Computing Teichmüller Maps between Polygons

Mayank Goswami[1], Xianfeng Gu[2], Vamsi P. Pingali[3], and Gaurish Telang[4]

1 **Algorithms and Complexity, Max-Planck Institute for Informatics Saarbrücken 66123, Germany**
gmayank@mpi-inf.mpg.de

2 **Department of Computer Science, Stony Brook University Stony Brook, NY 11794-4400, USA**
gu@cs.stonybrook.edu

3 **Department of Mathematics, Johns Hopkins University Baltimore, MD - 21218, USA**
vpingali@math.jhu.edu

4 **Department of Applied Mathematics and Statistics, Stony Brook University Stony Brook, NY 11794-3600, USA**
gaurish.telang@stonybrook.edu

Abstract

By the Riemann mapping theorem, one can bijectively map the *interior* of an n-gon P to that of another n-gon Q conformally (i.e., in an angle preserving manner). However, when this map is extended to the boundary it need not necessarily map the *vertices* of P to those of Q. For many applications it is important to find the "best" vertex-preserving mapping between two polygons, i.e., one that minimizes the maximum angle distortion (the so-called dilatation). Such maps exist, are unique, and are known as extremal quasiconformal maps or Teichmüller maps.

There are many efficient ways to approximate conformal maps, and the recent breakthrough result by Bishop computes a $(1+\epsilon)$-approximation of the Riemann map in linear time. However, only heuristics have been studied in the case of Teichmüller maps.

We present two results in this paper. One studies the problem in the continuous setting and another in the discrete setting.

In the continuous setting, we solve the problem of finding a finite time procedure for approximating Teichmüller maps. Our construction is via an iterative procedure that is proven to converge in $O(\text{poly}(1/\epsilon))$ iterations to a $(1+\epsilon)$-approximation of the Teichmüller map. Our method uses a reduction of the polygon mapping problem to the marked sphere problem, thus solving a more general problem.

In the discrete setting, we reduce the problem of finding an approximation algorithm for computing Teichmüller maps to two basic subroutines, namely, computing discrete 1) compositions and 2) inverses of discretely represented quasiconformal maps. Assuming finite-time solvers for these subroutines we provide a $(1+\epsilon)$-approximation algorithm.

1998 ACM Subject Classification I.3.5 Computational Geometry and Object Modeling

Keywords and phrases Teichmüller maps, Surface registration, Extremal Quasiconformal maps, Computer vision

Digital Object Identifier 10.4230/LIPIcs.SOCG.2015.615

1 Introduction

A foundational result in complex analysis, the Riemann mapping theorem, implies that the interiors of two n-gons P and Q can be mapped bijectively and conformally (i.e., in an angle

31st International Symposium on Computational Geometry (SoCG'15).
Editors: Lars Arge and János Pach; pp. 615–629
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

preserving way[1]) to one another. By a result of Caratheodory [4], such a map $f : P \to Q$ extends continuously to the boundary of P (the edges). Generally the vertices of P do not map to the vertices of Q under this extended mapping.

Consider the collection of functions f that map P to Q, and take the vertices of P to the vertices of Q. In general such an f is bound to stretch angles, and a classical way to measure this angle stretch by f at a point $p \in P$ is by means of a complex-valued function $\mu_f(p)$ called the Beltrami coefficient[2] of f. The Beltrami coefficient satisfies $||\mu_f||_\infty < 1$. If μ_f is identically zero, then f is conformal. The problem we consider is computing the "best" such map f_* in the above class, i.e., an f_* such that the norm of its Beltrami coefficient $||\mu_*||_\infty$ is the smallest amongst all (uncountably many) maps satisfying the above conditions. Bijective maps that stretch angles but by a bounded amount are called quasiconformal homeomorphisms (q.c.h.), and the best q.c.h. f_* is called the extremal q.c. map, or the Teichmüller map.

As an example consider two rectangles $R_i = [0, a_i] \times [0, b_i] (i = 1, 2)$ in the plane. Consider the space of all q.c.h. $f : R_1 \to R_2$ such that f takes the vertices to the vertices. It was shown by Grötzsch [12] that the affine map $f_*(x, y) = (a_2x/a_1, b_2y/b_1)$ with $\mu_*(x, y) = (1 - r)/(1 + r), r = b_2a_1/a_2b_1$ is the unique extremal q.c. map; any other map f would stretch angles at some point $p \in R_1$ more than g (i.e., $\exists p \in R_1 : |\mu_f(p)| > |\mu_*(p)|$). For the general n-gon case mentioned above, such a nice formula does not exist for the extremal map. However, the extremal map exists and is unique. These are the famous theorems of Teichmuller [22, 23], proven rigorously later by Ahlfors [1].

Algorithms for computing the Riemann map from a polygon to the disc [8, 7, 2] have gathered a lot of attention and found many applications. However, no such algorithm that approximates the extremal map is known. In contrast to the Riemann mapping theorem, where a constructive proof is known, the proof by Teichmüller/Ahlfors is an existence result only. In fact, to the authors' knowledge there does not exist a method that, given a starting f between P and Q, computes a g with $||\mu_g||_\infty < ||\mu_f||_\infty$ if one exists. We are motivated by the following question.

Question: Does there exist a finite-time approximation algorithm for computing the Teichmüller map between two n-gons?

We give the first results for theoretically constructing and algorithmically computing Teichmüller maps for the polygon problem above. Our procedure is iterative; we start with a q.c.h. that sends the vertices of P to those of Q in the prescribed order, improve on it, and then recurse on the improved map.

The need for an algorithm. Conformal geometry has found many applications in the fields of computer graphics [14], computer vision [24] and medical imaging [25, 13]. Computing Teichmüller maps generalizes almost all of these applications as q.c. maps allow boundary values to be prescribed. In [26], it was concluded that extremal q.c. maps have almost all the properties desired from an ideal surface registration algorithm, one of the biggest problems in computer vision.

---

[1] A homeomorphism f is angle preserving if it preserves oriented angles between curves: For any two curves γ_1 and γ_2 through a point p and oriented angle θ between them, $f(\gamma_1)$ and $f(\gamma_2)$ intersect at $f(p)$ at angle θ.

[2] For a function f between open sets in the complex plane $\mathbb{C}$, $\mu_f = \frac{f_{\bar{z}}}{f_z} = \frac{f_x + if_y}{f_x - if_y}$, where f_x and f_y denote partials w.r.t x and y, respectively.

An algorithm for computing Teichmüller maps would be a step forward in examining various questions in pure mathematics too. In [3] the author proposes how an algorithm for our problem would help us attack one of the most famous conjectures in geometric function theory – Brennan's conjecture. Teichmüller theory is an active area of research in mathematics, and it has connections to topology[3], dynamics, algebraic geometry, and number theory [15]. An algorithm for our problem helps one compute and visualize geodesics in the so-called Teichmüller space (w.r.t. Teichmüller's metric), which may be of independent interest.

Related work. Almost all algorithms in computational q.c. geometry have appeared mainly in graphics or vision venues. In many works (e.g. [19]) a q.c.h. is represented by its Beltrami coefficient, and softwares implementing basic subroutines (e.g. solving the Beltrami equation) in computational q.c. geometry have existed for some time.

The first paper addressing the problem of computing extremal q.c. maps was [26]. The authors propose an interesting heuristic based on Teichmüller's characterization; they formulate an energy function and minimize it using an alternate-descent method. Simulations showed that if the initial map is chosen correctly, the algorithm converges in many instances. Unfortunately, the energy obtained is "highly nonlinear" and non-convex. Even in the absence of numerical errors due to discretization, it is not known whether the minimization procedure converges to an approximation of the extremal map.

In [17] another heuristic was proposed using the connection to the theory of harmonic maps. This was simulated on a variety of examples and in many instances ended up with a good answer. However, no convergence proofs (continuous or discrete settings) were provided. Recently, in [18] it is argued that a procedure similar to that in [17] converges in the limit if certain parameters are chosen carefully manually. However, there are no bounds on the progress made in a step, and therefore it is not known if the procedure (even in the continuous setting) ends with an approximation in finite time.

Results. In comparison to all the previous work, we take a theoretical approach to constructing an algorithm for Teichmüller maps. In the continuous setting we have a procedure (Theorem 8) that converges in the limit to the exact extremal map and we also give bounds on the progress made in each step. Using this we can show that our procedure always, no matter what the starting map, gives an arbitrarily good approximation of the desired map in a finite number of iterations. A salient feature of our analysis is that we do not use an energy-based approach and work directly with the dilatation (the maximum angle stretch).

In the discrete setting, we state precisely all the subroutines needed for our algorithm and provide approximation guarantees. We present a novel subroutine `INF-EXT` that produces a type of Beltrami coefficient fundamental in the study of extremal maps, and prove (Theorem 9) that it produces an arbitrarily good approximation. We give error bounds on the discrete algorithm we propose, and show that (Theorem 10) modulo two basic subroutines[4], our algorithm produces a $(1+\epsilon)$-approximation of the extremal map.

[3] It had been used by Lipman Bers to give a simpler proof of Thurston's classification theorem for surface homeomorphisms.

[4] It is indeed surprising that tasks as basic as composing two q.c.h. (specified by their piecewise constant Beltrami coefficient), or finding the inverse of one, cannot be accomplished correctly yet. These two subroutines have been implemented in the past various times without error bounds, and as of now no approximation algorithms exist for them.

Because of space constraints in this extended abstract, all of our complete proofs can be found in the full version [11].

2 Informal discussion of results and techniques

As mentioned in the introduction, the aim is to compute the extremal q.c.h between two polygons. Intuitively, if μ_f is the Beltrami coefficient of f, f maps an infinitesimally small circle around p to something that roughly looks like a small ellipse at $f(p)$, with $(1 + |\mu_f(p)|)/(1 - |\mu_f(p)|)$ as the ratio between its major and minor axes.

Our strategy to tackle the polygon mapping problem is to first reduce it to the marked sphere problem. The marked sphere problem is: Given a q.c.h. f_0 from the sphere to itself taking a collection of given points (z_k) to another collection (w_k), compute the unique extremal q.c.h. f_* that not only takes z_k to w_k (for all k) but is also isotopic to f_0 (i.e. it can be "continuously deformed" to f_0 after pinning the values at z_k). We first prove that a solution to the marked sphere problem gives a solution to the polygon mapping problem (Theorem 7). For future reference we also note that the complex plane can be thought of as the sphere minus the north pole (the point at infinity).

Representation and complexity. In theory, a normalized q.c.h. f can be specified by specifying μ_f. For computational purposes, unless a closed form expression for f_* or μ_* is available, the best one can do is to evaluate f_* or μ_* on a dense mesh of points inside the domain. Our goal can be stated as follows.

Goal: Given a $\delta > 0$, compute the values of f_* on a given set of points inside the base polygon P, where the Beltrami coefficient μ_f of f satisfies $||\mu_f||_\infty < ||\mu_*||_\infty + \delta$.

Even if the polygons P and Q have rational coordinates, there is no known way to represent the extremal map with finite precision (for all we know, all representations may consist of transcendental numbers). In fact, we have found examples where this is true even for the Schwarz-Christoffel mapping [5]. Thus, it is not known whether the problem is in NP or not. We therefore straightaway aim towards an approximation algorithm. The model we consider is a real RAM model, where we are allowed to do exact basic arithmetic operations and take logarithms of complex numbers in constant time.

2.1 Continuous construction

One of our main results is constructing a sequence of q.c.h. f_i (which can all be continuously deformed to the starting map f_0) that converge to the desired extremal q.c.h. quickly (to get within ϵ of the extremal one we need $O(1/\epsilon^4)$ iterations). The map f_{i+1} is obtained from the previous one f_i by a composition $f_{i+1} = f_i \circ h_{i+1}$ where h_{i+1} fixes all the z_k and is obtained from f_i by convex optimization and solving (partial and ordinary) differential equations.

The main innovation in our approach is to "search" for the "best" map indirectly in a sense. One important result [1] in Teichmüller theory is the following : Given a complex-valued function $\mu_f(p)$ such that $\|\mu_f\|_\infty < 1$ there is an essentially unique q.c.h. f such that $\mu_f(p)$ is its "angle-stretch". In other words, the q.c.h. are "indexed" by their Beltrami coefficients.

[5] The Schwarz-Christoffel mapping is the "explicit" formula for the conformal map from the upper half plane $\mathbb{H}$ to a polygon, and, by composition, a formula for the conformal map between two polygons

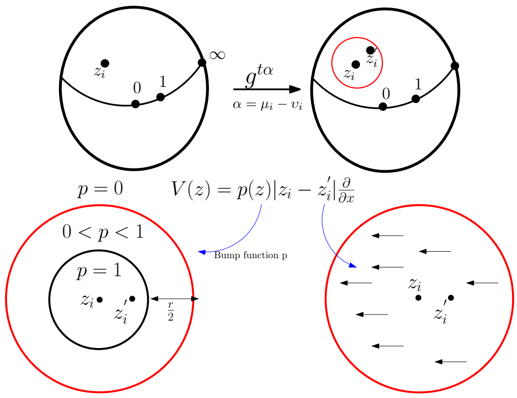

Figure 1 Construction of the self map h_i. Left: The map $g^{t\alpha}$ moves z_i to a point z_i' within $O(t^2)$. Right: The disk of radius $r = O(t^2)$ enlarged, showing the bump function p and the direction of the flow of the vector field.

One recovers f from μ_f by solving a partial differential equation called the *Beltrami equation*, $f_{\bar{z}}/f_z = \mu$.

Given the Beltrami coefficient μ_i of f_i, we search for the best (least L^∞ norm) Beltrami coefficient υ_i satisfying a certain technical condition called “infinitesimal equivalence” (Definition 3). This essentially boils down to a convex optimization problem. For a small $t > 0$, the q.c.h. g_i corresponding to the Beltrami coefficient $t(\mu_i - \upsilon_i)$ almost fixes the (z_k). It moves them only slightly (Figure 1 left). Then we correct for this motion by flowing the images $z_k' = g_i(z_k)$ back to (z_k) by solving a system of ordinary differential equations using a vector field, shown in the right side of Figure 1. We then compose these two maps.

This way, we get a map h_{i+1} which fixes the points (z_k). Moreover, we can prove that $f_{i+1} = f_i \circ h_{i+1}$ *has a smaller maximum angle-stretch (i.e., smaller dilatation)* than f_i. We iterate this process to converge to an arbitrarily good approximation of the desired extremal q.c.h. relatively quickly.

2.2 Approximation algorithm

We discretize the continuous construction given above in order to come up with an approximation algorithm `EXTREMAL` modulo two basic subroutines. Along the way we come up with a subroutine (which we call `INF-EXT`) that finds the best piecewise constant Beltrami coefficient υ that is infinitesimally equivalent to a given one μ. We believe that this is an interesting technical result in its own right.

Our input is a mesh of sample points on the sphere, a triangulation of the sphere, a piecewise constant Beltrami coefficient (corresponding to the starting map f_0), and an error tolerance δ. The desired output is the collection of images of these sample points under the extremal q.c.h. within the error tolerance.

We follow the same steps as in the continuous construction. There is a small technicality in that we need a special kind of triangulation, and might need to make this triangulation smaller each time we use any of our subroutines to control the errors. To this end, we use a subroutine `TRIANG` which is constructed using the Delaunay refinement algorithm. We take the piecewise constant Beltrami coefficient μ_i, feed it into `INF-EXT` and obtain a piecewise constant Beltrami υ_i. Just as in the continuous construction, we choose an appropriate $t > 0$ and find the q.c.h. g_i corresponding to the Beltrami coefficient $t(\mu_i - \upsilon_i)$. To obtain g_i we solve the Beltrami equation using a subroutine `BELTRAMI`.

The q.c.h. g_i moves the points (z_k) a bit. We remedy this by using the vector field method through a subroutine `VECT-FIELD`. The subroutines `BELTRAMI` and `VECT-FIELD` are standard. Then we compose the maps akin to the continuous construction. Here is where we need to assume the existence of two technical, basic subroutines `PIECEWISE-COMP` and `PIECEWISE-INV`. Once this composition is performed, we obtain a map f_{i+1} which has a smaller dilatation than f_i. We set f_{i+1} as the starting map and iterate; the algorithm terminates by producing an approximation of the desired extremal map f_*.

The issue with the two subroutines `PIECEWISE-COMP` and `PIECEWISE-INV` is as follows: Given piecewise constant Beltrami coefficients α and β (whose corresponding q.c.h. are F and G respectively) we want to compute a good piecewise constant approximation of the Beltrami coefficient corresponding to F^{-1} and to $F \circ G$. Any algorithm in computational q.c. geometry may require these subroutines. There are good candidates for such subroutines but the problem is to prove their correctness. We did not perform any complexity analysis of our algorithm simply because we do not know the complexity of the conjectural subroutines `PIECEWISE-COMP` and `PIECEWISE-INV`. But we expect our algorithm `EXTREMAL` (including the assumed subroutines) to run in polynomial time.

3 Preliminaries

In this section we present the main players from q.c. theory involved in our construction. Various eminent mathematicians (Teichmüller, Ahlfors, Bers, Reich, Strebel, Krushkal, Hamilton, etc.) have contributed to Teichmüller theory. We refer the reader to [10] and [15] for some excellent introductions to Teichmüller theory.

Quasiconformal maps and Beltrami coefficients/differentials. For a function f between two open sets in the complex plane, define partials $f_z = f_x - if_y$ and $f_{\bar{z}} = f_x + if_y$, where f_x and f_y are the partials with respect to (Euclidean coordinates) x and y. Let $\hat{\mathbb{C}}$ denote the Riemann sphere ($\mathbb{C}$ union the point at infinity). A homeomorphism $f\colon \hat{\mathbb{C}} \to \hat{\mathbb{C}}$ is quasiconformal provided that it satisfies the *Beltrami equation* $f_{\bar{z}} = \mu(z) f_z$ for some complex-valued function μ satisfying $||\mu||_\infty < 1$. μ is called the *Beltrami coefficient*, and is a measure of the non-conformality of f. In particular, the map f is conformal if μ is identically 0. The following theorem makes the notion of the Beltrami coefficients indexing the corresponding q.c.h. precise.

▶ **Theorem 1.** *The Beltrami equation gives a one to one correspondence between the set of quasiconformal homeomorphisms of $\hat{\mathbb{C}}$ that fix the points $0, 1$ and ∞ and the set of measurable complex-valued functions μ on $\hat{C}$ for which $||\mu||_\infty < 1$. Furthermore, the normalized solution f^μ of the Beltrami equation of μ depends holomorphically on μ and for any $r > 0$ there exists $\delta > 0$ and $C(r) > 0$ such that*

$$|f^{t\mu}(z) - z - tV(z)| \leq C(r)t^2 \ \textit{for}\ |z| < r \ \textit{and}\ |t| < \delta, \tag{1}$$

where $V(z) = -\frac{z(z-1)}{\pi} \int\int_{\mathbb{C}} \frac{\mu(\zeta)d\xi d\eta}{\zeta(\zeta-1)(\zeta-z)}$, *and* $\zeta = \xi + i\eta$.

We need some more definitions and concepts. They are summarized here:

Composition formula. Let μ, σ and τ be the Beltrami coefficients of quasiconformal maps f^μ, f^σ and f^τ with $f^\tau = f^\sigma \circ (f^\mu)^{-1}$. Then

$$\tau = \left(\frac{\sigma - \mu}{1 - \bar{\mu}\sigma} \frac{1}{\theta} \right) \circ (f^\mu)^{-1}, \quad \text{where} \quad p = \frac{\partial}{\partial z} f^\mu(z) \quad \text{and} \quad \theta = \frac{\bar{p}}{p}. \tag{2}$$

Quadratic differentials. For $R = \hat{\mathbb{C}}_{\{0,1,\infty,z_1,\dots z_{n-3}\}}$ (the Riemann sphere with n marked points, three of which are normalized to be 0, 1 and ∞), the complex vector space formed by the linear span of the $n-3$ functions

$$\phi_k(z) = \frac{1}{z(z-1)(z-z_k)}, \quad 1 \le k \le n-3, \tag{3}$$

is called the space of holomorphic quadratic differentials on R, denoted by $A(R)$.

Equivalence relations on Beltrami coefficients. Let $B(R)$ denote the set of all complex-valued measurable functions on R. Let $B_1(R) = \{\mu \in B(R) : ||\mu||_\infty < 1\}$. Given two coefficients μ and υ in $B_1(R)$, denote the solution to their respective normalized[6] Beltrami equations as f^μ and f^υ. Let R_0 and R_1 denote two marked spheres. The following definition concerns maps from R_0 to R_1.

▶ **Definition 2** (Global equivalence). μ and υ are called globally equivalent ($\mu \sim_g \upsilon$) if:
1. $f^\mu(z_i) = f^\upsilon(z_i)\ \forall i$.
2. The identity map from R_1 to R_1 is homotopic to $f^\upsilon \circ (f^\mu)^{-1}$ via a homotopy consisting of quasiconformal homeomorphisms.

A Beltrami coefficient υ is called trivial if it is globally equivalent to 0. A Beltrami coefficient with the least L_∞ norm in its global class is called globally extremal. *In other words, the marked sphere problem specifies as input a Beltrami coefficient μ, and asks to output the extremal Beltrami coefficient μ_* that is globally equivalent to μ.*

▶ **Definition 3** (Infinitesimal equivalence). μ and υ are infinitesimally equivalent (written $\mu \sim_i \upsilon$) if $\int_R \mu\phi = \int_R \upsilon\phi$ for all $\phi \in A(R)$, with $||\phi|| = 1$. A Beltrami coefficient υ is called infinitesimally trivial if it is infinitesimally equivalent to 0.

▶ **Definition 4** (Infinitesimally extremal). A Beltrami coefficient υ is called infinitesimally extremal if $||\upsilon||_\infty \le ||\mu||_\infty$ for all $\mu \sim_i \upsilon$.

Optimality condition. The importance of the infinitesimally extremal Beltrami coefficients is conveyed by the celebrated Hamilton-Krushkal, Reich-Strebel, necessary and sufficient condition for optimality. Informally, this theorem states that a Beltrami coefficient μ_* is globally extremal if and only if it is infinitesimally extremal and the corresponding q.c.h. takes the domain to the desired target. See [10] for a precise statement.

Another important fact is that for all the cases we are interested in, any globally extremal Beltrami coefficient is of Teichmüller form – it can be written as $\mu_* = k_*\bar{\phi}/|\phi|$, for a unique constant $k_* < 1$ and a unique quadratic differential $\phi \in A(R)$.

An important remark on the optimality condition. Note that given a starting μ, the ν that is extremal in the infinitesimal class of μ will be of Teichmüller form. However, *it will generally not be globally equivalent to μ*. This is why we have an iterative procedure – if ν was also globally equivalent to μ we would be done in one step. We use ν and μ to obtain μ_1, and inductively ν_1 to obtain μ_2 and so on, to get to the globally extremal μ^* which is in the same global class as μ and is infinitesimally extremal in its class, and hence is of Teichmüller form.

[6] Fixing the points 0,1 and ∞. Hence the freedom of Möbius transformation is accounted for.

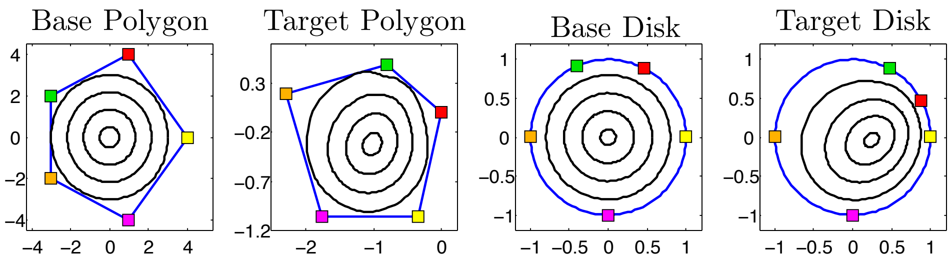

Figure 2 An example of a Teichmüller map between pentagons. If ϕ_1 and ϕ_2 are a basis of the space of quadratic differentials, the above map corresponds to the solution to the Beltrami equation of $\mu = \frac{\bar{\phi}}{8\phi}$, where $\phi = \frac{1}{3}\phi_1 + \frac{2}{3}\phi_2$. On the right is the same map when pulled to the unit disks via the Riemann mapping.

4 Problem statement and main theorems

In this section we first describe the polygon mapping and the marked sphere problems, and prove that the marked sphere problem is more general. We will then proceed to state our main results.

4.1 Problem statements and reduction

Let P and Q be two n-gons[7] in the plane. Let $\{v_i\}_{i=1}^n$ and $\{v_i^{'}\}_{i=1}^n$ be an ordering of the vertices of P and Q, respectively. The fact that the polygons are conformally equivalent to the upper half plane $\mathbb{H}$, and that composition by conformal maps does not change the dilatation imply that an n-gon is essentially the same as $\mathbb{H}$ with n marked points on the boundary $\partial\mathbb{H} = \mathbb{R}$.

▶ **Problem 5** (Polygon mapping problem). *Given* $\{z_1, ...z_n, w_1, ...w_n\} \in \partial\overline{\mathbb{H}}$*, find* $\tilde{f}_* : \overline{\mathbb{H}} \to \overline{\mathbb{H}}$ *(with Beltrami coefficient* $\mu_{\tilde{*}}$*) satisfying:*

1. $\tilde{f}_*$ *is a quasiconformal homeomorphism of* $\overline{\mathbb{H}}$ *to itself.*
2. $\tilde{f}_*(z_i) = w_i$, $i \in \{1, ...n\}$
3. $||\tilde{\mu}_*||_\infty \leq ||\mu_f||_\infty$ *for all* f *satisfying (1) and (2) above.*

Note that by Teichmüller's theorems the above $\tilde{f}_*$ exists and is unique. We state the marked sphere problem next, and show that it is in fact a generalization of the polygon mapping problem.

▶ **Problem 6** (Marked sphere problem). *Given* $\{z_1, ...z_{n-3}, z_{n-2} = 0, z_{n-1} = 1, z_n = \infty\}$, $\{w_1, ...w_{n-3},\ w_{n-2} = 0, w_{n-1} = 1, w_n = \infty\}$*, and* $f_0 : \hat{\mathbb{C}} \to \hat{\mathbb{C}}$ *such that* $f_0(z_i) = w_i$*, find* $f_* : \hat{\mathbb{C}} \to \hat{\mathbb{C}}$ *satisfying:*

1. f_* *is a quasiconformal homeomorphism of* $\hat{\mathbb{C}}$ *to itself.*
2. f_* *is isotopic to* f_0 *relative to the points* $\{0, 1, \infty, z_1, ..z_{n-3}\}$*, i.e.* $f_*(z_i) = w_i$*.*
3. $||\mu_*||_\infty \leq ||\mu_f||_\infty$ *for all* f *satisfying (1) and (2) above.*

We call the base z_i-marked sphere R and the target w_i-marked sphere S from now on. The reason why the marked sphere problem requires a starting map f_0 as input is that by Teichmüller's theorem, the extremal map exists and is unique within each isotopy class. The following theorem shows that Problem 6 is indeed general.

[7] We allow for ∞ to be a vertex of the polygon.

▶ **Theorem 7** (Reduction). *An algorithm for Problem 6 can be used to give a solution to Problem 5.*

Proof sketch. Consider an instance of the polygon mapping problem, and map the polygons conformally in linear time using [2] to the upper-half plane such that the vertices go to points on the real line. Then, using a piecewise affine function f_0 map the corresponding upper half-planes to one another taking the vertices to the vertices. Since f_0 is real on $\mathbb{R}$, we extend it by symmetry to the entire Riemann sphere. Call this extended map f. This then provides us a special instance of the marked sphere problem, where all the marked points are on the real line. We then prove that the extremal map f_* homotopic to f is symmetric, and that the restriction of f_* to the upper half plane solves the original polygon mapping problem. Full proof in [11]. ◀

4.2 Results

Denote the Beltrami coefficient of f_0 by μ_0. We want to obtain μ_* that is globally equivalent (Definition 2) to μ_0 and has the smallest L_∞ norm in this global class. We will obtain a sequence of q.c.h. f_i (and their Beltrami coefficients μ_i) that in the limit converge to the unique extremal map f_* (and the dilatations of μ_i will converge to the dilatation of μ_*). All the μ_i lie in the same global class – that of μ_0. The main difficulty we overcome is that since the global class of μ_0 does not have a "nice" structure (e.g. it is not convex in the generic case; in fact the only way to know whether two Beltrami coefficients μ_1 and μ_2 are in the same global class is to solve their Beltrami equations). To overcome this, we break up this minimization over the global class of μ_0 into a sequence of minimizations over the infinitesimal classes (Definition 3) of μ_i (that are convex domains) and solutions of differential equations.

We will first present our main theorem in the continuous setting. By the "continuous setting" we mean that we assume the existence of black boxes that solve all the sub-problems involved exactly; e.g. given a Beltrami coefficient μ, we can get $f^\mu(z)$ for any z exactly.

▶ **Theorem 8** (Limiting procedure for Marked Sphere Problem). *There exists a sequence of q.c.h. f_i s.t.:*

1. **Isotopic:** *f_i is isotopic to f_0, and $f_i(z_j) = w_j$, for all i and j.*
2. **"Explicit" construction:** *Let υ_i be the extremal coefficient in the infinitesimal class of μ_i. Then μ_{i+1} is an "explicit function" of μ_i and υ_i in that it can be obtained by solving two differential equations depending only on μ_i and υ_i.*
3. **Uniform Convergence:** *$f_i \to f_*$ uniformly and $\|\mu_i\|_\infty \to \|\mu_*\|_\infty$ as $i \to \infty$.*
4. **Fast convergence:** *There exist constants $C > 0$ and $\delta_0 > 0$ such that for all $\delta < \delta_0$ and for all $i \geq C/(\delta^4(1 - ||\mu_0||_\infty)^2)$ we have $||\mu_i||_\infty - k_* < \delta$.*

Basically, getting υ_i from μ_i is the convex optimization part, and getting μ_{i+1} from μ_i and υ_i requires solving differential equations.

Now we proceed to the discrete implementation of our procedure. We represent all Beltrami coefficients as piecewise constant coefficients[8] on a fine mesh. Every step of the continuous procedure mentioned above is shown to have a discrete analogue. The mesh we

[8] In fact, the existence of the solution to the Beltrami equation of an arbitrary $\mu \in L^\infty$ with $||\mu||_\infty < 1$ was shown by 1) first showing the existence of the solution to a piecewise constant μ', 2) sewing the individual piecewise q.c. maps along the boundary, and 3) taking a limit of such piecewise constant coefficients $\mu'_n \to \mu$ and showing that the maps converge.

will be working on depends on the error tolerance δ required. The first theorem tells us how to discretise the convex optimization part.

▶ **Theorem 9** (Discrete infinitesimally extremal). *Given an error tolerance $0 < \delta < 1$, a collection of n marked points $z_1, z_2, \dots z_n$, a triangulation Δ_ϵ and a piecewise constant Beltrami coefficient μ (where $|\mu| < 1$ on every triangle), there exists an algorithm* `INF-EXT` *that computes a piecewise constant Beltrami coefficient $\hat{\upsilon}$ such that $|\hat{\upsilon}| - ||\upsilon||_\infty < \delta$ everywhere.*

Now we proceed towards the other steps. Computational quasiconformal theory is a field still in its infancy, and very few error estimates on these widely-used discretizations are known. We introduce two subroutines `PIECEWISE-COMP` and `PIECEWISE-INV` (their precise definitions are in section 6) that concern the discretization of compositions and inverses of quasiconformal maps. Assuming the existence of the subroutines `PIECEWISE-COMP` and `PIECEWISE-INV` we construct an approximation algorithm for the Teichmüller map.

▶ **Theorem 10** (Teichmüller Map Algorithm). *Assume the existence of the aforementioned subroutines. Given a triangulation T_0 that includes n marked points $z_1, \dots z_n$, a mesh of sample points S, an error tolerance δ, and a piecewise constant Beltrami coefficient μ_0 whose corresponding q.c.h. f_0 satisfies $f_0(z_j) = w_j$, there exists an algorithm* `EXTREMAL` *that computes Δ_ϵ, and the images of S up to an error of δ under a q.c.h. F having a piece-wise constant (in the computed triangulation) Beltrami coefficient μ_F such that*

1. *$\|\mu_F\|_\infty - \|\mu_*\|_\infty < \delta$ where μ_* is the Beltrami coefficient of the extremal quasiconformal map on the marked sphere in the isotopy class of f_0.*
2. *$|F(z_i) - w_i| = O(\delta)$.*

Thus our main result in the discrete case is a reduction of this approximation problem to two basic subroutines. We do not address the complexity of our approximation algorithm and expect that (along with the two conjectural subroutines) our algorithm runs in polynomial time.

5 The continuous construction

We first summarize our construction of the sequence $\{f_i\}$ of q.c.h. that converge to the extremal map. At step i, given the q.c.h. f_i with Beltrami coefficient μ_i, let υ_i denote the infinitesimally extremal Beltrami coefficient in the infinitesimal class of μ_i. Let $k_i = ||\mu_i||_\infty$ and $k_i^0 = ||\upsilon_i||_\infty$. Observe that $\mu_i - \upsilon_i$ is infinitesimally trivial (Definition 3).

1. Choose t such that
$$t = \min\left(\frac{3}{4}, C_1, \frac{\epsilon}{4}, \sqrt{\frac{\epsilon}{2C_2}}, \frac{(k_i - k_i^0)^2(1-k_i^2)}{1-k_i^2+C_2}\right), \tag{4}$$
where $\epsilon \le \min(1/2, (k_i - k_i^0)/8)$, and C_1 and C_2 are two explicit constants derived in the full version[11].
2. Use Subsection 5.1 to construct a quasiconformal self-homeomorphism h_i of the base z_k-marked sphere such that
   - μ_h is globally trivial (hence $h_i(z_k) = z_k$ for all k).
   - $||\mu_h - t(\mu_i - \upsilon_i)||_\infty < C_2 t^2$, where C_2 is the same constant as in (4).
3. Form $f_{i+1} = f_i \circ (h_i)^{-1}$. It turns out that f_{i+1} has smaller dilatation than f_i (by Lemma 11).
4. Reiterate with f_{i+1} as the starting map.

The second to last step i.e., calculating the composition $f_{i+1} = f_i \circ (h_i)^{-1}$ is the main point of the construction. To our knowledge, this is the first "constructive" way to produce a map having a smaller dilatation than a given one. The heart of this step is the following crucial lemma (proof in [11]):

▶ **Lemma 11** (Decreasing dilatation). *Let υ_f be the infinitesimally extremal Beltrami coefficient in the infinitesimal class of μ_f. Let $\mu_h(t)$ be a curve of Beltrami coefficients with the following properties:*

1. *$\mu_h(t)$ is globally trivial.*
2. *$\mu_h(t) = t(\mu_f - \upsilon_f) + O(t^2)$.*

Denote the solution to the Beltrami equation of $\mu_h(t)$ by h_t. Then $\exists\delta > 0$ such that $\forall t < \delta$, the map $f_t = f \circ (h_t)^{-1}$ has smaller dilatation than f.

Proof sketch of Theorem 8. Assume for now that the map h_i produced in each step satisfies the conditions of Lemma 11. Let $k_i = ||\mu_i||_\infty$ be the L^∞ norm of the Beltrami coefficient of f_i(the starting map at step i), and $k_i^0 = ||\upsilon_i||_\infty$ where υ_i is infinitesimally extremal. We lower bound the decrease $d = k_i - k_{i+1}$ in the dilatation in step 3 in terms of $k_i - k_i^0$. This is bounded below further by an expression which is in terms of $k_i - k_*$ (the distance from the extremal map). This is accomplished using Teichmüller's contraction principle, which gives a quantitative version of the following fact: If a Beltrami coefficient μ is close to the infinitesimally extremal coefficient υ, then it is also close to the globally extremal coefficient μ_*. Once we have d in terms of $k_i - k_*$, a standard geometric series argument coupled with a theorem on uniform convergence of sequences of q.c.h. on the sphere completes the proof. ◀

5.1 Constructing the self homeomorphisms

Starting at the ith step with a q.c.h. f_i, we now show how to construct the self homeomorphism h_i required by Lemma 11. We simplify notation by suppressing the index i, keeping in mind that this is the ith step of the procedure. Thus μ and μ_h will denote the Beltrami coefficients of f_i and h_i, respectively. Also, υ is the infinitesimally extremal Beltrami coefficient in the infinitesimal class of μ.

Let $\alpha = \mu - \upsilon$, t be as in Equation (4), and let $g^{t\alpha}$ be the normalized solution to the Beltrami equation for $t\alpha$. Denote $g^{t\alpha}(z_k)$ by $z_k^{'}$. As a consequence of the mapping theorem Theorem 1 that $z_k^{'}$ is not very far from z_k (the "error" is $O(t^2)$).

We will first construct another homeomorphism $K_{\mathbf{v}}$ from $\hat{\mathbb{C}}$ to itself which satisfies $K_{\mathbf{v}}(z_k^{'}) = z_k$. We then define the required self homeomorphism $h = K_{\mathbf{v}} \circ g^{t\alpha}$. The construction of $K_{\mathbf{v}}$ will be via a vector field method. A summary of this vector field method is as follows.

Let $\{D_1, \cdots, D_{n-3}\}$ denote disjoint open disks centered at z_k. Choosing the radius of each disk to be $r = d/4$, where $d = \max\limits_{1 \le k,l \le n-3} |z_k - z_l|$ ensures disjointness. We will fix these disks once and for all.

We first construct a self homeomorphism $K_{\mathbf{v}}^k$ of $\hat{\mathbb{C}}$ which is the identity map outside D_k, and maps $z_k^{'}$ to z_k. By means of a rotation we can assume that $z_k^{'}$ is real and greater than z_k. Consider the vector field

$$X(z) = p(z)(z_k^{'} - z_k)\frac{\partial}{\partial x},$$

where $p(z)$ is a C^∞ function identically zero outside D_k, and identically 1 inside the disk of radius $r/2$ around z_k, denoted as $D_k^{'}$. Let γ be the one parameter family of diffeomorphisms associated with this vector field (i.e. the flow of this field). We denote the time parameter by s and note that the diffeomorphism γ_1 sends $z_k^{'}$ to z_k. We denote this diffeomorphism γ at $s = 1$ by $K_{\mathbf{v}}^k$. Now define $K_{\mathbf{v}} = K_{\mathbf{v}}^{n-3} \circ K_{\mathbf{v}}^{n-2} \cdots \circ K_{\mathbf{v}}^1$, and $h = K_{\mathbf{v}} \circ g^{t\alpha}$. This is the desired "correction" that ensures that the q.c.h. h is indeed a self map.

Using PDE theory of the Beltrami equation, we then prove that the Beltrami coefficient of h_i so obtained does satisfy the hypothesis of Lemma 11. This completes all the details of our continuous construction.

6 The approximation algorithm

Here we present details of our approximation algorithm. Near the marked points the mesh is made up of (triangulated) regular polygons, whose number of vertices and radii depend on δ. The mesh is a triangulation with edge lengths bounded above by an appropriate ϵ that depends on δ. We call this triangulation a canonical triangulation Δ_ϵ of size ϵ. Its precise definition can be found in the full version [11]. We describe the convex optimization part of our algorithm next.

6.1 `INF-EXT`

We want to discretize the operator $\mathcal{P}(\mu)$ which returns υ with the least L^∞ norm satisfying $\int_R \upsilon\phi_i = \int_R \mu\phi_i$ for all ϕ_i in Equation (3). Note that the starting μ is piecewise constant at the start of every iteration.

▶ **Observation 12.** *The integral of ϕ_i over any triangle t_j can be computed analytically. We note that this formula involves taking the logarithm of a complex number.*

We approximate υ by piecewise constant Beltrami coefficients. The constraints for infinitesimal equivalence become linear constraints of the form $Ax = b$, where $A(i,j)$th equals $\int_{t_j} \phi_i$, x is the vector of unknown values of the piecewise constant υ on a triangle, and b is the vector of $\int_{t_j} \mu_j\phi_i$, where μ_j is the value of μ on triangle t_j. If A, x and b are real, an L^∞ minimization can be formulated as a linear program. In our case, we break the vectors and matrices into their real and complex parts, and then we can formulate the program as a quadratically constrained quadratic program. Although in general they are NP-hard to solve, in the the full version [11] we show that our program involves positive semi-definite matrices, and it is known that such instances can be solved in polynomial time using interior-point methods [20].

▶ **Lemma 13** (`INF-EXT`). *There exists an algorithm* `INF-EXT` *that, given a piecewise constant μ on Δ_ϵ returns a piecewise constant $\hat{\upsilon}$ such that $\max_{t_j} \hat{\upsilon}(t_j) \leq \max_{t_j} \beta(t_j)$, where β is any piecewise constant (on Δ_ϵ) Beltrami coefficient that is infinitesimally equivalent to μ.*

With this, we are now in a position to prove Theorem 9, which says that this piecewise approximation $\hat{\upsilon}$ is not very far from the true infinitesimally extremal υ. The full proof is relegated to the journal version [11].

6.2 Description of `EXTREMAL`

Apart from the subroutine `INF-EXT` we require a few more subroutines to discretize our procedure.

- `TRIANG`. The input is a set of points $\mathcal{S}$, a size M, and a triangulation Δ_ϵ. The output of `TRIANG` is a triangulation $\Delta_{\epsilon'}$ of the given size M containing $\mathcal{S}$ such that $\Delta_{\epsilon'}$ is a refinement of Δ_ϵ.

- `BELTRAMI`. The input is a triangulation Δ_ϵ of the plane, a piecewise constant Beltrami coefficient μ, and error tolerance δ. The output of `BELTRAMI` is a triangulation $\Delta_\epsilon^{'}$ that is a refinement of Δ_ϵ, and the images $\hat{f}(v_i)$ of the vertices $v_i \in \Delta_\epsilon^{'}$ such that $|f^\mu(v_i) - \hat{f}(v_i)| < \delta$.
- `VECT-FIELD`. The input is a C^k (k sufficiently large, e.g. $k > 10$) vector field X (written as a formula in terms of elementary functions), a triangulation Δ_ϵ, and an error tolerance δ. The output is a triangulation $\Delta_\epsilon^{'}$ that is a refinement of Δ_ϵ, the images of $v_i \in \Delta_\epsilon$ up to error δ under a C^k diffeomorphism γ_x corresponding to the flow along X, and a piecewise smooth Beltrami coefficient that approximates μ_{γ_x} up to error δ.
- `PIECEWISE-COMP`. The input is a triangulation Δ_ϵ, two piece-wise constant Beltrami coefficients μ_1 and μ_2 (corresponding to q.c.h. f_1 and f_2 respectively), and error tolerances δ_1 and δ_2. The output is a triangulation $\Delta_{\epsilon'}$ that is a refinement of Δ_ϵ, a piecewise constant Beltrami coefficient μ_{comp} that approximates the Beltrami coefficient of the composition $f_3 = f_1 \circ f_2$ within error δ_1 in the L^∞ topology, and the images $f_3(v_a)$ of the vertices v_a of $\Delta_{\epsilon'}$ up to an error of δ_2.
- `PIECEWISE-INV`. The input is a triangulation Δ_ϵ, a piecewise constant Beltrami coefficient μ (corresponding to q.c.h. f), and error tolerances δ_1 and δ_2. The output is a triangulation $\Delta_{\epsilon'}$ that is a refinement of Δ_ϵ, a piecewise constant Beltrami coefficient μ_{inv} that approximates the Beltrami coefficient of f^{-1} within error δ_1 in the L^∞ topology, and the images $f^{-1}(v_a)$ of the vertices of $\Delta_{\epsilon'}$ up to an error of δ_2.

`EXTREMAL` The algorithm summarized below is based on Section 5.

- Use `TRIANG` to produce a triangulation of size required by `INF-EXT` to run within an error of δ^{10}.
- Loop $i = 1$ to N where N is the number of iterations in Theorem 8 to produce the result within an error of $\delta/2$.
  1. Use `INF-EXT` to produce υ_i from μ_i within an error of δ^{10}. If $\upsilon_i = \mu_i$ then stop.
  2. Find t_i by Equation (4), using k_0 as $\|\upsilon_i\|_\infty$.
  3. Invoke `BELTRAMI` for the coefficient $t_i(\mu_i - \upsilon_i)$ to find the images of the marked points within an accuracy of t_i^3.
  4. Define the vector field X as in the continuous construction using a piecewise polynomial version of the bump function (that is C^{10} for instance). Then call `VECT-FIELD` to find a piecewise constant Beltrami coefficient up to an error of t_i^3.
  5. Use `PIECEWISE-COMP` to compose the Beltrami coefficients of step 3 and step 4 within an error $(\|\mu_i\| - \|\upsilon_i\|)^5$ for the Beltrami coefficient and δ/i^2 for the q.c.h.
  6. Use `PIECEWISE-INV` to find the Beltrami coefficient of the inverse of the q.c.h. of step 5, up to the same error as that in step 5.
  7. Call `PIECEWISE-COMP` to compose μ_i and the Beltrami coefficient of step 6 to form μ_{i+1} (up to the same error as that in step 5).

Implementing `TRIANG`, `BELTRAMI` and `VECT-FIELD`

1. `TRIANG`. Given a set of n points, we can obtain the Delaunay triangulation in $O(n \log n)$ time. While implementing `TRIANG` we first compute the Delaunay triangulation of all the points falling inside a triangle of the given triangulation. The we connect the vertices on the convex hull of such a set of points to one of the three vertices of the triangle they lie in. If this complete triangulation is not yet size M, we make the mesh denser by adding points as in [21] (points are added to either the circumcenters of triangles or mid-points of edges), until we reach the desired size.

2. BELTRAMI. The solution to the Beltrami equation for μ can be expressed as a series of singular operators applied to μ. There are many efficient algorithms and implementations[6],[9] existing for BELTRAMI. Most of them can bound the L^p norm of the error, but the methods in [6] can be used to bound the L^∞ error too[5].
3. VECT-FIELD. The idea of deforming a surface by a vector field has been applied extensively in computer graphics. We refer the reader to [16] for an implementation of VECT-FIELD.

▶ Remarks. Using the composition formula for Beltrami coefficients (Equation (2)), we see that in principle one may attempt to compute a piecewise constant approximation of the Beltrami coefficient of the composition $f \circ g$ of two q.c.h. f and g, and of g^{-1} (by setting $\sigma = 0$). However, this requires the derivative of g to be well-approximated in a piecewise constant manner. Therein lies the difficulty. Basically, one needs a good way of "discretising" the definition of the Beltrami coefficient of a q.c.h. The algorithm terminates by producing μ_N. The proof of Theorem 10 is similar to that of Theorem 8 and is omitted.

7 Discussions and future work

Our algorithm for the marked sphere problem also solves as a special case what is known as the "landmark constrained" Teichmüller map problem, where the points z_i and w_i are in the interior of the polygons, and a starting map is provided that sends z_i to w_i. A reduction similar to Theorem 7 works.

Open problems abound. In addition to studying the two conjectural subroutines the extremal map problem can be further explored in many directions.

1. Most of the ideas presented here (notably Lemma 11) may be used to envision an algorithm for computing Teichmüller maps between other Riemann surfaces. The problem is challenging for multiple reasons – for instance, an explicit basis of holomorphic quadratic differentials may not be available.
2. The authors feel that building a discrete version of Teichmüller theory would be an important achievement. Given a triangulated Riemann surface, defining a discrete analog of dilatation that gives nice results (e.g. existence and uniqueness) about the extremal map would be the next step in this direction.

References

1 L. V. Ahlfors. *Lectures on quasiconformal mappings*, volume 38 of *University Lecture Series*. American Mathematical Society, Providence, RI, second edition, 2006. With supplemental chapters by C. J. Earle, I. Kra, M. Shishikura and J. H. Hubbard.
2 C. Bishop. Conformal mapping in linear time. *Discrete and Comput. Geometry*, 44(2):330–428, 2010.
3 Christopher Bishop. Analysis of conformal and quasiconformal maps. Results from prior NSF support, 2012. http://www.math.sunysb.edu/~bishop/vita/nsf12.pdf.
4 C. Carathéodory. Über die gegenseitige Beziehung der Ränder bei der konformen Abbildung des Inneren einer Jordanschen Kurve auf einen Kreis. *Mathematische Annalen*, 73(2):305–320, 1913.
5 P. Daripa and M. Goswami, 2014. Private communication.
6 Prabir Daripa. A fast algorithm to solve the beltrami equation with applications to quasiconformal mappings. *Journal of Computational Physics*, 106(2):355–365, 1993.
7 T. A. Driscoll and L. N. Trefethen. *Schwarz-Christoffel Mapping*. Cambridge Monographs on Applied and Computational Mathematics. Cambridge University Press, 2002.

8 T. A. Driscoll and S. A. Vavasis. Numerical conformal mapping using cross-ratios and delaunay triangulation. *SIAM J. Sci. Comput*, 19:1783–1803, 1998.

9 D. Gaidashev and D. Khmelev. On numerical algorithms for the solution of a beltrami equation. *SIAM Journal on Numerical Analysis*, 46(5):2238–2253, 2008.

10 F. P. Gardiner and N. Lakic. *Quasiconformal Teichmüler theory*. American Mathematical Society, 1999.

11 M. Goswami, X. Gu, V. Pingali, and G. Telang. Computing Teichmüller maps between polygons. arXiv:1401.6395 – `http://arxiv.org/abs/1401.6395`, 2014.

12 H. Grötzsch. Über die Verzerrung bei nichtkonformen schlichten Abbildungen mehrfach zusammenhngender Bereiche. *Leipz. Ber.*, 82:69–80, 1930.

13 X. Gu, Y. Wang, T. F. Chan, P. M. Thompson, and S. T. Yau. Genus zero surface conformal mapping and its application to brain surface mapping. *IEEE Transactions on Medical Imaging*, 23(7):949–958, 2004.

14 X. Gu and S.T. Yau. Global surface conformal parameterization. In *Symposium on Geometry Processing (SGP'03)*, volume 43, pages 127–137, 2003.

15 J. H. Hubbard. *Teichmüller theory and applications to geometry, topology, and dynamics*. Matrix Editions, 2006.

16 Ldmm – the large deformation diffeomorphic metric mapping tool. `http://cis.jhu.edu/software/lddmm-volume/tutorial.php`.

17 L. Lui, K. Lam, S. Yau, and X. Gu. Teichmüller Mapping (T-Map) and Its Applications to Landmark Matching Registration. *SIAM Journal on Imaging Sciences*, 7(1):391–426, 2014.

18 L. M. Lui, Xianfeng Gu, and Shing Tung Yau. Convergence of an iterative // algorithm for Teichmüller maps via generalized harmonic maps. arXiv:1307.2679 – `http://arxiv.org/abs/1307.2679`, 2014.

19 Lok Ming Lui, Tsz Wai Wong, Wei Zeng, Xianfeng Gu, Paul M. Thompson, Tony F. Chan, and Shing-Tung Yau. Optimization of surface registrations using beltrami holomorphic flow. *Journal of Scientific Computing*, 50(3):557–585, 2012.

20 P.M. Pardalos and M.G.C. Resende. *Handbook of applied optimization*, volume 1. Oxford University Press New York, 2002.

21 J. Ruppert. A delaunay refinement algorithm for quality 2-dimensional mesh generation. *J. Algorithms*, 18(3):548–585, 1995.

22 O. Teichmüller. Extremale quasikonforme Abbildungen und quadratische Differentiale. *Preuss. Akad. Math.-Nat.*, 1, 1940.

23 O. Teichmüller. Bestimmung der extremalen quasikonformen Abbildungen bei geschlossenen orientierten Riemannschen Flächen. *Preuss. Akad. Math.-Nat.*, 4, 1943.

24 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.

25 Y. Wang, J. Shi, X. Yin, X. Gu, T. F. Chan, S. T. Yau, A. W. Toga, and P. M. Thompson. Brain surface conformal parameterization with the ricci flow. *IEEE Transactions on Medical Imaging*, 31(2):251–264, 2012.

26 O. Weber, A. Myles, and D. Zorin. Computing extremal quasiconformal maps. *Comp. Graph. Forum*, 31(5):1679–1689, 2012.