3D FACE RECOGNITION

Charles Beumier

Charles.beumier@elec.rma.ac.be
Royal Military Academy, Elec. Eng. Dept., 30 Avenue de la Renaissance, 1000 Brussels, Belgium

ABSTRACT

Face recognition has recently received much attention as a biometric security means well accepted by users. 3D face recognition, exploiting the geometry of the facial surface, addresses the two most sensitive aspects of face recognition: viewpoint and illumination changes. This paper presents a 3D face recognition approach based on the geometrical comparison of a set of profiles. These profiles are obtained as parallel planar cuts of the facial surfaces. In a first step, normalised profiles are extracted from each face independently thanks to the intrinsic symmetry of faces. The homologue profiles of test and reference surfaces are then matched two by two. The individual distances of profile pairs are combined into a global distance. The global distance after optimisation is used as criterion to state the similarity of faces. Results are presented for a population of 81 persons of a 3D face database of intermediate quality.

1. INTRODUCTION

In the field of security, biometrics has gained a lot of popularity. It aims at replacing the traditional password or PIN code (Personal Identification Number) that are subject to loss or theft. Many biometric techniques are technically proven and commercialised, but they lack the acceptance by the users who find them intrusive.

In this context, face recognition is well accepted by users and appears as an appropriate alternative, provided that sufficient recognition rates are guaranteed. Unfortunately, face recognition from 2D images suffers from the negative influences of viewpoint and illumination changes on performances. 3D face recognition addresses these problems through the exploitation of facial surfaces, little dependent on viewpoint and illumination, and offering information complementary to the grey-level analysis.

The paper is organised as follows. Section 2 briefly presents the prototype that has been developed and used to capture facial surfaces for the BIOMET [1] project. Section 3 describes the proposed geometrical approach for 3D face recognition. Section 4 presents the results and section 5 concludes the paper.

2. 3D ACQUISITION

2.1. Principle

3D acquisition with structured light consists (Figure 1) in the projection of a particular light pattern (in this case parallel colour stripes) on the 3D scene that is captured by a camera. Any point P (X,Y,Z) of the scene is localised in the space at the intersection of a line issued from the camera and passing through the image point p, and a plane corresponding to the projected stripe ‘s’ on point P.

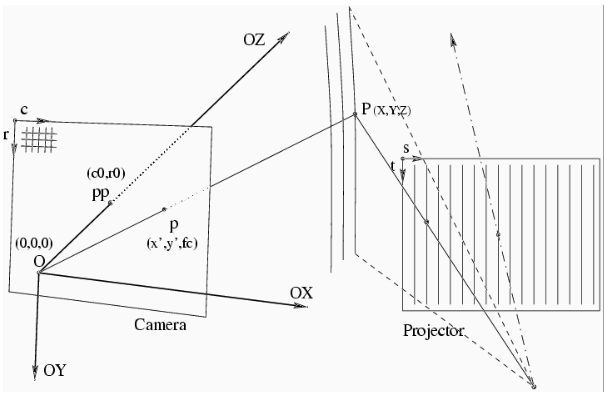

Figure 1: *structured light principle*

To obtain reliable measurements, the system is calibrated from a few presentations of a calibration object with known geometry. After this process, parameters concerning the camera, the projector and their arrangement are estimated so that an image position and a stripe index are correctly converted into the X, Y, Z coordinates of the 3D point.

Acquiring the 3D surface of an object consists in localizing points in the image with the index of the projected stripe. For the system referred to in this paper (and described in [2]), stripes are localised by edge detection and indexed from the colour of a few stripe neighbours whose colour sequence appears only once in the colour sequence of the projected pattern.

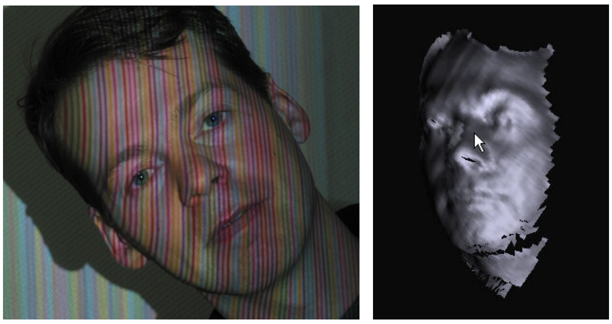

Figure 2: Face with projected stripes and extracted 3D surface

2.2. Prototype

The camera canon G2 and the projector, built from mechanical and optical elements of a low-cost slide projector, have been fixed on a stand to ensure that the system remains calibrated. The basis is rotated 40° to reduce the interference of the vertically projected stripes with the face features like the eyebrows and the mouth, while maintaining a similar density of stripes on the left and right parts of the face.

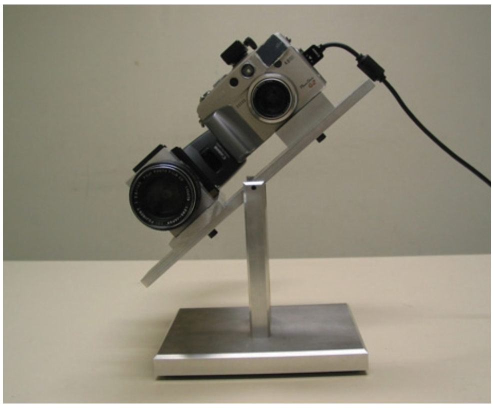

Figure 3: Acquisition prototype

The powerful (initially 150 W) incandescent lamp has been replaced by a flash lamp that does not require a fan cooler. The powerful instantaneous flash avoids motion blur and allows for larger depth of field.

A PC controls the acquisition through a USB port. Camera settings like the focal length and the exposure are accessible by a Canon driver program that allows for image storage and proper initialisation at the next start up of the camera. When an image is to be captured, the driver sends a request to the camera that triggers the flash, grabs the picture and sends it to the PC through the USB bus.

2.3. Database

The prototype has been used during the third campaign of the BIOMET project [1] that aimed at the collection and exploitation of biometric information like voice, fingerprint, face and signature. 81 persons showed up during the third campaign and six 3D shots of each person were taken with small rotation changes (frontal, left, right, up, down, and frontal poses).

To assess the quality of this database, each person was assigned one of the three classes: reject, fair, good. See an example for each class in Figure 4. About one person out of five belongs to class 'reject' in that database.

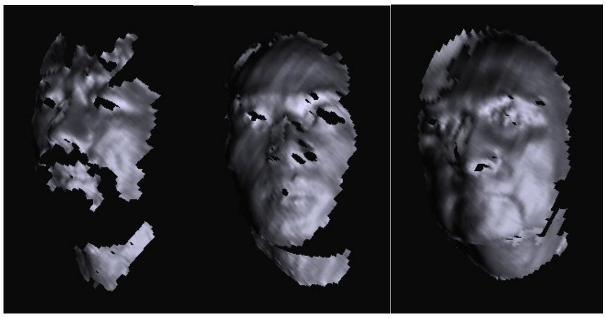

Figure 4: Reject, fair and good captures

The major limitation of the database is that it contains only one session. The major defect is due to the lack of contrast during capture, which reduced the area of capture to the surface sufficiently perpendicular to the system. We can expect from structured light to acquire more of the facial surface in one shot, and a better fidelity around the nose and eyes.

3. FACE RECOGNITION

3.1. Introduction

In our early developments about 3D face comparison [3], we rejected curvature analysis that highlighted the noise present in our 3D face capture. We preferred a geometrical approach, based on the Euclidean distance, which also has the advantage to show the visual quality of 3D acquisition and matching. Facial surfaces were compared along planar curves (profiles) obtained by surface cuts with parallel planes.

The recognition of a test face is assessed through its comparison with one or several references. Each comparison gives an error that is the minimum distance between the test and reference facial surfaces when rotation and translation parameters are optimised. If this error is below a given threshold, the test face is declared to be from the reference person.

Two types of error can occur: the test face of a client is rejected or the test face of an impostor is accepted. Expressed as a percentage of tests, we talk respectively about False Rejection Rate (FRR) and False Acceptance Rate (FAR). Each recognition system expresses a trade-off between low FAR and low FRR (see the ROC curve in Figure 5). Depending on the application, one can change the threshold to favour either low FAR or low FRR. In our results, we present Equal Error Rates, corresponding to threshold values for which FAR is equal to FRR.

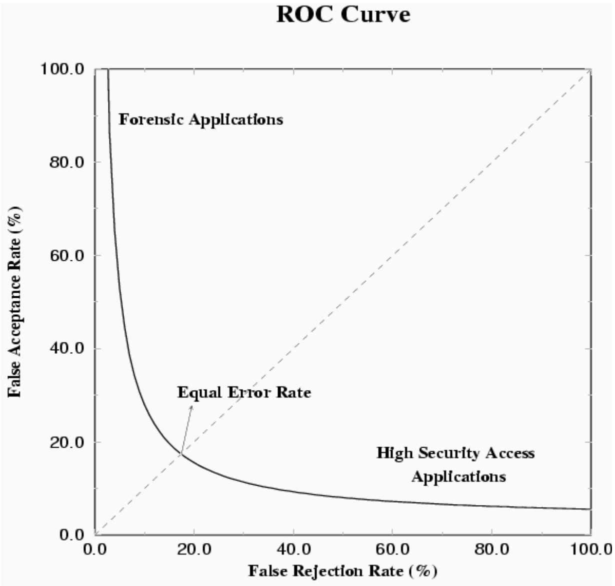

Figure 5: ROC curve and Equal Error Rate

3.2. Face Normalisation

Normalised curves are first extracted from the facial surfaces thanks to the natural symmetry of the faces. Parallel planes, distant of 1 cm, are adapted to capture up to 15 planar profiles that must be symmetric relatively to a vertical plane passing through the nose. Even if faces are not perfectly symmetric, a clear optimum exists. From the three rotation and three translation parameters, only three parameters intervene in the optimisation, leading to rapid processing.

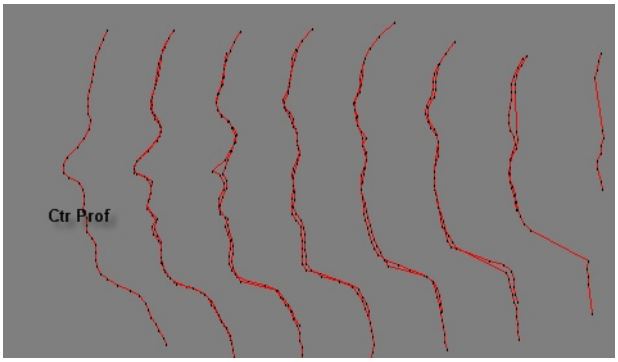

Figure 6: Normalisation of profiles from symmetry

3.3. Profile Comparison

Our face comparison approach issues a global distance from the matching of corresponding profiles extracted and normalised as explained in the previous section. Each profile is converted into a curve by computing the slope along the profile for two points separated by 4 cm. A rotation of the profile in its plane corresponds to a shift of the curve. The measure of similarity between two curves is given by the standard deviation of slope differences along the curves. This measure is independent from the rotation of the profiles. Only one translation parameter (shift along the curve) has to be tuned to find the optimum.

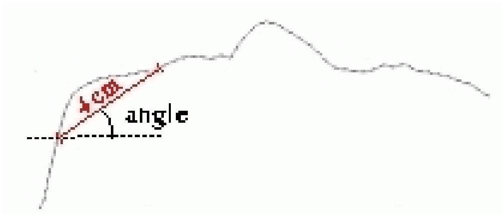

Figure 7: Slope along profile

3.4. Facial Surface Comparison

The global distance between two facial surfaces is computed as the mean value of the comparison scores of corresponding slope curves. If the global distance is lower than a threshold, the two faces are declared similar. The next section presents results with the threshold chosen to get equality between FAR and FRR, leading to the EER (Equal Error Rate).

4. RESULTS

The presented 3D face comparison algorithm has been tested on the BIOMET 3D database. As we dispose of 6 shots for nearly every person, we analysed the difference in recognition performance when shots of different orientations are considered.

	Client tests	Impostor tests	EER
Shots 123-456	651	51324	14.1 %
Shots 12-56	268	16688	8.9 %
Shots 1-6	70	5546	3.6 %

Table 1: EER for various shots of the BIOMET Database

Based on the results presented in Table 1, we see the good performance of comparing frontal shots (1-6). The EER of 3.6 % is particular promising, provided that nobody was rejected from the tests and that the 3D database is of median quality. The performance decreases when shots with different orientations are compared mainly because the contrast of the stripes was limited during the acquisition campaign. The area of facial capture was limited so that common area for facial surface comparison is limited for different orientations.

5. CONCLUSION

We presented a 3D face recognition approach based on the geometrical comparison of corresponding 2D profiles obtained by parallel cutting planes. The 6-dimensional matching problem has been split into a 3-dimensional normalisation that exploits the intrinsic facial symmetry and a 1-dimensional planar curve comparison procedure based on the local slope along profile (1 parameter).

In the results obtained for the BIOMET database, the recognition rate for frontal shots is very good. The comparison of non-frontal views is affected by the reduced common part of the face in face captures, especially for the BIOMET database that suffered from a reduction of facial coverage due to the lack of stripe contrast.

Generally speaking, performance would benefit from a larger representation of each person, disposing of more shots. Texture information, not yet integrated in BIOMET 3D tests, is expected to largely increase recognition performance.

6. REFERENCES

[1] S. Garcia-Salicetti, C. Beumier, G. Chollet, B. Dorizzi, J. Leroux-les-jardins, Jan Lunter, Yang Ni, D. Petrovska-Delacrétaz, "BIOMET: A Multimodal Person Authentication Database Including Face, Voice, Fingerprint, Hand and Signature Modalities", AVBPA 2003, Guildford, UK, June 9-11, 2003.

[2] C. Beumier, "Facial Surface Acquisition with Colour Striping", *ICANN/ICONIP 2003*, Istanbul, Turkey, pp. 548-551, June 2003.

[3] C. Beumier, M. Acheroy, "Automatic Face Authentication from 3D Surface", In *British Machine Vision Conference BMVC 98*, Univ. of Southampton UK, pp. 449-458, September 1998.