IRACST – International Journal of Computer Networks and Wireless Communications (IJCNWC), ISSN: 2250-3501
Vol.11, No 3, July– Sep 2021

A Review Of 3D Face Recognition Systems And Multi-Modal 3D+2D Systems

Kevin W. Bowyer, Kyong Chang, and Patrick Flynn
Department of Computer Science and Engineering
University of Notre Dame
Notre Dame, Indiana 46556 USA

kwb, kchang, flynn @cse.nd.edu

Abstract

The vast majority of face recognition research has focused on the use of two-dimensional intensity images, and is covered in existing survey papers. This survey focuses on face recognition performed by matching two three-dimensional shape models, either alone or in combination with two-dimensional intensity images. Challenges involved in developing more accurate three-dimensional face recognition are identified. These include the need for improved sensors, recognition algorithms, and experimental methodology.

Keywords: biometrics, face recognition, three-dimensional, multi-modal.

1. Introduction

Evaluations such as the Face Recognition Vendor Test 2002 [19] make it clear that the current state ofthe art in face recognition is not yet sufficient for the more demanding biometric applications. However, biometric technologies that currently offer greater accuracy, such as fingerprint and iris, require much greater explicit cooperation from the user. For example, fingerprint requires that the subject cooperate in making physical contact with the sensor surface. This raises issues of how to keep the surface clean andgerm-free in a high-throughput application. Iris imaging currently requires that the subject cooperate to carefully position their eye relative to the sensor. This can also cause problems in a high-throughput application. Thus it appears that there is significant potential application-driven demand for improved performance in face recognition systems.

The vast majority of face recognition research, and all of the major commercial face recognition systems, use normal intensity images of the face. We will refer to these as "2D images." In contrast,a "3D image" of the face is one that represents three-dimensional shape. A distinction can be made between representations that include only the surface of the face and those that include the whole head. In this distinction, the face surface would be "2.5-D" and the whole head would be 3D. We will ignore this distinction here, and refer to the shape of the face surface as 3D. The 3D shape of the face is often sensed in combination with a 2D intensity image. In this case, the 2D image can be thought of as a "texture map" overlaid on the 3D shape. An example of a 2D intensity image and the corresponding 3D shape are shown in Figure 1, with the 3D shape rendered both in the form of a range image and in the form of a shaded 3D model. A range image, a shaded model, and a wire-frame mesh are common alternatives for rendering 3D face data.

A recent survey of face recognition research is given in [23], but it does not include algorithms based on matching 3D shape. In this current survey, we focus specifically on face recognition algorithms that match the 3D shape of the face to be recognized against the enrolled 3D shape of the face(s) of the known person(s). That is, we are interested in systems that perform person recognition or authentication by matching two 3D face descriptions. We use "recognition" here to refer to one-to-many matching to

2D intensity 3D as range

image 3D as shaded
model

Figure 1. Example 2D Intensity and 3D Shape. The left and middle images are the cropped intensity and range images, respectively, as would be used by PCA style face recognition algorithms. The right image is a 3/4 view of the 3D shape from which the range image is created.

find the best match above some threshold, and "authentication" to refer to one-to-one matching used to verify or reject a claimed identity. A particular research group may present their results in the context of one type of application or the other, but the core 3D representation and matching issues are essentially the same. We do **not** consider here the family of approaches in which a generic, "morphable" 3D face model is used as an intermediate step in matching two 2D images for face recognition [5]. As commonly used, the term *multi-modal biometrics* refers to the use of multiple imaging modalities, such as 3D and 2D images of the face. We consider algorithms for multi-modal 3D and 2D face recognition along with those that use only 3D shape.

We are particularly interested in 3D face recognition because it is often thought that the use of 3D has the potential for greater recognition accuracy than the use of 2D face images. For example, one paper states - "Because we are working in 3D, we overcome limitations due to viewpoint and lighting variations" [14]. Another paper describing a different approach to 3D face recognition states - "Range images have the advantage of capturing shape variation irrespective of illumination variabilities" [10]. Similarly, a third paper states - "Depth and curvature features have several advantages over more tra- ditional intensity based features. Specifically, curvature descriptors 1) have the potential for higher accuracy in describing surface based events, 2) are better suited to describe properties of the face in a areas such as the cheeks, forehead, and chin, and 3) are viewpoint invariant" [9]. We will return to thisissue of 3D versus 2D later in the paper.

2 Survey of 3D Face Recognition Algorithms

Although early work on 3D face recognition was done over a decade ago, the number of published papers on 3D and multi-modal 2D+3D face recognition is small enough that we can cover essentially all such work. Often a research group has published multiple papers as they develop a line of work. In such cases, we discuss only the most recent and easily accessible publication from that line. Some important relevant properties of the published algorithms and results in 3D face recognition are summarized in Table 1.

Cartoux et al. [7] approach 3D face recognition by segmenting a range image based on principal curvature and finding a plane of bilateral symmetry through the face. This plane is used to normalize forpose. They consider methods of matching the profile from the plane of symmetry and of matching the face surface, and report 100% recognition for either in a small dataset.

Lee and Milios [12] segment convex regions in the range image based on the sign of the mean and Gaussian curvatures, and create an Extended Gaussian Image (EGI) for each convex region. A match between a region in a probe image and in a gallery image is done by correlating EGIs. A graph match-ing algorithm incorporating relational constraints is used to establish an overall match of probe image to gallery image. Convex regions are believed to change shape less than other regions in response to changes in facial expression. This gives this approach some ability to cope with changes in facial ex-pression. However, EGIs are not sensitive to change in object size, and so two similar shape but different size faces will not be distinguishable in this representation.

Gordon [9] begins with a curvature-based segmentation of the face. Then a set of features are extracted that describe both curvature and metric size properties of the face. Thus each face becomes a point in feature space, and matching is done by a nearest-neighbor match in feature space. Experiments are reported with a test set of three views of each of eight faces, and recognition rates as high as 100% are reported. It is noted that the values of the features used are generally similar for different images of the same face, “except for the cases with large feature detection error, or variation due to expression” [9].

Nagamine et al. [18] approach 3D face recognition by finding five feature points, using those feature points to standardize face pose, and then matching various curves or profiles through the face data. Experiments are performed for sixteen subjects, with ten images per subject. The best recognition rates are found using vertical profile curves that pass through the central portion of the face. Computational requirements were apparently regarded as severe at the time this work was performed, as the authors note that “using the whole facial data may not be feasible considering the large computation and hardware capacity needed” [18].

Achermann et al. [3] extend eigenface and hidden Markov model approaches used for 2D face recog- nition to work with range images. They present results for a dataset of 24 persons, with 10 images per person, and report 100% recognition using an adaptation of the 2D face recognition algorithms.

Tanaka et al. [20] also perform curvature-based segmentation and represent the face using an Ex- tended Gaussian Image (EGI). Recognition is then performed using a spherical correlation of the EGIs. Experiments are reported with a set of 37 images from a National Research Council of Canada range image dataset, and 100% recognition is reported.

Achermann and Bunke [2] report on a method of 3D face recognition that uses an extension of the Hausdorff distance matching. They report on experiments using 240 range images, 10 images of each of 24 persons, and achieve 100% recognition for some instances of the algorithm.

Hesher et al. [10] explore principal component analysis (PCA) style approaches using different num- bers of eigenvectors and image sizes. The image data set used has 6 different facial expressions for each of 37 subjects. The performance figures reported result from using multiple images per subject in the gallery. This effectively gives the probe image more chances to make a correct match, and is known to raise the recognition rate relative to having a single sample per subject in the gallery [16].

Medioni and Waupotitsch [14] perform 3D face recognition using iterative closest point (ICP) match- ing of face surfaces. Whereas most of the works covered here used 3D shape acquired through a structured-light sensor, this work uses a stereo-based system. Experiments with seven images each from a set of 100 subjects are reported, and an equal error rate of "better than 2%" is reported.

Moreno and co-workers [17] approach 3D face recognition by first performing a segmentation basedon Gaussian curvature and then creating a feature vector based on the segmented regions. They report results on a dataset of 420 face meshes representing 60 different persons, with some sampling of different expressions and poses for each person. They report 78% rank-one recognition on the subset of frontal views, and 93% overall rank-five recognition.

Lee and co-workers perform 3D face recognition by locating the nose tip, and then forming a feature vector based on contours along the face at a sequence of depth values [13]. They report 94% correct recognition at rank five, and do not report rank-one recognition.

Lao et al. [11] perform 3D face recognition using a sparse depth map constructed from stereo images. Iso-luminance contours are used for the stereo matching. Both 2D edges and iso-luminance contours are used in finding the irises. In this specific limited sense, this approach is multi-modal. However, there is no separate recognition result from 2D face recognition. Using the iris locations, other feature points are found so that pose standardization can be done. Recognition rates of 87% to 96% are reported using a dataset of ten persons, with four images taken at each of nine poses for each person.

Beumier and Acheroy [4] approach multi-modal recognition by using a weighted sum of 3D and 2D similarity measures. They use a central profile and a lateral profile, each in both 3D and 2D. Therefore they have a total of four classifiers, and an overall decision is made using a weighted sum of the similarity metrics. Results are reported for experiments using a 27-person gallery and a 29-person probe set. An equal error rate as low as 1.4% is reported for multi-modal 3D+2D recognition that merges multiple probe images per subject. In general, multi-modal 3D+2D is found to perform better than either 3D or 2D alone.

Wang et al. [22] use Gabor filter responses in 2D and "point signatures" in 3D to perform multi- modal face recognition. The 2D and 3D features together form a feature vector.

images from 50 subjects, six images per subject, with pose and expression variations. Recognition ratesexceeding 90% are reported.

Bronstein et al. use an isometric transformation approach to 3D face analysis in an attempt to better cope with variation due to facial expression [6]. One method they propose is effectively multi-modal 2D+3D recognition using eigendecomposition of flattened textures and canonical images. They show examples of correct and incorrect recognition by different algorithms, but do not report any overall quantitative performance results for any algorithm.

Tsalakanidou et al. [21] report on multi-modal face recognition using 3D and color images. The use of color rather than simply gray-scale intensity appears to be unique among the multi-modal work surveyed here. Results of experiments using images of 40 persons from the XM2VTS dataset [15] are reported for color images alone, 3D alone, and 3D + color. The recognition algorithm is PCA style matching, plus a combination of the PCA results for the individual color planes and range image. Recognition rates ashigh as 99% are achieved for the multi-modal algorithm, and multi-modal performance is found to be higher than for either 3D or 2D alone.

Chang et al. [8] report on PCA-based recognition experiments performed using 3D and 2D images from 200 persons. One experiment uses a single set of later images for each person as the probes. Another experiment uses a larger set of 676 probes taken in multiple acquisitions over a longer elapsed time. Results in both experiments are approximately 99% rank-one recognition for multi-modal 3D+2D, 94% for 3D alone and 89% for 2D alone. The multi-modal result was obtained using a weighted sum ofthe distances from the individual 3D and 2D face spaces. This work represents the largest experimental study yet reported in the literature either for 3D face alone or for multi-modal 2D+3D, in terms of the number of subjects, the number of gallery and probe images, and the time lapse between gallery and probe image acquisition.

Table 1. Summary Of Research On 3D and Multi-Modal 2D+3D Face Recognition

Reference	number of persons	number of images	image size	3D face data	reported performance	size variation	expression variation
Face Recognition Algorithms Using Only 3D Data							
Cartoux 1989 [7]	5	18	?	profile, surface	100%	yes	no
Lee 1990 [12]	6	6	256x150	EGI	none	no	some
Gordon 1992 [9]	26 train 8 test	26 train 24 test	?	feature vector	100%	yes	no
Nagamine 1992 [18]	16	160	256x240	multiple profiles	100%	yes	no
Achermann 1997 [3]	24	240	75x150	range image	100%	yes	no
Tanaka 1998 [20]	37	37	256x256	EGI	100%	no	no
Achermann 1997 [2]	24	240	75x150	point set	100%	yes	no
Hesher 2003 [10]	37	222 (6 expr. ea.)	242x347	range image	97%	yes	no
Medioni 2003 [14]	100	700 (7 poses ea.)	?	surface mesh	98%	yes	no
Moreno 2003 [17]	60	420 (3 expr., 2 poses)	avg 2,200 point mesh	feature vector	78%	yes	some
Lee 2003 [13]	35	70	320x320	feature vector	94% at rank 5	yes	no
Multi-Modal 3D + 2D Face Recognition Algorithms							
Lao 2000 [11]	10	360	480x640	surface mesh	91%	yes	no
Beumier 2001 [4]	27 gallery 29 probes	240 2D	?	multiple profiles	1.4% EER	yes	no
Wang 2002 [22]	50	300	128x512	feature vector	>90%	no	yes
Bronstein 2003 [6]	157	?	2250 avg. vertices	range image	not reported	yes	yes
Tsalakanidou 2003 [21]	40	80	100x80	range image	99% 3D+2D 93% 3D only	yes	no
Chang 2003 [8]	200 (+ 75 in training)	951	480x640	range image	99% 3D+2D 93% 3D only	yes	no

3 Challenges To Improved 3D Face Recognition

Reflecting on the algorithms and results reviewed in the previous section, we can identify three ma- jor areas in which advances are required in order for 3D face recognition to become practical for wide application. One area is 3D sensor technology. While there may be a germ of truth in the optimism about 3D face data relative to 2D face images, there are still significant limitations in current 3D sen- sor technology. A second area is improved algorithms. For example, most current algorithms for 3D face recognition do not handle variation in facial expression well. Additionally, current algorithms for multi-modal 3D+2D recognition are multi-modal only in a weak sense. A third area is experimental methodology. Most published results to date are not based on a large and challenging dataset, do not report statistical significance of observed differences in performance, and make a biased comparison between multi-modal results and the baseline results from a single modality.

Improved 3D Sensors.

Successful practical application of 3D face recognition would be aided by various improvements in 3D sensor technology. Among these are: (1) reduced frequency and severity of artifacts, (2) increased depth of field, (3) increased spatial and depth resolution, and (4) reduced acquisition time.

It is important to point out that while 3D shape is *defined* independent of illumination, it is not *sensed* independent of illumination. Illumination conditions **do** affect the quality of sensed 3D data. Even under ideal illumination conditions for a given sensor, it is common for artifacts to occur in face regions such as oily regions that appear specular, the eyes, and regions of facial hair such as eyebrows, mustache, or beard. The most common types or artifacts can generally be described subjectively as “holes” or “spikes.” A “hole” is essentially an area of missing data, resulting from the sensor being unable to acquire data. A “spike” is an outlier error in the data, resulting from, for example, an inter-reflection in a projected light pattern or a correspondence error in stereo. An example of “holes” in a 3D face image sensed with the Minolta scanner is shown in Figure 2. Artifacts are typically patched up by interpolating new values based on the valid data nearest the artifact.

Figure 2. Example of "Hole" and "Spike" Artifacts In Sensed 3D Shape.
The 3D data is rendered as a cropped, frontal view, range image on the left. The black regions are "holes" of missing data. The data is rendered as a side view of a shaded shape model on the right. Noise points in the data are readily apparent as "spikes" away from the face surface. Essentially all 3D scanners are subject to some level these sorts of artifacts in the raw data.

Another limitation of current 3D sensor technology, especially relative to use with non-cooperative subjects, is the depth of field for sensing data. The depth of field for acquiring usable data might range from about 0.3 meter or less for a stereo-based system to about one meter for a structured light system such as the Minolta Vivid 900 [1]. Larger depth of field would lead to more flexible use in application.

There is some evidence suggesting that recognition algorithms might benefit from 3D depth resolution accuracy below 1mm [8]. Many 3D sensors do not have this accuracy in depth resolution.

Lastly, the image acquisition time for the 3D sensor should be short enough that subject motion is not a significant issue. Acquisition time is generally a more significant problem with structured-light systems than with stereo systems. It may be less of an issue for authentication type applications in which the subjects can be assumed to be cooperative, than it is for recognition type applications.

Considering all of these factors related to 3D sensor technology, it seems that the optimism some- times expressed for 3D face recognition relative to 2D face recognition is still somewhat premature. The general pattern of results in the multi-modal 3D+2D studies surveyed here suggests that 3D face recog- nition holds the potential for greater accuracy than 2D. And existing 3D sensors are

3D allows greater accuracy than 2D also suggest that multi-modal recognition allows greater accuracy than either modality alone. Thus the appropriate issue may not be 3D versus 2D, but instead the best method to combine 3D and 2D.

Improved Algorithms.

One limitation to some existing approaches to 3D face recognition involves sensitivity to size variation. Approaches that use a purely curvature-based representation, such as extended Gaussian images, are not able to distinguish between two faces of similar shape but different size. Approaches that use a PCA-based or ICP-based algorithm can handle size change between faces, but run into problems with change of facial expression between the enrollment image and the image to be recognized.

Approaches that effectively assume that the face is a rigid object will not be able to handle expression change. Handling change in facial expression would seem to require at least some level of part-whole model of the face, and possibly also a model of the range of possible non-rigid motion of the face. The seriousness of the problem of variation in facial expression between the enrollment image and the image to be recognized is illustrated in the results shown in Figure 3. This experiment focuses on the effects of expression change. Seventy subjects had their gallery image acquired with “normal expression” one week, a first probe image acquired with “smiling expression” in another week, and a second probe image acquired with “normal expression” in still another week. Recognition was done with PCA-based 2D and 3D algorithms [8]. The upper CMC curves represent performance with time-lapse only between gallery and probe; the lower pair represents time lapse and expression change. With simple time lapse but no expression change between the gallery and probe images, both 3D and 2D result in a rank-one recognition rate around 90%. There is noticeable drop in performance when expression variation is introduced, to 73% for 2D and 55% for 3D. In this case, where the 3D recognition algorithm effectively assumes the face is a rigid shape, 3D performance is actually more negatively affected by expression change than is 2D performance. The relative degradation between 3D and 2D appears not to be a general effect, but instead is

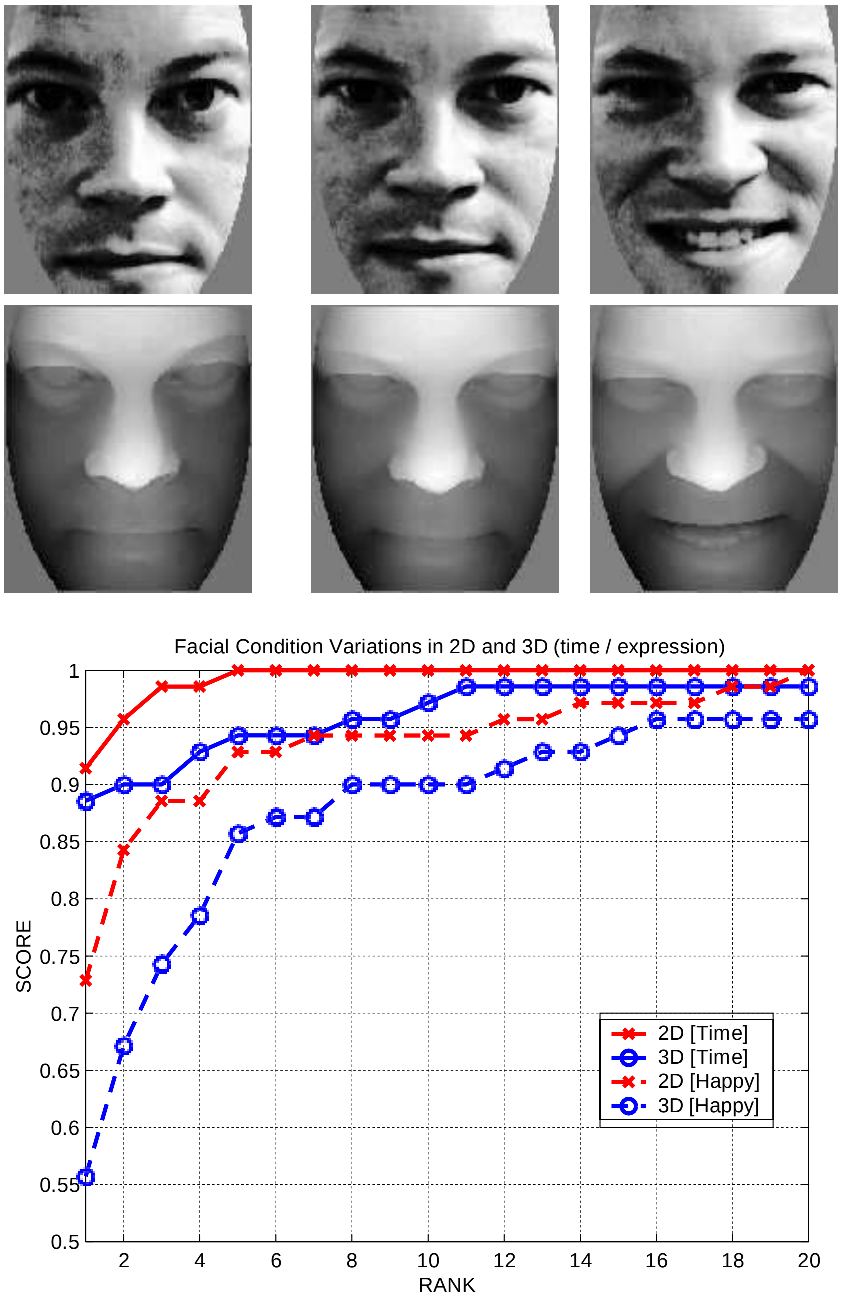

Figure 3. Effects Of Expression Change On 3D And 2D Recognition Rates.

In addition to a need for more sophisticated 3D recognition algorithms, there is also a need for more sophisticated multi-modal combination. Multi-modal combination has so far taken a fairly simple ap- proach. The 3D recognition result and the 2D recognition result are each produced without reference to the other modality, and then the results are combined in some way. It is at least potentially more powerful to exploit possible synergies between the the two modalities in the interpretation of the data. For example, knowledge of the 3D shape might help in interpreting shadow regions in the 2D image. Similarly, regions of facial hair might be easy to identify in the 2D image and help to predict regions ofthe 3D data which are more likely to contain artifacts.

Improved Methodology and Datasets.

One barrier to experimental validation and comparison of 3D face recognition is lack of appropriate datasets. Desirable properties of such a dataset include: (1) a large number and demographic variety of people represented, (2) images of a given person taken at repeated intervals of time, (3) images of a given person that represent substantial variation in facial expression, (4) high spatial resolution, for example, depth resolution of 1 mm or better, and (5) low frequency of sensor-specific artifacts in the data. Expanded use of common datasets and baseline algorithms in the research community will facilitate the assessment of the state of the art in this area. It would likely also improve the interpretation of research results if the statistical significance, or lack thereof, was reported for observed performance differences between algorithms and modalities. Another aspect of improved methodology would be the use, where applicable, of explicit, distinct training, validation and test sets. For example, the “face space” for a PCA algorithm might be created based on a training set of images, the number of eigenvectors used and the distance metric used then selected based on a validation set, and finally the performance estimated on atest set. The different sets of images would be non-overlapping with respect to the persons representedin each.

A more subtle methodological point is involved in the comparison of multi-modal results to baseline results from a single modality. In the context of this survey, there are several publications that compare the performance of multi-modal 3D+2D face recognition to the performance of 2D alone. The multi- modal 3D+2D performance is always observed to be greater than the performance of 2D alone. However, this comparison is too simple, and is effectively biased toward the multi-modal result. Enrolling a subject in a multi-modal system requires two images, a 3D image and a 2D image. The same is true of the information used to recognize a person in a multi-modal system. Therefore, a more appropriate comparison would be to a 2D recognition system that uses two images of a person both for enrollment and for recognition. When this sort of controlled comparison is done, the differences observed for multi-modal 3D+2D compared to "multi-sample" 2D are smaller than those for a comparison to plain 2D.

Summary.

As evidenced by the publication dates in Table 1, activity in 3D and multi-modal 3D+2D face recogni- tion has expanded dramatically in recent years. It is an area with important potential applications. At the same time, there are many challenging research problems still to be addressed. These include the development of more practical and robust sensors, the development of improved recognition algorithms, and the pursuit of more rigorous experimental methodology. The development of improved recognition al- gorithms will be spurred by more rigorous research methodology, involving larger and more challenging datasets, and more carefully controlled performance evaluations.

4 Acknowledgments

This work was supported by the DARPA Human ID program through Office of Naval Research N000140210410 and by the National Science Foundation through EIA 01-30839.

References

Konica Minolta 3D digitizer. *available at www.minoltausa.com/vivid/*, January 2004.

B. Achermann and H. Bunke. Classifying range images of human faces with Hausdorff distance. *15-th International Conference on Pattern Recognition*, pages 809–813, September 2000.

B. Achermann, X. Jiang, and H. Bunke. Face recognition using range images. *International Conference on Virtual Systems and MultiMedia*, pages 129–136, 1997.

C. Beumier and M. Acheroy. Face verification from 3D and grey level cues. *Pattern Recognition Letters*,22:1321–1329, 2001.

V. Blanz and T. Vetter. Face recognition based on fitting a 3D morphable model. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25:1063–1074, September 2003.

A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Expression-invariant 3D face recognition. *Audio- and Video-Based Person Authentication (AVBPA 2003)*, LCNS 2688, J. Kittler and M.S. Nixon, eds.:62–70, 2003.

J. Y. Cartoux, J. T. LaPreste, and M. Richetin. Face authentication or recognition by profile extraction from range images. *Proceedings of the Workshop on Interpretation of 3D Scenes*, pages 194–199, November 1989.

K. Chang, K. Bowyer, and P. Flynn. Face recognition using 2D and 3D facial data. *2003 Multimodal User Authentication Workshop*, pages 25–32, December 2003.

G. Gordon. Face recognition based on depth and curvature features. *Computer Vision and Pattern Recogni- tion (CVPR)*, pages 108–110, June 1992.

C. Hesher, A. Srivastava, and G. Erlebacher. A novel technique for face recognition using range images.
Seventh Int'l Symposium on Signal Processing and Its Applications, 2003.

S. Lao, Y. Sumi, M. Kawade, and F. Tomita. 3D template matching for pose invariant face recognition using 3D facial model built with iso-luminance line based stereo vision. *International Conference on Pattern Recognition (ICPR 2000)*, pages II:911–916, 2000.

J. C. Lee and E. Milios. Matching range images of human faces. *International Conference on Computer Vision*, pages 722–726, 1990.

Y. Lee, K. Park, J. Shim, and T. Yi. 3D face recognition using statistical multiple features for the local depth information. *16th International Conference on Vision Interface*, available at www.visioninterface.org/vi2003, June 2003.

G. Medioni and R. Waupotitsch. Face recognition and modeling in 3D. *IEEE International Workshop on Analysis and Modeling of Faces and Gestures (AMFG 2003)*, pages 232—233, October 2003.

K. Messer, J. Matas, J. Kittler, J. Luettin, and G. Maitre. XM2VTSDB: the extended M2VTS database. *Second International Conference on Audio- and Video-based Biometric Person Authentication*, pages 72–77, 1999.

J. Min, K. W. Bowyer, and P. Flynn. Using multiple gallery and probe images per person to improve performance of face recognition. *Notre Dame Computer Science and Engineering Technical Report*, 2003.

A. B. Moreno, Ángel Sánchez, J. F. Vélez, and F. J. Díaz. Face recognition using 3D surface-extracted
descriptors. *Irish Machine Vision and Image Processing Conference (IMVIP 2003)*, September 2003.

T. Nagamine, T. Uemura, and I. Masuda. 3D facial image analysis for human identification. *InternationalConference on Pattern Recognition (ICPR 1992)*, pages 324–327, 1992.

P. J. Phillips, P. Grother, R. J. Michaels, D. M. Blackburn, E. Tabassi, and J. Bone. FRVT 2002: overview and summary. *available at www.frvt.org*, March 2003.

H. T. Tanaka, M. Ikeda, and H. Chiaki. Curvature-based face surface recognition using spherical correlation principal directions for curved object recognition. *Third International Conference on Automated Face andGesture Recognition*, pages 372–377, 1998.

F. Tsalakanidou, D. Tzocaras, and M. Strintzis. Use of depth and colour eigenfaces for face recognition.
Pattern Recognition Letters, 24:1427–1435, 2003.

Y. Wang, C. Chua, and Y. Ho. Facial feature detection and face recognition from 2D and 3D images. *Pattern Recognition Letters*, 23:1191–1202, 2002.

W. Zhao, R. Chellappa, and A. Rosenfeld. Face recognition: a literature survey. *ACM Computing Surveys*, 35:399–458, December 2003.