

A Statistical Method for Robust 3D Surface Reconstruction from Sparse Data

V. Blanz,¹ A. Mehl,² T. Vetter³ and H.-P. Seidel¹

¹ Max-Planck-Institut für Informatik, Saarbrücken, Germany

² Department of Restorative Dentistry, Ludwig Maximilians University, München, Germany

³ University of Basel, Departement Informatik, Basel, Switzerland

Abstract—

General information about a class of objects, such as human faces or teeth, can help to solve the otherwise ill-posed problem of reconstructing a complete surface from sparse 3D feature points or 2D projections of points. We present a technique that uses a vector space representation of shape (3D Morphable Model) to infer missing vertex coordinates. Regularization derived from a statistical approach makes the system stable and robust with respect to noise by computing the optimal tradeoff between fitting quality and plausibility. We present a direct, non-iterative algorithm to calculate this optimum efficiently, and a method for simultaneously compensating unknown rigid transformations. The system is applied and evaluated in two different fields: (1) reconstruction of 3D faces at unknown orientations from 2D feature points at interactive rates, and (2) restoration of missing surface regions of teeth for CAD-CAM production of dental inlays and other medical applications.

I. INTRODUCTION

Dentists and dental technicians are highly-trained experts in restoring parts of teeth that are missing, for example due to caries. They design fillings and inlays that fit the boundaries and the overall shape of the remaining tooth, but also contain characteristic morphological patterns, such as cusps and grooves [1]. Human experience with shapes of teeth makes this otherwise underconstrained problem tractable, and guarantees morphologically adequate results. We present an automated system for 3D surface reconstruction that adopts this knowledge-based strategy.

A related problem from a different field is to reconstruct 3D shape of objects, such as human faces, from 2D image coordinates of a few feature points. This problem is underconstrained not only because we have to interpolate between sparse feature points, but also due to the 2D projection of 3D points into the image plane. Our algorithm assumes this projection to be orthographic, which is approximately the case for images of faces taken from distances of 2 meters or more. Pose, location and scale of the face may be arbitrary and unknown.




Fig. 1. Given the remaining material of a tooth and 10 hand-labeled feature points (left), our algorithm reconstructs missing surface regions. The reconstruction (right) does not only fit to the original material, but also contains typical features of dental morphology, which are represented statistically in a Morphable Model of 3D teeth.

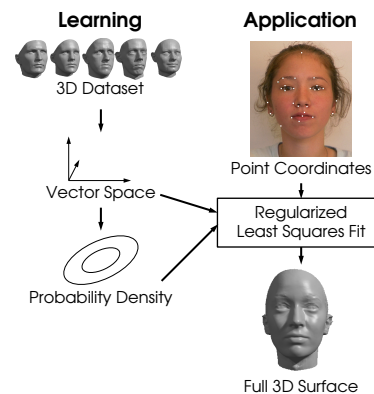


Fig. 2. Based on a vector space representation and an estimate of the probability density of 3D faces, the system computes a complete 3D surface from a sparse set of 2D or 3D point coordinates.

Our learning-based approach relies on a dataset of 3D scans, which are converted into a vector space representation (Morphable Model). Restricting possible solutions to the span of examples, we can infer missing information about additional points, using correlation between vector components which is implicitly stored in the Morphable Model. We show that additional regularization is necessary for robust and stable results, and use a statistically motivated regularization based on the probability density of the examples. The regularization finds a tradeoff between fitting the surface to the feature points, and producing a plausible solution in terms of prior probability. This optimum is computed in a direct, non-iterative way at interactive rates using a Singular Value Decomposition (SVD). By simultaneously estimating 3D shape and rigid alignment, our algorithm is robust with respect to unknown orientation.

The general technique can be used for all problems where sparse measurements are given and the full solution is in the linear span of a set of examples. The measurement can be any linear mapping from the full vector, such as a product of the following operators: a projection that selects a subset of vector coordinates (sparse features or surface regions), a 3D rotation, and a projection into the image plane. After a theoretical derivation of the algorithm, we present a compensation scheme for unknown rigid alignment, and introduce constraints that restrict points only in one direction while leaving them unconstrained along planes or lines, such as the eyebrows. Individual weights can be assigned to each feature point, and an iterative refinement paradigm captures details of the input data. We present an interactive program for model-based surface reconstruction, give an experimental evaluation that demonstrates the effect of regularization, and present a full system for restoration of teeth in dentistry.

II. PREVIOUS WORK

Completion of missing data with Eigenspace-approaches has been used for reconstructing an image of a face from a few pixels [2], and for reconstructing 2D shape and color from a set of 2D point coordinates [3], [4]. However, these methods applied a least-squares fit without regularization, which makes the solution very unstable. Our statistical approach solves this problem by computing an optimal tradeoff both in overconstrained and underconstrained situations and in the presence of noise. Moreover, these entirely image-based methods are restricted to frontal views, while our 3D approach and the automated 3D alignment scheme works for any unknown orientation. Filling-in missing information was also achieved as a side-effect in iterative reconstruction methods [5] and [6].

Regularization has been applied to surface reconstruction from sparse, contradictive or noisy measurements in computer vision in a technique that defines a unique solution for these otherwise ill-posed problems by minimizing an additional smoothness functional [7]. In contrast, we use a class-specific prior probability learned from examples [6], [8] for regularization in a statistically motivated approach based on a maximum posterior probability estimator [9], [10]. We present a direct solution without costly iterations, and without marginalization.

3D alignment of rigid shapes is an important problem in many fields, such as 3D shape acquisition, medical image registration and computer vision, and a number of solutions have been proposed (for an overview, see [11]). Among the non-iterative methods, some involve an SVD of the 3×3 matrix $\mathbf{H} = \frac{1}{p} \sum_{i=1}^p (\mathbf{q}_i - \bar{\mathbf{q}})(\mathbf{q}'_i - \bar{\mathbf{q}}')^T$ between two point sets \mathbf{q}_i and \mathbf{q}'_i , and their respective averages $\bar{\mathbf{q}}$, $\bar{\mathbf{q}}'$ [12], [13], [14]. Others use a linearization of product terms to reproduce the pairwise point distances of feature points [15]. Alignment of rigid surfaces is also achieved by iterative closest point algorithms [16], [17]. We adopt the strategy of iteratively updating point correspondences in Section X-A. A first-order Taylor expansion of the projected 2D coordinates and subsequent SVD is described in [14]. In contrast, we expand the rotation operator in 3D. This allows us to combine rigid alignment with shape estimation and regularization for fitting unknown, deformable surfaces.

From a variety of methods to reconstruct smooth surfaces from point-cloud data, we mention three related approaches: Fitting algebraic surfaces [18] uses a parameterized class of surfaces, such as spheres, but cannot handle more natural shapes. Implicit surfaces with radial basis functions [19] can describe any shape, achieving smooth results with a regularization functional based on second derivatives. Other methods use local polynomial approximations [20]. All of these methods can only interpolate between given points without inferring missing structures from class-specific information.

Deformable 3D models have been fitted to 2D and 3D data before. Reconstruction from 2D feature points has been achieved with Levenberg-Marquardt optimization [21], [22], and tree search [23]. For separate estimation of pose and shape, linearized approximations of 3D rotation similar to ours have been used [22]. An alternative solution that approximates 3D rotation by 2D affine transformations and pre-learned parallax has been proposed in an algorithm based on SVD, but without regularization [24]. A number of methods have analyzed frame-to-frame motion in video to recover pose and

non-rigid deformation for tracking of faces, using a linearized system based on a matrix factorization [25], or combinations of optical flow with constraints derived from deformable models [26], [27], [28]. Unlike these methods, our approach relies on a single image for estimating depth.

Reconstruction of 3D shape and texture from color values is described in [6], [8]. However, the iterative optimization in this analysis-by-synthesis system is computationally expensive. A similar method has also been used to fit the Morphable Face Model to 3D range scans in a cylindrical parameterization $r(h, \phi)$ [6]. The method that we present here is related to this approach, but uses only feature point coordinates. This simplified setting allows us to formulate a faster and mathematically more elegant solution that involves nothing more than an SVD of a matrix of moderate size.

III. CLASS-SPECIFIC SHAPE PROPERTIES

The characteristic shape properties of objects classes are derived from a dataset of 3D scans. Having defined a reference scan or mesh with p vertices, dense point-to-point correspondence with all other scans is established with an automated technique [6], so corresponding points such as the tip of the nose are identified on all examples. For each scan $i \in \{1, \dots, m\}$, the coordinates of these p corresponding surface points are concatenated to a vector

$$\mathbf{v}_i = (x_1, y_1, z_1, \dots, x_p, y_p, z_p)^T \in \mathbb{R}^n, \quad n = 3 \cdot p. \quad (1)$$

In this representation, any convex combination

$$\mathbf{v} = \sum_{i=1}^m a_i \mathbf{v}_i, \quad a_i \in [0, 1], \quad \sum_{i=1}^m a_i = 1 \quad (2)$$

describes a new element of the class. In order to remove the second constraint, we use barycentric coordinates relative to the arithmetic mean:

$$\mathbf{x} = \mathbf{v} - \bar{\mathbf{v}}, \quad \bar{\mathbf{v}} = \frac{1}{m} \sum_{i=1}^m \mathbf{v}_i, \quad (3)$$

so (2) becomes

$$\mathbf{x} = \sum_{i=1}^m b_i \mathbf{x}_i, \quad \mathbf{x}_i = \mathbf{v}_i - \bar{\mathbf{v}}. \quad (4)$$

Note that the set of \mathbf{x}_i is linearly dependent: $\sum_{i=1}^m \mathbf{x}_i = 0$. The constraint $a_i \in [0, 1]$ would define a sharp boundary of the object class. We prefer to describe the class in terms of a probability density $p(\mathbf{v})$ of \mathbf{v} being in the object class. $p(\mathbf{v})$ can be estimated by a Principal Component Analysis (PCA, see [29]): Let the data matrix \mathbf{X} be

$$\mathbf{X} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_m) \in \mathbb{R}^{n \times m}. \quad (5)$$

The covariance matrix of the data set is given by

$$\mathbf{C} = \frac{1}{m} \mathbf{X} \mathbf{X}^T = \frac{1}{m} \sum_{j=1}^m \mathbf{x}_j \mathbf{x}_j^T \in \mathbb{R}^{n \times n}, \quad (6)$$

PCA is based on a diagonalization $\mathbf{C} = \mathbf{S} \cdot \text{diag}(\sigma_i^2) \cdot \mathbf{S}^T$. Since \mathbf{C} is symmetrical, the columns \mathbf{s}_i of \mathbf{S} form an orthogonal set of eigenvectors. σ_i are the standard deviations within the data along the eigenvectors. The diagonalization can be calculated by a Singular Value Decomposition (SVD, [30]) of \mathbf{X} .

The space spanned by \mathbf{x}_i is at most $m' = (m - 1)$ dimensional, and the rank of \mathbf{X} and \mathbf{C} is at most m' , so $\sigma_m = 0$, and \mathbf{s}_m is not meaningful. If we use the scaled eigenvectors $\sigma_i \mathbf{s}_i$ as a basis, vectors \mathbf{x} are defined by coefficients c_i :

$$\mathbf{x} = \sum_{i=1}^{m'} c_i \sigma_i \mathbf{s}_i = \mathbf{S} \cdot \text{diag}(\sigma_i) \mathbf{c} \quad (7)$$

The variance of the original data is 1 in each coefficient c_i , so the estimated normal distribution takes the simple form

$$p(\mathbf{c}) = \nu_c \cdot e^{-\frac{1}{2}\|\mathbf{c}\|^2}, \quad \nu_c = (2\pi)^{-m'/2}. \quad (8)$$

$\|\mathbf{c}\|^2$ is the distance from the average, scaled according to the normal distribution (Mahalanobis Distance).

IV. MODEL-BASED SURFACE COMPLETION

Given the positions of a reduced number $f < p$ of feature points, our task is to find the 3D coordinates of all other vertices. Similar to Equation (1), we write the 2D or 3D coordinates of the feature points as vectors $\mathbf{r} \in \mathbb{R}^l$ ($l = 2f$ or $l = 3f$), and assume that \mathbf{r} is related to \mathbf{v} by

$$\mathbf{r} = \mathbf{L}\mathbf{v} \quad \mathbf{L} : \mathbb{R}^n \mapsto \mathbb{R}^l. \quad (9)$$

\mathbf{L} may be any linear mapping, such as a product of a projection that selects a subset of components from \mathbf{v} for sparse feature points or remaining surface regions, a rigid transformation in 3D, and an orthographic projection to image coordinates. Until Section VI, we assume that \mathbf{L} is known precisely. Let

$$\mathbf{y} = \mathbf{r} - \mathbf{L}\bar{\mathbf{v}} = \mathbf{L}\mathbf{x} \quad (10)$$

If \mathbf{L} is not one-to-one, the solution \mathbf{x} of (10) will not be uniquely defined. To reduce the number of free parameters, we restrict \mathbf{x} to the linear combinations of \mathbf{x}_i , which makes sure that \mathbf{x} is in the object class. Depending on the number l of feature coordinates and the dimension of the span of examples, the solution of (10) is still not unique. We address this issue at the end of this section.

As we cannot expect to find a linear combination of the examples that solves (10) exactly, we minimize

$$E(\mathbf{x}) = \|\mathbf{L}\mathbf{x} - \mathbf{y}\|^2. \quad (11)$$

Let $\mathbf{q}_i = \mathbf{L}(\sigma_i \mathbf{s}_i) \in \mathbb{R}^l$ be the reduced versions of the scaled eigenvectors, and

$$\mathbf{Q} = (\mathbf{q}_1, \mathbf{q}_2, \dots) = \mathbf{L}\mathbf{S} \cdot \text{diag}(\sigma_i) \in \mathbb{R}^{l \times m'}. \quad (12)$$

In terms of model coefficients c_i from (7), (11) is

$$E(\mathbf{c}) = \|\mathbf{L} \sum_i c_i \sigma_i \mathbf{s}_i - \mathbf{y}\|^2 = \|\mathbf{Q}\mathbf{c} - \mathbf{y}\|^2. \quad (13)$$

The optimum can be found by a Singular Value Decomposition [30]

$$\mathbf{Q} = \mathbf{U}\mathbf{W}\mathbf{V}^T \quad (14)$$

with a diagonal matrix $\mathbf{W} = \text{diag}(w_i)$, and $\mathbf{V}^T \mathbf{V} = \mathbf{V}\mathbf{V}^T = \text{Id}_{m'}$. The pseudo-inverse (see [30]) of \mathbf{Q} is

$$\mathbf{Q}^+ = \mathbf{V}\mathbf{W}^+ \mathbf{U}^T, \quad (15)$$

$$\mathbf{W}^+ = \text{diag} \left(\begin{array}{cc} w_i^{-1} & \text{if } w_i \neq 0 \\ 0 & \text{otherwise} \end{array} \right). \quad (16)$$

To avoid numerical problems, the condition $w_i \neq 0$ may be replaced by a threshold $w_i > \epsilon$. The minimum of $E(\mathbf{c})$ can be computed with the pseudo-inverse: $\mathbf{c} = \mathbf{Q}^+ \mathbf{y}$.

This vector \mathbf{c} has another important property [30]: If the minimum of $E(\mathbf{c})$ is not uniquely defined, \mathbf{c} is the vector with minimum norm $\|\mathbf{c}\|$ among all \mathbf{c}' with $E(\mathbf{c}') = E(\mathbf{c})$. In our case, this means that we obtain the vector with maximum prior probability (Equation 8).

By Equation (3) and (7), \mathbf{c} is mapped to \mathbb{R}^n :

$$\mathbf{v} = \mathbf{S} \cdot \text{diag}(\sigma_i) \mathbf{c} + \bar{\mathbf{v}}. \quad (17)$$

For solving Equation (10), it might seem more straightforward to compute $\mathbf{x} = \mathbf{L}^+ \mathbf{y}$ with the pseudo-inverse \mathbf{L}^+ of \mathbf{L} . However, the result would, in general, neither be in the span of examples, nor would it minimize $\|\mathbf{c}\|$, so none of the model information would be considered.

V. STATISTICAL SHAPE RECONSTRUCTION

The previous solution will ensure that E is minimized, and in particular that $E = 0$ whenever this is possible. Prior probability is only considered among solutions of equal $E(\mathbf{c})$.

However, it is well known in statistical learning that optimal generalization (in our case to the unknown surface regions) in terms of the expected error is not always achieved by reproducing the given data precisely [31]. In regression, for example, a smooth curve that passes the measured data points at some distance may lead to better generalization than a complicated curve that passes through all measurements and shows overfitting effects elsewhere. Regularization provides a powerful tool to resolve this problem. In the following, we derive a regularization framework from statistical assumptions and a maximum a posteriori approach.

There are a number of reasons why measurements \mathbf{y} should not be fully reproduced by the method described in the previous chapter: \mathbf{y} may be subject to noise in detecting the feature points, or wrong assumptions on the rotation involved in \mathbf{L} . Moreover, the linear span of the examples will, in general, only approximate the novel shapes \mathbf{x} encountered in the application of the system.

Minimizing $E(\mathbf{x}) = \|\mathbf{L}\mathbf{x} - \mathbf{y}\|^2$ may therefore lead to solutions that are too far from the average and heavily distorted, as demonstrated experimentally in Section IX-A. To avoid this overfitting, we propose a tradeoff between matching quality and prior probability of the solution. This tradeoff will be derived from a Bayesian approach in the following section.

A. Maximum Posterior Probability

Given the observed vector \mathbf{y} , we are looking for the coefficients \mathbf{c} with maximum posterior probability $P(\mathbf{c}|\mathbf{y})$. As an intermediate step, consider the likelihood of measuring \mathbf{y} , given \mathbf{c} : In the noiseless case, \mathbf{c} would define the vector

$$\mathbf{y}_{\text{model}} = \mathbf{L} \sum_i c_i \sigma_i \mathbf{s}_i = \sum_i c_i \mathbf{q}_i = \mathbf{Q}\mathbf{c} \quad (18)$$

We assume that each dimension j of the measured vector \mathbf{y} is subject to uncorrelated Gaussian noise with a variance σ_N^2 . Then, the likelihood of measuring $\mathbf{y} \in \mathbb{R}^l$ is given by

$$P(\mathbf{y}|\mathbf{y}_{\text{model}}) = \prod_{j=1}^l P(y_j|y_{\text{model},j}) \quad (19)$$

$$= \prod_{j=1}^l \nu_N \cdot e^{-\frac{1}{2\sigma_N^2}(y_{\text{model},j} - y_j)^2} = \nu_N^l \cdot e^{-\frac{1}{2\sigma_N^2}\|\mathbf{y}_{\text{model}} - \mathbf{y}\|^2} \quad (20)$$

with a normalization factor ν_N . In terms of the model parameters \mathbf{c} , the likelihood is

$$P(\mathbf{y}|\mathbf{c}) = \nu_N^l \cdot e^{-\frac{1}{2\sigma_N^2}\|\mathbf{Q}\mathbf{c} - \mathbf{y}\|^2}. \quad (21)$$

According to Bayes Rule [32], the posterior probability is

$$P(\mathbf{c}|\mathbf{y}) = \nu \cdot P(\mathbf{y}|\mathbf{c}) \cdot p(\mathbf{c}). \quad (22)$$

with a constant factor $\nu = (\int P(\mathbf{y}|\mathbf{c}') \cdot p(\mathbf{c}') d\mathbf{c}')^{-1}$.

Substituting (8) and (21) yields

$$P(\mathbf{c}|\mathbf{y}) = \nu \cdot \nu_N^l \cdot \nu_c \cdot e^{-\frac{1}{2\sigma_N^2}\|\mathbf{Q}\mathbf{c} - \mathbf{y}\|^2} \cdot e^{-\frac{1}{2}\|\mathbf{c}\|^2}, \quad (23)$$

which is maximized by minimizing the cost function

$$E = -2 \cdot \log P(\mathbf{c}|\mathbf{y}) = \frac{1}{\sigma_N^2} \|\mathbf{Q}\mathbf{c} - \mathbf{y}\|^2 + \|\mathbf{c}\|^2 + \text{const.} \quad (24)$$

B. Combined Cost Function

In this section, we show that the cost function (24) can be minimized in a single step. To simplify the calculation, we introduce a weight factor $\eta = \sigma_N^2 \geq 0$ and minimize

$$E = \|\mathbf{Q}\mathbf{c} - \mathbf{y}\|^2 + \eta \cdot \|\mathbf{c}\|^2. \quad (25)$$

In the optimum,

$$0 = \nabla E = 2\mathbf{Q}^T\mathbf{Q}\mathbf{c} - 2\mathbf{Q}^T\mathbf{y} + 2\eta\mathbf{c}, \quad (26)$$

so

$$\mathbf{Q}^T\mathbf{Q}\mathbf{c} + \eta\mathbf{c} = \mathbf{Q}^T\mathbf{y}. \quad (27)$$

Singular Value Decomposition $\mathbf{Q} = \mathbf{U}\mathbf{W}\mathbf{V}^T$ yields

$$\mathbf{Q}^T\mathbf{Q} = \mathbf{V}\mathbf{W}\mathbf{U}^T\mathbf{U}\mathbf{W}\mathbf{V}^T = \mathbf{V}\mathbf{W}^2\mathbf{V}^T. \quad (28)$$

since \mathbf{U} is orthogonal in all columns i with $w_i \neq 0$. From (27),

$$\mathbf{V}\mathbf{W}^2\mathbf{V}^T\mathbf{c} + \eta\mathbf{c} = \mathbf{V}\mathbf{W}\mathbf{U}^T\mathbf{y}. \quad (29)$$

Multiplying by \mathbf{V}^T , this can be solved for \mathbf{c} :

$$\text{diag}(w_i^2 + \eta) \cdot \mathbf{V}^T\mathbf{c} = \mathbf{W}\mathbf{U}^T\mathbf{y} \quad (30)$$

$$\mathbf{c} = \mathbf{V}\text{diag}\left(\frac{w_i}{w_i^2 + \eta}\right)\mathbf{U}^T\mathbf{y} \quad (31)$$

The overall result is obtained from (3) and (7):

$$\mathbf{v} = \bar{\mathbf{v}} + \mathbf{S}\text{diag}(\sigma_i)\mathbf{V}\text{diag}\left(\frac{w_i}{w_i^2 + \eta}\right)\mathbf{U}^T\mathbf{y}. \quad (32)$$

C. Special case $\mathbf{L} = id_n$

To find the closest element of the span of examples from a vector \mathbf{x} that is entirely known, or to approximate a given element of the span by a more plausible solution, we can set $\mathbf{L} = id_n$. The Singular Value Decomposition of \mathbf{Q} is then trivial, and Equation (32) reduces to

$$\mathbf{v} = \bar{\mathbf{v}} + \sum_i \frac{1}{1 + \frac{\eta}{\sigma_i^2}} \langle \mathbf{s}_i, \mathbf{y} \rangle \mathbf{s}_i. \quad (33)$$

The most relevant dimensions s_i with large standard deviation σ_i are affected less by η than those with small σ_i .

VI. TRANSLATION, SCALE AND ROTATION

Since real-world data are, in general, not aligned with the model, it is essential to compensate for unknown rigid transformations. 3D feature points could be pre-aligned with the corresponding points of the average surface using 3D3D Absolute Orientation [12], [13], [14]. The result would be sub-optimal, however, since the optimal alignment depends on the shape *after* reconstruction. In this section, we present a method to compensate for transformations and deformations simultaneously. In the case of 2D image coordinates, pose estimation of an unknown shape is even more difficult than in the 3D setting. Our solution makes the algorithm applicable to pictures of faces at arbitrary, unknown orientations.

The method includes additional coefficients c_i and vectors \mathbf{s}_i in the linear combination (7). *Shifting* is achieved simply by including vectors

$$\mathbf{s}_{tx} = (1, 0, 0, 1, 0, 0, \dots, 1, 0, 0)^T \quad (34)$$

and \mathbf{s}_{ty} , \mathbf{s}_{tz} in the set of \mathbf{s}_i . *Scaling* is allowed for by a vector $\mathbf{s}_s = \bar{\mathbf{v}}$: Let

$$\mathbf{v} = \sum_{i=1}^{m'} c_i \sigma_i \mathbf{s}_i + \bar{\mathbf{v}}, \quad \text{then} \quad (35)$$

$$\lambda \mathbf{v} = \sum_{i=1}^{m'} (\lambda c_i) \sigma_i \mathbf{s}_i + (1 - \lambda) \mathbf{s}_s + \bar{\mathbf{v}}, \quad (36)$$

so the scaled shape is represented in the model by adding $(1 - \lambda)\mathbf{s}_s$, and scaled coefficients λc_i . For λ far from 1, this may affect the regularization tradeoff in the cost function

(25), leading to more conservative estimates at a large scale than at a small scale because the same shape becomes more “expensive”. For $\lambda \approx 1$, this effect is not noticeable.

Rotation is compensated in a two-pass approach: Based on approximations for small angles $\gamma, \theta, \phi \ll 1$, the method estimates rotation and 3D shape. This estimate, which may be unprecise in case of larger rotations, is included in \mathbf{L} in the second pass, and the residual small rotation along with a refined 3D shape are computed.

For $\gamma, \theta, \phi \ll 1$, we can set the cosines to 1, and ignore products of sinus terms. For a product of 3 rotations around the z, x and y-axis, we obtain

$$\mathbf{R} = \mathbf{R}_\gamma \mathbf{R}_\theta \mathbf{R}_\phi \approx \begin{pmatrix} 1 & -\sin \gamma & \sin \phi \\ \sin \gamma & 1 & -\sin \theta \\ -\sin \phi & \sin \theta & 1 \end{pmatrix} \quad (37)$$

$$\mathbf{R} \begin{pmatrix} x \\ y \\ z \end{pmatrix} \approx \sin \gamma \begin{pmatrix} -y \\ x \\ 0 \end{pmatrix} + \sin \theta \begin{pmatrix} 0 \\ -z \\ y \end{pmatrix} + \sin \phi \begin{pmatrix} z \\ 0 \\ -x \end{pmatrix} + \begin{pmatrix} x \\ y \\ z \end{pmatrix}.$$

We combine the 3×3 matrices \mathbf{R} to a mapping \mathbf{R}_v of the full vectors $\mathbf{v} \in \mathbb{R}^n$ and perform a first-order Taylor expansion of $\mathbf{R}_v \mathbf{v}$, ignoring effects of rotation on the principal components, to obtain a linear combination

$$\mathbf{R}_v \mathbf{v} \approx \sum_{i=1}^{m'} c_i \sigma_i \mathbf{s}_i + c_\gamma \mathbf{s}_\gamma + c_\theta \mathbf{s}_\theta + c_\phi \mathbf{s}_\phi + \bar{\mathbf{v}} \quad (38)$$

$$\mathbf{s}_\gamma = (-\bar{y}_1, \bar{x}_1, 0, -\bar{y}_2, \bar{x}_2, 0, \dots)^T \quad (39)$$

$$\mathbf{s}_\theta = (0, -\bar{z}_1, \bar{y}_1, 0, -\bar{z}_2, \bar{y}_2, \dots)^T \quad (40)$$

$$\mathbf{s}_\phi = (\bar{z}_1, 0, -\bar{x}_1, \bar{z}_2, 0, -\bar{x}_2, \dots)^T. \quad (41)$$

The vectors \mathbf{s}_γ , \mathbf{s}_θ , \mathbf{s}_ϕ , \mathbf{s}_{tx} , \mathbf{s}_{ty} , \mathbf{s}_{tz} , \mathbf{s}_s are normalized to $\|\mathbf{s}\| = 1$ and treated just as the principal components \mathbf{s}_i in the optimization (Section V). To control their weights in regularization, their standard deviations should be set equal to the most dominant principal component σ_1 .

After the optimization, the angles are recovered from $\gamma = \arcsin c_\gamma$, $\theta = \arcsin c_\theta$, $\phi = \arcsin c_\phi$, and included in \mathbf{L} for a second pass. Due to the approximations, the estimate from the first pass may be wrong by several degrees if the initial orientation was far from the true orientation (Figure 4). The second pass, however, yields the final, stable estimate already, which justifies this simple and efficient method.

VII. DIRECTIONAL CONSTRAINTS AND WEIGHTS

Many facial features, such as the eyebrows and lips, are in fact lines or curves rather than points. The constraint in the fitting process should let these features move freely along the tangent line, and restrict them only in the orthogonal direction. With the method described so far, the user would face the ill-defined problem of selecting corresponding points along the curve, and the reconstruction would be restricted by the user’s arbitrary selection (Figure 3). In 3D fitting (e.g. reconstructing teeth from 3D features), we may want to restrict points to planes, for example the approximal contacts to adjacent teeth, but leave them unconstrained otherwise.

Consider the individual feature points in \mathbf{y} and \mathbf{y}_{model} : Let \mathbf{q} be the observed 2D or 3D coordinates of one point, and \mathbf{q}_{model} the coordinates predicted by the model. Then the distance $d = \|\mathbf{q} - \mathbf{q}_{model}\|^2$ which has contributed to (25) is replaced by




Fig. 3. Directional constraints leave lines such as eyebrows or lips unrestricted along the tangent. Clicking on feature points, the user cannot define correspondence along the lines properly (top row). Point-based constraints would distort the face (bottom, left), while directional constraints yield realistic results (center) by allowing for mismatches along the tangent (bottom right: blended images).

$$d_n = \langle \mathbf{n}, (\mathbf{q} - \mathbf{q}_{model}) \rangle^2 \quad (42)$$

$$= \|\mathbf{n}^T (\mathbf{q} - \mathbf{q}_{model})\|^2 = \|\mathbf{A}\mathbf{q} - \mathbf{A}\mathbf{q}_{model}\|^2 \quad (43)$$

with a normal vector \mathbf{n} of the contour line or boundary plane, $\|\mathbf{n}\| = 1$, and a matrix $\mathbf{A} = (n_x, n_y)$ or $\mathbf{A} = (n_x, n_y, n_z)$. The constraints used in previous sections, which restricted points in all directions, can be written in (43) with $\mathbf{A} = id_2$ or id_3 . Combining the matrices \mathbf{A} of all feature points to a block-diagonal matrix \mathbf{A}_y on vectors \mathbf{y} , (25) becomes

$$E = \|\mathbf{A}_y \mathbf{Q} \mathbf{c} - \mathbf{A}_y \mathbf{y}\|^2 + \eta \cdot \|\mathbf{c}\|^2. \quad (44)$$

The optimization problem can be solved in the same way as in Section V-B if we compute the Singular Value Decomposition of the product $\mathbf{A}_y \mathbf{Q} = \mathbf{U} \mathbf{W} \mathbf{V}^T$. The result is

$$\mathbf{v} = \mathbf{S} \mathbf{diag}(\sigma_i) \mathbf{V} \mathbf{diag}\left(\frac{w_i}{w_i^2 + \eta}\right) \mathbf{U}^T \mathbf{A}_y \mathbf{y} + \bar{\mathbf{v}}. \quad (45)$$

Individual weights for the dimensions of \mathbf{y} can be implemented in the same framework with a diagonal matrix \mathbf{A}_y formed by the weight factors. Weights can reflect the importance assigned to different feature points, the precision of fitting these points, or the noise σ_N of each measurement.

VIII. RECONSTRUCTION OF 3D FACES FROM IMAGES

From a small number of 2D positions of feature points, the algorithm can recover 3D shape of human faces at high resolution, inferring both depth and the missing vertex coordinates. The system is based on a Morphable Model [6] that has been built from laser scans of 200 faces, using a modified optical flow algorithm to compute dense point-to-point correspondence. Each face is represented by the coordinates of $p = 75972$ vertices at a spacing of less than 1mm. We use only the 140 most relevant principal components.

For shape reconstruction, the user clicks on feature points in the image and the corresponding points on the 3D reference model. Good results are achieved with 15 to 20 points. Due to the automated 3D alignment (Section VI), no estimate of pose, position and size is required. Computation time is 250ms on a 1.7GHz Intel Xeon processor for forming \mathbf{Q} from the large matrix \mathbf{S} , SVD of \mathbf{Q} in two passes for pose estimation, and computation of the full face (Equation 32).




Fig. 4. From an original image at unknown pose (top, left) and a frontal starting position (top, right), the algorithm estimates 3D shape and pose from 17 feature coordinates, including 7 directional constraints (second row). We used 140 principal components and 7 vectors for transformations (Section VI). The third row shows the texture-mapped result. Computation time is 250ms.



Fig. 5. Reconstruction and model-based deformation.

Figure 4 shows an example of 3D shape reconstruction from 17 feature points. 7 of the points are directional constraints, as indicated by tangent lines in the figure. The system successfully compensated for rotation, scaling and translation. The color values of the image are mapped as a texture on the surface (bottom row), and missing color values are reflected from visible parts or filled in with the average texture of the morphable model. Better texture values for filling in could be computed with the same method that we used for shape. Figure 5 shows another example of a reconstruction, and illustrates model-based image manipulation: After texture-mapping, the marker points can be moved, producing a novel shape from within the object class, which is rendered back into the image. Note that changes on one side, e.g. the width of the nose or the eyes, are transferred symmetrically.

IX. EXPERIMENTAL EVALUATION

To evaluate the predicted properties of regularization on the reconstruction setting of Section VIII, we split the database randomly into a training set and a test set of $m = 100$ faces each. The training set provides the examples \mathbf{v}_i available to




Fig. 6. The effect of η on average reconstruction results for 100 novel faces, given 50 feature points, and using 99 principal components. As η increases, the feature points are matched less precisely, so E_y grows. In contrast, $\|c\|$ decreases, as the results become more plausible. The overall 3D shape error E_{full} is lowest for a tradeoff between both criteria.

the system for PCA ($m' = 99$). The test set is used to assess performance on novel faces.

From the vertices of the full model, we selected sets of $f = 17, 50, \text{ or } 1000$ vertices as feature points. The vertices in the two smaller sets are salient points such as the corners of the mouth, those in the large set are selected randomly. Their image plane coordinates, computed by orthographic projection in frontal orientation, form synthetic test vectors $\mathbf{r} = \mathbf{L}\mathbf{v} \in \mathbb{R}^l$, $l = 2 \cdot f$ (Equation 9). The mapping \mathbf{L} is assumed to be known. Focussed on the properties of our regularization method, this evaluation does not include compensation for transformations (Section VI).

The evaluation is based on the following quantities, which are averaged across all 100 training or test faces:

- $E_y = \|\mathbf{Q}\mathbf{c} - \mathbf{y}\|$, the image plane matching error for all feature points (in units of pixels in a 300x300 image).
- $\|c\|$, the Mahalanobis distance from the average.
- The per-vertex average of distances in 3D space between




Fig. 7. Given the image coordinates of 50 feature points of a novel face (top left), 3D shape was reconstructed with 99 principal components. The result depends on a tradeoff between the precision of feature point matching, and prior probability. This tradeoff is controlled by the parameter η .




Fig. 8. The average shape reconstruction errors for 100 training faces depend on the level of noise σ_N added to each feature point coordinate. While noise-free training data are best analyzed with $\eta = 0$ (solid line), reconstruction quality for noisy data is best at $\eta = \sigma_N^2$.

reconstruction and original, computed over the entire set of vertices (except for neck and top of forehead):

$$E_{full} = \frac{1}{p} \sum_{i=1}^p \left\| \begin{pmatrix} x_{i, \text{reconst.}} \\ y_{i, \text{reconst.}} \\ z_{i, \text{reconst.}} \end{pmatrix} - \begin{pmatrix} x_{i, \text{orig.}} \\ y_{i, \text{orig.}} \\ z_{i, \text{orig.}} \end{pmatrix} \right\| \quad (46)$$

A. Reconstruction of Novel Faces

The effect of overfitting is clearly demonstrated in Figures 6 and 7: For small values of η , the face is heavily distorted, and E_{full} is large. Still, the feature point coordinates are precisely recovered, as indicated by the low error E_y .

As η increases, E_y grows, while $\|c\|$ decreases since the prior probability of the solution gains more weight in the optimization. As the shape becomes more smooth and more plausible, the overall reconstruction error E_{full} decreases, and reaches its minimum at $\eta = 2$. If η is too large, the output is too close to the average to fit the data, so both E_y and E_{full} are high. The values on the right in Figure 6 are the baseline obtained with the average head $\bar{\mathbf{v}}$. Figure 6 is a typical curve observed in many fields of statistical learning [31].

In Figure 6, the optimal reconstruction achieves an average error $E_{full} = 2.72\text{mm}$ on the test set. Reducing the number of principal components m' from 99 to 40 increases the error to $E_{full} = 2.81\text{mm}$. The error increases also if the number of feature points is reduced: For $f = 17$, $E_{full} = 3.16\text{mm}$, while for $f = 1000$, $E_{full} = 2.24\text{mm}$, using $m' = 99$.

B. Noisy Feature Point Coordinates

To investigate the effect of noise on the feature point coordinates, we evaluate the system on the 100 training faces with their $m' = 99$ principal components. Without noise, the problem $\mathbf{Q}\mathbf{c} = \mathbf{y}$ has a unique solution if $l = 2 \cdot f \geq m'$. For $f = 50$ and $f = 1000$ feature points, the system correctly recovers the training faces: $E_y < 0.004$ pixels and $E_{full} < 0.002\text{mm}$. The solid line in Figure 8 shows that this perfect solution at $f = 50$ would be impaired by $\eta > 0$. If only $f = 17$ points are available, the face is no longer uniquely defined, and we found that our system produces a solution with smaller $\|c\|$ that still solves the problem ($E_y < 0.0001$), but is different from the original ($E_{full} = 2.1\text{mm}$), as expected from theory (Section IV and V-A).

If Gaussian noise with a standard deviation of $\sigma_N = 0.1$ or $\sigma_N = 1$ pixels is added to the horizontal and vertical image




Fig. 9. A sample of the set of 166 teeth (occlusal surfaces of first upper molars).




Fig. 10. The average and the first principal components of the Biogeneric Morphable Model of Teeth.

coordinates of each feature point, \mathbf{y} becomes more and more difficult to recover, and overfitting occurs, as demonstrated by the large errors E_{full} . The error in Figure 8 takes its minimum at $\eta = \sigma_N^2$, so reconstruction is best for the vectors with maximum posterior probability, as expected from Section V-A.

X. A BIOGENERIC MORPHABLE MODEL OF TEETH

Given the remaining dental material after parts have been lost, for example due to caries, the goal of reconstruction is to find a 3D surface that fits the remaining material, taking into account the large individual variations in teeth, but is at the same time morphologically plausible: The technique has to insert typical structures, such as cusps and grooves, even if they were entirely missing in the remaining material. Clearly, this would not be possible with geometrical methods such as Splines, NURBS and GDM [33], [34]. Other methods require manual selection of an intact prototype tooth that is adjusted in shape using image warping [35]. In contrast, our method automatically deforms and adjusts the Morphable Model based on the shape of the remaining material and the morphological variations found within the set of examples.

From impressions of 166 carious-free, intact first upper molar teeth of children within the age of 6 to 9 years, we formed stone replicas and measured the occlusal surface (i.e. the surface predominantly used in chewing) with a 3D scanner at a resolution of about $30\mu m \times 30\mu m$ points [36], [37]. The pattern of cusps and grooves varies considerably across individuals, even though some morphological properties, such as the overall layout of the main cusps and fissures, are shared by all samples (Figure 9). These common features allow us to establish correspondence between the scans $z(x, y)$ with a modified optical flow algorithm in an automated procedure [6]. Based on this correspondence, the x, y, z - coordinates form a Morphable Model of teeth \mathbf{v}_i (Equation 1), with an average and principal components shown in Figure 10. Note that the morphable model is flexible not only along the z direction, but also along x and y , which is essential to represent the varying locations of cusps in these directions.

A. Deformable Iterative Closest Point Algorithm

The only manual interaction, after scanning the residual dental material, is to select some feature points on that scan, such as the peaks of cusps or centers of grooves, and identify the corresponding points on the average tooth. Our method does not require a standard set of features, so if entire structures are missing, the user selects whatever is available.

The subsequent steps are performed automatically by our system: The 3D coordinates of the selected feature points are used to estimate the entire shape according to the method from Section V-B and with a full compensation for rigid transformations (Section VI). Figure 11 shows results of this algorithm (second and third column). To illustrate the quality of the fit, we substitute the remaining points of the original tooth into the recovered shape \mathbf{v} : For each vertex x_k, y_k, z_k in \mathbf{v} , we replace z_k by $z_{original}(x_k, y_k)$ if that value is non-void.

Clearly, the reconstruction from a sparse set of points will not consider details of the original shape, such as round versus pointed cusps. We therefore propose a multi-pass method that refines the shape, adapting it to whatever surface information is available from the remaining surface: After the first iteration, corresponding points on the model and the original are already close together. The algorithm selects a large number of pairs of nearest neighbors: it samples the original scan $(x, y, z_{original}(x, y))$ at a fixed grid resolution, and finds the nearest neighbor in \mathbb{R}^3 among the vertices x_k, y_k, z_k of the reconstructed shape \mathbf{v} . With this extended set of corresponding feature point, a refined reconstruction is computed (Figures 1 and 11). Since the correspondence mapping between scan and model may drift slightly along the surface with each refinement pass, additional passes may further improve the result slightly. Residual discontinuities at the boundary may be removed by smoothing.

To evaluate the precision of the reconstruction with ground truth, we randomly selected 66 test scans, removed 43% of the surface according to the pattern in Figure 1, and identified 10 feature points. The other 100 scans were used for PCA. After 2 passes of the algorithm, the average error $|z_{reconst.}(x, y) - z_{original}(x, y)|$ per point x, y and per scan from the true surface was $143\mu m$ on missing regions and $85\mu m$ on remaining material. Measuring distance not along the z -direction, but to the closest neighbour of each point, the average is $104\mu m$ on missing regions and $62\mu m$ on remaining material. The root-mean-square errors between 3D positions of closest neighbours are $137\mu m$ and $95\mu m$, respectively.

XI. CONCLUSION

We have presented an algorithm that uses a 3D Morphable Model to infer vector components, such as the coordinates of mesh vertices, from incomplete measurements in a robust and efficient way. Regularization provides an optimal tradeoff between fitting quality and plausibility. Additional compensation for rigid transformations is an essential component for processing real-world data.

The algorithm was tested in two applications: The first is an estimation of 3D shape of a high resolution face model from 2D feature coordinates. Unlike non-linear, iterative systems [6], our algorithm does not rely on color values of images, and therefore reconstructs characteristic details of individual faces only if they are sufficiently defined by feature points. However, the fast and non-iterative computation makes the algorithm

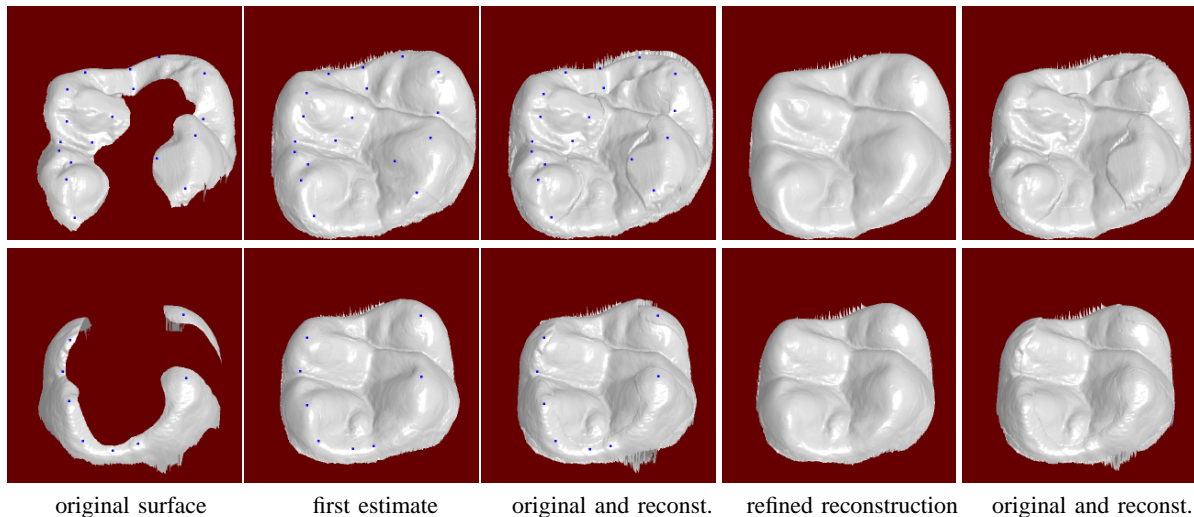


Fig. 11. From the remaining surface of two teeth (left, top and bottom) and a set of selected points (blue), the system produces a first estimate of the tooth. In the third column, we inserted the original surface into the reconstruction wherever available to demonstrate the quality of the fit. After a second pass, the reconstruction is refined (fourth and fifth columns).

an interesting low-cost alternative for interactive tools. The second application is a complete system for restoring 3D surfaces of teeth in medical contexts such as inlay design.

Our statistical approach and empirical evaluation reveal some of the general properties of surface estimation and restoration, and we provide an efficient algorithm for a wide range of practical applications.

REFERENCES

- [1] M. Ash and S. Ramfjord, *Occlusion*, 4th ed. W.B. Saunders, 1994.
- [2] R. Everson and L. Sirovich, "The karhunen-loeve transform for incomplete data," *J. Opt. Soc. Am., A*, vol. 12, no. 8, pp. 1657–1664, 1995.
- [3] B. W. Hwang, V. Blanz, T. Vetter, and S. W. Lee, "Face reconstruction from a small number of feature points," in *Int. Conf. on Pattern Recognition, ICPR2000*, Spain, 2000.
- [4] B. W. Hwang and S. W. Lee, "Reconstruction of partially damaged face images based on a morphable face model," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 25, no. 3, pp. 365–372, 2003.
- [5] M. Jones and T. Poggio, "Multidimensional morphable models: A framework for representing and matching object classes," *Int. Journal of Comp. Vision*, vol. 29, no. 2, pp. 107–131, 1998.
- [6] V. Blanz and T. Vetter, "A morphable model for the synthesis of 3D faces," in *Computer Graphics Proc. SIGGRAPH'99*, 1999, pp. 187–194.
- [7] D. Terzopoulos, "The computation of visible-surface representations," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 10, no. 4, pp. 417–438, 1988.
- [8] V. Blanz and T. Vetter, "Face recognition based on fitting a 3d morphable model," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 25, no. 9, pp. 1063–1074, 2003.
- [9] T. Poggio and F. Girosi, "Networks for approximation and learning," *Proc. of the IEEE*, vol. 78, no. 9, pp. 1481–1497, 1990.
- [10] J. Hornegger, D. Paulus, and H. Niemann, "Probabilistic modeling in computer vision," in *Handbook of Comp. Vision and Appl.*, B. Jähne and H. Haußecker, Eds. Academic Press, 2000, pp. 517–541.
- [11] M. A. Audette, F. P. Ferrie, and T. M. Peters, "An algorithmic overview of surface registration techniques for medical imaging," *Medical Image Analysis*, vol. 4, no. 3, pp. 201–217, 2000.
- [12] K. S. Arun, T. S. Huang, and S. D. Blostein, "Least-squares fitting of two 3-D point sets," *IEEE Trans. on Pattern Analysis and Mach. Int.*, vol. 9, no. 5, pp. 698–700, 1987.
- [13] B. K. P. Horn, H. M. Hilden, and S. Negahdaripour, "Closed-form solution of absolute orientation using orthonormal matrices," *Journal of the Optical Society of America A*, vol. 5, no. 7, pp. 1127–1135, 1988.
- [14] R. Haralick, H. Joo, C. Lee, X. Zhuang, V. Vaidya, and M. Kim, "Pose estimation from corresponding point data," *IEEE Trans on Systems, Man and Cyb.*, vol. 19, no. 6, pp. 1426–1445, 1989.
- [15] A. Ansar and K. Daniilidis, "Linear pose estimation from points or lines," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 25, no. 5, pp. 578–589, 2003.
- [16] P. J. Besl and N. D. McKay, "A method for registration of 3-D shapes," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 14, no. 2, pp. 239–256, 1992.
- [17] S. Rusinkiewicz and M. Levoy, "Efficient variants of the ICP algorithm," in *Proc. Conf. on 3D Digital Imaging and Modeling*, 2001, pp. 145–152.
- [18] V. Pratt, "Direct least-squares fitting of algebraic surfaces," *Comp. Graphics (SIGGRAPH 87)*, vol. 21, no. 4, pp. 145–152, 1987.
- [19] J. C. Carr, R. K. Beatson, J. B. Cherrie, T. J. Mitchell, W. R. Fright, B. C. McCallum, and T. R. Evans, "Reconstruction and representation of 3d objects with radial basis functions," in *Computer Graphics Proc. SIGGRAPH'01*, 2001, pp. 67–76.
- [20] A. Adamson and M. Alexa, "Approximating and intersecting surfaces from points," in *EG Symp. on Geometry Processing*, 2003, pp. 245–254.
- [21] D. G. Lowe, "Fitting parameterized three-dimensional models to images," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 13, no. 5, pp. 441–450, 1991.
- [22] F. Pighin, R. Szeliski, and D. Salesin, "Modeling and animating realistic faces from images," *Int. Journal of Computer Vision*, vol. 50, 2, pp. 143–169, 2002.
- [23] S. Umeyama, "Parameterized point pattern matching and its application to recognition of object families," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 15, no. 2, pp. 136–144, 1993.
- [24] B. Bascle and A. Blake, "Separability of pose and expression in facial tracing and animation," in *Int. Conf. Comp. Vision*, 1998, pp. 323–328.
- [25] M. Brand, "Morphable 3d models from video," in *Conf. on Computer Vision and Pattern Recognition*, 2001, pp. 456–463.
- [26] H. Li, P. Roivainen, and R. Forchheimer, "3-D motion estimation in model-based facial image coding," *IEEE Trans. on Pattern Analysis and Machine Intell.*, vol. 15, no. 6, pp. 545–555, 1993.
- [27] D. DeCarlo and D. Metaxas, "Adjusting shape parameters using model-based optical flow residuals," *IEEE Trans. on Pattern Analysis and Machine Intell.*, vol. 24, 6, pp. 814–823, 2002.
- [28] P. Eisert, T. Wiegand, and B. Girod, "Model-aided coding: A new approach to incorporate facial animation into motion-compensated video coding," *IEEE Trans. on Circuits and Syst for Video Techn.*, vol. 10, 3, pp. 344–358, 2000.
- [29] S. Haykin, *Neural Networks: A Comprehensive Foundation*. Prentice Hall, 1998.
- [30] J. Stoer, *Numerische Mathematik I*, 8th ed. Berlin: Springer, 1999.
- [31] V. N. Vapnik, *The Nature of Statistical Learning Theory*. New York: Springer-Verlag, 1995.
- [32] R. Duda, P. Hart, and D. Stork, *Pattern Classification*, 2nd ed. New York: John Wiley & Sons, 2001.
- [33] S. Adolph and S. Gürke, "Modeling of a fitting inlay from various information," in *VMV Stuttgart*, 2001, pp. 309–316.
- [34] L. Felber, T. Leemann, and W. H. Mörmann, "Computergestützte vollautomatische Konstruktion von Inlays," *Acta Med Dent Helv*, vol. 2, pp. 217–225, 1997.
- [35] D. Paulus, M. Wolf, S. Meller, and H. Niemann, "3d computer vision for tooth restoration," *Medical Image Analysis*, vol. 1.4, pp. 1–18, 1998.
- [36] A. Mehl, W. Gloger, K. H. Kunzelmann, and R. Hickel, "Entwicklung eines neuen optischen Oberflächenmessgerätes zur präzisen dreidim. Zahnvermessung," *Deutsche Zahnärztliche Z.*, vol. 51, pp. 23–27, 1996.
- [37] A. Mehl, W. Gloger, K. H. Kunzelmann, and R. Hickel, "A new optical 3-d device for the detection of wear," *J Dent Res*, vol. 76, pp. 1799–1807, 1997.