Fusion of Likelihood Ratio Classifier with ICP-based Matcher for 3D Face Recognition

Berk Gokberk   Luuk J. Spreeuwers   Raymond Veldhuis

University of Twente
Department of Electrical Engineering
Signals and Systems Group
P.O. box 217, 7500 AE Enschede, The Netherlands
B.Gokberk@utwente.nl L.J.Spreeuwers@utwente.nl R.N.J.Veldhuis@utwente.nl

Abstract

Three-dimensional (3D) face recognition systems have started to become popular in biometric systems recently. This is due to several factors: i) facial shape characteristics contain discriminative information, ii) availability of practical 3D acquisition devices, and iii) invariance of 3D facial information to several factors such as illumination and pose changes. It has been shown that classical texture-based 2D face classifiers have difficulties in identifying faces under such conditions. Therefore, taking advantage of 3D facial shape information either alone or together with 2D modality is considered to be a viable solution under such circumstances. In this work, we propose a novel 3D face recognition system using the combination of different individual 3D face classifiers; namely Linear Discriminant-based (LDA) Likelihood Ratio classifier with the Iterative Closest Point-based (ICP) matching algorithm. Both systems operate on aligned and normalized 3D facial surfaces. Alignment phase of the proposed system carries out absolute alignment such that all faces are in a specific position and direction, with non-facial parts removed. LDA-based system uses absolutely aligned faces and produce similarity scores using Likelihood ratio-based classifier. However, the ICP-based classifier performs additional surface matching between absolutely aligned faces, which can be considered as relative alignment. After pair-wise alignment of gallery and probe faces, the ICP algorithm produces dissimilarity scores, by measuring the quality of surface registration. With the use of different matching algorithms and different alignment methods, our approach tries to minimize the shortcomings of each individual method. Finally, the scores obtained by 3D face recognizers are fused to improve the verification accuracy. Our preliminary experiments conducted on the subset of FRGC v2 3D face database show promising performance improvement in verification simulations.

1   Introduction

In recent years, three dimensional (3D) face recognition systems has caught considerable attention due to their advantages over classical 2D face recognition systems. Most importantly, 3D surface characteristics of a human face offers more discriminatory information than using only 2D facial texture information. With the availability of state-of-the-art 3D facial acquisition devices, 3D facial shape can be obtained accurately. The invariance of facial shape to varying illumination conditions and it's potential to easily compensate for in-depth rotation variations make the use of 3D facial information attractive in many biometrics applications.

Most of the proposed solutions for 3D face recognition systems rely on i) efficient and accurate alignment/registration of facial surfaces, and ii) surface matching techniques to infer the similarity of two faces. One of the standard techniques to achieve these goals is to use the Iterative Closest Point [3] (ICP) algorithm to align two facial surfaces and to output the quality of alignment as the similarity score which then can be used as a matching criterion. Another popular class of algorithms use the projections of the aligned surfaces into 2D images, referred to as depth or range images, and then use classical image representation techniques to classify faces. Statistical feature extraction techniques such as Principal Component Analysis, Linear Discriminant Analysis (LDA) and Independent Component Analysis are frequently used with depth image representations [2].

It has been shown that the most critical part of any 3D face recognition systems is the accurate alignment and registration of human faces. Therefore most of the scientific research in the 3D face recognition systems has been carried out for finding precise alignment of faces. Since human faces resemble each other, as opposed to 3D object identification problem where the object classes are distinctively different from each other, transforming faces into a same coordinate system is crucial to extract discriminative information by the pattern classifier. In this paper, we propose a 3D face recognition system that i) encapsulates two different registration methods, absolute and relative registration, and ii) uses two different powerful pattern classifiers, Likelihood Ratio classifier and Linear Discriminant Analysis, to compute the similarity of any given face pairs. Proposed system automatically carries out initial absolute alignment, i.e., transforming faces into the same coordinate system and normalizes facial surface information, i.e., removing noise, filling incomplete data, and extracting the central facial region. Likelihood ratio-based classifier is used to classify absolutely aligned faces. The complementary pipeline encompasses relative registration via the ICP approach to obtain a finer registration between a given pair of facial surfaces after the absolute alignment phase and Linear Discriminant Analysis to extract discriminative features. Finally, the two similarity scores obtained from these two approached are combined at the fusion phase to achieve better recognition rates.

2   Proposed 3D Face Recognition System

The proposed 3D face recognition system is composed of the following steps: 1) Initial absolute alignment of faces, 2) two different facial matchers (likelihood ration classifier and ICP-based Linear Discriminant Analysis) working in parallel and 3) the score-level fusion algorithm. The overall diagram of our 3D face recognition system is illustrated in Figure 1.

2.1   Absolute Alignment

In the first phase, face normalization and absolute alignment, the aim is to translate, rotate and scale facial surfaces to a canonical position. In order to achieve absolute alignment, central vertical symmetry plane of a facial surface is automatically found by a multi-resolution (from low to high resolution) approach. Nose tip point is also automatically localized which enables us to translate 3D facial point cloud to the origin of the coordinate system. Face normalization algorithm also comprises of automatic noise removal and hole filling. Due to the acquisition sensor, 3D human facial surfaces mostly contain perturbations such as small surface bumps and also impulse-like spikes. Local averaging filters and median filters are used to remove such surface errors. In most of the 3D facial sensors operating by the principle of laser beams, non-reflective dark facial regions usually do not provide depth measurements. Due to this reason, some parts of the acquired faces such as eyes and eyebrows contain holes. If holes
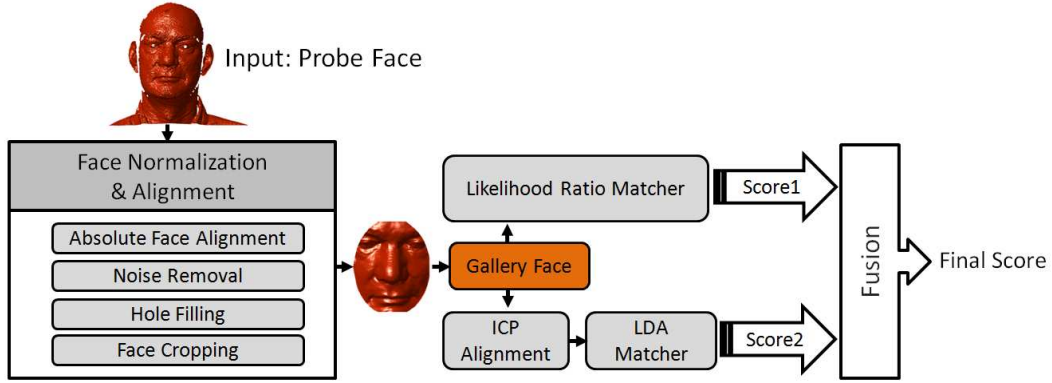
Figure 1: Overall diagram of the proposed approach.

are relatively small, they can be filled my filtering operations, such as using median filter. However, when they are large, then filtering operations may not be sufficient. Therefore, in our system, these type of situations are dealt with by using the facial symmetry property. In the proposed approach, since the vertical symmetry plane is computed automatically, we fill such large holes with the information present in the mirror side of the facial surface. Lastly, using the automatically determined nose tip location, we keep the central facial region by cropping out the other parts by a 3D spherical mask. Figure 2 shows an input raw facial surface, its cropped version, and the final facial surface after the smoothing and hole filling operations.
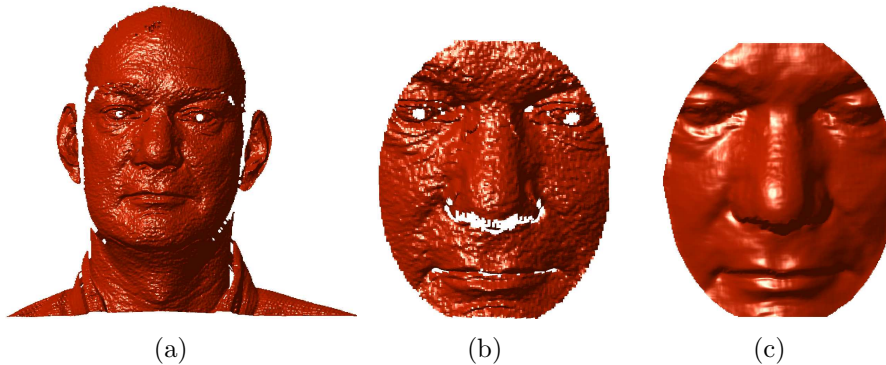

| (a) | (b) | (c) |

Figure 2: a) Input face, b) cropped/normalized face, and c) smoothed and filled facial image.

2.2 Likelihood Ratio Classifier

Biometric verification can be considered as a two class pattern recognition problem to separate genuine and impostor classes. Likelihood Ratio (LLR) classifier finds a linear decision boundary under the assumption that these two classes are Gaussian[4]. During training, each user i is modeled by his/her mean feature vector, μ_i, and covariance matrix Σ_i. The impostor class is represented by a global mean and covariance matrix, μ_I and Σ_I. Using the log-likelihood ratio, the matching score d of a given test sample x to the claimed identity i can be computed by: $d = -(x - \mu_i)^T \Sigma_i^{-1}(x - \mu_i) + (x -$

$\mu_I)^T \Sigma_I^{-1}(x - \mu_I)$. Given a verification threshold t, the user can either be accepted as a genuine user or rejected as an impostor. In order to make the estimation of user specific mean and covariance matrices easier, we assumed that all users share a common covariance matrix and applied dimensionality reduction by the Principal Component Analysis and Linear Discriminant Analysis methods.

2.3   Linear Discriminant Analysis after the ICP Alignment

In our system, likelihood ratio classifier operates on absolutely aligned depth images. Although absolute alignment establishes sufficient registration, i.e., correspondence between the related depth measurements across faces, it is also worthwhile to carry out pairwise (relative) alignment between face pairs. The intuition here is that with pairwise registration 3D facial surface information coming from the same identity may align better to each other compared to the others at the expense of more computational load. Computational load is mainly due to the need of registering the probe face to each gallery face. However, in a typical verification scenario, only one registration with the enrolled template is needed as opposed to identification. In our approach, we choose to use the Iterative Closest Point algorithm for pairwise rigid registration. The ICP algorithm finds the best rotation and translation parameters to align one surface to the other. If the facial surfaces do not exhibit significant deformations such as due to extreme expressions, the ICP approach performs sufficiently well. Given a test face and the claimed identity, we use the absolutely aligned facial surfaces as inputs to the ICP algorithm. Since both the probe and gallery face of the claimed identity is coarsely aligned beforehand, the convergence of the ICP algorithm is usually fast. After aligning the probe face to the gallery face, we compute the depth image and use Linear Discriminant Analysis to extract discriminative facial features.

2.4   Score-Level Fusion by the Logistic Discriminant

As a fusion algorithm, we choose to use a linear discriminant-based approach, namely the logistic discriminant. The attractive point of using a linear discriminant approach in fusion is to bypass the estimation of posteriors of the classes[5]. It is usually sufficient to estimate the parameters of the discriminant directly. Although non-linear discriminants can also be used, we choose to employ a linear decision boundary because of its effectiveness and simplicity. In logistic discrimination, the ratio of the class-conditional densities are are modeled, specifically, it is assumed that the log likelihood ratio is linear. The parameters of the logistic discriminant are learned by a gradient-descent mechanism. In a verification setting, the two classes to be separated correspond to the genuine and impostor cases. At the training phase, a linear separating line is learned by a training set of impostor and genuine scores. At the verification phase, the output of the logistic discriminant measures how likely a given test face is coming from the genuine and impostor classes. A typical decision line learned by the logistic discriminant is shown in Figure 4.

3   Experimental Results

We have tested our algorithm on a subset of the FRGC v.2 3D face database [6]. In the FRGC v.2 database, there are 4,007 face scans of 465 subjects. 3D shape data contains 30,000 to 40,000 3D coordinates. Although the quality of the scanned data is high, there are two types of noise affecting 3D faces: small protrusions and impulse noise-like jumps. In our experiments, we have used the Spring 2003 part of the FRGC dataset for the training of the Likelihood ratio classifier and Linear Discriminant Analysis-based classifier. Spring 2003 part of the FRGC database contains 943 3D scans of 275

subjects. An independent test, having 977 scans of the 100 subjects, is used to test the verification accuracy of the proposed 3D verification system. We have compared every possible pair of facial images in the test set, producing 476,776 comparisons. There are 6,092 genuine and 470,684 impostor comparisons in our experimental protocol.

Figure 3 shows the score distributions obtained by the LLR and ICP-based matchers. Figure 3(a) LLR and ICP scores are shown together in a 2D plot for genuine (red dots) and impostor (blue dots) comparisons. In order to better visualize the overlapping regions of the genuine and impostor distributions, score distributions of the LLR and ICP matchers are shown independently in Figure 3(b) and Figure 3(c), respectively. Corresponding to this experimental setup, the Equal Error Rates (EER) obtained by the LLR and ICP-based system are 4.45% and 2.71%, respectively. Based on this results, it is seen that improving the alignment by the ICP-based relative registration improves the accuracy of the verification system.
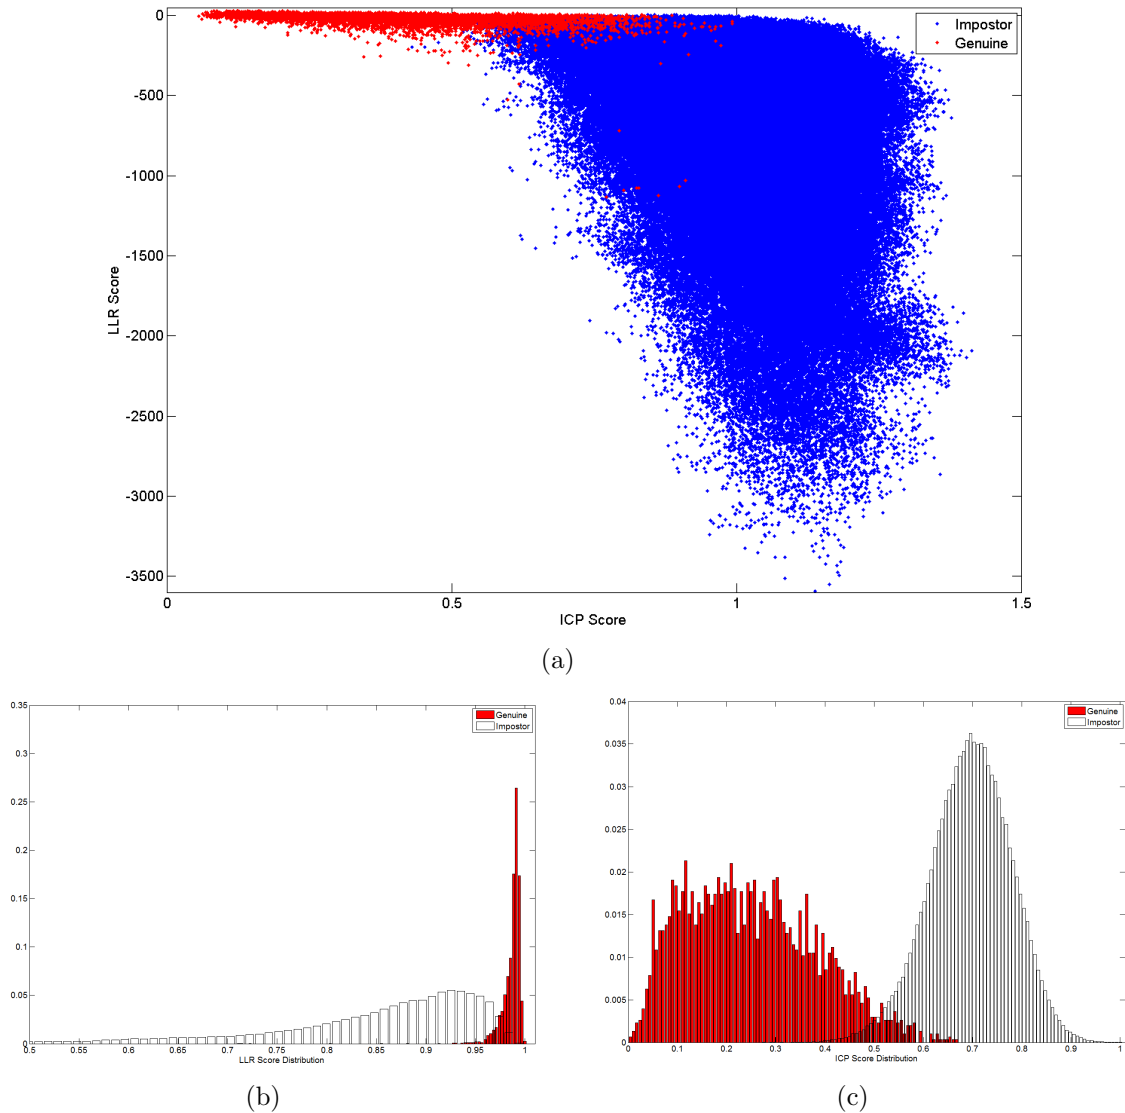

(a)

(b)

(c)

Figure 3: a) LLR and ICP scores in two dimensions, b) LLR scores, and c) ICP scores.

In the second part of our experiments, we investigate the effect of score-level fusion of LLR and ICP-based matchers. Since the logistic discriminant-based fusion algo-

rithm needs training before the fusion at the verification phase, we have formed two independent sets from the test part of the FRGC v.2 experiments. To train the fusion algorithm, we have used half of the genuine and impostor comparisons, i.e., 3,046 genuine and 235,342 comparisons are used for that purpose. The remaining genuine and impostor comparisons are used to test the fusion accuracy. In this particular fusion experimental protocol, LLR and ICP-based matchers obtained 4.40% and 2.62% EER which are similar to the EERs obtained by using whole genuine and impostor comparisons. The LLR and ICP scores for this experiment together with the decision line found by the logistic discriminant-based fusion algorithm (black line) are shown in Figure 4. As seen from Figure 4, decision line better separates the genuine and impostor classes compared to using each matcher alone. This results is confirmed by the EER obtained by the fusion system which is 1.54%. In order to analyze the performance of the proposed scheme under different operating points, i.e., at different FAR and FRR points, it is better to look at the Receiver Operating Characteristic (ROC) curves of the verification algorithms. In Figure 5, the ROC curves of the LLR, ICP, and the fusion algorithm are shown. Visual inspection of the ROC curve reveals that in all operating regions, fusion algorithm improves the performance of both LLR and ICP-based systems.
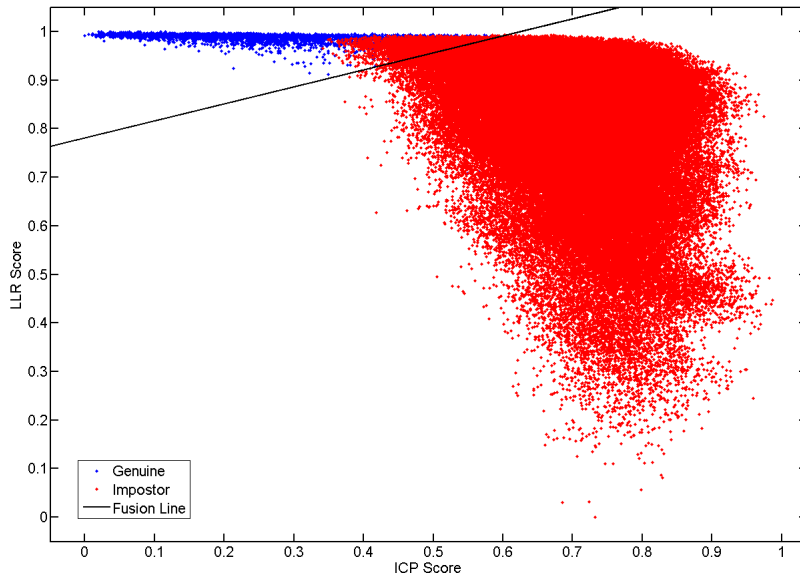

Figure 4: Decision line found by the fusion algorithm.

4 Conclusion

In this paper, fusion of two different 3D face recognition methods has been proposed. The first matching engine employs Likelihood Ratio classifier where the 3D facial surfaces are represented as depth images. Prior to verification, an automatic facial surface alignment method is used to transform faces into a canonical position. This absolute alignment step together with the use of LLR classifier has shown to be very effective in the verification experiments conducted on a subset of the FRGC v.2 database, obtaining 4.45% EER. The second matching engine further improves the registration of facial image pairs by the ICP algorithm and employs LDA-based classifier to compute the similarity scores. Experimental findings showed that the second matching engine
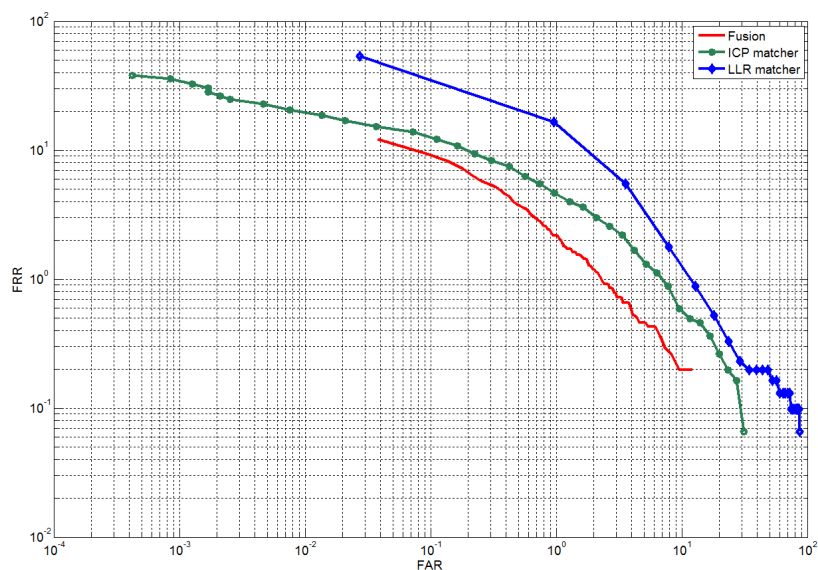
Figure 5: ROC curves for the LLR (blue), ICP (green) and the fusion system (red).

performs accurately by achieving an EER of 2.71%. Combining the two matchers at the score level by a linear discriminant-based approach finally produces the best verification accuracy on the FRGC v.2 database, producing an EER of 1.54%. As future work, it is worthwhile to investigate other fusion mechanisms different than the linear discriminants.

References

[1] Bowyer, K. W., Chang, K., Flynn, P., "A survey of approaches and challenges in 3D and multi-modal 3D + 2D face recognition," Computer Vision and Image Understanding, Vol.101, pp. 1-15, 2006.

[2] Gokberk, B., Dutagaci, H., Ulas, A., Akarun, L., Sankur, B., "Representation Plurality and Fusion for 3D Face Recognition, " IEEE Transactions on Systems, Man, and Cybernetics-Part B: Cybernetics, Vol.38, Issue.1, 2008.

[3] Besl, P., McKay, N., "A Method for Registration of 3D Shapes," IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol.14, pp. 239-256, 1992.

[4] Veldhuis, R., Bazen, A., Kauffman, J., Hartel, P., "Biometric verification based on grip-pattern recognition," SPIE, 2004.

[5] Alpaydin, E., "Introduction to Machine Learning," The MIT Press, October 2004, ISBN 0-262-01211-1.

[6] Phillips, J. P., Flynn, P. J., Scruggs, T., Bowyer K. W., Chang, J., Hoffman, K., Marques, J., Min, J., Worek, W., "Overview of the face recognition grand challenge," International Conference on Computer Vision and Pattern Recognition, pp. 947-954, 2005