Landmark Detection for Unconstrained Face Recognition

Panagiotis B. Perakis *

National and Kapodistrian University of Athens
Department of Informatics and Telecommunications
p.perakis@di.uoa.gr; takis@antinoos.gr

Abstract. In this dissertation a novel method for 3D landmark detection and pose estimation, suitable for both frontal and side 3D facial scans, is presented. It exploits 3D and 2D information by using local shape descriptors to extract candidate interest points that are subsequently identified and labeled as anatomical landmarks. Additionally, a novel generalized framework for combining facial feature descriptors that can be used for landmark detection is introduced, and several feature fusion schemes are proposed and evaluated. However, feature detection methods which use general purpose shape descriptors cannot identify and label the detected candidate landmarks. To this end, a 3D *Facial Landmark Model* (FLM) of facial anatomical landmarks is introduced. Candidate landmarks, irrespectively of the way they are generated, can be identified and labeled by matching them with the FLM. Finally, a novel method for unconstrained face recognition is introduced. It employs the 3D landmark detector to provide an initial pose estimation and to indicate occluded areas with missing data for each facial scan. Subsequently, a 3D *Annotated Face Model* (AFM) is registered and fitted to the scan using facial symmetry to complete the occluded areas. Using a biometric signature resulted from the wavelet representation of the fitted AFM, the proposed method can perform comparisons among interpose facial scans, unlike previously proposed methods that require frontal scans.

Keywords: Biometrics, Face Recognition, Landmark Detection, Shape Models, Shape Descriptors, Feature Extraction, Feature Fusion, Pose Estimation, Partial Matching, Deformable Models.

1 Introduction

Biometrics is the science of establishing the identity of a person based on the physical (e.g., fingerprints, face, hand geometry, and iris) or behavioral (e.g., gait, signature, and keyboard dynamics) attributes associated with an individual [1].

Face recognition is the procedure of recognizing an individual from their facial attributes or features and is one of the primary biometric modalities. Face

* Dissertation Advisor: Theoharis Theoharis, Professor.

recognition has several advantages over other biometric technologies: it is non-intrusive, since the facial region is generally exposed, and potentially easy to use [2]. Thus, research and development in face recognition followed naturally.

The performance of face recognition systems has improved significantly since the first automatic face recognition system was developed by Kanade [3] in 1973. Furthermore face recognition can now be performed in "realtime" for images captured under constrained situations. Although progress in face recognition has been encouraging, it has also turned out to be a difficult endeavor, especially for unconstrained tasks where view point, illumination, inter-object occlusions, facial expressions and facial accessories vary considerably [2].

2 Challenges & Motivation

Face recognition has proved to be a very challenging task due to the numerous sources of variation in 2D and 3D facial data. These variations can be environment-based (illumination conditions, occlusions by other objects or accessories), subject-based (pose and expression variations) and acquisition-based (image scale, distortion, noise, spikes and holes).

The main reason for using information from 3D data in face recognition systems is that the data acquired by 3D acquisition devices are invariant to pose and lighting conditions, these being the major challenges with which face recognition algorithms must cope [4].

With the increase in the availability of 3D data, several 3D face recognition approaches have been proposed. These approaches aim to overcome the limitations of 2D face recognition by offering pose invariance. However, although they claim pose invariance, they mostly utilize frontal 3D scans assuming that the entire face is visible to the sensor (see the surveys of Bowyer *et al.* [5] and Chang *et al.* [6]). This assumption is not always valid in real-world applications, since unconstrained acquisition may lead to facial scans with extensive occlusions that result in missing data due to pose variations.

Thus, existing 3D face recognition methods, fail to address large pose variations and to confront the problem of missing facial areas in an automatic way. The main assumption of these methods is that even though the head can be rotated with respect to the sensor, the *entire* face is always visible. However, this is true only for "almost frontal" scans or "reconstructed" complete face meshes aligned to frontal pose. *Side scans usually have large missing areas, due to self-occlusion, that depend on pose variations.* These scans are very common in realistic scenarios such as uncooperative subjects or uncontrolled environments. Therefore, to take advantage of the full pose invariance potential of 3D face recognition, the problem of missing data must be addressed. Thus, in a face recognition system, an initial registration step, based on landmark points' correspondence, is necessary in order to make the system pose invariant [7, 8].

However, facial landmark detection also suffers from the same sources of variation in 2D and 3D facial data that face recognition does [9–13]. Both 2D and 3D facial landmark detection suffer from occlusion, pose and expression variations.

In addition, 2D facial landmark detection also suffers from illumination variations. Thus, a landmark detection algorithm must be pose-invariant to address the problem of missing facial areas and, at the same time, expression-invariant in order to allow the registration of the various instances of the face liable to expression variations.

3 Aim & Methodology

The uncontrolled conditions of real-world biometric applications pose a great challenge to any 3D face recognition approach. The unconstrained acquisition of data from uncooperative subjects may result in facial scans with significant pose and expression variations.

In this dissertation, an integrated novel method is proposed, in order to automatically detect landmarks on 3D facial scans that exhibit pose and expression variations, and hence consistently register and compare any pair of facial datasets subjected to missing data due to self-occlusion in a pose- and expression-invariant face recognition system.

The proposed landmark detection and face recognition system employs an automatic pose- and expression-invariant landmark detector, using local facial feature descriptors and a deformable 3D *Facial Landmark Model* (FLM) to ensure global topological consistency of the detected landmarks [14, 8, 15, 16].

3.1 Training of Facial Landmark Models and Feature Templates

At the training phase, a Facial Landmark Model (FLM) is created by first aligning the training landmark sets and calculating a mean landmark shape using *Procrustes Analysis*, and then applying *Principal Component Analysis* (PCA) to capture the shape variations [17–19]. The FLM serves as a 3D geometric model of the landmark points. Also, templates for each shape descriptor that represents each landmark point are calculated from training facial datasets [14, 8, 15, 16].

The shape templates serve as feature descriptors for each landmark point. The feature descriptors that have been used, depending on the case, include the *Shape Index* [20], a continuous map of principal curvature values of a 3D object's surface, the *Spin Image* [21], a local descriptor of the object's 3D point distribution, the *Extruded Points* [15], a local descriptor of a 3D object's points that extrude most and the *Edge Response* [22] descriptor, a local descriptor of the 2D texture gradient of a 3D object.

3.2 Facial Landmark Detection

At the detection phase, the algorithm first detects candidate landmarks on the queried facial datasets according to the similarity of the extracted facial features

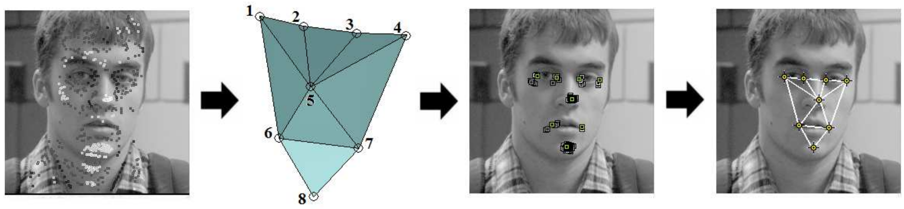

Fig. 1. Process pipeline of landmark detection: (a) extracted candidate landmarks using feature descriptors; (b) Facial Landmark Model (FLM); (c) landmark sets consistent with FLM; (d) resulting optimal landmark set.

with the feature templates. The extracted candidate landmarks are then filtered out and labeled by matching them with the FLM (Fig. 1) [14, 8, 15, 16].

During the research conducted under this dissertation, several versions of the presented generalized framework for facial landmark detection were applied. The most important are summarized in the following:

SISI–NPSS METHOD To locate landmark points, shape index target values for each landmark class (eye outer corner, eye inner corner, nose tip, mouth corner and chin tip) were searched for on the shape index map. Subsequently, the candidate landmark points of the five landmark classes that are obtained from the shape index map were further filtered out according to the similarity of their spin images with the spin image templates representing each landmark class. The resulting candidate landmark points of the five landmark classes were subsequently filtered out according to their consistency with the FLM. To find the optimum landmark set, the product of *normalized Procrustes distance* $\times$ *(1–mean spin similarity)* was used as a distance metric between the candidate landmark sets and the FLM.

Fusion METHOD In this method, fusion schemes for combining landmark features were incorporated into the landmark detection pipeline. To locate landmark points the shape index map, the spin image map and the edge response map were fused into a resultant similarity map, each for every landmark class. The candidate landmarks for each landmark class were searched on the corresponding resultant similarity map. Subsequently, the candidate landmark points of the five landmark classes were filtered out according to their consistency with the FLM. To find the optimum landmark set, the product of *normalized Procrustes distance* $\times$ *(1 – resultant similarity)* was used as a distance metric between the candidate landmark sets and the FLM.

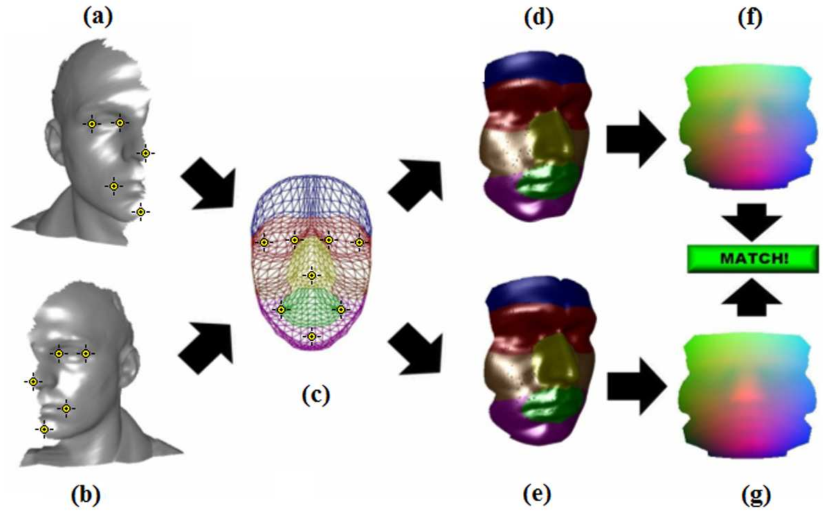

Fig. 2. Interpose matching using the proposed method: (a) and (b) opposite side facial scans with extensive missing data and detected landmarks; (c) generic Annotated Face Model (AFM); (d) and (e) registered and deformed AFM for each scan (facial symmetry used); (f) and (g) extracted geometry images.

3.3 Partial Face Recognition

The landmark detector provides an initial pose estimation (frontal, right, left) and indicates occluded areas with missing data for each facial scan resulting from pose variations. Facial landmark detection is a crucial first step for the registration of the facial datasets that have to be compared [8, 15].

Subsequently, a generic *Annotated Face Model* (AFM) [7] is registered and fitted to each facial probe scan, using a subdivision-based deformable model framework. During fitting, facial symmetry is used to complete the occluded areas of the face [8, 15]. Signature metadata are extracted using a *wavelet transformation* on the *geometry* and *normal images* of the fitted AFM (Fig. 2). A similarity measure between signature metadata of probe and gallery facial datasets provide the face recognition results.

4 Experimental Results

4.1 Landmark Detection

Test Databases For the performance evaluation of the proposed landmark detector, the largest publicly available 3D face and ear databases were combined. To evaluate the performance of the method against yaw variations, frontal, semi-profile and profile facial datasets were used. To evaluate the tolerance of the method against expression variations, subjects with varying degrees of expressions were included. To have a measure of the landmark detection error, the used facial datasets were manually annotated at the queried landmark points. For frontal facial scans, the FRGC v2 database [23, 24] was used. For side facial scans, the Ear Database from the University of Notre Dame (UND) [25] was used.

For the conducted experiments, the following collections of facial datasets were created:

- DB00F: Contains 975 frontal facial scans obtained from 149 different subjects, selected from the FRGC v2 database, including subjects with varying degrees of expressions (45.44% "neutral", 36.41% "mild" and 18.15% "extreme"), acquired under varying illumination conditions (e.g. half of the face shaded).
- DB00F45RL: a composite frontal-to-profile database with the datasets of 39 common subjects found in the FRGC v2 database and in the UND Ear database. This database contains 117 (3x39) facial scans having three poses, frontal (39 scans) and 45° left (39 scans) and right (39 scans).
- DB45L and DB45R: two semi-profile databases with 118 left and 118 right 45° side datasets, which come from 118 different subjects, obtained from the UND Ear database.
- DB60L and DB60R: two profile databases with 87 left and 87 right 60° side datasets, which come from 87 different subjects, obtained from the UND Ear database.

Landmark Detection Evaluation The performance evaluation of a landmark detector is generally presented by computing the following values, which represent the localization accuracy of the detected landmarks:

Absolute Distance Error: The Euclidean distance in physical units (e.g., *mm*) between the position of the detected landmark and the manually annotated landmark, which is considered ground truth.

Detection Success Rate: The percentage of successful detections of a landmark over a test database. Successful detection is considered as the detection of a landmark with Absolute Distance Error under a certain threshold (e.g., 10 *mm*).

Summary results for METHOD SISI–NPSS on all tested databases are presented in Table 1. The results clearly indicate that the proposed method exhibits high accuracy and robustness both to yaw and expression variations. The mean error is under 6.3 *mm*, with standard deviation under 2.6 *mm* on all tested facial scans. Also note that the mean error is under 10 *mm* for at least 90.4% of the tested facial scans and the facial side was correctly estimated on over 98.9% of the tested facial scans.

Table 1. Summary results for METHOD SISI–NPSS

Database	Mean Error mean (mm)	stdev (mm)	≤ 10 (mm)	Side Detection Rate
DB00F	5.00	1.85	97.85%	99.90%
DB00F-neutral	4.52	1.51	99.32%	100.00%
DB00F-mild	4.95	1.46	99.72%	100.00%
DB00F-extreme	6.28	2.60	90.40%	99.44%
DB00F45RL	4.97	1.92	97.44%	100.00%
DB45R	5.03	1.92	96.61%	100.00%
DB45L	4.75	1.91	97.46%	100.00%
DB60R	4.95	1.80	96.55%	98.85%
DB60L	5.30	2.49	93.10%	100.00%

Evaluation of Fusion Schemes The evaluation of the performance of the proposed distance to similarity mappings and fusion schemes for landmark detection is not a straight-forward task, since there are many factors that characterize performance. As already stated, fusion techniques are expected to improve system's *accuracy*, *efficiency* and *robustness*. An equally important characteristic of a fusion scheme is that of *monotonicity*, i.e., the addition of a new feature descriptor should improve prior results. A qualitative performance evaluation of the proposed fusion schemes according to the aforementioned characteristics is presented in Table 2.

Table 2. Qualitative evaluation of proposed fusion schemes

	Accuracy	Efficiency	Robustness	Monotonicity
L−L1	Fair	**High**	Fair	Fair
L−L2	Fair	Low	Fair	Fair
L−Lg	**High**	Fair	Fair	Fair
Q−L1	**High**	**High**	Fair	Fair
Q−L2	**High**	**High**	**High**	**High**
Q−Lg	**High**	Fair	Fair	Fair
G−L1	**High**	**High**	**High**	**High**
G−L2	**High**	**High**	Fair	Fair
G−Lg	**High**	Fair	Fair	Fair
L−Lmax	Low	Low	Low	Low
Q−Lmax	Low	Low	Low	Low
G−Lmax	Low	Low	Low	Low
L−Lmin	Unreliable	Fair	Fair	Low
Q−Lmin	Unreliable	Fair	Fair	Low
G−Lmin	Unreliable	Fair	Fair	Low

Current experimental results show that, in general, the Quadratic (Q) and Gaussian (G) mappings behave better than the Linear (L) mapping of distance measure to similarity measure. For the Linear mapping the product rule (Lg) behaves better than other rules. For the Quadratic mapping the rms rule (L2) behaves better than other rules. For the Gaussian mapping the sum rule (L1) behaves better than other rules. Quadratic and Gaussian mappings have almost the same performance.

Accuracy improvement is more dramatic when the information fused is correlated. In correlated features the performance of one descriptor predicts to some extent the performance of the other and strengthens the results. On the other hand highly uncorrelated features have similarity peaks that do not coincide and degrade the results. Efficiency improvement is achieved by excluding obvious non-matches, reducing the number of candidate landmarks, for each landmark class. Fusion, also, reduces system sensitivity to sample-specific, poor-quality or erroneous descriptors.

We can thus deduce that the best performance in terms of accuracy is exhibited by the Q-L2 and G-L1 fusion schemes, with the Q-L2 exhibiting a slight better performance than the G-L1 in landmarks' likelihood area reduction. Q-L2 and G-L1 also exhibit high robustness in yaw, expression and illumination variations, and strong monotonicity.

Also landmark localization using the Q-L2 fusion scheme improved the accuracy and robustness of the landmark detector (with $3.5-5.5$ mm mean landmark localization error), indicating the superiority of the fusion approach.

4.2 Partial Face Recognition

Test Databases

Combined UND Databases: To evaluate the performance of the proposed partial face recognition method, a combination of the largest publicly available 3D face and ear databases was used. For frontal facial scans, the FRGC v2 database [23, 24] was used. For side facial scans, the Ear Database from the University of Notre Dame (UND) [25] was used.

For the conducted experiments the following collections were defined:

- UND45LR: Contains 45° side scans from 118 subjects. For each subject, the left scan is considered gallery and the right is considered probe. *Total: 236 scans.*
- UND60LR: Contains 60° side scans from 87 subjects. For each subject, the left scan is considered gallery and the right is considered probe. *Total: 174 scans.*
- UND00LR: Gallery set has one frontal scan for each of the 466 subjects. Probe set has two 45° side scans (left and right) from 39 subjects and two 60° side scans (left and right) from 32 subjects. *Total: 608 scans.*

In all cases there is only one gallery scan per subject. Also, all subjects present in a probe set are also present in the gallery set (the opposite is not always true).

UH Databases: In addition to the UND databases a database with data collected at the University of Houston was used. The database contains 1,075 left and 1,075 right scans of 281 subjects. The novelty of this database is that each pair of left and right side scans was acquired simultaneously.

For the conducted experiments the following collection is defined:

– UHDB7LR-M: Contains multiple left and right side scan pairs from 281 subjects. For each subject, one left and one right scan are considered gallery and the rest are considered probes (1–6 left and 1–6 right scans per subject). *Total: 2,150 scans.*

In all cases there is one pair of gallery scans per subject. Also, all subjects present in a probe set are also present in the gallery set (the opposite is not always true).

The proposed method tackles the problem of matching arbitrary facial scans (left, right or frontal). This is considerably harder than matching only frontal scans, since a lot of the facial information is missing and it is not known a priori whether each scan is left, right or frontal.

Matching facial scans of arbitrary side In this experiment, the performance of the proposed partial face recognition method, using scans of arbitrary sides for gallery and probe sets, was evaluated. This is a realistic scenario, as the side scans (with extensive occlusions that lead to missing data) are very common in real world applications with unconstrained acquisition. The proposed method can match any combination of left, right or frontal facial scans with the use of facial symmetry. Moreover, the proposed method automatically detects the side of the scan by using the automatic landmark detector. For this experiment we utilized the UND45LR, UND60LR, UND00LR and UHDB7LR-M databases and the rank-one rates are given in Table 3.

Table 3. Rank-one Recognition Rate between facial scans of arbitrary side

	Rank-one Rate
UND45LR	86.4%
UND60LR	81.6%
UND00LR	76.8%
UHDB7LR-M	89.1%

In the cases of UND45LR and UND60LR, for each subject, the gallery set contains a single left side scan while the probe set contains a single right side scan. Therefore, facial symmetry is always used in order to perform identification. As expected, the 60° side scans yield lower results as they are considered more challenging compared to the 45° side scans.

In the case of UND00LR, the gallery set contains a frontal scan for each subject, while the probe set contains left and right side scans. This scenario is very common when the enrollment of subjects is controlled but the identification is uncontrolled. Compared to UND45LR and UND60LR, there is a decrease in the performance of the proposed method in UND00LR. One could argue that since the gallery set consists of frontal scans (without missing data), there should be an increase in performance. However, UND00LR has the largest gallery set, making it the most challenging database in current experiments.

In the case of UHDB7LR-M, for each subject, the gallery set contains a left and right side scan pair, while the probe set contains multiple left and right side scan pairs. As expected, since the gallery set has two scans per subject, the performance on this database is the highest among all databases. The performance difference is substantial compared to UND00LR (89.1% versus 76.8% rank-one). This indicates that one pair of left and right side scans is more descriptive than one frontal scan.

5 Conclusion

In this thesis an automatic facial landmark detection methodology has been proposed. It offers pose invariance and robustness to large missing (self-occluded) facial areas with respect to large yaw variations and high tolerance to large expression variations. The proposed approach consists of methods for landmark localization that exploit the 3D facial geometry and the modeling ability of trained landmark models. It has been evaluated using the most challenging 3D facial databases available, which contain scans with yaw variations of up to 80° and strong expressions. In these databases it achieved state-of-the-art accuracy (with $4.5-6.3$ mm mean landmark localization error), significantly outperforming existing methods.

Also, a novel generalized framework of fusion methods and their application to landmark detection has been presented. The proposed fusion scheme transforms features to similarities and then combines them to generate a resultant feature similarity, which is considered as the matching score for the detection of the queried landmarks. The proposed feature fusion framework is easily extensible to new feature-components, offers significant dimensionality reduction and works equally well for features extracted from 3D or 2D facial data.

For the proposed fusion scheme different distance to similarity mappings and different fusion rules have been evaluated. The results indicate that the quadratic distance to similarity mapping in conjunction with the rms rule for fusion (Q-L2) exhibits the best performance. Landmark localization using this fusion scheme achieved state-of-the-art accuracy (with $3.5-5.5$ mm mean landmark localization error), indicating the superiority of the fusion approach.

Finally, a novel 3D face recognition method suitable for real-world biometric applications was proposed. Unlike most previous methods that require frontal scans, the proposed method can perform partial matching among interpose facial scans, even when extensive data are missing. It exploits the 3D landmark detector to provide an initial pose estimation and to indicate occluded areas with missing data for each facial scan. By using facial symmetry to complete missing facial data, it can handle seamlessly frontal and side facial scans.

The presented method for partial face recognition is extensively evaluated against a variety of 3D facial databases, achieving state-of-the-art performance (with average rank-one recognition rate 83.7%), considerably outperforming existing methods, even when tested on the most challenging data, which contain scans with yaw variations up to 80° and strong expressions.

The proposed system is suitable for real-world scenarios as the only requirement is that half of the face is visible to the sensor, and its computational cost is low. Using a standard Intel Core 2 Duo 2.2 GHz PC, 18 *sec* on average are required to process a facial scan: 9 *sec* to localize the facial landmarks plus 9 *sec* to extract the biometric signature (geometry and normal images). The biometric signatures can be matched at a rate of 15,000 $matches/sec$.

Ackowledgement

This PhD research at the Department of Informatics and Telecommunications of the University of Athens has been co-financed by the European Union (European Social Fund - ESF) and Greek national funds through the Operational Program "Education and Lifelong Learning" of the National Strategic Reference Framework (NSRF) - Research Funding Program: Heracleitus II. Investing in knowledge society through the European Social Fund.

References

1. Ross, A., Jain, A.K.: Biometrics, overview. In Li, S., ed.: Encyclopedia of Biometrics. Springer, New York, NY (2009) 168–172
2. Li, S., Jain, A.K. In: Handbook of Face Recognition. Springer (2005) 1–11
3. Kanade, T.: Picture Processing by Computer Complex and Recognition of Human Faces. PhD thesis, Kyoto University (1973)
4. Kakadiaris, I., Passalis, G., Toderici, G., Perakis, P., Theoharis, T.: 3D-based face recognition. In Li, S., ed.: Encyclopedia of Biometrics. Springer, New York, NY (2009) 329–338
5. Bowyer, K., Chang, K., Flynn, P.: A survey of approaches and challenges in 3D and multi-modal 3D+2D face recognition. Computer Vision and Image Understanding **101**(1) (Jan. 2006) 1–15
6. Chang, K., Bowyer, K., Flynn, P.J.: An evaluation of multi-modal 2D+3D face biometrics. IEEE Transactions on Pattern Analysis and Machine Intelligence **27**(4) (Apr. 2005) 619–624
7. Kakadiaris, I., Passalis, G., Toderici, G., Murtuza, M., Lu, Y., Karampatziakis, N., Theoharis, T.: 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) (Apr. 2007) 640–649
8. Passalis, G., Perakis, P., Theoharis, T., Kakadiaris, I.: Using facial symmetry to handle pose variations in real-world 3D face recognition. IEEE Transactions on Pattern Analysis and Machine Intelligence **33**(10) (Oct. 2011) 1938–1951

9. Lu, X., Jain, A.: Multimodal facial feature extraction for automatic 3D face recognition. Technical Report MSU-CSE-05-22, Michigan State University (Oct. 2005)
10. Colbry, D., Stockman, G., Jain, A.: Detection of anchor points for 3D face verification. In: Proc. IEEE Computer Society Conference on Computer Vision and Pattern Recognition, San Diego, CA (Jun. 20-25 2005) 118
11. Lu, X., Jain, A.: Automatic feature extraction for multiview 3D face recognition. In: Proc. 7^{th} International Conference on Automatic Face and Gesture Recognition, Southampton, UK (Apr. 10-12 2006) 585–590
12. Lu, X., Jain, A., Colbry, D.: Matching 2.5D face scans to 3D models. IEEE Transactions on Pattern Analysis and Machine Intelligence **28**(1) (2006) 31–43
13. Colbry, D.: Human Face Verification by Robust 3D Surface Alignment. PhD thesis, Michigan State University (2006)
14. Perakis, P., Passalis, G., Theoharis, T., Kakadiaris, I.: 3D facial landmark detection under large yaw and expression variations. IEEE Transactions on Pattern Analysis and Machine Intelligence **35**(7) (July 2013) 1552–1564
15. Perakis, P., Passalis, G., Theoharis, T., Toderici, G., Kakadiaris, I.: Partial matching of interpose 3D facial data for face recognition. In: Proc. 3^{rd} IEEE International Conference on Biometrics: Theory, Applications and Systems, Arlington, VA (Sep. 28-30 2009) 439–446
16. Perakis, P., Theoharis, T., Passalis, G., Kakadiaris, I.: Automatic 3D facial region retrieval from multi-pose facial datasets. In: Proc. Eurographics Workshop on 3D Object Retrieval, Munich, Germany (Mar. 30 - Apr. 3 2009) 37–44
17. Dryden, I., Mardia, K.: Statistical Shape Analysis. Wiley (1998)
18. Stegman, M., Gomez, D.: A brief introduction to statistical shape analysis. Technical report, Technical University of Denmark (Mar. 2002)
19. Cootes, T., Taylor, C.: Statistical models of appearance for computer vision. Technical report, University of Manchester (Oct. 2001)
20. Dorai, C., Jain, A.K.: COSMOS - a representation scheme for 3D free-form objects. IEEE Transactions on Pattern Analysis and Machine Intelligence **19**(10) (Oct. 1997) 1115–1130
21. Johnson, A.E.: Spin Images: A Representation for 3-D Surface Matching. PhD thesis, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA (Aug. 1997)
22. Harris, C., Stephens, M.: A combined corner and edge detector. In: Proc. 4^{th} Alvey Vision Conference. (1988) 147–151
23. Phillips, P., Flynn, P., Scruggs, T., Bowyer, K., Chang, J., Hoffman, K., Marques, J., Min, J., Worek, W.: Overview of the Face Recognition Grand Challenge. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition, San Diego, CA (2005) 947–954
24. Phillips, P., Scruggs, T., O'Toole, A., Flynn, P., Bowyer, K., Schott, C., Sharpe, M.: FRVT 2006 and ICE 2006 large-scale experimental results. IEEE Transactions on Pattern Analysis and Machine Intelligence **32** (2010) 831–846
25. UND: University of Notre Dame Biometrics Data Sets. http:// www.nd.edu/ ~cvrl/ CVRL/ Data_Sets.html (2012)