

Image-based Aging Using Evolutionary Computing

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

Swansea University, United Kingdom

Abstract

Aging has considerable visual effects on the human face and is difficult to simulate using a universally-applicable global 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 person-specific development patterns of the individual. We use a data-driven framework for automatic image-based facial transformation in conjunction with a database of facial images. We build a novel parameterized model for encoding age-transformation in addition with the traditional model for face description. We utilize evolutionary computing to learn the relationship between the two models. To support this work, we also developed a new image warping algorithm based on non-uniform radial basis functions (NURBFs). Evolutionary computing was also used to handle the large parameter space associated with NURBFs. In comparison with several different methods, it consistently provides the best results against the ground truth.

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

1. Introduction

Visual modeling and simulation of facial age progression has applications in law enforcement, forensic science, visual psychology and the entertainment industry. 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 person-specific biological and lifestyle factors of individuals. There were previous attempts to formulate a global 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. Scandrett *et al.* [SSG06] first proposed a semi-automatic person specific approach by assuming a linear model that captures historical and consensus influence upon individuals, showing a promising way to address the shortcomings of global models.

This paper proposes a new method as further enhance-

ment to the person-specific approach. It considers that our current scientific understanding is not sufficient for establishing an adequate model for simulating person-specific age transformation, but assumes that there is a relationship between an individual's face and its aging patterns. We purposely do not attempt to establish an a-priori model for such a relationship, instead, we employ *genetic programming* to learn the functional mapping from a face model to an age transformation model in a person-specific manner.

In applications of aging simulation such as forensic science, missing person identification, lifestyle analysis, and online and mobile entertainment, obtaining a 3D model of an individual concerned poses great difficulty in practicality. While there are some 3D facial databases for estimating global aging trajectories [HBHP03, SSSB07], they have not yet reached the same scale as 2D facial databases in terms of the number of individuals and the temporal frequency and coverage at different ages. They are not suitable for person-specific aging as there are not sufficient person-specific links between different datasets in such 3D databases. Since the collection of individual temporal 3D facial data only started recently, it would take many years for a 3D database to reach the scale necessary for person-specific aging.

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

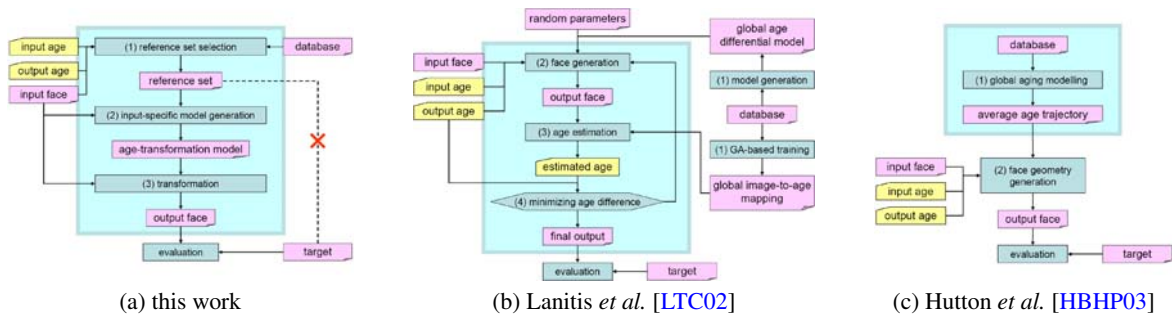


Figure 1: Three different data-driven frameworks for aging simulation.

We hence focus on image-based age transformation, utilizing the public domain FG-NET database [FG-05] with some 1500 facial images. In order to overcome the lack of 3D information of individual input images, we developed a new image warping algorithm based on non-uniform radial basis functions (NURBFs). The *genetic algorithm* technique (which is different from genetic programming) was used to handle the large parameter space associated with NURBFs.

Our results show that person-specific models performed better than global models in general, and our learning-based method (without the a-priori model of the relationship) consistently produced the closest matches in all comparison.

2. Related Work

Traditionally, the generation of age progressed images involves the skills of an artist working from photos of the individual concerned. 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 and modifies existing photographs of the child using interactive software tools.

Early attempts at age progression used geometric transformations to model craniofacial growth (e.g., [PS75, MT83]). A series of attempts were made to obtain 3D measurements of the growth of human faces (e.g., [ABC*00]). In particular, [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. [SSSB07] also collected 238 3D scans and established a global non-linear trajectory of growth. Because of the difficulties in collecting series of 3D scans of individual subjects, it is not a trivial task to establish person-specific aging patterns with such sparse data.

Image-based approaches can benefit from the existence of large collections of imagery data. [BP95, TSP05] employed image morphing for age progression that facilitates both geometric deformations and texture changes. 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 parametric space. [SSG06] introduced the notion of person-specific

aging model by moderating a global model with person-specific historical trends.

Attempts were also made for capturing typical aging features such as wrinkles and spots (e.g., [LWMT99, GMP*06]), focusing on visual realism of an artistic image transformation or surface details of a 3D virtual human.

Our work is also related to previous work in image and volume metamorphosis. The existing 2D algorithms (e.g., [Wol90, 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.

Considerations and Remarks. The work presented in this paper draws on 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, SSG06] is more consistent with the approach by forensic artists. However, [LTC02] relies on an age estimation function which becomes an accuracy bottleneck with a single scalar value. [SSG06] requires more than one input image in order to establish a historical axis, and assumption of the linear relationship between the historical and mean trajectories is also a major simplification.

We consider that the 3D approach by [HBHP03, SSSB07] can help our understanding of 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 solu-

tion 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 an a-priori model 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. 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. We built our database on the public domain database [FG-05] with approximately 1500 facial images of 100 individuals at various ages. The database was designed in a flexible and extensible manner, accommodating incompleteness in data entries. For each person, photographs are stored as an age progressive image set and tagged with the associated age in years. Each photo is associated with a set of 72 feature points. 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 covers every year over the 70 year span.

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 5.1), and some metadata such as gender, age, and ethnic group.

Reference Set. It is a set of face 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 5 and 6. In this work, the ground truth faces used for evaluation were always excluded from the reference set.

Normalized Texture Space. In order to account for a variety of variance, we compute all face textures in a normalized texture space. For the normalization process, we developed a new image warping algorithm (see Section 4) to account for the lacking in 3D information.

Reference Geometry. This is a standard geometry vector $G_{\mathcal{F}}$ 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, $G_{\mathcal{F}}$ was calculated as the mean geometry vector of a set of well-formed facial images selected manually.

Age-Transformation Model. In 5.2, we will define a person-specific age-transformation model and its relationship with

the traditional face model. The relationship is however an unknown attribute, and thereby in Section 6, we propose a new method for approximating this general model, and compare its effectiveness with several other approaches.

Genetic Algorithm and Genetic Programming. These are two of the most effective methods in *evolutionary computing* [Gol89, Koz92]. Both are search techniques for finding approximate or exact solutions to an optimization problem, but they are very different. The former specifies the solution space in the form of a genetic representation (e.g., an array of unknown coefficients of a polynomial). The latter assumes that the solution space be a computer program without an a-priori model. In this work, we use genetic algorithm for evolving solutions in our NURBF-based image warping since the solution space can be defined by a set of unknown parameters (Section 4). We use genetic programming for evolving the relationship between faces and aging patterns since we do not have an a-priori model of this functional mapping. A detailed description of how these two techniques are used can be found in [Hub07].

4. Non-Uniform Radial Basis Functions

Image warping is an essential component of this work, and in particular, it supports the texture space normalization, which is applied to every image in the database and every input image. It is also used to reverse the normalization for the output image. Note that the aging simulation in this work requires only the *warping* of an image under the influence of two sets of control features, but not the morphing between two images. We use only geometrical features for texture normalization, because (i) each photo in the FG-Net database [FG-05] has already been annotated with a standard set of 72 feature points, and (ii) it is not appropriate to use texture to influence its own warping, which would otherwise lead to inconsistent warping (e.g., changing image contrast could lead to a different warp). A number of recent algorithms designed for automatic control feature identification or smooth transformation between two images based on texture features (e.g., [CET01]) are thus neither required nor suitable for this work.

We adopted the feature-based field warping approach, which provides us with the necessary *flexibility* — in specifying a warping with a small set of key control features, and *controllability* — in directing the geometric deformation according to the defined control features. As the focus of existing work in image warping has largely been placed on the smooth transition from one image to another, the influence of each feature is normally uniformly defined in all directions proportionally to the proximity to the feature. The precise influence of each individual feature upon nearby pixels in a facial image is critical to the accurate modeling of age progression. Such influence depends on many factors, including the camera position and parameters corresponding to the image, and the 3D geometry of the featured face.

As our work is constrained to 2D projective space by the practicality of collecting a large number of facial images, most of which are family photographs taken over several decades, it is not feasible for obtaining a projection matrix for each image (as in view morphing [SD96]), or a 3D surface or volume (as in [VBPP05, LGL95]). We thereby developed a new warping algorithm that facilitates non-uniform influence of each control feature, which provides a generic means to encode the combined effect of various factors that determine the distribution of an influence function.

Let $\{C_1, C_2, \dots, C_u\}$ be a set of points each defining a control feature, and $\{D_1, D_2, \dots, D_u\}$ be the corresponding displacement of each control point. A traditional approach to the specification of the forward mapping of a point p in \mathbb{E}^2 under the combined influence of all control points is:

$$p' = p + \sum_{i=1}^n \omega_i \Theta_{RBF}(p, C_i) D_i, \quad \sum_{i=1}^n \omega_i = 1, \quad \omega_i \geq 0,$$

where ω_i is the weight of each control point, and Θ_{RBF} is a radial basis function (RBF), typically defined by a Gaussian function as:

$$\Theta_{RBF}(p, C_i) = \exp\left(-\frac{\|p - C_i\|^2}{R^2}\right), \quad R > 0.$$

R is the characteristic radius which affects the smoothness of the warping. It is common to assign different radii to individual control points to include additional nonlinear control over the weighting of each feature. Nevertheless, the influence of each control point remains uniformly distributed in all directions. In order to facilitate a non-uniform influence of each control point, we replace Θ_{RBF} with a non-uniform radial basis function (NURBF), Θ_{NURBF} , by using a Catmull-Rom spline in polar form.

4.1. Catmull-Rom Interpolation

For each feature point, $C_i, i = 1, \dots, u$, we define a set of v sub-control points $\{c_{i,1}, c_{i,2}, \dots, c_{i,v}\}$, each $c_{i,j} = \langle \theta_{i,j}, \rho_{i,j} \rangle$ is defined in polar coordinates relative to C_i . This enables us to replace the uniform R in Θ_{RBF} with a spline function.

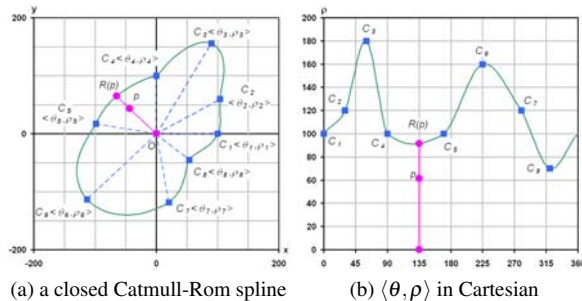


Figure 2: An example Catmull-Rom spline. (a) Its control points are specified in polar form. (b) The Catmull-Rom interpolation is applied to the spline by treating $\langle \theta, \rho \rangle$ as Cartesian coordinates.




Figure 3: A test image is warped by displacing a single control point vertically, (a) using a uniform RBF, (b) with a NURBF, which is defined by six sub-control points.

As illustrated in Figure 2, we define a closed Catmull-Rom spline [CR74] over the sub-control points sorted by θ . Other spline functions can also be used. However, the point-on-spline property associated with Catmull-Rom splines provides us a visually instinctive means for inspecting the formation process of each NURBF to be detailed in 4.2.

Given an arbitrary point $p \neq C_i$, we first calculate its polar coordinates $\langle \theta_p, \rho_p \rangle$ relative to C_i , then identify four consecutive sub-control points, $\langle \theta_{i,j}, \rho_{i,j} \rangle, j = k - 2, k - 1, k, k + 1$, such that $\theta_{i,k-1} \leq \theta_p < \theta_{i,k}$. The radius $R_{CR}(p, C_i)$ associated with p is thus the Catmull-Rom interpolation of $\rho_{i,j}, j = k - 2, k - 1, k, k + 1$ [CR74]. Thus, the NURBF corresponding to control point C_i is:

$$\Theta_{NURBF}(p, C_i) = \exp\left(-\frac{\|p - C_i\|^2}{R_{CR}(p, C_i)^2}\right).$$

Figure 3 shows a simple example of a NURBF in comparison with a RBF, demonstrating its flexibility and controllability in defining a non-uniform influence field.

4.2. Evolutionary Construction of NURBFs

Although NURBFs are able to capture the combined effect due to various factors that affect the influence of individual feature points, the number of sub-control points involved make it impractical to specify them manually, even for just one facial image. We again adopted a data-driven approach, and utilized a genetic algorithm to evolve NURBFs (i.e., their sub-control points) according to a target image, or a subset of images selected from the database based on an image to be transformed. The fitness evaluation is based on image difference.

For each image in the database, the NURBF is used to normalize the texture space according to the feature points of the image and those of the standard geometry vector $G_{\mathcal{R}}$. (For details, see [Hub07].)

Figure 4 shows an illustration of evolving a subset of 23 NURBFs (instead of 72 for visual clarity), each with 6 sub-control points, from a set of RBFs with the same radius. In this particular case, we used a known target image of a slightly rotated head to evolve the NURBFs. The results




Figure 4: Evolving an NURBF warping model for a specific input (a). Starting with a set of uniform RBFs in (b), the model gradually optimizes itself into a set of NURBFs in (e), which captures various factors that determine the influence of feature points. There is noticeable improvement after 1000 generation (η) with reduced MSE (ϵ) in (f) against the target (d). The MSE is calculated in the parameter space of MFM. For comparison, the results of two MATLAB warping methods are shown in (g) and (h).

clearly show that the evolution process is able to learn from the training image, and generates a set of NURBFs that capture the 3D nature of the transformation.

We have compared NURBF with a few traditional warping algorithms in Figure 4, including (g) *triangular mesh warping*, (h) *local weighted mean*, and (c) *RBF-based field warping*. Both (g) and (h) are supported by the MATLAB image-toolbox. (g) involves Delaunay triangulation and mapping between corresponding triangles using Barycentric coordinates. With (h), for each pixel in the target image, the algorithm finds the N closest feature points. It then uses the N feature points in both images to infer a second-order polynomial, which determines the coordinate mapping within the radius of this polynomial, which is the distance from the pixel to the furthest feature point used to infer the polynomial. (h) was generated with $N = 10$. In comparison with these three methods, as shown in Figure 4(f), NURBF exhibits significant improvement.

5. Modeling Faces and Age Transformation

5.1. Morphable Face Model (MFM)

The PCA-based face modeling has been used extensively in age simulation. We here give only a brief description for self-containment, and details can be found in [LTC02, HBHP03, SSG06] and [Hub07].

A facial image F is represented by two elementary components, namely *geometry* G and *texture* T . As illustrated in Figure 5(b), 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 .

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 PCA provides a highly effective means for modeling and transforming shapes and textures [BV99]. An arbitrary face $\langle G_a, T_a \rangle$ is thus encoded as:

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

where $\langle \bar{G}, \bar{T} \rangle$ be the mean of the basis faces in a reference set, and $\langle \Phi_g, \Phi_t \rangle$ be the two sets of n eigenvectors for geometry and texture respectively.

5.2. 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, $\langle \Delta G_i, \Delta T_i \rangle$, $i = 1, \dots, n$, 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 $\langle \Delta G_i - \overline{\Delta G}, \Delta T_i - \overline{\Delta T} \rangle$, $i = 1, \dots, n$, resulting in two sets of eigenvectors, $\{\psi_{g,i} \in \mathbb{R}^{2k} \mid i = 1, \dots, n\}$ and $\{\psi_{t,i} \in \mathbb{R}^{3l} \mid i = 1, \dots, n\}$. 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 person-specific 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. For person-specific aging simulation, it is necessary that each face pair is for the same person (or known to be closely related). Hence it is appropriate to use MATM for databases of connected datasets, such as FG-Net [FG-05], but not so for unconnected datasets, such as [HBHP03, SSSB07].

6. Approximating Age-Transformation

Eq. (1) in Section 5.2 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). In particular, we propose a completely new method for obtaining the missing information by learning from the relationship between faces and age transformation (in 6.3). In addition, several other variations of linear combination methods are described and compared in [Hub07].

6.1. Linear Combination of Age-Difference

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

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

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

Global Mean. Most 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 can 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*. As shown in Figure 5(c), the image resulting from the use of a global mean exhibits few facial features of the ground truth (unknown to the aging process). In comparison, Results (d) show the improvement in accuracy through localization. Using a smaller subset consisting of 20 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 ground truth. Figure 6 shows 10 such face pairs for the female example in Figure 5.

6.2. 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 face pair $(F_{a,i}, F_{b,i}), i = 1, \dots, s$ 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 $\langle \delta_{g,i}, \delta_{t,i} \rangle$, for $i = 1, \dots, s$.

One can linearly moderate each pair of $\langle \delta_{g,i}, \delta_{t,i} \rangle$ using the results of the above-mentioned comparison metric in 6.1. This approach is said to be *person-sensitive*, as different persons in the reference set (i.e., face pairs in \mathcal{F}) con-

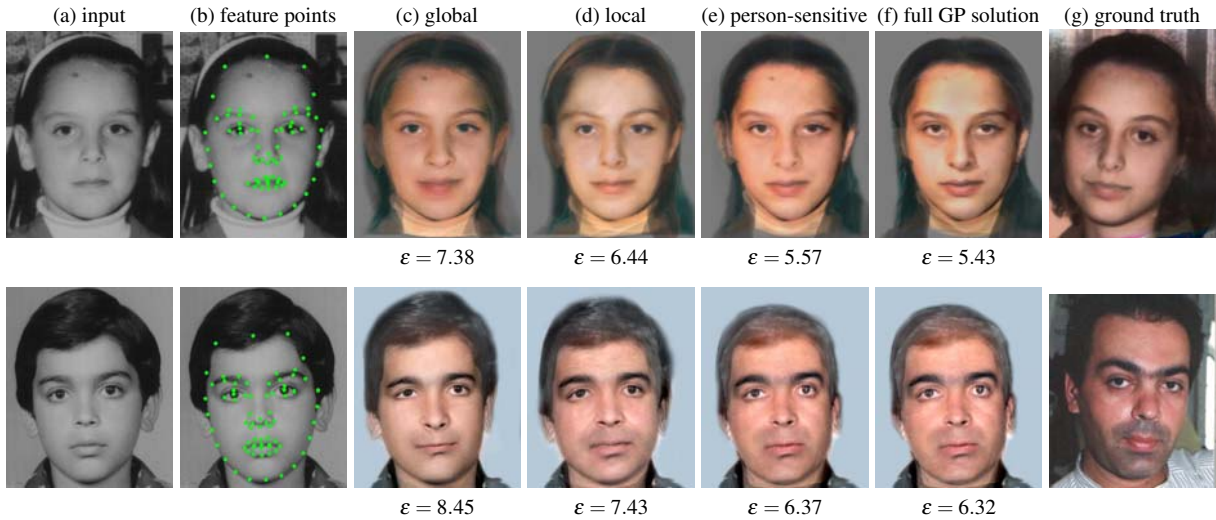


Figure 5: The data-driven framework enables an appropriate set of images to be selected according to an input image (a). Localization (d) show the relative merits over a global model (c). Person-sensitive linear combination (e) provides further improvement over localization. A full GP solution (f) gives the lowest level of MSE (ϵ) in comparison with the actual aged individual (g), which is not included in the training set.



Figure 6: A selection of 10 reference face pairs for the female example in Figure 5. The F_a and F_b sets are placed in two separate rows with corresponding face pairs in the each column.

tribute differently to the linear combination. Although the method is not exactly the same as the person-specific method by [SSG06], it fundamentally follows the same line of reasoning. While [SSG06] moderates the mean aging trajectory using the historical aging pattern prior to age a , we moderate the mean using the distances between the input and individuals in the reference set \mathcal{F} . The amount of the improvement over the straightforward local mean depends on the variation of the measured distances from the input face to faces in \mathcal{F} . For instance, the male (e) in Figure 5 shows more improvement over simple localization than the female (e), 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).

6.3. Evolving an MATM Parameter Set

Given $\langle G_t, T_t \rangle$, 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 face pairs selected according to $\langle G_t, T_t \rangle$. The two example results shown in Figure 5(f) clearly demonstrate the advantages of this approach.

6.3.1. Relation between Face and Aging

Consider a reference set, \mathcal{F} , of face pairs selected according to $\langle G_t, T_t \rangle$. 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 $\langle \alpha_{g,i}, \alpha_{t,i} \rangle$, for $i = 1, \dots, s$. Meanwhile from \mathcal{F} , we derive a local MATM, which is then used to encode the difference between each face pairs as $\langle \delta_{g,i}, \delta_{t,i} \rangle$, $i = 1, \dots, s$. We thus have s pairs of example mappings,

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




Figure 7: 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, $\eta = 100, 200, 500$.

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 F_g and F_t . Let:

$$\alpha_g = \{u_{g,1}, u_{g,2}, \dots, u_{g,m}\}, \quad \alpha_t = \{u_{t,1}, u_{t,2}, \dots, u_{t,m}\};$$

$$\delta_g = \{v_{g,1}, v_{g,2}, \dots, v_{g,m}\}, \quad \delta_t = \{v_{t,1}, v_{t,2}, \dots, v_{t,m}\}.$$

$m \leq s$ is the reduced dimension of MAM and MATM. We define F_g and 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 face pairs provide the evolution with a training set for evaluating the generated F_g and F_t in the process.

6.3.2. Genetic Programming

Genetic programming allows the automatic evolution of a computer program as a solution to an optimization problem, in our case the functional mapping from MAM to MATM. Such a computer program can be considered as a function (or a tree) with unknown formation and parameters. The goal of genetic programming is to establish the formation and parameters of the tree through a process of applying genetic operators to generate offspring. For example, 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. (For details, see [Hub07].)

In many ways, genetic programming is similar to an adaptive Monte Carlo method in computer graphics. Statistically, it will always reach a local best fit in a huge search space for functions. The resultant tree is thus a mathematical formula of functional mapping from MFM to MATM.

Figure 7 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.

7. Results and Discussions

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 8 shows one such test. In this example, the photo of a 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 9 shows a side-by-side comparison against the ground truth images. In addition to the direct comparison, we followed the practice of [LTC02] to extract the main facial features (i.e., *focus*) for comparison, and that of [SSSB07] to superimpose the clothing, hair and background (i.e., *context*) from the ground truth image onto the synthesized face. As shown in Figure 9, the focus extraction gives a perceptually better match than direct comparison, and the context addition gives the strongest impression of resemblance. Figure 9 also shows another set of the results for a female subject, and further example results can be found in [Hub07].

The performance of the learning-based approach 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 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 significant improvement in terms of time and resource requirements. The performance of other approaches, including person-sensitive, typically range from 1 to 5 minutes for a small reference set.

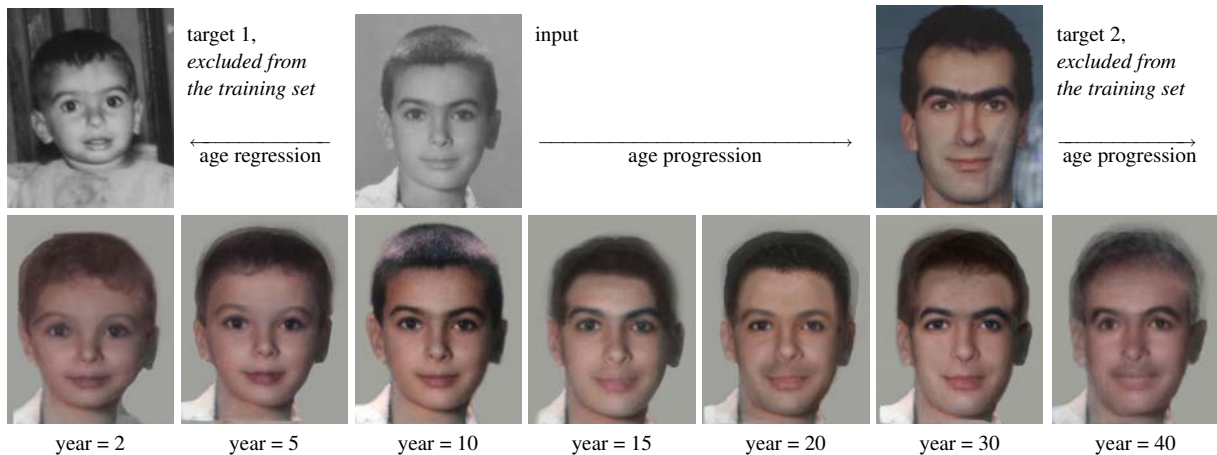


Figure 8: Age progression and regression for a given input of a male subject. Two target images are shown for comparison.

8. 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 to 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 [LTC02, HBHP03], the results obtained in this work represent a significant leap in terms of resemblance with target images. In comparison with [SSG06, SSSB07], our results are at least comparable if not noticeably better, while our methods do not involve the complexity of collecting historical imagery data, nor 3D models, of the subjects. These results also demonstrated that our approach is practically feasible, and conceptually close to the knowledge-based approach by human forensic artists. In fact, by not assuming an a-priori model for the mapping from MFM to MATM, genetic programming is the most appropriate and effective technique to be deployed in this work.

In this work, we found that age simulation for some race groups remains difficult due to the lack of suitable public domain data. However, the FG-Net consortium is continuing to increase its aging data repository, such an issue is expected to be alleviated in the near future. We also plan to implement more advanced computer vision techniques for processing images with more complex pose, and facial expression, and those with occlusions from hair, glasses and clothing.

Acknowledgment. Most photos 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 photos.

References

- [ABC*00] ANDRESEN P. R., BOOKSTEIN F. L., CONRADSEN K., ERSBÄYLL 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
- [BN92] BEIER T., NEELY S.: Feature-based image metamorphosis. *Computer Graphics* 26, 2 (1992), 35–42. 2
- [BP95] BURT D. M., PERRETT D. I.: Perception of age in adult caucasian male faces: Computer graphic manipulation of shape and color information. *Proc. 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, 5
- [CET01] COOTES T. F., EDWARDS G. J., TAYLOR C. J.: Active appearance models. *IEEE PAMI* 23, 6 (2001), 681–685. 3
- [CR74] CALMULL E., ROM R.: A class of local interpolated splines. In *Computer Aided Geometric Design*, Barnhill R., Riesenfeld R., (Eds.). Academic Press, 1974, pp. 317–326. 4
- [FG-05] FG-NET CONSORTIUM: Aging Database. <http://sting.cyccollege.ac.cy/~alanitis/fgnetaging/>, 2005. 2, 3, 6, 9
- [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. 2
- [Gol89] GOLDBERG D. E.: *Genetic Algorithms in Search, Optimization and Machine Learning*. Kluwer, 1989. 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 Transactions on Medical Imaging* 22, 6 (2003), 747–753. 1, 2, 3, 5, 6, 9
- [Hub07] HUBBALL D.: *Exploring the Relationship between Faces and Ageing in Image-based Age Transformation*. Master's thesis, Swansea University, 2007. Masters Thesis. 3, 4, 5, 6, 8

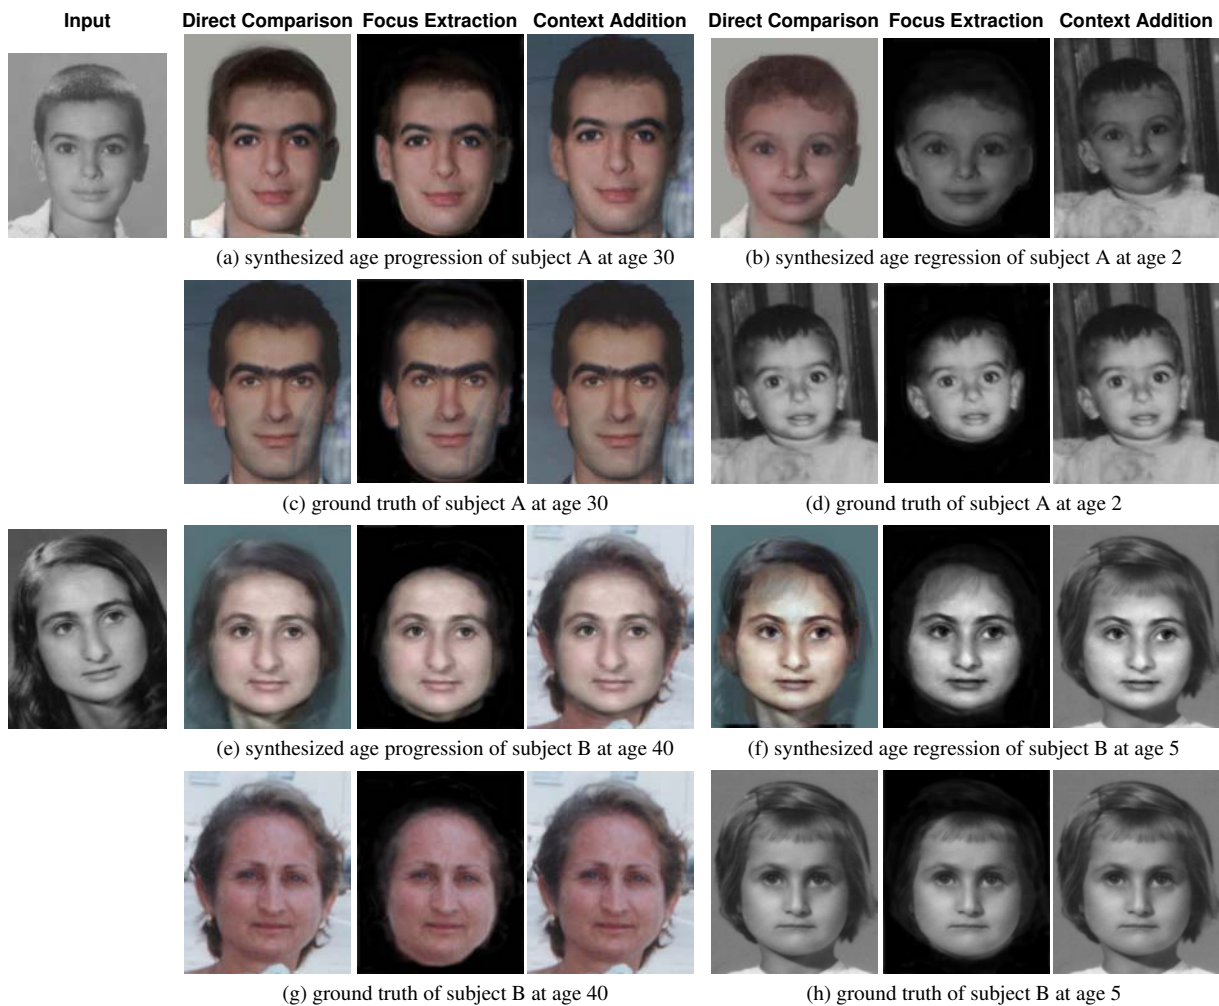


Figure 9: Comparison between synthesized images and the ground truth images in three different formats.

- [Koz92] KOZA J. R.: *Genetic Programming*. MIT Press, Cambridge, MA, 1992. 3
- [LGL95] LERIOS A., GARFINKLE C. D., LEVOY M.: Feature-based volume metamorphosis. In *Computer Graphics (Proc. SIGGRAPH 95)* (1995), Addison-Wesley, pp. 449–456. 2, 4
- [LTC02] LANITIS A., TAYLOR C. J., COOTES T. F.: Toward automatic simulation of aging effects on face images. *IEEE PAMI* 24, 4 (2002), 442–455. 2, 3, 5, 8, 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. 2
- [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
- [PS75] PITTENGER J. B., SHAW R. E.: Aging faces as viscoelastic events: Implications for a theory of nonrigid shape perception. *Journal of Experimental Psychology: Human Perception and Performance* 1, 4 (1975), 374–382. 1, 2
- [SD96] SEITZ S. M., DYER C. R.: View morphing. In *Proc. ACM SIGGRAPH* (1996), Addison-Wesley, pp. 21–30. 2, 4
- [SSG06] SCANDRETT C. M., SOLOMON C. J., GIBSON S. J.: A person-specific, rigorous aging model of human face. *Pattern Recognition Letters* 27 (2006), 1776–1787. 1, 2, 5, 7, 9
- [SSSB07] SCHERBAUM K., SUNKEL M., SEIDEL H.-P., BLANZ V.: Prediction of individual non-linear aging trajectories of faces. *Computer Graphics Forum* 26, 3 (2007), 286–294. 1, 2, 6, 8, 9
- [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. 4
- [Wol90] WOLBERG G.: *Digital Image Warping*. IEEE Computer Society Press, Los Alamitos, CA, 1990. 2