Distinctive Regions of 3D Surfaces

PHILIP SHILANE and THOMAS FUNKHOUSER
Princeton University

Selecting the most important regions of a surface is useful for shape matching and a variety of applications in computer graphics and geometric modeling. While previous research has analyzed geometric properties of meshes in isolation, we select regions that distinguish a shape from objects of a different type. Our approach to analyzing distinctive regions is based on performing a shape-based search using each region as a query into a database. Distinctive regions of a surface have shape consistent with objects of the same type and different from objects of other types. We demonstrate the utility of detecting distinctive surface regions for shape matching and other graphics applications including mesh visualization, icon generation, and mesh simplification.

Categories and Subject Descriptors: I.3.5 [**Computer Graphics**]: Computational Geometry and Object Modeling

General Terms: Algorithms

Additional Key Words and Phrases: Distinct features, partial matching, shape distinction, similarity, visualization

ACM Reference Format:

Shilane, P. and Funkhouser, T. 2007. Distinctive regions of 3D surfaces. ACM Trans. Graph. 26, 2, Article 7 (June 2007), 15 pages. DOI = 10.1145/1243980.1243981 http://doi.acm.org/10.1145/1243980.1243981

1. INTRODUCTION

Many problems in computer graphics, computer vision, and geometric modeling require reasoning about which regions of a surface are most important. For example, in an object classification system, a query might be compared only against the most important regions of a target object to provide more discriminating matches. Also in a mesh simplification algorithm, the importance of vertices might be used to determine the order in which they are decimated. In an icon generation system, the viewpoint might be chosen so that the important regions of an object are maximally visible.

Although there has been significant progress in algorithms for determining important regions of polygonal meshes, most prior work has focused on geometric properties of every mesh in isolation. For example, Lee et al. [2005] defined a measure of mesh saliency using a center-surround operator on Gaussian-weighted mean curvatures. Related measures of mesh importance have been defined by Gal et al. [2006], Li et al. [2005], Novotni et al. [2005], Gelfand et al. [2005], and others. However, almost all of these methods simply select regions where the curvature of a surface patch is different than in its immediate neighborhood.

Intuitively, the important regions of an object for recognition are not the ones with specific curvature profiles, but rather the ones that distinguish it from objects of other types, i.e., the *distinctive* regions. For example, consider the biplane shown in Figure 1. Any five-year-old will tell you that the important features of the biplane are the wings and tail (Figure 1). Those features are unique to biplanes and thus distinguish the biplane from other types of vehicles. Generalizing this idea, we define the *distinction* of a surface region as how useful it is for distinguishing the object from others of different types.

This definition has an interesting implication: in order to determine how distinctive a surface region is, we must not only consider its properties, but we must also consider how consistent those properties are within other instances of the same object type and how unique those properties are with respect to other object types under consideration. For example, if we consider the biplane among other types of airplanes, we find that the wings are the most distinctive features. However, if we consider it as a biplane among other types of objects (tables, animals, cars, etc.) many of which have large flat regions, we find that the tail is most distinctive (Figure 1). In general, the distinctive regions of a surface will be different depending on the granularity and range of object types under consideration.

In this article, we describe methods for defining, computing, and utilizing *surface distinction*. Our approach is motivated by shape-based retrieval measures. We define the distinction of a surface region based on a measure of how well a shape-based search of a database with that region as the query produces a ranked retrieval list with objects of the same type near the front of the list. To compute this measure of importance, we generate a random set of points on the surface of a mesh. Then, for each point, we compute Harmonic Shape Descriptors (HSD) [Kazhdan 2004] at four different scales and match them to all others of the same scale in a database of meshes partitioned into object classes. The distinction of every point at every scale is computed as the discounted cumulative gain (DCG) [Jarvelin and Kekalainen 2000] of the retrieval

The authors would also like to thank several funding agencies that provided partial funding for this work: National Science Foundation Grants IIS-0612231, CCR-0093343, CNS-0406415, and IIS-0121446; Air Force Research Laboratory Grant FA8650-04-1-1718; and the Google Research Grants Program.
Author's address: Computer Science Department, 35 Olden Street, Princeton, NJ, 08540; email: {pshilane,funk}@cs.princeton.edu.
Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.
© 2007 ACM 0730-0301/2007/06-ART7 $5.00 DOI 10.1145/1243980.1243981 http://doi.acm.org/10.1145/1243980.1243981

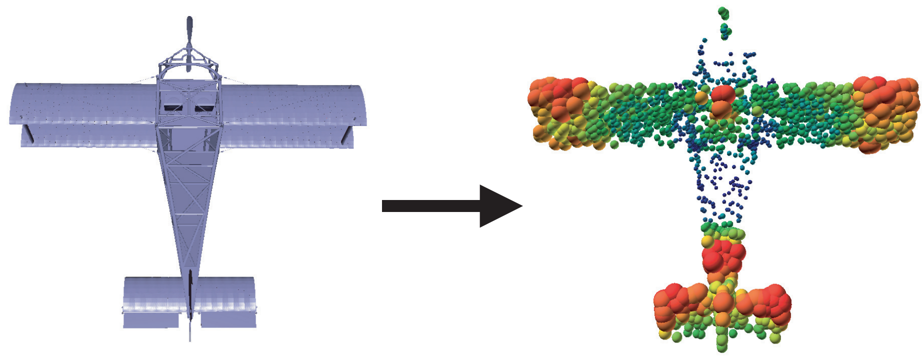

Fig. 1. Distictive regions of the plane correspond to the important regions that define the object type and distinguish the plane from other types of objects. Regions shown in red are the most distinctive, blue are least distinctive, and green are in the middle. This result corresponds to our intuition that the wings and tail are important features of a plane.

list generated by ranking the best matches from every object in the database and measuring how often objects from the same class appear near the front of the list [Shilane and Funkhouser 2006]. The distinction values are then optionally mapped onto the vertices of the mesh and used for a variety of visualization and processing applications.

The contributions of this article and its associated conference papers [Funkhouser and Shilane 2006; Shilane and Funkhouser 2006] are threefold. First, we provide a definition of surface distinction that captures how distinguishing every region of a surface is with respect to object classes in a database. Second, we describe a practical method for computing surface distinction using shape matching algorithms and evaluation methods derived from information retrieval. Third, we demonstrate the utility of our surface distinction measure for shape matching and several graphics applications including mesh visualization, icon generation, and mesh simplification.

2. RELATED WORK

There has been a long history of related work in cognitive psychology, computer vision, computer graphics, geometric modeling, statistics, and pattern recognition. In this section, we focus on techniques for selecting important local regions of 3D surfaces for shape matching and their applications for computer graphics.

Perceptual Criteria. There have been several attempts to select regions of 3D shapes that humans find visually important in object recognition, perceptual psychology, and computer vision. For example, Howlett et al. [2004] used an eye-tracker to record where a person looks at a 3D model and then used that information to assign importance to vertices in a mesh simplification algorithm. While this method captures a useful notion of surface importance, it is viewpoint-dependent and requires human analysis of every 3D mesh, which is impractical for the large databases of 3D meshes targeted by our system.

Several psychophysical experiments have found that the human visual system quickly processes regions of high curvature (e.g., Hoffman and Singh [1997]). These findings have been applied extensively for object recognition in computer vision [Lowe 1999; 2004]. For example, combinations of filters measuring edges and local maxima of curvature in 2D images have been used to focus scene recognition algorithms [Frintrop et al. 2004]. More recently, they have also been applied to define measures of saliency for mesh processing applications. For example, Lee et al. [2005] used a center-surround filter of curvature across multiple scales to select salient regions for mesh simplification and viewpoint selection. Similarly, Gal and Cohen-Or [2006] computed the saliency of a region based on its size relative to the whole object, its curvature, the variance of curvature, and the number of curvature changes within the region, and they used this measure to guide partial shape matching, self-similarity detection, and shape alignment. Li and Guskov [2005] computed surface signatures describing the curvature and other properties for local regions and only keep the ones with significantly nonzero magnitude. Novotni et al. [2005] selected points found as local extrema of the differences of Gaussian filters applied to the characteristic function of the surface.

While these approaches are able to select regions that may be visually noticable, they focus on curvature and other measures appropriate for manifold surfaces. Thus, they cannot be used effectively for the majority of 3D computer graphics models which often contain disjoint and intersecting polygons. More importantly, they measure how much a region sticks out from the rest of the object rather than how important the region is for defining the object type.

Statistical Criteria. Numerous techniques have been developed for selecting important features in the realm of statistical analysis and pattern recognition, which are covered in several classical books [Duda et al. 2001; Hastie et al. 2001; McLachlan 1992]. The classical problem is, given a set of feature vectors and labeled training data, to select a subset of features that is most useful for classification. Many techniques from this field can also produce real-valued importance scores for each feature. Discriminant analysis [Lachenbruch and Goldstein 1979] or analysis of variance (ANOVA) [Miller 1997] selects feature vectors that are consistent within a class and have a large separation from other classes of objects. Using regression analysis, a subset of features can be selected with stepwise selection that either grows or shrinks a subset of features to optimize a function [Copas 1983; Frank and Friedman 1993]. The problem we address differs from classical statistical analysis because these approaches assume a correspondence between features,

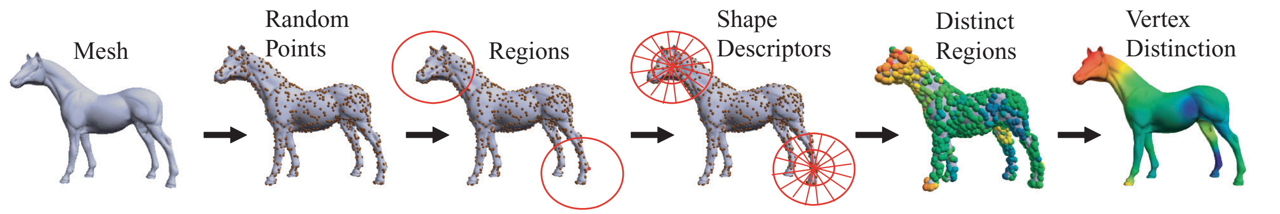

Fig. 2. Partitioning a surface into regions and computing how distinctive the regions are with respect to a set of object classes provided in a database.

and they select dimensions of a feature vector rather than positions on a surface.

In the shape matching literature, the relative rarity of local surface patches has been used as a measure of importance to guide several shape matching systems without requiring correspondences between patches. Typically, representations of local shape (shape descriptors) are computed for many regions of a surface, and then they are weighted according to a measure of uniqueness when point sets are aligned and/or matched. For example, Chua and Jarvis [1996] found selective points on a surface by comparing their descriptors to others in a local neighborhood and used the descriptor that was most unique for shape matching [Chua and Jarvis 1996]. Johnson [2000] computed the likelihood of each shape descriptor based on a Gaussian distribution of the descriptors for each mesh and then selected only the least likely ones to speed up surface matching. However, these methods only find descriptors that are rare—they do not specifically find ones that are distinctive of an object class.

Shan et al. [2004] used shape descriptor matching to define the importance of points on a mesh after calculating correspondences in a database. This method selects points based on how well they match multiple points on one object (e.g., a spherical region will be selected if there is an object with many spherical regions in the database), and thus it provides a measure of stability for shape matching rather than a measure of importance for object classification as is provided by our method.

This article provides an extended discussion of the methods first proposed in Shilane et al. [2006] and Funkhouser et al. [2006]. In Shilane and Funkhouser [2006], the idea of using retrieval performance metrics to measure shape distinction was first proposed. That paper focused on predicting the distinction of regions on a query shape, while this article focuses on the analysis of shapes in a classified database. In Funkhouser and Shilane [2006], the method described in this article was used to improve multipoint shape matching. This article provides further analysis of the benefit of using distinctive points for matching as well as visualization and mesh processing applications that utilize a measure of surface importance.

3. METHODS

In this article, we describe methods for defining, computing, and utilizing *surface distinction*, a measure of how precisely a local surface region indicates a particular class of objects. Intuitively, regions that are common among many object classes are not distinctive (e.g., planar regions, spherical regions, etc.), while others that are found in only one object class are very distinctive (e.g., the ears of a bunny, the wheels of a motorcycle, the handle of a guitar, the head of a wrench, etc.). In our system, we assign a continuous value of distinction to every surface region at multiple scales, with 0 indicating that the region is not distinctive at all (i.e., that region could be found equally well in any object class), 1 indicating that the region is perfectly distinctive (i.e., that region is found in only one object class), and values in between representing the degree to which the region distinguishes the object class (Figure 1).

Computation of surface distinction proceeds in our system as shown in Figure 2. Given a database of meshes partitioned into object classes, we first generate for each mesh a set of random points that are the centers of spherical regions covering its surface at multiple scales. Then, for every region, we compute a shape descriptor representing the distribution of surface area within that region. Next, we compare the difference between all pairs of shape descriptors to produce a ranked list of matches for each descriptor ordered from best to worst. The ranked lists are then analyzed to produce measures of how distinctive different local regions are, that is, how many descriptors from the same class of objects appear near the front of their ranked lists. These measures can be directly used to improve shape matching applications and can also be mapped from the regions back onto the vertices of the mesh and used to guide mesh visualization, processing, and analysis. The following sections describe each of these steps in detail.

3.1 Constructing Regions

The first step of our process is to define a set of local regions covering the surface of the object. In theory, the regions could be volumetric or surface patches; they could be disjoint or overlap; and, they could be defined at any scale.

In our system, we construct overlapping regions defined by spherical volumes centered on points sampled from the surface of an object (Figure 3). This choice supports robust processing of arbitrary surface meshes with degenerate topology, and it naturally supports overlapping regions at multiple scales. Formally, the database consists of a set of meshes $\{M_1, \dots, M_m\}$, each mesh M_j has a set of points $\{p_{1,j}, \dots, p_{n,j}\}$ where $p \in R^3$, and each point has a set of scales $\{s_1, \dots, s_r\}$, where a spherical region $r_{i,j,k}$ has center $p_{i,j}$ and radius s_k. We have experimented with two different point sampling methods, one that selects points randomly with uniform distribution with respect to surface area and another that selects points at vertices of the mesh with probability equal to the surface area of the vertices' adjacent faces. However, they do not give significantly different performance, and so we consider only random sampling with respect to surface area in the remainder of this article. Of course, other sampling methods that sample according to curvature, saliency, or other surface properties are possible as well. In most of our experiments, we consider four different scales for every point where the smallest scale has radius 0.25 times the radius of the entire object, and the other scales are 0.5, 1.0, and 2.0 times, respectively. Note that the biggest scale is just large enough to cover the entire object for the most extreme position on the surface, and the smallest scale is large enough to contain easily recognizable shape features.

Our implementation for selecting random points on a surface as centers for these spherical regions follows the ideas of Osada et al.

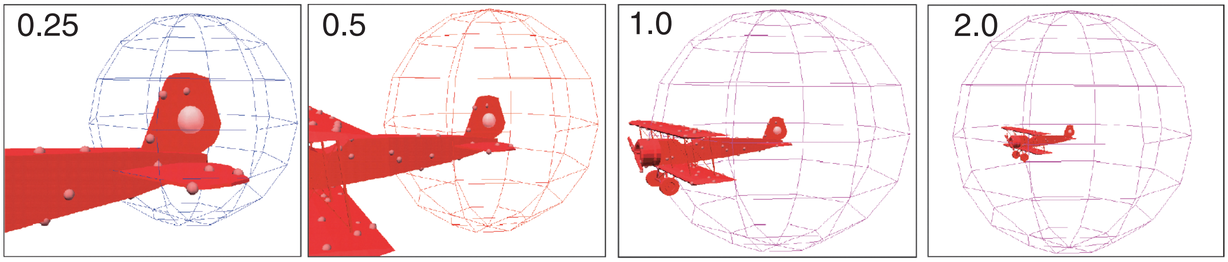

Fig. 3. Four regions are shown with the same center point but different scales. At the 0.25 scale, the tail is included, and at larger scales, progressively more of the plane is included. At the 2.0 scale, the entire plane is included, even when the region is centered on the end of the plane.

[2002]. We have modified their algorithm slightly to make sampling more efficient and to stratify samples in meshes with large triangles. Specifically, in the first stage, we allocate a number of points to every triangle in proportion to its surface area. Then, in the second stage, we sample the allocated number of points from every triangle uniformly with respect to its surface area. This method is faster than Osada's method, taking $O(n)$ rather than $O(n \log n)$ for a mesh with n triangles.

3.2 Describing Shapes

In the second step of our process, for each spherical region $r_{i,j,k}$, we generate and store a shape descriptor $x_{i,j,k}$ which has dimension d. There will be many such regions for every surface, so the shape descriptors must be quick to compute and concise to store. Since we will be matching all pairs of the shape representations, they must be indexable and quick to compare. Also, our methods should work for any input object representation, they must be independent of shape description, insensitive to topology, and robust to common input file degeneracies. Finally, since we aim to model how distinctive the shape of each surface region is, they must be discriminating of similar versus dissimilar shapes.

There are many shape descriptors that meet some or all of these goals (see surveys in Bustos et al. [2005]; Iyer et al. [2005]; Tangelder and Veltkamp [2004]). For example, Belongie et al. [2002] have used shape contexts for describing local regions of 2D images, and Kortgen et al. [2003] have extended their method to 3D. However, shape contexts are dependent on a particular orientation, and thus require alignment within a global coordinate system [Kortgen et al. 2003] or searching possible rotations as they are matched [Frome et al. 2004]. Johnson and Hebert [1999] have used spin images to represent the shapes of local regions with orientation dependence on just the normal to the surface at a sample point, and Vranic et al. [2003] have described Fourier descriptors that can be used to provide invariance to all rotations except those around the surface normal. However, those methods are sensitive to normal orientation, which is highly variable in sparsely sampled point sets considered in this article. Finally, Kazhdan et al. [2003] have described a Harmonic Shape Descriptor (HSD) that is invariant to all rotations. The main idea is to decompose a spherical region into concentric spherical shells of different radii, compute the spherical harmonic decomposition for a function describing the shape in each of those shells, and then store the amplitudes of the harmonic coefficients within every frequency (order) to form a feature vector for indexing and matching (see Funkhouser et al. [2003] for details).

In our system, we describe the shape of every spherical region by a Harmonic Shape Descriptor with 32 spherical shells and 16 harmonic frequencies. We chose this shape representation for several reasons. First, it requires only 512 floating point values (2048 bytes), and thus it is concise enough to store in memory for hundreds of points at several scales for thousands of meshes in a database. Second, it is rotation invariant, and thus two HSDs can be compared quickly without a priori alignment or searching all possible rotations. Third, it is fairly discriminating, according to tests by Shilane et al. [2004]. Finally, the L^2 (Euclidean) difference between any two HSDs provides a lower bound on the L^2 difference between their corresponding shapes, and thus HSDs can be combined with other shape descriptors as a conservative prefilter in a later version of this system.

Our method for computing the HSD for all regions of a single surface starts by computing a 3D grid containing a Gaussian function of the surface's Euclidian Distance Transform (GEDT) [Kazhdan 2004]. This function, which is one at the surface and falls off gradually with Euclidean distance, provides a soft penalty function for matching surfaces by comparison of volumes. The GEDT grid resolution is chosen to match the finest sampling rate required by the HSD for regions at the smallest scale; the triangles of the surface are rasterized into the grid, and the squared distance transform is computed and stored. Then, for every spherical region centered on a point sampled from the surface, a spherical grid is constructed by computing the GEDT at regular intervals of radius and polar angles; the Spharmonickit software [SpharmonicKit 2.5 1998] is used to compute the spherical harmonic decomposition for each radius, and the amplitudes of the harmonic coefficients within each frequency are stored as a shape descriptor.

Computing 128 shape descriptors at four scales (0.25, 0.5, 1.0, and 2.0) for a single mesh takes approximately three minutes overall. One minute is spent rasterizing the triangles and computing the squared distance transform at resolution sufficient for the smallest scale descriptors, and two minutes are spent computing the spherical grids and constructing the 512 harmonic shape descriptors.

3.3 Measuring Distinction

In the third step of our process, we compute how distinctive every shape descriptor is with respect to a database containing multiple classes of objects. Our goal is to compute a continuous measure that reflects how well the shape descriptor for a local region of a surface matches others within the same class of objects relative to how well it matches descriptors in other classes. Descriptors whose best matches are all from its own class are distinctive, while ones that match descriptors in a wide variety of classes equally well are not. While we would ideally like to calculate the distinction value for all combinations of local descriptors and at all scales, this is computationally infeasible. Instead, we make an independence assumption and calculate distinction for each descriptor independently, modeling distinction with an information retrieval metric.

Given the distance from the ith feature of M_j to the closest descriptor of every other mesh in the database,

$$dist(x_{i,j,k}, M_t) = \min_b |x_{i,j,k} - x_{b,t,k}|,$$

we sort the distances from smallest-to-largest to create the retrieval list for $x_{i,j,k}$. We then compute the distinction of the descriptor by evaluating a retrieval performance metric that measures how well meshes in the query's class appear near the front of the list. There are numerous evaluation metrics that could be used to convert a retrieval list into a numeric score including the nearest neighbor, precision for a fixed-length retrieval list, first or second tier, e-measure, f-measure, discounted cumulative gain (DCG), etc. that are described more thoroughly by Leifman et al. [2003] and Shilane et al. [2004]. Each of these measures has trade-offs in terms of how much of the retrieval list is included in the calculation (nearest neighbor uses the first retrieval result, while DCG requires the full list) versus the time necessary to calculate the results (nearest neighbor could be quickest using an indexing structure to find the closest result and DCG the slowest). We have selected the DCG [Jarvelin and Kekalainen 2000] retrieval measure because it has been shown to provide the most stable retrieval measure in previous studies [Leifman et al. 2003; Shilane et al. 2004].

Intuitively, the DCG measure of retrieval performance gives credit to matches within the same class at every position in the rank list but weights matches near the front of the list more than ones near the end of the list. More formally, to calculate the DCG for a query descriptor, a ranked list L of the best matches from each object is converted to a binary list G, where element G_i has value 1 if element L_i is in the same class as the query descriptor, and value 0 otherwise. Discounted cumulative gain is then computed with the following recursion:

$$DCG_i = \left\{ \begin{array}{ll} G_1, & i = 1 \\ DCG_{i-1} + \frac{G_i}{log_2(i)}, & otherwise \end{array} \right\}.$$

The result is then divided by the maximum possible DCG, which is achieved if the first C elements are in the correct class, where C is the size of the class:

$$DCG = \frac{DCG_N}{1 + \sum_{j=2}^{|C|} \frac{1}{log_2(j)}},$$

where N is the number of shapes in the database. Thus, the DCG for a descriptor is between zero and one with better retrieval performance corresponding to DCG values closer to one. The distinction score for a descriptor $x_{i,j,k}$ associated with position $p_{i,j}$ and scale s_k is the DCG score calculated for the retrieval list when using $x_{i,j,k}$ as a query.

$$D(x_{i,j,k}) = D(p_{i,j}, s_k) \equiv DCG.$$

In our system, we compute and store this measure of distinction for every shape descriptor of every object during an offline processing phase. Comparing two descriptors takes 2.5E-6 seconds on a 2.2GHz computer running Linux, and in general takes $O(Sn^2m^2)$ time to make all pairs of comparisons, where S is the number of scales, n is the number of descriptors per-mesh, and m is the number of meshes in the database. This process takes 37 hours for 128 points at four scales for 907 meshes in the Princeton Shape Benchmark [Shilane et al. 2004]. However, it must be done only once per-database during a batch processing phase, and thus we have not found computing time to be an issue—we computing the distinction for each descriptor once, and then stored it in a file for use multiple times in various applications.

3.4 Mapping to Vertices

The final step of the process is to map the computed measure of class distinction back onto the vertices of the mesh. While this step is not required of all applications (e.g., shape matching), it is useful for several mesh processing tasks (e.g., mesh simplification) that need to have a measure of importance associated directly with every vertex. An alternative to mapping distinction scores from samples on the surface is to calculate shape descriptors at each vertex and calculate distinction directly, but we have typically been working with a vastly smaller number of descriptors than the number of vertices per-mesh.

Our approach to this problem is quite straightforward. We simply model distinction as a mixture of Gaussians. For every vertex, we estimate how distinctive it is by computing a weighted average of the DCG values that have been computed for nearby shape descriptors for a given scale where the weights are determined by the value of a Gaussian function of the distance between the vertex and the middle of the surface region.

Consider mesh X consisting of a set of shape descriptors each with a center position $p \in R^3$ and distinction score defined for each scale, where $D(p, s)$ is calculated as described in Section 3.3. For every vertex v on the mesh of X, distinction is defined as follows.

$$D(v, s) = \frac{\sum_{p \in X} D(p, s) e^{\frac{-\|p-v\|^2}{2\sigma^2}}}{\sum_{p \in X} e^{\frac{-\|p-v\|^2}{2\sigma^2}}}.$$

While using the Euclidean distance instead of geodesic distance ignores connectivity information, it is robust to disconnected meshes. Also, since the regions selected on each shape are generally overlapping and nearby descriptors tend to be similar, it is reasonable to assume that distinction scores change smoothly across a mesh, and we have found that our technique for computing distinction per-vertex works well in practice. In all of the following results, we set $\sigma = 0.1$ times the mesh radius.

4. RESULTS

The methods described in the previous section have been tested on several databases of 3D meshes. In this section, we present images depicting mesh distinction for several examples and investigate (1) which regions of meshes are found to be most distinctive, (2) how mesh distinction compares to previous measures of importance (i.e., saliency), and (3) how sensitive our results are to different parameter settings.

4.1 Mesh Distinction Examples

To help the reader understand which regions are found distinctive by the proposed methods, we show a sampling of images depicting which regions are found to be distinctive for a variety of object classes in a variety of databases. In all images, regions shown in red are the most distinctive, blue regions are least distinctive, and green regions are in the middle. For example, in Figure 4, the ears of the Stanford Bunny are unique to rabbits (red) and thus distinguish the bunny from other classes of animals, while the shape of the body is not very distinctive (blue). Similarly, the head of the wrench, wheels of the vehicles, vase with the plant, and struts of the guitar are important parts for distinguishing each class of objects within the Princeton Shape Benchmark.

Our next example shows the distinctive regions found for three helicopters (all except the right-most image of Figure 5). In this case, shape descriptors were computed at the 0.25 scale at 1,024 points

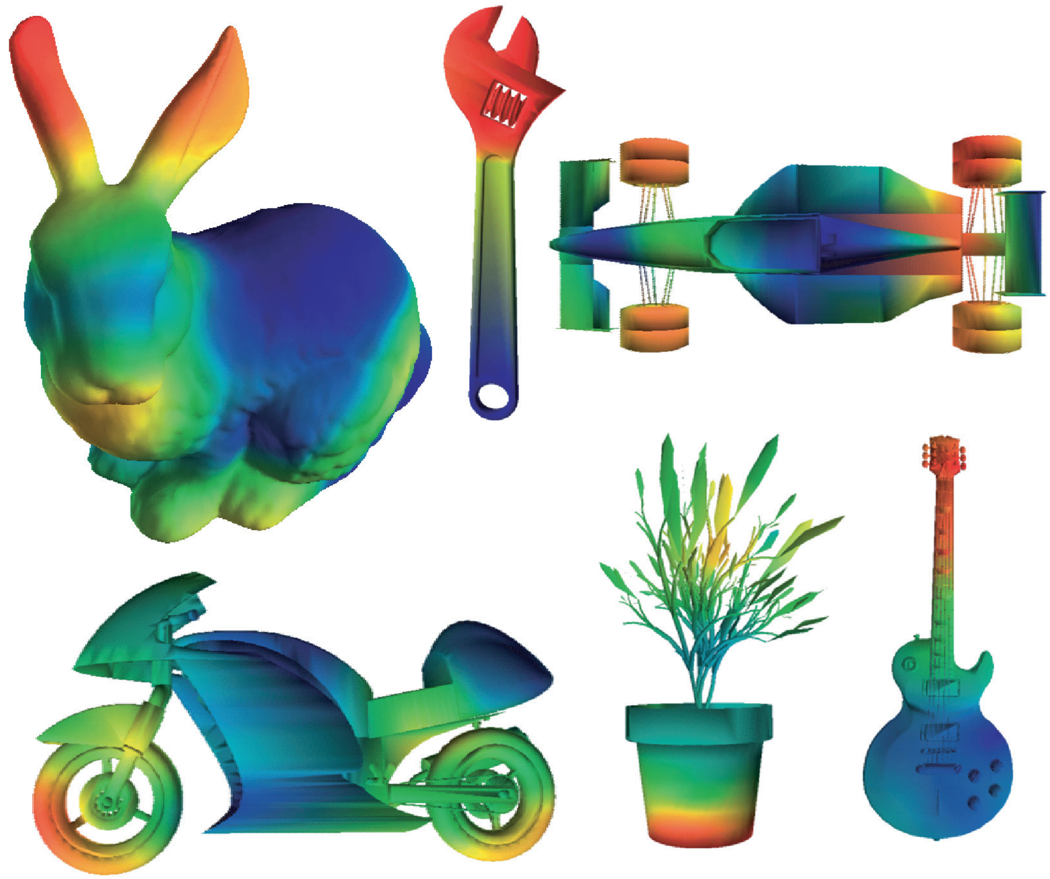

Fig. 4. Distinctive regions of meshes correspond to the important regions that define their class of objects. In all images, regions shown in red are the most distinctive, blue are least distinctive, and green are in the middle.

on each surface, and distinction was measured with respect to a database of flying vehicles dominated by airplanes. The propellers are red (the most distinctive region) in every case, which matches our intuition that the part that distinguishes a helicopter from other flying vehicles is its propeller. For comparison sake, we show the mesh saliency values computed by Lee et al. [2005] for the helicopter third from the left. The areas predicted to be salient by their approach highlight portions of the cockpit and tail due to variations of curvature there and do not detect the importance of the propellers.

A second example shows the regions that distinguish cars from other classes of vehicles (Figure 6). In this case, shape descriptors were computed at the 0.25 scale at 1,024 points, and distinction was measured with respect to a database containing cars, planes, and jeeps. The wheels are most distinctive (red) in all cases, although for the top-left car, the front and rear of the car are equally distinctive probably because of the different aspect ratio of this car relative to other cars. Again, for comparison sake, we show the mesh saliency values computed by Lee et al. [2005] for the top-right car—the regions predicted to be salient by their approach do not correspond as well to distinguishing parts.

A third example shows the regions found to be distinctive for a set of humans standing in a spread-eagle pose (Figure 7). In this case, only 256 shape descriptors were computed for each mesh at the 0.25 scale, and their distinction was measured with respect to all classes in the test database of the Princeton Shape Benchmark (PSB). For thirteen of the fifteen examples, the arms are found to be the most distinctive region. For the other two, the top of the head is most distinctive (those two people have wider arms and a slightly different pose than the others). This result is interesting because the region found to be most distinctive is not obviously what a human might choose as the distinguishing feature of this class at first glance. However, the PSB test database has 92 different classes, including "humans in a standing pose," "humans in a walking pose," "faces," "heads," "skulls," "hands," etc., and thus it is indeed the pose of the arms that best differentiates this class from the others. This example points out both an advantage and disadvantage of our approach: our method can automatically determine the differences between classes in a database, but the distinctive regions may not correspond to semantic parts.

4.2 Effect of Database

A key feature of our formulation for mesh distinction is that the results depend on the database under consideration and how it is partitioned into object classes. In this section, we investigate how the distinctive regions of a surface might be affected by changes in the database.

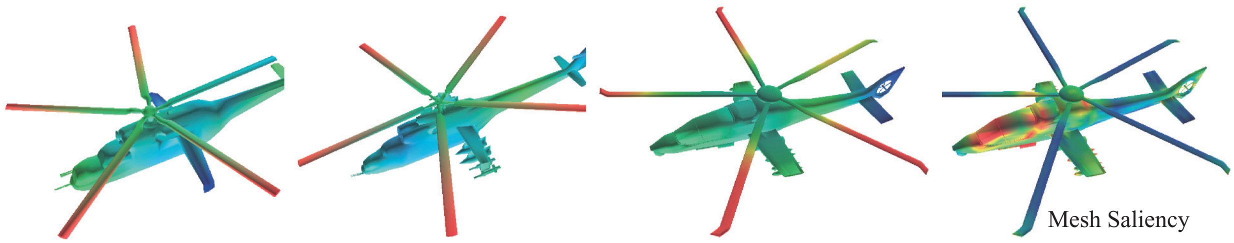

Fig. 5. Visualizations of distinctive regions (red) on three helicopters with respect to a database of flying vehicles. The two right-most images show a comparison of distinction to mesh saliency as computed by Lee et al. [2005] for the same helicopter model.

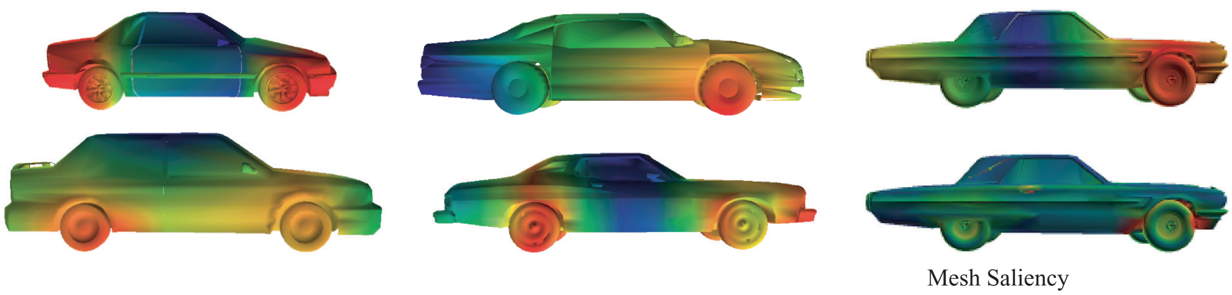

Fig. 6. Visualizations of mesh distinction for five cars as computed with our method at the 0.25 scale. The right-most column shows a comparison between distinction and saliency for the same car model. Note that the tires are consistently distinctive, but not particularly salient.

Fig. 7. Visualizations of mesh distinction for fifteen humans. Note that the distinctive regions for humans in this pose are typically around the elbow area. This region best differentiates this class of objects from other classes of human models in the Princeton Shape Benchmark.

Figure 8 shows four images of the same biplane, the first three of which are colored by mesh distinction as computed with our method at 1024 positions at the 0.25 scale, while the fourth image (bottom right) shows the mesh colored by mesh saliency. The difference between the first three images is only that different databases were used to evaluate the DCG measure during the computation of mesh distinction. The left-most image shows that the wings and tail are most distinctive with respect to the other 91 classes in the Princeton Shape Benchmark. The second image shows that the tail is most distinctive with respect to a smaller database containing other classes of vehicles (cars, jeeps). The third image from the left shows that the struts between the wings and cockpit are most distinctive with respect to a database containing different classes of planes (commercial jets, fighter planes, etc.).

In short, the distinctive area of the biplane changes depending on the database under consideration. This is a very useful feature of our method, as it allows the measure to adapt to finer differences in databases with more similar object classes.

4.3 Effect of Scale

Another factor affecting the distinction of surfaces is the scale (size of the region) covered by each spherical shape descriptor. Figure 9 compares the mesh distinction computed for a model of a dog with

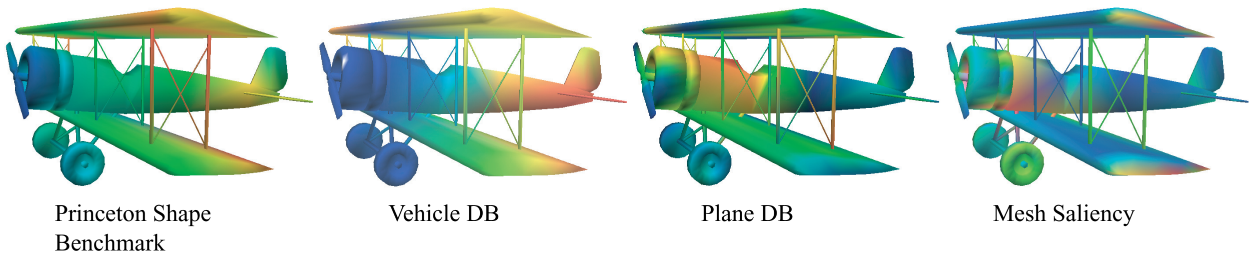

Fig. 8. The distinctive surfaces of the biplane depends on the database under consideration.

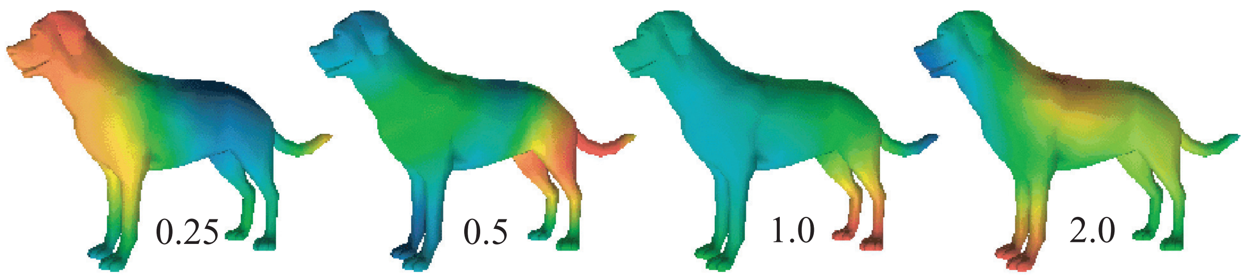

Fig. 9. As the scale of the shape descriptor increases, different surfaces become distinctive.

respect to other quadrupeds with shape descriptors covering 0.25, 0.5, 1.0, and 2.0 of the mesh radius. Please note that as the scale of the shape descriptors vary, the distinctive regions vary. At the smallest scale, the head is the most distinctive region, while at the largest scale, the most distinctive region is centered on the front feet. Compared to other quadrupeds in this database, at a small scale, the head is the most distinguishing local feature. At larger scales, the aspect ratio of dogs versus taller animals such as horses causes an extremity to be the center of the most distinguishing region. This result is typical, smaller scales usually choose a region with a small part having a distinctive shape, while larger scales usually choose an extremity that provides a distinctive center point for describing the global shape of the mesh.

These images highlight that distinction is dependent on the scale selected, and we have specifically preserved these differences as compared to combining distinction in a multiscale method as was calculated for saliency [Lee et al. 2005]. For shape matching purposes, descriptors can be calculated at multiple scales, so it is natural to focus a matching technique on distinctive regions at the appropriate scale. We previously explored shape matching by selecting descriptors using multiscale distinction [Funkhouser and Shilane 2006] and found better results when selecting descriptors at each scale independently.

4.4 Comparing Distinction to Alternative Heuristics

We next investigate how well distinction scores correspond to alternative heuristics across an entire database of shapes. Instead of using shape distinction, there are other possible techniques for selecting important regions on a shape by focusing on properties that are intrinsic to the shape.

Distance. Surfaces of a shape near the center of mass or near an extremity may represent important regions. We have noticed examples such as Figure 1 where positions on the extremity have high distinction which motivates this investigation.

Surface Area. The amount of surface area enclosed within each region varies across the shape depending on the curvature of the shape and scale of the descriptor. We might expect that regions that include a large amount of surface area are more distinctive, while regions that mostly enclose empty space are less distinctive.

Likelihood. Previous projects [Chua and Jarvis 1996; Johnson 2000] have treated shape descriptors as high dimensional feature vectors and selected the least likely descriptors for matching, and thus we have considered shape descriptor likelihood as a good indicator of distinction. We assumed a Gaussian distribution of shape descriptors to calculate the likelihood in a manner similar to [Shilane and Funkhouser 2006].

Saliency. Shape saliency finds the regions of shapes that stick out and are important for visual representation [Lee et al. 2005], so we considered saliency as a property similar to distinction. Saliency scores were calculated by the saliency.exe program (provided by Chang Lee [2005]) on the vertices of a mesh, and saliency scores were interpolated to the centers of the regions.

We compared each of these techniques to distinction for the 907 test models of the Princeton Shape Benchmark across all descriptor scales. We created 256 regions on each shape, created shape descriptors at multiple scales, and calculated distinction for each descriptor. We also calculated for each position the distance from the center of mass of the shape, the amount of surface area (for regions of each scale), the likelihood based on a Gaussian distribution, and the saliency score. We then calculated the correlation score [Pitman 1993] comparing distinction values to each alternative technique at each scale independently:

$$r = \frac{1}{(n-1)\sigma_x \sigma_{D(M_j,s_k)}} \sum_{i=1}^{n} (x_i - \overline{x})(D(p_{i,j}, s_k) - \overline{D(M_j, s_k)}),$$

where x_i is one of the alternative techniques calculated at position $p_{i,j}$ of mesh M_j, and $\overline{D(M_j, s_k)}$ is the average distinction score over the n regions of M_j at scale s_k. We found that, in all cases, the correlation score was between -0.04 and 0.07 where scores

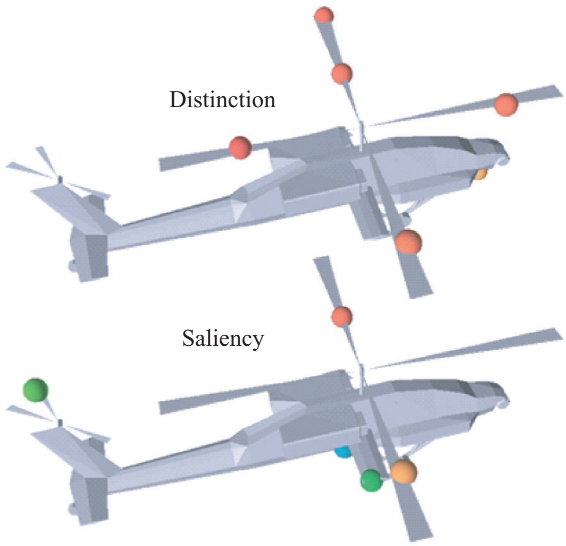

Fig. 10. Descriptors are selected based on distinction, likelihood, saliency, or are random selection. The coloring of the sphere is based on distinction scores, indicating that descriptors with poor distinction scores are selected with the other techniques. A similar number of descriptors are selected for all four techniques, although some appear on the backside of the mesh.

closer to either -1 or 1 indicate a linear relationship (negative or positive, respectively), and values close to zero indicate little or no association. These correlation scores for the 1.0 scale indicate that neither distance ($r = -0.04$), surface area ($r = 0.07$), likelihood ($r = 0.04$), nor saliency ($r = 0.03$) correlates well to distinction.

While this study only considers a linear relationship between distinction and other properties, it is clear that each property is unable to consistently predict which shape surfaces match within a class and to distinguish shapes from different classes.

5. SHAPE MATCHING APPLICATION

The motivating application for shape distinction is shape-based retrieval of 3D meshes, an important problem for many applications in computer graphics, mechanical CAD, molecular biology, medicine, and other fields. For this application, it is common to compute shape descriptors for many points on a surface, find correspondences between them, and model shape similarity based on the distances between corresponding shape descriptors (possibly with geometric deformation terms based on the alignment of correspondences) [Belongie et al. 2002; Berg et al. 2005; Chua and Jarvis 1996; Frome et al. 2004; Gelfand et al. 2005; Johnson and Hebert 1999; Novotni et al. 2005; Shan et al. 2004; Funkhouser and Shilane 2006]. In this section, we investigate whether it might be possible to improve the performance of local shape matching by focusing the match on distinctive surfaces of the meshes.

Our general strategy is to represent every target mesh in the database only by its most distinctive shape descriptors [Shilane and Funkhouser 2006]. By focusing on distinctive shape descriptors, there can still be thousands of descriptor correspondences to consider even though many shapes in the database are poor matches to the query. We use a priority-driven search algorithm to focus the search efforts on the most promising matches, which is more fully described in Funkhouser et al. [2006]. To briefly review, the system maintains a priority queue of partial matches between a subset of the descriptors on the query with descriptors from target shapes. The error metric incorporates not only the L^2 difference between corresponding descriptors across all scales, but also the differences between interfeature distances and surface normals. Only partial matches that can lead to the best possible matches are popped from the queue and extended to create a full correspondence of size K, and thus the system is able to find the optimal matches while focusing on only a small percentage of all possibilities.

Within this matching framework, we investigate methods for selecting a small subset of important descriptors to represent each target shape. The focus of our study is to compare distinction to alternative techniques for selecting important descriptors. Several previous projects [Frome et al. 2004; Mori et al. 2001] randomly selected descriptors on the mesh surface. The least likely shape descriptors have been used for matching by Chua and Jarvis [1996] and Johnson [2000] under the assumption that rare descriptors correspond to important shape features. Other projects [Gal and Cohen-Or 2006; Novotni et al. 2005] have focused on salient regions where the shape sticks out or has variable curvature. By analyzing the distinction of the target shapes' descriptors during an offline phase, we hope to improve the selection step.

We have implemented a simple feature selection algorithm that greedily chooses the most distinctive features on a mesh with the constraint that no two features are too close to one another. Our algorithm iteratively selects the most distinctive remaining descriptor whose position is not closer than a Euclidean distance threshold, D, to the center of any previously selected descriptor. This process avoids selecting many descriptors near each other on the mesh and provides an easy way to reduce the subset size by increasing the distance threshold D. While this algorithm is rather simple and could be improved, it provides a basis for comparing feature distinction with other measures of importance (e.g. saliency) in shape-based retrieval. We compare our technique of using the most distinctive descriptors against the alternatives of selecting descriptors randomly or based on the likelihood or saliency. Figure 10 shows regions selected on the same helicopter model using the four techniques. The

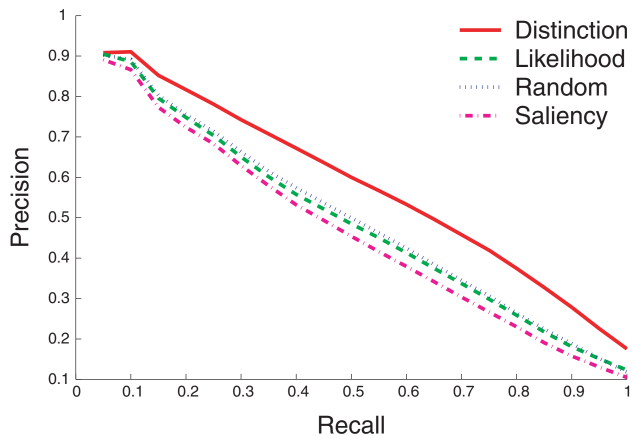

Fig. 11. Precision-recall plot for the training set of the Princeton Shape Benchmark.

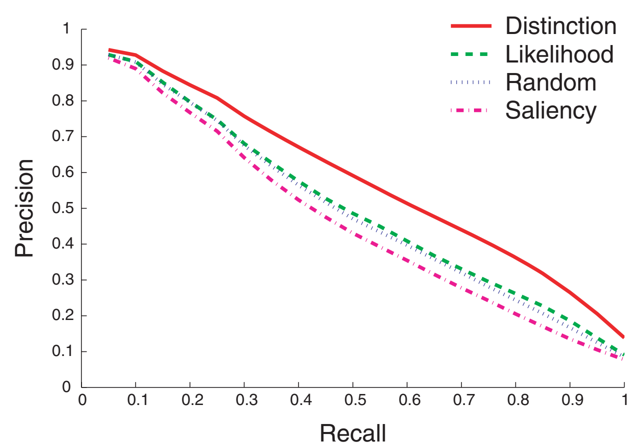

Fig. 12. Precision-recall plot for the test set of the Princeton Shape Benchmark comparing distinction to other techniques for selecting descriptors used in shape matching.

color of the sphere centered on each region indicates the distinction score associated with the region, where red spheres indicate higher distinction scores. Selecting regions based on likelihood, saliency, or random selection leads to representing meshes with regions that perform poorly in retrieval tasks. A similar number of regions were selected for the helicoper in all cases, but some selected regions are not visible in Figure 10 because they appear on the backside of the mesh.

In this study, 128 shape descriptors were computed for every mesh and matched against a subset of the descriptors from every other mesh. In our experiments, we explored a range of values for K and D and found that the relative performance of the three techniques was consistent across all settings. In the following discussion, we selected values for K and D that optimized the retrieval performance independently for descriptors selected based on likelihood, saliency, or distinction scores as well as descriptors selected randomly.

Figures 11 and 12 show precision-recall plots of retrieval results achieved with the proposed method during a leave-one-out study with the training and test sets of the Princeton Shape Benchmark (PSB). Retrieval statistics are also shown in Table I. Column 1 lists the method used for selecting descriptors for retrieval and Column 2 lists the number of descriptors K selected during matching. Columns 3–6 and 7–10 show several measures of retrieval performance, including the nearest neighbor classification rate (NN), the average percentage of correct matches with the first- and second- tiers of matches (1-Tier and 2-Tier, respectively), and the average DCG for each query (in all cases, higher scores are better, see Shilane et al. [2004] for details).

Looking at both the plots and tables, the first result to notice is that selecting features based on distinction provides better retrieval performance than selecting them based on saliency, likelihood, or at random. When considering multipoint matching ($K = 3$) with all four scales on the PSB test set, the nearest neighbor (NN) score for distinction is 74.3% as compared to 66.5% for both likelihood and random and 61.6% for saliency, and across all metrics distinction outperforms likelihood, random, and saliency. For single descriptor matching, $K = 1$ at the 1.0 scale, distinction (NN = 62.4%) also beats likelihood, random, and saliency with NN scores of 55.5%, 55.6%, and 53.4% respectively. Distinction shows a sizable improvement considering that differences between algorithms in other studies are often a few percentage points [Shilane et al. 2004]. While the numbers change somewhat for the PSB training set, the qualitative results are the same.

Besides investigating feature selection methods, we also compared retrieval with distinct descriptors to two other retrieval methods that provide an informative comparison for retrieval performance. The most common shape matching technique [Bustos et al. 2005; Tangelder and Veltkamp 2004] is to use a single shape descriptor centered at the centroid of each shape with a region size large enough to include the entire shape (centroid). Matching a single descriptor at the 1.0 scale on the surface of a shape has better retrieval performance than using the centroid with DCG scores of 54.4% for centroid as compared to 65.3%, 57.8%, 57.8%, and 57.4% for distinction, likelihood, random, and saliency, respectively, on the test set. The DCG score for Distinction increases to 70.6% when matching with three descriptors combined at each of four scales. Of course, this improved retrieval performance comes at some cost (retrieval time of 1.2 seconds versus 3 milliseconds), but we believe that surface descriptors are preferable when retrieval performance is critical.

We next compared our technique of selecting distinctive descriptors versus a best-case method where an oracle selects a single descriptor from the surface of the query shape across all scales. For each query shape, the single descriptor with the highest distinction score (calculated during preprocessing) was selected to use as the query, and the closest match to each target in the database was found. Although this process is not usually possible in a real application (since the class of the query is generally unknown), it provides an upper bound on the retrieval performance possible with surface descriptors. The oracle technique has a NN score of 89.5% on the test set, which dramatically outperforms all other selection techniques we have considered. This result suggests that future work should focus on improving the selection of descriptors from the query shape and that using surface descriptors for shape matching has the potential to achieve highly accurate retrieval results.

We also investigated how often the most distinctive region exists at a particular scale. Table II shows the percentage of time a descriptor from each scale was selected by the oracle technique. On both the PSB training and test sets, descriptors from every scale were selected as the most distinctive, though the 1.0 scale was selected most often.

Table I.
Selecting the most distinctive descriptors from the target set improves retrieval relative to selecting based on likelihood, saliency, or at random. Retrieval improves when using several local descriptors on the query (K = 3 in this result) as compared to using a single descriptor. Using the most distinctive descriptors improves over using a single global descriptor (centroid), while there is still room to improve upon these results since a single descriptor selected by an oracle outperforms any other technique. All experiments are with meshes from the training and test sets of the Princeton shape benchmark.

		PSB Training Set				PSB Test Set			
Descriptor Selection	K	NN (%)	1-Tier (%)	2-Tier (%)	DCG (%)	NN (%)	1-Tier (%)	2-Tier (%)	DCG (%)
Distinction	3	75.4	47.1	58.9	72.4	74.3	45.5	57.0	70.6
Likelihood	3	64.9	37.4	49.2	64.4	66.5	37.0	49.0	64.2
Random	3	68.2	39.0	50.0	65.8	66.5	36.0	47.7	63.4
Saliency	3	61.7	35.1	46.6	62.7	61.6	33.5	44.5	60.9
Distinction	1	68.0	42.4	54.5	68.6	62.4	39.0	51.3	65.3
Likelihood	1	55.8	31.3	43.6	59.5	55.5	29.7	40.9	57.8
Random	1	56.3	31.8	43.6	59.9	55.6	30.0	41.4	57.8
Saliency	1	54.6	30.6	42.8	59.1	53.4	29.2	40.6	57.4
Centroid	1	54.1	28.6	38.1	57.0	53.3	26.3	35.1	54.4
Oracle	1	92.6	54.6	63.4	81.1	89.5	53.5	63.3	79.7

Table II.
With the Oracle selection method, the most distinctive feature was selected to represent each query shape across all scales. Every scale was used for matching, though the 1.0 scale was selected most often for both the PSB training and test databases.

	Scale			
Database	0.25	0.5	1.0	2.0
PSB Training	18.3%	23.4%	38.5%	19.8%
PSB Test	20.9%	27.0%	34.4%	17.6%

6. OTHER APPLICATIONS

Finding distinctive regions on a mesh may be useful for applications beyond shape matching. In this section, we consider two applications: mesh simplification and icon generation.

6.1 Mesh Simplification

For many applications, it is important to create a simplified version of a complex mesh. For example, consider an online parts catalog where images of many tools are shown onscreen at the same time. To improve rendering times, the mesh representing each tool might be simplified. However, to preserve object recognition and emphasize differences within a large collection of meshes, the distinctive features of each tool should be simplified less than the rest of the mesh.

Most techniques for mesh simplification have focused on minimizing the geometric error in a simplified mesh (e.g., Garland and Heckbert [1997]), while others have attempted to minimize errors in the rendered images. In particular, Lee et al. [2005] used their estimation of mesh saliency to weight vertices in a quadrics-based simplification scheme. We follow this work by weighting vertices instead with our mesh distinction scores. Since surface distinction identifies parts that are consistent within a class and distinct from other classes, we expect the simplification algorithm to preserve distinctive features better than other approaches. While features that are salient to the human visual system may not necessarily be preserved with our approach, less distinguishing features will be simplified which, under extreme simplification, will produce a mesh caricature.

To review, quadric error simplification works by contracting the edge that has the least quadric error. The quadric error for each vertex is a measure of how far that vertex has moved during simplification. We augment this basic algorithm by adjusting the error for each edge based on how distinctive its two vertices are. If D_v is the distinction of mesh regions mapped onto vertex v as described in Section 3.4, then the new error for every edge e is $E_e = D_{v_1} Q_{v_1} + D_{v_2} Q_{v_2}$. To accentuate the difference between distinctive and nondistinctive regions, however, the D for the lower 65% of vertices was set to the minimal D value. After each edge is collapsed, the new vertex is assigned an error that is the maximum of the two vertices collapsed so that distinctive regions preserve their scores without being averaged with nearby areas.

Simplification results achieved with this method are shown in Figures 13 and 14. For the hammer example shown in Figure 13, descriptors were computed for 1024 regions at scale 0.25, and their distinction was computed within the context of a database containing four hammers among nineteen meshes representing other classes of tools (screwdrivers, wrenches, and shovels). For a database of tools, the distinguishing features are generally at the functional end of the object away from the handle. The mesh for this hammer was then simplified from 1,700 triangles to 300 triangles (Figure 13) using the distinction weighted error metric. Note that the head of the hammer is the most distinctive region of the mesh and remains well-preserved. For comparison sake, we show simplifications to the same triangle counts achieved using Garland's standard quadric error and using Lee's method of weighting the quadric error by mesh saliency in Figure 13. Note that our method preserves the head of the hammer, the most distinctive part, better than these other methods.

Figure 14 shows similar results achieved when simplifying the mesh for a horse. In this case, the head was found to be most distinctive in the context of a database containing four horses among five other classes of quadrupeds (rabbit, dog, feline, and pig). So, when the mesh was simplified from 97K triangles to 2K triangles, the fine details in and around the horse's head are well-preserved, while the body is greatly simplified. Since the competing methods do not identify these important regions of the horse as strongly, they provide more simplification to the head, while preserving detail in the creases of the body.

6.2 Icon Generation

With the increasing size of 3D model collections, quickly presenting models to a user is an ongoing problem. Whether the application is

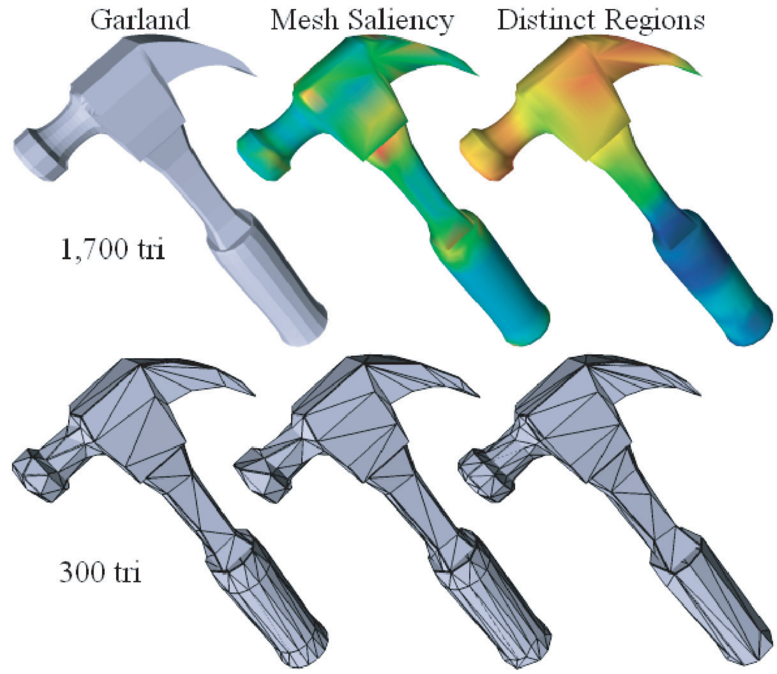

Fig. 13. Simplification results using Garland's method, mesh saliency, and distinctive regions. Notice that more details are preserved in the head of the hammer by focusing on distinctive surfaces.

viewing a catalog of objects or presenting retrieval results in a search engine, the important features of shapes must be shown clearly to the user, perhaps with icons. Focusing the icon on the features that distinguish different classes of shapes could help increase comprehension.

Most previous work on icon generation has focused on the problem of viewpoint selection, that is, positioning a camera oriented towards the center of the object. Vázquez et al. [2001] selected the position that maximized the entropy of the viewed mesh where the optimal view would see all of the faces of the mesh with the same relative projected area. Blanz et al. [1999] studied the preferences of users and found that views at an angle to the major axis were selected. Using their own definition of mesh saliency, Lee et al. [2005] selected views that maximized the visibly salient regions.

We have developed a method for generating icons that displays only the most distinctive region of a mesh. Our algorithm is quite simple. After computing shape descriptors at the 0.25 scale for many points on the mesh, we select the single most distinctive position with respect to a chosen database. We then produce an image of the mesh zoomed in such that the view frustrum exactly covers that most distinctive region. The camera angle is chosen by the computer automatically with one of many possible heuristics or by interactive user control.

We find that this simple method produces useful icons for many classes of objects. For example, Figure 15 shows automatically generated icons for six shapes in the Princeton Shape Benchmark. Note that the images show the region of the object that best distinguish it from other classes of objects, and thus might be more readily recognized than images of the entire object. For many meshes (such as the turtle, wrench, and car), large and recognizeable features are visible in the icon. Showing a limitation of our approach, the biplane icon is focused on the tail region because that region distinguished planes from many other classes of shapes, but perhaps the tail is not the most semantically important feature to humans. However, it should be clearly stated that our measure of distinction is based on 3D shape matching not 2D image matching, and thus it is not guaranteed that the regions determined to be most distinctive by our method will match the ones most visually recognizable by a human. Nonetheless, we find that our simple method based on mesh distinction produces good icons in most cases.

7. CONCLUSION AND FUTURE WORK

In summary, we have defined distinctive regions of a 3D surface to be those whose shape provides the best retrieval performance when matched against other shapes in a database of objects partitioned into classes. This definition produces measures of distinction that adjust to the types of classes in the database and provides shape information at multiple scales. For a number of examples, we have shown that the most distinctive parts are consistent within a class and typically correspond to identifiable parts of a surface. We have also shown how distinctive regions can be used to improve shape matching as well as guide several applications in computer graphics including mesh simplification and icon generation.

There are several strengths and weaknesses of our method that should be considered and addressed in future work. First, the shape descriptor (HSD) used in our implementation is not the most descriptive possible. Recent work by Funkhouser et al. [2006] has shown that the Fourier Shape Descriptor has better retrieval performance than the HSD. However, a strength of our approach is that it is independent of a particular shape descriptor. So, in future work,

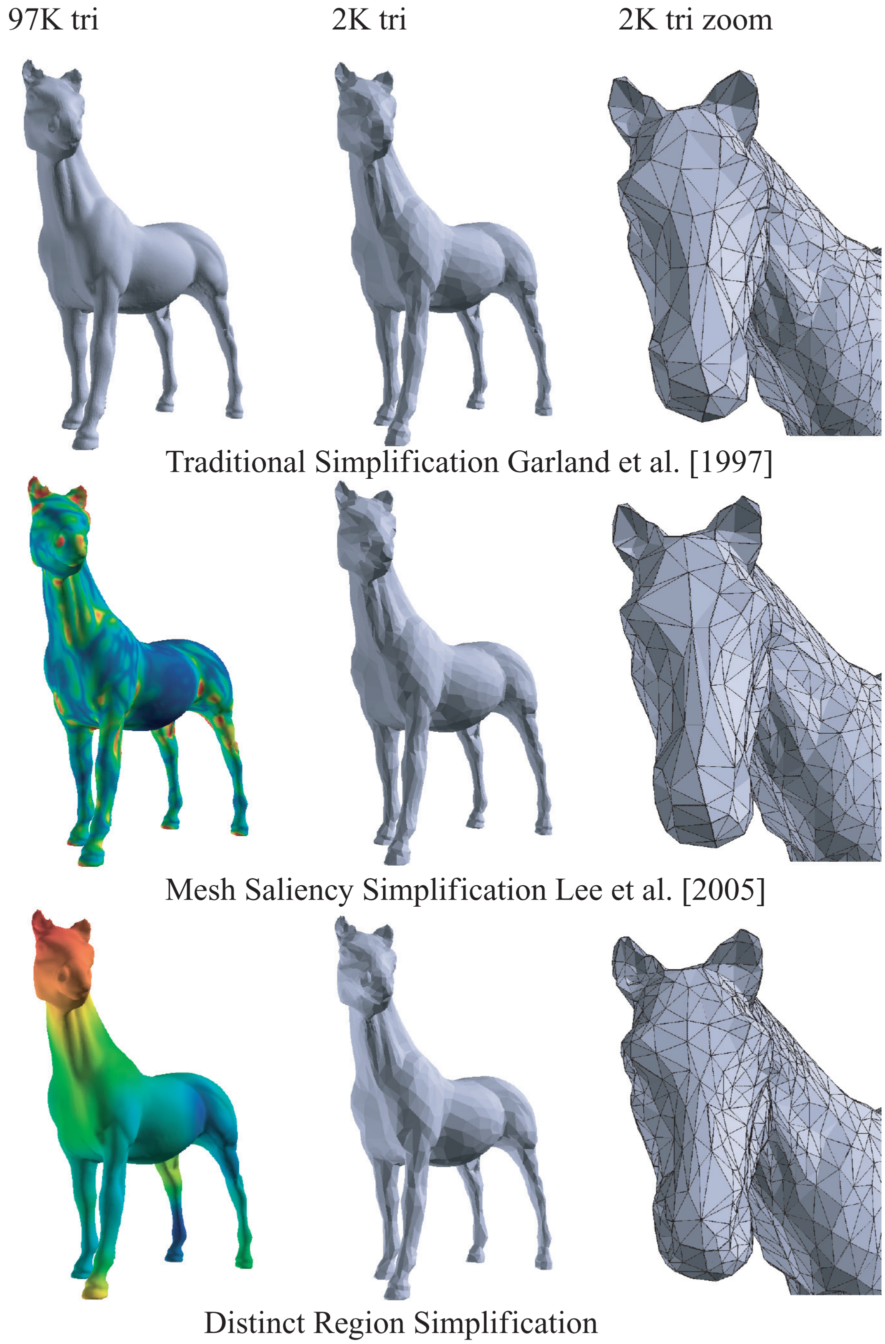

Fig. 14. Simplification results using Garland's method, mesh saliency, and distinctive regions. Notice that details of the eyes and nose are better preserved using our method, while using mesh saliency areas are preserved throughout the horse's body.

we intend to investigate using more descriptive (but perhaps more expensive) shape descriptors to define mesh distinction.

Second, the retrieval measure (DCG) used in our implementation is slow to compute; it requires a full retrieval list to be constructed for every shape descriptor. Calculating distinction for the shape database takes approximately 2.75 minutes per-shape, while calculating saliency takes 4.3 seconds per-model and less than a second is required to calculate likelihood or descriptors selected randomly. Since distinction can be calculated during the preprocessing of the database, we believe it is a worthwhile step to dramatically improve retrieval performance. As discussed in Section 3.3, an alternative to DCG is to use a measure that only requires a fixed number of nearest neighbors to be found (e.g., E-measure), whose computation can be accelerated by indexing.

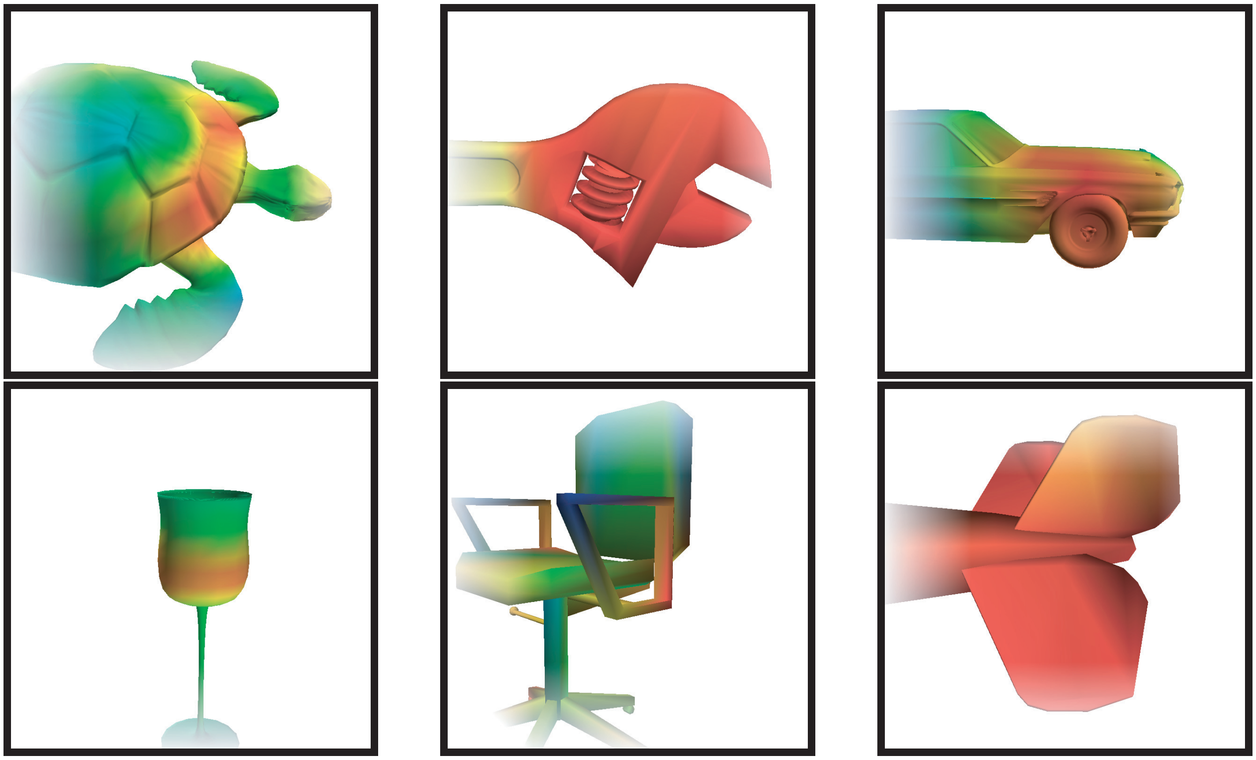

Fig. 15. Icons showing the most distinctive surface region for each mesh.

The third area for future work is to study methods for incrementally updating distinction values as a database changes. With our current definition, as shapes are added or removed, distinction scores must be recalculated for the entire database. By changing the definition of distinction (e.g. an approximation to DCG), it may be possible to dramatically speed up the time to insert or remove shapes while preserving accurate distinction scores. A distinction measure that only requires the first few retrieval results (such as First-Tier) could be updated efficiently using an indexing structure built on bounding regions that include the nearest neighbors for each descriptor. The trade-offs with an approximation scheme involve bounding the error as well as the possible overhead of storing extra information with each shape.

Fourth, our implementation has focused on distinction in 3D surface shape. While this is well-motivated for some applications (e.g., shape matching), perhaps other approaches based on distinction in 2D images of 3D shapes would be better for others (e.g., icon generation and mesh simplification). Our algorithms for icon generation and shape matching are first steps; we believe that there is a wealth of new ways to utilize mesh distinction in these and other applications in computer graphics.

ACKNOWLEDGMENTS

The authors would like to thank Doug DeCarlo, Szymon Rusinkiewicz, Joshua Podolak, and the Princeton Graphics group for useful discussions about this project.

REFERENCES

BELONGIE, S., MALIK, J., AND PUZICHA, J. 2002. Shape matching and object recognition using shape contexts. *IEEE Trans. Patt. Analy. Mach. Intell. 24,* 24, 509–522.

BERG, A., BERG, T., AND MALIK, J. 2005. Shape matching and object recognition using low distortion correspondence. In *IEEE Computer Vision and Pattern Recognition (CVPR)*.

BLANZ, V., TARR, M., BUELTHOFF, H., AND VETTER, T. 1999. What object attributes determine canonical views? *Perception*. 575–599.

BUSTOS, B., KEIM, D., SAUPE, D., SCHRECK, T., AND VRANIĆ, D. 2005. Feature-based similarity search in 3D object databases. *ACM Comput. Surv. 37,* 4, 345–387.

CHUA, C. AND JARVIS, R. 1996. Point signatures: A new representation for 3D object recognition. *Int. J. Comput. Vision 25,* 1, 63–85.

COPAS, J. 1983. Regression, prediction, and shrinkage. *J. Royal Statist. Soc.* Series B (Methodological) *45,* 3, 311–354.

DUDA, R., HART, P., AND STORK, D. 2001. *Pattern Classification*, 2nd ed. John Wiley & Sons, New York, NY.

FRANK, I. E. AND FRIEDMAN, J. H. 1993. A statistical view of some chemometrics regression tools. *Technometrics 35,* 2 (May), 109–135.

FRINTROP, S., NCHTER, A., SURMANN, H., AND HERTZBERG, J. 2004. Saliency-based object recognition in 3D data. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*. Sendai, Japan.

FROME, A., HUBER, D., KOLLURI, R., BULOW, T., AND MALIK, J. 2004. Recognizing objects in range data using regional point descriptors. In *European Conference on Computer Vision (ECCV)*. 224–237.

FUNKHOUSER, T., MIN, P., KAZHDAN, M., CHEN, J., HALDERMAN, A., DOBKIN, D., AND JACOBS, D. 2003. A search engine for 3D models. *Trans. Graph. 22,* 1, 83–105.

FUNKHOUSER, T. AND SHILANE, P. 2006. Partial matching of 3D shapes with priority-driven search. In *Symposium of Geometry Processing*.

GAL, R. AND COHEN-OR, D. 2006. Salient geometric features for partial shape matching and similarity. *ACM Trans. Graph. 25,* 1 (Jan.), 130–150.

GARLAND, M. AND HECKBERT, P. S. 1997. Surface simplification using quadric error metrics. In *Proceedings of SIGGRAPH 1997*. 209–216.

GELFAND, N., MITRA, N., GUIBAS, L., AND POTTMAN, H. 2005. Robust global registration. In *Symposium on Geometry Processing*. Vienna, Austria.

HASTIE, T., TIBSHIRANI, R., AND FRIEDMAN, J. 2001. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer-Verlag, Berlin, Germany.

HOFFMAN, D. D. AND SINGH, M. 1997. Salience of visual parts. *Cognition 63*.

HOWLETT, S., HAMILL, J., AND O'SULLIVAN, C. 2004. An experimental approach to predicting saliency for simplified polygonal models. *Proceedings of the 1st Symposium on Applied Perception in Graphics and Visualization*.

IYER, N., JAYANTI, S., LOU, K., KALYANARAMAN, Y., AND RAMANI, K. 2005. Three dimensional shape searching: State-of-the-art review and future trends. *Comput.-Aid. Design 37,* 5, 509–530.

JARVELIN, K. AND KEKALAINEN, J. 2000. IR evaluation methods for retrieving highly relevant documents. In *23rd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*.

JOHNSON, A. 2000. Surface landmark selection and matching in natural terrain. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. Vol. 2. 413–420.

JOHNSON, A. AND HEBERT, M. 1999. Using spin-images for efficient multiple model recognition in cluttered 3-D scenes. *IEEE PAMI 21,* 5, 433–449.

KAZHDAN, M. 2004. Shape representations and algorithms for 3D model retrieval. Ph.D. thesis, Department of Computer Science, Princeton University.

KAZHDAN, M., FUNKHOUSER, T., AND RUSINKIEWICZ, S. 2003. Rotation invariant spherical harmonic representation of 3D shape descriptors. In *Symposium on Geometry Processing*. 167–175.

KORTGEN, M., PARK, G.-J., NOVOTNI, M., AND KLEIN, R. 2003. 3D shape matching with 3D shape contexts. In *7th Central European Seminar on Computer Graphics*.

LACHENBRUCH, P. AND GOLDSTEIN, M. 1979. Discriminant analysis. *Biometrics 35,* 1 (March), 69–85.

LEE, C. H., VARSHNEY, A., AND JACOBS, D. W. 2005. Mesh saliency. In *ACM SIGGRAPH*. ACM Press, New York, NY, 659–666.

LEIFMAN, G., KATZ, S., TAL, A., AND MEIR, R. 2003. Signatures of 3D models for retrieval. In *the 4th Israel-Korea Bi-National Conference on Geometric Modeling and Computer Graphics*. 159–163.

LI, X. AND GUSKOV, I. 2005. Multi-scale features for approximate alignment of point-based surfaces. In *Symposium on Geometry Processing*. Vienna, Austria.

LOWE, D. 1999. Object recognition from local scale-invariant features. In *Proceedings of the International Conference on Computer Vision*. Vol. 2. 1150–1157.

LOWE, D. 2004. Distinctive image features from scale-invariant keypoints. *Int. J. Comput. Vision 60,* 2 (Nov.), 91–110.

MCLACHLAN, G. J. 1992. *Discriminant Analysis and Statistical Pattern Recognition*. Wiley-Interscience, New York, NY.

MILLER, R. G. 1997. *Beyond ANOVA: Basics of Applied Statistics*. Chapman & Hall/CRC, New York, NY.

MORI, G., BELONGIE, S., AND MALIK, J. 2001. Shape contexts enable efficient retrieval of similar shapes. In *IEEE Computer Vision and Pattern Recognition (CVPR)*.

NOVOTNI, M., DEGENER, P., AND KLEIN, R. 2005. Correspondence generation and matching of 3D shape subparts. Tech. rep. CG-2005-2, Universität Bonn.

OSADA, R., FUNKHOUSER, T., CHAZELLE, B., AND DOBKIN, D. 2002. Shape distributions. *ACM Trans. Graph. 21,* 4, 807–832.

PITMAN, J. 1993. *Probability*. Springer-Verlag, Berlin, Germany.

SHAN, Y., MATEI, B., SAWHNEY, H. S., KUMAR, R., HUBER, D., AND HEBERT, M. 2004. Linear model hashing and batch ransac for rapid and accurate object recognition. In *IEEE International Conference on Computer Vision and Pattern Recognition*.

SHILANE, P. AND FUNKHOUSER, T. 2006. Selecting distinctive 3D shape descriptors for similarity retrieval. In *Shape Modeling International*. 108–117.

SHILANE, P., MIN, P., KAZHDAN, M., AND FUNKHOUSER, T. 2004. The Princeton Shape Benchmark. In *Shape Modeling International*. 167–178.

SPHARMONICKIT 2.5. 1998. Fast spherical transforms: Spharmonickit. http://www.cs.dartmouth.edu/~geelong/sphere/.

TANGELDER, J. AND VELTKAMP, R. 2004. A survey of content based 3D shape retrieval methods. In *Shape Modeling International*. 145–156.

VÁZQUEZ, P.-P., FEIXAS, M., SBERT, M., AND HEIDRICH, W. 2001. Viewpoint selection using viewpoint entropy. In *Proceedings of the Vision Modeling and Visualization Conference (VMV'01)*. Aka GmbH, 273–280.

VRANIC, D. V. 2003. An improvement of rotation invariant 3D shape descriptor based on functions on concentric spheres. In *IEEE International Conference on Image Processing (ICIP'03)*. Vol. 3. 757–760.

Received October 2006; revised April 2007; accepted April 2007