

# Animated Statues

Jonathan Starck, Gordon Collins, Raymond Smith, Adrian Hilton, and John Illingworth

Centre for Vision, Speech and Signal Processing,  
University of Surrey,  
Guildford, GU2 7XH, UK.

July 8, 2002

In this paper we present a layered framework for the animation of high-resolution human geometry captured using active 3D sensing technology. Commercial scanning systems can now acquire highly accurate surface data across the whole-body. However, the result is a dense, irregular, surface mesh without any structure for animation. We introduce a model-based approach to animating a scanned data-set by matching a generic humanoid control model to the surface data. A set of manually defined feature points are used to define body and facial pose, and a novel shape constrained matching algorithm is presented to deform the control model to match the scanned shape. This model-based approach allows the detailed specification of surface animation to be defined once for the generic model and re-applied to any captured scan. The detail of the high-resolution geometry is represented as a displacement map on the surface of the control model, providing smooth reconstruction of detailed shape on the animated control surface. The generic model provides animation control over the scan data-set, and the displacement map provides control of the high-resolution surface for editing of geometry, or level of detail in reconstruction or compression.

**Key words:** Human animation – Deformable model – Displacement map – Model-based vision

## 1 Introduction

The goal of physical realism is one of the most demanding tasks in the field of human body modelling and animation. High quality and photo-realistic computer graphics models of people can now be obtained from commercial systems [10], based on dense surface measurement techniques such as laser scanning [3], stereo photogrammetry [1] and active light projection [4]. While this is sufficient for static objects, the dense and irregular surface data contains no structure for animation. In this paper we propose a framework to provide both control

over the surface detail in these high-resolution geometric models and control of the model structure for animation.

A key component to realism in animated models is to provide natural surface deformations. Currently the specification and animation of surface models is a lengthy process requiring frequent intervention for realistic results [33]. Here we introduce a model-based approach to animation of high-detail human surface data, following the functional model paradigm introduced by Terzopoulos et. al. [42]. The technique is based on fitting a low-resolution generic humanoid model to the surface data to provide a control layer for animation. The philosophy of the approach is that the detailed specification of the surface animation can be defined once for the generic control model and then re-applied to animate any human data-set.

Our model consists of 3 layers as described in [39, 41]: the control model skeleton providing animation from key-frame or motion capture data; the control model, matched to the scan surface; and the high-resolution scanned data. The high-resolution geometry is represented by a single displacement map from the surface of the control model, providing an efficient representation of the surface shape and control over the detail in the layered model. This layered approach facilitates smooth animation of surface detail, control for rendering, compression and editing of surface detail.

Figure (1) shows the pipeline for constructing our layered animation model. We use several recent advances in computer graphics and computer vision. We introduce a novel shape constrained deformable surface model to automatically match a generic control model to a scanned data-set. We then apply the normal-volume representation [19] to map the surface data onto the surface of the control model and obtain a single displacement map [37]. The control model and high-resolution surface detail can then be efficiently and seamlessly animated through manipulation of the control skeleton. We reconstruct the high-resolution surface detail on the animated control surface using either uniform, or non-uniform adaptive subdivision, inserting the detailed geometry from the displacement map.

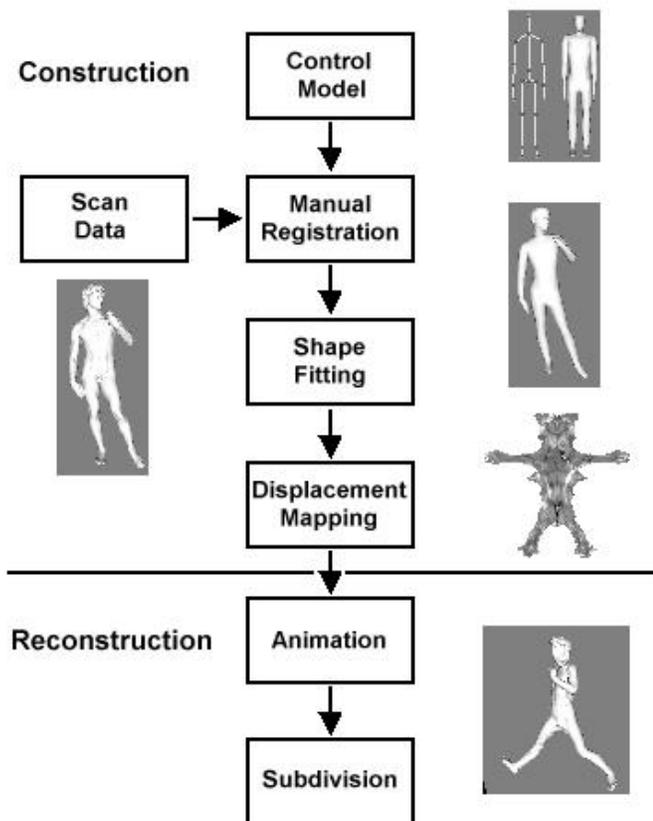


Fig. 1. Pipeline for generating layered model.

The novel features of our algorithm are:

- Registration of a generic control model with a 3D data-set to recover pose.
- Automatic fitting of the model to the data-set using shape constrained matching providing a close fit, interpolation of noisy/missing data, robust matching and preservation of mesh parameterisation for animation.
- Representation of human surface detail using a displacement map, giving an efficient and accurate representation of surface detail. Intuitive representation of detail using a single displacement map.
- Framework for representing a static whole-body scanned data-set for efficient animation, rendering, and control of surface detail.

The paper is organised as follows, in the next section we present previous work on human body modelling, and the representation of high-resolution surface geometry. We then layout the process of model registration in section 3, fitting in section 4, and displacement mapping in section 5, to construct a layered model of a full-body human scan. Finally we present results of reconstructed animated scans, mesh editing and Level of Detail control.

## 2 Background

### 2.1 Model Fitting

Modelling of the human body has been of extensive interest in computer graphics and computer vision. Our approach to model fitting is closely related to previous work on reconstructing geometric models of the human head [27, 6, 31, 20, 36] and whole body [18, 30, 21], using a series of key features to conform a generic model to match the shape of an individual. Here we combine these methods, using a set of feature points to register a generic control model with a scanned data-set, recovering the pose of the individual for both body and facial animation.

In order to match the shape of the human body in the scan, we treat the generic control model as a deformable model. “Deformable-models” were first introduced in computer vision by Terzopoulos et. al. [44], and have been used widely to reconstruct high quality geometric models from 2D image data and 3D volumetric data [34]. The technique treats a prior model as a free form shape that deforms to fit an object. The strength of the approach lies in the constrained shape deformation, allowing for the recovery of shape even in the presence of sparse or noisy data for highly variable and complex shapes.

Here we introduce a novel formulation for a shape constrained deformable model. Free-form shape deformation has a high degree of freedom and a model is free to reach a shape that is unrelated to the prior model. Techniques have been explored to constrain shape during deformation by parameterising a model with meaningful shape variables [43, 13, 40]. Parameterised shape models provide a compact representation, however the variability of the shape that can be described is limited. We therefore use a local shape constraint to allow for variability between individuals with clothing. We also formulate the local constraint to preserve the parameterisation of the control mesh under deformation. We are then able to apply the animation structure of the control model to the deformed shape of the control mesh. We combine this constrained deformable surface model with a coarse to fine matching technique to achieve robust matching of the control model to a scan data-set.

### 2.2 Displacement Mapping

Displacement mapping is an efficient way of representing a highly detailed surface with a lower resolution mesh. Instead of storing the detail as a triangulation (vector points and their connectivity) it is stored as a scalar image function mapped onto the low resolution polygonal surface (scalar displacements and texture coordinates). This representation of surface detail is analogous to texture mapping. The representation is efficient and enables rapid level of detail (LOD) control, compression and editing of detail via the displacement map image.

There have been several other research efforts towards the reconstruction of displacement maps from cap-

tured data [26, 39, 28]. Krishnamurthy and Levoy [26] ensured a continuous normal by calculating displacements from a B-Spline surface. They construct their control meshes from manually selected surface points. Smith et al. [39] introduced displacement maps on a polygonal mesh using the normal volume representation to achieve continuity of mapping between different polygons. This work is discussed further in section 5. Lee et al. [28] base their approach on subdivision surfaces and take their normals to be the normals of the limit function of their surface.

Displacement maps based on an underlying geometric model can only map surfaces that are injective to the model. A surface with many folds cannot be mapped uniquely onto an unfolded control model. The resulting mapping will be non-injective since multiple points on the detail surface will map to the same point on the control surface. This problem is countered by well chosen control meshes [28]. In general the closer the fit between control mesh and detail, the more complete the mapping. In this paper we introduce shape constrained fitting techniques to achieve a close fit between a control model and a detailed surface. Displacement mapping onto the polygonal control is then used to represent surface detail [39].

### 3 Control Model Registration

The task of model-based fitting is to deform a control model to fit the shape of the human scan data, capturing the pose for animation. The process is divided into the recovery of body and facial pose, then surface shape. In this section we describe the generic humanoid control model we use, and the registration of the model with the scan data-set.

#### 3.1 Parameterised Generic Humanoid Model

In this work we adopt the Humanoid Animation Working Group (H-Anim) specification for a humanoid model [5]. The model consists of a single seamless mesh defining the surface shape of the body, attached to an underlying skeleton structure for animation as shown in Fig. (2). The model is animated using a control skeleton with 16 articulated joints, capturing the gross pose of the human body.

Any appropriate deformation technique can be used to animate the model surface. Here we adopt a standard vertex weighting scheme widely used in current commercial software packages. The skeleton structure is animated as a rigid set of bone segments. Each bone is associated with a set of mesh vertices with a corresponding set of weights. The deformation of the vertices of the surface mesh is then given by Eq. (1).

$$\mathbf{x}_i = \sum_b w_{ib} \mathbf{T}_b \mathbf{x}_i^0 \quad (1)$$

In Eq. (1),  $\mathbf{x}_i^0$  is the global position of the  $i^{th}$  vertex in the default body pose,  $\mathbf{T}_b$  is the global homogeneous



Fig. 2. The generic humanoid control model.

transformation matrix from the default to the current pose for the  $b^{th}$  bone, and  $w_{ib}$  is the corresponding vertex weight associating vertex  $\mathbf{x}_i^0$  with the  $b^{th}$  bone. Where a bone does not influence a vertex, the vertex weight  $w_{ib} = 0$ , and for rigid body animation only one bone affects each vertex with a corresponding weight  $w_{ib} = 1$ .

The body pose of the generic model is defined by the joint rotations, limb lengths and translation of the control skeleton. We model the articulation of the control skeleton with a 3 degree of freedom (DOF) root rotation, 3DOF rotations at the vertebrae, hips and shoulder joints, 2DOF at the clavicles and wrists, and 1DOF at the elbows, knees and ankles [16, 17]. The dimensions of the skeleton are defined by 9DOF for the lengths of the spine, head, clavicles, upper-arm, forearm, hands, thigh, shank and foot, with left and right segments constrained to be symmetric.

The surface deformation scheme can be formulated in terms of the set of joint angles  $\theta$  and segment lengths  $\mathbf{l}$ , together with the global translation  $\mathbf{t}_{root}$ , Eq. (2). Here  $\mathbf{T}'_b$  is the homogeneous rigid body transformation of the local bone coordinate frame as given in Eq. (3), where  $\exp(\omega_i)$  is the rotation of the  $i^{th}$  joint in the hierarchy in an axis-angle representation [15],  $i = 0$  being the skeleton root,  $\hat{\mathbf{n}}_i^0$  is the unit offset of the bone joint from the joint of the parent bone in the default pose and  $l_i$  is the associated length of the parent bone segment. The local coordinate frames of the bone segments are defined to coincide with the orientation of the global coordinate frame in the default pose, and so the local coordinates of a vertex are given by the offset from the default location of the bone joint centre  $\mathbf{o}_b^0$ .

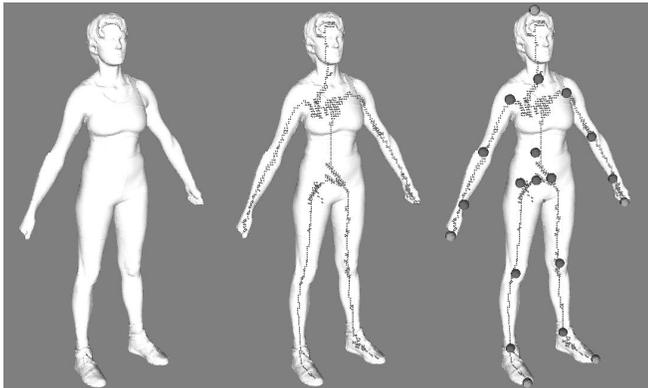
$$\mathbf{x}_i = \sum_b w_{ib} \mathbf{T}'_b (\mathbf{x}_i^0 - \mathbf{o}_b^0) \quad (2)$$

$$\mathbf{T}'_b = \begin{bmatrix} \exp(\mathbf{w}_0) & \mathbf{t}_{root} \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \exp(\mathbf{w}_1) & \hat{\mathbf{n}}_1^0 l_1 \\ 0 & 1 \end{bmatrix} \dots \begin{bmatrix} \exp(\mathbf{w}_n) & \hat{\mathbf{n}}_n^0 l_n \\ 0 & 1 \end{bmatrix} \quad (3)$$

### 3.2 Manual Feature Location

We manually define a set of feature points on the human data-set to register the parameterised control model and to define the specific shape of the face. An interactive tool is used to select points on both the control model and the scan data-set for manual registration.

Previously [14, 46] the discrete medial axis has been used to automate the placement of joints in 3D models. However, human skeletal joints such as the shoulder, do not necessarily lie on the medial axis. In this work we utilise the discrete medial axis purely as a visual guide to joint placement. We generate the medial surface by first converting the model to a volumetric voxel representation using a depth-buffer based voxelisation algorithm [23]. A topological thinning process is then applied to generate the set of voxels on the discrete medial surface [8]. Joint centres are located in the interface by user selection of joint positions on a rendered view of the model. Points on the image generate a ray in the model space and the joint positions are located in 3D as the closest point on each ray to the discrete medial surface. Figure (3) illustrates the manual placement of joint positions in a human scan. A 1cm voxel sized is used in generation of the discrete medial surface. We allow the user to adjust the joint positions along each ray to refine the estimated body pose if necessary.



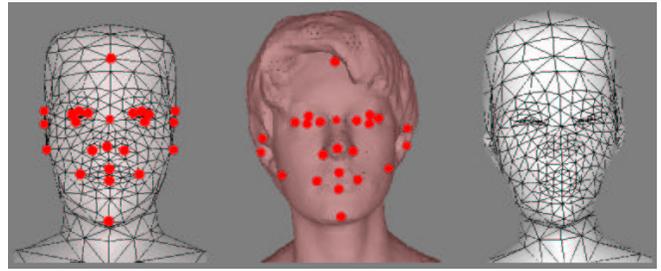
(a) Surface data (b) Medial surface (c) Manual placement

**Fig. 3.** Joint positions located using the volumetric discrete medial surface.

The correspondence between the face of the control model and the data-set is defined using a limited number of manually labelled feature points. A sub-set of the MPEG-4 facial definition points [25] are used to identify the prominent features of the face. Figure (4) illustrates the points located on the control model and the manual identification on a surface data-set.

### 3.3 Model Registration

We estimate the joint angles, dimensions and location of the parameterised generic model to match the features of the model with the joint and facial feature positions located on the data-set. We minimise the error between the points with respect to the model parameters in a



(a) Feature points (b) Manual points (c) Reshaped model

**Fig. 4.** Facial feature points located on the scan data-set.

standard least-squares framework. Equation (4) defines the least-squares problem, where  $\mathbf{p}_k^H$  indicates a feature point located on the scan data, and  $\mathbf{p}_k^C$  the corresponding feature point on the parameterised control model.

$$\min \sum_k \|\mathbf{p}_k^H - \mathbf{p}_k^C\|^2 \quad (4)$$

The problem is non-linear, requiring an iterative numerical solution. A good initialisation is important in optimisation, improving the convergence of the solver and the quality of the solution [29, 7]. We use an approximate analytic solution for the joint rotations of the model as an initialisation. The 6DOF position and orientation of the model is calculated from the location of the shoulder and hip joints under the assumption that the trunk is a single rigid body. For each limb segment we calculate the 3DOF rotation at the base joint, such as the shoulder, and 1DOF rotation at the intermediate joint, such as the elbow, from the location of the base, intermediate and distal joints. The model parameters are then refined using an iterative solution to the least-squares problem. We make use of an extensible bound-constrained BFGS solver [47], that can deal with a larger number of observed features and allows for the inclusion of constraints such as rotation and limb length limits on the model.

Once the skeleton of the control model is posed to match the features on the data-set, we update the shape of the model surface to match the geometry defined by the features to recover the facial pose. Here we use scattered data interpolation to smoothly interpolate the desired change in shape defined by the vector valued offset between the posed model features,  $\mathbf{p}_k^C$ , and the manually defined features on the scan data,  $\mathbf{p}_k^H$  [36, 21]. We fit a vector valued displacement function  $\mathbf{f}(\mathbf{p})$  to the known displacements from the model feature positions using the radial basis function  $\phi$  for scattered data interpolation, with an affine basis  $\mathbf{a}$  to model the global displacement field, Eq. (5). We use the radial basis function  $\phi(\mathbf{p}) = \|\mathbf{p}\|^3$ , minimising the thin-plate energy functional in 3D, a measure of the total bending in the hyper-surface, to give a globally smooth interpolation function [45].

$$\mathbf{f}(\mathbf{p}) = \sum_k \mathbf{c}_k \phi(\mathbf{p}_k^H - \mathbf{p}_k^C) + [\mathbf{p} \ 1] \mathbf{a} \quad (5)$$

The weights  $\mathbf{c}_k$  and the parameters of the affine transformation  $\mathbf{a}$  are solved by constructing a linear system of

equations in terms of the known vector displacements,  $\mathbf{f}(\mathbf{p}_k^C) = \mathbf{p}_k^H - \mathbf{p}_k^C$  and the constraints  $\sum_k \mathbf{c}_k = 0$ ,  $\sum_k \mathbf{c}_k^T \mathbf{p}_k^C = 0$  that remove the affine contribution from the radial basis functions [36]. The linear system is symmetric and is solved by LU decomposition. Finally the vertices of the posed control model are updated according to the interpolated vector displacement at each vertex,  $f(\mathbf{x}_i)$ . Figure (4c) shows an updated model face, matching the defined feature points.

#### 4 Matching Control Model to Scan

Once the control model is posed to match the scan data-set, the task is to deform the shape of the generic model so that it conforms closely to the scan surface without changing the topology and the parameterisation across the surface mesh. In this section we introduce a shape constrained deformable surface model to preserve the prior shape and parameterisation of the model, and formulate the matching of the model to the scan data-set to refine the control surface shape.

##### 4.1 Constrained Deformable Surface Model

The ‘‘deformable-model’’ problem is formulated as an energy minimisation task. We wish to minimise the potential energy of the model derived from the fit of the control model to the data. This is regularised by the internal energy of the model which penalises the deviation from the desired model properties. The deformable surface model  $\{\mathbf{x}\}$  minimises the energy function  $E(\mathbf{x})$  incorporating the potential energy from data fitting  $P(\mathbf{x})$ , and the internal energy from the shape of the model  $S(\mathbf{x})$ .

$$E(\mathbf{x}) = P(\mathbf{x}) + S(\mathbf{x}) \quad (6)$$

The potential energy is derived by integrating a data fitting error  $e(\mathbf{x})$  across the model surface  $\mathbf{x}(u, v)$ .

$$P(\mathbf{x}) = \int \int e(\mathbf{x}(u, v)) du dv \quad (7)$$

Traditionally the model is treated as a thin-plate material under tension, yielding a final model with a local area minima and local distortion minima that fits the data. The internal energy is given by the integral of the membrane and the thin-plate functional across the model surface where constants  $\mu, \nu$  define the trade-off between the area and distortion constraints, Eq. (8).

$$S(\mathbf{x}) = \mu \int \int (\|\mathbf{x}_u\|^2 + \|\mathbf{x}_v\|^2) du dv \quad (8)$$

$$+ \nu \int \int (\|\mathbf{x}_{uu}\|^2 + 2\|\mathbf{x}_{uv}\|^2 + \|\mathbf{x}_{vv}\|^2) du dv$$

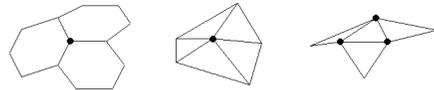
$$\gamma \frac{d^2 \mathbf{x}}{dt^2} + \lambda \frac{d\mathbf{x}}{dt} = -\frac{dE}{d\mathbf{x}} \quad (9)$$

A physics-based approach to energy minimisation is used whereby the model evolves to a local energy minima under the principles of Lagrangian dynamics. The effect is that of an elastic like body which dynamically deforms to fit the data. The model is treated as time dependent  $\mathbf{x}(u, v, t)$ , and the Lagrange equations of motion are solved for the evolution of the surface, Eq. (9).

Rather than treating our prior model as a thin-plate material under tension, we wish to formulate the local shape and parameterisation of the surface mesh in order to preserve these properties under model deformation. Local shape on a surface is described by 3D position, normal and curvature, with possible higher order terms such as change in curvature. We are interested in surface shape up to a rigid transformation, implying the use of curvature or higher order terms to formulate local shape. Such measures alone, however, do not describe the parameterisation of a mesh and would provide no constraint on the deviation from the generic parameterisation of the model.

Montagnat and Delingette [35] use a description of surface position in terms of local frames to give a rotation invariant estimate of local shape and parameterisation. The technique is based on the specific parameterisation of 2-simplex surface models, Fig. (5a). 2-Simplex meshes have 3-connected vertices, a vertex position can therefore be described by the barycentric coordinates in the plane of the 3 neighbour vertices and the simplex angle defining the offset from the plane. Montagnat and Delingette used a spring force between the current vertex positions during deformation, and the reconstructed default positions based on the neighbourhoods of the vertices, to constrain both shape and parameterisation of the mesh. We generalise this approach to the irregular triangular mesh commonly used for shape representation in computer graphics.

For an irregular mesh a vertex position is not well defined in relation to the vertex neighbourhood, Fig. (5b). With an irregular number of vertices in the 1-neighbourhood, it is not possible to obtain a consistent definition of a local frame to describe the position of the central vertex. We therefore consider a triangle face-based scheme as used by Kobbelt et. al. [24]. The vertex positions of a triangle face can be defined by the barycentric coordinates and height offset in the local frame of the vertices on the faces edge-connected to the central face, the vertices surrounding the central triangle as shown in Fig. (5c). The position of a mesh vertex is therefore constrained by the local position in the triangle face frames in the 1-neighbourhood of the vertex, leading to a 2-neighbourhood support structure for a vertex position.



(a) 2-Simplex mesh (b) Triangular mesh (c) Triangle frame  
**Fig. 5.** Vertex locations in a triangle frame.

We define the internal energy of a shape constrained model as the integral across the surface of the deviation in the local shape from the generic shape defined in each face based frame. This is given by the summation of the error at the mesh vertices,  $\mathbf{x}_i$ , preserving the local parameterisation and shape in the vertex positions. Equation (10) defines the internal energy where  $(\alpha_{if}^0, \beta_{if}^0, h_{if}^0)$  are the default barycentric coordinates  $(\alpha, \beta)$  and height offset  $h$  in the  $f^{th}$  face based frame for the  $i^{th}$  vertex with valence  $N_i$ .

$$S(\mathbf{x}) = \sum_i \sum_f \frac{\|\mathbf{x}_i - \mathbf{x}(\alpha_{if}^0, \beta_{if}^0, h_{if}^0)\|^2}{N_i} \quad (10)$$

We define the potential energy of the shape constrained model as the summation of the error  $e(\mathbf{x}_i)$  in data fitting at the mesh vertices. The constrained energy formulation is then given by Eq. (11). We solve the energy minimisation problem by steepest descent, equivalent to a zero mass dynamic system. Equation (12) gives the update equation for the vertices of the shape constrained model.

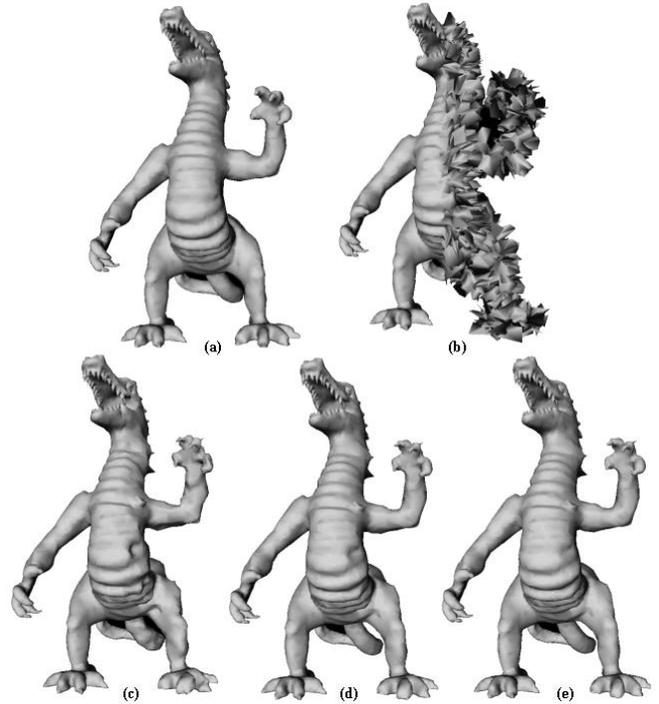
$$E(\mathbf{x}) = \sum_i e(\mathbf{x}_i) + \sum_i \sum_f \frac{\|\mathbf{x}_i - \mathbf{x}(\alpha_{if}^0, \beta_{if}^0, h_{if}^0)\|^2}{N_i} \quad (11)$$

$$\frac{d\mathbf{x}_i}{dt} = -\nabla e(\mathbf{x}_i) - \sum_f \frac{(\mathbf{x}_i - \mathbf{x}(\alpha_{if}^0, \beta_{if}^0, h_{if}^0))}{N_i} \quad (12)$$

The model evolves to minimise the error in fitting to the data, also minimising the deviation in the local shape and parameterisation of the model. Figure (6) illustrates the evolution of the shape constrained deformable model with no data fitting term  $e(\mathbf{x}_i)$ , in recovering the shape of an irregular triangular mesh with added noise. This demonstrates that the shape constrained model converges to the original shape even after severe distortion.

#### 4.2 Model to Surface Matching

The shape constrained deformable model is matched to the body data in order to deform the model to fit the shape of the body. Matching is a fundamental and well studied problem in computer vision in both 2D and 3D. A standard approach is to use a nearest neighbour assignment such as iterative closest point (ICP) [9]. However, this will lead to distortion in the matching according to the initial alignment of the model with the data. The approach has been extended by considering an all-neighbours assignment [22]. Matching then becomes inexact, where the weighting of the assignment to each neighbour defines the degree of certainty in the assignment. Here we use the robust point matching technique introduced by Rangarajan [38] where an all-neighbours assignment is combined with a coarse to fine refinement of the assignment through a deterministic annealing framework. We combine the framework with the deformation of the shape constrained model to obtain a coarse to fine match of the model to the data.



**Fig. 6.** Recovery of shape on an irregular triangular mesh (a) Original model (b) Model with noise (c),(d),(e) Convergence to original shape.

We define the error metric in fitting the data as the least-squared error between the model and the data set. The potential energy function is given by Eq. (13) where  $\mathbf{x}_i$  spans the set of model vertices and  $\mathbf{y}_j$  spans the set of scan data points.

$$P(\mathbf{x}) = \sum_{i=1}^{i=I} \sum_{j=1}^{j=J} m_{ij} \|\mathbf{y}_j - \mathbf{x}_i\|^2 \quad (13)$$

The all-neighbour assignment is defined by a matrix  $M$  of continuous match variables  $m_{ij}$  subject to the following constraints: a data point  $\mathbf{y}_j$  can match one model vertex  $m_{ij} = 1$ , partially match several vertices  $\sum_i m_{ij} \leq 1$  or may not match any vertices  $m_{ij} = 0$  for all  $i$ , and a model vertex  $\mathbf{x}_i$  must be assigned to the data set  $\sum_j m_{ij} = 1$ . The inequality constraint is transformed to an equality constraint by augmenting the match matrix with a set of slack variables, Eq. (14).

$$0 \leq m_{ij} \leq 1 \quad \sum_{i=1}^{i=I+1} m_{ij} = 1 \quad \sum_{j=1}^{j=J} m_{ij} = 1 \quad (14)$$

A deterministic annealing framework is used for the assignment by adding an entropy term to the energy function [38]. The temperature  $T$  of the entropy term defines the degree of fuzziness in the assignment, Eq. (15).

$$P(\mathbf{x}) = \sum_i \sum_j m_{ij} \|\mathbf{y}_j - \mathbf{x}_i\|^2 + T \sum_i \sum_j m_{ij} (\log(m_{ij}) - 1) \quad (15)$$

The final energy equation for the shape constrained deformable model is given by Eq. (16), consisting of the shape constraint, the data fit and the entropy term.

$$\begin{aligned}
 E(\mathbf{x}) = & \sum_i \sum_f \frac{\|\mathbf{x}_i - \mathbf{x}(\alpha_{if}^0, \beta_{if}^0, h_{if}^0)\|^2}{N_i} \\
 & + \sum_i \sum_j m_{ij} \|\mathbf{y}_j - \mathbf{x}_i\|^2 \\
 & + T \sum_i \sum_j m_{ij} (\lg(m_{ij}) - 1)
 \end{aligned} \quad (16)$$

For a fixed temperature  $T$  and fixed model configuration  $\mathbf{x}_i$  the match parameters  $m_{ij}$  can be derived by differentiating the energy function with respect to  $m_{ij}$ , Eq. (17).

$$\begin{aligned}
 \frac{dE}{dm_{ij}} = & \|\mathbf{y}_j - \mathbf{x}_i\|^2 + T \lg(m_{ij}) = 0 \\
 m_{ij} = & \exp\left(-\frac{\|\mathbf{y}_j - \mathbf{x}_i\|^2}{T}\right)
 \end{aligned} \quad (17)$$

The Sinkhorn balancing technique of alternating row and column normalisation of the match matrix  $M$  [11] can then be performed to satisfy the match constraints in Eq. (14). Equation (17) shows that a vertex assignment is weighted according to the relative distance to a surface point  $\|\mathbf{y}_j - \mathbf{x}_i\|$  with the temperature  $T$  defining the effective region of matching in space. As the temperature tends to zero this method approaches the ICP algorithm.

For a fixed assignment we can derive the gradient of the energy function with respect to the mesh vertices  $\mathbf{x}_i$ , Eq. (18), giving a gradient descent solution for the evolution of the mesh.

$$\frac{d\mathbf{x}_i}{dt} = \sum_j m_{ij} (\mathbf{y}_j - \mathbf{x}_i) - \sum_f \frac{(\mathbf{x}_i - \mathbf{x}(\alpha_{if}^0, \beta_{if}^0, h_{if}^0))}{N_i} \quad (18)$$

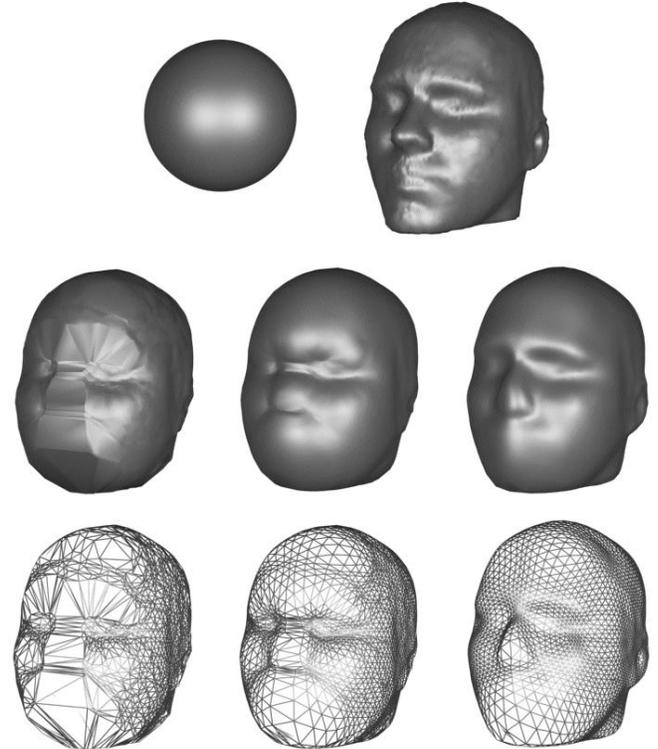
We fit the deformable model to the data-set by alternately updating the assignment to the 3D points  $m_{ij}$  using Eq. (17) and satisfying the constraints in Eq. (14), then updating the vertex positions of the model  $\mathbf{x}_i$  using Eq. (18). The temperature is reduced as the model is updated to refine the assignment, giving a coarse to fine approach to matching. We start at an initial temperature  $T_{init}$ . The match parameters are initialised according to Eq. (17), with the slack matches set to the minimum match  $m_{ij}$  for each data point  $j$ . The parameters are then balanced to satisfy the constraints. The model is updated using explicit Euler integration steps until convergence. The convergence criteria is set as the maximum component of the gradient falling within the current region of matching,  $\{\|\frac{dE}{d\mathbf{x}_i}\|\} \leq \sqrt{T}$ . We then alternately reduce the temperature, redefine the match parameters, and then update the model until a final temperature  $T_{final}$  is reached.

In practice the data set in a human scan can be large. A processed scan can consist of around  $10^5$  data points, with around  $10^3$  vertices in the control model there are in the order of  $10^8$  match parameters. Many of the parameters are redundant for each vertex, where the distance to

the corresponding data point exceeds the matching temperature and the match parameter is effectively zero. We therefore restrict the match for each vertex to a closest subset of points, reducing the number of match parameters to be updated and stored. An octree representation of the data set is constructed for efficient retrieval of closest points. In matching we also only retrieve points whose surface normal lies in the same half-plane as the vertex normal in order to avoid incorrect surface matches.

The initial temperature  $T_{init}$  is automatically defined as the the maximum squared error for the nearest neighbour of the initial model vertices, giving a region of matching equivalent to the worst-case nearest neighbour assignment. The final temperature  $T_{final}$  is set as the maximum desired error tolerance in the final model. We set  $T_{final} = 0.001^2$  corresponding to a tolerance of 1mm.

Figure (7) illustrates the performance of the proposed technique in comparison to standard ICP matching for the problem of mapping a regularly parameterised sphere to a human head data-set. Figure (7a) shows the match for ICP, (7b) shows ICP combined with the proposed local shape constraint and Fig. (7c) the constrained multi-point matching technique introduced in this paper. This demonstrates that the use of a coarse to fine, multi-point match, combined with the shape constraint improves the surface match over standard ICP.



(a) ICP match (b) Constrained ICP (c) Constrained multi-point  
**Fig. 7.** Fitting of uniformly triangulated sphere to a head surface.

## 5 Representation of Surface Detail

The model registration and shape fitting described deforms the control model to produce a close approximation to the shape of the scan data-set. The control model can be animated using the predefined animation structure of the generic model. We then accurately and efficiently represent the surface detail by constructing a displacement map. The displacement map represents the offset between the control model and the scan data across the surface of the control model. In previous work [19] the normal-volume was introduced to derive a continuous mapping between 3-space and a polygonal mesh. Here we briefly describe the details of the normal-volume mapping, displacement map generation and detail reconstruction on the control model.

We wish to describe surface points  $\mathbf{p}^H$  on the high-resolution scan surface in terms of an indexed triangle on the control model and barycentric coordinates  $(\alpha, \beta)$  which give a point  $\mathbf{p}^C(\alpha, \beta)$  and a displacement  $d$  along the interpolated normal  $\mathbf{n}^C(\alpha, \beta)$ , as in Eq. (19).

$$\mathbf{p}^H = \mathbf{p}^C(\alpha, \beta) + d\mathbf{n}^C(\alpha, \beta). \quad (19)$$

To do this we define the normal volume of a control triangle which is given by offsetting each control triangle vertex  $(\mathbf{v}_r, \mathbf{v}_s, \mathbf{v}_t)$  along its normal  $(\mathbf{n}_r, \mathbf{n}_s, \mathbf{n}_t)$  by distances  $\pm d$ , Fig. (8a). The resulting union of all normal volumes gives a continuous volumetric envelope which encloses the high-resolution scanned data. This ensures that any detailed point within this normal-volume can be mapped to the control mesh.

We can write the control point  $\mathbf{p}^C$  and its normal  $\mathbf{n}^C$  by a bi-linear interpolation of the control mesh vertices as given in Eq. (20). This interpolation ensures a continuous normal across the whole of the control surface. This both allows each point to be uniquely mapped, Fig. (8b), and also (as we shall see in section (6)) ensures that, once the surface is animated, the detailed reconstruction will deform smoothly.

$$\begin{aligned} \mathbf{p}^C(\alpha, \beta) &= \alpha\mathbf{v}_r + \beta\mathbf{v}_s + (1 - \alpha - \beta)\mathbf{v}_t \\ \mathbf{n}^C(\alpha, \beta) &= \alpha\mathbf{n}_r + \beta\mathbf{n}_s + (1 - \alpha - \beta)\mathbf{n}_t. \end{aligned} \quad (20)$$

The detailed mesh is parameterised with respect to the control mesh by solving Eq. (19) so that for each scan data point  $\mathbf{p}_i^H$  we obtain a displacement  $d$ , its corresponding control triangle and barycentric coordinates  $(\alpha, \beta)$ . The equation is first solved for  $d$  by obtaining the plane with normal  $(\mathbf{v}_t - \mathbf{v}_s) \times (\mathbf{v}_s - \mathbf{v}_r)$  in the normal volume which passes through  $\mathbf{p}_i^H$  as well as a displaced control vertex. Once  $d$  is known the barycentrics are given by Eq. (19).

We compute a displacement image by defining texture coordinates  $\mathbf{u}$  which uniquely map each control triangle to a texture plane. If a control triangle has texture coordinates  $(\mathbf{u}_r, \mathbf{u}_s, \mathbf{u}_t)$  then this mapping is defined by

$$\mathbf{u} = \alpha\mathbf{u}_r + \beta\mathbf{u}_s + (1 - \alpha - \beta)\mathbf{u}_t \quad (21)$$

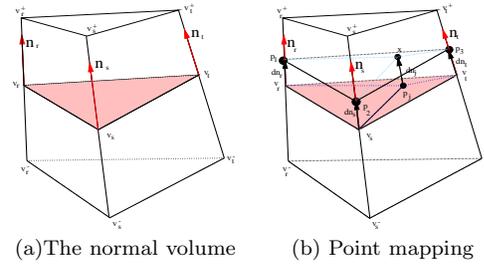


Fig. 8. Normal volume mapping.

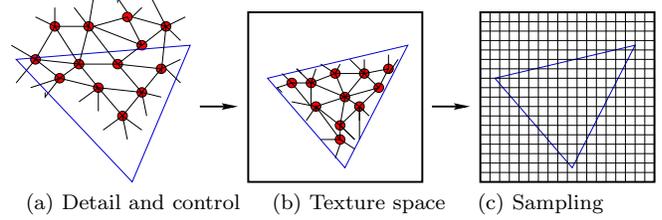


Fig. 9. Displacement map generation.

Figure (9) shows how we sample the texture image at regular intervals and record displacements by first mapping to barycentric coordinates given by Eq. (21) and then interpolating the solutions to Eq. (19) for each detailed triangle vertex.

Figure (10) shows the mapping stored as an image. Here we use the model pelting technique introduced by Piponi and Borshukov [37] for texture map generation to produce a single displacement map for our generic control model. This provides an intuitive representation of detail for mesh editing.

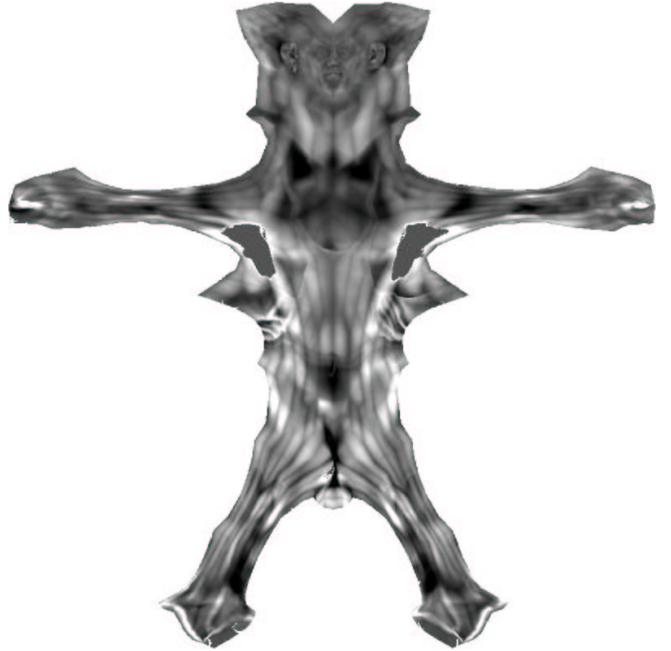


Fig. 10. A displacement map generated for a cyberware human scan, white indicates max +ve displacement, black max -ve displacement.

The detailed mesh can be reconstructed by subdividing the control mesh and displacing the new vertices.

Subdivision is performed using edge split operations in which each triangle has 3 new vertices inserted at the centre of each edge to produce four new triangles. The new vertices are displaced along the continuous normal by calculating the texture coordinates and sampling the displacement image at this point.

## 6 Results

The presented framework for the recovery of layered animation models has been tested on full body data-sets captured using the Cyberware laser scanning system [3]. We would like to thank Cyberware for the provision of 4 human data-sets and the Stanford Computer Graphics Laboratory for provision of Michelangelo’s David from the Digital Michelangelo Project [32]. The David data-set has been edited here to remove the plinth and the connection between the hands and the body.

### 6.1 Reconstruction

The high-resolution surface scans are efficiently represented by a displacement map on the control model. We assess the representation by assessing the error in reconstructing the original surface data-sets for each model, using the recovered model pose for each data-set. We measure the mean and the root mean square (RMS) error from the reconstructed model surface  $S_1$  to the original surface  $S_2$ . If  $d(\mathbf{x}, S_1)$  is defined as the minimum distance from a point  $\mathbf{x}$  on the surface  $S_1$  to the surface  $S_2$  then the mean distance from  $S_1$  to  $S_2$  is given by Eq. (22), and RMS distance is given by Eq. (23). For triangulated meshes this distance can be calculated using the Metro tool [12].

$$d(S_1, S_2)_{mean} = \frac{1}{area(S_1)} \int d(\mathbf{x}, S_2) d\mathbf{x} \quad (22)$$

$$d(S_1, S_2)_{RMS} = \left( \frac{1}{area(S_1)} \int d(\mathbf{x}, S_2)^2 d\mathbf{x} \right)^{\frac{1}{2}} \quad (23)$$

Table (1) gives the RMS error and Table (2) the mean error in reconstructing the original data-sets. Michelangelo’s David has been scaled here to a height of 2m for comparison with the Cyberware models. Mean error converges to the quantisation level in the displacement map. RMS error however is affected by outliers in the reconstruction and does not converge to the quantisation level.

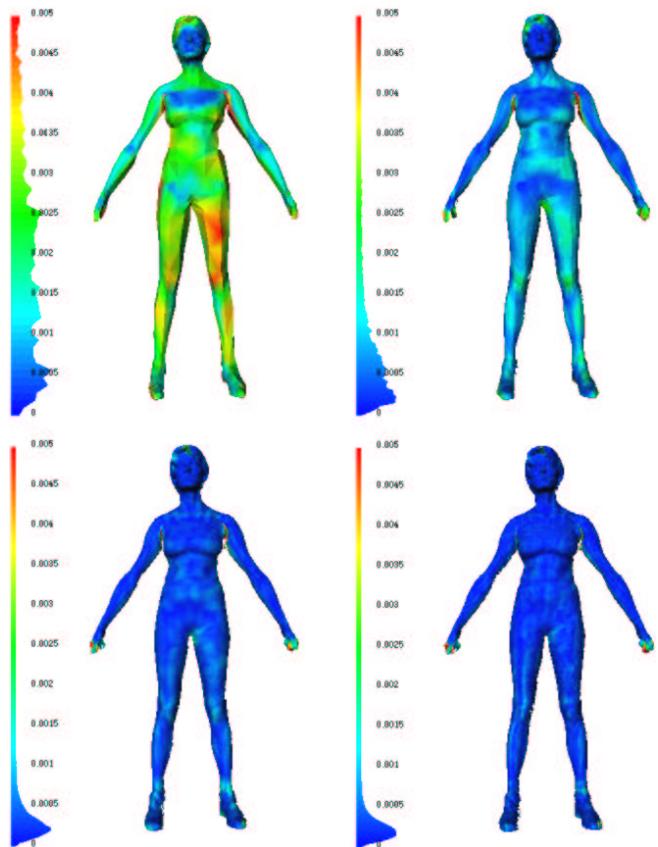
Outliers are caused by noninjectivity in the displacement map as described in section 2.2. Although the control mesh accurately fits the majority of the high-resolution surface, there are outliers where the fitting fails and the detail mesh is folded with respect to an unfolded control mesh. Where the surface is folded detail cannot be represented by a scalar displacement map. Figure 11 shows that errors occur at complex features of the scanned mesh (hands) or at folded areas (armpits). The problem may be countered by using a higher resolution control model and registering the control model with additional feature points on the complex areas of the body.

	#Vertices	Model1	Model2	Model3	Model4	David
L0	1056	3.67	3.82	5.02	4.50	5.97
L1	4216	2.27	2.27	4.06	3.54	4.73
L2	16840	1.93	1.97	3.89	3.43	4.47
L3	67344	1.90	1.93	3.91	3.42	4.52
L4	269336	1.90	1.93	3.91	3.43	4.55
e		0.48	0.47	0.40	0.38	1.46

**Table 1.** Root mean square error (/mm) to the original surface data for the reconstructed models at 0,1,2,3,4 uniform levels of subdivision (L), compared to the quantisation error in the displacement map (e).

	#Vertices	Model1	Model2	Model3	Model4	David
L0	1056	2.16	2.25	2.57	2.46	3.51
L1	4216	0.98	1.01	1.27	1.19	2.14
L2	16840	0.62	0.63	0.87	0.81	1.67
L3	67344	0.53	0.54	0.78	0.72	1.56
L4	269336	0.52	0.52	0.76	0.71	1.54
e		0.48	0.47	0.40	0.38	1.46

**Table 2.** Mean error (/mm) to the original surface data for the reconstructed models at 0,1,2,3,4 uniform levels of subdivision (L), compared to the quantisation error in the displacement map (e)



**Fig. 11.** Error distribution (/m) on reconstructed Cyberware model at 0,1,2, and 3 levels of uniform subdivision. Error figures generated using Metro.

## 6.2 Animation

Animation using the layered model structure is demonstrated in Fig. (12) for a Cyberware human scan. The control skeleton is driven from motion capture data and the control surface is smoothly animated. The continuous normal defined on the control model ensures that as the control mesh deforms, the reconstructed detail deforms smoothly. Animation from the complete motion capture sequence can be viewed at [2].

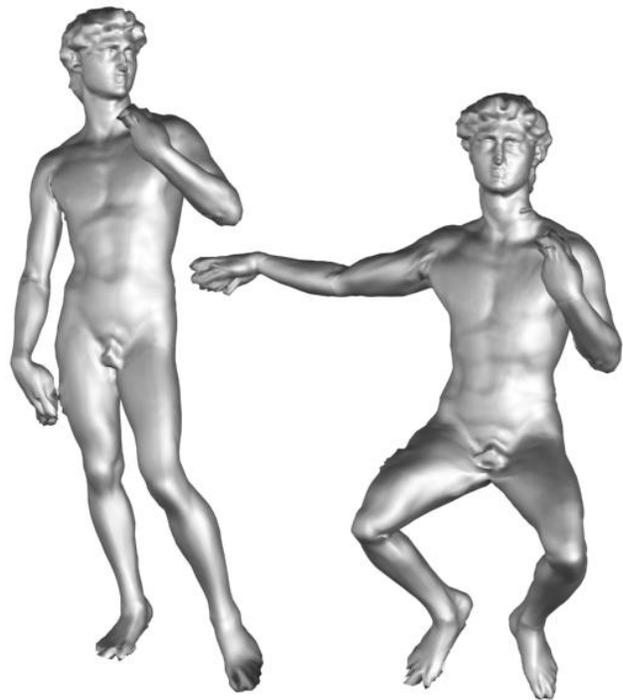


**Fig. 12.** Frames from an animation, original model courtesy of Cyberware. Animation can be viewed at <http://www.ee.surrey.ac.uk/Research/VSSP/3DVision/Animatedstatues/>

We also demonstrate the animation of Michelangelo's David, Fig. (13). The complex pose and detail in this data-set represent a demanding task in generating the layered model for animation. At the feet and hands the scan surface is folded with respect to the control model and the reconstruction cannot match the detail of the original scan. The reconstructed model in the original pose is shown in Fig. (13a). At the left elbow, the body pose produces a complex surface deformation that cannot be matched by the vertex-weighting scheme adopted in this work. We find that at the left elbow the control model becomes folded with respect to the scan surface. We therefore cannot animate the left elbow without distortion from incorrect surface mapping in the displacement map generated. Animation of the rest of the body is shown in Fig. (13b).

## 6.3 Level-of-Detail Control

The displacement map representation provides several desirable features for control of surface detail. Figure (14) shows the result of a simple edit of the displacement image with a standard image editing tool. Figure



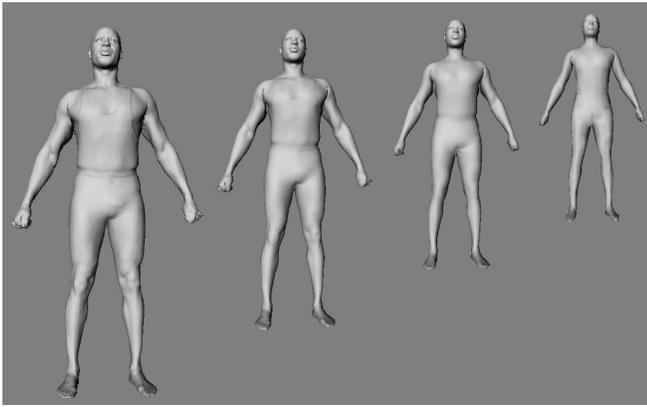
(a) Reconstructed model (b) Reconstructed animated model

**Fig. 13.** Animation of Michelangelo's David, original model courtesy of Stanford Computer Graphics Laboratory.

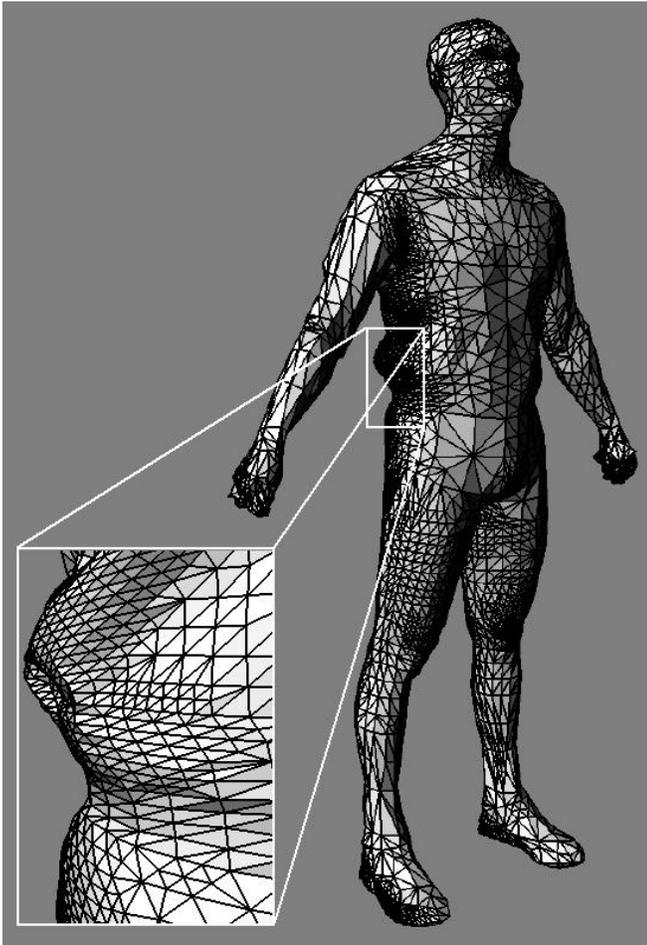
(15) gives a demonstration of static level of detail. The models which are further away require less levels of subdivision to give the impression of the same amount of detail. Adaptive level of detail is also possible and Fig. (16) shows the result of a nonuniform subdivision algorithm which subdivides at areas of greater displacement.



**Fig. 14.** Editing the displacement map, original model courtesy of Cyberware.



**Fig. 15.** Reconstructed model at different levels of detail (L3,L2,L1,L0), original model courtesy of Cyberware.



**Fig. 16.** Non-uniform subdivision, original model courtesy of Cyberware.

**7 Conclusions**

We have presented a framework for animating high-resolution human surface data captured from commercially available 3D active sensing technology. We use a model-based approach, matching a generic humanoid control model to the dense surface data to provide a control model for animation of the surface. We register the generic model with the data-set using a set of manu-

ally defined feature points and joint locations to recover the pose of the model. The model is then automatically matched to the data as a shape constrained deformable surface, preserving the original shape and parameterisation of the generic model. The shape constrained matching ensures robust matching in the presence of noise and missing data. Preservation of the mesh parameterisation ensures that the deformed control model can be animated using the predefined animation structure of the generic model.

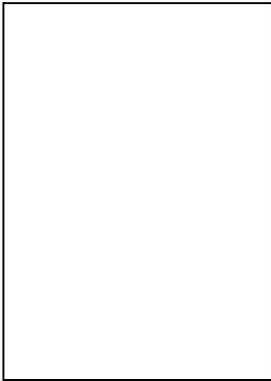
The technique is based on a generic humanoid control model that can be optimised for model animation. Surface detail is represented as a displacement map from the low-resolution surface of the control model, providing an efficient representation of the detailed geometry of the model and allowing control over surface detail. Here we use a single displacement map generated as a human pelt [37] giving an intuitive representation of the surface detail for editing. This approach provides a robust mechanism for animating whole-body scan data with minimal manual intervention. Work remains to provide our control model with a detailed surface deformation scheme to simulate surface dynamics and to address non-injective mapping of the high-resolution surface.

**10**

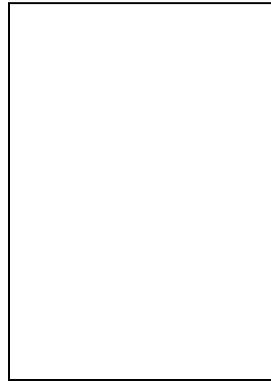
1. *3D-MATIC Laboratory*. <http://www.faraday.gla.ac.uk/>.
2. *Animated Statues*. <http://www.ee.surrey.ac.uk/Research/VSSP/3DVision/Animatedstatues/AnimatedStatues.html>.
3. *Cyberware Inc.* <http://www.cyberware.com/>.
4. *Wicks and Wilson Limited*. <http://www.wwl.co.uk/>.
5. *The Humanoid Animation Specification, Humanoid Animation Working Group (H-Anim)*. <http://www.h-anim.org/>, 2001.
6. T. Akimoto, Y. Suenaga, and R.S. Wallace. Automatic creation of 3D facial models. *IEEE Computer Graphics and Applications*, 13:16–22, 1993.
7. C. Barron and I.A. Kakadiaris. Estimating anthropometry and pose from a single image. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 669–676, 2000.
8. G. Bertrand. A parallel thinning algorithm for medial surfaces. *Pattern Recognition Letters*, 16:979–986, 1995.
9. P.J. Besl and N.D. McKay. A method of registration of 3-d shapes. *IEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):239–255, 1992.
10. B. Buxton, L. Dekker, I. Douros, and T. Vassilev. Reconstruction and Interpretation of 3D Whole Body Surface Images. In *Scanning 2000*, Paris, May 2000.
11. H. Chui and A. Rangarajan. A New Algorithm for Non-Rigid Point Matching. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 2, pages 44–51, 2000.
12. P. Cignoni, C. Rocchini, and R. Scopigno. Metro: measuring error on simplified surfaces. *Computer Graphics Forum*, 17(2):167–174, 1998.
13. T.H. Cootes, A. Hill, C. Taylor, and J. Haslam. The use of active shape models for locating structures in medical images. *Image Vision Computing*, 12(6):355–366, 1994.
14. N. Gagvani, D. Kenchammana-Hosekote, and D. Silver. Volume animation using the skeleton tree. In *IEEE Symposium on Volume Visualization*, pages 47–54, October 1998.
15. F.S. Grassia. Practical parameterization of rotations using the exponential map. *The Journal of Graphics Tools*, 3(3), 1998.

16. M.R. Grosso, R. Quach, E. Otani, J. Zhao, S. Wei, P.H. Ho, J. Lu, and N.I. Badler. *Anthropometry for computer graphics human figures. Technical Report MS-CIS-89-71*. University of Pennsylvania, Dept. of Computer and Information Science, Philadelphia, PA, 1989.
17. J. Hamill and M.K. Katheen. *Biomechanical Basis of Human Movement*. Williams and Witkins, 1995.
18. A. Hilton, D. Beresford, T. Gentils, R. Smith, W. Sun, and J. Illingworth. Whole-body modelling of people from multi-view images to populate virtual worlds. *The Visual Computer*, 16(7):411–436, 2000.
19. A. Hilton and J. Illingworth. Multi-resolution geometric fusion. In *International Conference on Recent Advances in 3D Digital Imaging and Modeling*, pages 181–188. IEEE, 1997.
20. H.H.S Ip and L. Yin. Constructing a 3d individualized head model from two orthogonal views. *The Visual Computer*, 12:254–266, 1996.
21. X. Ju and J.P Siebert. Conforming generic animatable models to 3d scanned data. In *Numerisation 3D - Scanning 2001*, May 2001.
22. I.A. Kakadiaris and D. Metaxas. Three-dimensional human body model acquisition from multiple views. *International Journal of Computer Vision*, 30(3):191–218, 1998.
23. E-A Karabassi, G. Papaioannou, and T. Theoharis. A fast depth-buffer-based voxelization algorithm. *Journal of Graphics Tools*, 4(4):1999, 5-10.
24. L. Kobbelt, S. Campagna, J. Vorsatz, and H.P. Seidel. Interactive Multi-Resolution Modeling on Arbitrary Meshes. In *SIGGRAPH Conference Proceedings*, pages 105–114, August 1998.
25. R. (Ed.) Koenen. *Coding of Moving Pictures and Audio: MPEG 4 Standard ISO/IEC JTC1/SC29/WG11 N4030*. <http://mpeg.telecomitalia.com/standards/mpeg-4/mpeg-4.htm>, 2001.
26. V. Krishnamurthy and M. Levoy. Fitting smooth surfaces to dense polygon meshes. pages 313–324, August 1996.
27. T. Kurihara and K. Arai. *A Transformation Method for Modeling and Animation of the Human Face from Photographs*. Springer-Verlag, Berlin, 1991.
28. A. Lee, H. Moreton, and H. Hoppe. Displaced subdivision surfaces. In *SIGGRAPH Conference Proceedings*, pages 85–94. ACM SIGGRAPH, July 2000.
29. J. Lee and S.Y. Shin. A hierarchical approach to interactive motion editing for human-like figures. In *SIGGRAPH Conference Proceedings*, pages 39–48, 1999.
30. W. Lee, J. Gu, and N. Magnenat-Thalmann. Generating animatable 3D virtual humans from photographs. In *EUROGRAPHICS*, volume 19(3), 2000.
31. Y. Lee, D. Terzopoulos, and K. Waters. Realistic modeling for facial animation. In *SIGGRAPH Conference Proceedings*. ACM SIGGRAPH, 1995.
32. M. Levoy, K. Pulli, B. Curless, S. Rusinkiewicz, D. Koller, L. Pereira, M. Ginzton, S. Anderson, J. Davis, J. Ginsberg, J. Shade, and D. Fulk. The digital michelangelo project: 3d scanning of large statues. In *SIGGRAPH Conference Proceedings*. ACM SIGGRAPH, 2000.
33. J.P. Lewis, M. Cordner, and N. Fong. Pose space deformation: A unified approach to shape interpolation and skeleton-driven deformation. In *SIGGRAPH Conference Proceedings*. ACM SIGGRAPH, 2000.
34. T. McInerney and D. Terzopoulos. Deformable models in medical image analysis: a survey. *Medical Image Analysis*, 1(2):91–108, 1996.
35. J. Montagnat and H. Delingette. Volumetric Medical Images Segmentation using Shape Constrained Deformable Models. In *First Joint Conference CVRMed-MRCAS*, volume Lecture Notes in Computer Science 1205, pages 13–22, March 1997.
36. F. Pighin, J. Hecker, D. Lischinski, R. Szeliski, and D.H. Salesin. Synthesizing realistic facial expressions from photographs. In *SIGGRAPH Conference Proceedings*. ACM SIGGRAPH, 1998.
37. D. Piponi and G. Borshukov. Seamless texture mapping of subdivision surfaces by model pelting and texture blending. In *SIGGRAPH Conference Proceedings*, pages 471–477. ACM SIGGRAPH, 2000.
38. A. Rangarajan, H. Chui, E. Mjolsness, S. Pappu, L. Davachi, P. Goldman-Rakic, and J. Duncan. A robust point matching algorithm for autoradiograph alignment. *Medical Image Analysis*, 4(1):379–398, 1997.
39. R. Smith, W. Sun, A. Hilton, and J. Illingworth. Layered animation using displacement maps. In *IEEE International Conference on Computer Animation*, pages 146–154, Philadelphia, Pennsylvania, May 2000.
40. L.H. Staib and J.S. Duncan. Model-Based Deformable Surface Finding for Medical Images. *IEEE Transactions on Medical Imaging*, 15(5):720–731, 1996.
41. W. Sun, A. Hilton, R. Smith, and J. Illingworth. Layered animation of captured data. *The Visual Computer*, 17:457–474, 2001.
42. D. Terzopoulos. From physics-based representation to functional modeling of highly complex objects. In *NSF-ARPA Workshop on Object Representation in Computer Vision*, pages 347–359. Springer Verlag, 1994.
43. D. Terzopoulos and D. Metaxas. Dynamic 3D models with local and global deformations: deformable superquadrics. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13:703–714, 1991.
44. D. Terzopoulos, A. Witkin, and M. Kass. Constraints on Deformable Models: Recovering Shape and Nonrigid Motion. *Artificial Intelligence*, 36(1):91–123, 1988.
45. G. Turk and J.F O'Brien. Shape transformation using variational implicit functions. In *SIGGRAPH Conference Proceedings*, pages 335–342. ACM SIGGRAPH, August 1999.
46. L. Wade and R.E. Parent. Fast, fully-automated generation of control skeletons for use in animation. In *Computer Animation*, pages 164–169, Philadelphia, Pennsylvania, May 2000.
47. C. Zhu, R. H. Byrd, and J. Nocedal. L-BFGS-B: Algorithm 778: L-BFGS-B, FORTRAN routines for large scale bound constrained optimization. *ACM Transactions on Mathematical Software*, 23(4):550–560, 1997.

JONATHAN STARCK received B.A. and M.Eng degrees in Engineering from Cambridge University, UK in 1996 and a MSc in Medical Physics and Clinical Engineering from Sheffield University, UK in 1997. He trained as a clinical scientist in the National Health Service, UK from 1996 to 2000. He is now working toward a Ph.D. as a research fellow at the University of Surrey on the capture of photo-realistic models of people.

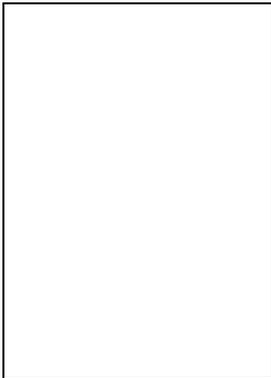


GORDON COLLINS graduated with a BSc. in Mathematics and Philosophy from Leeds University in 1993 and with an MSc. in Nonlinear Mathematics from Edinburgh University in 1994. He completed his PhD in Numerical Analysis at Bristol University in 1998. Since returning to academia from the world of banking, he has lectured in applied mathematics at Bristol University and been a postdoctoral research fellow at University of Surrey. His interests are in surface parameterisation, adaptivity and physics based deformation.

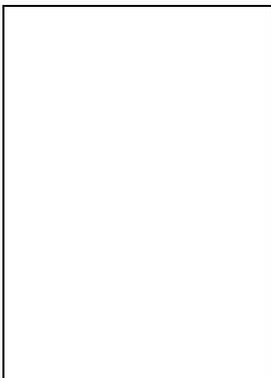


JOHN ILLINGWORTH BSc, DPhil, FIEE. Professor Illingworth has been active in image processing, computer vision and pattern recognition for 20 years and is co-author of nearly 150 conference and journal papers. He is a past Chairman of the British Machine Vision Association and is editor of IEE Proceedings on Vision, Image and Signal Processing.

This article was processed by the author using the  $\text{\LaTeX}$  style file *cljour2* from Springer-Verlag.



RAYMOND SMITH graduated from the University of Surrey in 1998 with a Master of Engineering degree in Information Systems Engineering. Since October 1998, he has been a Ph.D. Student in the Centre for Vision, Speech and Signal Processing at the University of Surrey, performing research into the automatic creation of 3D models suitable for animation. He also has interests in VRML, human modeling and animation.



ADRIAN HILTON received a B.Sc.(Hons.) in Mechanical Engineering and a D.Phil. degree from the University of Sussex in 1988 and 1992 respectively. In 1992 he joined the CVSSP at the University of Surrey where he has worked for the past decade on capture of shape, appearance and movement of real objects for realistic computer graphics and accurate reconstruction. He has published over fifty refereed articles and received Best Paper awards from the journal Pattern Recognition in 1996 and IEE Electronics and Communications in 1999. Industrial collaboration has resulted in the first hand-held 3D capture system and first booth for 3D capture of people which received two EU IT Awards for Innovation. He is currently leading a research team with the goal of developing computer vision technologies for visual media production in film and broadcast.

Industrial collaboration has resulted in the first hand-held 3D capture system and first booth for 3D capture of people which received two EU IT Awards for Innovation. He is currently leading a research team with the goal of developing computer vision technologies for visual media production in film and broadcast.