

RBF Shape Histograms and Their Application to 3D Face Processing

Nick Pears
Department of Computer Science
University of York, UK
nep@cs.york.ac.uk

Abstract

We present novel, pose invariant 3D shape descriptors and we test the performance of these descriptors, when applied to the problems of nose identification and localisation in 3D face data. We generate an implicit radial basis function (RBF) model of the facial surface and construction of our novel features is based on sampling this RBF model over a set of concentric spheres to give a spherically-sampled RBF (SSR) histogram. In addition to providing a feature identification mechanism, SSR histograms can be processed, with very little computational overhead, to estimate the volumetric intersection of the object (face) and a bounding sphere, centred on any object surface point. A minimisation of this volume, at an appropriate scale, can be used to both define and localise the facial nose tip to an arbitrary resolution. We test our descriptors on a subset of the particularly challenging University of York 3D face database. This data set consists of 1736 3D faces, with facial expression variations, pose variations, data spikes and missing parts. Noses vertices are identified at a rate of 99.6% on unseen subjects and our approach significantly outperforms three variants of spin images.

1. Introduction

Recently, there has been a lot of research interest in both 3D face processing [12], [16] and 2D/3D face processing [15], [6]. Appearance based methods have proved competitive in terms of achieving state-of-the-art performance in 2D face recognition. It is possible to adapt these methods, such as fisherface [2], to work with 3D data [9]. The results have been promising, because of the excellent background segmentation and explicit, discriminating 3D data. A requirement for such methods to work well is that all the data has a common alignment, which is usually rendered as a fronto-parallel view. We have developed a process for highly repeatable 3D face alignment, when that 3D face data is potentially noisy and has missing parts due to spectacles, beards and self-occlusion. The four steps of this

process are as follows: (i) filter the data automatically, (ii) identify the nose tip vertex and refine (interpolate) the nose tip location, (iii) find the face orientation and (iv) generate a pose-normalised depth map ready for recognition. Finding the nose tip location is a key step, because we can then align a face to a 3D upper face template using either ‘iterative closest points’ (ICP) [3] on nose-centred data for small pose angles (typically, less than 30 degrees), or the method of ‘isoradius contours’ [13] for arbitrary pose angles. In this approach, an isoradius contour allows pose alignment to be implemented as a simple process of 1D signal correlation.

This paper focuses on 3D feature *identification* and *localisation*, and exemplifies our novel 3D surface representations by describing step (ii) above, nose tip identification and localisation, subject to pose variations and noisy data. Historically, several researchers have sought to extract pose invariant 3D surface descriptors. Notable early examples from the 1990’s include Stein and Medioni’s ‘splash’ representations [14], Chua and Jarvis’ ‘point signature’ representations [7] [15] and Johnson and Herbert’s ‘spin images’ [10]. Of these methods, spin images have been taken up most widely by the research community [1], perhaps because they are intuitive and simple to compute. Recent work, for example, has focussed on matching multi-resolution pyramids of spin images [8] in order to speed up the matching process. Other researchers have used spin images to localise 3D facial features [5] and implement face recognition from 3D shape information derived from shading [11].

Our techniques are underpinned by a radial basis function (RBF) model of the 3D facial surface and so, in the following section, we provide a brief overview of RBF modelling. In section 3, we describe our 3D surface descriptors, while in section 4, we describe how these descriptors are used in a *cascaded filter and refinement* approach to nose localisation. In section 5, we test the performance of these descriptors on a 3D face dataset, which is noisy and has both pose variations and facial expression variations. The performance is compared to three variants of spin image representation. A final section is used for conclusions.

2. RBF surface modelling

RBFs have been used extensively for interpolating scattered 3D point cloud data [4]. To briefly recap this work, a *radial* function has a value at some point in n-dimensional space \mathbf{x} , which only depends on its 2-norm relative to another point, called a ‘‘centre’’. Hence, in our case of 3D space, the radial function value is constant over a sphere. A *radial basis function* uses a weighted sum of basis functions to implicitly model a surface, where the basis function may be linear, Gaussian, cubic spline or some other function, which is radial in form. Formally, we follow reference [4] and define a radial basis function, s , as:

$$s(x) = p(x) + \sum_{i=1}^{N_c} \lambda_i \Phi(x - x_i) \quad (1)$$

where p is a polynomial (linear or quadratic), λ_i are the RBF coefficients, Φ is a real valued function called the basis function and x_i are the N_c RBF centres.

In fitting a 3D surface, s is chosen such that $s(\mathbf{x}) = 0$ forms a surface that smoothly interpolates the data points x_i . Thus the RBF model parameters implicitly define the surface as the set of points where the RBF function is zero. This is called the zero isosurface of the RBF function. Note that one can not simply solve the equation $s(x_i) = 0$ for our N data points, as this yields a trivial solution of $s(\mathbf{x}) = 0$ everywhere. Constraints where s is non-zero need to be used. Since we may readily generate ‘off-surface points’ using surface normal data, s can be chosen to approximate a *signed distance to surface* function.

Thus we can generate normals, generate off surface points using the normals, and then fit the parameters of the RBF function. Once we have the RBF function parameters, we may evaluate the function at any arbitrary 3D point \mathbf{x} , which gives the signed distance to surface at that point. By convention points below the facial surface (inside the head) are negative and those above the facial surface are positive and those on the surface, zero.

3. Spherically sampled RBF (SSR) descriptors

In spin images [10], a surface point uses its associated surface normal to form a basis with which to encode neighbouring points. Neighbouring point positions are encoded in cylindrical coordinates, as the radius in the tangent plane and height above the tangent plane. All points are binned onto a fixed grid. Corresponding 3D points across a pair of similar objects can be matched by a process of correlation of spin images or any other matching metric. Issues in spin image generation include (i) noise affecting the computation of the local surface tangent plane and (ii) problems of appropriate bin size selection. Due to these problems, we were motivated to make use of an RBF model to generate




Figure 1. Sampling of $s(x)$ in a balloon image

invariant 3D surface descriptors, which we call *spherically sampled RBF (SSR) surface descriptors*.

3.1. SSR shape histograms (‘balloon images’)

Here we propose a new kind of local surface representation, which can be derived readily from the RBF model and we call this an *SSR shape histogram*. To generate such an SSR shape histogram, we first distribute a set of n sample points evenly across a unit sphere, centered on the origin. To do this, we employ the octahedron sub division method, which, for K iterations, generates $n = \alpha\beta^K$ points. The constants are $[\alpha, \beta]^T = [8, 4]^T$ and we use $K = 3$, which gives $n = 512$. The sphere is then scaled by q radii, r_i , to give a set of concentric spheres and their common centre is translated such that it is coincident with a facial surface point. (Note that this can be a raw vertex, but can also be anywhere between vertices, on the zero isosurface).

The RBF (‘signed distance-to-surface’) function, s , is then evaluated at the $N = nq$ sample points on the concentric spheres and these values are normalised by dividing by the appropriate sphere radius, r_i , giving a set of values in the range -1 to 1. If this normalised value $s_n = \frac{s}{r_i}$ is binned over p bins, then we can construct a (pxq) histogram of normalised RBF values, which may, for visualisation purposes, be rendered as a ‘balloon image’. Note that the balloon analogy comes from incrementally inflating a sphere through the 3D domain of the RBF function, as shown in figure 1, and this is consistent with Johnson and Herbert’s [10] ‘spin image’ naming of their cylindrically encoded shape histogram. Examples of balloon images for the (protruding) nose and (flat) forehead are given in figures 2 and 3 respectively.

3.2. SSR convexity values

Clearly, the convexity of the local surface shape around some point is related to the overall brightness of the balloon image. This motivates us to consider how SSR his-

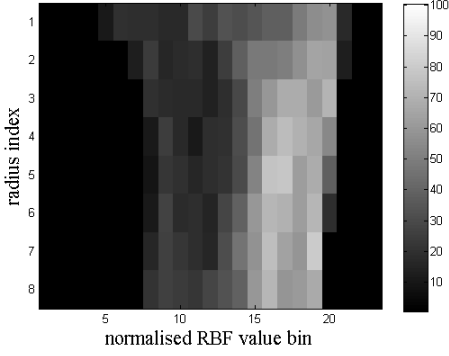


Figure 2. SSR histogram for the (protruding) nose




Figure 3. SSR histogram for the (flat) forehead (same subject)

tograms may be processed to give a pose invariant convexity value for high resolution, repeatable feature localisation. For example, if we wish to localise the nose tip, we may first define the nose tip as the point on the facial surface, where a sphere centred on that point and the face have minimum volumetric intersection. We then need to consider how to calculate the volumetric information from the SSR histogram and our approach is illustrated in figure 4. In this figure, the point p is on the object (face) surface, the upper left part of the figure is above the object surface ($s(\mathbf{x}) > 0$) and the lower right part of the figure is below the object surface ($s(\mathbf{x}) < 0$). We have illustrated three concentric spheres (solid lines) of radius (r_1, r_2, r_3), separated by Δr over which the RBF function is sampled and we consider three co-radial samples for each of these radii at $\mathbf{x}_1, \mathbf{x}_2$ and \mathbf{x}_3 respectively, noting that $s(\mathbf{x}_1) > 0, s(\mathbf{x}_2) < 0$ and $s(\mathbf{x}_3) > 0$. The dashed circles in the figure indicates the position of (non-sampling) concentric spheres that bound volumetric segments, and these have radii, ρ_i midway between the sampling spheres, namely at $\rho_i = \frac{(r_i + r_{i+1})}{2}$. In order to determine an estimate of the total volumetric intersection between the outer (dashed) sphere in the figure of radius $\rho_3 = r_3 + \frac{\Delta r}{2}$, we need to sum the all the volumetric contributions centred on radial sampling directions with $s(\mathbf{x}_1) < 0$, over all sampling radii and all sam-

pling spheres. In figure 4, the central shaded volumetric segment contributes to the intersection, but the two outer shaded volumes do not. Note that the segments centred on the larger radii have bigger volumes, and thus a weighting vector needs to be applied to the summation. Thus the volumetric intersection, V_p , at point p is given by:

$$V_p = \frac{k}{n} \mathbf{v}^T \mathbf{n}^- \quad (2)$$

where $k = \frac{4\pi}{3}$ is a constant related to the volume of a sphere, n is the total number of sample points on a sphere, \mathbf{v}^T is a vector containing the q volumetric weights (one for each radius) that increase with the cube of radius, and \mathbf{n}^- is a vector where each element is the count of the total number of sample points on a given sphere in which $s(\mathbf{x}) < 0$.

An equivalent, but more elegant approach, is to define a metric that is a relative measure of the volume of the sphere that is above the object surface compared with the volume of the sphere below the object surface. With this in mind, we define a SSR based convexity value for the point, p , as

$$C_p = \frac{k}{n} \mathbf{v}^T [\mathbf{n}^+ - \mathbf{n}^-] \quad (3)$$

where \mathbf{n}^+ is a vector in which each element is the count of the total number of sample points on a given sphere where $s(\mathbf{x}) > 0$. With this metric, a highly convex shape (a spike), will have a value close to 1.0, a highly concave shape (a pit) will have a value close to -1.0 and a flat area will have a value close to zero. This can be clearly seen from equation 3, where the elements in \mathbf{n}^+ and \mathbf{n}^- will be similar, giving a near zero vector on the right of the equation. In its simplest form, a very approximate SSR convexity value can be computed using a single sphere, which makes the volumetric weighting vector \mathbf{v} in equation 3 redundant. We use this form in this paper, which amounts to averaging the signs of the RBF function:

$$C_p = \frac{1}{n} \sum_{i=1}^n \text{sign}(s_i) \quad (4)$$

In order to illustrate the potential of this technique, a sampling sphere is moved over a facial surface (figure 5a), which illustrates the RBF distance-to-surface values by a colour mapping. The resulting SSR convexity map is shown in figures 5(b),(c),(d). A surface is rendered over this plot to aid visualisation, where the lighter areas have a convexity value near to +1 and the darker areas are close to -1 (i.e. concave). The figure indicates that, in this case, the nose is the peak convexity value in the map.

4. Nose tip identification and localisation

Computing an SSR convexity map over all vertices is computationally expensive, thus we *identify* the raw nose tip

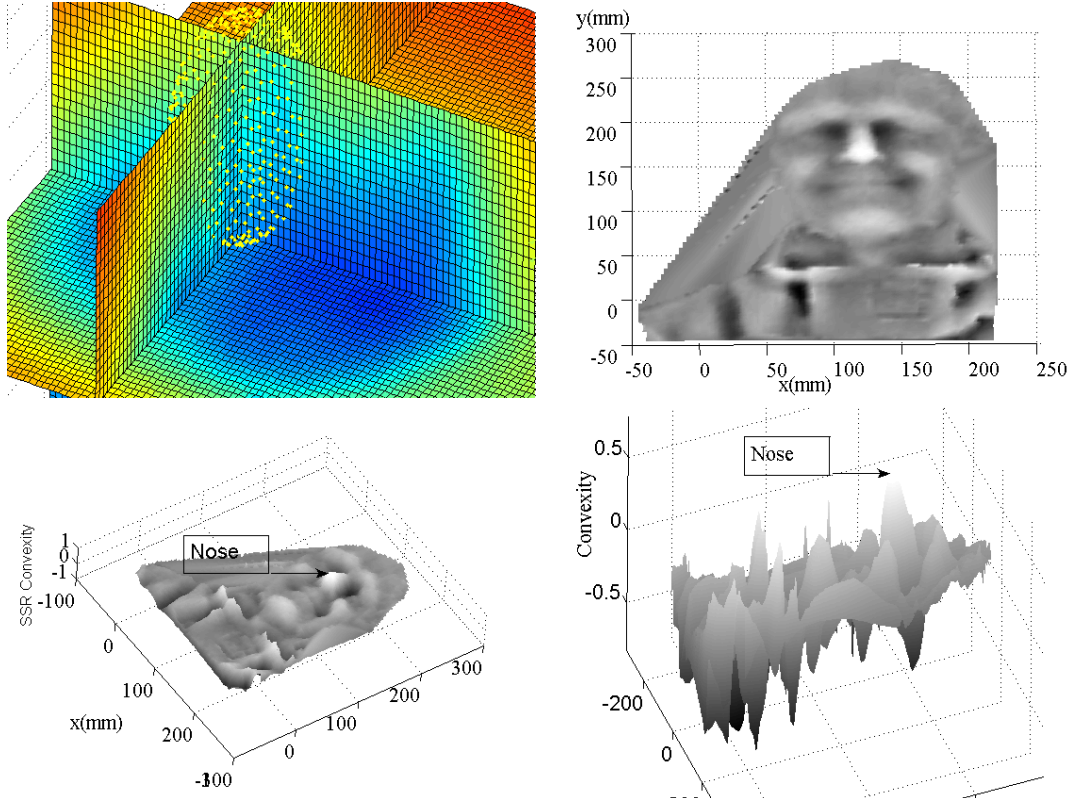


Figure 5. (a) Sampling the RBF function (b,c,d) The SSR convexity map from three views




Figure 4. Computing the SSR convexity value

vertex via a cascaded filtering process, as illustrated (from left-to-right) in figure 6. We then apply a *localisation* refinement by maximising the SSR convexity value, in the local vicinity of the identified raw vertex, using a local high density RBF-derived zero isosurface (see top-to-bottom in figure 6). The concept here is to use progressively more expensive operations to eliminate vertices. The constraints (thresholds) employed at each filtering stage are designed to

be weak, by examining trained nose feature value distributions, so that the nose tip itself is never eliminated. Conceptually, this amounts to considering every vertex as a candidate nose position, where all but one vertex are ‘false positives’. Then, at each stage, we apply a filter to reduce the number of false positives, until we have a small number of candidates at the final stage, at which point our most expensive/discriminating test is used to find the correct vertex.

The feature that we use in filter 1 is a *distance to local plane* (DLP), which has already been used to remove data spikes. The filter uses a weak threshold, which is four standard deviations below the average DLP value for trained noses. In filter 2, we compute SSR convexity values using a single sphere of radius 20mm and, again, we set a weak threshold. At this stage, we have multiple local maxima in SSR convexity value (see figure 5) and so we find these local maxima and eliminate all vertices that are not local maxima. Finally we use SSR shape histograms by finding the minimum Mahalanobis distance in ‘balloon image space’ to select the correct nose vertex. This raw vertex is refined by maximising the SSR convexity value over a local high density zero isosurface of the RBF function.

Figure 7 shows the nose candidates for each stage in the filtering process for a challenging test image. (3D vertices are mapped into the registered bitmap for visualisation.)

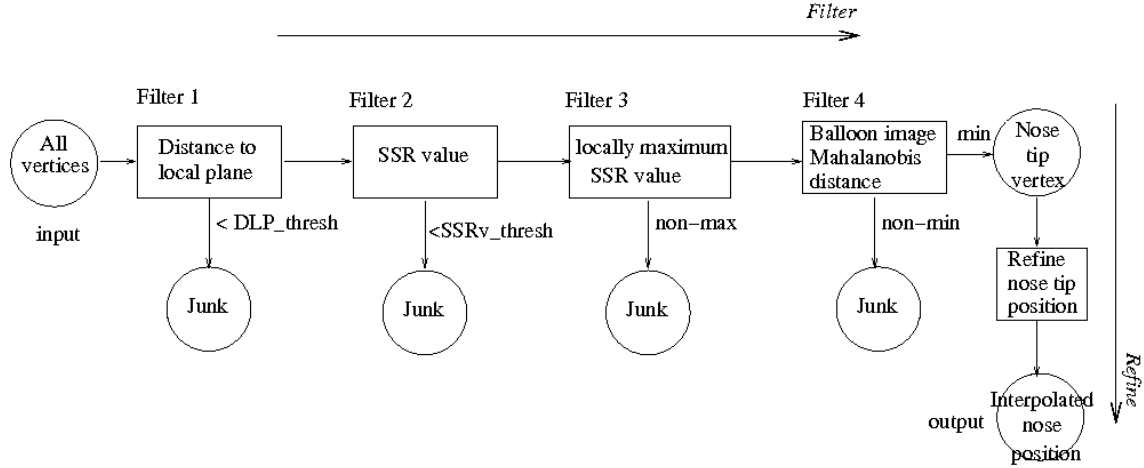


Figure 6. The cascade filter for nose-tip identification

5. Evaluation

We have evaluated our RBF derived shape descriptors on a subset of the particularly challenging University of York 3D face database. This data set has 1736 3D faces of 280 different people (subjects) and contains facial expression variations, pose variations and missing parts. The modal mesh resolution in the dataset is 3mm-4mm. We have found it convenient to split our evaluation into two categories of performance metric, namely:

1. A feature *identification* metric, measured as the percentage of correctly identified nose-tip features. This metric measures the performance of SSR shape histograms in a simple classification scheme, when compared to three variants of spin images.
2. A feature *localisation* metric, measured as the RMS repeatability of the localisation of the nose-tip. This metric measures the performance of the SSR convexity value in providing a repeatable nose localisation.


5.1. Nose identification experimental methodology

Examining the filtering stages in figure 6, one might reasonably ask: why not just take the nose candidate outputs from filter 3 (the local maxima of SSR value), compute the Mahalanobis distance to the training set of SSR values and select the minimum distance as the identified nose vertex? This is a good question, because if we can not improve on this nose identification performance, then filter 4 (using balloon images or spin images) is, at best, a waste of processing time and may even be detrimental to the overall identification performance. Therefore, we apply this metric in place of filter 4 as a baseline test (control).

Overall, we have applied five nose identification methods, each of which use the minimum Mahalanobis distance

as the nose identification criterion. The training and testing data, however, is different in each case, and is as follows: (1) Baseline test using SSR convexity values. (2) Standard spin images (spin-image type 1), where cylindrical polar coordinates, (r, h) , of local vertices are binned. (3) Our own variant of spin image (spin-image type 2), which bins a radius and angle above/below the local tangent plane $(r, \tan^{-1}(\frac{h}{r}))$. (4) A spin image which bins $(\log(r), h)$ (spin-image type 3). This is often used to give higher weight to closer vertices. (5) SSR shape histograms (balloon images). Our experimental methodology was as follows:


1. A registered bitmap for each of the 1736 images was displayed and a human operator was asked to click their best estimate of the nose tip position using a mouse, and the 2D mouse clicks were stored on disk.
2. Our nose vertex identification process, described by the filters in figure 6, was applied to the database, such that we found a set of candidate nose positions (filter 3 outputs), which were locally maximal values of SSR convexity. Our process uses weak thresholding and hence always finds the nose tip vertex (this was manually verified), but there are typically up to 10 other false positives, which occur on the chin, Adam's apple, shirt collars, quiffs of hair and spectacle frames.
3. We mapped each of these 3D nose candidates into their associated, registered 2D bitmap images and the bitmap position closest to the manual nose click (in step 1), was stored on disk as the correct nose vertex. This allowed us to collect training data for nose features and allowed us to establish a ground truth for the testing phase of nose identification.
4. We randomly selected 100 subjects (of the 280) and for each of these persons, we randomly selected a capture




(a) Filter 1 output (High DLP)



(b) Filter 2 output (High SSR convexity)



(c) Filter 3 output (Local convexity peaks)



(d) Raw nose (dot), refined nose (cross)

Figure 7. 3D vertex outputs of the *cascade filter and refine* process mapped into 2D image

condition to give 100 training 3D images.

- For each of these 100 training 3D images, we constructed a SSR shape histogram, using 8 radii of 10mm to 45mm in steps of 5mm and 23 bins for normalised RBF values. This gave SSR shape histograms (or balloon images) of dimension (8x23). We also constructed three variants of spin images, as described above. These were constructed to the same resolution as the balloon images, namely (8x23) resolution, using a maximum radius of 45mm and a height of ± 45 mm.

- We applied principal components analysis (PCA) to all four sets of training data, reducing the shape descriptor dimensionality from 184 to 64.
- For all nose candidates (filter 3 outputs) on all test images, we calculated the Mahalanobis distance to the trained data for all five methods above. For each test image, the vertex with the minimum Mahalanobis distance was identified as the nose and stored.
- We then counted, for each of the five methods, what percentage of noses were correctly identified.

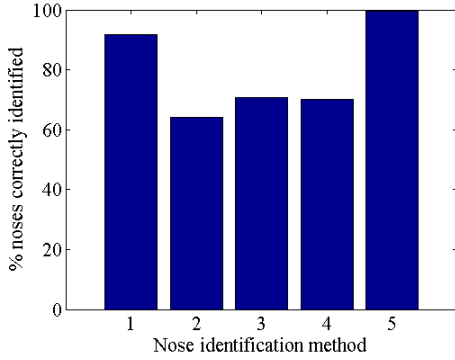


Figure 8. Nose ID performance: 1. SSR values, 2,3,4 Spin-image variants 5. SSR histograms

5.2. Nose identification results

In our database of 1736 3D images, we used 100 images of 100 individuals as training data, leaving a *test set A*, of 515 3D images, which contains the remaining images of these 100 individuals, not used in the training set, and *test set B*, which contains 1121 3D images of individuals who never appear in the 3D training set.

The results of nose identification are given in table 1 and these are displayed as a bar graph for test set B in figure 8. Note that we obtained a 91.7% rate of successful nose identification by using the SSR convexity values. Using SSR histograms improved this figure to 99.6%, whereas use of spin images degraded the system performance to around 70% and hence should be considered inappropriate for the approach adopted here. There are several reasons why SSR histograms outperformed spin images in this experiment:

1. Spin images require a normal estimate, and this normal varies greatly close to the nose tip, due to the high surface curvature. Furthermore, the normal estimate is corrupted by residual noise and missing parts. SSR histograms, however, do not require a normal estimation at all and hence are immune to such effects.
2. The data in our data set has missing parts, particularly around the eyes, when the subject is wearing spectacles. These missing parts corrupt spin images, but have little effect on SSR histograms, because the RBF function is defined everywhere in 3D space.
3. Spin images, in the form used here, use raw vertices and so the data density is a function of the raw mesh resolution. In contrast a SSR histogram can sample the RBF function to any required density. (Here we used 512 samples on each of 8 spheres, giving 4096 data elements in each SSR histogram).

5.3. Nose localisation experimental methodology

To make a preliminary evaluation of our nose localisation refinement (inter-vertex interpolation) approach, we used 80 3D facial scans in arbitrary poses, each of which had a registered 2D image. We compared our approach both with a simple automatic method and a manual method, in which a user was asked to select a raw 3D coordinate for each of the 80 images, by viewing the surface and rotating it in 3D. In the simple automatic method, the face is rotated through a raster scan of pan and tilt angles within a 45 degree cone and the *nearest point* to the camera acquires a vote. The vertex with the highest number of votes is chosen as the nose coordinate. This is called the NPH (nearest point histogram) method. Our experimental procedure was as follows:

1. Manually locate (by cursor click) three 2D features in the 2D bitmap image: we use the outer corner (*exocanthion*) of the left and right eyes and the midpoint of the upper vermillion line, which is the upper lip's junction with the face (*labiale superius*).
2. Interpolate to determine the corresponding 3D coordinates, using texture coordinates in the raw 3D file, and use these 3D locations to define a face frame (i.e object centred rather than camera centred frame).
3. Transform the computed nose position from the camera frame to the face frame.
4. Examine the within-class (single subject) repeatability of nose localisation in the face frame, using an RMS metric.
5. Use the average within-class RMS value to compare with the manual method and NPH methods.

5.4. Nose localisation results

The repeatability results of our SSR localisation method, the NPH method and the manual method are given in figure 9. We can clearly see that the NPH method is poor and that our SSR method slightly outperforms the manual method. In part, that is to be expected, since the manual method operates on raw vertices at the original mesh resolution (3-4mm), whereas the nose refinement method interpolates a higher density (2mm resolution) zero isosurface using the RBF model. The results do, however, inspire confidence in the method, and give repeatable results in the presence of noise. Finally, one has to remember that errors in manually locating face frame features and in 2D-to-3D registration appear across all of these results.

Test set	SSR values		Spin image 1		Spin image 2		Spin image 3		SSR histograms	
	Fails	% Pass	Fails	% Pass	Fails	% Pass	Fails	% Pass	Fails	% Pass
test A (515 images)	48	90.7%	185	64%	153	70.3%	152	70.5%	3	99.4%
test B (1121 images)	93	91.7%	400	64%	316	70.8%	339	70%	4	99.6%

Table 1: Nose identification results using five different methods




Figure 9. Nose location repeatability RMS(mm) in the three face frame dimensions

6. Conclusions

We have presented the spherically-sampled RBF (SSR) histogram, also called a balloon image, which is based on sampling RBF functions on concentric spheres, at arbitrary resolutions in 3D space. These representations are pose invariant and yield shape descriptors that do not require a potentially noisy local surface normal estimate. Also, they are relatively immune to missing parts, as the RBF function is defined everywhere in 3D space. Our experiments on nose vertex identification, using a challenging dataset of 1736 3D faces, indicate that these factors appear to be important when characterising high curvature surfaces in the presence of noise, missing parts and limited resolution data.

We have shown that it is possible to derive a SSR convexity value from the SSR histogram, which describes the volumetric intersection between a sphere and the object of interest (face). A notable issue here is that this feature, in essence, is derived as a summation, which has the effect of suppressing (averaging) noise, where many conventional 3D surface features, such as curvature, are based on differencing, whose effect is to amplify noise.

References

[1] J. Assfalg, A. D. Bimbo, and P. Pala. Spin images for retrieval of 3d objects by local and global similarity. In *Proc. 17th Int. Conf. on Pattern Recognition (ICPR'04)*, volume 3, pages 906–909, 2004.

[2] P. N. Belhumeur, J. Hespanha, and D. J. Kriegman. Eigenfaces vs. fisherfaces: Recognition using class specific linear

projection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(7):711–720, 1997.

- [3] P. Besl and N. D. McKay. A method for registration of 3D shapes. *IEEE Trans. Pattern Analysis and Machine Intelligence (PAMI)*, 14(2):239–256, 1992.
- [4] J. C. Carr, R. K. Beatson, J. B. Cherrie, T. J. Mitchell, W. R. Fright, B. C. McCallum, and T. Evans. Reconstruction and representation of 3d objects with radial basis functions. In *Proc. ACM Siggraph 2001*, pages 67–76, 2001.
- [5] C. Conde, R. Cipolla, L. J. Rodriguez-Aragon, A. Serrano, and E. Cabello. 3d facial feature location with spin images. In *IAPR Conf. on Machine Vision Applications (MVA'05)*, pages 418–421, 2005.
- [6] K. I. Chang, K. W. Bowyer, and P. J. Flynn. An evaluation of multimodal 2d+3d face biometrics. *IEEE Trans. PAMI*, 27(4):619–624, 2005.
- [7] C. S. Chua and R. Jarvis. Point signatures: A new representation for 3D object recognition. *Int. Journal of Computer Vision*, 25(1):63–85, 1997.
- [8] H. Q. Dinh and S. Kropac. Multi-resolution spin-images. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR'06)*, pages 863–870, 2006.
- [9] T. Heseltine, N. E. Pears, and J. Austin. Three-dimensional face recognition using combinations of surface feature map subspace components. *Image and Vision Computing*, 26(3):382–396, 2008.
- [10] A. E. Johnson and M. Herbert. Using spin images for efficient object recognition in cluttered 3d scenes. *IEEE Trans. PAMI*, 21(5):433–449, 1997.
- [11] Y. Li, W. A. P. Smith, and E. R. Hancock. Face recognition using patch-based spin images. In *18th Int. Conf. on Pattern Recognition (ICPR'06)*, pages 408–411, 2006.
- [12] X. Lu, A. K. Jain, and D. Colbry. Matching 2.5d face scans to 3d models. *IEEE Trans. PAMI*, 28(1):31–43, 2006.
- [13] N. E. Pears and T. D. Heseltine. Isoradius contours: New representations and techniques for 3d face matching and registration. In *3rd Int. Symposium on 3D Data Processing, Visualization and Transmission (3DPVT'06)*, University of North Carolina, USA, pages 176–183, 2006.
- [14] F. Stein and G. Medioni. Structural indexing: Efficient 3-d object recognition. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 14(2):125–145, 1992.
- [15] Y. Wang, C. Chua, and Y. Ho. Facial feature detection and face recognition from 2d and 3d images. *Pattern Recognition Letters*, 23(10):1191–1202, 2002.
- [16] C. Xu, T. Tan, Y. Wang, and L. Quan. Combining local features for robust nose location in 3d facial data. *Pattern Recognition Letters*, 27:1487–1494, 2006.