SHREC'08 Entry: 3D Face Recognition using Facial Contour Curves

Frank B. ter Haar* Remco C. Veltkamp†

Department of Information and Computing Sciences, Utrecht University, the Netherlands

ABSTRACT

In this work we compute the similarity of 3D faces using a set of eight contour curves. These contours were selected and matched using our 3D face matching framework. In previous work, we performed extensive research to the selection of distinctive facial curves for 3D face matching. To relate the performance of several of these curves to other face matching methods, we participated the SHape REtrieval Contest (SHREC) of 3D Face Scans. Within this contest we have used a set of eight C-contours and tested their face retrieval performance using two different distance measures. In an attempt to increase the expression invariance of these curves, we employed our 3D face matching framework to match either 100% of the selected features or the subset of the best 60% of the selected features. Results show that the selected distance measure can have a great influence on the distinctiveness of facial curves. In case of large variations in facial expressiveness, the subset of the best 60% of the features increases the overall performance. With a recognition rate of 91.1% and a mean average precision of 0.693 our method performs reasonably well compared to other methods.

1 INTRODUCTION

The first 3D face matching algorithms mainly focussed on comparing face scans with neutral expressions. However, for a 3D face matching algorithm to be widely applicable, expression invariant algorithms were necessary. With an emphasis on how to match a 3D neutral scan with a 3D expression scan, algorithms became more advanced. Different sets of various facial curves were investigated for their invariance to expressions [2, 3] and surface geodesics were put to use to obtain bending invariant representations [1]. With the continuously growing sizes of 3D face databases, such as the FRGC v.2 with over four thousand 3D face scans, the scalability of 3D face matching method is furthermore an important aspect. Our 3D face matching framework enabled the selection and comparison of different facial curves [3], allowing the selection of discriminative features for efficient face matching. In this work we briefly describe the settings of our framework, that we used to retrieve faces within the SHape REtrieval Contest of 3D Face Scans.

2 METHOD

Starting from the tip of the nose in a pose normalized face, our framework [3] extracts profile curves over the face surface in different directions. These sets of profile curves are used to determine the similarity of two faces. To match two profile curves, we match a set of samples along the curves. When combined, all samples with the same constraints builds up a face contour. Because our framework was implemented to work on surface meshes with a proper topology and without holes, we fitted a morphable face model [4] to each of the 427 face scans. After the fitting, we obtained a good model with the required properties for our face matching framework.

*e-mail: fhaar@cs.uu.nl

†e-mail: Remco.Veltkamp@cs.uu.nl

To compute the similarity of two faces $\mathbf{A}$ and $\mathbf{B}$, we extract N_c samples for each of the N_p profiles. Such a sample $\mathbf{A}_{ij}$ is defined as the intersection(s) of profile *i* and contour *j*. Because the profiles and their contour samples are extracted in a structured way, we can assume that these $N \leq N_p \cdot N_c$ samples correspond for faces $\mathbf{A}$ and $\mathbf{B}$. The distances between these correspondences introduce a dissimilarity. We use this information in a 3D face matching framework that consists of the generic formula

$$d(\mathbf{A},\mathbf{B}) = \frac{1}{N}\sum_{i=1}^{N_p} f_w \sum_{j=1}^{N_c} d_s(\mathbf{A}_{ij},\mathbf{B}_{ij}),$$

which must be instantiated with the following parameters:

- The number of profiles N_p.
- The number of contours N_c.
- The distance measure for two corresponding samples $d_s(\mathbf{A}_{ij},\mathbf{B}_{ij})$.
- The function f_w that determines the weight for a profile with N_c samples.

For the experiment we extracted from each face a set of N_p=90 profile curves. Along these curves we extracted N_c=8 C-samples, building a feature set of eight C-contours (720 sample points). Figure 1 shows the contour samples extracted from a face model that was fitted to scan data. We have used this set of features also in [3].

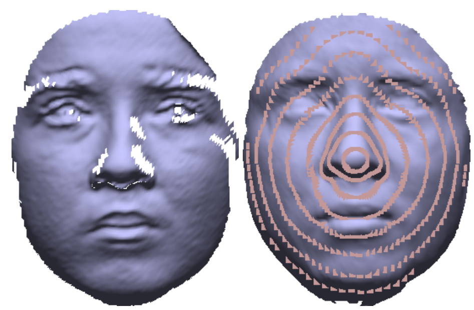

Figure 1: The set of eight contour curves extracted from a model (right) which was fitted to a scan (left). The line shown in black is one of the N_p=90 profile curves.

The information of each 3D face is now reduced to a small set of ($N_c \times N_p$) 3D sample points, with one-to-one correspondence to the same set of samples in a different face. In case two faces are identical, the distance of these extracted samples sum up to zero, otherwise they cause a dissimilarity. In our experiments we have used two different distance measures (d_{p1} and d_{p2}) within measure d_s.

$$d_s(\mathbf{A}_{ij},\mathbf{B}_{ij}) = min_{\forall p\in\mathbf{A}_{ij},\forall q\in\mathbf{B}_{ij}} d_p(p,q),$$

$$d_{p1}(p,q) = (e(p,p_{nt}) - e(q,p_{nt}))^2,$$

$$d_{p2}(p,q) = (e(p,q))^2,$$

where d_{p1} is the squared difference in Euclidean distance e to the tip of the nose (p_{nt}) and d_{p2} is the squared Euclidean distance.

	p	d_p	RR	MAP	MADP
ter Haar run1	100%	d_{p1}	76.1%	0.561	0.729
ter Haar run2	100%	d_{p2}	86.9%	0.668	0.809
ter Haar run3	60%	d_{p1}	81.5%	0.615	0.766
ter Haar run4	60%	d_{p2}	91.1%	0.693	0.823

Table 1: Retrieval results of our different runs

Our framework can be easily adapted, such that the matching of features becomes more invariant to facial expressions. By introducing function f_w, a weight can be assigned to different profile curves. In this experiment we use f_w to assign a weight of 1 to the p% best matching profiles and a weight of 0 to the other profiles. Values p=60 and p=100 were used. With this function, profiles in facial areas that changed because of an expression can be neglected for the actual face comparison.

3 Results

The results of four different runs on the SHREC'08 3D face database are shown in Table 1. These runs are based on the eight uniformly extracted C-contours and are matching according to,

1. p=100% of the profiles using measure d_{p1},
2. p=100% of the profiles using measure d_{p2},
3. p=60% of the profiles using measure d_{p1},
4. p=60% of the profiles using measure d_{p2}.

Based on the squared Euclidean distance (d_{p2}) the selected features perform approximately 10% better with recognition rates (RR) of 86.9% (*run2*) and 91.1% (*run4*) on this database. This can be explained as follows. When the curve distance of a profile on two different faces increases, the distance to the nose tip may remain the same while the Euclidean distance between the profiles does change. On the other hand, when the Euclidean distance between two samples is zero, the distance to the nose tip is zero too. Matching features using the Euclidean distance (d_{p2}) between samples is therefore more reliable, achieving high results.

The use of a smaller percentage of best matching profiles increases the recognition rates with approximately 5% to a maximum of 91.1% for *run4*. The matching of profiles in a more selective way allows the exclusion of data in facial regions which could have been changed by an expression. Since the SHREC'08 3D face database contains scans of subjects with various expressions, the matching of a subset of profiles improves the retrieval performance. Figure 2 shows an example of a difficult query (rank 1) which is recognized using p=60% of the profile curves and not recognized using p=100% of the profile curves.

To perform 427 times 427 face comparisons using the 720 predetermined correspondences, our face matching method requires less than one hundred seconds on a Pentium IV 2.8 GHz.

4 Concluding remarks

We have shown that a small set of contour curves can be used for 3D face retrieval in both an efficient and effective manner. Efficient because correspondence are determined before the actual 3D face matching, which allows the matching of 427 times 427 faces in less than one hundred seconds. Effective because a MAP up to 0.693 and a recognition rate up to of 91.1% were achieved resulting in a sixth place (out of 19) in the shape retrieval contest of 3D face scans. The different runs we experimented with, shows that C-contour curves are more distinctive when matched using the squared Euclidean distance compared to the squared difference in Euclidean distance to the nose tip. To increase the retrieval performance of these contour curves in case of facial expressions, our employed 3D face matching framework can be easily adapted to match only the samples that belong to a subset of best matching profile curves. In the future, we will investigate other types of contour sets and the optimal way to match their samples, for the task of expression invariant face retrieval.

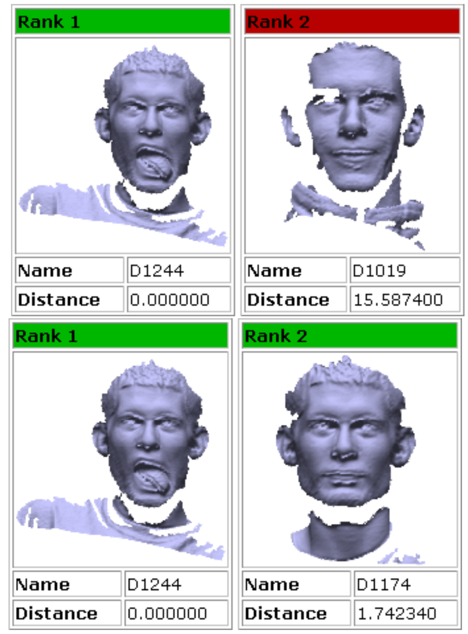

Figure 2: An arbitrary gesture not recognized using p=100% (left), but correctly recognized using p=60% (right)

References

[1] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Three-dimensional face recognition. *IJCV*, 64(1):5–30, 2005.

[2] C. Samir, A. Srivastava, and M. Daoudi. Three-Dimensional Face Recognition Using Shapes of Facial Curves. *TPAMI*, 28(11):1858–1863, 2006.

[3] F. B. ter Haar and R. C. Veltkamp. A 3D Face Matching Framework. In *Proc. Shape Modeling International (SMI'08)*, 2008.

[4] University of South Florida, Prof. Sudeep Sarkar. USF HumanID 3D Face Database.