Weight, Sex, and Facial Expressions: On the Manipulation of Attributes in Generative 3D Face Models

Brian Amberg, Pascal Paysan, and Thomas Vetter

University of Basel

Abstract. Generative 3D Face Models are expressive models with applications in modelling and editing. They are learned from example faces, and offer a compact representation of the continuous space of faces. While they have proven to be useful as strong priors in face reconstruction they remain to be difficult to use in artistic editing tasks. We describe a way to navigate face space by changing meaningful parameters learned from the training data. This makes it possible to fix attributes such as height, weight, age, expression or 'lack of sleep' while letting the infinity of unfixed other attributes vary in a statistically meaningful way.

We propose an inverse approach based on learning the distribution of faces in attribute space. Given a set of target attributes we then find the face which has the target attributes with high probability, and is as similar as possible to the input face.

1 Introduction

When producing movies or computer games it is common to use low dimensional generative 3D face models, which encode the space of possible faces with a few hundred parameters. These models encode the variability of the 3D shape, albedo, and reflectance properties. By varying the parameters of the model new 3D faces are created, which can then be rendered in the movie or game. The parameters of face models like 3D-MM [1] are a compact description of faces, but they are generally not meaningful. And even in manually constructed face models one can only change characteristics of the mesh like the size of the chin, but what the artist really wants is not to change the chin but to create a face which looks older, more male, or even more trustworthy. We propose a system which allows such manipulations directly in attribute space. To this end we learn an association between face and attribute space. We demonstrate our system with a 3D-Morphable Model (3D-MM) of Shape and Texture, but it is equally applicable to other models, which might be better at capturing statistics of wrinkles, hair, or reflectance properties of the skin. 3D-MMs are interesting because the starting face can be easily initialised from a real person using a reconstruction method as proposed for images in [1], videos in [2] or 3D scans in [3].

In this paper we use a relatively wide definition of *attribute*, ranging from physical attributes like crookedness of the nose or testosterone level over latent attributes like lack of sleep to cultural attributes like trustworthiness.

Sex:	*Female*	**Male**	**Male**	**Male**	**Male**
Height:	*164cm*	*188cm*	*185cm*	**176cm**	**176cm**
Weight:	*54kg*	*87kg*	**66kg**	**66kg**	**66kg**
Joy:	*0%*	*0%*	*0%*	*0%*	**100%**
	(a)	(b)	(c)	(d)	(e)

Fig. 1: *Generating new faces by changing some attributes of a face while keeping other attributes fixed. Starting from an initial face we can fix some attributes (typeset in bold) and let other attributes (typeset in italic) vary according to the statistical distribution of attributes in our training set. When an attribute is fixed we determine the face which simultaneously 1) is most similar to the starting face, 2) has the requested attributes with high probability and 3) is a likely face.*

Our method is not only useful as a tool for character generation, but also for the investigation and visualisation of the dependencies between attributes and faces. We can visualise the change of face features associated with attributes, or generate an infinity of psychological stimuly with precisely defined attributes, where it is possible to vary only a single attribute and keep others constant. We can for example generate images which vary in handsomeness while keeping perceived sex, height and weight constant.

2 Facial Attribute Manipulation

We state the problem as follows. For a given starting face and a new set of attributes, find the face which has the chosen attribute vector while being as close as possible to the input face in the natural distance in face space. With our method the artist generates in an intuitive way different distance functions in face space by specifying which attributes can vary and which should be fixed. This is explained in more detail in the next section.

We want to find the distribution $p(\bar{\boldsymbol{x}} \mid \boldsymbol{x}, \boldsymbol{a})$ of faces $\bar{\boldsymbol{x}}$ which are similar to a starting face $\boldsymbol{x}$ and have the attributes $\boldsymbol{a}$. Our training data consists of a collection of N example faces, described by their model coefficient $\boldsymbol{x}_i \in \mathbb{R}^k = \mathcal{F}$, and corresponding attribute vectors $\boldsymbol{a}_i \in \mathbb{R}^l = \mathcal{A}$. Attributes can be physical measurements like age, height, weight and sex which can be determined by measurement, subjective attributes like trustworthiness which are gauged with questionnaires, latent attributes like lack of sleep, and rapidly changing attributes like facial expressions, which were classified by asking the subjects to perform a specific expression. We model the likelihood as

$$p(\bar{\boldsymbol{x}} \mid \boldsymbol{x}, \boldsymbol{a}) \propto p(\boldsymbol{x}, \boldsymbol{a} \mid \bar{\boldsymbol{x}}) p(\bar{\boldsymbol{x}}) \tag{1}$$

$$= p(\boldsymbol{x} \mid \bar{\boldsymbol{x}}) p(\boldsymbol{a} \mid \bar{\boldsymbol{x}}) p(\bar{\boldsymbol{x}}), \tag{2}$$

where the probability $p(\boldsymbol{a} \mid \bar{\boldsymbol{x}})$ is learned from labelled training data. If enough data is available it is sensible to use a Gaussian Process for the regression, to get varying variances, or one can assume homoscedastic uncertainty and learn only a single σ_a and model

$$p(\boldsymbol{a} \mid \bar{\boldsymbol{x}}) = \mathcal{N}(\boldsymbol{a} \mid M(\bar{\boldsymbol{x}}), \sigma_a \boldsymbol{I}) \quad . \tag{3}$$

where $M(\bar{\boldsymbol{x}})$ is a regressed mapping from face space to attribute space. A normal distribution is used for the similarity measure, which assumes that face space is smooth in face appearance. This corresponds to the Mahalanobis distance in face space, which has shown good performance in face recognition experiments (e.g. [3]).

$$p(\boldsymbol{x} \mid \bar{\boldsymbol{x}}) = \mathcal{N}(\boldsymbol{x} \mid \bar{\boldsymbol{x}}, \sigma_d \boldsymbol{I}) \quad . \tag{4}$$

The choice of normal distributions is motivated by the ease of handling and the small number of parameters, but when enough data is available more sophisticated probability densities can be used. With the proposed distributions one arrives at the likelihood

$$\begin{aligned} p(\bar{\boldsymbol{x}} \mid \boldsymbol{x}, \boldsymbol{a}) \propto p(\boldsymbol{a} \mid \bar{\boldsymbol{x}}) p(\boldsymbol{x} \mid \bar{\boldsymbol{x}}) p(\bar{\boldsymbol{x}}) = \\ \frac{1}{2\pi\sqrt{\sigma_a \sigma_d}} \exp\left(-\frac{\|M(\bar{\boldsymbol{x}}) - \boldsymbol{a}\|^2}{\sigma_a^2} - \frac{\|\bar{\boldsymbol{x}} - \boldsymbol{x}\|^2}{\sigma_d^2} - \|\bar{\boldsymbol{x}}\|^2\right) \end{aligned} \tag{5}$$

We could now sample from this distribution, but our purpose is to generate a single predictable answer. Therefore we calculate the maximum likelihood solution.

$$T_{\mathrm{ML}}(\boldsymbol{x}, \boldsymbol{a}) = \arg\min_{\bar{\boldsymbol{x}}} \|M(\bar{\boldsymbol{x}}) - \boldsymbol{a}\|^2 + \gamma_1 \|\boldsymbol{x} - \bar{\boldsymbol{x}}\|^2 + \gamma_2 \|\bar{\boldsymbol{x}}\|^2 \tag{6}$$

where γ_1, γ_2 are trade-off parameters derived from σ_d, σ_a, which determine how closely the attributes should be matched.

Note that even though we used simple distributions the resulting model can be arbitrarily expressive – and the optimisation problem arbitrarily complicated – depending on the choice of the regression function. By taking the limit case of modelling $p(\boldsymbol{a} \mid \bar{\boldsymbol{x}})$ as a Dirac function we can also get a “hard” formulation of our problem, which when combined with linear regression leads to a very efficient optimisation problem. More details are given in section 2.2. We found that while the “hard” formulation works well for some attributes, it is for many attributes better to use the probabilistic formulation. The probabilistic version is harder to optimize but is still fast enough for interactive use. Our method is illustrated in figure 2.

2.1 Regression Methods

We evaluated a number of linear and nonlinear regressors and will describe them in this section. We found, that for the attributes attractiveness and, trustworthiness a *linear regression* worked best. For age, weight, height we got good results with Support Vector Regression (SVR) and the binary attribute sex and the expressions were best modelled by using the probability of one of the classes as the attribute. This is explained in section 2.1.

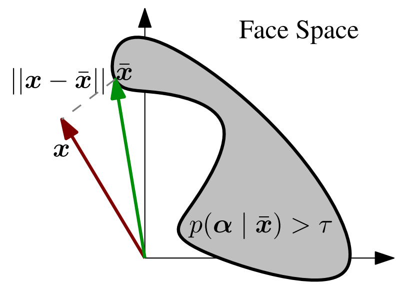

Fig. 2: *Given a current face $\boldsymbol{x}$ and a target attribute vector $\boldsymbol{a}$, we find the face $\bar{\boldsymbol{x}}$ which has the target attributes (with a high probability) as measured by the regression function $M(\bar{\boldsymbol{x}}) = \boldsymbol{a}$, while having the least change in identity as measured by the distance function $||\boldsymbol{x} - \bar{\boldsymbol{x}}||$. Additionally we incorporate the face-prior, which is not depicted here.*

Measured Attributes In addition to the attributes established with questionaires we included nonlinear measures on the shape model, namely distance between two points and angle between three points to allow changing the nose size and shape and to open and close the eyes and to fix the inter eye distance. These measures are (nonlinear) functions of the shape parameters and can be directly used instead of regressed attributes in our framework.

Linear Classifier The simplest classifier is a linear function of the input coefficients

$$M(\boldsymbol{x}) := \boldsymbol{M}\boldsymbol{x} + \boldsymbol{g} \quad . \tag{7}$$

Each row in $\boldsymbol{M}$ is essentially what has been used as an *attribute vector* in [1], it encodes how much a change in the attribute changes the coefficients of the face parameter.

Binary Attributes Some of the attributes we are using are of a discrete nature. Most prominently sex takes on only two distinct values, male and female.[1]. So while for height or attractiveness we can acquire continuous values on which a regression can be performed, this is impossible for discretely labeled attributes. Typically this was handled by determining the direction of main variance by linear regression where the classes c_1, c_2 are identified with the labels -1, 1. We found, that for this types of attributes a more intuitive and accurate method is to use the probability $p(c_1 \mid \boldsymbol{x})$ as the attribute. The probability is determined with Bayes theorem as

$$\begin{aligned} p(c_1 \mid \boldsymbol{x}) &= 1 - p(c_2 \mid \boldsymbol{x}) \\ &= p(\boldsymbol{x} \mid c_1)p(c_1)/p(\boldsymbol{x}) \\ &= p(\boldsymbol{x} \mid c_1)p(c_1)/(p(\boldsymbol{x} \mid c_1)p(c_1) + p(\boldsymbol{x} \mid c_2)p(c_2)) \quad . \end{aligned} \tag{8}$$

[1] Obviously there are other taxonomies, but we regard only this simple system.

The priors $p(c_i)$ are determined either from the samples, or in the case of sex we fixed them at $1/2$. The probabilities $p(\boldsymbol{x} \mid c_i)$ are modelled with parametric distributions, where we found that in the case of sex these probabilities were well modelled by normal distributions. A comparison between using linear regressor to model the change of sex and using equation 8 is shown in figure 3. Though the effect is subtle, it is visible that a change of sex does not correspond to a movement in the same direction in feature space for each face. Males and females are nonlinearly distributed, which is captured nicely by the mixture model. This concurs with the fact that when interpreted as a classifier, the probability $p(c_1 \mid \boldsymbol{x})$ does better in a cross-validation than the linear classifier. This method can be extended to n-ary attributes by splitting them into $n-1$ two valued attributes.

2.2 Solving Attribute Manipulation

We solve equation 6 with a Levenberg-Marquardt method. The optimisation was fast enough to allow us to build an interactive application.

An efficient algorithm for the linear regression A very efficient algorithm can be derived when using (1) the Mahalanobis Distance for face similarity and (2) a linear classifier with an associated dirac distribution for the attribute distribution and (3) a uniform face prior. In this case we are searching for the minimum of $\|\bar{\boldsymbol{x}} - \boldsymbol{x}\|^2$ under the constraint that $M(\bar{\boldsymbol{x}}) = \boldsymbol{a}$. For the linear classifier the solutions lie in a displaced subspace of face-space, such that the closest face fullfilling the attributes can be found by a projection. We can derive the solution from the Lagrange equation

$$L(\bar{\boldsymbol{x}}, \boldsymbol{\lambda}) = \|\bar{\boldsymbol{x}} - \boldsymbol{x}\|^2 + \boldsymbol{\lambda}^T(\boldsymbol{M}\bar{\boldsymbol{x}} + \boldsymbol{g} - \boldsymbol{a}) \quad , \tag{9}$$

which takes on its extremum at

$$\begin{aligned} 2\bar{\boldsymbol{x}} - 2\boldsymbol{x} - \boldsymbol{M}^T\boldsymbol{\lambda} &= \boldsymbol{0} \\ \boldsymbol{M}\bar{\boldsymbol{x}} + \boldsymbol{g} - \boldsymbol{a} &= \boldsymbol{0} \quad . \end{aligned} \tag{10}$$

This linear system can be written in matrix notation as

$$\begin{bmatrix} \boldsymbol{I} & -2^{-1}\boldsymbol{M}^T \\ \boldsymbol{M} & \boldsymbol{0} \end{bmatrix} \begin{bmatrix} \bar{\boldsymbol{x}} \\ \boldsymbol{\lambda} \end{bmatrix} = \begin{bmatrix} \boldsymbol{x} \\ \boldsymbol{a} - \boldsymbol{g} \end{bmatrix} \quad . \tag{11}$$

The solution is a linear function of input face and target parameters

$$\bar{\boldsymbol{x}} = \underbrace{\begin{bmatrix} \boldsymbol{I} \ \boldsymbol{0} \end{bmatrix} \begin{bmatrix} \boldsymbol{I} & -2^{-1}\boldsymbol{M}^T \\ \boldsymbol{M} & \boldsymbol{0} \end{bmatrix}^+}_{\text{constant}} \begin{bmatrix} \boldsymbol{x} \\ \boldsymbol{a} - \boldsymbol{g} \end{bmatrix} \quad , \tag{12}$$

where we use $\cdot^+$ to denote the pseudo inverse. Note that this is different from the method in [1], as it allows the simultaneous handling of multiple parameters, where attributes can be fixed or left free to vary at the discretion of the user. Also, we can prescribe absolute values, such as 30 years and 83 kilos, instead of only offsets from the current shape.

3 Evaluation

As this method is targeted at humans, it is inevitable to assess the performance of the system based on human judgement – *it has to look good*. But additionally, as the basis of this method are regression methods, it is possible to evaluate the regressions on a test set, which we used to determine the kind of classifier to use for each attribute. If the classifier is unable to correctly classify the test set, then the method must also fail human judgement. We will now present some results, more are shown in the accompanying online material.

3.1 Model

We constructed the underlying 3D model in the same way as [3]. We acquired 800 facial scans from over 300 IDs, which were brought into correspondence with a nonrigid ICP method similar to [4] but with a different regularisation term. To increase the database we added mirrored versions of all scans. Every ID had at least one neutral scan. From the neutral scans we build separate PCA models of shape and texture, which we call the identity models. Two more PCA models where built from the offsets between the expression scans and the corresponding neutral scans. We call this the expression model. The two models are concatenated, such that we have two sets of coefficients, the identity coefficients and the expression coefficients. For the experiments in this paper we used 80 neutral shape, 40 expression shape and 20 neutral texture components, resulting in a 140 dimensional face space.

3.2 Speed

The runtime of the optimisation depends on the evaluation of the regression functions and their derivatives. For the experiments used in this paper we have real-time results when using only linear constraints and the method proposed in section 2.2. When using the nonlinear classifiers presented here, we have typically a runtime of less than half a second for the iterative methods, allowing interactive exploration of attribute space. To speed up the iterative methods we first fit 25 steps from the reference face $\boldsymbol{a}$ and from the current estimate from the previous fitting, and continue then a full fitting from the position with the lower residual, which speeds up the search time in interactive use.

3.3 Binary Attributes

We show in figure 3 that changing the sex when using the nonlinear probability $p(\text{male} \mid \boldsymbol{x})$ as the attribute results in a face which has convincingly changed sex but is closer to the starting face than that resulting from linear regression. While fixing the target probability $p(\text{male} \mid \boldsymbol{x})$ to 0 or 1 results in convincing faces, it is more difficult to determine a suitable value in the linear scale, where -1 or +1 often correspond to too

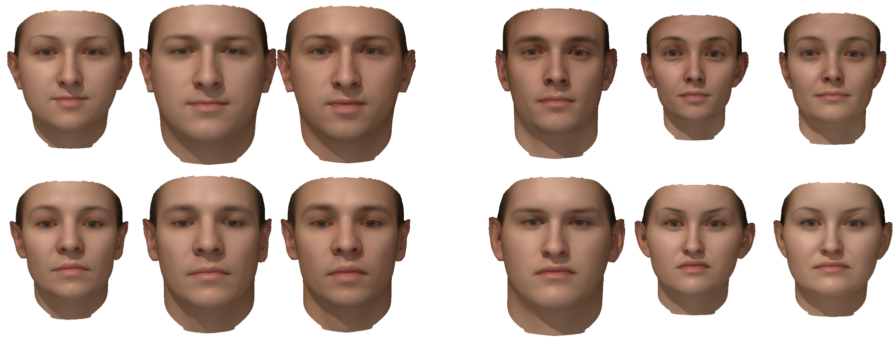

(a) Original (b) Linear (c) Nonlinear (a) Original (b) Linear (c) Nonlinear

Fig. 3: *Modelling the distribution of sex in face space as a Gaussian mixture is superior to determining just the direction of main variance. Using the probability of 'male' as an attribute as in equation 8 results in faces (c) which retain more of the characteristics of the source face (a) while undergoing a perceived change of sex than the result achievable with linear regression (a).*

pronounced or not enough pronounced changes in sex. For this comparison we fixed the value of the linearly regressed sex to that which we determined for the face generated from the nonlinear regression. This should result in a fair comparisons of the methods.

The second set of attributes that we applied this method to were expressions. Our dataset includes labeled examples of five different expressions, and we trained one classifier per expression, in a one against all scheme. In figure 5 we show a sequence, where the starting face has a sad expression, which we then remove by setting all expressions to neutral. By changing then the joy value to one and keeping the other expression attributes fixed we changed the initial sad face into a happy face. As we have relatively little data in this high dimensional face-space, we had to regularize the estimation of the Gaussian for each class. This was done by setting the small eigenvalues of the covariance to a constant.

3.4 Measured Attributes

Distance and angular measures as introduced in section 2.1 are another powerful editing method, which is closer to traditional mesh editing. The difference is, that as we determine the maximum likelihood solution of the model, a local change always has global influence. So changing the length of the nose will also change the size of the nose, and if that should not happen, then the size of the nose has to be constrained by another measure and set to fixed. And the strength of our approach is that both applications are possible, depending on the use case at hand. We demonstrate in figure 4 a number of changes to the size and shape of the nose of the input face, and show that this can be combined with other classifiers by finally changing the expression to sad.

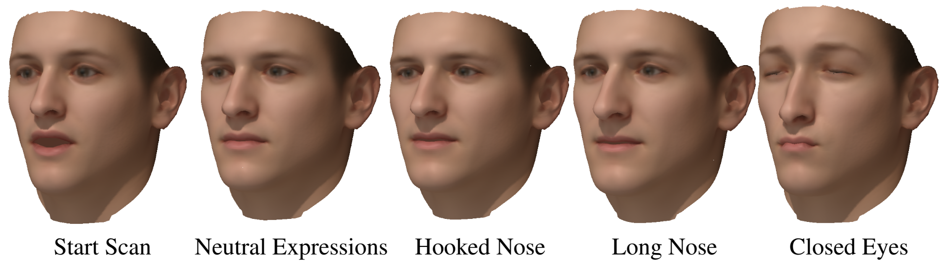

Fig. 4: *Face Space Measures allow the manipulation of length and angles in the face. They integrate naturally into our framework, such that we determine the most likely face for the given length and angle constraints. To demonstrate the integration of expressions and measures we first remove the expressions from the starting face using the expression attribute, then make its nose more hooked, enlarge the nose and finally close the eyes.*

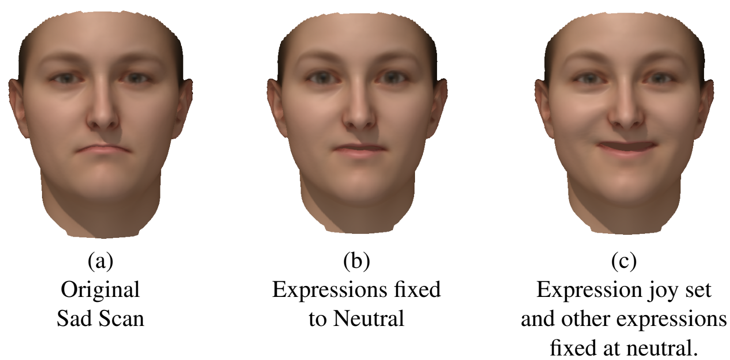

Fig. 5: *Our method allows the removal (b) of an expression from the input scan (a) and subsequent addition of other expressions (c).*

4 Prior Work

Manipulation of facial attributes with the help of statistics learned from example faces has been introduced already by [1] in the seminal paper on 3D-MMs. They learned the direction of maximum variance of an attribute from a labelled set of faces and changed the face parameters of a given identity according to that direction. In addition to not being able to manipulate "pure" attributes, i.e. changing only the weight but not the height, this approach can neither prescribe attributes in an absolute scale, i.e. to age a person to 43 years, nor does the change of attributes depend on the starting identity – every person ages in the same way. Later work [5] of the same group which is most closely related to our approach extended this work to absolute scales and multiple attributes by constraining the generated face to be in the solution space of a linear regression from faces to attributes. Their proposal is a special case of our method.

Other authors propose to learn a function from attributes to shapes [6,7,8,9,10,11], which addresses a different problem. Only the mapping from shapes to attributes is surjective, i.e. to a single face there is always only a single value of each attribute, but not the mapping from attributes to shapes, many faces share the same attributes. Therefore, these approaches can only generate an “average” face for each attribute. The mentioned papers explore combinations of anthropometric or subjective attributes, and linear or nonlinear mappings and linear, multilinear or nonlinear models, but all of them do not constrain the solution by similarity to a starting face, as we do. To generate faces conforming to anthropometric measurements [7] fits a linear function from measurements to faces, which is correct for distance measurements, but not for angle measurements. Similarly, [10] interpolate the example faces based on their distance in measurement space with RBF interpolation, which also does not guarantee correct distance. Our method works with correct measurements, and, as stated, overcomes the injectivity problem by further constraining the mapping by distance to a starting face.

Note that while multilinear models as in [12,11] separate meaningful groups of latent factors, they do not offer meaningful parameters inside of a group (tensor mode). Because multilinear models do not have meaningful parameters [11] suggested to drive face manipulation by motion capture data of actors, and combine identity and expression tensor parameters from different measurements for expression transfer. This is also an interesting application, but it is very different from our goal of attribute and expression editing in a meaningful parameters space. The same applies to [13] where a generative body model is fitted to motion capture data.

We expect, that the change in face parameters for a change of attributes would depend on the starting face. This was addressed in [14] They learned a nonlinear age regression, and directions of change which are parallel to the gradient of the regression function were used to determine a trajectory through face space. This addresses the problem that different individuals should change differently and makes it possible to set an absolute value for an attribute. The disadvantage of that method is, that it only handles a single attribute, and does not model the covariance of multiple attributes. Changing first a persons weight also changes the height, so it is impossible to change height and weight simultaneously to a fixed value. [15] recently proposed a method to manipulate attractiveness in 2D frontal photographs. The method is similar to [14]. They also perform a regression from faces described by the distance between feature points to an attractiveness rating. The input face is then morphed such that its attractiveness rating increases. Our approach, which is also based on a regression, allows simultaneous manipulation of many attributes, enables the artist to specify a precise set of target attributes, and makes it possible to explicitly choose the attributes which should covary and which should stay fixed.

While [14] assume that the direction of the gradient of the nonlinear age regression will leave the attributes which make up the identity the same, as this corresponds to the smallest Mahalanobis distance in face space which achieves the desired change of attribute. Correspondingly [15] assume that closer faces in a euclidean distance in their parametrization of face space corresponds to the smallest change in identity. In reality, when the age of a person changes it is for example the case that the inter-eye

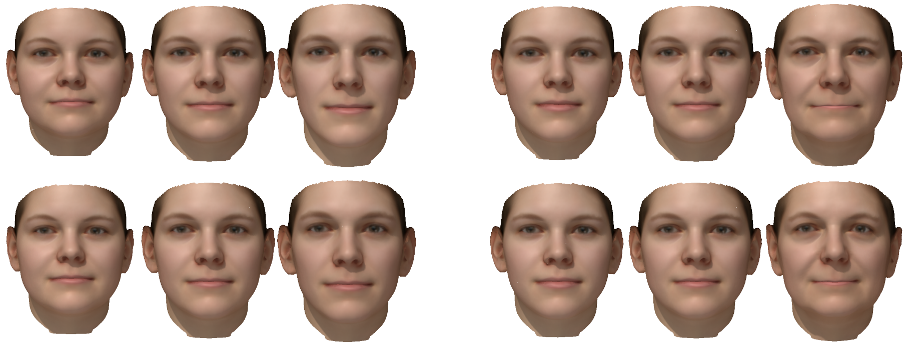

Fig. 6: *For some attributes we can achieve better control with support vector regression than with linear regression. The effect is relatively subtle in this presentation, but notice e.g. the mouth shape which stays more constant and also the more similar global shape.*

distance stays fixed and the sex of the person does not change, measures like this are not incorporated in the distance function used in the mentioned papers. We propose to learn many attributes and decide *depending on the task at hand* which attributes should stay the same, and which are allowed to vary. So with our method it is possible to specify the understanding of what makes two faces similar by deciding on attributes which should not vary, while previous methods only assumed that a short distance in face space leaves the identity unchanged.

5 Conclusion

We presented a method to manipulate attributes of faces in generative face models, where the attribute space is learned from labelled examples. We demonstrated the method using a 3D Morphable Model. There exist an infinity of possible attributes ranging from physical values like the curvature of the ears to cultural valuations like attractiveness. Attributes are distributed nonlinearly in face space, which we addressed by learning a nonlinear regression from face space to attribute space. And attributes co-vary, when navigating attribute space it is desirable to be able to choose which should vary, and which should be fixed. For example a change of attractiveness should not change the sex of the modelled face, but might be allowed to change the curvature of the nose. We addressed this in a probabilistic framework.

Our method has three applications. It enables psychologists to generate stimuli with systematically varying attributes, it gives artists a powerful and intuitive new tool for character design based on learned statistics of real faces, and additionally, it is fun to play with.

References

1. Blanz, V., Vetter, T.: A morphable model for the synthesis of 3d faces. In: SIGGRAPH '99, New York, NY, USA, ACM Press/Addison-Wesley Publishing Co. (1999) 187–194 1, 4, 5, 8
2. Ilić, S., Salzmann, M., Fua, P.: Implicit meshes for effective silhouette handling. Int. J. Comput. Vision **72** (2007) 159–178 1
3. Amberg, B., Knothe, R., Vetter, T.: Expression invariant face recognition with a morphable model. Face and Gesture 2008 (2008) 1, 3, 6
4. Amberg, B., Romdhani, S., Vetter, T.: Optimal step nonrigid ICP algorithms for surface registration. In: Computer Vision and Pattern Recognition, 2007. CVPR '07. IEEE Conference on. (2007) 1–8 6
5. Blanz, V., Albrecht, I., Haber, J., Seidel, H.P.: Creating face models from vague mental images. In Gröller, E., Szirmay-Kalos, L., eds.: EUROGRAPHICS 2006 (EG'06). Volume 25:3 of Computer Graphics Forum., Vienna, Austria, Blackwell (2006) 645–654 8
6. Allen, B., Curless, B., Popović, Z.: The space of human body shapes: reconstruction and parameterization from range scans. ACM Trans. on Graphics, SIGGRAPH 2003 **22** (2003) 587–594 9
7. Allen, B., Curless, B., Popović, Z.: Exploring the space of human body shapes : Data-driven synthesis under anthropometric control. In: Proc. Digital Human Modeling for Design and Engineering Conference. Volume 113:1., Rochester, MI, SAE International (2004) 245–248 9
8. Skorupski, J., Yee, J., McCoy, J., Davis, J.: Facial type, expression, and viseme generation. In: SIGGRAPH '07: ACM SIGGRAPH 2007 posters. (2007) 34 9
9. Allen, B., Curless, B., Popović, Z., Hertzmann, A.: Learning a correlated model of identity and pose-dependent body shape variation for real-time synthesis. In: SCA '06: Proceedings of the 2006 ACM SIGGRAPH/Eurographics symposium on Computer animation, Aire-la-Ville, Switzerland, Switzerland, Eurographics Association (2006) 147–156 9
10. Zhang, Y., Badler, N.I.: Face modeling and editing with statistical local feature control models. International Journal of Imaging Systems and Technology **17** (2007) 341–358 9
11. Vlasic, D., Brand, M., Pfister, H., Popovic, J.: Face transfer with multilinear models. ACM Trans. Graph. **24** (2005) 426–433 9
12. Vasilescu, M.A.O., Terzopoulos, D.: Multilinear analysis of image ensembles: Tensorfaces. In: In Proceedings of the European Conference on Computer Vision. Volume 1. (2002) 447–460 9
13. Anguelov, D., Srinivasan, P., Koller, D., Thrun, S., Rodgers, J., Davis, J.: Scape: shape completion and animation of people. ACM Trans. Graph. **24** (2005) 408–416 9
14. Scherbaum, K., Sunkel, M., Seidel, H.P., Blanz, V.: Prediction of individual non-linear aging trajectories of faces. In: Eurographics 2007. Volume 26:3 of Computer Graphics Forum., Prague, Czech Republic, Blackwell (2007) 285–294 9
15. Leyvand, T., Cohen-Or, D., Dror, G., Lischinski, D.: Data-driven enhancement of facial attractiveness. ACM Trans. on Graphics, SIGGRAPH 2008 **27** (2008) to appear 9