2004-01-2188

Exploring the space of human body shapes: data-driven synthesis under anthropometric control

Brett Allen, Brian Curless, Zoran Popović
University of Washington

ABSTRACT

In this paper, we demonstrate a system for synthesizing high-resolution, realistic 3D human body shapes according to user-specified anthropometric parameters. We begin with a corpus of whole-body 3D laser range scans of 250 different people. For each scan, we warp a common template mesh to fit each scanned shape, thereby creating a one-to-one vertex correspondence between each of the example body shapes. Once we have a common surface representation for each example, we then use principal component analysis to reduce the data storage requirements. The final step is to relate the variation of body shape with concrete parameters, such as body circumferences, point-to-point measurements, etc. These parameters can then be used as "sliders" to synthesize new individuals with the required attributes, or to edit the attributes of scanned individuals.

INTRODUCTION

Digital human characters have many applications in areas such as computer graphics, computer vision, ergonomic and clothing design, and communication (e.g., avatars for virtual conferencing or online fitting rooms). However, creating the 3D shape models for these applications is a difficult and time-consuming task, particularly if realism is desired. This modeling task is further multiplied for each different character that is needed for a particular application.

In this paper, we will demonstrate our method for creating human character models automatically, subject to desired body shape parameters.

Our approach is data-driven, that is, we will use a database of data collected from real people. Anthropometric studies measure distances and circumferences between landmark points on the body. These data have been used in the past to create computer models of shape, such as the work of DeCarlo et al. [DeCarlo98]. However, such sparse sets of measurements cannot provide the full level of realism that one has come to expect in modern computer graphics. Fortunately, recent anthropometric studies have collected dense surface data using 3D laser range scanners, such as the one shown in Figure 1(a). Although the surfaces captured by such scanners may be noisy or incomplete in some regions (see Figures 1(b) and (c)), 3D scanners provide an excellent source of data for creating computer models.

For example, Blanz and Vetter demonstrated how to use a collection of head scans to create an editable head model that can even be used to estimate 3D shape from photographs [Blanz99]. In the past, we have shown how to use 3D scans to learn how body shape varies with pose [Allen02]. In this paper, we will focus on using

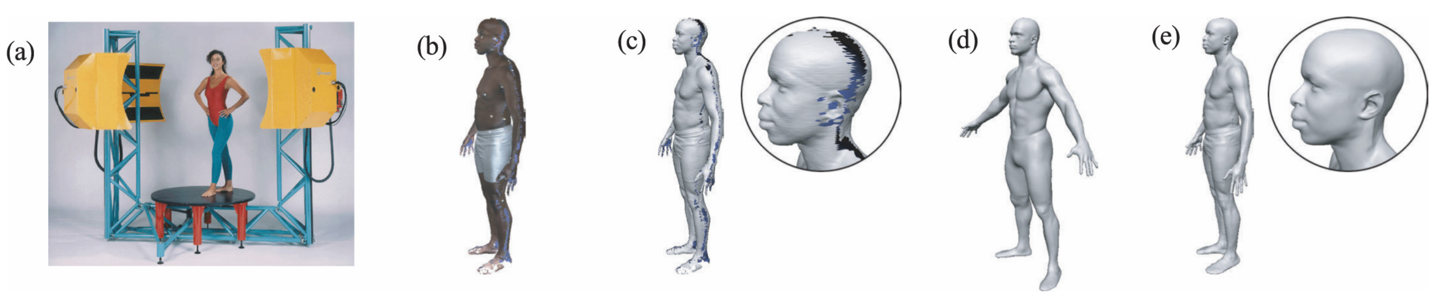

Figure 1: (a) Cyberware whole-body range scanner. (b) 3D scan of a subject. Notice the adhesive white markers. (c) Scan without surface color. The dark (blue) regions are holes in the scanned surface. (d) Artist-generated template surface, without holes. (e) The template surface after being deformed to match the scan in (c). Notice that the holes have been filled in.

whole-body scans to learn how body shape varies between individuals in all parts of the body, using the techniques introduced in our SIGGRAPH 2003 paper [Allen03]. A similar technique for working with whole-body scans has also been presented by Seo et al. [Seo03, Seo03a].

DATA

The input set of range scans used by our algorithm was collected as part of the Civilian American and European Surface Anthropometry Resource (CAESAR). The CAESAR Project collected traditional anthropometric measurements, demographic information, and 3D range scans of several thousand individuals in the United States and Europe. During the scanning, the subjects wore gray bicycle shorts, a nylon cap to cover the hair, and for the female subjects, a gray sports bra. In addition, 74 adhesive markers were placed on each subject at anthropometric landmarks. These markers are visible in the texture data associated with the scan, as shown in Figure 1(b).

The scans were taken using a Cyberware whole-body scanner, which acquires approximately 250,000 surface points with an accuracy of 2-5 mm. In occluded areas (such as under the arm and between the legs), and areas with grazing-angle views (such as the top of the head and shoulders) the scanner is unable to acquire the surface, resulting in holes in the reconstructed mesh. A detail of these holes is shown in Figure 1(c).

ALGORITHM

The key to analyzing the variation in body shape is to compare corresponding points on each surface. For example, the 74 landmark points on each scan give us information about where in 3D space certain bony landmarks were located. Using traditional anthropometric techniques, we can use the landmark positions to calculate certain linear measures about each individual, such as the length of major bones in the body. However, to analyze the full shape of an individual, we need to create a correspondence between many more points over the body – 60,000 in all. This labeling task goes beyond what is possible with human labor, and so we have developed an automatic algorithm to establish a correspondence between each scan.

Our approach begins with a template surface (see Figure 1(d)), which is a triangle mesh surface representation of a human that contains the 60,000 points that we would like to match. For each scanned individual, we will deform this template shape so that it is as close as possible to the scanned surface. In addition to matching the template, we also wish to minimize the distortion of the template surface, so that the template's nose will map to the scanned individual's nose, the eyes to the eyes, and so on.

This matching process uses the 74 landmark positions as an initialization, and proceeds in a coarse-to-fine manner. The mathematical details of this process are described in our 2003 paper [Allen03]. An example of the resulting matched mesh is shown in Figure 1(e).

Once the matching process is complete, we can fully describe the shape of any individual in our data set as a list of the x, y, and z coordinates of each vertex in the deformed template surface. We will refer to these shape vectors as $\mathbf{s}_i$, where i is the index of the example. Each shape vector contains 180,000 elements, and there are 125 shape vectors for the male scans, and 125 for the female scans. This quantity of data is quite cumbersome to work with, and because many individuals' body shapes are quite similar to each other, there is considerable redundancy. Consequently, before continuing, we will reduce the data using the technique of principal component analysis (PCA).

We compute the principal component analysis on the $\mathbf{s}_i$ vectors (separately for the male and female data sets), using the technique described by Turk and Pentland [Turk91]. The result is a mean shape vector $\mathbf{m}$, and 124 component vectors $\mathbf{v}_j$, where j is the component index. Also, for each example with index i, we have a weight vector, $\mathbf{w}_i$. To reconstruct example i, we simply take a linear combination of the components:

$$\mathbf{s}_i = \mathbf{m} + \sum_{j=1}^{124} (\mathbf{w}_i)_j \, \mathbf{v}_j$$

The key feature of PCA is that the components are organized such that $\mathbf{v}_1$ is the direction of greatest variation in shape-space, and $\mathbf{v}_{124}$ is the direction of least significant variation. To reduce the amount of data we need to store, we can throw away the $\mathbf{v}_i$ vectors above a certain value of i, and still reconstruct reasonable body shapes. For example, instead of keeping all 124 components in our male dataset, we can reduce the number of components to 40 or fewer.

However, the goal of this work is to synthesize new body shapes, not just reconstruct the example shapes. In order to accomplish this task, we need to find a relationship between intuitive attributes (such as height and other measurements of the body) and body shape. One simple way to find this relationship is to use linear regression to relate the component weights $\mathbf{w}_i$ of the observed individuals and their attributes. For example, if the height of the male subject with index i is h_i, then we could do a regression between h_i and $(\mathbf{w}_i)_1$ to find out how the first principal component weight relates to height. The resulting best-line fit is shown in Figure 2. If we repeat this procedure for the other principal components, then we can create an individual with any particular height using the learned linear relationships.

In Figure 3(a), we demonstrate using this regression technique for six different anthropometric measures: stature (height), bitragion breadth (head breadth),

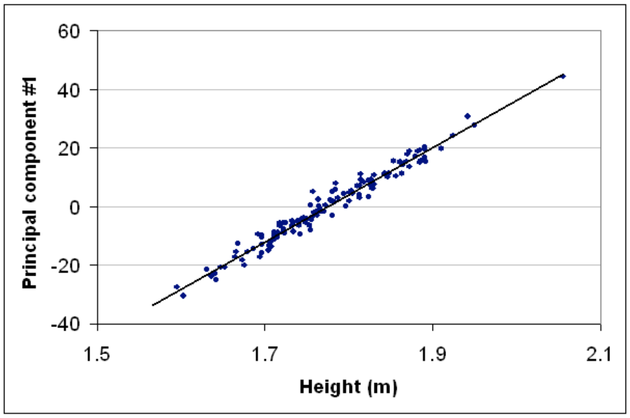

Figure 2: Graph of subject height versus first principal component weight for 125 male subjects.

shoulder breadth, arm length, bi-cristale breadth (hip breadth), and leg length. We calculate these measurements using the point-to-point distance between the landmark points in the CAESAR database. Figure 3(a) shows the average male body shape, and then six separate edits relative to this average. Notice however, that these edits are not independent; for example, increasing the bitragion breadth in column 2 results in a taller, larger individual. Consequently, this technique is useful for only specifying a single parameter, or for exploring how overall body shape relates to a single parameter as a general trend.

We would like to be able to specify several parameters in order to "dial up" a particular body type. For this task, we must perform our regression across several attributes. We can achieve a least-squares best fit using a matrix pseudoinverse calulation. Suppose we have p different attributes, and the subject i's attribute values are $\mathbf{a}_i = [x_1\ x_2\ \ldots\ x_p\ 1]^T$. If we combine the vectors $\mathbf{a}_i$ into a matrix $\mathbf{A}$, and the vectors $\mathbf{w}_i$ into a matrix $\mathbf{W}$, then we can compute the relationship between all of the attributes and the body shape using the following equation:

$$\mathbf{M} = \mathbf{W}\mathbf{A}^{+}$$

In the above equation, $^{+}$ denotes the pseudoinverse operation. Now given a new set of attributes, $\mathbf{a}$, we can compute the principal component weights as $\mathbf{w} = \mathbf{M}\mathbf{a}$.

In part (b) of Figure 3, we demonstrate editing all six measurements in tandem. Notice, for example, how the average subject's height is preserved in all cases except for column 1, where the height is edited. Similarly, the bitragion breadth is constant in all columns except for column 2, and so on. In this way, we can truly specify all six parameters to create a body type of the desired proportions.

Note that it is not necessary to begin with the average body shape. Indeed, due to the linearity of our calculations, we can start with any body shape, real or imagined, and edit various aspects of their physique.

CONCLUSION

To conclude, we have shown how we can synthesize and edit body shapes using intuitive anthropometric controls by training our system with raw range scan data. Our technique enables a user to create a wide variety of body shapes without any modeling expertise. The resulting meshes are complete and may be sampled at different resolutions for many different digital human applications.

Furthermore, our system may be used as a data exploration device, for visualizing and exploring relationships between various bodily measurements. Indeed, we have also experimented with using other kinds of attributes, including demographic information such as weight, shoe size, or income, to explore population trends with respect to these variables. Note that in these cases, our linear model can only be considered a rough approximation to the true relationships in the data.

A video demonstration of our algorithm, and additional applications and results may be downloaded from the following web page:

http://grail.cs.washington.edu/projects/digital-human/pub/allen03space.html

ACKNOWLEDGMENTS

We are grateful for Kathleen Robinette for providing the CAESAR data, Domi Pitturo for supplying the template mesh, and Daniel Wood for his multiresolution-mesh implementation. This work was supported by the University of Washington Animation Research Labs, NSF grants CCR-0098005 and EIA-0121326, the Natural Sciences and Engineering Research Council of Canada, and industrial gifts from Microsoft Research, Electronic Arts, and Sony.

REFERENCES

1. [Allen02] Brett Allen, Brian Curless, and Zoran Popović. 2002. Articulated body deformation from range scan data. *ACM Transactions on Graphics (ACM SIGGRAPH 2002)*, 21, 3, 612—619.
2. [Allen03] Brett Allen, Brian Curless, and Zoran Popović. 2003. The space of all body shapes: reconstruction and parameterization from range scans. *ACM Transactions on Graphics (ACM SIGGRAPH 2003)*, 22, 3, 587—594.
3. [Blanz99] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3D faces. In *Proceedings of ACM SIGGRAPH 99*, ACM Press/Addison-Wesley Publishing Co., Computer Graphics Proceedings, Annual Conference Series, 187—194.

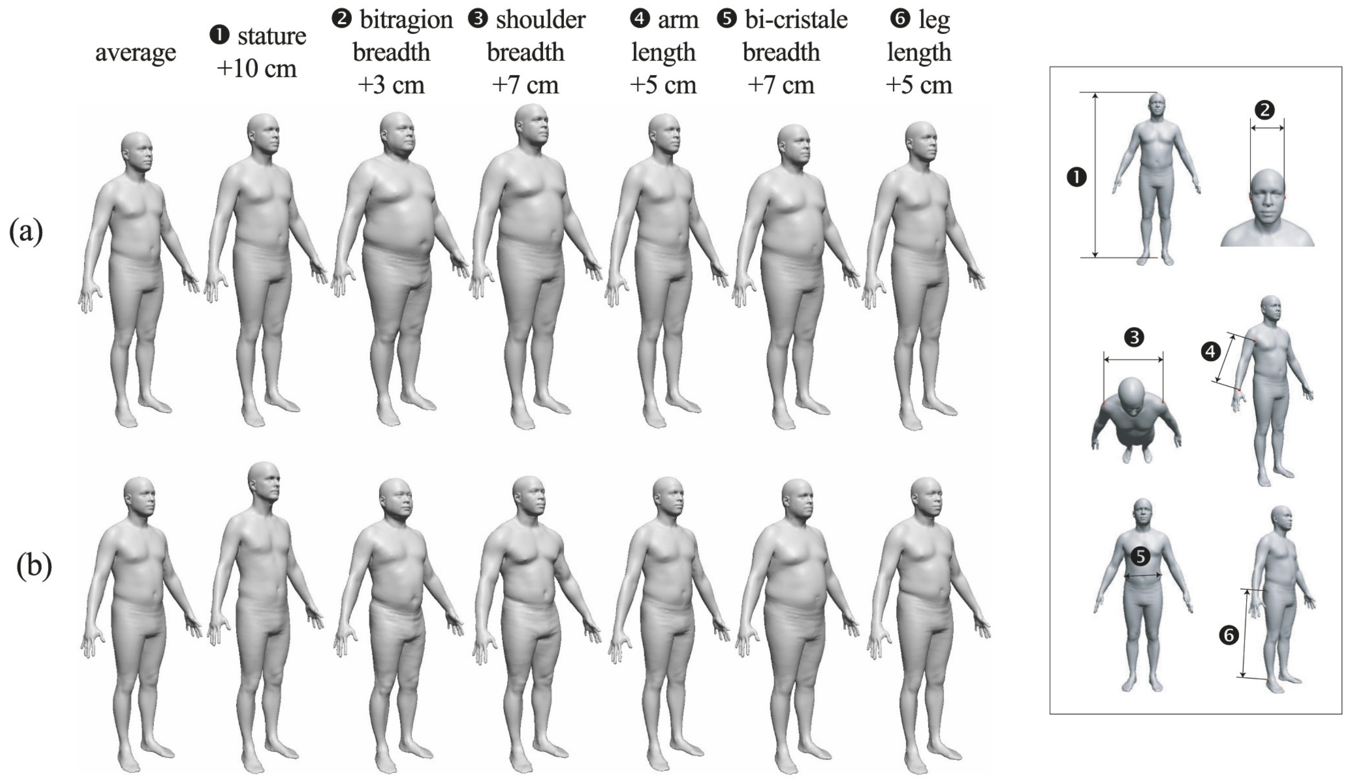

Figure 3: Editing attributes of an individual's body shape. On the top row, single attributes have been edited to explore how body shape varies relative to different parameters. On the bottom row, all six attributes are edited in tandem, so that as one attribute changes, the other five remain fixed.

4. [DeCarlo98] Douglas DeCarlo, Dimitris N. Metaxis, and Matthew Stone. 1998. An anthropometric face model using variational techniques. In *Proceedings of ACM SIGGRAPH 98*, ACM Press, Computer Graphics Proceedings, Annual Conference Series, 67—74.
5. [Seo03] Hye-Won Seo and Nadia Magnenat-Thalmann. 2003. An automatic modeling of human bodies from sizing parameters. In *Proceedings of the 2003 Symposium on Interactive 3D Graphics*, ACM Press, 19—26.
6. [Seo03a] Hye-Won Seo, Frederic Cordier, and Nadia Mangnenat-Thalmann. 2003. Synthesizing body models with parameterized shape modifications. In *Proceedings of the 2003 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, Eurographics Association, 120—125.
7. [Turk91] Matthew Turk and Alex Pentland. 1991. Eigenfaces for recognition. *Journal of Cognitive Neuroscience* 3, 1, 71—86.

CONTACT

Further information about the Digital Humans project at the University of Washington may be found at: http://grail.cs.washington.edu/projects/digital-human.