


Prototyping and Transforming Facial Textures for Perception Research



Bernard Tiddeman, Michael Burt, and David Perrett
University of St. Andrews, Scotland

We present new, wavelet-based methods for prototyping facial textures and for artificially transforming the age of facial images.

Transforming facial images along perceived dimensions (such as age, gender, race, or health) has application in areas as diverse as psychology, medicine, and forensics. We can use prototype images to define the salient features of a particular face classification (for example, European female adult or East-Asian male child). We then use the differences between two prototypes to define an axis of transformation, such as younger to older. By applying these changes to a given input face, we can change its apparent age, race, or gender.

Psychological investigations reveal a limitation with existing methods that's particularly apparent when changing the age of faces.¹

We relate the problem to the loss of facial textures (such as stubble and wrinkles) in the prototypes due to the blending process. In this article, we review the existing face prototyping and transformation methods and present a new, wavelet-based method for prototyping and transforming facial textures.

Previous methods

Francis Galton² first developed facial prototyping more than 120 years ago—he made multiple exposure photographic images of several faces after aligning the eye positions. The modern equivalents use image warping³ to distort the component face images into the mean shape before blending the color values. Aligning the facial features improves the facial prototype's clarity.

In previous work, our laboratory constructed each prototype by averaging a set of facial images in terms of 2D shape and pixel color⁴ (Figure 1). In this work, we delineate the shape of each face in the set with 179 points located along contours around the major facial features (eyes, nose, and mouth) and the facial border. We can place these points manually or automatically by

using Active Shape Models⁵ after training with a manually delineated face set. We find the average shape by averaging the position of each delineated point across the set—for example,

$$\bar{\mathbf{x}} = \frac{1}{N} \sum_{i=0}^N \mathbf{x}_i \quad (1)$$

where \mathbf{x}_i is the i th shape vector made from the x and y coordinates of the n delineated face points,

$$\mathbf{x}_i = (x_0^i, y_0^i, x_1^i, y_1^i, \dots, x_n^i, y_n^i) \quad (2)$$

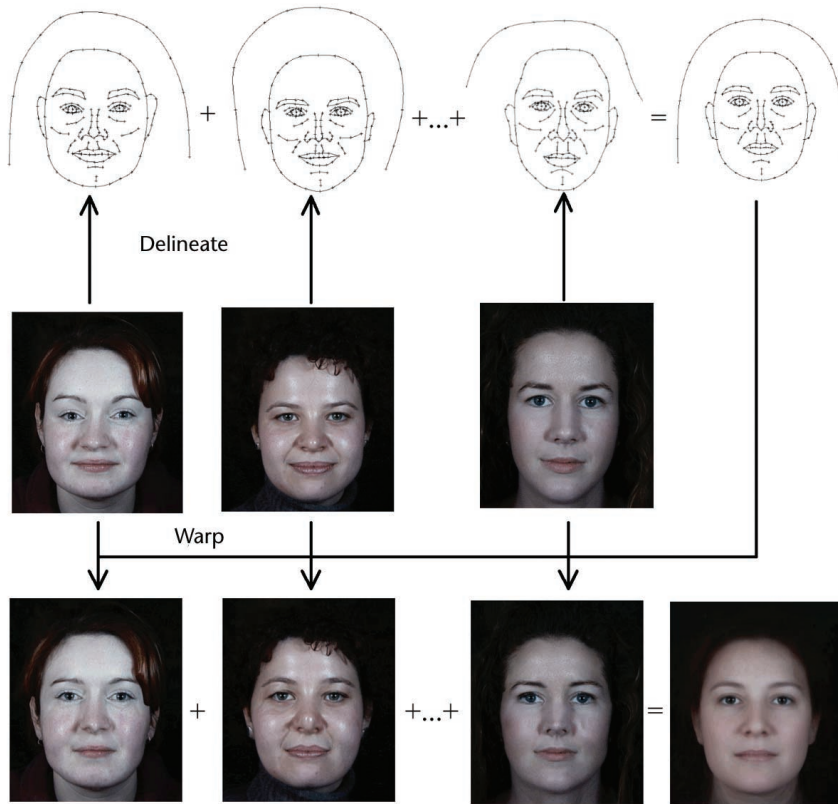
and $\bar{\mathbf{x}}$ is the mean shape vector of N delineated faces.

We find the color of each pixel in the prototype image by warping each component image into the average shape and calculating the mean color. For example,

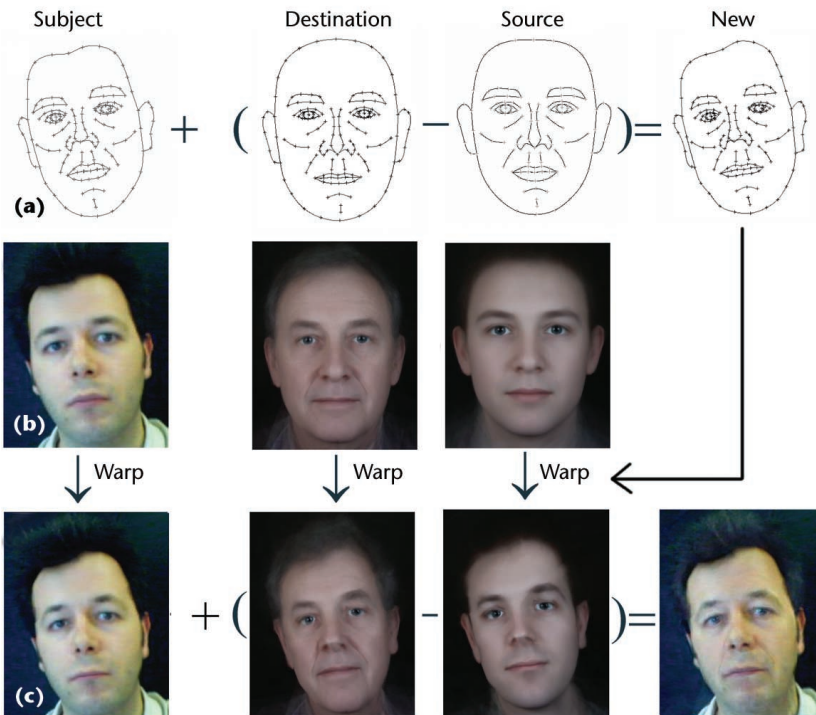
$$\bar{\mathbf{c}}(x, y) = \frac{1}{N} \sum_{i=0}^N \mathbf{c}_i(W_x^i(x, y), W_y^i(x, y)) \quad (3)$$

where $\mathbf{c}_i(x, y)$ is the red, green, and blue (RGB) color vector of image i at point (x, y) , (W_x^i, W_y^i) is the translation vector given by the warping function for image i to the average shape, and $\bar{\mathbf{c}}$ is the mean color over the N images. In this work we use linear warping over the feature points' Delaunay triangulation, which we can combine with a one-to-one constraint algorithm to prevent the warped image folding.⁶

We can use the facial prototypes just described to define the typical differences between two sets of images. We can apply these differences to a delineated subject's face image⁴ or a frame from a moving sequence⁷ in two principal steps (Figure 2). The first step calculates the face's new shape. We normalize the position, orientation, and scale of the two prototype's shape templates to the subject's using a least-squares rigid-body fit (translation, scaling, and rotation). We then cal-



1 Constructing a shape and color prototype. We delineate (top row) and average the facial features of each input face (center row) to define the mean shape. We then warp each component image into the mean shape (bottom row) and average it to produce the prototype image.




2 An illustration of the shape and color transformation process: (a) Define new shape, (b) warp subject and prototypes into new shape, and (c) transform colors at each pixel. (The prototype rigid-body normalization step isn't shown.)

culate and scale (if desired) each feature point's translation from the source prototype (usually of the same class as the image undergoing the transform) to the destination prototype and add it to the corresponding point on the subject's face. For example,

$$\mathbf{x}'_s = \mathbf{x}_s + (\mathbf{x}_{p2} - \mathbf{x}_{p1}) \quad (4)$$

where \mathbf{x}_s is the original shape vector, \mathbf{x}'_s is the transformed shape vector, and \mathbf{x}_{p1} and \mathbf{x}_{p2} are the normalized source and destination prototype shapes, respectively.

3 (a) Averaging a set of texture wavelet components produces a function with far smaller average magnitude. (b) Averaging the smoothed magnitudes of the wavelet components stores additional information about the texture. (c) We divide the mean wavelet components by their smoothed magnitude and multiply them by the sample's smoothed average magnitude to produce a texture function more representative of the sample. The horizontal axis indexes a range of 30 pixels in the horizontal wavelet decomposition at the second spatial scale, and the vertical axis shows the value of the wavelet coefficients.



Finally, we warp the subject's image into the new shape.

The transformation's second stage calculates the new facial colors. We also warp the source and destination prototype images into the new shape and calculate and scale (if desired) the color shift of each pixel and add it to the color of the subject's pixel

$$\mathbf{c}'_s = \mathbf{c}_s + (\mathbf{c}_{p2} - \mathbf{c}_{p1}) \quad (5)$$

where \mathbf{c}'_s is the (new) transformed image, \mathbf{c}_s is the warped subject's image, and \mathbf{c}_{p1} and \mathbf{c}_{p2} are the warped source and destination prototype images, respectively. This process resembles image morphing,⁸ but unlike morphing—which changes the identity of the subject during the morph—this method maintains the identity (or a family resemblance) across the transformation. For example, in Figure 2 we attempt to change the subject's apparent age from a younger adult (approximately 30 years old) to an older adult (approximately 60 years old) while maintaining identity.

Texture-enhanced facial prototyping

Experiments involving the rating of real and prototype faces reveal a difference between the perceived age of facial prototypes and the age of the component faces used to construct them.¹ The perceived age rat-

ings of male and female prototypes showed an age deficit compared to the original faces used to construct the prototypes. It's clear when looking at the prototype images of older faces that wrinkles aren't captured in the prototypes. The facial textures appear smoother than they do in the real faces. This is because the warping stage doesn't align the (undelineated) wrinkles precisely. Delineating facial wrinkles is difficult, if not impossible, due to the wide variability in their fine structure. Even automated warping methods (for example, as used in Blanz and Vetter⁹) impose a smoothness constraint.

To improve the textures of the facial prototypes, we propose a wavelet-based method (see Stollnitz et al.¹⁰ for an introduction on using wavelets in computer graphics). Heeger and Bergen,¹¹ for example, successfully used wavelet-based methods to synthesize new texture images from an example texture image. Also, wavelets play an important role in early visual processing in mammals,¹² which implies that wavelet use should give visually realistic results.

The aim here is to adjust the amplitude of edges (changes in pixel intensity) in the prototype image so that it matches the average edge amplitudes in the sample at different locations, orientations, and spatial scales. The edges in an image can be defined using a

Choice of Wavelet Analysis

The choice of wavelet analysis for a particular computer graphics application can be broken down into the choice of filter, orientations, and sampling. Here we explain how we selected the wavelet analysis used in this work.

Wavelet filter choice


In this work we decompose the image using a wavelet function (Figure A) that resembles the real part of a Gabor function (a Gaussian function modulated by a cosine function), which can be implemented efficiently using finite support filters. The attractiveness of Gabor filters for texture analysis and synthesis is partly because they're optimally localized in both the spatial and frequency domains, and also because they approximate some of the filters used in the early stages of mammalian visual processing.¹

Choice of orientations

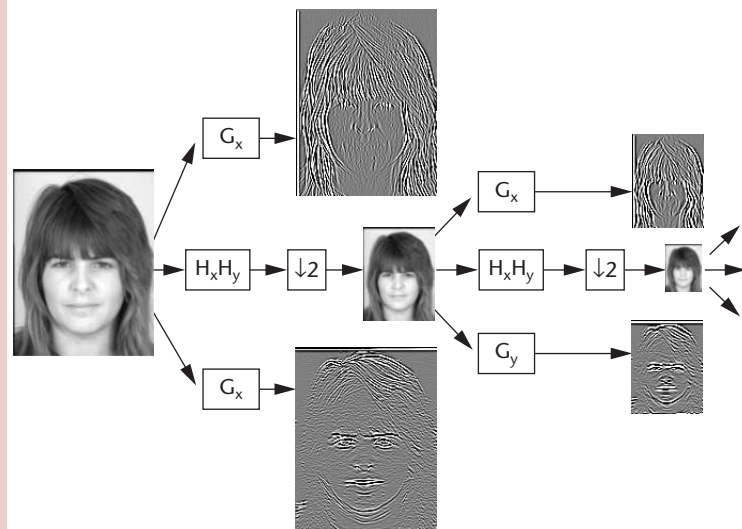
Using nonoriented wavelet pyramids (such as the Laplacian pyramid) presents a relatively simple wavelet decomposition. However, processing based on these wavelets doesn't let us amplify oriented edges by different amounts—for example, to amplify the horizontal crow's feet lines radiating out from the eye corners but not the circular line of the eye orbit. The human visual system uses many orientations separated by approximately 15 degrees, but the use of this many orientations adds considerably to the processing time and memory requirements. For efficiency, we use only two orientations at each scale.

Choice of sampling

Different samplings can affect the quality of reconstruction—even exact reconstruction algorithms can produce artifacts after processing. The wavelet pyramid must contain at least as much data as the original image (such as critical sampling) for exact reconstruction. We encountered star-shaped artifacts after processing when using critically sampled wavelet decompositions and found better results when using oversampled wavelet pyramids. In this article we subsample only the low-pass residual at



A We derived the low-pass filter (left) from H (a cubic B-spline function) and the wavelet filter (right) from G (an approximation to a real Gabor function) for the forward wavelet decomposition.



B A typical image (left) and the corresponding wavelet decomposition pyramid (right). The coefficients of the filters H and G are $(1, 4, 6, 4, 1)/16$ and $(1, -4, 6, -4, 1)/16$, respectively.

each scale (Figure B), which produces a pyramid $8/3$ times the original image's size.

Reference

1. J.G. Daugman, "Two-Dimensional Spectral Analysis of Cortical Receptive Field Profiles," *Vision Research*, vol. 20, 1980, pp 847-856.

variety of different edge-detecting wavelet filters, at different spatial scales and with different orientations. Biological considerations, the desire for computational efficiency, and avoiding visible reconstruction of artifacts after processing guides the choice of wavelets, orientations, and sampling (see the "Choice of Wavelet Analysis" sidebar for details).

Inspecting the original averaging process in the wavelet domain reveals the source of the prototyping texture problem (Figure 3a). If the edges of a set of functions aren't precisely aligned, their average is a function with much smaller edge amplitudes. To overcome this

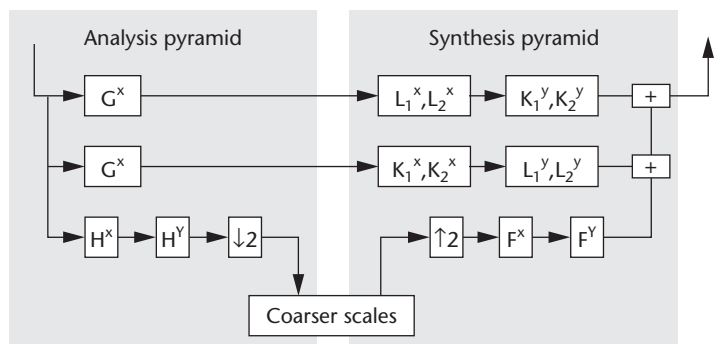
problem, we calculate a locally weighted measure of the edge strength σ at each point (x, y) in a particular wavelet subband w given by

$$\sigma_w(x, y) = h_x * h_y * |w(x, y)| \quad (6)$$

where h_x and h_y are the 1D cubic B-spline smoothing filters, used to filter (the $*$ operator) along x and y axes, respectively. In other words, we find the magnitude of the edges at each point and smooth the values to give a measure of the edge strength in a small region about that point in each wavelet subband.

Exact Reconstruction

Figure C shows the process of constructing and collapsing the wavelet pyramid. We first convolve



C Diagram of a single level of the forward and reverse wavelet decomposition in 2D.

Table A. Filter coefficients for the forward and inverse wavelet decomposition.

Index	0	±1	±2	±3	±4	±5
H(16)*	6	4	1			
G(16)	6	-4	1			
F(8)	10	5	-1	-1		
L ₁ (8)	22	8	1			
L ₂ (8)	48	27	8	1		
K ₁ (256)	186	36	4	-4	-1	
K ₂ (256)	168	34	16	-1	-4	-1

*The filters are symmetrical about 0 and the values in each row are divided by the factors shown in brackets.

the image with the 1D high-pass filter G in the x direction and keep it at full resolution. We also convolve the original signal with G in the y direction and keep it at full resolution. We then filter the original signal with the 1D low-pass filter H and subsample it by a factor of two in both the x and y directions.

To reconstruct an image from the wavelet pyramid, we first expand the low pass image by inserting zeros between pixels followed by convolution with the filter F in both x and y directions. The reduction and subsequent expansion of the low-pass filtered image is equivalent to a separable convolution of the original image with different filters at odd and even indexed pixels. Hence we require different reconstruction filters for odd and even pixels when filtering the high-pass components. We use 1D filter pairs of the form (F₁, F₂), in which we use the filter F₁ for odd indexed pixels in the output image and we use F₂ at even indexed pixels in the output image. We convolve the high-pass filtered images with the filter pair (L₁, L₂) in the same direction as the filter G and also with a second filter pair (K₁, K₂) in the orthogonal direction. Finally, we add the three filtered images to produce the reconstructed image at the next (finer) scale. Table A gives the coefficients of the decomposition filters, H and G, and the reconstruction filters, F, L₁, L₂, K₁, and K₂.

The mean values of σ across the set of N facial images

$$\bar{\sigma}_w(x, y) = \frac{1}{N} \sum_{i=0}^{N-1} \sigma_w^i(x, y) \tag{7}$$

gives a measure of the average edge size in each region of each wavelet subband (Figure 3b). Hence we can use these values to amplify the edges of the shape and color prototype \bar{w} to have more representative local values, such as

$$v(x, y) = \frac{\bar{w}(x, y) \bar{\sigma}(x, y)}{\sigma_{\bar{w}}(x, y)} \tag{8}$$

where $\sigma_{\bar{w}}$ is the smoothed magnitude of subband w of the shape and color prototype at pixel (x, y) and v is the wavelet transform of the resulting texture-enhanced prototype (Figure 3c). After processing the wavelet magnitudes, we use an exact reconstruction algorithm to collapse the wavelet pyramid (see the “Exact Reconstruction” sidebar for details). Figure 4 shows the results of the additional processing on the younger and older, male and female prototypes used in the percep-

tual experiments described in the “Experimental validation of age prototypes” section.

Texture-enhanced facial transformations

Our previous attempts to transform the apparent age of facial images have also had limited success,¹ partly because of the limitations in the prototypes. Age ratings of real and age-increased facial images showed that the original aging methods achieved only 50 percent of the target age shift. Age reduction wasn’t attempted previously because the shape and color transformation process didn’t remove textural cues to age (lines and wrinkles). We confirm in the following sections that an attempted 30-year age shift (from 55 to 25 years old) with the original age transformation method produces an average age reduction of only 2.6 years.

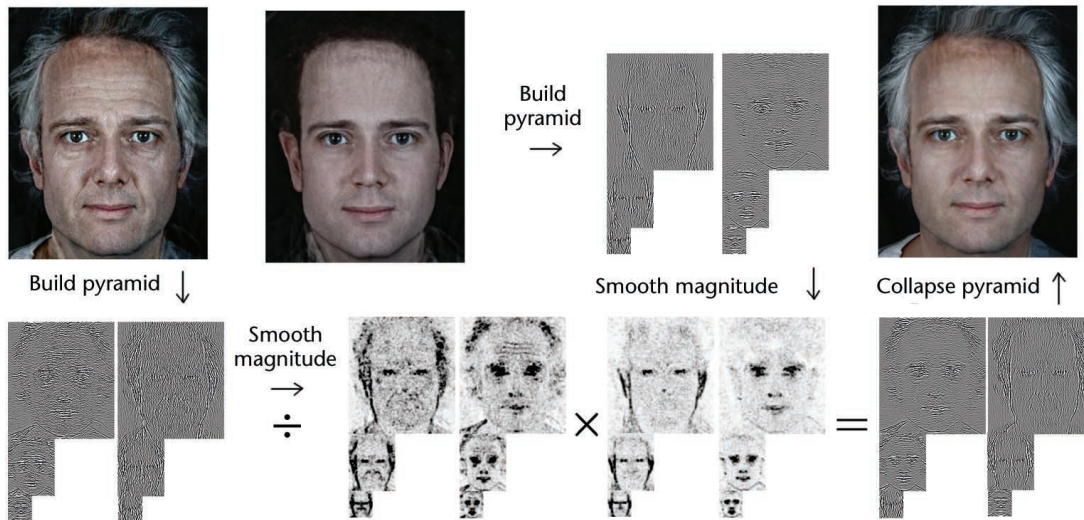
To improve aging and rejuvenating faces, we use the textured prototypes described previously to smooth or edge-enhance a subject’s facial image to an appropriate amount for the destination age group (Figure 5). Again, we’d like the smoothing or edge-enhancement to vary in different image regions, spatial scales, and edge orientations. After performing the shape and color transformation (equations 4 and 5), we calculate the wavelet transform w'_s of the transformed image c'_s and rescale the edges to have the same local magnitudes in each subband as the texture-processed destination prototype p_2 .



4 Example prototypes without (top) and with (bottom) the additional texture processing. From left to right: Males 18 to 45 years old, males over 45 years old, females 18 to 45 years old, and females over 45 years old.



5 The texture-enhanced transformation process. We transform the original image using the shape and color method (as in Figure 2) and warp the target prototype into the new shape. We then build wavelet pyramids from these two images and calculate their magnitudes. After rescaling, we collapse the subject's pyramid to give the new image.



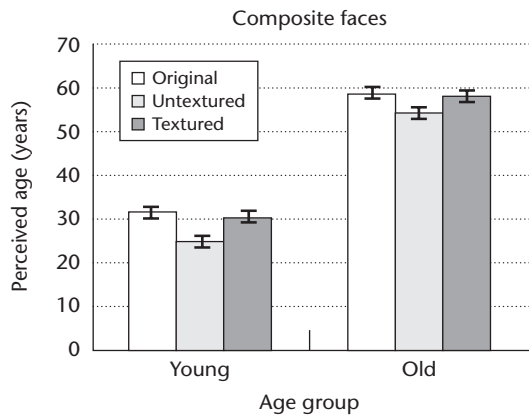
We can then independently process each wavelet sub-band w'_s of c'_s using

$$w''_s(x, y) = w'_s(x, y) \frac{\sigma_{p2}(x, y)}{\sigma'_s(x, y)} \quad (9)$$

where σ'_s and σ_{p2} are the smoothed magnitude of w'_s and w_{p2} . We find the new transformed image by collapsing the processed wavelet pyramid w''_s .

Experimental validation of age prototypes

We conducted two experiments to validate the wavelet processing methods' ability to capture age information. The first experiment was to validate the prototyping method and show that textured (shape-, color-, and wavelet-processed) prototypes capture the age of the sample faces more accurately than untextured (only shape- and color-processed) prototypes. We made



6 Average perceived age (± 1 standard error) of original face images, with textured and untextured prototypes. Subjects perceived the untextured prototypes as significantly younger than the average age of the original images in both groups (younger and older), but perceived the textured prototypes as having an age appropriate to the original face sample.

younger adult male and female prototypes from 23 male and 18 female faces (aged 18 to 45 years) and older adult prototypes from 15 male and 14 female faces (aged over 45 years). Figure 4 shows the untextured and textured male and female prototypes for the two different age groups.


We then presented all the face images (70 original faces, 4 untextured prototypes, and 4 textured prototypes) to 21 raters (10 males aged 19 to 55 years, 11 females aged 20 to 28 years) in random order, one image at a time, and asked them to estimate the age (in years) of each face. For each rater, we calculated the average perceived age of the 45 original younger adult face images, the 29 original older adult face images, and the textured and untextured male and female prototypes for the younger and older face groups. We compared the average perceived age of the younger and older faces to the average perceived age of each pair (male and female) of textured and untextured prototypes using Student's matched *t*-tests.

The results (Figure 6) show that the age ratings of the untextured prototypes are significantly younger than the average ratings of the original faces for both the younger and older face groups. The younger faces had an age difference of 6.6 years, where $t_{20} = 8.4$, and $p < 0.0005$. The older faces exhibited an age difference of 4.7 years, where $t_{20} = 3.4$, and $p < 0.0005$. The textured prototypes corrected this age deficit, which the raters didn't perceive as significantly different from the average original of the faces in either age group (for younger faces $t_{20} = 1.8$ and $p < 0.08$; for older faces $t_{20} = 0.7$ and $p < 0.52$). Thus, the additional wavelet prototyping accurately captures age information lost in the blending process.

Experimental validation of age transformations

The second experiment tested the effectiveness of the new age transformation technique. We transformed 10 female and 10 male faces evenly divided into the two age categories (18 to 45 and over 45 years old). We aged the 10 faces in the younger age group using the untextured method (shape and color transformation using the untextured pair of prototypes for the younger, 18 to 45 year-olds and older, over 45 year-olds) and with the new textured method (shape, color, and texture transformation using a textured destination prototype). We also used the same two methods to age reduce the 10 faces from the older age group. This

7 Transformation examples: original faces (left), without (center), and with (right) texture processing. (a) Male rejuvenation, (b) male aging, (c) female rejuvenation, and (d) female aging.



produced a total of 60 face images (20 of the original, 20 of the old method transformed, and 20 of the new method transformed). We presented each of the 60 face images individually in random order to 16 observers (6 males, aged 23 to 57, mean 40 years; 10 female, aged 22 to 38, mean age 27.5 years), who we asked to rate the age of the face to the nearest year.

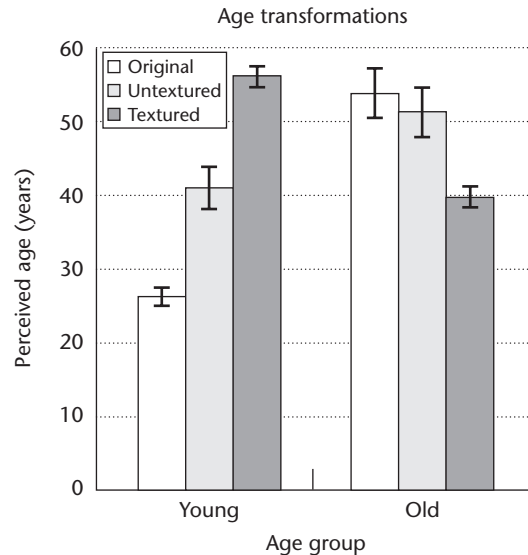
Figure 7 shows age transformations for male and female subjects. The texture-processed images appear to reflect the desired age change more accurately. The results of the age rating experiment in Figure 8 confirm this. The perceived age of the original younger adult faces is 25 years. Applying the color and shape aging transformation to this younger adult group of faces increases their apparent age by an average of 14.8 years. Transformation with the additional texture processing is twice as effective, aging the faces by an average of 30.0 years. The increase in age produced by transformation with texture processing (mean 30.0 years, standard error ± 1.3 years, Figure 8) corresponds to that expected from the age difference between the younger and older adult prototype faces (mean 27.9, standard error ± 1.4 years, Figure 6). Hence, the wavelet-enhanced age transformation projected the faces by the correct amount and reached the target age bracket.

Applying the shape and color transform to rejuvenate faces hardly affects the apparent age, achieving an average age reduction of only 2.6 years. Rejuvenation using the wavelet processing methods reduced the apparent age by 14 years. While texture-enhanced transformation is five times more effective, evidently improvements can be made. A possible explanation is that the average hair color change during rejuvenation isn't always sufficient to project the subject's hair color into the normal range for the younger age group. Any cue, such as slightly gray hair, may have a disproportionate influence on age perception.

Future work

We plan to use the new prototyping and transformation methods to investigate psychological theories of facial attraction related to aging. For example, like other animals humans probably learn certain key qualities from their parents, a process known as imprinting. Children may imprint on their parents' age, so we can test, for example, whether people born to older parents are attracted to older faces. We also hope to use the texture-enhanced age transformations to investigate how perceived personality attributes (such as perceived intelligence or perceived sociability) are affected by our stereotypic attitudes toward age.

We'll also investigate technical extensions to the texture processing algorithms. Our results show that previous statistical models of facial images in terms of shape and color are incomplete. The smoothed wavelet mag-



8 Average perceived age (± 1 standard error) of original and age-transformed face images. We age the younger adult faces (left) 15 years with the untextured (shape and color) prototype transformation and we age them 30 years with the texture-enhanced transformation. We rejuvenate the older faces (right) by less than three years using shape and color transformation and by more than 14 years using the additional texture processing.

nitude is an important additional facial image attribute, which is necessary for a complete statistical description of face images. We've started the statistical modeling of facial textures by constructing a representation of average wavelet magnitudes. In the future we hope to extend this model to include a description of the distribution of facial textures within a population of faces—for example, by employing Principal Component Analysis on the wavelet magnitudes. ■

Acknowledgments

Unilever Research funded this research through a grant. We'd like to thank all those who participated in the experiments presented here and who donated their face images to our collection.

References

1. D.M. Burt and D.I. Perrett, "Perception of Age in Adult Caucasian Male Faces: Computer Graphic Manipulation of Shape and Color Information," *Proc. Royal Soc. London*, series B, vol. 259, 1995, pp 137-143.
2. F.J. Galton, "Composite Portraits," *Nature*, vol. 18, 1878, pp. 97-100.
3. D. Ruprecht and H. Muller, "Image Warping with Scattered Data Interpolation," *IEEE Computer Graphics and Applications*, vol. 15, no. 2, Mar. 1995, pp. 37-43.
4. D.A. Rowland and D.I. Perrett, "Manipulating Facial Appearance through Shape and Color," *IEEE Computer Graphics and Applications*, vol. 15, no. 5, Sept. 1995, pp. 70-76.
5. A. Lanitis, C.J. Taylor, and T.F. Cootes, "Automatic Interpretation and Coding of Face Images Using Flexible Models," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 19, no. 7, July 1997, pp. 743-756.
6. B. Tiddeman, N. Duffy, and G. Rabey, "A General Method for Overlap Control in Image Warping," *Computers and Graphics*, vol. 25, no. 1, Feb. 2001, pp. 59-66.
7. B. Tiddeman and D. Perrett, "Moving Facial Image Transformations Based on Static 2D Prototypes," *Proc. 9th Int'l Conf. Central Europe on Computer Graphics, Visualization,*

PURPOSE The IEEE Computer Society is the world's largest association of computing professionals, and is the leading provider of technical information in the field.

MEMBERSHIP Members receive the monthly magazine **COMPUTER**, discounts, and opportunities to serve (all activities are led by volunteer members). Membership is open to all IEEE members, affiliate society members, and others interested in the computer field.



EXECUTIVE COMMITTEE

<p>President: BENJAMIN W. WAH * <i>University of Illinois Coordinated Sci Lab 1308 W Main St Urbana, IL 61801-2307 Phone: +1 217 333 3516 Fax: +1 217 244 7175 b.wab@computer.org</i></p>	<p>VP, Publications: RANGACHAR KASTURI † VP, Standards Activities: LOWELL G. JOHNSON* VP, Technical Activities: DEBORAH K. SCHERRER (2ND VP)*</p>
<p>President-Elect: WILLIS K. KING * Past President: GUYLAINE M. POLLOCK* VP, Educational Activities: CARL K. CHANG (1ST VP)* VP, Conferences and Tutorials: GERALD L. ENGEL * VP, Chapters Activities: JAMES H. CROSS †</p>	<p>Secretary: WOLFGANG K. GILOI* Treasurer: STEPHEN L. DIAMOND* <i>2000-2001 IEEE Division V Director: DORIS L. CARVER † 2001-2002 IEEE Division VIII Director: THOMAS W. WILLIAMS †</i> Acting Executive Director: ANNE MARIE KELLY †</p>

* voting member of the Board of Governors
† nonvoting member of the Board of Governors

BOARD OF GOVERNORS

Term Expiring 2001: Kenneth R. Anderson, Wolfgang K. Giloi, Haruhisa Ichikawa, Lowell G. Johnson, Ming T. Liu, David G. McKendry, Anneliese Amschler Andrews

Term Expiring 2002: Mark Grant, James D. Isaak, Gene F. Hoffnagle, Karl Reed, Kathleen M. Swigger, Ronald Waxman, Akibiko Yamada

Term Expiring 2003: Fiorenza C. Albert-Howard, Manfred Broj, Alan Clements, Richard A. Kemmerer, Susan A. Mengel, James W. Moore, Christina M. Schober

Next Board Meeting: 9 Nov 2001, Denver, CO

COMPUTER SOCIETY OFFICES

<p>Headquarters Office 730 Massachusetts Ave. NW Washington, DC 20036-1992 Phone: +1 202 371 0101 Fax: +1 202 728 9614 E-mail: hq.ofc@computer.org</p>	<p>European Office 13, Ave. de L'Aquilon B-1200 Brussels, Belgium Phone: +32 2 770 21 98 Fax: +32 2 770 85 05 E-mail: euro.ofc@computer.org</p>
<p>Publications Office 10662 Los Vaqueros Cir., PO Box 3014 Los Alamitos, CA 90720-1314 Phone: +1 714 821 8380 E-mail: help@computer.org Membership & Publication Orders: Phone: +1 800 272 6657 Fax: +1 714 821 4641 E-mail: help@computer.org</p>	<p>Asia/Pacific Office Watanabe Building 1-4-2 Minami-Aoyama, Minato-ku, Tokyo 107-0062, Japan Phone: +81 3 3408 3118 Fax: +81 3 3408 3553 E-mail: tokyo.ofc@computer.org</p>

EXECUTIVE STAFF

<p>Acting Executive Director : ANNE MARIE KELLY Publisher: ANGELA BURGESS Acting Director, Volunteer Services: MARY-KATE RADA</p>	<p>Chief Financial Officer: VIOLET S. DOAN Director, Information Technology & Services: ROBERT CARE Manager, Research & Planning: JOHN C. KEATON</p>
--	---

IEEE OFFICERS

President: JOEL B. SNYDER
President-Elect: RAYMOND D. FINDLAY
Executive Director: DANIEL J. SENESE
Secretary: HUGO M. FERNANDEZ VERSTAGEN
Treasurer: DALE C. CASTON
VP, Educational Activities: LYLE D. FEISEL
VP, Publications Activities: JAMES M. TIEN
VP, Regional Activities: ANTONIO BASTOS
VP, Standards Association: MARCO W. MIGLIARO
VP, Technical Activities: LEWIS M. TERMAN
President, IEEE-USA: NED R. SAUTHOFF



and *Computer Vision 2001* (WSCG 2001), Univ. of West Bohemia, Plzen, Czech Republic, 2001, pp. 260-266.

- T. Beier and S. Neely, "Feature Based Image Metamorphosis," *Computer Graphics* (Proc. Siggraph 91), ACM Press, New York, 1992, pp. 35-42.
- V. Blanz and T. Vetter, "A Morphable Model for the Synthesis of 3D Faces," *Computer Graphics* (Proc. Siggraph 99), ACM Press, New York, 1999, pp. 187-194.
- E. Stollnitz, T.D. DeRose, and D.H. Salesin, "Wavelets for Computer Graphics: A Primer (Parts 1 and 2)," *IEEE Computer Graphics and Applications*, vol. 15, no. 3, May 1995, pp. 76-84.
- D.J. Heeger and J.R. Bergen, "Pyramid-Based Texture Analysis/Synthesis," *Computer Graphics* (Proc. Siggraph 95), ACM Press, New York, 1995, pp. 229-238.
- J.G. Daugman, "Two-Dimensional Spectral Analysis of Cortical Receptive Field Profiles," *Vision Research*, vol. 20, 1980, pp. 847-856.



Bernard Tiddeman is a research fellow in the School of Psychology at the University of St. Andrews, Scotland. His research interests include developing new computer graphics methods for psychological and medical facial research. He received a BSc in applied mathematics from St. Andrews in 1992, an MSc in computing from Manchester in 1994, and a PhD in computing from Heriot-Watt in 1998.



Michael Burt is a doctoral student in the School of Psychology at the University of St. Andrews, where he uses computer graphics techniques to study facial perception, particularly facial attractiveness and individual differences in the perception of facial expressions. He received a BSc in psychology from the University of St. Andrews in 1994.



David Perrett is a professor in the School of Psychology at the University of St. Andrews. His research interests include understanding the brain mechanisms that underlie object recognition and facial perception. He received a BSc in psychology from the University of St. Andrews in 1976 and a DPhil in psychology from Oxford University in 1981.

Readers may contact the authors at Perception Laboratory, School of Psychology, University of St. Andrews, Scotland, UK, email {bpt, dmb, dp}@st-and.ac.uk.

For further information on this or any other computing topic, please visit our Digital Library at <http://computer.org/publications/dlib>.