

Face Recognition from Facial Surface Metric

Alexander M. Bronstein¹, Michael M. Bronstein¹,
Alon Spira², and Ron Kimmel²

¹ Technion - Israel Institute of Technology, Department of Electrical Engineering,
32000 Haifa, Israel

{alexbron, bronstein}@iee.org

² Technion - Israel Institute of Technology, Department of Computer Science,
32000 Haifa, Israel

{ron, salon}@cs.technion.ac.il

Abstract. Recently, a 3D face recognition approach based on geometric invariant signatures, has been proposed. The key idea is a representation of the facial surface, invariant to isometric deformations, such as those resulting from facial expressions. One important stage in the construction of the geometric invariants involves in measuring geodesic distances on triangulated surfaces, which is carried out by the *fast marching on triangulated domains* algorithm.

Proposed here is a method that uses only the metric tensor of the surface for geodesic distance computation. That is, the explicit integration of the surface in 3D from its gradients is not needed for the recognition task. It enables the use of simple and cost-efficient 3D acquisition techniques such as photometric stereo. Avoiding the explicit surface reconstruction stage saves computational time and reduces numerical errors.

1 Introduction

One of the challenges in face recognition is finding an invariant representation for a face. That is, we would like to identify different instances of the same face as belonging to a single subject. Particularly important is the invariance to illumination conditions, makeup, head pose, and facial expressions – which are the major obstacles in modern face recognition systems.

A relatively new trend in face recognition is an attempt to use 3D imaging. Besides a conventional face picture, three dimensional images carry all the information about the geometry of the face. The usage of this information, or part of it, can potentially make face recognition systems less sensitive to illumination, head orientation and facial expressions.

In 1996, Gordon showed that combining frontal and profile views can improve recognition accuracy [1]. This idea was extended by Beumier and Acheroy, who compared central and lateral profiles from the 3D facial surface, acquired by a structured light range camera [2]. This approach demonstrated some robustness to head orientations. Another attempt to cope with the problem of head pose was presented by Huang et al. using 3D morphable head models [3]. Mavridis

et al. incorporated a range map of the face into the classical face recognition algorithms based on PCA and hidden Markov models [4]. Their approach showed robustness to large variations in color, illumination and use of cosmetics, and it also allowed separating the face from a cluttered background.

Recently, Bronstein, Bronstein, and Kimmel [5] introduced a new approach which is also able to cope with problems resulting from the non-rigid nature of the human face. They applied the bending invariant canonical forms proposed in [6] to the 3D face recognition problem. Their approach is based on the assumption that most of human facial expressions are near-isometric transformations of the facial surface. The facial surface is converted into a representation, which is invariant under such transformations, and thus yields practically identical signatures for different postures of the same face.

One of the key stages in the construction of the bending invariant representation is the computation of the geodesic distances between points on the facial surface. In [5], geodesic distances were computed using the Fast Marching on Triangulated Domains (FMTD) algorithm [7]. A drawback of this method is that it requires a polyhedral representation of the facial surface. Particularly, in [5] a coded-light range camera producing a dense range image was used [9]. Commercial versions of such 3D scanner are still expensive.

In this paper, we propose 3D face recognition based on simple and cheap 3D imaging methods, that recover the local properties of the surface without explicitly reconstructing its shape in 3D. One example is the photometric stereo method, that first recovers the surface gradients. The main novelty of this paper is a variation of the FMTD algorithm, capable of computing geodesic distances given only the metric tensor of the surface. This enables us to avoid the classical step in shape from photometric stereo of integrating the surface gradients into a surface.

In Section 2 we briefly review 3D imaging methods that recover the metric tensor of the surface before reconstructing the surface itself; Section 3 is dedicated to the construction of bending-invariant canonical forms [6], and in Section 4 we present our modified FMTD algorithm. Section 5 shows how 3D face recognition works on photometric stereo data. Section 6 concludes the paper.

2 Surface Acquisition

The face recognition algorithm discussed in this paper treats faces as three-dimensional surfaces. It is therefore necessary to obtain first the facial surface of the subject that we are trying to recognize.

Here, our main focus is on 3D surface reconstruction methods that recover *local* properties of the facial surface, particularly the surface gradient.³ As we will show in the following sections, the actual surface reconstruction is not really needed for the recognition.

³ The relationship between the surface gradient and the metric tensor of the surface is established in Section 4 in equations (16) and (18).

2.1 Photometric Stereo

The photometric stereo technique consists of obtaining several pictures of the same subject in different illumination conditions and extracting the 3D geometry by assuming the Lambertian reflection model. We assume that the facial surface, represented as a function, is viewed from a given position along the z-axis. The object is illuminated by a source of parallel rays directed along l^i (Figure 1).

$$I^i(x, y) = \rho(x, y)n(x, y) \cdot l^i, \quad (1)$$

where $\rho(x, y)$ is the object albedo, and $n(x, y)$ is the normal to the object surface, given as

$$n(x, y) = \frac{(-z_x(x, y), -z_y(x, y), 1)}{\sqrt{1 + \|\nabla z(x, y)\|_2^2}}. \quad (2)$$

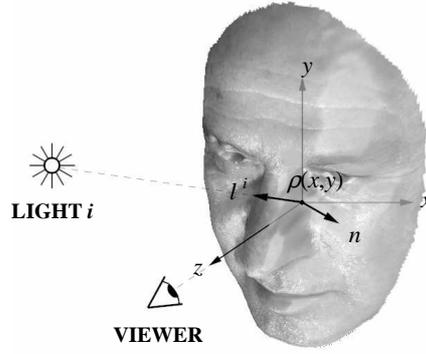


Fig. 1. 3D surface acquisition using photometric stereo

Using matrix-vector notation, Eq. (2) can be rewritten as

$$I(x, y) = Lv, \quad (3)$$

where

$$L = \begin{bmatrix} l_1^1 & l_2^1 & l_3^1 \\ \vdots & \vdots & \vdots \\ l_1^N & l_2^N & l_3^N \end{bmatrix}; \quad I(x, y) = \begin{bmatrix} I^1(x, y) \\ \vdots \\ I^N(x, y) \end{bmatrix}, \quad (4)$$

and

$$v_1 = -z_x v_3; \quad v_2 = -z_y v_3; \quad v_3 = \frac{\rho(x, y)}{\sqrt{1 + \|\nabla z\|_2^2}}. \quad (5)$$

Given at least 3 linearly independent illuminations $\{l^i\}_{i=1}^N$ and the corresponding observations $\{I^i\}_{i=1}^N$, one can reconstruct the values of ∇z by pointwise least-squares solution

$$v = L^\dagger I(x, y), \quad (6)$$

where $L^\dagger = (L^T L)^{-1} L^T$ denotes the Moore-Penrose pseudoinverse of L . When needed, the surface can be reconstructed by solving the Poisson equation

$$\tilde{z}_{xx} + \tilde{z}_{yy} = z_{xx} + z_{yy} , \quad (7)$$

with respect to \tilde{z} , which is the minimizer of the integral measure

$$\int \int ((\tilde{z}_x - z_x)^2 + (\tilde{z}_y - z_y)^2) dx dy.$$

Photometric stereo is a simple 3D imaging method, which does not require expensive dedicated hardware. The assumption of Lambertian reflection holds for most parts of the human face (except the hair and the eyes) and makes this method very attractive for 3D face recognition application.

2.2 Structured light

Proesmans et al. [11] and Winkelbach and Wahl [12] proposed a shape from 2D edge gradients reconstruction technique, which allows to reconstruct the surface normals (gradients) from two stripe patterns projected onto the object. The reconstruction technique is based on the fact that directions of the projected stripes in the captured 2D images depend on the local orientation of the surface in 3D. Classical edge-detecting operators can be used to find the direction of the stripe edges.

Figure 2 describes the relation between the surface gradient and the local stripe direction. A pixel in the image plane defines the viewing vector s . The stripe direction determines the stripe direction vector v' , lying in both the image plane and in the viewing plane. The real tangential vector of the projected stripe v_1 is perpendicular to the normal $c = v' \times s$ of the viewing plane and to the normal p of the stripe projection plane. Assuming parallel projection, we obtain

$$v_1 = c \times p . \quad (8)$$

Acquiring a second image of the scene with a rotated stripe illumination relative to the first one, allows to calculate a second tangential vector v_2 . Next, the surface normal is computed according to

$$n = v_1 \times v_2 . \quad (9)$$

In [13], Winkelbach and Wahl propose to use a single lighting pattern to estimate the surface normal from the local directions and *widths* of the projected stripes.

3 Bending-invariant representation

Human face can not be considered as a rigid object since it undergoes deformations resulting from facial expressions. On the other hand, the class of transformations that a facial surface can undergo is not arbitrary, and a suitable model

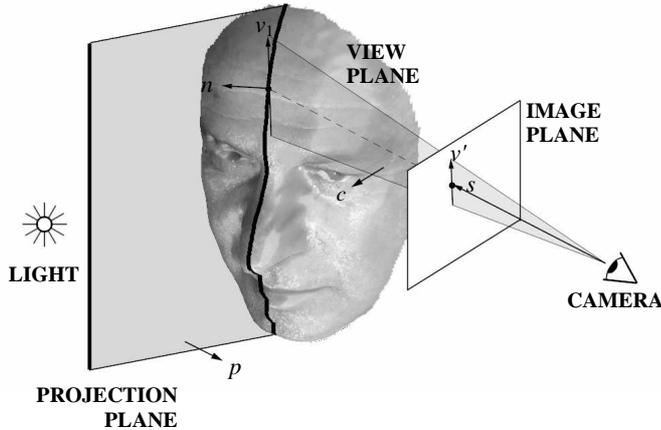


Fig. 2. 3D surface acquisition using structured light

for facial expressions is of *isometric* (or length-preserving) transformations [5]. Such transformations do not stretch or tear the surface, or more rigorously, preserve the surface metric. In face recognition application, faces can be thought of as an equivalence classes of surfaces obtained by isometric transformations. Unfortunately, classical surface matching methods, based on finding an Euclidean transformation of two surfaces which maximizes some shape similarity criterion (see, for example, [15], [16], [17]) usually fail to find similarities between two isometrically-deformed objects.

In [6], Elad and Kimmel introduced a deformable surface matching method, referred to as *bending-invariant canonical forms*, which was adopted in [5] for 3D face recognition. The key idea of this method is computation of invariant representations of the deformable surfaces, and then application of a rigid surface matching algorithm on the obtained invariants. We give a brief description of the method, necessary for the elaboration in Section 4.

Given a polyhedral approximation of the facial surface, S . One can think of such an approximation as if obtained by sampling the underlying continuous surface with a finite set of points $\{p_i\}_{i=1}^n$, and discretizing the metric associated with the surface

$$\delta(p_i, p_j) = \delta_{ij} . \quad (10)$$

We define the matrix of *squared mutual distances*,

$$(\Delta)_{ij} = \delta_{ij}^2 . \quad (11)$$

The matrix Δ is invariant under isometric surface deformations, but is not a unique representation of isometric surfaces, since it depends on arbitrary ordering and the selection of the surface points. We would like to obtain a geometric invariant, which would be unique for isometric surfaces on one hand, and will allow using simple rigid surface matching algorithms to compare such invariants on the other. Treating the squared mutual distances as a particular case of

dissimilarities, one can apply a dimensionality-reduction technique called multidimensional scaling (MDS) in order to embed the surface points with their geodesic distances in a low-dimensional Euclidean space \mathbb{R}^m [10], [14], [6].

In [5] a particular MDS algorithm, the *classical scaling*, was used. The embedding into \mathbb{R}^m is performed by first double-centering the matrix Δ

$$B = -\frac{1}{n}J\Delta J \quad (12)$$

(here $J = I - \frac{1}{n}U$; I is a $n \times n$ identity matrix, and U is a matrix of ones). Then, the first m eigenvectors e_i , corresponding to the m largest eigenvalues of B , are used as the embedding coordinates

$$x_i^j = e_i^j; \quad i = 1, \dots, n; \quad j = 1, \dots, m. \quad (13)$$

where x_i^j denotes the j -th coordinate of the vector x_i . The set of points x_i obtained by the MDS is referred to as the bending-invariant canonical form of the surface; when $m = 3$, it can be plotted as a surface. Standard rigid surface matching methods can be used in order to compare between two deformable surfaces, using their bending-invariant representations instead of the surfaces themselves. Since the canonical form is computed up to a translation, rotation, and reflection transformation, in order to allow comparison between canonical forms, they must be aligned. This can be done by setting the first-order moments (center of mass) and the mixed second-order moments of the canonical form to zero (see [18]).

4 Measuring geodesic distances

One of the crucial steps in the construction of the canonical form of a given surface, is an efficient algorithm for the computation of geodesic distances on surfaces, that is, δ_{ij} . A numerically consistent algorithm for distance computation on triangulated domains, henceforth referred to as Fast Marching on Triangulated Domains (FMTD), was used by Elad and Kimmel [6]. The FMTD was proposed by Kimmel and Sethian [7] as a generalization of the fast marching method [8]. Using FMTD, the geodesic distances between a surface vertex and the rest of the n surface vertices can be computed in $O(n)$ operations. Measuring distances on manifolds was later done for graphs of functions [19] and implicit manifolds [20].

Since the main focus of this paper is how to *avoid* the surface reconstruction, we present a modified version of FMTD, which computes the geodesic distances on a surface, using the values of the surface gradient ∇z only. These values can be obtained, for example, from photometric stereo or structured light.

The facial surface can be thought of as a parametric manifold, represented by a mapping $X : \mathbb{R}^2 \rightarrow \mathbb{R}^3$ from the parameterization plane $U = (u^1, u^2) = (x, y)$ to the manifold

$$X(U) = (x^1(u^1, u^2), x^2(u^1, u^2), x^3(u^1, u^2)); \quad (14)$$

which, in turn, can be written as

$$X(U) = (x, y, z(x, y)) . \quad (15)$$

The derivatives of X with respect to u^i are defined as $X_i = \frac{\partial}{\partial u^i} X$, and they constitute a non-orthogonal coordinate system on the manifold (Figure 3). In the particular case of Eq. (15),

$$X_1(U) = (1, 0, z_x(x, y)); \quad X_2(U) = (0, 1, z_y(x, y)) . \quad (16)$$

The distance element on the manifold is

$$ds = \sqrt{g_{ij}u^i u^j} , \quad (17)$$

where we use Einstein's summation convention, and the metric tensor g_{ij} of the manifold is given by

$$(g_{ij}) = \begin{bmatrix} g_{11} & g_{12} \\ g_{21} & g_{22} \end{bmatrix} = \begin{bmatrix} X_1 \cdot X_1 & X_1 \cdot X_2 \\ X_2 \cdot X_1 & X_2 \cdot X_2 \end{bmatrix} \quad (18)$$

The classical Fast Marching method [8] calculates distances in an orthogonal coordinate system. The numerical stencil for the update of a grid point consists of the vertices of a right triangle. In our case, $g_{12} \neq 0$ and the resulting triangles are not necessarily right ones. If a grid point is updated by a stencil which is an obtuse triangle, a problem may arise. The values of one of the points of the stencil might not be set in time and cannot be used. There is a similar obstacle in Fast Marching on triangulated domains which include obtuse triangles [7].

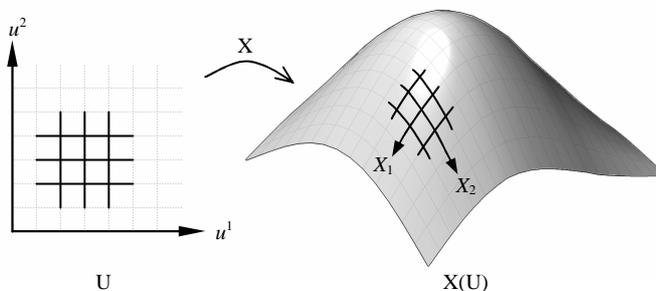


Fig. 3. The orthogonal grid on the parameterization plane U is transformed into a non-orthogonal one on the manifold $X(U)$

Our solution is similar to that of [7]. We perform a preprocessing stage for the grid, in which we split every obtuse triangle into two acute ones (see Figure 4). The split is performed by adding an additional edge, connecting the updated grid point with a non-neighboring grid point. The distant grid point becomes

part of the numerical stencil. The need for splitting is determined according to the angle between the non-orthogonal axes at the grid point. It is calculated by

$$\cos \alpha = \frac{X_1 \cdot X_2}{\|X_1\| \|X_2\|} = \frac{g_{12}}{\sqrt{g_{11}g_{22}}} . \quad (19)$$

If $\cos \alpha = 0$, the axes are perpendicular, and no splitting is required. If $\cos \alpha < 0$, the angle α is obtuse and should be split. The denominator in the rhs of Eq. (19) is always positive, so we need only check the sign of the numerator g_{12} . In order to split an angle, we should connect the updated grid point with another point, located m grid points from the point in the X_1 direction, and n grid points in the X_2 direction (m and n can be negative). The point is a proper supporting point, if the obtuse angle is split into two acute ones. For $\cos \alpha < 0$ this is the case if

$$\cos \beta_1 = \frac{X_1 \cdot (mX_1 + nX_2)}{\|X_1\| \|mX_1 + nX_2\|} = \frac{mg_{11} + ng_{12}}{\sqrt{g_{11}(m^2g_{11} + 2mng_{12} + n^2g_{22})}} > 0, \quad (20)$$

and

$$\cos \beta_2 = \frac{X_2 \cdot (mX_1 + nX_2)}{\|X_2\| \|mX_1 + nX_2\|} = \frac{mg_{12} + ng_{22}}{\sqrt{g_{22}(m^2g_{11} + 2mng_{12} + n^2g_{22})}} > 0 . \quad (21)$$

Also here, it is enough to check the sign of the numerators. For $\cos \alpha > 0$, $\cos \beta_2$ changes its sign and the constraints are

$$mg_{11} + ng_{12} > 0; \quad \text{and} \quad mg_{12} + ng_{22} < 0 . \quad (22)$$

This process is done for all grid points. Once the preprocessing stage is done, we have a suitable numerical stencil for each grid point and we can calculate the distances.

The numerical scheme used is similar to that of [7], with the exception that there is no need to perform the unfolding step. The supporting grid points that split the obtuse angles can be found more efficiently. The required triangle edge lengths and angles are calculated according to the surface metric g_{ij} at the grid point, which, in turn, is computed using the surface gradients z_x, z_y . A more detailed description appears in [22].

5 3D face recognition using photometric stereo without surface reconstruction

The modified FMTD method allows us to bypass the surface reconstruction stage in the 3D face recognition algorithm introduced in [5]. Instead, the values of the facial surface gradient ∇z is computed on a uniform grid using one of the methods discussed in Section 2 (see Figure 5). At the second stage, the raw data are preprocessed as proposed in [5]; in that paper, the preprocessing stage was limited to detecting the facial contour and cropping the parts of the face outside the contour.

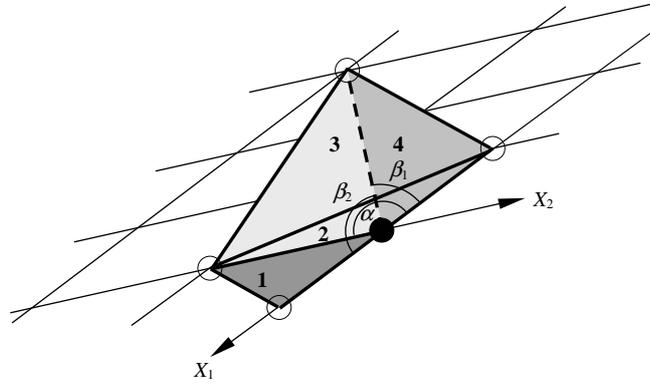


Fig. 4. The numerical support for the non-orthogonal coordinate system. Triangle 1 gives a proper numerical support, yet triangle 2 is obtuse. It is replaced by triangle 3 and triangle 4

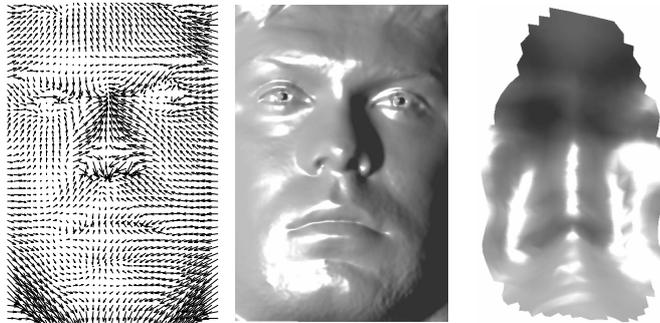


Fig. 5. Surface gradient field (left), reconstructed surface (center) and its bending-invariant canonical form represented as a surface (right)

Next, an $n \times n$ matrix of squared geodesic distances is created by applying the modified FMTD from each of the n selected vertices of the grid. Then, MDS is applied to the distance matrix, producing a canonical form of the face in a low-dimensional Euclidean space (three-dimensional in all our experiments). The obtained canonical forms are compared using a rigid surface matching algorithm. Texture is not treated in this paper.

As in [5], the method of moments described in [18] was used for rigid surface matching. The (p, q, r) -th moment of a three-dimensional surface is given by

$$M_{pqr} = \sum_n (x_n^1)^p (x_n^2)^q (x_n^3)^r, \quad (23)$$

where x_n^i denotes the i -th coordinate of the n -th point in the surface samples. In order to compare between two surfaces, the vector of first M moments $(M_{p_1 q_1 r_1}, \dots, M_{p_M q_M r_M})$, termed as the *moment signature*, is computed for each signature surface. The Euclidean distance between two moment signatures measures the dissimilarity between the two surfaces.

5.1 Experimental results

In order to exemplify our approach, we performed an experiment, which demonstrates that comparison of canonical forms obtained without actual facial surface reconstruction, is in some cases, better than reconstruction and direct (rigid) comparison of the surfaces. It must be stressed that the purpose of the example is not to validate the 3D face recognition accuracy (which has been previously performed in [5]), but rather to test the feasibility of the proposed modified FMTD algorithm together with photometric stereo.

The Yale Face Database B [21] was used for the experiment. The database consisted of high-resolution grayscale images of different instances of 10 subjects of both Caucasian and Asian type, taken in controlled illumination conditions (Figure 6). Some instances of 7 subjects were taken from the database for the experiment. Direct surface matching consisted of the retrieval of the surface gradient according to Eq. (6) using 5 different illumination directions, reconstruction of the surface according to Eq. (7), alignment and computation of the surface moments signature according to Eq. (23). Canonical forms were computed from the surface gradient, aligned and converted into a moment signature according to Eq. (23).

In order to get some notion of the algorithms accuracy, we converted the relative distances between the subjects produced by each algorithm into 3D proximity patterns (Figure 7). These patterns, representing each subject as a point in \mathbb{R}^3 , were obtained by applying MDS to the relative distances (with a distortion of less than 1%). The entire cloud of dots was partitioned into clusters formed by instances of the subjects C_1 – C_7 . Visually, the more C_i are compact and distant from other clusters, the more accurate is the algorithm. Quantitatively, we measured (i) the variance σ_i of C_i and (ii) the distance d_i between the centroid of C_i and the centroid of the nearest cluster. Table 1 shows

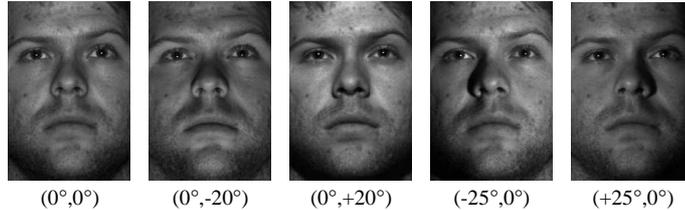


Fig. 6. A face from Yale Database B, acquired with different illuminations. Numbers in brackets indicate the azimuth and the elevation angle, respectively, determining the illumination direction

a quantitative comparison of the algorithms. Inter-cluster distances d_i are given in units of the variance σ_i . Clusters C_5 – C_7 , consisting of a single instance of the subject are not presented in the table. The use of canonical forms improved the cluster variance and the inter-cluster distance by about one order of magnitude, compared to direct facial surface matching.

Table 1. Properties of face clusters in Yale Database B using direct surface matching (*dir*) and canonical forms (*can*). σ is the variance of the cluster and d is the distance to the nearest cluster.

Cluster	σ_{dir}	d_{dir}	σ_{can}	d_{can}
C_1	0.1749	0.1704	0.0140	4.3714
C_2	0.2828	0.3745	0.0120	5.1000
C_3	0.0695	0.8676	0.0269	2.3569
C_4	0.0764	0.7814	0.0139	4.5611

6 Conclusions

We have shown how to perform face recognition according to [5], without reconstructing the 3D facial surface. We used a modification of the Kimmel-Sethian FMTD algorithm for computation of geodesic distances between points on the facial surface using only the surface metric tensor at each point. Our approach allows to use simple and efficient 3D acquisition techniques like photometric stereo for fast and accurate face recognition. Experimental results demonstrate feasibility of our approach for the task of face recognition.

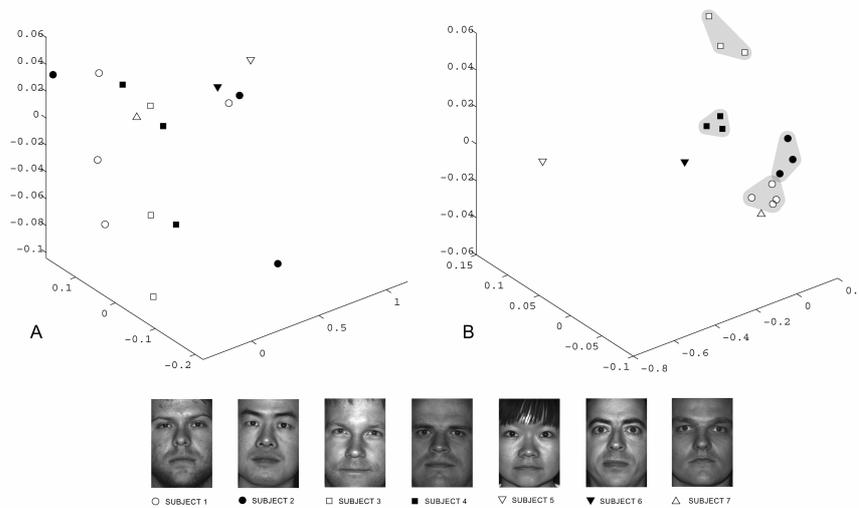


Fig. 7. Visualization of the face recognition results as three-dimensional proximity patterns. Subjects from the face database represented as points obtained by applying MDS to the relative distances between subjects. Shown here: straightforward surface matching (A) and canonical forms (B)

7 Acknowledgement

This research was supported by Dvora Fund of the Technion, Bar Nir Bergreen Software Technology Center of Excellence and the Technion V.P.R. Fund - E. and J. Bishop Research Fund.

References

1. Gordon, G.: Face recognition from frontal and profile views. Proc. Int'l Workshop on Face and Gesture Recognition (1997) 47–52
2. Beumier, C., Acheroy, M. P.: Automatic face authentication from 3D surface. In: Proc. British Machine Vision Conf. (1988) 449–458
3. Huang, J., Blanz, V., Heisele, V.: Face recognition using component-based SVM classification and morphable models. In: SVM (2002) 334–341
4. Mavridis, N., Tsalakanidou, F., Pantazis, D., Malassiotis S., Strintzis, M. G.: The HISCORE face recognition application: Affordable desktop face recognition based on a novel 3D camera. In: Proc. Int'l Conf. Augmented Virtual Environments and 3D Imaging, Mykonos, Greece (2001)
5. Bronstein, A., Bronstein, M., Kimmel, R.: Expression-invariant 3D face recognition. In: Kittler, J., Nixon, M. (eds.): Proc. Audio and Video-based Biometric Person Authentication. Lecture Notes in Computer Science, Vol. 2688. Springer-Verlag, Berlin Heidelberg New York (2003) 62–69
6. Elad, A., Kimmel, R.: Bending invariant representations for surfaces. In: Proc. CVPR (2001) 168–174 (1997) 415–438

7. Kimmel, R., Sethian, J. A.: Computing geodesic on manifolds. In: Proc. US National Academy of Science Vol. 95 (1998) 8431–8435
8. Sethian, J. A.: A review of the theory, algorithms, and applications of level set method for propagating surfaces. In: Acta numerica (1996) 309–395
9. Bronstein, A., Bronstein, M., Gordon, E., Kimmel, R.: High-resolution structured light range scanner with automatic calibration. Techn. Report CIS-2003-06, Dept. Computer Science, Technion, Israel (2003)
10. Borg, I., Groenen, P.: Modern multidimensional scaling - theory and applications. Springer-Verlag, Berlin Heidelberg New York (1997)
11. Proesmans, M., Van Gool, L., Oosterlinck, A.: One-shot active shape acquisition. In: Proc. Internat. Conf. Pattern Recognition, Vienna, Vol. C (1996) 336–340
12. Winkelbach, S., Wahl, F. M.: Shape from 2D edge gradients. In: Lecture Notes in Computer Science, Vol. 2191. Springer-Verlag, Berlin Heidelberg New York (2001) 377–384
13. Winkelbach, S., Wahl, F. M.: Shape from single stripe pattern illumination. In: Van Gool, L. (ed.): Pattern Recognition DAGM. Lecture Notes in Computer Science, Vol. 2449. Springer-Verlag, Berlin Heidelberg New York (2002) 240–247
14. Schwartz, E. L., Shaw, A., Wolfson, E.: A numerical solution to the generalized mapmaker’s problem: flattening nonconvex polyhedral surfaces. In: IEEE Trans. PAMI, Vol. 11 (1989) 1005–1008
15. Faugeras, O. D., Hebert, M.: A 3D recognition and positioning algorithm using geometrical matching between primitive surfaces. In: Proc. 7th Int’l Joint Conf. on Artificial Intelligence (1983) 996–1002
16. Besl, P. J.: The free form matching problem. In: Freeman, H. (ed.): Machine vision for three-dimensional scene. New York Academic (1990).
17. Barequet, G., Sharir, M.: Recovering the position and orientation of free-form objects from image contours using 3D distance map. In: IEEE Trans. PAMI, Vol. 19(9) (1997) 929–948
18. Tal, A., Elad, M., Ar, S.: Content based retrieval of VRML objects - an iterative and interactive approach. In: EG Multimedia Vol. 97 (2001) 97–108
19. Sethian, J., Vladimirsky, A.: Ordered upwind methods for static Hamilton-Jacobi equations: theory and applications. Techn. Report PAM 792, Center for Pure and Applied Mathematics, Univ. Calif. Berkeley (2001)
20. Memoli, F., Sapiro, G.: Fast computation of weighted distance functions and geodesics on implicit hyper-surfaces. In: Journal of Computational Physics, Vol. 173(2) (2001) 730–764
21. Yale Face Database B. <http://cvc.yale.edu/projects/yalefacesB/yalefacesB.html>
22. Spira, A., Kimmel, R.: An efficient solution to the eikonal equation on parametric manifolds, In: INTERPHASE 2003 meeting, Isaac Newton Institute for Mathematical Sciences, 2003 Preprints, Preprint No. NI03045-CPD, UK (2003)