Shape Measure for Identifying Perceptually Informative Parts of 3D Objects

Sreenivas Sukumar, David Page, Andrei Gribok, Andreas Koschan, and Mongi Abidi
Imaging, Robotics, and Intelligent Systems Laboratory, The University of Tennessee
{ssrangan, dpage, agribok, akoschan, abidi} @utk.edu

Abstract

We propose a mathematical approach for quantifying shape complexity of 3D surfaces based on perceptual principles of visual saliency. Our curvature variation measure (CVM), as a 3D feature, combines surface curvature and information theory by leveraging bandwidth-optimized kernel density estimators. Using a part decomposition algorithm for digitized 3D objects, represented as triangle meshes, we apply our shape measure to transform the low level mesh representation into a perceptually informative form. Further, we analyze the effects of noise, sensitivity to digitization, occlusions, and descriptiveness to demonstrate our shape measure on laser-scanned real world 3D objects.

1. Introduction

Perceptual organization refers to the basic capability of the human visual system to derive relevant groupings and structures from 3D objects without prior knowledge of its contents [1]. The traditional approach to such organization in computer vision research relies on formulation of perceptual heuristics of symmetry, clustering, connectivity, attention and many such features that considerably reduce the search space for image understanding. However, the main focus over the last few years in verifying perceptual theories and extending them to computer vision systems has been within the domain of 2D images. With the development of 3D range sensing hardware and processing capabilities, we no longer have to deal with insufficient information of the 3D world in 2D images. Our goal in this paper is to represent 3D objects into visual parts and use psychological evidence to identify among the segmented parts, the ones that convey more shape information towards the recognition of the object.

For example, consider the segmented crank model in Figure 1. If we assume that we do not have an *apriori* definition of a crank and think about which part of the crank grabs immediate visual attention, most observers would point to the curved handle. What characteristic of the handle considered as a 3D surface contributed to visual saliency? Can this idea of perceptual complexity or perceptual attention be quantified as a surface feature?

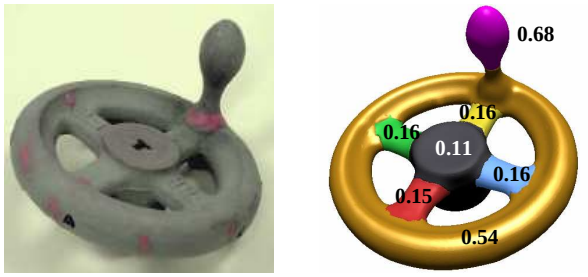

(a) (b)

Figure 1: (a) Photograph of the crank. (b) Segmented laser-scanned 3D mesh model of the crank with the CVM characterizing each part. We see the CVM identifying the curved handle and the toroidal body as informative parts in the model.

With the CVM, that can be understood as the entropy of surface curvature, we formulate a potential solution to quantify this idea of perceptual attention. We derive the definition of the CVM from Attneave's seminal work [2] that lists the informational aspects of visual perception. His conclusions based on the experimental study [3] on how human perception judges visual complexity along several variables such as curvedness, symmetry, compactness among others further emphasizes that variation in curvature is a rich source of visual information content. In using Attneave's observations towards the definition of the CVM as surface shape descriptor, we summarize the contributions of this paper as follows:

- Formulation of the CVM as a surface feature that combines surface curvature and information theory,
- Incorporation of automated bandwidth selection in curvature density estimation to avoid user selected parameters to compute entropy,
- Representation of 3D objects with the CVM identifying perceptually significant parts with the

help of a perception-driven mesh segmentation algorithm.

With that introduction, we very briefly discuss state-of-the-art surface features on 3D triangle mesh datasets in Section 2. We then proceed with the formulation and implementation details of the CVM in Section 3 and follow the theory with experimental results and analysis in Section 4. In Section 5, we use the CVM on laser-scanned real 3D objects and compare our approach to another graph-description method. Based on the experiments, we conclude by identifying potential applications of the CVM in Section 6.

2. Related Work

In this section, we will very briefly explain the psychophysical basis of our work and also establish the need for a surface feature based on perceptual principles towards describing 3D objects.

Psychophysics: The basis of the algorithm that we propose in this paper derives out of Attneave's psychophysical results [3]. His experiments on 2D contours, revealed a regression equation relating the perceived complexity of a polygonal contour to its symmetry, angular variability, and the logarithm of the number of turns of the contour. The logarithmic dependency reveals the similarity with Shannon's logarithmic relation of complexity towards description in information theory.

Further, Attneave, states that his psychological evidence is better explained using information theory treating perception as an economical description problem. Palmer includes Attneave's conclusions on 2D contours and also suggests structural information theory as the potential direction towards implementing Gestalt's ideas for 3D shape representation [4]. He also emphasizes that extending Attneave's observations for 3D objects is a non-trivial task. Recently, Todd tried to extend Attneave's experiments on random 3D surfaces to identify potential sources of perceptual attention [5]. His experiments lead to the conclusion that 3D objects are best characterized by the lines connecting local curvature extrema and the local variation in the curvature of the surface itself. His study further agrees with Huffman and Singh's [6] minima rule of part decomposition and Biderman's recognition by parts [7]. The use of curvature as the basis feature for segmentation and description is psychophysically valid. We are hence motivated to represent 3D objects by first decomposing objects into parts and then using the CVM as a shape measure to describe the relative complexity of these parts.

3D features: The vast majority of research in perceptual feature selection is with 2D images. Most 2D features do not extend directly to 3D triangle meshes because, with intensity images perceptual cues are defined on edges and boundary contours extracted from the image. But with 3D data, we have to consider the underlying surface shape for characterizing the visual cues. We do find methods for quantifying global mesh complexity [8] and mesh saliency [9] targeting shape indexing and level-of-detail analysis. But these methods are not able to describe local spatial characteristics required for object description. In the following paragraphs, we briefly discuss such 3D features proposed for object description and recognition.

On range images that encapsulate 3D structure, Dorai and Jain [10] have defined two curvature-based 3D features known as the shape index and curvedness for their Curvedness-Orientation-Shape Map on a Sphere (COSMOS). They represent ideal range images as maximal patches of constant shape index. They use their definition of a shape index that is also in a sense a measure of shape complexity towards a graph description of 3D objects. They further formulate a shape spectrum on these patches as an object description scheme.

We also find several other features in shape signatures like 3D Fourier descriptors [11], local feature histograms [12] and spin images [13]. Additionally, Kortgen et al. [14] extend the 2D shape contexts [15] to 3D features. Khotanzad and Hong [16] have extended Duda and Hart [17] to include a subset of 3D moments as features on point clouds. Following a similar trend, Cybenko et al. [18] use second order moments, spherical kernel moment invariants, surface area, and other metrics as features for their object recognition system. On the other hand, histogram-based methods [19, 20] generate features based on the statistical properties of the object. For example, Besl [19] has considered the crease angle at each edge between two triangles in a mesh as a feature and Osada et al. define several shape functions to define their shape distributions [20]. Specifically, their D2 shape function which is the distance between two random points chosen on the surface of the object acts as a dominant surface descriptor on their experiments with 3D models. Their histogram distributions from shape functions are similar to our curvature densities with the CVM.

Most other features in the literature are well summarized by Cordone et al. [21] in his recent survey on 3D shape similarity through feature matching. We identify that very few of these features cater to the part-based description for object recognition and hence direct the application of the CVM as a region feature specifically targeting the part-based perception model of 3D objects.

3. Our Algorithm

Our work uses triangle mesh datasets (commonly used output format from 3D scanners) that are simply a collection of vertices and triangles approximating the 3D surface of digitized 3D objects. Although this representation is useful for visualization, the low level description is often inadequate for computer vision tasks. Higher level symbolic and informative description hence is a necessity. In aiming to move to such a higher level description, Stankiewicz [22] proposes three models of perception: feature-based, alignment-based and part-based and further argues that the part-based approach to object description as the intuitive and fruitful attempt to object recognition. We benefit from his conclusions and the psychophysical evidence to implement a graph description scheme towards the demonstration of the CVM. As mentioned earlier, we are inclined towards segmentation based on perceptual principles. We have used a mesh segmentation algorithm similar to [23] that makes use of Hoffman and Singh's minima rule of perception to decompose an object into a graph network of surface patches. Our idea is to use these segmented surface patches and apply the CVM to characterize them.

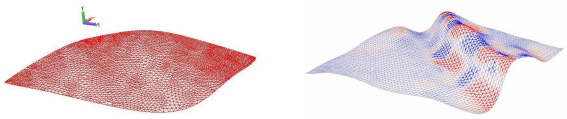

CVM = 0.022 CVM = 0.43

Figure 2: Which of these surfaces contains more shape information? (The surfaces are textured based on the Gaussian curvature.) We see that the CVM is able to numerically differentiate perceived shape complexity of 3D surfaces.

We now present the theory and formulation for the CVM algorithm. As with Attneave's [3] observations, the motivation for CVM is the observation that shapes with smoothly varying curvature attribute less to the informative attention-grabbing aspect of perception and hence appear less significant than shapes with significant variation in curvature. For instance, a flat surface has uniform constant curvature. Since there is no variation in its surface curvature, we say that such surfaces possess little or no shape information. However, a surface that models the terrain of a mountain for example, will have significant variation in curvature. We suggest that such surfaces possess more shape information. Our argument can be better understood using Figure 2. In addition, we suggest that surfaces with repetitive curvature patterns have lower shape information than that surface with no patterns. Our observations agree with Lowe's [1] and Attneave's theory [2] about how the probability of accidental occurrence contributes to visual significance.

We emphasize that these observations relate to the notions of Shannon entropy and his definition of entropy as measure of information in digital communication theory [24]. Shannon's established information framework assumes a closed alphabet set of symbols to compute entropy as the average length to describe a symbol. The symbols with low probability of occurrence convey more information and require more number of bits for description. We use this characteristic of Shannon's framework as a measure to describe complexity of description and apply it to 3D surfaces. Our approach to extend the framework to a 3D triangle mesh can be easily understood using the simple block diagram as shown in Figure 3.

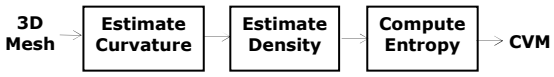

Figure 3: This block diagram shows how to compute CVM.

We base the CVM on curvature since curvature is invariant to choices in coordinate frames, viewing angles, and surface parameterizations, which are desirable properties for feature selection. As an infinitesimal local feature, curvature lacks sufficient generalization to serve as a 3D feature in our computer vision context. As a result, CVM aims to characterize curvature over a region (or patch) of a surface. To compute CVM, the first step is to estimate curvature on a triangle mesh. But, a triangle mesh is a singular surface with infinite curvature concentrated at the vertices and edges of the mesh and zero curvature on the faces of the triangles that make up the mesh. These singular values of curvature are not useful since our interest lies in the curvature values of the original smooth (non-singular) surface that the mesh approximates.

Of the several methods to estimate curvature on triangle meshes [25] that we tested for the CVM, we chose the Gauss-Bonnet approach as a trade-off between computation speed and estimation accuracy. This method uses the umbrella neighborhood of triangles immediately adjacent to a vertex to estimate the Gaussian curvature at that vertex. This method is also known as the loss of angle approach since we use the angles subtended by the edges emanating from the vertex of interest. If we consider a vertex v, then we can define the angle α_i for the corner of each triangle adjacent to v. From Gauss-Bonnet theorem of differential geometry, we can estimate the Gaussian

curvature of the underlying smooth surface at v as Equation 1, where the summation is over the umbrella neighborhood and A is the accumulated area of the triangles around v.

$$\kappa = \frac{3\left[2\pi - \sum_{i=0}^{a-1} \alpha_i\right]}{A} \tag{1}$$

Now, we have curvature estimates of a digitized continuous surface. These curvature values are not a quantized and normalized entity on surfaces. Also, the curvature estimates cannot be considered as a discrete symbol towards computation of the entropy. Histograms that assume an origin and bin width will not give us accurate estimates of the underlying curvature density for any arbitrary surface. Hence, we suggest a data driven approach to density estimation using Gaussian kernels.

We make use of kernel density estimators (KDE) [26] as a tool to compute the density function p of the curvature values over the entire mesh. Consider Equation 2 where p is the estimate of the density function, n is the number of vertices in the mesh, h is the bandwidth of interest, G is the kernel function and κ_i is the curvature at vertex v_i. We visualize KDE as a series of 'bumps' placed at each of the n estimates of curvature in the density space. The kernel function G determines the shape of these bumps while the bandwidth h determines their extent. With large data sets (n is large), the choice for G does not have a strong influence on the estimate. We recommend the Gaussian kernel although meshes provide large sample points to make the minimization of the mean-integrated squared error (MISE) for bandwidth estimation easier.

$$p(x) = \frac{1}{nh}\sum_{i=1}^{n} G\left(\frac{x-\kappa_i}{h}\right) \tag{2}$$

$$G(u) = \frac{1}{\sqrt{2\pi}} e^{-\frac{u^2}{2}} \quad \text{such that} \quad \int_{-\infty}^{\infty} G(x)dx = 1 \tag{3}$$

The more significant parameter for accurate and stable estimation is not the kernel but the bandwidth h. Data driven bandwidth optimization approaches [27] such as the distribution scale methods, cross validation, L-stage, plug-in and advanced bootstrap methods theoretically aim to minimize the MISE between the actual density and the computed density. For our algorithm, we use the plug-in method (Equation 4) for optimal bandwidth selection as this method sufficiently reduces the computational overhead compared with commonly used cross validation methods by making good approximations to minimize the MISE. Using a Gaussian kernel further gives us a closed form approximation for the data-driven bandwidth h on our curvature estimates as shown below.

$$h_{opt} = \left[\frac{243\,R(G)}{35\,\mu_2(G)^2 n}\right]^{\frac{1}{5}} \hat{\sigma} \approx 1.06\,\hat{\sigma} n^{-\frac{1}{5}} \tag{4}$$

where $R(G) = \int G(t)^2\,dt, \mu_2(G) = \int t^2 G(t)\,dt$ and $\hat{\sigma}$ is the absolute deviation of the curvature data κ_i. With the Gaussian kernel, the approximation in Equation 4 can be used as a quick method of choosing the bandwidth parameter.

The choice of h in the density estimation avoids the need for any user selected parameters in our algorithm in addition to providing scale invariance. On scale invariance, suppose, we were considering spherical objects of different radii, the value of h automatically selected would adjust itself proportional to the radius, there by quantifying the curvature variation among spheres of different radii to be similar and tending towards the delta function.

The final step in the CVM algorithm is the computation of entropy. Our choice of entropy as the measure of complexity towards description arises out of the fact that non-accidental occurrences, which in our case is unexpected variation in curvature, contributes to more structure information. As discussed earlier, we use Shannon's definition of entropy on the curvature density to quantify that shape information. However, a direct application of Shannon entropy leads to variability relative to different mesh resolutions. In other words, if we directly use Shannon entropy, we are unable to compare two meshes that have different vertex counts. As a result, we normalize entropy to account for resolution.

$$CVM = -\sum p(\kappa) \log_n p(\kappa) \tag{5}$$

where p is the probability density of curvature estimated using the KDE.

The normalization occurs through the base (n) of the log function that allows comparison of similar surfaces at different resolutions in addition to the convenience of the [0, 1] feature space. We have indirectly penalized the measure of complexity based on the number of samples that attributed to the complexity. Also, the CVM falls into the "interval" scale of measurements, where the difference is more significant instead of the ratio.

4. Experiments and Analysis

This section outlines the experimental analysis of CVM as a surface feature. We begin by demonstrating the descriptive capability of the CVM on synthetic hyperbolic patches. In Figure 4, we show 9 surface patches characterized by the CVM. We see that the CVM is able to arrange each of these patches in agreement with the human intuition. On the same set of hyperbolic surface patches in Figure 5, we compare the descriptiveness of the variance of curvature, range, mean of shape index [10] and the variance of the D1 and D2 shape functions [20].

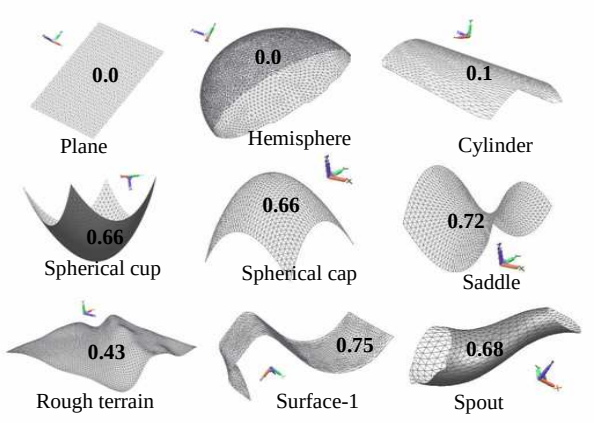

Figure 4: CVM follows human intuition in describing some synthetic surfaces.

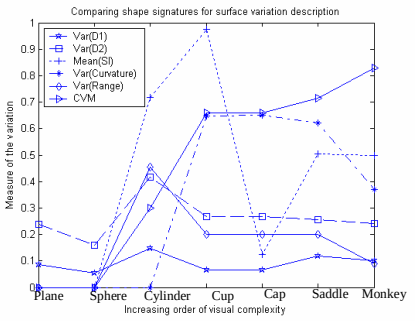

Figure 5: Comparison of 3D features on hyperbolic patches arranged in ascending order of human perceptual complexity.

We observe from Figure 5 that the CVM increases with perceived complexity while the variance of curvature suffers from curvature not being normalized. On the other hand, the shape index performs better than the variance of curvature but does not seem to agree with the human idea of perceptual complexity. Our shape measure tends to perform better in a perceptual sense on these surfaces compared to the D1 and D2 shape functions also. We show later that the CVM along with perceptual mesh segmentation can characterize an object as a minimal set of perceptual parts while the shape index requires more number of patches to describe the same object. We move on to analyze the CVM for noise, occlusions and digitization.

Effect of noise: We conducted two different experiments to analyze the effect of noise. One of them based on a known noise model to find what strength of noise significantly perturbs our measure, while the other was to analyze the effect of scanner noise in comparison with the results from the synthetic data. We conclude from the results in Figure 6, that the CVM is well behaved within acceptable levels of both random noise and scanner noise.

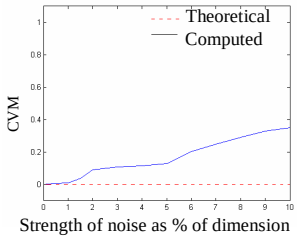

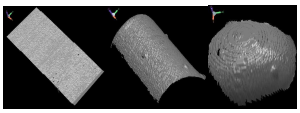

CVM on laser scanned surfaces:
Plane : 0.001 Cylinder : 0.13
Hemisphere: 0.06

(a) (b)

Figure 6: Effect of noise. (a) CVM appears to be sensitive to noise when the strength of the noise corrupting a planar surface increases as a percentage of the longest dimension. (b) CVM values on laser scanned 3D hyperbolic patches are still close to the values from the synthetic data in Figure 4.

Effect of digitization: In Figure 7, we present the variability of the CVM due to effects of surface digitization. The experiment was to find how the number of vertices and the irregularity in meshing affect the CVM. CVM appears to be quite robust to digitization as expected from its formulation. We make this claim from the experiments on a sphere at multiple vertex counts and also a multi-resolution irregularly triangulated saddle-like surface. With the sphere in Figure 7a, we see that as the number of vertices increases, the CVM also approaches the theoretical value. However, we do observe that at low resolution, the curvature estimates are unreliable and corrupt the CVM. The same phenomenon happens with increasing levels of noise as well, because curvature approximation methods on triangle meshes

are usually very sensitive to noise. We are encouraged by the stability of CVM to noise and digitization that we now analyze the effect of occlusions.

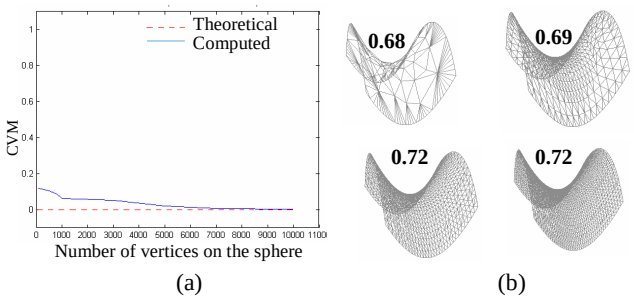

Figure 7: Sensitivity to digitization. (a) Effect of resolution on CVM. CVM converges to a theoretical value as the number of vertices is increased in the mesh. (b) CVM also is not sensitive to irregular triangulation as long as shape is preserved on the saddle-like surface.

Effect of occlusions: Having understood the effect of noise of the IVP SC386 range scanner in our previous experiment, we consider a real world object modeled and integrated from multi-view range images in Figure 8. We use the multi-view range images to understand the behavior of CVM with occlusions. As one should expect from curvature-based features, the CVM can only be a good descriptor as long as the occluded component of an object has no significant new information towards identifying the object. This limitation checks the extensibility of the CVM to scene analysis.

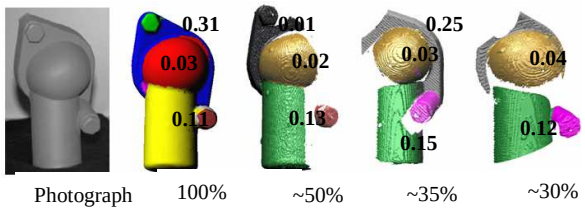

Figure 8: Effect of occlusions. The CVM characterizes segmented surfaces from range images of the water neck object. We indicate in percentage the actual overlap of the range image with the object. CVM reacts on the base of the water neck more than the cylindrical and spherical section as we lose information due to occlusions.

5. Object Description using the CVM

A key assumption in the above development of our CVM is that the triangle meshes model a smooth surface. Hence, we note that the CVM requires smooth patches as a result of mesh segmentation for a descriptive characterization of a 3D object. We show the CVM describing various parts of four objects (a cup, a disc brake, a crank and a water neck) in Figure 9. The higher magnitude of the shape measure corresponds to the most informative part of the object.

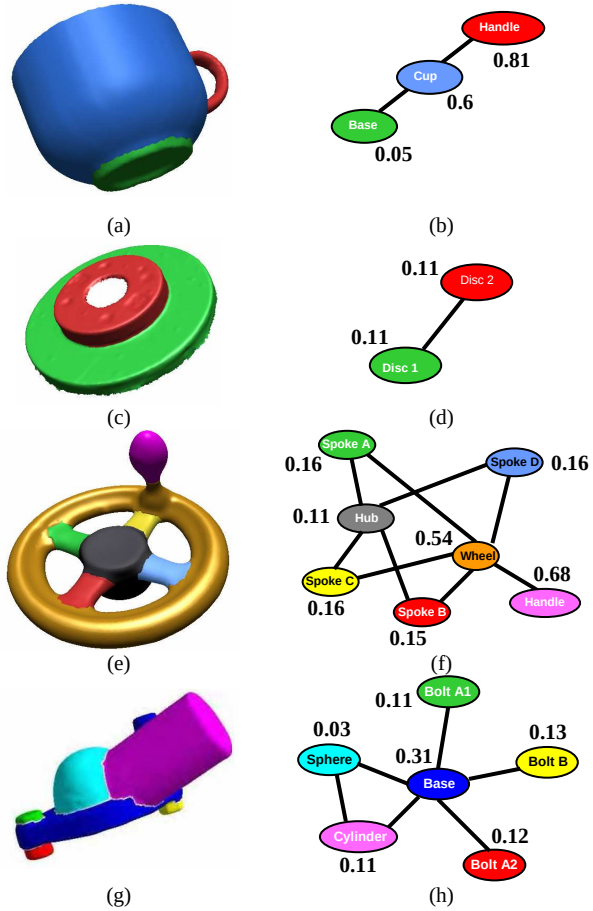

Figure 9: Object description using our method. (a) and (b) A cup: CVM identifies the handle as the informative part. (c) and (d) A disc brake: CVM values hint identical structures. (e) and (f) Crank: The handle is recognized as the visually complex shape. (g) and (h) Water neck: CVM says the base of the object is the one that gives maximum shape information.

We use these four examples also to demonstrate how our approach has reduced the search space compared to a different approach of representing an object as patches described using shape index [10] for the attributed graph matching step before recognition. We extended their framework that was demonstrated on range images to triangle meshes, by first segmenting along crease edges using their measure of

curvedness and further decomposing the object into regions of constant shape index. The regions of constant shape index basically correspond to different hyperbolic shape primitives that best fits the underlying surface. We summarize the comparison in Figure 10. Even on segmented parts, we observed that the shape index characterized larger patches as a combination of smaller patches better, while CVM was able to deal with sufficiently larger patches.

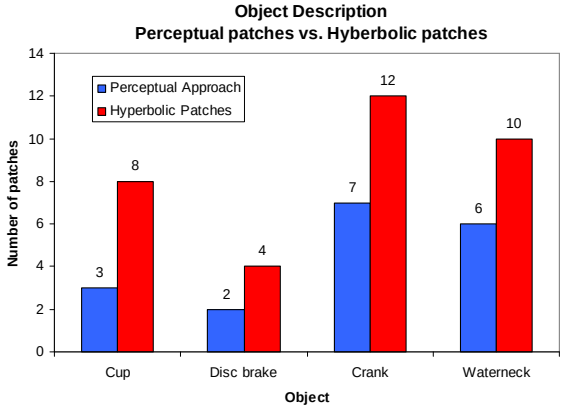

Figure 10: Comparing perceptual organization with constant shape index organization of 3D objects as hyperbolic patches.

5. Conclusions and Future Work

We have demonstrated our CVM as a surface feature on triangle mesh datasets that appears to agree with the human intuition of visual attention. By first segmenting the object into perceptual parts and using our definition of the CVM to characterize each one of these parts, we are able to transform a mesh-based representation of a 3D object to a more symbolic, descriptive and informative form. The graph network preserves the structural relationship of parts of the 3D object and the CVM describes the components in a relative sense of visual salience quantifying perceived complexity.

We would like to note that the CVM has assumed no prior information about the context of shape significance. Hence, CVM will not be able to characterize contextual saliency. For example, if a small hemispherical defect on the crank was more significant, in recognizing the defect, our approach to perceptual organization will only be able to identify the defect in its structural form (in the graph) but not by the value of the CVM.

Although CVM improves the description of patches larger in size than the constant-shape maximal patches of Dorai and Jain [10], we note that extremely large patches can lead to an over generalized density function. This drawback surfaces with some real world objects that do not easily segment out into perceptual parts. Also, we note that the CVM being a scale invariant feature requires simple scale preserving attributes such as bounding box dimensions, surface area etc. of the segmented parts for an efficient representation of the object. Our preliminary results using a graph matching algorithm [28] on a small database of objects encourages our future directions in transforming the proposed object description framework into a recognition framework.

Acknowledgements

This work was supported by the DOE URPR under grant DOE-DEFG02-86NE37968 and the DOD/RDECOM/NAC/ARC Program under grant W56HZV-04-2-0001.

References

[1] D. G. Lowe, *Perceptual Organization and Visual Recognition*, Kluwer Academic Publishers, Boston, 1985.

[2] F. Attneave, “Some informational aspects of visual perception”, *Psychology Review*, 61(3):183-193, 1954.

[3] F. Attneave, “Physical determinants of the judged complexity of shape”, *Journal of Experimental Psychology*, 53(4):221-227, 1957.

[4] S. E. Palmer, *Vision Science-Photons to Phenomenology*, Cambridge, MA: MIT Press, 1999.

[5] J. T. Todd, “The visual perception of 3D shape”, *Trends in Cognitive Science*, 8:115-121, 2004.

[6] D. D. Hoffman, M. Singh, “Salience of visual parts”, *Cognition*, 63:29-78, 1997.

[7] I. Biederman, “Recognition-By-Components: A Theory of Human Image Understanding”, Psychological Review, 94(2):115-147, 1987.

[8] D. L. Page, A. Koschan, S. Sukumar, B. Abidi, and M. Abidi, “*Shape analysis algorithm based on information theory*”, in Proc. of the Int’l. Conf. on Image Processing, 1: 229-232, 2003.

[9] C. H. Lee, A. Varshney, and David Jacobs, “Mesh Saliency”, *ACM SIGGRAPH*, 24(3):659-666, 2005.

[10] C. Dorai and A. K. Jain, “COSMOS-A representation scheme for 3D free-form objects”, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(10):1115-1130, 1997.

[11] D. V. Vranic, “An improvement of rotation invariant 3D shape descriptor based on functions on concentric spheres”, *in Proc. Int'l Conf. on Image Processing*, 3: 757-760, 2003.

[12] G. Hetzel, B. Leibe, P. Levi, and B. Schiele, "3D object recognition from range images using local feature histograms", *in Proc. Int'l Conf. on Computer Vision and Pattern Recognition*, 2:394-399, Dec. 2001.

[13] A. Johnson and M. Hebert, "Using spin images for efficient object recognition in cluttered 3D scenes", *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 21(5):433-449, 1999.

[14] M. Kortgen, G. J. Park, M. Novotni, and R. Klein, "3D shape matching with 3D shape contexts", *in Proc. Central European Seminar on Computer Graphics*, Apr. 2003.

[15] S. Belongie, J. Malik, and J. Puzicha, "Shape matching and object recognition using shape contexts", *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 24(4): 509-522, 2002.

[16] A. Khotanzad and Y. H. Hong, "Invariant image recognition by Zernike moments", *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 12(5):489-497, 1990.

[17] R. M. Duda and P. E. Hart, *Pattern Classification*, New York, NY: John Wiley and Sons, Inc., 1973.

[18] G. Cybenko, A. Bhasin, and K. Cohen, "Pattern recognition of 3D CAD objects", *Smart Engineering Systems Design*, 1:1-13, 1997.

[19] P. Besl, "Triangles as a primary representation: Object recognition in computer vision", *Lecture Notes in Computer Science*, 994: 191-206, 1995.

[20] R. Osada, T. Funkhouser, B. Chazelle, and D. Dobkin, "Shape distributions", *ACM Trans. on Graphics*, 21(4): 807-832, 2002.

[21] A. Cardone, S. K. Gupta, and M. Karnik, "A Survey of Shape Similarity Assessment Algorithms for Product Design and Manufacturing Applications", *ASME Journal of Computing and Information Science in Engineering*, 3(2): 109-118, 2003.

[22] B. J. Stankiewicz, "Models of the Perceptual System", *In the Encyclopedia of Cognitive Science*, Macmillan. 2002.

[23] D. L. Page, A. F. Koschan, and M. A. Abidi, "Perception-based 3D Triangle Mesh Segmentation Using Fast Marching Watersheds", *in Proc. Intl. Conf. on Computer Vision and Pattern Recognition*, 2:27-32, 2003.

[24] C. E. Shannon, "A mathematical theory of communication", *The Bell System Technical Journal*, 27:379-423, 623-656, July 1948.

[25] T. Surazhsky, E. Magid, O. Soldea, G. Elber, and E. Rivlin, "A comparison of Gaussian and mean curvature estimation methods on triangle meshes", *in Proc. Int'l Conf. on Robotics and Automation*, (Taiwan), 1021-1026, 2003.

[26] B. W. Silverman, *Density Estimation for Statistics and Analysis*, London, UK: Chapman and Hall, 1986.

[27] M. P. Wand and M. C. Jones, *Kernel Smoothing*, London, UK: Chapman and Hall, 1995.

[28] M. A. Eshera and K. S. Fu, "An image understanding system using attributed symbolic representation and inexact graph-matching", *IEEE Trans. on. Pattern Analysis and Machine Intelligence*, 8(5):604-618, 1986.