

Pose-Oblivious Shape Signature

Ran Gal, *Student Member, IEEE*, Ariel Shamir, *Member, IEEE*, and
Daniel Cohen-Or, *Senior Member, IEEE*

Abstract—A 3D shape signature is a compact representation for some essence of a shape. Shape signatures are commonly utilized as a fast indexing mechanism for shape retrieval. Effective shape signatures capture some global geometric properties which are scale, translation, and rotation invariant. In this paper, we introduce an effective shape signature which is also pose-oblivious. This means that the signature is also insensitive to transformations which change the pose of a 3D shape such as skeletal articulations. Although some topology-based matching methods can be considered pose-oblivious as well, our new signature retains the simplicity and speed of signature indexing. Moreover, contrary to topology-based methods, the new signature is also insensitive to the topology change of the shape, allowing us to match similar shapes with different genus. Our shape signature is a 2D histogram which is a combination of the distribution of two scalar functions defined on the boundary surface of the 3D shape. The first is a definition of a novel function called the *local-diameter* function. This function measures the diameter of the 3D shape in the neighborhood of each vertex. The histogram of this function is an informative measure of the shape which is insensitive to pose changes. The second is the *centricity* function that measures the average geodesic distance from one vertex to all other vertices on the mesh. We evaluate and compare a number of methods for measuring the similarity between two signatures, and demonstrate the effectiveness of our pose-oblivious shape signature within a 3D search engine application for different databases containing hundreds of models.

Index Terms—Shape-signature, shape-matching, pose-oblivious.

1 INTRODUCTION

THE signature of a shape is a concise representation of the shape that captures some of its essence. A signature does not fully represent the shape, and it is impossible to reconstruct the shape from it. However, if the signature succeeds in expressing some of the shape's properties well, it can be used as a succinct shape representative in various applications. A typical application area for using shape signatures is 3D shape similarity and matching. In these applications, signatures are extracted from 3D geometric objects and used to determine shape similarity. Instead of comparing the full 3D objects' models, only the signatures are compared. This technique accelerates the matching process in order of magnitude. More importantly, the semantic meaning of similarity is defined by the signature used for comparison. For this reason, the properties of the shape signature itself are of great importance to the success of fast and effective similarity measurements between shapes.

For example, rigid-body-transformation invariance is often desired for a 3D shape signature. Many of the recently proposed shape signatures aimed at capturing some essence of a shape while being rigid-body-transformation invariant and often also uniform-scale invariant [1], [2]. Nevertheless, frequently, 3D objects are not rigid as in cups or chairs, but flexible to change their spatial arrangement or pose. For instance, a human or an animal 3D model may

come in many different poses: standing, running, sitting, lying, etc., a pair of scissors or a box may be open or closed. These models represent the same object although their pose is different. Many shape signatures which are effective for matching rigid objects, do not handle pose differences of flexible objects well. In this work, we concentrate on the pose-invariance property which is important for shape signatures. We define a signature which is rigid-body-transformation invariant, and is expressive to identify and distinguish between different shapes similar to top-performing previously defined shape signatures. However, it also remains largely consistent through pose changes of the same shape (see Fig. 1), outperforming previous methods when the database contains objects with pose changes (see Table 2).

Let Ψ be a transformation that changes the pose of an object O , such as skeletal articulations, and $\sigma(O)$ the signature of O . Although our signature cannot guarantee that $\sigma(O) \equiv \sigma(\Psi(O))$, in general, the distance between $\sigma(O)$ and $\sigma(\Psi(O))$ is very small and certainly smaller than that of current 3D shape signatures. Therefore, we term our shape signature *pose-oblivious*. The key to our pose-oblivious signature definition is the use of pose-oblivious features of the shape. These features are in fact functions defined on the surface of the mesh, and they remain largely consistent when the pose of the object changes (Fig. 3). The first is a novel *local-diameter* function (DF), which captures the local shape of the object's volume. This new function examines the object's diameter in the neighborhood of each point on the surface.

The histogram of the DF function is an expressive signature regarding the distinction between shapes, and also carries the pose-oblivious property of the function itself (Fig. 1). Nevertheless, the spatial distribution of the function is completely lost. To alleviate this, we use a second measure, the *centricity* function (CF), that has been used

- R. Gal and D. Cohen-Or are with the School of Computer Science, Tel-Aviv University, Ramat-Aviv, Tel-Aviv 69978, Israel.
E-mail: galran2@gmail.com, dcor@post.tau.ac.il.
- A. Shamir is with the Efi Arazi School of Computer Science, The Interdisciplinary Center, PO Box 167, Herzlia 46150, Israel.
E-mail: arik@idc.ac.il.

Manuscript received 3 July 2005; revised 14 June 2005; accepted 15 Aug. 2006; published online 10 Jan. 2007.

For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org, and reference IEEECS Log Number TVCG-0081-0705.

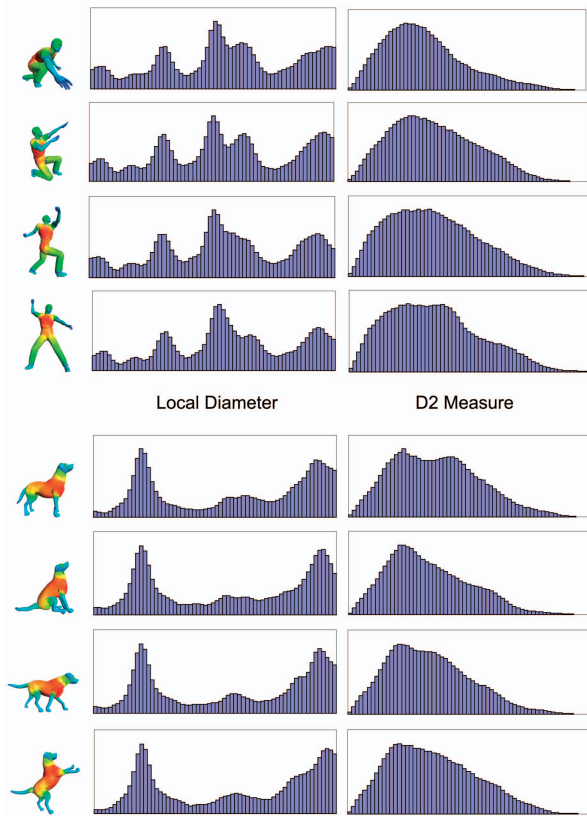


Fig. 1. The local diameter function (DF) signature compared, for instance, to D2 measure from [2]. The DF is more expressive regarding the distinction between shapes (a man or a dog), and is more oblivious to pose changes.

previously [3]. This function measures for each point on the boundary surface of the object, the average of the geodesic distances to all other points on the surface. We use the centrality function as a positional measure and create a 2D signature which is a histogram combining the two functions' distributions. This signature gives better discrimination results than each function on its own.

2 SHAPE SIGNATURE PROPERTIES

Shape signatures that characterize the global shape of a 3D model are often *rigid-body transform* and uniform *scale* invariant. These include the volume-to-surface ratio, statistical moments, and Fourier transform of the boundary or the volume of the model [1], [4], [5]. These signatures only use global properties to characterize the overall shape of the object. Hence, they are not very discriminative about object details. The concept of global feature-based similarity has been refined by comparing distributions of global features instead of the global features directly using shape distributions [2]. Global measures as well as shape distributions are easy to implement, they can be indexed efficiently and allow for very fast retrieval using the nearest neighbor algorithm and others. Nevertheless, shape distributions have their limitations as much of the shape information is lost by projecting onto distribution descriptors (see, e.g., [6]). In general, designing expressive global measures is not easy and there is still a constant effort to develop signatures that are rigid body and scale invariant.

The semantics of *pose-invariance* for articulated objects is similar to rigid-body transformation invariance in rigid objects. A human standing, walking or bending still represents a human. Although 3D models of such objects can be geometrically different, they should often be considered the same or close. Hence, providing a pose-oblivious shape signature is of major importance in applications such as search engines and shape-matching.

In the context of shape-matching, other works have shown pose-oblivious results. Specifically, the *topology* of 3D models is also an important shape characteristic and is often pose-invariant. Pure topological signatures such as the Genus, the number of connected components and, in general, the Betti numbers of the shape [7], are very crude descriptions and may sometimes even harm shape similarity measures. The term “topology” is often used to describe the overall structure of the shape. Toward this end, graph-based and skeleton-based methods attempt to extract a more succinct representation that characterizes the shape components and the way they are linked together [8], [3], [9], [10], [11], [12]. In such methods, the object signature is typically represented in the form of a relational data-structure such as a graph. Hence, the similarity estimation problem is transformed into a graph comparison problem. This facilitates articulated body matching, since topology is usually a *pose-oblivious* characteristic of the object. However, general graph matching is a very difficult problem, and there is a need to align the graphs or sometimes even subgraphs. The graph extraction process is often very sensitive to topology changes and noise. Furthermore, the cost of graph comparisons increases proportionally with graph size, resulting in relatively slow comparison and retrieval times.

In contrast, our proposed signature carries the efficiency, simplicity and robustness of shape distribution methods [2], and is not sensitive to topology changes. It is scale, rotation and translation invariant, in addition to it being pose-oblivious. This signature gives very good results using various metrics and models from the Princeton Shape Benchmark database [13], and performs even better when various poses of the same or similar objects are used.

3 RELATED WORK

The problem of similarity and matching of shapes has been extensively studied in numerous fields such as computer vision, robotics, molecular biology and others. Many have focused primarily on matching shapes in 2D images. Matching 3D models seems easier since it does not require recognition—the geometry is given, and there is no occlusion or disrupting external effects such as lighting and reflections. On the other hand, 3D models typically lack a simple parametrization domain and, thus, registration and feature correspondence are more difficult. For a broad introduction to shape-matching methods, please refer to any of several survey papers [14], [15].

In computer graphics and geometry processing fields, the matching of 3D shapes was developed mainly for shape retrieval. Recently, new methods were developed for the retrieval of 3D models in the context of a Web search engine, based on geometric properties rather than textual ones [1], [16], [17], [13], [18]. In large shape collections, it is inefficient to sequentially match all objects in the database

with the query object. For fast and efficient retrieval, efficient indexing search structures are needed. Numerous methods exist for analyzing 3D shapes and extracting different types of shape descriptors, or signatures, that can be compared to determine similarity between models. These employ geometric or topological attributes of the shape, or both.

Signatures which are primarily based on geometric properties of the shape either use a global measure or the distribution of a geometric property. Global properties of the 3D models include statistical moments of the boundary or the volume of the model [1], [19], and Fourier descriptors [16], [20], [21]. Histograms and shape distributions measure properties based on distance, angle, area, and volume measurements between random surface points [2], [22]. An additional harmonic-based representation was presented in [23], [24] which is intrinsically rotation invariant and shown to provide good matching performance. The Light Field Descriptor (LFD) represents a model as a collection of images rendered from uniformly sampled positions on a view sphere [25]. By measuring the L_1 -difference between all rotations and all pairings of positions, it can also be considered rotation-invariant. Nevertheless, none of these methods is pose-invariant and they cannot support articulated body matching.

Graph-based methods attempt to extract a structure from a 3D shape generalizing it to a graph showing how the shape components are linked together. Some of these methods use discretization based either on voxels [8], [10] or on Voronoi and Delaunay complexes [11], to extract a skeleton or partition the object to its components. These are then used to create the object's graph representation. Other methods use morse functions on the surface to characterize its topology building a multiresolution Reeb graph [3], [9]. The graph-based representations are often pose-oblivious. Nevertheless, they are complex and sometimes error prone due to discretization. They are susceptible to topology changes. They rely on graph matching which is a hard problem, and suffer from relatively slow comparison and retrieval times.

In [26], the authors present a bending invariant representation based on multidimensional scaling (MDS). This representation is an embedding of the geometric structure of a surface to a small dimensional Euclidean space, in which geodesic distances are approximated by Euclidean ones. The method aims at filtering out the "pose" of the object by bringing all objects to a canonical pose. This method gives good results on simple isometric surfaces that share the same geometric structure, but is too sensitive to modifications and hard to control on general 3D meshes (see examples in Section 6 for more details).

Other methods take into account local features on the boundary surface of the shape in the neighborhood of points. Usually, these techniques are based on matching local descriptors, such as Spin images or histograms [27], [28], [29], [30], [31]. Since they describe local surface measures, they may also be oblivious to global pose changes. However, they often do not perform well on global shape matching since their local nature does not




Fig. 2. Examples of the cone of rays shot to the opposite side of the mesh.

provide a good signature of the overall shape. Furthermore, these methods can be inefficient for global matching since they usually require large amounts of storage space.

In [13], a benchmark model database was compiled and a thorough comparison of different 3D matching methods was introduced. For instance, the ground-truth classification includes separate classes for humanoids that are standing, have their arms up, or are walking. Some applications expect all those models to be in the same category. For those applications, our shape signature can provide more effective retrieval, while maintaining the efficiency, simplicity, and robustness of shape distribution methods.

4 THE LOCAL DIAMETER SHAPE-FUNCTION

The idea of the local diameter shape-function is to create a type of low pass filtering to a diameter measure, which relates to the medial axis transform (MAT) [32]. In general, the medial axis of a 3D object is a collection of nonmanifold sheets. Computing the medial axis and the MAT of a surface mesh is an expensive process, and the medial axis itself is difficult to handle [33], [34]. Discrete approximations such as skeletons depend on voxelization and are often sensitive to noise. Therefore, we replace the local shape-radius by a measure of the local *shape-diameter*, and use it as a function on the boundary of the object.

We assume that a 3D object is defined using a boundary surface (e.g., a triangular mesh) which is almost watertight. On a smooth surface, the exact diameter can be defined by the distance to the antipodal surface point using the opposite direction of the normal. On a piecewise linear mesh, it is difficult to define the exact antipodal point. Moreover, we want to express the diameter of the object in the neighborhood of a point, which is different from the exact distance to the antipodal point.

The Local-Diameter Function: The local shape diameter at a point on the boundary of the object is defined using a robust statistics measure of the diameters in a cone around the direction opposite to the normal of the point (Fig. 2). By testing more than 1,000 meshes and checking the effect on the shape signature, we arrived at a procedure with hard-coded parameters (no manual tuning is required for various applications or data). First, we use a large opening angle of 120° for the cones. Second, we sample 50 rays for each cone. Third, we remove outliers for various reasons. The top 30 percent and bottom 10 percent of the values are discarded since some rays may reach parts which are too close or too distant (up to infinity if the mesh contains

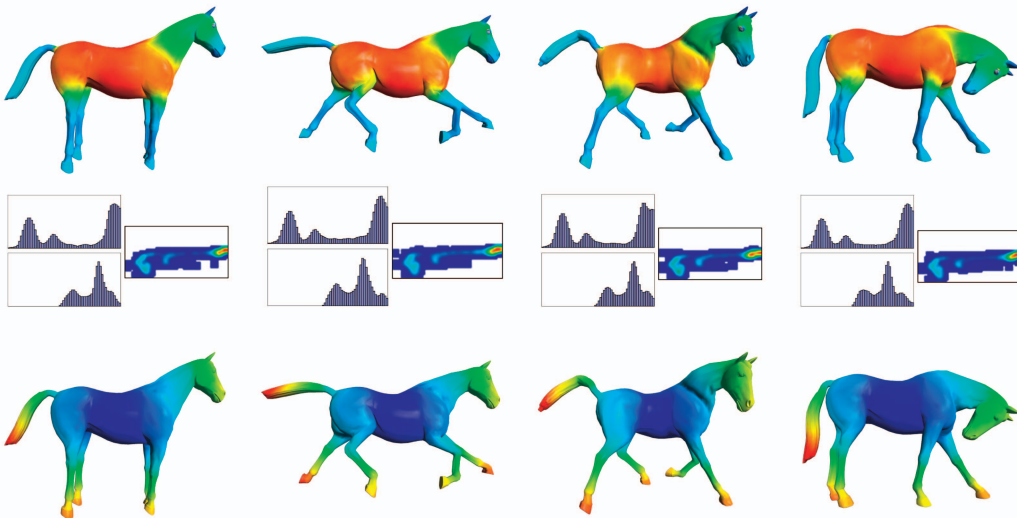


Fig. 3. Pose-oblivious functions and signatures. (a) Four horses colored by the DF (diameter) values (blue is low, red is high) are shown, and (c) shows the visualization of the CF (centricity) values. In (b), we present the signature histograms (DF on top, CF on bottom, and CDF in the middle).

holes). We also check the normal at the intersection point and ignore intersections where the normal is pointing in a similar direction to the origin-point of the ray (this may happen if there are self intersections or internal parts). The final shape diameter value is calculated as the average of the remaining rays."

This definition of the diameter shape-function (DF) is invariant to rigid body transformations. To create a function which is also scale independent we divide the function values by the maximum diameter of all measures. Furthermore, the diameter shape-function is insensitive to any deformation that does not alter the volumetric shape locally. This includes articulated character deformations, skeleton-based movements, or piecewise-rigid transformations. Still, there are positions on the mesh where the measure can change after such deformations. For instance, at the tip of the elbow of a person bending his arm, the measure can change considerably. To overcome this, we further smooth the function values on the mesh by averaging the value of each vertex with its neighbors (Fig. 3a).

5 THE POSE-OBVIOUS SIGNATURE

The local diameter shape-function expresses a good distinction between the different object parts, which is oblivious to the object's poses (see, e.g., [35]). However, a signature of a 3D shape must be a succinct representation of the shape, and there is a need to convert the function defined on the boundary of the object to a shape signature. This is done in a similar manner as shape distribution measures in [2]. The key idea is to create an approximation of the probability density function of the values on the mesh. For shape distributions, this is done by sampling the values of the measure (e.g., D1 is the distance of a point to the centroid of the object, and D2 is the distance between two random points), and then building the histogram of values using 64 bins.

Calculating the local diameter on a sample point on the mesh boundary involves ray shooting and averaging. To create a histogram of the values, we use a real approximation

of the function distribution on the mesh instead of random sampling. We calculate the function value for each of the vertices of the mesh and weight each of them according to its influence on the boundary. The influence of a vertex is defined as the area of the triangles surrounding it divided by the whole boundary area. We use a histogram of 64 entries and add the weighted values in each bin to define the signature vector.

This 64-entry DF vector is an expressive signature which is pose oblivious as can be seen in Figs. 1 and 3. Still, by converting the function to a one-dimensional histogram, a considerable amount of information is lost. One of the most valuable pieces of information which is lost relates to the spatial distribution of the values of the function. We seek to augment the volume-function values with some geometric positioning indicator. Nevertheless, the use of 3D positions or relative distances to the centroid will damage the pose-oblivious nature of the signature. Therefore, we use another relatively pose-oblivious measure of spatial positioning—the normalized centricity function (CF).

The Centricity Function: The centricity of each vertex is defined as the average geodesic distance to all other vertices. For geodesic distance calculations, we use a similar method to [3] including short-cut edges. We then divide the centricity value of each vertex by the maximum centricity value on the mesh to arrive at a CF function value between 0 and 1. The combined histogram of the CF and DF functions is a 2D array of scalar values between [0, 0] and [1, 1]. This 2D array is created by quantizing the values of the two functions (32 values for CF and 64 for DF). Hence, each bin with values (x, y) contains the approximated probability of a point on the boundary of the mesh to have a DF value of x and a CF value of y . This 2D rectangular histogram is visualized as a 2D image with entries mapped to colors (see Figs. 3 and 5).

Regarding implementation, the basic operation in the DF calculation is the ray-mesh intersections. This operation is well known in ray-tracing and can be accelerated accordingly using search structures. We used a spatial octree built around




Fig. 4. A model containing a number of connected components and in different tessellation does not affect the DF and CF signatures.

the mesh to assist in intersection finding. In general, this construction does not take more than a few seconds even on large meshes, and, consequently, computing the diameter function even on large meshes takes only a few minutes. As an example, computing the diameter function for more than 1,000 meshes with up to 20K vertices (an average of around 5,000), took around 24 hours, which is less than two minutes on average. The computation of the CF function is more expensive, but took at most 10 minutes for the large meshes (20K vertices). We computed both the DF and CF functions on all the vertices of a mesh. However, when preprocessing time is a limiting factor, one can successfully approximate both functions by computing exact values only on a subset of vertices and using averaging.

Another important issue of a shape signature is its robustness to the object's representation. The local diameter function is only meaningful on objects which define a closed volume. This means that objects which contain nonvolumetric parts, or interior parts, may need some preprocessing (e.g., plants in some of the PSB examples [13]). Other objects in the database may contain holes or missing parts. Furthermore, our approach relies on a good estimation of the normals of the processed surface. Noisy models or unconnected polygon soups can cause our signature to become less reliable. Our outlier removal method prevent averaging rays with no intersection at all (infinity-rays), and the smoothing helps to correct discrepant values. This introduces robustness to small cracks and holes in the boundary, meaning we can work with meshes that do not need to be truly watertight. Similarly, the centricity function calculation is suitable for connected models, and, hence, we use virtual linking-edges to connect different components in objects. Furthermore, by weighting the function values by the area of influence of each vertex, the definition of CF and DF signatures are not sensitive to the tessellation of the object (Fig. 4).

6 EXPERIMENTAL RESULTS

The signature of a shape is usually used as an index in a database of shapes and enables fast queries and retrieval.




Fig. 5. An example of different models in various poses with their 2D CDF signatures. It can be seen that models of the same class have very similar histograms.

Hence, to achieve accurate results there is a need to define the distance measure between two signatures. For one-dimensional vectors, such as shape distributions and also our DF and CF signatures, several options were investigated in [2]. These include Minkowski L_n norms, the χ^2 measure, and Earth Mover's distance [36]. For our 2D signature CDF, we tested L_1 and L_2 by unfolding the matrix as a 64×32 vector of values, and also correlation coefficient [29] and χ^2 measures. In fact, we found that using different metrics on different signatures may affect the query results and the success measures. In our experiments, we tested all different types of metrics for each signature when possible. Although we show all results (Tables 1, 2, and 3), we found that metrics such as χ^2 and the correlation coefficient usually give better results. Although this calls for further investigation, these metrics are more suitable for measuring distance between histograms as they give some global measure of difference as opposed to local point-to-point distance such as L_n norms.

In addition to our three signatures DF (1D histogram of the diameter function), CF (1D histogram of the centricity function), and CDF (2D combined histogram), we implemented the D1 and D2 signatures from [2] which seem to give best shape distribution results. Furthermore, we compared our signature to two out of the three top performing descriptors on the PSB as

TABLE 1
Various Quantitative Measures Evaluated on the Sub-PSB Database for Different Signatures

Sub-PSB	Nearest Neighbor	First Tier	Second Tier	E-Measure	DCG
LFD	86.72%	47.35%	62.19%	37.75%	79.61%
SH	76.82%	40.63%	56.87%	32.08%	73.98%
CDF (χ^2)	64.85%	39.63%	57.62%	32.11%	69.79%
CDF (L_1)	56.25%	30.29%	46.10%	24.73%	65.58%
CDF (L_2)	54.95%	31.39%	47.73%	25.07%	65.43%
CDF (CC)	55.99%	32.82%	49.06%	24.67%	65.65%
DF (χ^2)	61.20%	38.41%	57.06%	29.62%	69.48%
DF (L_1)	58.07%	36.03%	55.78%	28.56%	68.55%
DF (L_2)	57.81%	34.50%	53.37%	27.03%	67.69%
DF (CC)	50.78%	27.93%	45.95%	22.22%	64.16%
CF (χ^2)	33.07%	19.54%	34.93%	15.41%	55.83%
CF (L_1)	29.69%	18.31%	32.82%	14.74%	54.68%
CF (L_2)	29.43%	17.72%	32.33%	14.18%	54.39%
CF (CC)	27.86%	17.72%	31.66%	13.62%	54.03%
D2 (χ^2)	58.03%	33.98%	47.92%	26.31%	65.19%
D2 (L_1)	59.38%	32.70%	47.54%	25.99%	66.77%
D2 (L_2)	59.11%	32.25%	47.56%	25.28%	66.19%
D2 (CC)	55.47%	31.89%	48.44%	25.06%	65.89%
D1 (χ^2)	47.13%	28.31%	44.94%	22.17%	62.09%
D1 (L_1)	47.14%	26.60%	44.04%	20.71%	61.35%
D1 (L_2)	44.79%	26.08%	43.46%	19.88%	60.23%
D1 (CC)	41.41%	23.92%	42.68%	18.76%	59.24%

Our CDF signature is slightly lower in performance compared to top descriptors such as SH and LFD (see Fig. 6).

described in [13]: The Light Field descriptor (LFD) [25] (implementation taken from [37]), and the Spherical Harmonic descriptor (SH) [24] (implementation taken from [38]).

To compare the effectiveness of the proposed signatures, we executed a series of shape matching experiments with three different databases of 3D models:

- **Sub-PSB** is a database containing around 400 models from the PSB database. We used the given base classification, where we joined together classes like “humans” and “humans with arms outstretched.” We did not use some classes, such as “plants,” as

TABLE 2
Various Quantitative Measures for Different Signatures Evaluated on the ISDB Database, Which Includes Many Articulated Characters

ISDB	Nearest Neighbor	First Tier	Second Tier	E-Measure	DCG
LFD	72.64%	44.37%	62.44%	38.70%	72.83%
SH	78.30%	46.64%	63.52%	40.71%	74.78%
CDF (χ^2)	100.00%	98.34%	99.67%	61.90%	99.81%
CDF (L_1)	100.00%	96.64%	99.29%	61.05%	99.56%
CDF (L_2)	100.00%	96.18%	99.10%	61.45%	99.41%
CDF (CC)	100.00%	97.26%	99.69%	62.18%	99.63%
DF (χ^2)	100.00%	95.62%	99.49%	59.92%	99.53%
DF (L_1)	100.00%	96.03%	99.33%	60.20%	99.61%
DF (L_2)	100.00%	95.83%	99.31%	59.92%	99.44%
DF (CC)	100.00%	97.87%	99.48%	61.44%	99.71%
CF (χ^2)	97.17%	81.66%	87.30%	58.08%	91.92%
CF (L_1)	98.11%	79.60%	86.90%	58.30%	91.52%
CF (L_2)	97.17%	78.10%	83.91%	57.05%	90.29%
CF (CC)	97.17%	76.04%	83.35%	55.49%	89.72%
D2 (χ^2)	53.77%	29.93%	48.24%	30.12%	62.67%
D2 (L_1)	58.49%	31.14%	50.06%	30.22%	63.18%
D2 (L_2)	56.60%	32.11%	50.41%	30.75%	63.48%
D2 (CC)	60.38%	30.94%	48.22%	30.75%	63.20%
D1 (χ^2)	64.15%	44.94%	63.89%	39.44%	71.51%
D1 (L_1)	62.26%	42.92%	62.95%	38.17%	70.71%
D1 (L_2)	57.55%	41.31%	61.65%	37.42%	69.27%
D1 (CC)	56.60%	39.57%	60.42%	35.81%	67.56%

The CDF and DF signatures outperform all other descriptors (see Fig. 7).

TABLE 3
Various Quantitative Measures for Different Signatures Evaluated on the Combined Database CDB, Which Includes Only 20 Percent of New Models Compared to the Sub-PSB

CDB	Nearest Neighbor	First Tier	Second Tier	E-Measure	DCG
LFD	78.98%	42.15%	58.64%	33.29%	75.98%
SH	75.08%	36.93%	54.41%	29.42%	71.35%
CDF (χ^2)	70.82%	44.89%	62.56%	34.89%	74.88%
CDF (L_1)	65.91%	39.53%	55.53%	29.31%	71.96%
CDF (L_2)	64.68%	39.77%	56.31%	29.09%	71.77%
CDF (CC)	65.09%	41.31%	57.54%	28.79%	71.64%
DF (χ^2)	67.89%	42.73%	59.23%	30.62%	73.01%
DF (L_1)	66.94%	42.66%	59.19%	30.67%	73.73%
DF (L_2)	66.74%	42.08%	59.10%	29.20%	72.86%
DF (CC)	60.99%	37.56%	53.18%	25.86%	70.26%
CF (χ^2)	44.91%	26.55%	40.42%	18.92%	61.07%
CF (L_1)	42.92%	25.05%	38.71%	18.22%	60.01%
CF (L_2)	42.09%	24.52%	37.59%	17.26%	59.37%
CF (CC)	41.27%	24.25%	37.12%	16.89%	58.93%
D2 (χ^2)	57.29%	29.64%	43.67%	21.75%	63.34%
D2 (L_1)	55.24%	28.76%	43.54%	21.75%	63.12%
D2 (L_2)	54.21%	29.02%	44.05%	21.64%	62.87%
D2 (CC)	52.16%	28.77%	44.57%	21.29%	62.73%
D1 (χ^2)	44.64%	27.16%	43.16%	19.86%	60.37%
D1 (L_1)	43.94%	25.70%	41.88%	18.55%	59.66%
D1 (L_2)	42.09%	25.26%	41.25%	17.70%	58.58%
D1 (CC)	39.63%	22.57%	39.96%	16.42%	57.46%

Still, CDF is one of the best performing descriptors (see Fig. 8).

they included many nonvolumetric objects or many nonconnected parts which needed considerable preprocessing. Similarly, many models in the PSB contain internal structures, resulting in erroneous DF calculations.

- **ISDB** is a database of different articulated figures of animals and humans containing about 80 models.
- **CDB** is the union of the two previous databases. Note that in the combined database, similar classes from the two databases were merged together.

The models contained anywhere between 200 and 35,000 polygons. Not all models formed a single manifold surface or even a well-defined solid region. Some models contained cracks, self-intersections, and/or missing polygons—none of which caused significant artifacts during rendering with a z-buffer, but all of which are problematic for some 3D shape matching algorithms. The experiments were run on a PC with a 3GHz Pentium 4 processor and 1,024 MB of memory.

We evaluated several qualities of retrieval measurements. We used the same parameters as in [13]. For more details on these methods, the reader is referred to [13]:

- **Nearest neighbor:** The percentage of closest matches that belong to the same class as the query.
- **First-tier and second-tier:** The percentage of models in the query’s class that appear within the top K matches, where K depends on the size of the query’s class.
- **E-measure:** A composite measure of the precision and recall for a fixed number (32) of retrieved results.
- **Discounted Cumulative Gain (DCG):** A statistic that weights correct results near the front of the list more than correct results later in the ranked list under the assumption that a user is less likely to consider elements near the end of the list.

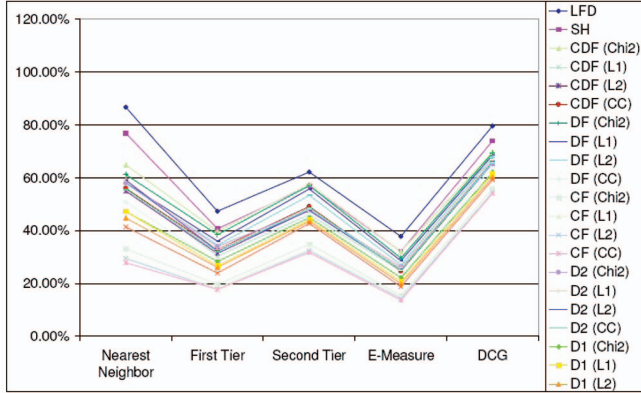


Fig. 6. On PSB models (Table 1), the CDF is still compatible with the best descriptors.




Fig. 7. For articulated characters (the ISDB from Table 2), the CDF and DF outperform all other descriptors.




Fig. 8. On the combined database (Table 3), the CDF is one of the best descriptors.

We summarize the results into three tables, one for each database. Table 1 (Fig. 6) presents the results for the sub-PSB database, Table 2 (Fig. 7) for the ISDB database, and Table 3 (Fig. 8) for the combined database. The two best measures for each quality measurement are shown in bold in each column. Examining the results, we can see that the CDF signature is compatible and only slightly worse than the best measures for PSB-type models (Table 1). However, for pose variations, CDF (and in fact also DF) is much better than the other measures (Table 2). On the combined CDB, even though the ISDB models consist of less than 20 percent of all

TABLE 4
A Summary of the Two Best Descriptors in Various Measurements on the Three Databases

Database	Nearest Neighbor	First Tier	Second Tier	E-Measure	DCG
Sub-PSB	LFD / SH	LFD / SH	LFD / CDF	LFD / CDF	LFD / SH
ISDB	CDF / DF	CDF / DF	CDF / DF	CDF / DF	CDF / SH
CDB	LFD / SH	CDF / DF	CDF / DF	CDF / LFD	CDF / LFD

We see that in most cases, the best results are achieved with the CDF signature.

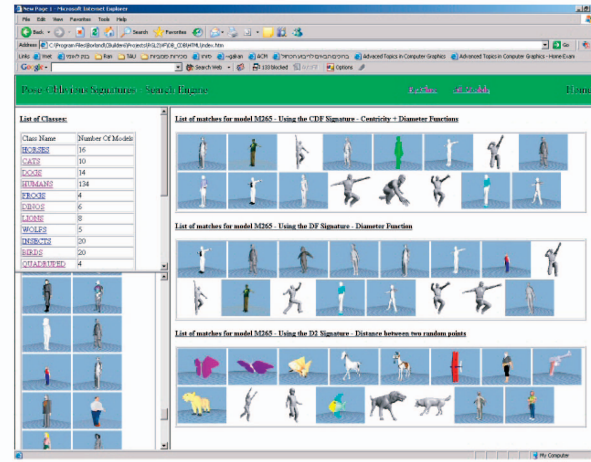


Fig. 9. A screen shot from the model search engine application.

the models, CDF remains one of the best measures (Table 3). In Table 4, we summarize the results by showing the two best signatures that achieve the highest results in each database and quality measure. In most cases, the best results are achieved using the CDF signature.

We also experimented with another “pose-invariant” signature based on multidimensional scaling (MDS) defined in [26]. Although this method gives good results on isometric surfaces that share the same geometric structure, it is too sensitive to topological changes and hard to control on general meshes. In Fig. 11 we compare various examples of our signature and MDS results.

7 SEARCH ENGINE

To study our pose-oblivious signatures, we have developed a simple search engine for 3D polygonal models (see Fig. 9). The motivation is to provide a tool with which users can retrieve models from a 3D model repository based on their shape attributes. In the current version, the user selects a 3D model from the database and the application computes the dissimilarity measure for all models in the database using the methods described in this paper. The application then shows the query model and the most similar models in the database.

Some example results obtained with this 3D search engine are shown in Fig. 10 at the end of the paper. The images in the leftmost column show the query 3D models, while the columns on the right show the closest matches among the 3D models in our CDB database using CDF signature. For instance, a query with a human model (top




Fig. 10. Examples of similarity retrieval results for shape matching queries using the CDF signature and CDB database. The first column on the left shows the query model and the other columns show the first eight closest models from the database in ascending order.

row) returns all humans in various poses; a query with an ant model returns a collection of ants and spiders with one helicopter and a turtle. A query of an airplane, which is a complex object with internal parts, returns three chairs.

In all examples, the query time on close to 500 models took under a second. For larger databases, more sophisticated indexing methods can further accelerate the performance. Even though this 3D search engine is rather simple, it shows the potential of pose-oblivious signatures for more intuitive search results. A more thorough examination of the results can be seen on our Web site at <http://www.faculty.idc.ac.il/arik/PoseOblivious/>. In this site, we have precalculated all results of queries on the CDB database and present them using images and html browsing. This site compares the results of the queries using CDF, LFD, SH, and D2 descriptors.

8 DISCUSSION AND CONCLUSIONS

In this paper, we have described a new type of 3D-model shape signature. This signature is expressive enough to discriminate between different models, and carries a number of attractive properties: it is rigid-body transformation and uniform scale invariant, it is not sensitive to topology changes of the model and, most importantly, it is pose-oblivious, i.e., it is insensitive to pose changes of the same object. The new signature achieves good performance and retrieval results for different classes of 3D models with the efficiency of comparing histogram signatures.

Evaluating the results on any database strongly depends on the definition of the classes and the number of models in each class. The PSB does not contain many articulated models in different poses. For this reason, we created a new database with 80 percent of its models from the PSB, and the rest, articulated models. Still, we show that the shape-oblivious




Fig. 11. Comparison of our CDF signature (middle row) and the MDS embedding in 3D (third row) for nine objects from our database. Examples (a)-(c) show that similar objects in similar poses can be mapped by MDS to different signature poses. The right foot is once in the front and once in the back. This can cause the matching to fail. In this case, the MDS embedding enhanced the problem of distinguishing between different poses. Examples (e) and (f) show that the MDS is extremely susceptible to topology changes. The geometry of these two examples differs by two triangles, while the MDS embedded results exhibit large differences as can be seen. Examples (g)-(i) show the inability of the MDS method to handle objects with complex geometry and a number of disconnected components. Last, examples (j)-(l) show how sometimes for relatively simple objects, MDS can map similar objects to different signatures and different objects to similar signatures. In this example using MDS, object (k) will match object (l) and would not match object (j). Note that the CDF signatures do not demonstrate such problems.

signature works well on the combined database and is much better on both natural or articulated objects. Its performance is reduced when the models include internal parts (Table 5). One possible solution to this is to combine several signatures together when building a 3D search engine.

A strong limitation of our approach is that the calculation of the diameter function is unreliable on nonvolumetric models, or on models which contain internal structures. For this reason, we could not reliably process many of the PSB models, and for the others, the results were not accurate. In many ways, the definition of a signature which is both effective and highly robust for object representation,

remains a challenge. We conclude that, although our shape signature achieves good results for a general-purpose database of models, it is best suited for a situation where articulated deformations are of importance.

In the future, we would like to enhance the calculations of the diameter function to cope with nonhollow objects as well. We would also like to investigate the issue of comparing two-dimensional signatures and its effect on the matching performance. Another possible direction for research is the development of pose-oblivious signatures that can deal with partial matching.

TABLE 5
Quality Results from Classes of Models in CDB of the CDF
and LFD (LFD was Taken as One of the Top Performing
Signatures on the CDB)

		Nearest Neighbor	First Tier	Second Tier	E- Measure	DCG
Humans (134)	CDF	87.57%	57.01%	83.76%	31.44%	89.15%
	LFD	86.30%	54.75%	80.97%	30.19%	87.64%
Human Hands (33)	CDF	81.82%	66.86%	76.28%	66.86%	86.90%
	LFD	81.82%	50.85%	69.60%	50.85%	81.16%
Horses (16)	CDF	93.75%	67.50%	85.00%	55.52%	91.58%
	LFD	75.00%	45.83%	75.42%	50.00%	74.29%
Insects (20)	CDF	71.00%	23.11%	33.78%	22.57%	61.96%
	LFD	75.00%	28.42%	32.68%	24.71%	63.71%
Airplanes (29)	CDF	44.83%	25.37%	43.47%	26.90%	62.10%
	LFD	96.55%	25.99%	31.28%	24.94%	70.70%
Guns (7)	CDF	71.43%	40.48%	50.00%	20.30%	68.01%
	LFD	100.00%	88.10%	92.86%	30.08%	97.63%
Chairs (33)	CDF	45.45%	20.45%	34.38%	20.45%	59.28%
	LFD	90.91%	43.37%	55.59%	43.37%	78.77%
Ships (21)	CDF	23.81%	11.43%	19.52%	12.82%	45.84%
	LFD	80.95%	39.52%	43.57%	32.23%	70.33%

Natural or articulated objects are matched better using CDF while, for some complex artificial objects, LFD can outperform CDF also because the diameter function is less reliable.

REFERENCES

- [1] M. Elad, A. Tal, and S. Ar, "Content Based Retrieval of VRML Objects: An Iterative and Interactive Approach," *Proc. Sixth Eurographics Workshop Multimedia*, pp. 107-118, 2001.
- [2] R. Osada, T. Funkhouser, B. Chazelle, and D. Dobkin, "Shape Distributions," *ACM Trans. Graphics*, vol. 21, no. 4, pp. 807-832, 2002.
- [3] M. Hilaga, Y. Shinagawa, T. Kohmura, and T.L. Kunii, "Topology Matching for Fully Automatic Similarity Estimation of 3D Shapes," *Proc. ACM SIGGRAPH*, pp. 203-212, Aug. 2001.
- [4] C. Zhang and T. Chen, "Efficient Feature Extraction for 2D/3D Objects in Mesh Representation," *Proc. Int'l Conf. Image Processing (ICIP '01)*, pp. 935-938, 2001.
- [5] J.R. Corney, H.J. Rea, D.E.R. Clark, J. Pritchard, M.L. Breaks, and R.A. MacLeod, "Coarse Filters for Shape Matching," *IEEE Computer Graphics and Applications*, vol. 22, pp. 65-74, 2002.
- [6] C.Y. Ip, D. Lapadat, L. Sieger, and W.C. Regli, "Using Shape Distributions to Compare Solid Models," *Proc. Seventh ACM Symp. Solid Modeling and Applications (SMA '02)*, pp. 273-280, 2002.
- [7] J. Munkres, *Elements of Algebraic Topology*. Addison-Wesley, 1984.
- [8] K. Siddiqi, A. Shokoufandeh, S.J. Dickinson, and S.W. Zucker, "Shock Graphs and Shape Matching," *Int'l J. Computing Vision*, vol. 35, no. 1, pp. 13-32, 1999.
- [9] D.-Y. Chen and M. Ouhyoung, "A 3D Object Retrieval System Based on Multi-Resolution Reeb Graph," *Proc. Computer Graphics Workshop*, pp. 16-20, June 2002.
- [10] H. Sundar, D. Silver, N. Gavvani, and S. Dickinson, "Skeleton Based Shape Matching and Retrieval," *Proc. IEEE Int'l Conf. Shape Modeling and Applications (SMI '03)*, pp. 130-139, May 2003.
- [11] T. Dey, J. Giesen, and S. Goswami, "Shape Segmentation and Matching with Flow Discretization," *Proc. Workshop Algorithms and Data Structures (WADS '03)*, pp. 25-36, 2003.
- [12] M. Mortara, G. Patané, M. Spagnuolo, B. Falcidieno, and J. Rossignac, "Blowing Bubbles for Multi-Scale Analysis and Decomposition of Triangle Meshes," *Algorithmica*, vol. 38, no. 1, pp. 227-248, 2003.
- [13] P. Shilane, P. Min, M. Kazhdan, and T. Funkhouser, "The Princeton Shape Benchmark," *Proc. IEEE Int'l Conf. Shape Modeling and Applications (SMI '03)*, pp. 167-178, June 2004.
- [14] J.W. Tangelder and R.C. Veltkamp, "A Survey of Content Based 3D Shape Retrieval Methods," *Proc. Int'l Conf. Shape Modeling and Applications*, pp. 145-156, June 2004.
- [15] 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*, 2004.
- [16] D.V. Vranić, D. Saupe, and J. Richter, "Tools for 3D Object Retrieval: Karhunen-Loeve Transform and Spherical Harmonics," *Proc. IEEE 2001 Workshop Multimedia Signal Processing*, pp. 293-298, 2001.
- [17] N. Iyer, Y. Kalyanaraman, K. Lou, S. Jayanti, and K. Ramani, "A Reconfigurable 3D Engineering Shape Search System Part I: Shape Representation," *Proc. ASME DETC 2003 Computers and Information in Eng. (CIE) Conf.*, 2003.
- [18] P. Min, "A 3D Model Search Engine," PhD dissertation, Dept. of Computer Science, Princeton Univ., 2004.
- [19] G. Cybenko, A. Bhasin, and K.D. Cohen, "Pattern Recognition of 3D CAD Objects," *Smart Eng. Systems Design*, vol. 1, pp. 1-13, 1997.
- [20] D. Saupe and D.V. Vranić, "3D Model Retrieval with Spherical Harmonics and Moments," *Proc. 23rd DAGM-Symp. Pattern Recognition*, pp. 392-397, 2001.
- [21] R. Ohbuchi, M. Nakazawa, and T. Takei, "Retrieving 3D Shapes Based on Their Appearance," *Proc. Fifth ACM SIGMM Int'l Workshop Multimedia Information Retrieval*, pp. 39-45, 2003.
- [22] R. Ohbuchi, T. Minamitani, and T. Takei, "Shape-Similarity Search of 3D Models by Using Enhanced Shape Functions," *Proc. Theory and Practice of Computer Graphics Conf.*, p. 97, 2003.
- [23] T. Funkhouser, P. Min, M. Kazhdan, J. Chen, A. Halderman, D. Dobkin, and D. Jacobs, "A Search Engine for 3D Models," *ACM Trans. Graph.*, vol. 22, no. 1, pp. 83-105, 2003.
- [24] M. Kazhdan, T. Funkhouser, and S. Rusinkiewicz, "Rotation Invariant Spherical Harmonic Representation of 3D Shape Descriptors," *Proc. Eurographics/ACM SIGGRAPH Symp. Geometry Processing (SGP '03)*, pp. 156-164, 2003.
- [25] D.-Y. Chen, X.-P. Tian, Y.-T. Shen, and M. Ouhyoung, "On Visual Similarity Based 3D Model Retrieval," *Computer Graphics Forum*, vol. 22, no. 3, pp. 223-232, 2003.
- [26] A.E. Elbaz and R. Kimmel, "On Bending Invariant Signatures for Surfaces," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 25, no. 10, pp. 1285-1295, 2003.
- [27] T. Zaharia and F. Prteux, "Three-Dimensional Shape-Based Retrieval within the MPEG-7 Framework," *Proc. SPIE Conf. Nonlinear Image Processing and Pattern Analysis XII*, pp. 133-145, 2001.
- [28] C.S. Chua and R. Jarvis, "Point Signatures: A New Representation for 3D Object Recognition," *Int'l J. Computer Vision*, vol. 25, no. 1, pp. 63-85, 1997.
- [29] A.E. Johnson and M. Hebert, "Using Spin Images for Efficient Object Recognition in Cluttered 3D Scenes," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 21, no. 5, pp. 433-449, May 1999.
- [30] D. Huber, A. Kapuria, R.R. Donamukkala, and M. Hebert, "Parts-Based 3-D Object Classification," *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR '04)*, June 2004.
- [31] A. Frome, D. Huber, R. Kolluri, T. Bulow, and J. Malik, "Recognizing Objects in Range Data Using Regional Point Descriptors," *Proc. European Conf. Computer Vision (ECCV)*, 2004.
- [32] H. Choi, S. Choi, and H. Moon, "Mathematical Theory of Medial Axis Transform," *Pacific J. Math.*, vol. 181, no. 1, pp. 57-88, 1997.
- [33] N. Amenta, S. Choi, and R. Kolluri, "The Power Crust, Unions of Balls, and the Medial Axis Transform," *Computational Geometry: Theory and Applications*, vol. 19, nos. 2-3, pp. 127-153, 2001.
- [34] T.K. Dey and W. Zhao, "Approximating the Medial Axis from the Voronoi Diagram with a Convergence Guarantee," *Algorithms—Proc. 10th Ann. European Symp. (ESA '02)*, pp. 387-398, 2002.
- [35] A. Shamir, D. Cohen-Or, and L. Shapira, "Consistent Partitioning Skeletonization and Mesh Analogies Using the Shape Diameter Function," submitted for publication.
- [36] Y. Rubner, C. Tomasi, and L.J. Guibas, "A Metric for Distributions with Applications to Image Databases," *Proc. Sixth Int'l Conf. Computer Vision (ICCV '98)*, p. 59, 1998.
- [37] D.-Y. Chen, X.-P. Tian, Y.-T. Shen, and M. Ouhyoung, Field Descriptors Web Page, <http://3d.csie.ntu.edu.tw/~dynamic/3DRetrieval/>, 2005.
- [38] M. Kazhdan, "Executables for Computing/Comparing the Spherical Harmonic Representations of 3D Models," <http://www.cs.jhu.edu/~misha/HarmonicSignatures/>, 2005.



Ran Gal received the BSc degree in mathematics and computer science from Bar-Ilan University in 1994, and the MSc degree in computer science from Tel Aviv University in 2006. Currently, he is a PhD student in the School of Computer Science at Tel Aviv University. His research interests are in computer graphics and include shape modeling, shape matching, mesh processing, and image synthesis. He is a student member of the IEEE.



He is a senior member of the IEEE and the IEEE Computer Society.

Daniel Cohen-Or received the BSc degree in both mathematics and computer science (1985), the MSc degree in computer science (1986) from Ben-Gurion University, and the PhD degree from the Department of Computer Science (1991) at the State University of New York at Stony Brook. He is an associate professor in the School of Computer Science at Tel Aviv University. His current research interests include rendering, visibility, shape modeling, and image synthesis.



Ariel Shamir received the BSc degree in mathematics and computer science from Hebrew University in 1991, and the PhD degree in computer science in 1999. He spent two years in the Center for Computational Visualization at the University of Austin, Texas, and then joined the computer science faculty at the Interdisciplinary Center in Herzliya, Israel. Dr. Shamir is currently a visiting scientist at Mitsubishi Electric Research Laboratory. His

research interests include geometric modeling, computer graphics, virtual reality, and visualization. He is a member of the IEEE and the IEEE Computer Society.

► **For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.**