

A Part-aware Surface Metric for Shape Analysis

Rong Liu¹, Hao Zhang¹, Ariel Shamir², and Daniel Cohen-Or³

¹Simon Fraser University, ²The Interdisciplinary Center, ³Tel Aviv University

Abstract

The notion of parts in a shape plays an important role in many geometry problems, including segmentation, correspondence, recognition, editing, and animation. As the fundamental geometric representation of 3D objects in computer graphics is surface-based, solutions of many such problems utilize a surface metric, a distance function defined over pairs of points on the surface, to assist shape analysis and understanding. The main contribution of our work is to bring together these two fundamental concepts: shape parts and surface metric. Specifically, we develop a surface metric that is part-aware. To encode part information at a point on a shape, we model its volumetric context – called the volumetric shape image (VSI) – inside the shape’s enclosed volume, to capture relevant visibility information. We then define the part-aware metric by combining an appropriate VSI distance with geodesic distance and normal variation. We show how the volumetric view on part separation addresses certain limitations of the surface view, which relies on concavity measures over a surface as implied by the well-known minima rule. We demonstrate how the new metric can be effectively utilized in various applications including mesh segmentation, shape registration, part-aware sampling and shape retrieval.

Categories and Subject Descriptors (according to ACM CCS): Computer Graphics [I.3.5]: Computational Geometry and Object Modeling—Boundary representations, Geometric algorithms, languages, and systems

1. Introduction

The fundamental geometric representation of 3D objects in computer graphics is surface-based. Solutions of many problems involving the analysis and understanding of a 3D object utilize a metric, which prescribes a distance function over the boundary surface of the object. Shape decomposition, construction of shape signatures, sampling, and surface parameterization are all examples where a surface metric would play a role. Typically, finding the right metric which captures the essence of the problem at hand is the key.

Well-known surface metrics include geodesic [Car76] and isophotic [PHHH04] distances, where the latter measures angles between surface normals; a combination of the two is also common [LZH*07]. Anisotropic geodesic metric based on curvature tensor is also considered in [SJC08] for parametric surfaces. Another metric receiving recent attention is diffusion distance [dGGV08], which accounts for the degree of connectedness, over the surface, between two points. Moreover, one can always derive a discrete distance func-

tion for mesh primitives by assigning attributes to primitives, defining edge weights between adjacent attributes, and finally computing distances over the shortest paths in the primal or dual graph of the mesh, such as the one employed in [KT03] for segmentation. We adopt this strategy in this paper. Since the resulting distance function is the shortest graph distances, it is indeed a *metric* that satisfies positive definiteness, symmetry, and the triangle inequality.

1.1. Motivation and contribution

In this paper, we develop a surface metric that is part-aware. Our work is motivated both by the wide-ranging use of surface distance functions in geometry processing and the importance of shape parts in object recognition [Heb49, HR84]. Beyond recognition tasks, the ability to infer part information from a shape has proven to be useful in a variety of other applications, e.g., segmentation [Sha08], skeleton extraction and animation [dGGV08, SSCO08], and mesh editing [GSP*07, KJS07]. However, no surface metric employed

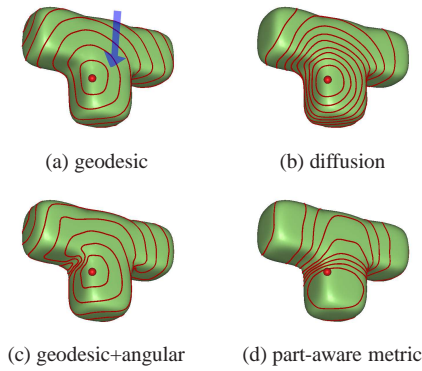


Figure 1: Iso-contours for four metrics: 10 contours sampled up to maximum distance from source (red dot). (a)-(b) Geodesic and diffusion [dGGV08] distances are insensitive to curvature and the perceived part boundary over the flat region indicated by the arrow. (c) Angular distance [KT03] senses concavity, but not part boundary over flat region. (d) Part-aware metric senses the whole part boundary.

so far explicitly encodes part information. Figure 1 presents a comparison between some previous surface metrics and ours, on a synthetic “T” shape. The iso-contour plots clearly show that only our metric fully senses the boundary of the bottom part of the shape.

Our part-aware metric is able to improve the performance of shape analysis algorithms which rely on a surface metric; its advantages are shown via several example applications.

1.2. Overview of approach

As the concept of a “part” is not well-defined, the definition of a part-aware metric is challenging. Popular techniques to capture parts are based on the *minima rule* [HR84], which induces part boundaries along *concave creases*. A standard realization of the minima rule is based on measuring surface properties, such as curvature or dihedral angles. However, these measures are limited by their local surface nature. In particular, local surface concavities do not completely characterize part boundaries, as exemplified by the flat region (indicated by the arrow) shown in Figure 1(a): we see no surface concavity over the perceived part boundary.

Our key observation is that although separation between parts inevitably involves concavity, to capture it one is not restricted to surface measurement. The inherent limitation of surface measurement can be addressed from a more global and volumetric view for part analysis. This leads to the realization that shape concavity can be manifested inside the shape’s volume via occlusion or visibility considerations.

Specifically, we examine a 3D object from within its enclosed volume and measure the visibility of the object surface from appropriately chosen reference points in the vol-

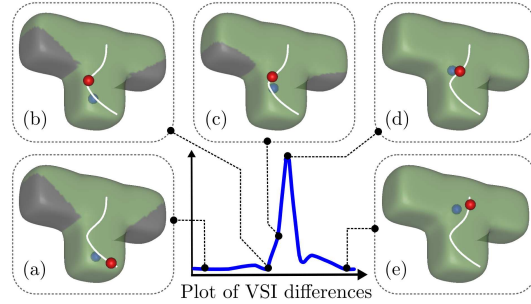


Figure 2: Plot (blue curve) of VSI differences as a point (red dot) moves along the white curve on the surface. Visible regions at reference points (blue dots) are painted in green. (a)-(b): Gradual change of visible region and VSI (difference ≈ 0) while in the bottom part. (c)-(d): Large change near part boundary. (d)-(e): Gradual change while in top part.

ume. Unlike a “part”, a visible region is clearly defined. Since parts are generally convex, the visible regions from within a part remain stable or change only gradually while moving inside the part. In contrast, when moving across a part boundary to another part, the visible region undergoes large change due to concave regions separating the two parts. To quantify the visible regions, we define a surface attribute, the *Volumetric Shape Image* (VSI), by linking surface elements to inner reference points and measuring surface visibility at those points. In Figure 2, we show how visible regions and VSIs change for the “T” shape along a path on the surface. Employing VSI as a signature to measure distances, we are able to build a part-aware surface metric.

2. Related work

One can find many uses of surface metrics in geometry processing. Gatzke et al. [GGGZ05] rely on geodesic distances to construct the curvature map signature for shape matching. Lai et al. [LZH*07] combine geodesic and isophotic distances into a feature-sensitive metric for sampling, remeshing, and multi-scale feature selection. In shape retrieval, different metrics have been used to derive global shape descriptors, such as shape distribution [FMK*03] and average geodesic field [GSCO07].

A classical application of surface metrics is mesh segmentation [Sha08]. Both region growing [PKA03] and clustering, e.g., fuzzy [KT03] or spectral clustering [LZ04], rely on a surface metric. Spectral clustering is representative of a class of algorithms where intrinsic metrics are used to set up an affinity or Laplacian matrix. Eigenvectors of the matrix then transform the original shape into particular embeddings suitable for the application at hand. Examples include bending-invariant signature for shape retrieval [EK03], spectral shape correspondence [JZvK07], and intrinsic symmetry detection [OSG08].

Recently, de Goes et al. [dGGV08] adopt diffusion distance to segment articulated bodies. The diffusion distance considers more global information than the geodesic and isophotic distances. Modeled using random walk, it measures the connectedness between points over the surface, making itself suitable for segmentation. However, it is still purely surface-based and does not account for volumetric information explicitly. The latter goal is partially accomplished by the local reach [DZM08] and the *shape diameter function* (SDF) [GSCO07, SSCO08]; both are scalar attributes. However, unlike VSI, local reach reflects only the *thickness* of local volume and SDF does not capture the general volumetric context either. An existing volume-oriented surface metric is inner distance [LJ07], which measures the shortest distance between two surface points inside the volume; it is pose-invariant but not part-aware.

Any work on shape segmentation has to face the question of how to define a part or part boundary. Besides requiring approximate convexity of parts [LA07], one of the best known rules for part analysis is the minima rule [HR84]; it characterizes part boundaries as negative curvature minima. By definition, the minima rule is typically realized via curvature measurements, which have so far been predominantly surface-based [KT03, LZ04, LLS*05, PKA03]. Volumetric considerations for part analysis on the other hand, e.g., the SDF [SSCO08], appear to be more closely linked to skeletal shape representations [CM07, SP08].

Shape parts can be identified from a 1-D skeleton based on its branching structures and local radii [LTTH01, ATC*08]. However, defining and computing such a skeleton are not easy tasks [CM07]. Also, the skeleton and radius function typically only induce a rough partition and the final cut still needs to be computed on the surface [ATC*08].

3. Geodesic and angular distances

Our part-aware metric combines three measures: geodesic, angular, and VSI distances. As we are interested in a surface metric, the input shape is assumed to be a manifold triangle mesh. However, our algorithm tolerates meshes with boundaries, as long as their volume is reasonably well-defined.

We first describe the angular and approximate geodesic distances, following the treatment of Katz and Tal [KT03]. In the following, we develop the concept by making use of the dual graph of a mesh, whose nodes are the mesh faces. Note that we can also work on the mesh primal graph similarly, when an application needs to deal with vertices.

Denote by \tilde{G} the dual graph and \tilde{e}_{ij} the edge incident to nodes p_i and p_j , which represent the centroids of face f_i and f_j , respectively. Let p_c be the midpoint of \tilde{e}_{ij} . The edge weight of \tilde{e}_{ij} is given by $\|p_i - p_c\| + \|p_j - p_c\|$. We call the resulting weighted graph the geodesic graph \tilde{G}_g . The geodesic distance between two faces is approximated by the graph distance between their corresponding nodes in \tilde{G}_g .




Figure 3: Distance field plots on a synthetic “snake” model, where red dots indicate sources. In (b), the blue curve shows how the shortest path between the two parts can bypass the concave region. This is the leakage problem and it makes angular distance “part-unaware”, in contrast to VSI distance.

Angular distances can be obtained similarly by defining an angular graph \tilde{G}_a . Given two adjacent faces f_i and f_j with normals n_i and n_j , respectively, the weight of \tilde{e}_{ij} is given by $\cos(n_i \cdot n_j)/\pi$, if f_i and f_j form a concave dihedral angle, or 0 otherwise. Thus edges only assume weights in concave regions: the more concave a region is, the larger the weights, roughly in accordance with the minima rule [HR84].

However, the angular distance suffers from what we call the *leakage* problem, caused by defining edge weights using only local surface properties. Figure 3 shows several distance field plots on a synthetic “snake” model. In this figure and throughout the paper we color plot scalar fields by linearly interpolating the hue values in the HSL color space; blue corresponds to small field values. Indicative of the leakage problem, geodesic and angular distances produce almost indistinguishable results. In contrast, part-aware distances measured by VSI exhibit clear separation between parts. The leakage problem also causes the iso-contours, in Figure 1(c), to cross the flat part boundary without “realizing” it. More comparisons between angular and VSI distances can be found in Figure 7.

4. Part-aware metric

To better capture part information, we resort to another distance between faces that is defined via volumetric shape images (VSIs). To this end, we rely on a third graph, the VSI graph \tilde{G}_v . It is defined similarly to \tilde{G}_g and \tilde{G}_a , with the exception that its edge weights are derived using the VSIs, from a more global and volumetric point of view. Once \tilde{G}_v is obtained, we blend the three graphs together to produce the final metric; see Section 4.4.

4.1. Overview

The main observation is that visible regions as observed inside a shape’s volume provide strong hints for part transition;




Figure 4: Plots of visible regions as a point (red dot) travels inside a shape. There is a large change in the shape of the visible region as the point crosses the part boundary (b)-(c).

see Figure 4 as well as Figure 2 for example. While moving inside a shape part, the visible region should remain stable or only change gradually. We do not capture visible regions from the surface, which may offer significantly different perspectives. We first normalize the perspective of these surface points by mapping them to reference points inside the shape. Such a normalization also alleviates problems of computing visible regions from the surface due to surface noise. We have found mapping surface points onto the medial sheet to be a good choice; see Section 4.2 for details.

To measure differences between visible regions, merely using their volumes is insufficient. We need a more descriptive signature. There are many possible shape descriptors one can consider [PSF04] or a straightforward approach is to compute the actual visible regions and examine how much they overlap. However a constraint we face is that the visible region, generally a complex 3D polyhedron, is expensive to compute. We thus resort to a sampling approach, to roughly capture the shape of the visible region and arrive at a simple signature that is efficient to handle.

Figure 5 summarizes the steps to compute the VSI distance between two adjacent faces, f_i and f_j . In step 1, both faces are mapped to their reference points r_i and r_j . In step 2, the visible regions at r_i and r_j are sampled into their VSIs, S_i and S_j . In step 3, S_i and S_j are used to derive the distance between f_i and f_j in graph G_v .

4.2. Reference point construction

Give a point on the surface, its corresponding medial center is the center of the maximal sphere inside the shape and tangent to the surface at the point. To compute the reference point r_i for face f_i , which is the medial center corresponding to the centroid of f_i , we adopt the simple ray-shooting technique of Shapira et al. [SSCO08].

Given a face f_i with normal n_i , as shown in Figure 5 (step 1), we use a cone with apex at the centroid c_i of f_i and center its axis along $-n_i$. Multiple rays inside the cone, including one along $-n_i$, are cast into the shape. Denote by l_j one of the rays, and $\|l_j\|$ its length inside the mesh. The approximate diameter of the maximal inscribed sphere touching c_i

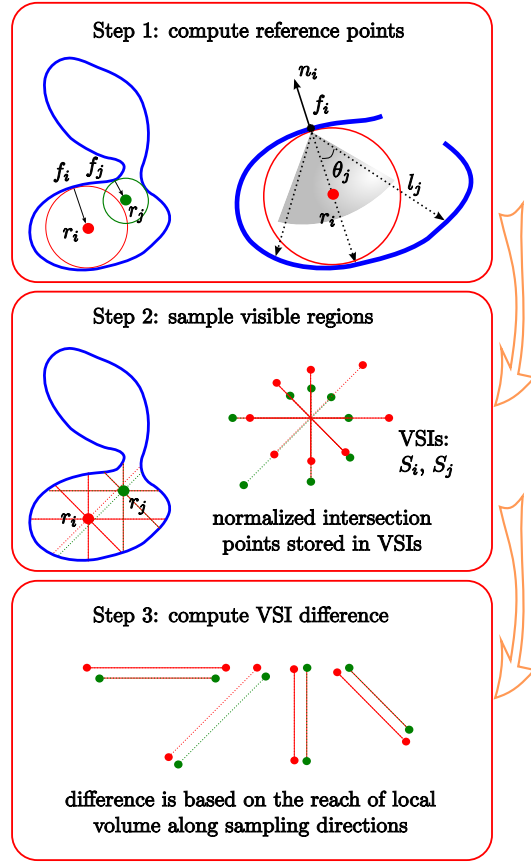


Figure 5: Three steps to compute VSIs and their difference.

at the surface is $d_i = \operatorname{argmin}_j \{ \|l_j\| / \cos \theta_j \}$, where θ_j is the angle between l_j and $-n_i$. The reference point r_i is the center of that sphere with radius $d_i/2$, and $r_i = c_i - 0.5 * d_i * n_i$.

The ray casting approach can be sensitive to noise. To alleviate this problem, we restrict the cone opening angle to 80° . In addition, the resulting reference points are smoothed using Laplacian smoothing over a neighborhood graph, which connects two reference points if their corresponding faces are adjacent on the mesh or the points themselves are close Euclidean neighbors to each other.

Note that we only compute the medial points roughly as the reference points need not to strictly reside on medial sheets. Computing rigorous medial sheets [CKM99, DZ02] of polygonal shapes is quite involved and is not necessary for our purpose. Figure 6 shows two reference point sets.

4.3. Volumetric shape images

We sample the visible regions as seen from the reference points and quantify their differences. From a reference point r_i , we send out m rays uniformly sampled on a Gaussian




Figure 6: Reference points (blue dots) computed for the horse (40K faces) and dragon (50K faces) models.

sphere and collect the intersection points, $s_i^{(j)}$'s, between the rays and the surface. Each intersection point is normalized with respect to its reference point, i.e., $s_i^{(j)} \leftarrow s_i^{(j)} - r_i$. The normalized intersection points are stored in a set $S_i = \{s_i^{(1)}, \dots, s_i^{(m)}\}$, called the volumetric shape image (VSI) of r_i (or of f_i). A larger m leads to better approximation of visible regions by VSIs, but at a higher complexity. We found that setting $m = 100$ achieves a good trade-off between accuracy and efficiency. Figure 5 (step 2) shows the VSIs, S_i (red) of r_i and S_j (green) of r_j . To avoid cluttering in the figure, only 8 sampling rays are drawn.

Since the sampling ray directions are fixed in a global coordinate system, two VSIs are naturally registered. We may define the difference between two VSIs simply as: $\text{diff}(S_i, S_j) = \frac{1}{m} \sum_{k=1}^m \|s_i^{(k)} - s_j^{(k)}\|^2$, which tends to possess large values near boundary regions due to dramatic visible region changes. Albeit meaningful, such a measure fails to accommodate accumulated translation errors. Consider a reference point moving within a large part. Though the slight translations between adjacent moves generate only small VSI differences, such small differences are easily accumulated up along the way. This can be observed from the significant difference between S_i and S_j , as illustrated in Figure 5 (step 2), despite that both r_i and r_j are within the same part.

We have observed that the reach of a local volume along a certain direction, measured from reference points, is in general a more stable measurement. Figure 5 (step 3) attests to this observation, in which we see that the reach of the volume of the lower part along the four directions are almost the same, even though r_i and r_j are distant. To exploit this property, we pick m sampling rays in opposite directions ($m/2$ ray pairs). Suppose that the two intersection points of a pair of rays are arranged consecutively in S_i , we define the difference between the VSIs as

$$\text{diff}(S_i, S_j) = \frac{1}{\sum_k w_k} \sum_{k=1}^{m/2} w_k \left(l_i^{(k)} - l_j^{(k)} \right)^2, \quad (1)$$

where $l^{(k)} = \|s^{(2k-1)} - s^{(2k)}\|$ is the local reach along the k -th direction. Weights w_k 's are designed to penalize outliers: as VSI is discrete, not all $(l_i^{(k)} - l_j^{(k)})^2$ may truthfully reflect the difference of the visible regions. To tackle this problem,




Figure 7: Angular (top row), SDF (middle row), and VSI (bottom row) distance fields on various models.

we fit a Gaussian to the distribution of such values and set

$$w_k = \begin{cases} e^{-(d_k - u)^2 / (2\sigma^2)} & \text{if } d_k < u + 2\sigma \\ 0 & \text{else } d_k \geq u + 2\sigma \end{cases}, \quad (2)$$

where $d_k = (l_i^{(k)} - l_j^{(k)})^2$, u is the mean, and σ is the standard deviation. Accordingly, values are weighted based on their frequencies. Outliers falling two deviations above the mean are ignored.

With edge \tilde{e}_{ij} in graph \tilde{G}_v weighted by $\text{diff}(S_i, S_j)$, the VSI distance is merely the shortest graph distance in \tilde{G}_v . Two such distance fields have already been shown in Figure 3(c, d). In Figure 7, we compare VSI distance to both the angular distance and the shape diameter function (SDF) [SSCO08]. Faces with similar SDF values tend to belong to the same part. Hence, it is natural to consider an SDF distance in a similar way, by defining the weight of \tilde{e}_{ij} as $|\text{sdf}(f_i) - \text{sdf}(f_j)|$, where $\text{sdf}(f_i)$ denotes the SDF value at face f_i . From these examples, we see that the VSI distance stays fairly stable within parts, and induces large jumps across part boundaries. It outperforms both the angular distance and the SDF distance.

Note that the VSIs can be computed for open meshes as well, so long as the mesh volume is reasonably well-defined. This is achieved by simply ignoring rays that do not intersect the surface. The hand model in Figure 7 is an open mesh.

To speed up the intersection tests required by computing reference points and VSIs, we voxelize the space. When a

input mesh is dense, we further speed up the processing by computing VSIs for only a small set of reference points and interpolating for nearby ones. For example, computing VSIs for the dragon model (Figure 6) takes about 15 seconds.

4.4. Combining distances

With the geodesic graph \tilde{G}_g , angular graph \tilde{G}_a and VSI graph \tilde{G}_v , we build the combined graph \tilde{G}_c as

$$\tilde{G}_c = (\alpha \otimes \tilde{G}_g) \oplus (\beta \otimes \tilde{G}_a) \oplus (\gamma \otimes \tilde{G}_v), \quad (3)$$

where $\alpha + \beta + \gamma = 1$ and $0 \leq \alpha, \beta, \gamma \leq 1$. Operator \otimes multiplies the edge weights of a graph by a scalar, and \oplus blends two graphs by summing up the weights of corresponding edges. Our final part-aware metric is defined as the shortest graph distances between the nodes of \tilde{G}_c . Note that as the three graphs encode different types of information, their edge weights must be normalized first. To this end, we linearly normalize the edge weights of each graph to $[0, 1]$.

For the part-aware metric, the parameters, α , β and γ , control the relative importance of the three distances, and their setting is application-dependent. For example, in mesh segmentation, we are primarily interested in a shape’s part structure; therefore we should emphasize the VSI distance. For shape registration, however, the importance of the geodesic distance ought to be increased, as both invariance to stretching (merit of VSI distance) and invariance to pose (merit of geodesic distance) are desired. Note that the edge weights in \tilde{G}_g are biased towards 1 after the normalization, since the majority of the edges or the faces of a mesh typically tend to have similar sizes. However, this is not true for the edge weights in \tilde{G}_a and \tilde{G}_v . For this reason, α will typically be set to a small value in order for the other two distances to take their effect.

5. Applications

In this section, we apply the part-aware metric to several geometry processing applications to illustrate its effectiveness. Note that in each application, the weights, α, β, γ , of the part-aware metric are fixed. Also, our goal herein is not to design new algorithms, but to evaluate the effectiveness induced by our part-aware metric to existing algorithms.

5.1. Shape segmentation

Shape segmentation [Sha08] is an important geometry processing operation that has received a great deal of attention in recent years. In this experiment, we show how our part-aware metric helps improve the performance of certain segmentation algorithms [KT03, LZ04].

Consider fuzzy clustering [KT03] as an example. The success of the algorithm depends on *fuzzy* regions that cover intended cut boundaries, as a cut is only extracted from within fuzzy regions. Figure 8 compares the fuzzy regions (red)

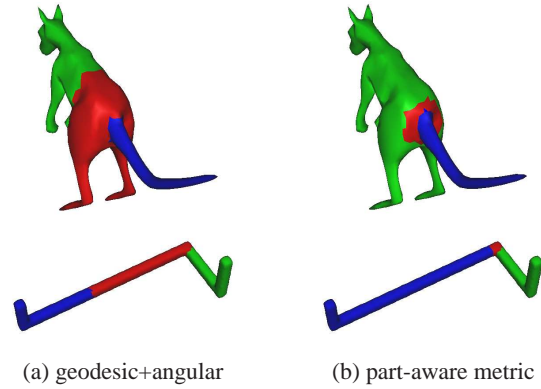


Figure 8: Fuzzy regions (red) produced by part-aware metric tend to enclose the intended part boundary more tightly.

computed on two models using the original distance metric from [KT03] and our part-aware one. In this experiment, we compute a smallest fuzzy region which covers any intended boundary. Thanks to its part-awareness, the fuzzy regions obtained using our metric are smaller and more constrained to the desired boundary regions. A smaller fuzzy region is preferred as it facilitates the subsequent graph cut operation.

Spectral clustering [Wei99] is a popular clustering approach and its success relies on the quality of the distance measure used. To test our metric in this aspect, we apply it to spectral mesh segmentation [LZ04]. Briefly, we derive pair-wise distances between mesh faces using our metric, and compute an affinity matrix from these distances. The eigen-structure of the affinity matrix is then utilized to build an Euclidean embedding – called a *spectral embedding* – of mesh faces, in which the standard *K*-means clustering algorithm is applied to obtain a segmentation result. Since our distance metric better reflects part structure in a shape, segmentation quality is expected to be improved.

Figure 9 shows several comparative segmentation results. For sub-plots (a), (b) and (c), one can refer to the models’ VSI distance fields in Figure 7. Pay attention to how the boundary regions, either between the small finger and the palm, between the horse’s front leg and the body, or between the dragon’s tail and the body, induce significant distance changes when VSI distance is used. Clearly, the part-aware metric helps improve the segmentation quality in those regions. For all these tests, we set $\alpha = 0.01$, $\beta = 0.195$, and $\gamma = 0.795$, to emphasize VSI distance. Geodesic distance is not completely discarded since it helps produce smoother part boundaries and prevents undesirably large parts from being formed as well.

5.2. Shape registration

Shape registration or correspondence [CH03, JZvK07] is another good example where a surface metric can become

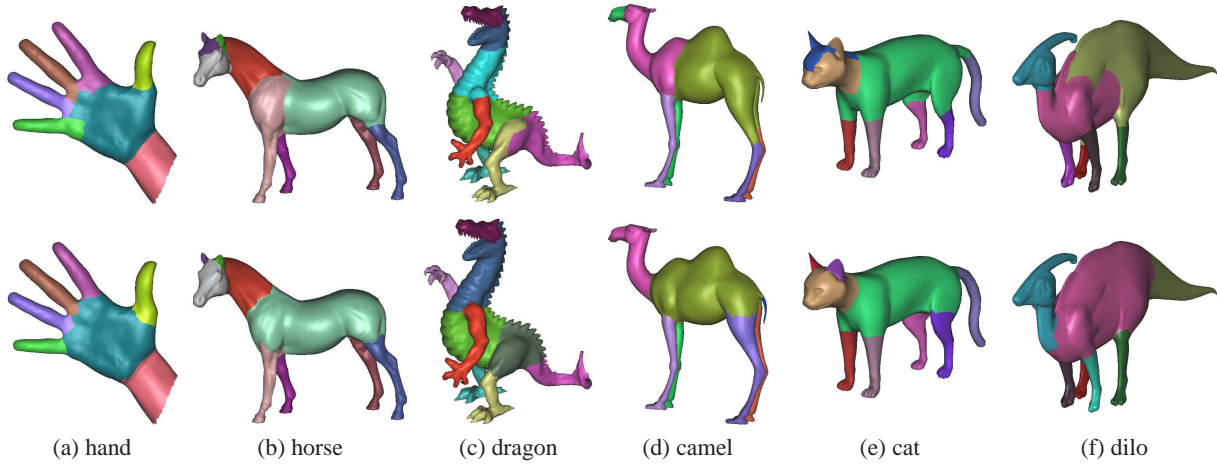


Figure 9: Segmentation results on various models. The top row shows results obtained using geodesic and angular distances as described in [LZ04]; while the bottom row shows results achieved using our part-aware metric.

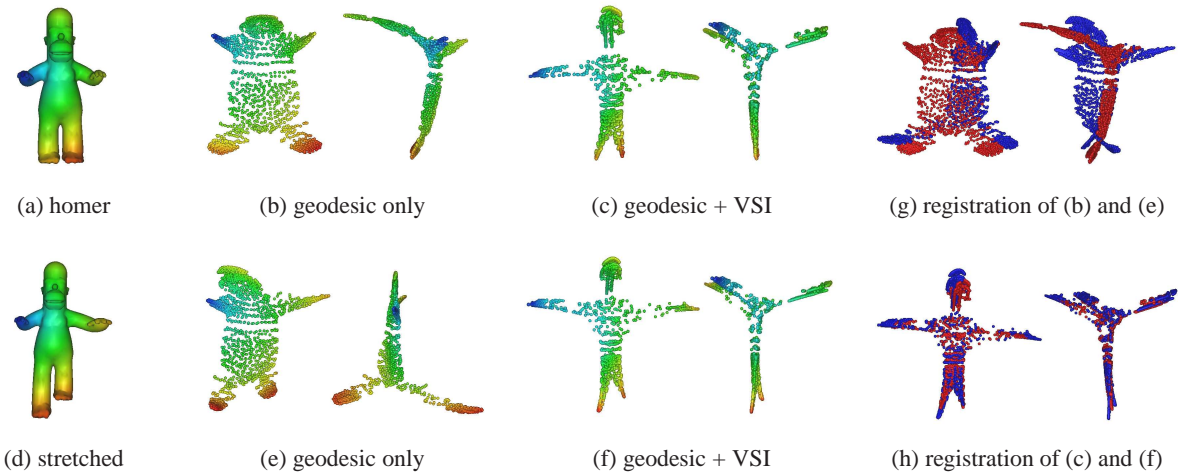


Figure 10: Spectral embeddings and registration of the original (a) and stretched (d) homer models. Plots (b), (c) and (e), (f) show, each from two different angles, the embeddings of the original and stretched homer with and without VSI distance. The coloring implies the mapping between original mesh vertices and their embeddings. We see the embeddings in (c) and (f) are much more similar to each other than those in (b) and (e) (especially for the feet), thanks to the stretch-invariance of VSI distances. This helps the subsequent registration task, whose results are shown in (g) and (h).

useful. For such a task, bending-invariance [EK03] has been successfully addressed, while stretch-invariance still remains challenging. We consider the spectral correspondence framework from [JZvK07] to investigate how the part-aware metric could potentially improve shape registration. We embed mesh vertices in Euclidean space and the registration is established by corresponding the embeddings using the iterative closest point algorithm described in [BM92].

Figure 10 shows the derived embeddings of two homer models. To focus on stretching, we purposely stretched the

left arm and the right leg of one homer. As stretching changes the geodesic distance between vertices, the resulting embeddings in (b) and (e) with geodesic distance are quite different. In contrast, the VSI distance is more stable against stretching. When it is combined with geodesic distance, the embeddings in (c) and (f) are made similar, facilitating the registration in sequel. We set $\alpha = 0.1$, $\gamma = 0.9$ and note that this setting places sufficient emphasis on geodesic distances as per the weight distribution in \tilde{G}_g (see Section 4.4).

In practice, we believe that both geodesic and VSI dis-

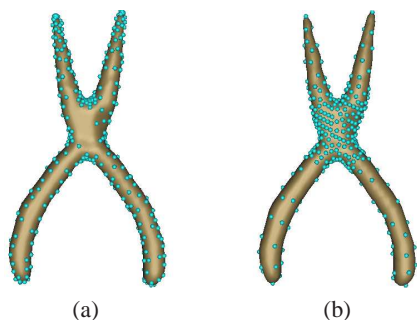


Figure 11: Feature-sensitive (a) vs. part-aware sampling (b). The sampling rate is 10%.

tances should be considered simultaneously for registration, since the former contributes to bending-invariance, while the latter contributes to stretch-invariance. Also, the geodesic distance imposes certain distances between neighboring vertices and keeps them adequately spread-out in the embedding. Their weights may be tuned to strike the balance between the two. It is worth investigating how their weights can be set automatically based on the input data.

5.3. Part-aware sampling

The max-min sampling scheme can be used for uniform sampling on surfaces. It iteratively places a sample that maximizes the minimum distance, measured by the geodesic metric, from all previously found samples. Beyond uniform sampling, it is often desirable to distribute samples based on certain feature criteria. In typical feature-sensitive remeshing, more samples are placed near high curvature regions, resulting in higher triangle density therein. This can be accomplished via max-min sampling, but with the isophotic metric [LZH*07]. Figure 11(a) shows a feature-sensitive sampling on the pliers model. When our part-aware metric is applied, the feature regions tend to be part boundaries, even if a boundary region is locally flat (different from feature-sensitive sampling). This is shown in Figure 11(b). In this test, we set $\alpha = 0.01$, $\beta = 0$, and $\gamma = 0.99$, similar to those for segmentation while ignoring angular distance, so that samples would not concentrate in concave regions.

A straightforward use of a part-aware surface tessellation is in neighborhood traversal: part-aware neighborhood can be traced simply by following topological distances in the mesh graph. Also, more granularity near part boundaries enables more refined deformation behavior over those non-rigid regions. In contrast, having large triangles which cut across part boundaries would be undesirable in this context.

5.4. Shape retrieval via part-aware shape distribution

We also study the potential benefits that our metric can offer for shape retrieval. Osada et al. [OFCD02] consider using

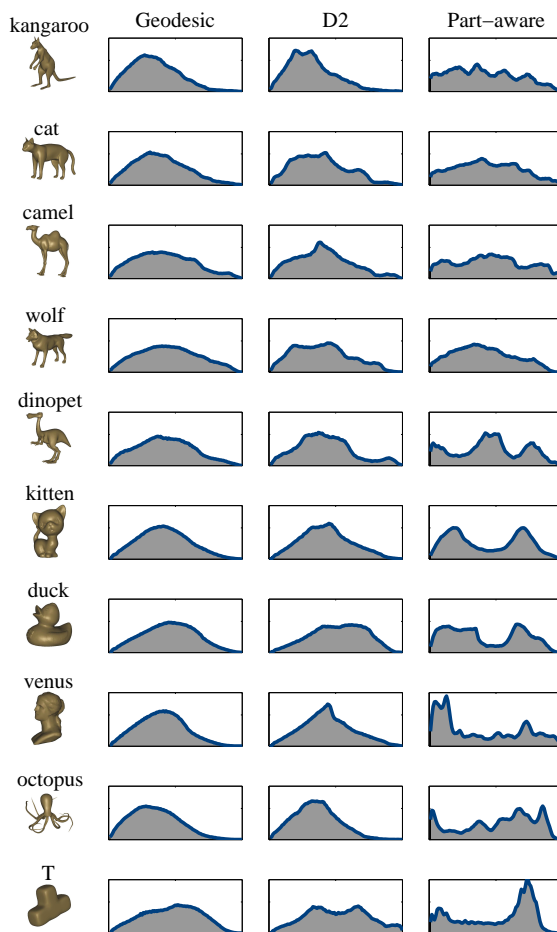


Figure 12: Ten models used for our shape retrieval test. The last three columns show the distance histograms of each model, obtained using geodesic, D2 and part-aware distances, respectively.

the probability distribution sampled from a shape function as the shape descriptor. It is suggested that the D2 shape function, namely the pair-wise Euclidean distances between surface points, serves as a good candidate. Geodesic distance shape function is investigated in [MS08]. Although it is invariant to bending, the probability distribution function of geodesic distances is not sufficiently discriminative in general. We plan to add more discriminating power to this descriptor, while preserving sufficient bending-invariance.

In our experiment, we select ten models shown in Figure 12. The first five of them (kangaroo, cat, camel, wolf and dinopet) are animals with similar part structures, but under different articulation and stretching. In retrieval, the ideal result should be that when using each of these five shapes to query, the other four come on top. We consider three shape




Figure 13: Retrieval results using geodesic, $D2$ and our part-aware distances. In each block, we run 6 queries and the query shapes sit in the first (shaded) row. In each column, the 4 most similar shapes are shown, and the non-optimally retrieved shapes are marked by squares.

functions: geodesic, $D2$, and our part-aware distances. For each sampled shape function, we compute its histogram, which serves as the shape descriptor. During retrieval, the χ -square distances between histograms are used to rank the matches. Figure 13 shows the retrieval results using the three shape functions. We see that only our part-aware metric is able to consistently achieve the ideal results for all the five

similar shapes. Note that in the last column, we query using the octopus simply for comparison purpose, in which the ideal ranking is hard to judge.

We set $\alpha = 0.02$, $\beta = 0$, and $\gamma = 0.98$, making $\gamma \gg \alpha$ to showcase our goal. From Figure 12, we observe that the histograms of geodesic and $D2$ distances stay relatively stable. Contrastingly, the part-aware metric produces more discriminating histograms. Meanwhile, due to the geodesic ingredient in the metric, bending invariance is also preserved to a certain extent; this is exemplified by the similarity among the first four histograms in the last column. This experiment highlights an advantage of our metric: by combining geodesic and VSI distances, it provides an effective compromise between bending-invariance and discriminating capability. However, this is a rather preliminary test and we plan to expand the database and investigate more in this issue.

6. Conclusion and future work

A metric defined on mesh surfaces is of fundamental interest when dealing with geometry. Well-known metrics include geodesic, isophotic, and diffusion distances; other extensions to these metrics also exist. Part is an important notion for shape analysis and has been studied in a variety of domains, ranging from 2D to 3D, from psychology to geometry processing. Although isophotic and diffusion distances have been successfully applied to mesh segmentation, the distances themselves do not explicitly take part information into consideration.

In this paper, we develop a surface metric specifically for encoding part information. Our key observation is that the visible region of a point inside the shape volume is sensitive to part boundaries. Such a property is realized into an effective part-aware metric, via sampling and volumetric shape image. The effectiveness of our metric is testified by several concrete experiments.

In future work, we plan to explore more applications that can exploit our part-aware metric, and study intelligent and application-dependent ways to tune the weights for the three distances. We also plan to further investigate into shape retrieval, expand the test database, and compare with other metric-based shape signatures [OFCD02, BCG08]. Finally, we wish to take the volumetric view for part analysis [AMSF08] further. Since the current metric still evaluates distances between points along surface, it does not recognize parts consisting of disconnected surface regions.

Acknowledgments: We would like to thank the anonymous reviewers for their helpful comments, Sagi Katz and Ayellet Tal for the help on the fuzzy clustering test. This work was supported in part by grants from an NSERC Grant (No. 611370), a MITACS Grant (No. 320026), the Israel Ministry of Science and Education Tashtiot Grant “Advanced Methods For Reconstruction Of Three-Dimensional Synthetic Object”, and the Israel Science Foundation.

References

- [AMSF08] ATTENE M., MORTARA M., SPAGNUOLO M., FALCIDIANO B.: Hierarchical convex approximation of 3d shapes for fast region selection. *Computer Graphics Forum (Special Issue of Symposium on Geometry Processing)* 27, 5 (2008), 1323–1333.
- [ATC*08] AU O. K.-C., TAI C.-L., CHU H.-K., COHEN-OR D., LEE T.-Y.: Skeleton extraction by mesh contraction. *ACM Trans. on Graphics* 27, 3 (2008), 1–10.
- [BCG08] BEN-CHEN M., GOTSMAN C.: Characterizing shape using conformal factors. In *Proc. of Eurographics Workshop on Shape Retrieval* (2008).
- [BM92] BESL P. J., MCKAY N. D.: A method for registration of 3-d shapes. *IEEE Trans. Pattern Anal. Mach. Intell.* 14, 2 (1992), 239–256.
- [Car76] CARMO M. P. D.: *Differential Geometry of Curves and Surfaces*. Prentice Hall, 1976.
- [CH03] CARCASSONI M., HANCOCK E. R.: Correspondence matching with modal clusters. *IEEE Trans. Pattern Anal. Mach. Intell.* 25, 12 (2003), 1609–1615.
- [CKM99] CULVER T., KEYSER J., MANOCHA D.: Accurate computation of the medial axis of a polyhedron. In *SMA '99: Proc. of the fifth ACM symposium on Solid modeling and applications* (1999), pp. 179–190.
- [CM07] CORNEA N. D., MIN P.: Curve-skeleton properties, applications, and algorithms. *IEEE Trans. Vis. & Comp. Graphics* 13, 3 (2007), 530–548.
- [dGGV08] DE GOES F., GOLDENSTEIN S., VELHO L.: A hierarchical segmentation of articulated bodies. *Computer Graphics Forum (Special Issue of Symposium on Geometry Processing)* 27, 5 (2008), 1349–1356.
- [DZ02] DEY T. K., ZHAO W.: Approximate medial axis as a voronoi subcomplex. In *SMA '02: Proceedings of the seventh ACM symposium on Solid Modeling and Applications* (2002), pp. 356–366.
- [DZM08] DYER R., ZHANG H., MÖLLER T.: Surface sampling and the intrinsic Voronoi diagram. *Computer Graphics Forum (Special Issue of Symposium on Geometry Processing 2008)* 27, 5 (2008), 1393–1402.
- [EK03] ELAD A., KIMMEL R.: On bending invariant signatures for surfaces. *IEEE Trans. Pat. Ana. & Mach. Int.* 25, 10 (2003), 1285–1295.
- [FMK*03] FUNKHOUSER F., MIN P., KAZHDAN M., CHEN J., HALDERMAN A., DOBKIN D., JACOBS D.: A search engine for 3D models. *ACM Trans. on Graphics* 22, 1 (2003), 83–105.
- [GGGZ05] GATZKE T., GRIMM C., GARLAND M., ZELINKA S.: Curvature maps for local shape comparison. In *Proc. IEEE Int. Conf. on Shape Modeling and Applications* (2005).
- [GSCO07] GAL R., SHAMIR A., COHEN-OR D.: Pose-oblivious shape signature. *IEEE Trans. Vis. & Comp. Graphics* 13, 2 (2007), 261–271.
- [GSP*07] GAL R., SORKINE O., POPA T., SHEFFER A., COHEN-OR D.: 3D collage: Expressive non-realistic modeling. In *Proc. of Symposium on Non-photorealistic Animation and Rendering* (2007), pp. 7–14.
- [Heb49] HEBB D. O.: *The Organization of Behavior*. John Wiley, 1949.
- [HR84] HOFFMAN D. D., RICHARDS W. A.: Parts of recognition. *Cognition* (1984), 65–96.
- [JZvK07] JAIN V., ZHANG H., VAN KAICK O.: Non-rigid spectral correspondence of triangle meshes. *Proc. IEEE Int. Conf. on Shape Modeling and Applications* 13, 1 (2007), 101–124.
- [KJS07] KREAVOY V., JULIUS D., SHEFFER A.: Model composition from interchangeable components. In *Proc. of Pacific Graphics* (2007), pp. 129–138.
- [KT03] KATZ S., TAL A.: Hierarchical mesh decomposition using fuzzy clustering and cuts. *ACM Trans. on Graphics* 22, 3 (2003), 954–961.
- [LA07] LIEN J.-M., AMATO N. M.: Approximate convex decomposition of polyhedra. In *Proc. of ACM Symposium on Solid and Physical Modeling* (2007), pp. 121–131.
- [LJ07] LING H., JACOBS D. W.: Shape classification using the inner-distance. *IEEE Trans. Pat. Ana. & Mach. Int.* 29, 2 (2007), 286–299.
- [LLS*05] LEE Y., LEE S., SHAMIR A., COHEN-OR D., SEIDEL H.-P.: Mesh scissoring with minima rule and part salience. *Comput. Aided Geom. Des.* 22, 5 (2005), 444–465.
- [LTTH01] LI X., TOON T., TAN T., HUANG Z.: Decomposing polygon meshes for interactive applications. In *Proc. of Symposium on Interactive 3D Graphics* (2001), pp. 35–42.
- [LZ04] LIU R., ZHANG H.: Segmentation of 3D meshes through spectral clustering. In *Proc. of Pacific Graphics* (2004), pp. 298–305.
- [LZH*07] LAI Y.-K., ZHOU Q.-Y., HU S.-M., WALLNER J., POTTMANN H.: Robust feature classification and editing. *IEEE Trans. Vis. & Comp. Graphics* 13, 1 (2007), 34–45.
- [MS08] MAHMOUDI M., SAPIRO G.: Three-dimensional point cloud recognition via distributions of geometric distances. In *Proc. IEEE Conf. on Comp. Vis. and Pat. Rec.* (2008), pp. 1–8.
- [OFCD02] OSADA R., FUNKHOUSER T., CHAZELLE B., DOBKIN D.: Shape distributions. *ACM Trans. on Graphics* 21 (2002), 807–832.
- [OSG08] OVSJANIKOV M., SUN J., GUIBAS L.: Global intrinsic symmetries of shapes. *Computer Graphics Forum (Special Issue of Symposium on Geometry Processing)* 27, 5 (2008), 1341–1348.
- [PHHH04] POTTMANN H., HOFER M., HAIDER C., HANBURY A.: The isophotic metric and its application to feature sensitive morphology on surfaces. In *Proc. Euro. Conf. on Comp. Vis.* (2004), pp. 560–572.
- [PKA03] PAGE D. L., KOSCHAN A., ABIDI M. A.: Perception-based 3d triangle mesh segmentation using fast marching watersheds. In *Proc. IEEE Conf. on Comp. Vis. and Pat. Rec.* (2003), pp. 27–32.
- [PSF04] P. SHILANE P. MIN M. K., FUNKHOUSER T.: The princeton shape benchmark. In *Proc. of Shape Modelling International* (2004).
- [Sha08] SHAMIR A.: A survey on mesh segmentation techniques. *Computer Graphics Forum* 27, 6 (2008), 1539–1556.
- [SJC08] SEONG J.-K., JEONG W.-K., COHEN E.: Anisotropic geodesic distance computation for parametric surfaces. In *Proc. IEEE Int. Conf. on Shape Modeling and Applications* (2008), pp. 179–186.
- [SP08] SIDDIQI K., PIZER S. E.: *Medial Representations*. Computational Imaging and Vision, Vol. 37. Springer, 2008.
- [SSCO08] SHAPIRA L., SHAMIR A., COHEN-OR D.: Consistent mesh partitioning and skeletonization using the shape diameter function. *Visual Computer* 24, 4 (2008), 249–259.
- [Wei99] WEISS Y.: Segmentation using eigenvectors: A unifying view. In *Proc. of International Conference on Computer Vision* (1999), pp. 975–982.