Isoradius Contours:
New representations and techniques for 3D face registration and matching

Nick Pears and Tom Heseltine
Department of Computer Science
University of York
York, YO10 5DD, UK.

Abstract

We propose a technique for 3D face registration and matching using a novel representation called "isoradius contours". An isoradius contour is the contour on the 3D facial surface that is a known fixed distance relative to some predefined reference point (the tip of the nose). A 3D face representation contains many isoradius contours with different radii and the first major benefit of the technique is that the shape of the contours is independent of the facial pose, due to the infinite rotational symmetry of a sphere. The second major benefit of the technique is that registration, alignment and matching can be implemented using a simple process of 1D correlation. Our results have shown that registration and alignment is of comparable accuracy to ICP (iterative closest points), but is fast, non iterative, and is robust to the presence of outliers.

1 Introduction

Successful automatic 2D face recognition is often confounded by variations in lighting, face orientation and facial expression, and by background clutter. Recently, several groups of researchers [1][2][3][4] have attempted to perform face recognition from 3D data, captured using stereo cameras and/or using structured lighting. This approach has several immediate benefits: the face is easily segmented from the background, the facial orientation can be normalised to a fronto-parallel view if the recognition algorithm requires it, captured shape is independent of lighting variations, 3D shape information is explicit rather than implicit and it is easier to model facial expression.

Given that we have 3D face data, and (possibly) also a standard 2D colour-texture image registered with this data, how do we implement a face recognition system? There are many ways of encoding 3D structure in the literature. Gordon [2] discusses face recognition based on depth and curvature features. In this work 3D facial features are extracted, such as nose ridge, nose bridge, and eye corners and comparison between two faces is based on their relationship in feature space. Also there are several examples of work in which both 2D (colour/intensity) 3D images are used for the purpose of face recognition/verification [3][5]. Appearance based methods using sub-space techniques have proved competitive over recent years in terms of achieving state-of-the-art performance in 2D face recognition. It has been shown that it is possible adapt these methods, such as eigenface [6] and fisherface [7], to work with 3D and 3D/2D data [8][1][9]. The results have been promising, not least because of the excellent background segmentation and explicit, discriminating 3D data. A requirement for sub-space methods to work well is that all the data has a common alignment, which is usually a fronto-parallel view.

The work presented here introduces our work on the multi-featured "isoradius contour" representation. This provides a fast (non-iterative) accurate and robust means of aligning faces, and hence is a useful preprocessing stage for the application of sub-space based 3D and 3D/2D face recognition, but also provides a mechanism for face matching in itself as face correlations are a by-product of the alignment process. In this paper, we discuss (i) an orientation invariant representation of 3D faces and their associated (registered) 2D colour-intensity image data, (ii) how to extract this representation, (iii) the registration of faces using this representation and (iv) the correlation-based matching of faces using this representation.

In the following section, a brief overview of the isoradius representation is given to set the context of the paper. This is followed by a section on related literature, where, in particular, two related concepts are discussed, namely Stein and Medioni's "Splash" representations [10] and Chua and Jarvis's "point-signature" representation [11]. Section 4 describes how to extract such a representation from the raw 3D camera data. Section 5 then describes our approach to 3D registration and compares it to other approaches, such as the popular ICP (iterative closest points) algorithm [12]. Section 6 describes our evaluation procedure and results of registration and direct correlation matching, which has very little processing overhead over and above registration.

2 Isoradius contours : an overview

We start with the following definition: *An isoradius contour is a space curve defined by the locus on a 3D surface that is a known fixed distance relative to some predefined reference point.* Thus an isoradius contour (IRAD) can be thought of as the intersection of a sphere (of given radius, centered on that reference point) with the object surface. In the case of faces, an obvious choice for this reference point is the tip of the nose. An IRAD face representation contains many isoradius contours with different radii. A good visualisation of this is to imagine a set of concentric spheres separated by a (typically uniform) radius step, with a face positioned such that the nose tip is at the centre of all spheres. Clearly the shape of the intersection of the spheres with the face is independent of the 3 DOF head orientation, due to the infinite rotational symmetry of the spheres. This orientation invariance is a major benefit of the representation. The face representation that we propose is multi-featured in the sense that many feature types can be measured along the isoradius contour. In particular, these features include

- The shape of the isoradius contour itself, for example, as defined by it's curvature.
- Properties of the 3D facial surface along the contour, such as face curvature. In this sense the isoradius contour serves as a 'marker' on the facial surface, over which to extract a signal.
- Properties of the registered 2D colour-texture image, such as intensity and colour chromatics. Again, the isoradius contour is acting as a marker on the facial surface.

Thus, for example, we can have a collection of forty isoradius contours (IRADs) spanning the face, each with five features (IRAD signals) associated with it. This gives us an ensemble of two hundred one-dimensional signals representing the face. If the IRAD signals are scanned out in a consistent way relative to a coordinate system centered on the nose tip and aligned with the camera axes, then the rotational shift of the signal is only dependent on the head orientation.

This suggests that we can align one head with another by a process of 1D signal correlation, where the correct alignment occurs at the maxima of IRAD signal correlations, which are consistent with some rotational shift. Furthermore, the magnitudes of these correlations (when rotationally consistent) gives an immediate measure of how well one face matches with another. Note that the method described can be applied to any object. However, it works particularly well for faces, as there is a single easily identified point of interest, the nose tip.

3 Related work

The iterative closest point (ICP) algorithm due to Besl and McKay has been used extensively to register 3D data [12]. There are a large number of variants on this technique, but the basic steps are as follows: Firstly, correspondences are established between pairs of features across the two objects, based on proximity. Then the rigid transformation that maps the first member of the pair onto the second is computed. That transformation is then applied to all features in the first structure to establish further correspondences. These three steps are then reapplied until convergence is detected. The algorithm is effective when given a good initial correspondences are given and hence it is required that an initial estimate of the transformation between the two surfaces is already known. This is generally achieved with a coarse correspondence scheme, such as that used by Lu [13], where heuristics applied to local, curvature based shape indices are used before application of ICP. In order to compare ICP with our algorithm, we initialise it with the translation between the two nose tips. We have found that ICP and IRADs have similar alignment accuracies, when there are no data outliers and when the initial angular displacement is sufficiently small for ICP to converge to the correct solution. This is not surprising since they are both based on a least squares transformation between corresponding points. However, our algorithm has a number of advantages: (i) it is robust to 3D data outliers, since it explicitly finds specific contours on the facial surface. In contrast ICP can not distinguish between outliers and the facial surface, and fails unless the data is preprocessed to remove outliers. (ii) Our algorithm is not iterative and hence is fast. (iii) An alignment solution is guaranteed even over very large rotational displacements, such as 180 degrees roll. This is not the case for ICP, particularly when we have the common combined effect of outliers and large rotational displacements. (iv) Finally our algorithm is easily adapted to cope with changes in facial expression. ICP can not distinguish which facial areas are more rigid and sufficiently reliable for registration and matching. However, we have found that certain portions of the IRAD signals are invariant to expression, namely those around the forehead and nose bridge, whereas IRADs over the mouth area move between facial expressions, and are less suitable for registration.

To our knowledge, there are two previous works that use the rotational invariance of the sphere to extract 3D representations invariant to orientation. The first is due to Stein and Medioni [10], who call their approach *Structural Indexing*, with their basic representation called a *Splash* used in a hash table 3D object indexing/retrieval approach. A similar, later approach, termed *Point Signature* representations has been presented by Chua and Jarvis [11]. Like our approach, both of these methods intersect a sphere around a

feature point or point of interest (POI) in order to achieve invariance of the representation to a three degree-of freedom (DOF) orientation. However, they are quite different from our approach for a number of reasons:

Firstly, compared to both splashes [10] and point signatures [11], we do not attempt to encode the shape around multiple points of interest (POI) using a single contour. Rather, we choose to encode the whole facial surface using a dense set of contours around a single POI. This seems appropriate for human faces, since there is indeed a single distinctive (high curvature) structure in the data, namely the tip of the nose, which is easily detected. In detecting a single POI there is an implicit one-to-one correspondence between the POI on the probe (test) data and the POI on each of the 3D models in the database and thus the 3 DOF translation between the two models is easily computed.

Secondly, we do not attempt to extract a set of structural features from the data around the contour. Breaking a softly curved organic structure such as a human face into a piecewise linear segments is not stable and indeed a complex multi-threshold approach is pursued by Stein and Medioni [10] in an attempt to circumvent this problem. In contrast, we extract signals, which can be matched by a straightforward process of one-dimensional signal correlation.

Thirdly, we do not use references that can be corrupted by missing data. One of the problems with using point signatures [11] over larger area surface patches is that is suffers from 'missing parts'. If there are missing parts in the 3D face data, due to glasses or the nose occluding part of the cheek in non-frontal view, then the plane fitting process in the point signature method will be corrupted, which in turn will corrupt the projection distances. The best fit plane in this method will also be sensitive to changes in facial structure, such as those caused by changes in expression, hence any local change in surface structure will affect the whole representation of the surface around the point of interest. In contrast, our method maintains a consistent signal for all rigid sections of the surface regardless of any structural changes in other sections. For example, the part of the contour passing through the rigid forehead is not affected by the same contour passing through the malleable mouth area.

Finally, we propose the use of multiple signals types around the contours. In addition to measuring surface shape along the contour, we measure the surface reflectance to create an ensemble of signals around each contour in a set of contours.

4 Extracting IRADs

Extraction of IRADs can be divided into the following broad stages: (i) a 'point-of-interest' (POI) is extracted (the nose tip) on which the 3D data is centered, (ii) a series of isoradius surface contours are extracted on the 3D object surface (iii) both 3D surface properties and aligned colour-intensity properties along each of those contours is extracted into an 'ensemble' of one-dimensional (1D) signals. We now describe these stages in more detail.

4.1 Extraction of the POI

The first stage in extracting the isoradius representation is to locate a single point of interest (POI) on the captured 3D facial surface. The advantage of using a single point-of-interest is that the matching from query POI to database POI is implicit as it is one-to-one. A crucial point is that the 3D position of this point should be detected reliably, such that there is high repeatability (low variance in 3D position) relative to the more rigid face components (forehead, nose-bridge, cheekbones). Good interest points are those that have a local maximum in 3D surface curvature. In the case of the human face, the tip of the nose is a good choice for interest point. We have localised the nose tip using a simple approach which assumes that the pan and tilt of the face is within a 45 degree cone relative to the fronto-parallel direction. In this simple approach, we pan and tilt the face throughout a 90 degree range and build a histogram representing the frequency that a 3D vertex is marked as nearest the camera (minimum z coordinate). The vertex with the modal score in this histogram is taken as the nose tip. We have found this approach to be adequate for our initial tests, although we are currently investigating more sophisticated approaches.

4.2 Extraction of the contour locus

In order to extract properties associated with IRADs, such as curvature, we need to define the IRAD contour as a set of evenly spaced points along the IRAD sphere. The methods we have used to extract IRAD contours aim to minimize the effects of noise and interpolation errors associated with sampling a noisy and non-uniformly spaced signal. To do this we force the computation of the IRAD contour position to occur at points of equal step length across a sphere, which we call an IRAD sphere, as shown on the left hand side of figure 1. The simplest method we have used intersects planar facets of the 3D face mesh with IRAD spheres. Recently we have implemented a more sophisticated approach as follows:

1. For each IRAD (radius, r) find the list of surface patches that intersect the IRAD sphere of radius r. Here a "surface patch" is bounded by four adjacent points on the x-y grid.

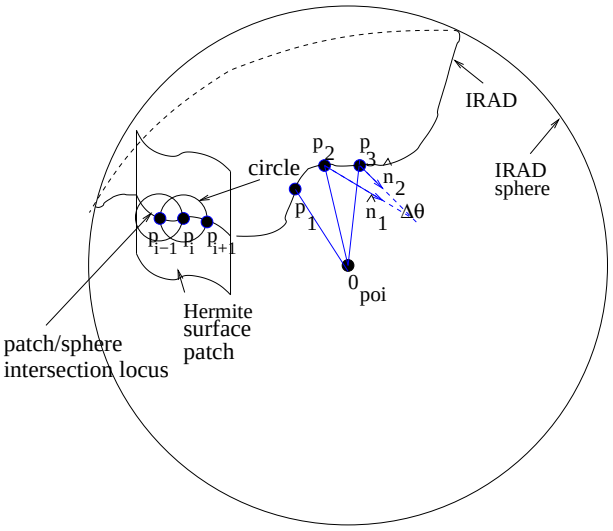

Figure 1: Extraction and encoding of an IRAD

2. Each surface patch, within a neighbourhood of four points on the x-y grid, is defined by fitting cubic Hermite patches using a neighbourhood of sixteen 3D points over the x-y grid.

3. After finding an initial point on the IRAD contour, the next point must be a fixed step away, and also on the IRAD sphere. This defines a circle on the IRAD sphere centered on the initial IRAD point, as shown in figure 1. In order to determine the next IRAD point, we must intersect this circle in 3D space, with the Hermite surface patch. We solve this equation with Brent's numerical root finding algorithm, which yields a pair of solutions for the two places in which the circle on the IRAD sphere pierces the Hermite surface patch. One of these solutions is very close to the previous IRAD point and so this is rejected to prevent the IRAD turning back on itself (see points $p_{i-1}, p_i, pi+1$ in figure 1). The remaining solution is accepted as the correct IRAD contour point to keep the IRAD growing in the correct direction.

4. This process is repeated and thus the IRAD grows using equal steps along the IRAD sphere, whilst interpolating across piece-wise continuous Hermite patches that define the facial surface at any arbitrary point on the x-y grid.

4.3 Extraction of the signal ensemble

The isoradius method requires the extraction of 3D shape properties and (possibly) colour-intensity properties of the face along isoradius contours, such that a set, or ensemble, of 1D signals is generated. Extraction of shape properties may be done in a number of ways, the most obvious of which are (i) using the facial curvature(change in facial surface normal) to encode shape and (ii) using the curvature of the IRAD contour itself. The first of these is straightforward, and requires measuring the change in face surface normal from one reference point to the adjacent one using a straightforward backward (or forward) difference operation. Obviously the operation should not be computed across breaks in the IRAD contour. Encoding the curvature of the IRAD contour itself is slightly more complicated. In this case we measure the curvature of the IRAD contour, which is due to the face shape, rather than the curvature which is simply due to the fact that the IRAD is distributed across the surface of a sphere. The process is illustrated at the centre of figure 1. Given that curvature, $\kappa = \frac{\Delta\theta}{\Delta s}$, if we maintain a constant step length, Δs, along the isoradius contour, then the angular changes, $\Delta\theta$, encode the contour shape. How do we actually compute $\Delta\theta$ along the contour? Consider three consecutive points (p_1, p_2, p_3) on the contour, separated by a fixed, but small Δs, as shown in figure 1. A normal to the contour in the direction direction $\hat{n}$ is computed as the cross-product of the two vectors Op_1 and Op_2, where p_1 and p_2 are two points on the isoradius contour and O is the POI centred origin. This vector can be recomputed for points p_2 and p_3 using the cross-product of Op_2 and Op_3. The change in angle of this normal vector, $\Delta\theta$, is the angle that we need.

5 Registration of 3D faces

In order to make a meaningful use of a whole range of IRAD signals associated with different radii from the nose tip, the IRAD signals must be moved into a common rotational alignment. It turns out that this process of achieving a common rotational alignment of all IRADs simultaneously achieves registration of one 3D facial surface with another, which is a necessary pre-requisite for the application of appearance based sub-space methods of facial recognition, such as Eigenface (PCA based) and Fisherface (LDA based). We now propose a new method to align a pair of 3D facial surfaces using IRAD signals. The pair of 3D images being matched may be two individual facial surfaces, either of the same subject or different subjects, or could be a match of an individual subject with an average facial surface, for the purposes of alignment to a common orientation throughout the full database population. This would be useful in the application of sub-space methods and requires a two-pass process to generate an average face in the first place.

Given that we have an ensemble of signals, 1D signal correlation suggests itself as a mechanism for alignment,

and, as a by product, a means of matching with little extra processing over and above the registration procedure. Much of the power of the isoradius method is that the dense, comprehensive multi-contour, multi-feature representation employed makes 1D signal correlation central to the registration and matching process. Of course, in the correlation process, we need to deal with IRAD signals of different sizes (this can occur even on two 3D images of the same person, particularly if the face is in a different expression) and IRAD signals that are fragmented due to holes in the data. For now, lets suppose that the two signals are the same size and both form a closed loop on the facial surface. We express these signals as discrete data sets: $\mathbf{x} = [x_1...x_n]^T$ and $\mathbf{y} = [y_1...y_n]^T$. The normalised cross correlation C is given as:

$$C = \frac{\mathbf{x^T y}}{\sqrt{\mathbf{x^T x} + \mathbf{y^T y}}}, \quad where \quad \mathbf{x^T x} + \mathbf{y^T y} > t^2 \quad (1)$$

for some threshold t. For n-1 rotational shifts of the $\mathbf{x}$ vector, we obtain n values of C, which yields a normalised cross correlation signal over n values. (Note that we talk of 'rotational shifts', because the $\mathbf{x}$ vector is being managed in a buffer where data is shifted and when it falls off the end of the buffer it is placed at the front of the buffer, much like a rotational shift register.)

The maximum value of the correlation signal suggests the correct alignment of the IRAD contour pair and we can generate a list of 3D correspondences along the matched pair of IRAD contours, as:

$$\mathbf{x}_q(i) \rightarrow \mathbf{x}_d(j) \quad , i = 1...n, j = i + k, \ modulo(n) \quad (2)$$

where $\mathbf{x}_q = (x, y, z)_q^T$ is a 3D point on the query surface, $\mathbf{x}_d = (x, y, z)_d^T$ is a 3D point on the database surface, n is the number of points on the IRAD signal pair, and k is the rotational shift (in contour steps) required to achieve the peak in correlation.

Thus, if we had a face with N IRADs we would have N correlation peaks, each of which has a (different sized) set of n 3D correspondences. Each of these sets of 3D correspondences defines a 3D rotation about the origin, which is positioned at the nose tip. We compute these rotations using the least squares method of Haralick et al [14]. First compute the cross covariance matrix, $\mathbf{K}$ given by:

$$\mathbf{K} = \Sigma_{i=1}^{n}(x_q(i) - \overline{x}_q)(x_d(j) - \overline{x}_d)^T \quad (3)$$

we then compute the singular value decomposition of K as

$$\mathbf{K} = \mathbf{USV}' \quad (4)$$

where $\mathbf{S}$ is the diagonal matrix of singular values and $\mathbf{V}$ and $\mathbf{U}$ are orthogonal matrices. The rotation matrix, $\mathbf{R}$, is then given by

$$\mathbf{R} = \mathbf{VU}' \quad (5)$$

We now have N rotation matrices and most of these matrices represent very similar 3D rotations, namely the 3D rotation that approximately aligns the two faces. However, some matrices, particularly those associated with very small or very large radius IRAD contours, can represent inconsistent alignments, the most common of which is an alignment which is approximately 180 degrees out of phase with the correct alignment. To remove these inconsistent IRAD correlations, we implement a simple rotation clustering process as follows. Rotation matrices are converted to angle-axis format (via quaternions) and we find the largest cluster within the set of N points in angle-axis space. The IRADs associated with this cluster are labelled 'inlier IRADs', while the remainder are deemed to be outliers. We then create a large list of 3D point correspondences, by concatenating all 3D correspondences associated with inlier IRADs. Finally, we reapply the LS process in equations 3 to 5 to this large correspondence list in order to compute the refined rotation matrix, $\mathbf{R}$.

6 Results

We now present our results, when applying the IRAD method to our data sets, in which each "3D image" consists of the 3D point cloud and mesh, and a registered 2D intensity image. Our testing proceeded initially with sanity checks on the method and then with tests of the performance of the method as follows:

1. We extracted some IRAD signals to check that the signal shapes were intuitively what we expected. IRADs extracted for a single 3D image are shown in figures 2, 3. These are shown from two different viewpoints to emphasise the three dimensionality. Notice how the front view gives a pattern similar to a fingerprint, and the IRAD pattern can be viewed as a kind of "faceprint".

2. We tested the ability of our system to align 3D images. For this experiment, we have used a subset of our data, namely four subjects, each with eleven 3D images to test our alignment procedure. Each person in this set is marked up with three small 2x2 chessboard patterns applied to rigid parts of the face in order to estimate the ground truth alignment after applying IRAD based registration. Each image of a given subject is registered with every other image of that subject to give a total of 220 registration operations. An example of IRAD based registration and alignment is shown in figures 4 to 7. The first two figures show images, as captured from the camera, but centered on the nose tip, the third

image is the first image registered and rotated to the second image using the computed rotation ($\mathbf{R}$), and the final image is the second image aligned to the first, using the inverse of the computed rotation. Figure 8 shows the IRAD contour curvature signal for image 1 at a radius of 30mm. The first peak in the signal is where the IRAD traverses the nose bridge and the second is the mouth area. The fact that the signal has a pair of dominant peaks explains why individual IRAD registration errors are usually manifested as 180 degree misalignments. Note that it is possible to correct such errors by selecting the second most dominant correlation peak, if the angular displacement associated with this second peak falls within the dominant angle-axis cluster.

3. We tested the correlation based recognition ability of *individual* facial curvature IRAD signals at different radii, using 30 subjects with 12 images each. (Note that the IRAD contours used in this experiment were generated directly from mesh facet intersection with IRAD spheres, i.e. there was no Hermite patch fitting). In this experiment 64,620 face matching operations were performed and a match was accepted or rejected based on thresholding a correlation score. (Note that the data set we used is very challenging, with 50% of the images in a non-standard pose/expression. The set comprised 6 neutral, one head facing up (about 45 degrees), one head facing down (45 degrees), one happy expression, one angry expression, one eyebrows raised, and one 3D image double the distance from the camera.) Although direct correlation is not necessarily a good way to use IRADs for face matching, this experiment provides useful information to try to understand which parts of the face, in terms of IRAD contours are the most discriminating, and hence what range of IRAD radii we may wish to incorporate in more sophisticated matching schemes. Direct correlation based recognition rate curves, expressing false acceptance ratio (FAR) against false rejection ratio (FRR) for a selection of six IRADs, with facial curvature signals, are shown in fig 9. The equal error rate (EER), when FAR is equal to FRR, expresses the performance of direct correlation when matching faces and is shown for each radius that we tested in fig 10. Note that the areas across the nose bridge are most discriminating across the data set, whereas those very close to the nose and those on the periphery of the facial surface as less discriminatory. When visualising IRADs, these results accord with intuition, since the area very close to the nose tip is near spherical and featureless, and peripheral contours can be quite uniform in depth and near circular in the x-y plane. In contrast the contours over the nose bridge and cheek areas tend to be much more varied across the data set. The best performance achieved was an EER of 21.91%, which is encouraging, given that the result is from a single signal around a single IRAD contour.

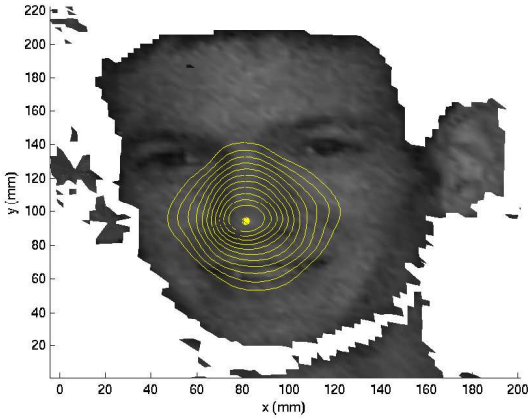

Figure 2: Irad contours on 3D image 00179-08: front view

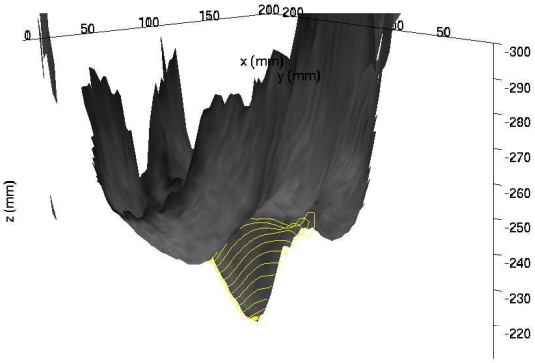

Figure 3: Irad contours on 3D image 00179-08: profile view

7 Conclusions

We have proposed a new representation for 3D face registration and matching, called the isoradius contour (IRAD) representation. It has proved to be effective for 3D face registration and has several advantages over the commonly used ICP method [12] and its variants, namely it is robust to outliers, non iterative and almost always provides a good alignment solution, irrespective of the initial displacement in pose. Furthermore, unlike ICP, the method has good potential to be made invariant to facial expression by analysing the variance of IRAD signals across different

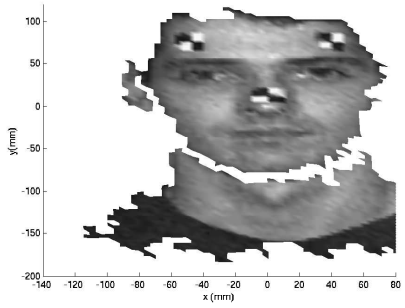

Figure 4: Marked subject frontal (image 1)

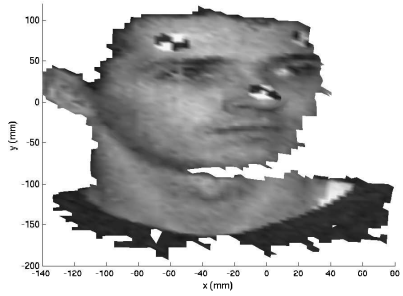

Figure 5: Marked subject looking to left (image 2)

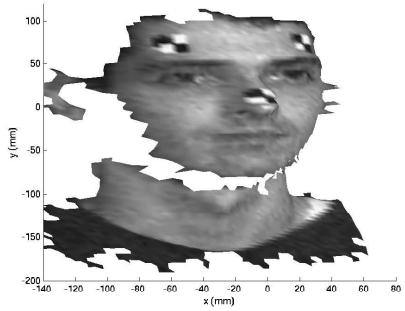

Figure 6: Image 1 rotated to align with image 2

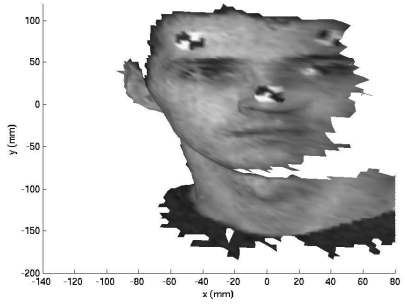

Figure 7: Image 2 rotated to align with image 1

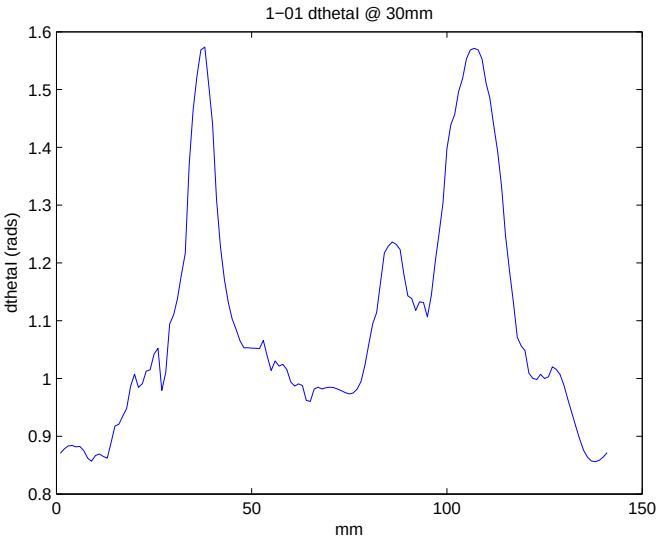

Figure 8: IRAD contour curvature signal for image 1 at radius=30mm

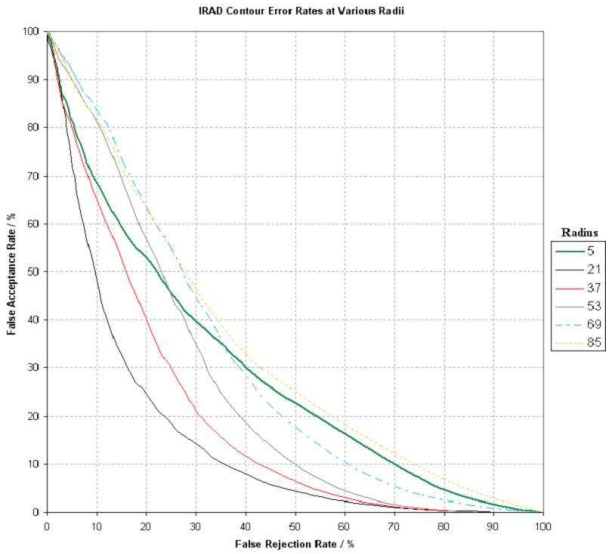

Figure 9: Error curves for six individual IRAD signals

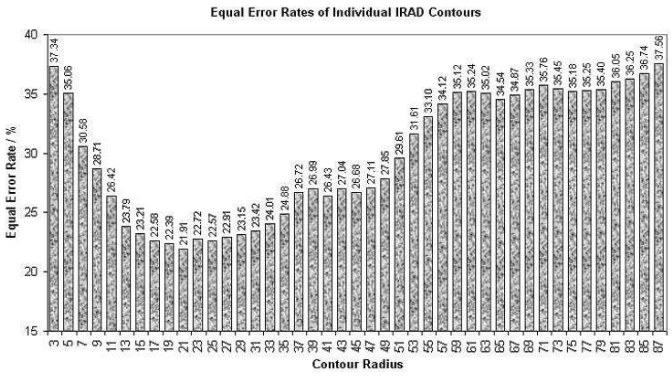

Figure 10: Bar chart of EER for face curvature IRAD signals at various radii.

parts of the face under a variety of facial expressions. A useful by-product of the IRAD registration procedure is that we have computed a set of maximum correlation scores that are consistent with some rigid rotational alignment. We have shown that these correlation scores may be used in a simple direct correlation matching process, to determine which IRAD signals are most discriminating and, overall, to give a reasonable face recognition performance, possibly as a correlation-based pre-filter to a more sophisticated matching method. Finally, we believe that the IRAD representation is appropriate to support simultaneous 3D face recognition and identification of facial expression. Modelling how IRADs deform over a range of facial expressions is one focus of our current research.

References

[1] T Heseltine N E Pears and J Austin, "Three-dimensional face recognition: A fi shersurface approach," in *Proc. Int. Conf. Image Analysis and Recognition. LCNS 3212, part II*, 2004, pp. 684–691.

[2] G. G. Gordon, "Face recognition based on depth and curvature features," in *Proc. IEEE Computer Society Conf. on: Computer Vision and Pattern Recognition*, 1992, pp. 808–810.

[3] Ho Y Wang Y, Chua C and Ren Y, "Integrated 2d and 3d images for face recognition," in *11th Int. Conf. on Image Analysis and Processing*, 2001, pp. 48–53.

[4] A Bronstein Bronstein M R Kimmel and A Spira, "3d face recognition without facial surface reconstruction," in *Proceedings of ECCV 2004*, 2004.

[5] Tsalakanidou F S Malassiotis and M G Strintzis, "Face localisation and authentication using color and depth images," *IEEE trans. Image Processing*, vol. 14, no. 2, pp. 152–168, 2005.

[6] M Turk and A Pentland, "Eigenfaces for recognition," *Journal of Cognitive Neuroscience*, vol. 3, no. 1, pp. 71–86, 1991.

[7] Peter N. Belhumeur, Joao Hespanha, and David J. Kriegman, "Eigenfaces vs. fi sherfaces: Recognition using class specifi c linear projection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 19, no. 7, pp. 711–720, 1997.

[8] T Heseltine N E Pears and J Austin, "Three-dimensional face recognition: An eigensurface approach," in *Proc. IEEE Int. Conf. Image Processing*, 2004, pp. 1–2.

[9] Chang K I K W Bowyer and P J Flynn, "An evaluation of multimodal 2d+3d face biometrics," *IEEE Trans. PAMI*, vol. 27, no. 4, pp. 619–624, 2005.

[10] F Stein and G Medioni, "Structural indexing: Effi cient 3-d object recognition," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 14, no. 2, pp. 125–145, 1992.

[11] C.S. Chua and R. Jarvis, "Point signatures: A new representation for 3D object recognition," *Int. Journal of Computer Vision*, vol. 25, no. 1, pp. 63–85, 1997.

[12] P.J. Besl and N.D. McKay, "A method for registration of 3D shapes," *IEEE Trans. Pattern Analysis and Machine Intelligence (PAMI)*, vol. 14, no. 2, pp. 239–256, 1992.

[13] Lu X A K Jain and D Colbry, "Matching 2.5d face scans to 3d models," *IEEE Trans. PAMI*, vol. 28, no. 1, pp. 31–43, 2006.

[14] Haralick R M Joo H Lee C Zhuang X Vaidya V G Kim M B, "Pose estimation from corresponding point data," *IEEE Trans. Sys. Man. Cybernetics*, vol. 19, no. 6, pp. 1426–1446, 1989.