Eurographics Workshop on 3D Object Retrieval (2010)
I. Pratikakis, M. Spagnuolo, T. Theoharis, and R. Veltkamp (Editors)

Robust Volumetric Shape Descriptor

Raif M. Rustamov

Drew University, Madison, NJ

Abstract

This paper introduces a volume-based shape descriptor that is robust with respect to changes in pose and topology. We use modified shape distributions of [OFCD02] in conjunction with the interior distances and barycentroid potential that are based on barycentric coordinates [RLF09]. In our approach, shape distributions are aggregated throughout the entire volume contained within the shape thus capturing information conveyed by the volumes of shapes. Since interior distances and barycentroid potential are practically insensitive to various poses/deformations and to non-pervasive topological changes (addition of small handles), our shape descriptor inherits such insensitivity as well. In addition, if any other modes of information (e.g. electrostatic potential within the protein volume) are available, they can be easily incorporated into the descriptor as additional dimensions in the histograms. Our descriptor has a connection to an existing surface based shape descriptor, the Global Point Signatures (GPS) [Rus07]. We use this connection to fairly examine the value of volumetric information for shape retrieval. We find that while, theoretically, strict isometry invariance requires concentrating on the intrinsic surface properties alone, yet, practically, pose insensitive shape retrieval still can be achieved/enhanced using volumetric information.

1. Introduction

Retrieval of deformable shapes has an array of practical applications, and different requirements can be imposed depending on the specific area. In computer vision, perhaps, one may limit attention to boundary surfaces and concentrate on shape descriptors extracted from the surface alone. In areas such as medical imaging, automated medical diagnosis, and protein bioinformatics important information can be conveyed by the volumes of shapes. Desirable in such cases is a pose/topology insensitive shape descriptor that is based on aggregating information from the entire volume of the shape.

A recent example of a shape descriptor that incorporates volumetric information is the recently introduced volumetric shape DNA [RWSN09]. However, this descriptor is not pose insensitive. Whereas local-diameter and centricity functions [GSCO07] gather information about the volume and are pose insensitive, these are only defined and aggregated over the surface, not the entire volume. Finally, inner distance based methods [LJ07, LFR09] can in principle be used for aggregation over the volume and are pose insensitive, yet they are sensitive to topological changes such as addition of small handles.

The main contribution of this paper is a volume-based shape descriptor that is robust with respect to changes in pose and topology. Our approach uses modified shape distributions of [OFCD02] in conjunction with the interior distances and barycentroid potential that are based on barycentric coordinates [RLF09]. To extract the descriptor, we first generate a uniform sampling of the shape's interior. Next, the sample is divided into classes depending on the value of the barycentroid potential. Finally, for each choice of two classes (that can be the same), histograms are generated to capture the distribution of pair-wise interior distances between points in these classes.

This method was chosen for three reasons. First, interior distances and barycentroid potential are practically insensitive to various poses/deformations [RLF09], and to non-pervasive topological changes (addition of small handles). Our shape descriptor inherits such insensitivity as well. Second, histograms are aggregated through the entire volume contained within the shape thus capturing information conveyed by the volumes of shapes. In addition, if any other modes of information (e.g. electrostatic potential within the protein volume) are available, they can be easily incorporated into the descriptor as additional dimensions in the his-

DOI: 10.2312/3DOR/3DOR10/001-005

tograms. Third, our descriptor has a connection to an existing surface based shape descriptor, the Global Point Signatures (GPS) [Rus07], which is pose and topology insensitive. This connection allows a fair examination of the value of volumetric information for shape retrieval.

This paper is organized as follows. We first review the concepts of interior distance and barycentroid potential in Section 3. The descriptor is described in Section 4. Its computation and connection to GPS are described in Section 5. We present the results of shape retrieval experiments using the descriptor in Section 6.

2. Related Work

Although a variety of shape descriptors, see excellent surveys [BKS*05, IJL*05, TV08], take the volume into the consideration or can be made so, these approaches are usually not pose insensitive. On the other hand, pose insensitive approaches such as canonical forms [EK03], shape DNA [RWP05], spectral embedding [JZ07], and GPS embedding [Rus07] are oblivious to the volumetric properties of the object.

The recently introduced volumetric shape DNA of [RWSN09] is based solely on volumetric information. This shape descriptor considers the 3D Laplace operator within the volume of the shape and computes its spectrum after imposing appropriate boundary conditions. This idea is similar to the surface based shape DNA [RWP05], with the difference that the Laplace-Beltrami operator on surface is replaced by the 3D Laplacian. While the former operator is pose (isometry) invariant, the latter is not, and so the resulting volumetric descriptor is not pose insensitive. This also shows that the similar modification of GPS embedding [Rus07] using the volumetric Laplacian will not be pose invariant either.

In [GSCO07], pose insensitive local-diameter and centricity functions are introduced, and their histograms are used as shape descriptors. Whereas local-diameter and centricity functions do gather information about the volume, they are only defined and aggregated over the surface, not the entire volume. This precludes their use with possible other modes of information in the volumes of shapes.

Inner distance is the length of the shortest path between two points through the interior of the object. Approaches based on this distance [LJ07, LFR09] can in principle be modified to aggregate over the volume. While inner distance is pose insensitive, it is sensitive to topological changes such as addition of small handles.

3. Interior Distance

The main ingredients of our approach come from the recently introduced framework for measuring distances in the interior of a surface mesh using barycentric coordinates [RLF09]. Thus, we quickly review the ideas of the interior distance and barycentroid potential.

Given a distance on a boundary mesh, the *interior distance* between points inside the mesh volume is obtained by propagating the boundary distance into the interior using the following three-step procedure. First, the boundary mesh is embedded into a high-dimensional space in a way that best preserves the prescribed distances on the boundary. This converts the given distances to Euclidean distances, similarly to multidimensional scaling. For distances such as diffusion and commute time distances, the embedding is known beforehand because these distances are defined in terms of Euclidean embeddings. Second, the embedding of the boundary is extended to the interior of the mesh using barycentric coordinates. This step is similar to cage-based editing, or image warping via barycentric coordinates with only difference that the boundary and interior are mapped to a high dimensional space. Third, the interior distance between any two points is defined simply as the Euclidean distance between their images under the described embedding.

The resulting distance can be computed by the formula

$$\hat{d}^2(p,q) = (\vec{w}(p) - \vec{w}(q))^T A (\vec{w}(p) - \vec{w}(q)),$$

where $\vec{w}(p) = (w_1(p), w_2(p), \ldots, w_n(p))^T$ is the column-vector of barycentric coordinates of point p with respect to the vertices of the boundary mesh, and A is the Gram matrix of the high-dimensional embedding, namely it contains pairwise dot products of the embedding coordinates. Another concept from [RLF09] that we will need is the *barycentroid potential* defined as

$$U^2(p) = \vec{w}(p)^T A \vec{w}(p),$$

which they used to define a center point, called barycentroid, of an object.

The main properties of the interior distance and the barycentroid potential that are relevant to volumetric shape description are as follows. First, both of them are defined in the entire volume contained within the mesh which makes possible their aggregation over the volume. Second, they are insensitive to deformation/pose, tessellation, and slight topological changes, and so the resulting shape descriptor can inherit these properties. Third, both can be computed fast after pre-processing, which makes the computation of descriptors fast as well.

4. Shape Descriptor

Our shape descriptor is a modification of d2 shape distributions [OFCD02], and is essentially the joint probability distribution of pair-wise interior distance and barycentroid potentials of the endpoints. The use of a joint probability distribution helps to recover some of the spatial information which would have been lost when using distributions of these functions separately.

In practice we approximate the probability distribution by a histogram. Let B and C be the number of bins for pair-wise interior distance and barycentroid potential respectively. We compute the descriptor as follows. First, we voxelize the interior of the object, and select one representative point per voxel to generate a uniform sampling of the object's interior. Second, we compute the barycentroid potential for each sample point, and divide the sample into C classes – corresponding to the C equally spaced bins for the potential. Third, for each pair of the obtained classes we generate the histogram (with B bins) of pair-wise interior distances between points one of which belongs to the first class, the other to the second class. Finally, the resulting $C \times C$ histograms are smoothed using a linear kernel method. Of course, because of the symmetry one only needs to keep somewhat more than half of the obtained histograms. During shape retrieval, we compare these histograms using chi-square distance.

In addition to the robustness properties inherited from the interior distance and barycentroid potential, this shape descriptor has the advantage of being easily extensible to cases when other modes of information are available within the volume of the mesh. For example, in bioinformatics applications one may be interested in comparing bio-molecules based on the spatial distribution of electrostatic potential or molecular properties. This information can be readily incorporated into our shape descriptor by adding extra dimensions into the histograms.

5. Computation and Connection to GPS

We discuss the evaluation of our shape descriptor, specifically of the interior distance and barycentroid potential. In addition, we uncover a relation with an existing shape descriptor: it turns out that if instead of sampling the interior, we were to only sample the surface we can obtain the GPS embedding based descriptor of [Rus07].

Remember that in order to instantiate the interior distance framework one needs to supply a distance on the boundary mesh, and barycentric coordinates; but the choice of the latter has no bearing on the rest of our discussion. We consider the case where the boundary distance is the commute-time distance. We now describe explicitly the steps involved in computing the interior distance. Consider a discrete Laplace-Beltrami operator of the mesh, and compute its eigendecomposition $\{\lambda_k, \phi_k\}_{k=0}^n$. First, the sought embedding of the boundary mesh into a high-dimensional space is given by

$$p^* = \left(\frac{\phi_1(p)}{\sqrt{\lambda_1}}, \frac{\phi_2(p)}{\sqrt{\lambda_2}}, \frac{\phi_3(p)}{\sqrt{\lambda_3}}, \cdots \right) \in R^n.$$

In fact, the commute time distance between to surface points p and q is defined as the Euclidean distance between their images p^* and q^*. The map $p \to p^*$ was called as the GPS embedding in [Rus07]. Second, the extension of this embedding to the interior of the mesh is given by the similar formula

$$p^* = \left(\frac{\hat{\phi}_1(p)}{\sqrt{\lambda_1}}, \frac{\hat{\phi}_2(p)}{\sqrt{\lambda_2}}, \frac{\hat{\phi}_3(p)}{\sqrt{\lambda_3}}, \cdots \right) \in R^n,$$

where this time p is any point within the volume contained in the surface. Here, $\hat{\phi}_k$ is the volumetric function obtained by extending the surface function ϕ_k via barycentric interpolation. Not that this new map $p \to p^*$ is defined on the volume within the mesh, and coincides with the GPS embedding on the mesh itself. Perhaps, one can call this volumetric embedding as the *interior GPS*, or *iGPS* for short. Third, the interior distance between any two interior points p and q is computed as the Euclidean distance between their images p^* and q^* under the described volumetric embedding. In addition, the barycentroid potential at p is simply the norm (the distance from the origin) of the image point p^*, namely $U^2(p) = ||p^*||^2$.

When these facts are put together we see that our shape descriptor in this case is a volumetric extension of the GPS based shape descriptor of [Rus07]. GPS descriptor is based on surface alone, and is invariant to isometries and insensitive to topology changes. This relationship between two descriptors is important for our experiments because it allows examining fairly the benefits of the volumetric information in shape retrieval.

6. Results

In order to investigate properties of our shape descriptor empirically, we have performed several experiments. The goal of these experiments is to study how the descriptor responds to changes in pose/topology, and to compare its performance to the GPS based distributions.

Pose insensitivity: Figure 1 shows the histograms corresponding to various articulations of Armadillo-man included in the Watertight Benchmark. Although there are 20 models in the same class as the Armadilloman, only 15 of these can be considered to be approximately isometric deformations, as the five models have a missing leg/arm. There are models with missing ears, but we kept them. The close clustering of these curves gives an evidence for pose insensitivity of our shape descriptor. In addition, we conducted a shape retrieval experiment (described later in this section), during which we obtained an average dynamic recall value of 87.1% for the Armadilloman class (including models with missing leg/arm) – a high value which means that the variation of our descriptor between various deformations of Armadilloman are discernibly less than the variations between Armadillomen and other object classes.

Topology insensitivity: We show that the interior distance and, consequently, our shape descriptor is insensitive

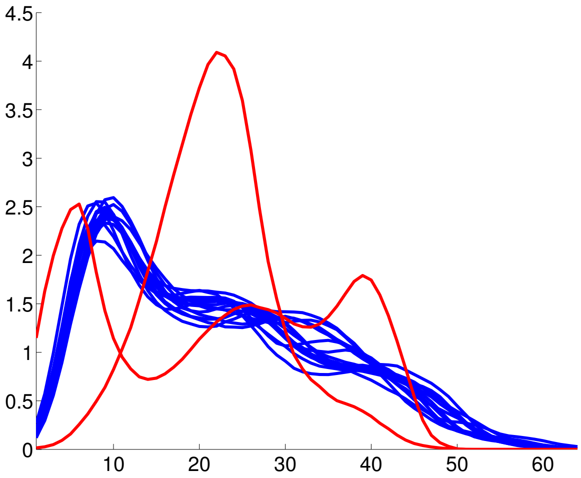

Figure 1: *Our shape descriptor ($C = 1$, $B = 64$) for the Armadillo and its articulations are shown in blue. The two curves in red are included for comparison; they represent a cup and an ant models.*

to topological changes. Our result complements the insensitivity results established in the original interior distance paper [RLF09].

The type of topological changes we consider here is when two very close but separate parts of an object are welded together by a surface patch of a relatively small area. Theoretically, the topological robustness of the interior distance and the related barycentroid potential is the result of two factors: the robustness of the Laplace-Beltrami eigenfunctions

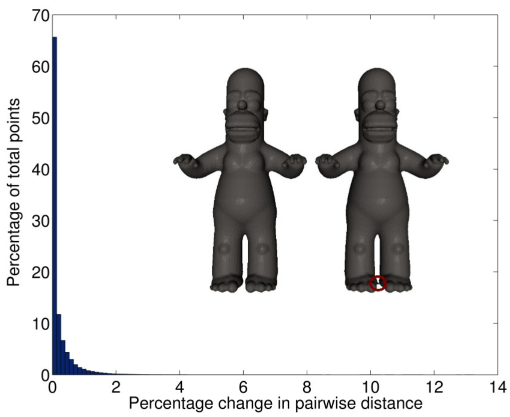

Figure 2: *The effect of topology change on the pairwise interior distances. The histogram of relative change in pairwise distances is shown. For 99.6% of pairs the change in the interior distance is less than 3%.*

Descriptor	NN	FT	ST	ADR	DCG
GPS	79.3%	45.1%	56.0%	66.1%	75.8%
This paper	81.3%	48.7%	62.0%	69.1%	77.9%

Table 1: *Shape retrieval statistics. For both descriptors the bin numbers are $C = 4$ and $B = 64$.*

(more precisely, of GPS embedding, [OSG08]) and the fact that mean-value interpolation formula contains (a discretization of) a surface integral.

We verify this theoretical prediction by an experimental comparison in the change of pairwise interior distances for the Homer model and its topological variation obtained by welding the feet together. The same uniform sample of points was sampled for both of the Homer models, and pairwise interior distances were considered. Instead of comparing the resulting histograms, we provide a more sensitive analysis by comparing corresponding pairwise distances directly. Figure 2 summarizes our findings. Notice that the relative change in pairwise distances was never greater than 14%; for 99.6% of the pairs it was less than 3%.

Comparison to GPS: For our experiments we use the Watertight Benchmark [GBP07] which consists of 400 closed surface models, divided into 20 equal object classes. Some of the classes include both different objects and articulations of the same object. We conduct a series of "leave-one-out" experiments: every model in the benchmark is queried against all other models. The ranked result lists generated by the queries are used to compute five retrieval statistics: nearest neighbor (NN), first tier (FT), second tier (ST), average dynamic recall (ADR), and discounted cumulative gain (DCG). Notice that since both our and GPS shape descriptors are translation, scale, and rotation invariant, no normalization is applied to the models. During shape retrieval, we compare descriptors using chi-square distance. We use $C = 4$ classes in order to match with the experiments in [Rus07]. The results of our experiments are summarized in Table 1, and show that our descriptor outperforms GPS in all of the statistics.

Our main conclusions from this comparative experiment are as follows. First, our shape descriptor is able to capture information that is different from GPS. Second, our descriptor displays a practically useful amount of robustness to deformations. Third, although strict isometry invariance requires concentrating on the intrinsic surface properties alone (as in the GPS), pose/deformation insensitive shape retrieval still can be achived/enhanced using volumetric information.

Implementation and performance: We used the standard cotangent weight Laplacian [PP93, MDSB02], with point areas evaluated using the function provided in Trimesh2 library [Rus08]. The generalized eigenvalue problem was set up and solved using MATLAB. The number of eigenpairs used was 250. The implementations of interior distance and barycentroid potential follow [RLF09], and was done in C++ and MATLAB. To instantiate the interior distance framework we adopt commute-time distance and mean-value coordinates. We used a $32 \times 32 \times 32$ voxel grid to obtain a uniform sampling of the volume.

The timing is reported on an Intel Core 2 Duo 2GHz laptop with 1GB RAM. For smaller models of about 10K vertices the eigenpair computation takes about 70-100 seconds. We note that for large models the method presented in [VL08] is more appropriate as it takes only about 80 seconds for a 60K vertex model, and about 9 minutes for a 250K model. The rest of the shape descriptor computation including voxelization takes less than 2 minutes on average.

As shown in [RLF09], the interior distance and barycentroid potential are fairly insensitive to tessellation. As a result, in order to save computational time, the involved meshes can be simplified without much loss of volumetric information. In our experiments, we did not simplify meshes in the benchmark in order to maintain fairness when comparing to the surface based method (GPS).

7. Conclusion

We have introduced a shape descriptor that captures volumetric information and is insensitive to pose and slight topology changes. Additionally we investigated empirically the performance of this descriptor on a common shape retrieval benchmark. Our main conclusion is that although strict isometry invariance requires concentrating on the intrinsic surface properties alone (as in the GPS), pose insensitive shape retrieval still can be achived/enhanced using volumetric information.

This work provides a small step and therefore has limitations that suggest topics for future work. For example, since our approach is based solely on histograms, we believe that the retrieval results can be improved considerably by extracting more informative shape descriptors from the volumetric extension of the GPS embedding discussed in Section 5. Another interesting question relates to the recent work of Mémoli [MÓ9] where similarity measured by the existing spectral descriptors is related to a notion of distance between shapes analogous to the Gromov-Wasserstein distance. It would be interesting to construct the corresponding framework for the volumetric invariant explored in this paper.

Acknowledgments

For the 3D models used in this paper we thank Daniela Giorgi and AIM@SHAPE. We are grateful to Szymon Rusinkiewicz for making publicly available the extremely useful Trimesh2 library.

References

[BKS*05] BUSTOS B., KEIM D. A., SAUPE D., SCHRECK T., VRANIĆ D. V.: Feature-based similarity search in 3d object databases. *ACM Comput. Surv. 37*, 4 (2005), 345–387.

[EK03] ELAD A., KIMMEL R.: On bending invariant signatures for surfaces. *IEEE Trans. Pattern Analysis and Machine Intelligence 25*, 10 (2003), 1285–1295.

[GBP07] GIORGI D., BIASOTTI S., PARABOSCHI L.: *Watertight Models Track*. Tech. Rep. IMATI-CNR-GE 9/07, September 2007.

[GSCO07] GAL R., SHAMIR A., COHEN-OR D.: Pose-oblivious shape signature. *IEEE Transactions on Visualization and Computer Graphics 13*, 2 (2007), 261–271.

[IJL*05] IYER N., JAYANTI S., LOU K., KALYANARAMAN Y., RAMANI K.: Three-dimensional shape searching: state-of-the-art review and future trends. *Computer-Aided Design 37*, 5 (2005), 509–530.

[JZ07] JAIN V., ZHANG H.: A spectral approach to shape-based retrieval of articulated 3d models. *Computer Aided Design 39* (2007), 398–407.

[LFR09] LIU Y.-S., FANG Y., RAMANI K.: Idss: deformation invariant signatures for molecular shape comparison. *BMC Bioinformatics 10*, 1 (2009), 157.

[LJ07] LING H., JACOBS D.: Shape classification using the inner-distance. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) 29*, 2 (2007), 286–299.

[MÓ9] MÉMOLI F.: Spectral gromov-wasserstein distances for shape matching. In *Workshop on Non-Rigid Shape Analysis and Deformable Image Alignment (ICCV workshop, NORDIA'09)* (october 2009).

[MDSB02] MEYER M., DESBRUN M., SCHRÖDER P., BARR A.: Discrete differential geometry operators for triangulated 2-manifolds. In *Proceedings of Visual Mathematics* (2002).

[OFCD02] OSADA R., FUNKHOUSER T., CHAZELLE B., DOBKIN D.: Shape distributions. *TOG 21*, 4 (2002), 807–832.

[OSG08] OVSJANIKOV M., SUN J., GUIBAS L.: Global intrinsic symmetries of shapes. In *Eurographics Symposium on Geometry Processing (SGP)* (2008).

[PP93] PINKALL U., POLTHIER K.: Computing discrete minimal surfaces and their conjugates. *Experiment. Math. 2*, 1 (1993), 15–36.

[RLF09] RUSTAMOV R., LIPMAN Y., FUNKHOUSER T.: Interior distance using barycentric coordinates. *Computer Graphics Forum (Symposium on Geometry Processing) 28*, 5 (July 2009).

[Rus07] RUSTAMOV R.: Laplace-beltrami eigenfunctions for deformation invariant shape representation. In *Symposium on Geometry Processing* (2007).

[Rus08] RUSINKIEWICZ S.: trimesh2, 2008.

[RWP05] REUTER M., WOLTER F.-E., PEINECKE N.: Laplace-spectra as fingerprints for shape matching. In *Solid and Physical Modeling* (2005), pp. 101–106.

[RWSN09] REUTER M., WOLTER F.-E., SHENTON M., NIETHAMMER M.: Laplace-beltrami eigenvalues and topological features of eigenfunctions for statistical shape analysis. *Computer-Aided Design 41*, 10 (2009), 739–755. doi:10.1016/j.cad.2009.02.007.

[TV08] TANGELDER J., VELTKAMP R.: A survey of content based 3d shape retrieval methods. *Multimedia Tools and Applications 39* (2008), 441–471.

[VL08] VALLET B., LÉVY B.: Manifold harmonics. *Computer Graphics Forum (Proceedings Eurographics)* (2008).