# AUTOMATIC AND ROBUST SEMANTIC REGISTRATION OF 3D HEAD SCANS

**David C. Schneider, Peter Eisert**

Fraunhofer Institute for Telecommunications—Heinrich-Hertz Institute
Einsteinufer 37, 10587 Berlin, Germany
{david.schneider | peter.eisert}@hhi.fraunhofer.de
Fax: +49-30-3927200

**Keywords:** 3D face processing, face registration, morphable head models, geometry reconstruction

**Abstract**

A robust and fully automatic algorithm for semantic registration of laser scans of human faces is introduced. The scans are converted into a new semantically annotated topology such that topologically equivalent vertices have the same meaning. The algorithm works in three steps. First, landmark points are placed using a multi-stage variant of the iterative closest points scheme and an annotated reference mesh. Second, the scan is matched to the reference mesh globally by a nonlinear transformation. Finally, the transformed mesh is resampled in the reference mesh topology and the vertex locations are tracked back to the original scan. The use of the algorithm for the construction of a morphable head model is described. Moreover, a method for repairing defective head scans (i.e. scans with missing or corrupted vertex data) using the registration algorithm and the model is introduced and experimentally investigated. It is shown that a face can be reconstructed convincingly using as little as 20 percent of its vertices.

## 1 Introduction

In this paper we introduce a robust method for full semantic annotation (or registration) of laser-scanned human faces that does not require manual intervention. Also, the scans are transformed to a new mesh topology more suitable for applications such as 3D graphics production and also for computer vision research.

Laser scanners are the prime tool for acquiring detailed 3D models of human faces. However, the meshes generated by laser scanners typically have a topology that reflects the operation of the scanner and is unsuitable for many applications. The data-set used for this work, for example, has a cylindrical grid topology with vertices of the form $(\phi_i, z_i, r_i)$ where $\phi_i$ are regularly spaced angles, $z_i$ are regularly spaced vertical offsets and $r_i$ are varying radii; see fig. 1 (left) for an example. The aim of the algorithm to be introduced is to construct meshes from arbitrary face scans with the following properties:

- All meshes have the same topology; different faces only vary in vertex locations.
- The mesh topology is semantically interpretable, i.e. topologically equivalent vertices in different faces have the same "meaning" such as tip of the nose, center of upper lip, etc.

Fig. 1 (right) shows the result of our algorithm applied to the laser-scan in the same figure. Meshes of this type are useful in various areas of application and research. Some examples:

(a) In 3D media production, a facial animation rig can be defined once in respect to a generic, semantically annotated mesh. Scanned faces can be converted to the topology of the generic mesh. The rig can then be automatically transformed to the new face's geometry and the face can be animated immediately without manual rigging.

(b) In biometry, a face needs to be semantically annotated in order to measure distance ratios between facial landmarks and other properties.

(c) In computer vision research, linear subspace models of face geometry, often called "morphable head models" [5, 15], are a popular and powerful tool for tracking, 3D reconstruction, image analysis and other tasks. To generate a morphable model a database of semantically annotated head scans is required. We briefly address the construction of a morphable model with the help of our algorithm in section 4.1.

(d) The registration algorithm can be used together with the morphable model to improve or repair the results of 3D face acquisition techniques. This is especially useful for error-prone vision techniques like stereo analysis but also for model based repairing of scanning errors in established methods like structured light or laser scanners. For that purpose, an acquired mesh must first be registered with the model before model information can be used to improve the scan. In section 4.2 we show how our algorithm and model can be used to reconstruct face geometry even from highly corrupted vertex data.

## 2 Related work

Face registration algorithms can roughly be divided into methods that exploit texture information and purely geometric

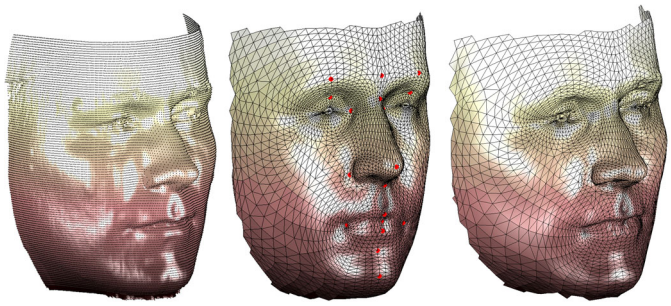

***Figure 1: Left: Typical topology of a typical laser-scan. Center: The reference mesh used for topology conversion with landmarks for the first algorithm stage. Right: The laser-scan semantically registered in the reference mesh topology with our algorithm.***

approaches. Texture based approaches were primarily developed in the context of data-driven head modelling, an application that is also the motivation behind our work; see section 4.1. On the side of geometric methods, some specialize exclusively on faces while others try to solve the general problem of registering arbitrary surfaces by nonrigid transformations. The algorithm proposed in this paper belongs to the class of face-specific purely geometric approaches.

On the side of texture-based methods Blanz and Vetter [5, 15] use an algorithm based on optical flow. Traditionally used to estimate 2D image deformation, optical flow is extended to use 2D texture as well as local 3D surface properties of the face in order to compute a matching deformation. However, the authors themselves [5] as well as Paterson and Fitzgibbon [11] report the method to be unreliable on "exotic" faces. Thus Paterson and Fitzgibbon [11] manually annotate 30 landmark points in the texture and use radial basis function interpolation to compute a matching warp. 3D points are found by inverse-warping the texture coordinates associated with the 3D points.

Purely geometric head registration is used by Kähler et al. [9] to animate facial expressions of laser scanned heads with an anatomical animation model. Their method is based on iteratively subdividing a rough face mesh and aligning the new vertices to a reference model. The initial model for the subdivision is obtained from manually placed landmarks.

Several geometric algorithms were proposed in the context of data driven modelling. E.g. Allen et al. [1] register full human body scans by using a generic nonlinear optimizer on a three-term error function. The function penalizes distances of closest points and dissimilarities between transforms of nearby points, thereby controlling the rigidity of the overall transformation. However, a third term penalizing distance between manually annotated landmarks is required to prevent the method from falling into local minima. Recently, Amberg et al. [2] used the Iterative Closest Points (ICP) scheme—which is also used in our algorithm—to minimize a similar error function in a stepwise optimal fashion. Their algorithm performs excellent on faces as well as general nonrigid registration problems; however, it still requires manual annotation of several landmarks.

Geometric methods were also developed in the field of face recognition and database retrieval. They typically aim at registering faces in a common frame of reference before extracting the features used for classification or retrieval of a prototype face in the database. For example, Li and Barreto [10] obtain profile curves by intersecting the laser scan mesh with axis parallel planes. More sophisticated, Haar and Veltkamp [13] use curves originating from the nose in all directions and achieve good retrieval results. These techniques, however, aim at retrieving a model of the same person rather than registering multiple faces in a semantically valid way.

## 3 Automatic semantic registration of head scans

The algorithm presented here is face specific; it works on 3D geometry alone and does not exploit or require texture information. Also, no manual landmark annotation is required.

The algorithm comprises three steps: First landmark points are found on the laser-scan mesh using a modification of the iterative closest points (ICP) scheme. Second, a nonlinear transform is computed to match the laser-scan with a reference mesh. Finally the transformed mesh is resampled in the topology of the reference mesh and vertex locations are tracked back to the original scan.

### 3.1 Automatic landmark placement

The ICP algorithm [3], which we modified for our purpose, computes a rigid transform that aligns two or more partially overlapping meshes—typically laser scans of a larger scene—such that the overlapping parts match as good as possible according to some error function. Typical error functions are distances of closest points or point-to-plane distances (e.g. [7]). The optimal transform depends on matching the right points in both meshes; to match the right points, however, the optimal transform must be known. Therefore, it is simply assumed that the closest points match, computes and applies the transform induced by these correspondences and iterates. ICP is guaranteed to converge to a minimum of the error function; see [3]. However, this minimum may be local and thereby the meshes must be roughly aligned before ICP is invoked to find a good match. There is a huge body of literature on ICP and numerous optimizations have been proposed, aiming mostly at improving stability and speed of convergence; see [12] for an overview.

The ICP algorithm is quite general and the way it is employed in our work differs in two aspects from its typical usage. Firstly, we don not use it on partially and exactly overlapping meshes but rather on fully overlapping but only roughly similar meshes. Secondly, while ICP is typically used with a rigid transform, we exploit the fact that it works—in principle—with any kind of transform that can be uniquely solved for from a set of corresponding points. However, the degrees of freedom of the transform directly correspond with the algorithm's stability:

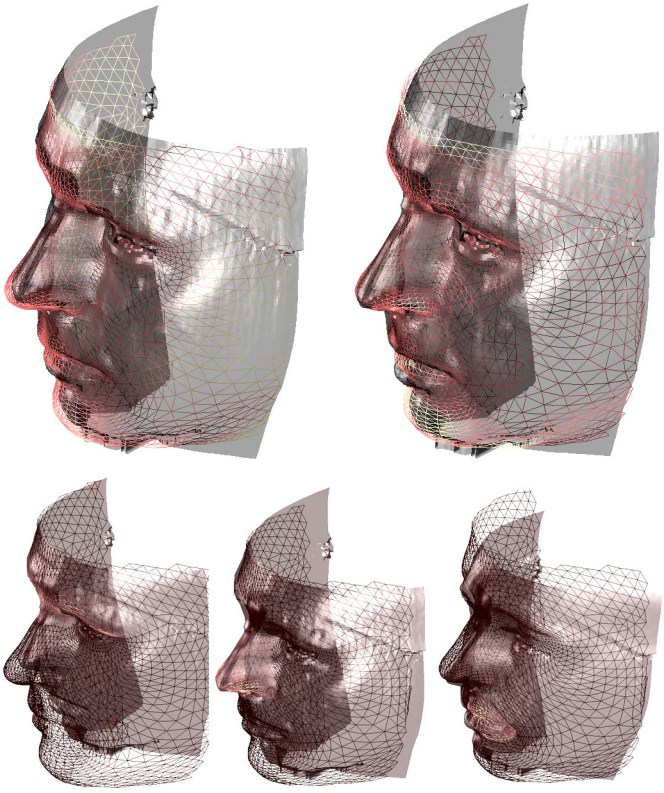

*Figure 2: Stage 1 rigid registration (top left), stage 2 affine registration (top right), stage 3 partial affine registrations of eyes, nose, left mouth (bottom row).*

The more degrees of freedom there are, the more likely ICP is to converge to an undesired local minimum.

In our algorithm, landmarks are found by matching a predefined reference mesh with annotated landmarks—depicted in fig. 1 (center)—to the laser-scan point cloud. The first approximation of this matching is computed with the common ICP method and results in a rigid transformation matrix $\mathbf{T}_r$. Due to the stability of rigid transform ICP it is sufficient for initialization to normalize the scale of reference mesh and laser-scan point cloud and align their principal axes. Clearly, after this first stage landmarks are far from their semantically corresponding points in the reference mesh due to the vastly different proportions of facial features in different people; see fig. 2 top left.

Therefore, in the second stage reference mesh and point cloud are aligned with an affine transform ICP thereby allowing the laser scan to deform; this yields an affine matrix $\mathbf{T}_a$. Note that only due to the rigid ICP before, the point cloud's initial alignment is good enough for the affine ICP to converge to a better match. Still, landmark correspondence is insufficient (fig. 2 top right).

Thus, in a third iteration, affine ICP is performed on parts of the model: Instead of using the full reference mesh as ICP target, only the vertices of $k$ different predefined facial parts are

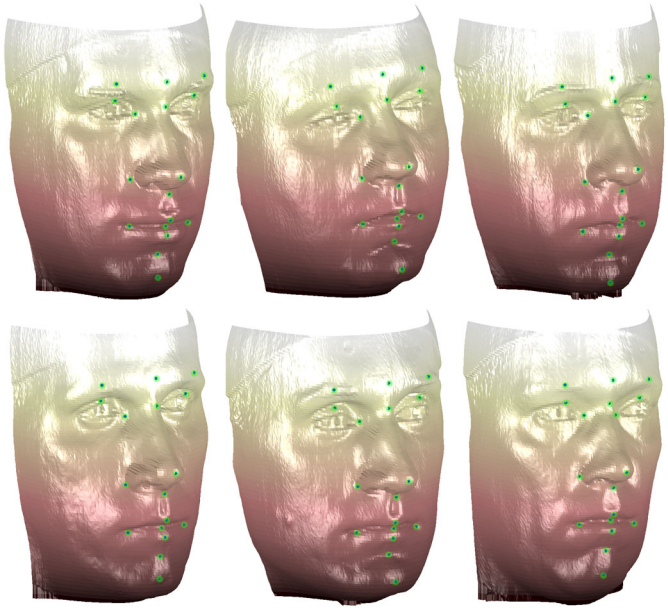

*Figure 3: Examples of automatically placed landmark points in six different face scans.*

employed. Three of these parts are illustrated in the bottom row of fig. 2. Each landmark point in the reference mesh is assigned one of the parts as its authoritative frame of reference (even if some landmarks appear in multiple parts due to overlap). The result of the third stage is a set of $k$ different affine transforms, $\mathbf{T}_{a,1}...\mathbf{T}_{a,k}$ (fig. 2 bottom).

After the third stage landmark correspondence between reference mesh and laser-scan point cloud is sufficiently precise. Thus let $\mathbf{l}$ be a landmark location in the reference mesh belonging to part $i$ and $\mathbf{p}$ the closest point in the laser-scan cloud to $\mathbf{l}$. Then the location $\mathbf{q}$ of the landmark point in the original, untransformed laser-scan cloud is given by

$$\mathbf{q} = \mathbf{T}_r^{-1}\mathbf{T}_a^{-1}\mathbf{T}_{a,i}^{-1}\mathbf{p}. \quad (1)$$

Examples for automatically placed landmarks in several head scans are shown in fig. 3.

### 3.2 Nonlinear transformation

In the second step of the algorithm, the laser-scan point cloud is matched to the reference model in a global, nonlinear fashion using the corresponding landmarks, $\mathbf{l}_i$ in the reference mesh and $\mathbf{q}_i$ in the laser-scan: We seek a transformation $\mathcal{T}(\cdot)$ such that $\mathcal{T}(\mathbf{q}_i) = \mathbf{l}_i$ for all $i$ and such that the points in between are interpolated in a natural fashion. Note that the transformations in the third ICP stage are local and linear and can thereby not be used for the global nonlinear matching; if, for example, $\mathbf{T}_{a,1}\mathbf{T}_a\mathbf{T}_r$ matches the noses it won't match the mouths. The global matching to be computed should be as conservative as possible in the nonlinear part in order to avoid awkward deformations of the head geometry. This is realized by the thin plate spline formalism (see, for example, Bookstein [6]), which yields a transformation minimizing a physically interpretable

bending energy.

From a thin plate spline perspective we solve a three dimensional scattered data interpolation problem. The known values are the displacement vectors from the laser-scan landmark points to the reference mesh landmarks,

$$\mathbf{d}_i = \mathbf{l}_i - \mathbf{q}_i. \tag{2}$$

The unknowns to be interpolated are the displacement vectors of all non-landmark points in the laser-scan mesh. Dealing with a three dimensional problem, the basis function to use for the spline is $u(x) = |x|$ according to [6] and thus, with $u_{i,j} = ||\mathbf{q}_i - \mathbf{q}_j||$, we get a thin plate spline matrix

$$\mathbf{S} = \begin{bmatrix} 0 & u_{1,2} & \dots & u_{1,k} & 1 & \mathbf{q}_1^T \\ u_{2,1} & 0 & & u_{2,k} & 1 & \mathbf{q}_2^T \\ \vdots & & \ddots & \vdots & \vdots & \vdots \\ u_{k,1} & u_{k,2} & \dots & 0 & 1 & \mathbf{q}_k^T \\ 1 & 1 & \dots & 1 & 0 & \\ \mathbf{q}_1 & \mathbf{q}_2 & \dots & \mathbf{q}_k & & \mathbf{0}_{3\times 3} \end{bmatrix}. \tag{3}$$

Hence the weight matrix for the mesh transform is

$$\mathbf{W} = \mathbf{S}^{-1} \begin{bmatrix} \mathbf{d}_1 & \dots & \mathbf{d}_k & \mathbf{0}_{3\times 3} \end{bmatrix}^T \tag{4}$$

and an arbitrary point $\mathbf{p}$ in the laser-scan transforms to

$$\mathbf{p}' = \mathbf{p} + \mathbf{W}^T \begin{bmatrix} u_1 & \dots & u_k & 1 & \mathbf{p}^T \end{bmatrix}^T \tag{5}$$

with $u_i = ||\mathbf{p} - \mathbf{q}_i||$.

### 3.3 Resampling

The final step is to resample the non-linearly deformed laser-scan—which now closely matches the reference mesh—in the topology of the latter. Therefore, for each point $\mathbf{v}$ in the reference mesh (fig. 1, right), its point normal $\mathbf{n}_v$ is computed and the intersection point $\mathbf{p}^d$ of the deformed laser-scan mesh with a line through $\mathbf{v}$ in the direction of $\mathbf{n}_v$ is determined. If there are multiple intersections the one closest to $\mathbf{v}$ is used; moreover there is a threshold on the distance $||\mathbf{v} - \mathbf{p}^d||$ to exclude deficient matches at the scan's border. The intersection point is taken to be the semantically equivalent point to $\mathbf{v}$ in the laser-scan point cloud. However, $\mathbf{p}^d$ is in the *deformed* laser-scan which is not what we ultimately aim at. Therefore, the vertices $\mathbf{t}_1^d$, $\mathbf{t}_2^d$, $\mathbf{t}_3^d$ of $\mathbf{p}^d$'s enclosing triangle in the deformed scan are determined as well as their topological equivalences $\mathbf{t}_1$, $\mathbf{t}_2$, $\mathbf{t}_3$ in the undeformed mesh. Finally, the equivalent $\mathbf{p}$ to $\mathbf{v}$ in the original laser scan mesh is given by tracking $\mathbf{p}^d$ back with barycentric transformation:

$$\mathbf{p} = \begin{bmatrix} \mathbf{t}_1 & \mathbf{t}_2 & \mathbf{t}_3 \end{bmatrix} \begin{bmatrix} \mathbf{t}_1^d & \mathbf{t}_2^d & \mathbf{t}_3^d \end{bmatrix}^{-1} \mathbf{p}^d. \tag{6}$$

To speed up the computation of the ray-mesh-intersection, the deformed laser scan mesh is represented as an axis aligned bounding box tree and the fast ray box intersection of Williams et al. [16] is used. Note that due to the size of a typical laser-scan it is not feasible to build the tree down to the level of individual triangles. In our implementation with approximately 35 000 points in a laser-scan, there are 200 triangles per leaf that have to be tested for each ray.

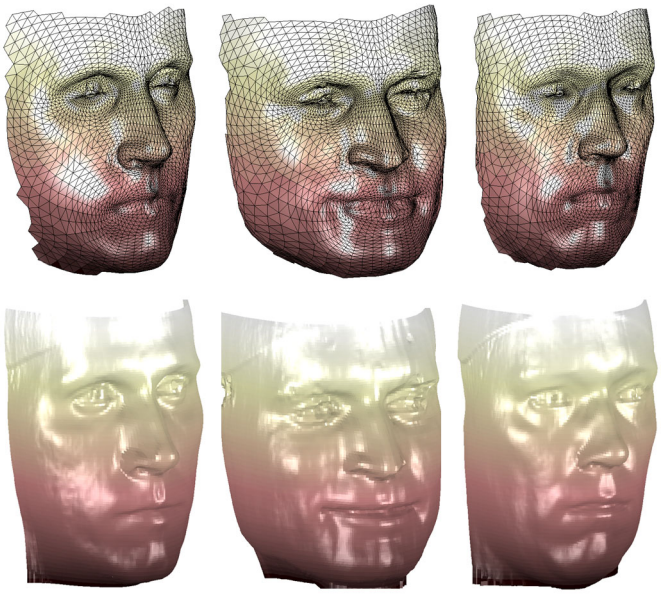

***Figure 4: Final results of the algorithm after resampling (top row). Original laser scans (bottom row).***

### 3.4 Reference mesh

We found that the success of the first stage—automatic landmark placement—partly relies on choosing a good reference mesh. The mesh we currently use successfully is itself created from a small number of laser-scans by an interactive variant of the algorithm described: The landmarks are placed manually in the scans and the second and third stage of the algorithm are computed as described above using an artificial (i.e. not scanned) CG head mesh as reference. The new reference mesh is obtained by aligning the registered meshes with rigid ICP and taking the mean location for each vertex.

## 4 Application: Repairing scans with a morphable model

In section 1 we sketched a number of applications for the registration algorithm. Here we describe how a morphable head model can be built from automatically registered scans. Together with the registration method we introduced the model can be used to repair defective parts of head scans from arbitrary acquisition techniques. We demonstrate experimentally that model based mesh repair is far more powerful than simple geometry interpolation.

### 4.1 Morphable model construction

Automatic semantic scan registration greatly facilitates the construction of a morphable head model in the style of Blanz and Vetter [5, 15] which we shall briefly sketch in the following, giving more detail only where we deviate from the typical methodology. We built a database of 180 male head scans which were processed with the method introduced above. The registered and resampled scans are aligned with

rigid ICP to prevent the model from capturing mere variances in location and orientation. The head model is split up in five parts—eyes, nose, mouth, chin and the rest—that overlap at their borders. For each part, each component of each vertex is treated as a variable and a principal component analysis (PCA) is performed using the algorithm for high dimensional data [14, 4]. The principal components obtained describe the most significant variations of the face parts in order of variance explained. An approximation of an arbitrary face can be described by a $n \times 5$ weight matrix $\mathbf{W}$ where $n$ is the number of principal components to use. The mesh geometry vector $\mathbf{g}_k$ of the $k$th part is then given by

$$\mathbf{g}_k = \phi_k + \mathbf{P}_k \mathbf{w}_k \tag{7}$$

where $\phi_k$ is the mean of the part, $\mathbf{P}_k$ its principal component matrix and $\mathbf{w}_k$ the $k$th column of the weight matrix. Conversely, the weight matrix $\mathbf{W}$ for a previously registered model with part geometry $\mathbf{m}_1 ... \mathbf{m}_k$ is given by

$$\mathbf{W} = \left[\mathbf{P}_1^+(\mathbf{m}_1 - \phi_1) \; ... \; \mathbf{P}_k^+(\mathbf{m}_k - \phi_k)\right] \tag{8}$$

where $\mathbf{P}^+$ is the pseudo inverse of $\mathbf{P}$.

To construct a full face mesh without creases, the parts have to be blended which is not addressed in detail in [5]. The blending weights we use only depend on a part's topology and can therefore be computed in advance for all scans: Treating a part as an undirected graph, let $\mathrm{d}(v, w)$ be the length of a shortest path (in the sense of graph theory) between vertices $v$ and $w$. Then a vertex is assigned the weight

$$w = \min_{b \in B} \{\mathrm{d}(v, b)\} + 1 \tag{9}$$

where $B$ is the set of border vertices of the face part in question. Whenever a point appears in multiple parts, say $\mathbf{v}_1, ..., \mathbf{v}_n$ with weights $w_1 ... w_n$, the point $\mathbf{p}$ in the blended mesh is given by

$$\mathbf{p} = \frac{1}{\sum_{i=1}^{n} w_i} \sum_{i=1}^{n} (w_i \mathbf{p}_i) \,. \tag{10}$$

An example of a face reconstructed with the model is shown in fig. 5; the first face in the first row is the original, its right neighbor is generated with the model using 50 eigenvectors per face part.

### 4.2 Repairing scans: Reconstruction from incomplete data

Repairing a defective mesh comprises the following steps.

First, defective parts are manually marked in the laser-scan.

Second, the scan is registered with the model using the algorithm introduced. The defective parts are left out at the ICP stage but included in the morphing and resampling stage. Thereby all vertices are present in the resampled model but some are marked as defective.

Third, the defective vertices are removed and the weight matrix $\mathbf{W}$ of equation (7) has to be determined from the incomplete data. Thus, in equation (8), the part geometry vectors $\mathbf{m}_i$ may have a multiple of three less components than required. A straightforward way of dealing with this is to simply remove the corresponding components from the part's principal component matrix $\mathbf{P}_i$ and mean vector $\phi_i$ and solve for the weights. Denoting the reduced principal component matrix by $\tilde{\mathbf{P}}_i$ and the reduced mean vector by $\tilde{\phi}_i$, the full and repaired mesh part geometry $\mathbf{m}_{rep}$ is given by

$$\mathbf{m}_{rep} = \phi_i + \mathbf{P}_i \left(\tilde{\mathbf{P}}_i^+(\mathbf{m}_i - \tilde{\phi}_i)\right) . \tag{11}$$

Finally, the reconstructed model parts are blended using eq. 10 and the defective vertices in the resampled model are replaced by the corresponding vertices of the model reconstruction.

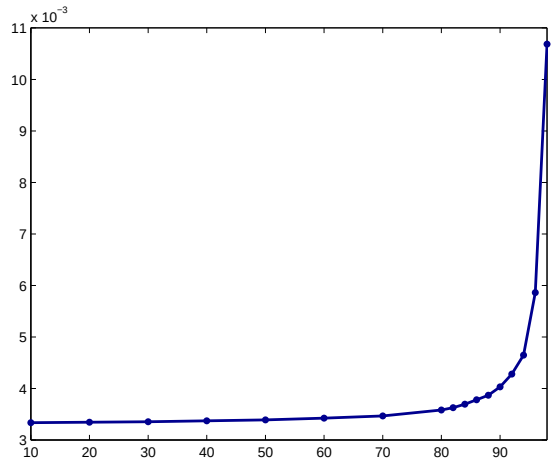


***Figure 6: Mean reconstruction error over percentage of missing vertices.***

To investigate the model's capabilities for reconstruction with missing data in a principled way, the following experiment is conducted: In ten randomly chosen head scans that are not included in the PCA model an increasing number of randomly chosen vertices—uniformly distributed over the face—is removed and scans are reconstructed with the method introduced above. Fifty eigenvectors are used per component. The missing data reconstructions are compared to the full data reconstructions of the faces. The error is measured as mean distance between corresponding points in the missing and the full data reconstruction. Each of the ten heads is reconstructed five times, resulting in 50 trials for each missing-data percentage.

The mean error curve of the experiment subject to missing data percentage is depicted in fig. 6 showing that the error only begins to rise sharply at over 90 percent missing vertex data. This is also illustrated in fig. 5 where several faces from the reconstruction experiment are shown. With 80 percent missing data the result is still visually convincing compared to the full data reconstruction. Only at 98 percent the reconstructed face clearly deteriorates. The experiment shows that with a good

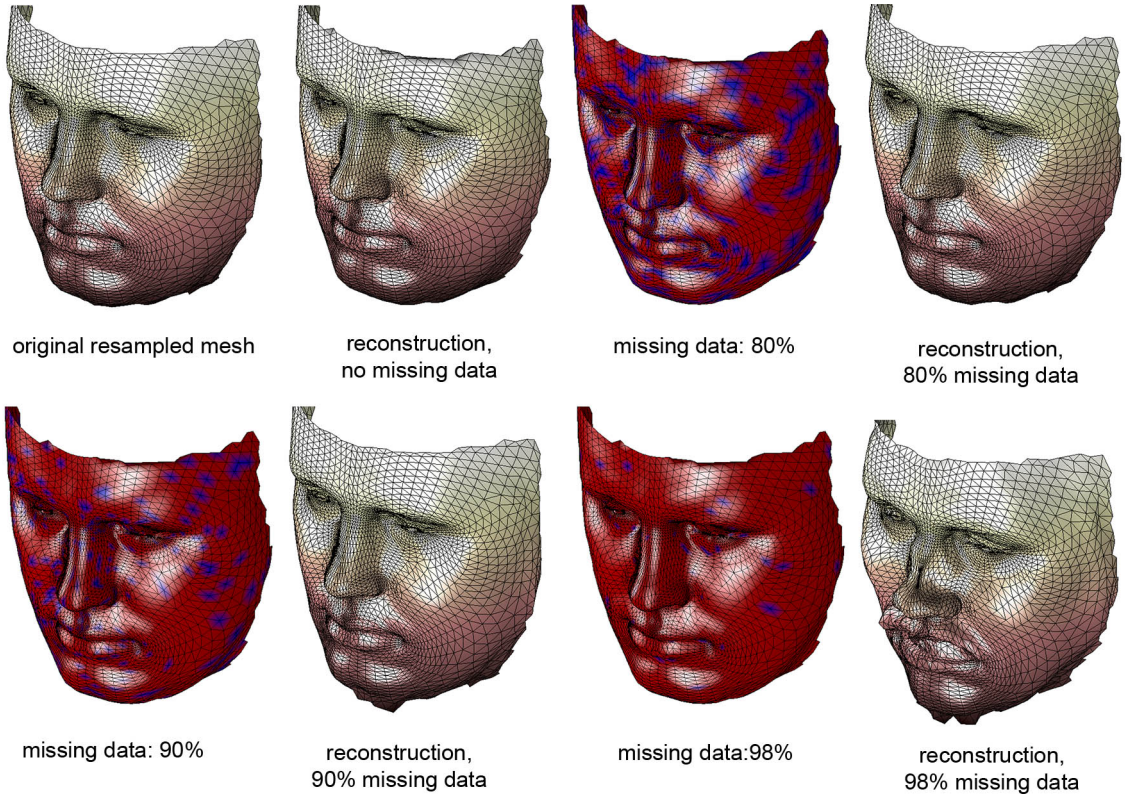


***Figure 5: Morphable model based reconstruction with and without missing data. The red/blue meshes show the vertices used (blue) and discarded (red).***

model very few correct vertices suffice to reconstruct a visually convincing face model.

## 5 Conclusion and outlook

We introduced a new robust method for registering laser-scans of human faces. The registered faces are transformed to a new semantically annotated topology defined by a reference mesh. The algorithm facilitates the construction of linear subspace models of face geometry. Together, the registration algorithm and the model can be used to repair defective face meshes from arbitrary 3D acquisition techniques.

In comparison to existing methods our algorithm is simple, robust and straightforward. As it is based on the traditional, well-understood ICP scheme the numerous optimizations that were developed therefor can be applied to speed up the registration process; indeed traditional ICP can be considered real-time capable [12]. Also our algorithm does not require manual landmark annotation. On the downside the nonlinear deformation scheme currently employed may not yield results as precise as in more refined schemes such as [2]. For face registration, the first stage of our algorithm could therefore be used for automatic landmark placement followed by a general non-rigid ICP method for warping.

Our future research will concentrate on the use of the algorithm and the morphable model in 3D acquisition. The model can not only be used for repairing a defective model but also at earlier stages of the acquisition process, e.g. for providing priors in a 3D reconstruction algorithm like the single-shot structured light method we introduced in [8]. Also, reconstruction from very few points as suggested by the experiment described is to be further investigated.

## References

[1] Brett Allen, Brian Curless, and Zoran Popovic. The space of human body shapes: Reconstruction and parameterization from range scans. In *ACM SIGGRAPH*, 2003.

[2] Brian Amberg, Sami Romdhani, and Thomas Vetter. Optimal step nonrigid icp algorithms for surface registration. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2007.

[3] Paul J. Besl and Neil D. McKay. A method for registration of 3-d shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14:239–256, 1992.

[4] Christopher M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2007.

[5] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3d faces. In *Proceedings of the 26th*

*annual conference on Computer graphics and interactive techniques*, 1999.

[6] F. L. Bookstein. Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11:567 – 585, 1989.

[7] Yang Chen and Gérard Medioni. Object modeling by registration of multiple range images. *Image and Vision Computing*, 10:145 – 155, 1992.

[8] Philipp Fechteler, Peter Eisert, and Jürgen Rurainsky. Fast and high resolution 3d face scanning. In *Proc. International Conference on Image Processing*, volume III, pages 81–84, San Antonio, USA, 2007.

[9] Kolja Kähler, Jöorg Haber, Hitoshi Yamauchi, and Hans-Peter Seidel. Head shop: Generating animated head models with anatomical structure. In *Proceedings of the 2002 ACM SIGGRAPH*, 2002.

[10] Chao Li and Armando Barreto. Profile-based 3d face registration and recognition. In *Information Security and Cryptology, ICISC 2004*, Lecture Notes in Computer Science. Springer, 2005.

[11] J. Paterson and A. Fitzgibbon. 3d head tracking using non-linear optimization. In *British Machine Vision Conference 03*, 2003.

[12] Szymon Rusinkiewicz and Marc Levoy. Efficient variants of the icp algorithm. In *Third International Conference on 3D Digital Imaging and Modeling*, 2001.

[13] Frank B. ter Haar and Remco C. Veltkamp. A 3d face matching framwork. In *Proceedings IEEE Shape Modeling International*, 2008.

[14] Matthew Turk and Alex Pentland. Face recognition using eigenfaces. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 1991.

[15] Thomas Vetter and Volker Blanz. Estimating coloured 3d face models from single images: An example based approach. In *Computer Vision - ECCV 98*, volume Volume 1407/1998 of *Lecture Notes in Computer Science*. Springer, 1998.

[16] Amy Williams, Steve Barrus, R. Keith Morley, and Peter Shirley. An efficient and robust ray-box intersection algorithm. In *International Conference on Computer Graphics and Interactive Techniques*, 2005.