Animating Scanned Human Models

João Fradinho Oliveira
Department of Computer Science
University College London
WC1E 6BT
London, UK
joao.oliveira@cs.ucl.ac.uk

Dongliang Zhang
CAD\CAM Center
Hong Kong University of
Science & Technology
Clear Water Bay, Kowloon,
Hong Kong

Bernard Buxton
Department of Computer Science
University College London
WC1E 6BT
London, UK
b.buxton@cs.ucl.ac.uk

Bernhard Spanlang
Department of Computer Science
University College London
WC1E 6BT
London, UK
b.spanlang@cs.ucl.ac.uk

ABSTRACT

We present techniques for automatically creating and animating models obtained from human whole body scanned data. A layered model is developed in which the underlying skeleton, simplified surface and mapping of the surface to the skeletal structure are generated without manual intervention. External body landmarks such as those employed in anthropometric surveys are used to obtain the skeleton, these landmarks are then reused to help map the model surface consistently to the skeleton. The main contribution of this work is to show how automatically or even manually extracted surface landmarks greatly simplify the task of skeleton generation and surface bone mapping, the later uses surface connectivity to *grow* surface regions as opposed to conventional volume searches. The techniques are illustrated by animation from motion capture data in which a vertex blending technique is used to create continuous and smooth deformation.

Keywords
character animation, skeleton generation, mesh simplification, vertex blending, deformation.

1. INTRODUCTION

Real-time computer generation and animation of realistic 3D human models presents several research challenges. For example, in an application where a small number of high fidelity models are to be used in close-up, whether it be in entertainment, communications, or a virtual environment, it is necessary that the models' appearance is personalised, in shape and look, to resemble real humans and that they move and deform smoothly and realistically. This is especially the case if the models are to be used to represent specific individuals, but it is also true whenever a number of distinct characters are required in order to populate a computerised world, and remains true to some extent even when many electronic characters are required in crowded scenes if undesirable artefacts are to be avoided.

However, creating and animating such models is often a time consuming task, requiring detailed input from skilled modellers, animators and creative artists. Although a number of systems have been developed for such applications in computer graphics, there are none that can deal with the high resolution 3D scans of real people that are becoming increasingly available. For example, human whole body scanners have been produced by Cyberware, TC2, Virtronics, Wicks and Wilson, Hamamatsu and others, often for anthropometric purposes or for applications in the fashion and clothing industries [8]. Such systems may be used to produce scans of particular individuals or of as many people as needed, for example for sizing surveys, but the data they provide is usually a static image of the body skin or of some close fitting undergarment.

Journal of WSCG, Vol.11, No.1., ISSN 1213-6972
WSCG'2003, February 3-7, 2003, Plzen, Czech Republic.
Copyright UNION Agency – Science Press

The aim of this work is thus to bring together and develop techniques that would enable such high resolution scans to be used automatically to produce the layered human models frequently used in animation. This requires that the scans be segmented into head, torso, arms and legs, that joints be identified and a skeleton-like structure inserted inside a scan, that surface points be mapped to each section or 'bone' of the skeleton, and that the surface be deformed appropriately as the skeleton is moved to animate the model. In addition, it is frequently necessary to simplify the models created from the scanned data in order to reduce the computational burden in many animation applications, especially those using a number of models and/or needing real-time operation as in a virtual environment. Level of detail representations of such models therefore also need to be produced, in particular [18], to reduce the number of triangles from over 100000 in the original model, to a few hundred triangles without loss of joint articulations or of limbs.

In this study, the human model is represented by a triangular mesh and described as indicated above, by a layered model [24]. The model is constructed from a point cloud, which is obtained from a Hamamatsu Body Line Scanner [13], and refined by using the techniques of surface reconstruction [7] and mesh simplification [18]. We use fully automatic techniques for segmenting the 3D body scans and locating their key landmarks as published in [8] and [4]. These surface landmarks are close to the joint locations that we wish to articulate as shown in Fig. 4.

Our main contribution from this study is to show how automatically or even manually extracted surface landmarks greatly simplify the task of skeleton generation (section 3.3) and surface bone mapping (section 3.4). Finally we apply a vertex blending technique to generate smooth human deformation.

We have effectively applied our method to a variety of human models, ranging from male and female adults to models of children (Fig. 10). This would not be possible if the techniques required manual intervention at any stage. The structure of this paper is as follows. In the next section, we briefly review previous work on human animation. In section 3, we introduce our layered human model, review the feature extraction and segmentation in section 3.1, mesh reduction in section 3.2, in section 3.3 we present the automatic skeleton generation method, in section 3.4 we present our novel surface bone mapping technique and in section 3.5 we describe the motion capture system. In section 4, a vertex blending technique is employed to deform the human model, and we present how to compute a weight function for the vertex blending. Finally, in section 5 we show some results. We summarise this paper and discuss future work in section 6.

2. RELATED WORK

Animation of realistic 3D human models, typically involves three main tasks: a) modelling of the character: skeleton creation, skin and muscle representation, possibly cloth; b) deformation: rigid deformation of the skeleton, soft deformation of the skin and possibly muscles; c) motion control: motion capture, key-frame interpolation, behaviour interpolation, control point guidance, and animation parameters or scripts.

Several human animation systems have followed a layered approach to represent skeleton, muscle, skin and control [5, 14]. For deformation of muscles, and skin, techniques derived from free form deformation (FFD) of spline control points have been used. Thalmann et al [21, 27] try to personalise avatars with a few thousand polygons by using metaballs and scaling the model via five or six parameters. For higher fidelity it would seem natural to use scanned body models. Nebel [17] constructs volumetric meshes of soft tissue above the muscle from photo realistic surface scans. The surface skin mesh is replicated three times to form epidermis, dermis and subcutaneous tissue. A finite element method (FEM) is used to model the different physical properties of the layers, and external forces are then applied to simulate skin deformation.

Chen et al [6] use a FEM to model muscle deformation, Scheepers et al [20] go one step further to model anatomy based muscles that deform the skin. Mapping the skin vertices to underlying bones or muscles, often relies on a skilled animator [15]. The task of surface to bone mapping can also be quite cumbersome for animators, due to the global nature of querying planes, we adopt a surface growing technique that does not inadvertently select other unrelated parts of the model (section 3.4). Teichmann [25] has developed a semi-interactive process to define the skeleton of an arbitrary object, based on a closed mesh. This method produces good articulated models, but requires filtering of medial axis results and still some manual intervention to help guide the skeleton creation process. In section 3.3 we show how automatic surface landmarks can find useful approximations of the medial axis for skeletons.

Sloan et al [23] combine shape blending with simple transform blending for soft deformation of the animated surface. In their system, muscles can contract via target shape position interpolation and at the same time move with the bones. Motion capture data can then be used to drive the skeleton animation.

For example Fua et al [10] detect and derive animation parameters from analysing motion video sequences. These motion capture sequences can then be interpolated with other sequences and altered to have expression such as happy or sad [19].

3. LAYERED HUMAN MODEL

In our work, the human model is considered as a deformable, articulated layered model. We use two layers: skeleton and skin. For the purpose of applying motion capture data to our model, the skeleton resembles the hierarchy of the source motion capture sensors. The structure and hierarchy of the skeleton that we wish to build is shown in Fig. 1. Subjects being scanned are typically scanned with a stance that is suitable for the initial pose matching of motion data such as the BioVision hierarchy (BVH) format [2]. Each skeleton segment is described by a local reference system used to define its local position and orientation as also discussed in section 3.4. The following subsections, describe how we build the matching layered model.

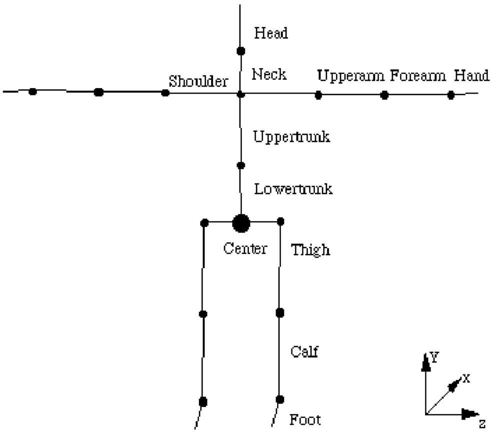

Figure 1. Skeleton structure, hierarchy and axes of the global co-ordinate system

3.1 Feature Extraction/Segmentation

In all the examples used here, the whole body data was obtained from a Hamamatsu scanner Horiguchi [13]. Data from this scanner is conveniently represented as a set of horizontal data slices made of 3D vertices [4, 13]. If data from some other system such as that developed by TC2, say, is to be used which is not represented in this way, one can use octrees to query and convert the object into equivalent slices. Once the data is in this 'slice' form, and provided it is of sufficiently high resolution to delineate the requisite details of the anatomy, the landmark detection and segmentation process discussed in [8, 4] may be used to divide the body scan into six groups of slices corresponding to the head, torso, arms and legs. We will thus only briefly review the process here. At a high level, primary landmarks are detected such as the top of the head, torso, neck, left and right armpits, crotch, and the ends of the arms and legs. These first landmarks are detected by algorithms such as that designed to locate the armpits from a re-entrant surface condition as illustrated in Fig. 2. The centroid of each horizontal data slice is calculated, and the vertices binned into sectors of angular width ß, which is related to the number of samples in the slice. The arms can be segmented from the torso by detecting the transition slice that indicates the branching point at the armpits. Sector bins will typically contain only a few vertices, but when the branching at the armpit occurs, a tolerance can be set to detect the increase in number of vertices in a bin. Similarly the neck and crotch can be detected by reasoning about the average distance to the centroid from slice to slice, and about changes in depth [8].

The detected primary landmarks define sub-search areas on the surface of the body scan in each of which a variety of discriminant functions use local shape characteristics like local maximum curvature to detect other landmarks such as: shoulder, elbow, wrist, waist, hip, knee and ankle. This algorithm is part of the system developed to automatically process the Hamamatsu body scans, for applications in medicine and in the retail and clothing industries [4].

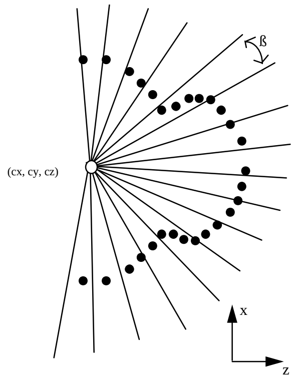

Figure 2 Underarm segmentation using re-entrant point tolerance criteria

Alternatively one can simply select manually vertices in the surface of the model that are close enough to the automatically detected landmarks mentioned above.

3.2 Mesh Simplification

Applying deformation matrices to all the 67595 vertices of the original model is computationally unfeasible for interactive animation rates and in addition, rendering the associated 135099 polygons would also prove to be a problem. To overcome these difficulties, we used the simplification method of Oliveira and Buxton [18]. In particular, since this method uses automatic edge orientation constraints to ensure a global mesh structure suitable for good normal interpolation at low levels of detail, we should be able to articulate and animate crowd scenes that require high quality, low-level of detail rendering of humans (190-250 polygons/ avatar), for example, by texture mapping facial details and clothing on to our models as in the work of Hilton et al [12]. In addition, this simplification method [18] yields a robust and clean approximation to the medial axes of the body parts, which can later be used for finding the appropriate joint positions of the skeleton. The medial axes approximation consists of a set of computed centroid vertices from each slice of the high resolution scan (Fig. 4). The following figure, 3, shows the original body scan surface on the left, and on the right, the decimated surface reduced to ~3.7% of the data quantity (4999 polygons), which is used in our animation.

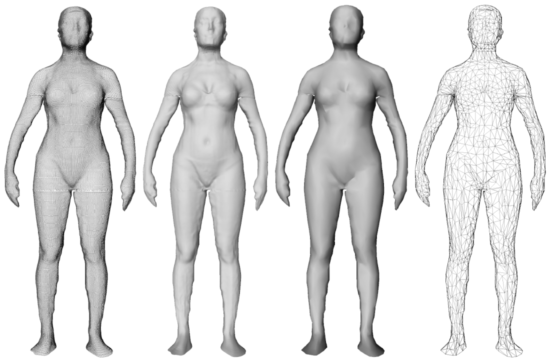

Figure 3 Mesh simplification (from left to right): original body scan model of 135099 triangles, 67595 vertices [wireframe/smooth shading],simplified model of 4999 triangles, 2535 vertices [smooth shading/wireframe]

3.3 Building the Skeleton

We can now use the detected surface landmarks of section 3.1 (Fig. 4, left) together with the axes which we regard as defining an approximate medial axis, from section 3.2 (Fig. 4, middle) to find the skeleton. For each surface landmark (neck, chest, waist, shoulder, elbow, wrist, knee and ankle etc.) we compute the distance to each vertex in the approximate medial axis, and we choose the vertex on the approximate medial axis with the shortest distance to be the corresponding joint position. The approximate medial axes described in section 3.2 was generated by joining the centroids of the original, high resolution horizontal data slices in each of the groups representing the six body segments. Since the model has been segmented as described in section 3.1, the search on the approximate medial axis set for the point that represents the joint corresponding to a surface landmark is reduced to under 200 distance calculations. To conform with the hierarchy of the motion capture data, some joints such as the uppertorso (uppertrunk) and abdominal (lowertrunk) joints (Fig. 1) are determined by making use of the ratios of skeleton segments [11]. Figure 4 (right), shows the skeleton generated automatically from the scanned human model, based on these key landmarks.

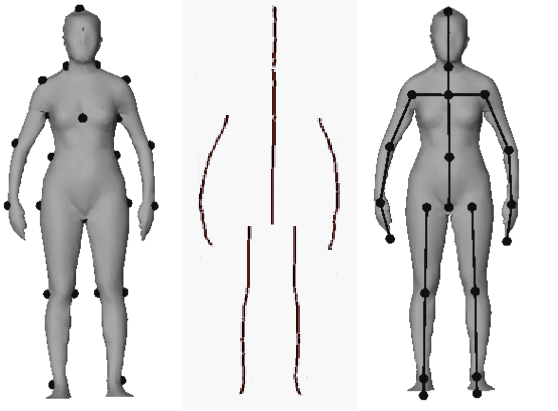

Figure 4 (from left to right) Surface landmarks, medial axis approximation, bone segments and joints

3.4 Mapping

After we have built the skeleton layer, we divide the surface mesh layer into several parts based on the skeleton, and then, map the vertices of the surface into their corresponding skeleton segments. In this process, each surface vertex (of the decimated model) is labelled as belonging to a particular segment referred to as a 'bone' [10]. This is done by placing at most joints, bi-sector separating planes, perpendicular to the plane formed by bone pairs, and testing the vertex position in relation to the plane/bone segment [24] as shown in Fig. 5. Finally we grow the surface region as illustrated in Fig. 5. To do so, we start by selecting a triangle closest to the ends of the bones at the tips of the hands and feet (Fig. 5). All triangles connected to these that do not cross the wrist and ankle bisector plane (Fig. 6-a) respectively belong to the hands and feet. A similar procedure is then applied to the lower arms and legs, etc upwards, until every vertex in the scan is assigned to a region. Two additional planes are created perpendicular to the shoulder bone (Fig. 5), approximately 3/5 of the distance from the chest joint to the shoulder joint. These are to accommodate complicated motion of the body in the shoulder region. In addition problems may arise when the orientation of a bisector plane would lead to it intersecting the interior of the body segment. In particular, planes at the shoulder do not bisect the joint angle but are constrained to contain the armpit landmarks as shown in Fig. 5 (open circles) and in Fig. 6-b. Vertices behind this plane and below some small vertical offset(shortest edge length of model) from the landmark are selected first, thus guaranteeing that no vertices beyond the armpit are selected, and then the remaining vertices from the upper shoulder are assigned by selecting vertices that are both above the offset and behind the plane.

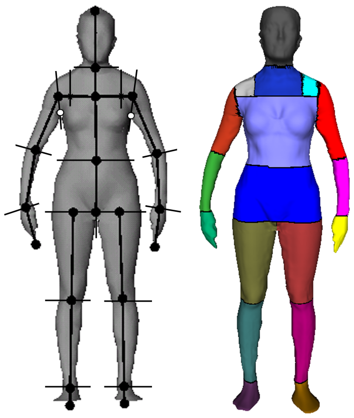

Figure 5 left) Surface separation planes, right) surface *grown* regions

Other methods would typically require the animator to tweak the separating planes so as to not select other parts of the body, or require the user to create additional constraining planes, which can be difficult to control and visualise. The mapping between the surface and skeleton, sets joint and bone proximity parameters for each vertex that will be used in the blending deformation process described in section 4. Specifically the mapping consists of calculating three parameters for each vertex: u, r and q, where u is the ratio along the skeleton segment and its value is between 0 and 1, r is the distance from the vertex to the skeleton segment, and q is the angle which represents the orientation related to the skeleton segment. From Fig. 6, we see that the mapping is based on a cylindrical coordinate system [24].

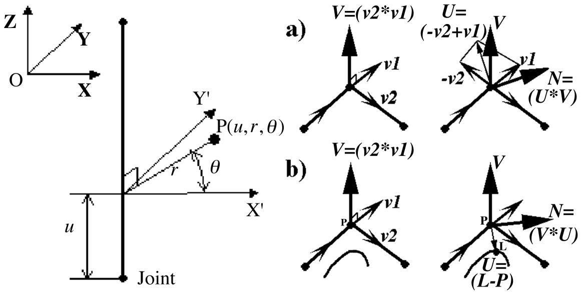

Figure 6 left) Mapping the vertex onto the bone, right) the types of separating planes: a) joint bisector, b) constrained by a landmark (rightarmpit)

3.5 Motion Capture Data

Finally we use the BVH format [2] for applying motion capture data available from the avatar modelling package, Poser 4.0 contents CD. In addition to the skeletal hierarchy of the bones and an initial pose, this format describes each motion frame as a sequence of x, y, z positions and the Euler angles of the skeleton joints. The sequence of parameters can then be applied to each bone segment's deformation matrix. We adjust the local matrices in the beginning to reflect the initial pose of the motion capture data, and then refresh the angle deformation to produce the different animation frames. We apply angular interpolation if the motion capture data available is not smooth enough.

4. VERTEX BLENDING

In order to ensure that the body surface does not break up as the skeleton moves, the above process requires that the surface is deformed and blended at each vertex as illustrated in Fig. 7.

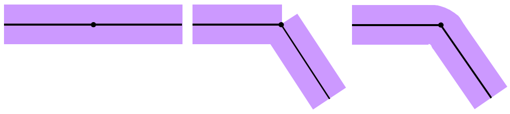

Figure 7 a) Original shape b) Rigid transformation c) Deformation using vertex blending

In our method, we use a vertex blending technique [9] to create continuous and smooth human deformation. To make this possible, several transformation matrices are used, each matrix exerting influence on its part of the object. For example, the first matrix can twist the left part of a limb in one direction while the second matrix does the same to the right part in the opposite direction. Apart from the transformation matrices themselves, one must set weight factors for each vertex. These weight factors determine the influence exerted by each matrix and are set as vertex attributes along with normal, colour, texture coordinates, etc. By way of an example let us consider the idealised model shown in Fig. 7-c. The model consists of two parts. Even though the two parts move independently, the model remains continuous and smooth. The model is transformed by two matrices - one for each model part. The matrices are based on their bones. Vertices belonging only to one bone are simply transformed along with the bone. Vertices belonging to two or more bones are transformed by several matrices after which the result obtained is averaged in proportion to the weight factors. The mathematical expression of the vertex blending for two transformation matrices may be described as:

$$\mathbf{x}=\mathbf{T}_0\mathbf{x}_0 w + \mathbf{T}_1\mathbf{x}_0 (1-w)$$

$$\mathbf{N}=\mathbf{T}_0\mathbf{N}_0 w + \mathbf{T}_1\mathbf{x}_0 (1-w)\,, \qquad (1)$$

where: x is vertex position after deformation, x_0 is the vertex position before deformation, N is vertex normal after deformation, N_0 is the vertex normal before deformation, T_0 is the transformation matrix for the first skeleton segment, T_1 is the transformation matrix for the second skeleton segment, and w is the weight assigned to the vertex. For the first segment, the weighting factor is w, and for the second 1-w. From (1), we see that a vertex's final position and normal are determined by linear interpolation, as the weight varies from 0 to 1 according to how the vertex's position is related to the skeleton segments and the skeleton structure.

In our deformation method, we define the vertex weight based on two structures two-link and multi-link structures. In a two link structure as shown in Fig. 8 (left), a joint is linked to two skeleton segments. For the human model, most pairs of skeleton segments are two-link structures, such as, the upperarm-lowarm, lowerarm-hand, thigh-calf, calf-foot, etc. However, a multi-link structure is also required as shown in Fig. 8 (right), for joints linked to more than two skeleton segments. Examples of such a linkage are the skeleton segments uppertrunk, leftshoulder, right shoulder and neck, as shown in Fig. 4 (right).

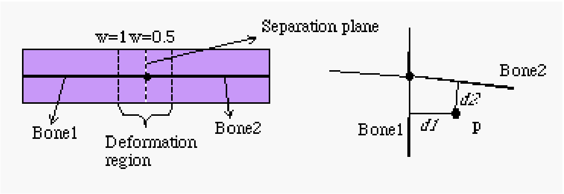

Figure 8 Computing weights for different structures: left) two-link, right) multi-link

For the first structure, we divide the surface part corresponding to a skeleton into two regions: deformable region and undeformable region. The vertices within the undeformable region are far from a skeleton joint and their weight equals 1. The vertices within the deformable region are near the joint. Their weights are between 0.5 and 1.0. As shown in Fig. 8 (left), the weight of the vertex on the separating plane is 0.5, and the weight of the vertex on the plane separating the deformable region and undeformed region is 1.0. The value of the weight for a vertex within the deformable region can be computed according to the mapping parameter u of the vertex.

For the multi-link structure, the computation of weights is more complicated. In some feature-based image morphing algorithms [1] and inverse distance weighted interpolation algorithms [22], the distance between two geometric elements is usually considered as a weighting factor. In our method, we also use the distance from the vertex to the skeleton segment to compute the vertex's weight. To reduce the computational cost of deformation, however, we only use the two skeleton segments which have the most influence on the vertex to compute the weight. As shown in Fig. 8 (right), the bones that have most influence are bone1 and bone2, so we use them to compute the weight of the vertex. Thus, if the distances from the vertex to bone1 and bone2 are d_1 and d_2, respectively, the weight w_1 of the vertex related to bone1 and the weight w_2 related to bone2 is calculated as

$$w_1 = \frac{{d_2}^2}{{d_1}^2 + {d_2}^2} \text{ and } w_2 = \frac{{d_1}^2}{{d_1}^2 + {d_2}^2}. \quad (2)$$

Vertices of two link structures in regions that always remain undeformed are mapped to a single bone. The vertices of multi-link structures are mapped to the two closest bones. Mapping vertices to more than one bone in the deformable regions of two link structures and multi-link structures, allows one to use mapping parameters such as the angle to each bone, to avoid artefacts that may arise in some kinds of animation [28].

5. RESULTS

In our human animation system, we apply BVH motion sequences to animate the scanned human model, as described in section 3.5. Fig. 9 shows some snapshots from results of such animation. In this figure, the model has 2531 vertices and 5000 triangles. Fig. 10 shows results that we were able to generate very quickly by applying our method to a variety of human models. We obtain the animation at about 100 frames/second on a PC with a 650 MHz Pentium III CPU, 128MB of RAM and a GeForce2 graphics card.

6. CONCLUSIONS

In this paper, the skeleton of a scanned human model is built automatically by making use of automatically extracted key landmarks. This method is very useful for generating articulated models without the interaction of the user. The results of our automatic system shown in Fig. 10 would not have been possible if the techniques were not automatic at every stage. The system however is flexible to cope with manually selected landmark. We selected slightly differently positioned landmarks (e.g. a knee landmark being at the back of the knee, instead of on the side of the knee), and the algorithm ensured that the corresponding nearest joint in the approximate medial axis was found, making it flexible and easily adaptable to different motion capture data.

Our surface to bone mapping technique has equally proved to be very robust in improving existing modelling paradigms. The human deformation algorithm based on the vertex blending techniques can generate smooth deformation and is very efficient. Further reducing the number of vertices in the model as described in [18] should allow for many avatars to be animated, if necessary simultaneously.

In this paper we have not addressed collision detection which is of course necessary whenever there is significant movement of the model/avatar. We are currently working on a two step strategy for collision detection in our animation system. We plan to use a method similar to [3], for the first stage of collision detection, which uses a cascade of tests, based on the intersection test of bounding volumes quickly to eliminate most of the object pairs from detailed considerations. If the higher-level intersection test of a pair of cylinders has an intersection, we will go further to detect accurate collisions. We only need to check collisions for the geometric elements that are inside the intersection space. For this second step we plan to use vertex-to-triangle collision tests [29]. Finally we plan to use spatial and temporal coherence to help predicting collisions before deformation.

In addition, further work remains to be done. A more general algorithm should be developed for automatically generating the skeleton for any deformable model, for example by applying the ideas of segmentation based on detection of branching points to more generic models. Second, an efficient deformation algorithm that can reflect the muscle deformation should be developed. In our deformation algorithm using the vertex blending technique, if we can improve the weighting function of the vertices by considering the influence of mapped muscle deformation, the results will be more realistic and equally fast, since the muscles could be modelled with the same three position parameters defined in our current mapping.

7. ACKNOWLEDGEMENTS

The authors would like to thank Hamamatsu Photonics UK for providing the Body Line Scanner. Also we would like to thank Prof. Mel Slater for comments on an earlier draft. Dongliang Zhang's contribution was funded by EPSRC grant reference GR/M46082/01. In addition, one of us [JO] would like to thank the Fundação Calouste Gulbenkian and JNICT/PRAXIS XXI, for financial support.

8. REFERENCES

[1] Beier, T., Neely, S., Feature-Based Image Metamorphosis, SIGGRAPH'92, pp.35-42.

[2] http://www.biovision.com/bvh.html, BVH format.

[3] Boulic, R., Capin, T., Huang, Z., Kalra, P., Linterrnann, B., Magnenat-Thalmann N., Moccozet, L., Molet, T., Pandzic, I., Saar, K., Schmitt, A., Shen, J., Thalmann, D.. The HUMANOID Environment for Interactive Animation of Multiple Deformable Human Characters, Computer Graphics Forum, Blackwell Publishers. Edited by Frits Post and Martin Göbel. ISSN 1067-7055, 14(3), pp.337-348, 1995.

[4] Buxton, B., Dekker, L., Douros, I., Vassilev, T., Reconstruction and Interpretation of 3D Whole Body Surface Images, Scanning 2000.

[5] Chadwick, J., Haumann, D.R., Parent, R.E., Layered construction for deformable animated characters, ACM Computer Graphics, pp. 234-243, 1989.

[6] Chen, T., D., Zelter, D., Pump It Up: Computer Animation of a Biomechanically Based Model of Muscle Using the Finite Element Method, SIGGRAPH'92, pp. 89-98.

[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, pp. 65-74 1998.

[8] Dekker, L., Douros, I., Buxton, B., Treleaven, P., Building Symbolic Information for 3D Human Body Modelling from Range Data, Proc. of the Second International Conference on 3-D Digital Imaging and Modelling, IEEE Computer Society, pp. 388-397, 1999.

[9] Dietrich, S., Vertex Blending under DirectX 7 for the GeForce 256, Technical Presentations 99

[10] Fua, P., Grün, A., Plänkers, R., D'Apuzzo, N., Thalmann, D., Human Body Modelling and Motion Analysis from Video Sequences, International Symposium on Real-Time Imaging and Dynamic Analysis, Hakodate, Japan, 1998.

[11] Henry Dreyfuss Associates., The Measure of Man and Woman, ISBN 0823030318

[12] Hilton, A., Beresford, D., Gentils, T., Smith, R., Sun, W., Virtual People: Capturing human models to populate virtual worlds. In IEEE International Conference on Computer Animation, pp.174-185, 1999.

[13] Horiguchi, C., Sensors that detect shape, J. Adv. Automation Technology V7No.3(95)pp210-216

[14] Kalra, P., Thalmann, N.M, Moccozet, L., Sannier, G., Aubel, A., Thalmann, D., Real-Time Animation of Realistic Virtual Humans, IEEE Computer Graphics & Applications, pp. 42-56, 1999.

[15] Lander, J., Skin them bones: game programming for the web generation, Game Developer, 1994.

[16] Laurent, M., Thalmann, N.M, Dirichlet free-form deformations and their application to hand simulation, Computer Animation, 1997.

[17] Nebel, J.-C., Soft Tissue Modelling from 3D Scanned Data, Deform2000, Switzerland, 2000.

[18] Oliveira, J., Buxton, B., Light weight virtual humans, EUROGRAPHICS-UK'01, pp. 45-52.

[19] Rose, C., Cohen, M., F., Bodenheimer, B.,Verbs and Adverbs: Multidimensional Motion

Interpolation, IEEE Computer Graphics and Applications, pp. 32-40, September, 1998.

[20] Scheepers, F., Parent, R., E., Carlson, W., E., May, S., F., Anatomy-Based Modelling of the Human Musculature,SIGGRAPH'97,pp163-172

[21] Shen, J., Thalmann, D., Interactive shape design using metaballs and splines, ACM Computer Graphics, SIGGRAPH'86, pp. 151-160.

[22] Shepard, D., A two-dimensional interpolation function for irregularly-spaced data, Proc. 23rd National Conference ACM, pp. 517-524, 1968.

[23] Sloan P. P., Rose C F III, Cohen M F, Shape and Animation by Example, Technical Report MSR-TR-2000-79, 2000

[24] Sun, W., Hilton, A., Smith, A. R., Illingworth, J., Layered Animation of Captured Data, Animation and Simulation '99,10thEurographics Workshop Proc., pp145-154, 1999.

[25] Teichmann, M., Teller, S., Assisted articulation of closed polygonal models, MIT Technical report and technical sketch at SIGGRAPH'98.

[26] Thalmann, D., Ronfard, R., Implicit simplical models for adaptive curve reconstruction, IEEE Trans. Pattern Analysis and Machine Intelligence, pp.321-325, 1996.

[27] Thalmann, D., Shen, J., Chauvineau, E, Fast Realistic Human Body Deformations for Animation and VR Applications, Computer Graphics International 1996, Pohang, Korea.

[28] Weber, J., Run-Time Skin Deformation, Game Developers Conference 2000.

[29] Zhang, D.L., Yuen, M.M.F., A heuristic collision detection method for cloth drape simulation, Proc. of IASTED International Conference on Computer Graphics and Imaging, Las Vegas, pp.297-302, 2000.

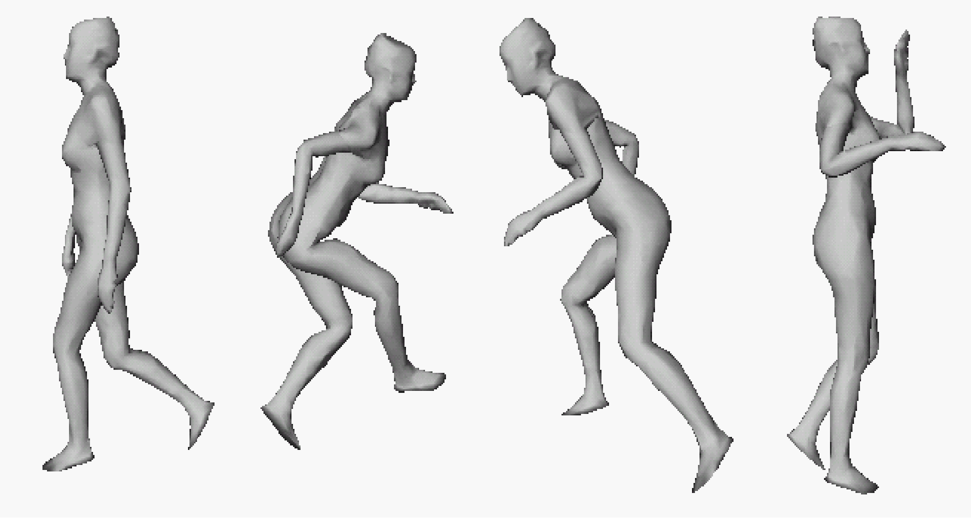

Figure 9 Several motions on a model (walking, sneaking, running and drinking)

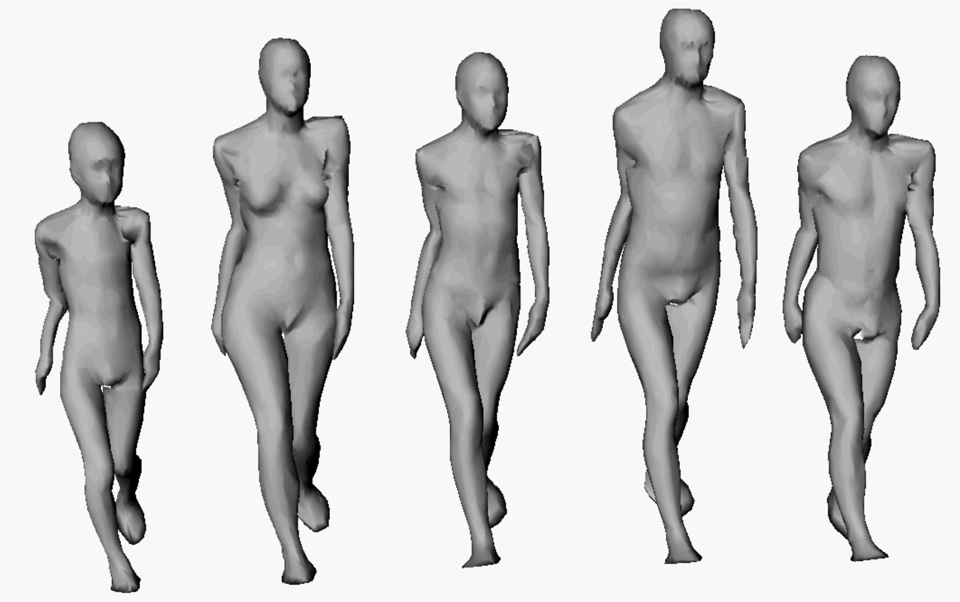

Figure 10 Motion capture data applied to different models