EG UK Theory and Practice of Computer Graphics (2007)
Ik Soo Lim, David Duce (Editors)

Towards Individualized Ageing Functions for Human Face Images

D. W. Hunter and B. P. Tiddeman

School of Computer Science. St. Andrews University. UK.

Abstract

We present an overview of work in progress to create ageing functions for the transformation of human face images. Our aims are twofold; Firstly, by extension to three dimensions, the problems of pose and lighting associated with photographs can be eliminated, and the shape changes associated with ageing modelled more accurately. Secondly, by using multi-variate statistical methods that can take individuality into account, the correlations between appearance and type of ageing can be learned and the most appropriate one automatically selected.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Physically based modeling

1. Introduction

Accurate prediction of how a person's appearance will vary with age has a variety of applications. Such as, aiding in the search for missing persons, planning cosmetic surgery, as well as applications in the movie industry and other visual arts. We intend, as part of current work in progress, to improve upon current two-dimensional methods by using a face-fitting technique to fit a statistical face-model, known as a Three Dimensional Morphable Model (3DMM) [BR02], to photographs of human faces. This will eliminate most of the problems associated with pose and lighting as well as approximate the three-dimensional shape of the subjects face. We intend to use this data to investigate various multi-variate statistical methods which we believe will provide improved ageing functions by finding correlation between appearance and the way in which an individual ages.

2. Literature Review

Researchers have attempted to model visual facial ageing in humans for some time now. The Cardioidal strain coordinate transform, has been used by a number of researchers [PS75] [PSM75] [MT83] [VB89]. In this method ageing is controlled by an arbitrary constant. Whereas more recent methods such as Lantis, Taylor and Cootes [LTC02] and Scandrett, Solomon and Gibson [SSG06] have used models that can be trained using a database of human subjects.

Benson and Perrett used a stratified method [RP95], dividing the face images into five separate age groups. They created an average for each strata, and for the overall population, using feature based warping and per-pixel blending. They used a simple warping technique applying the differences between the target and subject age-group average. In order to test the result, they asked a large number of volunteers to estimate the age of both the synthesised and original images. Experimental evaluation showed that both techniques produced a significant increase in the perceived age, although significantly less than the age difference between the original groups used to train the transform. Tiddeman et al [TBP01] later improved on both the capture of age in prototypes and the effectiveness of age transformations using a wavelet-based approach to retain facial textures.

Lantis et al. used a statistical appearance model to train an ageing function over a database of photographs. Their method involved delineating key features (eyes, ears, chin etc.) on a set of photographs. A shape average could then be computed by averaging the positions of the feature points. Intensity information was also sampled from with in the facial region. The shape and intensity data was combined into a single 'face vector.' Principle Component Analysis per-

© The Eurographics Association 2007.

formed on the covariance matrix of the face vector deviations was used to find the main axis of variations from the mean, leading to a compact parametric description of each face. They used a genetic algorithm to plot an ageing function through the PCA space. They couldn't carry out the inverse operation directly, that is, given an age and an ageing function estimate the appearance. Instead they adopted a simulation approach. They generated a look up table, by averaging all the face vectors corresponding to a certain age, so that the table contained the most typical face vector for that age range [LTC02].

More recently Scandrett et al. devised a method for incorporating a wide variety of different, but not necessarily independent factors, such as individual history or parental and sibling similarity. Ageing functions were modelled in a piece-wise linear fashion between age groups as vectors in the parameter-space. New functions were created as a linear combination of ageing vectors weighted according to a maximum-likelihood criteria. They validated their result visually as well as with mean-squared pixel error metric [SSG06].

Previous research into facial ageing functions have mostly concentrated on transformations in two dimensions. An exception to this is Hutton et al. [HBHP03], who used a linear warp between the average of two sets of 3D faces, using Kernel-smoothing to remove noise from data-set. The resulting 3D models where validated visually by the author.

3. Proposed Method

Our proposed method makes use of two databases for its calculations, one a set of 3D scans of individuals of a variety of ages, and the other a set of 2D images of individuals at multiple age points in their lives. The 3D scans are used to create a statistical model, containing the principle components or eigenfaces [TP91] of the scans. This model is used both to create 3D models from the 2D images and as a coordinate space to train the ageing function through.

3.1. Building a 3D Morphable Model

When working with two dimensional images a statistical model called an Active Appearance Model (AAM) is often employed. Introduced by Cootes et al. [CET98] a combined shape and colour model is created using PCA. In three-dimensions the separate 3D shape and 2D colour of the 3D Morphable Model has advantages as it is generally easier to fit to an image than a combined model.

The 3D Morphable Model (3DMM) developed by Blanz and Vetter, and explained in [BV99], describes the set of human faces using separate shape and colour components. The 3D positions of the vertices and RGB values of the texture are concatenated in two vectors:

$$\mathbf{S} = (X_1, Y_1, Z_1, X_2, Y_2, Z_2, \cdots, X_n, Y_n, Z_n)^T, \quad (1)$$

$$\mathbf{T} = (R_1, G_1, B_1, R_2, B_2, G_2, \cdots, R_m, G_m, B_m)^T \quad (2)$$

The shape S and colour T of a new face are generated as linear combination of weighted PCA vectors $\mathbf{S}_j$, $\mathbf{T}_j$ and the averages $\hat{\mathbf{S}}$ and $\hat{\mathbf{T}}$.

$$\mathbf{S} = \hat{\mathbf{S}} + \sum_{j=1}^{k} \alpha_j \mathbf{S}_j \quad (3)$$

$$\mathbf{T} = \hat{\mathbf{T}} + \sum_{j=1}^{k} \beta_j \mathbf{T}_j \quad (4)$$

The weights α_j and β_j form the parameter vectors α and β. New faces are created by varying these parameters.

This representation can also be used with other statistical representations such as Independent Component Analysis [LW99].

In order to perform any meaningful analysis a dense correspondence must be found between the faces in the 3D face model collection. Authors using depth-map based systems have been able to use 2D image based techniques such as optical-flow to find correspondences [BV99]. We are using a 3DMD face capture system [TDM] that outputs an unstructured polygon mesh. We adapt a 'standard' face model to each subject's face. The standard mesh topology we use is just an individual selected on the basis of the quality of the scan (i.e. all major features visible). We adapt this to each subject using a two stage warping process. The first stage is a feature based warping method, in which manually placed landmarks are used to drive a multi-level free-form deformation (MFFD), which is a hierarchy of B-spline interpolating functions with progressively finer resolution [LCS95] [LWS97]. We have implemented the MFFD warping using a space and search efficient octree data structure. Once the face meshes are in approximate alignment, rays are traced out of the standard mesh from each vertex in both directions and the first intersection (within a maximum radius) with the target face is found using an octree ray-tracing method. Not all vertices will find a target, and so these displacements are interpolated across the standard mesh. This brings the standard mesh into good alignment with the subject.

Rendering

The 3DMM's shape is projected onto a 2D surface using the standard method of graphics APIs such as OpenGL [Ope]

© The Eurographics Association 2007.

and Directx [Dir].

$$S_{2D} = M_{4x4}P_{4x4}(\hat{\mathbf{S}} + \sum_{j=1}^{k} \alpha_j \mathbf{S}_j) \quad (5)$$

where M is the combined rotation, scaling and translation, ‘Model-view‘ matrix and P is the perspective matrix as defined by the DirectX specification [Dir] incorporating focal distance and field of view. The 3D positions are converted to 4D vectors by padding the value 1 on the end of the vector. When the shape parameters α are altered by $\delta\alpha$ the resulting warp can be projected to a 2D surface similarly:

$$W(;\delta\alpha) = M_{4x4}P_{4x4}(\sum_{j=1}^{k} \delta\alpha_j \mathbf{S}_j) \quad (6)$$

3.2. Fitting a 3DMM to an image

In order to build a statistical model of how individuals age we need a number of snap-shots of each individual at various stages in their life. 3D scanners of sufficient quality have been developed only recently, so collections of scans of individuals at different ages do not exist. Waiting for individuals to age in order to rescan them is beyond the time frame of this project. Photography on the other hand has been in existence for over a century and photographs of individuals at different ages are relatively easily obtained. The proposed solution is to fit a generic 3D statistical model to the photographs and so obtain an approximate 3D model of each face at several different ages. A technique for obtaining these models has been developed by Blanz and Vetter [BR02], and has been extended by Zhang and Samaras to use spherical harmonics to describe the space of possible lighting conditions [ZS04]. Romdhani’s thesis [Rom05] and book chapter [SRV05] provides a detailed overview of the methods and introduces a method that uses a maximal-probability function. Thus the most probable result is taken each iteration. This system still offers the same advantages over 2D image analysis; variations in pose can be accounted for by using rotations and lighting effects on the 3D model. These rotation and lighting effects, which distort 2D image analysis, can help shape the 3D model.

Forward Additive Algorithm

In order adapt a 3DMM to an image we attempt to minimize a cost function

$$\chi^2 = \sum_{x \subset \Omega} (M(x;\mathbf{p}) - I(x))^2 \quad (7)$$

Where $I(x)$ is the image to fitted to $M(x;\mathbf{p})$ is image produced by rendering the 3DMM with parameters p. All parameters necessary for rendering the 3DMM, position, rotation, scale, perspective, shape and colour, are concatenated in p. Ω the subset of all samples in the image. This has previously been solved using an iterative gradient descent technique, with a suitable choice of update fucntion.

As a first attempt at solving (7) we took the same approach as Lucas and Kanade [LK81] and used a first-order Taylor series expansion of the update function:

$$M(\mathbf{p}+\delta\mathbf{p}) \approx M(\vec{p}) + \delta\vec{p} \cdot \nabla M(\vec{p}) \quad (8)$$

where

$$\nabla M(\mathbf{p}) = \left(\frac{\partial M}{\partial p_0}, \frac{\partial M}{\partial p_1}, \cdots, \frac{\partial M}{\partial p_N}\right) \quad (9)$$

is the gradient of the model with respect to each of the model parameters. By partially differentiating with respect to the parameters p_i we get the parameter update equation:

$$\delta\vec{p} = (\nabla^2 M)^{-1}(\nabla M \cdot (I - M)) \quad (10)$$

where

$$\nabla^2 M(\vec{p}) = \left\{\frac{\partial M}{\partial p_i} \cdot \frac{\partial M}{\partial p_j}\right\} \quad (11)$$

is an approximation of the *Hessian* matrix. This is a simplified form of the forward additive algorithm update step from [BM04]:

$$\Delta\mathbf{p} = H^{-1}\sum_x [\nabla I \frac{\partial \mathbf{W}}{\partial \mathbf{p}}]^T [M(x;\mathbf{p}) - I(x))] \quad (12)$$

where the Hessian H is defined as:

$$H = \sum_x [\nabla I \frac{\partial \mathbf{W}}{\partial \mathbf{p}}]^T [\nabla I \frac{\partial \mathbf{W}}{\partial \mathbf{p}}] \quad (13)$$

The difference is in the calculation of the derivative images $\frac{\partial M}{\partial p_i}$. A simple linear approximation can be used:

$$\frac{\partial M}{\partial p_i} = \frac{M(;\mathbf{p}+\delta_i\mathbf{p}) - M(;\mathbf{p}-\delta_i\mathbf{p})}{2} \quad (14)$$

where $\delta_i\mathbf{p}$ is a small change in the value of parameter i and zero elsewhere. Using the chain-rule to combine the image derivative ∇I with the parameter update warp (6) produces the *steepest-descent parameter update* images of [BM04]

$$\frac{\partial M}{\partial p_i} = \nabla I \frac{\partial W_i}{\partial p_i} \quad (15)$$

thus making (10) identical to (12).

In the case of parameters, such as lighting or colour the derivative images are not related to the image gradient and so the steepest-descent parameter updates are undesirable. In this case (14) is used. This method is applicable to 3DMM as the Hessian is updated each iteration, unfortunately this also makes the method very slow. In the case of two-dimensional AAMs a faster Inverse Compositional algorithm can be used, we are currently working on extending this to work with 3D Morphable Models.

Inverting the algorithm

As an on going project we are attempting to develop an inverted version of the fitting algorithm. The Inverse compositional algorithm, as described by Baker and Simon in [BM04], reverses the relationship between the target image

© The Eurographics Association 2007.

I and the rendered image M and attempt to warp the target image into the image rendered with the current parameters;

$$\chi^2 = \sum_{x \subset \Omega} [M(W(x;\mathbf{p})) - I(W(x;\delta\mathbf{p}))]^2 \tag{16}$$

where $W(x;\mathbf{p})$ is a warp function. The warp is the updated using a compositional update step;

$$W(x;\mathbf{p}) \leftarrow W(x;\mathbf{p}) \circ W(x;\delta\mathbf{p})^{-1} \tag{17}$$

Taking a first order Taylor series expansion and differentiating w.r.t to each of the parameters yields;

$$\delta\mathbf{p} = (\nabla^2 M(;0))^{-1} (\nabla M(;\mathbf{p})\dot{(}I - M(;\mathbf{p})) \tag{18}$$

The attraction of this algorithm is that the Hessian matrix, $\nabla^2 M(;0)$, is constant over all iterations, and thus can be pre-computed rather than updated each frame. In the case of 3DMM the composition step and it associated inversion of the warp presents difficulties. In [BPCM04] Baker et al. Showed that as the inverse composition algorithm does not hold for points off the surface of the model the algorithm is incorrect for surface based 3D models, such as a triangle mesh and thus needs to be modified. Ramdhani noted that (16) does not necessarily require that I be warped into $M(W(x;\mathbf{p}))$ but can be warped into another frame of reference. He also performed the inversion of the 3DMM warp by finding the triangle that would contain the point in the forward warp and recovering its relative position in the triangle on the reference frame.

GPU computation

The repetitive per pixel nature of the algorithm make it very suitable for parallel computation on Graphics Processor Units(GPUs) found in most modern PCs. The computation and rendering of the morphable model as well as the calculation of the derivative and error images are trivially achieved on a GPU as the only operations necessary are addition, subtraction and multiplication. The dot-products are harder to compute as a sum has to be computed over the entire image, and GPUs are designed for computation on individual pixels. Dot-products can be computed in logarithmic time using an algorithm called ping-ponging. The only part of the algorithm not computed on the GPU is the linear equation which is computed on the CPU using Singular Value Decomposition [PTVF92]. This means that the only data required to be read back from the GPU is the Hessian matrix and the dot-products of the error image and the derivative images. Notably shader versions 3.0 and below, GPUs have no ability to store the parameters used to update the morpable model and so these would have to be updated each iteration by the CPU.

3.3. Ageing Model

Here we look at two possible methods of impoving the generation of ageing functions. The first *Tensors* is aimed at finding correlations between identity on one hand, the face appearance at a given age. The second *Kernel PCA* is a method of creating non-linear regressions using methods similar to those for linear regressions.

Tensors

PCA is a two dimensional linear method, with the principle components being the eigen-vectors of a co-variance matrix. As the eigenvectors are data-driven they can miss correlations. Tenors are a n-dimensional method with each axis representing a different attribute, such as identity, age, expression etc. Tensors can be though of as an n-dimensional grid with each cell containing some data e.g. a face. A new face image can be synthesised by using a set of weights along each axis. The higher dimensionality of the tensor construction means that typically fewer parameters are required to describe the data to a desired level of accuracy. Although PCA cannot be trivially extended to n-dimensional tensors, Higher Order Singular Value Decomposition (HOSVD) acts as a reasonable approximation [VT05]. It has been used by Vlasic et al. [VBPP05] to build a model based on fourth order tensors comparing visemes, expressions and identity. It allows regression to be computed independently on each of the attributes in the model.

A tensor can be written as:

$$\sum_{i_N=1}^{J_N} \cdots \sum_{i_2=1}^{J_2} \sum_{i_1=1}^{J_1} a_{i_1,i_2,\ldots,i_N} \mathbf{u}_{i_1} \circ \mathbf{u}_{i_2} \circ \cdots \circ \mathbf{u}_{i_N} \tag{19}$$

and decomposed using HOSVD as

$$T \otimes_1 \mathbf{U}_1 \otimes_2 \mathbf{U}_2 \otimes_3 \cdots \otimes_N \mathbf{U}_N \tag{20}$$

where $\otimes_n$ is the mode-n product. New faces can be generated by replacing one of the $\mathbf{U}_n$ terms. For instance, if axis 1 contains expression data, replacing $\mathbf{U}_1$ will produce a new face image with a different expression but with none of the other parameters altered. Using one axis to represent identity and another to represent age would allow for an analysis of the relationship between two parameters independently of other attributes represented on other axis of the tensor. Finding suitable function to generate a replacement for the relevant $\mathbf{U}$ term is one of goals of this project.

The tensor construction assumes that there is some data in each cell and each new attribute-axis exponentially increases that amount of data that must be gathered. Adding an expression attribute to a tensor already containing age and identity, for example, would require a face model for each expression type at each age for each individual in the data-base. Furthermore the data-set we are working with has missing data for some of the individuals. A method of imputing the missing data is required.

Kernel Methods

It is known that the ageing function for any individual is not a straight line in PCA space [LTC02]. Scandrett et al. approximated this non-linear space in a piece-wise linear fasion, by finding ageing vectors between stratified age-groups. Kernel methods allow linear methods to be used in a non-linear context by providing a mapping from a non-linear data-space to a linear-space of higher or even infinite dimensionality. One such Kernel method, Support Vector Machines(SVM) has recently been used by [WL04] to synthesise shaped changes in a face due to ageing. Others such as Kernel-PCA could be used, with regression being formed on the data in a similar manor to linear PCA to find an ageing fucntion. Although care must be taken to choose an appropriate kernel.

Evaluation

Root mean squared pixel error (as used by Scandrett et al. [SSG06] for example) to determine the efficacy of an ageing model has the disadvantage that the two images of the same person at the same age can have quite different error values even after normalization. Perceptual evaluation, where a group of people are asked to rate an image according to its age or similarity to another image, provides an alternative to mathematical metrics. Human rating of the image has the advantage that differences between images that produce large pixel errors but are not significant to human perception can be discounted. Individual human raters can be somewhat subjective in the ratings they give, the effect of this can be reduced by using a number of different raters.

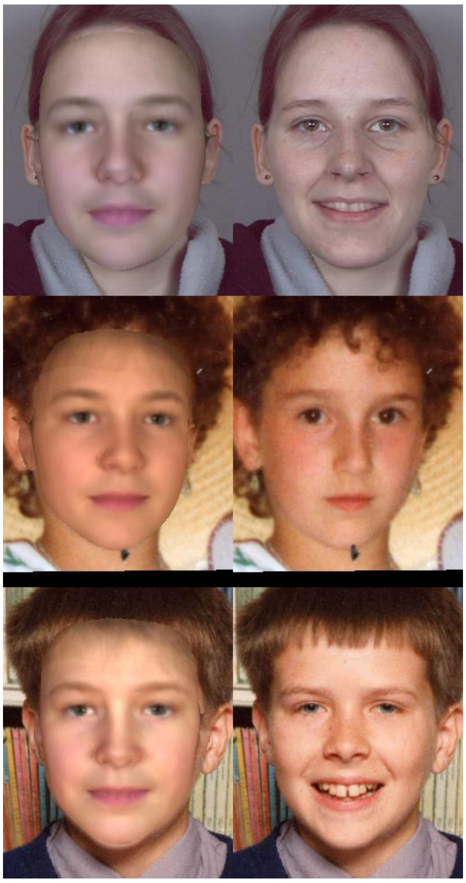

Results of fitting 3DMM to out-of-set face images using the forward additive algorithm. The left hand column shows the fitted 3DMM overlayed onto the original image. The right-hand column shows the original image. An affine warp has been used to bring the eyes of the central image into horizonal alignment.

Figure 1: *Fitting of 3DMM to images of out-of-set faces.*

4. Results

4.1. Face Fitting Results

In order to test the efficacy of the algorithms we are using we built a 3D Morphable Model using a training set of 167 individuals for varying age, sex and ethnicity (although this was moslty white European due to local demographics.) The first 20 principle components of each shape and colour were used, accounting for 90% of shape and 95% colour variation. The forward additive algorithm was then tested on out of set images (fig. 1). Provided the initial position and rotation parameters where positioned fairly close to the final parameters, the fitting proved relatively robust.

Tensor Ageing Results

Figure 2 shows the results of preliminary experiments using Tensors for face ageing on two-dimensional images. A set of face images was delineated using an Active Appearance Model (AAM) [CET98] and corrected by hand. The average of this set of points was calculated and the images warped into this average using a linear warp of the triangulated feature point template. The resulting image was masked so that only the face was visible, this was then vectorized and added to the image tensor. In order to find the parameters of an input face two methods where employed. Firstly a simple linear approximation, and secondly an Alternating Least Squares approach. In the ALS method one set of parameters is fixed and the remaining parameters are found using a Least Squares fit, this is then repeated, fixing a another set of parameters until the error between the tensor image and the input image is within a preset error margin. A face image was aged by setting the parameters of the tensor model relevant to ageing to the average of the target age group and warping the image accordingly. With the linear method the resulting image was too similar to the average image as shown in figure 2 as any individuality is smoothed out by the age transform. With the ALS method the individuality of the input image dominates over the desired age change. It was observed that a lack of standardization in the input image (e.g. lighting and pose) was a major contributing factor to these problems, so we hope that standardization using the 3D model fitting will help alleviate these problems.

© The Eurographics Association 2007.

Results of the synthesis of face images using Tensors. The top row is the reconstuction of a male image in the infant age group, the bottom row show the reconstuction of a female image in the student age group. The left column uses the linear method of reconstruction, the right column the ALS method.

Figure 2: *Ageing of faces using Tensor Method on 2D images.*

4.2. Conclusions and Future Work

Face fitting is slow and relatively inaccurate, at least for the purpose of ageing. An inverted method would be significanlty faster, possibly even real-time.

Much progress has been made in modeling facial ageing in humans and in devising statistical methods that can take into account the wide variety of factors that affect it. Functions based on linear warps between age-group strata provide a simple but effective method of synthsising the effects of ageing. Probability density functions to find the maximal-liklihood vectors can create new ageing functions from information about an individual, e.g personal growth history. Nevertheless, there is still significant room for improvement. An appropriate non-linear ageing function could improve the accuracy of the function when the desired age is between strata. New statistical methods such as HOSVD remain to be explored in this context, they may provide improved correlations between individual data and their ageing function.

References

[BM04] BAKER S., MATTHEWS I.: Lucas-kanade 20 years on: A unifying framework. *International Journal of Computer Vision 56*, 3 (March 2004), 221 – 255.

[BPCM04] BAKER S., PATIL R., CHEUNG K. M., MATTHEWS I.: *Lucas-Kanade 20 Years On: Part 5*. Tech. Rep. CMU-RI-TR-04-64, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, November 2004.

[BR02] BLANZ V., ROMDHANI S.: Face identification across different poses and illuminations with a 3d morphable model. In *FGR '02: Proceedings of the Fifth IEEE International Conference on Automatic Face and Gesture Recognition* (Washington, DC, USA, 2002), IEEE Computer Society, p. 202.

[BV99] BLANZ V., VETTER T.: A morphable model for the synthesis of 3d faces. In *SIGGRAPH '99: Proceedings of the 26th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 1999), ACM Press/Addison-Wesley Publishing Co., pp. 187–194.

[CET98] COOTES T. F., EDWARDS G. J., TAYLOR C. J.: Active appearance models. *Lecture Notes in Computer Science 1407* (1998), 484–??

[Dir] : *www.microsoft.com/windows/directx/default.mspx*.

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

[LCS95] LEE S.-Y., CHWA K.-Y., SHIN S. Y.: Image metamorphosis using snakes and free-form deformations. In *SIGGRAPH '95: Proceedings of the 22nd annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 1995), ACM Press, pp. 439–448.

[LK81] LUCAS B., KANADE T.: An iterative image registration technique with an application to stereo vision. In *IJCAI81* (1981), pp. 674–679.

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

[LW99] LIU C., WECHSLER H.: Comparative assessment of independent component analysis, 1999.

[LWS97] LEE S., WOLBERG G., SHIN S. Y.: Scattered data interpolation with multilevel b-splines. *IEEE Transactions on Visualization and Computer Graphics 3*, 3 (1997), 228–244.

[MT83] MARK L., TODD J.: The perception of growth in three dimensions. *Perception and Psychophysics 33*, 2 (1983), 193–196.

[Ope] www.opengl.org.

[PS75] PITTENGER J., SHAW R.: Aging faces as viscolelastic events: Implications for a theory of nonrigid shape perception. *J. Experimental Psychology: Human Perception and Performance 1*, 4 (1975), 374–382.

[PSM75] PITTENGER J., SHAW R., MARK L.: Perceptual information for the age level of faces as a higher order

© The Eurographics Association 2007.

invariant of growth. *J. Experimental Psychology: Human Perception and Performance 5*, 3 (1975), 478–493.

[PTVF92] PRESS W., TEUKOLSKY S., VETTERLING W., FLANNERY B.: *Numerical Recipes in C*, 2nd ed. Cambridge University Press, Cambridge, UK, 1992.

[Rom05] ROMDHANI S.: *Face Image Analysis using a Multiple Feature Fitting Stategy*. PhD thesis, University of Basel, 2005.

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

[SRV05] SAMI ROMDHANI J.-S. P., VETTER T.: *3D Morphable Face Model, a Unified Approach for Analysis and Synthesis of Images*. Elsevier, 2005, ch. 4.

[SSG06] SCANDRETT C. M., SOLOMON C. J., GIBSON S. J.: A person-specific, rigorous aging model of the human face. *Pattern Recogn. Lett. 27*, 15 (2006), 1776–1787.

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

[TDM] : *www.3dmd.com*.

[TP91] TURK M. A., PENTLAND A. P.: Face recognition using eigenfaces. In *Computer Vision and Pattern Recognition, 1991. Proceedings CVPR '91., IEEE Computer Society Conference on* (1991), pp. 586–591.

[VB89] V. BRUCE M. BURTON T. D.: Further experiment on the perception of growth in three dimensions. *Perception and Psychophysics 46*, 6 (1989), 528–536.

[VBPP05] VLASIC D., BRAND M., PFISTER H., POPOVIĆ; J.: Face transfer with multilinear models. *ACM Trans. Graph. 24*, 3 (2005), 426–433.

[VT05] VASILESCU M. A. O., TERZOPOULOS D.: Multilinear independent components analysis. *Proc. Computer Vision and Pattern Recognition Conf.* (2005).

[WL04] WANG J., LING C. X.: Artificial aging of faces by support vector machines. *Lecture Notes in Computer Science: Advances in Artificial Intelligence 3060* (2004), 499–503.

[ZS04] ZHANG L., SAMARAS D.: Pose invariant face recognition under arbitrary unknown lighting using spherical harmonics. In *ECCV Workshop BioAW* (2004), Maltoni D., Jain A. K., (Eds.), vol. 3087 of *Lecture Notes in Computer Science*, Springer, pp. 10–23.

© The Eurographics Association 2007.