Faller RV (ed): Assessment of Oral Health.
Monogr Oral Sci. Basel, Karger, 2000, vol 17, pp 73–129

Analysis of the Morphology of Oral Structures from 3-D Co-Ordinate Data

V. Jovanovski [a], *E. Lynch*[b]

[a] Department of Adult Oral Health, St. Bartholomew's and the Royal London School of Medicine and Dentistry, London, and
[b] Restorative Dentistry and Gerodontology, School of Clinical Dentistry, The Queen's University of Belfast, UK

Abstract

A non-intrusive method is described which can be used to determine the forms of oral structures. It is based on the digitising of standard replicas with a co-ordinate-measuring machine. Supporting software permits a mathematical model of the surface to be reconstructed and visualised from captured three-dimensional co-ordinates. A series of surface data sets can be superposed into a common reference frame without the use of extrinsic markers, allowing changes in the shapes of oral structures to be quantified accurately over an extended period of time. The system has found numerous applications.

Introduction

There is a clinical need for objective quantification of the shapes of three-dimensional oral structures such as the morphology of teeth, the position and contour of the gingivae, the surface characteristics of fillings and their adaptation to the teeth in which they are placed. This need is particularly evident in the field of conservative dentistry which is concerned with all aspects of the conservation and restoration of teeth as well as the whole care of the patient. Additionally, such care requires assessment not only of shape, but also of changes in shape due to, for example, the wear of teeth and restorations, plaque accumulation, gingival recession or gingival inflammation.

Much of this assessment has previously been subjective and reliant on the clinician's experience. Examples are the USPHS criteria for evaluating

occlusal restorations [1] by which several aspects of a restoration are classified as ideal, acceptable or unacceptable based on the observations of two or more evaluators, or the Leinfelder method [2] in which an experimental cast is compared visually to standard casts calibrated for wear in steps of 100 μm.

Subjective methods are insensitive to small changes and produce few, if any, quantitative results. Consequently, numerous techiques for obtaining objective measurements have been developed. Sometimes the features of interest can be presented in a planar (two-dimensional) form, for instance by producing sectioned silicone replicas which permit lengths and angles to be measured from photographs magnified by a known factor [3]. However, the most useful and widely applicable techniques are those which provide numerical data in the form of three-dimensional co-ordinates. Such techniques include stereophotogrammetry [4, 5], reflex microscopy [6–8], laser interferometry [9, 10], scanning electron microscopy [11–14], microtomography [15–17], conforcal microscopy [18], infra-red cameras [19] and mechanical digitisers [20]. A review can be found in Chadwick [21] and a table of methods and accuracies in Bayne et al. [22].

From the point of view of conservative dentistry, the techniques of greatest interest are those which can provide three-dimensional co-ordinate data acquired from an entire tooth surface with sufficient density and accuracy to permit the construction of a computer model. Additionally, the speed of data acquisition must be sufficient to permit practical application in clinical research studies involving large numbers of subjects.

Each of the techniques listed above is well suited for its particular area of application, but fails to meet these requirements in one or more aspects. A satisfactory solution has been provided by a category of measurement apparatus known as 'co-ordinate measuring machines' (CMMs), many of which were developed in the 1980s for industrial quality control of manufactured parts. Most CMMs consist of a horizontal platform upon which the object to be measured (a tooth replica, for instance) is placed and 3-D co-ordinates are acquired from it by an active data acquisition component in the form of a mechanical stylus or optical probe.

A number of systems based on CMMs or devices of equivalent functionality have been employed in dental research [23–30]. Their application generally involves three distinct stages of work. Firstly, a replica or model of the study surface is produced. Depending on the specific system used, there may be a requirement that the material be non-elastic, electrically conductive or of a particular colour and reflectivity. Secondly, co-ordinate data are acquired from the model surface. Finally, the data are analysed to produce quantitative results.

The published reports vary in the emphasis and amount of detail given relating to the methodology, but it is evident that the analysis of acquired co-ordinate data has frequently been problematic due to the unavailability of software which would be specifically tailored to the needs of dental research. The use of software intended for other application areas can render the process impracticably laborious. Where specialised software has been written, little is reported about the mathematical methods employed. Mitchell and Chadwick [31] write that 'although clinical results derived from such techniques are reported within the dental literature, none of the dental papers to date has described fully the underlying theory'.

With this in mind we present a detailed account of the methodology, supporting software and clinical application of the system for dental co-ordinate metrology developed at St. Bartholomew's and the Royal London School of Medicine and Dentistry [32]. Emphasis is given to description of the supporting mathematical models and the software which permits researchers to perform their own analyses by providing: (1) visualisation of surfaces with realistic rendering to enable the identification of anatomical features of interest; (2) construction of cross-sections and measurement of lengths and angles on such cross-sections; (3) quantification of changes in morphology by comparison of sequential replicas, particularly depth and volume changes on hard tissue and gingivae, and (4) output of results in graphical and numerical formats.

Particular attention is paid to the 'superposition' (or 'registration') of data sets obtained from a sequence of replicas of the same oral structure taken at different times. Such data sets do not generally lie in the same frame of reference and their comparison requires the application of co-ordinate transformations (rotations and translations) by which corresponding anatomical features are brought into coincidence.

The material is presented here in the following sequence, which reflects the progression from clinical replication to numerical results:

Replication of the study surface

↓

Co-ordinate data acquisition

↓

Construction of a computer model of the study surface

↓

Superposition of sequential data into a common reference frame

↓

Measurement

Replication of the Study Surface

Materials and Methods

The physical characteristics and mode of operation of CMMs preclude their use directly in the mouth. However, the provision of dental care invariably involves the replication of oral structures at some stage and dentistry has developed materials for the accurate and safe replication of oral structures, together with clinical expertise to use them effectively [33, 34]. Silicone impression materials are ideal for the purpose in hand, having been widely researched to determine their safety for use in the mouth, their elasticity, their ability to replicate very fine detail and their dimensional stability [35, 36].

The replicas in our studies were produced by a procedure similar to the standard routine followed in patient care but appropriately modified to address the special considerations that arise from the intended purpose of measuring changes in surface morphology. In studies involving sequential replicas of the same oral structure over a period of time it was desirable that the co-ordinate data capture be performed with all replicas having approximately the same orientation. In order to recreate not only the surface of interest but also its orientation, a special impression tube was designed which could be repositioned within the mouth in the same orientation each time.

An aspect of a tooth was selected as appropriate for the study to be undertaken, usually the buccal or occlusal surface of a tooth, and a plaster model was constructed from an alginate impression (fig. 1). A 9-mm length of accurately machined square section, rigid brass tubing of 12- or 16-mm nominal internal width was selected according to the size of the area to be studied. The surface of the model was lubricated and the impression tubing positioned on the plaster model over the surface to be replicated, so that its long axis was approximately normal to that surface (fig. 2). Its periphery was modified as necessary using a stone or a file so as to approximate the surface reasonably well. A polymethyl methacrylate resin was applied to the outside of the tubing to ensure that there were no gaps between the surface of the tooth and the tubing and to facilitate the correct locating of the tube in the mouth itself on one or many occasions.

On a subsequent patient visit, the quadrant containing the study surface was isolated and the tooth was gently dried with air. The tube was then located in its position in the mouth, held firmly in place and a polyvinyl siloxane dental impression material (Kerr Extrude, SDS Kerr, Orange, Calif., USA) was then injected onto the tooth surface within the tube to fill the latter to excess (fig. 3).

When the impression material had polymerised, the tube containing the silicone impression was removed from the mouth. Approximately 3 min after

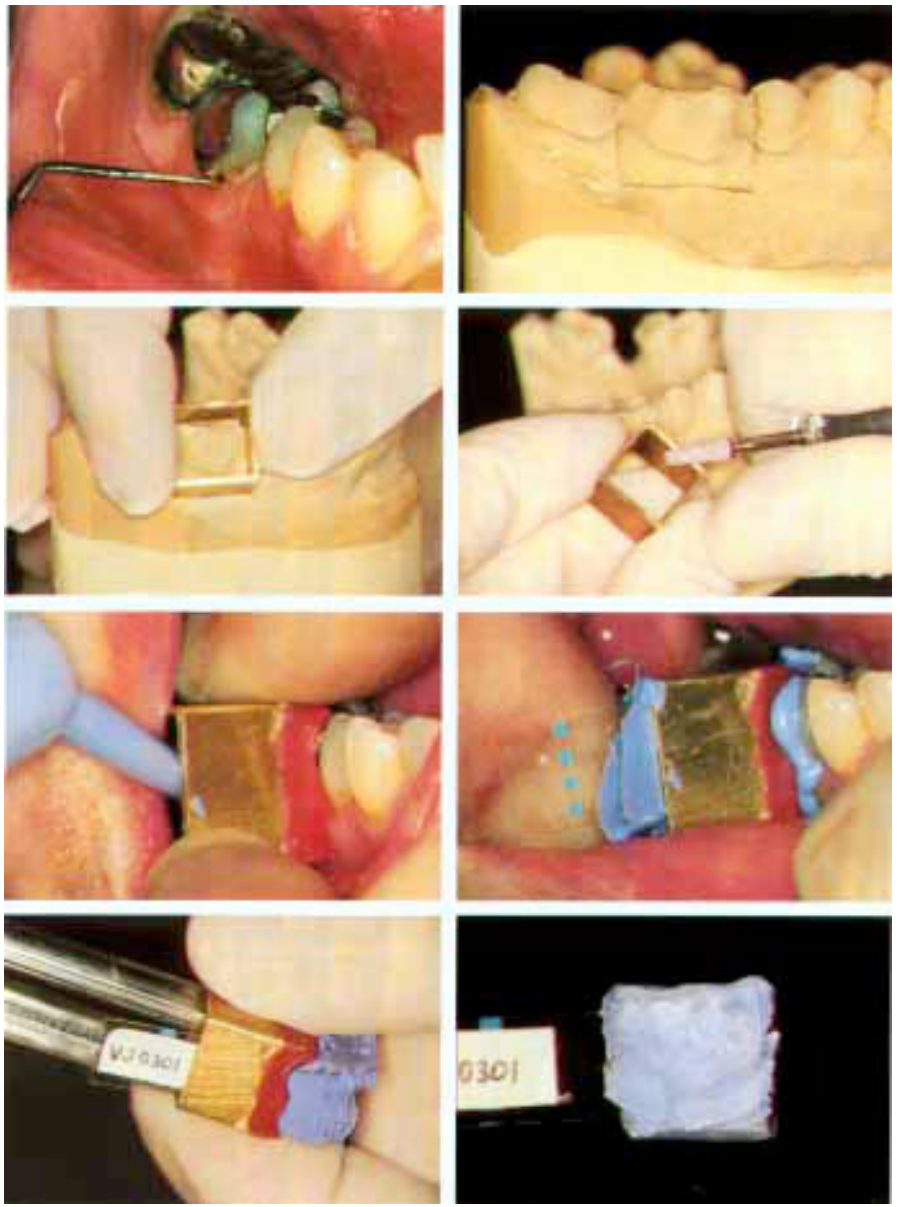

Fig. 1. The surface of interest identified in vivo and on the plaster model.
Fig. 2. Fabrication of the impression tube.
Fig. 3. Replication of the study surface.
Fig. 4. The completed replica.

Table 1. Repeatability of hard tissue replication

Study date	Number of pairs	RMS/point μm (mean ± SD)
1996	112	8.4 ± 2.1
1996	30	7.7 ± 1.8
1997	63	6.5 ± 2.5
1997	18	5.9 ± 1.4

this removal the base of the brass tube was trimmed flush with the edge of the tubing using a sharp blade and the block of silicone extruded from it. The tube was then available for further use on a subsequent occasion and the replica was ready for study (fig. 4).

Results

The physical properties of silicone-based impression materials have been researched extensively [36–40], showing their dimensional accuracy and long-term stability to be in the region of 0.1%. However, the replicas used in this work were not produced under the controlled conditions available in a laboratory, but in a clinical environment where they might be subjected to many additional influences. It was therefore considered necessary to perform our own evaluations of the practically achievable repeatability of sequential replicas of the same oral structure.

Sequences of repeat replicas were digitised and the data sets superposed into a common reference frame. Pairs of data sets were compared by evaluating the z co-ordinates of each of the data sets at 2,500 x–y locations distributed on a rectangular grid and recording the differences (which should ideally be zero). The measure of repeatability was defined as the root mean square (RMS) of the z differences, customarily termed the 'RMS/point' [30, 41]. Bearing in mind the procedures carried out to arrive at superposed mathematical models of the replica surfaces, the repeatability values obtained were all-inclusive, being influenced not only by the replication accuracy but also by the accuracy of the co-ordinate measurements, the superposition, as well as the instability of any soft tissue, e.g. the gingivae.

Hard Tissue

Table 1 shows the repeatability of replication of hard tissue over the course of several studies involving a total of 223 replica pairs taken at intervals ranging from 10 min to 6 months. Values which were found to differ significantly from these could invariably be attributed to the clinical procedure, for instance distortion due to failure of the replica material to polymerise fully.

Soft Tissue

Repeatability in the replication of gingivae was assessed on a series of 10 replicas taken in immediate succession. The pairwise RMS/point values were found to be between 29.0 and 64.5 (mean 41.8) μm. It is not currently clear what proportion of these differences in the sequential replicas of gingivae can be attributed to the physical distortion caused by the pressure applied during replication, and how much to changes in the gingival contour resulting from other factors including alterations in blood flow, desiccation or reactions to the chemicals within the impression material prior to its polymerisation. However, the repeatability of replication of the gingivae is sufficiently good to permit the investigation of changes in their contour caused by factors of interest in our studies, since such changes are generally considerably greater.

Plaque

The repeatability of replicas made in order to determine the location and thickness of plaque were comparable to those made on clean teeth. There was some concern that the plaque may be removed or disturbed by the replication process, but measurements of the repeat replicas indicated that any such effects, if they had occurred, would be too small to detect with our system. Furthermore, the clinical applications have confirmed that the accuracy is sufficient, for instance, to demonstrate the progressive accumulation of plaque over a period of 21 days and to detect statistically significant differences between the efficacy of two plaque-removal procedures.

Data Acquisition

CMM and Optical Probe

The central data acquisition equipment is a Merlin Mk II 750 Co-ordinate Measuring Machine (International Metrology Systems Ltd., Livingstone, UK) fitted with a Renishaw (Renshaw plc, Wotton-under Edge, UK) OP2 optical probe (fig. 5). The dimensions of the machine permit accommodation of a wide variety of differently sized objects, combining a large measuring volume ($500 \times 750 \times 500$ mm) with high accuracy throughout that volume. The worktable consists of a block of solid granite whose mass reduces the effects of environmental vibration. The position of the probe in each axis is determined to a resolution of 0.5 μm by a photocell which detects patterns on a stainless steel/glass grating. The accuracy of the position in each axis is 4.0 μm + L/275000 where L is the distance (in metres) travelled from an initial starting point. Control is provided by an IBM-PC compatible

Fig. 5. The CMM and OP2 probe.

computer interfaced to the electronic desk unit which integrates the functions of the CMM and probe.

The Renishaw OP2 optical probe operates on the principle of 'optical triangulation' (fig. 6). A laser diode generates infrared laser radiation with a wavelength of 830 nm which is focused onto the surface, producing a spot 25 μm in diameter. The light scattered from the surface is detected by a position-sensitive device from which the surface height is computed. Each probe is individually calibrated and the calibration data stored in the controller firmware. The vertical measurement range of the probe is ± 2 mm but this is effectively extended by the vertical motion of the probe column. Further flexibility in orienting the probe is provided by a multidirectional mounting which permits rotation of the probe around two axes. The maximum measure-

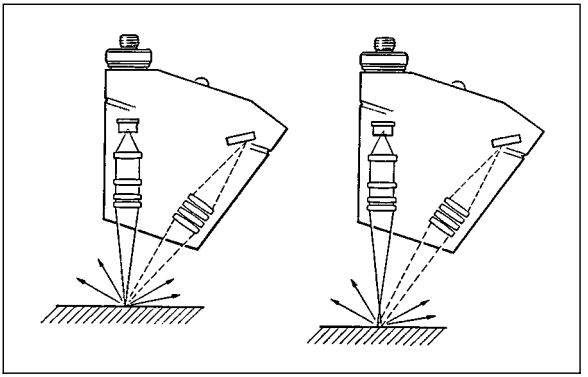

Fig. 6. Principle of operation of the OP2 probe. Copyright © 1988 Renishaw plc. Reproduced with permission.

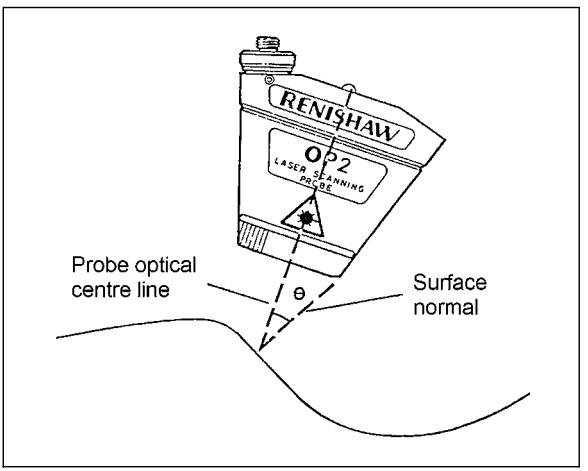

Fig. 7. The angle between the optical centre line and the surface normal. Copyright © 1988 Renishaw plc. Reproduced with permission.

ment frequency is 50 readings/s but in the integrated Merlin/OP2 configuration the optimal speed was determined to be 10 readings/s.

Two primary characteristics of a measurement device are its 'repeatability' (variations between repeated measurements) and 'accuracy' (how close the measurement is to the true value). The OP2 probe has a quoted repeatability of 2 μm. Its accuracy, however, depends on several factors, foremost of which is the angle between the probe's optical centre line and the surface normal (fig. 7), termed the 'gradient'. As the gradient increases, the intensity of the

scattered light which reaches the detector decreases and the focused spot becomes elongated to an ellipse, which reduces the detection reliability. An optimal arrangement is achieved when the surface normal coincides either with the optical axis of the laser diode or with that of the detector and in such cases its accuracy is 5 μm.

A replica is digitised by programming the probe to move across the surface tracing a pattern of parallel 'scan lines'. The data point spacing (or 'scanning pitch') is most commonly selected to be 100 μm, with dimensions of a typical region of interest being 10×10 mm. Such a region requires around 25 min of processing time. The worktable can accommodate up to 36 replicas, placed on a 6×6 grid of squares, which can be processed sequentially without operator intervention. On completion of a batch of replicas the data sets are transferred to shared storage on a Novell NetWare (Novell Inc., Utah, USA) server where they can be accessed by all authorised users.

Repeatability and Accuracy

Familiarity with the capabilities and, equally importantly, the limitations of a measurement system allows it to be used efficiently while avoiding any temptation to extract meaningful information from unreliable data. The repeatability and accuracy are a complex combination of many factors and although it is convenient to summarise them in terms of single values, such summaries are not always realistic indicators of performance.

A series of experiments to establish the repeatability and accuracy of measurements in practice were carried out under standard operating conditions and on surfaces which were representative of those encountered in the clinical applications.

Repeatability was assessed by performing multiple measurements of a single replica. Since the OP2 probe is sensitive to the surface gradient, a labial surface which contained both mild and severe gradients was chosen (fig. 8).The replica was digitised four times. Within each region of interest (A–D), the z co-ordinates were evaluated at 2,500 locations and the standard deviation was computed for each location. Repeatability for each region was defined as the average of the 2,500 standard deviations. The results are given in table 2. As expected, measurements on steeper gradients were less repeatable, except in the case of region D which was not located on the smooth hard tissue.

The accuracy was determined by measuring objects of known shape and dimensions and also the replicas of those objects made from standard impression material. On a sphere of known diameter, the acquisition errors manifested themselves as progressive systematic deviations from the true values in accordance with the graph in figure 9, plus an additional random component which

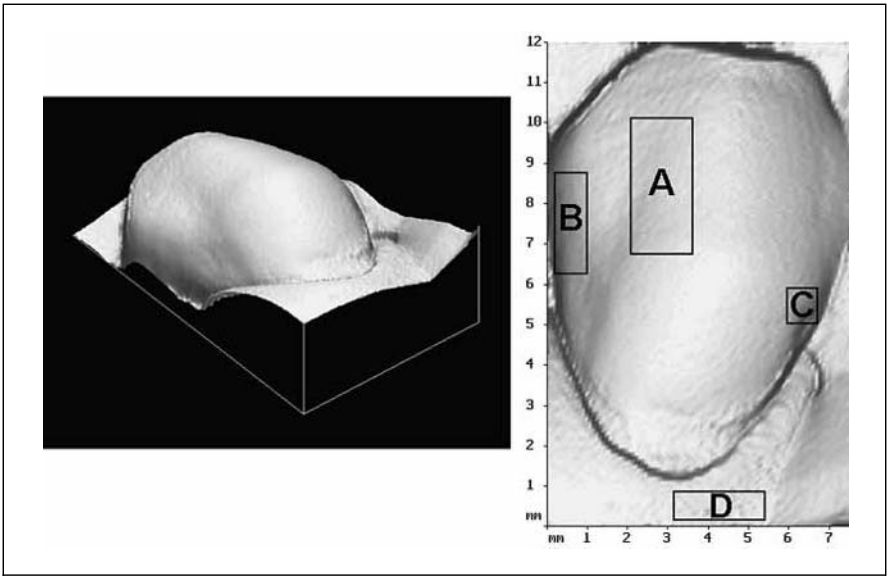

Fig. 8. Regions with different gradients.

Table 2. Repeatability on different gradients

	Region			
	A	B	C	D
Gradient	20°	52°	70°	20°
Repeatability, μm	2.4	4.0	5.3	4.5

is of a similar magnitude to the errors on flat, optimally oriented surfaces [42, 43]. The result of measuring a smooth surface was, therefore, again a smooth surface, although slightly distorted in regions with steep gradients. This had the desirable consequence of preserving reproducibility. Thus, when comparing sequential replicas of a single tooth, the difference between the two replicas could be determined to a greater accuracy (generally 5 μm) than could their true shapes.

It was also found that the probe's accuracy was not omnidirectional with respect to rotation around the optical centre line and that optimal results could be achieved if the probe travelled in a direction which was normal to the plane defined by the laser beam and the optical centre line [44]. The

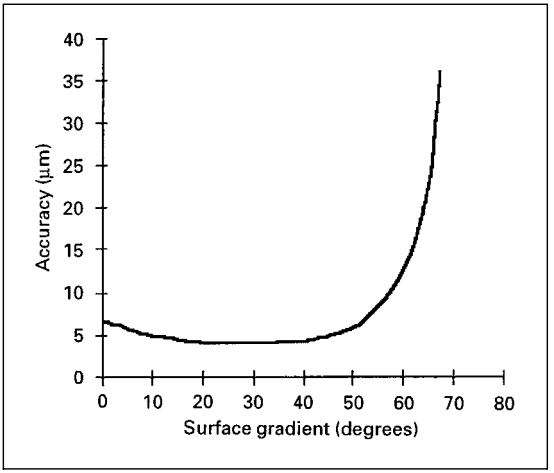

Fig. 9. OP2 probe accuracy vs. gradient.

standard procedure for digitising replicas was modified to take this aspect into account.

If a surface with a variable gradient is to be digitised without continually reorienting the probe, the best compromise is to orient the probe in such a way that the optical centre line is normal to the 'mean best fit plane' through the surface. Some surfaces can be processed using two or more probe orientations, but doing so requires additional operator effort and is not used routinely. In most applications involving single surfaces of teeth, the region of interest does not extend over too large a range of gradients and the replica can be positioned on the worktable in such a way that it is presented to the probe in a favourable orientation.

Although the error in z co-ordinates acquired on steep gradients may be relatively large (assuming that the probe is oriented so that its optical centre line coincides with the z axis), the errors of computed 'distances' are smaller when the dimension of interest is in a direction normal to the surface, for instance when measuring the loss of hard tooth tissue. If Δz is the acquisition error for a point on a flat surface inclined at angle α to the horizontal, then the error of measured distances Δd is only affected by that component of the acquisition error which lies in the direction of the dimension of interest ($\Delta d = \Delta z \cos \alpha$). If a replica is digitised in two different orientations, the average probe-to-surface angle will also differ and the measurements at any particular location on the replica will have been made under different accuracy conditions. Similar considerations apply when digitising two sequential replicas of the

same oral structure. For purposes of quantifying morphological change, the two data sets can be brought into a common reference frame but doing so does not compensate for inconsistencies caused by acquisition errors which arise from the different probe-to-surface gradients. In practice this aspect proved not to be a problem for misorientations of up to 10º, an accuracy of positioning which is easily achieved by the operator.

The repeatability of measurements was found to be lower on regions whose contour changes rapidly over an interval comparable to the scanning pitch. This is not due to acquisition errors, since it occurs even with a theoretically perfect digitiser, but a consequence of the sparse spacing of the data points relative to the complexity of the surface contour. In relation to oral structures, this effect can occur to a small extent on gingivae, shown by the reduced repeatability for region D of figure 8. However, changes in gingival contour are evaluated over a sufficiently large region (>1 mm^2) for the height uncertainty to be greatly reduced by averaging. Thus, even with high-accuracy digitising systems, on certain types of surface features the sampling density is the primary determinant of accuracy and repeatability. The scanning pitch therefore needs to be appropriate to the surface morphology and the desired detail of information. In the literature, assessments of accuracy and repeatability of individual measurement systems are generally produced by measuring a flat surface, perhaps at several different inclinations. Such assessments should not be taken to mean that the values so provided are valid for 'real' tooth replicas, or that they are constant over the entire surface.

Reconstruction of Surfaces from Co-Ordinate Data

The data set obtained by digitising a replica consists of individual 3-D points (fig. 10). For the purposes of visualisation and measurement such a representation is not the most convenient or practical. For instance, one method of computing the volume enclosed between two surfaces is based on evaluating the difference in their z co-ordinate for a given x and y. This approach is not possible if the two surfaces were digitised over different x–y grids. The co-ordinate data therefore need to be represented in a form appropriate to the intended analyses. Since the orientation and movement of the laser probe during scanning are such that only one z co-ordinate is sampled at each (x, y) location, the natural form of a continuous surface that models the acquired data is a single-valued function $z = f(x, y)$.

The aim of surface reconstruction is to find such a function $f(x, y)$ which passes through the measured points exactly (interpolation) or as closely as sensibly possible (approximation) and which models the underlying surface

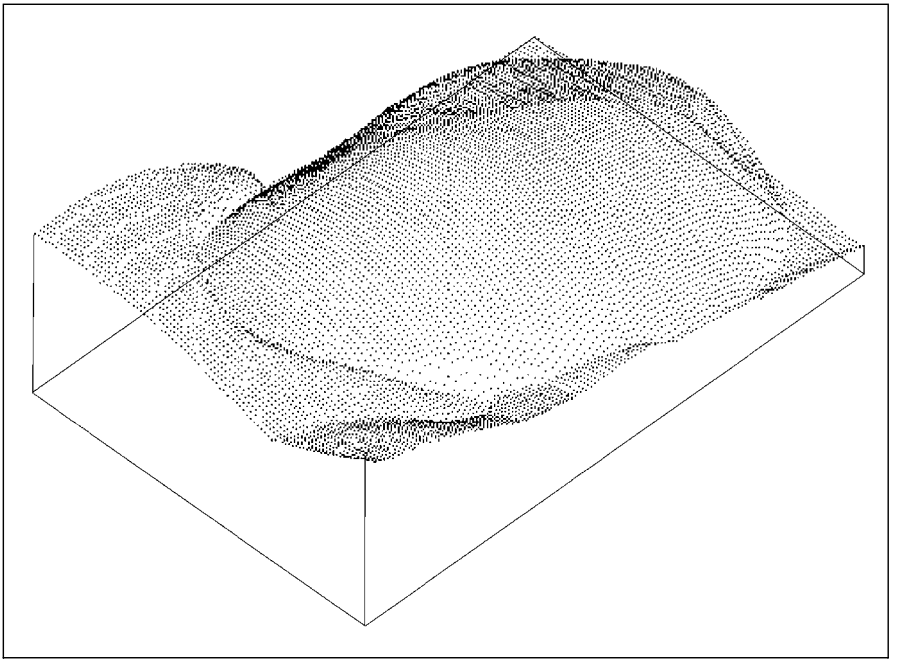

Fig. 10. A data set consisting of 11,432 points acquired from a labial tooth surface.

with sufficient accuracy in the regions between data points. The degree of success with which this can be done is dependent, amongst other things, on the accuracy of the measurements, their density (which should be sufficiently high to represent the smallest features of interest) and their (x, y) locations (more complex and time-consuming for non-uniform data). From a practical viewpoint, it is also desirable that the computation and subsequent evaluation of the function f be efficient and numerically stable.

Polynomial Splines

Consider the interpolation of a set of 2-D data points $\mathrm{A}=\{(x_1, y_1), (x_2, y_2), ..., (x_n, y_n)\}$ obtained from measurements of a curve defined on an interval $[x_{\min}, x_{\max}]$ satisfying $x_{\min}=x_1<x_2<\cdots<x_n=x_{\max}$. We wish to find a function $y=f(x)$ defined on the same interval which passes through all the points of A.

The simplest solution is provided by the piecewise-linear function whose graph can be constructed by joining the successive points of A with line segments. This function is continuous but it is not smooth and therefore it cannot be expected to model curved lines accurately.

On the other hand, the n-th degree polynomial

$$a_0 + a_1 x + a_2 x^2 + \cdots + a_n x^n \tag{1}$$

whose coefficients a_i can be determined formally by solving the system of n linear equations obtained by substituting each (x_i, y_i) in turn into expression (1), passes through all the points of A, but its value between the data points may often exhibit spuriously large oscillations. Its behaviour is non-local in the sense that each coefficient is dependent on all the data points and even a small change in the location of any of the points will in general alter the appearance of the entire curve.

An intermediate scheme which is both sufficiently flexible and stable is one whereby a function f is constructed by defining different polynomials on each subinterval $[x_i, x_{i+1}]$. These polynomials can be of a low degree (often cubic) and conditions can be specified at the joins in such a way that the resulting curve has continuity not only of value, but also of its first and possibly higher-order derivatives. Functions constructed in this way are known as 'polynomial splines', a term which derives from the analogy with the draughtsman's tool of the same name [45].

The concept of polynomial splines can be extended to higher dimensions. Polynomial spline surfaces have found wide application in a diversity of areas including computer-aided design [46], reverse engineering [47, 48], terrain mapping [49, 50] and medicine [51]. On this basis they were considered appropriate for modelling oral structures whose shapes are almost arbitrarily complicated and which belong to the category of 'free-form' surfaces about which very few assumptions can be made. A concise introduction to the mathematical theory and practical implementation of polynomial splines can be found in Cox [52, 53], and with more detail in de Boor [54].

Polynomial Spline Curves

Polynomial spline curves can be expressed as linear combinations of basis functions known as 'B-splines'. In defining a spline on an interval $[x_{\min}, x_{\max}]$, we will consider that the interval is partitioned by N points $\lambda_1, \ldots, \lambda_N$ termed 'interior knots':

$$x_{\min} < \lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_N < x_{\max}$$

and that additional points, termed 'exterior knots', are defined such that

$$\cdots \leq \lambda_{-1} \leq \lambda_0 \leq x_{\min}, \quad x_{\max} \leq \lambda_{N+1} \leq \lambda_{N+2} \leq \cdots$$

Most commonly, the exterior knots are selected such that $\lambda_j = x_{\min}$ if $j < 1$ and $\lambda_j = x_{\max}$ if $j > N$ and this will be assumed to apply here.

The interval $[x_{\min}, x_{\max}]$ is divided into $N+1$ disjoint subintervals $I_0, I_1, ..., I_N$ defined by

$$I_j = \begin{cases} [\lambda_j, \lambda_{j+1}) & j = 0, 1, ..., N-1, \\ [\lambda_j, \lambda_{j+1}] & j = N. \end{cases}$$

A B-spline of order n, denoted by $N_{n,j}(x)$, is defined on the knots $\lambda_{j-n}, \lambda_{j-n+1}, ..., \lambda_j$ by the recurrence relation due to Cox [55] and de Boor [56]:

$$N_{n,j}(x) = \left(\frac{x-\lambda_{j-n}}{\lambda_{j-1}-\lambda_{j-n}}\right) N_{n-1,j-1}(x) + \left(\frac{\lambda_j - x}{\lambda_j - \lambda_{j-n+1}}\right) N_{n-1,j}(x) \qquad \text{if } n > 1$$

and

$$N_{1,j}(x) = \begin{cases} 1 & \text{if } x \in I_{j-1}, \\ 0 & \text{otherwise.} \end{cases}$$

The B-spline $N_{n,j}(x)$ consists of polynomial 'pieces' of degree $n-1$ so, for instance, B-splines of order 4 are referred to as cubic B-splines. It is non-zero only on the interval $(\lambda_{j-n}, \lambda_j)$ and therefore if a function is represented as a linear combination of B-splines, altering any one of the coefficients will only alter the function locally, over the same interval.

A polynomial spline curve $s(x)$ of order n defined on a set of knots $\lambda_1, ..., \lambda_N$ is a linear combination of B-splines

$$s(x) = \sum_{j=1}^{q} c_j N_{n,j}(x)$$

where $q = N+n$ and $c_1, ..., c_q$ are 'B-spline coefficients'. Interpolation of a set of data points with a spline curve thus consists of selecting an appropriate grid and then determining the unknown B-spline coefficients from the conditions $s(x_i) = y_i$. Figure 11 shows a cubic spline curve

$$s(x) = \sum_{j=1}^{11} c_j N_{4,j}(x)$$

with $N = 7$ interior knots. The B-spline $N_{4,6}(x)$ is shown in bold.

Polynomial Spline Surfaces

A polynomial spline surface is defined over a rectangular region $[x_{\min}, x_{\max}] \times [y_{\min}, y_{\max}]$. Taking each axis separately, denote by $M_{n_x,i}(x)$ the B-splines of order n_x for a set of N_x interior knots $\lambda_1, ..., \lambda_{N_x}$ on $[x_{\min}, x_{\max}]$ and by $N_{n_y,j}(y)$ the B-splines of order n_y for a set of N_y interior knots $\mu_1, ..., \mu_{N_y}$ on $[y_{\min}, y_{\max}]$. A polynomial spline surface $s(x, y)$ is defined as the sum of products of B-splines

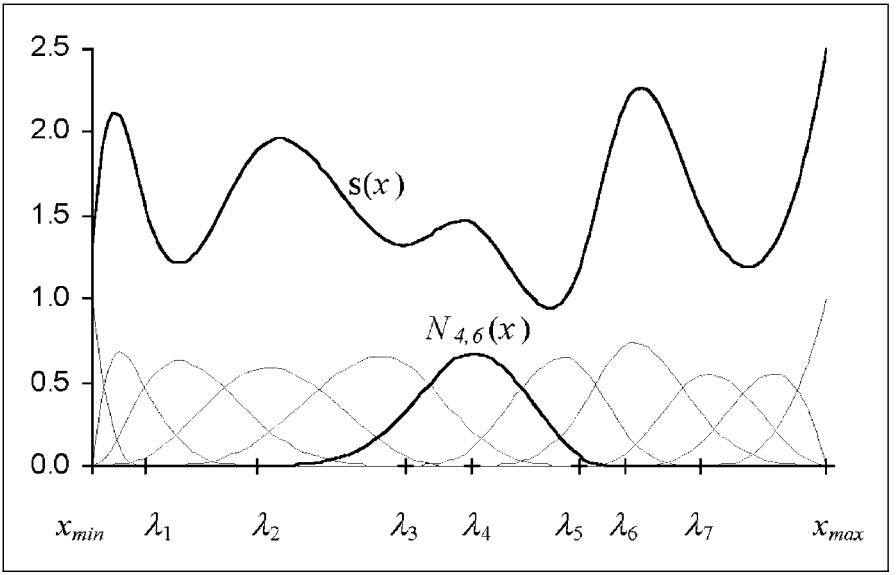

Fig. 11. A spline curve and the B-splines of which it is composed.

$$s(x) = \sum_{i=1}^{q_x} \sum_{j=1}^{q_y} c_{ij} M_{n_x,\, i}(x) N_{n_y,\, j}(y)$$

with coefficients $\{c_{ij} \mid i = 1, \ldots, q_x, j = 1, \ldots, q_y\}$ where $q_x = N_x + n_x$ and $q_y = N_y + n_y$. The process of constructing the interpolant of a set of data points (x_i, y_i, z_i) consists of finding the unknown coefficient matrix **C** by solving the system of linear equations resulting from the conditions $s(x_i, y_i) = z_i$. If the x–y locations of the data points lie on a rectangular $m_x \times m_y$ grid, then the coefficients **C** can be determined extremely rapidly and efficiently [53] by taking advantage of the separability of the system into the form $\mathbf{A}_x \mathbf{C} \mathbf{A}_y^{\mathrm{T}} = \mathbf{F}$ where $\mathbf{A}_x$ is the $m_x \times m_x$ matrix consisting of B-spline basis functions in x evaluated at the x gridlines, and $\mathbf{A}_y$ similarly for y.

Interpolation of CMM Data

As a consequence of the way in which a spline surface is defined as a sum of products of univariate B-splines, interpolation is straightforward only if the data points lie on a rectangular x–y grid. This is not the case with CMM data whose spacing is controlled by two parameters: the speed of motion of the probe column, and the frequency of readings of the OP2 probe. The two parameters are computed by the system software from a user-specified desired point spacing, but the actual achieved spacing is not guaranteed to be equal

to this value, or to be regular. Additionally the laser beam need not be oriented vertically in which case the scan line is not straight but dependent on the surface contour. Whilst some CMMs are capable of producing data with constant spacing, this is not a great advantage since the rectangular distribution is generally lost if a data set is rotated in order to superpose it onto another. It was therefore necessary to consider methods appropriate for data points which do not lie on a rectangular grid.

Although any of the many existing methods for gridding generally scattered data could be applied, advantage can be taken of the fact that the data points are distributed on approximately straight lines [57–60]. The data set is regularised first in the direction of x, then in y by two series of successive univariate interpolations, resulting in a new data set, representative of the original, but which lies on a rectangular grid. This new data set is then interpolated with a bicubic spline surface as described above. Details are given by Jovanovski et al. [32]. The method of arriving at a surface representation by successive interpolations may appear to be a rather indirect way of constructing an approximant but there were practical reasons for this approach, foremost of which was the speed of execution. The software related to this project was required to perform with an acceptable speed on a wide variety of standard, modestly specified computers and this aim was achieved. A data set consisting of 10,000 points requires just 5.6 s of processing time on a Pentium 166 processor.

Accuracy of the Spline Interpolant

The spline interpolant of the gridded data points does not, in general, pass exactly through the original data points. The question therefore arises whether it is a satisfactory model. The data points are subject to measurement errors and there is therefore no great advantage in reproducing them exactly. It is suggested by Cox [53] that 'an approximant should be sought that reproduces the data to the accuracy warranted by the data'. To determine if such an approximant was being attained, the residuals were evaluated on several types of replicas which were typical for the intended clinical applications and scanned with a pitch of 100 μm. It was apparent that the fit was worse in some regions than in others and that these regions were associated with detailed morphological features.

The root mean square residuals were found to be 0.09 μm on hard tissue, 1.95 μm on plaque and 1.06 μm on gingivae. These values were considered satisfactory. While it is possible to improve them further, for instance by applying the iterative method described by Anderson and Mason [59] and Anderson and Pink [61], it should be borne in mind that such improvements would be only in the numerical sense of fitting the surface model to the digitised

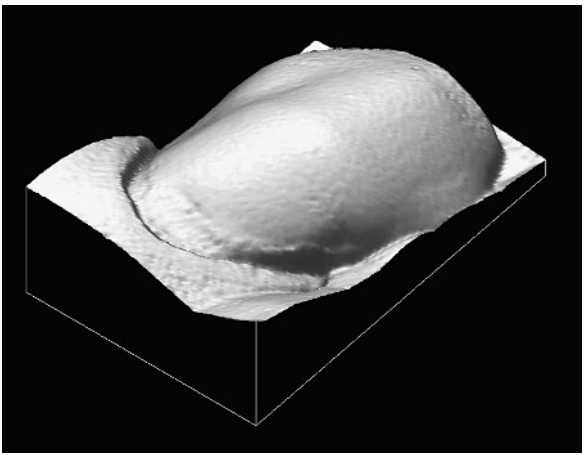

Fig. 12. The reconstructed labial tooth surface.

data and not necessarily produce a better representation of the underlying surface.

Visualisation and Measurement

When the surface has been represented as a continuous function $z = f(x, y)$, it can be manipulated in a variety of ways. Realistic rendering with Gouraud shading [62] permits the identification of anatomical features of interest and viewing of the surface in any desired orientation (fig. 12). Cross-sections may be defined by interactively marking a line segment which constitutes the x–y extent and measurements of lengths and angles may be made (fig. 13). The cross-sections need not necessarily be 'vertical' (i.e. parallel to the z axis), in which case the intersection between the plane and the surface is computed using contouring algorithms [63, 64].

Surface Representation in Other Dental Systems

Very little has been published about the surface representation methods used in other dental applications. Some systems permit digitising of a series of sequential replicas to be made on a common x–y grid by using external markers and mechanical repositioning. In studies of wear, it has been considered sufficient to evaluate the changes only at selected locations on the x–y grid and not on any intermediate points, which would require interpolation.

Cubic spline curves have been used in the development of the University of Minnesota system [65] to represent the surface as a series of parallel contours, and the intention of using bicubic splines for a 'full' surface representa-

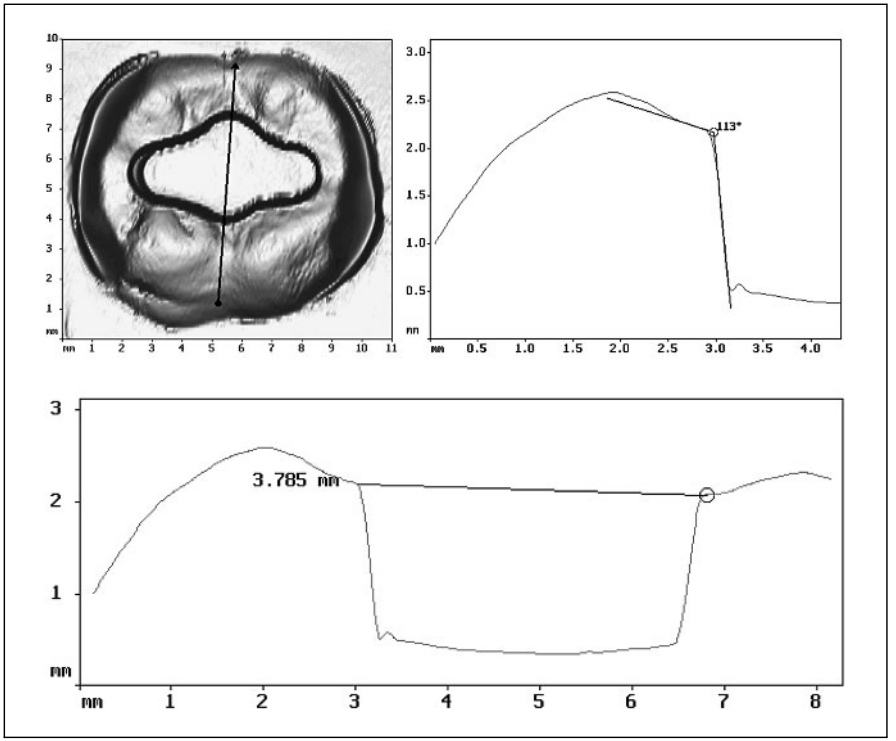

Fig. 13. Cross-sections, angles and lengths.

tion at a later stage was stated. In other systems, co-ordinate data sets are stored simply as an array of z values on an x–y grid. When a z value is required for point (x_0, y_0) not on the grid, a local interpolation is performed by fitting a polynomial through neighbouring points and evaluating it at (x_0, y_0). Most commonly, this has been bilinear interpolation (fitting a plane through three neighbouring grid points). Such is the case with the University of Munich system [30]. Mitchell and Chadwick [31] implement bilinear interpolation in the FORTRAN software, but refer to the possibility of using higher-order polynomials in the gridding of scattered data.

Although simple to implement, bilinear interpolation produces a less accurate surface model than quadratic or cubic interpolation. Furthermore, local interpolation can be inefficient since it entails solving a system of equations whenever a z co-ordinate is required. The speed of such computations is not critical in the gridding of scattered data since it is only performed once for each data set, but if z co-ordinates are to be evaluated at many locations,

for instance during measurement, visualisation or superposition of sequential data sets, the advantages of polynomial splines, whose coefficients are pre-computed, are evident.

Superposition of Sequential Data

Mathematical Background

Comparison of two sequential measurements of the same tooth replica can reveal any changes over time, such as those due to plaque formation, caries or restorations. For such a comparison to be possible, the two data sets must be in the same frame of reference in the sense that corresponding points (those which represent the same anatomical feature) should have the same numerical co-ordinates. However, the co-ordinate data from any two replicas will not generally lie in the same reference frame. The two major contributing factors are: (i) variations in orientation of the brass tube assembly when taking the replicas, and (ii) variations in the orientation of the replicas when placed on the worktable of the CMM. Attempting to reduce these variations by physically constraining the positioning of the brass tube assembly in the mouth and of the replicas on the measuring platform is impractical since it would still permit errors whose magnitude is considerably greater than the dimensions of the intended measurements. The corrections to the orientation can, instead, be made retrospectively by applying to the (x, y, z) co-ordinates of the digitised surfaces such transformations which correspond to the physical motion of rigid objects. This approach can deal with all the causes of disorientation simultaneously and with high potential accuracy.

The process of bringing two data sets into a common reference frame in this way is known as 'registration' (the term commonly used in medicine and computer vision) or 'superposition' (elsewhere). The registration problem occurs in a variety of application areas including automated inspection of manufactured components, 2-D and 3-D medical imaging (with data from sequential observations or multimodal imaging) and computer vision (to determine the 'pose' of an object and recognise it by fitting to models of known categories of objects).

A Generalised Method

Suppose that a single image is printed on two transparent sheets but not necessarily in an identical orientation (fig. 14). If a person was asked to bring the two images into superposition, their most likely course of action would be to keep one of them in a fixed position and move the other one by rotating and translating so as to either bring the corners into coincidence, or maximise

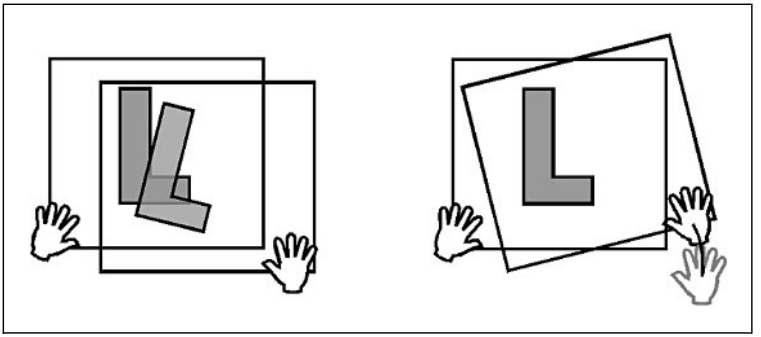

Fig. 14. Superposing two images 'by hand'.

the area of overlap. This procedure may be formulated mathematically as the following three-step 'generalised method': (1) express the position π of the second image relative to its initial position by a set of numerical parameters $u_1, u_2, ..., u_n$; (2) define a measure $E(\pi)$ of the positional disparity, and (3) determine the position of best fit by finding a vector of parameter values $\mathbf{u} = (u_1^*, u_2^*, ..., u_n^*)$ for which $E(u_1, u_2, ..., u_n)$ attains a minimum.

Unsurprisingly, these three steps form the essence of all registration methods encountered in practice. However, the specific choice of method for accomplishing each step depends on the nature of the problem as not all solutions will be equally appropriate to all situations. In the example of figure 14, two possible disparity functions were suggested: the first one expressed by the distances between the corresponding corners of the 'L' symbols, and the second one by the area of the non-overlapping regions. If the corners of the 'L' symbols were rounded, it would be difficult to identify corresponding points and the second function would be more appropriate despite its greater complexity.

Rigid Motion and its Parameterisation

Registration of two surfaces involves movement of one of the surfaces as a rigid body. In order to arrive at a parametric representation of the position of the surface relative to its initial position it is necessary to consider the class of those co-ordinate transformations which correspond to physical movement without deformation. Such transformations are termed 'rigid motion' [66, 67] and have the general form:

$$\tau(\mathbf{x}) = \mathbf{R}\mathbf{x} + \mathbf{v} \tag{2}$$

where $\mathbf{R}$ is a 'proper orthogonal' 3×3 matrix and $\mathbf{v}$ a fixed vector. A matrix $\mathbf{R}$ is proper orthogonal if $\mathbf{R}^T\mathbf{R} = \mathbf{I}$ (the identity matrix) and $|\mathbf{R}| = 1$. Such

matrices are also termed 'rotations' since they corespond to the standard geometrical notions of physical rotation around the co-ordinate origin. Stated in geometrical terms, the characterisation 2 is: every rigid motion can be represented as a rotation around the co-ordinate origin, followed by a translation.

The class of rigid motions forms a group in relation to the composition of transformations: (1) the identity transformation $\tau(\mathbf{x})=\mathbf{x}$ is a rigid motion; (2) if $\tau(\mathbf{x})=\mathbf{R}\mathbf{x}+\mathbf{c}$ is a rigid motion then its inverse $\tau^{-1}(\mathbf{x})=\mathbf{R}^{\mathrm{T}}\mathbf{x}-\mathbf{R}^{\mathrm{T}}\mathbf{c}$ is also a rigid motion, and (3) if $\tau_1(\mathbf{x})=\mathbf{R}_1\mathbf{x}+\mathbf{c}_1$ and $\tau_2(\mathbf{x})=\mathbf{R}_2\mathbf{x}+\mathbf{c}_2$ are rigid motions, then their composition $\tau_2(\tau_1(\mathbf{x}))=\mathbf{R}_2\mathbf{R}_1\mathbf{x}+(\mathbf{R}_2\mathbf{c}_1+\mathbf{c}_2)$ is also a rigid motion.

When registration is performed, the 'disparity' function E (termed the 'objective function'), can ultimately be represented as a multivariate function of the elements of the matrix $\mathbf{R}$ and of the translation vector $\mathbf{v}$. However, the nine elements of the rotation matrix are not independent, but connected by the relations arising from the orthogonality of $\mathbf{R}$ and as such they cannot serve as the independent variables of E.

In general the elements of an $n\times n$ orthogonal matrix are functions of $n(n-1)/2$ variables [68]. In three-dimensional space, rotation matrices can be represented as a product of three rotations $\mathbf{R}=\mathbf{R}_3\,\mathbf{R}_2\,\mathbf{R}_1$ of the form

$$\mathbf{R}_1=\begin{bmatrix}1 & 0 & 0\\ 0 & c_1 & -s_1\\ 0 & s_1 & c_1\end{bmatrix}\quad \mathbf{R}_2=\begin{bmatrix}c_2 & 0 & s_2\\ 0 & 1 & 0\\ -s_2 & 0 & c_2\end{bmatrix}\quad \text{and}\quad \mathbf{R}_3=\begin{bmatrix}c_3 & -s_3 & 0\\ s_3 & c_3 & 0\\ 0 & 0 & 1\end{bmatrix}$$

where $c_i^2+s_i^2=1$ $(i=1, 2, 3)$. Geometrically, $\mathbf{R}_1$ is a rotation around the x axis by an angle φ_1 where $c_1=\cos\varphi_1$ and $s_1=\sin\varphi_1$. Similarly, $\mathbf{R}_2$ and $\mathbf{R}_3$ are rotations around the y and z axes, respectively. The explicit form of the product $\mathbf{R}_3\,\mathbf{R}_2\,\mathbf{R}_1$ is

$$\mathbf{R}=\begin{bmatrix}c_2c_3 & s_1s_2c_3-c_1s_3 & c_1s_2c_3-s_1s_3\\ c_2s_3 & s_1s_2s_3+c_1c_3 & c_1s_2s_3-s_1c_3\\ -s_2 & s_1c_2 & c_1c_2\end{bmatrix}.$$

Thus, a rotation matrix can be parameterised by three rotation angles [69]. On occasions where it is necessary to convert between matrix and parameter representations, the rotation angles can be found from the relationships:

$$\tan\varphi_1=\frac{r_{32}}{r_{33}}\quad \tan\varphi_2=\frac{-r_{31}}{\sqrt{r_{32}^2+r_{33}^2}}\quad \text{and}\quad \tan\varphi_3=\frac{r_{21}}{r_{11}}.$$

In summary, every rigid motion $\mathbf{R}\mathbf{x}+\mathbf{v}$ in three-dimensional space can be represented by a six-parameter vector

$$\mathbf{p}=(\varphi_1, \varphi_2, \varphi_3, v_1, v_2, v_3)$$

where φ_1, φ_2 and φ_3 are rotation angles and v_1, v_2 and v_3 are the components of the translation vector $\mathbf{v}$.

As outlined above, determining the optimum transformation requires finding the minimum of an objective function E which contains a rotation matrix $\mathbf{R}$ whose elements r_{ij} are functions of the rotation angles. Since the methods for minimising E generally require the evaluation of its partial derivatives with respect to the rotation angles, it is useful at this stage to define $\partial\mathbf{R}/\partial\alpha$, the 'partial derivative of a rotation matrix $\mathbf{R}$ with respect to the rotation angle' α, to be the matrix whose (i, j)-th element is the partial derivative of $r_{ij}(\varphi_1, \varphi_2, \varphi_3)$ with respect to α.

Specifically,

$$\frac{\partial \mathbf{R}}{\partial \varphi_1} = \begin{bmatrix} 0 & c_1 s_2 c_3 + s_1 s_3 & -s_1 s_2 c_3 + c_1 s_3 \\ 0 & c_1 s_2 s_3 - s_1 c_3 & -s_1 s_2 s_3 - c_1 c_3 \\ 0 & c_1 c_2 & -s_1 c_2 \end{bmatrix}$$

$$\frac{\partial \mathbf{R}}{\partial \varphi_2} = \begin{bmatrix} -s_2 c_3 & s_1 c_2 c_3 & c_1 c_2 c_3 \\ -s_2 s_3 & s_1 c_2 s_3 & c_1 c_2 s_3 \\ -c_2 & -s_1 s_2 & -c_1 s_2 \end{bmatrix}$$

and

$$\frac{\partial \mathbf{R}}{\partial \varphi_3} = \begin{bmatrix} -c_2 s_3 & -s_1 s_2 s_3 - c_1 c_3 & -c_1 s_2 s_3 + s_1 c_3 \\ c_2 c_3 & s_1 s_2 c_3 - c_1 s_3 & c_1 s_2 c_3 + s_1 s_3 \\ 0 & 0 & 0 \end{bmatrix}.$$

This notation permits expressions to be written compactly and in a form similar to that for single-valued functions of one variable.

Methods Based on Identification of Corresponding Points

The position and orientation of an object in three-dimensional space are determined uniquely by the locations of at least three non-collinear points on its surface. If pairs of corresponding points are identified on two surfaces, then the surfaces can be superposed by finding the rigid motion which brings the second set of points into 'best' correspondence with the first. In practice an exact match will usually not be possible, due to errors arising from measurement or from the process of identifying the corresponding points, so the notion of 'best' is defined in a least squares sense. The formal statement of the problem is: given two sets of points $\mathrm{A} = \{\mathbf{a}_1, \mathbf{a}_2, ..., \mathbf{a}_m\}$ and $\mathrm{B} = \{\mathbf{b}_1, \mathbf{b}_2, ..., \mathbf{b}_m\}$, find a rigid motion $\tau(\mathbf{x}) = \mathbf{R}\mathbf{x} + \mathbf{v}$ which, when applied to B, minimises the sum of squares of distances

$$E(\mathbf{R}, \mathbf{v}) = \sum_{i=1}^{m} \| \mathbf{a}_i - \tau(\mathbf{b}_i) \|^2.$$

This problem has been considered by many researchers [69–73]. It can be shown that the optimal transformation has the form $\tau(\mathbf{x}) = \mathbf{R}(\mathbf{x} - \bar{\mathbf{b}}) + \bar{\mathbf{a}}$ where $\bar{\mathbf{a}}$ and $\bar{\mathbf{b}}$ are the centroids of A and B, respectively. Furthermore, if A and B are centred around the co-ordinate origin to obtain new data sets X and Y, then $\mathbf{R}$ is that rotation matrix which maximises the sum

$$\sum_{i=1}^{m} \mathbf{x}_i^{\mathrm{T}} \mathbf{R} \mathbf{y}_i.$$

The optimal rotation matrix can be found efficiently by singular value decomposition (SVD). The SVD of a real $m \times n$ ($m \geq n$) matrix $\mathbf{A}$ is its representation in the form

$$\mathbf{A} = \mathbf{U}\mathbf{S}\mathbf{V}^{\mathrm{T}}, \tag{3}$$

where $\mathbf{U}$ is an $m \times n$ column-orthogonal matrix[1], $\mathbf{V}$ is $n \times n$ orthogonal and $\mathbf{S} = \mathrm{diag}(\sigma_1, ..., \sigma_n)$ where $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_n \geq 0$. Details of the SVD can be found in Golub and Van Loan [74] and algorithms for its computational implementation in Press et al. [75].

The matrix $\mathbf{R}$ which produces the optimal fit of the two sets of points X and Y is found as follows: (1) form the $m \times 3$ matrices $\mathbf{X}$ and $\mathbf{Y}$ whose rows are the co-ordinates of the points in X and Y, respectively; (2) form the product $\mathbf{C} = \mathbf{Y}^{\mathrm{T}}\mathbf{X}$; (3) compute the SVD $\mathbf{C} = \mathbf{U}\mathbf{S}\mathbf{V}^{\mathrm{T}}$, and (4) the solution is $\mathbf{R} = \mathbf{U}\mathbf{V}^{\mathrm{T}}$.

Because of its simplicity, registration by point correspondence has been applied extensively. In relation to anatomical surfaces (including teeth), there are two approaches to the identification of corresponding points: (a) external markers can be attached to the surface under investigation [8, 24, 76, 77] or made on the surface itself [28, 78], and (b) characteristic points (landmarks) can be identified on the surface anatomy retrospectively by presenting the acquired data on a computer visualisation system [79].

Both approaches were attempted in our laboratory.

External Markers

Three ball-bearing were permanently attached to a tooth throughout the course of the investigations. The acquisition of a large number of points on the surface of these ball-bearings permitted the co-ordinates of their centres to be computed and these centres served as the three required reference points. The co-ordinates of the centres could be identified to within 8 μm in x and y, but only to within 36 μm in z [80].

[1] If $\mathbf{u}_i$ and $\mathbf{u}_j$ are two columns of $\mathbf{U}$ then $\mathbf{u}_i^{\mathrm{T}} \mathbf{u}_j = 1$ if $i = j$ and 0 otherwise.

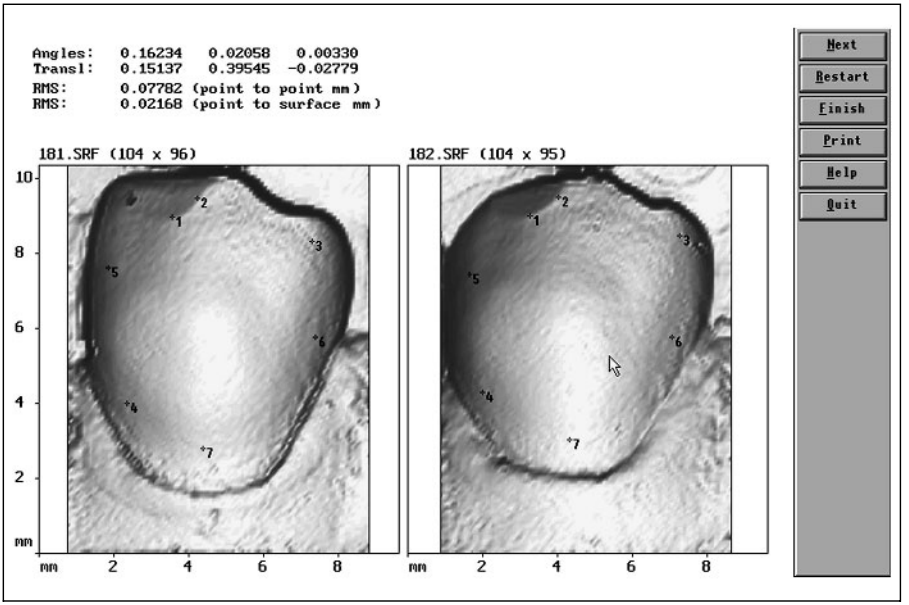

Fig. 15. Interactive identification of corresponding points.

This degree of accuracy was, however, insufficient. Such closely spaced reference points did not form a stable configuration for superposition since the errors in the computed transformation angles are inversely proportional (approximately) to the distances between the reference points. The clinical aspect was also problematic. External markers are uncomfortable for the subject and the oral environment may cause chemical changes to their surface. In a study of denture wear using a reflex microscope, Adams et al. [8] concluded that the markers 'should not be subject to any, even microscopic, changes of position, substance or definition during the monitoring period', since the registration accuracy could be seriously affected by any such changes.

Surface Landmarks

The software for surface reconstruction and visualisation developed for this project incorporates a facility to display a baseline and follow-up surface side-by-side and to mark points interactively on either of them (fig. 15). A point-to-point fit is performed and the computed transformation parameters are applied to the data points of the follow-up surface and re-interpolated to produce a bicubic spline surface superposed onto the first.

The accuracy ultimately attainable with this method was 10 μm [81], but only with great effort. The accurate identification of corresponding points was

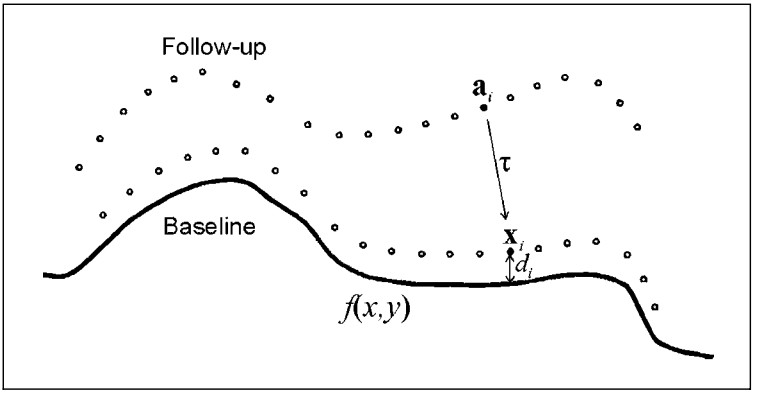

Fig. 16. Transforming a set of representative points.

limited by the resolution of the displayed image (30 μm/screen pixel) and by the relatively featureless morphology of some categories of surfaces, particularly anterior teeth. Partial compensation for both these sources of inaccuracy could be made by selecting a larger number of corresponding pairs, but the effect of one or two grossly mismatched pairs was serious [79]. However, this method is capable of quickly providing an initial approximation which can further be refined by other methods, described next.

Shape-Based Method

Methods based on the locations of landmarks, by their very nature, only use a small amount of the surface information. By utilising the additional information made available by the representation of the baseline surface as a continuous functions $f(x, y)$, it was possible to develop a method based on point-to-surface distances [82] which is robust and whose accuracy is operator-independent.

Suppose we are given a baseline surface $z = f(x, y)$ and a set of points $A = \{\mathbf{a}_1, \mathbf{a}_2, ..., \mathbf{a}_m\}$ selected from a nominally identical follow-up surface (fig. 16). We wish to fit the set of points A onto the baseline surface by applying a rigid motion $\tau(\mathbf{a}) = \mathbf{Ra} + \mathbf{v}$. Denoting by $\mathbf{x}_i = (x_i, y_i, z_i)$ the images of the points in A under the transformation τ (i.e. $\mathbf{x}_i = \mathbf{Ra}_i + \mathbf{v}$), the aim is to find such a vector of parameters $\mathbf{p} = (\varphi_1, \varphi_2, \varphi_3, v_1, v_2, v_3)$ of the transformation τ that minimises some measure of the disparity E between the transformed points X and the surface f.

Let the function E be defined as the sum of squares of distances d_i between the transformed points $\mathbf{x}_i$ and the surface f. Leaving the particular choice of

the distance function $d = d(x, y, z)$ unspecified for now, the following will hold true: (a) E is a function of the individual distances d_i: $E = E(d_i)$; (b) d_i is a function of the co-ordinates of the transformed i-th point: $d_i = d_i(x_i, y_i, z_i)$, and (c) x_i, y_i and z_i are functions of the transformation parameters: $\mathbf{x}_i = \mathbf{x}_i(\mathbf{p})$.

Therefore, if the set of points A is specified, then the value of E will depend only on the transformation parameters $\mathbf{p}$.

$$E(\mathbf{p}) = \sum_{i=1}^{m} d_i^2(\mathbf{p}) \tag{4}$$

Minimisation of the Objective Function

Expressions similar to equation 4 occur frequently when dealing with experimental data that contain normally distributed measurement errors, because a sum of squares is a maximum-likelihood estimator [75]. The problem of minimising such an expression is a 'least squares problem'.

Consider first the least squares problem where each function d_i is a *linear* function of its parameters $\mathbf{p} = (p_1, p_2, ..., p_n)$:

$$d_i(\mathbf{p}) = a_{i1}p_1 + \cdots + a_{ij}p_j + \cdots + a_{in}p_n - b_i \qquad (i = 1, ..., m).$$

The optimal vector $\mathbf{p}^*$ can be determined by the following method [74]: (1) form the matrix $\mathbf{A}$ whose (i, j)-th element is a_{ij} and the vector $\mathbf{b} = (b_1, b_2, ..., b_m)$, let $r = \text{rank}(\mathbf{A})$; (2) compute the SVD $\mathbf{A} = \mathbf{USV}^T$, and (3) let $\mathbf{u}_1, ..., \mathbf{u}_n$ and $\mathbf{v}_1, ..., \mathbf{v}_n$ denote the columns of $\mathbf{U}$ and $\mathbf{V}$, respectively. Then, the optimal vector $\mathbf{p}^*$ is given by

$$\mathbf{p}^* = \sum_{i=1}^{r} \frac{\mathbf{u}_i^T \mathbf{b}}{\sigma_i} \mathbf{v}_i.$$

The optimal vector $\mathbf{p}^*$ is also termed the solution of the overdetermined system of linear equations $\mathbf{Ap} = \mathbf{b}$ in the least squares sense because it has the property that $\| \mathbf{Ap}^* - \mathbf{b} \| \leq \| \mathbf{Ap} - \mathbf{b} \|$ for all vectors $\mathbf{p}$.

If the functions d_i are non-linear, then the minimisation of expression 4 presents a non-linear least squares problem, which may be solved by the Gauss-Newton method [83]. The iterations of the Gauss-Newton method are performed as follows.

GN1: Suppose that $\mathbf{p}_k$ is an approximation to the optimal solution $\mathbf{p}^*$. Let $\mathbf{J}$ denote the $m \times n$ matrix whose (i, j)-th element is the partial derivative

$$\frac{\partial d_i(p_1, ..., p_n)}{\partial p_j}.$$

GN2: use the SVD of $\mathbf{J}$, as outlined above, to find a vector $\mathbf{h}_k$ which is the solution of the system of linear equations $\mathbf{Jh}_k = -\mathbf{d}$ in the least squares sense, where $\mathbf{d} = (d_1, d_2, ..., d_m)$.

GN3: Set $\mathbf{p}_{k+1}=\mathbf{p}_k+\mathbf{h}_k$. The vector $\mathbf{p}_{k+1}$ is a better approximation to $\mathbf{p}^*$.
GN4: Repeat steps 1–3 until satisfactory accuracy has been attained.

The Gauss-Newton method can now be applied to the least squares problem presented by the fitting of a transformed set of sampled points to a surface. The vector of unknown parameters is $\mathbf{p}=(\varphi_1, \varphi_2, \varphi_3, v_1, v_2, v_3)$ and therefore the i-th equation of the system obtained in step GN2 will have the form:

$$\left[\frac{\partial d_i(\mathbf{p})}{\partial\varphi_1}\ \frac{\partial d_i(\mathbf{p})}{\partial\varphi_2}\ \frac{\partial d_i(\mathbf{p})}{\partial\varphi_3}\ \frac{\partial d_i(\mathbf{p})}{\partial v_1}\ \frac{\partial d_i(\mathbf{p})}{\partial v_2}\ \frac{\partial d_i(\mathbf{p})}{\partial v_3}\right]\mathbf{h}=-d_i(\mathbf{p}). \tag{5}$$

The partial derivatives of the distances d_i with respect to the transformation parameters $\mathbf{p}$ can be found by applying the chain rule (with indices omitted for notational convenience):

$$\frac{\partial d}{\partial \mathbf{p}}=\frac{\partial d}{\partial \mathbf{x}}\cdot\frac{\partial \mathbf{x}}{\partial \mathbf{p}}.$$

It is readily verified that

$$\frac{\partial d}{\partial \mathbf{x}}=\frac{\partial d}{\partial(x, y, z)}=\left[\frac{\partial d}{\partial x}\ \frac{\partial d}{\partial y}\ \frac{\partial d}{\partial z}\right]$$

and that

$$\frac{\partial \mathbf{x}}{\partial \mathbf{p}}=\left[\begin{array}{c|c|c|ccc} \dfrac{\partial \mathbf{R}}{\partial\varphi_1}\mathbf{a} & \dfrac{\partial \mathbf{R}}{\partial\varphi_2}\mathbf{a} & \dfrac{\partial \mathbf{R}}{\partial\varphi_3}\mathbf{a} & 1 & 0 & 0 \\ & & & 0 & 1 & 0 \\ & & & 0 & 0 & 1 \end{array}\right].$$

The explicit form of equation (5) is therefore:

$$\left[\frac{\partial d_i}{\partial x}\ \frac{\partial d_i}{\partial y}\ \frac{\partial d_i}{\partial z}\right]\left[\begin{array}{c|c|c|ccc} \dfrac{\partial \mathbf{R}}{\partial\varphi_1}\mathbf{a}_i & \dfrac{\partial \mathbf{R}}{\partial\varphi_2}\mathbf{a}_i & \dfrac{\partial \mathbf{R}}{\partial\varphi_3}\mathbf{a}_i & 1 & 0 & 0 \\ & & & 0 & 1 & 0 \\ & & & 0 & 0 & 1 \end{array}\right]\mathbf{h}=-d_i. \tag{6}$$

All that remains is to select an appropriate distance function $d(x, y, z)$. The 'true' distance from a point $\mathbf{x}=(x, y, z)$ to a surface $f(x, y)$ is the distance between the point $\mathbf{x}$ and that point on the surface f which is closest to $\mathbf{x}$, termed the 'foot point'. Finding the foot point is computationally expensive [84], but a first-order approximation can be used. If the surface f is considered to be represented by its tangent plane at the point $(x, y, f(x, y))$ which is the z-projection of $\mathbf{x}$ onto f, then the first-order approximation of the distance is given by:

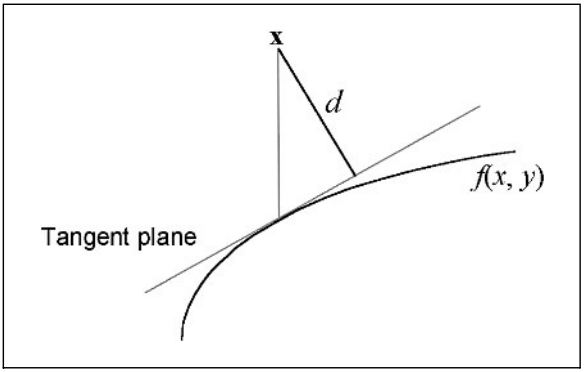

Fig. 17. A first-order approximation of the point-to-surface distance.

$$d = \frac{z - f(x, y)}{\sqrt{1 + f_x^2(x, y) + f_y^2(x, y)}}$$

where f_x and f_y denote the partial derivatives of f with respect to x and y (fig. 17).

Substituting this distance function into equation 6, the final explicit form of equation 5 in full is:

$$\mathbf{n}_i^{\mathrm{T}} \left[\frac{\partial \mathbf{R}}{\partial \varphi_1} \mathbf{a}_i \;\middle|\; \frac{\partial \mathbf{R}}{\partial \varphi_2} \mathbf{a}_i \;\middle|\; \frac{\partial \mathbf{R}}{\partial \varphi_3} \mathbf{a}_i \;\middle|\; \begin{matrix} 1 \\ 0 \\ 0 \end{matrix} \;\middle|\; \begin{matrix} 0 \\ 1 \\ 0 \end{matrix} \;\middle|\; \begin{matrix} 0 \\ 0 \\ 1 \end{matrix} \right] \mathbf{h} = \frac{f(x_i, y_i) - z_i}{\sqrt{1 + f_x^2(x_i, y_i) + f_y^2(x_i, y_i)}} \qquad (7)$$

where $(x_i, y_i, z_i) = \mathbf{R}\mathbf{a}_i + \mathbf{v}$, and $\mathbf{n}_i$ is the unit normal of the baseline surface at $(x_i, y_i, f(x_i, y_i))$, i.e.

$$\mathbf{n}_i = \frac{1}{\sqrt{1 + f_x^2(x_i, y_i) + f_y^2(x_i, y_i)}} \left[-f_x(x_i, y_i) \;\; -f_y(x_i, y_i) \;\; 1 \right]^{\mathrm{T}}.$$

Superposition of Non-Identical Surfaces

In the preceding section it was assumed that the baseline and follow-up surfaces represent the same underlying shape, except for random measurement errors. However, the ultimate purpose of registration is to enable the measurement of changes in shape, and that fact implies that in practice the two surfaces will differ to some extent. The differences may be small (wear, for instance) or large (cavity preparation) but in most cases there exists a region which remains relatively stable throughout, or is subject to considerably smaller changes than those in the region of interest. The registration of

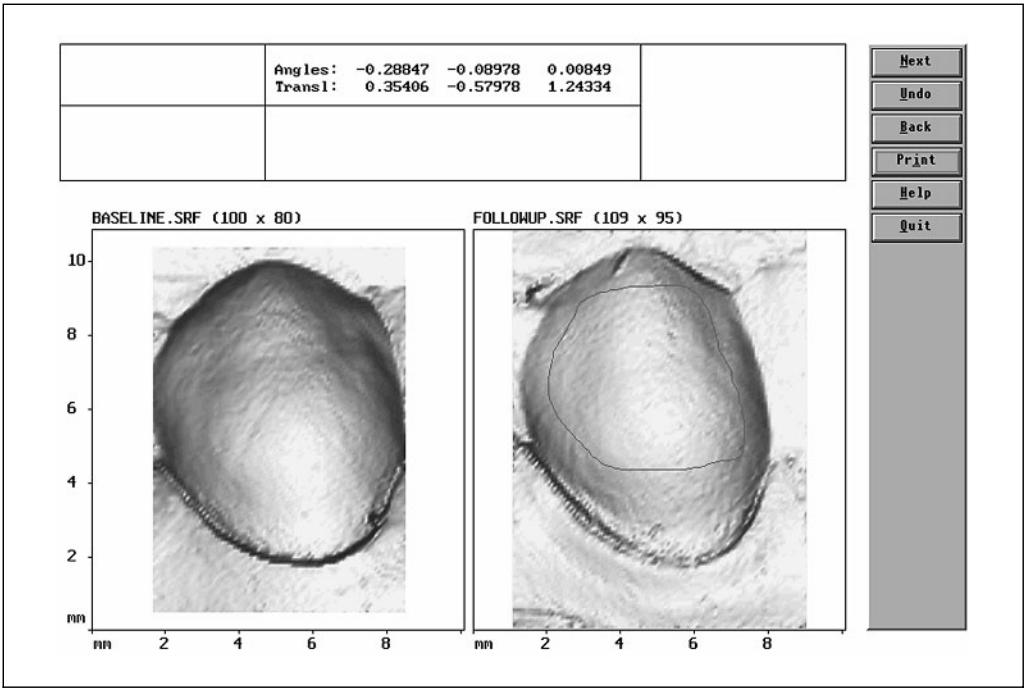

Fig. 18. The stable region outlined on the follow-up surface.

non-identical surfaces is achieved by restricting the selection of representative points to such stable regions. It is then possible to proceed as if the surfaces were identical.

Software Implementation

Registration is performed in three stages. First, an initial approximation of the transformation parameters is obtained by marking three or more pairs of corresponding points as in the landmark-based method, but with the difference that the points only need to be marked approximately because the method accuracy is not critically dependent on the choice of initial parameter values.

In the second stage, the operator marks the extent of the stable regions on the follow-up surface by outlining them with the mouse (fig. 18). The replicas in this example are from a study of the accumulation of plaque and changes in gingival contour. The crown was cleaned carefully using gauze, except for the region of interest in the proximity of the gingival margin.

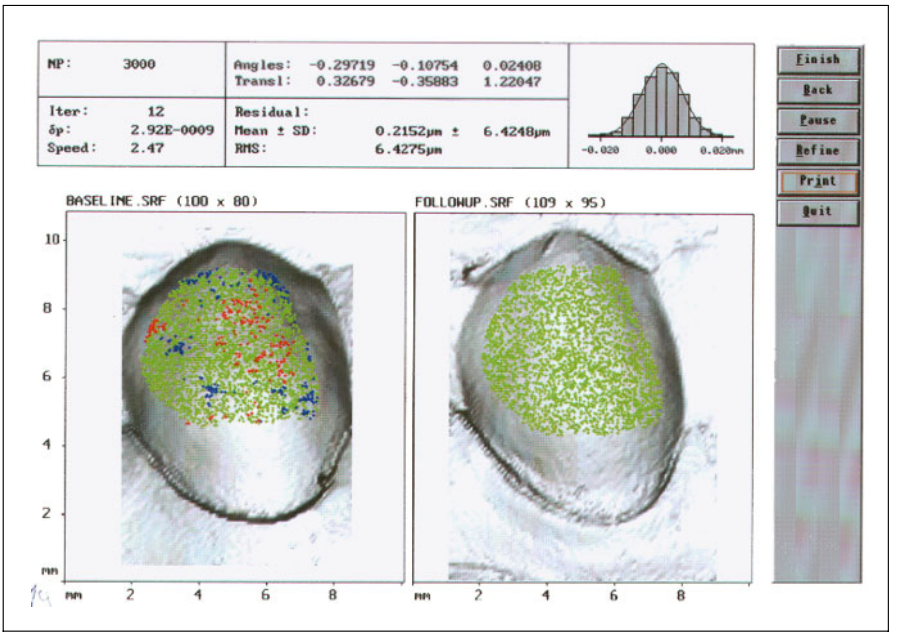

Fig. 19. The main processing phase.

On completion of the second stage, a set of 3,000 representative points which lie in the stable region of the follow-up surface is automatically generated by random uniform sampling (the number of representative points is chosen to be as high as possible, whilst maintaining an acceptable speed of operation). Following this, the main processing phase is initiated during which the representative points are fited to the baseline surface by the Gauss-Newton method. The display (fig. 19) is updated at each iteration to show: (1) the representative points, transformed by the parameter values obtained at the current iteration, overlaid onto the baseline surface: (2) a histogram of the distances (residuals) between the baseline surface and the transformed representative points of the follow-up surface; (3) statistics of the residuals (mean, standard deviation, RMS); (4) the values of the transformation parameters at the current iteration, and (5) the magnitude of their change $\delta\mathbf{p}$ from the previous iteration.

The primary indicator of a satisfactory fit is the RMS of the residuals (RMS/point). Additional verification is provided by the histogram. If the marked regions are indeed identical except for randomly distributed measurement errors, it is expected that the histogram should correspond closely to the overlaid Gaussian curve which has the same mean and standard deviation

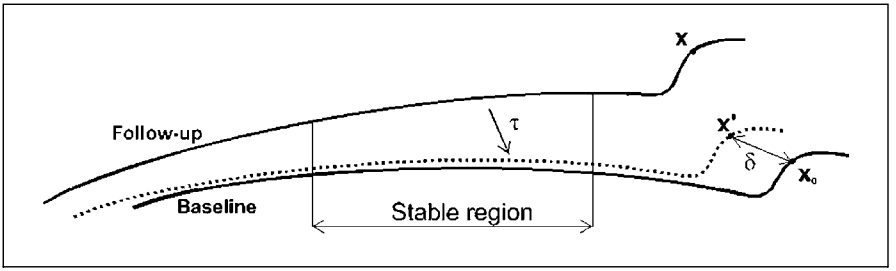

Fig. 20. Inaccurate registration, but low RMS.

as the residuals. Those transformed representative points which correspond to the worst 15% of residuals (outliers) are displayed in a different colour to the others (red for positive, blue for negative). The software incorporates a facility to exclude the outliers from further consideration and allows the iterative procedure to continue with the new 'refined' set of representative points. When the procedure has converged, the optimal transformation is applied to all the data points of the follow-up surface. From these transformed points a new, superposed follow-up surface is generated.

Results

The practical application of the shape-based method in a number of clinical studies involving many hundreds of replicas has shown that tooth surfaces can generally be superposed to within 6–9 μm RMS/point. The fit is not perfect because the registration is being performed on two surfaces which are not absolutely identical, as a consequence of the finite accuracy of replication and digitising. Furthermore, if a pair of surfaces are superposed repeatedly, the parameters of the computed optimal transformation will be subject to small variations, depending on the size and location of the marked stable regions, and also on the specific locations of the selected representative points. It was therefore of interest to quantify not only the overall goodness of fit, but also to determine which factors affect it, and to what extent.

Is some cases, the RMS/point deviation is not a complete indicator of the registration accuracy, a fact which has also been noted in other application areas [85, 86]. The example in figure 20 shows a situation where the stable region is an almost perfectly circular arc.

While the RMS/point value is low, the misregistration δ between corresponding anatomical features $\mathbf{x}_0$ and $\mathbf{x}'$ is high. From this it is clear that the evaluation of registration accuracy should be based on a measure capable of distinguishing between solutions which produce equal RMS/point values.

If τ^* is a known optimal registration and $\hat{\tau}$ is a solution obtained by some method, then the misregistration of a particular point $\mathbf{x}$ is the distance $\| \hat{\tau}^{-1}\tau^*(\mathbf{x}) - \mathbf{x} \|$, equal to the distance that a point $\mathbf{x}$ is displaced by the residual transformation $\tau_{err} = \hat{\tau}^{-1}\tau^*$. Clearly, the displacement is not constant for all points $\mathbf{x}$ due to the distance-dependent effects of rotation, but its average value can be estimated. Given a transformation τ whose parameters are $(\varphi_1, \varphi_2, \varphi_3, v_1, v_2, v_3)$ and a point $\mathbf{x}$, it can be shown that

$$\| \tau(\mathbf{x}) - \mathbf{x} \| \leq \| \mathbf{x} \| \sqrt{\varphi_1^2 + \varphi_2^2 + \varphi_3^2} + \sqrt{v_1^2 + v_2^2 + v_3^2} = \| \mathbf{x} \| \, \| \varphi \| + \| \mathbf{v} \|,$$

provided that the rotation angles are small, which is the case with τ_{err}.

The quantity $\| \mathbf{x} \| \, \| \varphi \| + \| \mathbf{v} \|$ is upper bound and scale-dependent. In our studies, with replica dimensions of approximately 10×10 mm centred around the co-ordinate origin, an appropriate 'mean' distance is 5 mm. We therefore define the mean displacement (MD) of a transformation τ as $\mathrm{MD}(\tau) = 5$ mm $\| \varphi \| + \| \mathbf{v} \|$. (Angles are dimensionless quantities, while vector norms are distances (lengths). The '5 mm' multiplier plays the role of $\| \mathbf{x} \|$ in the expression $\| \mathbf{x} \| \, \| \varphi \| + \| \mathbf{v} \|$. The MD is therefore also expressed in units of length.)

In this way, and assuming that the optimal transformation τ^* between two surfaces is known, the registration accuracy of a particular solution $\hat{\tau}$ can be defined as the MD of the residual transformation $\tau_{err} = \hat{\tau}^{-1}\tau^*$. In this context, the quantity $\mathrm{MD}(\hat{\tau}^{-1}\tau^*)$ serves as a measure of the 'distance' between the two transformations $\hat{\tau}$ and τ^*.

The repeatability of a sequence of solutions $(\tau_1, \tau_2, ..., \tau_n)$ is defined in an analogous way to the standard deviation of a random variable, as the root mean square of the 'distances' of each of the n solutions from their 'mean'. A transformation τ_{mean} is obtained by averaging the n vectors of transformation parameters, and the repeatability is defined as:

$$\text{Repeatability} = \mathrm{RMS}(\mathrm{MD}(\tau_k^{-1}\tau_{mean}), k = 1, ..., n).$$

A series of experiments to assess the accuracy and repeatability of registration were performed using data from actual digitised replicas of labial surfaces and also from surfaces generated synthetically. Each experiment was repeated 50 times with different initial transformation parameters which simulated the variations that would arise in practice from the landmark-based initial approximation phase.

Synthesised 'Typical' Case

Evaluation of registration accuracy requires that the exact solution (ground truth) be known, but this is not achievable with pairs of surfaces obtained by the standard procedure. Instead, a transformation τ was applied to a single data set P_1, obtained from a labial tooth surface, to create a new

data set P_2. The two sets of points were interpolated to create two continuous surfaces S_1 and S_2. These two surfaces were 'digitised' synthetically by sampling two sets of data points D_1 and D_2 in such a way that their x–y distribution resembled that which would be obtained with the CMM. Random errors which modelled the characteristics of the OP2 probe were added to their z co-ordinates. Finally, the data sets D_1 and D_2 were interpolated to create a pair of surfaces related by the known transformation τ. When the two surfaces were superposed in the standard way, the accuracy, repeatability and RMS/point were 9.8, 8.1 and 6.4 μm, respectively.

Effects of the Distribution of Acquired Data Points

When replicas are digitised repeatedly, the x–y distribution of each set of data points is different and the spline interpolants of those data sets therefore also differ to a small extent. The accuracy, repeatability and RMS/point of data sets obtained synthetically from a single surface, but sampled at different x–y locations, were found to be 1.25, 1.26 and 1.24 μm, respectively.

Effects of the Locations of Representative Points

Repeated registration of two 'real' data sets was performed with different initial parameters, but with the same set of representative points each time. The repeatability was found to be near-perfect (10^{-4} μm). From this, two conclusions follow.

(1) Different representative points produce different solutions. This is expected, since the 'follow-up' structure represented by these points is different each time to an extent dependent on the accuracy of replication and measurement.

(2) The same solution was obtained each time, despite the different initial parameters. Therefore, no local minima of the objective function were encountered. It is generally found in other application areas that 3-D registration of free-form surfaces is sensitive to local minima, particularly with structures whose morphology is detailed, such as human brains. In such situations the global minimum can be sought either by repeating the superposition many times and rejecting 'outlier' solutions, or by parameter space searching. However, in this application the surfaces were smooth and the measurement errors low in relation to the data point spacing and therefore the solution was not sensitive to the choice of initial transformation parameters.

Effects of the Number of Representative Points

Since the information used in superposing the follow-up surface consists only of its representative points, it was expected that selecting a larger number of points would provide a better description and therefore a more stable and

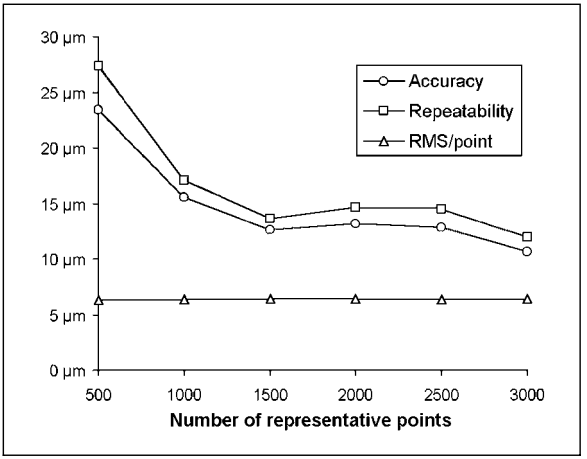

Fig. 21. Registration accuracy and repeatability vs. number of representative points.

accurate fit. Varying their number from 500 to 3,000 produced the results in figure 21. Better solutions were indeed obtained by selecting a greater number of points, but it was also apparent that an increase in their number above 1,500 yielded only slight improvements. These facts would not have been revealed by examining the RMS/point value alone, since it remained constant.

Effects of the Magnitude of Acquisition Errors

The sequence of operations employed to evaluate the 'typical' case were performed with different magnitudes of synthetic acquisition errors. The graph in figure 22 shows that a linear dependence exists between the acquisition errors and the RMS/point, but that the effect on accuracy and repeatability was proportionally greater. It is not yet clear why this should be so, but it indicates (as does the previous experiment) that the RMS/point value alone does not always provide all the information of interest.

Convergence Properties

The values of the six-dimensional vector $\mathbf{p}$ of transformation parameters at each iteration of the Gauss-Newton method form an iterative sequence $\mathbf{p}_1$, $\mathbf{p}_2$, ... which converges toward some optimal value $\mathbf{p}^*$ (within the bounds of the computing precision). The convergence properties of these sequences were of interest because this information might reveal deficiencies or indicate how the speed of operation could be improved.

A fundamental characteristic of convergent iterative sequences is the 'order of convergence', defined as the largest number r for which

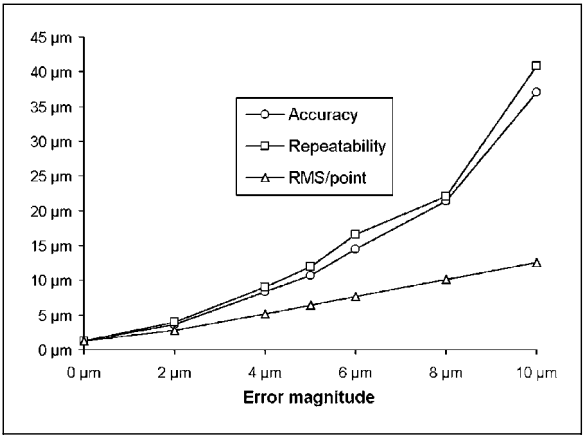

Fig. 22. Registration accuracy and repeatability vs. magnitude of acquisition errors.

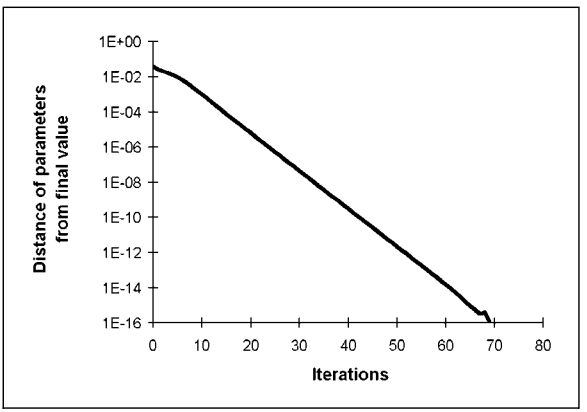

Fig. 23. The distances $\| \mathbf{p}_k - \mathbf{p}^* \|$.

$$0 \leq \lim_{k \to \infty} \frac{\| \mathbf{p}_{k+1} - \mathbf{p}^* \|}{\| \mathbf{p}_k - \mathbf{p}^* \|^r} \leq \infty.$$

The Gauss-Newton method is ultimately capable of achieving a quadratic rate of convergence, but only if the objective function tends to zero as $\mathbf{p}_k$ approaches $\mathbf{p}^*$. Otherwise its convergence is linear, and this was found to apply in the case of replica surfaces. The distances $\| \mathbf{p}_k - \mathbf{p}^* \|$ obtained during several runs of the registration procedure were recorded and plotted on a logarithmic scale against the iteration count. In each case, an almost exactly straight line was obtained (figure 23), showing that the ratio $\| \mathbf{p}_{k+1} - \mathbf{p}^* \| / \| \mathbf{p}_k - \mathbf{p}^* \|$ was almost constant.

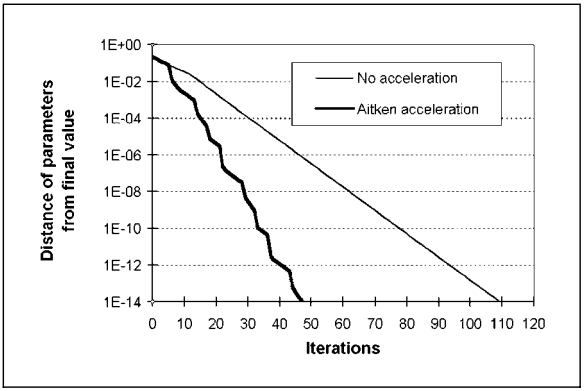

Fig. 24. Aitken acceleration of a slowly converging sequence.

The linear convergence permitted the application of a predictive technique known as Aitken acceleration to considerably reduce the number of required iterations. Consider a sequence $\{a_k\}$ which converges linearly towards a value a^* and in which each element is generated from the preceding one by some iterative procedure P. From the definition of the order of convergence it follows that for any three successive values a_{k-2}, a_{k-1} and a_k the following will hold true:

$$\frac{a_{k-1}-a^*}{a_{k-2}-a^*} \approx \frac{a_k-a^*}{a_{k-1}-a^*}. \tag{8}$$

The relationship 8 permits the approximate value of the unknown value a^* to be found as:

$$a^* \approx a_k - \frac{(a_k-a_{k-1})^2}{a_k-2a_{k-1}+a_{k-2}}. \tag{9}$$

The computed value of a^* is not exact, but if this value is used as the next element a_{k+1} in the sequence, instead of the value of a_{k+1} which would have been obtained by applying the procedure P, the convergence is greatly accelerated. Aitken acceleration has been implemented in the registration software and figure 24 shows the result of its application to a surface pair for which the convergence was particularly slow, resulting in an execution time of 30 s on a Pentium 166 processor.

Discussion

Accuracy and Repeatability

The experimental results showed that the method error was negligible when applied to identical surfaces, and that the registration accuracy depends primarily on the accuracy of the acquired data. In the absence of a known ground truth, all registration methods are ultimately statistical [87], and while it is not easy to assess the accuracy of any particular solution, repeated registration of synthetically generated data sets, which closely model the real data, has provided estimates of what may be expected in practice.

An additional factor, not considered in the experiments, relates to the size, shape and location of the stable regions from which the representative points are selected. This is a complex issue which proved very difficult to quantify rigorously since the stable regions may be either single or multiple, with varying size and patterns of spacing. This question is closely related to that of ill-posed problems such as the example in figure 20, where the RMS/point remains constant or changes very slowly relative to the changes in the transformation parameters. It is well known that some categories of surfaces cannot be superposed uniquely. For example, a cone is invariant under rotation around its axis and therefore the superposition of two cones has one rotational degree of freedom. Similar considerations apply to spheres, cylinders, planes and other forms which are self-similar, i.e. invariant under some rigid motion. When such surfaces are superposed, the matrix $\mathbf{J}$ of the associated least squares system $\mathbf{Jh} = \mathbf{d}$ is singular and permits an infinity of solutions. Such surfaces did not occur in our investigations, since the measurement errors introduce sufficient imperfections to disrupt the geometrical regularity, but it was noted that selecting very flat or cylindrical representative regions caused slower convergence and worse repeatability. Despite the potential for labial tooth surfaces to be self-similar, there are extenuating circumstances. Measurements of changes due to plaque or wear are generally made in the z direction which represents the direction of maximum constraint and therefore they are not strongly affected by registration errors in the direction of maximum freedom. Given the fact that motion-invariance is an inherent property of the surface shape, it would be desirable to develop a method which would provide an *a priori* estimate of registration accuracy and stability. Recent work in other application areas has produced some promising results [88, 89].

Identification of Stable Regions and Removal of Outliers

The registration of non-identical surfaces is based on the assumption that stable regions exist and that they can be identified. The locations and extents of such regions differ depending on the type of change being studied and

therefore some clinical knowledge of the problem is required, but there is always the possibility that the regions which have been marked as identical are, in fact, different. Possible causes are: (1) erroneous assumptions about the extent of the stable regions; (2) the presence of plaque or saliva on the tooth surface; (3) air bubbles on the replica surface, and (4) distortion of the replica caused by moving the impression tray at the critical stage while the impression material is setting.

The software provides facilities for the identification and rectification of such occurrences. If the differences between the marked regions are due only to measurement errors, then the RMS residual of the fitted points is expected to be in the range of 6–9 μm, the histogram of the residuals should exhibit a Gaussian distribution and the outlier points (both positive and negative) should be distributed uniformly.

If the outlier points are concentrated within a very small area, indicating a local fault, the operator can specify that they be automatically removed from the set of representative points and thereby excluded from further consideration. Larger regions of outlier concentration indicate more serious problems requiring further investigation which might lead either to the rejection of one or both of the replicas, or reconsideration of the possible locations of the stable regions.

Outlier removal has been used in other surface registration applications, including dental-related ones [30, 31] as a step towards automating the detection of stable regions, but there are difficulties with this approach. Namely, when the criteria for outlier detection are applied, the two surfaces are not yet correctly superposed, or there would be no need to take further action. The point-to-surface distances of the representative points will therefore differ from their final optimal values, permitting two types of classification errors to occur. Firstly, there might be failure to detect points which are not outliers according to the current relative attitude of the two surfaces, but would be if the surfaces were correctly superposed. Secondly, some 'good' points might be wrongly identified as outliers and rejected, thereby reducing the amount of available information. For these reasons, the registration software leaves the process of outlier removal under the control of the operator, to be used only where appropriate, such as in compensating for local imperfections.

Other Registration Methods

The development of surface registration methods has largely been motivated by the requirements of computer vision, industrial inspection and image-guided surgery. It can be seen from literature surveys [90, 91] that such methods essentially conform to the template of the 'generalised method', but that the specifics of each method depend on the properties of the acquired data. For

instance, computer-aided tomography produces a 3-D data set whereby an object is represented as a large number of small cubes, each of which is assigned a discrete density value. Clearly, such data and the errors associated with its acquisition are of a very different nature to the tooth surface co-ordinate data and therefore require different processing methods.

Of methods that are applicable to co-ordinate data and surfaces reconstructed from it, the best known is probably the iterative closest point (ICP) method [92]. The ICP method is based on the fitting of two sets of corresponding points, as follows: (1) select a set of representative points on the follow-up surface; (2) for each follow-up point, find the point on the baseline surface which is closest to it; (3) fit the follow-up points to the baseline points, and (4) repeat steps 2–3 until convergence.

This method is popular because it is equally applicable to 2-D curves, '$2^1/_2$-D' (single-valued) surfaces and 3-D surfaces in general. It does not require the derivatives of the surface function to be known or estimated and permits straightforward implementation, the main requirement being a function which returns the closest point and a routine for least squares fitting of two sets of points.

The ICP method was implemented in our data analysis software as an alternative, but found to be inefficient for the following reasons: (1) finding the closest point is time-consuming [84], and (2) the presence of flat horizontal regions tends to 'anchor' the surface vertically which causes very slow convergence.

The method of Chen and Medioni [93] is a close counterpart of the method used in this work. This method, too, is based on minimising the distances from points to tangent planes, but whereas our plane passes through the control point $(x, y, f(x, y))$ (fig. 17), the control point in Chen and Medioni is chosen at the intersection of the baseline surface and a line through (x, y, z) which is normal to the transformed follow-up surface. The computation of this surface-to-line intersection incurs a considerable computational expense.

A point-to-surface method frequently cited in the medical literature is the method of Pelizzari et al. [94, 95] which has been used in brain imaging. The baseline surface is represented as a series of parallel contour lines and the follow-up surface as a set of points. The minimisation is performed by a parameter search method, which is reflected in the lengthy execution time, quoted as 'several minutes' for a set of 300 points on a VAX-11/750 computer.

In dental-related applications, Bangerter et al. [27] have used the simplex algorithm (also known as Nelder and Mead's polytope method) [96] for superposing tooth surfaces, but no further particulars are given. The University of

Minnesota system [97] also employs the simplex algorithm with a reported registration accuracy of 16.9 μm RMA/point on buccal surfaces [98]. The method accuracy of the latter is quoted as 3.5 μm [99].

The simplex algorithm is a parameter search method. As such, it does not require the derivatives of the objective function to be known or approximated, but its convergence is very slow. It was considered in the early stages of this work, but rejected after tests which showed that the superposition of two identical curves required over 100 iterations.

Dastane et al. [41] have performed a comparative study of three registration methods using commercially available software. The best achievable result in terms of RMS/point was 9.49 μm on surfaces defined by points on an 81×81 grid with a spacing of 175 μm and 500 representative points. The surfaces were nominally identical, one being obtained from the other by a known rigid transformation. Under comparable conditions, our shape-based method achieved 0.013 μm. This difference perhaps best illustrates the shortcomings of applying 'off-the-shelf' solutions intended for general purpose use or for other specific application areas.

The registration method used in the University of Munich system [30] bears much resemblance to ours, having evidently been developed following a similar line of reasoning. An initial approximation is obtained by identifying three pairs of corresponding points, stable regions are defined by multiple polygons and the disparity function is defined as the sum of squares of point-to-surface distances. Minimisation is by gradient descent, conjugate gradient descent and the Levenberg-Marquardt method [83]. The RMS/point on registered occlusal surfaces was 10.1 μm, mainly due to measurment errors, whereas the intrinsic registration error is described as 'considerably less than 1 μm'.

Least squares minimisation of point-to-surface distances has also been applied to tooth surfaces by Mitchell and Chadwick [31]. Although the system is described as experimental, the reported particulars of its development are especially welcome since they show that the key issues identified by the authors as relevant to the accurate detection of changes in morphology are included within those that have been addressed in this work.

Clinical Applications

The validity of the practical and theoretical methods described above has been confirmed by their successful application in objective clinical research. This section presents examples of some of the types of our investigations which have been enhanced by co-ordinate data analysis.

Wear of Restorations

Quantification of the wear of restored occlusal anterior tooth surfaces has been a dominant area in the study of dental morphology by analysis of co-ordinate data using both CMMs and other measurement systems. Wear over long periods can be readily detectable and since this area has been extensively researched, it did not figure prominently in our investigations. It was of interest, due to the timescales involved, to determine if useful data could be obtained from 'historical' replicas taken for purposes unrelated to co-ordinate acquisition. Figure 25 shows a plane cross-section through a sequence of three replicas of a molar tooth followed longitudinally over a 10-year period taken at 5-year intervals. The restoration covered an area of 80% or more of the occlusal surface of the tooth and extended to one of the approximal surfaces and onto more than half the buccal surface. In the top image, the wear can particularly be seen adjacent to the cavo-surface margin of the restoration, which was more steeply inclined near the area where the restoration meets the tooth surface. The middle image shows the good superposition of the surfaces buccally and lingually, where relatively little wear had occurred, whilst the cusp tips, being related to occlusal contact areas, had worn relatively more. The centre half of the cross-sections shows the greater relative wear of the composite resin restorative material (P-10, 3M Dental Products, St. Paul, Minn., USA). In the lower image the wear at one point on the composite resin material has been quantified and shown to be 345 μm.

Plaque

Plaque is frequently present in sufficient quantity to enable its visual detection from the rendered images (fig. 26). Its distribution can be viewed either on individual cross-sections (fig. 27) or over the entire surface with colour-coded 'subtraction maps' (fig. 28) of the type commonly used in imaging applications [27, 100]. The distribution of plaque is frequently very irregular and therefore measurements of its thickness at individual locations are susceptible to considerable variations. For this reason, plaque thickness is evaluated by permitting the operator of the software to define a rectangular region of interest (usually 1 mm^2) over which the thickness is automatically evaluated at several thousand locations and averaged.

The presence of a stable region required for superposing is ensured by removing the plaque (with gauze) from a region located around the mid-line on the upper coronal half of the tooth prior to replication. Assessment of the methodology is reported by Yeganeh et al. [101, 102]. The method has been applied to compare the plaque removal ability of two toothpastes [103] and proved capable of revealing a statistically significant difference in plaque depth change between the two products (mean = 19 μm, SE = 6.53, $p < 0.01$).

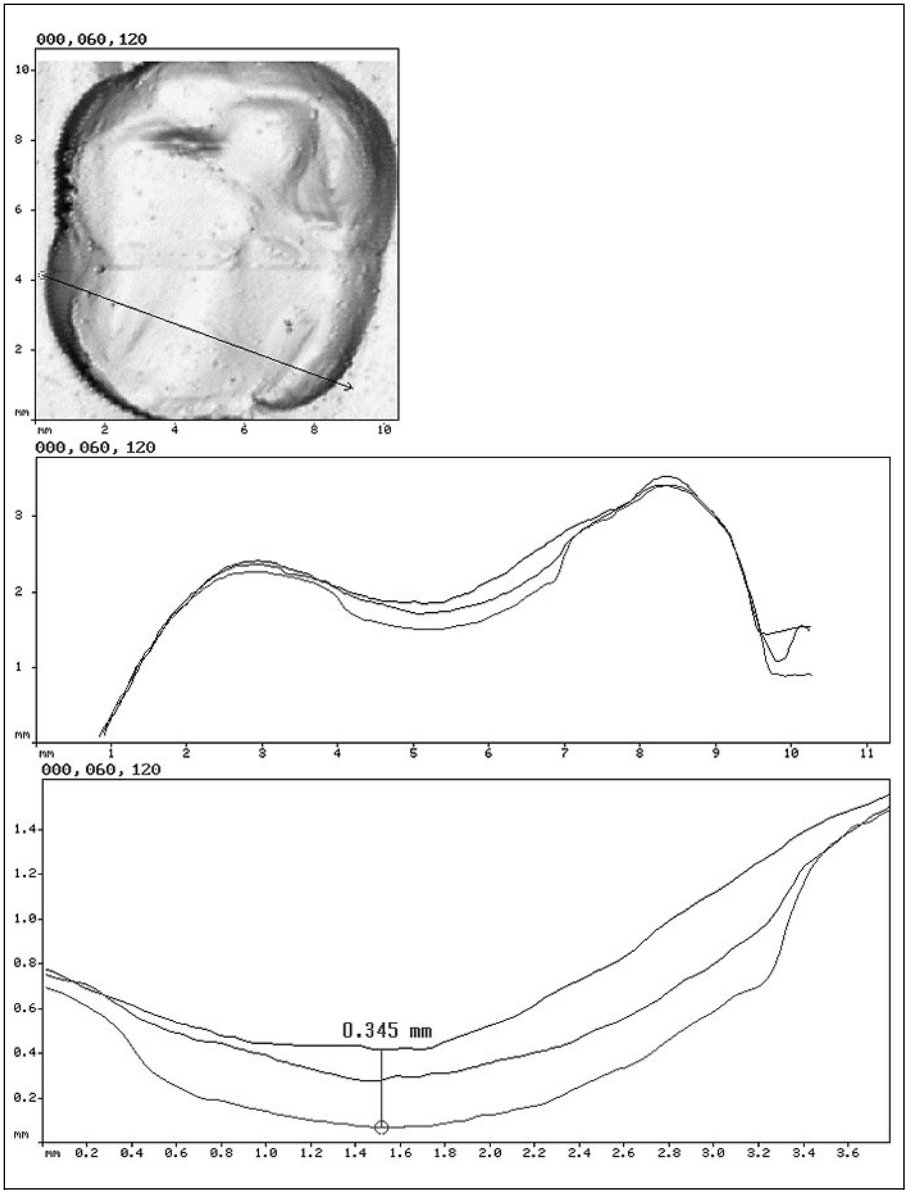

Fig. 25. Wear of a restoration at 5 and 10 years.

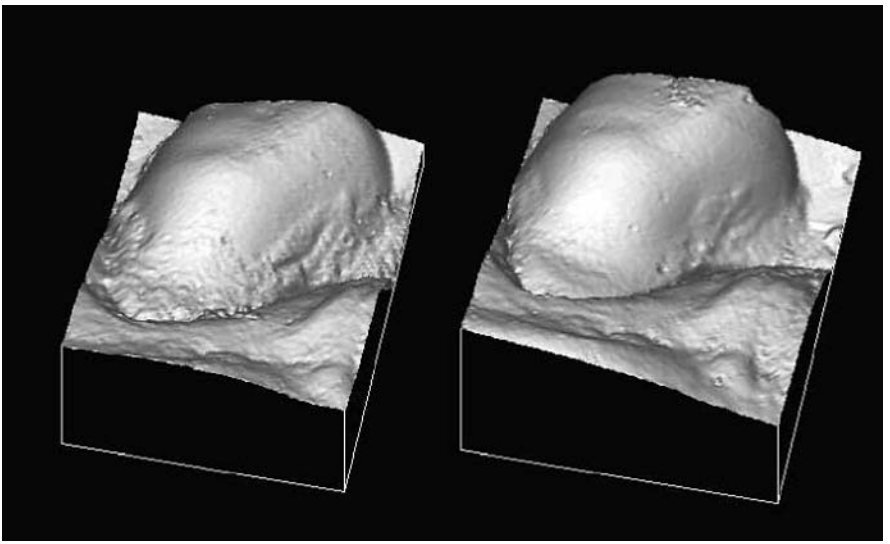

Fig. 26. Plaque before and after cleaning.

Gingivitis

Co-ordinate metrology is particularly well suited to the study of changes in gingival contour, since these are large in relation to the method accuracy and the entire tooth surface can be used as the stable region for superposing. Applications to date include a study of experimentally induced gingivitis over a period of 21 days [104] and canine gingivitis [105].

Effects of Two Toothbrushes on Hard Tissue

The system proved capable of discriminating between the effects of different toothbrushes on hard tissue loss during a 3-month study period, while clinical assessment during the same period did not produce statistically significant observations [106].

Three-Month Assessment of an Antimicrobial Root Sealant

Co-ordinate metrology has previously been used to assess the volume and area of sealant applied to occlusal surfaces [100, 107]. Our system was applied in an investigation of the thickness and wear of a protective root sealant on exposed root surfaces [108]. Impressions were taken before and after the placement of the sealant to assess its thickness (fig. 29) and also after 3 months to assess wear and failure (fig. 30).

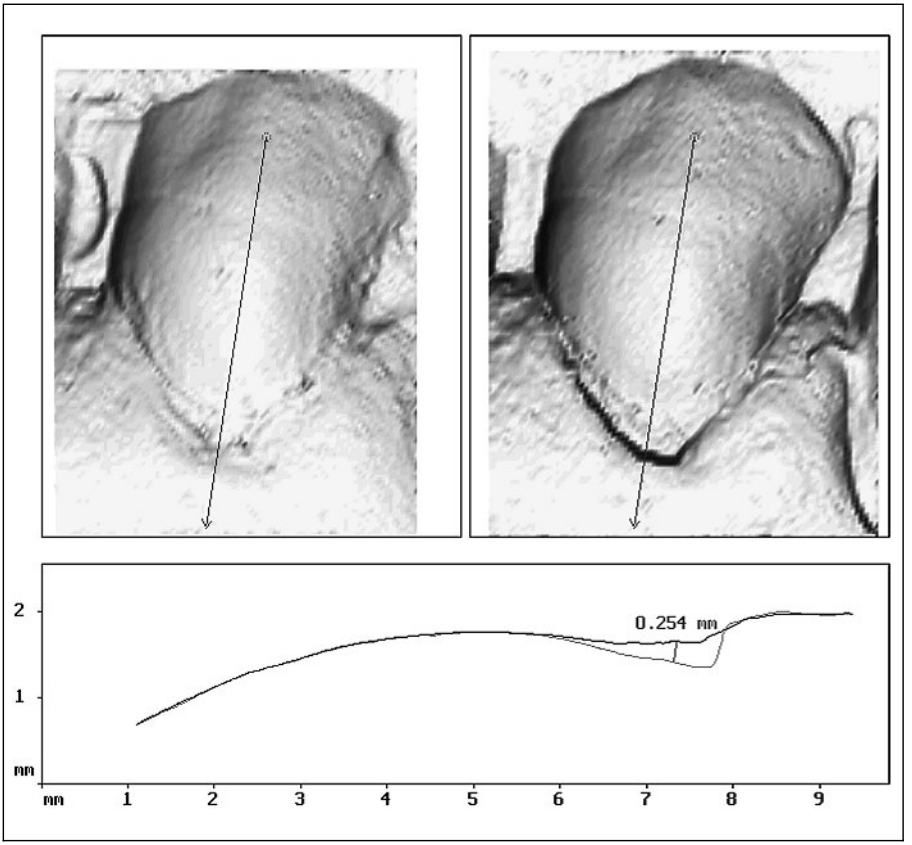

Fig. 27. Plaque cross-section.

Geometry of Crowns and Crown Preparations

The variations in geometry of labial shoulder preparations from a variety of operators and the effects that these variations have on the emergence profile of metal ceramic crowns have been assessed and reported by Seymour et al. [109, 110]. An example is shown in figure 31.

Other Application Areas – Treatment of Inflamed Joints

Replicas of a single inflamed proximal interphalangeal joint injected with a steroid to reduce swelling were made at baseline prior to injection and 2 and 7 days after the steroid injection. The scanning pitch was 300 μm (fig. 32). The RMS/point of stable regions was between 90 and 130 μm – remarkably accurate considering that any bending of the joint would alter the shape

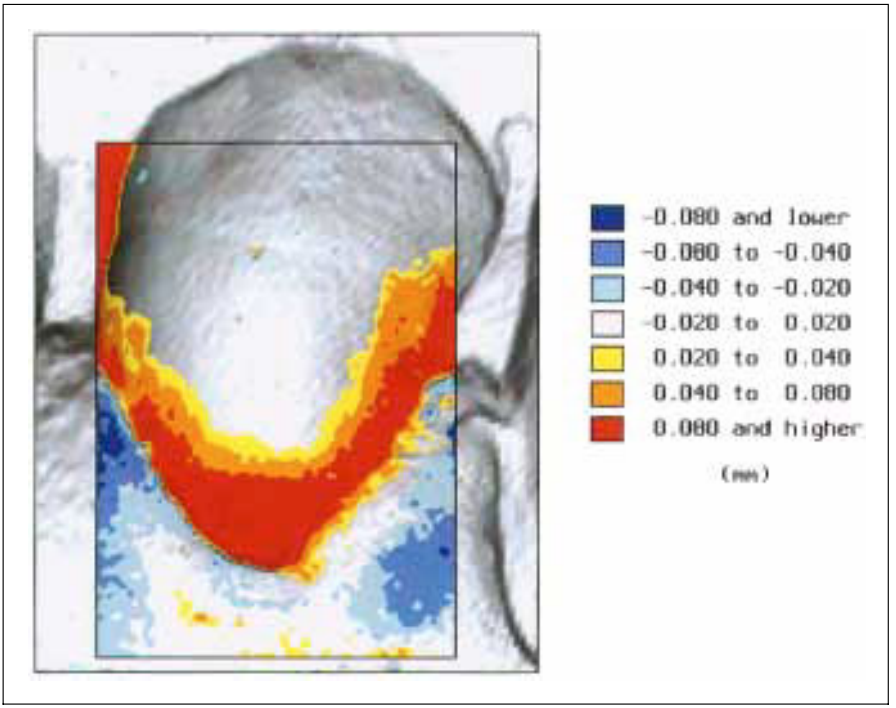

Fig. 28. Plaque–subtraction map.

considerably. The measured improvement (fig. 33) was not detected through clinical observation by the rheumatologist, nor by the patient.

Conclusions

Computer-based methods have in the past been described as 'extremely expensive, subject to superpositioning errors and impractical for evaluating casts in large clinical studies' [22]. With this new system, patient and clinician time consumption is minimal and involves the use of familiar techniques. Since replicas can be obtained easily, speedily and economically, and forwarded to a central laboratory for digitising, the cost implications, even when large clinical studies are involved, are not of the significance that has previously been the case. The replication method is highly repeatable for hard tissue and plaque, less so for soft tisue, but sufficient to permit the reliable detection of changes.

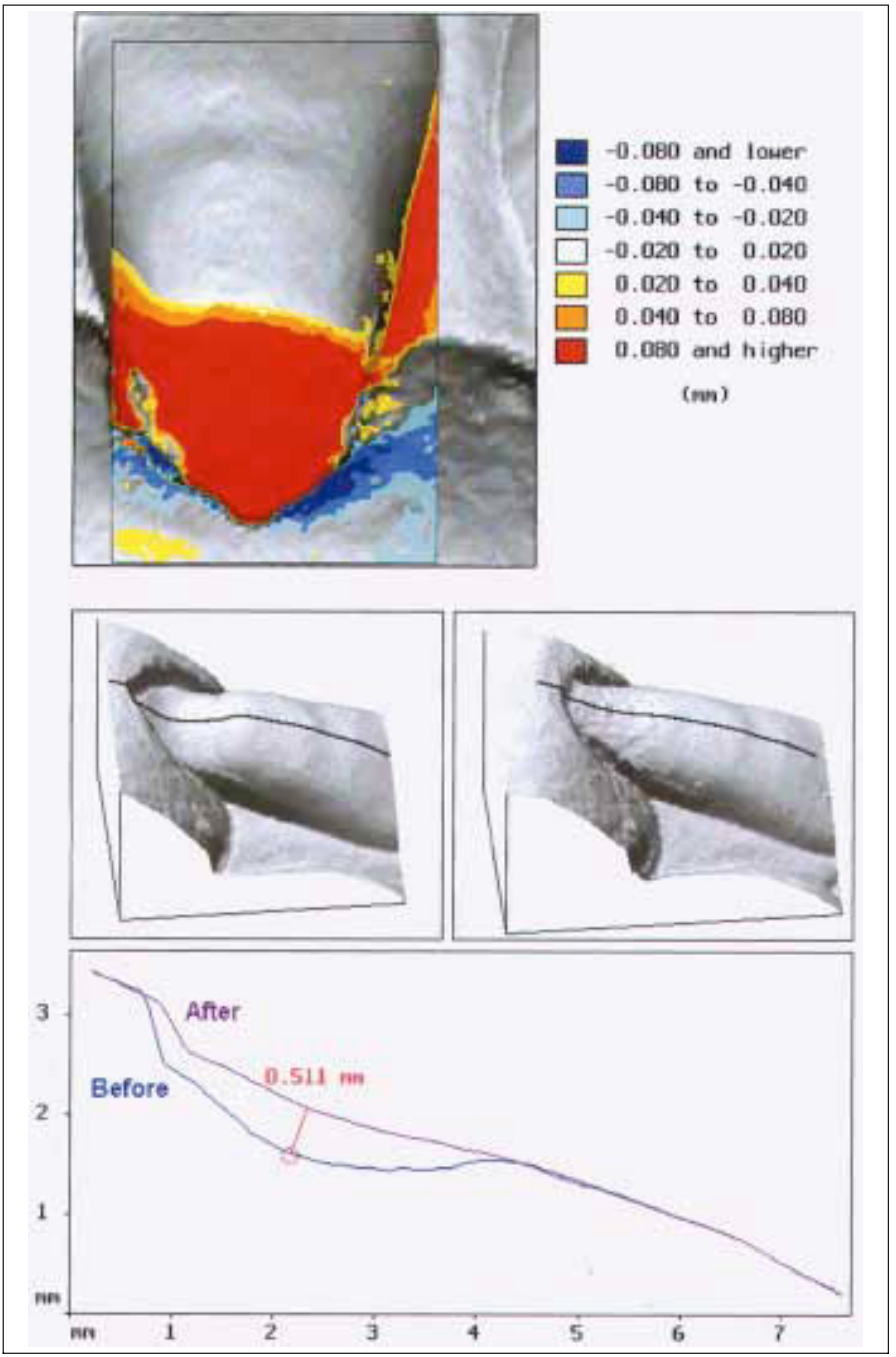

Fig. 29. Extent and depth of applied sealant.

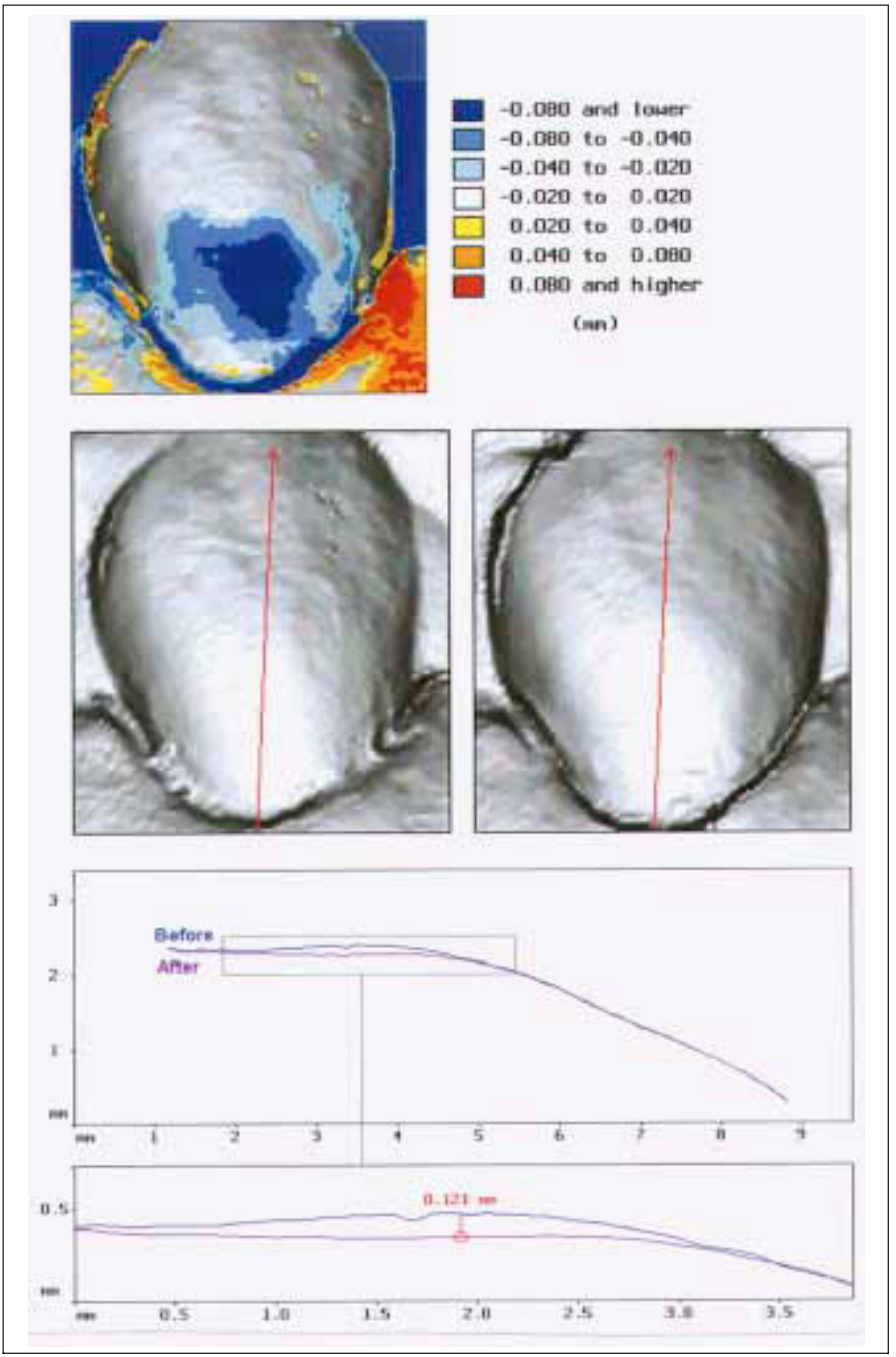

Fig. 30. Sealant wear after 3 months.

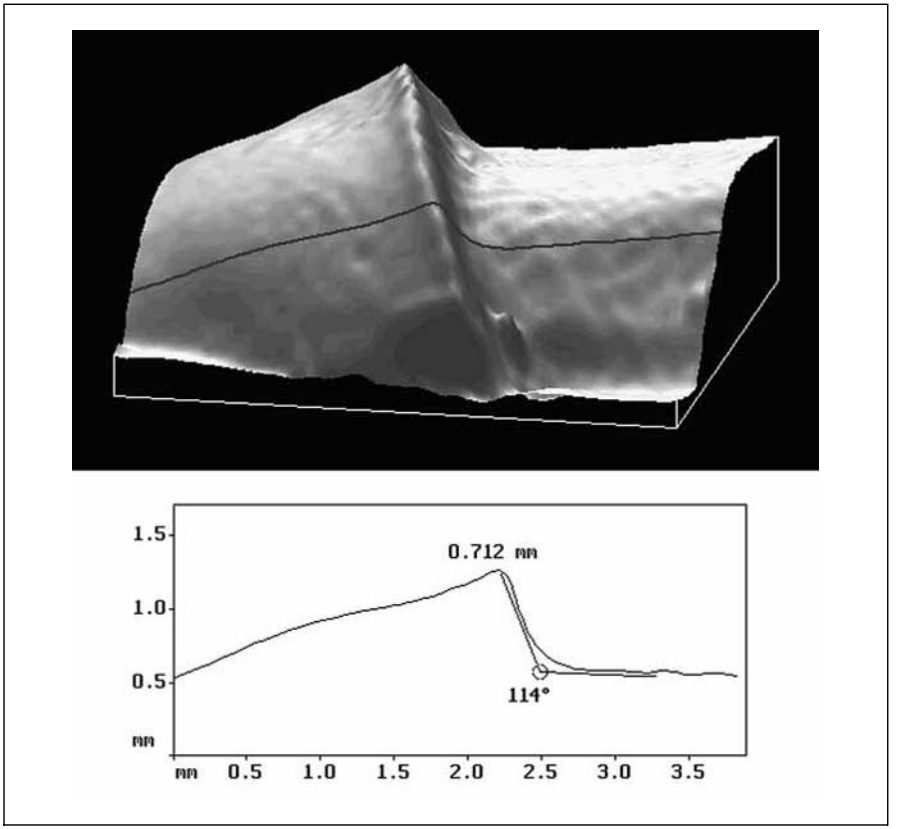

Fig. 31. Geometry of crown preparations.

The majority of co-ordinate data acquisition systems applied in dental research have employed probes which make physical contact with the surface of the specimen. Such probes are not ideal for measuring silicone impressions because they exert pressure which may cause deformation of the elastic impression material and therefore require rigid models to be produced. Additionally, post-processing of the data is necessary to compensate for the shape of the stylus, since the exact point with which it contacts the surface is gradient-dependent [111]. The ability to acquire co-ordinate data directly from standard silicone impressions eliminates the time, effort and potential for additional errors associated with the production of rigid models. The system has been applied over a sufficiently lengthy period to permit a more detailed assessment of its characteristics than has been the case with many other acquisition systems, thus enabling its optimal exploitation.

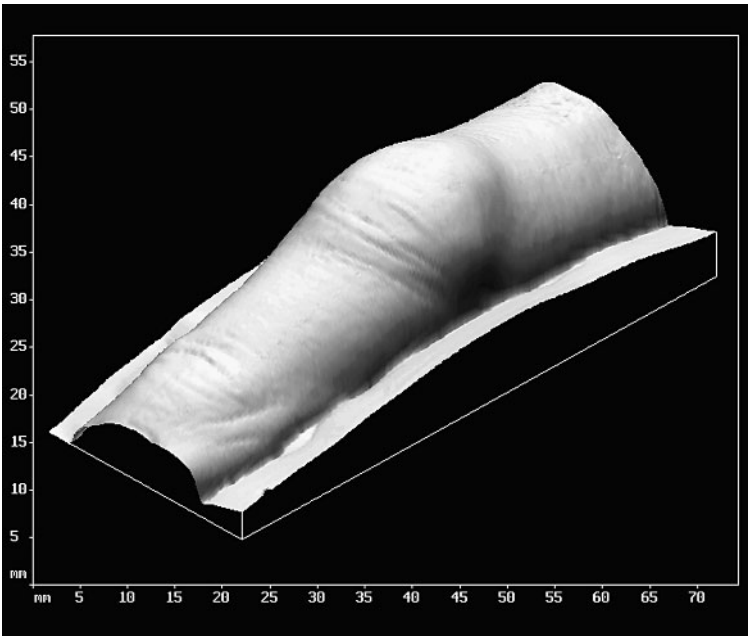

Fig. 32. The surface obtained from a replica of an inflamed joint.

The reference-free superposition of data sets obtained from sequential replicas permits investigations of changes in morphology to be performed which would otherwise be impractical if external markers were required. It can also be applied to replicas and models produced earlier, in studies which have not anticipated the use of co-ordinate metrology.

The registration method itself is based on the well-known principle of least squares optimisation. The convergence has been augmented by Aitken acceleration.

Extensive assessment of the method's accuracy and repeatability on synthetically constructed surfaces and on clinically acquired data has been performed. The registration accuracy of clinically acquired data was found to be 6–9 μm in terms of RMS/point. Operation of the supporting software is straightforward and efficient, requiring little operator training. Its visualisation capabilities permit the creation of images which can be easily perceived and understood by clinicians.

The study and application of co-ordinate data analysis over an extended period has enabled detailed, quantifiable 3-D measurements to be recorded in longitudinal studies on gingivae, plaque and root surfaces of teeth, which

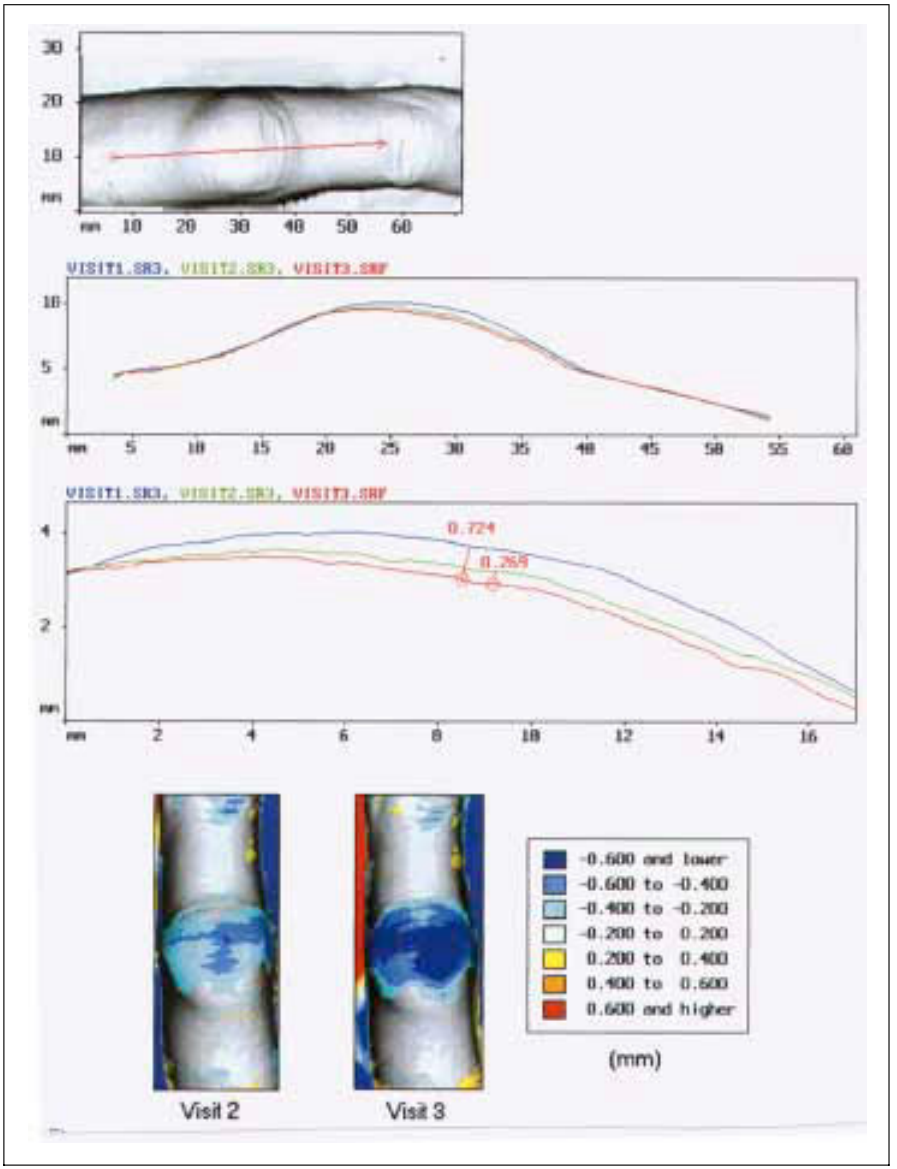

Fig. 33. Progressive improvement relative to baseline.

have been rarely investigated in the past. There is good correlation between the results obtained using the methods described and those determined by clinical examination alone, e.g. plaque indices vs. the measured plaque or changes in gingival contour, though the discriminating power of co-ordinate data analysis is far better. Areas of further clinical research in which the system might be applied are not limited to dentistry, but include a wide category of medical and other investigations of changes in shape.

Acknowledgements

We are grateful to: Prof. Harry Allred, founder of the Dental Metrology Unit; Prof. Maurice Cox of the National Physical Laboratory, for providing mathematical expertise; Mrs Lifong Zou, for digitising all replicas used in these studies; Ms Sharzad Yeganeh, Ms Aylin Baysan, Dr Kevin Seymour and Dr Carolyn Morris-Clapp, clinical investigators, and Prof. David Blake of the University of Bath for collaboration on the study of inflamed joints.

The authors would like to acknowledge the support of clinical studies by: Optiva Corporation, USA; The Scientific Foundation for Dental Diagnosis, Netherlands; Unilever Research, UK; Procter and Gamble, USA; Synthélabo, France and Dentsply, Germany.

References

1 Cvar JF, Ryge G: Criteria for the clinical evaluation of dental restorative materials. USPHS Publication No 790-244. San Francisco, US Government Printing Office, 1971.
2 Leinfelder KF, Taylor DF, Barkmeier WW, Goldberg AJ: Quantitative wear measurement of posterior composite resins. Dent Mater 1986;2:198–201.
3 Elderton RJ: An objective method for measuring the surface morphology of cavities and restorations in vivo. J Oral Rehabil 1977;4:323–334.
4 Chadwick RG, McCabe JF, Walls AW, Mitchell HL, Storer R: Comparison of a novel photogrammetric technique and modified USPHS criteria to monitor the wear of restorations. J Dent 1991;19: 39–45.
5 Chadwick RG: Close range photogrammetry – A clinical dental research tool. J Dent 1992;20: 235–239.
6 Scott PJ: Pepper's ghost observed in Cape Town. S Afr J Photogram Remote Sensing Cartog 1984; 14:89–95.
7 Adams LP, Wilding RJC: Tooth wear measurements using a reflex microscope. J Oral Rehabil 1988; 15:605–613.
8 Adams LP, Jooste CH, Thomas CJ: An indirect in vivo method for quantification of wear in denture teeth. Dent Mater 1989;5:31–34.
9 Atkinson JT, Groves D, Lalor MJ, Cunningham J, Williams DF: The measurement of wear in dental restorations using laser dual-source contouring. Wear 1982;76:91–104.
10 Williams DF, Cunningham J, Lalor MJ, Groves D, Atkinson JT: Laser techniques for the evaluation of wear in class II restorations. J Oral Rehabil 1983;10:407–414.
11 Farely DW, Shin D, Jones R, Agrawal CM: The wear of human enamel opposing dental porcelain (abstract). J Dent Res 1999;78:112.
12 Boyde A, Howell PGT, Franc F: Simple SEM stereophotogrammetric method for three-dimensional evaluation of features on flat substrates. J Microsc 1986;143:257–264.

13 Howell PGT, Reid SA: A microcomputer-based system for rapid on-line stereological analysis in the SEM. Scanning 1986;8:139–144.
14 Howell PGT, Boyde A, Ross HF: A three axis stereocomparator for scanning electron microscopic photogrammetry – RS3. Scanning 1986;8:182–186.
15 Anderson P, Davis GR, Elliott JC: Microtomography. Microsc Anal 1994;15:31–33.
16 Nielsen RB, Alyassin AM, Peters DD, Carnes DL, Lancaster J: Microcomputed tomography: An advanced system for detailed endodontic research. J Endod 1995;21:561–568.
17 Balto K, Müller R, Carrington DC, Stashenko P: Quantitation of periapical bone destruction in mice by micro-computed tomography (abstract). J Dent Res 1999;78:122.
18 Watson TF, Boyde A: Confocal light microscopic techniques for examining dental operative procedures and dental materials. Am J Dent 1991;4:193–200.
19 Mörmann WH, Bindl A: The new creativity in ceramic restorations: Dental CAD-CIM. Quintessence Int 1996;27:821–828.
20 Alcañiz M, Chinesta F, Monserrat C, Grau V, Ramón A: An advanced system for the simulation and planning of orthodontic treatments; in Höhne KH, Kikinis R (eds): Visualization in Biomedical Computing. 4th Int Conf. Berlin, Springer, 1996, pp 511–520.
21 Chadwick RG: A review: The assessment of the durability of composite resin restorative materials in vivo. Clin Mater 1989;4:241–253.
22 Bayne SC, Taylor DF, Rekow ED, Wilder AD, Heymann HO: Confirmation of Leinfelder clinical wear standards. Dent Mater 1994;10:11–18.
23 Roulet JF, Reich T, Lutz F: High precision occlusal mapping – A new method for measuring wear of posterior composite resins (abstract). J Dent Res 1983;62:220.
24 Roulet JF: Development of appropriate measuring devices; in Roulet J-F (ed): Degradation of Dental Polymers. Basel, Karger, 1987, pp 101–108.
25 Lambrechts P, Vanherle G, Vuylsteke M, Davidson CL: Quantitative evaluation of the wear resistance of posterior dental restorations. J Dent 1984;12:252–267.
26 DeLong R, Pintado MR, Douglas WH: Measurement of change in surface contour by computer graphics. Dent Mater 1985;1:27–30.
27 Bangerter V, Christensen R, Christensen G: Method for determination of in vivo wear (abstract). J Dent Res 1987;66:126.
28 McDowell GC, Bloem TJ, Lang BR, Asgar K: In vivo wear. Part I: The Michigan computer-graphic measuring system. J Prosthet Dent 1988;60:112–120.
29 Chadwick RG, Mitchell HL, Cameron I, Hunter B, Tulley M: Development of a novel system for assessing tooth and restoration wear. J Dent 1997;25:41–47.
30 Mehl A, Gloger W, Kunzelmann KH, Hickel R: A new optical 3-D device for the detection of wear. J Dent Res 1997;76:1799–1807.
31 Mitchell HL, Chadwick RG: Mathematical shape matching as a tool in tooth wear assessment – Development and conduct. J Oral Rehabil 1998;15:921–928.
32 Jovanovski V, Tay WM, Zou L, Anderson IJ, Cox MG, Forbes AB, Allred H, Morganstein SI, Lynch E: Objective assessment of three-dimensional structures in clinical dentistry using methods of co-ordinate metrology. Nanobiology 1996;4:55–61.
33 Starcke EN Jr: A historical review of complete denture impression materials. J Am Dent Assoc 1975;91:1037–1041.
34 Doubleday B: Impression materials. Br J Orthod 1998;25:133–140.
35 Chee WW, Donovan TE: Polyvinyl siloxane impression materials: A review of properties and techniques. J Prosthet Dent 1992;68:728–732.
36 Fano V, Gennari PU, Ortalli I: Dimensional stability of silicone-based impression materials. Dent Mater 1992;8:105–109.
37 Clancy JM, Scandrett FR, Ettinger RL: Long-term dimensional stability of three current elastomers. J Oral Rehabil 1983;10:325–333.
38 Williams PT, Jackson DG, Bergman W: An evaluation of the time-dependent dimensional stability of eleven elastomeric impression materials. J Prosthet Dent 1984;52:120–125.
39 Tjan AH, Whang SB, Tjan AH, Sarkissian R: Clinically oriented evaluation of the accuracy of commonly used impression materials. J Prosthet Dent 1986;56:4–8.

40 Corso M, Abanomy A, Di Canzio J, Zurakowski D, Morgano SM: The effect of temperature changes on the dimensional stability of polyvinyl siloxane and polyether impression materials. J Prosthet Dent 1998;79:626–631.
41 Dastane A, Vaidyanthan TK, Vaidyanthan J, Mehra R, Hesby R: Development and evaluation of a new 3-D digitization and computer graphic system to study the anatomic tissue and restoration surfaces. J Oral Rehabil 1996;23:25–34.
42 Zou L, Yang QP, Jovanovski V, Tay WM, Lynch E: Measurement of tooth morphology using a laser probe fitted on a co-ordinate measuring machine; in De Chiffre L (ed): Proc 4th Int IMEKO Symp on Laser Metrology for Precision Measurement and Inspection in Industry. Lyngby, Center for Geometrical Metrology, Institute of Manufacturing Engineering, Technical University of Denmark, 1996, pp 11–20.
43 Zou L, Jovanovski V, Tay WM, Lynch E: Error distribution on a digitised datum sphere (abstract). J Dent Res 1996;75:1165.
44 Zou L, Jovanovski V, Tay WM, Lynch E: The influence of three parameters on laser probe scanning of irregular surfaces (abstract). J Dent Res 1995;74:895.
45 Faux ID, Pratt MJ: Computational Geometry for Design and Manufacture. Chichester, Ellis Horwood, 1979.
46 Catley D, Whittle C, Thornton P: Applications of the general surface definition and manipulation system GENSURF; in Martin RR (ed): The Mathematics of Surfaces II. Oxford, Clarendon Press, 1987, pp 171–202.
47 Sarkar B, Menq C-H: Smooth-surface approximation and reverse engineering. Computer Aided Design 1991;23:623–628.
48 Weir DJ, Milroy MJ, Bradley C, Vickers GW: Reverse engineering physical models employing wrap-around B-spline surfaces and quadrics. Proc Inst Mech Eng [B] 1996;210:147–157.
49 McLain DH: Drawing contours from arbitrary data points. Computer J 1976;17:318–324.
50 MacCarthy BL, Handscomb DC: The approximation of hydrographic survey data using tensor-product B-spline surfaces; in Handscomb DC (ed): The Mathematics of Surfaces III. Oxford, Clarendon Press, 1989, pp 373–389.
51 Ateshian GA: A B-spline least-squares surface-fitting method for articular surfaces or diarthroidal joints. J Biomech Eng 1993;115:366–373.
52 Cox MG: Practical spline approximation; in Turner PR (ed): Lecture Notes in Mathematics 965: Topics in Numerical Analysis. Berlin, Springer, 1982, pp 79–112.
53 Cox MG: Algorithms for Spline Curves and Surfaces. NPL Report DITC 166/90. Teddington, National Physical Laboratory, 1990.
54 de Boor C: A Practical Guide to Splines. New York, Springer, 1978.
55 Cox MG: The numerical evaluation of B-splines. J Inst Math Appl 1972;10:134–149.
56 de Boor C: On calculating with B-splines. J Approx Theory 1972;6:50–62.
57 Sabin MA: Contouring – The state of the art; in Earnshaw RA (ed): Fundamental Algorithms for Computer Graphics. NATO ASI Series, vol F17. Berlin, Springer, 1985, pp 411–482.
58 Anderson IJ, Cox MG, Mason JC: Tensor-product spline interpolation to data on or near a family of lines. Numer Alg 1993;5:193–204.
59 Anderson IJ, Mason JC: Surface fitting by exploiting data distribution; in Mullineux G (ed): Mathematics of Surfaces VI. Oxford, Clarendon Press, 1996, pp 137–152.
60 Anderson IJ: A piecewise approach to piecewise approximation. Numer Alg 1997;15:132–152.
61 Anderson IJ, Pink CK: Spline surface fitting of rank deficient problems: A natural solution; in Goodman T, Martin R (eds): The Mathematics of Surfaces VII. Winchester, Information Geometers, 1997, pp 439–461.
62 Gouraud H: Continuous shading of curved surfaces. IEEE Trans Comp [C] 1971;20:623–628.
63 Heap BR, Pink MG: Three contouring algorithms. NPL Mathematical Report 81. Teddington, National Physical Laboratory, 1969.
64 Schumaker LL: Fitting surfaces to scattered data; in Lorentz GG, Chui CK, Schumaker LL (eds): Approximation Theory II. New York, Academic Press, 1976, pp 203–268.
65 DeLong R, Bramwell MC: Matching an initial surface to a worn surface subject to rotation and translation; in Martin RR (ed): The Mathematics of Surfaces II. Oxford, Clarendon Press, 1987, pp 137–149.

66 Mirsky L: An Introduction to Linear Algebra. Oxford, Clarendon Press, 1995.
67 Schreier O, Sperner E: Modern Algebra and Matrix Theory. New York, Chelsea Press, 1951.
68 Schwerdtfeger H: Introduction to Linear Algebra and the Theory of Matrices. Groningen, Noordhoff, 1950.
69 Hanson RJ, Norris MJ: Analysis of measurements based on the singular value decomposition. SIAM J Sci Stat Comp 1981;2:363–373.
70 Arun KS, Huang TS, Blostein SD: Least-squares fitting of two 3-D point sets. IEEE Trans Pattern Anal Machine Intell 1987;9:698–700.
71 Kanatani K: Analysis of 3-D rotation fitting. IEEE Trans Pattern Anal Machine Intell 1994;16: 543–549.
72 Horn BKP: Closed-form solution of absolute orientation using unit quaternions. J Opt Soc Am 1987;4:629–642.
73 Benjemaa R, Schmitt F: A solution for the registration of multiple 3D point sets using unit quaternions; in Burkhardt H, Neumann B (eds): Computer Vision – ECCV ′98. Berlin, Springer, 1998, vol 2, pp 241–254.
74 Golub GH, Van Loan CF: Matrix Computations, ed 2. London, Johns Hopkins University Press, 1989.
75 Press WH, Teukolsky SA, Vetterling WT, Flannery BP: Numerical Recipes in C, ed 2. Cambridge, Cambridge University Press, 1992.
76 Bartlett DW, Smith BGN: Measurment of tooth wear in patients with palatal erosion. Br Dent J 1997;182:179–184.
77 Arai K, Fujita M, Ishikawa H: Recording of casts in ICP with an optical 3D measuring system (abstract). J Dent Res 1999;78:162.
78 Bloem TJ, McDowell GC, Lang BR, Powers JM: In vivo wear. Part II: Wear and abrasion of composite restorative materials. J Prosthet Dent 1988;60:242-249.
79 Toennies KD, Udupa JK, Herman GT, Wornom IL, Buchman SR: Registration of 3D objects and surfaces. IEEE Comput Graphics Appl 1990;10:52–62.
80 Bedford J, Zou L, Jovanovski V, Lynch E: Reproducibility of data collected by a co-ordinate measuring machine (abstract). J Dent Res 1993;72:742.
81 Jovanovski V, Zou L, Tay WM, Lynch E, Cox MG: Superposition of co-ordinate data obtained from a sequence of replicas (abstract). J Dent Res 1995;74:895.
82 Jovanovski V, Zou L, Tay WM, Lynch E, Cox MG: Assessment of three-dimensional structures in clinical dentistry; in Ciarlini P, Cox MG, Pavese F, Richter D (eds): Advanced Mathematical Tools in Metrology III. Singapore, World Scientific, 1997, pp 263–265.
83 Gill PE, Murray W, Wright MH: Practical Optimization. London, Academic Press, 1981.
84 Anderson IJ, Turner DA, Mason JC: An efficient algorithm for solving the surface foot point problem; in Cripps R (ed): The Mathematics of Surfaces VIII. Winchester, Information Geometers, 1998, pp 89–103.
85 Grimson WEL, Ettinger GJ, White SJ, Lozano-Pérez T, Wells WM, Kikinis R: An automatic registration method for frameless stereotaxy, image guided surgery and enhanced reality visualization. IEEE Trans Med Imaging 1996;15:129–140.
86 Ellis RE, Fleet DJ, Bryant JT, Rudan J, Fenton P: A method for evaluating CT-based surgical registration; in Troccaz J, Grimson E, Mösges R (eds): Proc CVRMED-MRCAS ′97. Berlin, Springer, 1997, pp 141–150.
87 Pennec X, Thirion J-P: Validation of 3-D Registration Methods based on Points and Frames. INRIA Report 2470. Sophia-Antipolis, INRIA, 1995.
88 Simon DA, Kanade T: Geometric constraint analysis and synthesis: Methods for improving shape-based registration accuracy; in Troccaz J, Grimson E, Mösges R (eds): Proc CVRMED-MRCAS ′97. Berlin, Springer, 1997,pp 181–190.
89 Stoddart AJ, Lemke S, Hilton A, Renn T: Estimating pose uncertainty for surface registration. Image Vision Comput 1998;16:111–120.
90 van den Elsen P, Pol E-JD, Viergever MA: Medical image matching – A review with classification. IEEE Engineering Med Biol Mag 1993;12:26–39.
91 Lavalée S: Registration for computer-integrated surgery: Methodology, state of the art; in Taylor RH (ed): Computer-Integrated Surgery. Cambridge, MIT Press, 1995, pp 77–97.

92 Besl PJ, McKay ND: A method for registration of 3-D shapes. IEEE Trans Pattern Anal Machine Intell 1992;14:239–256.
93 Chen Y, Medioni G: Object modeling by registration of multiple range images. Image Vision Computing 1992;10:145–155.
94 Pelizzari CA, Chen GTY, Spelbring DR, Weichselbaum RR, Chen C-T: Accurate three-dimensional registration of CT, PET and/or MR images of the brain. J Comput Assist Tomogr 1989;13:20–26.
95 Pelizzari CA, Tan KK, Levin DN, Chen GTY, Balter J: Interactive 3D patient-image registration; in Colchester ACF, Hawkes DJ (eds): Information Processing in Medical Imaging, 12th International Conference. Berlin, Springer, 1991, pp 132–141.
96 Walsh GR: Methods of Optimisation. London, Wiley, 1975.
97 DeLong R, Douglas WH: An artificial oral environment for testing dental materials. IEEE Trans Biomed Eng 1991;38:339–345.
98 DeLong R, Pintado MR, Heinzen M, Ko CC, Douglas WH: A comparison of contact and optical methods for measuring abfraction wear (abstract). J Dent Res 1999;78:161.
99 DeLong R, Pintado MR, Douglas WH, Taylor J, Breuer MM: Quantification of clinical changes in the gingival crest height using computer graphics (abstract). J Dent Res 1993;72:330.
100 Pintado MR, Conry JP, Douglas WH: Measurement of sealant volume in vivo using image-processing technology. Quintessence Int 1988;19:613–617.
101 Yeganeh S, Zou L, Jovanovski V, Heath MR, Lynch E: Objective quantification of plaque thickness in vivo (abstract). J Dent Res 1996;75:84.
102 Yeganeh S, Lynch E, Jovanovski V, Zou L: Quantification of root surface plaque using a new 3-D laser scanning method. J Clin Periodontol 1999;26:692–697.
103 Waterfield PC, Shearer B, Lynch E, Jovanovski V, Morris-Clapp C: Metrology: A novel method to measure plaque removal by dentifrices (abstract). J Dent Res 1999;78:498.
104 Yeganeh S, Lynch E, Shearer B, Heath MR, Jovanovski V, Zou L: Clinical and metrology assessment of experimentally induced gingivitis (abstract). J Dent Res 1999;78:351.
105 Snider A, Buchanan W, Yeh C, Donovan Brand R, Mahony L, Jovanovski V, Zou L, Lynch E: Canine gingivitis assessment by computer aided 3 dimensional ccordinate metrology (abstract). J Dent Res 1999;78:351.
106 Savill G, Grigor J, Huntingdon E, Jackson R, Lynch E: Toothbrush design: Adapting for the future. Int Dent J 1998;48:519–525.
107 Conry JP, Pintado MR, Douglas WH: Measurement of fissure sealant surface area by computer. Quintessence Int 1990;21:27–33.
108 Baysan A, Lynch E, Jovanovski V, Brailsford S, Zou L, Beighton D: Three month assessment of an antimicrobial root sealant (abstract). J Dent Res 1999;78:444.
109 Seymour K, Zou L, Samarawickrama DYD, Lynch E: Assessment of shoulder dimensions and angles of porcelain bonded to metal crown preparations. J Prosthet Dent 1996;75:406–411.
110 Seymour KG: Variations in the Labial 'Shoulder' Geometry of Teeth Prepared to Receive Metal Ceramic Crowns; PhD thesis, University of London, 1998.
111 Hewlett ER, Orro ME, Clark GT: Accuracy testing of three-dimensional digitizing systems. Dent Mater 1992;8:49–53.

Dr. Vladimir Jovanovski, Department of Adult Oral Health,
St. Bartholomew's and the Royal London School of Medicine and Dentistry,
Turner Street, London E1 2AD (UK)
Tel. +44 20 7377 7000/ext. 2118, Fax +44 20 7377 7375, E-Mail V.Jovanovski@mds.qmw.ac.uk