

Stefanie Wuhrer · Pengcheng Xi · Chang Shu

Human Shape Correspondence with Automatically Predicted Landmarks

Abstract We consider the problem of computing accurate point-to-point correspondences among a set of human bodies in similar posture using a landmark-free approach. The approach learns the locations of the anthropometric landmarks present in a database of human models in similar postures and uses this knowledge to automatically predict the locations of these anthropometric landmarks on a newly available scan. The predicted landmarks are then used to compute point-to-point correspondences between a template model and the newly available scan. This study conducts a large-scale evaluation to examine the accuracy of the computed correspondences. Furthermore, we show that the correspondences are accurate enough for the application of motion transfer.

1 Introduction

We consider the problem of computing accurate point-to-point correspondences among a set of human shapes in similar posture without using landmarks. The shapes of the human bodies are represented by possibly incomplete triangular meshes, which can be acquired by 3D sensing devices such as laser or structured-light body scanners. This problem arises from building statistical models from a database of 3D scans. In order to conduct statistical analysis of 3D shapes, the raw scans have to be parameterized in such a way that likewise anatomical parts correspond across the models [13]. This problem is also known as groupwise registration in medical imaging [12]. A statistical model built from the dense correspondence allows us to study the shape variability of a population and has been proven to be a powerful tool for mesh deformation [4], human shape reconstruction [16], and ergonomic design [23]. The correspondence problem is also essential to object recognition [19].

While many approaches have been proposed to compute point-to-point correspondences [29], only few of them have

been applied to statistical model building and shape analysis. Here accurate correspondence is important because it directly affects the quality of the statistical model. Computing the correspondence among a set of human shapes is further complicated by the intrinsic symmetry of the human body shape and by the shape variation across a population. Hence, fully-automatic correspondence methods that exploit the intrinsic geometry of the shape [18, 33] may fail when corresponding humans that differ in shape.

In a typical 3D anthropometry survey, the human subjects are asked to maintain a standard posture like the one shown in Figure 2. However, in practice slight posture variations are inevitable. The correspondence problem is further complicated by the noisy and incomplete nature of the scan data. Allen et al. [2] proposed an effective method in which a template mesh is deformed to fit each data model, thus simultaneously solving the correspondence problem and filling the holes smoothly. Anthropometric landmarks placed on the human bodies prior to scanning are used to guide the fitting process. Since the anthropometric landmarks are stable positions on the human body, they provide accurate seed points for corresponding the rest of the points. This approach has been extended by various authors [4, 17]. However, all of these extensions require a set of manually placed marker positions on the body.

Unfortunately, manually placing the landmarks is a tedious task and it is impractical to use routinely in large surveys where several thousands of subjects are typically scanned. Ben Azouz et al. [6] proposed an automatic method to compute the landmarks from the scans. However, the method has not yet been demonstrated to be effective for computing dense correspondences. Detailed analysis and evaluation in terms of the quality of the correspondence have to be provided.

The purpose of this paper is to provide a fully automatic solution to the problem of computing point-to-point correspondences among a set of human shapes in similar posture. To the best of our knowledge, previous methods for analyzing the human body shape use known landmark positions when computing the correspondences. Furthermore, methods that compute correspondences for arbitrary shapes in a

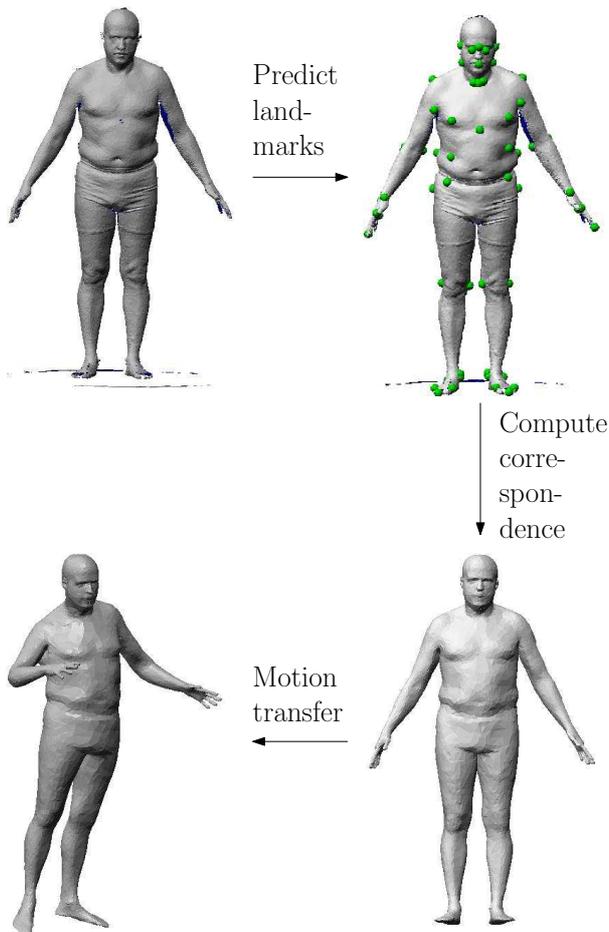


Fig. 1 Overview of the proposed method.

fully automatic manner [18, 33] take advantage of the intrinsic geometry of the shapes and may fail for human shapes due to intrinsic symmetries. We review previous methods in detail in Section 2.

We integrate an extension of the landmark prediction method [6] and the template fitting method [2]. The landmark prediction starts with learning the characteristics and locations of the landmarks on a human model in standing posture using a database of human models. This information is used to predict the landmark positions on a human shape in similar posture. Figure 2 shows a model in standing posture where the landmark positions are shown in red. Section 3 summarizes how to predict the landmarks. Special care is taken to outline our extensions to the previous method by Ben Azouz et al., which lead to a more efficient algorithm. The predicted landmarks are then used for deforming a template model to fit the mesh data [31] as outlined in Section 4. In our experiments, we use the Civilian American and European Surface Anthropometry Resource (CAESAR) database [26].

An area in which the literature is particularly weak is the evaluation of the quality of the correspondence. Often, only visual results are given and only small numbers of models

are tested. In Section 5, we conduct a large-scale evaluation of the accuracy of the computed correspondences in the context of shape analysis. Finally, in Section 6, to further demonstrate the quality of the correspondence, we apply this approach to motion transfer. Figure 1 summarizes the process.

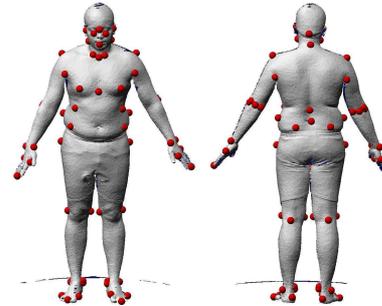


Fig. 2 Model of the CAESAR database in standing posture with landmarks shown in red.

2 Related Work

Computing dense point-to-point correspondences between two possibly deformed surfaces has received considerable attention in recent years [29]. Although many algorithms have been developed to solve the correspondence problem, only few of these algorithms are suitable for statistical model building and shape analysis.

When considering the application of statistical model building, we need to solve for the correspondence and for a smooth transformation between a group of shapes. The reason is that solving only for correspondence may result in high levels of noise, which in turn affects the statistical model significantly.

Approaches that solve the correspondence problem by aligning two shapes using a transformation that is approximately rigid [27, 15, 1] are not suitable to align a set of human shapes in similar posture due to non-rigid shape and slight posture changes. Hence, we focus our attention on approaches that take non-rigid transformations into account.

Blanz and Vetter [9] use a set of landmarks to compute the correspondence between pairs of human faces. Allen et al. [2] use a set of landmarks and a template model to deform a template model to human shapes in similar postures. Amberg et al. [3] improve upon this method by formulating the underlying optimization problem using an iterative linear framework that is shown to yield better convergence. Angelov et al. [4] extend Allen et al.'s approach to work for varying postures. The approach computes the correspondences between a large database of humans and uses the result to build an animated surface model of a moving person. The database contains one subject in multiple postures and the remaining subjects in the standing posture of the CAESAR database. Hasler et al. [17] improve the approach by

using fewer landmark positions and by using a database containing many subjects in multiple postures. While all of these approaches have been applied to shape analysis applications and yield animations of high quality, landmark positions are required. Pauly et al. [24] compute the correspondences and the transformation between multiple views of a scan for the application of scan completion using a small set of landmark positions.

Recently, several landmark-free approaches have been proposed. Huang et al. [18] proceeds by iteratively alternating between a correspondence optimization and a deformation optimization. The approach can be viewed as an extension of the Iterative Closest Point algorithm (ICP) [8] that is often used to solve the rigid correspondence problem. The method is shown to perform well if the two meshes are initially well aligned. If the alignment is poor, the method fails. Huang et al. show that the obtained correspondences yield visually pleasing shape interpolations. The main drawback of this method is that it relies heavily on non-intuitive user-defined parameters. This makes the method hard to use. Zhang et al. [33] propose a technique that solves the correspondence problem by finding a small set of features and by choosing the best feature correspondence as the one that minimizes a deformation energy. To improve the efficiency of the algorithm, the tree of all matching features is pruned if the features are too dissimilar. Nonetheless, the algorithm is not as efficient as the algorithm of Huang et al. [18]. Once the feature correspondences are computed, the full correspondence is found by deforming the full mesh based on the feature points. The main drawback of this method is the computational inefficiency. Results are only demonstrated for models with less than 4000 vertices. Furthermore, like the method of Huang et al., the tree pruning relies heavily on non-intuitive user-defined parameters. Chang and Zwicker [11] use a reduced deformable model to compute the correspondence and the transformation between two surfaces. Instead of operating on the surface directly, the approach needs to convert the surface into a voxel grid. This is computationally expensive. Furthermore, this step introduces the use of several input parameters. While all of these methods are landmark-free, they require a set of non-intuitive user-specified input parameters.

Methods that require neither landmark positions nor a set of user-specified input parameters have been proposed for motion capture [10, 14]. The methods assume that the same shape was captured in several gradually changing poses and use this information to learn a deformation model. In our application, this type of input data is not available. Li et al. [21] propose an approach to register pairs of range images without using any landmark positions or input parameters. While the method is shown to perform well, the method makes use of the fact that each surface is a terrain and can be parameterized by a projection to a plane. Since our aim is to register the surface of full human bodies, this method cannot be applied.

In medical imaging, two methods were proposed recently [12, 32] that solve the correspondence problem using a landmark

free approach that operated on the unit sphere. First, each shape is mapped to the unit sphere. Second, the charts on the unit sphere are aligned to minimize an energy. Davies et al. [12] align the charts until the description length is minimized while Yeo et al. [32] align the charts to minimize a diffeomorphic energy. Both approaches require as input manifold meshes with spherical topology. That is, surfaces with holes cannot be taken into consideration. Since we aim to compute correspondences between incomplete laser scans, these approaches cannot be applied.

We propose a landmark-free approach to compute the correspondence between a set of human shapes in similar posture using a template model. We apply the resulting correspondences to statistical model building and to animate the models using motion transfer.

3 Predicting Landmarks

This section outlines how to automatically predict a set of anthropometric landmarks on a 3D scan of a human body.

First, we summarize the approach by Ben Azouz et al. [6] to automatically locate 73 labelled anthropometric landmarks on 3D scans of human models in standing posture. The approach is based on statistical learning and proceeds by using an anthropometric database to learn the characteristics and the locations of the landmarks using a pairwise Markov network. The learned information is then used to predict the landmarks on a newly available scan S .

The 73 landmarks are connected by the manually chosen graph structure shown in Figure 3. In this graph, each landmark is a node and certain pairs of nodes are connected using a set of edges. The approach uses a database of human scans in standing posture to learn the location, Spin image [20], and surface curvature of each node and the length and orientation of each edge. In this step, the node and edge characteristics are modeled as independent Gaussian distributions and we learn their means and standard deviations using maximum likelihood estimation. The learned information is combined in a pairwise Markov network, where the learned Gaussian distributions are used as node and edge potentials.

The data set used for training consists of 200 male scans of the CAESAR database. Note that for each of the models in the CAESAR database, the locations of the 73 anthropometric landmarks are known. Before the learning step, the number of triangles of each scan is reduced to 50000.

When a new scan S becomes available, the approach predicts the 73 landmark positions by performing probabilistic inference on the learned Markov network. Before the prediction, the number of triangles of each scan is reduced to 50000. To perform probabilistic inference, the approach needs to limit the number of possible candidate locations, or labels, for each landmark. Let l_0, \dots, l_{72} denote the landmarks. To limit the number of labels for landmark l_i , the approach only considers vertices that are located inside the sphere with radius proportional to σ_i centered at μ_i ,

where μ_i is the learned mean location of l_i and where σ_i is the learned covariance of l_i . The inference problem is then solved using loopy belief propagation [25]. The loopy belief propagation algorithm is iterated at most 20 times.

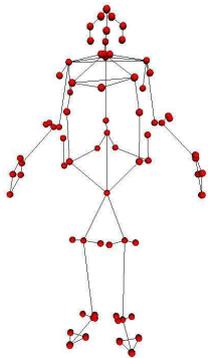


Fig. 3 Structure of the landmark graph.

Second, we outline our modifications to the approach by Ben Azouz et al. [6]. Our training phase proceeds in the same way as the previous approach and the same data is used for training. We modify the prediction phase of the algorithm to obtain a faster approach, which allows us to predict landmarks on a large database. To achieve this, we reduce the number of labels for landmark l_i . This has a large influence on the running time of the algorithm because the complexity of the algorithm is quadratic in the number of labels. Note that when considering as labels all of the vertices that are located inside the sphere with radius proportional to σ_i centered at μ_i , landmarks in areas of the mesh with high resolution that have a large variation in the training set have many associated labels. To avoid this problem, we only consider as labels for l_i the 200 vertices of the scan that are closest to μ_i . We find these labels using a k -nearest neighbor search in a kd-tree (we use the approximate nearest neighbor library [5] in our implementation). Furthermore, we found that iterating the loopy belief propagation algorithm five times yields sufficient results. Hence, we iterate the loopy belief propagation algorithm at most five times. These modifications yield a significant speed-up over the approach originally proposed by Ben Azouz et al. at the cost of a decrease in accuracy.

We predict the landmarks on 500 male scans of the CAESAR database that were not used for training. We know the true anthropometric landmarks for these scans. Hence, we can measure the accuracy of the predicted landmarks as the Euclidean distance between the predicted landmarks and the true landmarks. The accuracy is shown in Table 1. Note that the average error of most landmarks is less than $3cm$. The landmarks we use as reference are subject to uncertainties since different experts place the landmarks at slightly different positions and that for most traditional measurements, the acceptable error is around $1cm$ [6].

Note that the predictions of landmarks 26 and 27 have a large average error. The reason is that these landmarks are

located in an area of the lower back that has few surface features. Since these areas are not crucial during our subsequent correspondence step, we ignore these two landmarks in the following. Landmark 73, located at the crotch, also has a large average error. This landmark is located in an area of the body that is often not visible to the scanner. Hence, for many scans, this surface area is missing. This results in a large error since the statistical inference method only considers surface vertices as possible locations for the landmark. As this landmark is crucial in the subsequent correspondence step, we use it in spite of its large average error.

Figure 4 shows the ground truth landmarks and the predicted landmarks on two models of the CAESAR database. The predicted landmarks are shown in green and the ground truth landmarks are shown in red.

4 Correspondence

This section summarizes how to use the predicted landmarks to find a surface correspondence. We use the approach by Xi et al. [31] to compute a surface correspondence using an alignment procedure based on the predicted landmarks. The approach aims to align a template model T to the scan S using a two-step alignment procedure. The template model T and its landmarks are known. In our work, T contains 20000 triangles.

First, the approach computes a radial basis function (RBF) using the known landmarks $L^{(T)}$ of T and the predicted landmarks $L^{(S)}$ on S . Given two sets of n landmarks $L^{(T)}$ and $L^{(S)}$ (here $n = 71$), we aim to compute a set of three RBFs, one for each coordinate, such that $L_i^{(T)} = f(L_i^{(S)})$, $i = 1, \dots, n$, where $L_i^{(T)}$ is the i -th landmark on T and $L_i^{(S)}$ is the i -th landmark on S . For RBFs, the mapping has the property $f(L_i^{(S)}) = \sum_{j=1}^n w_j \Phi_j(L_i^{(S)})$, where Φ_j are a set of n basis functions and w_j are weights. We use Gaussian basis functions, which means $\Phi_j(r) = \exp(-r^2/c^2)$, where c is a constant (here $c = 1.0$). The approach deforms all of T 's vertices using the RBFs to obtain a rough alignment of T to S .

Second, the approach deforms T using a non-linear optimization method similar to the one used by Allen et al. [2]. The aim of this minimization is to move the vertices of T towards their closest neighbor on S while enforcing neighboring vertices to move using similar transformations. We use the limited-memory Broyden-Fletcher-Goldfarb-Shanno quasi-Newton approach [22] to solve the optimization problem.

5 Evaluation

In this section we conduct an evaluation of the accuracy of the computed correspondences over a large database of 500 scans. We run our single-threaded implementation on an Intel Duocore 3GHz machine with 8GB of RAM. For

Landmark	Average (mm)	Standard Deviation (mm)	Landmark	Average (mm)	Standard Deviation (mm)
1 Sellion	12	0.4	38 Rt. Dactylion	19	0.36
2 Rt. Infraorbitale	15	0.12	39 Rt. Ulnar Styloid	15	0.41
3 Lt. Infraorbitale	16	0.3	40 Rt. Metacarpal-Phal. V	16	0.31
4 Supramenton	17	0.18	41 Lt. Acromion	17	0.35
5 Rt. Tragion	17	0.15	42 Lt. Axilla, Ant	23	0.069
6 Rt. Gonion	18	0.1	43 Lt. Radial Styloid	14	0.24
7 Lt. Tragion	17	0.56	44 Lt. Axilla, Post.	18	0.6
8 Lt. Gonion	18	0.3	45 Lt. Olecranon	16	0.33
9 Nuchale	27	0.27	46 Lt. Humeral Lateral Epicon	18	0.27
10 Rt. Clavicale	16	0.32	47 Lt. Humeral Medial Epicon	18	0.41
11 Suprasternale	18	0.21	48 Lt. Radiale	18	0.041
12 Lt. Clavicale	17	0.55	49 Lt. Metacarpal-Phal. II	15	0.3
13 Rt. Thelion/Bustpoint	9.5	0.28	50 Lt. Dactylion	21	0.93
14 Lt. Thelion/Bustpoint	10	0.16	51 Lt. Ulnar Styloid	15	0.11
15 Substernale	21	0.44	52 Lt. Metacarpal-Phal. V	15	0.061
16 Rt. 10th Rib	31	0.95	53 Rt. Knee Crease	14	0.33
17 Rt. ASIS	24	0.19	54 Rt. Femoral Lateral Epicon	18	0.63
18 Lt. 10th Rib	29	0.62	55 Rt. Femoral Medial Epicon	17	0.077
19 Lt. ASIS	23	0.36	56 Rt. Metatarsal-Phal. V	12	0.038
20 Rt. Iliocristale	19	0.096	57 Rt. Lateral Malleolus	13	0.017
21 Rt. Trochanterion	17	0.36	58 Rt. Medial Malleolus	17	0.065
22 Lt. Iliocristale	15	0.29	59 Rt. Sphyrion	14	0.3
23 Lt. Trochanterion	12	0.76	60 Rt. Metatarsal-Phal. I	13	0.056
24 Cervicale	17	0.18	61 Rt. Calcaneus, Post.	13	0.066
25 10th Rib Midspine	23	0.73	62 Rt. Digit II	9.8	0.41
26 Rt. PSIS	38	1.2	63 Lt. Knee Crease	13	0.15
27 Lt. PSIS	36	0.69	64 Lt. Femoral Lateral Epicon	18	0.3
28 Waist, Preferred, Post.	21	0.78	65 Lt. Femoral Medial Epicon	20	0.17
29 Rt. Acromion	16	0.23	66 Lt. Metatarsal-Phal. V	11	0.39
30 Rt. Axilla, Ant	24	0.7	67 Lt. Lateral Malleolus	14	0.28
31 Rt. Radial Styloid	15	0.53	68 Lt. Medial Malleolus	11	0.009
32 Rt. Axilla, Post.	18	1.1	69 Lt. Sphyrion	10	0.18
33 Rt. Olecranon	16	0.19	70 Lt. Metatarsal-Phal. I	11	0.059
34 Rt. Humeral Lateral Epicon	18	0.56	71 Lt. Calcaneus, Post.	12	0.18
35 Rt. Humeral Medial Epicon	18	0.27	72 Lt. Digit II	12	0.34
36 Rt. Radiale	19	0.055	73 Crotch	44	0.071
37 Rt. Metacarpal Phal. II	17	0.33			

Table 1 Error of localization computed over 500 test human scans.

one scan, it takes about 148s to predict the landmarks and about 10s to compute the correspondence. While we know the ground truth of the landmark positions, no ground truth is available for the other parts of the surfaces. Hence, we cannot evaluate the accuracy of the correspondences with respect to a ground truth correspondence. However, we can still evaluate the quality of the correspondence by evaluating the statistical model created from the correspondence. We use three measures suggested by Styner et al. [28]: compactness, generalization, and specificity.

To compute these measures, we compute a statistical model of the corresponded data using principal component analysis (PCA). A compact model has little variance and it requires few principal components to describe the model. Compactness is defined as $C(m) = \sum_{i=1}^m \lambda_i$, where m is the number of principal components that are kept and where λ_i is the i -th eigenvalue. The compactness error is computed as $\sigma_{C(m)} = \sum_{i=1}^m \sqrt{2/n} \lambda_i$.

A general model has the ability to describe instances that are not in the training set. We compute a statistical model with all but one of the corresponded scans of the training set and fit the model to the excluded scan. Generalization $G(m)$ measures the accuracy of the description of the excluded scan. It is computed as the average error over the complete set of trials. The standard error is computed using the standard deviation σ of errors of the complete set of trials as $\sigma_{G(m)} = \sigma / \sqrt{n - 1}$.

A specific model only generates shapes that are similar to the ones in the training set. Specificity $S(m)$ is defined as the average distance of S_{rand} to its nearest neighbor in the training set, where S_{rand} is a shape computed using the statistical model from a uniformly distributed random sample in PCA shape space. The standard error is computed using the standard deviation σ of the distances as $\sigma_{S(m)} = \sigma / \sqrt{n}$.

We compute these measures for the correspondences obtained using the ground truth landmarks and for the correspondences obtained using the predicted landmarks. Note

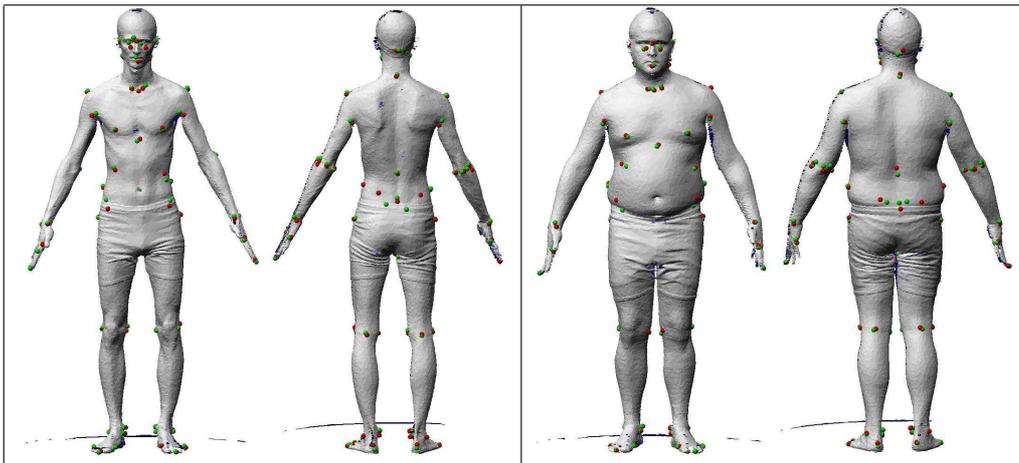


Fig. 4 Predicted landmarks (green) and ground truth landmarks (red) on two models of the CAESAR database.

that the correspondences obtained using the ground truth landmarks correspond to the result obtained by Allen et al.’s algorithm [2]. Since we focus on an automatic method, we do not compare to any other semi-automatic methods that require manually placed landmarks [4, 3, 17]. Furthermore, we do not compare to any automatic methods based on aligning intrinsic geometries [18, 33] since these methods may lead to misalignments due to the symmetries in intrinsic geometry of the human body and due to the variations between different human shapes.

Figure 5 shows the results. In all three graphs, the x -axis shows the number m of principal components and the y -axis shows the values of the error measures $\sigma_{C(m)}$, $\sigma_{G(m)}$, and $\sigma_{S(m)}$. For specificity, we consider 100 random samples. While the PCA space learned using the predicted landmarks is slightly more compact and more specific than the PCA space learned using the ground truth landmarks, it is slightly less general. Overall, both models are similar.

The quality of the correspondences is visualized in Figure 6. We manually applied a texture to the template model (left of Figure 6) and transferred the texture using the correspondences obtained with the predicted landmarks. We can see that the texture map is preserved for models with different body shape.

6 Application to Motion Transfer

This section shows that the accuracy of the computed correspondences yields visually pleasing results for the application of motion transfer. We fit a skeleton $B^{(T)}$ to T and we assign weights $W^{(T)}$ to each vertex on the surface of T using the approach by Baran and Popovic [7], where $W^{(T)}$ contains one weight per bone in the skeleton for each vertex of T . This allows to animate the template model T using a linear blend skinning method. That is, given all the transformations \mathbf{T}_i of the i -th bone of $B^{(T)}$, a vertex v_j of T is transformed to $v_j^* = \sum_i W_{j,i}^{(T)} \mathbf{T}_i(v_j)$, where $W_{j,i}^{(T)}$ is the weight

for the i -th bone and the j -th vertex of T . An animation is shown in the first row of Figure 7.

To animate a scan S , we first parameterize the scan by predicting landmarks and deforming T to fit the data in S as outlined above. We use the known point-to-point correspondences between the parameterized scan and T to compute an RBF deformation as in Section 4. We deform the skeleton $B^{(T)}$ using the RBF to obtain a new skeleton $B^{(S)}$ aligned to the parameterized scan. To animate the parameterized scan, we use the linear blend skinning method with the skeleton $B^{(S)}$ and the weights $W^{(T)}$. Some results of this animation are shown in Figure 7. We can see that the results are visually pleasing and that no artifacts occur.

7 Conclusions

We proposed an automatic method to compute accurate point-to-point correspondences between a set of human models in similar posture. We conducted a large-scale error evaluation in the context of shape analysis and we used the correspondences to perform motion transfer. It remains a challenging problem to compute an accurate point-to-point correspondence between human models in varying postures without the use of landmark positions. While some work has been conducted to predict landmarks on humans in arbitrary postures [30], the problem of computing accurate point-to-point correspondences remains open. When trying to solve this problem using the framework presented here, the descriptors, the landmark graph, and the registration algorithm would need to be modified in order to be posture-invariant.

References

1. Dror Aiger, Niloy Mitra, and Daniel Cohen-Or. 4-points congruent sets for robust surface registration. *ACM Transactions on Graphics*, 27(3):#85, 1–10, 2008. Proceedings of SIGGRAPH.
2. Brett Allen, Brian Curless, and Zoran Popović. The space of human body shapes: reconstruction and parameterization from range

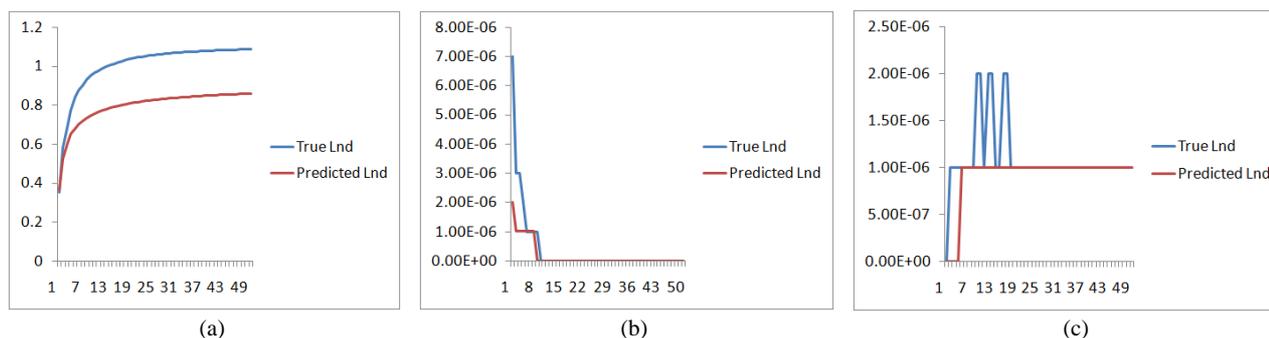


Fig. 5 (a) Compactness. (b) Generalization. (c) Specificity. The x -axis shows the number of principal components and the y -axis shows the values of the error measures.

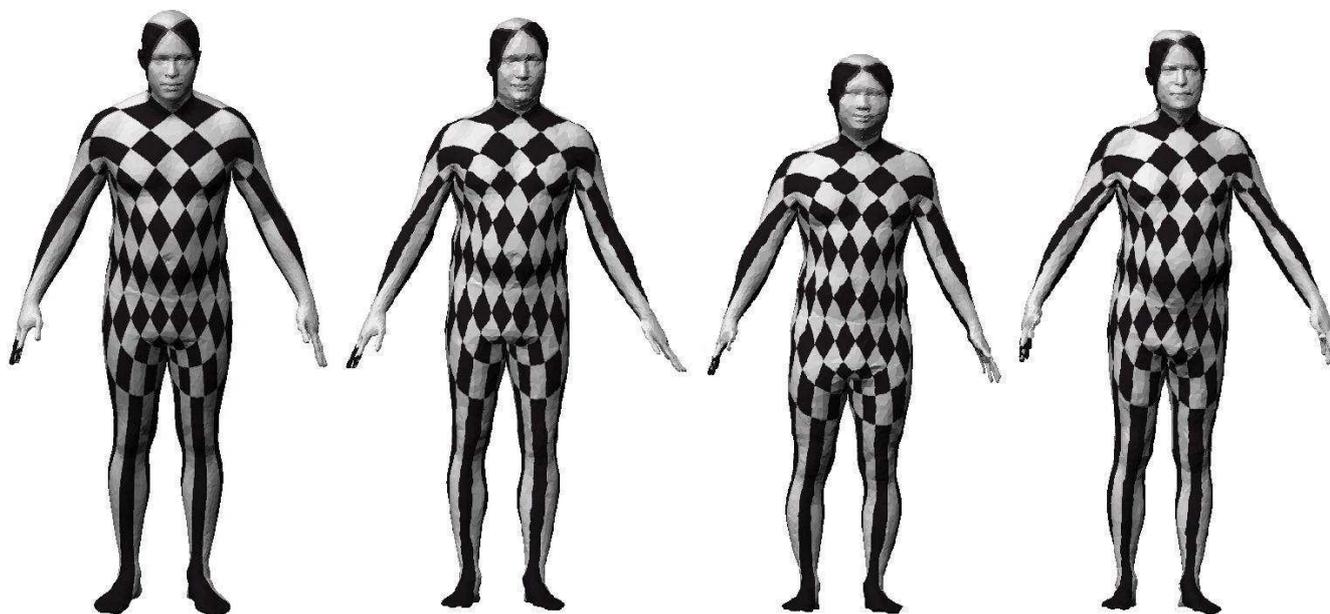


Fig. 6 Texture mapping of the corresponded models. The template mesh with a texture map is shown on the left. The texture map is transferred to the three models shown on the right using the computed correspondences.

- scans. *ACM Transactions on Graphics*, 22(3):587–594, 2003. Proceedings of SIGGRAPH.
- Brian Amberg, Sami Romdhani, and Thomas Vetter. Optimal step nonrigid icp algorithms for surface registration. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2007.
 - Dragomir Anguelov, Praveen Srinivasan, Daphne Koller, Sebastian Thrun, Jim Rodgers, and James Davis. Scape: shape completion and animation of people. *ACM Transactions on Graphics*, 24(3):408–416, 2005. Proceedings of SIGGRAPH.
 - Sunil Arya and David M. Mount. Approximate nearest neighbor queries in fixed dimensions. In *ACM-SIAM Symposium on Discrete Algorithms*, 271–280, 1993.
 - Zouhour Ben Azouz, Chang Shu, and Anja Mantel. Automatic locating of anthropometric landmarks on 3d human models. In *3D Data Processing, Visualization and Transmission*, 2006.
 - Ilya Baran and Jovan Popović. Automatic rigging and animation of 3D characters. *ACM Transactions on Graphics*, 26(3), 2007. Proceedings of SIGGRAPH.
 - Paul Besl and Neil McKay. A method for registration of 3-D shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):239–256, 1992.
 - Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3d faces. In *Proceedings of SIGGRAPH*, pages 187–194, 1999.
 - Will Chang and Matthias Zwicker. Automatic registration for articulated shapes. *Computer Graphics Forum (Special Issue of SGP 2008)*, 27(5):1459–1468, 2008.
 - Will Chang and Matthias Zwicker. Range scan registration using reduced deformable models. *Computer Graphics Forum (Special Issue of Eurographics 2009)*, To appear.
 - Rhodri H. Davies, Carole J. Twining, Tim F. Cootes, John C. Warterton, and Chris J. Taylor. 3D statistical shape models using direct optimization of description length. In *Proceedings European Conference Computer Vision*, 2002.
 - Ian Dryden and Kanti Mardia. *Statistical Shape Analysis*. Wiley, 2002.
 - Juergen Gall, Carsten Stoll, Edilson de Aguiar, Christian Theobalt, Bodo Rosenhahn, and Hans-Peter Seidel. Motion capture using joint skeleton tracking and surface estimation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2009.
 - Natasha Gelfand, Niloy J. Mitra, Leonidas J. Guibas, and Helmut Pottmann. Robust global registration. In *Symposium on Geometry Processing*, page 197, 2005.

16. Peng Guan, Alexander Weiss, Alexandru O. Bălan, and Michael J. Black. Estimating human shape and pose from a single image. In *IEEE International Conference on Computer Vision*, 2009.
17. Nils Hasler, Carsten Stoll, Martin Sunkel, Bodo Rosenhahn, and Hans-Peter Seidel. A statistical model of human pose and body shape. *Computer Graphics Forum (Special Issue of Eurographics 2008)*, 2(28), 2009.
18. Qixing Huang, Bart Adams, Martin Wicke, and Leonidas J. Guibas. Non-rigid registration under isometric deformations. *Computer Graphics Forum (Special Issue of SGP 2008)*, 27(5), 2008.
19. Varun Jain and Hao Zhang. A spectral approach to shape-based retrieval of articulated 3d models. *Computer Aided Design*, 39(5):398–407, 2007.
20. Andrew E. Johnson and Martial Hebert. Using spin images for efficient object recognition in cluttered 3d scenes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(5):433–449, 1999.
21. Hao Li, Robert W. Sumner, and Mark Pauly. Global correspondence optimization for non-rigid registration of depth scans. *Computer Graphics Forum*, 27(5), 2008.
22. Dong C. Liu and Jorge Nocedal. On the limited memory method for large scale optimization. *Mathematical Programming B*, 45:503–528, 1989.
23. Pierre Meunier, Chang Shu, and Pengcheng Xi. Revealing the internal structure of human variability for design purposes. In *IEA World Congress on Ergonomics*, 2009.
24. Mark Pauly, Niloy Mitra, Joachim Giesen, Markus Gross, and Leonidas Guibas. Example-based 3d scan completion. In *Symposium on Geometry Processing*, page 23, 2005.
25. Judea Pearl. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann, 1988.
26. Kathleen Robinette, Hans Daanen, and Eric Paquet. The CAESAR project: A 3-D surface anthropometry survey. In *Conference on 3D Digital Imaging and Modeling*, pages 180–186, 1999.
27. Szymon Rusinkiewicz and Marc Levoy. Efficient variants of the ICP algorithm. In *Conference on 3D Digital Imaging and Modeling*, June 2001.
28. Martin Styner, Kumar T. Rajamani, Lutz-Peter Nolte, Gabriel Zsemlye, Gábor Székely, Christopher J. Taylor, and Rhodri H. Davies. Evaluation of 3d correspondence methods for model building. In *Information Processing in Medical Imaging*, pages 63–75, 2003.
29. Oliver van Kaick, Hao Zhang, Ghassan Hamarneh, and Danial Cohen-Or. A survey on shape correspondence. In *Eurographics State-of-the-art Report*, 2010.
30. Stefanie Wuhrer, Zouhour Ben Azouz, and Chang Shu. Semi-automatic prediction of landmarks on human models in varying poses. In *Canadian Conference on Computer and Robot Vision*, 2010.
31. Pengcheng Xi, Won-Sook Lee, and Chang Shu. Analysis of segmented human body scans. In *Graphics Interface*, 2007.
32. Thomas Yeo, Mert Sabuncu, Tom Vercauteren, Nicholas Ayache, Bruce Fischl, and Polina Golland. Spherical demons: Fast diffeomorphic landmark-free surface registration. *IEEE Transactions on Medical Imaging*, 29(3):650–668, 2010.
33. Hao Zhang, Alla Sheffer, Daniel Cohen-Or, Qingnan Zhou, Oliver van Kaick, and Andrea Tagliasacchi. Deformation-driven shape correspondence. *Computer Graphics Forum (Special Issue of SGP 2008)*, 27(5), 2008.

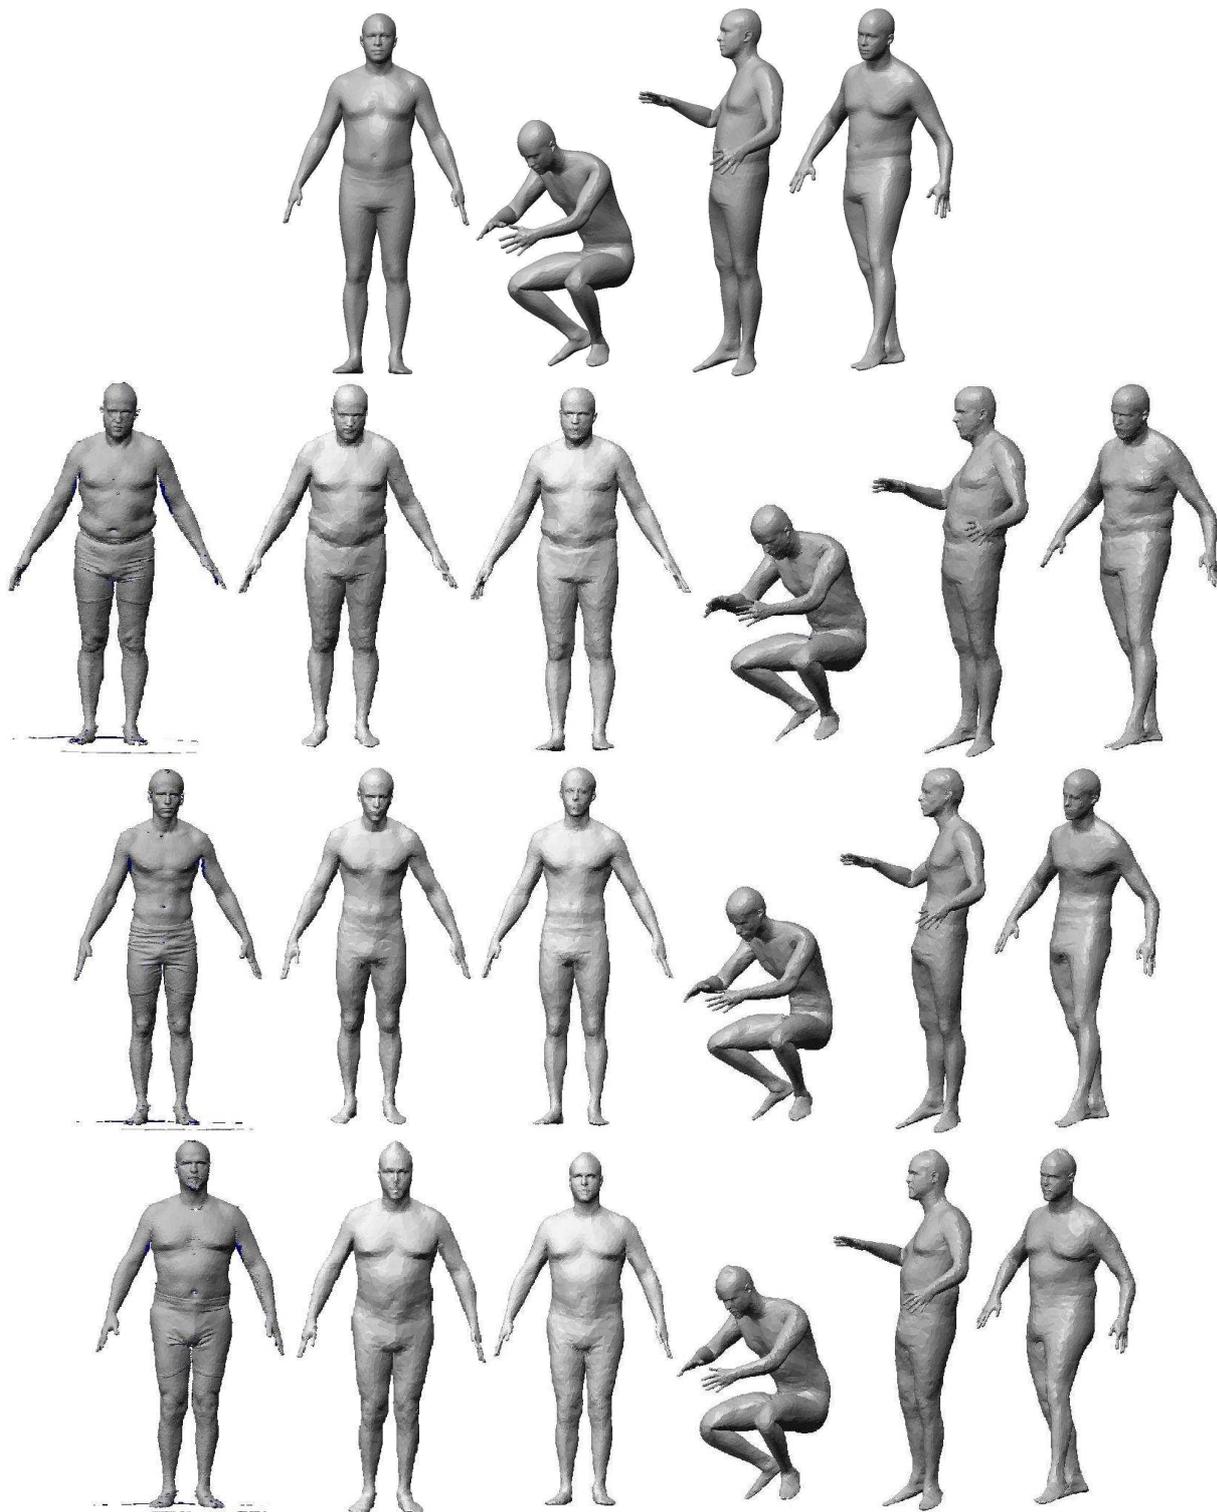


Fig. 7 Motion Transfer. Top row (left to right): template model T and three animated poses. Remaining rows (left to right): scan S of the CAESAR database, T deformed to S using ground truth landmarks, T deformed to S using predicted landmarks, and three animated poses obtained using motion transfer with the result that uses predicted landmarks.