NRC Publications Archive
Archives des publications du CNRC

Posture invariant surface description and feature extraction
Wuhrer, Stefanie; Ben Azouz, Zouhour; Shu, Chang

NRC Publications Record / Notice d'Archives des publications de CNRC:
https://nrc-publications.canada.ca/eng/view/object/?id=34062eb9-2100-42c1-8171-a5fe6b805d25
https://publications-cnrc.canada.ca/fra/voir/objet/?id=34062eb9-2100-42c1-8171-a5fe6b805d25

Questions? Contact the NRC Publications Archive team at
PublicationsArchive-ArchivesPublications@nrc-cnrc.gc.ca. If you wish to email the authors directly, please see the
first page of the publication for their contact information.

Vous avez des questions? Nous pouvons vous aider. Pour communiquer directement avec un auteur, consultez la
première page de la revue dans laquelle son article a été publié afin de trouver ses coordonnées. Si vous n'arrivez
pas à les repérer, communiquez avec nous à PublicationsArchive-ArchivesPublications@nrc-cnrc.gc.ca.

National Research
Council Canada

Conseil national de
recherches Canada

Canada

Posture Invariant Surface Description and Feature Extraction

Stefanie Wuhrer[1]   Zouhour Ben Azouz[2]   Chang Shu[1]

Abstract

We propose a posture invariant surface descriptor for triangular meshes. Using intrinsic geometry, the surface is first transformed into a representation that is independent of the posture. Spin image is then adapted to derive a descriptor for the representation. The descriptor is used for extracting surface features automatically. It is invariant with respect to rigid and isometric deformations, and robust to noise and changes in resolution. The result is demonstrated by using the automatically extracted features to find correspondences between articulated meshes.

1. Introduction

Surface meshes obtained through active and passive 3-D sensing such as range scanning, structured-light projection, and stereo have increasingly become a new media for representing the 3-D shapes of the real-world objects. In practice, we need certain processing techniques that are similar to what can be found in image processing. Automatically extracting and corresponding prominent surface features are examples of these techniques that arise in various applications such shape correspondence [23], object recognition [10], and segmentation [11].

Existing methods that extract large numbers of features based on the geometry of a surface are not invariant with respect to significant non-rigid deformations. Hence, features extracted for different poses of an articulated object do not correspond to the same intrinsic locations. In this paper, we consider the problem of posture invariant surface description and feature extraction. We derive a posture invariant surface descriptor and use it to extract features. This way, we can extract large numbers of features for meshes with significant articulation. Note that these corresponding features can be used to find dense point-to-point correspondences between articulated shapes [23].

More precisely, this paper proposes a descriptor that is invariant with respect to rigid and *isometric* transformations of an triangular manifold mesh S. A transformation of S is

isometric if and only if all of the lengths measured on S are preserved. Isometric transformations are relevant in practice because the locomotions of humans and many animals are approximately isometric.

The basic idea is to first map different shapes into a space in which the postures of the same shape under isometric transformations have similar representations. Feature extraction can subsequently be performed on the posture-invariant representation. One approach for this is to map the intrinsic geometry of S into $\mathbb{R}^3$ using Multi-Dimensional Scaling (MDS). This mapping removes the non-intrinsic shape variations and thus obtains a posture-invariant representation known as a *canonical form* [7].

To obtain a posture invariant descriptor, we propose to use a novel modification of the Spin image [10] computed on the canonical form of S. A spin image is a rotation-invariant surface descriptor that is obtained by projecting part of the surface into the local coordinate system of a vertex. By varying the projection area, it is possible to control the descriptor continuously from local to global.

Spin images are defined on oriented points, which are 3-D points with surface normal. However, the canonical form obtained through MDS does not preserve the surface orientation. We will show, in Section 3.3, that Spin images can be modified to allow their use on non-oriented surfaces.

Finally, we extract features on S by finding the vertices of S that have the most unusual descriptors. We demonstrate that the proposed feature finder is robust to surface noise and change of resolution. We also validate the approach by finding corresponding features between pairs of poses of articulated models.

2. Related Work

This section reviews approaches that describe a surface locally. First, we review approaches that are suitable for large posture change. Katz et al. [11] proposed to extract as feature points vertices that correspond to local maxima of the average geodesic distance function and that are located on the convex hull of the canonical form [7]. They used these feature points to compute a posture invariant segmentation. Due to the restriction that only vertices on the convex hull of the canonical form are selected, this approach can only extract features that are tips of prominent com-

---

[1]National Research Council of Canada, Ottawa, Canada.

[2]Institut Supérieur d'Informatique and Ecole Supérieur Privée d'Ingénieurs et de Technologies, Tunis.

ponents. Furthermore, no special care is taken to address possibly self-intersecting canonical forms. While this suffices to compute a segmentation, applications such as shape recognition or shape correspondence require more features. Tierny et al. [17] developed shape signatures to locally describe patches of the shape. The approach uses the Reeb graph to decompose a shape into patches (disks and annuli) and maps each patch to a plane in a canonical way. The surface is described by the amount of stretch required to map the patches to their canonical shape in the plane. Note that not every vertex has a descriptor. Hence, this approach cannot be used to extract features. Hilaga et al. [9] proposed a descriptor based on the average geodesic distance function on a surface. Zhang et al. [22] used this descriptor combined with Reeb graphs to find features. However, they did not show that the approach is suitable for posture invariant feature extraction.

Second, we review approaches that allow for restricted posture change. Gelfand et al. [8] proposed a surface descriptor based on local volume. Hence, this descriptor is invariant with respect to deformations that locally maintain the volume. Gelfand et al. did not demonstrate that the descriptor is posture invariant with respect to any deformation in practice. Zaharescu et al. [20] recently proposed a descriptor that can incorporate both geometric and photometric information. The descriptor is an extension of SIFT [14] in that it uses a scale space. Although the descriptor is applied to compute the correspondences between two non-rigid objects, all of the illustrated deformations are small. The descriptor proposed in this paper is applied to models with large deformations.

Third, we review approaches that can be adapted to allow for posture changes. Note that none of the following approaches allows for posture changes in its current form. Zhang and Hebert [21] computed a descriptor by computing a local parameterization of a surface patch using harmonic maps. The harmonic map is then transformed to an image by giving each pixel a color corresponding to the mean curvature at the corresponding surface point. Sun and Abidi [16] proposed a descriptor that projects geodesic contours at a vertex p to the tangent plane at p. While the geodesic contours are posture invariant, the orientation of the tangent plane at p is not.

3. Surface Descriptor

This section outlines a surface descriptor that is invariant with respect to rigid transformations and isometric deformations. Suppose S contain n vertices. Denote a vertex of S by p and denote the position vector of p by $\vec{p}$. For each vertex, we compute a descriptor that is posture invariant. The descriptor at a vertex p of S is computed as a modification of the Spin image at p of the canonical form of S. In the following, we review canonical forms and Spin images.

3.1. Canonical Form

Elad and Kimmel [7] define the *canonical form* X of a triangular manifold S as the mapping of S to $\mathbb{R}^3$, such that the Euclidean distances between the mapped vertices approximate the geodesic distances between the original vertices well. Let $\delta(p, q)$ denote the geodesic distance on S between two vertices p and q on S and let $d_X(p, q)$ denote the Euclidean distance between the two points in X that correspond to p and q. The canonical form of X is computed via multi-dimensional scaling with the geodesic distances between vertices on the triangular manifold as dissimilarities. That is, the canonical form is the embedding that minimizes $E = \sum_{p, q \in S} (\delta(p, q) - d_X(p, q))^2$. The canonical form has the desirable property that it is posture invariant. Note that the energy E is invariant with respect to translation, rotation, and reflection of X.

We use fast marching [12] to compute geodesic distances on S. We then use least-squares multi-dimensional scaling [2, p.146-155] to compute the canonical form in $\mathbb{R}^3$. For increased space efficiency, we compute the canonical form using a coarse-to-fine strategy as outlined by Wuhrer et al. [19]. The approach by Wuhrer et al. consists of two steps. First, n' vertices of S are used to compute a canonical form at low resolution. Second, the remaining vertices of S are added to the canonical form one by one by minimizing a least-squares energy function.

The canonical form X of S is a set of vertices in $\mathbb{R}^3$. We can enhance X by adding the edges and triangles present in S to their corresponding vertices in X. This yields a possibly self-intersecting triangular manifold mesh. On this mesh, we can estimate the normal at a vertex p as the average of the normals of the faces incident to p. In the following, when referring to X, we refer to this mesh. Since X is invariant with respect to reflections, the orientation of the normal of X at p is ambiguous.

3.2. Spin Image

Johnson and Hebert [10] define the *Spin image* of an oriented triangular manifold S at a vertex p as an image obtained by projecting part of S into a local coordinate system defined by the surface normal $\vec{n}(p)$ at p. The projection of a point q to the Spin image of p uses a reminiscent of cylindrical coordinates. That is, the x-coordinate of the Spin image corresponds to the distance from q to the line in direction $\vec{n}(p)$ passing through p and the y-coordinate of the Spin image corresponds to the signed distance from q to the tangent plane of p. It remains to define what part of S is projected to the Spin image of p. This depends on three parameters: the width w of the (square) Spin image, the bin size b of the Spin image, and the support angle α. A vertex q of S is projected to the Spin image of p if and only if $\|q - p\| < wbr$ and $\angle(\vec{n}(p), \vec{q} - \vec{p}) < \alpha$, where $\|q - p\|$

denotes the Euclidean distance between q and q, r denotes the resolution of the mesh, and $\angle(\vec{n}(p), \vec{q} - \vec{p})$ denotes the angle between the two vectors $\vec{n}(p)$ and $\vec{q} - \vec{p}$.

Spin images have a variety of desirable properties. They are invariant with respect to rotation, translation, and scaling and robust with respect to noise and clutter. By adjusting w or b, Spin images can go from a local representation at p to a global representation of S.

3.3. Descriptor

We combine canonical forms and Spin images to obtain a surface descriptor that is invariant with respect to rigid transformations and isometric deformations. This is achieved by eliminating posture dependence using the canonical form X of S and by eliminating dependence on rigid transformations using Spin images of X. That is, for a vertex p of S, we use as descriptor a modification of the Spin image of X at the vertex of X corresponding to p. We next describe the modification of the Spin images.

Spin images are defined for oriented surfaces. However, we wish to use Spin images of the canonical form X as surface descriptors. Recall that orientation of normals on X are ambiguous. By setting the support angle α to $180°$ and by folding the Spin image along the median y-coordinate, such that the y-axis corresponds to the unsigned distance, we can eliminate the dependence on orientation. This is required when using Spin images of canonical forms. Denote the modified Spin image by *folded Spin image*. In the following, we denote the folded Spin image of X at the vertex of X corresponding to p with parameters b and w by $D_{b,w}(p)$ and we use $D_{b,w}(p)$ as surface descriptor.

3.4. Dissimilarity

To extract features based on the surface descriptors, we need to compare the surface descriptors $D_{b,w}(p)$ and $D_{b,w}(q)$, where p and q are vertices of S. We use a dissimilarity based on the normalized linear correlation coefficient $R(D_{b,w}(p), D_{b,w}(q))$ between $D_{b,w}(p)$ and $D_{b,w}(q)$ as suggested by Johnson and Hebert [10]. The coefficient $R(D_{b,w}(p), D_{b,w}(q))$ takes values between -1 and 1. We define the dissimilarity $d(D_{b,w}(p), D_{b,w}(q))$ as $1 - R(D_{b,w}(p), D_{b,w}(q))$. Note that $d(D_{b,w}(p), D_{b,w}(q))$ is always positive. The smaller $d(D_{b,w}(p), D_{b,w}(q))$, the larger the similarity between $D_{b,w}(p)$ and $D_{b,w}(q)$.

3.5. Properties

The surface descriptor $D_{b,w}(p)$ has a number of desirable properties.

- Invariance with respect to rigid transformations since folded Spin images have this property.

- Invariance with respect to isometric deformations since canonical forms have this property.

- Can capture both local and global shape properties by varying b and w.

- Robustness with respect to noise. Although canonical forms are not invariant to noise, we demonstrate in Section 5 that features extracted using this descriptor are robust with respect to small amounts of noise that do not alter the global shape of the canonical form. This is due to the use of multiple levels of folded Spin images as outlined below.

- Robustness with respect to changes in the mesh resolution. Although folded Spin images are not invariant to the resolution of the mesh, we demonstrate in Section 5 that features extracted using this descriptor are robust with respect to changes in the mesh resolution.

4. Feature Extraction

This section introduces an algorithm to extract surface features based on the descriptor introduced in the previous section. We discuss a basic algorithm and its variant for improved efficiency.

4.1. Basic Feature Extraction

We aim to extract features that are robust with respect to noise and changes in resolution. When scalar surface descriptors are considered, reducing sensitivity to noise is often achieved using a scale-space representation [20].

The descriptor proposed in this paper is not a scalar function. However, since our surface descriptor can describe the surface both locally and globally, we use m *levels of Spin images* to make the feature extraction algorithm robust to noise as follows. We start with a local descriptor computed with bin size b (level zero) and gradually increase the bin size of the descriptor to capture more global shape information for higher levels. The descriptor at level i is $D_{ib,w}(p)$.

The approach extracts features in each of the m levels independently and combines the features in different levels. We extract k features in level i as the k vertices with the most unusual descriptors $D_{ib,w}(p)$.

To find the vertices that have the most unusual descriptors in level i, we use a modification of the approach by Gelfand et al. [8]. We proceed in two steps.

First, we compute the descriptors at level i for all vertices of S. Our goal is to efficiently compute a reminiscent of a histogram of the descriptors based on the pairwise dissimilarities $d(D_{ib,w}(p), D_{ib,w}(q))$, where p and q are two vertices of S. Instead of computing a histogram in a high-dimensional space, we compute for each vertex p the average of the dissimilarities between the descriptor of p and its

l most similar descriptors. Using an average of l most similar descriptors instead of the most similar descriptor makes the approach more robust to noise. We set $l = \lceil \frac{n}{100} \rceil$ in our experiments.

Second, we extract as features the vertices with the largest average values. These vertices have a descriptor that is most dissimilar from its most similar descriptors. If two vertices have the same average value, we cannot distinguish them and we need to either extract both vertices as features or none of them. Our algorithm repeatedly finds sets of features that have the largest average value not yet considered. We add the set to the feature set as long as the total number of features does not exceed k. Otherwise, the algorithm terminates.

To avoid double-counting of features, when extracting features, we do not accept a new feature that is located within a 2-ring neighborhood of an existing feature.

In a last step, the features found over m levels are combined as follows. A feature is reported if it was extracted in at least two consecutive levels. That is, if p was extracted as feature at level i and at least one vertex in 2-ring neighborhood of p was extracted as feature at level $i - 1$ or level $i + 1$, then we consider p to be a feature. As before, we do not accept a new feature that is located within a 2-ring neighborhood of an existing feature.

Finally, we analyze the running time for the feature extraction. Since both m and k are constants, it takes on average $O(n^2 \log n)$ time to extract the features.

Note that the feature extraction algorithm has the desirable property that if we extract a set F_1 of k_1 features and if we extract a set F_2 of $k_2 > k_1$ features using the same parameters b, w, and m, then $F_1 \subseteq F_2$.

4.2. Sampling for Improved Efficiency

Computing folded Spin images of all of the vertices of X at each level is time consuming. We therefore propose to use only a sample set of the original vertices for feature extraction. That is, we only consider a sample set P of n'' vertices of S as possible features. This has the effect that we only need to compute n'' folded Spin images at each level. Note that we do not change the resolution of S.

While this approach makes the algorithm more time efficient, it produces results that are less accurate, thereby giving a way to trade off between time and accuracy.

In our experiments, we set $n'' = n'$. This way, we only need to compute one set of samples. When computing the folded Spin image of a sample vertex v, we use the average of all the normals on X corresponding to vertices in v's 2-ring neighborhood.

5. Results

The experiments were conducted on an Intel Pentium D with 3.5 GB of RAM using a C++ implementation. The cat models contain 7207 vertices and the horse models contain 8431 vertices. Both data sets were created by Sumner et al. [15]. The models of a female dancer and a male dancer contain 3400 vertices and were created by Bronstein et al. [5]. The models of a female dancer in a skirt were created by Vlasic et al. [18]. We down sampled these models to contain about 4000 vertices.

For the experiments in this paper, if each vertex is a possible feature, we set $w = 10$, $b = 0.25$, $m = 5$. If the sampling strategy is used, we set $w = 10$, $b = 0.25$, and $m = 3$. In all of the figures, features are shown in red.

5.1. Properties

We demonstrate the properties of using the proposed descriptor for feature extraction. Namely, we demonstrate that descriptor is robust with respect to non-rigid deformations, Gaussian noise, and changes in resolution. We set $k = 20$. For the following experiments, we do not use the sampling strategy.

Figures 1 (top) and 2 show the effect of non-rigid deformations of the mesh. Although the poses vary significantly, similar features are found.
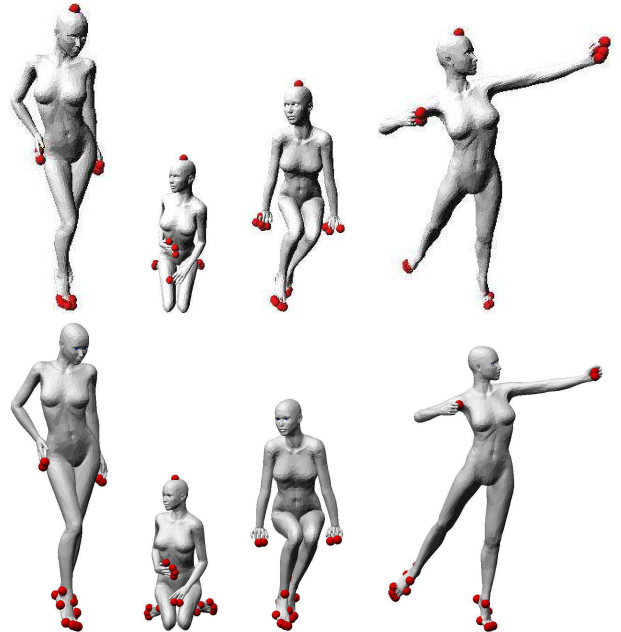

Figure 1. *Top: Effect of non-rigid deformation model. Bottom: Effect of the sampling strategy.*

Figure 3 shows the effect of adding noise to the mesh. We add random Gaussian noise to the model of a cat. The noise has mean $0.005r$ (middle) and $0.01r$ (right), where r
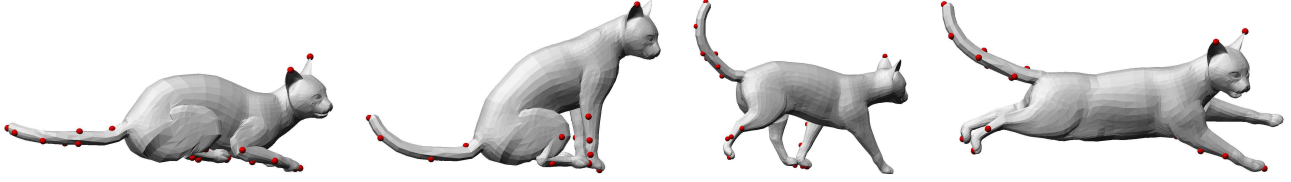
Figure 2. *Effect of non-rigid deformation model.*

Figure 3. *Effect of adding Gaussian noise to the model.*

is the average edge length of the mesh. We can see that most of the features are preserved.

Figure 4 shows the effect of changing the resolution of the mesh. Although the proposed descriptor is not invariant with respect to the resolution of the mesh in theory, we can see that similar features are found for all resolutions. The reason is the use of multiple levels of Spin images.

The bottom row of Figure 1 shows the features found for the dancer model when using the sampling strategy. Note that the features are similar to the ones shown in the top row of Figure 1. With the sampling strategy, it takes about 39 seconds to extract the features for each of the poses. When each vertex is considered as a possible feature, it takes about 140 seconds to extract the features for each of the poses.

5.2. Comparison

We compare the features extracted using our approach to the features extracted by Katz et al. [11]. For our approach, we set $k = 60$. For Katz et al., the number of extracted features is fixed. We use qhull for the convex hull computation [1].

Figure 5 shows the features extracted on different poses of a horse. The top row shows the result using our algorithm and the bottom row shows the result using the algorithm by Katz et al. While the features extracted by Katz et al.'s algorithm are posture invariant, there are few features because it only extracts features at extremities of prominent components. That is, the algorithm cannot extract features on joints of an animal. In the example shown in Figure 5, not every extremity contains features. For instance, no features are extracted on the front legs of the horse. In contrast, our algorithm is able to extract many posture invariant features located at extremities and joints of prominent components.

5.3. Application to Surface Correspondence

We validate our approach by computing correspondences between the features extracted on articulated models. The correspondences between features can be used to find dense point-to-point correspondences [23]. In all of the

experiments, corresponding vertices are shown in the same color.

The correspondence is computed greedily by corresponding the vertices with the most similar descriptor values. We compute the similarity between two feature points by averaging the descriptor similarities between the two features over all m levels. Furthermore, we compute the dissimilarity between two features as the minimum dissimilarity between descriptor values in a 2-ring neighborhood of the features. If the most similar descriptor value is not symmetric for a feature, then no correspondence is assigned.

Note that this greedy correspondence is intended for validation only. We do not intend to solve the correspondence problem efficiently in this paper.

We first compute the correspondences for the horse dataset with known ground truth. The experiment computes the correspondences between the features extracted on two pairs of poses. We set $k = 60$. Figure 6 shows the feature correspondences that are obtained with and without the sampling strategy. Note that not all of the features are visible in this figure. Figure 7 shows the histograms of the errors of the feature correspondences with respect to the ground truth correspondences. The error is measured as the number of edges on the shortest path between the ground truth correspondence and the computed feature correspondence. This error measure is suitable for nearly uniform meshes. Note that this means that errors obtained by matching features to symmetric parts of the horse are potentially huge. However, since a surface descriptor cannot distinguish locally between symmetric parts, these symmetric mismatches should not be counted as errors.

We discuss the results obtained with and without using the sampling strategy. First, we discuss the results obtained when all of the vertices are possible features. It takes about 10.5 minutes to compute the features for each of the poses. For the pair of poses on the right, all but four of the mismatched features are matched to symmetric body parts. For the pair of poses on the left, all but two of the mismatched features are matched to symmetric body parts. Hence, the

Figure 4. *Effect of changing the resolution of the model. The left model contains* 14410 *triangles, the model in the middle contains* 5000 *triangles, and the right model contains* 1024 *triangles.*
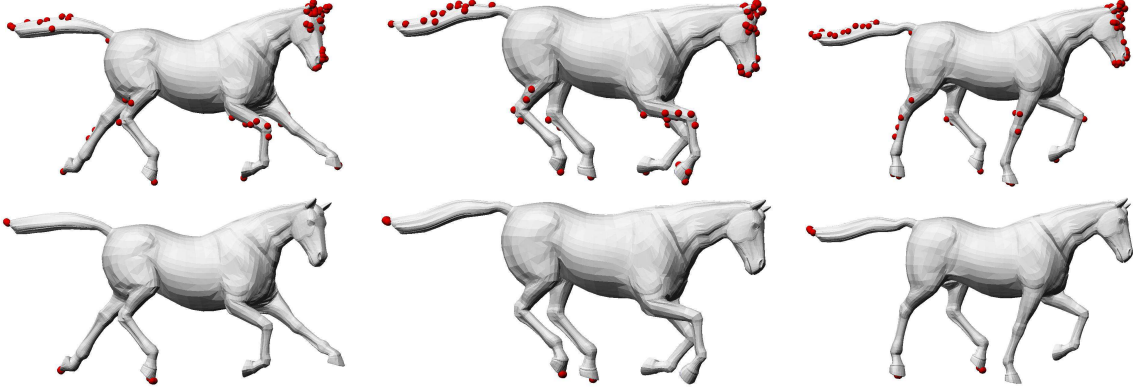

Figure 5. *Comparison. Top: Result of our algorithm with* $k = 60$. *Bottom: Result by Katz et al. [11].*

accuracy of the computed correspondences is high.

Second, we discuss the results obtained when the sampling strategy is used. It takes about 4 minutes to compute the features for each of the poses. For the pair of poses on the right, all but two of the mismatched features are matched to symmetric body parts. For the pair of poses on the left, all but three of the mismatched features are matched to symmetric body parts. The accuracy of the result is lower than before, as can be seen in Figure 7. Furthermore, fewer features are extracted when the sampling strategy is used.

All of the following experiments use all of the vertices as possible features. Figure 8 shows the correspondences between the features extracted with $k = 20$ on three pairs of dancers. Note that all of the correspondences are either close to the true correspondence or on symmetric body parts.

Finally, we use the method to establish correspondences between surfaces with similar, but not identical intrinsic geometry. Figure 9 shows the correspondences between the features extracted with $k = 20$ on a male and a female dancer. Although both models are humans, they are not isometric to each other. Note that all but one of the correspondences are either close to the true correspondence or on symmetric body parts of the dancers. Figure 10 shows the correspondences between the features extracted with $k = 60$ on two female dancers wearing a skirt. Note that the surface joining the skirt and the legs of the dancer moves in a non-isometric fashion during the motion of the dancer. The left of Figure 10 shows the front view of the pair of dancers and the right of Figure 10 shows the back and side view of the same pair of dancers. Note that most of the

correspondences are either close to the true correspondence or on symmetric body parts of the dancers. Hence, the approach is suitable for surfaces with similar, but not identical intrinsic geometry.
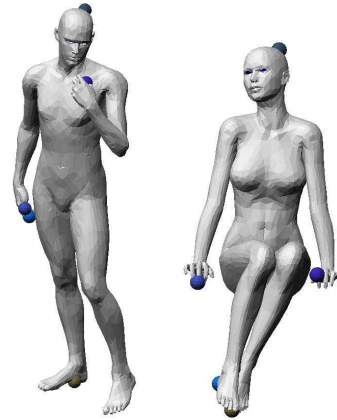

Figure 9. *Correspondences.*

6. Conclusion

This paper introduces a posture invariant shape descriptor and applies it to feature extraction and feature correspondence. The descriptor is invariant with respect to rigid and isometric transformations and it is robust with respect to noise and changes in mesh resolution. Furthermore, the descriptor can describe the shape locally and globally.

The user can set the following parameters: the number of features to be extracted, the bin size b, the width w, and the number of levels m. We leave it for future work to find
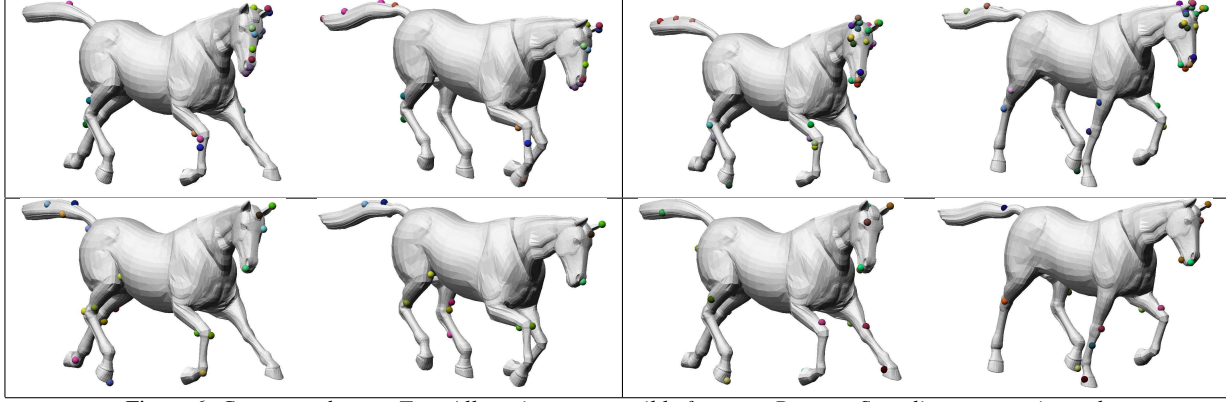
Figure 6. *Correspondences. Top: All vertices are possible features. Bottom: Sampling strategy is used.*
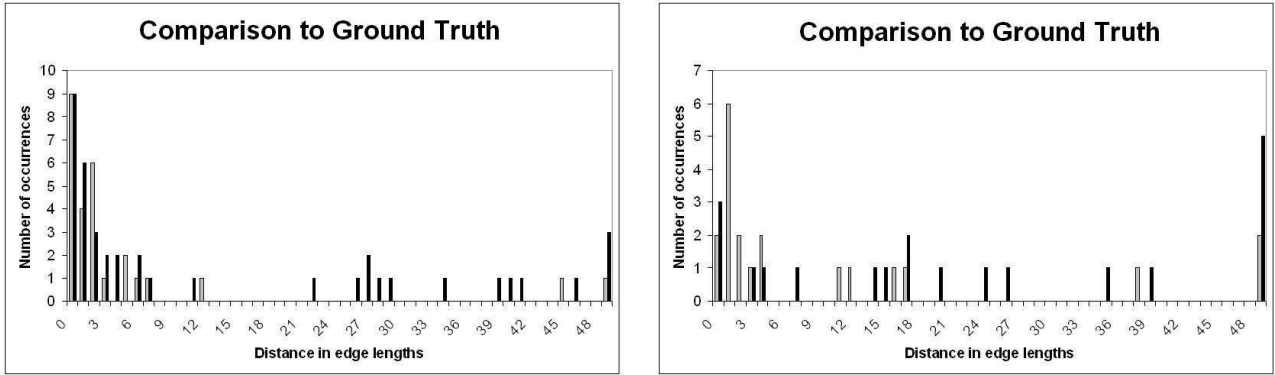

Figure 7. *Comparison to ground truth. Grey columns correspond to the poses in Figure 6(left). Black columns correspond to the poses in Figure 6(right). Left: All vertices are possible features. Right: Sampling strategy is used.*

the best b, w, and m automatically.

We leave the following limitations for future work.

- The descriptor is not suitable for surfaces with non-Euclidean geometry. This limitation can be remedied if prior knowledge about the surface's intrinsic geometry is available by using generalized multi-dimensional scaling [4].

- The descriptor is sensitive with respect to topological noise because geodesics are used to compute the canonical forms.

- For meshes with boundaries, many features are extracted along boundaries of the mesh because vertices along a boundary have fewer close neighbors.

- The approach is not suitable for partial matching because two globally different shapes have different canonical forms.

References

[1] C. Barber, D. Dobkin, and H. Huhdanpaa. The Quickhull algorithm for convex hulls. *ACM Trans. Math. Software*, 22(4):469–483, 1996.

[2] I. Borg and P. Groenen. *Modern Multidimensional Scaling Theory and Applications*. Springer, 1997.

[3] A. Bronstein, M. Bronstein, and R. Kimmel. *Numerical Geometry of Non-Rigid Shapes*. Springer, 2008.

[4] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *PNAS*, 103(5):1168–1172, 2006.

[5] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Calculus of non-rigid surfaces for geometry and texture manipulation. *IEEE TVCG*, 13(5):902–913, 2007.

[6] D. Burnsides and K. Robinette. 3-D landmark detection and identification in the CAESAR project. In *3DIM*, pp. 393–298, 2001.

[7] A. Elad and R. Kimmel. On bending invariant signatures for surfaces. *IEEE TPAMI*, 25(10):1285–1295, 2003.

[8] N. Gelfand, N. J. Mitra, L. J. Guibas, and H. Pottmann. Robust global registration. *SGP*, p. 197, 2005.

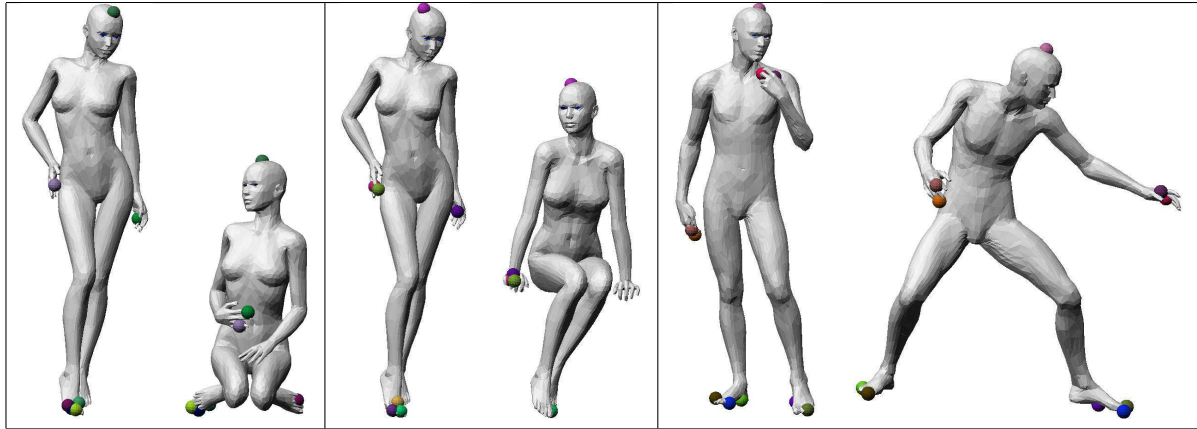[9] M. Hilaga, Y. Shinagawa, T. Kohmura, and T. L. Kunii. Topology matching for fully automatic similarity
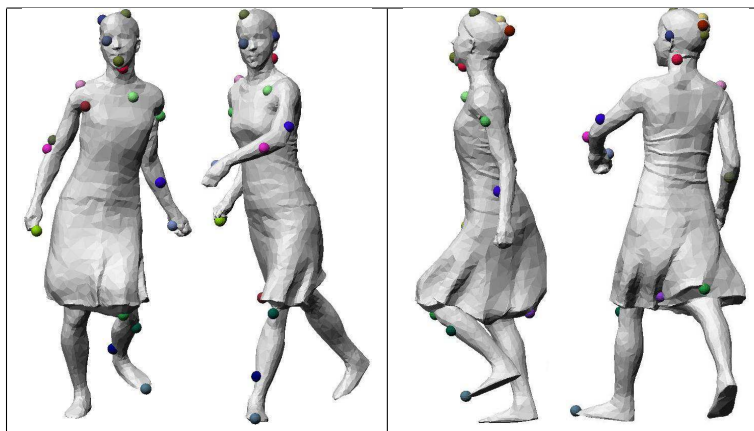
Figure 8. *Correspondences.*

Figure 10. *Correspondences. Left: front view of dancers. Right: back and side view of dancers.*

estimation of 3d shapes. *SIGGRAPH*, pp. 203–212, 2001.

[10] A. E. Johnson and M. Hebert. Using spin images for efficient object recognition in cluttered 3d scenes. *IEEE TPAMI*, 21(5):433–449, 1999.

[11] S. Katz, G. Leifman, and A. Tal. Mesh segmentation using feature point and core extraction. *Vis. Comp.*, 21(8–10):649–658, 2005.

[12] R. Kimmel and J. Sethian. Computing geodesic paths on manifolds. *PNAS*, 95:8431–8435, 1998.

[13] C. H. Lee, A. Varshney, and D. W. Jacobs. Mesh saliency. *ACM TOG*, 24(3):659–666, 2005.

[14] D. G. Lowe. Object recognition from local scale-invariant features. *ICCV*, p. 1150, 1999.

[15] R. W. Sumner and J. Popović. Deformation transfer for triangle meshes. *ACM TOG*, 23(3):399–405, 2004.

[16] Y.-Y. Sun and M. A. Abidi. Surface matching by 3d point's fingerprint. *ICCV*, 2001.

[17] J. Tierny, J.-P. Vandeborre, and M. Daoudi. Reeb chart unfolding based 3D shape signatures. *Eurographics*, pp. 13–16, 2007.

[18] D. Vlasic, I. Baran, W. Matusik, and J. Popović". Articulated Mesh Animation from Multi-view Silhouettes. *ACM TOG*, 27(3):97, 2008.

[19] S. Wuhrer, C. Shu, Z. B. Azouz, and P. Bose. Posture invariant correspondence of incomplete triangular manifolds. *IJSM*, 13(2):139–157, 2007.

[20] A. Zaharescu, E. Boyer, K. Varanasi, and R. P. Horaud. Surface feature detection and description with applications to mesh matching. *CVPR*, 2009.

[21] D. Zhang and M. Hebert. Harmonic shape images: A representation for 3d free-form surfaces based on energy minimization. *EMMCVPR*, pp. 30–43, 1999.

[22] E. Zhang, K. Mischaikow, and G. Turk. Feature-based surface parameterization and texture mapping. *ACM TOG*, 24(1):1–27, 2005.

[23] H. Zhang, A. Sheffer, D. Cohen-Or, Q. Zhou, O. van Kaick, and A. Tagliasacchi. Deformation-Driven Shape Correspondence. *CGF*, 27(5):1431-1439, 2008.