A Non-generative Approach for Face Recognition Across Aging*

Soma Biswas†, Gaurav Aggarwal††, Narayanan Ramanathan† and Rama Chellappa†
Center for Automation Research, UMIACS
†Dept. of ECE, ††Dept. of Computer Science
University of Maryland, College Park
{soma, gaurav, rama}@cfar.umd.edu

***Abstract*— Human faces undergo a lot of change in appearance as they age. Though facial aging has been studied for decades, it is only recently that attempts have been made to address the problem from a computational point of view. Most of these early efforts follow a simulation approach in which matching is performed by synthesizing face images at the target age. Given the innumerable different ways in which a face can potentially age, the synthesized aged image may not be similar to the actual aged image. In this paper, we bypass the synthesis step and directly analyze the drifts of facial features with aging from a purely matching perspective. Our analysis is based on the observation that facial appearance changes in a coherent manner as people age. We provide measures to capture this coherency in feature drifts. Illustrations and experimental results show the efficacy of such an approach for matching faces across age progression.**

I. INTRODUCTION

Face recognition is one of the most successful applications of decades of research on image analysis and understanding [27]. Research in this area has traditionally focused on analyzing and modeling changes in facial appearance due to variations in illumination conditions, facial pose, expressions, etc. Other than these commonly occurring variations, aging is another phenomenon that affects facial appearance significantly. Though effects of aging on facial appearance have been studied for a long time, it is only recently that efforts have been made to recognize faces across age progression. Automatic matching of faces as people age is particularly useful for tasks like passport/visa renewal, identifying missing persons where authorities need to verify if the old and new photographs belong to the same person. Unlike other variates like illumination conditions and viewpoint, there is no simple geometric/statistical model to analyze appearance changes due to aging. Changes in facial appearance due to aging typically depend on quite a few factors like race, geographical location, eating habits, stress level, etc., that makes the problem of matching faces across aging extremely difficult.

Most existing works [10], [3], [26], [25], [11], [12], [21], [23], [24], [8], [22] on facial aging focus primarily on modeling and simulating aging effects on human face and report impressive simulation results. Though aging simulation has important graphics applications, this approach has certain limitations from the perspective of face recognition across aging. Given the infinite different ways in which a person can age depending on his/her surroundings, habits, etc., it is difficult to predict how a person will appear at a different age. Also, simulating face images at target age assumes that both the base and target age are known or can be estimated which by itself is a difficult problem. But in spite of this large variability, humans are quite good at matching faces across age progression. This may mean that irrespective of the exact manner in which a person ages, there is a certain pattern in the way facial appearance changes with age.

So for matching age-separated faces, an alternate approach would be to analyze whether the change in appearance between two faces can be attributed to the aging effects. Facial aging effects manifest in the form of facial shape change and change in skin texture like appearance of wrinkles, in addition to other intangibles. The relative effects of these factors also depend on the age being considered, eg., shape variations are more pronounced in children while textural changes are more prominent in adults [16]. In this paper, of all the different facial aging effects, we analyze the drift of facial features due to shape variation or sagging of underlying muscles and propose measures to capture the coherent patterns in these variations computationally. Thus given a pair of face images, matching is performed based on the coherency of feature drifts. If the two images belong to the same subject, we show that the drifts in features follow a coherent pattern which may not be the case if the images belong to different subjects. We propose measures to capture this feature drift coherency. Illustrations and experimental results show the efficacy of such a non-generative approach for matching faces across age progression.

A. *Organization of the paper*

The rest of paper is organized as follows: The following section discusses a few related works from the literature on facial aging simulation and age estimation. The proposed approach to measure coherency in facial feature drifts is detailed in Section III. Results of experiments performed to evaluate the efficacy of the approach are shown in Section IV. The paper concludes with a brief discussion and directions for future research in Section V.

II. RELATED WORK

Facial aging has been an area of interest for decades [19], [20], [17], [16], but it is only recently that efforts have

*Partially funded by the DISA ENCORE contract DCA200-02-D-5014 with Unisys Corporation serving as the prime on behalf of the Government.

been made to address problems like age estimation, age transformation, etc. from a computational point of view [10], [3], [26], [25], [11], [12], [21], [23], [24], [8], [22]. Burt and Perrett [3] investigate visual cues to age using facial composites that blend facial shape and color from multiple faces. Kwon and Lobo [10] classify input images as babies, young adults and senior adults based on cranio-facial development and skin wrinkle analysis. Wu *et al.* [26] describe a skin deformation model to simulate face wrinkles using an elastic process assembled with visco and plastic units. Tiddeman *et al.* [25] present a wavelet-based method for prototyping and transforming facial features to increase the perceived age of the input images. Lanitis *et al.* [12] use PCA-based transformation models to explain the effects of aging on facial appearance. The proposed statistical model is used for tasks like 1) age estimation from new face images, 2) simulating aging effects, and 3) face recognition across age variations. A similar PCA-based statistical face transformation model is used in [11] to obtain a compact parametric representation of an input face image for the task of automatic age estimation. Different classifiers are designed and compared that predict age given the parametric description of the input image.

In [21], Ramanathan and Chellappa study the effect of age progression on facial similarity between a pair of images of the same individual. A Bayesian age difference classifier is proposed to classify images based on age differences and perform face verification across age progression. In [22], they propose a craniofacial growth model to characterize growth related facial shape variations in children. This model makes use of anthropometric evidences to predict appearance across years and to perform face recognition using the synthesized images. Geng *et al.* [8] propose a subspace based approach for automatic age estimation. Given a previously unseen image, its aging pattern is determined by projecting it onto a subspace obtained using training data consisting of several time-separated images of individuals. Suo *et al.* [24] simulate aging effects using a dynamic Markov process on a multi-layer AND-OR graph integrating the effects of global appearance changes in hair style and shape, deformation and aging effects of different facial components, and wrinkle appearance. In [23], Scandrett *et al.* propose linear and piecewise models that rely on average developmental trends, to predict aging effects on human faces. In a recent paper, Ling *et al.* [14] use gradient orientation pyramid in a Support Vector Machine (SVM) based framework to verify images across age progression.

Most existing works address the problem of face verification across age progression from a simulation point of view (other than [14]). Given the difficulty of simulating effects of different factors that can affect the way a person ages, an alternative is to analyze if the difference in two input images can be attributed to aging. In this paper, we propose an approach to perform face verification across age progression based on the coherency of feature drifts in the input images.

III. Feature drifts

Aging brings about a wide variety of changes in the appearance of human faces. In children, this is mainly manifested in the form of significant shape changes due to growth of bones, muscle, etc. Though the overall shape does not change significantly for adults, there is subtle drift in facial features due to various factors like muscle sagging, weight gain or loss, etc. The features do not drift independent of each other. Depending on the underlying shape and muscle structure of the individual, there is some coherency among these drifts. We represent facial shape by a small set of points on the interior region of the face. Points on the facial contour are not considered since their positions change considerably with changes in pose. Fig. 1 illustrates such coherency in feature drifts for a few images from the FGnet aging database [1]. For images of different subjects, due to different shape and muscle structure, the feature drifts computed this way may not be coherent (Fig. 1). Such a coherency in feature drifts has been also been used in [13] for the task of face verification in videos. We try to formalize this observation and propose a measure for the coherency in the feature drifts with the goal of matching faces across age separation. The drift maps shown in Fig. 1 are obtained directly using the manually marked feature points available in the dataset.

Fig. 1. Drifts in facial features for a few age-separated face images from the FGnet aging database. The drifts across images of same individuals appear coherent (top two rows) while they are somewhat incoherent (third row) when the images belong to different individuals. (Best viewed in color)

If 3D range data of age-separated individuals are available, then they can also be analyzed directly to determine any coherency pattern due to feature drift/shape change. In the absence of a controlled 3D aging dataset, we recover the 3D shapes from the images using a Shape-from-Shading algorithm [2]. The shapes recovered from age-separated images of both same and different individuals are shown in

Fig. 2. We see that the recovered shapes reflect the changes that occur due to aging, for example, sagging of muscles under the eye regions, etc. Obviously, these effects will be more reliable and pronounced if the 3D data is used directly, as any shape from shading (SFS) algorithm used to obtain the shapes from the images tend to smooth the recovered shapes. The maps of the difference of the surface normals in the x, y and z directions show the changes that have occurred due to aging. We note that misregistration of age-separated faces may produce some changes that are not related to aging. The coherency in the difference maps when the input images are of the same individual is less visible in the second set of difference maps when the input images are of different individuals.

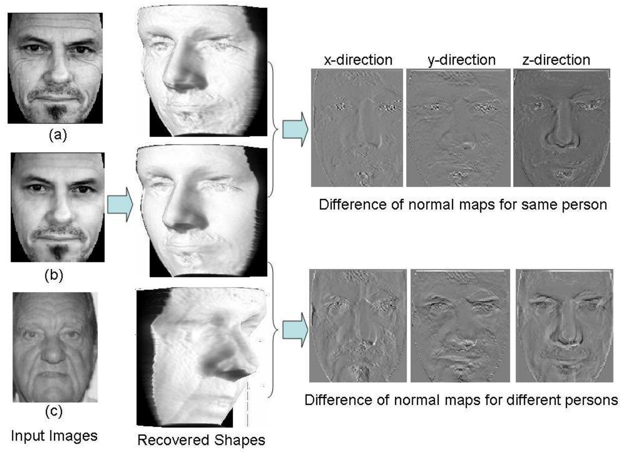

Fig. 2. Coherency in shape change with aging. (a) Input images; (b) Recovered 3D shapes of the corresponding input images; (c) Difference between the surface normal maps for the x, y and z directions.

A. Coherency of feature drifts

Now the question arises 'How can this coherency in the feature drift maps be captured computationally?'. We draw inspiration from the theory of electrostatics to compute this measure of coherency. Consider a system of K charged particles with charges denoted by q_i, $i = 1, 2, \cdots, K$. The potential energy between two charges q_i and q_j separated by distance r_{ij} is given by

$$U_{\mathrm{E}} = k_e \frac{q_i q_j}{r} \tag{1}$$

where k_e is known as the Coulomb's constant. The combined potential energy of the system of charges follows superposition principle and is given by the sum of the potential energies for each pair of charges

$$U_{\mathrm{E}} = k_e \left(\frac{q_1 q_2}{r_{12}} + \frac{q_1 q_3}{r_{13}} + \frac{q_2 q_3}{r_{23}} + \ldots \right) \tag{2}$$

From the principle of minimum total potential energy, the stable configuration or equilibrium state of a system is the one with minimum potential energy.

The drift maps we compute are treated as the sparse charge distributions. The potential energy of a system of charges is equivalent to the incoherency in the feature drift map. The intuition behind this analogy is that higher incoherency in the feature drifts implies that the configuration of feature drifts is not stable and so it is more likely that the two images do not come from the same individual. Similar to the electrostatic potential energy, the drift incoherency is inversely proportional to the distance between the two feature locations in consideration. This follows from the fact that neighboring drifts resulting from the sagging of the same underlying muscle will be coherent. In addition, this allows for far-off regions of the face with different underlying muscles to drift relatively independent of each other, without adding to the incoherency of the drift map. As potential energy increases with the magnitude of the charges, the incoherency of the feature drift map increases with the difference in the magnitude and direction of the drift vectors constituting the drift map.

Based on this analogy, we define a measure of incoherency between two feature drifts as

$$U_{\mathrm{ij}} = \frac{\| a_i - a_j \|}{r_{ij}} \tag{3}$$

Here $\| a_i - a_j \|$ is the magnitude of the vector difference between the two feature drifts a_i and a_j while r_{ij} is the distance between the corresponding feature locations. Following the superposition principle, the combined potential energy of the drift map characterized by K feature drifts is given by

$$C = \sum_{i=1}^{K} \sum_{j=i+1}^{K} U_{ij} \tag{4}$$

The feature drift maps thus obtained can be utilized to create a transformation map that interpolates the drifts at all points of the image using the available drifts. The transformation map can be used to normalize for the drifts for further analysis.

IV. Experimental evaluation

Fig. 3 shows a schematic of the proposed non-generative approach for face verification across age progression. We now present results of the experiments performed to evaluate the effectiveness of the proposed method.

A. Matching Age-separated Face Images in Children

Since we are dealing mainly with shape changes, we first evaluate our approach for matching face images of children. Aging effects in children are mainly manifested in the form of considerable shape changes as a result of growth of bones, loss of baby fat, etc. that makes the problem very challenging. For this analysis, we use 350 genuine pairs and randomly chosen 1100 impostor image pairs of children between the ages of $1-18$ from the FGnet dataset [1].

For computing the drift map, one needs to locate these facial features reliably to obtain the drift maps. In our implementation, we use SIFT features [15], which are extrema in scale space of the image. These features are highly distinctive and are invariant to image scale and have proved to be useful in several applications. Fig. 4 shows the SIFT features

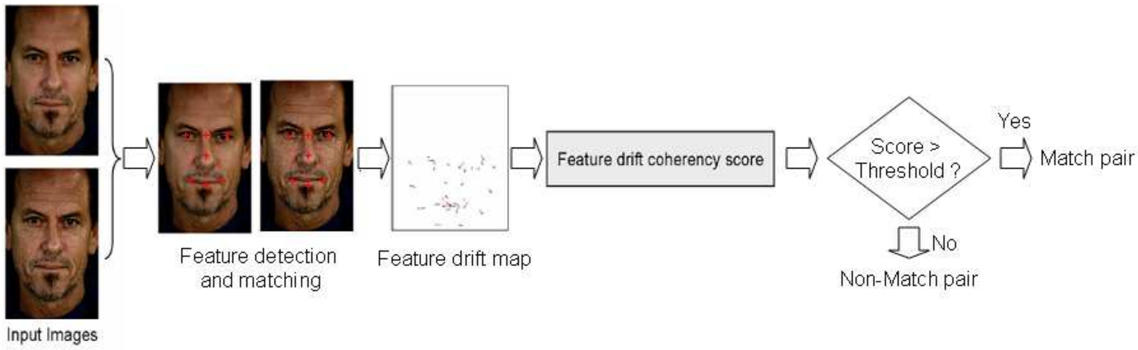

Fig. 3. Block diagram of the proposed method for matching face images across aging.

computed for a few images from the FGnet aging database. As seen in the figure, one can obtain quite a few distinctive facial features like the locations of the eyes, eyebrows, nose, mouth, etc. using this approach. Each feature computed this way is characterized by a 128-dimensional vector that represents the gradient distribution around the feature point. Feature correspondence across two images is computed using these vectors.

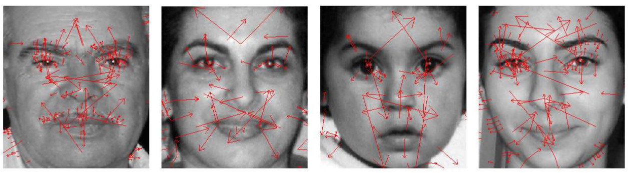

Fig. 4. SIFT features obtained for a few images from the FGnet dataset. (Best viewed in color)

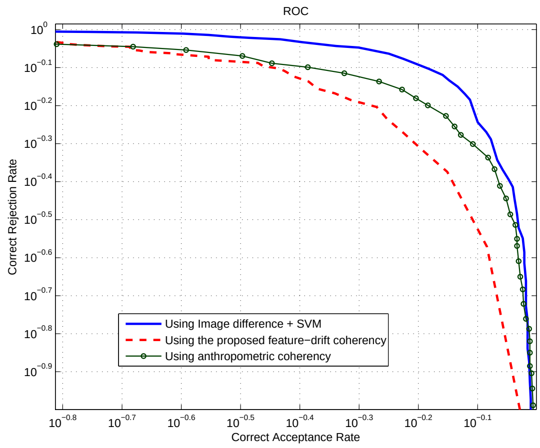

Fig. 5. Verification performance obtained using the proposed feature drifts coherency and anthropometric coherency on a challenging dataset consisting of age-separated images of children.

The overall incoherency of the drifts is measured using (4). The ten best correspondences provided by the SIFT algorithm are used to obtain the drift map. The absence of sufficient correspondences is taken care of by incorporating a penalty term. In our experiments, we add a fixed penalty term for each missing correspondence. The intuition is that fewer correspondences imply that the images are not similar and thus are more incoherent. Fig. 5 shows the Receiver Operator Characteristic (ROC) curve obtained in a verification experiment using this incoherency measure as the distance between the input pair of images. The plot compares correct rejection rate against correct acceptance rate. The correct rejection rate is the fraction of correctly rejected impostor pairs while the correct acceptance rate is the fraction of correctly accepted genuine pairs. Ideally, one would want to have both these quantities close to one simultaneously. As desired, despite being extremely simple (the incoherency measure depends on just 10 correspondences), the proposed measure is able to reasonably separate the genuine pairs from impostors.

Fig. 5 also comapares the performance of the proposed approach with **SVM+diff** [18] which uses differences of normalized images using an SVM classifier. For SVM+diff implementation, the images are first aligned with the help of the eye locations, cropped using an elliptic region and then resized to 80×70 for efficient computation and subsequent training using SVM classifier. Each image is normalized to have zero mean and unit variance before computing the image difference. SVM classifier with RBF kernel is used. The proposed method performs better than its counterpart probably because the feature drift analysis is able to account for large shape variations present in the input image pairs.

Global Anthropometric Coherency: As far as analyzing aging effects in young face images are concerned, there has been a recent interest in using facial anthropometric data [5] [9] [22]. Face anthropometric studies provide a quantitative description of the human craniofacial growth in terms of measurements taken between key landmarks on faces across different ages in children [6] [7]. These measurements indicate the typical growth patterns of human faces as they age. Anthropometric measurements have recently been successfully utilized to simulate realistic facial aging [9] [22].

In this work, for matching faces of children, we also analyze the usefulness of facial anthropometric data to evaluate the global consistency of the facial feature drifts with respect to the general growth patterns. Since we mainly deal with frontal face images, we use only those facial landmarks that can be reliably located using photogrammetry. In a synthesis framework, the face anthropometric data is used to constrain the facial growth model so that the simulated face at the target age is consistent with the data. In our case, the two input face images individually adhere to the anthropometric data. So the goal of matching two age-separated faces boils down to evaluating the degree to which the change from one face to the other deviates from the typical growth pattern(s). In our study, we use the following 11 facial measurements (Fig. 6): 1) Width of the face (zy-zy); 2) Height of the face (n-gn); 3) Height of the mandible (sto-gn); 4) Width of the mandible (go-go); 5) Intercanthal Width (en-en); 6) Biocular Width (ex-ex); 7) Eye fissure width (ps-pi); 8) Nose Width (al-al); 9) Nose height (n-sn); 10) Width of the mouth (ch-ch); 11) Height of the upper lip (sn-sto).

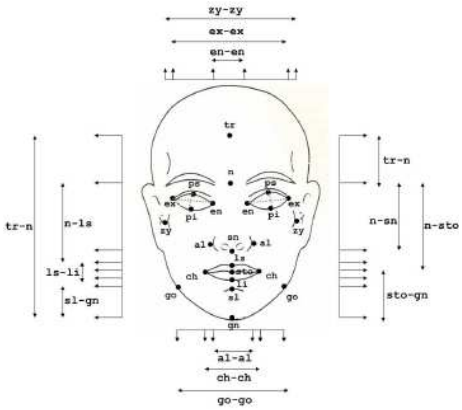

Fig. 6. Anthropometric locations [22].

From the available facial measurements at different ages, the average growth rate of different parts of the face are calculated. Let $d_1, d_2, \cdots, d_{11}$ denote the average growth rates for the 11 measurements we use in our task. The relative growth with respect to one of the measurements, say d_{11} is given by

$$F_i = \frac{d_i}{d_{11}}; \qquad i = 1, 2, \cdots, 10 \tag{5}$$

If the observed relative growth rate corresponding to the two input face images are denoted by $f_1, f_2, \cdots, f_{10}$, the global incoherency of the growth pattern is given as

$$C_{\text{global}} = \sum_{i=1}^{10} (F_i - f_i)^2 \tag{6}$$

Fig. 5 also shows the verification performance obtained using the global anthropometry based coherency. Interestingly, such a measure performs worse than the proposed feature-drift based coherency measure in the experiment. This may be due to the fact that the anthropometric studies provide typical growth patterns. Though these measurements are quite useful for simulating aging effects, they may be too generic to provide required discriminability as far as matching task is concerned.

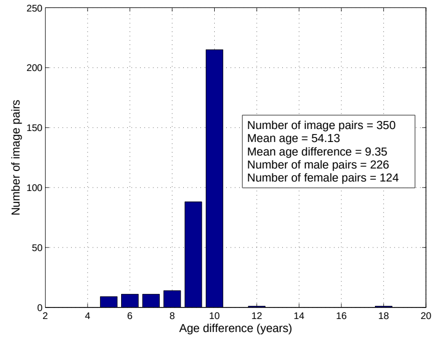

Fig. 7. Distribution of age difference between the pairs of images.

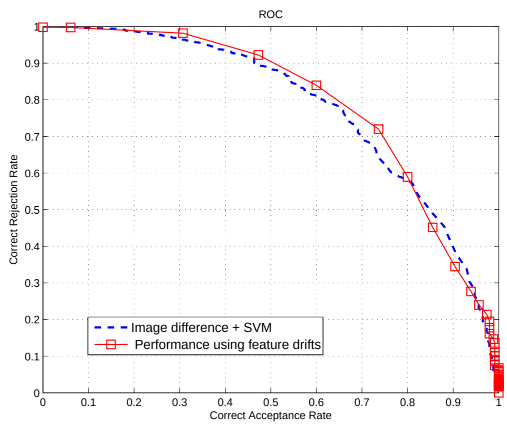

Fig. 8. Verification performance using incoherency of feature drifts.

B. *Matching Age-separated Face Images in Adults*

Shape change in adult faces is mainly manifested in the form of subtle drift of facial features due to various factors like muscle sagging, weight gain or loss, etc. In this section, we analyze the usefulness of the proposed feature-drift based coherency estimation to match age-separated face images of adults. For this purpose, we use a subset of a private Passport dataset which consists of age separated pairs of adult face images. The images are in general frontal with slight pose variations. All the images were aligned using the eye locations and resized to the same size. In our experiments, we use 350 genuine pairs of images with average age separation of over 9 years.

Fig. 7 provides the distribution of age separation for these pairs. The verification performance obtained using the proposed feature-drift coherency measure is shown in Fig. 8. For matching adult face images across aging, we see that the performance of our method is similar to that of SVM+diff. This is probably due to the fact that shape variations are not as prominent in adults as in children. The distributions of the genuine and impostor scores obtained are shown in Fig. 9.

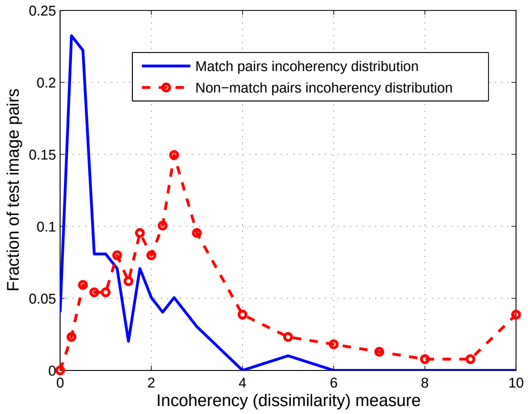

Fig. 9. Genuine and impostor incoherency score distributions obtained in a verification experiment using the proposed feature drift analysis.

V. Discussion and Future Work

The aging pattern of an individual depends on a variety of different factors that are difficult to model in a computational framework. But humans are quite good at matching faces across age progression. This may mean that irrespective of the exact manner in which a person ages, there is a coherency in the way facial appearance changes with age. This motivates us to capture and utilize this coherency to recognize age-separated faces. Specifically, we analyze the coherency of the drifts in various facial features to verify whether two age-separated images belong to the same individual or not.

Preliminary results presented in this paper show the effectiveness of such a non-generative approach even with simple measures of capturing coherency in aging. There are several possible extensions of the approach presented in this paper. We use SIFT features to match the facial features across two images in order to evaluate the feature drift coherency. Since SIFT features are not specific to human faces, it does not always locate all the facial features. Also appearance changes like growth of facial hair or weight gain/loss will affect the SIFT feature detection. For different image pairs, different number of features at different locations may be extracted. Though a suitable penalty term (as used in our experiments) can take into account such differences, more specific feature detectors like the ones based on Active Appearance Models (AAM) [4] may be more effective in capturing coherency of facial feature drifts. Moreover, since the drifts of features depend on the underlying facial muscle structure, this information may be used to obtain a better measure of drift coherency. Also, measures to capture textural variations with aging may be useful for matching age-separated images in adults.

References

[1] Fg-net aging database: Face and gesture recognition working group. http://www-prima.inrialpes.fr/FGnet/.

[2] S. Biswas, G. Aggarwal, and R. Chellappa. Robust estimation of albedo for illumination-invariant matching and shape recovery. In *Proceedings of IEEE Conference on Computer Vision*, 2007.

[3] D. Burt and D. Perrett. Perception of age in adult caucasian male faces: computer graphic manipulation of shape and colour information. *Proceedings of the Royal Society of London B*, B-259:137–143, 1995.

[4] T. F. Cootes, G. J. Edwards, and C. J. Taylor. Active appearance models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(6):681–685, 2001.

[5] D. DeCarlo, D. Metaxas, and M. Stone. An anthropometric face model using variational techniques. In *Proceedings SIGGRAPH*, pages 67–74, 1998.

[6] L. G. Farkas. *Anthropometry of the Head and Face*. Raven Press, New York, 1994.

[7] L. G. Farkas and I. R. Munro. *Anthropometric Facial Proportions in Medicine*. Charles C. Thomas, Springfield, Illnois, 1987.

[8] X. Geng, Z. Zhou, and K. Smith-Miles. Automatic age estimation based on facial aging patterns. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(12):2234–2240, December 2007.

[9] K. Kahler. A head model with anatomical structure for facial modeling and animation. Master's thesis, der Universitat des Saarlendes, 2003.

[10] Y. H. Kwon and N. da Vitoria Lobo. Age classification from facial images. *Computer Vision and Image Understanding*, 74(1):1–21, 1999.

[11] A. Lanitis, C. Draganova, and C. Christodoulou. Comparing different classifiers for automatic age estimation. *IEEE Transactions on Systems, Man, and Cybernetics, Part B*, 34(1):621–628, February 2004.

[12] A. Lanitis, C. Taylor, and T. Cootes. Toward automatic simulation of aging effects on face images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(4):442–455, April 2002.

[13] B. Li and R. Chellappa. Face verification through tracking facial features. *Journal of the Optical Society of America A*, 18:2969–2981, December 2001.

[14] H. Ling, S. Soatto, N. Ramanathan, and D. Jacobs. A study of face recognition as people age. In *Proceedings of IEEE International Conference on Computer Vision*, 2007.

[15] D. G. Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2):91–110, November 2004.

[16] L. S. Mark, J. B. Pittenger, H. Hines, C. Carello, R. E. Shaw, and J. T. Todd. Wrinkling and head shape as coordinated sources of age level information. *Journal of Perception and Psychophysics*, 27(2):117–124, 1980.

[17] L. S. Mark and J. T. Todd. The perception of growth in three dimensions. *Journal of Perception and Psychophysics*, 33(2):193–196, 1983.

[18] P. J. Phillips. Support vector machines applied to face recognition. In *Proceedings of the Conference on Advances in Neural Information Processing Systems II*, pages 803–809, 1999.

[19] J. B. Pittenger and R. E. Shaw. Aging faces as viscal-elastic events: Implications for a theory of nonrigid shape perception. *Journal of Experimental Psychology: Human Perception and Performance*, 1(4):374–382, 1975.

[20] J. B. Pittenger, R. E. Shaw, and L. S. Mark. Perceptual information for the age level of faces as a higher order invariant of growth. *Journal of Experimental Psychology: Human Perception and Performance*, 5(3):478–493, 1979.

[21] N. Ramanathan and R. Chellappa. Face verification across age progression. *IEEE Transactions on Image Processing*, 15(11):3349–3361, November 2006.

[22] N. Ramanathan and R. Chellappa. Modeling age progression in young faces. In *Proceedings of IEEE International Conference on Computer Vision and Pattern Recognition*, pages 387–394, 2006.

[23] C. Scandrett, C. Solomon, and S. Gibson. A person-specific, rigorous aging model of the human face. *Pattern Recognition Letters*, 27(15):1776–1787, November 2006.

[24] J. Suo, F. Min, S. Zhu, S. Shan, and X. Chen. A multi-resolution dynamic model for face aging simulation. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2007.

[25] B. Tiddeman, M. Burt, and D. Perrett. Prototyping and transforming facial textures for perception research. *IEEE Computer Graphics and Applications*, 21(5):42–50, 2001.

[26] Y. Wu, N. Thalmann, and D. Thalmann. A dynamic wrinkle model in facial animation and skin aging. *Journal of Visualization and Computer Animation*, 6:195–205, 1995.

[27] W. Zhao, R. Chellappa, P. J. Phillips, and A. Rosenfeld. Face recognition: A literature survey. *ACM Computing Surveys*, 35(4):399–458, 2003.