Experimental Analysis of Harmonic Shape Images

Dongmei Zhang and Martial Hebert
The Robotics Institute, Carnegie Mellon University

Abstract*

In our early work, a representation called Harmonic Shape Images for 3D free-form surfaces was proposed and applied to solve the surface-matching problem. In this paper, extensive experiments using real data are conducted to analyze the performance of Harmonic Shape Images with respect to discriminability, stability and robustness to resolution and occlusion. The results show that Harmonic Shape Images are an appropriate representation for 3D surface comparison. Examples of surface comparison using real data are presented in the paper.

1. Introduction

Surface comparison is a fundamental issue in computer vision. A large amount of work has been done regarding this issue. The approaches to solving the problem can be classified into two categories: model-based matching[5]-[17] and matching by registration[18]-[21].

In our early paper[23], a representation called Harmonic Shape Images was proposed to represent and compare 3D free-form surfaces. The basic idea of Harmonic Shape Images is to map a 3D surface patch with disc topology to a 2D domain and encode the shape information of the surface patch into the 2D image. This simplifies the surface-matching problem to a 2D image-matching problem. When constructing Harmonic Shape Images, a mathematical tool called harmonic maps is employed to solve the mapping problem between a 3D surface patch with disc topology and a 2D domain.

In this paper, we focus on the experimental analysis of Harmonic Shape Images with regard to the following issues: how discriminative Harmonic Shape Images are, how stable Harmonic Shape Images are and how robust Harmonic Shape Images are with respect to resolution and occlusion. Those issues are important for obtaining consistently good performance when applying Harmonic Shape Images to surface comparison.

A library of twenty surface patches is constructed by extracting patches from real free-form objects. The discriminability of Harmonic Shape Images is evaluated by how different the patches are from one another in terms of shape. Using surface patches of similar shapes, the stability of Harmonic Shape Images is evaluated by how similar those patches are. Patches of the same shape but with different resolutions are compared to one another in order to evaluate the robustness of Harmonic Shape Images with respect to resolution. Similar experiments are conducted to evaluate the robustness to occlusion.

This paper is organized as follows: Harmonic Shape Images are briefly reviewed in Section 2; in Section 3, experimental analysis of Harmonic Shape Images is presented in detail; applying Harmonic Shape Images in surface matching is discussed in Section 4 along with experimental results. Conclusions and future work will be presented at the end of the paper.

2. Harmonic Shape Images

In this section, we will briefly review the concept of Harmonic Shape Images, the generation process and the properties. Please refer to [23] for details.

2.1. General Concept

The 3D free-form surfaces studied in this paper are represented by polygonal meshes. According to [18], a free-form surface S is defined to be a smooth surface such that the surface normal is well defined and continuous almost everywhere, except at vertices, edges and cusps. Comparing two such meshes directly is difficult due to the following reasons: the topology may be different for different objects; the sampling may be different even for the same object; the surfaces may not be complete because of occlusion and clutter in the scene.

The development of Harmonic Shape Images was motivated by the above difficulties. Given a 3D surface S as shown in Figure 1(a), let v denote an arbitrary vertex in S. Let $D(v, R)$ denote the surface patch which has the central vertex v and radius R and has disc topology. $D(v, R)$ is connected and consists of all the vertices in S whose surface distance is less than, or equal to, R. The overlaid region in Figure 1(a) is an example of $D(v, R)$. Its amplified version is shown in Figure 1(b). If the unit disc is selected to be the domain and $D(v, R)$ is mapped onto the domain using certain strategy, then the resultant image $h(D(v, R))$ is called harmonic image as shown in Figure 1(c). The harmonic image preserves the shape and continuity of $D(v, R)$. Because correspondences can be established between the vertices in $D(v, R)$ and the vertices in $h(D(v, R))$, the Harmonic Shape Image of $D(v, R)$, $HSI(D(v, R))$, can be obtained by associating shape attributes, e.g., curvature, at each vertex of $D(v, R)$ with the corresponding vertex in $h(D(v, R))$. Figure 1(d) shows the Harmonic Shape Image of the surface patch in Figure

* This work was supported by NSF Grant IRI-9711853.

0-7695-0062-5/99 $10.00 © 1999 IEEE

1(b). The curvature values are gray-coded. High intensity values correspond to high curvature values.

Harmonic Shape Images can be generated for any vertex on a given surface as long as there exists a valid surface patch at that vertex. Here, a valid surface patch means a connected surface patch with disc topology.

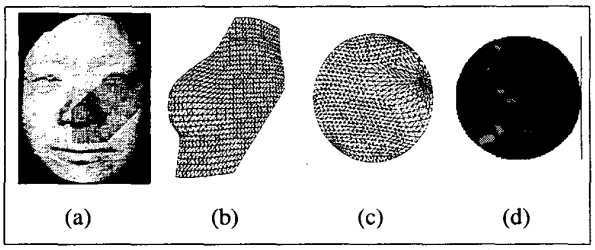

Figure 1: (a) A surface patch (overlaid region) *D(v, R)* on a given surface; (b) Amplified version of *D(v, R)* in (a); (c) The harmonic image of *D(v, R)*; (d) The Harmonic Shape Image of *D(v, R)*.

2.2. Generation of harmonic images

The generation of harmonic images is based on a mathematical tool called harmonic maps[1][2]. Harmonic maps solve the problem of mapping between different metric manifolds by solving partial differential equations. Due to the expensive computational cost in solving partial differential equations and the discrete nature of surfaces we deal with in practice, it is natural to look for an approximation of harmonic maps.

Eck et al proposed an approximation approach to harmonic maps in [4]. Eck's approximation consists of two steps. At the first step, the boundary of the 3D surface patch is mapped onto the boundary of an equilateral triangle that is selected to be the 2D target domain. At the second step, the interior of the surface patch is mapped onto the interior of the equilateral triangle with the boundary mapping as a constraint. Our approach uses the same interior mapping strategy as that of Eck's approach but a different target domain and a different boundary mapping strategy.

2.2.1. Interior Mapping

Let D be a 3D surface patch with disc topology and P be a unit disc in 2D. We use $D(v, R)$ to denote that the central vertex of D is v and the radius of D is R. Let ∂D and ∂P be the boundary of D and P, respecitively. Let v_i^i, $i=1,...,n^i$, be the interior vertices of D. The interior mapping ϕ maps v_i^i, $i=1,...,n^i$, onto the interior of the unit disc P with a given boundary mapping $b:\partial D \rightarrow \partial P$. $\phi(i)$ is obtained by minimizing the following energy functional.

$$E(\phi)=\frac{1}{2}\sum_{\{i,j\}\in Edges(D)} k_{ij}\left\|\phi(i)-\phi(j)\right\|^2 \qquad (1)$$

in which $\phi(i)$ and $\phi(j)$ are the images of the interior vertices i and j of D on P. The values of $\phi(i)$ define the mapping ϕ. k_{ij} serve as spring constants[4]. The minimum of the energy functional $E(\phi)$ can be found by solving a sparse linear least-square system of (1) for the values $\phi(i)$[4].

The 2D image of D after being mapped onto P is named the harmonic image $h(D)$ of D. Examples of D and $h(D)$ are shown in Figure 1(b) and (c) respectively.

2.2.2. Boundary Mapping

The constuction of the boundary mapping $b:\partial D \rightarrow \partial P$ is illustrated in Figure 2.

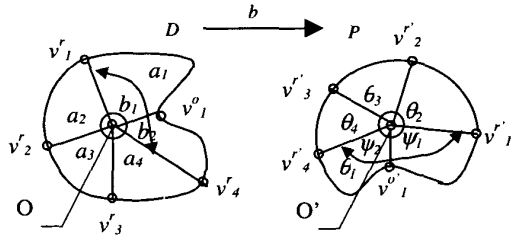

Figure 2: Illustration of the boundary mapping between the surface patch and the 2D domain.

First of all, let us define the vertices and vectors in Figure 2. O is the central vertex of D and O' is the center of P. v_i, $i=1,...,5$ are the boundary vertices of D. D is said to have radius R when the surface distance from any vertex in D to the central vertex O is less than, or equal to, R. For some boundary vertices, e.g., v_i^r, $i=1,...,4$, the surface distance between any of them and the central vertex O is equal to R; for other boundary vertices, e.g., v_1^o, the surface distance is less than R. The vertices in the former case are called radius boundary vertices and the vertices in the later case are called occluded boundary vertices. Radius boundary vertices are determined by the size of the surface patch, while occluded boundary vertices are determined by occlusion (either self occlusion or occlusion by other objects). The vector from the central vertex O to a radius vertex v_i^r is called a radius vector, while the vector from O to an occluded boundary vertex v_j^o is called an occlusion vector.

Now let us define the angles in Figure 2. Angles a_i, $i=1,...,4$ are the angles between two adjacent radius vectors $\overrightarrow{v_i^r O}$ and $\overrightarrow{v_{i+1}^r O}$. Angles b_j, $j=1, 2$, are the angles between two adjacent occlusion vectors, or one occlusion vector and one adjacent radius vector, in an occlusion range. An occlusion range is a consecutive sequence of occlusion boundary vertices except for the first and last ones. For example, (v_4^r, v_1^o, v_1^r) is an

occlusion range. The sum of b_j over an occlusion range is the angle a_i formed by the first and last radius vectors of this occlusion range.

The construction of the boundary mapping consists of two steps. At the first step, the radius boundary vertices are mapped onto the boundary of the unit disc P, which is a unit circle. In Figure 2, v_i^r, $i=1,...,4$, are mapped to $v_i^{r'}$, $i=1,...,4$, respectively. It can be seen that once the angles θ_i are determined, the positions of $v_i^{r'}$ are determined. θ_i is computed as follows:

$$\theta_i = \frac{a^i}{\sum_{k=1}^{n} a_k} 2\pi \tag{2}$$

At the second step, the occlusion boundary vertices in each occlusion range are mapped onto the interior of the unit disc P. For example, in Figure 2, v_1^o, which is in the occlusion range (v_4^r, v_1^o, v_1^r), is mapped onto $v_1^{o'}$. Once the angles ψ_j and the radii r_j are determined, the position of v_1^o is determined. ψ_j are computed as follows.

$$\psi_j = \frac{bj}{\sum_{m=1}^{n} b_m} a_i \tag{3}$$

in which n is the number of angles within the occlusion range and a_i is the angle corresponding to the occlusion range. r_j is defined to be

$$r_j = \frac{dist(r_j, O)}{R} \tag{4}$$

in which $dist(r_j, O)$ is the surface distance between the occlusion boundary vertex v_j^o and the central vertex O. R is the radius of the surface patch D.

It should be noted that Harmonic Shape Images generated using different starting vertices are different by a planar rotation[23]. The rotation difference will be found later by the comparison process. Therefore, the starting vertex can be selected randomly.

2.3. Generation of Harmonic Shape Images

In Section 2.2, we have shown that, given a surface patch D, its harmonic image $h(D)$ can be created using harmonic maps. There is one-to-one correspondence between the vertices in D and the vertices in $h(D)$. Harmonic Shape Images, $HSI(D)$, are generated by associating a shape attribute at each vertex of $h(D)$. In our current implementation, an approximation of the curvature at each vertex is used to generate Harmonic Shape Images. For details about the curvature approximation, please refer to [5].

Figure 1(d) is an example of Harmonic Shape Image in which the curvature values are gray-scale coded. High intensity values represent high curvature values.

2.4. Properties of Harmonic Shape Images

Due to the way that Harmonic Shape Images are generated and the application of harmonic maps in the generation, Harmonic Shape Images have many good properties. First of all, this representation is local. It does not depend on the overall topology of the underlying object. Second of all, the mapping between a given surface patch D and the 2D target domain P is one-to-one and onto. Therefore, when two Harmonic Shape Images match, the correspondences between vertices on the two patches being compared can be established immediately. Furthermore, this mapping constructs a parameterization of D in P. This parameterization allows us to resample the original surface patch, if necessary. In practice, resampling such as raster scanning makes it easy to compare two Harmonic Shape Images. Lastly, Harmonic Shape Images are intrinsic to the shape of the underlying surfaces. Therefore, they are invariant to the pose of the underlying surfaces. They are unique and their existence is guaranteed by the existence of harmonic maps.

3. Experimental Analysis of Harmonic Shape Images

Harmonic Shape Images, their generation and properties have been discussed in Section 2. Considering the application of Harmonic Shape Images, which is surface comparison, the following issues need to be investigated because they are directly related to the performance of surface comparison.

- Discriminability
- Stability
- Robustness to resolution
- Robustness to occlusion

These issues will be discussed in detail in this section.

3.1. Discriminability

For a surface patch in a given scene, its Harmonic Shape Image needs to be representative or discriminative enough in order for its best match to be identified from a library of surface patches. The following experiment has been conducted to illustrate how discriminative Harmonic Shape Images are.

A library of surface patches is created which consists of twenty patches extracted from real free-form objects. Their Harmonic Shape Images are computed and stored in the library as well. Two examples of the surface patches in the library and their Harmonic Shape Images are shown in Figure 3.

Pair comparison of the Harmonic Shape Images has been done among the twenty surface patches in the library. The comparison of two Harmonic Shape Images is implemented as the combination of a planar rotation plus the normalized correlation between the two images(5).

$$R(D_1, D_2) = \max_{0 \le \theta \le 2\pi} r_\theta(HSI(D_1), HSI_\theta(D_2)) \quad (5)$$

High values of $R(D_1, D_2)$ indicate high similarity between surface patches, while low values indicate low similarity. Figure 4 shows the pair comparison result. It can be seen from Figure 4 that the values on the diagonal are 1.0. This is not surprising because all the surface patches are the same as themselves. The values elsewhere are much lower than the values on the diagonal. This indicates that all the twenty patches are quite different from one another in terms of shape.

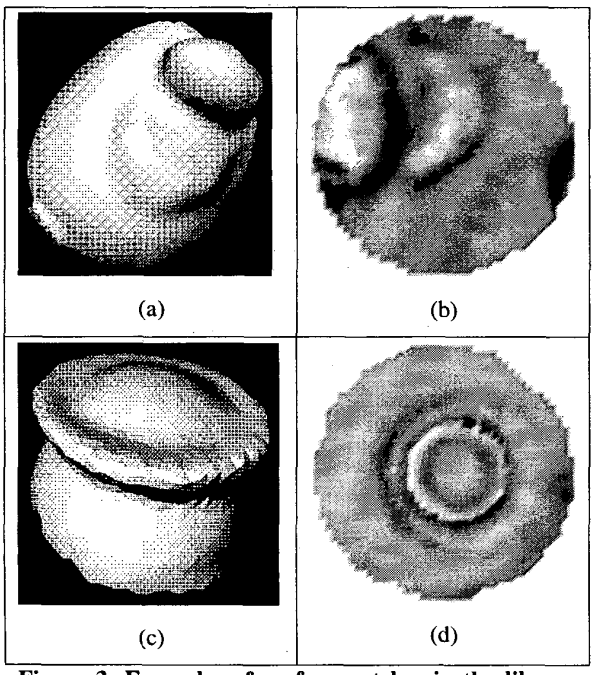

Figure 3: Examples of surface patches in the library and their Harmonic Shape Images.

Given a set of correlation coefficients, in order to quantitatively show how distinctive one Harmonic Shape Image is from other Harmonic Shape Images based on those correlation coefficients $R(D_1, D_2)$, a statistic approach is used to detect the distinctive one as an outlier in that data set[22]. Considering the correlation coefficients $R(D_1, D_2)$ as values of a random variable, the first step of the statistic approach is to transform the set of correlation coefficients $R(D_1, D_2)$ into a normalized distribution according to the following formula:

$$C(D_1, D_2) = \ln \frac{1 + R(D_1, D_2)}{1 - R(D_1, D_2)} \quad (6)$$

in which $R(D_1, D_2)$, which is defined in (5), is the correlation coefficient between surface patches D_1 and D_2. $C(D_1, D_2)$ is defined to be the similarity value between the two patches. The second step is to plot the histogram of the similarity values as shown in Figure 5.

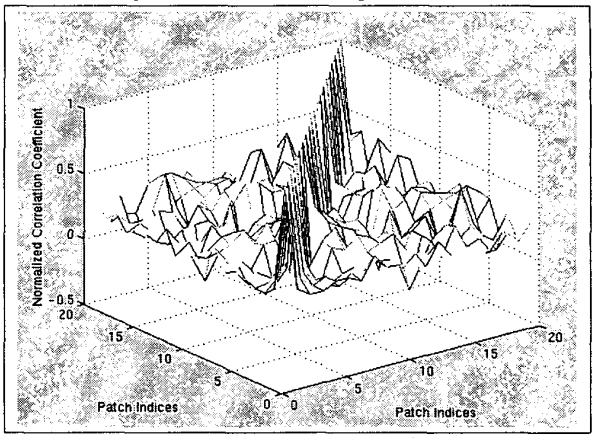

Figure 4: Pair comparison of Harmonic Shape Images among the twenty surface patches in the library. High correlation coefficient values on the diagonal indicate the surface patches are the same as themselves. Low correlation coefficient values elsewhere indicate the patches are quite different from one another.

The third step is to find the outliers in the histogram using the following method[22]: The fourth spread f_s of the histogram is defined as the median of the largest $N/2$ measurements (upper fourth) minus the median of the smallest $N/2$ measurements (lower fourth). Statistical moderate outliers are $1.5fs$ units above (below) the upper (lower) fourth. Extreme outliers are $3fs$ units above (below) the upper (lower) fourth.

Figure 5 shows the histogram of the similarity values of the surface patch D_1(Figure 3(a)) to all twenty patches in the library. The line in Figure 5 indicates the value for extreme outliers. One extreme outlier is found in Figure 5 and that is the similarity value of the surface patch D_1 to itself. Therefore, D_1 is sufficiently different from other patches in the library based on the comparison of their Harmonic Shape Images. This illustrates how discriminative Harmonic Shape Images are. Figure 6 shows another example.

3.2. Stability

In addition to being discriminative, Harmonic Shape Images should be stable as well. This means that surface patches of similar shapes should have similar Harmonic Shape Images. For example, on smooth free-form surfaces, the surface patches of two neighboring vertices should have similar shapes. Therefore, their Harmonic Shape Images should also be similar to each other. This observation is important in surface comparison because in

practice, no two discrete surfaces are sampled exactly the same way. In this case, the surfaces should still be matched using their Harmonic Shape Images.

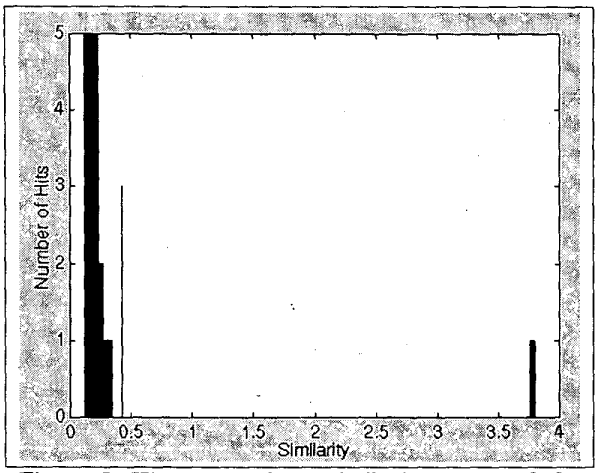

Figure 5: Histogram of the similarity values of the surface patch D_1(Figure 3(a) to all twenty patches in the library.

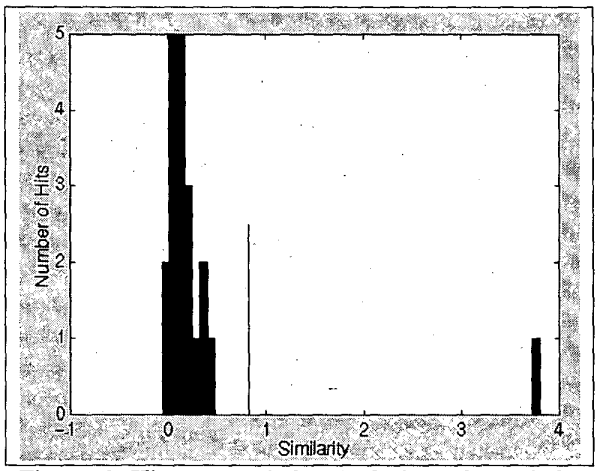

Figure 6: Histogram of the similarity values of the surface patch D_2(Figure 3(c) to all twenty patches in the library.

The following experiment is conducted to illustrate the stability of Harmonic Shape Images. One object model is shown in Figure 7(a). Surface patches are created for each vertex on the model. For example, the patch D_1 in Figure 7(b) is centered at the vertex shown on the head of the duck in Figure 7(a). The patch in Figure 7(c) is centered at the vertex shown on the wing of the duck. Harmonic Shape Images of all the patches are computed and compared to the Harmonic Shape Image of D_1. The correlation coefficients are color-coded and displayed in Figure 7(d), (e) and (f). High brightness corresponds to high correlation coefficient values. It can be seen that the central vertices of those patches that are similar to D_1 are in the neighborhood of the central vertex of D_1(Figure 7(d) and 8). The histogram in Figure 9 shows how those similar patches differ from other patches.

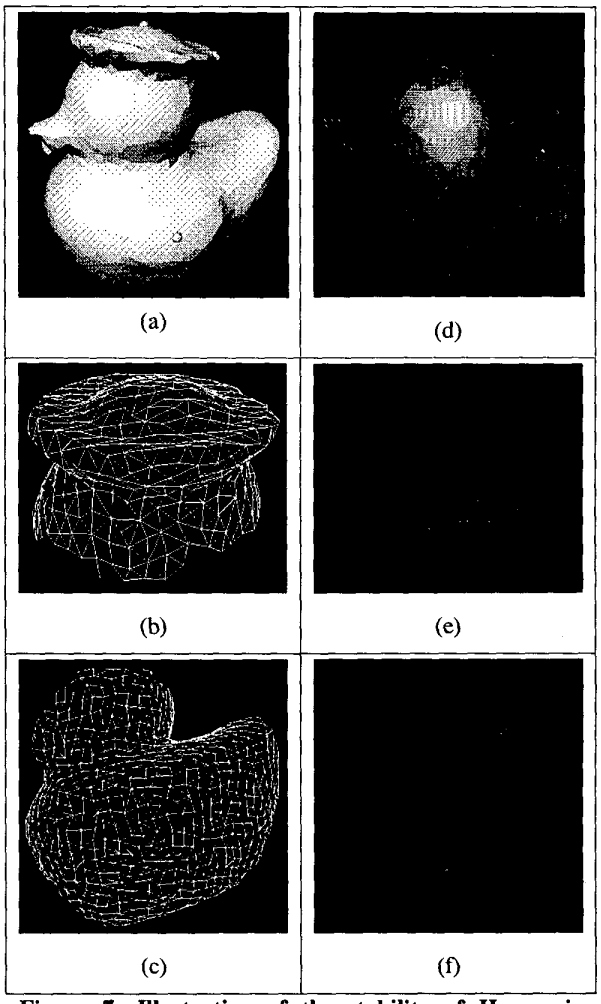

Figure 7: Illustration of the stability of Harmonic Shape Images. (a) An object model; (b) The surface patch centered at the vertex shown on the head of the duck in (a); (c) The surface patch centered at the vertex shown on the wing of the duck in (a); (d) Top view of the color-coded object using the correlation coefficients; (e) Side view; (f) Front view.

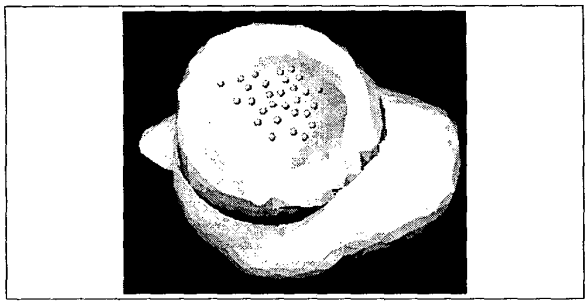

Figure 8: The central vertices of those patches that are similar to the patch in Figure 7(b).

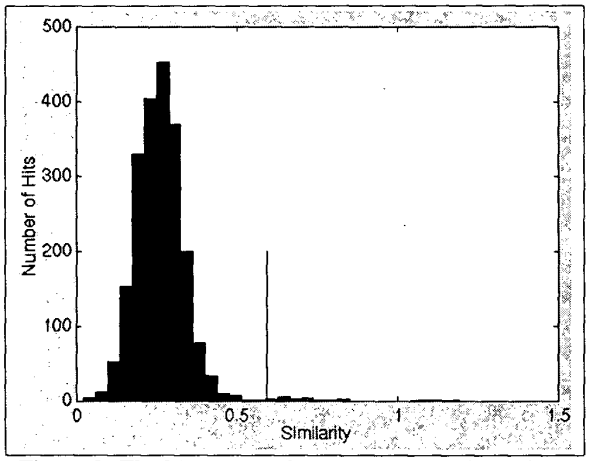

Figure 9: The histogram of the similarity values of the patch D_1 to all the patches on the object shown in Figure 7(a). The extreme outliers have their central vertices in the highlighted area of Figure 7(d) and marked in Figure 8.

3.3. Robustness to Resolution

As discussed in previous sections, Harmonic Shape Images do not depend on any specific sampling strategy, e.g., uniform sampling. For a given surface, as long as the sampling rate is high enough such that the shape of the surface can be sufficiently represented, its Harmonic Shape Image is also good enough for surface matching. It should be noted that the comparison of Harmonic Shape Images does not require the two surface patches have the same sampling frequency. In practice, it is rare for discrete surfaces to have exactly the same frequency. Although different resolutions may introduce noise in creating Harmonic Shape Images, it can be seen from the following experiment that Harmonic Shape Images are robust to this kind of noise.

Seven surface patches of different resolutions are created. They all have the same shape as the patch shown in Figure 3(a). Two of them are shown in Figure 10(b) and (c), respectively. The Harmonic Shape Images of the seven patches are computed and added to the library. Pair comparison of any two Harmonic Shape Images in the library is conducted and the result is shown in Figure 11.

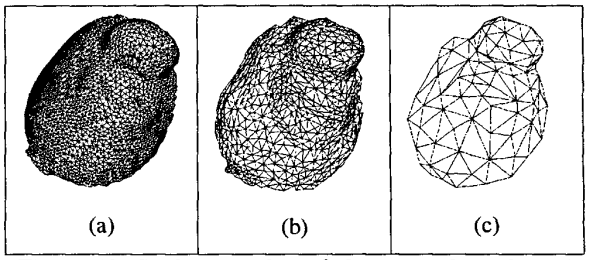

Figure 10: Surface patches of the same shape but of different resolutions. (a) The original patch. Its rendered version is shown in Figure 3(a); (b) The patch with resolution ratio 3:1; (c) The patch with resolution ratio 25:1.

In Figure 11, High correlation coefficient values on the back center indicate that the newly added seven patches are similar to one another. The high values on both the front sides indicate that the seven patches are similar to the first patch in the library. Low values everywhere else indicate that the seven surfaces are quite different from other surfaces in the library.

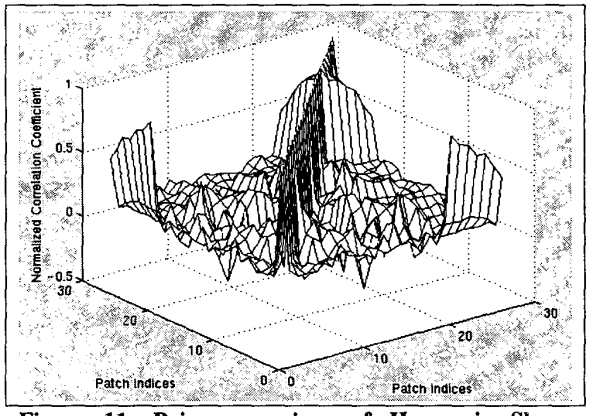

Figure 11: Pair comparison of Harmonic Shape Images among the twenty-seven surface patches in the library. High correlation coefficient values on the back center indicate that the newly added seven patches are similar to one another. The high values on both the front sides indicate that the seven patches are similar to the first patch in the library. Low values everywhere else indicate that the seven surfaces are quite different from other surfaces in the library.

The statistic approach discussed in section 3.1 is used to determine the best match of the seven patches among the original twenty patches in the library. The first patch (Figure 10(a)) is found to be the best match for all seven patches, which is the correct result. Figure 12 shows the recognition result for the patch shown in Figure 10(c).

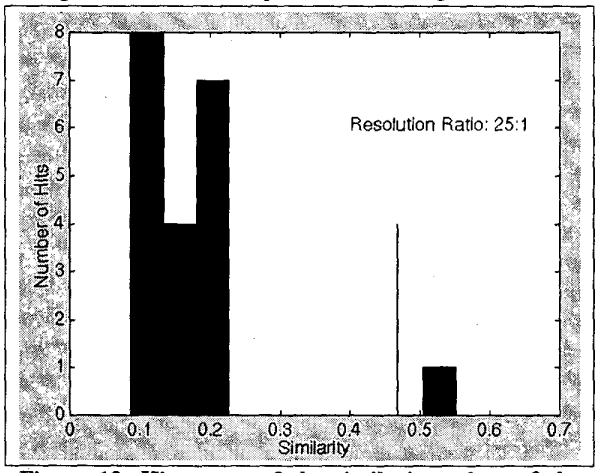

Figure 12: Histogram of the similarity values of the surface patch(Figure 10(c)) to the original twenty patches in the library. The patch shown in (Figure 10(a)) is found to be the best match.

3.4. Robustness to Occlusion

One important issue for surface representations is their robustness to occlusion, i.e., correct matching result should still be obtained even when the surfaces being compared are not complete. In this section, we first explain why Harmonic Shape Images are robust to occlusion and then, using real data, demonstrate their robustness.

The reason for the robustness of Harmonic Shape Images with respect to occlusion lies in the way in which the boundary mapping is constructed. Recall that in Section 2.2.2, the boundary vertices are classified into radius boundary vertices and occluded boundary vertices. It is the radius boundary vertices that determine the angles a_i, which then determine the overall boundary mapping. The effect of occlusion is limited within the occlusion range; therefore, it does not propagate much outside of the occlusion range. This means that, as long as there are enough radius boundary vertices present in the surface patch, the overall harmonic image will remain approximately the same in spite of occlusion.

In order to verify the above observation, the following experiment is conducted. Take the surface patch D_1 shown in Figure 13(a) and cut one part off it to simulate occlusion. The resulted patch D_{1c} is shown in Figure 13(b). Because the vertices on D_{1c} are a subset of the vertices on D_1, we can compare their harmonic images to see how much difference the occlusion causes. Figure 13(c) shows the disparity map. The brightness of a pixel is proportional to the norm of the difference vector between the two harmonic images. It can be seen that the largest differences appear at the occlusion boundary, and the differences decrease gradually as moving away from the occlusion boundary. This shows that harmonic image does not change significantly in the presence of occlusion. Therefore, the Harmonic Shape Image does not change significantly, either. Figure 14 shows the Harmonic Shape Images with and without occlusion.

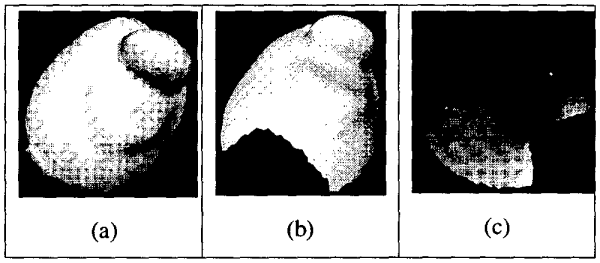

Figure 13: Simulation of occlusion. (a) The original patch; (b) The same patch as (a) with a part cut off; (c) The disparity map between the harmonic images of the patches with and without occlusion.

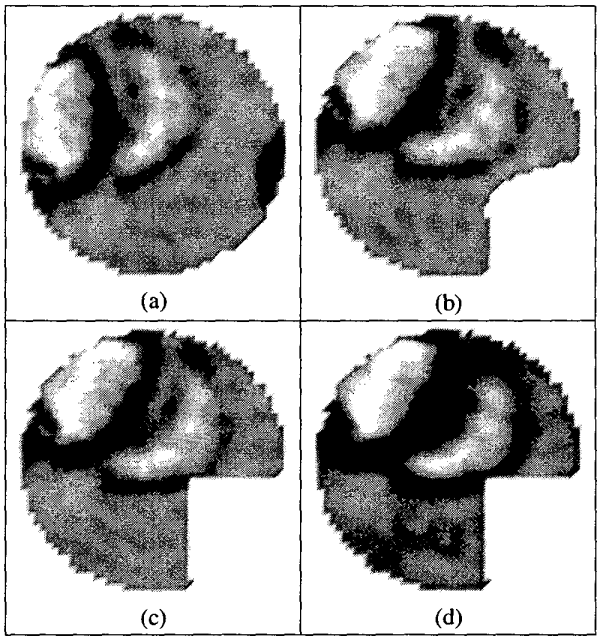

Figure 14: Harmonic Shape Images of surface patches with and without occlusion. (a) The Harmonic Shape Image of the original patch; (b) The Harmonic Shape Image of the patch with occlusion; (c) The Harmonic Shape Image in (b) with pixels close to the occlusion boundary cut off; (d) The Harmonic Shape Image in (c) after Gaussian smoothing using $\sigma = 0.25$.

By comparing the Harmonic Shape Images in Figure 14(a) and (b), it can been that the difference is fairly small. It should be noted that the rotational difference is caused by the selection of different starting vertices for the boundary mapping. The two Harmonic Shape Images will be aligned in the same orientation by the comparison process. This has already been discussed in section 2.2.2. It is not surprising to see that the two Harmonic Shape Images in Figure 14(a) and (b) still have high correlation coefficient 0.86. This value is good enough to find the best match in the library of twenty patches. By cutting off the pixels that are close to the occlusion boundary (Figure 14(c)) and smoothing both Harmonic Shape Images using a Gaussian kernel with $\sigma = 0.25$ (Figure 14(d)), the correlation coefficient can be improved to 0.9.

Another example is shown in Figure 15 and Figure 16 to illustrate the robustness of Harmonic Shape Images with respect to occlusion. The patches in Figure 15(b) has three parts occluded compared to the patch in Figure 15(a). The effect of occlusion can be seen from the disparity map (Figure 15(c)) of their harmonic images. Again the effect is mostly localized in the neighborhood of the occlusion boundaries. Figure 16 shows the Harmonic Shape Images with and without occlusion. The correlation coefficient of the Harmonic Shape Images in Figure 16(a) and (b) is 0.78. It is good enough for the best match to be identified among the patches in the library. This value can be improved to 0.82 by cutting off close-to-occlusion-boundary pixels(Figure 16(c)) and smoothing the Harmonic Shape Images(Figure 16(d)) with a Gaussian kernel of $\sigma = 0.25$.

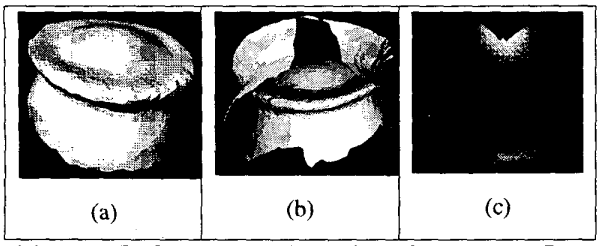

Figure 15: The simulation of occlusion. (a) The original surface patch; (b) The same patch as (a) with three parts occluded; (c) The disparity map between the harmonic images of the patches in (a) and (b).

A more realistic example is shown in Figure 17 to further demonstrate the robustness of Harmonic Shape Images with respect to occlusion. The surface patch in Figure 17(b) has the same shape as the one in Figure 17(a). However, their sampling vertices are different. The resolution ratio between them is 2:1. The patch in Figure 17(b) has an occlusion part. The Harmonic Shape Image of this patch is compared to the twenty patches in the library to find the best match. The correct result is obtained and the histogram of the similarity values is shown in Figure 18.

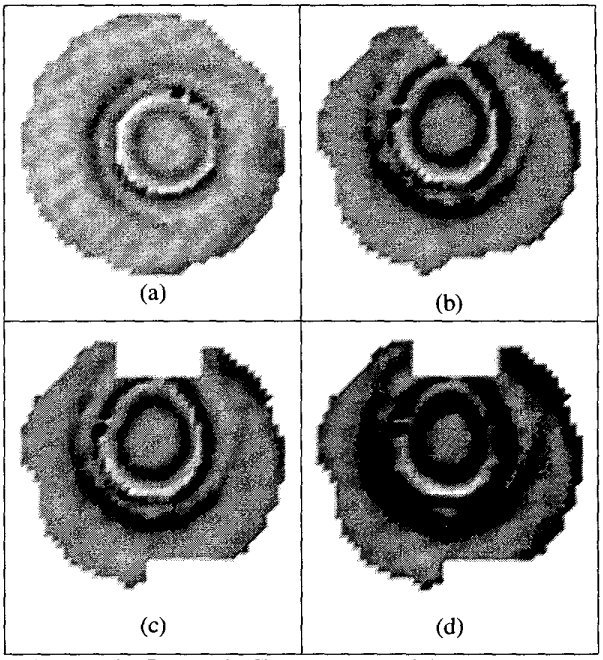

Figure 16: Harmonic Shape Image of (a) the original patch; (b) the original patch with occlusion. (c) Harmonic Shape Image of the occluded patch with pixels that are close to the occlusion boundary cut off; (d) (c) with Gaussian smoothing using $\sigma = 0.25$.

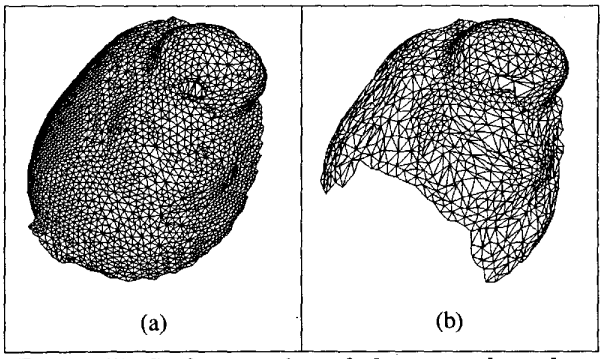

Figure 17: Surface patches of the same shape but different resolutions and occlusion parts.

4. Surface Matching Using Harmonic Shape Images

We have discussed the generation and properties of Harmonic Shape Images in the previous sections. In this section, we will apply Harmonic Shape Images to solve the surface-matching problem. The matching algorithm will be explained first and then some experimental results using real data will be presented.

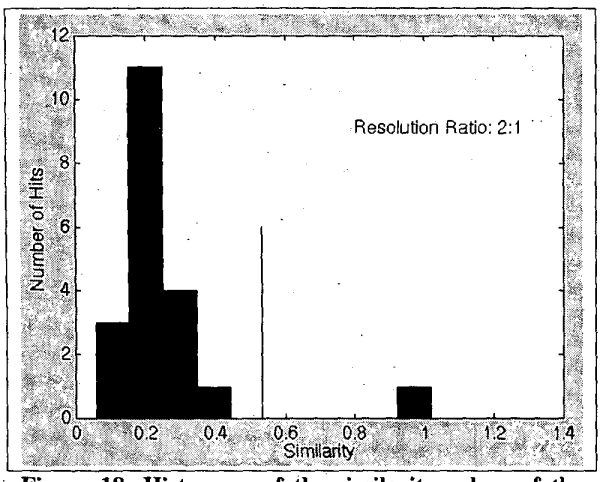

Figure 18: Histogram of the similarity values of the surface patch(Figure 17(b)) to the original twenty patches in the library. The patch shown in (Figure 17(a)) is found to be the best match.

4.1. The Matching Algorithm

The surface matching algorithm using Harmonic Shape Images is as follows: Given two 3D surfaces S_1 and S_2 to be matched, the Harmonic Shape Images $HSI(D_i(v_i, R))$ are generated for each surface patch of S_1. Then a surface patch $D_j(v_j, R)$ is randomly selected on S_2 and its Harmonic Shape Image $HSI(D_j(v_j, R))$ is computed. At the third step, $HSI(D_i(v_j, R)$ is compared to $HSI(D_i(v_i, R))$, $i=1,...,n$, and the similarity value $C(D_j(v_j, R)$ is computed. The best match is identified based on the histogram of all the similarity values using the statistical approach that has been discussed in section 3.1.

Recall the stability property of Harmonic Shape Images in section 3.2, the comparison of Harmonic Shape Images described in the above surface-matching algorithm can be improved using a coarse-to-fine approach. The Harmonic Shape Image, $HSI(D_j(v_j, R))$, of the randomly selected patch, $D_j(v_j, R)$, on S_2 does not need to be compared to every patch on S_1. In fact, only a certain number of randomly selected patches are enough for the coarse comparison. Once a match is found, the fine comparison is conducted by comparing the patch on S_2 to the newly selected patches on S_1. Those newly selected patches are in the neighborhood of the best-matched patch found in the previous comparison. The computation of comparing Harmonic Shape Images can be greatly reduced using this coarse-to-fine approach.

4.2. Surface Matching Examples

The first example is used to illustrate the coarse-to-fine surface-matching algorithm explained in the previous section. Two surfaces, S_1 and S_2, to be registered are shown in Figure 19(a) and (b), respectively. A patch on S_2 is selected as shown in Figure 19(c). The dot marks the central vertex of the patch. For the coarse comparison step, 100 patches are randomly selected on S_1. The central vertices of those patches are marked by dots in Figure 19(d). The circled dot indicates the central vertex of the best-matched patch. At the fine comparison step, the patches that are in the neighborhood of the best-matched patch at the coarse comparison step are candidates to match the patch on S_2. The central vertices of those patches are marked by dots in Figure 19(e). The dot pointed by the arrow indicates the central vertex of the best-matched patch(Figure 19(f)).

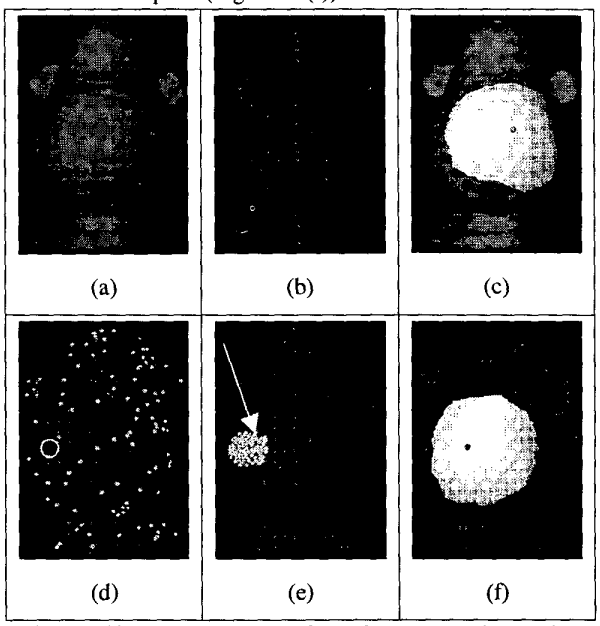

Figure 19: An example of surface matching using Harmonic Shape Images. (a), (b) Surfaces S_1 and S_2 to be registered; (c) A selected patch on S_1 with its central vertex marked; (d) The central vertices of the randomly selected patches on S_2 for coarse comparison; the circled vertex is the center of the best-matched match; (e) The central vertices of the patches for fine comparison; the vertex pointed by the arrow is the center of the best-matched patch in the fine comparison; (f) The best-matched patch with its central vertex marked.

After the best-matched patch is found, the correspondences between the two patches are established (Figure 20(a)). The transformation between the original surfaces is computed using those correspondences. The registered surfaces are shown in Figure 20(b).

It should be noted that for the above example, the best-matched surface patch is found by one iteration of fine comparison. In practice, more iterations may be needed to find the best-matched patch. On the other hand, from the registration point of view, the result of coarse comparison

may be good enough to provide an initial estimate of the transformation for some registration algorithms such as ICP. In fact, the coarse comparison result in the above experiment does provide a good initial estimate for the ICP algorithm. The registration result using this approach is shown in Figure 20(c). Another surface registration example is shown in Figure 21.

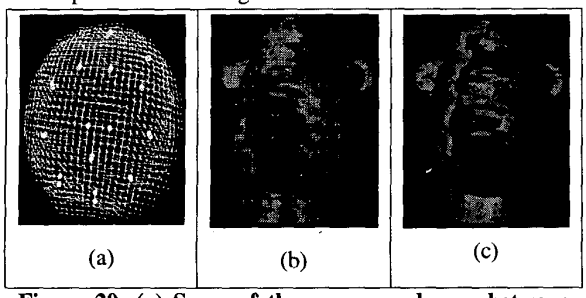

Figure 20: (a) Some of the correspondences between the two matched patches; (b) The registered surfaces using the coarse-to-fine surface matching scheme; (c) The registered surfaces using the coarse comparison and the ICP algorithm.

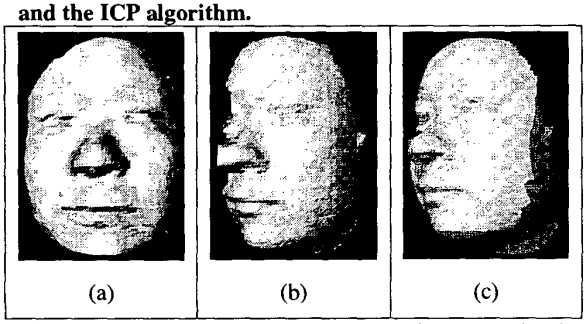

Figure 21: An example of Surface registration. (a) A front view of a person's face; (b) A side view; (c) The registered surfaces.

5. Conclusion and Future Work

In this paper, the experimental analysis of Harmonic Shape Images has been conducted with respect to discriminability, stability and robustness to resolution and occlusion. It has been shown that Harmonic Shape Images are discriminative and stable. Consistently good performance of surface comparison using Harmonic Shape Images can be obtained in the presence of occlusion and different resolution.

There are two directions for the future work. The first one is to investigate the possibility of theoretical analysis of the robustness of Harmonic Shape Images. The second one is to apply the proposed surface matching approach to more applications such as object recognition and classification.

Reference

[1] Y. Xin, "Geometry of Harmonic Maps", Birkhauser, 1996.

[2] J. Eells and L.H, Sampson, "Harmonic mappings of Riemannian manifords. Amer. J. Math., 86:109-160, 1964.

[3] B. O'Neill, "Elementary differential geometry", Academic Press, Inc., 1996.

[4] Matthias Eck, Tony DeRose, Tom Duchamp, Hugues Hoppe, Michael Lounsbery, and Werner Stuetzle, "Multi-resolution Analysis of Arbitrary Meshes", University of Washington, Technical Report, 95-01-02, January 1995.

[5] M. Hebert, K. Ikeuchi and H. Delingette, "A spherical representation for recognition of free-form surfaces", IEEE Transactions on Pattern Analysis and Machine Intelligence, 17(7): 681-689, July 1995.

[6] K. Higuchi, M. Hebert and K. Ikeuchi, "Building 3D models from unregistered range images", CVGIP-Image Understanding, Vol. 57. No. 4. July 1995.

[7] H. Shum, M. Hebert and K. Ikeuchi, "On 3D shape similarity", Proc. CVPR'96, pp. 526-531. June 1996.

[8] D. Zhang, M. Hebert, "Multi-scale classification of 3D objects", Proc. CVPR'97, pp.864-869, July, 1997.

[9] D. Zhang, M. Hebert, A. Johnson and Y. Liu, "On Generating Multi-resolution Representations of Polygonal Meshes", ICCV'98 Workshop on Model-based 3-D Image Analysis, January 3, 1998, Bombay, India.

[10] O.D. Faugeras and M. Hebert, "The representation, recognition and locating of 3-D objects", Int'l J. of Robotics Research, vol. 5, No. 3, pp. 27-52, Fall 1986.

[11] C. Dorai, A. Jain, "COSMOS - a representation scheme for 3D free-form objects", IEEE Transaction Pattern on Pattern Analysis and Machine Intelligence, 19(10): pp. 1115-1130, 1997.

[12] P.J. Besl, "Triangles as a primary representation", Object Representation in Computer Vision, M. Hebert, J. Ponce, T. Boult and A. Gross, eds., pp. 191-206, Berlin, Springer-Verlag, 1995.

[13] T. Joshi, J. Ponce, B. Vijayakumar and D.J. Kriegman, HOT curves for modeling and recognition of smooth curved 3D objects, Proc. IEEE Conf. Computer Vision and Pattern Recognition, Seattle, Wash., pp.876-880, June, 1994.

[14] F. Stein and G. Medioni, "Structural indexing: efficient 3-D object recognition", IEEE Transactions Pattern on Pattern Analysis and Machine Intelligence, 14(2): pp. 125-145, 1992.

[15] D. Keren, K. Cooper and J. Subrahmonia, "Describing complicated objects by implicit polynomials", IEEE Transactions on Pattern Analysis and Machine Intelligence, 16(1): pp. 38-53, 1994.

[16] C.S. Chua and R. Jarvis, "3D free-form surface registration and object recognition", Int'l J. of Computer Vision, vol. 17, pp. 77-99, 1996.

[17] A. Johnson, "Spin-images: a representation for 3-D surface matching", CMU-RI-TR-97-47.

[18] P.J. Besl, "The free-form surface matching problem", Machine Vision for Three-dimensional Scenes, H. Freeman, ed., pp. 25-71, Academic Press, 1990.

[19] P.J. Besl and N.D. Mckay, "A method for registration of 3-D shapes", IEEE Transactions on Pattern Analysis and Machine Intelligence, 14(2): pp. 239-256, 1992.

[20] Y. Chen and G. Medioni, "Object modeling by registration of multiple range images", Image Vision Computing, 10(3): 145-155, 1992.

[21] R. Bergevin, D. Laurendeau and D. Poussart, "Estimating the 3D rigid transfromation between two range views of a complex object", 11th IAPR, Int'l Conf. Patt. Recog., pp. 478-482, The Hague, The Netherlands, Aug. 30 - Sep. 3, 1992.

[22] J. Deovre, "Probability and statistics for engineering and science", Brooks/Cole, Belmont, CA, 1987.

[23] D. Zhang and M. Hebert, "Harmonic maps and their applications in surface matching", Proc. CVPR'99.