EUROGRAPHICS 2007 / D. Cohen-Or and P. Slavík
(Guest Editors)

Volume 26 (2007), Number 3

Prediction of Individual Non-Linear Aging Trajectories of Faces

K. Scherbaum[1], M. Sunkel[1], H.-P. Seidel[1] and V. Blanz[2]

[1] MPI Informatik, Saarbrücken, Germany
[2] Universität Siegen, Germany

Abstract

Represented in a Morphable Model, 3D faces follow curved trajectories in face space as they age. We present a novel algorithm that computes the individual aging trajectories for given faces, based on a non-linear function that assigns an age to each face vector. This function is learned from a database of 3D scans of teenagers and adults using support vector regression.

To apply the aging prediction to images of faces, we reconstruct a 3D model from the input image, apply the aging transformation on both shape and texture, and then render the face back into the same image or into images of other individuals at the appropriate ages, for example images of older children. Among other applications, our system can help to find missing children.

Categories and Subject Descriptors (according to ACM CCS): I.3.6 [Computer Graphics]: Methodology and Techniques–Interaction techniques I.4.10 [Image Processing and Computer Vision]: Image Representation–Hierarchical, Multidimensional, Statistical J.m [Computer Applications]: Miscellaneous–Forensic Sciences

1. Introduction

Police investigators who search for children that have been missing for several years have to predict the children's current looks from images taken at an earlier age. Today, much of this work is done by forensic artists, based on their experience and artistic skill. In order to simplify and automate this task, we present a method that learns from a large dataset of 3D scans of faces how children grow, and applies this transformation to 3D faces and to images.

In our approach, individual faces are represented as face vectors in a 3D Morphable Model of faces [BV99]. Over the years, each face will transform along a curved trajectory in this high-dimensional space. Ideally, we would like to measure these trajectories in a longitudinal study with a dense set of time samples of a number of individuals. We could then transfer the aging trajectories to new individuals. However, such data are difficult to collect, so we are facing a significantly more difficult problem: *Given a database that contains one single scan for each individual, plus the age of each person, predict the effect of aging on novel faces.* This involves two challenges: (1) Learn how an individual face would change over time (non-linear dependency on time), and (2) learn how the change depends on the individual face of the person (non-linear dependency on the position in face space).

To solve this learning problem, we proceed in two steps: we learn the function that assigns an age value to each face vector, and compute individual aging curves to obtain new face vectors at given ages.

This strategy has been proposed in an entirely linear approach [BV99], where a linear regression has been applied to describe facial attributes, such as gender or body weight, from individual, annotated scans. Then, the gradient of this function was used to change the attributes. The rationale behind this was that the gradient defines the shortest path, i.e. the minimal change necessary to obtain the desired attribute value. It has been shown that the computation of the gradient depends critically on the scalar product [BV99, BAHS06], and that PCA-based Mahalanobis-distance is more appropriate than a simple L_2 norm in shape or texture space.

If a linear age function is used, the constant gradient will shift all faces along the same, straight trajectory as they age. In contrast, the algorithm described here is based on a non-

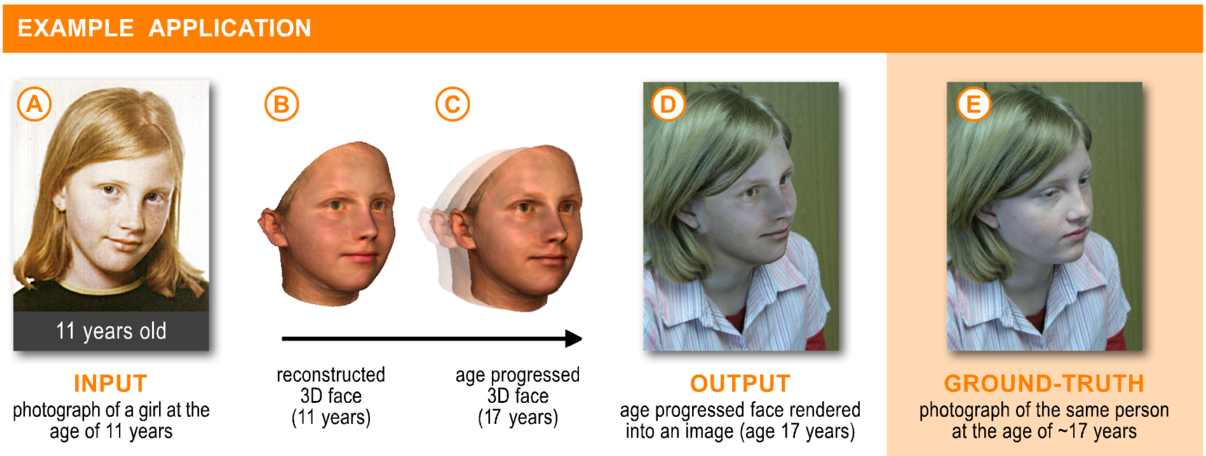

Figure 1: *The example shows a picture of an 11 year old girl (A). We reconstructed a 3D model of her face by fitting the Morphable Model to her face (B). Our growth algorithm then transformed the 3D face into a face at an appropriate target age (here: 17 years) (C). Finally we rendered the age progressed 3D face into an arbitrary background image (here: the ground-truth image was chosen as background) (D). For comparison, (E) shows a real picture of the girl at the target age.*

linear age function. Aging curves are computed by following the gradient of this function, which involves solving a differential equation using Runge-Kutta integration. As a result, aging trajectories depend in a non-linear way both on the age and on the individual face.

We embed our 3D aging transformation in a Morphable Model framework that includes 3D shape reconstruction from images, using an analysis-by-synthesis approach [BV99], and a method for inserting faces into existing images [BSVS04]. These two elements are necessary for the typical setting in forensic applications: The image material of missing children is usually only a set of snapshots at random imaging conditions, so the algorithm has to be very robust in terms of input data. On the other hand, it is useful to be able to render predicted faces into pictures of other children at the current age of the missing child to maintain consistent age of the face and the rest of the body. Moreover, the system that we propose can produce images with different hairstyles.

In summary, the contributions of this paper are:

1. Non-linear aging curves, which capture the different phases of growth and aging of facial tissue,
2. Individual aging effects, which may distinguish, for example, between the aging of obese and skinny faces,
3. A general approach for both 3D models and 2D images, and
4. The method is applicable to images at any given pose and illumination. This is a crucial feature for most real-world applications, such as police work where only a limited set of snapshots of the missing children's faces are available.

2. Related Work

In the image domain, the first algorithms for aging faces were developed in the early 80s [TMSP80, BS81]. In the work of Burt, Rowland and Perrett [BP95, RP95], database images are warped for 2D registration, based on the locations of feature points. Then, average faces of two different age groups are computed, and the difference warp field and color change is applied to novel faces. In 2001 this method was extended such that the textures of faces are represented in a wavelet pyramid. By manipulating their wavelet magnitudes locally, one may enhance edges and wrinkles and map them to new faces to simulate effects of aging. [TBP01]

Scandrett et al. present a semi-automatic 2D approach where a PCA is applied to a set of photographs of faces. Differences in pose up to 30°are compensated and average aging trajectories are computed based on statistical measures. Two different age progression algorithms are proposed: a linear and a piecewise method. The piecewise algorithm may also involve average developmental trends and consider familial correlations, if available. In contrast to our work some of the statistics rely on information gained from several pictures of one individual at previous ages. [SSG06b]

In a 2D Active Appearance Model and a dataset of age progressive images of 45 individuals, Lanitis et al. [LTC02] explored linear, quadratic and cubic age functions. Intermediate ages of these faces were generated by computing random face vectors in the region of interest, estimating their age and averaging groups that have the same age. For simulating the aging of novel faces, they searched for the nearest neighbour in the database and used the aging curve of that face, or they formed a weighted sum of curves of similar faces. The system is restricted to frontal views.

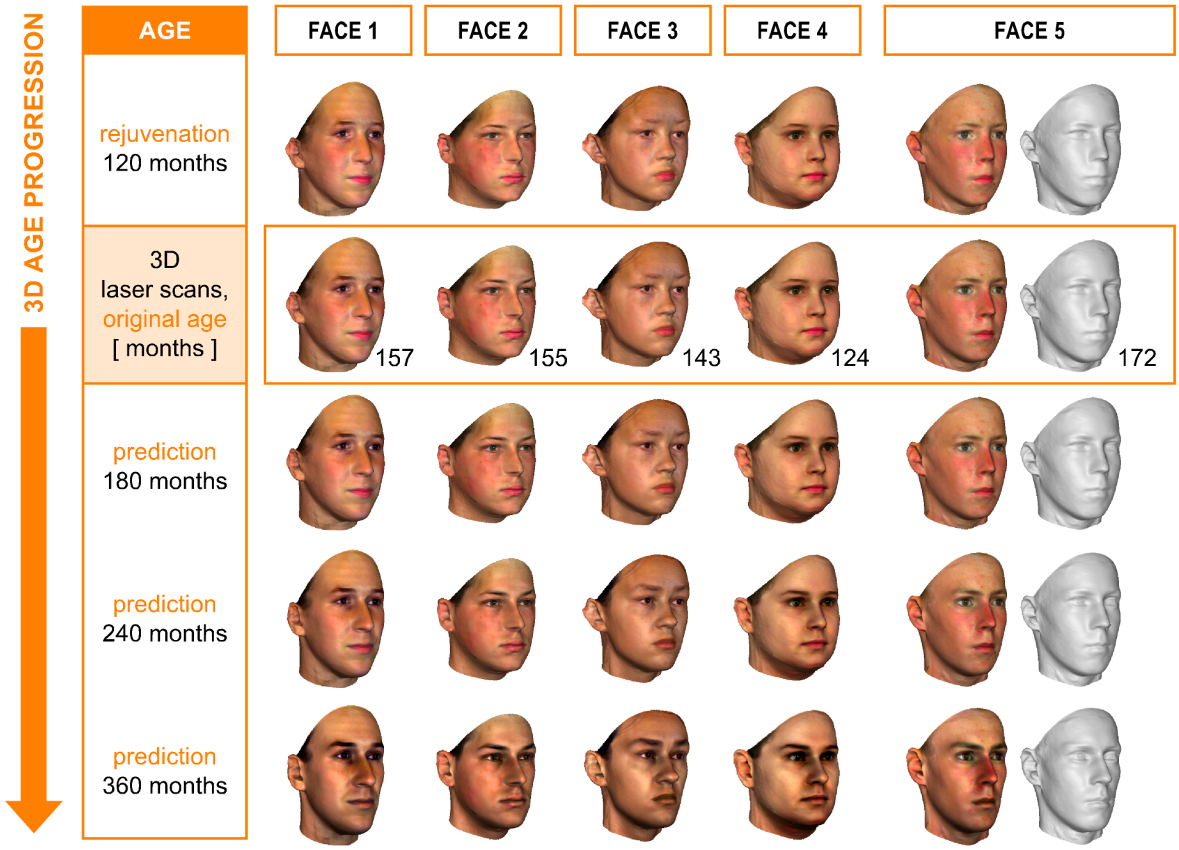

Figure 2: *The figure shows different example faces of our database (from left to the right). Each face is transformed along its individual non-linear aging trajectory. Individual characteristics of the faces are retained by the age transformation.*

For three-dimensional faces, learning facial attributes from 3D scans of faces and attribute values, such as gender or body weight, has been achieved with a linear method in the context of 3D Morphable Face Models [BV99]. Recently, several new PCA-based methods have been presented [SSG06a, HSG04]. Wang et al. proposed a non-linear multi-resolution method to extract and transfer subtle expression details of individuals [WHL*04]. Discriminating characteristics of a person's expression style may be used to synthesize new expressions. In a different approach, the growth of faces was predicted from anthropometric measurements of a sparse set of landmark points [KHYS02, Ram06]. For older faces, wrinkles are an important age cue that has been controlled by a number of techniques [LWMT99, WKMMT99, KHYS02].

On a 3D model, Hutton proposed a non-linear aging trajectory that was computed by fitting a single smooth curve to the faces at different ages [Hut04, HBHP03]. Each point of this curve is a weighted sum of the input faces (kernel smoothing), and it reflects the aging trajectory of an average face. The same curve is then applied to all individual faces. Hutton also performed support vector regression for predicting the age of novel faces. In our paper, we go beyond that by presenting an algorithm that employs support vector regression for the simulation of aging. For age estimation, a number of other methods have been proposed recently, which are not closely related to our Support Vector Regression approach [LDC04, GZZ*06].

For capturing individual differences in facial expressions, multilinear models use a tensor representation that provides separate model parameters for identity and expression [VT02, VBPP05]. Given a new face, the multilinear model can predict the person's facial expressions. For aging, however, a straight-forward transfer of the multilinear approach would not be able to produce curved aging trajectories, so the same change of facial features would be applied over the entire interval of ages that are studied. In contrast, our approach captures individual differences and non-linear trajectories at the same time.

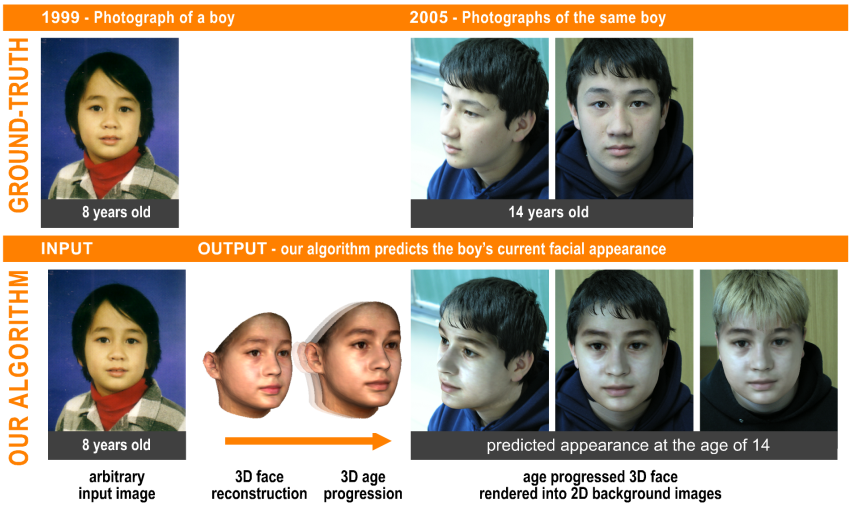

Figure 3: *Processing steps to generate age progressed images of a face: reconstructing the 3D shape and texture from a single image of a child leads to a 3D face model. By following its individual age trajectory in face space, we transform the face to an older age (bottom row). Rendering the results into different images of the same person in its actual age allows us to compare the result with the actual appearance of the face (ground-truth, top row).*

Beyond a certain age, aging of faces involves the increase of high spatial frequency structures such as wrinkles. In a recent study based on high-resolution scans, a local statistical model simulates aging of adult faces by adapting the statistics of the face to the target age [GMP*06]. In our paper, we restrict ourselves to lower spatial frequency effects and to facial growth of children and young adults.

3. A Morphable Model of Faces

The Morphable Model of 3D faces [VP97, BV99] provides a powerful representation for faces, but it is also used for model-based registration and for 3D shape reconstruction in our system. In the Morphable Model, faces are represented by shape vectors $\mathbf{S}$ and texture vectors $\mathbf{T}$ such that each linear combination of different faces is a new, realistic face

$$\mathbf{S} = \sum_{i=1}^{m} a_i \mathbf{S}_i, \quad \mathbf{T} = \sum_{i=1}^{m} b_i \mathbf{T}_i \tag{1}$$

within a few standard deviations from the average. The components of $\mathbf{S}$ and $\mathbf{T}$ are the 3D coordinates and RGB texture values of the vertices of a polygon mesh,

$$\mathbf{S} = (x_1, y_1, z_1, x_2, \dots, x_n, y_n, z_n)^T \tag{2}$$

$$\mathbf{T} = (R_1, G_1, B_1, R_2, \dots, R_n, G_n, B_n)^T. \tag{3}$$

In our model, $n = 75972$. To make sure that the linear combinations do not blend features of faces on different locations of the surface, it is essential to establish point-to-point correspondence between all vectors, so a given vector component describes the same point of the face in each face vector. In an initial model of 200 adult faces, this correspondence has been computed automatically with an optical flow algorithm [BV99]. The multiple, incomplete surface scans of each face described in this paper required a more robust method which is described in Section 6.

In order to reduce the dimensionality of the subsequent learning problem, we perform a Principal Component Analysis PCA on the shape and texture vectors. This defines an orthogonal set of basis vectors $\mathbf{s}_i$, $\mathbf{t}_i$, and with the average shape $\bar{\mathbf{s}}$ and texture $\bar{\mathbf{t}}$, we can define shape and texture coefficients c_i^s, c_i^t such that

$$\mathbf{S} = \bar{\mathbf{s}} + \sum_{i=1}^{m} c_i^s \cdot \mathbf{s}_i, \quad \mathbf{T} = \bar{\mathbf{t}} + \sum_{i=1}^{m} c_i^t \cdot \mathbf{t}_i. \tag{4}$$

For the simulation of aging in teenagers, we relied on two datasets. One is a set of 200 adult faces recorded with a Cyberware PS3030 scanner [BV99]. The other is a set of 238 scans of teenager faces that we acquired with a new setup and processing pipeline that we describe in Section 6.

4. Predicting the Age of Faces

Before we can actively change the age of 3D faces with our algorithm, we have to learn the function $f : \mathbb{R}^k \rightarrow \mathbb{R}$ that maps any face $\mathbf{x}$ to a scalar age value. In the following, let $\mathbf{x}$ be either shape coefficients $(c_1^s, ..., c_k^s)^T$, $k \leq m$, or texture coefficients $(c_1^t, ..., c_k^t)^T$. To reduce the computational complexity of the learning problem, we do not use all principal components, but experiment with values $k = 20, 40, 80$, which can be justified by the fact that most changes of the

	Shape			Texture			
	20 PCs	40 PCs	80 PCs	20 PCs	40 PCs	80 PCs	
$RBF-SVR$	38.90	36.33	33.26	29.35	28.24	30.66	msE_{cross}
$linear$	67.66	67.53	64.55	62.87	58.53	55.36	msE_{cross}
$RBF-SVR$	32.68	14.90	11.83	18.12	12.15	2.28	msE_{train}
$linear$	66.14	58.30	45.43	60.05	50.85	39.59	msE_{train}

Table 1: *Mean squared generalization- and training errors (in months) on a data set of 393 faces aged between 95 and 360 months. Regression was performed on 20,40 and 80 principal components, respectively. The generalization errors were obtained by cross-validation with 39 randomly chosen test samples. The training error resulted from training and testing with the complete data set.*

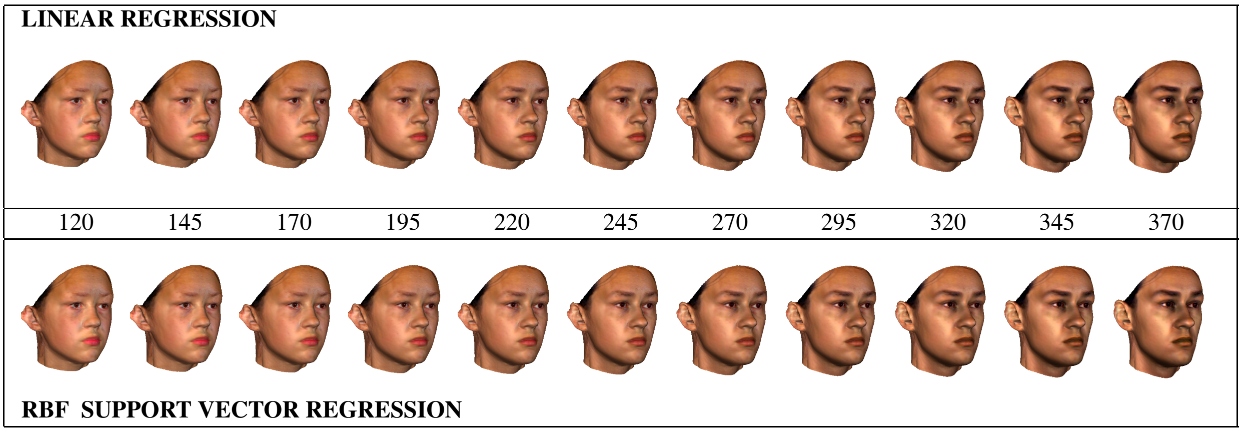

Figure 4: *Linear and non-linear age progression for 80 principal components. We show the comparison for face 3 of Figure 2 in the range of 120 to 370 months. The subtle differences between the two curves are difficult to see. On the original face, the age estimation error was less than one month, while the linear estimation was 18 months off in shape, and 9 months in texture.*

overall face shape are found in the first principal components.

To learn the age function, we use non-linear Support Vector Regression [Vap95, SS02] on training sets of l pairs $(\mathbf{x}_i, y_i)$, $i = 1, ..., l$ with y_i denoting the age of each example face i. As kernel functions, we use RBF kernels and polynomials. On our data, the regression results are consistently better for RBF than for polynomials, so we will only refer to RBF kernels in the rest of this section. The support vector regression function for RBF kernels is given by

$$f(\mathbf{x}) = \sum_{i=1}^{l} \alpha_i \cdot y_i \cdot e^{-\gamma \cdot \|\mathbf{x}_i - \mathbf{x}\|^2} + b \qquad (5)$$

α_i and b are real numbered values that are determined by support vector training.

We use the LIBSVM implementation for ε-Support Vector Regression [CL01], and perform grid search using cross-validation on the training sets to find the optimal parameters, such as the width γ of the RBF kernel.

We trained and evaluated the regression functions separately on shape and on texture coefficients, using 20, 40 and 80 principal components. The evaluation is done on the combined dataset of teenagers and adults, but without any faces older than 30 years, because our data set contains only a few sparse examples in that age, and errors on these faces would mask the effects in the age period that we are interested in. The dataset, therefore, contains 393 persons aged between 95 and 360 months. For comparison, we also perform linear regression with a straight-forward least squares fit. Table 1 shows training errors on the full dataset, and generalization errors obtained by cross validation. In cross validation, we split the dataset in 11 different random ways into 90% training and 10% test faces to measure the performance on previously unseen faces.

Both in terms of training error and in generalization per-

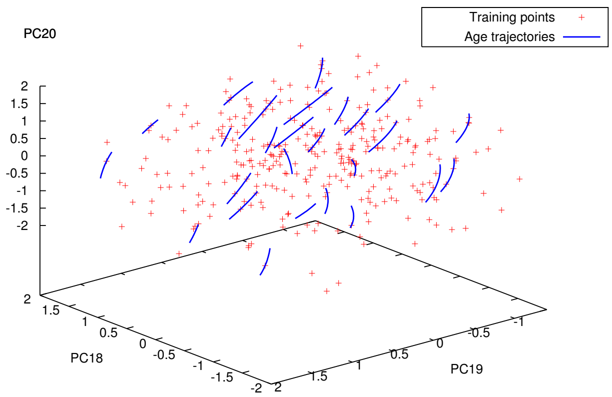

Figure 5: *Projected into the space of the last 3 out of 20 principal components, the curvature of the age trajectories (blue curves) is clearly visible.*

formance, the Support Vector function is superior to the linear function. The training error shows that the non-linear function can be adapted to the given data more closely than the linear function, and the generalization shows that is does not suffer from overfitting, but produces an appropriate estimate of the true age function.

The results indicate that there is a significant non-linear component in the age function, and therefore a simple linear transformation of faces in aging simulation will not capture all of the aging effects found in the database. We use the superior performance of the Support Vector regression as a motivation to rely on this function for the computation of aging trajectories.

5. Aging Trajectories

The main idea of the novel algorithm that we present in this paper is that the aging trajectory $\mathbf{z}(t)$ should be parallel to the gradient of the age function $f(\mathbf{x})$ at each moment in time, and that it should pass through the starting vector $\mathbf{x}_0$, so we are looking for a function $\mathbf{z}: I\!R \to I\!R^k$ such that at any age t

$$\frac{d\mathbf{z}}{dt}(t) = \nabla f(\mathbf{z}(t)), \tag{6}$$

with the starting condition

$$\mathbf{z}(t_0) = \mathbf{x}_0, \tag{7}$$

where $\mathbf{x}_0$ is the starting face at an age t_0. This defines a family of individual trajectories $\mathbf{z}_{\mathbf{x}_0,t_0}(t)$. The motivation of (6) is that for all $\mathbf{x}$ and t the gradient of f describes the direction of minimal change in $\mathbf{x}$ to achieve a given change in age t, so the characteristic features of the face are retained in the best way possible. This criterion has been applied successfully with linear functions f and attributes such as gender and body weight [BV99, BAHS06].

In order to compute $\mathbf{z}_{\mathbf{x}_0,t_0}(t)$, we have to integrate the differential equation (6). We do this by a fourth order Runge-Kutta algorithm [PTVF92], which is the most widely used integration algorithm. Runge-Kutta integration performs small steps along the gradient ∇f in an interleaved manner. ∇f can be computed from the regression function (5):

$$\nabla f(\mathbf{x}) = 2\cdot\gamma\cdot\sum_{i=1}^{l}\alpha_i\cdot y_i\cdot e^{-\gamma\cdot\|\mathbf{x}_i-\mathbf{x}\|^2}\cdot(\mathbf{x}_i-\mathbf{x}) \tag{8}$$

Figure 5 visualizes some of the simulated age trajectories (blue) in shape space. We selected a set of principal components (coordinate axes in Figure 5) that shows the curvature of the trajectories. In the first principal components, the trajectories look almost straight.

Given a face $\mathbf{x}_0$, we can now compute the face vector for a target age t_{target}. Let $t_{est} = f(\mathbf{x}_0)$ be the estimated age of the starting face. Then we have to simulate the trajectory along a time period $\Delta t = t_{target} - t_{est}$ to find the predicted face $\mathbf{z}_{\mathbf{x}_0}(t_{target})$. In some cases, the true age t_0 of the starting face may be known, and in general, this will be slightly different from t_{est}. To resolve this conflict, we simulate the intended time period $\Delta t = t_{target} - t_0$ and compute $\mathbf{z}_{\mathbf{x}_0}(t_{est} + \Delta t)$. The justification for this is that the starting face may have looked younger or older than it really was, and we may assume a continued aging process rather than stagnation, which could be the result if we would insist on the correctness of f and output $\mathbf{z}_{\mathbf{x}_0}(t_{target})$.

In order to investigate the non-linear components in $\mathbf{z}_{\mathbf{x}_0,t_0}(t)$, consider the angles between gradients of f in different positions in face space (Table 2). The first row lists the mean angles $\measuredangle\left(\nabla f(x_i), \nabla f(x_j)\right)$ between the support vector gradients for all pairs of training data. The results, which were measured on the same dataset as in Table 1, indicate that the proposed system generates different trajectories for different faces, so it does capture the individual variation in aging. The second row lists the mean angles $\measuredangle(\nabla f(x_i), \nabla_{lin})$ between the gradients at the training data and the constant gradient of the linear estimation, demonstrating that the trajectories using Support Vector regression deviate considerably from the simple linear approach.

The third row in Table 2 gives the internal angles $\measuredangle(\nabla f(x_{start}), \nabla f(x_{end}))$ between the tangent vectors in the starting and end points (ages 95 and 360 months) of each individual trajectory $\mathbf{z}_{\mathbf{x}_0,t_0}(t)$. The angles show that the trajectories are curved, as we expected. Although the trajectories straighten with increasing numbers of principal components, they still go beyond a linear regression approach.

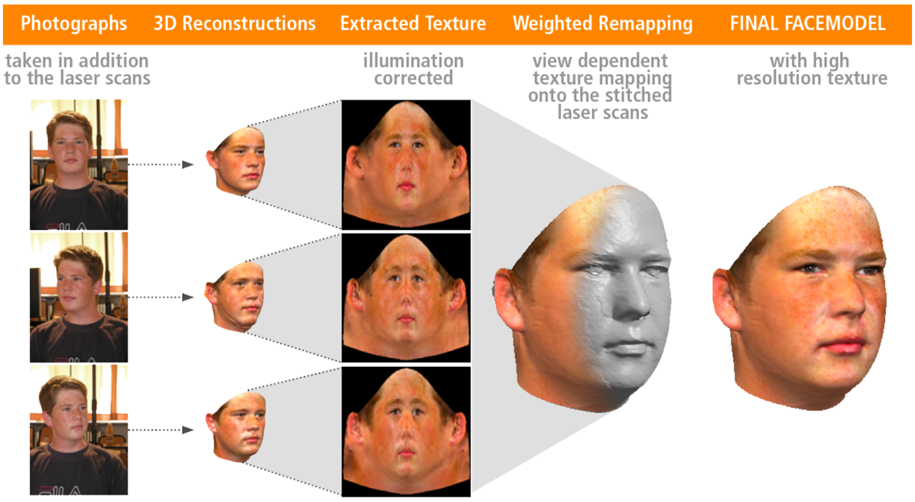

Figure 6: *The final texture of each facemodel consists of three separately acquired texture images. By fitting the 3D Morphable Model to each texture image consecutively, three illumination corrected textures are extracted. Combining them and mapping them to the reconstructed shape of the same face yields the final face model.*

	Shape			**Texture**		
	20 PCs	40 PCs	80 PCs	20 PCs	40 PCs	80 PCs
$\measuredangle(\nabla f(\mathbf{x}_i), \nabla f(\mathbf{x}_j))$	15.7°	22.5°	9.1°	33.5°	21.2°	14.8°
$\measuredangle(\nabla f(\mathbf{x}_i), \nabla_{lin})$	56.6°	49.4°	38.3°	55.0°	46.4°	37.2°
$\measuredangle(\nabla f(\mathbf{x}_{start}), \nabla f(\mathbf{x}_{end}))$	10.3°	10.9°	4.0°	30.0°	12.7°	6.2°

Table 2: *1st row: Mean angles between gradients $\measuredangle(\nabla f(\mathbf{x}_i), \nabla f(\mathbf{x}_j))$ in all pairs $(\mathbf{x}_i, \mathbf{x}_j)$ of training samples. 2nd row: Mean angles $\measuredangle(\nabla f(\mathbf{x}_i), \nabla_{lin})$ between gradients of the non-linear and the linear function. 3rd row: Mean angles $\measuredangle(\nabla f(\mathbf{x}_i), \nabla f(\mathbf{x}_j))$ between tangents in the start- and end points of 30 age trajectories. The training set consisted of 393 faces represented by 20, 40 and 80 principal components, with ages between 95 and 360 months.*

Figure 4 shows an example face, transformed along the linear and the non-linear trajectory. The visual differences between the curves are quite subtle. Based on the theoretical background and the precision of age predictions (Table 1), we can argue that the non-linear trajectory captures more of the true effects of facial growth.

6. A Database of 3D Scans of Teenagers

As we have mentioned above, the dataset of 200 adult faces was augmented by 3D scans of of 238 teenager faces, covering an age range from a minimum age of 96 months (8 years) to a maximum age of 191 months (almost 16 years). Among them, 125 were male faces, and 113 were female faces.

In a mobile setup that we operated at schools, we used Konica Minolta VI-910 laser scanners to record textured depth maps from different viewing angles. For each face, we recorded three partial 3D laser scans, each showing the face from a different viewing angle (0°, 35°and −35°) at a resolution of 76 800 vertices per scan (320×240). Unlike the cylindrical measurements by the stationary Cyberware PS3030 scanner, however, we had to combine the separate scans to a single surface.

For processing the individual scans, we fit the 3D Morphable Model of adults to the facial surface. We modified an algorithm that was previously described in [BV99] for cylindrical scans such that it operates on the depth maps provided by the Minolta scans. The fitting algorithm minimizes the difference in depth and texture between the linear combination of shapes (Equation (4) and the depth map of the scan. The result, in our case, is a best fit of the adult face model to each new scan. The model fitting establishes point-to-point correspondence to the Morphable Model, so we can form new shape and texture vectors by sampling the scans at each vertex of the reconstructed surface. For those parts of the face that were visible to the Minolta scanner, we replace the best-fit vertex position by the precise scan data, while all oth-

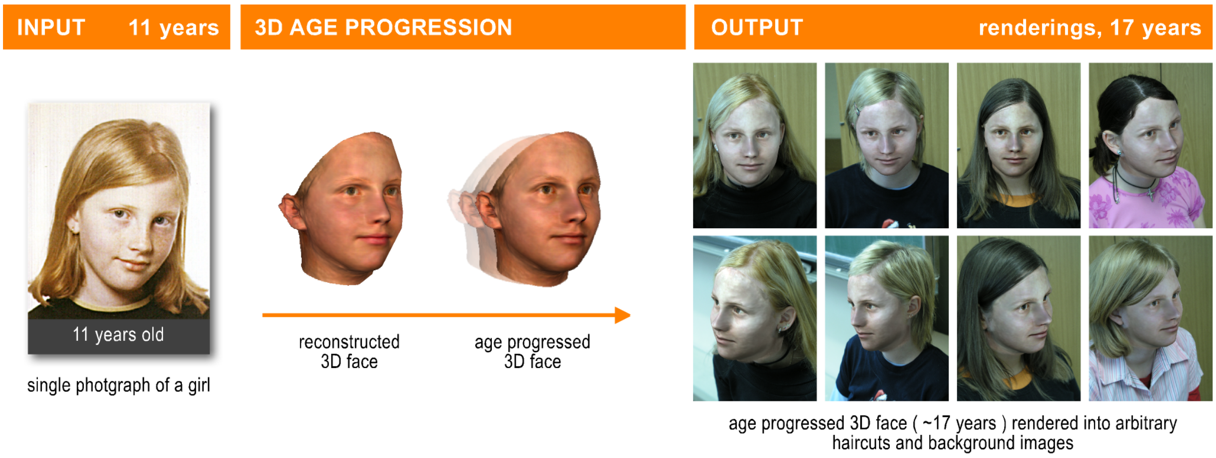

Figure 7: *Once a face is transformed to another age, it can be rendered into arbitrary images. This method allows us to apply different haircuts to a face.*

ers remain at the estimated best-fit position. To combine the different scans of the same face, we form a weighted sum of vertex positions wherever more than one scan direction gives a valid measurement. Vertices of the Morphable Model that are visible in none of the scans, for example holes or edges near the ears or on the neck, remain at the best-fit position.

To improve the resolution of the textures, we took digital images of the children. By fitting the model to these images and extracting the texture [BV99], we obtained a high resolution texture from each image. The fitting algorithm finds the linear combination of shape and texture vectors (Equation (4)), the rigid parameters of 3D pose and the parameters of illumination such that a synthetic image of the face model is as similar as possible to the input image [BV99]. After fitting, the color of each vertex is sampled in the images, and illumination effects are compensated automatically.

From the photographs taken of each person, we select a frontal view and a side view from 45 °to the left and right, respectively. The 3D reconstruction algorithm is applied to all 3 images separately. Each texel from each view is assigned a weight value that is computed according to the following scheme: If the point is occluded or if the angle ϕ between the viewing direction and the surface normal is larger than $\phi_1 = 80^\circ$, then the weight is $w = 0$. If $\phi < \phi_2 = 40^\circ$, $w = 1$. In the interval between, the function

$$w = \frac{1}{2}(1 + cos(\frac{\phi - \phi 2}{\phi_1 - \phi_2} \cdot 180^\circ)) \quad (9)$$

generates a continuous and smooth transition (Figure 6).

In a first step, the weights w are used in each view separately to combine the texture values extracted from the image with the value estimated by the fitting algorithm based on the linear combination of examples. The second step combines the textures from the 3 views with these weights. If all 3 weights are 0 in a texel, which means that all 3 textures are based on the best fit value only, these texture values are weighted equally. Figure 2 shows examples of aging 3D scans taken from our new database of face models.

7. Application to Image Data

In most applications, 3D scans of children whose age has to be changed are not available, so it is essential to be able to run the algorithm on image data. Moreover, the desired output of a system would not be an isolated 3D face mesh, but an image of a child in a new scene context. Our framework accounts for both aspects of the problem.

In order to reconstruct a 3D face from an image, we use the model-based algorithm [BV99] that was mentioned in the previous section, which estimates shape, texture and all relevant scene parameters. For the reconstructed face, the age trajectory can be computed.

For rendering the modified face into a novel target image, we rely on the fact that the scene parameters are estimated by the fitting algorithm [BSVS04]. The user runs the fitting algorithm also on the target image, and obtains the appropriate pose and illumination parameters that he needs for rendering the age-transformed face into the image. The rigid alignment of faces within the Morphable Model and the consistent illumination of the scans makes sure that we can exchange faces in this way. Examples of aging in images are shown in Figure 1, 3, 7 and 8.

Being able to render the age-transformed face into new images makes it possible to account for the growth of the rest of the body as well as potential changes in hairstyle and in the environment. We simply render the face into a photo of another child at the target age. By rendering the age-transformed face into images of children with different

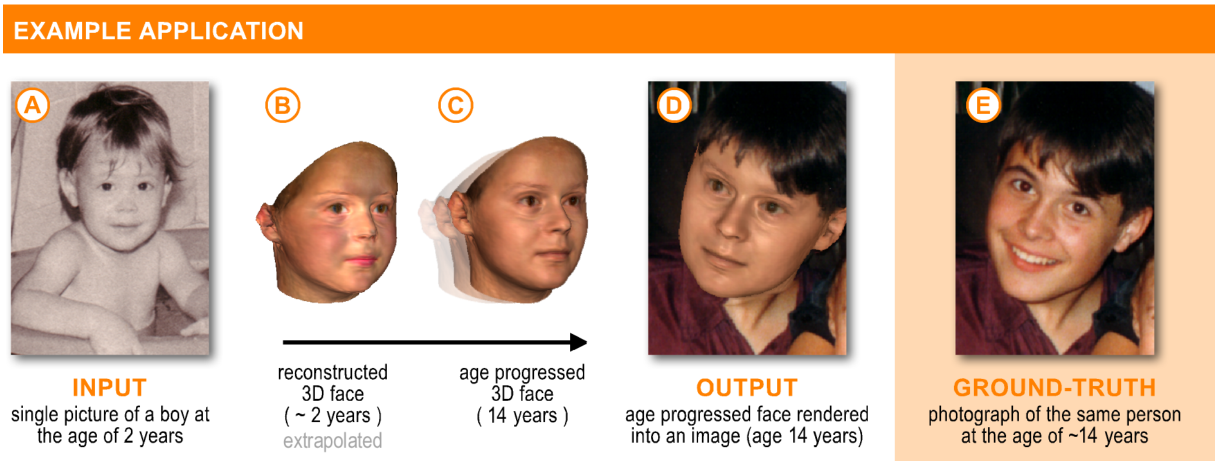

Figure 8: *The example shows a picture of a 2 year old boy (A). We reconstructed a 3D model of his face by fitting the Morphable Model to his face (B). Our growth algorithm then transformed the 3D face into a face at an appropriate target age (here: 16 years) (C). Finally we rendered the age progressed 3D face into an arbitrary background image (here: the ground-truth image was chosen as background) (D). For comparison, (E) shows a real picture of the boy at the target age.*

hairstyles, our system generates a variety of possible appearances (Figure 7).

8. Conclusion

We have addressed the problem of aging in a very general setting: From training data that contain no longitudinal measurements, we estimate how aging of faces depends both on age and on the individual face, and we apply our synthetic aging to images at any given pose and illumination.

A straight-forward application of our method is in law enforcement, where missing children can be found based on their predicted facial appearance. Our system supports the work of forensic artists in this field by providing a solid empirical ground for altering the features of faces, due to the example-based approach.

Additional applications are in the fields of character design, animation and morphing for special effects. Moreover, the system can be a high-level tool for image processing. Both the processing of scan data after acquisition, and the application of the algorithm to images of children are mostly automated, making the framework easy and convenient to use. Manual interaction is reduced to clicking a small number of feature points. This may be an important factor in any future deployment of the system.

In the best case, the prediction made by our algorithm can only give the most likely image of what a child may look like in the future. Environmental factors, such as nutrition, exposure to sunlight or physical activity contribute to the development of the face. Diseases, injuries or psychological stress may also leave visible traces. In addition to these factors, it may well be that some genes start to determine the visual appearance only at a certain age without having influenced the face before. Some of these latent genetic dispositions can be observed in images of parents or other relatives. Within the system presented in this paper, it is easy to import features from photographs of parents after a 3D reconstruction.

Even though linear models are still an appropriate basis in many fields of learning-based graphics, in particular if only a small set of training data are available, new non-linear have a tremendous potential to capture all the mutual dependencies between parameters in an elegant and unified theoretical framework.

References

[BAHS06] Blanz V., Albrecht I., Haber J., Seidel H.-P.: Creating face models from vague mental images. In *Computer Graphics Forum, Vol. 25, No. 3 EUROGRAPHICS 2006* (Vienna, Austria, 2006), Gröller E., Szirmay-Kalos L., (Eds.). 1, 6

[BP95] Burt D. M., Perrett D.: Perception of age in adult caucasian male faces: computer graphic manipulation of shape and colour information. *Proc. R. Soc. Lond. B 259* (1995), 137–143. 2

[BS81] Burson N., Schneider T.: Method and Apparatus for Producing an Image of a Person's Face at a Different Age. *Proceedings: Biological Sciences United States Patent 4276570*, 1355 (June 1981), 137–143. 2

[BSVS04] Blanz V., Scherbaum K., Vetter T., Seidel H.-P.: Exchanging faces in images. In *Computer Graphics Forum, Vol. 23, No. 3 EUROGRAPHICS 2004*

(Grenoble, France, 2004), Cani M.-P., Slater M., (Eds.), pp. 669–676. 2, 8

[BV99] BLANZ V., VETTER T.: A morphable model for the synthesis of 3D faces. In *Computer Graphics Proc. SIGGRAPH'99* (1999), pp. 187–194. 1, 2, 3, 4, 6, 7, 8

[CL01] CHANG C.-C., LIN C.: LIBSVM: a library for support vector machines. Software available at http://www.csie.ntu.edu.tw/~cjlin/libsvm, 2001. 5

[GMP*06] GOLOVINSKIY A., MATUSIK W., PFISTER H., RUSINKIEWICZ S., FUNKHOUSER T.: A statistical model for synthesis of detailed facial geometry. In *SIGGRAPH '06* (2006), pp. 1025–1034. 4

[GZZ*06] GENG X., ZHOU Z.-H., ZHANG Y., LI G., DAI H.: Learning from facial aging patterns for automatic age estimation. In *MULTIMEDIA '06: Proceedings of the 14th annual ACM international conference on Multimedia* (New York, NY, USA, 2006), ACM Press, pp. 307–316. 3

[HBHP03] HUTTON T. J., BUXTON B. F., HAMMOND P., POTTS H. W. W.: Estimating average growth trajectories in shape-space using kernel smoothing. *IEEE Trans. Med. Imaging 22*, 6 (2003), 747–753. 3

[HSG04] HILL C., SOLOMON C., GIBSON S.: Plausible aging of the human face using a statistical model. In *Proc. 7th International Conference on Computer Graphics and Imaging* (2004), IASTED Acta Press, pp. 56–59. 3

[Hut04] HUTTON T. J.: *Dense Surface Models of the Human Face*. PhD thesis, University College, London, 2004. 3

[KHYS02] KÄHLER K., HABER J., YAMAUCHI H., SEIDEL H.-P.: Head shop: Generating animated head models with anatomical structure. In *Proc. ACM SIGGRAPH Symposium on Comp. Anim. (SCA) 2002* (2002), pp. 55–64. 3

[LDC04] LANITIS A., DRAGANOVA C., CHRISTODOULOU C.: Comparing different classifiers for automatic age estimation. *IEEE Transactions on Systems, Man and Cybernetics, Part B 34*, 1 (February 2004), 621–628. 3

[LTC02] LANITIS A., TAYLOR C. J., COOTES T. F.: Toward automatic simulation of aging effects on face images. *IEEE Trans. Pattern Analysis and Machine Intelligence 24*, 4 (2002), 442–455. 2

[LWMT99] LEE W.-S., WU Y., MAGNENAT-THALMANN N.: Cloning and aging in a vr family. In *VR '99: Proceedings of the IEEE Virtual Reality* (Washington, DC, USA, 1999), IEEE Computer Society, p. 61. 3

[PTVF92] PRESS W. H., TEUKOLSKY S. A., VETTERLING W. T., FLANNERY B. P.: *Numerical Recipes in C*. Cambridge University Press, Cambridge, 1992. 6

[Ram06] RAMANATHAN N.AND CHELLAPPA R.: Modeling age progression in young faces. In *Computer Vision and Pattern Recognition* (2006), pp. 387–394. 3

[RP95] ROWLAND D., PERRETT D.: Manipulating facial appearance through shape and color. *IEEE Computer Graphics and Applications 15*, 5 (1995), 70–76. 2

[SS02] SCHÖLKOPF B., SMOLA A. J.: *Learning with Kernels*. MIT Press, Cambridge, Massachusetts, 2002. 5

[SSG06a] SCANDRETT C., SOLOMON C., GIBSON S.: Towards a semi-automatic method for the statistically rigorous ageing of the human face. *IEEE Proc. - Vision, Image, and Signal Processing 153*, 5 (October 2006), 639–649. 3

[SSG06b] SCANDRETT C. M., SOLOMON C. J., GIBSON S. J.: A person-specific, rigorous aging model of the human face. *Pattern Recognition Letters 27*, 15 (2006), 1776–1787. 2

[TBP01] TIDDEMAN B., BURT M., PERRETT D.: Prototyping and transforming facial textures for perception research. *IEEE Computer Graphics and Applications 21*, 5 (2001), 42–50. 2

[TMSP80] TODD J., MARK L., SHAW R., PITTENGER J.: The perception of human growth. *Scientific American 242* (Feb. 1980), 106–114. 2

[Vap95] VAPNIK V. N.: *The Nature of Statistical Learning Theory*. Springer-Verlag, New York, 1995. 5

[VBPP05] VLASIC D., BRAND M., PFISTER H., POPOVIC J.: Face transfer with multilinear models. In *Computer Graphics Proc. SIGGRAPH'05* (2005), pp. 426–433. 3

[VP97] VETTER T., POGGIO T.: Linear object classes and image synthesis from a single example image. *IEEE Trans. on Pattern Analysis and Machine Intelligence 19*, 7 (1997), 733–742. 4

[VT02] VASILESCU M. A. O., TERZOPOULOS D.: Multilinear analysis of image ensembles: Tensorfaces. In *ECCV (1)* (2002), pp. 447–460. 3

[WHL*04] WANG Y., HUANG X., LEE C.-S., ZHANG S., LI Z., SAMARAS D., METAXAS D., ELGAMMAL A., HUANG P.: High resolution acquisition, learning and transfer of dynamic 3-d facial expressions. In *Computer Graphics Forum* (2004), pp. III: 677–686. 3

[WKMMT99] WU Y., KALRA P., MOCCOZET L., MAGNENAT-THALMANN N.: Simulating wrinkles and skin aging. *The Visual Computer 15*, 4 (1999), 183–198. 3