

Reconstruction of the surface of the human body from 3D scanner data using B-splines

Ioannis Douros^{*}, Laura Dekker, Bernard F. Buxton

Dept. of Computer Science, University College London, London WC1E 6BT, UK

ABSTRACT

There are an increasing number of applications that require the construction of computerised human body models. Recently, Hamamatsu Photonics have developed an accurate and fast scanner based on position-sensitive photon detectors, capable of providing in a few seconds a dense (256x400 points) representation of the body. The work presented here is designed to exploit such a scanner's capabilities. An algorithm is introduced that deals with the surface-from-curves problem and can be combined with an existing curves-from-points algorithm to solve the surface-from-points problem. The algorithm takes as input a set of B-spline curves and uses them to drive a fast and robust surface generation process. This is done by adequately sampling the curves, in a manner that incorporates explicit assumptions about human body geometry and topology. The result is a compound, multi-segment, and yet entirely smooth surface, that may be used to calculate body volume and surface area.

Keywords: Human body, automated surface reconstruction, B-splines, interpolation, segmentation, noise removal

1. INTRODUCTION

Recent advances in three-dimensional scanning technology have enabled the generation of high-density point data sets that describe the surfaces of real objects, including animate objects such as the human body. This means that it is now possible to produce computer-based models that describe in detail the topology and the geometry of an actual human body. Such models are of increasing importance in a number of application areas, in particular in the clothing industry and in medical research where they may be used to perform fast and accurate automatic measurements, for example, to monitor body growth, or to design clothes customised to an individual's body shape. However, such measurements cannot be made on raw point data. A technique must be devised in order to produce a proper surface (skin), given the data. This requires an algorithm that can analyse the data and from it infer the topology (and subsequently the geometry) of the body.

The Department of Computer Science at University College London has been loaned a Body Lines Scanner manufactured by Hamamatsu Photonics¹, capable of collecting whole body scans in a few (~12) seconds. The scanner's main mechanism consists of 8 cameras, each one measuring reflections from an array of 32 infra-red LED's. The scanner thus captures a 256-point horizontal cross-section of the body every 5 millimetres, as the cameras slide down the 2 metres of the scanner's working height. Data quantity is thus the order of 200 Kbytes/image and its range precision is approximately 12 bits, sufficient to cater for the range accuracy of approximately 2mm in the 1m working width of the machine.

A number of algorithms are presented that deal with the surface-from-curves problem. Previous work on the extraction of curves-from-points has been exploited, so that their merging comprises a solution to the surface-from-points problem. To handle the surface-from-curves problem, it is assumed that the procedure proposed takes as input B-spline curves that have been fitted to the initial point data within each horizontal cross-section and uses these curves to drive a fast and robust surface generation process. This is done by adequately resampling the curves, then interpolating a NURBS surface through the samples. An example of such a technique is the one devised by Jones et al.² at Loughborough University.

2. BACKGROUND

The 'skinning problem' (surface from points) is not a new one. Previous attempts fall into two categories: those that require

^{*} Correspondence: Email: i.douros@cs.ucl.ac.uk; WWW: <http://www.cs.ucl.ac.uk/staff/i.douros>; Telephone: ++44 171 419 3695; Fax: ++44 171 387 1397

explicit information about the structure of the data points (connectivity), and those that automatically or intelligently infer the structure of the point set from the points themselves. The methods in the first category are the most generic and can be applied to any object, but they do not necessarily give the best approximation to the actual object, and are also time consuming, since providing connectivity information is often computationally very demanding and, in some cases requires manual assistance. On the other hand, methods in the second category are less generic (usually each one of them applies to a specific type of object), but they usually give better results because they can utilise all implicit and assumed information that is 'hidden' in the data.

2.1. Previous work on generic skinning algorithms

The scheme devised by Hughes Hoppe is a generic method for fitting a smooth surface comprised of B-spline patches on unstructured point data of arbitrary topology. It is extremely accurate if given noiseless data without gaps, and also has the advantage that it requires absolutely no prior knowledge of data topology in order to generate the surface³. However, the point set given by the scanner is not noiseless, and it contains gaps due to occlusion. Hoppes algorithm thus cannot be used because it tends to infer incorrect topologies and join areas of the body that should not join, e.g. hands to thighs.

Stoddart and Hilton have developed of a generic surface reconstruction technique, based on the notion of 'deformable surfaces'⁴. Their method, which is commonly known as 'Slime', generates G^1 continuous deformable surfaces using a generalisation of the concept of biquadratic B-splines. The surfaces produced can describe any surface topology and local discontinuities (step edges, roof edges or cusps), The method has low computational complexity and can adaptively generate control points of variable density in order to describe surfaces that are very detailed at some places and very smooth and bland at others. However, the 'Slime' method works on voxel data, which means that an initial set of point data has to be converted to voxel data before being processed. This might initially sound good, as it also combines well with the idea of using a marching cubes algorithm⁵ on the voxels in order to calculate enclosed volume, but in practice this is hard to do. Conversion from point to voxel form is not a straightforward task, especially with body scanner data, where noisy and missing data can affect conversion and subsequently ruin the fit. In particular missing data, e.g. at the armpits and inner thighs, is a problem that is very difficult to overcome using Slime surfaces at their current state of development.

2.2. The Least Squares B-spline Curve fitting algorithm

The work presented here is based on the techniques developed by Dekker and West^{6,7}. Their work focused on tackling problems specific to fitting surfaces to bodies given initial data from the Hamamatsu Body Lines scanner, including: (i) noisy data (owing to reflections), (ii) missing data (owing to occlusion), (iii) outliers (points not belonging to body surface), and (iv) inconsistent data ('ripples' at areas imaged by more than one camera). Without this previous work, the method presented here could not be applied to the body scanner data. They take an unstructured point cloud and from it produce a body segmentation and structured, well defined data. The techniques developed consist of two stages, data preprocessing and a least-squares fitting stage. During the first stage, data is prepared in order to acquire slices of valid data points, so that a curve can be fitted around each one of them. This stage includes cleaning and segmentation. The steps of this stage are the following:

Data cleaning: Each scan produces a total of 102,500 points, many of which are noisy and inaccurate, and therefore have to be removed from the dataset. There are six types of such points, each one handled in a different way. These are: (a) low intensity false points, removed by intensity thresholding, (b) reflections from the scanner structure (internal walls), removed by spatial thresholding since their locations are known, (c) points from the first/last LED of each camera that are ignored because they fall on the body almost tangentially, (d) high intensity outliers, removed by using a nearest neighbour filter, (e) multiple reflections, removed by checking whether a point lies in the same quadrant as the camera that detected it, and (f) inconsequential slices above the head, removed using a thresholding on the number of valid points per slice that have remained after all other cleaning procedures.

Body growing and segmenting: A number of techniques are applied to detect which points belong to which part of the body. This is done using a growing procedure that starts from the slice located at 15% of the body height (guaranteed to be legs) and keeps detecting leg points upwards until the groin is reached. Then, it continues growing up the torso until underarms are detected, and body growing continues downwards to the arms and upwards to shoulders and head, until all valid points are marked as leg points, arm points or torso points. Local criteria of proximity and topology are applied to determine groin and underarm location.

Data association: In each slice of data there are now a number of remaining valid points. Before a curve can be fitted to

them, their order must be determined. For simple slices, such as through an individual arm or leg, this is done by calculating the centroid and then ordering the points by angle. For complex slices, a neighbourhood scheme is followed starting from a seed point.

The second fitting stage consists of two steps. First, a knot sequence is generated, and for each data point p_i a parameter value w_i is determined. The objective is then to minimise the distance between p_i and the point $S(w_i)$ on the curve. Error minimisation is carried out by iteratively moving the control points (a least-squares problem solved using Cholesky decomposition) and then the parameter values are updated (adjusted) so as to produce a smaller error value. The process is iteratively repeated until the total error drops below a certain threshold or until the number of iterations exceeds a maximum limit.

The result is a set of cubic B-splines that are closed around horizontal slices of the body parts. For each curve, all control points are coplanar and the distance between successive slices is 5mm (the ‘sampling frequency’ of the scanner). No other assumption can be made about the curves. The starting point may differ from slice to slice and can be either on the front or the rear or elsewhere on each slice. In addition, each curve may be directed clockwise or counterclockwise. This lack of explicit knowledge about each curve is dealt with by an appropriate sampling strategy.

3. DESIGN AND METHODOLOGY

In this section, an overview of the surfaces-from-curves algorithm is presented, together with the concepts that have driven the design of the method. The main stages of the algorithm are presented with particular reference to skinning the torso, legs and feet.

3.1 Overview of the proposed algorithm

The main steps of the algorithm are as follows:

- Fit horizontal least-squares B-spline curves along the data slices
- Detect key slices (feet, groin, underarms, top of head)
- Resample the groin slice at uniform intervals
- Re-order and calibrate the groin-slice sample sequence
- Produce torso samples by growing upwards from the groin slice by using a chord bisector scheme
- Produce leg samples by growing downwards from the groin slice by using a chord bisector OR bounding box scheme
- Divide the sample set into secondary subsegments
- Fit a B-spline surface on each subsegment
- Smoothly join the B-spline surfaces

A detailed explanation of each step is presented in the sections that follow. In the above, we note that *slices* are groups of co-planar data points that lie on horizontal cross sections of the body, and within which the curve fitting is carried out. *Key slices* are slices that contain key landmark points of the body, such as the groin and armpits, that are important for the purposes of surface reconstruction. For example, the *groin slice* is the set of points that lie at the same horizontal plane as the branching point of the groin. Finally, a *sample sequence* is a sequence of points acquired by evaluating the parametric expression of a B-spline that has been fitted around a slice, for non-decreasing values of the spline parameter, a *segment* (or *primary segment*) is any group of data or sample points, and a *subsegment* (or *secondary segment*) is any subset of a segment.

3.2. Assumptions

The surface reconstruction technique proposed is not generic. For it to work, a number of assumptions are made. First, it is assumed that the object whose surface is to be reconstructed is a human body that is intact, belongs to an adult and that there are no serious deformities. It is further assumed that the subject is standing upright with feet on the scanner floor or on a couple of rectangular, dark coloured blocks of equal and known height, length and width, used to reduce contamination of the body data by reflections from the scanner floor. Additionally, it is assumed that the height of the subject plus the height of the standing blocks (if any) does not exceed 2 metres and that the subject is wearing light coloured underwear that is neither too baggy, nor too tight, so that the actual form of the body is subject to minimal distortion, ideally of less than one




Figure 1: The primary segmentation of the body

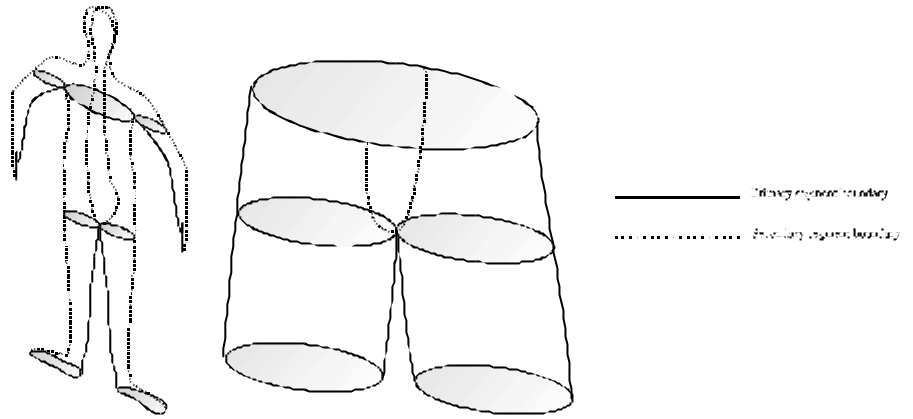


Figure 2: Primary and secondary segments of a human body (left) and an object with a single simple branching (right)

or two millimetres. The subject is also assumed to be wearing a light coloured swimming hat such that hairstyle features are reduced as much as possible and the shape of the scalp is preserved to a reasonable degree. As for posture, the feet are placed apart such that the legs do not touch each other at any point below the groin. The arms are positioned pointing downwards, so that there is a visible branch point at each armpit and the arms/hands do not touch each other or any other part of the body below the armpits. Both armpits are located at approximately the same horizontal level and the ends of both arms/hands are located at approximately the same horizontal level, which in turn is located higher than the level of the groin. However, no other information regarding body geometry, size, or measurements is assumed.

3.3. Body Resegmentation

The output of the fitting to the horizontal slices is a set of curves structured in such away that it is very easy to detect the boundaries of the segments by simply counting the number of curves for each slice. This number reveals which part of the body the slice belongs to. The scanning posture ensures that the tips of the fingers are at a level higher than the groin, so we have: (i) from floor to groin: 2 curves/slice, (ii) from groin to finger tips: 1 curve/slice, (iii) from finger tips to armpits: 3 curves/slice, and (iv) from armpits to top of head: 1 curve/slice. This sort of segmentation is convenient for the skinning scheme that follows, as all segments produced have the topology of a deformed cylinder. We have six primary segments (fig. 1): Left leg, Right leg, Torso, Left arm, Right arm and Shoulders/head (one segment). In addition, we further divide each primary segment into secondary segments, defined to have the properties described below.

3.4. Further Segmentation

Each primary segment joins to another one, in the sense that e.g. the ‘top edge’ of the leg joins with the ‘bottom edge’ of the torso, the shoulders to the top of the torso and arms etc. We now regard each secondary segment as a quadrilateral such that the left (right) edge of a secondary segment can only join with the right (left) edge of a secondary segment that belongs to the same primary segment. In addition, if a primary segment joins with another primary segment and there is no branching, we insist they have the same number of secondary segments, which must also join in a one-to-one manner. Similarly, if a primary segment (‘stem’) joins with another two primary segments in a branching manner, then the number of the secondary segments of the ‘stem’ must be equal to the sum of the numbers of secondary segments of the ‘branches’. Finally, when the top (bottom) edge of a secondary segment joins with the bottom (top) edge of a secondary segment that belongs to a different primary segment, their left (right) edges must join smoothly to produce a smooth curve.

If these requirements are met, then the whole body surface can be divided into quadrilateral surface patches that follow an orthogonal connectivity, even if the body has three branches, and joining the segments smoothly is conceptually easy. Figure 2 illustrates these ideas, and indicates how the body should be topologically segmented.

4. CURVE RESAMPLING

Resampling of the least-square B-spline curves fitted to the horizontal slices is necessary in order to produce a set of points




Figure 3: Re-ordering and calibration of samples. A is their centroid. B is the one taken as a starting point. Then C is the fourth sample, and if the determinant of triangle ABC is negative, the order of samples must be reversed.

through which the surface will be interpolated. In the sections that follow, a number of curve sampling strategies are discussed, and the advantages and disadvantages of each are presented. They have all been implemented and tested in order to discover which strategy it is best to use in each part of the body according to the parts' different topological characteristics.

4.1. Uniform parametric intervals

The simplest idea was to sample each slice by evaluating it at equidistant values of the curve parameter. This is straightforward and, for the B-splines used, has the advantage that samples are more dense at parts of the curve where detail is higher. However, samples require further processing in order to produce a proper dataset. For example, we need all sample sequences for the torso slices to begin at the same azimuth, (say at the spine) and to continue clockwise. With the uniform parameter sampling scheme, sequences will begin where the spline curve begins and will continue in the same sense as for each curve. Therefore, after such a sample sequence is acquired, it has to be re-ordered and calibrated as illustrated in figure 3. The centroid of the samples in each slice is calculated and an axis of symmetry is defined. The sample on the positive side of the x-axis (front to back) and closest to the y-axis (left to right) is then taken as the starting point for each resampled curve and the whole sequence is shifted cyclically. To calibrate the sense of the sample sequence we use three points: the centroid, the first sample and a later one (in practice if we have about 15 samples/slice then the third or fourth sample works well) are then used to determine the sense of the resampled curve and it is reversed if necessary.

The above works well for most parts of the body, but produces a poor skinning result in areas such as the legs where the least-square B-splines are poorly fitted because of missing data. Also, since there is no functional relationship between samples in successive slices, there may be a helical connectivity between them that results in ripples or twists in the final surface. This problem indicates that, in order to get a proper surface, a scheme has to be followed, in which the sampling on one slice depends on the location of the respective sample in the previous slice. The methods that follow utilise such a relationship and, in the final implementation, uniform sampling is only carried out on the slice sampled first (groin slice) in order to produce a 'seed sequence' from which samples on the other slices are grown.

4.2. Torso slice growing based on closest distance from previous sample

In order to obtain a set of samples free of such twists, we compare samples on each slice with those on the previous slice. The i -th sample on the current slice is defined to have the minimum distance from the i -th sample on the previous slice (fig. 4). This also handles the case of very poorly fit curves, when, as typically found for the least square curve fitting algorithm, parts of the curve represent the correct form of the slice, but there are, in addition, spurious loops in noisy regions or where data is sparse. The closest distance method ensures that samples are not taken on these spurious loops. Hence, provided the seed slice is well-fitted, sampling of other slices is fault tolerant (fig 5). However, a disadvantage of this method is that there is no constraint to guarantee a uniform angular distribution of samples. As a result, it sometimes happens that neighbouring columns of samples tend to converge as we move away from the seed slice, producing vertical folds and ripples in the final surface. This convergence is mainly observed where there are concavities, such as the groin and sternum areas. In the following sections, modifications that overcome this problem by using alternative continuity relationships are described.




Figure 4: Torso slice growing based on closest distance



Figure 5: Fault tolerance of the 'closest distance' scheme




Figure 6: Torso slice growing based on intersection with curve normal

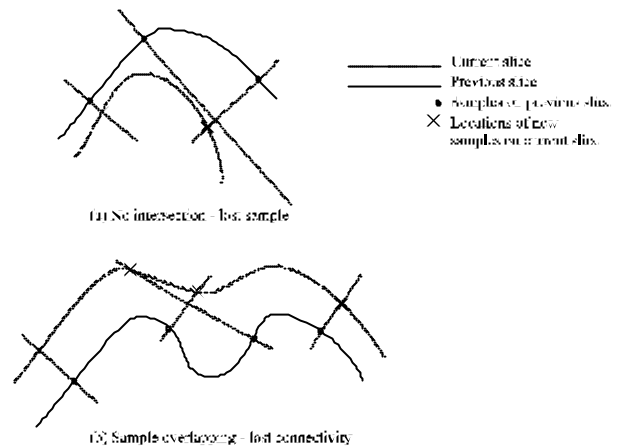


Figure 7: Problems of the 'curve normal' scheme

4.3. Torso slice growing based on intersection with curve normal

The idea here is to calculate from the DeCasteljau triangle of the B-spline the plane normal to the curve in the previous slice at the i -th sample, find where it intersects the current slice, and set the point of intersection to be the i -th sample on the current slice (fig. 6). If there is more than one such point of intersection, the one that is closest to the i -th sample on the previous slice is chosen. However, where curvature is high: (i) either intersection is at a point such that the proper sequence (as determined by hand) is overridden (fig. 7 - bottom), or (ii) there is no intersection point (fig.7 - top). A further refinement using a chord bisector was therefore used.

4.4. Torso slice growing based on intersection of chord bisector

The idea here is that, in order to get the i -th sample on the current slice, instead of intersecting the slice with the normal at the i -th sample to the previous slice, we intersect with the bisector of the line segment defined by the $(i-1)$ th and the $(i+1)$ th sample on the previous slice (fig. 8). This ensures that the sample distribution is taken into account and, given the types of curves encountered in this application, significantly reduces the likelihood of overlapping the ordering of samples. It has also been empirically observed that the "no intersection" case is much less frequent than if the normal is used as in 4.3 above. For example, when we used the two schemes on the same number of subjects, their average "no intersection" rate was 0.3% when the normal was used, whereas the chord bisector scheme always gave intersections. Fault tolerance is also higher, because this scheme is not affected by spurious cusps on poorly fit curves that would deflect the normal.

4.5. Leg slice growing based on bounding box

The geometry of the legs is slightly different, and even the chord bisector scheme does not work there. The problem is that, as we grow from the groin slice downwards, sample column convergence begins at the knees because of the concavities on




Figure 8: Torso slice growing based on intersection with chord bisector




Figure 9: Leg slice growing based on bounding box




Figure 10: A geometrically feasible solution to the smooth curve join problem

the sides of the knee-caps. As growing moves further down, samples on the front side of the leg tend to diverge sideways, and finally the shape of the foot is lost. For this reason, an alternative method has to be used for the legs. Here, we force sampling all around the circumference of leg slices by calculating a bounding box of the slice. The bounding box is created globally with respect to the scanner's frame of reference, allowing us to set central sample points wherever the bounding box is tangent to the slice (fig. 9). Other sample points are then obtained by constructing chord bisectors as above. Although this method guarantees that samples will be taken from all around the slice, there is no sense of 'growing', and again helical topology might still emerge as the central sample points shift from slice to slice.

5. SURFACE FITTING

In this section we describe the procedure that generates the surface. The problem breaks down to interpolating a B-spline curve through a sequence of points. Once this is achieved, the method can be expanded to generate surfaces that interpolate a grid of orthogonal topology, as it can easily be shown⁸ that the problem is separable. Furthermore, we want the subsequent surface to be such that each of its edges goes through one of the boundary rows or columns of the grid given above. We may thus start working with each row of data points first, and then process each column of the result matrix.

5.1. Ensuring smooth continuity between successive cubic B-Splines during interpolation

Given that we have two sets of data points in the same slice or same column, and that we want to interpolate a cubic B-spline to each one of them, such that the curves are successive and join together smoothly (G^1 continuity), the problem can be formulated as follows: We have a sequence of data points p_1, p_2, \dots, p_n and another sequence of data points d_1, d_2, \dots, d_l such that the two sequences are successive, i.e. $p_n = d_1 = p$. First of all, we want the two interpolated curves to join, so the last control point of the first curve, v_{n-1} , must coincide with the first control point of the second curve, w_{-2} , which in turn must coincide with the first data point of the second curve: $v_{n-1} = p = w_{-2}$. In addition to this constraint, we want them to join smoothly with G^1 continuity so that the tangent direction at the end of the first curve is the same as the tangent direction at the start of the second curve. Points v_{n-2} and w_{-1} are thus chosen such that v_{n-2} , w_{-1} and p are collinear (fig. 10).

However, the problem remains underconstrained. We can set the three points v_{n-2} , w_{-1} and p to be collinear in direction \underline{g} , but this direction is not yet determined. Moreover, the magnitudes of the tangent vectors are also not yet determined. They




Figure 11: A semantically feasible solution to the smooth curve join problem




Figure 12: The result of the smooth join for various values of factor f : (a) $f=0.1$, (b) $f=0.3$, (c) $f=0.5$, (d) 0.8

depend solely on the distances of v_{n-2} and w_{-1} from p . Therefore, a strategy is required empirically to set the control points v_{n-2} and w_{-1} such that the form of the join is semantically feasible and the resulting curves do not over-oscillate or produce ripples. We thus looked at the behaviour of the neighbouring data points that define the two curves and noticed that the join appears to be optimal when \underline{g} is parallel to the direction of the line defined by sample points p_{n-1} and d_2 on either side of p . This guarantees (fig. 11) that the direction of \underline{g} lies between the directions of the tangent vectors at p_{n-1} and d_2 in such a way that tangent direction from p_{n-1} to d_2 changes in a ‘monotonic’ manner. As for the magnitudes of the tangents, we set the distances between p and the neighbouring control points according to the distances from p to the neighbouring sample points, which provide an estimate of the curvature of the ideal smooth curve around the join. In the final implementation, these distances are calculated and multiplied by a normalisation factor f . Having calculated these normalised distances and the direction vector \underline{g} , control points v_{n-2} and w_{-1} are uniquely defined. The best value for f has been empirically found to be about 0.3. Higher values give very large distances and generate a 2-lobed ‘stretch’, while lower values generate concavities around the join that result in a ‘pointed lobe’. These effects are illustrated in figure 12. In cases where we have only one curve, that we want to be a smoothly closed loop, this technique can be used without change.

5.2. B-spline surface fitting on a smoothly closed deformed cylinder

Given the above, smoothly joining the seam of a deformed cylinder implemented as a quadrilateral B-spline surface is




Figure 13: The general case of four quadrilateral segments joining together with an orthogonal connectivity




Figure 14: Four quadrilaterals joining to form a branching object. It can be seen that it is only a special case of the connectivity shown in figure 13

straightforward. We interpolate row-wise first regarding each row of sample points as a point sequence for interpolation of a smooth closed loop by duplicating the first data point at the end of the sequence. This is repeated for each row of data points, after which we can proceed to interpolate in the same way along each column.

5.3. Smooth segment joining

The last problem to be solved is to make the surfaces that represent each segment join smoothly with each other. For simplicity, the case presented here is the simplest one of two quadrilateral segments joining edge-to-edge, but it is not difficult to see that this can be extended to any sort of quadrilateral segments with orthogonal connectivity. Thus, we let: $\{a_{ij}; i = 1, \dots, m; j = 1, \dots, n\}$ be a grid of data points that define a quadrilateral surface segment A and also, similarly, $\{b_{ij}; i = 1, \dots, m; j = 1, \dots, l\}$ a grid of points that define a second segment B. We want them to join in such a way that the 'right' edge of A joins smoothly with the 'left' edge of B. This means that $a_{pj} = b_{jl}$ for $j = 1, 2, \dots, m$, and that we have to take the j -th row of A and the j -th row of B and interpolate through the two sequences so as to create two curves that join smoothly as specified above, after which we interpolate column-wise as before to get the surfaces. In case we have four segments to be joined smoothly, with a connectivity like the one illustrated in figure 13, then we proceed similarly. First, we interpolate each segment row-wise, so as to ensure a smooth join between 1 and 2 and between 3 and 4, and then we interpolate each segment column-wise, so as to ensure a smooth join between 1 and 3 and between 2 and 4. In the case that the segment has an edge where it does not join to another segment, we set the 'free' control points of each row/column to coincide with the endpoints of the data point sequence(s). This eliminates the degree of freedom that contributes to the curvature at the boundary of the segment, so that the curvature is uniquely defined by the rest of the data points. A typical example of such edges are the lower edges of the leg segments, which define the shape of the feet at the bottom-most slice.

Figure 14 illustrates a more complex case, of four segments with a branching topology. We have a 'torso' that is broken into two secondary segments, the 'left torso' and the 'right torso' (according to the segmentation rules described earlier) as well as two 'legs'. This is analogous to the simpler cases, so we: (i) interpolate row-wise on the left torso and the right torso setting the degrees of freedom so that the left edge of the right torso joins smoothly with the right edge of the left torso along the 'front seam', and also so that the right edge of the right torso joins smoothly with the left edge of the left torso along the 'rear seam'; (ii) interpolate row-wise on the left leg and on the right leg using the closed cylinder strategy; (iii) interpolate column-wise on the left leg and left torso so that they join smoothly; (iv) interpolate column-wise on the right leg and right torso so that they join smoothly. It must be noted that in the last two steps, we do not enforce any smoothness constraints on the curves that meet right on the saddle point.

6. RESULTS

6.1. Output of the algorithm, analysis and critique of the results

The body surface produced by the aforementioned procedures is shown in figure 15. The sampling scheme employed in constructing figure 15 was the one using uniform parametric intervals with re-ordering and calibration. The sampling schemes used were: (i) for the torso: intersection with chord bisector, and (ii) for the legs: bounding box. The shoulder area has been included with the torso, with the arms ignored so that both shoulders and torso were incorporated into a single primary segment with the topology of a deformed cylinder. This was done in order to investigate the behaviour of the sampling schemes in the bust area, since this is of particular interest for clothing design applications. It also shows the




Figure 15: Result of the algorithm (final surface) after interpolation and smooth joining

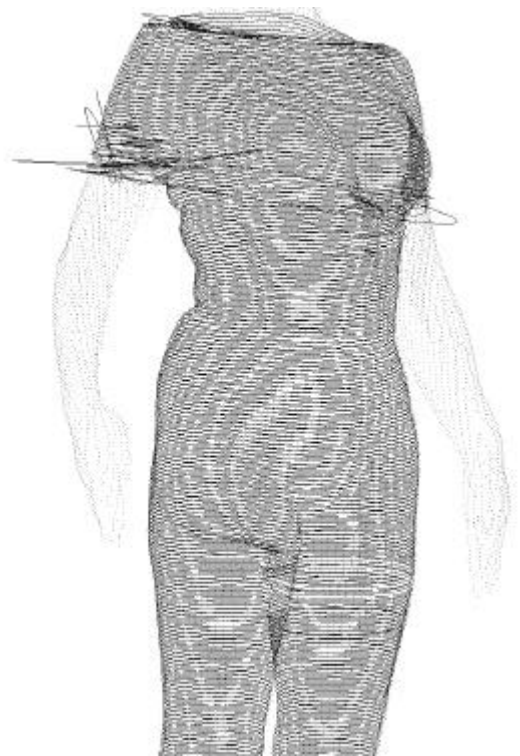


Figure 16: The set of least-squares fitted slice curves from which the surfaces of figure 15 were produced

tolerance of the algorithm to poorly fit curves. Figure 16 shows the set of curves from which the surface shown on figure 15 was generated. These curves are not well fitted around the armpits. However, the final surface is still true to the original data. Although the surface is generally unaffected by outliers, meaningful details such as the 'pit' around the navel and the concavities around the knees, are preserved. The effect of outliers, although being rather strong for some of the initial slice curves, is significantly reduced in the final surface. The concavity along the spine is also a feature that is preserved in the final surface. It is only interrupted at the armpit level, because the subject was scanned in her underwear. In general, surface quality tends to be better for the torso than for the legs, where it is difficult to sample the data appropriately. All segments join seamlessly and the surface is smooth everywhere. The only features that have been completely lost are the natural cusps of body folds (more frequent on obese subjects), but this behaviour was predictable since this method emphasises surface smoothness and is not capable of modelling discontinuities. To achieve the latter, the algorithm would need to be redesigned so that discontinuities could be detected at the early stages of processing and related information retained throughout the process, in order to contribute to the final surface interpolation stage.

6.2. Further work

We have already indicated some aspects of the system that need improvement. There are, in addition, a number of areas of work that were not discussed but are ultimately required for the completion of our body models. The resampling proves to be very effective as it greatly simplifies the points-to-surface problem when the initial point set is unorganised. The ideal solution would be to sample each slice at equal angles around a principal axis. This is the approach of the LASS method², which works well if only the torso is of interest and manual stages are included in the procedure. Although such an approach does not scale easily, if at all, to the case of a branching object, it suggests a solution that involves locating sample points where all slices intersect a single plane. This would guarantee that no helical geometries will occur, and it would therefore be interesting to investigate how this plane should be defined for the case of a complex, multisegmented object such as the human body. In addition, the arms and head are currently excluded from the procedure. However, provided we exclude fingertips and the very top for the head, the arms and head can be processed in exactly the same way as the legs and torso, and surfaces could be generated in a similar manner. These surfaces can be easily connected to the existing skinned segments, in the same way as legs are connected to the torso.

When the above have been achieved, the only problem remaining to be solved is the closing of the surface at the top of the head and at the ends of the arms. This cannot be done using quadrilateral B-spline patches, because singularities occur. One of the approaches that seems to be the most appropriate is to model these parts of the body as deformed hemispheres. This can be done by, for example, locating the topmost point of the head and constructing curves that start from the top and extend radially to the sample points of the next, lower slice. These curves, together with the respective segments of the top slice curve can define triangular B-spline surface patches, and smooth joining of those patches can be guaranteed since triangular patches also incorporate degrees of freedom that can be manipulated in order to achieve the desired tangency at their boundaries.

7. EXAMPLE APPLICATION: VOLUME AND SURFACE AREA CALCULATION

This section explains how the system developed can be conveniently used in order to calculate body volume and surface area. These calculations are important in medical research and also, for some applications, in the clothing industry⁷. To evaluate the accuracy of the proposed methods, a detailed comparison with underwater weighing and body pod data is currently being carried out as implementation of the method for the whole body is approaching completion.

7.1. Volume calculation

Since two neighbouring slices are parallel and the distance between them is known, we can easily calculate the volume between the two slices from the area of the curves defining them and sum the slice volumes to estimate the whole body volume. In particular, we can use the fact that slices are defined by 3rd degree curves and calculate the areas using an iterative subdivision method⁹. Obtaining the volume is then a problem of numerical integration which may be carried out in a manner analogous to the trapezoid rule for integrating a function. Alternatively, numerical integration may be carried out in a manner analogous to the Simpson's rule⁹, if we take into account that the wall of the deformed cylinder is also a 3rd degree surface that interpolates between the slices.

7.2 Surface area calculation

The above idea can extend easily to calculation of the surface area of the body. In this case, the first step is to calculate the length of each slice by iterative subdivision of the B-spline curves. Obtaining the surface area is then again a problem of numerical integration. The latter may be performed geometrically by approximating the surface between slices as a ribbon or, more accurately when it is very slanted as at the shoulders, by a frustum of a cone⁹. Alternatively, an adaptive subdivision algorithm can be used that recursively calculates the surface area of each (secondary) segment, and adds the results to obtain an estimate of the surface area of the whole body. The use of parametric B-splines makes this method fairly easy to implement, since such a subdivision is used in standard algorithms for rendering purposes⁸.

8. CONCLUSIONS

This work has shown that in general, B-spline surfaces in conjunction with least square B-spline curves are very useful for skinning body scanner data as they enable us to deal with problems of noise, minor artefacts, contradictory scanner head

data, poor sampling due to occlusion, smooth segment joining and branching topologies. The current implementation has proved to be successful, and no unexpected problems were encountered. It is also reasonably fast and takes only a few minutes on a Silicon Graphics O₂ workstation, provided that the least-squares slice curves are given as input. However, the least squares curve fitting procedure requires up to a few hours of computer time on a similar machine. This work demonstrates that the surface-from-point-set problem is greatly simplified if the body is first segmented and curves are fitted to subsets of data in order to drive the actual surface patch fitting problem. It also makes it clear that non-generic methods that incorporate assumptions that are related to the structure and geometry of the human body have a strong potential and are appropriate for expansion and further development.

ACKNOWLEDGEMENTS

The authors would like to thank Hamamatsu Photonics UK for providing the Body Lines Scanner, Jonathan Wells and Nigel Fuller from the Dunn Nutrition Research Centre, Addenbrookes Hospital, Cambridge, for their assistance in the acquisition of volumetric data, and Elliot West for providing the design and implementation of the Least Squares B-spline Curve Fitting algorithm. They would also like to express their thanks to the anonymous volunteers who offered to have themselves scanned for the purposes of this research. The work presented was partially funded by the Wellcome Trust.

REFERENCES

1. Horiguchi,C., 'BL (Body Line) Scanner - The development of a new 3D measurement and Reconstruction System', Hamamatsu Photonics, International Archives of Photogrammetry and Remote Sensing, Col. XXXII, Part 5, Hakodate 1998.
2. Jones,P.R.M., Brooke-Wavell,K., West,G., 'Format for Human Body Modelling from 3D Body Scanning', HUMAG Research Group, Univ. of Loughborough, International Journal of Science and Technology, Vol.7, No.1, 1995 (pp. 7-16).
3. Eck,M., Hoppe,H., 'Automatic Reconstruction of B-Spline Surfaces of Arbitrary Topological Type', Computer Graphics (SIGGRAPH '96 Proceedings), pages 325-334.
4. Stoddard,A., Hilton,A., Illingworth,J., 'Slime : A new deformable surface', Department of Electronic and Electrical Engineering, University of Surrey.
5. Lorensen,W., Cline,H., 'Marching Cubes: A High Resolution 3D Surface Construction Algorithm', Computer Graphics Vol. 21, No. 4, July 1987.
6. West,E., 'B-Spline Surface Skinning for Body Scanner Data', MRes Project Report, Department of Computer Science, University College London, September 1997.
7. Dekker,L., Khan,S., West,E., Buxton,B., Treleaven,P., 'Models for Understanding the 3D Human Body Form', IEEE International Workshop on Model-Based 3D Image Analysis, in conjunction with ICCV '98, January 1998.
8. Piegl,L., Tiller,W., 'The NURBS book', Second edition, Springer-Verlag Berlin, 1997.
9. Douros,I., 'B-spline surface reconstruction of the human body from 3D scanner data', MRes project report, Department of Computer Science, University College London, 1998.