

Probabilistic Modeling and Visualization of the Flexibility in Morphable Models

M. Lüthi, T. Albrecht, and T. Vetter

Computer Science Department, University of Basel, Switzerland
{marcel.luethi,thomas.albrecht,thomas.vetter}@unibas.ch

Abstract. Statistical shape models, and in particular morphable models, have gained widespread use in computer vision, computer graphics and medical imaging. Researchers have started to build models of almost any anatomical structure in the human body. While these models provide a useful prior for many image analysis task, relatively little information about the shape represented by the morphable model is exploited. We propose a method for computing and visualizing the remaining flexibility, when a part of the shape is fixed. Our method, which is based on Probabilistic PCA, not only leads to an approach for reconstructing the full shape from partial information, but also allows us to investigate and visualize the uncertainty of a reconstruction. To show the feasibility of our approach we performed experiments on a statistical model of the human face and the femur bone. The visualization of the remaining flexibility allows for greater insight into the statistical properties of the shape.

1 Introduction

Morphable models, i.e. statistical shape models based on dense point-to-point correspondence, have become a widely used tool in computer vision, computer graphics and medical imaging. The main idea behind a morphable model is to span a space of shapes (3D surfaces) by taking linear combinations of example shapes [1]. A probability distribution is estimated from the example shapes, quantifying the probability of observing each linear combination. The most common use of morphable models is to restrict the solution-space of ill posed problems by penalizing unlikely instances of the shape. Typical examples include image segmentation [2,3,4], registration [5,6] or 2D-3D surface reconstruction [7,8,9]. In this context, the model is used to answer the following question:

- Given a shape, how likely is it that the shape belongs to the object class represented by the morphable model?

These applications exploit only the fact that the variability of the shape as a whole can be represented and quantified by the morphable model.

In this paper we are trying to get a deeper understanding of the information a morphable model represents and how one part of the model influences the rest. The central question we are trying to answer is:

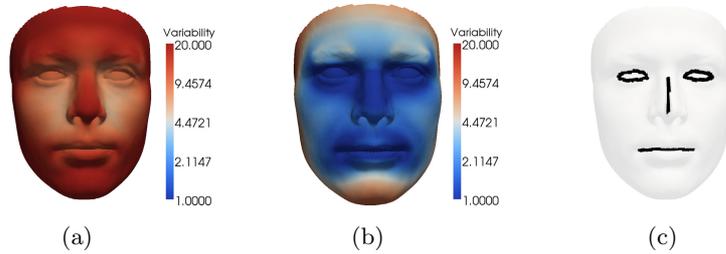


Fig. 1. Flexibility of a morphable model of the human face. The colors represent the variability (in mm) for each point. Figure (a) shows the full flexibility of the morphable model. In (b), the most likely reconstruction of the sketch depicted in (c) is shown, together with the remaining variability.

- Given only a part of a shape, what is the most likely completion of this shape and how much variance remains in the model given this partial information?

We illustrate this with an example. Assume we are given a morphable model of the human face. Let s be a random variable representing the surface, with its distribution given by the morphable face model. Figure 1(a) shows the mean face $E[s]$ and the variability represented by the model (which is, loosely speaking, the variance $\text{var}(s)$). Now suppose we are given a rough sketch of the eyes, nose and mouth in form of the black lines in Figure 1(c) and wish to reconstruct a full face from these lines. Denoting the given lines by s_b , we are interested in the distribution of the random variable $s|s_b$. Figure 1(b) shows the most likely reconstruction $s^* := \arg \max_s p(s|s_b)$ of the full shape, as well as the remaining variability $\text{var}(s|s_b)$. Naturally, the shape variability is much lower than in Figure 1(a) because the sketch s_b is observed. Hence, knowledge of the distribution of $s|s_b$ not only leads to an approach for the reconstruction of a face from the sketch, but, equally interesting, indicates how well the face is determined by the sketch. In the remainder of the paper we show how these quantities can be computed, under the assumption that the shapes follow a normal distribution.

One main assumption of morphable models is that the shapes, which are usually represented as very high dimensional vectors, lie on an embedded linear manifold (i.e. plane) within this high dimensional shape space. This manifold is found by performing Principal Component Analysis (PCA) of the sample covariance matrix. Unfortunately, standard PCA does not provide a probability model in the shape space [10]. In particular, in our “high-dimension - small sample” case, the covariance matrix becomes singular, which leads to various problems in statistical reasoning. In the approach presented in this paper, we use Probabilistic PCA (PPCA) [11], which defines a proper covariance structure in the shape space. The PPCA approach also directly implies a method for reconstructing a full shape given only partial information by using the posterior mean as the best reconstruction.

Knowledge of the best reconstruction and remaining variability of a shape is important in many application domains. We are particularly interested in the area of reconstructive medicine, where the problem arises frequently that a traumatised structure has to be reconstructed to fit the remaining parts. Being able to assess the remaining flexibility in a statistically sound way is important for the planning of reconstructive surgical interventions and the prediction of the outcome. Further, it is of general interest to know how much the different parts of an anatomical structure determine its variability. This knowledge can for example give important clues for designing implants and prostheses.

Several authors have proposed very similar methods for reconstruction of a full surface from partial information [12,13,14,15,16]. In fact, the reconstruction method resulting from the PPCA approach encompasses the one proposed by Blanz and Vetter [12], and Basso and Vetter [13] as a special case. A similar model based on factor analysis, which strongly resembles our PPCA model, was proposed by Machade et al. in [17]. However, to the best of our knowledge, the PPCA framework for model based reconstruction and the use of the full posterior distribution for modeling and visualizing the remaining flexibility has not been considered before. Indeed, the only work we are aware of that explicitly tries to model the remaining flexibility is the one by Albrecht et al. [18]. In contrast to our work it does not admit a probabilistic interpretation and requires a separate algorithm for shape reconstruction.

2 Background

Before we present our model, we briefly review the mathematical concepts we will use in the remainder of the paper. In order to apply statistical methods to shapes we need to be able to represent them as random variables. A particularly simple approach, is to sample the shape and organize the sampled points as a vector in Euclidean space. Such a vector is usually referred to as a shape vector.

Two types of statistical shape models are distinguished in the literature. In the *Active Shape Models* [19], the shape is given as a 2D contour and is relatively sparsely sampled. In contrast, in the *Morphable Model* [1], the shape is given as a 3D surface and the sampling is dense. From a mathematical point of view, the concepts are the same. An important difference is, however, that in the case of the Morphable Model the dimensionality of the shape vectors is much larger than the number of observations. It is this property that motivates our work, and we will therefore focus only on this case in the remainder of this paper.

2.1 Shape Vectors and Registration

Let $\{T_i \subset \mathbb{R}^3 | i = 1, \dots, n\}$ be n surfaces, given in some suitable representation. Define an arbitrary surface, say T_1 , as a reference surface. We assume that each surface T_i was obtained by warping the reference surface T_1 with a smooth vector field $\phi_i : T_1 \rightarrow \mathbb{R}^3$. That is

$$T_i = \{x + \phi_i(x) | x \in T_1\}.$$

Let $\hat{\Gamma}_1$ be a suitable discretization of Γ_1 of N points (e.g. $\hat{\Gamma}_i$ is represented as a triangle mesh with N vertices). Note, that the same discretization is induced by the mapping ϕ for each surface Γ_i . We define the shape vector $s_i \in \mathbb{R}^{3N}$ as

$$s_i = (v_x^{i,1}, v_y^{i,1}, v_z^{i,1}, \dots, v_x^{i,N}, v_y^{i,N}, v_z^{i,N})^T,$$

where the vector $v^{i,j} = (v_x^{i,j}, v_y^{i,j}, v_z^{i,j})$ represents the x, y, z coordinates of the j -th vertex of $\hat{\Gamma}_i$.

Usually we are given the surfaces $\Gamma_1, \dots, \Gamma_n$ rather than a reference surface Γ_1 and the corresponding vector fields $\{\phi_i\}_{i=2}^n$. Finding a vector field ϕ that maps between a given pair of surfaces is a central problem in medical imaging and computer vision and is referred to as the *registration* or *correspondence* problem. Many algorithms for surface registration have been described in the literature (see e.g. [20] for an overview).

2.2 Principal Component Analysis and Statistical Shape Models

PCA is a well known and widely used method for dimensionality reduction and data visualization. From n data sets, represented by vectors $s_i \in \mathbb{R}^m$ the mean $\mu = \frac{1}{n} \sum_{i=1}^n s_i$ and covariance matrix $\Sigma \in \mathbb{R}^{m \times m}$ with $\Sigma = \frac{1}{n} \sum_{i=1}^n (s_i - \mu)(s_i - \mu)^T$ can be estimated. PCA consists of an eigenvalue decomposition or *principal component* transformation of Σ :

$$\Sigma = UD^2U^T, \quad (1)$$

where $U \in \mathbb{R}^{m \times m}$ is the orthonormal matrix of the eigenvectors of Σ , ordered according to the size of the corresponding eigenvalues, which make up the diagonal of the matrix $D^2 = \text{diag}(\sigma_1^2, \dots, \sigma_m^2)$. Note that if $n < m$, we have $\sigma_i = 0$ for all $i > n$.

When building a PCA-based shape model, it is assumed that the training datasets s_i and linear combinations thereof form a linear space of shapes that can be modelled by a multivariate normal distribution $\mathcal{N}(\mu, \Sigma)$. With the help of a coefficient vector α , each shape can be represented as:

$$s = s(\alpha) = UD\alpha + \mu. \quad (2)$$

Thanks to this representation the probability density function $\mathcal{N}(\mu, \Sigma)$ evaluated at $s(\alpha)$ takes the form:

$$p(s(\alpha)) = \frac{1}{z} \exp(-\|\alpha\|_2^2), \quad (3)$$

where $z = \sqrt{(2\pi)^m \det(D)}$ is a normalization factor [12].

2.3 Linear Gaussian Models

The PPCA model considered in this paper is a linear Gaussian model of the form

$$y = Ax + b + \varepsilon, \quad (4)$$

where $A \in \mathbb{R}^{m \times n}$ and $b \in \mathbb{R}^m$ are parameters, $x \sim \mathcal{N}(\mu, \Lambda)$ and $\varepsilon \sim \mathcal{N}(0, L)$ are normally distributed random variables. For this model, the predictive distribution $p(y)$ and posterior distribution $p(x|y)$ have a simple analytic form, as summarized in the following theorem [21]:

Theorem 1 (Theorem for Linear Gaussian Models). *Given a marginal Gaussian distribution for x and a conditional distribution for $y|x$ in the form*

$$p(x) = \mathcal{N}(\mu, \Lambda) \quad (5)$$

$$p(y|x) = \mathcal{N}(Ax + b, L) \quad (6)$$

the marginal distribution of y and the conditional distribution of x given y are given by

$$p(y) = \mathcal{N}(A\mu + b, L + A\Lambda A^T) \quad (7)$$

$$p(x|y) = \mathcal{N}(\Sigma A^T L^{-1}(y - b) + \Lambda^{-1}\mu, \Sigma) \quad (8)$$

where

$$\Sigma = (\Lambda + A^T L A)^{-1}. \quad (9)$$

It is the fact that we have an analytic form of the posterior distribution that allows us in the following to model and visualize the remaining variability.

3 Shape Modeling Using Probabilistic PCA

In this section we present our PPCA based method for modeling the shape variability and show how it leads to a natural approach for shape reconstruction. Further, we show how the resulting posterior distribution can be used to visualize effectively the remaining flexibility in the model.

3.1 Mathematical Model

The mathematical model we use for shape modeling is obtained by applying standard Probabilistic Principal Component Analysis, as proposed by Tipping and Bishop [11], to a set of surfaces represented as shape vectors.

Let $\{s_i \in \mathbb{R}^{3N}\}_{i=1}^n$ be n shape vectors as defined in section 2.1. The main assumption in PPCA is that the high dimensional observations can be explained by a mapping from a low dimensional latent space plus some additional Gaussian noise. Let α be a d -dimensional latent variable

$$p(\alpha) = \mathcal{N}(0, \mathcal{I}_d). \quad (10)$$

We model the conditional distribution of observing the shape vector s as

$$p(s|\alpha) = \mathcal{N}(W\alpha + \mu, \sigma^2 \mathcal{I}_{3N}) \quad (11)$$

where $W \in \mathbb{R}^{3N \times d}$ is a linear mapping and $\mu \in \mathbb{R}^{3N}$ a shape vector. We can interpret this as a generative model, where the shape s is given by the mapping

W of the latent variable α plus some additive Gaussian noise $\varepsilon \sim \mathcal{N}(0, \sigma^2 \mathcal{I}_{3N})$. That is

$$s = W\alpha + \mu + \varepsilon. \quad (12)$$

According to Theorem 1 the predictive distribution

$$p(s) = \int p(s|\alpha)p(\alpha) d\alpha. \quad (13)$$

is again Gaussian with mean μ and covariance matrix $WW^T + \sigma^2 \mathcal{I}_{3N}$. In summary, we obtain the distribution

$$p(s) = \mathcal{N}(\mu, WW^T + \sigma^2 \mathcal{I}_{3N}). \quad (14)$$

Theorem 1 also provides us with an expression for the posterior distribution:

$$p(\alpha|s) = \mathcal{N}(M^{-1}W^T\sigma^{-2}(s - \mu), M^{-1}), \quad (15)$$

where

$$M = \sigma^{-2}W^TW + \mathcal{I}_d.$$

Tipping and Bishop [11] showed that the maximum likelihood solution for the parameters μ, W, σ^2 is

$$\mu_{\text{ML}} = \frac{1}{n} \sum_{i=1}^n s_i \quad (16)$$

$$W_{\text{ML}} = U_d(D_d^2 - \sigma^2 \mathcal{I}_d)^{\frac{1}{2}} \quad (17)$$

$$\sigma_{\text{ML}} = \frac{1}{3N - d} \sum_{i=d+1}^{3N} D_{ii}^2. \quad (18)$$

Here, $U_d \in \mathbb{R}^{3N \times d}$ and $D_d \in \mathbb{R}^{d \times d}$ are the sub-matrices corresponding to the d largest eigenvalues of the covariance matrix decomposition in (1). Using these maximum likelihood estimates in the generative model (12) yields a striking similarity with the PCA model (2), which gives the method its name. However, in contrast to the standard PCA, PPCA provides a fully probabilistic model. This allows for the computation of the full posterior distribution and to deal with missing data in a principled way.

3.2 Missing Data

Assume now that a part of the model is given. Without loss of generality, the model can be partitioned as $s = (s_a, s_b)$ with s_b given and s_a unknown. We would like to reconstruct the full shape $s \in \mathbb{R}^{3N}$ from the partial shape $s_b \in \mathbb{R}^{3\tilde{N}}$. Equation 14 can be written as

$$p(s) = p(s_a, s_b) = \mathcal{N}\left(\begin{bmatrix} \mu_a \\ \mu_b \end{bmatrix}, \begin{bmatrix} W_a W_a^T & W_a W_b^T \\ W_b W_a^T & W_b W_b^T \end{bmatrix} + \sigma^2 \mathcal{I}_{3\tilde{N}}\right). \quad (19)$$

Using the well known formula for the conditional distribution of a multivariate normal distribution, we have

$$p(s_a|s_b) = \mathcal{N}(\mu_{s_a|s_b}, \Sigma_{s_a|s_b}) \quad (20)$$

with

$$\mu_{s_a|s_b} = \mu_a + W_a W_b^T (W_b W_b^T + \sigma \mathcal{I}_{3\tilde{N}})^{-1} (s_b - \mu_b) \quad (21)$$

and

$$\Sigma_{s_a|s_b} = (W_b W_b^T + \sigma \mathcal{I}_{3\tilde{N}}) - W_a W_b^T (W_b W_b^T + \sigma \mathcal{I}_{3\tilde{N}})^{-1} W_b W_a^T. \quad (22)$$

The above equations give us a simple recipe for shape reconstruction. Unfortunately, it requires the inversion of the matrix $(W_b W_b^T + \sigma \mathcal{I})$, which is potentially huge. Using the fact that the shape is determined by the latent variables, we instead try to find an expression for $p(\alpha|s_b)$:

$$p(s_b|\alpha) = \mathcal{N}(W_b \alpha + \mu_b, \sigma^2 \mathcal{I}_{3\tilde{N}}) \quad (23)$$

$$p(\alpha) = \mathcal{N}(0, \mathcal{I}_d) \quad (24)$$

Again, we are in the position to apply Theorem 1:

$$p(\alpha|s_b) = \mathcal{N}(M^{-1} W_b^T \sigma^{-2} (s_b - \mu_b), M^{-1}), \quad (25)$$

with

$$M = \sigma^{-2} W_b^T W_b + \mathcal{I}_d. \quad (26)$$

In all practical cases, $W_b^T W_b$ will be small and can easily be computed. Once α has been determined, the most likely shape s^* is given by

$$s^* = \arg \max_s p(s|\alpha) = W \alpha + \mu. \quad (27)$$

This reconstruction is sufficient for the majority of shape modeling applications. Hence, we hardly ever need to compute the full covariance matrix $\Sigma_{s_a|s_b}$. It is, however, interesting to write down the distribution $p(s_a|s_b)$ in terms of the latent variables:

$$p(s|s_b) = p(s_a|s_b) = \int p(s_a|\alpha, s_b) p(\alpha|s_b) d\alpha = \int p(s_a|\alpha) p(\alpha|s_b) d\alpha \quad (28)$$

where we used the fact that s_a and s_b are conditionally independent given α . This can be interpreted as a projection of the observation onto the latent space, followed by the reconstruction of the full shape for the given α .

3.3 Reconstruction of Partial Shapes

We show now how the results from Section 3.2 can be used to reconstruct missing parts of any shape that can be modeled by a morphable model. In order to model the partial shape s_b as a part of a given complete morphable model, it has to be in correspondence with the reference shape (cf. Section 2.1).

The latent variable $\alpha|s_b$ is distributed according to Equation (25). The most probable reconstruction is obtained by reconstructing the full shape from the mean according to Equation (27). In addition to providing us with the most probable reconstruction, $p(\alpha|s_b)$ describes the distribution for all possible reconstructions. By considering how strongly this distribution is concentrated around its mean, we see exactly how reliable the reconstruction with the mean is. In effect, $p(\alpha|s_b)$ models the remaining flexibility of the morphable model given the fixed part s_b . In the next section, we will show how this flexibility can be explored visually.

3.4 Visualizing the Remaining Variability

We reconstruct the shapes from the latent variable α using Equation (27), i.e. $s^* := W\alpha + \mu$. According to the standard formula the covariance matrix of s^* under this affine transformation becomes

$$WM^{-1}W^T. \quad (29)$$

Note that this is a simpler distribution than the $p(s|s_b)$ given in Equation (20), as we can ignore the Gaussian noise ε of the original model (12) here.

In order to visualize the main modes of variation of this distribution, we perform an additional PCA, i.e. an eigenvalue decomposition of the covariance matrix $WM^{-1}W^T$. When we choose the maximum likelihood estimator $W = W_{\text{ML}}$ from Equation (17), the covariance matrix decomposes as follows:

$$W_{\text{ML}}M^{-1}W_{\text{ML}}^T = U_d(D_d^2 - \sigma^2\mathcal{I}_d)^{\frac{1}{2}}M^{-1}(D_d^2 - \sigma^2\mathcal{I}_d)^{\frac{1}{2}}U_d^T. \quad (30)$$

By computing a $(d \times d)$ -dimensional eigenvalue decomposition of the inner part

$$(D_d^2 - \sigma^2\mathcal{I}_d)^{\frac{1}{2}}M^{-1}(D_d^2 - \sigma^2\mathcal{I}_d)^{\frac{1}{2}} =: \tilde{U}\tilde{D}^2\tilde{U}^T, \quad (31)$$

we get the eigenvalue decomposition

$$W_{\text{ML}}M^{-1}W_{\text{ML}}^T = (U_d\tilde{U})\tilde{D}^2(U_d\tilde{U})^T, \quad (32)$$

without having to calculate a prohibitively large $(3N \times 3N)$ -dimensional PCA.

The main modes of variation of the shapes s^* are the eigenvectors corresponding to the largest eigenvalues. They model those deformations of the shapes causing a maximum deformation of the full shape, while keeping the given part s_b virtually fixed.

By visualizing these modes of variation, we can see how much flexibility remains in the model after fixing s_b and thus how well s_b determines the rest of the shape. For instance, the eigenvector v_1 corresponding to the largest eigenvalue σ_1^2 is the unique deformation with unit norm that changes the full model s as much as possible, while keeping s_b fixed within the limits of the noise modeled by ε . By visualizing $\mu \pm 3\lambda_1v_1$ we can observe 99 % of the variation along this first mode of variation.

The Parameter σ was introduced in Equation (11) as the variance of additive noise assumed to be present in the model. A maximum likelihood estimator was given in Equation (18). When reconstructing partial shapes and examining the remaining flexibility, σ controls how strictly the given part of the model s_b is matched.

The larger σ is chosen, the more the shape is allowed to deviate from s_b . Of course this results in a larger remaining flexibility as even the parts of the shape constrained by s_b are allowed to move slightly.

The maximum likelihood estimator for σ given in Equation (18) is:

$$\sigma_{\text{ML}} = \frac{1}{3N - d} \sum_{i=d+1}^{3N} D_{ii}^2.$$

The number of non-zero eigenvalues is usually small compared to the dimensionality. Therefore, the maximum likelihood solution σ_{ML} becomes small or even zero, which will lead to the covariance matrix M in equation (26) being (close to) singular. This results in little or no remaining flexibility as well as possible overfitting in the reconstructed shapes.

Letting $\sigma \rightarrow 0$ in Equation (17) and (14), leads to $W = U_d D_d$ and thus the use of the sample covariance matrix $U_d D_d^2 U_d^T$ as our covariance estimator. It is well known in statistics that in the “small sample, large dimension” case, the maximum likelihood estimator of the covariance matrix provides a poor estimate of the real covariance matrix. Letting $\sigma^2 > 0$ be a parameter corresponds to the usual shrinkage approach for covariance estimation and can be shown to improve the estimate in such cases (see e.g. Schäfer and Strimmer [22]).

In a practical setting, σ is a very sensible parameter and has the natural interpretation as controlling the balance between matching accuracy and flexibility. So instead of the very small σ_{ML} , it can for instance be chosen to match the measurement error of the capturing device.

4 Results and Medical Applications

We conducted a number of experiments using statistical models of the human face and the femur bone. The experiments show how the model can be used to answer typical questions that arise in clinical practice.

4.1 Experimental Setup

For the femur experiments we used a statistical model built from 50 surfaces of normal femurs. To establish correspondence among the surfaces, we used the non-rigid registration algorithm proposed by Albrecht et al. [5]. For the face experiments we used the data from the Basel Face Model¹, which consists of 200 registered faces, acquired with a structured light scanner. All the experiments have been performed with the parameter $\sigma = 10$. This value has deliberately been chosen relatively large, to make the variations clearly visible in the paper.

¹ Basel face model: <http://faces.cs.unibas.ch>

4.2 Results

For our first experiment we considered the case where the nose is missing in a face and has to be reconstructed. This is a case that actually occurs in clinical practice, for example where a large tumor has to be removed. With our method the reconstruction can be efficiently computed, requiring only a surface scan of the patient. Figure 4.2 shows the reconstruction results as well as the variability represented by the first mode of variation. It can be seen that the reconstruction closely resembles the original nose. This is an indication that the shape of the nose is rather well constrained, given the remaining facial surface. Figure 2(f) shows an extremely unlikely reconstruction (with probability less than 10^{-13}). However, even such an unlikely reconstruction still looks natural.

Our second experiment shows that a valid reconstruction is also possible when only a small part of the face is fixed. Figure 3 shows the reconstruction and the variation captured by the first two principal components. In Figure 3(b) the variability that remains for each point is color coded. The variability σ_{v^i} for the point v^i is defined as

$$\sigma_{v^i} = \sqrt{\text{var}(v_x^i)^2 + \text{var}(v_y^i)^2 + \text{var}(v_z^i)^2},$$

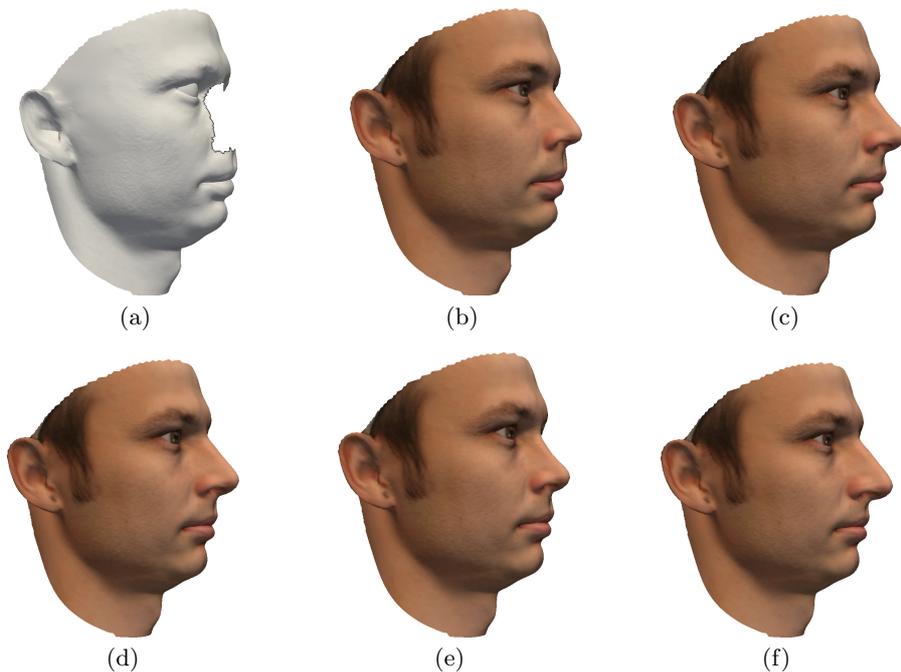


Fig. 2. Reconstruction of a nose: (a) shows the surface with the nose removed. (b) shows the real face while (c) shows the reconstructed nose. In (d) and (e) we see the $\pm 3\sigma$ of the main mode of variation. (f) shows a nose where the first 7 components are 3σ from the mean.

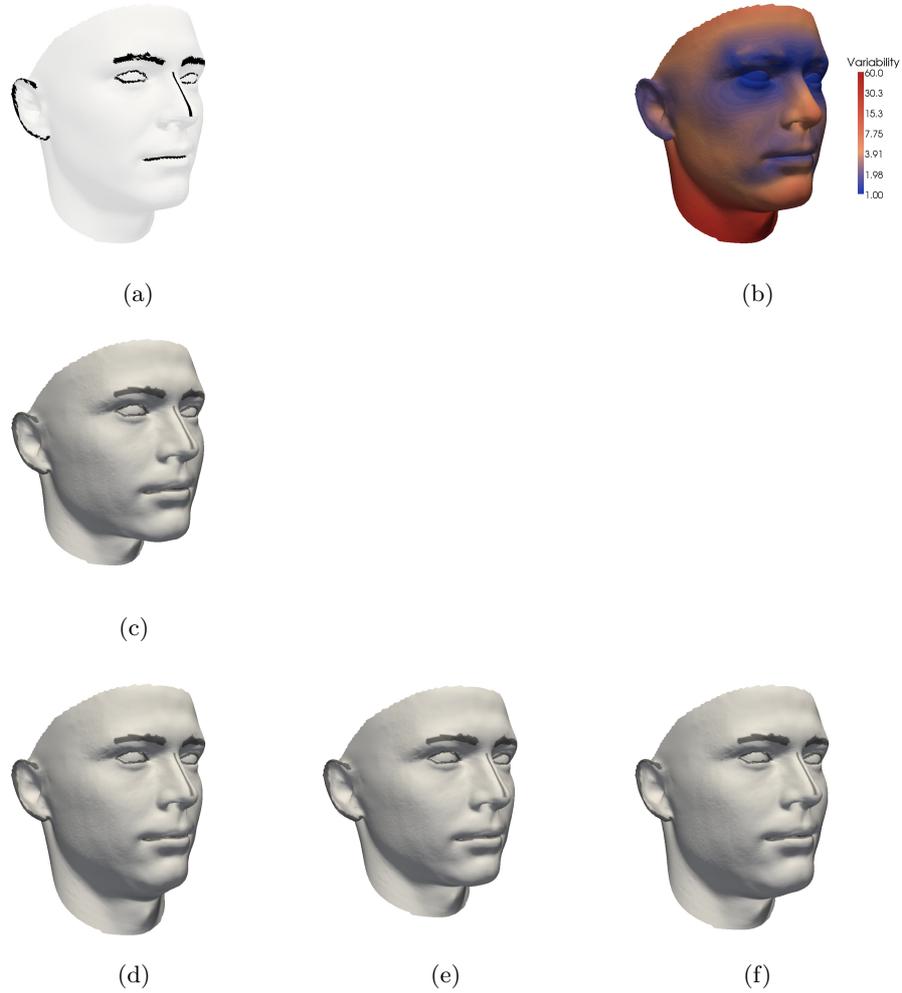


Fig. 3. Reconstruction of the face where only a sketch (a) is given. (b) shows the best reconstruction. The colors encode the variability (in *mm*) at the given point. (d), (c) show $\pm 3\sigma_1$ in the first mode of variation. (f), (e) show $\pm 3\sigma_2$ in the second mode of variation.

i.e. it is the norm of the variance measured in each direction. Of course, the reconstruction from only a sketch shows much more variability than what can be observed in the nose example. The last experiment shows how the model can be used to investigate how well the femur bone is determined by the joint surfaces. This variability can be helpful in prosthesis design. Figure 4 shows the variability in the direction of the first two principal components.

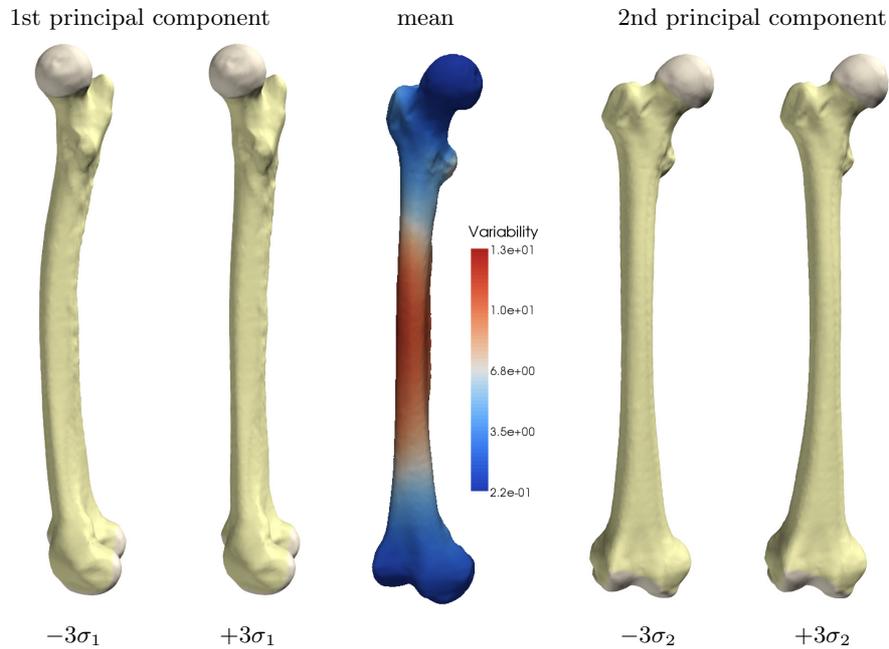


Fig. 4. When a statistical model of the human femur bone is fitted to given joint surfaces (grey), considerable flexibility remains, visualized here by the first two principal components with standard deviation $\sigma_{1,2}$. The joint surfaces are taken from the mean, seen in the middle, colored according to the remaining variability (in *mm*).

4.3 Reconstruction in Practice

For all reconstruction examples that we presented in this section the surfaces were already in correspondence with the model. In practice, establishing the correspondence is a pivotal step that both influences the reconstruction error and the remaining variability. As this paper's main focus is on modeling the remaining flexibility and not the reconstruction of missing parts, we refer the reader to the article by Basso and Vetter [13] for a more thorough evaluation of the actual reconstruction using an equivalent method.

5 Conclusion

We presented a method for computing and visualizing the remaining flexibility in statistical shape models, when part of the shape is known. To model the shape variability we use a probabilistic model based on PPCA. The flexibility remaining in the model can be computed from the posterior distribution arising from the PPCA model. We proposed to compute the remaining flexibility by evaluating the principal components of this posterior distribution and showing

the effect that changing the corresponding coefficients has on the shape. We presented experiments that illustrate typical applications of our model. Our results for shape reconstruction are similar to those achieved in comparable reconstruction approaches. However, in contrast to other methods, we also showed in our experiments the full range of possible reconstructions that complete the given partial shape. Furthermore, our method allows us to assign a probability to every reconstruction.

The main application of our method is to gain more insight into the information represented by a morphable model, and learn more about the statistical properties of the surfaces. The model allows us to investigate how strictly a given part determines a shape. This is of particular interest in the medical domain, where such questions frequently arise in the planning of reconstructive surgeries as well as the designing of prosthesis.

In future work we will investigate the question whether it is possible to automatically find the parts of the surface that best determine its shape. A particular application we have in mind is to use this information in face modelling, for investigating which parts of the face determine the identity of a person and which parts constitute the “remaining flexibility”.

Acknowledgements

This work was funded by the Swiss National Science Foundation in the scope of the NCCR CO-ME project 5005-66380 and the Hasler Foundation in scope of the HOVISSE project. The face data is courtesy of Pascal Paysan (Figure 2(b)), University of Basel.

References

1. Blanz, V., Vetter, T.: A morphable model for the synthesis of 3D faces. In: SIGGRAPH 1999: Proceedings of the 26th annual conference on Computer graphics and interactive techniques, pp. 187–194. ACM Press, New York (1999)
2. Leventon, M.E., Grimson, W.E.L., Faugeras, O.: Statistical shape influence in geodesic active contours. In: CVPR, vol. 1, p. 1316 (2000)
3. Tsai, A., Yezzi, A., Wells, J.W., Tempany, C., Tucker, D., Fan, A., Grimson, W.E., Willsky, A.: A shape-based approach to the segmentation of medical imagery using level sets. *IEEE Transactions on Medical Imaging* 22(2), 137–154 (2003)
4. Cremers, D., Rousson, M., Deriche, R.: A Review of Statistical Approaches to Level Set Segmentation: Integrating Color, Texture, Motion and Shape. *International Journal of Computer Vision* 72(2), 195–215 (2007)
5. Albrecht, T., Lüthi, M., Vetter, T.: A statistical deformation prior for non-rigid image and shape registration. In: *IEEE Conference on CVPR*, pp. 1–8 (2008)
6. Wang, Y., Staib, L.: Physical model-based non-rigid registration incorporating statistical shape information. *Medical Image Analysis* 4(1), 7–20 (2000)
7. Benameur, S., Mignotte, M., Destrempes, F., De Guise, J.: Three-dimensional bi-planar reconstruction of scoliotic rib cage using the estimation of a mixture of probabilistic prior models. *IEEE Transactions on Biomedical Engineering* 52(10), 1713–1728 (2005)

8. Lamecker, H., Wenckeback, T., Hege, H.: Atlas-based 3D-shape reconstruction from X-ray images. In: Proceedings of the 18th International Conference on Pattern Recognition (ICPR 2006), vol. 1, pp. 371–374 (2006)
9. Fleute, M., Lavallee, S.: Nonrigid 3-D/2-D Registration of Images Using Statistical Models. LNCS, pp. 138–147. Springer, Heidelberg (1999)
10. Roweis, S.: EM Algorithms for PCA and SPCA. In: NIPS, pp. 626–632 (1998)
11. Tipping, M.E., Bishop, C.M.: Probabilistic principal component analysis. *Journal of the Royal Statistical Society* 61, 611–622 (1999)
12. Blanz, V., Mehl, A., Vetter, T., Seidel, H.P.: A statistical method for robust 3d surface reconstruction from sparse data. In: International Symposium on 3D Data Processing, Visualization and Transmission (2004)
13. Basso, C., Vetter, T.: Statistically Motivated 3D Faces Reconstruction. In: Proceedings of the 2nd International Conference on Reconstruction of Soft Facial Parts, p. 71 (2005)
14. Hwang, B., Lee, S.: Reconstructing a Whole Face Image from a Partially Damaged or Occluded Image by Multiple Matching. In: Lee, S.-W., Li, S.Z. (eds.) ICB 2007. LNCS, vol. 4642, pp. 692–701. Springer, Heidelberg (2007)
15. Jensen, K., Sparring, J.: Reconstructing Teeth with Bite Information. In: Ersbøll, B.K., Pedersen, K.S. (eds.) SCIA 2007. LNCS, vol. 4522, pp. 102–111. Springer, Heidelberg (2007)
16. Gong, X., Wang, G.: A Novel Deformation Framework for Face Modeling from a Few Control Points. In: Wang, G., Li, T., Grzymala-Busse, J.W., Miao, D., Skowron, A., Yao, Y. (eds.) RSKT 2008. LNCS, vol. 5009, pp. 434–441. Springer, Heidelberg (2008)
17. Machado, A., Gee, J., Campos, M.: Visual data mining for modeling prior distributions in morphometry. *IEEE Signal Processing Magazine* 21(3), 20–27 (2004)
18. Albrecht, T., Knothe, R., Vetter, T.: Modeling the Remaining Flexibility of Partially Fixed Statistical Shape Models. In: Workshop on the Mathematical Foundations of Computational Anatomy, MFCA 2008, New York, USA (2008)
19. Cootes, T., Taylor, C., Cooper, D., Graham, J., et al.: Active shape models-their training and application. *Computer Vision and Image Understanding* 61(1), 38–59 (1995)
20. Audette, M.A., Ferrie, F.P., Peters, T.M.: An algorithmic overview of surface registration techniques for medical imaging. *Medical Image Analysis* 4, 201–217 (2000)
21. Bishop, C.: *Pattern recognition and machine learning*. Springer, Heidelberg (2006)
22. Schäfer, J., Strimmer, K.: A Shrinkage Approach to Large-Scale Covariance Matrix Estimation and Implications for Functional Genomics. *Statistical Applications in Genetics and Molecular Biology* 4(1), 32 (2005)