



PRIFYSGOL CYMRU ABERTAWE
UNIVERSITY OF WALES SWANSEA

UNIVERSITY OF WALES SWANSEA

REPORT SERIES

**Exploring the Relationship between Faces and Aging
in Image-based Age Transformation**

by

Daniel Hubball, Min Chen and Phil W. Grant

Report # CSR 4-2007

 **Computer Science**
Gwyddor Cyfrifiadur

Exploring the Relationship between Faces and Aging in Image-based Age Transformation

Daniel Hubball, Min Chen and Phil W. Grant [†]

Swansea University, United Kingdom

Abstract

Aging has considerable effects on the appearance of the human face and is difficult to simulate using a universally-applicable global aging model. In this paper, we focus on the hypothesis that the patterns of age progression (and regression) are related to the face concerned, as the latter implicitly captures the characteristics of gender, ethnic origin, and age group, as well as possibly the lifestyle and health, of the individual. We present a data-driven framework for automatic image-based facial transformation in conjunction with a database of facial images. In addition to a parameterized model for face description, we also build a parameterized model for age-transformation. We introduced and compared several different methods to obtain a person-specific mapping in the model space from an encoded face description to an encoded age-transformation, and we adopted different algorithmic approaches including localization, person-sensitive linear combination, feature-sensitive linear combination, and learning the functional mapping from example transformations using genetic programming. We found that genetic programming is particularly suited for aging simulation where the current understanding is not yet adequate for producing an effective global aging model. In comparison with the existing techniques for age transformation, we have achieved a noticeable improvement in terms of the resemblance between the output images and the actual target images (which are unknown to either the age transformation algorithm or the training process), demonstrating the effectiveness and usability of this new approach.

Categories and Subject Descriptors (according to ACM CCS): I.2.8 [Artificial Intelligence]: Problem Solving, Control Methods and Search; I.3.3 [Computer Graphics]: Picture/Image Generation; I.3.6 [Computer Graphics]: Methodology and Techniques; I.4.5 [Image Processing and Computer Vision]: Reconstruction.

1. Introduction

Visual modeling and simulation of facial age progression has applications in law enforcement, forensic science, visual psychology and the entertainment industry. Age progressed images seen in the media today, such as those used for finding missing persons, typically involve skilled manual work to a large extent. A completely automated approach, which is still an elusive goal in these applications, presents an interesting challenge to computer graphics.

Aging is a relatively individual-specific process and is influenced by factors associated with some generic groupings such as gender, ethnic origin, growth phases and geo-environmental conditions, as well as factors associated with

the lifestyle and health of individuals. There were previous attempts to formulate a generic aging model, based on craniofacial growth patterns [PS75, MT83], and average facial images [BP95, TSP05]. The main shortcoming of relying on one or a few pre-defined global models is the difficulty for the models to accommodate the complexity and diversity of individuals' facial age progression.

The *data-driven* approach, which has been successfully deployed in areas such as facial animation (e.g., [PHL*98, VBPP05]) can address the shortcomings of the global model approach. It is practically feasible to establish and expand a database that contains facial images of a large number of individuals captured at different ages, together with some metadata about their gender, ethnic origin, lifestyle, etc. Given an input image of a face to be aged and some information about the person, our objective is to automatically

[†] e-mail: {m.chen, p.w.grant}@swansea.ac.uk

generate an age-transformation model that is specifically appropriate for this particular individual.

Based on the common practice of forensic artists, we hypothesized that a facial image implicitly captures the characteristics of gender, ethnic origin, and age group, as well as possibly the lifestyle and health, of the individual concerned. Hence, the patterns of age progression (and regression) are related to the facial image, though such a relationship is yet to be fully understood. This work is not intended to provide a scientific understanding of such a relationship in craniofacial growth. Instead, we focus on establishing an empirical model for capturing such a relationship from a database, and using the empirical model to synthesize new images for age-progression and regression.

In this paper, we present a data-driven framework for automatic facial age progression (and regression) in conjunction with a database of about 1500 facial images, most of which are in the public domain [FG-05]. We concentrate on image-based age transformation, since acquiring 3D facial data of a large number of individuals, each at different ages, poses a completely different problem in terms of practicality. Nevertheless, the extrapolation of the models and algorithms presented to 3D is a relatively straightforward exercise.

Our approach differs from previous work on facial age progression in several ways:

- We do not rely on a global model, and instead, we build person-specific models for face description and age-transformation based on a subset of imagery data selected according to an input face and associated metadata.
- In addition to the traditional approach for using the *morphable face model* to encode faces, we also compute age progression through a new *morphable aging model* for encoding differential specifications of age-transformation. Both parameterized models are constructed using *principal component analysis*.
- We focus on the hypothesis that the patterns of age progression (and regression) are related to the face concerned. We introduced and compared several different methods to obtain a person-specific mapping in the model space from an encoded face description to an encoded age-transformation, by adopting different algorithmic approaches including *localization*, *person-sensitive* linear combination, *feature-sensitive* linear combination, and *learning* the functional mapping from example transformations using genetic programming.
- The last approach using genetic programming is particularly significant as it does not presuppose a functional model that specifies a parameterized mathematical relationship between the model for face description and that for age-transformation. Hence most optimization methods, which operate only in the parameter space of a presupposed functional model, cannot be deployed to establish such a relationship. In order to obtain a person-specific mapping in the model space from an encoded

face description to an encoded age-transformation, we employed genetic programming to automatically generate a solution (function as well as parameters) by learning from example transformations in the training set.

- In comparison with the existing techniques for age transformation, we have achieved a noticeable improvement in terms of the resemblance between the output images and the actual target images (which are excluded from any training set), demonstrating the effectiveness and usability of this new approach, especially to such a complex problem for which the scientific community is yet to gain sufficient understanding in order to formulate a universally-applicable global model.

2. Related Work

Facial Age Progression. The traditional approach to generation of age progressed images typically involves the skills of an artist working from photos of the individual concerned. Semi-automated techniques are now common, which allow forensic artists trained in facial anthropometry to modify existing photographs using interactive software tools and impose appropriate facial changes based on a known aging trend. For example, for generating an aged image of a missing child, a forensic artist normally uses photographs of the close relatives of the child to identify a common aging trend. Such a process is time consuming and the quality of the results depends on the skills and interpretations of the artist as well as the availability of suitable information for identifying an aging trend.

Early attempts at age progression used geometric transformations to model craniofacial growth. Pittenger *et al.* [PS75, PSM79] applied affine shear and cardioidal strain transformations to facial outlines to determine the effects of the transformations on the perceived age. Mark and Todd [MT83] extended this method to face images and 3D facial data, showing that growth transformations were applicable to more detailed facial representations. A series of attempts were made to obtain 3D measurements of the growth of human faces (e.g., [ABC*00]). In particular, Hutton *et al.* [HBHP03] reported a study with 3D surface scans of 400 subjects. Using the 3D parametric model similar to [BV99], they constructed a global age trajectory representing the change of facial geometry.

Burt and Perrett [BP95] employed image morphing for age progression that facilitates both geometric deformations and texture changes. The method is based on a composite face created from a set of images. Further extensions to this work [TBP01, TSP05] investigated wavelet-based texture enhancement to compensate for fine texture detail which is often lost in the composite images.

Lanitis *et al.* [LTC02] proposed a statistical approach to age progression, which is centered around an age function that estimates an age value from a given facial image in para-

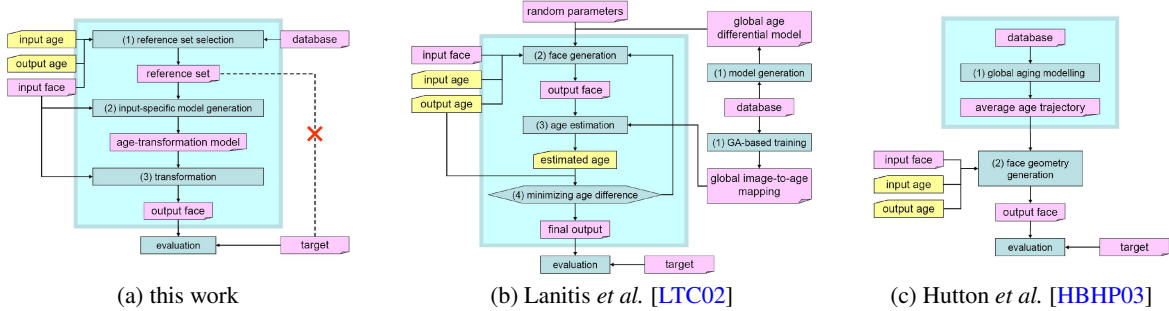


Figure 1: Data-driven frameworks for aging simulation.

metric space. The parameters of this age function is optimized using a genetic algorithm with age-tagged facial images. Given an age, [LTC02] generates a large number of plausible combinations of parameters of a face (though the paper did not explain how plausible combinations were decided). The age function is then invoked to see if each combination is equal to the given age. Based on this mechanism, age progression of a face is obtained. The results obtained by [LTC02] represent the state of the art of aging simulation. However, the perception of the results is undermined by the artificial clipping of the resultant images, which often did not follow the face boundaries.

In addition, attempts were made for capturing typical aging features such as wrinkles and spots from examples and superimposing such features onto a model (e.g., [BBTN04, MAK*02, LWMT99]). Recently, [GMP*06] used a statistical model to synthesize facial textures for simulating aging effects. Nevertheless, the focus on these attempts was normally placed on visual realism of an artistic image transformation or surface details of a 3D virtual human. The work in this domain usually relies on establishing a generic global model, hence limits its application mainly to virtual environments, arts and entertainment.

Metamorphosis of Imagery Data. Our work is also related to previous work in image and volume metamorphosis, which is used extensively in computer animation and for manipulating 2D and 3D imagery data. The existing 2D algorithms (e.g., [Wo190, BN92]) provided us with a set of valuable considerations in algorithm design and feature specification, while their 3D extensions (e.g., [LGL95, SD96, BV99]) prompted us to consider the non-uniform effects of 3D shape transition in the projective image space.

Evolutionary Computing. This is a branch of Artificial Intelligence [BFM97], and it has been successfully deployed in many areas of science, engineering, and industrial design. In computer graphics, there has been a series of investigations into the use of evolutionary computing methods for model creation and optimization. These include optimizing procedurally generated images [Sim91], designing sculptures [TLH91] generating motion for articulated figures [NM93, vF93, AFP*95], solving sampling problems in global

illumination [LH93], creating virtual creatures [Sim94], and modeling facial expression [Lim95]. Genetic programming [Koz92, Koz94] was a common method employed in such development (e.g., [Sim91, WO93, Lim95]).

Considerations and Remarks. The work presented in this paper draws various successful experience from previous work, while making significant new contributions to overcome the shortcomings in the existing approaches.

In particular, we consider that *Principal Component Analysis* (PCA), offers a cost-effective approach to the modeling of human faces, which has been demonstrated in a huge collection of previous work, including [LTC02, HBHP03, BV99]. In addition, we propose to extend the use of PCA for modeling age transformation, resulting in a differential model for aging.

We consider that using local models for age progression as in [LTC02] is more consistent with the approach by forensic artists, but believe that their approach for building aging simulation on an age estimation function introduces an accuracy bottleneck with a single scalar value. Their approach also makes the general aging model difficult to evaluate directly, and a trial-and-test mechanism has to be used instead.

We consider that the 3D approach by [HBHP03] may have helped understand the general trend of aging. In the short or medium term, the lack of temporal 3D data of individual faces hinders the construction of local models. By focusing on 2D imagery data, we are able to explore the relationship between individual faces and their aging patterns, while devising an aging simulation method that can be deployed in many practical applications.

We consider that machine learning offers a practical solution to address the complexity in modeling age progression. This was partially demonstrated by [LTC02, HBHP03]. We propose to take this learning process further by not assuming a fixed function formation, such as polynomials, in modeling the relationship between face description and age-transformation. Instead, we use genetic programming to learn the formation and parameters of a mapping function that captures such a relationship.

3. A Data-Driven Framework

We adopted a *data-driven approach* for modeling and simulation of facial age progression. We found that the design principles of some existing data-driven frameworks (e.g., [BV99]) can be adapted for this work. Figure 1 illustrates the main algorithmic steps of the data-driven framework developed in this work, which are juxtaposed with those of [LTC02] and [HBHP03]

Database. The collection of a large volume of suitable facial images and associated metadata is an inherently difficult task to be undertaken within a relatively small research project. The database was thereby designed in a flexible and extensible manner, accommodating incompleteness in data entries. The current database contains approximately 1500 facial images of 100 individuals at various ages. For each person, photographs are stored as an age progressive image set and tagged with the associated age in years. The database contains entries from both genders, and of individuals from various ethnic origins, with age ranging between 1 and 70 years. Unsurprisingly, none of the image sets collected has a complete cover over the 70 year span. As we rely heavily on family photographs, they inevitably vary in quality, resolution, illumination, 3D head pose, and facial expression, and often contain occlusions from hair, glasses, makeup, clothing and other people or objects in the scene. Some images are in grayscale.

Input Face and Faces in the Database. In Figure 1(a), the term ‘face’ is referred to as a placeholder for an original image, the associated feature points, its PCA encoded parameter set (see also 4.1), and some metadata such as gender, age, and ethnic group.

Reference Set. A reference set is a set of image pairs selected from the database according to the input and output ages, and the input face. It is used to construct an age-transformation model dynamically. The selection and use of a reference set will be further discussed in Sections 4.2 and 5. In this work, the target faces used for evaluation were always excluded from the reference set.

Age-Transformation Model. In Section 4.2, we will describe a general morphable age-transformation. However, this general model contain unknown attributes, and thereby is not adequate for performing age-transformation. In Section 5, we will propose and discuss several different approaches for approximating this general model.

4. Modeling Faces and Facial Transformation

A facial image F is represented by two elementary components, namely *geometry* G and *texture* T . G is a set of annotated 2D points, which specifies the key geometrical features of a face, and is stored in the database as a vector with the corresponding face F .

A *reference geometry* is a standard geometry vector $G_{\mathcal{R}}$

associated with the database, and it normally remains unchanged throughout the life-cycle of the database. Its main use is for defining a *normalized texture space*, and can be obtained from an abstract facial representation, a typical human face, or as the average of a set of geometry vectors. For our database, $G_{\mathcal{R}}$ was calculated as the mean geometry vector of a set of well-formed facial images selected manually.

The *texture component*, T , is a texture image generated by warping the original facial image F from the corresponding geometry component G to the reference geometry $G_{\mathcal{R}}$. The texture component, which is a color image, is also stored alongside the original face F . For consistency in the database, all grayscale textures are converted to color textures, via the *HSV* color space. We use the grayscale texture as the V component, and take the H and S components from a mean color texture computed from a set of well-formed face images selected for the corresponding ethnic group. We then transform the *HSV* texture back to the *RGB* color space, and normalize the *RGB* texture using histogram equalization, resulting in an approximated color texture.

Encoding G and T in a parameterized model provides a highly effective means for modeling and transforming shapes and textures [BV99]. In particular, PCA enables us to build a parameterized *morphable face model* (MFM), for geometry and texture respectively, from a set of basis facial images, and to build a parameterized *morphable age-transformation model* (MATM), for modeling age-transformation of facial geometry and texture, from a set of basis transformations.

4.1. Morphable Face Model (MFM)

Let the geometry component of a facial image be a column vector of k feature points, that is, $G = (x_1, y_1, \dots, x_k, y_k)^T \in \mathbb{R}^{2k}$, and the texture component be a column vector of l pixels, that is, $T = (r_1, g_1, b_1, \dots, r_l, g_l, b_l)^T \in \mathbb{R}^{3l}$. Let $\langle G_i, T_i \rangle$, $i \in [1, n]$ (always assuming $i \in \mathbb{N}$) be a set of n basis faces, and $\langle \bar{G}, \bar{T} \rangle$ be the mean of these basis faces. We can apply PCA to $\langle G_i - \bar{G}, T_i - \bar{T} \rangle$, resulting in two sets of n eigenvectors, $\{\phi_{g,i} \in \mathbb{R}^{2k} \mid i \in [1, n]\}$ and $\{\phi_{t,i} \in \mathbb{R}^{3l} \mid i \in [1, n]\}$, both in column form.

Without losing generality, we assume that eigenvectors in each set are in descending order according to the corresponding eigenvalues. We select the first $m < n$ eigenvectors from each set, resulting in a *parameterized face model*, with reduced dimension, in the form of a tuple $\langle \Phi_g, \Phi_t \rangle$, where

$$\Phi_g = [\phi_{g,1}, \phi_{g,2}, \dots, \phi_{g,m}], \quad \Phi_t = [\phi_{t,1}, \phi_{t,2}, \dots, \phi_{t,m}]$$

correspond to geometry and texture respectively. An arbitrary face $\langle G_a, T_a \rangle$ can thus be encoded using the model as:

$$\alpha_g = \Phi_g^T \cdot (G_a - \bar{G}), \quad \alpha_t = \Phi_t^T \cdot (T_a - \bar{T}).$$

The model is morphable as it allows continuous metamorphosis of facial images in the model space. In other words,

a ‘numerical’ morph from $\langle \alpha_g, \alpha_t \rangle$ to an arbitrary parameter set $\langle \beta_g, \beta_t \rangle$ resulting in a morph from $\langle G_a, T_a \rangle$ to $\langle G_b, T_b \rangle$, such that $G_b = \overline{G} + \Phi_g \cdot \beta_g$ and $T_b = \overline{T} + \Phi_t \cdot \beta_t$.

4.2. Morphable Age-Transformation Model (MATM)

Given two facial images, F_{t_1} and F_{t_2} at ages t_1 and t_2 respectively, a transformation from F_{t_1} to F_{t_2} can be defined by the corresponding geometry and texture transformation, that is, $\Delta G = G_{t_2} - G_{t_1}$ and $\Delta T = T_{t_2} - T_{t_1}$. Given a set of n example transformations captured in the database, $\langle \Delta G_i, \Delta T_i \rangle$, $i \in [1, n]$, we can construct an aging model, from these basis transformations, for modeling the geometry and texture transformation respectively.

Let $\langle \overline{\Delta G}, \overline{\Delta T} \rangle$ be the mean of these basis transformations. Similar to the construction of the face model, we first apply PCA to $\langle \Delta G_i - \overline{\Delta G}, \Delta T_i - \overline{\Delta T} \rangle$, $i \in [1, n]$, resulting in two sets of eigenvectors, $\{\Psi_{g,i} \in \mathbb{R}^{2k} \mid i \in [1, n]\}$ and $\{\Psi_{t,i} \in \mathbb{R}^{3l} \mid i \in [1, n]\}$. We then reduce the dimension of the model by selecting $m < n$ eigenvectors with the largest eigenvalues from each set, which yields a *parameterized age-transformation model*, $\langle \Psi_g, \Psi_t \rangle$, where

$$\Psi_g = [\Psi_{g,1}, \Psi_{g,2}, \dots, \Psi_{g,m}], \quad \Psi_t = [\Psi_{t,1}, \Psi_{t,2}, \dots, \Psi_{t,m}].$$

The model is also morphable as it allows continuous metamorphosis of differential representations of age-transformation. Using $\langle \Psi_g, \Psi_t \rangle$, we can encode a given age-transformation $\langle \Delta G, \Delta T \rangle$ as a parameter set, $\langle \delta_g, \delta_t \rangle$, in the model space. In reverse, we can decode a given parameter set $\langle \delta_g, \delta_t \rangle$ into an age-transformation $\langle \Delta G, \Delta T \rangle$. Hence, an arbitrary parameter set $\langle \delta_g, \delta_t \rangle$ can be used to age-transform a face $\langle G_{t_1}, T_{t_1} \rangle$ to $\langle G_{t_2}, T_{t_2} \rangle$ as:

$$\begin{aligned} G_{t_2} &= G_{t_1} + \Delta G = G_{t_1} + \overline{\Delta G} + \Psi_g \cdot \delta_g, \\ T_{t_2} &= T_{t_1} + \Delta T = T_{t_1} + \overline{\Delta T} + \Psi_t \cdot \delta_t. \end{aligned} \quad (1)$$

Since we are interested in age progression, we restrict each face pair to be the facial images of the same person at two different ages, though the model defined here can in general be used in conjunction with other differential specifications of facial transformation.

5. Approximating Age-Transformation

Eq. (1) in Section 4.2 defines an age-transformation from an input age t_1 to a target age t_2 . However, in applications of age progression, the main challenge is that neither $\langle \delta_g, \delta_t \rangle$ nor $\langle \Delta G, \Delta T \rangle$ in Eq. (1) is usually known. In this section, we examine different approaches for approximating Eq. (1). Apart from the *global mean age-difference* method, all other methods were proposed in this work.

5.1. Linear Combination of Age-Difference

With this approach, age-transformation is realized either directly in terms of geometry and texture components of faces

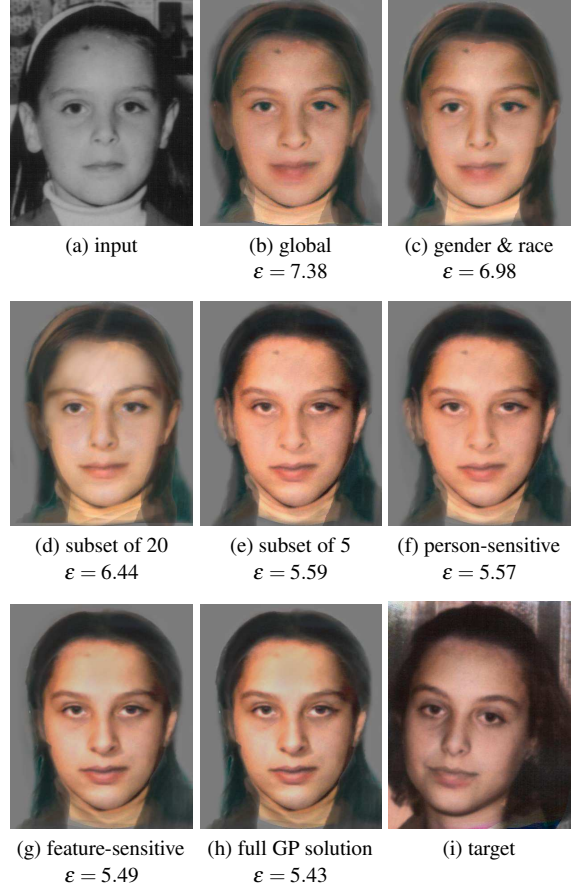


Figure 2: The data-driven framework enables an appropriate set of images to be selected according to an input image (a). Different levels of localization (c)-(e) show the relative merits over the use of a global model (b). Person-sensitive (f) and feature-sensitive (g) linear combination provides further improvement over localization. A full GP solution (h) gives the lowest level of MSE (ϵ) in comparison with the actual aged individual (g), which is not included in the training set.

or indirectly via the parameter space of MFM. The former implies simplifying Eq. (1) with

$$G_{t_2} \approx G_{t_1} + \overline{\Delta G}, \quad T_{t_2} \approx T_{t_1} + \overline{\Delta T}. \quad (2)$$

The latter utilizes MFM to encode the input face F at time t_1 and a set of s reference face-pairs

$$\mathcal{F} = \{(F_{a,1}, F_{b,1}), (F_{a,2}, F_{b,2}), \dots, (F_{a,s}, F_{b,s})\}.$$

Here we use indices a and b instead of t_1 and t_2 intentionally, as it is not necessary to make an exact match for the input and target ages. With \mathcal{F} , we can compute the average age-difference $\langle \overline{\Delta \alpha_g}, \overline{\Delta \alpha_t} \rangle$ in the parameter space, and obtain the output face $\langle \alpha_{g,t_2}, \alpha_{t,t_2} \rangle$ as

$$\alpha_{g,t_2} \approx \alpha_{g,t_1} + \overline{\Delta \alpha_g}, \quad \alpha_{t,t_2} \approx \alpha_{t,t_1} + \overline{\Delta \alpha_t}, \quad (3)$$

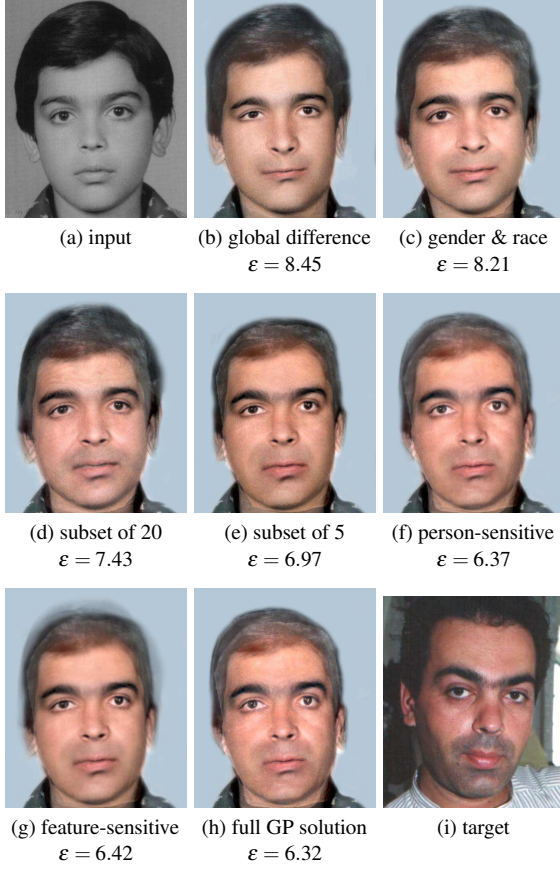


Figure 3: A different example that demonstrates a similar pattern of improvements to Figure 2. Noticeably, the person-sensitive approach (f) shows greater improvement over local mean differential specification, while the feature-sensitive approach (g) does not have any advantage over (f). The full GP solution (h) remains the best approach in this test.

from which $\langle G_{t_2}, T_{t_2} \rangle$ can be reconstructed.

Global Mean. Almost all existing work on facial age progression focused on the use of a *global mean differential specification* between faces at two different ages, or several of such mean specifications for different gender and ethnic groups (e.g., [MT83, HBHP03, TSP05]). While this method may be adequate for compiling some global statistics, the hypothesis that all males, all females, or each ethnic group will age in a similar manner is likely to be an over-generalization when considering a specific individual.

Local Mean. As mentioned in Section 2, forensic artists commonly use appropriate photographs of close relatives of a person to be aged to identify an age progression trend. This suggests that using a local model appropriate to the input data, including imagery and meta-information, should normally result in more accurate age progression images. Hence

we adopt an approach to enable the dynamic computation of localized models based on input data.

Given an encoded input face $\langle \alpha_{g,t_1}, \alpha_{t,t_1} \rangle$, we select a subset of persons from the database using a comparison metric. For each person in the database, let a and b be the ages of two available facial images of the person such that $|t_1 - a|$ and $|t_2 - b|$ are minimal. The metric computes a notional distance between the input and the person as a weighted sum of several normalized measurements including (i) gender, (ii) ethnic origin, (iii) $|t_1 - a|$, (iv) $|t_2 - b|$, and (v) imagery difference, in the parameter space, between the input and the person’s image at time a . In general it is not difficult to include other meta-information in the metric if such information is available in the database.

This allows us to establish a set of reference face-pairs \mathcal{F} ‘related’ to the input face, and obtain an output face based on a *local mean differential specification*. Results (b)-(e) in both Figures 2 and 3 indicate the improvement of age progression results through localization. As shown in (b), the image resulting from the use of a global mean exhibits few facial features of the target, which is unknown to the aging simulation process. Using a smaller subset consisting of images of those who are a closer match with $\langle G_{t_1}, T_{t_1} \rangle$, we can obtain age progression images that are usually closer to the target image.

Weighted Average. We can further improve the localization in the calculation of $\langle \Delta G, \Delta T \rangle$ or $\langle \Delta \alpha_g, \Delta \alpha_t \rangle$ by linearly moderating the age-difference of each pair of images in \mathcal{F} using the results of the above-mentioned comparison metric. This approach is said to be *person-sensitive*, as different persons in the reference set (i.e., image pairs in \mathcal{F}) contribute differently to the linear combination.

5.2. Linear Combination of MATM Parameter Set

Although localization clearly shows its advantage over the global approach, it still represents a coarse approximation by omitting the last term in Eq. (1), which parameterizes the age-transformation using MATM. Since for each image pair $(F_{a,i}, F_{b,i}), i \in [1, s]$ in \mathcal{F} we can obtain its corresponding $\langle \delta_{g,i}, \delta_{t,i} \rangle$, we can replace Eq. (1) with:

$$\begin{aligned} G_{t_2} &= G_{t_1} + \Delta G \approx G_{t_1} + \overline{\Delta G} + \Psi_g \cdot \overline{\delta_g}, \\ T_{t_2} &= T_{t_1} + \Delta T \approx T_{t_1} + \overline{\Delta T} + \Psi_t \cdot \overline{\delta_t}, \end{aligned} \quad (4)$$

where $\langle \overline{\delta_g}, \overline{\delta_t} \rangle$ is the average of $\langle \delta_{g,i}, \delta_{t,i} \rangle, \forall i \in [1, s]$.

Similar to the linear combination of age-difference in 5.1, one can compute $\langle \overline{\delta_g}, \overline{\delta_t} \rangle$ as a global mean, a local mean or a weighted average. Results (f) in both Figures 2 and 3 show the application of Eq. (4) with a *person-specific* linear combination of $\langle \delta_{g,i}, \delta_{t,i} \rangle, \forall i \in [1, s]$ moderated by the results of the comparison metric as in 5.1. Such an age-transformation is sensitive to individual persons in the reference set \mathcal{F} , and therefore the amount of the improvement over straightforward local mean specification depends on the variation of the

measured distances from the input face to faces in the reference set. For instance, in comparison with (e), (f) in Figure 3 shows more improvement than (g) in Figure 2, as the distance variation of the male reference set (STD=0.13) is much greater than that of the female reference set (STD=0.05). Our tests also show that there is a consistent, but minor improvement of using linear combination of the MATM parameter set over that of age-difference described in 5.1.

5.3. Feature-Sensitive MATM Parameter Set

We can extend the above-mentioned *person-sensitivity* to *feature-sensitivity* by considering that different features of each person in \mathcal{F} may affect the age-transformation differently. Those features that resemble the corresponding features of the input face should influence the age-transformation more than those different. Hence we would like to moderate the linear combination of $\langle \delta_{g,i}, \delta_{t,i} \rangle, \forall i \in [1, s]$ differently for each individual feature.

Let G be the geometry vector of the input face, and $\{G_{a,1}, G_{a,2}, \dots, G_{a,s}\}$ be the geometry vectors of the ‘ a ’ faces in the reference set \mathcal{F} . We can measure the difference between the j^{th} ($j \in [1, 2k]$) component of G and that of $G_{a,i}, i \in [1, s]$ as

$$\mu_{g,i}[j] = 1 - \tanh(|G[j] - G_{a,i}[j]|), \quad \omega_{g,i}[j] = \frac{\mu_{g,i}[j]}{\sum_{i=1}^s \mu_{g,i}[j]}$$

where $\tanh()$ is the hyperbolic tangent function that maps domain $[0, \infty]$ to $[0, 1]$. We organize all ω measurements into an $s \times 2k$ *feature distance* matrix as:

$$\Omega_g = \begin{bmatrix} \omega_{g,1}[1] & \omega_{g,1}[2] & \cdots & \omega_{g,1}[2k] \\ \omega_{g,2}[1] & \omega_{g,2}[2] & \cdots & \omega_{g,2}[2k] \\ \dots & \dots & \dots & \dots \\ \omega_{g,s}[1] & \omega_{g,s}[2] & \cdots & \omega_{g,s}[2k] \end{bmatrix}.$$

Given the encoded geometry parameters $\{\delta_{g,1}, \delta_{g,2}, \dots, \delta_{g,s}\}$ for the age-transformations in \mathcal{F} , we compute an encoding for each feature, $j \in [1, 2k]$ as the weighted linear combination of $\delta_{g,i}, i \in [1, s]$, yielding $2k$ sets of feature-sensitive geometry encoding as:

$$\Theta_g = [\delta_{g,1}, \delta_{g,2}, \dots, \delta_{g,s}] \cdot \Omega_g.$$

Similarly, for the texture vectors, we can use a texture comparison metric (e.g., [YS04]) to obtain an $s \times 3l$ *feature distance* matrix Ω_t , and $3l$ sets of feature-sensitive texture encoding Θ_t . We can now introduce feature-sensitivity to the last term of Eq. (4) as:

$$\begin{aligned} G_{t_2} &= G_{t_1} + \Delta G \approx G_{t_1} + \overline{\Delta G} + \text{diag}(\Psi_g \cdot \Theta_g), \\ T_{t_2} &= T_{t_1} + \Delta T \approx T_{t_1} + \overline{\Delta T} + \text{diag}(\Psi_t \cdot \Theta_t) \end{aligned} \quad (5)$$

where the function $\text{diag}()$ extracts the diagonal elements of a matrix, and places them in a column vector.

5.4. Evolving an MATM Parameter Set

Given $\langle G_{t_1}, T_{t_1} \rangle$, which is encoded as $\langle \alpha_g, \alpha_t \rangle$ using model $\langle \Phi_g, \Phi_t \rangle$, we want to find a relation between $\langle \alpha_g, \alpha_t \rangle$ and $\langle \delta_g, \delta_t \rangle$, that is, $\delta_g = \mathbf{F}_g(\alpha_g), \delta_t = \mathbf{F}_t(\alpha_t)$, such that \mathbf{F}_g and \mathbf{F}_t are two computable functions. Considering that such a relation is not yet well understood, and there is no suggestion for any computable formulae in the literature, we make no assumption about the precise format of \mathbf{F}_g and \mathbf{F}_t . Instead, we use generic programming to evolve suitable functions, driven by the example relations in the subset of image pairs selected according to $\langle G_{t_1}, T_{t_1} \rangle$. The two example results shown in Figures 2(h) and 3(h) clearly demonstrate the advantages of this approach.

5.4.1. Relation between Face and Aging

Consider a reference set, \mathcal{F} , of image pairs selected according to $\langle G_{t_1}, T_{t_1} \rangle$. Using $\{F_{a,1}, F_{a,2}, \dots, F_{a,s}\}$ in \mathcal{F} , we build a local MAM, which is subsequently used to encode $F_{a,i}$ as $\langle \alpha_{g,i}, \alpha_{t,i} \rangle, \forall i \in [1, s]$. Meanwhile from \mathcal{F} , we derive a local MATM, which is then used to encode the difference between each image pairs as $\langle \delta_{g,i}, \delta_{t,i} \rangle, i \in [1, s]$. We thus have s pairs of example mappings,

$$\langle \alpha_{g,i}, \alpha_{t,i} \rangle \mapsto \langle \delta_{g,i}, \delta_{t,i} \rangle, i = 1, 2, \dots, s.$$

We can learn from these example mappings, and evolve a new mapping $\langle \alpha_g, \alpha_t \rangle \mapsto \langle \delta_g, \delta_t \rangle$ for $\langle G_{t_1}, T_{t_1} \rangle$, which is defined by two computable functions \mathbf{F}_g and \mathbf{F}_t . Let:

$$\begin{aligned} \alpha_g &= \{u_{g,1}, u_{g,2}, \dots, u_{g,m}\}, & \alpha_t &= \{u_{t,1}, u_{t,2}, \dots, u_{t,m}\}; \\ \delta_g &= \{v_{g,1}, v_{g,2}, \dots, v_{g,m}\}, & \delta_t &= \{v_{t,1}, v_{t,2}, \dots, v_{t,m}\}. \end{aligned}$$

$m \leq s$ is the reduced dimension of MAM and MATM. We define \mathbf{F}_g and \mathbf{F}_t by using a set of sub-functions:

$$v_{g,j} = \mathbf{f}_{g,j}(u_{g,1}, u_{g,2}, \dots, u_{g,m}) \quad v_{t,j} = \mathbf{f}_{t,j}(u_{t,1}, u_{t,2}, \dots, u_{t,m}),$$

where $j = 1, 2, \dots, m$. As each $\mathbf{f}_{g,j}$ is a scalar function, we can formulate the sub-function as an arithmetic expression of m variables $\{u_{g,1}, u_{g,2}, \dots, u_{g,m}\}$. We also define $\mathbf{f}_{t,j}$ in the same manner. Therefore, the task for obtaining each sub-function becomes the evolution of an expression tree that represents the sub-function. The example image pairs provide the evolution with a training set for evaluating the generated \mathbf{F}_g and \mathbf{F}_t in the process.

5.4.2. Genetic Programming

Genetic programming [Koz92, Koz94] is one of the most effective methods for evolutionary computing, and differs most noticeably from other evolutionary methods (e.g., genetic algorithms) by allowing the automatic evolution of the formation and parameters of a function in the form of trees.

Terminal and Function Sets. Each individual of the population is a tree built from a user defined set of terminals T_{set} and functions F_{set} . The terminal set, which is used for defining leaves of the tree, consists of the variables of the




Figure 4: An MATM is evolved based on 22 examples, and applied to transform the input image. Four stages are shown, including the initial generation, and the best of three example generations, $\eta = 100, 200, 500$.

problem, and random real constants. Elements of the function set can be assigned to internal nodes of the tree. For our application, T_{set} and F_{set} are defined as:

$$T_{set} = \{\mathbb{R}, u_{g,1}, u_{g,2}, \dots, u_{g,m}\},$$

$$F_{set} = \{+, -, *, \div, \text{sqrt}, \text{log}, \text{abs}, \wedge^2, \wedge^3\}.$$

Here, as an example, we consider only a terminal set for subfunctions in F_g . The same process can be applied to F_t . The arity of the functions are also defined.

Initialization. The process begins by building an initial population of trees by randomly selecting elements from T_{set} and F_{set} and by setting a user-definable tree depth limit, D_{max} . Two methods for constructing the trees are possible. Beginning at the root, the *Full* method selects nodes from T_{set} until depth D_{max} is reached, and then nodes are taken from F_{set} . The *Grow* method stochastically chooses nodes from T_{set} and F_{set} at all depths before D_{max} . Here, we used the *ramped half-and-half method* [Koz92], where each individual tree of the initial population is constructed by choosing either of the two methods with equal probability.

Genetic Operators. The crossover genetic operator takes two parents and produces two offspring. Offspring are formed by randomly selecting nodes and by the exchange of subtrees. Mutation takes a single parent and produces one offspring. Mutation is by replacing randomly chosen nodes or by growing a new subtree rooted at the chosen node.

Fitness Evaluation. The evaluation function determines the fitness of each individual in the population. Selection is proportionate to fitness, giving fitter individuals a higher probability of becoming a parent of the next population. We form fitness cases from the subset of image pairs \mathcal{F} . The fitness evaluation is carried out in the model space constructed from $\{F_{b,1}, F_{b,2}, \dots, F_{b,s}\}$, we take the error between two parameter sets as the fitness value.

Figure 4 demonstrates the evolution process with a training set of 22 fitness cases. The example shows the application of Eq. (1) with three different generations of $\langle \delta_g, \delta_t \rangle$, to the input, indicating a gradual improvement of the results.

6. Results

In our data-driven framework, there is no fundamental difference between *age progression* and *age regression*. Hence, to evaluate the methods developed in this work, we conducted a series of tests of age-transformation between different ages. All the results shown in this section were produced with the approach for evolving the MATM using genetic programming. Figure 5 shows one of such tests. In this example, the photo of 10 year old boy was used as the input, with the two other photos of the same person at year 2 and year 30 used as targets for evaluation. For each age-transformation, on average 15 training examples were used. In comparison with the targets, the results bear good resemblance to the target images, specially in terms of facial outline, relative positions and sizes of the main facial features. Together with those without comparative targets, they show an overall trend of age progression from 2 to 40. In general, the quality of each transformation reflects the quality of the corresponding training set.

Figure 6 shows another example with a female subject. On the whole, the transformation of facial features is apparent, though the overall impression of age progression is not as obvious as Figure 5. In our tests, this shortcoming is common with most female subjects. We identified that the main reason is due to the more noticeable changes of hairstyle in this gender group. As long as one uses a training set involving a reasonably large number of individuals, the average hairstyle, as shown in Figure 6, often dilutes the overall impression of the transformed facial features. Because hairstyle plays a significant part in human perception of faces, especially female faces, it is useful to take this issue into account during the initial search for the training set. One approach is to obtain several training sets, each containing individuals with relatively similar hairstyles in addition to other search criteria. Figure 7 shows the use of three different training sets to derive three different optional results.

The performance of the approach based on genetic programming depends mainly on the number of generations in the evolutionary computation. Typically, each generation takes about 29 seconds and 500 generations (about 4 hours) can normally yield some satisfactory results. The

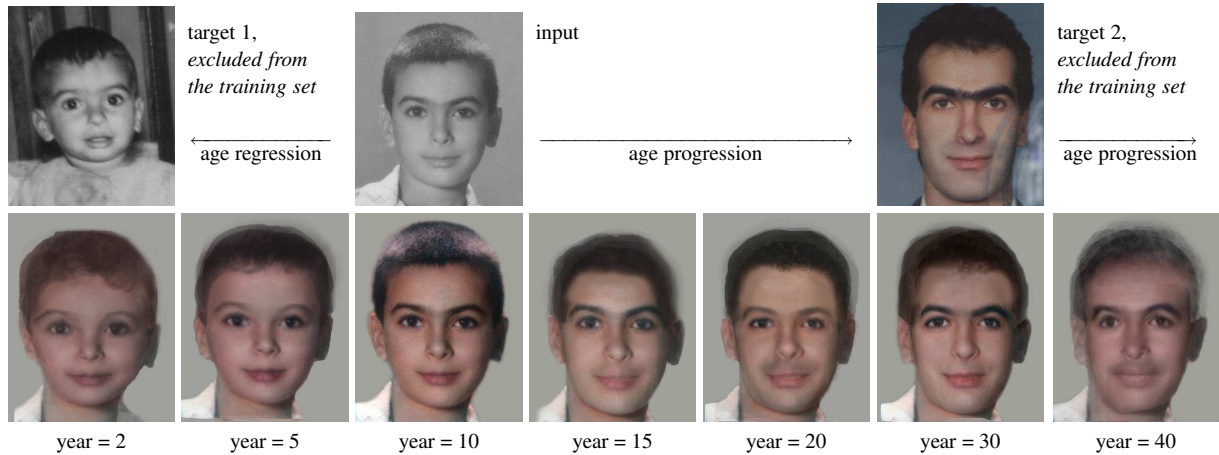


Figure 5: Age progression and regression for a given input of a male subject. Two target images are shown for comparison.

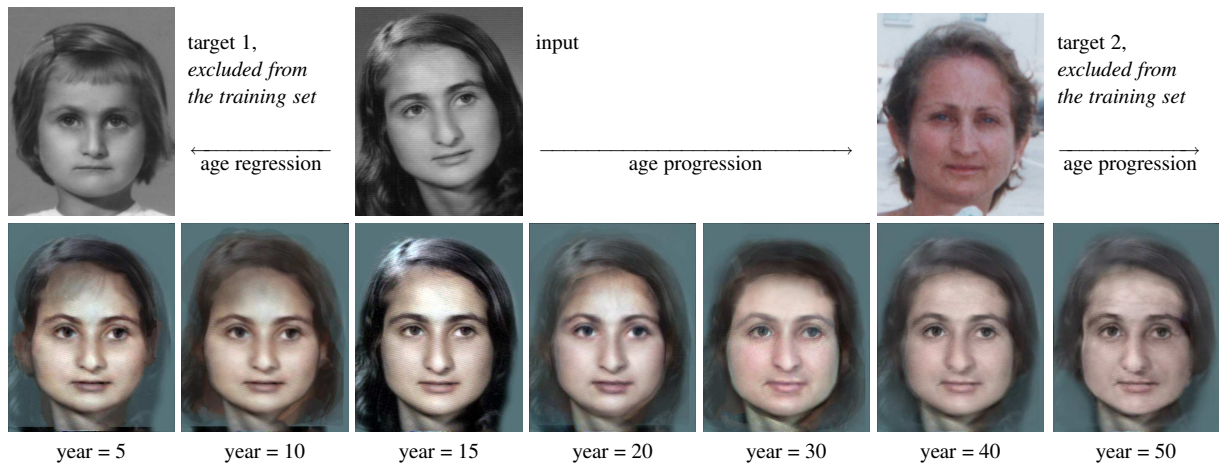


Figure 6: Age progression and regression for a given input of a female subject. Two target images are shown for comparison.

larger the training set, the more generations may be required for the evolution process to converge. In comparison with the semi-automatic approach available currently, which typically takes a few days work of a skilled forensic artist, this is a significant improvement in terms of time and resource requirements.

The performance of other approaches, including person-sensitive and feature-sensitive age-transformation, typically range from 1 to 5 minutes for a relatively small reference set.

7. Conclusions and Future Work

We have presented a data-driven framework for visual modeling and simulating age progression (and regression). To overcome the difficulty in formulating effective global models for age-transformation, we introduced the morphable age-transformation model (MATM), which enables us to explore the relationship between a face and its aging patterns.

We considered several approaches for establishing such a relationship according to a reference set in order obtain an approximated MATM. Among different approaches, using genetic programming to evolve a mapping from an encoded face to an encoded age-transformation clearly shows its advantage in producing the best results. Other approaches, such as person- and feature sensitive age-transformation, are also shown to be more effective than the traditional methods based on one or a few global models. These approaches can be deployed cost-effectively in interactive applications.

In comparison with the state of the art (e.g., [LTC02, HBHP03]), the results obtained in this work represent a significant leap in terms of resemblance with target images. In particular, we chose to show the results without any post-processing, such as hair removal and face clipping, avoiding any favorable bias in the perception of the results. These results demonstrated that our approach is practically feasible as it does not require historical 3D data of individuals, and




Figure 7: Using three different subsets of training examples, we obtain three different ‘de-aged’ hairstyles for a female subject.

conceptually close to the knowledge-based approach by human forensic artists.

We believe that it is beneficial, even just for scientific research on aging, to establish a substantial database provided that it is in compliance with appropriate legislation and regulations. With such a database, it may be possible to gain a quantitative understanding of the relationship between some major characteristics of aging and the major visual features of human faces, as well as the relationship between age-transformation and various clusters of metadata. Such quantitative understanding could in return serve applications such as computer animation and virtual environments.

Acknowledgment. Most image sets used in this work were provided by the FG-Net consortium [FG-05]. Authors are also grateful to many friends who have given us the permission to use their family photos.

References

- [ABC*00] ANDRESEN P. R., BOOKSTEIN F. L., CONRADSEN K., ERSBÅYLL B. K., MARSH J. L., KREIBORG S.: Surface-bounded growth modeling applied to human mandibles. *IEEE Transactions on Medical Imaging* 19, 11 (2000), 1053–1063. 2
- [AFP*95] AUSLANDER J., FUKUNAGA A., PARTOVI H., CHRISTENSEN J., HSU L., REISS P., SHUMAN A., MARKS J., NGO J. T.: Further experience with controller-based automatic motion synthesis for articulated figures. *ACM Transactions on Graphics* 14, 4 (1995). 3
- [BBTN04] BASTANFARD A., BASTANARD O., TAKAHASHI H., NAKAJIMA M.: Towards anthropometrics simulation of face rejuvenation and skin cosmetic. *Computer Animation and Virtual Worlds* 15 (2004), 347–352. 3
- [BFM97] BACK T., FOGEL D. B., MICHALEWICZ Z. (Eds.): *Handbook of Evolutionary Computation*. IOP Publishing Ltd., Bristol, UK, 1997. 3
- [BN92] BEIER T., NEELY S.: Feature-based image metamorphosis. *Computer Graphics (Proc. SIGGRAPH 92)* 26, 2 (1992), 35–42. 3
- [BP95] BURT D. M., PERRETT D. I.: Perception of age in adult caucasian male faces: Computer graphic manipulation of shape and color information. *Proceedings of the Royal Society of London, Series B – Biological Sciences* 259 (1995), 137–143. 1, 2
- [BV99] BLANZ V., VETTER T.: A morphable model for the synthesis of 3d faces. In *Computer Graphics (Proc. SIGGRAPH 99)* (1999), ACM Press, pp. 187–194. 2, 3, 4
- [FG-05] FG-NET CONSORTIUM: Aging Database. <http://sting.cyccollege.ac.cy/~alanitis/fgnetaging/>, 2005. 2, 10
- [GMP*06] GOLOVINSKIY A., MATUSIL W., PFISTER H., RUSINKIEWICZ S., FUNKHOUSER T.: A statistical model for synthesis of detailed facial geometry. *ACM Transactions on Graphics (Proc. SIGGRAPH 2005)* 25, 3 (2006), 1025–1034. 3
- [HBP03] HUTTON T. J., BUXTON B. F., HAMMOND P., POTTS H. W. W.: Estimating average growth trajectories in shape-space using kernel smoothing. *IEEE Transactions on Medical Imaging* 22, 6 (2003), 747–753. 2, 3, 4, 6, 9
- [Koz92] KOZA J. R.: *Genetic Programming*. MIT Press, Cambridge, MA, 1992. 3, 7, 8
- [Koz94] KOZA J. R.: *Genetic programming II: automatic discovery of reusable programs*. MIT Press, Cambridge, MA, 1994. 3, 7
- [LGL95] LERIOS A., GARFINKLE C. D., LEVOY M.: Feature-based volume metamorphosis. In *Computer Graphics (Proc. SIGGRAPH 95)* (1995), Addison-Wesley, pp. 449–456. 3
- [LH93] LANGE B., HORNUNG C. : A solution to global illumination by genetic algorithms. In *Proc. Int. Conf. on Artificial Neural Networks and Genetic Algorithms* (Innsbruck, Austria, 1993), pp. 595–601. 3
- [Lim95] LIM I. S.: Evolving facial expressions. In *Proc. IEEE Int. Conf. on Evolutionary Computation* (1995), vol. 2, pp. 515–552. 3
- [LTC02] LANITIS A., TAYLOR C. J., COOTES T. F.: Toward automatic simulation of aging effects on face images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 24, 4 (2002), 442–455. 2, 3, 4, 9

- [LWMT99] LEE W.-S., WU Y., MAGNENAT-THALMANN N.: Cloning and aging in a VR family. *The Visual Computer* 5, 1 (1999), 32–39. 3
- [MAK*02] MUKAIDA S., ANDO H., KINOSHITA K., KAMACHI M., CHIHARA K.: Facial image synthesis using age manipulation based on statistical feature extraction. In *Proc. VIIP 2002* (2002), pp. 12–17. 3
- [MT83] MARK L. S., TODD J. T.: The perception of growth in three dimensions. *Perception and Psychophysics* 33, 2 (1983), 193–196. 1, 2, 6
- [NM93] NGO J. T., MARKS J.: Spacetime constraints revisited. In *Computer Graphics (Proc. SIGGRAPH 93)* (1993), ACM Press and Addison-Wesley, pp. 343–350. 3
- [PHL*98] PIGHIN F., HECKER J., LISCHINSKI D., SZELISKI R., SALESIN D. H.: Synthesizing realistic facial expressions from photographs. In *Computer Graphics (Proc. SIGGRAPH 98)* (1998), ACM Press, pp. 75–84. 1
- [PS75] PITTENGER J. B., SHAW R. E.: Aging faces as viscal-elastic events: Implications for a theory of nonrigid shape perception. *Journal of Experimental Psychology: Human Perception and Performance* 1, 4 (1975), 374–382. 1, 2
- [PSM79] PITTENGER J. B., SHAW R. E., MARK L. S.: 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 (1979), 478–493. 2
- [SD96] SEITZ S. M., DYER C. R.: View morphing. In *Computer Graphics (Proc. SIGGRAPH 96)* (1996), Addison-Wesley, pp. 21–30. 3
- [Sim91] SIMS K.: Artificial evolution for computer graphics. *Computer Graphics (Proc. SIGGRAPH 91)* 25, 4 (1991), 319–328. 3
- [Sim94] SIMS K.: Evolving virtual creatures. In *Computer Graphics (Proc. SIGGRAPH 94)* (1994), ACM Press and Addison-Wesley, pp. 15–22. 3
- [TBP01] TIDDEMAN B., BURT D. M., PERRETT D. I.: Prototyping and transforming facial textures for perception research. *IEEE Computer Graphics and Applications* 21, 5 (2001), 42–50. 2
- [TLH91] TODD S., LATHAM W., HUGHES P.: Computer sculpture design and animation. *Journal of Visualization and Computer Animation* 2 (1991), 98–105. 3
- [TSP05] TIDDEMAN B. P., STIRRAT M. R., PERRETT D. I.: Towards realism in facial image transformation: Results of a wavelet mrf method. *Computer Graphics Forum* 24, 3 (2005), 449–456. 1, 2, 6
- [VBPP05] VLASIC D., BRAND M., PFISTER H., POPOVIĆ J.: Face transfer with multilinear models. *ACM Transactions on Graphics (Proc. SIGGRAPH 2005)* 24, 3 (2005), 426–433. 1
- [vF93] VAN DE PANNE M., FIUME E.: Sensor-actuator networks. In *Computer Graphics (Proc. SIGGRAPH 93)* (1993), ACM Press and Addison-Wesley, pp. 335–342. 3
- [WO93] WATABE H., OKINO N.: A study on genetic shape design. In *Proc. 5th Int. Conf. on Genetic Algorithms* (1993), pp. 445–450. 3
- [Wol90] WOLBERG G.: *Digital Image Warping*. IEEE Computer Society Press, Los Alamitos, CA, 1990. 3
- [YS04] YUAN X., SHI P.: Microcalcification detection based on localized texture comparison. In *Proc. International Conference on Image Processing* (2004), vol. 5, pp. 2953–2956. 7