Noname manuscript No.
(will be inserted by the editor)

Comparing local shape descriptors

Paul Heider, Alain Pierre-Pierre, Ruosi Li, Rolf Mueller, and Cindy Grimm

Received: date / Accepted: date

Abstract Local shape descriptors can be used for a variety of tasks, from registration to comparison to shape analysis and retrieval. There have been a variety of local shape descriptors developed for these tasks, which have been evaluated in isolation or in pairs, but not against each other. We provide a survey of existing descriptors and a framework for comparing them. We perform a detailed evaluation of the descriptors using real data sets from a variety of sources. We first evaluate how stable these metrics are under changes in mesh resolution, noise, and smoothing. We then analyze the discriminatory ability of the descriptors for the task of shape matching. Finally, we compare the descriptors on a shape classification task. Our conclusion is that sampling the normal distribution and the mean curvature, using 25 samples, and reducing this data to 5-10 samples via Principal Components Analysis, provides robustness to noise and the best shape discrimination results. For shape classification, mean curvature sampled at the vertex or averaged, and the more global Shape Diameter Function, performed the best.

Keywords Spin images · shape descriptors · mean curvature · Gaussian curvature · feature detection

1 Introduction

Local shape descriptors are used in 3D shape matching to find unique points on a surface, to match up points on different models, and to define feature vectors that can be used to classify shapes. They can also be used to speed up searches by reducing the model to a small number of features which are easily compared. Given the plethora of descriptors out there, what works best? We evaluate that question in two contexts: shape feature matching and shape classification. Shape feature matching comes down to the following question: given sets of "similar" points, how well does the descriptor do at clustering similar points while distinguishing between points in different sets? Shape classification looks at the same problem, but for the *distribution* of shape descriptor values across an entire shape.

We first survey existing local shape descriptors, grouping them by type (Section 2). Next, we provide a framework for comparing these descriptors against each other. In the process of doing this, we develop several new variations of existing descriptors (Section 3). In specific, we create a rotation-independent version of Point or Spin descriptors [3] and their variants.

To evaluate the descriptors we perform three studies. The first study (Section 4) is a straightforward analysis of behavior under changing mesh quality, noise, and smoothing. The second study (Section 5) uses handpicked similar feature points to determine which descriptors are both sensitive (can determine if features are the same) and specific (can distinguish one feature from another). The third study looks at how the shape descriptors perform for a shape classification task (Section 6). We also provide a correlation analysis on the descriptors (Section 7).

We compare features for both more standard manmade objects [18] and biological data sets [17], and extend the analysis in these previous studies to the classification task.

C. Grimm
Washington University in St. Louis
Tel.: +314-935-4576
E-mail: cmg@wustl.edu
Rolf Mueller
Virginia Tech
E-mail: rolf.mueller@vt.edu

Our conclusion is that sampling mean curvature or the normal distribution at roughly 25 samples per local neighborhood, followed by Principal Components Analysis to reduce the data to 7-10 numbers, are the two best descriptors in terms of robustness and discrimination power at the feature level. At the object level (classification task) the Shape Diameter Function or mean curvature, sampled locally, are the two best choices.

Contributions: 1) A survey of existing local descriptors. 2) A systematic evaluation of a variety of shape descriptors on different data sets. 3) Rotationally-invariant modification of Point signatures. 4) Normalized comparison functions that allow for direct comparison of all descriptors.

2 Local descriptors

A good local descriptor is one that is invariant to "unimportant" geometric changes, typically rotation and translation, sometimes scaling, and sometimes bending (such as posing an articulated character). It should take into account the local shape of the surface surrounding a given point. It should also be robust to noise and sampling errors: Geometric noise (vertices moving), Mesh topology noise (the mesh connectivity changes), and Global topology noise (the creation of handles and tunnels). Local descriptors should also have a meaningful comparison function, one that scales roughly linearly with perceived shape change and is robust to noise.

To handle invariance, most descriptors measure geometric properties that are invariant to translation and rotation, such as curvature, length, volume, and angle. Scale invariance requires a relative measure — for example, the length of a curve over a radius — or scaling the object to a default size.

There are, broadly, two ways to determine a local neighborhood around a point. The first is to use Euclidean distance, for instance, all of the (connected) surface that is contained within a sphere of a given radius. The second is to use Geodesic distance, for instance, by walking out a set distance on the surface. Empirically, we have determined that for small-scale features the two are qualitatively similar. The difference really only matters when considering large-scale features such as entire limbs. At that scale, pose-invariance (a bent limb should be the "same" as a straight one) can only be achieved by considering geodesic distances. We divide existing local descriptors into three classes. The first two only look at local data, the third at global data. The two local data classes split on whether they sample a metric locally or fit a model to the local neighborhood.

2.1 Ring-based descriptors (sample metric locally)

Blowing bubbles [24] intersects the surface with a set of concentric spheres and extracts information about the surface in two steps. First, they simply count the number of closed contours, ignoring curves that are far away (1, 2, or more than 2). They further classify the contours using the length of the contours and a local concavity measure that determines if the curve centroid is above or below the point. Altogether, these measures can be used to classify the surface into eleven different groups. These descriptors (length of curve, centroid) are generalized to an integrative framework in [26].

Geodesic fans [14] sample a metric (such as curvature) on the mesh using concentric geodesic rings instead of spheres. Geodesic fans re-sample the metric into evenly-spaced samples in the radial and angular direction. To compare two fans, every possible rotation is tried, and the one with the minimal error is kept. They also introduce a 1D "curve", where the values around a ring are collapsed into a single number (average, minimum, and maximum values). This eliminates the need to try all possible rotations but does result in a loss of information. A modified version of this, which combined curvature along geodesics with normal variation, was recently used to do polyp detection [25].

Splash descriptor Stein and Medioni [31] sample the normal at regularly-spaced intervals. The normals are mapped to a spherical coordinate system using the normal n_0 at the point and a tangent vector t, with t chosen by finding the point P_i which has the maximum value of $\sqrt{< n_0, n_i >^2 + < t, n_i >^2}$. They then map this 3D curve to 2D by computing the curvature and torsion along it. It is not clear how necessary (or useful) this last step is; it was primarily motivated by compatibility with the rest of their system.

Point descriptor Chua and Jarvis [3] used a similar sampling, but recorded the distance from the contour to a plane fit to the contour and passing through the point. Yamany and Farag [33,34] proposed a modified version of this where they took a line through the point on the ring. They then stored the line length and angle with respect to the normal. They convert this to a 2D image by using the first angle and length as the axes, and the normal angle as the image's intensity value. They used all points in the mesh for this descriptor; by weighting the points by their distance to the center point [1,20] this descriptor can be localized.

2.2 Expanding descriptors (fit model to region)

Fitting polynomials: Cipriano et. al [4] grow their disks using geodesics. Each disk is then treated as a

height field over the tangent plane. They fit a quadratic polynomial to the height field, weighting the points by their surface area and how close to the center they are. Because they project their points directly to the tangent plane they may get folding; this is mitigated by a post-processing step which detects these cases as outliers. Finally, they define two descriptors. The first is the curvature of the fitted quadratic. The second descriptor treats the height field as an intensity image, and measures the anisotropy of the image. This can be computed directly from the polynomial coefficients, but is somewhat less specific than the curvature measures.

Mesh saliency: Mesh saliency [21] uses the concept of center-surround from perception to measure how the center of a disk differs from a disk twice as big. Essentially, they sum up a metric (usually mean curvature) using a Gaussian weighted sum centered at the point. They then repeat this sum using a kernel twice the size. The value for the disk is then the ratio of the two sums, normalized.

Volume and surface area: Instead of measuring derivative information (eg, curvature) on the surface, an alternative is to measure integrative information, such as volume [7] or surface area [26]. These measures, being integrated, are nominally more stable in the presence of noise than derivative measures. To extract more information about the local patch it is possible to apply Principal Components Analysis [5,6]. Pottman et al [26] provide a nice summary and comparison of these integrative, invariant geometric measures.

2.3 Global and iterative operator descriptors

Smoothing: As a mesh is smoothed, vertices change their positions. Vertices in high-curvature regions tend to move more than in low-curvature ones. Essentially, record the distance moved by each vertex for each smoothing iteration, for some number of iterations. This descriptor is used by Li and Guskov [23] to find interesting points and by the brain mapping community [11] to identify sulcal folds on cortical surfaces.

Shape Diameter Function: Essentially, measure how far away the surface is in the *opposite* direction of the surface normal [30]. This gives a local measure of the diameter of the surface at that point.

Geodesics and diffusion: The techniques in this class do not directly measure the geodesics, but instead measure a diffusion process flowing along the geodesics. The Laplace-Beltrami operator [28] is used to compute the diffusion because it essentially records the mesh connectivity.

Bronstein et. al [2] use diffusion geometry (how long does it take, on average, to walk from one point to another on the surface using a random walk?) to compare two surfaces. To turn this into a local descriptor, measure the average probability of walking from the point P to all the neighboring points at time t, for increasing values of t [15]. An alternative is to measure the average distance to all of the points in the neighborhood [12]. Sun et. all [32] perform a heat diffusion operation on the surface, and track the accumulated heat at the point over time.

Comparing two diffusion signatures requires some care because the values change rapidly for small time values, then smooth out. For this reason, Sun et. all [32] sample time using a logarithmic scale and normalize by dividing by the area under the curve.

3 Shape descriptors

The previous section surveyed a variety of existing methods for computing local shape descriptors. In order to compare them, we have developed a unified framework that attempts to preserve the original metrics and sampling strategies while producing descriptors that can be meaningfully compared to each other. To address the comparison problem, we introduce a normalization step that converts the raw descriptor to the range $[0, 1]$, accounting for non-linearities in the metric as best as possible. Additionally, several descriptors have inherent in them either picking a canonical tangent direction [3, 34] or comparing all possible rotations [35] to produce a rotation-invariant descriptor. We adopt a different approach, which is to use a histogram. This captures the distribution of metric values, but does result in a loss of ordering information.

We formally define a local shape descriptor as a mapping from a small subset of the surface centered at a point P to a vector d of numbers. These vectors can then be compared to determine how similar two points are, or used in a classifier. By varying the amount of surface used, the descriptors can pick out smaller or larger features. *Metric* here refers to a value, such as curvature, calculated at every point on the surface. The general algorithm for computing a descriptor is as follows:

1a, Ring and expanding descriptors: For each point P, sample the metric at R concentric rings or patches spaced a distance r apart. This produces a vector of dimension RS, where S is the number of metric samples per ring.

1b, Iterative operator descriptors: Apply the iterative operator. Record the values at R iterations, where each iteration corresponds (approximately) to the operator expanding a distance r. This produces a vector of

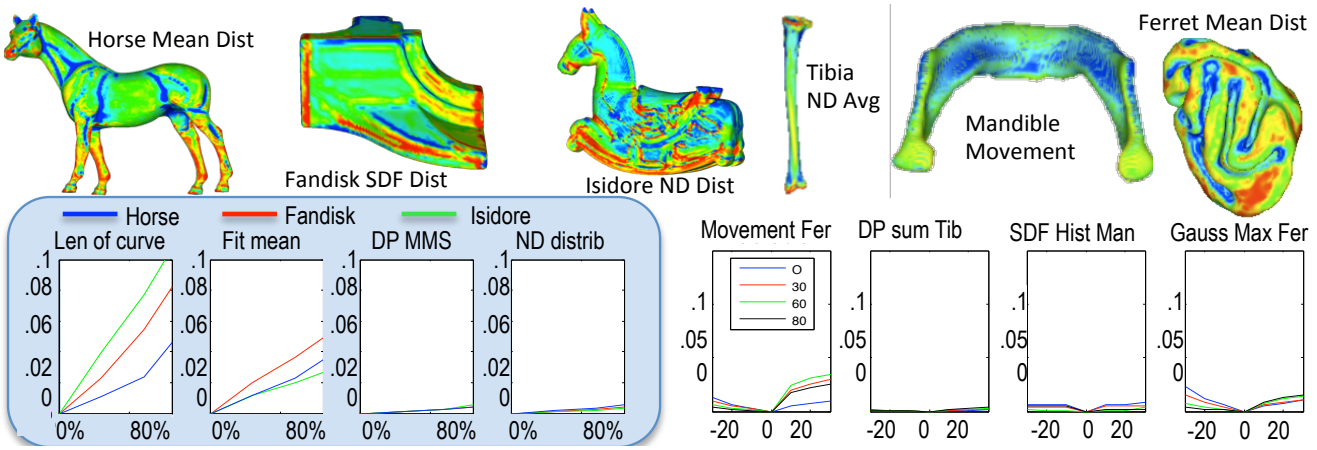

Fig. 1 Top: The six meshes used in the stability study. Bottom left: Mesh reduction. x-axis is the mesh reduction amount, y-axis is the difference in the descriptors, averaged across the surface (|| reduced - original ||). Bottom right: Noise and smoothness. x-axis goes from noisy to smooth (see Sec 4). The meshes were compared to their original mesh of the same resolution. Curve color indicates level of mesh reduction.

dimension RS for each point P, where S is the dimension of the output of the operator at point P (usually $S = 1$).

2, Reduction: Apply Principle Component Analysis (PCA) or Multi-dimensional Scaling (MDS) as appropriate to all of the vectors, and re-project the vectors onto this coordinate system.

3, Normalize: Normalize the distribution of values in each dimension, using a non-linear mapping as necessary, so the values lie between zero and one and are (relatively) evenly distributed across the range.

Section 3.1 defines the specific descriptors used in our study. Section 3.2 discusses the coordinate system change and normalization step.

Sampling the surface: We experimented with three different methods for sampling the surface: 1) Intersecting spheres [24] (Euclidean), 2) Growing disks, and 3) Exponential maps [29] (Geodesic). Empirically, they all produce qualitatively similar results. For this paper, we use 1) with $r = 0.0375B/R$, where B is the diagonal of the bounding box of the surface and $R = 5$. In general, for surfaces with small, thin structures or substantial noise we recommend using the intersecting spheres approach. For surfaces without these features we recommend the Exponential map approach because it is faster and does not require a second parameterization step. We resample metric data evenly along the rings, at a spacing of $2\pi r/20$. For all three methods the computation time is dependent on the number of surface points and the average number of points k that lie within the maximum radius r for each point. Ideally, k should be between 25 and 50 for $n \approx 5$. If the surface sampling is denser than this, significant computation time is wasted because the distance between surface samples is much smaller than the distance between the rings.

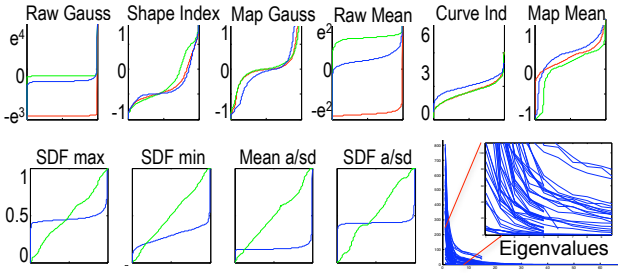

Fig. 2 Top row: Curvature values, sorted in increasing order (raw, index, mapped). Red - Mandible, Blue - Tibia, Green - Ferret. Bottom row: Applying mapping (Section 3.2). Blue curve is original data, sorted and linearly scaled to the range [0,1]. Green curve is normalized data. Bottom right: Eigen values of all descriptors.

3.1 Specific descriptors

An implementation of these descriptors is available at `https://sourceforge.net/projects/meshprocessing/`.

3.1.1 Ring-based descriptors:

[DP] Distance to plane: Fit a plane to the ring then calculate the signed distance of each point to the plane.
[ND] Normal distribution (2 values): 1 - Fit a plane to the vertex, a point on the ring, and the vertex's normal. Project the ring point's normal onto that plane, and find the angle with the normal. 2 - Fit a plane to two consecutive points on the ring and the first point's normal. Project the second point's normal onto this plane. This sampling is rotation-independent.

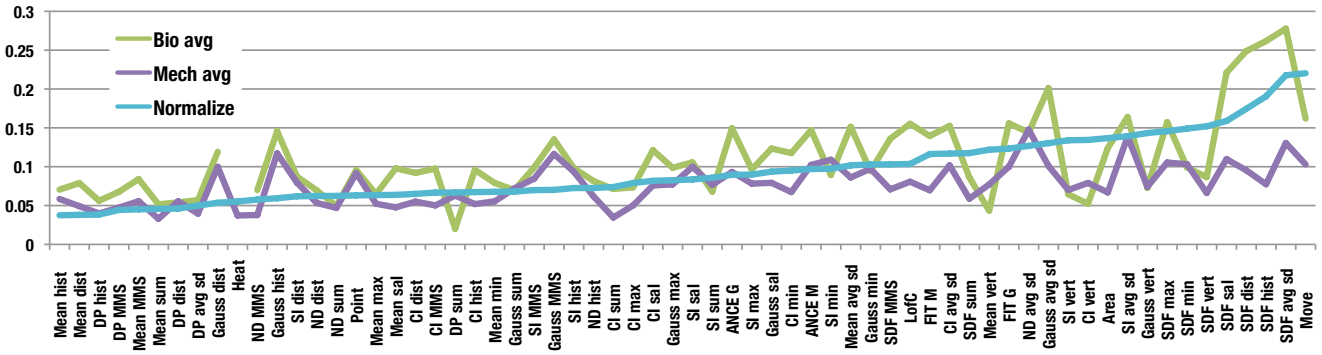

Fig. 3 Stability score. For each data set we averaged the L^2 norm of the difference between the original and the altered shape descriptor at each vertex across all 28 meshes, then averaged the results. The descriptors are sorted by this score, normalized by dividing by the average descriptor norm. We also show the raw score for the three mechanical models and the three biomedical models.

Curvatures: We use four curvature values, Mean, Gaussian, Shape and Curvature index (SI, CI) [19]. SI and CI map curvature values to a reasonable range using `arctan` and `log`. We apply an ad-hoc curvature normalization step to the curvatures (see Section 3.2) because otherwise large values swamp the other ones (Figure 2).
[SDF] Shape Diameter function: Essentially, the distance to the opposite side surface [30,12].

For the above descriptors we experiment with eight different approaches to sampling the data on the rings:
[VERT] $R = 1, S = 1$. Just the data at the vertex.
[SAL] $S = 1$. Saliency sampling [21]. We used the radii as the center sampling size.
[MIN/MAX/AVG] $S = 1$. Minimum, Maximum, or Average of the values on the ring.
[MMA] $S = 3$. Minimum, maximum, and averaged values.
[DIST] $S = 5$. Sort the values. Take the values that lie at the 0%, 10%, 50%, 90%, and 100% places in the sorted list.
[HIST] $S = 7$. Values at 0%, 10%, 30%, 50%, 70%, 90%, and 100%.
[AS] $S = 2$. Take the average and the standard deviation of the values on the ring.

3.1.2 Expanding descriptors:

[LEN:] Length of the ring over the radius.
[AREA:] Area of the volume in the sphere.
[ANCE:] Uses the ANCE [16] method to calculate the Mean and Gaussian curvature, using the vertex location and the re-sampled points of the ring as input.
[FIT:] Fit a degree two polynomial to all of the surface points inside the ring, plus the resampled ring points, weighted by distance. Calculate the Mean and Gaussian curvature from the polynomial. We use Desbrun's'[8] approach to parameterize the surface.

3.1.3 Iterative operator descriptors:

[MOV] Movement. The distance of the vertex from its original position after applying Laplacian smoothing [9] three times for each ring.
[HEAT] Heat diffusion [32].

Taking into account the different ways of sampling the rings, we have $8 + 2 \times 5 + 5 \times 9 + 1 = 67$ descriptors (we include a descriptor which is all of the curvature metrics plus SDF sampled at the vertex).

3.2 Coordinate systems, normalization, and reduction

Naively comparing the raw descriptors described above using an L^n norm has problems. First, the distribution of values may not be linear, or even purely exponential (see Figure 2c,d). This means that a delta difference of, eg, 0.5 may mean nearly identical for vectors with values at the extremes, but not at all the same for vectors with values at the center. Scaling by the length of d can help some, but does not really address the problem. Also, the vectors themselves may not be uniformly distributed. We us a combination of dimension reduction and normalization to address these problems.
Dimension reduction: After computing the raw values we perform a coordinate system transformation by applying Principal Components Analysis (PCA) or Multidimentional Scaling (MDS) as appropriate to all descriptor values. This has an added advantage that the first eigenvector carries the bulk of the information (see Figure 2, bottom right). In this new coordinate system we apply an ad-hoc normalization to map the values in each dimension to the range $[0, 1]$. For studies one and

two (Sec. 4 and 5) we keep all of the dimensions. In practice, we have empirically determined that we can drop the remaining coefficients when the Eigenvalues drop below 10% of the first one without much loss. This happens around the 3rd - 10th coordinate, depending on the original dimensionality of the data (MMS - 5, DIST - 7, HIST - 10, all others 3). We used the reduced dimensionality data for study three (Sec. 6).

We use PCA for all but the Average and Standard Deviation sample method. For this one, the correct comparison method is the *Kullback-Leibler Divergence (KLD)*. We compute a distance matrix using KLD, then apply multi-dimensional scaling (MDS) to that matrix. Because the KLD is expensive to compute, in practice we apply MDS to a $5R \times 5R$ matrix. We select the $5R$ vectors by adding in the vector that is furthest from any of the vectors currently selected.

Directionality: The eigenvectors can point in one of two ways. For visualization purposes, it is nice if the positive direction of the first eigenvector corresponds to positively curved regions, as best as possible. We determine if the first eigenvector should be flipped by comparing the coordinate directionality to either the Gaussian (Gaussian-based metrics) or the Mean (all other) curvature direction.

Normalization: This is a purely ad-hoc solution to the comparison problem. Our only justification is that the descriptor values "look" evenly distributed after the mapping (see Figure 2). Sort the coordinate values. Divide the sorted values into 5 bins, at percentages $(0, 0.1, 0.3, 0.7, 0.9, 1)$. Within each bin, map the values to the range of the bin, optionally applying Eq. 1) one to three times. We determine how many times by choosing the mapping that is closest to a line (ie, minimizes the sum $|x - y|$).

$$y(x) = e^{x^2} \text{ or } \quad y(x) = 1 - e^{(1-x)^2} \tag{1}$$

Note that this equation has two versions — one for when the y values lie below the line $x = y$, one for when they lie above. The values x are assumed to be linearly mapped to the range $[0, 1]$ before applying this mapping.

We scale all objects to be unit size and map the raw mean and Gaussian curvature values as follows: Mean curvature: bin boundaries at 0 (-500), 0.05 (-20), 0.35 (-8), 0.5 (0), 0.65 (10), 0.95 (60) and 1 (700), with 2, 1, 0, 0, 0, 2 applications of Eq. 1. Gaussian curvature: 0 (-40,000), 0.05 (-2000), 0.15 (-500), 0.5 (0), 0.85 (400), 0.95 (3000) and 1 (40,000) with 2, 1, 0, 0, 2, 3 applications of Eq. 1. These values were found by experimentation on the curvature values produced by the data sets.

4 Stability study

For this study we evaluate the effects of mesh resolution, noise, and smoothing on the descriptors. We started with six meshes (fandisk, Isidore rocking horse, horse, ferret brain, mandible, and tibia, see Figure 1), from which we generated a total of 28 meshes each with different resolutions and noise or smoothing.

Reduction: For each mesh we generated three meshes at different resolutions by applying QSlim [13] with a 30%, 60%, and 80% reduction in the percentage of faces. Even the 80% reduction did not result in a noticeable visual change. To establish the correspondence, we project the original mesh vertices onto the reduced meshes, interpolating the values in the faces. We then averaged the difference between the original and reduced descriptors across all vertices. We plot the increase in error as the mesh is reduced. For the summary plot, we used the 80% reduction values.

Noise: We generated clamped Gaussian noise $\delta = e^{[-2,2]/2}$ and shifted each point along its normal $\hat{n}$. Let B be the diagonal of the bounding box containing the surface. We created three meshes with three noise levels, $\pm\delta B$ (0.001, 0.0005, 0.00025) for each of the four mesh reduction levels. To generate the noise plots, we compared the noisy mesh to the unaltered mesh of the same resolution, producing four plots per descriptor.

Smoothing: We applied area-normalized Laplacian smoothing [9]. At each iteration we moved each point 1/3 of the way to its Laplacian average, then area-normalized. We generate meshes and plots as for the Noise case.

Figure 3 summarizes the results across all descriptors and all studies. For all measures, we used the L^2 norm.

We saw the same general trends in all three studies, in terms of the behavior of the descriptors. Distance to plane and Normal distribution were the most stable descriptors, with Mean, CI, SI, and Gauss roughly similar. SDF was the next most stable, followed by Movement and Point (combining Mean, Gauss, SI, CI, SDF, and Movement at the point). Length of Curve, followed by the four fitting descriptors (ANCE and FIT), were not that stable.

For the sampled descriptors, the best sampling strategy was the histogram one (seven samples) followed closely by the distribution descriptor (five samples) then MMA (three samples). AVG, Saliency, and MIN/MAX were, on average, the same, with AVG slightly outperforming the others. Average and Standard deviation (AS) and Vertex were not very stable, with Vertex being very unstable.

We analyzed the effects of mesh resolution, noise, and sampling independently by looking at the individ-

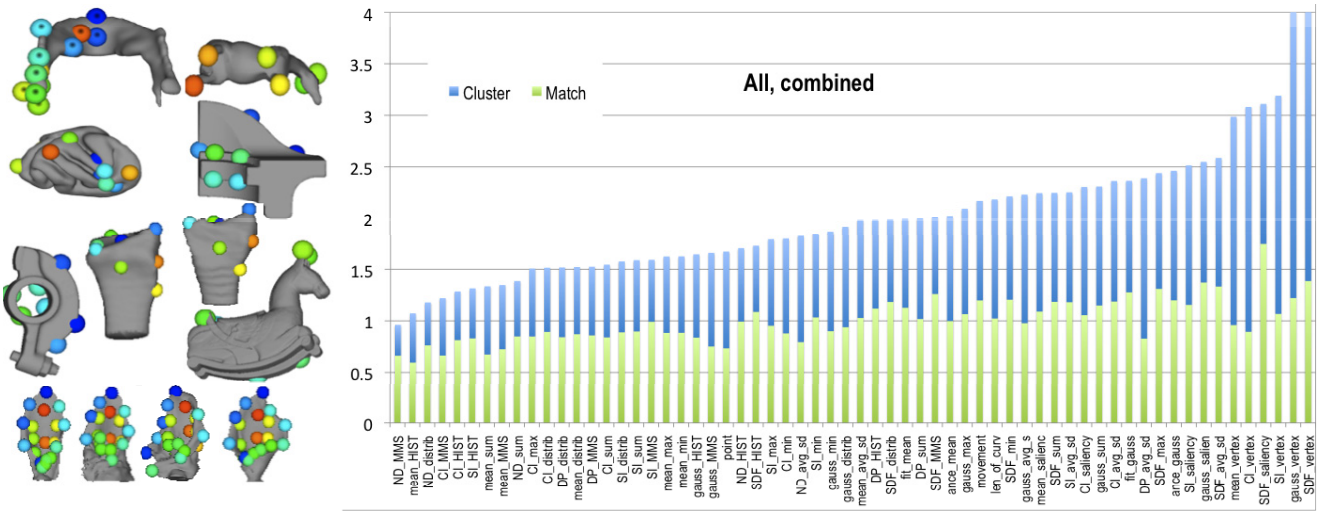

Fig. 4 Local feature selectivity. a) Examples of selected points. Note: Ball size is the actual size of the largest ring. b) Yellow bars: Average spread of the shape descriptor values from the mean. A smaller spread is better. Blue bars: How likely a point was to match to another point in the same set (Sec 5). Zero is better. Both measures were divided by their mean in order to combine them in one plot.

ual plots (examples in Figure 1, bottom right). In summary, we saw the following: 1) Adding noise affects the higher resolution meshes more than it does the lower resolution ones. 2) Smoothing tends to affect the lower resolution meshes more than the higher resolution ones. 3) The smoother or "nicer" the mesh is to start with the less adding noise or smoothing makes a difference. 4) Using multiple rings produced an order of magnitude improvement over simply using the descriptor calculated at the vertex. 7) Additional averaging (saliency sampling, fitting) was, surprisingly, more prone to error than sampling using rings. 8) SI and CI are slightly more stable than our (normalized) Mean and Gauss curvature, with Gauss being the most unstable.

5 Sensitivity study

The previous study looked at the stability of the descriptors — how much they were influenced by noise, mesh sampling, and smoothing. In these studies, we evaluate the descriptors by how well they can distinguish features. For each data set (mandible (10), ferret (8), bat ears (35), radius and ulna (10), mechanical parts (4)) we hand-picked sets of points which should be similar (the "same" point on different instances of similar meshes). Each set was chosen to be sufficiently different from the other ones, eg the tips of the ears versus the edge. This is, admittedly, a human-perception biased method of creating an evaluation set. However, given that we want to work with real data sets, not artificially generated ones, we believe it is justified. Randomly moving the selected feature points a small amount ($< 0.2r$) produced qualitatively similar results.

Data analysis: Ideally, points in one set should have descriptors that are similar, while points in other sets should have descriptors that are dissimilar. We picked two measures to evaluate this. The first measure is simply the spread of a given shape descriptor's feature value over all of the points in the set. These values are summed up over all of the sets to yield a score for each shape descriptor. A zero average and narrow standard deviation are good.

The second measure evaluates how distinctive the shape descriptors are, i.e., how likely a point from one set is to match to points in another set. To calculate this, we compare each point to all of the other points, and sort the results. We then count how far down the list we have to go to find a point in the same set. Ideally, the count should be zero. For this study, we only compared points that came from the same data set (eg, the bat ears), not to all of the points in all of the data sets.

We summarize the results in Figure 4. We combine both scores into a single plot by plotting the first measure, and the second measure normalized to the first measure. The mean and normal descriptors are clearly the best, followed by distance to the plane and Gauss. Not too surprising, sampling each ring with 3-7 samples (MMS, DIST, HIST) performed better than taking a single sample on the ring. For the mechanical models alone, MMS was the better choice).

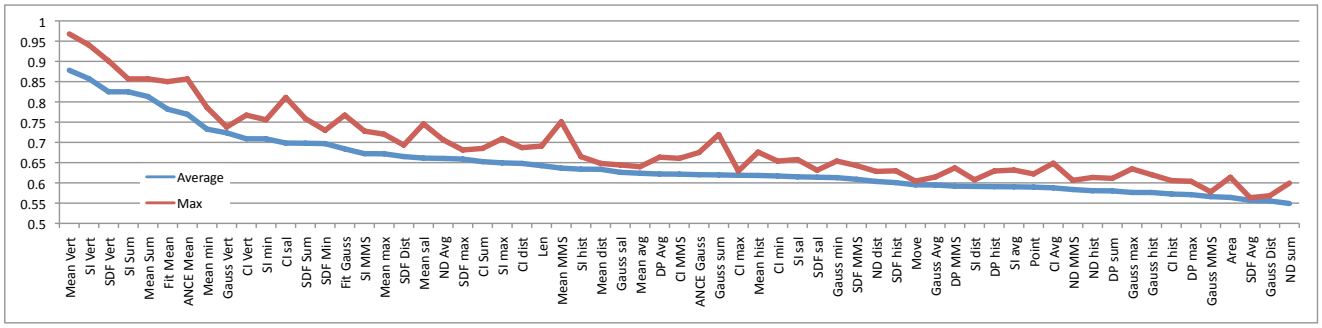

Fig. 5 RAND classification scores, sorted by average RAND score.

6 Classification study

We examine the effectiveness of the shape descriptors for a bag-of-words style classification [22,10]. For this study we used the four mechanical models (horse, Isidore, fandisk, and rocker arm), the ferret brains (13), the tibias (13), and a subset of the bat ears (35 from 7 different genera). Using two samples, chosen at random from each data set, we computed an initial set of $k \in [3, 10]$ shape descriptor clusters using K-Means. The shape descriptor vectors were clipped at the 10% Eigenvalue (Section 3.2). For each mesh in the data set we then computed an area-weighted histogram based on the distribution of the k clusters across the surface. We then ran K-Means clustering again, this time on the histograms, and compared the classification results to the known classification using the RAND index [27]. We limited the number of words (k) to 10 because K-Means became unstable around 10 clusters. Results are shown in Figure 5. The mean curvature metrics, either sampled at the vertex or averaged around the vertex, performed the best, followed by the shape diameter function, which is a more global metric. Given the stability results of the previous section, the best metric would be the average of the mean curvature sampled around the vertex, followed by the average of the shape diameter function sampled around the vertex.

Feature-based species classification: The previous study looked at the entire shape descriptor distribution; in this study we looked at classification using just the shape descriptor values at hand-picked locations (see bottom left of Figure 4). We used a total of 35 meshes from 7 genus. We formed a feature vector by concatenating the descriptors at each feature point, and used K-Means to produce a classification. We ranked the descriptors by how well they matched the known classification. Distance to plane and CI were the clear winners, followed by Normal to Plane. Again, using more samples (Hist, MMS, Dist) was better than using fewer samples. Since these two metrics are not strongly correlated, combining them might produce better outcomes.

7 Correlation study

An obvious question to ask is, if one descriptor is good, would two be better? We did not explicitly compare combing descriptors, but we did perform a correlation study (see Figure 6). The strongest correlations are, of course, between different samplings of the same descriptor. SDF correlates the most strongly with itself, followed by Mean curvature, Normal Distribution, and Distance to Plane. Gaussian curvature did not show much correlation. Of the remaining descriptors, Movement and Length of Curve did not correlate strongly with much else. Of the two, Movement is better for both noise and sensitivity, making it a good candidate to combine with either the Normal Distribution or the Mean curvature.

8 Results and discussion

The studies clearly show that sampling data just at a vertex is bad; this is not surprising. The general trend, for all of the descriptors, is that more samples is better, both for stability *and* discrimination, regardless of the descriptor used. We did see evidence of over-fitting for the highest level of sampling in the sensitivity tests, which indicates that the ideal sampling is somewhere around 3-5 samples per ring, which can be reduced to 7-10 samples total via PCA. The Normal distribution descriptor is consistently the best, followed by Mean curvature or the Curvature index.

The non-sampled descriptors (Movement, Length of curve) did about as well as the one sample per ring version of the other descriptors, which indicates that it

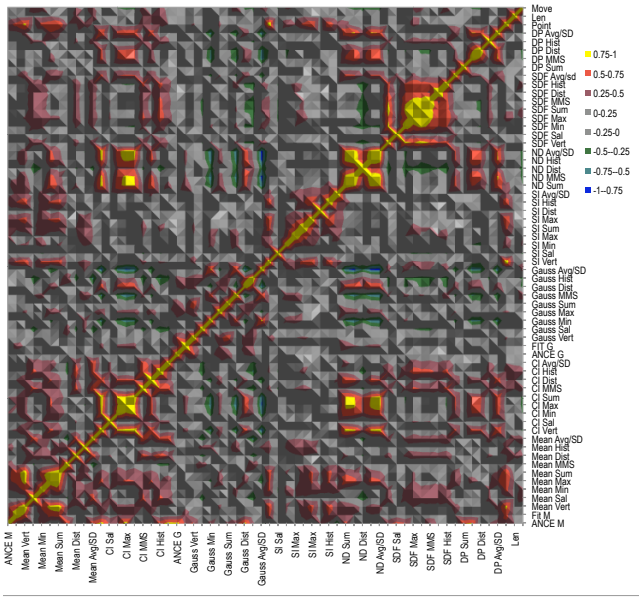

Fig. 6 Correlations of descriptors over all data sets.

is primarily the increased sampling that is of benefit. Not too surprisingly, the Shape Diameter Function is less discriminating locally than other descriptors since, in a sense, it is measuring mid-scale features in the form of the local medial axis.

One surprising result is that the ranking for stability is similar for the ranking for discrimination power. This hints that the ability to filter out small geometry changes is related to the ability to cluster similar local shapes.

Limitations of the study: Obviously, there are many ways that local shape descriptors can be constructed, and there may be unintended biases in the particular implementations we use (for example our ad-hoc normalization). We also only evaluated the feature matching task at small scales; for segmentation a different descriptor may be more appropriate.

Timings: For smaller meshes ($< 30,000$ vertices) the descriptors are roughly equivalent, at 2-10 seconds per descriptor (Movement being much faster, ANCE and Fit being slowest). For bigger meshes, finding the rings dominates the calculations (up to an hour for 90,000 vertices). Area and SDF also do not scale well. Saliency and avg/sd are the slowest ring sampling methods, with the others about equal. The PCA and normalization calculations are dominated by the length of the vector; a few seconds for most of the descriptors, up to a couple of minutes for the histogram sampling on large ($> 70,000$) meshes.

Acknowledgements: Funded in part by NSF grants CCF 0702662 and DBI 1053171.

9 Conclusion

We have presented a systematic evaluation of local shape descriptors for the task of local feature matching on real data sets, both biological and man-made. This was accomplished by creating a unifying framework for the disparate local shape descriptors.

Acknowledgements Funded in part by National Science Foundation grants CCF 0702662, DBI 1053171, DMS 0540701, NIH T90 DA022871, Shandong Taishan Fund, NNSF of China, Ministry of Education, People's Republic of China (985 & 211), Shandong University, and the EU CILIA Project.. Thanks to Dr. Bayly for the ferret brains, Dr. Daniel Low for the mandibles, and Dr. Crisco for the radius and ulna bones.

References

1. Belongie, S., Malik, J., Puzicha, J.: Shape matching and object recognition using shape contexts. IEEE Trans. Pattern Anal. Mach. Intell. **24**(4), 509–522 (2002). DOI http://dx.doi.org/10.1109/34.993558
2. Bronstein, A.M., Bronstein, M.M., Kimmel, R., Mahmoudi, M., Sapiro, G.: A gromov-hausdorff framework with diffusion geometry for topologically-robust non-rigid shape matching. Int. J. Comput. Vision **89**, 266–286 (2010). DOI http://dx.doi.org/10.1007/s11263-009-0301-6. URL `http://dx.doi.org/10.1007/s11263-009-0301-6`
3. Chua, C.S., Jarvis, R.: Point signatures: A new representation for 3d object recognition. Int. J. Comput. Vision **25**(1), 63–85 (1997). DOI http://dx.doi.org/10.1023/A:1007981719186
4. Cipriano, G., Phillips Jr., G.N., Gleicher, M.: Multi-scale surface descriptors. IEEE Trans. on Viz. and CG **15**, 1201–1208 (2009). DOI http://dx.doi.org/10.1109/TVCG.2009.168. URL `http://dx.doi.org/10.1109/TVCG.2009.168`
5. Clarenz, U., Griebel, M., Rumpf, M., Schweitzer, M.A., Telea, A.: Feature sensitive multiscale editing on surfaces. Vis. Comp. **20**, 329–343 (2004). DOI 10.1007/s00371-004-0245-3. URL `http://portal.acm.org/citation.cfm?id=1014497.1014501`
6. Clarenz, U., Rumpf, M., Telea, A.: Robust feature detection and local classification for surfaces based on moment analysis. IEEE Trans. on Viz. and CG **10**, 516–524 (2004). DOI 10.1109/TVCG.2004.34. URL `http://portal.acm.org/citation.cfm?id=1009231.1009305`
7. Connolly, M.L.: Measurement of protein surface shape by solid angles. J. Mol. Graph. **4**, 3–6 (1986). DOI 10.1016/0263-7855(86)80086-8. URL `http://portal.acm.org/citation.cfm?id=15940.15941`
8. Desbrun, M., Meyer, M., Alliez, P.: Intrinsic parameterizations of surface meshes. Computer Graphics Forum **21**, 209–218 (2002)
9. Desbrun, M., Meyer, M., Schroeder, P., Barr, A.H.: Implicit fairing of irregular meshes using diffusion and curvature flow. In: SIGGRAPH '99, pp. 317–324 (1999). DOI http://doi.acm.org/10.1145/311535.311576
10. Fehr, J., Streicher, A., Burkhardt, H.: A bag of features approach for 3d shape retrieval. In: Proceedings of the 5th International Symposium on Advances in Visual Computing: Part I, ISVC '09, pp. 34–43. Springer-Verlag, Berlin, Heidelberg (2009). DOI http://dx.doi.

org/10.1007/978-3-642-10331-5_4. URL http://dx.doi.org/10.1007/978-3-642-10331-5_4
11. Fischl, B., Sereno, M.I., Dale, A.M.: Cortical surface-based analysis: Inflation, flattening, and a surface-based coordinate system. NeuroImage **9**(2), 195 – 207 (1999). DOI DOI:10.1006/nimg.1998.0396. URL http://www.sciencedirect.com/science/article/B6WNP-45FCP7R-2V/2/e5571c1b8fda5a564cc17632e56bb785
12. Gal, R., Shamir, A., Cohen-Or, D.: Pose-oblivious shape signature. IEEE Trans. on Viz. and CG **13**, 261–271 (2007). DOI http://dx.doi.org/10.1109/TVCG.2007.45. URL http://dx.doi.org/10.1109/TVCG.2007.45
13. Garland, M., Heckbert, P.S.: Surface simplification using quadric error metrics. In: SIGGRAPH '97, pp. 209–216 (1997). DOI http://doi.acm.org/10.1145/258734.258849
14. Gatzke, T., Grimm, C., Garland, M., Zelinka, S.: Curvature maps for local shape comparison. In: Shape Modeling and Applications, pp. 246–255 (2005). DOI 10.1109/SMI.2005.13. URL http://portal.acm.org/citation.cfm?id=1097876.1098476
15. de Goes, F., Goldenstein, S., Velho, L.: A hierarchical segmentation of articulated bodies. In: SGP '08, SGP '08, pp. 1349–1356 (2008). URL http://portal.acm.org/citation.cfm?id=1731309.1731315
16. Goldfeather, J., Interrante, V.: A novel cubic-order algorithm for approximating principal direction vectors. ACM Trans. Graph. **23**(1), 45–63 (2004). DOI http://doi.acm.org/10.1145/966131.966134
17. Grimm, C., Li, R., Heider, P., Pierre-Pierre, A., Mueller, R.: Poster: A comparison of local shape descriptors for biological applications. Computational Advances in Bio and Medical Sciences, IEEE International Conference on **0**, 245 (2011). DOI http://doi.ieeecomputersociety.org/10.1109/ICCABS.2011.5729898
18. Heider, P., Pierre-Pierre, A., Li, R., Grimm, C.: Local shape descriptors, a survey and evaluation. In: 3DOR, pp. 49–56 (2011)
19. Koenderink, J.J., van Doorn, A.J.: Surface shape and curvature scales. Image Vision Comput. **10**, 557–565 (1992). DOI 10.1016/0262-8856(92)90076-F. URL http://portal.acm.org/citation.cfm?id=145375.145386
20. Kortgen, M., Park, G.J., Novotni, M., Klein, R.: 3d shape matching with 3d shape contexts. In the 7th Central European Seminar on Computer Graphics (2003)
21. Lee, C.H., Varshney, A., Jacobs, D.W.: Mesh saliency. ACM Trans. Graph. **24**, 659–666 (2005). DOI http://doi.acm.org/10.1145/1073204.1073244. URL http://doi.acm.org/10.1145/1073204.1073244
22. Li, X., Godil, A.: Exploring the bag-of-words method for 3d shape retrieval. In: Image Processing (ICIP), 2009 16th IEEE International Conference on, pp. 437 –440 (2009). DOI 10.1109/ICIP.2009.5414415
23. Li, X., Guskov, I.: Multi-scale features for approximate alignment of point-based surfaces. In: SGP 2005 (2005). URL http://portal.acm.org/citation.cfm?id=1281920.1281955
24. Mortara, M., Patane, G., Spagnuolo, M., Falcidieno, B., Rossignac, J.: Blowing bubbles for multi-scale analysis and decomposition of triangle meshes. Algorithmica **38**(1), 227–248 (2003). DOI http://dx.doi.org/10.1007/s00453-003-1051-4
25. Ong, J.L., Seghouane, A.K.: From point to local neighbourhood: Polyp detection in ct colonography using geodesic ring neighbourhoods. IEEE Trans Image Process (2010). URL http://www.biomedsearch.com/nih/From-Point-to-Local-Neighbourhood/20840898.html
26. Pottmann, H., Wallner, J., Huang, Q.X., Yang, Y.L.: Integral invariants for robust geometry processing. Comput. Aided Geom. Des. **26**, 37–60 (2009). DOI 10.1016/j.cagd.2008.01.002. URL http://portal.acm.org/citation.cfm?id=1464516.1464795
27. Rand, W.M.: Objective criteria for the evaluation of clustering methods. Journal of the American Statistical Association **66**(336), pp. 846–850 (1971)
28. Rustamov, R.M.: Laplace-beltrami eigenfunctions for deformation invariant shape representation. In: SGP 2007, pp. 225–233 (2007). URL http://portal.acm.org/citation.cfm?id=1281991.1282022
29. Schmidt, R., Grimm, C., Wyvill, B.: Interactive decal compositing with discrete exponential maps. In: SIGGRAPH '06, pp. 605–613. ACM (2006). DOI http://doi.acm.org/10.1145/1179352.1141930
30. Shapira, L., Shamir, A., Cohen-Or, D.: Consistent mesh partitioning and skeletonisation using the shape diameter function. Vis. Comput. **24**, 249–259 (2008). DOI 10.1007/s00371-007-0197-5. URL http://portal.acm.org/citation.cfm?id=1356722.1356724
31. Stein, F., Medioni, G.: Structural indexing: Efficient 3-d object recognition. IEEE Trans. Pattern Anal. Mach. Intell. **14**(2), 125–145 (1992). DOI http://dx.doi.org/10.1109/34.121785
32. Sun, J., Ovsjanikov, M., Guibas, L.: A concise and provably informative multi-scale signature based on heat diffusion. In: SGP '09, SGP '09, pp. 1383–1392 (2009). URL http://portal.acm.org/citation.cfm?id=1735603.1735621
33. Yamany, S.M., Farag, A.A.: Free-form surface registration using surface signatures. In: ICCV '99, p. 1098. IEEE Computer Society (1999)
34. Yamany, S.M., Farag, A.A.: Surfacing signatures: An orientation independent free-form surface representation scheme for the purpose of objects registration and matching. IEEE Trans. Pattern Anal. Mach. Intell. **24**(8), 1105–1120 (2002). DOI http://doi.ieeecomputersociety.org/10.1109/TPAMI.2002.1023806
35. Zelinka, S., Garland, M.: Similarity-based surface modelling using geodesic fans. In: SGP '04, pp. 204–213. ACM (2004). DOI http://doi.acm.org/10.1145/1057432.1057460