

Matching 3D Faces with Partial Data

Prathap Nair and Andrea Cavallaro
Queen Mary University of London
Multimedia and Vision Group
London, E1 4NS, UK
{prathap.nair, andrea.cavallaro}@elec.qmul.ac.uk

Abstract

We present a novel approach to matching 3D faces with expressions, deformations and outliers. The matching is performed through an accurate and robust algorithm for registering face meshes. The registration algorithm incorporates prior anthropometric knowledge through the use of suitable landmarks and regions of the face to be used with the Iterative Closest Point (ICP) registration algorithm. The localization of landmarks and regions is achieved through the fitting of a 3D Point Distribution Model (PDM) and is independent of texture, pose and orientation information. We show that the use of expression-invariant facial regions for registration and similarity estimation outperforms the use of the entire face region. Evaluation is performed on the challenging GavabDB database and we achieve 93.7% rank-1 recognition with an overall retrieval accuracy of 91.1%.

1 Introduction

The matching of 3D face meshes requires accurate comparison of surface properties from different meshes. This task can be impaired when meshes have deformations (due to expressions or medical conditions) or outliers (due to the acquisition process). A common approach to match 3D meshes is through the Iterative Closest Point (ICP) algorithm [3] for rigid registration. As ICP is based on the closest point associations from one mesh to the other, in the presence of deformations and outliers its performance degrades due to global minima. Moreover, ICP requires roughly aligned meshes in order to converge in terms of the mean square error (MSE). Many variants to the original ICP algorithm have been proposed to improve speed and convergence [2, 22], but without removing the above mentioned limitations. Other global registration methods exist [8, 9, 11], some of which use the ICP, but are also inappropriate in the presence of deformations and outliers.

The localization of specific anthropometric landmarks and regions on face meshes often plays an important part in these applications. Landmarks can aid the ICP algorithm in achieving rough alignment of meshes, and by themselves provide valuable semantic information. In biometric applications, landmarks are often instrumental in the generation of signatures for faces [19] and isolation of expression invariant regions for matching [4]. However, the dependence on prior knowledge of feature map thresholds, orientation and pose is evident in most existing methods for landmark localization on meshes [4, 5, 14]. The segmentation of faces is also important prior to analysis when the mesh includes other

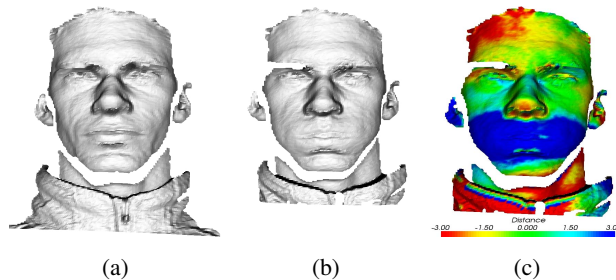


Figure 1: Example highlighting limitations of the conventional ICP approach: (a-b) Sample person scans with different expressions and amount of non-facial regions; (c) result of distance estimation between the meshes post-registration, showing inaccurate regions of change and introduction of errors due to non-facial regions (red and blue regions).

body parts and outliers. In [12] an ICP based 3D face recognition approach is presented, where the face is first detected and segmented prior to registration. The face detection is performed by finding the nose tip and segmentation is done through a cropping sphere centred at the nose tip. This approach is highly restrictive to the database used as each input mesh is assumed to contain only one frontal face. Moreover, the cropping sphere has a fixed radius and hence the segmentation is not robust to scale variance.

In this paper, we propose a robust algorithm to perform effective rigid registration of face meshes. We localise landmarks on the face based on the fitting of a Point Distribution Model (PDM) [17] which allows us to work independently of pose and orientation information, and eliminates the need of a texture map. Face segmentation is then performed with a cropping sphere that is scaled and positioned using the localised landmarks. The 3D face registration algorithm then incorporates prior anthropometric knowledge in achieving fine registration of meshes. The prior knowledge in the registration is included using suitable landmarks and regions of the face to be used in the ICP registration.

2 Problem definition

We represent a 3D face dataset as $\Phi = \{(\Phi_{N_1}^1, C^1), (\Phi_{N_2}^2, C^2), \dots, (\Phi_{N_M}^M, C^M)\}$, where $\Phi_{N_i}^i = \{\Psi_1^i, \Psi_2^i, \dots, \Psi_{N_i}^i\}$ are the N_i faces of a person C^i with M people in the dataset. Then given any face $\Psi_j^i \in \Phi$ we aim to retrieve the identity and set of faces $(\Phi_{N_i}^i, C^i)$ of the person, using only geometric information while being independent of texture, orientation and pose information.

Figure 1 (a-b) shows two face meshes of a person with different expressions and amount of *non-facial* regions captured. Non-facial regions here refer to other body parts and clothing. We applied the conventional ICP algorithm using vertices sampled from the entire mesh and computed the distance between the meshes post-registration (a detailed explanation of the distance estimation algorithm is presented in Sec. 5). The results (Fig. 1 (c)) show that the presence of deformations (due to expression change) and non-facial regions causes a poor registration. Inaccurate regions of change can be noticed on the forehead, with a large amount of error introduced by the non-facial regions. We aim to overcome these limitations of the conventional ICP approach as described below.

3 Facial model construction

We use a model based on a Point Distribution Model (PDM) [6] to detect landmarks and segment the face and facial regions for the registration algorithm. The PDM is a parameterized model, $\mathbf{\Omega} = \Upsilon(\mathbf{b})$, where $\mathbf{\Omega} = \{\omega_1, \omega_2, \dots, \omega_N\}$, with $\omega_i = (x_i, y_i, z_i)$ representing each landmark. The vector \mathbf{b} holds the parameters which can be used to vary the shape and Υ defines the function over the parameters. To this end, a training set of L samples with N landmarks representing the region of interest are used. We then have L training shapes where each k^{th} shape is a $3 \times N$ element vector, $\mathbf{\Omega}^k = \{\omega_1^k, \omega_2^k, \dots, \omega_N^k\}$.

Training shapes are aligned and scaled to eliminate global transformations, so that statistical analysis is carried out only on the shape variations. Procrustes analysis [10] is used to align the training shapes to their mutual mean in a least-squares sense via similarity transformations. This minimises \mathbf{D} , which is the sum of distances of each shape $\mathbf{\Omega}^k$ to the mean $\bar{\mathbf{\Omega}} = \frac{1}{L} \sum_{k=1}^L \mathbf{\Omega}^k$, i.e., $D = \sum_{i=1}^N |\omega_i^k - \bar{\omega}_i|^2$. At each iteration, $\bar{\mathbf{\Omega}}$ is scaled such that $|\bar{\mathbf{\Omega}}| = 1$. Using PCA, the variations of the shape cloud formed by the training shapes in the $(3LN)$ - dimensional space are estimated along the principal axes. The principal axes and corresponding variations are represented by the eigenvectors and eigenvalues obtained from the covariance \mathbf{Z} of the data, computed using

$$\mathbf{Z} = \frac{1}{L-1} \sum_{k=1}^L (\mathbf{\Omega}^k - \bar{\mathbf{\Omega}})(\mathbf{\Omega}^k - \bar{\mathbf{\Omega}})^T. \quad (1)$$

If ϕ contain the t eigenvectors corresponding to the largest eigenvalues, then any shape similar to those in the training set can be approximated using

$$\mathbf{\Omega} \approx \bar{\mathbf{\Omega}} + \phi \mathbf{b}, \quad (2)$$

where $\phi = (\phi_1 | \phi_2 | \dots | \phi_t)$ and \mathbf{b} is a t dimensional vector given by $b = \phi^T (\mathbf{\Omega} - \bar{\mathbf{\Omega}})$. The value of t is chosen such that the model represents 98% of the shape variance, ignoring the rest as noise [6]. The mean shape is obtained when all parameters are set to zero.

4 Model fitting

The PDM $\mathbf{\Omega}$ is fitted onto a new mesh Ψ^i by performing similarity transformations of the model, estimated using three control points of the mean shape $\bar{\mathbf{\Omega}}$, which are the inner eye points (ω_r and ω_l) and the nose tip point (ω_n), with $\{\omega_r, \omega_l, \omega_n\} \in \bar{\mathbf{\Omega}}$. We isolate candidate vertices on a face mesh using curvature-based feature maps and select the inner eye and nose tip areas as they can be robustly isolated from other vertices. In this section we will discuss the feature extraction and preprocessing to isolate the candidate vertices. In order to characterize the curvature property of each vertex on the face mesh we compute two feature maps, namely the *shape index* and the *curvedness index* [7]. These maps are derived based on the principal curvature values, $\kappa_1(\cdot)$ and $\kappa_2(\cdot)$, at all the vertices of the mesh using differential geometry. The shape index, ρ , at a vertex v_i , is defined as

$$\rho(v_i) = \frac{1}{2} - \frac{1}{\pi} \tan^{-1} \left(\frac{\kappa_1(v_i) + \kappa_2(v_i)}{\kappa_1(v_i) - \kappa_2(v_i)} \right), \quad (3)$$

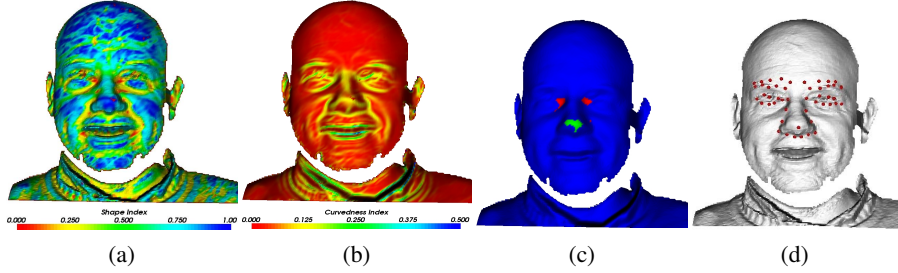


Figure 2: Example scan showing feature maps used to isolate candidate vertices and example of model fitting: (a) shape index, (b) curvedness index, (c) candidate vertices (regions in green are nose tip vertices and regions in red are eye vertices), (d) fitted model.

where $\kappa_1(v_i) \geq \kappa_2(v_i)$ and $\rho(\cdot) \in [0, 1]$. While $\rho(\cdot)$ can describe subtle shape variations from concave to convex thus providing a continuous scale between salient shapes, it does not give an indication of the scale of curvature. For this reason the curvedness of a surface is introduced. The curvedness of a surface, $\gamma(\cdot)$, at a vertex v_i , is defined as

$$\gamma(v_i) = \frac{\sqrt{\kappa_1^2(v_i) + \kappa_2^2(v_i)}}{2}. \quad (4)$$

We compute the low-level feature maps of Eq. 3 and Eq. 4 after Laplacian smoothing to reduce outliers arising from the scanning process.

To reduce the computational overhead through the reduction of outlier candidate vertices, the original mesh is first decimated [18]. Then the feature maps are averaged across vertex neighbors according to

$$\tilde{\rho}(v_i) = \frac{1}{P} \sum_{p \in \mathcal{P}(v_i)} \rho(v_p), \quad \tilde{\gamma}(v_i) = \frac{1}{P} \sum_{p \in \mathcal{P}(v_i)} \gamma(v_p), \quad (5)$$

where $\mathcal{P}(v_i)$ is the set of P neighboring vertices of v_i .

If $\tilde{\gamma}(\cdot) > \gamma_s$, then v_i is in a salient high-curvature region. The condition $\tilde{\rho}(\cdot) < \rho_e$ identifies concave regions; while $\tilde{\rho}(\cdot) > \rho_n$ identifies convex regions. We can therefore relax thresholds to segregate candidate inner eye vertices from the nose tip ones (Fig. 2 (c)). The thresholds $\gamma_s = 0.1$, $\rho_e = 0.3$ and $\rho_n = 0.7$ were found to be adequate for the entire database, with second-order neighborhoods for feature averaging and a decimation of 80%. A further reduction of outlier candidate combinations is performed by checking the triangle formed by each combination of 2 inner eye and 1 nose tip vertices. A plausible inner eye-nose triangle should be acute angled with the squared length of each side being smaller than the sum of the squares of the two other sides.

To transform the model Ω we use as target points plausible combinations of the candidate inner eye vertices and candidate nose tip vertices on Ψ^i . Next the remaining points of Ω are moved to the closest vertices on Ψ^i , Ω is then projected back into the model space and the parameters \mathbf{b} are updated. Based on this exhaustive search over the isolated candidate vertices, the transformation exhibiting the minimum deviation from the mean shape is chosen as the fit for the model (Fig. 2 (d)).




Figure 3: Face segmentation results on the subject scans from Fig. 1, showing effective removal of outlier body parts and clothing.

5 Face registration and similarity calculation

Prior to registration, the face needs to be segmented from other outlier body parts or clothing. *Segmentation* is performed by placing a sphere of radius r with center at the midpoint of the line joining the nasal bridge and nose tip. The intersection boundary of the face and the sphere is then used to segment the face [12]. To account for different sizes of faces, the value of r is set in proportion to the size of the fitted model. A value of $r = 2.6l$ was found to be appropriate for all scans in our database, where l is the Euclidean distance between the nose tip and nasal bridge landmarks (nose length). Examples of face segmentation results is shown in Fig. 3, corresponding to the segmented faces of the subject shown in Fig. 1.

Face *registration* is based on rigid registration with prior anthropometric knowledge, which utilizes an adaptation of the ICP algorithm. The algorithm uses the detected landmarks from the PDM fitting to first perform a coarse registration. Coarse registration is based on the best fit mapping in a least squares sense. The landmarks are then used to segment specific *stable* regions on the face, which are robust to expressions and facial deformations. These regions are finally used to achieve fine registration.

Stable regions (R_s) include the region around the inner eye points, the nasal bridge between these points and around the eyebrow region (Fig. 4 (a)). The nose region (Fig. 4 (b)) is also relatively stable to most natural expressions. Vertices from these regions are localised using the fitted model and selected for ICP registration. We will refer to this approach as Selective-ICP (S-ICP).

Once the reference scan and the test scan are registered, the next step is to evaluate the distance between the two meshes as a measure of *similarity*. To this end, the symmetric Hausdorff distance [1] is used. Let Ψ and Ψ' be the two facial meshes and $\partial(v_i, \Psi')$ be the distance between a vertex $v_i \in \Psi$ and Ψ' . If we define

$$\partial(v_i, \Psi') = \min_{v'_i \in \Psi'} (\|v_i - v'_i\|), \quad (6)$$

then the Hausdorff distance, $\partial(\Psi, \Psi')$, is given by

$$\partial(\Psi, \Psi') = \sum_{v_i \in \Psi} \max[\partial(v_i, \Psi')], \quad (7)$$

and the symmetric Hausdorff distance, ∂_s , is then given by

$$\partial_s(\Psi, \Psi') = \max[\partial(\Psi, \Psi'), \partial(\Psi', \Psi)]. \quad (8)$$




Figure 4: Regions used for registration and distance estimation: (a) stable regions R_S composed of the region around the inner eye points, the nasal bridge and eyebrow region; (b) nose region; (c) combination of R_S and the nose region; (d) Regions $R_1 - R_6$ used in the distance estimation.

Table 1: Comparison of rank-1 retrieval accuracy on using R_1 , R_2 and R_3 in the similarity estimation, post-ICP registration

Regions	R_1	R_2	R_3	$R_1 + R_2$	$R_1 + R_3$	$R_2 + R_3$	$R_1 + R_2 + R_3$
Rank-1 recognition	74.9%	69.1%	85.2%	78.0%	84.5%	82.7%	85.9%

The Hausdorff distance estimation is based on different regions of the face towards finding the most robust facial regions to be used in the similarity estimation. The face was divided into 6 regions ($R_1 - R_6$), separating the forehead, eyes, nose, cheek, mouth and chin, as shown in Fig. 4 (d). For a given selection of P regions $S_R \in \Psi$ and $S'_R \in \Psi'$, the similarity Δ is defined as,

$$\Delta = \max \left[\sum_{v_i \in S_R} \max [\partial(v_i, \Psi')], \sum_{v'_i \in S'_R} \max [\partial(\Psi', v'_i)] \right]. \quad (9)$$

We first tested the rank-1 recognition accuracy on applying the ICP algorithm over the cropped face and using the 6 facial regions in the similarity estimation. The single region that lead to the worst retrieval is the mouth region (R_5), which is the most affected by variations in expressions. The best results were provided by the forehead, eyes and nose regions (R_1 , R_2 and R_3) and their combinations. Table 1 shows the rank-1 recognition accuracy obtained with the use of these regions.

6 Experimental results

We evaluate the performance of the proposed registration approach for 3 region configurations (Fig. 4 (a-c)), and compare it with the conventional ICP approach. We refer to the configuration with R_S as *S-ICP1*, with R_3 as *S-ICP2* and with the combination of both regions as *S-ICP3*. The GavabDB database [13] is used to evaluate the proposed approach. The dataset consists of 427 face scans, corresponding to 7 scans of 61 subjects, and provides a challenging collection with samples having varying degrees of expressions, poses and presence of other body parts and clothing. For the detection of landmarks on the face meshes, the PDM Ω is generated from the BU-3DFE [21] database and the corresponding




Figure 5: Rank-1 recognition (left) and retrieval accuracy (right) obtained using the forehead (R_1), eyes (R_2) and nose (R_3) regions in the similarity estimation.

ground-truth landmarks provided with the database. We use 48 ground-truth landmarks from the eyebrows, eyes and nose regions and include an additional landmark at the nose-tip. Our training set is composed of landmarks from 150 faces, corresponding to 25 scans from 6 individuals each. The 25 scans corresponds to five expressions (anger, disgust, fear, happiness, sadness and surprise) with four degrees (intensities) of expressions and one neutral expression.

The performance of the approaches is evaluated in terms of recognition and retrieval accuracies. Recognition here refers to the accuracy of the retrieved rank-1 identity, while retrieval refers to accuracy of retrieving faces of the same person with most similarity. The retrieval accuracy is measured using the average dynamic precision (ADP) [20]. The ADP is defined as

$$ADP = \frac{1}{S} \sum_{i=1}^S \frac{T_i}{i}, \quad (10)$$

where T_i is the number of true positives, with $T_i \leq i$, and S is the scope size which refers to the total number of expected true positives. S is set to 7 in our experiments since we have 7 samples per person. For example, if for a given query the retrieved results correspond to [1, 1, 0, 1, 1, 0, 1] until rank-7 (where 1 is a true positive and 0 is a true negative), the $ADP = 1 + 1 + 0.67 + 0.75 + 0.8 + 0.67 + 0.71 = 5.56/7 = 0.794$.

Figure 5 shows the rank-1 recognition and retrieval accuracy in terms of the ADP of the 4 approaches. The overall best results were obtained on using the nose region for both the registration (S-ICP2) and similarity estimation (R_3), with a rank-1 recognition rate of 93.7% and an ADP of 91.1%. For ICP the nose region is the most robust with a recognition rate of 85.2% and an ADP of 83.1%. In the S-ICP1, the region combination of the forehead, eyes and nose ($R_1 + R_2 + R_3$) gives the best result with recognition rate 87.4% and ADP 83.9%. Finally, for S-ICP3 we again achieve most robustness with regions ($R_1 + R_2 + R_3$), with recognition rate 87.8% and ADP 83.9%. We observe a correlation between the regions used in the registration and subsequently in the similarity estimation. In S-ICP2 the nose is used to register the faces and in turn proves to be the most robust in estimating the distances. In S-ICP1 and S-ICP3, parts of the forehead, eyes and nasal bridge is used to register the faces and these regions also proves most robust in the similarity measure. However, recognition and retrieval using R_3 outperformed the

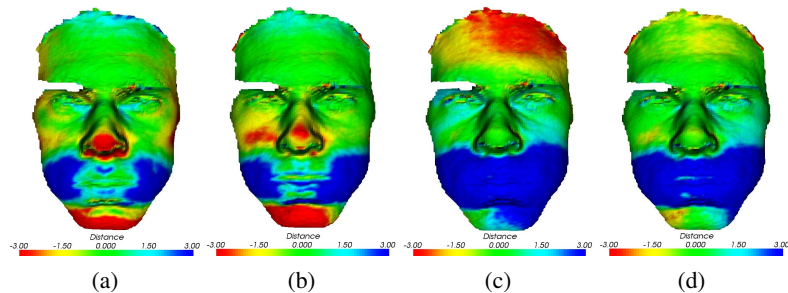


Figure 6: Registration results using various region configurations and comparison with ICP: (a) ICP; (b) S-ICP1; (c) S-ICP2; (d) S-ICP3.

other regions in terms of robustness.

Figure 6 shows example results with estimated distances, post-registration, with the 4 approaches. In contrast to the previously shown results (ICP over the entire mesh), we see here that using ICP on the cropped face (Fig. 6 (c)) provides a better registration. However, erroneous regions of change still exist around the nose, cheek and chin areas. The same is noticed with the S-ICP1 approach. S-ICP2 also results in inaccurate distances estimated around the forehead region in this case. The best registration is achieved through the S-ICP3 approach with regions of change only found around the cheek region which is consistent with the expected deformation. Figure 7 shows the retrieval results with an example query exhibiting an exaggerated expression. The rank-1 retrieved face corresponds to the query in all cases. ICP performs the worst with only 1 true positive within the top-6 retrieved faces. S-ICP2 outperforms the other approaches with all true positives in the top-6 faces.

On the database discussed here, Moreno *et al.* [14, 15] reported 78% rank-1 recognition (in 2003 and 2005), while Mousavi *et al.* [16] report 91% in recent work (in 2008). Our approach, S-ICP2, achieves 93.7% rank-1 recognition. The main failure mode of the S-ICP2 approach occurs in faces with a large amount of noise or holes around the nose region, leading to incorrect PDM fitting and consequently poor registration. Figure 8 shows an example query face which resulted in poor retrieval due to the entire nose region being absent. This limitation could be overcome through hierarchical model fitting with subsets of the PDM and dynamic stable region selection, and is included in our future work.

7 Conclusions

We presented a 3D face retrieval approach based on a coarse-to-fine registration using anthropometric knowledge and facial landmarks localised using a 3D Point Distribution Model (PDM). Facial regions are used in the registration and similarity estimation and we demonstrated that the use of expression invariant regions outperforms use of the entire face. The nose is the most robust in all region configurations. Future work includes applying a weighted combination of distances from different regions to further improve discrimination and hierarchical model fitting with subsets of the PDM in the case of missing key-landmarks. The influence of mesh pre-processing on registration, such as smoothing and hole filling, will also be investigated.




Figure 7: Retrieval results of an example query showing robustness of using the nose region R_3 in the S-ICP2 approach (faces highlighted in green indicate true positives, while red indicate false positives): (a) ICP; (b) S-ICP1; (c) S-ICP2; (d) S-ICP3.

References

- [1] N. Aspert, D. Santa-Cruz, and T. Ebrahimi. MESH: Measuring errors between surfaces using the Hausdorff distance. In *Proc. IEEE International Conference in Multimedia and Expo*, pages 705–708, Lausanne, Switzerland, Aug 2002.
- [2] R. Bergevin, M. Soucy, H. Gagnon, and D. Laurendeau. Towards a general multi-view registration technique. *IEEE Trans. Pattern Anal. Machine Intell.*, 18(5):540–547, May 1996.
- [3] P. Besl and N. McKay. A method for registration of 3D shapes. *IEEE Trans. Pattern Anal. Machine Intell.*, 14:239–256, Feb 1992.
- [4] K. Chang, K. Bowyer, and P. Flynn. Multiple nose region matching for 3D face recognition under varying facial expression. *IEEE Trans. Pattern Anal. Machine Intell.*, 28(10):1695–1700, Oct 2006.
- [5] D. Colbry, G. Stockman, and A. Jain. Detection of anchor points for 3D face verification. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, pages 118–125, New York, NY, Jun 2006.
- [6] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham. Active shape models: their training and application. *Computer Vision and Image Understanding*, 61(1):38–59, Jan 1995.



Figure 8: Failure mode: Example scan showing missing nose region which causes failure of the region based approaches due to incorrect PDM fitting and region localisation.

- [7] C. Dorai and A. K. Jain. Cosmos - a representation scheme for 3D free-form objects. *IEEE Trans. Pattern Anal. Machine Intell.*, 19(10):1115–1130, Oct. 1997.
- [8] A. Fitzgibbon. Robust registration of 2D and 3D point sets. In *Proc. British Machine Vision Conference*, pages 411–420, Manchester, UK, Sept 2001.
- [9] N. Gelfand, N. J. Mitra, L. J. Guibas, and H. Pottmann. Robust global registration. In *Proc. Eurographics Symposium on Geometry Processing*, pages 197–206, Vienna, Austria, Jul 2005.
- [10] C. Goodall. Procrustes methods in the statistical analysis of shape. *Journal of the Royal Statistical Society*, 53(2):285–339, 1991.
- [11] A. Makadia, A. Patterson, and K. Daniilidis. Fully automatic registration of 3D point clouds. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, pages 1297–1304, New York, NY, Jun 2006.
- [12] A. Mian, M. Bennamoun, and R. Owens. An efficient multimodal 2D-3D hybrid approach to automatic face recognition. *IEEE Trans. Pattern Anal. Machine Intell.*, 29(11):1927–1943, Nov 2007.
- [13] A. B. Moreno and A. Sanchez. GavabDB: a 3D face database. In *Proc. Workshop on Biometrics on the Internet COST275*, pages 77–85, Vigo, Spain, Mar 2004.
- [14] A.B. Moreno, A. Sanchez, J.F. Velez, and F.J. Diaz. Face recognition using 3D surface-extracted descriptors. In *Proc. Irish Machine Vision and Image Processing Conference*, Coleraine, Ireland, Sept 2003.
- [15] A.B. Moreno, A. Sanchez, J.F. Velez, and F.J. Diaz. Face recognition using 3D local geometrical features: PCA vs. SVM. In *Proc. 4th International Symposium on Image and Signal Processing and Analysis*, pages 185–190, Zagreb, Croatia, Sept 2005.
- [16] M. H. Mousavi, K. Faez, and A. Asghari. Automatic feature extraction for multiview 3D face recognition. In *Proc. 7th International Conference on Computer and Information Science*, pages 208–213, Oregon, USA, May 2008.
- [17] P. Nair and A. Cavallaro. Region segmentation and feature point extraction on 3D faces using a point distribution model. In *Proc. IEEE International Conference on Image Processing*, Texas, USA, Sept 2007.
- [18] W. J. Schroeder, J. A. Zarge, and W. E. Lorensen. Decimation of triangle meshes. In *Proc. 19th annual conference on Computer graphics and interactive techniques*, pages 65–70, Chicago, Illinois, Aug 1992.
- [19] J. Shi, A. Samal, and D. Marx. How effective are landmarks and their geometry for face recognition? *Computer Vision and Image Understanding*, 102(2):117–133, May 2006.
- [20] R. Typke, R. C. Veltkamp, and F. Wiering. A measure for evaluating retrieval techniques based on partially ordered ground truth lists. In *Proc. International Conference on Multimedia and Expo (ICME)*, pages 1793–1796, Toronto, Canada, Jul 2006.
- [21] L. Yin, X. Wei, Y. Sun, J. Wang, and M. Rosato. A 3D facial expression database for facial behavior research. In *Proc. 7th International Conference on Automatic Face and Gesture Recognition*, pages 211–216, Southampton, UK, Apr 2006.
- [22] Z. Zhang. Iterative point matching for registration of free-form curves and surfaces. *International Journal of Computer Vision*, 13(2):119–152, Oct 1994.