

Analysis of Segmented Human Body Scans

Pengcheng Xi
National Research Council of Canada
Ottawa, Ontario, Canada
pengcheng.xi@nrc-cnrc.gc.ca

Won-Sook Lee
University of Ottawa
Ottawa, Ontario, Canada
wslee@uottawa.ca

Chang Shu
National Research Council of Canada
Ottawa, Ontario, Canada
chang.shu@nrc-cnrc.gc.ca

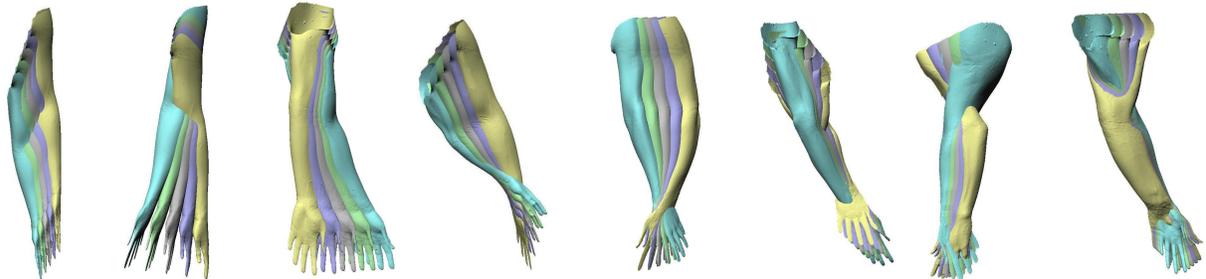


Figure 1: Variations on first eight principal components of left arm (along each component, five different weights are selected to create these variations; different views of these eight images are used for a better observance of the variations)

ABSTRACT

Analysis on a dataset of 3D scanned surfaces have presented problems because of incompleteness on the surfaces and because of variances in shape, size and pose. In this paper, a high-resolution generic model is aligned to data in the Civilian American and European Surface Anthropometry Resources (CAESAR) database in order to obtain a consistent parameterization. A Radial Basis Function (RBF) network is built for rough deformation by using landmark information from the generic model, anatomical landmarks provided by CAESAR dataset and virtual landmarks created automatically for geometric deformation. Fine mapping then successfully applies a weighted sum of errors on both surface data and the smoothness of deformation. Compared with previous methods, our approach makes robust alignment in a higher efficiency. This consistent parameterization also makes it possible for Principal Components Analysis (PCA) on the whole body as well as human body segments. Our analysis on segmented bodies displays a richer variation than that of the whole body. This analysis indicates that a wider application of human body reconstruction with segments is possible in computer animation.

CR Categories: I.4.6 [Computing Methodologies]: Image Processing and Computer Vision—Segmentation; I.6.4 [Computing Methodologies]: Simulation and Modeling—Model Validation and Analysis

Keywords: CAESAR, Radial Basis Function, consistent parameterization, principal components analysis, segmentation

1 INTRODUCTION

During the past two centuries, anthropometry has been applied to measuring human bodies for analysis of body variations. Traditional measurement instruments include tape (for circumference and length), spreading caliper (for head breadth), anthropometer,

sliding caliper, large skinfold caliper, scale and so on. These instruments were only able to record some of the human body information.

Today, anthropometry is playing an important role in industrial design, clothing and vehicle designs. Distributions of body dimensions are important information in these areas for designing better products for people. Take face mask as an example, where designers need to know a distribution of the curvatures and surface features for certain population. This cannot be measured with traditional instruments and thus poses higher requirements for better approaches of recording human bodies. One possible solution is 3D scanning technology. But processing noisy and incomplete surfaces in the raw dataset takes a lot of computation and no general methodologies are available for processing different objects. Another problem is that scanning results can tell the geometry of disparate subjects but not of a uniform space. So the scanning does not really produce a plausible model. Our assumption is that if we could build a morphable model on a set of data and combine this with statistical approaches, the subjects created by the model would be more interesting.

The difficulties in building a morphable model of this kind are three-fold: first, each dataset has a different number of vertices; second, to be consistent, a model needs to be able to match the features on different data; and third, the model has to fill the holes on the scanned surface that result from grazing angles during scanning.

A representative work on building a morphable model is from [3] where the authors applied a multi-resolution approach to conduct first rough deformation and then fine alignment. We found it possible to speed up the rough deformation and make the fine alignment easier for almost the same-quality result.

This paper introduces a two-step alignment framework that combines rough transformation using the RBF network as the first step with fine mapping as the second step. The RBF network helps align the high-resolution generic model to the target in a very efficient way and removes the markers' error from the second step. Simplified error definition makes easier the control over the selection of weights and parameters. Combining these two steps produces a high degree of efficiency. It also proves to be robust through the alignment of hundreds of targets. The starting point was the Civilian American and European Surface Anthropometry Resource

Graphics Interface Conference 2007
28-30 May, Montréal, Canada
Copyright held by authors. Permission granted to
CHCCS/SCDHM to publish in print form, and ACM
to publish electronically.



(CAESAR) project that collected thousands of range scans of volunteers aged from 18 to 65 in Europe and the United States.

The good alignment of all the training data selected from the CAESAR database makes it possible to conduct statistical analysis. A principal components analysis on all the aligned whole bodies present principal variations in the training set. To construct more variable human bodies, we also analyze the possibility of segmented reconstructions by looking into the variations for each body segment.

The rest of the paper is organized as follows. The next section reviews the relevant literature on the practical problems of aligning different scans and of repairing holes on scanned surfaces. Analyzing only currently available approaches, we introduce our two-step approach to alignment in Section 3. In Section 4, PCA analysis on whole body and segments are discussed. Our conclusion is presented in Section 5.

2 RELATED WORK

In this section, we focus on the literature of modeling body shape variation, finding mutually consistent surface representations, filling holes in scanned data, and registering non-rigid surfaces.

A representative work on modeling the variation of human body shape is by DeCarlo et al. [7]. The authors described a system capable of automatically generating distinct, plausible face geometries. They introduced two steps in constructing a face: first, the generation of a random set of measurements that characterize the face, and second, the construction of the best surface that satisfies the geometric constraints imposed by those measurements. The output surface minimizes a measure of fairness, which formalizes how much the surface bends and stretches away from the kind of shape that faces normally have. This is basically an early work on anthropometrical statistics for likely face measurements, which are later used for constraints on a parameterized surface.

Blanz and Vetter's work in [22] is a milestone for human modeling and reconstruction. It explores two key problems of computer-aided face modeling. First, new face images or new 3D face models can be registered automatically by making a large number of comparisons with an internal model. Second, the approach regulates the naturalness of the faces modeled and avoids those with an "unlikely" appearance. This work successfully includes statistical approaches into computer graphics as a way of obtaining visual information from a given database.

Also in [22], Blanz and Vetter aligned the features using a modified version of 2D optical flow. This is because their head scans have a cylindrical parameterization. The system also enables users to create new characters and to modify facial attributes by varying the model's coefficients. Within the constraints imposed by prior probability, the range of possible faces varies greatly, and all linear combinations look natural.

Whole-body parameterizations are more complicated because the human body does not easily lend itself to cylindrical mapping [3]. Praun et al. [9] therefore introduced an approach that builds an n -way correspondence between meshes of the same topological type by the use of feature markers. This gives immediate correspondences between models and allows remeshes.

The approach from [9], however, can only apply to complete surfaces. In datasets where some surface data is missing, the parameterization will not work well. Also, since this parameterization relies greatly on landmarks, problems arise when markers are missing. A latest work from [8] relies heavily on a large number of landmarks for alignment in the first stage and on thin-plate splines for warping non-landmark vertices.

Some work on filling holes on the scanned surface has been done. In [12], the authors started with constructing a signed distance function, the zero set of which is defined only in the vicinity of observed

surfaces. A diffusion process then extends this function through the volume until its zero set bridges any holes that may be present. Other hole-free reconstruction methods, such as [14] and [19], have the desirable feature of filling holes smoothly. These methods are good for areas like the top of the head; but for the soles of the feet, the approach fails to make them stable. In addition, not all the approaches are able to fill the holes between two legs.

In [2], Allen and his colleagues started from a subdivision template that resembles the range surface. They re-parameterized the surface by sampling it along the template normals to construct a set of displacement maps. They were then able to fill in displacements smoothly. Their work defines smoothness relative to the template surface. The soles of the feet, for example, will be filled in flat. A related displacement-mapped approach was also developed in [1]. But a pre-condition for avoiding the crossing of sample rays is that the template surface is already a fairly close match to the original [1]. The work that follows applies generic models to alignment as well as to filling holes. One of the latest contributions comes from Kahler et al. on the "Head Shop" [17], which presents a versatile construction and deformation method for head models with anatomical structure, suitable for real-time physics-based facial animation. The paper employs this technique to fit a generic head model to imperfect scan data, and to simulate head growth from early childhood to adult age. Their deformation uses volumetric radial basis functions. Though these approaches work well for convex objects, volumetric deformations are not suitable for entire bodies.

Anguelov et al. [5] presented a framework for determining complex shape models from range scan data. Their framework consists of algorithms based on the theory of probabilistic graphical models. This allows complex shape models of different objects and object classes to be learned with minimal human intervention.

RBFs for elastic alignment was first used in [6] where intermediate objects are constructed given two or more objects of general topology. Another RBF-based approach is from [16], where dense surface correspondences are computed by volume morphing with RBF followed by a cylindrical projection. Volume morphing roughly aligns features of the two models such as eye sockets, nose ridge, lip corners, and chin points.

Seo and Magnenat-Thalmann [11] designed a system that is close to Thomas Vetter's face system [22]. Starting from 3D data, they derived by RBF the deformation functions that generate the appropriate shape and proportion of body geometry [11]. They applied a skeleton-driven approach to the rigid deformation of a generic model to target body surfaces.

Hitherto, the generic model approach has been commonly used for human body modeling. In [23], Lee et al. successfully applied a similar method to face modeling and animation.

The matching technique used in [3] is based on an energy-minimization framework, and is similar to that of Marschner et al. [21]. Instead of using the surface smoothness term in [21], the authors in [3] tried to minimize variation of the deformation itself, so that the holes in the mesh are filled in with detail from the template surface. This keeps local warping stable so that the whole surface gets little deformation in local areas.

In an earlier work, Feldmar and Ayache [13] described a registration technique with matching surface points, normals, and curvature that maintain a similar affine transformation within spherical regions of space. The authors in [3] have the same idea by using "locally affine deformations". The difference is that they do not use surface normals or curvature, but directly define the smoothness term over the surface rather than over a spherical volume [13].

In this paper, we apply RBF, and revise the definition of weighted errors, from previous work, and try to combine them to achieve same-quality alignment results with a higher efficiency in a template-based approach.



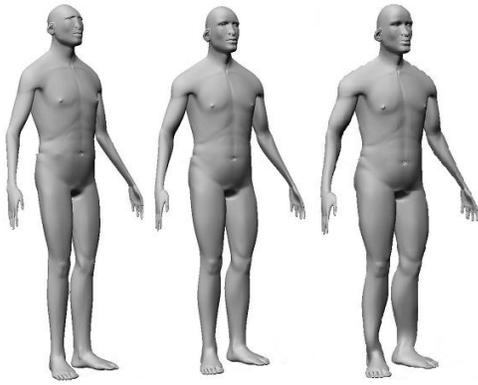


Figure 4: Deformed generic models with a multi-quadrics radial function (from left to right, the value of s selected is 0, 0.005 and 0.01)

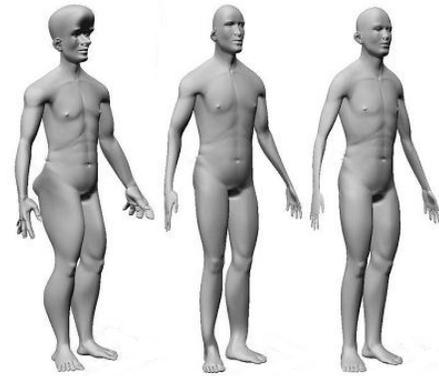


Figure 6: Deformations with different radial basis functions (from left to right: thin-plate, multi-quadrics and Gaussian)

The third radial function to be introduced is a Gaussian function:

$$\Phi_j(r) = \exp(-r^2/c^2) \quad (5)$$

where c is the only parameter. Experimental results on different values for c are demonstrated in Figure 5. It is found that while increasing the value of c , deformation on pose is attenuated and the deformed surface becomes coarser. This tendency could be explained by the definition of radial basis function. With the increment of c value, the radial function also increases correspondingly. The increment of the function value means there is a long “Gaussian” distance between current vertex and the influencing landmarks, which results in that the deformation effect of these landmarks will be attenuated.

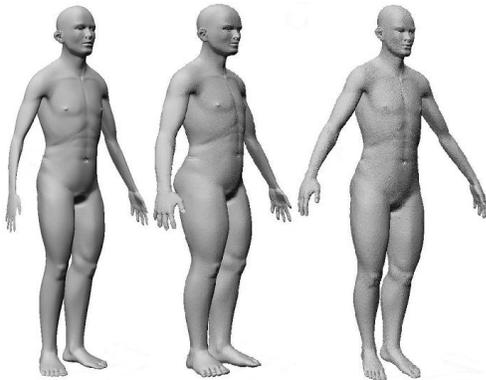


Figure 5: Deformed generic models with a Gaussian radial function (from left to right, the value of c selected is 1.0, 10.0 and 100.0)

By comparison of these three radial basis functions in Figure 6, the Gaussian function with the parameter value of “1.0” is finally selected for the project. One main reason is that Gaussian function has the best performance in maintaining human-like shape during the rough deformation; however, based on the RBF-transformed generic models with another two radial basis functions, there will be more artifacts on the fine-mapped result.

After RBF transformation, the generic model has been deformed closer to the target model in shape, pose, and height (see comparisons in Figure 7 for difference before and after RBF-deformation). RBF ensures a good registration of landmarks in that all landmarks have been mapped to corresponding targeting positions. Therefore, in the second step, the fine mapping, marker information is discarded and only data and smoothness errors are counted.

RBF has several practical advantages. First, if the original generic model is significantly different from the target, the RBF can provide translation, rotation and scaling for the rough alignment in an efficient way. Second, the normal information of most vertices can be discarded because the transformed surface is already close to the target. Experimental results will be shown to prove that RBF moves the generic model faster than the marker error did in [3].

3.2 Fine alignment

After the rough transformation, we only need to concentrate on a greatly reduced distance between the transformed generic model and the target. It must be emphasized that a still high difference between the two models will result in a local minimum in the fine mapping process. By discarding markers’ errors, the error definitions in [3] are simplified to two: (a) the difference between the source surface and target surface (data error) and (b) the difference between transformations on neighboring vertices (smoothness error).

Suppose the transformed generic model is U , and the scanned target surface is Γ . For each vertex on U , we define a 4×4 transformation matrix T_i . Each vertex has 12 degrees of freedom to define the transformation. Our intent is to find a set of transformations that move each vertex to its closest point on Γ . This process needs to be repeated several times in order to align the RBF-transformed generic model to the target surface.

3.2.1 Error minimization scheme

The evaluation of these repetitions depends on two errors: data error and smoothness error. Data error is defined as:

$$E_d = \sum_{i=1}^m \text{dist}^2(T_i v_i, \Gamma), \quad (6)$$

where m is the number of non-marker vertices in U , and $\text{dist}()$ function computes the distance between a vertex on U and its closest vertex on Γ .

Since the generic model has 328K triangles, an “Approximate Nearest Neighbour” searching library [20] is used to identify the closest vertices.

Smoothness error is defined as:

$$E_s = \sum_{\{i,j\} \in \text{edges}(U)} \|T_i - T_j\|^2 \quad (7)$$

The smoothness error is not defined for the smooth surface but for the actual deformations applied to the generic model. The diffu-

sion process of transformations extends this deformation to neighboring vertices. The result is that deformations for an area of nearby vertices are maintained as close as possible.

The definition of smoothness error minimizes the deformation over the template surface and thus prevents adjacent parts of the template surface from becoming aligned to disparate parts of the target surface.

The overall error is defined as a weighted sum of these two errors:

$$E = aE_d + bE_s \quad (8)$$

where a and b are two weights for the data and smoothness errors.

By setting two different weights for data and smoothness error, and by minimizing the overall error in (8) after running the optimization using L-BFGS-B [4], we get the fine mapping results in Figure 8. From three different viewing angles, the fine mapping result looks the same as the target.

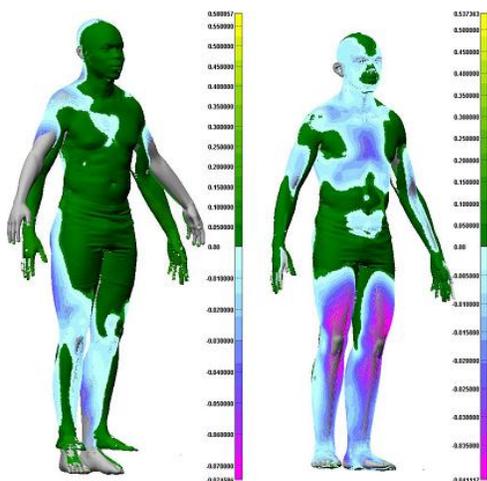


Figure 7: Comparisons between “csr0232a” and the generic model (left: original generic model and the target, right: RBF-deformed generic model and the target)

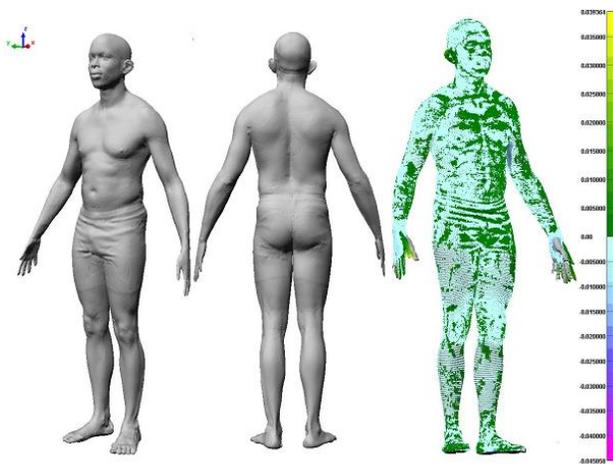


Figure 8: Fine mapping results for the target “csr0232a” (from left to right: diagonal view, back view, and a comparison between fine mapping result and the target)

In implementation of this approach, some special treatments need to be taken into consideration. To avoid the two legs get

bridged, normal information of each vertex should be calculated for finding a compatible targeting vertex. When aligning the generic surface to the target where data is missing, we get an evenly distributed result by keeping vertices unmoved if their nearest neighbors lie at the boundary of holes. Also both of the hands and ears are specially treated by finding their index in the whole model and assigning a lower weight (e.g., 0.1) for the data error during the fine mapping process.

3.2.2 Analysis of results

With different sets of weights, Figure 9 shows a comparison of results. By setting a higher weight for the data error, the artifact of scratches on the surface is removed. A higher weight for the smoothness error will create a very smooth surface while a higher weight for the data error will keep more surface details. Thus a balance needs to be kept by selecting a set of proper weights. And this control is easier compared with that in [3], where they need to adjust weight value selections for more errors.

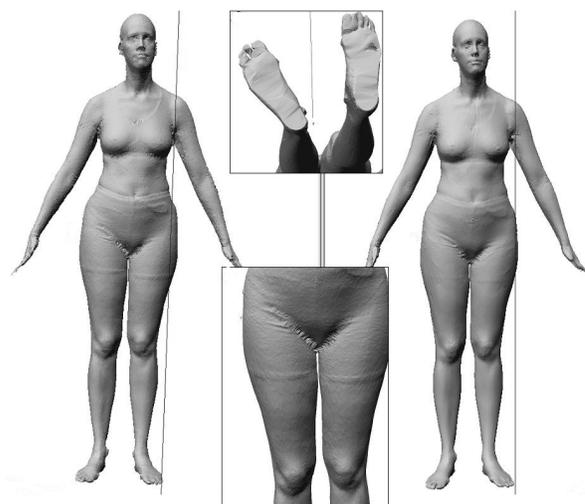


Figure 9: Fine mapping results from two sets of weights: higher data weight (left: $a=0.96$, $b=0.04$) removes some artifacts in the result where the smoothness weight is higher (right: $a=0.71$, $b=0.29$); in the middle: feet soles repaired and smoothed, two legs get de-bridged.

We implement part of the approach in [3] and try to find the deformation ability of the defined smoothness error. Taking into account the marker and smoothness errors only, a result is shown as the middle image in Figure 10. A close look at the image patch shows that, in a high-resolution model, the markers’ error plus the defined smoothness error have a limited local deformation ability only. Vertices being far from markers will not get deformed as expected. In Figure 10, since the marker at the crotch between two legs does not help deform neighboring vertices, a fine-mapped result by adding the data error displays artifacts between the legs.

Figure 10 can also explain the reason for using a multi-resolution approach in [3] because there will be a better deformation performance if the generic model is first deformed in a lower-resolution. However, this will create another problem on deciding the resolution: a simplified generic model with a still high resolution will have the same problem in disseminating the deformation from markers; in a generic model with a too low resolution, deformations of certain vertices affected by nearby markers will get mixed up and the result will not be as good as expected. Also for a different generic model, there is no general rule that can decide a proper resolution for sub-sampling of the generic model.



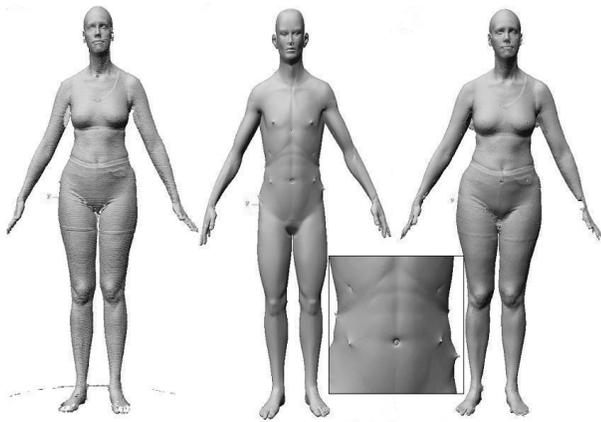


Figure 10: Deformation results (the left picture is a female body target; the middle image is an alignment result by using smoothness and marker errors; after including the data error, a finally aligned result is shown on the right)

Table 1: Summary of the comparison on final errors with two approaches when the final errors in both methods make almost no changes (for our approach: 200 iterations are needed; for the method in [3]: at least 800 iterations are needed)

Statistics	our method	method in [3]
Num_of_points	143522	143522
Mean error (m)	-0.000153	0.000154
Standard deviation (m)	0.005998	0.006129
Root mean squared error (m)	0.006000	0.006131

For comparison, we also fully implement the multi-resolution method from Allen et al. [3] to the same target. Table 1 compares the two approaches, and shows that, after convergence, the RBF-based fine matching gives almost the same accurate results as the two-resolution approach in [3].

By setting a threshold of the root mean squared error to be 0.006200 m, we measure the time and iterations needed for each approach. Table 2 shows a great advance over the method in [3]. Both experiments are tested on a PC with 3.1GHz CPU and 1GB memory.

4 PCA ON THE PARAMETERIZED CAESAR DATA

4.1 Mode variation of CAESAR dataset

The parameterized representation of each data in the CAESAR has a well-corresponded same number of vertices. By putting coordinates of all the vertices into a vector S_i , we have a dataset $\{S_i\}, i = 1 \dots m$, here m is the number of 3D data in the database.

Table 2: Summary of the comparison results with two approaches for the time and iterations needed to reach a same precision (root mean squared error = 0.006200m)

Statistics	our method	method in [3]
Time needed (minutes)	6.5	12.5
Iterations needed	160	600

By applying the PCA theory [18], new 3D human body can be reconstructed from:

$$S = \bar{S} + \Phi \cdot B, \quad (9)$$

where \bar{S} is the average of all the S_i . Φ is a matrix composed of eigen-vectors corresponding to a set of descending eigen-values. B is a variable in column vector, which can be changed to reconstruct different 3D data based on the eigen-space. Setting all the cells in vector B to be zero, the reconstructed 3D data will be an average model of all the 3D human bodies.

Based on the PCA theory, we expect that changing value of the first cell in B will create results of a most obvious variance among the database. Figure 11 shows an obvious change in height.

By changing the second element in B , a sequence of reconstructed 3D bodies is shown in Figure 12, where the most change lies in body weight. We add that in each mode variation, there will not be a single difference among all the models. Arm pose variance is small compared with the weight factor in Figure 12.

By increasing $B[3]$, there is a change in the body pose from backward to forward as shown in Figure 13. Less obvious mode changes along other eigen-vectors mainly correspond to pose and weight.

There is an earlier work from [24] on analyzing the human spaces with a volumetric representation. Our results on whole-body analysis agrees to those in [24] and we are going to perform segmented PCA on the training dataset.

4.2 PCA analysis on segmented CAESAR data

Our motivation was to find whether there is certain relation between variations of 3D whole body and segments. Therefore, we first discuss the segmentation on the generic model and apply this segmentation to all the training data for PCA analysis on each segment.

4.2.1 Segmentation on the generic model

The generic model is manually separated into six parts: head, left arm, right arm, torso, left leg and right leg. Neighboring segments have overlapping parts. Figure 14 displays the segments with and without the torso.

After segmentation on the generic model, we are able to find index files for each segment in the whole generic model. These index files are useful in two aspects: they help find the segments for each aligned CAESAR data, and help finding the overlapping band patches between two neighboring segments. The overlapping is selected for stitching neighboring segments in our future work.

4.2.2 PCA analysis on human body parts

Since each segment is part of the whole body, we can expect some similar changes as those from the whole body. For human head, it is found the first three mode variations correspond mainly to the height, pose change along forward-backward and left-right directions. After removing these changes, we also find the shape changes in Figure 15, where the middle column pictures are an average head. These pictures are another proof that along each eigen-vector, there is always a combination of variations on height, pose, or scale.

PCA analysis on the torso also creates interesting results as shown in Figure 16. Besides the first mode change in height, the second variation lies in weight, and the third mode corresponds to a volume-preserving deformation. The next few mode changes relate to pose and shape, e.g. bending of the torso.

As Figure 1 shows, the first mode variation controls the height of the left arm, while other seven modes correspond to variations in poses and scales. The similar height variation also exists in the first mode of changes for right leg. Figure 17 displays the other four changes corresponding mainly to pose and scale of the right leg.



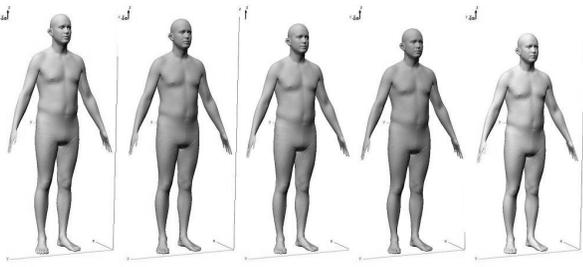


Figure 11: First mode variation among the selected database: height (from left to right: $B[1] = -20, -10, 0, 10, 20$. Other cells in B are all set to be zero)

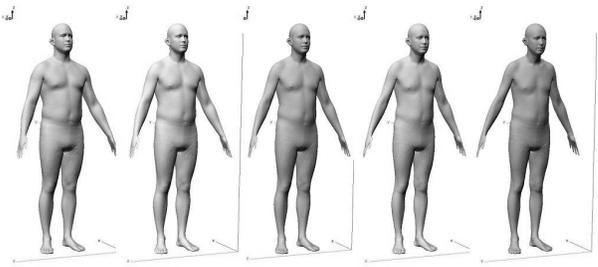


Figure 13: Third mode variation among the selected database: pose from backward to forward (from left to right: $B[3] = -16, -8, 0, 8, 16$. Other cells in B are all set to be zero)

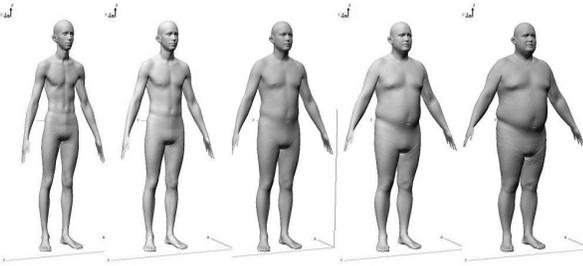


Figure 12: Second mode variation among the selected database: weight and arm pose (from left to right: $B[2] = -20, -10, 0, 10, 20$. Other cells in B are all set to be zero)

We conclude that there is a reasonable interval for selection of weights along each eigen-vector at about $[-20, 20]$ and this interval shrinks as the eigen-value increases.

We find some good correspondences of variation from 3D whole body and segments, e.g. the first principal component all represents a variation of height. We also find from experiments that body parts present a richer variation than that of the whole body.

The analysis on segmented body parts is important in that it could help reconstructing more interesting results. The CAESAR database is mainly composed of body scans of people from Europe and North America, where people have different body ratios from those of Asian people. This analysis on segmented human bodies will make it possible to reconstruct Asian people by first reconstructing body parts with specific selection of coefficients and then stitching them for a complete body.

5 CONCLUSION

We introduced an efficient and robust two-step template-based approach for the alignment of the CAESAR datasets. This solves a number of problems such as modeling shape variation, finding mutually consistent surface representations, filling holes in scanned data, and registering non-rigid surfaces.

Our innovation saves a good deal of time in alignment while still providing a good basis for fine mapping. Since this two-step approach simplifies error definition, manipulations of weight selection become easier. Also experimental results prove the robustness of our approach that it works very well for both male and female bodies of different shapes, heights and of slightly different poses.

Artifacts that result from the irregularity in the topology of the generic model can be removed in two ways: (a) by applying a high weight for the data error, or (b) by re-triangulating the generic model to get random topologies.

After aligning 964 male bodies from the CAESAR database, we are able to analyze variations of the whole body along different

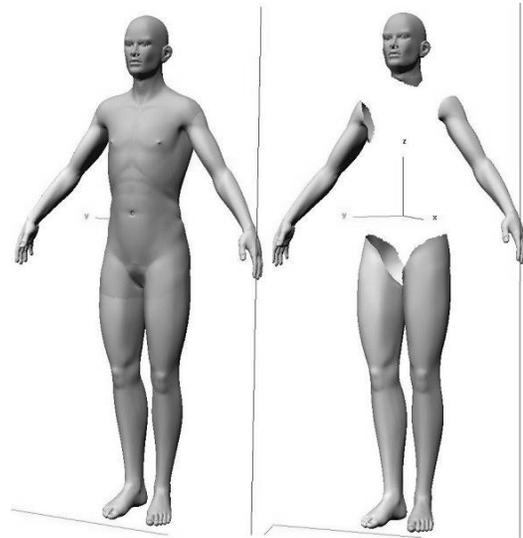


Figure 14: segmented generic model (with and without the torso)

principal components. A segmentation on the generic model also makes it possible to analyze these variations for each body segment. We are pleased to find some common variations from the whole body and its segments. Moreover, a more interesting finding is that variations of body segments are much richer than those of the whole body. These variations will be valuable for creating more variable virtual humans in computer animation.

REFERENCES

- [1] Hilton A, Starck J, and Collins G. From 3d shape capture to animated models. In *Proceedings of the First International Symposium on 3D Data Processing, Visualization, and Transmission (3DPVT 2002)*, pages 246–257, June 2002.
- [2] Allen B, Curless B, and Popovic Z. Articulated body deformation from range scan data. *ACM Transactions on Graphics (TOG)*, 21(3):612–619, 2002.
- [3] Allen B, Curless B, and Popovic Z. The space of human body shapes: reconstruction and parameterization from range scans. *ACM Transactions on Graphics (TOG)*, 22(3):587–594, 2003.
- [4] Zhu C, Byrd RH, Lu P, and Nocedal J. Algorithm 778. l-bfgs-b: Fortran subroutines for large-scale bound constrained optimization. *ACM Transactions on Mathematical Software*, 23(4):550–560, 1997.
- [5] Angelov D, Srinivasan P, Koller D, Thrun S, Rodgers J, and Davis J. Scape: Shape completion and animation of people. *ACM Transactions on Graphics (TOG)*, 24(3):408–416, 2005.



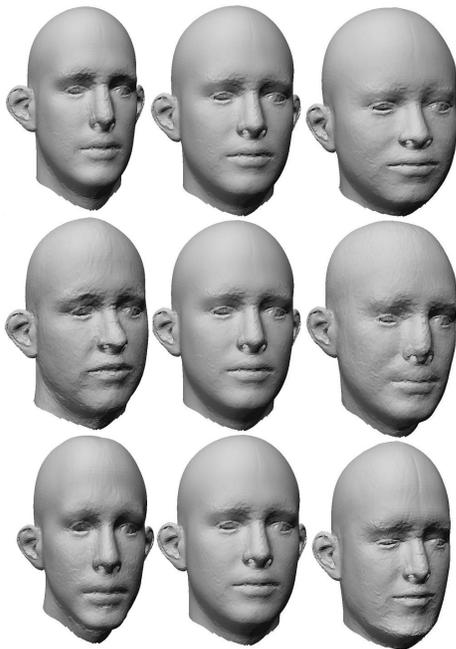


Figure 15: First three mode variations of head shape after discarding their difference in pose (first row also corresponds to height change, second row corresponds to forward-backward pose change, and third row has a change in pose along left-right direction)

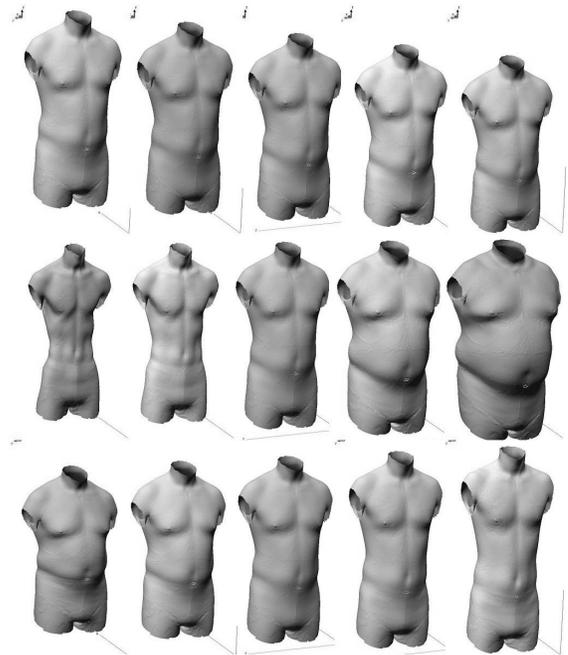


Figure 16: First three mode variations of the torso (the middle column torsos are the average one)

- [6] Cohen-Or D, Solomovic A, and Levin D. Three-dimensional distance field metamorphosis. *ACM Transactions on Graphics (TOG)*, 17(2):116–141, 1998.
- [7] DeCarlo D, Metaxas D, and Stone M. An anthropometric face model using variational techniques. In *Computer Graphics (Proceedings of SIGGRAPH 98)*, pages 67–74, 1998.
- [8] Wiley DF, Amenta N, Alcantara DA, Ghosh D, Kil YJ, Delson E, Harcourt-Smith W, Rohlf FJ, John C, and Hamann B. Evolutionary morphing. In *Proceedings of IEEE Visualization*, pages 431–438, 2005.
- [9] Praun E, Sweldens W, and Schroder P. Consistent mesh parameterization. In *Computer Graphics (Proceedings of SIGGRAPH 2001)*, pages 179–184, 2001.
- [10] Bookstein FL. *Morphometric tools for landmark data*. Cambridge University Press, 1997.
- [11] Seo H and Magnenat-Thalmann N. An automatic modeling of human bodies from sizing parameters. In *ACM SIGGRAPH 2003 Symposium on Interactive 3D Graphics*, pages 19–26, 2003.
- [12] Davis J, Marschner SR, Garr M, and Levoy M. Filling holes in complex surfaces using volumetric diffusion. In *Proceedings of the First International Symposium on 3D Data Processing, Visualization, and Transmission (3DPVT)*, June 2002.
- [13] Feldmar J and Ayache N. Rigid and affine registration of smooth surfaces using differential properties. In *Proceedings of 3rd European Conference on Computer Vision (ECCV)*, pages 397–406, 1994.
- [14] Carr JC, Beatson RK, Cherrie JB, Mitchell TJ, Fright WR, McCallum BC, and Evans TR. Reconstruction and representation of 3d objects with radial basis functions. In *Computer Graphics (Proceedings of SIGGRAPH 2001)*, pages 67–76, 2001.
- [15] Noh JY, Fidaleo D, and Neumann U. Animated deformations with radial basis functions. In *ACM Virtual Reality and Software Technology (VRST)*, pages 166–174, 2000.
- [16] Noh JY and Neumann U. Expression cloning. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*, pages 277–288, 2001.
- [17] Kahler K, Haber J, Yamauchi H, and Seidel HP. Head shop: Generating animated head models with anatomical structure. In *Proceedings*

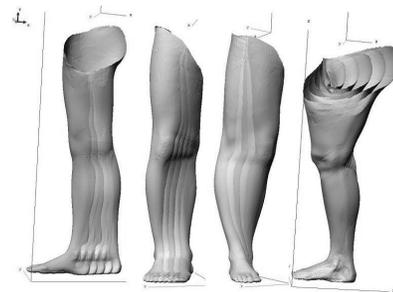


Figure 17: 2nd to 4th mode variation of the right leg (in each picture, five different weights are selected for creating the variant models)

- of the ACM SIGGRAPH Symposium on Computer Animation*, pages 55–64, 2002.
- [18] Smith LI. A tutorial on principal components analysis. maintained by Cornell University, 2002.
- [19] Whitaker R. A level-set approach to 3-d reconstruction from range data. *International Journal of Computer Vision*, 29(3):203–231, 1998.
- [20] Arya S, Mount M, Netanyahu NS, Silverman R, and Wu AY. An optimal algorithm for approximate nearest neighbor searching. *Journal of the ACM*, 45(6):891–923, 1998.
- [21] Marschner SR, Guenter B, and Raghupathy S. Modeling and rendering for realistic facial animation. In *Proceedings of 11th Eurographics Workshop on Rendering*, pages 231–242, 2000.
- [22] Blanz V and Vetter T. A morphable model for the synthesis of 3d faces. In *Computer Graphics (Proceedings of SIGGRAPH 99)*, pages 187–194, 1999.
- [23] Lee WS and Magnenat-Thalmann N. Fast head modeling for animation. *Journal Image and Vision Computing*, 18(4):355–364, 2000.
- [24] Azouz Z, Rioux M, Shu C, and Lepage R. Characterizing human shape variation using 3d anthropometric data. *The Visual Computer*, 22(5):302–314, 2006.

