Robust 3D Face Recognition by Using Shape Filtering

Liang Cai Student
cailiang-23@163.com
Feipeng Da Prof
dafp@seu.edu.cn

Research Institute of Automation
Southeast University
Nanjing, China

Abstract

Achieving high accuracy in the presence of expression variation remains one of the most challenging aspects of 3D face recognition. In this paper, we propose a novel recognition approach for robust and efficient matching. The framework is based on shape processing filters that divide face into three components according to its frequency spectral. Low-frequency band mainly corresponds to expression changes. High-frequency band represents noise. Mid-frequency band is selected for expression-invariant feature which contains most of the discriminative personal-specific deformation information. By using shape filter, it offers a dramatic performance improvement for both accuracy and robustness. We conduct extensive experiments on FRGC v2 databases to verify the efficacy of the proposed algorithm, and validate the above claims.

1 Introduction

Face recognition is highly attractive in a wide range of applications including surveillance, access control and machine-human interaction, mainly because of its non-intrusive nature. Due to that 3D facial shape is invariant to illumination and makeup, it is believed that 3D face recognition has the potential for more accuracy. Moreover, with the development of 3D scan technology, the acquisition of 3D shape become more accurate and less intrusive, which allows for 3D facial recognition with low cost by measuring geometry on the face. However, since expression deformations bring difficulty to classifier for differentiating interpersonal disparities, expression is the greatest challenge to real world application.

In this paper, we are interested in the issue of robust 3D face recognition under expression variation. Our aim is to do this without overly constraining the conditions in which facial data is acquired. Most often, this means that the amount of scan data used for enrollment is only one and the test data used for authentication can take various expressions. Motivated by multi-scale deformation, in this paper, we describe a novel framework for accurate recognition under varying expression, based on manifold shape filtering techniques. The main contribution is a frequency division approach for expression challenge. Facial shape is firstly transformed into geometry image based on mesh parametrization. Then mid-frequency band of shape, which contains most of the discriminative personal information, is used for recognition.

BMVC 2010 doi:10.5244/C.24.65

The remainder of this paper is organized as follows: in the next section, we will review relevant previous work. An overview of the usage of shape filters in 3D face recognition is given in Section 2.1. Section 3 describes the methods in detail. The performance of the proposed algorithm is reported in Section 4. Section 5 discusses the advantages and limitations of our approach. The paper is concluded with a summary and an outline of directions for future research in Section 6.

2 Related Work

3D face recognition has been extensively studied and many approaches have been proposed to overcome expression influence. Some recently surveys are given by Bowyer et al. [5], Scheenstra et al. [17] and Zhao et al. [23]. There are two kinds of approaches to handle expression variation, namely, local-based and global-based methods. Local-based methods rigidly match sub-facial regions of least expression deformations. Global-based methods try to extract expression-invariant feature or apply deformations to the 3D facial data to counteract expression deformations for differentiating interpersonal disparities. Despite the fact that high performance can be achieved by using curves or local rigid region on face [12, 15, 16, 20], the global approaches are more promising in handling the expression challenge because it is possible to extract more information from the face. This section reviews global approaches that address expression challenge in two different ways: embedding and deformable model approaches.

Bronstein et al. [6] who demonstrate that facial expressions can be modeled as isometries of the surface propose the embedding approach. Face under various expressions is deformed by embedding it into a multi-dimensional space using Multi-Dimensional Scaling (MDS). This preserves the intrinsic geometries of face and eliminates the extrinsic geometries. Thus, an expression invariant representation (canonical form) of face is obtained, as the facial expressions are mainly extrinsic geometries. Finally, the canonical forms are rigidly matched against each other.

3D morphable model is a well-known approach proposed by Blanz et al. [3] who use statistical deformable model to estimate the 3D face shape. Amberg et al. [2] describe a method for face recognition by statistically fitting this model to shape data. In the work by Al-Osaimi et al. [1], deformations are also learned in a Principal Component Analysis (PCA) subspace, which is called expression deformation model. In the work by Wang et al. [22], the learned deformation fields are applied to deform the non-neutral 3D scans to neutral ones. Thin Plate Spline (TPS) is used by Lu et al. [11] to learn expression deformation from a control group of neutral and non-neutral scans. A common feature of these approaches is that they tend to handle model fitting by statistical method. The advantage is its robustness to noise and missing data. However, statistical method is just linear approximation to non-rigid deformation. Moreover, it requires training on more than one sample, which raises questions on generalizing performance. Hence, non-statistical models are proposed to handle these problems.

Huang et al. [9] propose a global to local deformation framework to deform a shape to a new one of the same class. They also show their framework's applicability to 3D faces. In the work by Kakadiaris et al. [10] and Passalis et al. [14], an Annotated Face Model (AFM) is deformed to scan data. Their approach consists of multistage alignment algorithms for performing model fitting and mesh parametrization for converting shape information to 2D geometry image. However, it does not specify expression-invariant features for indi-

vidual identification. Botsch et al. [4] present a multi-scale deformation approach. They decompose spatial shape into two parts of different frequency bands. The key is that detail information preserves well during face deformation. This idea motivates us to extract expression-invariant features by shape filtering from the view of frequency domain. This improvement presented in the paper allows us to enhance recognition performance under expression variation.

2.1 Overview

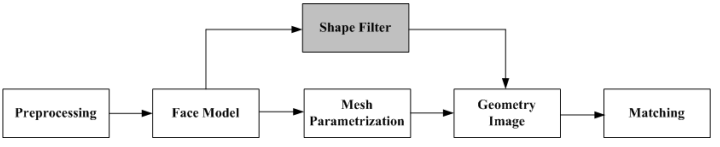

Figure 1: Overview of our proposed shape filtering method

Classical framework of geometry image analysis presented by Kakadiaris et al. [10] and Passalis et al. [14] is shown in Fig.1. Deformable face model is firstly fitted to preprocessing scan data. After that, deformation and normal map are assigned to the model's planar parametrization mesh to obtain their geometry images. Then three channels (X, Y, and Z) of them are analyzed using a wavelet transform and the coefficients are stored for individual features. In our approach, a shape filter is applied to the deformed face model. The filter is constructed by spectral method, which outputs shape signal of mid-frequency band. After that, it is also assigned to the model's parametrization mesh in order to acquire shape-filtered geometry image called normal variation image. As the work by Kakadiaris et al. [10] and Passalis et al. [14], deformation field is combined with the normal variation to recognize individual.

3 The Detail of Our Approach

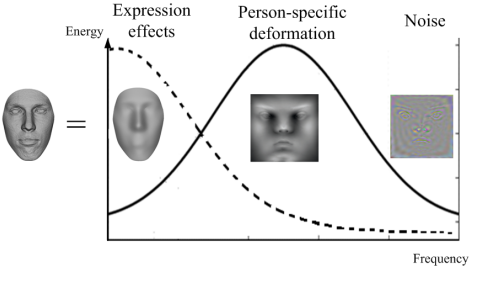

Figure 2: The generative model used for our 3D face recognition

Our main assumption is that each individual face consists of three parts (See Fig. 2). Low-frequency band mainly corresponds to expression changes. Mid-frequency band contains most of the discriminative personal-specific deformation information. High-frequency

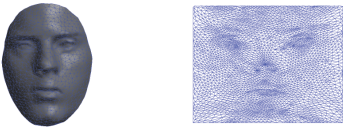

(a) (b)

Figure 3: (a) face mesh (b) its square parametrization on planar domain

band represents noise. The fundamental idea behind our method is to apply shape filter to obtain mid-frequency information, and then convert it into a 2D parametrization space that retains the geometry information. We will give detailed description of our method and perform comparison experiment to verify its efficiency.

3.1 Preprocessing

The purpose of preprocessing is twofold: to crop face area and to unify facial data from scanner.

Face cropping: The input data acquired by 3D scanner has some areas that should be removed in the next deformation extraction, such as shoulder, hair, and ears. Our segmentation process contains basically two steps. Initially, the nose tip is detected according to depth information. Then, the face region of interesting covered by an ellipse is cropped.

Smoothing: As most high-resolution scanners produce noisy data in real-world conditions, another important task in this stage is to smooth the facial surface. There are many smoothing algorithms based on triangulation mesh, which could fulfill this task (e.g. laplacian filters. [7]). Hence, in our approach, we utilize implicit surface method called partition of unity (PU) of implicit surface. [13]. It is a detail-preserving noise robust surface reconstruction algorithm based on diffusion technique. It is also a efficient tool for extrapolating missing data. Since laser scanners usually produce holes or data missing in certain areas (e.g., eyes, eyebrows and chins), we can rely on PU to get regular data.

3.2 Mesh Parametrization

After rigidly registered to face model using Iterative Closest Point (ICP), each individual is represented by the model and its deformation. The deformed model is then converted to geometry image, based on mesh parametrization.

Parametrization is defined as mapping a 3D surface $\mathscr{M}$ onto a suitable planar domain $\mathscr{D} \subset \mathscr{R}^3$. The mapping ψ is a bijection. Applying the similar definition to the discrete deformable model whose vertices are X, we can compute a 2D position $\psi(x_i)$ for all the vertices. More detail on mesh parametrization is given by Floater et al. [8] and Sheffer et al. [18]. In our method, we just rely on linear methods for mesh parametrization. Precisely, $\psi = (\psi_1, \psi_2)$ is the solution to linear equation:

$$\begin{cases} L\psi_1(x_i) = L\psi_2(x_i) = 0, \forall x_i \notin \partial\mathscr{M} \\ \psi(x_i) = \psi^0(x_i) \in \partial\mathscr{D}, \forall x_i \in \partial\mathscr{M} \end{cases} \tag{1}$$

where L is laplacian matrix whose detailed description will be given in next section. $\partial\mathscr{M}$ is

(a) (b)

Figure 4: Shape filtering on manifold. (a) the original face. (b) its reconstruction results using m=900, 300, 100, respectively.

the boundary of the mesh. The boundary condition $\psi^0(x_i)$ describes a 1D piecewise linear curve in the plane. The curve should be convex for the parameterization to be bijective. In order to get regular sample result, the curve is fixed as square's boundary (See Fig. 3).

3.3 Shape Filtering

Previous works on 3D face feature extraction rely on deformation information and normal map[10][14]. They are encoded on a 2D grid. In this paper, we extract mid-frequency band of shape as feature for geometry image analysis. We select spectral method as our shape filter.

Based on graph Laplacian operators, Taubin[19] firstly introduces the spectral analysis as a digital geometry signal processing tool to characterize the classical approximations of filter. This approach is based on the similarity between the eigenvectors of the graph Laplacian and the basis functions used in the discrete Fourier transform. The basic idea is to compute the eigenfunctions and eigenvalues of the Laplacian on a general manifold surface.

Given a mesh with n vertices, its graph Laplacian operators $L \in \mathbb{R}^{n\times n}$ is a matrix where $L_{i,j} = \omega_{i,j} > 0$, whenever (i,j) is an edge, otherwise $L_{i,j} = 0$. $L_{i,i} = -\Sigma\omega_{i,j}$. The coefficients $\omega_{i,j}$ are weights associated to the edges of the graph. One may use the uniform weight $\omega_{i,j} = 1$ or more elaborate weights, computed from the embedding of the graph, e.g. distance weight with $\omega_{i,j} = \frac{1}{\|x_i - x_j\|}$ and cotangent weight with $\omega_{i,j} = cot(\alpha_{ij}) + cot(\beta_{ij})$. In order to avoid the influence of irregular mesh, we will use fully mesh-independent symmetric weight[21]:

$$\omega_{i,j} = \frac{cot(\alpha_{ij}) + cot(\beta_{ij})}{\sqrt{A_i A_j}}, \tag{2}$$

where α_{ij}, β_{ij} denote the two angles opposite to the edge (i,j) and A_i, A_j are the areas of the Voronoi cell of vertex i, j.

The solution to this eigen problem after mapping them into the canonical basis yields a series of eigenpairs (H^k, λ_k) called the Manifold Harmonics Bases (MHB). Smaller eigenvalues of the spectrum are correlated to low-frequency signals which account for the global feature while higher eigenvalues correlated to high-frequency signals represent the details. Using the MHB, we can define the Manifold Harmonics Transform (MHT) to convert the signal on the manifold surface into the spectral domain.

Considering a function defined on the geometry vertex position of triangulated surface based on $X = \sum_{j=1}^{n} x_j \phi_j$, where x_j denotes the function value at vertex j and $\phi_j(x_k) = \delta_{j,k}$,

(a) (b) (c) (d)

Figure 5: Geometry images of different attributions. (a) geometry coordinate image (b) normal image (c) normal variation image (d) deformation image.

we can reconstruct X at vertex j using the first m frequencies:

$$x_j = \sum_{k=1}^{m} \widetilde{x}_k H_j^k, \widetilde{x}_k = \langle X, H^k \rangle = \sum_{j=1}^{n} x_j D_{jj} H_j^k. \tag{3}$$

D is a diagonal matrix defined by mesh structure. $H^k(m = 1,2,\ldots,n)$ are the orthogonal manifold harmonics bases defined on the surface. In practice, X could be geometry shape coordinate or appearance color. The reconstruction for shape using m low-frequency bases are visualized in Fig. 4. This process can be seen as applying an ideal low-pass filter to the signal on manifold. Similarly, we can define high-pass filter using the frequencies started from m. In our experiments, we set the value of m to be 7.

3.4 Geometry Image

Each individual facial shape X can be represented by $X = \mathscr{B} \oplus \mathscr{D}$, where $\mathscr{B}$ is a smooth base surface. It can be computed based on MHT with the m leading eigenvectors. The geometric detail information $\mathscr{D}$ is a vector function that associates a displacement vector $\mathbf{h}_i$ with each point $\mathbf{b}_i$ on the base surface. According to Botsch's work[4], it is also called multi-scale deformation. Hence, the per-vertex displacement vector can be written as:

$$\mathbf{h}_i = \mathbf{x}_i - \mathbf{b}_i, \mathbf{h}_i \in \mathbb{R}^3 \tag{4}$$

The straightforward representation for the geometric details is a displacement vector. On the other side, two normal directions for facial shape X and its base surface $\mathscr{B}$ can be calculated. If we denote them using $\mathbf{n}_{\mathbf{x_i}}$ and $\mathbf{n}_{\mathbf{b_i}}$, the variation of normal direction can be written as:

$$\Delta = \mathbf{n}_{\mathbf{x_i}} - \mathbf{n}_{\mathbf{b_i}} \tag{5}$$

In this paper, we just focus on the normal variation. It is the output of the system after applying spectral filter to facial shape. Then it is assigned to the vertices of planar mesh. After regular sample on this grid, we can get 2D image containing geometry information. Fig. 5 shows us four types of geometry images. They are obtained by assigning different attributions to mesh vertices, namely geometry coordinate image, normal image, normal variation image and deformation image, respectively. Then, cosine similarity measure is used during matching step.

Table 1: Performance comparison

Method	Neutral Only	Expressions Only	Full Database
With Filter	98.4%	96.1%	97.5%
Without Filter	96.3%	90.2%	93.7%

4 Performance

In this section, we present experiments on 3D database for face recognition, which validate the shape filtering algorithm. We will demonstrate its robustness to expressions by conducting comprehensive experiments on FRGC face database[5]. The FRGC v2 includes FRGC v1 and images acquired during the fall of 2003 and spring of 2004. It includes 4007 3D face images of 466 distinct human subjects with from 1 to 22 images per subject acquired by the Minolta Vivid scanner. The spatial resolution of the scanner is 480×640. However, the resolution of faces in the database varies because they are scanned at different distances from the scanner. For each subject, there are also non-neutral expressions such as smile, astonish, puffy cheeks and angry.

4.1 Performance Metrics

Two different scenarios are used for experiments: identification and verification. In an identification scenario, the database are divided into two sets: probe and gallery. We register only one 3D data for each person using the earliest scan in the gallery and the rest as probes. In the probe set, there may be more than one instance of each person. Further, there is an easy division on probe into two subsets: one containing only data sets where facial expressions are present, the other containing only data sets with neutral expressions. In identification experiment, the system outputs a rank of the most similar face based on the similarity distance metrics. Then, the identification performance is measured using a Cumulative Match Characteristic (CMC) curve.

In verification scenario, a 4007×4007 similarity matrix is computed and normalized by matching all the 3D scans. Using the masks provided by the FRGC protocol, the Receiver Operating Characteristic (ROC) curve is obtained to evaluate the verification performance.

4.2 Effect of Shape Filtering

The purpose of this identification experiment is to evaluate the filtering stage which adopted in our work. In this test, the resolution of geometry image is set to 256×256. Note that the geometry image chosen for identification without filtering stage is normal image.

The performance on different subsets without filtering stage is measured and compared to performance on these sets using shape filter. Fig.6.a shows the results without shape filtering stage. The rank-1 identification rate drops more than 5 percentage when expression only probe is employed. After using shape filter (See Fig.6.b), the decrease of rank-1 identification rate on two subsets becomes no more than 2.3%. There is a 2.7% increase of recognition rate between them. On the other side, shape filter raises the performance by 2.1% for neutral only probe set, 5.9% for expression only set and 3.8% for full database. The comparisons are concluded in Tab.1. It is obvious that using shape filter can improve both accuracy and robustness.

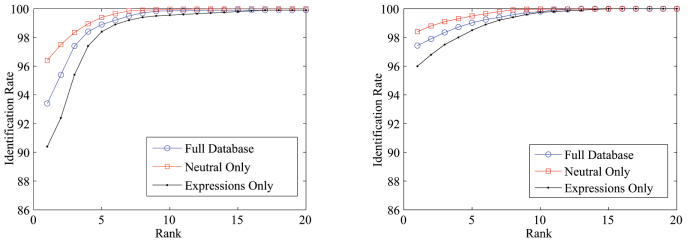

Figure 6: Performance comparison. (a) shows the results without shape filter. (b) is the results using shape filter

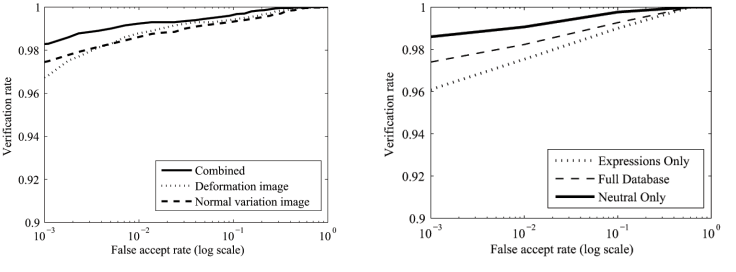

Figure 7: Verification performance. (a) ROC III curve for different images (b) ROC curve for different probe sets

4.3 Geometry Images and Their Fusion

The two geometry images, normal variation image, deformation image and their fusion are tested in this section. Fig.7.a indicates the verification result by the ROC curves.

In our experiment, the model-based deformation are extracted for each probe to characterize an individual. In order to disregard unreliable facial areas, we use area weight to remove the effects of intensive expression, just like AFM. Then, these deformation is assigned to planar mesh of the model. Pyramid transform coefficient of deformation field is computed for individual feature together with the coefficient of multi-scale deformation. Complex wavelet structural similarity metric (CW-SSIM) index algorithm is also employed to measures the similarity between two sliding windows placed in the same positions of the two images. From Fig. 7. a, we can see that normal variation image feature has higher performance than deformation image feature. The underlying reason of such result is as follows. There are still registration errors using ICP, especially in case of intense facial expressions. Normal variation image does not depend on registration result that can get rid of this errors. The fusion of the two features gives the higher verification rate than the individual features. With the FAR of 0.1% in the fusion, the verification rate is 98.2%. The results demonstrate that normal variation and deformation images can be used in describing the individual characteristics and the combination of them achieves higher performance.

4.4 Effect of Expressions Variation

This experiment focuses on the effect of facial expressions in verification. The FRGC v2 database provides a categorization of the expressions that each individual assumes, allowing an easy division on two subsets: one containing only data sets where facial expressions are present, the other containing only data sets with neutral expressions.

The performance on the two subsets is measured and compared to the performance on the full database as shown in Fig.7.b. The verification rates at 0.001 FAR are 97.6%, 96.2% and 98.7% respectively on the full database, a subset containing only non-neutral facial expressions and a subset containing only neutral expressions. The decrease of 2.5% of the verification rate is very modest. This can be concluded that the proposed approach is robust to expressions variation.

4.5 Comparison with Other Methods

We compare our algorithm with ICP and the state of the art methods including the AFM[10], the EDM[1] and the MMH[12]. AFM is also a non-statistical method based on rigid registration. EDM is a statistical deformable model focusing on expression variation. MMH is a 2D-3D hybrid approach. Each method uses all data in FRGC v2 database. A comparison of the identification rate by rank-one recognition and the verification rate at 0.001 FAR are presented in Table 2. It clearly indicates that our proposed scheme performs better than others. We can also see that non-statistical approaches obtain higher performance evaluation. Note that the results are quoted from the original papers. NA denotes that the data is not given by their works.

Table 2: Performance comparison

Method	ICP	AFM	EDM	MMH	This paper
Rank-one	75.7%	97%	96.1%	96.2%	97.5%
ROC III	NA	97%	94.05%	NA	98.2%

5 Discussion

From an accuracy point of view, we show that mid-frequency band information of facial shape could be used for discriminative personal-specific deformation feature. Our results in face recognition demonstrate the accuracy of our method. It becomes apparent that we can use shape filter to increase performance. This is shown in the identification scenario, where, by using shape filter, there was a significant increase in performance (Fig. 6).

A detailed error analysis in the face recognition experiments revealed that the majority of the failure cases are related to face segment step. This happens because the detection of nose tip unavoidably affects the deformation extraction and the final geometry image analysis steps. In cases of hair interference, the segment area does not cover the whole face. In addition, filter based on manifold harmonics transform used in this paper is step-type filter. It is different from our generative model in Fig. 2. Designing gaussian-type filter is within our future work.

From an efficiency point of view, the proposed method has a very efficient retrieval step but has a relatively heavy preprocessing step for each new test. In the identification scenario, the average preprocessing time for a facial scan was approximately 20.6s. The average retrieval time was 10^{-4}s. These tests were carried on matlab with 2.33 GHz CPU and 2 GB RAM. Registration is the most time-consuming step. A detailed analysis in shape filter and mesh parametrization revealed that registration method has little affect on the most geometry images except deformation image. Hence, shape filtering has the potential to handle this problem.

6 Conclusion

We have proposed a novel solution to expression challenge by using shape filtering. Compared with previous statistical-based recognition algorithms, the proposed method enjoys several major advantages. Firstly, our method is efficient without registering several expressions for learning. Secondly, it can match face without considering generalizing problem. Compared with other non-statistical approaches, we have improved both accuracy and robustness. We have conducted experimental evaluations on FRGC v2 database. The experimental results show that the proposed method is efficient and achieves better recognition performance.

Although promising experimental results have validated the efficiency of our algorithm, some limitations should be addressed in future work. First of all, our face segment algorithm is based on nose tip detection. Depth information based approach is not robust. To address the issue, we may consider employing multi-scale feature points detection algorithm. Secondly, reducing computation time also need to be considered. Finally, we may add appearance filter to enhance the performance.

References

[1] F. Al-Osaimi, M. Bennamoun, and A. Mian. An expression deformation approach to non-rigid 3d face recognition. *IJCV*, 81(3):302–316, 2009.

[2] B. Amberg, R. Knothe, and T. Vetter. Expression invariant face recognition with a morphable model. *IEEE International Conference on Automatic Face and Gesture Recognition*.

[3] V. Blanz and T. Vetter. Face recognition based on fitting a 3d morphable model. *TPAMI*, 25(9):1063–1074, 2003.

[4] M. Botsch. and O. Sorkine. On linear variational surface deformation methods. *IEEE Transactions on Visualization and Computer Graphics.*, 14(1):213–230, 2008.

[5] K.W. Bowyer, K. Chang, and P. Flynn. A survey of approaches and challenges in 3d and multi-modal 3d+2d face recognition. *CVIU*, 101(1):1–15, 2006.

[6] A. Bronstein, M. Bronstein, and R. Kimmel. Three-dimensional face recognition. *IJCV*, 64(1):5–30, 2005.

[7] L. Bruno. and H. Zhang. Sepctral mesh processing. In *SIGGRAPH ASIA 2009*, 2009.

[8] M. S. Floater and K. Hormann. Surface parameterizatio: a tutorial and survey. *In Advances in Multiresolution for Geometric Modelling*.

[9] Xiao. Huang, N. Paragios, and D.N. Metaxas. Shape registration in implicit spaces using information theory and free form deformations. *TPAMI*, 28(8):1303–1318, 2006.

[10] I. Kakadiaris, G. Passalis, G. Toderici, N. Murtuza, Y. Lu, and T. Theoharis. Three-dimensional face recognition in the presence of facial expressions: An annotated deformable model approach. *TPAMI*, 29(4):640–649, 2007.

[11] X. Lu and A. Jian. Deformation modeling for robust 3d face matching. *TPAMI*, 30(8): 1346–1356, 2008.

[12] A. Mian, M. Bennamoun, and R. Owens. An efficient multimodal 2d-3d hybrid approach to automatic face recognition. *TPAMI*, 29(11):1927–1943, 2007.

[13] Y. Nagai, Y. Ohtake, and H. Suzuki. Smoothing of partition of unity implicit surfaces for noise robust surface reconstruction. *Computer Graphics Forum*, 28(5):1339–1348, 2009.

[14] G. Passalis, I. Kakadiaris, and T. Theoharis. Intraclass retrieval of nonrigid 3d objects: Application to face recognition. *TPAMI*, 29(2):218–229, 2007.

[15] C.C. Queirolo, L. Silva, O.R. Bellon, and M.P. Segundo. 3d face recognition using simulated annealing and the surface interpenetration measure. *TPAMI*, 32(2):206–219, 2009.

[16] C. Samir, A. Srivastava, and M. Daoudi. Three-dimensional face recognition using shapes of facial curves. *TPAMI*, 28(11):1858–1863, 2006.

[17] A. Scheenstra, A. Ruifrok, and R.C. Veltkamp. A survey of 3d face recognition methods. *Lecture Notes in Computer Science*, 3546(2005):891–899, 2005.

[18] A. Sheffer, E. Praun, and K. Rose. Mesh parameterization methods and their application. *Foundations and Trends in Computer Graphics and Vision*, 2(2):105–171, 2006.

[19] G. Taubin. Signal processing approach to fair surface design. *In:SIGGRAPH*.

[20] F. terHaar and R.C. Veltkamp. A 3d face matching framework for facial curves. *Graphical Models*, 71(2):71–91, 2009.

[21] B. Vallet and B. levy. Spectral geometry processing with manifold harmonics. *Computer Graphics Forum*, 27(2):251–260, 2008.

[22] Y. Wang, G. Pan, and Z. Wu. 3d face recognition in the presence of expression: A guidance-based constraint deformation approach. *In:CVPR*.

[23] W. Zhao, R. Chellappaand P. Phillips, and A. Rosenfeld. Face recognition: A literature survey. *ACM Computing Surveys*, 35(4):399–458, 2003.