UR3D-C: Linear Dimensionality Reduction for Efficient 3D Face Recognition

Omar Ocegueda[1], Georgios Passalis[1,2], Theoharis Theoharis[1,2],
Shishir K. Shah[1], Ioannis A. Kakadiaris[1]

Abstract

We present a novel approach for computing a compact and highly discriminant biometric signature for 3D face recognition using linear dimensionality reduction techniques. Initially, a geometry-image representation is used to effectively resample the raw 3D data. Subsequently, a wavelet transform is applied and a biometric signature composed of 7,200 wavelet coefficients is extracted. Finally, we apply a second linear dimensionality reduction step to the wavelet coefficients using Linear Discriminant Analysis and compute a compact biometric signature. Although this biometric signature consists of just 57 coefficients, it is highly discriminant. Our approach, UR3D-C, is experimentally validated using four publicly available databases (FRGC v1, FRGC v2, Bosphorus and BU-3DFE). State-of-the-art performance is reported in all of the above databases.

1. Introduction

During the last decade, several 3D face recognition methods have claimed to have overcome the limitations of 2D face recognition. However, only a few of these methods have been thoroughly validated. This lack of validation was mainly attributed to the limited availability of publicly available 3D facial databases. In the past few years several such databases were released, alleviating this problem. The 3D facial databases currently available, consist of thousands of datasets that include a variety of facial expressions, pose variations and even partial occlusions.

In this paper, we propose a novel 3D face recognition method, UR3D-C, and validate it using the most challenging databases available: FRGC v1, FRGC v2, Bosphorus and BU-3DFE. These databases, with the exception of FRGC v1, are considered very challenging as they each have several thousand datasets. Moreover, they can be considered complementary as they have different characteristics (e.g., FRGC v2 has a large number of subjects while Bosphorus has a large variety of facial expressions).

The main contribution of this paper is the introduction of a biometric signature for 3D face recognition that is very compact (only 57 coefficients), and highly discriminant (state-of-the-art performance in challenging databases). In fact, it outperforms all previously published methods in both efficiency and accuracy.

To extract the geometry images from the raw 3D data, we follow the subdivision-based deformable model framework, UR3D, proposed by Kakadiaris *et al*. [7]. In UR3D, after each geometry image is extracted, a pair of biometric signatures is computed based on two separate wavelet transforms (Haar and Pyramid). In order to measure the similarity between a probe and a gallery, the pairs of signatures are compared (using a different metric for each type of signature) and the results are fused. In UR3D-C, we derive an initial signature composed of 7,200 coefficients by applying the Haar wavelet transform. Subsequently, we apply the LDA algorithm proposed by Yu and Yang [18]. The result is a compact biometric signature of only 57 coefficients that is used to measure the similarity between a probe and a gallery image. Recently, we proposed UR3D-B [10], an extension to UR3D, consisting of a feature selection step (Gauss-Markov Posterior Marginals) that reduces the dimensionality of the biometric signature to 360 coefficients. By applying LDA to the compact signatures, state-of-the-art performance on FRGC v2 was obtained. UR3D-C is more efficient and more accurate when compared to [7] and [10] (discussed in detail in Section 3.3).

Additional state-of-the-art methods that use FRGC v2 for experimental validation were proposed by Queirolo *et al*. [14], Faltemier *et al*. [5] and Al-Osaimi *et al*. [1]. For a broader survey of the field the reader is referred to Bowyer *et al*. [3] and Chang *et al*. [4]. Queirolo *et al*. [14] used Simulated Annealing to match facial parts derived from a segmentation step. Subsequently, the individual scores were combined using a sum rule to compute the final score. Faltemier *et al*. [5] considered 38 regions on the face and compared the corresponding regions between the probe and gallery datasets using the Iterative Closest Point (ICP) algorithm. A voting scheme was then used to fuse the results from the individual regions. Finally, Al-Osaimi *et al*. [1] used an expression deformation model to transfer facial expressions from one dataset to another. A method based on Principal Component Analysis (PCA) was then applied to compare the datasets. Note that UR3D-C outperforms all

---

[1]Computational Biomedicine Lab, Department of Computer Science, University of Houston, Houston, TX 77004, USA

[2]Department of Informatics and Telecommunications, University of Athens, Ilisia, Attiki 15784, Athens, Greece
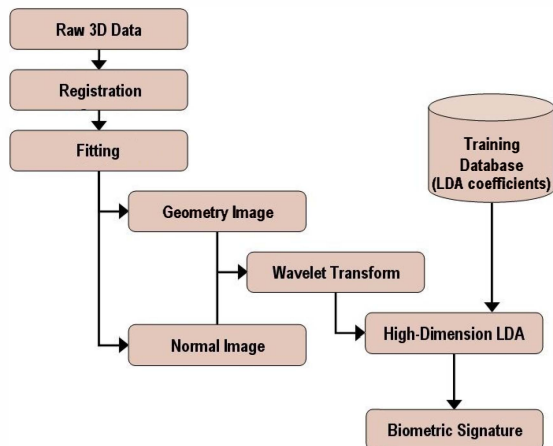
1

Figure 1. Depiction of the algorithmic pipeline of UR3D-C.

the above methods, as demonstrated in Section 3.3.

2. Methods

The basic step of UR3D-C is the application of LDA on the wavelet transform of the geometry images. A geometry image is the result of the injective mapping of all vertices of a 3D object (x, y and z coordinates) to a 2D grid representation (u, v coordinates) [6]. Thus, a geometry image is a regular continuous sampling of a 3D model represented as a 2D image, with each u, v pixel corresponding to the original x, y, z coordinates. There are several approaches in the literature for the creation of geometry images. In this paper, we use the approach presented by Kakadiaris *et al.* [7]. The algorithmic pipeline, which is based on UR3D [7], is briefly described in Section 2.1 while the novel high-dimensional LDA approach is discussed in Section 2.2.

2.1. Overview

The main step involves fitting an Annotated Face Model (AFM) to the raw 3D data using a subdivision-based deformable model framework [7]. The geometry images are then created from the fitted AFM. The basic steps of this pipeline are:

1. The raw 3D data obtained by the optical or laser scanner are converted to a polygonal representation. Simple filters (e.g., smoothing, median cut) are applied to the data to alleviate scanner specific problems (e.g., holes or noise).

2. The raw 3D data are registered to the AFM using a coarse-to-fine approach: an initial registration is provided by the ICP algorithm [2] while a more accurate registration is provided by the Simulated Annealing algorithm [16].

3. The AFM is fitted to the registered data to acquire the shape of the individual. For the fitting, a deformable model framework [9] is used that is based on subdivision surfaces (specifically, Loop's subdivision scheme [8]).

4. A geometry image is created from the fitted AFM. This is possible since the AFM has a native injective mapping from $\mathcal{R}^3$ to $\mathcal{R}^2$ (this property is not violated by the fitting process). The first derivative of the geometry image is also computed, thus creating a normal image of similar size.

5. Each of the geometry and normal images has a spatial resolution of 256×256 and 3 channels (393,216 coefficients). A level-four Haar wavelet transform is applied to the images, producing 256 wavelet packets of size 16×16. Only the 40 most discriminant packets are kept, resulting in 61,440 coefficients. Using a selection mask derived from the AFM, only 30 out of the 256 coefficients inside each packet are retained (the remaining are disregarded as they are considered less useful for biometric purposes). The resulting biometric signature consists of 7,200 coefficients ($2 \times 3 \times 40 \times 30$).

6. With the application of high-dimensional LDA, the final biometric signature's size is reduced from 7,200 to just 57 coefficients, as presented in Section 2.2.

2.2. Linear Dimensionality Reduction

The "weighted" L_1 norm, which is the metric used to compare pairs of wavelet signatures, can be written as:

$$d(\mathbf{g}, \mathbf{p}) = \sum_{i=1}^{n} a_{ii}|g_i - p_i| = |\mathbf{A}(\mathbf{g} - \mathbf{p})|_1, \qquad (1)$$

where $\mathbf{g}$ and $\mathbf{p}$ are n-dimensional vectors, $\mathbf{A}$ is a diagonal $n \times n$ matrix whose elements are the wavelet weights assigned to each coefficient (it contains three weights for a coefficient assigned to a pixel), and $|\cdot|_1$ represents the L_1 norm. If we relax the constraint over $\mathbf{A}$ to be diagonal, and consider any $m \times n$, $1 < m < n$ matrix instead, we will, at the same time, increase the degrees of freedom of $\mathbf{A}$ from n to $n \times m$, and reduce the dimensionality of the data from n to m. Given a set of training data $\mathcal{T} = \{(\mathbf{x_i}, y_i)\}_{i=1}^{N}$, where N is the number of samples in the database, $\mathbf{x_i} \in \mathcal{R}^n$ are $n-$dimensional biometric signatures, $y_i \in \{1, 2, ..., K\}$ are the labels assigned to the signatures (the identifier of each subject), LDA provides a set of discriminant directions $\{\mathbf{w_i}\}_{i=1}^{K-1}$ that maximize the Fisher's criterion:

$$\arg\max_{\mathbf{w} \in \mathcal{R}^n} \frac{\mathbf{w^T B w}}{\mathbf{w^T W w}}, \qquad (2)$$
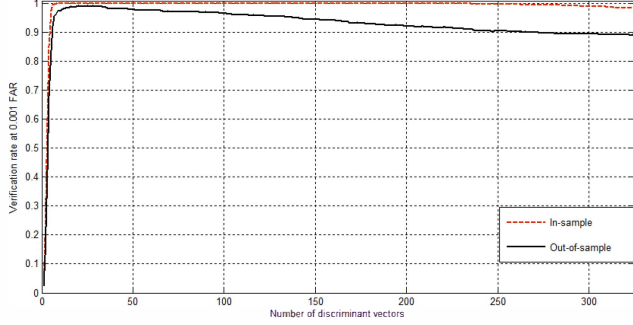
2

Figure 2. Verification performance as a function of the number of discriminant vectors used for recognition.

where $K \leq N$ is the number of subjects in the training set, $\mathbf{B}$ is the between class scatter matrix, and $\mathbf{W}$ is the within class scatter matrix. It has been shown that using only the leading discriminant directions can often improve the generalization performance of the classifier [19]. In the context of face recognition, this suggests that there exists a subset of discriminant directions $\mathcal{S} = \{\mathbf{w_i}\}_{i=1}^{t}$ that captures the most discriminant information to recognize different identities, while its complement $\mathcal{S}' = \{\mathbf{w_i}\}_{i=t+1}^{K-1}$ captures the discriminant information related to the training set. Thus, using $\mathcal{S}'$ for recognition leads to over-fitting. This phenomenon is illustrated in Fig. 2.

We divided the FRGC v2 database into two subsets: $\mathcal{T}_1$ (70 % of the subjects) and $\mathcal{T}_2$ (30 % of the subjects), computed the $K - 1$ discriminant vectors from $\mathcal{T}_1$ using LDA ($K = 326$) and projected $\mathcal{T}_2$ on the first t vectors, $0 < t < K$. We then defined the distance between any two samples in $\mathcal{T}_2$ as the L_2 distance between their projected t−dimensional vectors. Next, we computed the verification rate at 0.001 FAR on $\mathcal{T}_1$ and $\mathcal{T}_2$ for all values of t. The verification rate on $\mathcal{T}_1$ is the *in-sample* verification rate and the verification rate on $\mathcal{T}_2$ is the estimated *out-of-sample* verification rate. Note that the out-of-sample verification rate decreases as a function of the number of discriminant vectors after reaching its maximum at $t = 20$, while the in-sample verification rate remains at its maximum up to 213 vectors, after which the performance decreases slightly. This suggests that the first 20 vectors capture the most discriminant information for face recognition, while the rest capture information about the training set. The procedure for computing the t discriminant vectors is described in Algorithm 1.

3. Experimental Results

In the following experiments, we will follow the methodology described in Section 2.2 to define the set of discriminant vectors and the number t of vectors to be used for

---

Algorithm 1 LDA-Train

Require: A training set $\mathcal{T} = \{(\mathbf{x_i}, y_i)\}_{i=1}^{N}$ of n-dimensional vectors (defined by the selected wavelet coefficients, as described in Section 2.1) labeled according to subject identity, $y_i \in \{1, 2, ..., K\}$

1: Split $\mathcal{T}$ into two subsets $\mathcal{T}_1$, $\mathcal{T}_2$ such that each subject in $\mathcal{T}$ appears in $\mathcal{T}_1$ or $\mathcal{T}_2$ but not both
2: Let K_1 and K_2 be the number of subjects in $\mathcal{T}_1$ and $\mathcal{T}_2$ respectively
3: Apply LDA to $\mathcal{T}_1$, and let $\mathbf{A}$ be the resulting $(K_1 - 1) \times n$ projection matrix
4: For each $0 < r < K_1$, define $\mathbf{A_r}$ as the $r \times n$ matrix formed using the first r rows of $\mathbf{A}$
5: **for all** $r \in \{1, 2, ..., K_1 - 1\}$ **do**
6:     Compute the projection $\hat{\mathbf{x}}_\mathbf{i} = \mathbf{A_r} \mathbf{x_i}$ for all vectors $\mathbf{x_i} \in \mathcal{T}_2$
7:     Compute the verification rate v_r at 0.001 FAR on $\mathcal{T}_2$ using the projected r-dimensional vectors $\hat{\mathbf{x}}_\mathbf{i}$ as biometric signatures and the L_2 norm as metric
8: **end for**
9: Set $t = \arg\max_r \{v_r\}$
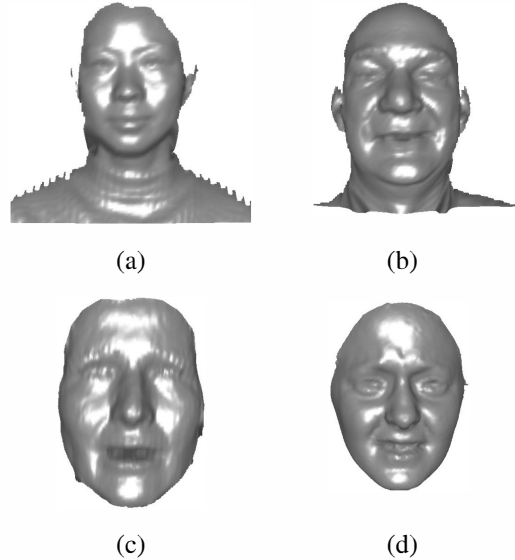10: **return** $(t, \mathbf{A_t})$

Figure 3. Example datasets from databases: (a) FRGC v1, (b) FRGC v2, (c) Bosphorus and (d) BU-3DFE.

recognition.

3.1. Databases

To thoroughly validate the performance of UR3D-C we used four publicly available databases. Example datasets from each of these databases are depicted in Fig. 3 while their description is given below.

- **FRGC v1**: The FRGC v1 database [11] contains a total of 943 range images acquired in 2003. These data were obtained from 275 subjects, all exhibiting neutral expressions. The hardware used to acquire these range data was a Minolta Vivid 900 laser range scanner, with a resolution of 640×480.

- **FRGC v2**: The FRGC v2 database [13, 12] contains a total of 4,007 range images, acquired between 2003 and 2004. These data were obtained from 466 subjects and contain various facial expressions (e.g., happiness and surprise).

- **Bosphorus**: The Bosphorus database [15] consists of 3D facial data with facial expressions, pose variations and partial occlusions. To acquire these data, an Inspeck Mega Capturor II 3D was used. For the experiments in this paper, only the frontal datasets that did not have any occlusions were used. As a result, there were a total of 2,902 datasets from 105 subjects. We call this subset $\mathcal{B}_1$. Compared to FRGC v2, this database has a larger variety of facial expressions but a smaller number of subjects.

- **BU-3DFE**: The BU-3DFE database [17] contains a total of 2,500 3D datasets from 100 subjects. Each subject has 25 different facial expression. The database consists of 56% female and 44% male subjects, with age ranging from 18 to 70 years and a variety of ethnic/racial ancestries.

3.2. Cross Database Validation

Although cross validation is the standard way to assess the performance of systems that require training, using cross validation would make the results not directly comparable to other face recognition systems that, for example, reported their performance on FRGC v2. Therefore, we tested the performance using each database for training and testing. The results are summarized in Tables 1 and 2. In Table 1, we report the rank-1 recognition rate, while in Table 2, we report the verification rate at 0.001 FAR for each database by taking the first sample of each subject (in lexicographical order) as gallery and the rest as probes.

3.3. Face Recognition Grand Challenge Experiments

To compare UR3D-C with the state-of-the-art, we used the experiments defined for FRGC v2 during the Face Recognition Grand Challenge (FRGC). To this end, we report the verification rate at 0.001 FAR for experiments I, II and III using each of the available databases for training (Table 3). Additionally, the rank-1 recognition rate is also reported. The best results are obtained by training with $\mathcal{B}_1$. Since different databases exhibit different characteristics, it

Rank-1, First vs. Rest					
		Training			
		FRGC v1	BU-3DFE	$\mathcal{B}_1$	FRGC v2
Testing	FRGC v1	100.0 %	96.7 %	99.9 %	100.0 %
	BU-3DFE	97.3 %	99.9 %	99.7 %	99.3 %
	$\mathcal{B}_1$	96.4 %	98.2 %	99.6 %	98.6 %
	FRGC v2	98.1 %	96.1 %	97.9 %	99.7 %

Table 1. Rank 1 recognition rate using each database under consideration for training and testing. The first sample of each individual forms the gallery set and the rest of the samples are used as probes.

Verification Rate at 0.001 FAR, First vs. Rest					
		Training			
		FRGC v1	BU-3DFE	$\mathcal{B}_1$	FRGC v2
Testing	FRGC v1	99.7 %	95.7 %	99.1 %	100.0 %
	BU-3DFE	91.3 %	99.1 %	97.2 %	97.0 %
	$\mathcal{B}_1$	89.6 %	93.8 %	96.9 %	94.1 %
	FRGC v2	97.5 %	95.5 %	97.5 %	99.7 %

Table 2. Verification Rate at 0.001 FAR using each database under consideration for training and testing. The first sample of each individual forms the gallery set and the rest of the samples are used as probes.
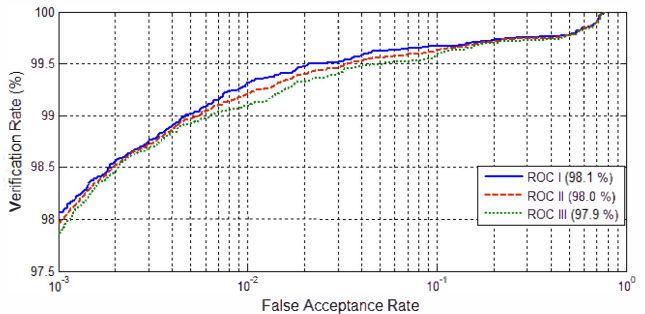

Figure 4. Verification performance on FRGC v2 using the $\mathcal{B}_1$ and FRGC v1 databases for training.

is natural to expect that the discriminant vectors computed using LDA with different databases, captured different discriminant information. This information can be combined in different ways to further improve the recognition performance.

The last column of Table 3 summarizes the results obtained by averaging the similarity matrices trained with $\mathcal{B}_1$ and FRGC v1, which corresponds to a signature size equal to 57 coefficients. The recognition performance of UR3D-C, along with some of the most recently published 3D-3D face recognition algorithms, is presented in Table 4, and the corresponding Receiver Operating Characteristic (ROC) curves are depicted in Fig. 4. Note that UR3D-C outperforms all previously published methods.

4

	BU-3DFE	FRGC v1	$\mathcal{B}_1$	FRGC v1 +$\mathcal{B}_1$
Signature dimension	21	30	27	57
ROC I	95.4 %	96.5 %	97.4 %	98.1 %
ROC II	95.4 %	96.4 %	97.3 %	98.0 %
ROC III	95.4 %	96.1 %	97.2 %	97.9 %
Rank-1	96.1 %	98.1 %	97.9 %	99.0 %

Table 3. Recognition performance for all FRGC v2 experiments and rank-1 recognition rate using different training sets. Different training sets result in different signature length.

Method	ROC I	ROC II	ROC III	Rank-1
UR3D-C	98.1 %	98.0 %	97.9 %	99.0 %
UR3D-B [10] (2011)	97.5 %	97.1 %	96.8 %	-
Queirolo *et al.* [14] (2010)	-	-	96.6 %	98.4 %
UR3D [7] (2007)	97.3 %	97.2 %	97.0 %	97.0 %
Faltemier *et al.* [5] (2008)	-	-	94.8 %	97.2 %
Al-Osaimi *et al.* [1] (2009)	-	-	94.1 %	96.5 %

Table 4. Recognition performance on FRGC v2 of UR3D-C compared against the most recently published 3D-3D face recognition methods.

4. Discussion

Some important aspects of our Dimensionality Reduction approach must be mentioned. At first, it might appear too simplistic to use a linear transformation for dimensionality reduction, since there are several non-linear approaches reported in the literature. However, Fig. 2 and Tables 1 and 2 demonstrate that a nearly perfect in-sample performance (i.e., training and testing on the same database) can be obtained using a linear transformation, which means that the training set was nearly linearly separable, and therefore, the application of non-linear techniques is not necessary under these conditions. Furthermore, the use of more sophisticated techniques bears the risk of over-fitting the training set.

Although the homoscedasticity assumption (i.e., the within covariance matrix is the same for all classes) used in LDA may be unreasonable, relaxing this assumption is not likely to improve the recognition performance. The reason is that our classification problem is conceptually different from other classification tasks. In a standard classification problem, the categories of the data are defined in advance, so that the variability within each class may be used to improve the classification accuracy. However, in face recognition, we are only applying LDA for dimensionality reduction. This is because the categories (i.e., the subject identities) and even the cardinality of the categories are not known in advance. In addition, those categories are, in general, different from the categories in the training set.

Thus, using the variability of the categories in the training set (i.e., the variability associated to each subject) will not provide information about the variability of a subject outside the training set. Therefore, allowing heteroscedasticity is more likely to cause over-fitting.

The problem of face recognition can be formulated as a binary classification problem, in which the objective is to distinguish between pairs of faces corresponding to the same individual and pairs of faces corresponding to different individuals. However, this formulation is problematic because it is highly unbalanced (there are many more negative examples than positive examples). Under this formulation, it is more appropriate to use a non-linear classifier and a dimensionality reduction of the biometric signature (e.g., a signature composed of only 360 coefficients as proposed in [10]) is mandatory in order to apply more sophisticated classifiers. However, in this paper we demonstrated that it is possible to use a larger signature, composed of 7,200 coefficients, since LDA algorithms specially designed for high-dimensional data are readily available [18].

Finally, notice that the resulting biometric signature, composed of 57 coefficients (456 bytes using double precision floating point values), allows us to compare a large number of signatures in a short period of time, making the identification in large databases very efficient in terms of speed and storage (it is possible to store all the biometric signatures of the FRGC v2 database in less than 2 *MB*). In our experiments, we are able to perform 999,504 comparisons per second using an AMD Opteron processor at 2.1 *GHz*.

5. Conclusion

A 3D face recognition method, UR3D-C, using linear dimensionality reduction was presented in this paper. It uses an automated algorithmic pipeline to derive a wavelet transform from geometry images. LDA is applied to a training dataset and subsequently used to compress the biometric signature to just 57 coefficients. UR3D-C was experimentally validated using four publicly available databases. The importance of LDA's training was thoroughly investigated as the performance was computed for all possible combinations of training and test databases (sixteen combinations for the four databases). Finally, the verification rates at 0.001 FAR in the three standard experiments defined for FRGC v2 are the highest to date (98.1 %, 98.0 % and 97.9 %, respectively).

Acknowledgments

References

[1] F. Al-Osaimi, M. Bennamoun, and A. Mian. An expression deformation approach to non-rigid 3D face recognition. *International Journal of Computer Vision*, 81(3):302–316, March 2009. 1, 5

[2] P. Besl and N. McKay. A method for registration of 3D shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):239–256, February 1992. 2

[3] K. Bowyer, K. Chang, and P. Flynn. A survey of approaches and challenges in 3D and multi-modal 3D+2D face recognition. *Computer Vision and Image Understanding*, 101(1):1–15, January 2006. 1

[4] K. Chang, K. Bowyer, and P. J. Flynn. An evaluation of multi-modal 2D+3D face biometrics. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(4):619–624, April 2005. 1

[5] T. Faltemier, K. Bowyer, and P. Flynn. A region ensemble for 3D face recognition. *IEEE Transactions on Information Forensics and Security*, 3(1):62–73, March 2008. 1, 5

[6] X. Gu, S. Gortler, and H. Hoppe. Geometry images. In *Proc. 29^{th} International Conference on Computer Graphics and Interactive Techniques*, pages 355–361, San Antonio, TX, July 2002. 2

[7] I. A. Kakadiaris, G. Passalis, G. Toderici, M. N. Murtuza, Y. Lu, N. Karampatziakis, and T. Theoharis. Three-dimensional face recognition in the presence of facial expressions: An annotated deformable model approach. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(4):640–649, April 2007. 1, 2, 5

[8] C. Loop. Smooth subdivision surfaces based on triangles. Master's thesis, Department of Mathematics, University of Utah, 1987. 2

[9] C. Mandal. *A dynamic framework for subdivision surfaces*. PhD thesis, University of Florida, 1998. 2

[10] O. Ocegueda, S. K. Shah, and I. A. Kakadiaris. Which parts of the face give out your identity? In *Proc. IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 641–648, Colorado Springs, CO, June 2011. 1, 5

[11] U. of Notre Dame. University of Notre Dame Biometrics Database. http://www.nd.edu/∼cvrl, 2008. 4

[12] J. Phillips, T. Scruggs, A. O'Toole, P. Flynn, K. Bowyer, C. Schott, and M. Sharpe. FRVT 2006 and ICE 2006 large-scale results. *National Institute of Standards and Technology Internal Report*, 7408, March 2007. 4

[13] P. J. Phillips, P. J. Flynn, T. Scruggs, K. W. Bowyer, J. Chang, K. Hoffman, J. Marques, J. Min, and W. Worek. Overview of the face recognition grand challenge. In *Proc. IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 1, pages 947–954, San Diego, CA, June 20-25 2005. 4

[14] C. Queirolo, L. Silva, O. Bellon, and M. Segundo. 3D face recognition using simulated annealing and the surface interpenetration measure. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(2):206–219, February 2010. 1, 5

[15] A. Savran, N. Alyüz, H. Dibeklioğlu, O. Çeliktutan, B. Gökberk, B. Sankur, and L. Akarun. Bosphorus database for 3D face analysis. In *Proc. 1^{st} COST 2101 Workshop on Biometrics and Identity Management*, pages 47–56, Roskilde, Denmark, May 2008. 4

[16] P. Siarry, G. Berthiau, F. Durbin, and J. Haussy. Enhanced simulated annealing for globally minimizing functions of many-continuous variables. *ACM Transactions on Mathematical Software*, 23(2):209–228, 1997. 2

[17] L. Yin, X. Wei, Y. Sun, J. Wang, and M. Rosato. A 3D facial expression database for facial behavior research. In *Proc. 7^{th} International Conference on Automatic Face and Gesture Recognition*, pages 211–216, Southampton, UK, April 10-12 2006. 4

[18] H. Yu and J. Yang. A direct LDA algorithm for high-dimensional data with application to face recognition. *Pattern Recognition*, 34:2067–2070, October 2001. 1, 5

[19] M. Zhu and T. J. Hastie. Feature extraction for nonparametric discriminant analysis. *Journal of Computational and Graphical Statistics*, 12(1):101–120, March 2003. 3