



3D Shape Matching by Geodesic Eccentricity

Adrian Ion, Nicole M. Artner, Gabriel Peyré, Salvador B. Lopez Marmol,
Walter G. Kropatsch, Laurent D. Cohen

► To cite this version:

Adrian Ion, Nicole M. Artner, Gabriel Peyré, Salvador B. Lopez Marmol, Walter G. Kropatsch, et al.. 3D Shape Matching by Geodesic Eccentricity. Search in 3D 2008, Jun 2008, Anchorage, Alaska, United States. pp.1-8, 10.1109/CVPRW.2008.4563032 . hal-00419826

HAL Id: hal-00419826

<https://hal.science/hal-00419826>

Submitted on 25 Sep 2009

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

3D Shape Matching by Geodesic Eccentricity*

Adrian Ion

PRIP, Vienna University of Technology
ion@rip.tuwien.ac.at

Gabriel Peyré

CEREMADE, Université Paris-Dauphine
peyre@ceremade.dauphine.fr

Walter G. Kropatsch

PRIP, Vienna University of Technology
krw@rip.tuwien.ac.at

Nicole M. Artner

ARC, Smart Systems Division, Austria
nicole.artner@arcs.ac.at

Salvador B. López Mármol

PRIP, Vienna University of Technology
salva@rip.tuwien.ac.at

Laurent Cohen

CEREMADE, Université Paris-Dauphine
cohen@ceremade.dauphine.fr

Abstract

This paper makes use of the continuous eccentricity transform to perform 3D shape matching. The eccentricity transform has already been proved useful in a discrete graph-theoretic setting and has been applied to 2D shape matching. We show how these ideas extend to higher dimensions. The eccentricity transform is used to compute descriptors for 3D shapes. These descriptors are defined as histograms of the eccentricity transform and are naturally invariant to euclidean motion and articulation. They show promising results for shape discrimination.

1. Introduction

Shape recognition is a central topic in computer vision. It requires to set up a signature that characterizes the properties of interest for the recognition [27]. The invariance of the signature or the matching method to changes in orientation and scale, and to local deformations such as articulation is important for the identification of 2D and 3D shapes.

Applications of shape recognition can be found in Computer Aided Design/Computer Aided Manufacturing (CAD/CAM), virtual reality (VR), medicine, molecular biology, security, and entertainment [3].

Existing approaches can be divided into [3]: **Statistical** descriptors, like for example geometric *3D moments* employed by [6, 19] and the *shape distribution* [17, 12]. **Extension-based** descriptors, which

are calculated from features sampled along certain directions from a position within the object [30, 28]. **Volume-based** descriptors use the volumetric representation of a 3D object to extract features (examples are *Shape histograms* [1], *Model Voxelization* [29] and point set methods [26]). Descriptors using the **surface geometry** compute curvature measures and/or the distribution of surface normal vectors [20, 32]. **Image-based** descriptors reduce the problem of 3D shape matching to an image similarity problem by comparing 2D projections of the 3D objects [2, 5, 4]. Methods matching the **topology** of the two objects (for example *Reeb* graphs, where the topology of the 3D object is described by a graph structure [8, 23]). **Skeletons** are intuitive object descriptions and can be obtained from a 3D object by applying a thinning algorithm on the voxelization of a solid object like in [25]. Descriptors using **spin images** work with a set of 2D histograms of the object geometry and a search for point-to-point correspondences is done to match 3D objects [13].

The majority of shape descriptors is quite complex and not invariant to the deformation or articulation of object parts. In this article we extend the simple and efficient approach to shape matching with the eccentricity transform [11] to 3D. Based on the computation of geodesic lengths, eccentricity histograms are robust with respect to changes in orientation, scale and articulation of the object. The presented approach could be fitted to either of the categories **extension-based** or **volume-based**. In Section 2 the eccentricity transform is recalled. Section 3 presents the eccentricity histogram descriptors. Experimental results and discussion follow in Section 4.

* Partially supported by the Austrian Science Fund under grants S9103-N13 and P18716-N13.

2. Eccentricity Transform

The following definitions and properties follow [15, 10], and are extended to n -dimensional domains.

Let the shape \mathcal{S} be a closed set in \mathbb{R}^n . A (geodesic) path π is the continuous mapping from the interval $[0, 1]$ to \mathcal{S} . Let $\Pi(\mathbf{p}_1, \mathbf{p}_2)$ be the set of all paths between two points $\mathbf{p}_1, \mathbf{p}_2 \in \mathcal{S}$ within the set \mathcal{S} . The geodesic distance $d(\mathbf{p}_1, \mathbf{p}_2)$ between two points $\mathbf{p}_1, \mathbf{p}_2 \in \mathcal{S}$ is defined as the length $\lambda(\pi)$ of the shortest path $\pi \in \Pi(\mathbf{p}_1, \mathbf{p}_2)$ between \mathbf{p}_1 and \mathbf{p}_2

$$d(\mathbf{p}_1, \mathbf{p}_2) = \min\{\lambda(\pi(\mathbf{p}_1, \mathbf{p}_2)) | \pi \in \Pi(\mathbf{p}_1, \mathbf{p}_2)\} \quad (1)$$

where

$$\lambda(\pi(t)) = \int_0^1 |\dot{\pi}(t)| dt$$

and $\pi(t)$ is a parametrization of the path from $\mathbf{p}_1 = \pi(0)$ to $\mathbf{p}_2 = \pi(1)$. The *eccentricity transform* of \mathcal{S} can be defined as, $\forall \mathbf{p} \in \mathcal{S}$

$$ECC(\mathcal{S}, \mathbf{p}) = \max\{d(\mathbf{p}, \mathbf{q}) | \mathbf{q} \in \mathcal{S}\}, \quad (2)$$

i.e. to each point \mathbf{p} it assigns the length of the shortest geodesics to the points farthest away from it.

In [15] it is shown that this transformation is quasi-invariant to articulated motion and robust against salt and pepper noise (which creates holes in the shape). An analysis of the variation of geodesic distance under articulation can be found in [16].

The definition above accommodates n -dimensional objects embedded in \mathbb{R}^n as well as n -dimensional objects embedded in higher dimensional spaces (e.g. the 2D manifold given by the surface of a closed 3D object). A definition for graphs is given in [15].

This paper considers the class of 6-connected discrete shapes \mathcal{S} defined by points on a square grid \mathbb{Z}^3 . Paths need to be contained in the area of \mathbb{R}^3 defined by the union of the support cubes for the voxels of \mathcal{S} . The distance between any two voxels, whose connecting segment is contained in \mathcal{S} , is computed using the ℓ_2 -norm. (For easy understanding, some images with 2D shapes are also shown). The eccentricity transform is computed for:

- the *whole object*, i.e. \mathcal{S} in Equations 1 and 2 is the 3D object itself (all voxels of the object);
- the *border voxels* of the object, i.e. \mathcal{S} in Equations 1 and 2 is the 6 connected voxel surface in 3D, made out of voxels of the object that are 26 connected to a background voxel.

An *eccentric point* is a point that reaches a maximum in Equation 2, and all eccentric points lie on the border of \mathcal{S} [15]. The *center* is the set of points that have the smallest eccentricity.

2.1. Computation

The straight forward approach is: for each point of \mathcal{S} , compute the distance to all other points and take the maximum. In [10] faster computation and efficient approximation algorithms are presented. For this paper, the fastest one, algorithm ECC06, has been used (see [10] for an analysis of the speed/error performance).

ECC06 relies on the computation of the *shape bounded single source distance transform*¹ $DT(\mathcal{S}, \mathbf{p})$ (Figure 1.b), which is computed for estimated eccentric point candidates in an iterative manner. The shape bounded single source distance transform $DT(\mathcal{S}, \mathbf{p})$ associates to each point $\mathbf{q} \in \mathcal{S}$ the geodesic distance to \mathbf{p} . The 3D DT can be computed using Fast Marching [21] which allows for an efficient computation in $O(n \log(n))$ steps for a grid with n points.

Figure 1 shows a comparison of the geodesic and Euclidean distances. Figures 2 and 3 show the eccentricity transform of two shapes. For the case of the 3D shape, the eccentricity transform is presented for the whole object and for the border voxels.




Figure 1. An example of (a) Euclidean and (b) geodesic distance function, for some starting point p_0 .

3. Eccentricity Histogram Matching

To match two binary shapes, we first create a shape descriptor for each of them, and then match these descriptors to obtain a similarity measure. A first ap-

¹Also called geodesic distance function.




Figure 2. Left: a 2D binary shape. Right: eccentricity transform ECC - same gray value assignment as in Figure 1 (the point with smallest eccentricity is marked).

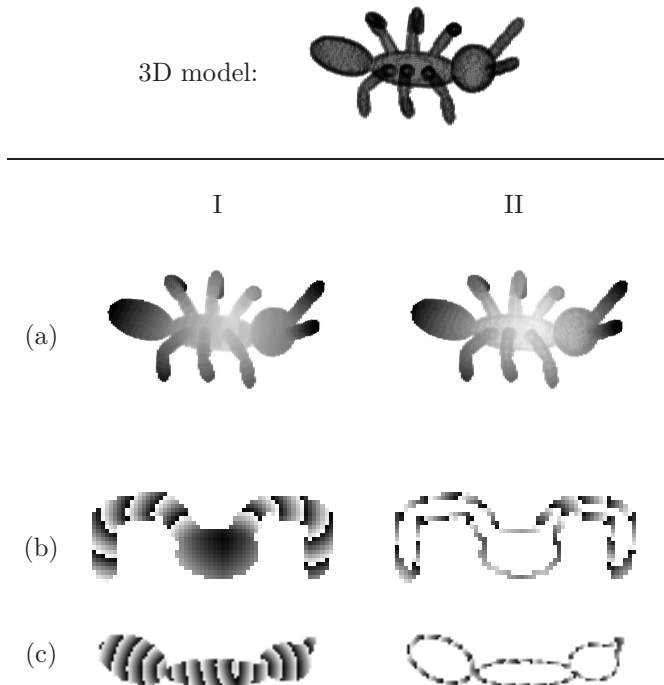


Figure 3. Top: 3D model of an ant. Column I: eccentricity transform of the ant. Column II: eccentricity transform of the border voxels of the ant. (a) Volume visualization of the eccentricity transform - darker means higher eccentricity, left and right use different gray value scales. (b) Cut along plane $x = \text{const}$. (c) Cut along plane $z = \text{const}$. (a) and (b) use the same gray value assignment as in Figure 1.

proach for using *ECC* histograms to match 2D Shapes can be found in [11].

ECC histogram descriptor. The basic building block of our shape descriptor is the histogram \mathbf{h} of the eccentricity transform *ECC* of the shape \mathcal{S} . We use k bins for the histogram. Then, the histogram descriptor is the vector $\mathbf{h} \in \mathbb{R}^k$ defined by: $\forall i = 1, \dots, k$

$$\mathbf{h}(i, \mathcal{S}) = \frac{1}{|\mathcal{S}|} \# \left\{ \mathbf{p} \in \mathcal{S} \mid \frac{i-1}{k} \leq \frac{ECC(\mathcal{S}, \mathbf{p}) - m}{M - m} < \frac{i}{k} \right\},$$

where $|\mathcal{S}|$ is the number of pixels in \mathcal{S} , and m and M are the smallest, respectively largest, eccentricity values. The obtained histogram contains only bins for the values actually existing in the eccentricity transform i.e. from minimum to maximum eccentricity, and the sum of all bins is 1. Fig. 4 shows examples of histograms for shapes with different geometric features. We note that the histogram \mathbf{h} is invariant under euclidean transformations, scaling and isometric bending of \mathcal{S} . In numerical applications, we use $k = 200$ and each descriptor is normalized such that the sum of all bins is one.




Figure 4. Top: some shapes \mathcal{S} used during experiments. Bottom: their eccentricity histograms \mathbf{h} (*ECCobj*).

Comparison of histograms. In order to match the descriptors of the two shapes \mathcal{S} and $\tilde{\mathcal{S}}$, we need to compute a meaningful distance between histograms. Let $\mathbf{h}, \tilde{\mathbf{h}} \in \mathbb{R}^k$ be the two histograms of \mathcal{S} and $\tilde{\mathcal{S}}$ computed as above. We propose to use the simple ℓ_2 -norm defined by

$$\delta(\mathbf{h}, \tilde{\mathbf{h}}) \stackrel{\text{def.}}{=} \sqrt{\sum_{i=1}^k (\mathbf{h}(i) - \tilde{\mathbf{h}}(i))^2}.$$

One could use more elaborate metrics such as the χ^2 metric or those defined in [18], but we found in numerical experiments that all these metrics give results similar to δ , which is the easiest to compute.

We can compute the distance $\Delta(\mathcal{S}, \tilde{\mathcal{S}})$ between two shapes \mathcal{S} and $\tilde{\mathcal{S}}$ as the distance of their histogram descriptors

$$\Delta(\mathcal{S}, \tilde{\mathcal{S}}) \stackrel{\text{def.}}{=} \delta(\mathbf{h}, \tilde{\mathbf{h}}). \quad (3)$$

4. Experiments

One of the most widely used 3D object retrieval databases is the Princeton Shape Benchmark [22]. It contains 1,814 3D object models organized by class and is effective for comparing the performance of a variety of methods. However, the majority of the models corresponds to rigid, man-made objects. Only a limited number of objects in the database have articulated parts. As one of the main advantages of using eccentricity is its robustness with respect to articulation, we have turned to the McGill Shape Benchmark [33]. It contains several models from the Princeton repository and others added by the authors. The main advantage of this benchmark is that from its 455 3D objects, 255 have significant part articulation. These articulated objects are grouped into 10 classes (see Figure 5). This paper shows the results on the 10 enumerated classes. See [24] for results of other methods for the same benchmark.

Two descriptors were used (see Figure 3):

1. **ECCobj** - eccentricity of the whole object;
2. **ECCborder** - eccentricity of the border voxels.

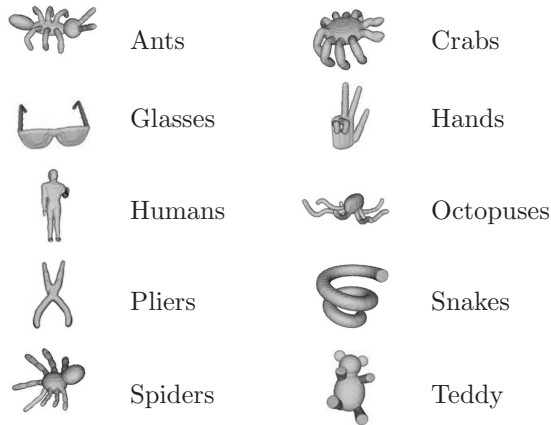


Figure 5. The object classes from the McGill 3D shape database having significant part articulation.




Figure 6. Percentage recall for several rank thresholds. Top: ECCobj, Bottom: ECCborder.

ECCborder is an attempt to get a feeling for how the eccentricity transform of the boundary relates to the eccentricity transform of the object. In the ideal case, the former should be computed on the 2D manifold defined by the boundary of the 3D objects. Computing in the “volume” of the 6 connected voxel boundary is an approximation that is possible with the available

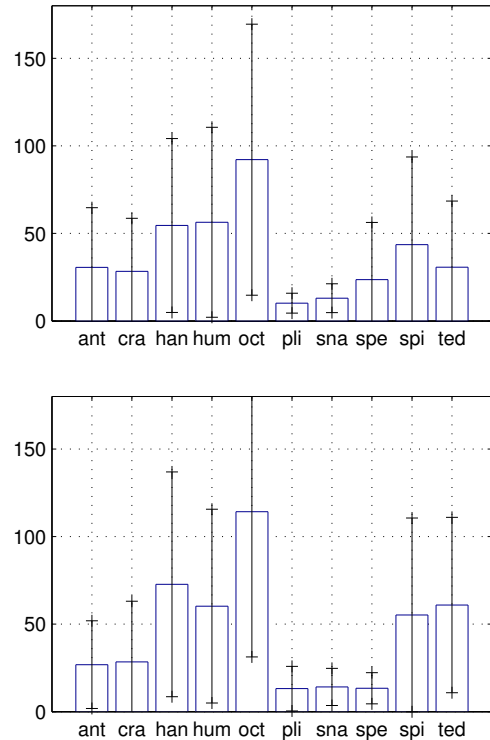


Figure 7. Average ranks for each class. Top: ECCobj, Bottom: ECCborder. (The first three letters of each class name are printed.)

tools.

Figure 6 shows the percentage recall for several rank thresholds ($t = 10, 20, \dots$). For each rank t , the percentage of models in the database in the same category as the query (not including the query itself) with indexing rank $\leq t$ is given. The average results and standard deviation, over all classes, are given.

Figure 7 shows the average ranks for each class. For all queries in a class, the ranks of all other objects in that class are computed. The average and standard deviation for each class are given.

Tables 1 and 2 show the average score (see Equation 3) for all pairs of classes. Each object in the database is matched against all other objects and each cell shows the average score between all combinations of objects of the two classes defined by the row and column. See Section 4.1 for a discussion on the results.

Figure 8 shows the *precision-recall* curves for each of the 10 classes. Precision and recall are common in information retrieval for evaluating retrieval performance. They are usually used where static document sets can be assumed. However, they are also used in dynamic environments such as web page retrieval [7]. Given a query object Q and retrieval result R contain-

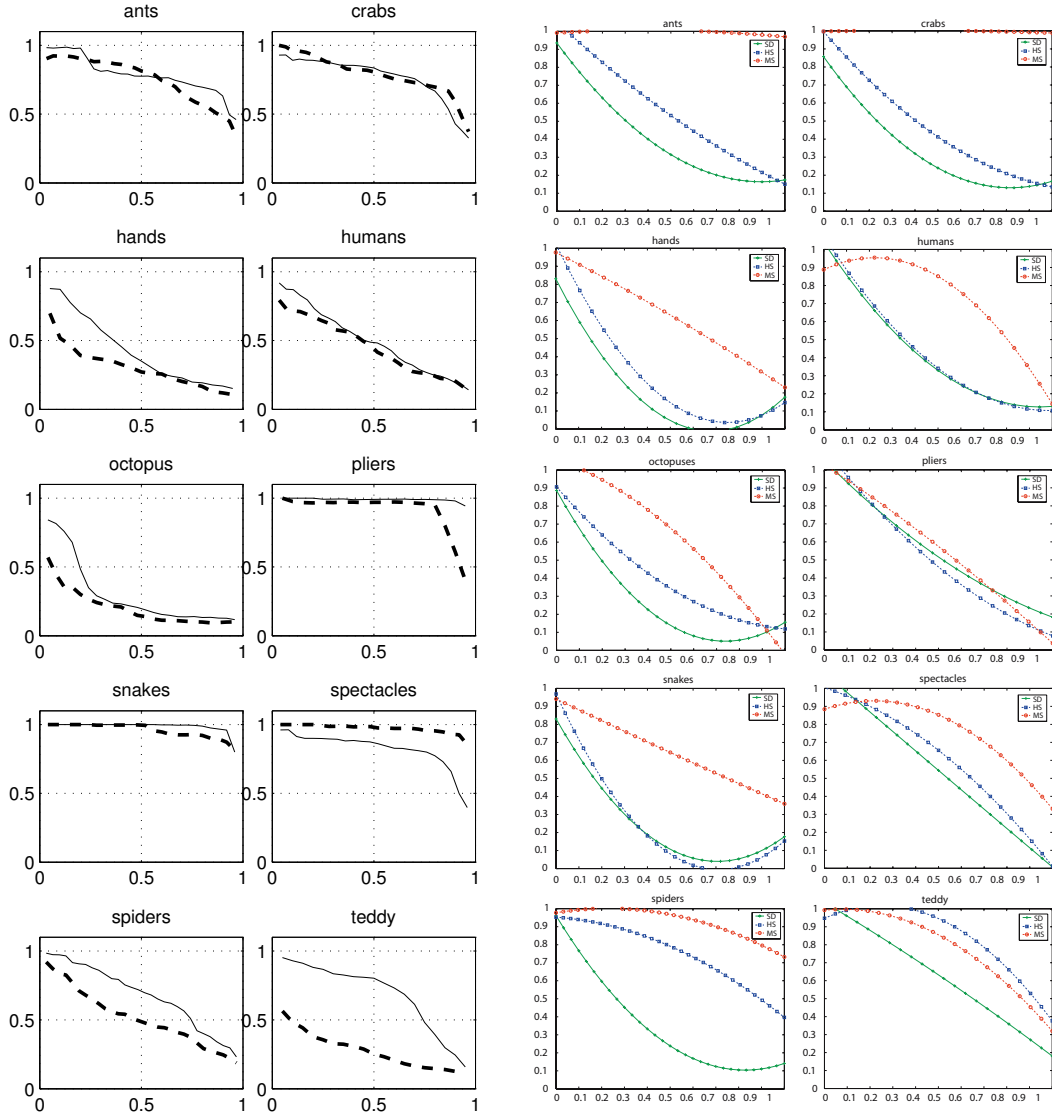


Figure 8. Left two columns: Precision-recall for the ten classes. Thin continuous line: ECCobj, Thick interrupted line: ECCborder. Right two columns (image taken from [24], with kind permission of Springer Science and Business Media): results of three other methods on the same database: medial surfaces (MS) [24], harmonic spheres (HS) [14], and shape distributions (SD) [17]. Precision: horizontal axis, recall: vertical axis.

ing t objects, precision refers to the ratio of the relevant objects retrieved to the total number retrieved, i.e. the number of objects in R of the same class as Q divided by t . Recall refers to the ratio of the relevant objects retrieved to the total number of relevant objects in the database, i.e. the number of objects in R of the same class as Q divided by the total number of objects the same class as Q , that could be returned. In our case, the latter one is equal to the number of objects of the class of Q in the database minus 1 (the query itself is not a possible answer). Precision-recall curves are produced by varying the parameter t . Better results are

characterized by curves closer to the top i.e. recall = 1 for all values of precision.

As can be seen in Figures 6, 7, and 8, and Tables 1 and 2, ECCobj does a better job than ECCborder. The percentage recall of the two methods is comparable, with slightly better results from ECCobj. With respect to the average ranks, ECCobj does clearly better with the hands, humans, octopus, pliers, spectacles, spiders, teddy, and is close or slightly worse with the ants, crabs, and snakes. Both, ECCobj and ECCborder, have the smallest average class distance (highest similarity) correct for 8 out of 10 classes, with ECCobj

	ants	crabs	hands	humans	octopus	pliers	snakes	spectacles	spiders	teddy
ants	1.75	5.33	3.72	3.53	7.20	2.95	2.91	7.25	5.43	3.72
crabs	5.33	1.55	3.73	3.50	3.67	3.02	4.36	3.99	3.43	2.53
hands	3.72	3.73	2.30	3.04	5.60	2.71	3.76	6.19	4.53	2.51
humans	3.53	3.50	3.04	2.19	5.03	2.13	3.15	5.04	3.52	2.62
octopus	7.20	3.67	5.60	5.03	3.90	5.02	6.22	4.04	4.06	4.58
pliers	2.95	3.02	2.71	2.13	5.02	0.55	1.82	4.97	3.64	2.11
snakes	2.91	4.36	3.76	3.15	6.22	1.82	0.80	5.73	4.83	3.55
spectacles	7.25	3.99	6.19	5.04	4.04	4.97	5.73	2.24	3.97	5.13
spiders	5.43	3.43	4.53	3.52	4.06	3.64	4.83	3.97	2.25	3.50
teddy	3.72	2.53	2.51	2.62	4.58	2.11	3.55	5.13	3.50	1.46

Table 1. Average matching results using ECCobj multiplied by 100 (smaller means more similar). For each row, the first and second smallest value are printed in bold.

	ants	crabs	hands	humans	octopus	pliers	snakes	spectacles	spiders	teddy
ants	1.00	2.45	2.16	1.60	3.09	1.61	2.42	5.62	2.47	1.66
crabs	2.45	1.41	2.30	3.01	3.42	2.88	3.64	6.92	3.08	2.38
hands	2.16	2.30	1.94	2.65	2.98	2.48	3.49	6.05	2.78	2.29
humans	1.60	3.01	2.65	1.57	3.19	1.54	2.16	5.12	2.54	1.93
octopus	3.09	3.42	2.98	3.19	2.97	2.72	3.82	4.80	2.69	2.80
pliers	1.61	2.88	2.48	1.54	2.72	0.65	1.75	4.51	2.00	1.47
snakes	2.42	3.64	3.49	2.16	3.82	1.75	0.85	4.86	3.21	2.44
spectacles	5.62	6.92	6.05	5.12	4.80	4.51	4.86	1.67	4.69	5.26
spiders	2.47	3.08	2.78	2.54	2.69	2.00	3.21	4.69	1.76	2.07
teddy	1.66	2.38	2.29	1.93	2.80	1.47	2.44	5.26	2.07	1.45

Table 2. Average matching results using ECCborder multiplied by 100 (smaller means more similar). For each row, the first and second smallest value are printed in bold.

having the correct class as the second smallest one for the humans and octopus (see Tables 1 and 2).

Figure 8 shows comparative precision-recall results of ECCobj and ECCborder and three other methods. Again ECCobj and ECCborder are comparable, except for the teddy bears, where ECCobj is clearly superior to ECCborder. With respect to the other methods, both eccentricity based methods score better on the pliers, spectacles and snakes, while with the octopus, two of the three other methods have considerably better results.

4.1. Discussion

The histogram of the eccentricity transform characterizes the compactness of the shape (e.g. a flat histogram characterizes a very elongated object, a histogram with monotonically decreasing values characterizes a rather compact object - see Figure 4). This measure is used to compute similarities between shapes, robust with respect to scaling, rotation, and part articulation. The matching results are promising, especially when considering the straightforward approach.




Figure 9. Examples of similar eccentricity histograms corresponding to objects of different classes. From left: octopus and spectacles, pliers and human.

The major current limitations of our approach include: (1) eccentricity histograms do not capture the topology of the shape, and thus histograms of different shapes can be very similar. (2) histogram “matching” (whether using the ℓ_2 -norm or more sophisticated methods) is inherently low level and does not consider the higher level context in which it is applied. We discuss each of these issues below.

1. Connectivity of the isoheight lines/surfaces of the eccentricity transform does capture the part structure of a shape [9], but the histograms “throw away” this information. Figure 9 shows examples of similar histograms belonging to objects of different classes.

2. In Figure 10, histograms of tree different shapes are presented. Histogram (a) is considered by Equa-




Figure 10. Example where histogram matching fails, as it is unable to capture the context of the match. Shown are ECC histograms of: (a) snake, (b), (c) human.

tion 3 more similar to (b) and to (c), than (b) to (c). Keeping in mind that these are distances, one can see that (b) should represent an object more similar to (c) than to (a), as both (b) and (c) have more “medium” and “short” distance points (more compact), while (a) is a less compact object. Please note that we have also considered other histogram matching methods (χ^2 statistic and diffusion distance) and none of them produced overall significantly different results. A possible way to improve this would be to learn a distance metric [31] for matching eccentricity histograms - in this way the matching speed would not suffer too much.

In the case of 2D shapes the eccentricity of the border is a constant. In 3D it manages to capture some of the properties of the shape, but it looks more unstable. The eccentricity transform of a simply connected volume has a single, stable center (minimum), while the eccentricity transform of its border will have a disconnected center or at least one with a more complex structure. The fact that ECCobj produced better results than ECCborder can be related to the compactness of the shapes (e.g. teddy). In our approximation of the border, for very thin parts, most of the voxels of the part are also border voxels (e.g. snakes).

Compared to other approaches (e.g. [24]), one can identify the aspects discussed above (see Figure 8 and Tables 1 and 2). For classes with simple topology (e.g. snakes and spectacles), the results are very good. For classes where part decomposition and structure play an important role (e.g. octopus v.s. spiders and crabs), the discrimination capabilities are reduced.

5. Conclusion

This paper presents matching of articulated 3D shapes using the eccentricity transform. The descriptors are defined as histograms of the eccentricity transform and matched using the ℓ_2 -norm. A big advantage of the presented approach is that it is simple and efficient. Experimental results on a database of articulated 3D shapes are given and compared to results of three other methods. We plan to study decomposition of articulated shapes using the eccentricity transform and extend the current approach to part based shape matching.

References

- [1] M. Ankerst, G. Kastenmiller, H.-P. Kriegel, and T. Seidl. 3d shape histograms for similarity search and classification in spatial databases. In *6th Int. Symp. on Advances in Spatial Databases*, pages 201–226, London, UK, 1999. Springer.
- [2] T. Ansary, J.-P. Vandeborre, S. Mahmoudi, and M. Daoudi. A bayesian framework for 3d models retrieval based on characteristic views. *2nd Int. Symp. on 3D Data Processing, Visualization and Transmission, 2004. 3DPVT 2004*, pages 139–146, 6-9 Sept. 2004.
- [3] B. Bustos, D. A. Keim, D. Saupe, T. Schreck, and D. V. Vranić. Feature-based similarity search in 3d object databases. *ACM Comput. Surv.*, 37(4):345–387, 2005.
- [4] D.-Y. Chen, M. Ouhyoung, X.-P. Tian, Y.-T. Shen, and M. Ouhyoung. On visual similarity based 3d model retrieval. In *Eurographics*, pages 223–232, Granada, Spain, 2003.
- [5] C. M. Cyr and B. B. Kimia. A similarity-based aspect-graph approach to 3d object recognition. *International Journal of Computer Vision*, 57(1):5–22, 2004.
- [6] M. Elad, A. Tal, and S. Ar. Content based retrieval of vrml objects: an iterative and interactive approach. In *Eurographics Workshop on Multimedia*, pages 107–118, New York, 2001. Springer.
- [7] T. Fawcett. An introduction to roc analysis. *Pattern Recognition Letters*, 27(8):861–874, 2006.
- [8] M. Hilaga, Y. Shinagawa, T. Kohmura, and T. L. Kunii. Topology matching for fully automatic similarity estimation of 3d shapes. In *SIGGRAPH '01: Proc. of the 28th conf. on Computer graphics and interactive techniques*, pages 203–212, New York, USA, 2001. ACM.
- [9] A. Ion and W. G. Kropatsch. Mapping a coordinate system to a non-rigid shape. In *annual workshop of the Austrian Association for Pattern Recognition (OAGM/AAPR)*. OCG, 2008.
- [10] A. Ion, W. G. Kropatsch, and E. Andres. Euclidean eccentricity transform by discrete arc paving. In *14th Int. Conf. on Discrete Geom. for Comp. Imagery (DGCI)*, LNCS, France, 2008. Springer.
- [11] A. Ion, G. Peyré, Y. Haxhimusa, S. Peltier, W. G. Kropatsch, and L. Cohen. Shape matching using the geodesic eccentricity transform - a study. In *annual workshop of the Austrian Association for Pattern Recognition (OAGM/AAPR)*. OCG, 2007.

- [12] C. Y. Ip, L. Regli, W. C. Sieger, and A. Shokoufandeh. Automated learning of model classifications. In *8th ACM Symposium on Solid Modeling and Applications*, pages 322–327, New York, 2003. ACM Press.
- [13] A. 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, May 1999.
- [14] M. Kazhdan, T. Funkhouser, and S. Rusinkiewicz. Rotation invariant spherical harmonic representation of 3d shape descriptors. In *SGP '03: Proc. of the 2003 Eurographics/ACM SIGGRAPH symposium on Geometry processing*, pages 156–164. Eurographics Association, 2003.
- [15] W. G. Kropatsch, A. Ion, Y. Haxhimusa, and T. Flanitzer. The eccentricity transform (of a digital shape). In *13th Int. Conf. on Discrete Geom. for Comp. Imagery (DGCI)*, pages 437–448, Hungary, October 2006. Springer.
- [16] H. Ling and D. W. Jacobs. Shape classification using the inner-distance. *IEEE Trans. Pattern Anal. Mach. Intell.*, 29(2):286–299, 2007.
- [17] R. Osada, T. Funkhouser, B. Chazelle, and D. Dobkin. Shape distributions. *ACM Trans. Graph.*, 21(4):807–832, 2002.
- [18] R. Osada, T. Funkhouser, B. Chazelle, and D. Dobkin. Shape distributions. *ACM Trans. Graph.*, 21(4):807–832, 2002.
- [19] E. Paquet, A. Murching, T. Naveen, A. Tabatabai, and M. Rioux. Description of shape information for 2-d and 3-d objects. *Signal Processing: Image Communication*, 16(1–2):103–122, September 2000.
- [20] E. Paquet and M. Rioux. Nefertiti: A tool for 3-d shape databases management. *SAE transactions*, 108:387–393, 1999.
- [21] J. Sethian. *Level Sets Methods and Fast Marching Methods*. Cambridge Univ. Press, 2nd edition, 1999.
- [22] P. Shilane, P. Min, M. M. Kazhdan, and T. A. Funkhouser. The princeton shape benchmark. In *Int. Conf. on Shape Modeling and Applications (SMI), Genova, Italy*, pages 167–178. IEEE Computer Society, 2004.
- [23] Y. Shinagawa, T. Kunii, and Y. Kergosien. Surface coding based on morse theory. *Computer Graphics and Applications, IEEE*, 11(5):66–78, September 1991.
- [24] K. Siddiqi, J. Zhang, D. Macrini, A. Shokoufandeh, S. Bouix, R. Chen, and S. Dickinson. Retrieving articulated 3-d models using medial surfaces. *Machine Vision and Applications*, 2007.
- [25] H. Sundar, D. Silver, N. Gagvani, and S. Dickinson. Skeleton based shape matching and retrieval. *Shape Modeling International, 2003*, pages 130–139, 12-15 May 2003.
- [26] J. W. H. Tangelder and R. C. Veltkamp. Polyhedral model retrieval using weighted point sets. In *Shape Modeling International*, pages 119–129, Seoul, Korea, 2003. IEEE.
- [27] R. C. Veltkamp and L. Latecki. Properties and performance of shape similarity measures. In *Proc. of the 10th IFCS Conf. Data Science and Classification*, Slovenia, July 2006.
- [28] D. Vranic and D. Saupe. 3d model retrieval with spherical harmonics and moments. In *23rd DAGM-Symposium on Pattern Recognition*, pages 392–397, London, UK, 2001. Springer.
- [29] D. Vranic and D. Saupe. 3d shape descriptor based on 3d fourier transform. In *EURASIP Conference on Digital Signal Processing for Multimedia Communications and Services*, pages 271–274. Comeinius University, 2001.
- [30] D. Vranic and D. Saupe. Description of 3d-shape using a complex function on the sphere. In *International conference on Multimedia and Expo*, pages 177–180. IEEE, 2002.
- [31] L. Yang, R. Jin, R. Sukthankar, and Y. Liu. An efficient algorithm for local distance metric learning. In *The Twenty-First National Conference on Artificial Intelligence and the Eighteenth Innovative Applications of Artificial Intelligence Conference, July 16-20, 2006, Boston, Massachusetts, USA*. AAAI Press, 2006.
- [32] T. Zaharia and F. Preux. Three-dimensional shape-based retrieval within the mpeg-7 framework. In *SPIE Conference on Nonlinear Image Processing and Pattern Analysis XII*, pages 133–145, 2001.
- [33] J. Zhang, K. Siddiqi, D. Macrini, A. Shokoufandeh, and S. J. Dickinson. Retrieving articulated 3-d models using medial surfaces and their graph spectra. In *5th Int. Workshop Energy Minimization Methods in Computer Vision and Pattern Recognition, EMMCVPR 2005*, volume 3757 of *LNCS*, pages 285–300. Springer, 2005.