

Intrinsic 3D Dynamic Surface Tracking based on Dynamic Ricci Flow and Teichmüller Map

Xiaokang Yu
Dept of Comp Sci
Qingdao Univ
Qingdao, PR China
xyu@qdu.edu.cn

Na Lei
Dept of Soft and Tech
Dalian Univ of Tech
Dalian, PR China
nalei@dlut.edu.cn

Yalin Wang
Comp.Sci.& Engin
Arizona State Univ
Arizona, USA
yalin.wang@asu.edu

Xianfeng Gu
Dept of Comp Sci
Stony Brook Univ
Stony Brook, USA
gu@cs.sunysb.edu

Abstract

3D dynamic surface tracking is an important research problem and plays a vital role in many computer vision and medical imaging applications. However, it is still challenging to efficiently register surface sequences which has large deformations and strong noise. In this paper, we propose a novel automatic method for non-rigid 3D dynamic surface tracking with surface Ricci flow and Teichmüller map methods. According to quasi-conformal Teichmüller theory, the Teichmüller map minimizes the maximal dilation so that our method is able to automatically register surfaces with large deformations. Besides, the adoption of Delaunay triangulation and quadrilateral meshes makes our method applicable to low quality meshes. In our work, the 3D dynamic surfaces are acquired by a high speed 3D scanner. We first identified sparse surface features using machine learning methods in the texture space. Then we assign landmark features with different curvature settings and the Riemannian metric of the surface is computed by the dynamic Ricci flow method, such that all the curvatures are concentrated on the feature points and the surface is flat everywhere else. The registration among frames is computed by the Teichmüller mappings, which aligns the feature points with least angle distortions. We apply our new method to multiple sequences of 3D facial surfaces with large expression deformations and compare them with two other state-of-the-art tracking methods. The effectiveness of our method is demonstrated by the clearly improved accuracy and efficiency.

1. Introduction

Many imaging tracking algorithms have been proposed in the last few decades [27]. Several approaches have relied on patterns of intensity variation to identify the trajectories of objects (or major features) over time by locating its position in the video sequences. However, for some specific applications, e.g. expression transfer in motion capture [3, 22] and patient setup for critical organ protection in radiation

therapy [18], 2D (image-based) tracking systems somewhat are sensitive to image noises and cannot provide high resolution tracking solution with refined details. In contrast, 3D geometric tracking [33, 18] can capture detailed object geometric information and thus may offer better tracking performance that is desired by these applications.

Recently, the rapid development of 3D reconstruction techniques and the increasing need for automated 3D image data analysis has generated a great deal of interest in 3D image dynamic tracking research [6, 19, 34]. Often the registration among surface frames is required, and a correspondence field must be computed to register one surface frame nonlinearly onto the other. Although surface registration was well studied in computer vision and medical imaging research [23, 28], but dense and accurate registration between dynamic noisy image frames is still challenging to compute. One common way to achieve is to first map each of the 3D surfaces to canonical parameter spaces via isometry [1, 17] or conformal mappings [23]. The surface correspondence problem can be addressed by computing a flow in the parameter space of the two surfaces, which induces a correspondence field in 3D. Furthermore, specific feature points or landmark curves are enforced either by matching differences [26] or boundary conditions [20].

Quasi-conformal mapping is a generalization of conformal mapping and can be used to describe general surface mappings. By the quasiconformal Teichmüller theory, there exists a unique Teichmüller map (T-Map) among all quasi-conformal mappings between two surfaces. The obtained T-map minimizes the maximal dilation. In this paper, we show how to use T-map together with dynamic surface Ricci flow method to assist in matching of 3D face geometric surface video when there are noises, large deformations and lack of features. Fig. 1 illustrates a tracking sequence obtained using our algorithm. We can find that the template (represented by the quadrilateral mesh) is automatically transferred to other frames with different expressions while enforcing




Figure 1: 3D Facial tracking results: the quadrilateral mesh on the first frame is the template mesh. The vertices with 3 or 5 valences are singularities, the eye corners, nose hole and mouth corners are feature points, the boundary points of the quad mesh are also treated as feature points. We concentrate all the curvatures on the singularities and make the rest part of the face to be flat by using Ricci flow. We calculate the Teichmüller map between two consecutive surfaces with flat-cone metrics, all the singularities and feature points are treated as constraints. The template is automatically transformed to each frame in the dynamic 3D facial surface sequence and the feature points are matched consistently across surfaces.

feature point matching on all surfaces.

In our work, we propose two key techniques: dynamic Ricci flow and Teichmüller map. The Ricci flow method can deform the surface Riemannian metric by prescribing the target Gaussian curvature and surface registration is more efficient based on flat Riemannian metrics than general metrics with complicated curvatures. The connectivity of the mesh is updated to be Delaunay during the whole flow, which greatly improves the robustness of the method, such that it can handle meshes with noises, low quality triangulations. The other key tool is Teichmüller map, which models general diffeomorphic mappings between surfaces with large deformations. It exists and is unique, especially aligns the landmarks and minimizes distortion.

1.1. Related work

The literature for surface registration and tracking is vast. In the following, we only review some most related works. The registration/tracking methods are classified by the types of mappings they rely on, including near-isometric maps, harmonic/conformal maps and quasi-conformal maps.

Near-isometry Map Ovsjanikov et al. [17] show under mild genericity conditions, a single correspondence can be used to recover an isometry defined on entire shapes, and thus the space of all isometries can be parameterized by one correspondence between a pair of points. Huang et al. [7] treated non-rigid registration as an optimization problem and solved it by alternating between correspondence and deformation optimization. This algorithm can register objects undergo large deformations. Ovsjanikov et al. [16] proposed the functional maps method, which decompose any function on the surface by its eigen functions of

the Laplace-Beltrami operator, and match the spectrum of the function. The representation is compact, suitable for global inference and supports algebraic operations.

Harmonic/Conformal Map Wang et al. [24, 31] applied harmonic maps for the matching the 3D surfaces, because harmonic maps are diffeomorphic under mild conditions and minimizes the stretching energy of the map. The work is further generalized in [23], and conformal mapping is used to replace the harmonic maps, because conformal geometric method is capable of handling surfaces with complicated topologies. Later, Zeng et al. introduced Ricci flow method in [30], which can produce conformal mappings for general surfaces, especially it can realize the Riemann uniformization [4]. Uniformization converts 3D geometric problems to 2D ones and simplifies registration, matching. Especially, when two topological disks are isometric, in uniformization domain, they differ by a Möbius transformation, which can be easily found by only searching for a few correspondences [8, 9]. Zeng et al. [32] formulated a high-order graph matching problem to search for the optimal dense matching result by combining multiple matching criteria. Later in [33], by incorporating the uniformization idea, they defined a robust intrinsic distance function for measuring the cost of matching two points and a unified framework for intrinsic 3D surface tracking even when the two surfaces have inconsistent boundaries and are not isometrically deformed.

Quasi-conformal Maps Quasi-conformal mapping has been applied for registering 3D surfaces with large deformations in [28]. The method extracts features and estimates the Beltrami coefficient, and using curvature flow method to

solve the Beltrami equation. This method finds one quasi-conformal map, whereas our method find the optimal map, the Teichmüller map. Teichmüller map has been applied for surface registration with landmark constraints in [11], which handles topological disks. The method achieves the diffeomorphism with least angle distortion. The method is generalized to handle more complicated topological types in [14]. Both method focus on solving surface registration problem using iterative methods. In our current work, we aim at tracking a sequence dynamic surfaces. If each Teichmüller map is constructed from scratch, it will be extremely time-consuming. Instead, we utilize the information from previous Teichmüller map, and use its Beltrami coefficients as the initial guess for the next step computation. Our design improves the efficiency greatly.

1.2. Contribution

This work tackles the 3D dynamic surface tracking problem by applying dynamic discrete Ricci flow and Teichmüller mapping, which is

1. Rigorous: the method has solid theoretic foundation, which guarantees the existence and uniqueness of solutions to the dynamic Ricci flow as long as the target curvature satisfies the Gauss-Bonnet condition.
2. General: the method applies the Teichmüller mapping, which can handle large deformations, instead of those close to conformal mapping or isometric mappings, and enforce the alignment of the landmarks. The resulting mapping is guaranteed to be homeomorphic and has the least angle distortions.
3. Efficient: the method concentrates the curvature on landmarks and makes everywhere else to be flat. This solves the 3D surface tracking/registration problems on 2D planar domain, hence greatly reduces the time complexity, and improves the efficiency.

2. Theoretic Background

This section introduces the framework for dynamic unified discrete surface Ricci flow and Teichmüller map. The theory is based on the variational principle on discrete surfaces utilizing the derivative cosine law [13]. The fundamental concepts and basic schemes can be found in [12, 29].

2.1. Smooth Surface Ricci Flow

Suppose (S, \mathbf{g}) is a surface with a Riemannian metric \mathbf{g} , one can choose the isothermal coordinates, such that $\mathbf{g} = e^{2u(x,y)}(dx^2 + dy^2)$. The Gaussian curvature is given by $K(x, y) = \frac{1}{e^{2u(x,y)}} \Delta u(x, y)$, where Δ is the Laplace operator $\Delta = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}$.

Definition 2.1 (Conformal Mapping) Given two Riemannian surfaces (S, \mathbf{g}) and (T, \mathbf{h}) , a smooth mapping $\varphi : S \rightarrow T$ is called conformal, if the pull-back metric induced by φ and the original metric differ by a scalar function: $\varphi^* \mathbf{h} = e^{2u} \mathbf{g}$, where $u : S \rightarrow \mathbb{R}$ is a scalar function.

Intuitively, a conformal mapping preserves angles and maps infinitesimal circles to infinitesimal circles. In practice, Hamilton's Ricci flow is a powerful tool to compute conformal deformation by prescribing target curvatures.

Definition 2.2 (Hamilton's Ricci Flow) Given a closed Riemannian surface (S, \mathbf{g}) , the normalized Ricci flow is defined as

$$\frac{dg_{ij}}{dt} = 2\left(\frac{2\pi\chi(S)}{A(0)} - K\right)g_{ij}, \quad (1)$$

where $\chi(S)$ is the Euler characteristic number of the surface, $A(0)$ is the initial total area.

Hamilton and Chow proved that the surface Ricci flow converges to the constant curvature metric.

2.2. Discrete Surface Ricci Flow

Practically, surfaces are represented as triangular meshes. A discrete metric on a mesh is the edge length function, denoted as $l : E \rightarrow \mathbb{R}^+$, which satisfies the triangle inequality. The discrete Gauss curvature is the angle deficit, defined on vertices, $K : V \rightarrow \mathbb{R}$,

$$K(v) = \begin{cases} 2\pi - \sum_{jk} \theta_i^{jk}, & v \notin \partial M \\ \pi - \sum_{jk} \theta_i^{jk}, & v \in \partial M \end{cases}, \quad (2)$$

where θ_i^{jk} is the corner angle at v_i in the face $[v_i, v_j, v_k]$, and ∂M represents the boundary of the mesh.

The discrete Gaussian curvature are determined by the discrete Riemannian metric via the cosine law,

$$l_i^2 = l_j^2 + l_k^2 - 2l_j l_k \cos \theta_i. \quad (3)$$

The Gauss-Bonnet theorem still holds in the discrete case. The total curvature equals to the product of 2π and the Euler characteristic number χ ,

$$\sum_{v \notin \partial \Sigma} K(v) + \sum_{v \in \partial \Sigma} K(v) = 2\pi\chi(\Sigma). \quad (4)$$

The cotangent edge weight plays an important role. Given an interior edge $[v_i, v_j]$ adjacent to two faces $[v_i, v_j, v_k]$ and $[v_j, v_i, v_l]$, the cotangent weight is defined as

$$w_{ij} = \cot \theta_k^{ij} + \cot \theta_l^{ji}, \quad (5)$$

if the edge is on the boundary, adjacent to the face $[v_i, v_j, v_k]$, then the cotangent weight is

$$w_{ij} = \cot \theta_k^{ij}. \quad (6)$$

A triangulation of the mesh is called *Delaunay*, if all cotangent edge weights are non-negative.

Given a triangular mesh M , the discrete conformal factor is a function defined on each vertex $\mathbf{u} : V \rightarrow \mathbb{R}$, the length of an edge $[v_i, v_j]$ is given by

$$l_{ij} = \exp(u_i) \beta_{ij} \exp(u_j), \quad (7)$$

where β_{ij} is the initial edge length.




Figure 2: Conformal mapping (top row) preserves infinitesimal circles; quasi-conformal mapping (bottom row) maps infinitesimal ellipses to infinitesimal circles.

Definition 2.3 (Discrete Surface Ricci Flow) The discrete surface Ricci flow is defined as

$$\frac{du_i(t)}{dt} = \bar{K}_i - K_i(t), \quad (8)$$

where \bar{K}_i is the target curvature at the vertex v_i , and the discrete metric is given by Eqn.7. During the flow the triangulation is updated to be Delaunay.

The existence to the Ricci flow is recently proven,

Theorem 2.4 (Discrete Uniformization[5]) Given a target curvature \bar{K} satisfying the Gauss-Bonnet condition in Eqn.4, and for each vertex $\bar{K}_i \in (-\infty, 2\pi)$, then there exists a solution to the Ricci flow Eqn.8. The solution is unique up to a constant.

Furthermore, the discrete Ricci flow is the negative gradient flow of the discrete Ricci energy:

$$E_{\Sigma}(\mathbf{u}) = \int^{\mathbf{u}} \sum_{i=1}^n (\bar{K}_i - K_i) du_i. \quad (9)$$

The gradient of the Ricci energy is $(\bar{K}_i - K_i)^T$, the Hessian matrix consists of cotangent edge weights Eqn. 5,

$$\frac{\partial K_i}{\partial u_j} = \frac{\partial K_j}{\partial u_i} = w_{ij}, \quad (10)$$

and the diagonal elements are

$$\frac{\partial K_i}{\partial u_i} = - \sum_j w_{ij}. \quad (11)$$

The first row in Fig. 2 shows a conformal mapping from a human face to a disk, where infinitesimal circle shapes are preserved on both surfaces.

2.3. Teichmüller Mapping

Quasi-conformal mapping is a generalization of conformal mapping, which preserves orientation between surfaces with bounded conformality distortions. Suppose $\varphi : \mathbb{C} \rightarrow \mathbb{C}$ is a diffeomorphism, then the Beltrami coefficient μ of the φ is given by

$$\frac{\partial \varphi}{\partial \bar{z}} = \mu(z) \frac{\partial \varphi}{\partial z}. \quad (12)$$

The quasi-conformal mapping maps infinitesimal circles to infinitesimal ellipses. The ratio between the longer axis and the shorter axis of the infinitesimal ellipse is denoted by $(1 + |\mu|)/(1 - |\mu|)$, which describes the eccentricity of the infinitesimal ellipse. The orientation of infinitesimal ellipse is given by the angle between the longer axis and the real axis, i.e., $\theta = 1/2 \arg \mu$. Hence, μ measures the local conformality distortion of φ , and if φ is conformal then $\|\mu\|_{\infty}$ is 0, φ is diffeomorphic then $\|\mu\|_{\infty} < 1$. The maximal dilation of φ is given by $\mathcal{K}(\varphi) = (1 + \|\mu\|_{\infty})/(1 - \|\mu\|_{\infty})$. The second row of Fig. 2 shows a quasi-conformal mapping between a face surface and a disk. Note the circles on the disk are mapped to ellipses on the face surface.

Definition 2.5 (Teichmüller Mapping (T-Map)) Suppose $\varphi : S_1 \rightarrow S_2$ is a quasi-conformal mapping. φ is a Teichmüller mapping associated with quadratic differential $q = f dz^2$, where $f : S_1 \rightarrow \mathbb{C}$ is a holomorphic function, if its associated Beltrami coefficient is of the form: $\mu(\varphi) = k \bar{f}/|f|$ for some constant $0 \leq k \leq 1$.

Given two homeomorphic surfaces, the Teichmüller map exists and is unique and it minimizes the maximal dilation $\mathcal{K}(\varphi)$. Fig. 3 illustrates a Teichmüller mapping between two human facial surfaces with markers. The Teichmüller mapping maps infinitesimal circles on the male face to infinitesimal ellipses on the female face, such that all ellipses are with the same eccentricity. In our work, the facial surfaces are punctured at the feature points, Ricci flow is applied to compute surface conformal flattening with special curvature settings on the feature points. Later, we pursue Teichmüller maps between the punctured surfaces which is the diffeomorphism with the least conformality distortion.

3. Algorithms

Fig. 4 illustrates the algorithm pipeline of our method. Frame (a) is the input surface; (b) shows the identified facial features and generated cut graph; (c) shows the flat-cone metric computed by surface Ricci flow method; (d) is the T-map between it and the template surface, and (e) illustrates the final dynamic facial surface mapping result.

Input Data They are raw geometric data obtained from a 3D scanning system [25], without any digital geometric processing, such as denoising, smoothing etc. The input surfaces are with gray level or color texture image.

Feature Extraction Using facial texture images as input, we applied machine learning methods similar to [21][35]




Figure 4: Algorithm pipeline.




Figure 3: Teichmüller mapping between two human facial surfaces with markers (white dots), all ellipses are of the same eccentricity.

to extract major facial features points (part of the feature points are shown in Frame (b)). They include eye corners, mouth corners, nose hole boundaries, apexes of the cheekbone and so on. The contours of the eye regions, the mouth region and the whole face region are obtained by connecting the sequential feature points along them. Further the eye regions and the mouth region are cut off along their contours and the region exterior to the whole face region is removed by cutting along the face contour. Then we compute a spanning tree of all feature points, and find the shortest paths from the tree to all the boundaries. The union of the tree and the shortest paths is called the *cut graph* (the blue lines in Fig. 4 (b)). By slicing the surface along the cut graph, the facial surface becomes a topological disk.

Metric Flattening By Ricci Flow Algorithm 1 sketches the outline of metric flattening process with dynamic Ricci flow. Surface Ricci flow provides a powerful tool to generate canonical spaces between different surface frames, even with large deformations. We adopt Ricci flow to map all surfaces to a canonical space where each feature point has a designed Gaussian curvature while Gaussian curvature are zero for other points. Frame (c) shows the flat metric with

Algorithm 1: Dynamic Discrete Surface Ricci Flow.

Input: The input mesh M and the target curvature \bar{K} , threshold ϵ

Output: The edge length which realizes the target curvature

- 1 Compute the initial edge lengths $\{\beta_{ij}\}$;
 - 2 Initialize the conformal factor to be zeros;
 - 3 **while** *true* **do**
 - 4 Compute the edge lengths using Eqn.7;
 - 5 Update the triangulation to be Delaunay by edge swapping;
 - 6 Compute the corner angles using Eqn.3;
 - 7 Compute the cotangent edge weights using Eqn. 5 and Eqn. 6;
 - 8 Compute the vertex curvature using Eqn.2;
 - 9 **if** $\forall |\bar{K}_i - K_i(\mathbf{h})| < \epsilon$ **then**
 - 10 | Break;
 - 11 **end**
 - 12 Compute the gradient of the Ricci flow using Eqn.8;
 - 13 Compute the Hessian of the Ricci energy using Eqn.10 and Eqn.11;
 - 14 Solve the linear system $Hess(\mathbf{u})\delta\mathbf{u} = \nabla E(\mathbf{u})$
 - 15 $\mathbf{u} \leftarrow \mathbf{u} + \delta\mathbf{u}$;
 - 16 **end**
 - 17 **return** the edge length $\{l_{ij}\}$
-

cone singularities using dynamic Ricci flow method where the curvatures are concentrated on the feature points. In Figure 1 frame (a), the quadrilateral mesh is treated as the template mesh. Each feature point corresponds to a vertex in the template, for each interior feature point, if the valence is 3, then the target curvature is $\pi/2$, if the value is 5, then the target curvature is $-\pi/2$; for feature points on the boundaries, the target curvatures are zeros. The target

Algorithm 2: Teichmüller Mapping.

Input: Two triangular meshes S and T , corresponding interior landmark constraints $\{p_i\}_{i=1}^n \subset S$, $\{q_i\}_{i=1}^n \subset T$, and a threshold δ of Beltrami coefficient difference

Output: A Teichmüller map $\varphi : S \rightarrow T$ which matches landmarks

- 1 Compute initial edge length and edge weight in S , set $n \leftarrow 0$
 - 2 Initially set Beltrami coefficient μ_0 to be zero everywhere (or copy from previous frames),
 - 3 Set $n \leftarrow n + 1$, and compute a constrained harmonic map $\varphi_n : (S, \mu_n) \rightarrow T$, such that φ_n aligns the corresponding landmarks given above;
 - 4 Compute the Beltrami coefficient ν_n define on each face using Eqn. 12
 - 5 Project the Beltrami coefficient ν_n to obtain μ_{n+1} using Eqn. 13
 - 6 Repeat step 3 through 5, until $\|\mu_n - \mu_{n+1}\| < \delta$ for each face. Then the final Teichmüller map is given by φ_n .
-

curvatures for all the other vertices are zeros by default. After obtaining the flat metric, we slice the surface along the cut graph, and flatten the surface on the plane. The regular grids on the plane are pulled back to the surface, the red curves are the horizontal lines, the blue curves are the vertical lines. Since the cut graph is mapped to the domain boundary on the parametric space, we can see on Frame (c) that these curves are broken across the cut graph. On the parameter domain, Dulaunay triangulation is applied to refine the mesh quality.

Teichmüller Map The template in Figure 1 frame (a) is a quadrilateral mesh. We treat each face as a unit planar square, then the template has a flat metric with cone singularities. We use T to represent the template quad-mesh, \mathbf{d} the flat-cone metric. For each surface S_i in the sequence, with the flat-cone metric \mathbf{g}_i , we can compute the Teichmüller map $\varphi_i : (S_i, \mathbf{g}_i) \rightarrow (T, \mathbf{d})$. Algorithm 2 briefs the procedure and a detailed description is next.

Generalized Harmonic Map The Teichmüller map is obtained via the generalized harmonic map with a given Beltrami coefficient μ . In our implementation, the Beltrami coefficient is defined on faces $\mu : F \rightarrow \mathbb{C}$. We locally isometrically embed each neighborhood of the surface (S, \mathbf{g}) on the complex plane, and obtain complex local coordinates z . For each face $[v_i, v_j, v_k]$, we define a local linear map, $w = z + \mu\bar{z}$, then compute the cotangent edge weight w_{ij} according to Equation 5. The mapping φ is harmonic, if and only if

$$\varphi(v_i) = \frac{\sum_j w_{ij} \varphi(v_j)}{\sum_j w_{ij}}.$$

By isometrically embedding (T, \mathbf{d}) on the plane locally, we can choose any local coordinates to represent the mapping φ , and the above equation holds independent of the choice of the local coordinates. The above equation means, the image of each vertex is the weighted center of the images of its neighbors, which is called the mean value property of the harmonic map [20]. For a boundary vertex v_i , the image $\varphi(v_i)$ is the projection of the weighted center to the boundary. In the computation, each time we move the image of an interior vertex to its weighted center, and the image of a boundary vertex to the projection of its weighted center to the boundary until the algorithm converges.

Beltrami Coefficient Projection In order to obtain the Teichmüller map, at each iteration, we update the Beltrami coefficient, such that its norm is constant everywhere. In the n -th iteration, After calculating a generalized harmonic map $\varphi_n : (S, \mathbf{g}, \mu_n) \rightarrow (T, \mathbf{d})$, we compute the Beltrami coefficient of φ_n using the formula of Eqn.12, denoted as ν_n . Then we set μ_{n+1} as

$$\mu_{n+1} = c_n \frac{\nu_n}{|\nu_n|}, \quad c_n = \frac{\sum_i |\nu_n(f_i)| A_i}{\sum_i A_i}, \quad (13)$$

where A_i is the area of the i -th face, c_n is the mean value of the norm of the Beltrami coefficient on all faces. By iteratively perform Beltrami coefficient projection, and compute generalized harmonic map, we obtain a sequence of ν_n (Beltrami coefficient of φ_n), μ_{n+1} (projection of ν_n), $\varphi_{n+1} : (S, \mu_{n+1}, \mathbf{g}) \rightarrow (S, t)$ (generalized harmonic map), eventually the algorithm converges to the unique T-map.

Quadrilateral Mesh Generation The Teichmüller maps between the template surface and each frame are obtained by the above methods. The template surface is a canonical quadrilateral mesh, the mappings push forward the quad-mesh to all frames, this induces a consistent quadrilateral tessellation among all the frames.

4. Experimental Results

We implemented our method using generic C++, on a Windows platform with 3.5GHz Intel Core i5 processor and 16GB memory with NVIDIA GTX960 graphics card with 1024 CUDA cores. The Ricci flow and Teichmüller map methods are implemented with cuSparse linear algebra library [15]. The 3D human facial surfaces were scanned from a high speed and high resolution, phase shifting scanner, as described in [25]. The acquisition speed is 30 frames per second, each frame consists of about 100k sampling points with both gray-level and color texture images.

4.1. Dynamic Human Facial Expression Tracking

Figure 5 illustrates the registration result between two frames with a large facial deformation. The quadrilateral mesh on the smiling face is the template mesh, where the vertices with valences 3 or 5 are treated as registration constraints. There are a total of 28 singularities, including




Figure 5: Each sampling point and its corresponding point are connected by a straight line.

features points such as the eye corners, the nose holes, the mouth corners and the boundary of the template region obtained using machine learning method [35]. We first use Ricci flow to concentrate all the curvatures on these singularities and make the rest part of the surface to be flat. Then we compute the Teichmüller map between two surfaces with flat-cone Riemannian metrics, and the feature constraints. The mapping minimizes the angle distortion and produces a diffeomorphism. We randomly select a set of sampling points on the source surface (front), namely the smiling face, and display their images on the target surface (rear), the face with curious expression. We connect each pair of corresponding points with straight lines. We can find that these feature points are well mapped. Note that the two facial surfaces differ by a dramatic expression change. It demonstrates the accuracy of our surface matching method.

Figure 6 further demonstrates the result for tracking a sequence of 3D faces with large expression deformations (a tracking demonstration video is provided in our supplementary material). From the results, we can observe that the template connectivity (visualized by the quadrilateral mesh) is consistently attached to the facial surface and all the interior vertices of the template are tightly glued to right points on the target surfaces without any sliding effect. In contrast, the sliding effect is obvious in the result of a conventional method [33] (as in the supplementary material). By a simple visual checking, we can find that our method achieves better tracking results.

4.2. Robustness and Correctness Validation

We tested our method for 4 actors, more than 20 sequences of 3D facial surfaces with large expression deformations. Fig. 6 shows on tracking result of a male facial surface with different expressions. The expression deformation is relatively large, including mouth and eye open-



Figure 6: 3D surface tracking for facial surfaces with large expression changes.




Figure 7: The norms of the Beltrami coefficients.

ing/closing, non-symmetric facial twisting. The boundaries of the surfaces are noisy, with partial occlusions, incomplete geometric data. The geometry of the eye lids are bumpy, because the scanning speed is slower comparing to the eye blinking. Note that all the computations are carried out on the raw data, directly obtained by the 3D scanning system, without any digital geometry processing, such as mesh simplification, denoising, or smoothing. Our dynamic Ricci flow and Teichmüller mapping methods are robust enough to process the raw data, produce the desired flat metric with cone singularities and generate smooth tracking sequences.

To demonstrate the correctness of our facial tracking results, we compute the Beltrami coefficient μ_n for each map-

ping φ_n between the first frame S_0 and every other frame S_n in the sequence and compute the L^∞ norm of the Beltrami coefficient. We draw the norms of the Beltrami coefficients in Fig. 7, the horizontal axis is the index of the frame number n , the vertical axis is $\|\mu_n\|_\infty$. It is obvious that all the L^∞ norms are less than 1. It means that all the mappings are homoeomorphisms. Besides, all the norms are close to 0. It means that the conformality distortions are small. Furthermore, from the curve we can observe that the mapping between faces with similar expressions has less distortion than the mapping between faces with large expression deformation. The experimental results match our intuition.

4.3. Tracking Accuracy and Efficiency Comparison

In order to quantify the registration accuracy, we apply the following metric. The n -th facial surface is denoted as S_n , the mapping between S_0 and S_n is $\varphi_n : S_0 \rightarrow S_n$. We manually select a set of markers $\{p_1, p_2, \dots, p_k\} \subset S_0$, and their corresponding points on S_n as $\{q_1^n, q_2^n, \dots, q_k^n\} \subset S_2$. Then we measure the distance between the automatic computational result $\varphi_n(p_i)$ and the manually prescribed the image q_i^n ,

$$d_n = \frac{1}{L_n} \sqrt{\frac{\sum_{i=1}^k d(q_i^n, \varphi_n(p_i))^2}{k}}$$

where L_n is the length of the diagonal of the bounding box of the surface S_n and $k = 100$.

We study two other methods for comparison, optimal flow method [10, 2] and Markov Random Field (MRF) method [33] and compare our cost functions with theirs.

We implemented the optical flow algorithm in [10] based on [2]. We project the left part of the face to a 640×480 perspective view selected to maximize visibility and apply the optical flow algorithm to establish correspondences between two frames. For the template points that belong to the projected part, we compute the cost function defined in [33] and choose the correspondence with lowest cost as the matching result. We linearly interpolate the correspondence of other points within the template that are visible. For MRF approach, they consider the set of all possible 3D surface matchings defined by specifying triplets of correspondences in the uniformization domain, then they introduce a new matching cost between two 3D surfaces. The lowest feature differences across this set of matchings that cause two points to be corresponded, become the matching cost of that particular correspondence. The matching cost is computed in the uniformization domain. This matching cost is then combined with regularization terms that enforce spatial and temporal motion consistencies, into a maximum a posteriori (MAP) problem which they approximate using a Markov Random Field (MRF).

We compared the average distance error of facial sequence tracking, using our algorithm, optical flow and MRF, respectively. Fig. 8 shows the comparison results, the




Figure 8: Comparison with the MRF method in [33] and optical flow method in [10].

horizontal axis is the frame index n and the vertical axis is the normalized distance error. It is obvious that the optical flow method in [10] produces the largest error, and when the deformation between two frames is large, the MRF in [33] degrades more significantly than our method.

Furthermore, from a practical point of view, MRF method may get stuck at the local optimas. It makes the coding and debugging more challenging. In contrast, our algorithms can reach the unique global optimum easily, therefore the implementation and the debugging are much easier. We further compare the time complexity of our method and that of MRF method [33]. Both methods are speeded up by using GPUs and tested with same surfaces sequence. It takes more than 30s to register two surfaces with 100k vertices by MRF method, takes about 5s by our algorithm. It is obvious that our method is much faster.

In [33], Zeng *et al.* combine MRF optimization and conformal geometry so they assume the expression deformation is near-isometric or conformal. However, by examining the curves in Figure 7, we see that the mapping is most likely to be quasi-conformal. Especially for faces with large expression changes, the mappings are far from conformal. Hence the assumption in [33] seems incorrect. Instead, we use Teichmüller mapping theory, which is able to handle general mappings, especially with large deformations. Although further investigation is warranted, it may partially explain why we achieved more accurate results in Fig. 8.

5. Conclusion and Future Work

This work introduces a novel method for 3D dynamic surface tracking based on discrete surface Ricci flow and Teichmüller mapping. The current method depends on machine learning based feature extraction, and focuses on human facial surfaces. In future, we will also study tracking general dynamic surfaces with larger deformations.

Acknowledgement: This work was partially supported by grants from: Qingdao Application Foundation 14-2-4-40-jch, NSF grants: IIS-1421165, DMS-1413417, NSF DMS-1418255, AFOSR FA9550-14-1-0193.

References

- [1] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Three-dimensional face recognition. *Int. J. Comput. Vision*, 64(1):5–30, Aug. 2005.
- [2] A. Bruhn, J. Weickert, and C. Schnörr. Lucas/kanade meets horn/schunck: Combining local and global optic flow methods. *International Journal of Computer Vision*, 61(3):211–231, 2005.
- [3] R. J. Campbell and P. J. Flynn. A survey of free-form object representation and recognition techniques. *Computer Vision and Image Understanding*, 81(2):166–210, 2001.
- [4] H. M. Farkas and I. Kra. *Riemann Surfaces*. Springer, 2004.
- [5] X. Gu, F. Luo, J. Sun, and T. Wu. A discrete uniformization theorem for polyhedral surfaces. *Journal of Differential Geometry*, 2017 in press.
- [6] C. Hernández, G. Vogiatzis, G. J. Brostow, B. Stenger, and R. Cipolla. Non-rigid photometric stereo with colored lights. In *Computer Vision, 2007. ICCV 2007. IEEE 11th International Conference on*, pages 1–8. IEEE, 2007.
- [7] Q.-X. Huang, B. Adams, M. Wicke, and L. J. Guibas. Non-rigid registration under isometric deformations. In *Proceedings of the Symposium on Geometry Processing*, SGP '08, pages 1449–1457, Aire-la-Ville, Switzerland, Switzerland, 2008. Eurographics Association.
- [8] Y. Lipman and I. Daubechies. Surface comparison with mass transportation. Technical report, Princeton University, 2010.
- [9] Y. Lipman and T. Funkhouser. Möbius voting for surface correspondence. *ACM Transactions on Graphics*, 28(3):1–12, 2009.
- [10] C. Liu. *Beyond pixels: exploring new representations and applications for motion analysis*. PhD thesis, Citeseer, 2009.
- [11] L. M. Lui, K. C. Lam, S.-T. Yau, and X. Gu. Teichmüller mapping (t-map) and its applications to landmark matching registration. *SIAM J. Imaging Sciences*, 7(1):391–426, 2014.
- [12] F. Luo. Rigidity of polyhedral surfaces. *Journal of Differential Geometry*, 96(1):241–302, 2014.
- [13] F. Luo, X. Gu, and J. Dai. *Variational Principles for Discrete Surfaces*. Advanced Lectures in Mathematics. High Education Press and International Press, 2007.
- [14] T. C. Ng, X. Gu, and L. M. Lui. Teichmüller extremal map of multiply-connected domains using beltrami holomorphic flow. *Journal of Scientific Computing*, 60(2):249–275, 2014.
- [15] NVIDIA. cusparse library. <http://docs.nvidia.com/cuda/cusparse/>, 2016.
- [16] M. Ovsjanikov, M. Ben-Chen, J. Solomon, A. Butscher, and L. Guibas. Functional maps: A flexible representation of maps between shapes. *ACM Trans. Graph.*, 31(4):30:1–30:11, July 2012.
- [17] M. Ovsjanikov, Q. Mrigot, F. Mmoli, and L. Guibas. One point isometric matching with the heat kernel. *Computer Graphics Forum*, 29(5):1555–1564, 2010.
- [18] P. J. Schoffel, W. Harms, G. Sroka-Perez, W. Schlegel, and C. P. Karger. Accuracy of a commercial optical 3D surface imaging system for realignment of patients for radiotherapy of the thorax. *Phys Med Biol*, 52(13):3949–3963, Jul 2007.
- [19] A. Shaji, A. Varol, L. Torresani, and P. Fua. Simultaneous point matching and 3d deformable surface reconstruction. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 1221–1228. IEEE, 2010.
- [20] R. Shi, W. Zeng, Z. Su, J. Jiang, H. Damasio, Z. Lu, Y. Wang, S.-T. Yau, and X. Gu. Hyperbolic harmonic mapping for surface registration. *IEEE Trans Pattern Anal Mach Intell*, 2016.
- [21] Y. Sun, X. Wang, and X. Tang. Deep convolutional network cascade for facial point detection. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2013.
- [22] D. Vlasic, M. Brand, H. Pfister, and J. Popović. Face transfer with multilinear models. In *ACM Transactions on Graphics (TOG)*, volume 24, pages 426–433. ACM, 2005.
- [23] S. Wang, Y. Wang, M. Jin, X. D. Gu, and D. Samaras. Conformal geometry and its applications on 3d shape matching, recognition, and stitching. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(7):1209–1220, July 2007.
- [24] 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 *International Conference on Computer Vision*, 2005.
- [25] Y. Wang, M. Gupta, S. Zhang, S. Wang, X. Gu, D. Samaras, and P. Huang. High resolution tracking of non-rigid motion of densely sampled 3d data using harmonic maps. *International Journal of Computer Vision*, 76(3):283–300, 2008.
- [26] Y. Wang, L. M. Lui, T. F. Chan, and P. M. Thompson. Optimization of brain conformal mapping with landmarks. *Med Image Comput Assist Interv*, 8(Pt 2):675–683, 2005.
- [27] A. Yilmaz, O. Javed, and M. Shah. Object tracking: A survey. *ACM Comput. Surv.*, 38(4), Dec. 2006.
- [28] W. Zeng and X. Gu. Registration for 3d surfaces with large deformations using quasi-conformal curvature flow. In *CVPR*, 2011.
- [29] W. Zeng and X. Gu. *Ricci Flow for Shape Analysis and Surface Registration*. SpringerBriefs in Mathematics. Springer New York, 2013.
- [30] W. Zeng, D. Samaras, and X. D. Gu. Ricci flow for 3d shape analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32:662–677, 2010.
- [31] W. Zeng, Y. Zeng, Y. Wang, X. Yin, X. D. Gu, and D. Samaras. 3D non-rigid surface matching and registration based on holomorphic differentials. In *European Conference on Computer Vision*, 2008.
- [32] Y. Zeng, C. Wang, Y. Wang, X. Gu, D. Samaras, and N. Paragios. Dense non-rigid surface registration using high-order graph matching. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2010.
- [33] Y. Zeng, C. Wang, Y. Wang, X. Gu, D. Samaras, and N. Paragios. Intrinsic dense 3d surface tracking. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1225–1232, 2011.
- [34] L. Zhang, N. Snavely, B. Curless, and S. M. Seitz. Spacetime faces: High-resolution capture for modeling and animation. In *Data-Driven 3D Facial Animation*, pages 248–276. Springer, 2008.
- [35] Z. Zhang, P. Luo, C. C. Loy, and X. Tang. Facial landmark detection by deep multi-task learning. In *European Conference on Computer Vision*, 2014.