Curvature Maps For Local Shape Comparison

Timothy Gatzke
and Cindy Grimm
Washington University in St. Louis
St. Louis, Missouri, USA
Telephone: (314) 935–4576

Michael Garland
and Steve Zelinka
University of Illinois
Urbana-Champaign, Illinois, USA
Telephone: (800) 555–1212

***Abstract*— The ability to identify similarities between shapes is important for applications such as medical diagnosis, object registration and alignment, and shape retrieval. In this paper we present a method, the *Curvature Map*, that uses surface curvature properties in a region around a point to create a unique signature for that point. These signatures can then be compared to determine the similarity of one point to another. To gather curvature information around a point we explore two techniques, rings (which use the local topology of the mesh) and Geodesic Fans (which trace geodesics along the mesh from the point). We explore a variety of comparison functions and provide experimental evidence for which ones provide the best discriminatory power. We show that Curvature Maps are both more robust and provide better discrimination than simply comparing the curvature at individual points.**

I. Introduction

In this paper we address the problem of local surface similarity, *i.e.*, is a region of a surface the same "shape" as another region? This information is useful for a variety of applications. For example, identifying corresponding regions between two similar surfaces is a necessary first step toward alignment and registration of those surfaces. Previous approaches to local surface matching have either focused on man-made objects, where features are easy to find, or required some type of user interaction to select features. Manual selection of corresponding features and subjective determination of the difference between objects are both time consuming processes requiring a high level of expertise. Our approach automates this process, while still providing the user with control over what aspects of the surface match are important.

A. Approach

Our approach uses curvature, which is an intrinsic property of the surface, as a base metric. Because curvature is a point metric, it does not provide information about the region around the point. To incorporate local shape information, we define a *curvature map* around a given vertex v. This curvature map accumulates curvature information from a region around v, and can take one of two forms: A one-dimensional (1-D) map, which only considers the distance from v, or a two-dimensional (2-D) map that uses both the distance and the orientation information. Note that using just the curvature at v is the 0-D form of the curvature map function.

We investigate various methods of building curvature maps from both mean and Gaussian curvature, and the effect of the size of the region. We then define a similarity function that compares two curvature maps.

B. Contribution

In this paper we develop the curvature map and comparison functions for local shape similarity. Curvature maps are robust with respect to grid resolution and mesh regularity. Both the 1-D and 2-D comparison functions yield a high degree of discrimination for local shapes, compared to the 0-D methods which have been used previously. Curvature calculation on discrete meshes is often noisy [1] and not always accurate [2]. Because curvature maps combine curvature information over a region, they are less susceptible to these issues.

Section II discusses previous work. In Section III we define curvature maps, including how we calculate curvature, define a local region on the surface, and different similarity measures. In Section IV, we evaluate the various similarity measures using both a test shape with known curvature, and several common meshes. Section V summarizes the conclusions of this study and outlines possible areas for future work.

II. Previous Work

Similarity measures based on distances between sets of points, feature vectors, histograms, signatures, and graph representations can be found in object recognition, three-dimensional model matching, computer vision, feature detection, correspondence, registration, and pose estimation. These methods are primarily global rather than local in nature, *i.e.*, they match entire surfaces. A few of these techniques have been applied to local surface matching; we discuss these in more detail.

Shum et al. [3] use the L_p distance between local curvature functions mapped to a semi-regular triangulation of the unit sphere as a local measure; unfortunately, this technique is only applicable to closed surfaces which are topologically spherical. A number of segmentation methods also use curvature, particularly the sign of the curvature [4] [5], isosurfaces and extreme curvatures [6], or watersheds of a curvature function [7] [8] [9]. Watershed algorithms show sensitivity to noise and to the user-specified watershed depth threshold. Splitting the surface into regions still gives only coarse information about the differences between local regions, and small changes to the shape can make large changes in the segmentation.

There have been a few attempts to create local signatures. Planitz et al. [10] propose a signature based on a local region around select vertices. However, the use of distances and angles between normals for points in a local support region makes this method sensitive to point distributions. Shape contexts [11] represent the shape of an object, with respect to a particular point on the object, as a 2-D histogram of the relative coordinates of other points sampled from the surface. The sampling of points limits this method for detailed shape matching.

These similarity measures are applicable to coarse shape matching for shape retrieval, but generally provide limited discrimination between similar shapes. Moreover, in general, methods based on distances between points, such as Hausdorff distance, multi-resolution Reeb graphs [12], shape distributions [13] [14], and spin images [15], are sensitive to the distribution of the points.

III. Comparing Local Surface Shape

This section describes the Curvature Map, and how it is used to identify regions of similar shape. We first define two methods for creating samples around the point, one based on the mesh topology and one based on geodesic sampling. Next, we describe how we calculate curvature on the mesh. Finally, we define the comparison function itself.

A. Defining Rings of a Mesh

Given a specified vertex of the mesh, we can define a set of "rings" around the vertex using the existing mesh structure. The i^{th} ring around Vertex v_0 is defined as the set of vertices $v \in V$ such that there exists a shortest path from v_0 to v containing i edges. The set of rings $R_i, i \leq N$ defines the N-Ring neighborhood about v_0.

Figure 1 shows the first nine rings around a selected vertex of the mesh. The ring structure can be extended an arbitrary distance from any point; however, as the distance increases, the shape of the ring may become irregular.

B. Geodesic Fans

Geodesic fans [16] represent a local surface resampling that provides a uniform neighborhood structure around a vertex. In particular, a geodesic fan consists of a set of spokes, and a set of samples on each spoke. The spokes are geodesics marched out across the surface from the neighborhood center, equally spaced in the conformal plane of the neighborhood's 1-ring. With the samples equally spaced along each spoke, they form a local geodesic polar map around the vertex. Each set of points equi-distant from the neighborhood center is treated as a ring. Following Zelinka and Garland, we use interpolated normal geodesics [17] where possible, reverting to straightest geodesics [18] if the smoothness criterion for interpolated normal geodesics is not met.

We use this procedure to generate fans at each vertex of the mesh. Sample fans at two vertices are shown in Figure 2. Each fan point is defined in terms of the Barycentric coordinates in some triangular face in the original mesh. These Barycentric

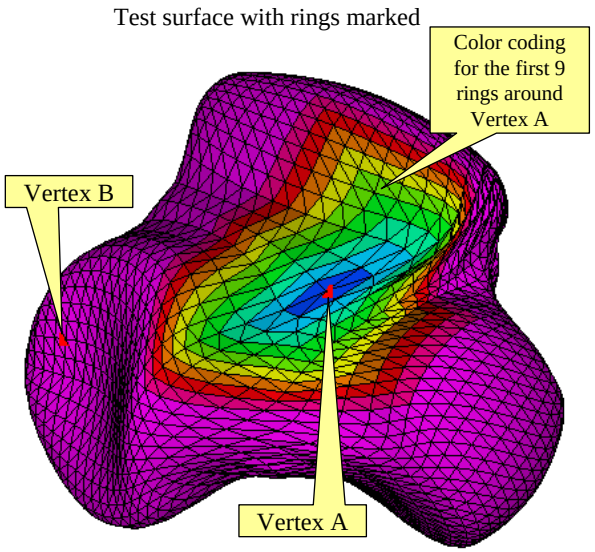

Fig. 1. Test surface with Vertices A and B highlighted. The first nine rings defined around Vertex A are color coded. The mesh is fairly uniform except for blending between sections. Note that the ring structure is still well-defined in spite of the skewness near its right edge.

coordinates are used to interpolate curvature values defined on the mesh to the fan point. This forms a uniform sampling of curvature data around each vertex. As the sampling increases, more overhead is required to store the fan data.

The regularity of geodesic fans can break down as the distance from the point increases, due to a) stretching of the circumferential spacing while the radial spacing remains uniform, and b) issues in constructing geodesics over longer distances. As a result, the fan resolution may be locally finer, coarser, or both, when compared to the mesh resolution. If the sampling is coarser than the mesh triangle size, then the geodesic fan will not incorporate all of the curvature data available.

C. Estimating Curvature

Gatzke and Grimm [2] evaluate various curvature estimation methods for triangular meshes. Based on their results, we choose an algorithm that fits a 2-Ring neighborhood using a natural parameterization of the input mesh [19]. This method is reasonably robust with respect to noise as well as mesh irregularity, and provides consistent accuracy of the curvature values. Gaussian curvature and mean curvature are plotted as scalar properties on the surface of the test shape in Figure 3.

D. 1-D Curvature Maps

The 1-D form of the curvature map is defined over M rings, where the rings come from either the mesh structure or the geodesic fan structure. Each point p_i in the map is constructed from data accumulated along the ring R_i. The point p_i can have one or more data values; this allows us to compare, for example, both the Gaussian and the mean

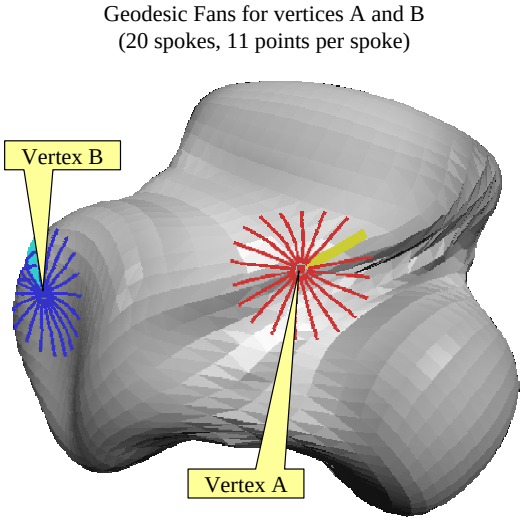

Fig. 2. Geodesic fans at two vertices. The first spoke of each fan is highlighted and used to track the relative orientation for 2-D fan comparisons. Fan parameters include the number and length of spokes, and the number of points per spoke.

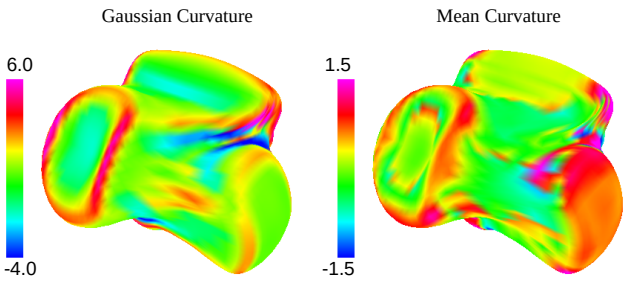

Fig. 3. Gaussian curvature (left) and mean curvature (right). Note that the Gaussian curvature ranges over $[-4, 6]$, while the mean curvature ranges over $[-1.5, 1.5]$.

curvature simultaneously (see Figure 4). Each element of p_i generates a curve as a function of the ring distance [1].

Because the Gaussian curvature is a product of the principal curvatures and the mean curvature is an average, the Gaussian curvature magnitudes will be roughly proportional to the square of the mean curvatures. A square root function applied to the Gaussian curvatures gives a more equal weighting. Similarly, we use a logarithmic function to reduce the effect of large variations in the peak curvature values, which tend to dominate over areas with lower curvature magnitudes. The emphasis here is to match the shape of curves rather than just the magnitude of its peaks.

More formally, the curvature map κ_{map} at a vertex v is a set of N piecewise linear functions defined over the rings R_i:

[1]The curvature map is formulated for a discrete mesh, but the same concept can be applied to an analytic surface, where the curve values for discrete increments would be replaced by a continuous function on the surface.

$$\kappa_{map} = \{f^j : r_i \rightarrow \Re\}_{0<j<N},\ 0 < i < M \quad (1)$$

$$r_i = \sqrt{A_i/\pi} \quad (2)$$

$$g(\kappa) = \begin{cases} \frac{1}{N_i}\sum_{w\in R_i}\kappa(w) & or \\ \max_{w\in R_i}\kappa(w) & or \\ \min_{w\in R_i}\kappa(w) \end{cases} \quad (3)$$

$$h_1(x) = \begin{cases} x & or \\ sign(x)\ sqrt\|x\| \end{cases} \quad (4)$$

$$h_2(x) = \begin{cases} x & or \\ sign(x)\log(1+\|x\|) \end{cases} \quad (5)$$

$$f^j(\kappa) = h_2 \circ h_1 \circ g(\kappa) \quad (6)$$

where A_i the area of the i-Ring neighborhood. The functions f^j can be applied to Gaussian κ_g or mean κ_m curvature. r_i is used to normalize the parameterization of the f^j curves with respect to the area covered by the region.

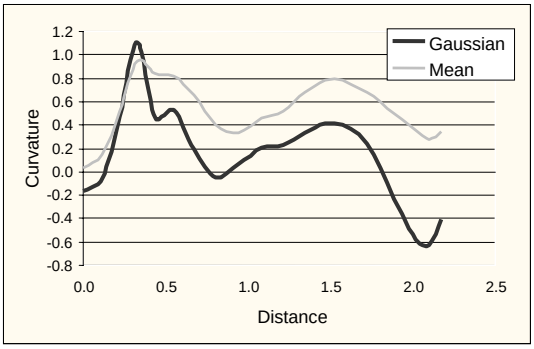

Fig. 4. Maps of Gaussian Curvature and Mean Curvature as a function of distance from the point. Peaks and valleys tend to be more pronounced for the Gaussian curvature curve, which is the product of the principal curvatures ($\kappa_g = \kappa_1\kappa_2$), compared to the mean curvature, which is an average ($\kappa_m = \frac{\kappa_1+\kappa_2}{2}$).

To compare the shape at two points, such as those shown in Figure 1, we compare the corresponding curvature map functions (see Figure 5). The shape similarity, S is a function of the difference between the individual curves. Let f_A be the set of curves for one point, and f_B the curves for the second point.

$$S = \sum_j \int_0^R \left(\|(f^j)_A(r) - (f^j)_B(r)\|\right) dr \quad (7)$$

Note that the difference we compute is actually a dissimilarity measure, with zero indicating high similarity and positive values indicating the relative difference between shapes. The user can also specify the radial distance over which the curvature maps are compared. This provides a parameter to control the size of the region used to compute similarity between points.

E. 2-D Curvature Maps

The 2-D curvature map is similar to the 1-D map, except that we maintain the angular (θ_k) information and accumulate

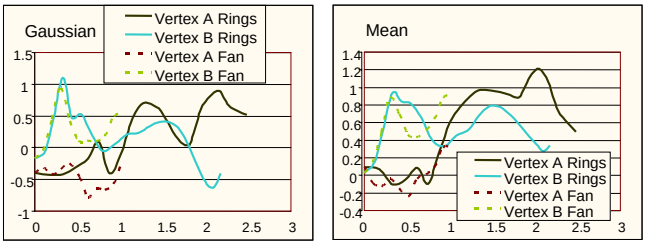

Fig. 5. Map of Gaussian (left) and mean (right) curvature as a function of distance from the selected vertex. The ring-based and fan-based curves for a particular point start out at the same value and initially have similar shape, but diverge due to a) non-uniformity of the rings, and b) fans sampling only a subset of the data, as the distance from the center increases. In this case, the fans cover a smaller area than the rings.

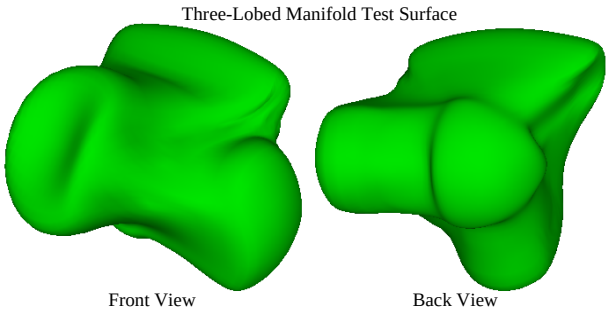

Fig. 6. Two views of the test surface used for shape comparison. The left and right lobes in the front view are the same except for the addition of a dent (concave region) in the end of the left lobe.

data only along a single spoke (*i.e.*, there is one f^j per spoke). Let N_s be the number of spokes:

$$\kappa_{map} \quad = \quad \{(f_k^j)\}_{0<j<N, 0<k<N_s)}\} \tag{8}$$

The comparison metric sums up the curve differences along each spoke. There are N_s possible alignments between two fans; we calculate $S2$ for each alignment and chose the smallest value.

$$S2 \quad = \quad \sum_j \sum_k \int_0^R \left(\|(f_k^j)_A(r) - (f_k^j)_B(r)\| \right) dr \tag{9}$$

It is important that the fans are generated with the same number of spokes. By checking all possible relative orientations of the fans, the 2-D form can also provide information about the relative orientation of the points. As with the 1-D curvature map, the user can chose the size of the region to compare over by selecting R.

IV. Discussion

To evaluate our metrics we created a test shape with known curvature properties (see Figure 6). Because this manifold surface is defined parametrically, we can easily generate a range of cases for testing that cover curvatures found in realistic applications. We also applied the curvature map to standard meshes such as the Stanford Bunny mesh.

We first look at the discrimination power of the 0,1, and 2-D curvature maps, using the "best" f^j functions for each case. Next, we describe our study to determine which f^j functions have the best discrimination power. Finally, we look at evaluation times for each of the techniques.

A. Comparing 0-,1-, and 2-D Curvature Maps

We compare the 0-, 1-, and 2-D curvature maps for our test shape and the bunny. The top (Vertex A) and bottom (Vertex B) rows of images in Figure 7 show which points on the surface are most similar to the selected vertex. For all of these images, we apply the square root and logarithmic functions to the average Gaussian curvature, and the logarithmic function to the average of the mean curvature. As expected, the number of similar points decreases as we increase the dimension of the curvature map.

The ring and fan-based 1-D methods are similar in discriminatory power, but differ slightly in which points they mark as similar. Small differences may be due to differences in the size and shape of the regions covered by the rings and fans. We also varied the size of the region covered by the fans, keeping the same number of spokes and number of points along each spoke. The results remained similar as long as we adjusted the number of rings to match the approximate region sizes.

1) Choice of Comparison Functions: The visualization of similarity as a scalar function plotted on the surface of the object gives an indication of the improved ability to differentiate based on shape, but is not as useful in determining which of our 1-D curvature map functions, and associated comparison functions works best. To test these options, we identify groups of points that we expect to be similar, based on our intuition. The similarity for each pair of points is used to form a distance grid. Distance grids for 0-D, 1-D ring-based, 1-D fan-based, and 2-D methods are shown in Figure 10. We chose eight groups, where each group contains three vertices. Group A is located in the concave region of one lobe. Groups B, C, and D are in three saddle regions occurring between pairs of lobes. Group E is in the crease along the rounded back of the main body. Groups F, G, and H are on convex regions of the body and two lobes respectively.

Comparing the distance grids allows us to evaluate various combinations of comparison functions. The comparison function having the most similarity between points of the same group (darkest 3×3 boxes along the diagonal), with much less similarity (lighter) for dissimilar groups, was deemed best. The first five groups include concave regions, while the last three are primarily convex, so similarity between certain groups is expected.

Average mean curvature with the square root function applied to the average Gaussian curvature gave the best discrimination in our tests. The logarithmic function has a less significant effect, but this importance may depend on the nature of the curvature peaks. We varied the number of

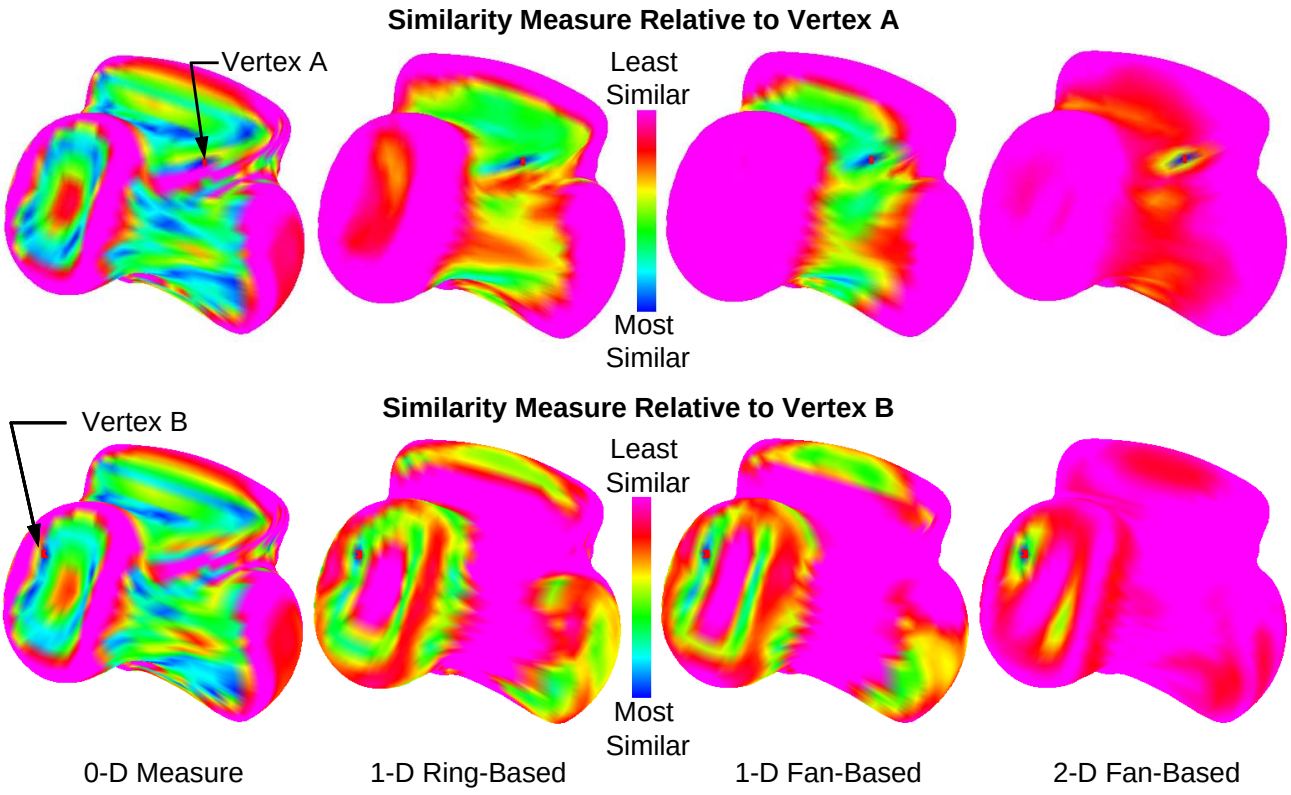

Fig. 7. Top: Similarity measure relative to Vertex A plotted on the test surface. Bottom: Vertex B. The color scale ranges from blue (high similarity to the selected point) to magenta (most dissimilar). Nine rings were used in the ring-based calculation. 20 spokes, 11 samples per spoke, were used in the 1- and 2-D fan-based calculation; the surface area is approximately the same as the ring version. Note that the 0-D measure (far left) is very noisy compared to the 1-D ring-based (center left) and fan-based (center right) measures. The 2-D measure (far right) shows few points with similarity to the selected vertex.

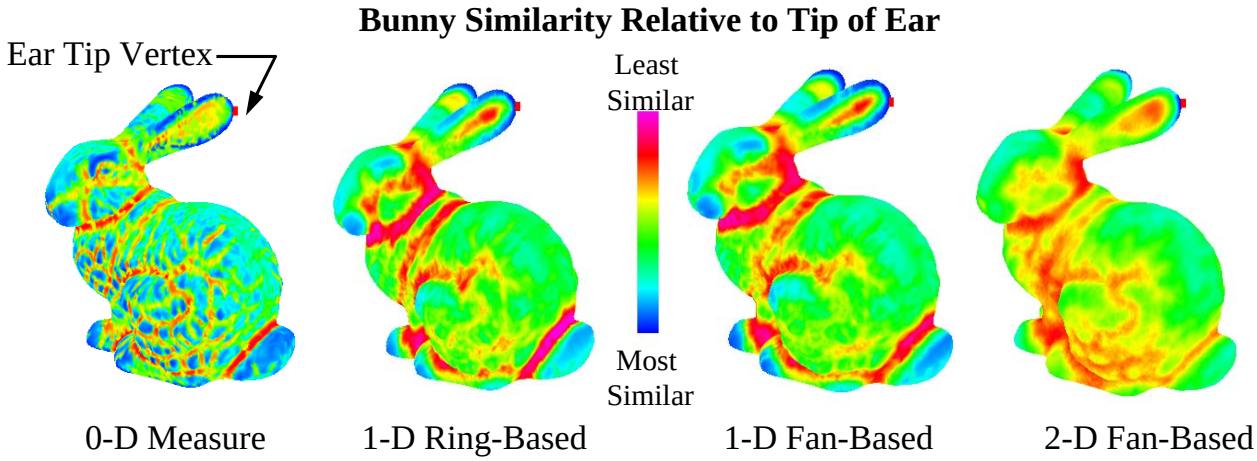

Fig. 8. Similarity measure relative to a vertex on the tip of the ear of the Stanford Bunny. The color scale ranges from blue (high similarity to the selected point) to magenta (most dissimilar). The 0-D similarity has significant noise, while the 1-D methods isolates the tips of the ears much more cleanly. The 2-D method is even more discriminating, with similarity limited to the tip of the other ear.

rings over a wide range, but for our test case, there was little change after about eight rings. Using fewer rings caused more degradation as we approach the 0-D curvature map. Using the minimum curvature or maximum curvature, instead of the average over the ring, performed poorly. Using a vector of both the minimum and maximum curvature in a ring did much better, but was not quite as effective as the average.

The 1-D ring-based method generates the highest degree of self-similarity within the groups. The 1-D fan-based method

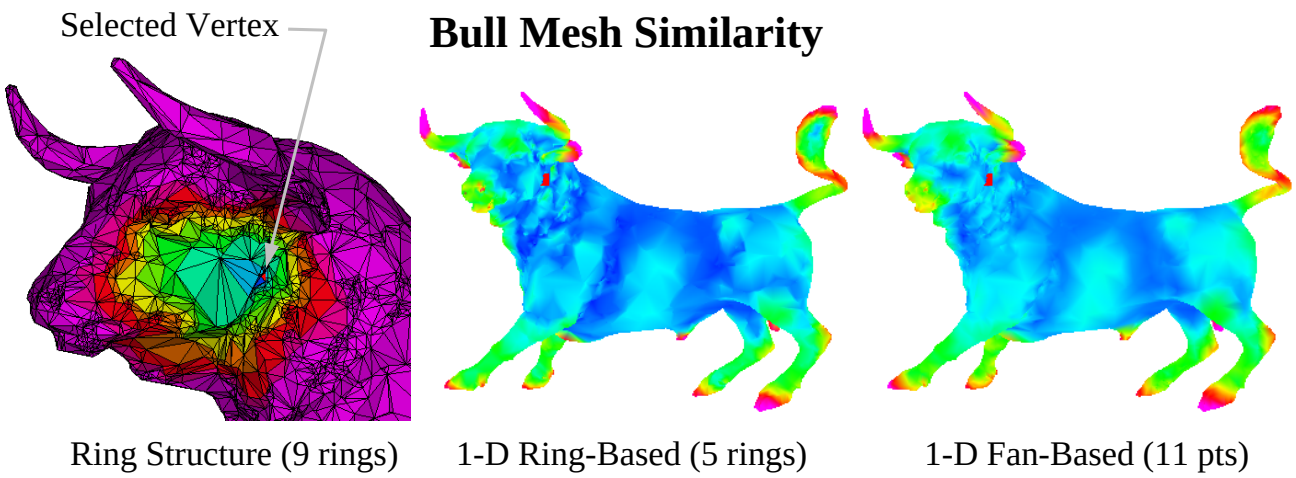

Fig. 9. Similarity measure for the Bull mesh. The color scale ranges from blue (high similarity to the selected point) to magenta (most dissimilar). The view on the left shows that the ring structure is very non-symmetric about the selected vertex, due to the irregularity of the bull mesh. Even so, the ring-based and fan-based 1-D methods provide comparable similarity measures.

does not do quite as well within groups, but is good at distinguishing between the groups. The 0-D method does not differentiate between Groups A and B, and has poor self-similarity for Groups C through E. All three methods have just subtle differences for the last three groups. Overall, the ring-based 1-D method most consistency indicates more similarity withing the group than between groups.

B. Applying Curvature Maps to Other Objects

To test how our new similarity measure works in practice, we apply it to a mesh of the Stanford Bunny. The bunny has a much more irregular surface, with regions of similar curvature, but quite a bit of local curvature variation. As Figure 8 (far left) shows, the 0-D (point curvature) similarity is very noisy due to these local curvature variations. The 1-D ring-based similarity measure (second from left) was generated from Gaussian curvature with log and square root functions and mean curvature with a log function applied, compared over eight rings. The same functions applied to the 1-D similarity based on eleven fan points and 2-D similarity are shown in the second from the right and far right images of Figure 8 respectively. The results are consistent with our test surface, i.e., the 0-D method is extremely noisy, both 1-D methods identify much smaller and more consistent regions of similarity. The 2-D method has even more differentiation between ear tip points and points not on the ear tip, with similarity indicated only for the tip of the other ear.

We also apply curvature maps to the mesh of a bull. This mesh is highly irregular, causing the ring structure to be asymmetric about the selected vertex, as shown in Figure 9. However, the ring-based and fan-based 1-D methods still provide similar results.

C. Efficiency Comparison

We also made comparisons of the speed of the methods for the test shape and the bunny mesh. Table 1 contains pre-processing times for computing curvature on the mesh, creating a ring-based curvature map, and creating a fan-based curvature map. All times are per mesh vertex. Identifying the ring structure around each vertex is included in the ring-based map times, and fan generation time is added to the map creation time for the fan-based maps. Table 2 shows the times for computing the similarity of each point of the mesh relative to a selected point, normalized by the number of vertices. The 1-D and 2-D methods were timed for four, eight, and eleven rings/points. All times were computed on a 1.7 GHz Pentium M processor. Some inaccuracy in the smaller times for the test shape is due to approaching the resolution of our timing algorithm. The comparison functions are much faster than the pre-processing step, with the 0-D and 1-D methods a few orders of magnitude faster than the 2-D comparisons.

TABLE I
PREPROCESSING TIMES (PER VERTEX) FOR 1.7 GHZ PENTIUM M PROCESSOR

Preprocessing Times - msec/vertex		
	Test Shape	Bunny
Compute Curvature	1.5	1.6
Ring-based Map	1.8	30.0
Fan-based Map	3.4 to 10.4†	5.2 to 26.1†

† Time is proportional to the physical length of fan spokes

D. Curvature Maps for Finding Unique Features

In order to look for key features in the mesh, we look for the groups of points that are least similar to the remaining points. For each point, we compute its similarity to all other points, and then sort these by decreasing similarity. A Gaussian function is applied to the sorted similarity curves, and the resulting contribution, which represents a non-parametric kernel density estimate, quantifies how many other points the given point is similar to. The smallest values indicate the points most

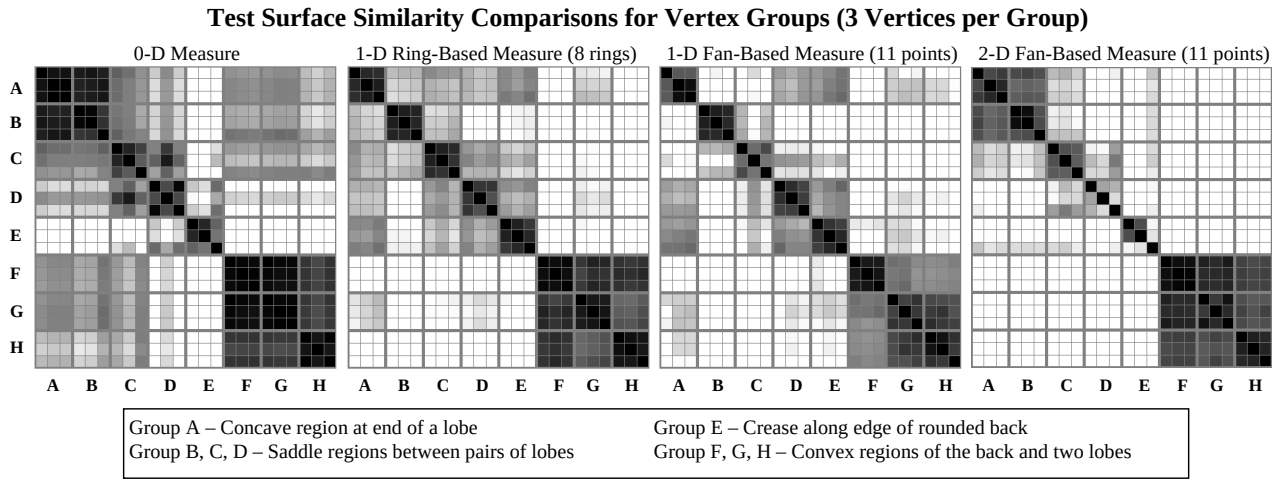

Fig. 10. Distance grids for select points. The similarity within groups, indicated by the darkest 3×3 boxes along the diagonal, and dissimilarity between groups, based on lighter off-diagonal squares, was most consistent for the 1-D ring-based measure.

TABLE II
COMPARISON TIMES (PER VERTEX) FOR 1.7 GHZ PENTIUM M PROCESSOR

Comparison Times - µsec/vertex		
Comparison Method	Test Shape	Bunny
0-D (Point curvature)	2.9	1.1
1-D Ring-based Map (4 pts)	6.4	4.6
1-D Ring-based Map (8 pts)	10.1	8.6
1-D Ring-based Map (11 pts)	12.1	12.1
1-D Fan-based Map (4 pts)	7.5	4.4
1-D Fan-based Map (8 pts)	11.0	8.5
1-D Fan-based Map (11 pts)	12.4	11.5
2-D Map (4 pts)	672	671
2-D Map (8 pts)	1283	1287
2-D Map (11 pts)	1584	1597

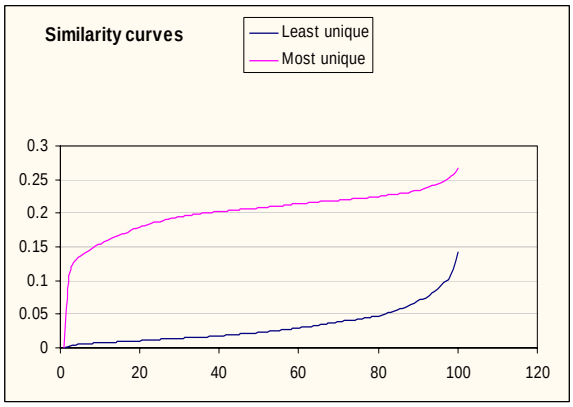

Fig. 11. The similarity curves for the most and least similar vertices.

different from the general population. The sorted similarity curves for the highest and lowest similarity density are shown in Figure 11 (the kernel value was set to be 0.05 at the 100th point). The three-hundred most unique points are highlighted for the 0-D (left) and 1-D (right) methods in Figure 12. The 1-D mehtod picks up more consistent point groupings than the 0-D method. This is apparent in both the neck region and on the tail.

E. Recommendations

The ring-based and fan-based 1-D methods are comparable in ability to discriminate, comparison times, and setup times. Ring-based methods are more appropriate for larger regions, provided the mesh is fairly uniform. If storage space is not an issue, the fan-based 1-D method provides a more consistent comparison for smaller local regions.

The 1-D ring-based method can also be used to pre-process a mesh to identify regions that are similar. The more expensive, but more exact, 2-D method can then be applied just to these regions.

In summary, it is preferable to use the ring-based 1-D method for comparing larger local regions, as long as the mesh quality is reasonable. The slower 2-D methods can be reserved for the final stage when exact matching is required.

V. CONCLUSIONS AND FUTURE WORK

We have presented the curvature map, a new method for comparing local shape based on surface curvature. It has been applied as a 1-D method on N-ring neighborhoods and as a 1-D or 2-D method on Geodesic fans. Point curvature (0-D) methods do a poor job of distinguishing between local regions. Curvature maps demonstrate improved capability to discriminate shape as compared to these 0-D methods.

Determining how far out to go when comparing local shape is still an open issue, and is likely case dependent. Assessing which comparison settings consistently produce good results for a wide range of shapes is another area that could benefit from further research, but the methods used here have a

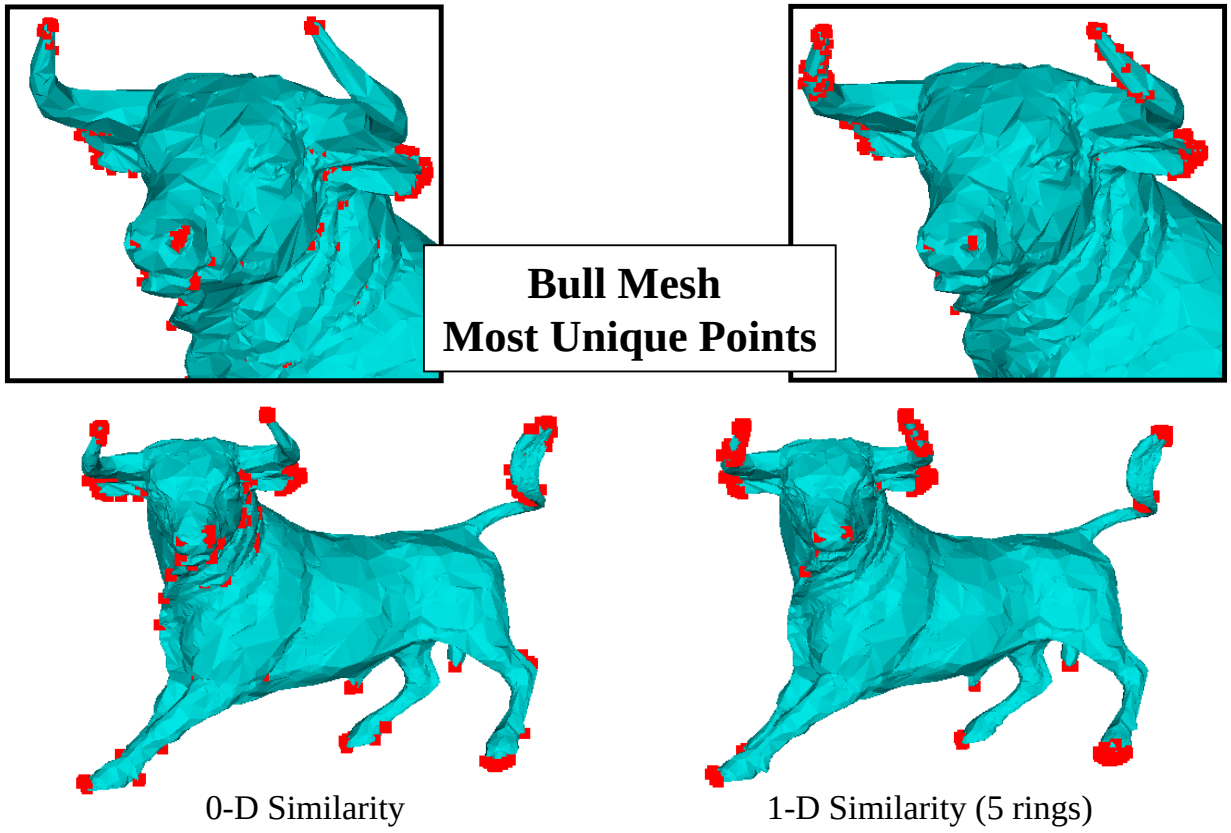

Fig. 12. The three-hundred most unique points based on similarity to to all other points. The 0-D method (left) picks up most of the peak curvatures, but finds a lot of isolated points in the neck and face region. The 1-D method (right) finds consistent groups of points reflecting key features in the mesh.

reasonable grounding in intuition, and appear to be a good place to start.

Curvature maps offer a valuable capability to differentiate local shapes. These methods will be applied to the shape matching problem to identify corresponding points based on shape similarity. Additional work will extend the comparison methods to account for shape similarity when objects or portions of objects are scaled differently.

REFERENCES

[1] A. Hertzmann and D. Zorin, "Illustrating smooth surfaces," in *Siggraph 2000, Computer Graphics Proceedings*, K. Akeley, Ed. ACM Press / ACM SIGGRAPH / Addison Wesley Longman, 2000, pp. 517–526.

[2] T. D. Gatzke and C. M. Grimm, "Assessing curvature metrics on triangular meshes," Washington University, St. Louis Missouri, Tech. Rep., jun 2003.

[3] H.-Y. Shum, M. Hebert, and K. Ikeuchi, "On 3d shape similarity," in *Proc. IEEE Computer Vision and Pattern Recognition*, 1996, pp. 526–531.

[4] A. M. McIvor, D. W. Penman, and P. T. Waltenberg, "Simple surface segmentation," in *DICTA/IVCNZ97*, Massey University, New Zealand, Dec. 1997, pp. 141–146.

[5] R. Wilson and E. R. Hancock, "Consistent surface curvature labelling."

[6] F. Vivodtzev, L. Linsen, G.-P. Bonneau, B. Hamann, K. I. Joy, and B. A. Olshausen, "Hierarchical isosurface segmentation based on discrete curvature," in *Proceedings of VisSym'03- Data Visualization 2003*. New York, New York: Association for Computing Machinery, 2003.

[7] A. Mangan and R. Whitaker, "Surface segmentation using morphological watersheds," in *IEEE Visualization '98: Late Breaking Topics*, Oct. 1998, pp. 29–32.

[8] A. P. Mangan and R. T. Whitaker, "Partitioning 3d surface meshes using watershed segmentation," in *IEEE Transactions on Visualization and Computer Graphics 5(4)*, 1999, pp. 308–321.

[9] S. Pulla, A. Razdan, and G. Farin, "Improved curvature estimation for watershed segmentation of 3-dimensional meshes," Apr. 2001.

[10] B. M. Planitz, A. J. Maeder, and J. A. Williams, "Intrinsic correspondence using statistical signature-based matching for 3d surfaces," 0000.

[11] M. G., S. Belongie, and H. Malik, "Shape contexts enable efficient retrieval of similar shapes," in *CVPR 1*, 2001, pp. 723–730. [Online]. Available: citeseer.nj.nec.com/article/mori01shape.html

[12] M. Hilaga, Y. Schinagawa, T. Kohmura, and T. L. Kuni, "Topology matching for fully automatic similarity estimation of 3d shapes," in *Computer Graphics (Proc. SIGGRAPH 2001)*, 2001, pp. 203–212.

[13] R. Osada, T. Funkhouser, B. Chazelle, and D. Dobkin, "Matching 3D models with shape distributions," 2001, pp. 154–166. [Online]. Available: citeseer.nj.nec.com/osada01matching.html

[14] ——, "Shape distributions," in *ACM Transactions on Graphics, 21(4)*, Oct. 2002.

[15] A. Johnson and M. Hebert, "Recognizing objects by matching oriented points," 1997. [Online]. Available: citeseer.nj.nec.com/article/johnson97recognizing.html

[16] S. Zelinka and M. Garland, "Similarity-based surface modelling using geodesic fans," in *Proc. Second Eurographics Symposium on Geometry Processing*, 2004, pp. 209–218.

[17] H. Biermann, I. Martin, F. Bernardini, and D. Zorin, "Cut-and-paste editing of multiresolution surfaces," in *Proceedings of the 29th annual conference on Computer graphics and interactive techniques*. ACM Press, 2002, pp. 312–321.

[18] K. Polthier and M. Schmies, “Straightest geodesics on polyhedral surfaces,” *Mathematical Visualization*, p. 391, 1998.
[19] M. Desbrun, M. Meyer, and P. Alliez, “Intrinsic parameterizations of surface meshes,” *Eurographics 2002, Vol. 21, Number 2*, 2002.