

Fully Automated and Highly Accurate Dense Correspondence for Facial Surfaces

Carl Martin Grewe^(✉) and Stefan Zachow

Mathematics for Life and Materials Sciences, Zuse Institute Berlin, Berlin, Germany
{grewe,zachow}@zib.de




Fig. 1. Two facial expressions (a,b) from our database set into dense correspondence using the proposed framework. High geometric and photometric details are accurately morphed between both expressions via a dense corresponding mesh.

Abstract. We present a novel framework for fully automated and highly accurate determination of facial landmarks and dense correspondence, *e.g.* a topologically identical mesh of arbitrary resolution, across the entire surface of 3D face models. For robustness and reliability of the proposed approach, we are combining 2D landmark detectors and 3D statistical shape priors with a variational matching method. Instead of matching faces in the spatial domain only, we employ image registration to align the 2D parametrization of the facial surface to a planar template we call the *Unified Facial Parameter Domain* (UFPD). This allows us to simultaneously match salient photometric and geometric facial features using robust image similarity measures while reasonably constraining geometric distortion in regions with less significant features. We demonstrate the accuracy of the dense correspondence established by our framework on the BU3DFE database with 2500 facial surfaces and show, that our framework outperforms current state-of-the-art methods with respect to the fully automated location of facial landmarks.

Keywords: Dense face matching · Face shape and appearance models · Markerless motion capture

Electronic supplementary material The online version of this chapter (doi:[10.1007/978-3-319-48881-3_38](https://doi.org/10.1007/978-3-319-48881-3_38)) contains supplementary material, which is available to authorized users.

1 Introduction

The fully automated matching of sparse or dense facial landmarks in unconstrained 2D or 3D measurement data, *e.g.* the semantic annotation of facial images captured *in the wild*, is of great interest in various fields, ranging from entertainment to affective computing. When dealing with conventional cameras, the loss of information due to the perspective projection requires sophisticated techniques for robust estimation of pose or facial landmarks. Even more demanding is the ill-posed inverse problem of estimating the 3D shape from 2D images. Knowledge about plausible variations in facial shape and appearance as well as their correlation are learned from training samples and used to constrain results to desired solutions especially in unconstrained environments. Similarly, for the semantic annotation and tracking of facial features from 3D data, statistical shape and appearance models (SSAM) of faces improve the reliability and robustness of automated approaches as has been shown recently [1].

Facial morphology varies between individuals due to factors like sex, age, or ethnicity, while significant intra-individual changes are caused by facial expressions. Although 3D databases including a wide variety of both, inter- and intra-individual factors, are publicly available (*e.g.* [2,3]), the training samples used to construct statistical face models are restricted to face scans in neutral position (see [4,5]). Only few models include expressions, for instance the work published by Brunton, Bolkart, and Wuhler in [6] or Cao *et al.* [7]. Unfortunately, these models do not include appearance and solely capture 3D shape variation. They are thus limited for applications in computer vision.

A reason for the rare availability of statistical models of facial shape and appearance lies in the challenging problem of dense correspondence estimation for faces. Many generic shape matching methods as well as approaches specifically tuned to estimate dense correspondence for faces have been proposed, but they either lack accuracy, robustness or automation. Nevertheless, approaches satisfying all these characteristics are needed to establish the next generation of 3D face models, and in order to handle improved geometric and photometric resolution of new scanning devices, growing 3D databases, and applications requiring highly accurate semantic annotation of faces in raw measurement data.

With applications for fully automated processing of large-scale databases in mind, we propose a new framework for dense 3D face matching (see Fig. 2). To ensure robustness of the automated processing, we extract reliable prior knowledge on facial shape and appearance from the input data using 2D facial landmark detectors and non-rigid fitting of 3D face models. Highly-accurate dense correspondence, even for fine facial structures (see Fig. 1), is obtained by combining the prior knowledge with a variational approach for the matching of geometric and photometric facial features. We evaluate and compare the performance of our method in localizing facial landmarks on 2500 scans of the publicly available BU3DFE dataset [2] as well as on 400 high-resolution 3D face models acquired using our own prototypic stereophotogrammetric setup. The accuracy of the dense correspondence established by our method can not only be used to improve various applications such as the retargeting of facial shape and texture




Fig. 2. The proposed framework includes two stages for initialization and dense matching of accurate correspondence. The matching allows to transfer semantic annotations and a reference mesh to the input data.

or the detection of facial expressions. By providing the basis for fully automated computation of individual blendshape rigs as well as large-scale statistical face models, our framework opens up new directions for computer vision tasks, particularly in the emerging field of consumer devices equipped with 3D sensors.

2 Related Work

Semantic face annotation has been subject of active research in different communities during the last twenty years. The general problem can be stated as the definition of inter-individually corresponding facial landmarks, ranging from few landmarks at clear anatomical structures to an arbitrary number of points covering the entire face, and their identification in raw measurement data. However, the measurement device and its sensor characteristics affect how accurately significant features can be located and distinguished from other landmarks as well as from surrounding facial and non-facial parts.

In the case of 2D images taken with conventional cameras, a great variety of algorithms exist for the detection of sparse facial landmarks [8]. Usually, locally significant, intensity-based features around landmark points are extracted from facial images contained in a database and used to train landmark predictors. Current methods are able to locate the silhouette of a face, as well as a number of sparse landmarks reliably from the frontal view in presence of a wide range of inter- and intra-individual variation even in unconstrained situations [9–12]. Similarly, for the detection of sparse landmarks on 3D measurement data, knowledge about characteristic geometric properties is gathered from an annotated database. For instance, in [13–15] local quantities like geodesic length, surface area or curvature measures are employed to learn the relevant features of distinct facial landmarks for later prediction.

More challenging is the problem of establishing dense correspondence across the entire face where landmarks cannot be clearly defined by local photographic or geometric features. Instead, correspondence estimation in regions like the cheeks or the forehead is usually constrained by means of mathematical objectives. In the case of 2D warping techniques, the topological subdivision of the

facial region into geometric primitives allows the definition of dense correspondence. For example in [16, 17], triangular patches covering the facial region are established and affinely warped to match new landmark positions. These approaches yield continuous correspondence mappings for the entire face varying inter- and intra-individually, but suffer from the strong assumption of affine warps.

Recent methods that directly operate in the 3D domain take advantage of the ability to measure distortion of the surface or the embedding space when deforming a shape into another. In the computer graphics community, general non-rigid shape matching approaches have been developed that are based, for instance, on the *as-isometric-as-possible* assumption or by measuring the deformation energy (e.g. [18–22] for a comprehensive survey). A common strategy often adopted in computer vision tasks is to use spatial warping techniques, like non-linear variants of the well-known Iterative Closest Point algorithm (ICP) [23], that deform a template surface into the target e.g. by locally constraining coherent deformation of the surface (see *Coherent Point Drift* for a general technique [24], and [4, 5] particularly for faces).

When matching objects of a specific type, like human faces, methods significantly benefit from additional prior knowledge that is incorporated into the matching process. For instance in [7, 14], a 3D statistical shape model (SSM) of the face is fitted to the target. Non-linear ICP is then used to warp the template to the target in order to project dense correspondence. Similarly in [15], Gilani, Shafait, and Ajmal combined feature detectors based on geodesic curves with the fitting of a deformable model to assign dense corresponding points to unseen faces. The advantages of these methods are their reliability and robustness, which make them ideally suited for the automated processing of large databases. However, the accuracy of the established correspondence is limited, mainly because of two reasons: (1) The constraints derived from the prior knowledge are not flexible enough to match individual features, and (2) most approaches only use the facial geometry for matching.

Alternatively, the problem of face matching can be casted into an image registration task. Using the continuous parametrization of the target and the template surfaces, their features can be commonly mapped into the plane (see Fig. 3). In [25], annotated surface patches were matched to a template by mapping both to the unit circle. Via the common parametrization, dense correspondence was established to build a statistical shape model of anatomical structures that was successfully applied in medical image processing [26, 27]. Additionally, methods like Optical Flow can be used to improve dense correspondence between the flattened photographic textures. As appearance varies heavily between individuals, the method of [28, 29] applied a smoothing filter to the estimated flow field to obtain valid correspondences. By exploiting the temporal dependency between successive scans recorded with professional 3D video setups, recent work has shown that highly accurate dense correspondence can be established over entire facial performances of an actor [30–32]. Kaiser *et al.* [33] as well as Savran and Sankur [34] proposed variational registration methods employing robust similar-




Fig. 3. Matching of a facial surface S to the reference R : Parametrizations Φ_S and Φ_R are computed and photometric as well as geometric features are mapped to the plane. The dense correspondence mapping $\Psi_{\Phi_S \rightarrow \Phi_R}$ accurately registers photographic and geometric features from S and R .

ity measures on photographic and geometric features. They showed that image registration methods can be used to successfully establish accurate dense correspondence between individuals. However, variational approaches are typically prone to converge to undesired local minima or require additional user interaction, which prevents them from being used in a fully automated processing of large-scale face databases.

3 Challenges and Overview

The estimation of dense correspondences on human faces is particularly challenging, mainly due to the following reasons: (1) global facial morphology significantly varies between individuals, (2) facial expressions cause large intra-individual changes in shape and appearance, and (3) large regions like the cheeks or the forehead provide little information on correspondence between individuals.

To build a fully automated method for accurate dense correspondence estimation on human faces, we propose a novel pipeline that addresses these challenges by combining the reliability of methods using prior knowledge with the accuracy of variational matching based on image registration (see Fig. 2). Our key contribution can be divided into two subsequent processing stages:

1. Given a photographically textured raw 3D surface S , we estimate reliable initial correspondence using 2D facial landmark detectors and non-rigid fitting of a 3D SSAM to S . The initial correspondence estimates are used to compute a parametrization for the surface $\Phi_S : S \subset \mathbb{R}^3 \mapsto \mathbb{R}^2$ that maps features into the plane as reliable initial values for variational correspondence matching.
2. We employ an image registration approach to estimate a mapping $\Psi_{\Phi_S \rightarrow \Phi_R} : \mathbb{R}^2 \mapsto \mathbb{R}^2$ that optimizes dense correspondence accurately by matching individual photographic and geometric features from Φ_S to a reference template Φ_R which we call *Unified Facial Parameter Domain* (UFPD) (see Fig. 3).

By using prior information to compute Φ_S , our framework accounts for challenges (1) and (2). The combination of photometric and geometric features with reasonable constraints which penalize non-isometric deformations during image




Fig. 4. Left: Our prototypical stereophotogrammetric setup. Right: Facial landmarks detected in a frontal view using implementations of [35] and [11]. The set of landmarks used for SSAM fitting and parametrization are marked in green. (Color figure online)

registration further helps to define correspondence according to (3) and to accurately match intra- and inter-individual features that have roughly been aligned in the first stage.

A central concept of our approach is the definition of the planar UFPD $\subset [0, 1]^2$ (see Subsect. 4.1). We propose the UFPD as the reference template domain Φ_R aggregating all relevant information for robust and reliable optimization of the dense correspondence mapping $\Psi_{\Phi_S \rightarrow \Phi_R}$. Similar to the template provided by face SSMs, we learn significant geometric and photometric features in the UFPD from our high-resolution face database, that serves as the reference during the variational matching stage. Using the inverse of $\Psi_{\Phi_S \rightarrow \Phi_R}$, we are able to transfer a dense corresponding mesh to the surface S via $\Phi_S^{-1} \circ \Psi_{\Phi_S \rightarrow \Phi_R}^{-1} \circ \Phi_R$.

4 Method

This section describes the main parts of our matching framework. The data is acquired with our prototypical stereophotogrammetric setup using eight DSLR cameras (six Nikon D800E, two Nikon D810, 36 MP each) in four stereo-pairs, and two flashes (Elinchrom 1000) arranged in a semicircular arc around a common focal point. We employed the method of Beeler *et al.* [36] for stereo-matching and *Poisson surface reconstruction* [37] to obtain detailed facial surfaces. High-resolution photographic textures are seamlessly composed by *Poisson image editing* [38]. Before describing the processing stages for new facial surfaces in detail, we define the UFPD as follows.

4.1 The Unified Facial Parameter Domain

A key strategy of our framework is a *Unified Facial Parameter Domain* (UFPD), that serves as a flattened facial template during variational matching similar to [28]. As the UFPD $\subset [0, 1]^2$ represents inter- and intra-individually varying faces, we propose to learn significant photometric and geometric facial features from a representative database.




Fig. 5. The *unified facial parameter domain*: average photographic texture with the set of sparse landmarks (a), the same texture overlaid with its corresponding weight map (b) and the curvature where values are mapped to a normalized gray scale ranging from black (min) to white (max) (c).

Initially, a reference parametrization Φ_R of the average face of the *Basel Face Model* (BFM, see [4]) has been computed employing the *QuadCover* method presented in [39] as it minimizes isometric distortion. Using Φ_R , the set of sparse facial landmarks provided by the 2D landmark detectors as described in Subsect. 4.2 is marked on the average face and mapped to the UFPD accordingly. As the BFM only provides low-resolution vertex colors, we have computed an average photographic texture (16 MP resolution) from the high-resolution face database acquired with our own setup by projecting the initial parametrization of the fitted BFM to the facial surfaces (see Subsect. 4.2). The rendered textures were then median-averaged to retain sharp edges and salient features around anatomical structures like the eyes. We use mean curvature as the geometric feature of each surface during surface matching. To avoid unwanted influence of non-corresponding high-frequency features like facial hair or small wrinkles, all surfaces were filtered using Laplacian surface smoothing. According to the generation of the photographic texture, the averaged mean curvature images were mapped to the UFPD (see Fig. 5).

To account for the specific value of the photometric and geometric features in various facial regions during dense correspondence optimization, we have defined weight maps in the UFPD. The photometric texture is particularly informative in the regions around the eyes, eye-brows, and mouth because they clearly separate skin from other anatomical structures. Similarly, the color of the nostrils is highly valuable for matching due to its high contrast to the skin tone. The geometric features are matched on the entire facial surface except the outer hairline, because this region varies heavily between individuals and disrupts the matching procedure.

Together the set of sparse landmarks, the textures and the weight maps define the UFPD as shown in Fig. 5. Note that the particular definition of UFPD is done once in advance and is independent of the proposed approach. In principle, the parametrization of the UFPD is extensible and can simply be adopted to different scenarios. Additional features or weight maps used for dense matching can easily be integrated.




Fig. 6. For model fitting to (a), the SSAM is first rigidly aligned using 3D landmarks (b) and fitted to the data (c). The resulting SSAM instance as shown in (d) roughly matches the facial features (mouth, nose or eyes).

4.2 Initial Estimation of Facial Landmarks

The detection of sparse facial landmarks is done on the frontal view of a face using two state-of-the-art algorithms (see Fig. 4). The method of Kazemi and Sullivan [11] employs cascades of weak learners, which showed to be more accurate in detecting landmarks with respect to individual morphological features and facial expressions. STASM provided by Milborrow and Nicolls [35] fits a statistical shape and appearance model to the image data and appeared to be more robust. Both methods detect 68 and 77 facial landmarks in frontal faces, respectively. We combined a set of well-defined landmarks faithfully predicted by both approaches for further processing (Exo- and Endocanthion, Pronasale, and Cheilion).

Dense correspondence is further estimated using an SSAM fitted to the raw 3D data similar to [7] (see Fig. 6). Since we desire a combination of shape and color information for better alignment of significant structures like the eyes or mouth, we implemented a fitting routine using the BFM. The sparse facial landmarks are used to estimate the initial parameters of the similarity transform aligning the SSAM with the raw data by performing a single ICP step [23]. Starting with the average face of the BFM, new shape and intensity parameters $P = (P_S \in \mathbb{R}^m, P_I \in \mathbb{R}^n)$ are obtained as the maximum-a-posteriori estimates employing a centered isotropic Gaussian prior with hyper-parameter $\sigma = (\sigma_S, \sigma_I)$ according to

$$p(P|C) \sim p(C|P) p(P|0, \sigma), \quad (1)$$

where $C \subset \mathbb{N} \times \mathbb{N} \times \mathbb{R}$ is our set of robust landmarks between the SSAM and the point cloud with an additional weight assigned. For color estimation, correspondence is established by nearest-neighbour lookup, while for shape estimation, points are also matched by similar colors. Data likelihoods $p(C|P)$ are defined as isotropic Gaussians according to their Euclidean distance by

$$p(C|P) = \prod_{(i,j,\beta) \in C} \mathcal{N}(x_i | m_j, \beta), \quad (2)$$




Fig. 7. Initial parametrizations Φ_S for a sad expression computed without any soft-constraints (a), with constraints for the inner vertices V° (b), and additionally combined with facial landmarks L (c). Note that the latter already aligns features to the UFPD nicely (d).

where x_i are positions or colors of the point cloud and m_j of model vertices. Varying point density in both, the BFM as well as the target, introduces a bias into the data likelihood $p(C|P)$ (e.g. high vertex density around the cheeks in the BFM). We therefore determined the correspondence weights β to be the inverse sum of both frequencies, in which each point occurs in the set of correspondences. Correspondences are determined in each optimization step and new parameters for shape and intensity are estimated using the solution of the system of linear equations with Tikhonov regularization (see [40]) equivalent to Eq. 1.

4.3 Computation of the Initial Parametrization Φ_S

Because variational methods are prone to convergence to local minima, we propose a method to estimate the initial parametrization Φ_S , such that it roughly aligns with features of the UFPD. We constrain the computation of Φ_S using the sparse facial landmarks and the reference parametrization Φ_R projected from the fitted SSAM.

The set of sparse facial landmarks on S is used to match the corresponding positions x_i as defined in the UFPD. Nearest vertices on the surface mesh $S = (V, E)$ are determined and a set of labeled correspondences $L = \{(v_i, x_i) \mid v_i \in V, x_i \in \text{UFPD}\}$ is assembled. Using the projected reference parametrization Φ_R , the facial region on S is segmented and non-facial parts that map outside the UFPD are discarded. We fix the boundary of the facial region as defined by Φ_R to its corresponding position in the UFPD. Similarly, the inner vertices are soft-constrained to their positions as defined by Φ_R . Two separate sets are determined by $K^\partial = \{(v_i, y_i) \mid v_i \in V^\partial, y_i \in \text{UFPD}\}$ for the boundary vertices V^∂ of S and K° for the inner vertices $V^\circ = V \setminus V^\partial$.

To compute Φ_S while accounting for the soft constraints defined by the landmarks L and the inner vertices K° as well as the fixed boundary K^∂ of the facial surface, we adopt the approach of convex-combination maps [41]. Here, the mapping of vertices is expressed as a weighted sum of its 1-ring neighbors:

$$u(v_i) = \sum_{j \in N_1(v_i)} \lambda_{ij} u(v_j). \quad (3)$$

To keep geometric distortion minimal, λ_{ij} is calculated as the mean value weight defined in [42] while the boundary vertices are constrained according to K^∂ . By rewriting Eq. 3 as a linear least squares problem in the mapped coordinates of the inner vertices $u(V^\circ)$, the soft constraints

$$\frac{\alpha}{|L|} \sum_{(v_i, x_i) \in L} \|u(v_i) - x_i\|^2 + \frac{\beta}{|K^\circ|} \sum_{(v_i, y_i) \in K^\circ} \|u(v_i) - y_i\|^2 \quad (4)$$

can conveniently be added, where α, β are weighting factors accounting for the influence of the soft constraints. The solution of the equivalent sparse system of linear equations in V° gives the desired mapping Φ_S (see Fig. 7).

4.4 Variational Matching for Accurate Dense Correspondence

Dense correspondence of S is improved using a variational approach for surface matching inspired by the work presented in [18, 33, 34]. The initial surface parametrization Φ_S allows us to map arbitrary features of S into the plane. We can then use off-the-shelf image registration frameworks that allow us to combine robust similarity measures with reasonable regularization terms into a common objective for optimization of the correspondence mapping $\Psi_{\Phi_S \rightarrow \Phi_R}$.

To measure the similarity between photographic and geometric features, we use two data terms accordingly defined to the weight maps in the UFPD (see Fig. 5b and c). Several image metrics have been investigated (*e.g.* sum of squared differences, flavors of mutual information, gradient metrics) and we found an advanced version of Normalized Cross Correlation (NCC, see [43]) to be well suited for our purpose. As a correlation measure, the advantage of NCC is its robustness to changes *e.g.* in lightning or exposure as well as individual facial characteristics like skin tone that vary significantly with respect to the UFPD.

To regularize the correspondence mapping $\Psi_{\Phi_S \rightarrow \Phi_R}$ in regions with less significant facial features, we use a regularization term similar to [34]. This term, called orthonormality criterion P_{oc} as defined by Eq. (7) in [44], employs the Green-Lagrange strain to measure isometric distortion. Additionally, local foldings of $\Psi_{\Phi_S \rightarrow \Phi_R}$ are avoided by the bending energy P_{be} as defined by Klein and Staring [43].

We discretized the correspondence mapping $\Psi_{\Phi_S \rightarrow \Phi_R} \in F$, where F is the space of cubic B-spline transformations with 128^2 basis functions located on a uniform regular grid covering the UFPD. The objective used to solve the image registration therefore becomes:

$$\begin{aligned} O(u) = & \int_{\text{UFPD}} w_{be} P_{be}(u, x) + w_{oc} P_{oc}(u, x) dx \\ & + \int_{\text{UFPD}} m_P(x) NCC(P_S(u(x)), P_R(x)) dx \\ & + \int_{\text{UFPD}} m_G(x) NCC(G_S(u(x)), G_R(x)) dx, \end{aligned} \quad (5)$$




Fig. 8. Close-up of the photometric (eyes) as well as the geometric features (nose area) before (upper row), and after (lower row) dense correspondence has been optimized. UFPD and the individual features are overlayed using a chessboard pattern as indicated by S, R . Note the accurate correspondence between characteristic morphological structures.

where $m_P(x), m_G(x)$ are weight maps of the photometric (P_S, P_R) and the geometric (G_S, G_R) features from S and the UFPD. The optimization of Eq. 5 was implemented using *elastix*, a framework for rigid and non-rigid image registration [45]. A multi-scale approach for both, the discretization of the correspondence mapping and image resolution is used during optimization. We employ quasi-Newton L-BFGS optimizer including line-search for faster convergence.

5 Experiments and Results

We built a database consisting of 400 facial surfaces acquired with our stereophotogrammetric setup as described in the beginning of Sect. 4. We tested the proposed framework on our database because it contains highly detailed reconstructions including high-resolution photographic textures comparable to 3D models acquired with state-of-the-art stereophotogrammetric devices. We refer the reader to the supplementary material provided with this paper for a collection of representative surfaces from our database.

We also run extensive experiments on all 2500 cases of the BU3DFE database [2]. This database contains 3D models of 100 persons varying in sex, age and ethnicity. The faces are captured in neutral position as well as 6 basic emotions of the *Facial Action Coding System* [46] in 4 levels of intensity. An initial surface reconstruction using [37] was done to close holes or remove meshing artifacts frequently contained in the raw data (*e.g.* below the chin). All data was processed in a fully automatic fashion.

During initial correspondence estimation in the first stage of our framework, sparse facial landmarks were reliably located at the expected positions in the 2D images. In some cases, especially in presence of extreme expressions or when the camera perspective significantly differs from the frontal view, facial landmark detection was less accurate. However, the combination of landmarks




Fig. 9. Results for several expressions from the BU3DFE. Note the accurate dense correspondence established over the entire surface. The red dots indicate the location of landmarks used for quantitative evaluation. (Color figure online)

from both detectors is generally reliable and serves as valuable information in further processing. During fitting of the SSAM, the incorporation of color information improves the registration to structures like the mouth, eyes, and eyebrows where geometric information is less significant. Unfortunately, the BFM is built from a dataset containing neutral expressions only and it fails to adapt to strong variations in shape of the mouth or the eyes. To avoid implausible results in case of expressions, we gave high weight to the prior distribution in shape space and chose $\sigma = (10, 1)$ in Eq. 1.

For the same reason, the initial surface parametrization Φ_S was computed with higher weight given to the set of landmarks L than to the inner vertices V° with $\alpha = 1, \beta = 0.001$. Theoretically, the constraints that we have added to Eq. 4 might lead to non-injective parametrizations (see [41] for a discussion). In practice, we did not find any cases where this occurred. The initial parametrizations obtained are roughly aligned with the UFPD and served as suitable starting values for the optimization of dense correspondence (see Fig. 7).

The dense correspondence mapping accurately registers photographic and geometric features with the UFPD (see Fig. 8). We fixed weights $w_{be} = 150$ and $w_{oc} = 2$ to ensure bijectivity of $\Psi_{\Phi_S \rightarrow \Phi_R}$ and to reasonably constrain matching in regions with less significant features. To quantitatively evaluate our approach, we measured the deviation of landmarks distributed with the BU3DFE. Corresponding landmarks were defined in the UFPD and identified accordingly on the original surfaces after matching. Landmark-wise Euclidean distance was computed and averaged (see Table 1). Using the proposed framework, we were able to predict the landmarks with higher accuracy than previous approaches (except Pronasale in [13] where about 200 cases have been discarded). Moreover as depicted by the standard deviations, the prediction-uncertainty has been significantly reduced.

Table 1. Localization error on the BU3DFE database (for landmarks see Fig. 9). The improvement with respect to the best result from previous work is reported in the last column.

	Segundo <i>et al.</i> [13]		Salazar <i>et al.</i> [14]		Gilani <i>et al.</i> [15]		This paper		
	Mean	SD	Mean	SD	Mean	SD	Mean	SD	impr.
Ex(L)	-	-	9.63	6.12	4.42	2.74	2.95	1.93	33.3 %
En(L)	6.33	4.82	6.75	4.54	4.75	2.64	3.04	1.75	36 %
Ex(R)	-	-	8.49	5.82	4.35	2.70	3.22	2.18	26.0 %
En(R)	6.33	5.04	6.14	4.21	3.29	2.67	3.23	1.86	1.80 %
Sn	-	-	-	-	3.90	3.26	1.97	1.06	49.49 %
Prn	1.87	1.12	5.87	2.70	2.91	2.03	2.05	1.21	-9.63 %
Sbal(L)	-	-	-	-	4.86	2.80	2.37	1.37	51.23 %
Sbal(R)	-	-	-	-	3.57	2.59	2.47	1.29	30.81 %

In fact, using the surface mapping established with our approach, we are able to predict any number of landmarks or mesh vertices that are identified in the UFPD. Here, we used a low-level reference mesh of about 15k vertices as it is sufficient for the resolution available in BU3DFE. We have segmented facial regions in the UFPD and generated a color-coded texture overlaid with a chessboard pattern. The result for several facial expression scans of a single individual is shown in Fig. 9.

Finally, the reference mesh was transferred to all surfaces of the BU3DFE. To demonstrate the suitability of our approach for morphological analysis and generation of statistical face models, we have build two SSAMs using Principal Component Analysis (PCA) on the vertex positions and the photographic textures of the 3D face models. The first SSAM contains the geometric variation related to the inter-individual morphology using the neutral scans only. The second model captures the intra-individual variations due to facial expressions. We have simply computed displacement vector fields for the expressions with respect to the neutral scan of each subject and applied it to the average face of the first SSAM using vertex correspondence. Note the morphological variation captured by the shape parameters in Fig. 10. The chessboard pattern is accurately morphed when the shape varies.

6 Limitations and Future Work

In some rare cases of extreme expressions, we found that matching in the mouth and the forehead region is disturbed by folds, *e.g.* by matching them to other features like the eyebrows. Special detectors could be used to remove these features from textures. Similar strategies could be integrated to handle severe changes in surface area/topology by an open mouth or closed eyes. In the future, we will use the BU3DFE-SSAM instead of the BFM because it already contains several expressions and thus better adapts to an individual morphology. A Riemannian variant of the BU3DFE-SSAM will be established, as non-linear shape




Fig. 10. The BU3DFE-SSM. Left: the face shape according to $\pm 2SD$ of the first three shape parameters of the neutral model. Middle: The average neutral face. Right: The first three shape parameters ($\pm 2SD$) of the expression model.

spaces have been shown to be superior to PCA based models *e.g.* for learning relationships between shape and expressions.

The variational matching in the plane comes at the cost of geometric distortions introduced by the parametrization. As the proposed framework is independent of the actual UFPD, improved definitions will be investigated in further experiments. Similarly, we aim in learning the parameters used in our framework from an annotated ground truth database to further improve accuracy and robustness of automated data processing. The run time of the framework highly depends on the resolution of the input data. In our experiments, we measured times between 0.5 and 3 min on a standard workstation without optimizing our code. We believe that the computation time could be significantly reduced if certain routines are implemented more efficiently and by employing computational parallelism.

7 Conclusions

We have presented a framework for the fully automated determination of highly-accurate dense correspondence for facial surfaces. We showed that the proposed approach works well on a wide range of textured 3D face models varying inter- and intra-individually. Our approach outperforms state-of-the-art methods as confirmed by our experiments. To the best of our knowledge, no SSAM of faces based on a variety of facial expressions with dense correspondence has been released to the research community yet. We are aiming to publish the BU3DFE-SSAM and believe, that this model including geometric as well as photometric variation will help researchers to understand the complex nature of facial morphology. The proposed framework will help to build the next generation of highly-detailed 3D face models on a large scale basis and thus opens up new directions for applications in computer vision and computer graphics.

Acknowledgments. Special thanks go to Olaf Hellwich, Gabriel Le Roux, The Anh Pham, Sven-Kristofer Pilz, Honglei Wang, and Martin Zänker for their valuable contribution and support. Our research is funded by the *Image Knowledge Gestaltung. Cluster of Excellence at the Humboldt-Universität zu Berlin*, with financial support from the German Research Foundation as a part of the Excellence Initiative.

References

1. Brunton, A., Salazar, A., Bolkart, T., Wuhler, S.: Review of statistical shape spaces for 3D data with comparative analysis for human faces. *Comput. Vis. Image Underst.* **128**, 1–17 (2014). Elsevier
2. Yin, L., Wei, X., Sun, Y., Wang, J., Rosato, M.J.: A 3D facial expression database for facial behavior research. In: IEEE International Conference Automatic Face and Gesture Recognition (FGR), pp. 211–216 (2006)
3. Savran, A., Alyüz, N., Dibeklioglu, H., Çeliktutan, O., Gökberk, B., Sankur, B., Akarun, L.: Bosphorus database for 3D face analysis. In: Schouten, B., Juul, N.C., Drygajlo, A., Tistarelli, M. (eds.) BIOD 2008. LNCS, vol. 5372, pp. 47–56. Springer, Heidelberg (2008)
4. Paysan, P., Knothe, R., Amberg, B., Romdhani, S., Vetter, T.: A 3D face model for pose and illumination invariant face recognition. In: IEEE International Conference On Advanced Video and Signal Based Surveillance (AVSS), pp. 296–301. IEEE (2009)
5. Booth, J., Roussos, A., Zafeiriou, S., Ponniah, A., Dunaway, D.: A 3D morphable model learnt from 10,000 faces. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2016)
6. Brunton, A., Bolkart, T., Wuhler, S.: Multilinear wavelets: a statistical shape space for human faces. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) ECCV 2014, Part I. LNCS, vol. 8689, pp. 297–312. Springer, Heidelberg (2014)
7. Cao, C., Weng, Y., Zhou, S., Tong, Y., Zhou, K.: Facewarehouse: a 3D facial expression database for visual computing. *IEEE Trans. Vis. Comput. Graph.* **20**, 413–425 (2014). IEEE
8. Zafeiriou, S., Zhang, C., Zhang, Z.: A survey on face detection in the wild: past, present and future. *Comput. Vis. Image Underst.* **138**, 1–24 (2015)
9. Cootes, T.F., Ionita, M.C., Lindner, C., Sauer, P.: Robust and accurate shape model fitting using random forest regression voting. In: Fitzgibbon, A., Lazebnik, S., Perona, P., Sato, Y., Schmid, C. (eds.) ECCV 2012, Part VII. LNCS, vol. 7578, pp. 278–291. Springer, Heidelberg (2012)
10. Ren, S., Cao, X., Wei, Y., Sun, J.: Face alignment at 3000 fps via regressing local binary features. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 1685–1692 (2014)
11. Kazemi, V., Sullivan, J.: One millisecond face alignment with an ensemble of regression trees. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 1867–1874 (2014)
12. de la Torre, F., Chu, W.S., Xiong, X., Vicente, F., Ding, X., Cohn, J.: Intraface. In: IEEE International Conference Automatic Face and Gesture Recognition (FGR), vol. 1, pp. 1–8. IEEE (2015)
13. Segundo, M.P., Silva, L., Bellon, O.R.P., Queirolo, C.: Automatic face segmentation and facial landmark detection in range images. *IEEE Trans. Syst. Man Cybern. B: Cybern.* **40**, 1319–1330 (2010). IEEE

14. Salazar, A., Wuhrer, S., Shu, C., Prieto, F.: Fully automatic expression-invariant face correspondence. *Mach. Vis. Appl.* **25**, 859–879 (2014). Springer
15. Gilani, Z.S., Shafait, F., Mian, A.: Shape-based automatic detection of a large number of 3D facial landmarks. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4639–4648 (2015)
16. Matthews, I., Baker, S.: Active appearance models revisited. *Int. J. Comput. Vis.* **60**(2), 135–164 (2004)
17. Theobald, B.J., Matthews, I., Mangini, M., Spies, J.R., Brick, T.R., Cohn, J.F., Boker, S.M.: Mapping and manipulating facial expression. *Lang. Speech* **52**(2–3), 369–386 (2009)
18. Litke, N., Droske, M., Rumpf, M., Schröder, P.: An image processing approach to surface matching. *Symp. Geom. Process. (SGP)* **255**, 207–216 (2005)
19. Bronstein, A.M., Bronstein, M.M., Kimmel, R.: Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *Proc. Natl Acad. Sci.* **103**, 1168–1172 (2006)
20. Windheuser, T., Schlickewei, U., Schmidt, F.R., Cremers, D.: Geometrically consistent elastic matching of 3D shapes: a linear programming solution. In: *2011 International Conference on Computer Vision (ICCV)*, pp. 2134–2141. IEEE (2011)
21. Dubrovina, A., Kimmel, R.: Approximately isometric shape correspondence by matching pointwise spectral features and global geodesic structures. In: *Advances in Adaptive Data Analysis*, vol. 3, pp. 203–228. World Scientific (2011)
22. Van Kaick, O., Zhang, H., Hamarneh, G., Cohen-Or, D.: A survey on shape correspondence. *Comput. Graph. Forum* **30**(6), 1681–1707 (2011)
23. Besl, P.J., McKay, N.D.: Method for registration of 3-D shapes. In: *Robotics-DL tentative*, International Society for Optics and Photonics, pp. 586–606 (1992)
24. Myronenko, A., Song, X.: Point set registration: coherent point drift. *IEEE Trans. Pattern Anal. Mach. Intell.* **32**, 2262–2275 (2010)
25. Lamecker, H., Seeba, M., Hege, H.C., Deuhard, P.: A 3D statistical shape model of the pelvic bone for segmentation. In: *Proceedings SPIE Medical Imaging*, vol. 5370, pp. 1341–1351 (2004)
26. Duy, N.T., Lamecker, H., Kainmueller, D., Zachow, S.: Automatic detection and classification of teeth in CT data. In: Ayache, N., Delingette, H., Golland, P., Mori, K. (eds.) *MICCAI 2012*. LNCS, vol. 7510, pp. 609–616. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-33415-3_75](https://doi.org/10.1007/978-3-642-33415-3_75)
27. Ehlke, M., Ramm, H., Lamecker, H., Hege, H.C., Zachow, S.: Fast generation of virtual x-ray images for reconstruction of 3D anatomy. *IEEE Trans. Vis. Comput. Graph.* **19**(12), 2673–2682 (2013)
28. Blanz, V., Vetter, T.: Face recognition based on fitting a 3D morphable model. *IEEE Trans. Pattern Anal. Mach. Intell.* **25**, 1063–1074 (2003). IEEE
29. Blanz, V., Vetter, T.: A morphable model for the synthesis of 3D faces. In: *Proceedings of the Annual Conference on Computer Graphics and Interactive Techniques*, pp. 187–194. ACM Press/Addison-Wesley Publishing Co. (1999)
30. Bradley, D., Heidrich, W., Popa, T., Sheffer, A.: High resolution passive facial performance capture. *ACM Trans. Graph. (TOG)* **29**, 1–10 (2010). ACM
31. Beeler, T., Hahn, F., Bradley, D., Bickel, B., Beardsley, P., Gotsman, C., Sumner, R.W., Gross, M.: High-quality passive facial performance capture using anchor frames. *ACM Trans. Graph. (TOG)* **30**, 75–85 (2011). ACM
32. Fyffe, G., Jones, A., Alexander, O., Ichikari, R., Debevec, P.: Driving high-resolution facial scans with video performance capture. *ACM Trans. Graph. (TOG)* **34**, 8 (2014). ACM

33. Kaiser, M., Störmer, A., Arsić, D., Rigoll, G.: Non-rigid registration of 3D facial surfaces with robust outlier detection. In: IEEE Winter Conference on Applications of Computer Vision (WACV), pp. 1–6 (2009)
34. Savran, A., Sankur, B.: Non-rigid registration of 3D surfaces by deformable 2D triangular meshes. In: IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), pp. 1–6 (2008)
35. Milborrow, S., Nicolls, F.: Active shape models with SIFT descriptors and MARS. In: International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISAPP), pp. 380–387 (2014)
36. Beeler, T., Bickel, B., Beardsley, P., Sumner, B., Gross, M.: High-quality single-shot capture of facial geometry. *ACM Trans. Graph. (TOG)* **29**, 41–50 (2010). ACM
37. Kazhdan, M., Bolitho, M., Hoppe, H.: Poisson surface reconstruction. In: Proceedings of the Fourth Eurographics Symposium on Geometry Processing, vol. 7, pp. 61–70 (2006)
38. Pérez, P., Gangnet, M., Blake, A.: Poisson image editing. *ACM Trans. Graph. (TOG)* **22**, 313–318 (2003). ACM
39. Kälberer, F., Nieser, M., Polthier, K.: Quadcover-surface parameterization using branched coverings. *Comput. Graph. Forum* **26**, 375–384 (2007). Wiley
40. Bishop, C.M.: *Pattern Recognition and Machine Learning*. Springer, New York (2006)
41. Floater, M.S.: Parametrization and smooth approximation of surface triangulations. *Comput. Aided Geom. Des.* **14**, 231–250 (1997). Elsevier
42. Floater, M.S.: Mean value coordinates. *Comput. Aided Geom. Des.* **20**, 19–27 (2003). Elsevier
43. Klein, S., Staring, M.: *elastix the manual v4.7*. Technical report, Image Sciences Institute, University Medical Center, Utrecht (2014)
44. Staring, M., Klein, S., Pluim, J.P.: A rigidity penalty term for nonrigid registration. In: *Medical Physics*, vol. 34, pp. 4098–4108. American Association of Physicists in Medicine (2007)
45. Klein, S., Staring, M., Murphy, K., Viergever, M., Pluim, J.P., et al.: *elastix: a toolbox for intensity-based medical image registration*. *IEEE Trans. Med. Imaging* **29**, 196–205 (2010). IEEE
46. Ekman, P., Friesen, W.V.: *Unmasking the face: a guide to recognizing emotions from facial clues* (2003)