Harmonic Maps and Their Applications in Surface Matching

Dongmei Zhang and Martial Hebert
The Robotics Institute, Carnegie Mellon University
{dzhang, hebert}@ri.cmu.edu

Abstract*

The surface-matching problem is investigated in this paper using a mathematical tool called harmonic maps. The theory of harmonic maps studies the mapping between different metric manifolds from the energy-minimization point of view. With the application of harmonic maps, a surface representation called harmonic shape images is generated to represent and match 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.

Due to the application of harmonic maps in generating harmonic shape images, harmonic shape images have the following advantages: they have sound mathematical background; they preserve both the shape and continuity of the underlying surfaces; and they are robust to occlusion and independent of any specific surface sampling scheme.

The performance of surface matching using harmonic maps is evaluated using real data. Preliminary results are presented in the paper.

1. Introduction

Surface matching is a fundamental issue in computer vision. Generally speaking, given two free-form surfaces in 3D space, the goals of surface matching are to find the rigid transformation between the two surfaces, to establish the correspondences, and to determine whether the two surfaces are similar in terms of shape. Surface matching is of both theoretical interest and practical importance because it is closely related to the basic problem of surface representation and it has wide applications in surface registration, object recognition and object classification.

A large amount of research has been done on surface matching. The approaches to solving the problem can be classified into two categories: model-based matching[5]-[17] and matching by registration[18]-[21].

Among the surface representations proposed so far, the one that uses local shape signature[17] is one of those that have been quite successful in real applications. This representation is independent of surface topology and easy to compute. The matching performance using this representation decreases gracefully as the occlusion and clutter increases. However, the major limitation of this representation is that it only provides sets of individual point correspondences that have to be grouped into sets of mutually consistent correspondences. This limitation comes from the fact that the local signatures capture the shape of the surface only partially.

In this paper, in our continuing effort to develop the data-level representations for surface matching, we investigate the surface-matching problem using a mathematical tool called harmonic maps with the goal of addressing the limitation of the representation in[17]. The harmonic map theory studies mappings between different manifolds from an energy-minimization point of view. With the application of harmonic maps, a surface representation called harmonic shape images is created and used for surface matching. Furthermore, owing to the properties of harmonic maps, harmonic shape images are able to provide all the point correspondences once two regions are matched. This will be shown in detail in the paper.

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, harmonic maps provide a mathematical solution to the mapping problem between a 3D surface patch with disc topology and a 2D domain.

This paper is organized as follows: The mathematical background of harmonic maps is introduced in Section 2; in Section 3, an approximation of harmonic maps is discussed in detail; applying harmonic maps in surface matching is discussed in Section 4 followed by the experimental results and analysis in Section 5. Conclusions and future work will be presented at the end of the paper.

2. Harmonic Maps

According to [2], the concept of harmonic maps is closely related to the concept of geodesics. Geodesics are the shortest connection between two points in a metric

* This research was supported in part by the National Science Foundation under grant IRI9711853.

continuum, e.g., a Riemannian manifold. Geodesics are critical points of the following length integral

$$L(c) = \int_0^1 \left| \frac{\partial}{\partial t} c \right|^2 dt \qquad (1)$$

where $c : [0,1] \to N$ is the parametrization of the curve propotional to arc length. The generalization of energy integral in (1) for maps between Riemannian manifolds leads to the concept of harmonic maps. Harmonic maps are critical points of the corresponding integral where energy density is defined in terms intrinsic to the geometry of the source and target manifolds and the map between them.

Formally, harmonic maps are defined as follows[1][2]. Let *(M, g)* and *(N, h)* be two smooth manifolds of dimensions *m* and *n*, respectively and let $\phi : (M, g) \to (N, h)$ be a smooth map. Let (x^i), $i = 1, \ldots, m$ and (y^α), $\alpha = 1, \ldots, n$ be local coordinates around *x* and $\phi(x)$, respectively. Take (x^i) and (y^α) of *M* and *N* at corresponding points under the map ϕ whose tangent vectors of the coordinate curves are $\partial / \partial x^i$ and $\partial / \partial y^\alpha$ respectively. Then the energy density of ϕ is defined as

$$e(\phi) = \frac{1}{2} \sum_{i,j=1}^{m} g^{ij} \sum_{\alpha,\beta=1}^{n} \frac{\partial f^\alpha}{\partial x^i} \frac{\partial f^\beta}{\partial x^j} h_{\alpha\beta} \qquad (2)$$

in which g^{ij} and $h^{\alpha\beta}$ ($h_{\alpha\beta}$ is the inverse of $h^{\alpha\beta}$) are the components of the metric tensors in the local coordinates on *M* and *N*. The energy of ϕ in local coordinates is given by the number

$$E(\phi) = \int_M e(\phi) dM \qquad (3)$$

If ϕ is of class C^2, $E(\phi) < \infty$, and ϕ is an extremum of the energy, then ϕ is called a harmonic map and statisfies the corresponding Euler-Lagrange equation.

The above is a general definition of harmonic maps. Now let us look at a special case in which both the source manifold and the target manifold are surfaces in the 3D Euclidean space. To be more specific, let *D* be a surface of disc topology and *P* be a planar region. According to the results in the theory of harmonic maps[2], the following problem has a unique solution: given a homeomorphism *b* between the boundary of *D* and the boundary of *P*, there exists a unique harmonic map $\phi : D \to P$ that agrees with *b* on the boundary of *D* and minimizes the energy functional of *D*.

Furthermore, the harmonic map ϕ has the following properties[2]: it always exists; it is unique and continuous; it is one-to-one and onto and it is intrinsic to *D* and *P*. All the above properties show that the harmonic map ϕ is a well-behaved mapping.

3. Approximation of Harmonic Maps

As we have already seen in the previous section, harmonic maps are solutions of 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. In this section, we will first discuss the interior mapping with a given boundary mapping and then discuss our boundary mapping in detail.

3.1. 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, $i=1,\ldots,n^i$, be the interior vertices of *D*. The interior mapping ϕ maps v_i^i, $i=1,\ldots,n^i$, onto the interior of the unit disc *P* with a given boundary mapping $b : \partial D \to \partial P$. $\phi(i)$ is obtained by minimizing the following energy functional.

$$E(\phi) = \tfrac{1}{2} \sum_{\{i,j\} \in Edges(D)} k_{ij} \left\| \phi(i) - \phi(j) \right\|^2 \qquad (4)$$

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 with the definition in (5)

$$k_{ij} = ctg\theta(e_{mi}, e_{mj}) + ctg\theta(e_{li}, e_{lj}) \qquad (5)$$

in which $\theta(e_{mi}, e_{mj})$ and $\theta(e_{li}, e_{lj})$ are defined in Figure 1. If e_{ij} is only associated with one triangle, then there will be only one term on the right of (5).

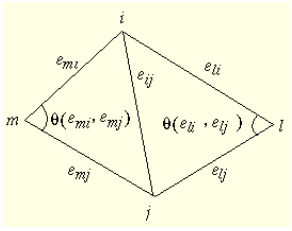

Figure 1: Definition of spring constants.

An instance of the functional $E(\phi)$ can be interpreted as the energy of a spring system by associating each edge in

D with a spring. Then the mapping problem from D to P can be viewed as adjusting the lengths of those springs when squeezing them down onto P. If the energy of D is considered to be zero, then the energy is increased when squeezing the springs down to P because all the springs are compressed. Different ways of adjusting the spring lengths correspond to different mappings ϕ. The best ϕ minimizes the energy functional $E(\phi)$. By defining the spring constants in (5), the best ϕ best preserves the ratios of edge lengths in D, therefore the shape of D, under the boundary mapping b. The minimum of the energy functional $E(\phi)$ can be found by solving a sparse linear least-square system of (4) for the values $\phi(i)$[4]. This ϕ is an approximation of the harmonic maps discussed in the previous section.

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 2(a) and (b) respectively.

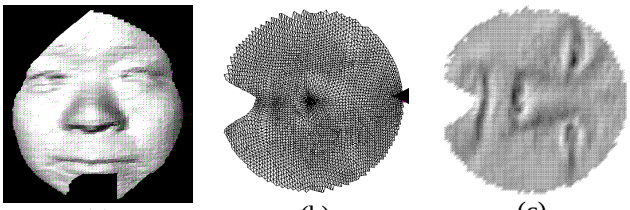

(a) (b) (c)

Figure 2: (a) An example of surface patch; (b) its harmonic image; (c) its harmonic shape image. (c) is the same as (b) except that it has shape attribute color-coded at each vertex. It will be introduced later in the paper.

3.2. Boundary Mapping

The constuction of the boundary mapping $b : \partial D \to \partial P$ is illustrated in Figure 3.

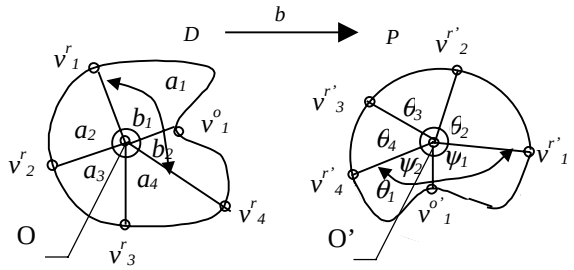

Figure 3: Illustration of the boundary mapping between the surface patch and the 2D domain.

First of all, let us define the vertices and vectors in Figure 3. O is the central vertex of D and O' is the center of P. v_i, $i=1,\ldots,5$ 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, $i=1,\ldots,4$, 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 3. Angles a_i, $i=1,\ldots,4$ 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. For example, the sum of b_j over (v_4^r, v_1^o, v_1^r) is a_1.

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 3, v_i^r, $i=1,\ldots,4$, are mapped to $v_I^{r'}$, $i=1,\ldots,4$, 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:

$$\theta_i = \frac{a^i}{\sum_{k=1}^{n} a_k} 2\pi \tag{6}$$

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 3, v_1^o, which is in the occlusion rangle (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.

$$\psi_j = \frac{bj}{\sum_{m=1}^{n} b_m} a_i \tag{7}$$

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

$$r_j = \frac{dist(r_j, O)}{R} \tag{8}$$

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.

Two issues need to be mentioned regarding the above boundary mapping. The first is that the boundary vertices of *D* need to be ordered in either a clock-wise or counter-clock-wise manner before constructing the boundary mapping. Similarly, when mapping vertices onto the boundary of *P*, either clock-wise or counter-clock-wise order needs to be determined. The two orders must remain consistent for all surface patches for the convenience of matching.

The second issue is how to select the starting vertex among the boundary vertices of *D*. If the starting vertex is always mapped to the same vertex on the boundary of the unit disc, then different starting vertices will result in different boundary mappings. This will, in turn, result in different interior mappings. For example, the harmonic image shown in Figure 2(b) has the starting vertex indicated by the black mark, while in Figure 4(a), the harmonic image of the same surface patch is different due to a different starting vertex.

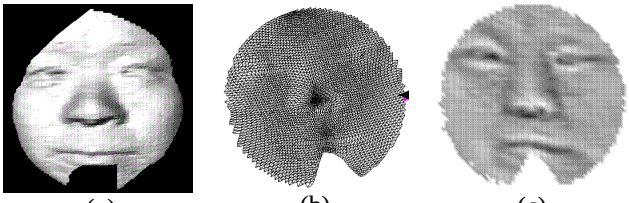

(a) (b) (c)

Figure 4: The harmonic image(b) and harmonic shape image(c) of the same surface patch(a) shown in Figure 2(a) with a different a starting vertex for boundary mapping.

In fact, the harmonic images with different starting vertices are different by a planar rotation. The reason for this is that neither the angles θ_i(in (6)) nor the angles ψ_i (in (7)) will change with respect to different starting vertices. Nor will the radius r_j in (8) change. Therefore, the starting vertex can be selected randomly. The rotation difference will be found later by the matching process which will be discussed in the next section.

4. Surface Matching Using Harmonic Maps

The key role that harmonic maps play in surface matching is to help create a surface representation called harmonic shape images. Matching 3D surfaces can then be simplified to matching harmonic shape images. In this section, we first discuss the generation of harmonic shape images and then discuss how to match them.

4.1. Harmonic Shape Images

In Section 2 and Section 3, 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, *hs(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 3(c) and Figure 4(c) are examples of harmonic shape images in which curvature values (from high to low) are proportional to the intensity values.

Because harmonic shape images are generated based on harmonic maps, they have the same properties as that of harmonic maps, e.g., existence and uniqueness. The generation process shows that harmonic shape images preserve both the shape and the continuity of the corresponding surface patches. It also shows that harmonic shape images are independent of the sampling of the surface patches. As we discussed in Section 1, these properties are the goals we set for harmonic shape images.

Similar to harmonic shape images, more images can be generated by associating other properties, e.g., color, to each vertex in harmonic images. This shows that harmonic maps provide a general framework for storing surface-related information.

4.2. Matching Harmonic Shape Images

In Section 3, we have already discussed the planar rotation difference between harmonic images under different starting boundary vertices. As a result, the corresponding harmonic shape images have rotation difference as well (Figure 3(c) and Figure 4(c)). This rotation difference can be found by template matching plus a planar rotation.

The similarity criterion proposed in [17] is used to evaluate the matching result of two harmonic shape images. It is defined as follows.

$$C(D_1, D_2) = (\ln \frac{1+R}{1-R})^2 - \lambda \frac{1}{N-3} \qquad (9)$$

in which R is the normalized correlation coefficient between $hs(D_1)$ and $hs(D_2)$. N is the number of corresponding vertices. λ is a weight controlling the contribution of N to $c(D_1, D_2)$ and is set to one half. More details about this similarity coefficient can be found in [17].

4.3. Matching Surfaces by Matching Harmonic Shape Images

Given two 3D surfaces S_1 and S_2 to be matched, harmonic shape images $hs(D_i(v_i, R))$ are generated for each surface patch of S_1. Then randomly select a surface patch $D_j(v_j, R)$ on S_2 and compute its harmonic shape image $hs(D_i(v_j, R))$. At the third step, match $hs(D_i(v_j, R)$ to $hs(D_i(v_i, R))$, $i=1,\ldots,n$, and compute the similarity criterion $C(D_j(v_j, R), D_i(v_i, R))$. The best match is identified based on the

histogram of all the similarity values using a statistical approach[24].

5. Performance and Analysis

In this section, harmonic shape images are used in the application of surface registration. We first present some results on real data and then conduct some analysis on harmonic shape images.

5.1. Surface Registration Using Harmonic Shape Images

Two surfaces S_1 and S_2 shown in Figure 5(a), (b) are to be registered. The surface patch in wireframe overlayed on S_1 is a selected surface patch $D_i(v_i, R)$ on S_1. Following the procedure of surface matching discussed in Section 4, the surface patch that best matches $D_i(v_i, R)$ is found. Figure 5(c) shows the correspondences between the two matched surface patches. The transformation between S_1 and S_2 can be computed using those correspondences. The registered surfaces are shown in Figure 5(d).

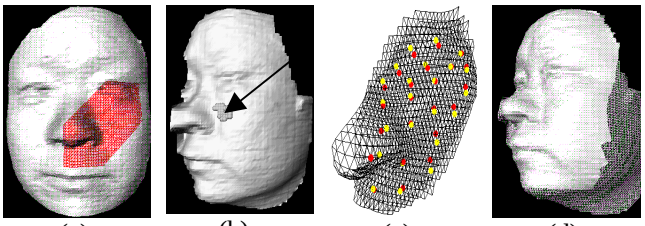

(a) (b) (c) (d)

Figure 5: Surface registration using harmonic shape images. (a), (b) surfaces to be registered; (c) Some of the correspondences between the best matched surface patches; (d) The registered surfaces. For the surface patches on S_2 that are good matches of the surface patch in (a), their central vertices are marked. The center of the best matched patch is indicated by the arrow.

Figure 6 is the histogram of the similarity values of matching $D_i(v_i, R)$ on S_1 to all surface patches on S_2. The following statistical approach is used to determine the best match.

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.

In Figure 6, good matches correspond to the extreme outliers above the upper fourth. They are on the high end of the horizontal axis. The good matches are also displayed in Figure 5 with their central vertices marked in green. The best match is marked in red. The fact that all the good matches are in the neighborhood of the best match shows that the harmonic shape images are an effective and discriminating representation.

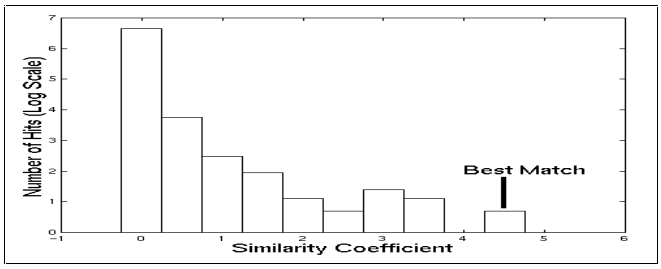

Figure 6: Histogram of the similarity values. Good matches are on the high end of the horizontal axis.

5.2. Robustness of Harmonic Shape Images

One important issue for surface representations is their robustness because robustness ensures the correct surface matching in the presence of noise, different resolution and occlusion. In this section, we first explain why harmonic shape images are robust to occlusion and then, using real data, demonstrate the robustness under occlusion, different resolution and noise.

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 3.2 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.

The purpose of the first experiment is to demonstrate the robustness of harmonic shape images under occlusion. Let us use the surface patches D_1 and D_2 in Figure 7(1) and (2) as an example to illustrate the experiment. D_2 is the same as D_1 except for the occluded region. Their harmonic shape images are shown in Figure 7(5) and (6), respectively. It can be seen that the harmonic shape images of D_1 and D_2 are similar in spite of the occlusion on D_2. This means that the occlusion part of D_2 does not affect much of the shape representation for the non-occluded regions, thus making it possible to match D_2 to D_1 by matching the harmonic shape images of their non-occluded regions. The normalized correlation coefficient of $hs(D_1)$ and $hs(D_2)$ is 0.9878, which verifies that D_2 can still be matched correctly to D_1 with occlusion.

A sequence of surface patches with different occlusion parts on D_1 is matched to D_1 to further test the robustness of harmonic shape images under occlusion. Figure 7(3) and (4) are two examples in the sequence. Their harmonic

shape images are shown in Figure 7(7) and (8), respectively. Figure 8 shows the matching result.

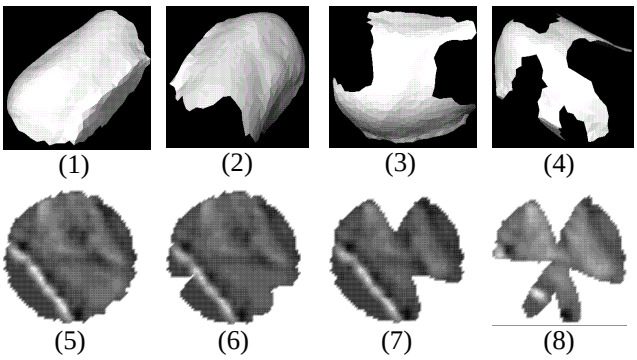

Figure 7: A surface patch and its variations with different parts cut off to simulate occlusion.

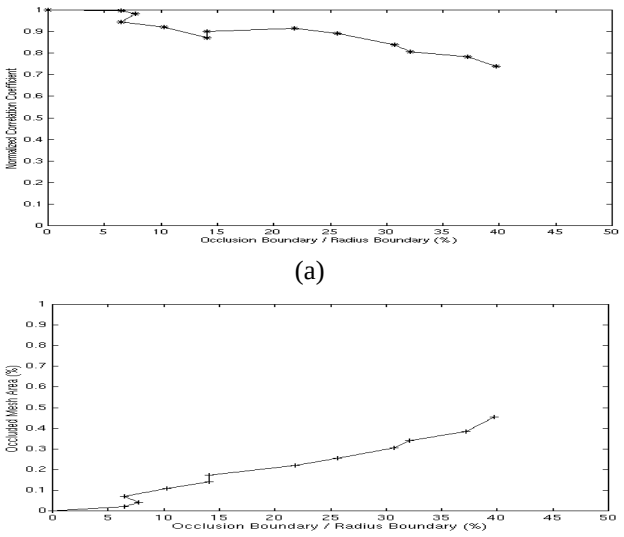

Figure 8: Matching results under occlusion. The original surface is shown in Figure 7(1). Some of the surfaces with occlusion are shown in Figure 7(2)-(4). (a) Normalized correlation coefficient; (b) Percentage of occluded area.

It can be seen from Figure 8 that as the percentage of occlusion boundary increases, the occluded mesh area also increases. However, the normalized correlation coefficient remains stable in the range of [0.8, 1.0]. Figure 7(8) shows the harmonic shape image of the surface patch in Figure 7(4). In spite of severe occlusion, the harmonic shape image for the non-occluded regions is still similar to that in Figure 7(5).

The second experiment is to test the robustness of harmonic shape images to the presence of occlusion, different resolution and noise. In Figure 9, surface patches (2)-(8) and the one in (1) have the same shape but different resolution and occlusion parts. The surface patches in (2)-(8) are matched to the one in (1) by matching their harmonic shape images. The result in Figure 10 shows a stable matching performance.

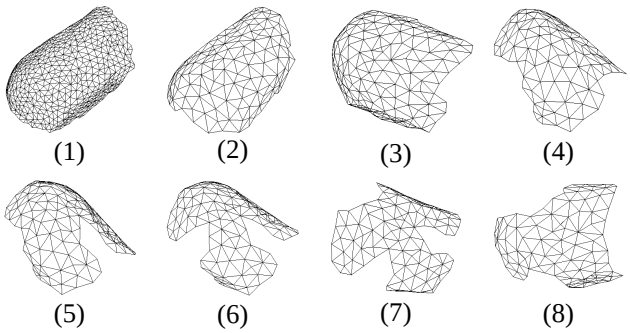

Figure 9: Surface patches with the same shape but different resolution and occluded parts. (1) A surface patch of high resolution; (2) A surface patch of the same shape as (1) but with low resolution; (3)-(8) occluded surface patches with different parts cut off on (2).

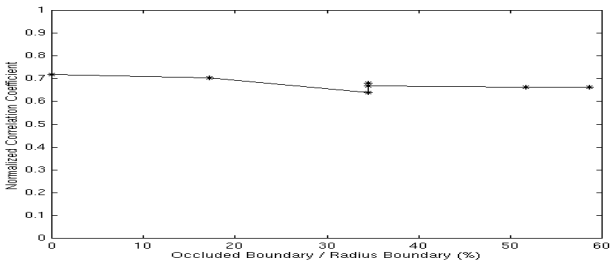

Figure 10: Matching results for the surface patches in Figure 9. Compared to Figure 8(a), the drop of the normalized correlation coefficient is due to the use of a coarse discrete curvature approximation. The effect of this curvature approximation becomes significant in this experiment because the two surface patches are of different resolution and the resolution for the surface patches (2)-(8) is low.

The above experiments are further conducted on a more complex surface patch shown in Figure 11(1). Some of the occluded surface patches are shown in Figure 11(2)-(4). Their harmonic shape images are shown in Figure 11(6)-(8). The curve in Figure 12 shows that the overall matching performance of harmonic shape images is stable under occlusion. The fluctuation of the curve shows that the robustness of harmonic shape images with respect to occlusion does not only depend on the ratio of occluded region to the entire surface patch, but also on the shape of the occluded region.

6. Conclusion and Future Work

In this paper, the problem of surface matching is investigated using a mathematical tool called harmonic maps. With the application of harmonic maps, harmonic shape images are generated to represent and match 3D free-form surfaces. Harmonic shape images have sound mathematical background. They preserve both the shape and continuity of the underlying surfaces. Preliminary results have shown that harmonic shape images are robust to occlusion and independent of surface sampling.

There are two directions for the future work. The first one is to further analyze the properties of harmonic shape images both theoretically and experimentally. The second one is to apply the proposed surface matching approach to more applications such as object recognition and classification.

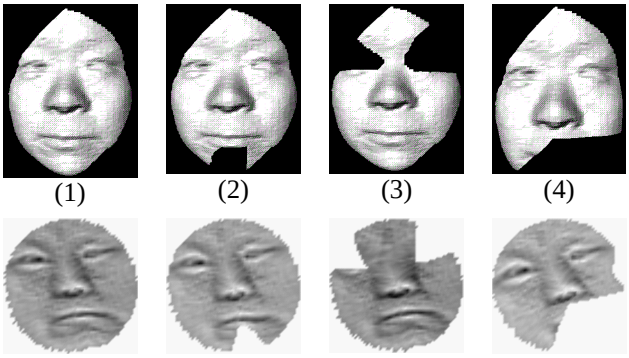

Figure 11: A complex surface patch(1) with different parts cut off to simulate occlusion(2), (3), (4). Their harmonic shape images are shown in (5)-(8).

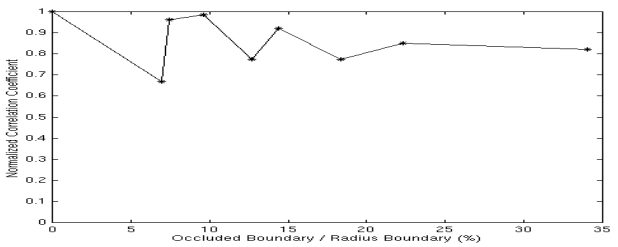

Figure 12: Matching results under occlusion. The original surface is shown in Figure 11(1). Some of the surfaces with occlusion are shown in Figure 11(2)-(4).

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. Schwartz and M. Sharir, "Identification of partially obscured objects in two and three dimensions by matching noisy characteristic curves", The Int. J. Robotics Research, 6(2): pp. 29-44, 1987.

[23] S. Sclaroff and A. Pentland, "Object recognition and categorization using modal matching", Proc. 2nd CAD-Based Vision Workshop, pp. 258-265. Champion, Pennsylvania, Feb. 8-11, 1994.

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