Using the CSM Correspondence Calculation Algorithm to Quantify Differences Between Surfaces

Elizabeth Guest[1,a] , Elizabeth Berry[1], David Morris[2]

[1]Centre of Medical Imaging Research, Leeds University
[2]Leeds Dental Institute, Leeds

contact email: e.guest@lmu.ac.uk

Abstract

In this paper, we apply the CSM (Correspondences by Sensitivity to Movement) algorithm to the problem of quantifying differences between surfaces. The basic principle of the CSM algorithm is to test the reliability of a match by moving the point of interest to see if it always matches to the same place. This method is used to compare pre- and post-operative range scans of the faces of patients, who have undergone surgery to their jaws. We show that, although locating closest points between surfaces is less sensitive to noise in the data, the CSM method is useful for indicating shape changes, especially in regions where the two surfaces are close.

1 Introduction

Accurate and robust methods for calculating correspondences between similar surfaces are necessary for a number of applications, such as evaluation of the effects of aesthetic surgery to the face and growth measurement. Point correspondence methods are also employed in surface registration algorithms and for developing point distribution models [1]. In this paper, we are particularly interested in methods for evaluating the effects of aesthetic surgery to the face.

Methods developed specifically for evaluating patients' faces before and after surgery include those proposed by McCance et al [2,3] and Coombs [4]. McCance measures the distances between two surfaces by firing radials from the centroid and calculating the intersection of each radial with both surfaces. This method has the disadvantage that the results can be misleading if the actual displacement is in a different direction to the radial direction. Coombes has developed a method that classifies each point on both surfaces into types of surface patch. Unfortunately, no method has been developed for automatically calculating correspondences between patches.

Several other surface correspondence algorithms have been proposed in the literature. Many of these simply take as corresponding point, the point giving the highest values of a match function [5-7]. These methods are unsatisfactory for evaluating changes in a surface because the point giving the best match will often not give a "sensible" corresponding

[a] Present address: School of Computing, Leeds Metropolitan University, Leeds

BMVC 2000 doi:10.5244/C.14.46

point. For example, a point in a featureless region of the surface should match well to many points in a large region on the other surface. In this situation, if the surfaces are registered, the most "sensible" corresponding point would be the nearest point in this region to the point being matched. Best match methods, however, will often choose a different point, which will often be in a different location in the region. This has the effect that neighbouring points will often not have close corresponding points.

Other methods include matching crest lines, defined as lines of maximum principal curvature [8] and spin images [9]. "Spin images" are obtained by projecting the image onto coordinate systems, which are unique for each point and determined by the surface normal at that point. Since each point has a unique spin image associated with it, and which is invariant to a rigid transformation of the object, this technique does not require that the images are registered before correspondences are calculated. The drawback with both of these methods is that they will only locate correspondences for distinctive points; correspondences for all other points have to be interpolated.

In this paper, we describe a surface correspondence algorithm, which uses the CSM (Correspondences by Sensitivity to Movement) algorithm along with a surface match function. This algorithm was initially employed within the ICP (Iterative Closest Point) registration algorithm [10] and shown to give better results than other correspondence algorithms [11]. It has the advantages that there is no segmentation, a reliability of the correspondence is calculated, and the corresponding points of points that are close, will in general, also be close. The performance of the CSM algorithm is evaluated with respect to quantifying the changes to patients' faces after corrective surgery to the mandible and maxilla, and compared with the performance of a method that locates the closest point on the surface to the point being matched.

2 Matching and Registration

The surface correspondence algorithm described in this paper is made up of two parts: a match function for comparing points on different surfaces; and the CSM (Correspondences by Sensitivity to Movement) algorithm which uses the match values given by the match function to calculate a corresponding point along with a measure of reliability. In the following, surfaces A and B are assumed to be triangular meshes and *nodes* on surface A will be matched to *nodes* on surface B. Note that the nodes on surface A can correspond to any *point* on surface B, although they are only matched to nodes on surface B.

There are three stages to the CSM algorithm:

1. Match a node (the matchnode) of surface A to all the nodes within a certain radius in surface B. (The radius defines the search region.)
2. Move the matchnode around a little, and for each position of the matchnode, calculate a "tentative corresponding point" on surface B. This produces a scatter of corresponding points that should be clustered on surface B, and that may be clustered along a line or around a point, or widely scattered on the surface.
3. Calculate a corresponding point and a measure of reliability by analysing the distribution of the tentative corresponding points. If the tentative corresponding points are scattered along a line, the corresponding point is the closest point on the line to the matchnode; otherwise the centroid of the scatter of points is used. The reliability is derived from the sum of squared distances of the second principal axis for the scatter of points.

2.1 The Match Function

The match function for matching a node on surface A to a node on surface B is based on three quantities: the curvedness [12], shape [12], and "relative angle". The curvedness and shape are standard measures based on the principal curvatures; the relative angle is the angle between the normal and the vector from the main axis of symmetry of the object to the node. This is a new measure that imposes a "texture" on the surface normals and thus helps to localise nodes in otherwise featureless regions. At first, the surface normals were compared to those of a sphere, but as the face resembles a sphere quite well, this provides little additional information about the surface. It was found that using a cylinder as a template works much better, giving better localisation to the nodes. Note that the relative angle measure, unlike surface normals, is independent of the pose of the object.

We define a separate measure for each of these quantities, and then multiply these measures together. Each component is matched using the following formulae:

Curvedness: $$r = \min\left(\frac{1+1000\times R_1}{1+1000\times R_2}, \frac{1+1000\times R_2}{1+1000\times R_1}\right)$$

where R_1 and R_2 represent the curvedness of the points being matched. The factor 1000 is used to stop the values being swamped by 1.0, which has been included to prevent division by zero. Note that although the curvedness varies with scale, this match function is independent of scale.

Shape: $$c = \exp\{-\frac{9}{2} \mid S_1 - S_2 \mid^2\}$$

where S_1 and S_2 represent the shape of the points being matched. The value 9/2 effectively gives a half width at half maximum for this gaussian as 1/3.

Relative angle: $$g = \exp\{-\frac{162}{\pi^2} \mid A_1 - A_2 \mid^2\}$$

where A_1 and A_2 are the relative angles of the points being matched. This gaussian has a half width at half maximum of pi/18, which is 10 degrees.

These formulae have been designed (via the use of gaussian functions for the latter two quantities) to give values in the range [0,1] to ensure that they have equal weight in the final matchvalue given by: $m = rcg$

2.2 Calculating Tentative Corresponding Points

If the matchnode from surface A is matched to all the nodes within a search region in surface B, a matchmap is obtained. A "tentative corresponding point" for the matchnode can be calculated by summing over the matchmap. This is achieved by the formula:

$$\underline{q}_i = \frac{\sum K_i \underline{x}_i}{\sum K_i}$$

where $\underline{x}_i$ is the vector from the matchnode to node i in the search region, and K_i is the "attraction", inversely proportional to distance, of node i given by

$$K_i = \frac{m_i}{1+|\underline{x}_i|^b} \times \frac{meandist_i}{n_i} \quad , \quad b \geq 2$$

where m_i is the match value for point *i*, b is a constant, $meandist_i$ is the mean distance from point *i* to all the nodes connected to it, and n_i is the number of nodes connected to point *i.* The second part of this formula is included to prevent areas with densely packed nodes having more influence than areas where the nodes are less densely packed. The constant *b* was set equal to 3.0 for this application.

2.3 Generating a Corresponding Point and a Measure of Reliability

The formulae given above enable us to calculate a single "tentative corresponding point" for each matchnode on surface A. If we move the matchnode a little and sum over the matchmap with the matchnode in this new position, we will in general, get a different tentative corresponding point. By repeating this procedure for fixed displacements of the matchnode, a scatter of tentative corresponding points will be generated. The distribution of this scatter of points is analysed to provide a corresponding point and a measure of reliability of the correspondence.

This distribution of points is analysed by calculating the principal axes and corresponding eigenvalues of the tentative corresponding points. As long as surface B can be locally approximated by a plane (which can be achieved for smooth surfaces by increasing the scale, or by decreasing the perturbation distance of the matchnode), the third eigenvalue will approximate the sum of squared distances of the tentative corresponding points to the surface. The first principal axis provides a line along which the points are scattered; and the second eigenvalue provides a measure of how closely the points are scattered along this line. If the ratio of the second and first eigenvalues is greater than a threshold (currently set at 0.5), the corresponding point is set to the closest point on the line to the matchnode; otherwise it is given by the centroid of the scatter of tentative corresponding points.

The measure of reliability of the match is based on the second eigenvalue since this is large when the points are widely scattered over surface B and small both when they are scattered along a line and when they are clustered near a point. To give the reliability value, the second eigenvalue is converted to the range [0,1] using the following formula:

$$w = \exp\left\{\frac{32v^2}{t^2 d^2}\right\}$$

where *d* is the maximum displacement undergone by the matchnode, *v* is the second eigenvalue, and *t* is the number of tentative corresponding points. For this application, the displacement *d* was set to 2.5mm.

3 Experiments

A Cyberware laser scanner 3030HRC (Cyberware Inc. Monterey, California) was used to scan two patients before and after surgery to their jaws. The maxilla of patient 1 was moved 11 mm forwards and 3.5 mm upwards; and the mandible was moved 7 mm backwards. Patient 2 had their mandible moved back by 5-6 mm and their maxilla moved 5 mm forwards and tilted so that the anterior and posterior regions were moved upwards by 2 mm and 4 mm respectively. Several scans of each patient were taken at each session. In addition, in order to provide data for checking that large differences are not found when the face has not been surgically altered, 10 volunteers, who were not undergoing surgery, were scanned twice on the same day.

The face was manually segmented from each scan, converted to a triangular mesh, and decimated so that the number of nodes was reduced to 15% of the original number. Scans were automatically registered using the CSM correspondence algorithm, combined with the ICP rigid registration algorithm [11]. The scans of the volunteers were registered using the whole face, whereas the pre- and post-operative scans were registered by registering the forehead and half of the eye orbits from the post operative scan to the whole face from the pre-operative scan.

Two methods were used to compare the differences between pairs of scans: the CSM correspondence algorithm described above and the closest point method. The closest point method simply calculates the vector from a matchnode on surface A to the closest point on surface B to the matchnode. The displacement vectors generated by these methods were labelled according to their direction compared to the normals of surface A into inward and outward pointing vectors. The displacements were displayed on surface A using warm (yellow-orange-red) colours for outward displacements and cold (green-blue-purple) colours for inward displacements. Note that although these methods produce displacement vectors, they do not give an indication of soft tissue movements as a result of surgery, but instead, quantify the differences between two surfaces.

Three sets of experiments were carried out:

1. The scans of the 10 volunteers were evaluated using the CSM and closest point methods in order to check that these methods do not give large differences when the face has not been surgically altered.
2. A pre-operative scan from each patient was matched to a post-operative scan using both methods. The resulting displacement vectors were compared to the known bone displacements by recording and averaging the actual displacements for 4 points in the regions of each of the 16 points shown in figure 1.
3. In order to evaluate the accuracy of the registration and reproducibility of the calculated displacements, four post operative scans of patient 1 were registered to a single pre-operative scan, also from this patient. Displacement vectors were calculated from the pre-operative to each post-operative scan and then each of these sets of displacement vectors was compared pairwise with the others.

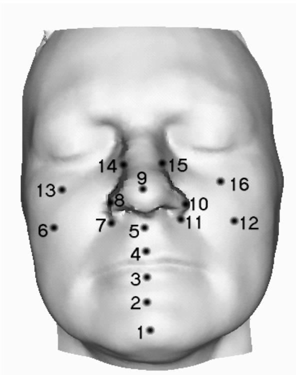

Figure 1: The 16 positions at which the vector displacements were measured.

	mean	std
CSM	1.91	1.38
Closest point	0.66	0.57

Table 1: The mean and standard deviation (in mm) of the differences between the surfaces for the scans of normal faces.

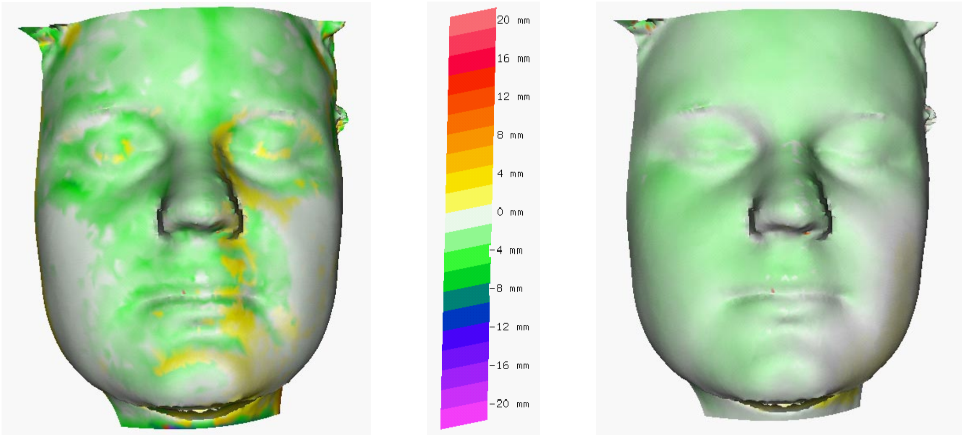

CSM Closest Point

Figure 2: The results of registering two scans of the same person taken on the same day.

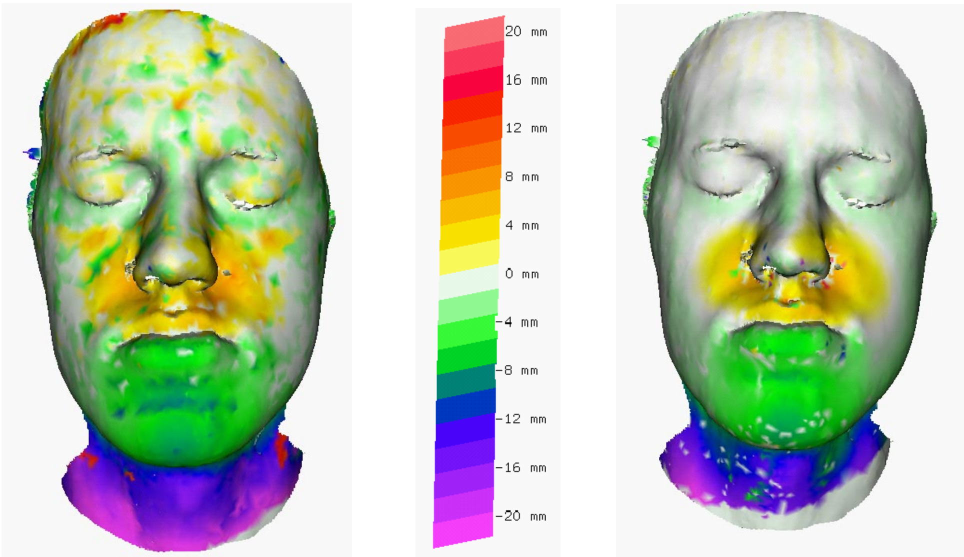

CSM Closest Point

Figure 3: The results of using the two methods to compare the pre and post-operative facial surfaces for patient 1.

4 Results

Table 1 shows that the closest point method gives the smallest surface displacement values. These values are comparable to those found in other work based on landmarks [13,14]. The displacement maps for one of the volunteers are shown in figure 2. The distribution of non-zero displacements shows that the scans are well registered and that the errors are probably due to measurement errors by the scanning equipment.

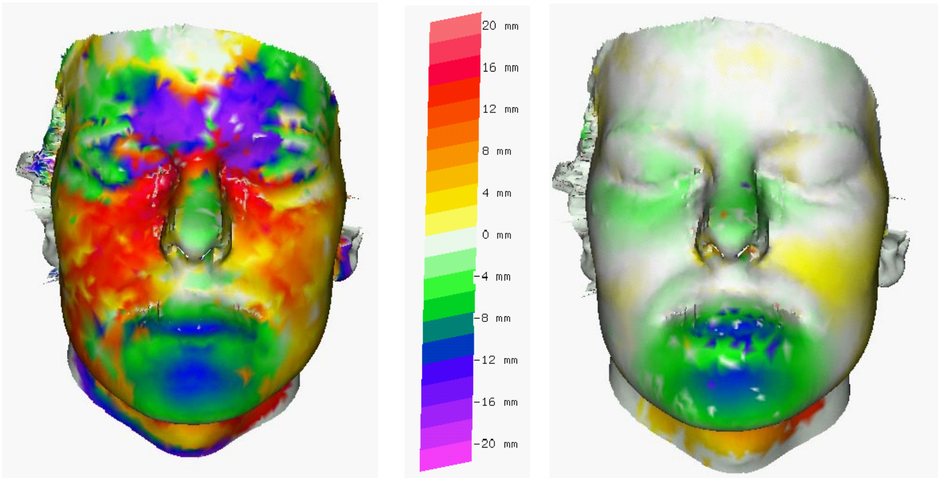

Figure 4: The results of using the two methods to compare the pre and post-operative facial surfaces for patient 2.

CSM								
	Post 1		Post 2		Post 3		Post 4	
	mean	std	mean	std	mean	std	mean	std
Post 1	0.00	0.00	3.58	3.55	2.48	3.14	2.77	3.54
Post 2			0.00	0.00	3.86	3.05	2.71	2.37
Post 3					0.00	0.00	3.07	3.33
Post 4							0.00	0.00
Closest distance								
Post1	0.00	0.00						
Post2	1.84	1.94	0.00	0.00				
Post3	1.24	1.60	2.22	2.26	0.00	0.00		
Post4	1.44	1.88	1.36	1.63	1.70	2.12	0.00	0.00

Table 2. The mean and standard deviation (in mm) for the displacements, calculated for four post-operative images that had been registered to a single pre-operative image. The displacements were calculated using the closest point method.

Figure 3 shows the results of using both methods to compare the pre- and post-operative scans of patient 1. The images in figure 3 show that the two scans are well registered because the foreheads show little displacement. Both methods show that the mandible was moved back by about 7 mm (point 1), which agrees closely with the given bone movement. Both methods also show that the maxilla was moved forwards, but the soft tissue displacements are less than the given bone displacement. The CSM method gives greater displacement in the cheek regions, suggesting that although the two surfaces are very close in this region, there has been some shape change.

The results of patient 2 are shown in figure 4. Both methods show that the chin was moved back, the bridge of the nose has moved down and that the mandible has dropped. Despite the large and complicated movement of the maxilla, the closest point method shows little difference between the pre- and post-operative scans in the cheek region, although there is some (asymmetric) indication that the area under the eyes has moved

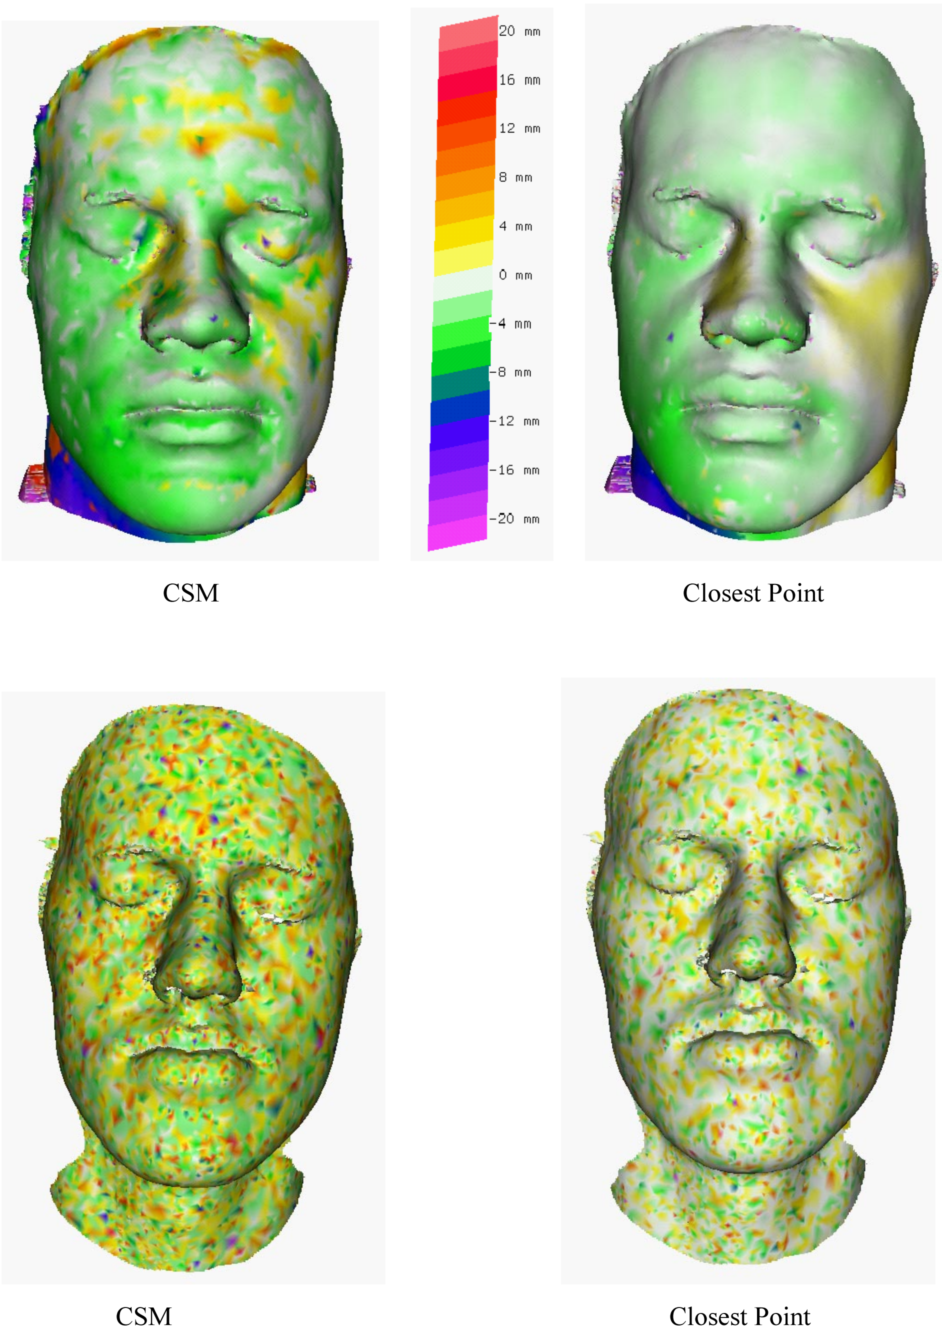

Figure 5. The results of experiment 3 for scans post 2 and post 3 (see table 4). The top row shows the results of calculating the displacements between post 2 and post 3. The bottom row shows the results of comparing the displacements calculated from the pre-operative image to each post-operative image.

down, while the cheek regions level with the alae have moved out by 2-3 mm. The CSM method, in contrast, shows major changes in the cheek regions. When the surfaces of the scans are examined, the shape of the cheeks looks quite different. They seem to have dropped and to have moved forwards relative to the nose. This movement is shown clearly only by the CSM method, which when the displacement vectors are displayed, shows a shape change in a downwards circular movement.

Table 2 shows the results of comparing displacements between a single pre-operative scan and four post-operative scans that have been registered to this pre-operative scan. Figure 5 shows the results for the worst pairwise comparisons. The top row of this figure, which shows the differences between two post-operative scans, shows that the registration is good, but that the nose appears to have shifted in the opposite direction to the rest of the face.

The second row of figure 5 shows the differences in the displacement vectors calculated between the pre-operative scan and two post-operative scans. Since these images have a mottled appearance rather than a systematic pattern, we conclude that any registration errors are not affecting the displacement calculation, and that the errors must be due to noise in the images. The argument for this conclusion is further strengthened by the fact that the differences in the displacement vectors are reduced when gaussian smoothing is applied to the triangular surface meshes before the displacements are calculated.

All the results for experiment 3 show that the CSM method gives larger discrepancies between the displacement vectors than the closest point method.

5 Conclusion

The results of the above experiments show that both the CSM and closest point algorithms are useful for quantifying the differences between surfaces, but that ideally, the information from each should be combined. Since the CSM method measures distances between corresponding points, where a corresponding point is a point on the surface with similar shape, this method can give an indication of shape change. This is seen particularly in the cheek regions of patient 2. Here, the closest point method shows little difference between the two surfaces, but the CSM method gives large displacements, suggesting that the shape of the cheeks has changed. This is corroborated by a visual examination of the pre- and post-operative scans.

Note that, although the CSM method provides a weight for each correspondence, these weights were not taken into account in these calculations. If these weights were to be taken into consideration, probably combined with some kind of smoothing/interpolation algorithm, it is likely that effects due to artefacts in the scans will have less effect, allowing better results to be achieved.

Acknowledgements

This research was performed in the Brite Euram project PISA (NR. BRPR CT97 0378). Partners in this project are: Materialise NV, Belgium (coordinator); Philips Medical Systems Nederland BV, EasyVision Modules – Advanced Development, the Netherlands; DePuy International Ltd, United Kingdom; Ceka NV, Belgium; Katholieke Universiteit Leuven, Laboratory for Medical Imaging Research and Division of Biomechanics and Engineering Design, Belgium; University of Leeds, Centre of Medical Imaging Research, United Kingdom; OBL, Paris, France.

The Cyberware scanner is supported by the Wellcome Trust (047808/Z/96/Z); and we thank Gillian Walker for acquiring the scans of volunteers with funding from the Wellcome Trust (VS98/LEE/003/CH/TG/JG).

References

1. Brett A, Taylor C. A framework for automated landmark generation for automated 3D statistiscal model construction. Information Processing in Medical Imaging: Proc 16th Int.Conf.(IPMI'99), Visegrad, Hungary, June 1999;376-81.
2. McCance AM, Moss JP, Fright WR, James DR, Linney AD. A three dimensional analysis of soft and hard tissue changes following bimaxillary orthognathic surgery in skeletal III patients. British Journal of Oral & Maxillofacial Surgery 1992;30(5):305-12.
3. McCance AM, Moss JP, Wright WR, Linney AD, James DR. A three-dimensional soft tissue analysis of 16 skeletal class III patients following bimaxillary surgery. British Journal of Oral & Maxillofacial Surgery 1992;30(4):221-32.
4. Coombes AM, Moss JP, Linney AD, Richards R, James DR. A mathematical method for the comparison of three-dimensional changes in the facial surface. European Journal Of Orthodontics 1991;13(2):95-110.
5. Feldmar, J. and Ayache, N. Rigid, Affine and Locally Affine Registration of Free-Form Surfaces. Sophia Antipolis. INRIA. 1994; 2220.
6. Pito R. A Registration Aid. Proceedings of International Conference on Recent Advances in 3-D Digital Imaging and Modeling 1997;85-92.
7. Weik S. Registration of 3-D Partial Surface Models Using Luminance and Depth Information. Proceedings International Conference on Recent Advances in 3-D Digital Imaging and Modeling 1997;93-100.
8. Subsol G, Thirion J-P, Ayache N. A Scheme for Automatically Building Three-Dimensional Morphometric Anatomical Atlases: Application to a Skull Atlas. Medical Image Analysis 1998;2(1):37-60.
9. Johnson AE, Hebert M. Using Spin Images for Efficient Object Recognition in Cluttered 3D Scenes. IEEE Transactions On Pattern Analysis And Machine Intelligence 1999;21(5):433-49.
10. Besl P, McKay N. A Method for Registration of 3-D shapes. IEEE Transactions On Pattern Analysis And Machine Intelligence 1992;14(2):239-56.
11. Guest E, Fidrich M, Kelly S, Berry E. Robust Surface Matching for Registration. Proceedings International Conference on 3-D Digital Imaging and Modeling 1999;169-77.
12. Koenderink J. Solid Shape. MIT Press; 1990.
13. Aung SC, Ngim RC, Lee ST. Evaluation of the laser scanner as a surface measuring tool and its accuracy compared with direct facial anthropometric measurements. British Journal of Plastic Surgery 1995;48(8):551-8.
14. Bush K, Antonyshyn O. Three-dimensional facial anthropometry using a laser surface scanner: validation of the technique. Plastic & Reconstructive Surgery 1996;98(2):226-35.