

3D Non-rigid Surface Matching and Registration Based on Holomorphic Differentials

Wei Zeng¹, Yun Zeng¹, Yang Wang², Xiaotian Yin¹,
Xianfeng Gu¹, and Dimitris Samaras¹

¹ Stony Brook University, Stony Brook NY 11790, USA

² Carnegie Mellon University, Pittsburgh PA 15213, USA

Abstract. 3D surface matching is fundamental for shape registration, deformable 3D non-rigid tracking, recognition and classification. In this paper we describe a novel approach for generating an efficient and optimal combined matching from multiple boundary-constrained conformal parameterizations for multiply connected domains (i.e., genus zero open surface with multiple boundaries), which always come from imperfect 3D data acquisition (holes, partial occlusions, change of pose and non-rigid deformation between scans). This optimality criterion is also used to assess how consistent each boundary is, and thus decide to enforce or relax boundary constraints across the two surfaces to be matched. The linear boundary-constrained conformal parameterization is based on the holomorphic differential forms, which map a surface with n boundaries conformally to a planar rectangle with $(n - 2)$ horizontal slits, other two boundaries as constraints. The mapping is a diffeomorphism and intrinsic to the geometry, handles an open surface with arbitrary number of boundaries, and can be implemented as a linear system. Experimental results are given for real facial surface matching, deformable cloth non-rigid tracking, which demonstrate the efficiency of our method, especially for 3D non-rigid surfaces with significantly inconsistent boundaries.

1 Introduction

In recent decades, there has been a lot of research into surface representations for 3D surface analysis, which is a fundamental issue for many computer vision applications, such as 3D shape registration, partial scan alignment, 3D object recognition, and classification [1,2,3]. In particular, as 3D scanning technologies improve, large databases of 3D scans require automated methods for matching and registration. However, matching surfaces undergoing non-rigid deformation is still a challenging problem, especially when data is noisy and with complicated topology. Different approaches include curvature-based representations [4,5], regional point representations [2,6], spherical harmonic representations [7,8], shape distributions [9], multi-dimensional scaling[10], local isometric mapping [11], summation invariants [12], landmark-sliding [13], physics-based deformable models [14], Free-Form Deformation (FFD) [15], and Level-Set based methods [16]. However, many surface representations that use local geometric

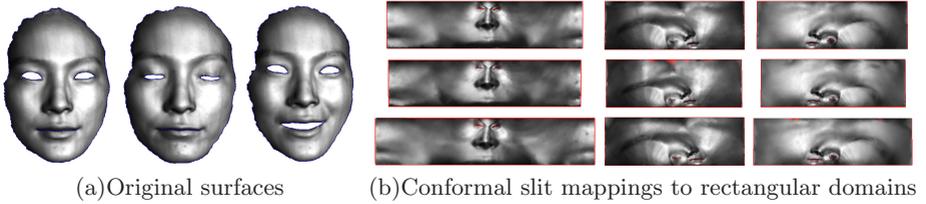


Fig. 1. Multiple slit mappings. The columns in (b) correspond to mouth, left-eye and right-eye boundary condition from left to right, where the optimal mouth, left-eye and right-eye area distortions are induced respectively. Optimal conformal parameterizations are chosen for different regions on the surface.

invariants can not guarantee a global convergence and might suffer from local minima in the presence of non-rigid deformations. To address this issue, many global parameterization methods have been developed recently based on conformal geometric maps [17,18,19,20,21,22]. Although the previous methods have met with a great deal of success in both computer vision and graphics, there are three major shortcomings in conformal maps when applied to matching of real discrete data such as the output of 3D scanners: 1) *complicated topology of the inputs*, 2) *area distortions* and 3) *inconsistent boundaries*. In this paper we will address the above three issues by introducing a novel linear algorithm for multiply connected surfaces based on holomorphic differentials.

Most existing conformal mapping methods can only handle surfaces with the simplest topology, namely, genus zero surface with a single boundary. In reality, due to partial occlusion, noises, arbitrary surface patch acquired by a single scan by a camera-based 3D scanner, eg. face frontal scan, cloth, machine parts etc, is a genus zero surface with arbitrary number of holes, which are called multiply connected domains. The only previous existing conformal geometric method that can handle such surfaces is the Ricci Flow (RF) method in [23]. But RF is highly non-linear, hence much slower than linear methods. All linear methods, such as harmonic maps [18] and Least-Squares Conformal Maps (LSCMs) [20], are not guaranteed to generate diffeomorphisms between multiply connected domains. In this work, we introduce a novel method to handle multiply connected domains, without any restriction on the number of holes. The mappings generated by the harmonic maps and LSCMs can have flip overs. The mapping computed by RF is globally one-to-one, whereas our current method preserves the bijectiveness and improves the speed by tens of times faster.

The second major disadvantage of conformal mapping is that while it guarantees no angle distortions, it may introduce large area distortions. If a large portion of the surface shrinks to a tiny area on the parameter domain, which would make matching problematic. In order to tackle this problem, we propose to *combine multiple mappings*. For a given surface, there are infinite conformal mappings to flatten it onto the plane. For a given area on the surface, the area distortions induced by different mappings vary drastically. For each part of the surface, we can pick a specific conformal mapping, which would enlarge this part

and shrink the remaining parts. For example, Fig. 1 shows different conformal mappings for the same face surface. Three mappings enlarge the areas surrounding the mouth, the left eye and the right eye respectively. By combining matching results, we can get a better result and overcome the shrinkage problem.

The third issue we address is that conventional conformal geometric methods cannot handle surfaces with unreliable boundaries. In reality, many surfaces have multiple boundaries, where matching on the boundaries is crucial. Even more important, boundaries on 2.5D scans are sensitive to motion, noise, object pose, etc, and hence methods like harmonic maps [17,18,19] fail when the requirement for a constant boundary changes.

Hence we present three contributions:

The first contribution of this paper is to develop an algorithm which combines the results of several matches based on different conformal maps of the surfaces to be matched. Although a conformal mapping has no angle distortion, it introduces highly non-uniform area distortion. In order to avoid aliasing problems while matching, by selecting multiple conformal maps, we can ensure that all parts of the surface map to areas of equivalent size, albeit not on the same map. We can then combine all these locally optimal matching results, thus improving overall matching accuracy.

Our second contribution is to design flexible boundary conditions for matching. Some boundaries of the scanned surfaces are inconsistent among frames whereas other boundaries are more reliable. In our method, we map reliable boundaries across surfaces to be matched and allow the image of a boundary to slide on the target boundary. If the boundary is not reliable, we set it free. We also integrate feature constraints in the algorithm. Taking advantage of meaningful features is essential for any matching or registration method. In the case of large non-rigid deformations, matched features allow accurate description of the deformations. In the extreme case matching can be achieved based on feature constraints only, without any boundary constraints.

Our final contribution is to introduce a novel matching method for multiply connected domains based on canonical conformal mappings in conformal geometry. All multiply connected domains can be conformally mapped to canonical planar domains, which are annuli with concentric circular slits or rectangles with horizontal slits, as shown in Fig. 2. The resulting map does not have any singularities and is a diffeomorphism, i.e., *one-to-one* and *onto*. These maps are stable, insensitive to resolution changes, and robust to noise. Hence, the original 3D surface-matching problem simplifies to a 2D image-matching problem of the conformal geometric maps, which is better understood [24,25]. This is the first time this mapping method is applied in the computer vision field. Multiply connected domains are the most difficult cases in conformal geometry. This first application of the following theorem leads to a highly efficient method handling genus 0 surfaces with multiple holes with linear complexity.

Theorem 1 ((Canonical Mappings of Multiply Connected Domain)).
The function ϕ effects a one-to-one conformal mapping of S onto the annulus minus $n - 2$ concentric arcs.



(a) Multi. conn. domain S (b) Circular slit map ϕ (c) Conformal text. mapp. by ϕ (d) Parallel slit map $\log \phi$ (e) Conformal text. mapp. by $\log \phi$.

Fig. 2. Conformal mapping of multiply connected domains (best viewed in color). This scan of the human face surface is a multiply connected domain with 4 boundaries. It is conformally mapped to an annulus with concentric circular slits by a circular slit mapping ϕ_1 , shown in (b). The boundaries γ_0, γ_1 are mapped to outer and inner circles, γ_2, γ_3 are mapped to circular slits. The parallel slit map $\log \phi_1$ is shown in (d). (c) and (e) show that the circular and parallel slit mappings are conformal.

The rest of the paper is organized as follows: the theoretic background and linear algorithm for computing the slit mapping of multiply connected domains are introduced in Sect. 2. The algorithms for 3D surface matching and reliable boundary selection are proposed in Sect. 3. Experimental results are presented in Sect. 4, and we conclude with discussion and future work in Sect. 5.

2 Algorithm for Slit Mapping

All surfaces embedded in \mathbb{R}^3 have the induced Euclidean metric \mathbf{g} . A *conformal structure* is an atlas, such that on each local chart, the metric can be represented as $\mathbf{g} = e^{2u}(dx^2 + dy^2)$. A surface with a conformal structure is a *Riemann surface*, therefore, all surfaces in \mathbb{R}^3 are Riemann surfaces.

A *harmonic 1-form* τ on a Riemann surface can be treated as a vector field with zero circulation and divergence. A *holomorphic 1-form* ω can be treated as a pair of harmonic 1-forms, $\omega = (\tau_1, \tau_2)$, such that τ_2 can be obtained by rotating τ_1 about the normal by 90° . We say τ_2 is *conjugate* to τ_1 , and denoted as $^*\tau_1 = \tau_2$. It is convenient to use the complex representation of holomorphic 1-forms, $\omega = \tau_1 + \sqrt{-1}\tau_2$. All the holomorphic 1-forms form a group, which is isomorphic to the first cohomology group $H^1(S, \mathbb{R})$.

A topological genus zero surface S with multiple boundaries is called a *multiply connected domain*. Suppose the boundary of S is a set of loops $\partial S = \{\gamma_0, \gamma_1, \dots, \gamma_n\}$, where γ_0 is the exterior boundary. Then a set of basis of holomorphic 1-forms can be found, $\omega_1, \omega_2, \dots, \omega_n$, such that the integration of ω_i along γ_j equals to δ_{ij} . Special holomorphic 1-forms can be found, such that

$$\text{Im}\left(\int_{\gamma_i} \omega\right) = \begin{cases} 2\pi & i = 0 \\ -2\pi & i = 1 \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

then if we choose a base point p_0 on the surface, for any point p , we choose arbitrary path γ on the surface, define a complex function $\phi(p) = e^{\int_{\gamma} \omega}$, which

maps the surface to an annulus. γ_0 is mapped to the outer boundary, γ_1 to the inner boundary, and all other boundaries are mapped to the concentric circular slits. Then the (complex) logarithm of ϕ maps the surface periodically to a rectangle, with all the circular slits mapped to horizontal slits. We call ϕ a *circular slit mapping*, $\log \phi$ a *horizontal slit mapping*.

The algorithm for computing slit mapping is straight forward, first we compute a set of holomorphic 1-form basis of the surface, $\{\omega_i\}$. Then we find a holomorphic 1-form represented as the linear combination of the basis $\omega = \sum \lambda_i \omega_i$, such that Equation (1) holds. Details can be found in [26].

3 Algorithm for Surface Matching

Suppose S_1 and S_2 are two multiply connected domains, then our goal is to find an one to one map $f : S_1 \rightarrow S_2$, which is as close to an isometry as possible. Instead of matching the two surfaces in \mathbb{R}^3 directly, we find two conformal maps $\phi_k : S_k \rightarrow D_k, k = 1, 2$ which map the surface to the planar domains D_k , then we compute a planar map between the planar domains $\tilde{f} : D_1 \rightarrow D_2$. Then the matching is the composition $f = \phi_2^{-1} \circ \tilde{f} \circ \phi_1$, also described in [23].

In theory, if f is isometric and the boundary conditions are consistently set, then \tilde{f} is an identity. In our applications for face matching and cloth tracking, the mappings are close to isometry. Therefore, the planar mappings are near to the identities. This greatly simplifies the matching process.

We first explain the metric to measure the matching quality of a simple matching, then how to choose good parameterizations for optimal area distortions and how to detect the consistency between boundaries. Then we explain in detail the matching process based on a single parameterization, and finally the algorithm to fuse multiple matching results.

3.1 Optimality Criterion

Since we want to combine matching results, a metric is needed to measure the quality of the match. Let S_1, S_2 be the surfaces, $f : S_1 \rightarrow S_2$ be the match (non-rigid in general), we want our criterion to measure the distance between the match f and a rigid motion. Our proposed criterion is based on the theorem [27]:

Theorem 2. *Suppose S is a surface with conformal parameters (u, v) , the Riemannian metric is represented as $e^{2\lambda(u, v)}$, the mean curvature is $H(u, v)$. Then S is determined by λ and H uniquely up to a rigid motion in \mathbb{R}^3 .*

where $e^{2\lambda}$ measures the area distortion, and called the conformal factor. Suppose S_1, S_2 have been conformally mapped to the plane with conformal factors λ_1, λ_2 , and their mean curvatures are H_1 and H_2 , $p \in S_1$ and mapped to $f(p) \in S_2$. Therefore, we define the match energy as our optimality criterion:

$$E(f) = \int_{S_1} |\lambda_1(p) - \lambda_2(f(p))|^2 + |H_1(p) - H_2(f(p))|^2 dp. \quad (2)$$

If the matching energy is 0, then the match must be rigid motion. Thus, the smaller the energy the better the match.

The conformal factor can be approximated in the following way: Suppose $v \in S$ is a vertex, $\phi : S \rightarrow D$ is a conformal map, $\lambda(v) = \frac{\sum_{v \in f} A(f)}{\sum_{v \in f} A(\phi(f))}$, where f is a face adjacent to v , $A(f)$ is the original area of the face f , $A(\phi(f))$ is the area of the planar image of f .

The mean curvature can be approximated as $H(p)\mathbf{n}(p) = \frac{1}{2}\Delta\mathbf{r}(p)$, where $\mathbf{n}(p)$ is the normal vector at p , $\mathbf{r}(p)$ is the position vector, Δ is the Laplace-Beltrami operator, which can be approximated using the cotan formulae [19].

3.2 Choice of Conformal Parameterization

In theory, there are infinite conformal mappings for a given surface. We can only afford to compute few of them. The following is our method to choose optimal parameterizations for different regions on the surface.

First, we partition the surface to regions using the curved surface equivalent of Voronoi Diagrams. We compute the shortest distance from each vertex to all interior boundaries, choose the closest boundary, and label the vertex. We put all the vertices with the same label in the same region. The outer boundary of the scanned surface is usually noisy and inconsistent, therefore, we do not compute the Voronoi region adjacent to it.

Second, for each interior boundary γ_i , we choose a conformal parameterization, which maps the γ_i as the exterior boundary on the parameter domain. Then the region associated with γ_i has the optimal area distortion on the parameterization.

For example, for a human face with eyes and mouth open, we partition the face into 3 regions, surrounding the left, right eyes and the mouth respectively. As shown in Fig. 1, the conformal parameterizations in the 2nd, 3rd and 4th columns have optimal area distortions for the mouth region, the left eye and the right eye regions respectively.

3.3 Boundary Consistency Checking

Suppose $S_k, k = 1, 2$ are two surfaces with the same number of boundaries, $\partial S_k = \{\gamma_1^k, \gamma_2^k, \dots, \gamma_n^k\}$, where γ_i^k is the i -th boundary loop on S_k . Also we assume the correspondence between boundaries (i.e. mouth to mouth, left eye to left eye, etc.) is known. Then we choose two boundary loops and without loss of generality, assume they are $\gamma_k \in \partial S_k$. The following procedure detects whether the two boundaries are consistent or not.

1. Compute conformal mapping $\phi_k : S_k \rightarrow D_k$, which maps γ_k to the outer boundary of planar domain D_k .
2. Compute the conformal factor λ_k and mean curvature H_k for S_k .
3. Find an affine mapping $f : D_1 \rightarrow D_2$. Measure the matching energy of Equation (2) in the neighborhood of the outer boundary of D_k , and if this energy is greater than a given threshold, then the boundary loops γ_1 and γ_2 are inconsistent, otherwise, they are consistent.

According to conformal geometric theory, the image of each conformal mapping must be a rectangle; affine mapping can match different rectangles with minimal stretching energy. Our process can detect inconsistencies by trying all boundary combinations, and picking the one with least matching energy.

3.4 Simple Matching

The following steps explain the process for a matching based on a single parameterization:

1. Compute a conformal mapping $\phi_k : S_k \rightarrow D_k$.
2. Let the corresponding feature points on S_k be $F_k = \{p_1^k, p_2^k, p_3^k, \dots, p_m^k\}$. Compute a harmonic map $\tilde{f} : D_1 \rightarrow D_2$, such that $\Delta f = 0$, with the following constraints:
 - (a) Feature constraints: $\tilde{f}(\phi_1(p_i^1)) = \phi_2(p_i^2), \forall p_i^k \in F_k$.
 - (b) Boundary constraints: $\tilde{f}(\phi_1(\gamma_i^1)) = \phi_2(\gamma_i^2)$, if $\gamma_i^1 \in \partial S_1$ and $\gamma_i^2 \in \partial S_2$ are consistent.
3. The simple matching is given by $f = \phi_2^{-1} \circ \tilde{f} \circ \phi_1$.

The algorithm is applied on discrete meshes. We refer readers to [18] for details of computing harmonic maps, which is equivalent to solve a Dirichlet problem with boundary conditions. The final mapping is represented as follows: Suppose $v \in S_1$ is a vertex on the first mesh, $f(v)$ is a point $p \in S_2$ on the second surface, and p is on a face $[v_1, v_2, v_3] \in S_2$, such that $p = \mu_1 v_1 + \mu_2 v_2 + \mu_3 v_3$, where (μ_1, μ_2, μ_3) are the barycentric coordinates of p in $[v_1, v_2, v_3]$. Then we represent $f(v)$ as a pair

$$f(v) = ([v_1, v_2, v_3], (\mu_1, \mu_2, \mu_3)), \quad (3)$$

and call it a *natural representation* of the matching f .

Figure 3 illustrates a simple matching. The face surfaces in the first column need to be matched. The second column shows a conformal mapping to a rectangular domain, where γ_0^k are mapped to the top, γ_1^k are mapped to the bottom. Our algorithm found that γ_0^1 and γ_0^2 are inconsistent, γ_1^1 and γ_1^2 are consistent.

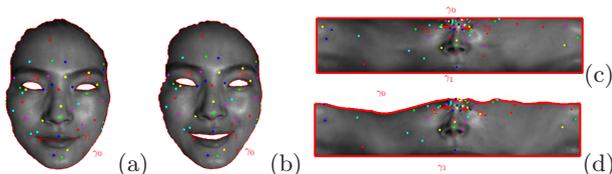


Fig. 3. Matching using boundary and feature constraints (best viewed in color). SIFT points [28] on the texture are color encoded. (a) is mapped to (b). (c) slit map for (b). (d) simple matching result. In (d), boundary γ_1 (mouth boundary) is deemed consistent using the algorithm in Sect. 3.3 and enforced to match, whereas boundary γ_0 (outer boundary) is deemed inconsistent and left free. Hence the quality of the match is due to the combination of boundary and interior feature point constraints.

Therefore, the harmonic map aligns the feature points as well as $\phi_1(\gamma_1^1)$ with $\phi_2(\gamma_1^2)$, but $\phi_1(\gamma_0^1)$ remains free.

3.5 Fusion of Matches

Suppose we have computed several simple matchings $f_i : S_1 \rightarrow S_2$, with natural representations given by Equation 3. By selecting multiple conformal maps, we can ensure that all parts of the surface map to areas of equivalent size, albeit not on the same map. We can then combine all these locally optimal matching results, using the following algorithm: Given a vertex $v \in S_1$,

1. Find a neighborhood of v on S_1 , $N_k(v) = \{[v_1, v_2, v_3] \in S_1 | d(v, v_i) < k, i = 1, 2, 3\}$, where $d(v, v_i)$ is the number of edges of the shortest path from v to v_i . In our implementation, we choose k around 4.
2. Compute the matching energy of f_i restricted on $N(v)$, denoted as $E(f_i|_{N(v)})$. If there is no image of v in f_i , we set $E(f_i|_{N(v)})$ be ∞ . The energy is represented as an integration and computed by conventional Monte Carlo method.
 - (a) Randomly generate N sample points $\{p_1, p_2, \dots, p_N\}$ in $N_k(v)$ with uniform distribution.
 - (b) Compute their images $\{q_1, q_2, \dots, q_N\}$, $q_j = f_i(p_j)$.
 - (c) Estimate the energy as

$$\frac{\sum_{j=1}^N [\lambda(p_j) - \lambda(q_j)]^2 + [H(p_j) - H(q_j)]^2}{N}.$$

3. Set the combined map as $f(v) = f_k(v)$, $k = \min_j E(f_j|_{N(v)})$.

In extensive testing of the combined matching algorithm for fine enough input meshes, the combined mappings are one-to-one and onto.

4 Experimental Results

This work handles 3D moving deformable data with complex topologies, which are very difficult to acquire. Data availability limits our experimental datasets, which however are still very challenging for existing methods. We thoroughly tested our algorithms on 50 facial scans (Face1-4, each set including more than 10 faces) with different expressions and posture, and 15 scans for deformable clothes. The face surfaces are topological three-hole annuli, the cloth surfaces are topological disks. These surfaces are representative because of their general topologies, with big distortions and very inconsistent boundaries. Such difficult experiments sufficiently support the generality and effectiveness of our method.

Tracking Deformable Cloth Surfaces. We tested our algorithm for tracking deformable cloth surfaces. The cloth surfaces are captured by the 3D scanner introduced by [29]. Each frame has about 10K vertices and 20K faces. The cloth surface is a quadrilateral (i.e., only one simple boundary). Its parallel slit mapping can be automatically computed through the preprocessing of double covering [19], which makes the surface a topological annulus (with two boundaries).

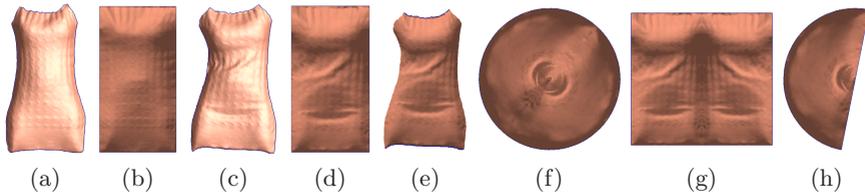


Fig. 4. Matching deforming cloth surfaces. (a) and (c) are two cloth surfaces, (b) and (d) are each rectangular parameter domain. The computing process for (c) is shown by (e-h): (e) double-covered surface, (f) circular map, (g) rectangular map, (h) half-circular map, then get the rectangular domain (d).

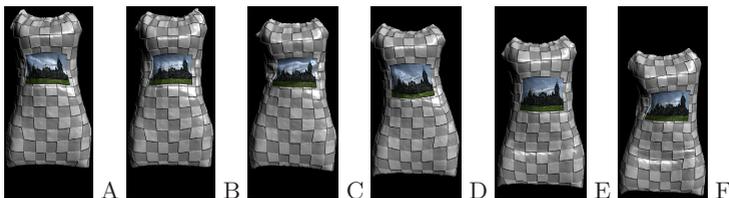


Fig. 5. Cloth tracking sequences (A-F) with consistent texture coordinates. The consistently deformable tracking can be observed from the checker board texture motion.

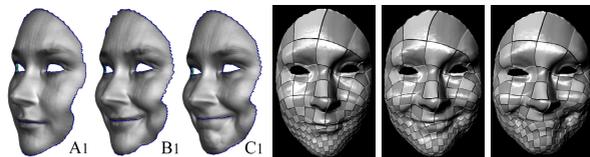


Fig. 6. Non-rigid matching and registration for Face1(A₁-C₁ with expression change and mouth closed, visualized by consistent checkerboard texture)

Fig. 4 illustrates the process of computing the rectangular mapping. It also demonstrates the simple matching between two frames. The tracking is based on the matchings frame by frame. In order to visualize the tracking results, we put checkerboard textures to the first frame, and propagated the texture parameters to the other frames, through the correspondences from tracking, see Fig. 5. The checkerboard textures are consistent across frames, without oscillating effects, or checker collapse. Thus we demonstrate that the matching between two frames is a diffeomorphism, and the tracking is stable and automatic.

Matching Facial Surfaces. Figures 6, 7 and 8 illustrate our experimental results on matching and registering two human faces with different expressions and inconsistent boundaries, acquired by the method in [18] with greyscale texture. The feature points were computed directly using SIFT [28] algorithm on the textures. The figures show how the boundaries are noisy and inconsistent.

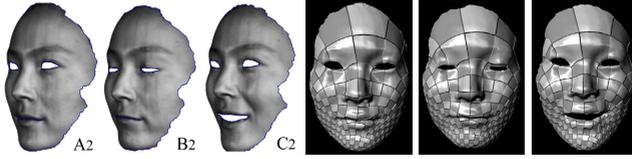


Fig. 7. Non-rigid matching and registration for Face2(A₂-C₂ with eye, mouth motion)

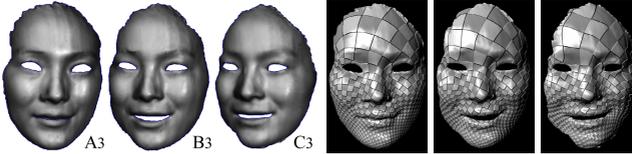


Fig. 8. Non-rigid matching and registration for Face3(A₃-C₃ with mouth motion, pose change). The rotation is about 30°. The snapshots in the left and right are taken from the original and frontal view respectively. The boundaries are significantly different.

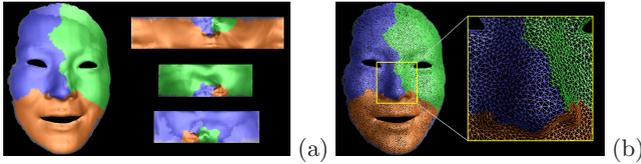


Fig. 9. Combined matching (best viewed in color): (a) the combined map, where the image of each vertex is selected from three simple matchings. The choices are color encoded; (b) the complete and zoomed in meshes.

Figure 9 shows the combined matching, where the colors indicate the choices from simple matchings. The selection metric depends on the matching energy defined in Equation 2. Intuitively, the energy measures the matching distortion. According to differential geometry, if the number is 0, then the matching must be a rigid motion. Thus the smaller the energy the better the matching. Table 1 demonstrates that the combined matching has the optimal accuracy.

Matching for Faces with Pose Change. Figure 8 shows our matching results for three faces scanned from different poses. The rotation is about 30°. Even though the exterior boundaries are very different as can be seen in the right, the matches are consistent in the mapped textures.

For the surfaces with significantly inconsistent outer boundaries, we tested two kinds of simple matching conditions on them, enforced and relaxed matching. Figure 10 illustrates both of the matching results.

We compared our current method based on holomorphic differentials (HD) with several other existing methods.

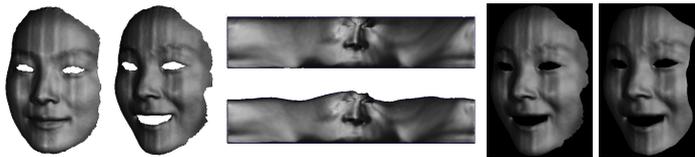
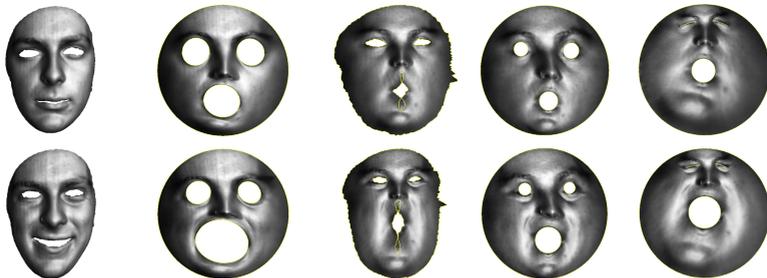


Fig. 10. Matching of significantly inconsistent boundaries of Face4



(a) Surfaces (b) Harmonic Map (c) LSCM (d) Ricci Flow (e) Holo. Diff.

Fig. 11. Comparison of geometric mappings for multiply connected domains. (b) HM makes the boundary areas much stretched. (c) LSCM generates self-intersections depending on two prescribed feature points. (d) and (e) are one-to-one and onto.

Table 1. Matching energy: R1,2,3 are the areas around the mouth, the left eye and the right eye respectively, and Map1,2,3 are the slit maps induced by each area. The number represents the matching energy of each region under each simple matching.

Face1	R1	R2	R3	Total	Face2	R1	R2	R3	Total
Comb	0.439	0.217	0.076	0.732	Comb	0.138	0.094	0.117	0.349
Map1	0.439	18.788	21.422	40.649	Map1	0.138	37.634	29.920	67.692
Map2	13.645	0.217	13.995	27.857	Map2	35.151	0.094	14.228	49.473
Map3	24.259	8.646	0.076	32.981	Map3	34.990	24.672	0.117	59.779

Matching Accuracy. We tested the matching accuracy of three methods: Iterative Closest Point (ICP) [30], Holomorphic Differentials (HD) and Ricci Flow (RF) [23]. As a baseline system, ICP is one of the most popular 3D shape matching method and has relatively good performance. RF is the only previous method to compute the matching for multiple connected domains. Our current method combines multiple mappings, which overcomes the area distortion issue, whereas RF only uses one mapping. The matching error is measured for facial surfaces and cloth surfaces by computing the relative Hausdorff average distance (RHAD) under ICP, HD and RF. We matched the first frame to others within each class and got the average match error as follows: Face1(0.028, 0.009, 0.007), Face2(0.021, 0.014, 0.010), Face3(0.089, 0.020, 0.016) and Face4(0.074, 0.015, 0.012) for (ICP, RF, HD); Cloth(0.0258, 0.0003) for (ICP, HD). For all the tested experiments, our method HD outperforms both the ICP and RF methods.

Table 2. Computational time

Name	Face1	Face2	Face3	Face4	Cloth	Time(s)	Face1	Face2	Face3	Face4	Cloth
Faces	15,000	28,800	14,387	28,148	20,998	HD	21	140	20	138	41
Verts	7,698	14,778	7,385	14,376	10,667	RF	610	4412	580	4236	N/A

Table 3. Performance comparison of geometric mapping methods

	Harmonic Map	LSCM	Ricci Flow	Holo. Diff.
Is one-to-one map	No	No	Yes	Yes
Time complexity	Linear	Linear	Non-linear	Linear
Boundary occlusion	Sensitive	Not sensitive	Sensitive	Not sensitive
Boundary constraint	Needed	Not needed	Needed	Needed
Feature constraint	Not needed	Two points needed	Not needed	Not needed
Resolution change	Not sensitive	Not sensitive	Not sensitive	Not sensitive
Topology limited	Topological disk	Topological disk	Arbitrary surface	Multiply conn. domain

Efficiency. We implemented our algorithm using generic C++ on Windows XP and used conjugate gradient optimization without using any linear package. Table 2 reports the computational time on a Laptop with CPU 2.00 GHZ, RAM 3.00 GB. RF method is non-linear, while our work is LINEAR and 30 times FASTER. Other mapping methods for computing such surfaces, like HM and LSCM, can not guarantee one-to-one map and may generate intersections (see Fig. 11). The comparison among each mapping method is illustrated in Table 3.

Automaticity and Uniqueness. The method is completely automatic. The set of holomorphic differential generators are computed automatically. The choice of holomorphic 1-form for matching is also automatic; basically, each pair of boundaries uniquely determines a unique 1-form and a unique conformal map. There is no ambiguity. The best choice of the mapping for each region on the surface is automatically determined by our algorithm as the one minimizing Equation (2). Therefore, the choice is unique and rigorous.

5 Conclusion

This work introduces novel methods for 3D surface matching based on multiple conformal mappings using holomorphic differentials (HD) for multiply connected domains. We propose a method to choose conformal parameterizations which minimize area distortions for each region and then combine the local optimal parameterization to cover the whole surface. An optimality criterion is designed to assess how consistent each boundary is across the two surfaces to be matched. That allows us to enforce matching between consistent boundaries while relaxing constraints between inconsistent ones. Compared with harmonic maps and LSCMs, HD method can generate one-to-one mapping and is linear. We tested

our matching and tracking algorithm on a large amount of 3D facial surfaces and deformable cloth surfaces with significantly inconsistent boundaries. The experiments demonstrated that our combined matching method got more accuracy and more efficiency than Ricci Flow method, which is the only previous one-to-one mapping for multiply connected domains. In the future, we will explore linear methods for surfaces with more complicated topologies, and the matching and tracking among surfaces with inconsistent topologies.

References

1. Campbell, R.J., Flynn, P.J.: A survey of free-form object representation and recognition techniques. *Computer Vision and Image Understanding* 81, 166–210 (2001)
2. Ruiz-Correa, S., Shapiro, L., Meila, M.: A new paradigm for recognizing 3d object shapes from range data. In: *ICCV*, pp. 1126–1133 (2003)
3. Huber, D., Kapuria, A., Donamukkala, R., Hebert, M.: Parts-based 3d object classification. In: *CVPR*, vol. II, pp. 82–89 (June 2004)
4. Vemuri, B., Mitiche, A., Aggarwal, J.: Curvature-based representation of objects from range data. *Image and Vision Computing* 4, 107–114 (1986)
5. Xiao, P., Barnes, N., Caetano, T., Lieby, P.: An mrf and gaussian curvature based shape representation for shape matching. In: *CVPR* (2007)
6. Sun, Y., Abidi, M.: Surface matching by 3d point's fingerprint. In: *ICCV*, vol. II, pp. 263–269 (2001)
7. Frome, A., Huber, D., Kolluri, R., Bulow, T., Malik, J.: Recognizing objects in range data using regional point descriptors. In: Pajdla, T., Matas, J(G.) (eds.) *ECCV 2004*. LNCS, vol. 3023, pp. 224–237. Springer, Heidelberg (2004)
8. Funkhouser, T., Min, P., Kazhdan, M., Chen, J., Halderman, A., Dobkin, D., Jacobs, D.: A search engine for 3d models. In: *ACM TOG*, pp. 83–105 (2003)
9. Osada, R., Funkhouser, T., Chazelle, B., Dobkin, D.: Shape distributions. In: *ACM TOG*, vol. 21, pp. 807–832 (2002)
10. Bronstein, A.M., Bronstein, M.M., Kimmel, R.: Expression-invariant representations of faces. *IEEE Trans. Image Processing* 16(1), 188–197 (2007)
11. Starck, J., Hilton, A.: Correspondence labelling for wide-timeframe free-form surface matching. In: *ICCV* (2007)
12. Lin, W.Y., Wong, K.C., Boston, N., Yu, H.H.: Fusion of summation invariants in 3d human face recognition. In: *CVPR* (2006)
13. Dalal, P., Munsell, B.C., Wang, S., Tang, J., Oliver, K., Ninomiya, H., Zhou, X., Fujita, H.: A fast 3d correspondence method for statistical shape modeling. In: *CVPR* (2007)
14. Terzopoulos, D., Witkin, A., Kass, M.: Constraints on deformable models: Recovering 3d shape and nonrigid motion. *Artificial Intelligence* 35, 91–123 (1988)
15. Huang, X., Paragios, N., Metaxas, D.: Establishing local correspondences towards compact representations of anatomical structures. In: Ellis, R.E., Peters, T.M. (eds.) *MICCAI 2003*. LNCS, vol. 2878, pp. 926–934. Springer, Heidelberg (2003)
16. Malladi, R., Sethian, J.A., Vemuri, B.C.: A fast level set based algorithm for topology-independent shape modeling. *JMIV* 6(2/3), 269–290 (1996)
17. Zhang, D., Hebert, M.: Harmonic maps and their applications in surface matching. In: *CVPR 1999*, vol. II, pp. 524–530 (1999)
18. Wang, Y., Gupta, M., Zhang, S., Wang, S., Gu, X., Samaras, D., Huang, P.: High resolution tracking of non-rigid 3d motion of densely sampled data using harmonic maps. In: *ICCV 2005*, vol. I, pp. 388–395 (2005)

19. Gu, X., Wang, Y., Chan, T.F., Thompson, P.M., Yaun, S.: Genus zero surface conformal mapping and its application to brain surface mapping. *TMI* 23(7) (2004)
20. Levy, B., Petitjean, S., Ray, N., Maillot, J.: Least squares conformal maps for automatic texture atlas generation. In: *SIGGRAPH*, pp. 362–371 (2002)
21. Sharon, E., Mumford, D.: 2d-shape analysis using conformal mapping. In: *CVPR 2004*, vol. II, pp. 350–357 (2004)
22. Wang, S., Wang, Y., Jin, M., Gu, X.D., Samaras, D.: Conformal geometry and its applications on 3d shape matching, recognition, and stitching. *PAMI* 29(7), 1209–1220 (2007)
23. Gu, X., Wang, S., Kim, J., Zeng, Y., Wang, Y., Qin, H., Samaras, D.: Ricci flow for 3d shape analysis. In: *ICCV* (2007)
24. Lowe, D.: Distinctive image features from scale-invariant keypoints. *IJCV* 60(2), 91–110 (2004)
25. Athitsos, V., Alon, J., Sclaroff, S., Kollios, G.: Boostmap: A method for efficient approximate similarity rankings. In: *CVPR 2004*, vol. II, pp. 268–275 (2004)
26. Yin, X., Dai, J., Yau, S.T., Gu, X.: Slit map: Conformal parameterization for multiply connected surfaces. In: *Geometric Modeling and Processing* (2008)
27. Gu, X., Vemuri, B.C.: Matching 3d shapes using 2d conformal representations. In: Barillot, C., Haynor, D.R., Hellier, P. (eds.) *MICCAI 2004*. LNCS, vol. 3216, pp. 771–780. Springer, Heidelberg (2004)
28. Lowe, D.: Object recognition from local scale-invariant features. In: *ICCV 1999*, pp. 1150–1157 (1999)
29. Hernández, C., Vogiatzis, G., Brostow, G.J., Stenger, B., Cipolla, R.: Non-rigid photometric stereo with colored lights. In: *ICCV*, vol. 1 (2007)
30. Besl, P.J., McKay, N.D.: A method for registration of 3-D shapes. *PAMI* 14(2), 239–256 (1992)