDF fits.

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 scale of vi-

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 Beckman distribution parameters obtained for different facial regions across two subjects*

Description	Subject 1	Subject 2
forehead	m1=0.175, m2=0.150, w=0.60	m1=0.225, m2=0.100, w=0.60
temple	m1=0.200, m2=0.125, w=0.95	m1=0.225, m2=0.200, w=0.30
cheek	m1=0.225, m2=0.200, w=0.60	m1=0.275, m2=0.125, w=0.45
nose	m1=0.175, m2=0.125, w=0.95	m1=0.125, m2=0.050, w=0.85
chin	m1=0.225, m2=0.200, w=0.75	m1=0.325, m2=0.150, w=0.70

Table 3: *Cross validation of microscale two-lobe distibution done at mesoscale resolution*

asperity scattering [KP03] 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.

(a) (b)

Figure 12: *(a) 15mm × 10mm forehead patch from a neutral expression, with marked reference point. (b) The same forehead patch with a raised-eyebrows expression, exhibiting anisotropic microstructure, with submillimeter furrows.*

Since actors are scanned in a variety of expressions, it would be of interest to synthesize consistent microstructure across expressions. This may not be straightforward, as we can observe that microstructure changes dramatically as skin stretches and contracts (Fig. 12) in the course of expression formation. We believe that recording skin microstructure under calibrated amounts of stress and shear will be useful 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 produces 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.

8. Acknowledgments

We would like to thank Rupa Upadhyay and Connie Siu for sitting as subjects; Kathleen Haase, Bill Swartout, Randy

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

Figure 11 provides renderings of faces with synthesized microstructure at 16K resolution. The renderings were 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.


(a) Mesoscale scan   (b) Mesoscale scan with Microscale BRDF   (c) Microscale scan BRDF

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).*

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 [LH06] and PatchMatch [BSFG09]. We do not yet address real-time rendering, for which GPU-accelerated normal distribution functions [TLQ*05, HSRG07] will be required. The complexity of skin microstructure may benefit from techniques developed for rendering especially high-frequency normal distributions such as car paint [RSK09].

We currently ignore the face's velvety fine hair and the

Figure 11: *Renderings from 16K displacement maps with synthesized microstructure. See supplemental material for high resolution images*

References

[ANCN10] ACEVEDO G., NEVSHUPOV S., COWELY J., NORRIS K.: An accurate method for acquiring high resolution skin displacement maps. In *ACM SIGGRAPH 2010 Talks* (New York, NY, USA, 2010), SIGGRAPH '10, ACM, pp. 4:1–4:1. 2

[APS00] ASHIKHMIN M., PREMOZE S., SHIRLEY P. S.: A microfacet-based BRDF generator. In *Proceedings of ACM SIGGRAPH 2000* (2000), pp. 65–74. 2

[Ash01] ASHIKHMIN M.: Synthesizing natural textures. In *Proceedings of the 2001 symposium on Interactive 3D graphics* (New York, NY, USA, 2001), I3D '01, ACM, pp. 217–226. 5

[BBB*10] BEELER T., BICKEL B., BEARDSLEY P., SUMNER B., GROSS M.: High-quality single-shot capture of facial geometry. *ACM Trans. Graph. 29* (July 2010), 40:1–40:9. 2, 4

[BSFG09] BARNES C., SHECHTMAN E., FINKELSTEIN A., GOLDMAN D. B.: Patchmatch: a randomized correspondence algorithm for structural image editing. *ACM Trans. Graph. 28* (July 2009), 24:1–24:11. 8

[CDMR05] CULA O. G., DANA K. J., MURPHY F. P., RAO B. K.: Skin texture modeling. *International Journal of Computer Vision 62* (2005), 97–119. 10.1023/B:VISI.0000046591.79973.6f. 2

[CGS06] CHEN T., GOESELE M., SEIDEL H. P.: Mesostructure from specularities. In *CVPR* (2006), pp. 1825–1832. 2

[CT82] COOK R. L., TORRANCE K. E.: A reflectance model for computer graphics. *ACM TOG 1*, 1 (1982), 7–24. 2, 4

[DHT*00] DEBEVEC P., HAWKINS T., TCHOU C., DUIKER H.-P., SAROKIN W., SAGAR M.: Acquiring the reflectance field of a human face. In *ACM SIGGRAPH* (2000). 2

[DI11] D'EON E., IRVING G.: A quantized-diffusion model for rendering translucent materials. In *ACM SIGGRAPH 2011 papers* (New York, NY, USA, 2011), SIGGRAPH '11, ACM, pp. 56:1–56:14. 1

[DJ05] DONNER C., JENSEN H. W.: Light diffusion in multilayered translucent materials. *ACM TOG 24*, 3 (2005), 1032–1039. 1

[DWT*10] DONG Y., WANG J., TONG X., SNYDER J., LAN Y., BEN-EZRA M., GUO B.: Manifold bootstrapping for svbrdf capture. *ACM Transactions on Graphics 29*, 4 (July 2010), 98:1–98:10. 2

[EF01] EFROS A. A., FREEMAN W. T.: Image quilting for texture synthesis and transfer. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 2001), SIGGRAPH '01, ACM, pp. 341–346. 2

[EL99] EFROS A. A., LEUNG T. K.: Texture synthesis by non-parametric sampling. In *Proceedings of the International Conference on Computer Vision-Volume 2 - Volume 2* (Washington, DC, USA, 1999), ICCV '99, IEEE Computer Society, pp. 1033–. 2

[FJP02] FREEMAN W., JONES T., PASZTOR E.: Example-based super-resolution. *Computer Graphics and Applications, IEEE 22*, 2 (mar/apr 2002), 56 –65. 2

[GFT*11] GHOSH A., FYFFE G., TUNWATTANAPONG B., BUSCH J., YU X., DEBEVEC P.: Multiview face capture using polarized spherical gradient illumination. In *Proceedings of the 2011 SIGGRAPH Asia Conference* (New York, NY, USA, 2011), SA '11, ACM, pp. 129:1–129:10. 2, 3, 4

[GHP*08] GHOSH A., HAWKINS T., PEERS P., FREDERIKSEN S., DEBEVEC P.: Practical modeling and acquisition of layered facial reflectance. *ACM Transactions on Graphics 27*, 5 (Dec. 2008), 139:1–139:10. 2, 3, 6

[GMP*06] GOLOVINSKIY A., MATUSIK W., PFISTER H., RUSINKIEWICZ S., FUNKHOUSER T.: A statistical model for synthesis of detailed facial geometry. *ACM Transactions on Graphics 25*, 3 (July 2006), 1025–1034. 3

[HB95] HEEGER D. J., BERGEN J. R.: Pyramid-based texture analysis/synthesis. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 1995), SIGGRAPH '95, ACM, pp. 229–238. 2

[HJO*01] HERTZMANN A., JACOBS C. E., OLIVER N., CURLESS B., SALESIN D. H.: Image analogies. In *Proceedings of*

the 28th annual conference on Computer graphics and interactive techniques (New York, NY, USA, 2001), SIGGRAPH '01, ACM, pp. 327–340. 2, 5

[HK93] HANRAHAN P., KRUEGER W.: Reflection from layered surfaces due to subsurface scattering. In *Proceedings of SIGGRAPH 93* (1993), pp. 165–174. 1

[HSRG07] HAN C., SUN B., RAMAMOORTHI R., GRINSPUN E.: Frequency domain normal map filtering. *ACM Trans. Graph. 26* (July 2007). 8

[JCRA11] JOHNSON M. K., COLE F., RAJ A., ADELSON E. H.: Microgeometry capture using an elastomeric sensor. *ACM Trans. Graph. 30* (Aug. 2011), 46:1–46:8. 2

[JMLH01] JENSEN H. W., MARSCHNER S. R., LEVOY M., HANRAHAN P.: A practical model for subsurface light transport. In *Proceedings of ACM SIGGRAPH 2001* (2001), pp. 511–518. 1

[KP03] KOENDERINK J., PONT S.: The secret of velvety skin. *Mach. Vision Appl. 14* (September 2003), 260–268. 8

[LH05] LEFEBVRE S., HOPPE H.: Parallel controllable texture synthesis. In *ACM SIGGRAPH 2005 Papers* (New York, NY, USA, 2005), SIGGRAPH '05, ACM, pp. 777–786. 2

[LH06] LEFEBVRE S., HOPPE H.: Appearance-space texture synthesis. *ACM Trans. Graph. 25* (July 2006), 541–548. 2, 5, 8

[LKG*03] LENSCH H. P. A., KAUTZ J., GOESELE M., HEIDRICH W., SEIDEL H.-P.: Image-based reconstruction of spatial appearance and geometric detail. *ACM TOG 22*, 2 (2003), 234–257. 2

[LSZ01] LIU C., SHUM H.-Y., ZHANG C.-S.: 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* (2001), vol. 1, pp. I–192 – I–198 vol.1. 3

[LZS04] LIU Z., ZHANG Z., SHAN Y.: Image-based surface detail transfer. *Computer Graphics and Applications, IEEE 24*, 3 (may-june 2004), 30 –35. 3

[MGW01] MALZBENDER T., GELB D., WOLTERS H.: Polynomial texture maps. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 2001), SIGGRAPH '01, ACM, pp. 519–528. 3

[MHP*07] MA W.-C., HAWKINS T., PEERS P., CHABERT C.-F., WEISS M., DEBEVEC P.: Rapid acquisition of specular and diffuse normal maps from polarized spherical gradient illumination. In *Rendering Techniques* (2007), pp. 183–194. 2, 6

[MWAM05] MARSCHNER S. R., WESTIN S. H., ARBREE A., MOON J. T.: Measuring and modeling the appearance of finished wood. *ACM Trans. Graph. 24* (July 2005), 727–734. 3

[MWL*99] MARSCHNER S. R., WESTIN S. H., LAFORTUNE E. P. F., TORRANCE K. E., GREENBERG D. P.: Image-based BRDF measurement including human skin. In *Rendering Techniques* (1999). 2

[RB07] RAMANARAYANAN G., BALA K.: Constrained texture synthesis via energy minimization. *IEEE Transactions on Visualization and Computer Graphics 13* (January 2007), 167–178. 2

[RR08] RAMELLA-ROMAN J. C.: Out of plane polarimetric imaging of skin: Surface and subsurface effect. In *Optical Waveguide Sensing and Imaging*, Bock W. J., Gannot I., Tanev S., (Eds.), NATO Science for Peace and Security Series B: Physics and Biophysics. Springer Netherlands, 2008, pp. 259–269. "10.1007/978-1-4020-6952-9_12". 2

[RSK09] RUMP M., SARLETTE R., KLEIN R.: Efficient resampling, compression and rendering of metallic and pearlescent paint. In *Vision, Modeling, and Visualization* (Nov. 2009), Magnor M., Rosenhahn B., Theisel H., (Eds.), pp. 11–18. 8

[RTG97] RUSHMEIER H., TAUBIN G., GUÉZIEC A.: Applying shape from lighting variation to bump map capture. In *Rendering Techniques* (1997), pp. 35–44. 2

[TLQ*05] TAN P., LIN S., QUAN L., GUO B., SHUM H.-Y.: Multiresolution reflectance filtering. In *Rendering Techniques 2005: 16th Eurographics Workshop on Rendering* (June 2005), pp. 111–116. 8

[TOS*03] TSUMURA N., OJIMA N., SATO K., SHIRAISHI M., SHIMIZU H., NABESHIMA H., AKAZAKI S., HORI K., MIYAKE Y.: Image-based skin color and texture analysis/synthesis by extracting hemoglobin and melanin information in the skin. *ACM TOG 22*, 3 (2003), 770–779. 2

[TS67] TORRANCE K. E., SPARROW E. M.: Theory of off-specular reflection from roughened surfaces. *J. Opt. Soc. Am. 57* (1967), 1104–1114. 2, 4

[TZL*02] TONG X., ZHANG J., LIU L., WANG X., GUO B., SHUM H.-Y.: Synthesis of bidirectional texture functions on arbitrary surfaces. *ACM Trans. Graph. 21* (July 2002), 665–672. 2

[WGP*10] WILSON C. A., GHOSH A., PEERS P., CHIANG J.-Y., BUSCH J., DEBEVEC P.: Temporal upsampling of performance geometry using photometric alignment. *ACM Trans. Graph. 29* (April 2010), 17:1–17:11. 3

[WL01] WEI L.-Y., LEVOY M.: Texture synthesis over arbitrary manifold surfaces. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 2001), SIGGRAPH '01, ACM, pp. 355–360. 2

[WLKT09] WEI L.-Y., LEFEBVRE S., KWATRA V., TURK G.: State of the art in example-based texture synthesis. In *Eurographics 2009, State of the Art Report, EG-STAR* (2009), Eurographics Association. 2

[WM04] WANG L., MUELLER K.: Generating sub-resolution detail in images and volumes using constrained texture synthesis. In *Proceedings of the conference on Visualization '04* (Washington, DC, USA, 2004), VIS '04, IEEE Computer Society, pp. 75–82. 2

[WMP*06] WEYRICH T., MATUSIK W., PFISTER H., BICKEL B., DONNER C., TU C., MCANDLESS J., LEE J., NGAN A., JENSEN H. W., GROSS M.: Analysis of human faces using a measurement-based skin reflectance model. *ACM Transactions on Graphics 25*, 3 (July 2006), 1013–1024. 2

[Woo78] WOODHAM R. J.: Photometric stereo: A reflectance map technique for determining surface orientation from image intensity. In *Proc. SPIE's 22nd Annual Technical Symposium* (1978), vol. 155. 2

[XYZ] XYZRGB: 3D laser scanning - XYZ RGB Inc. http://www.xyzrgb.com/. 2, 4

[YHBZ01] YING L., HERTZMANN A., BIERMANN H., ZORIN D.: Texture and shape synthesis on surfaces. In *Proceedings of the 12th Eurographics Workshop on Rendering Techniques* (London, UK, 2001), Springer-Verlag, pp. 301–312. 2

[ZJMB11] ZHAO S., JAKOB W., MARSCHNER S., BALA K.: Building volumetric appearance models of fabric using micro ct imaging. *ACM Trans. Graph. 30* (Aug. 2011), 44:1–44:10. 3

[ZREB06] ZICKLER T., RAMAMOORTHI R., ENRIQUE S., BELHUMEUR P. N.: Reflectance sharing: Predicting appearance from a sparse set of images of a known shape. *PAMI 28*, 8 (2006), 1287–1302. 2