

Building Superquadric Men from 3D Whole-Body Scan Data

Paper id: JICS-SI-paper#6

Abstract – *This paper describes a scheme for reconstructing 3D human models from whole-body scan data. The 3D human models are composed of superquadric primitives each corresponding to a functional body part. The reconstruction starts with primary segmentation of scan data, which defines the location of the articulation link in the human scan and results in topologically simple (primary) body parts. The primary body parts are further decomposed into sub-parts suitable for superquadric representation. These segmented sub-parts are refined to achieve global compatibility and then fitted with superquadrics to complete our process. Being based on compact volumetric primitives, the whole reconstruction scheme is able to deal directly with raw scan data comprising 3D point clouds sampled on human body surfaces. Both qualitative and quantitative evaluations of the method are conducted.*

Keywords: 3D shape modelling, human body scanning, superquadric fitting, segmentation, reconstruction, articulated models.

1 Introduction

The recent advances in 3D imaging technology have enabled us to acquire dense 3D point clouds sampled on whole human body surfaces (human body scan data) with reasonable measurement accuracy and fast capture time [1]. Ever since the emergence of such technology, the need for quantitative analysis of human body shape has been continuously increasing, with applications ranging from anthropometrical research [2], clothing design [3], to virtual human animation [4]. However, due to the large diversity of human body shapes and postures and the ill-conditioned properties of scan data, processing and modelling of human scan data in those applications have proved a very challenging problem.

Methods for human scan data processing and modelling previously reported in the literature mainly fall in two categories: surface reconstruction and surface parameterization. In the surface reconstruction, a surface representation of the human body shape is recovered from the point cloud sampling of the human body [5-8]; in the surface parameterization, the recovered surface

representation is regularized to have the right semantic for its parameter domain [9-12]. However, due to the great complexity of the problem, the previously-reported methods have been designed under severe constraints. First of all, in order to reduce the body shape variation caused by postures, a standard anthropometry posture (standing upright with arms held at sides) has to be assumed [5, 6, 12]. The standard posture makes the shape space of human body simplified and therefore reduces the complexity in characterising human body shapes. Secondly, many approaches require strict data properties, e.g., complete body surface coverage [10, 11], detailed specification of body scanning equipments [5, 8, 9], etc. How to systematically generalize those researches is still an open issue. Finally, the use of manual intervention [7, 11] is usually involved in human body scan data processing and modelling, which seriously limits the throughput and efficiency.

The objective of this research is to relax the constraints of the previous methods for 3D scan data analysis, which have narrowed the range of tasks to which 3D human scan data can be applied. To this end, a general human scanning scenario is assumed so that no constraints on subject postures and data properties will be imposed. Under such assumptions, we have devised a fully automatic scheme for reconstructing human body shapes directly from 3D scan data up to a gross level of structure. The human models employed are composed of *superquadrics*, which are compact volumetric primitives that can represent a large variety of shapes with just a few parameters. The reconstructed superquadric human models alone can be directly utilised in applications where only gross human geometry is required such as computer games, virtual environments, manikin manufacturing, etc. Moreover, while providing a simple parameterization of the human body and having the potential to be extended to represent surface details, the superquadric human models could serve as an intermediate step in the computation of realistic human shapes.

The reminder of the paper is organized as follows. Section 2 describes the superquadric model we use to reconstruct human body shape. Section 3 discusses the reconstruction scheme. Some experimental results are presented in Section 4, followed by conclusions in Section 5.

2 Superquadric human model

The superquadric is a well-known part-level model in the field of computer vision and graphics. Being an extension of the quadric surface, the superquadric incorporates two shape control parameters ε_1 and ε_2 to adjust the curvature of its surface (see Eq (1)). Therefore, a superquadric is able to represent a much larger range of shapes than the traditional quadric surface. The implicit form of the superquadric is defined as follows.

$$f(x, y, z) = \left(\left(\frac{x}{a_1} \right)^{\frac{2}{\varepsilon_2}} + \left(\frac{y}{a_2} \right)^{\frac{2}{\varepsilon_2}} \right)^{\frac{\varepsilon_2}{\varepsilon_1}} + \left(\frac{z}{a_3} \right)^{\frac{2}{\varepsilon_1}} = 1 \quad (1)$$

The shape of the superquadric (1-isosurface of the function f) is controlled by parameters ε_1 and ε_2 . When $\varepsilon_1, \varepsilon_2$ vary, the shape smoothly changes. In the special case of $\varepsilon_1 = \varepsilon_2 = 1$, the superquadric degenerates to a common ellipsoid. a_1, a_2, a_3 are scaling factors on x -, y -, z - axis respectively.

While the implicit form of the superquadric is convenient for the tasks of fitting and collision detection, the superquadric surface can also be represented in an explicit way using spherical product [11], which not only provides a direct method for evaluating the superquadric surface but also produces a simple parameterisation for the surface that can effectively facilitate some further operations on the surface such as surface registration and displacement construction.

The representational strengths of the superquadric make it suitable for representing blob-like shapes [13, 14], especially shapes formed by natural processes [15]. Hypothesising that human body parts belong to this category, we propose to employ superquadrics to build up a human body model, with each superquadric representing a functional body part. The superquadric itself is symmetric along its three principal axes. Therefore, we further extend the superquadrics formulation with a *tapering deformation* [16] applied to the dominant principal axis of each superquadric, which allows for scale change along this axis. Such scale change is evident when human body parts are “thick” and “thin” at different areas along their axes. The formulation of the tapering deformation is as follows.

$$\mathbf{X}_D = T(\mathbf{X}) \Leftrightarrow \begin{cases} x_D = \left(\frac{k_x}{a_3} z + 1 \right) x \\ y_D = \left(\frac{k_y}{a_3} z + 1 \right) y \\ z_D = z \end{cases} \quad (2)$$

In Eq. (2), $\mathbf{X} = [x, y, z]$ denotes a point on the original superquadric surface and $\mathbf{X}_D = [x_D, y_D, z_D]$ represents a point on the deformed superquadric surface.

The whole superquadric human model comprises of 15 superquadrics each corresponding to a functional body part and a kinematic link between body parts. The 15 functional body parts include 1 head, 1 upper torso, 1 lower torso, 2 upper arms, 2 lower arms, 2 upper legs, 2 lower legs, 2 hands, and 2 feet. The partitioning of the body parts is a compromise between accuracy and efficiency, as the resolution of whole-body scan data is normally not sufficient to capture the detailed shape of finer body parts such as fingers and toes. The kinematics of the functional body parts, are expressed as rigid transformations between the parts and a reference frame called world coordinate frame.

$$\mathbf{X}_w = K(\mathbf{X}_D) = \mathbf{R}\mathbf{X}_D + \mathbf{T} \quad (3)$$

Note that Eqs (1), (2), and (3) are all differentiable, which endows the proposed human model with strong analytic properties which facilitate robust recovery of a stable model from the data (see Section 3.1).

It is worth mentioning that similar superquadric-based human models have been reported in the literature, mostly for tracking human figures in image sequences [17-19], in which superquadrics are applied to represent the shape of the human figures. However, the context of human tracking is significantly different from that of this research. The objective of human tracking is to resolve motion information of human figures from 2D image data, while the objective of this research is to capture human shapes from 3D point clouds. Therefore techniques involved for human tracking can not be applied directly in this research.

3 Reconstruction of superquadric men

3.1 Part fitting

Given data points sampled on a body part, we fit a superquadric to this data set and thereby represent the body part by a corresponding superquadric primitive and its associated parameters. Due to the differentiability of the proposed superquadric models, it is possible to fit the models analytically to the underlying data. Employing the algebraic distance proposed by Solina [20], superquadric fitting can be formulated as follows:

$$\min_{\mathbf{P}} (a_1 a_2 a_3) \sum_j (f^{\varepsilon_1}(T^{-1}(K^{-1}(\mathbf{X}_w^j))) - 1)^2 \quad (4)$$

Here \mathbf{X}_w^i represents the i -th data point and \mathbf{P} is a vector comprising all parameters of the superquadric including its shape parameters, scaling parameters, tapering parameters and kinematic parameters. Because f , T and

K are all differentiable, the fitness function to be minimised (4) is also differentiable and this property allows us to utilise derivative-driven methods to optimise the minimisation process (4). Specifically, we chose the Levenberg-Marquardt approach [21] because it is designed specifically to optimise a fitness function in the form of square sum. The shape control parameters are subject to boundary conditions of: $0.5 \leq \varepsilon_1 \leq 1.5, 0.5 \leq \varepsilon_2 \leq 1.5$ to avoid over-fitting during optimisation. Fitting progresses iteratively and in each iteration step the parameter vector \mathbf{P} is updated using the following equation:

$$[\mathbf{J}(\mathbf{P})^T \mathbf{J}(\mathbf{P}) + \lambda \mathbf{I}] \delta \mathbf{P} = -\mathbf{G}(\mathbf{P}) \quad (5)$$

Since both the Jacobian matrix $\mathbf{J}(\mathbf{P})$ and the gradient vector $\mathbf{G}(\mathbf{P})$ in Eq. (5) can be computed analytically, a close-form optimisation solution can be found.

The iterative fitting process requires that the parameters in \mathbf{P} be initialised. To achieve, on average, a fast convergence rate we initialise $\varepsilon_1, \varepsilon_2$ to their mid-range values, i.e., $\varepsilon_1 = \varepsilon_2 = 1$ and the tapering parameters are set initially to induce no tapering deformation, i.e., $k_x = k_y = 0$. The rotation, translation and scaling parameters are initialised using Principal Component Analysis (PCA) which fits an ellipsoid to the 3D point cloud data based on statistical criteria [22]. The data distributions projected on the principal axes of the ellipsoid recovered by PCA are assumed to be Gaussian and their associated standard deviations are used to initialise the superquadrics scaling parameters a_1, a_2, a_3 as follows:

$$a_1 = 3\sigma_x; a_2 = 3\sigma_y; a_3 = 3\sigma_z \quad (6)$$

where $\sigma_x, \sigma_y, \sigma_z$ denote standard deviations on the principal axes (denoted as x, y, z -axes) respectively. The rigid transformation parameters (rotation and translation) of the superquadric are initially computed by aligning the axes of the superquadric to the principal axes of the ellipsoid fitted to the data using PCA.

3.2 Body shape decomposition

Fitting a superquadric to a specific body part requires that the data points corresponding to the body part under analysis be segmented out of the whole scan dataset. This task is known as body shape decomposition and proves difficult when large body form variations and posture change are involved and when the scan data itself is deficient (as described in section 1). We tackled this problem in three stages. In the first stage a coarse segmentation (called primary segmentation) is applied to decompose the whole body data into primary body parts, i.e., torso and four limbs. The second stage comprises a finer segmentation (called secondary segmentation) using a ‘‘fit-and-segment’’ scheme, which iteratively fits

and segments the data to obtain a more detailed shape by adding more superquadric primitives into the representation. In the third stage, the superquadric primitives are relaxed and re-fitted to achieve global compatibility, while the corresponding segmentation is attained simultaneously.

3.2.1 Primary segmentation

The coarse segmentation involves a topology-based decomposition technique on point sets [23]. The whole dataset is first processed by a connectivity test where a connectivity graph is built upon the data points. Geodesic distances between data points are then computed using wavefront-propagation. Computing the Morse function based on Geodesic distance, the whole dataset is quantised into different levels (see Fig. 1(a)).

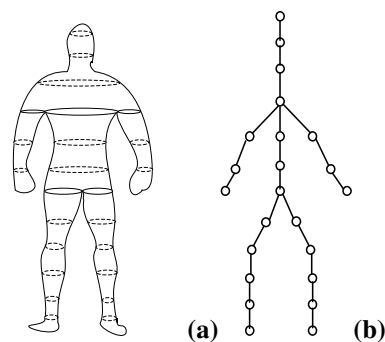


Fig. 1 Primary segmentation : (a) Human scan and level-set curves ; (b) Discrete Reeb graph

The data points in each level are grouped into discrete curves based on their connectivity. These discrete curves are linked between adjacent levels if there are connecting points between the curves, resulting in a discrete Reeb graph (Fig. 1(b)) where the discrete curves are represented by nodes. This set of nodes is then decomposed into subsets by detecting nodes corresponding to human joints, which is achieved by analysing singularities in the discrete Reeb graph and then applying a heuristic criterion that captures basic human typography. The data points contained in the different subsets of nodes approximately correspond to the 5 primary body parts, i.e., 1 torso (including head) and 4 limbs (see Fig. 4(a)).

3.2.2 Secondary segmentation

A primary body part, obtained from the topology-based segmentation, still consists of several functional body parts, which can be extracted using a ‘‘fit-and-split’’ algorithm. A single superquadric is first fitted to the primary body part, and a *consistency measure* is then computed along z -axis of the superquadric (see Fig. 2). The consistency at any z -level is a weighted combination of the normalized area difference and the centre location difference between the data and the superquadric defined as follows:

$$\xi = w_A(A_s - A_d)/A_s + w_C|\mathbf{C}_s - \mathbf{C}_d|/\sqrt{a^2 + b^2} \quad (7)$$

where ξ represents consistency, A_s , A_d denote areas and \mathbf{C}_s , \mathbf{C}_d denote the centres of the superquadric slice and the data slice respectively, a and b are the semimajor and semiminor axes of the superquadric slice (see Fig. 2), w_A , w_C are weights for terms representing area difference and centre difference. Clearly, if the superquadric is fitted well to the data, the value of ξ should be small.

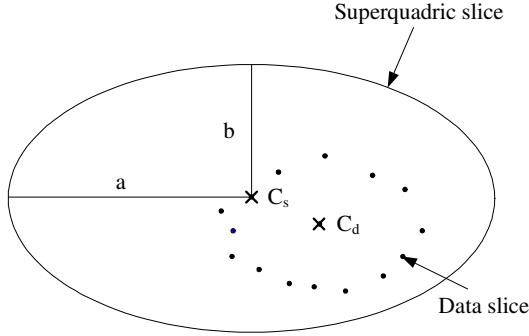


Fig. 2 Consistency between data and superquadric

A primary body part is split based on the detection of *salient points* which lie on a *consistency graph* defined by plotting the consistency function against the z-axis distance of the body part superquadric. Salient points are detected as zero-crossings of first order and second order derivatives of the consistency graph, which correspond to salient changes on this plot. Fig. 3 demonstrates a primary body part (an arm), its fitted superquadric and the consistency graph between the data and the superquadric. Two salient points are labelled on the consistency graph corresponding to the wrist and elbow areas on the arm.

Having detected the salient points, the data is split at the point that minimizes the fitting error of the bisected data (e.g., the salient point corresponding to elbow in Fig. 3). This splitting process is applied twice (the second split is applied to the data segment produced by the first splitting operation that has the larger fitting error) to obtain three segments of a primary body part.

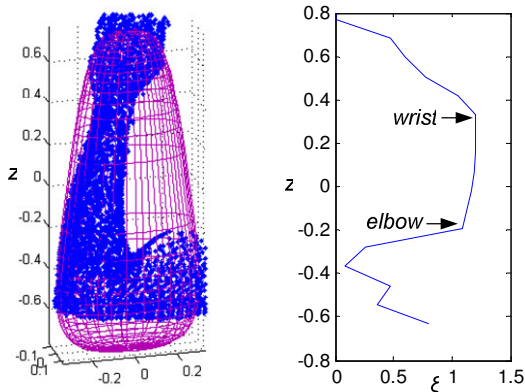


Fig. 3 Primary body part, fitted superquadric and consistency graph

3.3 Part refining

Since the fine segmentation process described in Section 3.2.2 is based on local analysis of primary body parts, it is difficult to attain global compatibility amongst all of the segments used to represent the whole body. We have devised a relaxation labelling technique to tackle this problem. The underlying idea is to adjust the labels of data points associated with potentially corresponding superquadrics such that global compatibility is improved. First, we assign each possible label a probability:

$$P(\mathbf{x}_j : S_i) = \frac{1/\delta(\mathbf{x}_j, S_i)}{\sum_i 1/\delta(\mathbf{x}_j, S_i)} \quad (8)$$

Here \mathbf{x}_j denotes a data point, S_i represents a superquadric primitive, and $\delta(\mathbf{x}_j, S_i)$ means the radial distance between \mathbf{x}_j and S_i . Such a definition of probability satisfies the condition $\sum_i P(\mathbf{x}_j : S_i) = 1$, and

guarantees that the label for each data point has the higher probability associated with the closer superquadric. We then define a compatibility coefficient for each possible labelling of each pair of data points:

$$c((\mathbf{x}_j : S_i), (\mathbf{x}_l : S_k)) = 1 - d(S_i, S_k) / \max_{i,k} (d(S_i, S_k)) \quad (9)$$

Here $d(S_i, S_k)$ is the Euclidean distance between centres of superquadrics S_i, S_k . Using this definition, we can calculate the compatibility support for a possible labelling of a data point from all data points:

$$Q(\mathbf{x}_j : S_i) = \sum_l w_{jl} \sum_k c((\mathbf{x}_j : S_i), (\mathbf{x}_l : S_k)) P(\mathbf{x}_l : S_k) \quad (10)$$

Here w_{jl} are weights that reflect the relative importance of data point pairs in supplying each other a context. We simply assign a value of 1.0 to the weights of data points that are considered as neighbours ($d(\mathbf{x}_j, \mathbf{x}_l) < 3\sigma$, σ is the resolution of scan data) and a value of 0.0 otherwise. The relaxation labelling progresses iteratively. In each iterative step, probabilities of data points associated with superquadrics are updated as follows:

$$P^{s+1}(\mathbf{x}_j : S_i) = \frac{P^s(\mathbf{x}_j : S_i) Q^s(\mathbf{x}_j, S_i)}{\sum_k P^s(\mathbf{x}_j : S_k) Q^s(\mathbf{x}_j, S_k)} \quad (11)$$

The above formula assigns a new probability to each label by taking into account compatibility support from all data points, thereby providing a global adjustment mechanism for the labelling. Upon termination of the iteration process, each data point is

assigned with a label that has the highest probability, amongst all possible labels, for that data point. The iteration length is mainly justified by observation: we discovered that the label probabilities exhibit no significant change within in a few iteration steps. Therefore we chose 3-10 steps as the normal range of iterations in the relaxation labelling process.

4 Experiments

In this section, we first use an example to demonstrate the whole body scan segmentation and fitting scheme, and then evaluate the quality of superquadric human models achieved using a group of scan datasets.




Figure 4. An example of segmentation and fitting : (a) primary segmentation ; (b) secondary segmentation ; (c) relaxation labelling ; (d) Superquadric fitting

Fig. 4 illustrates data output from different stages of the scheme. Fig. 4 (a) is the result achieved by the primary segmentation, where the whole dataset has been clearly segmented into 5 primary parts, although the segmentation boundaries are not strictly faithful to those found in human anatomy. These primary parts are further decomposed, as shown in Fig. 4 (b), where basically 15 segments of body have been recovered. Because of the influence of the primary segmentation, the locations of the segments are not adequately aligned to the true body parts. The relaxation labelling process has improved the segmentation quality, as shown in Fig. 4 (c), by making the segments more compatible with each other. The final fitting result is displayed in Fig. 4 (d), which has captured an explicit “summary” of the main shape of the human body as implied in the scan data.




Fig. 5 Segmentation and fitting results of different scan samples

In order to understand the quality of superquadric human model achieved by our proposed segmentation and fitting scheme, we processed a group of real scan datasets, which were acquired from different types of scanner, representing different human bodies in a variety of typical postures. In Fig. 5, scan samples (a,b,c,d) were acquired from Cyberware body scanners [24], which are based on laser scanning technology. The samples comprises of two women (a,c) and two men (b,d), who were in two different postures, i.e., standing and sitting. Scan samples (e,f,g,h) are all men, but in three different postures, acquired from Wicks and Wilson body scanners, which are different from Cyberware scanners, adopting Moiré fringing 3D imaging principles. Despite all of the differences (shape, imaging approach, resolution and posture) in the samples in Fig. 5, the datasets have been segmented reasonably well and superquadric human models have been obtained to represent principal shape of the bodies. Almost all the samples have been decomposed into 15 body parts, except for samples (b,e), in which the arms were not segmented into upper parts and lower parts, because the arms are in straight pose and exhibit no detectable salient change in the elbow areas.

Basically the reconstructed superquadric models have captured the gross shape of the human bodies as judged from their visual appearance. Since ground truth knowledge of the human body shapes is not available,

we are unable to evaluate rigorously how well the models represent the body shapes. However, we believe that the magnitude of the fitting errors between the data points and the models provides a good indicator of the quality of the reconstructed models. Table 1 lists four statistics (the mean, maximum, minimum values and standard deviations) of the fitting errors for all the models shown in Fig. 5. The radial distances between the data points and the models have been adopted to represent fitting errors, as radial distance (Euclidean distance along a radius) affords a more intuitive interpretation, based on its physical meaning, than the algebraic distance used commonly in minimisation (4).

Table 1 Fitting errors of superquadric human models in Fig. 5 (unit: meter)

Model	mean	maximum	minimum	std
(a)	0.0117	0.0622	5.197×10^{-7}	0.0102
(b)	0.0101	0.0853	1.620×10^{-6}	0.0090
(c)	0.0187	0.1089	1.587×10^{-7}	0.0180
(d)	0.0121	0.0621	2.229×10^{-8}	0.0102
(e)	0.0116	0.1046	5.990×10^{-7}	0.0101
(f)	0.0116	0.0840	3.664×10^{-7}	0.0106
(g)	0.0118	0.0850	2.817×10^{-7}	0.0099
(h)	0.0115	0.0724	5.603×10^{-7}	0.0095

The statistics in Table 1 suggest that superquadrics can represent human body shapes reasonably well. The average distance between the data point and the fitted superquadric human model is of the order of 0.01 meters. To allow fair comparison, each dataset has been normalised to 2 meters high on its longest axis. The fitting error standard deviations also remain relatively low, implying that we have achieved reasonable fitting consistency between the reconstructed superquadric models and the corresponding scan data. The error levels indicated by the data in Table 1 are tolerable in applications for which the coarse shape of the human body is sufficient, such as in virtual environments and real-time computer games. Moreover, the above statistics exhibit no significant change over the data samples, despite their considerable differences in characteristics (shape, pose, resolution, etc.). This observation suggests that the methods adopted in this work should be sufficiently stable to support the applications cited above.

5 Conclusion

This paper reports a scheme to reconstruct analytical human models from human 3D scan data. Superquadric primitives are proposed to model human

shape, which compress the high dimensional human shape captured with point-cloud data into the much lower dimensional space of superquadric model parameters. The compact nature and analytic properties exhibited by our representation confer advantages in applications such as 3D human shape retrieval, real-time computer games, virtual environment, object recognition, human body tracking, etc. Moreover, the simple parameterisation of the superquadric model can potentially facilitate further analysis of the human shape, e.g., body part registration. Finally, the reconstruction scheme itself is fully automatic and able to deal with human scan point clouds in arbitrary postures, thus offering a prospect of complete machine-interpretation of whole-body human shapes in the future.

In addition to proposing the concept of recovering an articulated parts-based building superquadric human model from 3D scan data, this paper also makes the following technical contributions: the part splitting algorithm based on a consistency graph is believed to be novel, and can be used to segment other cylinder-like objects with only slight adaptation of splitting strategy. Moreover, the relaxation labelling technique to improve compatibility of individual superquadrics is also new in the field of superquadric fitting.

Since a displacement map can be attached to single superquadric [11], we propose to add the original surface details on top of the superquadric human models. Moreover, we believe that the superquadric-based segmentation and fitting scheme can be extended further to allow other articulated objects to be represented and reconstructed from 3D scan data.

Acknowledgement

The authors wish to acknowledge funding from the UK Engineering and Physical Sciences Research Council and the support of the Imaging Faraday Partnership.

References

- [1] J. P. Siebert and S. J. Marshall, "Human body 3D imaging by speckle texture projection photogrammetry," *Sensor Review*, vol. 20, pp. 218-226, 2000.
- [2] P. R. M. Jones and M. Rioux, "Three dimensional surface anthropometry: applications to human body," *Optics and Lasers in Engineering*, vol. 28, pp. 89-117, 1997.
- [3] R. P. Pargas, N. J. Staples, and J. S. Davis, "Automatic measurement extraction for apparel from a three-dimensional body scan," *Optics and Lasers in Engineering*, vol. 28, pp. 157-172, 1997.
- [4] J. Starck, G. Collins, R. Smith, A. Hilton, and J. Illingworth, "Animated statues," *Journal of Machine Vision Applications*, pp. 248-259, 2003.

- [5] I. Douros, L. Dekker, and B. F. Buxton, "Reconstruction of the surface of the human body from 3D scanner data using B-splines," presented at SPIE conference on Three-Dimensional Image Capture and Applications II, 1999.
- [6] J. H. Nurre, "Locating landmarks on human body scan data," presented at International Conference on Recent Advances in 3D Digital Imaging and Modeling, NJ USA, 1997.
- [7] P. R. M. Jones, P. Li, K. Brook-Wavell, and G. M. West, "Format of Human Body Modelling from 3-D Body Scanning," *International Journal of Clothing Science*, vol. 7, pp. 7-16, 1995.
- [8] B. Buxton, L. Dekker, I. Douros, and T. Vassilev, "Reconstruction and Interpretation of 3D Whole Body Surface Images," presented at Scanning 2000, Paris, 2000.
- [9] L. Dekker, I. Douros, B. F. Buxton, and P. Treleaven, "Building Symbolic Information for 3D Human Body Modeling from Range Data," presented at International Conference on 3-D Digital Imaging and Modeling, Ottawa, Canada, 1999.
- [10] X. Ju and J. P. Siebert, "Individualising Human Animation Models," presented at Eurographics, Manchester, UK, 2001.
- [11] B. Allen, B. Curless, and Z. Popovic, "The space of human body shapes: reconstruction and parameterization from range scans," presented at SIGGRAPH, 2003.
- [12] X. Ju, N. Werghi, and J. P. Siebert, "Automatic Segmentation of 3D Human Body Scans," presented at IASTED International Conference on Computer Graphics and Imaging 2000, Las Vegas, USA, 2000.
- [13] A. H. Barr, "Superquadrics and angle-preserving transformations," *IEEE Computer graphics and Applications*, vol. 1, pp. 11-23, 1981.
- [14] C. W. Chen, T. S. Huang, and M. Arrot, "Modelling, analysis, and visualization of left ventricle shape and motion by hierarchical decomposition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 16, pp. 342-356, 1994.
- [15] A. Pentland, "Perceptual Organization and the representation of natural form," *Artificial Intelligence*, vol. 28, pp. 293-331, 1986.
- [16] A. H. Barr, "Global and Local deformations of solid primitives," *Computer Graphics*, vol. 18, pp. 21-30, 1984.
- [17] C. Sminchisescu and B. Triggs, "Mapping Minima and Transitions in Visual Models," *International Journal of Computer Vision*, vol. 61, 2005.
- [18] C. Sminchisescu and A. Telea, "Human Pose Estimation From Silhouettes: A Consistent Approach Using Distance Level Sets," presented at WSCG International Conference on Computer Graphics, Visualization and Computer Vision, 2002.
- [19] C. Sminchisescu, D. Metaxas, and S. Dickinson, "Improving the Scope of Deformable Model Shape and Motion Estimation," presented at International Conference on Computer Vision and Pattern Recognition, 2001.
- [20] F. Solina and R. Rajcsy, "Recovery of parametric models from range images: the case of superquadric with global deformations," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 12, pp. 131-147, 1990.
- [21] W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling, *Numerical Recipes in C : The Art of Scientific Computing*: Cambridge University Press, 1986.
- [22] I. T. Jolliffe, *Principal Components Analysis*. New York: Springer-Verlag, 1986.
- [23] Y. Xiao, J. P. Siebert, and N. Werghi, "Topological Segmentation of Discrete Human Body Shapes in Different Postures Based on Geodesic Distance," presented at International Conference on Pattern Recognition, Cambridge, UK, 2004.
- [24] www.cyberware.com