

A Data-driven Segmentation for the Shoulder Complex

Q Youn Hong¹, Sang Il Park^{2,3} and Jessica K. Hodgins^{1,3}

¹School of Computer Science, Carnegie Mellon University

²Department of Digital Contents, Sejong University

³Disney Research, Pittsburgh

Abstract

The human shoulder complex is perhaps the most complicated joint in the human body being comprised of a set of three bones, muscles, tendons, and ligaments. Despite this anatomical complexity, computer graphics models for motion capture most often represent this joint as a simple ball and socket. In this paper, we present a method to determine a shoulder skeletal model that, when combined with standard skinning algorithms, generates a more visually pleasing animation that is a closer approximation to the actual skin deformations of the human body. We use a data-driven approach and collect ground truth skin deformation data with an optical motion capture system with a large number of markers (200 markers on the shoulder complex alone). We cluster these markers during movement sequences and discover that adding one extra joint around the shoulder improves the resulting animation qualitatively and quantitatively yielding a marker set of approximately 50 markers for the complete skeleton. We demonstrate the effectiveness of our skeletal model by comparing it with the ground truth data as well as with the recorded video. We show its practicality by integrating it with the conventional rendering/animation pipeline.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation

1. Introduction

The shoulder complex is one of the most complicated joints in the human skeletal structure as it consists of three bones: the clavicle, the scapula and the humerus, as well as their associated muscles, tendons and ligaments. The shoulder bones form three synovial joints: the sternoclavicular (SC) joint, the acromioclavicular (AC) joint and the glenohumeral (GH) joint. The combination of these joints allows the shoulder complex to have both a very wide range of motion and subtle internal movements such as a shrug [vdH97].

Unlike the complexity of the anatomical shoulder mechanism, animated characters often have shoulders that are modeled as a single ball and socket joint. The shoulder joint is defined as three rotational degrees of freedom between two bones, one attached to the spine and the other representing the upper arm. This discrepancy between the complexity of anatomical system and the simple animated model impacts the quality of the resulting animation, especially when the motion is rendered on a full human figure with a skinned surface.

There have been various approaches to reduce the gap between human shoulder deformations and those created in computer graphics. In particular, there has been a great deal of work in skinning algorithms aimed at volume preservation [MG03, KZ05, WPP07, RHC09] as well as detailed anatomical models [SPCM97, NT00, ZCCD04, PCLS05, LST09]. In this paper, we approach this problem by learning the structure of the shoulder from data. We collect the data using densely spaced surface markers and reconstruct the mesh surface with the technique presented by Park and Hodgins [PH06]. From these example poses, we determine the appropriate skeletal structure including the number of joints and the position of each joint. We integrate this new skeletal structure into a traditional motion capture and rendering pipeline by identifying a small number of markers (four/segment) which approximate the motion of the skeleton generated from the large marker set and estimating the skin weight of each vertex on a high resolution mesh with respect to the computed skeleton. We compare our results to

video of the capture session and quantitatively to the mesh created from the large marker set.

Our main contribution is a new skeletal model which reduces artifacts and creates more visually pleasing motions for the shoulder in animation rendered in the conventional capture/animation/rendering pipeline. Although our model is not anatomically based, it is data-driven and therefore provides a good match to the actual deformations seen in the skin and muscle around the shoulder. We further contribute to the literature by placing our collected data and model parameters on the web (<http://graphics.cs.cmu.edu/projects/shoulder>) for others to use in their captures or research.

The remainder of the paper is organized as follows. In the next section, we describe the existing research in skin deformation and techniques for constructing the skeletal structure based on data. Section 3 describes our approach: collecting data and then determining the skeleton and skin weights. In section 4, the experimental results are presented. The limitations and future work are discussed in section 5.

2. Background

Skeleton-driven techniques are the prevalent techniques for generating human character animation. The most popular method is skeleton subspace deformation or skinning, in which each vertex on the skin surface is connected to underlying bones with proper weight values. The position of the vertices are determined by the connected bones as described by Lewis and his colleagues [LCF00]. Most commercial animation tools such as 3ds Max and Maya support this method, and we use this approach to render our final results. However, a simple implementation of this method suffers from artifacts such as the “collapsing-elbow” and the “candy-wrapper” effect (described in [LCF00]). Many recent research efforts have been aimed at overcoming these problems through the use of more sophisticated models or by carefully determining the parameters and coefficients of the method based on the examples.

The improved models include such techniques as changing the interpolation scheme [KZ05], using swept ellipsoids for approximating the model [HYC*05], mimicking the volume effects of muscles and fat by adding deformable geometries under the skin [GW05], building a simple anatomical model from the outside in [PCLS05], and applying a post-processed volume correction [RHC09].

The second approach uses sample skin data in various poses as a training set and determines a more accurate set of parameters. Our work falls into this class of approaches in that we use an optimization framework to find an additional joint. Mohr and Gleicher add extra joints in the region of complex deformation observed in the data [MG03]. The number of extra joints and their positions were determined by the user. The movement of the extra joints is in-

ferred from the neighboring joints rather than captured directly. Wang and Phillips also use sample surface data to determine the skinning parameters [WP02]. However, rather than modifying the skeletal structure, they break down the single weight value for a bone into 12 sub-weight values, where each corresponds to a component of the transformation matrix of a bone. Similar to Wang and Phillips’ method, Merry and his colleagues present a method called Animation space, in which a vertex has 4 sub-weight values for a bone [MMG06]. Park and Hodgins model dynamic skin deformation based on captured skin data [PH08]. They focus on representing the non-linearity of the skin deformation given the skeletal structure while here we present a method to find an effective skeletal structure when the skinning algorithm is given.

Recently, mesh animation data has become more readily available with new surface acquisition methods. Using mesh animation data is not easy due primarily to the difficulty of editing such a large unstructured dataset. One solution is to convert the surface data into skeleton-based skinned data. James and Twigg suggested a method for representing mesh animation as a set of near-rigid objects based on the mean-shift segmentation method [JT05]. Angelov and his colleagues compute a skeletal structure from a set of 3D scanned surface data [ASK*05], and de Aguiar and his colleagues compute it from mesh animation [dATT08]. We basically follow those approaches in estimating the skeletal structure and add an automatic method for determining the number of bones and joint locations.

3. Approach

Our algorithm for computing a skeletal model and the skinning weights has five steps: database construction, rigid segmentation, skeleton extraction, marker selection and skin weight computation. We now describe those steps in detail.

3.1. Database Construction

We build a database of captured skin surfaces around the shoulder area from a variety of poses. There have been




Figure 1: Database construction: (a) marker placement on the subject’s body; (b) captured markers and marker surface.

Table 1: List of motions used for building the shoulder motion database.

Training set		Testing set	
Name	Number of frames	Name	Number of frames
range of motion	2523	body building	1114
stretching	1766	golfing	776
various poses	1602	jumping jack	717
		swimming	1580


many methods developed for this purpose including scanners [ACP02, ASK*05] and optical motion capture devices [SMP03, PH06]. We use an optical motion capture system for capturing the surface because the optical motion capture device captures motion in realtime. We use a Vicon commercial motion capture system consisting of 16 near infra-red cameras at a rate of 120 frames/sec and place approximately 200 markers on the subject's shoulder area (Figure 1(a)). The markers are positioned with approximately a 1 inch spacing.

In order to cover the deformation space of the shoulder effectively, we captured a variety of motions including a range-of-motion exercise, stretching, and golfing. We divide the motion into a training set and a testing set (Table 1). The captured data are represented as marker trajectories. Markers which are missing due to occlusions are filled in using the method of Park and Hodgins [PH06]. In this technique a triangular mesh is constructed and used to identify the neighboring markers based on the connectivity embedded in the mesh (Figure 1(b)). The motion of the neighbors is used to reconstruct the position of an occluded marker.

3.2. Rigid Segmentation

The conventional skinning process for a virtual character is a process of approximating the deformation of the human body as an articulated set of rigid bodies. Thus, for effective skinning, the character surface should be divided in such a way that its deformation space is modeled well. In this work, we achieve this segmentation by finding sets of triangles from the marker surface that move together in the marker mesh animation. Thus, we transform the problem of rigid segmentation into a clustering problem.

Inspired by the work of James and Twigg [JT05], we cluster the triangles on a marker surface based on the rotational displacement of each triangle because the triangles have the same rotational displacement during rigid transformations. Given a time instance t , the rotational displacement of the i -th triangle is represented as a unit-quaternion $\mathbf{q}_i^t \in \mathbb{SO}^3$, which is computed by finding the absolute orientation by solving a least-squares problem [Hor87]. In clustering, our feature vector for the i -th triangle is $\mathbf{F}_i = [\mathbf{v}_i^1 \ \mathbf{v}_i^2 \ \cdots \ \mathbf{v}_i^T]^T$, where T is the total number of frames in the training set and $\mathbf{v}_i^t \in \mathbb{R}^3$ is the rotational displacement vector corresponding to \mathbf{q}_i^t computed by the logarithm map [LS02]. By using \mathbf{v}_i^t instead of \mathbf{q}_i^t , we assume that two rotational changes are sim-

**Figure 2:** Segmentation results: (a) three segments; (b) four segments; (c) five segments; (d) six segments

ilar if their corresponding displacement vectors are close in the Eulerian space.

Many methods have been proposed for clustering. For example, mean shift is used by James and Twigg [JT05] for a similar purpose. However, in our case, the number of clusters is finite and small. Therefore, we test all reasonable candidate number of clusters using standard K-means clustering. We choose three as the smallest number to test because it is comparable to the common skeletal structure used in games consisting of the right chest, the upper arm, and the lower arm. The largest number we test is six, which we believe is the maximum number that makes sense for rigging. We test each integer value, and evaluate how many clusters are needed. Because the clustering result from K-means is sensitive to the choice of the initial guesses for the cluster centers, we use a randomized approach: we repeat the clustering many times (1000 times) with a random initial seed and choose the best. We measure the quality of the clustering by summing the distance between feature vectors and their cluster center. On an Intel®Core™2 CPU at 2.66GHz, it takes less than 30 minutes to execute the K-mean algorithm 1000 times.

Figure 2 shows our segmentation results with a varying number of segments. In the three segments result (Figure 2(a)), the segmentation matches the conventional skeletal structure. The result with four segments reveals the shoulder area as a separate segmentation (Figure 2(b)). Five segments refines the shoulder area more (Figure 2(c)). The additional segment provided by six clusters appears around the tip of the wrist and does not contribute to the shoulder refinement.




Figure 3: Four cluster results (a) segmentation result; (b) skeleton extraction; (c) skinning weight computation; (d) registered detailed surface and skinning weight transfer.

3.3. Skeleton Extraction

After segmentation, we optimize the skeletal structure. Each segment represents a bone in the skeleton. We assume that our skeletal structure branches from the spine around the upper chest and forms a single chain until the wrist joint. This assumption does not have an anatomical basis. However, the number of segments we choose is too small to build a complex skeletal structure. More importantly, this single chain assumption makes the skeleton easy to build, which is crucial for a practical animation pipeline. We set each joint between two adjacent bones as a 3 degree-of-freedom (dof) rotational ball joint for generality. Based on this assumption, defining the skeletal structure is dependent on determining the joint location.

A ball joint between two adjacent bones can be found as a fixed point on one bone which also remains stationary across time when described in the other bone's local coordinate system. We employ the optimization method presented by Angelov and colleagues [AKP*04], which models computation of the joint location as a weighted linear regression.

3.4. Marker Selection

In the conventional motion capture pipeline, a minimal set of markers are used to generate the kinematic constraints of the bones and therefore determine their movement. The number of markers required to sufficiently constrain a bone depends on the type of joint. However, four markers for each segment is enough to solve for either three dof or six dof joints. Based on an analysis of the data from the dense marker set, we provide a method to guide the placement of the small number of markers so that the rigid motion of a segment can be computed for a new motion.

We choose those representative markers based on two conditions: the selected markers should show a similar transformation to that of the associated segment (computed from




Figure 4: Selected markers of the 4-cluster case: (a) front view; (b) rear view.

the dense marker set) and the distance between the selected markers should be sufficient to prevent any confusion in the motion capture process. We measure how a marker satisfies the first condition by the deviation of its real position from its estimated position when we assume it follows the transformation of the segment rigidly:

$$e_r = \sum_{t=0}^T \| \mathbf{p}^t - \mathbf{T}^t \hat{\mathbf{p}} \|^2, \quad (1)$$

where e_r is the deviation value of a marker, \mathbf{p}^t is the captured position of the marker at time t , \mathbf{T}^t is the computed transformation of the segment, and $\hat{\mathbf{p}}$ is the position of the marker at a reference instance.

To enforce the second condition, we define a distance energy e_d of a marker such that e_d increases when the distance between the marker and any selected representative markers is reduced and goes to zero when the distance value is greater than a predefined threshold, that is

$$e_d = \sum_{i=0}^{N_s} d(\hat{\mathbf{p}}, \hat{\mathbf{p}}_i), \quad (2)$$

where N_s is the number of existing representative markers, $\hat{\mathbf{p}}_i$ is the position of the i -th selected marker, and $d(\cdot, \cdot)$ is a function defined as follows:

$$d(\hat{\mathbf{p}}, \hat{\mathbf{p}}_i) = \begin{cases} 1 - \frac{\|\hat{\mathbf{p}} - \hat{\mathbf{p}}_i\|}{d_t} & \text{if } \|\hat{\mathbf{p}} - \hat{\mathbf{p}}_i\| < d_t \\ 0 & \text{otherwise} \end{cases}, \quad (3)$$

In our experiment, we assign d_t as one fourth of the longest distance between any two markers on the segment.

Putting these two conditions together, we choose the representative markers to have the minimal energy value $e = \omega \times e_r + (1 - \omega) \times e_d$, where ω is the weight between those two policies (0.5 in our experiments). Markers are chosen independently for each cluster using a greedy algorithm: after computing the energy value for each marker, we pick the marker with the lowest value. Then, we repeat this process until we have selected four markers. The first marker of a segment is selected solely based on its e_r value. Figure 4 shows the representative markers for four clusters.

3.5. Skin Weight Computation

Given the new skeletal model, the final process is to bind a detailed surface model for rendering to the skeleton by assigning weight values to each of the vertices. We call this surface the *detailed surface model* in order to avoid confusion with the marker surface. When the detailed surface model is animated, the rendered results depend heavily on the skin weight values. Skilled artists assign the weight values manually. In our work, we automatically assign the weight values of the detailed surface. To compute the skin weights of the detailed surface, we first compute the weight values for the vertices on the marker surface. We exploit the method of Park and Hodgins [PH06]: the weight values of markers inside the segment are only dependent on the segment and the weight values of markers in the boundaries between neighboring segments are dependent on the likelihood that a marker belongs to each segment, which is computed from the distance between captured vertex and vertex rigidly deformed by each segment. After computing the weight values for the markers, we transfer the weight values of the markers to the detailed surface by registering the detailed surface onto the marker surface and then by applying barycentric interpolation to the weight values.

4. Results

The primary purpose of this work is to determine an effective skeletal configuration for the shoulder in a standard motion capture animation pipeline. We have tested skeletal structures with varying number of segments. Table 2 shows the errors for the synthesized motion.

For evaluation purposes, we use the marker surface as a skin to be animated by the skeleton. The surface is deformed along with the movement of the skeleton by using a standard skinning algorithm for single weight enveloping. Weight values used in the skinning algorithm are computed from our skeletal model. The joint angles of the skeleton are computed either using the full marker set or the reduced marker set.

We measure the quality of the skeletal model by comparing the positions between deformed vertices and their corresponding captured markers. Figure 5 summarizes the reconstruction error for the training set. The values shown in the graph are the average error values over the whole frame of the training data as well as over all the markers. In the three segments model (most similar to the conventional skeletal model), there exist larger errors. However, adding one more segment around the shoulder as shown in the four segments model greatly reduces the error values. Interestingly, adding more segments does not improve the error values further. The graph also shows the effectiveness of our marker selection (red line). Although the skeletal motion estimated using the dense markers shows consistently better results than the selected small marker set, the global tendency over the number of clusters is consistent.

Table 2: The summary of the synthesis error (mm) when using the dense full marker set and the reduced marker set for estimating the skeletal motion, respectively.

Motion		Number of Segments					
		3		4		5	
		avg (std)	max (min)	avg (std)	max (min)	avg (std)	max (min)
Training	Full	15.87 (7.44)	34.13 (0.00)	13.30 (6.17)	27.80 (0.00)	13.20 (6.04)	28.57 (0.00)
	Reduced	20.79 (9.45)	50.30 (0.00)	16.01 (7.95)	41.51 (0.00)	17.04 (8.14)	35.92 (0.00)
Body-building	Full	11.90 (6.74)	24.48 (0.00)	10.37 (5.79)	21.48 (0.00)	10.74 (5.98)	22.90 (0.00)
	Reduced	17.70 (12.28)	45.96 (0.00)	13.87 (8.08)	25.54 (0.00)	12.98 (6.04)	22.30 (0.00)
Golfing	Full	16.84 (6.23)	27.33 (0.00)	13.99 (4.36)	20.01 (0.00)	13.52 (3.96)	18.66 (0.00)
	Reduced	26.17 (8.95)	42.51 (0.00)	21.21 (7.50)	35.31 (0.00)	19.04 (5.74)	28.52 (0.00)
Jumping-jack	Full	16.38 (6.20)	24.82 (0.00)	13.69 (4.90)	20.92 (0.00)	13.61 (4.66)	19.55 (0.00)
	Reduced	24.83 (10.54)	47.83 (0.00)	20.11 (9.46)	39.21 (0.00)	19.97 (7.12)	28.73 (0.00)
Swimming	Full	14.21 (5.76)	34.89 (0.00)	12.21 (5.11)	31.55 (0.00)	12.39 (5.38)	32.55 (0.00)
	Reduced	18.82 (7.80)	44.45 (0.00)	14.60 (6.73)	36.27 (0.00)	16.42 (8.03)	45.37 (0.00)




Figure 5: Reconstruction error with varying number of segments is plotted when using the dense full marker set and the reduced marker set for the training set.

We also synthesize the marker surface data for the test motion and measure the errors for the full and reduced marker sets (Figure 6 and Figure 7). The size of the error is dependent on the type of motion but the error values show similar trends. We conclude that inserting one additional segment between the chest and the upper arm greatly improves the representation of the shoulder but that adding additional segments does not further improve the quality.

Figure 8 shows a screen shot of the results for reconstructing the training set when the different estimated skeletal models are used to animate a detailed surface. Because of the small number of degrees of freedom, the three segments model has a shrinking shoulder when the actor stretches his arm forward. The four and five segments models look similar, and they appear to conserve volume. They also better match the photo of Figure 8(d). Figure 8(e) shows the com-




Figure 6: Synthesis error with varying number of segments is plotted when using the dense full marker set for the test motion.




Figure 7: Synthesis error with varying number of segments is plotted when using the reduced marker set for the test motion.

parison between the three segments model (red) and the four segments model (blue), in which the shrinking of the three segments model is obvious. On the other hand, as shown in Figure 8(f), the four segments model (blue) agrees closely with that of the five segments model (green). Figure 9 and Figure 10 illustrate the results when the models are applied to the test motion. The four segments model is visually superior to the three segments model but adding additional segments does not show much further improvement.

We apply our suggested four segments shoulder model to a full body character and compare it with the three segments shoulder model. We first adopt a conventional full body skeletal model and adjust its upper body to include the extra shoulder segment. The test motions are captured using 69 optical markers, with the markers on the chest, the shoulders and the arms placed at the positions suggested by our method. We also transfer the skin weight values of the shoulder area obtained in Section 3.5 to the full body surface model by aligning the shoulder-only surface onto the

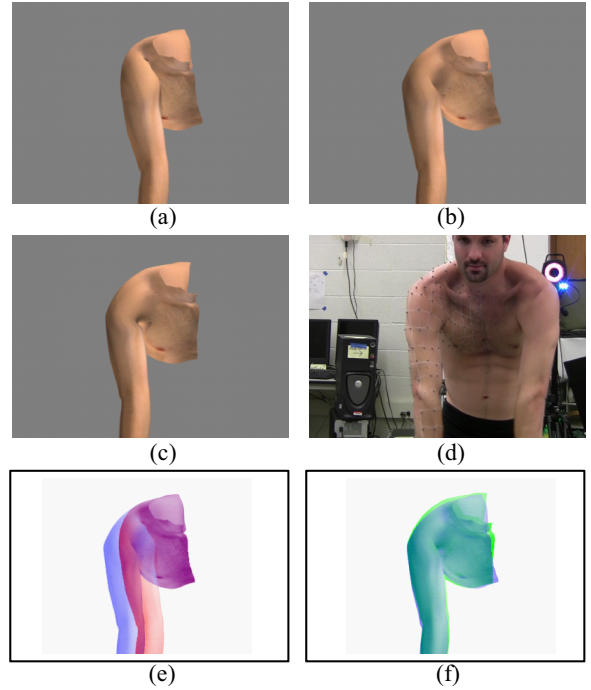


Figure 8: Visual evaluation of a pose from the training set after rendering a detailed surface: (a) three segments; (b) four segments; (c) five segments; (d) photo; (e) comparison between three segments result (red) and four segments result (blue); (f) comparison between four segments (blue) and five segments (green).

Table 3: Comparison of our method and the method of MG03: the synthesis error (mm) when using the reduced marker set to synthesize in our model and reduced marker set in MG03's model, respectively.

Motion	Methods			
	Ours		MG03	
	avg (std)	max (min)	avg (std)	max (min)
Training	16.01 (7.95)	41.51 (0.00)	20.55 (10.62)	59.90 (0.00)
Body-building	13.87 (8.08)	25.54 (0.00)	17.67 (11.65)	47.73 (0.00)
Golfing	21.21 (7.50)	35.31 (0.00)	31.10 (12.52)	53.25 (0.00)
Jumping-Jack	20.11 (9.46)	39.21 (0.00)	24.10 (15.55)	61.00 (0.00)
Swimming	14.60 (6.73)	36.27 (0.00)	18.46 (7.00)	41.25 (0.00)

full body surface. Figure 11 shows the comparison results between the two skeletal models. The artifact of the collapsing shoulder is alleviated by our model.

Finally, we compare our method with the skinning method proposed by Mohr and Gleicher [MG03], in which they add extra joints procedurally to the locations where the linear blending skinning generates artifacts. For testing our




Figure 9: Visual evaluation of a pose from body-building motion in the test set: (a) three segments; (b) four segments; (c) five segments; (d) six segments.




Figure 10: Visual evaluation of a pose from swimming motion in the test set after rendering a detailed surface: (a) three segments; (b) four segments; (c) five segments; (d) photo.

method, we use the four segments model built from the reduced marker set, whereas for testing their method, we use the three segments model and add an extra joint to the shoulder joint, assigning a halfway interpolated angle to the added joint. The weight values in each model are computed by following each paper. Table 3 shows the synthesis error of each method. The comparison results show that our method yields the smaller average error per marker, especially for golfing and jumping jacks.




Figure 11: Visual evaluation of a pose from cowboy in a full body animation: (a) skeleton without the extra shoulder segment; (b) skeleton with the extra shoulder segment

5. Discussion

We propose a new shoulder model which creates a visually more pleasing skin surface when combined with standard skinning algorithms. The new model is data-driven and approximates the skin surface data captured with an optical motion capture system. We discover that adding one extra joint to the common shoulder model improves the quality of the resulting animation, whereas adding more joints does not provide further improvement. We demonstrate the effectiveness of our model by comparing with ground truth data as well as with recorded video. We also show its practicality by integrating it with the conventional rendering/animation pipeline.

Our method has certain limitations. Because we use a data-driven approach, the results are inherently subject-specific, that is, the provided guideline in this paper is most effective when applied to an animated subject having a similar body type to that of our human test subject. It is possible that a different body segmentation would provide better results for different body types. We would like to explore more body types and possible generalization between body types.

Our results are also dependent on the choice of the training motion. We choose a range of motion exercise as the training set. Because our optimization method relies on a least-squares framework, the results can be biased with respect to the number of the frames of similar poses. We can overcome this problem by collecting only distinct poses from the examples and making sure that our data include all common shoulder poses.

We assume that the skeletal structure from the clavicle to the arm forms a single chain structure without any branches. The method of using the minimum spanning tree in [OBBH00] could be a good alternative for determining the skeletal hierarchy in general.

We chose to study the shoulder complex because it often creates problems in capturing and skinning of realistic human characters. Other body parts may also require more sophisticated models than the conventional skeletal model. For example, it is well known that twisting the fore-arm shows the candy-wrapper artifact in the conventional skeletal mode. In this case, the simple rotary joint model does not accurately present the skin surface.

6. Acknowledgements

The authors would like to thank Moshe Mahler for his help in modeling and Justin Macey for his assistance in the motion capture. The authors also would like to thank Autodesk for their donation of Maya software. This work has been partially supported by NSF grant CCF-0702556. Partial support for Sang Il Park was provided by Disney Research, Pittsburgh and Seoul Research and Business Development Program (10557).

References

- [ACP02] ALLEN B., CURLESS B., POPOVIĆ Z.: Articulated body deformation from range scan data. *ACM Transactions on Graphics* 21, 3 (2002), 612–619.
- [AKP*04] ANGUELOV D., KOLLER D., PANG H., SRINIVASAN P., THRUN S.: Recovering articulated object models from 3d range data. In *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence* (2004), pp. 18–26.
- [ASK*05] ANGUELOV D., SRINIVASAN P., KOLLER D., THRUN S., RODGERS J., DAVIS J.: SCAPE: shape completion and animation of people. *ACM Transactions on Graphics* 24, 3 (2005), 408–416.
- [dATTS08] DE AGUIAR E., THEOBALT C., THRUN S., SEIDEL H.-P.: Automatic Conversion of Mesh Animations into Skeleton-based Animations. *Computer Graphics Forum* 27, 2 (4 2008), 389–397.
- [GW05] GUO Z., WONG K. C.: Skinning with deformable chunks. *Computer Graphics Forum* 24, 3 (2005), 373–382.
- [Hor87] HORN B. K. P.: Closed-form solution of absolute orientation using unit quaternions. *Journal of the Optical Society of America A* 4, 4 (1987), 629–642.
- [HYC*05] HYUN D.-E., YOON S.-H., CHANG J.-W., SEONG J.-K., KIM M.-S., JÜTTLER B.: Sweep-based human deformation. *The Visual Computer* 21, 8–10 (2005), 542–550.
- [JT05] JAMES D. L., TWIGG C. D.: Skinning mesh animations. *ACM Transactions on Graphics* 24, 3 (2005), 399–407.
- [KZ05] KAVAN L., ZARA J.: Spherical blend skinning: A real-time deformation of articulated models. In *Proceedings of ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games* (2005), pp. 9–16.
- [LCF00] LEWIS J. P., CORDNER M., FONG N.: Pose space deformations: A unified approach to shape interpolation and skeleton-driven deformation. In *Proceedings of ACM SIGGRAPH 2000* (2000), pp. 165–172.
- [LS02] LEE J., SHIN S. Y.: General construction of time-domain filters for orientation data. *IEEE Transactions on Visualization and Computer Graphics* 8, 2 (2002), 119–128.
- [LST09] LEE S.-H., SIFAKIS E., TERZOPOULOS D.: Comprehensive biomechanical modeling and simulation of the upper body. *ACM Transactions on Graphics* 28, 4 (2009), 1–17.
- [MG03] MOHR A., GLEICHER M.: Building efficient, accurate character skins from examples. *ACM Transactions on Graphics* 22, 3 (2003), 562–568.
- [MMG06] MERRY B., MARAIS P., GAIN J.: Animation space: A truly linear framework for character animation. *ACM Transactions on Graphics* 25, 4 (2006), 1400–1423.
- [NT00] NEDEL L. P., THALMANN D.: Anatomic modeling of deformable human bodies. *The Visual Computer* 16, 6 (2000), 306–321.
- [OBBH00] O'BRIEN J. F., BODENHEIMER R. E., BROSTOW G. J., HODGINS J. K.: Automatic joint parameter estimation from magnetic motion capture data. In *Proceedings of Graphics Interface 2000* (May 2000), pp. 53–60.
- [PCLS05] PRATSCHER M., COLEMAN P., LASZLO J., SINGH K.: Outside-in anatomy based character rigging. In *Proceedings of ACM SIGGRAPH / Eurographics Symposium on Computer Animation* (2005), pp. 329–338.
- [PH06] PARK S. I., HODGINS J. K.: Capturing and animating skin deformation in human motion. *ACM Transactions on Graphics* 25, 3 (2006), 881–889.
- [PH08] PARK S. I., HODGINS J. K.: Data-driven modeling of skin and muscle deformation. *ACM Transactions on Graphics* 27, 3 (2008), 1–6.
- [RHC09] ROHMER D., HAHMANN S., CANI M.-P.: Exact volume preserving skinning with shape control. In *Proceedings of the 2009 ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (2009), ACM, pp. 83–92.
- [SMP03] SAND P., McMILLAN L., POPOVIĆ J.: Continuous capture of skin deformation. *ACM Transactions on Graphics* 22, 3 (2003), 578–586.
- [SPCM97] SCHEEPERS F., PARENT R. E., CARLSON W. E., MAY S. F.: Anatomy-based modeling of the human musculature. In *Proceedings of ACM SIGGRAPH 1997* (1997), pp. 163–172.
- [vdH97] VAN DER HELM F. C. T.: A standardized protocol for motion recordings of the shoulder. In *Proceedings of the First Conference of the International Shoulder Group* (1997), pp. 27–28.
- [WP02] WANG X. C., PHILLIPS C.: Multi-weight enveloping: Least-squares approximation techniques for skin animation. In *Proceedings of ACM SIGGRAPH / Eurographics Symposium on Computer Animation* (2002), pp. 129–138.
- [WPP07] WANG R. Y., PULLI K., POPOVIĆ J.: Real-time enveloping with rotational regression. In *SIGGRAPH '07: ACM SIGGRAPH 2007 papers* (New York, NY, USA, 2007), ACM, p. 73.
- [ZCCD04] ZORDAN V. B., CELLY B., CHIU B., DiLORENZO P. C.: Breathe easy: model and control of simulated respiration for animation. In *Proceedings of ACM SIGGRAPH / Eurographics Symposium on Computer Animation* (2004), pp. 29–37.