Skin Microstructure Deformation with Displacement Map Convolution

Koki Nagano[†] * Graham Fyffe[†] Oleg Alexander[†] Jernej Barbič[‡] Hao Li[‡] Abhijeet Ghosh[⋆] Paul Debevec[†]
[†]USC Institute for Creative Technologies [‡]University of Southern California [⋆]Imperial College London

(a) no microstructure (b) static microstructure (c) dynamic microstructure (d) photograph

Figure 1: *A rendered facial expression with (a) mesostructure only (b) static microstructure from a neutral expression (c) dynamic microstructure from convolving the neutral microstructure according to local surface strain compared to a reference photograph of the similar expression. The insets show detail from the lower-left area.*

Abstract

We present a technique for synthesizing the effects of skin microstructure deformation by anisotropically convolving a high-resolution displacement map to match normal distribution changes in measured skin samples. We use a 10-micron resolution scanning technique to measure several *in vivo* skin samples as they are stretched and compressed in different directions, quantifying how stretching smooths the skin and compression makes it rougher. We tabulate the resulting surface normal distributions, and show that convolving a neutral skin microstructure displacement map with blurring and sharpening filters can mimic normal distribution changes and microstructure deformations. We implement the spatially-varying displacement map filtering on the GPU to interactively render the effects of dynamic microgeometry on animated faces obtained from high-resolution facial scans.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation;

Keywords: skin microgeometry, deformation, measurement, displacement map, convolution

1 Introduction

Simulating the appearance of human skin is important for rendering realistic digital human characters for simulation, education, and entertainment applications. Skin exhibits great variation in color, surface roughness, and translucency over different parts of the body, between different individuals, and when it's transformed by articulation and deformation. But as variable as skin can be, human perception is remarkably attuned to the subtleties of skin appearance, as attested to by the vast array of makeup products designed to enhance and embellish it.

Advances in measuring and simulating the scattering of light beneath the surface of the skin [Jensen et al. 2001; Weyrich et al. 2006; d'Eon et al. 2007] have made it possible to render convincingly realistic human characters whose skin appear to be fleshy and organic. Today's high-resolution facial scanning techniques (e.g. [Ma et al. 2007; Beeler et al. 2010; Ghosh et al. 2011] record facial geometry, surface coloration, and surface *mesostructure* details at the level of skin pores and fine creases to a resolution of up to a tenth of a millimeter. By recording a sequence of such scans [Beeler et al. 2011] or performing blendshape animation using scans of different high-res expressions (e.g. [Alexander et al. 2010; Fyffe et al. 2014], the effects of dynamic mesostructure – pore stretching and skin furrowing – can be recorded and reproduced on a digital character.

Recently, [Graham et al. 2013] recorded skin *microstructure* at a level of detail below a tenth of a millimeter for sets of skin patches on a face, and showed that texture synthesis could be used to increase the resolution of a mesostructure-resolution facial scan to one with microstructure everywhere. They demonstrated that skin microstructure makes a significant difference in the appearance of skin, as it gives rise to a face's characteristic pattern of spatially-varying surface roughness. However, they recorded skin microstructure only for static patches from neutral facial expressions, and did not record the dynamics of skin microstructure as skin stretches and compresses.

Skin microstructure, however, is remarkably dynamic as a face makes different expressions. Fig. 2 shows a person's forehead as they make surprised, neutral, and angry expressions. In the neutral expression (center), the rough surface microstructure is relatively isotropic. When the brow is raised (left), there are not only mesostructure furrows but the microstructure also develops a pattern of horizontal ridges less than 0.1 mm across. In the perplexed expression (right), the knitted brow forms vertical anisotropic structures in its microstructure. Seen face to face or filmed in closeup, such dynamic microstructure is a noticeable aspect of human expression, and the anisotropic changes in surface roughness affect the appearance of specular highlights even from a distance.

Dynamic skin microstructure results from the epidermal skin layers being stretched and compressed by motion of the tissues under-

*e-mail:nagano@ict.usc.edu

Figure 2: *Three real forehead expressions (surprised, neutral, and perplexed) made by the same subject showing anisotropic deformations in microstructure.*

neath. Since the skin surface is relatively stiff, it develops a rough microstructure to effectively store a reserve of surface area to prevent rupturing when extended. Thus, parts of the skin which stretch and compress significantly (such as the forehead and around the eyes) are typically rougher than parts which are mostly static, such as the tip of the nose or the top of the head. When skin stretches, the microstructure flattens out and the surface appears less rough as the reserves of tissue are called into action. Under compression, the microstructure bunches up, creating micro-furrows which exhibit anisotropic roughness. Often, stretching in one dimension is accompanied by compression in the perpendicular direction to maintain the area of the surface or the volume of tissues below. A balloon provides a clear example of roughness changes under deformation: the surface is diffuse at first, and becomes shiny when inflated.

While it would be desirable to simulate these changes in appearance during facial animation, current techniques do not record or simulate dynamic surface microstructure for facial animation. One reason scale: taking the facial surface to be 25cm $\times$ 25cm, recording facial shape at 10 micron resolution would require real-time Gigapixel imaging beyond the capabilities of today's camera arrays. And simulating a billion triangles of skin surface, let alone several billion tetrahedra of volume underneath, would be computationally very expensive using finite element techniques.

In this work, we approximate the first-order effects of dynamic skin microstructure by performing fast image processing on a high-resolution skin microstructure displacement map obtained as in [Graham et al. 2013]. Then, as the skin surface deforms, we blur the displacement map along the direction of stretching, and sharpen it along the direction of compression. On a modern GPU, this can be performed at interactive rates, even for facial skin microstructure at ten micron resolution. We determine the degree of blurring and sharpening by measuring *in vivo* surface microstructure of several skin patches under a range of stretching and compression, tabulating the changes in their surface normal distributions. We then choose the amount of blurring or sharpening to affect a similar change in surface normal distribution on the microstructure displacement map. While our technique falls short of simulating all the observable effects of dynamic microstructure, it produces measurement-based changes in surface roughness and anisotropic changes in surface microstructure orientation consistent with real skin deformation. For validation, we compare renderings using our technique to real photographs of faces making similar expressions.

2 Background and Related Work

Our skin is a complex multilayered organ that plays numerous roles in protection, heat regulation, sensing, and hydration. As the principal surface we see when we look at each other, it is also key in human communication, telling others about our health, our age, our physical state, and our emotions.

Skin is elastic, with an ability to safely stretch an average of 60 to 75 percent [Edwards and Marks 1995]. The top layers of the skin, the epidermis, are relatively stiff and achieve much of their elasticity through a network of fine-scale ridges and grooves which provide a reserve of tissue which can flatten when pulled [Montagna and Parakkal 1974]. This microstructure varies in scale and texture throughout the body and is responsible for the specular BRDF of the skin. As seen in Fig. 2, skin's texture orientation and BRDF change dramatically as the skin is subjected to stretching and compression.

The physical and mechanical properties of skin have been studied significantly, including the relationship between deformation and changes in specular reflectance. [Ferguson and Barbenel 1981] measured skin roughness under varying strain and quantified how stretching reduces the surface roughness along that direction (Fig. 3). They also showed a relationship between the surface profile length and the distribution of surface orientations. [Federici et al. 1999], [Guzelsu et al. 2003], and [Schulkin et al. 2003] use surface roughness of soft tissue materials as an indicator of mechanical effects including applied stresses and strain. [Federici et al. 1999] and [Guzelsu et al. 2003] proposed a noninvasive method to measure the stretch of soft tissues including skin based on specular reflectivity. They used polarized light to measure specular reflectance and observed that the reflected light increases with stretching as the surface becomes smoother. [Schulkin et al. 2003] extended the measurement to characterize how subsurface reflectance changes in response to mechanical effects.

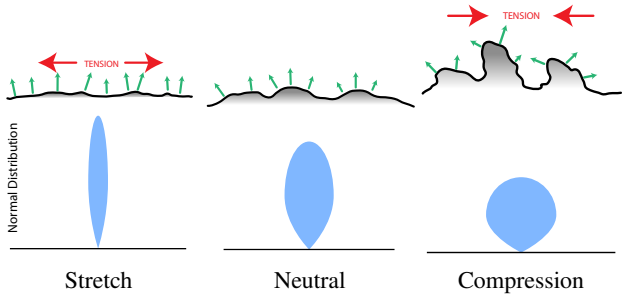

Figure 3: *A normal distribution changing under deformation*

In computer graphics, light reflection from rough surfaces is often modeled by physically-based microfacet distribution models such as [Torrance and Sparrow 1967] and [Cook and Torrance 1982]. A microfacet distribution model simulates roughness by symmetric V-grooves at the microscopic level, called microfacets, which are assumed to behave like a perfect mirror. More recently, Gaussian random microfacet models are favored over V-groove models [Heitz 2014]. The change of microfacet orientations at a micro level changes the resulting surface BRDF and alters the appearance of the surface. Therefore, it is important to take into account such dynamic distributions in simulating a surface BRDF.

Physically-based simulation techniques have been applied to facial animation at a range of scales from overall facial shape (e.g. [Platt and Badler 1981; Terzopoulos and Waters 1993; Sifakis et al. 2005; Bickel et al. 2009]) to surface mesostructure on the order of forehead furrowing and crow's feet around the eyes (e.g. [Bickel

et al. 2012; Rémillard and Kry 2013; Li and Kry 2014]). However, due to computational complexity such techniques have not yet been applied to skin microstructure in the 10 micron scale over an entire face. Data driven techniques have been employed to synthesize facial details onto novel face poses using e.g. polynomial functions [Ma et al. 2008] or statistical models [Golovinskiy et al. 2006], but only at the scale of mesostructure. Recent work in material acquisition has noted the importance of capturing surface microstructure using techniques such as computational tomography [Zhao et al. 2011] and an elastomeric sensor [Johnson et al. 2011]. In our work, we wish to capture *in vivo* skin microstructure without contacting the surface and adapt the microstructure acquisition process of [Graham et al. 2013], which showed that constrained texture synthesis could be used to create microstructure for an entire facial model based on a set of discrete microstructure patches. [von der Pahlen et al. 2014] showed a real-time implementation of skin microstructure using procedural noise functions tailored to match measured skin samples.

Accurately modeling and efficiently rendering the subtle reflection effects of surface microstructure has been an area of significant recent interest [Dupuy et al. 2013; Jakob et al. 2014; Yan et al. 2014] for man-made materials. The technique proposed by [Dupuy et al. 2013] admits a microstructure scaling factor which can uniformly reduce or increase the amplitude of surface microstructure as a material deforms. Our work models a richer set of deformation-based reflectance effects by directionally blurring and sharpening the surface microstructure based on local surface strain, and we tune this filtering to match measurements of real skin patches. However, we do not address efficient anti-aliased rendering techniques.

3 Basic Approach

Our approach to synthesizing skin microstructure deformation is to directionally blur and sharpen a high-resolution surface displacement map according to the amount of stretching or compression. We can visualize this along a one-dimensional cross-section as in Fig. 4. In the center is a 3mm wide cross-section of ventral forearm skin measured in [Liew et al. 2011] with optical coherence tomography (OCT). To its left and right are compressed and stretched versions with no modification to surface height, decreasing the surface length by 20% and increasing it by 35%, respectively. More realistically, the surface would deform to minimize strain. We can approximate this effect by smoothing the height map of the surface under stretching as seen at the right, and sharpening the height map under compression as seen at the left. In both cases, the surface length of the neutral profile is now maintained, causing a greater effect on the distribution of surface normals.

Fig. 5 shows this microstructure convolution technique applied to the surface of a deforming sphere. In the top row, the sphere is rendered with a microstructure displacement map generated with a volume noise function. When shrunk, an isotropic sharpening filter is applied to the displacement map, making it bumpier and giving the sphere a rougher surface reflectance. When expanded, the sphere's displacement map is blurred, giving it a smoother, shinier appearance, as when inflating a balloon. The effects are better seen in the expanded insets of areas near the specular reflections. In the bottom row, the sphere is textured with a displacement map from a real skin microgeometry sample, expanded to the sphere using texture synthesis. The sphere is squashed and stretched to produce anisotropic surface strain, which causes anisotropic filtering of the displacement map: blurring in one dimension and sharpening in the other. This results in anisotropic micro-ridges, similar to those seen in real human skin in Fig. 2 during facial expression.

For the sphere example, the amount of stretching or blurring proportional to the surface strain was chosen by the user to create an appealing dynamic appearance. For rendering a realistic human face, it would be desirable for the filter kernel to be driven according to the behavior of real skin. To this end, we use a measurement apparatus to record the behavior of skin microstructure under stretching and compression as described in the next section.

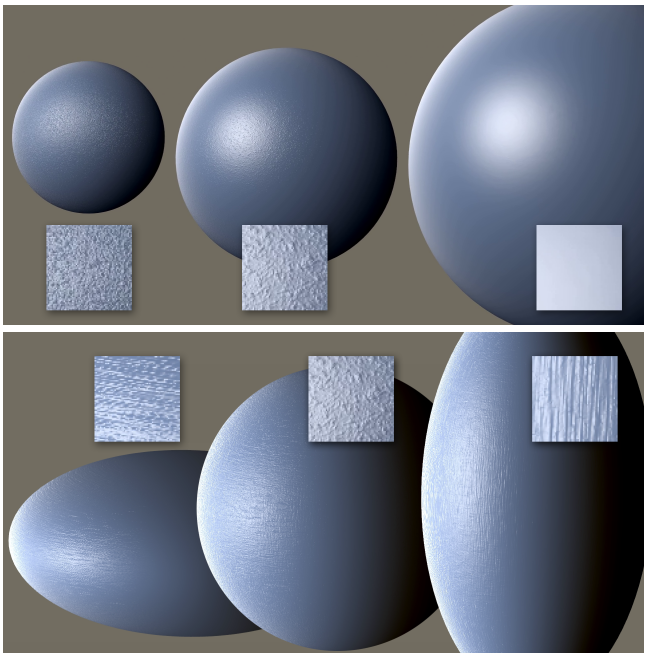

Figure 5: *Deforming sphere with dynamic microgeometry. (Top) The microstructure becomes rougher through displacement map sharpening when shrunk, and smoother through blurring when expanded. The insets show details of the specular highlights. (b) Anisotropic compression and stretching yields anisotropic microstructure.*

4 Measurement

We record the surface microstructure of various skin patches at 10 micron resolution with a setup similar to [Graham et al. 2013] which uses a set of differently lit photos taken with polarized gradient illumination [Ma et al. 2007]. The sample patches are scanned in different deformed states using the lighting apparatus with a custom stretching measuring device consisting of a caliper and a 3D printed stretching aperture. The aperture of the patch holder is set 8 mm for the neutral deformation state and is set 30 cm away from a Ximea machine vision camera which records monochrome 2048 by 2048 pixel resolution images with Nikon 105 mm macro lens at f/16, so that each pixel covers a 6 micron square of skin. The 16 polarized spherical lighting conditions allow the isolation and measurement of specular surface normals, resulting in a per-pixel surface normal map. We integrate the surface normal map to compute a displacement map and use a high pass filter to remove surface detail greater than the scale of a millimeter to remove surface bulging.

Each skin patch, such as part of the forehead, cheek, or chin, is coupled to the caliper aperture using 3M double-sided adhesive tape, and each scan lasts about half a second. After performing the neutral scan, the calipers are narrowed by 0.8mm and the first compressed scan is taken; this continues until the skin inside the aperture buckles significantly. Then, the calipers are returned to neutral,

Figure 4: *Stretching and compressing a measured OCT skin profile, with and without convolution filters to maintain surface length.*

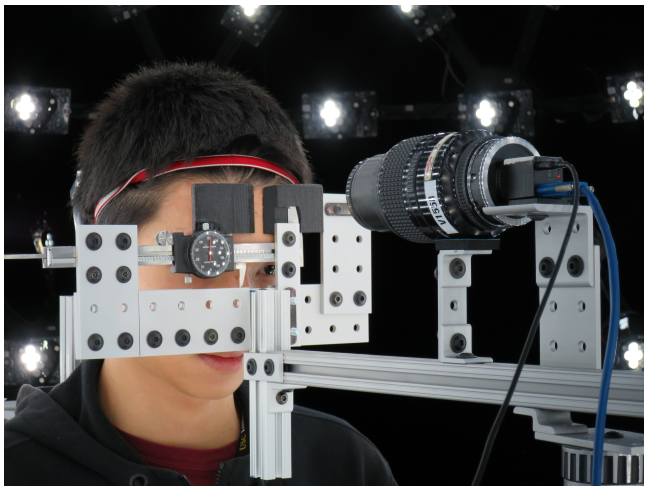

Figure 6: *Microstructure acquistion in a polarized LED sphere with macro camera and articulated skin deformer.*

and scans are taken with progressively increased stretching until the skin detaches from the double-stick tape. Fig. 7 shows a skin sample in five different states of strain. The calipers can be rotated to different angles, allowing the same patch of skin to be recorded in up to four different orientations, such as the forehead sample seen in Fig. 8.

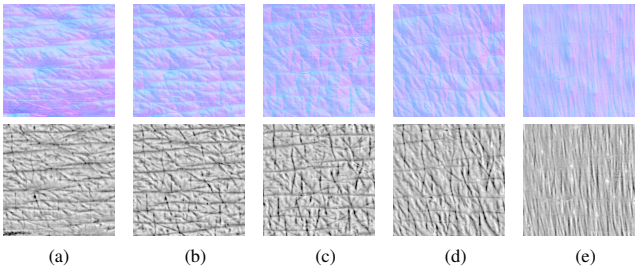

Figure 7: *Texture-aligned surface normal (top) and displacement (bottom) maps of a skin patch under vertical compression and stretching. (a) full compression, (b) medium compression, (c) neutral, (d) medium stretching, and (e) full stretching.*

With skin patch data acquired, we now wish to characterize how surface microfacet distributions change under compression and stretching. After applying a denoising filter to the displacement maps to reduce camera noise, we create a histogram of the surface orientations observed across the skin patch under its range of strain. Several such histograms are visualized in Fig. 8 next to their corresponding skin samples, and can also be thought of the specular lobe which would reflect off the patch. As can be seen, stretched skin becomes anisotropically shinier in the direction of the stretch, and anisotropically rougher in the direction of compression. For some samples, such as the chin in Fig. 8(g,h), we observed some dependence on the stretching direction to the amount of change in normal distributions. However, we do not yet account for the effect of the stretching direction in our model.

The variance in x and y of the surface normal distribution quantify the degree of surface smoothing or roughening according to the amount of strain put on the sample. Fig. 9 plots the changes in surface normal distribution in the direction of the strain for several facial skin patches. Again, stretched skin becomes shinier, and compressed skin becomes rougher.

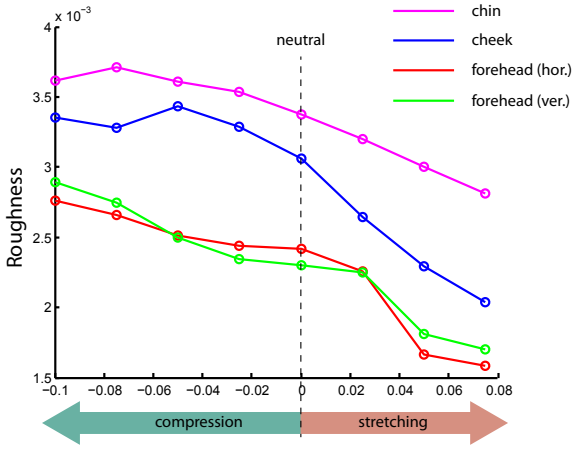

Figure 9: *Surface normal distributions plotted against the amount of strain for several skin patches.*

5 Microstructure Analysis and Synthesis

Based on the skin patch data obtained as in Section 4, we model the relationship between the measured surface normal distribution and surface stretching or compression (collectively *deformation*). We observe that the change in surface detail in Fig. 7 from neutral (c) to stretched (e) qualitatively resembles a blurring filter in the direction of stretch, and from neutral (c) to compressed (a) qualitatively resembles a sharpening filter in the direction of compression, perhaps with some blurring in the perpendicular direction. These qualitative observations are consistent with the surface normal distribution plots in Fig. 8. Such filters are also inexpensive to compute on GPU hardware. Hence, we design a method to synthesize the microgeometry of skin under deformation using a microgeometry displacement map of the skin in a neutral state and a parametric family of convolution filters ranging continuously from sharpening to blurring. We then synthesize spatially and temporally varying microstructure for faces undergoing dynamic deformation at interactive rates by driving the filter parameters by local deformation fields. Our framework is generic in the sense that the neutral microgeometry displacment map can come from any source, such as noise functions [von der Pahlen et al. 2014] or data-diven image analogy techniques [Graham et al. 2013].

Deformation Model We formulate our model as a convolution in the 2D texture coordinate space associated with the surface geometry. Given a high-resolution microstructure displacement map D for the neutral pose, we synthesize the deformed microdisplacement map D' as:

$$D'(u,v) = (D * K_{u,v})(u,v), \tag{1}$$

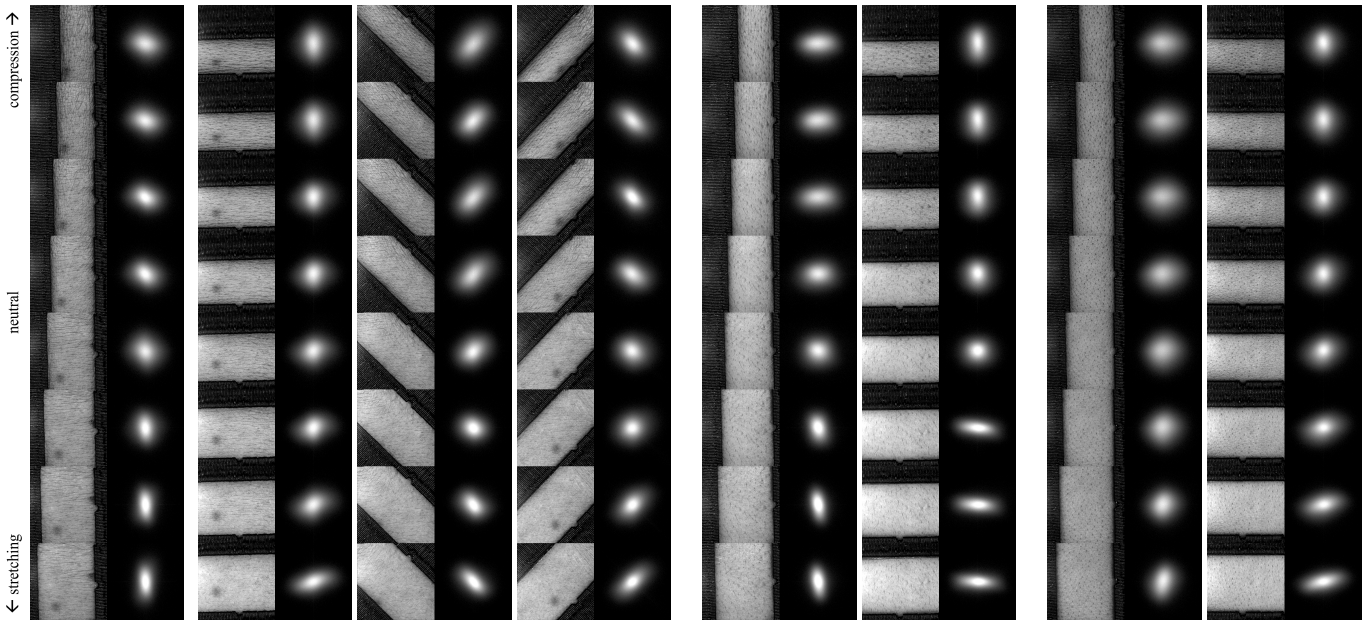

(a) forehead horizontal (b) forehead vertical (c) forehead 45° (d) forehead 135° (e) cheek horizontal (f) cheek vertical (g) chin horizontal (h) chin vertical

Figure 8: *Each column shows measured 8mm wide facial skin patches under different amounts of stretching and compression, with a histogram of the corresponding surface normal distributions shown to the right if each sample.*

where (u, v) are texture space coordinates, $K_{u,v}$ is the convolution kernel for coorinate u, v, and $*$ represents discrete convolution:

$$(D * K)(u, v) = \sum_i \sum_j D(u - i, v - j)K(i, j). \quad (2)$$

Supposing K were constant over the entire surface, and linearly seperable over some perpendicular axes $\mathbf{s}$ and $\mathbf{t}$, we could write $D' = D*K = (D*[\mathbf{s}]\,\mathbf{k}^s)*[\mathbf{t}]\,\mathbf{k}^t$, where $\mathbf{k}^s$ and $\mathbf{k}^t$ are 1D kernels and $*[\mathbf{a}]$ represents convolution along some axis $\mathbf{a} = (\mathbf{a}_\mathrm{u}, \mathbf{a}_\mathrm{v})$:

$$(D*[\mathbf{a}]\,\mathbf{k})(u, v) = \sum_i D(u - i\mathbf{a}_\mathrm{u}, v - i\mathbf{a}_\mathrm{v})\mathbf{k}(i). \quad (3)$$

This allows efficient computation as a sequence of two 1D convolutions, reducing the computational cost from $O(N^2)$ to $O(2N)$ for an $N \times N$ kernel. Unfortunately, in general, the kernel is spatially varying. Nonetheless, if the kernel varies gradually, we may still approximate the spatially varying 2D convolution as a sequence of two spatially varying 1D convolutions, as illustrated in Fig. 10:

$$D'(u, v) \approx (D^s *[\mathbf{t}_{u,v}]\,\mathbf{k}^t_{u,v})(u, v), \quad (4)$$
$$\text{where } D^s(u, v) = (D*[\mathbf{s}_{u,v}]\,\mathbf{k}^s_{u,v})(u, v). \quad (5)$$

In practice we align the axis $\mathbf{s}$ to the direction of stretch (if any) and $\mathbf{t}$ to the direction of compression (if any), where $\mathbf{s}$ and $\mathbf{t}$ are always mutually perpendicular. We employ a 2-parameter family of 1D convolution kernels encompassing sharpening and blurring:

$$\mathbf{k} = (1 - \alpha)\delta + \alpha\mathcal{G}(\cdot, \sigma), \quad (6)$$

where δ is the discrete delta function, $-2 \le \alpha \le 1$ is the filter strength, and σ is the standard deviation of the normalized discrete Gaussian kernel $\mathcal{G}$. With $\alpha > 0$ the filter blurs the signal; with $\alpha < 0$ it sharpens; and with $\alpha = 0$ it preserves the signal.

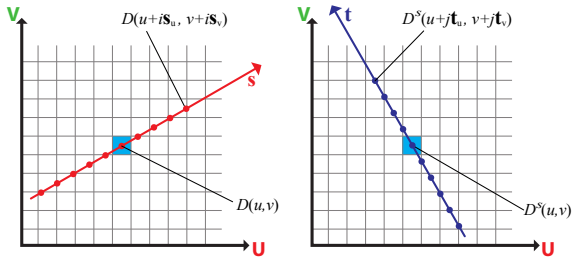

Figure 10: *Displacement map pixels are sampled along the principal directions of strain with a separable filter for convolution. D^s is computed from D, then D' is computed from D^s.*

Parameter Fitting We estimate the parameters (α^s, σ^s) defining $\mathbf{k}^s$ and (α^t, σ^t) defining $\mathbf{k}^t$ for each measured skin patch using a brute-force approach, where the principle directions $\mathbf{s}$ and $\mathbf{t}$ are known for each patch. We assume these parameters to be constant over the extent of each patch. We search for $\alpha \in (-2 \ldots 1)$ with a granularity of 0.01 and $\sigma \in (1\mu\mathrm{m} \ldots 100\mu\mathrm{m})$ with a granularity of 0.5μm. We find the parameters that minimize the total variation between surface normal histograms of the ground truth displacement map and the convolved neutral displacement map. The surface normals are computed in the target deformation coordinates, meaning the neutral displacement map is stretched or compressed after convolution to match the deformed shape. We compute the surface normal histograms on the GPU by splatting the surface normal computed at each pixel in the displacement map into a grid of buckets, based on the u and v components of the normal in tangent space. Care must be taken that enough samples are available for the number of buckets used, in order to avoid bucket aliasing which may misguide the optimization. We used patches with up to 2000×2000 pixels and 64×64 buckets. For efficiency, we optimize only (α^s, σ^s) or (α^t, σ^t) for the kernel in the direction of caliper movement, omitting the kernel in the the other direction. The total variation metric is simply the absolute difference between the

ground truth histogram and the histogram of the convolved neutral patch, summed over the buckets:

$$\alpha, \sigma = \underset{\hat{\alpha},\hat{\sigma}}{\operatorname{argmin}} \sum_{b \in B} \left| H_b - \hat{H}_b \right|, \tag{7}$$

where B is the set of histogram buckets, H is the ground truth histogram, and $\hat{H}$ is the histogram obtained by 1D convolution of the neutral skin patch using parameters $\hat{\alpha}, \hat{\sigma}$ in the direction of movement. Fig. 11 tabulates the fitted parameters for a patch of forehead skin undergoing a range of stretching and compression.

r	0.73	0.79	0.88	0.94	1.08	1.15	1.22	1.28
α	-2	-1.76	-0.2	-0.02	0.54	0.27	1	1
σ	6.5	6.5	13	22	14.5	19	17	24

Figure 11: *Fitted kernel parameters α and σ (μm) for a patch of forehead skin undergoing a range of stretching and compression. r is the stretch ratio, with $r > 1$ stretching and $r < 1$ compressing.*

Kernel Table Construction In order to apply our model to skin undergoing a novel deformation, we must first establish a relationship between the deformation and the kernel parameters, using the table of parameters fitted to sample patches. For each measured skin patch, the principal axes are known and also the stretch ratio r is known, defining the deformation. When $r > 1$, the patch is undergoing stretching, and when $r < 1$, the patch is undergoing compression. Fig. 12 plots the fitted kernel parameters α and σ against the stretch ratio r along the primary axis of stretch or compression. Based on these data points, we fit piecewise linear models relating α and σ to r. We partition the domain of r into line segments by manually inspecting the data. For α, we use connected line segments to enforce smooth transitions between compression and stretching, and we constrain one segment to pass through neutral case ($r = 1, \alpha = 0$) exactly. For σ, we fit one line segment to the compression samples and another line segment to the stretch samples, with a discontinuity at $r = 1$. For example, the forehead patch yields the following model:

$$\alpha = \min(1, 15.4r - 13.8, 3.09(r-1)); \tag{8}$$

$$\sigma = \begin{cases} (38.2r - 26.5)\mu\text{m} & \text{if } r \geq 1 \\ (70.5r - 46.9)\mu\text{m} & \text{otherwise.} \end{cases} \tag{9}$$

As we are not assigning any physically meaningful interpretation to α and σ, the fitted functions serve merely as a rapid lookup to compute parameters from stretch ratios. Other functions could be employed if desired, but it is important that $\alpha = 0$ in the neutral case, and hence a single line fit to α would not fit both the stretching behavior and compression behavior well.

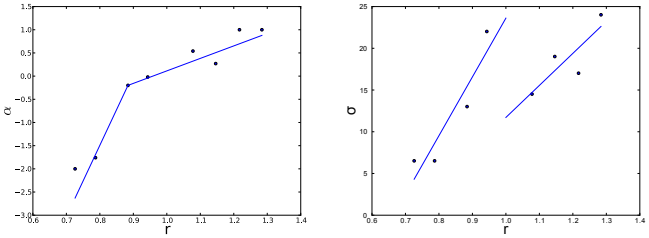

Figure 12: *Fitted kernel parameters α and σ plotted against the stretch ratio r. Dots represent parameters fitted to sample patches, and lines represent the piecewise linear fits in* (8)*,* (9)*.*

Applying to a Novel Surface Animation Given a triangular mesh undergoing novel deformation, we may estimate the principal directions and stretch ratios of the local deformation at every vertex on the mesh, and in turn compute the associated kernel parameters, allowing the deformed displacement map to be synthesized using the two-pass separable convolution technique in a GPU fragment shader. Our generic framework allows deformation from physical simulation, keyframe animation, or facial performance capture. Consider a triangle P with v_0, v_1, and v_2 being vertices on the neutral face, and the deformed counterpart P' likewise with v_0', v_1', and v_2' being the vertices. We compute the rotation R that maps the triangle P in 3D space to the so-called tangent space. Tangent space may be determined by first computing the affine 2×3 transform Q that maps world coordinates to UV texture coordinates, and then computing R as the closest rotation matrix to Q (using e.g. SVD or QR decomposition), first concatenating to it a third row that is the cross product of the first two rows. As the last component in tangent space is irrelevant, we truncate the last row of R again leaving a 2×3 matrix. We likewise compute R' to map P' into 2D tangent space. Then we compute a linear transformation T that maps the 2D neutral triangle RP onto the deformed triangle $R'P'$:

$$\begin{bmatrix} e_1' & e_2' \end{bmatrix} = T \begin{bmatrix} e_1 & e_2 \end{bmatrix}, \tag{10}$$

where e_i is an edge from Rv_0 to Rv_i, and the rest is analogous, and the 2 by 2 linear transformation matrix T can be trivially computed. Such T can be found per deformed triangle. In practice, we want as smooth a 2D deformation field as possible in order to drive the per-pixel displacement map without visible seams. To that end, we average the linear transformation T of each face attached to a vertex, and then interpolate the per-vertex T within a GPU fragment shader, and perform Singular Value Decomposition of the form:

$$T = U \Sigma V^{\mathsf{T}}, \tag{11}$$

per-fragment with

$$\Sigma = \begin{bmatrix} r_s & 0 \\ 0 & r_t \end{bmatrix}, \tag{12}$$

where r_s, and r_t are the stretch ratios ($r_s > r_t$). With a 2 by 2 matrix, such a SVD can be trivially computed in closed form in a GPU fragment shader, providing smooth spatially varying 2D deformation fields as illustrated in Fig.13. (Alternatively, if the mesh animation is produced using a physics simulation such as a finite element model, the stretch and strain could be obtained more directly from the simulation [Irving et al. 2004].) The transform from principle deformation axes to UV coordinates is then:

$$S = QR^{\mathsf{T}}V, \tag{13}$$

and hence the principle axes $\mathbf{s}$ and $\mathbf{t}$ in UV coordinates are the first and second columns of S, normalized, and the magnitudes of the columns serve as the conversion factor required for converting σ from world distance to UV distance for convolution. Indeed, the convolution may be performed without conversion if $\mathbf{s}$ and $\mathbf{t}$ are taken as the un-normalized columns of S. Substituting the values r_s and r_t into the parametric kernel model (8), (9) produces the kernels $\mathbf{k}^s$ and $\mathbf{k}^t$ at every point on a deforming surface, allowing the deformed displacement map to be synthesized for the entire face.

6 Results

Fig. 5 shows two deforming spheres with microstructure convolved according to local surface deformation, as described in Sec. 3.

Fig. 14 shows frames from a sequence of a 1cm wide digitized skin patch being deformed by an invisible probe. It uses a relatively

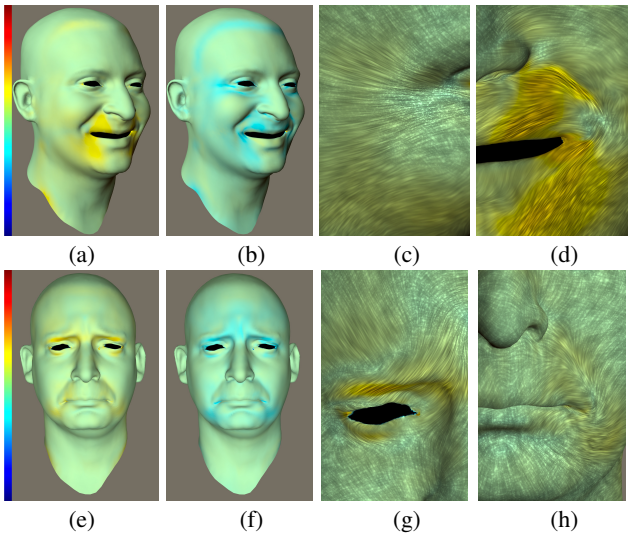

Figure 13: *Strain field visualization for a smile expression (top row) and a sad expression (bottom row) with the first stress eigenvalue (a), (e), and the second eigen value (b), (f), and strain direction visualization (c), (d), (g), and (h).*

low-resolution finite element volumetric mesh with 25,000 tetrahedra to simulate the mesostructure which in turns drives dynamic microstructure convolution. The neutral microstructure was recorded using the system in Fig. 6 at 10 micron resolution from the forehead of a young adult male, and its microstructure is convolved with parameters fit to match its own surface normal distribution changes under deformation as described in Sec. 5. The rendering was made using the V-Ray package to simualte subsurface scattering. As seen in the accompanying video, the skin microstructure bunches up and flattens out as the surface deforms at a resolution much greater than the FEM simulation.

Fig. 1 highlights the effect of using no microgeometry, static microgeometry, and dynamic microgeometry simulated using displacement map convolution with a real-time rendering. Rendering only with 4K resolution mesostructure from a standard facial scan produces too polished an appearance at this scale. Adding static microstructure computed at 16K resolution using a texture synthesis technique [Graham et al. 2013] increases visual detail but produces conflicting surface strain cues in the compressed and stretched areas. Convolving the static microstructure according to the surface strain using normal distribution curves from a related skin patch as in Sec. 5 produces anisotropic skin microstructure consistent with the expression deformation and a more convincing sense of skin under tension as can be observed in the reference photograph of a similar expression on the right. Renderings are made with a real-time hybrid normal shading technique [Ma et al. 2007].

Fig. 15 compares the real-time renderings with static (left) and dynamic microstructure (right) from a facial performance sequence of a female subject. Again, the dynamic microstructure rendering creates anisotropic dynamic microstructure, providing a visceral sense of surface tension around the cheek region undergoing a smile expression. The details can be seen in motion (and, ideally, full-screen) in the accompanying video.

Fig. 16 (a) through (c) show qualitative validation of our technique (b) compared with a reference photograph of a similar facial region and expression. Facial details with the specular channel on its own (a) highlight the dynamic microstructure which introduces qualitatively similar varying angles of anisotropic surface texture in the

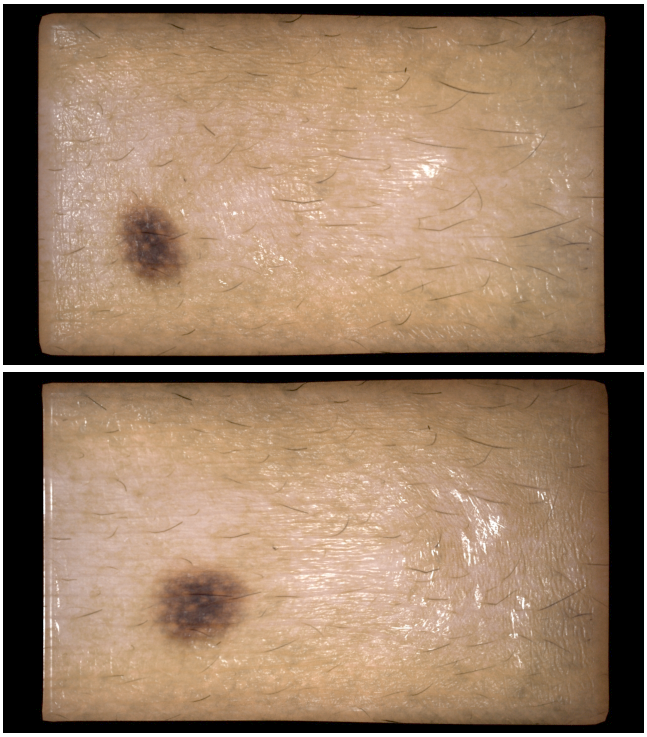

Figure 14: *A sampled skin patch is deformed with FEM which drives microstructure convolution, rendered with path tracing.*

areas undergoing stretching and compression. Fig. 16 (d) through (f) show additional renderings of an older subject, exhibiting plausible dynamic surface details under natural facial expressions.

7 Discussion

While in several ways our skin microstructure deformation technique produces plausible results, it is important to note that it is at best an approximation to the complex tissue dynamics which occur at the microscale. It does, however, get two perceptually important aspects correct, which is the local orientation of the anisotropic microstructure, and, by construction, the anisotropic surface normal distributions. As a result, given the relative efficiency of image convolution, this technique may prove useful for increasing the realism and skin-like quality of virtual characters for both interactive and offline rendering.

8 Future Work

Our technique of simulating dynamic skin microstructure leaves open several areas for future work. First, we currently employ simple linear kernels to filter the displacement map to approximate the effects of stretching and compression. A limitation of linear filters is that they apply the same amount of smoothing or sharpening to both ridges and grooves, whereas a ridge, filled with tissue, is likely to deform less than a groove, filled with air. One can imagine developing nonlinear filters for microstructure deformation which might better simulate the local shape of the deformed skin surface.

Second, while the results in this paper were made using supersampling for antialiasing, efficient rendering of dynamic microstructure should leverage multiresolution surface detail rendering techniques such as LEAN [Olano and Baker 2010] or LEADR [Dupuy et al. 2013] mapping. LEADR mapping may be especially rele-

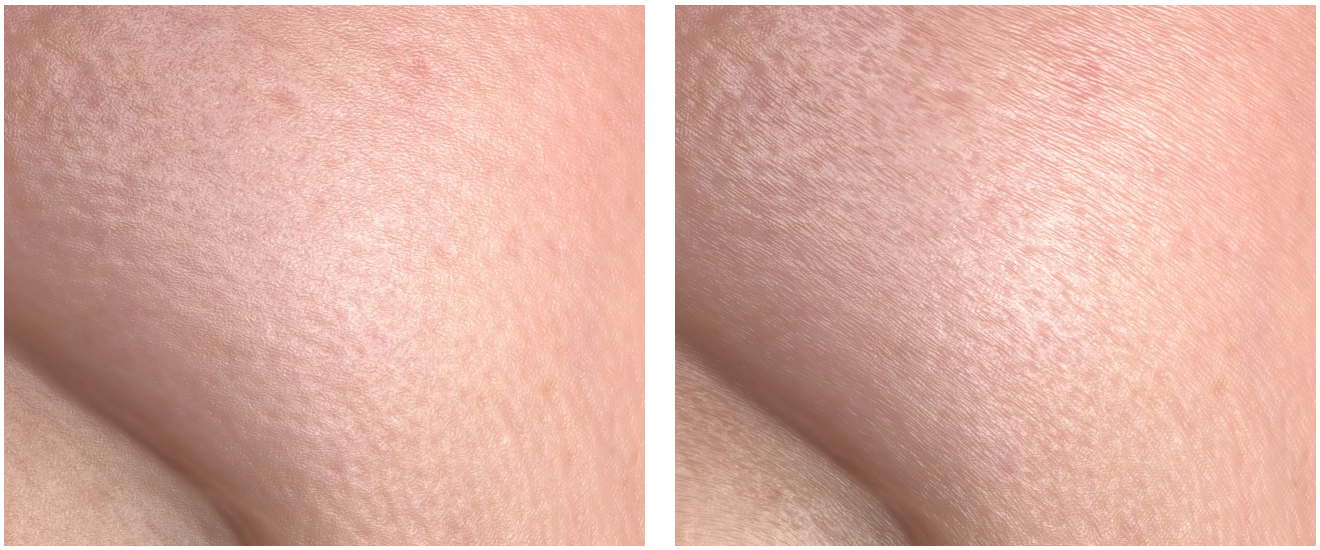

Figure 15: *Real-time rendering of the cheek region from a facial performance animation with enhanced dynamic surface details (right) compared to the static microstructure rendering (left). The dynamic microstructure provides an additional indication of the deformation on the cheek when the subject makes a smile expression.*

vant since it is designed to work with deformable animated surfaces and includes a microfacet BRDF model with masking and shadowing. Ideally, the efficient resampling schemes of LEADR mapping could be augmented to accomodate the spatially-varying convolutions performed in our microstructure deformation process.

9 Conclusion

We have presented a fast approximate approach to simulating the effects of deforming surface microstructure under compression and stretching where a high-resolution displacement map is blurred and sharpened according to local surface strain. We measured normal distributions of real skin samples under stretching and compression to drive the amount of blurring and compression for animating faces. The results show a greater visual indication of surface tension seen in the surface reflections than using a static microstructure displacement map, and more skin-like behavior for deforming surfaces. Since the technique can be implemented on the GPU at interactive rates, it may be useful for rendering high-quality animated characters both for pre-rendered and interactive applications.

Acknowledgements

The authors would like to thank Randal Hill, Kimberly Lu, Ari Shapiro, Cary Peng, Bill Phelps, Emily OBrien, Jay Busch, Xueming Yu, Etienne Danvoye, Javier von del Pahlen, the Digital Human League, Valerie Dauphin, and Kathleen Haase for their assistance and support. This research was sponsored by the U.S. Army Research Laboratory (ARL), the Funai Foundation for Information Technology, and in part by the National Science Foundation (CAREER-1055035, IIS-1422869) and the Sloan Foundation, and Royal Society Wolfson Research Merit Award. The content of the information does not necessarily reflect the position or the policy of the US Government, and no official endorsement should be inferred.

References

ALEXANDER, O., ROGERS, M., LAMBETH, W., CHIANG, J.-Y., MA, W.-C., WANG, C.-C., AND DEBEVEC, P. 2010. The Digital Emily Project: Achieving a photoreal digital actor. *IEEE Computer Graphics and Applications 30* (July), 20–31.

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.

BEELER, T., HAHN, F., BRADLEY, D., BICKEL, B., BEARDSLEY, P., GOTSMAN, C., SUMNER, R. W., AND GROSS, M. 2011. High-quality passive facial performance capture using anchor frames. *ACM Trans. Graph. 30*, 4 (July), 75:1–75:10.

BICKEL, B., BÄCHER, M., OTADUY, M. A., MATUSIK, W., PFISTER, H., AND GROSS, M. 2009. Capture and modeling of non-linear heterogeneous soft tissue. *ACM Trans. Graph. 28*, 3 (July), 89:1–89:9.

BICKEL, B., KAUFMANN, P., SKOURAS, M., THOMASZEWSKI, B., BRADLEY, D., BEELER, T., JACKSON, P., MARSCHNER, S., MATUSIK, W., AND GROSS, M. 2012. Physical face cloning. *ACM Trans. Graph. 31*, 4 (July), 118:1–118:10.

COOK, R. L., AND TORRANCE, K. E. 1982. A reflectance model for computer graphics. *ACM Trans. Graph. 1*, 1 (Jan.), 7–24.

D'EON, E., LUEBKE, D., AND ENDERTON, E. 2007. Efficient rendering of human skin. In *Proceedings of the 18th Eurographics Conference on Rendering Techniques*, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, EGSR'07, 147–157.

DUPUY, J., HEITZ, E., IEHL, J.-C., POULIN, P., NEYRET, F., AND OSTROMOUKHOV, V. 2013. Linear efficient antialiased displacement and reflectance mapping. *ACM Trans. Graph. 32*, 6 (Nov.), 211:1–211:11.

EDWARDS, C., AND MARKS, R. 1995. Evaluation of biomechanical properties of human skin. *Clinics in Dermatology 13*, 375–380.

FEDERICI, J. F., GUZELSU, N., LIM, H. C., JANNUZZI, G., FINDLEY, T., CHAUDHRY, H. R., AND RITTER, A. B. 1999.

Noninvasive light-reflection technique for measuring soft-tissue stretch. *Appl. Opt. 38*, 31 (Nov), 6653–6660.

FERGUSON, J., AND BARBENEL, J. 1981. Skin surface patterns and the directional mechanical properties of the dermis. In *Bioengineering and the Skin*, R. Marks and P. Payne, Eds. Springer Netherlands, 83–92.

FYFFE, G., JONES, A., ALEXANDER, O., ICHIKARI, R., AND DEBEVEC, P. 2014. Driving high-resolution facial scans with video performance capture. *ACM Trans. Graph. 34*, 1 (Dec.), 8:1–8:14.

GHOSH, A., FYFFE, G., TUNWATTANAPONG, B., BUSCH, J., YU, X., AND DEBEVEC, P. 2011. Multiview face capture using polarized spherical gradient illumination. *ACM Trans. Graph. 30*, 6 (Dec.), 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 Trans. Graph. 25*, 3 (July), 1025–1034.

GRAHAM, P., TUNWATTANAPONG, B., BUSCH, J., YU, X., JONES, A., DEBEVEC, P., AND GHOSH, A. 2013. Measurement-based synthesis of facial microgeometry. *Computer Graphics Forum 32*, 2pt3, 335–344.

GUZELSU, N., FEDERICI, J. F., LIM, H. C., CHAUHDRY, H. R., RITTER, A. B., AND FINDLEY, T. 2003. Measurement of skin stretch via light reflection. *Journal of Biomedical Optics 8*, 1, 80–86.

HEITZ, E. 2014. Understanding the masking-shadowing function in microfacet-based brdfs. *Journal of Computer Graphics Techniques (JCGT) 3*, 2 (June), 32–91.

IRVING, G., TERAN, J., AND FEDKIW, R. 2004. Invertible finite elements for robust simulation of large deformation. In *Proceedings of the 2004 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, SCA '04, 131–140.

JAKOB, W., HAŠAN, M., YAN, L.-Q., LAWRENCE, J., RAMAMOORTHI, R., AND MARSCHNER, S. 2014. Discrete stochastic microfacet models. *ACM Trans. Graph. 33*, 4 (July), 115:1–115:10.

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.

LI, P., AND KRY, P. G. 2014. Multi-layer skin simulation with adaptive constraints. In *Proceedings of the Seventh International Conference on Motion in Games*, ACM, New York, NY, USA, MIG '14, 171–176.

LIEW, Y. M., MCLAUGHLIN, R. A., WOOD, F. M., AND SAMPSON, D. D. 2011. Reduction of image artifacts in three-dimensional optical coherence tomography of skin in vivo. *Journal of Biomedical Optics 16*, 11, 116018–116018–10.

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.

MA, W.-C., JONES, A., CHIANG, J.-Y., HAWKINS, T., FREDERIKSEN, S., PEERS, P., VUKOVIC, M., OUHYOUNG, M., AND DEBEVEC, P. 2008. Facial performance synthesis using deformation-driven polynomial displacement maps. *ACM Trans. Graph. 27*, 5 (Dec.), 121:1–121:10.

MONTAGNA, W., AND PARAKKAL, P. F. 1974. *The structure and function of skin*. Academic Press, New York, NY, USA, ch. 2.

OLANO, M., AND BAKER, D. 2010. Lean mapping. In *Proceedings of the 2010 ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*, ACM, New York, NY, USA, I3D '10, 181–188.

PLATT, S. M., AND BADLER, N. I. 1981. Animating facial expressions. *SIGGRAPH Comput. Graph. 15*, 3 (Aug.), 245–252.

RÉMILLARD, O., AND KRY, P. G. 2013. Embedded thin shells for wrinkle simulation. *ACM Trans. Graph. 32*, 4 (July), 50:1–50:8.

SCHULKIN, B., LIM, H. C., GUZELSU, N., JANNUZZI, G., AND FEDERICI, J. F. 2003. Polarized light reflection from strained sinusoidal surfaces. *Appl. Opt. 42*, 25 (Sep), 5198–5208.

SIFAKIS, E., NEVEROV, I., AND FEDKIW, R. 2005. Automatic determination of facial muscle activations from sparse motion capture marker data. *ACM Trans. Graph. 24*, 3 (July), 417–425.

TERZOPOULOS, D., AND WATERS, K. 1993. Analysis and synthesis of facial image sequences using physical and anatomical models. *IEEE Trans. Pattern Anal. Mach. Intell. 15*, 6 (June), 569–579.

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

VON DER PAHLEN, J., JIMENEZ, J., DANVOYE, E., DEBEVEC, P., FYFFE, G., AND ALEXANDER, O. 2014. Digital ira and beyond: Creating real-time photoreal digital actors. In *ACM SIGGRAPH 2014 Courses*, ACM, New York, NY, USA, SIGGRAPH '14, 1:1–1:384.

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 Trans. Graph. 25*, 3 (July), 1013–1024.

YAN, L.-Q., HAŠAN, M., JAKOB, W., LAWRENCE, J., MARSCHNER, S., AND RAMAMOORTHI, R. 2014. Rendering glints on high-resolution normal-mapped specular surfaces. *ACM Trans. Graph. 33*, 4 (July), 116:1–116:9.

ZHAO, S., JAKOB, W., MARSCHNER, S., AND BALA, K. 2011. Building volumetric appearance models of fabric using micro ct imaging. *ACM Trans. Graph. 30*, 4 (July), 44:1–44:10.

Figure 16: *Real-time comparison renderings from a blendshape animation (middle) compared to the reference photographs of a similar expression (right). Specular-only real-time renderings show anisotropic dynamic microstructure at different orientations in the expressions (left). Top two rows: young male subject's crow's feet region with the eyes shut tightly, and the stretched cheek region when the mouth is pulled left. Middle two rows: young female subject crow's feet region, and the nose under the squint expression. Bottom row: older subject's mouth with a smile expression (d), raised down forehead (e), and the stretched cheek and mouth regions (f).*