

A Study of Shape Similarity for Temporal Surface Sequences of People

Peng Huang, Jonathan Starck, Adrian Hilton
Centre for Vision, Speech and Signal Processing
University of Surrey, Guildford, GU2 7XH, UK
{P.Huang, J.Starck, A.Hilton}@surrey.ac.uk

Abstract

The problem of 3D shape matching is typically restricted to static objects to classify similarity for shape retrieval. In this paper we consider 3D shape matching in temporal sequences where the goal is instead to find similar shapes for a single time-varying object, here the human body. Local-feature distribution descriptors are adopted to provide a rich object description that is invariant to changes in surface topology. Two contributions are made, (i) a comparison of descriptors for shape similarity in temporal sequences of a dynamic free-form object and (ii) a quantitative evaluation based on the Receiver-Operator Characteristic (ROC) curve for the descriptors using a ground-truth data set for synthetic motion sequences. Shape Distribution [25], Spin Image [15], Shape Histogram [1] and Spherical Harmonic [17] descriptors are compared. The highest performance is obtained by volume-sampling shape-histogram descriptors. The descriptors also demonstrate relative insensitivity to parameter setting. The application is demonstrated in captured sequences of 3D human surface motion.

1. Introduction

The development of modelling, digitising and visualising techniques for computer generated shapes has led to a rapidly growing amount of available models. Three-dimensional (3D) shape matching has been widely investigated [31, 4, 10, 13] as a means of effective and efficient object retrieval. However, shape matching techniques typically only consider a single static shape and are designed to classifying similar from dissimilar objects. In this paper we consider the problem of 3D shape matching in temporal sequences where the goal is to find similar shapes for a single time-varying object.

The requirement for temporal shape similarity is motivated by recent work in animation synthesis where sequences of captured motions are simply re-ordered to construct highly-realistic animated content. In particular we

concentrate on human surface motion [29, 32] in which a database of 3D video clips must be compared and concatenated at points where the 3D shape similarity is maximised. Previous work has considered only the similarity in pose of the human skeleton [19] which does not account for surface deformations in clothing and hair. Surface similarity has been defined either manually [29] or through a shape descriptor [32]. Similarity requires a shape descriptor that is sufficiently distinct to differentiate articulated pose while tolerant to changes in surface topology for similar poses. In this paper we review and compare current techniques from the shape retrieval literature for the problem of human surface shape similarity.

The primary contribution of this work is a quantitative evaluation of 3D shape similarity in time-varying sequences using a synthetic ground-truth dataset in which surface correspondence is predefined. A comparison is made for local feature distribution techniques including: Shape Distribution [25], Spin Image [15], Shape Histogram [1] and Spherical Harmonics [17] assuming unknown correspondence. Performance is evaluated by comparing the Receiver Operating Characteristic (ROC) for the techniques, showing the trade-off between correctly and incorrectly classified similarity. The best performance is demonstrated by a shape histogram volume descriptor. These techniques are demonstrated on real data from a public database [28].

In Section 2, we present a review of the state-of-the-art in 3D shape retrieval methods. Section 3 describes the implementation details for the 3D shape descriptors used in our evaluation. In Section 4 the quantitative ground-truth evaluation is described for a synthetic dataset. In Section 5, we present the experimental results along with the application to real data. Finally, we draw conclusions in Section 6.

2. Background

The problem of shape similarity has been widely studied in the 3D shape retrieval literature. There are several techniques: feature-based, graph-based, view-based and bending-invariant methods.

2.1. Feature based methods

Global features are used to characterise the overall shape of 3D models. Typical global features include: volume, surface area, moments, Fourier coefficients and Wavelet coefficients. Zhang and Chen[34] propose an algorithm to efficiently calculate these global features of a 3D model directly from a surface mesh representation. Paquet *et al.*[26] provide three global feature-based descriptors for 3D shape matching, a cord-based descriptor, moment-based descriptor and wavelet-based descriptor. Corney *et al.*[7] use three convex-hull based indices such as hull crumpliness, hull packing and hull compactness for coarsely filtering candidates prior to a more detailed analysis. Kazhdan *et al.*[16] present a reflective symmetry descriptor that extracts the global symmetry information. Such global features are relatively simple to compute but do not provide discrimination at a local level.

Local features can give a more distinctive similarity measure. Shum *et al.*[27] define similarity as the L_2 distance between the local curvature distribution over the mesh representation for two 3D objects. Zaharia and Preteux [33] present the 3D Shape Spectrum Descriptor (3D SSD), which is defined as the distribution of a shape index over the entire mesh, to provide an intrinsic shape description of a 3D mesh. Chua and Jarvis [6] provide a point signature to describe 3D free-form surfaces that is invariant to rotation and translation. Johnson and Hebert[15] present a 3D shape-based object recognition system using Spin Image. These features provide local shape information to improve discrimination between similar shapes.

Local features are compared using a descriptor of the feature distribution. Osada *et al.*[25] introduced a Shape Distribution as a signature to discriminate similar and dissimilar models. A Similarity Measure is computed as the difference between Shape Distributions, which is invariant to translation, rotation and tessellation of the 3D polygonal model. Ankerst *et al.*[1] use a 3D Shape Histogram as a shape signature to classify a molecular database. Kortgen *et al.*[18] attach a 3D Shape Context descriptor to each surface sample point. Shape Context was introduced by Belongie *et al.*[3] for 2D matching. Ohbuchi *et al.* [24] introduce two further shape descriptors, Angle Distance (AD) and Absolute Angle Distance (AAD) histograms for 3D matching.

Another popular approach is a transform-based representation which describes shapes in a transformation invariant manner. Kazhdan *et al.*[17] propose Spherical Harmonic Descriptors that are invariant to rotation for 3D shape retrieval. However, the representation has a potential ambiguity problem. The frequency decomposition is performed independently in concentric spheres, such that two different shapes can have the same spherical harmonic representation. Novotni and Klein [23] provide 3D Zernike Descrip-

tors for 3D shape retrieval, which is an extension of the Spherical Harmonic Representation. A set of descriptors are obtained that are orthonormal, complete and rotation invariant. However, 3D Zernike Descriptors suffer the same ambiguity problem.

2.2. Graph based methods

In the CAD industry, the most common used graph-based representations are Boundary Representation (B-rep) and Constructive Solid Geometry (CSG). El-Mehalawi and Miller [8] construct an attributed graph from a B-rep and measure similarity by using an inexact graph matching algorithm. Similarly, McWerther *et al.*[20] compare models based on shape using information extracted from B-rep into Model Signature Graphs. However, these methods are limited to the CAD community, for example, in matching mechanical parts, and cannot apply to commonly used 3D mesh representations. Sundar *et al.*[30] use a skeletal graph which encodes both the geometric and topological information in the surface to match and compare 3D models. Menier *et al.*[21] propose a method to find the skeleton pose which can best fit the body pose in their motion capture work. However, a skeletal articulated model whose dimensions are manually set is required, and their skeleton sequences lack smooth temporal continuity. Hilaga *et al.*[12] propose a method based on Multi-resolutional Reeb Graphs(MRGs) to estimate a measure of similarity and correspondence between 3D shapes. The similarity is calculated with a coarse-to-fine strategy using the attributes of nodes in the MRG and topological consistency.

2.3. View-based methods

View-based methods represent objects by their image-plane projection. Heczko *et al.* [11] present a silhouette descriptor to characterise 3D objects in terms of their silhouettes obtained by parallel projections. This method requires a PCA-normalised preprocessing, so it is rotation and scale invariant. A view is represented as a feature vector of the fourier coefficients for a regularly sampled set of points on the silhouette contour. Heczko *et al.* [11] propose a similar descriptor called the depth buffer descriptor in which similarity is computed as the distance between the depth buffer feature vectors of two 3D models. Chen *et al.*[5] introduce the LightField Descriptor in which the appearance of an object is characterised by the projected appearance in a set of camera views. Similarity is computed by rotating the camera system surrounding each model until the highest overall similarity (cross-correlation) between the two models from all viewing angles is reached. The similarity between two 3D models is defined as the summation of the similarities across all the corresponding 2D images.

2.4. Bending-invariant methods

Bending-invariant techniques have been proposed to retrieve similar objects independent of changes in articulated pose. Elad and Kimmel [9] present a method to construct a bending invariant signature for these models. They utilise the geodesic distance between surface points as an invariant to surface bending. A bending invariant surface is generated by transforming the geodesic distances between points into Euclidean ones (via an MDS procedure). They translate the problem of matching non-rigid objects in various postures into a simpler problem of matching rigid objects. Jain and Zhang [14] present an approach to robust shape retrieval from databases containing articulated 3D models. Each shape is represented by the eigenvectors of a shape affinity matrix defining the geodesic surface distance between model points. This gives a spectral embedding which achieves normalisation against rigid-body transformations, uniform scaling, and shape articulation.

3. 3D Shape Descriptors

Evaluation of shape similarity is restricted to local feature distributions. Global features provide only a coarse descriptor that will be insufficient to distinguish similarity in a time varying sequence where an object can have the same global properties for a relatively large proportion of time. Graph methods are based on matching surface topology and so do not necessarily handle changes in topology that can occur in captured sequences of articulated motion. Deformation-invariant techniques explicitly match similarity independent of object deformations. Finally view-based techniques could be considered, although in practise with articulated objects there can be large changes in the projected shape of an object with only small changes in articulated pose. In this section we describe the feature distribution descriptors considered.

3.1. Shape Distribution [25]

Shape Distribution provides a shape signature as a probability distribution of a shape function that measures some geometric properties of a 3D model. Typical shape functions are the angle, distance and area for randomised points on the model surface. Here we adopt the $D2$ measure, the distance between two random points on the surface, as proposed by Osada *et al.* [25]. Similarity is measured as the L_2 distance between the distribution $D2$ defined for two meshes.

3.2. Spin Image [15]

A Spin Image is a 2D histogram that encodes the density of mesh vertices projected onto an object-centred space. A spin image is created for an oriented point at a vertex in the

surface mesh as follows:

- An object-centred coordinate (α, β) is computed for each vertex according to the distance α along and the distance β from the principal axis of the object.
- A 2D accumulator indexed by (α, β) is created and the accumulator is incremented for each vertex within the support of the spin image.
- The final histogram is normalised.

The centre of mass and the first axis of the Principal Component Analysis (PCA) of the distribution of mesh vertices is used to define the object-centred coordinate system for a mesh.

3.3. Shape Histogram [1]

Shape Histograms are based on a partitioning of the space containing a 3D model. The complete space is decomposed into disjoint cells, which correspond to the bins of the histogram. Two space partitions are considered:

- Shell model: the space is decomposed into concentric shells around the centre of mass that is invariant to rotation.
- Spherical coordinate model: the space is transformed to a spherical coordinate system (r, ϕ, θ) around the centre of mass, requiring a rotational alignment in (ϕ, θ) for comparison.

Two types of histograms are compared describing first the model surface, then the volume. The surface descriptor accumulates mesh vertices in the histogram and the volume descriptor accumulates voxels in a volumetric representation of the internal region of the surface. Both histograms are normalised. The spherical coordinate histogram is compared invariant of rotation using two approaches, firstly by aligning the model to the principal axes obtained through PCA and secondly by testing similarity for all feasible rotations in ϕ and θ . Instead of rotating a 3D mesh, we generate a fine histogram first, shift it with 1° resolution, and re-bin to a coarse histogram. A similarity measure is computed as L_2 distance between the coarse histograms.

3.4. Spherical Harmonics [17]

Spherical Harmonics represent the volume of an object by a set of spherical basis functions and a descriptor is derived by measuring the energy contained in different frequency bands. The frequency components in the descriptor are rotation invariant. A spherical harmonic descriptor is constructed as follows.

- The volume of an object is divided into a set of concentric shells.

- The frequency decomposition in each shell is computed directly from the mesh surface [22].
- The norm for each frequency component at each radius is concatenated into a 2D histogram indexed by radius and frequency.

The resolution of the shape descriptor is defined by the number of shells defining the radii (r) and the preserved bandwidth (bw) in the spherical harmonics. A similarity measure is computed as L_2 distance between the histograms.

4. Performance Evaluation

The performance of the shape descriptors is evaluated using a ground-truth data-set from simulated data. Temporal mesh sequences are constructed for different motions and the classification of correct and incorrect similarity is assessed using the Receiver-Operator Characteristic (ROC) curves for each technique. The performance of each technique is first evaluated separately to determine optimal parameter settings for the descriptors. The techniques are then compared in deriving similar shapes within motion sequences.

4.1. Ground-truth data

A simulated data-set is created using an articulated character model animated from motion capture data. The following motions are considered: running in a circle; rock&roll dancing; and sprint, with each motion 150 frames long. See Figure 1 for an example. The simulated data is a surface mesh with 20k vertices and 35k triangles. The animated model has a single mesh and so the surface correspondence is known at all frames and rigid-body registration can be performed to align the frames for ground-truth assessment of similarity. This correspondence is only used to compute the “true” surface distance, and is not used in computing the shape similarity measures.

The distortion between two surfaces is measured using the Hausdorff Distance [2], characterising the maximum distance between the surfaces. Let X and Y be the set of mesh vertices for two surfaces, if $d(x, y)$ denotes the Euclidean Distance between one vertex $x \in X$ and one vertex $y \in Y$, we calculate the Hausdorff Distance as follows:

$$D_H(X, Y) = \max\{\max_{y \in Y} \min_{x \in X} \{d(x, y)\}, \max_{x \in X} \min_{y \in Y} \{d(x, y)\}\}$$

Ground-truth shape similarity is defined where the maximum distance $D_H(X, Y)$ falls below a fixed predefined threshold. The character model used is scaled to height of 170cm and the threshold for acceptable similarity is set at $D_H(X, Y) < 10cm$.




Figure 1. Simulated Data. (a) run circle (b) rock&roll (c) sprint at frame 1,16,31,46,61.

4.2. Evaluation Criterion

The performance of the shape descriptors is evaluated using the ROC curve, showing the true-positive rate (TPR) or *sensitivity* in correctly defining similarity against the false-positive rate (FPR) or *one-specificity* where similarity is incorrect.

$$TPR = \frac{ts}{ts + fd}; \quad FPR = \frac{fs}{fs + td}$$

where ts, fs, fd, td are from the confusion matrix, where ts denotes the number of true-similar predictions, fs the false similar, td true dissimilar and fd false dissimilar in comparing the predicted similarity between two frames to the ground-truth similarity.

	<i>TrueSim.</i>	<i>TrueDis.</i>
<i>Pred.Sim.</i>	ts	fs
<i>Pred.Dis.</i>	fd	td

A variable threshold is defined on the similarity score for a shape descriptor. Where similarity falls below the threshold the surfaces are classified as similar, otherwise they are dissimilar. A perfect shape descriptor with a perfect threshold setting would yield a result of 100% sensitivity (all true similar shapes are found) and 100% specificity (no false similar shapes are found). In practise this is not achieved and the performance of the descriptor is assessed by varying the similarity threshold and evaluating sensitivity-specificity trade-off, the ROC curve.

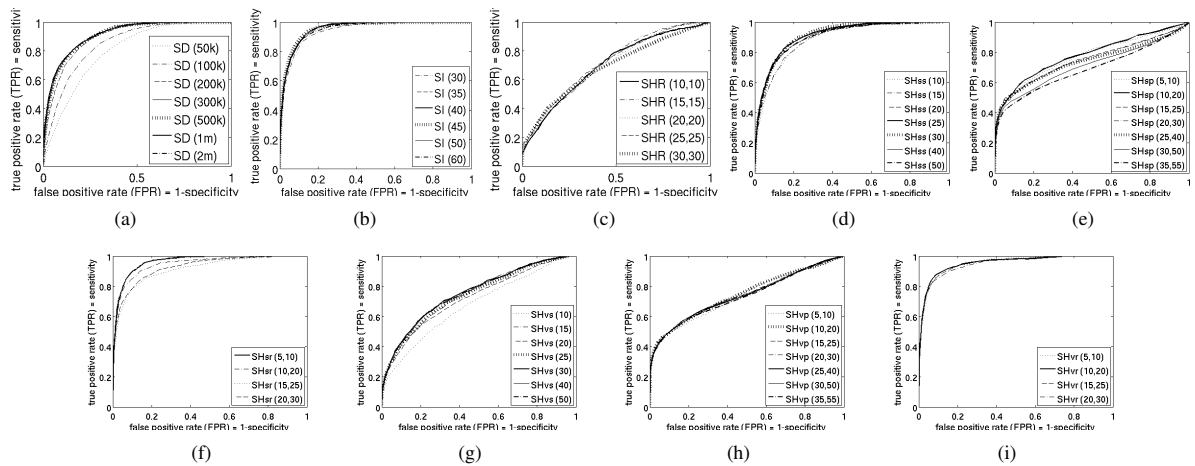


Figure 2. Parameter Setting. (a) Shape Distribution(SD); (b) Spin Image(SI); (c) Spherical Harmonics(SHR); Shape Histogram(SH): (d) surface-shell(ss); (e) surface-spherical histogram with PCA-alignment(sp); (f) surface-spherical histogram with rotation(sr); (g) volume-shell(vs); (h) volume-spherical histogram with PCA-alignment(vp); (i) volume-spherical histogram with rotation(vr).

5. Experimental Results

In this section, we present and discuss the experimental results of the evaluation. Parameter setting in the shape descriptors is first evaluated using self-similarity in the run-circle motion. The techniques are then compared for intra-motion similarity using the best performing parameter settings. Finally the best-performing techniques are applied to a real-data set captured from multiple-view video sequences of a person.

5.1. Parameter Setting

5.1.1. Shape Distribution

Shape Distribution has only one parameter, N , the number of samples to use in constructing the $D2$ distribution. Each sample corresponds to the distance between two random points on the surface of a 3D model. Figure 2(a) shows the performance with different value of N , as the number of samples increases the performance converges to a limit. A near-optimal setting of $N = 1\text{million}$ is chosen where no further increase in performance is apparent.

5.1.2. Spin Image

The Spin Image has two parameters, Nb_α , the number of bins for α , and Nb_β , the number of bins for β in the 2D histogram. The bins correspond to rectangular regions and an equal spatial resolution is adopted $Nb = Nb_\alpha = Nb_\beta$. Figure 2(b) shows the performance with Nb . A limit is found at

$Nb = 40$ where a subsequent increase in resolution reduces performance. As Nb increases, the resolution in resampling a surface in the histogram is increased. As resolution increases surfaces that are spatially close will become resampled into different bins in the histogram and so the overlap between the histograms is in turn reduced.

5.1.3. Shape Histogram

Six variants of the shape histogram are considered. A shell histogram with surface and volume samples, and a spherical coordinate histogram with surface and volume samples using PCA alignment and by evaluating similarity up to a rotation. The shell histogram is defined by the number of shells Ns . The spherical coordinate histogram by the number shells Ns , and angular bins Nb_ϕ, Nb_θ . The resolution of the angular bins is set equal at $Nb = Nb_\phi = 0.5Nb_\theta$. Figure 2(d)-2(i) shows them.

It can be seen that the spherical-coordinate histogram with PCA alignment performs least-well. This is a result of the ambiguity in consistently defining the principal axes used to align two surfaces. As an object becomes more symmetric, the eigenvalues approach making the definition of a consistent set of axes ambiguous. The surface-shell histogram outperforms the volume-shell histogram and is equivalent to the spin-image descriptor with the loss of one-degree of freedom. The highest performance is obtained using the rotated spherical-coordinate histogram which introduces a third degree-of-freedom in the descriptor and derives the best alignment between two surfaces to evalu-




Figure 3. Evaluation for Shape Descriptors. The test sequence: (a) run circle; (b) rock&roll; (c) sprint.

ate similarity. It is interesting to note that performance is improved by using a relatively coarse histogram comparison and the volume-descriptor is more robust to changes in the parameters N_s, N_b . The spherical volume-descriptor is adopted with bins $N_s = 10, N_b = 20$.

5.1.4. Spherical Harmonics

The spherical harmonic descriptor is controlled by two parameters N_s , the number of shells and N_{bw} , the bandwidth, that is the number of frequencies preserved in the harmonics. The greater the value of N_s and N_{bw} the higher the resolution of the descriptor. For simplicity we set $N_b = N_s = N_{bw}$. Figure 2(c) shows that performance is relatively stable with N_b and a low $N_b = 10$ is selected in practise due to the computational cost of deriving the descriptor.

5.2. Comparison of Shape Descriptors

The relative performance of the descriptors is now shown in evaluating the self-similarity for the run-circle motion, rock&roll motion, sprint motion in Figure 3. Similarity curves at the first frame in the sprint motion are shown in Figure 4. The near-optimal parameter settings for the descriptors selected in the run-circle motion are adopted. The results demonstrate that the best performance is obtained using the spherical-coordinate histogram technique. The performance of Spin Image drops when applied to sprint motion. This may be explained by an inherent reflective ambiguity in the descriptor: an object and its mirror have the same descriptor. For example, in Figure 4, frame 29 is quite dissimilar to frame 1 (query frame) but very similar to its mirror, due to the reflective ambiguity, it obtains a valley in the similarity curve. The Shape Distribution descriptor is also invariant to a mirror transformation [25] and has a relatively coarse discrimination. The poorest performance is obtained using spherical harmonics. This is a surprising result considering the performance demonstrated

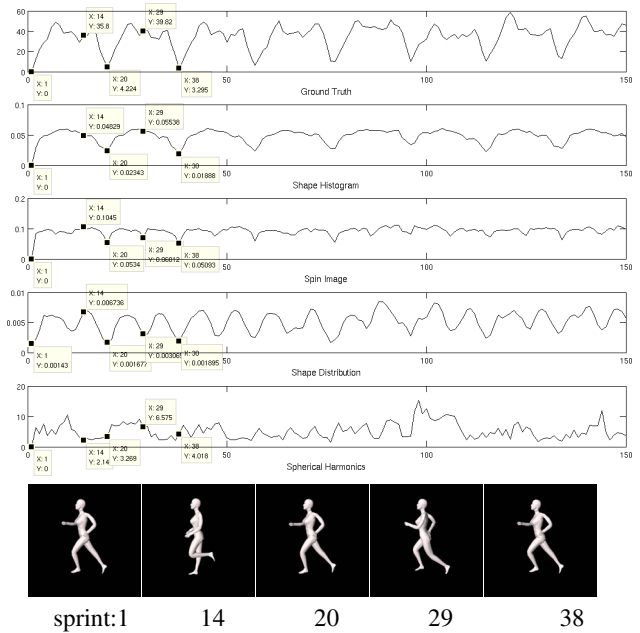


Figure 4. Comparison for Shape Descriptors of 1st frame with sequence for sprint motion. Points marked on the similarity curves indicate the frame shown below.

in the 3D shape retrieval literature. This may be explained in part by the inherent “information loss” in the descriptor [17]: the descriptor is unchanged if we apply different rotations to the different frequency components of a spherical function; the harmonic representation only stores the energy in each frequency component which does not provide enough information to reconstruct the component up to rotation. This means dissimilar shapes may have the same descriptors, the descriptor is therefore more likely to produce false-positive shape matches. For example, in Figure

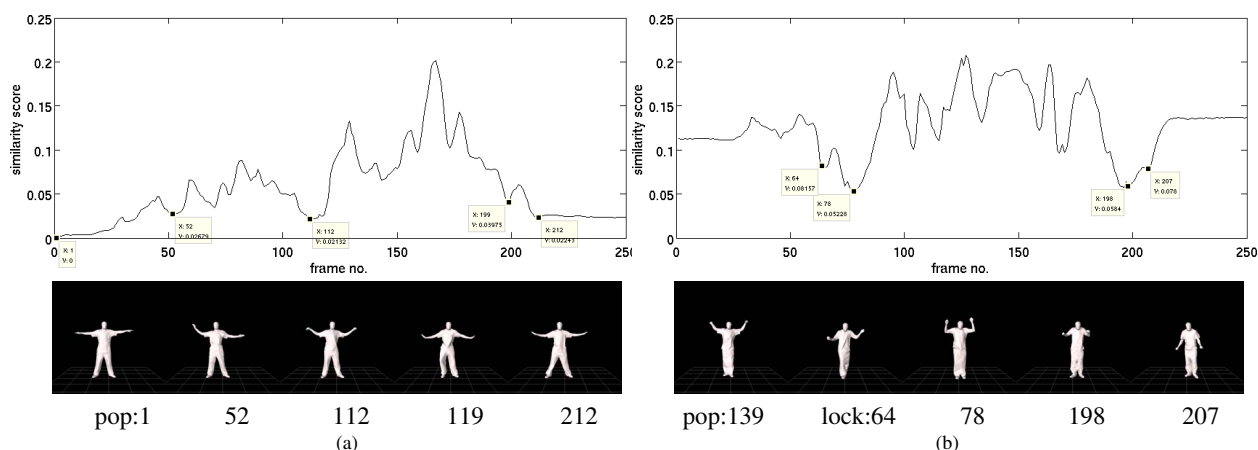


Figure 5. Similarity Measure on Real Data. Intra-motion of pop: (a) similarity curve of frame 1 with sequence. Inter-motion between pop and lock: (b) similarity curve of frame 139 of pop with lock sequence.

4, frame 14 is quite dissimilar to frame 1 but frames 14-20 all form a local minima in the similarity curve. We conclude that for the problem of matching articulated human motion, the spherical volume-histogram descriptors provide the best performance. The descriptors are also demonstrated to be relatively insensitive to parameter setting as shown in Figures 2(i).

5.3. Similarity Measure on Real Data

Finally, we apply a rotated volume-spherical Shape Histogram (SHvr) descriptor to a captured 3D surface sequences of a street-dancer from a public database [28]. Each mesh contains around 140k vertices and 280k triangles. Similar shapes are selected by looking for the valley of the similarity curve. Examples are shown in Figure 5(a) and Figure 5(b), respectively for intra-motion pop dancing and inter-motion between pop and lock dancing.

6. Conclusion

In this paper we have presented a quantitative evaluation of shape descriptors applied to the problem of finding 3D shape similarity in sequences of a temporally varying articulated freeform object, the surface of a clothed person. Local-feature distribution descriptors were adopted from the literature to provide a sufficiently distinct object description and allow for changes in captured surface topology. The Shape Distribution [25], Spin Image [15], Shape Histogram [1] and Spherical Harmonic [17] descriptors were compared. Performance was evaluated using the Receiver-Operator Characteristic in classifying correct versus incor-

rect similarity against ground-truth using a dataset synthesised from a 3D character model animated using motion capture data.

The highest performance is obtained from volume-sampling shape-histogram descriptors. The descriptors also demonstrate relative insensitivity to parameter setting. Principle component based alignment of surfaces demonstrated poor performance in comparison to testing similarity up to a rotation. Shape Distribution and Spin Images demonstrate an ambiguity to reflective symmetry. Spherical harmonics provide an ambiguous descriptor that greatly reduced the specificity defining similarity. While the evaluation is not exhaustive due to breadth of the literature on 3D shape similarity the work clearly demonstrates the advantage of volume sampling shape-histogram in the context of shape matching in temporal surface sequences of people.

Acknowledgements

This work was funded by EPSRC Grant EP/E001351.

References

- [1] M. Ankerst, G. Kastentmiller, H. P. Kriegel, and T. Seidl. 3D shape histograms for similarity search and classification in spatial databases. *Advances in Spatial Databases, 6th International Symposium, SSD'99*, 1651:207–228, 1999.
- [2] N. Aspert, D. Santa-Cruz, and T. Ebrahimi. Mesh: measuring errors between surfaces using the hausdorff distance. *ICME '02. Proceedings*, 1:705–708 vol.1, 2002.
- [3] S. Belongie, J. Malik, and J. Puzicha. Shape context: A new descriptor for shape matching and object recognition.

- In *NIPS*, pages 831–837, 2000.
- [4] B. Bustos, D. A. Keim, D. Saupe, T. Schreck, and D. V. Vrani. Feature-based similarity search in 3d object databases. *ACM Comput. Surv.*, 37(4):345–387, December 2005.
- [5] D. Y. Chen, M. Ouhyoung, X. P. Tian, and Y. T. Shen. On visual similarity based 3d model retrieval. *Eurographics*, pages 223–232, 2003.
- [6] C. S. Chua and R. Jarvis. Point signatures: A new representation for 3d object recognition. *International Journal of Computer Vision*, 25(1):63–85, October 1997.
- [7] J. Corney, H. Rea, D. Clark, J. Pritchard, M. Breaks, and R. Macleod. Coarse filters for shape matching. *Computer Graphics and Applications, IEEE*, 22(3):65–74, 2002.
- [8] M. El-Mehalawi. A database system of mechanical components based on geometric and topological similarity. part ii: indexing, retrieval, matching, and similarity assessment. *Computer-Aided Design*, 35(1):95–105, January 2003.
- [9] A. Elad and R. Kimmel. On bending invariant signatures for surfaces. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 25(10):1285–1295, 2003.
- [10] T. Funkhouser, M. Kazhdan, P. Min, and P. Shilane. Shape-based retrieval and analysis of 3d models. *Commun. ACM*, 48(6):58–64, June 2005.
- [11] M. Heczko, D. A. Keim, D. Saupe, and D. V. Vranic. Verfahren zur hnlichkeitssuche auf 3d-objekten. *Datenbank-Spektrum*, 2:54–63, 2002.
- [12] M. Hilaga, Y. Shinagawa, T. Kohmura, and T. L. Kunii. Topology matching for fully automatic similarity estimation of 3d shapes. In *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 203–212, New York, NY, USA, 2001. ACM Press.
- [13] N. Iyer, S. Jayanti, K. Lou, Y. Kalyanaraman, and K. Ramani. Three-dimensional shape searching: state-of-the-art review and future trends. *Computer-Aided Design*, 37(5):509–530, April 2005.
- [14] V. Jain and H. Zhang. A spectral approach to shape-based retrieval of articulated 3d models. *GMP '06*, 2006.
- [15] A. E. Johnson and M. Hebert. Using spin images for efficient object recognition in cluttered 3d scenes. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 21(5):433–449, 1999.
- [16] M. Kazhdan, B. Chazelle, D. P. Dobkin, A. Finkelstein, and T. A. Funkhouser. A reflective symmetry descriptor. In *ECCV*, volume 2, pages 642–656, 2002.
- [17] M. Kazhdan, T. Funkhouser, and S. Rusinkiewicz. Rotation invariant spherical harmonic representation of 3d shape descriptors. *SGP '03: Proceedings of the 2003 Eurographics/ACM SIGGRAPH symposium on Geometry processing*, pages 156–164, 2003.
- [18] M. Kortgen, G. J. Park, M. Novotni, and R. Klein. 3d shape matching with 3d shape contexts. In *the 7th Central European Seminar on Computer Graphics*, April 2003.
- [19] L. Kovar, M. Gleicher, and F. Pighin. Motion graphs. In *SIGGRAPH '02: Proceedings of the 29th annual conference on Computer graphics and interactive techniques*, volume 21, pages 473–482, New York, NY, USA, July 2002. ACM Press.
- [20] D. Mcwherter, M. Peabody, W. C. Regli, and A. Shokoufandeh. Solid model databases: Techniques and empirical results. *Journal of Computing and Information Science in Engineering*, 1(4):300–310, 2001.
- [21] C. Menier, E. Boyer, and B. Raffin. 3d skeleton-based body pose recovery. *Proceedings of the 3rd International Symposium on 3D Data Processing, Visualization and Transmission, Chapel Hill (USA)*, June 2006.
- [22] M. Mousa, R. Chaine, and S. Akkouche. Direct spherical harmonic transform of a triangulated mesh. *Journal of graphics tools*, 11(2):17–26, 2006.
- [23] M. Novotni and R. Klein. 3d zernike descriptors for content based shape retrieval. *The 8th ACM Symposium on Solid Modeling and Applications*, June 2003.
- [24] R. Ohbuchi, T. Minamitani, and T. Takei. Shape-similarity search of 3d models by using enhanced shape functions. *Theory and Practice of Computer Graphics, 2003. Proceedings*, pages 97–104, 2003.
- [25] R. Osada, T. Funkhouser, B. Chazelle, and D. Dobkin. Shape distributions. *ACM Trans. Graph.*, 21(4):807–832, October 2002.
- [26] E. Paquet. Description of shape information for 2-d and 3-d objects. *Sig. Proc.: Image Comm.*, 16:103–122, September 2000.
- [27] H. Y. Shum, M. Hebert, and K. Ikeuchi. On 3d shape similarity. *Proceedings of the 1996 Conference on Computer Vision and Pattern Recognition (CVPR '96)*, pages 526–531, June 1996.
- [28] J. Starck and A. Hilton. Surface capture for performance-based animation. *IEEE Computer Graphics and Applications*, 27(3):21–31, 2007.
- [29] J. Starck, G. Miller, and A. Hilton. Video-based character animation. In *SCA '05: Proceedings of the 2005 ACM SIGGRAPH/Eurographics symposium on Computer animation*, pages 49–58, New York, NY, USA, 2005. ACM Press.
- [30] H. Sundar, D. Silver, N. Gagvani, and S. Dickinson. Skeleton based shape matching and retrieval. *SMI '03: Proceedings of the Shape Modeling International 2003*, pages 130–130, 2003.
- [31] J. W. H. Tangelder and R. C. Veltkamp. A survey of content based 3d shape retrieval methods. *Shape Modeling Applications, 2004. Proceedings*, pages 145–156, 2004.
- [32] J. Xu, T. Yamasaki, and K. Aizawa. Motion editing in 3d video database. *3dprt*, 0:472–479, 2006.
- [33] T. Zaharia and F. Preteux. Three-dimensional shape-based retrieval within the mpeg-7 framework. volume 4304, pages 133–145, January 2001.
- [34] C. Zhang and T. Chen. Efficient feature extraction for 2d/3d objects in mesh representation. pages 935–938, 2001.