

Estimating 3D Human Shapes from Measurements

Stefanie Wuhrer · Chang Shu

Received: date / Accepted: date

Abstract Recent advances in 3-D imaging technologies give rise to databases of human shapes, from which statistical shape models can be built. These statistical models represent prior knowledge of the human shape and enable us to solve shape reconstruction problems from partial information. Generating human shape from traditional anthropometric measurements is such a problem, since these 1-D measurements encode 3-D shape information. Combined with a statistical shape model, these easy-to-obtain measurements can be leveraged to create 3D human shapes. However, existing methods limit the creation of the shapes to the space spanned by the database and thus require a large amount of training data. In this paper, we introduce a technique that extrapolates the statistically inferred shape to fit the measurement data using non-linear optimization. This method ensures that the generated shape is both human-like and satisfies the measurement conditions. We demonstrate the effectiveness of the method and compare it to existing approaches through extensive experiments, using both synthetic data and real human measurements.

Keywords Human models · statistical prior · three-dimensional reconstruction

1 Introduction

Realistic 3D human shapes are required in various applications, such as the design of products that fit a target population. Typically, these shapes need to have certain characteristics or be samples from a population. Although it is possible to digitize humans using 3D imaging technologies, it is impractical to find and scan suitable human subjects for each individual application. On the other hand, there is a long history of using anthropometric measurements to describe the human shape. These measurements are linear and curvilinear distances between anatomical landmarks or circumferences at predefined locations. In spite of providing limited shape information, they are readily available, and therefore widely used.

Intuitively, the linear measurements should encode information about the body shape. If an accurate 3D human shape can be inferred from a simple set of measurements, traditional anthropometric data can be leveraged to create 3D datasets without resorting to expensive scanning equipment.

The early attempts at solving this problem deform a generic human model to fit the measurement data. However, the models thus obtained are not necessarily human forms because generating a human shape from a sparse set of measurements is an under-constrained problem. More sophisticated methods use knowledge of the human shape learned from a database of 3D scans. These methods build a statistical

Stefanie Wuhrer
Saarland University and Max-Planck Institut Informatik
Campus E 1.7, Saarbrücken, Germany
E-mail: swuhrer@mmci.uni-saarland.de
Tel.: +49-681-302-70751 Fax: +49-681-302-70155

Chang Shu
National Research Council of Canada
1200 Montreal Road, Ottawa, Canada
E-mail: chang.shu@nrc-cnrc.gc.ca

model from a set of registered 3D scans and establish a relationship between the measurements and the shape space, which allows to predict detailed 3D shapes from sparse and partial shape data. Apart from predicting 3D shapes from measurements [21], [1], other partial shape information, such as marker positions [3] and 2D images [15], [8], has also been used to reconstruct 3D shapes.

Existing methods that predict the 3D shapes using a statistical model limit the generated shapes to the space spanned by the shape model and local variations of human shapes that are outside this space cannot be predicted accurately. This means that these methods require a 3D database that accurately represents a target population. For instance, human subjects of different ethnicities and different weight classes must be present in the database. To apply these techniques, one must carefully choose a target population and acquire a representative database. In reality, the acquisition of a 3D anthropometry data that represents a target population well is complicated and expensive. For instance, the acquisition of the Civilian American and European Surface Anthropometry Resource took 4 years and cost \$6 Million¹.

In this paper, we address the problem of estimating 3D human body and face shapes given a set of anthropometric measurements. Our main goal, in contrast to previous work, is to predict shapes that are inside the space of human shapes, but still account for local variations not captured by the training data. This extrapolation allows us to create human shapes based on relatively small 3D databases.

As most of the previous approaches, our method does not take posture or expression changes into account. That is, we assume that all of the training data is in a standard pose. This scenario is commonly assumed since in a typical 3D anthropometry survey, the human subjects are asked to maintain a standard posture.

The rest of the paper is organized as follows. Section 2 reviews previous work on estimating 3D body or face shapes from partial data. Section 3 gives an overview of our approach, and Section 4 outlines the details of a shape refinement step, which constitutes the main contribution of this paper. We aim for design applications in engineering where accuracy is essential. We show in Section 5, through experiments, that straightforward application of existing techniques does not produce satisfactory results. Finally, Section 6 concludes the work.

2 Related Work

This section reviews existing work that aims to estimate 3D human shapes based on partial input information. We categorize the related literature by the type of input data.

2.1 Estimating 3D Body or Face Shapes from Measurements

DeCarlo et al. [10] presented an early attempt at using anthropometric measurements to compute human face shapes. The method generates a random set of measurements based on a set of rules that capture, for instance, typical facial proportions. To create a 3D face shape, the method deforms a template model to fit these measurements using a free-form deformation technique called variational modeling. Variational modeling allows to give the measurements as constraints when deforming the template. To ensure that the estimated shape remains in the shape space of human shapes, much care is taken when creating the measurements that are used as constraints. While this method generates human-like face shapes using a statistical model, the accuracy of the shapes created by this method is limited by the sparse set of measurements used to represent the training data. That is, the results are overly smooth and do not contain realistic details.

The following approaches estimate human body shapes from a discrete set of measurements, all proceeding by learning a linear or near-linear mapping between the training set of human shapes and the space of measurements and then using this mapping to predict a shape based on a new set of measurements. While these approaches work well when the predicted shape is inside the shape space spanned by the training set, they do not allow for extrapolations from this shape space.

Wang [25] uses a parameterized database of human shapes consisting of *feature patches* to find a new human shape based on a given set of measurements. Feature patches are initialized to smooth patches that are subsequently refined to capture geometric details of the body shape. While this refinement does not maintain the smoothness of the patches, the final results are visually smooth and do not contain realistic

¹ <http://www.sae.org/standardsdev/tsb/cooperative/caesumm.htm>

localized shape details. The approach finds models in the database with measurements similar to the given set of measurements and computes the new human shape as a linear combination of these models. Wei et al. [26] use a similar approach, but the models are represented using a layer-based representation.

Allen et al. [1], [2] start from a parameterized database of human bodies in similar poses and perform Principal Component Analysis (PCA) of the data. The training database is used to learn a linear mapping called *feature analysis* from the set of measurements of the training data to the PCA space. Feature analysis can be used to compute a new PCA weight based on a new set of measurements, and the learned PCA model allows to compute a new body shape from this PCA weight. Chu et al. [9] perform feature analysis on a database of human shapes consisting of feature patches.

Seo and Megnenat-Thalmann [21] represent the human body as a triangular mesh with an associated skeleton model. As with the approach of Allen et al., this approach reduces the dimensionality of the data using PCA. The approach learns a mapping from the set of measurements of the training data to the PCA space using an interpolating radial basis function with Gaussian kernel [12]. Hasler et al. [17] apply the two previously reviewed approaches to a new representation of human body shapes that is posture invariant. Their method simultaneously models body pose and shape.

Recently, Baek and Lee [4] presented a technique that uses hierarchical clustering to build a statistical model of the training data. The approach proceeds by clustering the training database and by performing a multi-cluster analysis. To predict a body shape based on a set of input measurements, the approach finds the shape within the learned shape space that best describes the measurements using an optimization of shape parameters. It is shown experimentally that accurate and visually pleasing body shapes are estimated when the input body sizes are inside the shape space spanned by the training data. This approach is conceptually similar to the first optimization step of our algorithm. Hence, we expect that this approach does not allow to model shape variations that are outside the shape space spanned by the training data.

2.2 Estimating 3D Body or Face Shapes from Markers

Anguelov et al. [3] aim to estimate a 3D human body shape based on a sparse set of marker positions. This technique is useful when motion capture data is available. Anguelov et al.'s SCAPE model represents the human body as a triangular mesh with an associated skeleton model. As with the approach of Allen et al., this approach reduces the dimensionality of the data using PCA. Anguelov et al. compute a new triangular mesh based on a set of marker positions by adjusting the PCA weights to solve a non-linear optimization problem.

2.3 Estimating 3D Body or Face Shapes from Images

Finally, we review approaches that aim to estimate human body shapes based on a set of input images. The following approaches proceed by learning a correlation between a training set of 3D face or body shapes and a set of derived 2D images and by using this learned correlation to predict a shape based on a new set of 2D images. These approaches work well when the predicted shape is inside the shape space spanned by the training set. However, they do not handle optimizations outside the learned shape space of 3D models.

Blanz and Vetter [6] estimate a 3D face shape from a single input image in neutral expression. They start by building a parameterized database of textured 3D faces and by performing PCA on the shape and texture data. Given an input image, the learned PCA space is searched to find the textured shape that best explains the input image.

Seo et al. [22] estimate a body shape of a human in standard posture from two images. Starting from a parameterized database of human meshes in similar poses, the approach performs PCA of the 3D data. Given the two images, the learned PCA space is searched to find a set of PCA weights that corresponds to a 3D shape that matches the input images well. Boisvert et al. [7] solve the same problem using an optimization that aims to find the optimal set of PCA weights.

Chen and Cipolla [8] aim to estimate the human body shape in a fixed pose based on a given silhouette. Starting from a parameterized database of human meshes in similar poses and a set of corresponding silhouettes, the approach performs PCA of the 3D and 2D data separately. The approach then computes

a mapping from the PCA space of the silhouette data to the PCA space of the 3D data using a Shared Gaussian Process Latent Variable Model (SGPLVM) [23]. Given a new silhouette, the approach maps the silhouette into silhouette PCA space and uses the SGPLVM to map to the PCA space of the 3D meshes. Ek et al. [13] use a similar approach to estimate the pose of a human body based on a given silhouette. Recently, statistical prior models allow to track the pose of a human body even when occlusions and self-occlusions are present [20].

Guan et al. [15] estimate both the shape and pose of a human body shape from a single photograph based on the SCAPE model. When adjusting the PCA weights, the shape is deformed to best match the image in terms of a shape-from-shading energy. Weiss et al. [27] propose to use the SCAPE model to compute a 3D body scan from noisy Kinect data. Hasler et al. [16] estimate both the shape and pose of a human body from a photograph of a dressed person.

3 Approach

This section outlines the proposed approach. As input to the method, we are given a database of triangular manifold meshes X_0, \dots, X_{n-1} of human bodies or faces with similar posture or expression and a set of measurements \mathcal{S} . Let P_i in \mathcal{S} denote the measurements corresponding to X_i . Furthermore, we are given a set of distances P_{new} . Our aim is to estimate a shape X_{new} that interpolates the distances P_{new} .

As previous methods, the approach proceeds by learning a correlation between the shapes and the measurements. When predicting a new shape, our approach finds an initial solution based on the learned correlation. Unlike previous approaches, however, our approach refines this solution to fit the measurements using two steps of non-linear optimization. First, we optimize the shape of the model with respect to the learned shape space. That is, we aim to find the point in the learned shape space that best describes the measurements. This gives a realistic human shape drawn from the distribution fitted to the training data, which can only contain local shape variations present in the training data. To account for other shape variations, we perform a second mesh-based optimization. This optimization deforms the shape to fit the measurements as close as possible while satisfying a smoothness constraint without using prior knowledge of the body shape, and therefore predicts shapes with local variations not present in the training database. This second optimization step is the main reason why our approach, unlike feature analysis, can model shape variations not present in the training database.

When using a mesh-based deformation that optimizes the sought measurements, we can only expect to obtain a realistic human body shape if we start with a body shape that is close to the solution. For this reason, we start from the shape in the learned shape space that best describes the sought measurements as opposed to starting directly from the point in shape space predicted by the learned correlation.

3.1 Measurements

Anthropometric measurement represents a valuable source of human shape information, from which we can infer 3D shapes for design applications. To learn a relationship between anthropometric measurements and a set of 3D models, the measurements computed from the models should be equivalent to the measurements conducted by a human operator. The measurements can be one of three types: Euclidean, geodesic, and circumference. A Euclidean or geodesic distance can be easily computed from two vertices on the model. A circumference can be computed by intersecting (part of) the model with a plane, finding the (possibly closed) polygonal chain of the intersection that contains a specified vertex, and measuring the circumference of the convex hull of this chain. This type of measurement is shown for the hip and waist circumferences in Figure 3. The reason we measure the length of the convex hull instead of the length of the chain itself is that many anthropometric measurements, such as the chest circumference, measure the length of a convex hull. We specify the intersecting plane π using a vertex p and a normal direction \mathbf{n} . Furthermore, we specify a part of the model that is to be intersected with the plane π . This is necessary for some measurements, such as for the chest or hip circumference, where we wish to only intersect the torso (and not the arms) of the model with π . Note that the intersection of π with a subset of X consists of a set of polygonal chains, where every vertex of the chain is either a vertex of X or the intersection of an edge of X with π .

3.2 Training

We now consider learning a correlation between the set of shapes X_i and the corresponding measurements P_i . To do this, the database of the human shapes needs to be parameterized, a process that computes point-to-point correspondences among the shapes. This is in general a difficult problem [24]. In practice, anthropometric markers identifying salient anatomic positions are often used to guide the correspondence process. Such markers are provided in 3D anthropometric surveys, for example, the Civilian American and European Surface Anthropometry Resource (CAESAR) database [19]. In this work, the known marker positions are used to deform a template shape to each subject of the database [1], [29], [28].

The Euclidean and geodesic measurements are specified by their two endpoints and the type of distance to be considered and the circumference measurements are given by one plane (specified by a vertex and a normal vector) and the part of the body to be intersected. Since the database is parameterized, the vertices and body parts can be given in terms of their vertex and triangle numbers on the mesh. With this information, a set of distances P_0, \dots, P_{n-1} can be computed for X_0, \dots, X_{n-1} .

We learn a mapping between \mathcal{S} and the space of human body shapes using feature analysis. We choose this model because it describes both the shape space of X_i and the mapping from the space of measurements to the learned shape space using linear functions. This allows us to compute the gradients of the energy functions used in Section 4 explicitly, and hence, to use an efficient numerical solver to optimize the energy functions. To apply feature analysis, we first perform PCA on the meshes X_i . In PCA space, each shape X_i is represented by a vector W_i of PCA weights. PCA yields a mean shape μ and a matrix A that can be used to compute a new shape X_{new} based on a new vector of PCA weights W_{new} as $X_{new} = AW_{new} + \mu$. Recall that the aim is to create a shape based on a point P_{new} in \mathcal{S} . To achieve this goal, feature analysis learns a linear mapping from P_i to $W_i, i = 0, \dots, n - 1$. This yields a matrix B that can be used to compute a vector of PCA weights W_{new} based on P_{new} as $W_{new} = BP_{new}$, which generates a shape X_{new} based on P_{new} as $X_{new} = ABP_{new} + \mu$. To give all PCA coordinates the same weight, we normalize each entry of the PCA weights by its corresponding PCA eigenvalue before performing feature analysis.

3.3 Prediction

The mapping between the measurements and the PCA weights of the 3D shapes learned in the previous section allows us to find an initial shape X_{new}^{init} given P_{new} . However, straightforward application of feature analysis has the disadvantage that the shapes obtained for atypical P_{new} are not contained in the space of human shapes (e.g. see Figure 8). Since the subsequent refinement process depends on this initialization, it is important to restrict X_{new}^{init} to be within the range of human shapes. Hence, we restrict $W_{new} = BP_{new}$ such that each dimension $W_{new}[i]$ is at most l times the standard deviation of the PCA space along dimension i . This choice is based on the assumption that the shape space is modeled as independent Gaussian distributions and most of the shapes are located within l standard deviations of the mean for a suitable parameter l . This step restricts the reconstructed shape to stay within the space of human shapes that was learned using PCA. In our implementation, we use the normalized weight vector to find an initial shape X_{new}^{init} .

The shape generated from the PCA weights can only serve as an initial estimation, because the shape space spanned by PCA is limited by the sample size and may not account for all human shape variations. To find a shape that respects the required measurements while staying in the learned shape space, we need to refine X_{new}^{init} . The following section presents our novel method for shape refinement in detail.

4 Shape Refinement

This section outlines a novel approach to refine the initial estimate X_{new}^{init} to respect the required measurements while staying in the learned shape space. We formulate the refinement problem as an energy minimization problem.

Let p_0, \dots, p_{m-1} denote the vertices of X_{new} and let $\mathbf{p}_0, \dots, \mathbf{p}_{m-1}$ denote their position vectors. Furthermore, let $W_{new} = A^+(X_{new} - \mu)$ be the vector of PCA weights corresponding to X_{new} , where A^+ is the pseudo-inverse of A .

For Euclidean measurements, we are given the target length $l_t(d)$ of the Euclidean distance of the segment d between two vertices p_i and p_j . The shape X_{new} that satisfies all given Euclidean distance constraints minimizes

$$E_e = \sum_{d \in \mathcal{D}} ((\mathbf{p}_i - \mathbf{p}_j)^2 - (l_t(d))^2)^2,$$

where \mathcal{D} is the set of all Euclidean distances that have a desired length, and p_i and p_j are the endpoints of d .

For geodesic measurements, we are given the target length $l_t(P)$ of the geodesic path P between two vertices p_i and p_j . Let $l_g(P)$ denote the geodesic length of P . In the following, we assume that the relative length of each edge e of P with respect to $l_g(P)$ does not change during the deformation of the mesh. This is a reasonable assumption as we wish to preserve relative edge lengths during the deformation. Furthermore, this assumption translates into an easy optimization problem. With this assumption, we can compute the target length $l_t(e)$ of e as $l_t(e) = \frac{l_t(P)}{l_g(P)} l_g(e)$, where $l_g(e)$ is the current length of e . Then, the shape X_{new} that satisfies all given geodesic distance constraints minimizes

$$E_g = \sum_{e \in \mathcal{P}} ((\mathbf{p}_k - \mathbf{p}_l)^2 - (l_t(e))^2)^2,$$

where \mathcal{P} is the set of all geodesic paths that have a desired length, and p_k and p_l are the endpoints of e .

For circumference measurements, we are given the target length $l_t(C)$ of the circumference C of the convex hull of the polygonal chain passing through vertex p_i and contained in the intersection between the mesh and the plane π through p_i with normal direction \mathbf{n} . We compute the length $l_g(C)$ of the circumference on the mesh as outlined in Section 3.1. Let q_i denote all of the points on the convex hull of the polygonal chain. With a similar reasoning as above, in the following, we assume that the relative length of each edge e of the convex hull with respect to $l_g(C)$ does not change during the deformation of the mesh. Hence, we can compute the target length $l_t(e)$ of e as $l_t(e) = \frac{l_t(C)}{l_g(C)} l_g(e)$, where $l_g(e)$ is the current length of e . Then, the shape X_{new} that satisfies all given circumference constraints minimizes

$$E_c = \sum_{e \in \mathcal{C}} ((\mathbf{q}_i - \mathbf{q}_j)^2 - (l_t(e))^2)^2,$$

where \mathcal{C} is the set of all circumferences that have a desired length, and q_i and q_j are the endpoints of e . Recall that q_i and q_j are not necessarily vertices of the mesh.

Our aim is to deform X_{new}^{init} , such that $E_m = E_e + E_g + E_c$ is minimized. We solve this problem using two consecutive steps by minimizing E_m with respect to W_{new} followed by minimizing E_m with respect to \mathbf{p}_i . The first minimization finds the shape in the learned shape space that best describes the measurements. This step finds a realistic human body shape whose measurements are close to the sought ones. The second minimization then further deforms the shape using a mesh-based optimization without using prior knowledge on the shape. This minimization allows the shape to deform locally in ways not present in the training data. During this minimization, a smoothness term is used to maintain a realistic body shape.

4.1 Minimization with respect to W_{new}

We minimize E_m with respect to W_{new} using a quasi-Newton approach. This approach has the advantage of having a near-quadratic convergence rate if the initial solution is close to the minimum. Since this approach requires analytic gradients, we next discuss the derivatives of the different energy terms.

The derivative of E_e with respect to \mathbf{p}_i is

$$\nabla_{\mathbf{p}_i} E_e = \sum_{d \in D(p_i)} 4((\mathbf{p}_i - \mathbf{p}_j)^2 - (l_t(d))^2)(\mathbf{p}_i - \mathbf{p}_j),$$

where $D(p_i)$ is the set of distances in D with endpoint p_i . The derivative of E_e with respect to W_{new} is $\nabla_{W_{new}} E_e = A^+ \nabla_{\mathbf{p}_i} E_e$. The derivatives of E_g with respect to \mathbf{p}_i and W_{new} are similar to $\nabla_{\mathbf{p}_i} E_e$ and $\nabla_{W_{new}} E_e$. To compute the derivative of E_c with respect to \mathbf{p}_i , recall that each vertex q_i of the convex hull of the polygonal chain is either a vertex of the mesh or the intersection of an edge of the mesh with

π . Hence, every q_i can be expressed as a convex combination of at most two vertices of the mesh. This formulation allows us to compute the derivatives of E_c with respect to \mathbf{p}_i and W_{new} in a similar way to $\nabla_{\mathbf{p}_i} E_e$ and $\nabla_{W_{new}} E_e$.

Note that the geodesic path between p_i and p_j and the circumference measurement specified by p_i and \mathbf{n} may change when the mesh is deformed. To remedy this problem, we update the geodesic paths and the circumference measurements and solve the resulting optimization problem s times. This finds the point W_{new} in PCA space corresponding to the shape that has measurements closest to the desired ones. As before, we normalize W_{new} such that each dimension $W_{new}[i]$ of W_{new} is at most l times the standard deviation of the PCA space along dimension i to ensure that the shape stays within the learned shape space. Finally, W_{new} is used to compute the shape X_{new}^{pca} as $X_{new}^{pca} = AW_{new} + \mu$.

4.2 Minimization with respect to \mathbf{p}_i

Second, we minimize E_m with respect to the vertex positions \mathbf{p}_i to ensure that local variations not present in the training data are predicted correctly by locally deforming the shape to fit the required measurements. However, simply minimizing E_m may result in meshes that are not smooth. Hence, we encourage close-by parts of the mesh to deform similarly by considering the smoothness energy

$$E_s = \sum_{p_i \in X_{new}} \sum_{p_j \in N(p_i)} (\Delta \mathbf{p}_i - \Delta \mathbf{p}_j)^2,$$

where $\Delta \mathbf{p}_i$ is the translation vector by which \mathbf{p}_i is moved and $N(p_i)$ is the one-ring neighborhood of p_i . The derivative of E_s with respect to \mathbf{p}_i is

$$\nabla_{\mathbf{p}_i} E_s = \sum_{p_j \in N(p_i)} 2(\Delta \mathbf{p}_i - \Delta \mathbf{p}_j).$$

Note that the smoothness term ensures that the predicted shape stays within the shape space of human shapes. The vertex positions are initialized to the vertex positions of X_{new}^{pca} . We minimize the combined energy $E = (1 - \lambda)E_m + \lambda E_s$ with respect to \mathbf{p}_i , where λ is the weight for the smoothing term. As before, we update the geodesic paths and the circumference measurements and solve the resulting optimization problem s times, yielding the final result X_{new} .

4.3 Discussion of parameters

The proposed method depends on the parameters l, λ, s , and on the termination criteria of the energy minimization. In our experiments, we stop the energy minimization if the relative change of the energy is smaller than $1e8$ times the machine accuracy, if the norm of the gradient is smaller than $1e - 8$, or if 100 iterations have been performed. The parameter l determines the size of the acceptable shape space for reconstruction. This parameter provides a way to trade off high measurement accuracy and likely human shapes. If the measurements given to the method are known to be accurate measurements of a real human, l can be set to a high value. If the measurements given to the method are randomly generated, however, l needs to be low to avoid non-human shapes. Hence, in our experiments with real data, we set $l = 10$ and in our experiments with synthetic data, we set $l = 3$. The parameter λ determines the relative weight of the smoothing term. In our experiments, we set $\lambda = 0.1$. The parameter s determines how often the geodesic paths and circumferences are recomputed. We need to set s large enough to achieve the required accuracy, yet small enough to obtain an efficient approach. In our experiments, we set $s = 3$ unless specified otherwise.

We leave it for future work to find the optimal parameter settings automatically. This is a challenging problem because the result depends on the parameter values, and because including the parameter values in the set of parameters to be optimized gives a highly nonlinear optimization problem.

5 Experiments

This section demonstrates the effectiveness of our method using both synthetic experiments and experiments with real data.

We implemented the proposed algorithm using C++. The implementation uses Dijkstra’s algorithm [11] to compute geodesic distances, the quickhull algorithm [5] to compute the convex hull to determine circumferences, and the limited-memory Broyden-Fletcher-Goldfarb-Shanno scheme [18] to solve the optimization problems. We choose these standard algorithms for their ease of use. To parameterize the database used for training, the efficient approach by Xi et al. [29] is employed.

We compare our results to the results obtained using feature analysis and to the results obtained by using a SGPLVM mapping from the space of the measurements to the PCA space of the models. Both approaches are reviewed in Section 2. We use the code by Ek et al. [13] to compute the SGPLVM.

5.1 Input Measurements

The synthetic experiments aim to reconstruct human face and body shapes. For each synthetic experiment, we learn a multivariate Gaussian distribution $\mathcal{N}(\mu, \Sigma)$ from the available sets of measurements. First, 200 sets of measurements are generated by randomly sampling points from $\mathcal{N}(\mu, \Sigma)$. Note that since these samples are drawn from the learned distribution, most of the measurements are similar to the ones in the training set. We denote this set of measurements by S_{close} . Second, 200 sets of measurements are generated by randomly sampling points that are located on the ellipsoidal surface $x^T \Sigma^{-1} x$, where $x = \mu + k \Sigma_d$, k is a constant, and Σ_d is the vector that contains the diagonal elements of Σ . Note that the larger k is, the farther the measurements are from the training set. Let S_k denote this set of measurements. In our experiments, we set $k = 2, 4$.

The experiments based on real data aim to reconstruct human body shapes from a set of measurements. We consider two types of input data. For the first type of input data, we digitally measure the distances in \mathcal{S} on a number of parameterized 3D human body scans and use these measurements as input to our algorithm. We denote these measurements by S_{real} . Note that in this case, the ground truth of the 3D body shape is known and can be used to evaluate the results.

For the second type of input data, we asked a number of volunteers to measure each other. This results in a set of real-world measurements of real subjects. We denote these measurements by S_{meas} . As in this case, no ground truth 3D body shape is known, we validate the accuracy of the 3D reconstruction visually by comparing the 3D shape to silhouettes of photographs of the volunteers. This experiment shows the accuracy we can expect when a person who is not an expert in anthropometry tries to use our method to build a digital clone of herself.

5.2 Face Experiments

This section validates our approach using a set of synthetic experiments on 3D human face shapes.

Training Data

The training set consists of 50 faces from the CAESAR database [19]. After parameterization, each face contains 6957 triangles. This experiment considers the seven measurements shown in Figure 1. Here, each dimension measures the geodesic distance between the pair of points shown in the same color. Note that several geodesic paths may overlap.

Synthetic Experiments

We generate new face shapes that aim to interpolate the measurements in S_{close} and S_k using our approach, feature analysis, and SGPLVM. Table 1 shows the errors of the generated faces for each method. We define an error for each dimension as the absolute value between the distance on the generated shape and the input measurement for this dimension, and use this to compute the error in mm as the average

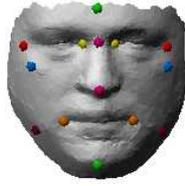


Fig. 1 Measurements used to define shape space for the face experiments.

	S_{close}	S_2	S_4
Feature Analysis	0.67	1.63	6.76
SGPLVM	5.35	5.28	10.88
Our Approach	0.68	2.00	7.27

Table 1 Comparison between our approach, feature analysis, and SGPLVM for the synthetic face experiments. The error in mm is the average error over all dimensions.

error over all dimensions. Note that feature analysis yields the lowest average errors and that SGPLVM yields the highest average errors for all experiments.

Several of the shapes predicted using feature analysis are clearly outside the shape space of human faces, as shown on the right of Figure 2. This is undesirable because human perception is sensitive with respect to unrealistic variations in face shape. In fact, most of the faces predicted using feature analysis in S_4 are clearly outside of the shape space of human faces.

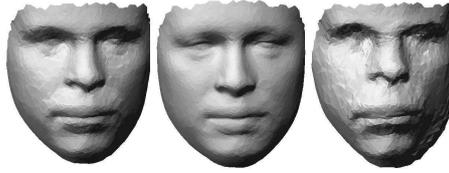


Fig. 2 Face shape predictions from a random sample in S_4 . Estimates from left to right using our approach, SGPLVM, and feature analysis.

It is desirable to have variability among the computed shapes. Both feature analysis and our approach predict significantly different shapes for all of the experiments. SGPLVM produces slightly different face shapes for different samples of S_{close} and S_2 , but for all samples in S_4 , the same face shape (visually similar to the mean shape) is produced.

In summary, our approach is the only one of the compared approaches that yields a variety of face shapes while always estimating shapes that are in the space of human faces. Furthermore, our approach yields lower errors than SGPLVM.

5.3 Body Experiments

This section validates our approach using a set of synthetic experiments and experiments based on real data.

Training Data

The training set consists of 360 bodies of the CAESAR database. After parameterization, each body contains 60000 triangles. This experiment considers 34 measurements: 14 Euclidean measurements, 4 circumference measurements (hip, waist, chest, and head circumferences) and 4 circumferences defined by four geodesic measurements each (knee and arm circumferences). Figure 3 shows the measurements on one shape.

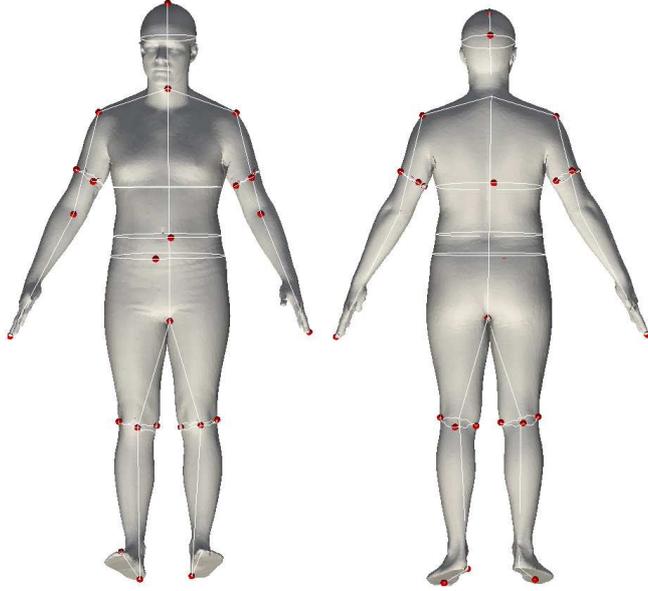


Fig. 3 Measurements used to define shape space for the body experiments. Red spheres show the specified points of the measurements and red lines show the measurements for this specific shape.

Synthetic Experiments

We generate new body shapes that aim to interpolate the new measurements in S_{close} and S_k using our approach, feature analysis, and SGPLVM. Table 2 shows the errors of the generated bodies for each method. The errors are computed in the same way as in the previous experiment. Note that our approach yields the lowest average errors for all experiments.

	S_{close}	S_2	S_4	S_{real}
Feature Analysis	3.29	14.26	36.86	3.09
SGPLVM	4.91	12.19	22.88	4.55
Our Approach	1.74	6.39	15.49	1.10

Table 2 Comparison between our approach, feature analysis, and SGPLVM for the body experiments. The error in mm is the average error over all dimensions.

For most samples in S_4 and for several samples in S_2 , feature analysis predicts shapes that are outside of the shape space of human bodies. For the samples in S_{close} , feature analysis yields visually pleasing shapes. SGPLVM always yields body shapes that are in the shape space of human bodies. However, for the samples in S_4 , the variation of the predicted body shapes is low. Some of the shapes predicted using S_4 with our approach contain small local artifacts due to the optimization with respect to the vertex coordinates. Note that setting λ to a higher value will reduce these artifacts at the cost of increased error. For S_2 and S_{close} , our approach always predicts globally and locally realistic body shapes. Figure 4 shows four of the body shapes in S_{close} predicted using our approach. We can see that a large variety of shapes can be computed.

Experiments Using Digital Measurements of Real Subjects

For this experiment, we consider 50 additional parameterized body shapes from the CAESAR database and use each of these models to digitally measure the 34 relevant measurements. We call this set of measurements S_{real} .

For the body shapes that aim to interpolate the measurements in S_{real} , Table 3 shows the errors of the Euclidean measurements and circumferences of the generated bodies. The error of the knee and arm circumferences is computed as the sum of the errors of the individual dimensions that contribute to

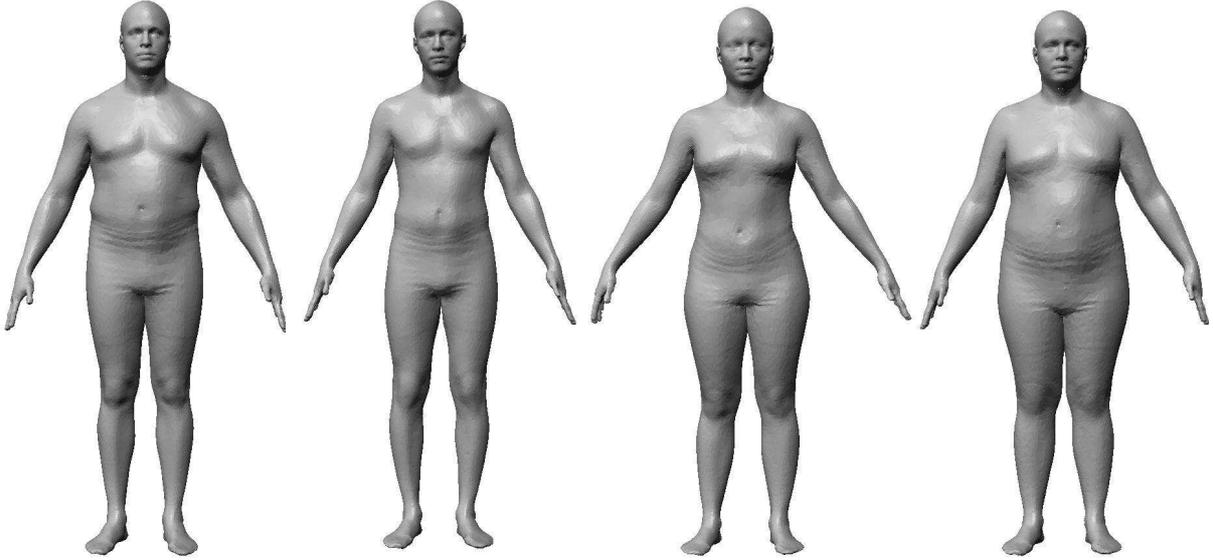


Fig. 4 Four of the shapes computed using our approach from random samples in S_{close} .

the circumference. We compute the average, standard deviations and maximum errors in mm for each dimension. Note that our approach yields the lowest average errors for all dimensions and the lowest maximum errors for most dimensions. Figure 5 shows the predictions for three different subjects. Each mesh is shown using a color coding corresponding to the signed distance from the parameterized body scan that was used to compute the measurements in S_{real} . Note that SGPLVM yields predictions with the highest errors and also predictions that are visually most dissimilar from the ground truth. The predictions obtained using feature analysis and our approach are of similar quality. Note however that the predictions obtained using our approach are more accurate in localized areas of the body such as the waist area of the subject shown in the first row. This is due to the deformation of the mesh that aims to obtain the correct waist circumference.

	Feature Analysis			SGPLVM			Our Approach		
	Avg.	Std.	Max.	Avg.	Std.	Max.	Avg.	Std.	Max.
Length Left Foot	3.31	2.54	10.01	4.04	2.91	11.47	2.38e-4	3.82e-4	2.02e-3
Length Left Lower Leg	1.90	1.67	6.50	4.34	2.74	11.30	4.93e-4	4.28e-4	1.82e-3
Length Left Upper Leg	2.08	1.57	6.79	4.09	3.44	12.07	7.79e-4	8.27e-4	3.48e-3
Length Right Upper Leg	2.38	1.77	7.67	4.14	4.02	25.71	9.56e-4	8.30e-4	3.62e-3
Length Right Lower Leg	1.99	1.43	5.48	5.00	3.99	18.91	5.57e-4	5.52e-4	2.49e-3
Length Right Foot	4.37	3.11	14.33	4.84	3.92	16.70	2.09e-4	3.90e-4	2.28e-3
Length Upper Body	1.71	1.22	4.90	6.93	6.16	27.67	2.99e-4	2.81e-4	9.91e-4
Right Shoulder-Neck Distance	1.75	1.60	6.99	3.06	2.31	9.95	1.71e-3	1.49e-3	6.65e-3
Length Right Upper Arm	2.21	1.66	7.78	3.73	2.65	10.98	6.11e-4	4.35e-4	1.74e-3
Length Right Lower Arm	3.17	1.99	8.89	4.36	4.39	16.95	4.42e-4	4.31e-4	1.79e-3
Left Shoulder-Neck Distance	2.37	1.93	9.85	2.82	2.38	12.77	2.12e-3	1.97e-3	8.50e-3
Length Left Upper Arm	4.24	2.82	12.36	5.16	3.87	19.17	7.16e-4	6.06e-4	2.20e-3
Length Left Lower Arm	2.46	1.86	8.35	4.89	3.36	14.14	4.79e-4	4.93e-4	2.21e-3
Head-Neck Distance	4.03	3.04	12.58	4.47	3.04	12.57	1.14e-3	1.09e-3	5.38e-3
Circumference Right Knee	6.61	2.84	14.12	6.78	3.16	12.82	5.48	2.85	15.36
Circumference Left Knee	6.97	4.75	24.40	7.12	3.74	18.40	6.02	4.22	24.44
Circumference Right Arm	7.98	3.57	18.35	8.55	4.02	22.62	4.62	2.56	11.17
Circumference Left Arm	7.38	3.42	16.66	8.42	4.23	18.13	5.07	2.92	16.36
Hip Circumference	10.20	6.16	26.83	18.97	18.51	110.30	5.29	6.77	40.01
Waist Circumference	8.99	6.99	32.13	15.51	11.08	45.19	1.24	1.18	6.11
Chest Circumference	8.32	5.12	22.50	15.87	14.61	89.53	1.29	1.05	4.08
Head Circumference	10.56	7.14	27.99	11.67	7.38	28.75	8.41	6.99	26.76

Table 3 Average, standard deviation and maximum errors (in mm) of measurements over 50 samples of S_{real} .

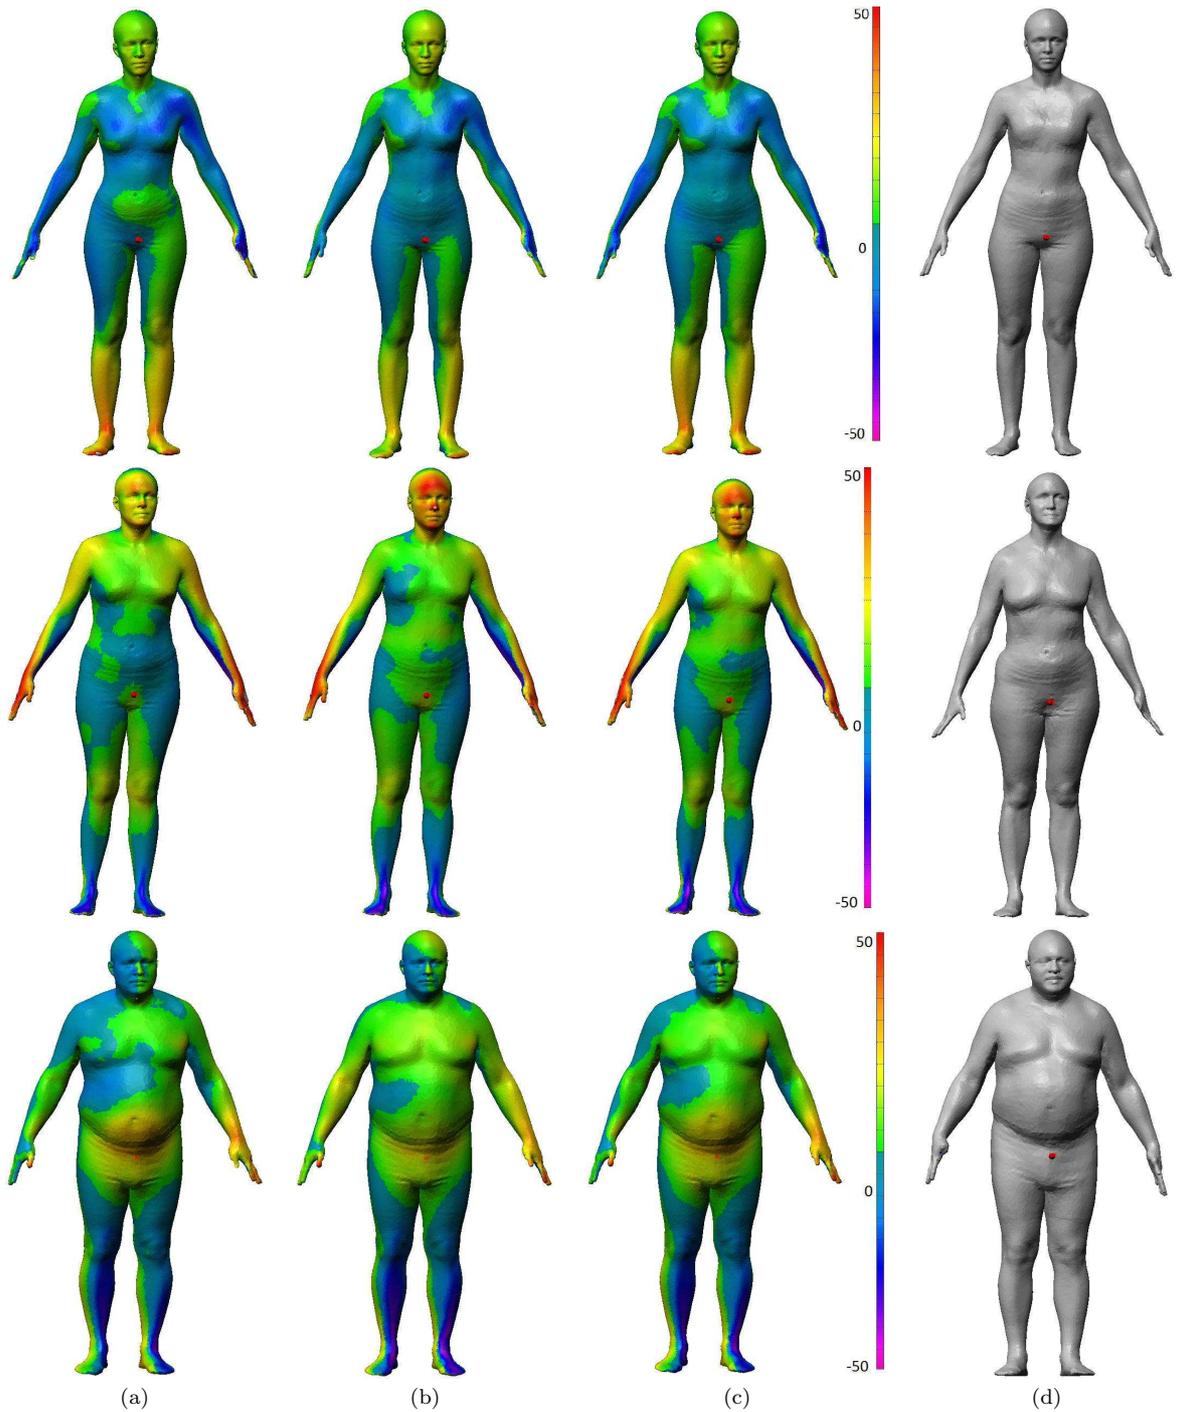


Fig. 5 Predicted body shapes: (a) shows the result using our approach, (b) shows the result using SGPLVM, (c) shows the result using feature analysis, and (d) shows the ground truth. Each mesh is shown using a color coding corresponding to the signed distance (in mm) from the parameterized body scan that was used to compute the measurements in S_{real} .

Experiments Using Real Measurements

This section presents an experiment that aims to predict the body shape from a set of measurements that were measured on real subjects using a tape measure. For this experiment, we consider the measurements of six volunteers who measured each other. Since the measurements of interest are difficult to measure for non-experts, we expect errors in the input data. The parameters are therefore set to $l = 3$ to restrict the shapes to the space of the training data and to $\lambda = 1.0$ to avoid fine fitting to inaccurate measurements.

Furthermore, we found that the measurements of the knee and arm circumferences were unreliable in practice, so we did not use these measurements to generate the reconstructions.

Figure 6 shows the results. Each column shows for one subject screen shots of the predicted body shape and silhouette images extracted from photographs of the clothed subject taken from similar viewpoints. Note that since the silhouettes are taken from clothed subjects, the silhouettes only approximate the silhouettes of the true body shapes (in fact, the silhouettes in the images should be wider than the true silhouettes). Furthermore, the 3D models and the silhouettes may show the subjects in slightly different postures. We can observe that even with the limited accuracy of both the measurements and the silhouettes, we obtain 3D body shape predictions that are visually consistent with the silhouette images.

Table 4 shows the errors of the measurements. Note that the errors are overall consistent with the errors reported for digitally measured distances in Table 3, albeit significantly higher. There are two main reasons for the higher errors. First, the measurements were conducted by non-experts. Gordon et al. [14] reported that measurements can be repeated with an accuracy of about 1cm when conducted by trained experts. In our scenario, we expect significantly higher measurement errors. Second, some of the measurements do not pass through salient surface points and are difficult to measure. Hence, the manual measurements and the digital measurements (i.e. the measurements computed from the 3D model by the algorithm) may not measure exactly the same distances. In spite of these problems, all measurement errors are lower than 6cm. This accuracy may suffice for applications, such as virtual games or virtual try-on.

	Subject 1	Subject 2	Subject 3	Subject 4	Subject 5	Subject 6
Length Left Foot	2.21	8.26	4.31	18.36	6.42	4.80
Length Left Lower Leg	8.38	0.67	4.60	3.11	6.37	2.58
Length Left Upper Leg	1.27	3.36	8.47	1.14	8.97	14.98
Length Right Upper Leg	6.51	0.86	12.02	4.21	13.05	7.52
Length Right Lower Leg	8.19	3.34	7.57	2.64	9.84	6.74
Length Right Foot	7.91	0.64	0.23	10.95	1.56	6.28
Length Upper Body	1.02	0.50	1.13	5.76	0.84	6.35
Right Shoulder-Neck Distance	4.68	2.45	3.22	14.81	4.58	5.12
Length Right Upper Arm	14.65	2.41	5.19	13.92	0.61	5.32
Length Right Lower Arm	3.84	0.47	1.21	4.50	0.18	13.10
Left Shoulder-Neck Distance	4.13	5.62	5.72	6.87	3.23	11.89
Length Left Upper Arm	12.93	0.87	5.74	15.89	2.02	15.22
Length Left Lower Arm	5.53	4.67	2.97	5.79	1.53	4.61
Head-Neck Distance	12.24	5.98	11.12	3.70	10.15	30.29
Hip Circumference	29.86	7.56	24.99	17.71	28.83	27.20
Waist Circumference	3.53	12.75	2.89	24.34	7.94	15.56
Chest Circumference	53.85	4.02	13.64	14.23	15.07	25.94
Head Circumference	39.66	6.91	40.02	3.57	30.48	19.71

Table 4 Errors (in mm) of measurements for each of the six subjects shown in Figure 6.

Experiments With a Small Training Set

Finally, we demonstrate the ability of our approach to predict shapes that are not covered well by the training data. Recall that this is relevant in practice because predicting human shapes based a small training sample avoids the complicated and expensive acquisition of large anthropometric databases. We use as training data a set of 35 of the bodies of the CAESAR database. All of the bodies in the training set are male bodies with typical body shape, see Figure 7.

We choose five additional male bodies with body shape variations not covered by the training data and we use each of these models to digitally measure the 34 relevant measurements. We then use these measurements for prediction. For this experiment, we set $l = 3$ since the shapes are not represented well by the training data. Furthermore, we set $s = 10$ to increase the accuracy of the prediction.

Figure 8 shows the results. The first row shows the bodies of the CAESAR database used for prediction, the second row shows the result using our approach, the third row shows the result using SGLVM, and the last row shows the result using feature analysis. Note that our approach predicts body shapes

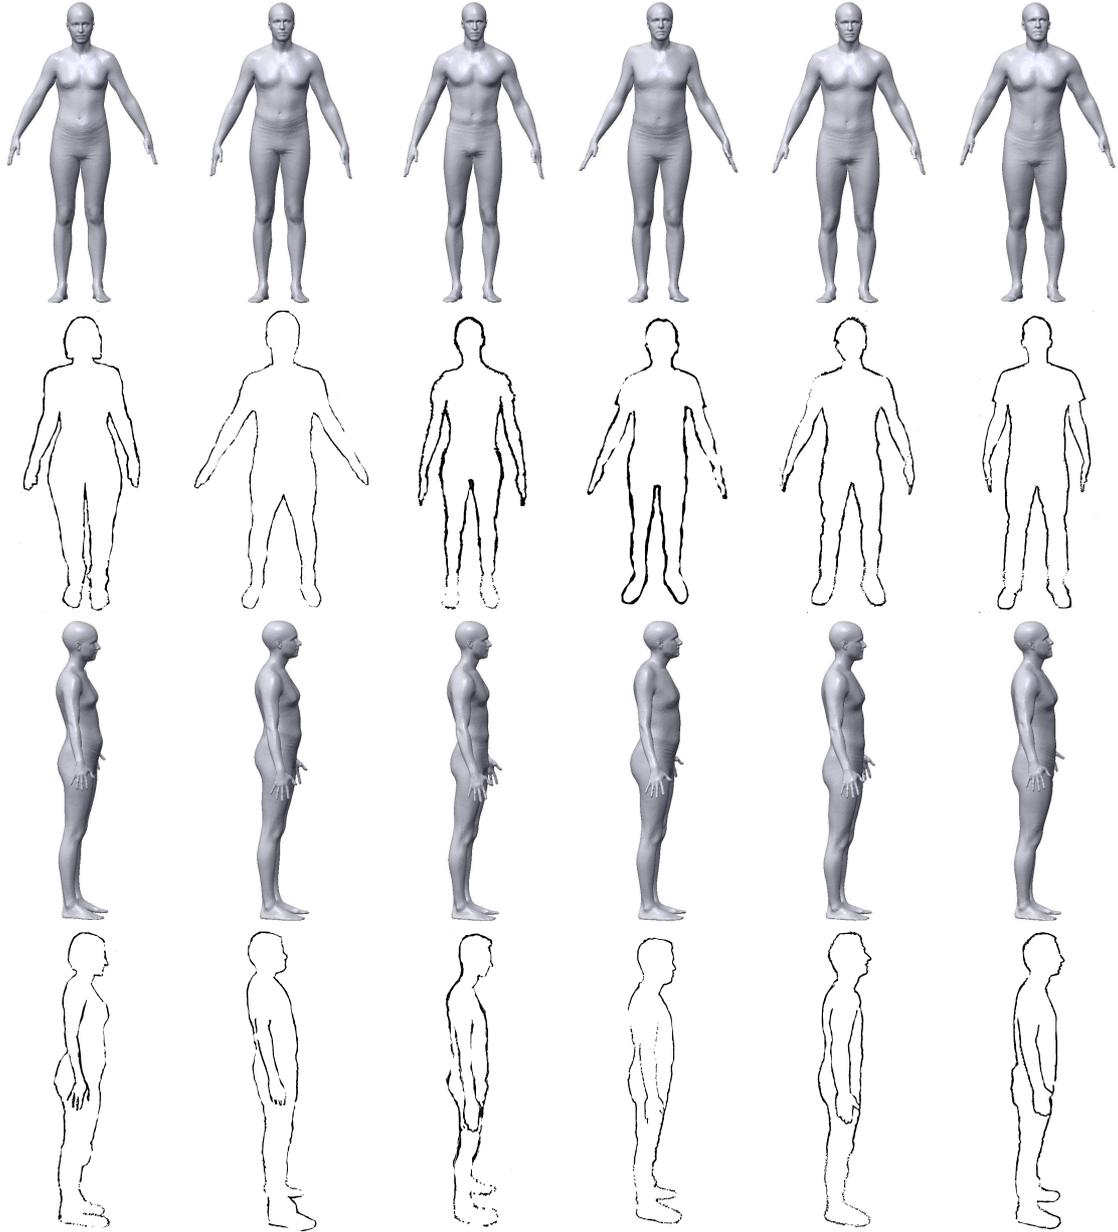


Fig. 6 Predicted body shapes from real measurements acquired by non-experts. Each column shows for one subject the reconstructed 3D model and two silhouette images extracted from photographs of the clothed subject.

that are similar to the ground truth for most shapes although a limited training set was used. For the body in the last column, the result using our approach is not as large as the ground truth. The reason is that the input body shape is far from the training data. The results using SGPLVM are all close to the mean shape and far from the ground truth. The results using feature analysis are outside of the shape space of human body shapes. This shows that our approach best handles the case where the predicted shapes have variations that are outside of the shape space spanned by the training data.

5.4 Summary

We conducted experiments based on human face and body shapes using both real and synthetic data. We conducted an extensive evaluation by estimating 600 face shapes and over 650 body shapes. In summary, we demonstrated experimentally that our method has the following properties.

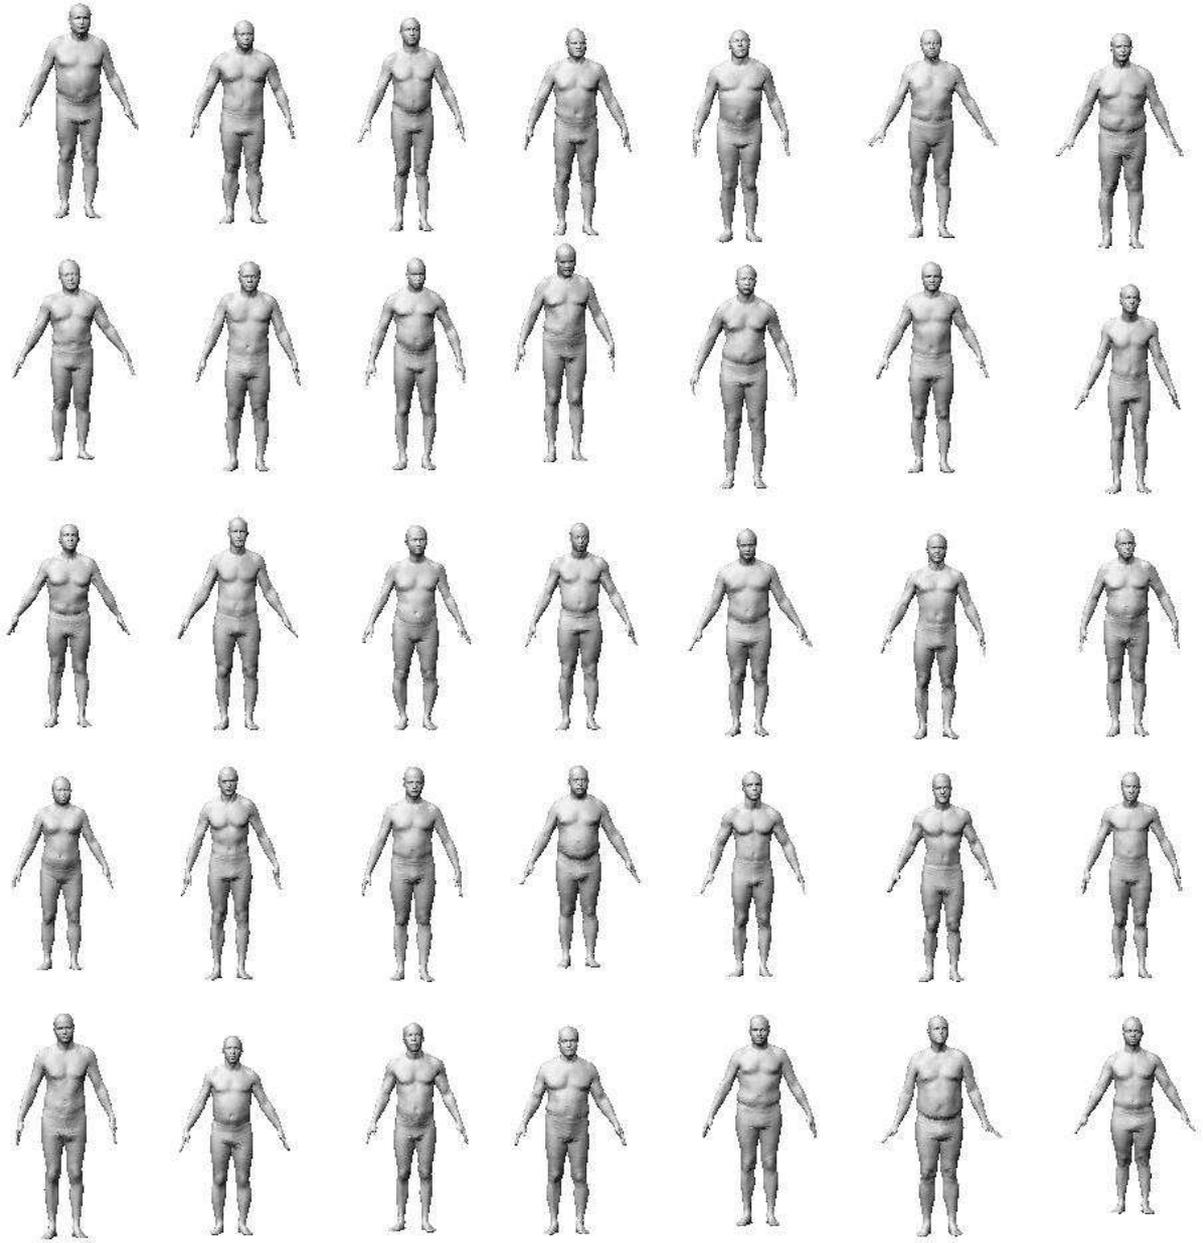


Fig. 7 35 bodies used for training.

- Our method estimates realistic body or face shapes that are outside of the learned shape space.
- Our method yields results that are more accurate than SGPLVM, and that, unlike feature analysis, always stay within the shape space of human body or face shapes.
- Our method accurately predicts local shape variations not present in the training set (which accounts for accurate predictions of the waist shapes in Figure 5).
- Our method yields visually pleasing results even when a small database is used for training. This is important in practice because the acquisition of a large database for training is costly.
- Even for inaccurate input measurements acquired by non-experts, our method is able to find visually consistent results.
- Our method is computationally efficient. The running time of the optimization steps is linear in the number of vertices. To predict one human body shape, our (non optimized) implementation takes about 35 seconds on a standard PC with a dual core CPU and 8GB of RAM.

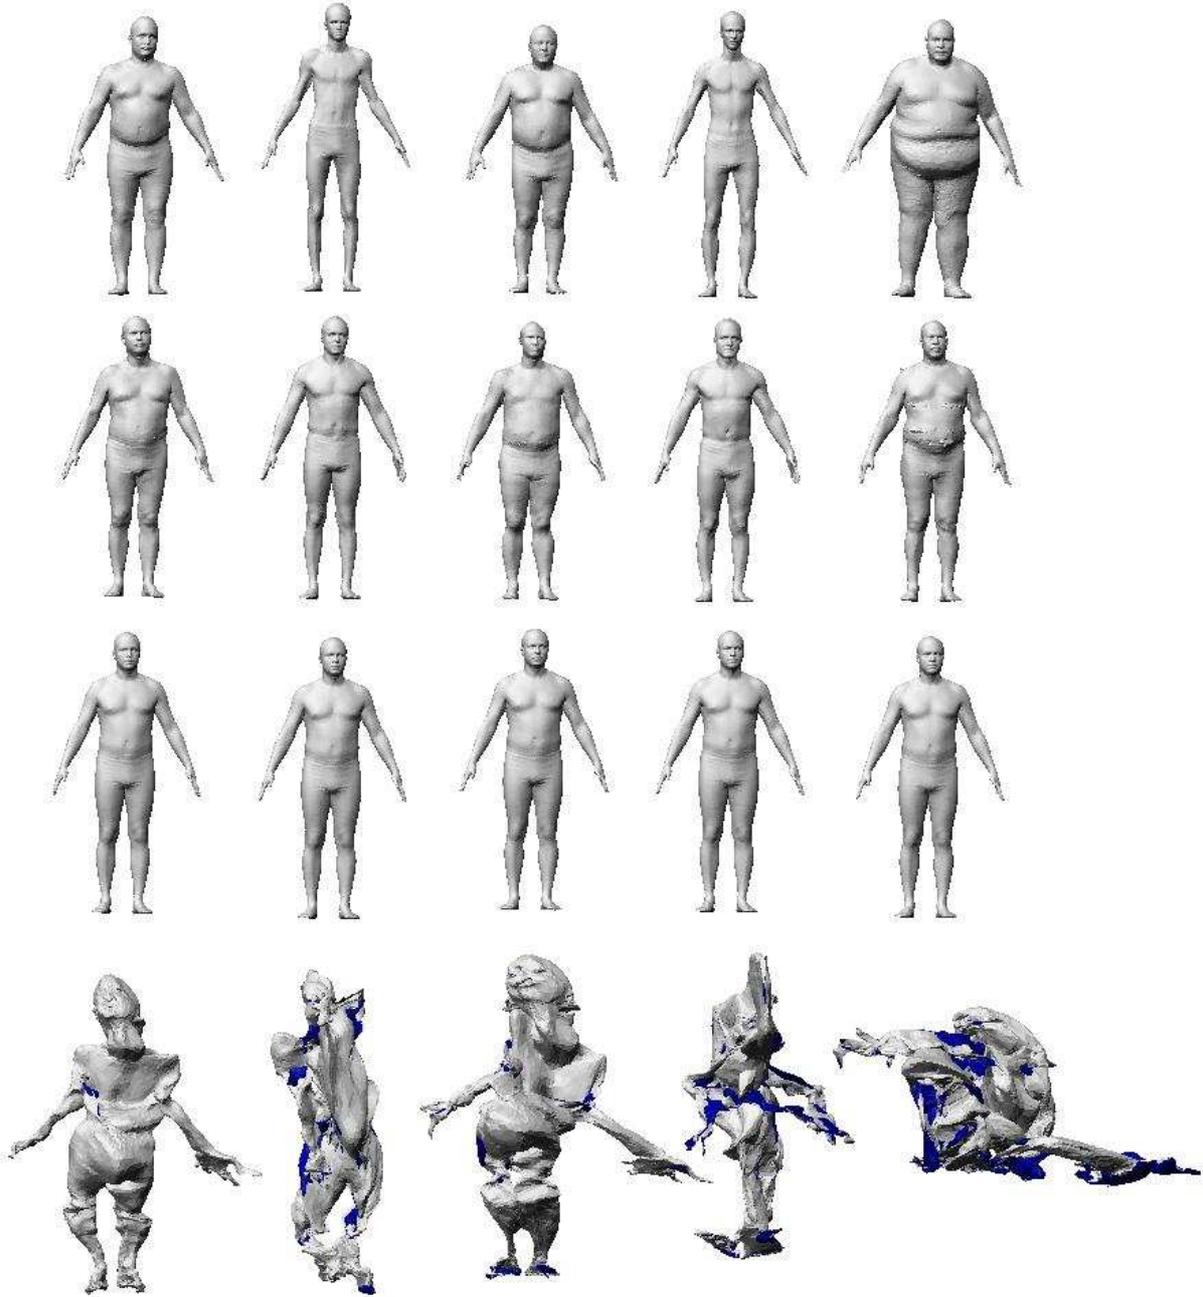


Fig. 8 Predicted body shapes based on the training data set shown in Figure 7. The first row shows the ground truth, the second row shows the result using our approach, the third row shows the result using SGPLVM, and the last row shows the result using feature analysis.

6 Conclusion

This paper presented a novel approach to estimate human body or face shapes based on simple linear measurements. The resulting shapes can be used in applications, such as ergonomic design or virtual try-on. Extensive experimental evaluation demonstrated that our proposed method estimates realistic body or face shapes that accurately represent the given measurements and that may have variations outside of the learned shape space. That is, unlike previous approaches, our approach can predict shape variations not present in the training data while maintaining a realistic human shape, which allows us to generate accurate shape estimates. Furthermore, we demonstrated that our approach interpolates

the given measurements more accurately than feature analysis or SGPLVM, especially if the database available for training is small.

Acknowledgments

We thank the volunteers for participating in the experiment. Furthermore, we thank Pengcheng Xi for helpful discussions and for providing us with the training data, Neil Lawrence for providing us with the SGPLVM code, and the anonymous reviewers for helpful comments. This work has partially been funded by the Cluster of Excellence *Multimodal Computing and Interaction* within the Excellence Initiative of the German Federal Government.

References

1. B. Allen, B. Curless, and Z. Popović. The space of human body shapes: reconstruction and parameterization from range scans. *ACM Transactions on Graphics*, 22(3):587–594, 2003. Proceedings of SIGGRAPH.
2. B. Allen, B. Curless, and Z. Popović. Exploring the space of human body shapes: Data-driven synthesis under anthropometric control. In *SAE Symposium on Digital Human Modeling for Design and Engineering*, paper# 2188, 2004.
3. D. Anguelov, P. Srinivasan, D. Koller, S. Thrun, J. Rodgers, and J. Davis. Scape: shape completion and animation of people. *ACM Transactions on Graphics*, 24(3):408–416, 2005. Proceedings of SIGGRAPH.
4. S.-Y. Baek and K. Lee. Parametric human body shape modeling framework for human-centered product design. *Computer-Aided Design*, 44(1):56–67, 2012.
5. C. B. Barber, D. P. Dobkin, and H. Huhdanpaa. The quickhull algorithm for convex hulls. *ACM Transactions on Mathematical Software*, 22(4):469–483, 1996.
6. V. Blanz and T. Vetter. A morphable model for the synthesis of 3d faces. In *Proceedings of SIGGRAPH*, pages 187–194, 1999.
7. J. Boisvert, C. Shu, S. Wuhrer, and P. Xi. Three-dimensional human shape inference from silhouettes: Reconstruction and validation. *Machine Vision and Applications*, To appear.
8. Y. Chen and R. Cipolla. Learning shape priors for single view reconstruction. In *Workshop on 3D Imaging and Modelling*, pages 1425–1432, 2009.
9. C.-H. Chu, Y.-T. Tsai, C. C. Wang, and T.-H. Kwok. Exemplar-based statistical model for semantic parametric design of human body. *Computers in Industry*, 61(6):541–549, 2010.
10. D. DeCarlo, D. Metaxas, and M. Stone. An anthropometric face model using variational techniques. In *Proceedings of SIGGRAPH*, pages 67–74, 1998.
11. E. W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1(1):269–271, 1959.
12. I. Dryden and K. Mardia. *Statistical Shape Analysis*. Wiley, 2002.
13. C. H. Ek, P. H. S. Torr, and N. D. Lawrence. Gaussian process latent variable models for human pose estimation. In *Workshop on Machine Learning for Multimodal Interaction*, pages 132–143, 2007.
14. C. C. Gordon, T. Churchill, C. E. Clauser, B. Bradtmiller, J. T. McConville, I. Tebberets, and R. A. Walker. Anthropometric survey of U.S. army personnel: Methods and summary statistics 1988. Technical report, U.S. Army Natick Research, Development, and Engineering Center, 1989.
15. P. Guan, A. Weiss, A. O. Balan, and M. J. Black. Estimating human shape and pose from a single image. In *International Conference on Computer Vision*, pages 1381–1388, 2009.
16. N. Hasler, H. Ackermann, B. Rosenhahn, T. Thormählen, and H.-P. Seidel. Multilinear pose and body shape estimation of dressed subjects from image sets. In *Conference on Computer Vision and Pattern Recognition*, pages 1823–1830, 2010.
17. N. Hasler, C. Stoll, M. Sunkel, B. Rosenhahn, and H.-P. Seidel. A statistical model of human pose and body shape. In P. Dutré and M. Stamminger, editors, *Computer Graphics Forum*, 28(2):337–346, 2009.
18. D. C. Liu and J. Nocedal. On the limited memory method for large scale optimization. *Mathematical Programming*, 45(3):503–528, 1989.
19. K. Robinette, H. Daanen, and E. Paquet. The CAESAR project: A 3-D surface anthropometry survey. In *3-D Digital Imaging and Modeling*, pages 180–186, 1999.
20. C. Schmaltz, B. Rosenhahn, T. Brox and J. Weickert. Region-based pose tracking with occlusions using 3D models. In *Machine Vision and Applications*, 23(3):557–577, 2012.
21. H. Seo and N. Magnenat-Thalmann. An automatic modeling of human bodies from sizing parameters. In *Proceedings of the 2003 Symposium on Interactive 3D Graphics*, pages 19–26, 2003.
22. H. Seo, Y. I. Yeo, and K. Wohn. 3d body reconstruction from photos based on range scan. *Technologies for E-Learning and Digital Entertainment*, pages 849–860, 2006.
23. A. P. Shon, K. Grochow, A. Hertzmann, and R. P. N. Rao. Learning shared latent structure for image synthesis and robotic imitation. In *Neural Information Processing Systems*, pages 1233–1240, 2005.
24. O. van Kaick, H. Zhang, G. Hamarneh, and D. Cohen-Or. A survey on shape correspondence. *Computer Graphics Forum*, 30(6):1681–1707, 2011.
25. C. C. Wang. Parameterization and parametric design of mannequins. *Computer Aided Design*, 37(1):83–98, 2005.
26. W. Wei, X. Luo, and Z. Li. Layer-based mannequin reconstruction and parameterization from 3d range data. In *Advances in Geometric Modeling and Processing*, pages 498–504, 2008.

27. A. Weiss, D. Hirshberg, and M. Black. Home 3d body scans from noisy image and range data. In *International Conference on Computer Vision*, pages 1951–1958, 2011.
28. S. Wuhrer, P. Xi, and C. Shu. Human Shape Correspondence with Automatically Predicted Landmarks. In *Machine Vision and Applications*, 23(4):821–830, 2012.
29. P. Xi, W.-S. Lee, and C. Shu. Analysis of segmented human body scans. In *Graphics Interface*, pages 19–26, 2007.