Modeling the Remaining Flexibility of Partially Fixed Statistical Shape Models

Thomas Albrecht, Reinhard Knothe, and Thomas Vetter

Computer Science Department, University of Basel, Switzerland
{Thomas.Albrecht,Reinhard.Knothe,Thomas.Vetter}@unibas.ch

Abstract. Statistical shape models are widely used to model the variability of biological shapes. They can be used to reconstruct missing information given partial or noisy data. In case of partial data, many different reconstructions are possible, and one is not only interested in a plausible reconstruction but also the remaining flexibility within the model and the reliability of the reconstruction. We present a method to model the remaining flexibility when some part of a statistical shape model is fixed. Using such a flexibility model, we can give answers to questions like: Does one half of a human femur bone determine the other half? or How much is the shape of a face determined by its contour?

1 Introduction

Statistical shape models are widely used in medical image analysis, computational anatomy, and computer vision to model the variability of biological shapes, see [1–6] for instance. The variability of a certain class of shapes is deduced from a representative set of example data from this class.

If the example data sets represent the class of shapes well, the model can be fitted to virtually any individual shape within the class. It is even possible to fit the model to partial data from an individual, [7, 8]. The missing remaining data is automatically reconstructed by the model, yielding the most plausible reconstruction.

However, there may be many other possible reconstructions which fit the partial data equally well. The partial data only determines a part of the model, while the rest of the model may remain flexible. This paper focuses on modeling this remaining flexibility of the partially determined model.

Keeping a part $\mathbf{x}_b$ of the complete model $\mathbf{x}$ fixed, how much flexibility remains for the remaining part $\mathbf{x}_a$? In principle, as PCA models can be statistically interpreted by a multivariate Gaussian distribution with probability density function p, we can model the remaining flexibility by the conditional distribution $p(\mathbf{x}_a|\mathbf{x}_b)$. We will see however, that there is no nontrivial conditional distribution if more components are kept fixed than there are degrees of freedom in the model, as is usually the case in models built from a small set of examples.

In this case, we propose a method which models the remaining flexibility of the variable points when the fixed points are allowed to move slightly instead of

being completely fixed. The method leads to a generalized eigenvalue problem which can be solved efficiently.

The remaining paper is organized as follows. In Section 2 we will review the well-known concept of PCA-based statistical shape models. In Section 3 we will derive an expression of the conditional distribution $p(\mathbf{x}_a|\mathbf{x}_b)$ and investigate when this can actually be used. In Section 4, we will model the remaining flexibility when the fixed points are allowed to move slightly. In Section 5 we will use the proposed method to model the remaining flexibility of a shape model of the human femur bone when its distal part is fixed, and the remaining flexibility of a face model, when its contour is fixed.

Related Work. Virtually all cited papers deal with statistical shape models with applications in medicine or computer vision, [1–8]. Most notably, [5] uses Canonical Correlation Analysis to model the correlation between different parts of a statistical model, computing model coefficients which maximize the correlation between these parts. However, they do not address our problem of modeling the remaining flexibility when one part of the model is fixed.

2 Principle Component Analysis

The 3D statistical shape models used in this paper are based on a Principle Component Analysis (PCA) of a set of training data comprised of n 3D surfaces. Each surface is represented by a triangular mesh with the same number $m \in \mathbb{N}$ of vertices, which are stacked into a data vector $\mathbf{x} = (x_1, y_1, z_1, ...x_m, y_m, z_m)^T \in \mathbb{R}^M$ with $M = 3m$. The training surfaces need to be in correspondence.

As in all PCA models, the samples are assumed to be i.i.d. samples drawn from a multivariate normal distribution that is approximated by the estimated multivariate normal distribution $\mathcal{N}(\bar{\mathbf{x}}, \mathbf{\Sigma})$. The mean is estimated by the arithmetic mean $\bar{\mathbf{x}} = \frac{1}{n}\sum_{i=1}^{n} \mathbf{x}_i$ of all samples. The covariance matrix $\mathbf{\Sigma} \in \mathbb{R}^{M\times M}$ can be estimated from the mean-free data matrix $\mathbf{X} := [\mathbf{x}_1 - \bar{\mathbf{x}}, \ldots, \mathbf{x}_n - \bar{\mathbf{x}}] \in \mathbb{R}^{M\times n}$ as $\mathbf{\Sigma} = \frac{1}{n}\mathbf{X}\mathbf{X}^T$.

When $\mathbf{X}$ is decomposed with a (reduced) Singular Value Decomposition $\mathbf{X} = \mathbf{U}\mathbf{W}\mathbf{V}^T$ into the product of a column-orthonormal matrix $\mathbf{U} \in \mathbb{R}^{M\times n}$, a diagonal matrix $\mathbf{W} \in \mathbb{R}^{n\times n}$, and an orthonormal matrix $\mathbf{V} \in \mathbb{R}^{n\times n}$, the covariance matrix can be expressed as $\mathbf{\Sigma} = \frac{1}{n}\mathbf{U}\mathbf{W}^2\mathbf{U}^T$.

The columns $\mathbf{u}_i$ of the matrix $\mathbf{U}$ are the eigenvectors of $\mathbf{\Sigma}$. They are known as the principal components of the model and describe the main modes of variation of the training data. Their corresponding eigenvalues $\sigma_i^2 := \frac{1}{n}w_i^2$ describe the variance of the model projected onto these eigenvalues. They are arranged according to size so that $\mathbf{u}_1$ is the direction with maximal projected variance σ_1^2.

One individual $\mathbf{x}$ in the object class modeled by the PCA model can be identified by its coefficients $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_n)^T \in \mathbb{R}^n$:

$$\mathbf{x} = \bar{\mathbf{x}} + \sum_{i=1}^{n} \alpha_i \sigma_i \mathbf{u}_i = \bar{\mathbf{x}} + \tfrac{1}{\sqrt{n}}\mathbf{U}\mathbf{W}\boldsymbol{\alpha}. \tag{1}$$

Under the assumption that the data $\mathbf{x}$ is distributed according to a multivariate normal distribution $\mathcal{N}(\bar{\mathbf{x}}, \mathbf{\Sigma})$, the coefficient vector $\boldsymbol{\alpha}$ is distributed according to $\mathcal{N}(\mathbf{0}, \mathbf{I}_n)$.

In this paper two PCA models are used: A model of the femur bone, built from 21 CT scans of femur bones, which were hand-segmented and brought into correspondence with [9], and a model of the human face built from 100 face scans that were brought in correspondence with a modification of [10].

3 Conditional Distribution

In a PCA-based statistical model, we now wish to fix a certain number $l \in \mathbb{N}$ of points in the model in order to investigate how flexible the model remains with respect to the remaining $m - l$ variable points. This means that we fix $L = 3l$ components of the model vector $\mathbf{x}$. Without loss of generality, we can assume that they are the last l components and $\mathbf{x}$ can be partitioned as $\mathbf{x} = (\mathbf{x}_a, \mathbf{x}_b)^T$.

As $\mathbf{x}$ is distributed according to a multivariate normal distribution, the conditional distribution is a also a multivariate normal distribution $\mathcal{N}(\boldsymbol{\mu}_{a|b}, \mathbf{\Sigma}_{a|b})$. Its mean and covariance can be calculated from $\bar{\mathbf{x}}$ and $\mathbf{\Sigma}$.

The matrix of principal components $\mathbf{U} \in \mathbb{R}^{M \times n}$ can be partitioned according to $\mathbf{x}$. For simplicity's sake, we define the matrix of the principal components scaled by the diagonal matrix $\mathbf{W} \in \mathbb{R}^{n \times n}$ as $\mathbf{Q} := \frac{1}{\sqrt{n}}\mathbf{U}\mathbf{W} \in \mathbb{R}^{M \times n}$. We have

$$\mathbf{U} = \begin{pmatrix} \mathbf{U}_a \\ \mathbf{U}_b \end{pmatrix}, \qquad \mathbf{Q} = \tfrac{1}{\sqrt{n}}\mathbf{U}\mathbf{W} = \tfrac{1}{\sqrt{n}} \begin{pmatrix} \mathbf{U}_a\mathbf{W} \\ \mathbf{U}_b\mathbf{W} \end{pmatrix} = \begin{pmatrix} \mathbf{Q}_a \\ \mathbf{Q}_b \end{pmatrix} \tag{2}$$

The covariance matrix $\mathbf{\Sigma} \in \mathbb{R}^{M \times M}$ can be calculated from $\mathbf{Q}$ by $\mathbf{\Sigma} = \frac{1}{n}\mathbf{U}\mathbf{W}^2\mathbf{U}^T = \mathbf{Q}\mathbf{Q}^T$. We can partition $\mathbf{\Sigma}$ as follows:

$$\mathbf{\Sigma} = \begin{pmatrix} \mathbf{\Sigma}_{aa} & \mathbf{\Sigma}_{ab} \\ \mathbf{\Sigma}_{ba} & \mathbf{\Sigma}_{bb} \end{pmatrix} = \begin{pmatrix} \mathbf{Q}_a\mathbf{Q}_a^T & \mathbf{Q}_a\mathbf{Q}_b^T \\ \mathbf{Q}_b\mathbf{Q}_a^T & \mathbf{Q}_b\mathbf{Q}_b^T \end{pmatrix}. \tag{3}$$

Note that it is only possible to calculate a nontrivial conditional distribution if the matrix $\mathbf{\Sigma}_{bb}$ is invertible. A necessary condition for $\mathbf{\Sigma}_{bb}$ to be invertible is that the number of fixed degrees of freedom L is less the number of training examples n of the statistical model. Let us, for the moment, assume that $\mathbf{\Sigma}_{bb}$ is invertible and calculate the conditional distribution according. According to [11], the covariance matrix $\mathbf{\Sigma}_{a|b}$ of the can be expressed as:

$$\begin{aligned} \mathbf{\Sigma}_{a|b} &= \mathbf{\Sigma}_{aa} - \mathbf{\Sigma}_{ab}\,\mathbf{\Sigma}_{bb}^{-1}\,\mathbf{\Sigma}_{ba} && (4)\\ &= \mathbf{Q}_a\mathbf{Q}_a^T - \mathbf{Q}_a\mathbf{Q}_b^T\,(\mathbf{Q}_b\mathbf{Q}_b^T)^{-1}\,\mathbf{Q}_b\mathbf{Q}_a^T && (5)\\ &= \mathbf{Q}_a\left(\mathbf{I}_n - \mathbf{Q}_b^T(\mathbf{Q}_b\mathbf{Q}_b^T)^{-1}\mathbf{Q}_b\right)\mathbf{Q}_a^T, && (6) \end{aligned}$$

where $\mathbf{I_n}$ denotes the $(n \times n)$ identity matrix. While the original matrix $\mathbf{U}$ is column-orthonormal, the sub-matrix $\mathbf{U}_b$ defined in Equation (2) in not. Therefore, we perform an additional a singular value decomposition $\mathbf{Q}_b = \bar{\mathbf{U}}_b \mathbf{W}_b \mathbf{V}_b^T$ with a column-orthonormal $\bar{\mathbf{U}}_b$. Thus, the expression can be expanded to:

$$\begin{aligned} \mathbf{\Sigma}_{a|b} &= \mathbf{Q}_a \left(\mathbf{I}_n - \mathbf{V}_b \mathbf{W}_b \bar{\mathbf{U}}_b^T \bar{\mathbf{U}}_b \mathbf{W}_b^{-2} \bar{\mathbf{U}}_b^T \bar{\mathbf{U}}_b \mathbf{W}_b \mathbf{V}_b \right) \mathbf{Q}_a^T && (7) \\ &= \mathbf{Q}_a \left(\mathbf{I}_n - \mathbf{V}_b^T \mathbf{V}_b \right) \mathbf{Q}_a^T. && (8) \end{aligned}$$

Where, in the last step, we have used $\bar{\mathbf{U}}_b^T \bar{\mathbf{U}}_b = \mathbf{I}_n$. Similarly, the mean $\boldsymbol{\mu}_{a|b}$ is given as:

$$\begin{aligned} \boldsymbol{\mu}_{a|b} &= \bar{\mathbf{x}}_a + \mathbf{\Sigma}_{ab} \mathbf{\Sigma}_{bb}^{-1} (\mathbf{x}_b - \bar{\mathbf{x}}_b) && (9) \\ &= \bar{\mathbf{x}}_a + \mathbf{Q}_a \mathbf{V}_b \mathbf{W}_b^{-1} \bar{\mathbf{U}}_b^T (\mathbf{x}_b - \bar{\mathbf{x}}_b), && (10) \end{aligned}$$

with $\bar{\mathbf{x}} = (\bar{\mathbf{x}}_a, \bar{\mathbf{x}}_b)^T$.

Provided that $\mathbf{\Sigma}_{bb}$ is invertible, $\mathbf{x}_a$ can be reconstructed from $\mathbf{x}_b$ as $\boldsymbol{\mu}_{a|b}$. This is the reconstruction presented in [7] as the maximum a posteriori reconstruction. All other reconstructions which fit $\mathbf{x}_b$ are modeled by $\mathcal{N}(\boldsymbol{\mu}_{a|b}, \mathbf{\Sigma}_{a|b})$. The more flexibility this distribution allows, the less reliable the reconstruction by $\boldsymbol{\mu}_{a|b}$ is.

For each of these reconstructions, the fixed values $\mathbf{x}_b$ are matched equally well. Indeed, if we take a closer look at the expression for the covariance matrix $\mathbf{\Sigma}_{a|b}$ in Equation (8) we notice that the inner part $(\mathbf{I}_n - \mathbf{V}_b^T \mathbf{V}_b)$ is a projection onto the orthogonal complement of the column space of $\mathbf{V}_b$. As $\mathbf{V}_b$ is the "input matrix" of the SVD of $\mathbf{Q}_b$, this projection is a projection onto the kernel of $\mathbf{Q}_b$. This means that the distribution $\mathcal{N}(\boldsymbol{\mu}_{a|b}, \mathbf{\Sigma}_{a|b})$ models only linear combinations $\mathbf{Q}_a \boldsymbol{\alpha}$ of the scaled principal components for which $\boldsymbol{\alpha}$ is in the kernel of $\mathbf{Q}_b$, i.e. $\mathbf{Q}_b \boldsymbol{\alpha} = 0$. So for all $\boldsymbol{\alpha} \sim \mathcal{N}(\boldsymbol{\mu}_{a|b}, \mathbf{\Sigma}_{a|b})$, the deformation of the fixed points is zero.

Invertibility of $\mathbf{\Sigma_{bb}}$. So far, we have assumed that $\mathbf{\Sigma}_{bb} = \mathbf{Q}_b \mathbf{Q}_b^T \in \mathbb{R}^{L \times L}$ is invertible, i.e. $\operatorname{rank}(\mathbf{\Sigma}_{bb}) = L$. However, if $\mathbf{Q}_b$ has less than L linearly independent columns we have $\operatorname{rank}(\mathbf{\Sigma}_{bb}) < L$. In particular, if the statistical model is built from less than L examples or less than L principal components are used, the rank of $\mathbf{\Sigma}_{bb}$ will be less than L and the above calculations involving an inversion of $\mathbf{\Sigma}_{bb}$ are not valid. In this case there is no nontrivial conditional distribution $p(\mathbf{x}_a, \mathbf{x}_b)$.

In this case the mean $\boldsymbol{\mu}_{a|b}$ can still be approximated by using the pseudo-inverse of $\mathbf{\Sigma}_{bb}$ instead of its inverse, [7]. However, if we try to use the pseudo-inverse for calculating the covariance matrix according to Equation (4), we get $\mathbf{\Sigma}_{a|b} = 0$.

The problem is that the conditional distribution $\mathcal{N}(\boldsymbol{\mu}_{a|b}, \mathbf{\Sigma}_{a|b})$ models only coefficient vectors from the kernel of $\mathbf{Q}_b$. Without a nontrivial kernel of $\mathbf{Q}_b$, it is not possible to calculate a nontrivial conditional distribution. Therefore, we propose that instead of looking for coefficient vectors $\boldsymbol{\alpha} \in \ker \mathbf{Q}_b$, i.e. $\mathbf{Q}_b \boldsymbol{\alpha} = 0$,

we will look for coefficients for which $\mathbf{Q}_b\boldsymbol{\alpha}$ is small. This means that we relax the constraint of keeping the fixed points completely fixed to allowing them to move slightly.

It also has to be noted that even if there is a nontrivial kernel of $\mathbf{Q}_b$, it will most likely still be of interest to consider deformations which allow slight changes of the fixed points $\mathbf{x}_b$ when considering the remaining flexibility of the model.

4 Generalized Eigenvalues

The aim is to model the flexibility of the variable points $\mathbf{x}_a$ when the fixed points $\mathbf{x}_b$ are allowed to move slightly. The deformations are given as $\mathbf{Q}_a\boldsymbol{\alpha}$ and $\mathbf{Q}_b\boldsymbol{\alpha}$. A measure for the change caused by these deformations is the squared Euclidean norm of these vectors:

$$\|\mathbf{Q}_a\boldsymbol{\alpha}\|^2 = \boldsymbol{\alpha}^T\mathbf{Q}_a^T\mathbf{Q}_a\boldsymbol{\alpha}, \quad \|\mathbf{Q}_b\boldsymbol{\alpha}\|^2 = \boldsymbol{\alpha}^T\mathbf{Q}_b^T\mathbf{Q}_b\boldsymbol{\alpha}. \tag{11}$$

As we are first and foremost interested in the coefficients $\boldsymbol{\alpha}$ which change the variable points $\mathbf{x}_a$ as much as possible, we can formulate our aim as a constrained maximization problem:

$$\max_{\boldsymbol{\alpha}\in\mathbb{R}^k} \boldsymbol{\alpha}^T\mathbf{Q}_a^T\mathbf{Q}_a\boldsymbol{\alpha} \tag{12}$$

$$\text{subject to } \boldsymbol{\alpha}^T\mathbf{Q}_b^T\mathbf{Q}_b\boldsymbol{\alpha} = c, \tag{13}$$

where $c \in \mathbb{R}^+$ quantifies the amount of change allowed in the fixed coefficients.

Introducing a Lagrangian multiplier λ and differentiating with respect to $\boldsymbol{\alpha}$ leads to the generalized eigenvalue problem:

$$\mathbf{Q}_a^T\mathbf{Q}_a\boldsymbol{\alpha} = \lambda\,\mathbf{Q}_b^T\mathbf{Q}_b\boldsymbol{\alpha}. \tag{14}$$

Both matrices $\mathbf{Q}_a^T\mathbf{Q}_a$ and $\mathbf{Q}_b^T\mathbf{Q}_b$ are positive definite and symmetric. The generalized eigenvalue problem can be solved efficiently with standard software (LAPACK, MATLAB), yielding a set of generalized eigenvectors $\{\boldsymbol{\alpha}_1, \ldots, \boldsymbol{\alpha}_k\}$ arranged according to the size of their corresponding generalized eigenvalues $\{\lambda_1, \ldots, \lambda_k\}$. The eigenvectors are scaled so that $\|\mathbf{Q}_b\boldsymbol{\alpha}_i\|^2 = \boldsymbol{\alpha}_i^T\mathbf{Q}_b^T\mathbf{Q}_b\boldsymbol{\alpha}_i = 1$ for $i = 1, \ldots, k$. If we pre-multiply Equation (14) by $\boldsymbol{\alpha}^T$, we see that for an eigenvector $\boldsymbol{\alpha}_i$ and its eigenvalue λ_i, we have

$$\|\mathbf{Q}_a\boldsymbol{\alpha}_i\|^2 = \lambda_i\|\mathbf{Q}_b\boldsymbol{\alpha}_i\|^2. \tag{15}$$

This means that, measured in the squared Euclidean norm, the deformation determined by the coefficient vector $\boldsymbol{\alpha}_i \in \mathbb{R}^k$ changes the variable model points λ_i times as much the fixed ones. Therefore, the eigenvector $\boldsymbol{\alpha}_1$ corresponding to the largest eigenvalue λ_1 is the coefficient vector which causes the largest change on the variable points (with a squared Euclidean norm of λ_1), changing the fixed points only slightly (with a deformation with a squared Euclidean norm of 1). The last eigenvectors change the fixed points more than the variable ones. In

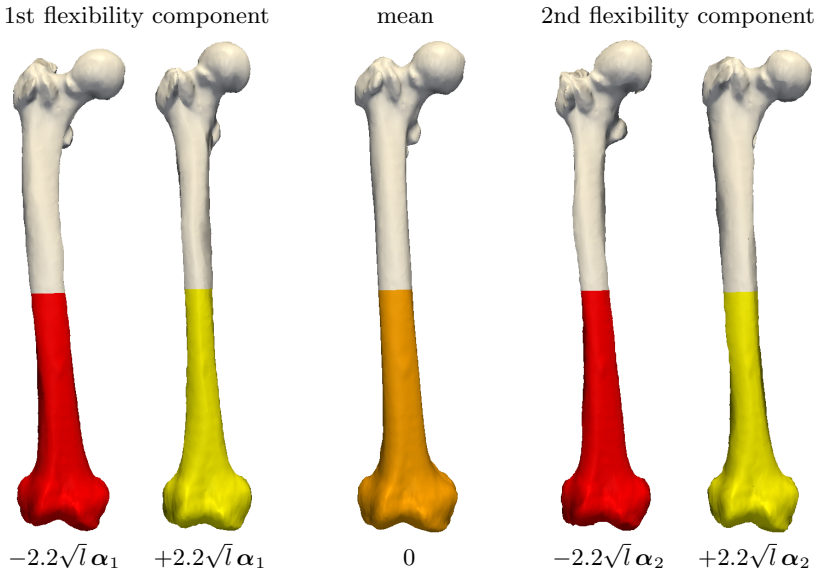

Fig. 1. In a statistical shape model of the human femur bone, the two first flexibility components $\boldsymbol{\alpha}_1$ and $\boldsymbol{\alpha}_2$ model those deformations that change the proximal part (gray) as much as possible, while changing the distal part (colored) as little as possible.

fact, exchanging fixed and variable points leads to the same eigenvalues, only inverted.

We call the coefficient vectors $\boldsymbol{\alpha}_i$ the flexibility components. $\boldsymbol{\alpha}_1$ is the first flexibility component. Just like the principal components of the PCA model, they can be scaled and added together to achieve different deformations from the mean. When a flexibility component $\boldsymbol{\alpha}_i$ is scaled by $\sqrt{r}$ it causes a deformation of squared Euclidean norm r of the fixed points and of $\lambda_i\, r$ of the variable points. The larger the eigenvalues λ_i are, the more flexibility remains for the fixed points and the less reliable a reconstruction of $\mathbf{x}_a$ from $\mathbf{x}_b$ is considered. In order to meat the constraint from Equation (13), the vectors have to be scaled by $\sqrt{c}$.

The squared Euclidean norm is not a very intuitive measure of the deformation as it implicitly depends on the number l of fixed points. A more intuitive measure is the mean squared norm, averaged over all l fixed points:

$$\frac{1}{l}\sum_{k=1}^{l}(\mathbf{Q}_b\boldsymbol{\alpha}_i)_k^2 = \frac{1}{l}\|\mathbf{Q}_b\boldsymbol{\alpha}_i\|^2. \tag{16}$$

Therefore, if we scale the coefficient vectors $\boldsymbol{\alpha}_i$ by $\sqrt{l}$, and the coordinates of the model are given in millimeters, the squared deformation of one model point is 1 millimeter on average.

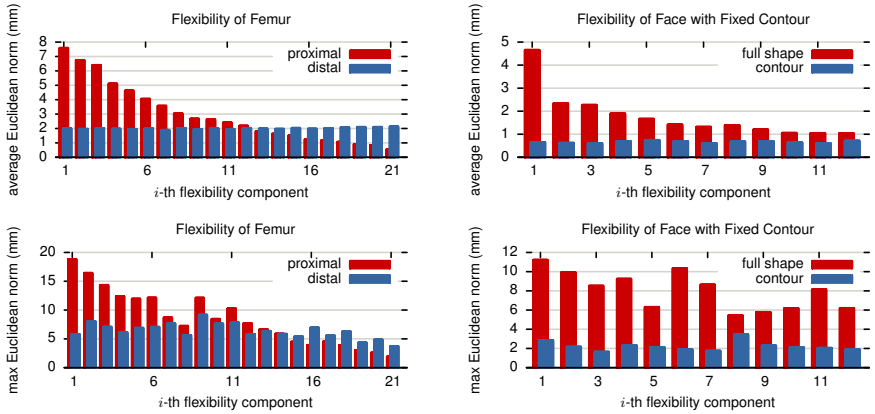

Fig. 2. For the examples shown in Figures 4 and 1, the deformation of the fixed vertices (blue) is compared with the flexibility of the full shape (red). The first flexibility components strongly affect the full shape, while they hardly alter the fixed vertices.

Regularization. The flexibility components $\boldsymbol{\alpha}_i$ are calculated as generalized eigenvectors. As outlined in Section 2, the coefficients $\boldsymbol{\alpha}$ of the PCA model are distributed according to a multivariate normal distribution $\mathcal{N}(0, \mathbf{I}_n)$. Without regularization, it is possible that some of the entries of $\boldsymbol{\alpha}_i$ can be extremely large. In terms of the distribution $\mathcal{N}(0, \mathbf{I}_n)$, such extremely large values are very unlikely and in practice they cause unnatural deformations of the modeled shape.

Up to a normalizing factor, the probability of a deformation caused by $\boldsymbol{\alpha}$ can be calculated by $e^{-\frac{1}{2}\|\alpha\|^2}$. By minimizing $\|\boldsymbol{\alpha}\|^2 = \boldsymbol{\alpha}^T\boldsymbol{\alpha}$, the probability of the deformation is maximized. Therefore, we propose replacing the constraint in Equation (13) by the following constraint:

$$\boldsymbol{\alpha}^T \mathbf{Q}_b^T \mathbf{Q}_b \boldsymbol{\alpha} + \eta\, \boldsymbol{\alpha}^T \boldsymbol{\alpha} = c, \tag{17}$$

with a regularizing parameter $\eta \in \mathbb{R}^+$. In this way, we limit not only the deformation of the fixed points given by $\|\mathbf{Q}_b\boldsymbol{\alpha}\|^2$ but also the improbability of the deformation, given by $\|\boldsymbol{\alpha}\|^2$. The corresponding generalized eigenvalue problem is a regularized version of Equation (14) and is given by:

$$\mathbf{Q}_a^T \mathbf{Q}_a \boldsymbol{\alpha} = \lambda\, (\mathbf{Q}_b^T \mathbf{Q}_b + \eta \mathbf{I}_n)\boldsymbol{\alpha}. \tag{18}$$

This system no longer admits extremely large values in the generalized eigenvectors $\boldsymbol{\alpha}_i$. The parameter η controls the balance between the original problem and the regularizing effect of allowing only probable shapes.

Note that when we deform not the mean but a certain individual $\mathbf{x}$, which is defined by coefficients $\boldsymbol{\beta}$, the deformation coefficients $\boldsymbol{\alpha}$ are not distributed according to $\mathcal{N}(0, \mathbf{I}_n)$ but according to $\mathcal{N}(-\beta, \mathbf{I}_n)$. Therefore it would make more sense to minimize $\|\boldsymbol{\alpha} + \boldsymbol{\beta}\|^2 = (\boldsymbol{\alpha} + \boldsymbol{\beta})^T(\boldsymbol{\alpha} + \boldsymbol{\beta})$. However, it is not obvious how to include such a constraint into a generalized eigenvalue problem like

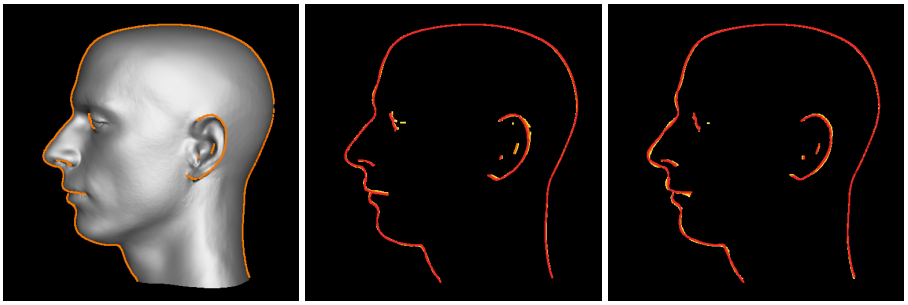

(a) (b) (c)

Fig. 3. The vertices on the contour in profile view are fixed (a). The 1st (b) and 2nd (c) flexibility components (see Figure 4) don't change the contour in profile view.

Equation (18). Therefore, we content ourselves with the proposed regularization, which penalizes large deformations, even though its statistical motivation is perfectible.

Nontrivial kernel of $\mathbf{Q}_b$. We have introduced the generalized eigenvalue problem Equation (14) in order to model the remaining flexibility of the variable points when it is not possible to calculate the conditional distribution $p(\mathbf{x}_a|\mathbf{x_b})$, which is the case if and only if $\mathbf{Q}_b$ has only the trivial kernel $\ker \mathbf{Q}_b = \{\mathbf{0}\}$.

But what happens to the generalized eigenvalue problem if $\mathbf{Q}_b$ does have a nontrivial kernel? In this case, we can split up the space of all coefficients $\boldsymbol{\alpha}$ into the kernel of $\mathbf{Q}_b$ and its orthogonal complement. In the kernel, we have all deformations which do not change the fixed points at all, which can be modeled by the conditional distribution. The covariance matrix in the kernel is simply $\mathbf{Q}_a\mathbf{Q}_a^T$, as in the kernel, the projection term from Equation (8), $(\mathbf{I}_n - \mathbf{V}_b^T\mathbf{V}_b) = \mathbf{I}_n$. In the complement, we can compute the generalized eigenvalue problem in order to additionally allow deformations which change the fixed points slightly.

From a practical point of view, when $\ker \mathbf{Q}_b \neq \{\mathbf{0}\}$ and therefore $\mathbf{Q}_b^T\mathbf{Q}_b$ is singular, the aforementioned LAPACK or MATLAB routines return a basis for the kernel as generalized eigenvectors with eigenvalues infinity, as for these vectors we have $\mathbf{Q}_b^T\mathbf{Q}_b\boldsymbol{\alpha}_i = 0$. The remaining eigenvalues span the complement of the kernel and are computed as usual.

If we use the regularized form of the generalized eigenvalue problem Equation (18), there will be no infinite eigenvalues, as the matrix $(\mathbf{Q}_b^T\mathbf{Q}_b + \eta\mathbf{I}_n)$ is always nonsingular. However, for vectors $\boldsymbol{\alpha} \in \ker \mathbf{Q}_b$, we have $(\mathbf{Q}_b^T\mathbf{Q}_b + \eta\mathbf{I}_n)\boldsymbol{\alpha} = \eta\boldsymbol{\alpha}$. Therefore, on $\ker \mathbf{Q}_b$, we effectively solve the eigenvalue problem:

$$\boldsymbol{\alpha}^T\mathbf{Q}_a^T\mathbf{Q}_a\boldsymbol{\alpha} = \lambda\,\eta\boldsymbol{\alpha}, \tag{19}$$

whose eigenvectors corresponding to the largest eigenvalues maximize $\|\mathbf{Q}_a\|^2$. In this regularized case, there is no strict decomposition into $\ker \mathbf{Q}_b$ and its complement.

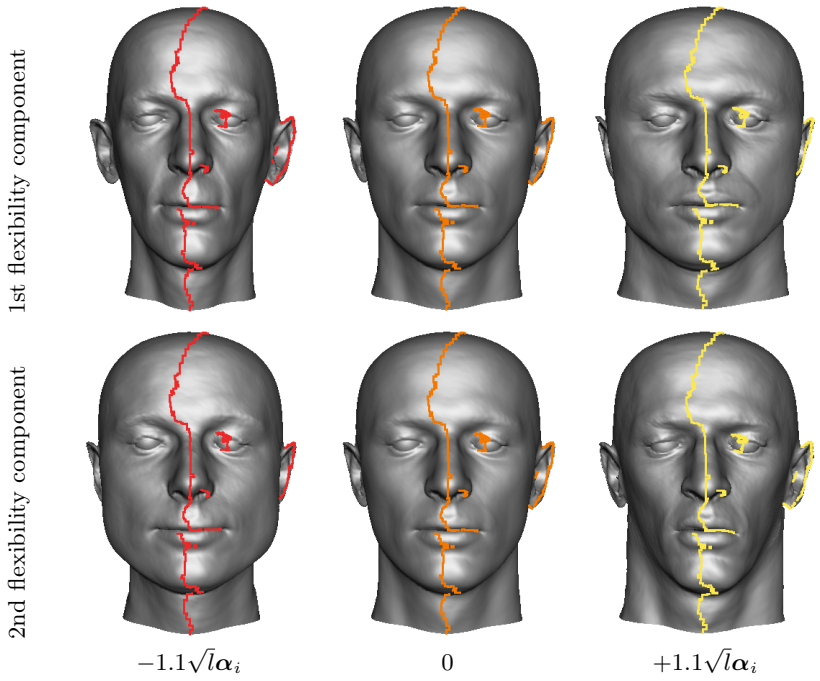

Fig. 4. A statistical shape model of human faces is used to model the remaining flexibility with fixed contour. The vertices on the contour in profile view (see also Figure 3) are fixed. The flexibility components model the remaining variability within the model. The 1st/2nd flexibility components is shown in the 1st/2nd row.

5 Experiments and Conclusion

We demonstrate the use of the flexibility models for two different scenarios with two different models, one for the shape of human faces and one for the human femur bone.

In the case of the femur bone, we are interested in determining how well the distal (bottom) part of the femur determines the proximal (top) part, in order to estimate how reliable the model can be for reconstruction of missing or injured parts of the bone. In Figure 1, the distal part of the bone is colored. It corresponds to the fixed points $\mathbf{x}_b$. From the generalized eigenvalue problem Equation (18), the flexibility components $\boldsymbol{\alpha}_i$ are calculated. The regularization parameter was chosen as $\eta = 10$. Figure 1 illustrates the effect of deforming the mean by $\pm 2.2\sqrt{l}$ times the first two flexibility components α_1 and α_2. Clearly, both the distal and the proximal part are changed, but the proximal part is much more heavily deformed. In Figure 2, the amount of deformation is plotted, measured in the mean and the maximum of the Euclidean norm of the deformation at each point. For the mean norm, the ratio is approximately equal to

the corresponding eigenvalue. For the mean squared norm, which is not plotted here, it is of course exactly equal to this eigenvalue.

In case of the face model we are interested in the question: Given the vertices on the contour, how much is the shape of the face determined? Here, the fixed vertices $\mathbf{x}_b$ are the vertices of the occluding contour in profile view, as shown in Figure (3 a)). This is the contour for one individual represented by its model coefficients $\boldsymbol{\beta}$. Again, the flexibility components $\boldsymbol{\alpha}_i$ are computed using Equation (18), with $\eta = 10$. To visualize the result, the surface with the coefficients $\boldsymbol{\beta} \pm 1.1\sqrt{l}\boldsymbol{\alpha}_i$ is shown in Figure 4. We see that the first two flexibility components heavily deform the model, while the vertices at the contour are almost fixed, see Figure 3 b,c). The resulting deformations measures are plotted in Figure 2. The ratio is much higher than in the femur case as less points are kept fixed and the model is built from many more examples, making it more expressive.

Conclusion. We have introduced a way to model the remaining flexibility of a statistical shape model when a part of the model is kept as fixed as possible, even in the absence of a nontrivial conditional distribution. In future work, we will apply this technique in the fitting of statistical models to partial data.

References

1. Fleute, M., Lavallee, S.: Building a complete surface model from sparse data using statistical shape models: Application to computer assisted knee surgery. Medical Image Computing and Computer-Assisted Intervention-MICCAI 98 880–7
2. Zheng, G., Rajamani, K., Nolte, L.: Use of a dense surface point distribution model in a three-stage anatomical shape reconstruction from sparse information for computer-assisted orthopaedic surgery: a preliminary study. ACCV 6 52–60
3. Subsol, G., Thirion, J., Ayache, N.: A scheme for automatically building three-dimensional morphometric anatomical atlases: application to a skull atlas. Medical Image Analysis 2(1) (1998) 37–60
4. Dam, E., Fletcher, P., Pizer, S.: Automatic shape model building based on principal geodesic analysis bootstrapping. Medical Image Analysis (2007)
5. Rao, A., Babalola, K., Rueckert, D.: Canonical Correlation Analysis of Sub-cortical Brain Structures Using Non-rigid Registration. Lecture Notes in Computer Science 4057 (2006) 66
6. Shen, D., Herskovits, E., Davatzikos, C.: An adaptive-focus statistical shape model for segmentation and shape modeling of 3-D brain structures. IEEE Transactions on Medical Imaging 20(4) (2001) 257–270
7. Blanz, V., Vetter, T.: Reconstructing the complete 3d shape of faces from partial information. Informationstechnik und Technische Informatik 44(6) (2002) 1–8
8. Keller, M., Knothe, R., Vetter, T.: 3D Reconstruction of Human Faces from Occluding Contours. Lecture Notes in Computer Science 4418 (2007) 261
9. Dedner, A., Lüthi, M., Albrecht, T., Vetter, T.: Curvature guided level set registration using adaptive finite elements. In: Pattern Recognition. (2007) 527–536
10. Amberg, B., Romdhani, S., Vetter, T.: Optimal Step Nonrigid ICP Algorithms for Surface Registration. Computer Vision and Pattern Recognition, 2007. CVPR'07. IEEE Conference on (2007) 1–8
11. Bishop, C.: Pattern recognition and machine learning. Springer (2006)