Conformal Surface Alignment with Optimal Möbius Search

Huu Le, Tat-Jun Chin and David Suter
School of Computer Science, The University of Adelaide
{huu.le, tat-jun.chin, david.suter}@adelaide.edu.au

Abstract

Deformations of surfaces with the same intrinsic shape can often be described accurately by a conformal model. A major focus of computational conformal geometry is the estimation of the conformal mapping that aligns a given pair of object surfaces. The uniformization theorem enables this task to be acccomplished in a canonical 2D domain, wherein the surfaces can be aligned using a Möbius transformation. Current algorithms for estimating Möbius transformations, however, often cannot provide satisfactory alignment or are computationally too costly. This paper introduces a novel globally optimal algorithm for estimating Möbius transformations to align surfaces that are topological discs. Unlike previous methods, the proposed algorithm deterministically calculates the best transformation, without requiring good initializations. Further, our algorithm is also much faster than previous techniques in practice. We demonstrate the efficacy of our algorithm on data commonly used in computational conformal geometry.

1. Introduction

Given two 3D shapes, the problem of shape correspondence is to find a meaningful relation (or mapping) between the elements of the shapes [35]. Solving such a problem is fundamental to many vision and graphics applications, such as object recognition, 3D shape retrieval, shape morphing and attribute transfer - to name a few. A plethora of variants also exist for the general problem, where the variants differ based on the type of input data (point cloud, mesh, etc.), type of alignment function, partial or full overlap, whether pre-identified landmarks are available, etc.

Within the broad literature on shape correspondence, our work belongs to the class of conformal geometric methods [38]. A conformal mapping preserves angles locally, and is thus insensitive to surface deformations [36]. It has been observed that object instances with the same intrinsic shape (e.g., faces with different expressions [36], deformable 2D shapes [12], brains between different individuals [15]) can be aligned well conformally.

The uniformization theorem [7] states that all surfaces that are topological spheres or discs can be conformally embedded to a canonical 2D domain, e.g., a unit sphere, a hyperbolic disc. In Fig. 1(a), these embeddings are represented by $\Phi_{\tilde{\mathcal{M}}}$ and $\Phi_{\tilde{\mathcal{B}}}$ for respectively two surfaces $\tilde{\mathcal{M}}$ and $\tilde{\mathcal{B}}$. The embedded surfaces can be aligned conformally by a Möbius transformation f. The direct mapping between the surfaces can then be composed as $\Phi_{\tilde{\mathcal{B}}}^{-1} \circ f \circ \Phi_{\tilde{\mathcal{M}}}$. Note that in practice, discrete analogues of the mappings are used.

Conformal shape correspondence thus amounts to finding $\Phi_{\tilde{\mathcal{M}}}$, $\Phi_{\tilde{\mathcal{B}}}$ and f. Many authors first calculate the embeddings $\Phi_{\tilde{\mathcal{M}}}$ and $\Phi_{\tilde{\mathcal{B}}}$ in a "flattening" step, before estimating f, e.g., [17, 15, 28, 9, 27, 24]. Various methods have been devised for flattening [6, 22, 16, 17, 33]. In this paper, we focus on the second step, i.e., Möbius alignment.

A class of existing methods for Möbius alignment are correspondence-based. Correspondences between the surfaces can be obtained by identifying landmarks (meaningful locations such as corners of eyes or tips of noses) or matching salient keypoints [17, 15, 36, 9, 29]. Given a sufficient number of correspondences, a Möbius transformation can be directly estimated [28]. The effectiveness of correspondence-based methods hinges on the veracity of the correspondences and their coverage of the surfaces.

Correspondence-free techniques as we consider in this paper must directly estimate the Möbius transformation, and implicitly the surface correspondence. Current methods include randomized search [28], iterative closest points (ICP) [8], gradient-based local optimization [24], brute force enumeration [27], and graph matching [40]. While not affected by inaccurate or insufficient pre-identified correspondences, these methods suffer, however, from other weaknesses; namely, no guarantee of success [28], dependence on good initialisation [8, 24], and very high computational cost [27, 40].

Contributions We propose a novel *globally optimal* algorithm for correspondence-free Möbius alignment. We focus on surfaces that are topological discs, i.e., f acts on the hyperbolic disc. Based on branch-and-bound (BnB) [20], our algorithm guarantees global optimality, thus obviating the need for good intializations. Further, our method is also

much more efficient than previous techniques. Note that there have been previous usages of BnB for point set alignment, but these are mostly for the rigid case [10, 31, 19].

Conformally aligning sufaces that are topological discs has many practical applications, e.g., facial expression analysis [36], shape similarity analysis [27], and brain mapping [23]. Our work thus presents a useful tool to the very important area of computational conformal geometry [18].

1.1. Correspondence-free Möbius alignment

On the unit sphere, a Möbius transformation has 6DOF and is uniquely defined by three distinct correspondences. Lipman and Funkhouser [28] randomly sample triplets from both surfaces to create three correspondences, estimate a Möbius transformation from each sample, and let the samples vote for the best correspondences. Although it usually produces good results, the algorithm is random in nature and does not guarantee optimality or "goodness". Another disadvantage is its high cost for large input shapes.

In the context of Möbius alignment, ICP [8] iterates between point correspondence assignment and updates to the Möbius transformation. While it guarantees local convergence, the quality of the alignment depends on the initialzation. Note that ICP has been used in other works on isometric shape correspondence (e.g., [21, 11]) without explicitly estimating Möbius transformations. Koehl and Hass [24] used gradient descent to estimate Möbius transformations. Similar to ICP, their algorithm is only locally convergent thus necessitating good initializations.

Lipman *et al*. [27] introduced the continuous Procrustes distance to measure shape similarity. A key step in the distance calculation requires Möbius alignment. To estimate a 3DOF Möbius transformation on the hyperbolic disc, a brute force search over all possible correspondences was applied. This is generally a very costly procedure.

Zeng *et al*. [40] proposed higher-order graph matching for shape correspondence, where they construct the higher-order energy terms based on Möbius transformation. It is well-known, however, that graph matching is intractable in general, thus their method is suitable only for very small input sizes. A subsequent heuristic is used to "upgrade" the result to a dense correspondence.

2. The Möbius search problem

Given two surfaces $\tilde{\mathcal{M}}$ and $\tilde{\mathcal{B}}$ with disk topology in 3D (e.g., 3D scans of faces), we wish to estimate the conformal mapping that aligns them. The surfaces are first conformally flattened to yield $\mathcal{M} = \Phi_{\tilde{\mathcal{M}}}(\tilde{\mathcal{M}})$ and $\mathcal{B} = \Phi_{\tilde{\mathcal{B}}}(\tilde{\mathcal{B}})$, where $\mathcal{M}$ and $\mathcal{B}$ both lie in the hyperbolic disc $\mathbb{D}$ [1]. In our work, we use mid-edge conformal flattening [32] (as implemented by [27]) to generate $\mathcal{M}$ and $\mathcal{B}$.

Conformal flattening is unique up to an automorphism of the hyperbolic disk, which is a Möbius transformation $f : \mathbb{D} \mapsto \mathbb{D}$ with the definition

$$f(\mathbf{x}|\mathbf{z}, \theta) = e^{i\theta} \frac{\mathbf{x} - \mathbf{z}}{1 - \mathbf{x}\bar{\mathbf{z}}}. \tag{1}$$

Here, $\mathbf{x}, \mathbf{z}$ are complex numbers corresponding to points in $\mathbb{D}$, $\bar{\mathbf{z}}$ denotes the complex conjugate of $\mathbf{z}$, and $\theta \in [-\pi; \pi]$. Observe that (1) can be decomposed as $f = g \circ h$, where $g : \mathbb{D} \mapsto \mathbb{D}$ and $h : \mathbb{D} \mapsto \mathbb{D}$, with the definitions

$$g(\mathbf{x}|\theta) = e^{i\theta}\mathbf{x} \;\text{ and }\; h(\mathbf{x}|\mathbf{z}) = \frac{\mathbf{x} - \mathbf{z}}{1 - \mathbf{x}\bar{\mathbf{z}}}. \tag{2}$$

Intuitively, parameter $\mathbf{z}$ specifies the point that is mapped by h to the center of $\mathbb{D}$, while θ is a rotation angle.

The task of conformally aligning $\tilde{\mathcal{M}}$ and $\tilde{\mathcal{B}}$ is thus reduced to finding the f that aligns $\mathcal{M}$ and $\mathcal{B}$.

2.1. Objective function

In practice the surfaces are discrete and noisy, thus we must search for the f that is the "best" in some sense. Let $\mathcal{M} = \{\mathbf{m}_j\}_{j=1}^{N_1}$ and $\mathcal{B} = \{\mathbf{b}_k\}_{k=1}^{N_2}$ be the set of points after flattening. Recall that the points lie in $\mathbb{D}$ and are expressed as complex numbers. We adopt the objective function of [10], originally for rigid registration, to our case:

$$Q(\mathbf{z}, \theta) = \sum_j \max_k \mathbb{I}\left(|f(\mathbf{m}_j|\mathbf{z}, \theta) - \mathbf{b}_k| \leq \epsilon\right). \tag{3}$$

Here, $|\mathbf{x}|$ denotes the magnitude of a complex number, ϵ is a matching threshold, and $\mathbb{I}(\cdot)$ is an indicator function that returns 1 if the input statement is true, and 0 otherwise. The constant threshold ϵ can be changed to ϵ_j to make it point specific and dependent on scaling effects of the flattening on $\mathbf{m}_j$.

In words, (3) evaluates the number of points that are aligned under $f(\mathbf{x}|\mathbf{z}, \theta)$. The inner $\max$ checks if there is a point in $\mathcal{B}$ that matches $\mathbf{m}_j$ — thus, *a priori* identified correspondences are not assumed. Further, a match is declared only if the distance between the points is within ϵ — thus, it is not expected that each point in $\mathcal{M}$ has a valid match in $\mathcal{B}$. This is crucial if the surfaces only partially overlap. Note that this objective function does not guarantee one-to-one matching, but that does not hurt the accuracy of the applications we tested.

2.2. Problem definition

Using (3), the Möbius search problem is defined as

$$q^* = \max_{\mathbf{z},\theta} Q(\mathbf{z}, \theta), \tag{4}$$

which equates to finding the Möbius transformation that aligns as many points from $\mathcal{M}$ with $\mathcal{B}$ as possible. The problem can be re-expressed as

$$u^* = \max_{\mathbf{z}} U(\mathbf{z}), \tag{5}$$

$$U(\mathbf{z}) = \max_{\theta} \sum_j \max_k \mathbb{I}\left(|f(\mathbf{m}_j|\mathbf{z}, \theta) - \mathbf{b}_k| \leq \epsilon\right). \tag{6}$$

The purpose of this rearrangement is to exploit the fact that, given $\mathbf{z}$, finding θ can be done very efficiently (Sec. 3), such that solving for θ can seen as "evaluating" $U(\mathbf{z})$. This enables the formulation of a BnB algorithm (Sec. 4) that optimizes $\mathbf{z}$ explicitly and θ implicitly.

3. Solving for rotation angle

Defining $\mathbf{m}'_j = h(\mathbf{m}_j|\mathbf{z})$, (6) can be rewritten as

$$U(\mathbf{z}) = \max_{\theta} \sum_j \max_k \mathbb{I}\left(\left|e^{i\theta}\mathbf{m}'_j - \mathbf{b}_k\right| \leq \epsilon\right). \quad (7)$$

Recall that $e^{i\theta}$ specifies (via Euler's equation) a rotation of the complex plane $\mathbb{C}$ about the origin. Define

$$\Omega^j_{\mathbf{z}} = \{e^{i\theta}\mathbf{m}'_j \mid \theta \in [-\pi, \pi]\} \quad (8)$$

as the circle resulting from rotating $\mathbf{m}'_j$ by 2π radians, and

$$\mathbb{O}^k_\epsilon = \{\mathbf{x} \mid \mathbf{x} \in \mathbb{C}, |\mathbf{x} - \mathbf{b}_k| \leq \epsilon\}. \quad (9)$$

as the disc centered at $\mathbf{b}_k$ of radius ϵ; see Fig. 1(b).

Let $[\alpha^{jk}_{\mathbf{z},1}, \alpha^{jk}_{\mathbf{z},2}]$ be the range of angles θ, such that rotating $\mathbf{m}'_j$ with any θ from the range will cause the point to fall into $\mathbb{O}^k_\epsilon$. Intuitively, any $\theta \in [\alpha^{jk}_{\mathbf{z},1}, \alpha^{jk}_{\mathbf{z},2}]$ will yield

$$\mathbb{I}(|e^{i\theta}\mathbf{m}'_j - \mathbf{b}_k| \leq \epsilon) = 1. \quad (10)$$

The range limits $\alpha^{jk}_{\mathbf{z},1}$ and $\alpha^{jk}_{\mathbf{z},2}$ can be obtained in closed form via circle-to-circle intersections [2], see Fig. 1(b). In the case where $\mathbb{O}^k_\epsilon$ does not intersect $\Omega^j_{\mathbf{z}}$, the range is empty, implying that no θ can cause $\mathbf{m}'_j$ to match $\mathbf{b}_k$. For details of calculating the range limits, see the supplementary material.

For each $\mathbf{m}'_j$, let the set of angular ranges be

$$\mathcal{S}^j_{\mathbf{z}} = \left\{[\alpha^{jk}_{\mathbf{z},1}, \alpha^{jk}_{\mathbf{z},2}]\right\}_{k=1}^{N_2}. \quad (11)$$

Note that overlapping ranges in $\mathcal{S}^j_{\mathbf{z}}$ are merged, while ranges that extend beyond $[-\pi, \pi]$ are "wrapped around"; see Fig. 1(c). Function (7) can then be re-expressed as

$$U(\mathbf{z}) = \max_{\theta} \sum_j \max_{[\alpha_1,\alpha_2]\in\mathcal{S}^j_{\mathbf{z}}} \mathbb{I}(\theta \in [\alpha_1, \alpha_2]). \quad (12)$$

In words, evaluating $U(\mathbf{z})$ amounts to finding the θ that intersects as many as possible the angular ranges across $S^j_{\mathbf{z}}$, $j = 1, \ldots, N_1$; see Fig. 1(c). Such a problem can be solved exactly and efficiently in $\mathcal{O}(N \log N)$ [13, Chapter 10]; see the supplementary material for the detailed algorithm.

4. Main algorithm

The idea of BnB is to recursively partition and prune the search space until the globally optimal solution is found. In the context of maximizing U (5), the search space is the hyperbolic disc $\mathbb{D}$ (as mentioned in Sec. 2.2, the search for θ is done implicitly). The main "design choices" are how to partition $\mathbb{D}$, and how to construct an upper bounding function $\hat{U}$ for pruning subregions of $\mathbb{D}$. Algorithm 1 summarizes our algorithm, and details are provided in the following.

4.1. Partitioning the hyperbolic disc

Algorithm 1 is initialized by enclosing $\mathbb{D}$ with the tightest bounding square (a subset of the complex plane $\mathbb{C}$). The square is recursively divided into four equal sub-squares (Line 13). In each sub-square $\mathbb{S}$, we attempt to update our current best solution (Line 11), or to prune $\mathbb{S}$ using the bounding function $\hat{U}$ (Line 15). A sub-square $\mathbb{S}$ that cannot be pruned is inserted into a priority queue $\mathcal{Q}$ for further processing. Note that since we actually partition the unit square that encompasses $\mathbb{D}$, a square $\mathbb{S}$ may lie outside of $\mathbb{D}$. Thus if $\mathbb{S} \cap \mathbb{D} = \emptyset$, $\mathbb{S}$ is discarded (Line 5). The above partitioning and pruning steps effectively explores the space $\mathbb{D}$. Intuitively, a tighter bounding function $\hat{U}(\mathbb{S})$ will prune more aggressively, thus leading to fewer iterations. In the following, we describe our bounding function.

4.2. Bounding function

Given a square region $\mathbb{S}$, we must derive an upper bounding function $\hat{U}(\mathbb{S})$ that satisfies

$$\hat{U}(\mathbb{S}) \geq \max_{\mathbf{z}\in\mathbb{S}} U(\mathbf{z}) \quad (13)$$

to enable pruning in Algorithm 1.

We begin by seeking to bound the region

$$\mathbb{M}^j_{\mathbb{S}} = \{h(\mathbf{m}_j|\mathbf{z}) \mid \mathbf{z} \in \mathbb{S}\} \quad (14)$$

arising from the uncertainty of $\mathbf{z} \in \mathbb{S}$ for each $\mathbf{m}_j$. For simplicity, we approach this via the tightest bounding disc $\mathbb{R}$ of $\mathbb{S}$ instead; see Fig. 2(a). Clearly, since $\mathbb{S} \subseteq \mathbb{R}$, then $\mathbb{M}^j_{\mathbb{R}}$ is a bound over $\mathbb{M}^j_{\mathbb{S}}$. Now, $\mathbb{M}^j_{\mathbb{R}}$ can itself be bounded within a "wedge" $\mathbb{W}^j_{\mathbb{R}}$ defined by 4 parameters: the bounding radii

$$r^j_{\mathbb{R},1} = \min_{\mathbf{z}\in\mathbb{R}} |h(\mathbf{m}_j|\mathbf{z})|, \quad r^j_{\mathbb{R},2} = \max_{\mathbf{z}\in\mathbb{R}} |h(\mathbf{m}_j|\mathbf{z})|, \quad (15)$$

and the bounding angles

$$\theta^j_{\mathbb{R},1} = \min_{\mathbf{z}\in\mathbb{R}} \angle h(\mathbf{m}_j|\mathbf{z}), \quad \theta^j_{\mathbb{R},2} = \max_{\mathbf{z}\in\mathbb{R}} \angle h(\mathbf{m}_j|\mathbf{z}); \quad (16)$$

see Fig. 2(a). Hence, to bound the region $\mathbb{M}^j_{\mathbb{S}}$, we determine $\mathbb{W}^j_{\mathbb{R}}$ that is defined by the four parameters above.

Bounding radii Based on standard identities of complex numbers [3], we observe the symmetry

$$|h(\mathbf{m}_j|\mathbf{z})| = \left|\frac{\mathbf{m}_j - \mathbf{z}}{1 - \mathbf{m}_j\bar{\mathbf{z}}}\right| = \frac{|\mathbf{m}_j - \mathbf{z}|}{|1 - \mathbf{m}_j\bar{\mathbf{z}}|} \quad (17)$$
$$= \frac{|-(\mathbf{m}_j - \mathbf{z})|}{|\overline{1 - \mathbf{m}_j\bar{\mathbf{z}}}|} = \frac{|\mathbf{z} - \mathbf{m}_j|}{|1 - \mathbf{z}\bar{\mathbf{m}}_j|} = |h(\mathbf{z}|\mathbf{m}_j)|.$$

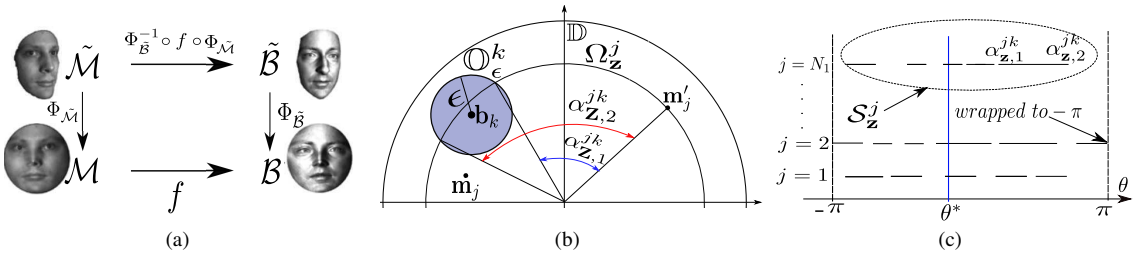

(a) (b) (c)

Figure 1. (a) Surface alignment using conformal mapping. (b) Intersection of $\Omega_{\mathbf{z}}^j$ with $\mathbb{O}_\epsilon^k$. (c) Illustration of problem (12) .

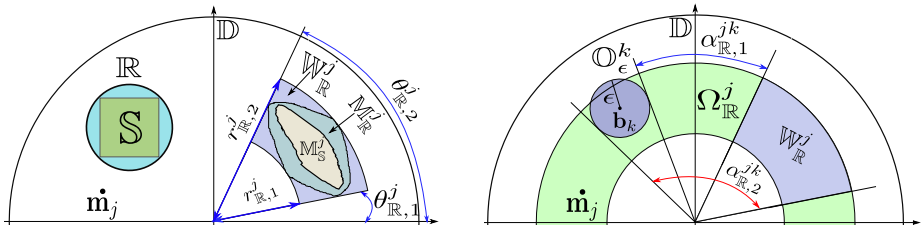

(a) (b)

Figure 2. (a) Uncertainty area with bounding parameters. (b) Intersection of the annulus $\Omega_{\mathbb{R}}^j$ with $\mathbb{O}_\epsilon^k$.

Algorithm 1 Möbius Search

Require: Point sets $\mathcal{M}, \mathcal{B} \subseteq \mathbb{D}$, threshold ϵ

$\mathcal{Q} \leftarrow$ empty priority queue, $\mathbb{S} \leftarrow$ tightest bounding square of $\mathbb{D}$, $u^* \leftarrow 0$, $\mathbf{z}^* \leftarrow null$
Insert $\mathbb{S}$ into $\mathcal{Q}$ with priority $\hat{U}(\mathbb{S})$
while $\mathcal{Q}$ is not empty **do**
  Obtain a square $\mathbb{S}$ with the highest priority from $\mathcal{Q}$
  if $\mathbb{S} \cap \mathbb{D} \neq \emptyset$ **then**
    $\mathbf{z}_0 \leftarrow$ center point of $\mathbb{S}$
    if $U(\mathbf{z}_0) = u^*$ **then**
      $\mathbf{z}^* \leftarrow \mathbf{z}_0$
      return u^*
    else if $U(\mathbf{z}_0) > u^*$ **then**
      $\mathbf{z}^* \leftarrow \mathbf{z}_0$, $u^* \leftarrow U(\mathbf{z}_0)$
    end if
    Subdivide $\mathbb{S}$ into four squares $\{\mathbb{S}_l\}_{l=1}^4$
    for all $l = 1, \ldots, 4$ **do**
      if $\hat{U}(\mathbb{S}_l) > u^*$ **then**
        Insert $\mathbb{S}_l$ into $\mathcal{Q}$ with priority $\hat{U}(\mathbb{S}_l)$
      end if
    end for
  end if
end while
return u^* and $\mathbf{z}^*$; obtain θ^* solving $U(\mathbf{z}^*)$ (12).

The bounding radii (15) can thus also be obtained as

$$r_{\mathbb{R},1}^j = \min_{\mathbf{z}\in\mathbb{R}} |h(\mathbf{z}|\mathbf{m}_j)|, \quad r_{\mathbb{R},2}^j = \max_{\mathbf{z}\in\mathbb{R}} |h(\mathbf{z}|\mathbf{m}_j)|. \tag{18}$$

Let the center and radius of $\mathbb{R}$ be $\mathbf{c}_{\mathbb{R}}$ and $r_{\mathbb{R}}$. The range of $h(\mathbf{z}|\mathbf{m}_j)$ for all $\mathbf{z} \in \mathbb{R}$ is defined as

$$\mathbb{N}_{\mathbb{R}}^j = \{h(\mathbf{z}|\mathbf{m}_j) \mid \mathbf{z} \in \mathbb{R}\}. \tag{19}$$

Now, it is known that, if $\mathbb{R}$ is a disc, $\mathbb{N}_{\mathbb{R}}^j$ is also a disc [30, Chapter 3]. Further, the center and radius of $\mathbb{N}_{\mathbb{R}}^j$ are

$$\mathbf{c}_{\mathbb{N}_{\mathbb{R}}^j} = \frac{\boldsymbol{\gamma} - \mathbf{m}_j}{1 - \bar{\mathbf{m}}_j \boldsymbol{\gamma}}, \tag{20}$$

$$r_{\mathbb{N}_{\mathbb{R}}^j} = \left| \mathbf{c}_{\mathbb{N}_{\mathbb{R}}^j} - \frac{\mathbf{c}_{\mathbb{R}} + r_{\mathbb{R}} - \mathbf{m}_j}{1 - \bar{\mathbf{m}}_j(\mathbf{c}_{\mathbb{R}} + r_{\mathbb{R}})} \right|, \tag{21}$$

$$\text{where} \quad \boldsymbol{\gamma} = \mathbf{c}_{\mathbb{R}} - r_{\mathbb{R}}^2 / \overline{(-1/\bar{\mathbf{m}}_j + \mathbf{c}_{\mathbb{R}})}. \tag{22}$$

Note that the region $\mathbb{M}_{\mathbb{R}}^j$, obtained by reversing the role of $\mathbf{z}$ and $\mathbf{m}_j$ in (19), is not a disc in general.

The bounding radii (18) can then be calculated as

$$r_{\mathbb{R},1}^j = |\mathbf{c}_{\mathbb{N}_{\mathbb{R}}^j}| - r_{\mathbb{N}_{\mathbb{R}}^j}, \quad r_{\mathbb{R},2}^j = |\mathbf{c}_{\mathbb{N}_{\mathbb{R}}^j}| + r_{\mathbb{N}_{\mathbb{R}}^j} \tag{23}$$

in closed form, where we offset the former to 0 if negative, and clamp the latter to 1 if greater than 1.

Bounding angles Manipulating $h(\mathbf{m}_j|\mathbf{z})$ again by

$$h(\mathbf{m}_j|\mathbf{z}) = \frac{(\mathbf{m}_j - \mathbf{z})(\overline{1 - \mathbf{m}_j\bar{\mathbf{z}}})}{(1 - \mathbf{m}_j\bar{\mathbf{z}})(\overline{1 - \mathbf{m}_j\bar{\mathbf{z}}})} = \frac{(\mathbf{m}_j - \mathbf{z})(1 - \bar{\mathbf{m}}_j\mathbf{z})}{|(1 - \mathbf{m}_j\bar{\mathbf{z}})|^2},$$

we can express $h(\mathbf{m}_j|\mathbf{z})$ as a multiplication and scaling of two complex numbers. Using the identity

$$\mathbf{ab} = |\mathbf{a}||\mathbf{b}|e^{i(\angle\mathbf{a}+\angle\mathbf{b})} \quad \forall \mathbf{a}, \mathbf{b} \in \mathbb{C}, \tag{24}$$

of complex numbers [3], we can surmise that

$$\angle h(\mathbf{m}_j|\mathbf{z}) = \angle(\mathbf{m}_j - \mathbf{z}) + \angle(1 - \bar{\mathbf{m}}_j\mathbf{z}). \quad (25)$$

Define the regions

$$\mathbb{A}^j_{\mathbb{R}} = \{\mathbf{m}_j - \mathbf{z} \mid \mathbf{z} \in \mathbb{R}\}, \quad \mathbb{B}^j_{\mathbb{R}} = \{1 - \bar{\mathbf{m}}_j\mathbf{z} \mid \mathbf{z} \in \mathbb{R}\}.$$

Clearly $\mathbb{A}^j_{\mathbb{R}}$ a disc; obtained by reflecting disc $\mathbb{R}$ and translating the result by $\mathbf{m}_j$. Its center and radius are respectively

$$\mathbf{c}_{\mathbb{A}^j_{\mathbb{R}}} = -\mathbf{c}_{\mathbb{R}} + \mathbf{m}_j, \quad r_{\mathbb{A}^j_{\mathbb{R}}} = r_{\mathbb{R}}. \quad (26)$$

Since multiplying two complex numbers serves to multiply their respective magnitudes (24), multiplying $\mathbb{R}$ with $\bar{\mathbf{m}}_i$ expands the disc by a factor of $|\mathbf{m}_j|$. Thus $\mathbb{B}^j_{\mathbb{R}}$ is also a disc with center and radius respectively

$$\mathbf{c}_{\mathbb{B}^j_{\mathbb{R}}} = -\bar{\mathbf{m}}_j\mathbf{c}_{\mathbb{R}} + 1, \quad r_{\mathbb{B}^j_{\mathbb{R}}} = |\mathbf{m}_j| r_{\mathbb{R}}. \quad (27)$$

The bounding angles (16) can then be calculated as

$$\theta^j_{\mathbb{R},1} = \min_{\mathbf{a}\in\mathbb{A}^j_{\mathbb{R}}, \mathbf{b}\in\mathbb{B}^j_{\mathbb{R}}} \angle\mathbf{a} + \angle\mathbf{b}, \quad (28)$$

$$\theta^j_{\mathbb{R},2} = \max_{\mathbf{a}\in\mathbb{A}^j_{\mathbb{R}}, \mathbf{b}\in\mathbb{B}^j_{\mathbb{R}}} \angle\mathbf{a} + \angle\mathbf{b}. \quad (29)$$

These values can be obtained in closed form, since the angular ranges of $\mathbb{A}^j_{\mathbb{R}}$ and $\mathbb{B}^j_{\mathbb{R}}$ are known.

Bound calculation Given the wedge $\mathbb{W}^j_{\mathbb{R}}$, we are now ready to compute the upper bound (13). Our strategy here is a generalization of the technique in Sec. 3.

First, generalizing (8), we define the annulus

$$\Omega^j_{\mathbb{R}} = \{e^{i\theta}\mathbf{x} \mid \mathbf{x} \in \mathbb{W}^j_{\mathbb{R}}, \theta \in [-\pi, \pi]\} \quad (30)$$

obtained by rotating the wedge $\mathbb{W}^j_{\mathbb{R}}$ by 2π radians; see Fig. 2(b). Continuing the idea in Sec. 3, for each pair (j, k), we obtain the angular range $[\alpha^{jk}_{\mathbb{R},1}, \alpha^{jk}_{\mathbb{R},2}]$, that bounds the rotation angle θ that allows a point from $\mathbb{W}^j_{\mathbb{R}}$ to "touch" $\mathbb{O}^k_\epsilon$; see Fig. 2(b) for an intuitive example. The range limits can also be obtained in closed form; for brevity, we leave the details in the supplementary material.

Of course, if $\mathbb{O}^k_\epsilon$ does not intersect with $\Omega^j_{\mathbb{R}}$, then the range is empty. This implies that $f(\mathbf{m}_j|\mathbf{z}, \theta)$ cannot match with $\mathbf{b}_k$ under all $\mathbf{z} \in \mathbb{S}$ and $\theta \in [-\pi, \pi]$.

Define now $\mathcal{S}^j_{\mathbb{R}}$ to be the set of angular intervals

$$\mathcal{S}^j_{\mathbb{R}} = \{[\alpha^{jk}_{\mathbb{R},1}, \alpha^{jk}_{\mathbb{R},2}]\}_{k=1}^{N_2}. \quad (31)$$

Again, overlapping and out-of-bound ranges are preprocessed as in Sec. 3. The upper bound is evaluated as

$$\hat{U}(\mathbb{S}) = \max_\theta \sum_j \max_{[\alpha_1,\alpha_2]\in\mathcal{S}^j_{\mathbb{R}}} \mathbb{I}(\theta \in [\alpha_1, \alpha_2]), \quad (32)$$

which again can be solved exactly and efficiently as a line intersection problem; cf. (12).

4.3. Algorithm convergence

Here, we establish the proofs required [20] to guarantee that Algorithm 1 converges to the globally optimal result.

Lemma 1 *$\hat{U}(\mathbb{S})$ obtained according to (32) satisfies (13).*

Proof By design, the relationship

$$\mathbb{M}^j_{\mathbb{S}} \subseteq \mathbb{M}^j_{\mathbb{R}} \subseteq \mathbb{W}^j_{\mathbb{R}} \quad (33)$$

always holds. Thus, the annulus $\Omega^j_{\mathbb{R}}$ bounds the location of $f(\mathbf{m}_j|\mathbf{z}, \theta)$ for all $\mathbf{z} \in \mathbb{S}$ and $\theta \in [-\pi, \pi]$. The angular intervals $\mathcal{S}^j_{\mathbb{R}}$ are also optimistic since they are constructed by aligning $\mathbb{W}^j_{\mathbb{R}}$ with $\mathbb{O}^k_\epsilon$ for all k. This establishes that $\hat{U}(\mathbb{S})$ cannot underestimate $U(\mathbf{z})$ for all $\mathbf{z} \in \mathbb{S}$. ∎

Lemma 2 *As $\mathbb{S}$ collapses to a single point $\mathbf{z}$,*

$$\hat{U}(\mathbb{S}) = U(\mathbf{z}). \quad (34)$$

Proof If $\mathbb{S}$ is a single point $\mathbf{z}$, then $\mathbb{M}^j_{\mathbb{S}}$, defined in (14), equates to the singleton set $\{h(\mathbf{m}_j|\mathbf{z})\}$. Since $\mathbb{R}$ is the tightest bounding disc of $\mathbb{S}$, $\mathbb{M}^j_{\mathbb{R}}$ also equates to $\mathbb{M}^j_{\mathbb{S}}$. Now, based on definitions (15) and (16), $\mathbb{W}^j_{\mathbb{R}}$ also collapses to a single point $\{h(\mathbf{m}_j|\mathbf{z})\}$, thus yielding $\mathbb{M}^j_{\mathbb{S}} = \mathbb{M}^j_{\mathbb{R}} = \mathbb{W}^j_{\mathbb{R}}$.

The annulus $\Omega^j_{\mathbb{R}}$ thus becomes the circle $\Omega^j_{\mathbf{z}}$, and the angular ranges $\mathcal{S}^j_{\mathbb{R}}$ and $\mathcal{S}^j_{\mathbf{z}}$ are equal. Thus, $\hat{U}(\mathbb{S})$ as defined in (32) reduces to $U(\mathbf{z})$ as defined in (12). ∎

5. Results

We benchmarked the performance of our algorithm (Möbius Search, henceforth represented as MS) against the previous methods surveyed in Sec. 1.1, namely

- Möbius voting (MV) [28]. We used the code provided by the authors [4]. However, since in this paper we focus on aligning surfaces that are topological discs (f is 3DOF), we modified the code such that a random sample consists of two randomly chosen correspondences.
- Brute force method (BF). Following [27], we implemented BF as follows: all possible pairings between $\mathcal{M}$ and $\mathcal{B}$ are considered. Each pairing is sufficient to estimate $\mathbf{z}$. For each $\mathbf{z}$, θ is enumerated across a sufficiently fine grid along $[-\pi, \pi]$ to find the best combination.
- Iterative closest points (ICP). The original method of [8] was modified as alluded in [28] for Möbius alignment.
- Graph matching (GM) [40]. We used the implementation of [26] for graph matching. Since our paper focuses on the 3DOF Möbius transform f, we included up to binary energies only in the cost function. Note, however, that this does not significantly simplify the problem, since graph matching is intractable even for binary graphs [5].

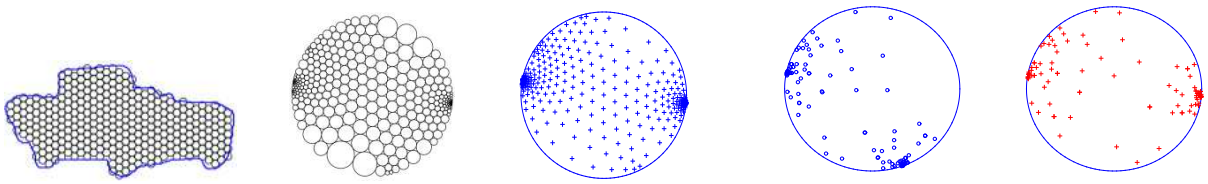

(a) (b) (c) (d) (e)

Figure 3. Steps for generating synthetic data: (a) car-01 shape; (b) conformal map of (a); (c) centers of circles in (b); (d) $\mathcal{M}$ being sub-sampled from (c); (e) $\mathcal{B}$ was generated by applying a random Möbius transformation to $\mathcal{M}$ then added with noise and outliers.

Although we did not compare against the gradient descent method of [24], as a locally convergent method, we expect its performance to be similar to ICP. Also, GM is only feasible for small input sizes N_1, N_2. In our experiments, GM was run with $N_1, N_2 \leq 20$ (in [40], input sizes of at most 15 was tested for the true graph matching part).

All experiments were run on a standard PC with 3.5GHz processor and 8GB of main memory. Due to page limits, only representative results can be shown here; see the supplementary material for more results.

5.1. Comparison metrics

Given a pair of conformally flattened surfaces $\mathcal{M}$ and $\mathcal{B}$, each method above was executed to estimate f. Apart from recording the runtime, we also obtained the following quality measures of the estimated f:

- Qbnb: the value of (3) for f.
- Qmv: the number of mutually closest pairs under f, where (j, k) is a mutually closest pair if $\mathbf{b}_k$ is the nearest neighbor of $f(\mathbf{m}_j)$ among $\mathcal{B}$ and vice versa.

Note that Qmv is as defined and used in [28, 27] for assessing deformation errors of Möbius alignment.

Where ground truth correspondences $\{\mathbf{m}_t, \mathbf{b}_t\}_{t=1}^{T}$ (from landmarks etc.) were available, we used them to calculate the following quality metric:

- Qtruth: the number of ground truth correspondences that are mutually closest pairs under f.

5.2. Synthetic data experiment

The purpose of this experiment is to evaluate the performance and accuracy of the methods under controlled settings. The steps to generate input point sets $\mathcal{M}$ and $\mathcal{B}$ are summarized in Fig. 3: first, a 2D shape from the MPEG7 dataset [25] (specifically, car-01) was chosen and conformally mapped to $\mathbb{D}$ using the circle packing technique [34]. A number of N_1 points were then randomly sampled to produce the set $\mathcal{M}$. A random Möbius transformation f was generated (by randomly choosing $\mathbf{z}$ and θ) and applied on $\mathcal{M}$ to yield the set $\mathcal{B}$. Gaussian noise of $\sigma = 0.01$ was afflicted on $\mathcal{B}$ to increase realism. Further, to simulate outliers and partially overlapping data, $\rho\%$ of points on both $\mathcal{M}$ and $\mathcal{B}$ were randomly chosen and re-sampled to lie uniformly in $\mathbb{D}$. In our experiment, we used $N_1 = \{100, 50, 20\}$ and $\rho = \{0, 25, 50\}$. For MS and (3), ϵ =0.01 was used. Again, note that GM is only feasible for $N_1 = 20$.

Fig. 4 shows the results. It can be seen that MS always achieved the theoretical maximum $(1 - \rho)N_1$ of the objective function (3). In the presence of low oulier rates (ρ is 0% or 25%), the estimated f of MS, MV, BF and GM were of similar quality. When the outlier percentage ρ was 50%, however, only MS could produce good (in fact, optimal) results. Note that in Fig. 4 the quality of ICP was much lower than the others due to the lack of good initializations

To accommodate ICP, the experiment was repeated by producing a "milder" randomized Möbius transformation that relates $\mathcal{M}$ to $\mathcal{B}$, specifically by choosing the parameters such that $|\mathbf{z}| \leq 0.1$ (close to the center of $\mathbb{D}$) and $\theta \leq 10°$. The results for this repeated experiment were displayed as ICP2 in Fig. 4. It can easily be seen that with good initialization, ICP gave acceptable quality when there were no outliers - however, the quality degraded rapidly as the outlier rate increased. Interestingly, the runtime of MS increased marginally when the true f was close to the identity mapping - possibly because as the true $\mathbf{z}$ is closer to the center of $\mathbb{D}$, a deeper search must be conducted to before a good suboptimal solution $\mathbf{z}_0$ is retrieved to enable effective pruning.

In terms of runtime, MS and MV could terminate well within 1 minute, though MS occasionally took longer than MV. Note that MS provides guaranteed global optimality, unlike MV. The runtime for BF, as anticipated, was too long to be practical, e.g., more than 3 hours for $N = 100$. To view the actual numerical values of the results above, see the supplementary material.

5.3. Conformal teeth alignment experiment

We followed the experiment by Boyer et al. [9] to perform surface alignment on 3D scans of teeth. While the original aim of Boyer et al. was shape comparison, here, we focus on the alignment step. In our experiment, we chose three pairs of teeth originally used in [27], specifically, Human01 and Human04, Chimpanzee51376 and Chimpanzee51379, and Gorilla167335 and Gorilla16736.

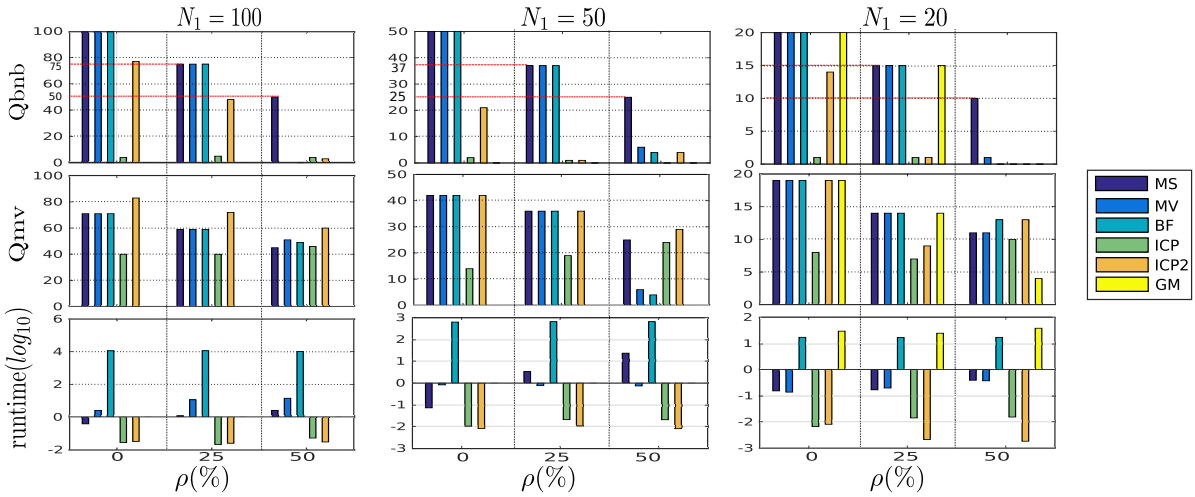

Figure 4. Results for synthetic data. Columns represent experiments for different values of N_1. Rows represent measurements of Qbnb (defined in (3)), Qmv (defined in 5.1) and runtime (in seconds). Note that the runtime is quoted in log_{10} scale. ICP2 is explained in 5.2. GM is only feasible for $N_1 = 20$.

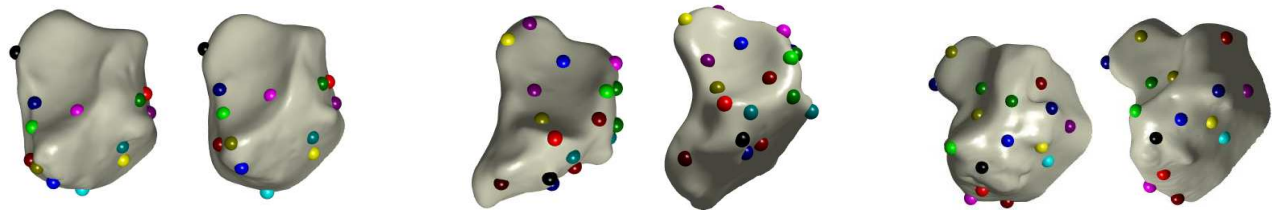

(a) Human 01 (b) Human 04 (c) Chimp. 51376 (d) Chimp.51379 (e) Gorilla 167335 (f) Gorilla 167336

Figure 5. Correspondences found by MS for three pairs of teeth.

The meshes of the teeth were flattened to $\mathbb{D}$ using the program of Lipman et al. [27]. On the flattened surfaces, we also conducted the sampling process of [28, 9] to create the point sets $\mathcal{M}$ and $\mathcal{B}$. Specifically, N_1 and N_2 points ($N_1, N_2 = \{100, 50, 20\}$) were chosen using the farthest point sampling (FPS) algorithm [14]. Note that 13 ground truth correspondences $\{\mathbf{m}_t, \mathbf{b}_t\}_{t=1}^{13}$ (manually annotated landmarks) were available per problem instance, thus Qtruth value could be obtained for each method.

Table 1 summarizes the quantitative results, while Fig. 5 shows qualitative results for MS. As expected, due to the global optimality guarantees, MS returned the solution with the highest Qbnb value. Also, MS demonstrated typically superior accuracy in terms of Qtruth, as compared to BF and MV. However, when the input size was small ($N_1, N_2 = 20$), none of the methods were able to obtain satisfactory Qtruth values. This was due to the overly impoverished structural information after excessive sampling. Due to the lack of good initializations (the intialized state of $\mathcal{M}$ and $\mathcal{B}$ depends on the implementation of the conformal flattening procedure), ICP generally could not find good estimates of f, and it was able to align about half of the ground truth correspondences. While GM was feasible on $N_1, N_2 = 20$, it is apparent that the estimated f was far from ideal due to the overly sparse input data. In terms of runtime, all the methods except BF were able to terminate in about or less than 1 minute.

5.4. Conformal face alignment experiment

The previous experiment was repeated for conformal face alignment, following [39]. While the previous works aimed at applications such as facial expression recognition, in our experiment, we focused on the task of Möbius alignment. From a practical standpoint, our MS algorithm can be used to automatically and deterministically find landmark correspondences, which is a crucial step in facial processing applications [39].

Again, following the previous works, we used data from the BU-3DFE face dataset [37], specifically, we chose three pairs of faces with the same expression but at different degrees: Happy 01-Happy 04, Disgust 01-Disgust 02, Sad 01 - Sad 02. The same steps as in Sec. 5.3 were used for flattening and subsampling; see Fig. 6 for the resulting data. For this dataset, since the ground truth landmarks were not available, we manually annotated 13 landmarks on the faces to create ground truth correspondences. Table 1 summa-

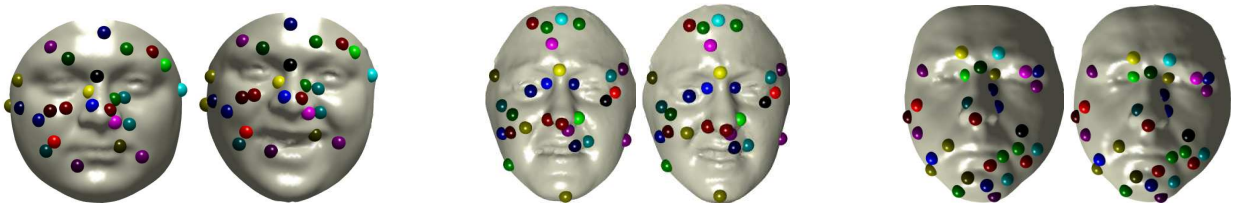

(a) Happy 01 (b) Happy 04 (c) Digust 01 (d) Disgust 02 (e) Sad 01 (f) Sad 02

Figure 6. Correspondences found by MS between three pairs of face

		N_1, N_2	Methods	Qbnb	Qmv	Qtruth	Time (sec)
Human01	Human04	100	MS	**40**	80	**13**	45.056
			BF	30	**84**	**13**	10185.000
			MV	35	81	**13**	29.160
			ICP	22	76	6	**0.027**
		50	MS	**8**	38	**13**	8.780
			BF	3	**41**	12	629.000
			MV	3	39	5	6.891
			ICP	4	37	6	**0.009**
		20	MS	**11**	**18**	11	0.809
			BF	6	**18**	8	17.265
			MV	**11**	**18**	**12**	0.408
			GM	5	11	3	26.280
			ICP	3	16	7	**0.005**
Gorilla167335	Gorilla167336	100	MS	**68**	**83**	**12**	12.370
			BF	66	81	**12**	10659.000
			MV	61	80	**12**	5.651
			ICP	8	78	6	**0.052**
		50	MS	**26**	**39**	6	3.234
			BF	22	38	5	634.000
			MV	24	**39**	6	5.970
			ICP	1	36	**7**	**0.005**
		20	MS	**11**	**16**	5	0.180
			MV	10	15	4	0.347
			BF	10	**16**	**6**	17.208
			GM	8	10	1	77.513
			ICP	0	14	**6**	**0.001**
Chimpanzee51376	Chimpanzee51379	100	MS	**54**	82	**12**	26.913
			BF	52	**85**	**12**	10365.000
			MV	47	84	**12**	28.045
			ICP	22	48	1	**0.027**
		50	MS	**25**	**42**	**13**	1.774
			BF	19	**42**	8	691.000
			MV	20	38	8	3.684
			ICP	7	22	3	**0.005**
		20	MS	**13**	15	0	0.857
			BF	7	**17**	**4**	17.056
			MV	7	12	2	0.260
			GM	6	10	0	47.236
			ICP	4	9	1	**0.068**

		N_1, N_2	Methods	Qbnb	Qmv	Qtruth	Time (sec)
F0001_DI01WH_F3D	F0001_DI02WH_F3D	100	MS	**75**	81	**13**	1.999
			BF	71	**85**	**13**	10656.000
			MV	74	66	7	2.853
			ICP	46	61	1	**0.021**
		50	MS	**22**	37	**13**	1.792
			BF	19	**41**	**13**	635.000
			MV	18	40	**13**	1.870
			ICP	12	26	1	**0.005**
		20	MS	**10**	15	**9**	0.220
			BF	6	14	8	17.241
			MV	6	14	7	0.278
			GM	8	**16**	**9**	45.059
			ICP	3	11	0	**0.003**
M0044_HA01IN_F3D	M0044_HA04IN_F3D	100	MS	**53**	83	**13**	15.768
			BF	34	**84**	**13**	10269.000
			MV	40	80	12	28.648
			ICP	18	61	1	**0.022**
		50	MS	**27**	42	**13**	4.109
			BF	25	**44**	**13**	641.041
			MV	23	**44**	**13**	1.549
			ICP	3	32	1	**0.010**
		20	MS	**16**	**19**	**13**	0.373
			BF	14	**19**	**13**	17.544
			MV	**16**	**19**	**13**	0.286
			GM	7	12	9	55.105
			ICP	11	14	0	**0.002**
M0021_SA01WH_F3D	M0021_SA02WH_F3D	100	MS	**45**	85	**13**	15.081
			BF	42	**88**	**13**	9695.200
			MV	28	56	0	56.793
			ICP	32	63	0	**0.028**
		50	MS	**24**	45	**13**	14.474
			BF	23	**46**	**13**	623.966
			MV	3	19	4	1.526
			ICP	3	34	0	**0.006**
		20	MS	**14**	14	**13**	0.231
			BF	**14**	**15**	**13**	17.445
			MV	**14**	14	**13**	0.140
			GM	4	13	9	69.509
			ICP	5	12	0	**0.006**

Table 1. Results from conformally aligning three pairs of teeth (left) and three pairs of faces (right). In each problem instance, the best quality measure and runtime obtained among all the methods are bolded. MS: Möbius Search, MV: Möbius voting, BF: brute force, GM: graph matching. See Sec. 5.1 for definitions of the quality measures.

rizes the quantitative results, while Fig. 6 illustrates qualitative results of MS. It can be seen that MS generally outperformed the other methods in terms of both accuracy and runtime.

6. Conclusions and future work

We proposed a novel approach for conformal surface alignment with guaranteed global optimum. Our experiments showed that this algorithm is much more efficient than state-of-the-art techniques for conformally aligning topological disc surfaces.

This work opens up a new direction for further research on global optimization methods in the field of computational conformal geometry. One notable expansion which can be studied in the future is 6DOF Möbius search for genus zero surfaces with spherical topology.

Acknowledgements This work was supported by ARC grants DP160103490 and DP130102524.

References

[1] https://en.wikipedia.org/wiki/Poincare_disk_model. 2

[2] http://mathworld.wolfram.com/Circle-CircleIntersection.html. 3

[3] http://mathworld.wolfram.com/ComplexNumber.html. 3, 5

[4] https://www.cs.princeton.edu/~vk/projects/CorrsCode/. 5

[5] https://en.wikipedia.org/wiki/Quadratic_assignment_problem. 5

[6] S. Angenent, S. Haker, A. Tannenbaum, and R. Kikinis. Conformal geometry and brain flattening. In *MICCAI*, 1999. 1

[7] L. Bers. Uniformization, moduli and Kleinian groups. In *Bull. London Math. Soc.*, volume 4, pages 257–300, 1972. 1

[8] P. J. Besl and N. D. MacKay. A method for registration of 3-d shapes. *IEEE TPAMI*, 14(2):239–256, 1992. 1, 2, 5

[9] D. Boyer, Y. Lipman, E. StClair, J. Puente, B. Patel, T. Funkhouser, J. Jernvall, and I. Daubechies. Algorithms to automatically quantify the geometry similarity of anatomical surfaces. In *Proc. Nat'l Academy of Sciences*, 2011. 1, 6, 7

[10] T. Breuel. Implementation techniques for geometric branch-and-bound matching methods. *CVIU*, 90(3):258–294, 2003. 2

[11] A. Bronstein, M. Bronstein, and R. Kimmel. Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. In *Proc. Nat'l Academy of Science*, 2006. 2

[12] A. M. Bronstein, M. M. Bronstein, A. M. Bruckstein, and R. Kimmel. Analysis of two-dimensional non-rigid shapes. *IJCV*, 78(1):67–88, 2008. 1

[13] M. De Berg, M. Van Kreveld, M. Overmars, and O. C. Schwarzkopf. *Computational geometry*. Springer, 2000. 3

[14] Y. Eldar, M. Lindenbaum, M. Porat, and Y. Y. Zeevi. The farthest point strategy for progressive image sampling. *Image Processing, IEEE Transactions on*, 6(9):1305–1315, 1997. 7

[15] X. Gu, Y. Wang, T. Chan, P. Thompson, and S.-T. Yau. Genus zero surface conformal mapping and its application to brain surface mapping. *IEEE Trans. Medical Imaging*, 23(8):949–958, 2004. 1

[16] X. Gu and S.-T. Yau. Global conformal surface parametrization. In *Proc. Eurographics Symp. Geometry Processing*, 2003. 1

[17] X. D. Gu and B. C. Vemuri. Matching 3D shapes using 2d conformal representations. In *MICCAI*, 2004. 1

[18] X. D. Gu and S.-T. Yau. *Computational conformal geometry*. Advanced Lectures in Mathematics. Higher Education Press, 2008. 2

[19] R. I. Hartley and F. Kahl. Global optimization through rotation space search. *IJCV*, 82:64–79, 2009. 2

[20] R. Horst and H. Tuy. *Global optimization*. Springer, 1996. 1, 5

[21] Q.-X. Huang, B. Adams, M. Wicke, and L. J. Guibas. Non-rigid registration under isometric deformations. *Computer Graphics Forum*, 27(5):1449–1457, 2008. 2

[22] M. Hurdal, P. Bowers, K. Stephenson, D. Sumners, K. Rehm, K. Shaper, and D. Rotenberg. Quasiconformally flat mapping the human cerebellum. In *MICCAI*, 1999. 1

[23] M. K. Hurdal and K. Stephenson. Discrete conformal methods for cortical brain flattening. *Neuroimage*, 45(1):S86–S98, 2009. 2

[24] P. Koehl and J. Hass. Automatic alignment of genus-zero surfaces. *IEEE TPAMI*, 36(3):466–478, 2014. 1, 2, 6

[25] L. J. Latecki, R. Lakämper, and U. Eckhardt. Shape descriptors for non-rigid shapes with a single closed contour. In *Computer Vision and Pattern Recognition, 2000. Proceedings. IEEE Conference on*, volume 1, pages 424–429. IEEE, 2000. 6

[26] M. Leordeanu, M. Hebert, and R. Sukthankar. An integer projected fixed point method for graph matching and map inference. In *Advances in neural information processing systems*, pages 1114–1122, 2009. 5

[27] Y. Lipman, R. Al-Aifari, and I. Daubechies. The continuous Procrustes distance between two surfaces. *Communication on Pure and Applied Mathematics*, 66(6):934–964, 2011. 1, 2, 5, 6, 7

[28] Y. Lipman and T. Funkhouser. Möbius voting for surface correspondence. In *ACM Transactions on Graphics (TOG)*, volume 28, page 72. ACM, 2009. 1, 2, 5, 6, 7

[29] H. Lu, L.-P. Nolte, and M. Reyes. Interest points location for brain image using landmark-annotated atlas. *Int'l J. Imaging Systems Technology*, 22:145–152, 2012. 1

[30] D. Mumford, C. Series, and D. Wright. *Indra's pearls: the vision of Felix Klein*. Cambridge University Press, 2002. 4

[31] C. Olsson, F. Kahl, and M. Oskarsson. Branch-and-bound methods for Euclidean registration problems. *IEEE TPAMI*, 31(5):783–794, 2009. 2

[32] U. Pinkall and K. Polthier. Computing discrete minimal surfaces and their conjugates. *Experimental mathematics*, 2(1):15–36, 1993. 2

[33] B. Springborn, P. Schröder, and U. Pinkall. Conformal equivalence of triangle meshes. In *SIGGRAPH Asia*, 2008. 1

[34] K. Stephenson. *Introduction to circle packing: The theory of discrete analytic functions*. Cambridge University Press, 2005. 6

[35] O. Van Kaick, H. Zhang, G. Hamarneh, and D. Cohen-Or. A survey on shape correspondence. *Computer Graphics Forum*, 30(6):1681–1707, 2011. 1

[36] S. Wang, Y. Wang, M. Jin, X. D. Gu, and G. Samaras. Conformal geometry and its applications on 3D shape matching, recognition, and stitching. *IEEE TPAMI*, 29(7):1209–1220, 2007. 1, 2

[37] L. Yin, X. Wei, Y. Sun, J. Wang, and M. J. Rosato. A 3d facial expression database for facial behavior research. In *Automatic face and gesture recognition, 2006. FGR 2006. 7th international conference on*, pages 211–216. IEEE, 2006. 7

[38] W. Zeng and D. X. Gu. Conformal geometric methods in computer vision. In *CEWIT*, 2011. 1

[39] W. Zeng, H. Li, L. Chen, J.-M. Morvan, and X. D. Gu. An automatic 3d expression recognition framework based on sparse representation of conformal images. In *Automatic*

Face and Gesture Recognition (FG), 2013 10th IEEE International Conference and Workshops on, pages 1–8. IEEE, 2013. 7

[40] Y. Zeng, C. Wang, Y. Wang, X. Gu, D. Samaras, and N. Paragios. Dense non-rigid surface registration using high-order graph matching. In *CVPR*, pages 382–389. IEEE, 2010. 1, 2, 5, 6