

Conformal Geometry and Its Applications on 3D Shape Matching, Recognition and Stitching

Sen Wang, Yang Wang, Miao Jin, Xianfeng Gu, Dimitris Samaras

The authors are with State University of New York at Stony Brook, Stony Brook, NY 11794, USA.

Manuscript received Dec 21, 2005; revised May 20, 2006.

Abstract

3D shape matching is a fundamental issue in computer vision with many applications such as shape registration, 3D object recognition and classification. However, shape matching with noise, occlusion and clutter is a challenging problem. In this paper, we analyze a family of quasi-conformal maps including harmonic maps, conformal maps and least squares conformal maps with regards to 3D shape matching. As a result, we propose a novel and computationally efficient shape matching framework by using least squares conformal maps. According to conformal geometry theory, each 3D surface with disk topology can be mapped to a 2D domain through a global optimization and the resulting map is a diffeomorphism, i.e., *one-to-one* and *onto*. This allows us to simplify the 3D shape-matching problem to a 2D image-matching problem, by comparing the resulting 2D parametric maps, which are stable, insensitive to resolution changes and robust to occlusion and noise. Therefore, highly accurate and efficient 3D shape matching algorithms can be achieved by using the above three parametric maps. Finally, the robustness of least square conformal maps is evaluated and analyzed comprehensively in 3D shape matching with occlusion, noise and resolution variation. In order to further demonstrate the performance of our proposed method, we also conduct a series of experiments on two computer vision applications, i.e., 3D face recognition and 3D non-rigid surface alignment and stitching.

Index Terms

Shape Representations, Shape Matching, Conformal Geometry, 3D Face Recognition.

I. INTRODUCTION AND PREVIOUS WORK

3D shape matching is a fundamental issue in 3D computer vision with many applications, such as shape registration, partial scan alignment, 3D object recognition and classification [8], [50], [37], [23]. As 3D scanning technologies improve, large databases of 3D scans require automated methods for matching. However, matching 3D shapes in noisy and cluttered scenes is a challenging task. Moreover, since most 3D shape scanners can only capture 2.5 D data of the target surfaces, aligning and stitching partial 3D surfaces is a fundamental problem in many research areas, such as computer vision, mechanical engineering, and molecular biology.

Generally, the crux of 3D shape matching is finding good shape representations, allowing us to match two given free-form surfaces by comparing their shape representations. Different approaches include curvature-based representations [45], regional point representations [26], [37], [43], [10], spherical harmonic representations [27], [17], [18], shape distributions [34], spline

representations [7] and harmonic shape images [51]. However, many shape representations that use local shape signatures are not stable and cannot perform well in the presence of noise. In this paper, we propose to use a family of quasi-conformal maps, including harmonic maps, conformal maps and least squares conformal maps, that does not suffer from such problems. According to conformal geometry theory, each 3D shape with disk topology can be mapped to a 2D domain through a global optimization and the resulting map is a diffeomorphism, i.e., *one-to-one* and *onto*. Consequently the 3D shape-matching problem can be simplified to a 2D image-matching problem of the quasi-conformal maps. These maps are stable, insensitive to resolution changes and robust to occlusion and noise. The 2D maps integrate geometric and appearance information and 2D matching is a better understood problem [31], [4]. Therefore, highly accurate and efficient 3D shape matching algorithms can be achieved using quasi-conformal maps.

The robustness and easy use of the technique we proposed allow us to cope with more challenging problems such as surface alignment and stitching, when only two parts of surfaces could be matched. There has been a lot of research on 3D surface alignment and stitching in recent decades, such as identification and indexing of surface features[15], [42], computing principal axes of scans[12], exhaustive search for corresponding points[9], or iterative closest point(ICP) methods[28], [36], [38], [6]. Compared to matching, there are other additional issues in surface stitching, such as registration and integration[44]. 3D surface alignment and stitching is still a challenging task especially when the transformation between the surfaces to be aligned is non-rigid, e.g., when taking successive scans of humans that might not be standing still. Based on conformal geometry theory, an important property of Least Squares Conformal Maps(LSCMs) is that they can map a 3D surface to a 2D domain in a continuous manner with minimized local angle distortion. This implies that *LSCMs are not sensitive to surface deformations*, which leads to a natural solution to 3D non-rigid surface alignment and stitching.

Quasi-Conformal maps including harmonic maps, conformal maps and least squares conformal maps have been used in several applications of computer vision and graphics. In [51], Zhang et al. proposed harmonic maps for surface matching. In [49], Wang et al. use harmonic maps to track dynamic 3D surfaces. In [20], [19], [48], [47], conformal maps are used for face and brain surface matching. Moreover, Sharon et al. [40] use conformal maps to analyze similarities of 2D shapes. Least squares conformal maps are introduced by Levy et al. [29] for texture atlas generation and used by Wang et al. [46] to do 3D surface matching. In order to calculate

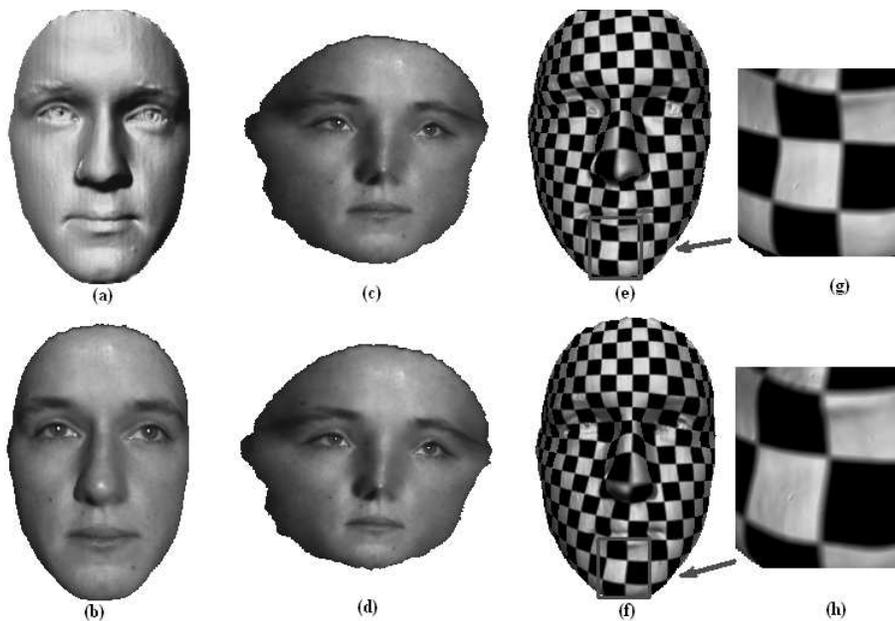


Fig. 1. Distortion comparison between a conformal map and a harmonic map. (a) Original surface with texture. (b) Original surface without texture. (c) The 2D conformal map of the surface with texture. (d) The harmonic map of the surface with texture. (e) Checkerboard textured surface by conformal mapping. (f) Checkerboard textured surface by harmonic mapping. Because of angle-preservation, (c) and (e) have less distortions than (d) and (f), which can be clearly seen in the close-up views (g) and (h) of the chin areas in the gray boxes respectively.

harmonic maps, the surface boundary needs to be identified and a boundary mapping from 3D surfaces to the 2D domain needs to be created, which can be a difficult problem especially when part of the surface is occluded. However, the two other quasi-conformal maps we discuss in this paper, conformal maps and least squares conformal maps, do not need boundary information and so lend themselves as a natural choice to solve this problem. Moreover, in addition to the advantages of harmonic maps, such as sound mathematical basis and preservation of continuity of the underlying surfaces, conformal maps are also angle preserving, which leads to less distortion and robustness to noise. The differences between conformal maps and harmonics maps are shown in Fig. 1.

In this paper, we make the following contributions:

- 1) We analyze a family of quasi-conformal maps, including harmonic maps, conformal maps and least squares conformal maps, when applied to 3D shape matching and compare their properties comprehensively.

- 2) We propose a novel 3D shape matching framework, using least squares conformal maps.
- 3) We systematically evaluate the robustness of least squares conformal maps on 3D shape matching for different challenges such as occlusion, noise and resolution variation. For completeness purposes, we also provide a full comparison between different quasi-conformal maps for 3D shape matching.
- 4) We demonstrate the performance of least squares conformal maps in practice through 3D face recognition and 3D non-rigid surface alignment and stitching.

The rest of the paper is organized as follows: The mathematical background of the harmonic and conformal maps is introduced and compared in Section II. A framework for 3D shape matching using least squares conformal maps is proposed in Section III. Experimental results and performance analysis are presented in Section IV, and we conclude with discussion and future work in Section V.

II. THEORETICAL BACKGROUND OF CONFORMAL GEOMETRY

An important merit of quasi-conformal maps, including harmonic maps, conformal maps and least squares conformal maps, is to reduce the 3D shape-matching problem to a 2D image-matching problem, which has been extensively studied. Quasi-conformal mappings, which are almost conformal, do not distort angles arbitrarily and this distortion is uniformly bounded throughout their domain of definition[3]. We are dealing with 3D surfaces, but since they are manifolds, they have an inherent 2D structure, which can be exploited to make the problem more tractable using conformal geometry theory [20], [40]. Most work using conformal geometry theory is done in surface parameterization, which can be viewed as an embedding from a 3D surface \mathbf{S} with disk topology to a planar domain \mathbf{D} . Following the introduction of the notions of harmonic maps, conformal maps and least squares conformal maps, these three parametric maps will be compared in a comprehensive manner.

A. Harmonic maps

As described in [51], a harmonic map $H : \mathbf{S} \rightarrow \mathbf{D}$ is a critical point for the harmonic energy functional,

$$E(H) = \int_{\mathbf{S}} |\nabla H|^2 d\mu_{\mathbf{S}}, \quad (1)$$

and can be calculated by minimizing $E(H)$. The norm of the differential $|\nabla H|$ is given by the metric on \mathbf{S} and \mathbf{D} , and μ_S is the area element on 3D surface \mathbf{S} [39], [33], [13], [14]. Since the source surface mesh \mathbf{S} is in the form of a *discrete* triangular mesh, we approximate the harmonic energy as [13], [51], [20],

$$E(H) = \sum_{[v_0, v_1]} k_{[v_0, v_1]} |H(v_0) - H(v_1)|^2, \quad (2)$$

where $[v_0, v_1]$ is an edge connecting two neighboring vertices v_0 and v_1 , and $k_{[v_0, v_1]}$ is defined as

$$\frac{1}{2} \left(\frac{(v_0 - v_2) \cdot (v_1 - v_2)}{|(v_0 - v_2) \times (v_1 - v_2)|} + \frac{(v_0 - v_3) \cdot (v_1 - v_3)}{|(v_0 - v_3) \times (v_1 - v_3)|} \right) \quad (3)$$

where $\{v_0, v_1, v_2\}$ and $\{v_0, v_1, v_3\}$ are two adjacent triangular faces.

By minimizing the harmonic energy, a harmonic map can be computed using the Euler-Lagrange differential equation for the energy functional, i.e.,

$$\Delta E = 0, \quad (4)$$

where Δ is the Laplace-Beltrami operator [39], [33], [13], [14]. This will lead to solving a sparse linear least squares system for the mapping H of each vertex v_i [13], [51], [49], [20]. If the boundary condition

$$H|_{\partial \mathbf{S}} : \partial \mathbf{S} \rightarrow \partial \mathbf{D} \quad (5)$$

is given, the solution exists and is unique.

Although harmonic maps are easy to compute, they require satisfaction of the above boundary condition, which becomes unreliable when there are occlusions in the 3D original data. To overcome this problem, the missing boundaries can be approximated [51], which might be enough for rough surface matching. However, since interior feature points are often more robust to occlusion, it is desirable to replace the boundary condition with feature constraints. This can be achieved by conformal maps, another mathematical tool in conformal geometry theory, which only require several feature constraints as an input and obviate the need to specify the boundary condition.

B. Conformal maps

It can be proven that there exists a mapping from any surface with a disk topology to a 2D planar domain [21], which is one-to-one, onto, and angle preserving. This mapping is called *conformal mapping* and keeps the line element unchanged, except for a local scaling factor [16].

Conformal maps have many appealing properties, one of which is their connection to complex function theory [16], [29]. Consider the case of mapping a planar region \mathbf{S} to the plane. Such a mapping can be viewed as a function of a complex variable, $d = \mathcal{U}(s)$. Locally, a conformal map is simply any function \mathcal{U} which is analytic in the neighborhood of a point s and such that $\mathcal{U}'(s) \neq 0$. A conformal mapping \mathcal{U} thus satisfies the Cauchy-Riemann equations, which are

$$\frac{\partial u}{\partial x} = \frac{\partial v}{\partial y}, \quad \frac{\partial u}{\partial y} = -\frac{\partial v}{\partial x}. \quad (6)$$

where $d = u + iv$ and $s = x + iy$.

Differentiating one of these equations with respect to x and the other with respect to y , we obtain the two Laplace equations

$$\Delta u = 0, \quad \Delta v = 0. \quad (7)$$

where $\Delta = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}$. Any mapping which satisfies these two Laplace equations is called a harmonic mapping. Thus a conformal mapping is also harmonic. However, unlike the harmonic maps described in the previous section, which need the boundary mapping $H|_{\partial S}$ fixed in advance, conformal maps can be calculated without demanding the mesh boundary to be mapped onto a fixed shape. For a discrete mesh, the main approaches to achieve conformal parameterizations are: harmonic energy minimization[11], [20], [19], [48], [47], Cauchy-Riemann equation approximation[29], Laplacian operator linearization[21], circle packing[24], most isometric parameterizations(MIPS)[22] and angle-based flattening method[41]. In this paper, we compute conformal maps using the harmonic energy minimization method[20].

Riemann's theorem states that for any surface \mathbf{S} homeomorphic to a disc, it is possible to find a parameterization of the surface satisfying Eq. 6 [29], which can be uniquely determined by two points on surface \mathbf{S} . However, to better handle the errors caused by noise in the data and the inaccuracy of finding feature points, we introduce additional feature constraints, indicating that the corresponding features on two 3D surfaces should be mapped onto the same locations in the 2D domain. However, with these additional constraints, it is not always possible to satisfy

the conformality condition. Hence, we seek to minimize the violation of Riemann's condition in the least squares sense.

C. Least squares conformal maps

The Least Squares Conformal Map(LSCM) parameterization algorithm generates a discrete approximation of a conformal map by adding more constraints. Here we give a brief description (see [29] for details using different constraints).

Given a discrete 3D surface mesh \mathbf{S} and a smooth target mapping $\mathcal{U} : S \rightarrow (u, v)$, then, as described in section II-B, \mathcal{U} is conformal on \mathbf{S} if and only if the Cauchy-Riemann equation,

$$\frac{\partial \mathcal{U}}{\partial x} + i \frac{\partial \mathcal{U}}{\partial y} = 0 \quad (8)$$

holds true on the whole of \mathbf{S} . However, in general this conformal condition cannot be strictly satisfied on the whole triangulated surface \mathbf{S} , so the conformal map is constructed in the least squares sense:

$$MinC(S) = \sum_{d \in S} \int_d \left| \frac{\partial \mathcal{U}}{\partial x} + i \frac{\partial \mathcal{U}}{\partial y} \right|^2 dA \quad (9)$$

where d is a triangle on the mesh \mathbf{S} . If we suppose the mapping \mathcal{U} is linear on d then

$$C(S) = \sum_{d \in S} \left| \frac{\partial \mathcal{U}}{\partial x} + i \frac{\partial \mathcal{U}}{\partial y} \right|^2 A(d) \quad (10)$$

where $A(d)$ is the area of the triangle d . Furthermore let $\alpha_j = u_j + iv_j$ and $\beta_j = x_j + iy_j$, so $\alpha_j = \mathcal{U}(\beta_j)$ for $j = 1, 2, \dots, n$. Then, we rearrange the vector α such that $\alpha = (\alpha_f, \alpha_p)$ where α_f consists of $n - p$ free coordinates and α_p consists of p constraint point coordinates. Therefore, Eq. 10 can be rewritten as

$$C(S) = \|M_f \alpha_f + M_p \alpha_p\|^2 \quad (11)$$

where $M = (M_f, M_p)$, a sparse $m \times n$ complex matrix (m is the number of triangles and n is the number of vertices). The least squares minimization problem in Eq. 11 can be efficiently solved using the Conjugate Gradient Method. Thus we can map a 3D surface to a 2D domain with multiple correspondences as constraints by using the LSCM technique.

Since LSCMs have almost all the properties of conformal maps and also provide more correspondences as additional constraints, we expect them to be very useful in 3D shape matching and recognition.

D. Comparison of quasi-conformal maps

Based on conformal geometry theory, harmonic maps, conformal maps and least squares conformal maps (LSCMs) between two topological disks preserve continuity of the underlying surfaces, with minimal stretching energy and angle distortion. All the above quasi-conformal maps are invariant for the same source surface with different poses, thus making it possible to account for global rigid transformations. A very important property, which governs our matching algorithm, is that all the maps can establish a common 2D parametric domain for the two surfaces. Therefore we can simplify the 3D shape-matching problem to a 2D image-matching problem. However, they vary in performance for 3D surface matching as can be seen in table I.

TABLE I
PERFORMANCE COMPARISON OF QUASI-CONFORMAL MAPS.

	Harmonic Maps	Conformal Maps	Least Squares Conformal Maps
Resolution changes	Not sensitive	Not sensitive	Not sensitive
Boundary constraint	Needed	Not needed	Not needed
Boundary occlusion	Difficult to handle	No significant impact	No significant impact
Interior feature points used in mapping	Do not use	Use 2 Points (from Riemann's theorem)	Use more feature constraints
Error of interior feature points detection	Not sensitive	Sensitive	Not sensitive
Computational Complexity	Linear	Nonlinear (with linear approximation available)	Linear

Compared to the exact solutions for harmonic maps and conformal maps, LSCMs are generated by minimizing the violation of Riemann's condition in the least squares sense. This optimization-based parameterization method has the following properties:

- 1) LSCMs have the same properties as conformal maps, e.g., existence and uniqueness which have already been proven in [29].
- 2) LSCMs can map a 3D shape to a 2D domain in a continuous manner with minimized local angle distortion.
- 3) LSCMs can handle missing boundaries and occlusion and also allow multiple constraints.
- 4) LSCMs are independent of mesh resolution.

- 5) The least squares minimization problem in calculating LSCMs has the advantage of being linear.

For actual 3D surfaces, it is very likely to have noise and missing data. From the above comparison, we can see that LSCMs are the best candidate among all three parametric maps to perform 3D shape matching efficiently. LSCMs do not require the boundary condition explicitly which means they can handle missing boundaries and occlusions. Also, they take multiple feature constraints as input, which allows them to better handle noise introduced by the feature point detection. We confirm this experimentally in the experiment section by analyzing the robustness of the three parametric maps for 3D shape matching with occlusion, noise and resolution variation. In the remainder of this paper, we propose a framework of 3D shape matching using LSCMs.

III. SHAPE MATCHING FRAMEWORK USING LEAST SQUARES CONFORMAL MAPS

To match 3D shapes accurately and efficiently, a new 2D representation, least squares conformal shape images, is developed in our framework using LSCMs. Therefore, we simplify the original 3D shape-matching problem to a 2D image-matching problem. In particular, our shape matching framework includes two steps: First, interior feature correspondences are detected by using spin-images[26]; After that, we generate and match least squares conformal shape images.

A. Correspondence detection using spin-images

In order to use least squares conformal mappings, we need to establish interior feature constraints between the 3D shapes. For this purpose, we first select candidate points with curvature larger than a threshold T_c , and then compare their spin-images to detect feature correspondences. The spin-image is a well-known technique that has been proven useful for 3D point matching[26]. It encodes the surface shape surrounding an oriented point p by projecting nearby surface points into a 2D histogram, which has cylindrical coordinates of radius r and height h centered at p , with its axis aligned with the surface normal of p . The number of bins and support size in the spin-image histograms are parameters fixed at generation. It has been shown that the matching results using spin-images are insensitive to the choice of the above parameters [23]. In our experiments, the highest confidence feature correspondences are used.

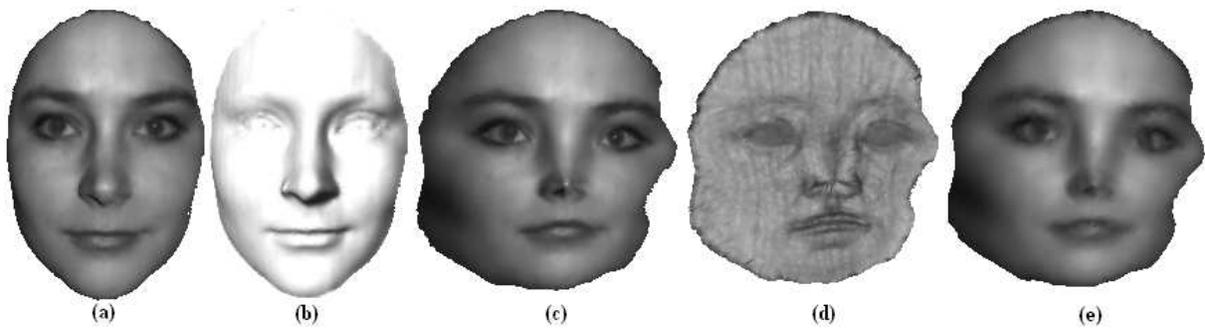


Fig. 2. Least Squares Conformal Shape Image: (a) Original surface with texture. (b) Original surface without texture. (c) Least squares conformal maps with texture. (d) Least squares conformal shape image. (e) Least squares conformal maps of the same surface, subsampled by a factor of 4, still very similar to (c).

The typical number of selected feature points is 5-6 for 3D face surfaces and 10-12 for brain surfaces.

B. Least squares conformal shape images

In this section, we will introduce a method to describe 3D surfaces using least squares conformal shape images (LSCSIs). In section II-C, we have shown that there exists a least squares conformal mapping that can map each 3D surface with disk topology to the canonical 2D domain. The LSCSIs are generated by associating a shape attribute with each vertex. Mean curvature is a useful geometric attribute that depends only on the surface's intrinsic geometry. In our method, the mean curvature is computed in the same way as in [20]. Moreover, least squares conformal maps can also help generate additional shape representations by associating other attributes, e.g. texture, which leads to a natural solution of combining multiple important cues for 3D surface matching and recognition, such as shape and texture. In our current framework, these cues are weighted equally for surface matching. More elaborate schemes to combine different cues can be done in the future work.

As an example, Fig. 2(d) shows the LSCSI of the surface Fig. 2(b), with darker color representing larger mean curvature. Fig. 2(a) is the original surface with texture information and Fig. 2(c) is its LSCM. Fig. 2(e) is the LSCM of a lower resolution (25%) version of the original surface. The similarity between Fig. 2(c) and Fig. 2(e) shows that LSCMs are independent to resolution variation.

C. Matching surfaces by matching LSCSIs

Given two general surfaces S_1 and S_2 with disk topology, we first detect high curvature correspondences using spin-images. Then, by incorporating interior correspondences as constraints, LSCSIs are generated for both surfaces as described in section II-C. After that, the normalized correlation coefficient M_{S_1, S_2} and the similarity criterion $S(S_1, S_2)$ introduced in [25] are computed on the two resulting LSCSIs by

$$M_{S_1, S_2} = \frac{N \sum p_i^{S_1} p_i^{S_2} - \sum p_i^{S_1} \sum p_i^{S_2}}{\sqrt{(N \sum (p_i^{S_1})^2 - (\sum p_i^{S_1})^2)(N \sum (p_i^{S_2})^2 - (\sum p_i^{S_2})^2)}} \quad (12)$$

$$S(S_1, S_2) = \left(\ln \frac{1 + M_{S_1, S_2}}{1 - M_{S_1, S_2}} \right)^2 - \frac{1}{2N} \quad (13)$$

where N is the number of overlapping points in the LSCSIs of 3D surface S_1 and S_2 , and $p_i^{S_k}$ is the value (e.g., the mean curvature or the texture) of point i in the LSCSI of surface S_k ($k = 1, 2$). In the case of matching surfaces with different resolutions, N is the number of overlapping points in the LSCSIs of the surface with the lower resolution.

According to section II, an important property of Least Squares Conformal Maps (LSCMs) is that they can map a 3D shape to a 2D domain in a continuous manner with minimized local angle distortion. This implies that *LSCSIs are not sensitive to surface deformations*, e.g., if there is not too much stretching between two faces with different expressions, they will induce similar LSCSIs. As an example, Fig. 3 shows a comparison between the LSCSIs of faces with different expressions and of different faces. More specifically, the first, the second, and the third columns of Fig. 3 correspond to face scans of one subject with different expressions while the fourth column corresponds to another subject. For each column in Fig. 3, the bottom row represents the LSCSIs of the surfaces (shown in the middle row), with darker color representing larger mean curvature. The original surfaces with texture information are also shown in the top row of Fig. 3. Based on Eq. 12, the normalized correlation coefficient ($M_{i,j}$) between Fig. 3(i) and Fig. 3(j) and the normalized correlation coefficient ($M_{i,k}$) between Fig. 3(i) and Fig. 3(k) are 0.92 and 0.86, respectively, while the normalized correlation coefficient ($M_{i,l}$) between Fig. 3(i) and Fig. 3(l) is only 0.65. As is evident, the normalized correlation coefficients of LSCSIs between the face scans of the same person with different expressions are much larger than the coefficients between face scans of different persons, thus making it possible to match surfaces with small

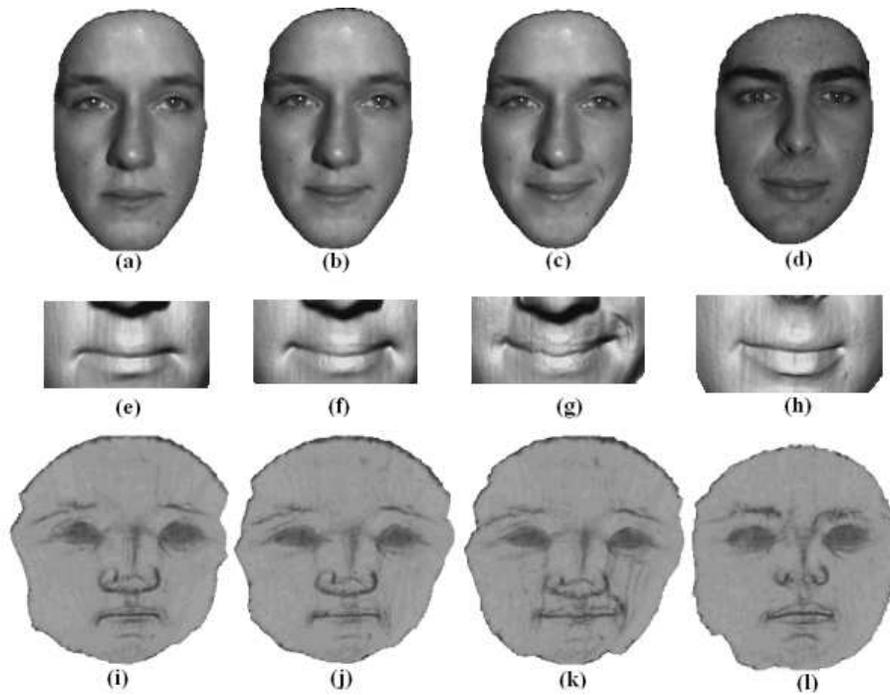


Fig. 3. Surface matching with deformation: The original 3D surfaces with texture are in the top row. The detail of the deformed mouth areas are shown in the second row and the LSCSIs of the original surfaces are in the last row. In each row, the first, the second, and the third surfaces are from the same person with different expressions and the fourth one is another person. The normalized correlation coefficient ($M_{i,j}$) between (i) and (j) and the normalized correlation coefficient ($M_{i,k}$) between (i) and (k) are 0.92 and 0.86, respectively, while the normalized correlation coefficient ($M_{i,l}$) between (i) and (l) is only 0.65.

deformations using LSCSIs. This relative expression-invariance is also an important property for shape representations used in face recognition.

However, for 3D surfaces with holes, which violate the disk topology assumption, we can not calculate the LSCMs directly. To overcome this problem, we can simply fill in the holes through interpolation[30] and then use our method to generate the LSCSIs of the new surfaces. The filled-in regions are masked out when we compute the normalized correlation coefficient using Eq. 12. As discussed in section II-D, LSCMs depend on the geometry in a continuous manner, which leads to robustness to local perturbation. Fig. 4 demonstrates the robustness of our method to holes on surfaces. The normalized correlation coefficient of the LSCSIs shown in Fig. 4(b,f) is 0.99, which means a very good match between the two surfaces of Fig. 4(a,e) after hole filling. If we desire to preserve the non-disk topology of the object during matching, then the object should be partitioned into simpler parts with disk topology[29] which could then

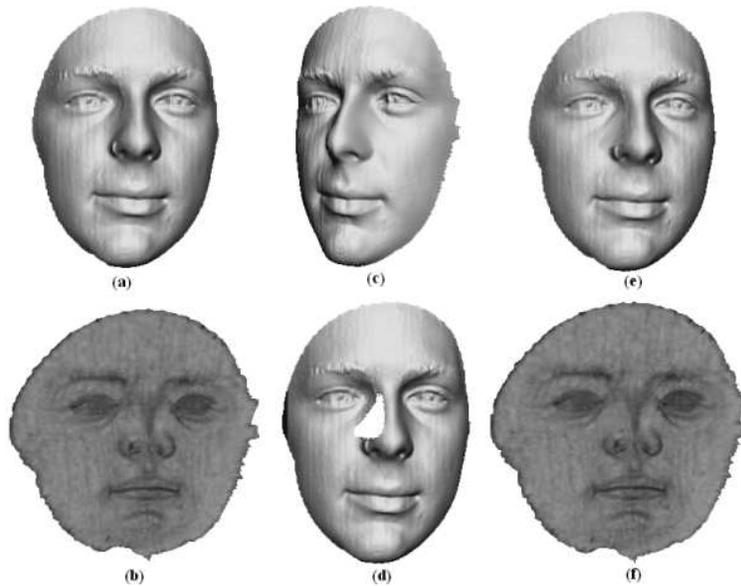


Fig. 4. An example of surface matching with holes : (a) A frontal 3D scan. (b) The LSCSI of (a). (c) A side 3D scan of the same subject as in (a), which has a hole illustrated in (d). (e) The same surface of (c,d) after hole filling. (f) The LSCSI of (e).

be matched. Optimal partitioning will be studied in future work.

IV. EXPERIMENTAL RESULTS AND PERFORMANCE ANALYSIS

In this section we analyze the robustness of our proposed 3D shape matching method using least squares conformal maps on real data with occlusion, noise and resolution variation. Furthermore, we demonstrate the performance of our method through two applications: 3D face recognition and 3D non-rigid surface alignment and stitching.

A. Robustness analysis

In this section we use two surface types: brains (4 instances) and faces (6 instances) to analyze the performance of our proposed 3D shape matching method. We present three experiments in which 3D surface matching is performed under occlusion, noise and resolution variation using least squares conformal maps, followed by a full comparison between several related work of quasi-conformal maps including harmonic maps, conformal maps and least squares conformal maps.

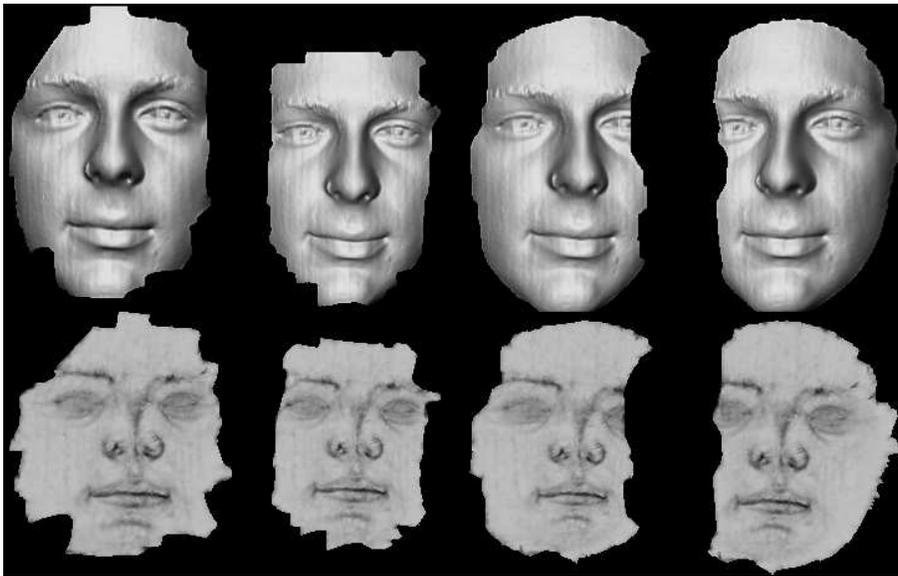


Fig. 5. 3D face surfaces and their LSCSIs under occlusion. The original 3D face surfaces with different occlusions are in the top row. Their LSCSIs are in the bottom row.

1) *Experiment on data occlusion:* In this experiment, we test the robustness of Least Squares Conformal Maps (LSCMs) under occlusion for both face and brain surfaces. Such occlusions might be caused by rotation of the object in front of the scanner. Figures 5 and 7 show examples of 3D face and brain surfaces respectively, under different occlusions with their least squares conformal shape images (LSCSIs). For each original surface, partially occluded surfaces were generated with occlusion rates between 5% and 45%. Average matching results of these face and brain surfaces using LSCMs are shown in Fig 6 and 8, respectively. In the video clip accompanying this paper, we superimpose the matched surfaces with significant occlusions (only 60% of area is common to both). Matching error is very hard to detect visually, which suggests that our framework could be useful for partial scan alignment.

2) *Experiment on noisy data:* The second experiment tests the robustness of Least Squares Conformal Maps (LSCMs) in the presence of noise. We add gaussian noise($\mathcal{N}(0, \sigma)$) on each vertex of the face and brain surfaces. σ increases from 0.0 mm to 2.0 mm while the window size for computing the curvatures of 3D face and brain surfaces is 10.0 mm. Example surfaces with noise under different σ are shown in Fig 9. We match the various noisy surfaces to the original noise-free surface and the average matching results of the face and brain surfaces are

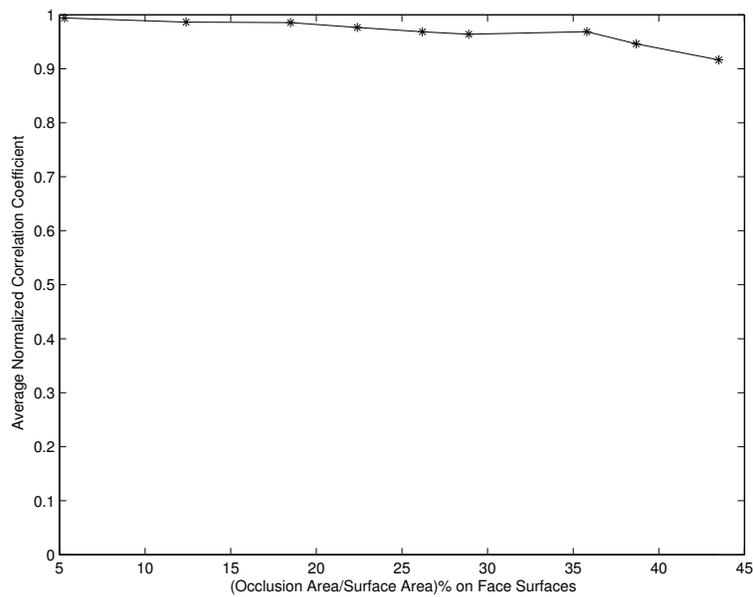


Fig. 6. Average matching results of the face surfaces under occlusion using LSCMs.

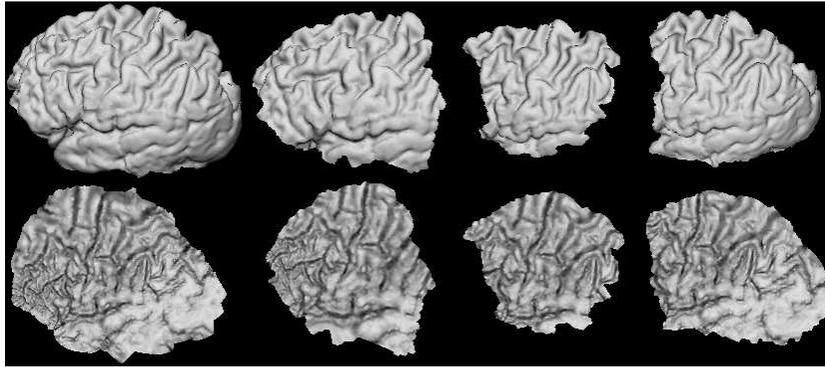


Fig. 7. 3D brain surfaces and their LSCSIs under occlusion. The original 3D brain surfaces with different occlusions are in the top row. Their LSCSIs are in the bottom row.

shown in Fig. 10 for various σ values. From the results we can see that LSCMs appear robust to gaussian noise.

3) *Experiment on resolution variation:* The third experiment tests the robustness of Least Squares Conformal Maps (LSCMs) to resolution changes. Fig. 11 shows examples of 3D face and brain surfaces with resolution variation, where all the meshes have the same shape but different resolution. The surfaces with low resolution are matched to the original surfaces and average

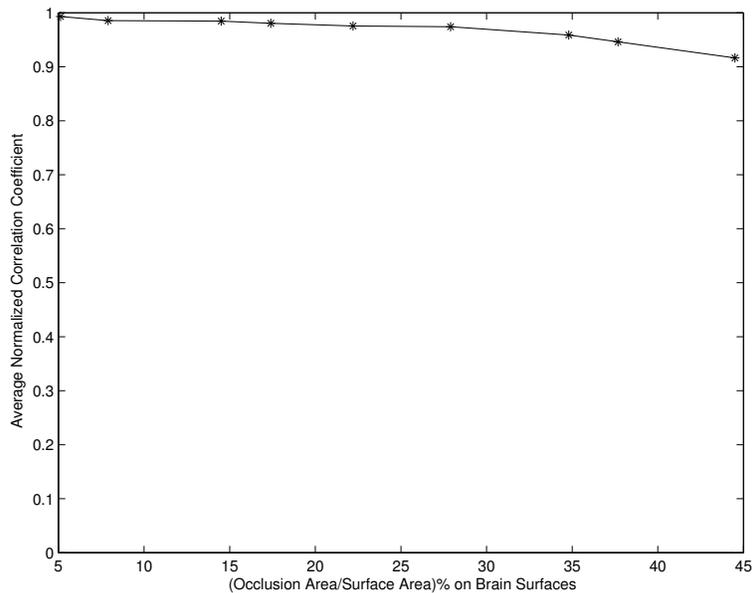


Fig. 8. Average matching results of the brain surfaces under occlusion using LSCMs.

matching results using the LSCMs are shown in Fig. 12. Results show that LSCMs achieve fairly stable matching results and impervious to resolution changes. A small deterioration of the matching results is due to the use of a discrete curvature approximation, since approximation error increases as the resolution drops.

4) *Comparison between quasi-conformal maps:* For completeness purposes, we also performed comparison experiments between several related work of quasi-conformal maps, including least squares conformal maps, conformal maps[20] and harmonic maps[51], [49], to confirm the conclusion in Section II-D. Average matching results of the face and brain surfaces using the above three parametric maps under occlusion, noise and resolution variation are shown in Fig 13, 14 and 15, respectively. In Fig 13, since the harmonic maps require satisfaction of the surface boundary condition as discussed in section II-A, the performance of harmonic maps is more impacted than the performance of conformal maps and least squares conformal maps. Instead, changes of boundary have very small effects on both conformal maps and least square conformal maps. From the results in Fig 14 we can see that all three maps appear robust to gaussian noise. However, since conformal maps depend on 2 feature points only, which might be detected with errors caused by the noise, they have lower matching rates than the harmonic maps and the least

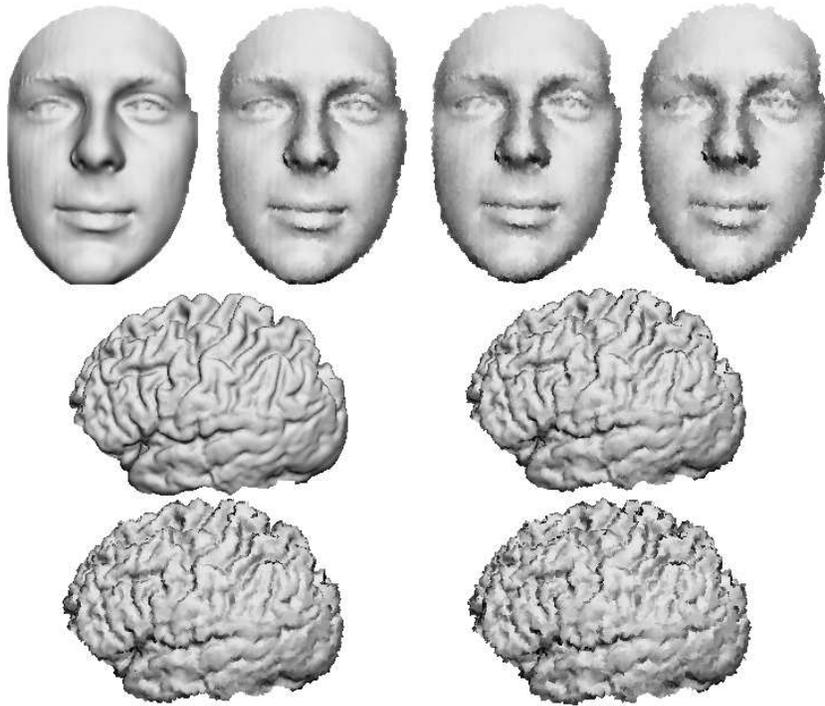


Fig. 9. Examples of face and brain surfaces under Gaussian noise with different σ set to 0.0, 0.4, 1.0 and 2.0 mm, respectively.

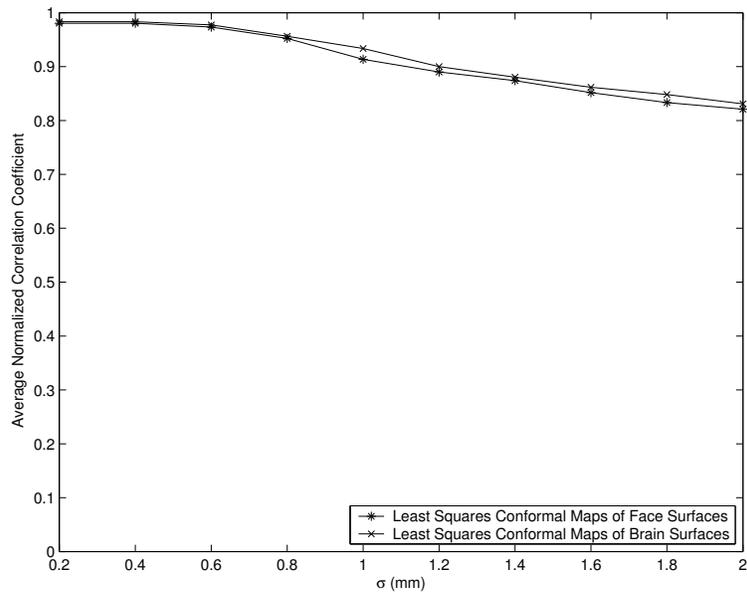


Fig. 10. Average matching results of LSCMs under Gaussian noise increases. The window size for computing the curvatures of faces surfaces and brain surfaces is 10.0 mm and the σ increases from 0.0 mm to 2.0 mm.

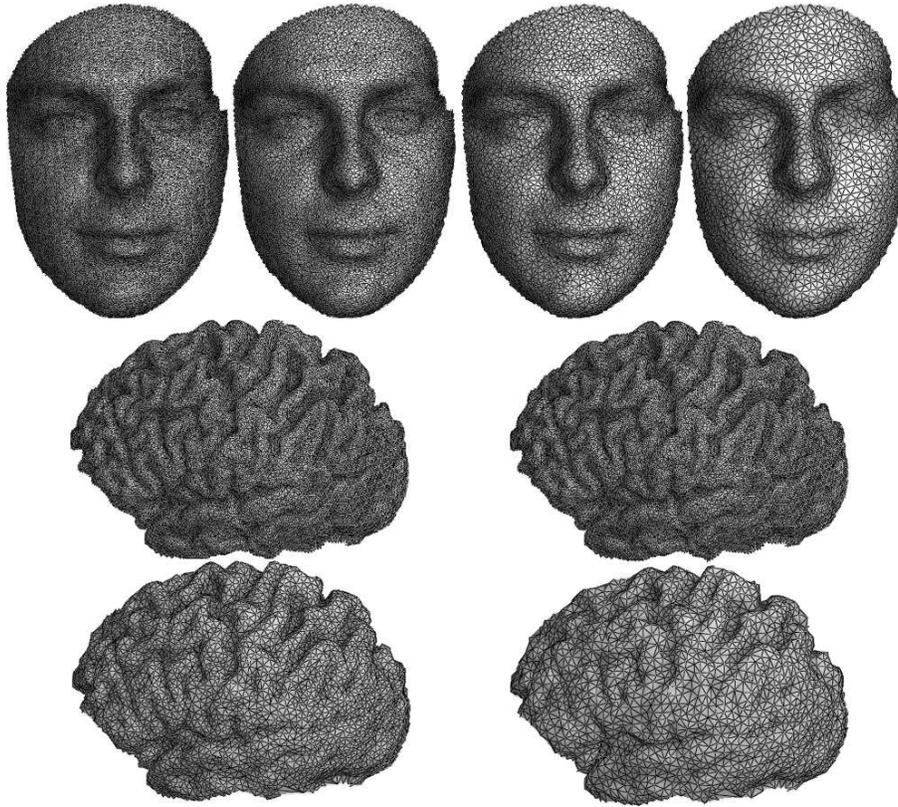


Fig. 11. 3D face and brain surfaces with 1, 1/2, 1/4 and 1/8 of the original resolution, respectively.

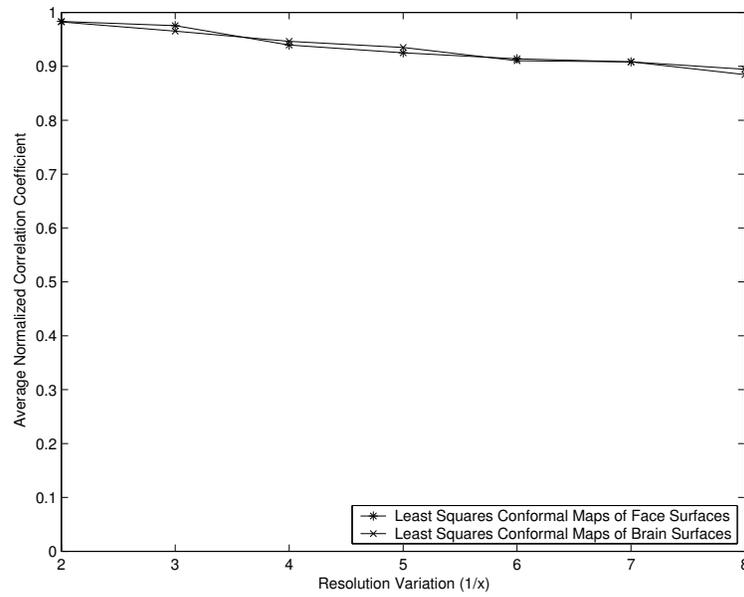


Fig. 12. Average matching results of LSCMs under resolution variation.

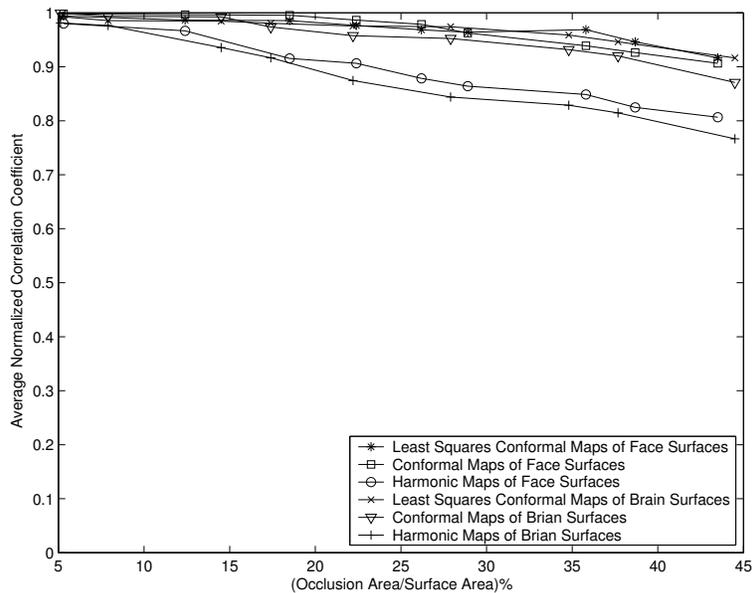


Fig. 13. Average matching results of the face and brain surfaces under occlusion using all three parametric maps.

square conformal maps. Finally, Fig 15 shows that the above three parametric maps achieve fairly stable matching results and all of them are impervious to resolution changes.

B. Recognition of 3D faces

In this section, we apply Least Squares Conformal Maps (LSCMs) to 3D face recognition on a 3D face database which contains 100 3D face scans from 10 subjects. The data are captured by a phase-shifting structured light ranging system in different time[52]. Each face has approximately 80K 3D points with both shape and texture information available (example face data from two subjects in the database are shown in Fig. 16). In order to further evaluate our recognition method, we also perform a comparison with other existing methods, including the surface curvature technique[45] and the spherical harmonic shape contexts [17]. For the computation of curvatures from 3D surfaces we had to chose the size of the neighborhood for the surface fit. Clearly, choosing the mask size is a trade-off between reliability and accuracy. When choosing a small mask curvature computation will be strongly affected by noise, due to the small number of points considered for regression. The reliability of the curvature estimation can be improved by increasing the size of the mask. However, a large mask size will produce an incorrect result in the area curvature changes quickly. In our experiments, we used a mask size of 10×10 . The

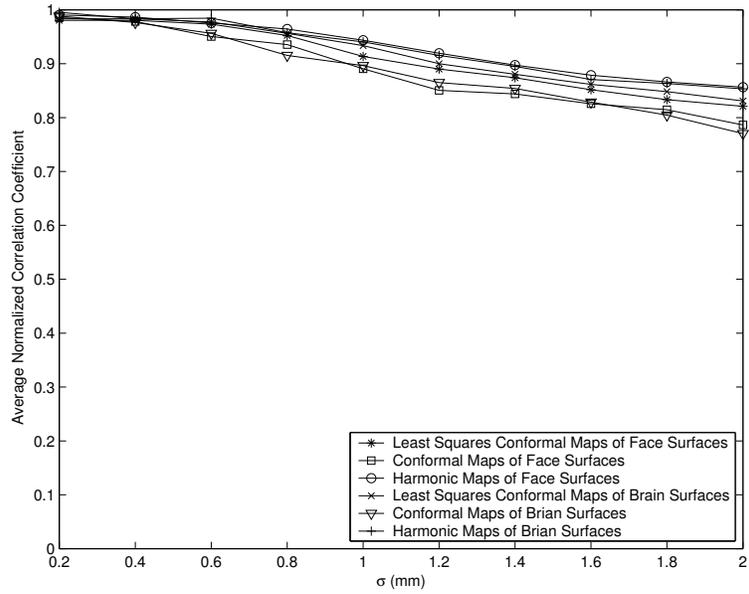


Fig. 14. Average matching results of all three parametric maps under gaussian noise increases. The window size for computing the curvatures of faces surfaces and brain surfaces is 10.0 mm and the σ increases from 0.0 mm to 2.0 mm.

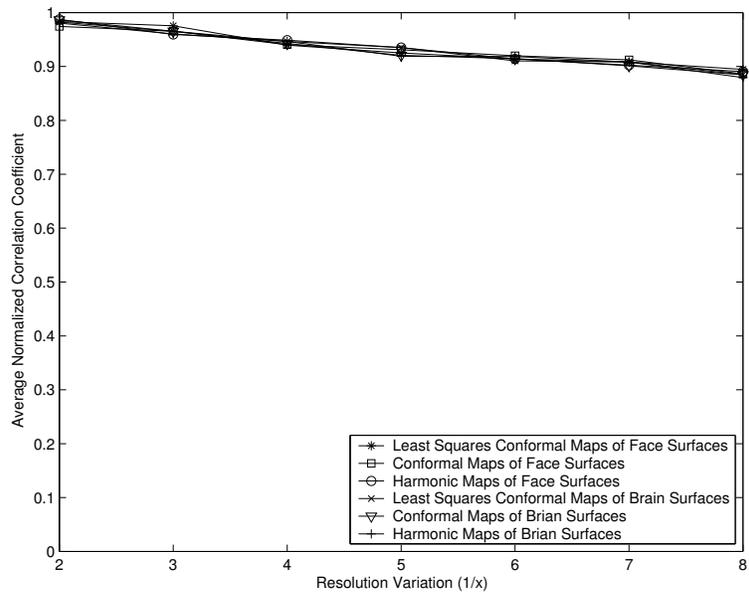


Fig. 15. Average matching results of all three parametric maps under resolution variation.

spherical harmonic shape contexts descriptor is computed using the method developed in [17], based on 3D shape contexts. The 3D shape contexts technique is the straightforward extension of 2D shape contexts[32], to three dimensions. The support region for a 3D shape contexts is a sphere centered on the basis point p and its north pole oriented with the surface normal estimate N for p . The support region is divided into bins by equally spaced boundaries in the azimuth and elevation dimensions and logarithmically spaced boundaries along the radial dimension. Based on the histogram from 3D shape contexts, we use the bin values as samples to calculate a spherical harmonic transformation for the shells and discard the original histogram. The descriptor is a vector of the amplitudes of the transformation, which are rotationally invariant in the azimuth direction, thus removing the degree of freedom. We compute the spherical harmonic shape contexts representations in 64×64 grids sampled evenly along the directions of longitude and latitude with bandwidth $b = 16$.

In each experiment, we randomly select a single face from each subject for the gallery and use all the remaining faces as the probe set. The average recognition results from 15 experiments (with different randomly selected galleries) are reported in Table II. From the recognition results, we can see that the least squares conformal maps perform 10.7% better than the spherical harmonic shape contexts and 14.3% better than the surface curvature technique even if only the shape information is used. Moreover, least squares conformal maps allow to combine both shape and texture information, which improves the accuracy of 3D face recognition.

TABLE II

RECOGNITION RESULTS OF LEAST SQUARES CONFORMAL MAPS, SPHERICAL HARMONIC SHAPE CONTEXTS AND SURFACE CURVATURE TECHNIQUE.

Recognition Result	Least Squares Conformal Maps	Spherical Harmonic Shape Contexts	Surface Curvature
Using shape information only	97.3%	86.6%	83.0%
Using texture information only	98.0%	N/A	N/A
Using both shape and texture	98.4%	N/A	N/A

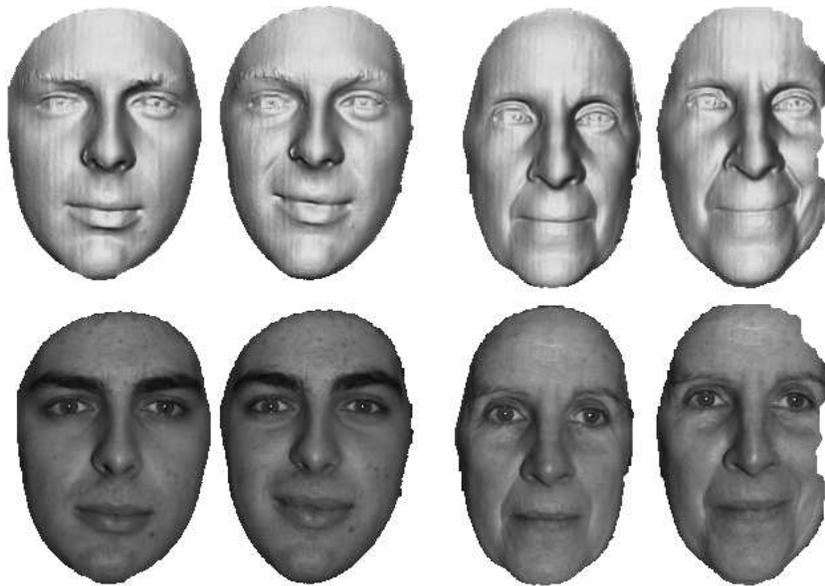


Fig. 16. Two subjects in the 3D face database. Shape information is in the first row and texture information is in the second row.

C. Non-rigid surface alignment and stitching

In this section, we apply the Least Squares Conformal Maps(LSCMs) to another application: 3D non-rigid surface alignment and stitching. A very important property, which governs our alignment and stitching algorithm, is that the LSCMs can establish a 2D common parametric domain for the 3D surfaces. Therefore we can simplify the 3D surface alignment and stitching problem to a 2D registration and stitching problem. Furthermore, because the LSCMs is a diffeomorphism, i.e., one-to-one and onto, we can detect and remove the duplicated regions in the original 3D surfaces by removing the overlapping areas in the resulting 2D common parametric domain. After that, we can stitch the 3D surface patches by connecting the exclusive regions in the resulting LSCMs. There is a lot of research on 3D surface remeshing[2], [5], [1], [35], but in our case the problem is simplified to a 2D triangulation problem by connecting the neighboring patches in the 2D common parametric domain. As an example, Fig. 17 demonstrates the alignment and stitching of two 3D surfaces undergoing non-rigid deformations. 3D faces are captured by a phase-shifting structured light ranging system [52] and each face has approximately 80K 3D points with both shape and texture information available. The subjects were not asked to keep their head and facial expression still during the 3D face scanning.

Furthermore, Fig. 18 shows another example of the accurate face alignment and stitching result of our method on two 3D scans of one face undergoing different transformations and deformations. The leftmost column shows the two input 3D face scans with texture. The same 3D face scans without texture information are shown in the second column. The Least Squares Conformal Shape Images (LSCSIs) of both 3D scans are in the third column. Their aligned LSCSIs and the resulting stitched 3D faces are in the fourth column. Because of the one-to-one mapping between the LSCSI and original face, we can align and stitch 3D faces by registering and stitching 2D LSCSIs.

In order to demonstrate the performance of our method, we also compare our results to the results from the Iterative Closest Point (ICP) method [38] in Fig. 19. Fig. 19(a) shows a 3D scan of a neutral face, while Fig. 19(b) shows a 3D scan of the same face undergoing a large deformation in the mouth area. From Fig. 19(c) and (d) which are the front view of (a) and (b), we can see the occlusion area clearly. The face alignment and stitching result of our method is in Fig. 19(f) with the close up view of mouth area in Fig. 19(h). The result of the ICP method is in Fig. 19(e) with the close up view in Fig. 19(g). As we can see, in the close up view Fig. 19(g), there is a redundant region in the result because the ICP method failed to detect the overlapping areas between deformed surfaces and can only register two surface with rigid transformations. However, as can be seen in Fig. 19(g) and (h), our method correctly aligns even at areas of significant local deformations.

V. CONCLUSIONS AND FUTURE WORK

In this paper, we presented a family of quasi-conformal maps, including harmonic maps, conformal maps and least squares conformal maps, and proposed a fully automatic and novel 3D shape matching framework using least squares conformal shape images – a new shape representation which simplified the 3D surface matching problem to a 2D image matching problem. The performance of least squares conformal maps was evaluated vis-a-vis other existing techniques in 3D face recognition and 3D non-rigid surface alignment and stitching. Furthermore, our comparison results have shown that all above three parametric maps are robust to occlusion, noise and different resolutions and that the least squares conformal mapping is the best choice for 3D surface matching.

In future work, we will continue to exploit the properties of conformal maps and further

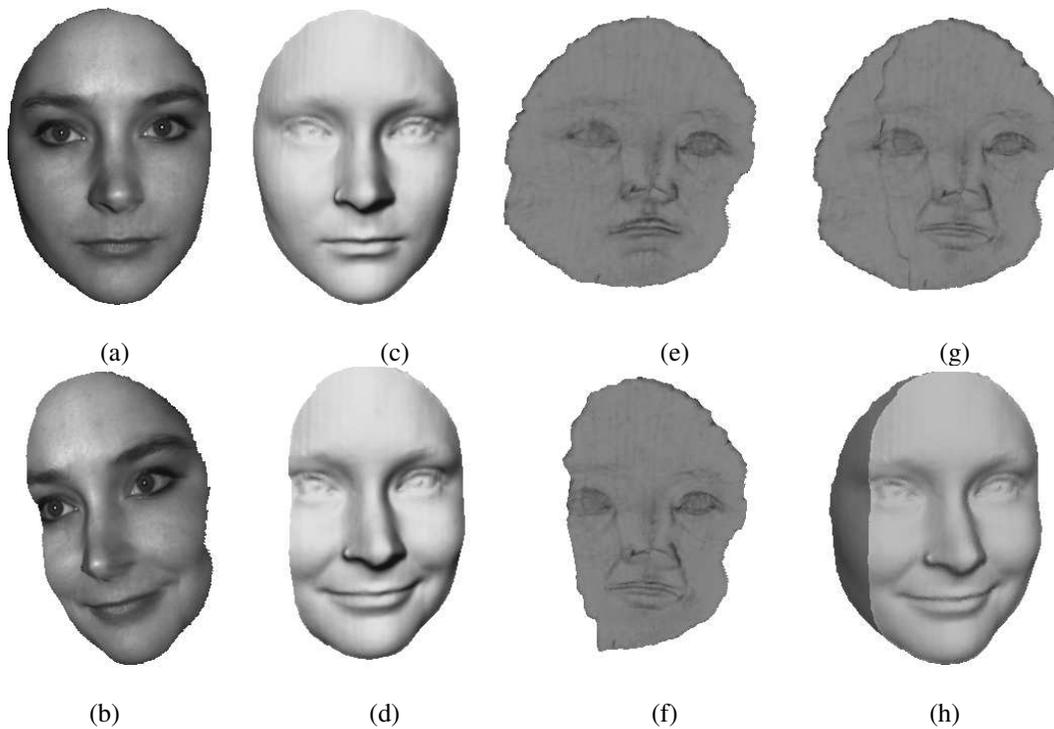


Fig. 17. An example of surface alignment and stitching: (a,b) Two original 3D faces with texture in different poses and deformations. (c,d) Original 3D faces without texture. (e,f) The Least Squares conformal Shape Images (LSCSI) of the faces. (g) The aligned LSCSI of the two faces. (h) The resulting 3D face by stitching a part of (c) into (d). Because of the one-to-one mapping between the LSCSI and original face, we can align and stitch 3D faces by registering and stitching 2D LSCSI.

analyze the properties of conformal shape representations for surfaces with non-disk topology. We plan to use our framework for applications such as 3D object classification and registration under non-rigid deformations.

VI. ACKNOWLEDGEMENTS

We are grateful to Peisen Huang and Song Zhang for using their phase-shifting structured light ranging system to capture the 3D face data presented in this paper. This work was partially supported by grants: NSF ACI-0313184, NSF CMS-9900337, NSF CAREER Award CCF-0448339, NIH RR13995, and DOJ 2004-DD-BX-1224.

REFERENCES

- [1] P. Alliez, D. Cohen-Steiner, O. Devillers, Bruno Levy, and Mathieu Desbrun. Anisotropic polygonal remeshing. *ACM Transactions on Graphics*, pages 485–493, 2003.

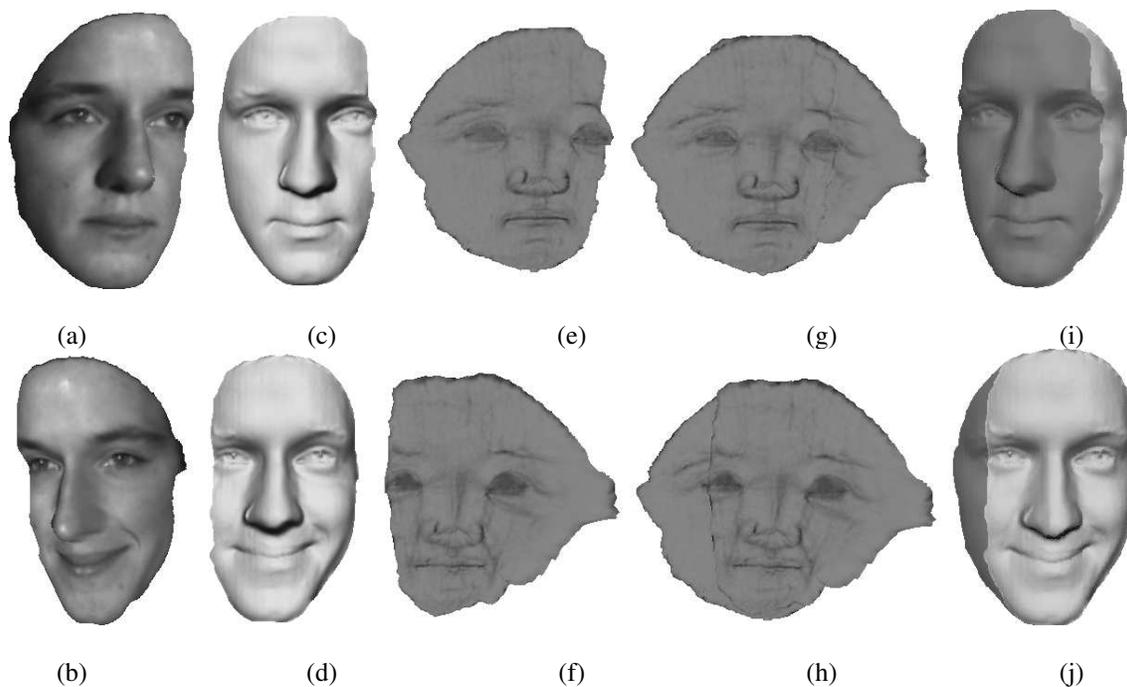


Fig. 18. Another example of surface alignment and stitching: (a,b) Two original 3D faces with texture in different poses and deformations. (c,d) Original 3D faces without texture. (e,f) The Least Squares conformal Shape Images (LSCSI) of the faces. (g) The aligned LSCSI of the two faces by connecting the non-overlapping area in (f) into (e). (h) The aligned LSCSI of the two faces by connecting the non-overlapping area in (e) into (f). (i) The resulting 3D face by stitching a part of (d) into (c). (j) The resulting 3D face by stitching a part of (c) into (d). Because of the one-to-one mapping between the LSCSI and original face, we can detect and remove the duplicated regions in the original 3D surfaces by removing the overlapping areas in the resulting 2D common parametric domain. The user can decide which of the two expressions to keep on the final stitched mesh. In this case (i) has the expression of original (a) and (j) of original (b).

- [2] P. Alliez, M. Meyer, and M. Desbrun. Interactive geometry remeshing. In *SIGGRAPH '02*, pages 347–354, 2002.
- [3] G. D. Anderson, M. K. Vamanamurthy, and M. K. Vuorinen. *Conformal Invariants, Inequalities, and Quasiconformal Mappings*. Wiley-Interscience, March 1997.
- [4] V. Athitsos, J. Alon, S. Sclaroff, and G. Kollios. Boostmap: a method for efficient approximate similarity rankings. In *CVPR04*, pages II: 268–275, 2004.
- [5] M. Botsch, C. Rössl, and L. Kobbelt. Feature sensitive sampling for interactive remeshing. In *Proceedings of the 2000 Conference on Vision Modeling and Visualization*, pages 129–136, 2000.
- [6] B. Brown and S. Rusinkiewicz. Non-rigid range-scan alignment using thin-plate splines. In *Symposium on 3D Data Processing, Visualization, and Transmission*, 2004.
- [7] V. Camion and L. Younes. Geodesic interpolating splines. In *EMMCVPR01*, pages 513–527, 2001.
- [8] R.J. Campbell and P.J. Flynn. A survey of free-form object representation and recognition techniques. *Computer Vision*

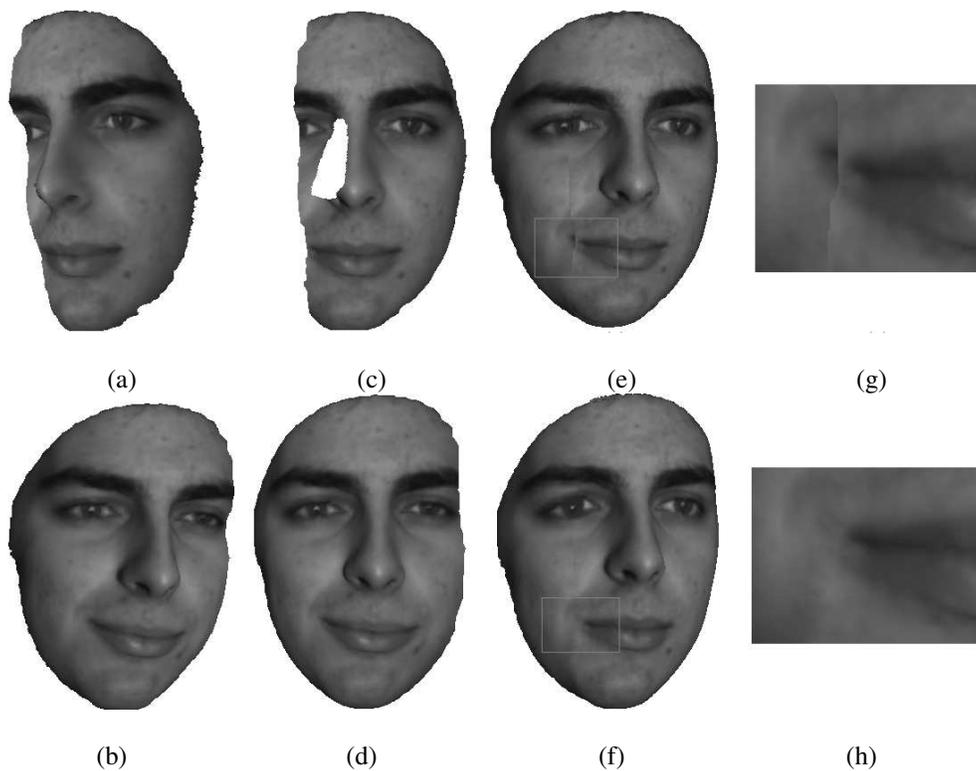


Fig. 19. A comparison between the alignment and stitching result of our method and of the ICP method: (a) A 3D scan of a neutral face. (b) A 3D scan of the same face undergoing a large deformation in the mouth area. (c) and (d) are the front view of (a) and (b) with the occlusion area shown clearly. (e) The face alignment and stitching result of the ICP method. (g) The close up view in the mouth area of (e). (f) The face alignment and stitching result of our method. (h) The close up view in the mouth area of (f).

and Image Understanding, 81:166–210, 2001.

- [9] C.S. Chen, Y.P. Hung, and J.B. Cheng. RANSAC-based DARCES: A new approach to fast automatic registration of partially overlapping range images. *Trans. PAMI*, 21(11):1229–1234, 1999.
- [10] C.S. Chua and R. Jarvis. 3d free-form surface registration and object recognition. *IJCV*, 17:77–99, 1996.
- [11] M. Desbrun, M. Meyer, and P. Alliez. Intrinsic parameterizations of surface meshes. In *Eurographics02*, pages 209–218, 2002.
- [12] C. Dorai, J. Weng, and A. K. Jain. Optimal registration of object views using range data. *Trans. PAMI*, 19(10):1131–1138, 1997.
- [13] M. Eck, T. DeRose, T. Duchamp, H. Hoppe, M. Lounsbery, and W. Stuetzle. Multiresolution analysis of arbitrary meshes. In *SIGGRAPH'95*, pages 173–182, 1995.
- [14] J. Eells and J. H. Sampson. Harmonic mappings of riemannian manifolds. *Amer. J. Math.*, 86:109–160, 1964.
- [15] O. Faugeras and M. Hebert. The representation, recognition and locating of 3d object. *Int. Journal of Robotics Research*, 5(3):27–51, 1986.
- [16] M. S. Floater and K. Hormann. Surface parameterization: a tutorial and survey. In *Advances in Multiresolution for Geometric Modelling*, pages 157–186. Springer, 2004.

- [17] A. Frome, D. Huber, R. Kolluri, T. Bulow, and J. Malik. Recognizing objects in range data using regional point descriptors. In *ECCV04*, May 2004.
- [18] T. Funkhouser, P. Min, M. Kazhdan, J. Chen, A. Halderman, D. Dobkin, and D. Jacobs. A search engine for 3d models. In *ACM Transactions on Graphics*, pages 83–105, 2003.
- [19] X. Gu and B. C. Vemuri. Matching 3d shapes using 2d conformal representations. *MICCAI, Lectures Notes in Computer Science*, 3216:771–780, 2004.
- [20] X. Gu, Y. Wang, T. F. Chan, P. M. Thompson, and S.T. Yaun. Genus zero surface conformal mapping and its application to brain surface mapping. *IEEE Transaction on Medical Imaging*, 23(7), 2004.
- [21] S. Haker, S. Angenent, A. Tannenbaum, R. Kikinis, G. Sapiro, and M. Halle. Conformal surface parameterization for texture mapping. *IEEE Transactions on Visualization and Computer Graphics*, 6:181–189, 2000.
- [22] K. Hormann and G. Greiner. MIPS: An efficient global parametrization method. In *Curve and Surface Design: Saint-Malo 1999*, pages 153–162. 2000.
- [23] D. Huber, A. Kapuria, R.R. Donamukkala, and M. Hebert. Parts-based 3d object classification. In *CVPR04*, pages II: 82–89, June 2004.
- [24] M.K. Hurdal, K. Stephenson, P.L. Bowers, D.W.L. Sumners, and D.A. Rottenberg. Coordinate systems for conformal cerebellar flat maps. *NeuroImage*, 11:S467, 2000.
- [25] A. Johnson. *Spin-Images: A Representation for 3D Surface Matching*. PhD thesis, Robotics Institute, CMU, 1997.
- [26] A. Johnson and M. Hebert. Using spin images for efficient object recognition in cluttered 3d scenes. *PAMI*, 21:433–449, 1999.
- [27] M. Kazhdan, T. Funkhouser, and S. Rusinkiewicz. Rotation invariant spherical harmonic representation of 3d shape descriptors. In *Eurographics/ACM SIGGRAPH symposium on Geometry processing*, pages 156–164, 2003.
- [28] M. Levoy, K. Pulli, B. Curless, S. Rusinkiewicz, D. Koller, L. Pereira, M. Ginzton, S. Anderson, J. Davis, J. Ginsberg, J. Shade, and D. Fulk. The digital michelangelo project: 3d scanning of large statues. In *ACM SIGGRAPH 2000*, pages 131–144, 2000.
- [29] B. Levy, S. Petitjean, N. Ray, and J. Maillot. Least squares conformal maps for automatic texture atlas generation. In *SIGGRAPH'02*, pages 362–371, 2002.
- [30] P. Liepa. Filling holes in meshes. In *SGP '03: Proceedings of the symposium on Geometry processing*, pages 200–205, 2003.
- [31] D.G. Lowe. Distinctive image features from scale-invariant keypoints. *IJCV*, 60(2):91–110, 2004.
- [32] G. Mori, S. Belongie, and J. Malik. Efficient shape matching using shape contexts. *PAMI*, 27(11):1832–1837, 2005.
- [33] B. O’Neill. *Elementary Differential Geometry*. 1997.
- [34] R. Osada, T. Funkhouser, B. Chazelle, and D. Dobkin. Shape distributions. In *ACM Transactions on Graphics*, volume 21, pages 807–832, 2002.
- [35] E. Praun and H. Hoppe. Spherical parametrization and remeshing. *ACM Trans. Graph.*, 22(3):340–349, 2003.
- [36] K. Pulli. Multiview registration for large data sets. In *3DIM99*, pages 160–168, 1999.
- [37] S. Ruiz-Correa, L.G. Shapiro, and M. Meila. A new paradigm for recognizing 3d object shapes from range data. In *ICCV03*, pages 1126–1133, 2003.
- [38] S. Rusinkiewicz, O. Hall-Holt, and M. Levoy. Real-time 3d model acquisition. In *SIGGRAPH02*, pages 438–446, 2002.
- [39] R. Schoen and S. T. Yau. *Lectures on Harmonic Maps*. 1997.
- [40] E. Sharon and D. Mumford. 2d-shape analysis using conformal mapping. In *CVPR04*, pages II: 350–357, 2004.

- [41] A. Sheffer and E. de Sturler. Parameterization of faceted surfaces for meshing using angle-based flattening. *Engineering with Computers*, 17(3):326–337, 2001.
- [42] F. Stein and G. Medioni. Structural indexing: Efficient 3d object recognition. *Trans. PAMI*, 14(2):125–145, 1992.
- [43] Y. Sun and M.A. Abidi. Surface matching by 3d point's fingerprint. In *ICCV01*, pages II: 263–269, 2001.
- [44] G. Turk and M. Levoy. Zippered polygon meshes from range images. In *SIGGRAPH '94*, pages 311–318, 1994.
- [45] B.C. Vemuri, A. Mitiche, and J.K. Aggarwal. Curvature-based representation of objects from range data. *Image and Vision Computing*, 4:107–114, 1986.
- [46] S. Wang, Y. Wang, M. Jin, X. Gu, and D. Samaras. 3d surface matching and recognition using conformal geometry. In *CVPR06*, pages II: 2453– 2460, 2006.
- [47] Y. Wang, M.-C. Chiang, and P. M. Thompson. Mutual information-based 3d surface matching with applications to face recognition and brain mapping. *ICCV05*, 1:527–534, 2005.
- [48] Y. Wang, X. Gu, K.M. Hayashi, T.F. Chan, P.M. Thompson, and S.T. Yau. Surface parameterization using riemann surface structure. In *ICCV05*, pages II: 1061–1066, 2005.
- [49] Y. Wang, M. Gupta, S. Zhang, S. Wang, X. Gu, D. Samaras, and P. Huang. High resolution tracking of non-rigid 3d motion of densely sampled data using harmonic maps. In *ICCV05*, pages I: 388–395, 2005.
- [50] J.V. Wyngaerd, L.V. Gool, R. Koch, and M. Proesmans. Invariant-based registration of surface patches. In *ICCV99*, pages 301–306, 1999.
- [51] D. Zhang and M. Hebert. Harmonic maps and their applications in surface matching. In *CVPR99*, pages II: 524–530, 1999.
- [52] S. Zhang and P. Huang. High resolution, real time 3d shape acquisition. In *CVPR04 Workshop on Real-time 3D Sensors and Their Use*, page 28, 2004.



Sen Wang received B.S. degree in electrical engineering department from Shandong University in 2000 and M.S. degree in computer sciences department from Stony Brook University in 2006. He is currently a Ph.D. candidate in computer science department at Stony Brook University. His research interests include surface matching, face modeling and analysis, face synthesis and recognition, as well as illumination modeling and estimation for face.



Yang Wang is a Ph.D. candidate in the Department of Computer Science at the State University of New York at Stony Brook, since August 2000. He received a B.S. and an M.Sc. in Computer Science from Tsinghua University in 1998 and 2000 respectively. He specializes in illumination modeling and estimation, 3D non-rigid motion tracking and facial expression synthesis and analysis. He is a member of ACM and IEEE.



Miao Jin is a Ph D. candidate in Computer Sciences Department at Stony Brook University. She obtained her B.S. degree in computer sciences department from Beijing University of Telecom. and Post in 2000, and M.S. degree in Computer Sciences department from Stony Brook University in 2006. Her research interests include computational geometry, surface parameterization, hyperbolic and projective structures on surface.



Xianfeng David Gu is an Assistant professor in Computer Science at Stony Brook University. He earned his Ph.D. in Computer Science from Harvard University in 2003. He won the NSF CAREER award in 2004. His research interests are computer graphics, computer vision, geometric modeling and medical imaging. His major works include geometry images, global conformal surface parameterization, manifold splines and computational conformal geometry. Recently, he introduces Ricci flow method to computer science. For more information see <http://www.cs.sunysb.edu/~gu>.



Dimitris Samaras received a diploma degree in computer science and engineering from the University of Patras in 1992, the MSc degree in computer science from Northeastern University in 1994, and the PhD degree from the University of Pennsylvania in 2001. He is an associate professor in the Department of Computer Science at Stony Brook University, where he has been working since September 2000. He specializes in deformable model techniques for 3D shape estimation and motion analysis, illumination modeling and estimation for recognition and graphics, and biomedical image analysis. He is a member of the ACM and IEEE.