Photo-Realistic Exemplar-Based Face Ageing

Andreas Schneider[1]*, Ghazi Bouabene[1], Ayet Shaiek[2], Sandro Schönborn[1]†, Frédéric Flament[2], Ghislain François[2], Virginie Rubert[2], Thomas Vetter[1]
[1] Departement of Mathematics and Computer Science, University of Basel, Switzerland
[2]L'Oréal Research and Innovation, Chevilly-Larue, France

***Abstract*— We propose a photo-realistic method for artificially ageing facial photographs by combining learned shape deformations with skin detail transfer between a donor and a receiver face. Facial ageing is a complicated process that most existing face models, such as 3d Morphable Models fail to express, due to lacking correspondence between the wrinkles of different individuals. We propose an exemplar-based approach to face ageing, where we transfer high-frequency details from an older face texture to a younger receiver. By warping the resulting image according to a learned shape ageing deformation, we obtain photo-realistic aged photographs. We evaluate the simulator with human perception experiments showing that we can indeed create results that are perceived to be real.**

I. INTRODUCTION

As we all know the passing of time influences the shape and texture of our face. The cheeks start to sag, more and deeper wrinkles appear, the amount of pigmentation spots increases. These effects dramatically change the appearance of a face. They are caused by a combination of environmental, metabolic and genetic factors[1], [2]. Some of them we can influence with behaviour, choice of life-style, cosmetics, and medical interventions. However, the effect they have on face appearance is unknown. Artificially simulating the course of ageing would give an ageing baseline to compare the influence of different interventions on face ageing. By comparing face ageing according to simulation with ageing under the influence of an intervention allows us to determine its efficacy. In this work we are concerned with creating the baseline face ageing simulation that ages an input face such that it is perceived older or younger.

Creating a photo-realistic face age progression model is difficult because it requires the modelling of details. The more faces age, the more details tend to appear such as wrinkles, pores, and pigmentation. Many of these features have in common that they do not appear at the same locations on the face of different individuals. For this reason it is difficult to capture them in models that are based on population averaging such as 3d Morphable Models and Active Appearance Models. Figure 2 highlights this by showing an example of face ageing of Morphable shape and texture models compared to ageing where we also transfer details. The approach in this work is to model the details in an exemplar-based way where we use exemplars of older faces that already contain details, and then transfer them to the receiving target face. Vice-versa in the case of rejuvenation: we transfer the details of a younger face to the older face. The approach of transferring details has the advantage, that we can create photo-realistic faces. This choice of method reflects our decision for photo-realism in favour of being person-specific.

978-1-7281-0089-0/19/$31.00

We have a collection of donor faces consisting of 3d shape and texture. Depending on the desired target age, an older or younger donor face is selected out of the collection. To transfer the details, we first establish correspondence between donor and receiving image by fitting the 3d Morphable Model to the input image. Then, by blending the higher-frequencies of the corresponding donor and receiver textures, skin details such as wrinkles and pigmentation either appear or disappear from the receiver face, depending on the target age. Identity preservation is ensured by excluding the eyes and mouth regions from the blending process. We then simulate the coarse shape deformation associated with ageing as a linear shape deformation in Morphable Model space. We learn the deformation from age-annotated, corresponding face scans using linear regression.

We evaluate the model with respect to perceived artificiality and perceived age. To show, that the model is realistic we ask human judges to distinguish between real and simulated images. To show that we can age faces by a specified amount, we ask our judges to estimate the age of the original and manipulated faces.

The resulting ageing simulation is capable of modelling the coarse shape deformation associated with ageing as well as the fine details that tend to appear with age. We show that the simulations are difficult to distinguish from real images and their target age correlates well with the perceived age.

II. RELATED WORK

An early work on artificial face ageing by Burt and Perrett[3] combines a shape and texture transformation without having an explicit face model by building a set of mean face images of different age groups called face prototypes. To age an input face, they add the difference between a younger and an older prototype to it. They validated their approach with a perception experiment and found that the faces manipulated to look older are indeed perceived as such. However, they also found that the perceived age is systematically underestimated and attributed this to the smoothing of the skin when building the mean faces which resulted in a loss of fine details such as small wrinkles. The same problem of detail loss occurs in all approaches that

*andreas.schneider@unibas.ch
†sandro.schoenborn@ch.abb.com, ABB Corporate Research, Baden, Switzerland

build a statistic over corresponding points of the face, such as Morphable Models and Active Appearance Models (AAM). They are able to model coarse ageing deformations on the face shape and texture such as the deepening of the naso-labial folds, growth of the nose and ears and sagging of the cheeks and neck. The first Morphable Model paper[4] formulates face ageing as a direction in face space that can be determined by linear regression. Non-linear shape ageing methods include the piecewise linear approach with AAMs[5] and some with more complex functions using support vector regression [6]. Lanitis et. al.[7], [8] investigate person specific polynomial ageing functions for artificial face ageing with Active Appearance Models. Independently of the regression method, all these models fail to express facial details such as wrinkles, pores, pigmentation spots, and other higher-frequency details that are not corresponding between subjects. Ageing methods capable of expressing details either transfer higher-frequency parts between images or have a different assumption on correspondence between faces. The quotient image approach [9] implicitly transfers details by building ratios between aligned images. These ratios can then be used to transfer expressions[10], makeup[11] and age faces [12], [13]. For ageing they either build ratios between images of the same person aged or between the mean image and an older face. Methods that explicitly transfer high-frequency details decompose an exemplar donor image into higher and lower-frequency bands. The exemplar-based methods for makeup transfer[14] and for ageing [15] warp a donor face image to a receiver image in 2d and then transfer the high-frequency details between them. Similarly, details can be removed by reducing high-frequency parts to rejuvenate faces[16].

Statistical models that do not assume strict point-to-point correspondence are able to model facial details according to a learned statistic. A parametric model for freckles[17] models freckle position on the face according to a stochastic process with learned parameters. Using a hierarchical face model[18] Suo et al. [19], [20] model the age progression of face details in multiple resolutions on 2d images. More recently, neural network based approaches to image synthesis and face ageing have appeared [21], [22], [23]. For example, Generative Adversarial Networks conditioned on the target age with identity preserving loss functions create a relatively realistic aged version of an input face[24], [25].

In contrast to GAN based methods, our pipeline is composed of a small number of clear steps giving the user full control over the ageing manipulation. Many of its parameters have semantic meaning, where we can for example limit the effect to a particular region of the face. Also, changes in our method can be attributed to either shape or texture, whereas this is more difficult to discern in GANs. Finally, our method can easily accommodate facial photographs taken in different conditions such as pose or illumination, using the same models.

III. PHOTO-REALISTIC AGEING PIPELINE

Different types of changes happen in the face due to ageing. We distinguish between detail changes that tend to be local and coarse changes that tend to be more global. We model the coarse changes of the shape with the Morphable Model shape model. To model the details, we transfer the higher-frequencies from a selected donor texture to the target face to be aged.

Figure 1 shows the different steps involved in ageing a face depicted on a single input photograph. First, we reconstruct pose, shape, texture, illumination and camera parameters of the face image. Knowing all these parameters gives us exact correspondence between the pixels and the 3d shape. To age an input face by the desired amount of years we transfer the skin details of a selected donor. We blend the high-frequency part of the donor texture with the high frequency part of the receiver texture. We render the original shape with the combined texture according to reconstructed pose and camera parameters. We then age the shape by the desired amount of years by projecting the 3d ageing deformation onto the image plane and warping the rendering accordingly. Additionally, we shade the image to better account for the shape difference.

In the following we describe all these steps in detail, starting with shape ageing.

IV. MORPHABLE MODEL AND SHAPE AGEING

3d Morphable Models (3DMM) [4], [26] model the variations of the shape and texture of human faces by leveraging established correspondences between a reference face and the remaining faces in the training data set. The training data set in our case is comprised of 323 three dimensional scans of female subjects of all ages containing both shape and texture.

Having such correspondences allows us to compute the deformation fields warping the reference face into each face in the training set. By applying Principal Component Analysis (PCA) on the obtained deformation fields, we can retrieve the basis functions defining the linear spaces representing facial variation in terms of shape and texture. Each face in our training data set can then be represented by its coefficients in the linear shape and texture spaces.

In addition to a quantified exploration of the shape and texture variability of human faces, such a model has several applications. One typical application is to reconstruct the full 3d shape and texture of a target face from a single 2d photograph. This process, called model fitting, consists of searching for the best shape and texture coefficients as well as the face pose, camera and illumination parameters. Such a model fitting can be performed using Markov Chain Monte Carlo sampling[27].

Another example application is to perform data analysis, or regression in the shape and texture space when provided with some meta-data associated with each face in the training set. In the following section, we will perform such a regression to model the shape variation according to the face's perceived age.

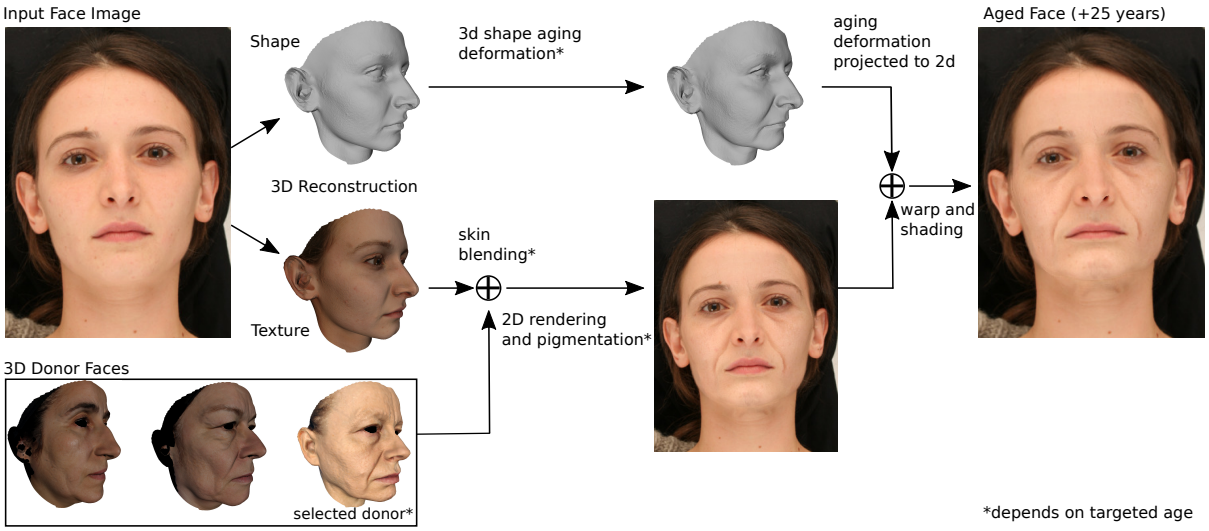

Fig. 1: Face ageing simulation overview. Ageing of an input face image by 25 years. Ageing of shape according to Morphable Model statistic and of texture according to selected exemplar face. 1) Reconstructing input face shape and illumination. 2) deforming shape by amount of years 3) blending of the input skin texture with an individual at target age. 4) add pigmentation details to the blended texture and render in 2d 4) 2d warping of rendered image according to ageing deformation. 5) adding shading difference caused by ageing deformation.

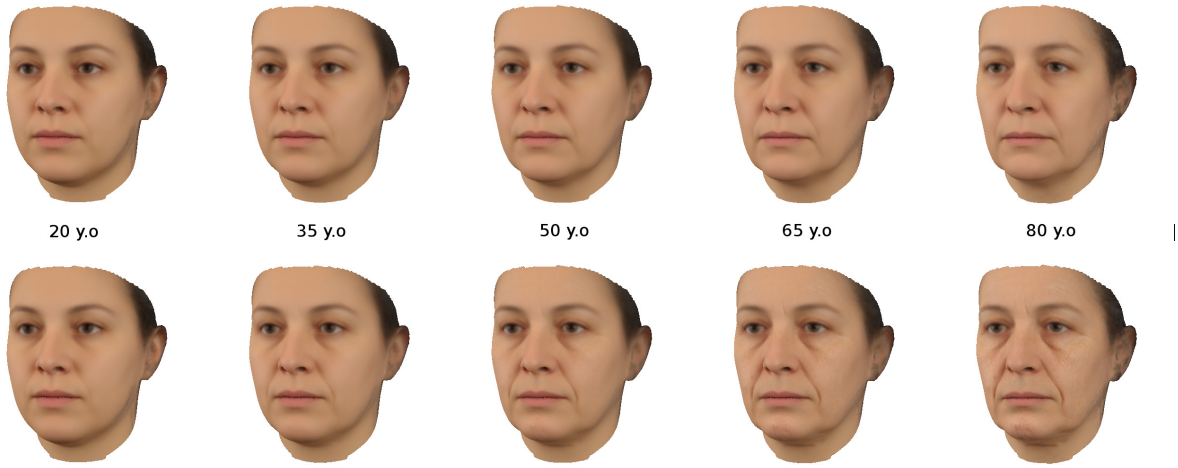

Fig. 2: Top row: Age progression in Morphable Model space. Only coarse shape deformations and colour changes of the texture can be expressed within the Morphable Model representation. Details such as wrinkles, skin irregularities and pigmentation spots are missing. Bottom row: exemplar-based ageing allows to age a face with fine skin details.

To build the Morphable Model used in this study, we used 323 three dimensional scans, captured with DI3D stereo-photogrammetry system, composed of three stereo pairs of 18 mega-pixel digital stills cameras. The scanned subjects are all Caucasian females aged between 18 and 70 years old, living in the same region. The focus on female subjects is motivated by the general scope of the research project that is cosmetics applications for the Caucasian female population. This choice should however not affect the generalizability of the methods presented here to more diverse data sets.

A. *Ageing Deformation*

In addition to the the shape and texture coefficients of each face in the training data set, we are provided with the perceived age associated with each face obtained by averaging the age estimations of a cohort of judges. Note that the perceived age was used here over the real one since the goal of the study is to artificially change the perception of age in manipulated pictures. Given the strong deviations between the perceived and real age for many subjects of our data set, the perceived age was chosen as the better reference to evaluate the delta in age perception incurred by our method.

Given this data set, we can find the hyper-plane fitting at

best the age values in the shape coefficients space, therefore giving us the linear direction in shape space corresponding to a change in the perceived age. To this end, we perform a ridge regression where we look for a $\vec{w}$ minimising:

$$\|\mathbf{X}\vec{w} - \vec{\alpha}\|^2 + \lambda\|\vec{w}\|^2 \quad \lambda \in R \tag{1}$$

where $\mathbf{X}$ is the matrix containing the shape coefficients of our training data, $\vec{\alpha}$ is the vector containing the perceived age labels, I the identity matrix and λ is a parameter determining the degree of regularization.

The minimizing $\hat{\vec{w}}$ is then given by:

$$\hat{\vec{w}} = (X^T X + \lambda I)^+ X^T \vec{\alpha} \tag{2}$$

Having obtained the hyperplane's orientation, we can now age a face with shape coefficients $\vec{x}$ by β years by adding multiples of the normalised vector $\hat{\vec{w}}$ to it.

$$\vec{x}' = \vec{x} + \beta \frac{\hat{\vec{w}}}{\|\hat{\vec{w}}\|^2} \tag{3}$$

where the normalisation is required to obtain the shape deformation corresponding to an ageing of exactly one year in terms of perceived age.

Figure 1 show the effect of applying 25 years of ageing to an example 3d scan. As you can see in the figure, this shape deformation mainly captures the coarse effects of ageing such as the deepening of the naso-labial folds, growth of the nose and ears and the sagging of the cheeks and neck but fails to model fine details such as small wrinkles.

V. Photo-realistic Shape Ageing

Given a 2d photograph, the first step in our pipeline is to retrieve its 3d shape and texture coefficients by fitting the 3DMM to the target image as depicted in Figure 1. With the obtained 3d shape, we can now artificially age the face shape according to a desired age change as we did in the previous section.

In order to obtain a 2d photograph matching the age manipulation, one option would be to render the aged 3d shape in 2d. This is indeed possible as we estimated all the parameters of shape, texture, pose, camera and illumination during the fitting phase that are necessary to render such an image [27]. The result of such a rendering would however lack photo-realism as the 3DMM texture instances tend to be too smooth therefore giving an artificial appearance. Instead, we project the 3d ageing deformation onto the 2d image plane using a transform mapping the mesh vertices onto image coordinates. To perform the projection, we follow transforms that are typically used in computer graphics where we project the vertices according to a pinhole camera model onto the sensor of the camera (more details in Schönborn[27]).

By warping the original photograph according to the projected warp field we can then obtain the desired shape deformation while maintaining photo realism.

Figure 3.d shows the projected deformation field obtained from ageing the target image in Figure 1. The results of warping the target image according to the projected 2d deformation field can be seen in Figure 3.e. As you can see, some artefacts appear on the boundary of the face due to the fact that the 2d deformation field is defined only on the face region as pictured in Figure 3.d.

To avoid such a rough change in the warping field, we extrapolate the projected deformation field using the method of function interpolation described in Poisson Image Editing [28]:

$$\min_f \int\int_\Omega |\nabla f|^2 \quad \text{so that} f|_{\delta\Omega} = f^*|_{\delta\Omega} \tag{4}$$

Where Ω is the image region outside the face region, $\delta\Omega$ is the boundary of Ω consisting in the union of the border of the face region with the border of the image, f^* is the projected deformation field defined over the face region and f is the interpolated deformation field over Ω. The obtained deformation f will then be a smooth deformation field with values at the border $\delta\Omega$ equal to those of f^* at the same border. Figure 3 shows the extrapolated deformation field and the result of warping the image accordingly therefore containing less artefacts. To finish the 2d shape warping we need to account for the difference in shading between the two shapes. Given that we have changed the 3d shape, the resulting shading of the original and the aged shape differs in a rendering with the same illumination setting.

Accounting for this shading difference is very important in our case since it is an important cue for shape in 2d. A corrected shading would in our case provide better visual hints to changes such as the deepening of the naso-labial folds, growth of the nose or sagging of the neck and cheeks. To compute the shading difference, we render the reconstructed face under the retrieved illumination, and, we render the deformed reconstructed face under the same illumination. We have a Lambertian reflectance model and we approximate the environment map with the first three bands of spherical harmonics. We do not model self-shadowing or inter-reflections as would be possible with Morphable Model PRT[29]. We then compute the difference between the two renderings therefore only retrieving the difference in shading between the two images. Finally, we add the difference image to the warped 2d image by adding the pixel values.

Figure 3 shows the result of adding the shading difference to the warped image. As you can see, the resulting image already shows a good coarse effect of the ageing. Compared to a rendering, the face now has the details of the input face. However, it still lacks the details that would appear due to ageing.

VI. Texture Ageing with Skin Detail Transfer

Until now we have neither modelled fine details of shape nor texture. Since we cannot use the Morphable Model to directly add such details, we will obtain them from the texture of an older face containing the desired details. To create an age progression series for a face, we first select a person whose age is close to the desired maximum target age of the series. This person's face will donate the details

encoded in the high-frequency part to the face we want to age. We assume, that the details appear gradually with increasing age. Therefore, as the age is increasing in the age progression, we increase the amplitude of the higher frequencies of the details, taking the full amount at the maximum target age. Additionally, to keep the identity, we do not transfer details at the position of the eyes, nose, and mouth.

The texture transfer process has four components. The selected donor texture t_d, the receiver texture t_r, the amount of donor amplitude to transfer $q(\alpha)$ that depends on age α, and the binary mask m indicating positions in the texture x that receive transfer.

The resulting texture t_α is obtained according to the formula below where the operator $\oplus$ represents blending of the image left and right of it. This operator will be detailed in the upcoming section.

$$t_\alpha(x) = \big(t_{\mathrm{r}}(x)(1 - q(\alpha)m(x))\big) \oplus \big(q(\alpha)m(x)t_{\mathrm{d}}(x)\big) \quad (5)$$

where

$$q(\alpha) = \left| \frac{\alpha - \alpha_r}{\alpha_d - \alpha_r} \right| \quad (6)$$

with α_r being the age of the receiver and α_d the age of the donor.

A. Detail Blending

Here we explain the details of the operator $\oplus$, blending the details of the donor into the receiver image. First, we decompose both texture images into a Laplacian pyramid[30][31], P_r for the receiver and P_d for the donor, where the size of the pyramid levels is halved each time. Each level of the pyramid represents a frequency band from which we can reconstruct the input image.

To transfer the details, we mix the two pyramids at the corresponding levels representing the higher frequencies according to the blending weights in equation 5. The resulting pyramid P of the blended texture image t_α looks like this at level l:

$$P_l(x) = P_{r_l}(x)(1 - q(\alpha)m(x)) + P_{d_l}(x)(q(\alpha)m(x)) \quad (7)$$

where P_{r_l} and P_{d_l} are the level l of the Laplacian pyramid of the receiver and donor textures, respectively. Since the details we are interested in are high frequency ones, only the five largest levels of the two pyramids are merged. For the remaining levels of P, the original corresponding levels of P_r are taken. The final texture is then obtained by collapsing the resulting pyramid P.

B. Pigmentation

In addition to wrinkles, ageing spots are part of the details we wish to add to our manipulated image. The amount of such pigmentation spots tends to increase with age since the face is exposed to the sun for a longer time. The blending approach described earlier already transfers a certain amount of such spots from the donor to the receiver and gradually blends them in. However, not only the strength, but also the amount of blending spots increases with age. To make the transfer more realistic, we add pigmentation spots in an amount proportional to the target age. To this end, we use a variant of the freckle model of Schneider et. al[17], where, instead of freckles, we have a dictionary of ageing spots, that we then sample at random positions on the receiver face. We determine the relation between the age and the amount of spots by counting detected ageing spots from a data set of 70 scans varying by age. We found a roughly linear relationship between the age α and the amount of ageing spots c.

$$c = \alpha * 1.59 + 143.11 \quad (8)$$

The large offset can be explained by the difficulty to distinguish between ageing spots and freckles as they are similar in colour and shape.

VII. EXPERIMENTS

Figure 4 and Figure 5 show examples of simulating ageing and rejuvenation of face images by different amounts. We now validate the ageing model with perception experiments. We measure how real synthesised images are perceived, and how accurately the proposed model follows age perception.

A. Experimental Setup

1) Dataset: In both experiments we use the same data set. It consists of 200 faces in total, of which 27 are real and 173 artificially aged. We age 25 face images and transfer the details from different donor images aged between 20 and 70 years. To validate the model's capability of simulating ageing and rejuvenation, we age them by 9 steps of 5 years each, between the ages of 20 and 65, depending on their starting perceived age. To make sure that the participants can concentrate on the face itself, and that they are not influenced by hair and clothing, we crop the images at the face region boundary and set everything else to white. Figure 6 shows an example of the interface shown to the judges, where the eyes have been anonymised for the purpose of the paper.

2) Judges: Human perception was collected from 80 judges. 46% of the judges were aged between 20 and 39 y.o and 54% aged between 40 and 59 y.o, all presenting good visual acuity. Images of the data set are displayed randomly. Two sessions are proposed with 100 images per session and a break of 2 minutes between the two blocks.

B. Perception of Realism

First, we measure how real the simulated images are perceived. Images generated from a realistic ageing simulation should be difficult to distinguish from real images. We show the images to the 80 judges and ask them to tell if the image is real or not. They select from two options either real or fake. An average of 69,5% of the simulated images and 77,4% of the non simulated images are found to look real by the 80 judges. In the following, we test if the two averages are significantly different. The null-hypothesis H_0 is that the two averages p_1 and p_2 are equal and the alternative H_A that they are different. We use a binomial test and compute the

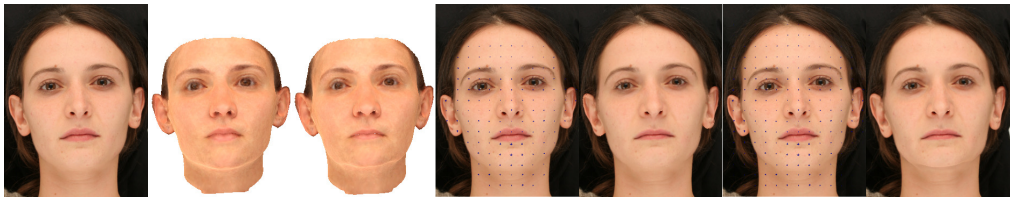

(a) (b) (c) (d) (e) (f) (g)

Fig. 3: Different steps of shape ageing. From the target image we reconstruct face shape and extract the texture. (a) input, (b) 3d rec., (c) shape aged, (d) warp field on face, (e) warped input, (f) warp field extrapolated outside the face, (g) warped.

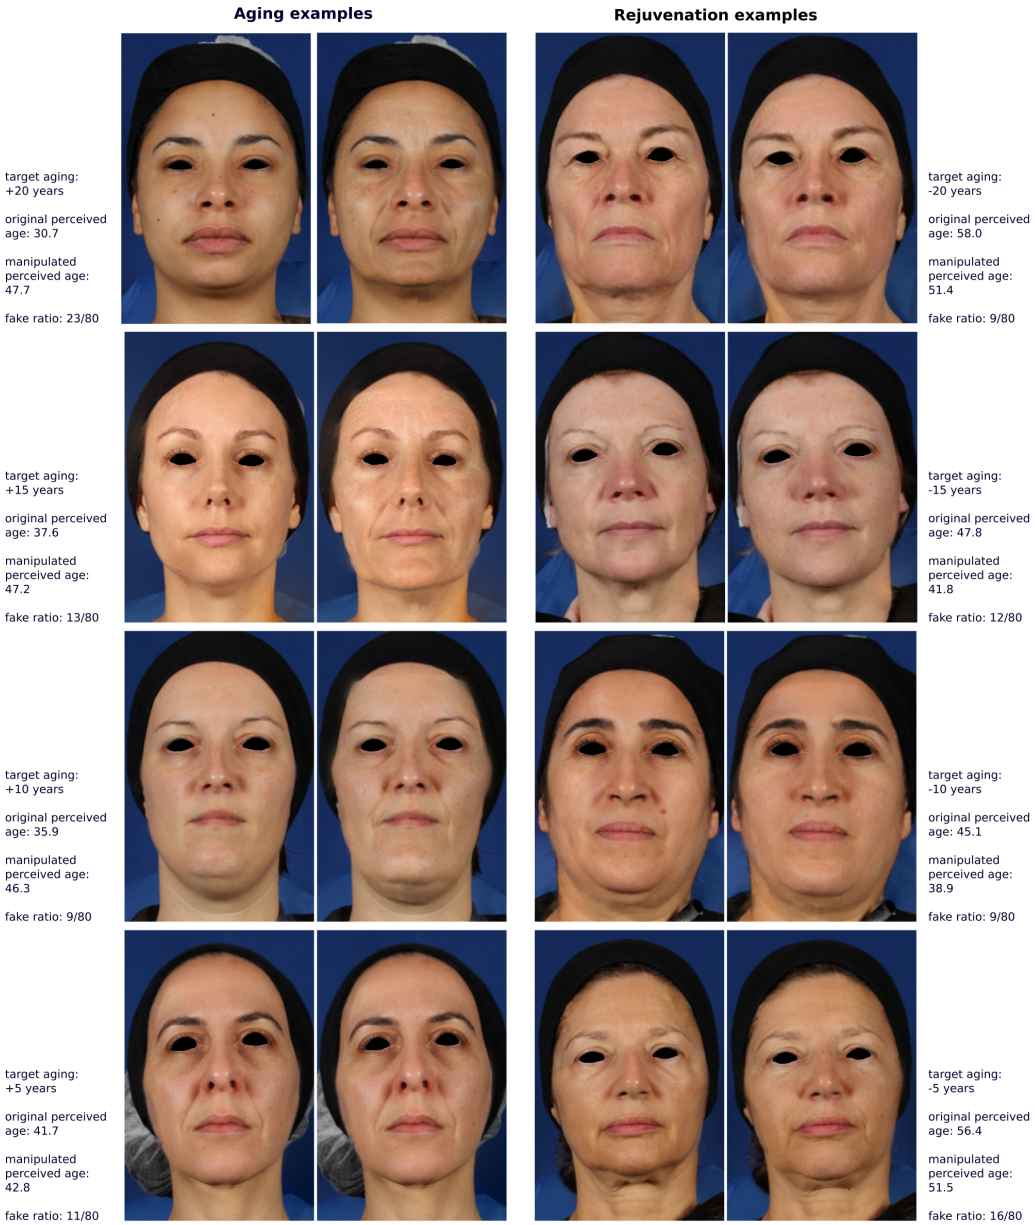

Fig. 4: Examples of ageing and rejuvenation manipulation with corresponding perception results. Left image of each pair is the original image. On the right the simulated image. The fake ratio indicates how many of the judges recognised them as being artificial.

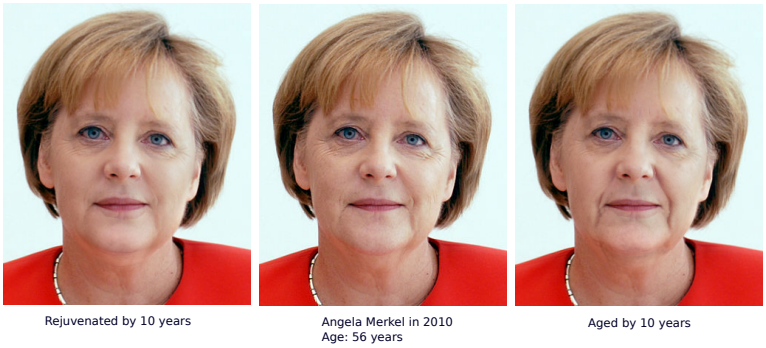

Fig. 5: Examples of ageing and rejuvenation manipulation of a known face.

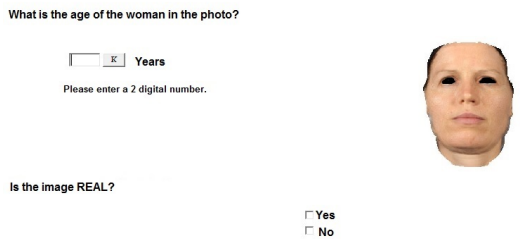

Fig. 6: Interface showed to judges.

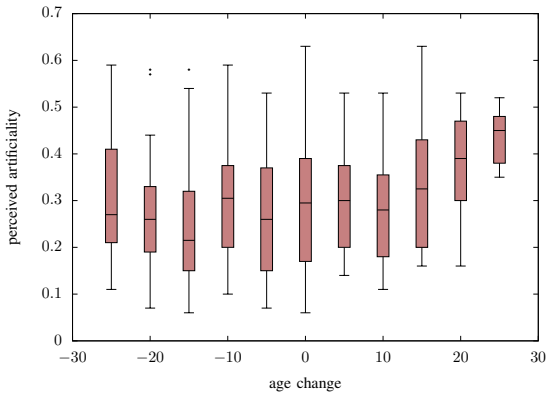

Fig. 7: Influence of age change on perceived artificiality. Values of zero are assigned to images that are perceived as real and values of one to the ones perceived as unreal.

z-score and compare it with the value of 1.96 corresponding to a significance level of $\alpha = 5\%$.

$$z = \frac{\hat{p_1} - \hat{p_2}}{\sqrt{\hat{p}(1-\hat{p})(\frac{1}{n_1} + \frac{1}{n_2})}}, \quad \hat{p} = \frac{n_1\hat{p_1} + n_2\hat{p_2}}{n_1 + n_2} \tag{9}$$

We obtain $z = 0.79$ which is smaller than $z_{\alpha/2} = 1.96$. This means, that we cannot reject the null-hypothesis. This shows that there is no obvious difference between how the simulations and the real images are perceived. Figure 7 shows the distribution of realism score as a function of the simulated age. The realism score is obtained by averaging the responses of the 80 judges to each image assigning a score of 1 to a fake image and a score of 0 to a real one.

C. Perception of Age in Simulated Ageing

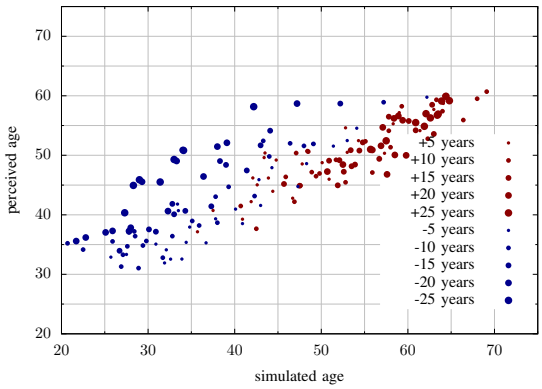

Fig. 8: Perceived age of artificially aged faces. Size of points indicates the amount of ageing applied to the receiver face. Blue points for rejuvenation, red points for ageing.

In this experiment, we validate, how well the simulated images agree with perception. We validate the precision of the ageing simulation by asking the judges to assign an age to the image. They answer the question by entering a two digit number. Figure 8 shows how the age of the simulated images is perceived. In an ideal case, the simulated and perceived ages should be the same, with a correlation coefficient of $R = 1$. We observe $R = 0.88$. Interesting to note is that the judges do not systematically underestimate the age of the simulations, as was reported in work of Burt and Adelson [3]. However, we observe the age to be overestimated in the case of rejuvenation and underestimated in the case of ageing. Therefore, recalibrating the ageing simulation according to the results of the study might be beneficial. To measure the model quality, we compute the difference between simulated age and perceived age. We compare it to the uncertainty of the judges, that is the standard deviation of perceived age between judges. A small model error and a larger uncertainty of the judges is indicative of a model producing images faithful to perceived age. The mean absolute difference between perceived age and targeted age is $5,6$ years. This is slightly below $6,7$ years, the standard deviation of perceived age between judges. Figure 9 shows the mean absolute difference

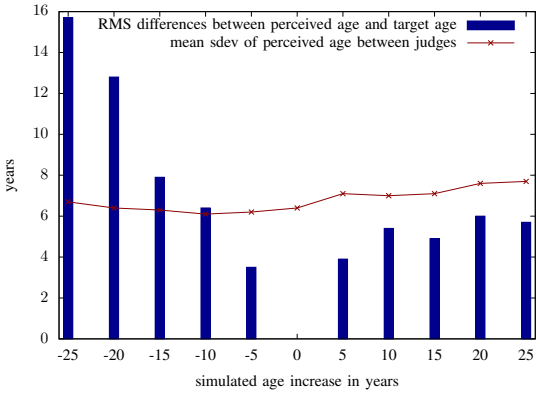

Fig. 9: Model accuracy of different amounts of ageing. We relate the difference between how old the faces are perceived and how they should be perceived according to the simulation with the standard deviation of perceived age between judges.

between the simulated age and the perceived age for different amounts of ageing. Also, it shows the mean variance of the perceived age between judges. There is a large mismatch for rejuvenation of more than 10 years. Astonishingly, in the case of ageing, the model error is relatively small compared to the standard deviation. We can say that the model works best on the range of -10 to 25 years of ageing.

VIII. CONCLUSION

We introduced an automatic face ageing approach to modify the age of any face on a picture by transferring face details from a set of donor faces. With this simple approach, we can create high-resolution photo-realistic results that agree with age perception. The modified faces are difficult to distinguish from real ones and they agree to human perception of ageing.

IX. FUTURE WORK

There are different avenues for future work to improve the ageing simulation. Currently, signs of ageing already present in the input face, such as wrinkles, are not handled differently than the rest of the skin. It would require a face representation for such features. Another area of improvement is related to surface light-interaction. For example, considering self-occlusion[29] because the amount of concavities increases with age, modelling wrinkles in the shape instead of the texture, and improving the illumination compensation of textures are all promising. We propose two improvements. One is to systematise the donor selection process. Selecting donors that show the same signs of ageing as the input, and that have a similar kind of illumination would address several of the ideas aforementioned. The other is to recalibrate the ageing model parameters according to the results of the perception experiments so that the correlation between simulated age and perceived age, as in Figure 8, increases.

REFERENCES

[1] F. Flament, R. Bazin, S. Laquieze, V. Rubert, E. Simonpietri, and B. Piot, "Effect of the sun on visible clinical signs of aging in Caucasian skin," *Clin Cosmet Investig Dermatol*, 2013.

[2] F. Flament and R. Bazin, *Skin Aging Atlas Volume 2: Asian Type*, 11 2010.

[3] D. M. Burt and D. I. Perrett, "Perception of age in adult caucasian male faces: Computer graphic manipulation of shape and colour information," *Proc. R. Soc. Lond. B*, vol. 259, pp. 137–143, 1995.

[4] V. Blanz and T. Vetter, "A morphable model for the synthesis of 3d faces," in *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*. ACM Press/Addison-Wesley Publishing Co., 1999, pp. 187–194.

[5] C. M. Scandrett, C. J. Solomon, and S. J. Gibson, "A person-specific, rigorous aging model of the human face," *Pattern Recognition Letters*, vol. 27, no. 15, pp. 1776–1787, 2006.

[6] K. Scherbaum, M. Sunkel, H.-P. Seidel, and V. Blanz, "Prediction of individual non-linear aging trajectories of faces," in *Computer Graphics Forum*, vol. 26, no. 3. Wiley Online Library, 2007, pp. 285–294.

[7] A. Lanitis, C. J. Taylor, and T. F. Cootes, "Toward automatic simulation of aging effects on face images," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 24, no. 4, pp. 442–455, 2002.

[8] ——, "Modeling the process of ageing in face images," in *Computer Vision, 1999. The Proceedings of the Seventh IEEE International Conference on*, vol. 1. IEEE, 1999, pp. 131–136.

[9] A. Shashua and T. Riklin-Raviv, "The quotient image: Class-based re-rendering and recognition with varying illuminations," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 23, no. 2, pp. 129–139, 2001.

[10] Z. Liu, Y. Shan, and Z. Zhang, "Expressive expression mapping with ratio images," in *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*. ACM, 2001, pp. 271–276.

[11] W.-S. Tong, C.-K. Tang, M. S. Brown, and Y.-Q. Xu, "Example-based cosmetic transfer," in *Computer Graphics and Applications, 2007. PG'07. 15th Pacific Conference on*. IEEE, 2007, pp. 211–218.

[12] F. Yun, Z. Nanning, L. Jianyi, and Z. Ting, "Facetransfer: A system model of facial image rendering," in *Systems, Man and Cybernetics, 2004 IEEE International Conference on*, vol. 3. IEEE, 2004, pp. 2180–2185.

[13] Y. Fu and N. Zheng, "M-face: An appearance-based photorealistic model for multiple facial attributes rendering," *IEEE Transactions on Circuits and Systems for Video technology*, vol. 16, no. 7, pp. 830–842, 2006.

[14] D. Guo and T. Sim, "Digital face makeup by example," in *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*. IEEE, 2009, pp. 73–79.

[15] Z. Liu, Z. Zhang, and Y. Shan, "Image-based surface detail transfer," *IEEE Computer Graphics and Applications*, vol. 24, no. 3, pp. 30–35, 2004.

[16] G. Guo, "Digital anti-aging in face images," in *Computer Vision (ICCV), 2011 IEEE International Conference on*. IEEE, 2011, pp. 2510–2515.

[17] A. Schneider, B. Egger, and T. Vetter, "A parametric freckle model for faces," in *Automatic Face & Gesture Recognition (FG 2018), 2018 13th IEEE International Conference on*. IEEE, 2018, pp. 431–435.

[18] Z. Xu, H. Chen, and S.-C. Zhu, "A high resolution grammatical model for face representation and sketching," in *null*. IEEE, 2005, pp. 470–477.

[19] J. Suo, F. Min, S. Zhu, S. Shan, and X. Chen, "A multi-resolution dynamic model for face aging simulation," in *Computer Vision and Pattern Recognition, 2007. CVPR'07. IEEE Conference on*. IEEE, 2007, pp. 1–8.

[20] J. Suo, S.-C. Zhu, S. Shan, and X. Chen, "A compositional and dynamic model for face aging," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 32, no. 3, pp. 385–401, 2010.

[21] H. Yang, D. Huang, Y. Wang, and A. K. Jain, "Learning face age progression: A pyramid architecture of gans," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.

[22] L. Jia, Y. Song, and Y. Zhang, "Face aging with improved invertible conditional gans," in *2018 24th International Conference on Pattern Recognition (ICPR)*, Aug 2018, pp. 1396–1401.

[23] Z. Wang, X. Tang, W. Luo, and S. Gao, "Face aging with identity-preserved conditional generative adversarial networks," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.

[24] G. Antipov, M. Baccouche, and J.-L. Dugelay, "Face aging with conditional generative adversarial networks," in *Image Processing (ICIP), 2017 IEEE International Conference on*. IEEE, 2017, pp. 2089–2093.

[25] S. Liu, Y. Sun, D. Zhu, R. Bao, W. Wang, X. Shu, and S. Yan, "Face aging with contextual generative adversarial nets," in *Proceedings of the 2017 ACM on Multimedia Conference*. ACM, 2017, pp. 82–90.

[26] P. Paysan, R. Knothe, B. Amberg, S. Romdhani, and T. Vetter, "A 3d face model for pose and illumination invariant face recognition," in *Advanced video and signal based surveillance, 2009. AVSS'09. Sixth IEEE International Conference on*. Ieee, 2009, pp. 296–301.

[27] S. Schönborn, B. Egger, A. Morel-Forster, and T. Vetter, "Markov chain monte carlo for automated face image analysis," *International Journal of Computer Vision*, vol. 123, no. 2, pp. 160–183, 2017.

[28] P. Pérez, M. Gangnet, and A. Blake, "Poisson image editing," *ACM Trans. Graph.*, vol. 22, no. 3, pp. 313–318, Jul. 2003. [Online]. Available: http://doi.acm.org/10.1145/882262.882269

[29] A. Schneider, S. Schönborn, L. Frobeen, T. Vetter, and B. Egger, "Efficient global illumination for morphable models," 2017.

[30] P. J. Burt and E. H. Adelson, "The laplacian pyramid as a compact image code," in *Readings in Computer Vision*. Elsevier, 1987, pp. 671–679.

[31] ——, "A multiresolution spline with application to image mosaics," *ACM Transactions on Graphics (TOG)*, vol. 2, no. 4, pp. 217–236, 1983.