VISUAL FACIAL AGEING USING PLS

Visual Ageing of Human Faces in Three Dimensions using Morphable Models and Projection to Latent Structures

D. W. Hunter and B. P. Tiddeman
School of Computer Science, Jack Cole Building, North Haugh, St. Andrews, KY16 9SX, U.K.

Keywords: Face, Facial ageing, Morphable models, Statistical modelling, Partial least squares.

Abstract: We present an approach to synthesising the effects of ageing on human face images using three-dimensional modelling. We extract a set of three-dimensional face models from a set of two-dimensional face images by fitting a Morphable Model. We propose a method to age these face models using Partial Least Squares to extract from the data-set those factors most related to ageing. These ageing related factors are used to train an individually weighted linear model. We show that this is an effective means of producing an aged face image and compare this method to two other linear ageing methods for ageing face models. This is demonstrated both quantitatively and with perceptual evaluation using human raters.

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 film industry and other visual arts. Since most researchers have concentrated on manipulating 2D images, 3D statistical models are a relatively recent innovation. In this paper we develop 3D models of ageing by fitting a Morphable Model to a set of photographs and introduce a new statistical ageing model based on Projection to Latent Structures (PLS) also known as Partial Least Squares.

2 LITERATURE REVIEW

Most previous methods for ageing a facial image have concentrated on transforming a 2D image. Cardioidal Strain was an early method that relied on the similarity between the mathematical function and facial ageing in children (Pittenger and Shaw, 1975; Pittenger et al., 1975; Mark and Todd, 1983; V. Bruce, 1989). This was later used in a modified form by Ramanathan and Challappa (Ramanathan and Chellappa, 2006). Rowland and Perrett used Triangulated Linear Warping to define an ageing trajectory by the average prototypes for two age-groups (Rowland and Perrett, 1995). Lanitis et al. trained a statistical model over a set of face images parametrised by a Principle Components Analysis model (Lanitis et al., 2002). Scandrett et al. also used PCA on a set of 2D images, ageing them using a piecewise linear model, combining the ageing trajectories between age-groups with an historical ageing trajectory from younger images of the individual (Scandrett et al., 2006). Suo et al. explored a different approach by describing the face using a Grammatical Model, (Xu et al., 2005) consisting of a hierarchical set of face components. An input face was aged using a Dynamic Markov Chain (Suo et al., 2007).

The idea of Modelling ageing using 3D models has been around for some time. Mark and Todd applied Cardioidal strain to a 3D model (Mark and Todd, 1983), Hutton and Buxton used Kernel Smoothing to create an ageing model of a set of 3D models (Hutton et al., 2003). More recently Scherbaum el al. (Scherbaum et al., 2007) fitted a Three-dimensional Morphable Model to a database of laser scanned cylindrical depth-maps. They used these models to train a Support Vector Regression model, synthesized a new face mode from by 'stepping' through the curved SVR space using a fourth order Runge-Kutta algorithm. The curved nature of the SVR model meant that the ageing paths were different depending on the parameters of the input face thus creating a semi-individualised model. However they had only one sample per subject to train the model

W. Hunter D. and P. Tiddeman B. (2009).
VISUAL FACIAL AGEING USING PLS - Visual Ageing of Human Faces in Three Dimensions using Morphable Models and Projection to Latent Structures.
In *Proceedings of the Fourth International Conference on Computer Vision Theory and Applications*, pages 340-345
DOI: 10.5220/0001793403400345
Copyright © SciTePress

and so captured population variations and not necessarily the variations due to ageing in a particular individual. Park et al. (Park et al., 2008) fitted a three-dimensional Morphable Model to a set of face images by fitting an Active Appearance Model and extracting a three-dimensional model from the AAM. Ageing was performed by calculating a set of weights between an input face and exemplar faces in the same age group. These weights are then used to build an aged face as a weighted sum of the corresponding faces at the target age.

Since many of the statistical methods used lost textural detail such as wrinkles, a few researchers developed methods that attempted to create appropriate textural detail in aged images. Tiddeman et al. used a wavelet transform (Tiddeman et al., 2001) and Markov Models (Tiddeman et al., 2005), Hussein used Bidirectional Reflectance Distribution Functions (Hussein, 2002) and Gandhi used Gaussian filters (Gandhi et al., 2004). These methods work by attempting to replace or adjust the high-frequency components of the image to match the high-frequency components of a prototype at the target age.

3 OVERVIEW

Our aim is to be able to take an image of a particular person and to create an ageing trajectory specific to that particular individual. Using a set of 3D face models we first separate those factors most related to ageing. Given a training-set of 3D models containing a 'snap-shot' of a number of individuals at various age points, from childhood to early-adulthood, we then train a set of ageing trajectories for each individual. Finally these trajectories are applied as a weighted sum of trajectories from the training-set.

3D data-sets featuring face models from the same individual at various age points are rare and incomplete, however 2D data-sets are more readily available. We therefore opted to use a face-fitting method to extract a 3D Morphable Model (Blanz and Vetter, 1999) from a two-dimensional image. We obtained a set of photographs by asking some student volunteers to supply images from a number of key ages. The resulting image set was divided into three strata, Mid Child containing individuals aged 5 to 8 year, Late Child covering 8 to 12 year-olds and a Student age-group between 17 and 23 years. The data-set contained 35 individuals, with one face model per individual in each strata.

A set of three-dimensional face models is required to construct the Morphable Model. We captured a set of 106 face models from volunteers using a 3dMD scanner (http://www.3dmd.com). The individuals ranged in age from 2 to 65 years old, the average age was 22.7 with a standard deviation of 17.45 years.

In this paper, we first briefly outline a process by which we generate the 3D models. We then describe and compare three ageing mechanisms. One, based on average Prototypes is the 3D analog of the 2D method used by Rowland and Perrett (Rowland and Perrett, 1995). The second, an Individualised Linear model, is the 3D analog of work by Lanitis at al (Lanitis et al., 2002) and is similar to the method of Park et al. (Park et al., 2008). We introduce a new technique based on Partial Least Squares (Wold, 1966).

4 THREE DIMENSIONAL MORPHABLE MODELS

Three-dimensional Morphable Models introduced by Blanz and Vetter use Principle Components Analysis to describe the space of human faces as a set of orthogonal basis vectors. Given a set of 3D dimensional face models with a one-to-one correspondence between vertices, we vectorise the vertex positions and colour values and centre each face by subtracting the mean of all the faces. PCA is then performed on the shape and colour values separately to produce a set of basis vectors. A reduced set of 40 eigenvectors for each of shape and colour were used to describe the face space, denoted $\mathbf{s}_j$, $\mathbf{t}_j$ respectively. The shape $\mathbf{s}$ and colour $\mathbf{t}$ of a new face are generated as,

$$\mathbf{s} = \hat{\mathbf{s}} + \sum_{j=1}^{k} \alpha_j \mathbf{s}_j, \ \mathbf{t} = \hat{\mathbf{t}} + \sum_{j=1}^{k} \beta_j \mathbf{t}_j \qquad (1)$$

where $\hat{\mathbf{s}}$ and $\hat{\mathbf{t}}$ are the averages of the shape and colour respectively. The weights α_j and β_j form the parameter vectors α and β, which we concatinate to form $\mathbf{p} = \alpha, \beta$. New faces are created by varying these parameters. This process is described in more detain in (Blanz and Vetter, 1999).

4.1 Fitting a Morphable Model to a Face Image

Three-dimensional scanning equipment is a relatively recent invention, and so databases of three-dimensional models of the same individual taken over a period of many years have yet to be built. However two-dimensional images, in the form of photographs are widely available. In order to build a set of face models we attempt to extract three-dimensional information from these images. Our fitting method was a simple adaptation of the Lucas Kanade Tomasi algorithm (Baker and Matthews, 2002) from two-dimensional face models to three-dimensional models, this method is similar to that detailed by Blanz and Vetter (Blanz and Vetter, 1999). We use a Taylor series expansion of the l^2-norm of the pixel difference between an input image and the rendered Morphable Model to find the parameters that minimise this difference. To improve the accuracy of the fitting a set of delineated feature points on the two-dimensional image are also matched to their corresponding points on the Morphable Model using the l^2-norm of their separating distance when projected onto the image plane. The result of the fitting operation is a set of vectorised shape and colour parameters $\mathbf{p}$ that describes the face contained in the two-dimensional input image as a three-dimensional Morphable Model. Figure 1 shows an example of results of fitting a Morphable Model using this technique.

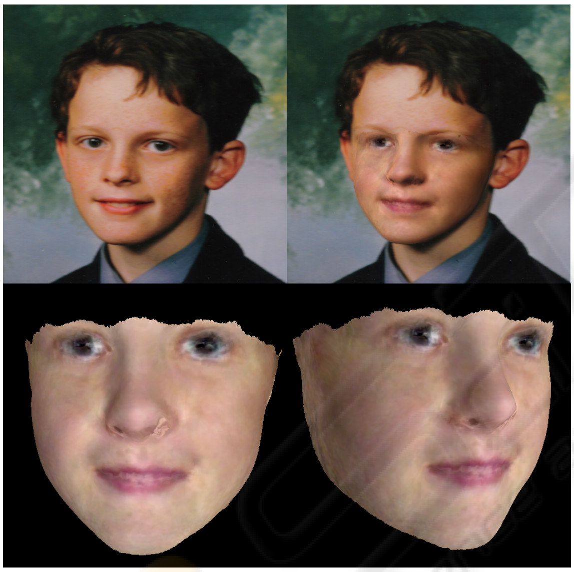

Figure 1: An example of a Three-dimensional Morphable Model fitted to a face image. The image on the top left is the original photograph. The image on the top right shows the results of the fitting, rendered in-situ. The images on the bottom row are of the same model rendered under neutral lighting conditions from different angles.

5 AGEING METHOD

We applied the face-fitting method outlined in the previous section to the photographs in the training set, to produce a set of 3D models of each individual at multiple age-points. We now use this training set to create an ageing model.

5.1 Age Prototypes

Prototype face-models were created for each age-stratum by averaging the parameters over all faces in the stratum.

$$\hat{\mathbf{f}}_{\mathbf{s}} = \sum_{i}^{m} \mathbf{p}_i \tag{2}$$

where $\hat{\mathbf{f}}_{\mathbf{s}}$ are the parameters of the averaged face model of all the faces in the stratum s. Here $\mathbf{p}_i$ is the vector of parameters for the i^{th} face model in the stratum and m is the number of faces in the stratum.

A linear transform is defined between from stratum j to stratum k we take as,

$$\mathbf{t} = \frac{\hat{\mathbf{f}}_{\mathbf{k}} - \hat{\mathbf{f}}_{\mathbf{j}}}{\hat{a}_k - \hat{a}_j} \tag{3}$$

where $\hat{a}_j$ and $\hat{a}_k$ are the average ages of the individuals within strata j and k respectively. An input $\mathbf{f}_{in}$ in stratum j is aged towards the age group of stratum k by moving it in the direction of the vector t and multiplying t by the desired number of years.

$$\mathbf{f}' = \mathbf{f}_{in} + (a_t - a_s)\mathbf{t} \tag{4}$$

where $\mathbf{f}'$ is the set of model parameters at the target age, a_s and a_t are ages of the input face and the target age respectively. Clearly this transform is most valid if the target age is within the range of years of the target stratum k.

5.2 Individualised Linear Transform

It is well known that faces do not age in an identical manner. In order to generate an ageing trajectory for an unseen individual we exploit the relationship between appearance and ageing trajectory. For each individual in the data-set a linear ageing path is defined as a vector from one sample face in the starting stratum to another in the target stratum containing the end age. If no suitable pair of sample faces can be found the individual is excluded from the data-set. We denote s,e as the start and end ages of the transform respectively, and $\mathbf{p}_i$ and $\mathbf{q}_i$ as the parameters of the face models of the i^{th} individual taken from the start and end strata respectively. We define a single linear ageing function such that the j^{th} parameter of the face model of the individual i at time t is,

$$\mathbf{f}(\mathbf{t})_j = t.\mathbf{a}_{i,j} + \mathbf{b}_{i,j} \tag{5}$$

where $\mathbf{a}$ and $\mathbf{b}$ are sets of weights and $\mathbf{a}_{i,j}$ and $\mathbf{b}_{i,j}$ are the j^{th} weights for the i^{th} individual in the training set. $\mathbf{a}$ defines the gradient of the path in $\Re^n$ and $\mathbf{b}$ the parameters of the face at time $t = 0$. These are defined as,

$$\mathbf{a} = \frac{\mathbf{q} - \mathbf{p}}{e - s}, \mathbf{b} = \mathbf{p} - s\mathbf{a} \tag{6}$$

These functions can be parametrised using $\mathbf{a}_i$ and $\mathbf{b}_i$ to describe the ageing function f_i for the i^{th} individual. A new ageing path for an unseen individual can be created using a linear weighted sum of the parameters of the ageing functions for each individual in the training set.

$$\mathbf{f}' = \sum_i^n \rho_i \mathbf{f}_i, \qquad \sum_i \rho_i = 1 \tag{7}$$

where ρ_i are a set of weights relating the unseen individual to the ageing path of the i^{th} individual in the data-set. The ρ_i's sum to one, so that that function does not add a scaling factor to the ageing path.

As in (Lanitis et al., 2002) the weighting ρ is defined using the probability distribution of the PCA space of the face model.

$$\rho(\mathbf{p}_{in}, \mathbf{p}_i) = e^{-\sum_j^n \frac{(p_{in,j} - p_{i,j})^2}{2\sigma_j^2}} \tag{8}$$

where $\mathbf{p}_{in}$ and $\mathbf{p}_i$ are the parameters of the input and i^{th} face model respectively. $p_{in,j}$ is the j^{th} parameter of the input face model. σ_j^2 is the variance of the PCA space in the j^{th} dimension.

This is similar to the method by (Park et al., 2008). Equation (7) can be combined with equation (1) to derive their method. Ours differs in that the weights are based on the PDF of the Morphable Model rather than linear interpolation.

5.3 Partial Least Squares Ageing

The data-set of parameters contains a significant amount of information that is not relevant to ageing. Any statistical analysis needs to separate those factors related to ageing from those that are not related either explicitly or implicitly.

Partial Least Squares (Wold, 1966) also known as a *Projection to Latent Structures* is a statistical distribution similar to Principle Components Analysis that describes mean centred data as a weighted linear combination of basis vectors. Unlike PCA, which finds directions of maximum variance in the data, PLS attempts to describe a set of dependent variables from a set of predictors. It works by extracting a set of latent vectors that decompose both the dependent and the independent matrices in such a manner as to explain as much of their covariance as possible.

We take the parameters of the face models in the data-set $\mathbf{f}_i$ and use them to build the matrix $X = [\mathbf{f}_1, \mathbf{f}_2, \ldots, \mathbf{f}_n]^T$ such that each row contains the parameters of an individual face model. We define $Y = [age_1, age_2, \ldots, age_n]^T$ where age_i is the corresponding ages to the i^{th} face. The rows of both X and Y are then mean centred and scaled by the inverse standard deviation $\frac{1}{\sigma}$

As described by (Abdi, 2007), we aim to decompose the independent variables as $X = TP^T$ with $T^T T = I$. T is the *score* matrix and P is the loading matrix. We *estimate* Y as $\hat{Y} = TBC^T$. The diagonal matrix B holds the regression weights, and C is the weight matrix of the dependent variables. See (Abdi, 2007) for further details on what these mean in practice. The columns of T are the latent vectors that form an exact decomposition of X but only an approximation to Y. The decomposition is found using an iterative algorithm whereby, each iteration, a latent vector is found that maximizes the covariance between X and Y and is then subtracted from both. The proportion of variance explained by this vector is found by dividing the sum of squares of the residuals by the the sum of squares of the input matrices X and Y.

PLS, like PCA, can be truncated such that a smaller number of basis vectors are found that approximately span the space of X. We found that the first 6 latent vectors explained 56.3% of the variance and showed little improvement in accuracy thereafter. So we trucated the PLS space to 6 latent vectors.

We separated the parameters into two components; the components most related to ageing and a remainder. As the data used to train the PLS model has been converted to Z-scores by centring the data on the mean and scaling by the standard deviation, we must convert the parameters of the input face $\mathbf{f}$ to Z-scores also. We denote the Z-score converted face as $\bar{\mathbf{f}}$. The parameters of a face model in PCA space are related to the parameters of the face in PLS space as $\dot{\mathbf{f}} \approx \mathbf{g}P$. Since the loading matrix P is not generally orthogonal in PLS regression, $\mathbf{g}$ is approximated using least squares regression,

$$\mathbf{g} = (P^T P)^{-1} P \bar{\mathbf{f}} \tag{9}$$

The PCA face model parameters can be recovered from the PLS space as $\bar{\mathbf{f}}' = \mathbf{g}P$ and converted from Z-scores to the original PCA parameter space using $\mathbf{f}' = \bar{\mathbf{f}}'\sigma + \hat{\mathbf{f}}$.

In general the recovered $\bar{\mathbf{f}}' \neq \bar{\mathbf{f}}$, so we compute the residual $\mathbf{r}$ as $\bar{\mathbf{f}} = \mathbf{g}P + \mathbf{r}$. Ageing is performed using the Individualised Linear ageing Transform described earlier on the PLS model parameters ($\mathbf{g}$) instead of the PCA model parameters ($\mathbf{p}$). After the face is aged the residuals $\mathbf{r}$ are added back in.

The results of ageing a face model using these methods is shown in figure 2.

5.4 Quantitative Evaluation

In order to determine the comparative effectiveness between different methods of ageing we used the Ma-

Figure 2: Examples of aged face images. The fisrt column shows the original model, the second the same individual at the target age, the remaining columns show, left to right, the original model aged using, Prototyping, Individual Linear transforms and the PLS method.

halanobis distance to measure the similarity between the aged face model and a face model captured by fitting the Morphable Model to a image of the same individual at the target age. We used leave-one-out cross-validation to evaluate the methods. Figure 1 shows the results of ageing using the prototyping, individualised linear and PLS methods. We can clearly see that the 'Individual Linear' method gives an improvement in accuracy over the 'Prototyping' method with a lower average error and the PLS method shows a marked improvement over both.

Table 1: RMSE between shape and colour parameters of aged face model and a known ground-truth model for each individual in the data-set. With 93 subjects for each method.

Ageing Method	RMSE	Standard Deviation
Prototyping	8.86	1.84
Individual Linear	8.69	1.92
PLS	7.4	1.4

5.5 Perceptual Evaluation

Quantitative measures may miss ageing cues that human raters would be able to detect. We performed a series of tests with human raters to evaluate the ability of the methods to produce images of the required age.

Each user was shown a single image of a rendered face model at a time and asked to estimate the age of the face shown. The age is selected from a range between 5 and 30 to the nearest year. The stimuli are a selection of mid-child faces aged to student age by the three-methods, prototyping, individualised linear and PLS, together with the rendered face-models of the individuals at the source and target age. The images were presented with uniform lighting and pose on a black background and in random order. No peripheral details such as hair were on display, limiting ageing cue to those in the face. The images were presented on public website which generated a significant amount of traffic, with an average of 105 age estimations per image, and just under 5000 for each ageing method being trialled.

Table 2 shows the mean perceived age in years for the face models aged by the different methods, as well as the mean ages of the rendered models of the original face models. Table 3 shows the mean age difference in years between the perceived age of the individual after the ageing method is applied and the target age the algorithm was attempting to recreate. We can see that all the methods succeed in ageing the faces towards the target age, but vary in how much they age the face model. The PLS method achieved the closest results to the target age of all the ageing methods. The original student and mid-child age groups showed a significant error implying that some age related information was lost by the fitting process.

Table 2: Mean (μ) and standard deviation (σ) of the human rated ages for faces ages by each method.

Ageing Method	μ	σ	Count
Prototyping	17.048	6.7605	5090
Individual Linear	16.801	6.8449	5092
PLS	17.115	6.6780	4987
Student	17.026	5.9044	6205
Mid Child	12.762	6.0626	4678

Table 3: Mean (μ) and standard deviation (σ) of the error in years in human rated ages for faces ages by each method.

Ageing Method	μ	σ	Count
Prototyping	-3.6614	6.1888	4596
Individual Linear	-3.8674	6.1688	4646
PLS	-3.5643	6.1098	4551
Student	-3.3737	5.4273	5855
Mid Child	6.2135	6.0815	4678

6 CONCLUSIONS

We have described a method of ageing 3D Morphable Models by a method based on *Projection to Latent Structures* or *Partial Least Squares*. This method shows an improvement over the others tested both in quantitative measures, in terms of similarity to a known ground-truth, and in perceptual evaluation by human raters. Due to its reliance on face-fitting methods the success of this method depends on the quality of the face model produced in the fitting stage.

Improved fitting techniques or a database of three-dimensional scans of the same person over several year, would improve the accuracy of these ageing methods. Other authors have used Quadratic and Cubic functions (Lanitis et al., 2002) in two-dimensions, or non-linear Kernel methods such as Support Vector Regression (Scherbaum et al., 2007) in three-dimensions, so an obvious extension is to examine non-linear individualised ageing paths.

ACKNOWLEDGEMENTS

This work was supported by Unilever PLC and the EPSRC. We would also like to thank David Perrett for providing the 2D photographic data-set, and for his advice on conducting perceptual evaluation.

REFERENCES

Abdi, H. (2007). Partial least square regression (pls regression). In *In N.J. Salkind (Ed.): Encyclopedia of Measurement and Statistics.*, pages 740–744. Thousand Oaks (CA): Sage.

Baker, S. and Matthews, I. (2002). Lucas-kanade 20 years on: A unifying framework: Part 1. Technical Report CMU-RI-TR-02-16, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA.

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

Gandhi, M. R., Levine, M. D., and G, M. R. N. (2004). A method for automatic synthesis of aged human facial images. Technical report, Masters thesis, McGill University,2004. 1.

Hussein, H. K. (2002). Towards realistic facial modeling and re-rendering of human skin aging animation. In *SMI '02: Proceedings of the Shape Modeling International 2002 (SMI'02)*, page 205, Washington, DC, USA. IEEE Computer Society.

Hutton, T. J., Buxton, B. F., Hammond, P., and Potts, H. W. W. (2003). Estimating average growth trajectories in shape-space using kernel smoothing. *Medical Imaging, IEEE Transactions on*, 22(6):747–753.

Lanitis, A., Taylor, C. J., and Cootes, T. F. (2002). Toward automatic simulation of aging effects on face images. *IEEE Trans. Pattern Anal. Mach. Intell.*, 24(4):442–455.

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

Park, U., Tong, Y., and Jain, A. K. (2008). Face recognition with temporal invariance: A 3d aging model. In *FGR06*.

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

Pittenger, J., Shaw, R., and Mark, L. (1975). Perceptual information for the age level of faces as a higher order invariant of growth. *J. Experimental Psychology: Human Perception and Performance*, 5(3):478–493.

Ramanathan, N. and Chellappa, R. (2006). Modeling age progression in young faces. In *IEEE Computer Vision and Pattern Recognition or CVPR*, pages I: 387–394.

Rowland, D. A. and Perrett, D. I. (1995). Manipulating facial appearance through shape and color. *IEEE Computer Graphics and Applications*, 15(5):70–76.

Scandrett, C. M., Solomon, C. J., and Gibson, S. J. (2006). A person-specific, rigorous aging model of the human face. *Pattern Recogn. Lett.*, 27(15):1776–1787.

Scherbaum, K., Sunkel, M., Seidel, H.-P., and Blanz, V. (2007). Prediction of individual non-linear aging trajectories of faces. In *The European Association for Computer Graphics, 28th Annual Conference, EUROGRAPHICS 2007*, volume 26 of *Computer Graphics Forum*, pages 285–294, Prague, Czech Republic. The European Association for Computer Graphics, Blackwell.

Suo, J., Min, F., Zhu, S., Shan, S., and Chen, X. (2007). A multi-resolution dynamic model for face aging simulation. In *IEEE Computer Vision and Pattern Recognition or CVPR*, pages 1–8.

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

Tiddeman, B., Stirrat, M., and Perrett, D. I. (2005). Towards realism in facial image transformation: Results of a wavelet mrf method. *Comput. Graph. Forum*, 24(3):449–456.

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

Wold, H. (1966). Estimation of principal components and related models by iterative least squares. *Multivariate Analysis*, pages 391–420.

Xu, Z., Chen, H., and Zhu, S.-C. (2005). A high resolution grammatical model for face representation and sketching. In *CVPR '05: Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05) - Volume 2*, pages 470–477, Washington, DC, USA. IEEE Computer Society.